

Network Working Group
Internet-Draft
Intended status: Informational
Expires: March 18, 2012

M. Boucadair
P. Levis
France Telecom
G. Bajko
T. Savolainen
Nokia
T. Tsou
Huawei Technologies (USA)
September 15, 2011

Huawei Port Range Configuration Options for PPP IPCP
draft-boucadair-pppext-portrange-option-09

Abstract

This document defines two Huawei 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 or NAT-based ones for port delegation and forwarding purposes.

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 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 18, 2012.

Copyright Notice

Copyright (c) 2011 IETF Trust and the persons identified as the document authors. All rights reserved.

Internet-Draft

Port Range IPCP Options

September 2011

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
1.1.	Use Cases	3
1.2.	Terminology	3
2.	Port Range Options	4
2.1.	Description of Port Range Value and Port Range Mask	4
2.2.	Description of Cryptographically Random Port Range option	7
2.2.1.	Random Port Delegation Function	7
2.2.2.	Description of Cryptographically Random Port Range Option	9
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	11
2.3.3.	Port Range Option Not Supported by the Server	12
2.3.4.	Port Range Option not Supported by the Client	14
3.	IANA Considerations	15
4.	Security Considerations	15
5.	Contributors	15
6.	Acknowledgements	15
7.	References	15
7.1.	Normative References	15
7.2.	Informative References	16
	Authors' Addresses	17

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 (a.k.a., Carrier Grade NAT (CGN, [[I-D.ietf-behave-lsn-requirements](#)])) or port range based ones (e.g., [[RFC6346](#)] [[I-D.boucadair-port-range](#)][I-D.despres-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 IP-Address IPCP Option.

This document adheres to the consideration defined in [[RFC2153](#)].

This document is not a product of 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 avoid involving several NAT levels 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).

[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:

Boucadair, et al.

Expires March 18, 2012

[Page 3]

Internet-Draft

Port Range IPCP Options

September 2011

- o Contiguous Port Range: a set of port values which form a contiguous sequence.
- o Non Contiguous Port Range: a set of port values which does 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 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.
- o Forwarded port or forwarder 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 [\[RFC1661\]](#).

[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 provided 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: The content of this field is specified in [Section 2.1](#) and [Section 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 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:

- 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 the other bits (a.k.a., 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 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 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 set of ports conveyed in an IPCP Port Range Option applies to all transport protocols.

The set of ports conveyed in a IPCP Port Range Option are revoked when the link is not any more 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" field of the option defined in [\[RFC2153\]](#) when conveying Port Range IPCP Option is provided in Figure 1.



MSB network order is used for encoding Port Range Value and Port Range Mask fields.

Figure 1: Format of the Port Range IPCP Option

- o 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.

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

```

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
+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+
      |   |
      |   | (two significant bits)
      v   v
```

```

+---+---+---+---+---+---+---+---+---+---+
|0 0 0 0 0 1 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| 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 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.

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 [\[RFC6056\]](#).

2.2.1. Random Port Delegation Function

Delegating random ports can be achieved by defining a function which takes as input a key 'k' and an integer 'x' within the range (1024, 65535) and produces an output 'y' also within the port range (1024, 65535).

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

The cryptographical mechanism ensures that the entire 64k port range can be efficiently distributed to multiple nodes in a way that when nodes calculate the ports, the results will never overlap with ports 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 range (1024, 65535), and produces an output y (the ciphertext), also an integer in range (1024, 65535). 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 e.g. 2048 random ports, it selects a starting point ' a ' ($1024 \leq a \leq 65536-2048$) in a way 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 different K for each IPv4 address it allocates to make attacks as difficult as possible. This way, learning the K used in IPv4 address IP1 would not help in attacking IPv4 address IP2 that 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 range (1024, 65535).

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 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 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" field of the option defined in [\[RFC2153\]](#) when conveying cryptographically Random Port Range IPCP Option is illustrated in Figure 3

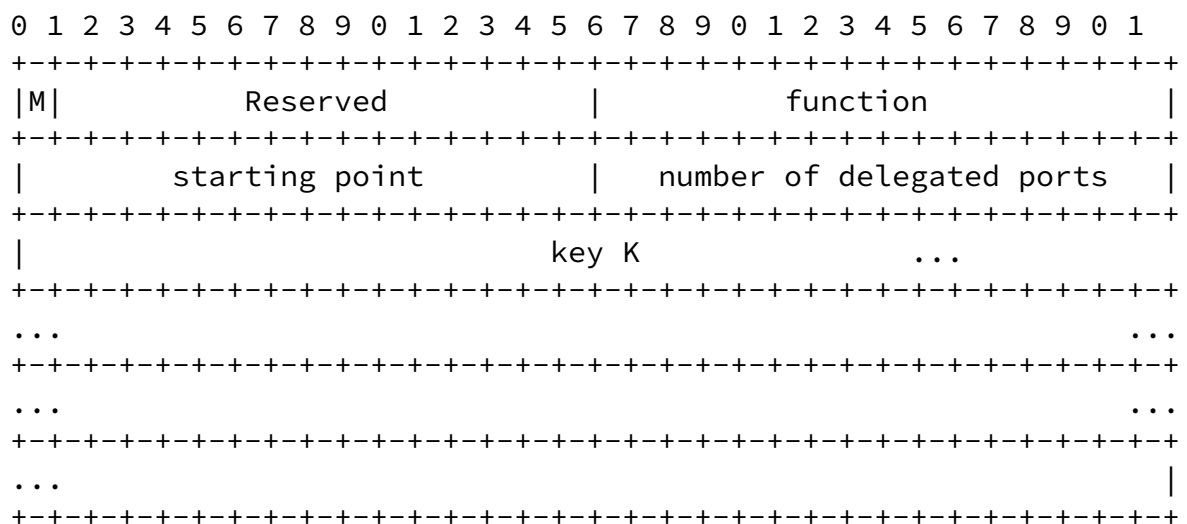


Figure 3: Format of the cryptographically Random Port Range option

o M: mode bit. 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.

- 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 '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 2.1](#).

[2.3](#). Illustration Examples

[2.3.1](#). Overview

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

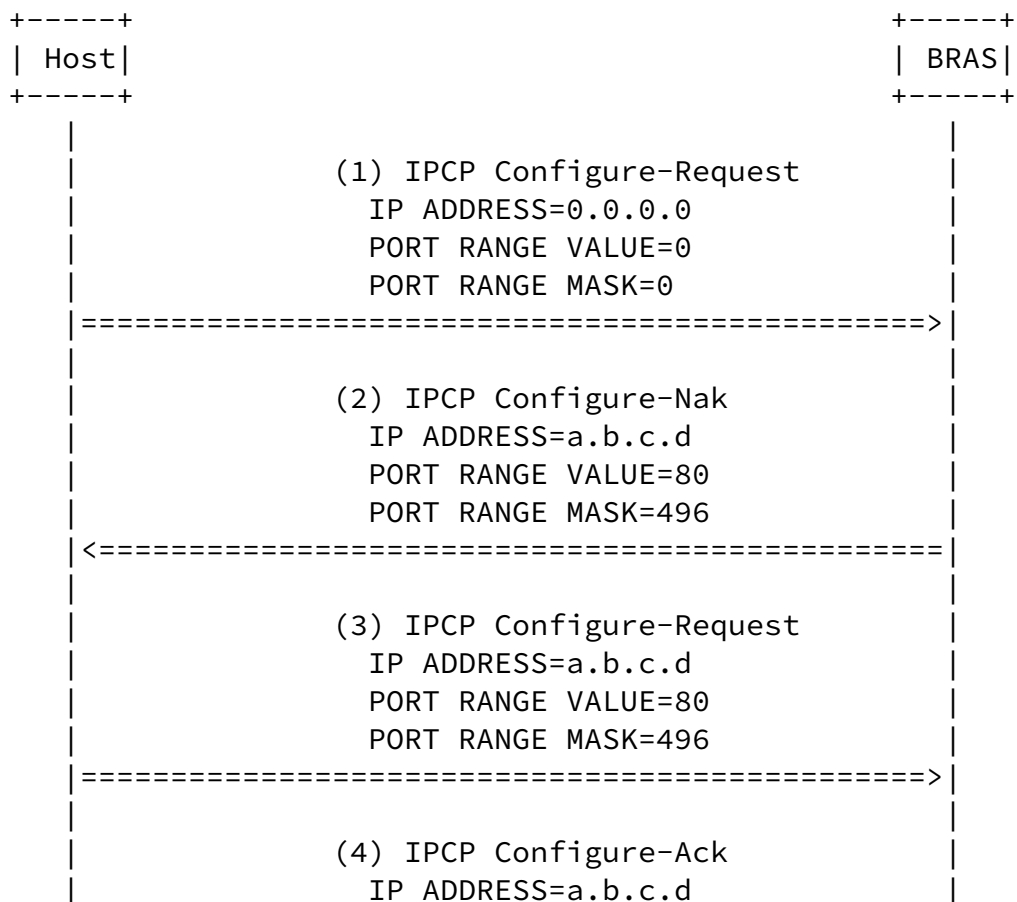
1. The first example illustrates a successful IPCP exchange;
2. The second example shows the IPCP exchange that occurs when Port Range Option is not supported by the server;
3. The third example shows the IPCP exchange that occurs when Port

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

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



2.3.3. Port Range Option Not Supported by the Server

This example (Figure 5) depicts an exchange of messages when the BRAS does not support IPCP Port Range Option.



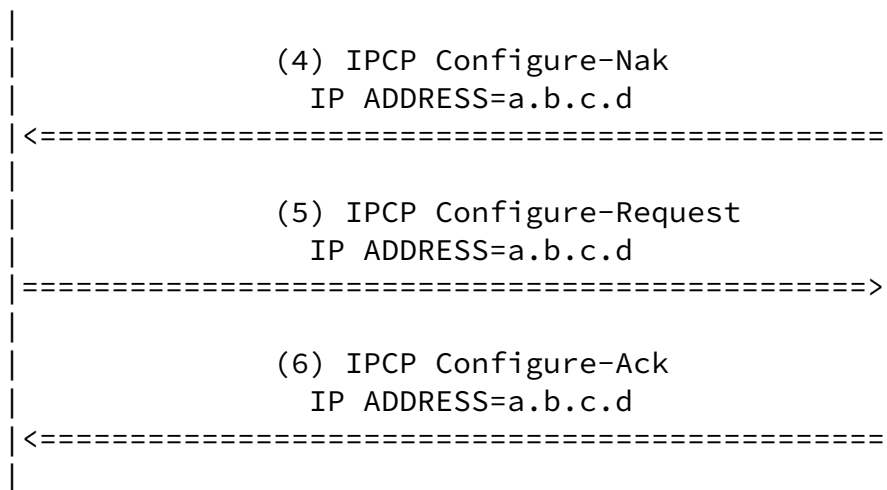


Figure 5: Failed flow: Port Range Option not supported by the server

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

This example (Figure 6) 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.




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

No action is required from IANA since this document adheres to [\[RFC2153\]](#).

4. Security Considerations

This document does not introduce any security issue in addition to those related to PPP. Service providers should use authentication mechanisms such as CHAP [\[RFC1994\]](#) or PPP link encryption [\[RFC1968\]](#).

The use of small and non-random port range may increase host exposure to attacks described in [\[RFC6056\]](#). This risk can be reduced by using larger port ranges, by using Random Port Range Option or by activating means to improve the robustness of TCP against Blind In-Window Attacks [\[RFC5961\]](#).

5. Contributors

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

6. Acknowledgements

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

7. References

7.1. Normative References

- [RFC1661] Simpson, W., "The Point-to-Point Protocol (PPP)", STD 51, [RFC 1661](#), July 1994.
- [RFC1968] Meyer, G. and K. Fox, "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. and K. Fox, "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.

[7.2.](#) Informative References

- [CIPHERS] Black, J. and P. Rogaway, "Ciphers with Arbitrary Finite Domains Topics in Cryptology", 2002, < CT-RSA 2002, Lecture Notes in Computer Science vol. 2271>.
- [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.
- [I-D.despres-sam]
Despres, R., "Scalable Multihoming across IPv6 Local-Address Routing Zones Global-Prefix/Local-Address Stateless Address Mapping (SAM)", [draft-despres-sam-03](#) (work in progress), July 2009.
- [I-D.ietf-behave-lsn-requirements]
Perreault, S., Yamagata, I., Miyakawa, S., Nakagawa, A., and H. Ashida, "Common requirements for Carrier Grade NAT (CGN)", [draft-ietf-behave-lsn-requirements-03](#) (work in progress), August 2011.
- [RFC6056] Larsen, M. and F. Gont, "Recommendations for Transport-Protocol Port Randomization", [BCP 156](#), [RFC 6056](#), January 2011.
- [RFC6269] Ford, M., Boucadair, M., Durand, A., Levis, P., and P. Roberts, "Issues with IP Address Sharing", [RFC 6269](#), June 2011.
- [RFC6346] Bush, R., "The Address plus Port (A+P) Approach to the IPv4 Address Shortage", [RFC 6346](#), August 2011.

Internet-Draft

Port Range IPCP Options

September 2011

Authors' Addresses

Mohamed Boucadair
France Telecom
Rennes 35000
France

Email: mohamed.boucadair@orange-ftgroup.com

Pierre Levis
France Telecom
Caen
France

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

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

Phone: +1 408 330 4424

Email: tina.tsou.zouting@huawei.com

