BGP Flow Specification Deployment Experience

Derek Gassen, Raul Lozano
Time Warner Telecom

Danny McPherson, Craig Labovitz
Arbor Networks

Agenda

- Flow Spec Overview
- About TWTC
- DDOS problem and Observations
- Previous Mitigation Approaches
- Experiences with Flow Spec
- Futures

Flow Spec Overview

The Problem: Providers have limited options for mitigating DDoS attacks internally and inter-provider, e.g:

- – BGP destination black-holes: completes the attack
- – BGP src/uRP: difficult for some spoofed attacks and/or support large numbers of sources
- – ACLS: difficult to maintain and occasionally dangerous to install

Flow Spec

Basic idea: Use BGP to distribute flow specification filters and dynamically filter on routers.

Flow Spec

- Encode flow specification rules as new BGP NLRI address family.
- BGP itself treats the FlowSpec NLRI as an opaque key to an entry in its database.
- Use extended communities to specify action (accept, discard, rate-limit, sample, redirect).
- Match on combination of source/dest prefix, source/dest port, ICMP type/code, packet size, DSCP, TCP flag, fragment encoding, etc.
  - Example: all TCP port 80..90 packets to 192.168.0/24.

Flow Spec Trust Model

- Unicast routing advertisements control traffic gets forwarded.
- Consider a filter as a “hole” in the aggregate of traffic that is being forwarded to a destination prefix.
- Accept filter when advertised by next-hop for the destination prefix.

Flow Spec Benefits

- Fine grain specification of filters with BGP’s ease of deployment/management
  - BGP solves distribution and trust problem
- Leverage ASIC filtering in routers.
- Available today

Flow Spec Limitations

- BGP lacks update level security
- No well defined application level acknowledgement nor statistics
- Only works for BGP enabled nodes
- BGP payload needs to be overloaded beyond “routing” (i.e. IDS signature update, matching of traffic beyond IP header)
- Possible/perceived operational issues between Network (Routing) Operations (NetOps) and Security Operations (SecOps)
- No centralized gathering of threat information

Flow Spec Status

IETF draft available at:

– http://www.tcb.net/draft-marques-idr-flow-spec-03.txt

- Implemented as of JunOS 7.2 (but not documented)
- At least three tier1/2 providers in process of production deployment
- Several security vendors announced intregration
- Cisco complimentary TIDP proposal

Time Warner Telecom (TWTC)

About Time Warner Telecom Inc.

- Based in Littleton, CO. 2,105 employees (June 30, 2006)
- Became separate entity from Time Warner, Inc in 1998
- National IP, Transport, and Switched Services Provider
  - 44 markets in 23 states
  - Metro Ethernet presence to 6,400+ buildings
  - Multiple business IP services & managed services
  - Managed Security & VoIP Services
  - All Juniper Backbone
  - Cisco Powered Metro

TWTC Backbone

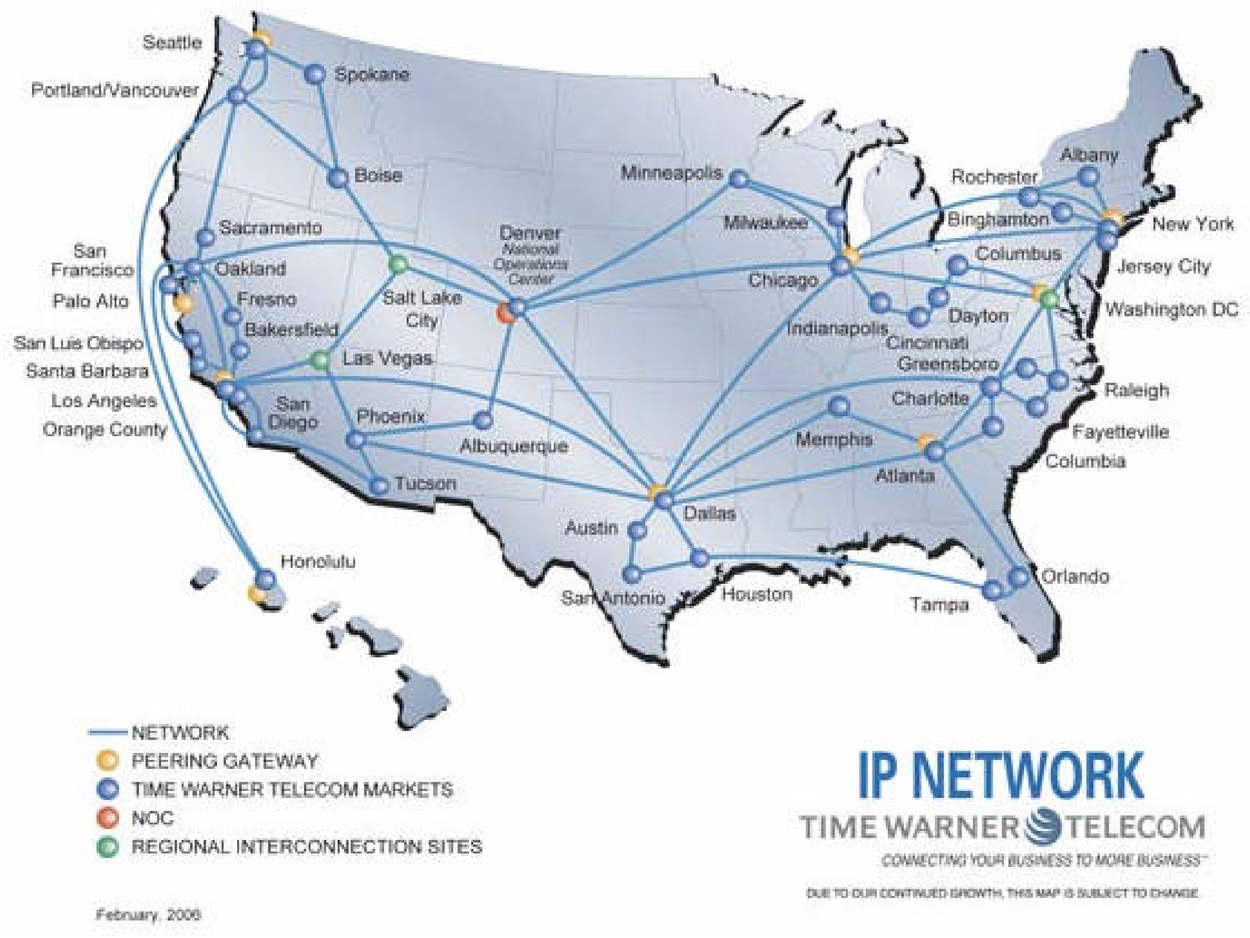

DDoS Problem

- Large customers often experience attacks
- Previously ad-hoc response/mitigation
  - Customer calls NOC for installation of filters/ACLs
  - Recently began letting customers announce own blackholes
- Massive DDoS attacks no longer just affect the end customer; infrastructure is at risk.
  - Congestion on backbone and at exchange
  - Transit circuit congestion and added burst cost if attack is a long duration
  - POP isolation depending on the size of attack
  - VoIP
- Economic
  - Increases the cost of operations – detrimental to the business

DDoS Observations

- Large DDoS Attacks
  - Organized crime using bot-net armies for extortion
  - Script kiddies launching attacks
- Commonly TCP SYN flood attacks are relatively small in actual bits per second and do their damage with actual packets per second. We have experienced SYN attacks greater than 1Mpps on our backbone due to large bot-nets
  - Most recent attack was ~ 2Mpps
- UDP fragment attacks by far most damaging by traffic load because a fewer number of PPS are required due to packet size
  - TWTC has seen attacks 2gb/s+ in size on a more frequent basis

DDoS Observations Cont.

- ICMP attacks are mitigated on our edge with policing filters and seem to be the least common attack method on our network
- Significant percentage of major attacks originating from APNIC. Attacks from Chinese IP space difficult to track due to national NAT gateways

Example of a DDoS Attack

Recent SYN attack to a customer

26,495.5% of 10 Kpps ~1.5gb/s

Maximum sustained > 5MPPS

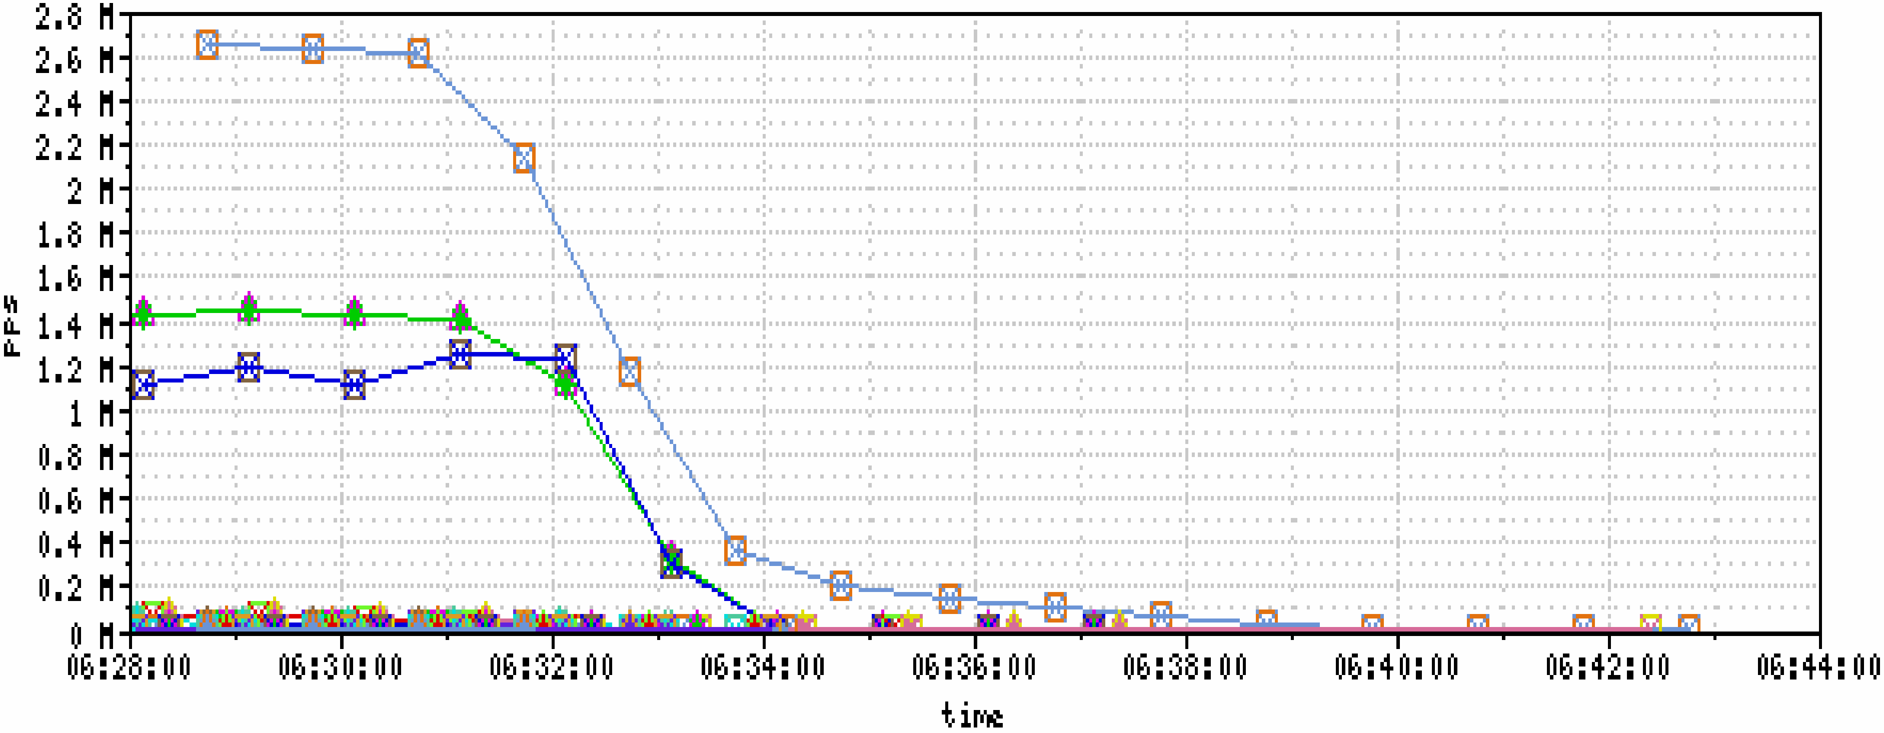

Previous Mitigation Approaches

1) Traditional destination based blackhole methodology.   (0/0 next-hop discard)
   - Withdraw specific customer prefix if the attack is extremely large. Restricted to BGP customers and infrastructure

2) Arbor Fingerprint sharing with upstream providers and peers

Mitigation Challenges

Destination BGP blackholing & firewall filtering is insufficient:

- Slow to generate since it requires login to the devices and configuration changes
- Terminates the traffic to a destination affecting availability
- Constant configuration changes to add and remove blackhole routes

Mitigation Challenges (cont)

Effort (NOC man hours, second-level engineering, etc) to handle mitigation:

- It can take from 15 minutes to approximately an hour to identify and mitigate
- Depends on the experience of the NOC personnel

Mitigation Today

Deploying flow-spec:

- JunOS 7.4R3
- Arbor 3.5
- Flowspec on Peers & Transits
- BGP community architecture (Granularity)
- Security Operation Center initiates FLOW SPEC mitigation using Arbor
- Uses BGP propagation to start mitigating the attack

Flow Spec Deployment

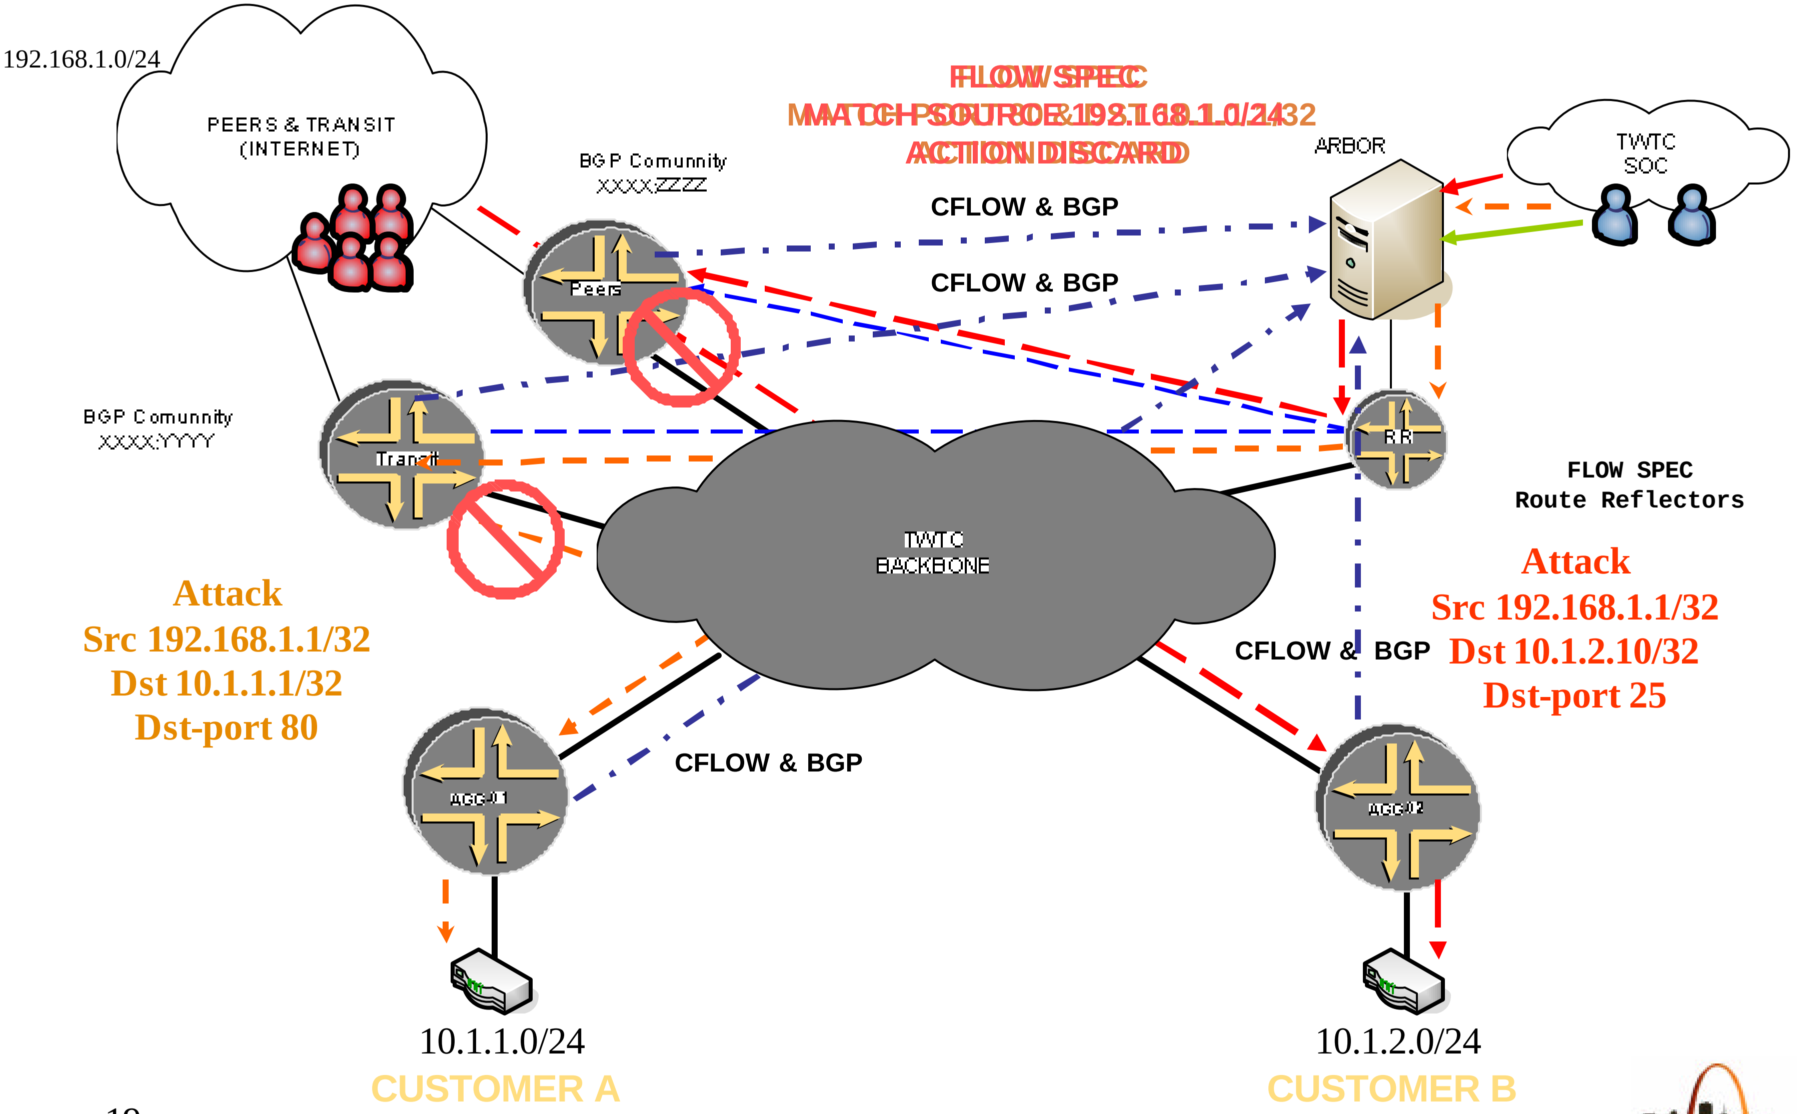

Attack
Src 192.168.1.1/32
Dst 10.1.1.1/32
Dst-port 80

Attack
Src 192.168.1.1/32
Dst 10.1.2.10/32
Dst-port 25

Experience with FlowSpec

- Early versions had bugs
  - Multiple dst/src will mitigate all ports
  - Port mapping from FLOWSPEC to the CFchip was incorrectly installed

- Performance and other limitations
  - Inability to count on discarded traffic
  - Arbor support for all FLOWSPEC actions

Experience with FlowSpec (CONT)

TWTC planning/testing:

- LAB testing using traffic generators creating flows with different profiles.
- Mitigating them with normal FLOW SPEC propagation
- Using tools such as Arbor to identify anomalies in the network and mitigate them using FLOW SPEC.
- FLOWSPEC works
- Controlled deployment
- Currently mitigating external ingress point (peers and transits)

Experience with Flow Spec (CONT)

Lab testing (Other use cases):

- Verified routing-instance for cleaning station
- Verified rate-limiting

Other Applications (LAB)

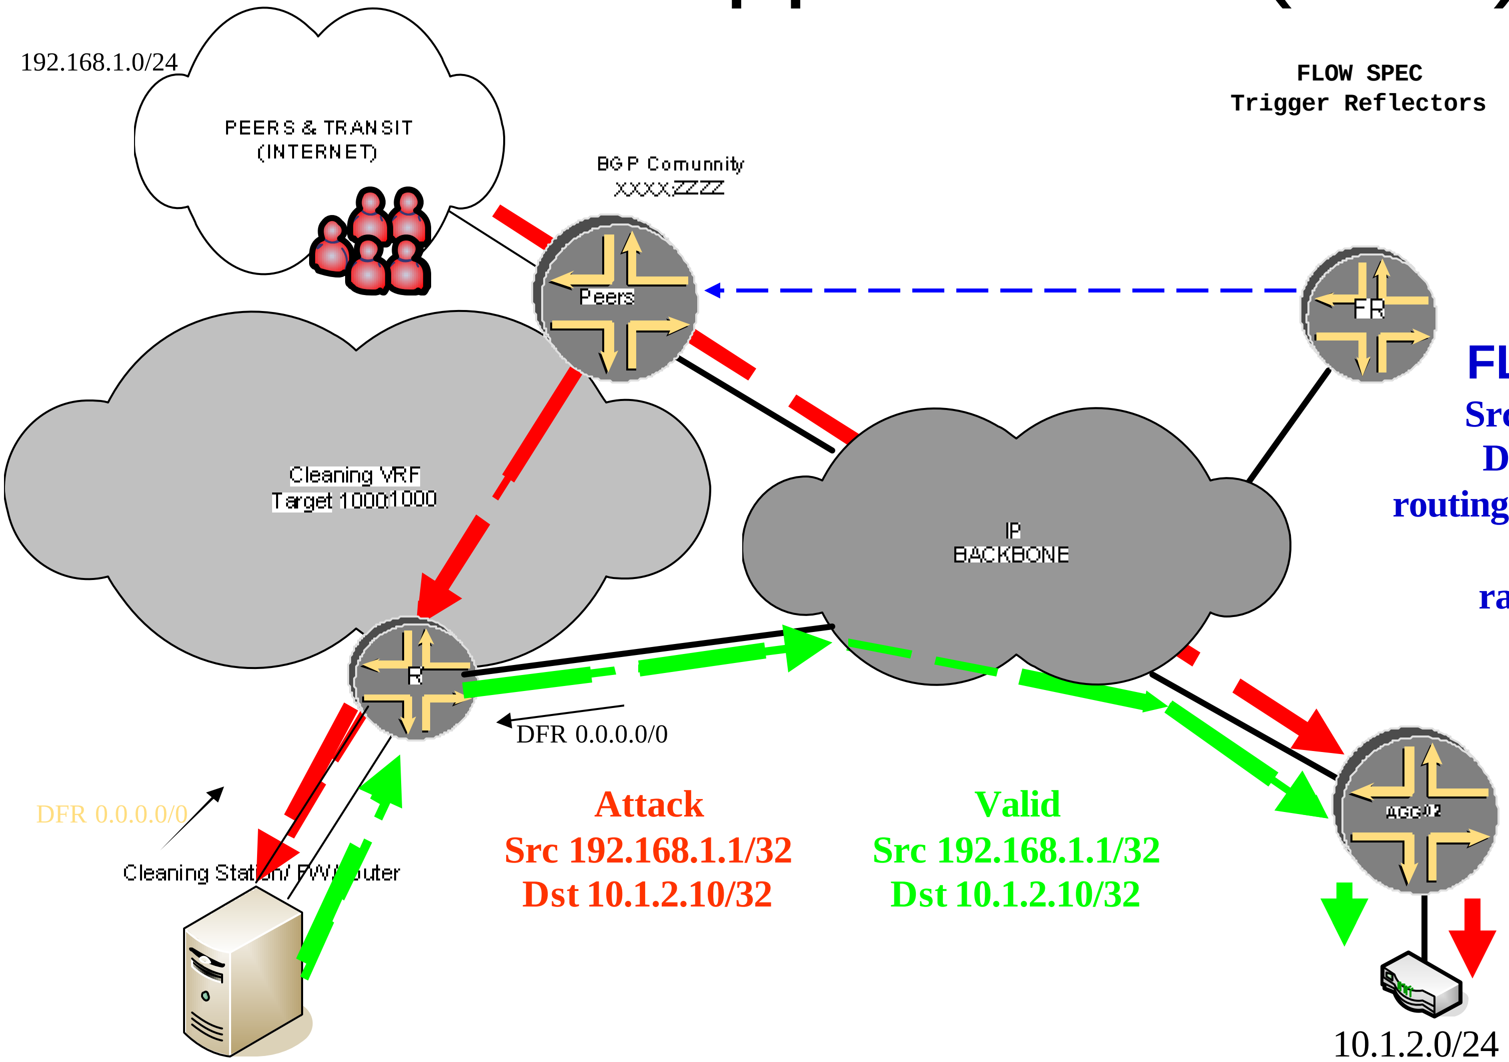

FLOW SPEC

Src 192.168.1.1/32

Dst 10.1.2.10/32

routing-instance 1000:1000

rate-limit 100kbps

Attack

Src 192.168.1.1/32

Dst 10.1.2.10/32

Valid

Src 192.168.1.1/32

Dst 10.1.2.10/32

10.1.2.0/24

CUSTOMER B

What is Missing?

- Multi-vendor support.
- Juniper
  - Support for rich routing techniques
    - Change destination address
    - Support for prefix lists/policy statements
    - Reporting discarded flows to cflowd
- Arbor
  - Better reporting on attacks mitigated by FLOWSPEC
  - Support for more actions such as routing instanceAutomatic mitigation
  - Automatic FLOWSPEC mitigation for well known threats signatures

What is Missing? (cont)

- Inter-carrier support FLOWSPEC
  - Mitigate the attack at the source.
  - Eliminate collateral damage for both carriers.

- Support for changing matching criteria such as DSCP code.

- Support of FLOWSPEC in policy actions

Conclusion

- Recommendations to other providers thinking of using FlOWSPEC:
    - Consider single vendor support
    - Controlled deployment
    - BGP tier design
    - Analysis of traffic
- Low-Cost by deploying flow spec on existing infrastructure
- Distributed attacks (worms) can be mitigated with a faster response.
- Granularity is powerful

THANK YOU