attempted to enforce against the Bank of Tokyo. The Court said that the recipient does not have to enquire into the authenticity of the message they receive unless he or she is on notice of dishonesty. The sender must keep secure any key, in this case a private telex line used to authenticate a message.

Meta-Firewall: A Sixth Generation Firewall — Part 2

Oliver Lau

Terminology in the firewall area is still confusing. Proxies, packet filters, 'stateful' filters, hybrid approaches, fifth generation firewalls and many more rule the market, and thus rule the user's mind of what is good and what is bad. But few people have thought about the relationships between all those technologies, how they can interact, and how they can be integrated to increase security on a perimeter network to a maximum. Let us call this approach a 'meta-firewall', designed to provide maximum security for dedicated purposes. All of the issues involved in planning for a solution for any network cannot be discussed, but it is an approach to a new way of thinking what can be done with firewalls and the like. This concluding part of a two-part article continues to build the layers of security to make a 'meta-firewall'.

Adding a second obstacle

Figure 1 shows how it is done. The proxy has been relocated to the DMZ. It no longer serves as a transition from the corporate network to the untrusted environment but as a single-homed facility to control traffic on the application level. A proxy server placed in this manner is commonly known as a bastion host.

If we don't have a proxy to automatically translate IP addresses, then we need to have another solution to make seamless interaction between public and private networks possible. Reminder: although it is possible to route a packet from a network with private IP addresses through a gateway outwards to a public network (the Internet), there are no routing table entries on any of the Internet routers that point back to private networks, so answers would never travel the other way round, instead replies would vanish silently.

A facility to fulfill the task of network address translation (NAT) could be a firewall that is able to treat packets to and from a private network, referred to as a stub network, conforming to RFC1631 (1). Some NAT facilities are only able to 'hide' internal addresses by mapping them to a single IP address (or sometimes a pool of them) which is the so-called 'masquerading' or 'single-IP-resolution'. True NAT might also map ingress traffic to internal addresses depending on the original destination address and/or destination port. With this in mind, we should consider an RFC1631-compliant firewall for our fictitious company, for it provides a maximum of scalability and a minimum of configuration changes. The traffic is now as follows.

Mail

Incoming E-mail from the Internet is relayed through the proxy to the internal mail server. For packets to be routed inwards through the firewall, an entry to the NAT table like the following is added:

```
untrusted                trusted
194.163.133.195, 25  192.168.0.1, 25 TCP
```

This means that packets arriving on TCP port 25 (SMTP) at the external interface of the firewall (194.163.133.195) are redirected to the internal mail server (192.168.0.1) on port 25.

Outgoing E-mail is processed by the internal mail server which decides whether to deliver mail locally or to relay to a host on the Internet.

Information services

For incoming HTTP and FTP requests the situation has not changed. Outbound requests now first have to pass the firewall,

then are served by the proxy.

DNS

DNS traffic has not changed significantly either. Requests from the trusted network are served by the internal DNS, requests from the untrusted side are served by the bastion host.

SQL

Still, there are connections to the internal network from the untrusted side of the firewall: from the information server to the SQL database and from the mail relay to the internal mail server. As we have stated above, this is not very prudent from a security perspective.

Only inside out

Professional database systems provide a function to mirror portions of the stored data on a remote system. That means, the remote database contains a subset of the information located in the main database. You can then restrict access to this duplicated information to 'read-only' except for the internal database server which needs to update the information on the mirror database regularly. This has been done in the topography sketched in *Figure 2*. The advantage is that the information server no longer has to initiate connections to the internal database mirror in the DMZ, which — provided that filtering rules on the outer router are set properly — no one on the Internet has direct access to. A drawback could be, that information update on the mirrored database could be delayed a little bit in time.

How could the relationship between the mail relay on the bastion host and the internal mail server be modified to shift connection initiation from the untrusted side to the trusted side? One simple solution to this problem might be to move the internal mail server into the DMZ. But you have to be very careful with that. The host has to be protected strongly, because of the user-related information it holds in the form of private mail accounts including the necessary passwords, whether encrypted or in plaintext.

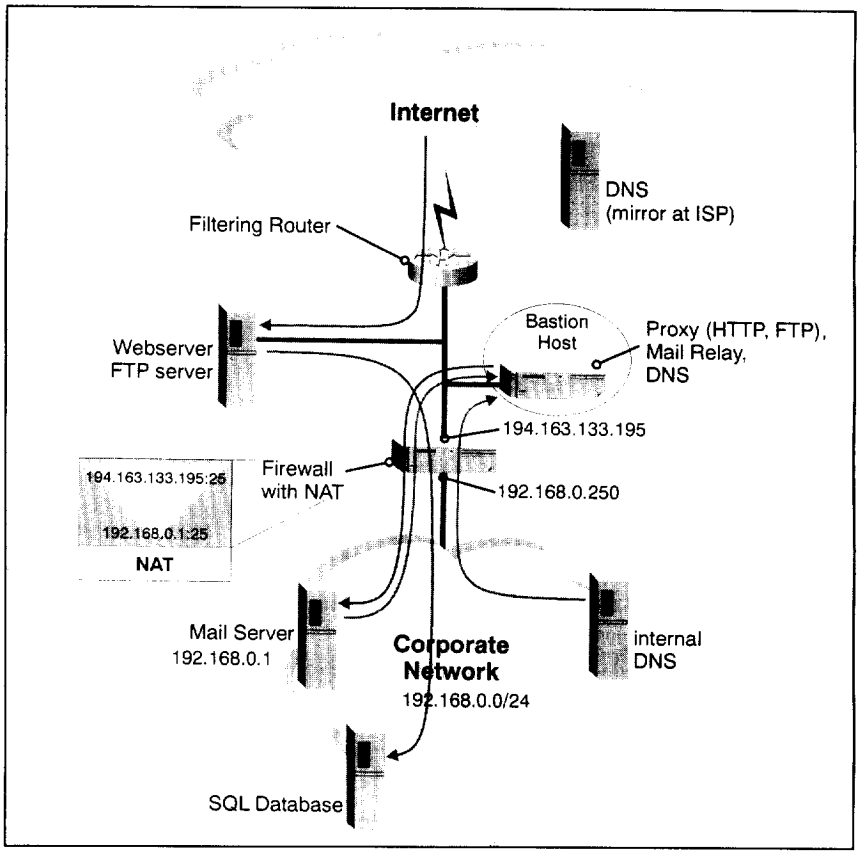

Figure 1: The second obstacle.

Having all connections originating from the trusted side of the corporate network, we don't need inbound NAT anymore. Instead, we have to establish rules on the firewall to map the private addresses to public ones. This is done via outbound NAT.

```
trusted              untrusted
192.168.0/24,*       194.163.133.195,*
```

This means that all connections from the trusted side of the firewall (network 192.168.0.0/24) to the Internet should be masqueraded with the public IP address of the firewall. For every outgoing connection the firewall holds an entry in its state table that describes how destination addresses of the reply packets should be retranslated to internal addresses.

Due to the benefits of contextual (stateful) filtering, the firewall

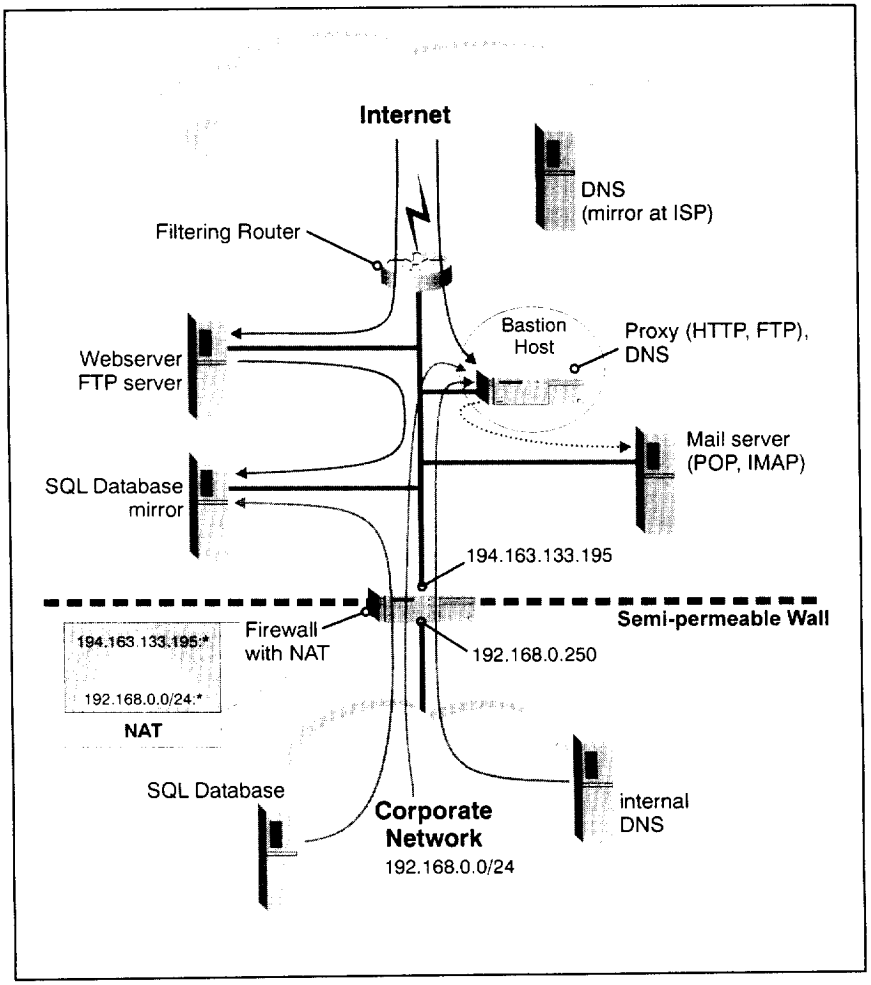

Figure 2: Outbound only.

protecting the corporate network can now act as a semi-permeable wall. Exclusively outbound connections are allowed, reducing the risk of misconfiguration.

Adding content security

Until now we have channelled traffic which is tantamount to controlling data flow up to the application layer. But inside the applications transported with the different protocols sometimes harmful or nauseating data resides. The omni-present SPAM mail, malicious Multipurpose Internet Mail Extensions (MIME) attachments, URLs pointing to questionable resources, vicious Java applets, malignant ActiveX code, viruses within downloaded files have to be recognized, too.

A technique commonly called 'content vectoring' is helpful with this task. Proxies are aware of application-specific details in data flow, they are the points at which traffic has to be detoured to facilities which are able to prove the data for validity and benignity. How this works is depicted in *Figure 3*. As all mail, all file transfers and all WWW requests are handled by the bastion host, it's easy to divert those information from there to a screening host which does all the above inspections, remove illicit data, alerts the system's manager in case of peril ahead, and kindly forwards everything that's safe.

Ultimate?

The last step is to segregate services, so that there is one machine for one service, to add a further line of defence in the form of firewall capable of NAT, and reconfigure the inner firewall to be 'invisible' on the network, a mode sometimes called 'stealth' or transparent mode. The functionality of such an approach is as follows.

The outer router provides a first line of defence by means of simple packet filtering based on source address, destination address, source port, destination port and the presence/absence of the ACK bit. If a packet reaches the outer firewall which normally should have been blocked by the router, alarm bells should ring immediately.

The outer firewall's outer interface has four IP addresses. Depending on the destination address and port (the combination of IP address and port is called a socket) of an incoming packet, NAT rules map that destination socket to a new one targeting to a certain host in the DMZ, as in the following table:

untrusted	trusted
194.163.133.194 25	192.168.0.2, 25 TCP
194.163.133.195 53	192.168.0.1, 53 UDP
194.163.133.196 21	192.168.0.3, 21 TCP
194.163.133.197 80	192.168.0.4, 80 TCP

```
trusted              untrusted
192.168.0.0/24,*     194.163.133.195,**
```

That means: the mail server is accessible at 194.163.133.194 with its default port 25 for any ingress connection, and that those packets are redirected to host 192.168.0.2 on port 25. The DNS is accessible at 194.163.133.195 on UDP port 53; packets to that socket are redirected to host 192.168.0.1 on UDP port 53; and so on. So the main task of the outer firewall is to channel inbound traffic to the certain target hosts, and to masquerade the source address of outbound connections with its own IP address.

Diversifying services onto separate machines has a simple rationale: less services, less code; less code, less flaws; less flaws, less security holes; less security holes, less exploits. Besides that, stability and performance are enhanced. If you interconnect the hosts and firewalls on the perimeter network with a (highly configurable) switch instead of a hub, traffic can be channelled even more granular, for you may even control broadcast propagation and IP-to-MAC-address resolution. In combination with static ARP tables on the hosts this type of control also provides ARP 'anti-spoofing', i.e. unauthorized additional hosts could in fact be attached to the switch, but they would never ever be able to seize frames from the line that are directly sent to it on the Ethernet layer. The only chance to read traffic that is not destined to it is through passive eavesdropping.

The internal firewall is supposed to be 'transparent' i.e. that its filtering engine implements

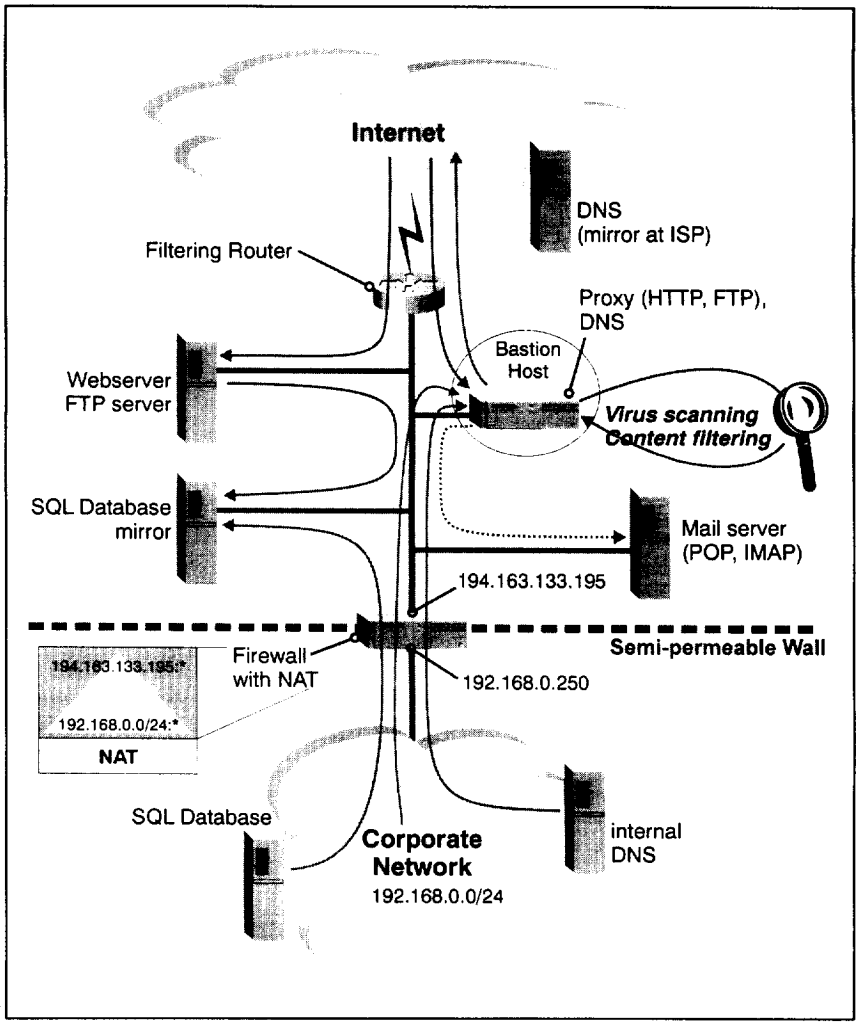

Figure 3: Content filtering.

an own stack for the supported protocols or overlays an existing stack with its own. A stateful filter for instance ideally relies on its own code to sequence and route packets. Because the internal firewall knows about the context of a communication process and only egress connections are accepted, it provides a semi-permeable wall (see *Figure 4*).

Adding speed and security with a proxy

It is often heard that proxies are slow. But this is not always true. Of course, proxies have to process packets on a high level, the so-called application level. They have to emulate the behaviour both of the sender and the receiver of the packets. You can imagine that this costs a lot of CPU cycles. But there are other approaches which could tremendously

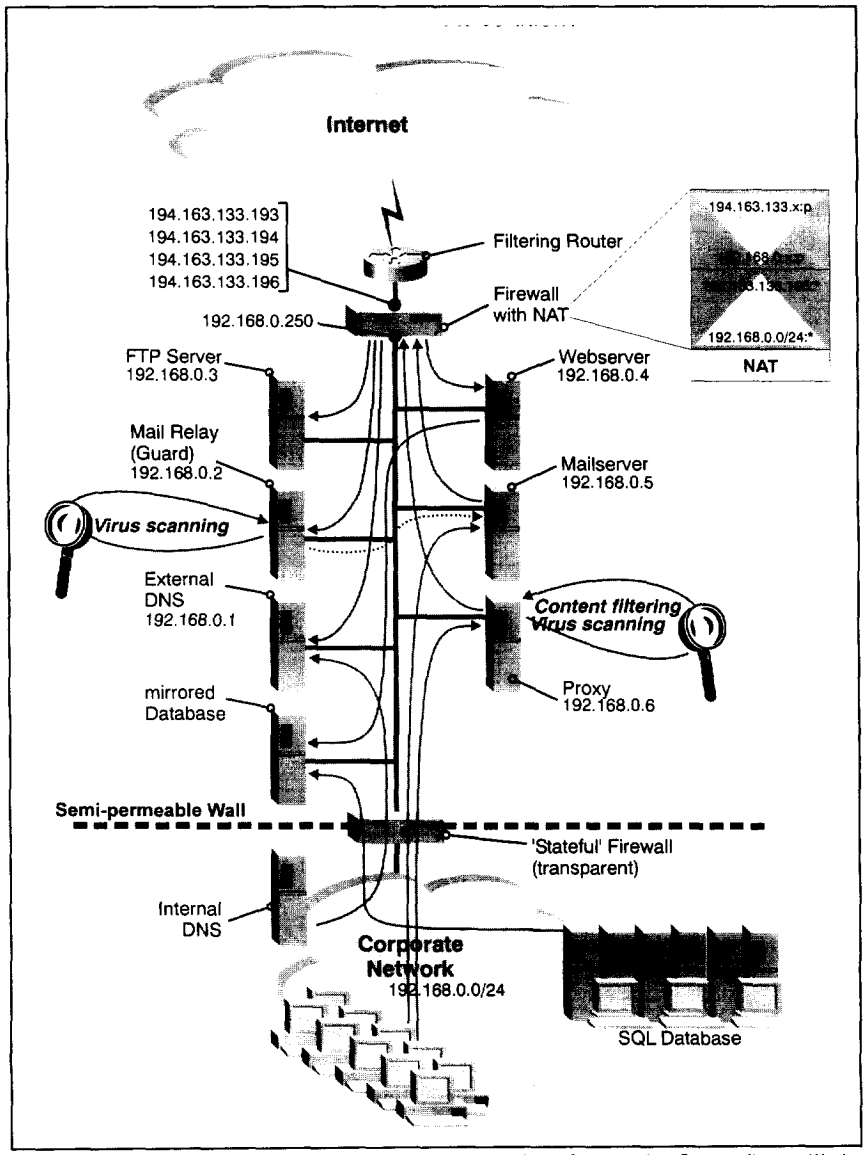

Figure 4: Meta-firewall. Deployment example of a meta-firewall: multiple 'moats' on the way to the internal network have been established. Services have been distributed to distinct machines, to minimize the risk of both a packet 'fall-through' from the untrusted side to the trusted side and of downtime due to a successful attack against any target. Deployment of proxy firewall strengthens the accompanying filtering measures.

increase network performance especially over slow links: caching proxies. The term of caching has a long history and is used in a variety of situations.

Our main interest lies in caching information that is found on the Web. Any time a caching proxy is requested to access a certain page on the Web or to download a file via FTP, it has a look in its more or less temporary storage to find out if the information has already been retrieved. If so, the data can be directly transferred from the cache (in best case the volatile memory, in worst case the hard disk drive) to the requester, saving time and bandwidth. As a rule of thumb, the more users share a cache the larger it should be. The larger a cache is, the more users have access to it and the smaller the assortment of information on the Web, the more bandwidth (and sometimes money) is saved.

Conclusion

Anyway, there are situations in which an attacker may traverse the security perimeter in an unlawful way. For instance, the encapsulation of an unknown and thus dangerous protocol in a well-known protocol (tunnelling) is very hard to detect. Neither proxies nor contextual filters are yet capable of responding to threats caused by protocols like Remote Procedure CAll (RPC), which in fact maps TCP/UDP ports to RPC services dynamically.

The situation aggravates dramatically, when the lines are encrypted, for instance through IPsec or worse within Secure Sockets Layer (SSL), Secure HyperText Transfer Protocol (S-HTTP) and Transport Layer Security (TSL) connections, where the firewall has no management via Internet Security Association Key Management Protocol (ISAKMP) or X.509 has to be extended to the firewalls. But the continual spreading of Internet Protocol Version 6 (IPv6) will most likely coerce developers to

take care of the evolving problems concerning confidentiality, integrity and authenticity.

Moreover, some sites cannot afford to restrict access to their sensitive or even mission-critical resources due to political reasons. There are many situations which allow attackers to eavesdrop lines, to erase and modify data (sometimes even unrecognized), in brief: to wreak havoc on your IT systems. Regardless of how meticulously firewalls have been adapted to your networked environment, they are not able to protect you 100%. Meta-firewalls are not the ultimate solution. But the multi-tiered approach is a huge step forward towards it.

We have learned to take a lot of parameters into account to plan for a secure gateway to the Internet and to provide a sound basis for further development. These parameters are:

- functional specification
- firewall management issues (not discussed in this article)
- level of security to be achieved as a result of proper asset assessment and risk analysis
- requirements in terms of reliability/stability/availability
- traffic volume (which has not been discussed in this article)
- future aims

However, the best technical measures and the best preparation is futile, when supporting methods like a strict and thought-over security policy, procedural papers, and training for users and administrative personnel, are not established.

References

(1)Egevang, K. and Francis, P., 1994. The IP Network Address Translator (NAT), RFC1631, The Networking Group, May 1994.

Glossary of terms

Black — a term commonly used in cryptography which describes a trustworthy channel; counterpart to red.

Cache — a repository for information that can be quickly accessed.

Perimeter — the outer boundary of a network, ideally the only physical connection to the outside world.

Red — a term commonly used in cryptography which describes an untrustworthy channel; counterpart to black.

Useful URLs

http://www.checkpoint.com/

http://www.cisco.com/offer/cf/d528b166xa.html

http://www.cisco.com/univercd/cc/td/doc/product/iaabu/pix/pix_v41/pixrn414.htm

http://www.digital.com/internet/solutions/have_web_brief.html

http://www.network-1.com/

http://www.raptor.com/

http://www.securecomputing.com/

http://www.tis.com/

http://www.data.com/cgi-bin/dynamic/lab_tests/firewalls97_extras2.txt

http://www.cert.dfn.de/fwl