Network Address Translation (NAT) Tutorial and Avaya™ Communication Manager 1.3 NAT Shuffling Feature

ABSTRACT

This document is a simplified tutorial on network address translation (NAT) and network address port translation (NAPT), or port address translation (PAT). This document also explains in detail the new NAT shuffling feature in Avaya™ Communication Manager 1.3.

Application Note

April 2003

COMPAS ID 97779

All information in this document is subject to change without notice. Although the information is believed to be accurate, it is provided without guarantee of complete accuracy and without warranty of any kind. It is the user's responsibility to verify and test all information in this document. Avaya shall not be liable for any adverse outcomes resulting from the application of this document; the user must take full responsibility.

NAT Tutorial and
Avaya™ Communication Manager 1.3 NAT Shuffling Feature

Table of Contents

1 NAT and NAPT/PAT

Network address translation (NAT) is a function, typically in a router or firewall, by which an internal IP address is translated to an external IP address. The terms "internal" and "external" are generic and ambiguous, and they are more specifically defined by the application. For example, the most common NAT application is to facilitate communication from hosts on private networks to hosts on the public Internet. In this case the internal addresses are private addresses, and the external addresses are public addresses.

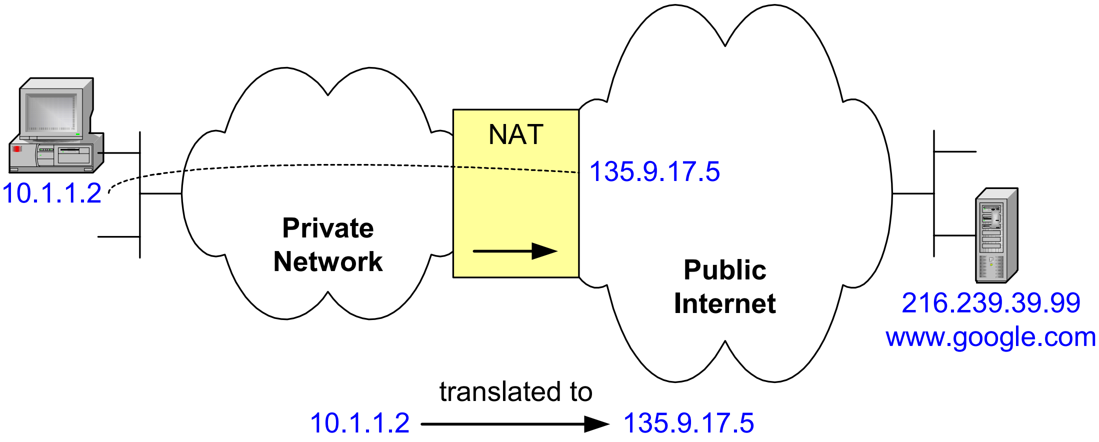

The figure above shows the private address 10.1.1.2 being translated to the public address 135.9.17.5. When the private network host makes a request to the public web server www.google.com, the request appears to the server to be coming from 135.9.17.5. The server replies to 135.9.17.5, and the NAT device (router or firewall) reverses the translation and forwards the reply to 10.1.1.2. *Note that this scenario does not utilize a web proxy server, which would be an entirely different scenario.*

NAT looks like this at the IP packet level.

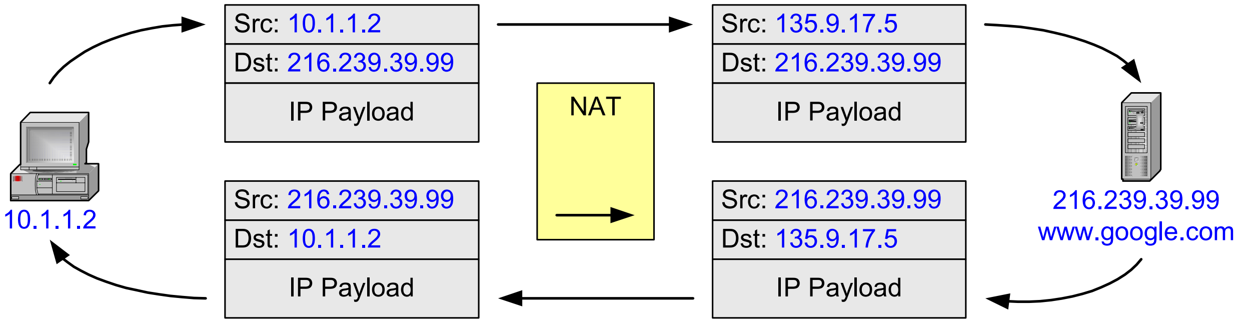

Another common NAT application is for some VPN clients. The internal address in this case is the physical address, and the external address is the virtual address.

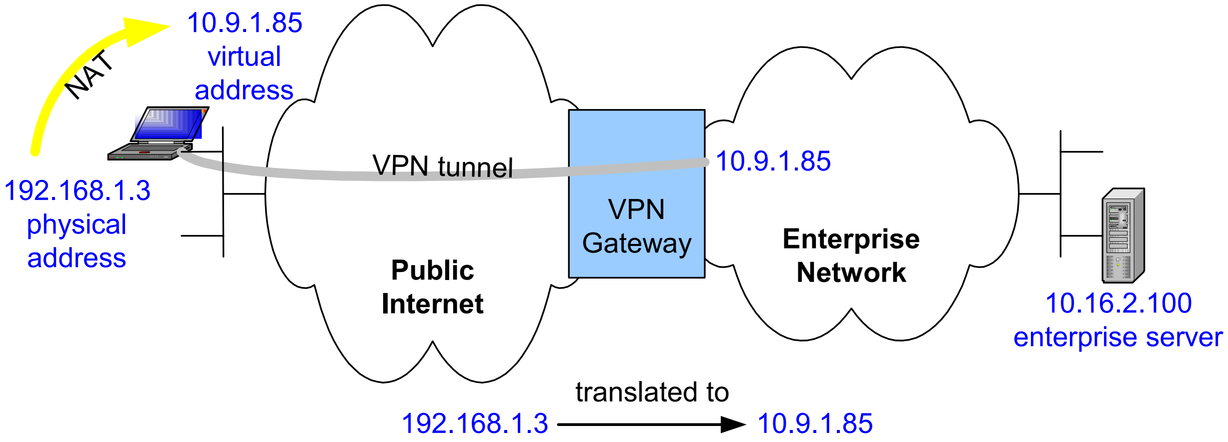

In the figure above the laptop has a physical address of 192.168.1.3. This could be the address given by the broadband (DSL, cable, etc.) service provider, or it could be a private small-office/home-office (SOHO) network address. This address does not necessarily have to be a private address as shown here, as the subscriber could pay for a public address from the broadband service provider. But regardless of the nature of the physical address, the point is that it cannot be used to communicate back to the enterprise through a VPN tunnel. Once the tunnel is established, the enterprise VPN gateway assigns a virtual address to the VPN client application on the laptop. This virtual address is part of the enterprise IP address space, and it must be used to communicate back to the enterprise.

The application of the virtual address varies among VPN clients. Some VPN clients integrate with the operating system in such a way that packets from IP applications (ie, FTP, telnet, etc.) on the laptop are sourced from the virtual IP address. That is, the IP applications inherently use the virtual IP address. With other VPN clients this does not occur. Instead, the IP applications on the laptop inherently use the physical IP address, and the VPN client performs a NAT to the virtual IP address. In this case the VPN client translates the physical address 192.168.1.3 to the virtual address 10.9.1.85.

This NAT is no different than if a router or firewall had done the translation. All requests coming from the laptop appear to the enterprise server to be coming from 10.9.1.85. The server replies to 10.9.1.85, and the VPN gateway forwards the replies through the tunnel to the VPN client, which then translates the destination address back to 192.168.1.3.

There are three main types of NAT, and each one is covered below.

1.1 Static 1-to-1 NAT

Static 1-to-1 NAT is what has already been covered up to this point. For every internal address there is an external address, with a static 1-to-1 mapping between internal and external addresses. It is the simplest yet least efficient type of NAT, in terms of address preservation, because every internal host requires an external IP address. This limitation is often impractical when the external addresses are public IP addresses. Sometimes the primary reason for using NAT is to preserve public IP addresses, and for this case there are two other types of NAT: many-to-1 and many-to-a-pool.

1.2 Dynamic Many-to-1 NAT

Dynamic many-to-1 NAT is as the name implies. Many internal addresses are dynamically translated to a single external address.

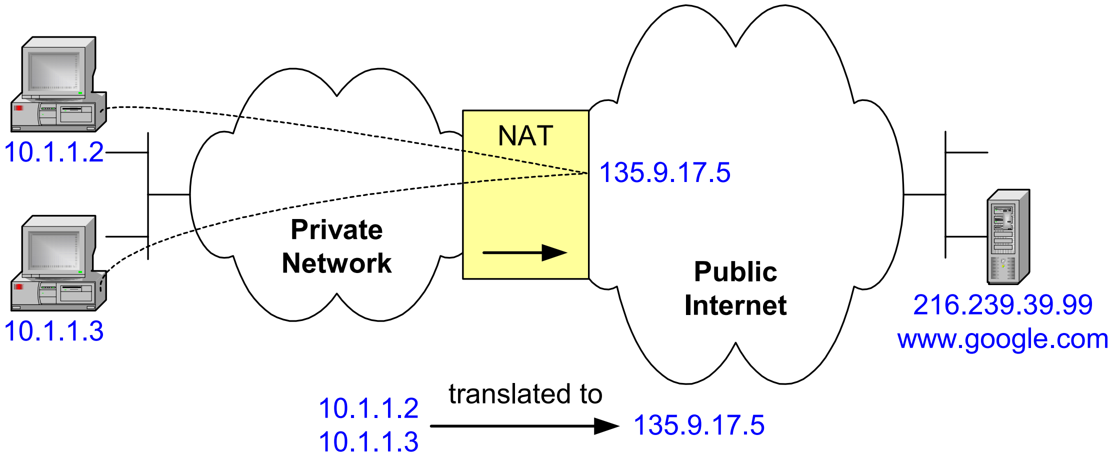

Multiple internal addresses can be translated to the same external address when the TCP/UDP ports are also translated, in addition to the IP addresses. This is known as network address port translation (NAPT) or simply port address translation (PAT). Now it appears to the external server that multiple requests are coming from 135.9.17.5 from different TCP/UDP ports. The NAT device remembers which internal source ports were translated to which external source ports.

NAPT looks like this at the IP packet level.

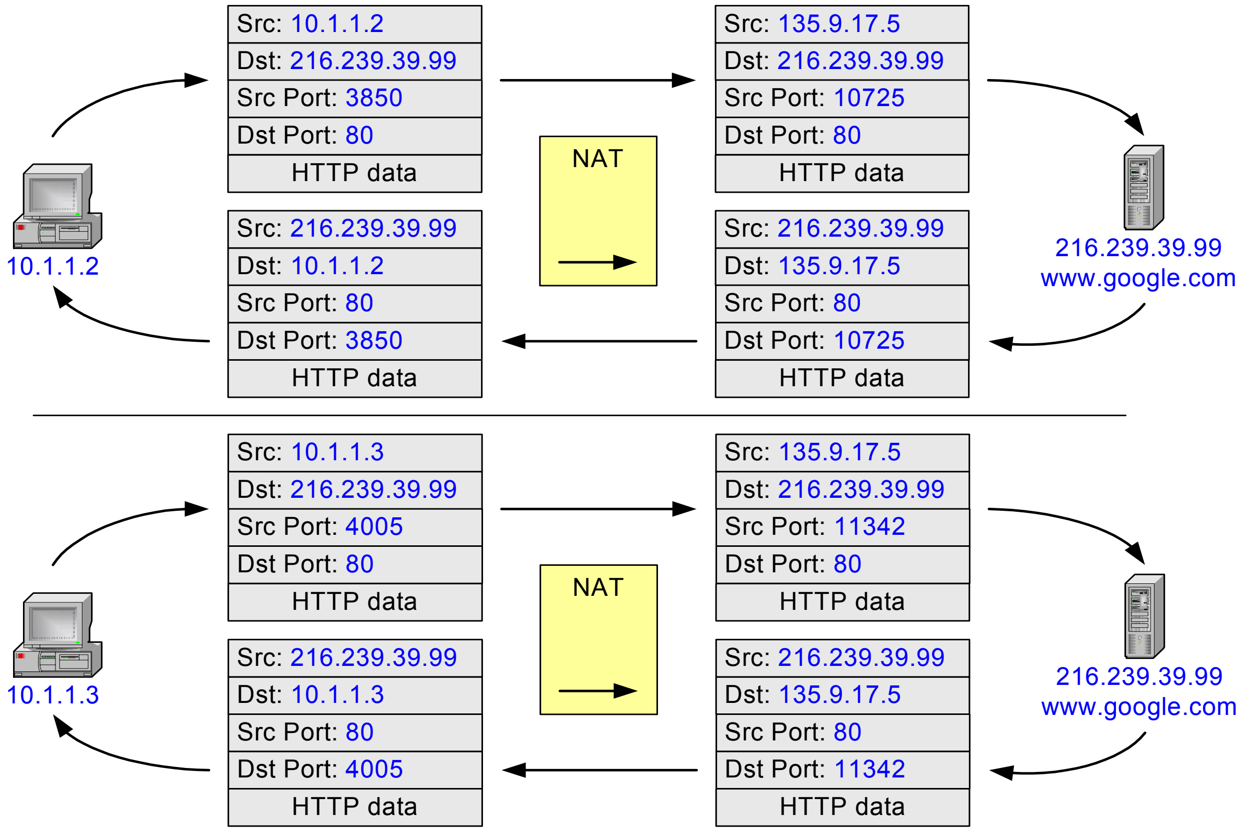

Src: 10.1.1.2
Dst: 216.239.39.99
Src Port: 3850
Dst Port: 80
HTTP data

Src: 135.9.17.5
Dst: 216.239.39.99
Src Port: 10725
Dst Port: 80
HTTP data

Src: 216.239.39.99
Dst: 10.1.1.2
Src Port: 80
Dst Port: 3850
HTTP data

Src: 216.239.39.99
Dst: 135.9.17.5
Src Port: 80
Dst Port: 10725
HTTP data

Src: 10.1.1.3
Dst: 216.239.39.99
Src Port: 4005
Dst Port: 80
HTTP data

Src: 135.9.17.5
Dst: 216.239.39.99
Src Port: 11342
Dst Port: 80
HTTP data

Src: 216.239.39.99
Dst: 10.1.1.3
Src Port: 80
Dst Port: 4005
HTTP data

Src: 216.239.39.99
Dst: 135.9.17.5
Src Port: 80
Dst Port: 11342
HTTP data

Many-to-1 NAT is commonly used for SOHO networks, and it is also used to some extent for enterprise networks. This is primarily because many internal hosts can use a single external address, which has to be purchased from the service provider, to access the Internet.

In the simplest form of many-to-1 NAT the internal host must initiate the communication to the external host, which then generates a port mapping within the NAT device, allowing the external host to reply back to the internal host. It is a paradox with this type of NAT (in its simples form) that the external host cannot generate a port mapping to initiate the communication with the internal host, and without initiating the communication there is no way to generate the port mapping. This condition does not exist with 1-to-1 NAT, as there is no mapping of ports.

Much can be said of the advantages and disadvantages of 1-to-1 and many-to-1 NAT, and there are many methods to nullify the advantages and compensate for the disadvantages of either. These topics are beyond the scope of this document.

1.3 Dynamic Many-to-a-Pool NAT

Many-to-a-pool NAT combines some of the characteristics of both 1-to-1 and many-to-1 NAT. The general idea behind many-to-a-pool NAT is that a 1-to-1 mapping is not desired, but there are too many internal hosts to use a single external address. Therefore, a pool of multiple external addresses is used for NAT. There are enough external addresses in the pool to support all the internal hosts, but not nearly as many pool addresses as there are internal hosts.

Depending on the router or firewall manufacturer, the actual functionality of this type of NAT can vary with regard to NAPT. For example, the NAT engine might employ NAPT always, or it might employ NAPT only after the entire pool of addresses is used at least once. Without some crafty configurations, which are beyond the scope of this document, NAPT must be employed when an external address is used by more than one internal host.

1.4 Issues with NAT and H.323

Here are some of the hurdles that NAT presents to H.323.

- H.323 messages, which are part of the IP payload, have embedded IP addresses in them. NAT translates the IP address in the IP header, but not the embedded addresses in the H.323 messages. This is a problem that can be and has been addressed with H.323-aware NAT devices. It has also been addressed with Avaya Communication Manager 1.3.

- When an endpoint (IP telephone) registers with the gatekeeper (call server), that endpoint's IP address must stay the same for the duration of the registration. This rules out almost all current implementations of many-to-a-pool NAT.

- TCP/UDP ports are involved in all aspects of IP telephony – endpoint registration, call signaling, and RTP audio transmission. These ports must remain unchanged for the duration of an event – duration of the registration or duration of a call. Also, the gatekeeper must know ahead of time which ports will be used by the endpoints for audio transmission, and the ports can vary per call. These requirements make it very difficult for H.323 to work with NAPT, which rules out almost all current implementations of many-to-1 and many-to-a-pool NAT.

2 Avaya Communication Manager 1.3 NAT Shuffling Feature

The Avaya Communication Manager 1.3 NAT Shuffling feature permits IP telephones and Softphones to work behind a NAT device. This feature was available prior to ACM 1.3, but it did not work with shuffled calls (Direct IP-IP Audio enabled). The ACM 1.3 feature now works with shuffled calls.

Terms: These terms are used to describe the NAT Shuffling feature.

- Native Address – The original IP address configured on the device itself. (internal address)
- Translated Address – The IP address after it has gone through NAT, as seen by devices on the other side of the translation. (external address)
- Gatekeeper – The Avaya device that is handling call signaling. It could be a portal to the gatekeeper, such as a C-LAN, or the gatekeeper itself, such as an S8300 server.
- Gateway – The Avaya device that is handling media conversion between TDM and IP, such as a MedPro board or G700 VoIP media module.

The essence of this feature is that the Avaya Communication Manager keeps track of the native and translated IP addresses for every IP station (IP telephone or Softphone). If an IP station registration appears with different addresses in the IP header and the RAS message, the call server stores the two addresses and alerts the station that NAT has taken place.

This feature works with static 1-to-1 NAT. It does not work with NAPT, so the TCP/UDP ports sourced by the IP stations must not be changed. Consequently, this feature does not work with many-to-1 NAT. This feature may work with many-to-a-pool NAT if a station's translated address remains constant for as long as the station is registered, and there is no port translation.

The NAT device must perform plain NAT – not H.323-aware NAT. Any H.323-aware feature in the NAT device must be disabled, so that there are not two independent devices trying to compensate for H.323 at the same time.

Rules: These rules govern the NAT Shuffling feature. The Direct IP-IP Audio parameters are configured on the SAT **ip-network-region** form.

- Rule1 – When Direct IP-IP Audio is enabled (default) and a NATed station and a non-NATed station talk to one another, the translated address is always used.
- Rule2 – When two NATed stations talk to one another, the native addresses are used (default) when **Yes** or **Native (NAT)** is specified for Direct IP-IP Audio, and the translated addresses are used when **Translated (NAT)** is specified.
- Rule3 – The Gatekeeper and Gateway must not be NATed. As long as this is true, they may be assigned to any network region.

Scenarios:
The following scenarios illustrate the NAT Shuffling feature. These are the tested and supported scenarios, except for scenario 2, which is given for explanatory purposes only. Scenario 1 is the preferred and recommended scenario because of its simplicity.

The diagrams should be interpreted as follows.

- Each letter in an IP address represents a different IP subnet, whether on a physical LAN segment or a logical VLAN. The numeral following the dot is the host address on that subnet.
- Each network cloud represents a physical or logical network separation. If two clouds are touching each other or connected via a NAT device, there is direct routing between them. Otherwise, they must route to each other via an intervening cloud(s).

Scenario 1

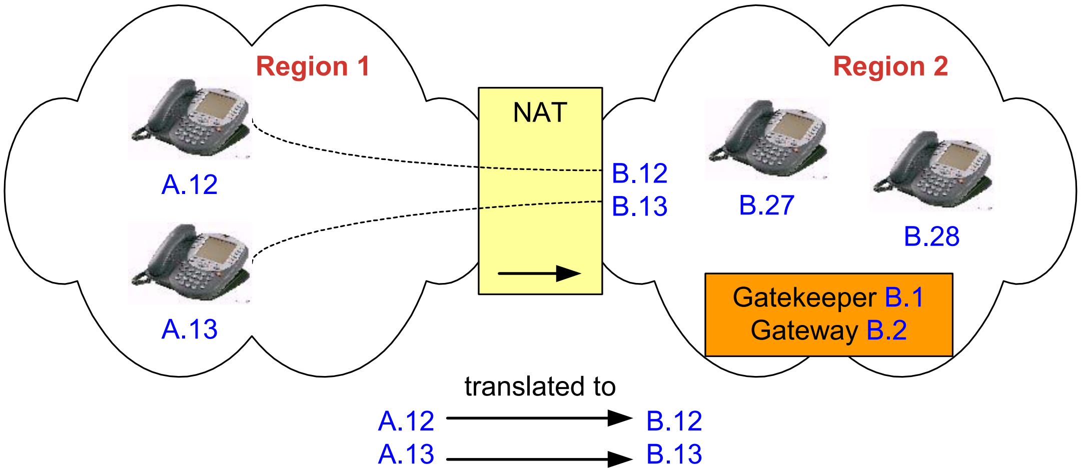

In this scenario the A.x devices talk directly to each other and to the B.x devices. The B.x devices talk directly to each other, but have no notion of A.x as an address space. The B.x devices talk to the A.x devices via the translated B.x addresses. Avaya Communication Manager knows that A.12 and A.13 are translated to B.12 and B.13.

Region 1 configuration:
Region 1 is behind a NAT device.

Intra-region Direct IP-IP: Yes
When A.12 and A.13 talk to each other, they direct their audio streams to each other's native addresses (Rule2). **Native (NAT)** could be explicitly configured.

Inter-region Direct IP-IP: Yes
B.27 and B.28 direct their audio streams to B.12 and B.13, not A.12 or A.13 (Rule1). **Translated (NAT)** could be explicitly specified.

Region 2 configuration:
Region 2 is not behind a NAT device.

Intra-region Direct IP-IP: Yes
Permit intra-region shuffling. No NAT instructions are required.

Inter-region Direct IP-IP: Yes
Permit inter-region shuffling. No NAT instructions are required for this region. They are already covered in the configuration for Region 1.

Note: This scenario shows all the Region 2 devices on the same IP subnet, but this is not a requirement. The Region 2 devices may be scattered across the corporate network on various subnets that can route to one another without NAT.

Scenario 2

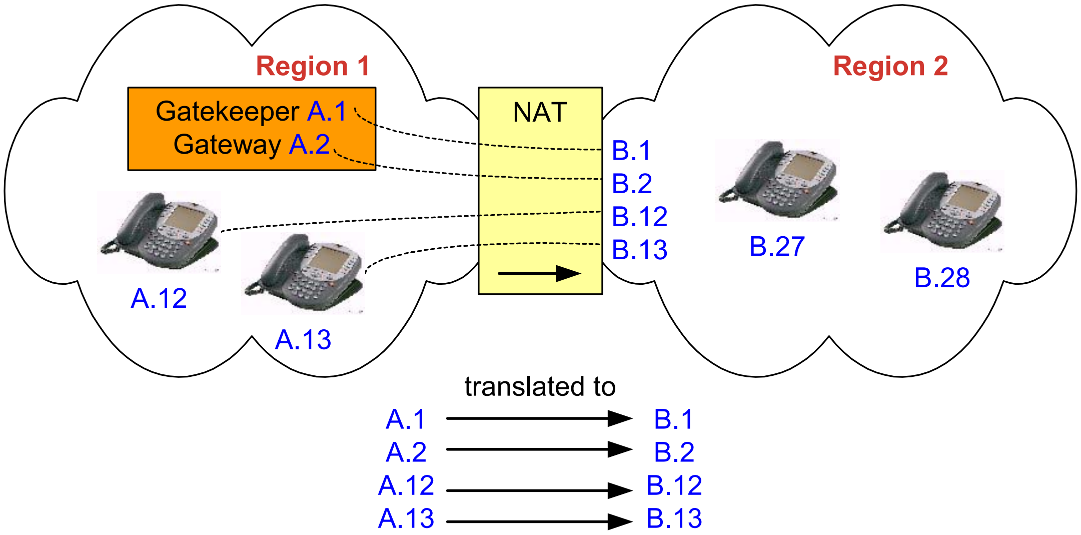

Same as scenario 1, except now ACM is on the other side of the NAT. This scenario can't work because ACM doesn't know what the translated addresses are. ACM will tell B.27 to send its audio stream to A.12, for example. But B.27 needs to be told to send its audio stream to B.12. Unless ACM knows the translations and has the intelligence, this can't be done. This solution is currently in the early stages of investigation; there are no estimates for completion.

Scenario 3a

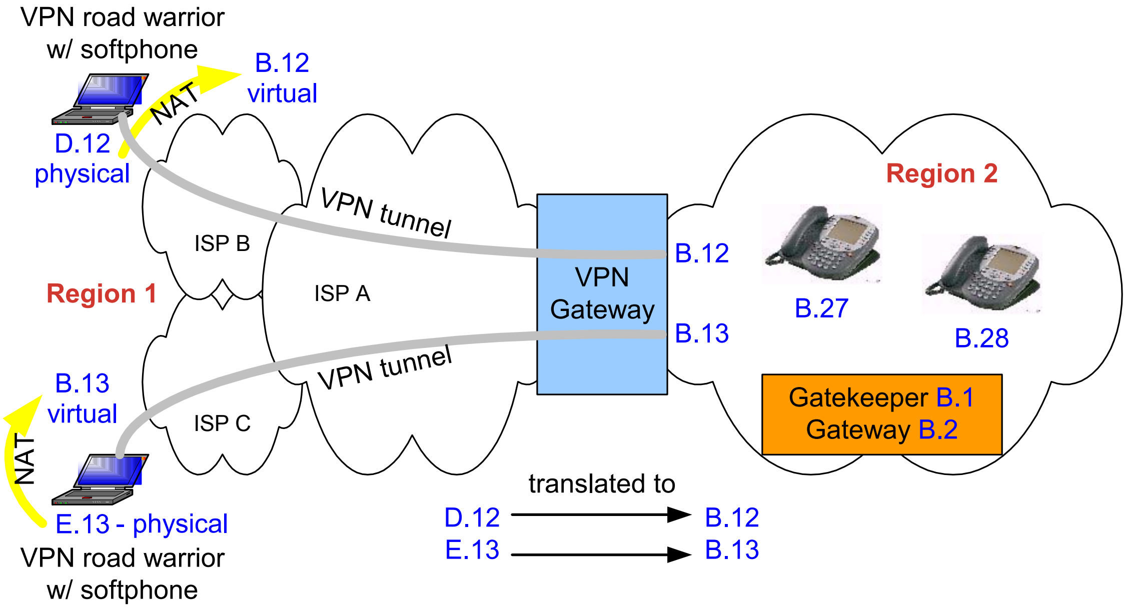

In this scenario the road warriors are accessing the corporate network, via VPN, through different Internet Service Providers (ISP). D.12 and E.13 can talk directly to the B.x devices, but they cannot talk directly to each other. One reason could be that D.12 and E.13 are private addresses, but there are other possible reasons. The road warriors cannot talk to each other using their translated addresses either, because they both terminate on the same VPN gateway. The B.x devices talk directly to each other, but talk to the road warriors via the translated B.x addresses. ACM knows that D.12 and E.13 are translated to B.12 and B.13.

Region 1 configuration:
Region 1 is behind a NAT device.

Intra-region Direct IP-IP: No
When D.12 and E.13 talk to each other, they must go through the Gateway.

Inter-region Direct IP-IP: Yes
B.27 and B.28 direct their audio streams to B.12 and B.13, not D.12 or E.13 (Rule1). **Translated (NAT)** could be explicitly configured.

Region 2 configuration:
Region 2 is not behind a NAT device.

Intra-region Direct IP-IP: Yes
Permit intra-region shuffling. No NAT instructions are required.

Inter-region Direct IP-IP: Yes
Permit inter-region shuffling. No NAT instructions are required for this region. They are already covered in the configuration for Region 1.

Note: This scenario shows all the Region 2 devices on the same IP subnet, but this is not a requirement. The Region 2 devices may be scattered across the corporate network on various subnets that can route to one another without NAT.

Scenario 3b

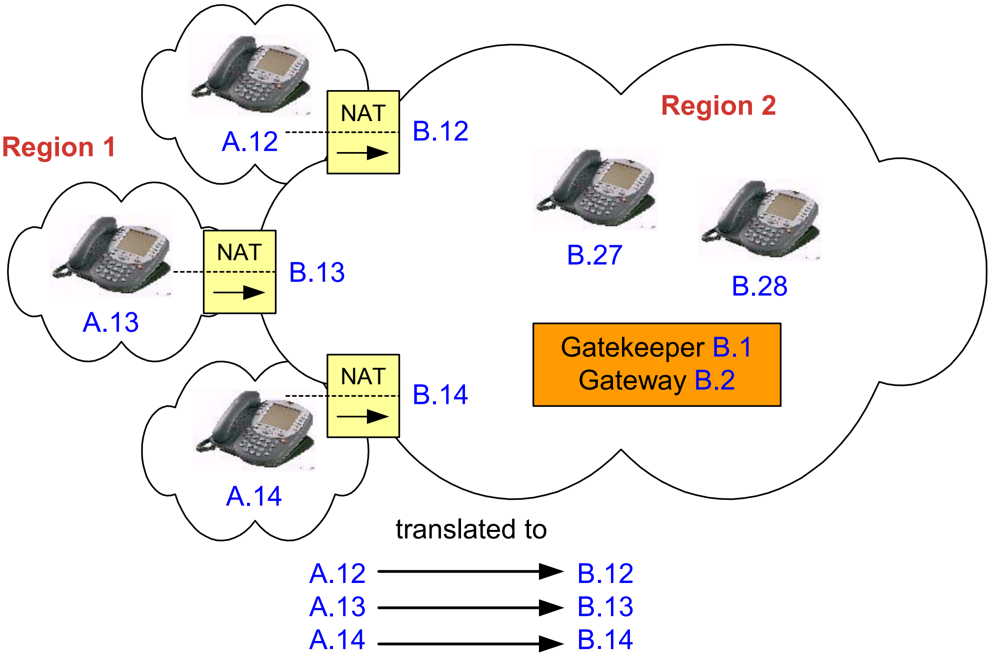

This scenario is similar to scenario 3a, with the main exception that each endpoint goes through a separate NAT device, as opposed to all endpoints terminating on the same VPN gateway. This is a rare scenario, but it can be accomplished. The A.x devices cannot talk directly to each other. They talk to each other using their translated addresses, and to the B.x devices directly. The B.x devices talk directly to each other, but to the A.x devices via the translated B.x addresses. ACM knows that A.12, 13, and 14 are translated to B.12, 13, and 14.

Region 1 configuration:
Region 1 is behind multiple NAT devices.

Intra-region Direct IP-IP: Translated (NAT)
When the A.x devices talk to each other, they direct their audio streams to each other's translated addresses.

Inter-region Direct IP-IP: Yes
B.27 and B.28 direct their audio streams to B.12, 13, and 14; not A.12, 13, or 14 (Rule1). **Translated (NAT)** could be explicitly configured.

Region 2 configuration:
Region 2 is not behind a NAT device.

Intra-region Direct IP-IP: Yes
Permit intra-region shuffling. No NAT instructions are required.

Inter-region Direct IP-IP: Yes
Permit inter-region shuffling. No NAT instructions are required for this region. They are already covered in the configuration for Region 1.

Note: This scenario shows all the Region 2 devices on the same IP subnet, but this is not a requirement. The Region 2 devices may be scattered across the corporate network on various subnets that can route to one another without NAT.

Scenario 3c

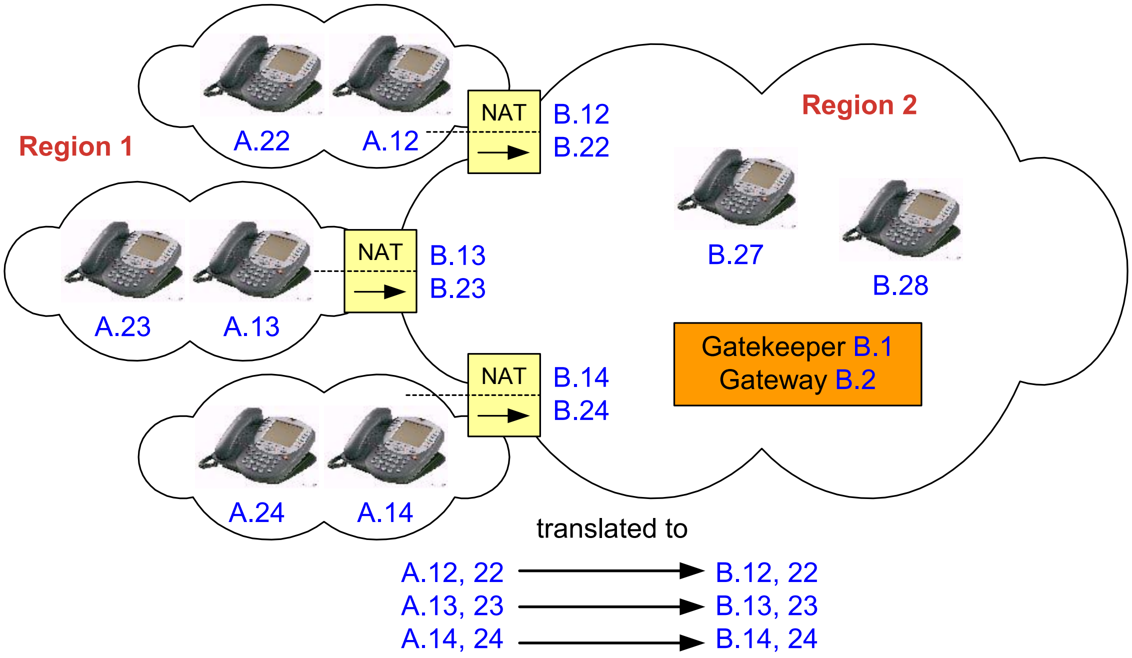

Same as scenario 3b, except now there are multiple endpoints behind each NAT device. Endpoints behind different NAT devices cannot talk directly to each other. The only way this scenario could work as illustrated is to NOT permit Intra-region Direct IP-IP in Region 1. This would cause all Region 1 endpoints to communicate through the Gateway when talking to one another.

An alternative is to create four separate regions instead of two, putting each set of enpdoints behind the same NAT device in its own region. Then scenario 1 would apply to each NATed region.

As illustrated with only two regions, ACM cannot know that endpoints talking to each other behind the same NAT device should use their native addresses, but endpoints talking to each other behind different NAT devices should use their translated addresses.

Scenario 4a

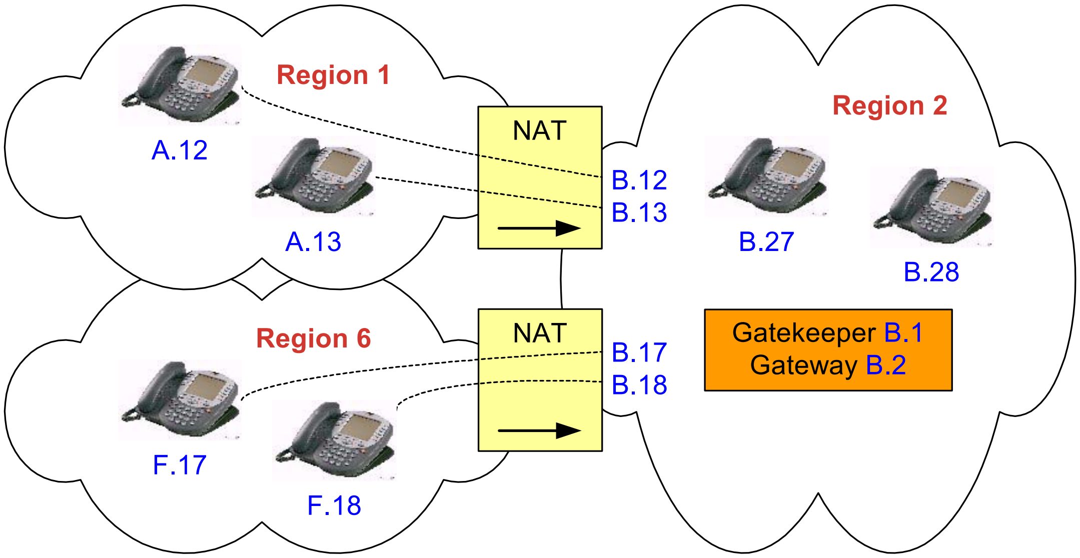

translated to

A.12 → B.12
A.13 → B.13
F.17 → B.17
F.18 → B.18

In this scenario the A.x devices talk directly to each other and to the B.x devices. The F.x devices also talk directly to each other and to the B.x devices. The A.x and F.x devices talk directly to each other. The B.x devices talk directly to each other, but have to use the translated B.x addresses to talk to A.x and F.x devices. ACM knows that A.x and F.x are NATed to their respective B.x addresses.

Region 1 configuration:
Region 1 is behind a NAT device.

Intra-region Direct IP-IP: Yes
When A.12 and A.13 talk to each other, they direct their audio streams to each other's native addresses (Rule2). **Native (NAT)** could be explicitly configured.

Inter-region Direct IP-IP: Yes
When A.x and F.x talk to one another, the native addresses are used (Rule2). When A.x and B.x talk to one another, B.x directs its audio stream to A.x's translated address (Rule1).

Region 6 configuration:
Region 6 is behind a NAT device.

Intra-region Direct IP-IP: Yes
When F.17 and F.18 talk to each other, they direct their audio streams to each other's native addresses (Rule2). **Native (NAT)** could be explicitly configured.

Inter-region Direct IP-IP: Yes
When A.x and F.x talk to one another, the native addresses are used (Rule 2). When F.x and B.x talk to one another, B.x directs its audio stream to F.x's translated address (Rule1).

Note: This scenario shows all the Region 2 devices on the same IP subnet, but this is not a requirement. The Region 2 devices may be scattered across the corporate network on various subnets that can route to one another without NAT.

Scenario 4b

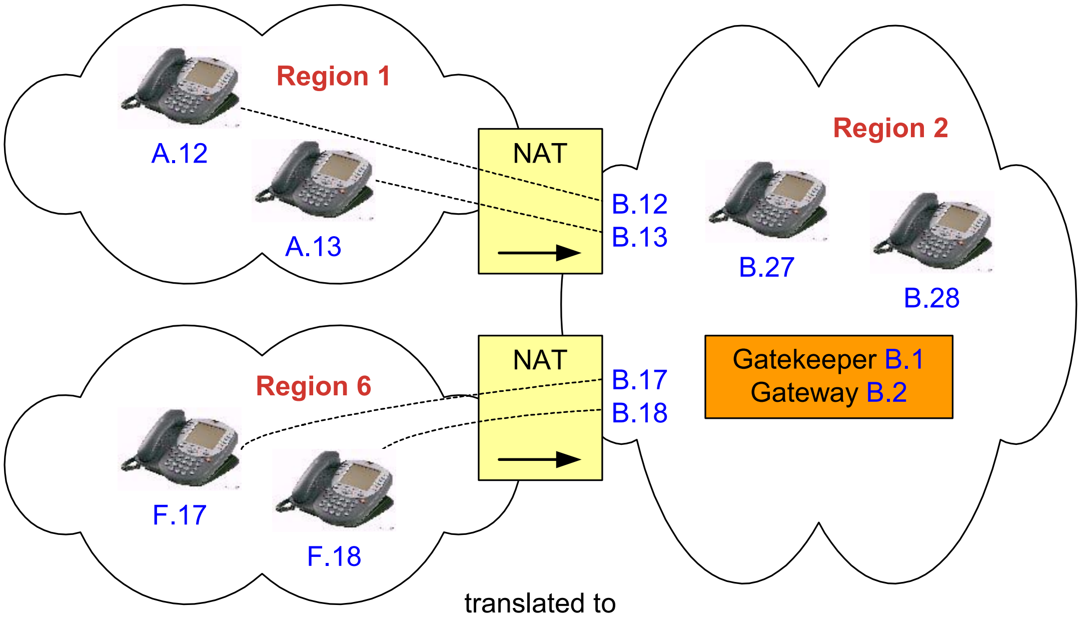

translated to

A.12 → B.12
A.13 → B.13
F.17 → B.17
F.18 → B.18

Same as scenario 4a, except regions 1 and 6 cannot talk directly to each other. In this scenario the A.x devices talk directly to each other and to the B.x devices. The F.x devices also talk directly to each other and to the B.x devices. The B.x devices talk directly to each other, but have to use the translated B.x addresses to talk to A.x and F.x devices. ACM knows that A.x and F.x are NATed to their respective B.x addresses.

Region 1 configuration:
Region 1 is behind a NAT device.

Intra-region Direct IP-IP: Yes
When A.12 and A.13 talk to each other, they direct their audio streams to each other's native addresses (Rule2). **Native (NAT)** could be explicitly configured.

Inter-region Direct IP-IP: Translated (NAT)
When A.x talks to B.x or F.x, the translated address(es) is used.

Region 6 configuration:
Region 6 is behind a NAT device.

Intra-region Direct IP-IP: Yes
When F.17 and F.18 talk to each other, they direct their audio streams to each other's native addresses (Rule2). **Native (NAT)** could be explicitly configured.

Inter-region Direct IP-IP: Translated (NAT)
When F.x talks to B.x or A.x, the translated address(es) is used.

Note: This scenario shows all the Region 2 devices on the same IP subnet, but this is not a requirement. The Region 2 devices may be scattered across the corporate network on various subnets that can route to one another without NAT.

Scenario 5a

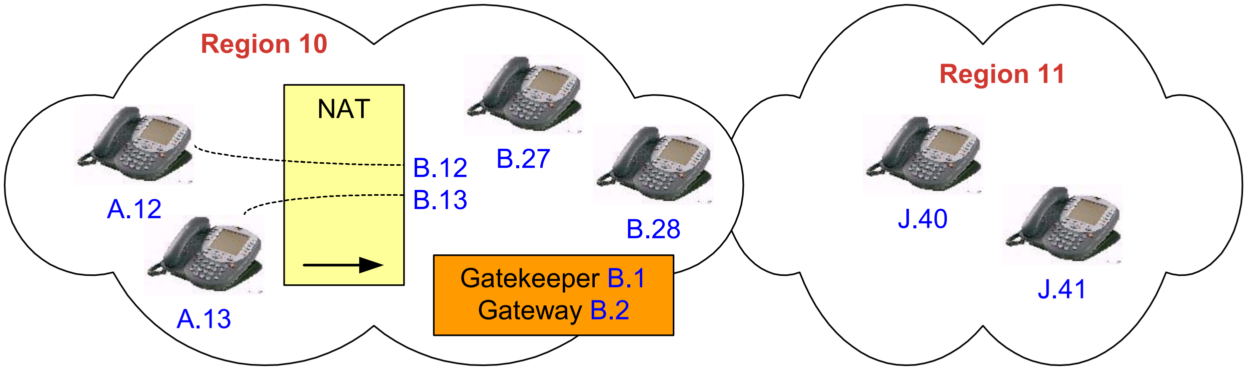

translated to

A.12 → B.12

A.13 → B.13

In this scenario the A.x devices talk directly to each other and to the B.x and J.x devices. The B.x and J.x devices talk directly to each other, but have no notion of A.x as an address space. The B.x and J.x devices talk to the A.x devices via the translated B.x addresses. ACM knows that A.12 and A.13 are translated to B.12 and B.13.

Region 10 configuration:
Part of Region 10 is behind a NAT device.

Intra-region Direct IP-IP: Yes
When A.x and B.x talk to one another, B.x directs its audio stream to A.x's translated address (Rule1). When A.12 and A.13 talk to each other, they direct their audio streams to each other's native addresses (Rule2).

Inter-region Direct IP-IP: Yes
When A.x and J.x talk to one another, the translated address is used (Rule1). When B.x and J.x talk to one another, there is no NAT involved.

Region 11 configuration:
Region 11 is not behind a NAT device.

Intra-region Direct IP-IP: Yes
Permit intra-region shuffling. No NAT instructions are required.

Inter-region Direct IP-IP: Yes
Permit inter-region shuffling. No NAT instructions are required for this region. They are already covered in the configuration for Region 10.

Scenario 5b

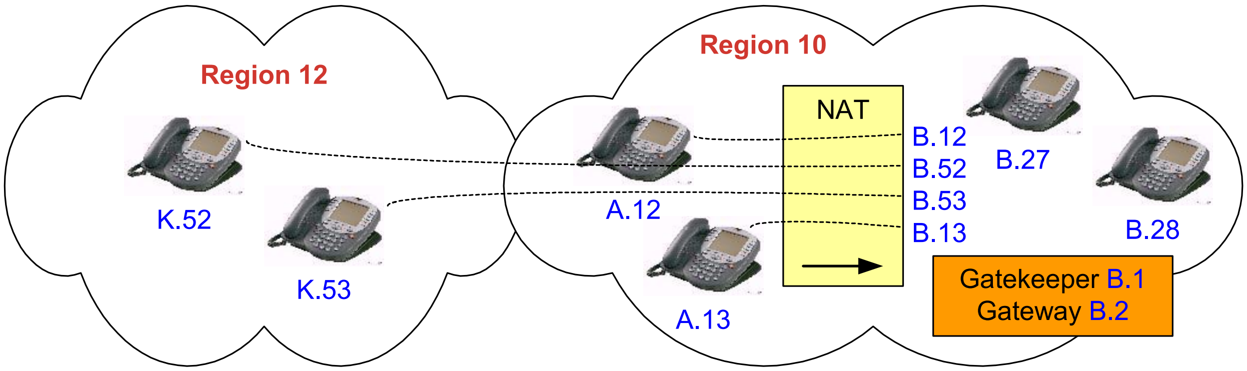

	translated to	
A.12	→	B.12
A.13	→	B.13
K.52	→	B.52
K.53	→	B.53

The difference between this scenario and scenario 5a is that region 12 is NATed, whereas region 11 in scenario 5a was not. In this scenario the A.x and K.x devices talk directly to each other and to the B.x devices. The B.x devices talk directly to each other, but talk to the A.x and K.x devices via their translated B.x addresses. As always, ACM is aware of the translations.

Region 12 configuration:
Region 12 is behind a NAT device.

Intra-region Direct IP-IP: Yes
When K.x devices talk to one another, they direct their audio streams to each other's native addresses (Rule2). **Native (NAT)** could be explicitly configured.

Inter-region Direct IP-IP: Yes
When K.x and A.x talk to one another, the native addresses are used (Rule2). When K.x and B.x talk to one another, the translated address is used (Rule1).

Region 10 configuration:
Part of Region 10 is behind a NAT device.

Intra-region Direct IP-IP: Yes
When A.x and B.x talk to one another, B.x is instructed to direct its audio stream to A.x's translated address (Rule1). When A.12 and A.13 talk to each other, they direct their audio streams to each other's native addresses (Rule2). Same as scenario 5a.

Inter-region Direct IP-IP: Yes
When A.x and K.x talk to one another, the native addresses are used (Rule2). When B.x and K.x talk to one another, the translated address is used (Rule1).

Scenario 5c

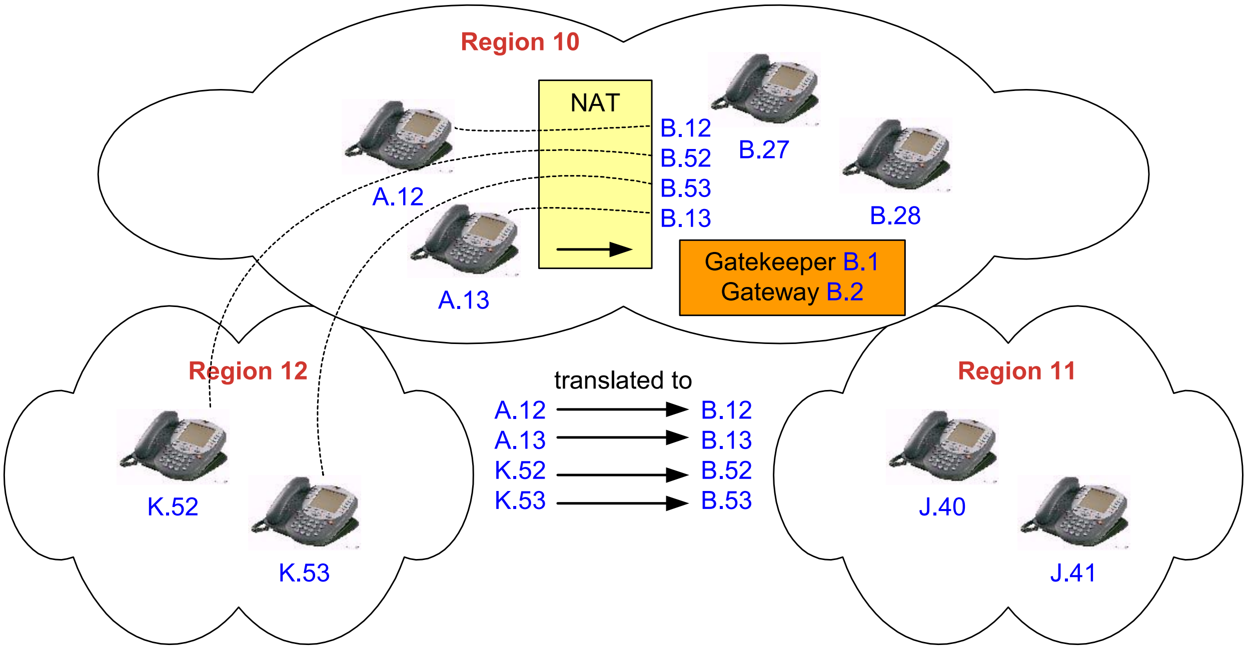

This is simply scenarios 5a and 5b put together. This scenario works using the same configurations already given for scenarios 5a and 5b.

Region 12 configuration:
Region 12 is behind a NAT device.

Intra-region Direct IP-IP: Yes
When K.x devices talk to one another, they direct their audio streams to each other's native addresses (Rule2). **Native (NAT)** could be explicitly configured.

Inter-region Direct IP-IP: Yes
When K.x and A.x talk to one another, the native addresses are used (Rule2). When K.x and B.x talk to one another, the translated address is used (Rule1). When K.x and J.x talk to one another, the translated address is used (Rule1).

Region 10 configuration:
Part of Region 10 is behind a NAT device.

Intra-region Direct IP-IP: Yes
When A.x and B.x talk to one another, B.x is instructed to direct its audio stream to A.x's translated address (Rule1). When A.12 and A.13 talk to each other, they direct their audio streams to each other's native addresses (Rule2).

Inter-region Direct IP-IP: Yes
When A.x and J.x talk to one another, the translated address is used (Rule1). When B.x and J.x talk to one another, there is no NAT involved. When A.x and K.x talk to one another, the native addresses are used (Rule2). When B.x and K.x talk to one another, the translated address is used (Rule1).

Region 11 configuration:
Region 11 is not behind a NAT device.

Intra-region Direct IP-IP: Yes
Permit intra-region shuffling. No NAT instructions are required.

Inter-region Direct IP-IP: Yes
Permit inter-region shuffling. No NAT instructions are required for this region. They are already covered in the configurations for Regions 10 and 12.