

Dynamic Virtual LANs for Adaptive Network Security

Diego Merani, Alessandro Berni, Michel Leonard

NATO Undersea Research Centre
Viale S. Bartolomeo 400
19138 La Spezia
Italy

netcentric@saclantc.nato.int

SUMMARY

The NATO Undersea Research Centre (formerly SACLANTCEN), the research establishment of the Allied Command Transformation (ACT) strongly relies on Network Centric technologies and capabilities to improve the effectiveness of its scientific research. This requires architectures for the interconnection and data sharing that are flexible, scalable, and built on open standards, to ensure transparent interoperability between shore laboratories (both NATO and national) and assets located at sea (research vessels, buoys, autonomous vehicles, sensors and acquisition systems), all connected using a wide range of communications media (e.g. SATCOM, wireless ad-hoc networks, acoustical undersea communications). In addition to that, to fulfil its mission, the Centre has an extensive cooperation program with scientists and researchers, consultants and contractors, civil and military personnel coming, for a limited time period, from several NATO nations. It is a common requirement for them to be temporary connected to the Intranet and to the external Internet: this requirement presents important issues about the security within the internal network; access to networking resources must be controlled while preserving the relative “openness” of a research centre. This paper presents some of the concepts and architectures developed to control the access to network resources and to react to internal attacks.

1.0 INTRODUCTION

The NATO Undersea Research Centre, located in La Spezia, Italy, is dedicated to fulfilling NATO's Operational Requirements in undersea warfare science and technology. Its Scientific Programme of Work, currently organized along three main thrust areas (Antisubmarine Warfare, Mine Countermeasures, and Rapid Environmental Assessment) has resulted during the past 40 years in several scientific and technical contributions that are now part of the set of standard capabilities of all NATO navies.

The execution of the Scientific Programme of Work is performed by an interdisciplinary team that covers different disciplines, such as acoustics, oceanography, ocean engineering, real-time processing, and signal processing. Over the past years, a continuously increasing focus has been put on Network-Centric technologies and capabilities, which have emerged as essential tools to enable and improve the effectiveness of its scientific research

The development of Network-Enabled capabilities in support of undersea research requires architectures for the interconnection and data sharing that are flexible, scalable, and built on open standards. This is essential to ensure transparent interoperability between shore laboratories (both NATO and national) and assets located at sea (research vessels, buoys, autonomous vehicles, sensors and acquisition systems). Also, a wide range of communications media needs to be supported (e.g. SATCOM, wireless at-hoc networks, acoustical undersea communications).

Paper presented at the RTO IST Symposium on “Adaptive Defence in Unclassified Networks”, held in Toulouse, France, 19 - 20 April 2004, and published in RTO-MP-IST-041.

The efforts in the development of Network-Enabled concepts are specifically oriented towards the definition of new generations of scientific instruments. The network will be used to improve the transmission of data from sensors to processors, increasing the capabilities of individual instruments. The resulting increase in efficiency will be larger than the sum of the individual instruments efficiencies.

1.1 The need for adaptive network security

The unclassified network of the Centre is connected to the Internet, and provides the standard services that are requested to a modern enterprise network: office automation, e-mail, Internet access and workgroup file sharing.

The Centre has an extensive cooperation program with scientists and researchers, consultants and contractors, civilian and military personnel visiting the Centre for periods of limited duration. It is a common requirement for them to be connected to the Internet, to use e-mail, download files. They also need to be able to exchange files with Centre staff and print their documents. For these reasons, the Centre has to provide connection to his network or part of it.

The impact of having external people working in collaboration with Centre staff is more evident because the internal infrastructure relies on Windows 2000 Active Directory standardization. Each new computer is automatically installed by Windows 2000 RIS (Remote Installation Server): this brings a high level of standardization in the operating system and software installed. Before connecting a new computer to the network, it is also checked against a checklist, to maintain the highest level of standardization and quality of service.

Visitors that come with their own laptop brings some configuration management issues: IP address and routing need to be reconfigured, and the same has to be done with network cards and printers. Also, when the visitors require participating in a workgroup with Centre staff members, file sharing and mutual authentication need to be tuned.

Having external computers connecting to the network introduces important threats: viruses or Trojan horses can be already active on the host; misconfigurations or previously installed software can unpredictably interact with the network; also, malicious user behaviour cannot be excluded (download of malicious applications from the Internet, unauthorized attempt to access Centre resources, DoS, eavesdropping etc.).

Centre policies for visitors connecting the network are able to mitigate, partially, the risks associated with viruses or malicious code. However, a wide range of threats comes and spreads using the network infrastructure as a vehicle. Also, identifying and remove the origin of an internal attack often involves a time lapse that can be used by the malicious agent (a virus or an attacker) to gain the control or spread across the network.

In this perspective, it is critical to have an infrastructure that is able to self-reconfigure and isolate portions of the network, when they are recognized as the source of an attack: therefore the network has to be deployed so that the impact of an attack can be quickly confined with the minimum impact to the availability of network facilities to the other users.

2.0 REQUIREMENTS AND TECHNICAL APPROACH

Undersea research, regardless of the specialty area being considered, requires the provision of network-enabled capabilities to a wide range of shared systems, ranging from acquisition systems and experimental sonars, to command and control information systems.

Those systems are in most cases experimental prototypes that are built in the Centre laboratories ashore, while the actual testing at sea is conducted onboard the two NATO Research Vessels, NRV Alliance and NRV Leonardo.

The computing facilities are rarely shared in between workgroups, this meaning that network utilization is well defined and confined inside subnetworks; on the other hand, a high rate of computer mobility leads to additional issues, like the automatic reconfiguration of a computer when moved from one network to the other (i.e. from the Centre to the Ship).

Scientific departments often use “non standard” computers. Those systems are not centrally managed and not automatically reconfigured when moved. This adds a level of complexity, because misconfigurations or human errors coming from those systems could interfere with the MIS or other computing facilities, thus reducing the QoS perceived by other users.

2.1 Technical approach

The operational environment of the NATO Undersea Research Centre is characterized by a centralized network management, which covers all locations where work is being conducted. The network is almost fully switched, being the access layer switches cascaded with a fiber optic link (trunk). Remote locations (the research vessels) are connected with the main network through a VSAT satellite link.

The network has been divided into logical subnets, corresponding to various workgroups. Each subnet is assigned a reserved address space. Usually, the address space assigned to a workgroup includes four class C networks, thus allowing further subnetting for special project-related cells inside a workgroup¹.

This architectural approach allows a high Quality of Service, limiting broadcast traffic, and confining undesired protocols within the boundaries of the subnet. In addition, the use of subnets increases security, allowing the discrimination of traffic at gateways, or limitations of access to resources or servers on a need-to-know basis.

The use of a plain subnetting scheme does not necessarily match with the physical positioning of workplaces in our building (this is also a problem for a large number of enterprises). Departments are often spread throughout a building, thus creating physically non-contiguous subnets. In addition, users and computers are often moved when reassigned from one project to another. This situation induces an overhead of re-configuration in a standard network.

The use of VLANs (Virtual Local Area Networks) provides a solution to problems associated with mobility and physical positioning of computers, by dynamically assigning the hosts to their subnet regardless of their physical position.

2.2 Static VLAN vs Dynamic VLAN

Two different approaches can be used when assigning computers to VLANs: static and dynamic. Normally, both methods are used, following the implementation schema suggested by Cisco Systems.

Static VLAN assignment means that each port of the switch is (statically) associated to a certain VLAN: a computer participates in the VLAN according to the port to which is connected. This approach is useful when the network is relatively stable and computers are rarely moved from one socket to another: once the switches are properly configured, the configuration overhead is low. In such an environment, security can be further improved by securing the MAC addresses on switch ports, thus limiting the number of MAC addresses allowed on a specific switch port. In addition, to get further selectivity and security, it is possible to manually assign MAC addresses to a switch port.

¹ See “Architectures for Network Centric Operations in Undersea Research [3]”

The configuration of static VLANs requires a perfect knowledge of network physical topology, since the network administrator needs to specify the VLAN number for every port of the access switches. On existing networks, this could create a relevant initial configuration overhead. In any case, switch ports must be reconfigured whenever a computer changes position or VLAN.

The dynamic VLAN approach on the other hand is more complex, but it does not require any reconfiguration when computers change their physical position in the network (for example, when they are moved from one office to another office). The VLAN assignment is made on the basis of physical (MAC) address of the connecting computer.

A database is maintained centrally on the VLAN Management Policy Server (VMPS); this database includes MAC addresses, and their association to the corresponding VLAN. When a computer is moved to a switch port configured as “dynamic port”, the switch observes the MAC address of the connecting computer. A query is then sent to the VMPS server, using the queried MAC address as keyword, obtaining the VLAN assignment in return. The port is then inserted in the VLAN specified in the configuration file. All these operations are performed in few seconds, in a manner transparent to the user.

2.3 VLAN Management Policy Server²

Using the VMPS, it's possible to assign switch ports to VLAN dynamically, on the basis of the source MAC address of the device connected to the port. When a computer is moved from one switch in the network to another, it receives the new port to the proper VLAN for that host dynamically. The VMPS opens a User Datagram Protocol (UDP) socket to communicate and listen to client requests. When the VMPS server receives a valid request from a client, it searches its database for a MAC address-to-VLAN mapping, as illustrated in the following figure.

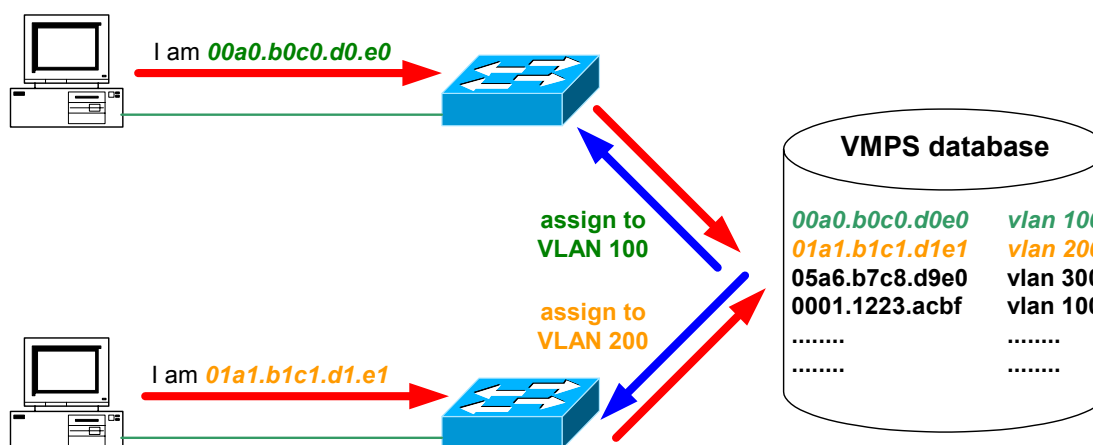


Figure 1 – VMPS initialization dialogue

If the assigned VLAN is restricted to a group of ports, VMPS verifies the requesting port against this group. If the VLAN is allowed on the port, the VLAN name is returned to the client. If the VLAN is not allowed on the port and VMPS is in open mode, the host receives an "access denied" response. If VMPS is in *secure mode*, the port is shut down.

² Cisco Systems document no. 78-14904-01 [4]

If a VLAN in the database does not match the current VLAN on the port and active hosts are on the port, VMPS sends an “access denied” or a “port shutdown” response based on the VMPS secure mode.

It is possible to configure a fallback VLAN name. When a device whose MAC address is not listed in the database, VMPS inserts the client in the fallback VLAN. If a fallback VLAN is not configured, and the MAC address is not listed in the database, VMPS sends an “*access denied*” or a “*port shutdown*” depending whether the VMPS is in open mode or in secure mode. If VMPS is in secure mode, it sends a response.

It is possible to make an explicit entry in the configuration table to deny access to specific MAC addresses for security reasons: this is done by specifying the --NONE-- keyword as the VLAN name. In this case, VMPS sends an “access denied” or “port shutdown” response.

2.4 VMPS and configuration management

The Centre’s network, as previously specified, is based on Windows 2000 Active Directory. Windows 2000 Active Directory DHCP server, configured in “static mode”, maintains the IP addressing. Static DHCP rely on the MAC address to assign the IP address. With this assumption, each computer on the network needs to have a reservation inside the DHCP server in order to get the IP address; moreover, the IP address assigned to a computer will be always the same. Administering the DHCP server requires manual insertion of the MAC address and manual assignment of a free IP address.

To summarize, these three different databases, DHCP, DNS and VLAN, have to be maintained to describe the configuration of the network. The data intersection scheme is given in table 1. Another database, the “Administrative Property Database”, has to be maintained to track administrative properties of the asset, since none of the operational databases is detailed enough for this purpose.

	RECORDS INSIDE EACH TABLE					
DHCP table	IP	MAC		COMPUTER NAME	USER/ CONTACT	
DNS table	IP			COMPUTER NAME		
VLAN table		MAC	VLAN NAME			
ADMIN.VE PROPERTY table		MAC			USER/ CONTACT	ASSET No.

Table 1

The VMPS configuration file is a text file based on a simple and effective syntax. It contains a brief set of common commands specifying the VMPS domain, ports groups and the policy to be used for unknown MAC addresses (shutdown or fallback VLAN).

The main section of the configuration file is a list of MAC addresses and VLAN names. This simplicity is also a weakness if seen from a network management perspective: inserting, modifying or deleting an item is not a user-friendly task, because the operations have to be performed after manual inspection of the MAC address.

To consolidate the configuration management issues, a new application has been written to generate the VMPS configuration file as a function of the other databases. The application, written in Java, stores all data associated with each computer owned by the Centre and offers a user-friendly graphical interface.

A planned extension of the software will also allow the definition of DHCP and DNS parameters, thus it will be possible to configure the participation of a new computer to the network, and also modifications or deletions, from a single management console.

3.0 ATTACK AND RESPONSE

3.1 Network Intrusion Detection

Network intrusion detection is assured by Cisco Secure Intrusion Detection Appliances. Each appliance “sniffs” the traffic looking for malicious or unauthorized traffic. Updated signatures are loaded on the IDS on a weekly basis, to compare the observed network traffic against those signatures. The system is able to detect atomic attacks (those identified by the single packet) and also composite attacks (those that requires the de-encapsulation of data flow).

The real challenge associated with a network IDS in a VLAN switched network is the actual positioning of the appliance. The success of IDS detection relies on the ability to snoop the traffic through the sensor network card: therefore the effectiveness is directly related to the traffic flow that can be physically observed.

A simple corporate network can be divided into at least three zones: the main access layer, the server farm, and the WAN/Internet zone. Normally the three zones are implemented using separate subnets and routing is required between them.

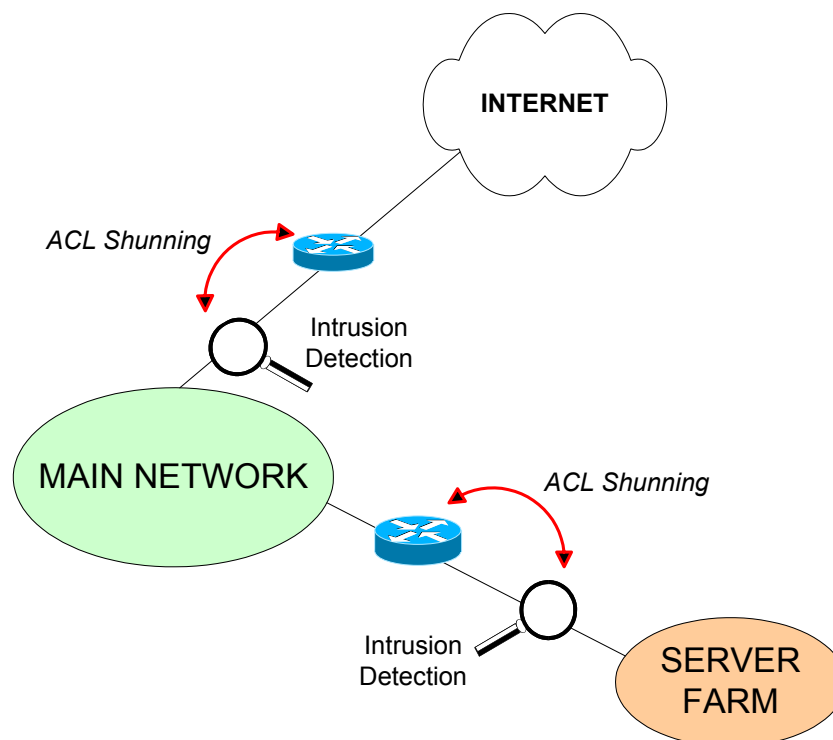


Figure 2 – IDS deployment

Good practice is to position the IDS sensor just above the links between two different zones: in the case depicted in picture 2 this means one IDS behind the firewall and another behind the server farm router. The sensors deployed following this pattern can detect attacks directed to the server farm or to the Internet.

Not only a sensor is able to detect malicious traffic, but also to send alert messages and take corrective actions. The IDS can shun the traffic using routers' ACL. Giving the proper coordinates to the IDS system, he can shun the traffic coming from IP address of an attacker, also blocking the subnet to which the attacker belong.

The shunning can stop at the perimeter router attacks coming from the Internet. Therefore, relying on the IP address, the configuration has to be fine tuned to avoid self Denial of Service: the source address could be spoofed and the shunning action could block a legitimate user or network. For this reason, a literature exists that it's against the use of shunning on Internet routers.

Shunning, instead, can be very useful if used in the intranet. Since attacks need to bypass defined boundaries to be effective, a segmented network can be modified in real-time or near to real time to isolate and confine malicious activities or code.

Far from discussing all the possible targets and methods to conduct a network attack, it is enough to consider that an insider attacker (malicious user or code) will, with all likelihood, target the Internet or the server farm. The Internet is the preferred target of mass-mailer worms, for example; but also malicious software may also perform other Internet activities, such as connections to pre-defined remote sites. Nevertheless, a user who wants to launch an attack to a corporate network will target the server farm, and possibly will use the Internet to download tools or send leaked information.

With these assumptions, the possibility of isolating attackers both from the server farm and from the Internet when an attack is detected would be a powerful countermeasure.

3.2 Access Control Lists and Shunning

It's common to consider a Local Network as an open space with no traffic shaping or authentication gates. This is however a violation of the good practice of *need-to-know* policy access. In our case, for example, there is no need for scientific workgroups to have permanent access to administrative computers; on the same basis, there is no need for having a permanent flow of traffic between groups, unless they are involved in a joint research or project.

This profile can be well translated into access-lists between subnets. File sharing is the main threat today, as viruses "jump" across network shares; indeed the NETBIOS protocol, used for Windows file sharing, can be blocked at the gateway between those groups that do not have the requirement to use it.

Traffic flows from the main network to the server farm have also to be profiled, since not all users need to connect to all servers.

A structure in which all the traffic is filtered and parsed by special-purpose systems, brings a higher level of security. Anyway, the evaluation of risk assessment and the need-to-know policy can unpredictably change when an attacker enters the network. Would this be a malicious user or a viral code, the need-to-know associated to the host of threat should decrease to zero.

Response to attack is a process that can be divided into logical steps:

Dynamic Virtual LANs for Adaptive Network Security

- evaluation and correlation
- automated action
- incident analysis
- managed action
- resolution

During the evaluation and correlation of an IDS event, the IDS engine decides the alarm level of the detected event. The level of a detected alarm is, usually, automatically classified on the basis of the signature file that constitute the IDS knowledge base. Anyway, the alarm level classification can be tuned to fit the needs of the environment to which it is applied. When the alarm is critical (obviously, the thresholds of the classification can be configured), an action can be triggered. Cisco IDS rely on two reactions: TCP resets and ACL shunning.

TCP reset acts on the TCP session that originated the alarm. The TCP reset sent by the IDS appliance is able to drop the session in which the malicious activity is detected.

IDS shunning blocks the host or network originating the attack for a configurable amount of time. Shunning is performed through “on-fly” reconfiguration of a router or firewall, which has to be identified during the initial setup of the IDS. The blocking device gets modified from the IDS, which inject ACL statements to shun the attack.

This powerful feature can therefore create self Denial Of Service attacks, if not properly configured: in particular, attacks coming from the external world can have a spoofed address as a source address, thus causing a legitimate IP address to be blocked; nevertheless, the attacker IP address could be used by a remote firewall to make Network Address Translation: blocking the address could result in preventing all the network behind the NATted address to reach the site protected by the shunner.

IDS shunning can better be used internally to isolate the subnets in which a suspicious activity is detected. It acts on the gateway that protect the server farm to limit the access to servers from the VLAN in which a suspicious signature is detected. This action can be performed real time and in an automated fashion.

The incident analysis is a survey to determine the effects of an attack. For example, if the attack is the result of a virus infection, the survey can estimate the entity of the virus infection; further actions can be issued, if needed, to isolate the attack sources by VLAN ACL (that can block traffic within the VLAN) or banning the MAC addresses of the attackers from the network.

This last stage of incident handling is required since the current IDS versions cannot shun connections at layer 2, leaving the possibility of connections within the subnet. In addition to that, the escalation of counter-attack actions should not be conducted automatically beyond a certain threshold to avoid the risk that false positives create a chain reaction and bring self-DoS.

4.0 CONCLUSIONS

The solution proposed in this paper applies to small-medium business enterprises, where the use of a Catalyst series 5000 or above (that supports IDS onboard, and is able to snoop inside VLANs) is not a cost-effective choice.

The segmentation of the network, combined with intrusion detection performed on the local traffic, gives the opportunity of implementing a simple, heuristic mechanism of automatic reaction to anomalies or attacks, nevertheless maintaining a level of service for the portions of the network that are not interested or involved in the attack.

The improvement to the overall security posture can be implemented adapting a network architecture based centrally managed virtual LANs. VLANs provide not only better quality of service to different workgroups, but also allow the quick isolation of network users and segments, which appear to be subject to an attack, thus limiting the impact to the whole of the network.

The examples given are based on the Cisco Systems Safe framework, but the concept can be applied to other environments when the necessary functions are supported in the network hardware and software.

5.0 REFERENCES

- [1] Leonard, M., Berni, A., Merani, D., “Architectures for Network Centric Operations in Undersea Research”, presented at the RTO SCI-137 Symposium on “Architecture for Network-Centric Operations”, held in Athens, Greece, 20 to 22 October 2003
- [2] Leonard, M., Berni, A., Merani, D., “INSC applications for undersea research”, presented at the Interoperable Networks for Secure Communications (INSC) symposium, held at NC3A, The Hague, The Netherlands, 4 to 6 November 2003
- [3] Berni, A., Leonard, M., “Network Architecture for real-time contact sharing in multistatic sonar operations”, SM-367 (NATO RESTRICTED), NATO Undersea Research Centre, La Spezia, 2003
- [4] Cisco Systems, “Configuring dynamic VLAN membership with VMPS”, Software Configuration Guide for Catalyst 4000 family, document no. 78-14904-01

