Characterization and Measurements of Enterprise Network Traffic with RMON2

Luciano Paschoal Gaspary, Liane Rockenbach Tarouco

Federal University of Rio Grande do Sul
Informatics Institute
Av. Bento Gonçalves, 9500 - Agronomia - CEP 91591-970 - Porto Alegre, Brazil
paschoal@inf.ufrgs.br, liane@penta.ufrgs.br

Abstract. The increasing growth of the amount and complexity of applications and protocols executed on computer networks has hindered the work of their administrators. They need to justify the ever-increasing investments accomplished on network equipment acquisition and on communication links leasing. For such, they must: identify who, and with which purpose, most consumes these resources, know if users and resources are located so that the presence of bottlenecks in the network is minimized and detect if some intruder, by means of a high-layer protocol is trying to invade it. An appropriate and current solution capable to answer these subjects is the use of RMON2, MIB that operates above the link layer, providing information needed to monitor client-server applications and end-to-end communications. This work presents the results of a study accomplished on this MIB, aiming at extracting from it means to control the users' activities, to monitor protocols and applications, to optimize the localization of users and resources and to accomplish security management.

1 Introduction

The investments accomplished in the expansion and maintenance of computer networks have surprisingly grown in the last years. Their popularization brought about the appearance of a high number of distributed applications and protocols, leading managers and administrators to investigate effective management solutions in order to reduce costs, provide high availability and, at least, maintain the quality of service once noticed.

These objectives were first reached through an immediate adhesion to new technologies. A real example was the recent migration that occurred in many companies, which left the old plane Ethernet standard for technologies such as Gigabit Ethernet and ATM. As most of the time users regard the network as an inexhaustible resource, they incorporate more and more applications and protocols to their daily routine, which for the administrator means the need for constant alterations in the network infrastructure.

These modifications involve costs and, therefore, they need to be justified. This is possible if the administrator can answer simple questions such as: which users or departments use the network? When is it most used? Which applications are executed? What are the activities of a certain user? Do users perceive an appropriate level of service? Are resources correctly allocated?

R. Stadler and B. Stiller (Eds.): DSOM'99, LNCS 1700, pp. 229-242, 1999.
© Springer-Verlag Berlin Heidelberg 1999

Answers to these subjects can be obtained with the use of accounting mechanisms. Until recently, the accounting process was precariously accomplished, once manual counting took place. The input and output traffic considering ports, connections and users was acquired by polling each network device. This method presents imprecise results [1], besides being very onerous once it generates a lot of management traffic on the network.

A good alternative to accomplish this task is the use of RMON2, MIB that operates with protocols above the data link layer, providing necessary information to monitor high-layer protocols and distributed applications [2][3]. The information collected by this MIB allows administrators to have a detailed view of the behavior of applications and protocols being executed as well as the resources usage rates and the users who most consume them. With such information, these administrators can redefine network traffic flows aiming at better resources usage. Besides, they can observe who communicates with whom and which applications are being used, which makes possible the establishment of policies to guarantee the appropriate use of the network.

This work presents how the information provided by RMON2 MIB may benefit the maintenance of network control and its usage profile discovery, which is an important task for the company to evaluate if investments on networking technology converge or not for the business interests. The paper is organized as follows. Section 2 presents an overview of RMON and RMON2 [4]. Sections 3,4,5 and 6 present the results of the study. Section 3 describes how to monitor user activities on the network. Section 4 presents how to trace the global usage profile of the network. Section 5 treats of the procedures to be accomplished to determine if users and resources are appropriately positioned. Section 6 describes how to use RMON2 to detect non-authorized users. Finally, in section 7 the final considerations are presented.

2 Overview of RMON and RMON2

In recent years the SNMP standard and the specification of MIB-II have been the dominant mechanism for network management. Software agents embedded in network equipment collect information about network traffic and statistics such as the number of input and output frames. With such information a manager knows the amount of traffic that enters and leaves each monitored device, but he is unaware of the behavior of this traffic in the local network as a whole.

The largest contributions to the group of SNMP standards are RMON and RMON2 specifications [5]. Their use has been increasing the efficiency and reducing costs in remote network monitoring and in protocol analysis. While RMON aims at identifying physical problems in the network looking at traffic from router to router, RMON2 monitores network usage patterns, observing the content of the packets of high-layer protocols and applications.

2.1 Characteristics and Evolution of RMON

The RMON specification is essentially the definition of a MIB. The efforts for its standardization began in 1990 with the creation of the workgroup RMON by IETF.

The proposed standard, RFC1271, was published in November 1991. The first RFC was specifically projected to Ethernet local networks. Later, in 1993, the workgroup proposed extensions for Token Ring in RFC1513. Due to the growing interest of the market, several manufacturers started to implement solutions considering the future standard. In 1995 the RMON MIB was standardized (RFC1757)[4].

RMON solutions operate in agreement with the client-server paradigm [6]. Client is the application executed in the central management station, which presents the obtained results to the user (administrator of the network). On the other hand, server is the agent software embedded in network devices or dedicated monitoring devices called probes, which collect information defined in RMON MIB and analyze the packets from the network. Agents and management application communicate using the SNMP protocol.

With RMON MIB, administrators can collect information from remote segments of the network in order to monitor its performance and detect possible problems. Figure 1 presents the groups of RMON MIB. It provides:

- traffic statistics for a network segment, for a certain host and for host pairs;
- a versatile mechanism of alarms and events that makes possible the configuration of thresholds and the notification of the administrator on changes of the network behavior;
- mechanisms for packet capture.

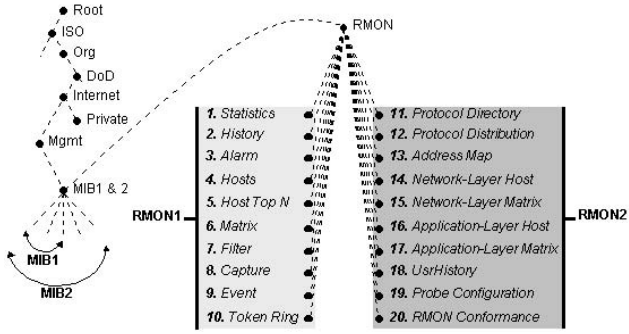

Fig. 1. RMON and RMON2 MIB groups

The RMON standard is defined so that the collection of data and its processing are accomplished by the agents (RMON servers). It can collect and store the value of objects defined in any MIB. When the administrator is interested in such information, he polls the agent only once and receives all information stored in *history* group. This decision in RMON standard contributes to the reduction of SNMP traffic and to decrease the processing accomplished by the management station.

RMON solutions have additional advantages. Without leaving the room, a manager can verify the traffic of a local or remote network segment. Therefore, expense of time and resources are almost quite eliminated because, in most cases, it is not necessary to send specialists to remote sites. A specialist in a central site can work in several problems requesting information to several probes located in remote places. Besides, several specialists can cooperate in the resolution of a certain problem requesting information to the same probe.

It is important to highlight that an RMON probe can monitor the whole traffic

within the LAN segment to which it is attached. It can capture all the MAC-layer frames and identify their source and destination addresses. Thus, the probe is able to provide detailed information about the amount of frames sent and received by each host of the LAN segment. However, if a router is attached to the LAN, one can neither determine the real source of the traffic that arrives from it nor the real destination of frames leaving the LAN through the router. This restriction is solved by RMON2 standard [6].

2.2 High-Layer Protocols Monitoring

The works to extend RMON MIB and include mechanisms to monitor higher-layer protocols began in 1994. This initiative, called RMON2, resulted in the creation of RFC2021 in January 1997 [2]. When monitoring high-layer protocols such as network and application-layer protocols, it is possible to visualize the whole corporate network instead of individual segments. The groups defined in RMON2 MIB are showed in figure 1. Briefly, they are:

- protocol directory (*protocolDir*): repository that indicates all the protocols that the probe is capable to interpret;
- protocol distribution *(protocolDist):* statistics about the amount of traffic generated by each protocol observed by the probe;
- address map *(addressMap):* associates each network-layer address to the respective MAC address, storing them in a table;
- network-layer host *(nlHost):* collects statistics about the amount of input/output traffic of the hosts based on their network-layer addresses;
- network-layer matrix *(nlMatrix):* provides statistics about the amount of traffic between host pairs based on their network-layer addresses;
- application-layer host *(alHost):* collects statistics about the amount of input/output traffic of the hosts based on their application-layer addresses;
- application-layer matrix *(alMatrix):* provides statistics about the amount of traffic between host pairs based on their application-layer addresses;
- user-history collection *(usrHistory):* periodically samples objects specified by the user (manager) and stores the collected information in accordance with parameters also defined by the user;
- probe configuration *(probeConfig):* defines configuration parameters for RMON probes;
- rmon conformance *(rmonConformance):* describes conformity requirements for RMON2 MIB.

3 Administration and Statistics on User Activities

If the administrator is interested in tracing the network usage profile it is very important to obtain information about how a certain user or department uses it. What users most use the network? With whom do these users communicate? Which applications and protocols do they execute? Such information enable the administrator to observe users' activities and verify if they fit the company interests.

3.1 Volume of Accesses

Information related to the volume of network accesses accomplished by a certain user can be obtained through requests to RMON2 *network-layer host* group. It decodes packets based on their network-layer addresses. Thus, administrators can observe beyond the routers that connect the sub-networks and identify the real hosts that are communicating [7]. The *network-layer host* group is composed of a control table (*hlHostControl*) and a data table (*nlHost*) (see figure 2).

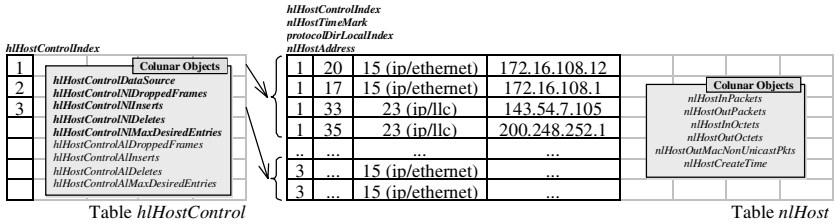

Fig. 2. *nlHost* group tables

The control table has an entry for each interface (sub-network) being monitored. Each entry contains information such as the number of frames received from the interface that the probe decides not to count, number of rows added and removed from the data table and maximum number of data rows (*nlHost* table) acceptable for this interface.

Table *nlHost* provides basic network statistics for each network-layer address seen on a monitored network segment. When a new row is added to table *hlHostControl*, the probe begins to observe packets and to collect network-layer addresses in the respective interface. A new row in table *nlHost* is created for each new discovered address.

Figure 2 shows table *hlHostControl* with three rows. It means that the probe is analyzing packets in three different network segments. As one can observe in table *nlHost*, four hosts were identified in the first segment (first index of the table). The second index (*nlHostTimeMark*) of this table denotes when the row was created or last updated. It allows the management station to retrieve only the updated entries each polling cycle, reducing traffic between this station and the agents. The third index, *protocolDirLocalIndex*, identifies the protocol encapsulation observed. The encapsulations that the probe is able to decode are defined in *protocol directory* group. Finally, the fourth index of this table contains the network-layer address of the monitored device. An example of information retrieved when polling the *nlHost* group is presented in table 1.

Table 1. Example of information retrieved from *nlHost* group

HostAddress	*InPkts*	*OutPkts*	*InOctets*	*OutOctets*	*OutMacNonUnicastPkts*
172.16.108.12	1.000	345	80.345	25.367	33
172.16.108.1	2.350	733	97.334	33.292	125
143.54.7.105	5.930	299	112.445	5.293	0
200.248.252.1	100	30	49.238	3.777	12

Requests to this group may help administrators to find out which users most use

the network and when it occurs. In this context, the *user history* group can be used to store object values in diferent time instants. Thus, if a management station wants to retrieve the value of an object in ten different time instants, only three messages are required: one for configuring the RMON2 probe, one for requesting the report and another for retrieving it (agent response). The savings of this approach are huge when compared to the convencional SNMP MIB, when the same task requires twenty messages (ten requests plus ten responses).

A simple formula to calculate the network usage rate (n.u.r), in percentage, by a certain user (host) during a time interval between t_1 and t_2 is presented below. *ifSpeed* denotes the speed of the network technology of the segment to where the probe is attached.

$$n.u.r = \left(\frac{\left[\left(nlHostInOctets_{t_2} + nlHostOutOctets_{t_2} \right) - \left(nlHostInOctets_{t_1} + nlHostOutOctets_{t_1} \right) \right] \bullet 8}{ifSpeed} \right) \bullet 100 \quad (1)$$

Figure 3 shows an example of a graph, which can be generated by polling the *nlHost* group and applying the formula just presented. In order to minimize management traffic in the network, one may configure a report in *user history* group and retrieve the *nlHost* object values sampled in different time instants only once. The disadvantage of this approach is that the graph can not be gradually generated and therefore partial views of it are not possible.

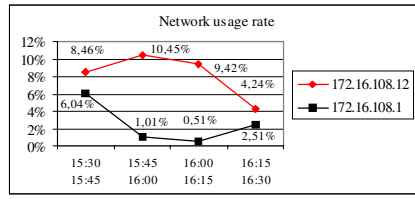

Fig. 3. Example of graph depicting the network utilization rate by two hosts

3.2 Applications and Protocols Used

Determining user network usage patterns relies on the knowledge of the administrator about the protocols and applications that each user executes as well as when it happens [6]. Such information can be obtained in *application-layer host* group.

The *application-layer host* group comprises a control table (*hlHostControl*), which is the same as the *network-layer host* group, and a data table (*alHost*), as shown in figure 4. Table *alHost* provides the administrator with input/output traffic statistics of hosts, considering application-layer protocols. The term application layer refers to all protocols above the MAC layer.

There is one or more entries in table *alHost* for each application-layer protocol discovered. These entries are also organized according to their network-layer address, so that information such as the HTTP traffic generated from or addressed to a certain host may be easily retrieved. Table *alHost* is indexed by the following objects:

- *hlHostControlIndex:* denotes the segment to where the probe is attached;
- *alHostTimeMark:* time filter;

- *protocolDirLocalIndex:* network-layer protocol identification;
- *nlHostAddress:* network-layer address;
- *protocolDirLocalIndex:* application-layer protocol identification.

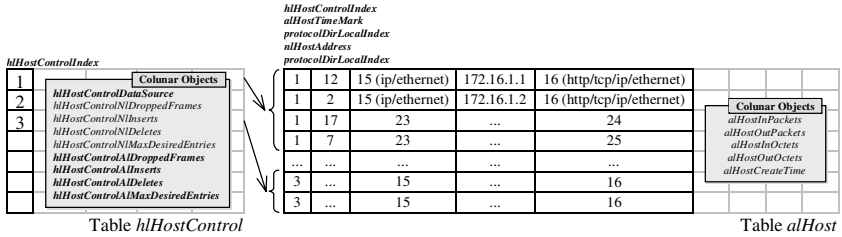

Fig. 4. *alHost* group tables

Observe that the object *protocolDirLocalIndex is* used twice to index the table. In fact, two different instances of this object are used: one of them identifies the network-layer protocol and the other consists of an application-layer protocol [7].

Table 2 illustrates an example of data retrieved from *alHost* group. By periodically polling this group or configuring remote reports to store *alHost* group objects in different time instants, the administrator can find out which applications and protocols users most employ and when exactly it occurs.

Table 2. Example of information retrieved from *alHost* group

HostAddress	*Protocol*	*InPkts*	*OutPkts*	*InOctets*	*OutOctets*
172.16.108.12	HTTP	830	342	45.311	42.543
172.16.108.12	SMTP	250	24	17.900	2.406
172.16.108.12	FTP	567	158	32.193	19.765
172.16.108.5	HTTP	2037	411	209.312	56.927

The information provided by *alHost* group can be used to draw graphs that indicate both the protocols and applications in use by a certain host and the amount of traffic generated by each of them. Figure 5a shows an example of a chart that can be obtained by polling the probe in two time instants (t_1 and t_2). It presents the usage rate of each protocol by the host 172.16.108.1 based on the total volume of packets observed coming from and addressed to it during the interval t_2-t_1. The formula to calculate it follows:

$$a.u.r = \frac{\left(alHostInOctets_{t2\ prot} + alHostOutOctets_{t2\ prot}\right) - \left(alHostInOctets_{t1\ prot} + alHostOutOctets_{t1\ prot}\right)}{\sum_{i}\left(alHostInOctets_{t2\ prot\ i} + alHostOutOctets_{t2\ prot\ i}\right) - \left(alHostInOctets_{t1\ prot\ i} + alHostOutOctets_{t1\ prot\ i}\right)} \bullet 100 \quad (2)$$

Figure 5b shows the protocol usage rate considering four consecutive time intervals. This type of historical graph aids the administrator to determine user access patterns. In this graph, protocol usage rate was calculated based on the speed of the monitored segment. Thus, it is possible to observe the impact that a protocol in use by a certain host is causing to the network. The formula to obtain this information is:

$$a.u.r = \frac{\left(\frac{\left[\left(alHostInOctets_{t2} + alHostOutOctets_{t2}\right) - \left(alHostInOctets_{t1} + alHostOutOctets_{t1}\right)\right] \bullet 8}{ifSpeed}\right)}{t_2 - t_1} \bullet 100 \quad (3)$$

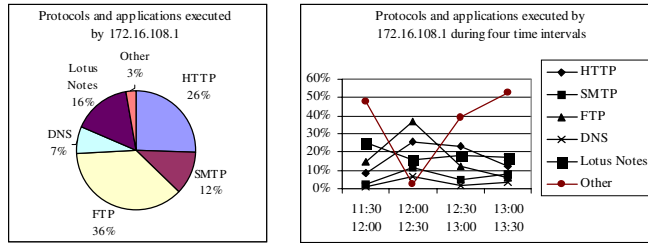

Fig. 5. Example of graphs depicting protocol usage rates

3.3 Established Communications

It is important to identify who are the local/remote peers of each established communication to further understand the behavior of network users. The *application-layer matrix* group has an important role in this process. It collects traffic statistics between communicating host pairs based on their network-layer addresses. The group consists of several tables. Two of them are depicted in figure 6.

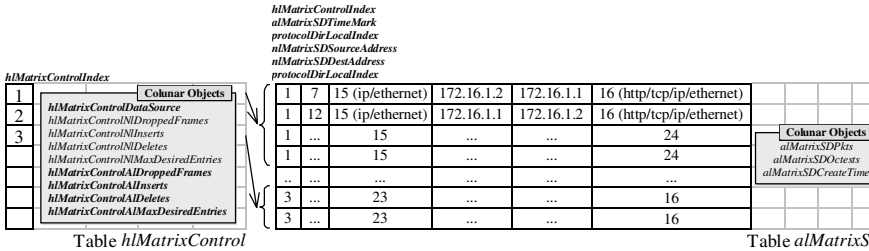

Fig. 6. *alMatrix* group tables

The control table (*hlMatrixControl*) is similar to the control tables already presented. The data table (*alMatrixSD*) stores information about the amount of traffic observed between host pairs. It is indexed by the following objects [2]:

- *hlMatrixControlIndex:* denotes the segment to where the probe is attached;
- *alMatrixSDTimeMark:* time filter;
- *protocolDirLocalIndex:* network-layer protocol identification;
- *nlMatrixSDSourceAddress*: source address;
- *nlMatrixSDDestAddress*: destination address;
- *protocolDirLocalIndex:* application-layer protocol identification.

Table 3 illustrates an example of data that is retrieved when table *alMatrixSD* is polled. As one can observe, this table counts the traffic from source to destination and vice-versa in two different entries. Through periodic requests to the RMON2 probe or by means of configured reports in *user history* group, it is possible to determine with whom a certain host is communicating, the protocols being used and the amount of traffic generated by them. In table 3, the user on host 172.16.108.12 is currently accessing Alta Vista® and the web site of the own company. Besides, it maintains an FTP connection with Microsoft®.

Table 3. Example of information retrieved from table *alMatrixSD*

SDSourceAddress	*SDDestAddress*	*Protocol*	*SDPkts*	*SDOctets*
172.16.108.12	altavista.digital.com	16 (http/tcp/ip/ethernet)	15	578
altavista.digital.com	172.16.108.12	16 (http/tcp/ip/ethernet)	237	17.900
172.16.108.12	ftp.microsoft.com	17 (ftp/tcp/ip/ethernet)	29	2.193
ftp.microsoft.com	172.16.108.12	17 (ftp/tcp/ip/ethernet)	12.033	409.312
172.16.108.12	172.16.108.1	16 (http/tcp/ip/ethernet)	49	5.971
172.16.108.1	172.16.108.12	16 (http/tcp/ip/ethernet)	14.987	1.000.329

4 Network Global Usage Profile

In the previous section, mechanisms to control and to monitor host-based (user) network activities were presented. However, in many situations the administrator must further investigate the network usage profile taking into consideration the whole company or some departments. For instance, such information is essential if the administrator wants to find out the departments that most consume network resources, when the network is overloaded and which users/departments contribute to intensify this problem.

The *protocol distribution* group counts the number of octets and packets monitored by the probe for each supported protocol encapsulation. It consists of two tables: *protocoloDistControl,* which controls the collection of basic statistics, and *protocolDistStats,* which stores the collected data [2]. Each entry of table *protocolDistControl* denotes one network segment and controls a group of entries in table *protocolDistStats*, one for each protocol recognized on this segment (see figure 7).

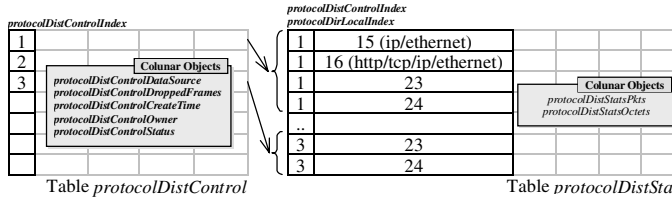

Fig. 7. *protocolDist* group tables

Table 4 shows an example of information retrieved from *protocolDist* group. The accomplishment of periodic requests to this group or the retrieval of historical reports from it makes possible the observation on the variation of protocols usage rate during a certain time period (see figure 8).

Table 4. Example of information retrieved from *protocolDist* group

Protocol	*protocolDistStatsPkts*	*protocolDistStatsOctet*
15 (ip/ethernet)	3.417	1.034.587
16 (http/tcp/ip/ethernet)	1.644	459.100
23 (smtp/tcp/ip/ethernet)	1.290	345.923
24 (ftp/tcp/ip/ethernet)	483	229.564

To determine the network usage rate of each department the administrator must

know which hosts belong to each department. In this case, the group to be polled is network-layer host. The methodology to be used is the following: the probe counts input/output packets and octets for each identified host. At the end of an accounting interval, the network usage rate is calculated for each host using the formula presented in (1). Afterwards, hosts of each department are grouped and their network usage rates are added.

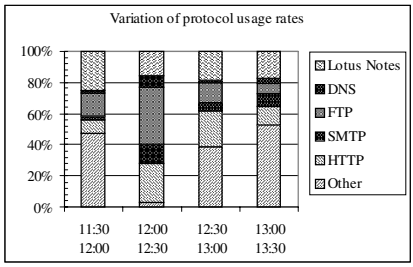

Fig. 8. Protocol distribution during a certain period

The graph in figure 9 was drawn based on object values collected in two different time instants. In this case, the resulting information reflects the peculiar situation of the network in that specific interval.

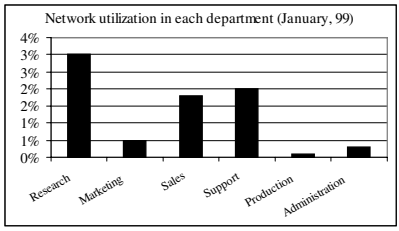

Fig. 9. Departmental network usage rates

Another possibility is to accomplish such measurements in several moments of the day, for several days, and to calculate the average of the obtained rates. This methodology provides the administrator with accurate information about network utilization in each department.

The type of information commented above can be decisive in cost allocation, once it helps the administrator to decide where to invest. In this context, the identification of the hosts that accomplish most of the accesses is also important. Such information can be found in *topN* tables from *network-layer matrix* group.

The *topN* tables provide an efficient way for a management station to obtain a ranked list of matrix table entries based on a chosen network statistic [2][7]. Again, a control table (*nlMatrixTopNControl*) and a data table (*nlMatrixTopN*) are used. Each entry in control table refers to a report under construction and has some configurable objects such as classification criterion (packets or octets) and sampling interval (in seconds).

Figure 10 illustrates the two mentioned tables. As one can see, there are two reports under construction. When the administrator requests the generation of a new report, he informs the probe how big it should be using the object

nlMatrixTopNControlRequestedSize. The probe evaluates its resources (memory, processor) and, if not overloaded, accepts the administrator's solicitation. Otherwise, it arbitrarily informs him the number of entries it will admit (*nlMatrixTopNControlGrantedSize).* In figure 10, the first report has four entries.

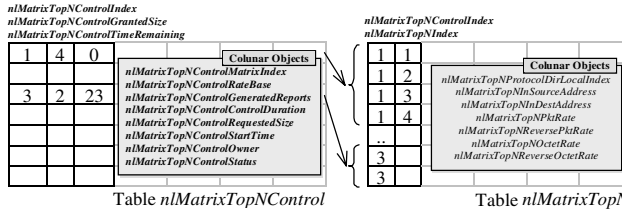

Fig. 10. *topN* tables from network-layer matrix group

It is important to highlight that the object *nlMatrixTopNControlTimeRemaining* denotes the time remaining for the next report to be ready. In figure 10, the first report has just been updated (*TimeRemaining=0*). As soon as the report is concluded, the probe automatically begins a new sampling period, whose duration is defined by the object *nlMatrixTopNControlControlDuration.*

Table 5 presents data retrieved from table *nlMatrixTopN.* These data make possible the creation of graphs as in figure 11, which shows the users who most consume network resources.

Table 5. Example of information retrieved from table *nlMatrixTopN*

Protocol	*Source Address*	*Dest Address*	*TopN PktRate*	*TopNReverse PktRate*	*TopN OctetRate*	*TopNReverse OctetRate*
15 (ip/ethernet)	172.16.108.12	172.16.108.1	213	32	40.065	6.023
15 (ip/ethernet)	172.16.108.45	172.16.108.23	156	17	23.913	2.194
15 (ip/ethernet)	172.16.106.25	200.248.252.1	89	29	12.882	6.745
15 (ip/ethernet)	172.16.109.7	172.16.108.1	67	13	5.294	968

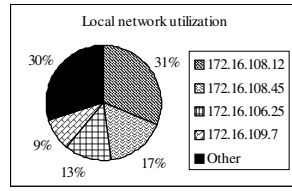

Fig. 11. Users who most consume network resources

5 Optimization of Users and Resources Distribution

An important contribution of RMON2 is the possibility to verify if users and resources are adequately positioned in the network in order to maximize traffic confinement in each department of the company [1]. The group that provides this information is application-layer matrix, previously presented in section 3.3. The procedure in figure 12 can be applied to optimize the location of users and resources.

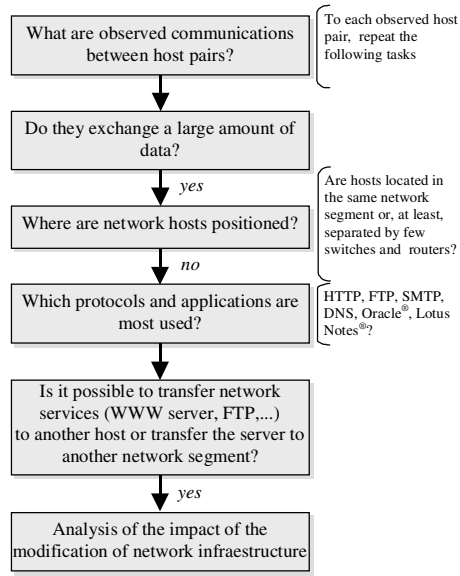

Fig. 12. Algorithm for optimization of users' and resources' location

In some cases, users and resources are appropriately positioned, minimizing traffic between different network segments. However, some resources may be overloaded, which affects the response times of protocols and applications. In such cases, load balancing is needed. To do that, the administrator must measure the usage rate of a certain resource. This information is provided by *application-layer matrix* group.

The methodology used to determine the current activities on a certain resource is the following. The probe counts input/output packets and octets for each identified host pair. The monitoring can be accomplished in two time intervals (instants t_1 and t_2) or periodically. At the end of an accounting interval, all the entries whose source or destination address is the same of the monitored resource are selected. The selected entries are then grouped according to the application-layer protocol. Afterwards, for each group the input/output octets are added.

Table 6 illustrates the methodology just presented. In the monitored resource, one can observe HTTP, SMTP and Oracle traffic. Thus, the rows can be arranged in three groups. In HTTP group, the number of input (34.567) and output (125.954) octets is added in instant t_1(160.521). The same calculation is repeated in instant t_2 (1.060.521). The amount of HTTP traffic observed on the resource is obtained by subtracting the result of the sum of input/output octets in instants t_1 and t_2. The same calculation must be accomplished in the other groups.

Table 6. Network activities on a certain host

Source Address	*Dest. Address*	*Protocol*	*Octets (t_1)*		*Octets (t_2)*		*t_2-t_1*
172.16.108.12	172.16.108.1	HTTP	34.567	160.521	192.224	1.060.521	900.000
172.16.108.1	172.16.108.12	HTTP	125.954		864.297		
172.16.109.5	172.16.108.1	SMTP	37.234	58.123	220.211	314.123	256.000
172.16.108.1	172.16.109.5	SMTP	20.889		93.912		
172.16.108.12	172.16.108.1	Oracle	45.082	68.863	341.090	534.834	465.971
172.16.108.1	172.16.108.12	Oracle	23.781		193.744		

Figure 13 shows an example of a graph, depicting network activities on a network server host along the day.

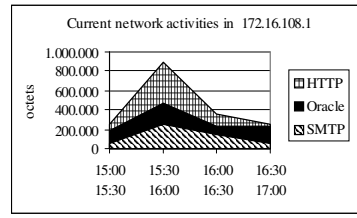

Fig. 13. Activities observed on a network server host

6 Security Management

Security is an essential issue to corporate networks. Nowadays, more and more mechanisms such as firewalls have been incorporated into the network in order to keep intruders away from strategic data of the company [7].

RMON2 can be used as a tool to detect the presence of intruders in the network. As already presented, the *application-layer matrix* group shows host pairs that are communicating (see section 3.3). Therefore, if the administrator periodically monitors this group, he can identify non-authorized users trying to establish communications with network hosts. It is also important to observe the protocol being used. Depending on the security policies of the company, a telnet may represent an attempt to invade the corporate network.

When a suspicious network address is observed, one can start capturing packets generated from this host and analyze the network operations that it is executing. This operation is supported by *capture group* from RMON MIB.

The graph in figure 14 shows the users that are accessing a certain resource. It is adequate to monitor strategic hosts where database, www and mail servers reside. The graph is generated with information retrieved from *application-layer matrix* group.

In some cases, it may be useful to identify which protocols each observed user executes. This information can be obtained in *alHost* group, already presented in section 3.2. Figure 15 shows an example of a graph that can be generated with the retrieved data.

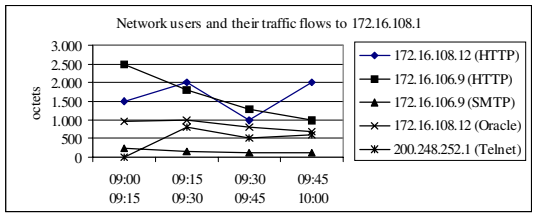

Fig. 14. Users and their traffic flows to a certain host

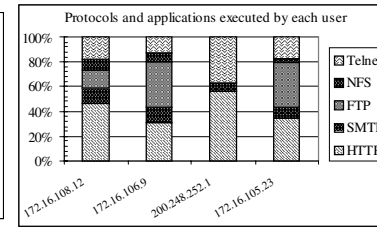

Fig. 15. Protocols executed by hosts

7 Conclusions

RMON2 represents a huge increase in capabilities. In many cases, its functionalities supplant protocol analyzers, trend generation tools and other [3]. This paper presented how to use this MIB to control user activities, to trace the network usage profile, to optimize the location of users and resources and also to accomplish security management. The contribution of this work is to show the administrator how to benefit from this MIB, which is powerful yet very complex.

Some of the objects and tables in RMON2 MIB were designed so that a simple display of their contents provides meaningful information. Most of them, however, must be organized in easy-to-view formats; otherwise they are of little use. Depending on the company investments on network technology, it is not always possible to buy management applications, which do treat such information and automatically convert it to charts and other diagrams. In this context, this work helps network managers to create their own management applications. Through *script* languages or high-level management libraries, it is possible to develop systems adapted to the company's needs without many investments.

The authors believe that RMON2 based solutions may be applied in other scenes such as ISPs. However, the huge amount of information that the probes need to handle requires the use of powerful monitoring devices. Even so, in situation of heavy network load existent probes have faced problems to accurately count the packets observed in the network [8].

References

1. Engel, Fred. Application Behavior and Statistics through RMON2. In Third IEEE International Workshop on Systems Management. New Port, Rhode Island, April 22-24, 1998
2. Waldbusser, S. Remote Network Monitoring Management Information Base Version 2 using SMIv2. Request for Comments 2021, January 1997
3. Gaspary, Luciano. Study on RMON2 Standard. Individual Work. Porto Alegre: Federal University of Rio Grande do Sul, 1998
4. Waldbusser, S. Remote Network Monitoring Management Information Base. Request for Comments 1757, 1995
5. Ulbrich, Luís Roberto. Active Remote Monitoring. Individual Work. Porto Alegre: Federal University of Rio Grande do Sul, 1997
6. Stallings, William. SNMP, SNMPV2 and RMON. Practical Network Management. Second Edition. USA: Addison Wesley, 1996
7. Perkins, David T. RMON - Remote Monitoring of SNMP-Managed LANs. First Edition. USA: Prentice Hall, 1998
8. Newman, D.; Giorgis, T.; Melson, B. Probing RMON2. Data Communications (May 21, 1998)