

Independent Submission
Request for Comments: 6431
Category: Informational
ISSN: 2070-1721

M. Boucadair
P. Levis
France Telecom
G. Bajko
T. Savolainen
Nokia
T. Tsou
Huawei Technologies (USA)
November 2011

Huawei Port Range Configuration Options for PPP
IP Control Protocol (IPCP)

Abstract

This document defines two Huawei IPCP (IP Control Protocol) options used to convey a set of ports. These options can be used in the context of port range-based solutions or NAT-based solutions for port delegation and forwarding purposes.

Status of This Memo

This document is not an Internet Standards Track specification; it is published for informational purposes.

This is a contribution to the RFC Series, independently of any other RFC stream. The RFC Editor has chosen to publish this document at its discretion and makes no statement about its value for implementation or deployment. Documents approved for publication by the RFC Editor are not a candidate for any level of Internet Standard; see Section 2 of RFC 5741.

Information about the current status of this document, any errata, and how to provide feedback on it may be obtained at <http://www.rfc-editor.org/info/rfc6431>.

Copyright Notice

Copyright (c) 2011 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.

Table of Contents

1. Introduction	2
1.1. Use Cases	3
1.2. Terminology	3
1.3. Requirements Language	4
2. Port Range Options	4
2.1. Description of Port Range Value and Port Range Mask	4
2.2. Cryptographically Random Port Range Option	6
2.2.1. Random Port Delegation Function	6
2.2.2. Description of Cryptographically Random Port Range Option	8
2.3. Illustration Examples	10
2.3.1. Overview	10
2.3.2. Successful Flow: Port Range Options Supported by Both the Client and the Server	10
2.3.3. Port Range Option Not Supported by the Server	11
2.3.4. Port Range Option Not Supported by the Client	13
3. Security Considerations	14
4. Contributors	14
5. Acknowledgements	14
6. References	14
6.1. Normative References	14
6.2. Informative References	15

1. Introduction

Within the context of IPv4 address depletion, several solutions have been investigated to share IPv4 addresses. Two flavors can be distinguished: NAT-based solutions (e.g., Carrier-Grade NAT (CGN) [CGN-REQS]) and port range-based solutions (e.g., [RFC6346] [PORT-RANGE-ARCH] [SAM]). Port range-based solutions do not require an additional NAT level in the service provider's domain. Several means may be used to convey port range information.

This document defines the notion of "Port Mask", which is generic and flexible. Several allocation schemes may be implemented when using a Port Mask. It proposes a basic mechanism that allows the allocation of a unique port range to a requesting client. This document defines Huawei IPCP options to be used to carry port range information.

IPv4 address exhaustion is only provided as an example of the usage of the PPP IPCP options defined in this document. In particular, Port Range options may be used independently of the presence of the IP-Address IPCP Option.

This document adheres to the considerations defined in [RFC2153].

This document is not a product of the PPPEXT working group.

Note that IPR disclosures apply to this document (see <https://datatracker.ietf.org/ipr/>).

1.1. Use Cases

Port Range options can be used in port range-based solutions (e.g., [RFC6346]) or in a CGN-based solution. These options can be used in a CGN context to bypass the NAT (i.e., for transparent NAT traversal, and to avoid involving several NAT levels in the path) or to delegate one or a set of ports to the requesting client (e.g., to avoid the ALG (Application Level Gateway), or for port forwarding).

Section 3.3.1 of [RFC6346] specifies an example of usage of the options defined in this document.

1.2. Terminology

To differentiate between a port range containing a contiguous span of port numbers and a port range with non-contiguous and possibly random port numbers, the following denominations are used:

- o Contiguous Port Range: A set of port values that form a contiguous sequence.
- o Non-Contiguous Port Range: A set of port values that do not form a contiguous sequence.
- o Random Port Range: A cryptographically random set of port values.

Unless explicitly mentioned, "Port Mask" refers to the tuple (Port Range Value, Port Range Mask).

In addition, this document makes use of the following terms:

- o Delegated port or delegated port range: A port or a range of ports that belong to an IP address managed by an upstream device (such as NAT) and that are delegated to a client for use as the source address and port when sending packets.
- o Forwarded port or forwarder port range: A port or a range of ports that belong to an IP address managed by an upstream device such as (NAT) and that are statically mapped to the internal IP address of the client and same port number of the client.

This memo uses the same terminology as [RFC1661].

1.3. 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].

2. Port Range Options

This section defines the IPCP Option for port range delegation. The format of vendor-specific options is defined in [RFC2153]. Below are the values to be conveyed when the Port Range Option is used:

- o Organizationally Unique Identifier (OUI): This field is set to 781DBA (hex).
- o Kind: This field is set to F0 (hex).
- o Value(s): The content of this field is specified in Sections 2.1 and 2.2.2.

2.1. Description of Port Range Value and Port Range Mask

The Port Range Value and Port Range Mask are used to specify one range of ports (contiguous or non-contiguous) pertaining to a given IP address. Concretely, the Port Range Mask and Port Range Value are used to notify a remote peer about the Port Mask to be applied when selecting a port value as a source port. The Port Range Value is used to infer a set of allowed port values. A Port Range Mask defines a set of ports that all have in common a subset of pre-positioned bits. This set of ports is also referred to as the port range.

Two port numbers are said to belong to the same port range if and only if they have the same Port Range Mask.

A Port Mask is composed of a Port Range Value and a Port Range Mask:

- o The Port Range Value indicates the value of the significant bits of the Port Mask. The Port Range Value is coded as follows:
 - * The significant bits may take a value of 0 or 1.
 - * All of the other bits (i.e., non-significant ones) are set to 0.
- o The Port Range Mask indicates, by the bit(s) set to 1, the position of the significant bits of the Port Range Value.


This IPCP Configuration Option provides a way to negotiate the Port Range to be used on the local end of the link. It allows the sender of the Configure-Request message to state which port range associated with a given IP address is desired, or to request that the peer provide the configuration. The peer can provide this information by NAKing the option, and returning a valid port range (i.e., (Port Range Value, Port Range Mask)).

If a peer issues a request enclosing the IPCP Port Range Option and the server does not support this option, the Port Range Option is rejected by the server.

The set of ports conveyed in an IPCP Port Range Option applies to all transport protocols.

The set of ports conveyed in an IPCP Port Range Option is revoked when the link is no longer up (e.g., when Terminate-Request and Terminate-Ack are exchanged).

The Port Range IPCP option adheres to the format defined in Section 2.1 of [RFC2153]. The "Value(s)" field of the option defined in [RFC2153] when conveying the Port Range IPCP Option is provided in Figure 1.



Most significant bit (MSB) network order is used for encoding the Port Range Value and Port Range Mask fields.

Figure 1: Format of the Port Range IPCP Option

- o M: mode bit. The mode bit indicates the mode for which the port range is allocated. A value of zero indicates that the port ranges are delegated, while a value of 1 indicates that the port ranges are port-forwarded.
- o Port Range Value (PRV): The PRV indicates the value of the significant bits of the Port Mask. By default, no PRV is assigned.

- o Port Range Mask (PRM): The Port Range Mask indicates the position of the bits that are used to build the Port Range Value. By default, no PRM value is assigned. The 1 values in the Port Range Mask indicate by their position the significant bits of the Port Range Value.

Figure 2 provides an example of the resulting port range:

- The Port Range Mask is set to 0001010000000000 (5120).
- The Port Range Value is set to 0000010000000000 (1024).

```

0                                     1
0 1 2 3 4 5 6 7 8 9 0 1 2 3 4 5
+-----+
| 0 0 0 1 0 1 0 0 0 0 0 0 0 0 0 0 | Port Range Mask
+-----+
      |      |
      v      v (two significant bits)
+-----+
| 0 0 0 0 0 1 0 0 0 0 0 0 0 0 0 0 | Port Range Value
+-----+

+-----+
| x x x 0 x 1 x x x x x x x x x x | Usable ports
+-----+ (x may be set to 0 or 1)

```

Figure 2: Example of Port Range Mask and Port Range Value

Port values belonging to this port range must have the fourth bit from the left set to 0, and the sixth bit from the left set to 1. Only these port values will be used by the peer when enforcing the configuration conveyed by PPP IPCP.

2.2. Cryptographically Random Port Range Option

A cryptographically random Port Range Option may be used as a mitigation tool against blind attacks such as those described in [RFC6056].

2.2.1. Random Port Delegation Function

Delegating random ports can be achieved by defining a function that takes as input a key 'K' and an integer 'x' within the 1024-65535 port range and produces an output 'y' also within the 1024-65535 port range.

The cryptographic mechanism uses the 1024-65535 port range rather than the ephemeral range, 49152-65535, for generating a set of ports to optimize IPv4 address utilization efficiency (see "Appendix B. Address Space Multiplicative Factor" of [RFC6269]). This behavior is compliant with the recommendation to use the whole 1024-65535 port range for the ephemeral port selection algorithms (see Section 3.2 of [RFC6056]).

The cryptographic mechanism ensures that the entire 64k port range can be efficiently distributed to multiple nodes such that when nodes calculate the ports, the results will never overlap with ports that other nodes have calculated (property of permutation), and ports in the reserved range (smaller than 1024) are not used. As the randomization is done cryptographically, an attacker seeing a node using some port X cannot determine which other ports the node may be using (as the attacker does not know the key). Calculation of the random port list is done as follows:

The cryptographic mechanism uses an encryption function $y = E(K, x)$ that takes as input a key K (for example, 128 bits) and an integer x (the plaintext) in the 1024-65535 port range, and produces an output y (the ciphertext), also an integer in the 1024-65535 port range. This section describes one such encryption function, but others are also possible.

The server will select the key K. When the server wants to allocate, for example, 2048 random ports, it selects a starting point 'a' ($1024 \leq a \leq 65536 - 2048$) such that the port range (a, a+2048) does not overlap with any other active client, and calculates the values $E(K, a)$, $E(K, a+1)$, $E(K, a+2)$, ..., $E(K, a+2046)$, $E(K, a+2047)$. These are the port numbers allocated for this node. Instead of sending the port numbers individually, the server just sends the values 'K', 'a', and '2048'. The client will then repeat the same calculation.

The server SHOULD use a different key K for each IPv4 address it allocates, to make attacks as difficult as possible. This way, learning the key K used in IPv4 address IP1 would not help in attacking IPv4 address IP2 where IP2 is allocated by the same server to different nodes.

With typical encryption functions (such as AES and DES), the input (plaintext) and output (ciphertext) are blocks of some fixed size -- for example, 128 bits for AES, and 64 bits for DES. For port randomization, we need an encryption function whose input and output is an integer in the 1024-65535 port range.

One possible way to do this is to use the 'Generalized Feistel Cipher' [CIPHERS] construction by Black and Rogaway, with AES as the underlying round function.

This would look as follows (using pseudo-code):

```
def E(k, x):
    y = Feistel16(k, x)
    if y >= 1024:
        return y
    else:
        return E(k, y)
```

Note that although E(k,x) is recursive, it is guaranteed to terminate. The average number of iterations is just slightly over 1.

Feistel16 is a 16-bit block cipher:

```
def Feistel16(k, x):
    left = x & 0xff
    right = x >> 8
    for round = 1 to 3:
        temp = left ^ FeistelRound(k, round, right)
        left = right
        right = temp
    return (right << 8) | left
```

The Feistel round function uses:

```
def FeistelRound(k, round, x):
    msg[0] = round
    msg[1] = x
    msg[2...15] = 0
    return AES(k, msg)[0]
```

Performance: To generate a list of 2048 port numbers, about 6000 calls to AES are required (i.e., encrypting 96 kilobytes). Thus, it will not be a problem for any device that can do, for example, HTTPS (web browsing over Secure Sockets Layer/Transport Layer Security (SSL/TLS)).

2.2.2. Description of Cryptographically Random Port Range Option

The cryptographically random Port Range IPCP Option adheres to the format defined in Section 2.1 of [RFC2153]. The "Value(s)" field of the option defined in [RFC2153] when conveying the cryptographically random Port Range IPCP Option is illustrated in Figure 3.

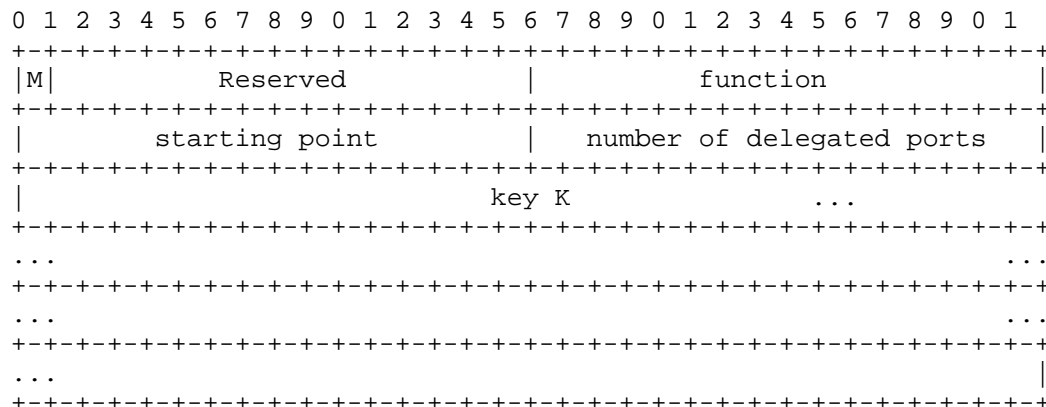


Figure 3: Format of the Cryptographically Random Port Range Option

- o M: mode bit. The mode bit indicates the mode for which the port range is allocated. A value of zero indicates that the port ranges are delegated, while a value of 1 indicates that the port ranges are port-forwarded.
- o Function: A 16-bit field whose value is associated with predefined encryption functions. This specification associates value 1 with the predefined function described in Section 2.2.1.
- o Starting Point: A 16-bit value used as an input to the specified function.
- o Number of delegated ports: A 16-bit value specifying the number of ports delegated to the client for use as source port values.
- o Key K: A 128-bit key used as input to the predefined function for delegated port calculation.

When the option is included in the IPCP Configure-Request, the "Key K" and "Starting Point" fields SHALL be set to all zeros. The requester MAY indicate in the "Function" field which encryption function the requester prefers, and in the "Number of Delegated Ports" field the number of ports the requester would like to obtain. If the requester has no preference, it SHALL also set the "Function" field and/or "Number of Delegated Ports" field to zero.

The usage of the option in IPCP message negotiation (Request/Reject/Nak/Ack) follows the logic described for Port Mask and Port Range options in Section 2.1.

2.3. Illustration Examples

2.3.1. Overview

The following flows provide examples of the usage of IPCP to convey the Port Range Option. As illustrated in Figures 4, 5, and 6, IPCP messages are exchanged between a Host and a BRAS (Broadband Remote Access Server).

2.3.2. Successful Flow: Port Range Options Supported by Both the Client and the Server

The following message exchange (Figure 4) depicts a successful IPCP configuration operation where the Port Range IPCP Option is used.

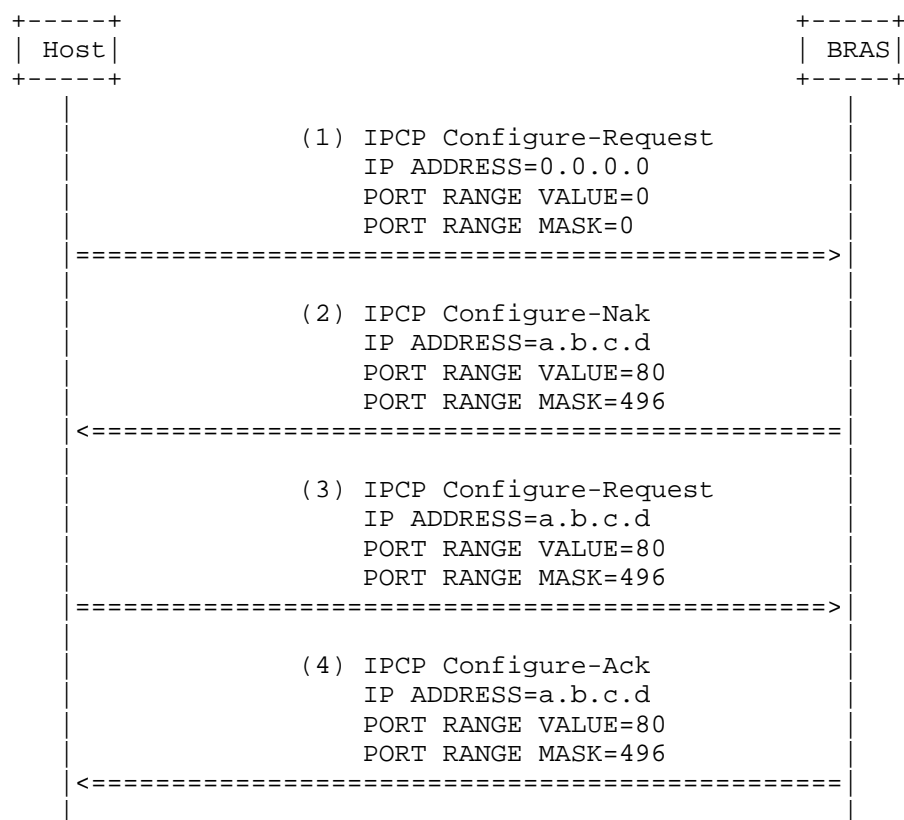


Figure 4: Successful Flow

The main steps of this flow are listed below:

- (1) The Host sends a first Configure-Request, which includes the set of options it desires to negotiate. All of these configuration options are negotiated simultaneously. In this step, the Configure-Request carries information about the IP address, the Port Range Value, and the Port Range Mask. The IP-Address Option is set to 0.0.0.0, the Port Range Value is set to 0, and the Port Range Mask is set to 0.
- (2) The BRAS sends back a Configure-Nak and sets the enclosed options to its preferred values. In this step, the IP-Address Option is set to a.b.c.d, the Port Range Value is set to 80, and the Port Range Mask is set to 496.
- (3) The Host re-sends a Configure-Request requesting that the IP-Address Option be set to a.b.c.d, the Port Range Value be set to 80, and the Port Range Mask be set to 496.
- (4) The BRAS sends a Configure-Ack message.

As a result of this exchange, the Host is configured to use a.b.c.d as its local IP address, and the following 128 contiguous port ranges resulting from the Port Mask (Port Range Value == 0, Port Range Mask == 496):

- from 80 to 95
- from 592 to 607
- ...
- from 65104 to 65119

2.3.3. Port Range Option Not Supported by the Server

Figure 5 depicts an exchange of messages where the BRAS does not support the IPCP Port Range Option.




Figure 5: Failed Flow: Port Range Option Not Supported by the Server

The main steps of this flow are listed below:

- (1) The Host sends a first Configure-Request, which includes the set of options it desires to negotiate. All of these configuration options are negotiated simultaneously. In this step, the Configure-Request carries the codes of the IP-Address, Port Range Value, and Port Range Mask options. The IP-Address Option is set to 0.0.0.0, the Port Range Value is set to 0, and the Port Range Mask is set to 0.
- (2) The BRAS sends back a Configure-Reject to decline the Port Range Option.

- (3) The Host sends a Configure-Request, which includes only the codes of the IP-Address Option. In this step, the IP-Address Option is set to 0.0.0.0.
- (4) The BRAS sends back a Configure-Nak and sets the enclosed option to its preferred value. In this step, the IP-Address Option is set to a.b.c.d.
- (5) The Host re-sends a Configure-Request requesting that the IP-Address Option be set to a.b.c.d.
- (6) The BRAS sends a Configure-Ack message.

As a result of this exchange, the Host is configured to use a.b.c.d as its local IP address. This IP address is not a shared IP address.

2.3.4. Port Range Option Not Supported by the Client

Figure 6 depicts exchanges where only shared IP addresses are assigned to end-users' devices. The server is configured to assign only shared IP addresses. If Port Range options are not enclosed in the configuration request, the request is rejected, and the requesting peer will be unable to access the service.




Figure 6: Port Range Option Not Supported by the Client

The main steps of this flow are listed below:

- (1) The Host sends a Configure-Request requesting that the IP-Address Option be set to 0.0.0.0, and without enclosing the Port Range Option.
- (2) The BRAS sends a Protocol-Reject message.

As a result of this exchange, the Host is not able to access the service.

3. Security Considerations

This document does not introduce any security issues in addition to those related to PPP. Service providers should use authentication mechanisms such as the Challenge Handshake Authentication Protocol (CHAP) [RFC1994] or PPP link encryption [RFC1968].

The use of small and non-random port ranges may increase host exposure to attacks, as described in [RFC6056]. This risk can be reduced by using larger port ranges, by using the random Port Range Option, or by activating means to improve the robustness of TCP against blind in-window attacks [RFC5961].

4. Contributors

Jean-Luc Grimault and Alain Villefranque contributed to this document.

5. Acknowledgements

The authors would like to thank C. Jacquenet, J. Carlson, B. Carpenter, M. Townsley, and J. Arkko for their review.

6. References

6.1. Normative References

- [RFC1661] Simpson, W., Ed., "The Point-to-Point Protocol (PPP)", STD 51, RFC 1661, July 1994.
- [RFC1968] Meyer, G., "The PPP Encryption Control Protocol (ECP)", RFC 1968, June 1996.
- [RFC1994] Simpson, W., "PPP Challenge Handshake Authentication Protocol (CHAP)", RFC 1994, August 1996.
- [RFC2119] Bradner, S., "Key words for use in RFCs to Indicate Requirement Levels", BCP 14, RFC 2119, March 1997.
- [RFC2153] Simpson, W., "PPP Vendor Extensions", RFC 2153, May 1997.
- [RFC5961] Ramaiah, A., Stewart, R., and M. Dalal, "Improving TCP's Robustness to Blind In-Window Attacks", RFC 5961, August 2010.

6.2. Informative References

- [CGN-REQS] Perreault, S., Ed., Yamagata, I., Miyakawa, S., Nakagawa, A., and H. Ashida, "Common requirements for Carrier Grade NAT (CGN)", Work in Progress, October 2011.
- [CIPHERS] Black, J. and P. Rogaway, "Ciphers with Arbitrary Finite Domains. Topics in Cryptology", CT-RSA 2002, Lecture Notes in Computer Science, vol. 2271, 2002.
- [PORT-RANGE-ARCH] Boucadair, M., Ed., Levis, P., Bajko, G., and T. Savolainen, "IPv4 Connectivity Access in the Context of IPv4 Address Exhaustion: Port Range based IP Architecture", Work in Progress, July 2009.
- [RFC6056] Larsen, M. and F. Gont, "Recommendations for Transport-Protocol Port Randomization", BCP 156, RFC 6056, January 2011.
- [RFC6269] Ford, M., Ed., Boucadair, M., Durand, A., Levis, P., and P. Roberts, "Issues with IP Address Sharing", RFC 6269, June 2011.
- [RFC6346] Bush, R., Ed., "The Address plus Port (A+P) Approach to the IPv4 Address Shortage", RFC 6346, August 2011.
- [SAM] Despres, R., "Scalable Multihoming across IPv6 Local-Address Routing Zones Global-Prefix/Local-Address Stateless Address Mapping (SAM)", Work in Progress, July 2009.

Authors' Addresses

Mohamed Boucadair
France Telecom
Rennes 35000
France

EMail: mohamed.boucadair@orange.com

Pierre Levis
France Telecom
Caen
France

EMail: pierre.levis@orange.com

Gabor Bajko
Nokia

EMail: gabor.bajko@nokia.com

Teemu Savolainen
Nokia

EMail: teemu.savolainen@nokia.com

Tina Tsou
Huawei Technologies (USA)
2330 Central Expressway
Santa Clara, CA 95050
USA

Phone: +1 408 330 4424
EMail: tina.tsou.zouting@huawei.com

