Minimal-Intrusion Traffic Monitoring And Analysis In Mission-Critical Communication Networks

Dr. Alberto Domingo Ajenjo
NATO Consultation, Command and Control Agency
PO Box 174, 2501 CD, The Hague. The Netherlands

and

Dr. Hermann Wietgrefe
NATO Consultation, Command and Control Agency
PO Box 174, 2501 CD, The Hague. The Netherlands

ABSTRACT

A good knowledge of expected and actual traffic patterns is an essential tool for network planning, design and operation in deployed, mission-critical applications. This paper describes those needs, and explains the Traffic Monitoring and Analysis Platform (TMAP) concept, as developed in support of NATO deployed military headquarters Communications and Information Systems. It shows how a TMAP was deployed to a real NATO exercise, to prove the concept and baseline the traffic needs per application, per user community and per time of day. Then, it analyses the obtained results and derives conclusions on how to integrate traffic monitoring and analysis platforms in future deployments.

Keywords: Traffic Capture, Traffic Analysis, Mission-Critical Networks, Protocol Analysis

1. INTRODUCTION

Modern communication networks carry a variety of traffic types, ranging from best-effort e-mail or file transfers to real-time multimedia data flows. In most of the cases, instantaneous bandwidth requirements are not easy to predict, and this becomes more evident as the number of users and services in the network increases. On the other hand, the user perceived communications quality depends mostly on how well the network can accommodate the changes in bandwidth requirement over time.

The above is especially applicable to mission-critical communication systems, such as those which provide connectivity to military, air-traffic control or industrial and utilities networks, to name a few. In those cases, it is essential that the network planning, design and implementation process ensures that users and services will get the required bandwidth, when and where needed.

Unfortunately, very often the actual requirements for such “missions” can not be accurately predicted in advance [1]. In those cases, the usage of data traffic patterns from similar existing systems constitute extremely valuable information for the different stages of communication systems development.

This paper describes the approach to the development of a Traffic Monitoring and Analysis Platform (TMAP) for a data communications network. It shows the results of a capture and measurement campaign [2] conducted over a complex, real-life, mission-critical operational network. Such data communications network was deployed to provide connectivity for a military (NATO) exercise. This scenario, while fully operational, constituted a controlled environment. Such a controlled environment served to baseline the captured traffic statistics with the Exercise concept of operation. The correlation allowed to derive real user data transmission patterns, and also to assess the reasonability of the captured data. Such reasonability is very difficult to verify in other, larger and more complex operational networks, where there exists limited knowledge of the users behaviour with respect to the usage of applications requiring communications.

2. USER REQUIREMENTS

There are a number of basic requirements that a traffic monitoring and analysis platform must meet. In general, it should be able to provide basic performance parameters, including network utilization levels (instantaneous, peak, average or trends), statistics on the collected traffic (protocol distribution, to/from information, packet lengths or errors), network response times, number of error packets and measured throughputs, to name a few.

Then, and as a typical constraint of mission critical networks, security must be guaranteed, and the monitoring platform must adhere to Confidentiality, Integrity and Authenticity (CIA) enforcement policies. It is also frequent that in this type of networks a number of different media (shared or switched Ethernet, FDDI, wireless or WAN trunks) have be simultaneously analysed, and that important changes to the network configuration (to which the monitoring platform must react) are frequent in relatively short periods of time.

Finally, the traffic monitoring tool should be as non-intrusive as possible, with negligible impact over the real traffic and network performances of the monitored network. The system stability, availability and reliability should ideally not be affected at all by the inclusion of the traffic monitoring platform.

3. TRAFFIC CAPTURE AND ANALYSIS PROCESS

The traffic capture and analysis process herein proposed follows a life cycle that combines both proactive and reactive phases. The process starts by base-lining a typical deployment, inserting the monitoring platform into an existing live network. Traffic is captured and later on analysed off-line. This permits to characterize the typical traffic for such a deployment, and generate a baseline (per user, per application, per protocol or per time of day) which will be used during subsequent phases of the life-cycle.

During the early specification of the communications for a future deployment, the acquired baseline data traffic patterns allows to estimate and quantify user and mission requirements, consistent with real operational needs. This turns into easier to implement requirements, at lower cost.

During the design phase, the baseline and the operational needs support decision making on how to avoid bottlenecks and provide communication resources where they are actually needed.

During the network deployment and operation, the theoretical data traffic patterns help as a reference model against which real-time measurements on the system can be compared. These comparisons uncover in many occasions improper use of the communication resources, security threats, or failures in a communications component or in an application. They can also reveal unexpected or unpredicted user needs and cater for them before congestion occurs.

Finally, when planning for technology/components upgrades or replacements, traffic models provide good information for analysis and simulation, both of which are by far more economical ways to proceed than rolling out a real system for evaluation and measurement purposes. This, in turn, permits to derive good assumptions about what the traffic expected from a different network layout would be, and how would bandwidth demand vary in the future, as new services appear and others are abandoned.

4. TYPICAL NETWORK LAYOUT

The typical network layout of a deployed headquarter (HQ) is depicted in figure 1. It normally consists on a number of deployment clusters, interconnected through wide-area-network links.

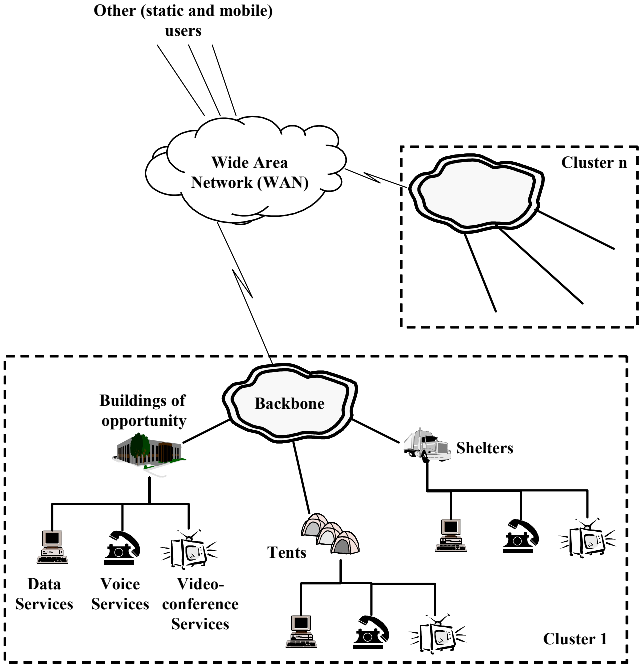

Figure 1. Typical deployment of a military headquarter

Each deployed cluster, which can currently support between a few tens and around a thousand users, provides a variety of voice and data services through a high capacity backbone, based on a fibre optic ring. This ring can be based on a variety of technologies, ranging from FDDI to switched Ethernet, or even proprietary solutions, as needed. As shown in figure 2, the backbone does not provide direct connectivity to final users but, instead, it only interconnects functional modules.

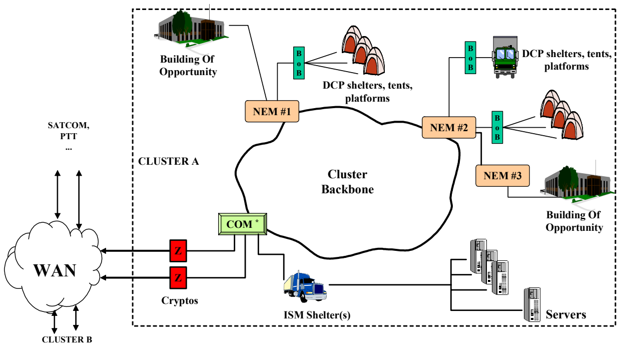

Figure 2. Typical layout of a deployed cluster

The Communications Module (COM) is hosted on a transportable shelter, and groups the main voice and data equipment in the cluster, the transport media (ring) control and management equipment, security equipment and access to the WAN links equipment. It provides the borderline between the WAN and the users in the cluster, and ensures homogeneous access of users to communication services, regardless of the communication infrastructure available in the deployment area (PTT, satellite, radio, etc.). Figure 3 shows the latest COM shelter developed by NC3A for test-bedding purposes.

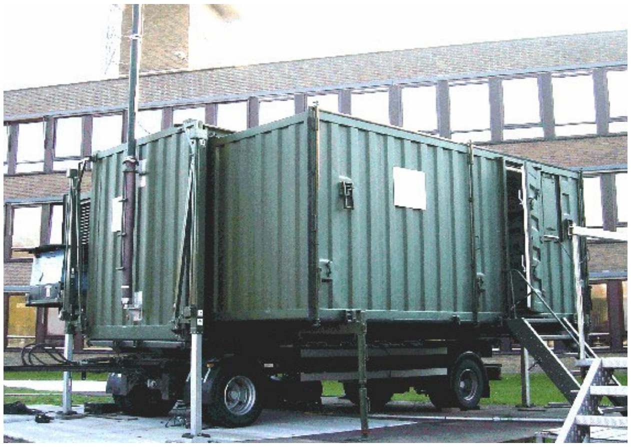

Figure 3. A test-bed COM shelter

Information Systems Modules (ISMs) and Office Modules (OMs) connect directly to the COM shelter or to the cluster backbone. They host, respectively, main IS servers and work positions. The bulk of users, however, are connected to the backbone through Network Extension Modules (NEMs). NEMs extend the backbone to serve large number of users, employing LAN technologies such as switched Fast Ethernet and small PABXs to provide data and voice services in nearby shelters, tents or buildings of opportunity. Figure 4 shows a typical layout of a NEM, which connects to the backbone via a combination of wired and wireless connections. Wireless NEMs can also be used.

The data traffic across the networks serves a variety of user applications, supported by a number of protocols of diverse nature with quite different requirements. Those requirements range from standard Windows messaging to real-time data delivery. This traffic is supported both in the high-capacity backbone and in the local area segments and, bandwidth permitting, can also be sent to the WAN connections.

5. THE TRAFFIC MONITORING AND ANALYSIS PLATFORM

NC3A has developed a (data) network Traffic Monitoring and Analysis Platform (TMAP) in support of NATO deployed command and control systems and, in particular, in support of Combined Joint Task Force Headquarters (CJTF HQ). The TMAP [3] is a tool for continuous traffic monitoring over a variety of physical media, different topologies and diverse protocols and applications. It supports the characterization of traffic in the network, in terms of traffic volume, source and destination, time distribution and protocol/application types, and to help on the definition of a traffic baseline. It eases the detection and characterization of traffic bursts and other non-standard traffic conditions, caused by increase in demand by user/applications, faulty conditions, etc. In particular, the tool allows the detection, identification and isolation of congestion situations. Finally, the TMAP supports network planning, network design and user/application segmentation before deployment, and facilitates network maintenance and reconfiguration during operations.

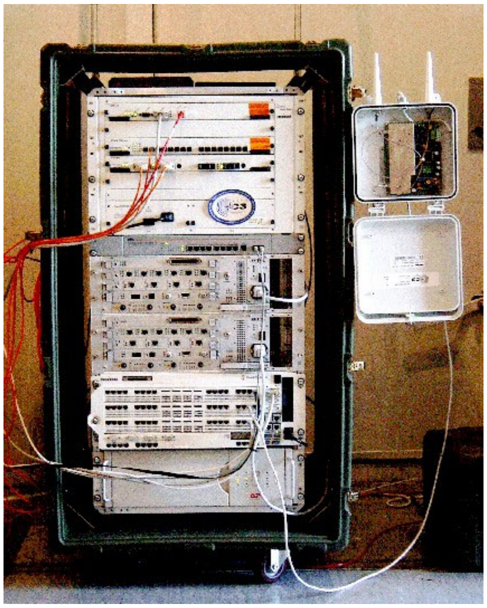

Figure 4. A Network Extension Module (NEM)

The TMAP is based on a Central Monitoring Console (CMC), plus a number of Remote Monitoring Devices (RMDs). Each RMD can be configured in advance to capture, monitor and analyse the traffic on a given segment of the network, and report it back to the CMC. RMDs must, then, be able to deal with a variety of physical interfaces and data protocols.

The CMC, shown in figure 5, is based on a rack-mounted PC with a fault-tolerant FDDI interface to connect to and monitor the cluster backbone ring, plus an Ethernet adapter to exchange monitoring data with the remote probes. It runs a traffic analysis application (Observer, from Network Instruments) that eases the visualization of traffic trends and the generation of traffic reports. The CMC is enabled with basic SNMP and RMON 1/2 capabilities.

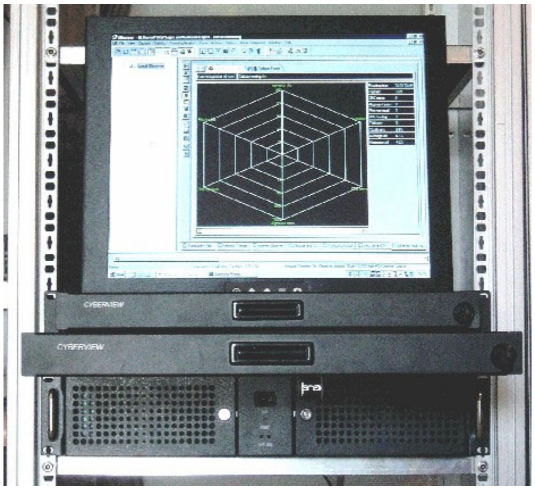

Figure 5. The Central Monitoring Console (CMC)

The RMD's are low cost, PC-based small-format computers that run basic capture applications. They are fitted with the appropriate network interface before being deployed, and non-intrusively capture traffic across the network segment to which they are connected to, and report the statistics to the CMC on a timely basis or upon request (off-line analysis is also possible). Figure 6 shows the typical components fitted into a RMD.

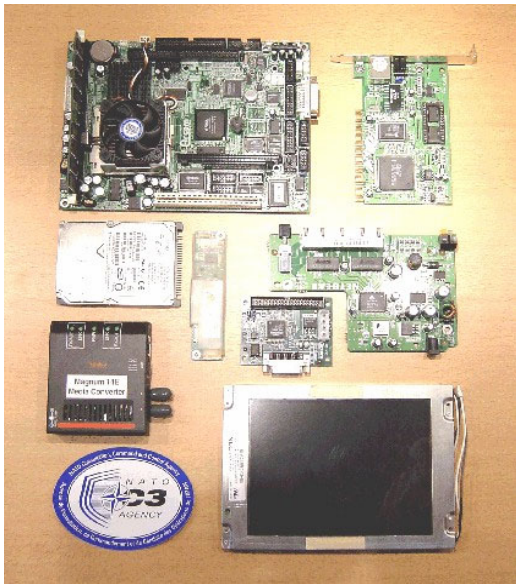

Figure 6. Typical parts in a Remote Monitoring Device (RMD)

Figure 7 shows a RMD mounted on a ruggerdised, small-factor enclosure, to survive hostile environments in an unattended manner. Currently, Ethernet, wireless Ethernet (802.11b), FDDI and Token ring interfaces are available for the RMDs. The port span capabilities of the monitoring application are used to cope with switched Ethernet environments.

The TMAP is at present a distributed passive monitoring tool. The analysis is based on real traffic across the network segments, that the system monitors and analyses. No synthetic traffic is generated by the platform and, therefore, the level of intrusion is limited to the exchange of the captured traffic statistics from the RMDs to the CMC, if such option is selected.

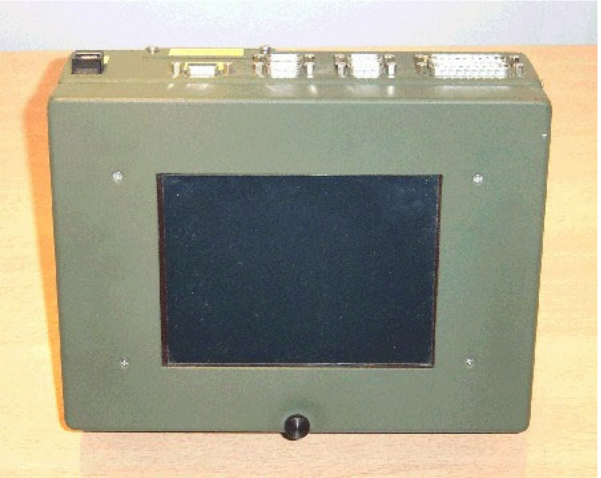

Figure 7. A Remote Monitoring Device (RMD)

On a later release of the TMAP, the implementation of active monitoring to measure applications response times, and automatic anomaly and alarm reporting to PDA devices will be further developed.

6. MEASUREMENT AND ANALYSIS RESULTS

A limited-scope traffic monitoring and analysis platform was deployed to the Allied Effort 2001 (AE'01) NATO Exercise, which was aimed to test the deployed headquarters concept in a Peace Support Operation. The Exercise took place in Wroclaw (Poland), during November 2001.

Figure 8 shows a simplified layout of the packet switched (data) network deployed during the Exercise. Cluster A hosted the majority of users, which were served using an FDDI optical ring, to which main Exercise data servers where attached. Three NEMs were also used to provide connectivity to several hundred users located in buildings of opportunity, shelters and tents. The WAN connections, from the COM module, provided connectivity with the second cluster (hosting the Land and the Maritime Component Commands) and with the static NATO assets, through commercial and satellite, 2x256 Kbps encrypted links.

For this exercise, a RMON 2 probe was attached to the FDDI backbone ring, and a console running traffic monitoring software was attached to one of the LAN segments, located behind a NEM. The main purpose was to test the monitoring concept on a real-life network operation, and to baseline the user traffic. Only one of the typically deployed networks (NATO Secret, Mission Secret and Unclassified), the Mission Secret network, was monitored.

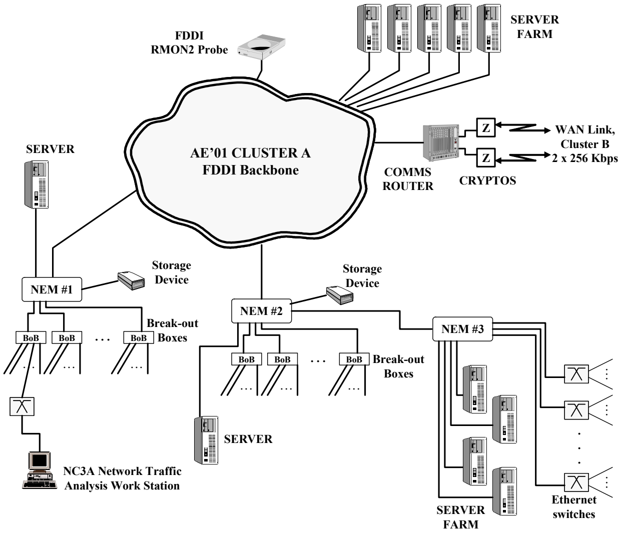

Figure 8. Layout of the data network in AE'01 Exercise (only cluster A is shown)

During the whole duration of the Exercise (7 days, plus the set-up and tear-down phases), the RMON 2 probe was left running, and reporting the captured traffic statistics to the central console every 5 minutes. Mapping of physical to network addresses and names was done through a combination of access to DHCP configuration files and automatic discovery by polling, to minimize the intrusion on the network. Traffic present on the ring was segmented based on user community, on protocol/application type, on packet length and on time of day. Traffic to/from external (to the cell) users was monitored and logged separately. In all cases, statistics with a resolution of one hour were collected. No local probes were attached to LAN segments, except for the one where the console was attached.

The initial traffic captures helped consolidating the pre-operational knowledge of functional communities deployed to the Exercise, and their respective network addresses. Then, a protocol analysis phase was started. A list of captured protocols and protocol/port combinations was obtained, and matched against the expected ones from the list of running user applications. Mismatches, which can highlight a wrong set-up, a not well defined application behaviour or a security hole, were dynamically resolved.

Collected data analysis was done using 3Com's Trascend Traffix Manager software, running over a Sun SparcStation. Figure 9 shows the overall traffic profile in the FDDI backbone ring during the Exercise, segmented per protocol type and time of day. The horizontal axis represents the collection time (around 9 days), while the vertical axis shows the total amount of traffic, in bytes per second, per protocol type (only the most representative protocols are shown).

In the figure, the set-up phase of the exercise can be seen, with an initial traffic capture, followed by an idle period, from the set-up to the actual Exercise start. Then, a 6 days collection period can be observed. During this period, day and night shifts can be differentiated (which would not necessarily be the case on a real operation). Peaks around shift changes, which correspond to assessment reporting and tasking can be identified. Other traffic peaks, corresponding to data backups can also be seen. Then, right before the last day, a massive system backup was performed, followed by a slow-down on the last evening, which can be observed as a traffic reduction, up to the last morning, when final data transfers for the generation of the test runs were performed, prior to disconnection.

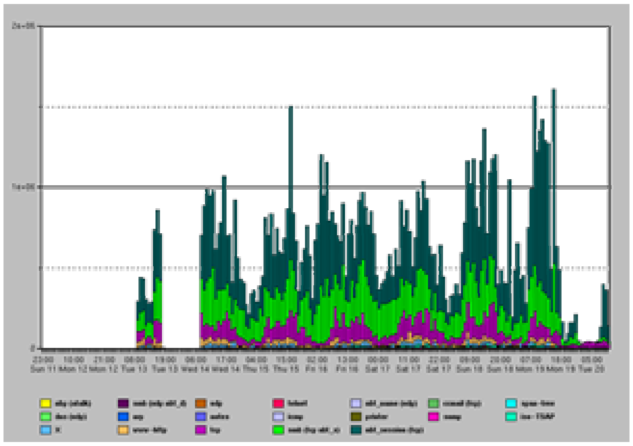

Figure 9. Overall traffic profile, as collected during AE'01 Exercise

Several segmentations based on the community which generated/received the traffic were also done. Figure 10 shows the traffic coming into/out-of the Cluster A from the Air Component Command user community, which was located in Cluster B. Due to the nature of the application used for Air Command and Control, most of the traffic was e-mail and HTTP, with some traces of file sharing. The figure shows also two different shades of grey for ISO-TSAP for the first and second half of the Exercise, as this port number was reclassified in the middle of the capture campaign, for easy of classification purposes.

Other segmentations based on user location, protocol, top-N traffic generators, etc. were also run. They helped to identify especially busy or idle network segments or users. They suggested better server/user arrangements for future deployments, and highlighted several problems related with the Information Systems, which were inappropriately attributed to the communication network. They also showed several other relevant facts, including that the average volume of traffic on the backbone was less than 10% of the backbone capacity, with peaks up to a 70% of the maximum volume. A need to replace the legacy FDDI ring technology was, therefore, not an urgent requirement, traffic-wise. Most of the traffic in the backbone (72%) was originated and destined to users in the cell, with no impact on the WAN links. Traffic to static headquarters was less than 1% of the total. Dominant traffic type was TCP for application sessions, mail Exchange, file and printer sharing. Packet sizes were distributed between 500 and 800 bytes, except for HTTP traffic, where requests and responses averaged 200 and 1,400 bytes per packet, respectively. The traffic generated by the monitoring platform was negligible.

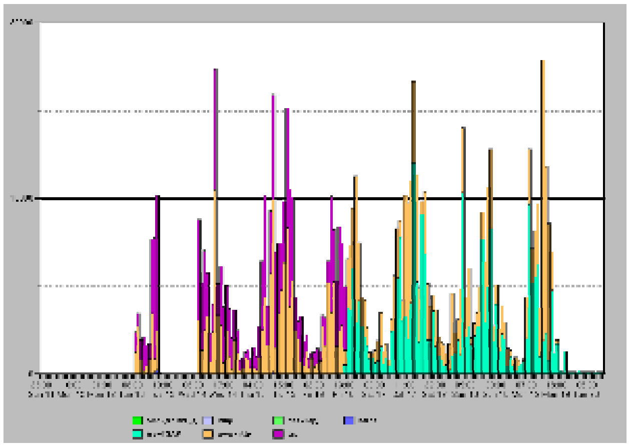

Figure 10. Profile of traffic to/from the Air Component Command user community

7. FUTURE DEVELOPMENTS

The traffic capture and analysis campaign described in this paper satisfied most of the initial expectations, and assessed the TMAP concept and its usability. It also served to identify a number of areas for improvement, which are currently under investigation.

The first area of improvement is based on enhancements to the operational concept. They include better (than five minutes) analysis resolutions, to help detecting and analysing traffic peaks and transients. Also in this area, the possibility to perform some basic processing in the RMDs, to minimize non-essential traffic on the network and cope with periods of segment disconnection is being analysed. How to efficiently monitor the typical three networks deployed in an operation (Unclassified, Mission Secret and NATO Secret) is also of concern.

Another area of particular interest is the monitoring of wireless networks. How to practically monitor this medium, taking into account the implications on security, is being considered.

Finally, other areas of interest include the possibility to upgrade the TMAP with capabilities to generate synthetic traffic and measure server response times when required, or the reporting of alarms and events to personal communication devices (GSM phones or PDAs).

8. CONCLUSIONS

Traffic monitoring and analysis in critical-mission networks is a very valuable tool in the complete operational life-cycle, from network planning to operations. A traffic monitoring and analysis platform can be easily deployed using Commercial-Off-The-Shelf (COTS) products, with a reduced integration effort, and at a reasonable cost. The configuration can be scaled and adapted to the network to analyse, and to the level of required monitoring detail.

While users initially react with reserve to the attachment of monitoring devices, they later on appreciate the benefices of a monitored communications system. Network managers and operators, on the other hand, find it a very useful tool to detect, explain and correct network problems in near-real-time. These tools, on the other hand, should be perceived as a complement to network planning and management applications (such as OpenView or SolarWind), and not as a replacement option.

Finally, the preparation of a good data capture and analysis plan is essential to ensure the success of the monitoring campaign. It should describe in detail what type of monitoring, when and how is going to be done, and it should let the network planners to seamlessly integrate the TMAP into the general network layout.

9. REFERENCES

[1] A. Domingo, C. Dumas, M. van Selm, R. Wik, A. van der Zanden. "*Deployable CIS modules – Cluster simulation model, user traffic model and initial backbone capacity simulation in OPNET*". NATO C3 Agency, Technical Note 863, August 2001.

[2] M. Bommezijn, A. Domingo, H. Wietgrefe. "*Exercise Allied Effort 01 CJTF HQ Backbone Traffic Monitoring and Analysis Results*". NATO C3 Agency, Technical Note 956, July 2002.

[3] A. Domingo, M. Bommezijn. "*Data traffic monitoring and analysis in support of NATO deployable CJTF networks*". NATO C3 Agency, Technical Note 966. In preparation.