



HURRICANE ELECTRIC
INTERNET SERVICES

Jumbo Frame Deployment at Internet Exchange Points (IXPs)

draft-mlevy-ixp-jumboframes-00.txt

Hurricane Electric

IPv6 Native Backbone – Massive Peering!

9,000 Bytes – It's a good idea.

LINX 75

London, UK – 22th November 2011

Martin J. Levy, Director IPv6 Strategy

Hurricane Electric

draft-mlevy-ixp-jumboframes-00

NATIVE IPv6
EVERYWHERE

Abstract

This document provides guidelines on how to deploy Jumbo Frame support on Internet Exchange Points (IXP). Jumbo Frame support allows packets larger than 1,500 Bytes to be passed between IXP customers over the IXPs layer 2 fabric. This document describes methods to enable Jumbo Frame support and keep in place existing 1,500 Byte communications.

This document strongly recommends that IXP operators choose 9,000 Bytes for their Jumbo Frame implementation.

<http://tools.ietf.org/html/draft-mlevy-ixp-jumboframes-00>



draft-mlevy-ixp-jumboframes-00

NATIVE IPv6
EVERYWHERE

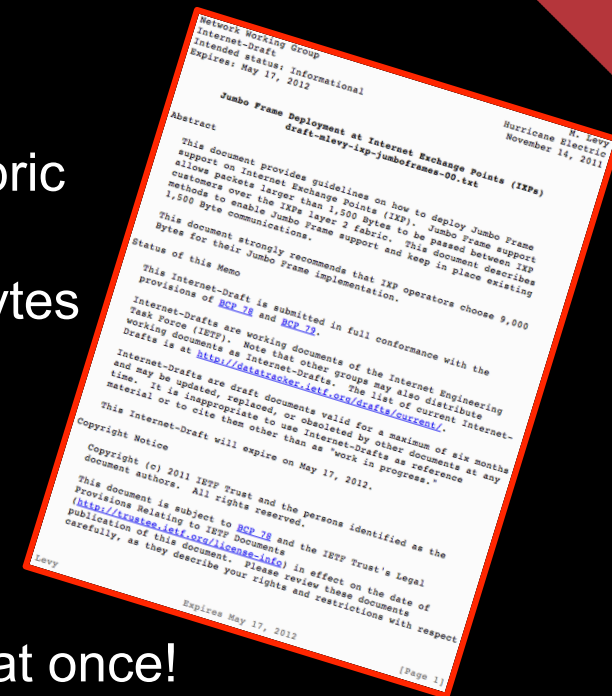
- Internet Exchange Points (IXPs) are vital IP backbones interconnect
 - All IXPs are Ethernet layer-2 based these days
 - 99% of IXPs default to 1,500 Bytes MTU (1,514 Frame)
 - A few IXPs that run high MTUs (NETNOD & NASA AIX)
- No documentation exists on how IXPs should provide large MTU services
 - No standard on large MTU size
 - Even the name “Jumbo Frames” is agreed upon
- Draft has written with input from various operators
 - Been “sitting around for nearly a year” mulling it over
 - Sent for review within IXP industry before submitting to IETF



draft-mlevy-ixp-jumboframes-00

NATIVE IPv6
EVERYWHERE

- Points covered:
 - Force MTU to be consistent over layer-2 fabric
 - Propose a consistent MTU value – 9,000 Bytes
 - Propose various solutions to implement
 - VLAN based (NETNOD & NASA AIX works)
 - Duplicate fabric hardware (too expensive!)
 - Explain how to do a “flag day” – all change at once!
 - Explain how to do BGP – prefer 9,000 Byte path
 -
 - Explain pitfalls – misconfigured layer-2 failures



Why 9,000 Bytes?

- Too many choices for size based on h/w vendor or h/w version ...
 - 9,000 Bytes
 - 9,170 Bytes
 - 9,174 Bytes
 - 9,180 Bytes
 - 9,192 Bytes
 - 9,216 Bytes


} Confusion or misconfiguration or both
- The number 9,000 is easy to remember – less confusion
 - 9,000 Bytes – handle at least 8,192 Bytes user data plus TCP/IP header
- See References in draft [JET2007]

Why 9,000 Bytes?

- Large packets are needed for:
 - Mass Data Replication – datacenter X to datacenter Y
 - Storage – Amazon S3, MS Azure Storage, iCloud by Apple
 - NNTP, DNS zone xfer's, etc
- IXPs are in the path for any-to-any communications
 - This draft addresses only ONE element within a data path
- Also allows BGP sessions at IXPs to operate with large MSS

■ Summary:

- Large MTU end-to-end traffic can optimize traffic flow and reduce packet overhead
- IXPs are a key part of the global routing system
- Someone has to document how to do this
- Not an end-to-end panacea – but a good start



Contact:

Martin J. Levy
Director, IPv6 Strategy
Hurricane Electric
760 Mission Court
Fremont, CA 94539, USA
<http://he.net/>

martin at he dot net
+1 (510) 580 4167