Network Working Group	M. Boucadair	
Internet-Draft	P. Levis	
Intended status: Experimental	France Telecom	
Expires: March 25, 2011	G. Bajko	
	T. Savolainen	
	Nokia	
	September 21, 2010	

TOC

Port Range Configuration Options for PPP IPCP
draft-boucadair-pppext-portrange-option-04

Abstract

This document defines two IPCP (IP Configuration Protocol) Options used to convey a set of ports. These options can be used in the context of port range-based solutions (port range delegation) or NAT-based ones (port delegation or port forwarding). Architectural considerations are out of scope of this document.

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 (Bradner, S., “Key words for use in RFCs to Indicate Requirement Levels,” March 1997.) [RFC2119].

Status of this Memo

This Internet-Draft is submitted to IETF 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 March 25, 2011.

Table of Contents

---

1. Introduction

TOC

Within the context of IPv4 address depletion, several solutions have been investigated to share IPv4 addresses. Two flavours can be distinguished: NAT-based solutions (a.k.a., Carrier Grade NAT(CGN, [I-D.shirasaki-nat444-isp-shared-addr] (Shirasaki, Y., Miyakawa, S., Nakagawa, A., Yamaguchi, J., and H. Ashida, “NAT444 addressing models,” July 2010.)) or port range based ones ([I-D.boucadair-port-range]

(Boucadair, M., Levis, P., Bajko, G., and T. Savolainen, “IPv4 Connectivity Access in the Context of IPv4 Address Exhaustion: Port Range based IP Architecture,” July 2009.) and [I-D.ymbk-aplusp] (Bush, R., “The A+P Approach to the IPv4 Address Shortage,” October 2009.) are examples of solutions which propose to share the same (public) IP address among several devices and to constrain the values used as port sources to a limited set of values). 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 new IPCP options to be used to carry Port Range information. IPCP has been widely used to convey configuration information such as IP Compression Protocol [RFC3241] (Bormann, C., “Robust Header Compression (ROHC) over PPP,” April 2002.)[RFC3544] (Koren, T., Casner, S., and C. Bormann, “IP Header Compression over PPP,” July 2003.) or IP-Address [RFC1332] (McGregor, G., “The PPP Internet Protocol Control Protocol (IPCP),” May 1992.).
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 IP-Address IPCP Option.
This document adheres to the consideration defined in [RFC2153] (Simpson, W. and K. Fox, “PPP Vendor Extensions,” May 1997.).

1.1. Use Cases

TOC

Port Range Options can be used in port range-based solutions (e.g., [I-D.boucadair-port-range] (Boucadair, M., Levis, P., Bajko, G., and T. Savolainen, “IPv4 Connectivity Access in the Context of IPv4 Address Exhaustion: Port Range based IP Architecture,” July 2009.)) or in a CGN-based solution to bypass the NAT (i.e., for transparent NAT traversal and avoid involving several NAT in the path) or to delegate one or a set of ports to the requesting client (e.g., avoid ALG (Application Level Gateway) or for port forwarding).
For improved security an option for delegating cryptographically random port range is defined.

TOC

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:

- *Contiguous Port Range: a set of port values which form a contiguous sequence.

- *Non Contiguous Port Range: a set of port values which does not form a contiguous sequence.

- *Random Port Range: a cryptographically random set of port values.

Unless explicitly mentioned, Port Mask refers to the couple (Port Range Value, Port Range Mask).
In addition, this document makes use of the following terms:

- *Delegated port or port range: a port or a range of ports belonging to an IP address managed by an upstream device (such as NAT), which are delegated to a client for use as source address and port when sending packets.

- *Forwarded port or port range: a port or a range of ports belonging to an IP address managed by an upstream device such as (NAT), which is/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 per [RFC1661] (Simpson, W., “The Point-to-Point Protocol (PPP),” July 1994.).

---

2. Port Range Options

TOC

This section defines the IPCP Option for Port Range delegation.

2.1. Description of Port Range Value and Port Range Mask

TOC

The Port Range Value and Port Range Mask are used to specify one range of ports (contiguous or not contiguous) pertaining to a given IP address. Concretely, 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 called 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:

*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 the other bits (a.k.a., non significant ones) are set to 0.

*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 the peer to 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)).
When the server assigns only shared IP addresses, the peer MUST include Port Range Option in its request. If not, Protocol-Reject sent by the server.
When a peer issues a request enclosing IPCP Port Range Option, and if the server does not support this option, the Port Range Option is rejected by the server.
The Port Range IPCP option adheres to the format defined in Section 1.1 of [RFC2153] (Simpson, W. and K. Fox, “PPP Vendor Extensions,” May 1997.).
The "value" field of the option defined in [RFC2153] (Simpson, W. and K. Fox, “PPP Vendor Extensions,” May 1997.) when conveying Port Range IPCP Option is provided in Figure 1 (Format of the Port Range IPCP Option).

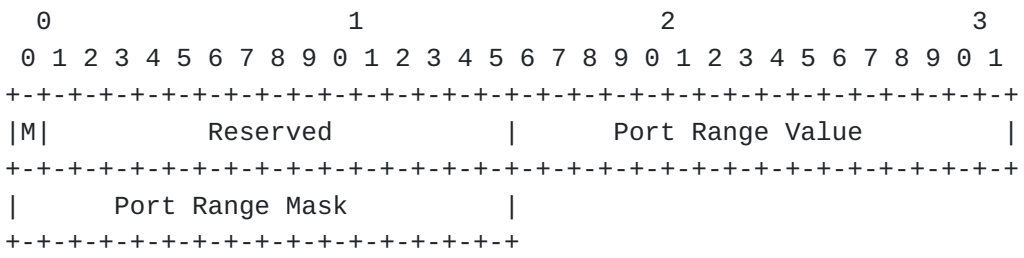

Figure 1: Format of the Port Range IPCP Option

*M: mode bit. It indicates the mode the port range is allocated for. A value of zero indicates the port ranges are delegated, while a value of 1 indicates the port ranges are port forwarded.

*Port Range Value (PRV): PRV indicates the value of the significant bits of the Port Mask. By default, no PRV is assigned.

*Port Range Mask (PRM): Port Range Mask indicates the position of the bits which 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 (Example of Port Range Mask and Port Range Value) provides an example of the resulting Port Range:
- Port Range Mask is set to 0001010000000000 (5120) and
- Port Range Value is set to 0000010000000000 (1024).

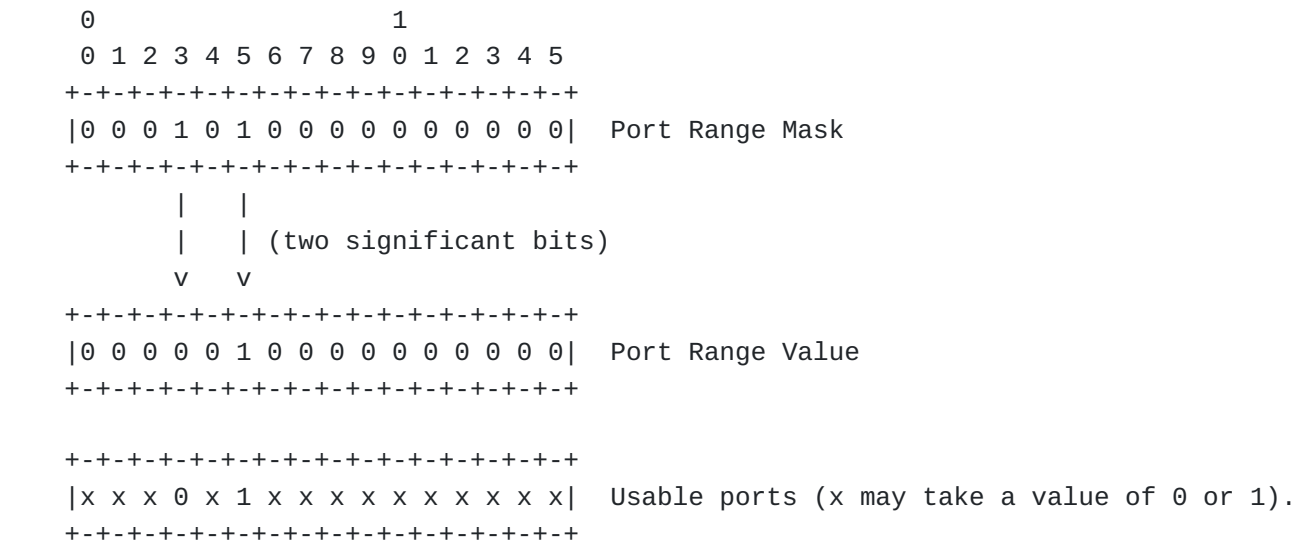

Figure 2: Example of Port Range Mask and Port Range Value

Port values belonging to this Port Range must have the 4th bit (resp. the sixth one), from the left, set to 0 (resp. 1). Only these port values will be used by the peer when enforcing the configuration conveyed by PPP IPCP.

TOC

2.2. Description of Cryptographically Random Port Range option

A cryptographically random Port Range Option may be used as a mitigation tool against blind attacks described in [I-D.ietf-tsvwg-port-randomization] (Larsen, M. and F. Gont, “Transport Protocol Port Randomization Recommendations,” August 2010.). The benefits of the approach and the method to calculate the delegated ports set are described in [I-D.bajko-pripaddrassign] (Bajko, G., Savolainen, T., Boucadair, M., and P. Levis, “Port Restricted IP Address Assignment,” October 2009.).

The cryptographically Random Port Range IPCP Option adheres to the format defined in Section 1.1 of [RFC2153] (Simpson, W. and K. Fox, “PPP Vendor Extensions,” May 1997.). The "value" field of the option defined in [RFC2153] (Simpson, W. and K. Fox, “PPP Vendor Extensions,” May 1997.) when conveying cryptographically Random Port Range IPCP Option is illustrated in Figure 3 (Format of the cryptographically Random Port Range option)

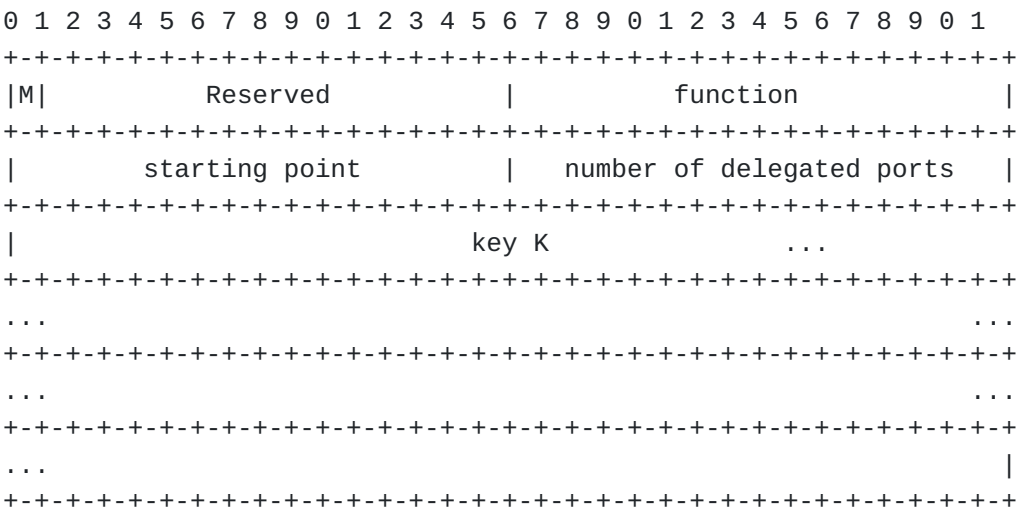

Figure 3: Format of the cryptographically Random Port Range option

---

*M: mode bit. It indicates the mode the port range is allocated for. A value of zero indicates the port ranges are delegated, while a value of 1 indicates the port ranges are port forwarded.

*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 5 of [I-D.bajko-pripaddrassign] (Bajko, G., Savolainen, T., Boucadair,

M., and P. Levis, “Port Restricted IP Address Assignment,” October 2009.).

*Starting Point: A 16 bit value used as an input to the specified function

*Number of delegated ports: A 16 bit value specifying the number of ports delegated to the client for use as source port values.

*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 'key field' and 'starting point' field SHALL be set to all zeros. The requester MAY indicate in the 'function' field which encryption function requester prefers, and in the 'number of delegated ports' field the number of ports the requester would like to obtain. If requester has no preference it SHALL set also 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 at section 2.3.

2.3. Illustration Examples

TOC

2.3.1. Overview

TOC

These flows provide examples of the usage of IPCP to convey the Port Range Option. As illustrated in Figure 4 (Successful flow), IPCP messages are exchanged between a Host and a BRAS (Broadband Access Server).

1. The first example illustrates a successful IPCP exchange;

2. The second example shows the IPCP exchange that occurs when Port Range Option is not supported by the server;

3. The third example shows the IPCP exchange that occurs when Port Range Option is not supported by the client;

4. The fourth example shows the IPCP exchange that occurs when Port Range Option is not supported by the client and a non null

IP (i.e., an address different from 0.0.0.0) address is enclosed in the first configuration request issued by the peer.

2.3.2. Successful Flow: Port Range Options supported by both the Client and the Server

TOC

The following message exchange (i.e., Figure 4 (Successful flow)) provides an example of successful IPCP configuration operation when the Port Range IPCP Option is used.

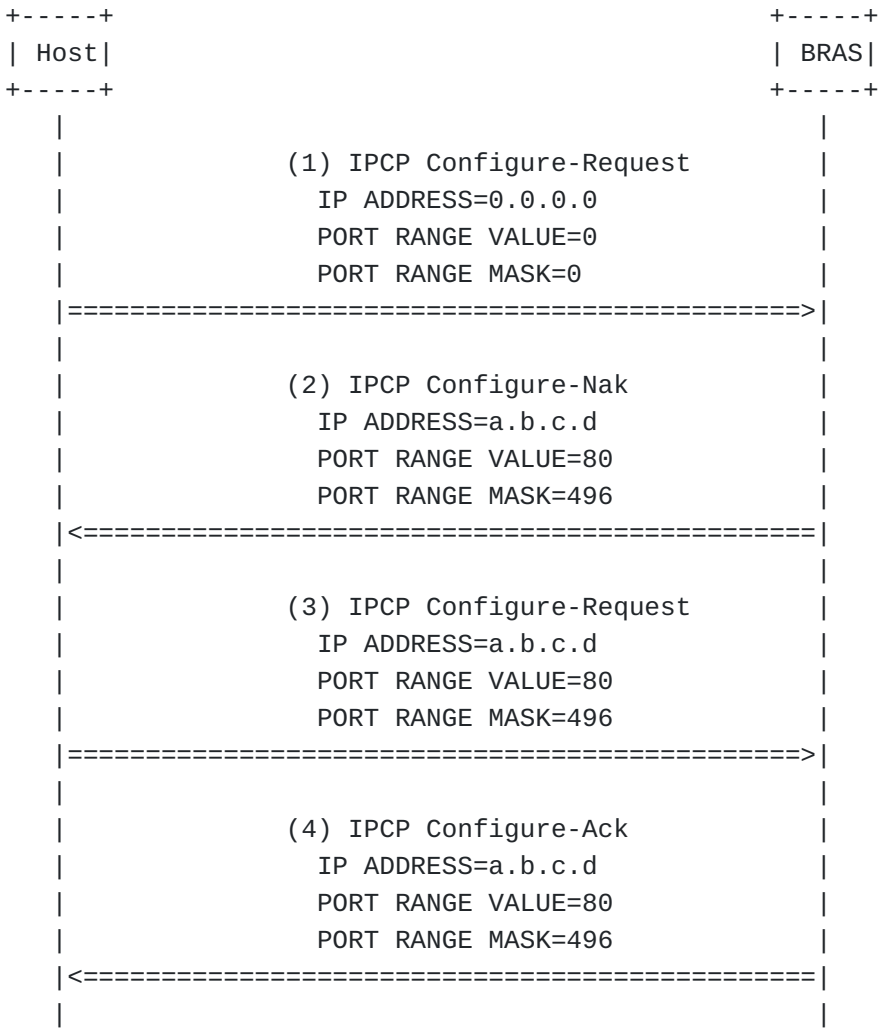

Figure 4: Successful flow

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 these Configuration Options are negotiated simultaneously. In this example, Configure-Request carries information about IP-address, Port Range Value and Port Range Mask. In this example, IP-address Option is set to 0.0.0.0, Port Range Value is set to 0 and Port Range Mask is set to 0.
>
> (2) BRAS sends back a Configure-Nak and sets the enclosed options to its preferred values. In this example: IP-Address Option is set to a.b.c.d, Port Range Value is set to 80 and Port Range Mask is set to 496.
>
> (3) The Host re-sends a Configure-Request requesting IP-address Option to be set to a.b.c.d, Port Range Value to be set to 80 and Port Range Mask to be set to 496.
>
> (4) BRAS sends a Configure-Ack message

As a result of this exchange, Host is configured to use as local IP address a.b.c.d and the following 128 contiguous Port Ranges resulting of 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

TOC

This example (Figure 5 (Failed flow: Port Range Option not supported by the server)) depicts an exchange of messages when the BRAS does not support IPCP Port Range Option.

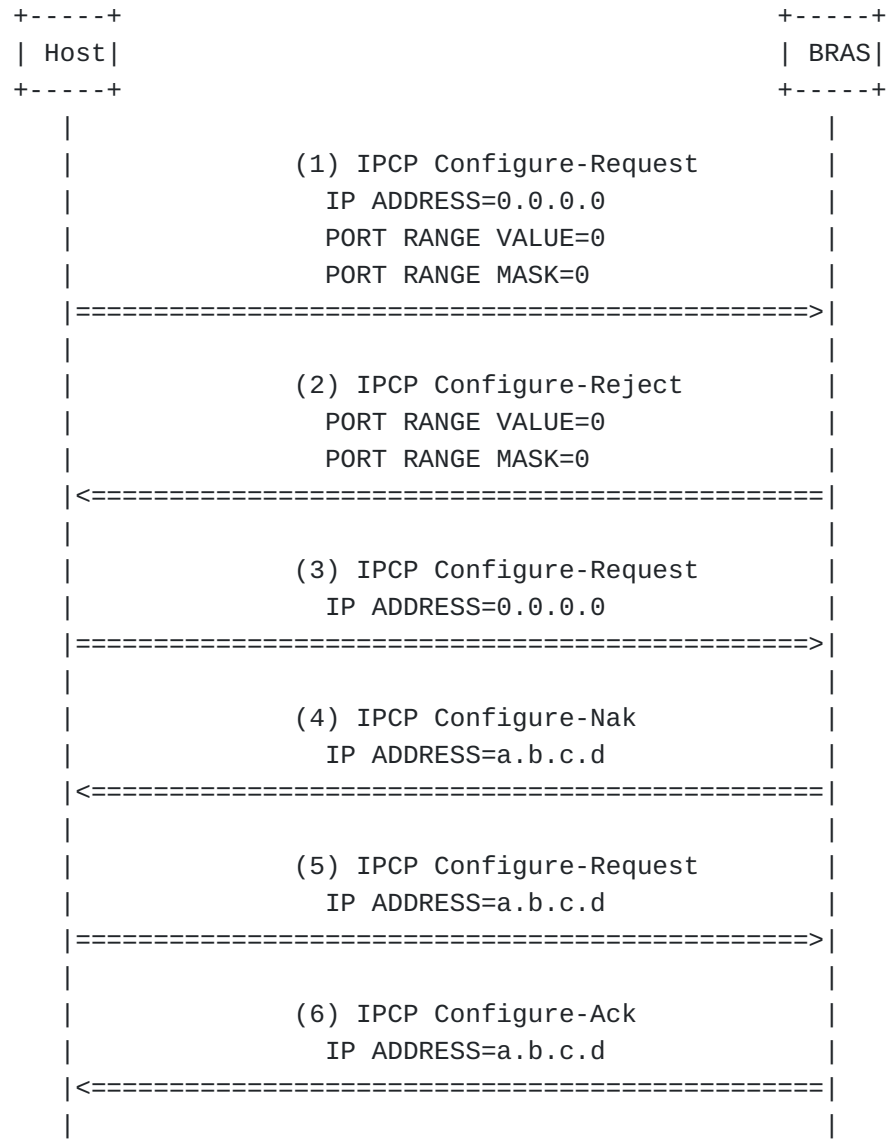

Figure 5: Failed flow: Port Range Option not supported by the server

---

The main steps of this flow are listed hereafter:

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

(3) The Host sends a Configure-Request which includes only the codes of IP-Address option. In this example, IP-Address Option is set to 0.0.0.0.

(4) BRAS sends back a Configure-Nak and sets the enclosed option to its preferred value. In this example: IP-Address Option is set to a.b.c.d.

(5) The Host re-sends a Configure-Request requesting IP-Address Option to be set to a.b.c.d.

(6) BRAS sends a Configure-Ack message.

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

2.3.4. Port Range Option not Supported by the Client

TOC

This example (Figure 6 (Port Range Option not supported by the Client)) depicts exchanges when only shared IP addresses are assigned to end-user's 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 as depicted in Figure 6 (Port Range Option not supported by the Client).

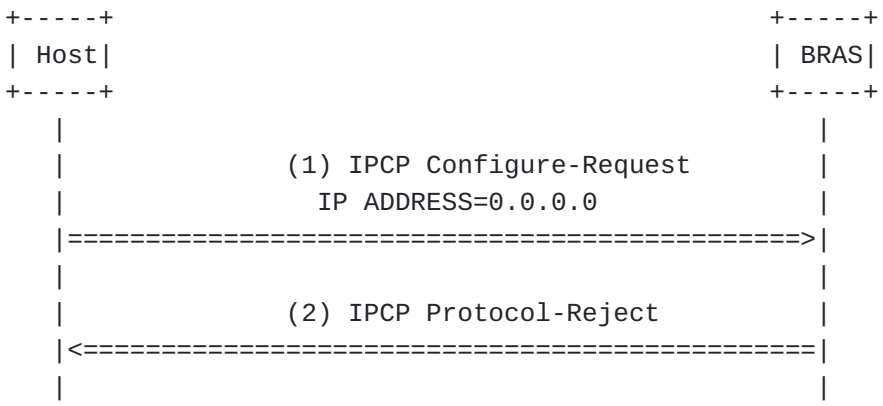

Figure 6: Port Range Option not supported by the Client

The main steps of this flow are:

(1) The Host sends a Configure-Request requesting IP-Address Option to be set to 0.0.0.0 and without enclosing the Port Range Option.

(2) BRAS sends a Protocol-Reject message.

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

---

3. IANA Considerations

TOC

No action is required from IANA since this document adheres to [RFC2153] (Simpson, W. and K. Fox, “PPP Vendor Extensions,” May 1997.).

---

4. Security Considerations

TOC

This document does not introduce any security issue in addition to those related to PPP. Service providers should use authentication mechanisms such as CHAP [RFC1994] (Simpson, W., “PPP Challenge Handshake Authentication Protocol (CHAP),” August 1996.) or PPP link encryption [RFC1968] (Meyer, G. and K. Fox, “The PPP Encryption Control Protocol (ECP),” June 1996.).
Use of small and non-random port range may increase host exposure to attacks described [I-D.ietf-tsvwg-port-randomization] (Larsen, M. and F. Gont, “Transport Protocol Port Randomization Recommendations,” August 2010.). This risk can be mitigated by using larger range or by using Random Port Range Option.

---

5. Contributors

TOC

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

---

6. Acknowledgements

TOC

The authors would like to thank Christian Jacquenet and James Carlson for their review.

TOC

7. References

7.1. Normative References

TOC

[RFC1332]	McGregor, G., "The PPP Internet Protocol Control Protocol (IPCP)," RFC 1332, May 1992 (TXT).
[RFC1661]	Simpson, W., "The Point-to-Point Protocol (PPP)," STD 51, RFC 1661, July 1994 (TXT).
[RFC1968]	Meyer, G. and K. Fox, "The PPP Encryption Control Protocol (ECP)," RFC 1968, June 1996 (TXT).
[RFC1994]	Simpson, W., "PPP Challenge Handshake Authentication Protocol (CHAP)," RFC 1994, August 1996 (TXT).
[RFC2119]	Bradner, S., "Key words for use in RFCs to Indicate Requirement Levels," BCP 14, RFC 2119, March 1997 (TXT, HTML, XML).
[RFC2153]	Simpson, W. and K. Fox, "PPP Vendor Extensions," RFC 2153, May 1997 (TXT).

7.2. Informative References

TOC

[I-D.bajko-pripaddrassign]	Bajko, G., Savolainen, T., Boucadair, M., and P. Levis, "Port Restricted IP Address Assignment," draft-bajko-pripaddrassign-02 (work in progress), October 2009 (TXT).
[I-D.boucadair-port-range]	Boucadair, M., Levis, P., Bajko, G., and T. Savolainen, "IPv4 Connectivity Access in the Context of IPv4 Address Exhaustion: Port Range based IP Architecture," draft-boucadair-port-range-02 (work in progress), July 2009 (TXT).
[I-D.ietf-tsvwg-port-randomization]	Larsen, M. and F. Gont, "Transport Protocol Port Randomization Recommendations," draft-ietf-tsvwg-port-randomization-09 (work in progress), August 2010 (TXT).
[I-D.shirasaki-nat444-isp-shared-addr]	Shirasaki, Y., Miyakawa, S., Nakagawa, A., Yamaguchi, J., and H. Ashida, "NAT444 addressing models," draft-shirasaki-nat444-isp-shared-addr-04 (work in progress), July 2010 (TXT).
[I-D.ymbk-aplusp]	Bush, R., "The A+P Approach to the IPv4 Address Shortage," draft-ymbk-aplusp-05 (work in progress), October 2009 (TXT).
[RFC3241]	Bormann, C., "Robust Header Compression (ROHC) over PPP," RFC 3241, April 2002 (TXT).
[RFC3544]	

	Koren, T., Casner, S., and C. Bormann, “IP Header Compression over PPP,” RFC 3544, July 2003 (TXT).

Authors' Addresses

TOC

	Mohamed Boucadair
	France Telecom
	3, Av François Château
	Rennes 35000
	France
Email:	mohamed.boucadair@orange-ftgroup.com
	Pierre Levis
	France Telecom
Email:	pierre.levis@orange-ftgroup.com
	Gabor Bajko
	Nokia
Email:	gabor(dot)bajko(at)nokia(dot)com
	Teemu Savolainen
	Nokia
Email:	teemu.savolainen@nokia.com