

Network Working Group
Internet-Draft
Intended status: Standards Track
Expires: December 14, 2014

D. Cheng
Huawei
J. Korhonen
Broadcom
M. Boucadair
France Telecom
S. Sivakumar
Cisco Systems
June 12, 2014

RADIUS Extensions for IP Port Configuration and Reporting
draft-ietf-radext-ip-port-radius-ext-01

Abstract

This document defines three new RADIUS attributes. For devices 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 (Carrier Grade NAT), NAT64, Provider WLAN 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 December 14, 2014.

Copyright Notice

Copyright (c) 2014 IETF Trust and the persons identified as the document authors. All rights reserved.

This document is subject to BCP 78 and the IETF Trust's Legal Provisions Relating to IETF Documents (<http://trustee.ietf.org/license-info>) in effect on the date of publication of this document. Please review these documents carefully, as they describe your rights and restrictions with respect to this document. Code Components extracted from this document must include Simplified BSD License text as described in Section 4.e of the Trust Legal Provisions and are provided without warranty as described in the Simplified BSD License.

Table of Contents

1. Introduction	3
2. Terminology	4
3. Extensions of RADIUS Attributes and TLVs	5
3.1. Extended Attributes for IP Ports	6
3.1.1. Extended-Type and IP-Port-Type TLV	6
3.1.2. IP-Port-Limit Attribute	7
3.1.3. IP-Port-Range Attribute	9
3.1.4. IP-Port-Forwarding-Map Attribute	12
3.2. RADIUS TLVs for IP Ports	14
3.2.1. IP-Port-Limit TLV	14
3.2.2. IP-Port-Ext-IPv4-Addr TLV	15
3.2.3. IP-Port-Int-IP-Addr TLV	16
3.2.4. IP-Port-Int-Port TLV	17
3.2.5. IP-Port-Ext-Port TLV	17
3.2.6. IP-Port-Alloc TLV	18
3.2.7. IP-Port-Range-Start TLV	19
3.2.8. IP-Port-Range-End TLV	20
3.2.9. IP-Port-Local-Id TLV	20
4. Applications, Use Cases and Examples	21
4.1. Managing CGN Port Behavior using RADIUS	21
4.1.1. Configure IP Port Limit for a User	22
4.1.2. Report IP Port Allocation/De-allocation	24
4.1.3. Configure Forwarding Port Mapping	25
4.1.4. An Example	27
4.2. Report Assigned Port Set for a Visiting UE	28
5. Table of Attributes	29
6. Security Considerations	30
7. IANA Considerations	30
8. Acknowledgements	31
9. References	31
9.1. Normative References	31

9.2. Informative References	32
Authors' Addresses	33

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 case, when IP packets traverse a CGN device, 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 be subject to port quota (see Section 5 of [RFC6269]).

The attributes defined in this document may also be used to report the assigned port range in some deployments such as Provider WLAN [I-D.gundavelli-v6ops-community-wifi-svcs]. For example, a visiting host can be managed by a CPE (Customer Premises Equipment) which will need to report the assigned port range 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:

1. IP-Port-Limit: This attribute may be carried in RADIUS Access-Accept, Access-Request, 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, associated with an IPv4 address.

2. IP-Port-Range: This attribute may be carried in RADIUS Accounting-Request packet. The purpose of this attribute is to report by an address sharing device (e.g., a CGN) to the RADIUS server the range of TCP/UDP ports and/or ICMP identifiers that have been allocated or deallocated associated with a given IPv4 address for a subscriber.
3. IP-Port-Forwarding-Map: This attribute may be carried in RADIUS Access-Accept, Access-Request, Accounting-Request or CoA-Request packet. The purpose of this attribute is to specify how a TCP/UDP port (or an ICMP identifier) mapping to another TCP/UDP port (or an ICMP identifier), and each is associated with its respective IPv4 address.

This document was constructed using the [RFC2629] .

2. Terminology

This document makes use if the following terms:

- o IP Port: refers to the port numbers of IP transport protocols, including TCP port, UDP port and ICMP identifier.
- o IP Port Type: refers to one of the following: (1)TCP/UDP port and ICMP identifier, (2)TCP port and UDP port, (3) TCP port, (4) UDP port, or (5)ICMP identifier.
- o IP Port Limit: denotes the maximum number of IP ports for a specific port type, that a device supporting port ranges can use when performing port number mapping for a specific user.
- o IP Port Range: specifies a set of contiguous IP ports, indicated by the smallest numerical number and the largest numerical number, inclusively.
- o Internal IP Address: refers to the IP address that is used as a source IP address in an outbound IP packet sent towards a device supporting port ranges in the internal realm. In the IPv4 case, it is typically a private address [RFC1918].
- o External IP Address: refers to 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 the IPv4 case, it is typically a global routable IP address.

- o 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.
- o 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.
- o External realm: refers to the networking segment where IPv4 public addresses are used in respective of the device supporting port ranges.
- o Internal realm: refers to the networking segment that is behind a device supporting port ranges and where IPv4 private addresses are used.
- o Mapping: 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.
- o Port-based device: a device that is capable of providing IP address and IP port mapping services and in particular, with the granularity of one or more subsets within the 16-bit IP port number range. A typical example of this device is a CGN, CPE, Provider WLAN 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. Extensions of RADIUS Attributes and TLVs

These three new attributes are defined in the following sub-sections:

1. IP-Port-Limit Attribute
2. IP-Port-Range Attribute
3. IP-Port-Forwarding-Map Attribute

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

3.1. Extended Attributes for IP Ports

3.1.1. Extended-Type and IP-Port-Type TLV

This section defines a new Extended-Type and an IP-Port-Type TLV (see Figure 1).

The IP port type may be one of the following:

- o TCP port, UDP port, and ICMP identifier
- o TCP port and UDP port
- o TCP port
- o UDP port
- o ICMP identifier




Figure 1

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 bytes of all fields this attribute, including the Type, Length, Extended-Type, and the embedded TLVs.

Extended-Type:

TBA2.

TLV1-Type:

Type field of IP-Port-Type TLV. This one byte field indicates the IP port type as follows:

TBA2-1:

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

TBA2-2:

Refer to TCP port and UDP port as a whole.

TBA2-3:

Refer to TCP port only.

TBA2-4:

Refer to UDP port only.

TBA2-5:

Refer to ICMP identifier only.

TLV1-Length:

Length field of IP-Port-Type TLV. This field indicates the total length in bytes of the TLV1, including the field of TLV1-Type, TLV1-Length, and the Value.

Value:

Value field of IP-Port-Type TLV. This field contains one or more TLVs, refer to Section 3.1.2, Section 3.1.3, Section 3.1.4 for details.

The interpretation of this field is determined by the identifier of "TBA1.TBA2.{TBA2-1..TBA2-5}" along with the embedded TLVs.

3.1.2. IP-Port-Limit Attribute

This attribute contains the Extended-Type and IP-Port-Type TLV defined in Section 3.1.1, along with the embedded IP-Port-Limit TLV and IP-Port-Ext-IPv4-Addr TLV, defined in Section 3.2.1 and Section 3.2.2, respectively. It specifies the maximum number of IP ports, as indicated in IP-Port-Limit TLV, of a specific port type, and associated with a given IPv4 address, as indicated in IP-Port-Ext-IPv4-Addr TLV for an end user. Note that when IP-Port-Ext-IPv4-Addr TLV is not included as part of the IP-Port-Limit Attribute,

the port limit is applied to all the IPv4 addresses managed by the port device, e.g., a CGN or NAT64 device.

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

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

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

The format of the IP-Port-Limit Attribute is shown in Figure 2. The fields are transmitted from left to right.




Figure 2

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 bytes of all fields of this attribute, including the Type, Length, Extended-Type, and the entire length of the embedded TLVs.

Extended-Type:

TBA2 - This one byte field contains a value that indicates the IP port type, refer to Section 3.1.1 for detail.

TLV1-Type:

TBA2-1, TBA2-2, TBA2-3, TBA2-4, or TBA2-5. Refer to Section 3.1.1 for detail.

TLV1-Length:

This field indicates the total length in bytes of the TLV1, including the field of TLV1-Type, TLV1-Length, and the entire length of the embedded TLVs.

Value:

This field contains a set of TLVs as follows:

IP-Port-Limit TLV:

This TLV contains the maximum number of IP ports of a specific IP port type and associated with a given IPv4 address for an end user. This TLV must be included in the IP-Port-Limit Attribute. Refer to Section 3.2.1.

IP-Port-Ext-IPv4-Addr TLV:

This TLV contains the IPv4 address that is associated with the IP port limit contained in the IP-Port-Limit TLV. This TLV is optionally included as part of the IP-Port-Limit Attribute. Refer to Section 3.2.2.

IP-Port-Limit attribute is associated with the following identifier:
Type(TBA1).Extended-Type(TBA2).IP-Port-Type TLV{TBA2-1..TBA2-5}.[IP-Port-Limit TLV(TBA3), {IP-Port-Ext-IPv4-Addr TLV (TBA4)}].

3.1.3. IP-Port-Range Attribute

This attribute contains the Extended-Type and IP-Port-Type TLV defined in Section 3.1.1, along with a set of embedded TLVs defined in Section 3.2.7 (IP-Port-Range-Start TLV), Section 3.2.8 (IP-Port-Range-End TLV), Section 3.2.6 (IP-Port-Alloc TLV), Section 3.2.2 (IP-Port-Ext-IPv4-Addr TLV), and Section 3.2.9 (IP-Port-Local-Id TLV). It contains a range of contiguous IP ports of a specific port type and associated with an IPv4 address that are either allocated or deallocated by a device for a given subscriber, and the information is intended to send to RADIUS server.

This attribute can be used to convey a single IP port number; in such case IP-Port-Range-Start and IP-Port-Range-End conveys the same value.

Within an IP-Port-Range Attribute, the IP-Port-Alloc TLV is always included. For port allocation, both IP-Port-Range-Start TLV and IP-Port-Range-End TLV must be included; for port deallocation, the inclusion of these two TLVs is optional and if not included, it implies that all ports that are previously allocated are now deallocated. Both IP-Port-Ext-IPv4-Addr TLV and IP-Port-Local-Id TLV are optional and if included, they are used by a port device (e.g., a CGN device) to identify the end user.

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

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

The format of the IP-Port-Range Attribute format is shown in Figure 3. The fields are transmitted from left to right.




Figure 3

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 bytes of all fields of this attribute, including the Type, Length, Extended-Type, and the entire length of the embedded TLVs.

Extended-Type:

TBA2 - This one byte field contains a value that indicates the IP port type, refer to Section 3.1.1 for detail.

TLV1-Type:

TBA2-1, TBA2-2, TBA2-3, TBA2-4, or TBA2-5. Refer to Section 3.1.1 for detail.

TLV1-Length:

This field indicates the total length in bytes of the TLV1, including the field of TLV1-Type, TLV1-Length, and the entire length of the embedded TLVs.

Value:

This field contains a set of TLVs as follows:

IP-Port-Alloc TLV:

This TLV contains a flag to indicate that the range of the specified IP ports for either allocation or deallocation. This TLV must be included as part of the IP-Port-Range Attribute. Refer to Section 3.2.6.

IP-Port-Range-Start TLV:

This TLV contains the smallest port number of a range of contiguous IP ports. To report the port allocation, this TLV must be included together with IP-Port-Range-End TLV as part of the IP-Port-Range Attribute. Refer to Section 3.2.7.

IP-Port-Range-End TLV:

This TLV contains the largest port number of a range of contiguous IP ports. To report the port allocation, this TLV must be included together with IP-Port-Range-Start TLV as part of the IP-Port-Range Attribute. Refer to Section 3.2.8.

IP-Port-Ext-IPv4-Addr TLV:

This TLV contains the IPv4 address that is associated with the IP port range, as collectively indicated in the IP-Port-Range-Start TLV and the IP-Port-Range-End TLV. This TLV is optionally included as part of the IP-Port-Range Attribute. Refer to Section 3.2.2.

IP-Port-Local-Id TLV:

This TLV 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. This TLV is

optionally included as part of the IP-Port-Range Attribute. Refer to Section 3.2.9.

The IP-Port-Range attribute is associated with the following identifier: Type(TBA1).Extended-Type(TBA2).IP-Port-Type TLV{TBA2-1..TBA2-5}.[IP-Port-Alloc TLV(TBA8), {IP-Port-Range-Start TLV (TBA9), IP-Port-Range-End TLV (TBA10)}, {IP-Port-Ext-IPv4-Addr TLV (TBA4)}, {IP-Port-Local-Id TLV (TBA11)}].

3.1.4. IP-Port-Forwarding-Map Attribute

This attribute contains the Extended-Type and IP-Port-Type TLV defined in Section 3.1.1, along with a set of embedded TLVs defined in Section 3.2.4 (IP-Port-Int-Port TLV), Section 3.2.5 (IP-Port-Ext-Port TLV), Section 3.2.3 (IP-Port-Int-IP-Addr TLV), Section 3.2.9 (IP-Port-Local-Id TLV) and Section 3.2.2 (IP-Port-Ext-IP-Addr TLV). The attribute contains a 2-byte IP internal port number that is associated with an internal IPv4 or IPv6 address, or a locally significant identifier at the customer site, and a 2-byte IP external port number that is associated with an external IPv4 address. The internal IPv4 or IPv6 address, or the local identifier must be included; the external IPv4 address may also be included.

The IP-Port-Forwarding-Map Attribute MAY appear in an Access-Accept packet. It MAY also appear in an Access-Request packet as a hint by the device supporting port mapping, 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-Forwarding-Map Attribute MAY appear in a CoA-Request packet.

The IP-Port-Forwarding-Map Attribute MAY also appear in an Accounting-Request packet.

The attribute MUST NOT appear in any other RADIUS packet.

The format of the IP-Port-Forwarding-Map Attribute is shown in Figure 4. The fields are transmitted from left to right.




Figure 4

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 bytes of all fields of this attribute, including the Type, Length, Extended-Type, and the entire length of the embedded TLVs.

Extended-Type:

This one byte field contains a value that indicates the IP port type, refer to Section 3.1.1 for details.

TLV1-Type:

TBA2-1, TBA2-2, TBA2-3, TBA2-4, or TBA2-5. Refer to Section 3.1.1 for detail.

TLV1-Length:

This field indicates the total length in bytes of the TLV1, including the field of TLV1-Type, TLV1-Length, and the entire length of the embedded TLVs.

Value:

This field contains a set of TLVs as follows:

IP-Port-Int-Port TLV:

This TLV contains an internal IP port number associated with an internal IPv4 or IPv6 address. This TLV must be included together with IP-Port-Ext-Port TLV as part of the IP-Port-Forwarding-Map attribute. Refer to Section 3.2.4.

IP-Port-Ext-Port TLV:

This TLV contains an external IP port number associated with an external IPv4 address. This TLV must be included together with IP-Port-Int-Port TLV as part of the IP-Port-Forwarding-Map attribute. Refer to Section 3.2.5.

IP-Port-Int-IP-Addr TLV:

This TLV contains an IPv4 or IPv6 address that is associated with the internal IP port number contained in the IP-Port-Int-Port TLV. Either this TLV or IP-Port-Local-Id TLV must be included as part of the IP-Port-Forwarding-Map Attribute. Refer to Section 3.2.3.

IP-Port-Local-Id TLV:

This TLV 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. Either this TLV or IP-Port-Int-IP-Addr TLV must be included as part of the IP-Port-Forwarding-Map Attribute. Refer to Section 3.2.9.

IP-Port-Ext-IPv4-Addr TLV:

This TLV contains an IPv4 address that is associated with the external IP port number contained in the IP-Port-Ext-Port TLV. This TLV may be included as part of the IP-Port-Forwarding-Map Attribute. Refer to Section 3.2.2.

The IP-Port-Forwarding-Map attribute is associated with the following identifier: Type(TBA1).Extended-Type(TBA2).IP-Port-Type TLV{TBA2-1..TBA2-5}.[IP-Port-Int-Port TLV(TBA6), IP-Port-Ext-Port TLV(TBA7), {IP-Port-Int-IP-Addr TLV (TBA5)}, {IP-Port-Ext-IPv4-Addr TLV (TBA4)}].

3.2. RADIUS TLVs for IP Ports

3.2.1. IP-Port-Limit TLV

This TLV (Figure 5) uses the format defined in [RFC6929]. Its Value field contains a 2-byte integer called IP-Port-Limit, which indicates the maximum number of ports of a specified IP-Port-Type and associated with a given IPv4 address assigned to a subscriber.

IP-Port-Limit TLV is included as part of the IP-Port-Limit Attribute (refer to Section 3.1.2).

Note that IP-Port-Limit TLV is embedded within IP-Port-Type TLV (refer to Section 3.1.1) for detail.




Figure 5

TLV2-Type:

TBA3: The type field for IP-Port-Limit TLV.

TLV2-Length:

This field indicates the total length in bytes of the TLV2, including the field of TLV2-Type, TLV2-Length, and the Value field, i.e., IP-Port-Limit.

IP-Port-Limit:

2-byte integer. This field contains the maximum number of IP ports of which, the port type is specified by container IP-Port-Type TLV.

3.2.2. IP-Port-Ext-IPv4-Addr TLV

This TLV (Figure 6) uses the format defined in[RFC6929]. Its Value field contains a 4-byte External IPv4 address.

IP-Port-Ext-IPv4-Addr TLV can be included as part of the IP-Port-Limit Attribute (refer to Section 3.1.2), IP-Port-Range Attribute (refer to Section 3.1.3), and IP-Port-Forwarding-Map Attribute (refer to Section 3.1.4).

Note that IP-Port-Ext-IPv4-Addr TLV is embedded within IP-Port-Type TLV (refer to Section 3.1.1) for detail.




Figure 6

TLV3-Type:

TBA4: The type field for IP-Port-IPv4-Addr TLV.

TLV3-Length:

6. The Length field for IP-Port-IPv4-Addr TLV.

IP-Port-Ext-IPv4-Addr:

4-byte integer. This field contains the IPv4 address that is associated with the range of IP ports.

3.2.3. IP-Port-Int-IP-Addr TLV

This TLV (Figure 7) uses format defined in [RFC6929]. Its Value field contains an internal IPv4 or IPv6 address.

IP-Port-Int-IP-Addr TLV can be included as part of the IP-Port-Forwarding-Map Attribute (refer to Section 3.1.4).

Note that IP-Port-Int-IP-Addr TLV is embedded within IP-Port-Type TLV (refer to Section 3.1.1) for detail.




Figure 7

TLV4-Type:

TBA5: The type field for IP-Port-Int-IP-Addr TLV.

TLV4-Length:

6 or 18 bytes. The Length field for IP-Port-Int-IP-Addr TLV.

IP-Port-Int-IP-Addr:

4 byte integer for IPv4 address or 16 byte for IPv6 address.

3.2.4. IP-Port-Int-Port TLV

This TLV (Figure 8) uses format defined in [RFC6929]. Its Value field contains an internal IP port number that is associated with an internal IPv4 or IPv6 address.

IP-Port-Int-Port TLV is included as part of the IP-Port-Forwarding-Map Attribute (refer to Section 3.1.4).

IP-Port-Int-Port TLV is embedded within embedded within IP-Port-Type TLV (refer to Section 3.1.1) for detail.




Figure 8

TLV5-Type:

TBA6: The type field for IP-Port-Int-Port TLV.

TLV5-Length:

4 bytes. The Length field for IP-Port-Int-Port TLV.

IP-Port-Int-Port:

2 byte integer. The internal IP port number that is associated with an IPv4 or IPv6 address.

3.2.5. IP-Port-Ext-Port TLV

This TLV (Figure 9) uses format defined in [RFC6929]. Its Value field contains an external IP port number that is associated with an external IPv4 address.

IP-Port-Ext-Port TLV is included as part of the IP-Port-Forwarding-Map Attribute (refer to Section 3.1.4).

IP-Port-Ext-Port TLV is embedded within IP-Port-Type TLV (refer to Section 3.1.1) for detail.




Figure 9

TLV6-Type:

TBA7: The type field for IP-Port-Ext-Port TLV.

TLV6-Length:

4 bytes. The Length field for IP-Port-Ext-Port TLV.

IP-Port-Ext-Port:

2 byte integer. The external IP port number that is associated with an IPv4 address.

3.2.6. IP-Port-Alloc TLV

This TLV (Figure 10) uses format defined in [RFC6929]. Its Value field contains a 2-byte integer called IP-Port-Alloc, which indicates either the allocation or deallocation of a range of IP ports.

IP-Port-Alloc TLV is included as part of the IP-Port-Range Attribute (refer to Section 3.1.3).

Note that IP-Port-Alloc TLV is embedded within IP-Port-Type TLV (refer to Section 3.1.1) for detail.




Figure 10

TLV7-Type:

TBA8: The type field for IP-Port-Alloc TLV.

TLV7-Length:

- 4. The Length field for IP-Port-Alloc TLV.

IP-Port-Alloc:

2-byte integer. This field indicates the allocation or deallocation of a range of IP ports as follows:

0:

Allocation

1:

Deallocation

3.2.7. IP-Port-Range-Start TLV

This TLV (Figure 11) uses format defined in [RFC6929]. Its Value field contains a 2-byte integer called IP-Port-Range-Start, which indicates the smallest port number of a range of contiguous IP ports.

IP-Port-Range-Start TLV is included as part of the IP-Port-Range Attribute (refer to Section 3.1.3).

Note that IP-Port-Range-Start TLV is embedded within IP-Port-Type TLV (refer to Section 3.1.1) for detail.




Figure 11

TLV8-Type:

TBA9: The type field for IP-Port-Range-Start TLV.

TLV8-Length:

- 4. The Length field for IP-Port-Range-Start TLV.

IP-Port-Range-Start:

2-byte integer. This field contains the smallest port number of a range of contiguous IP ports.

3.2.8. IP-Port-Range-End TLV

This TLV (Figure 12) uses format defined in [RFC6929]. Its Value field contains a 2-byte integer called IP-Port-Range-End, which indicates largest port number of a range of contiguous IP ports.

IP-Port-Range-End TLV is included as part of the IP-Port-Range Attribute (refer to Section 3.1.3).

Note that IP-Port-Range-End TLV is embedded within IP-Port-Type TLV (refer to Section 3.1.1) for detail.




Figure 12

TLV9-Type:

TBA10: The type field for IP-Port-Range-End TLV.

TLV9-Length:

4. The Length field for IP-Port-Range-End TLV.

IP-Port-Range-End:

2-byte integer. This field contains the largest port number of a range of contiguous IP ports.

3.2.9. IP-Port-Local-Id TLV

This TLV (Figure 13) uses format defined in [RFC6929]. Its Value field contains an identifier with local significance.

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.

IP-Port-Local-Id TLV can be included as part of the IP-Port-Range Attribute (refer to Section 3.1.3) and IP-Port-Forwarding-Map Attribute (refer to Section 3.1.4).

Note that IP-Port-Local-Id TLV is embedded within IP-Port-Type TLV (refer to Section 3.1.1) for detail.




Figure 13

TLV10-Type:

TBA11: The type field for IP-Port-Local-Id TLV.

TLV10-Length:

Variable number of bytes. The Length field for IP-Port-Local-Id TLV.

IP-Port-Local-Id:

This is 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 is the value contained in TLV7-Length field minus 2.

4. Applications, Use Cases and Examples

This section describes some applications and use cases to illustrate the use of the attributes proposed 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.

4.1.1.1. 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 3.1.2.

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 3.1.2), 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 14 illustrates how RADIUS protocol is used to configure the maximum number of TCP/UDP ports for a given subscriber on a NAT44 device.




Figure 14: 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 15 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 15: 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 3.1.3). 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 16 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 16: 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 17 illustrates how RADIUS protocol is used to configure a forwarding port mapping on a NAT44 device by using RADIUS protocol.




Figure 17: 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 18 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 18: 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 19 illustrates an example of the flow exchange which occurs when a visiting UE connects to a CPE offering WLAN 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 19: 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:

- o IP-Port-Limit: TBA1.TBA2.{TBA2-1..TBA2-5}.[TBA3, {TBA4}]
- o IP-Port-Range: TBA1.TBA2.{TBA2-1..TBA2-5}.[TBA8, TBA9, TBA10, {TBA4}, {TBA11}].
- o IP-Port-Forwarding-Map: TBA1.TBA2.{TBA2-1 .. TBA2-5}.[TBA6, TBA7, TBA5, {TBA4}]

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+	0+	0	0	0+	TBA	IP-Port-Limit
0	0	0	0	0+	TBA	IP-Port-Range
0+	0+	0	0	0+	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.

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 assignments for the new RADIUS attributes as follows:

- o TBA1 (refer to Section 3.1.1): 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].
- o TBA2 (refer to Section 3.1.1): This value is for the Extended-Type field and should be allocated from the Short Extended Space per [RFC6929].
- o TBA2-1, TBA2-2, TBA2-3, TBA2-4, and TBA2-5 (refer to Section 3.1.1): These values are for the Type field of IP-Port-Type TLV that is within the TBA2 container, and they should be allocated as TLV data type and effectively extend the attribute tree as TBA1.TBA2.{TBA2-1, TBA2-2, TBA2-3, TBA2-4, TBA2-5}.
- o TBA3 (refer to Section 3.1.2): This value is for the type field of IP-Port-Limit TLV. It should be allocated as TLV data type and it extends the attribute tree as TBA1.TBA2.{TBA2-1, TBA2-2, TBA2-3, TBA2-4, TBA2-5}.TBA3.
- o TBA4 (refer to Section 3.2.2): This value is for the Type field of IP-Port-Ext-IPv4-Addr TLV. It should be allocated as TLV data type and it extends the attribute tree as TBA1.TBA2.{TBA2-1, TBA2-2, TBA2-3, TBA2-4, TBA2-5}.[TBA4...].

- o TBA5 (refer to Section 3.2.3): This value is for the Type field of IP-Port-Int-IP-Addr TLV. It should be allocated as TLV data type and it extends the attribute tree as TBA1.TBA2.{TBA2-1, TBA2-2, TBA2-3, TBA2-4, TBA2-5}.[TBA5...].
- o TBA6 (refer to Section 3.2.4): This value is for the Type field of IP-Port-Int-Port TLV. It should be allocated as TLV data type and it extends the attribute tree as TBA1.TBA2.{TBA2-1, TBA2-2, TBA2-3, TBA2-4, TBA2-5}.[TBA6...].
- o TBA7 (refer to Section 3.2.5): This value is for the Type field of IP-Port-Ext-port TLV. It should be allocated as TLV data type and it extends the attribute tree as TBA1.TBA2.{TBA2-1, TBA2-2, TBA2-3, TBA2-4, TBA2-5}.[TBA7...].
- o TBA8 (refer to Section 3.2.6): This value is for the Type field of IP-Port-Alloc TLV. It should be allocated as TLV data type and it extends the attribute tree as TBA1.TBA2.{TBA2-1, TBA2-2, TBA2-3, TBA2-4, TBA2-5}.[TBA8...].
- o TBA9 (refer to Section 3.2.7): This value is for the Type field of IP-Port-Range-Start TLV. It should be allocated as TLV data type and it extends the attribute tree as TBA1.TBA2.{TBA2-1, TBA2-2, TBA2-3, TBA2-4, TBA2-5}.[TBA9...].
- o TBA10 (refer to Section 3.2.8): This value is for the Type field of IP-Port-Range-End TLV. It should be allocated as TLV data type and it extends the attribute tree as TBA1.TBA2.{TBA2-1, TBA2-2, TBA2-3, TBA2-4, TBA2-5}.[TBA10..].
- o TBA11 (refer to Section 3.2.9): This value is for the Type field of IP-Port-Local-Id TLV. It should be allocated as TLV data type and it extends the attribute tree as TBA1.TBA2.{TBA2-1, TBA2-2, TBA2-3, TBA2-4, TBA2-5}.[TBA11..].

8. Acknowledgements

Many thanks to Dan Wing, Roberta Maglione, Daniel Derksen, David Thaler, Alan Dekok, Lionel Morand, and Peter Deacon 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-lw4over6-10 (work in progress), June 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