Workgroup: idr
Internet-Draft:
draft-cheng-idr-redirection-risks-ps-02
Published: 13 March 2023
Intended Status: Standards Track
Expires: 14 September 2023
Authors: W. Cheng, Ed. D. li Y. liu
China Mobile Tsinghua University China Mobile
M. Huang F. Gao S. Chen
Huawei Zhongguancun Laboratory Huawei

Problem statement of Inter-domain Traffic Redirection Risks

Abstract

Redirection of network traffic on the Internet is a common technology. In operator network, there are complex scenarios, such as multi-domain interconnection and large scale network topology. Typo, limitation of out-of-band tool capabilities for configuration verification, network adjustment or failure may cause risks, such as traffic detour, traffic exposure, traffic black hole and traffic loops. This draft describes these risks.

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). Note that other groups may also distribute working documents as Internet-Drafts. The list of current Internet-Drafts is at https://datatracker.ietf.org/drafts/current/.

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."

This Internet-Draft will expire on 14 September 2023.

Table of Contents

1. Introduction

There are many network traffic redirection technologies[RFC1102], such as Policy Based Routing[RFC1104] and BGP Flow Specification Rules[RFC8955] etc. These technologies are widely used in carrier networks.

For example, BGP Flow Specification provides various filtering conditions and processing actions to implement traffic control[RFC9117]. This function is used not only to protect the device against denial-of-service (DoS) and distributed DoS (DDoS) attacks, but also used for network traffic optimization.

As reliability and effectiveness of traffic redirection are artificially guaranteed, there are risks, such as traffic detour, traffic exposure, traffic black hole, traffic loop, as well as inconsistent traffic paths between the control plane and data plane.

Currently, operators have applied the redirection technologies, such as BGP Flowspec, on a large scale. How to deal with these risks needs further discussion.

2. Risks Description

The operation and maintenance of redirection in the long run is a big challenge. Typo, limitation of out-of-band tool capabilities for configuration verification, network adjustment or failure may cause potential problems without system awareness.

The following figure shows the risks of traffic redirection.

Take the topology in Figure 1 as an example, In this application scenario, one campus network inter-connects to two providers Network. AS 65003 and AS 65500, AS 65105 and AS 65500 form a

provider-customer adjacency relationship. Assume that the user needs to transmit data to the server. According to the routing information on the control plane, the traffic is transmitted through the path [User -- AS65001 -- AS65003 -- AS65104 -- AS65106 -- Server].

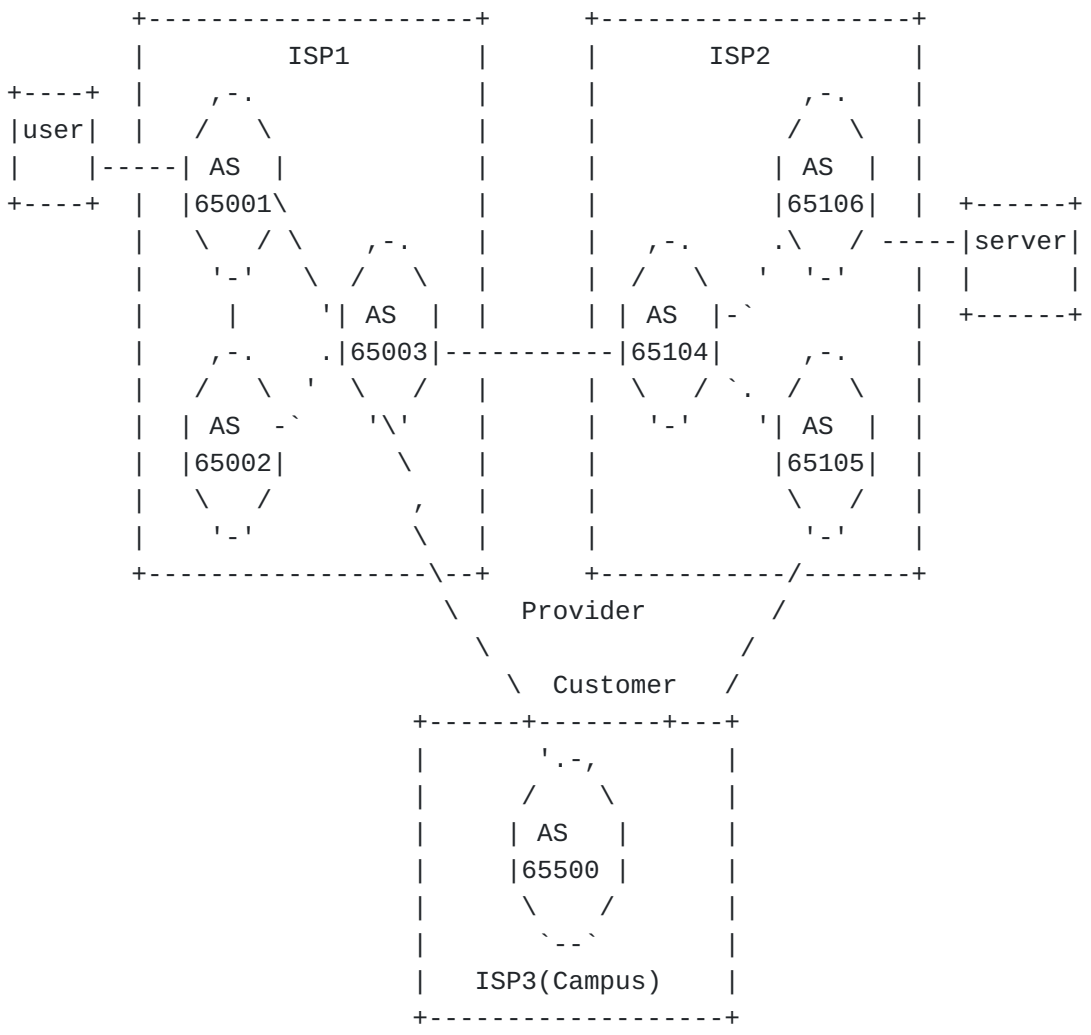

Figure 1: Example of the risks of traffic redirection

Risk 1: Violation of the valley-free principle[RFC7908] leads to traffic detour and exposure.

Assume ISP1 configures wrong traffic redirection rule, as a result, AS 65003 redirects traffic from AS 65104 to AS 65500. In this case, valley-free principle is violated as AS 65003 and AS 65500 form a provider-customer adjacency relationship. Traffic passes through the AS 65500 and exposes itself to the campus network.

Risk 2: The traffic is redirected to a network node that has no route, which leads to a traffic black hole

After traffic is redirected from AS 65003 to the AS 65500, the AS 65500 may not have a route to the destination server. In this case, the traffic is discarded, resulting in a traffic black hole.

Assume traffic is redirected from AS65003 to AS65500, and AS65500 learns the route to server from AS65105. After the traffic is redirected to the AS65500, it is forwarded to the server according to the route. If the link between AS 65500 and AS 65105 fails, as a result, the route is lost. The traffic is redirected to AS 65500 anyway, which also causes a traffic black hole.

Risk 3: One or more traffic redirections on the traffic transmission path may cause a traffic loop.

Assume that traffic is redirected from AS 65104 to AS 65105. In AS 65105, the traffic is transmitted to AS 65500, or the traffic is redirected to AS 65500. In this case, if the traffic on AS 65500 is transmitted to AS 65003, then traffic loop occurs.

Assume traffic is redirected from AS 65003 to AS 65500, and AS 65500 learns the route to server from AS 65105. After receiving the redirected traffic, the AS65500 forwards the traffic to the server according to the route. If the link between AS6500 and AS65105 fails, AS 65500 learns that the next hop of the route to the server is AS65003 through routing protocols. If the traffic is still redirected to AS 65500 at this time, AS 65500 will send packets back to AS65003, then traffic loop occurs.

Risk 4: inconsistent paths on the control plane and data plane may cause O&M risks.

The traffic owner expects traffic to be transmitted along the AS path carried in the route, but the actual transmission path is different from the AS path.

If the network O&M control system does not obtain traffic redirection information on the network, unpredictable risks may occur during traffic optimization, for example, network congestion.

For the risks mentioned above, it is not sufficient to rely on manual operation or automated management. Traffic redirection increases management difficulties and effectiveness requirements. It is necessary to explore technical solutions, such as redirection configuration verification, protocol extension, or path visualization, to reduce risks mentioned above.

3. Valuable Scenarios and Potential Solutions

Since BGP cannot perceive the AS_path generated by the inter-domain traffic redirection, the security of traffic redirection only relies

on human decision, which inevitably increases the risks. Below is three potential solutions.

Solution 1: AS path validation

BGP flowspec only considers the verification of the source of the flow specification.A potential solutions is to add a redirection path perception capability to BGP, and verify the validity of the redirection path before actually configuring the redirection. If the redirection could lead to the above risks, alert the operator explicitly.

Solution 2: Considerate routing

Traffic redirection may generate unexpected traffic models, causing some SLAs to be unguaranteed. A potential solution is to extend BGP to allow redirection rules and the next hop to be advertised to other peers for routing decision.

Solution 3: Network visualization

Inter-domain redirection also hinders network visualization, making it impossible to determine the forwarding path of packets solely based on routing protocols. A potential solution is to extend the BMP protocol to allow redirection rules and redirected forwarding paths to be sent to the BMP server, improving administrators' ability to visualize and plan their own traffic.

4. IANA Considerations

This document makes no request of IANA.

5. Security Considerations

TBD

6. Acknowledgements

The authors would like to thank Hao Li.

7. Normative References

[RFC1102] Clark, D., "Policy routing in Internet protocols", RFC 1102, DOI 10.17487/RFC1102, May 1989, <https://www.rfc-editor.org/info/rfc1102>.

[RFC1104] Braun, H., "Models of policy based routing", RFC 1104, DOI 10.17487/RFC1104, June 1989, <https://www.rfc-editor.org/info/rfc1104>.

[RFC7908] Sriram, K., Montgomery, D., McPherson, D., Osterweil, E., and B. Dickson, "Problem Definition and Classification of BGP Route Leaks", RFC 7908, DOI 10.17487/RFC7908, June 2016, <https://www.rfc-editor.org/info/rfc7908>.

[RFC8955] Loibl, C., Hares, S., Raszuk, R., McPherson, D., and M. Bacher, "Dissemination of Flow Specification Rules", RFC 8955, DOI 10.17487/RFC8955, December 2020, <https://www.rfc-editor.org/info/rfc8955>.

[RFC9117] Uttaro, J., Alcaide, J., Filsfils, C., Smith, D., and P. Mohapatra, "Revised Validation Procedure for BGP Flow Specifications", RFC 9117, DOI 10.17487/RFC9117, August 2021, <https://www.rfc-editor.org/info/rfc9117>.

Authors' Addresses

Weiqiang Cheng (editor)
China Mobile
China

Email: chengweiqiang@chinamobile.com

Dan Li
Tsinghua University
China

Email: tolidan@tsinghua.edu.cn

Yasi Liu
China Mobile
China

Email: liuyasi@chinamobile.com

Mingqing Huang
Huawei
China

Email: huangmingqing@huawei.com

Fang Gao
Zhongguancun Laboratory
China

Email: gaofang@zgclab.edu.cn

Shuanglong Chen
Huawei

China

Email: chenshuanglong@huawei.com