

IntelliMagic

Storage Intelligence



The IntelliMagic White Paper:

SMI-S for Data Collection of Storage Performance Metrics

December 2010

Summary:

This paper examines the fundamentals of SMI-S by providing a definition of SMI-S, an overview of components and functionality, and a description of information collected. This content is most applicable to enterprise storage in non-mainframe environments.

This reference guide was prepared by:

IntelliMagic BV
Leiden, The Netherlands
Phone: +31 71 579 6000

IntelliMagic Inc
Southlake, Texas USA
Phone: 1 214 432 7920

Email: info@intellimagic.net
Web: www.intellimagic.net

Disclaimer

This document discusses SMI-S data collection.

IntelliMagic products can be used to support all phases of Storage Performance Management processes. Appropriate usage and interpretation of the results of IntelliMagic products are the responsibility of the user.

Support

Please direct support requests to support@intellimagic.net

Please direct requests for general information to info@intellimagic.net.

Trademarks

All trademarks and registered trademarks are the property of their respective owners.

© 2010 IntelliMagic BV

Table of Contents

Chapter 1 Overview.....	4
Chapter 2 Data Collection Approaches	5
Chapter 3 What is SMI-S?	6
Chapter 4 How does SMI-S work?	7
Chapter 5 What type of data can be collected with SMI-S?	8
Performance Data	8
Configuration Data.....	9
Chapter 6 What are the SMI-S Provider configuration requirement	11
6.1 Key Documents	11
6.2 Required Software	12
6.3 Key Configuration Variables	12
Network Ports	12
Security and Privileges.....	13
Collection Intervals	13
Chapter 7 Conclusion.....	14

Chapter 1 Overview

To easily improve storage service levels and cost efficiency, most sites would like to use an integrated solution for managing storage performance. Except for rare cases, there is no longer a significant technical barrier preventing the adoption of a best of breed storage performance management solution such as the IntelliMagic Suite.

The reason for this new flexibility is the market acceptance and maturity of the vendor neutral Storage Management Initiative Specification (SMI-S). Initially released in 2002, SMI-S has been through several iterations, and now, eight years later, is on the 4th generation of the specification. Storage giants such as EMC, IBM, and Hitachi have all been key sponsors of the efforts in addition to other well known companies including 3PAR, NetApp, LSI, Quest Software, Symantec, and Fujitsu.

The need for cross-vendor management tools has driven the adoption of this specification. Initial adoption of the specification by hardware vendors was lukewarm and implementations were immature. Fortunately for users, many of the hardware vendors are now very supportive of the SMI-S standards. Based on the current trend, it is likely that most new storage hardware platforms will include native SMI-S support.

This paper examines the fundamentals of SMI-S including the following topics:

[Chapter 2: Data collection approaches](#)

[Chapter 3: What is SMI-S?](#)

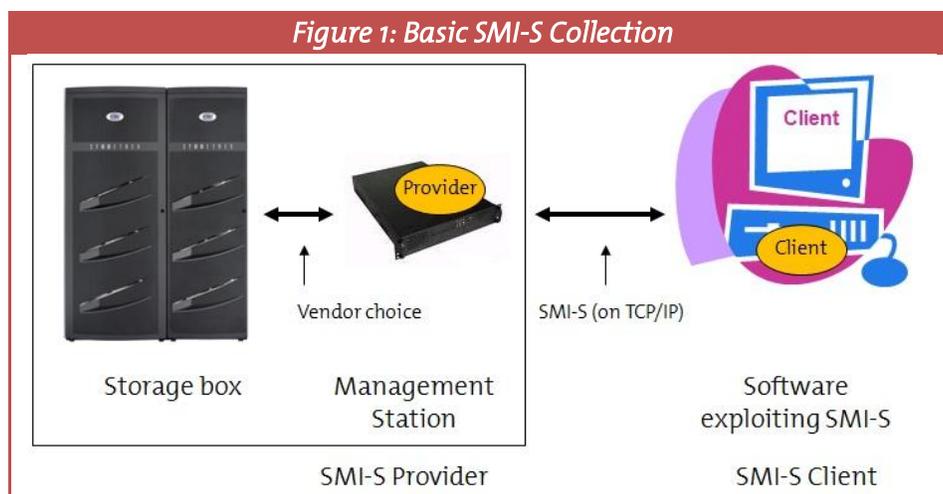
[Chapter 4: How does SMI-S work?](#)

[Chapter 5: What type of data does SMI-S Collect?](#)

[Chapter 6: What are the basic requirements for configuring an SMI-S provider?](#)

[Chapter 7: Conclusion](#)

The intended audience of this paper includes storage administrators, storage architects, and storage managers who are looking to gain a high level understanding of SMI-S. Figure 1: Basic SMI-S Collection, provides a high level example of the SMI-S data collection flow.



Chapter 2 Data Collection Approaches

In order to obtain storage performance and configuration data, tool vendors have, until recently, relied on platform specific APIs. While there are some definite advantages to this approach, there are also some drawbacks.

With SMI-S support becoming pervasive, it is now viable and strategic to consider SMI-S as an alternative for implementing storage management solutions. This chapter discusses the pros and the cons of each approach as illustrated in Table 1: Vendor Proprietary Versus Vendor Neutral.

<i>Table 1: Vendor Proprietary Versus Vendor Neutral</i>		
Feature	Vendor Specific	SMI-S (Vendor Neutral)
Communication between management tool and storage device	Vendor proprietary fast protocol	Standard XML based communication for all exchanges
Support for platform specific metrics	No distinction between standard and platform specific metrics	Protocol requires minimum number of fields and supports vendor extensions for extra fields and components.
Compatibility with other hardware	Specific to one hardware platform only	Any SMI-S supporting hardware
Third party access	Few documented interfaces, different interfaces for each platform	Open specification to any tool vendor
Initial Hardware Provisioning Support	Designed for specific hardware	Difficult to support with SMI-S due to hardware specific implementations.

There are several benefits of using proprietary solutions, including speed and the rich metrics provided. Conversely, there are some drawbacks, including inflexibility, maintenance challenges, expensive vendor specific licensing, and the tendency to be reliant on a hardware vendor for management solutions.

Vendor neutral solutions on the other hand, such as SMI-S, provide standard management interfaces and lower the barriers of entry for software vendors wishing to enter the storage management market. This is an advantage for users, who will benefit from competition in the market place.

One challenge associated with using a standard, is the necessary reliance on common communication models which are often slower. Another challenge is that certification for compliance utilizes a lowest common denominator approach. This means that the minimum requirements for compliance are set at a very low bar. Fortunately, all major vendors have implemented extensions of the SMI-S standard to report the additional metrics that are important for their hardware.

Chapter 3 What is SMI-S?

Storage Management Interface Specification, SMI-S, was created by the Storage Networking Industry Association (SNIA) in conjunction with the Distributed Management Task Force (DMTF) to develop and standardize interoperable storage management technologies.

SMI-S is a common, standards-based management specification that permits third party applications the ability to configure and manage a storage array. Using the “**provider**” (the actual software library), a management application doesn’t require knowledge of the specific architecture or infrastructure requirements of the particular storage platform.

In the SMI-S architecture, client applications communicate with Storage Management Interface Specification (SMI-S) providers, or Common Information Model (**CIM**) agents, to obtain performance and configuration information from storage area networking components such as systems, fabric, and host elements.

SMI-S providers can report about asset, alerts, and performance information, as well as facilitate storage provisioning activities. SMI-S also provides reporting for switch and tape libraries. Each vendor provides a unique provider that facilitates SMI-S based reporting and management for their device.

SMI-S providers can be implemented either as proxies to the devices or as embedded software within the actual storage platform. Most legacy storage platforms have implemented their SMI-S providers as proxies. The proxies are software libraries external from the storage platforms that accept SMI-S queries and commands, and translate them into vendor specific commands which they send to the storage platforms.

As the name implies, the embedded SMI-S providers are included on the storage platforms and do not require the installation or maintenance of a separate software package to provide an SMI-S interface to the storage platform. The trend for the newer platforms is to embed the SMI-S providers within the storage system as evidenced by the latest IBM DS8000 platforms and the EMC V-Max platforms.

For more information on SMI-S, please refer to the following URL:
http://www.snia.org/tech_activities/standards/curr_standards/smi

Chapter 4 How does SMI-S work?

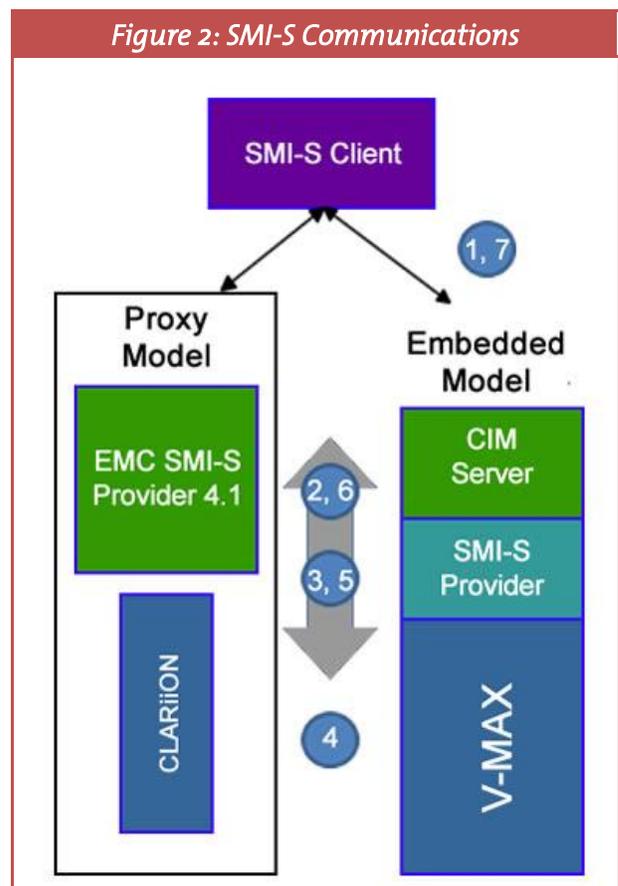
SMI-S uses standard XML based communication methods to facilitate exchanges between the client application, the SMI-S provider, and the storage device.

Figure 2: SMI-S Communications illustrates how an SMI-S client communicates with a CIM Server in an EMC environment.

1. The SMI-S client sends management requests to the CIM server. In an EMC environment, the CIM Server and the SMI-S provider components are both included in the EMC SMI-S Provider 4.1.
2. The CIM Server receives these management requests from the client and communicates with the SMI-S provider.
3. In turn the provider communicates with the storage device using native commands. The provider translates the CIM-XML protocol originating from the CIM Client/CIM Server into commands understood by the storage platform. Upon the completion of a command (4), a confirmation message is returned to the provider (5), then to the server (6), and finally to the client (7).

In most implementations there is no need to be aware of the CIM Server layer. It is almost always packaged with the vendor's provider software, and is often referred to as a part of the collective provider rather than a distinct layer.

EMC's implementation is a perfect example of this transparency. To a storage administrator you just need to know that you are installing the SMI-S Provider, which includes both the CIMOM Server and the SMI-S provider. These complexities will likely become more transparent over time.



Chapter 5 What type of data can be collected with SMI-S?

In addition to management functions, SMI-S enables the collection of both performance and configuration data. This chapter discusses the type of data that can be collected and reported on using SMI-S.

Performance Data

SMI-S uses a template of sorts, called the Block Server Performance Profile (BSP), to define the performance metrics that providers can make available to a management application. The standard template only requires storage subsystem level metrics. Additional components can be accommodated through vendor specific extensions.

The metrics described are:

- Block Server (Top level Computer System), I/O Ports (e.g., FCPorts)
- Front-end Ports and Back-end Ports
- Individual Controllers, Front-end and Back-end controller(s)
- Volumes or Logical Disks
- Extents with association to Pools
- Disk Drives

The metrics defined in the Block Server Performance Profile (BSP) provide much more detailed information about the storage systems than what is available from native distributed platform tools (e.g., *iostat*, *vmstat*, *sar*), but they do not provide performance information from a server perspective.

Due primarily to component differences between the platforms, not all vendors provide the complete set of metrics defined by the BSP. At a storage system level, the metrics in Table 2: BSP Metrics are those metrics required for all platforms.

<i>Table 2: BSP Metrics</i>		
Required Metrics	Definition	Type
Total I/Os	Total of Read I/Os and Write I/Os	Throughput
Kbytes transferred	Total amount of data transferred in Kbytes	Throughput
Read I/Os	Total number of Read I/Os	Throughput
Read hit I/Os	Total number of Read I/Os satisfied by Cache	Cache
Write I/Os	Total number of Write I/Os	Throughput
Write hit I/Os	Total number of Write I/Os satisfied by Cache	Cache

Additional metrics that are commonly available are noted in Table 3: Commonly Provided Optional Metrics.

Table 3: Commonly Provided Optional Metrics

Required Metrics	Definition	Type
Service Time	Refers to the amount of time required to service I/O requests. Sometimes this is provided for both read and writes.	Duration
Transfer Size	Average transfer size in Kbytes.	Throughput

Table 4: Vendor Component Support shows the optional metric categories that some vendors provide.

Table 4: Vendor Component Support

Component	IBM		EMC	
	DS8000	DS4000	Symm/DMX	CX
Storage system	👍	👍	👍	👍
Front-end Adapter		👍	👍	👍
Peer-to-peer			👍	
Back-end Adapter			👍	
Front-end Ports	👍	👍	👍	
Back-end Ports			👍	
Volume	👍	👍	👍	👍
Storage Pool/Rank	👍	👍		
Disk Drive		👍	👍	👍
Arbitrary LU				
Remote mirroring	👍		👍	👍

Configuration Data

In order to provide meaningful context for analysis, SMI-S provides configuration information that shows the relationships between the different logical and physical storage system components. These relationships include, but are not limited to, the relationship between the volumes and the back-end disk devices.

These relationships are critical during the analysis stage particularly as some metrics for higher level constructs (i.e., RAID Group or Extent Pool) are not available from all hardware platforms. By using the relationships between the higher level logical construct and the lower level logical constructs (i.e., Logical Volumes), metrics can be estimated for the higher level logical constructs by aggregating statistics from the lower level constructs. While this technique is imperfect, it provides a reasonable method for “filling in the gaps” and illustrates the criticality of the configuration data.

Logical Component Definitions

Each hardware platform has one or more of the following entities that need to be understood and related to the underlying physical components.

Extent: An extent consists of a collection of physical blocks. The extents are grouped together and assigned to an Extent Pool. In some platforms, extents are transparent in the sense you cannot configure them directly. Extent sizes vary depending on the platform.

Storage Pools: These are groupings of Extents from a set of RAID Parity Groups. These are sometimes referred to as extent pools.

Logical Volume: They refer to a specific grouping of extents from a single extent pool, or in cases where extent pooling is not employed, they consist of physical blocks located on one or more physical drives or RAID Parity Groups. They can be “provisioned” directly to hosts or indirectly as part of a larger Logical Volume Set.

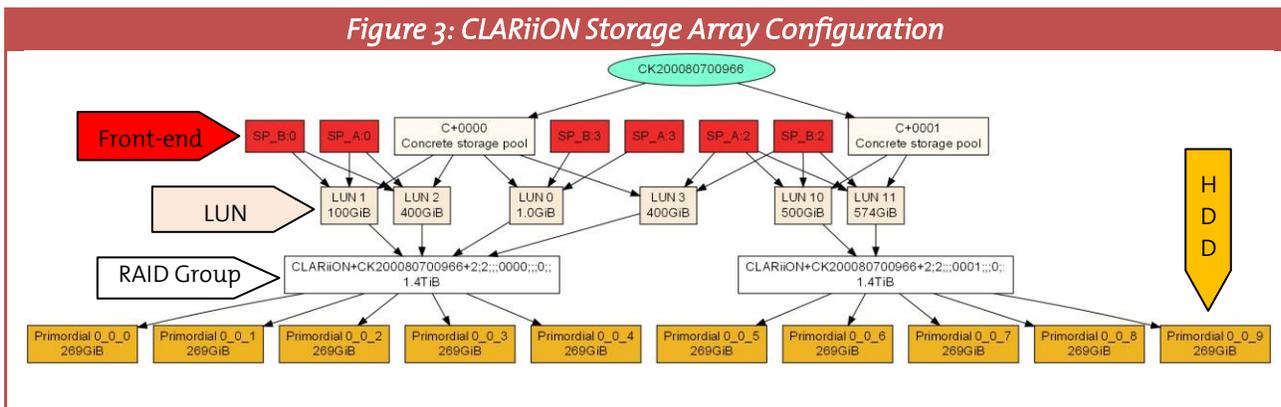
Logical Volume Set: On some storage platforms multiple logical volumes can be coalesced to create a larger volume. These volumes are provided as a discrete unit to the host who sees the Logical Volume Set as a single entity (for example as one LUN). On the back-end disks, the logical volumes can be concatenated or striped to create the larger Logical Volume Set.

RAID Parity Group: A Raid Parity Group is a grouping of RAID formatted physical devices.

The purpose of the configuration details is to describe the relationships between one logical entity and other logical entities or other physical entities such as:

- Which **Physical Drives** make up a **Raid Parity Group**?
- Which **RAID Parity Group** is part of a particular **Extent Pool**?
- Which **Logical Volumes** are defined on which **Extent Pools**?
- Which **Logical Volume** is part of a **Logical Volume Set**?

These relationships are often many layers deep and quite complex, but necessary in order to accurately monitor and identify the root causes of performance constraints. Figure 3: CLARiiON Storage Array Configuration illustrates the complexity of the relationships for a fairly simple CLARiiON storage array. In the case of the CLARiiON platform, a RAID Group consists of two or more RAID-formatted physical disks. From the RAID Group, logical volumes are formed and provisioned to the hosts.



Chapter 6 What are the SMI-S Provider configuration requirements?

The purpose of this chapter is to illustrate some of the key SMI-S provider configuration options, and to provide references to storage vendor specific documentation.

In order for any SMI-S client to collect data from your storage systems, an SMI-S provider must be configured for each storage system that you want to collect data from. Each of the storage vendors supporting SMI-S furnishes installation guides and instructions for their SMI-S provider.

As a general rule the configuration of the SMI-S provider consists of the following key steps:

1. Install the required storage vendor SMI-S provider software
2. Discover and confirm connectivity and authorization between the provider software and the end devices (e.g., Storage system).

URLs for vendor specific documentation can be found in Table 5: Vendor Documentation Matrix. These references provide a starting point for finding the vendor specific documentation:

<i>Table 5: Vendor Documentation Matrix</i>	
Storage Platforms	Documentation URL
EMC	http://Powerlink.EMC.com
HP	http://www.hp.com/support/manuals
Hitachi	http://www.hds.com/assets/pdf/hitachi-storage-command-suite-hitachi-device-manager-software.pdf
IBM DS8000/DS6000/ESS	http://www-01.ibm.com/support/docview.wss?rs=o&uid=ssq1S4000557
IBM DS3000/DS4000/DS5000	http://www.engenio.com/products/smi_provider_archive.html

Note: Hardware vendor URLs change frequently. If the URLs in Table 5: Vendor Documentation Matrix do not work, then try going to the parent URLs and searching for the storage platform and SMI-S, SMIS or CIM agent.

6.1 Key Documents

Each vendor provides documentation for the installation of their SMI-S providers. We have found the documents specified in Table 6: Key Vendor Documents to be helpful in guiding the installation of the SMI-S providers.

<i>Table 6: Key Vendor Documents</i>	
Storage Platforms	Document Name/Reference
EMC	EMC® SMI-S Provider Version X.X Release Notes
HP	HP StorageWorks Command View XP
Hitachi	Hitachi Storage Command Suite Hitachi Device Manager Software SMI-S Provider Installation and User Guide
IBM DS8000/DS6000/ESS	IBM TotalStorage Productivity Center: The Next Generation

Table 6: Key Vendor Documents

IBM DS3000/DS4000/DS5000	LSI SMI-S Provider Installation Guide for Version XX.xx.XX
--------------------------	--

Most of these documents can be found by searching the vendors' web sites displayed in Table 5: Vendor Documentation Matrix.

6.2 Required Software

Each of the SMI-S providers has its own required software depending on the storage platform. Table 7: SMI-S Providers and Storage Platforms demonstrates the necessary provider software for each of the storage platforms listed. Additional pre-requisites may be required depending on the vendor and platform.

Table 7: SMI-S Providers and Storage Platforms

Storage Platform	Provider Software
EMC CLARiiON	CLARiiON Navisphere Release 19, 22, 24, 26, 28, 29 & EMC Solutions Enabler V7.1.0 or later
EMC Symmetrix, and VMAX	Enginuity version 5668 or higher & EMC Solutions Enabler V7.1.0 or later
HP XP 24000	HP StorageWorks XP Command View AE CLI SMI-S
Hitachi USP-V	Device Manager & Device Manager agent 6.3 or later
IBM DS8000	CIM Agent for DS Open
IBM DS3000/DS4000/DS5000	SMI-S Provider (new) Version 10.10.GG.35 or later.

6.3 Key Configuration Variables

There are several common configuration variables that one should be aware of when configuring SMI-S providers. This section attempts to identify the key configuration variables and their default values.

Network Ports

The SMI-S Provider listens for requests from an SMI-S client like the IntelliMagic Vision for SMI-S collector on a default set of network ports. By default, the ports are 5988 for http and 5989 for https. These ports can be changed in the case where there is a conflict with some other service running on the same server. Consult the vendor specific documentation for instructions on modifying ports.

The SMI-S provider or CIM agent may require ports for communication with the storage devices. The ports required for this communication are in addition to the ports used by the CIM agent to SMI-S provider communication. Consult vendor specific documentation for determining the proper ports for the SMI-S provider to end device communications.

Security and Privileges

When configuring an SMI-S provider, it is important to understand the required user ids and their associated access privileges. In order to configure a provider to communicate with a storage device, a valid user ID and password must exist. Table 8: User IDs describes user ID information for known providers.

<i>Table 8: User IDs</i>		
Platform	Provider Software	Required Userids/Groups
EMC CLARiiON	CLARiiON Navisphere EMC Solutions Enabler with SMI-S Provider	ID used by NaviCLI to log on to CLARiiON
EMC Symmetrix, DMX and VMAX	Enginuity version 5668 or higher & EMC Solutions Enabler with SMI-S Provider	ID to access array, Monitor ID can be used.
HP XP	HP Command View XP	Administrator Userid/Password for device
Hitachi	Device Manager	View access to each storage system
IBM DS8000/DS6000	CIM Agent for DS Open (Embedded in R4.1 FW. Prior to R4.1 resides on HMC) DSCIMCLI	Administrator Userid/Password for device (Prior to R4.1 needed additional userid on HMC)
IBM DS3000/DS4000/DS5000	SMI-S Provider (new)	Defaults to authentication turned off, but can be enabled.

Collection Intervals

SMI-S provides an interval mechanism for data collection such that statistics can be queried at any point during the next interval. This allows agents to query information once every interval (3-60 minutes). The SMI-S standard does not allow true real-time collection. The ability to configure the provider collection interval will vary from platform to platform.

Chapter 7 Conclusion

Storage Performance Management (SPM) is growing in importance as data center managers seek to improve storage hardware efficiency and service levels. Mature SPM processes start with the collection of the right performance metrics.

SMI-S provides a robust, vendor neutral standard for collecting storage performance and configuration data. As the SMI-S standard continues to mature, and hardware vendors further integrate SMI-S support into their platforms, more management tools will utilize SMI-S data.

The reduced effort required to collect storage performance data will allow management software vendors to spend more effort innovating their products and solutions. No longer will collecting the data and providing generic charts be sufficient.

The innovations will likely take the form of advanced analytical capabilities with vendor specific analysis of performance and configuration metrics. Best of breed products are differentiated for example, by providing views of the data that quickly and easily identify performance and configuration issues at early stages, often before end-users ever experience degradation. Customers will expect qualitative analysis of the performance of their hardware investments.

IntelliMagic seeks to provide the highest quality Storage Performance Suite on the market through comprehensive data collection and sophisticated analysis that enable data centers to easily progress from re-active SPM processes to more efficient, effective, and lower cost pro-active and predictive processes.

To learn how IntelliMagic can help you achieve your specific storage performance objectives, please contact sales@intellimagic.net.