BIER BOF

IETF 91 – Honolulu, Hawaii

BoF

Chairs:

Greg Shepherd gjshep@gmail.com

Jeffrey Tantsura jeff.tantsura@ericsson.com


NL
1st BoF
BIER
BIT INDEXED EXPLICIT REPLICATION
bier@ietf.org - http://trac.tools.ietf.org/bof/trac/
IETF Nov 2015
Honolulu
Hawaii

Agenda

Welcome, note well, house keeping, agenda bash, Chairs - 5min

Purpose and Scope of this BoF, Chairs - 5min

draft-shepherd-bier-problem-statement, Greg Shepherd - 15min

Open discussion on problem statement – Room

Challenged multicast use cases – Chairs

To be revisited throughout the BoF

Proposed solutions:

draft-wijnands-bier-architecture, IJsbrand Wijnands - 30min (or as needed)

Open discussion questions:

What is the cost of a new data-plane behavior?

Does this problem warrant a new data-plane behavior?

Do we feel the potential benefits justify the work?

draft-kumar-bier-use-cases, xuxiaohu@huawei.com - 10min

draft-wijnands-mpls-bier-encapsulation, Jeffrey Zhang, zzhang@juniper.net - 15min

draft-rosen-l3vpn-mvpn-bier, Mahesh Sivakumar - 15min

draft-psenak-ospf-bier-extensions, Peter Psenak - 10min

draft-przygienda-bier-isis-ranges, A. Przygienda - 10min

IETF Notewell

Any submission to the IETF intended by the Contributor for publication as all or part of an IETF Internet-Draft or RFC and any statement made within the context of an IETF activity is considered an "IETF Contribution". Such statements include oral statements in IETF sessions, as well as written and electronic communications made at any time or place, which are addressed to:

The IETF plenary session

The IESG, or any member thereof on behalf of the IESG

Any IETF mailing list, including the IETF list itself, any working group or design team list, or any other list functioning under IETF auspices

Any IETF working group or portion thereof

Any Birds of a Feather (BOF) session

The IAB or any member thereof on behalf of the IAB

The RFC Editor or the Internet-Drafts function

All IETF Contributions are subject to the rules of RFC 5378 and RFC 3979 (updated by RFC 4879).

Statements made outside of an IETF session, mailing list or other function, that are clearly not intended to be input to an IETF activity, group or function, are not IETF Contributions in the context of this notice. Please consult RFC 5378 and RFC 3979 for details.

A participant in any IETF activity is deemed to accept all IETF rules of process, as documented in Best Current Practices RFCs and IESG Statements.

A participant in any IETF activity acknowledges that written, audio and video records of meetings may be made and may be available to the public.

Housekeeping

- WebEx

  Monday, November 10, 2014

  1:00 pm | Hawaii Time (Honolulu, GMT-10:00) | 2 hr

  Join WebEx meeting:
  https://workgreen.webex.com/workgreen/j.php?MTID=mdfb183b4cd0d2cfaeceacf3950895a00

  Meeting number: 825 308 148

  Meeting password: 1234

- Scribe? (Bier?)

Purpose of BoF

- Community input for and against please!
- Is this a problem worth solving?
- Does this problem warrant a new data-plane behavior?
- What is the cost of a new data-plane behavior?
- Do we feel the potential benefits justify the work?
- Is there a critical mass of people willing to work on the solution?

BIER BOF

IETF 91 – Honolulu, Hawaii

BIER Problem Statement
draft-shepherd-bier-problem-statement

Greg Shepherd gjshep@gmail.com

Arkadiy Gulko arkadiy.gulko@thomsonreuters.com

A Brief History of Multicast

- Steven Deering, 1985, Stanford University

  Yeah, he was way ahead of his time and too clever for all of us.

- A solution for layer2 applications in the growing layer3 campus network

  Think overlay broadcast domain

- Broadcast Domain

  all members receive

  all members can source

  members dynamically come and go

A Brief History of Multicast

RFC966 - 1985

```
Multi-destination delivery is useful to several applications,
including:

      - distributed, replicated databases [6,9].

      - conferencing [11].

      - distributed parallel computation, including
distributed gaming [2].
```

All inherently many-to-many applications

No mention of one-to-many services such as Video/IPTV

A Brief History of Multicast

Overlay Broadcast Domain Requirements

- Tree building and maintenance
- Network-based source discovery
- Source route information
- Overlay mechanism – tunneling

The first solution had it all

Distance Vector Multicast Routing Protocol

DVMRP, RFC1075 – 1988

A Brief History of Multicast

PIM – Protocol Independent Multicast

"Independent" of which unicast routing protocol you run

It does require that you're running one. ☺

Uses local routing table to determine route to sources

Router-to-router protocol to build and maintain distribution trees

Source discovery handled one of two ways:

1) Flood-and-prune PIM-DM, Dense Mode

2) Explicit Join w/ Rendezvous Point (RP) PIM-SM, Sparse Mode - The Current Standard

Brief History of Multicast

PIM-SM – Protocol Independent Multicast Sparse Mode

- Tree building and maintenance
- Network-based source discovery
- ~~Source route information~~
- ~~Overlay mechanism – tunneling~~

Long, Sordid IETF history

RFC4601 – 2006 (original draft was rewritten from scratch)

Primary challenges to the final specification were in addressing Network-based source discovery.

Brief History of Multicast

Today's dominant applications are primarily one-to-many

IPTV, Contribution video over IP, etc.

Sources are well known

SSM – Source Specific Multicast

RFC3569, RFC4608 – 2003

- Tree building and maintenance
- ~~Network-based source discovery~~
- ~~Source route information~~
- ~~Overlay mechanism – tunneling~~

Very simple and the preferred solution for one-to-many applications

Where are we?

- Multicast solutions have successfully devolved into a more simple subset-solution from the original architecture.
- Many operators jaded by the original complexity still think of multicast as failing the cost/benefit analysis.
- Operators with must-have multicast requirements maintain robust, stable multicast networks which continue to grow, exposing the limitations of the current architecture.

Problem Statement

- Current multicast methods all require explicit tree building protocols, thereby incurring a lot of state in the transit nodes.
- Current multicast methods, if they are to provide optimal delivery of multicast packets, require one explicitly built tree per multicast flow; there is no reasonable way to aggregate flows (having one state for multiple flows) without sacrificing optimal delivery.
- Comments?
- Challenged use cases?
  - MVPN
  - IPTV
  - Financial
  - DataCenter


NL
1st BoF
BIER
BIT INDEXED EXPLICIT REPLICATION
bier@ietf.org - http://trac.tools.ietf.org/bof/trac/
IETF Nov 2015
Honolulu
Hawaii