Dissemination of Flow Specification rules

Jean-Marc Uzé

TF-NGN, Zürich, 15/04/05

Capabilities of routers

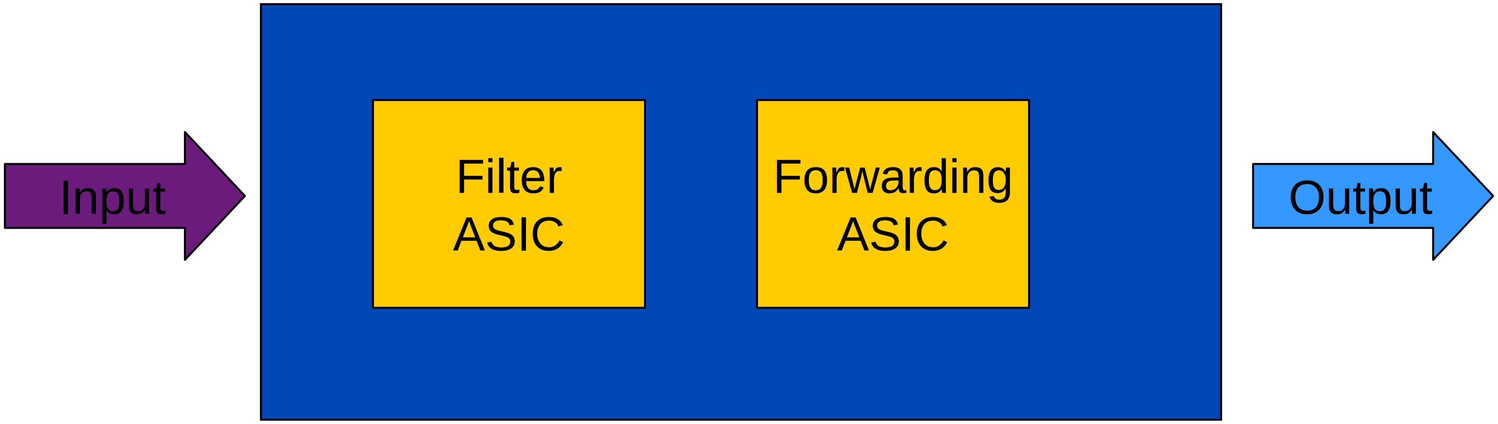

- Forwarding ASIC: maps destination prefix to next-hop.
- Traffic filtering ASIC: packet header pattern to action (drop, rate-limit, etc).

Filtering engine

- Manually configured with static entries.
- Is that a problem ?
- Why is the forwarding ASIC not programmed in the same way with static routes ?
  - i.e. to advertise new customer prefix we could pick up the phone and call peer networks.
  - Really safe.

Dynamic firewall filtering

- Routing is dynamic and inter-domain
- Filtering is static and intra-system

- Need to coordinate Flow filtering for intra and inter-domain
- More generally there is a need to Extend routing information with Flow Specification

Distributed DoS Attacks

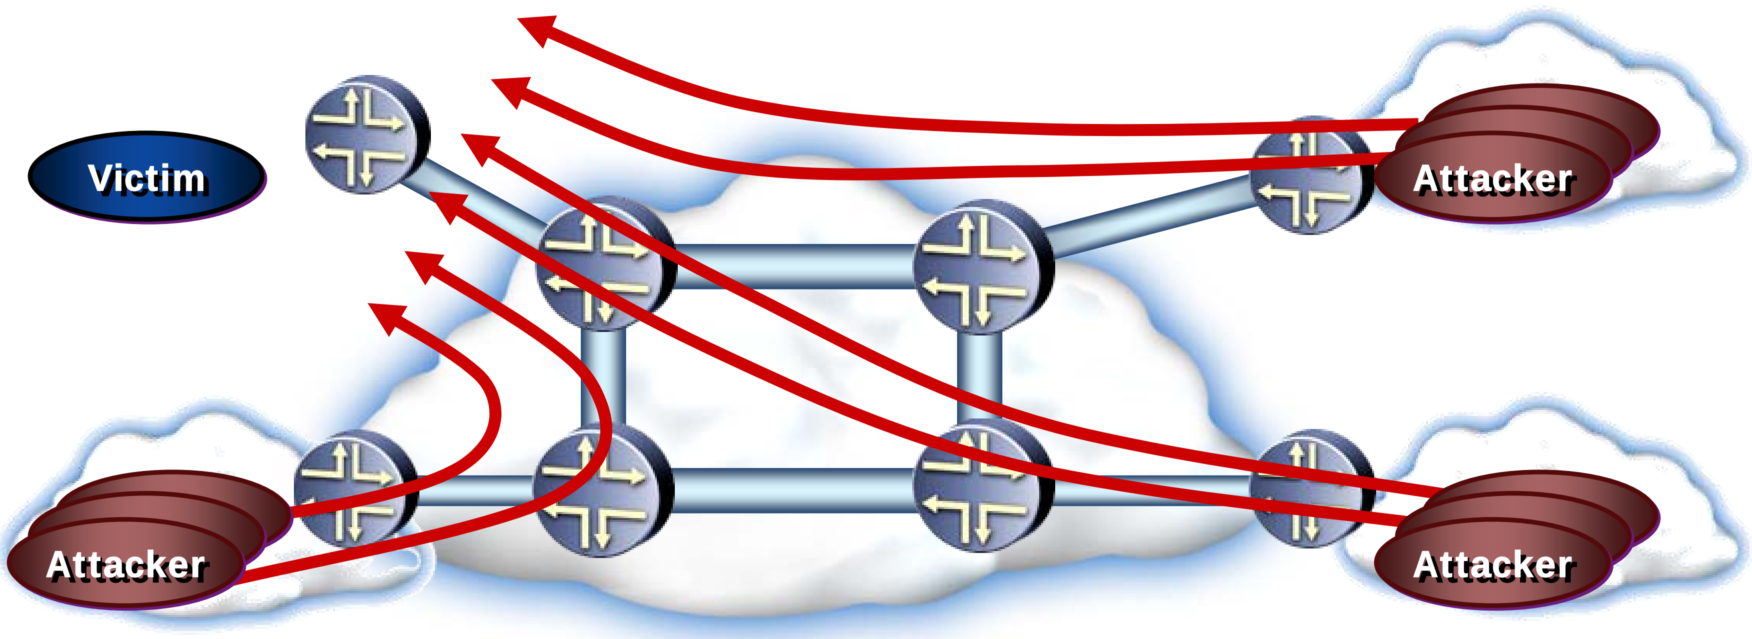

- Attacker compromises hosts in multiple networks, using them to launch a coordinated attack
- Attack can't simply be stopped at one point
- Achieves bandwidth leverage from multiple sources

Users @ 256k: 100 = 25Meg, 500 = 128Meg

DDoS Attack Architecture

Users @ 256k:

100 = 25Meg

500 = 128Meg

2,000 = .5Gig

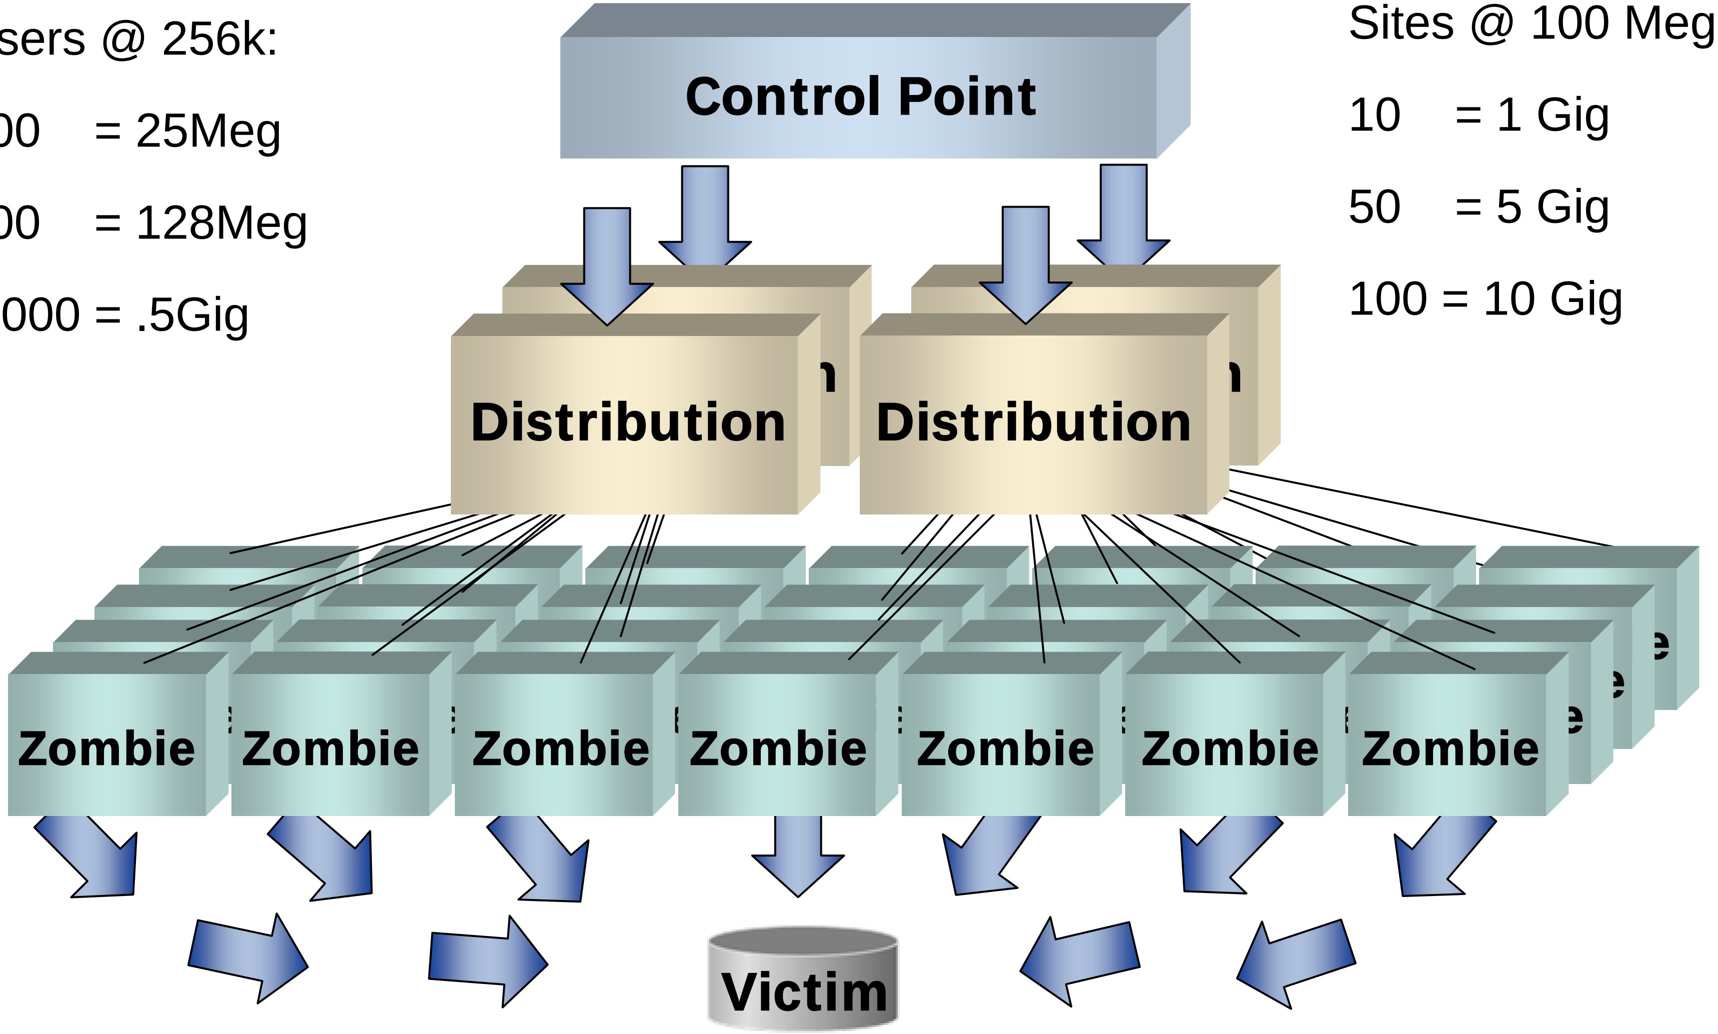

Sites @ 100 Meg:

10 = 1 Gig

50 = 5 Gig

100 = 10 Gig

Need to coordinate traffic filtering

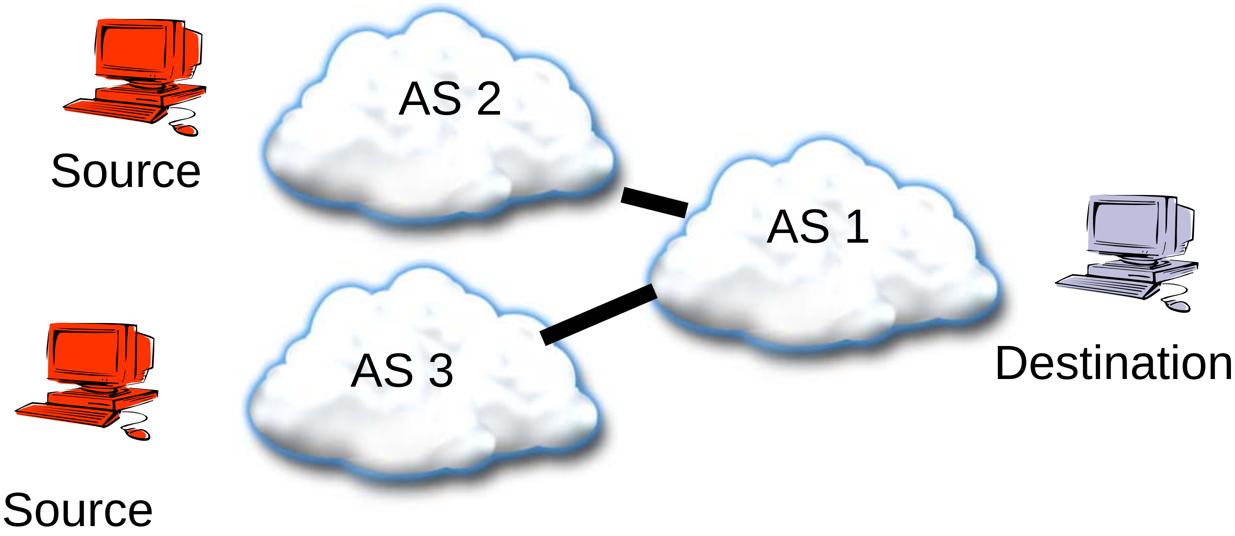

- Filter close to the source(s).
  - Traffic rates may be too large for AS 1 to handle without impact.
  - E.g. AS 2 to AS 1 interconnect can get congested.

Main current approaches

- Advertise /32s specific route
  1. with black-hole community in BGP
     - marks such route advertisements with a community that gets translated into a discard next-hop by the receiving router
  2. that attracts traffic to a particular node that serves as a deterministic drop point
  3. with a BGP community linked to a feature counting packet/bytes destinated to it (e.g. Destination Class Usage/DCU), on customer-facing ingress interfaces counters retrievable via SNMP => DOS identification

Real-time DoS Identification with Destination Class Usage

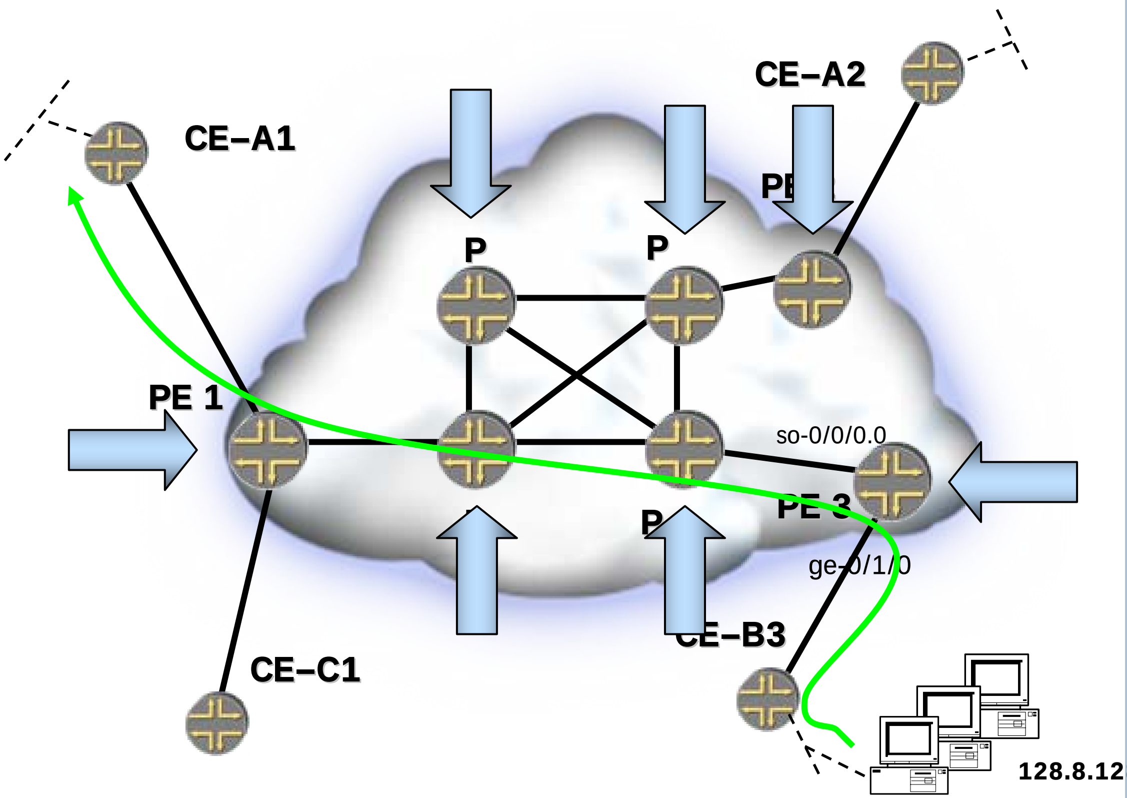


```
policy-options {
 community victim members 100:100;
 policy-statement set-dest-class
  term 1 {
   from {
    protocol bgp;
    community victim;
   }
   then {
    destination-class dcu-victim;
    accept;
   }
  }
 }
}

interfaces {
 so-2/0/1 {
  unit 0 {
   family inet {
    address 192.168.4.1/32;
    accounting {
     destination-class-usage;
    }
   }
  }
 }
}

routing-options{
 forwarding-table{
  export set-dest-class;
 }
}
```

Real-time DoS Identification with Destination Class Usage

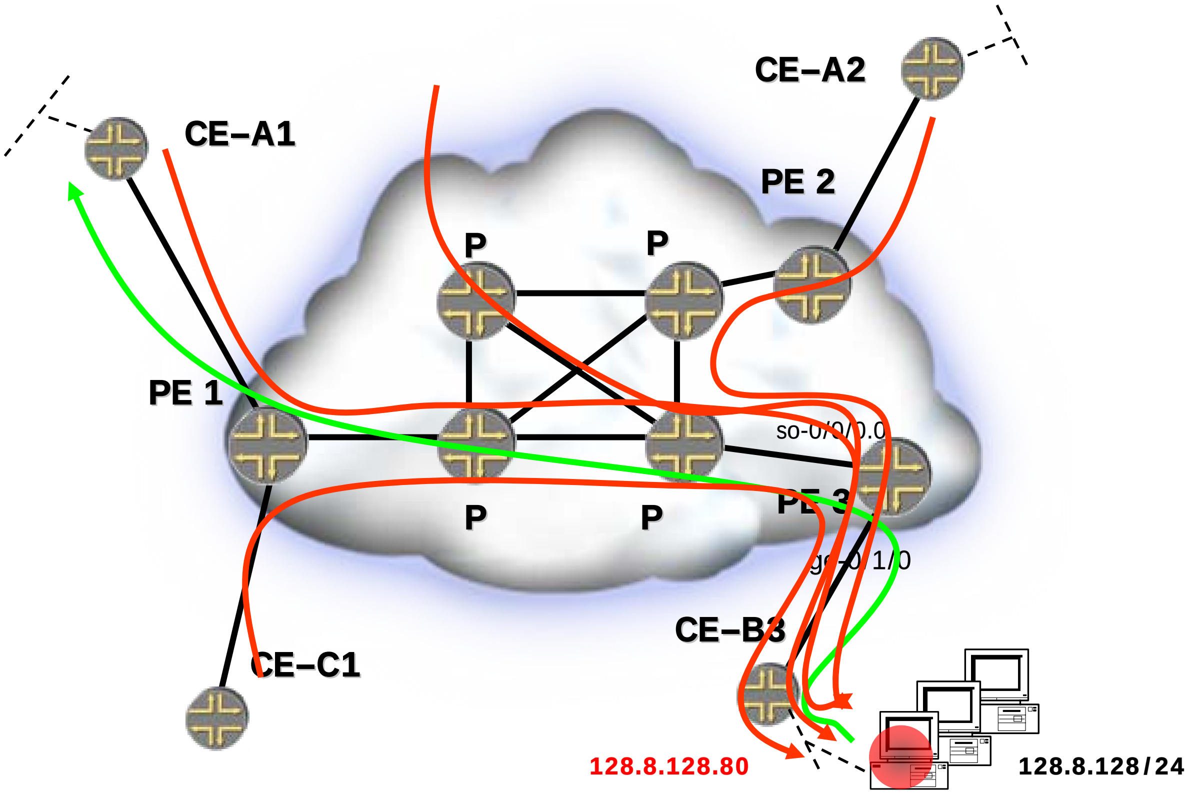

Real-time DoS Identification with Destination Class Usage

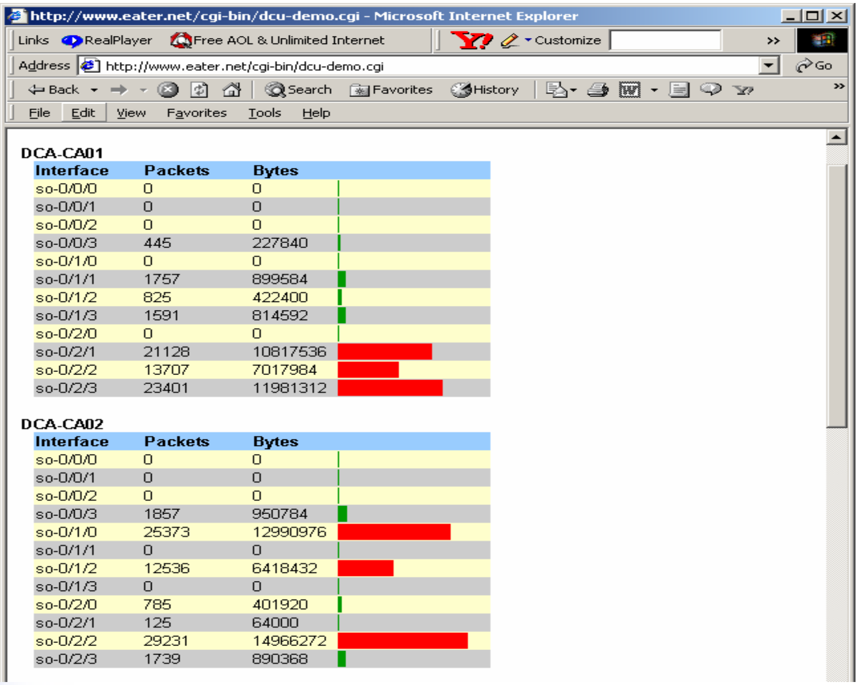

DCA-CA01

Interface	Packets	Bytes
so-0/0/0	0	0
so-0/0/1	0	0
so-0/0/2	0	0
so-0/0/3	445	227840
so-0/1/0	0	0
so-0/1/1	1757	899584
so-0/1/2	825	422400
so-0/1/3	1591	814592
so-0/2/0	0	0
so-0/2/1	21128	10817536
so-0/2/2	13707	7017984
so-0/2/3	23401	11981312

DCA-CA02

Interface	Packets	Bytes
so-0/0/0	0	0
so-0/0/1	0	0
so-0/0/2	0	0
so-0/0/3	1857	950784
so-0/1/0	25373	12990976
so-0/1/1	0	0
so-0/1/2	12536	6418432
so-0/1/3	0	0
so-0/2/0	785	401920
so-0/2/1	125	64000
so-0/2/2	29231	14966272
so-0/2/3	1739	890368

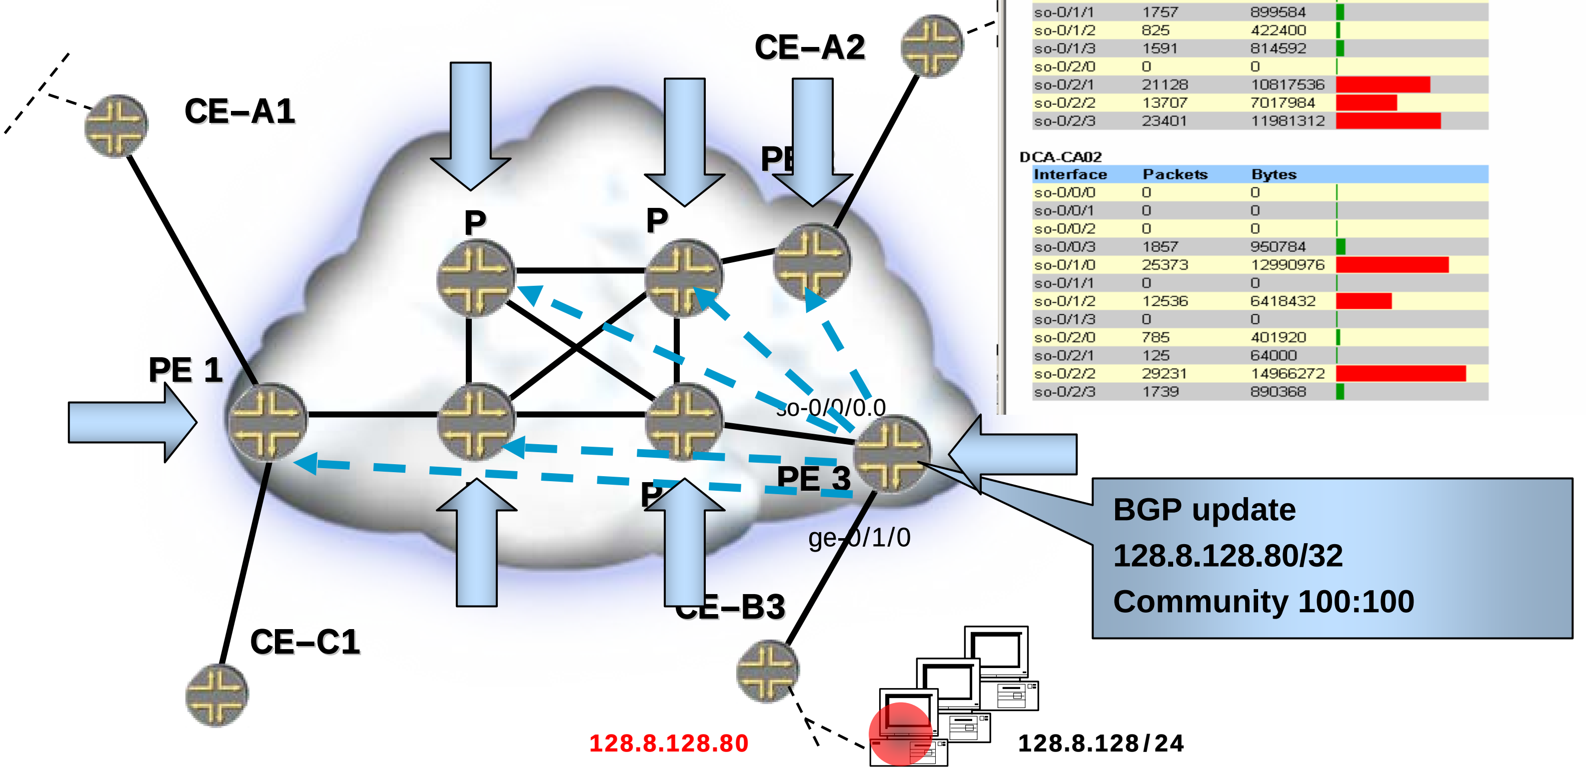

Common Problems with Current approaches

- We just advertise /32s in BGP unicast routing…
- Problems:
  - Need to open up policy filters to allow more specifics.
  - Mixed up with unicast routing.
  - Traffic filtering engines can deal with more granularity.

Solution: Dissemination of flow specification rules with BGP (1)

- Allow BGP to propagate an n-tuple
  - matching could be a combination of source/dest prefix, source/dest port, ICMP type/code, packet size, DSCP, TCP flag, fragment encoding, etc…, E.g.:
    - all packets to 10.0.1/24 and TCP port 25
    - all packets to 10.0.1/24 from 192/8 and port {range [137, 139] or 8080 (NLRI length of 16 bytes)

Solution: Dissemination of flow specification rules with BGP (2)

- Information is kept independently of unicast routing.
- But it is automatically validated against unicast routing.
- Filtering actions could be a combination of accept, discard, rate-limit, sample, redirect, etc...
  - Accomplished by mapping a user defined community value to platform network specific behavior via user configuration.

Juniper your Net

Trust model

- Unicast routing advertisements control where 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.

Validation

- Compare destination address of traffic filtering rule with best match unicast route for this prefix.
    1. Originator of filter and unicast route must be same.
    2. No more specifics from a different AS.

Flow spec Dissemination

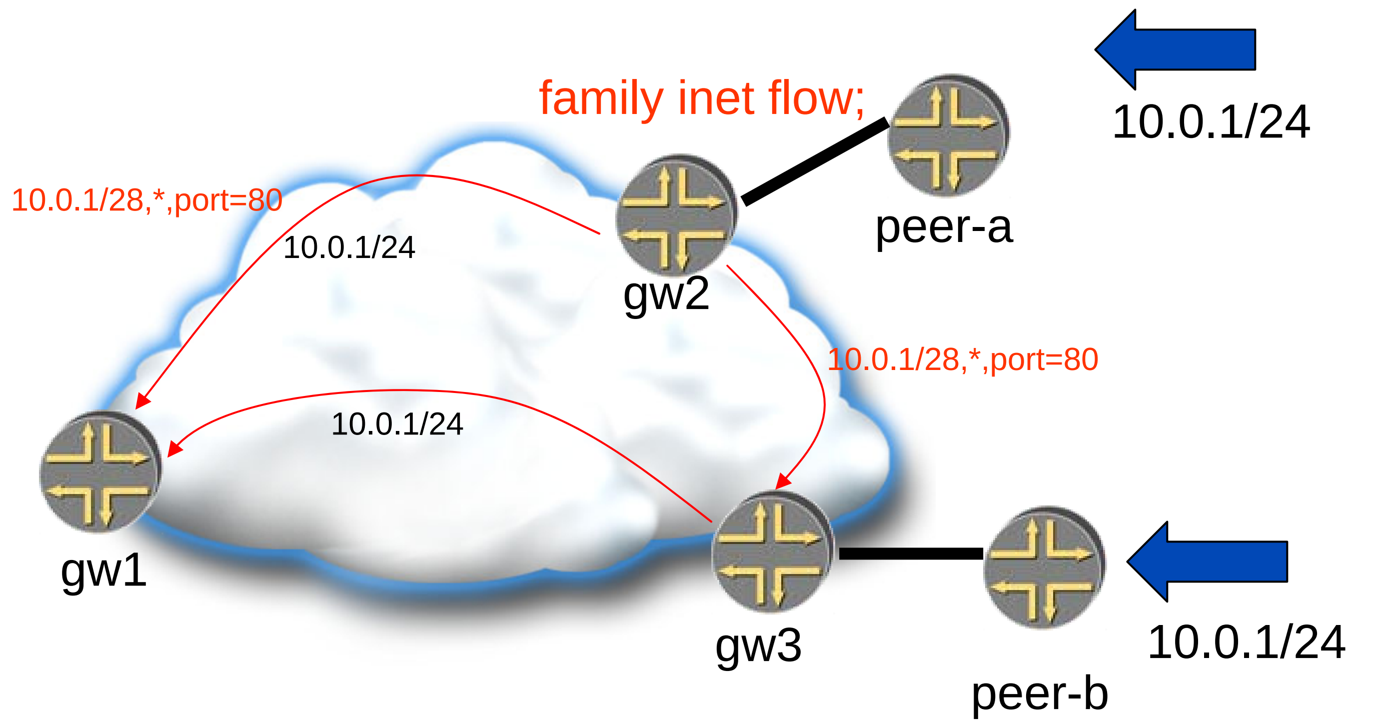

- gw2 accepts traffic filtering rule from a.
- gw3 prefers unicast via b; rejects filter.
- gw1 may (or not) accept the filter based on the "originator" of the BGP route.

Flow spec Dissemination

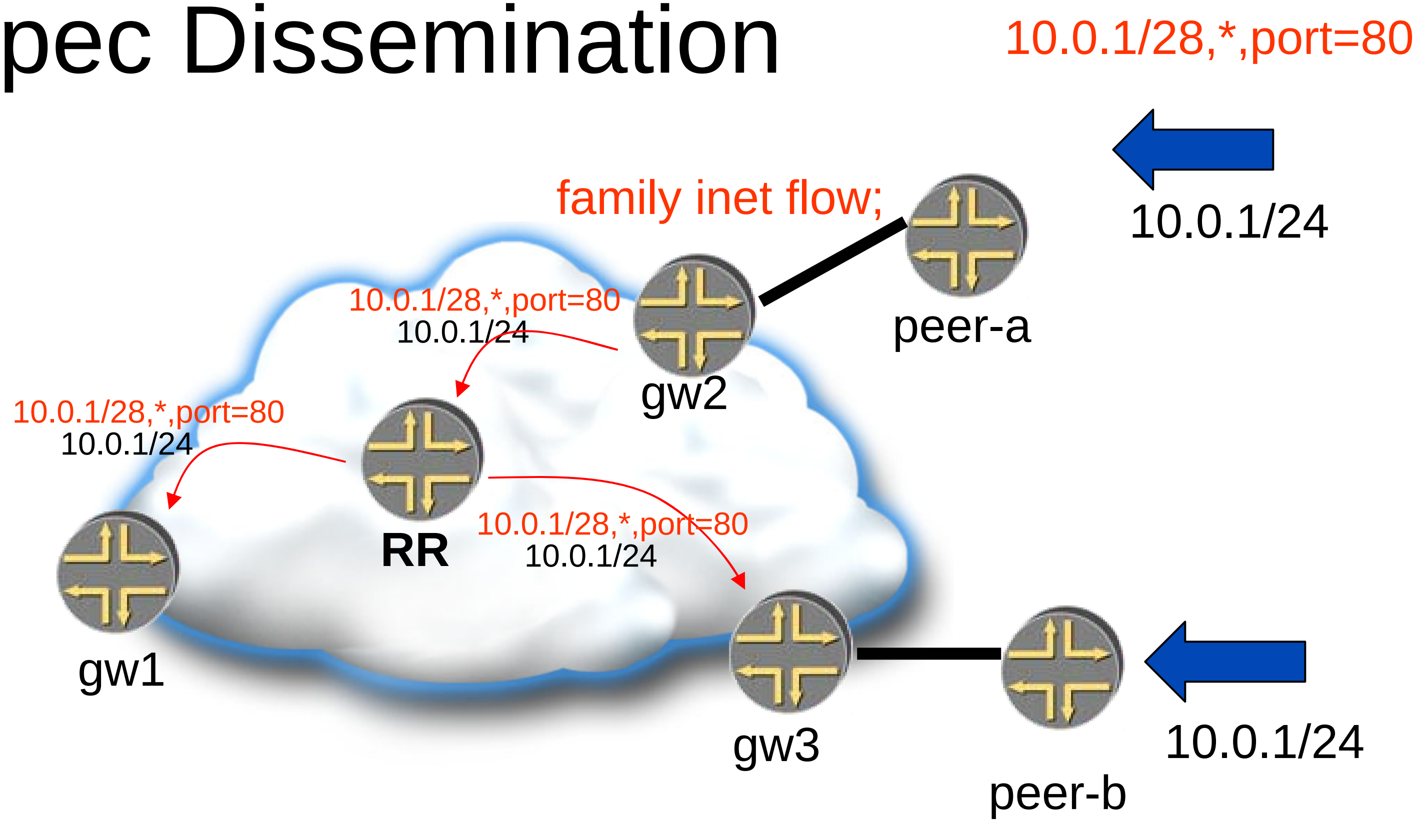

- With a Route Reflector:
  - The RR may (or not) accept the filter based on the “originator” of the BGP route.
  - gw3 prefers unicast via b; rejects filter.
  - gw1 will follow the RR decision (accept or not the filter)

Why BGP ? (1)

- BGP offers the substantial advantage of being an incremental addition to deployed mechanisms.
- The key issues in terms of complexity are problems which are common to unicast route distribution and have already been solved in the current environment
    - From an algorithmic perspective, the main problem that presents itself is the distributed loop-free distribution of <key, attribute> pairs. The key, in this particular instance, being a flow specification.

Why BGP ? (2)

- Allows a network operator to reuse:
  - internal route distribution infrastructure (e.g.: route reflector or confederation design)
  - existing external relationships (e.g.: inter-domain BGP sessions to a customer network)
- Proven Scalability and Flexibility of BGP in adding new services
  - Multicast, IPv6, L3 VPN, L2 VPN, VPLS

Summary

- Extension to routing information.
- Different traffic filters accepted in different parts of the network according to different unicast routing decisions.
- Inter-domain solution to coordinate traffic filtering
  - Open to several potential applications
- IETF draft: draft-marques-idr-flow-spec-02
- Mailing list:

  http://www.cqr.org/mailman/listinfo/flow-spec

Thank You

juze@juniper.net