BGP Failure Propagation for Enhancing Control-Plane Convergence

Problem statement and request for WG feedback

draft-li-idr-bgp-failure-propagation-convergence-01

Yuhang Li	Tsinghua University (presenter)
Yahui Li	Tsinghua University
Xia Yin	Tsinghua University
Han Zhang	Tsinghua University
Xingang Shi	Tsinghua University

IDR IETF125 Mar 2026

Motivation

- BGP convergence is a long-standing operational challenge.
- A valid BGP withdrawal does not always lead to immediate convergence.
- After a route is withdrawn, downstream ASes may successively select other currently visible paths, only to later discover that those paths are also invalid.
- This can lead to:
  - multiple rounds of path exploration (path hunting)
  - extra UPDATE/WITHDRAW churn
  - slower final convergence
  - transient loss and latency before a stable state is reached

Understanding this behavior remains an important operational concern in inter-domain routing.

Background

- BGP is a path distance-vector protocol.
- When a route is withdrawn:
  - each AS only knows the paths it has already learned
  - routers select the current next-best path
  - alternative paths may later also be withdrawn or invalidated
  - path exploration can therefore continue for multiple rounds

This behavior is a known consequence of incremental, asynchronous inter-domain routing.

Prior Work and Operational Evidence

- This behavior has been studied and observed for decades:
  - early operational experience already identified slow convergence as a real BGP issue
  - research literature documented path exploration after route withdrawal
  - operator communities continue to observe churn, repeated updates, and delayed convergence after failures
- Examples include RFC discussion, measurement studies, RIPE operator reports, NANOG discussions, and recent operational case studies.

1. [RFC 4984] Report from the IAB Workshop on Routing and Addressing.
2. C. Labovitz, A. Ahuja, A. Bose and F. Jahanian, "Delayed Internet routing convergence," in *IEEE/ACM Transactions on Networking,* June 2001.
3. Smith, Philip. “[routing-wg] Draft doc proposing Route Flap Damping Obsolesence.” ,RIPE Routing Working Group Mailing List, 24 Apr. 2006,
4. Jaeggli, Joel. “Re: Soliciting Your Opinions on Internet Routing: A Survey on BGP Convergence.” NANOG Mailing list , 9 Jan. 2017.
5. Tuber, David, Dzevad Trumic. “Cloudflare Outage on February 20, 2026.” Cloudflare Blog, 21 Feb. 2026,

Failure Scenarios

- A route is correctly withdrawn according to normal BGP semantics
- The receiving AS removes the route and runs best-path selection again
- However, other candidate paths may still appear usable at that moment
- Those candidates can later also be invalidated as more updates arrive
- Convergence is eventually reached, but only after several rounds of exploration

This is not a stuck-route or zombie-route case.

Signaling is invalid. The problem is delayed elimination of invalid alternatives.

It is a normal but expensive convergence process after a valid withdrawal.

Failure Scenarios — Examples

- AS1 announces a prefix1 p/23 to AS3
- AS1 announces a more specific prefix2 p/24 to AS2
- When AS1 withdraws prefix2 p/24
- AS3 will have path exploration (hunt):

  <2,1> <5,2,1> <4,5,2,1>
- AS3 still has reachability to p/23

before p/24 converges, AS3 continues to prefer AS2.

This will affect the routing of prefix p

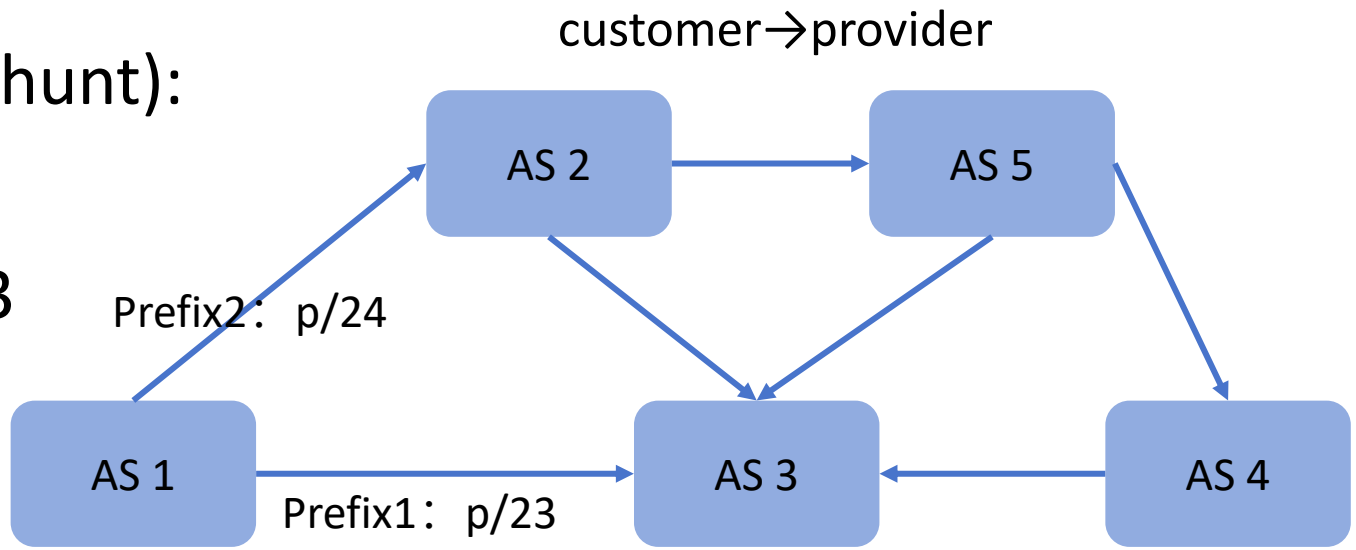

Prefix p/24
Path 1： <2,1>
Path 2： <5,2,1>
Path 3： <4,5,2,1>

Problem Scope

How can we reduce unnecessary path exploration after a valid route withdrawal?

- We do not focus on:
  - failed session liveness detection
  - permanently stuck or zombie routes
  - policy-correct route rejection
  - general security filtering behavior
- We focus on the case where:
  - withdrawal is valid
  - routing state is updated normally
  - but convergence still takes multiple rounds

Goal of This Work

The goals of this exploration are:

- Characterize multi-round path exploration after valid withdrawals
- Understand when alternative paths are explored unnecessarily
- Evaluate whether earlier invalidation knowledge could reduce convergence cost
- Do so without changing the core semantics of BGP withdrawals

I am not asking the WG to evaluate a protocol design today.

I am asking whether this scoped convergence problem is real, important, and worth further study in IDR, and what operational scenarios should ground it.

Design Question

After a valid withdrawal, can downstream ASes avoid exploring alternate paths that are already doomed to fail?

More specifically:

- Can invalid alternatives be recognized earlier?
- Can convergence reach the final state with fewer exploration rounds?
- Can this be achieved with operationally realistic mechanisms?

This is the design question we want to discuss with the WG, rather than proposing a fixed mechanism today.

Discussion Points

We would appreciate feedback on the following questions:

- Is this a meaningful and well-scoped convergence problem?
- Are multi-round exploration events operationally important enough to study further?
- Which real-world scenarios best represent this problem?
- What kinds of mechanisms would be operationally realistic?

Next Steps

- Refining the problem definition
- Studying how often valid withdrawals trigger multiple exploration rounds
- Evaluating candidate mechanisms for reducing redundant exploration

We welcome suggestions and guidance from the WG.