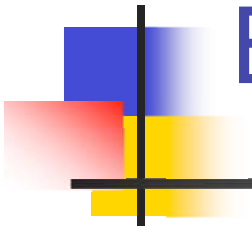

Customer-Triggered Real-Time Blackholes

Tim Battles tmbattles@att.com
Danny McPherson danny@arbor.net
Chris Morrow chris@uu.net

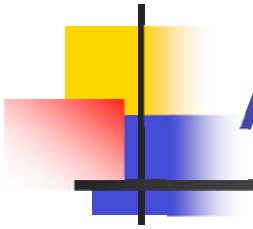

Agenda

- About Blackhole Routing
- Preparing the Tools
- Customer-Triggered Blackholes
- BGP Flow Specification

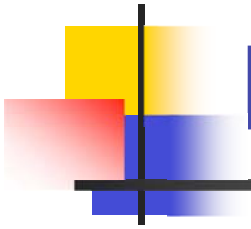

Before We Begin…

- How many folks in the room are responsible for network security at an ISP or enterprise?
- How many folks here employ blackhole routing today?
- How many employ source-based blackhole routing?
- How many folks here currently support customer-triggered blackhole routing?

About (D)DOS

- *It could come from anywhere; be prepared!*

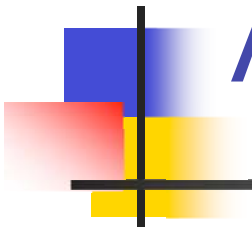

About Blackhole Routing

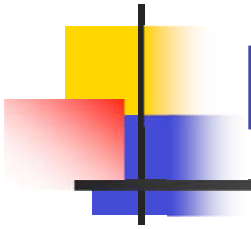

Remote-Triggered Blackholes

- Remote-triggered Blackhole filtering is the foundation for a whole series of techniques to traceback and react to (D)DOS attacks on an ISP's network.
- Preparation is key and does not impact ISP operations or network performance.
- Adds significant capabilities to an ISP's security toolkit!

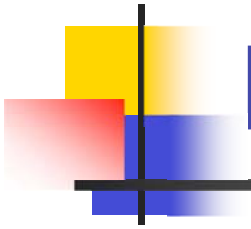

Miscellaneous

- Detection & Traceback
- ACLs are difficult to deploy (e.g., augment, deployment time, configuration management, performance, hardware support, etc..)
- NetFlow
- IP Accounting
- Raw Interface Stats

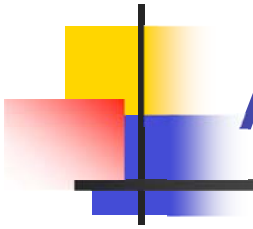

About Blackhole Routing

- *Blackhole Routing* or *Blackhole Filtering* results in packets being forwarded to a router's *bit bucket*, also know as:
  - Null 0
  - Discard Interface
- Initially worked only based on destination address, per it's exploit of a routers forwarding logic
- Typically results in desired packets being dropped with minimal or no performance impact.

Exploits Forwarding Logic

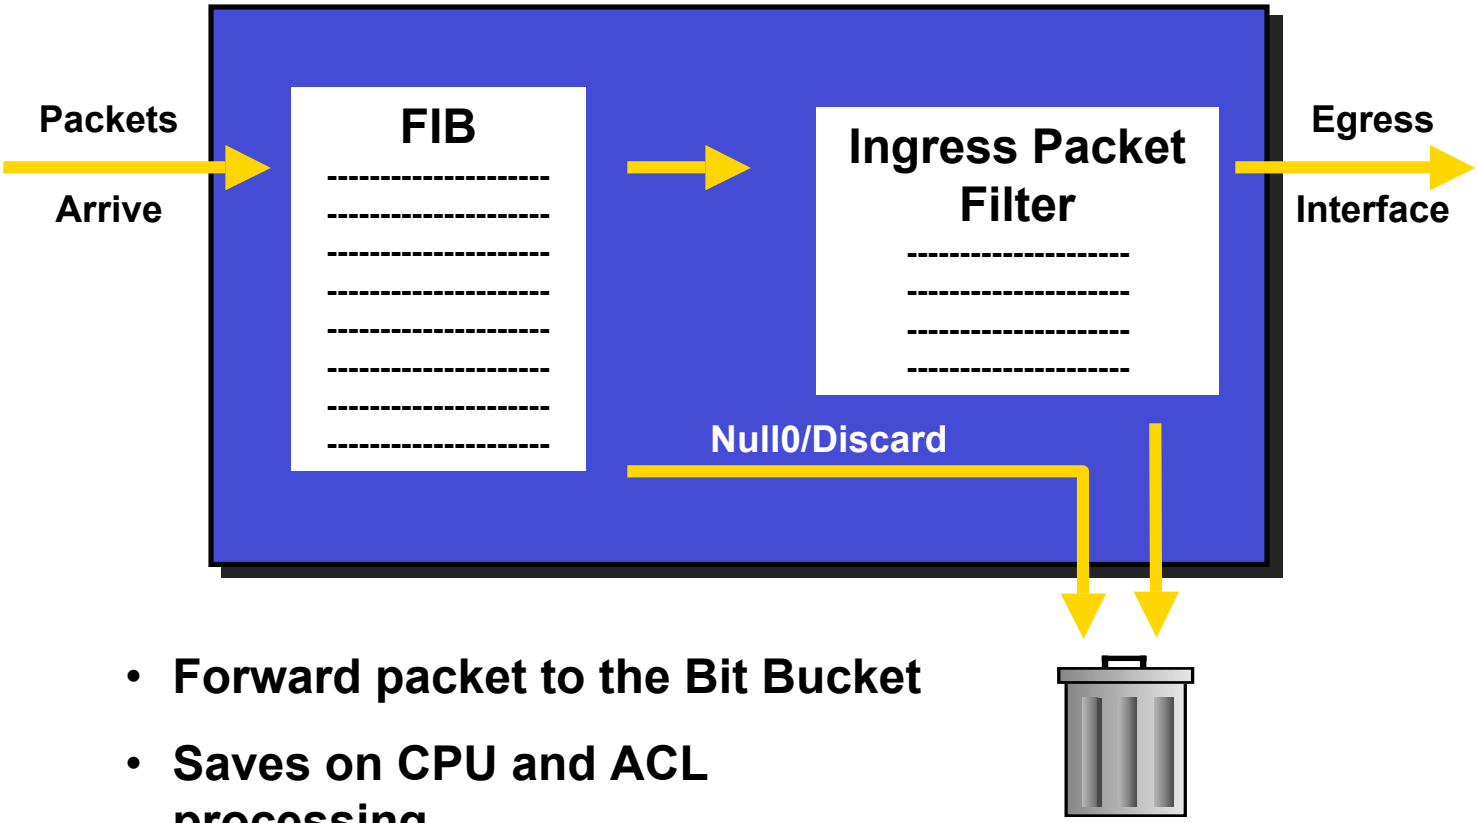

- **Forward packet to the Bit Bucket**
- **Saves on CPU and ACL processing**

Customer is DOSed – Before

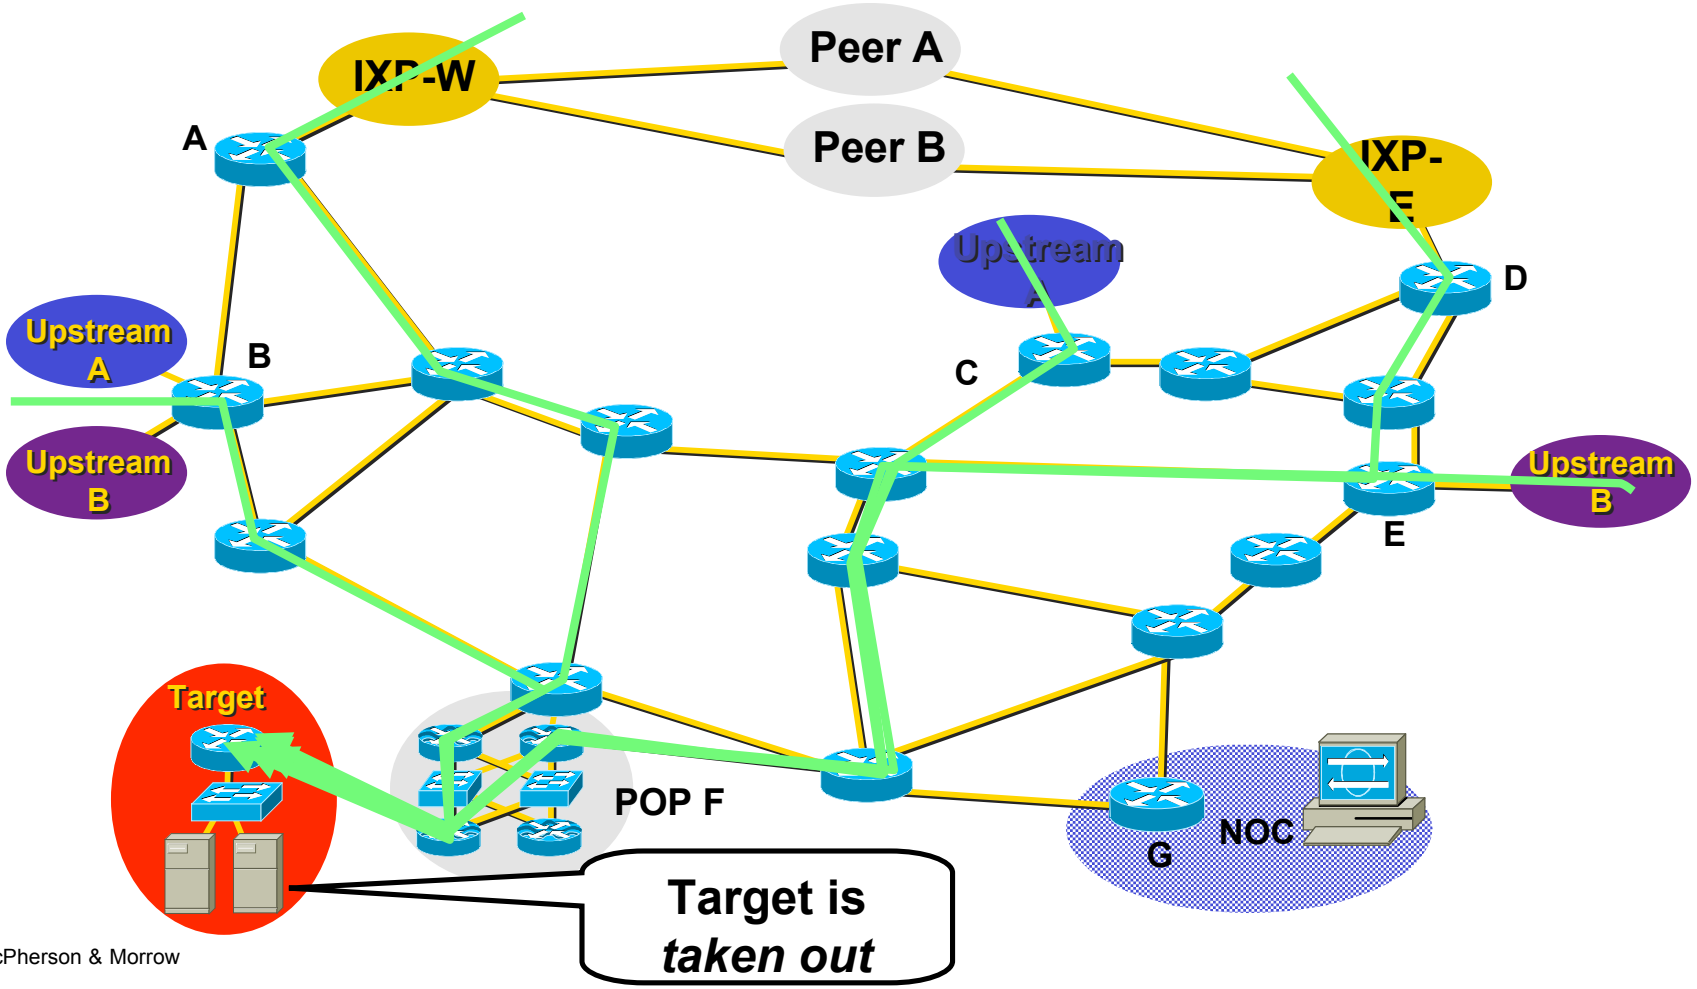

Customer is DOSed – Before – Collateral Damage

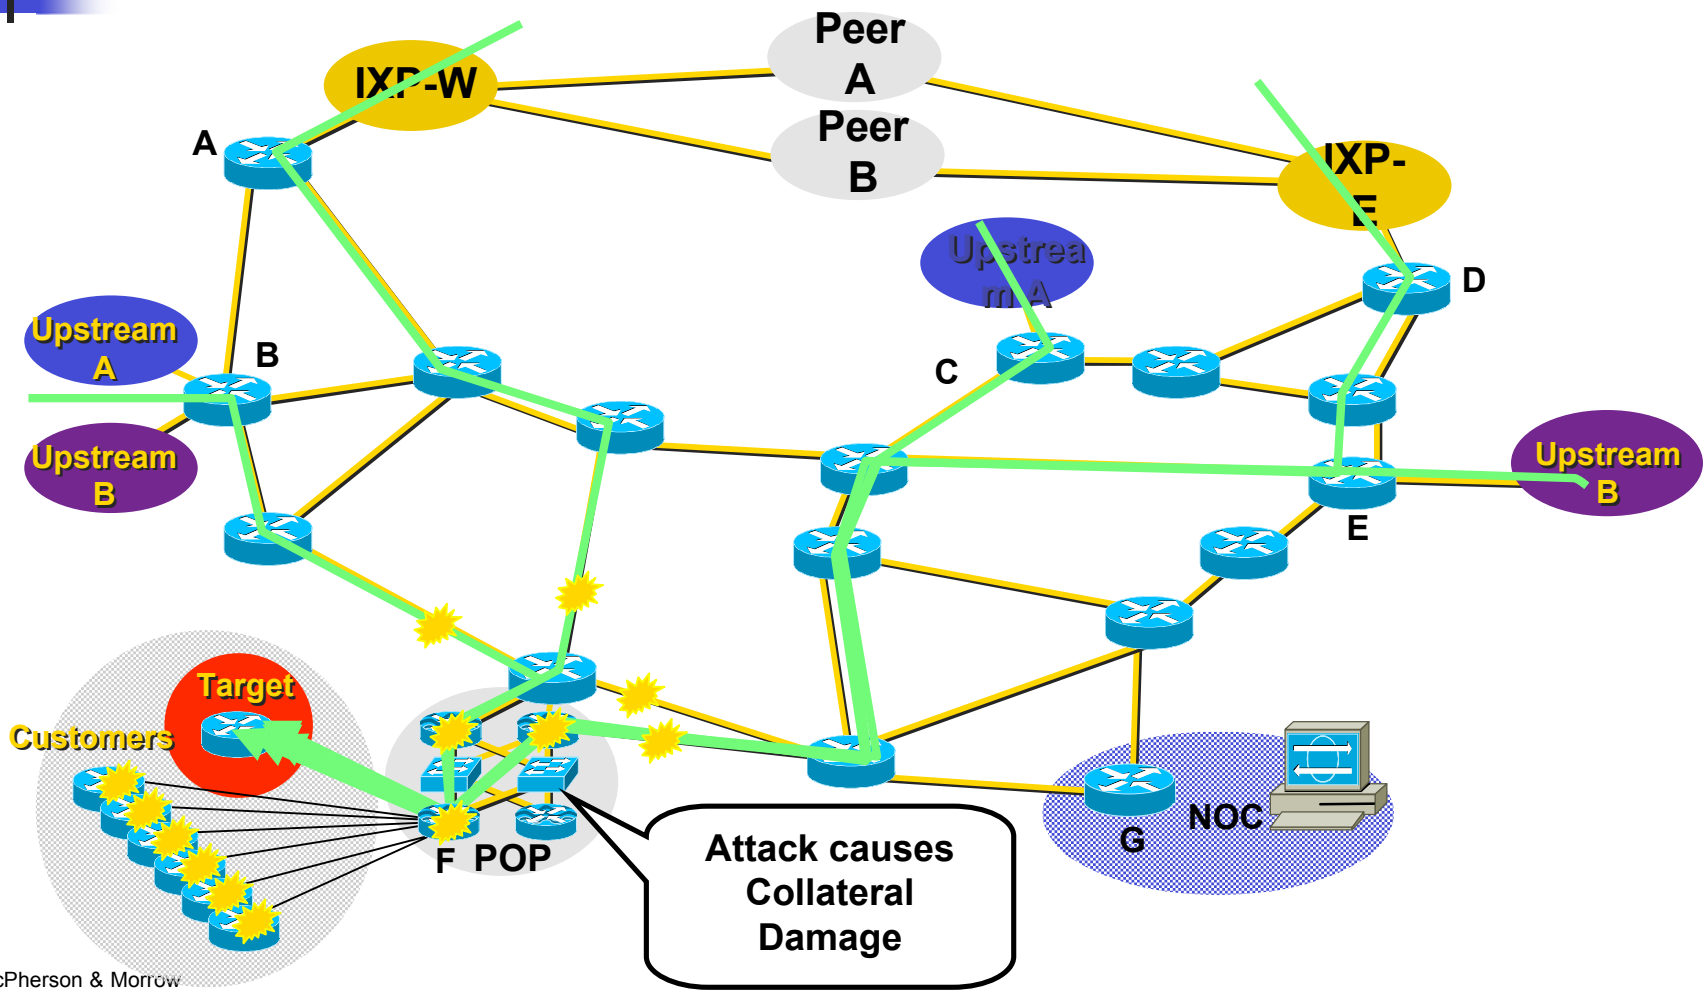

Customer is DOSed – After – Packet Drops Pushed to the Edge

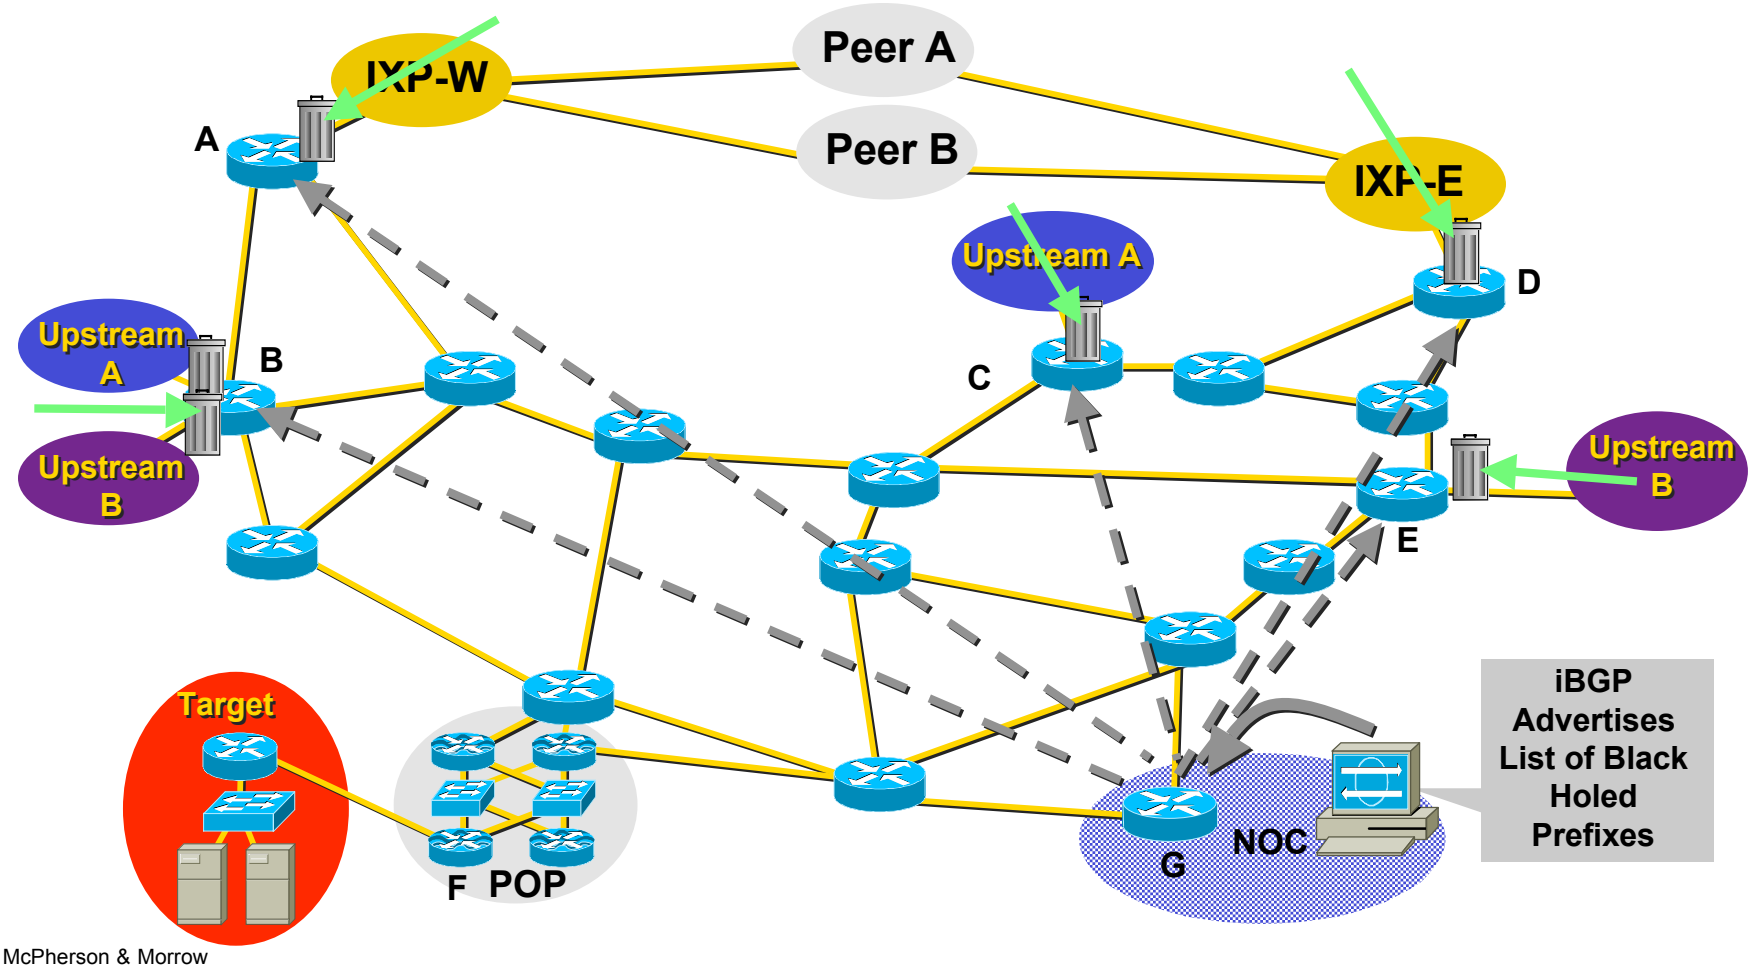

Preparing for Blackhole Routing

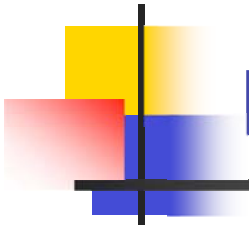

Remotely Triggered Blackhole Filtering

- Use BGP to trigger a network-wide response to a multi-source attack flow
- A static route and BGP will allow an ISP to trigger network-wide destination address blackholes as quickly as iBGP converges through the network.
- Provides ISPs a tool that can be used to respond to distributed denial of service events or employ techniques such as Backscatter Traceback[backscatter]

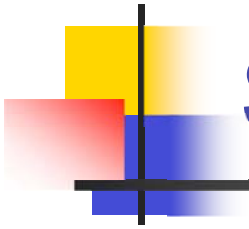

Step 1: Prepare all the Routers

- Allocate a small block of address space (e.g., RFC 1918 space or IANA reserved space) to be dedicated for black hole filtering. TEST-NET [RFC 3330], 192.0.2.0/24 is a potential option.
- Configure a static route on each router with your selected route, pointing to Null 0 or the discard:

```
ip route 192.0.2.1 255.255.255.255 Null0 255
ip route 192.0.2.2 255.255.255.255 Null0 199
ip route 192.0.2.3 255.255.255.255 Null0 50
```

Step 1- Prepare all the Routers w/ Trigger

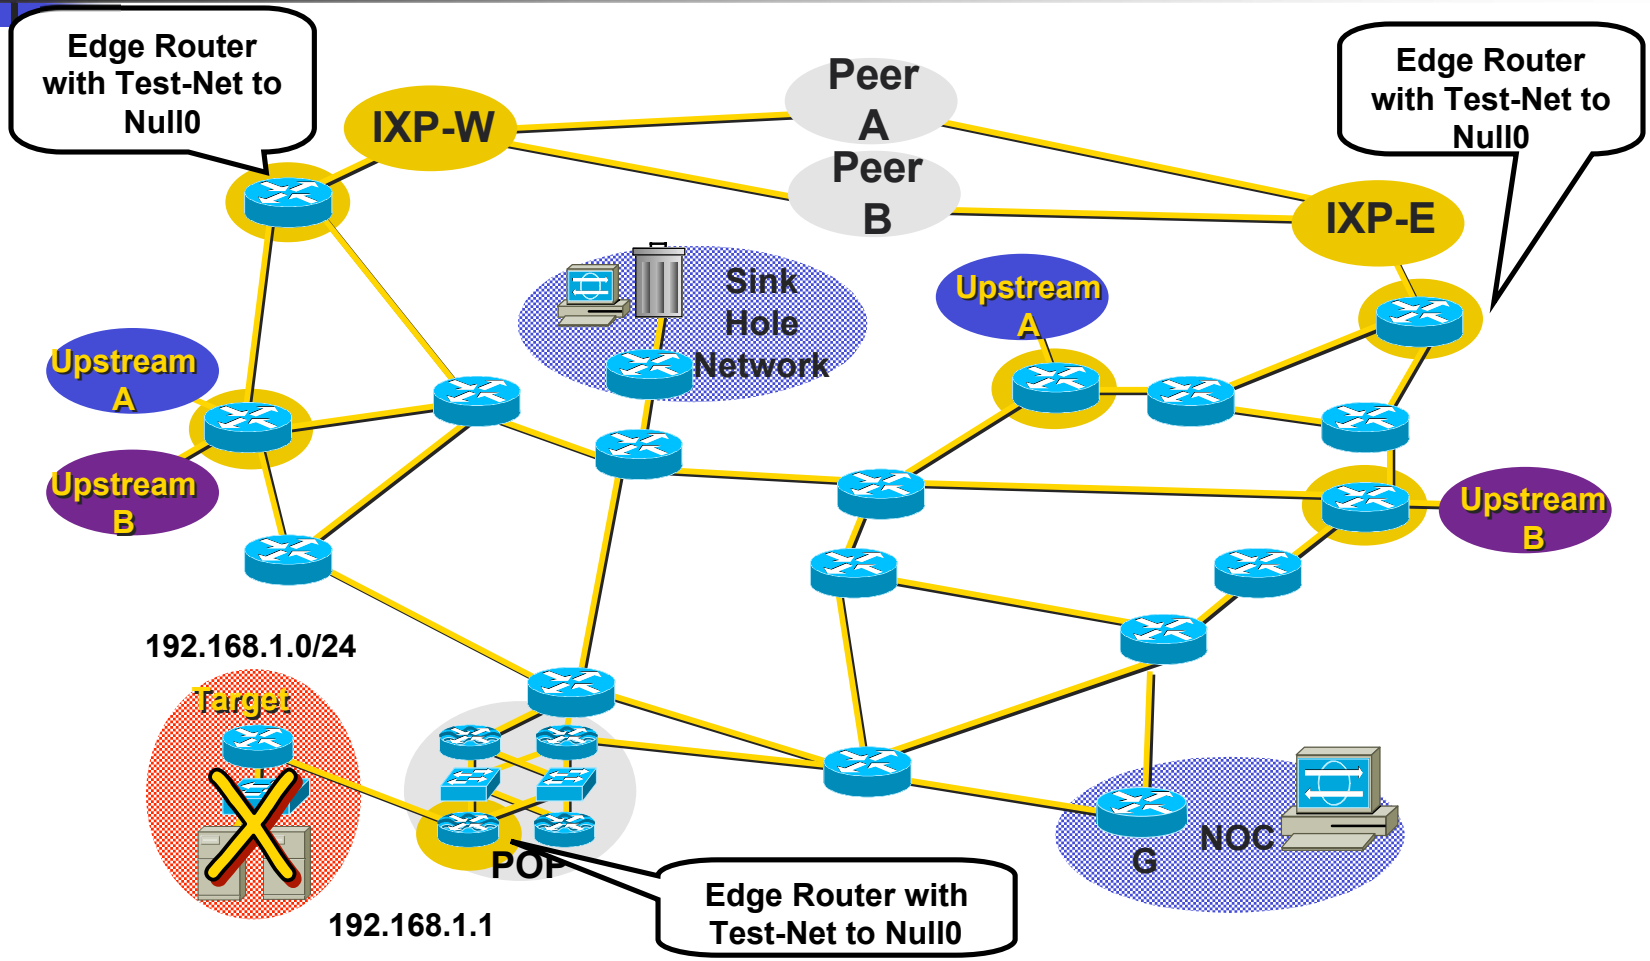

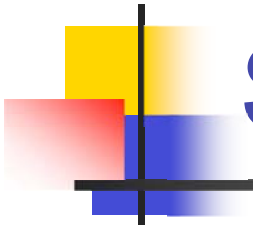

Sample TEST-NET Allocation

Address Block	Purpose
192.0.2.1/32	All iBGP routers for “Drop to NULL0”
192.0.2.2/32	All Peering Edge routers drop
192.0.2.3/32	All Customer Edge routers drop
192.0.2.4/30	Monitor Link addresses NOTE: provision these addresses in all Sinkholes
192.0.2.254	ANYCAST Sinkhole Address
192.0.2.8 -> balance	Sinkhole Diversion Addresses

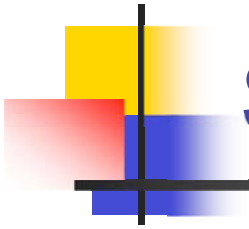

Step 2: Prepare the Trigger Router

- The trigger router is the device that will inject the iBGP announcement into the ISP's network
  - Should be part of iBGP mesh, but need not accept routes
  - Can be a separate router (or security tool)
  - Can be a production router
  - Can be a workstation with Zebra/GateD (interface with PERL scripts or other tools)
  - Commercial tools such as Arbor's Peakflow

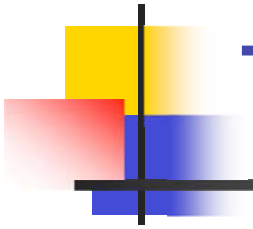

Tools Integration

- Could use to redirect to sinkhole analysis network:
  - Watch for consistent SYN-ACK storms
  - Worm detector (watch for scans, collect intel based on ports and signatures)
  - Background noise classification
  - Dark address monitoring & packet analysis
- Backscatter trigger
- No additional work after initial policy is implemented

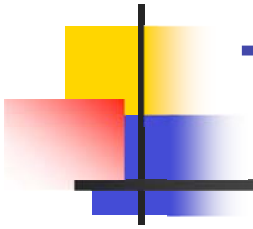

Tools Integration (cont..)

- Recommend dedicated trigger device, via routers with AAA & OTP, etc.. or a commercial tool.
- Couple with NetFlow or SNMP Data collection tools to identify scope, scale, duration and other characteristics of an attack and provide post-mortem/forensics data analysis functions, clarify billing disputes, etc..

Attack Detection Tools..

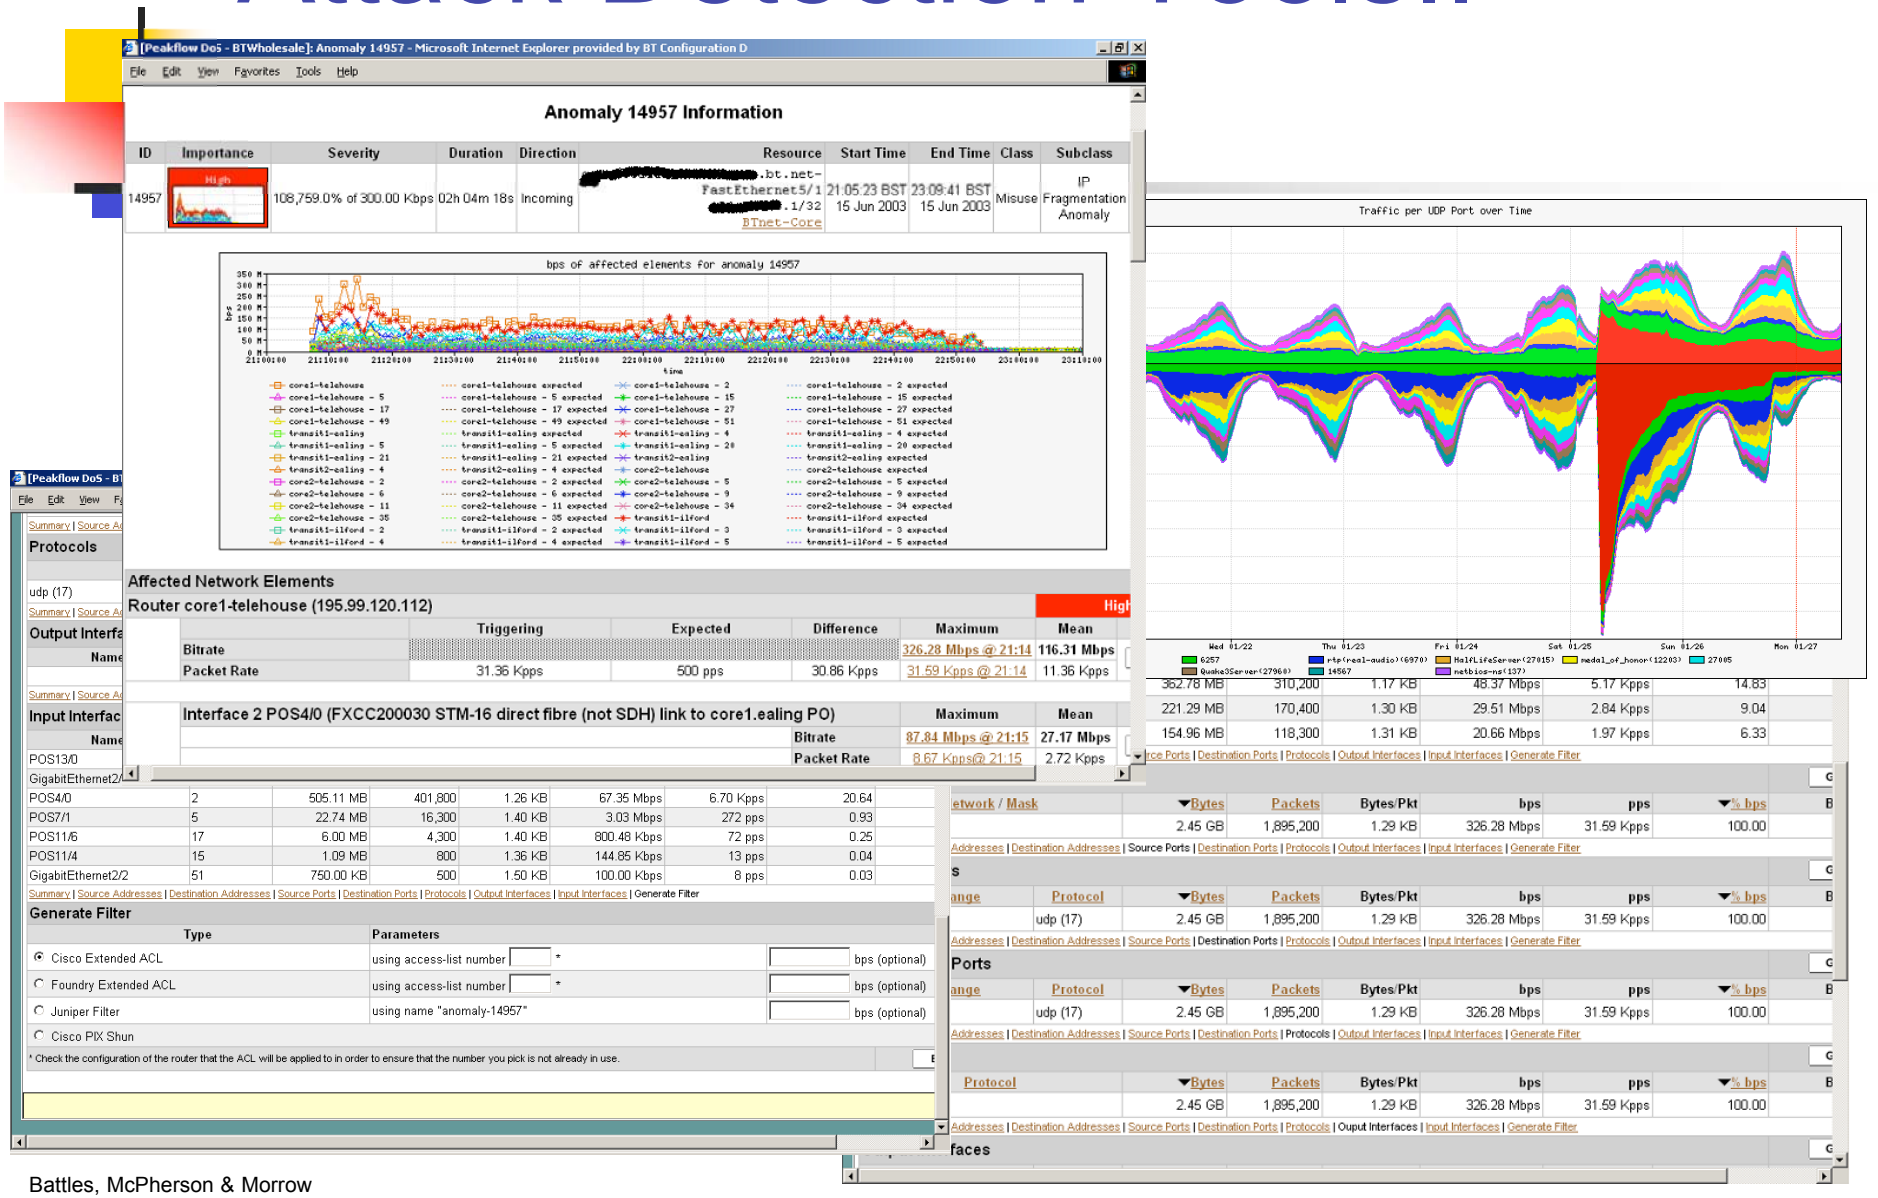

Commercial Tools…

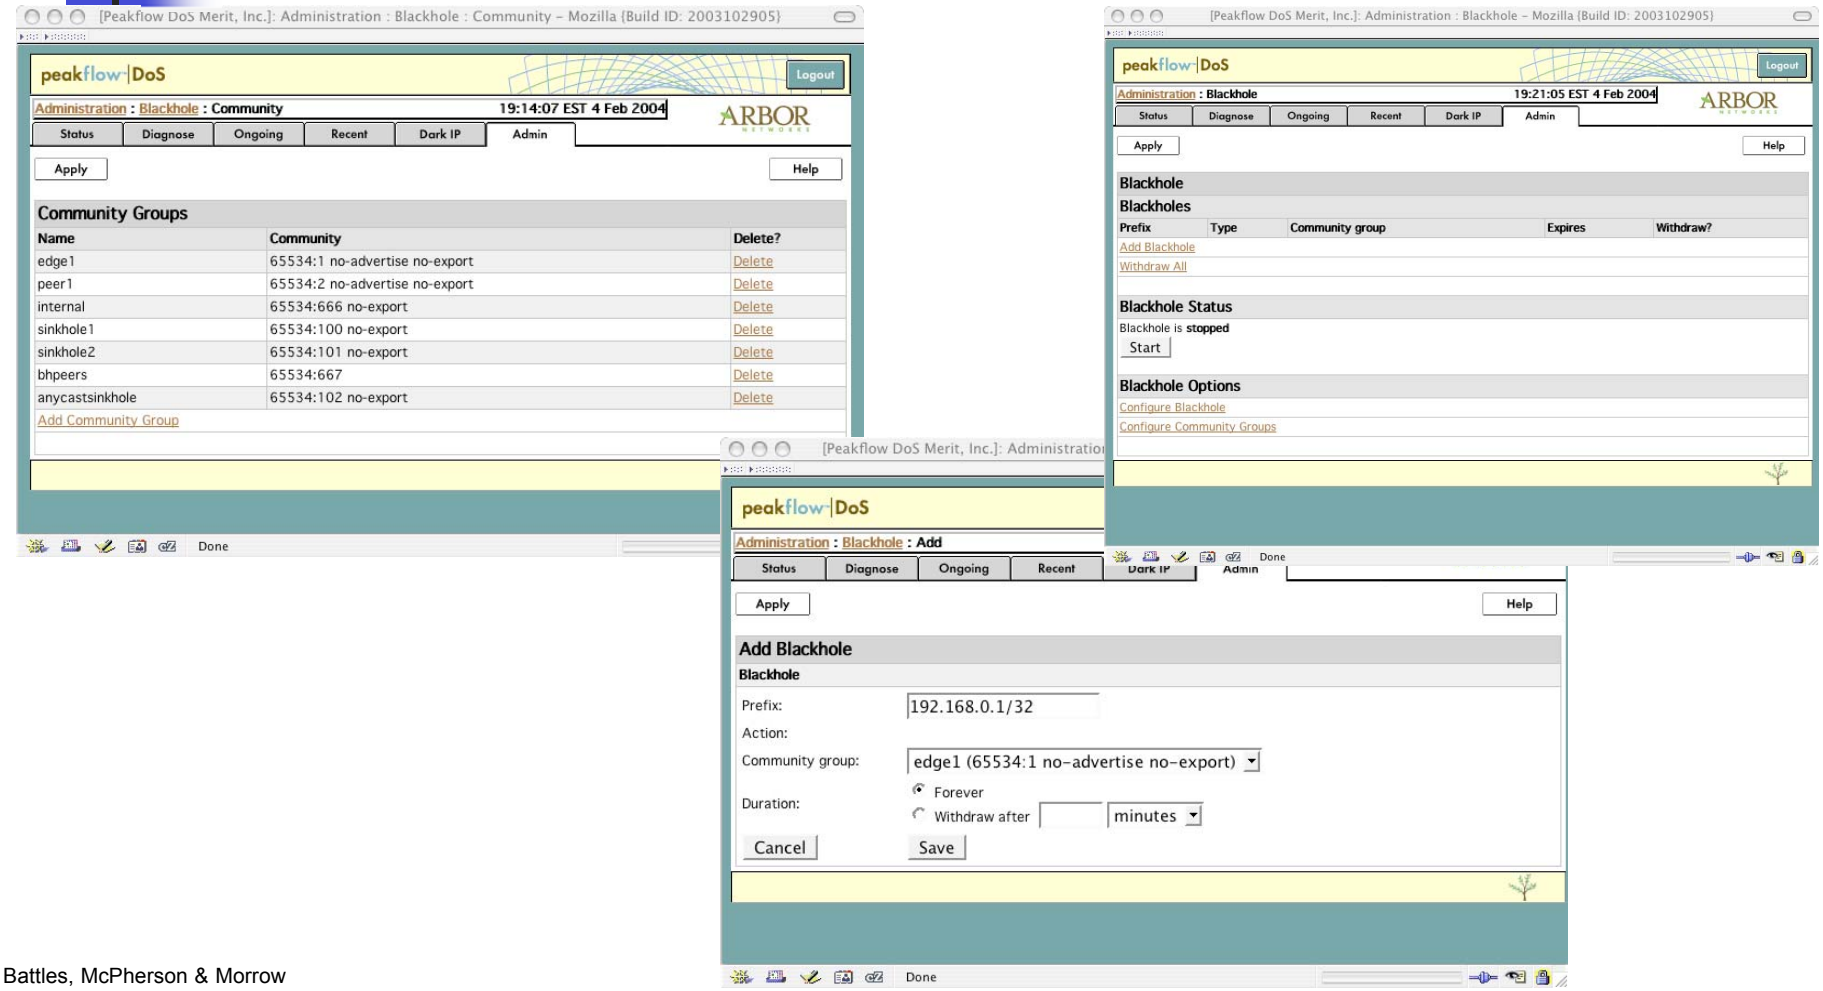

Trigger Router's Configuration

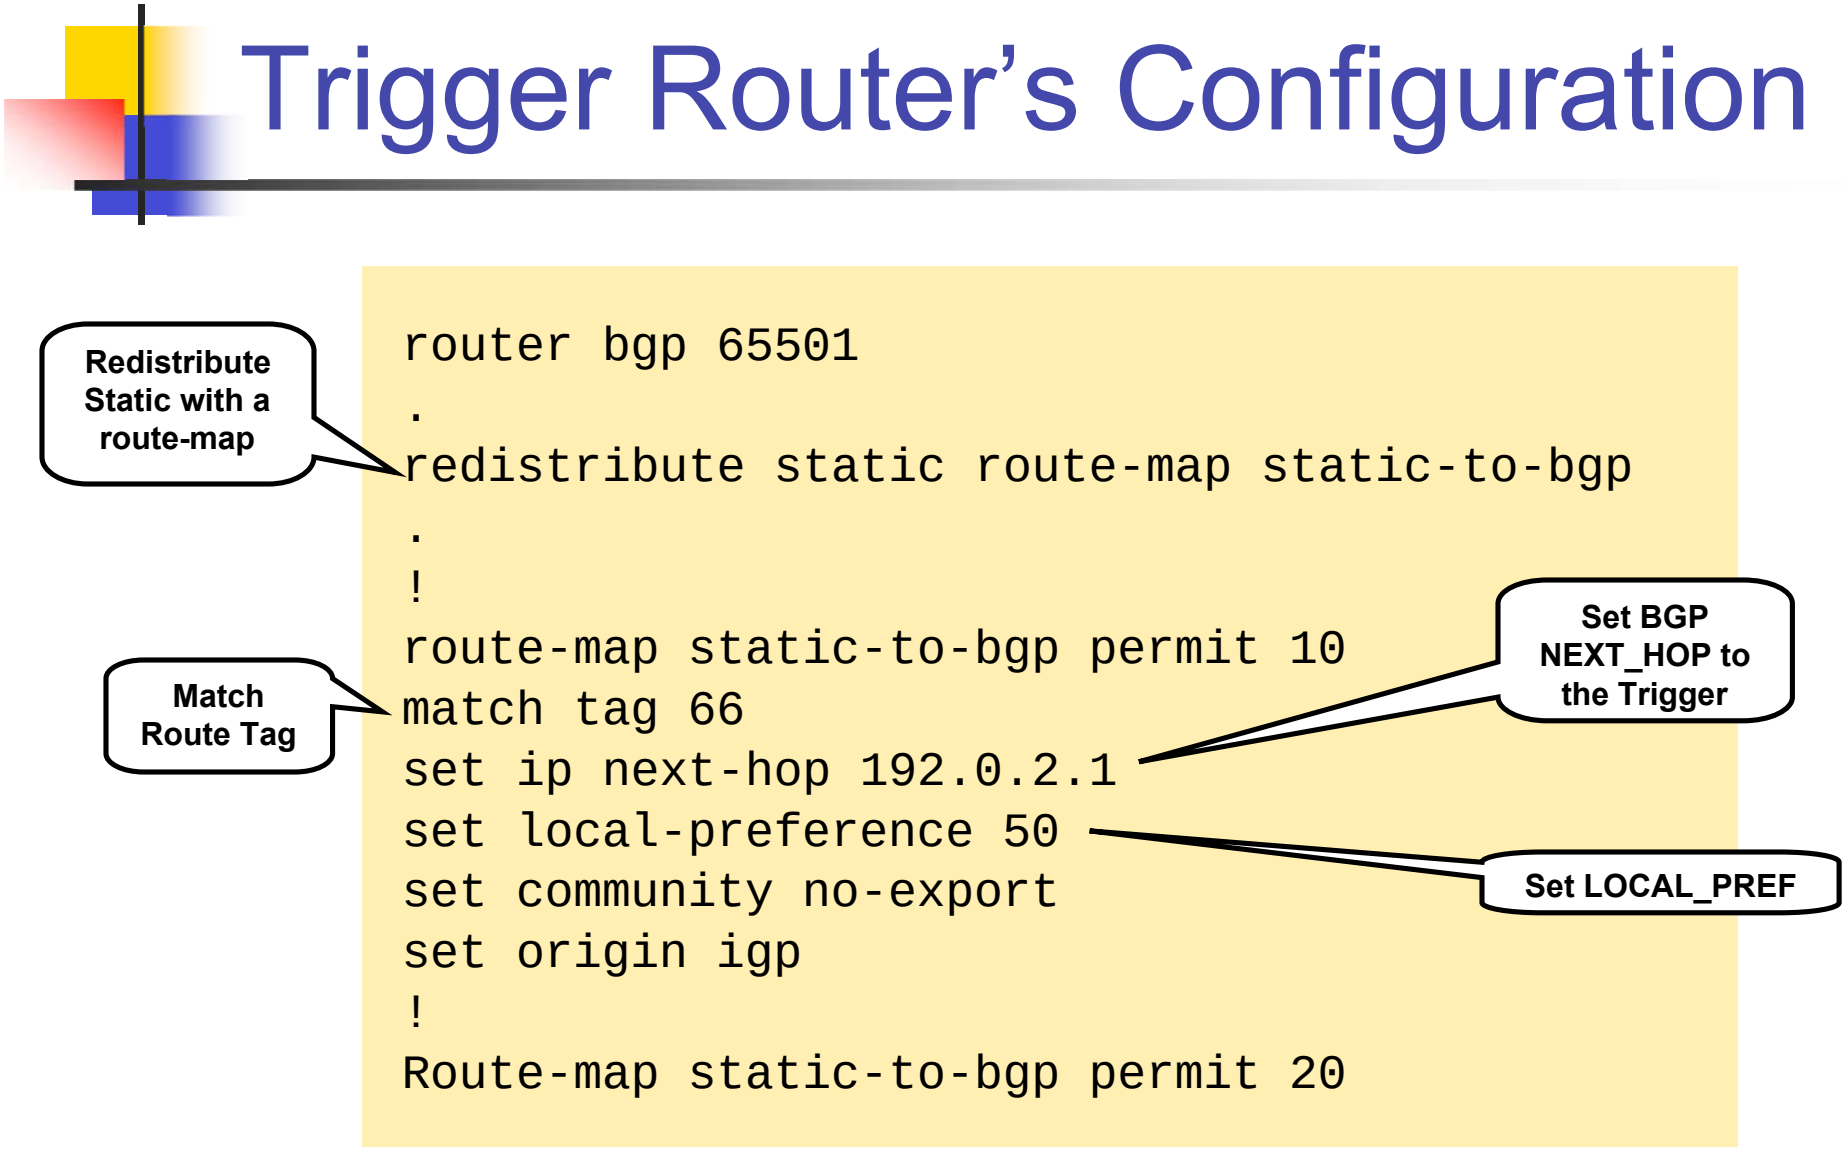

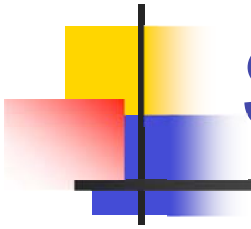

Step 3: Activate the Blackhole

- ISP adds a static route to the advertising router for the destination address they wish to blackhole. The route is added with *tag 66* to keep it separate from other static routes on the router.

```
ip route 192.168.1.1 255.255.255.255 Null0 Tag 66
```

- BGP Advertisement goes out to all BGP speaking routers
- Routers hear the announcement, glues it to the existing static route on the route, changes the BGP NEXT_HOP for the advertised route to Null 0
- Packets bound for destination are forwarded to Null 0/discarded

Step 3 – Activate the Black Hole

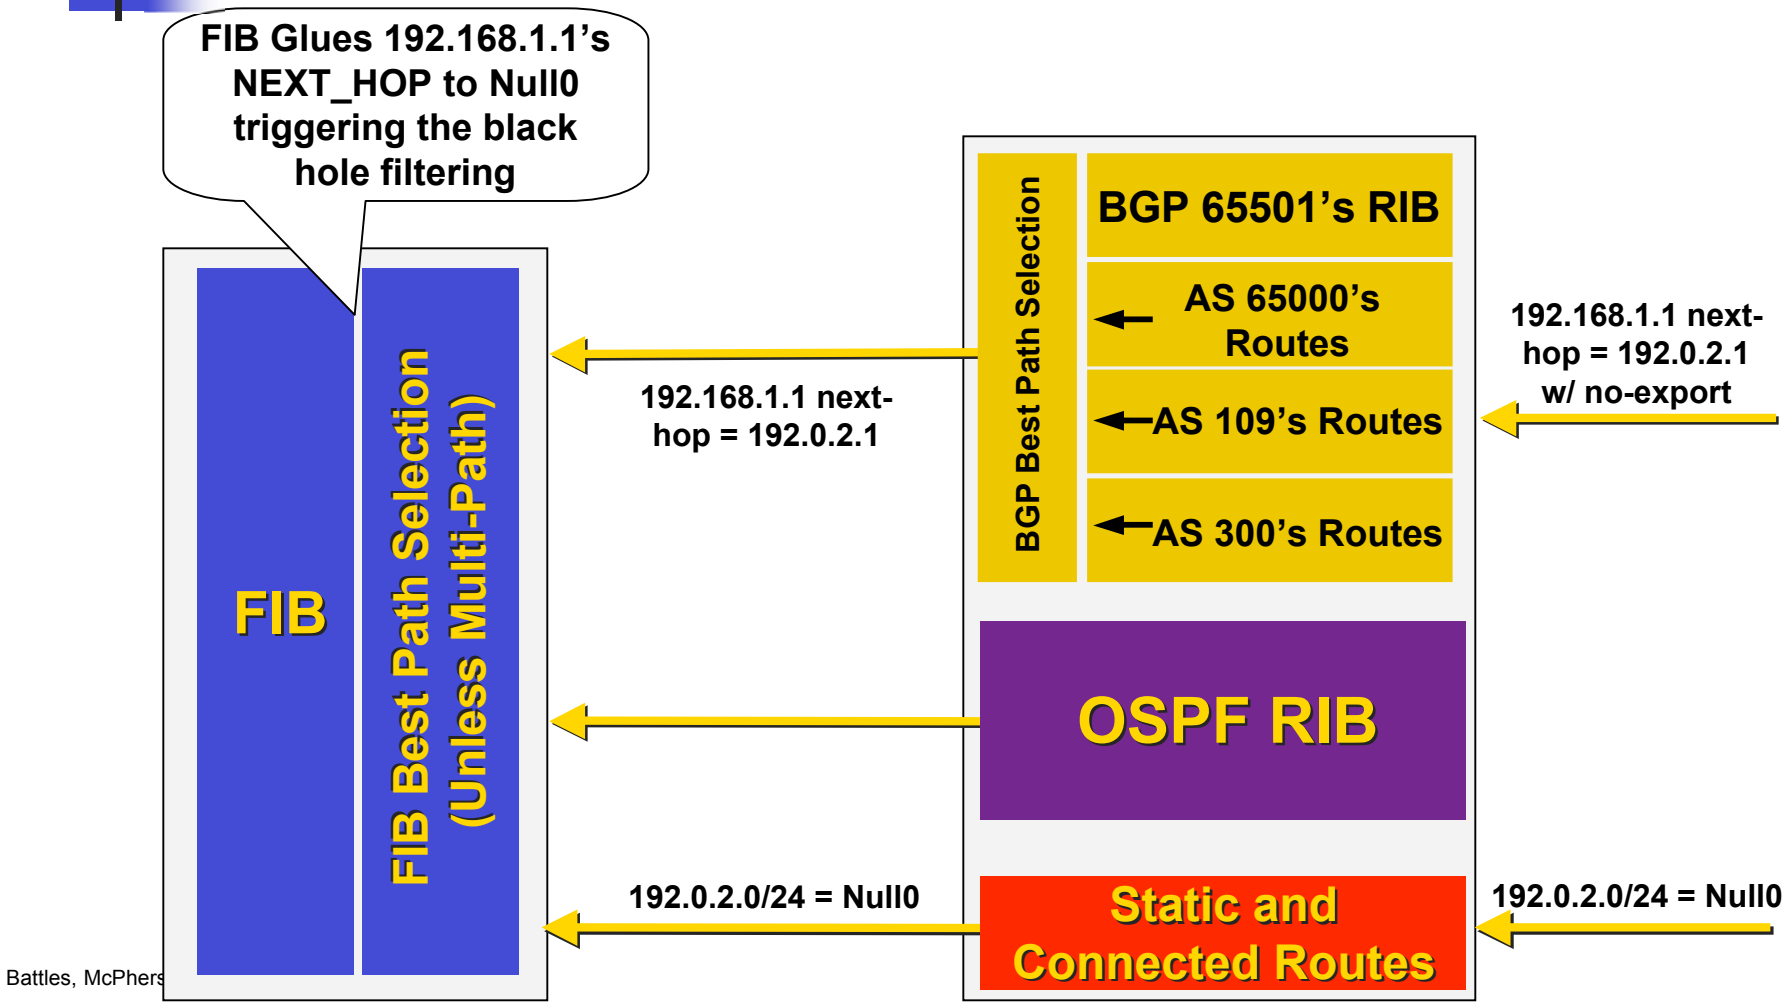

Step 3 – Activate the Black Hole

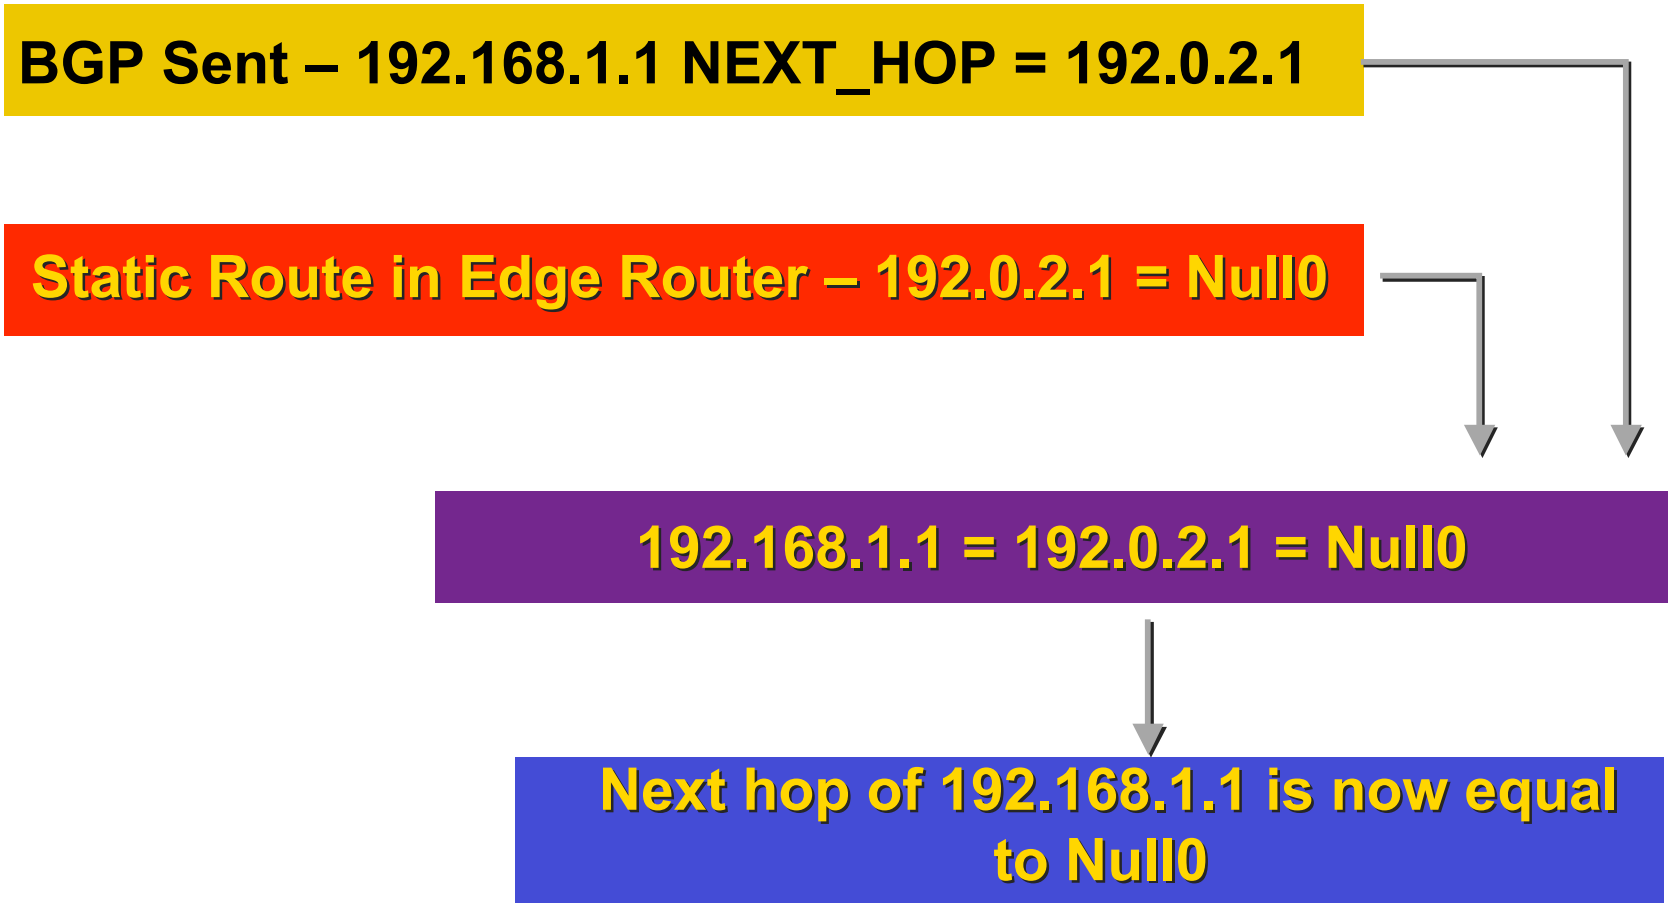

Step 3 – Activate the Black Hole

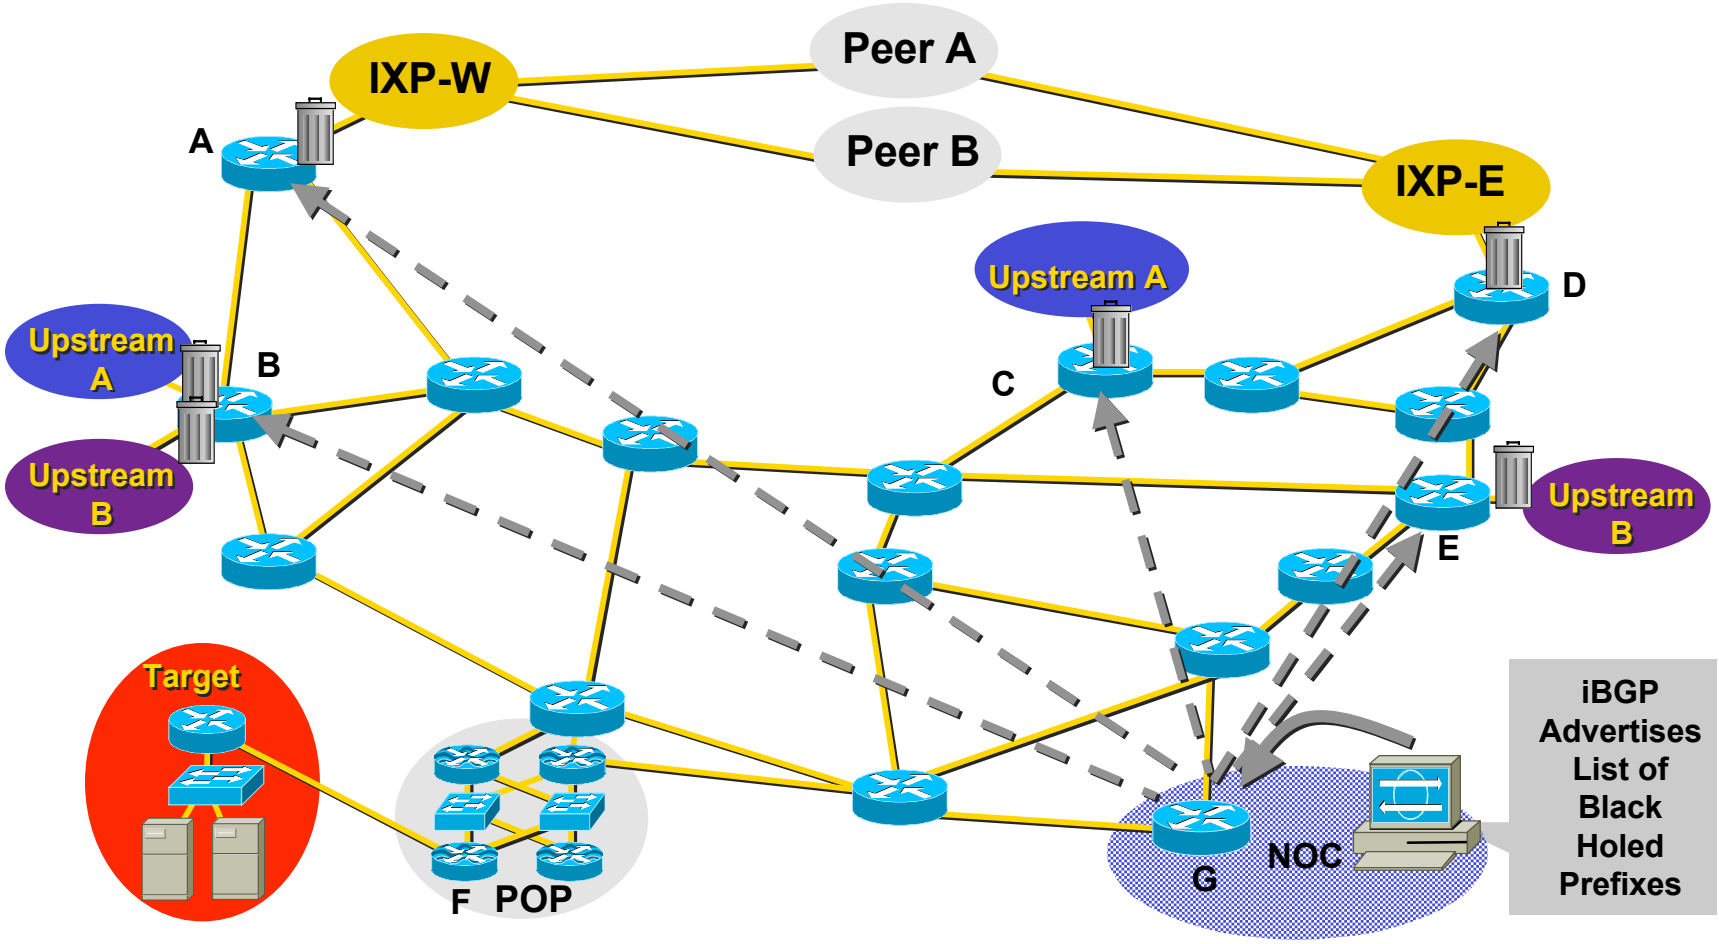

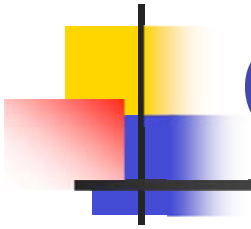

Community Based Trigger

- BGP Community-based triggering allows for more granular control over where you drop the packets.
- Three parts to the trigger:
  - Static routes to Null 0 on all the routers.
  - Trigger router sets the community and advertises the BGP update.
  - Reaction Routers (on the edge) matches community and sets the next-hop to the static route which maps to Null0.

Why Community Based Triggering?

- Flexibility, allows for more control on the DOS/DDOS reaction:
  - Community #1 can be for all routers in the network.
  - Community #2 can be for all peering routers. No customer routers – Preserves customer-customer connectivity if the victim is within your AS.
  - Community #3 can be for all customers (e.g., to push a inter-AS traceback to the edge of your network).
  - Trigger Communities per ISP Peer can be used to only black hole on one ISP Peer's connection. Allows for the DOSed customer to have partial service.

Inter-Provider Mitigation

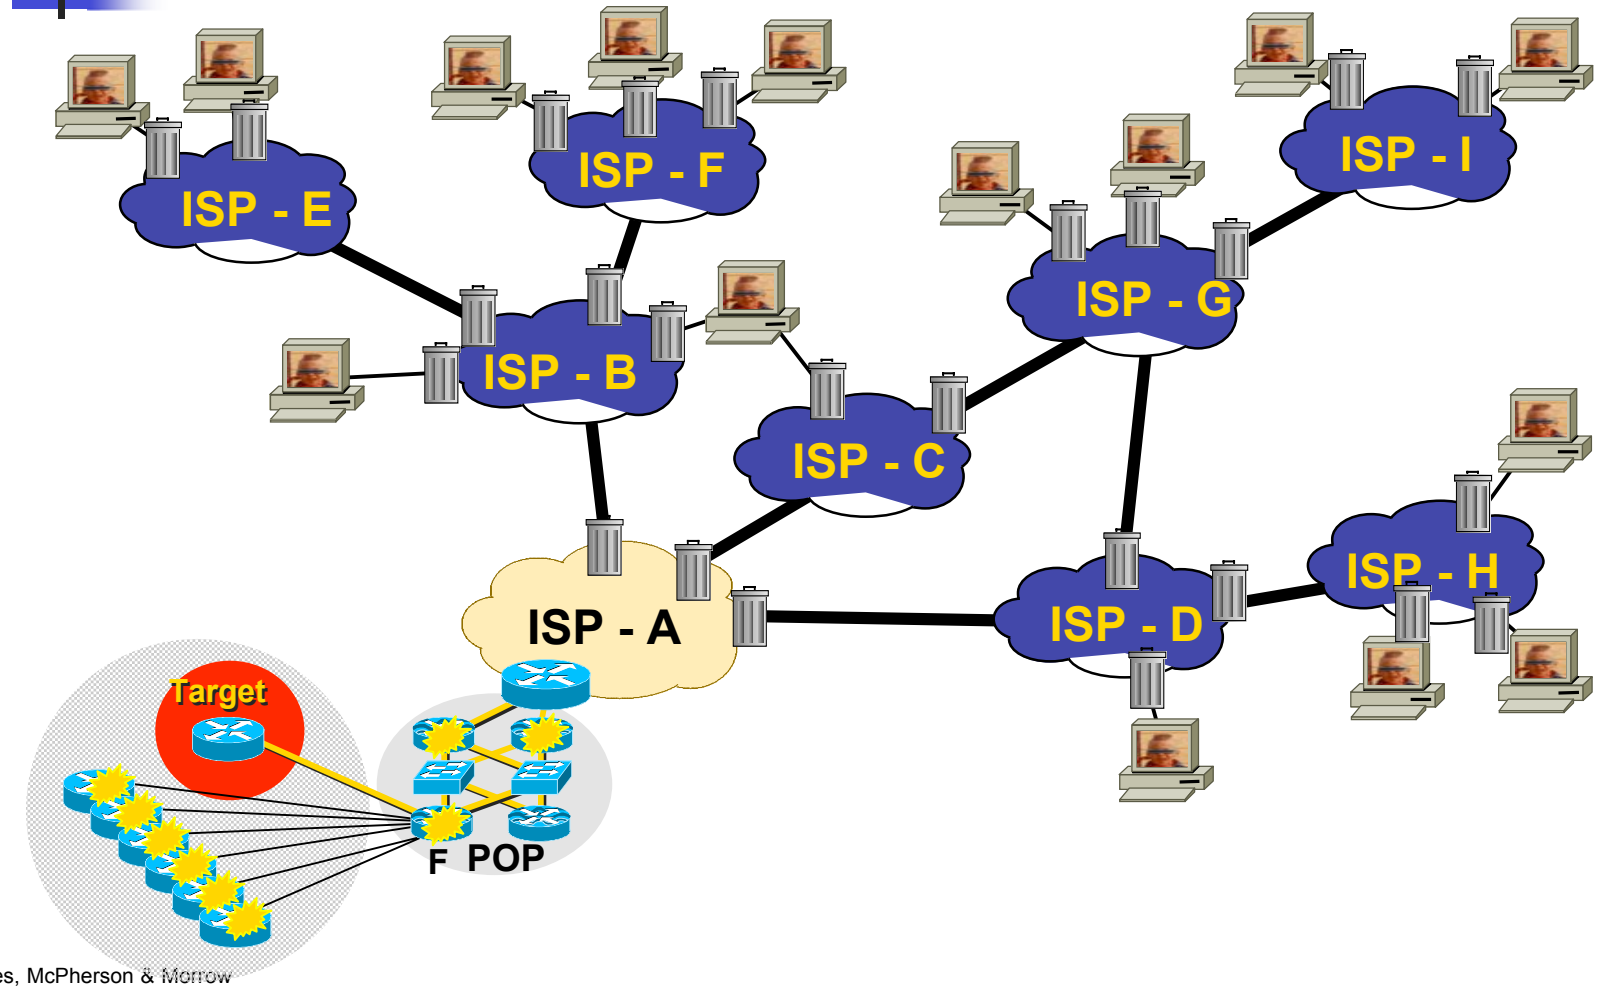

Gotchas with Black Hole Filtering

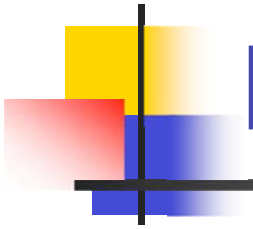

- Routers were designed to forward traffic, not drop traffic.
- ASIC Based Forwarding can drop traffic at line rate.
- Processor Based Forwarding can have problems dropping large amounts of data, especially architectures that require exception path punts for dropped packets.
- BGP RIB and subsequent FIB entries utilize CPU and memory resources and should be tracked.
- Remember the old shunt technique ….

Gotchas with Black Hole Filtering

- Back in the days when this was in the core of the Internet …..

- All “drops” to Null0 were process switched.
- Fast Drops fixed the problem for a while, but traffic loads increased to where they could not drop at line rate anymore.
- Bottom-line – Software based forwarding routers (any vendor) can forward faster then they can drop.

uRPF & Source-based Blackholes

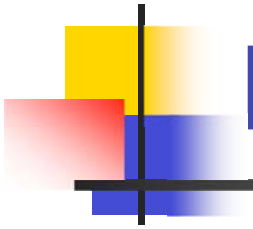

- Source-based blackholes are achievable as well, though likely don't make sense on the customer-facing front.

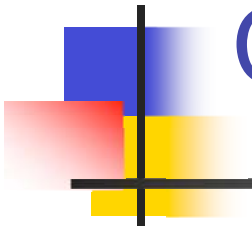

Customer-Triggered Blackholes

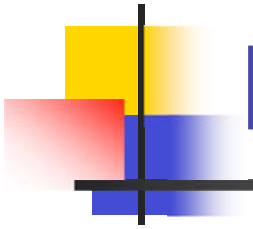

Deploy BGP Policy Set

- Accept more-specifics of customer routes with destination-based BGP blackholing community attached.
- No source-based blackholing
- Only accept more-specifics of customer prefixes

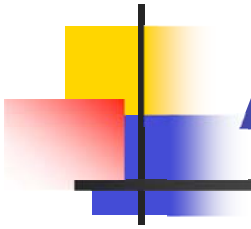

Accepting Longer Prefixes

- Only accept more-specifics of customer-allocated/advertised space.
- Policy depends on ingress prefix filtering policies
  - Explicit filters and any mask-length filters require preceding more-specific & community colored route-acceptance
  - Looser policies are perhaps less work but leave more room for errors
  - Define prefix-length acceptance criteria

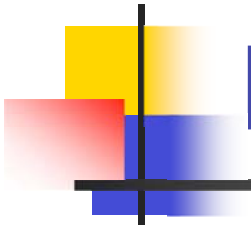

More on customer-triggered..

- MTTR decrease
- Customer driven, removes some liability
- Customer:
  - When you want
  - Where you want it
  - Your timeline, not the ISPs!
- Tag received routes with NO_EXPORT community (and likely, NO_ADVERTISE, though a direct BGP session with the peer is then required)
- Policies and announcement authority should be verified regularly, exception reporting should be automated

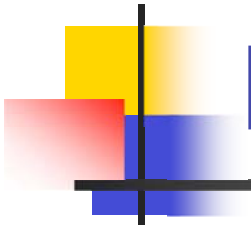

Enhanced Policy Language

- Specifies explicit prefix filters with exception policy that matching defined communities for blackhole or other.
- Complements explicit filtering without adding twice the configuration overhead to introduce acceptance of more-specifics for blackholing.

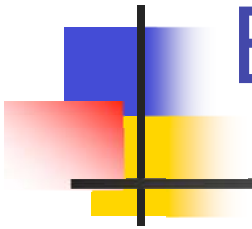

BGP Flow Specification

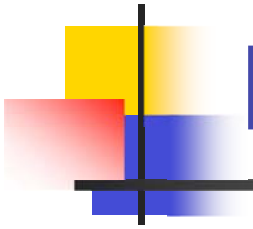

Draft Information

- Available at:
  - http://www.tcb.net/draft-marques-idr-flow-spec-00.txt
  - Currently expired from IETF Internet-Drafts directory, hope to post new version soon.
- Authors:
  - Jared Mauch
  - Danny McPherson
  - Robert Raszuk
  - Pedro Marques
  - Nischal Sheth

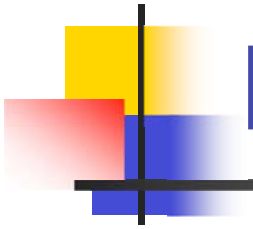

Draft Overview

- Specifies procedures for the distribution of flow specification rules via BGP.
- Defines application for the purpose of packet filtering [other] in order to mitigate (distributed) denial of service attacks
- Defines procedure to encode flow specification rules as BGP NLRI which can be used in any why the implementer desires.

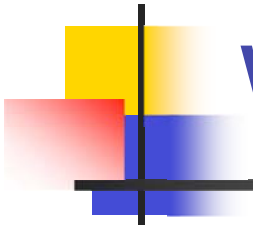

What's A Flow Specification?

- A flow specification is an n-tuple consisting of several matching criteria that can be applied to IP packet data.
- May or May not include reachability information (e.g., NEXT_HOP).
- Well-known or AS-specific COMMUNITIES can be used to encode/trigger a pre-defined set of actions (e.g., blackhole, PBR, rate-limit, divert, etc..)
- Application is identified by a specific (AFI, SAFI) pair and corresponds to a distinct set of RIBs.
- BGP itself treats the NLRI as an opaque key to an entry in its database.

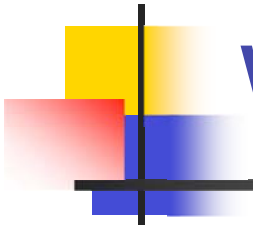

What's it for?

- Primarily: DDOS Mitigation
- Continue evolution from:
  - Destination-based blackhole routing
  - uRPF/source-based BGP blackhole routing
- To:
  - Much more precise mechanism that contains all the benefits of it's predecessors

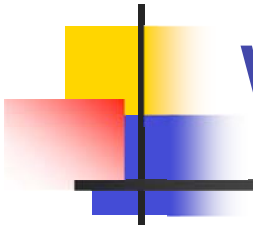

We Need Operator Feedback

- Is this useful?
- What's missing (e.g., more flexible specification language)
- Does this belong in BGP?
- What are our alternatives?
- Comments to authors are welcome!
  - flow-spec@tcb.net

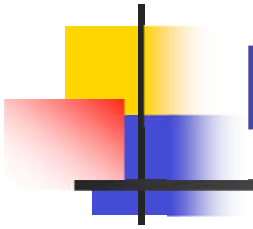

References

- [backscatter]
- [RFC 3330]
- ftp://ftpeng.cisco.com/cons/isp/security
- Other?

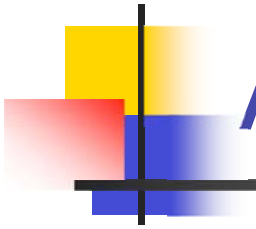

Acknowledgements

- Barry Greene
- Brian Gemberling

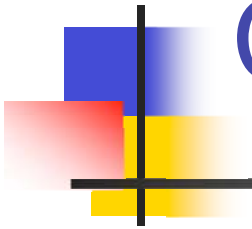

Comments/Questions/Other?