



A Case Study of Changes Made to a Nokia/Check Point[®] Firewall-protected H.323 Gatekeeper/Proxy Topology at Wrexham County Borough Council

Tim Fullard (Wrexham CBC), Geoff Constable (Welsh Video Network)

Reference: GD/VTAS/026
Version: 1
Date: 03/03/2007

Table of Contents

1. Stakeholder Statements	5
2. Purpose and Summary	6
3. Introduction.....	6
4. Existing Topology	7
5. Possible Topology Configurations	8
6. Topology Changes	10
6.1 Address Spoofing.....	11
6.2 Nokia Cluster ARP	12
6.3 NAT	12
6.4 H.323 Protocol	12
6.5 Service Session Timeouts.....	13
6.6 Quality of Service	13
6.7 Firewall Rules.....	14
7. Summary of Steps	15
8. Conclusion.....	15

Background

This document describes the network topology and firewall configuration changes involved in placing a firewall traversal solution so that H.323 traffic is routed through a corporate firewall rather than around it. The firewall is a Nokia/Check Point® firewall cluster located at a local authority's connection with the regional network. The firewall traversal solution is the Cisco® Multimedia Conference Manager. The work described was carried out by Tim Fullard of Wrexham County Borough Council with the assistance of staff at the Welsh Video Network Support Centre. The authors would like to thank their colleagues and management at both organisations for their help in making and documenting these changes.

1. Stakeholder Statements

The WVN-recommended topology for deploying the Cisco® MCM (Multimedia Conference Manager) is in co-edged mode. The standard co-edged topology and configuration of the Cisco® MCM is covered in the VTAS document *Configuring an H.323 gatekeeper for use with the JANET Videoconferencing Service* (<http://www.video.ja.net/gkconfig.pdf>). This topology is described in more detail in section 4, **Existing Topology**.

For its own security management reasons, Wrexham CBC (County Borough Council) elected to work with WVN (Welsh Video Network) to relocate the MCM into a DMZ (De-Militarised Zone) behind its firewall. This topology and the issues arising are described in this report.

WVN would like to point out that this is not their recommended topology. This document does not imply that the co-edged deployment mode is any less secure than the topology identified within. The co-edged topology is a known solution which has been deployed successfully for more than five years and is detailed in the Cisco® documentation for this product. Where a Welsh UA (Unitary Authority) decides to implement an alternative scenario, the service can no longer be guaranteed by WVN and the investigation and resolution of problems that may be encountered will be the responsibility of the UA concerned.

Where non-WVN managed devices are placed between the MCM Gatekeeper/Proxy and the UA PoP router, the WVN demarcation point becomes the WVN interface on the UA PoP router. The onus will fall on the organisation managing any non-WVN managed device to show that their equipment is not at fault before arranging a WVN field service visit. In order to ensure that a consistent level of service is provided to all supported organisations, WVN may be forced to charge for site visits that are necessary as a result of changes made by another party or issues with non-WVN managed equipment. Any UA considering deploying this topology should consult the WVN Support Centre prior to making any changes.

Should issues or restrictions be identified with the topology or issues affecting service become apparent at any time, WVN reserves the right to request that the topology is restored to the standard recommended deployment. WVN will not make such a request unreasonably, without due cause, or without consultation with the UA. Any such changes will be formally requested and should be implemented within a reasonable timescale.

Should any organisation decide to deploy the topology described herein, Wrexham CBC wish to point out that security policies and topologies are the responsibility of the organisation concerned and it bears no responsibility for any changes the organisation may wish to make. This document outlines the changes Wrexham CBC made to reflect Wrexham CBC's own policies.

2. Purpose and Summary

As part of the CYDAG (Cymdeithas Ysgolion Dros Addysg Gymraeg) videoconferencing pilot for the Welsh Language Videoconferencing Project, participating schools were each supplied by Education and Learning Wales with Cisco® 3725 routers running MCM software. With the help and co-operation of the UAs involved, the MCMs were sited in WNL (Welsh Networking Limited) LLNW (Lifelong Learning Network for Wales) communications racks. These are typically located within the main communications room in the headquarters of each UA. The MCMs are under the administration and management (administrative domain) of the WVN, who also installed full videoconferencing studios at each of the schools. Through funding from DELLS (the Welsh Assembly Government Department of Education, Lifelong Learning and Skills), the WVN Support Centre has supported studio users and the gatekeepers since the initial installation in 2004.

The MCM is a border security device that provides a H.323 gatekeeper and a H.323 proxy server. Each MCM deployed at the UAs has the capability to manage numerous endpoints and simultaneous calls. This makes it logical to locate this device centrally as the gatekeeper for all videoconferencing endpoints within each UA.

This report outlines the changes made to the border security topology at Wrexham CBC in order to bring the MCM into the Council's preferred topology. This summary of the work involved is published as an aid to other UAs with similar equipment who may wish to configure a similar topology. The information herein is offered in good faith but with no guarantees, and Wrexham CBC and WVN accept no responsibility for the consequences of emulating the work described.

3. Introduction

Wrexham CBC currently has two Nokia IP series firewall enforcement platforms running Check Point Firewall NGX™ (R60) configured in a Nokia IPSO load-balanced format. This platform forms the main entry and exit point to the Council's network – all traffic passing through this point has an enforced policy implemented. While there are two physical devices, both act as a single entity; the policies are implemented on both devices and there is no real logical distinction between them. With this in mind this document will refer to these two devices as a single entity using the term 'the firewall'.

Videoconferencing uses the H.323 protocol to transport video and audio traffic over an IP network. H.323 is an umbrella protocol containing a multitude of other protocols for call setup, capability negotiation and exchange of video and audio information. The way that H.323 calls are set up results in a set of unique challenges for a network administrator. These are inherent within the H.323 protocol and cannot be altered.

In particular, any traditional IP protocol will place the IP addresses of the source and destination machines in the header of the IP packets. H.323 has packets embedded in packets which results in IP addresses being scattered within the payload of the overall IP packet. Normally within an organisation this does not present any problem,

but can do so where any NAT (Network Address Translation) is to take place. NAT translates a private set of addresses to public Internet addresses for communications with Internet hosts outside the organisation. Usually this translation will be done at the logical border of the organisation's network, where network management passes to some other organisation.

To overcome this potential problem, any device performing NAT will need to be fully H.323-aware and able to perform full stateful protocol inspection. Each packet has to be deconstructed and then reconstructed with the appropriate IP addresses throughout. This process must be carried out for packets both entering and leaving the network. The Cisco[®] MCM is a dedicated unit designed to be deployed as a border device to proxy H.323 calls, and as a result can be placed at the periphery of a network that utilises NAT. The Cisco[®] MCM at Wrexham CBC has been in place for some time, at the edge of the network, and in parallel to the firewall (co-edge mode). The firewall is placed at such a border and performs traditional IP NAT; both the co-edged design and the one described below have left the Check Point H.323 NAT capability untested.

Additionally, the H.323 protocol does not specify well-known ports for the exchange of packets containing audio and media (and associated control signalling), but instead defines a range of ports that can be used for this purpose. The endpoints concerned negotiate precisely which ports are opened on a per call basis during the negotiation and setup of any H.323 call. When a call begins, four dynamically-assigned ports are negotiated in each direction. Among others there are ports for audio, audio control, video, and video control. The protocol allows for dynamic allocation of 64,000 ports, which is an unacceptably high number of ports to be left open for inbound traffic. A firewall that is H.323 aware needs to be able to learn the ports that have been negotiated by inspecting the exchanged call setup information and then open the ports for the duration of the call and close them afterwards.

All real time traffic, including H.323 data, is particularly sensitive to packet loss, latency and jitter, so it is therefore imperative that these are kept to a minimum.

4. Existing Topology

Until Check Point NGX version R60, which was released after the initial implementation of this deployment at Wrexham CBC, Check Point Firewall-1[®] did not have the capability to perform H.323 NAT. Videoconferencing traffic does not stay within the UA network and its addressing scheme. To traverse external networks the private addresses used on the internal network need to be converted to public addresses (using NAT), so a topology was chosen that was tried and tested at FE and HE organisations in Wales. This topology is the standard topology normally supported by the WVN. The MCM was implemented in co-edged mode (see Figure 1) which enabled the IP service to the CYDAG school to be commissioned. A dedicated port for videoconferencing was made available on the WNL router, and an Internet IP subnet with addresses on the LLNW network used. These subnets are allocated and managed by WVN.

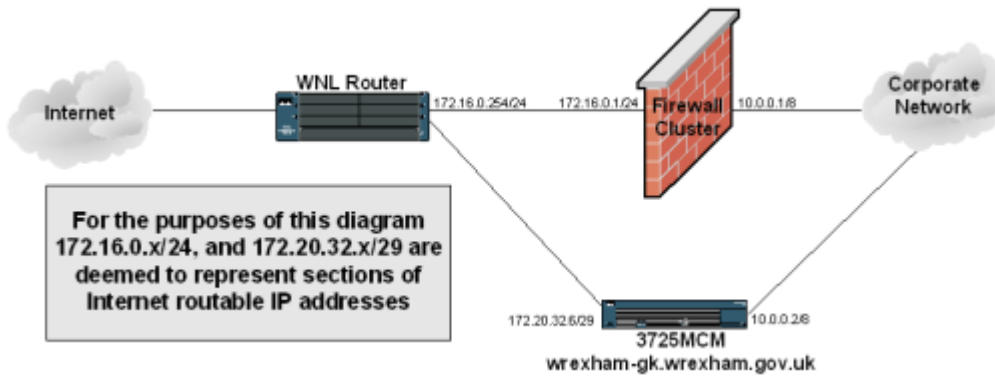


Figure 1: Co-edged Topology

While this topology avoids the problems associated with H.323 and NAT, traffic does not traverse the security solutions that the UA has implemented in order to protect the internal corporate network. The device is within the administrative domain of WVN and as such the ability to categorically define what traffic is allowed or not allowed via the MCM is beyond the control of UA personnel. Although functionally secure, this management issue is contrary to the Council security policy and for this reason the UA wished to test an alternative topology.

5. Possible Topology Configurations

The Council's organisational requirement is to provide a secure solution under the administrative control of the organisation that provides the same level of functionality as the co-edged design. Wrexham CBC considered that any change to topology should have as little impact as possible on the working practices of WVN and not create any obstacles to the support and management of the gatekeeper.

There are two possible topologies which ensure there is a single entry/exit point to the UA network. These are illustrated in Figures 2 and 3.




Figure 2: Option 1 Protected Gatekeeper Topology

Option 1 has the MCM with an internal interface connected to a DMZ on the firewall. The external interface has a public IP address allocated to it and is connected directly to an interface on the WNL router. While any traffic inbound to the corporate network is protected by the firewall, anything destined for the gatekeeper from the Internet does not have to traverse the same security solution.

While this would seem to preserve the single point of entry/exit into the UA network, it does not completely fulfil the Council's requirements. As outlined previously, the Cisco® MCM is under the administrative domain of WVN in line with installations at other UA sites in Wales. While this is the case under the current funding arrangements, there is the possibility that in the longer term support and administration could pass to the UA. If this does happen, then the above topology would again fail to adhere to the Council's security policy. Currently, the MCM device is owned by the Welsh Assembly Government and managed by WVN.




Figure 3: Option 2 Protected Gatekeeper Topology

The Council's preferred solution would allow the MCM to perform H.323-NAT by translating the internal IP addresses into public addresses, but still site the MCM behind the firewall. A certain amount of H.323 awareness is required from the firewall (to allow the dynamic opening and closing of ports used by H.323) but the solution illustrated in Figure 3 allows the gatekeeper to be incorporated on the edge of the network without avoiding the firewall itself. This meets all the functional requirements of the UA and adheres to its security policy.

The option 2 topology has the gatekeeper with two interfaces connected directly to the firewall, one on an internal DMZ using private IP addresses and another connected to an interface using a subnet of public IP addresses. To connect to the gatekeeper from either an external internet address or via the internal corporate network, traffic would now have to pass through the firewall.

Figure 4 illustrates further details of the chosen topology and is used as the basis for discussion for the rest of this document.

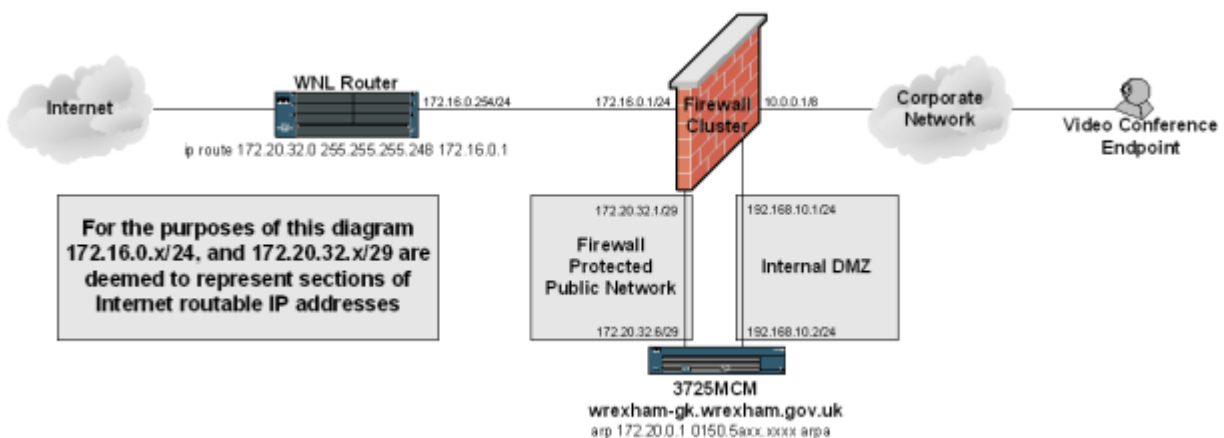


Figure 4: Protected Gatekeeper Topology

6. Topology Changes

In the new topology, the MCM has two separate connections to the firewall – one in the protected public IP address range (shown as Protected Public Network in Figure 4) and one in a firewall-protected DMZ (shown as Internal DMZ in Figure 4). The protected public address range deployed in this instance used IP addresses on the same WVN subnet as the existing MCM external connection. However, the model was also successfully tested using another public IP subnet.

When a videoconference is in progress the MCM acts as a gateway between IP networks. It appears to terminate the call to the local endpoint but relays messages and data out to the public network and the remote endpoint at the other end of the call. In this topology a datagram from the local endpoint with a media payload will be addressed to the IP address of the proxy – its external address. So packets sent to the conference will take the following route:

- the videoconference endpoint connected to the UA network will traverse the internal LAN and be routed towards the firewall on its internal port
- the firewall passes the H.323 packet out of the Internal DMZ interface to the internal address of the MCM
- the MCM will then pass the packet internally to the external IP port (this port has an external, public address and is part of the Protected Public Network – the packet effectively has NAT performed here)
- the packet is passed to the firewall and out of the same exit port that all traffic from the UA network uses to reach the Internet
- the packet then passes to the same port on the WNL router that is used for all traffic entering and exiting.

For the Council, the benefit of this topology is that the MCM is protected and has enforced policies defined by its own personnel. It is also able to use a public IP address for the H.323 proxy, performing the NAT itself and avoiding any issues associated with the firewall performing this task. Any standard IP traffic (i.e. non-H.323) bound for the internal corporate network passes through the firewall and so is subject to the local policy, and the same can now be said for any H.323 traffic.

This topology depends on the fact that routers use the most specific and applicable route in their routing tables to forward each IP datagram. This allows for a section of IP addresses to be routed via the firewall which will itself have an interface configured to connect to this external address range using a public IP address.

For the described topology to be successfully implemented the following factors need to be considered.

6.1 Address Spoofing

The Check Point firewall performs checks on traffic that reaches any of its interfaces to ensure that it is arriving from an expected source address. For example, traffic using a public IP address should only be received on an interface that is connected externally (not the internal LAN); similarly, packets with private addresses should be seen from interfaces connected internally.

With the topology illustrated in Figure 4, the MCM has two IP addresses, one of which is a public address that is visible to the Internet at large. As mentioned above, the H.323 proxy is bound to this public address so that it may perform NAT translation on any traffic going out towards the Internet.

When an endpoint registers with the gatekeeper it may initially communicate with the internal IP address but the H.323 registration will occur with the H.323 proxy using the external IP address. Any reply traffic from the MCM may actually come from this public address via the internal interface of the gatekeeper as this is the most direct route to the internal network. Figure 5 illustrates a simplified registration traffic flow.




Figure 5: Endpoint Registration Traffic Flow

When this occurs, the firewall will receive traffic on an internal interface that is from an external IP address and under normal circumstances, it will drop the packet. The Check Point firewall can be configured to expect traffic on the internal interface that has a source from either the internal DMZ subnet or the public IP address of the gatekeeper.

However, the firewall would not normally expect to receive packets from the same IP source address on both an internal and an external interface. Following the change outlined above, the firewall does not expect traffic on the external Protected Public Network interface from the MCM's external public IP. To overcome this, the firewall can be configured to perform the anti-spoofing checks as normal on that interface, ignoring source addresses for the subnet that is immediately connected to it (and hence the MCM's public IP Address). These rules are summarised in Table 1 (which refers to the topology and terminology outlined in Figure 4).

Firewall Interface	Expected Traffic
Internal DMZ	DMZ Subnet + MCM public IP address
Protected Public Network	Any external source + don't check anything from Protected Public Network (which includes the MCM public IP address)

Table 1. Anti-spoofing configuration for the new topology

6.2 Nokia Cluster ARP

As outlined previously, the Check Point solution is based around Nokia IPSO clustered enforcement modules. The clustering works by using addressing assigned to a virtual interface shared between all of the active enforcement modules. There seems to be a specific issue with Cisco[®] devices and their ability to correctly detect the IPSO cluster network address. It may be necessary to configure a static ARP entry on the Cisco[®] MCM relating to the Nokia cluster MAC and IP addresses. The Cisco[®] IOS Command takes the format:

```
arp [Firewall cluster IP address] [Firewall cluster MAC address] arpa
```

which in Figure 4 is shown as:

```
arp 172.20.0.1 0150.5axx.xxxx arpa
```

6.3 NAT

In the new topology, the MCM retains the functionality that it already had from the co-edged design with regards to NAT. The firewall should be configured not to perform NAT on any traffic coming from the internal network to the gatekeeper as there would be no benefit in doing so. The firewall would normally expect to perform NAT for any traffic originating inside the network and going to a public IP address, even if the device is already protected via the firewall. However, in this instance the firewall should be configured not to perform any NAT, as this is handled by the MCM. In this topology the firewall merely acts as a router, albeit with enhanced logging and protocol inspection.

6.4 H.323 Protocol

As discussed earlier, the H.323 protocol presents unique challenges when it comes to allowing traffic to pass over a firewall. Generally, IP protocols define well-known ports for particular services. For example, a web server listens for requests on port 80 (ports in this sense refer to logical or virtual ports rather than physical hardware). In relation to videoconferencing traffic, specific well-known ports will be used for certain elements of the flow: i.e. for the endpoint registration to the gatekeeper. However, as mentioned in section 3, **Introduction**, H.323 voice and video traffic will use ports that have been dynamically allocated during the call setup.

This problem is further exacerbated by the fact that most of this traffic (and all of the media and media control messages) uses UDP (rather than TCP). While this allows the TCP handshake overhead to be avoided, it means the firewall has a harder job to keep track of what flow is allowed or not, as the firewall uses TCP handshakes to follow the conversation between the devices effectively.

A level of H.323 awareness is thus required from the firewall to allow this to be handled efficiently by opening and closing ports dynamically, on demand. The alternative is to create firewall rules that will open up a wide range of ports on a permanent basis, which Wrexham CBC decided was an unacceptable potential compromise to security.

6.5 Service Session Timeouts

With regards to stateful inspection, the Check Point firewall will keep a record of the connections it has allowed to pass to ensure that any subsequent connection decisions are processed quicker. This also affects the returning traffic and leads to an element of statefulness for stateless protocols (such as UDP). By default, the Check Point device is configured to timeout TCP sessions after 3600 seconds (one hour) and a UDP virtual session after 40 seconds.

H.323 uses both TCP and UDP for both Call Setup / Signalling and Voice / Video feeds. UDP is used for the video and audio which are continuously sent so avoiding the UDP session timeout. However, if a videoconference should last longer than an hour these default settings may cause a premature disconnection of the conference.

Altering the setting in the Global Properties may not be desirable as these new settings would affect all of the services being used across the firewall. However, it is possible to alter the configuration of the defined H.323 services to increase these timers and virtually eliminate the problem. The change to these timers would then only affect the H.323 services. In this case, the H.323 services were altered to allow the firewall to remember sessions for 14,400 seconds (TCP) and 80 seconds (UDP).

6.6 Quality of Service

Voice and video traffic is highly susceptible to packet loss, latency and jitter. The above design does result in all traffic having to traverse an additional device and all possible steps should be taken to eliminate any cause for delay. The assumption here is that all devices in the H.323 path have been set to full speed and full duplex (and are not auto-sensing, which is known to cause issues that severely impact the quality of H.323 sessions).

It should also be noted that in the original configuration there were 2 x 100Mbit/s connections to the UA PoP router: one of these handled non-videoconferencing traffic to and from the firewall, the other handled videoconferencing traffic to and from the gatekeeper. Now, of course, there is only one connection handling all traffic, aggregated onto one 100Mbit/s link. Another effect of the changes to the original topology is that should the firewall go out of service (or run slowly) then videoconferencing is also affected, whereas before it would have continued undisturbed. Wrexham CBC's link to the LLNW has sufficient capacity for the aggregation of traffic not to present any quality issues, and there is built-in resilience within the firewall cluster. However, these issues do need to be considered if any changes to topology are being proposed.

The Check Point firewall, using Floodgate, will allow you to guarantee bandwidth for specific traffic types. Any other intermediary device should be configured to prioritise the videoconferencing traffic. While the new topology introduces extra stages in the path and inevitably adds some latency to each datagram's journey, in practice (thus far) the increased latency is not enough to have any appreciable effect on the call quality. While the new topology has been in place for testing, four endpoints (three schools and a council location) have all passed JVCS (JANET Videoconferencing Service) Quality Assurance tests, which should identify an unacceptable increase in latency. The topology is currently working successfully with a limited number of endpoints, but neither party can guarantee that as the number of simultaneous videoconferences or the general load on the firewall cluster increases that issues may not arise. In line with Wrexham CBC's standard practice, traffic levels of all

kinds will continue to be monitored so that any contention issues are identified and dealt with in good time.

6.7 Firewall Rules

Assuming that the firewall has sufficient H.323 awareness for the ports to be dynamically opened and closed, it should now be possible to generate policy rules to allow traffic flows to and from the MCM.

These rules can be broken down into two sets to cater for the external and internal requirements of the MCM. All JVCS calls are scheduled through the Booking Service and then dialled out from the central MCU (Multipoint Control Unit) infrastructure. This means that any JVCS conference originates from specific devices with known IP addresses. Non-JVCS *ad hoc* point-to-point calls are also possible, between the gatekeeper and individual external endpoints as illustrated in Figure 6.

Before firewall rules are created a policy decision needs to be made on whether or not to allow *ad hoc* point-to-point calls to take place. Preventing *ad hoc* calls would result in forcing point-to-point calls to be booked with JVCS (and as a result being routed through an MCU on JANET). To further enhance the security in place, Wrexham CBC made the policy decision to restrict external access from the gatekeeper to the known IP addresses of the JVCS MCUs. The firewall rules listed in Table 2 show that inbound connections to the MCM from the Internet have been restricted to WVN devices and the JVCS MCUs. WVN gatekeepers have been included to allow *ad hoc* test and support calls to be made.

This ensures that any point-to-point or multipoint calls are arranged through the JVCS Booking Service. As a result, external accessibility is required for H.323 traffic only to and from specific locations, reducing the possibility of unauthorised connections to the MCM being established (and unauthorised videoconferences from taking place).

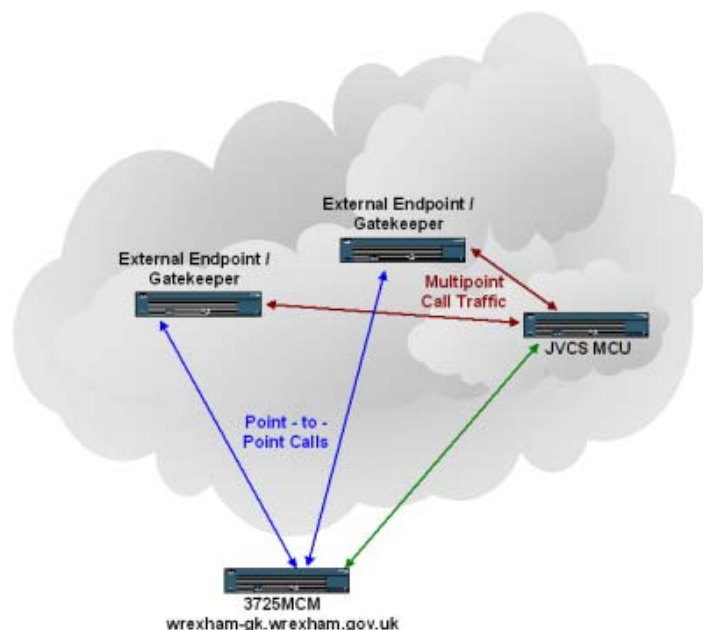


Figure 6: Point-to-point and JVCS H.323 Traffic Flows

When considering the nature of the internal traffic requirement, the situation is just as simple. Any internal endpoint will need to communicate with both the internal and external IP addresses of the gatekeeper for H.323. As the endpoints are known, the rules are configured to specify these.

The resulting rules are summarised in Table 2.

Source	Destination	Traffic
JANET MCUs, WVN	Gatekeeper External Address	H.323, etc ¹
Gatekeeper External Address	JANET MCUs, WVN	H.323
Internal endpoints	Gatekeeper (internal / external address)	H.323
Gatekeeper (internal / external address)	Internal endpoints	H.323

Table 2. Firewall rules for the new topology

7. Summary of Steps

The steps necessary to make the changes described in this document can be summarised as follows:

- Connect protected Internet address range via firewall
- Configure anti-spoofing
- Disable NAT
- Alter H.323 service session timeouts
- Ensure QoS implemented
- Generate firewall rules.

8. Conclusion

Because of different organisational requirements, WVN and Wrexham CBC had what appeared to be mutually conflicting requirements for the topology of the traversal of H.323 packets across the border of the UA and LLNW networks. However, through mutual trust, co-operation and gradual development, and by capitalising on developments in functionality of the firewall product in place, a solution has been found and implemented which maintains the functionality and IP address of the MCM, while bringing it within the security policy deployed by this Council.

1. The precise list of protocols and ports that are needed is available from the WVN Support Centre.

Copyright:

This document is copyright The JNT Association. Parts of it, as appropriate, may be freely copied and incorporated unaltered into another document unless produced for commercial gain, subject to the source being appropriately acknowledged and the copyright preserved. The reproduction of logos without permission is expressly forbidden. Permission should be sought from the JANET Service Desk.

Trademarks:

JANET[®] is a registered trademark of the Higher Education Funding Councils for England, Scotland and Wales. The JNT Association is the registered user of this trademark.

Check Point[®] and FireWall-1[®] are registered trademarks and NGX[™] is a trademark of Check Point Software Technologies Ltd. or its affiliates.

Cisco[®] and IOS are registered trademarks of Cisco Systems, Inc. and/or its affiliates in the U.S. and certain other countries.

Nokia and IPSO are trademarks or registered trademarks of Nokia Corporation.

Disclaimer:

The information contained herein is believed to be correct at the time of issue, but no liability can be accepted for any inaccuracies.

The reader is reminded that changes may have taken place since issue, particularly in rapidly changing areas such as internet addressing, and consequently URLs and e-mail addresses should be used with caution.

The JNT Association cannot accept any responsibility for any loss or damage resulting from the use of the material contained herein.



© The JNT Association 2007

003(07/07)

