Published on *Jisc community* (https://community.jisc.ac.uk)

Firewalls and proxies

It has become increasingly popular over the last few years for Janet sites to deploy firewalls; many sites have realised that their users do not require full, open access to all workstations and servers on a campus. By controlling access, staff time can be saved in chasing up ?hacking? incidents, and the types of service used can be kept under control. Bandwidth usage is also rising dramatically, so firewalls now have to be able to operate at gigabit speeds. If transport mode IPsec were used, firewalls would need to trust encrypted sessions passing through them, which some site security policies may not allow. The growth in peer-to-peer applications ? not all of the Moving Pictures Expert Group Layer-3 (MP3) piracy variety ? means new consideration has to be given in security policy documents to enabling end-to-end services, while preserving site security levels.

A guide written in 1998 ?The Use of Firewalls in Academic Environments? [ECSFW] has been made available to sites joining the Janet community. This document does not cover H.323 in any detail, and is, at the time of writing, under revision by the authors for republication by Janet-CERT.

Some of the issues involved are discussed in a now-expired IETF Internet Draft [SHORE].

Firewall considerations

Many firewalls are deployed in a ?default deny? configuration, for connections inbound to the site they protect. The ?default deny? mode means that services are blocked unless they are explicitly allowed. If a user wants to have an external service connect to their workstation or server inside the firewall, they need to consult the appropriate administrator and ask for an explicit rule to be added to the firewall to allow that service in.

Firewall rules generally specify IP addresses and port numbers that connections can be made to/from. For many services, this is a simple addition, e.g. port 80 for HTTP, port 25 for SMTP. For H.323, it is not so straightforward.

It is worth noting that not all H.323 deployments will require firewall use. As described previously in this document an H.323 studio terminal may be directly connected to a campus edge router, and be a hardware device with no services open to external access (or these may be filtered by simple access control lists on the edge router). If this model is followed, i.e. the H.323 device has no or limited firewall protection, then that device should not have any connectivity to other parts of the internal site network.

Dynamic use of TCP/UDP port numbers

H.323 uses a number of well-known TCP and UDP ports for communications, as standardised by the Internet Assigned Numbers Authority (IANA)[IANA]. While initial communications do occur on fixed TCP/IP ports, subsequent port numbers are usually negotiated dynamically as part of the call set-up procedure. These ports fall into the ?user? space, i.e. port 1024 through to 65535. Because these port numbers are arbitrary, and change with each session, it is not generally possible to open up specific ports on the firewall in advance of the H.323 session. The port numbers used cannot be determined without inspection and parsing of the data in the H.323 session initialisation exchange.

There are three general solutions for this problem:

1. The firewall monitors the set-up traffic to discover which ports to open up for the session.
2. The firewall has all UDP and TCP ports from 1024 upwards opened for access, which has severe implications for the security of other services running on high ports when on non H.323-dedicated equipment.
3. H.323 endpoints are used that allow fixed ports to be preset for the communications.

The third option is desirable, where possible.

Figure 6: TCP/UDP ports used by H.323 services

Port	Protocol	Description	Terminal	MCU	Gatekeepe
389	TCP	LDAP (NetMeeting)	X		
1300	TCP	H.235 secure signalling	X	X	
1503	TCP	T.120 data	X	X	
1718	UDP	Gatekeeper discovery	X	X	X
1719	UDP	Gatekeeper Registration, Admission and Status (RAS)	X	X	X
1720	TCP/UDP	H.323 call set-up	X	X	
1731	TCP	Audio call control	X	X	
> 1024	TCP	H.245	X	X	
> 1024	UDP	RTP (video)	X	X	
> 1024	UDP	RTP (audio)	X	X	
> 1024	UDP	Real-Time Control Protocol (RTCP) (control)	X	X	

Many firewall vendors have recognised this problem, and as a result have implemented H.323 modules that sniff the call set-up packet exchange to be able to open ports dynamically.

Such products that claim H.323 support include:

- CheckPoint® Firewall-1
- Nokia IP series
- Symantec Enterprise Firewall
- Cisco® PIX®
- NetScreen®

The Firewall-1 product, available as software and also on the (more expensive) Nokia hardware platform, handles H.323 traffic well, based on tests done in the VIP Project [VIP]. However, some problems with these products have been reported. Future versions of this guide will include summaries of such reports and open issues. Readers are invited to relate their experiences to the authors for inclusion in such revisions.

The presence of H.323 support in a firewall should be required in any university or campus procurement exercise.

A ?default deny? inbound policy can be configured such that an H.323 rule can be added for the specific hosts or terminals that should be allowed access. It is not recommended that site firewalls be run in ?default allow? mode.

Radvision has published a cookbook on H.323 and firewalls that makes good reading for those deploying firewalls and proxies in H.323 environments [RVFIRE].

In the JVCS-IP, sessions will be initiated outbound from the Janet MCUs, thus a firewall could be configured to only allow H.323 traffic from the known IP addresses of those MCUs. However, that would preclude any H.323 conference sessions arranged outside of the Janet service.

As with any ?middlebox? in a network path, one of the considerations for firewall deployment is the delay (or jitter) it imposes in traffic passing through it. Such delays will depend on the performance of the firewall device, the speed it can run at and the level of traffic observed at any one time. In general, there will be no significant delay imposed by such a firewall, e.g. a Solaris?-based CheckPoint? firewall on an Ultra 10 operating at potential Fast Ethernet speeds is unlikely to see traffic delayed more than a few milliseconds when background traffic is of the order of low tens of megabits. If campus traffic is bursty, or much higher for example due to Grid use, then hardware-enhanced firewalls (such as the Nokia platforms) should be investigated.

Proxy deployment

The topology considerations for proxy deployment are described above in Section 3.1. The proxy can reside inside the site, outside the site, or in the site DMZ. Use of such a proxy limits the requirements of trust in a firewall configuration, although it does not eliminate the security threat altogether. The proxy solution is being developed throughout Wales and most of Scotland.

Network Address Translation (NAT)

A proxy has an additional benefit where a site is using NAT and private IP addresses inside its site. While most UK universities have enough IPv4 address space to not require NAT, some sites do use it. In such cases the proxy can act as an application layer gateway relaying connections from (internal) private IP address space to (external) public IP address space, and vice-versa.

The presence of NAT hinders use of many applications, because the lack of a globally routable IP address for a host in the NATed site means that external hosts cannot communicate directly with it; some kind of proxy has to be used (and pre-configured). While some see this as a useful security feature, in practice NAT breaks the generally well-recognised end-to-end principle of the Internet (RFC1958, RFC2775).

Source URL: https://community.jisc.ac.uk/library/videoconferencing-booking-service/firewalls-and-proxies