
ECC V6.0 Performance Manager Solutions Guide

GDSE-08012-SG

Project Certification

Version1.0

Page 1 of 69

1 Version Control

Version	Date	Author	Reviewed	Comments
Draft 0.0		Colin N Bradley	Vivek Jain	
Draft 0.1	27 August 2008	Colin N Bradley	Vivek Jain	Updated document after feedback from Vivek
Draft 0.2	03 September 2008	Colin N Bradley/Vivek Jain		
RFC 1.0	26 th October 2008	Colin N Bradley/Vivek Jain		RFC

Legal Notice

Information contained herein is for Internal Use and may be used only for business purposes authorized by Citigroup.

CITI, CITIBANK, CITICORP, CITIGROUP and the Umbrella Device are trademarks and service marks of Citicorp and are used and registered throughout the world. Citibank, Citicorp, Citigroup and their subsidiaries also claim rights in certain other trademarks and service marks contained in these pages.

Contact

Please address any comments or questions regarding the information contained in this document via email to the “*GT Global Distributed Storage Engineering” distribution list or alternatively to gdse@citigroup.com.

Latest Versions

It is recommended that the latest version of the document be used from the GDSE website <http://engineering.citigroup.net/gdse/>. Newer versions may include certification on additional platforms, or comments on issues discovered after deployment.

2 Table of Contents

ECC V6.0 Performance Manger Solutions Guide	1
1 Version Control	2
2 Table of Contents	4
3 Introduction	8
3.1 Overview	8
3.2 Certification Scope	8
4 Performance Manager Architecture	9
4.1 Functionality Overview	9
4.2 Architecture Overview	9
4.2.1 Performance Manager Viewer	9
4.2.2 WLA Archiver Agent	10
4.2.3 ECC Agents	12
5 Sizing of the Performance Manager Environment	13
5.1 Sizing Case Study	15
5.1.1 London ECC Environment	15
6 Best Practice Recommendations for Deploying Performance Manager	20
7 Installation of Performance Manager	21
7.1 Installation Pre-Requisites	21
7.2 Install Performance Manager	22
7.3 WLA Archiver System Requirements	22
7.4 Installing the WLA Archiver	23
7.5 SYMACL	24
8 Set-Up of WLA Data Collection Policies	25
8.1 Pre-Requisites for the set-up of the WLA Data Collection Policies	25

8.2	WLA Daily DCP	26
8.2.1	WLA Daily DCP Set-Up	26
8.3	WLA Revolving DCP	29
8.3.1	WLA Revolving DCP Set-Up	30
8.4	WLA Analyst DCP	33
8.4.1	WLA Analyst DCP Set-Up	33
9	Performance Manager	36
9.1	Launching Performance Manager	36
9.2	Loading a BTP File	37
9.3	Performance Manager View	40
9.3.1	The Views pane	40
9.3.2	The Analysis toolbar	40
9.3.3	The Data view pane	40
9.4	Data Views	42
9.4.1	Factory data views	42
9.4.2	Public data views	42
9.4.3	Private Data Views	42
9.5	Automated Reports	42
9.6	Working with Graphs	43
9.6.1	Ad-hoc time graphs and histograms	43
9.6.2	Trend graph	43
9.6.3	Other graph representations	43
9.6.4	Creating Ad-hoc time graphs	43
9.7	Performance metrics	44
9.7.1	Vital signs	44

Appendices.....	45
9.8 Appendix A - Gathering Sizing Data for your Environment.....	45
9.8.1 Storage Array Sizing	45
9.8.2 McData Switch Port Sizing	47
9.8.3 Host Sizing.....	47
9.9 Appendix B - Symmetrix Metrics.....	51
9.9.1 Symmetrix Device Metrics	51
9.9.2 Symmetrix disk metrics	56
9.9.3 Symmetrix disk director metrics	57
9.9.4 Symmetrix host director metrics	59
9.9.5 Symmetrix port metrics	60
9.9.6 Symmetrix SRDF metrics	60
9.9.7 Symmetrix SRDF/A metrics	61
9.9.8 Symmetrix system metrics.....	62
9.10 Appendix C – Switch Metrics	64
9.10.1 Connectivity device port metrics	64
9.11 Appendix D – HDS Metrics	65
9.11.1 HDS Device metrics	65
9.11.2 HDS port metrics	66
9.12 Appendix E - Host metrics.....	67
9.12.1 Host CPU Metrics	67
9.12.2 Host device metrics.....	67
9.12.3 Host logical volume metrics.....	67
9.12.4 Host HP device metrics	68
9.13 Appendix F – HDS Metric Collection	68

9.13.1	HDS LUN metrics	68
9.13.2	HDS/Hitachi port metrics	69

3 Introduction

3.1 Overview

EMC ControlCenter (ECC) Performance Manager previously known as Workload Analyser (WLA) is an individually licensed and installed component of the ECC infrastructure. Performance Manager allows the collection of performance data from “objects” within the ECC managed environment. A managed object is a generic term used to include storage arrays, hosts, virtual machines, databases, fabric switches etc. The collection of these performance metrics from the environments allows the storage administrator to analyse performance metrics over a period of time. The Performance Manager interface can be launched directly from the ECC console or alternatively, if installed, can be launched from the Start menu of the ECC console server. A web access is also provided where a small selection of automated reports are available for view.

The performance data used by Performance Manager is gathered by the many ECC agents deployed in the ECC environment. For each managed object in the environment there is an applicable agent. As part of the Data Collection Policy for each agent performance data is captured and sent to a special Performance Manager agent known as the WLA¹ Archiver agent. The WLA Archiver agent acts on the sent data and converts this data into a file format² that can then be read by Performance Manager.

The real world uses of Performance Manager data can be arranged into two separate streams:

- Proactive Performance Management
- Reactive Performance Management

Within Citi Performance Manager has already had specific benefit in providing the base data required to properly size SRDF/A environments in EMEA. It also has been used to tackle some performance issues on the EMC DMX arrays as experienced by the business.

3.2 Certification Scope

The certification scope is detailed in the document ECC V6.0 Performance Manager Certification Project Scope GDSE-08012-PS.

¹ Even though Workload Analyser has been re-marketed as Performance Manager some legacy references are still made to WLA.

² The file is known as a *BTP* file.

4 Performance Manager Architecture

4.1 Functionality Overview

Performance Manager is an optional install component of the EMC ControlCenter infrastructure. As already mentioned it is separately licensed therefore a license must be obtained from EMC so that this software can be used.

4.2 Architecture Overview

Fundamentally Performance Manager is a separate piece of software that relies on ECC agents to provide it with the performance metrics. The Performance Manager application is comprised of a number of components as shown below:

- Performance Manager (Viewer)
- WLA Archiver Agent
- ECC Host agents

The following sections deal with each of these components in turn.

4.2.1 Performance Manager Viewer

The Performance Manager Viewer (PMV) allows the ECC administrator to review the collected performance statistics graphically. Analysis of the performance data is the responsibility of the ECC administrator, performance viewer will not inform you of performance bottlenecks. It will simply present the data graphically and the interpretation of the data lies with the administrator.

When launched the PMV queries the ECC Repository in order to determine the servers where WLA Archiver has been installed and is running. Once the PMV obtains the WLA Archiver server details further communication is handled directly with the WLA Archiver server via a HTTP request to the WLA Archiver (WlaHttpServer)³. In scenarios where multiple WLA Archivers are installed in the ECC environment, PMV will connect to each WLA Archiver in turn and retrieves their providers.xml file which is used to drive the PMV selection dialog.

The PMV allows the ECC administrator to select a .btp file of the managed object they wish to review and to load this .btp file into the PMV application. Once the file has been loaded the administrator can select from a multitude of different views and metrics and allow them to build up a picture of the performance activity taking place on the managed object, see Figure 1.

³ Details of this communication can be found in the WlaMonitor.log





Table 1: Performance Viewer

It is worth mentioning at this point that the level of detail that is on view within performance viewer is dependent on the managed object that is being analysed. For example the Symmetrix/DMX statistics are much more detailed than the switch or host statistics.

4.2.2 WLA Archiver Agent

The WLA archiver is the engine behind the performance manager application. It is deployed as an agent from the ECC console in the typical fashion i.e. pushed onto a server that has master/host agents installed. The WLA archiver is unique in the way in which it communicates with the rest of the ECC infrastructure and agents. Typically an agent (symm/sdm, host, FCC etc.) will communicate with the ECC infrastructure store servers or directly with the ECC Server repository directly depending on what action or information is being transferred. However with performance data the agents communicate directly with the WLA Archiver and do not interact with the stores or repository.

**Table 2: WLA Archiver Data Flow**

The data is transferred from the host agent⁴ to the WLA Archiver in ASCII format. These files are known as .ttp files and presented in the following format:

- 20080619_xx.ttp

Where xx is an incremental number in sequence of when the file was sent from the agent to the WLA archiver i.e. for a 24 hour period you would expect to see files numbered 20080619_00.ttp → 20080619_xx.ttp. These files are held in the <Interval> directory under the managed object i.e

EMC → WLAARCHIVE → <Managed Object> → <Managed Object ID> → Interval

Once the data for a specified time period has been transferred to the WLA Archiver server the WLA Archiver agent converts the .ttp files to a single .btp file. The .btp file can be loaded into Performance Viewer to analyse the performance statistics.

Having multiple WLA Archivers deployed in an environment allows for the sharing of workload between the archivers. However, the method of workload sharing is rudimentary at best. It is very important to work out the number of WLA Archivers in the environment prior to enabling any of the policies. The reason

⁴ I am using the term host agent generically here to cover all managed object agents.

being is that if 1 WLA Archiver is deployed in the environment and the WLA daily policy is enabled for all MO's in the environment then that 1 WLA Archiver will attempt to manage all the MO's. Adding an additional archiver does not mean that there will be a re-distribution of workload between the 2 Archivers.

This reinforces the need to properly understand the number of MO's that will be managed in the environment.

4.2.2.1 How a BTP file is Created

The BTP file is read by Performance Manager and contains all the performance information for the related MO. BTP files are generated once a day at 00:00 for each MO, the generation of these files is based on the already discussed TTP data file being available. The process for the creation of the BTP file is as follows:

- Each data provider transfers .ttp data to the WLA Archiver at pre-defined intervals i.e. for the WLA Daily_00 DCP for Symmetrix the collection schedule is set for 15mins. Therefore every 15mins ttp information will be sent to WLA Archiver.
- On the Archiver server in the EMC → WLAARCHIVE → <Managed Object> → <MO #> → interval. A .ttp file of the format <yyyymmdd.ttp> is appended with the incoming ttp data from the data provider.
- This .ttp file is converted to a .btp file at 00:00 each day. This then becomes the Daily btp file.
- At the last day of an archiving week one point per every shift is added to Weekly.btp. Every point is produced by averaging the corresponding data in Daily.btp over the last week.
- At the last day of an archiving month one point per every shift is added to Monthly.btp. Every point is produced by averaging the corresponding data in Daily.btp over the last month.

The revolving data collection also produces a .btp file as does the analyst collection however these do not follow the same process for creation.

4.2.3 ECC Agents

In the ECC environment there is the concept of the *managed object*. The physical and logical elements managed by ControlCenter are known as managed objects—they include storage arrays, hosts, virtual machines, switches, databases, and fabric zones, etc. For each managed object within the ECC environment there is a corresponding master agent and an *data provider*⁵. It is these *data providers* which at pre-defined time intervals gather performance data and transfer this to the WLA Archiver. The agents do perform other tasks and provide data directly back to the Stores and Repository, but discussion on this is outside the scope of this document.

The following table lists the agents that are capable of collecting performance data for their respective managed object and whether they are certified for use by CATE Storage engineering.

⁵ The data provider is usually the managed object agent i.e. Symm agent, FCC agent, Host agent etc.

Agent Name	WLA Daily	WLA Revolving	WLA Analyst	Agent Citigroup Certified	Performance Collections Certified
Database Agent for Oracle	Y	Y	n/a	Y	N
Fibre Channel Connectivity Agent	Y	Y	Y	Y	Y ⁶
Host Agent for AIX	Y	Y	n/a	Y	Y
Host Agent for HP-UX	Y	Y	n/a	Y	Y
Host Agent for Linux	Y	Y	n/a	Y	Y
Host Agent for Solaris	Y	Y	n/a	Y	Y
Host Agent for Windows	Y	Y	n/a	Y	Y
Storage Agent for Clariion	Y	Y	n/a	Y	Y
Storage Agent for HDS	Y	Y	n/a	Y	N
Storage Agent for NAS	Y	Y	n/a	N	N
Storage Agent for Symmetrix	Y	Y	n/a	Y	Y

Table 1: Agent WLA Support

5 Sizing of the Performance Manager Environment

The ECC infrastructure agents such as WLA Archiver need to be sized appropriately to reflect the size of the environment being managed. The word “sized” in the context of ECC means to have adequate agents and resources to properly manage the elements to which those agents relate. EMC have detailed these sizings for all agents in what is known as the Performance and Scalability Guidelines⁷ or PnS Guidelines as to which it is more commonly referred.

The WLA Archiver is fed data from a number of data providers within the environment, a full list of which is included in the [Table 1](#). The number of WLA Archiver agents deployed in the environment is dependent on the number and type of elements being managed and the size of each element being managed. The following tables specify the number of WLA Archiver required based in the PnS guidelines.

⁶ The WLA collection for switches has only limited support.

⁷ This solution guide has referenced the EMC ControlCenter 6.0 Performance and Scalability Guidelines. The document is available from the EMC Powerlink website.

Managed Object Type (medium-sized objects)	Dedicated	Shared
Storage Arrays	75	50
Hosts	300	200
Databases	75	50
Switches	50	30
Total of all listed managed objects	500	330

Table 3: Managed Objects per WLA Archiver

The table can be interpreted as

The table above makes the distinction between a server that is dedicated to a WLA Archiver and one that is shared with other infrastructure components i.e. a Store. EMC have based this sizing on the following server configuration:

- Shared: 2x2.4GHz CPUs, 2GB of RAM and 36GB of disk space.
- Dedicated: 2.2.4GHz CPUs, 2GB of RAM and 146GB of disk space.

Under the managed object type EMC have used the term “medium-sized objects”, the following table specifies the sizing difference between small, medium and large objects.

Managed Object	Resource	Small	Medium	Large
Symmetrix	Logical Volumes	1-800	801-2000	2001-8192
	Front-End Mappings	1-1600	1601-4000	4001-16000
Clariion	Disks	1-60	61-80	81-120
	LUNs	1-256	257-512	513-1024
Switch	Ports	1-16	17-64	65-256
Oracle or DB2 database	Sum of data files and tablespaces	1-100	101-200	201-400
SQL Server or Sybase	Sum of databases, data files, log files and file groups	1-100	101-200	201-400

Agent host	Host devices	1-16	17-128	129-512
	Logical volumes	1-16	17-64	65-128
	File Systems	1-16	17-64	65-128
	Redundant paths	1-2	3-4	5-32
ESX Server	Virtual machines	8	16	32
	Storage array devices	16	32	64
	File System (vmfs3)	8	16	32
	SAN Mapped Lun	8	16	32
	Redundant paths	2	2	2

Table 4: Managed Object Sizes

So to properly determine the number of WLA Archivers required for the environment the number of managed objects within the environment will need to be understood.

5.1 Sizing Case Study

5.1.1 London ECC Environment

As part of this certification a case study was performed on the London RDC ECC environment. In the London RDC environment the managed objects that were selected to have WLA switched on were:

- Storage Arrays
- Core Switches
- Hosts

It is recommended that a combination of ECC, StorageScope and Symcli be used to determine each number of managed objects in the environment and also the size of each managed object in the ECC environment. Details of how the tools mentioned can be used in managed object sizing can be found in the appendices.

The following sections provide a real world example of how to complete the sizing of an ECC environment and how to determine the number of WLA Archivers required for the environment.

5.1.1.1 Storage Array Sizing

As per Table 5 there are 2 metrics that need to be gathered to determine the correct sizing category of Symmetrix storage arrays in the environment, these are:

1. The number of Logical Volumes per array

2. The number of Front-End mappings per array

To determine the number of logical volumes per Symmetrix storage array in your environment, the easiest way to obtain this information is to run the following symcli command:

```
#symcfg list | grep local
```

000287750096	Local	DMX2000S	5670	32768	14	1711	Medium
000287750107	Local	DMX2000S	5670	36864	14	1253	Medium
000287750200	Local	DMX2000S	5670	32768	14	1407	Medium
000287750442	Local	DMX2000S	5670	32768	14	1117	Medium
000287751808	Local	2000S-M2	5670	65536	14	2567	Large
000287751819	Local	2000S-M2	5670	65536	14	2521	Large
000287890275	Local	3000-M2	5670	65536	14	2137	Large
000287890293	Local	3000-M2	5670	65536	14	3237	Large
000287890300	Local	3000-M2	5670	65536	14	3259	Large
000287890581	Local	3000-M2	5670	65536	14	3461	Large
000287890596	Local	3000-M2	5670	65536	14	3331	Large
000287890602	Local	3000-M2	5670	65536	14	3839	Large
000290100953	Local	DMX3-24	5771	65536	14	3823	Large
000290101065	Local	DMX3-24	5771	65536	14	2134	Large
000290101136	Local	DMX3-24	5771	81920	14	1947	Medium
000290101150	Local	DMX3-24	5771	65536	14	3823	Large
000290102819	Local	DMX3-24	5771	98304	14	3410	Large
000290102839	Local	DMX3-24	5771	98304	14	3029	Large

Table 5: Symmetrix Devices

The highlighted column specifies the number of logical volumes per Symmetrix storage array. This figure should be referenced with Table 5 for correct sizing of each array.

The second metric that needs to be captured is the number of front-end addresses per storage array. This can easily be obtained from StorageScope by generating a report that shows the “# LUNs on Port” per array and simply generating a pivot table that specifies the total LUNs per port for the entire array. See below:

Symmetrix ID	# of Front End Mappings	Category
287750096	1718	Medium
287750107	838	Small
287750200	918	Small
287750442	790	Small
287751808	1948	Medium
287751819	1878	Medium
287890275	1302	Small
287890293	1741	Medium

287890300	1886	Medium
287890581	1978	Medium
287890596	2072	Medium
287890602	2616	Medium
290100953	1946	Medium
290101065	3306	Medium
290101136	3466	Medium
290101150	1958	Medium
290102819	6218	Large
290102839	2802	Medium

Table 6: Symmetrix Front End Mappings

At the end of this process you will have two figures the number of devices per storage array and the number of front-end mappings per storage array. Sizing your environment should always be based on the “largest” category that the two figures fall into.

For example: If the number of devices per storage array falls into the Large array category, but the front-end mappings fall into the Medium array category as specified in table 5, the sizing should be based on the Large array category.

5.1.1.2 McData Switches

As indicated in Table 5 only one metric is needed to determine the switch sizing category and this is the number of ports per switch.

Name	Vendor	# Ports	# Used Ports	Firmware Revision	Model	Category
SW94R_SRDF_BLACK_RDC4	McData	40	26	09.00.00	ED-64M	Medium
SW93R_SRDF_GREEN_RDC4	McData	40	26	09.00.00	ED-64M	Medium
SW92R_SRDF_BLACK_RDC1	McData	32	12	09.00.00	ED-64M	Medium
SW91R_SRDF_GREEN_RDC1	McData	32	12	09.00.00	ED-64M	Medium
SW26R_CORE_BLUE_RDC4	McData	140	125	09.00.00	ED-140M	Large
SW25R_CORE_RED_RDC4	McData	140	125	09.00.00	ED-140M	Large
SW16R_CORE_BLUE_RDC4	McData	140	137	09.00.00	ED-140M	Large
SW15R_CORE_RED_RDC4	McData	140	137	09.00.00	ED-140M	Large
SW12R_CORE_BLUE_RDC1	McData	128	62	09.00.00	ED-140M	Large
SW11R_CORE_RED_RDC1	McData	128	62	09.00.00	ED-140M	Large
SW02R_CORE_BLUE_RDC1	McData	96	51	09.00.00	ED-140M	Large
SW01R_CORE_RED_RDC1	McData	96	52	09.00.00	ED-140M	Large
Total		1152	827			

Table 7: ECC Switch Port Data

The Table 6 data has been obtained from the StorageScope but is also readily available in the ECC Console⁸. The categorization of the switches has been based on the “#Ports” metric rather than the “# Used Ports” metric.

There are tight restrictions on the number of ports that the FCC agent can gather WLA daily data from in a specific time period. These restrictions are detailed in the table below.

Switch Vendor	McData						
Data Collection Policy Name	Number of Agent	1 to 512 Ports	513 to 1024 Ports	1025 to 2048 Ports	2049 to 3072 Ports	3073 to 4096 Ports	4097 to 8192 Ports
Fabric Validation	1	15 minutes	20 minutes	30 minutes	1 hour	2 hour	3 hours
	2	15 minutes	15 minutes	30 minutes	45 minutes	1 hour	2 hours
Device Validation	1	15 minutes	30 minutes	1 hour	2 hours	3 hours	4 hours
	2	15 minutes	20 minutes	45 minutes	1 hour	2 hours	3 hours
Discovery Scan	1	10 minutes	15 minutes	20 minutes	30 minutes	45 minutes	1 hours
	2	10 minutes	10 minutes	15 minutes	20 minutes	30 minutes	45 minutes
WLA Daily*	1	30 minutes	30 minutes	1 hour	Disable*	Disable*	Disable*
	2	30 minutes	30 minutes	1 hour	Disable*	Disable*	Disable*
WLA Revolving	1 or 2	Disable					
Performance Statistics	1 or 2	Monitor up to 10 switches at a time.					

Table 8: Recommended FCC Agent DCP Frequencies for McData Switches

- * The WLA daily DCP setting must be a factorial of 60 minutes i.e. 10min, 15min, 30min or 1 hour. This is because WLA processing occurs on hourly intervals and requires at least 1 data point.

It is important to note that the above example is based purely on McData switches. A similar table referencing the Cisco switches can be found in the EMC ControlCenter 6.0 Performance and Scalability Guidelines.

Due to the restrictions around the number of ports that data can be collected in for Large environments i.e. 1000+ ports it is recommended that only performance statistics for the core switches or switches that are of particular interest are collected.

5.1.1.3 Host Sizing

There are four different metrics as per Table 5 on which the host sizing is based. These are:

- Number of Host Devices
- Number of Logical Volumes

⁸ For details on how to do this please refer to the Appendices

- Number of Filesystems
- Number of HBA connections

Sizing data for each of these metrics can be combined into one table by using a StorageScope query and outputting the returned data to a csv file. As there are 4 different metrics that need to be taken into account to correctly categorize each host into *small*, *medium*, *large* and *extra large* hosts some data manipulation is still needed before this can be done. The details of how to generate this query and perform the data manipulation can be found in the Appendices.

Once the metrics have been gathered and the data manipulation performed a simple pivot table⁹

Row Labels	Count of Host Name
Extra Large	11
Large	36
Medium	259
Small	265
Grand Total	571

Table 9: Host Size Categorization

It is important to remember that the ECC PnS guidelines base the sizing on medium sized objects. As a rough guideline:

- * 1 Large Host = 20 Medium Hosts
- * 1 Extra Large Host = 40 Medium Hosts
- * 20 Small Hosts = 1 Medium Host

This is an approximate approach to the sizing of the number of WLA Archivers that will be required. The total of 571 hosts above increases to 3292 medium sized hosts excluding the 146 small hosts, when adjusted to the conversion metric above.

5.1.1.4 WLA Archiver Sizing

Based on the data collections performed in the previous 3 sections and the categorization of the managed objects into their various sizes the following table can be produced.

	Small	Medium	Large	Extra Large	Calculated Medium MO's
DMX	4	13	1	0	17
Hosts	265	259	36	11	1432

⁹ A pivot table is a means of presenting data within excel. It can be sorted and tailored to meet the needs of the creator.

Switches	0	4	8	0	20
----------	---	---	---	---	----

Table 10: Summary of Managed Object Sizes

Based on the managed object sizing shown above and applying the information as outlined in the ECC PnS guidelines (detailed in Table 4), the number of WLA Archivers required for this environment would be, **7**. It is important to note that EMC do not have a recommendation on the maximum number of WLA Archivers in a single ECC environment.

In most circumstances only the key managed objects in an environment storage arrays and switches will be of interest to the SAN administrator. Occasionally it may be necessary to collect performance data for a small number of servers but this should only be performed on an ad-hoc temporary basis for the duration of the performance issue. There are other tools in the Citi environment i.e. BMC and Quest that collect performance data for the host and database respectively.

6 Best Practice Recommendations for Deploying Performance Manager

The following section outlines the best practice recommendations for the deployment and configuration of Performance Manager, its agents and DCP's.

1. Always perform accurate sizing of your managed objects to obtain the correct amount of WLA Archivers for your environment
2. Deploy the WLA Archivers into the environment before enabling any WLA DCPs. Failure to do so could result in an unbalanced distribution of data providers to Archiver.
3. As a rule do not collect host statistics unless when there is a particular interest in a small number of servers. In this situation only enable the WLA Daily collection for the period of the performance issue.
4. In large SAN (1000+ ports) environments limit the WLA daily collection of switch statistics to the core.
5. Revolving DCP's should only be used in situations whereby there is a performance issue and higher frequency of data points in the .btp is required.
6. Due to the additional overhead the Revolving DCP has on the data provider server limit the collection to 2 managed objects for 3-6 hours per data provider or only as long as is necessary to capture the relevant information.
7. Analyst DCP's should only be used in situations where there is a performance issue or window of interest and a higher frequency of data points is required for the duration of that window.
8. Due to the additional overhead the Analyst DCP has on the data provider server limit the collection to 2 managed objects for 3-6 hours per data provider or only as long as is necessary to capture the relevant information.

9. Data Collection policy schedule for switch WLA data should be based on the tables provided in the PnS guidelines.
10. Data Collection policy schedule for DMX storage array WLA data is not defined in the PnS guidelines but is by default set to 15mins. This time interval is recommended for all DMX Storage arrays.
11. Data Collection policy schedule for Clariion storage array WLA data is not defined in the PnS guidelines but is by default set to 15mins. This time interval is recommended for all Clariion Storage arrays.
12. Data Collection policy schedule for HDS storage array WLA data is not defined in the PnS guidelines but is by default set to 15mins. This time interval is recommended for all HDS Storage arrays.

7 Installation of Performance Manager

7.1 Installation Pre-Requisites

Review and verify each of the following points prior to launching the InstallShield to help avoid problems during installation:

- Ensure the Performance Manager license key is entered through the Console.
- The target host must be running a supported version of the Windows operating system, and have the required amount of memory and processor speed (refer to *EMC ControlCenter Support Matrix* on Powerlink for details).
- The target host must have access to the ControlCenter Repository host through TCP/IP.
- The target host must use DNS to resolve names and IP addresses.
- Log into the target host as the Local User Administrator (*or other supported language equivalent*) and not as a Domain Account with Local Administrator Group rights.
- Close all other applications.
- Stop any previously installed ControlCenter Services, except the ECC Server and Repository.
- Stop and set to Manual the Distributed Transaction Coordinator Service.
- Stop and set to Manual any Anti-Virus software.
- Stop and Disable Windows Terminal Services (*requires a re-boot of the host*).
- Close the Services Window before starting the ControlCenter InstallShield.
- Have the ControlCenter installation disks at hand or map to a share where the installation disks have been loaded.

7.2 Install Performance Manager

1. Insert Disk #1 or map to the share where the installation disks have been loaded.
2. If installing from disk, InstallShield starts automatically, if installing from a share, you may need to run **setup.exe** to start the InstallShield.
3. If you are presented with an *Open File - Security* screen, click **Run**.
4. *Preparing Setup* screen appears briefly.
5. *InstallShield Warning* dialog appears, warning to close all other applications, click **OK**.
6. *InstallShield Warning* dialog appears, warning to turn anti-virus software off, click **OK**.
7. *Welcome* screen appears, click **Next**.
8. License Agreement screen appears, (if no other ControlCenter components are installed on this host, otherwise this screen is skipped). Review License Agreement, select **I accept the terms of the license agreement**, click **Next**.
9. *Evaluating Network Configuration* message appears, wait for it to clear.
10. *Choose Destination Location* screen appears, click **Change** to enter a different install location, otherwise click **Next**.
11. *ControlCenter Repository Options* screen appears, select **Specify existing ControlCenter Repository**, click **Next**.
12. *ECCRepository Data* screen appears, enter the **Hostname** and **Port Number** of the ControlCenter Repository Host, click **Next**.
13. *Contacting the Repository on Hostname* message appears, wait for it to clear.
14. *Select Features* screen appears, select **Performance Manager**, click **Next**.
15. *Ready to Install the Program* screen appears, click **Install**.
16. InstallShield proceeds to install Performance Manager.
17. *Install Complete* screen appears, click **Finish**.

7.3 WLA Archiver System Requirements

The minimum system requirements as outlined by the EMC PnS guidelines are:

- **Shared:** 2x2.4GHz CPUs, 2GB of RAM, and 36GB of disk space. WLA Archiver is deployed on a stand-alone Store host or with other agents.
- **Dedicated:** 2x2.4GHz CPUs, 2GB of RAM, and 146GB of disk space. No other infrastructure components or agents are installed on the host.

For VMWARE the system requirements are:

7.4 Installing the WLA Archiver

Before installation of the WLA Archiver agent the following tasks will need to have been completed:

- Determine the number of WLA Archiver agents required based on your requirements for performance reporting.
- Identify a server on which the agent can be installed.
- Performance License Manager has been installed.
- All WLA DCPS are disabled.
- Master and Host agents are installed and available on the target hosts.

Once you are ready please follow the following steps to install the WLA Archiver agent.

1. From the tree panel on the left hand side of the console screen select the server you wish to install the WLA Archiver and right click, select Agents and Install as shown in the Figure 1 below:




Figure 1: Installing an agent

2. Click Next in the dialogue box which appears and proceed to the Agent Installation dialogue box. As shown below in Figure 2.




Figure 2: Agent Installation Screen

3. Select the Workload Analyser Archiver agent and Click Next.
4. Click Next after the pre-requisite checks dialogue box returns a green tick mark.
5. On the Performance Manager Archiver Installation dialogue box enter the following information
 - Archiver Path: D:\emc\wlaarchiver
 - Report Path: D:\emc\wlaareports
6. Leave all other fields as default, and click OK to continue.
7. Enter a Task name on the next dialogue box and Click OK to install the WLA Archiver.
8. Wait for the installation to return a success before continuing to set-up the DCP's.

7.5 SYMACL

In order to gather Performance statistics from a DMX system the SYMACLs must be updated to allow connection of the Symm agent to the DMX array. For ECC 6.0 the SYMACL definitions have changed since the ECC 5.2 SP4 version. The changes are detailed below:

Changes:

- ECC_APP group is removed
- ECC_HOST ALL_DEVS VLOGIX access is removed
- **ECC_HOST ALL_DEVS ALL access is added**

Existing permissions are set up as below:

- | | | |
|------------|----------|--------|
| • ECC_HOST | ALL_DEVS | BASE |
| • ECC_HOST | ALL_DEVS | ECC |
| • ECC_HOST | ALL_DEVS | VLOGIX |
| • ECC_APP | ALL_DEVS | ALL |

New permissions are as follows:

- | | | |
|------------|----------|------|
| • ECC_HOST | ALL_DEVS | BASE |
| • ECC_HOST | ALL_DEVS | ECC |
| • ECC_HOST | ALL_DEVS | ALL |

These changes apply to SE 6.4.2.3 and above for ECC 6.0 Symm agent and will be documented in the Solutions Enabler Solutions guide.

The following section details the correct method for setting up of data collection policies for WLA in the ECC environment.

8 Set-Up of WLA Data Collection Policies

There are 3 WLA data collection policies (DCP) that can be enabled, although not all MO's have these data collection policies available. The 3 collection policies are:

- WLA Daily DCP
- WLA Revolving DCP
- WLA Analyst DCP

The WLA Daily DCP default polling cycles have been documented extensively in the EMC ControlCenter 6.0 Performance and Scalability Guidelines. The default polling cycles for each MO can be found in the Agent Guidelines section of the PnS Guidelines.

8.1 Pre-Requisites for the set-up of the WLA Data Collection Policies

Before starting on the set-up of the Data Collection Policies (DCPs) for WLA, you need to determine the following:

1. The managed objects you wish to monitor in your environment i.e. storage arrays, hosts, switches etc.
2. The number of each managed object you wish to monitor.
3. The size of each managed object you wish to monitor.

4. WLA Archivers are installed and are active.
5. Performance Manager is installed
6. StorageScope, Symcli and ECC are available to act as datasources for the sizing exercise,

8.2 WLA Daily DCP

The daily DCP collections are disabled by default. The daily WLA DCP collections are the primary collection policies for the managed objects in the environment. WLA Daily data collections are used to create the performance archives for historical performance analysis. Settings defined in the WLA Daily data collection policy determine the resolution of the data provider's performance archives. The default polling cycle for MO's in the environment is outlined in the Agent Guidelines section of the PnS guide Table 16.

Based on the sizing exercise that has been completed the following section details how to properly set up the WLA Daily DCP for the environment.

8.2.1 WLA Daily DCP Set-Up

You will need to logon to the ECC Console with sufficient privileges to edit Data Collection Policies.

1. From the right tree panel select Administration and expand the tree.
2. Select Data Collection Policies and expand the tree.
3. Select Policy Definitions and expand the tree.
4. Expand each one of Policy Definitions and ensure that the WLA collection policies are disabled.






Figure 3: WLA Daily Tree

5. Expand Storage Agent for Symmetrix and right-click WLA Daily_00 → Edit Data Collection Policy as shown above in Figure 2.

Figure 4: Policy Definition



6. Create a new Schedule as shown in Figure 3 above.
7. Dependent on the number of Storage Arrays that you wish to collect for se the interval appropriately see Figure 4 below.



The image shows a Windows-style dialog box titled "Schedule Definition: (New Schedule)". It has a "Properties" tab selected. Inside the dialog, there is a "Name" field with the text "WLA_Daily_Prod". Below it is an "Interval" field set to "15" and a unit dropdown menu set to "Min". Under the "Days" section, there are checkboxes for "Mon", "Tues", "Wed", "Thurs", "Fri", "Sat", and "Sun", all of which are checked. There are also "SetAll" and "ClearAll" buttons. Below the days section, there are time selection fields for "From" and "To". The "From" time is set to 00:00:00 and the "To" time is set to 23:59:59. At the bottom of the dialog are "OK", "Cancel", and "Help" buttons.

Figure 5: Schedule Definition

8. Once you have set the schedule definitions select OK.
9. Click on the "Apply To" tab and un-check the "Apply this policy to all applicable Symmetrix Objects"¹⁰, as per Figure 5.

¹⁰ You want to make sure that you have complete control over what MO's the policy is enabled for. By un-checking this box each additional MO which is added to the environment needs to be manually added to the DCP.




Figure 6: Policy Definition

10. Add the DMX arrays that you wish to monitor by using the Add> radio button.
11. Check the box titled “Policy Enabled?”.
12. Press OK to save the changes to the Policy.

The steps above demonstrate how the WLA Daily policy should be set up for the storage Agent for Symmetrix. The WLA Daily policy will need to be completed for the other MO's that wish to be monitored.

8.3 WLA Revolving DCP

As per guideline **D3** in the PnS guidelines the WLA Revolving DCP should be **disabled** to conserve resource utilization on the agent host. As per the guideline:

“Performance statistics data being collected by the WLA Daily, Revolving and Analyst DCP's are identical but they operate at different frequencies. Typically, WLA Revolving DCP would collect statistics data for the last 2 hours and save it on the local disk of agent host. Data collected by WLA Revolving DCP is used during the day whereas WLA daily is for short and long term trending etc.”

The revolving DCP when configured can collect granular performance data for a defined time interval. Once the time interval has been reached e.g. the default is 120 minutes, the oldest interval is removed from the WLA Revolving and the newest data is appended to the collection. WLA Revolving data is available for immediate use in Performance Manager providing that some data points have been collected.

Purpose The WLA Revolving data collections are used to:

- Create collections with resolutions that are higher than the WLA Daily collections and can be viewed on demand, such as shortly after a performance issue is discovered.
- Troubleshoot short-term performance-related issues.
- Identify a performance-related problem as soon as it is detected.
- Monitor the most recent status of the workload.

Revolving Performance data is stored on the agent host i.e. the host that is being used as the data provider. So, in the case of revolving data being collected for a DMX array the revolving data will be stored on the Symm/SDM agent host, as this host is acting as the data provider to WLA Archiver. In order to view the revolving performance data the revolving data needs to be actively pulled back from the data provider. This can be done via the console.

8.3.1 WLA Revolving DCP Set-Up

Warning – WLA Revolving Data Collection is resource intensive and can have a negative impact on the Data Provider performance. The running of WLA Revolving collection concurrently with either the WLA Daily or WLA Analyst Collections can lead to missing or incomplete data in some or all .btp files. You will also require a minimum of 100MB of free space on the data provider for the WLA Revolving Collection.

You will need to logon onto the ECC infrastructure with sufficient privileges to edit the Data Collection Policies.

1. From the right tree panel select Administration and expand the tree.
2. Select Data Collection Policies and expand the tree.
3. Select Policy Definitions and expand the tree.




Figure 7: WLA Revolving Tree

4. Expand Storage Agent for Symmetrix and right-click WLA Daily_00 → Edit Data Collection Policy as shown above in Figure 5.




Figure 8: Configuring Revolving DCP

5. Enter the name and interval fields dependent on the activity you are hoping to achieve.
6. Click OK once you have finished.
7. Click on the “Apply To” tab and select the DMX array which you are interested in.




Figure 9: WLA Revolving Policy Definition

8. Click the “Policy Enabled” box and then click OK to save your changes.
9. From the navigator panel click ECC_Administration and select “Policies”





Figure 10: ECC Administration > Policies

10. Expand “Storage Systems” in the tree view and select the array you have set the WLA Revolving policy for.
11. Add the array to the policies view as shown in Figure 9 below.



000290102819	Symmetrix SDM Agent	Discovery_00	Host	rdcsan-mgt2	Enabled	Hour_12
000290102838	Storage Agent for Symmetrix	Historical Data_00	Host	rdcsan-mgt1	Enabled	Day_01_0000
000290102839	Storage Agent for Symmetrix	Historical Data_00	Host	rdcsan-mgt2	Enabled	Day_01_0000
000290103452	Storage Agent for Symmetrix	Local Discovery_00	Host	rdcsan-mgt1	Enabled	Day_01_0000
000290103471	Storage Agent for Symmetrix	Local Discovery_00	Host	rdcsan-mgt2	Enabled	Day_01_0000
000290103544	Symmetrix SDM Agent	Masking Configuration_00	Symmetrix	000290103452	Enabled	Hour_08
000290103550	Storage Agent for Symmetrix	Performance Statistics_00	Symmetrix	000290103452	Enabled	Minute_02
Centera_CH_01E	Storage Agent for Symmetrix	Real-time BCVRDF Status_00	Symmetrix	000290103452	Enabled	Minute_01
Centera_CH_01R	Host Agent for Solaris	vLA Daily_00	Host	rdcsan-mgt1	Disabled	Minute_15
Centera_DE_01E	Host Agent for Solaris	vLA Daily_00	Host	rdcsan-mgt2	Disabled	Minute_15
Centera_DE_02E	Storage Agent for Symmetrix	vLA Revolving_00	Symmetrix	000290103452	Enabled	Revolving_Test
Centera_PL_01G						
Centera_PL_01S						

Figure 11: Data Provider Policy View

12. Right click on the WLA Revolving_00 for the relevant data provider in the right hand pane and select WLA → Get Data. This will pull back the revolving performance data from the data provider. Dependent on the duration of the collection and the number of data points it may take several minutes before this data is available for view in Performance Manager.

8.4 WLA Analyst DCP

As per guideline **D3** the Analyst DCP is disabled by default. Its purpose is to collect granular performance statistics over a set time period e.g. a backup is running between 01:00 → 04:00 and the effect of the backup on DMX performance needs to be determined. By setting the Analyst DCP to run for 1 or 2 min intervals over the duration of the backup then detailed performance metrics can be obtained.

The analyst DCP will run for the duration specified and will not run again unless it is enabled to do so. Both the Analyst and the Revolving DCPs are resource intensive and therefore should only be run when necessary and should not be left enabled indefinitely.

8.4.1 WLA Analyst DCP Set-Up

Warning – WLA Analyst Data Collection is resource intensive and can have a negative impact on the Data Provider performance. The running of WLA Analyst collection concurrently with either the WLA Daily or WLA Revolving Collections can lead to missing or incomplete data in some or all .btp files.

You will need to logon onto the ECC infrastructure with sufficient privileges to edit the Data Collection Policies.

1. From the right tree panel select Administration and expand the tree.
2. Select Data Collection Policies and expand the tree.
3. Select Policy Templates and expand the tree.




Figure 12: WLA_Analyst Template

4. Expand Storage Agent for Symmetrix and right-click WLA Analyst → Edit Data Collection Policy as shown above in Figure 10.
5. In the unique descriptor field enter a meaningful name i.e. Test and click OK when done.
6. Next Expand Policy Definitions → Storage Agent for Symmetrix.
7. Right Click on WLA Analyst_TEST and Click Disable. It is important to do this straight away else it will try and apply the policy to all available objects and you definitely don't want to do that.
8. Verify that the policy has been disabled → the little rosette on the policy turns from yellow to white.
9. Right click on WLA Analyst_Test → Edit Data Collection Policy.
10. Click on the Actions Tab and click New to enter a new schedule for collection.
11. In the dialogue box enter a meaningful name and then configure your schedule based on your requirements and Click Ok when done.
12. Click the Apply to tab and Remove the arrays that you are not interested in and add the arrays that you are interested in.
- 13. Make sure to un-check the “Apply this policy to all applicable Symmetrix objects” box.**
14. Check the Policy Enabled box and Click OK to continue.
15. The WLA Analyst Policy will be available for re-view in Performance Manager at the end of the Data Collection period as outlined in the schedule you defined.




Figure 13: Policy Definition

16. Click the “Policy Enabled” box and then click OK to save your changes.
17. From the navigator panel click ECC_Administration and select “Policies”




Figure 14: ECC Administration

18. Expand “Storage Systems” in the tree view and select the array you have set the WLA Revolving policy for.
19. Add the array to the policies view as shown in Figure x below.

Symmetrix SDM Agent	Discovery_00	Host	rdcsan-mgt2	Enabled	Hour_12
Storage Agent for Symmetrix	Historical Data_00	Host	rdcsan-mgt1	Enabled	Day_01_0000
Storage Agent for Symmetrix	Historical Data_00	Host	rdcsan-mgt2	Enabled	Day_01_0000
Storage Agent for Symmetrix	Local Discovery_00	Host	rdcsan-mgt1	Enabled	Day_01_0000
Storage Agent for Symmetrix	Local Discovery_00	Host	rdcsan-mgt2	Enabled	Day_01_0000
Symmetrix SDM Agent	Masking Configuration_00	Symmetrix	000290103452	Enabled	Hour_08
Storage Agent for Symmetrix	Performance Statistics_00	Symmetrix	000290103452	Enabled	Minute_02
Storage Agent for Symmetrix	Real-time BCVRDF Status_00	Symmetrix	000290103452	Enabled	Minute_01
Host Agent for Solaris	vLA Daily_00	Host	rdcsan-mgt1	Disabled	Minute_15
Host Agent for Solaris	vLA Daily_00	Host	rdcsan-mgt2	Disabled	Minute_15
Storage Agent for Symmetrix	vLA Revolving_00	Symmetrix	000290103452	Enabled	Revolving_Test

Figure 15: Policies View

9 Performance Manager

9.1 Launching Performance Manager

Performance Manager can be launched from 2 different places within the infrastructure, these are:

- From the ECC Console
- From the desktop or server on which Performance Manager is installed.

To launch Performance Manager from the ECC console do the following:

1. Select File → Launch → Performance Manager




Figure 16: Launching Performance Manager from ECC Console

2. On the dialogue box which appears enter your SOEid and password.

The above steps will enable you to launch Performance Manager from the ECC Console. To launch it locally from the server or desktop please do the following.

1. From the server Click Start → Programs → EMC → EMC ControlCenter → Performance Manager, as shown below in Figure 17.





Figure 17: Launching Performance Manager

2. On the dialogue box which appears enter your SOEid and password.

9.2 Loading a BTP File

Once successfully logged onto Performance Manager, the next step in the process is to load a .btp file for the MO that is of interest. To do this follow these steps:

1. In the Data Selection dialogue box select the Class and the identifier that you are interested in:

**Figure 18: Data Selection**

2. Select the date range that you are interested in and click OK to load the BTP file.
3. You should not be presented with the following screen. See below:




Figure 19: Performance Manager Loaded BTP file

9.3 Performance Manager View

Performance View (PV) is the manual ad hoc analysis tool used to view the performance and configuration data stored in the WLA archives. It can create a multitude of views (graphs, tables) of the data for specified managed objects in order to more easily identify performance problems. Every graph and table derives its statistics from the collections. The Performance Manager GUI consists of three areas:

- The Views Pane
- The Analyser Toolbar
- The Data View Pane

The following sections cover these in detail.

9.3.1 The Views pane

It contains three tabs at the bottom of the screen:

- Views – A top-level metrics overview of the array. Views are programmed graphs by EMC with definitions and characteristics that are stored in a file and are loaded when the data file opens.
- Metrics – Detailed statistics about all components of the array. Most of the ad-hoc analysis of data is done using Metrics tab. Functions performed here are creating the graph, modifying the graph, saving the graph as a view. Navigating to/from the configuration and links views to see relationships among components, and generally conducting detailed investigations.
- Config – Configuration details about the array, disks, LUNs, and associated SPs. The configuration of the storage array describes the relationship among logical and physical components in the system. e.g. Within the Symmetrix, devices are the logical components and directors' busses and disks are the physical components. The relationships among these components define the I/O paths taken to serve an I/O command.

9.3.2 The Analysis toolbar

It contains tools that perform data manipulation. The Data pane, on the right, contains all requested statistics, correlations and configurations.

9.3.3 The Data view pane

It is the graphical area, on the right, that contains all requested statistics, correlations and configurations.

<i>Tool icon</i>	<i>Tool name</i>	<i>Tool details</i>
	Vital Signs	Creates graphs for all Vital Signs listed under the selected item in the Views. When new views are created, they may be assigned a characteristic known as “vital sign” and shown in the Views tab with the red heart icon. This allows s to be accessed quickly.
	Graph per Object	This is used to create a single Object graph for multiple metrics for comparison purposes.
	Graph per Metric	This is used to create a single Metric graph for multiple objects (such as devices).
	Histogram	Presents the selected data in a historical bar graph.
	Add to Current Graph	Allows you to add a Nother object to an open Per Metric graph or a Nother metric to an open per Object graph.
	Insert This Graph Into Views	Saves the open and selected graph to the Views pane.
	Auto Correlate	Uses the items currently selected in the metric panes to create a Correlation table. Correlation tables present statistical data on the correlation between selected items. A Correlation table demonstrates how one object within a system affects a Nother.
	Correlate	Opens the Correlate Selections dialog box. You can then create graphs that display the correlation between selected items.
	Host to Symmetrix Disk Configuration	Displays the fields in the Configuration table from Host Device to Symmetrix disk information. This button is only enabled when viewing host or Symmetrix storage array configuration.
	Symmetrix Disk to Host Configuration	Displays the fields in the Configuration table from Symmetrix disk to Host Device information. This button is only enabled when viewing host or Symmetrix storage array configuration.
	Set Selection in Configuration	Transfers objects selected in the Configuration table into the Metrics panes.

	Set Selection in Configuration	Adds the objects selected in the Metrics Panes to the Configuration table.
	Compare Configurations	Displays, in a table, the differences between selected configuration files. This button is only enabled when viewing host or Symmetrix storage array configuration.
	Links	Activates the Links view.
	Enable Session Mode	Turns on the session mode. After turning on the session mode all subsequent graph are saved into the active session.

9.4 Data Views

Data views are graphs and histograms of interval, daily, weekly, monthly, or revolving performance data. They are created from Performance Manager for examining the performance and configuration data. They are opened, created, and displayed directly in the Performance Manager interface. The data views are also used to define the content of the Web views provided in the Performance Manager Automated Reports.

9.4.1 Factory data views

They are provided with Performance Manager as default and are available in the Views pane. Vital Signs are included with the factory data views. Vital signs are considered key indicators of the behavior of the data provider. Vital Signs are identified by the heart icon.

9.4.2 Public data views

They are performance and configuration data views defined by the user that are saved to the Repository where they can be shared with other users. This allows remote analysis of data without providing full access to the local archives.

9.4.3 Private Data Views

They are performance and configuration data views defined by the user that are saved locally to the host directories.

9.5 Automated Reports

Performance Manager Automated Reports are defined through automation jobs. The automation jobs are managed from the Automation Job Scheduler. Two types of reports are available:

- Factory Reports
- Job Scheduler

Factory Reports are EMC-defined views that come with the Performance Manager. These views come with pre-set threshold values for you to see and immediately recognize that a threshold was exceeded

somewhere within this report category. These are viewed via the PM view as well as Web browser. Job Scheduler lets you define and control your own views in addition to the factory ones.

9.6 Working with Graphs

9.6.1 Ad-hoc time graphs and histograms

Ad-hoc time graphs and histograms are created from the Metrics panes. Ad-hoc time graph is a time-series graph showing the metric change over a period of time. Ad hoc graphs are created once, based on the items selected. After closing an ad-hoc graph, the formula used to create the graph is lost. However, it is possible to save the graph formula before closing it.

A histogram is a bar graph of historical data of selected items. Histogram gives a cumulative count of the metric over a time period. It does Not show variation.

9.6.2 Trend graph

A trend graph is used to compare system performance over different time periods. By default all time graphs are created as trend graphs.

9.6.3 Other graph representations

Ad-hoc (line) graph can be converted to other type of graphs such as Scatter, Area, Time-Bar and Ribbon. This is done via the graph wizard. Also, the graph background can be changed via the graph wizard. Ribbon graph provides a 3D visualization of data.

****Note** The appearance of a ribbon graph is affected by the size of the graph being viewed. When the total number of data points times the number of files times the number of objects in the graph exceeds 2000 the ribbon graph will start to have problems.

9.6.4 Creating Ad-hoc time graphs

1. Open the Metrics panes and select the data provider identifier and category from the data provider list.
2. Select one or more objects from the objects list.
3. Select one or more metrics from the metrics list.
4. Click  (Per Objects) to create one Time graph for each object (and multiple metrics) selected in the objects list. Each line in the graph represents metrics selected from the metric list.
5. Click  (Per Metrics) to create one Time graph for each metric (and multiple objects) selected in the objects list. Each line in the graph represents objects selected from the object list.
6. Click  (Histogram) to create a bar graph of historical data for all selected objects and metrics (cumulative).
7. Click  (Add to Graph) to add more metrics or objects to an existing graph in the analysis

9.7 Performance metrics

9.7.1 Vital signs

EMC has defined certain important Performance Metrics for easy reference. They are shown with a sign



in the views pane. By clicking this sign the following views will be automatically generated.

- System and Device Write Pending Events per sec
- All Host Directors % Utilisation
- Disk Utilisation
- Disk Directors Utilisation
- System Cache Utilisation
- SCSI & Fibre Ports Utilisation
- Sampled Device Write Response Times
- System Total IO Rate
- Fibre Directors UtilisationSampled Read Response Times

For many users this is the first point of call when beginning a performance analysis.

A detailed listing of performance metrics obtainable for the following MO's can be found in the Appendices:

- DMX
- Clariion
- Cisco Switches
- McData Switches
- HDS Storage Arrays

For all other metrics the help inherent within PM provides details on these.

Appendices

9.8 Appendix A - Gathering Sizing Data for your Environment

9.8.1 Storage Array Sizing

As specified in the Sizing section of this document you will need to spend some time determining what you want to manage and the size of the objects you wish to manage.

For a DMX this is relatively straightforward. The symcli command:

```
#symcfg list | grep local
```

An example of which has been shown in section [5.1.1.1](#).

The front-end mapping information can easily be obtained by running a StorageScope report. To run the report do the following:

1. Open ECC Console, Click on ECC Administration.
2. You will see the menu options at the very top of the console screen change. Click on Reports → Launch StorageScope.
3. Enter your login information in the dialogue box that appears.
4. Click Analysis → Queries.
5. On the Queries screen that appears select the “Array Configuration” report. This is a built-In report and is available in ECC 6.0.
6. Click Run to run the report.
7. Once the report completes this can be cut and pasted into an excel spread sheet.
8. In order to quickly determine what is a small, medium or large array you can set up a formula within excel to do this for you. This is especially useful if you have a large number of arrays.

N	O	P	Q	R	S
Arrays: # LUNs	Arrays: # Array Ports	Arrays: # Physical Hosts	Arrays: Last Discovered	Array Size	
64	8	0			
1,711	28	28	8/5/2008 4:05		
1,253	30	4	8/5/2008 4:05		
1,407	44	3	8/5/2008 4:06		

Figure 20: Excel New Field

9. In excel create a new field called "Array Size".
10. In the cell just below where you have typed the field name "Array Size" enter the following formula:

```
=IF(N2<=1600,"Small",IF(N2<=4000,"Medium",IF(N2<=16000,"Large",IF(N2>16000,"Extra Large","Problem"))))
```

11. Note that the "N2" value is a field reference, you may need to edit this depending on the position of the Arrays:# LUNs values.

N	O	P	Q	R
Arrays: # LUNs	Arrays: # Array Ports	Arrays: # Physical Hosts	Arrays: Last Discovered	Array Size
64	8	0		Small
1,711	28	28	8/5/2008 4:05	Medium
1,253	30	4	8/5/2008 4:05	Small
1,407	44	3	8/5/2008 4:06	Small
1,327	30	2	8/5/2008 4:04	
1,253	30	0	8/5/2008 4:06	

Figure 21: Dragging Formula

12. Once you have entered the formula you can drag it to other cells in the same column as shown in Fig 12, and this will automatically populate the cell with Small, Medium, Large etc.

9.8.2 McData Switch Port Sizing

The best way I found to gather switch port information is to obtain it via the ECC console.

1. Launch the ECC Console.
2. In the tree panel on the left hand side Click Connectivity and expand → Connectivity Devices
3. Right Click on the switches you are interested in and select properties.
4. Click on the properties pane on the right hand side of the Console
5. On the menu bar at the top of the screen click File → Export
6. Rename the file to switches.csv or other appropriate title.
7. Once exported it can be loaded into excel and you can use the following formula in the same way it was used in the previous section to determine your sizing

```
=IF (N2<=16, "Small", IF (N2<=64, "Medium", IF (N2<=256, "Large", IF (N2>256, "Extra Large", "Problem"))))
```

8. You will need to change the N2 value to reflect the position of the # Ports value

9.8.3 Host Sizing

Host sizing is little more involved than the switch or storage sizing. In order to accurately perform the host sizing you will need StorageScope and agents installed on each of the servers you wish to size.

1. From the ECC Console click ECC Administration on the ECC Toolbar.
2. From the menu select Reports → Launch StorageScope
3. In the dialogue box that appears enter your SOEID and Password.
4. Select Analysis → Queries, as shown below in Figure 16.




Figure 22: StorageScope Queries

5. On the following screen scroll to the end of the page and click Create.
6. On the Define Initial Settings screen, fill in the Name field and Description Field as shown below in Figure 17.

1. Define Initial Settings
Steps: 1 > 2 > 3 > 4 > 5 > 6

Set identification information for the query and choose a period of time for which the query will return data.

Name

Description

Report to assist in the sizing of hosts for WLA

Query Categories

Select a category for this query

Capacity Planning

Select Time Period for Data

☒ Current Data

☐ Historical Data

☒ Date Range: From
(mm/dd/yyyy) to

☒ Time Period

Last Week

Select Objects to include

Current objects only

Select Frequency:

Daily Values (up to 42 days)

Figure 23: Define Initial Settings

7. In the dialog box that appears select the areas of interest as shown in Figure 18 and click Next to continue.

Hosts

- ☒ Hosts
- ☐ Chargeback
- ☐ Host Devices
- ☐ Shared Devices
- ☒ HBAs
- ☐ Ports
- ☐ Host Arrays
- ☐ Volume Groups
- ☐ Logical Volumes
- ☐ File Systems
- ☐ Components
- ☐ Groups
- ☐ ESX Server

Connectivity

- ☐ Switches
- ☐ Ports
- ☐ Port Connections
- ☐ Fabrics
- ☐ Active Zone Sets
- ☐ Active Zone Set Members
- ☐ Unzoned End Ports
- ☐ Groups

Arrays

- ☐ Arrays
- ☐ LUNs
- ☐ Meta Devices
- ☐ Device Allocation
- ☐ Disks
- ☐ Ports
- ☐ Port Connections
- ☐ LUN Masking
- ☐ Replicas
- ☐ RAID Groups
- ☐ Storage Groups
- ☐ Storage Pools
- ☐ Groups

Applications

[Clear all](#)

Figure 24: Areas of interest

8. On the select columns screen add the columns of interest to the selected columns field as shown in Figure 19 below and click Next to continue.




Figure 25: Select Columns.

9. Click Next on the Create Filters screen.
10. Click Next on the Select sorting, grouping & total options screen.
11. On the Select Run Settings screen select Where = File and Format = CSV. Uncheck the two limit to boxes and click Save.
12. Next Click Run Now.
13. Once the report has completed, select the option to “Click here to view this file”. Then save the file to a place you can easily access.
14. Open the file in excel
15. Hold down Ctrl+F and click on the replace tab in the dialogue box that appears.
16. In the Find What box enter “=” Leave the Replace with box blank and click Replace All.
17. Do another replacement this time replace “ ” and leave the Replace with box blank.
18. Create a pivot table by doing the following.
19. Click in the cell that says Host Name and press Ctrl+Shift+End. This will select all the data for entry into the Pivot Table.
20. Click on Insert and Pivot Table as shown below in Figure 20.





Figure 26: Inserting a Pivot Table

21. Select OK when the Create Pivot Table dialogue box appears.
22. On the right hand side of the screen, select Host Name and drag the field into the Row Labels box, and add the #File Systems, #Logical Volumes, #Host Devices and Adapter No. to the Σ Values box.



23. Ensure that the values show as sums except for the Adapter No. as this needs to be shown as Count.
24. Once you have done this you will need to divide each of File systems, Logical Volumes and Hosts devices values by the corresponding Count of Adpater Number.
25. Once you have done that you can apply the if statements to the categorization of the hosts as has been seen in the previous sections.

9.9 Appendix B - Symmetrix Metrics

9.9.1 Symmetrix Device Metrics

The following table lists and defines metrics for Symmetrix devices.

% hit	$100 * (\text{total hits per sec} / \text{total ios per sec})$ Percentage of I/O operations, performed by the Symmetrix device, that were immediately satisfied by cache.
% miss	$100 - (\% \text{ hit})$ Percentage of read and write miss operations performed each second by the Symmetrix device that were misses. A miss occurs when the requested read data is not found in cache or the write operation had to wait while data was destaged from cache to the disks.
% random read hit	$100 * (\text{read hits per sec} / \text{total ios per sec})$
% random read miss	$100 * (\text{read misses per sec} / \text{total ios per sec})$
% read	$100 * (\text{total reads per sec} / \text{total ios per sec})$ Percentage of read I/O operations performed by the Symmetrix device.
% read hit	$100 * (\text{total read hits per sec} / \text{total reads per sec})$ Percentage of read hit I/O operations performed each second by the Symmetrix device. Read hits occur when the read operation is satisfied directly from cache.
% read miss	$100 * (\text{total read misses per sec} / \text{total reads per sec})$ Percentage of read miss I/O operations performed by the Symmetrix device. Read misses occur when the data requested by the read I/O operation was not in cache, causing a wait while it was fetched from disk.
% sequential io	$100 * (\text{seq ios per sec} / \text{total ios per sec})$ Use this metric for mainframe performance analysis only.
% Sequential read	$100 * (\text{seq reads per sec} / \text{total ios per sec})$ Use this metric for mainframe performance analysis only.
% sequential writes	$100 * (\text{seq writes per sec} / \text{total ios per sec})$ Use this metric for mainframe performance analysis only.

% write	$100 * (\text{writes per sec} / \text{total ios per sec})$ Percentage of total write I/O operations performed each second by the Symmetrix device.
% write hits	$100 * (\text{write hits per sec} / \text{writes per sec})$ Percentage of write hit I/O operations performed by the Symmetrix device. A write hit occurs when the I/O write operation is received, it is immediately staged in cache and will be written to disk at a later time.
% write miss	$100 * (\text{write misses per sec} / \text{writes per sec})$ Percentage of write I/O operations performed each second by the Symmetrix device that were write misses. A write miss occurs when the I/O write operations are delayed because the system or device write-pending limit was reached and the cache had to destage slots to the disks before the writes could be written to cache.
average io size in Kbytes	$(\text{HA Kbytes transferred per sec}) / (\text{total ios per sec})$ Average size of an I/O operation performed by the Symmetrix device.
average read size in Kbytes	$(\text{Kbytes read per sec}) / \text{total reads per sec}$ Average size of a read I/O operation performed by the Symmetrix device.
average write size in Kbytes	$(\text{Kbytes written per sec}) / (\text{writes per sec})$
DA Kbytes read per sec	Number of kilobytes read by the disk directors from the disk each second.
DA Kbytes transferred per sec	$(\text{DA Kbytes read per sec}) + (\text{DA Kbytes written per sec})$ Sum of the kilobytes read and written each second on all disk directors associated with the Symmetrix device.
DA Kbytes written per sec	Number of kilobytes written to the disk from the disk director each second.
DA prefetched tracks per sec	Total prefetched tracks each second from the disk directors to the cache.
DA prefetched tracks Used per sec	Number of prefetched tracks used each second from the disk directors to the cache.
DA read requests per sec	Number of read requests each second performed by the disk directors to the cache.
DA write requests per sec	Number of write requests each second performed by the disk directors to the cache.
HA Kbytes transferred per sec	$(\text{Kbytes read per sec}) + (\text{Kbytes written per sec})$ Total number of kilobytes read and written each second through host directors associated with the Symmetrix device.
hits per sec	Number of I/O operations performed each second by the Symmetrix device that were immediately satisfied by cache.
ios per sec	Number of I/O operations performed each second by the Symmetrix

	device, including writes and random reads. In contrast, the total ios per sec metric includes writes, random reads, and sequential reads.
Kbytes read per sec	Number of kilobytes read by the Symmetrix device each second.
Kbytes written per sec	Number of kilobytes written by the Symmetrix device each second.
max write pending Threshold	Maximum number of write-pending slots available for the Symmetrix device. Max Write Pending Threshold is not a static number. It depends on Symmetrix activity. Each Symmetrix device is assigned a limit of write-pending slots that can dynamically change between a base value and a value three times the base (the maximum value). Once the Max Write Pending Threshold has reached three times the base value, writes to the device are delayed so that the cache can destage, which frees the cache slots. As cache slots are freed, the writes resume. While the write-pending limit is reached, disk directors operate in a priority destage write mode. This gives write data higher priority than usual. During the delay, writes to this Symmetrix device are counted as write misses.
read hits per sec	Number of random read hit I/O operations performed each second by the Symmetrix device. Read hits occur when the read I/O operation is satisfied directly from cache. The read hits per sec metric for the Symmetrix device statistic does not include sequential read hits. In contrast, the metric total read hits per sec includes random and sequential read hits per sec.
read misses per sec	(reads per sec) - (read hits per sec) Read misses occur when the data requested by the read I/O operation was not in cache, causing a wait while it was fetched from disk.
random reads per sec	Number of I/O operations performed each second by the Symmetrix device that were random reads. This Random Reads per sec metric for the Symmetrix device statistic does not include sequential reads. In contrast, the metric total read hits per sec includes random and sequential read hits per sec.
sampled average reads missed time (ms)	Completion time of a read-miss as measured by the host director. Measurements are taken for a sample set of approximately 30% of the I/Os. Read misses occur when the data requested by the read I/O operation was not in cache, causing a wait while it was fetched from disk. Metadevices: This is approximately (within milliseconds) the same response time that the host sees at the head of the meta. The value is calculated by dividing the total sampled elapsed time by the number of sampled events. The number of sampled events counter, and sampled elapsed time counter are maintained for all metamembers under the statistics for the head of the meta.
sampled average read time (ms)	Completion time of a read as measured by the host director. Measurements are taken for a sample set of approximately 30% of the I/Os.

	Metadevices: This is approximately (within milliseconds) the same response time that the host sees at the head of the meta. The value is calculated by dividing the total sampled elapsed time by the number of sampled events. The number of sampled events counter, and sampled elapsed time counter are maintained for all metamembers under the statistics for the head of the meta.
sampled average write time (ms)	Completion time of a write as measured by the host director. Measurements are taken for a sample set of approximately 30% of the I/Os. metadevices: This is approximately (within milliseconds) the same response time that the host sees at the head of the meta. The value is calculated by dividing the total sampled elapsed time by the number of sampled events. The number of sampled events counter, and sampled elapsed time counter are maintained for all metamembers under the statistics for the head of the meta.
sampled average WP disconnect time (ms)	Completion time of a write-miss as measured by the host director. Measurements are taken for a sample set of approximately 30% of the I/Os. A write miss occurs when a write I/O is held while data is destaged from cache. Metadevices: This is approximately (within milliseconds) the same response time that the host sees at the head of the meta. The value is calculated by dividing the total sampled elapsed time by the number of sampled events. The number of sampled events counter, and sampled elapsed time counter are maintained for all Metamembers under the statistics for the head of the meta.
seq ios per sec	(seq reads per sec) + (seq writes per sec) Use this metric for mainframe performance analysis only.
seq read hits per sec	Portion of the explicitly counted sequential reads that were satisfied from the cache on request. Information about the type of environment you are working in. Mainframe environment: The Sequential read hits per sec is useful as a sequential counter. Open System environment: The SCSI I/O command does not have indicators to identify whether the I/O is part of a sequential stream, so the counts are not actual indicators of full sequential activity. This metric is needed for open systems to calculate the metric total reads per sec. This metric must be added to the read hits per sec to get the value for the Total read hits per sec.
seq reads per sec	Information about the type of environment you are working in.

	Mainframe environment: The seq reads per sec is useful as a sequential counter. Open System environment: The SCSI I/O command does not have indicators to identify whether the I/O is part of a sequential stream, so the counts are not actual indicators of full sequential activity. In an open-system environment the Sequential Read metric indicates to the back end to start pre-fetch activity after two sequential read misses occurred in a row. In the case where the data is in cache and all reads are hits, the back end does not get involved and therefore there is no indication that the read is sequential. This metric is still needed for open systems to calculate the total reads per sec metric. Since sequential read misses are not included in the read counters the seq reads per sec must be added to the Random Reads per sec to get the value for the Total Reads per sec.
seq writes per sec	Number of write I/O operations performed each second that were sequential. Use this metric for mainframe performance analysis only.
system bus Kbytes per sec	$(\text{HA Kbytes transferred per sec}) + (\text{DA Kbytes transferred per sec})$
total DA req per sec	$(\text{DA read requests per sec}) + (\text{DA write requests per sec})$
total hits per sec	$(\text{total read hits per sec}) + (\text{write hits per sec})$ Total number of I/O operations performed each second by the Symmetrix device that were immediately satisfied by cache.
total ios per sec	$(\text{total reads per sec}) + (\text{writes per sec})$ Total number of read I/O and write I/O operations performed each second by the Symmetrix device.
total misses per sec	$(\text{total ios per sec}) - (\text{total hits per sec})$ Total number of missed I/O operations performed each second by the Symmetrix device.
total read hits per sec	$(\text{read hits per sec}) + (\text{seq read hits per sec})$ Total number of read hits and sequential read hit operations performed each second by the Symmetrix device.
total reads per sec	$(\text{random reads per sec}) + (\text{seq reads per sec})$ Total number of random read and sequential read operations performed each second by the Symmetrix device.
write hits per sec	Number of write hit operations performed each second by the Symmetrix device. A write hit occurs when the I/O write operation is received, it is immediately staged in cache and will be written to disk at a later time.

write misses per sec	$(\text{writes per sec}) - (\text{write hits per sec})$ This number of write misses that occurred for the Symmetrix device each second. A write miss occurs when the I/O write operations are delayed.
write pending count	Number of cache slots that were write pending for the logical volume at a point in time. This number changes according to the cache destage activity rate and the number of writes. A write is pending when it has been written to cache but has not yet been written to the disk.
writes per sec	Number of write I/O operations performed each second by the Symmetrix device.

9.9.2 Symmetrix disk metrics

The following table lists and defines metrics for Symmetrix disks.

% utilization	This metric shows how close the disk performance is to 100% utilization. Percentage is based on the number of SCSI commands being performed by the disk.
average hypers per seek	$(\text{seek distance per sec}) / (\text{seeks per sec})$ Average number of hypervolumes that the disk head crossed during one seek.
average Kbytes per read	$(\text{Kbytes read per sec}) / (\text{read commands per sec})$
average Kbytes per write	$(\text{Kbytes written per sec}) / (\text{write commands per sec})$
average read Time	$(\text{total read time per sec}) / (\text{read commands per sec})$
average write Time	$(\text{total write time per sec}) / (\text{write commands per sec})$
Kbytes read per sec	Number of kilobytes read by the Symmetrix disk each second.
Kbytes written per sec	Number of kilobytes written by the Symmetrix disk each second.
read commands per sec	Number of read operations performed each second by the Symmetrix disk.
seek distance per sec	Number of hypervolumes that the disk head crossed (during all seeks) each second.
seeks per sec	Number of times each second that the disk head moved to find data.
skip mask commands per sec	Number of skip mask commands performed each second by the disk. Skip Mask enables a fragmented track, which is chunks of data with gaps in between, to be read from disk or written to disk in one I/O operation. Without Skip Mask, each chunk of data is read or written in a separate I/O operation, causing a fragmented track to require multiple I/O operations to read or write.
total SCSI commands per sec	Total number of read commands, write commands, skip mask commands, verify commands, XOR write commands, and XOR write-read commands performed by the Symmetrix disk each second.

verify commands per sec	Number of verify commands performed each second by the disk. Verify commands are background scrubbing operations performed periodically to check data integrity on the drives.
write commands per sec	Number of write operations performed each second by the Symmetrix disk.
XOR write commands per sec	Number of exclusive OR (XOR) write commands performed each second by the disk. XOR commands are used to establish parity protection in RAID-S and RAID 5 configurations. In a RAID-S or RAID 5 environment, 3 physical I/Os for each write operation are generated: write, XOR-write-read, and XOR-write. When a write is issued to a disk, which is a member of a RAID-S or RAID 5 set, the write is immediately followed by an XOR-write-read command. The XOR-write-read command extracts the necessary information for parity protection. The information is then written to the parity disk using the XOR-write command.
XOR write-read commands per sec	Number of exclusive OR (XOR) write-read commands performed each second by the disk. XOR commands are used to establish parity protection in RAID-S and RAID 5 configurations. In a RAID-S or RAID 5 environment, 3 physical I/Os for each write operation are generated: write, XOR-write-read, and XOR-write. When a write is issued to a disk, which is a member of a RAID-S or RAID 5 set, the write is immediately followed by an XOR-write-read command. The XOR-write-read command extracts the necessary information for parity protection. The information is then written to the parity disk using the XOR-write command.

9.9.3 Symmetrix disk director metrics

The following table lists and defines metrics for disk (back-end, or dadir) directors.

% read	$100 * (\text{reads per sec} / \text{requests per sec})$ Percentage of read requests performed by the disk director.
% seq reads	$100 * (\text{prefetched tracks per sec} / (\text{prefetched tracks per sec} + \text{reads per sec}))$ Percentage of read requests that were sequential. A sequential read occurs when the tracks brought into cache by prefetch tasks are read in the order in which they were put in cache.
% utilization	This metric shows how close the disk director (back-end director) is to

	100% performance utilization. Percentage is based on the number of I/O operations being performed by the disk director.
% write	$100 * (\text{writes per sec} / \text{requests per sec})$ Percentage of write requests performed each second by the disk director.
average lru0 fall through time	Average time it takes a cache slot in LRU0 to be freed up. It is the average time from the first use of the contents to its reuse by another address. This is a key metric that Symmetrix algorithms use to manage the cache. It is a measure, on average, in seconds, of how long it takes a slot to go from the top of the LRU to the bottom of the LRU for reselection. The time <i>only</i> measures from the last time the slot is placed at the top of the LRU to the time it reaches the bottom of the LRU. Each time a slot is placed at the top of the LRU, the time value used to calculate the Fall Through Time is reset to current time (effectively, zero). Average is calculated using between 4000 - 8000 slots. This metric is a sometimes misunderstood Symmetrix statistic. The cache algorithms may have placed the slot at the top of the LRU many times before the slot finally reaches the bottom of the LRU. The total time the slot is in the cache would have to account for all of the times the slot was placed at the top of the LRU. This would equal cache residency time (time of reselection minus the time of entry). A common mistake is to think that Fall Through Time is cache residency time. Cache residency time does not exist in Symmetrix statistics.
ios per sec	Number of I/O operations performed each second between the Symmetrix disk director and the disk.
long misses per sec	Number of times each second that requested data was not in cache, and was not in the process of being fetched.
permacache requests per sec	Number of requests between permacache and the disk director each second. Permacache is the part of cache reserved for a particular application. A request is the protocol used by directors (host and disk) to marshal data to and from the cache.
prefetched tracks per sec	Number of tracks prefetched to cache each second by the disk director.
reads per sec	Number of read requests performed each second by the Symmetrix disk director.
requests per sec	Number of request operations performed by the disk director. A request is the protocol used by directors (host and disk) to marshal data to and from the cache. This metric is made up of reads and writes.
short misses per sec	Number of times each second that requested data was not in cache, but was being fetched when the request was executed.

tracks not used per sec	Number of prefetched tracks per sec that have not been used to satisfy a read request. For more information, see prefetched tracks per sec
tracks used per sec	Number of prefetched tracks per sec that were used to satisfy read requests. For more information, see prefetched tracks per sec
writes per sec	Number of write requests performed each second by the Symmetrix disk director.

9.9.4 Symmetrix host director metrics

The following table lists and defines metrics for ESCON, fibre, Ficon, and SCSI host (front-end, or sadir) directors.

% hit	$100 * (\text{hits per sec} / \text{requests per sec})$ Percentage of requests performed by the host director and immediately satisfied by cache.
% read	$100 * (\text{reads per sec} / \text{requests per sec})$ Percentage of read requests performed by the host director.
% read hit	$100 * (\text{read hits per sec} / \text{reads per sec})$ Percentage of read hit requests performed by the Symmetrix host directors. Read hits occur when the read I/O operation is satisfied directly from cache.
% utilization	This metric shows how close the host director (front-end director including Escon, fibre, or SCSI) is to 100% performance utilization. Percentage is based on the number of I/O operations being performed by the host director.
% write	$100 * (\text{writes per sec} / \text{requests per sec})$ Percentage of write requests performed each second by the host director.
device write pending events	Number of times, each second, that the write-pending limit for a specific Symmetrix device was reached. When the limit is reached, additional write I/O operations are deferred while waiting for data in cache to be destaged to the disk.
hits per sec	Number of read and write requests performed each second by the host director that was immediately satisfied by cache.
ios per sec	Number of I/O operations performed each second by the Symmetrix host director. This metric represents activity between the Symmetrix and the host or SAN device.
read hits per sec	$(\text{requests per sec}) - (\text{writes per sec}) - (\text{read misses per sec})$ Number of read hit requests performed each second by the Symmetrix host directors. Read hits occur when the read requests is satisfied directly from cache.

read misses per sec	Number of read miss requests performed each second by Symmetrix host director. Read misses occur when the data requested by the read I/O operation was not in cache, causing a wait while it was fetched from disk.
reads per sec	Number of random read requests performed each second by Symmetrix host director.
requests per sec	Number of cache slots accessed each second by host directors. A request is the protocol used by directors (host and disk) to marshal data to and from the cache. Requests are made up of reads and writes.
slot collisions per sec	Number of slot collisions each second. A slot collision occurs when two or more directors try to access the same cache slot and the slot happens to be locked for an update operation by one of the directors.
system write pending events	Number of times each second that write activity was heavy enough to use up the system limit set for write tracks occupying cache. When the limit is reached, writes are deferred until data in cache is written to disk.
writes per sec	Number of write requests performed each second by the host directors.

9.9.5 Symmetrix port metrics

The following table lists and describes metrics for Symmetrix ports.

% utilization	This metric shows how close the port performance is to 100% utilization. Percentage is the throughput in kilobytes each second being performed by the port divided by a fixed number based on the type of port.
average io size in Kbytes	$(\text{throughput in Kbytes per sec}) / (\text{ios per sec})$ Average number of kilobytes transferred through the Symmetrix port per I/O operation.
ios per sec	Number of I/O operations performed each second by the Symmetrix port.
throughput in Kbytes per sec	Number of kilobytes transferred through the Symmetrix port (entering the Symmetrix system) each second.

9.9.6 Symmetrix SRDF metrics

The following table lists and describes metrics for the RDF directors (RA1s, RA2s, and RFBs).

% utilization	This metric shows how close the RA1, RA2, and RFB directors are to 100% performance utilization. Percentage is based on the number of I/O operations being performed by the directors.
device write pending events per sec	Number of times each second that the write-pending limit for a specific Symmetrix device was reached. When the limit is reached, additional write I/O operations are deferred while waiting for data in cache to be destaged to the disk.
hits per sec	Number of requests performed each second by the host RA or RFB directors for which: Read data was in cache, or Write operation was sent directly to cache without waiting for data to be

	destaged from cache to the disks.
ios per sec	Number of I/O operations performed each second by the RA1, RA2, or RFB director.
Kbytes received per sec	Number of kilobytes received each second from the RLDs (Remote Link Director). If you are analyzing a source (R1) volume: Number of kilobytes received each second from the target (R2) volume.
Kbytes sent per sec	Number of kilobytes sent each second between the RLD (remote link directory). If you are analyzing a source (R2) volume: Number of kilobytes sent each second from the target (R1) volume.
read misses per sec	Number of read miss requests performed each second by the RA or RFB director. Read misses occur when the data requested by the read I/O operation was not in cache, causing a wait while it was fetched from disk.
read requests per sec	Number of read requests performed each second by the RA or RFB director.
requests per sec	Number of request operations performed by the RA1 or RA2 director. A request is the protocol used by directors (host and disk) to marshal data to and from the cache. This metric is made up of reads and writes.
slot collisions per sec	Number of slot collisions each second. A slot collision occurs when two or more directors try to access the same cache slot and the slot happens to be locked for an update operation by one of the directors.
system write pending events per sec	Number of times each second that write activity was heavy enough to use up the system limit set for write tracks occupying cache. When the limit is reached, writes are deferred until data in cache is written to disk.
write requests per sec	Number of write requests to the RA or RFB director each second.

9.9.7 Symmetrix SRDF/A metrics

The following table lists and describes metrics for SRDF/A sessions.

Active Cycle Size	Number of tracks in the active cycle.
Average Cycle Time	Number of seconds for the average time of all cycles in the session.
Average Restore Time	Average number of seconds since the last restore.
Cycle Number	Number identifying the SRDF/A cycle.
Duration of Last Cycle	Number of seconds for the last cycle in the session.
HA Writes	Number of tracks written by HA to the R1.
HA Repeat Writes	Number of repeat write operations by HA to the R1.
High Water Mark	Number of times each second that I/O activity met the session limit setting.
Host Write IO/s	Number of I/O write requests each second for all devices in the group.
Host Write MB/s	Number of host I/O write requests (in megabytes) written each second for all devices in the group.
Inactive Cycle Size	Number of tracks in the inactive cycle.
Last Restore Time	Number of seconds since the last restore.
Local Write Pending Count	Number of cache slots with writes pending to a local device.

Max Host Throttle	Number of seconds that host writes to SRDF/A devices will be throttled before dropping SRDF/A sessions when cache memory is full. Throttling will delay write operations with the host until a cache track is available. Possible values 0 through 65535.
Max Write Pending Slots	Number of system write pending cache slots available for the SRDF/A session.
Minimum Cycle Time	Minimum number of seconds to wait before attempting an SRDF/A cycle switch. Value ranges from 5 to 59 seconds.
Number of Device in Session	Number of user devices in the SRDF/A session.
R1 to R2 Lag Time	Number of seconds that the data on the R2 is behind the data on the R1.
RA Group	Number of the RDF (RA) Group.
RA Group IO/s	Number of I/O operations each second across the link for the RA Group.
RA Group MB/s	Number of I/O operations each second (in megabytes) across the link for the RA Group.
RDFA Active Indicator	Number indicating that the SRDF/A session is active (1) or inactive (0).
RDFA Max Cache Usage %	Percentage of system write pending cache available for the SRDF/A session. Value ranges from 0 to 100 percent.
RDFA Session	Number of the active SRDF/A session. Inactive sessions do not have an assigned type.
Session Number	Number of the SRDF/A session.
Session Priority	Number of the priority used to determine which SRDF/A session to drop when cache memory becomes full. Value ranges from 1 to 64, with 1 representing the highest priority SRDF/A session (last to be dropped).
System Write Pending Count	Current write pending count. Number of cache slots in the Symmetrix array with a write pending flag.
Time Since Last Switch	Number of seconds since the last switch from active to inactive or from inactive to active. The Metrics view title includes SRDF/A <Symmetrix Identifier/ RA Group / Session Number>; this view also includes the date and time that the collection began.
Total Cache Slots In Use	Number of cache slots in the active and inactive cycle for the SRDF/A session.
Uncommitted Tracks	Number of tracks for all of the SRDF/A devices in the session that have not been committed to the R2 device.

9.9.8 Symmetrix system metrics

The following table lists and describes metrics for the total of Symmetrix systems or units.

% hit	100 * (hits per sec / total ios per sec) Percentage of I/O operations performed by all of the Symmetrix devices, for which the read data was in cache and the write operation could be sent directly to cache without having to wait for data to be destaged from cache to the disks.
% reads	100 * (total reads per sec / total ios per sec) Percentage of total read I/O operations performed each second by all of

	the Symmetrix devices.
% sequential io	$100 * (\text{seq reads per sec} / \text{total ios per sec})$ Percentage of sequential I/O operations. Sequential I/O operations are a combination of the read and write operations. A sequential read occurs when the tracks brought into cache by prefetch tasks are read in the order in which they were put in cache. A sequential write occurs when the data is written to cache sequentially.
% writes	$100 * (\text{writes per sec} / \text{total ios per sec})$ Percentage of total write I/O operations performed by all of the Symmetrix devices.
deferred writes per sec	A deferred write is a write hit. A deferred write occurs when the I/O write operations are staged in cache and will be written to disk at a later time.
delayed dfw per sec	A delayed deferred fast write (DFW) is a write-miss. A delayed DFW occurs when the I/O write operations are delayed because the system or device write-pending limit was reached and the cache had to destage slots to the disks before the writes could be written to cache.
destaged tracks per sec	Number of tracks written per sec from cache to disks.
device max wp limit	Maximum number of write-pending slots for the entire Symmetrix. System write-pending limit is equal to 80% of the available cache slots. Symmetrix write-pending limit is not simply a sum of all Symmetrix device write-pending slots. It depends on other factors such as cache size and the Symmetrix configuration.
hits per sec	Number of I/O operations performed each second by all of the Symmetrix devices for which the read data was in cache and the write operation could be sent directly to cache without having to wait for data to be destaged from cache to the disks.
ios per sec	Number of I/O operations performed each second by all Symmetrix devices including writes and random reads. In contrast, the total ios per sec metric includes writes, random reads, and sequential reads.
Kbytes read per sec	Number of kilobytes read by all of the Symmetrix devices each second.
Kbytes transferred per sec	$(\text{Kbytes Written per sec}) + (\text{Kbytes Read per sec})$ Number of kilobytes written and read by all of the Symmetrix devices each second.
Kbytes written per sec	Number of kilobytes written by all of the Symmetrix devices each second.
number format pending tracks	Number of format pending tracks.
number free permacache slots	Number of available reserved cache slots.
number used permacache slots	Number of use reserved cache slots.
number write pending tracks	Number of tracks in cache that are waiting to be destaged to disk and cannot be overwritten.
prefetched tracks per sec	Number of tracks prefetched to cache each second across the Symmetrix system.
read hits per sec	Number of read hit I/O operations performed each second by all of the Symmetrix devices. Read hits occur when the read I/O operation is

	satisfied directly from cache.
reads per sec	Number of random reads performed each second by all of the Symmetrix devices. The reads per sec metric for the Symmetrix device statistic does not include sequential reads. In contrast, the total reads per sec metric includes the random and sequential reads each second.
seq reads per sec	Number of sequential read I/O operations each second. A sequential read occurs when the tracks brought into cache by prefetch tasks are read in the order in which they were put in cache.
total ios per sec	(Random Reads per sec) + (seq reads per sec) + (writes per sec) Total number of read I/O and write I/O operations performed each second by all of the Symmetrix devices.
total reads per sec	(Random Reads per sec) + (seq reads per sec) Total number of random read and sequential read operations performed each second by all of the Symmetrix devices.
write hits per sec	Number of write hits performed each second by all of the Symmetrix devices. A write hit occurs when the I/O write operation is received, it is immediately staged in cache and will be written to disk at a later time.
writes per sec	Number of total writes performed each second by all of the Symmetrix devices.

9.10 Appendix C – Switch Metrics

9.10.1 Connectivity device port metrics

The following table lists and defines Connectivity device port metrics, which include Connectivity device port errors.

****Note** McDATA ED-1032 switches use 32-bit counters for statistic values. These values cause gaps in Workload Analyzer (WLA) Archiver reports. To work around this issue, use a newer version of the switch with 64-bit counters.

BB Credit Zeros per sec	Total buffer-to-buffer activity detected each second by the port on the switch.
Class3 Discards	Number of Class 3 discards received each second by the port on the switch.
Class2 Frames per sec	Number of Class 2 frames received each second by the port on the switch.
Class3 Frames per sec	Number of Class 3 frames received each second by the port on the switch.
Error Frames	Number of error frames received each second for all the ports on the switch.
Frames Truncated	Number of truncated frames received each second for all the ports on the switch.
Frames per sec	Number of frames received each second for all the ports on the switch.

Invalid CRCs	Number of invalid cyclic redundancy checks received each second for all the ports on the switch.
Invalid Tx Words	Total number times each second that frames containing invalid words are transmitted for all the ports on the switch.
Kbytes Received per sec	Number of kilobytes received each second by the switch or switch port.
Kbytes Transferred per sec	(Kbytes Received per sec) + (Kbytes Transmitted per sec) Number of kilobytes transferred each second by the switch port.
Kbytes Transmitted per sec	Number of kilobytes transmitted each second by the switch or switch port.
Link Failures	Total number of link failures for all the ports on the switch.
Link Resets	(Rx Link Resets) + (Tx Link Resets) Total number of link resets received and transmitted by all the ports on the switch.
Loss of Signal	Number of times each second that loss of signal occurs for all the active ports on the switch.
Loss of Synchronization	Number of times each second that loss of synchronization occurs for all the ports on the switch.
Offline Sequences	Number of offline sequence occurrences for all the ports on the switch.
Receive Link Utilization	$100 * (\text{Kbytes Received per sec} / \text{Port Speed Kbytes per sec})$ Percentage of the receiving link that is used from port to port. Port Speed Kbytes per sec: Maximum number of kilobytes available to the port for I/O operations.
Received Frames per sec	Number of frame counts received each second by the port on the switch.
Rx Link Resets	Number of link resets received by all the ports on the switch.
Transmit Link Utilization	$100 * (\text{Kbytes Transmitted per sec} / \text{Port Speed Kbytes per sec})$ Percentage of the transmitting link that is used from port to port.
Transmitted Frames per sec	Number of frames transmitted each second by the switch and switch port.
Tx Link Resets	Number of link resets transmitted each second by all the ports on the switch.

9.11 Appendix D – HDS Metrics

9.11.1 HDS Device metrics

Not all of the metrics are collected for the HDS 9570V Series (Thunder) or Tagmastore USP family arrays. [Click here](#) to find out which metrics are not collected for the HDS 9570V Series (Thunder) and Tagmastore USP family arrays. If you are viewing an HDS 9570V Series (Thunder) or Tagmastore USP family array, the unavailable data views and metrics are grayed out.

There may be a mismatch between the HDS devices shown in Performance Manager as compared to those in the Console. The Storage Agent for HDS retrieves performance data for the HDS array from the Hitachi HiCommand Tuning Manager (HTM) server, while it retrieves configuration data from Hitachi HiCommand server. If the number of devices reported back from the Hitachi HTM server does not match the number of devices reported back from Hitachi HiCommand server, then HDS devices available in Performance Manager will be different than the HDS devices seen in the Console.

% read hits	$100 * (\text{read hits per sec} / \text{reads per sec})$ Percentage of reads that are hits.
% reads	$100 * (\text{reads per sec} / \text{ios per sec})$ Percentage of I/O operations that are reads.
% write hits	$100 * (\text{write hits per sec} / \text{writes per sec})$ Percentage of writes that are hits.
% writes	$100 * (\text{writes per sec} / \text{ios per sec})$ Percentage of I/O operations that are writes.
average Kbytes per io	$(\text{Kbytes transferred per sec} / \text{ios per sec})$ Average I/O size in kilobytes for the interval.
average Kbytes per read	$(\text{Kbytes read per sec} / \text{reads per sec})$ Average read size in kilobytes for the interval.
average Kbytes per write	$(\text{Kbytes written per sec} / \text{writes per sec})$ Average write size in kilobytes for the interval.
ios per sec	$(\text{reads per sec} + \text{writes per sec})$ Number of I/O operations performed each second.
Kbytes read per sec	Number of kilobytes read each second.
Kbytes transferred per sec	$(\text{Kbytes read per sec} + \text{Kbytes written per sec})$ Total kilobytes transferred each second.
Kbytes written per sec	Number of kilobytes written each second.
read hits per sec	Number of read hits each second.
reads per sec	Number of reads each second.
write hits per sec	Number of write hits each second.
writes per sec	Number of writes each second.

9.11.2 HDS port metrics

Port metrics are not collected for the HDS 9570V Series (Thunder) or Tagmastore USP family arrays. If you are viewing an HDS 9570V Series (Thunder) or the Tagmastore USP family array, the unavailable data views and metrics are grayed out.

average ios per sec	Average number of I/O operations performed each second by the HDS port.
average Kbytes per io	$\text{Kbytes transferred per sec} / \text{ios per sec}$ Average I/O size in kilobytes each second.
average Kbytes per sec	Average number of kilobytes transferred each second through the HDS port.
maximum ios per sec	Maximum number of I/O operations performed each second by the HDS port.
minimum ios per sec	Minimum number of I/O operations performed each second by the HDS port.
maximum Kbytes per sec	Maximum number of kilobytes transferred each second through the HDS port.
minimum Kbytes per sec	Minimum number of kilobytes transferred each second through the HDS port.

9.12 Appendix E - Host metrics

9.12.1 Host CPU Metrics

The following table lists and defines Host CPU Metrics.

% idle CPU	Percentage of time that the host CPU had been idle.
% io wait CPU	Percentage of time that the host CPU spent waiting for an I/O operation to complete.
% processor CPU	(% sys CPU) + (% user CPU) + (% io wait CPU) Percentage of time the host CPU is busy.
% sys CPU	Percentage of time that the host CPU spent in system mode.
% user CPU	Percentage of time that the host CPU spent in user mode.

9.12.2 Host device metrics

The following table lists and describes the Host device metrics.

average queue	Average length of the queue for the device.
ios per sec	Number of I/O operations performed each second by the host device.
Kbytes read per sec	Number of kilobytes read each second by the host device.
Kbytes transferred per sec	(Kbytes read per sec) + (Kbytes written per sec) Total number of kilobytes read and written each second by all of the host devices.
Kbytes written per sec	Number of kilobytes written each second by the host device.
reads per sec	Number of read I/O operations performed each second by the host devices.
response time (ms)	(service time (ms) + wait time (ms)) Total of service and wait time in milliseconds required by the host device.
service time (ms)	Amount of time in milliseconds spent servicing a request.
wait time (ms)	Average amount of time in milliseconds spent waiting before an I/O operation began.
writes per sec	Number of write I/O operations performed each second by the host device.

9.12.3 Host logical volume metrics

The following table lists and defines the host logical volume metrics.

The host logical volume statistics are only available if the Veritas Volume Manager is installed.

avg read time (ms)	Average time in milliseconds that it took the host logical volume to perform a read.
avg write time (ms)	Average time in milliseconds that it took the host logical volume to perform a write.
io per sec	(reads per sec) + (writes per sec)
Kbytes read per sec	Number of kilobytes read by the host logical volumes each second.

Kbytes transferred per sec	(Kbytes read per sec) + (Kbytes written per sec) Number of kilobytes transferred each second to and from the host logical volume.
Kbytes written per sec	Number of kilobytes written each second to the host logical volume.
reads per sec	Number of read I/O operations performed each second by the host logical volume.
writes per sec	Number of write I/O operations performed each second by the host logical volumes.

9.12.4 Host HP device metrics

The following table lists and defines the HP device metrics.

average queue	Average length of the queue for the device.
ios per sec	Number of I/O operations performed each second.
Kbytes transferred per sec	Number of kilobytes transferred each second to and from the host logical volume.
response time (ms)	(service time (ms) + wait time (ms))
service time (ms)	Amount of time in milliseconds spent servicing a request.
wait time (ms)	Average amount of time in milliseconds spent waiting before an I/O operation began.

9.13 Appendix F – HDS Metric Collection

The following tables list the Performance Manager metrics provided for the HDS 9570V array series and HDS Tagmastore USP array family [LUNs](#) and [ports](#).

9.13.1 HDS LUN metrics

Metrics	HDS 9570V Series (Thunder)	HDS Tagmastore UPS Family
% read hits	Yes	Yes
% reads	Yes	Yes
% write hits	Yes	Yes
% writes	Yes	Yes
average Kbytes per io	No	Yes
average Kbytes per read	No	Yes
average Kbytes per write	No	Yes
ios per sec	Yes	Yes
Kbytes read per sec	No	Yes
Kbytes transferred per sec	No	Yes
Kbytes written per sec	No	Yes
read hits per sec	Yes	Yes
reads per sec	Yes	Yes

write hits per sec	Yes	Yes
writes per sec	Yes	Yes

9.13.2 HDS/Hitachi port metrics

Metrics	HDS 9570V Series (Thunder)	HDS Tagmastore UPS Family
average ios per sec	No	No
average Kbytes per io	No	No
average Kbytes per sec	No	No
maximum ios per sec	No	No
maximum Kbytes per sec	No	No
minimum ios per sec	No	No
minimum Kbytes per sec	No	No