Network Address Translation (NAT)

NAT should be familiar to network managers – it is widely deployed in large private networks. NAT is described fully in RFC1918 “Address Allocation for Private Internets”. NAT was introduced partly as a means of conserving real or public (also sometimes called 'routable') IP addresses. The deployment of NAT allows large organisations to give every computer a unique Internet address without diminishing the available pool of public IP addresses. It does this by defining a set of addresses that should not be used on the public Internet and should only be used within the private network. Thus these addresses are 'unroutable'. Using NAT in a network also has the potential to add a layer of security – if the addresses within an organisation are not routable from the Internet it may be harder to 'attack' them. By definition, there must be some kind of translation machine between the inner and outer network. Whether this actually helps security is open to debate, but the practice of deploying NAT'd private networks has certainly helped conserve Internet addresses.

NAT is usually overcome by deploying a NAT server at the network boundary, which maintains mappings between private (NAT) addresses and public IP addresses. These mappings may be static (i.e. permanent) or dynamic (i.e. ad hoc and temporary). The problem that NAT presents to a deployment of H.323 is due to the fact that both H.225 messages and H.245 call setup messages bury their network (IP) address deep in the data payload of the IP packet. This payload is examined by the equipment at the other end and the source address within is used as the return address. If the return address is one that is behind a NAT boundary then the packet will never reach its destination, as NAT addresses are unroutable in the public Internet and the call will eventually time-out and fail. A 'naïve' NAT server will only change the addresses in the UDP or IP datagram header and footer, and not alter anything deeper in the protocol stack or in the data payload itself. This explains why H.323 works perfectly well when the two endpoints are within the same network, even when that network is using private (NAT) addresses. As long as the two endpoints have a route to each other then the call will succeed, as the packets exchanged never go beyond their particular NAT domain and so no translation is effected on the IP addresses used.