POINT SOLUTIONS IN AN ENTERPRISE MANAGEMENT WORLD: THE NETWORK MANAGEMENT APPLIANCE

LTC S. Mills
U.S. Army V Corps

C. Valentine
H. Werchan
The MITRE Corporation
Army Information Systems Department

Heidelberg, Germany

ABSTRACT

This paper will discuss development and testing of the Network Management Appliance (NMA), a protocol and bandwidth management device that has been developed and used by the U.S. Army V Corps as part of the tactical network management strategy. Previous capabilities have focused on aggregate bandwidth management using SNMP management tools such as the Multi-router Traffic Grapher (MRTG). Unfortunately, in the highly bandwidth-constrained tactical environment, tools such as MRTG only reveal the obvious—saturated links. V Corps required more visibility into what protocols, applications, and users were saturating the links. Based on this requirement, the NMA—a cross between a protocol analyzer and an RMON probe—was developed from COTS hardware and open source software. The information provided by the NMA has enabled V Corps to take a proactive stance in shaping bandwidth utilization and increasing the quality of service to the Warfighter.

BACKGROUND

In the tactical network environment bandwidth is at a premium. The bandwidth requirements associated with the growing number of deployed PCs has far outstripped the available bandwidth, particularly over reach-back circuits that tie the forward deployed networks back to the strategic, garrison and commercial networks. To better manage bandwidth utilization, in 1999 V Corps began routinely using bandwidth-monitoring tools such as MRTG and Round Robin Database (RRD) to measure WAN bandwidth utilization. Upon using these tools, a not-too-surprising trend was observed:

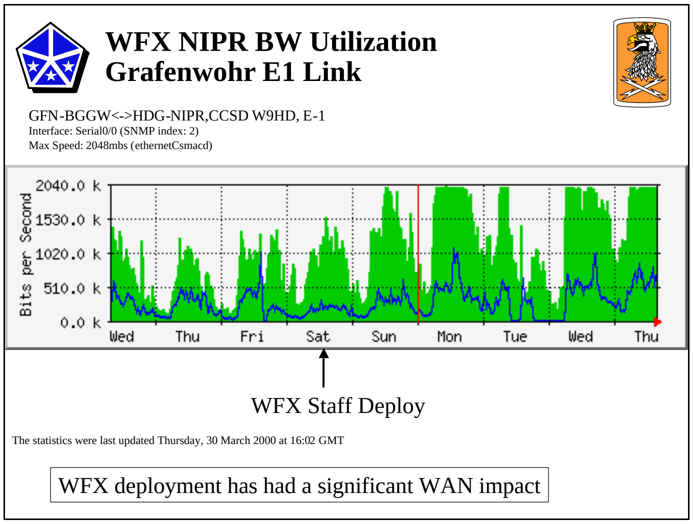

WFX NIPR BW Utilization
Grafenwohr E1 Link

GFN-BGGW<->HDG-NIPR,CCSD W9HD, E-1
Interface: Serial0/0 (SNMP index: 2)
Max Speed: 2048mbs (ethernetCsmacd)

The statistics were last updated Thursday, 30 March 2000 at 16:02 GMT

WFX deployment has had a significant WAN impact

It was not surprising to observe that immediately after deployment of the V Corps headquarters staff, bandwidth was routinely saturated on a [week]daily basis. Here is a similar picture on a subsequent deployment:

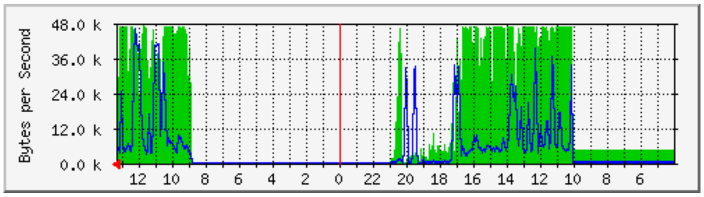

In each of the preceding pictures, the blue graph indicates outgoing bandwidth, whereas the green annotates incoming bandwidth.

As might be expected, the bandwidth utilization graphs only revealed the obvious—saturated links. In order to proactively manage bandwidth more visibility and detail was required. At this point, V Corps began investigating LAN protocol analyzer technology to gauge which users and which applications were consuming the most bandwidth. This initial effort used a commercial product called WebXRay running on a laptop that was collocated on a shared hub that interfaced the local LAN to the WAN router (this was later moved to a mirrored switch interface). Using WebXRay immediately provided the level of visibility to ascertain trends as can be seen in the following charts:[1]

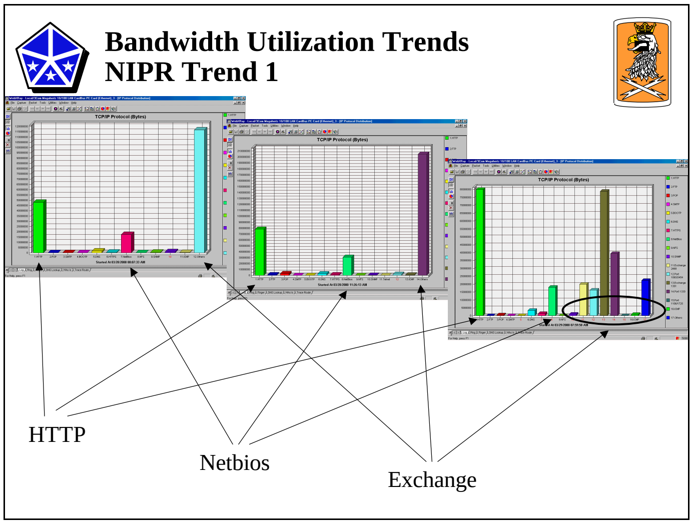

Bandwidth Utilization Trends
NIPR Trend 1

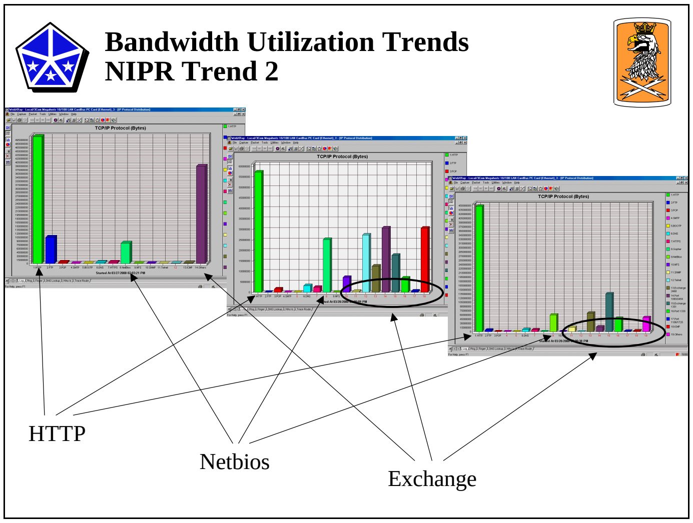

Bandwidth Utilization Trends
NIPR Trend 2

During a recent exercise, WebXRay was used to give us visibility into what protocols were placing the largest bandwidth load on the WAN interface. Unfortunately, as previously mentioned, this approach required manual intervention to collect and archive the data. Consequently, not enough time granularity was captured to be of significant use in post analysis. Towards the end of pre-WFX01, we experimented with NTOP, a server based protocol analyzer that can run automated and unattended. This open source product was fully leveraged in support of the warfighter and gave us the best-detailed view of bandwidth utilization and protocol distribution to date.

In addition to establishing protocol distribution trends, WebXRay was very helpful in identifying network configuration issues and problems (e.g., users using CONUS-based DNS servers for name resolution).

The problem with the WebXRay approach to monitoring protocol distribution is that it is not readily automated and therefore requires manual intervention to gather data and save reports (screen captures). Also, the laptop-based configuration was not readily suited for proper integration into the network data package. Based on these deficiencies we began looking at server-based tools that could more readily be automated. These initial looks including an application called TCPdump and eventually focused on an application called NTOP. NTOP, by design, provides the required reporting capability without user intervention. Its chronology capability is especially noteworthy for V Corps , making both short term and historical analysis much more practical.

WHAT IS THE NMA?

In short, the NMA consists of the following components: A commercial off-the-shelf (COTS) personal computer (PC) running Linux with dual network interface cards (NICs)[2]. In addition, the following open source software was used: Apache Web Server, NTOP, MRTG, the RRD, Sendmail for email notifications, SSH (Secure Shell) and Putty (a Windows SSH client) for secure remote access, Webmin (for remote web-based administration of server), and SNORT for IDS. Some NMA products were "published" via the web using Frontline Team Server (FTS).

The dual NIC PC is placed near the root switch that provides the interface between the LAN and the WAN (usually a router or perhaps a firewall). This switch is configured to mirror the WAN interface to an otherwise unused port to which one of the NMA NICs is attached. The other NMA NIC is attached to a separate switch interface. The use of dual NICs configuration isolates true WAN network traffic from any local, high-speed network traffic that would be generated when retrieving reports from the NMA. In fact, the promiscuous mode NIC monitoring the mirrored switch interface does not need to have a TCP/IP configured.

[1] It is interesting to note that "Trend 1" is a daytime snapshot of protocol distribution revealing Exchange email to be the top bandwidth-consuming application, whereas "Trend 2" is a night-time snapshot revealing web browsing as the top.

[2] The NICs must be capable of promiscuous mode operation

The availability of the open-source NTOP software was the real enabler for the NMA. It was conceived as a network parallel to the Unix TOP application, which is used by systems administrators as a quick glance at which software processes are consuming the most system resources (e.g., CPU and memory). NTOP was originally developed to provide the systems administrator a local perspective of network utilization (e.g., traffic destined to and from the host on which NTOP is running). However, by placing the NMA on a mirrored WAN/LAN interface using a promiscuous mode NIC, a global perspective (from a LAN perspective) of network utilization can be achieved.

NMA CAPABILITIES

Our NTOP-based protocol distribution prototype has been extremely useful as a tool in managing the available V Corps MAIN command post WAN bandwidth.

As an example of this powerful capability, consider the following MRTG screen shot for the V Corps MAIN command post 256Kbps satellite reach-back link.

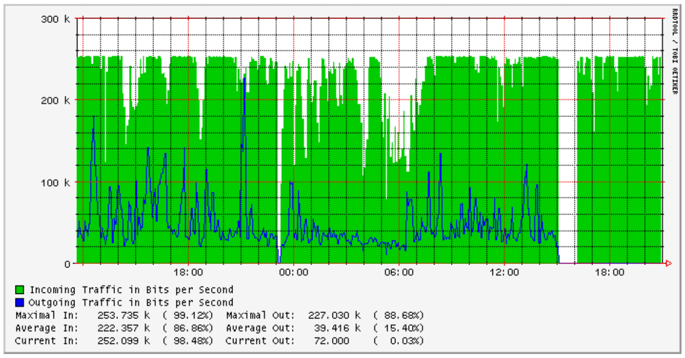

As can be seen from the graph, the link's incoming bandwidth is nearly continuously saturated at 87%, while the outgoing bandwidth averages 15%. This information alone does not provide nearly enough information to be proactive and control the bandwidth utilization by 'shaping' the traffic distribution. For instance, it could be quite possible that nearly all of the bandwidth is being used for web browsing, at the expense of SMTP email transfers, but with this level of information, this is only speculation.

Using NTOP as a complement to MRTG removes this speculation. Below is a graph of the NTOP data for the same approximate time frame as above.

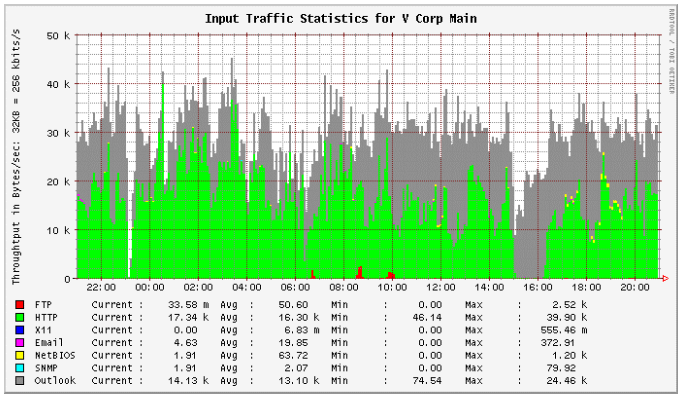

Clearly, this view of the link utilization removes any speculation that web browsing is causing the majority of the load, with email retrieval (gray) being in second. Also note that there is a small amount of FTP traffic (red). Interestingly, there is little-to-no SMTP (server-to-server) email (purple). At several points during the exercise, the NTOP view of protocol distribution clued us into other problems.

As further examples of the usefulness of NTOP data, consider the following graphs with brief comments.

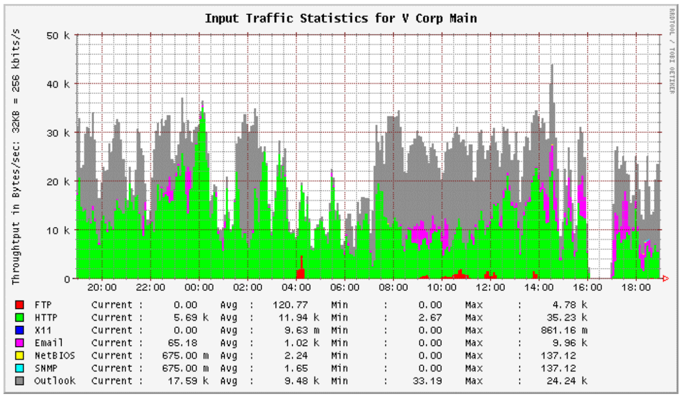

The above graph is a fairly typical snapshot of protocol distribution. Note that from 03:00 to 07:00, the bandwidth utilization is almost exclusively HTTP web browsing. Also note how much Outlook email traffic there is from 07:30 to 15:30. Lastly, the purple SMTP server-to-server email exchanges indicate email being sent to or received from the local exchange server from other servers. The fact that the amount of gray is so much greater than the amount of purple clearly indicates that a majority of users are reaching back to garrison for email.

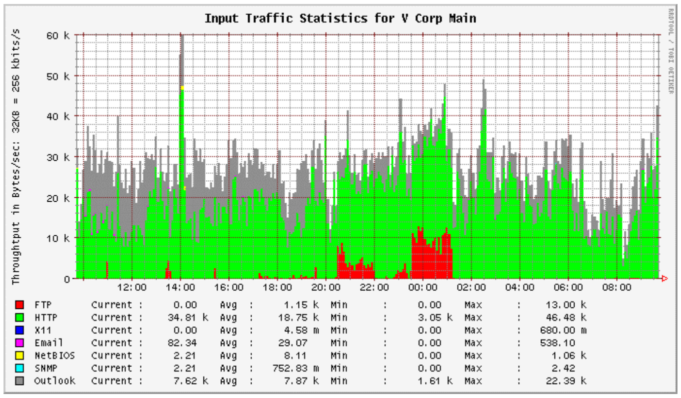

In the above graph, there is a large amount of ftp traffic in the middle of the night. Based on this information, it is easy to bring up other NTOP reports and further investigate the source and destination of the ftp traffic. Note the previous three graphs were generated by a small Perl program using NTOP data; data storage and graphing was done with RRD.

For example the following native NTOP report shows that the majority of the ftp traffic in the graph above is incoming versus outgoing (111.3 Mb versus 18.9 Mb).

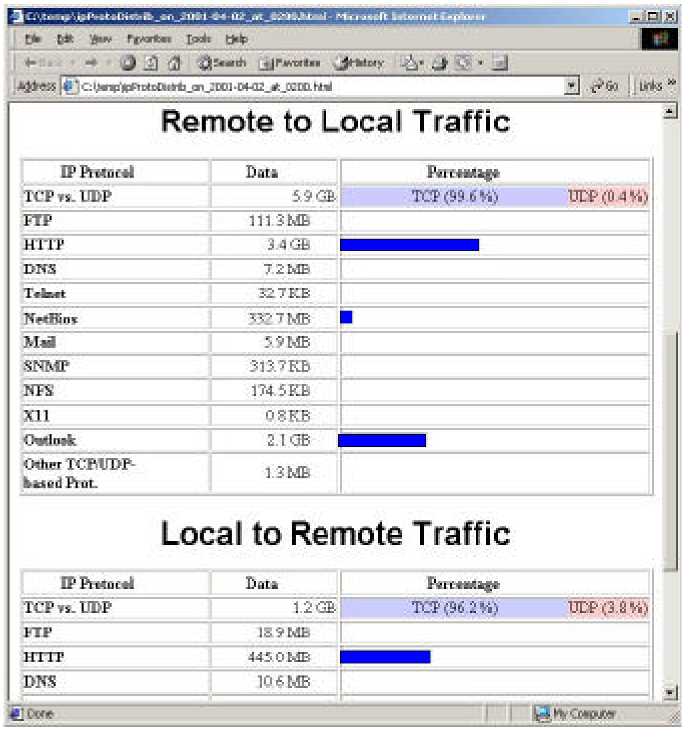
C:\temp\ipProtoDistrib_on_2001-04-02_at_0200.html - Microsoft Internet Explorer
File Edit View Favorites Tools Help
Back Search Favorites History
Address C:\temp\ipProtoDistrib_on_2001-04-02_at_0200.html
Go Links
Done
My Computer

Remote to Local Traffic

IP Protocol	Data	Percentage
TCP vs. UDP	5.9 GB	TCP (99.6 %) UDP (0.4 %)
FTP	111.3 MB	
HTTP	3.4 GB	
DNS	7.2 MB	
Telnet	32.7 KB	
NetBios	332.7 MB	
Mail	5.9 MB	
SNMP	313.7 KB	
NFS	174.5 KB	
X11	0.8 KB	
Outlook	2.1 GB	
Other TCP/UDP-based Prot.	1.3 MB	

Local to Remote Traffic

IP Protocol	Data	Percentage
TCP vs. UDP	1.2 GB	TCP (96.2 %) UDP (3.8 %)
FTP	18.9 MB	
HTTP	445.0 MB	
DNS	10.6 MB	

The following NTOP report gives a minute by minute accounting of the top three traffic sessions. In these charts it can be seen that Class.hq.c5.army.mil hostname is accessing ftpsj.zdnet.net (almost certainly an ftp server based on the hostname). Whether or not this ftp transfer is a legitimate requirement is a different consideration. However, the fact that reports such as these are collected and can be accessed tends to enforce better constraint among users. If it were not for the fact that AR 380-19 personally sensitive data were available in these reports, it might be recommended that they be automatically posted to the G6 TacWeb reports page.

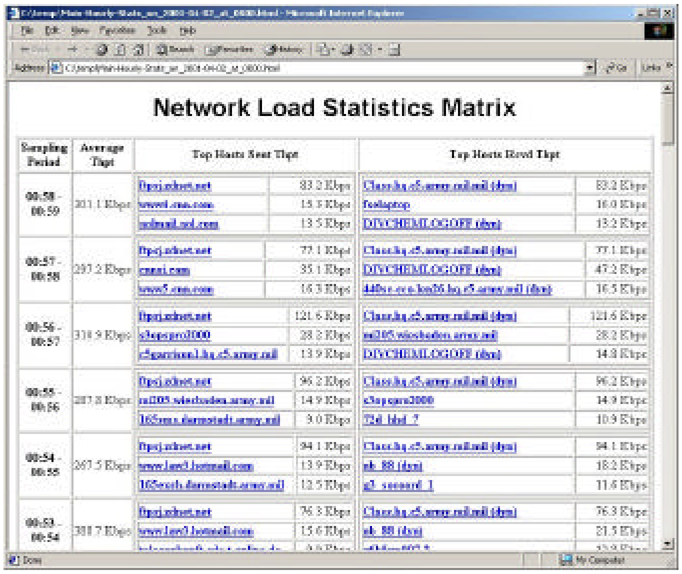
C:\temp\Main-Hourly-Stats_on_2001-04-02_at_0800.html - Microsoft Internet Explorer
File Edit View Favorites Tools Help
Back Search Favorites History
Address C:\temp\Main-Hourly-Stats_on_2001-04-02_at_0800.html
Go Links
Done
My Computer

Network Load Statistics Matrix

Sampling Period	Average Thpt	Top Hosts Sent Thpt		Top Hosts Rcvd Thpt	
00:58 - 00:59	301.1 Kbps	ftpsj.zdnet.net	83.2 Kbps	Class.hq.c5.army.mil.mil (dyn)	83.2 Kbps
		www1.cnn.com	15.5 Kbps	fosilaptop	16.0 Kbps
		webmail.aol.com	13.5 Kbps	DIVCHEMLOGOFF (dyn)	13.2 Kbps
00:57 - 00:58	297.2 Kbps	ftpsj.zdnet.net	77.1 Kbps	Class.hq.c5.army.mil.mil (dyn)	77.1 Kbps
		cnnsi.com	35.1 Kbps	DIVCHEMLOGOFF (dyn)	47.2 Kbps
		www5.cnn.com	16.3 Kbps	440sc-cn-hsn36.hq.c5.army.mil (dyn)	16.5 Kbps
00:56 - 00:57	318.9 Kbps	ftpsj.zdnet.net	121.6 Kbps	Class.hq.c5.army.mil.mil (dyn)	121.6 Kbps
		s3opspro2000	28.2 Kbps	mil205.wiesbaden.army.mil	28.2 Kbps
		c5garrison1.hq.c5.army.mil	13.9 Kbps	DIVCHEMLOGOFF (dyn)	14.8 Kbps
00:55 - 00:56	287.8 Kbps	ftpsj.zdnet.net	96.2 Kbps	Class.hq.c5.army.mil.mil (dyn)	96.2 Kbps
		mil205.wiesbaden.army.mil	14.9 Kbps	s3opspro2000	14.9 Kbps
		165mus.darmstadt.army.mil	9.0 Kbps	72d_hhd_7	10.9 Kbps
00:54 - 00:55	267.5 Kbps	ftpsj.zdnet.net	94.1 Kbps	Class.hq.c5.army.mil.mil (dyn)	94.1 Kbps
		www.law3.hotmail.com	13.9 Kbps	nk_88 (dyn)	18.2 Kbps
		165coh.darmstadt.army.mil	12.5 Kbps	g3_sennord_1	11.6 Kbps
00:53 - 00:54	381.7 Kbps	ftpsj.zdnet.net	76.3 Kbps	Class.hq.c5.army.mil.mil (dyn)	76.3 Kbps
		www.law3.hotmail.com	15.6 Kbps	nk_88 (dyn)	21.5 Kbps
		[illegible]	[illegible]	[illegible]	[illegible]

Restricting the viewability of NTOP data was of prime importance, because AR 380-19 disallows disseminating information about individuals' computer and network usage. However, we felt some information could be used by the whole community, and the point needed to be made that usage was being tracked. In order to accomplish this, network access controls were used to restrict NMA access to a single server running FTS, which automatically loaded the publicly viewable NTOP screens every 30 minutes and posted them to a web site. The end result was that the private data remained hidden, while the users were provided current network situational awareness.

Finally, note that in each of the above NTOP graphs, our bandwidth appears to exceed the maximum 32KBytes per second. This is due to the fact that the NTOP graphs are of the combined incoming and outgoing bandwidth. In addition, due to the fact that there is a web cache in-between the NTOP server and the satellite link, it is quite probable that the incoming bandwidth to the NTOP server is in fact greater than what is incoming over the satellite link.

STATUS

The NMA was successfully tested as a proof of concept during the Corps Warfighter exercise. As previously mentioned, it provided a level of visibility into the network utilization that had never before been seen.

There are plenty of opportunities to extend the functionality of the NMA as will be discussed in the next section. However, the near term immediate focus is to consolidate our progress to date and focus our activities on making the NMA as easy as possible to operate and maintain. This will include an automated CD recovery/installation mechanism, along with a web interface menu-driven configuration process (assign IP address, host name, etc.). This may also include a modem or terminal server access mechanism so that technical support personnel can remotely access a deployed system for troubleshooting.

In the meanwhile, the NMA will undergo another proof of concept demonstration in an exercise later this year.

FUTURE ENHANCEMENTS

It is very easy to become overwhelmed with the amount of data that the NMA can produce. Built-in data archiving functions store data for subsequent analysis. Nevertheless, additional mechanisms are needed for the user to be able to tailor the NMA to watch for specific patterns and trends.

One possible enhancement is to extend the web-interface menu driven configuration to allow the user to specify protocols and hosts to monitor. The NMA is pre-configured with the standard TCP/IP protocols, but there are a number of legacy applications that may be desirable to monitor. It would also be very useful to allow the user to define particular threshold conditions (e.g., web utilization over 50% of overall bandwidth). The NMA could then notify the network manager of the exception condition via email, pager, telephonic, or other means, at which point the network manager could further investigate with some of the more detailed NMA reports.

Some organizations may be adverse to the use of Linux as the base OS. Due to the nature of the products used, and the desire to operate in a true “headless” fashion, a Unix-based solution greatly eases the implementation. We may look at hosting the NMA on Solaris x86 to mitigate these concerns, but have no plans to pursue a Windows NT/2000 solution.

The built-in web-reporting interface is more than adequate for ad-hoc viewing and real-time monitoring, but it would desirable to use more powerful data analysis packages. There is a ODBC plug-in module to NTOP which may satisfy this requirement. It may be relatively straightforward to develop a standardized XML interface. Lastly, it may even be possible to instrument the NMA with an SNMP agent that would provide standard RMON-II data to an SNMP manager. This would be an interesting capability in that it would position the NMA to fit into an enterprise RMON-based management network, without sacrificing visibility by the local nodal manager. In lieu of an enterprise management reporting framework, it may be useful to be able to configure the NMAs to be cognizant of one another, thereby providing a rudimentary network management portal to gather bandwidth utilization information from anywhere within the network.

While extending functionality may tend to be more desirable, there is also a need to further enhance the security-in-depth posture of the NMA.

The current NMA configuration uses a robust combination of host and network level security and access controls. These need to continue to be re-looked as part of the overall network security architecture. For example, it may be desirable to move access control and authentication functions to a separate mechanism such as a RADIUS or LDAP server (perhaps in conjunction with router access controls, PKI certificates, etc.) Lastly, as additional functionality is added and more data is made available, the legal ramifications of how the data is used need to continue to be assessed (AR 380-19).

Server consolidation via virtual PCs (VMWare): The placement of the NMA in between the command post LAN and the tactical WAN is an ideal location for a number of other services. We have routinely operated the applications within a virtual PC environment using a COTS product called VMware. This was done primarily to simplify configuration management, but also presents a unique opportunity to host multiple virtual servers at this LAN/WAN network boundary location. As an example, consider a single 2U size dual processor, dual NIC server operating all off the following applications: NMA, Firewall, Web Cache, Content Filter, IDS. All on one box!

REFERENCES

[1] Deri1, L., Suin, S., Improving Network Security Using Ntop, http://jake.unipi.it/~deri/RAID.pdf

[2] Deri1, L., Suin, S., Effective Traffic Measurement Using Ntop, http://jake.unipi.it/~deri/ntop_IEEE.pdf.gz

[3] Deri1, L., Suin, S., Ntop: beyond Ping and Traceroute, *http://jake.unipi.it/~deri/ntop_DSOM99.pdf.gz*

[4] Lopez, R., Monitoring your Network with Linux, http://www.ntop.org/Monitoring.html