Management of the World Wide Web

J. Schönwälder, TU Braunschweig, Bültenweg 74/75, 38104 Braunschweig, Germany
M. Toet, University of Twente, P.O. Box 217, 7500 AE Enschede, Netherlands

Abstract

This paper discusses the state-of-the-art of World Wide Web management. We first identify a set of requirements for the management of the Web. Based on these requirements, we present the relevant Internet standards defined by the Internet Engineering Task Force (IETF) and we compare them with approaches that use the Web itself to make management information available. We evaluate both approaches against the identified requirements and we conclude with remarks how both approaches should evolve in the future.

1. Introduction

The World Wide Web is in wide-spread use today and many people and organizations, among them the organizers of this workshop, rely on its correct operation. The success of the Web is largely based on the development of a common naming scheme which allows to link documents of arbitrary types together. The complexity of the technology underlying the Web has increased in the last years, while at the same time the accessibility, usability and reliability of the Web became more and more important. To ensure that the Web can continue its evolution into the most important world wide information system, standardized facilities to efficiently manage the Web are needed.

Many Web service provider rely on a set of ad-hoc tools to assist in the management and control of resources and applications needed for the provision of Web services. These tools are often made up of a set of special-purpose programs or scripts to analyze log-files, to check the availability of applications and processes, to monitor the usage of network bandwidth, etc. These tools are usually not integrated into management stations, nor do they allow to understand and to deal with the relationships between the various components that influence the accessibility of the Web.

This paper looks at the problem of managing the World Wide Web from several different perspectives. These perspectives take into account not only the Web servers, but also the underlying protocols and systems that are needed for the provision of Web services. In Section 2, we identify the management information requirements from the various groups involved in running the Web. Section 3 describes how the Internet management framework can be applied to manage the Web. It also describes a proposed World Wide Web Management Information Base (MIB) which is currently under development within the IETF application management working group. Section 4 discusses some approaches that use the Web itself to make management information available. We evaluate the approaches and the proposed MIBs in Section 5 by relating them to the requirements identified in Section 2. We conclude with some remarks how the various approaches should evolve in Section 6.

2. Management Information Requirements

This section identifies management information that should be accessible in order to manage the Web. We tackle this problem by first identifying the groups of people that are involved in providing Web services. The relationships between these groups are than used to identify the different roles every group can perform. The third step is to identify the management information requirements for these groups and their roles.

2.1 Involved Groups

The primary goal of the Web is to make information accessible to human users. This immediately leads to the first group, the *users* (U) of the Web. Users can access information which is often available free of any charges.

An organization which makes information accessible to a user is called a *content provider* (CP). The content can be contained in multiple documents of various types. A content provider is usually interested to distribute the content to a large audience via the Web. He wants to get demographic information about the usage of his documents, so that he can improve the content and the presentation of the content.

A content provider is usually not running the Web servers himself. Instead, he has a contract with a *service provider* (SP), who has to ensure that the Web servers are providing the quality of service that has been agreed upon between the content provider and the service provider. A service provider has to control all the resources that are required to process Web requests. He is therefore interested to gather usage statistics for capacity planning, billing and fault management purposes.

A service provider has contracts with one or more *network providers* (NP). A network provider is responsible to supply IP connectivity. A network provider maintains its own collection of network devices and has contracts with other network providers. Many network provider buy transmission links from e.g. telecommunication companies. However, we do not further subdivide a network provider, because there are far more different technological and organizational possibilities than can be handled in this paper.

2.2 Relationships and Roles

The relationship between a network provider and a service provider or a service provider and a content provider is defined by a contract. It should be noted that the division into these groups and the contracts between them are not always explicit, because many Internet sites do not clearly distinguish between these groups. However, many commercial Web services are based on this model and large companies also tend to organize their internal structure according to this model, even if there are no written contracts.

The contracts between the various groups lead to two different roles. The first role is called the *customer role*. A group in a customer role buys services from another group. The provider of this service must make information available which allows the customer to control whether the delivered service meets the requirements defined in the contract.

Groups that provide services to other groups act in a *provider role*. This role requires management information in order to ensure that the provided service meets the customer demands defined in the contract and that the internal resources are used efficiently.

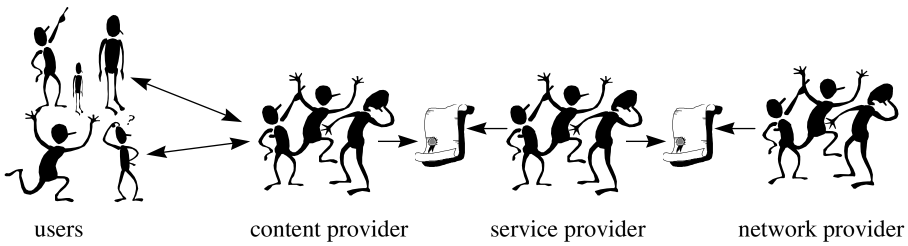

Figure 1: Groups, relationships and roles

The groups, their relationships and their roles are summarized in Figure 1. The relationship between users and content provider can also be seen as a contract, especially if the user has to pay for the contents. However, we focus here on the operational aspects of managing the Web. Questions that deal with user satisfaction are therefore outside the scope of this paper.

2.3 Requirements

We identify management information requirements by looking at the various groups and their roles. This approach results in a better understanding which information should be accessible to whom and who is responsible to provide this information.

CP1	Document access statistics for selected documents (page views).
CP2	Document access pattern for selected documents (path statistics).
CP3	Browser capabilities (browser types, browser capabilities).

Table 1: Content provider requirements (provider role)

We start by looking at the content provider. The management information identified for a content provider in a provider role (Table 1) gives the content provider feedback about the usage of the provided contents. This feedback allows the content provider to tune the content or the presentation of the content to the user needs

CP4	Service availability and downtime distribution.
CP5	Average service response times and error rates.

Table 2: Content provider requirements (customer role)

The management information identified for the customer role (Table 2) allows the content provider to monitor whether the services provided by the service provider are acceptable.

SP1	CPU usage for capacity planning and fault detection.
SP2	Disk space usage for capacity planning and fault detection.
SP3	Network usage for capacity planning and fault detection.
SP4	Availability of the overall server system(s).
SP5	Availability of the processes that make up the server(s).
SP6	Number of transactions handled by the server(s).

SP7	Number of successful and failed transactions.
SP8	Authentication and access control monitoring.
SP9	Error notification and error logging.

Table 3: Service provider requirements (provider role)

The service provider (SP) has more complex needs for management information (Table 3). He has to control a set of network elements and application processes and he wants to use management tools in order to detect failures or resource shortages automatically. The data retrieved from network elements can also be used for trend analysis in order to plan future investments.

SP10	Availability of the local IP network (percentage).
SP11	Connectivity to selected IP network provider.
SP12	Network usage and network utilization.
SP13	Response times to selected IP network provider.

Table 4: Service provider requirements (customer role)

The management information defined for the customer role (Table 4) allows the service provider to monitor the quality of the IP network service delivered by the network provider. The network usage and utilization measures can be used to understand traffic pattern which might result in improved Web server configuration (e.g. number of server processes).

NP1	Availability of the network and the network elements.
NP2	Connectivity to other IP network provider.
NP3	Traffic flows and traffic characteristics.
NP4	Network utilization and protocol specific error rates.
NP5	Subnetwork and network element specific status information.

Table 5: Network provider requirements (provider role)

The management information required by a network provider (NP) in the provider role is shown in Table 5. It allows to monitor the behavior of the network controlled by the network provider. The data can also be used to derive traffic statistics that help to plan the evolution of the network infrastructure.

3. Applying the Internet Management Framework

An approach to tackle the management problem is the application of Internet network management technology [1]. Network management platforms, that are based on this technology, provide integrated environments to detect error conditions, to maintain usage statistics and to assist in troubleshooting and configuration tasks. An integration of Web management capabilities into these management platforms is especially attractive for organizations that already run network management systems.

3.1 Existing IETF Standards

This section briefly describes existing IETF standards that can be used to obtain management information relevant for the management of the Web.

- The *Internet Control Message Protocol* (ICMP) [2] is still the simplest and most reliable way to test the reachability of Internet hosts. In addition, it allows to trace routes and to estimate network response times.
- The *TCP/IP MIB* (MIB-2) [3] contains basic usage and error statistics for the base protocols implemented by an IP network element. MIB-2 lists existing TCP connections together with their state. Network traffic which is sent or received by the network element is counted. This data allows to calculate interface loads which can be used to derive the load on the attached network, depending on the underlying network topology.
- The *Remote Network Monitoring MIB* (RMON-MIB) [4] allows to measure the network traffic on a network segment. The original RMON specification focuses on the network and the media access layers. An extension called RMON-2 [5] has been published recently, which allows to monitor traffic at the application layer.
- The *Host Resources MIB* (HR-MIB) [6] contains information about the hardware devices attached to a host system and the storage devices (including file-systems and main memory). In addition, the installed software and the processes currently running on the host is listed. The information exported by the HR-MIB is based on the abstraction level provided by the operating system.
- The *Network Services Monitoring MIB* (NSM-MIB) [7] describes management information about applications that provide network services. It lists all active associations (which usually refers to TCP connections) with statistics about the total number of associations since application startup time. More detailed information about active associations (used application protocol, description of the remote application, association type and duration) is optionally available.

3.2 Proposals from the Application MIB Working Group

The Application MIB working group of the IETF is working on additional standards for the management of applications. The problem of managing the World Wide Web is addressed as part of this effort. The application management MIBs use a three-level approach. The first level is the system-wide view, which can be instrumented without any changes to the applications running on the system. The second level is the application view, which requires instrumentation in the applications, but uses a MIB structure which is common to all applications. The third level are application-specific MIBs that deal with details that are not covered by the first and the second level.

- The *System Application MIB* (SYSAPPL-MIB) [8] defines the system view (first level) and contains management information about the files that make up an application and the currently running or recently finished applications. The MIB is based on a different abstraction level compared to the HR-MIB since its focuses on applications instead of single processes.

- The *Application Management MIB* (APPL-MIB) [9] defines the application view (second level) and extends the SYSAPPL-MIB by making resource management information (e.g. open files and network connections) available. In addition, status information about the application processes and transaction statistics are provided.

- The *World Wide Web MIB* (WWW-MIB) [10] extends the SYSAPPL-MIB and the APPL-MIB with information specific for the management of Web services (third level). It is organized into three groups, which we will discuss in some more detail.

  The *system group* describes the WWW entities known by the WWW-MIB. An entity can be a Web client, server or a proxy server. Administrative information (e.g. contact person, entity address, protocol used by the entity) is available for every entity. Virtual hosts are represented as separate entities.

  The *statistics group* contains protocol statistics. These statistics are based on an abstract document transfer protocol (DTP) which uses request/response interactions. The abstract DTP can be mapped onto real protocols like HTTP [11] or FTP [12]. The notion of an abstract DTP keeps the MIB definitions independent from protocol details and allows the support of new protocols by providing additional mappings. The statistics group contains detailed statistics about the requests and responses received or generated by an entity. In addition, summary statistics like the total number of requests and responses and the total number of bytes transferred are maintained.

  The *document group* contains information related to the documents that have been transferred by a document transfer protocol. The main purpose of this group is to make document access statistics available for troubleshooting purposes. The document group contains a table which lists details about the last N attempts to access documents. Additional tables list the top N accessed documents over a configurable time interval sorted by the number of accesses or the number of transferred bytes.

One problem, which showed up in the Application MIB working group, is the relationship between a process-oriented and a service-oriented viewpoint. The SYSAPPL-MIB and the APPL-MIB are in large parts based on a process-oriented viewpoint, where information is organized around processes. The process-oriented view allows an efficient implementation of shared tables in a subagent environment like AgentX [14] since processes can be easily uniquely identified by a process identifier. However, many services are realized by a collection of processes and a single process is sometimes responsible for providing multiple services. An example is the Apache HTTP server, where a single set of processes serves multiple virtual host addresses. This results in a many-to-many relationship between the process-oriented view and the service-oriented view. The service-oriented view leads to implementation problems in a subagent environment because there is no simple and general way to uniquely identify services. This means that independent subagent implementations must agree on a method to identify a service.

Some attempts have been made to make WWW statistics available which span a longer time frame than the uptime of the agent. This turned out to be difficult due to limitations of the SNMP language to describe managed objects (the MIB definition language does not allow to express the time base for a counter in machine readable format) and due to performance limitations of the SNMP protocol. Some experiments have been made for using RMON-2 time-filter tables [5] in order to deal with long-term access statistics. However, these tables have properties that can cause a normal SNMP walk to stay endlessly in a time-filter table. It was finally decided that this behavior is not acceptable for the WWW-MIB.

Other constraints for the design of the WWW-MIB are the expected performance impacts of adding instrumentation to high-end Web server implementations and the privacy implications of collecting demographic data and making it available to network managers.

4. Applying Web-Based Management

This section looks at approaches that use the Web itself to make management information available (web-based management). There are several implementations of this idea. We look at two solutions in more detail:

1. *Log-File Analyzer:* Web servers write log-files in order to log server activities. HTTP servers use a common log-file format, which allows to create log-file analyzers that can be used with different server implementations. One such analyzer is the Analog [15] program, which can create document access statistics, host and domain statistics, browser and referrer reports, error and status reports and summary statistics. About 180 options control which set of reports is generated. Information is filtered from the log-files before calculating the statistics. The default output format uses the Hypertext Markup Language (HTML) [16]. An HTML form interface allows to start the analyzer from a Web browser, which enables normal Web users to use this program from remote hosts.
2. *Server Management Modules:* Another approach to make management information available through the Web is the integration of special management modules into the server, which generates status information on the fly. A well-known example of this technique is the status page of the Apache HTTP server [17], which displays the status of the server processes that are used to handle incoming requests. The main difference to the log-file analyzer is that integrated server management modules usually generate documents with short-term information and that these status pages are closely bound to the server implementation.

Both implementations described above do not make full use of the Web technology available. For example, the Analog log-file analyzer could benefit from a Java-based front-end which allows to display and post-process the statistics in a more convenient way. However, it must be noted that collecting acceptable hit or page view statistics is currently not well supported by the underlying protocols, since Web caches do not report cache hit counts back to the originating server. Some people argue that it is practically impossible to collect accurate statistics by counting all page views on busy Web sites and that statistical methods should be used instead to obtain demographic data [18].

5. Evaluation

This section presents an evaluation of the relevant IETF standards, working group drafts and the approaches described in the previous section against the requirements identified in Section 2. In the tables, the abbreviations refer to the different approaches and the requirements discussed earlier. Marks in brackets indicate partial support.

	CP1	CP2	CP3	CP4	CP5
SYSAPPL				✓	
WWW	(✓)				
ANALOG	(✓)	(✓)	✓		(✓)
APACHE					(✓)

Table 6: Evaluation of content provider requirements

Table 6 shows the information requirements of a content provider. Three of the five requirements are currently not covered by IETF standards. It is currently not possible to obtain information about access pattern (CP2), about browser capabilities (CP3) and service response times (CP5). Even support for accurate page view statistics is limited. Table 6 also shows that a log-file analyzer like Analog may fill this hole. The built-in server statistic module of the Apache server only contributes little information.

	SP1	SP2	SP3	SP4	SP5	SP6	SP7	SP8	SP9	SP10	SP11	SP12	SP13
ICMP										✓	✓		✓
MIB-2			✓							✓		✓	
RMON			✓									✓	
HR	✓	✓		✓	✓								
NSM						✓							
SYSAPPL				✓	✓								
APPL				✓	✓	✓							
WWW						✓	✓		✓				
ANALOG			(✓)			✓	✓		(✓)				
APACHE					(✓)	✓							

Table 7: Evaluation of service provider requirements.

Table 7 shows the coverage of the information requirements identified for the service provider. One requirement needs special attention: the requirement for authentication and access control monitoring (SP8) has not been addressed by any of the MIBs and approaches described in this paper. It is also interesting to note that the NSM-MIB only satisfies one information requirement, while all the other MIBs usually address multiple requirements.

Log-file analyzers like Analog or management modules built into Web servers make statistics about the number of transactions and some error information accessible, but do not provide much information about the 'health' of the host systems running the servers.

	NP1	NP2	NP3	NP4	NP5
ICMP	✓	✓			
MIB-2	✓			✓	✓
RMON			✓	✓	✓

Table 8: Evaluation of network provider requirements

The coverage of the information requirements identified for the network provider is shown in Table 8. The coverage by network management standards is quite good, which is not surprising because these requirements touch classic network management issues. It should also be noted that log-file analyzers or management modules built into Web servers do not contribute in this area.

6. Conclusions

The application of Internet management technology to manage the Web is a good choice for network provider and service provider. This paper looked at three approaches to make management information available.

The existing and proposed IETF MIBs provide quite a bit of the information that is needed to manage the Web infrastructure. Some information can be found in different MIBs, which indicates some overlap. It is however surprising, that none of the existing specifications deals with authentication and access control monitoring. We therefore conclude that this is an important area for future work. There is also not much support for the information requirements of a content provider, due to limitations of the SNMP framework.

Information requirements of a content provider are nowadays best satisfied by log-file analyzers that make results available via the Web itself. This approach has several advantages. Content providers usually do not have a management system and it is thus much easier and cheaper to make management information accessible through the Web. In addition, the content provider usually has little needs to post-process the information, which allows to use human readable formats (e.g. HTML) instead of machine readable formats. Access to long-term statistics via the Web allows to pass complex queries to a server in order to select the relevant statistics, which is usually difficult to implement in the traditional Internet management environment.

The approach to integrate management modules into Web servers as implemented by the Apache server can be extended in the future to better address the needs of content providers as well as service providers. However, it is of little help for a network provider. We believe that the real power of this approach lies in configuration management of Web servers, a topic which has not been touched by this paper. We expect that configuration management of Web servers will be done via Web interfaces and built-in configuration management modules in the future.

Our final conclusion is that all the approaches discussed in this paper have their specific strengths and weaknesses. It is therefore necessary to combine the various MIBs in order to efficiently manage the Web. The analysis presented in this paper should give guidance to designers and developers of Web management systems regarding the strength of existing approaches. We also identified areas that need more attention in the future.

7. Acknowledgments

This paper is based on experiences obtained by a proof of concept study for the Centre for Earth Observation (CEO) of the European Commission [19] and has been heavily influenced by the discussions on the HTTP-MIB and the APPL-MIB mailing lists. We like to thank Harrie Hazewinkel, Eric van Hengstum, Carl Kalbfleisch, Aad von Moorsel, Aiko Pras and Hans Scholtes for their comments on earlier drafts of this paper.

References

[1] M. T. Rose, "The Simple Book - An Introduction to Management of TCP/IP based Internets", Prentice Hall, 1994.

[2] J. Postel, "Internet Control Message Protocol", RFC 792, Sept. 1981.

[3] K. McCloghrie and M. Rose, "Management Information Base for Network Management of TCP/IP-based Internets: MIB-II", RFC 1213, Mar. 1991.

[4] S. Waldbusser, "Remote Network Monitoring Management Information Base", RFC 1757, Feb. 1995.

[5] S. Waldbusser, "Remote Network Monitoring Management Information Base Version 2 using SMIv2", RFC 2021, Jan. 1997.

[6] P. Grillo and S. Waldbusser, "Host Resources MIB", RFC 1514, Sept. 1993.

[7] N. Freed and S. Kille, "Network Services Monitoring MIB", RFC 1565, Jan. 1994.

[8] C. Krupczak and J. Saperia, "Definitions of System-Level Managed Objects for Applications", Internet Draft, Apr. 1997.

[9] C. Kalbfleisch, C. Krupczak, R. Presuhn, and J. Saperia, "Application Management MIB", Internet Draft, July 1997.

[10] H. Hazewinkel, C. Kalbfleisch, and J. Schönwälder, "Definitions of Managed Objects for WWW Services", Internet Draft, July 1997.

[11] R. Fielding, J. Gettys, J. Mogul, H. Frystyk, and T. Berners-Lee, "Hypertext Transfer Protocol - HTTP/1.1", RFC 2068, Mar 1997.

[12] J. Postel and J. Reynolds, "File Transfer Protocol (FTP)", RFC 959, Oct. 1985.

[13] C. Kalbfleisch, "Applicability of Standards Track MIBs to Management of World Wide Web Servers", RFC 2039, Nov. 1996.

[14] M. Daniele, B. Wijnen, and D. Francisco, "Agent Extensibility (AgentX) Protocol Version 1", Internet Draft, June 1997.

[15] S. Turner, "Analog WWW Logfile Analysis", University of Cambridge Statistical Laboratory, 1997.

[16] T. Berners-Lee and D. Connolly, "Hypertext Markup Language - 2.0", RFC 1866, Nov. 1995.

[17] R. Thau, "Design considerations for the Apache Server API", in *Proc. 5th World Wide Web Conference*, May 1996.

[18] J. Pitkow, "In Search of Reliable Usage Data on the WWW", in *Proc. 6th World Wide Web Conference*, Apr. 1997.

[19] H. Hazewinkel, F.P.H. van Hengstum, and A. Pras, "Results of the CEO Project - WWW Management", CTIT Technical Report, University of Twente, Dec. 1996.