

Applicability of Standards Track MIBs to Management of World Wide Web Servers

Status of this Memo

This memo provides information for the Internet community. This memo does not specify an Internet standard of any kind. Distribution of this memo is unlimited.

1. Abstract

This document was produced at the request of the Network Management Area Director following the HTTP-MIB BOF at the 35th IETF meeting to report on the applicability of the existing standards track MIBs to management of WWW servers.

Requirements for management of a World Wide Web (WWW) server are presented. The applicable existing standards track MIBs are then examined. Finally, an analysis of the additional groups of MIB attributes that are needed to meet the requirements is presented.

Table of Contents

1.	Abstract.....	1
2.	Overview.....	2
3.	Requirements.....	3
3.1	Operational Model Requirements.....	3
3.1.1.	Host specific and Application Monitoring.....	3
3.1.2.	Dependencies among applications.....	3
3.1.3.	Error generation and reporting.....	3
3.1.4.	Capacity planning.....	4
3.1.5.	Log Digester.....	4
3.2.	Service Model Requirements.....	4
3.2.1.	Retrieval services.....	4
3.2.2.	Document information store -- managing documents.....	4
3.2.3.	Server configuration.....	4
3.2.4.	Server Control.....	4
3.2.5.	Quality of Service.....	4
4.	Relationship to existing IETF efforts.....	5
4.1.	MIB-II [2].....	5
4.2.	Host Resources MIB [3].....	5
4.3.	Network Services Monitoring MIB [4].....	6
4.4.	Application MIB [5].....	7

- 5. Summary of Existing Standards Track MIBs.....8
- 6. Definition of additional attributes.....9
- 7. Usage Scenarios.....11
- 8. Conclusion.....11
- 9. References.....13
- 10. Acknowledgments.....13
- 11. Further Information.....14
- 12. Security Considerations.....14
- 13. Authors' Address.....14

2. Overview

The World Wide Web (WWW) is a network of information, accessible via a simple easy to use interface. The information is often presented in HyperText or multi-media. The information is provided by servers which are located all around the world. The usability of the web depends largely on the performance of these servers. WWW servers are typically monitored through log files. This becomes a difficult task when a single organization is responsible for a number of servers. Since many organizations currently use the Internet Standard SNMP to manage their network devices, it is desirable to treat these WWW servers as additional devices within this framework. This will allow a single Network Management Station (NMS) to automate the management of a number of WWW servers as well as the entire enterprise. Defining a standard for this purpose allows a single management application to manage a number of servers from a variety of vendors. Additionally, a formal definition of what has to be managed and how to manage it tends to lead to integrated and improved performance and fault management.

Content providers are interested in the access statistics and configuration of their sites. The content provider may be the same or a different organization than the one that maintains the server as a whole. It may be possible to realize the new paradigm of "Customer Network Management" to provide this information to the content provider. This means that there exists a distinct organization different than the network operations center that is also interested in the management information from a device. Customer network management is desirable to allow each content provider on a server to access information about his own documents independent of the rest.

Various organizations may be interested in SNMP manageable WWW clients and proxies as well. At this time, our focus is on WWW servers. A natural extension to this work could be a framework for managing WWW Clients and general information retrieval systems like WWW proxies, NNTP, GOPHER, FTP and WAIS. The focus of this document remains the management of WWW servers.

3. Requirements

WWW servers can be viewed from several perspectives when assigning management responsibilities. For the sake of discussion, these perspectives are named the Operational Model and the Service Model. The Operational Model views WWW servers as computers with hardware, disk, OS and web server software. This model represents the actual resources that make up the machine so that it can be monitored from the perspective of resource utilization. The Service Model views the WWW server as a black box that simply handles the responses to requests from clients located on the web.

The two models compliment each other while providing distinct information about the server. Members of the organization responsible for the WWW server, may be interested in one and/or both of the management models. For this reason, the management information should be scalable, for one or both models to be implemented independent of the other.

With this in mind, the requirements for WWW server management can be summarized below by expanding upon those generated at the HTTP-MIB BOF.

3.1 Operational Model Requirements

3.1.1. Host specific and Application Monitoring

This includes monitoring the utilization of CPU, disk and network capacity.

3.1.2. Dependencies among applications.

Some systems implement a number of services within a single piece of code. Others use multiple pieces of code to implement the same set of services. Because of this, dependencies develop among processes. These dependencies become critical when a particular process needs to be stopped, restarted or reconfigured. These dependencies need to be defined within the management information so that management applications can operate the systems correctly.

3.1.3. Error generation and reporting

The WWW server generally reports errors via logging facilities. The format of the log file is not well defined. It is required that a standard facility for error reporting be utilized.

3.1.4. Capacity planning

It is required to obtain statistics which can be used for capacity planning purposes. This includes planning for increased network bandwidth, computing power, disk space, number of concurrent server threads, etc.

3.1.5. Log Digester

WWW servers generally report status information by data generated in Common Log Format [1]. This information needs to be preserved as attributes in a MIB to facilitate remote monitoring providing a standard way to represent and retrieve the management information.

3.2. Service Model Requirements

3.2.1. Retrieval services

Retrieval services are an abstract decoupling the information space from the underlying transport mechanism. The goal at this time is to focus on the requirements for management of WWW servers. There may be considerable overlap with other types of servers like (FTP, NNTP, GOPHER and WAIS). The term "retrieval services" is used here to retain this abstraction. It is required to get statistics about the usage and performance of the retrieval services.

3.2.2. Document information store -- managing documents.

Information from a WWW server can be static (a file) or dynamic (the output of some processing). Management of these two types of information sources range from maintaining access statistics and access permissions to verifying the operational status of all applications that provide the dynamic information.

3.2.3. Server configuration.

It is desirable to be able to centralize configuration management of the servers within an enterprise.

3.2.4. Server Control.

WWW servers generally need to be controlled in regards to starting and stopping them as well as rotating log files.

3.2.5. Quality of Service

Provide an indication of the quality of service the WWW server is providing.

4. Relationship to existing IETF efforts

In general, a WWW server is made up of or depends upon the following components:

- a general purpose workstation running some operating system
- http server software to answers requests from the network
- various support routines like CGI programs or external applications (like DBMS) used to access information
- a document store on one or more storage devices

The health and performance of each of the above components is of interest when managing a WWW server.

There are a number of standards track MIB modules that are of interest to the above list of items. This list includes MIB-II [2], Host Resources MIB [3], Network Service Monitoring MIB [4] and Application MIB [5].

This creates an impressive list of attributes to be implemented. A definition of various levels of management of a WWW server is desired so that the implementor may scale his implementation in chunks which may include various components of each section. For instance, this may allow customer network management without requiring the other groups being implemented.

4.1. MIB-II [2]

MIB-II defines the managed objects which should be contained within TCP/IP based devices.

The WWW server should support the applicable portions of MIB-II. This set probably includes, as a minimum, the following groups: system, interfaces, udp, icmp, tcp and snmp.

4.2. Host Resources MIB [3]

This MIB defines a uniform set of objects useful for the management of host computers independently of the operating system, network services, or any software application.

The MIB is structured as six groups; each specified as either "mandatory" or "optional". If ANY "optional" group of the MIB is implemented, then ALL "mandatory" groups of the MIB must also be implemented. This may cause implementation problems for some developers since many of these attributes require intimate knowledge of the OS.

The groups defined by the MIB are:

-System Group	Mandatory
-Storage Group	Mandatory
-Device Group	Mandatory
-device types	
-device table	
-processor table	
-network table	
-printer table	
-disk storage table	
-partition table	
-file-system table	
-file-system types	
-Running Software Group	Optional
-Running Software Performance Group	Optional
-Installed Software Group	Optional

The system group provides general status information about the host. The storage and device groups define the information about the configuration and status of the resources which compose the host. It defines the resources which make up a generic host system and how they relate to each other. Much of this information is useful for managing various aspects of a WWW server, like the file system and CPU utilization. This information is useful for meeting the operational requirements. Much of this information is however more detailed than many WWW server managers require for service level requirements.

The remaining groups define software components which are installed and/or running on the host. Performance information is defined which extends that defined for each running process. Unfortunately, the mapping between running software and installed software is difficult since it is related by a foreign key (Product ID) which does not appear to be required to exist in either table [6]. There is no provision to represent a group of processes which together perform some task (IE an application made up of multiple processes). The Applications MIB WG plans to address these deficiencies.

4.3. Network Services Monitoring MIB [4]

This MIB is one of three documents produced by the MADMAN (Message And Directory MANagement) Working group. It defines a set of general purpose attributes which would be appropriate for a range of applications that provide network services. This definition is from the perspective of the service without considering the implementation in terms of host computers or processes. Attributes provide

statistics and status on the in-bound and out-bound associations that are currently active, and which have been active.

This MIB is intended to be the minimum set of attributes common across a number of Network Service Applications. Additional attributes are to be defined as necessary to manage specific network service applications. WWW servers clearly fall into the category of network service applications. All attributes in this MIB are relevant to WWW servers.

The MIB consists of two tables:

-applTable	Mandatory
-assocTable	Optional

The applTable describes applications that provide network services and keeps statistics of the current number of active associations and the total number of associations since application initialization. The assocTable contains more detailed information about active associations.

The other two MIBs defined by MADMAN, MTA MIB [7] and DSA MIB [8], are not relevant to the management of WWW services. They do, however, demonstrate how to extend the Network Services Monitoring MIB for a specific set of applications.

4.4. Application MIB [5]

The Application MIB WG is defining two separate MIBs: the sysApplMib and the applMib. The first defines attributes that can be monitored without instrumenting the applications. The second will define additional attributes requiring application instrumentation.

The sysApplMIB allows for the description of applications as a collection of executables, and files installed and executing on a host computer. The objects support configuration, fault and performance management of some of the basic attributes of application software.

The groups defined in the sysApplMIB are:

-System Application Installed Group	Mandatory
-sysApplInstalledTable	
-sysApplCfgElmtTable	
-System Application Run Group	Mandatory
-sysApplRunTable	
-SysApplPastRunTable	
-sysApplElmtRunTable	
-sysApplElmtPastRunTable	

The sysApplInstalledTable captures what applications are installed on a particular host and the sysApplCfgElmtTable provides information regarding the executables and non executable files which collectively compose the application. The sysApplRunTable contains the application instances which are currently running and the sysApplPastRunTable contains a history about applications which have previously executed on the host. The sysApplElmtRunTable contains the process instances which are currently running and sysApplElmtPastRunTable contains a history about processes which have previously executed on the host.

It should be noted that two implementations of the same set of network services may each define a different set of processes and files within this MIB. Ultimately enough management information is needed so that these different implementations can at least be managed similarly.

WWW servers fall into the general category of application software. Therefore the attributes of this MIB are applicable if the process level detail is requested to meet the Operational Model requirements.

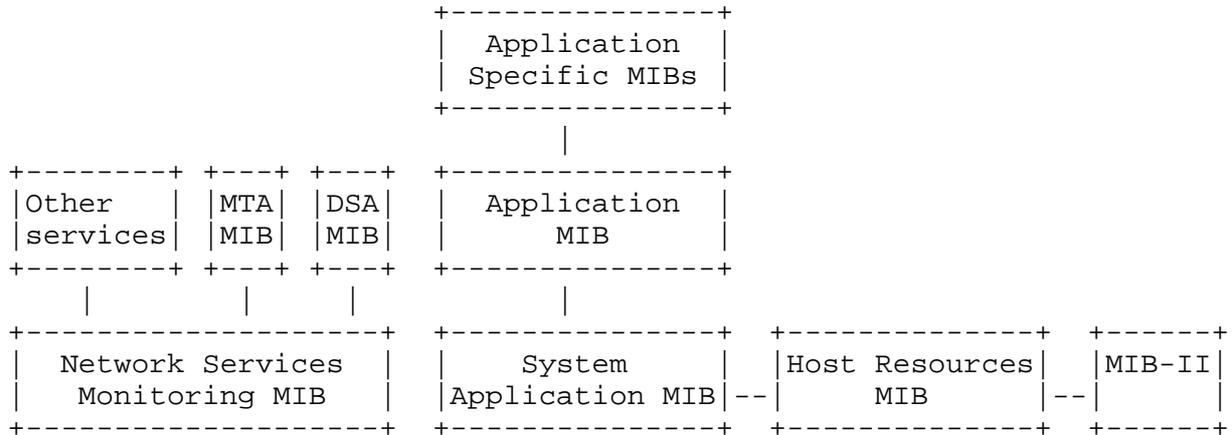
The Application MIB WG is to resolve the problems described above with the relationship between the running and installed software of the Host Resources MIB.

5. Summary of Existing Standards Track MIBs

The existing MIBs are largely orthogonal as demonstrated by the diagram below. Host Resources relates network information to the interfaces defined in MIB-II. The system application MIB relates its running element table to the equivalent entry in the Host Resources running software table.

It should be noted that the running software of the Host Resources includes ALL software running on the host, while the running element table of the system application MIB only includes "interesting" processes of monitored applications.

In the diagram below, "Other Services", "Application Specific MIBs" and "Application MIB" represent work to be done or in progress.



The stack of MIBs above "Network Services Monitoring MIB" represent monitoring from the Service Model. The other stacks represent monitoring from the Operational Model. Neither of these stacks goes to the level of specific detail for any application. The author is of the opinion that HTTP or Web Server specific MIBs would exist at the top of each stack to represent the service and implementation view of the server respectively. There should be a relationship between these two perspectives defined so that the correlations between the two perspectives is possible. This relationship would be useful for general application and service monitoring in addition to just web servers. However, it is not of specific interest to either the MADMAN WG or the Application MIB WG. It is therefore suggested that such a relationship is defined in a general case outside of either of those groups that would be applicable for WWW servers as well as for other application to service mappings.

6. Definition of additional attributes

The existing MIB attributes meet the Operational Model Requirement for tracking information specific to a host. Specifically, MIB-II, Host Resources and the Applications MIB address these items. The Network Services MIB addresses a portion of the service model requirement for the decoupling of the information space from the transport mechanism.

Several sets of additional attributes are needed to meet the remaining requirements. These additional attributes may be generally applicable to other network information retrieval services (like FTP, NNTP, GOPHER and WAIS) as well as client and proxy management. Management of these services is not the scope of this document.

These additional attributes can be classified as:

- 1) Definition of relationship between the Network Services Monitoring and Application MIBs. This allows the functional organization of the server to be known. It allows the management application to understand the effect of restarting specific processes on the services provided. This addresses the Operational Model requirement to model dependencies between applications.
- 2) Additions to generic Network Services Monitoring MIB. A draft [9] has already been circulated due to the work of a mailing list and a sample implementation. These attributes list a summary at the service level of the configuration and the health of the server. From this, performance metrics can be observed. In addition, the health of the server in terms of data timeouts is known. These attributes address the requirement for Operational Model tracking of specific activity and the requirement for Service Model retrieval services.
- 3) Document storage and access statistics are needed to address service model requirements.
- 4) Additions to Application MIB are required to address server configuration requirements in the service model.
- 5) Error and fault management attributes are required to address requirements for tracking specific activity of the web server.
- 6) Configuration and Control are items that may be able to be defined in a general way within the applications MIB. If not, a specific definition would be required here.

Of the items listed above, (1) is needed on a general basis. The others appear to the author as WWW server specific unless the scope of the work is opened to WWW clients and proxies as well as other services (like NNTP, FTP, GOPHER and WAIS).

7. Usage Scenarios

The example scenario will be a single host computer which implements WWW services using the "virtual domain" concept. In this model, a single host performs as the WWW server for one or more addresses. For the purpose of example, we will specify that there are three domains being serviced from this host whose WWW servers are:

- www.a.com
- www.b.com
- www.c.com

Some implementations may implement these services as one set of processes that handle requests for each of the addresses. Others may implement these services as a set of processes for each address. This means that the relationship defined between the Network Services Monitoring MIB and Application MIB components of the management information may vary between different implementations of the same configuration.

MIB-II and Host Resources would provide the information about the host including the CPU, disk and network. The Host Resource running table provide information on the processes in the system.

There would be an entry in the Network Services Monitoring applTable for each virtual domain. In addition, the assocTable shows which connections are currently active. An extension to the association table would be helpful to provide information as to what is being transmitted.

The sysApplMib would have entries in its installed software tables for the web server software and each "interesting" component. This should include the server binary, CGI programs, configuration files and possibly the server log files. Depending on the implementation of the server, the processes for each domain may show up in the same or different running software tables.

Additional information as described in the previous section would round out the management information that would be available for the WWW server.

8. Conclusion

A number of currently defined attributes are useful for management of a WWW server. Specifically, MIB-II and Host Resources should be considered for monitoring the health of the machine in terms of host and network configuration and capacity. The Network Services Monitoring MIB and the Application MIBs provide a general framework

to represent the components of the WWW server from both a service and implementation perspective. The Network Services Monitoring MIB suggests that extensions are necessary to cover specific network application monitoring. A set of such attributes can be well defined to provide status information of the WWW server. The Application MIB suggests similar extensions. Some of these attributes may be generic to all applications, and thus be implemented within the scope of the applMib. It is the opinion of this author that there will still remain specific instrumentation for WWW servers that can not, and should not, be covered in the Network Services Monitoring and Application MIBs.

Since the Network Services Monitoring MIB and the Applications MIB represent orthogonal efforts of management, it is desirable to define the relationship between the two in a standard way. This definition is probably more than a simple pointer from one table to another. Since it is outside the scope of either of those efforts, it is this author's opinion that that definition could and should be addressed within the scope of defining management of a specific application (IE WWW servers). This definition although defined for a particular application, should be useful in a general way to describe the relationship between the Network Services Monitoring MIB and the Applications MIB.

Additional attributes are needed in order to meet all of the requirements specified in this document. An IETF standard would prevent independent developments of this effort in many enterprise MIBs. It also allows management applications to control servers from multiple vendors. It is likely that as the work in this area progresses, the management information will be useful for other Network Information Retrieval services (like FTP, GOPHER, WAIS and NNTP) as well.

Finally, the Operational Model and Service Model Requirements lead to two main uses of the management information. Design of the MIB including the usage of the existing MIBs should allow one or the other or both of these models to be implemented in a standard way. This may be desirable depending specifically on the audience of the data, the cost of instrumentation and the resources of the system.

9. References

- [1] Anonymous, "Logging in the W3C httpd", <http://www.w3.org/hypertext/WWW/Daemon/User/Config/Logging.html>, W3C, July 1995.
- [2] McCloghrie, K., and M. Rose, Editors, "Management Information Base for Network Management of TCP/IP-based internets: MIB-II", STD 17, RFC 1213, Hughes LAN Systems, Performance Systems International, March 1991.
- [3] Grillo, P., and S. Waldbusser, "Host Resources MIB", RFC 1514, Network Innovations, Intel Corporation, Carnegie Mellon University, September 1993.
- [4] Kille, S., and N. Freed, "Network Services Monitoring MIB", RFC 1565, ISODE Consortium, Innosoft, January 1994.
- [5] Saperia, J., C. Krupczak, R. Sturm, and J. Weinstock, "Definition of Managed Objects for Applications", Work in Progress.
- [6] Krupczak, C. and S. Waldbusser, "Applicability of Host Resources MIB to Application Management", Empire Technologies, Inc., International Network Services, October 1995.
- [7] Kille, S., and N. Freed, "Mail Monitoring MIB", RFC 1566, ISODE Consortium, Innosoft, January 1994.
- [8] Mansfield, G., and S. Kille, "X.500 Directory Monitoring MIB", RFC 1567, AIC Systems Laboratory, ISODE Consortium, January 1994.
- [9] Hazewinkel, H., E. van Hengstum, A. Pras, "Definitions of Managed Objects for HTTP", Work in Progress.

10. Acknowledgments

This document was produced at the request of the Network Management Area Director following the HTTP-MIB BOF at the 35th IETF meeting to report on the applicability of the existing standards track MIBs to management of WWW servers.

The author gratefully acknowledges the comments of the following individuals:

Ned Freed, ned@innosoft.com
Innosoft, Inc.

Harrie Hazewinkel, hazewink@cs.utwente.nl
University of Twente

Cheryl Krupczak, cheryl@empiretech.com
Empire Technologies, Inc.

Rui Meneses, rui.meneses@jrc.it
Centre for Earth Observation

Jon Saperia, saperia@bgs.com
BGS Systems, Inc.

Juergen Schoenwaelder, schoenw@cs.utwente.nl
University of Twente

Chris Wellens, chrisw@iwl.com
InterWorking Labs, Inc.

11. Further Information

The current status of the HTTP-MIB standardization can be found on the World Wide Web at <URL:http://http-mib.onramp.net/>. An email list is in operation for discussion of this topic. To subscribe, send email to "http-mib-request@onramp.net" with the message body of "subscribe HTTP-MIB".

12. Security Considerations

Security issues are not discussed in this memo.

13. Authors' Address

Carl W. Kalbfleisch
OnRamp Technologies, Inc.
Email: cwk@onramp.net
1950 Stemmons Frwy
2026 INFOMART
Dallas, TX 75207, USA
cwk@onramp.net

Tel: (214) 672-7246
Fax: (214) 672-7275

