do not do anything about other IP denial-of-service attacks nor do they help with generalized packet 'clogging' attack techniques.

- Stateful attacks. There are still methods to attack a session in a stateful manner if the AH facilities are not used in conjunction with ESP or other solutions.

- Logging and accounting. These facilities are critical to evidence gathering and information analysis of potential attacks. None of the IPV6 facilities provide for logging of activities and accounting of user, session, component, application or other traffic generating or receiving facilities.

Of course, there are other methods to attack the facilities as well as other manners in which to use IPV6 facilities to defeat attacks. The illustration above serves, however, to show that there is not an all-encompassing approach in the offerings in IPV6 and customers should be careful to temper their enthusiasm with the reality of what the facilities can and cannot accomplish (properly implemented, of course).

Overall, IPV6 security facilities dramatically improve connection security, but the main items to remember are that they are not all-encompassing and lack many specific security facilities and logging facilities critical to overall security management. And, while many vendors are offering the facilities retrofitted to IPV4, they still require many other security technologies to properly secure network connectivity and session connections.

Meta-Firewall: A Sixth Generation Firewall — Part 1

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 (1) 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. The first part of this two-part article begins to build the layers of security in a firewall for an imaginary company.

What exactly is a firewall?

A firewall is any component that is capable of controlling and filtering traffic between at least two networks. As such, a firewall is both ingress and egress, depending on the side you are looking at it. When the term 'firewall' is addressed throughout this article it is in the sense of a dual-homed gateway (i.e. the firewall could be attached to only two networks) controlling traffic between a trusted and an untrusted network, where the trusted network mostly is the corporate network or a local area network, and the Internet mostly belongs to the untrusted side.

Although arguing based on an analogy could be dangerous, because it could lessen the accuracy of the real situation, here is a trial on firewall concepts: firewalls are the doors between networks, meticulously controlled by doorkeepers (the filtering engines), who look at you, inspecting your ID card, asking you where you want to go, where you come from and whom do you want to talk to, and either say, "Ok, you may pass!" or, "Stop, access prohibited!".

Some of the doorkeepers are really dumb, they can't remember if you have already entered or not. Others are 'stateful', e.g. when they see you twice attempting entrance without exiting, this is not allowed and you are refused. A third group of doorkeepers would never let you pass, instead they ask you for your message and whom it should be delivered to,

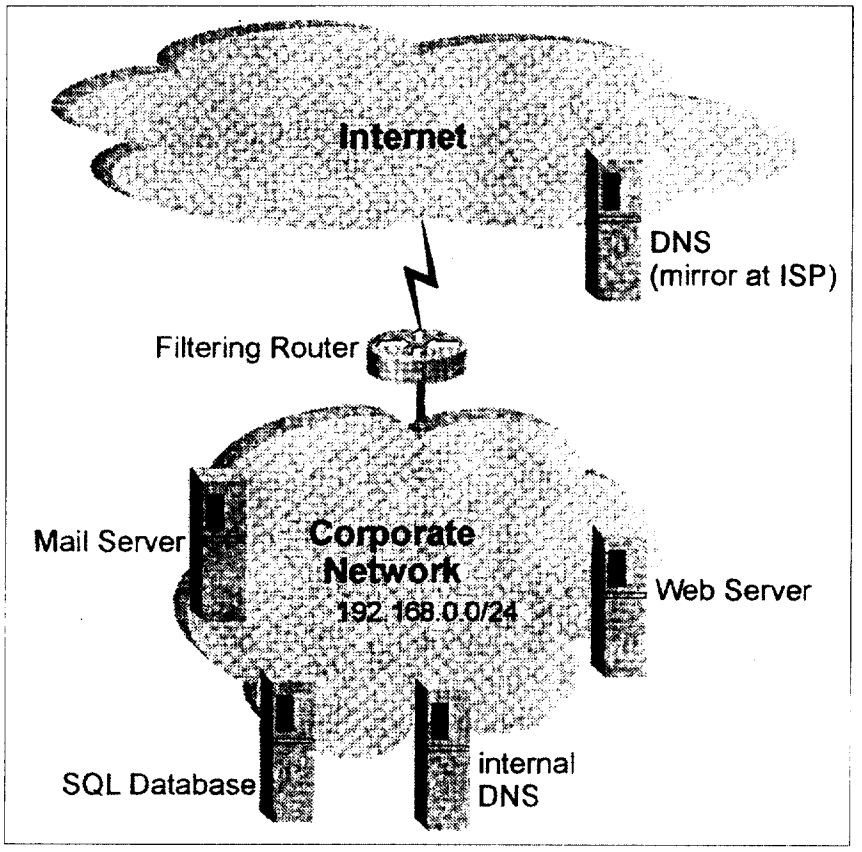

Figure 1: A very simple firewall. A firewall router is put in between the corporate network and another untrusted network, e.g. the Internet. The firewall router is equipped with packet filters, which restrict access to certain services ("Everything is allowed which is not explicitly denied"), but also to open up holes in the router to allow only specific traffic ("Everything's denied that is not explicitly allowed!"). This approach provides a single line of defence, which can relatively easily be broken. If it is broken, the attacker has his way open to do anything he wants.

then stalk away to inform the destination about your request and come back with a reply. When you are faced with a doorkeeper of the latter kind you have encountered a proxy.

Basically, firewalls have to:

- prevent private information flowing out
- protect against malicious data coming in
- protect against 'red' interlopers
- control what can be accessed by 'trusted' users
- prevent denial-of-service attacks on publicly accessible machines

In summary, firewalls are used to protect internal resources against attacks from the outside and to minimize the risk of information leakage to the outside. Certainly, you will have heard of multi-layered defence, of hybrid firewalls, all of which should increase security on your network. But how does it work? How do you achieve enhancements through layering? What is to be taken into consideration so that the multiple layers are used effectively? How do the technologies compare? How could they be integrated?

A simple approach

Consider *Figure 1*. This example deployment of a dedicated firewall router has been the state of the art in network security at many sites. Fortunately, there were experts like Cheswick and Bellovin (2, 3) who were paranoid enough to enhance it and they have taught us how to make our networks more secure. *Figure 1* shows a dual-homed gateway as a 'semi-intelligent' gateway between a corporate network (trusted) and the Internet (untrusted). Such a construction fits most of the purposes where the firewall is used as an entrance from a local network.

Ingress and egress

What does that mean? Basically, there are two directions a packet can travel — inbound and outbound — and there are two possibilities from where the connection between two hosts could have been established — inside or outside. When egress traffic is addressed, then we talk about traffic that belongs to a connection initiated on the local (trusted) network, when ingress traffic is addressed, then we talk about a connection that originated from the untrusted side of the firewall. For TCP/IP that means that ingress traffic originates from the untrusted side and the first packet travelling inwards carries the synchronize sequence numbers (SYN) bit. Of course, there are packets travelling outwards, but nonetheless we speak of ingress

traffic. Remember, the first packet determines the direction of the connection.

Fiction

Thoroughly discussed case studies help teach security so I am going to introduce you to a fictitious company whose operational demands will show what should be taken into account in establishing a 'smart' boundary between a corporate network and the Internet.

Assets

About 2000 people are currently employed in this company, hence presenting an enormous potential for helpdesk hassles, computer misuse and many more issues that could mess up security-related and administrative advances. Files are kept on tens of thousands of business partners, including credit card numbers, purchasing statistics, contracts and much much more data related to individuals. Moreover, their central computing database holds dozens of patents, research and development papers, and of course, the whole finance and accounting data.

Functional specification

Our fictitious company has long-term experience in running local networks, but has never heard of the security implications coming with access to the Internet with its population of more than 30 million hosts. The administrators are pretty informed about host security they think, because they know about various authentication methods and how to make 'sandboxes' in their computers. Although statistically more than 80% of all security incidents are inside jobs, nothing really bad had happened on their network except some accidental erasure of database entries or some loss of files, which have been quickly recovered thanks to prudent backup policies.

Now the time has come to get connected to the Internet. Here is what they want to do:

- Internal and external E-mail.
- Present themselves on the Web with their own Web server.
- Share some files with their customers and the public domain, so they would like to install an FTP server.
- Quickly obtain patches from vendors, necessitating access to various FTP sites on the Web.
- Visit Web sites.
- They have a lot of customers who were happy to see what status their orders have reached. The necessary information for order tracking resides on a huge clustered SQL database server on their local network. To ease the users' task of information retrieval the database should be enhanced with an HTML browser front-end.
- Until now only IPX and IP packets have been seen on their network, the range of IP addresses is one of the privates mentioned in (4). For

The Nature of Risk

Security is defined as "freedom from danger, risk etc." (*Webster's Encyclopaedic Unabridged Dictionary of the English Language*, 1989). But security also has a subjective meaning: "certainty of being protected against danger, risk etc." (Martin Sauer). Before it becomes clear what kind of security is achievable in the real world, you have to know about danger, risk, vulnerability and threat.

There are threats all over in our lives, mostly threats against the following IT targets: confidence, availability, integrity and authenticity. The threats are in stored data or data in transit. Threat can be divided into three classes:

- careless action (misorganization, misconception, wrong products, accidental erasure of data, accidental infiltration of viruses etc.)
- deliberate action (theft of seizable objects, theft of virtual objects like information, modification of accounting data for your own benefit etc.)
- flaws in hardware and software (hard disk drive failures, error-prone operating systems, software applications etc.)
- environmental influence (loss of power, power peaks, fire, storms, acts of God etc.)

Vulnerabilities do not depend on extraneous effects, but are immanent to a system, either technically or by operation.

Careless and deliberate actions are triggered either outside or inside an IT system. In the same way errors occur in any part of the networked hard and software, e.g. in switches, routers and protocol stacks, but they also can come from the outside in the form of broadcast storms, unusual heavy packet fragmentation, misled frames, broken connections etc.

The classical mythology speaks about a Greek warrior called Achilles who was killed when Paris wounded him in the heel, his only vulnerable spot. So, there are times when a threat meets a certain vulnerability and when this happens you are in danger.

To sum up, when we talk about Achilles' heels, we talk about vulnerabilities; when we talk about spears, we talk about threats; when the spear hits the heel, you are probably at risk.

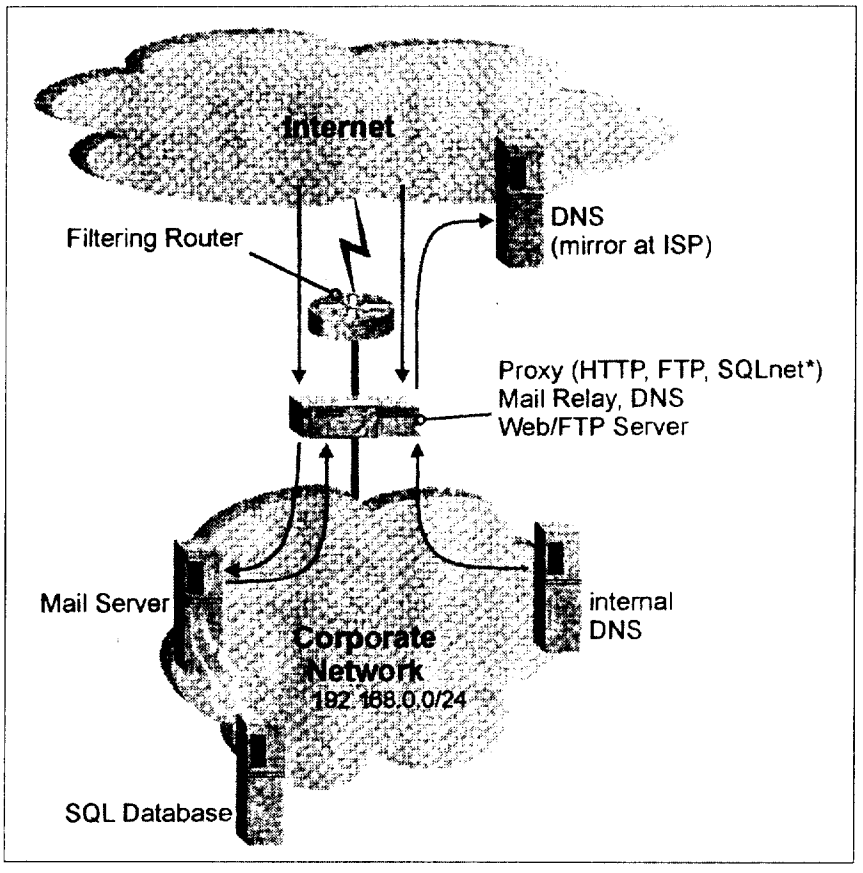

Figure 2: A simple proxy configuration.

both security reasons and convenience they decide to keep address assignments. This requires some kind of network address translation.

- On the Internet nothing works smoothly without a Domain Name Server (DNS), so they have to install one locally besides the mirror at their Internet Service Provider (ISP).

Technical specification

The most imperative task is to convert the internal addresses to official ones that can be used on the Internet. Following the notion of "keep it small and simple" you could take an ordinary Linux box into account. A nice tool capable of address hiding, i.e. unilateral address translation where any internal addresses are mapped to a single official address, is shipped with most distributions. This tool, known as ipfawdm (IP firewall administration) can also act as a packet filter down to the bit level, i.e. it's possible to inspect the value of the SYN, ACK (acknowledge) and TOS (Type of Service) bits and to establish filtering rules according to them.

Another way to convert IP addresses is by using proxies. Due to their capability of sitting independently in between the two interfaces, packets have no longer to be handed over (routed, forwarded) from one network adapter to the other. Instead, two separate connections are established on either side.

But most of the proxies are very limited in handling traffic due to their nature of being specific to applications. That means, for any protocol to be handled by the proxy a specific proxy application must exist. Happily, our company has no extraordinary applications to service so a proxy may be a good choice for them.

Authentication at the proxy server

Two of the firewall technologies (circuit-relays and application level proxies) allow it to authenticate users, which could be very helpful in terms of audit trails (who accessed what and when?), billing (not applicable for our fictitious company), authorization, (Accounting and Finance should not be granted access to several Internet resources. Some R&D workstations need full access) and privacy of information (only specially trained and authorized users might use the Internet for business purposes).

The more authentication methods are integrated in the proxy or circuit-relay, the more flexible the system is. Common implementations provide user identification by means of one-time passwords (OTPs), Kerberos, RADIUS, TACACS, challenge-response mechanisms, plaintext passwords, X.509 and some more. Our fictitious company decides not to establish authentication at the proxy at the moment, but is open for further discussions about it.

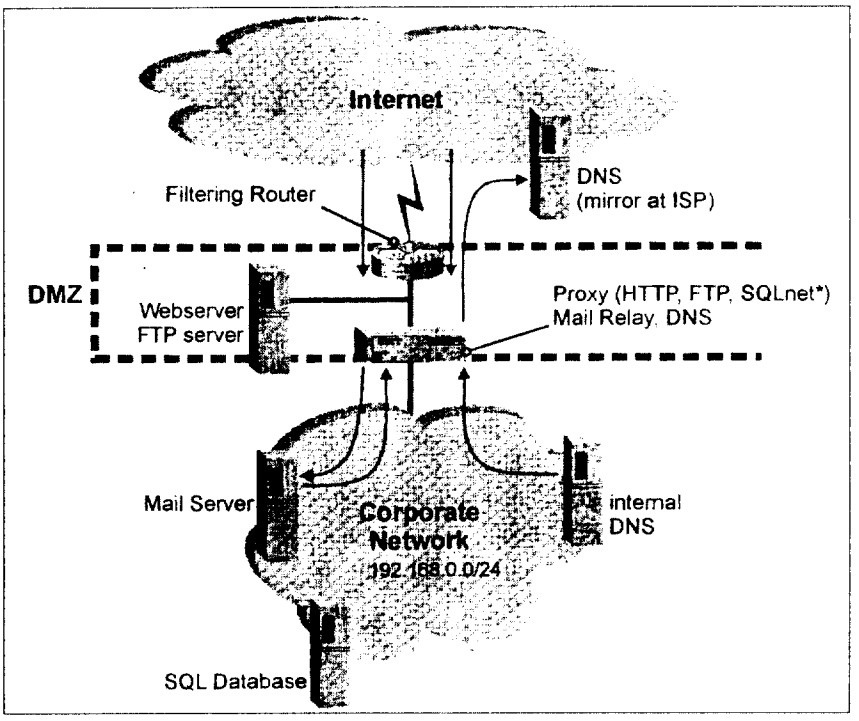

Figure 3: A simple proxy configuration II.

A simple proxy configuration

The proxy in *Figure 2* has a task similar to the one of the firewall in *Figure 1*. Instead of letting packets through directly from the client on the trusted network (the black network, a term borrowed from the cryptographers) to the server on the untrusted (red) network, the proxy establishes a separate connection to the red server. The client in turn actually connects to the proxy, but 'thinks' it is talking to the red server.

The company needs one proxy for HTTP, one for FTP and one for structured query language (SQLnet).

Also, the proxying host has to be aware of the DNS requests and zone transfers and has to act as a mail relay for inbound and outbound E-mail, i.e. mail should be forwarded to their internal mail server and outgoing mail has to pass the proxy. Their internal DNS could be used as a primary for their domain and as forwarder to the external DNS which in turn services inbound requests and could sit on the firewall, too, but it should not be aware of IP-to-name resolutions of the internal addresses. Providing resolution for corporate hosts is what the internal DNS is in charge of. The only solution to host a DNS is on the proxy itself, because traffic flow directly through the proxy is impossible without a separate address translation facility.

A big problem remains: where to place the Web and FTP server? There are two possibilities: on the proxy or on a separate machine on the network between the router and the proxy server. In the first case, one of the sixth of the 'Ten Commandments of Security' (5) is violated; the rule of 'diversified defence', which tells us that it is best to split the services onto different machines and not to let many of them reside on a single computer. In the latter case, the router can be used as an additional packet filter to prevent harmful access to all unused services on the local network and thus to provide a second line of defence. *Figure 3* shows the results of our studies. Let us have a look at the flow of traffic.

Mail

Anyone sending a message out to a colleague on the corporate network utilizes the internal mail server to connect to its SMTP port. The mail is stored there for retrieval by the addressee via Post Office Protocol (POP) or a comparable protocol like Internet Message Access Protocol (IMAP). Sending out a mail message to a participant of the Internet community works as follows: first a connection is established via SMTP to the internal mail server, then the server relays the mail through the proxy to the destination address. Inbound mail is handled similarly: because the external DNS server is configured with a Mail eXchanger (MX) entry for the mail relay, all mail traffic from the Internet is sent to the proxy, then forwarded to the internal mail server with a separate connection. Mail is then delivered locally.

Information services

Web and file transfer traffic is easy to manage. Anybody on the corporate network wishing to access Web or FTP sites connects with their browser to the HTTP and FTP proxy, which itself retrieves the requested data from the server on the Internet. Inbound Web traffic is even less difficult. Anybody on the Internet may directly access their information

server located in the so-called demilitarized zone (DMZ) between the router and the proxy server. Some packet filters on the router restrict incoming traffic, so that only the necessary ports on the information server are 'open'.

DNS

Hosts on the corporate network have to resolve names into IP addresses by asking the internal DNS. The internal DNS is configured both as a caching DNS and a primary for their domain. Any request it has no information about is forwarded to the external DNS which in turn tries to reply. Just to remind you: we think of all the internal addresses as being trusted and the external ones as being untrusted, so the first request for anything should go to an internal server. And so it is for the proxy. Of course, the proxy also has to resolve IP addresses, to accomplish this it has to ask the internal DNS first.

SQL

In the beginning we have claimed that the fictitious company likes to present portions of their database to their customers by means of dynamic creation Web pages. That means, the Web server has to be granted access through the SQL proxy to the internal database. Do you think this is acceptable? Just to remind you of something already stated: egress and ingress. We have said that ingress traffic originates from the untrusted side. As for proxying SQL, this is the case when the Web server requests access to the internal database, although it is possible to restrict access to the SQL proxy to a single way, this was the same thing with incoming mail (see above). In the second part of this article we will see how connections originating from the internal network could provide the same functionality with a higher level of security.

What does the company have now? The first two things that catch the eye are 'single points of defence': the proxy for the corporate network and the (weakly protecting) router for the information server. Although the proxy could also be protected by the router, packet filtering is no longer state of the art for a number of reasons, and trusting it as a measure narrowing the window of a potential attack is only recommended for sites with a relatively low demand for security. Besides, there are many services running on the proxy, so there are a lot of ports open for connection. With every open port the potential for a successful attack is augmented.

So, what is more prudent, to split services onto separate machines or to create a further line of defence? In the first case, you enhance the reliability/availability of the system as a whole. If one machine fails (perhaps as a result of a denial-of-service attack), the others are still up and running. You have added safety. In the second case, the extra measure could increase security significantly, depending on the technology you choose. But that's for later.

Using a second proxy would be unwise because one of the notions of network security states that it is best to use a layered defence that is built from different technologies and different operating systems, so that the exploitation of a flaw in one machine could not be repeated with the second one. So our company should install a distinct product like a (stateful) packet filter, ideally (from a security standpoint) running on top of an operating system distinct from the one of the proxy.

References

(1)Schultz, E.E., 1998. Securing Third-Party Connections, *Network Security*, January 1998.

(2)Bellovin and Cheswick, 1994. *Firewalls and Internet Security*, O'Reilly, 1994.

(3)Chapman, Z., 1996. *Building Internet Firewalls*, O'Reilly, 1996.

(4)Rekhter, B. et al., 1996. Address Allocation for Private Internets, RFC1918, The Networking Group, February 1998.

(5)Lau, O., 1998. The Ten Commandments of Security, *Computers & Security*, Vol. 14, No. 2, Elsevier Science, 1998.

(6)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/cgibin/dynamic/lab_tests/firewalls97_extras2.txt

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

Managing Network Security: Technical Protection for the Joint Venture

Fred Cohen

Over the last few years, computing has changed to an almost purely networked environment, but the technical aspects of information protection have not kept up. As a result, the success of information security programmes has increasingly become a function of our ability to make prudent management decisions about organizational activities. This series of articles takes a management view of protection and seeks to reconcile the need for security with the limitations of technology.

In today's enterprise computing environment, there is increasing need and pressure to form temporary collaborative computing environments wherein a limited portion of the expertise available in two or more companies is joined for the purposes of some project. As a further complication, the joint venturers are often competitors in closely related fields. In this environment, there is a substantial need for rapidly deployed collaborative computing environments that are mutually secure and yet permit authorized sharing.

Many companies have been so pressed to provide adequate cooperation that they have cut corners in terms of security and simply interlinked the partner internal networks, but this leaves a gaping hole in which the team members can intentionally or accidentally harm each other. An example of accidental harm that may result in substantial liability is the exploitation of one partner's network by a third-party attacker to break into the other partner's systems. It seems clear that it is in the best interest of both parties to secure both their network and that of their partner.

Many companies have tried to create internal enclaves in which team members can use select services, but again this creates many difficulties. For example, within the enclave, which may be at the other partner's site, you may not be able to access your corporate E-mail. The goal of the joint venture is to take full advantage of each partner's assets in order to produce the result, and any inhibition of this capability is a potential blockade to your success. Furthermore, users whose job is to get the joint venture to work are rarely willing to sacrifice much effort in order to accomplish desired security.

Recently, I was at a client's site (the client is also a regular reader of this series) discussing this very problem, and I thought it would be beneficial to discuss some of the solutions we discussed.

The styles of collaboration

There are several distinct styles of collaboration in widespread use. They are, to a first approximation:

- An embassy — a location in the other organization's site where your employees work during the joint effort (or the other way around). In the embassy, you are like the foreign office of a nation-state, typically (hopefully) located in a