BGP: Good MEDs Gone Bad!

NANOG 29

Danny McPherson

danny@arbor.net

Agenda

- Goals
- Potato Terminology
- What are MEDs?
- Where MEDs Make Sense
- MED Deployment Considerations
- Conclusion

Goals of this Talk

- Increase awareness of MED deployment considerations
- Increase awareness of MED-related protocol constraints
- Encourage operators to better understand their vendor(s) MED-related implementation
- Nothing new or Earth-shattering here…

Before We Begin…

- How many folks here know exactly what your MED policy is?
- How many folks here accept MEDs from customers? By default?
- How many folks here accept MEDs from peers? By default?
- How many folks here don't know?

Potato Terminology

- Hot Potato == *Closest-Exit Routing*; default shortest path routing
- Cold Potato == *Best-Exit Routing*; shortest hops, reflect IGP topology, route around congestion, marketing, other..
- Mashed Potato == “Less than Ideal” Routing; unintentional, often results from intended Best-Exit Routing

What Are MEDs?

- BGP MULTI_EXIT_DISC (MED), formerly known as INTER_AS_METRIC
- Optional non-transitive BGP attribute used to discriminate among multiple exit or entry points into *the same* neighboring AS
- All preceding selection criteria being equal, prefer path with LOWEST MED.

Where MEDs Make Sense

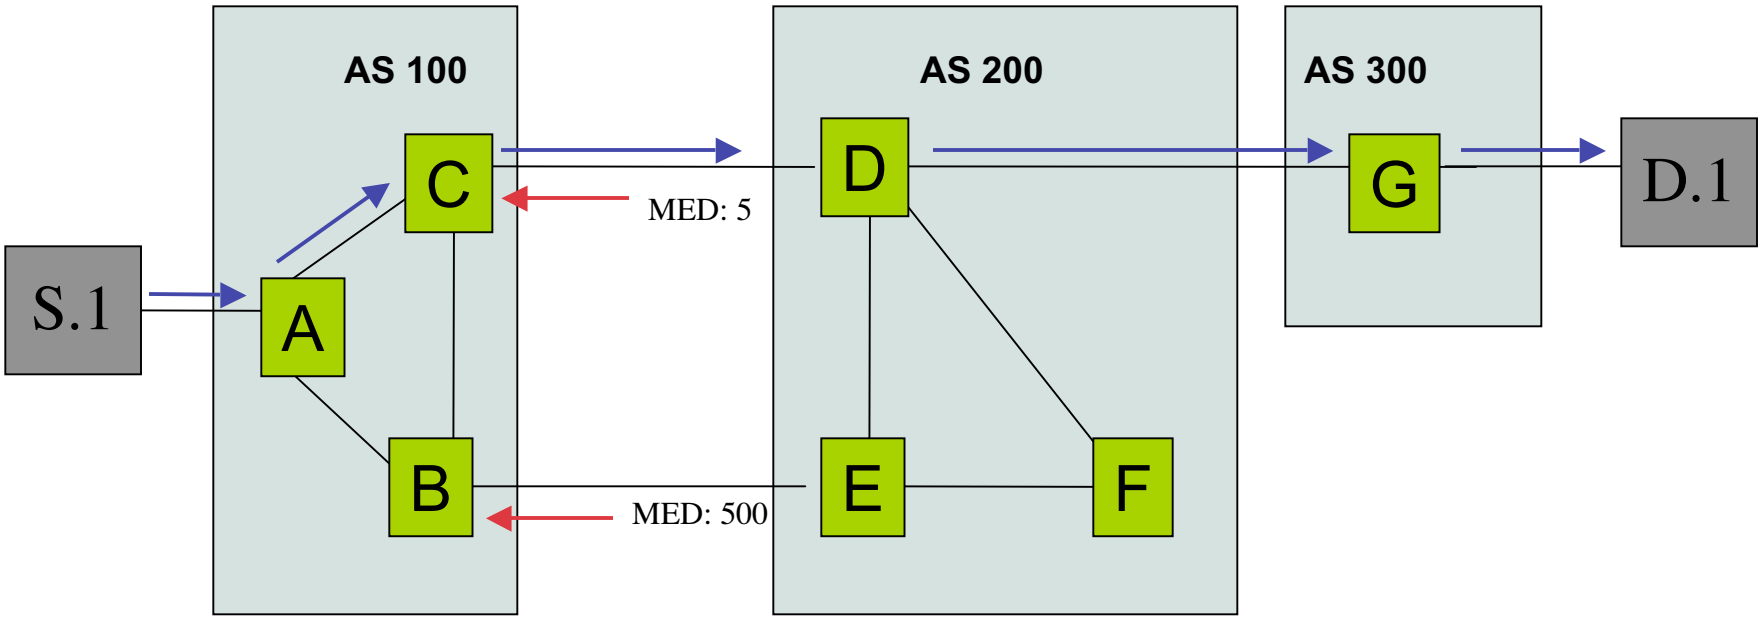

- Preferred S.1 --> D.1 path is A->C->D->G per advertised MEDs
- If MEDs weren't advertised AS100 would have no way to know that AS 200's D is optimal path

MED Deployment Considerations

- MEDs Break With Aggregation
- Inconsistent Vendor Behavior
- Persistent Route Oscillation Condition
- Route Flap Dampening and MED Churn
- Comparing Between Different Autonomous Systems
- Security Considerations
- BGP Update Packing

MEDs & Aggregation

- Aggregates are often generated from multiple locations within an AS
- When MEDs are derived from IGP metrics associated with said aggregates VERY sub-optimal routing may result

MEDs & Aggregation (cont.)

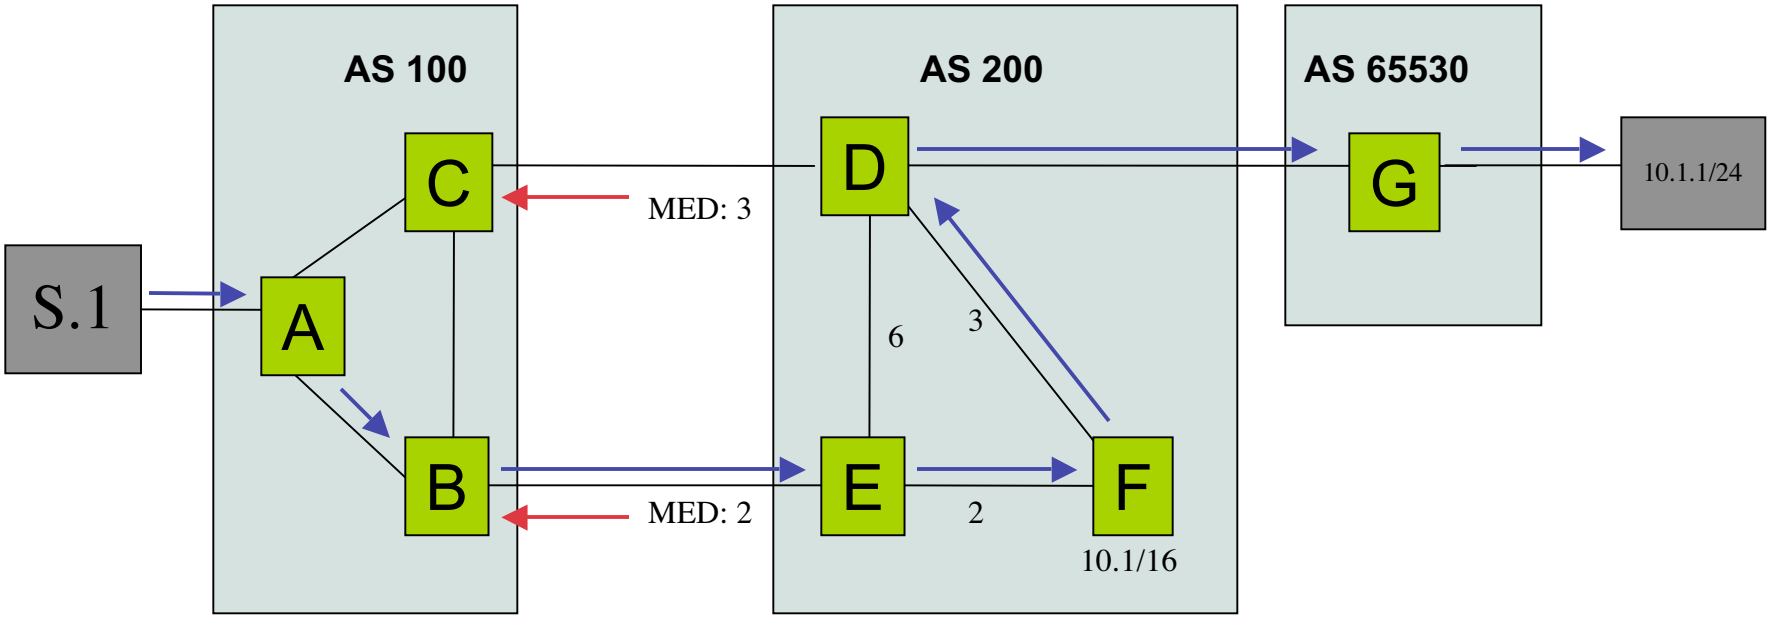

- Only 10.1/16 aggregate is advertised to AS 100. MEDs are derived from IGP metrics associated with aggregate source router F as BGP NEXT_HOP.
- Preferred S.1 --> D.1 path is A->B->E->F->D->G per advertised MEDs. AS 200 more-specific makes no difference.

Inconsistent Vendor Behavior

- Does your router vendor:
  - advertise MEDs to IBGP peers as a default behavior?
  - advertise MEDs to EBGP peers as a default behavior?
  - advertise MEDs to confederation peers by default?
  - compare MEDs between confederation peers and EBGP peers?
  - prefer no MED over MED of zero over …?
  - consider max MED (2^32-1) as unfeasible?
  - compare MEDs between different autonomous systems by default?
  - impose temporal route selection behavior to MEDs?

Persistent Route Oscillation

- MEDs are primary trigger for persistent route oscillation
- See RFC 3345 for details
- Alternatively, see Daniel Walton's FEB '01 talk on this topic.

Flap Dampening & MED Churn

- MEDs are often derived from IGP metrics (generally, this is a good idea to ensure BGP path selection is aligned with IGP)
- However, it means that IGP instabilities within an AS, or on even a single link, result in BGP route updates/withdraws
- Results in significant churn; may result in routes suppression. Transit AS IGP instabilities affect downstream prefixes.
- Some implementations do [arguably] clever things in a attempt to scope such behaviors -- Does your vendor? Have you disabled it?

Flap Dampening & MED Churn (cont.)

Origins of Internet Routing Instability (1999)

Craig Labovitz, G. Robert Malan, Farnam Jahanian

http://citeseer.nj.nec.com/labovitz99origins.html

Comparing MEDs Between Different Autonomous Systems

- MEDs values are derived from many different policies:
  - Static/Explicit
  - IGP Metrics:
    - Additive or local?
    - Do your peers use the same IGP? Is the available metric space the same?
  - Are your peers aware they're sending MEDs?
  - Are they sending MEDs at all?

Security Considerations

- MEDs may be used to manipulate a peer's route selection criteria in order to gain some advantage over that peer, usually via traffic diversion
- Do you accept MEDs from peers (or customers) as a default behavior?
- Use your imagination…

MEDs & Update Packing

- BGP Update packing allows prefixes with like attribute sets to be packed into a single update message.
- Provides an array of benefits!
- Lots of [potentially useless] MEDs lessens the benefits update packing provides.

Conclusions

- MEDs work in lots of places
- MEDs break in lots of places
- You should be aware of the difference!
- draft-mcpherson-grow-bgp-med-considerations-00.txt will be posted to internet-drafts soon and provides more detailed discussions of this topic.

Acknowledgements…

- Craig Labovitz
- John G. Scudder
- Vijay Gill

Thank You!