

draft-dickson-idr-last-resort

LAST_RESORT: A New BGP Well-Known Community

Brian Dickson
briand@ca.afilias.info

Problem Statement

- There is a well-known need for Networks to advertise “back-up only” paths
- This can be seen empirically by observing the techniques used: All manner of BGP communities, AS-path prepending (up to 9 times in some cases), and de-aggregation
- These techniques arose because of lack of direct relationships to third parties, lack of consensus, and lack of prevailing standard to address the problem

Problem Statement (cont)

- The real problem that continues to elude solution, is that of “BGP Wedgies” - non-deterministic, stable network states which are unintended
- Wedgies exist even when parties involved wish them to not exist
- Primarily these are caused by local policy mechanisms over-riding intended global policy on announced prefixes

Proposed Solution

What is being proposed, is a new Well-Known BGP Community.

There are two reasons for it to be Well-Known:

- It is very difficult to ask third parties to standardize on something that isn't truly standard
- The optimal goal, universal support, is only realistically achievable via router software enforcement

Counter-Arguments and Answers

- Operators don't like being told what to do
 - They aren't being told; they already want to achieve the results
- Vendors won't support it
 - Vendor support isn't strictly necessary
- Will break things
 - No! Incremental deployment is pain-free
- Requires everyone to support it
 - ...Before all Networks can remove existing hacks

Incremental Deployment

Two things are necessary for “Last Resort”:

- Prefix is announced with `LAST_RESORT`
- Receiver applies policy (either manually, or via router code)

If a prefix is not tagged, adding the policy will do nothing.

If a prefix is tagged, but no policy exists, nothing will happen.

If a “hack” is already in place, nothing will happen.

If “none of the above”, tag+policy = demote route

=> But this can only happen *deliberately!*

Incremental Deployment (cont)

Once a prefix is tagged with the community, the scope of the effect is limited by the topology of deployed policy changes (manual or software upgrade)

Once prefixes are tagged, operators will have visibility on the need to deploy

Incremental deployment will benefit operators, as they can reduce the “hacks” that don't scale well

“Peer pressure” is likely to influence deployment

Standardization is Key

Currently, the mechanisms used are a result of attempting to bypass local policy mechanisms

Too many proprietary systems exist

The result is a “tragedy of the commons”; nobody wants the current situation, but there's no easy fix without a standard mechanism.

Having a standard solution is like having a hammer: When you have a hammer, everything starts to look like a nail.

Comparison with other Well Knowns

Of the existing Well-Known Communities, only the NO_EXPORT community is widely used.

Only one major ISP (DT) blocks NO_EXPORT, and then only by ignoring announcements which have NO_EXPORT.

No one *strips* the NO_EXPORT community deliberately (to the best of my knowledge)

Supported directly by all major routing vendors

Comparison (cont)

LAST_RESORT:

- Has very clear semantics
- Meets a demonstrated need in the community
- Fixes BGP Wedgies
- May reduce prevalence of prepending and deaggregation over the long term
- Applies a global-scope (transitive) mechanism to a global-scope problem
- Is the right tool for the right job
- Mirrors NO_EXPORT in most regards

Summary

- Incremental to current BGP standard
- Standardizes a universally-usable solution
- Meant to replace a plethora of ad-hoc, marginally-effective, unscalable, resource-intensive hacks
- When used, solves BGP Wedgies problem
- Next steps?

Thank you

- Presentation on IETF 72 site
- Draft updates (if needed) will continue to be prompt
- Author: Brian Dickson, briand@ca.afilias.info