Network Working Group
Internet-Draft
Intended status: Standards Track
Expires: October 21, 2014

D. Cheng
Huawei
J. Korhonen
Broadcom
M. Boucadair
France Telecom
S. Sivakumar
Cisco Systems
April 19, 2014

RADIUS Extensions for IP Port Configuration and Reporting
draft-cheng-radext-ip-port-radius-ext-00

Abstract

This document defines three new RADIUS attributes. For device that implementing IP port ranges, these attributes are used to communicate with a RADIUS server in order to configure and report TCP/UDP ports and ICMP identifiers, as well as mapping behavior for specific hosts. This mechanism can be used in various deployment scenarios such as CGN, NAT64, Provider WiFi Gateway, etc.

This document does not make any assumption about the deployment context.

Requirements Language

The key words "MUST", "MUST NOT", "REQUIRED", "SHALL", "SHALL NOT", "SHOULD", "SHOULD NOT", "RECOMMENDED", "MAY", and "OPTIONAL" in this document are to be interpreted as described in RFC 2119 [RFC2119].

Status of This Memo

This Internet-Draft is submitted in full conformance with the provisions of BCP 78 and BCP 79.

Internet-Drafts are working documents of the Internet Engineering Task Force (IETF). Note that other groups may also distribute working documents as Internet-Drafts. The list of current Internet-Drafts is at http://datatracker.ietf.org/drafts/current/.

Internet-Drafts are draft documents valid for a maximum of six months and may be updated, replaced, or obsoleted by other documents at any time. It is inappropriate to use Internet-Drafts as reference material or to cite them other than as "work in progress."

This Internet-Draft will expire on October 21, 2014.

Table of Contents

1. Introduction

In a broadband network, customer information is usually stored on a RADIUS server [RFC2865] and at the time when a user initiates an IP connection request, the RADIUS server will populate the user's configuration information to the Network Access Server (NAS), which is usually co-located with the Border Network Gateway (BNG), after the connection request is granted. The Carrier Grade NAT (CGN)

function may also be implemented on the BNG, and therefore CGN TCP/ UDP port (or ICMP identifier) mapping behavior can be configured on the RADIUS server as part of the user profile, and populated to the NAS in the same manner. In addition, during the operation, the CGN can also convey port/identifier mapping behavior specific to a user to the RADIUS server, as part of the normal RADIUS accounting process.

The CGN device that communicates with a RADIUS server using RADIUS extensions defined in this document may perform NAT44 [RFC3022], NAT64 [RFC6146], or Dual-Stack Lite AFTR [RFC6333] function.

For the CGN example, when IP packets traverse a CGN, it would perform TCP/UDP source port mapping or ICMP identifier mapping as required. A TCP/ UDP source port or ICMP identifier, along with source IP address, destination IP address, destination port and protocol identifier if applicable, uniquely identify a session. Since the number space of TCP/UDP ports and ICMP identifiers in CGN's external realm is shared among multiple users assigned with the same IPv4 address, the total number of a user's simultaneous IP sessions is likely to subject to port quota.

The attributes defined in this document may also be used to report the assigned port set in some deployment such as Provider Wi-Fi [I-D.gundavelli-v6ops-community-wifi-svcs]. For example, a visiting host can be managed by a CPE which will need to report the assigned port set to the service platform. This is required for identification purposes (see WT-146 for example).

This document proposes three new attributes as RADIUS protocol's extensions, and they are used for separate purposes as follows:

- IP-Port-Limit:This attribute may be carried in RDIUS Acces-Accept, Accounting-Request or CoA-Request packet. The purpose of this attribute is to limit the total number of TCP/UDP ports and/or ICMP identifiers that an IP subscriber can use..

- IP-Port-Range:This attribute may be carried in RADIUS Access-Accept, Accounting-Request or CoA-Request packet. The purpose of this attribute is to specify the range of TCP/UDP ports and/or ICMP identifiers that an IP subscriber can use associated with an IPv4 address.

- IP-Port-Forwarding-Map:This attribute may be carried in RADIUS Access-Accept, Accounting-Request or CoA-Request packet. The purpose of this this attribute is to specify how a TCP/UDP port (or an ICMP identifier) mapping to another TCP/UDP port (or an ICMP identifier).

This document was constructed using the [RFC2629] .

2. Terminology

Some terms that are used in this document are listed as follows:

- IP Port - This term refers to IP transport protocol port, including TCP port, UDP port and ICMP identifier.

- IP Port Limit - This is the maximum number of TCP ports, or UDP ports, or the total of the two, or ICMP identifiers, or the total of the three, that a device supporting port ranges can use when performing mapping on TCP/ UDP ports or ICMP identifiers for a specific user.

- IP Port Range - This specifies a set of TCP/UDP port numbers or ICMP identifiers, indicated by the port/identifier with the smallest numerical number and the port/identifier with the largest numerical number, inclusively.

- Internal IP Address - The IP address that is used as a source IP address in an outbound IP packet sent toward a device supporting port ranges in the internal realm. In IPv4 case, it is typically a private address [RFC1918].

- External IP Address - The IP address that is used as a source IP address in an outbound IP packet after traversing a device supporting port ranges in the external realm. In IPv4 case, it is typically a global and routable IP address.

- Internal Port - The internal port is a UDP or TCP port, or an ICMP identifier, which is allocated by a host or application behind a device supporting port ranges for an outbound IP packet in the internal realm.

- External Port - The external port is a UDP or TCP port, or an ICMP identifier, which is allocated by a device supporting port ranges upon receiving an outbound IP packet in the internal realm, and is used to replace the internal port that is allocated by a user or application.

- External realm - The networking segment where IPv4 public addresses are used in respective of the device supporting port ranges.

- Internal realm - The networking segment that is behind a device supporting port ranges and where IPv4 private addresses are used.

- Mapping - This term in this document associates with a device supporting port ranges for a relationship between an internal IP address, internal port and the protocol, and an external IP address, external port, and the protocol.

- Port-based device - A device that is capable of providing IP address and TCP/UDP port mapping services and in particular, with the granularity of one or more subsets within the 16-bit TCP/UDP port number range. A typical example of this device can be a CGN, CPE, Provider Wi-Fi Gateway, etc.

Note the terms "internal IP address", "internal port", "internal realm", "external IP address", "external port", "external realm", and "mapping" and their semantics are the same as in [RFC6887], and [RFC6888].

3. RADIUS Attributes

[Discussion: Should we define a dedicated attribute (port_set_policies) to configure the following policies: (1) enforce port randomization, (2) include/exclude the WKP in the port assignment, (3) preserve parity, (4) quota for explicit port mapping, (5) DSCP marking policy, (6) Port hold down timer, (7) port hold down pool, etc. Perhaps we don't need to cover all these parameters. - The discussion should be in a separate draft allowing this draft dedicated to RADIUS extension only.]

In this section, we define the details of the following three new attributes:

- IP-Port-Limit Attribute

- IP-Port-Range Attribute

- IP-Port-Forwarding-Map Attribute

All these attributes are allocated from the RADIUS "Extended Type" code space per [RFC6929].

3.1. Extended-Type for IP-Port-Type

This section defines a new Extended-Type for IP port type. The IP port type may be one of the following:

- Refer to TCP port, UDP port, and ICMP identifier

- Refer to TCP port and UDP port

- Refer to TCP port
- Refer to UDP port
- Refer to ICMP identifier

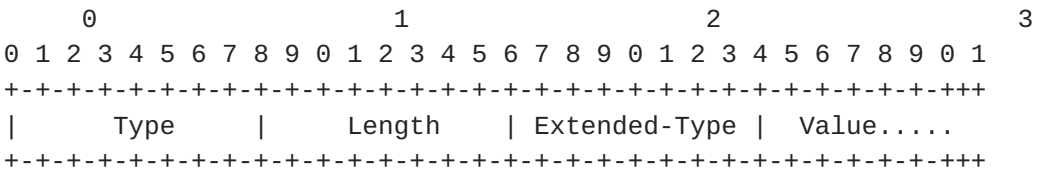

Type:

TBA1 - Extended-Type-1 (241), Extended-Type-2 (242), Extended-Type-3 (243), or Extended-Type-4 (244) per [RFC6929].

Length:

This field indicates the total length in octets of all fields this attribute, including the Type, Length, Extended-Type, and Value.

Extended-Type:

This one octet filed indicates the IP port as follows:

TBA1-1:

Refer to TCP port, UDP port, and ICMP identifier as a whole.

TBA1-2:

Refer to TCP port and UDP port as a whole.

TBA1-3:

Refer to TCP port only.

TBA1-4:

Refer to UDP port only.

TBA1-5:

Refer to ICMP identifier only.

Value:

This field contains one or more octects, and the data format MUST be a valid RADIUS data type.

The interpretation of this field is determined by the identifier of "TBA1.{TBA1-1..TBA1-5} along with the embedded TLV.

3.2. IP-Port-Limit Attribute

This attribute contains an Extended-Type along with a TLV data type with format defined in [RFC6929]. It specifies the maximum number of IP ports for a user.

The IP-Port-Limit MAY appear in an Access-Accept packet, it MAY also appear in an Access-Request packet as a hint by the device supporting port ranges, which is co-allocated with the NAS, to the RADIUS server as a preference, although the server is not required to honor such a hint.

The IP-Port-Limit MAY appear in an CoA-Request packet.

The IP-Port-Limit MAY appear in an Accounting-Request packet.

The IP-Port-Limit MUST NOT appear in any other RADIUS packets.

The format of the IP-Port-Limit RADIUS attribute format is shown below. The fields are transmitted from left to right.

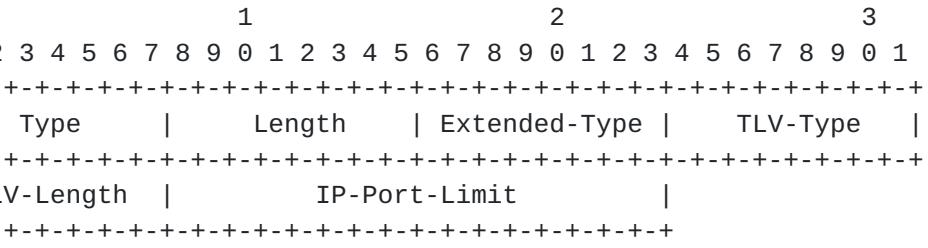

Type:

TBA1 - Extended-Type-1 (241), Extended-Type-2 (242), Extended-Type-3 (243), or Extended-Type-4 (244) per [RFC6929].

Length:

This field indicates the total length in octets of all fields of this attribute, including the Type, Length, Extended-Type, and the entire length of the embedded TLV.

Extended-Type:

This one octet field contains a value that indicates the IP port type, refer to Section 3.1 for details.

TLV-Type:

TBA2: for IP-Port-Limit TLV.

TLV-Length:

4.

IP-Port-Limit:

This field contains the maximum number of IP ports of which, the port type is specified by the value contained in the Extended-Type field.

Note this field is semantically associated with the identifier "TBA1.{TBA1-1..TBA1-5}.

3.3. IP-Port-Range Attribute

This attribute contains an Extended-Type along with a TLV data type with format defined in [RFC6929]. It contains a range of numbers for IP ports allocated by a device supporting port ranges for a given subscriber along with an external IPv4 address.

In some CGN deployment scenarios as described such as L2NAT [I-D.miles-behave-l2nat], DS-Extra-Lite [RFC6619] and Lightweight 4over6 [I-D.ietf-softwire-lw4over6], parameters at a customer premise such as MAC address, interface ID, VLAN ID, PPP session ID, IPv6 prefix, VRF ID, etc., may also be required to pass to the RADIUS server as part of the accounting record.

The IP-Port-Range MAY appear in an Accounting-Request packet.

The IP-Port-Range MUST NOT appear in any other RADIUS packets.

The format of the IP-Port-Range RADIUS attribute format is shown below. The fields are transmitted from left to right.

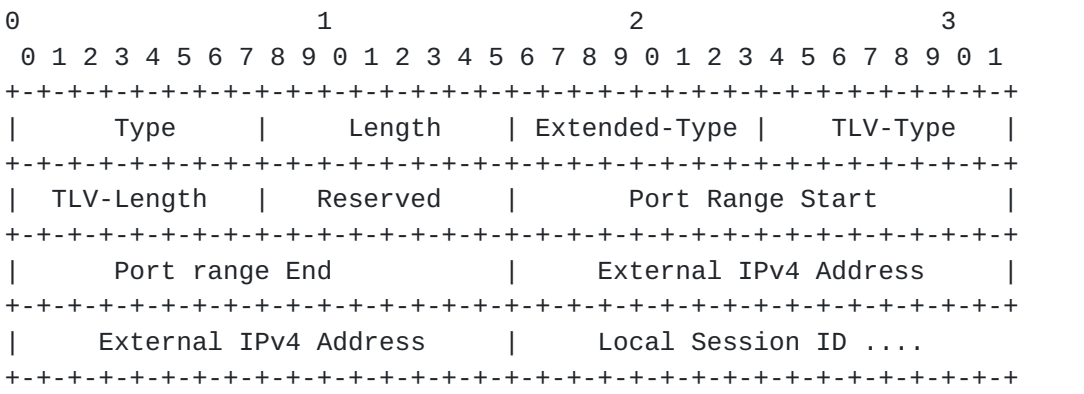

Type:

TBA1 - Extended-Type-1 (241), Extended-Type-2 (242), Extended-Type-3 (243), or Extended-Type-4 (244) per [RFC6929]

Length:

This field indicates the total length in octets of all fields of this attribute, including the Type, Length, Extended-Type, and the entire length of the embedded TLV.

Extended-Type:

This one octet field contains a value that indicates the IP port type, refer to Section 3.1 for details.

TLV-Type:

TBA3:

Allocation for IP-Port-Range TLV.

TBA4:

De-allocation for IP-Port-Range TLV.

TLV-Length:

>=11.

Reserved:

This field MUST be set to zero by the sender and ignored by the receiver.

Port Range Start:

This field contains the smallest IP port number, as specified in the Extended-Type, in the IP port range.

Port Range End:

This field contains the largest IP port number, as specified in the Extended-Type, in the IP port range.

External IPv4 Address:

This field contains the IPv4 address assigned to the associated subscriber to be used in the external realm. If set to 0.0.0.0, the allocation address policy is local to the device supporting port ranges.

Local Session ID:

This is an optional field and if presents, it contains a local session identifier at the customer premise, such as MAC address, interface ID, VLAN ID, PPP sessions ID, VRF ID, IPv6 address/ prefix, etc. The length of this field equals to the value in the TLV Length field minus 11 octets. If this field is not present, the port range policies must be enforced to all subscribers using a local subscriber identifier.

Note the data group in the "TLV Value" field above (i.e., "Port Range Start", "Port Range End", "External IPv4 Address", and "Local Session ID") is indicated by the identifier TBA1.{TBA1-1..TBA1-5}.{TBA3..TBA4}.

3.4. IP-Port-Forwarding-Map Attribute

This attribute contains an Extended-Type along with a TLV data type with format defined in [RFC6929]. It contains a 16-bit Internal Port that identifies the source TCP/UDP port number of an IP packet sent by the user, or the destination port number of an IP packet destined to the user, and in both cases, the IP packet travels behind the NAT device. Also it contains a 16-bit Configured External Port that identifies the source TCP/UDP port number of an IP packet sent by the user, or the destination port number of an IP packet destined to the user, and in both cases, the IP packet travels outside of the NAT device. In addition, the attribute may contain a 32-bit IPv4 address or a 128-bit IPv6 address, respectively, as their respective NAT mappings internal IP address. Together, the port pair and IP address determine the port mapping rule for a specific IP flow that traverses a NAT device.

The attribute MAY appear in an Access-Accept packet, and may also appear in an Accounting-Request packet. In either case, the attribute MUST NOT appear more than once in a single packet.

The attribute MUST NOT appear in any other RADIUS packets.

The format of the Port-Forwarding-Map RADIUS attribute format is shown below. The fields are transmitted from left to right.

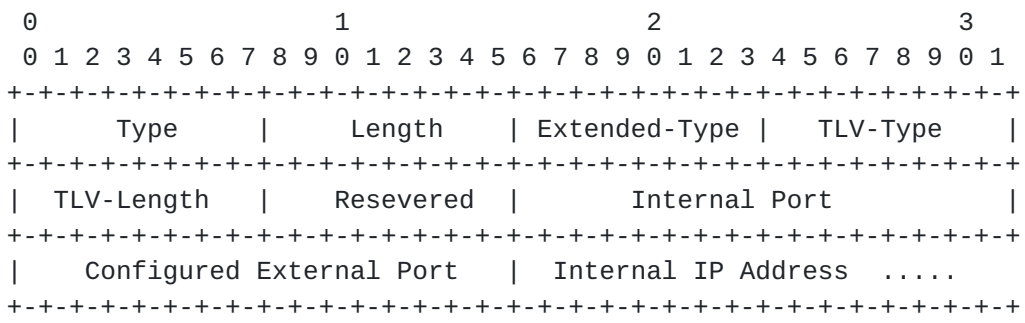

Type:

Type:

TBA1 - Extended-Type-1 (241), Extended-Type-2 (242), Extended-Type-3 (243), or Extended-Type-4 (244) per [RFC6929]

Length:

This field indicates the total length in octets of all fields of this attribute, including the Type, Length, Extended-Type, and the entire length of the embedded TLV.

Extended-Type:

This one octet field contains a value that indicates the IP port type, refer to Section 3.1 for details.

TLV-Type:

TBA5 - It indicates IP port mapping, and the associated internal IP address is an IPv4 or IPv6 address, or not included.

TLV-Length:

>=7.

Reserved:

This field is set to zero by the sender and ignored by the receiver.

Internal Port:

This field contains the internal port for the CGN mapping.

Configured External Port:

This field contains the external port for the CGN mapping.

Internal IP Address:

This field may or may not present, and when it does, contains the internal IPv4 or IPv6 address for the CGN mapping. Its length equal to the value in the TLV Length field minus 7.

Note the data group in the "TLV Value" field above (i.e., "Internal Port", "Configured External Port", and "Internal IP Address") is indicated by the identifier TBA1.{TBA1-1..TBA1-5}.TBA5.

4. Applications, Use Cases and Examples

This section describes some applications and use cases to illustrate the use of the attributes propsoed in this document.

4.1. Managing CGN Port Behavior using RADIUS

In a broadband network, customer information is usually stored on a RADIUS server, and the BNG hosts the NAS. The communication between the NAS and the RADIUS server is triggered by a subscriber when the user signs in to the Internet service, where either PPP or DHCP/DHCPv6 is used. When a user signs in, the NAS sends a RADIUS Access-Request message to the RADIUS server. The RADIUS server validates the request, and if the validation succeeds, it in turn sends back a RADIUS Access-Accept message. The Access-Accept message carries configuration information specific to that user, back to the NAS, where some of the information would pass on to the requesting user via PPP or DHCP/DHCPv6.

A CGN function in a broadband network would most likely reside on a BNG. In that case, parameters for CGN port/identifier mapping behavior for users can be configured on the RADIUS server. When a user signs in to the Internet service, the associated parameters can be conveyed to the NAS, and proper configuration is accomplished on the CGN device for that user.

Also, CGN operation status such as CGN port/identifier allocation and de-allocation for a specific user on the BNG can also be transmitted back to the RADIUS server for accounting purpose using the RADIUS protocol.

RADIUS protocol has already been widely deployed in broadband networks to manage BNG, thus the functionality described in this specification introduces little overhead to the existing network operation.

In the following sub-sections, we describe how to manage CGN behavior using RADIUS protocol, with required RADIUS extensions proposed in Section 3.

4.1.1. Configure IP Port Limit for a User

In the face of IPv4 address shortage, there are currently proposals to multiplex multiple subscribers' connections over a smaller number of shared IPv4 addresses, such as Carrier Grade NAT [RFC6888], Dual-Stack Lite [RFC6333], NAT64 [RFC6146], etc. As a result, a single IPv4 public address may be shared by hundreds or even thousands of subscribers. As indicated in [RFC6269], it is therefore necessary to impose limits on the total number of ports available to an individual subscriber to ensure that the shared resource, i.e., the IPv4 address remains available in some capacity to all the subscribers using it, and port limiting is also documented in [RFC6888] as a requirement.

The IP port limit imposed to a specific subscriber may be on the total number of TCP and UDP ports plus the number of ICMP identifiers, or with other granularities as defined in Section 3.2.

The per-subscriber based IP port limit is configured on a RADIUS server, along with other user information such as credentials. The value of these IP port limit is based on service agreement and its specification is out of the scope of this document.

When a subscriber signs in to the Internet service successfully, the IP port limit for the subscriber is passed to the BNG based NAS, where CGN also locates, using a new RADIUS attribute called IP-Port-Limit (defined in Section 3.2), along with other configuration parameters. While some parameters are passed to the subscriber, the IP port limit is recorded on the CGN device for imposing the usage of TCP/UDP ports and ICMP identifiers for that subscriber.

Figure 1 illustrates how RADIUS protocol is used to configure the maximum number of TCP/UDP ports for a given subscriber on a NAT44 device.

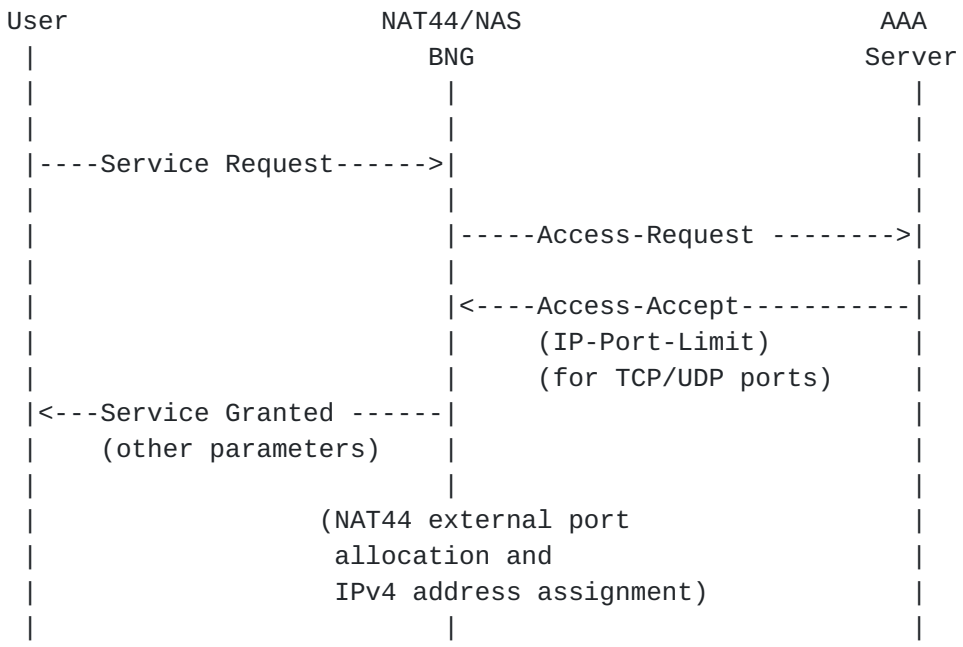

Figure 1: RADIUS Message Flow for Configuring NAT44 Port Limit

The IP port limit created on a CGN device for a specific user using RADIUS extension may be changed using RADIUS CoA message [RFC5176] that carries the same RADIUS attribute. The CoA message may be sent from the RADIUS server directly to the NAS, which once accepts and sends back a RADIUS CoA ACK message, the new IP port limit replaces the previous one.

Figure 2 illustrates how RADIUS protocol is used to increase the TCP/UDP port limit from 1024 to 2048 on a NAT44 device for a specific user.

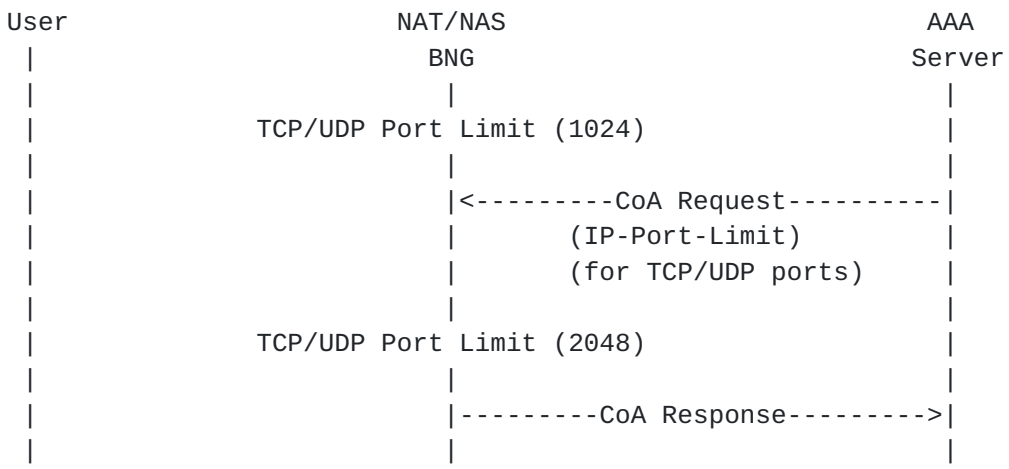

Figure 2: RADIUS Message Flow for changing a user's NAT44 port limit

4.1.2. Report IP Port Allocation/De-allocation

Upon obtaining the IP port limit for a subscriber, the CGN device needs to allocate a TCP/UDP port or an ICMP identifiers for the subscriber when receiving a new IP flow sent from that subscriber.

As one practice, a CGN may allocate a bulk of TCP/UDP ports or ICMP identifiers once at a time for a specific user, instead of one port/identifier at a time, and within each port bulk, the ports/identifiers may be randomly distributed or in consecutive fashion. When a CGN device allocates bulk of TCP/UDP ports and ICMP identifiers, the information can be easily conveyed to the RADIUS server by a new RADIUS attribute called the IP-Port-Range (defined in Section 3.3). The CGN device may allocate one or more TCP/UDP port ranges or ICMP identifier ranges, or generally called IP port ranges, where each range contains a set of numbers representing TCP/UDP ports or ICMP identifiers, and the total number of ports/identifiers must be less or equal to the associated IP port limit imposed for that subscriber. A CGN device may choose to allocate a small port range, and allocate more at a later time as needed; such practice is good because its randomization in nature.

At the same time, the CGN device also needs to decide the shared IPv4 address for that subscriber. The shared IPv4 address and the pre-allocated IP port range are both passed to the RADIUS server.

When a subscriber initiates an IP flow, the CGN device randomly selects a TCP/UDP port or ICMP identifier from the associated and pre-allocated IP port range for that subscriber to replace the original source TCP/UDP port or ICMP identifier, along with the replacement of the source IP address by the shared IPv4 address.

A CGN device may decide to "free" a previously assigned set of TCP/UDP ports or ICMP identifiers that have been allocated for a specific subscriber but not currently in use, and with that, the CGN device must send the information of the de-allocated IP port range along with the shared IPv4 address to the RADIUS server.

Figure 3 illustrates how RADIUS protocol is used to report a set of ports allocated and de-allocated, respectively, by a NAT44 device for a specific user to the RADIUS server.

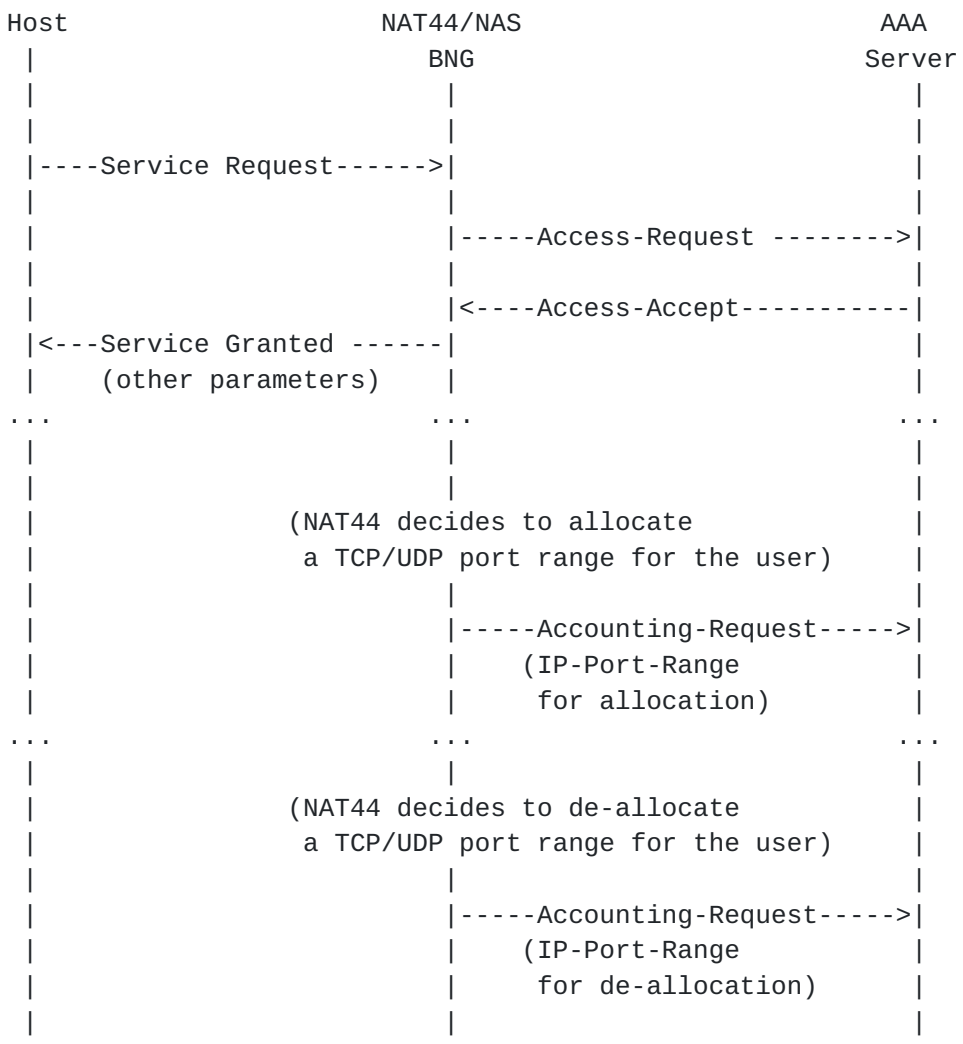

Figure 3: RADIUS Message Flow for reporting NAT44 allocation/de-allocation of a port set

4.1.3. Configure Forwarding Port Mapping

In most scenarios, the port mapping on a NAT device is dynamically created when the IP packets of an IP connection initiated by a user arrives. For some applications, the port mapping needs to be pre-defined allowing IP packets of applications from outside a CGN device to pass through and "port forwarded" to the correct user located behind the CGN device.

Port Control Protocol [RFC6887], provides a mechanism to create a mapping from an external IP address and port to an internal IP address and port on a CGN device just to achieve the "port forwarding" purpose. PCP is a server-client protocol capable of creating or deleting a mapping along with a rich set of features on a CGN device in dynamic fashion. In some deployment, all users need is

a few, typically just one pre-configured port mapping for applications such as web cam at home, and the lifetime of such a port mapping remains valid throughout the duration of the customer's Internet service connection time. In such an environment, it is possible to statically configure a port mapping on the RADIUS server for a user and let the RADIUS protocol to propagate the information to the associated CGN device.

Figure 4 illustrates how RADIUS protocol is used to configure a forwarding port mapping on a NAT44 device by using RADIUS protocol.

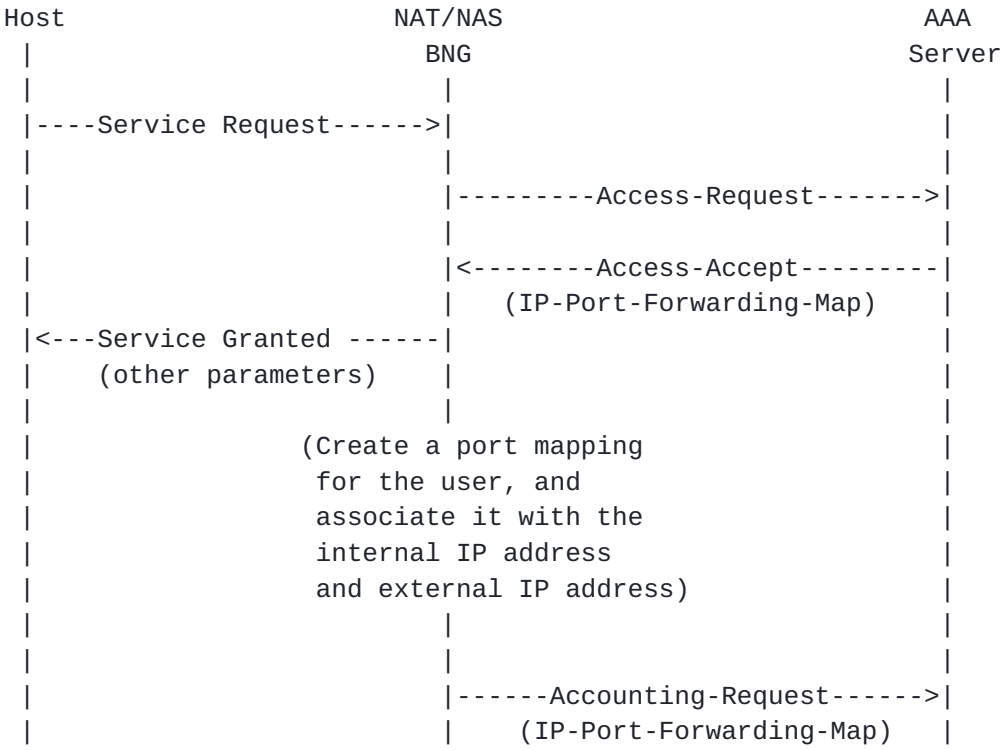

Figure 4: RADIUS Message Flow for configuring a forwarding port mapping

A port forwarding mapping that is created on a CGN device using RADIUS extension as described above may also be changed using RADIUS CoA message [RFC5176] that carries the same RADIUS associate. The CoA message may be sent from the RADIUS server directly to the NAS, which once accepts and sends back a RADIUS CoA ACK message, the new port forwarding mapping then replaces the previous one.

Figure 5 illustrates how RADIUS protocol is used to change an existing port mapping from (a:X) to (a:Y), where "a" is an internal port, and "X" and "Y" are external ports, respectively, for a specific user with a specific IP address

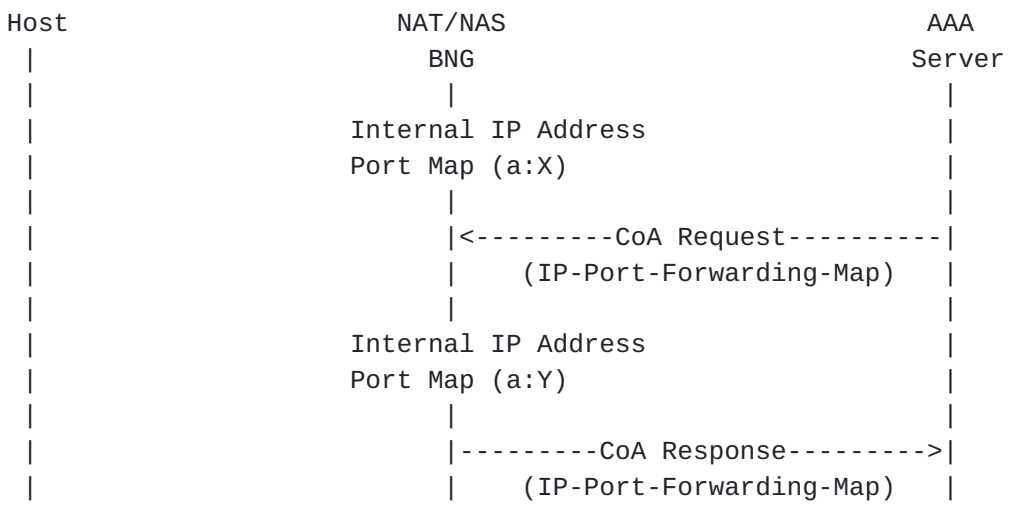

Figure 5: RADIUS Message Flow for changing a user's forwarding port mapping

4.1.4. An Example

An Internet Service Provider (ISP) assigns TCP/UDP 500 ports for the subscriber Joe. This number is the limit that can be used for TCP/UDP ports on a NAT44 device for Joe, and is configured on a RADIUS server. Also, Joe asks for a pre-defined port forwarding mapping on the NAT44 device for his web cam applications (external port 5000 maps to internal port 80).

When Joe successfully connects to the Internet service, the RADIUS server conveys the TCP/UDP port limit (1000) and the forwarding port mapping (external port 5000 to internal port 80) to the NAT44 device, using IP-Port-Limit attribute and IP-Port-Forwarding-Map attribute, respectively, carried by an Access-Accept message to the BNG where NAS and CGN co-located.

Upon receiving the first outbound IP packet sent from Joe's laptop, the NAT44 device decides to allocate a small port pool that contains 40 consecutive ports, from 3500 to 3540, inclusively, and also assign a shared IPv4 address 192.0.2.15, for Joe. The NAT44 device also randomly selects one port from the allocated range (say 3519) and use that port to replace the original source port in outbound IP packets.

For accounting purpose, the NAT44 device passes this port range (3500-3540) and the shared IPv4 address 192.0.2.15 together to the RADIUS server using IP-Port-Range attribute carried by an Accounting-Request message.

When Joe works on more applications with more outbound IP sessions and the port pool (3500-3540) is close to exhaust, the NAT44 device

allocates a second port pool (8500-8800) in a similar fashion, and also passes the new port range (8500-8800) and IPv4 address 192.0.2.15 together to the RADIUS server using IP-Port-Range attribute carried by an Accounting-Request message. Note when the CGN allocates more ports, it needs to assure that the total number of ports allocated for Joe is within the limit.

Joe decides to upgrade his service agreement with more TCP/UDP ports allowed (up to 1000 ports). The ISP updates the information in Joe's profile on the RADIUS server, which then sends a CoA-Request message that carries the IP-Port-Limit attribute with 1000 ports to the NAT44 device; the NAT44 device in turn sends back a CoA-ACK message. With that, Joe enjoys more available TCP/UDP ports for his applications.

When Joe travels, most of the IP sessions are closed with their associated TCP/UDP ports released on the NAT44 device, which then sends the relevant information back to the RADIUS server using IP-Port-Range attribute carried by Accounting-Request message.

Throughout Joe's connection with his ISP Internet service, applications can communicate with his web cam at home from external realm directly traversing the pre-configured mapping on the CGN device.

When Joe disconnects from his Internet service, the CGN device will de-allocate all TCP/UDP ports as well as the port-forwarding mapping, and send the relevant information to the RADIUS server.

4.2. Report Assigned Port Set for a Visiting UE

Figure 6 illustrates an example of the flow exchange which occurs when a visiting UE connects to a CPE offering Wi-Fi service.

For identification purposes (see [RFC6967]), once the CPE assigns a port set, it issues a RADIUS message to report the assigned port set.

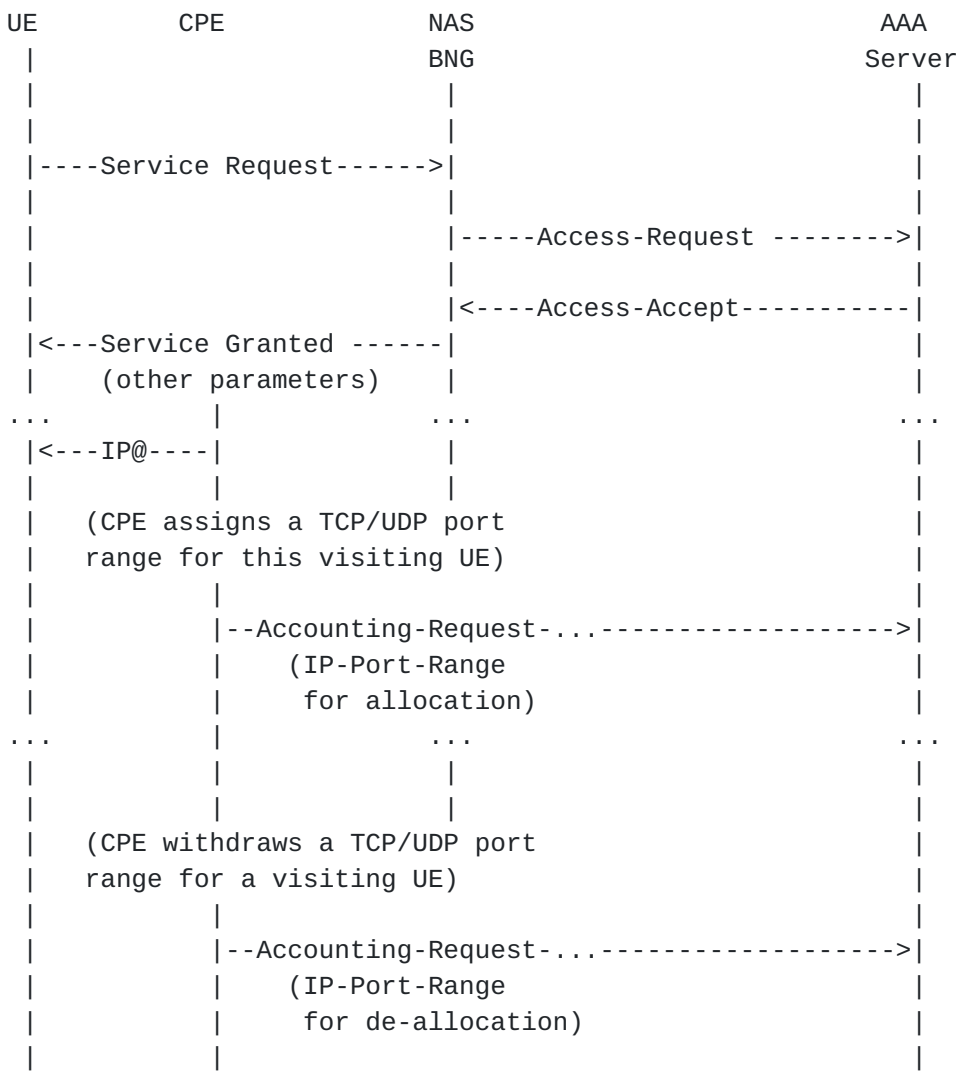

Figure 6: RADIUS Message Flow for reporting CPE allocation/de-allocation of a port set to a visiting UE

5. Table of Attributes

This document proposes three new RADIUS attributes and their formats are as follows:

- IP-Port-Limit: TBA1.{TBA!-1 .. TBA1-5}.TBA2

- IP-Port-Range: TBA1.{TBA1-1 .. TBA1-5}.{TBA3 .. TBA4}

- IP-Port-Forwarding-Map: TBA.1{TBA1-1 .. TBA1-5}.TBA5

The following table provides a guide as what type of RADIUS packets that may contain these attributes, and in what quantity.

Request	Accept	Reject	Challenge	Acct. Request	#	Attribute
0-1	0-1	0	0	0-1	TBA	IP-Port-Limit
0	0	0	0	0-1	TBA	IP-Port-Range
0-1	0-1	0	0	0-1	TBA	IP-Port-Forwarding-Map

The following table defines the meaning of the above table entries.

0 This attribute MUST NOT be present in packet.
0+ Zero or more instances of this attribute MAY be present in packet.
0-1 Zero or one instance of this attribute MAY be present in packet.

6. Security Considerations

This document does not introduce any security issue than what has been identified in [RFC2865].

7. IANA Considerations

This document requires new code point assignment for the new RADIUS attributes as follows:

- TBA1 (refer to Section 3.1): This value is for the Radius Type field and should be allocated from the number space of Extended-Type-1 (241), Extended-Type-2 (242), Extended-Type-3 (243), or Extended-Type-4 (244) per [RFC6929].

- TBA1-1, TBA1-2, TBA1-3, TBA1-4, and TBA1-5 (refer to Section 3.1): These values are for the Radius Extended Type field that are associated with TBA1.

- TBA2 (refer to Section 3.2): This value is for the TLV field and specifies the limit of the IP port imposed to a user.

- TBA3 (refer to Section 3.3): This value is for the TLV field and specifies the allocation action of IP ports by a port device (e.g., a CGN) for a user.

- TBA4 (refer to Section 3.3): This value is for the TLV field and specifies the de-allocation action of IP ports by a port device (e.g., a CGN) for a user.

- TBA5(refer to Section 3.4): This value is for the TLV field and specifies the mapping action on IP port by a port device (e.g., a CGN) for a user.

8. Acknowledgements

Many thanks to Dan Wing, Roberta Maglione, Daniel Derksen, David Thaler, Alan Dekok, and Lionel Morand for their useful comments and suggestions.

9. References

9.1. Normative References

[RFC1918] Rekhter, Y., Moskowitz, R., Karrenberg, D., Groot, G., and E. Lear, "Address Allocation for Private Internets", BCP 5, RFC 1918, February 1996.

[RFC2119] Bradner, S., "Key words for use in RFCs to Indicate Requirement Levels", BCP 14, RFC 2119, March 1997.

[RFC2629] Rose, M., "Writing I-Ds and RFCs using XML", RFC 2629, June 1999.

[RFC2865] Rigney, C., Willens, S., Rubens, A., and W. Simpson, "Remote Authentication Dial In User Service (RADIUS)", RFC 2865, June 2000.

[RFC5176] Chiba, M., Dommety, G., Eklund, M., Mitton, D., and B. Aboba, "Dynamic Authorization Extensions to Remote Authentication Dial In User Service (RADIUS)", RFC 5176, January 2008.

[RFC6929] DeKok, A. and A. Lior, "Remote Authentication Dial In User Service (RADIUS) Protocol Extensions", RFC 6929, April 2013.

9.2. Informative References

[I-D.gundavelli-v6ops-community-wifi-svcs]
Gundavelli, S., Grayson, M., Seite, P., and Y. Lee, "Service Provider Wi-Fi Services Over Residential Architectures", draft-gundavelli-v6ops-community-wifi-svcs-06 (work in progress), April 2013.

[I-D.ietf-softwire-lw4over6]
Cui, Y., Qiong, Q., Boucadair, M., Tsou, T., Lee, Y., and I. Farrer, "Lightweight 4over6: An Extension to the DS-Lite Architecture", draft-ietf-softwire-lw4over6-08 (work in progress), March 2014.

[I-D.miles-behave-l2nat]
Miles, D. and M. Townsley, "Layer2-Aware NAT", draft-miles-behave-l2nat-00 (work in progress), March 2009.

[RFC3022] Srisuresh, P. and K. Egevang, "Traditional IP Network Address Translator (Traditional NAT)", RFC 3022, January 2001.

[RFC6146] Bagnulo, M., Matthews, P., and I. van Beijnum, "Stateful NAT64: Network Address and Protocol Translation from IPv6 Clients to IPv4 Servers", RFC 6146, April 2011.

[RFC6269] Ford, M., Boucadair, M., Durand, A., Levis, P., and P. Roberts, "Issues with IP Address Sharing", RFC 6269, June 2011.

[RFC6333] Durand, A., Droms, R., Woodyatt, J., and Y. Lee, "Dual-Stack Lite Broadband Deployments Following IPv4 Exhaustion", RFC 6333, August 2011.

[RFC6619] Arkko, J., Eggert, L., and M. Townsley, "Scalable Operation of Address Translators with Per-Interface Bindings", RFC 6619, June 2012.

[RFC6887] Wing, D., Cheshire, S., Boucadair, M., Penno, R., and P. Selkirk, "Port Control Protocol (PCP)", RFC 6887, April 2013.

[RFC6888] Perreault, S., Yamagata, I., Miyakawa, S., Nakagawa, A., and H. Ashida, "Common Requirements for Carrier-Grade NATs (CGNs)", BCP 127, RFC 6888, April 2013.

[RFC6967] Boucadair, M., Touch, J., Levis, P., and R. Penno, "Analysis of Potential Solutions for Revealing a Host Identifier (HOST_ID) in Shared Address Deployments", RFC 6967, June 2013.

Authors' Addresses

Dean Cheng
Huawei
2330 Central Expressway
Santa Clara, California 95050
USA

Email: dean.cheng@huawei.com

Jouni Korhonen
Broadcom
Porkkalankatu 24
FIN-00180 Helsinki
Finland

Email: jouni.nospam@gmail.com

Mohamed Boucadair
France Telecom
Rennes
France

Email: mohamed.boucadair@orange.com

Senthil Sivakumar
Cisco Systems
7100-8 Kit Creek Road
Research Triangle Park, North Carolina
USA

Email: ssenthil@cisco.com