Traffic Diversion Techniques for DDoS Mitigation using BGP Flowspec

Leonardo Serodio

leonardo.serodio@alcatel-lucent.com

May 2013

Alcatel·Lucent

Distributed Denial of Service (DDoS) Attacks

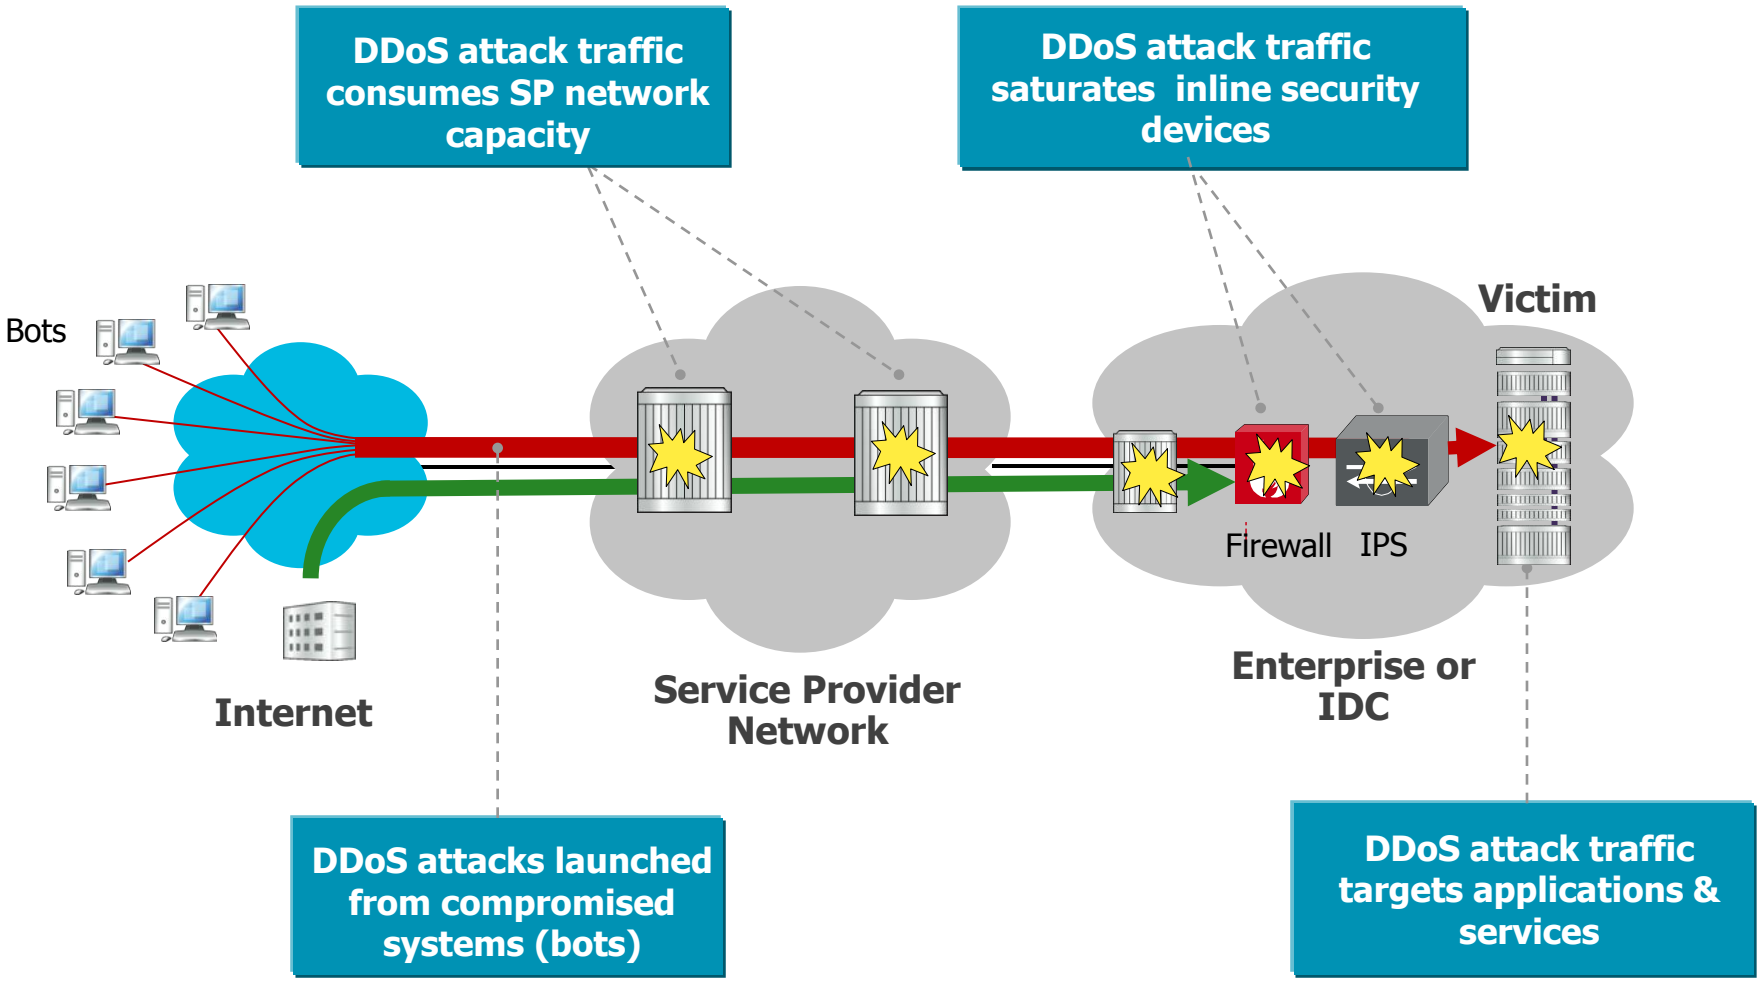

Cloud-based DDoS Defense with RTBH

BGP Announcement Originated by the Victim

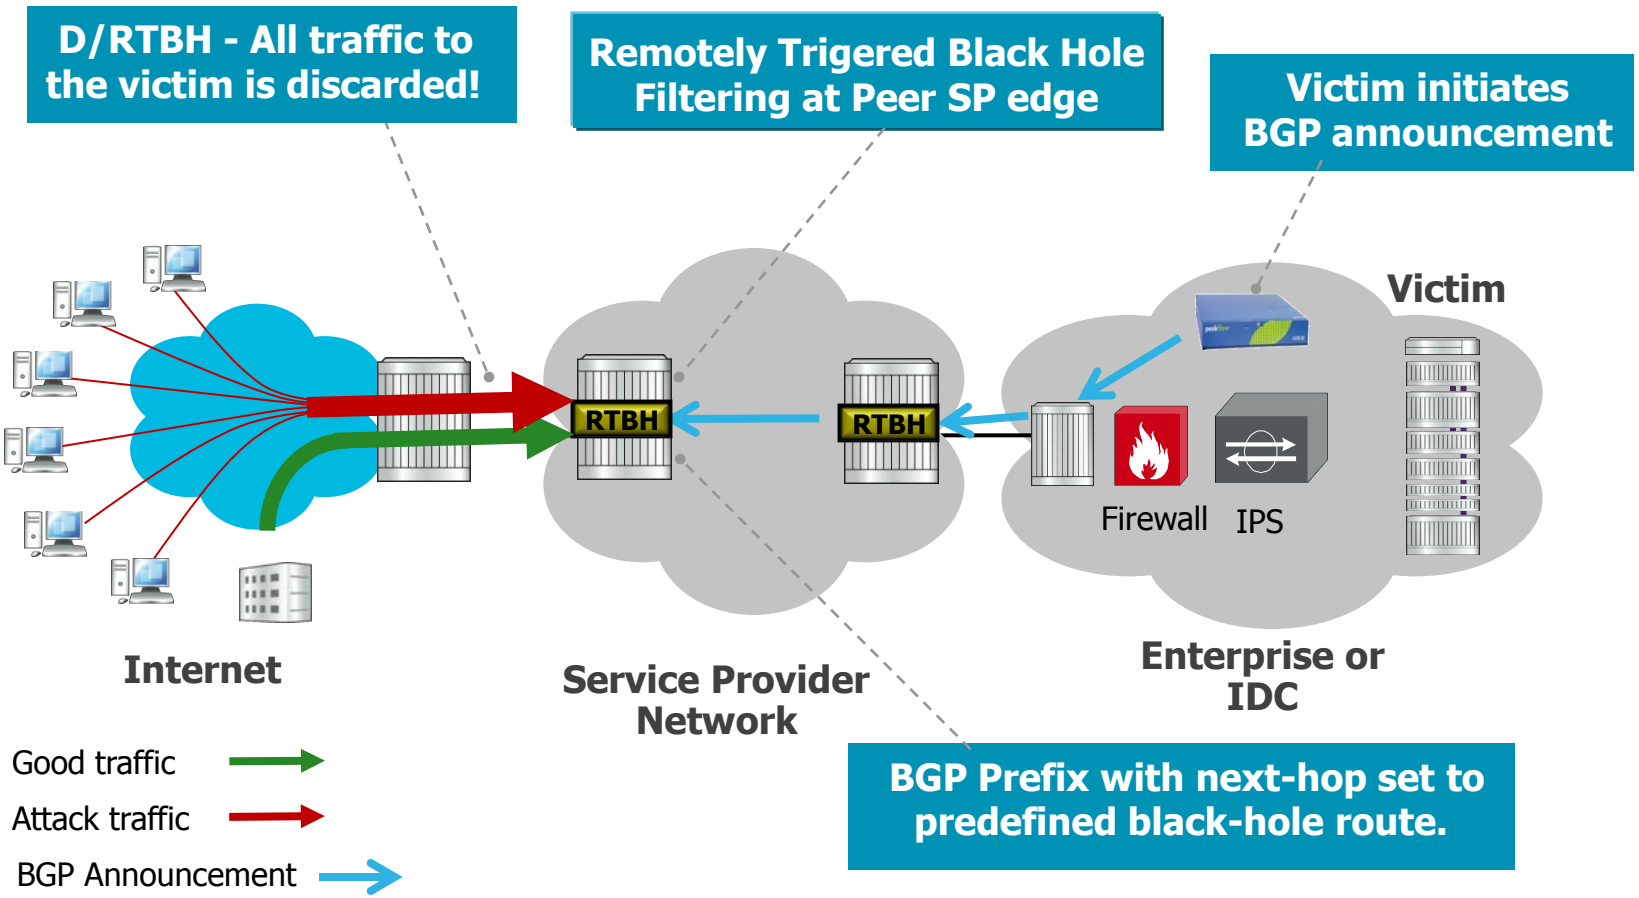

- Fourth most used tool to mitigate DDoS attacks [5]
- RFCs: RFC 3882, RFC 5635 (includes D/RTBH and S/RTBH)

Cloud-based DDoS Defense with RTBH

BGP Announcement Originated by the SP

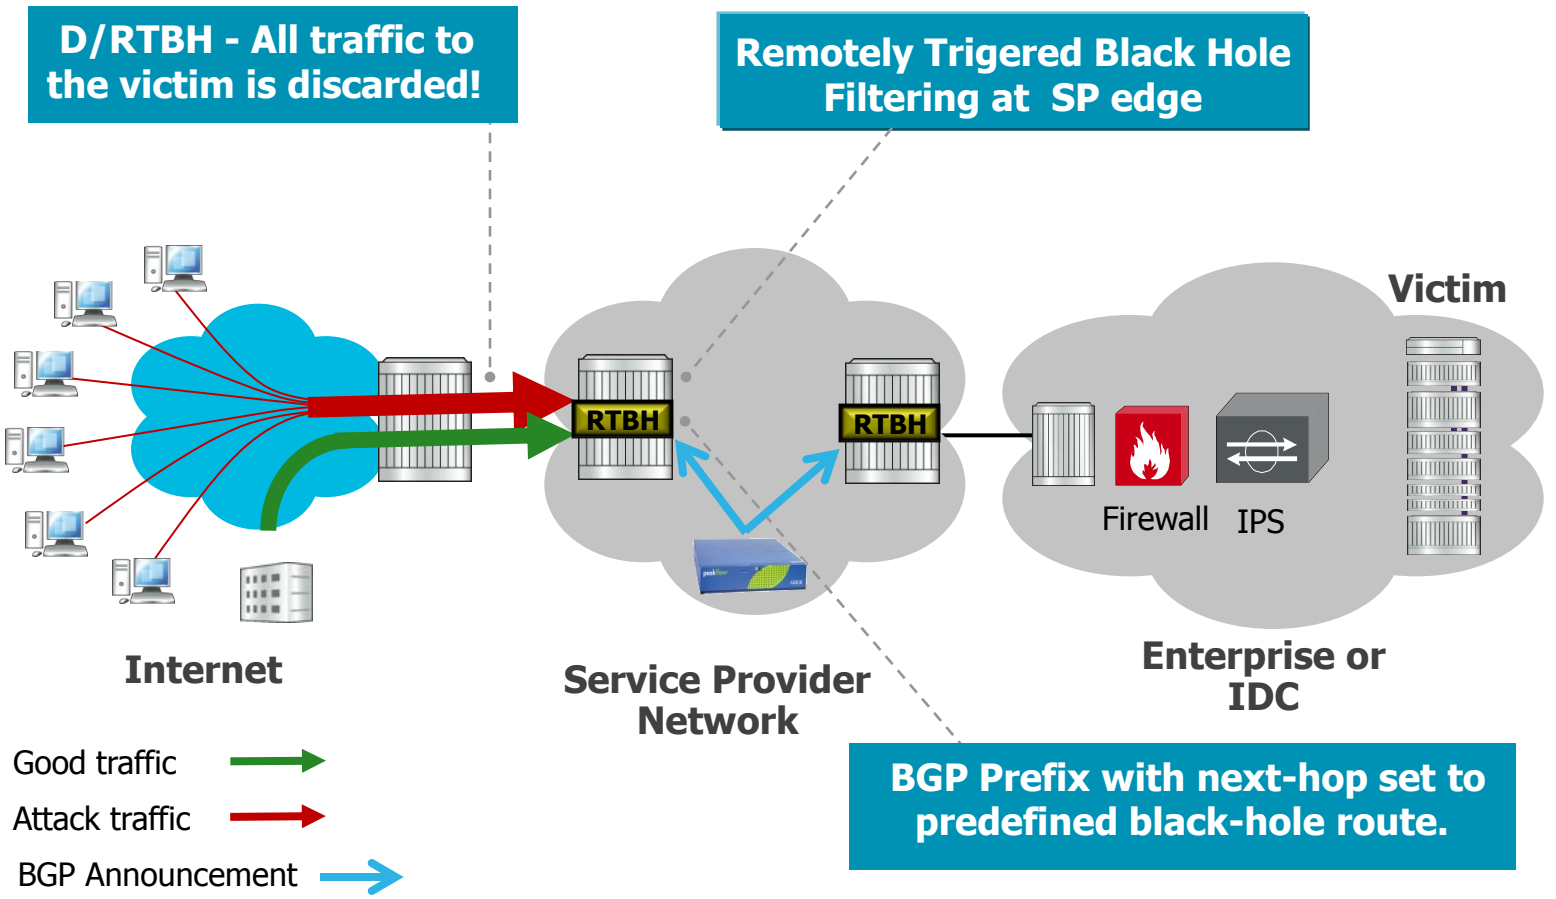

BGP Flow Specification

- BGP Flowspec defines a new BGP Network Layer Reachability Information (NLRI) format used to distribute traffic flow specification rules.
  - NLRI (AFI=1, SAFI=133): IPv4 unicast filtering
  - NLRI (AFI=1, SAFI=134): BGP/MPLS VPN filtering
- Specified by RFC 5575 [1], extended to IPv6 in [2]
- Main application today is to automate the distribution of traffic filter lists to routers for the mitigation of DDoS attacks.
  - Selectively drop traffic flows based on L3/L4 information.
  - Intelligent control platform builds filter rules to drop harmful traffic, encodes them as BGP flowspec routes and advertises them to BGP peers.

BGP Flow Specification

- The Flow specification can match on the following criteria:
  - Source / Destination Prefix
  - IP Protocol (UDP, TCP, ICMP, etc.)
  - Source and/or Destination Port
  - ICMP Type and Code
  - TCP Flags
  - Packet Length
  - DSCP (Diffserv Code Point)
  - Fragment (DF, IsF, FF, LF)
- Actions defined using Extended Communities:
  - 0x8006: traffic-rate (rate 0 discards all traffic for the flow)
  - 0x8007: traffic-action (sample)
  - 0x8008: redirect to VRF
  - 0x8009: traffic-marking (DSCP value)

BGP Flow Specification – Why Use It for ACLs?

- ACLs are still the single most widely used tool to mitigate DDoS attacks [5].
  - But…ACLs are demanding in configuration & maintenance.
- BGP Flowspec leverages the BGP Control Plane to simplify the distribution of ACLs, greatly improving operations:
  - Inject new filter rules to all routers simultaneously without changing router configuration.
  - Reuse existing BGP operational knowledge and best practices.
  - Control policy propagation via BGP Communities.
- Improve response time to mitigate DDoS attacks.
- Route validation performed for eBGP sessions, see draft [3] for revised validation procedure for iBGP sessions.

Cloud-based DDoS Defense with BGP Flowspec

Inter-domain flowspec injection

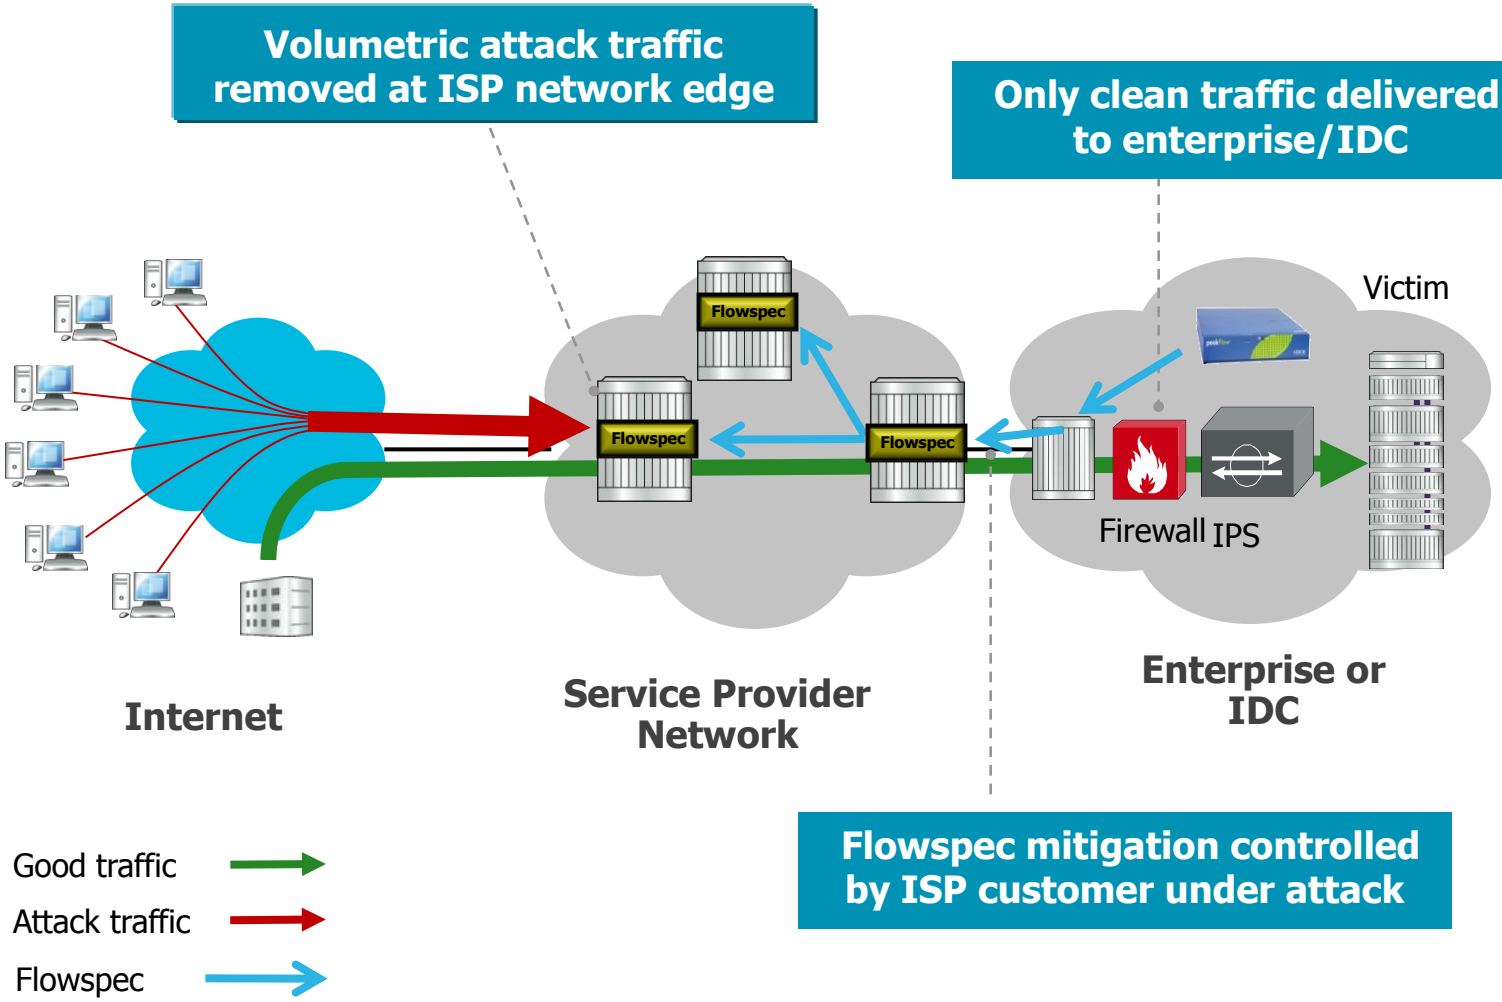

Cloud-based DDoS Defense with BGP Flowspec

Intra-domain flowspec injection

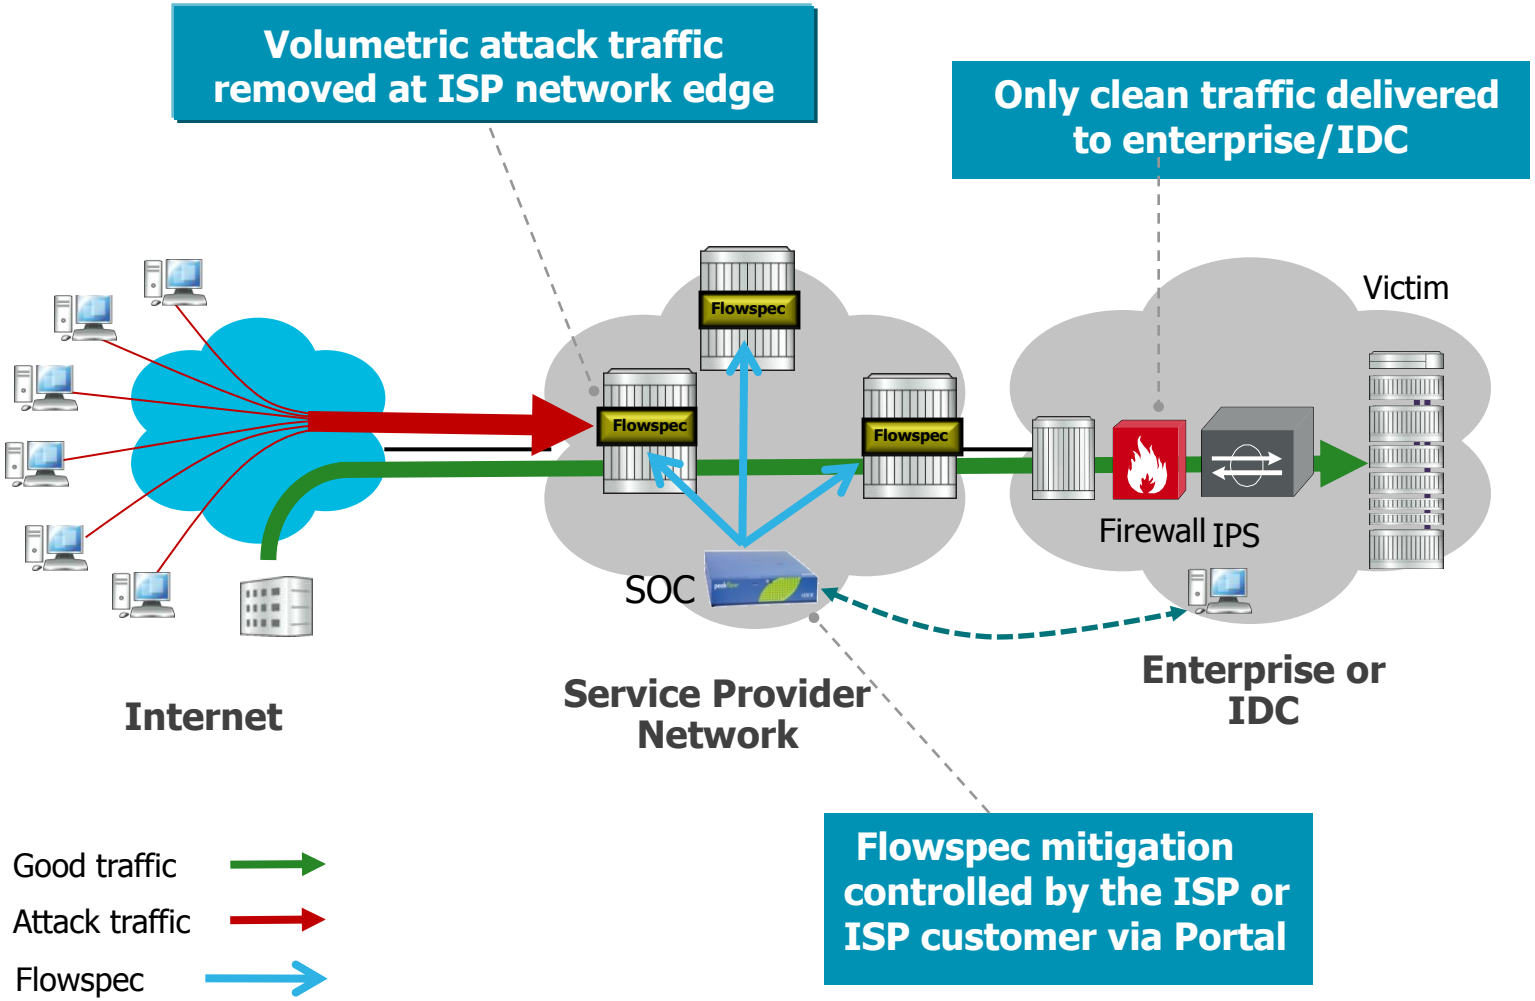

BGP Flow Specification – Vendors & Users

- Router vendors supporting BGP Flowspec:
  - Alcatel-Lucent 7750 SROS 9.0R1
  - Juniper JunOS 7.3
- DDoS mitigation vendors:
  - Arbor Peakflow SP 3.5
- BGP Tools:
  - ExaBGP Injector [7]
- Users:
  - North America: TW Telecom (TWTC) [6], other Tier 1, Tier 2
  - Europe: Tier 1, Tier 2
  - Latin America & Caribbean: RNP (Brasil) [8]
  - Flowspec itself is the 8th most used tool to mitigate attacks [5]

DDoS attacks: Increasing Scale & Sophistication

DDoS Attack Bandwidth (Gbps)

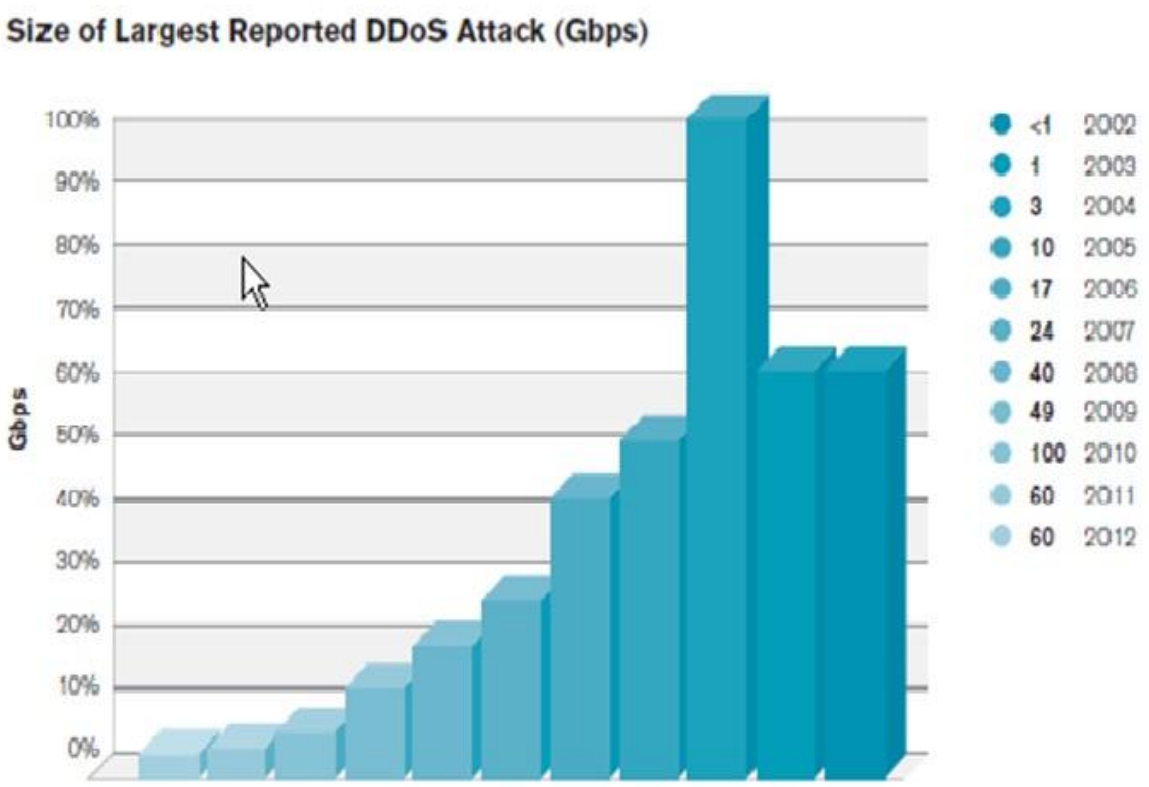

Source: Arbor Networks – 2012 Worldwide Infrastructure Security Report, Volume VIII

Volumetric Attacks
Network Infrastructure

Application Attacks
IDCs & Services

DDoS Protection with Mitigation Appliance (IDMS)
"Surgical Mitigation"

- Traffic anomaly is scrubbed by a DPI-capable mitigation appliance that **surgically** removes the attack traffic only.
- Mitigation appliances are also known as Intelligent DDoS Mitigation Systems (IDMS). IDMS are the second most used tool for DDoS protection [5].
- Able to mitigate complex, application-layer DDoS attacks without completing the attack.
- Typically a shared resource in the network infrastructure.
- Traffic anomalies need to be redirected in the network to go through the IDMS before reaching the intended destination:
  - Traffic Diversion or Offramping
  - Traffic Reinjection or Onramping

IDMS-based DDoS Mitigation Terminology

- Diversion or Offramping: rerouting of traffic destined to the victim to the DDoS mitigation appliance for scrubbing.
- Reinjection or Onramping: redirection of scrubbed (clean) traffic back to its intended destination.
- Typically, traffic diversion takes place through more specific BGP prefix announcements (victim addresses), usually in the GRT (called diversion/offramp route):
  - Easier to control & manipulate routes (NH, Communities)
  - Can be signaled across AS boundaries if required
- Traffic Reinjection usually requires tunneling or an alternate routing domain (VRF) to get clean traffic back to its intended destination without looping.

IDMS-based DDoS Mitigation
Scrubbing Center Design

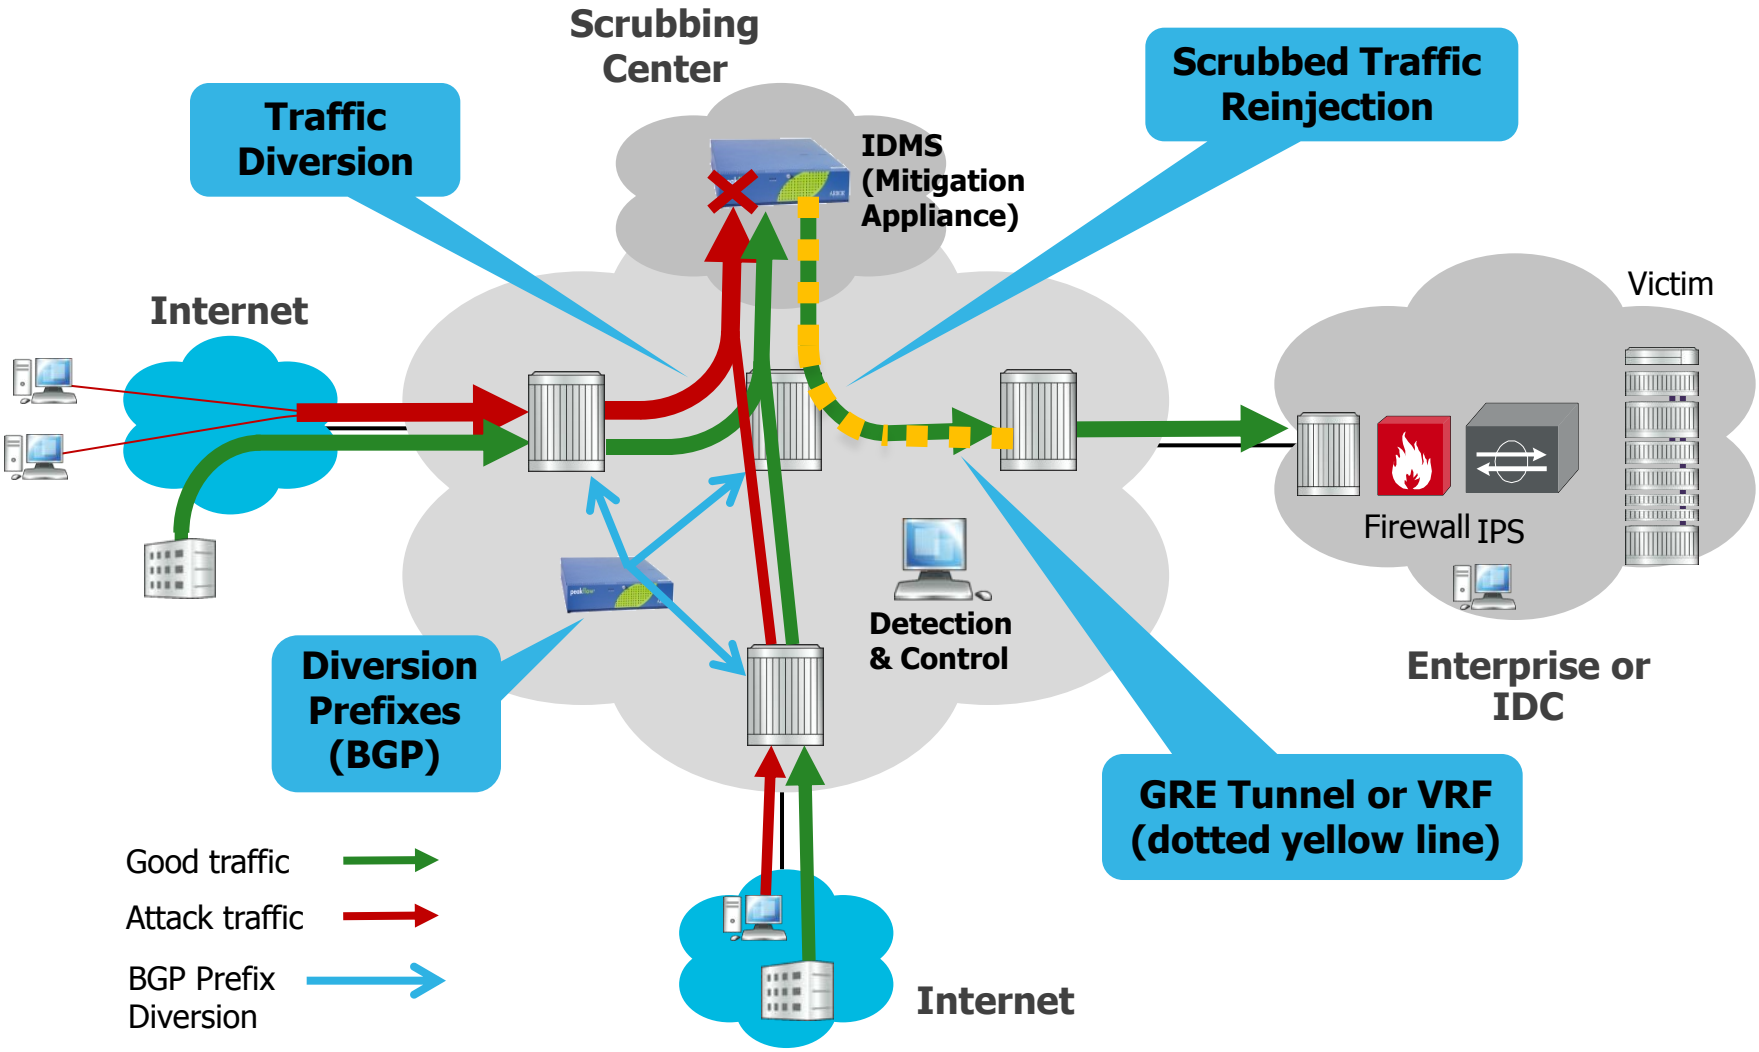

IDMS-based DDoS Mitigation
Distributed Design – Embedded IDMS

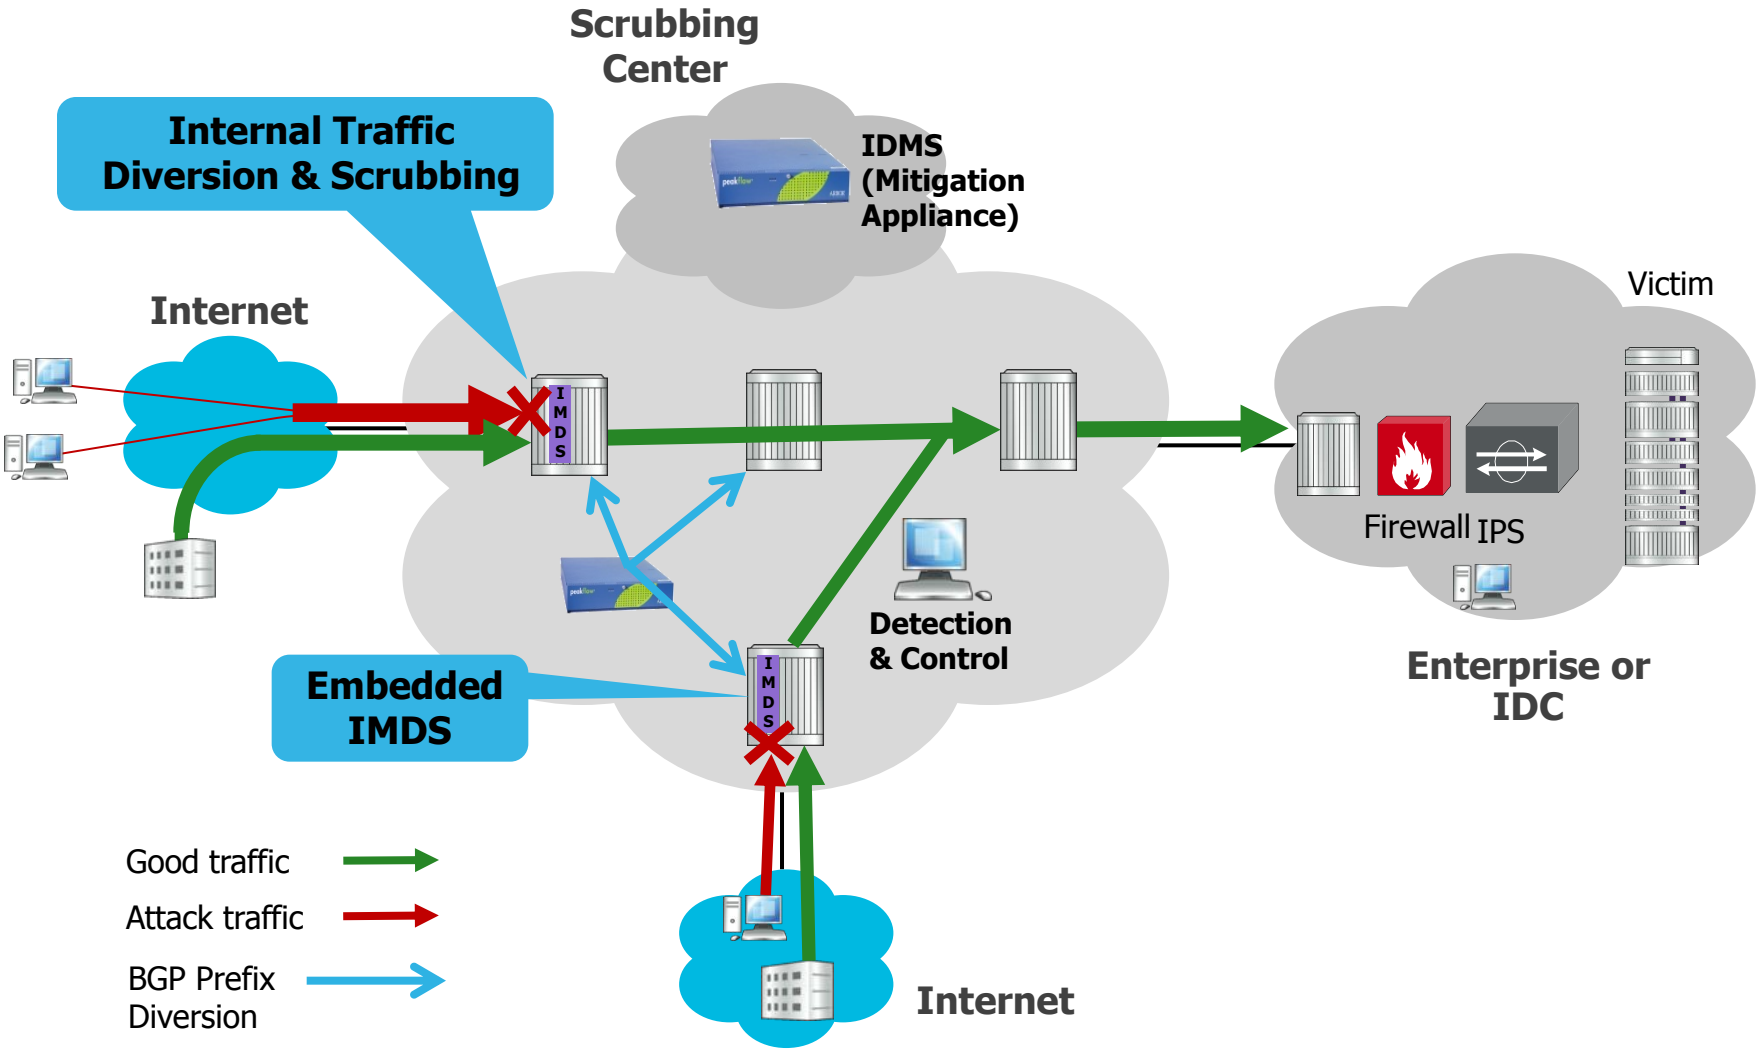

Traffic Diversion with BGP Prefixes – The Good & The Bad

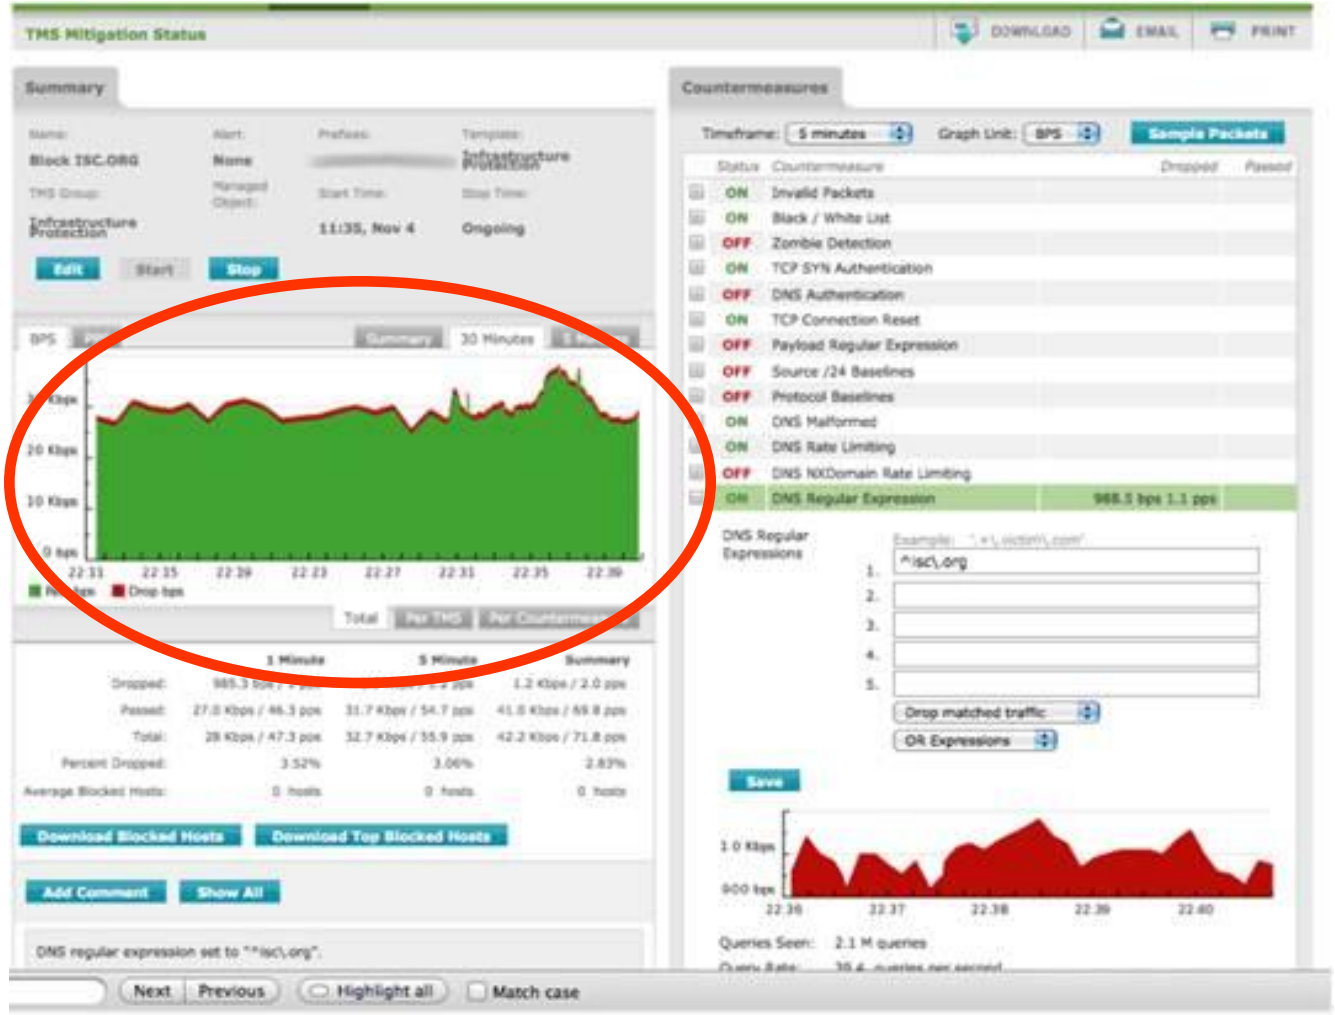

- Real mitigation of DNS attack

DDoS Mitigation Appliance – "Surgical Diversion" Using BGP Flowspec "Redirect to VRF" Action

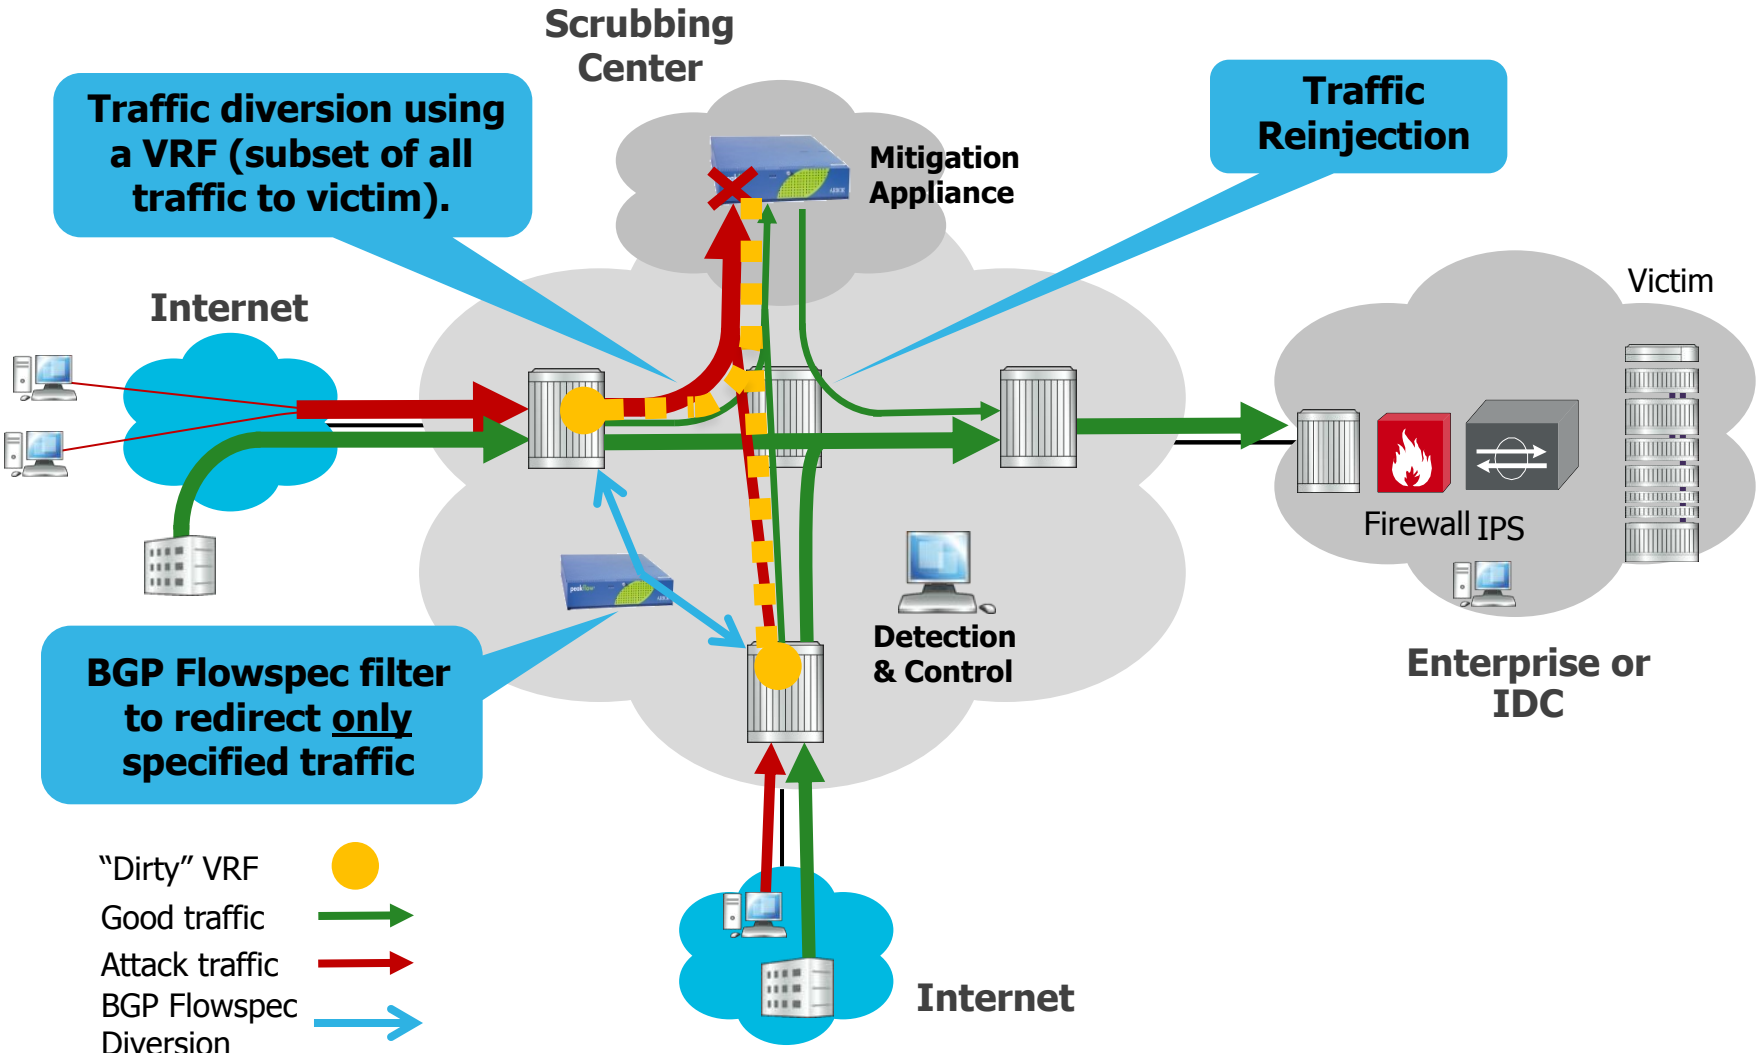

"Surgical Diversion" Using BGP Flowspec – Optimized Design & Operation

- No changes to the Global Routing Table (GRT)
  - Diversion performed by Flowspec NLRI
  - Flowspec filter Action configured to "Redirect to VRF"
    - Extended Community 0x8008.
    - Less intrusive to the routing system
- No need for a tunneling design for reinjection/onramping
  - Clean traffic can simply be sent back to the GRT
- More granular control of diverted traffic
  - Allows for the redirection of only a subset of the traffic to the victim: specific protocols, ports, source prefix, destination prefix
  - Less traffic overhead for Mitigation Appliance to deal with

"Surgical Diversion" Using BGP Flowspec – Enabling New Workflows

- Facilitates the implementation of new mitigation workflows for demanding use cases:
  - "Always on" Mitigations for critical resources:
    - HTTPS traffic only (normal web traffic follows on-demand mitigation model)
    - ICMP & UDP traffic
  - Victims with very large traffic volume
    - Divert just traffic from a certain block, or geographical region (based on IP Location)

"Surgical Diversion" for Data Centers –
Optimizing the Mitigation of Application-Layer DDoS Attacks

- On-demand or continuous mitigation of selective traffic
- Benefits of continuous DPI traffic scrubbing without the risks and demands of in-line deployments.

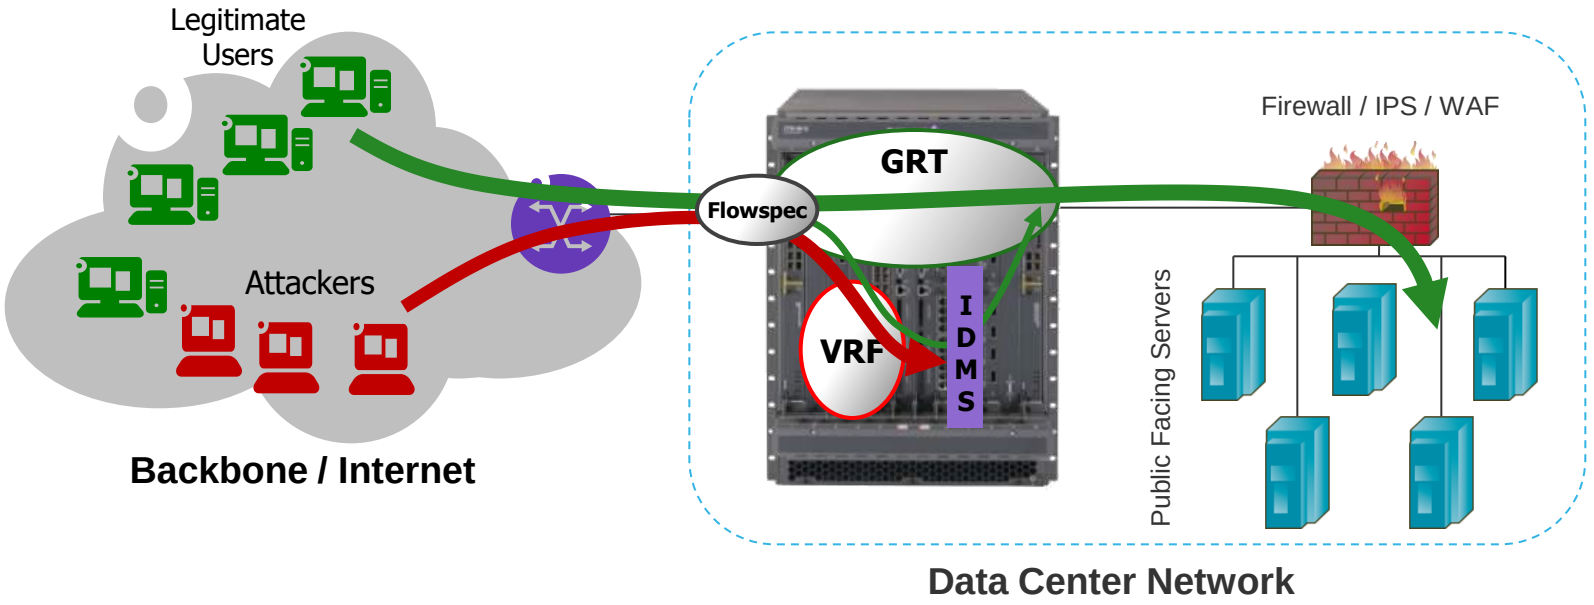

BGP Flowspec Diversion – Work in Progress

- New "Redirect to IP Next-Hop" Flowspec Action
  - Described in draft-simpson-idr-flowspec-redirect-02.txt [4]
  - New Extended Community value 0x0800
- Enables BGP flowspec redirection using encapsulations other than VRFs, such as GRE or MPLS tunnels.

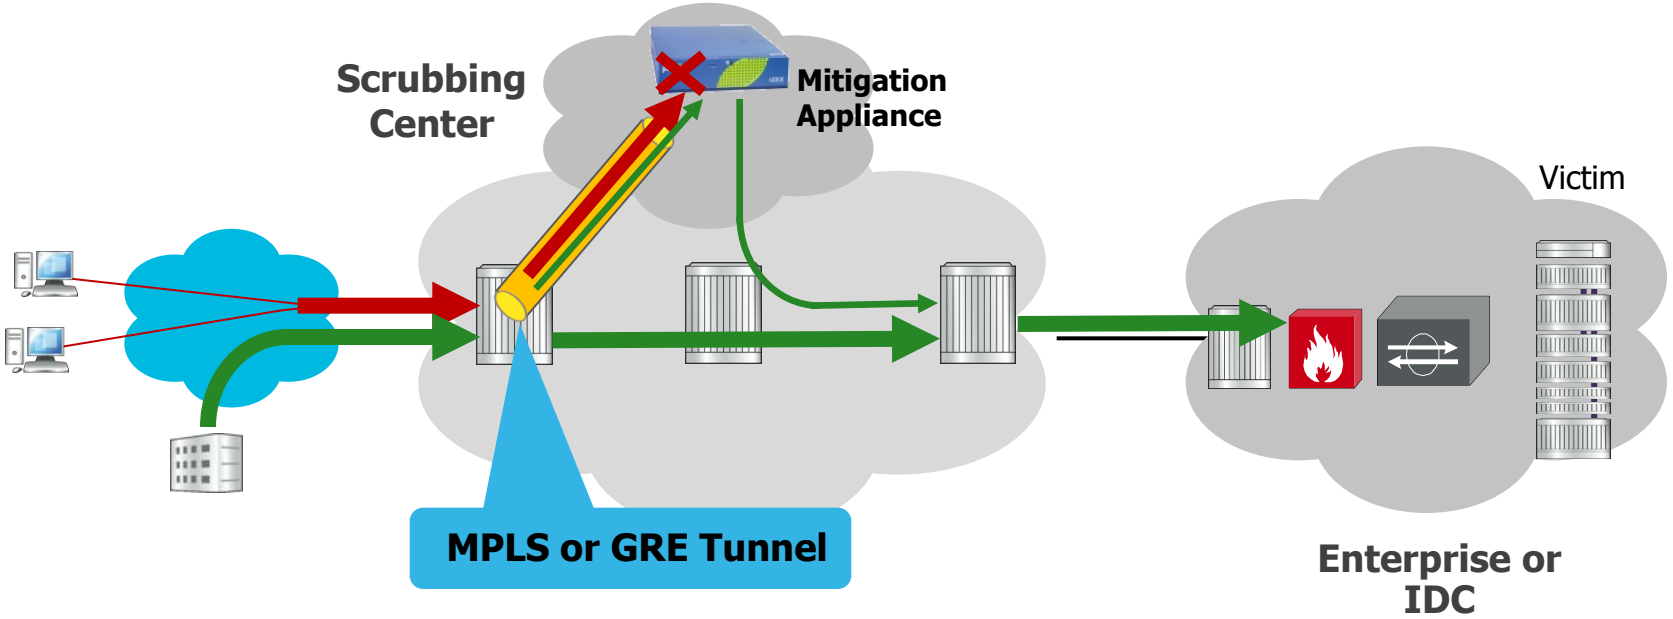

Summary – "Surgical Diversion" Benefits Using BGP Flowspec:

- Greatly simplifies traffic diversion design and operation:
  - Less intrusive to the routing system – no changes to the Global Routing Table (GRT).
  - On the reinjection side, there is no need to use tunneling (GRE) or VRF designs.
- Optimizes the benefits of a DDoS Mitigation Appliance ("surgical mitigation") with precise diversion:
  - Allows for a better optimization of the shared mitigation capacity of the mitigation appliance.
  - Addresses demanding mitigation use cases.
  - On-demand & continuous scrubbing per application or other criteria.

References:

- [1] RFC 5575, Dissemination of Flow Specification Rules
- [2] draft-ietf-idr-flow-spec-v6-03 – Dissemination of Flow Specification Rules for IPv6
- [3] draft-ietf-idr-bgp-flowspec-oid-01 – Revised Validation Procedure for BGP Flow Specifications
- [4] draft-simpson-idr-flowspec-redirect-02.txt – BGP Flow-Spec Extended Community for Traffic Redirect to IP Next Hop
- [5] Arbor Networks – 2012 Worldwide Infrastructure Security Report, Volume VIII
- [6] 2006 - NANOG 38, D. Gassen, R. Lozano (Time Warner Telecom), D. McPherson, C. Labovitz (Arbor Networks), "BGP Flow Specification Deployment Experience"
- [7] 2010 - LINX69, Thomas Mangin (Exa Networks), Andy Davidson (NetSumo), "BGP Route Injection" http://www.andyd.net/media/talks/BGPRouteInjection.pdf
- [8] GTER/GTS 2007, Raniery Pontes (RNP), "Flowspec em ação - Experiência de uso no backbone da RNP"

AT
THE
SPEED
OF
IDEAS™

Alcatel·Lucent

www.alcatel-lucent.com