N. Elkins
Inside Products
L. Kratzke
IBM
M. Ackermann
BCBS of Michigan
K. Haining
US Bank

February 2013

IPv6 IPID Needed
draft-elkins-6man-ipv6-ipid-needed-00.txt

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), its areas, and its working groups. Note that other groups may also distribute working documents as Internet-Drafts.

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."

The list of current Internet-Drafts can be accessed at http://www.ietf.org/ietf/1id-abstracts.txt

The list of Internet-Draft Shadow Directories can be accessed at http://www.ietf.org/shadow.html

This Internet-Draft will expire on September 4, 2013.

Copyright Notice

Copyright (c) 2013 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.

Abstract

The IPv4 main header contained a 16-bit IP Identification (IPID) field used for fragmentation and reassembly. In practice, this field was commonly used by network diagnosticians for tracking packets. In IPv6, the IPID has been moved to the Fragment header, and would only be used when fragmentation is required. Thus, the IPID field in IPv6, is no longer able to be utilized in the valuable role it played in IPv4, relative to diagnostics and problem resolution. This causes great concern in particular for end users and large enterprises, for whom Network/Application availability and performance can directly and profoundly affect bottom line financials. Several viable solutions to this situation exist. One potential solution is included in later sections of this RFC, but the primary intention of this RFC is to initiate action by the IETF on this issue, perhaps in the form of a Working Group Subcommittee.

Table of Contents

1. Introduction

In IPv4, the 16 bit IP Identification (IPID) field is located at an offset of 4 bytes into the IPv4 header and is described in RFC791 [RFC791]. In IPv6, the IPID field is a 32 bit field contained in the Fragment Header defined by section 4.5 of RFC2460 [RFC2460]. Unfortunately, unless fragmentation is being done by the source node, the packet will not contain this Fragment Header, and therefore will have no Identification field.

The intended purpose of the IPID field is to enable fragmentation and reassembly, and as currently specified is required to be unique within the maximum segment lifetime (MSL) on all datagrams. The MSL is often **2 minutes.**

In practice, the IPID field is used for more than fragmentation. During network diagnostics, packet traces may be taken at multiple places along the path, or at the source and destination. Then, packets can be matched by looking at the IPID.

Obviously, the time at each device will differ according to the clock on that device; so another metric is required. This method of taking multiple traces along the path is of special use on large multi-tier networks to see where the packet loss or packet corruption is happening. Multi-tier networks are those which have multiple routers or switches on the path between the sender and the receiver.

The inclusion of the IPID makes it easier for a device(s) in the middle of the network, or on the receiving end of the network, to identify flows belonging to a single node, even if that node might have a different IP address. For example, if the sending node is a mobile laptop with a wireless connection to the Internet.

For its de-facto diagnostic mode usage, the IPID field needs to be available whether or not fragmentation occurs. It also needs to be unique in the context of the entire session, and across all the connections controlled by the stack.

This document will present information that demonstrates how valuable and useful the IPID field has been (in IPv4) for diagnostics and problem resolution, and how not having it available (in IPv6), could be a major detriment to new IPv6 deployments and contribute to protracted downtimes in existing IPv6 operations. A possible solution to this situation will be suggested, but the primary intention of this document is to highlight the existence of this issue and to ask that the IETF research and recommend an optimal solution, perhaps by the formation of a Working Group Subcommittee. End users and Large Enterprises involved in this initiative to date seem to agree on the need to retain this valuable diagnostic facility, as well as the preference that it be able turned on and off as needed, if it cannot be statically included in the base header, as it is in IPv4.

2. Conventions used in this document

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 RFC2119 [RFC2119].

3. Applicability

The base IPv6 standard, RFC2460, [RFC2460] allows the use of extension headers, such as the Destination Options Header, in order to encode optional destination information in an IPv6 packet. Extended diagnostic information such as this MUST be sent by implementations upon request. The example solution is an implementation of the Destination Options Header. Once again, this approach assumes the IPID field cannot be included in the base header, as it is in IPv4.

BASIC RATIONALE FOR RETENTION OF THE IPID FIELD

1. **The ability to utilize the IPID has enhanced problem diagnosis** efforts and significantly reduced problem resolution time.

2. **Several actual use case examples are shown below.** These demonstrate how use of the IPID has reduced problem resolution time in very valuable production networks of Large Enterprises/End Users. In general, if a problem or performance issue with an application or network component can be fixed in minutes, as opposed to hours, this can mean significant dollar savings to large enterprises. The IPID can be used extensively when debugging involves traces or packet captures. Its absence in IPv6 may lead to protracted problem diagnosis and extended problem resolution time.

3. **This value/perspective may be unique to tech support organizations** of large enterprises. Other functional areas may not share this concern/perspective, as packets could continue to flow, but service levels may not be acceptable to end users during the extended problem resolution time.

4. **Although very situation dependent, the use cases below clearly** illustrate the value of network availability, and the need to keep problem resolution time to an absolute minimum.

5. **Another benefit of using the IPID to expedite problem resolution** is reducing the cost of associated resources being consumed during extended problem resolution, such as storage, CPU and staff time.

6. **Will IPID be critical in most problem resolutions?** NO! But if it even helps in a few per year, significant money and/or lost business could be saved.

7. A facility such as IPID, that has proven field value, should not be eliminated as an effective diagnosis tool!

USE CASE EXAMPLES.

USE CASE #1 --- Large Insurance Company
- (estimated time saved by use of IPID: 7 hours)

PERFORMANCE TOOL PRODUCES EXTRANEOUS PACKETS?
- Issue was whether a performance tool was accurately replicating session flow during performance testing?
- Trace IPIDs showed more unique packets within same flow from performance tool compared to IE Browser.
- Having the clear IPID sequence numbers also showed where and why the extra packets were being generated.
- Solution: Problem rectified in subsequent version of performance tool.
- Without IPID, it was not clear if there was an issue at all.

USE CASE #2 --- Large Bank
- (estimated time saved by use of IPID: 4 hours)

BATCH TRANSFER DURATION INCREASES 12X
- A 30 minute data transfer started taking 6-8 hours to complete.
- Possible packet loss? All vendors said no.
- Other Apps were working OK.
- 4 trace points used, and then IPIDs compared.
- Showed 7% packet loss.
- Solution: WAN hardware was replaced and problem fixed.
- Without IPID, no one would agree a problem existed

USE CASE #3 --- Large Bank
- (estimated time saved by use of IPID: 6 hours)

VERY SLOW INTERACTIVE PERFORMANCE.
- All network links looked good.
- Traces showed duplicated small packets (which can be OK).
- Saw that IPID was equal but TTL was always + 1.
- Network device was "Splitting" small packets only.(2 interfaces).
- The small packets were control info, telling other side to slow down.
- Erroneously looked like network congestion.
- Solution: Network Device replaced and good interactive performance restored.
- Without IPID, flows would have appeared OK.

USE CASE #4 --- Large Government Agency
- (estimated time saved by use of IPID: 9 hours)

VPN DROPS
- Cell phone connection to law enforcement were being dropped. Going through a VPN.
- All parties (both sides of VPN connection, application, etc.) said it was not their problem. Problem went on for weeks.
- Finally, when we were called in as consultants, we took a trace which showed packet with IPID and TTL that did not match others in the flow AT ALL was coming from router nearest application server end of VPN.
- Solution: Provider for VPN for application server changed. Problem resolved.
- Without IPID, much harder to diagnose problem.
- (Same case also happened with large corporation. Again, all parties saying not their fault until proven via packet trace.)

4. IPv6 Diagnostic Option Format

One Possible Implementation - Example Solution

4.1 Destination Options Header

The Destination Options Header is used to carry optional information that need be examined only by a packet's destination node(s). The Destination Options Header is identified by a Next Header value of 60 in the immediately preceding header and is defined in RFC2460 [RFC2460].

4.2. IPv6 Diagnostic Option

The IPv6 Diagnostic Option is used in a packet sent by a node to facilitate diagnostics by informing the recipient and passive viewers of the packet, such as packet capture facilities, of the packet's IP Identifier.

The IPv6 Diagnostic Option is encoded in type-length-value (TLV) format as follows:

```
 0                   1                   2                   3
 0 1 2 3 4 5 6 7 8 9 0 1 2 3 4 5 6 7 8 9 0 1 2 3 4 5 6 7 8 9 0 1
+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+
|  Option Type  | Option Length |
+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+
|                                                               |
+                                                               +
|                                                               |
+                          IP Identifier                        +
|                                                               |
+                                                               +
|                                                               |
+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+
```

Option Type

TBD = 0xXX (TBD) [To be assigned by IANA] [RFC2780]

Option Length

8-bit unsigned integer. Length of the option, in octets, excluding the Option Type and Option Length fields. This field MUST be set to 64.

IP Identifier

The IP Identifier of the packet for 64 bits.

The alignment requirement for the IP Identifier option is 8n+6.

The two highest-order bits of the Option Type field are encoded to indicate specific processing of the option; for the IP Identifier option, these two bits MUST be set to 00. This indicates the following processing requirements:

- 00
  - skip over this option and continue processing the header.
  - The data within the option cannot change en route to the packet's final destination.

The IPv6 Diagnostic Option MUST be placed as follows:

- After the Routing Header, if that header is present

- Before the Fragment Header, if that header is present

- Before the AH Header or ESP Header, if either one of those headers are present.

For each IPv6 packet header, the IPv6 Diagnostic Option MUST NOT appear more than once. However, an encapsulated packet MAY contain a separate IPv6 Diagnostic Option associated with each encapsulating IP header.

The inclusion of a IPv6 Diagnostic Option in a packet affects the receiving node's processing only for this single packet. No state is created or modified in the receiving node as a result of receiving a IPv6 Diagnostic Option in a packet.

4.3. Implementation Considerations

In implementation, a given OS/Stack may send this additional header for all connections, per higher level protocol, or in a more sophisticated usage, for a single connection or flow only.

The initiation of this header would preferably be done via a 'Debug on'/'Debug off' switch. That is, a diagnostician might decide that this header is required for a certain timeframe or for a certain set of packets after a network problem is encountered. The diagnostician would then issue a command to the stack indicating that addition of the IP Identifier header should now begin. This is the 'Debug on' state. After a certain amount of time, then 'Debug off' should be issued as a command. Alternatively, the stack might have a fixed time (for example, 5 minutes), after which debug mode will automatically be turned off. In their default configuration, IPv6 nodes SHOULD NOT include this option in IPv6 packets that they originate. That is, the switch or state SHOULD default to 'Debug off'.

Additionally, implementers MUST permit a system administrator to enable or disable including this option in originated IPv6 packets on at least a per-Upper-Protocol basis (e.g. at least provide enable/disable separate knobs for ICMP, UDP, TCP, or SCTP, in addition to a knob to enable or disable for all IPv6 packets), and SHOULD permit it also to be enabled or disabled on a per-flow basis. This configuration flexibility would increase the potential value of this new option, and would not increase implementation complexity unduly.

Please note that the ability to turn extended diagnostic information on or off, as appropriate, is very prevalent in many situations: servers, applications, network devices, etc.

Additionally, we suggest that stacks implement a separate IPID counter per connection, rather than a single counter for all connections.

5. Backward Compatibility

The example solution described in the previous sections of this document is backward compatible with all the currently defined IPv6 extension headers. According to RFC2460 [RFC2460], if the destination node does not recognize this option, it should skip over this option and continue processing the header.

6. Security Considerations

The example solution might be useful to security gateways seeking to identify operational security issues or in preventing Replay Attacks.

7. IANA Considerations

A Destination Option requires an IPv6 Option Number or Type [RFC2460] which is 8 bits.

```
HEX.........act..chg..rest (5 bits)
----        ---- ---  ----
TBD          00   0    TBD   IPv6 Diagnostic Option
```

For the IPv6 Option Number (Type), the first two bits indicate that the IPv6 node should skip over this option and continue processing the header if it does not recognize the option type. The third bit indicates that the Option Data must not change en-route, which permits this option to be protected by the IP Authentication Header.

10. References

10.1. Normative References

[RFC791] Postel, J., "Internet Protocol", RFC 791 / STD 5, September 1981.

[RFC2460] Deering, S. and R. Hinden, "Internet Protocol, Version 6 (IPv6) Specification", RFC 2460, December 1998.

[RFC2119] Bradner, S., "Key words for use in RFCs to Indicate Requirement Levels", BCP 14, RFC 2119, March 1997.

[RFC2780] Bradner, S., Paxson, V. "IANA Allocation Guidelines For Values In the Internet Protocol and Related Headers", RFC 2780, March 2000.
See also:
http://www.iana.org/assignments/ipv6-parameters

11. Acknowledgments

The authors would like to thank Fred Baker, Bill Jouris, Jose Isidro, **R. J. Atkinson, and James Ashton for their reviews and suggestions** that made this document better.

This document was prepared using 2-Word-v2.0.template.dot.

Authors' Addresses

Nalini Elkins
Inside Products, Inc.
36A Upper Circle
Carmel Valley, CA 93924
United States

Phone: +1 831 659 8360
Email: nalini.elkins@insidethestack.com

Lawrence Kratzke
IBM
8121 Glenbrittle Way
Raleigh, NC 27615
United States

Phone: +1 800-876-8801
Email: kratzke@us.ibm.com

Michael Ackermann
Blue Cross Blue Shield of Michigan
P.O. Box 2888
Detroit, Michigan 48231
United States

Phone: +1 310 460 4080
Email: mackermann@bcbsmi.com

Keven Haining
US Bank
16900 W Capitol Drive
Brookfield, WI 53005

Phone: +1 262-790-3551
Email: keven.haining@usbank.com