Foreword

Network Management Systems (NMS) are ubiquitous in Data Centers. The Internet protocols defined in the various Simple Network Management Protocol RFC documents, in short SNMP, are the de facto standard of every NMS. Powerful SNMP Consoles (also known as Managers) like the enterprise grade products CA Unicenter, HP OpenView or Tivoli from IBM, can immediately manage not only the networking gear (routers, firewalls,..) but also the vital parameters of most servers.

Open source consoles, like for instance OpenNMS, are also a viable and pocket friendly solution.

Application Management leverages on the Network Management Infrastructure. By simply defining new cluster of objects to be managed (MIBs), complex applications can be monitored and handled.

While other projects provide good anomaly handling, only a real SNMP console can deliver a proactive monitoring, by means of the historical gathering of data. Trends in the values, comfortably shown in an appealing graphic form, can provide early warnings on the system. Disk space approaching the full state, increasing sluggish responsiveness of an application, growing number of rows in a database can be seen at a glance.

MOSTdoc SNMP Agent

Today MOSTdoc, in addition to syslog and Nagios interfaces, has a SNMP Agent, that collects relevant information about the system, and answers to standard queries from SNMP managers. Like most agents (linux kernel agent, net-snmp,..), data is collected during idle time, and conveniently stored in a file structure. Collect frequency changes according to the usefulness of the information, and to the computational cost. Extracting the number of rows in a table can be a heavy task, so it is performed sparingly. Data collection is performed at best, taking advantage of internal routines, and avoiding unnecessary computations.

When a request arrives, the overhead is negligible, and the answer is sent back to the requester.

Decoupling the data gathering from the protocol oriented module is cheap: most parts are data independent, and so new parameters can be added with ease.

How it was built

The first thing to be defined is the MIB: written according to the SNMP standards, this file describe the parameters of a MOSTdoc environment: systemwide parameters, application data, DB information, valuable timings extracted from the run time logs. The cluster of objects can be linked in any place of the MIB hierarchy, though his natural place is beneath the ***enterprise*** tree.

The second module is a collection of routines that is in charge of collecting the needed information from the system, and storing this data in file structure. Any time a new MIB variable is added, the corresponding piece of code must be added. The file structure is a clone of the MIB tree, scalar and table value places are imposed from the MIB definition. These routines are called during the periodic housekeeping of the system (each fifteen minutes and at midnight)

The third module is the information broker, that answers the queries from the external requesters, reading the information from the run time repository. This routines are completely driven from the MIB definition, so when a new variable is added, no changes are are needed.

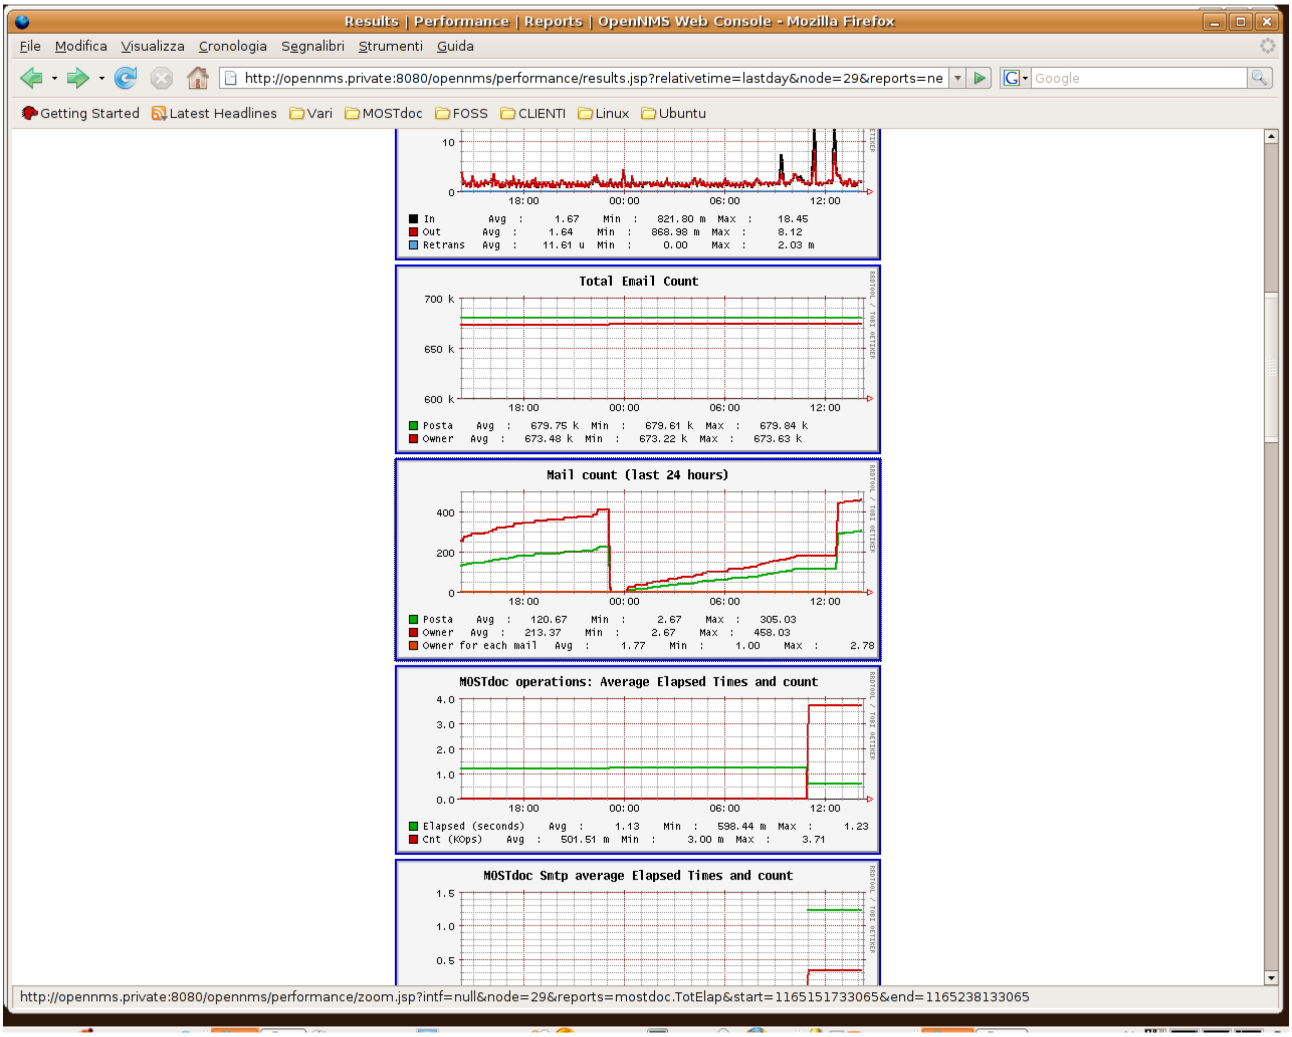
Results | Performance | Reports | OpenNMS Web Console - Mozilla Firefox
File Modifica Visualizza Cronologia Segnalibri Strumenti Guida
http://opennms.private:8080/opennms/performance/results.jsp?relativetime=lastday&node=29&reports=ne
Google
Getting Started Latest Headlines Vari MOSTdoc FOSS CLIENTI Linux Ubuntu
10
0
18:00 00:00 06:00 12:00
700 k
650 k
600 k
400
200
4.0
3.0
2.0
1.0
0.0
1.5
0.5
RRDTOOL / TOBI OETIKER
http://opennms.private:8080/opennms/performance/zoom.jsp?intf=null&node=29&reports=mostdoc.TotElap&start=1165151733065&end=1165238133065

```
■ In        Avg  :    1.67     Min  :   821.80 m  Max  :    18.45
■ Out       Avg  :    1.64     Min  :   868.98 m  Max  :    8.12
■ Retrans   Avg  :   11.61 u   Min  :     0.00    Max  :    2.03 m
```

Total Email Count

```
🟩 Posta    Avg  :   679.75 k  Min  :   679.61 k  Max  :   679.84 k
🟥 Owner   Avg  :   673.48 k  Min  :   673.22 k  Max  :   673.63 k
```

Mail count (last 24 hours)

```
🟩 Posta    Avg  :   120.67    Min  :     2.67    Max  :   305.03
🟥 Owner   Avg  :   213.37    Min  :     2.67    Max  :   458.03
🟥 Owner for each mail   Avg  :     1.77    Min  :     1.00    Max  :      2.78
```

MOSTdoc operations: Average Elapsed Times and count

```
🟩 Elapsed (seconds)    Avg  :     1.13    Min  :   598.44 m  Max  :     1.23
🟥 Cnt (KOps)     Avg  :   501.51 m  Min  :     3.00 m  Max  :     3.71
```

MOSTdoc Smtp average Elapsed Times and count

SNMP Traps

MOSTdoc conveys alarms, warnings and any asyncronous notification to external handlers: syslog, Nagios and SNMP. While syslog is generally available in every unix/linux server, Nagios and Snmp require a working console.