

Module 4: Analyzing Performance of Symmetrix Front-end Adapters

Upon completion of this module, you should be able to:

- Relate front-end architecture to performance capabilities
- Use Performance Viewer to view key front-end measures
- Propose solutions to performance problems



This module focuses on analyzing the performance of Symmetrix front-end directors.

CPU and Port Comparison – VMAX 20K Example

VMAX
Directors

2.3 GHz CPU core, up to 8 Gbps ports

DMX Family
Directors

DMX-1
DMX-2
DMX-3,4

Single 500 MHz CPU
Single 1 GHz CPU
Dual 1.3 GHz CPU, up to 4 Gbps ports

EMC²

Copyright © 2014 EMC Corporation. All Rights Reserved.


Module 4: Analyzing Performance of Symmetrix Front-end Adapters 2

Every VMAX 20k director card has dual, quad-core 2.3 GHz Xeon processors. Although there are fewer processing cores per director in a VMAX, the increased processing speed produces significant improvements in I/O handling. Additionally, the architecture allows multiple processors to combine effort on one heavily-utilized director's tasks. Benchmarking shows an almost doubling of I/O rate with the VMAX arrays.

Naturally, the processor and port speeds have been increasing with each new generation. The latest directors can have 8 Gb/s ports. The latest VMAX 40k and 10K use 6 core 2.8 Ghz Intel processors.

DMX-1 and DMX-2 directors had a single processor that handled all tasks. It was responsible for negotiating I/Os with the host on both ports, and also performed tasks like looking up cache locations and allocating new slots. The DMX-3 and DMX-4 directors had two processors. One handled the host communications and one handled the cache communications. This sharing of work and the faster processor speed in these models lets them process up to 70% more I/Os than the DMX-2 models.

VMAX Host Director Architecture – VMAX 20K Example



- One Engine has two director cards
- Each director card has four independent host directors (E,F,G,H)
- Each director has one (iSCSI, GigE, Ficon, SRDF) or two (Fibre) host ports (0,1)
- Each director card has eight connections to two redundant Virtual Matrix fabrics for communication with components on other cards

EMC²

Copyright © 2014 EMC Corporation. All Rights Reserved.


Module 4: Analyzing Performance of Symmetrix Front-end Adapters 3

Architecture of a VMAX director is very similar to the DMX architecture. A VMAX engine has two director cards that have both front-end and back-end directors. Each director card also has Global Cache Memory. Each director card has four host directors that have one or two ports each, just as in the DMX series. Another four directors in the same card handle disk operations. There is one processor core per director.

Directors can communicate with other components on other cards through the Virtual Matrix fabric. Each card has eight connections to two redundant Virtual Matrix fabrics (four connections to each fabric). As in the DMX family architecture, directors share these channels. Even communication to the other director card on the same engine requires a Matrix communication. The VMAX 40k doubles the matrix connections.

As in the DMX series, EMC software refers to each VMAX processing unit of one or two ports as a “director.” Throughout the rest of this course, the word “director” will be used consistently in the same way the software uses the term.

How Directors Process I/Os



- Director processor negotiates requests; must divide time on both ports
- Can negotiate on a port while data is transferring on the other port
- Each port can transfer data simultaneously

EMC²

Copyright © 2014 EMC Corporation. All Rights Reserved.

Module 4: Analyzing Performance of Symmetrix Front-end Adapters 4

Symmetrix VMAX host directors manage two I/O ports and must share time negotiating requests among both. While a processor is negotiating a transfer at one port, it cannot negotiate a new transfer on the other port. It will have to block that I/O until the negotiations are complete for the first one. Naturally, this applies only to a single director; all of the independent directors on a card or engine could be negotiating I/Os simultaneously.


The actual byte-by-byte data transfer is handled by subordinate processors and buffers dedicated to each port. Once an I/O has been negotiated and started on one port, the director can turn its attention to the other port and negotiate another I/O. Data transfer can be occurring on both ports simultaneously. In many cases, all the ports on a card, engine, or array can be transferring I/Os simultaneously.

Measuring Front-end System-Wide Traffic

- System-Wide KPIs

- ▶ Kbytes read per sec
- ▶ Kbytes written per sec
- ▶ reads per sec
- ▶ seq reads per sec
- ▶ writes per sec
- ▶ ios per sec
- ▶ total ios per sec
- ▶ total reads per sec
- ▶ % reads
- ▶ % writes
- ▶ read hits per sec

Kpi	All
Host IOs/sec	
Host MBs/sec	
Host MBs Written/sec	
Avg Read Response Time (ms)	
Avg Write Response Time (ms)	
% Reads	
% Writes	
% Hit	
BE IOs/sec	
System WP Events/sec	
System Max WP Limit	
% Cache WP	
XtremSW Cache IOs/sec	
% XtremSW Cache Read Hits	
XtremSW Cache MBs/sec	
XtremSW Cache Response Time	



- **Host Ios Per Sec** = reads + seq reads + writes
- **Prefetched Tracks/Sec** = reads triggered by sequential read activity
- **FE Reqs/Sec** = Data transfers between cache and the FE directors

EMC²

Copyright © 2014 EMC Corporation. All Rights Reserved.

Module 4: Analyzing Performance of Symmetrix Front-end Adapters 5

Several I/O per second and throughput measures are recorded at the array level. You can see from the KPI list on this slide that metrics are divided by reads and writes, and that some percentage-of-total metrics are given. These are all system-wide summaries of traffic through the host director ports. The “Host IOs per sec” measure is actually measuring the total number of I/Os being serviced by all volumes, and includes writes, and random and sequential reads.


An additional metric that can be found in the ALL list, is recorded at the System and Device level: prefetched tracks per second (“Prefetched Tracks/sec”). It is the number of tracks per second prefetched from disk to cache triggered by a detection of a sequential read stream.

An example of a sequential read stream is where track 27 of a device is already in cache, and you request a read that is addressed in track 28, but track 28 is not already in cache. Further sequential reads will be hits, but will be counted under such metrics as the %Hit and FE Read Hit Reqs/sec metrics as well.

Front end requests per second (“FE Reqs/sec”) is a measure of the data transfers between the directors and the cache. An IO may require multiple requests depending on IO size, alignment, or both.

Note: Other helpful system metrics include “% reads”, “% writes”, and “% Hit”.

Measuring Front-end System-Wide Traffic (Continued)




EMC²

Copyright © 2014 EMC Corporation. All Rights Reserved.

Module 4: Analyzing Performance of Symmetrix Front-end Adapters 6

This Diagnostic chart shows the host I/Os per second (“Host I/Os/sec”) across a whole array in a time graph with the trend line turned on and 7 forecast points added.

Example – System Measurements




Shown here are a few examples of some of the Performance Viewer measures available for system-wide activity. On the left, is a chart showing host writes per second, host reads per second, and host IOs per second.

The Host IOs/sec in this case does not equal host reads/sec + host writes/sec because of the inclusion of sequential activity.

The center graph shows the host megabytes written per second (“Host MBs written/sec”), host MBs read per second (“Host MBs Read/sec”) and host megabytes per second (“Host MBs/sec”) metrics. In this case, the host MBs Read + host MBs written does equal host MBs per second. The right graph shows the host reads per second (“Host Reads/sec”) and the front-end read requests per second (“FE Read Reqs/sec”) metrics.

Measuring I/O Per Second and Utilization by Director



EMC²

Copyright © 2014 EMC Corporation. All Rights Reserved.

Module 4: Analyzing Performance of Symmetrix Front-end Adapters 8

Shown here, on the top left is a diagnostic chart with directors FA-8F and FA-7F specifically selected and showing the host I/Os per second.

The bottom left is the summary dashboard front-end directors summary chart and the right chart is just one chart from the all front-end directors dashboard ("All FE Directors"). They are all showing similar things, just in different presentation formats.

Measuring I/O Per Second and Utilization by Director




To view performance for an individual front-end director, choose the front-end director dashboards and select the individual director you are investigating. In this example, we selected the front-end director FA-7F.

The front-end director dashboard depicted on this slide shows average write response time (in milliseconds), average read response time, percent of hit requests, host megabytes per second, host IOs per second, and percent busy.


Measuring I/O Per Second and Utilization by Director

- FE Director

- ▶ Host IOs / sec
- ▶ reads per sec
- ▶ writes per sec
- ▶ % Busy
- ▶ % Idle
- ▶ Port 0|1 IOs / sec



IOs Per Sec = Reads + Writes



EMC²

Copyright © 2014 EMC Corporation. All Rights Reserved.

Module 4: Analyzing Performance of Symmetrix Front-end Adapters 10

Measuring I/O rate and throughput is straightforward at the director level. I/Os per second, reads, and writes are available for a director under the “FE Director” tab. I/Os per second are also available for each port of a host director.


There are also utilization measures based on clock ticks (% Busy, % Busy Port 0|1). While these are not directly related to the number of I/Os per second, they accurately measure the utilization of the director. As a general guideline exceeding 50% of a director should raise concerns for growth in capacity and performance.

This chart is showing both the average and maximum Host IOs per second. This is done by selecting the chart properties dropdown, and selecting them from the Data Format options.

A few of the 15 KPIs are listed here. If you select the All tab (as opposed to KPI), you will see several pages of various metrics that can be shown.

What to Look For

- I/O per second imbalance among directors
- High director utilization— especially when I/O sizes are small




When examining the I/O per second for front-end directors, look first for balance. As we know from the Little's Law curve, a component can start having unacceptable delays well before reaching 100% utilization. An imbalance means that some components are more likely to be seeing these delays.

Also, look for high percentage busy. It is much more likely that a director will reach high utilization numbers when the I/O sizes are relatively small (<8 KB). A busy application will have to send a large number of small I/Os to transfer a large amount of data. This high number of I/Os might test the director's limits.

Neither of the graphs shown here show a great deal of balance across the directors. Some are processing a large number of I/Os, while others are nearly idle. Both graphs show some evidence of load balancing software, such as PowerPath. All of the directors appear to be working in pairs.

Solutions

- Balance I/O across directors
 - ▶ Use one port of each before beginning to use the other port
- Reduce load on any heavily utilized directors
- Get faster director processors



EMC²

Copyright © 2014 EMC Corporation. All Rights Reserved.

Module 4: Analyzing Performance of Symmetrix Front-end Adapters 12

If imbalance is the problem, then obviously balancing the I/O across all components is the solution. This is harder to do in an established environment, so take steps to balance the workload as new hosts and applications are being added. As the array is first brought online, use just one port of each director initially, then start using the second port. This helps spread the I/O around all the available directors before adding additional workload.

The illustration shows four front-end directors and lists the average I/O per second measure on each port.


Are the directors balanced?

Yes, each of these directors is processing 1000 I/O per second. Balancing I/O across a single director's ports will not reduce the utilization of the director.

Further examination might show that the port throughput is a problem on a director that has very uneven traffic on its ports, but port balance can be handled as a separate performance factor.

If balance has already been achieved and there is no way to reduce the workload, consider getting directors with faster processors if they are available for your array model. The maximum I/O per second is largely affected by the processor speed. It is not affected by the port speed!

Port Usage




EMC²

Copyright © 2014 EMC Corporation. All Rights Reserved.


Module 4: Analyzing Performance of Symmetrix Front-end Adapters 13

This is an example of the way the ports should be utilized. If you need to use the 2nd port of a director, you might choose lower utilized directors first.


Host Director I/O Per Second Examples




Oracle/SAP (UNIX) OLTP



MS Exchange (NT) Mail



SQL Server/SAP (NT) OLTP



Oracle (UNIX) Banking-OLTP

EMC²


Copyright © 2014 EMC Corporation. All Rights Reserved.

Module 4: Analyzing Performance of Symmetrix Front-end Adapters 14

Pictured above are several examples of host director I/O per second graphs that were taken from EMC customers. Note the time of day traffic changes and notable lack of balance in these examples.

Issue: Traffic Volume – Throughput

- Storage channels typically rated by maximum data transfer volume (Gb/s)
- Subordinate chips on each port are used to transfer data simultaneously
- Smaller I/Os have higher overhead to data ratio (smaller throughput)
- Larger I/Os have lower overhead to data ratio (larger throughput)



EMC²

Copyright © 2014 EMC Corporation. All Rights Reserved.

Module 4: Analyzing Performance of Symmetrix Front-end Adapters 15

Any host port has a maximum rating (usually rated in Gb/sec), but under normal traffic loads a lower actual transfer rate can be expected due to the overhead involved with transferring the I/Os. Some of the possible throughput of the channel will be consumed by the overhead. A greater percentage of the throughput will be consumed when small I/Os are transferred, since little data is transferred with each I/O. Large I/O transfers typically achieve much higher throughput.


I/O sizes are not set by the storage array; the host software always controls this I/O characteristic. When using filesystems, the filesystem manager (or the operating system kernel) typically controls the I/O size. Raw database applications that bypass filesystem managers set their own I/O size. In all cases, the size is usually chosen to help optimize the data transfers. Filesystem I/Os will typically be a multiple of the minimum file allocation size (the standard data block storage size for the filesystem). Databases will use an I/O size that closely matches the typical transaction size.

The size is never firmly fixed in an open systems environment. An occasional large transfer will be sent with very large I/O sizes if the filesystem or application thinks it will be more efficient. It is not unusual to see an occasional 100 KB I/O in the open systems environment.

As a result of the I/O overhead, the maximum realistic throughput for a channel is hard to pinpoint. Response time will grow larger as the amount of data transfers increases. Maximum, actual throughput is reached when the response time has become unacceptably large for the situation.

Measuring Throughput by Director

- FE Directors
 - Analyze> FE Director>Director(s) >Host MBs/Sec
 - Analyze> FE Director>Director(s) >Port 0|1 MBs/sec
 - Analyze> FE Director>Sort by Host MBs/sec



EMC²

Copyright © 2014 EMC Corporation. All Rights Reserved.

Module 4: Analyzing Performance of Symmetrix Front-end Adapters 16

You can measure director and port throughput in a variety of ways.


Use an Analyze view with the FE Directors folder selected to show directors. You can then sort this by any of the columns by clicking on the title.

You can select an individual director or many and then graph the host megabytes per second. Optionally, add a trend line to the chart.

If you are comparing directors, a snapshot chart may work as well.

Measuring Throughput by Port

- FE Director
 - ▶ Throughput in Mbytes per sec
 - ▶ % Busy / % Idle
 - ▶ Average IO size in Kbytes



EMC²

Copyright © 2014 EMC Corporation. All Rights Reserved.


Module 4: Analyzing Performance of Symmetrix Front-end Adapters 17

Throughput in megabytes per second can be measured for front-end ports.

The average I/O size for the port can also be shown. The average I/O size is often useful in detecting potential throughput issues. Note that this measure is an average value, calculated by dividing the throughput measure over the I/O per second measure for that time period. If your application is using a mix of 4 KB, 8 KB, and 32 KB I/Os, the calculated I/O size could be anywhere between 4 and 32.

What to Look For

- High Port utilization (~100%) especially when I/O sizes are large
- Imbalance among Ports



EMC²

Copyright © 2014 EMC Corporation. All Rights Reserved.

Module 4: Analyzing Performance of Symmetrix Front-end Adapters 18

Check for heavy port utilization when I/O sizes are large (>8 KB). Large I/Os tend to use the ports very efficiently, testing the limits of the port's bandwidth. When I/Os are large, the port throughput is more likely to become the bottleneck than the director I/O per second.

Imbalance is also a concern, since it always means that some channels are closer to their limits, while other have some idle time to spare. This might be a lower concern than raw utilization, however, since it is so difficult to balance both I/O per second and throughput in the open systems environment. This is discussed more on the next slide.

Solutions

- Reduce workload on any over-utilized ports
- Balance port throughput if possible
- Get faster ports (port speed)



If over utilization is the problem, then the obvious solution is to reduce the load on those ports. This can be done by moving some of the work to another port or to a completely different array.

As with nearly all parallel components in the Symmetrix environment, you would like to balance the throughput across all ports. This might be attempted by incrementally adding workload in a balanced way, as was discussed earlier. In actuality though, a typical open systems environment will make balancing by both I/O per second and throughput difficult. With dozens of hosts having different operating systems and applications all mapped to the front-end ports, the environment will probably be too chaotic to manually balance. Making some adjustments to perfect the port balance might completely throw off the director balance if the I/O sizes are not consistent across all applications (they never are in open systems environments). Take on a challenge like this only if port utilization is at a critically high level.

If workload cannot be reduced or rebalanced and utilization is still too high, getting faster ports might be a solution. Here, the port speed itself is what must improve. Increasing the director's processor speed will not affect the throughput in any way, but going from a 1 Gb port to a 2 Gb port will.

I/O Per Second and Throughput Check – Example

In a particular archive, the following measurements were recorded for a Fibre Director and its one active Port:

Time	I/O Per Sec	Dir % Busy	Through KB/Sec	Port % Busy	Average I/O Size (KB)
14:00	726	8.26%	3,166	2.55%	4.38
16:00	296	3.36%	13,667	11.01%	45.86

EMC²

Copyright © 2014 EMC Corporation. All Rights Reserved.

Module 4: Analyzing Performance of Symmetrix Front-end Adapters 20

In a particular archive, the values shown on this slide were measured for a Fibre director and its one active port (the other port was idle). Note that the I/O per second is highest at 14:00, and the corresponding director utilization is also highest at this time. Contrary to intuition, the port utilization is lower at this point, but higher at 16:00.

Why?

Remember that port utilization is based on the measure that affects port performance; throughput. Throughput was actually more than four times higher at 16:00.


How could this happen with a lowered I/O per second rate?

Throughput is larger when I/O size increases, so the I/O size must have gone up. Checking the numbers shows that I/O size went up by a factor of 10 at that time.

This shows how the port and director utilization measures can be independent of each other, and yet tied together by the I/O size.

Response time

- Length of time taken to process an I/O
- Symmetrix response time measured by Front-end Directors
- Host response time measured by Host Operating System



EMC²

Copyright © 2014 EMC Corporation. All Rights Reserved.

Module 4: Analyzing Performance of Symmetrix Front-end Adapters 21

Response time is another measure of performance that is often used in parallel with measures like I/O per second and throughput. Response time measures the total time taken to process an I/O. This is often very useful in detecting application slowdowns. Applications that make decisions on an I/O-by-I/O basis or serve human clients typically require low response times. Bulk data loading programs, like backup applications, are not affected by low response time.

Symmetrix response time is measured by the front-end directors. The time of the start of the I/O negotiation by the director to the time of the final acknowledgement is the Symmetrix response time. Therefore, these measures omit any time taken in SAN connectivity devices, HBA delays, or host queuing.

Host response times are measured by the host operating system. The time of the start of the application's request to the time of the operating system's final acknowledgement to the application is the host response time. This measure includes queuing, HBA and connectivity delays, as well as the entirety of the Symmetrix time.

Symmetrix response time can only be reported by Symmetrix-based tools. Host response times are reported by host-based tools: Windows Perfmon, UNIX sar, iostat, and many others.

Having two different ways of reporting response time is very useful, as the numbers can help isolate a problem. If the average Symmetrix response time is 23ms and the average host response time is 25ms, what is contributing the most to the delay? The Symmetrix is the problem here.


What if the Symmetrix time is 5ms and the host time is 25ms? Look for problems on the host or in the SAN network.

Symmetrix-Based Response Time Measures

000192606304 > Performance > Analyze > Diagnostic > 000192606304 > esx251_SG-BC

- Symmetrix Device

- ▶ Read response time (ms)
- ▶ Read miss time (ms)
- ▶ Write response time (ms)
- ▶ Write miss time (ms)



EMC²

Copyright © 2014 EMC Corporation. All Rights Reserved.

Module 4: Analyzing Performance of Symmetrix Front-end Adapters 22

Symmetrix front-end directors measure response time by individual device, so you will find the four measures listed on this slide in the Volume Metrics of Performance Viewer. Times for read, read miss, and write operations are currently available.

I/O Queuing




- If requested I/O rate is greater than the device response time (wait+service) will allow, host will queue I/Os for the device
- Multiple processes are required to queue I/Os
 - ▶ Multiple applications writing to a disk or filesystem
 - ▶ Multi-threaded applications writing to raw disks or filesystems



If the number of I/Os requested per second exceeds the number of I/Os that can be serviced by the device per second, the host will queue I/Os for the device. Queuing adds to the wait time, increasing the overall response time. High queues are generally considered indications of an I/O bottleneck. However, we will see later that some host queuing helps the Symmetrix DA processors optimize read misses by re-arranging the order of execution to minimize seek times.

Using Response and Queue Measures to Isolate Problems

- **High Service:** Storage subsystem slow
- **Low Service:** Storage subsystem fast
- **Wait > Service:** Host is sending more I/Os than storage can handle
- **High Queue:** Many host threads competing for resources



EMC²

Copyright © 2014 EMC Corporation. All Rights Reserved.

Module 4: Analyzing Performance of Symmetrix Front-end Adapters 24


The response time and queue measures can help isolate problems as being caused by either the host or the storage array. Service time measures the time required for the array to process an I/O; if it is the highest component of the response time, the storage is slowing the I/O rate. If service time is the lowest component of the response time, the storage subsystem is processing I/Os quickly.

Wait time increases when hosts are forced to wait for I/O resources to become available. If wait time is higher than the service time, then the host is sending more I/Os than the storage subsystem can handle. Although this situation could be improved if the array responded faster, all this situation really indicates is that the host requires a lot of I/O and is not doing much in the way of non-I/O activity (wait time is counted only when the process is waiting on I/O and has nothing else to do).

High queues happen in similar situations; the I/Os cannot be completed as fast as they are issued. Since queuing can only happen when multiple processes or threads issue I/Os in parallel, high queuing is also an indicator of heavy multi-tasking. High queuing and high wait times often happen at the same time.

VMware NMP Fixed vs. Round Robin

Fixed vs. Round Robin Multipathing in Vmware:



EMC²

Copyright © 2014 EMC Corporation. All Rights Reserved.

Module 4: Analyzing Performance of Symmetrix Front-end Adapters 25

This graph shows the dramatic difference that can be achieved when using Round Robin vs. Fixed in VMware native multipathing.

PowerPath Solution

- Measures I/O traffic on all paths
- When a single device can be accessed through multiple paths, PowerPath dynamically selects the least busy path
- Adjusts for spikes and changes in workload
- Provides an inherent failover capability among paths




Copyright © 2014 EMC Corporation. All Rights Reserved.

Module 4: Analyzing Performance of Symmetrix Front-end Adapters 26

PowerPath software can mitigate these issues by intelligently sending I/O down the least busy path. This ensures that the customer will consistently have the lowest overall channel utilization across all paths. PowerPath is an EMC host-resident software product that works with any storage system to deliver intelligent I/O path management, and intelligent I/O balancing. If storage is visible to multiple paths, PowerPath will intelligently balance the workload across the paths and will dynamically tune performance as workloads change. PowerPath also automatically detects and repairs server-to-storage path failures.

How PowerPath Works



EMC²

Copyright © 2014 EMC Corporation. All Rights Reserved.

Module 4: Analyzing Performance of Symmetrix Front-end Adapters 27

Without PowerPath, applications direct I/O to devices as needed. In the illustration shown on this slide, some of the devices have more I/Os in their queues, leading to longer delays on those devices. If devices that have more I/O are crowded onto one HBA, that channel will be utilized more than the others.

PowerPath balances I/O over multiple paths to the same device. Host applications write to devices, just as before. PowerPath drivers intercept the I/O, however, and directs it to the physical device that is attached to the least busy port.

Performance Viewer Roadmap

I/O per second by Front-End Director:

Metric	Issue
Analyze>Array>FE Director >KPI> Host IOs/sec	Are directors balanced?
Analyze>Array>FE Director >KPI> % Busy	Are directors over-utilized?
% Busy & Host IOs/sec	When does peak traffic happen?

Throughput by Front-end Port:

Metric	Issue
Analyze>Array>FE Director >KPI> Port 0 & 1 MBs/sec	Are ports balanced?
% Busy Port 0 & 1	When does peak traffic happen?

EMC²

Copyright © 2014 EMC Corporation. All Rights Reserved.

Module 4: Analyzing Performance of Symmetrix Front-end Adapters 28

Shown here are some KPIs for front-end adapters in Performance Viewer.

Performance Viewer Roadmap (Continued)

Response Time by Front-End Director

Metric	Issue
Analyze>Array>FE Director »KPI» Avg Read Response Time Avg Write Response Time	Is response time higher than demanded by application? Balanced? Use other measures to isolate issue.

EMC²

Copyright © 2014 EMC Corporation. All Rights Reserved.

Module 4: Analyzing Performance of Symmetrix Front-end Adapters 29

Shown here are some additional KPIs for front-end adapters in Performance Viewer.

Host I/O Limits

- Introduced in Enginuity 5876 Q4 2012 SR
- Provides the ability to limit application bandwidth and/or IOPS at the Storage Group level
 - ▶ Devices contained within an SG will not exceed the maximum IOPS and/or bandwidth set for the SG's quota (limit)
- Quota is defined on the SG, but is enforced at the Port Level
- Limits can be set on the Parent SG as well as on the Child SG (Cascaded Quotas)
 - ▶ Requires Enginuity 5876 Q2 2013 SR and above
- Quota is applied collectively to all devices in the SG
- By default, the defined quota is statically distributed between the directors contained in the Port Group
 - ▶ Dynamic redistribution of unused director quota can be enabled (disabled by default); Requires Enginuity 5876 Q2 2013 SR and above
- Limits can be defined using Unisphere for VMAX or SYMCLI



Copyright © 2014 EMC Corporation. All Rights Reserved.

Module 4: Analyzing Performance of Symmetrix Front-end Adapters 30

Host I/O Limits was introduced in Enginuity 5876 Q4 2012 SR. It provides the ability to limit the bandwidth and/or the IOPS consumed by devices in a Storage Group. The limits (also called Quotas) are set at the SG level, but are applied/enforced at the port level. By default, the limit is statically distributed equally among all the directors whose ports are part of the Port Group. For example, if a limit of 500 MB/s is set and there are two ports, FA:7E:0 and FA:8E:0, in the PG, then each director, 7E and 8E, will be allocated 250 MB/s for the SG that is using them. For the limits to be effective, the SG and PG must be a part of a Masking View. With Enginuity 5876 Q2 2013 SR, the allocation can be changed to dynamic. This is useful in cases where a director might fail. In this instance, the surviving director will get all of the limit that has been set on the SG. In the example shown on this slide, with dynamic distribution, if FA:8E was to fail, then FA:7E will get 500 MB/s as its limit. It should be noted that the limit does not guarantee a minimum bandwidth and/or IOPS. Limits can be set at the Parent SG level and/or the Child SG level.

Set Host I/O Limits

- Limit by MB/Sec or IO/sec (or both)

Set Host I/O Limits

Selected Storage Group: esx251_SG-BC

Set Host I/O Limit (MB/Sec) ☐ No Limit

Set Host I/O Limit (IO/Sec) ☒ 2000

Set Dynamic Distribution OnFailure ▼

Add to Job List ▼ Cancel Help

EMC²

Copyright © 2014 EMC Corporation. All Rights Reserved.

Module 4: Analyzing Performance of Symmetrix Front-end Adapters 31

The limit of 2000 IOPS is being set, which will be 1000 IOPS per port on both directors. It will be applied to all directors in the storage group.

Host I/O Limits – Example




EMC²

Copyright © 2014 EMC Corporation. All Rights Reserved.

Module 4: Analyzing Performance of Symmetrix Front-end Adapters 32

Here, you can see at about the 11:00 time frame where the Host I/O Limit was set. In this case, it only impacts the port 0 I/O since port 1 was already below the limit.

Host I/O Limits – Example (Continued)



EMC²

Copyright © 2014 EMC Corporation. All Rights Reserved.

Module 4: Analyzing Performance of Symmetrix Front-end Adapters 33

Graphing the Host IO Limit Avg Delayed Time (in MS) and % Quota Delayed IOs will shed light on how long the IOs are being delayed and what percentage of IOs are being delayed. Refer to the Technical Notes document, “Host I/O Limits for Symmetrix VMAX Family Arrays” for more information.

https://support.emc.com/docu43710_Host-I-O-Limits-for-Symmetrix-VMAX-Family-Arrays.pdf?language=en_US

Module 4: Summary

Key points covered in this module:

- Front-end architecture and scalability
- Key front-end measures
- Front-end performance issues



This module covered analyzing the performance of Symmetrix front-end directors.