RAT: A Quick (And Dirty?) Push for Mobility Support *

Rhandeev Singh Y. C. Tay W. T. Teo S. W. Yeow
Mobile Computing Group
National University of Singapore
Lower Kent Ridge Cresent, Singapore 119260
rhandeev@comp.nus.edu.sg tay@acm.org {teoweetu, yeowshin}@comp.nus.edu.sg

Abstract

The goal of IP mobility support is to provide the means by which applications on distinct computers are able to communicate when one or both computers have changed their physical network location. Mobile IP (MIP) tries to provide such support with a solution at the network layer.

To date, MIP deployment is insignificant and this is likely to remain so, unless there is a breakthrough. In contrast, Network Address Translation (NAT) is a network technology that is widely deployed. The Reverse Address Translation protocol (RAT) attempts to deploy mobility support by riding on this popular, proven technology.

Registration of the mobile node in RAT is done through existing applications (e.g. browsers) and traffic delivery is by address translation rather than tunneling; thus, the mobile node does not require operating system support nor RAT-specific modifications. RAT separates the registration and forwarding functions of the home agent, and does not support handover of connections.

RAT is designed to interoperate with MIP, since one of its goals is to act as a bootstrap for encouraging MIP deployment.

1. Introduction

After years of discussion, the Mobile IP working group of the Internet Engineering Task Force finalized a network protocol (MIP) for supporting mobile hosts [4]. Although another two years have since passed (a long time in current Internet development), MIP deployment is still negligible.

In contrast, Network Address Translation (NAT) is a network technology that is spreading fast among many organizations — from small entreprenuer firms to large global corporations — despite reservations about its architectural implications [10] and its lack of standardization [13, 14].

We believe that a breakthrough in MIP deployment can be achieved by riding on the rising popularity of NAT (in this paper, NAT includes both address translation and port translation). The Reverse Address Translation (RAT) protocol [16] is an attempt in this direction: it proposes some modifications to the proliferating NAT boxes, so that they can act as the infrastructure for mobility support (in place of the missing mobility agents in a MIP infrastructure).

RAT offers only limited mobility support, and is not meant to be a replacement for MIP. Rather, our intention is for RAT to serve as a bridge to eventual and widespread MIP deployment. We believe that MIP is caught in a chicken-and-egg trap: deployment is insignificant because there are no compelling applications, and such applications are not developed because there is, in turn, no wide-spread deployment of MIP.

Although NAT was originally intended to be a temporary solution to the IP address exhaustion problem [1], its popularity has generated numerous novel applications — virtual servers, port forwarding, load balancing, firewalling, etc. [7, 8, 11].

Similarly, even though RAT offers only limited mobility support, we expect its deployment to bring forth mobile applications that will help MIP break the chicken-and-egg trap. RAT is thus a bootstrap for MIP deployment.

Another motivation for RAT lies in RFC2002's requirement that the mobile node be MIP-aware: it must implement agent discovery, location registration and tunneling (otherwise, foreign agent support is necessary). These requirements make MIP dependent on operating system vendors for deployment. To remove this dependency, RAT replaces tunneling with address translation, and registration is through applications (like Web browsers and Java servlets). An off-the-shelf mobile node (we call it a 0MN or *zero-knowledge mobile node*) can thus use RAT for mobility sup-

* RAT research and development is funded in part by the National University of Singapore under ARF grant RP960683.

port, and this should further facilitate deployment.

As a tradeoff, RAT (unlike MIP) does not support handovers, i.e. it does not attempt to maintain transport and higher-layer connections when the 0MN changes network location. Moreover, RAT inherits NAT's weaknesses, such as those on security [14] and interoperability with other protocols [12].

We begin in Section 2 with an overview and example of how RAT works. Section 3 then discusses the motivation and issues in our protocol design. The details of RAT and its relation with MIP is described in Section 4. We then discuss the implementation and performance issues in Section 5. Section 6 concludes with a summary of this work, as well as related work on MIP deployment.

2. RAT: An Overview

In the RAT protocol, a mobility-enabled network (known as the *home network*) has the following functional components:

- At least one *registration server* (RS)

  0MNs register themselves at a registration server in order to obtain mobility support while away from the home network.

- At least one *RAT device* (RD) per home network to perform network address translation

  The RAT device functions as an intermediary between 0MNs and their clients — known as *correspondent nodes* (CNs) — translating packets and relaying them to make mobility possible.

  The RAT device may or may not be physically distinct from the registration server.

- *Zero-knowledge mobile nodes* (0MNs)

  A 0MN may move from one network location to another. However, it is always associated with a fixed IP address in the home network, known as its *home address*.

  The network address associated with its current physical location is known as its *care-of address*.

Once a 0MN obtains an IP address in a foreign network, the following sequence of events typically occurs (see Figure 1):

1. The 0MN registers itself at a RS.

2. The RS sends a Reverse Translation Binding Request to a RD.

3. The RD begins address translation for the 0MN.

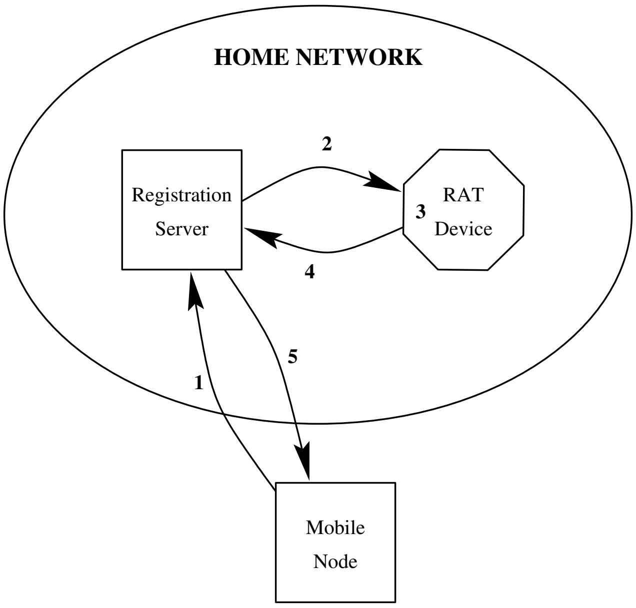

Figure 1. 0MN Registration Procedure

4. The RD sends a Reverse Translation Binding Reply to the RS.

5. The RS replies to the 0MN if desired.

The RD will perform address translation for the 0MN for a limited time interval, after which the Reverse Translation Binding Request (RTBR) expires, or until the RTBR is renewed or updated. During this time, correspondent nodes may establish connections with the 0MN (see Figure 2):

1. CN sends a datagram to the 0MN's home address.

2. The RD intercepts the datagram on behalf of the 0MN, makes it look like it came from the RD, and sends it to the 0MN's care-of address.

3. The 0MN replies to the RD.

4. The RD receives the 0MN's reply, makes it look like it came from the 0MN's home address, and sends it back to the CN.

5. The CN receives the reply thinking it came from the 0MN residing at its home address.

If the 0MN wishes to establish a connection with a CN, it will do so using its care-of address, without making use of mobility support (see Figure 2). In this way, mobility support is only required for CN-initiated sessions, but not for 0MN-initiated sessions.

As a result, the 0MN itself does not require any knowledge of the RAT protocol, which is a major advantage of RAT over MIP (see Section 4).

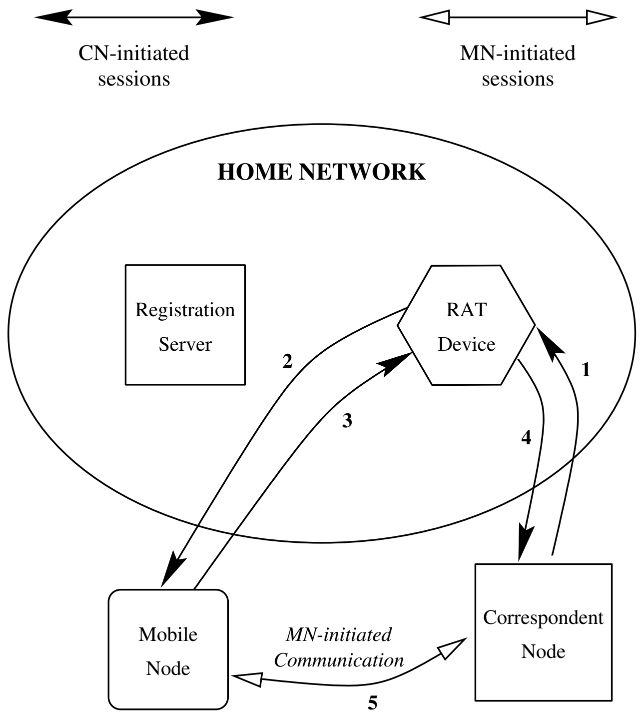

Figure 2. 0MN Mobility Support

3. Design Considerations

3.1. Functionality vs Simplicity

Our general approach is to keep the basic RAT protocol as simple as possible. To this end, we have chosen to sacrifice seamless mobility for simplicity. As argued in Section 3.5 below, this is not such a big price to pay because a true seamless mobility solution that operates at the network layer or above is either not possible now or not necessary.

Furthermore, a quick and easy MIP alternative leveraging on widely available network solutions could quickly gather the momentum needed to spearhead global MIP deployment, given RAT's ability to interoperate with the latter.

We believe that easy deployment of a simple mobility solution will promote the demand for MIP by stimulating the development of mobile applications and increasing the demand for seamless mobility support in the Internet, breaking the applications development-infrastructure deployment deadlock.

Hence, the basic RAT protocol has been kept simple. There are, in fact, no RAT protocol-specific requirements at the mobile nodes. This means that it is immediately possible to provide mobility support for any mobile node.

3.2. Flexibility and Scalability

We have several reasons for separating the registration process from the translation process in the RAT protocol.

Firstly, the process of translating and relaying IP datagrams is functionally distinct from the registration of mobile nodes. Since the choice of registration protocol is out of the scope of RAT, implementors will be free to use existing technology to make the registration process as convenient and secure as possible for end users, while relying on standard RAT device implementations.

It would be possible to use either a user-centered or a machine-centered model for registration. For example, a mobile computer could register via an encrypted channel, or an end user could register conveniently with an SSL-capable browser.

Secondly, address translation is computationally intensive. Separating RAT device from registration server allows us to reduce the computational load on the machine that is performing address translation. In addition, the mobile node identifies a registration server using a fixed IP address. Separating the RAT device from the registration server makes it possible to employ a variety of load-balancing setups without requiring the mobile nodes to know about this.

For example, we could have a registration server choose from a pool of RAT devices upon receiving a registration request from a mobile node, and send a corresponding Reverse Translation Binding Request to the least loaded RAT device.

A third reason for separating the RAT device from the registration server is that, in practice, high-volume address translation is likely to require dedicated hardware. Each network segment would require a separate RAT device. A central registration server allows us to modify the registration protocol without having to change every RAT device. This allows flexibility within a RAT installation.

Lastly, in the design of the protocol between the registration server and the RAT device, provisions have been made for the RAT device to be given the ability to service RFC2002-compliant registration requests. This means that it will be possible to modify a RAT device to interoperate with MIP mobility agents and service MIP mobile nodes if this becomes a necessity.

3.3. Migrating from RAT to MIP

The basic RAT protocol is meant to provide a simple mobility solution. Its main advantage is any node can be mobile if desired as no RAT-specific protocol support is needed for a 0MN. To push forward the deployment of a MIP infrastructure, we have designed RAT so that it can be incrementally enhanced to support a gradual installation of MIP entities. This is possible mainly because RAT adopts the

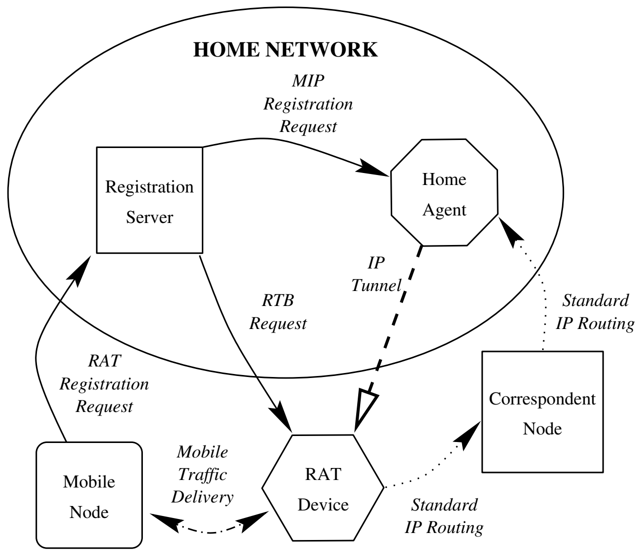

Figure 3. RAT-MIP Migration Step 1

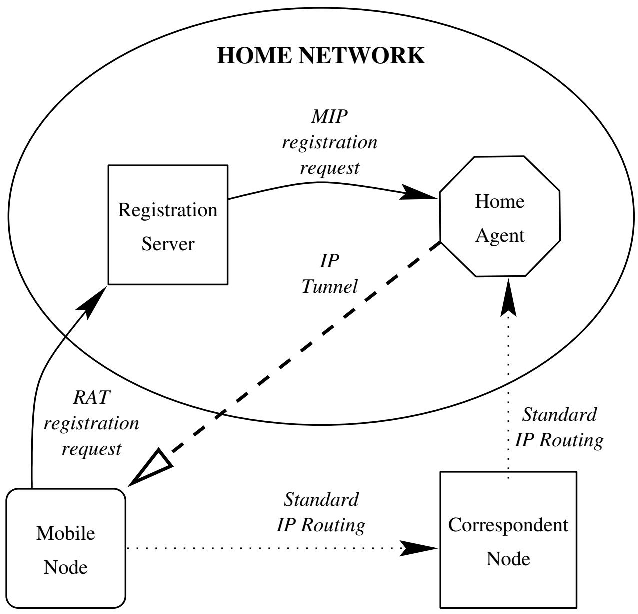

Figure 4. RAT-MIP Migration Step 2

same format and semantics of MIP registration messages. In the following, we describe how mobility support can migrate from RAT to MIP.

To begin supporting MIP, we need to deploy home agents instead of RAT devices at every network segment. Since the home agents' operating platform and hardware do not affect the mobile user, any available feasible implementation can be selected.

3.3.1. Migration Step 1 (Figure 3)

To interoperate with home agents, the RAT device will need to add support for IP tunnel decapsulation. After accepting a registration request from a 0MN, the registration server will now additionally send a MIP registration request to the home agent on behalf of the 0MN, with the tunnel endpoint as the RAT device. Mobile traffic delivery is now tunneled from the home agent to the RAT device and the RAT device will then do address translations to deliver the traffic to the 0MN.

As a result, a single RAT device can provide mobility support for multiple home networks by functioning as an intermediate agent between home agents and 0MNs (see Section 2).

3.3.2. Migration Step 2 (Figure 4)

The next transition phase is to support IP decapsulation at the mobile node. The movement detection and registration process remain unchanged, but mobile traffic delivery will no longer require a RAT device. The registration server will send MIP registration requests to the home agents as above but now the tunnel endpoint is not the RAT device but a mobile node. In this model, seamless mobility is now achievable.

3.3.3. Migration Step 3 (Figure 5)

The final evolution to complete MIP compatibility will require the mobile node to support movement detection and perform registration as specified in MIP. In this last stage, the registration server is no longer required; however the benefits of a central entity for managing mobility policies may argue for retaining the registration server. In such a configuration, the registration server acts as a trusted registration proxy between the mobile node and its home agent. Mobile nodes need only know its registration server address. All home address to home agent address associations are maintained at the registration server, allowing transparent handover to alternative home agents in the event of a home agent failure or migration.

3.4. Security

There are implications of RAT concerning existing host access control policies. Any policy based solely on a mobile node's home address and not the host's or user's identity will fail under RAT. This is because a mobile node may be located outside any network boundary defined by its home address; therefore, such security measures should deny access to mobile nodes in order to maintain expected trusted network boundaries.

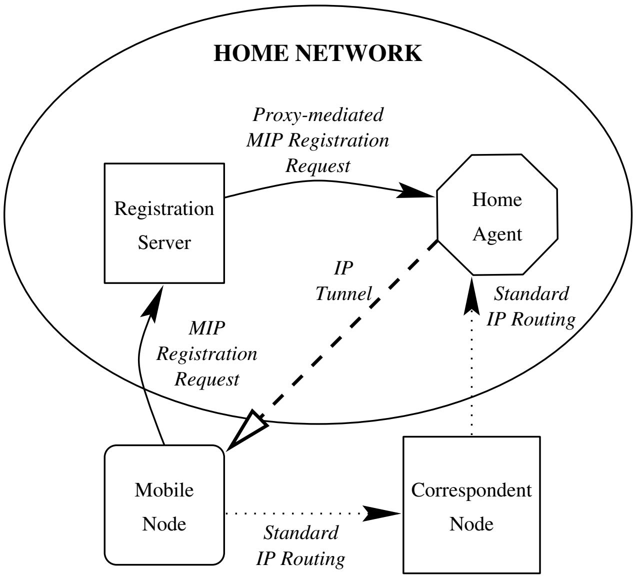

Figure 5. RAT-MIP Migration Step 3

Internal to the protocol, security can be divided into three categories. Firstly, there are the security issues arising from data delivery, that is, those involving address translation in the case of RAT, or packet forwarding through an IP tunnel in the case of MIP. In terms of data delivery, both MIP and RAT perform clear text routing. They provide no security protection for their payload. For additional security, external security measures are required.

In the case of RAT, the end-to-end tenet of IP networking is broken. The choice of external security measure is thus constrained. For example, security models for protocols such as IPSEC ESP will not work with RAT [2]. In contrast, such address-sensitive end-to-end security measures can work with MIP.

Secondly, we wish to ensure the authenticity of mobile nodes as they register at a registration server. As noted above, the RAT protocol does not attempt to standardize the registration process. Thus, a variety of existing user-centered as well as machine-centered authentication methods may be used, and the implementor is free to choose the most appropriate solution. Means that are even more secure than those provided by the MIP protocol may be used, such as public-key cryptographic methods.

Finally, RAT's Reverse Translation Binding (RTB) protocol specifies the communication between the RAT device and the registration server. This protocol must be made secure in order to prevent exploitation of the RAT device. Due to interoperation concerns with MIP, and because there is little need to redesign what we see as an extensible registration protocol which provides the same functionality in MIP, the RTB protocol employs the same security mechanisms and message formats as specified in MIP.

3.5. The Practical Reality

The fact that RAT is an extension of NAT has several important implications. Firstly, NAT is a tried and proven technology which is already widely researched and accepted. RAT development can potentially draw from this pool of resources.

This reality is not easy to ignore. NAT has become undeniably popular in an amazingly short span of time due at least in part to the pressures of IP address space exhaustion. The initial deployment of address translation solutions has already taken place.

Today, the existing NAT infrastructure can be extended into RAT to provide limited mobility in an incremental manner, whereas the deployment of a seamless mobility solution from scratch is more expensive. RAT has the potential to rapidly make limited mobility as ubiquitous as NAT itself.

Limited-area seamless mobility can be supported by link-layer handovers. In RFC2002, it was pointed out that link-layer mobility solutions are more efficient for micro-level mobility than MIP.

On the other hand, MIP aims to provide macro-level mobility. Unfortunately, where wide-area wireless coverage is provided by different incompatible link-layer solutions, seamless mobility across them may not be possible. This was pointed out in [9]. As a result, RAT should provide the same level of mobility support as MIP over current technology.

4. RAT and MIP in Operation

Both RAT and MIP achieve IP mobility by delivering datagrams (destined to the mobile node's globally recognized home address) arriving at the mobile node's home network to the mobile node's corresponding network address (care-of address known locally within the home network) in its current foreign network.

To achieve the above IP mobility, there must be some cooperation between the redirected traffic endpoints — from the home network to the mobile node. This cooperation can be classified into three main phases — the movement detection, the registration process and the mobile traffic delivery.

4.1. Phase 1: Movement Detection

In movement detection, a mobile node needs to determine changes in its location. The mobile node must detect that it has moved into another network segment or remains at its current location.

The MIP protocol introduces protocol-specific extensions in current ICMP router advertisements. These advertisements provide a mobile node with a list of IP addresses that are topologically correct in the advertised network link. A mobile node listening to these advertisements can thus determine its location. However, this is not all: movement from one foreign network to another foreign network, home network to foreign network and foreign network back to home network need to be distinguished and handled differently.

In the case of RAT, this movement detection is not handled transparently by the protocol. The onus lies on the human mobile user to determine the change in network locality. Since seamless mobility support (at the network layer) is not catered for, the mobile user needs only know he has moved to or from his home network. This is obvious when the typical home network is a fixed wire network. Such an approach simplifies network support for visiting mobile nodes and avoids placing additional protocol requirements on the mobile node. The same reasons motivated the choice of the mobile traffic delivery mechanism in Phase 3.

4.2. Phase 2: Registration Process

A mobile node must inform an entity in the home network of its current location in the Internet. This is to ensure the latter delivers traffic destined for the mobile node to the correct location in Phase 3.

Before the registration process can proceed, the mobile node must acquire a care-of address that can signify its current location. RAT operates in what is known as co-located mode in MIP. In co-located mode, the mobile node must acquire an IP address using an external mechanism available in the foreign network. This external mechanism is beyond both protocol specifications. For dial-up network access, the point-to-point (PPP) protocol is commonly used, while for other wired and wireless networks, dynamic host configuration (DHCP) protocol is popularly adopted.

After acquiring the care-of address, the mobile node must register this address with an entity in the home network. In MIP, this entity is known as a home agent. The home agent is also the same entity that will intercept traffic destined to any registered mobile node and deliver them to the care-of address in Phase 3. Therefore, the home agent is bound to the same network segment as the mobile node's home address.

For RAT, a registration server may handle the registration requests of every mobile node in an organization. The registration server is a separate entity from a mobile traffic delivery device. The registration server is solely for registration purposes, unlike the home agent in MIP. This provides a central body for the administration of mobility support policies. Since the registration process is not a frequent event, a registration server can easily support a large number of mobile nodes on multiple network segments.

MIP has specified the registration message formats for mobile nodes to request mobility service using the User Datagram Protocol (UDP) to port 434 of the home agents. This ensures interoperability between different implementations of mobile nodes and home agents.

RAT takes a very different approach to ensure interoperability and ease of deployment. The actual registration mechanism is not part of the RAT protocol. RAT uses applications standardized by other protocols to perform registration. As in Phase 1, RAT generally depends on the mobile user to perform registration. An organization can employ multiple types of application servers on the registration server. Using any common network application available on the mobile node such as secured shell clients and web browsers, the mobile user can register the care-of address with the corresponding application server.

The registration server discovery is dependent on the application protocol used. In the simplest configuration, the registration server address can be statically configured as in MIP. The complete movement detection and registration process is mobile user-oriented, since the user and not the machine is the best agent to determine if mobility is desired in the current location.

An example of an application protocol widely available on multiple platforms and easily deployable for RAT is the Secure Hypertext Transfer Protocol (https). A mobile user needs only a web browser to access a bookmarked URL hosted by a WWW server on the registration server. Java servlets may be downloaded from the registration server to request user authentication. After verification of user identity, the servlets can then transmit time-stamped "keep alive" beacons back to the registration server to re-confirm the mobile node's location at regular intervals. The registration messages sent can be encrypted or authenticated using a private key provided by the mobile user. The method of encryption or authentication need not be known to any entity other than the registration server.

4.3. Phase 3: Mobile Traffic Delivery

Wide Area Network IP mobility is only possible by the delivery of traffic arriving at the home network to the registered care-of address of the mode node which the traffic was meant for.

The mechanism employed to deliver a datagram to another location that would otherwise not be selected by its destination address in MIP is encapsulation. The default encapsulation protocol is IP Encapsulation within IP [5]. This requires that the home agent be the IP tunnel encapsulator and the mobile node be the decapsulator for co-located mode. Since the end hosts addresses never changes, MIP could provide seamless mobility. However, in reality this is still generally not achievable (see Section 3.5). Furthermore, to achieve seamless mobility, all traffic destined to the mobile node must be tunneled from the home agent, incurring an indirect, longer routing path for correspondent node to mobile node communication.

RAT leverages on the popularity of NAT, in particular Twice NAT [14], to provide mobile traffic delivery. One major attraction for NAT is end-node transparency, i.e. both parties communicating need not be aware that NAT takes place. The result is that both the mobile and correspondent nodes do not need to support additional protocols to communicate. For all purposes, mobile traffic are no different from the normal traffic on the Internet to the end nodes. All the re-routing complexity is hidden in a NAT box, the RAT device.

The RAT device does twice NAT to deliver traffic between the mobile and correspondent nodes. After Phase 2, the registration server informs the RAT device of the mobile node's latest care-of address. The mobile node's home address is translated to the care-of address for traffic from the correspondent node and vice-versa. Both the source and destination addresses of the datagrams which the RAT device intercepts are translated so that the latter will always appear as the virtual sender to the mobile node. The actual identity of the correspondent node is replaced by the RAT device. Traffic between the actual communicating end hosts must therefore pass through the RAT device. This incurs an even longer routing path than conventional MIP since both directions of the traffic now goes through a dog-legged route.

It should be pointed out that, unlike MIP, the mobile node always uses the acquired care-of address as its source address when outside its home network. Since the source address used is topologically correct, Ingress filtering [6] does not affect RAT mobility support in contrast to standard MIP operation.

For a typical mobile node, network communication is mostly initiated by the mobile node and not the correspondent node. In this case, the correspondent node communicates directly with the mobile node's care-of address (not the mobile node's home address) and the RAT device is not involved in the routing of any mobile traffic at all. Therefore, in comparison to MIP, 0MNs do not always incur unnecessary routing delay.

5. Implementation Issues

5.1. Existing Prototype

We have implemented a RAPT prototype (see 5.3.2) completely in Linux user space (http://cram.comp.nus.edu.sg:8080/cram/rat). Application level support is currently lacking for popular protocols such as IRC, FTP, HTTP and Real Network's protocols, but it is easy to adapt existing NAT implementations of application level gateways (ALGs) for these protocols to our RAT implementation.

We have successfully tested this implementation with HTTP, but more tests are needed in the case where there are proxy servers. Tests with FTP also worked; here, in the absence of an ALG, the data channel bypassed the RAT device.

We have also found it necessary for the RAT device to reset the TTL value of translated packets, to facilitate operation over long distances. This is a result of the long, indirect routing paths introduced by the protocol.

Since a user-space implementation excludes in-situ processing of the IP datagrams, we plan to port the address translation code into the Linux kernel for performance reasons. We have therefore put off performance measurements till after that code is moved into the kernel; meanwhile, what follows is a performance analysis of address translation in our prototype.

5.2. Performance

In our implementation of the RAT device, registration and binding information is stored in doubly-hashed dynamic tables. Whereas registration information is housed in a single table, each registration entry itself has a corresponding translation table which stores all the information required to perform address translation for a particular mobile node. Given a good pair of hashing functions, this setup gives constant time amortized lookups, insertions and deletions over the tables, which is important since the information in the tables is needed to translate every single packet.

The expiry of registrations and bindings, however, is implemented using two priority queues; this keeps the tables small and improves performance. Expired entries are removed whenever an incoming packet arrives. At best, it takes constant time to translate each packet; at worst, it takes $O(n \log n)$ time, where n is the maximum of the number of registration entries in the registration table and the number of bindings in the largest translation table. The worst case happens when a packet to be translated arrives and all entries in the largest table expire at about the same time; this is not likely to happen often, especially if the registration server uses multiple RAT devices.

5.3. Scalability

The RAT device must be able to determine which CNs and 0MNs are associated with each IP datagram that it encounters, in order to perform address translation correctly.

This can be carried out in two ways, and both methods require the RAT device to possess a pool of IP addresses that will be used during delivery. The scalability limitations of the protocol depend on which method is employed.

5.3.1. RAT using Network Address Translation

Using the first method, the RAT device examines the source and destination IP addresses of each packet to determine which CN and 0MN the packet is associated with. We will refer to this simply as RAT within this section.

To eliminate ambiguity, the RAT device must associate a unique IP address from its pool of addresses with each CN that is communicating with a 0MN at any one time. The number of CNs that are able to connect to 0MNs during any period of time is thus limited by the size of the RAT device's IP address pool.

This approach is independent of the protocols above the network layer. Therefore, RAT is able to function even when transport layer information is absent from packets.

5.3.2. RAT using Network Address and Port Translation

The second method of mapping an IP datagram to its associated CN and 0MN makes use of three additional fields: the transport protocol type and the transport layer source and destination connection identifiers (e.g. TCP port number or ICMP echo request/reply ID field). We call this *Reverse Address and Port Translation* (RAPT) within this section.

RAPT does not have the same limitations on scalability as does RAT, since the upper limit on the number of concurrent connections between CNs and 0MNs is larger in RAPT than in RAT by several orders of magnitude.

Instead, RAPT scalability is limited by its inability to support protocols that do not employ a transport layer connection identifier. For example, IP-in-IP tunnels cannot be translated by a RAT device that employs RAPT.

5.3.3. Application Level Gateways (ALGs)

Most application protocols assume that network and transport layer addresses do not change during transit. Some of these protocols embed network or transport layer address information in the payload. By default, such protocols will not function with a NAT device (e.g. IRC, RTSP, FTP) [12].

RAT and RAPT inherit this limitation from NAT. Therefore, each RAT device requires ALGs which ensure that information in the payload of each datagram remains meaningful as network and/or transport layer addresses are translated.

6. Conclusion

MIP deployment is stalled by a chicken-and-egg trap. We believe that RAT can break this "no deployment/no application" cycle by leveraging on the fast-spreading infrastructure of NAT boxes. Specifically, RAT

(1) uses applications for location registration and

(2) replaces tunneling by address translation for traffic forwarding.

Thus, instead of deploying a MIP infrastructure (of home and foreign agents), mobility support can be obtained by modifying the NAT boxes and adding registration servers to home networks.

Furthermore, (1) and (2) remove the need for operating system support on the mobile node; this makes mobility support available to any Internet-ready mobile node, thus facilitating RAT deployment.

RAT inherits NAT's weaknesses and does not support connection handovers. However, it interoperates with MIP, so RAT's weaknesses can be removed by a migration to MIP. Indeed, RAT is designed so that such a migration can be achieved incrementally.

RAT development grew out of the focus on deployment in the research done by the Mobile Computing Group at the National University of Singapore. Other work in this focus includes:

- A billing protocol for wireless traffic. This protocol [17] specifies how billing can be done for network traffic between a wireless base station and a mobile node. Such a facility can encourage retailers (cafes, hotels, airports, etc) to install wireless local area networks, since they can then profit from the customers' network traffic. The protocol thus facilitates a bottom-up deployment of wireless infrastructure for mobile computing.

- A Mobile IP Extension for Private Internets Support (MPN). MIP deployment requires an infrastructure of mobility agents. It is unlikely that the public sector will bear the costs of initiating such an infrastructure. On the other hand, corporations increasingly use private addresses for their private networks, and although Mobile IP can work within a private routing realm, the case for them to spend on a MIP infrastructure is strengthened if the mobility support can be extended to movement across routing realm boundaries. MPN is a MIP extension that supports movement between a private network and the global Internet, and between private networks [15]. By enlarging the scope of mobility support, MPN increases the incentive for corporate deployment of MIP.

7. References

[1] K. Egevang and P. Francis, "The IP Network Address Translator (NAT)", RFC1631, May 1994.

[2] R. Atkinson, IP Encapsulating Security Payload (ESP) RFC1827, August 1995.

[3] Y. Rekhter, B. Moskowitz, D. Karrenberg, G.J. de Groot and E. Lear, "Address Allocation for Private Internets", RFC1918, February 1996.

[4] C. Perkins, "IP Mobility Support", RFC2002, October 1996.

[5] C. Perkins, "IP Encapsulation within IP", RFC2003, October 1996.

[6] P. Ferfuson and D. Senie, "Network Ingress Filtering: Defeating Denial of Service Attacks which employ IP Source Address Spoofing", RFC2267, January 1998

[7] P. Srisuresh and D. Gan, "Load Sharing using IP Network Address Translation (LSNAT)", RFC2391, August 1998.

[8] M. Allman, S. Ostermann and C. Metz, "FTP Extensions for IPv6 and NATs", RFC2428, September 1998.

[9] Mary G. Baker, "Changing Communication Environments in MosquitoNet", Proceedings of the IEEE Workshop on Mobile and Computing Systems and Applications, December 1994.

[10] T. Hain, "Architectural implications of NAT", draft-iab-nat-implications-02.txt, October 1998.

[11] Michael Hasenstein, "IP Network Address Translation", ¡http://www.csn.tu-chemnitz.de/ mha/linux-ip-nat/diplom/nat.html¿, May 1997.

[12] M. Holdrege and P. Srisuresh, "IP Network Address Translator (NAT) Protocol Issues", draft-ietf-nat-protocol-01.txt, November 1998.

[13] P. Srisuresh and K. Egevang, "Traditional IP Network Address Translator (Traditional NAT)", draft-ietf-nat-traditional-01.txt, October 1998.

[14] P. Srisuresh and M. Holdrege, "IP Network Address Translator (NAT) Terminology and Considerations", draft-ietf-nat-terminology-01.txt, October 1998.

[15] W. T. Teo and Y. Li, "Mobile IP Extension for Private Internets Support (MPN)", draft-teoyli-mobileip-mvpn-01.txt, November 1998.

[16] W. T. Teo, S. W. Yeow and R. Singh, "Reverse Network Address Translators (RAT)", draft-teoyeow-mip-rat-01.txt, December 1998.

[17] W. L. Wang and C. W. Wu, "A Protocol for Billing Mobile Network Access Devices Operating in Foreign Networks", Proceedings of the IEEE Workshop on Enabling Technologies - Infrastructure for Collaborative Enterprises, Stanford, CA, June 1998, 262–268