

UNDERSTANDING HOW I/O WORKLOAD PROFILES RELATE TO PERFORMANCE

Stan Lanning
Principal Systems Engineer
Enterprise Storage Division Specialist
EMC Corporation
stan.lanning@emc.com



Table of Contents

Introduction	3
Why is an understanding of workload profile important?	3
Key Storage Workload Characteristics	4
IO Size	4
IOPS	5
Throughput.....	7
Response Time	9
Read / Write Ratio.....	9
Random vs. Sequential	10
Response Times by Media Type	12
You Say Potato, I say Pitato	12
Bits vs. Bytes	12
Decimal vs. Binary	14
Workload Profiles	16
What are Workload Profiles?	16
Workloads and Performance.....	18
How does a workload profile affect performance?	18
Workload Variations	18
Sustained Workload	18
Averages.....	18
Peaks and Bursts	19
Identifying Workloads Variations	20
Identifying Workload Profiles.....	22
Windows Perfmon utility	22
Linux / Unix 'iostat' command.....	24
VNX Monitoring and Reporting	25
Unisphere for VNX – Analyzer.....	26
Unisphere for VMAX – Performance Monitor.....	28
Conclusion	29

Disclaimer: The views, processes, or methodologies published in this article are those of the author. They do not necessarily reflect EMC Corporation's views, processes, or methodologies.

Introduction

When implementing a new application there is usually a requirements section or sizing guide that indicates minimum number of server processors or cores, the amount of memory, and storage capacity required. In more complex application systems such as Microsoft Exchange, SAP, or Oracle E-Business Suite there may be sizing calculators to help determine the number of servers, network bandwidth, database requirements, and so on. However, there are relatively few tools for storage performance sizing, and they tend to be specific to individual products or vendors. Plus, the inputs to these tools are often not well understood or simply not known.

As a storage solutions specialist involved in sizing IT infrastructure, one of the first questions I ask customers and sales teams is, “what are the IO workload profiles that the infrastructure needs to support?” Often, I receive a response of “100 Terabytes”, “2700 IOPS”, “350 Megabytes per second”, or simply “we don’t know”. Then I ask a series of questions trying to understand what applications the customer is running or planning to run, how data flows in and out of the system, and hopefully gather enough detail to construct the “IO workload profiles”. This information is necessary to design a storage solution that will meet performance and cost requirements.

What are “IO workload profiles”?

Why are they important?

How do we identify them?

How do they relate to performance sizing?

This article will explain the fundamentals of the characteristics of IO workload profiles – such as IOPS, MB/s, read/write ratios, response times, and queue lengths. It will also explore how different profiles can impact performance sizing, and some common tools that can be used to identify workload profile characteristics.

Why is an understanding of workload profile important?

If we do not understand the workload profile there is significant risk that the storage solution will not meet desired performance levels which can slow or even halt a business function. Or the solution may be “over sized” leading to higher than necessary acquisition, implementation, and operating costs.

Key Storage Workload Characteristics

There are many characteristics of an IO workload that affect performance – in fact, too many to cover in this article. We will focus on several key characteristics that describe and affect all storage devices, including:

IO Size

IOPS

Throughput

Response Time

Read/Write Ratio

Random vs. Sequential

A future article will go into more depth on additional aspects of workload profiles including “locality of reference”, “cache friendliness”, “burstiness”, and so forth

IO Size

In computing systems, the size of the Input or Output request (IO size) is typically expressed in Bytes (B), Kilobytes (KB), or sometimes Megabytes (MB).

The smaller an IO, the less time it takes to transmit. Similarly, the larger an IO, the longer it takes to transmit. However, there are advantages to using small and large IO sizes depending on the type of work being performed. Using too small or too large a size can lead to additional complications.

Here is an example:

Imagine your parents are moving to a new house and you need to move your collection of vintage comic books out of their basement before they are thrown away. Each comic book is similar to a file containing text and media content.

If you take one comic book at a time, there is little effort required to move each one and you can quickly move it from basement to the trunk or boot of your waiting car. However, you will have to make many trips to move the entire collection. This is similar to using a small IO size – very quick to move, but lots of round trips to complete the work.

If you pack 80 comic books into a large box and then move the box, it takes longer to pack, movement is slower due to the heavy box, and navigating the narrow basement

staircase may be more difficult. But the overall move may go faster – this is similar to large IO sizes – because lots of comic books are moved in each trip.

In most situations, the most efficient method involves packing the comic books into several medium size boxes because there will be fewer round trips than taking one comic book at a time, but you also avoid the long packing time and difficult maneuvering with a very large box. This is the compromise that most applications need to trade-off – moving data efficiently in the right size containers with as few round trips as possible.

Consider another example:

A roadway provides passage for cars and trucks of different sizes. In most countries the roadway or each lane has a standard width. Vehicles must comply with width and height limits in order to avoid obstacles such as surrounding traffic, overhead wires, bridges, and so on. Length is the primary variable in determining how much stuff – people and/or cargo – a vehicle can transport.

A car is generally short in length, width, and height – easier to maneuver and economical to operate, but does not hold very much cargo. A large truck is wider, taller, and much longer than most cars, and it can carry many times the amount of cargo.


The size of the vehicle does not necessarily determine how fast it can go. Most vehicles today can travel faster than the posted speed limit. Also keep in mind the size of the vehicle often affects how much work (fuel) is required to move the vehicle, but for now the analogy will avoid the complexities of acceleration, friction, and so on.



IOPS

IOPS measures the number of **I**ntput / **O**utput operations (or requests) **P**er **S**econd. This is sometimes referred to as requests per second.


Consider a road with cars and trucks going through toll booths. Assuming all vehicles are traveling at the same speed, the number of small cars that pass through the toll plaza in a given period of time will be much higher than the number of large trucks that pass through the toll plaza – mostly due to the length of the vehicle. Plus, the amount of cargo that passes through the toll booth will vary even more due to the carrying capacity of cars vs. trucks.



The number of vehicles that pass through the toll booth in one hour is analogous to the number of bytes that pass through an interface each second – IOPS.

Most roads have vehicles traveling in two directions, and the number of vehicles going in either direction is similar to the total IOPS of a storage or network interface. However, we may want to know how many vehicles travel in one particular direction in a given time period.

Four vehicles passing through within one second equals 4 IOPS – regardless of which direction they are traveling.



In network and storage environments, we often need to know the number of read and write requests (*inputs and outputs*) per second as separate values.

IOPS may be expressed together – total IOPS, or separately as read or write IOPS:


- IOPS = combined reads + writes per second
- rIOPS (or r/s) = reads per second
- wIOPS (or w/s) = writes per second

Throughput

Throughput measures the amount of work performed in a given time period.

Returning to the toll road example – while IOPS measures the number of vehicles that pass through the toll booths in an interval, throughput measures the amount of cargo that passes through the toll booths in a given period of time.

Assume the large truck and three small cars pass through the toll booth within one second – the total cargo is 128KB + 8KB + 8KB + 8KB within one second, or a throughput of 152KB/s.



In computer network and storage environments, throughput –sometimes called bandwidth - refers to the amount of data processed in one second, minute, or hour.

Here we introduce the critical three-way relationship - ***IOPS * IO Size = Throughput***. For example:

- An application processing 100 requests per second where each request is 4KB in size has a total throughput of 400KB/s
 - $100 \text{ IOPS} * 4\text{KB} = 400\text{KB/s}$
- A server with four HBAs – where each HBA processes 1,200 reads and 700 writes per second with an average IO size of 32KB – has an average throughput of 243,200KB/s
 - $(1200 \text{ reads} + 700 \text{ writes}) * 4 \text{ HBAs} * 32\text{KB} = 243,200\text{KB/s}$

If we know two of the values we can compute the third:

- A workload of 40 requests per second and 800KB/s throughput:
 - $800\text{KB/s} / 40 \text{ IOPS} = 20\text{KB}$ average IO size

Note that we do not know that all of the IOs are 20KB, only that the average is 20KB per IO for that measurement. In most cases, the average is good enough for sizing a solution.

However, if we do not know at least two of the three values we cannot accurately describe the basic parameters of the workload. For example:


- My application requires 15,000 IOPS
 - What is the IO size?
 - What is the throughput?

Response Time

The time required to send, process, and receive acknowledgement (ACK) for a request is known as Response Time (RT). A read request involves sending a request for data, processing time to retrieve the data, and returning the data payload. A write request involves sending a data payload, processing time to store the data, and returning confirmation (ACK) that the write has completed.

The size of a storage IO request impacts how long it takes to process and, in many cases, how much work is required to process the request. The larger the request (IO Size), the longer it will take to process – which increases response time.

Consider our toll road example. A small car can pass through the toll booth in less time than a large truck. Response time is analogous to the vehicle entering, paying the toll, and then exiting the toll both.



In the example above, the response time for a small IO request is 1 millisecond (1ms) and response time for a large IO request is 4ms.

- The average response time is $1\text{ms} + 4\text{ms} + 1\text{ms} / 3 = 2\text{ms}$ (avg. RT)

Note: if provided an average response time we cannot tell if all of the response times are the same. There may be peaks or “spikes” hidden in the average.


Read/Write Ratio

How many requests are reads compared to writes? A workload that is all reads may be expressed as 100% Reads, or 0% Writes. Another workload that has three write requests for every read request may be expressed as 75% Writes, 25% Reads, or a 1:3 Read/Write Ratio.

The read/write ratio is important because in most cases there is less effort required to process reads than to process writes. Thus, the more writes there are, the more work is done to process the writes. This is especially true for storage systems where a write request can require from 1.1x to 6x the amount of work that the same size read request may require.


Random vs. Sequential

Sequential IO is a sequence of requests accessing data blocks that are adjacent to each other:



Reference: http://en.wikipedia.org/wiki/File:Random_vs_sequential_access.svg

Random IO is the opposite, a sequence of requests that are not adjacent to each other and may not have any regular pattern:



Linear Media


Tape devices such as LTO, DAT, DDS, and even old 9-track reel-to-reel systems can store data in relatively compact space, typically at lower cost than hard drives or memory devices.

However, access to a specific piece of data requires loading the correct tape and then scanning or seeking to a specific location on the tape which can take seconds or even minutes to perform. Because of the nature of the linear storage format, tape drives are best used for sequential IO patterns such as backups and restores. However, tape media tends to break

down over time, and the combination of reliability concerns and slower performance have led to hard drive-based replacements for most backup and archive applications.

Rotating Media

Traditional hard drives contain one or more platters that rotate at high speeds, and one or more heads that perform reads and writes to a magnetic media on the surface of the platters. An actuator arm moves the head in order to reach different regions on the platter.



Each move of the head to a new position – known as “seek time” – takes time away from reading or writing data. For this reason, rotating media drives perform best with sequential IO where the data is adjacent and lined up requiring little to no head movement.

Traditional hard drives also tend to work very well with larger IO sizes. The best sequential access performance for a rotational hard drive typically occurs with IO sizes of 256KB to 512KB, and sometimes as large as 4MB for video applications.

Random Access Media (RAM)

In contrast, most forms of memory – DRAM, SRAM, NVRAM, and Flash – process random requests much easier because there is no head that needs to be moved before IO can be processed. In other words, memory devices do not incur seek as do rotational devices.

Memory devices also process IO requests much faster than hard drives or tape devices. The memory devices have lower response times and higher throughput, but they also cost much more on a \$/GB basis than hard drives or tape devices.

Response Times by Media Type

Recall that Response Time (RT) measures the full round trip time required to send, process, and acknowledge an IO request. The type of media being accessed has a significant effect on the response time, as does the IO size.

Most digital media – including RAM, Flash, and Hard Drives – measure response time performance for small block, random IO requests, typically 4KB or 8KB in size. The response time comparison between these device types is dramatic – sometimes one or two orders of magnitude difference:

Device Type	Average Response Time (ms)
Server SDRAM	Less than 0.001
PCIe Flash Card	0.02 – 0.10
Flash Drive	0.10 – 1.0
15K RPM HD	3 – 4
10K RPM HD	5 – 6
7.2K RPM HD	10 – 12

- Note: average response time for 4KB to 8KB request

You Say Potato, I say Pitato...

Before we go further, we need to address two frequent points of confusion regarding units of measure: bits vs. bytes, and decimal vs. binary.

Bits vs. Bytes

Bits are referenced with a lower-case 'b' and Bytes are referenced by an upper-case 'B'.

Eight bits equals one byte: 8b = 1B

FUN FACT

- 4 bits is called a “nybble” (aka “nibble” - a morsel or small bite)

File sizes and media capacity are measured in Bytes or multiples.

The storage and networking disciplines use both bits per second and bytes per second measures. Most interface metrics report Throughput as Bits per Second (lower case “b”/s):

- 10Gb/s Ethernet Network Interface Card (NIC)
- 16Gb/s Fibre Channel Host Bus Adapter (HBA)
- 40Gb/s Infiniband Host Channel Adapter (HCA)

Most capacity metrics refer to Throughput as Bytes per second (upper case “B”/s):

- 120MB/s – uncompressed speed of LTO-4 tape drive
- 477MB/s – 4K video playback with RGB 4:4:4 codec
- 1TB/hr – backup speed for large data warehouse

We can express Throughput in either form. For example:

- A modern 16Gb/s Fibre Channel Host Bus Adapter (HBA) is rated for 16Gb/s
 - $16\text{Gb/s} \div 8\text{b/B} = 2\text{GB/s}$
- The HBA can process 16Gb/s or 2GB/s

While individual network interfaces typically measure and report throughput in bits per second, it is very common to refer to the total workload in bytes per second – probably because I/O Size (in Bytes) times I/Os per second = Throughput in Bytes per second.

This distinction is important because storage and networking teams frequently need to work together and sometimes do not realize they are talking about different units of measure. A network interface may provide 10Gb/s, but if the application requires 2GB/s (16Gb/s) the system will need at least two 10Gb/s connections.

Decimal vs. Binary

You just purchased a new computer with a 2TB hard drive, and we know that 2 Terabytes (TB) equals 2000 Gigabytes (GB) of capacity. But after installing the operating system the computer shows the hard drive has roughly 1,862GiB of usable capacity. Where did the rest of the space go? And what is that 'i' in the middle of GiB?

Ironically they are the same capacity – the difference is units of measure: TB vs. TiB, GB vs. GiB, etc.

Look closer at the packaging on a hard drive and you will see the drive manufacturer lists capacity measured as 1 GB = 1,000,000,000 Bytes. Thus, the 2TB SATA drive is 2,000,000,000,000 Bytes.

The storage industry uses a decimal (base10) system for raw capacity using multiples of 1000 (10^3):

Decimal			
1 kilobyte (KB)	=	10^3	= 1000 ¹ = 1000 bytes
1 megabyte (MB)	=	10^6	= 1000 ² = 1000 kilobytes (KB)
1 gigabyte (GB)	=	10^9	= 1000 ³ = 1000 megabytes (MB)
1 terabyte (TB)	=	10^{12}	= 1000 ⁴ = 1000 gigabytes (GB)
1 petabyte (PB)	=	10^{15}	= 1000 ⁵ = 1000 terabytes (TB)
1 exabyte (EB)	=	10^{18}	= 1000 ⁶ = 1000 petabytes (PB)

In other words, 1 KB vs. 1 KiB is a ~2.4% difference, while 1 MB vs. 1 MiB is a ~4.63% difference – it is a compounding effect.

The discrepancy grows as the prefixes increase. Below is the table showing this effect:

Prefix	Decimal	Binary	Dec to Bin
Kilo	1,000	1,024	97.656%
Mega	1,000,000	1,048,576	95.367%
Giga	1,000,000,000	1,073,741,824	93.132%
Tera	1,000,000,000,000	1,099,511,627,776	90.949%
Peta	1,000,000,000,000,000	1,125,899,906,842,624	88.818%
Exa	1,000,000,000,000,000,000	1,152,921,504,606,846,976	86.736%

Reference: <http://en.wikipedia.org/wiki/Kibibyte>

Workload Profiles

What are Workload Profiles?

In the most generic sense, a workload profile describes the behavior of an application, or portion of an application with respect to how it performs work.

A storage workload profile describes several aspects of how an application accesses data on a storage device – such as a hard drive, USB stick, memory (e.g. DRAM or NVRAM), or external storage array (e.g. SAN, NAS, etc.).

The characteristics of a workload profile affect how much effort is required to perform the work. Typically, less effort required means less equipment and lower cost to provide the desired performance. Conversely, more effort typically means more equipment and higher cost.

Workload profiles for storage systems are typically described and grouped by key “Workload Characteristics” – i.e. attributes they have in common:

- Are the IO requests large or small?
- Are the IO requests reads, writes, or mixed reads and writes?
- Is the access pattern sequential or random?

Often, applications have multiple workload profiles. For example, an order processing system involves several operations:

- Order Entry – primarily data input (**writes**) of short line items (**small**), often many unrelated orders being processed at the same time (**random**)
- Packing and Shipping – data input and output to read order details, update inventory and shipping records (**mixed read/write**)
- Reporting and Data Backups – primarily data output (**reads**) of grouped information (**large**), often using indexes or tables scans (**sequential**)

Workloads and Performance

How does a workload profile affect performance?

Consider the following guidelines:

1. Large requests take longer to process than small requests.	<ul style="list-style-type: none">• The larger the IO size, the more time required to send and process the request.
2. In most situations, writes require more work than reads.	<ul style="list-style-type: none">• The more writes that a workload performs, the more work a storage system needs to do.• Most media types have higher response times for writes and often have additional overhead due to data protection, such as RAID.
3. Some storage technologies perform better with random IO patterns. Others perform better with sequential IO patterns.	<ul style="list-style-type: none">• Traditional hard drives generally perform best with sequential IO patterns by avoiding moving the drive head (seek) to a new location.• Memory and Flash devices tend to perform better with random IO patterns because there is no head movement and IOs can be parallelized to multiple cells.

Workload Variations

Sustained Workload

Activity or a level of performance that is maintained for several minutes or hours is known as a sustained workload. The importance of the sustained workload is that it is the best indicator of activity that lasts for a meaningful period of time and accurately represents a level of performance that needs to be supported by infrastructure.

When sizing for performance we typically focus on Sustained Peak. This represents a performance level that must be maintained for several minutes or a few hours without the extreme highs from bursts or lows of inactivity that can be hidden by averages.

The key focus of performance sizing is to ensure that a system can support sustained workload peaks without significant performance degradation.

Averages

A workload average typically represents the arithmetic mean for a set of values over a known time period. The average gives us an idea of the general behavior over time, but it does not provide any detail about the lows and highs (peaks) in a time series. An average can mask very

low or no activity that may be planned – such as maintenance work. The average can also hide very high activity or short bursts. By definition, averages change based on the time period that is examined.

A big problem with averages is that picking the wrong time period – too short, too long, or wrong start and stop times – can result in an average that is far different from the important period for applications and business activity. Often, it is not known when the important activity is occurring vs. activity that is not important.

The main risk of sizing to a workload average is that the storage infrastructure may not be able to deliver adequate performance during workload peaks.

Peaks and Bursts

A workload peak represents the highest level of performance achieved by an application, device, or infrastructure during a period of time. There can be separate peaks for different workload characteristics. For example:

- Peak IOPS during month end order processing
- Peak Throughput (MB/s) during database backups
- Peak Response Time during nightly batch processing

Very short peaks are sometimes referred to as bursts, typically lasting less than one second or as long as a couple of minutes. In storage and network environments it is common to observe very short bursts of activity that can be several times the sustained peak workload, and may be orders of magnitude higher than the average workload.

The primary question to consider is: How long does the burst of activity last?


If a burst is very short, it may be a distracter or anomaly that can be ignored for sizing purposes. The risk of sizing to very short bursts is that they do not accurately represent the required level of performance and can lead to a solution that may be too costly or too complex to operate and maintain.

In most cases, a solution should be sized to support the sustained peak workload which ensures that performance is acceptable during the busiest sustained period of work.

Identifying Workloads Variations

The example below is a JetStress test that simulates a Microsoft Exchange workload. The graph comes from Unisphere[®] Analyzer showing total IOPS for a VNX[®] array that is providing SAN storage for the test environment. The JetStress test runs in two phases:

4. **Mailbox Tests** – simulating messages being sent and received. This portion performs a mix of random reads and writes and best approximates performance that users will experience.
5. **Database Verify** – validates the integrity of the Exchange DB files. This simulates a batch process that typically runs at night or between time zones and has less impact on users.



- During phase 1 of the test there is an initial peak over 18,750 IOPS that lasts for three minutes – outlined in **red**.
- The remainder of phase 1 – outlined in **green** – shows an average of approximately 9,500 IOPS for nearly one hour, with short bursts and lows during that hour.
 - The short peak at the beginning of the test is due to all of the JetStress threads starting and is not normal workload.
- During phase 2 of the test we can see a short burst – outlined in **purple** – to over 30,000 IOPS, but it occurs only once during a sustained peak of ~26,000 IOPS – outlined in **orange** – that lasts for approximately fifteen minutes.

- In most scenarios it is acceptable to ignore the burst of 30,000 IOPS, and instead focus on sizing the solution to support ~26,000 IOPS during the DB Verify portion of the test.

Typically, the storage solution will be configured in order to satisfy the average of 9,500 IOPS. In many scenarios, sizing of the storage will include a growth factor of 10%, 20%, or possibly more to account for business growth over time.

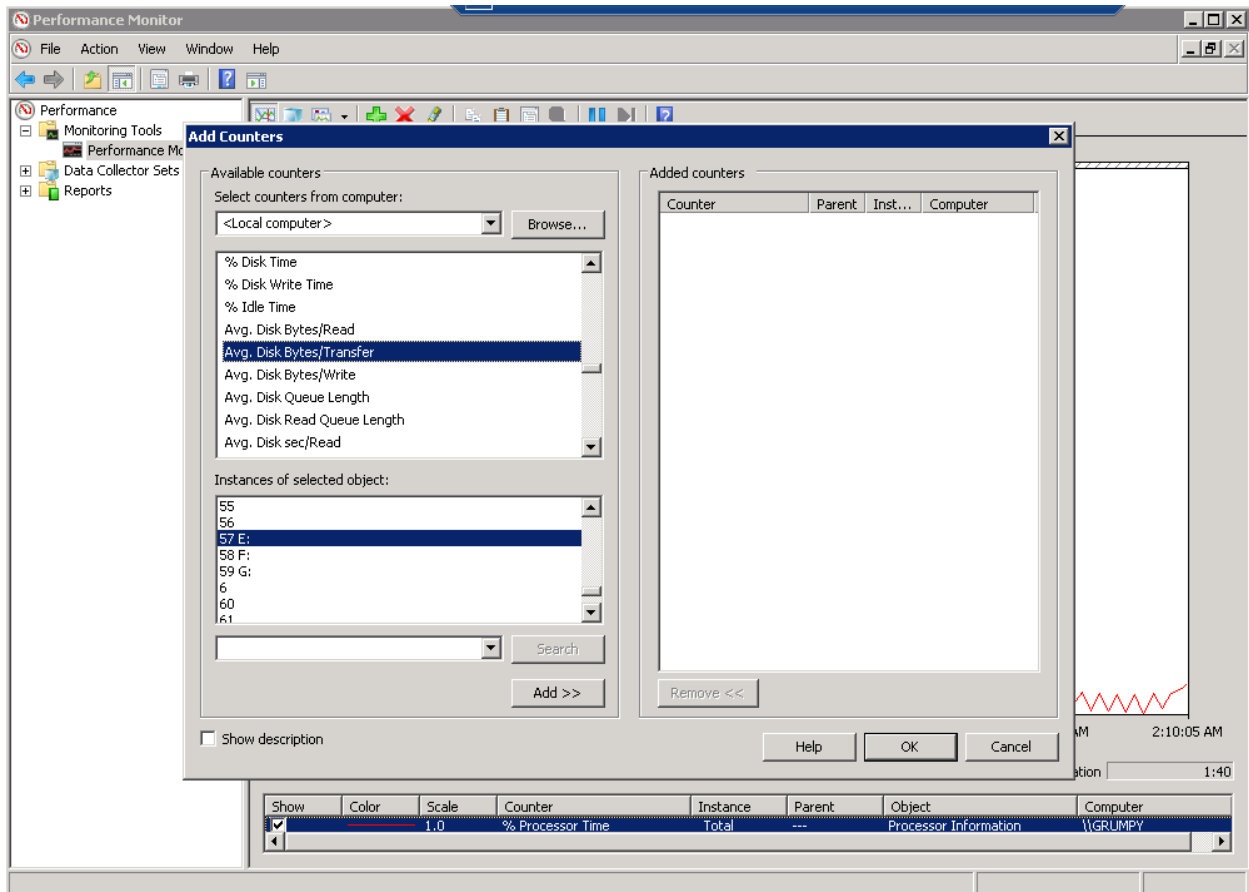
Identifying Workload Profiles

Windows Perfmon utility

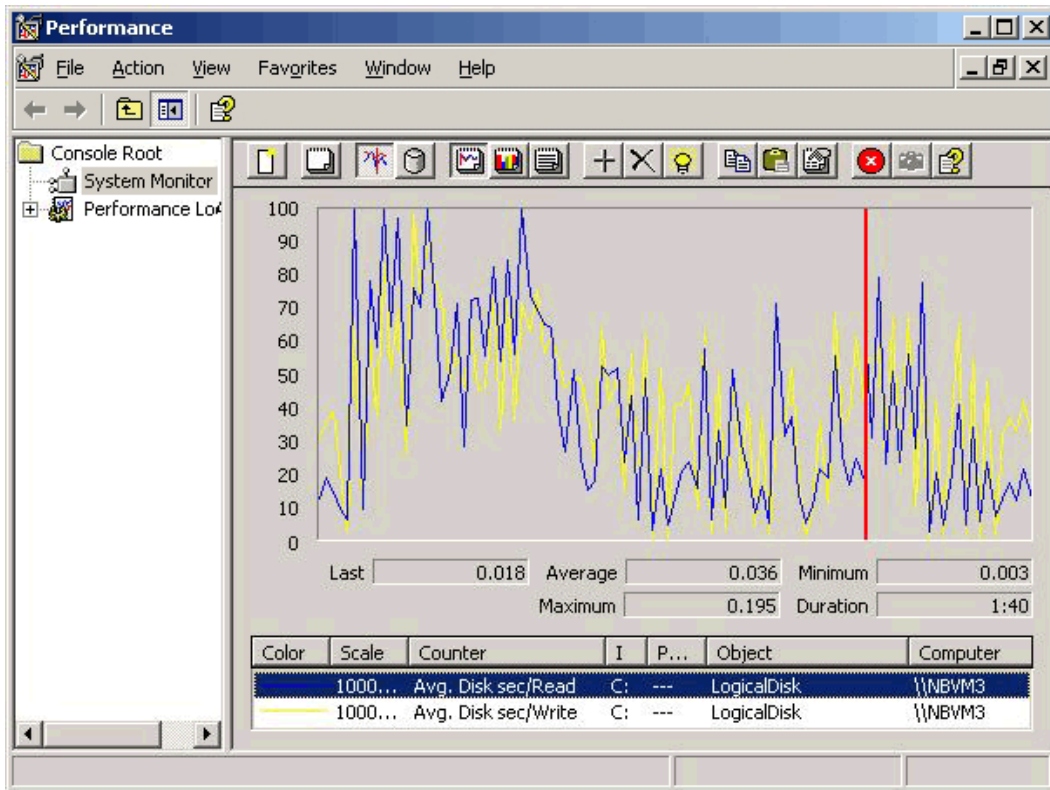
Perfmon is a utility included in most versions of Microsoft Windows operating system. It can report on a variety of metrics for CPU, memory, network, disk, and application resources.

For storage environments, there are several counters available for individual drives, or aggregate counters for all drives connected to a server. The image below shows an example of some key metrics in the Physical Disk category:

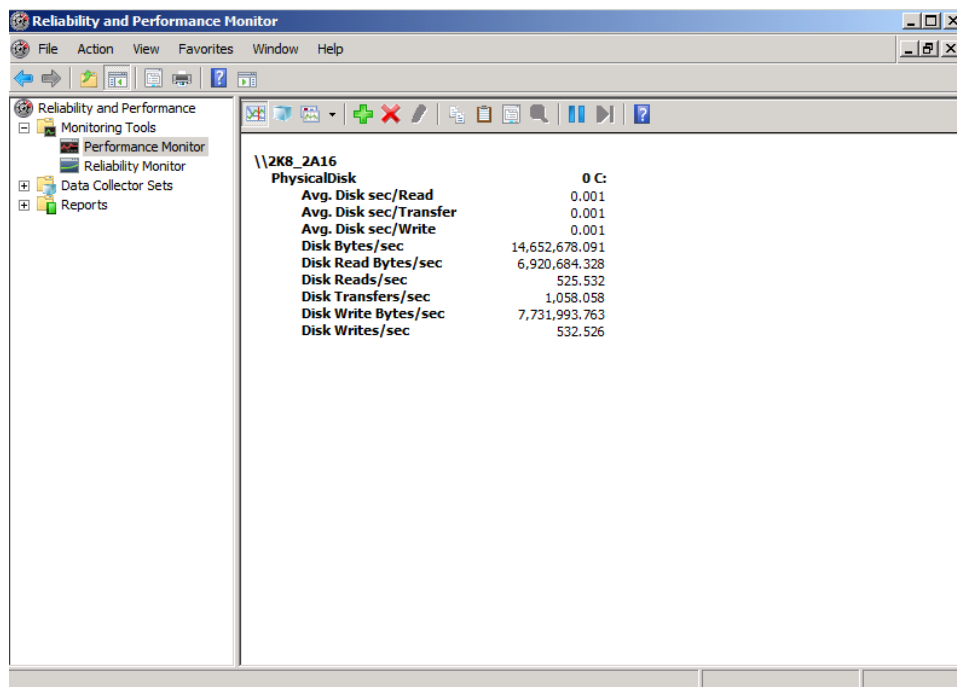
- IO Size = Disk Bytes/Transfer (also: Bytes/Read, Bytes/Write)
- IOPS = Disk Transfers/sec (also: Reads/sec, Writes/sec)
- Response Time = Avg. Disk sec/Transfer (also: sec/Read, sec/Write)



Perfmon metrics can be graphed over a period of time:



Or show current values in text format:



Linux / Unix 'iostat' command

The *iostat* command available in several Linux and Unix distributions provides performance metrics for CPU, memory, network, and storage. The utility is invoked via command line with an interval in seconds, and optionally the number of intervals to report. There are also optional flags to format the metrics in easier to read units such as KB or MB. For example, the following will run *iostat* every 10 seconds for 5 minutes:

```
iostat -xk 10 5
```

The output format varies from version to version, but they generally report the following metrics – averages over the time interval specified in the command:

- Read IOPS = r/s, and Write IOPS = w/s
- Read Throughput = rKB/s, and Write Throughput = wKB/s
- IO Size = avgrq-sz
- Response Time = await

The example below shows performance comparison of Fibre Channel (FC) drives vs. Flash drives in a storage array – as observed from the Linux host.

The image displays two terminal screenshots of the *iostat* command output. The top screenshot is for an FC disk, and the bottom is for a Flash disk. Red boxes highlight key metrics in both. To the right of each screenshot is a blue callout box summarizing the performance metrics.

Device	rrqm/s	wrqm/s	r/s	w/s	rKB/s	wKB/s	avgrq-sz	avgrq-ws	await	svctm	util
FC_1	0.00	1.00	131.00	55.00	1046.00	616.00	17.89	9.29	52.76	5.26	97.75
FC_2	0.00	0.00	109.00	74.50	876.00	344.25	16.85	5.57	36.31	4.43	92.45
FC_3	0.00	0.00	170.00	70.00	1424.00	562.00	16.02	11.47	45.11	3.62	96.90
FC_4	0.00	0.00	181.00	70.50	1449.00	572.00	16.06	9.97	39.57	3.84	96.55
FC_5	0.00	0.00	168.00	74.50	1344.00	616.00	16.16	8.55	38.60	3.92	95.10
FC_6	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00
FC_7	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00
FC_8	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00

Device	rrqm/s	wrqm/s	r/s	w/s	rKB/s	wKB/s	avgrq-sz	avgrq-ws	await	svctm	util
FLASH1	0.00	0.50	746.50	415.00	5964.00	2914.00	14.11	2.24	1.78	0.58	70.80
FLASH2	0.00	1.30	648.50	404.50	5360.50	1889.25	13.51	2.36	2.39	0.67	71.50
FLASH3	0.00	0.00	752.00	397.00	6016.00	3406.00	16.40	2.39	2.09	0.52	59.40
FLASH4	0.00	0.00	775.50	392.00	6204.00	3420.00	16.49	2.73	2.34	0.53	61.90
FLASH5	0.00	0.00	777.00	408.50	6216.00	3800.00	16.90	2.70	2.28	0.50	58.75
FLASH6	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00
FLASH7	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00
FLASH8	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00

FC Disk	
Reads/s	767
Writes/s	303
iops	1070
Queue	10
Wait	43 ms
Svtn	3.8 - 7 ms
Util	92-97 %

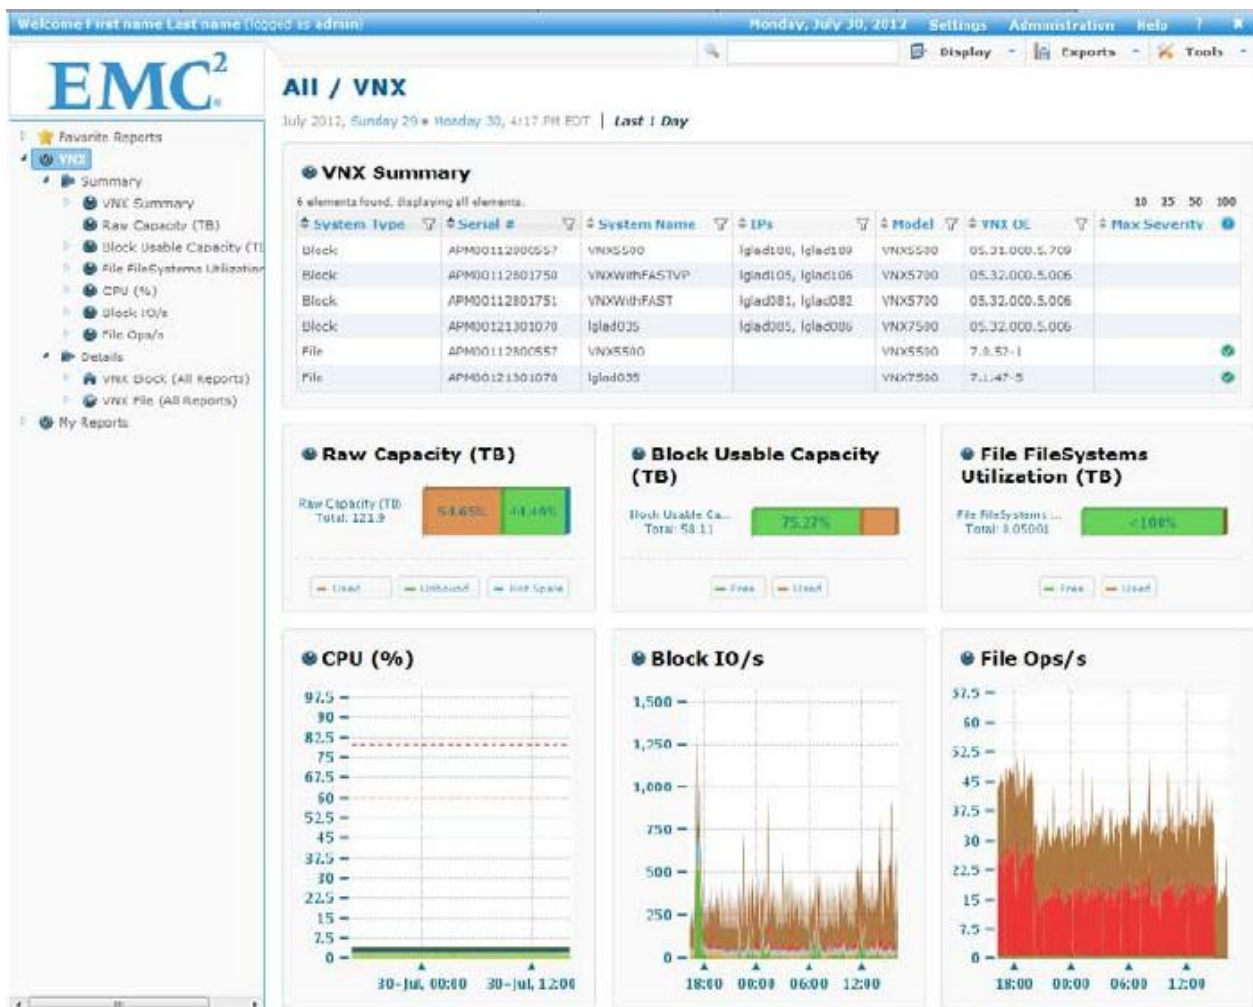
Flash Disk	
Reads/s	3718
Writes/s	2116
iops	5834 (5.5x)
Queue	2-3
Wait	2 ms (5%)
Svtn	< 0.7 ms (10x)
Util	58 - 71%

Reference: <http://bartsjerps.wordpress.com/2011/03/04/io-bottleneck-linux/>

VNX Monitoring and Reporting


Monitoring and Reporting (M&R) for VNX arrays is a lightweight version of EMC Storage Resource Management (SRM) Suite. VNX M&R focuses on capacity and performance metrics for VNX arrays with historical trends and alerting for one to several VNX arrays.

Below is an example of the VNX summary reporting from Monitoring and Reporting package. Available metrics include IOPS, Throughput, Response Time, and several others for SAN and NAS functionality.

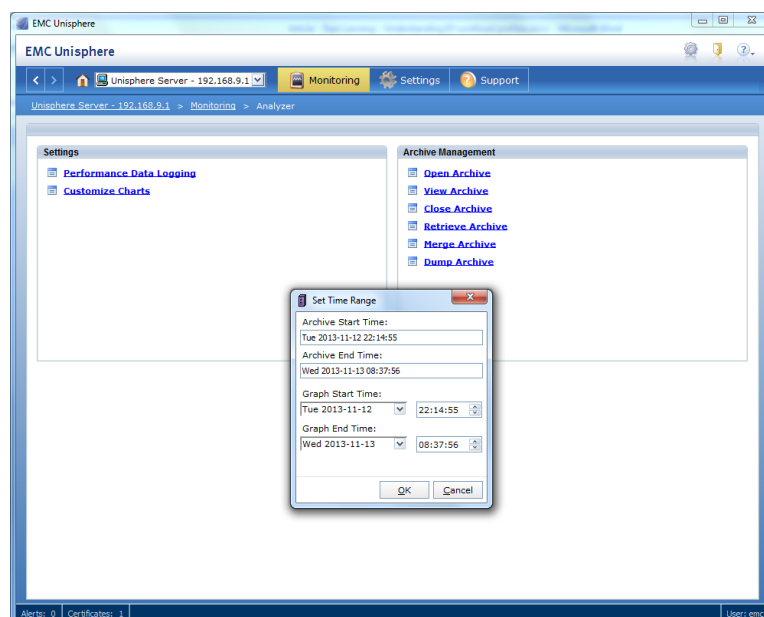


Unisphere for VNX – Analyzer

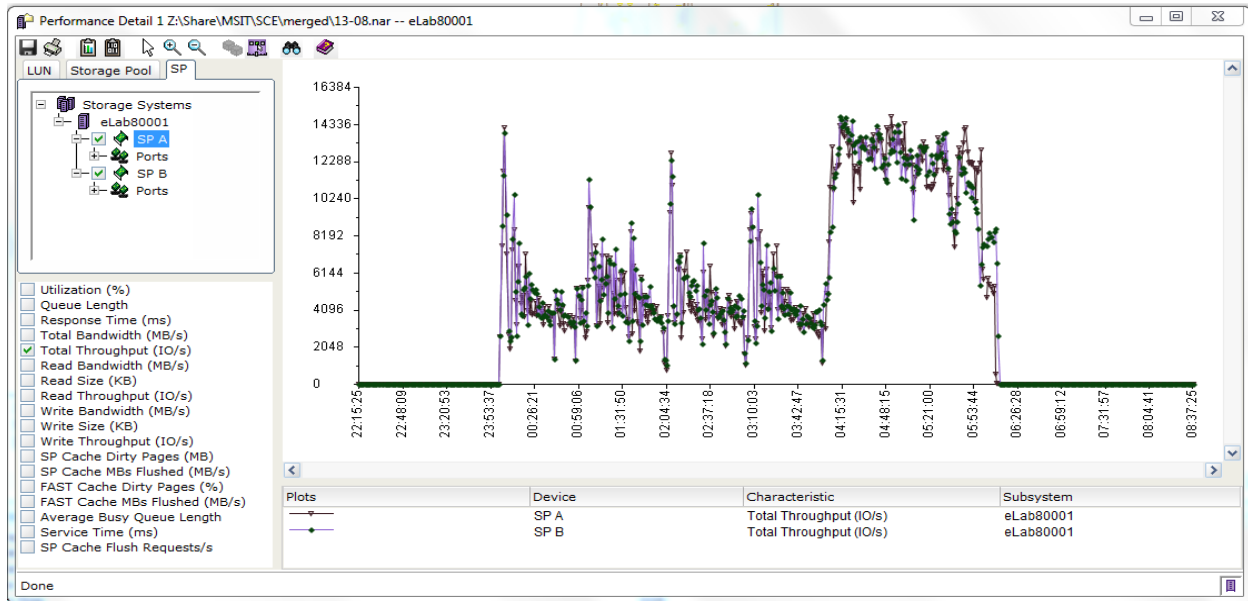
Unisphere Analyzer is a performance reporting feature for EMC VNX arrays. It is launched from within the Unisphere for VNX Graphical User Interface (GUI) and can provide near real-time or historical views of a wide variety of array performance metrics. Access Unisphere Analyzer via the Monitoring -> Analyzer menu option.



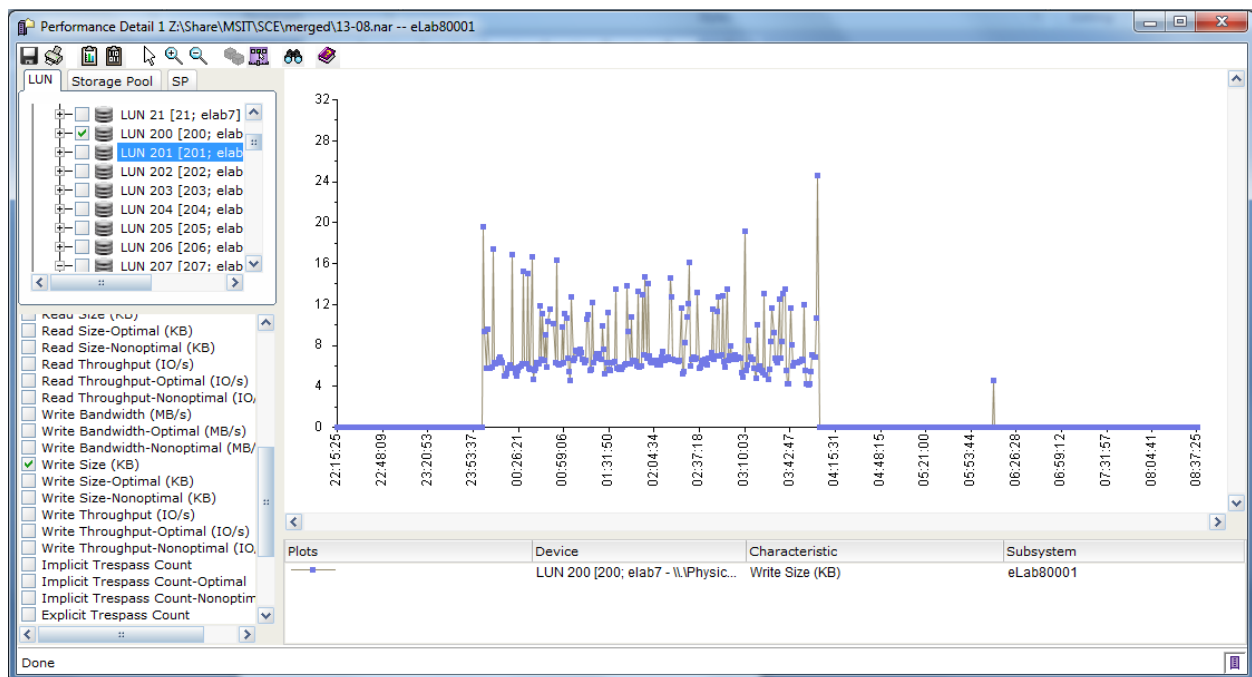
Historical metrics are collected in NAR (compressed archives) or NAZ (encrypted archives) files and later reviewed within Unisphere Analyzer. Here is an example opening a NAR file – a window pops up to show the date and time frame of the metrics within the file.



Several views are available including Performance Overview, Performance Summary, and Performance Detail. Below is the Performance Detail view showing total IOPS for each of the Storage Processors (SPs) on a VNX array.



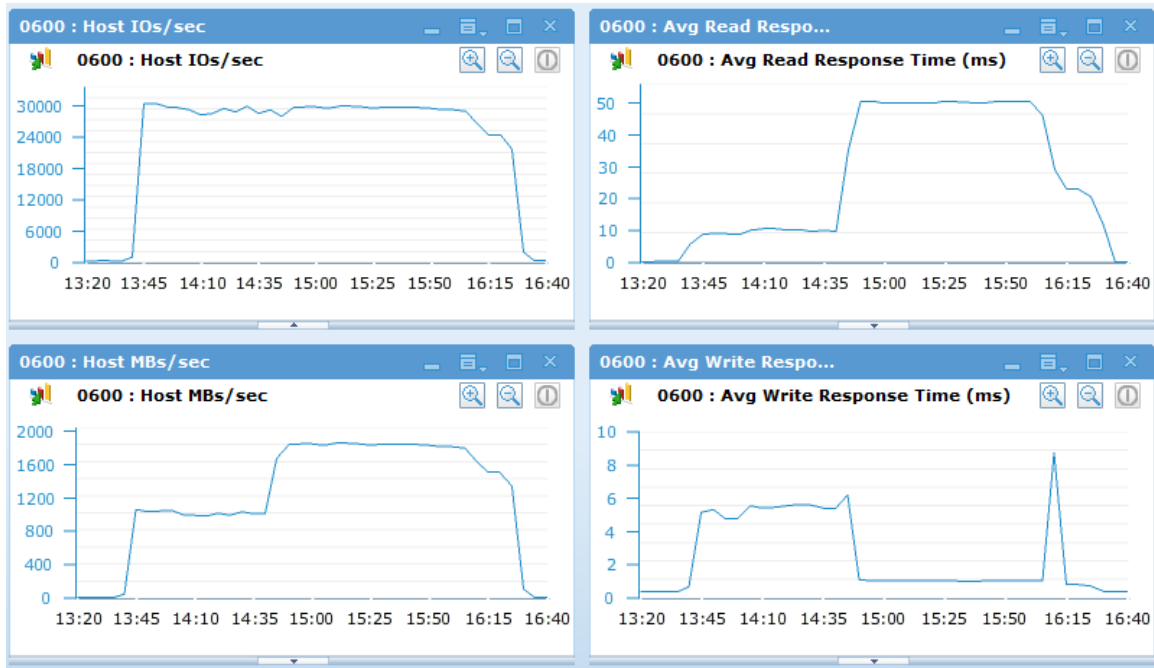
Unisphere Analyzer can drill down to individual LUNs and disks in the array. Here is an example of write IO size during a 4-hour application test – average is 8KB with short bursts as high as 24KB.



Unisphere for VMAX – Performance Monitor

Unisphere for VMAX includes a performance monitoring section that can view near real-time and historical performance metrics for one or more Symmetrix VMAX arrays. The graphs below show several key metrics during a test workload preparing an array to run Microsoft Exchange.

- IOPS = Host IOs/sec
- Throughput = Host MBs/sec
- Response Time = Avg Read Response Time (ms), and Avg Write Response Time (ms)



Conclusion

It is important to understand the key characteristics of an application or process – the workload profile – which includes IO size, IOPS, Throughput, Response Time, Read/Write Ratio, and several other metrics. By understanding the workload profiles on a host or storage system we can ensure that solutions are correctly sized in order to optimize acquisition and operating costs for computing infrastructure.

For more information on performance data gathering and reporting reference:

- EMC VNX Monitoring and Reporting Demonstration Video
 - <http://www.youtube.com/watch?v=SggOtZcqn-A>
- EMC Unisphere Analyzer for VNX Demonstration Video
 - http://www.youtube.com/watch?v=yCMZ_N7-p7A
- EMC Unisphere for VMAX Product Guide
 - https://support.emc.com/docu46997_Unisphere-for-VMAX-1.6-Product-Guide.pdf?language=en_US

EMC believes the information in this publication is accurate as of its publication date. The information is subject to change without notice.

THE INFORMATION IN THIS PUBLICATION IS PROVIDED “AS IS.” EMC CORPORATION 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 EMC software described in this publication requires an applicable software license.