Top-down Integrated Security Architecture for Heterogeneous IP Networks

G. Balis[1] E. Bouras[1] G. Diakonikolaou[1] G. Samara[1] D. Kaglis[2]

[1] Multimedia Laboratories, OTE Research
Hellenic Telecommunications Organization S.A.
15122 Marousi, Athens, Greece
Email: {gbalis,ebouras,gdiako,gsam}@oteresearch.gr
[2]Telecommunications Lab, Department of Electrical & Computer Engineering
National Technical University of Athens
15773 Zografou Athens, Greece
Email: kaglis@telecom.ece.ntua.gr

Abstract: The need of data and transaction protection has made access control and security a necessity. This has led to extensive research in the academic community that resulted in the specification of various protocols used to implement it. These protocols differ greatly and are rarely compatible with one another or towards the security policy they are designed to implement. There also exist proprietary protocols usually compatible only with a certain number of products, most of which are not widely recognized and adopted. Thus appeared the need for a comprehensive and compatible-oriented security policy design. This policy should support features as: Authentication, Authorization and Accounting, DMZ architecture, Intrusion detection, Service Level Agreement mechanisms and Disaster recovery of critical data. We shall take a case-study approach by examining the experimental network of the OTE Multimedia Labs designed in the course of several research programs conducted in the European community's research framework programs.

1 Introduction

A sound and effective security policy is of great importance for corporations today given the increased popularity and business opportunities of public wired and wireless Networks. Corporations today engage more often than before in on-line transactions and e-business solutions as a means of maximizing their profit as well as getting a hold of a new and promising audience. The Internet as a new market seemed very promising but soon it was realised that care should be taken so that on-line transactions could become as much safe as the old legacy ones. There was also the issue of data confidentiality and protection as well as identity proof. These needs led a vast number of researchers from the academic community as well as the industry to propose possible solutions in the form of new protocols, policies and network design architectures that could be deployed in order to achieve the above goals. Thus in the previous decade there was a flood of new protocols, access control policies and new network architectures proposed to address the need for security in public networks.

They were mostly the product of individual works with little or no central guidance or setting of requirements and this is the prime reason for the major incompatibilities amongst all these new promising products. Later there were some efforts to accommodate many of these new products into cohesive and concrete framework that would be widely adopted and implemented [5]. There was little if any success in these efforts and this brings us today where the situation has become a total labyrinth consisting of many incompatible protocols and policies. At the Multimedia Labs of OTE Telecommunication Company through the course of many research activities, both European as well as self-funded a rather large network was produced to be used as a test-bed for research purposes and measurement of various network variables. Various security protocols and policies were implemented and tests of interoperability amongst them were performed. The outcome of this research has lead to a network design that incorporated some well-known protocols as well some new emerging ones. The main effort was to produce a system that would be manageable and its components would interoperate to provide the desired results. The main components of this system are distinguished into three main categories: Access and Authorization Control, Packet filtering and Firewalls and Network Management Services (Central management and Disaster recovery).

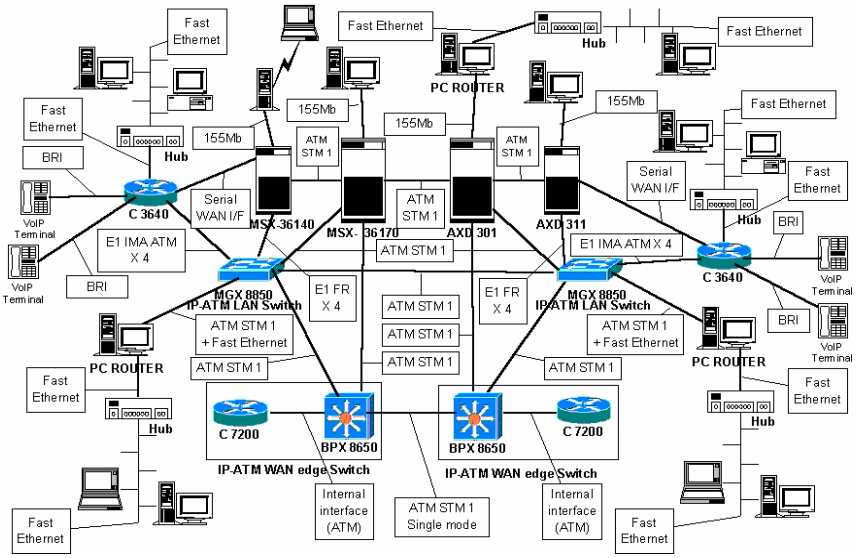

Figure1. OTE Multimedia Labs Network

2 Access and Authorization Control

Access control is the process of controlling who is allowed access to a network and what services they are allowed to use once they have access, as well as keeping a record of what actions each user performed [13].

Authentication, Authorization, and Accounting (AAA) is the architectural framework for configuring these three independent functions in a consistent manner [14].

Authentication: Provides the method of identifying a user, including login and password dialog, challenge and response. Identification of a user takes place prior to being allowed access to the network and its resources.

Authorization: Provides the method for remote access control, including one-time authorization or authorization for each service, per-user account list and profile and user group support. Authorization works by assembling a set of attributes that describe what the user is authorized to perform. These attributes are compared to the information contained in a database for the given user and the result is returned to the AAA server to determine the user's actual capabilities and restrictions. The database can be located locally on the access server or router or it can be hosted remotely on a RADIUS or TACACS+ security server. Remote security servers, such as RADIUS and TACACS+, authorize users for specific rights by associating attribute-value (AV) pairs, which define those rights, with the appropriate user.

Accounting: Provides the method for collecting and sending security server information used for billing, auditing, and reporting, such as user identities, start and stop times, executed commands (such as PPP), number of packets, and number of bytes. Accounting gives the ability to track the services users are accessing as well as the amount of network resources they are consuming. When AAA accounting is activated, the network access server reports user activity to the TACACS+ or RADIUS security server in the form of accounting records. Each accounting record is comprised of accounting AV pairs and is stored on the access control server. This data can then be analyzed for network management, client billing, and/or auditing [11].

AAA provides many benefits, including:

- Increased flexibility and control
- Scalability and fast deployment
- Standardized authentication protocols, such as RADIUS, TACACS+ [10], and Kerberos, NIS+
- Central accounts management and monitoring

NIS+: NIS+ is a hierarchical and secure network information service system for UNIX Operating Systems. In concept, NIS+ is basically a "simple" distributed database which allows information managers and system administrators to manage the network information for complex and heterogeneous computer systems. NIS+ replaces NIS as the default name service for Solaris.
NIS+ servers are able to "speak" the NIS protocol and can serve NIS client requests. Backward compatibility is provided as a mechanism to ease transition.
Advantages of the NIS+ over other services like NIS or /etc files are:

- SPEED: NIS+ propagates incremental updates in a much shorter period of time than NIS (which perform complete map updates).

- AVAILABILITY: NIS+ clients can soft bind to any host from set of NIS+ servers rather than hard binding to a particular server.
- SCALABILITY: NIS+ supports multiple, hierarchical domains rather than a single flat domain. Also, NIS+ tables are multi-columned, in contrast to NIS maps.
- SECURITY: Neither NIS nor /etc/ files provide security. NIS+ provides client authentication and access control.

The choice of the protocol to be used depends on the load expected on authentication servers as well as the security level requirements and costs involved n deployment.
In our labs we chose Cisco TACACS+ as the primary AAA protocol in favour of RADIUS. The choice was mainly based on some basic facts [12]: RADIUS encrypts only the header of the authentication packet while TACACS+ can encrypt the whole packet body. TACACS+ also uses TCP as its transport while RADIUS uses UDP which is faster but not so reliable. Dual-redundant RADIUS authentication servers provide the fallback in case the also Dual-redundant TACACS+ servers fail and also provide authentication for users' dial-up access since it has become the de-facto in ISPs worldwide [9]. The Authentication servers in turn authenticate users centrally through a NIS+ server. The NIS+ server provides centralized management of accounts and AAA servers based on the information contained on the NIS+ maps (database) provide users' access rights, detailed logging of actions performed and the ability to respond quickly in an intrusion attempt on such a large network. The following diagram shows the basic architecture selected for the Access Control part of the design in our labs.

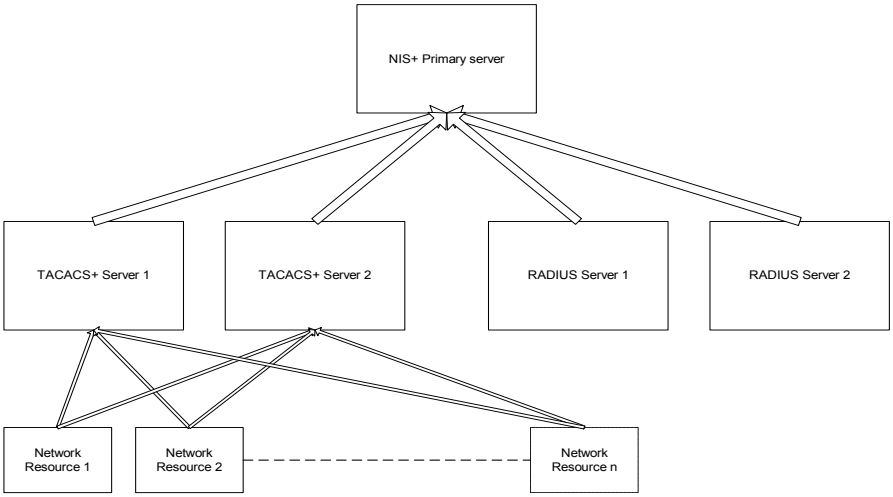

Figure2. AAA Process in a Top-down design

Access Control was designed to in a round-robin fashion, meaning that the authentication servers are used in turn in order to reduce load on the servers and maximize AAA services availability at any given time. To accomplish this round-robin architecture, new features that are present in latest Cisco IOS software were used and specifically AAA broadcasting that supports the use of multiple AAA servers in a somewhat distributed way.

3 Packet Filtering and Firewalls

Packet filtering is probably the most common element in a network security design. As such a basic part of the design extra care must be taken to avoid common errors and pitfalls such as: source address verification, anti-spoofing techniques, denial of broadcast type ICMP requests (also known as smurf Denial of Service DoS attacks). There are several proposed architectures for implementing a network protected by packet filtering devices. The one that balances strong security, availability and ease of use was deemed to be the Screened Subnet Architecture (or DMZ for Demilitarized Zone).

Screened Subnet (or DMZ) Architecture
This architecture is an extension of the screened host architecture which is basically the classical firewall setup of a packet filter between the outside hostile Internet and the corporate LAN. DMZ modifies this design by segregating the available Network into three distinct areas: a "semi-secure" or De-Militarized Zone (DMZ) subnet where the proxies lay, the protected corporate LAN and the public Internet [7]. This allows only the outside restricted access services in the DMZ Zone. The DMZ is further separated from the internal network by another packet filtering device that only allows connections to/from the proxies. The benefits of the specific architecture are:

- The filters are "intelligent" with logging capabilities.
- All incoming and outgoing services between the Internet and the internal networks pass via proxy servers in the DMZ.
- Services offered to the Internet (such as WWW or ftp) run on a dedicated machine which has no access to the inside (since this machine must be considered as potentially compromised - a "sacrificial lamb"). A server providing WWW or (writeable) ftp to the Internet is difficult to fully secure. If a site offers secure services such as Web commerce, it is advisable to install a second, specialized web server with strong SSL, a B1 (TCSEC) approved OS (like Argus Pit-Bull) and highly restrictive usage/monitoring/auditing/change management etc.
- The DMZ can be a switched LAN, or two switched LANs with dual homed bastion hosts between them. The latter is more secure since only proxied connections will be allowed through and prevents a software error in the filters. Direct inside <-> outside socket connections are denied and traffic becomes uni-directional between the LAN and DMZ areas.
- The specific architecture is modular & flexible.
- For maximum diversity of defence, two different firewall techniques should be used for the "packet filters".
- For very high availability, the DMZ with front & back end filters can be duplicated and hooked together by routers (2 on the inside and 2 on the outside) that support redundant routing.

The above architecture was implemented at the Multimedia laboratories at OTE as a means to enhance security. There are two dedicated software routers running FreeBSD that handle all the traffic flow between the public Network, the protected LAN and the DMZ area. The router itself does not provide any other application-level network services in order to enhance it’s security. Each router routes outbound traffic

through a different Internet connection and there is a third router that routes traffic based on availability and network load. Available routes with the packet filtering routers are exchanged using BGP4+. Note also that direct traffic between DMZ and private LAN is not permitted since this would compromise the design. In this manner load balancing and fail-over are accomplished.

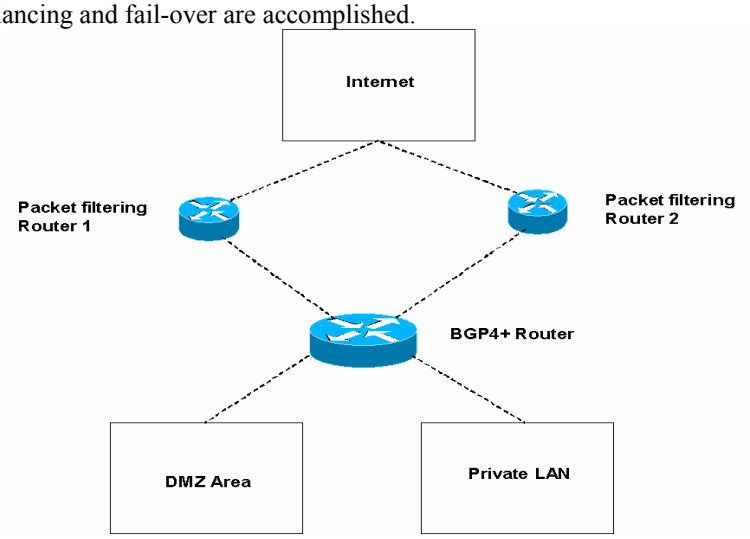

Figure3. Multimedia Labs Dual-redundant DMZ architecture design

NAT

NAT was mainly adopted as a means of preserving the fast depleting IPv4 addresses [3]. It is also used to enhance a sites security in various ways:

- It can divide a large network into several smaller ones. The smaller parts expose only one IP address to the outside. With inbound mapping, it is even possible to move services (such as Web servers) to a different computer without having to do any changes on external clients.

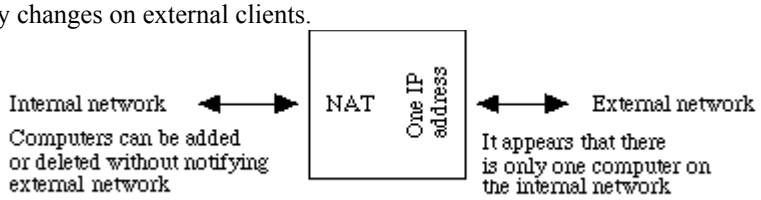

- NAT gateways usually provide a way to restrict access to the Internet.
- Another useful feature is traffic logging; since all the traffic to and from the Internet has to pass through a NAT gateway, it can record all the traffic to a log file. This file can be used to generate various traffic reports, such as traffic break-down caused by user, by a site, by network connection etc.
- Since NAT gateways operate on IP packet-level, most of them have built-in internetwork routing capability. The internetwork they are serving can be divided into several separate sub networks (either using different backbones or sharing the same backbone) which further simplify network administration and allow more computers to be connected to the network.

To summarize, a NAT gateway can provide the following benefits:

- Firewall protection for the internal network; only servers specifically designated with "inbound mapping" will be accessible from the Internet
- Protocol-level protection
- Automatic client computer configuration control
- Packet level filtering and routing

NAT Operation

The basic purpose of NAT is to multiplex traffic from the internal network and present it to the Internet as if it was coming from a single computer having only one IP address.

The TCP/IP protocols include a multiplexing facility so that any computer can maintain multiple simultaneous connections with a remote computer. It is this multiplexing facility the key to a single address NAT. To multiplex several connections to a single destination, client computers label all packets with unique "port numbers" [2].

A modern NAT gateway must change the Source address on every outgoing packet to be its single public address. Therefore also renumbers the Source Ports to be unique, so that it can keep track of each client connection. The NAT gateway uses a port mapping table to remember how it renumbered the ports for each client's outgoing packets. The port mapping table relates the client's real local IP address and source port plus its translated source port number to a destination address and port. The NAT gateway can therefore reverse the process for returning packets and route them back to the correct clients.

This process is completely dynamic. When a packet is received from an internal client, NAT looks for the matching source address and port in the port mapping table. If the entry is not found, a new one is created, and a new mapping port is allocated to the client. So the procedure goes as below:

- Incoming packet received on non-NAT port
- Look for source address, port in the mapping table
- If found, replace source port with previously allocated mapping port
- If not found, allocate a new mapping port
- Replace source address with NAT address, source port with mapping port

Packets received on the NAT port undergo a reverse translation process:

- Incoming packet received on NAT port
- Look up destination port number in port mapping table
- If found, replace destination address and port with entries from the mapping table
- If not found, the packet is not for the internal network and should be rejected

In our labs we employed NAT as e both a means of enhancing security and researching various services potential of working over NAT translators. Such sensitive services include but are not limited to IPSec Virtual Private Networks [4] and some real time applications [1].

Transparent proxies

A proxy is any device that acts on behalf of another. Application proxies (or application gateways) normally offer logging and access control at the application layer, but

at the cost of performance (all traffic must pass via the proxy) and complexity (the proxy needs to run on a special host). The term is most often used to denote Web proxying. Proxy technology is often seen as an alternative way to provide shared access to a single Internet connection.

It may also be necessary to modify the client program. Ideally, proxies and filtering should be used together to maximize security. Proxies are used also:

- For all services that must pass between the inside & outside networks.
- To configure such that access to the proxy from the outside is forbidden, except for where strong authentication mechanisms are in place, or for email & news.
- To use IP addresses rather than subnet/host names in access control lists.

A Web proxy acts as a "half-way" Web server: network clients make requests to the proxy, which then makes requests on their behalf to the appropriate Web server. The main benefits of Web proxying are:

- Local caching: a proxy can store frequently-accessed pages on its local hard disk; when these pages are requested, it can serve them from its local files instead of having to download the data from a remote Web server.
- Network bandwidth conservation: if more than one client requests the same page, the proxy can make one request only to a remote server and distribute the received data to all waiting clients.

Web proxying has the following disadvantages:

- Web content is becoming more and more dynamic, with new developments such as streaming video & audio being widely used. Most of the new data formats are not cacheable, eliminating one of the main benefits of proxying.
- Clients have to be explicitly set to use Web proxying; whenever there is a change (e.g. proxy is moved to a new IP address) each and every client has to be set up again.
- A proxy server operates above the TCP level and uses the machine's built-in protocol stack. For each Web request from a client, a TCP connection has to be established between the client and the proxy machine, and another connection between the proxy machine and the remote Web server. This puts lot of strain on the proxy server machine.

Both these benefits only become apparent in situations where multiple clients are very likely to access the same sites and so share the same data. Most of the times network administrators face up the following situations:

- It is necessary to force clients on the internal network to use proxy, whether they want to or not.
- Clients should use proxy, but don’t want them to know that they are being proxied.
- Finally clients have to be proxied, but don’t want to go to all the work of updating all the settings in the hundreds or thousands of web browsers.

This is where transparent proxying comes in. A web request can be interpreted by the proxy, transparently. That is, as far as the client software knows, it is talking to the origin server itself, when it is really talking to the proxy server [8]. There are two (2) general methods this works:

- The first is when the web proxy is not transparent proxy aware. For this purpose it can be used a little daemon called transproxy that sits in front of the web server and takes care all of the details for the network.
- A cleaner solution is to get a web proxy that is aware of the transparent proxying itself (one of them is squid, an open source cashing proxy).
- Alternatively, instead of redirecting the connections to the local ports it is possible to redirect the connections to remote ports.

NAT and Proxies
The difference between proxies and NAT is that proxies act at the application level whereas NAT acts at the network level. This means that the proxy must be application aware, and conversely all applications using it must be proxy aware. Where NAT is used as the basis of the Internet sharing, access to the Internet is transparent and users do not need to know it is in place.

4 Network Management Services

Central Management and Monitoring
Central management is always a requirement for large Networks that provide services to customers. Central Management is the process of handling network resources and services from a central interface in order to reduce administrative costs and at the same time minimize response times to requests and unforeseen events. By the term Network resources we mean every parameter that a Network's performance and availability is dependent on. This includes factors such as Quality of Service (QoS), uptime, Service Level Agreements, Independent service monitoring and many others. We decided to segregate the management and monitoring process to the basic OSI Layers [6]. So network level services such as IP connectivity, host uptime and network usage are monitored using Hewlett Packard's Openview that that provides a well-tested solution for Large-scale Networks and is modular and extendible. For data-link layer monitoring and management Agilent Advisor is used which has the ability to analyze many low level protocols such as: Ethernet, ATM, MPLS, Frame Relay, Routing protocols and isolate potential problems that may appear. Finally for application level services such as web-server uptime, Disk Usage, bandwidth usage per service, we use two open-source tools Multi Router Traffic's Graph (MRTG) and Nagios. Both provide the management console with a high degree of available information regarding service status, application level alarms, administrative notification of failures and graphical representation of the network services' operation.

Disaster Recovery
There is no doubt that disaster avoidance can not always be controlled, but through diligent planning and preparation, disaster response can be controlled. Disaster recovery refers to the restoration and continuance of critical IT infrastructure. Companies today follow the continuous business paradigm, which combines high-availability solutions with advanced disaster recovery techniques. The goal is to manage un-

planned situations with minimal or zero disruption. The ideal scenario when an unplanned event does occur is:

- **Recovery Time Objective (RTO)** to be minimal. RTO or maximum allowable downtime, describes the time within which business functions or applications must be restored.
- **Recovery Point Objective (RPO)** to be minimal. RPO describes the point in time to which data must be restored to successfully resume processing (often thought of as time between last backup and when outage occurred).
- **Costs** of solution and resources to be minimal.

The above three factors (RTO, RPO and cost) are often used as the basis for the development of recovery strategies (Disaster Recovery Plan - DRP) and as determinants as to whether or not to implement the recovery strategies during a disaster situation.

The primary objectives of a DRP are to guide an organization in the event of a disaster and to effectively re-establish critical business operations within the shortest possible period of time with a minimal loss of data. The goals of the planning project are to assess current and anticipated vulnerabilities, define the requirements of the business and IT communities, design and implement risk mitigation procedures and provide the organization with a plan that will enable it to react quickly and efficiently at the time of disaster.

In the event of a malicious attack to a network, data may be irretrievably lost. This could be catastrophic for a company. In this case, the best cure is data recovery from backups. When looking for improvements, special focus must be given to a recovery component that will have the greatest impact on recovery objectives.

Advanced Recovery Technologies

The main advanced recovery technologies used in today's networks are:

- **Electronic vault.** Electronically forwarding backup data to an offsite server or storage facility. Vaulting eliminates the need for tape shipment and therefore significantly shortens the time required to move the data offsite.
- **Standby Operating System.** The restore and recovery of the operating system is the most basic recovery step. Maintaining a copy of the customer's specific system on disk directly attachable to the recovery processor provides the ability to bring systems up immediately at time of disaster. The components of a standby operating system are weekly backups of the system used for production, a method of transporting the backups and a weekly restore.
- **Remote Journaling.** If the main focus of improving recovery is on the recovery point rather than the recovery time, the best choice is remote journaling. Remote journaling entails intercepting the writes to a local log or journal and transmitting a copy of those writes off-site in real time mode, providing for recovery to a point extremely close to the point of failure. The components of a remote journaling solution are a local application that creates local logs or journals, software to intercept and transport the writes to a remote location, and software to be used at time of test or disaster to apply the remote updates to the last valid backup data.
- **Remote Mirroring.** Remote Mirroring is the duplication of data in real time to ensure continuous availability, currency and accuracy. True mirroring will enable a zero recovery point objective. Organizations with the most to gain from remote

mirroring are those that can easily segregate critical applications for mirroring and recovery, or those whose applications are critical enough to warrant the added expense of remote mirroring.

- **System Replication.** System replication is simply the capability to provide a continuous operating environment by duplicating systems, data and network at a remote location. System replication is the most comprehensive solution for addressing RPO and RTO.
- **Hot Network Node.** In any recovery situation, users need access to their data. One way of preparing for this connection of recovery data to the production network is to locate a hot production node in the same location as the recovery capability. The benefit of having a hot network node is that it is continually monitored and in use, thereby minimizing failure potential. Establishing network communications at time of disaster can be complex and time consuming; prestaging of the configuration eliminates error and excess recovery time impact. Circuits into the hot node can also be leveraged to support data transmission for other advanced recovery solutions.

In the following table, the impact of each recovery solution to the three critical factors (RTO, RPO and cost) is depicted. We used a combination of System Replication and Hot Network Node in order to balance cost and recovery in case of failure. We also employed several legacy approaches such as regular backups. Tests were conducted under stress conditions and the result was a surprising minimization of RPO and a fair reduce in RTO as well.

	Electronic Vault	Standby Operating System	Remote Journaling	Database Shadowing	Remote Mirroring	System Replication	Hot Network Node
Impact to RTO	↓	↓	↓↓	↓↓↓	↓↓↓↓	↓↓↓↓↓	↓↓
Impact to RPO	↑↑	↑	↑↑↑↑	↑↑↑↑	↑↑↑↑	↑↑↑↑↑	↑
Cost	$	$$$	$$	$$$$	$$$$	$$$$$	$$

5 Conclusions

In this paper we describe a custom network security design that has the properties of being: designed top-down (with the network element in mind), using compatible protocols and techniques and being modular. The proposed network integrates both new and well-known elements. In future work, we intend to continue development of the proposed architecture as well as define an independent model for tying together these elements in a comprehensive framework that would be ready for deployment by corporate customers.

References

1. Aboba B., “IPSec-NAT Compatibility Requirements”, IETF Informational draft (draft-aboba-nat-ipsec-04.txt), (2001)
2. Egevang, K., Francis P., “The IP Network Address Translator (NAT)”, IETF RFC 1631, (1994)
3. Groot G., Karrenberg D., Lear E., Moskowitz B., Rekhter Y., “Address Allocation for Private Internets”, IETF RFC 1918, (1996)
4. Herscovitz, E., Secure Virtual Private Networks: “The Future of Data Communications”, *International Journal of Network Management* 9:213-220 (1999).
5. Kent S., Atkinson R., “Security Architecture for the Internet Protocol”, IETF RFC 2401, (1998)
6. Stevens, R., “*The Protocols (TCP/IP Illustrated, Volume 1)*”, Addison-Wesley, (1994)
7. Mc Kee, Jason, “DMZ Architectures”, SANS Institute, (2003)
8. Kiracofe, Daniel, “Transparent Proxy with Linux and Squid”, (2002)
9. Rigney, C., Willens, S., Rubens, A., Simpson, W. "Remote Authentication Dial In User Service (RADIUS)", IETF RFC 2865, (2000)
10. Finseth, C., "An Access Control Protocol, Sometimes Called TACACS", IETF RFC 1492, (1993)
11. Zseby, T., Zander, S., Carle, C., "Policy-Based Accounting", IETF RFC 3334, (2002)
12. Beadles, M., Mitton, D., "Criteria for Evaluating Network Access Server Protocols" IETF RFC 3169, (2001)
13. Vollbrecht, J., Calhoun, P., Farrell, S., Gommans, L., Gross, G., de Bruijn, B., de Laat, C., Holdrege, M., Spence, D., "AAA Authorization Framework", IETF RFC 2904, (2000)
14. de Laat, C., Gross, G., Gommans, L., Vollbrecht, J., Spence, D., "Generic AAA Architecture", IETF RFC 2903, (2000)

Web Resources

- http://www.business.att.com/content/whitepaper/3118web.pdf
  “A Guide to Business Continuity Planning: Protection and Recovery Services for Your Communications Infrastructure “, *An IDC Executive Brief, October 2001*
- http://www.itpapers.com/cgi/PSummaryIT.pl?paperid=8104&scid=76
  “Disaster Tolerance: The Technology of Business Continuity”
- http://www.drj.com/glossary/drjglossary.html
  “Business Continuity Glossary”, Disaster Recovery Journal
- http://www.comp-soln.com/BCP_whitepaper.pdf
  “Follow-on to “Disaster Recovery Planning-Process & Options” White Paper, April 2001
- http://www.comp-soln.com/DRP_whitepaper.pdf
  “An Overview of Disaster Recovery Planning Process-From Start to Finish”, March 1999
- http://www.highend-consulting.com/web/Arch/Brochures/data%20recovery.pdf
  “Data Recovery”, E-Net
- http://www.cisco.com/warp/public/cc/so/neso/wnso/power/disre_wp.pdf
  “Disaster Recovery Planning: Don’t Forget the WAN”