

Network Working Group
Internet-Draft
Intended status: Informational
Expires: July 21, 2016

G. Chen
China Mobile
W. Li
China Telecom
T. Tsou
J. Huang
Huawei Technologies
T. Taylor
PT Taylor Consulting
JF. Tremblay
Viagenie
January 18, 2016

Analysis of NAT64 Port Allocation Methods for Shared IPv4 Addresses
draft-ietf-sunset4-nat64-port-allocation-02

Abstract

This document enumerates methods of port assignment in Carrier Grade NATs (CGNs), focused particularly on NAT64 environments. Different NAT port allocation methods have been categorized and described. A series of port allocation design principle has been proposed to facilitate the implementations and deployment.

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 July 21, 2016.

Copyright Notice

Copyright (c) 2016 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	2
2.	Description of Port Allocation Methods	3
2.1.	Specific Feature of NAT64 Port Consumption	3
2.2.	Classification of Port Allocation Models	3
2.2.1.	Stateful vs. Stateless	3
2.2.2.	Dynamic vs. Static	4
2.2.3.	Centralized vs. Distributed	5
2.3.	Port Allocation Solutions	6
2.3.1.	Stateful Technologies	6
2.3.2.	Stateless Technologies	6
2.3.3.	Port Control Protocol (PCP)	7
3.	Port Allocation Design Principles	7
3.1.	Log Volume Optimization	7
3.2.	Connectivity State Optimization	9
3.3.	Port Randomization	9
3.4.	Port-range Implementation Recommendation	10
3.4.1.	Port Randomization and Port-Range Deallocation	10
3.4.2.	Issues Of Traceability	11
3.4.3.	Other Considerations	12
4.	Security Considerations	12
5.	IANA Considerations	13
6.	Acknowledgements	14
7.	References	14
7.1.	Normative References	14
7.2.	Informative References	15
	Authors' Addresses	17

1. Introduction

As a result of the depletion of public IPv4 addresses, Carrier Grade NAT (CGN) has been adopted by ISPs to share the available IPv4 resources. Overall, a CGN function maps IP addresses from one address realm to another, relying upon a mechanism of multiplexing multiple subscribers' connections over a number of shared IPv4 addresses to provide connectivity services to end hosts. A network-based NAT is implied by several approaches to IPv4 service continuity

over an IPv6 network including DS-Lite [RFC6333], NAT64 ([RFC6145] and [RFC6146]), etc.

In this memo, Section 2 has described the category of port allocation mechanism and relevant solutions. Section 3 looks more closely at different port allocation design principles, including log volume consideration, connectivity state optimization, Port-range implementation and port randomization. The proposals made in this section are applicable to the CGN environment in general, independently of the particular flavor of translation being used.

2. Description of Port Allocation Methods

2.1. Specific Feature of NAT64 Port Consumption

There was a test comparison of port consumption [RFC7269] on NAT64 and NAT44. Top100 websites (referring to Alexa statistics) were assessed to evaluate status of port usage on NAT44 and NAT64 respectively. It has been observed that the port consumption per session on NAT64 is roughly only half that on NAT44. 43 percent of top100 websites have AAAA records, therefore the NAT64 didn't have to assign ports to the traffic going to those websites. The results may be different if more services (e.g. game, web-mail, etc) are considered. But it is apparent that the effects of port saving on NAT64 will be amplified by increasing native IPv6 support.

Apart from the above observation, port allocation can be tuned according to the phase of IPv6 migration. As more content providers and services become available over IPv6, the utilization of NAT64 goes down since fewer destinations require translation progressing. Thus as IPv6 migration proceeds, it will be possible to relax the multiplexing ratio of IPv4 address sharing (see Appendix B of [RFC6269]).

2.2. Classification of Port Allocation Models

This section lists several categories to allocate the port information in NAT64. It also describes example cases for each allocation model.

2.2.1. Stateful vs. Stateless

- o Stateful

The stateful NAT can be implemented either by static address translation or dynamic address translation.

In the case of static address assignment, a one-to-one address mapping for hosts between a IPv6 network address and an IPv4 network address is pre-configured on the NAT operation. This case normally occurs when a server is deployed in an IPv6 domain. The static configuration ensures stable inbound connectivity.

Dynamic address assignment would periodically free the binding so that the global address could be recycled for later use. This increases the efficiency of usage of IPv4 resources.

- o Stateless

Stateless NAT is performed in compliance with [RFC6145]. The public IPv4 address is required to be embedded in the IPv6 address. Thus the NAT64 can directly extract the address and has no need to record mapping states.

A promising usage of stateless NAT may appear in a data centre environment where IPv6 server pools receive inbound connections from IPv4 users externally [I-D.ietf-v6ops-siit-dc]. NAT usage in other cases may be controversial. First off, the static one-to-one mapping does not address the issue of IPv4 depletion. Secondly, it introduces a dependency between IPv4 and IPv6 addressing. That creates other limitations since a change of IPv4 address will cause renumbering of IPv6 addresses.

2.2.2. Dynamic vs. Static

Port assignments can be dynamic (ports allocated on demand) or static (ports allocated as part of the configuration process).

- o Dynamic assignment

NAT64 uses dynamic assignment, since this achieves higher port utilization. Port allocations can be made with per-session or per-customer granularity. Per-session assignment is configured on the NAT64 by default since it maximizes port utilization. However, if only individual port numbers are assigned, this can result in a heavy log volume that may have to be recorded for legal data retention systems. To mitigate that concern, the NAT64 may dynamically allocate a port range for each connected subscriber or upon receipt of a first outgoing packet from an IPv6 host. This will significantly reduce log volume.

A proper port-range configuration should take user experiences into the considerations. A subscriber normally uses multiple applications simultaneously, e.g. maps applications, online video or game. The number of concurrent sessions is essential to

determine the number of ports the subscriber needs. A study from China Mobile has revealed that the average number of sessions consumed by one user's device was around 200 to 300 ports. Several devices may appear behind a CPE. Based on this observation, 1000 ports per subscriber household will provide enough room for multiple active users. Administrators should monitor usage to adjust this number if users are being limited by this number, or if usage is so low that fewer ports would be sufficient.

- o Static assignment

Static assignment makes port reservations in bulk for each internal address before subscriber connection. The assigned ports can be in either a contiguous port range or a non-contiguous port range for the sake of defense against port-guessing attacks (see Section 3.4.1). Log recording for each port assignment may not be necessary due to the stable mapping relations. Considerations of the interaction between port-range allocation and capacity impact are also applicable in the case of static assignment. [RFC7422] describes a deterministic algorithm to assign a port range for an internal IP address pool in a sequence.

2.2.3. Centralized vs. Distributed

There is an increasing need to connect NAT64 with downstream NAT46-capable devices to support IPv4 users/applications on an IPv6-only path. Several solutions have been proposed in this area, e.g., 464xlat [RFC6877], MAP-T [RFC7599] and 4rd [RFC7600]. Port allocation can be categorized as a centralized assignment on NAT64 or as a port delegation distributed to downstream devices (e.g, Customer Edge connected with NAT64).

- o Centralized Assignment

A centralized method makes port assignments once IP flows come to the NAT64. The allocation policy is enforced on a centralized point. Either a dynamic or static port assignment is made for received sessions.

- o Distributed Assignment

NAT64 can also delegate the pre-allocated port range to customer edge devices. That can be achieved through additional out-of-band provisioning signals (e.g., [I-D.ietf-pcp-port-set], [RFC7598]). The distributed model normally is performed A+P style [RFC6346] for static port assignment. The NAT64 should also hold the corresponding mapping in order to validate port usage in the

outgoing direction and route inbound packets. Delegated port ranges shift NAT64 port computations/states into downstream devices. The detailed benefits of this approach are documented in [I-D.ietf-softwire-stateless-4v6-motivation].

2.3. Port Allocation Solutions

2.3.1. Stateful Technologies

[RFC6146] describes a process where the dynamic binding is created by an outgoing packet, but it may also be created by other means such as a Port Control Protocol request (see Section 2.3.3). Looking beyond NAT64 for the moment, DS-Lite [RFC6333] refers to the cautions in [RFC6269] but does not specify any port allocation method. Both techniques DS-Lite and NAT64 assume a centralized model.

The specifications for both transition methods thus allow implementations to use the proposals made in individual port allocation and port range allocation

2.3.2. Stateless Technologies

The port allocation solutions that are being specified at the time of writing of this document are all variations on the static distributed model, to minimize the amount of state that has to be held in the network. That work includes:

- o Light-weight 4over6 (LW4o6 [RFC7596]), which requires the CPE to be configured explicitly with the shared IPv4 address and port set it will use on the WAN side of its NAT44 function. The border router is configured with the same information, reducing the state it must hold from per-session to per-subscriber amounts.
- o Mapping of Address and Port with Encapsulation (MAP-E [RFC7597]) and the experimental specifications Mapping of Address and Port with Translation (MAP-T [RFC7599]) and 4rd [RFC7600], already mentioned. These rely on an algorithmic embedding of WAN-side IPv4 address and assigned port set within the IPv6 prefix assigned to each CPE. Both the CPE and the border router must be configured with this information. However, the algorithm is designed to aggregate routing information such that the amount of state carried by the border router is of a lower order of magnitude than even the per-subscriber level.

All A+P variants support a 1-1 mapping mode, where the IPv4 and IPv6 addresses assigned to a CPE are independent. This can be helpful in transition, but, as with LW4o6, raises the amount of state in the network back to the per-subscriber level.

For a packet destined to a host outside the MAP domain from which the packet originated: MAP-E and 4rd treat the packet as an IPv4 over IPv6 tunnel via the border router.

MAP-T uses stateless mapping in the sense of Section 2.2.1 by embedding the destination IPv4 address within the IPv6 address of the packet sent to the border router.

2.3.3. Port Control Protocol (PCP)

The Port Control Protocol (PCP, [RFC6887]) can be used to reserve a single port or a port set [I-D.ietf-pcp-port-set] for applications. It requires that the NAT be controlled by a PCP server function. PCP provides an out-of-band signalling mechanism for coordinating dynamic allocation of ports between hosts and the border router, removes the need for ALGs, allows for successful incoming connections, etc.

3. Port Allocation Design Principles

3.1. Log Volume Optimization

[RFC6269] provides a thoughtful analysis on the issues of IP address sharing. It points out that IP address sharing may impact law enforcement since source address information will be lost during the translation. In order to identify a specific user associated with an IP address in a particular time slot, network administrators have to log the mapping status for each connection in a dedicated logging server. The storage of log information may pose a challenge to operators, since it requires additional resources and data inspection processes to identify users. For concrete details of what should be logged, see Section 3.1 of [I-D.ietf-behave-syslog-nat-logging]. The actual logging may use either IPFIX [RFC7011] or Syslog [RFC5424] depending on the operator's requirements.

It is desirable to reduce the volume of the logged information. Referring to the classification of port allocation methods given above, dynamic assignments can be managed on either a per-session or per-customer granularity. The coarser granularity will lead to lower log volume storage. A test was made by recording the log information from 200,000 subscribers in the Chinese network for 60 days. The volume of recorded information reached up to 42.5 terabytes with per-session logging in the raw format. The volume could be reduced to 10.6 terabytes with gzip format. Compared with that, it only occupied 40.6 gigabytes, three orders of magnitude smaller volume, with per-customer logging in the raw format (A port range sized by 1000 ports have been used). With static allocation, of course, no logs for port assignment are required, but a record of the configuration change is still required.

On the other hand, the lower logging volumes are associated with lower efficiency of port utilization. A port allocation based on per-customer granularity has to retain vacant ports in order to avoid traffic overflow. The efficiency can be evaluated by port utilization rate, and will be even lower if the static port allocation method is used. Inactive users may also impact the efficiency.

Table 1 summarizes the test results using Syslog. The ports were pre-allocated to customers regardless of online or offline status.

Port Allocation Method	Log Granularity	Estimated Log Volume	Port Utilization
Dynamic NAPT	Per-session	42.5 terabytes	100%
Dynamic port-range	Per-customer	40.6 Gigabytes	75%
Deterministic NAT, MAP-T, 4rd	None	None	(60% * 75%) = 45%

Table 1: Estimated Log Volumes For 200,000 Users Over 60 Days

Note: 75% is the estimated port utilization ratio per active subscriber. 60% is the estimated ratio of active subscribers to the total number of subscribers.

The data shown in Table 1 roughly demonstrates the tradeoff between port utilization and log volume reduction. Administrators may consider the following factors to make their design choice that would meet their deployment requirements:

- o average connectivity per customer per day;
- o peak connectivity per day;
- o the number of public IPv4 addresses available to the NAT64;
- o application demands for specific ports;
- o processing capabilities of the NAT64;
- o tolerable log volume.

3.2. Connectivity State Optimization

It has been observed that port consumption is significantly increased once subscribers land on a web page for video on demand, an online game, or map services. In those cases, multiple TCP connections may be initiated to optimize the performance of data transmissions for video download and message exchange. Given the video traffic growth trend, this likely presents a challenge for network operators who need to optimize connectivity states and avoid port depletion. Those optimizations may even affect the method of port-range allocation, because a subscriber is only allowed to use a pre-configured port resource.

Two optimizations may be considered:

- o Reducing the TIME-WAIT state. It is rather common that users change video channels often. Investigations have shown that 60% of videos are watched for less than 20% of their duration. The user's access patterns may leave a number of the TIME-WAIT states. Therefore, acceleration of TIME-WAIT state transitions could increase the efficiency of port utilization. [RFC6191] defines a mechanism for reducing TIME-WAIT state by proposing TCP timestamps and sequence numbers.

[I-D.ietf-tsvwg-behave-requirements-update] recommended applying [RFC6191] and PAWS (Protect Against Wrapped Sequence numbers, described in [RFC1323]) to NAT. This may also be a way to improve port utilization.

- o Another possibility is to use Address-Dependent Mapping or Address and Port-Dependent Mapping [RFC4787] to increase port utilization. This feature has already been implemented on a vendor-specific basis. However, it should be noted that REQ-7 and REQ-12 in [RFC6888] may reduce the incentive to use anything but the Address-Independent Mapping behaviour recommended by [RFC4787].

3.3. Port Randomization

Port randomization is a feature to enhance the defense against hijacking of flows. [RFC6056] specifies that:

"A NAT that does not implement port preservation ([RFC4787], [RFC5382]) should obfuscate selection of the ephemeral port of a packet when it is changed during translation of that packet."

A NAT based on per-session allocation normally follows this recommendation.

See Section 4 for a fuller discussion of port randomization.

3.4. Port-range Implementation Recommendation

Allocating a range of N ports at once reduces the log volume by a factor of N , while also reducing port utilization by a factor which varies with the address sharing ratio and other configuration parameters. This provides a clear motivation to use dynamic allocation of port-ranges rather than individual ports when it is possible to do so while maintaining a satisfactory level of port utilization (and by implication, shared global IPv4 address utilization). Dynamic allocation of port ranges may be used either as the sole strategy for port allocation on the NAT, or as a supplement to an initial static allocation. This section will provide specific consideration to the implementation.

3.4.1. Port Randomization and Port-Range Deallocation

When the user sends out the first packet, a port resource pool is allocated for the user, e.g., assigning ports 2001~2300 of a public IP address to the user's resource pool. Only one log should be generated for this port block. When the NAT needs to set up a new mapping entry for the user, it can use a port in the user's resource pool and the corresponding public IP address. If the user needs more port resources, the NAT can allocate another port block, e.g., ports 3501~3800, to the user's resource pool. Again, just one log needs to be generated for this port block.

Cryptographically random port assignment is discussed in Section 2.2 of [RFC6431]. Indeed, [RFC6431] takes this idea further by allocating non-contiguous sets of ports using a pseudorandom function. Scattering the allocated ports in this way provides a modest barrier to port guessing attacks. The use of randomization is discussed further in Section 4.

Suppose now that a given internal address has been assigned more than one block of ports. The individual sessions using ports within a port block will start and end at different times. If no ports in some port block are used for some configurable time, the NAT can remove the port block from the resource pool allocated to a given internal address, and make it available for other users. In theory, it is unnecessary to log deallocations of blocks of ports, because the ports in deallocated blocks will not be used again until the blocks are reallocated. However, the deallocation may be logged when it occurs to add robustness to troubleshooting or other procedures.

The deallocation procedure presents a number of difficulties in practice. The first problem is the choice of timeout value for the

block. If idle timers are applied for the individual mappings (sessions) within the block, and these conform to the recommendations for NAT behaviour for the protocol concerned, then the additional time that might be configured as a guard for the block as a whole need not be more than a few minutes. The block timer in this case serves only as a slightly more conservative extension of the individual session idle timers. If, instead, a single idle timer is used for the whole block, it must itself conform to the recommendations for the protocol with which that block of ports is associated. For example, REQ-5 of [RFC5382] requires an idle timer expiry duration of at least 2 hours and 4 minutes for TCP. The suggestions made in Section 3.2 may be considered for reducing this time.

The next issue with port block deallocation is the conflict between the desire to randomize port allocation and the desire to make unused resources available to other internal addresses. As mentioned above, ideally port selection will take place over the entire set of blocks allocated to the internal address. However, taken to its fullest extent, such a policy will minimize the probability that all ports in any given block are idle long enough for it to be released.

As an alternative, it is suggested that when choosing which block to select a port from, the NAT should omit from its range of choice the block that has been idle the longest, unless no ports are available in any of the other blocks. The expression "block that has been idle the longest" designates the block in which the time since the last packet was observed in any of its sessions, in either direction, is earlier than the corresponding time in any of the other blocks assigned to that internal address. As [RFC6269] points out, port randomization is just one security measure of several, and the loss of randomness incurred by the suggested procedure is justified by the increased utilization of port resources it allows.

3.4.2. Issues Of Traceability

Section 12 of [RFC6269] provides a good discussion of the traceability issue. Complete traceability given the NAT logging practices proposed in this draft requires that the remote destination record the source port of a request along with the source address (and presumably protocol, if not implicit) [RFC6302]. In addition, the logs at each end must be timestamped, and the clocks must be synchronized within a certain degree of accuracy. Here is one reason for the guard timing on block release, to increase the tolerable level of clock skew between the two ends.

Where source port logging can be enabled, this memo strongly urges the operators to do so. Similarly, intrusion detection systems

should capture source port as well as source address of suspect packets.

In some cases [RFC6269], a server may not record the source port of a connection. To allow traceability, the NAT device needs to record the destination IP address of a connection. As [RFC6269] points out, this will provide an incomplete solution to the issue of traceability because multiple users of the same shared public IP address may access the service at the same time. From the point of view of this draft, in such situations the game is lost, so to speak, and port allocation at the NAT might as well be completely dynamic.

The final possibility to consider is where the NAT does not do per-session logging even given the possibility that the remote end is failing to capture source ports. In that case, the port allocation strategy proposed in this section can be used. The impact on traceability is that analysis of the logs would yield only the list of all internal addresses mapped to a given public address during the period of time concerned. This has an impact on privacy as well as traceability, depending on the follow-up actions taken.

3.4.3. Other Considerations

[RFC6269] notes several issues introduced by the use of dynamic as opposed to static port assignment. For example, Section 12.2 of that document notes the effect on authentication procedures. These issues must be resolved, but are not specific to the dynamic port-range allocation strategy.

4. Security Considerations

The discussion which follows addresses an issue that is particularly relevant to the strategies described in Section 3 of this document. The security considerations applicable to NAT operation for various protocols as documented in, for example, [RFC4787] and [RFC5382] also apply to this proposal.

[RFC6056] summarizes the TCP port-guessing attack, by means of which an attacker can hijack one end of a TCP connection. One mitigating measure is to make the source port number used for a TCP connection less predictable. [RFC6056] provides various algorithms for this purpose.

As Section 3.1 of that RFC notes: "...provided adequate algorithms are in use, the larger the range from which ephemeral ports are selected, the smaller the chances of an attacker are to guess the selected port number." Conversely, the reduced range sizes proposed by the present document increase the attacker's chances of guessing

correctly. This result cannot be totally avoided. However, mitigating measures to improve this situation can be taken both at port block assignment time and when selecting individual ports from the blocks that have been allocated to a given user.

At assignment time, one possibility is to assign ports as non-contiguous sets of values as proposed in [RFC6431]. However, this approach creates a lot of complexity for operations, and the pseudo-randomization can create uncertainty when the accuracy of logs is important to protect someone's life or liberty.

Alternatively, the NAT can assign blocks of contiguous ports. However, at assignment time the NAT could attempt to randomize its choice of which of the available idle blocks it would assign to a given user. This strategy has to be traded off against the desirability of minimizing the chance of conflict between what [RFC6056] calls "transport protocol instances" by assigning the most-idle block, as suggested in Section 3. A compromise policy might be to assign blocks only if they have been idle for a certain amount of time whenever possible, and select pseudorandomly between the blocks available according to this criterion. In this case it is suggested that the time value used be greater than the guard timing mentioned in Section 3, and that no block should ever be reassigned until it has been idle at least for the duration given by the guard timer.

Note that with the possible exception of cryptographically-based port allocations, attackers could reverse-engineer algorithmically-derived port allocations to either target a specific subscriber or to spoof traffic to make it appear to have been generated by a specific subscriber. However, this is exactly the same level of security that the subscriber would experience in the absence of CGN. CGN is not intended to provide additional security by obscurity.

While the block assignment strategy can provide some mitigation of the port guessing attack, the largest contribution will come from pseudo-randomization at port selection time. [RFC6056] provides a number of algorithms for achieving this pseudo-randomization. When the available ports are contained in blocks which are not in general consecutive, the algorithms clearly need some adaptation. The task is complicated by the fact that the number of blocks allocated to the user may vary over time. Adaptation is left as an exercise for the implementor.

5. IANA Considerations

This document makes no request of IANA.

6. Acknowledgements

This document is the result of a merger of the original draft-chen-sunset4-cgn-port-allocation and draft-tsou-behave-natx4-log-reduction. Version -02 of draft-chen contains the following acknowledgements:

The author would like to thank Lee Howard and Simon Perreault for their helpful comments.

Many thanks to Wesley George and Marc Blanchet encourage the author to continue this work.

The authors of draft-tsou-behave-natx4-log-reduction have their own thanks to give. Mohamed Boucadair reviewed the initial document and provided useful comments to improve it. Reinaldo Penno, Joel Jaeggli, and Dan Wing provided comments on the subsequent version that resulted in major revisions. Serafim Petsis provided encouragement to publication after a hiatus of two years.

The present version of the document benefited from further comments by Lee Howard, Mohamed Boucadair and Alberto Leiva.

7. References

7.1. Normative References

- [RFC6056] Larsen, M. and F. Gont, "Recommendations for Transport-Protocol Port Randomization", BCP 156, RFC 6056, DOI 10.17487/RFC6056, January 2011, <<http://www.rfc-editor.org/info/rfc6056>>.
- [RFC6145] Li, X., Bao, C., and F. Baker, "IP/ICMP Translation Algorithm", RFC 6145, DOI 10.17487/RFC6145, April 2011, <<http://www.rfc-editor.org/info/rfc6145>>.
- [RFC6146] Bagnulo, M., Matthews, P., and I. van Beijnum, "Stateful NAT64: Network Address and Protocol Translation from IPv6 Clients to IPv4 Servers", RFC 6146, DOI 10.17487/RFC6146, April 2011, <<http://www.rfc-editor.org/info/rfc6146>>.
- [RFC6269] Ford, M., Ed., Boucadair, M., Durand, A., Levis, P., and P. Roberts, "Issues with IP Address Sharing", RFC 6269, DOI 10.17487/RFC6269, June 2011, <<http://www.rfc-editor.org/info/rfc6269>>.

7.2. Informative References

- [I-D.ietf-behave-syslog-nat-logging]
Chen, Z., Zhou, C., Tsou, T., and T. Taylor, "Syslog Format for NAT Logging (Work in Progress)", January 2014.
- [I-D.ietf-pcp-port-set]
Sun, Q., Boucadair, M., Sivakumar, S., Zhou, C., Tsou, T., and S. Perrault, "Port Control Protocol (PCP) Extension for Port Set Allocation (Work in Progress)", October 2015.
- [I-D.ietf-softwire-stateless-4v6-motivation]
Boucadair, M., Matsushima, S., Lee, Y., Bonness, O., Borges, I., and G. Chen, "Motivations for Carrier-side Stateless IPv4 over IPv6 Migration Solutions (Expired work in Progress)", November 2012.
- [I-D.ietf-tsvwg-behave-requirements-update]
Penno, R., Perrault, S., Boucadair, M., Kamiset, S., and K. Naito, "Network Address Translation (NAT) Behavioral Requirements Updates (Work in Progress)", November 2015.
- [I-D.ietf-v6ops-siit-dc]
Anderson, T., "SIIT-DC: Stateless IP/ICMP Translation for IPv6 Data Centre Environments (Work in progress)", October 2015.
- [RFC1323] Jacobson, V., Braden, R., and D. Borman, "TCP Extensions for High Performance", RFC 1323, DOI 10.17487/RFC1323, May 1992, <<http://www.rfc-editor.org/info/rfc1323>>.
- [RFC4787] Audet, F., Ed. and C. Jennings, "Network Address Translation (NAT) Behavioral Requirements for Unicast UDP", BCP 127, RFC 4787, DOI 10.17487/RFC4787, January 2007, <<http://www.rfc-editor.org/info/rfc4787>>.
- [RFC5382] Guha, S., Ed., Biswas, K., Ford, B., Sivakumar, S., and P. Srisuresh, "NAT Behavioral Requirements for TCP", BCP 142, RFC 5382, DOI 10.17487/RFC5382, October 2008, <<http://www.rfc-editor.org/info/rfc5382>>.
- [RFC5424] Gerhards, R., "The Syslog Protocol", RFC 5424, DOI 10.17487/RFC5424, March 2009, <<http://www.rfc-editor.org/info/rfc5424>>.
- [RFC6191] Gont, F., "Reducing the TIME-WAIT State Using TCP Timestamps", BCP 159, RFC 6191, DOI 10.17487/RFC6191, April 2011, <<http://www.rfc-editor.org/info/rfc6191>>.

- [RFC6302] Durand, A., Gashinsky, I., Lee, D., and S. Sheppard, "Logging Recommendations for Internet-Facing Servers", BCP 162, RFC 6302, DOI 10.17487/RFC6302, June 2011, <<http://www.rfc-editor.org/info/rfc6302>>.
- [RFC6333] Durand, A., Droms, R., Woodyatt, J., and Y. Lee, "Dual-Stack Lite Broadband Deployments Following IPv4 Exhaustion", RFC 6333, DOI 10.17487/RFC6333, August 2011, <<http://www.rfc-editor.org/info/rfc6333>>.
- [RFC6346] Bush, R., Ed., "The Address plus Port (A+P) Approach to the IPv4 Address Shortage", RFC 6346, DOI 10.17487/RFC6346, August 2011, <<http://www.rfc-editor.org/info/rfc6346>>.
- [RFC6431] Boucadair, M., Levis, P., Bajko, G., Savolainen, T., and T. Tsou, "Huawei Port Range Configuration Options for PPP IP Control Protocol (IPCP)", RFC 6431, DOI 10.17487/RFC6431, November 2011, <<http://www.rfc-editor.org/info/rfc6431>>.
- [RFC6877] Mawatari, M., Kawashima, M., and C. Byrne, "464XLAT: Combination of Stateful and Stateless Translation", RFC 6877, DOI 10.17487/RFC6877, April 2013, <<http://www.rfc-editor.org/info/rfc6877>>.
- [RFC6887] Wing, D., Ed., Cheshire, S., Boucadair, M., Penno, R., and P. Selkirk, "Port Control Protocol (PCP)", RFC 6887, DOI 10.17487/RFC6887, April 2013, <<http://www.rfc-editor.org/info/rfc6887>>.
- [RFC6888] Perreault, S., Ed., Yamagata, I., Miyakawa, S., Nakagawa, A., and H. Ashida, "Common Requirements for Carrier-Grade NATs (CGNs)", BCP 127, RFC 6888, DOI 10.17487/RFC6888, April 2013, <<http://www.rfc-editor.org/info/rfc6888>>.
- [RFC7011] Claise, B., Ed., Trammell, B., Ed., and P. Aitken, "Specification of the IP Flow Information Export (IPFIX) Protocol for the Exchange of Flow Information", STD 77, RFC 7011, DOI 10.17487/RFC7011, September 2013, <<http://www.rfc-editor.org/info/rfc7011>>.
- [RFC7269] Chen, G., Cao, Z., Xie, C., and D. Binet, "NAT64 Deployment Options and Experience", RFC 7269, DOI 10.17487/RFC7269, June 2014, <<http://www.rfc-editor.org/info/rfc7269>>.

- [RFC7422] Donley, C., Grundemann, C., Sarawat, V., Sundaresan, K., and O. Vautrin, "Deterministic Address Mapping to Reduce Logging in Carrier-Grade NAT Deployments", RFC 7422, DOI 10.17487/RFC7422, December 2014, <<http://www.rfc-editor.org/info/rfc7422>>.
- [RFC7596] Cui, Y., Sun, Q., Boucadair, M., Tsou, T., Lee, Y., and I. Farrer, "Lightweight 4over6: An Extension to the Dual-Stack Lite Architecture", RFC 7596, DOI 10.17487/RFC7596, July 2015, <<http://www.rfc-editor.org/info/rfc7596>>.
- [RFC7597] Troan, O., Ed., Dec, W., Li, X., Bao, C., Matsushima, S., Murakami, T., and T. Taylor, Ed., "Mapping of Address and Port with Encapsulation (MAP-E)", RFC 7597, DOI 10.17487/RFC7597, July 2015, <<http://www.rfc-editor.org/info/rfc7597>>.
- [RFC7598] Mrugalski, T., Troan, O., Farrer, I., Perreault, S., Dec, W., Bao, C., Yeh, L., and X. Deng, "DHCPv6 Options for Configuration of Software Address and Port-Mapped Clients", RFC 7598, DOI 10.17487/RFC7598, July 2015, <<http://www.rfc-editor.org/info/rfc7598>>.
- [RFC7599] Li, X., Bao, C., Dec, W., Ed., Troan, O., Matsushima, S., and T. Murakami, "Mapping of Address and Port using Translation (MAP-T)", RFC 7599, DOI 10.17487/RFC7599, July 2015, <<http://www.rfc-editor.org/info/rfc7599>>.
- [RFC7600] Despres, R., Jiang, S., Ed., Penno, R., Lee, Y., Chen, G., and M. Chen, "IPv4 Residual Deployment via IPv6 - A Stateless Solution (4rd)", RFC 7600, DOI 10.17487/RFC7600, July 2015, <<http://www.rfc-editor.org/info/rfc7600>>.

Authors' Addresses

Gang Chen
China Mobile
29, Jinrong Avenue
Xicheng District,
Beijing 100033
China

Email: phdgang@gmail.com, chengang@chinamobile.com

Weibo Li
China Telecom
109, Zhongshan Ave. West, Tianhe District
Guangzhou 510630
P.R. China

Email: mweiboli@gmail.com

Tina Tsou
Huawei Technologies
Bantian, Longgang District
Shenzhen 518129
P.R. China

Email: tina.tsou.zouting@huawei.com

James Huang
Huawei Technologies
Bantian, Longgang District
Shenzhen 518129
P.R. China

Email: James.huang@huawei.com

Tom Taylor
PT Taylor Consulting
Ottawa, Ontario
Canada

Email: tom.taylor.stds@gmail.com

Jean-Francois Tremblay
Viagenie
246 Aberdeen
Quebec, QC G1R 2E1
Canada

Phone: +1 418 656 9254
Email: jean-francois.tremblay@viagenie.ca
URI: <http://viagenie.ca>