NOCTION

You control the announcement, not the decision

Practical Inbound BGP engineering

Today’s Agenda

01 Outbound vs Inbound

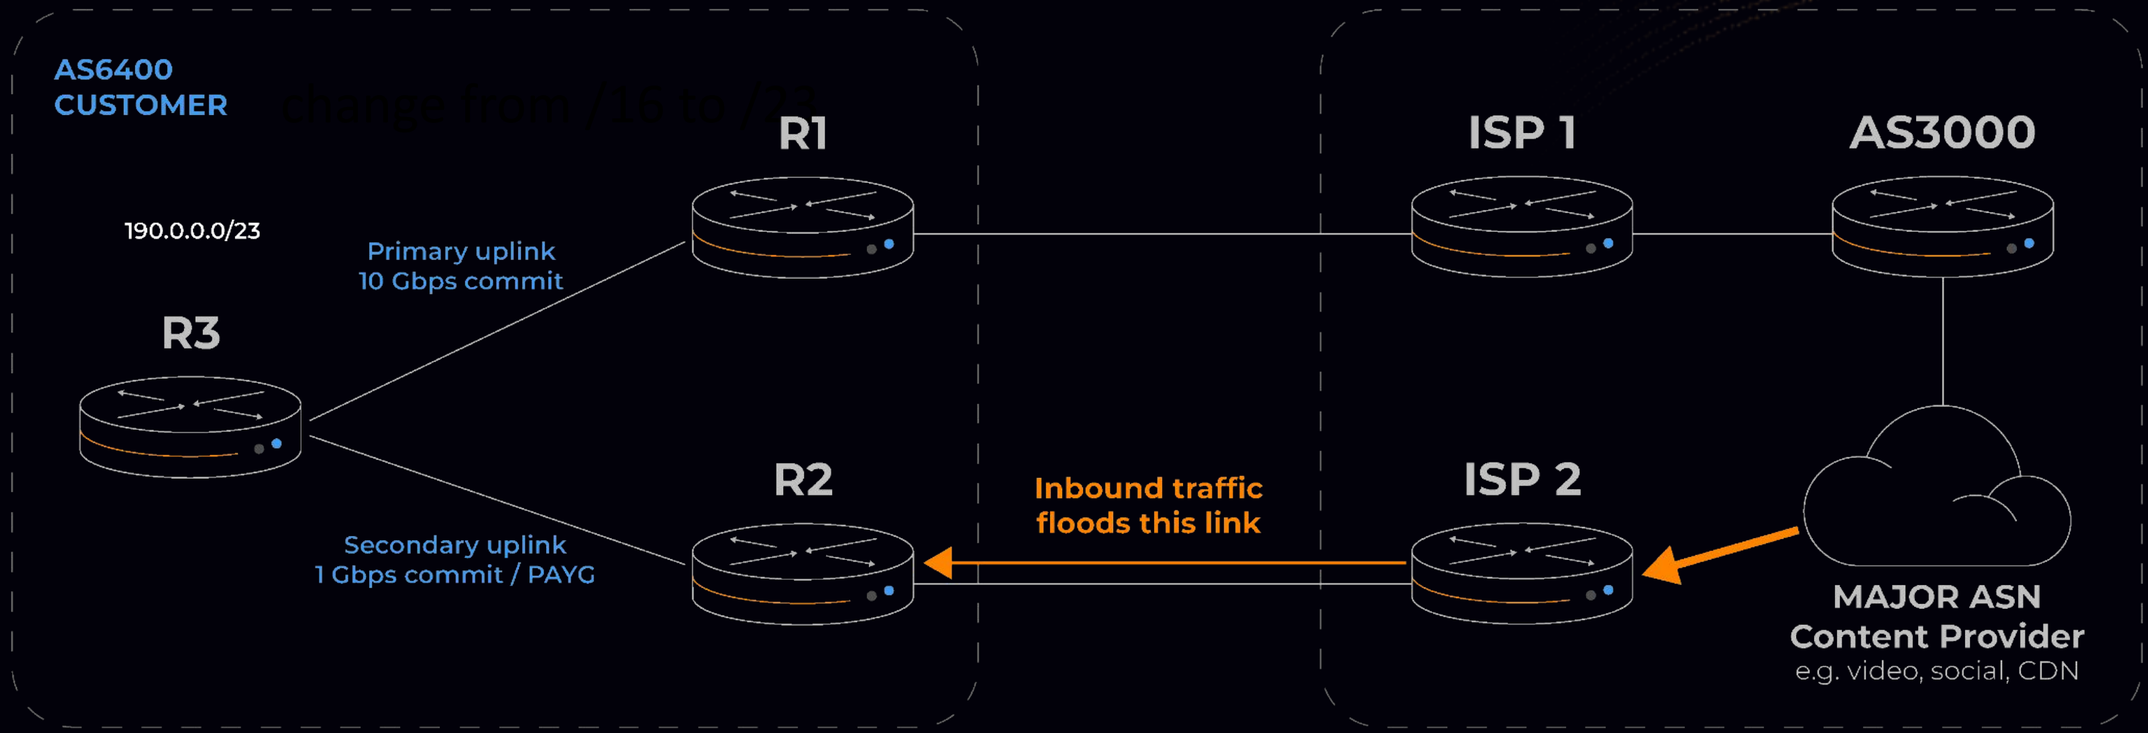
02 Inbound bandwidth control

03 Inbound performance optimization

04 Verification, guardrails and automation

05 Questions & Answers

The Inbound Challenge



Result: 95th-percentile billing overage

Inbound traffic prefers the secondary path due to normal BGP path selection.

Tactic 1 - AS_PATH Prepending (*The Sledgehammer*)



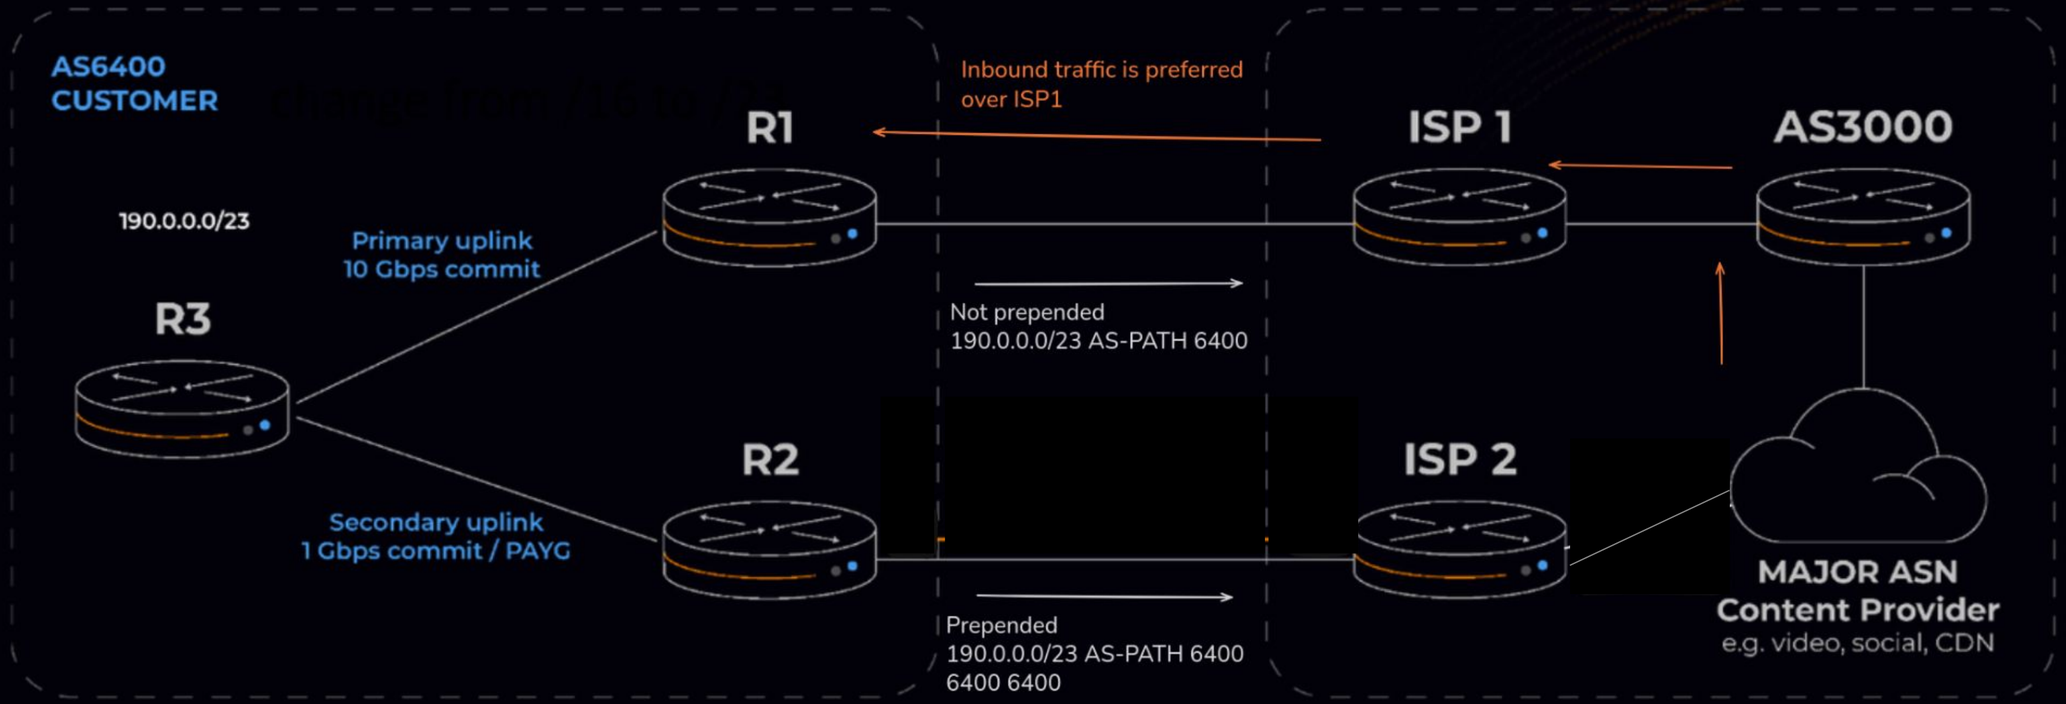
Prefix-set CUST-PS that is matching the prefix 190.0.0.0/16.

```
prefix-set CUST-PS
  190.0.0.0/16
 end-set
```

Route-policy AS-PATH-RP matching the prefix-set CUST-PS.

```
route-policy AS-PATH-RPL
  if destination in CUST-PS then
    prepend as-path 6400 2
  else
    pass
  endif
 end-policy
```

Inbound control with AS_PATH Prepending

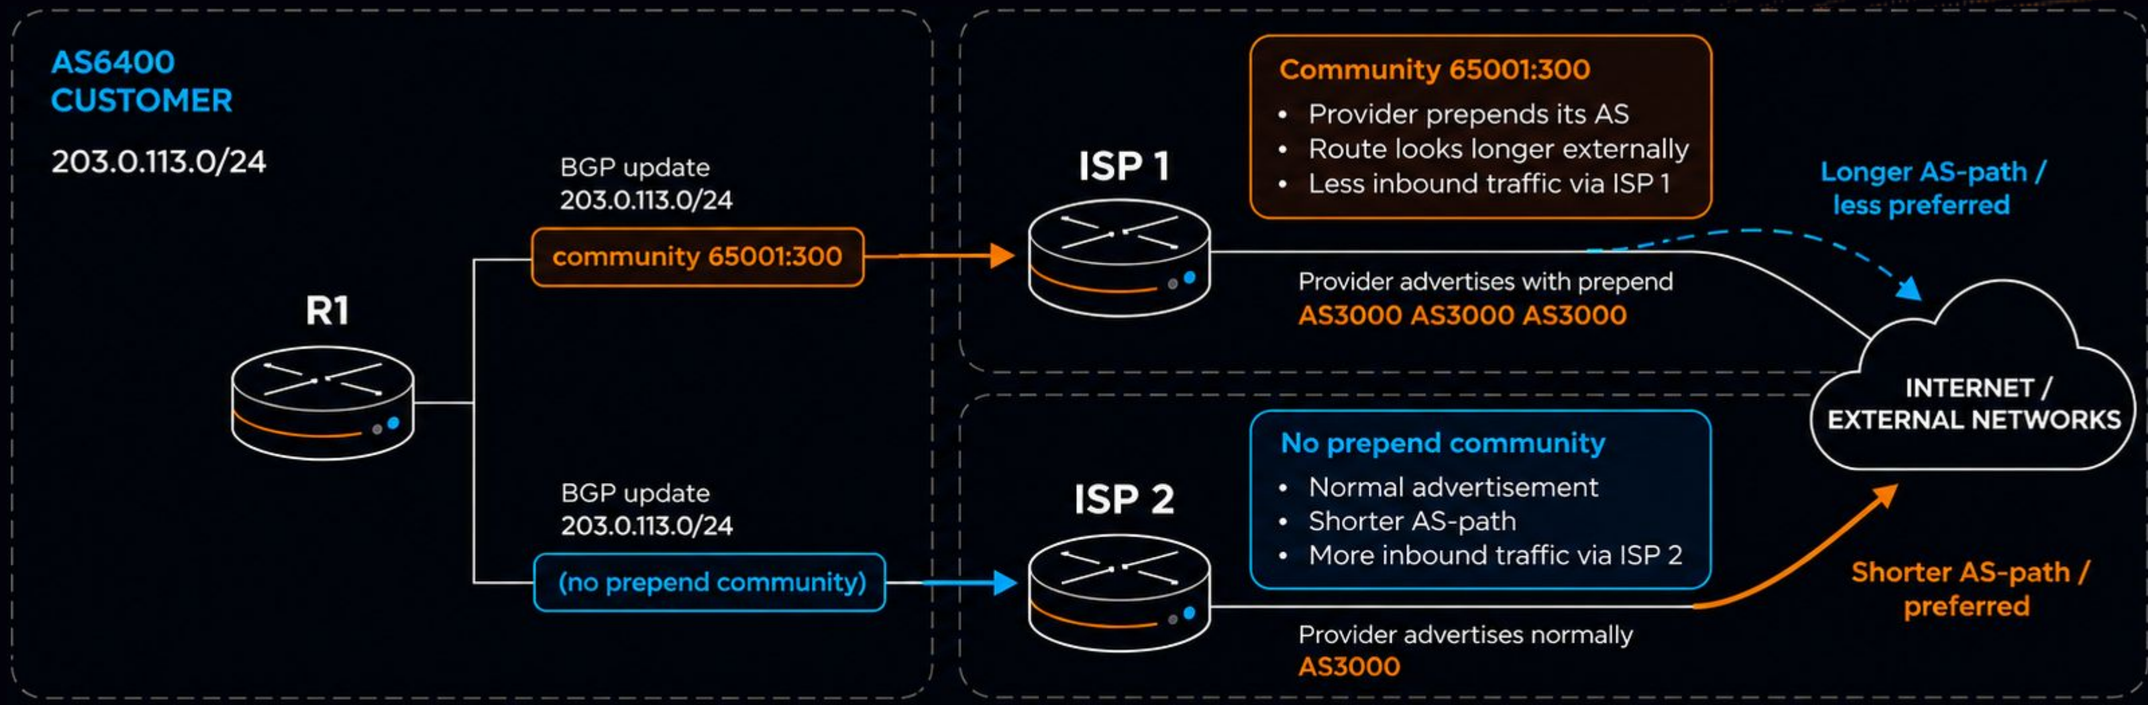


Tactic 2 - BGP Communities (*The Scalpel*)



Community Type	Typical Effect	Example
Lower Local Preference	Makes your route less preferred inside Provider A, potentially causing traffic to reach you via Provider B.	NTT: 2914:450 → Local-Pref 96 (fallback) / Arelion: 1299:50 → Local-Pref 50 (backup)
Prepend to Selected Peer	Reduces traffic arriving through Provider A from a specific external network	NTT: 65501:<peer-AS> / 65502:<peer-AS> / 65503:<peer-AS> → prepend 1× / 2× / 3× toward that peer

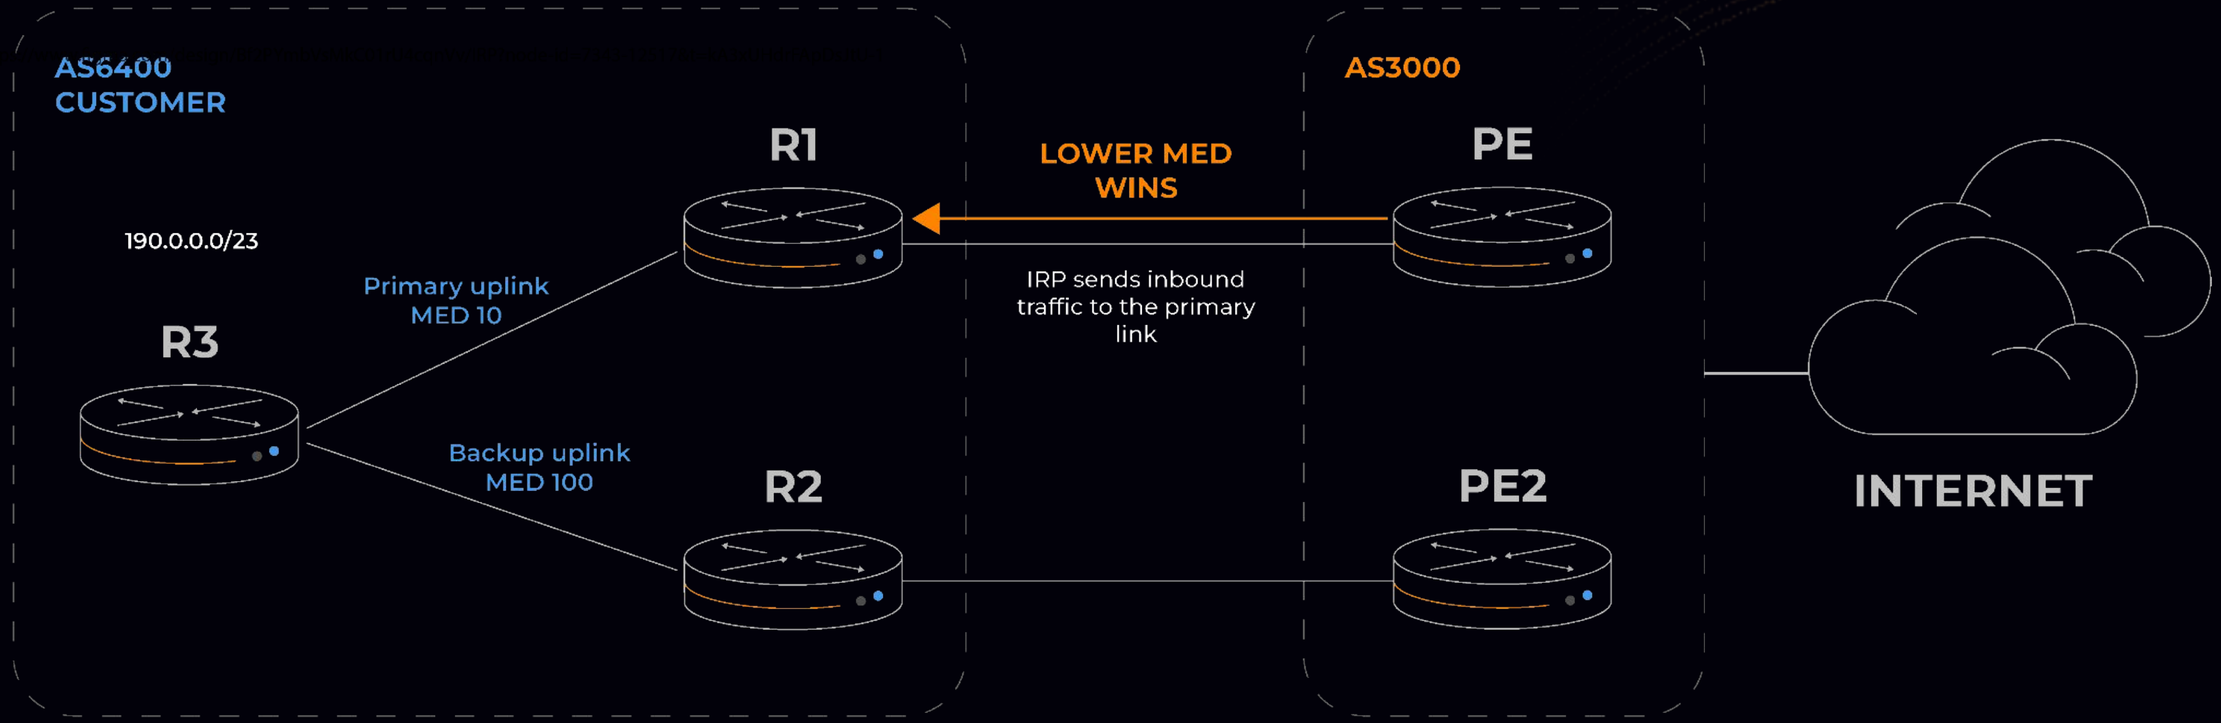
Inbound control with BGP communities



Result: Shift inbound traffic away from ISP 1

The customer tags the route with a BGP community that tells ISP 1 to prepend its ASN when advertising the prefix externally. Because the route through ISP 1 appears longer, many external networks prefer the shorter path via ISP 2.

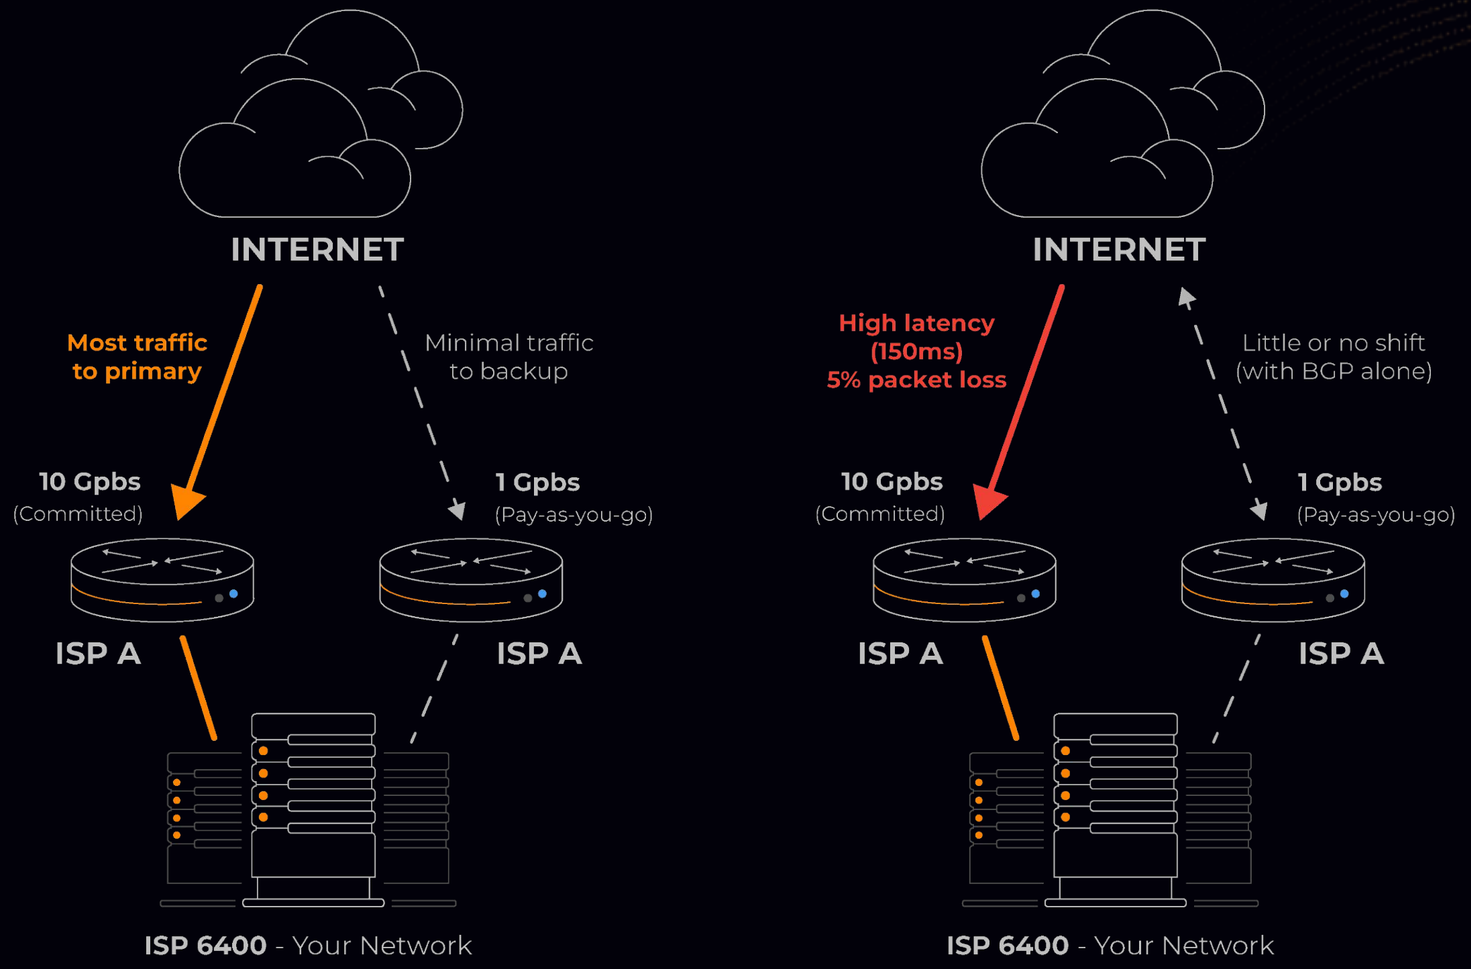
Tactic 3 - Multi-Exit Discriminator (*MED*)



Result: MED steers traffic only within the same provider AS

Use MED for primary/backup control with one ISP.
It normally does not influence traffic between different ISPs.

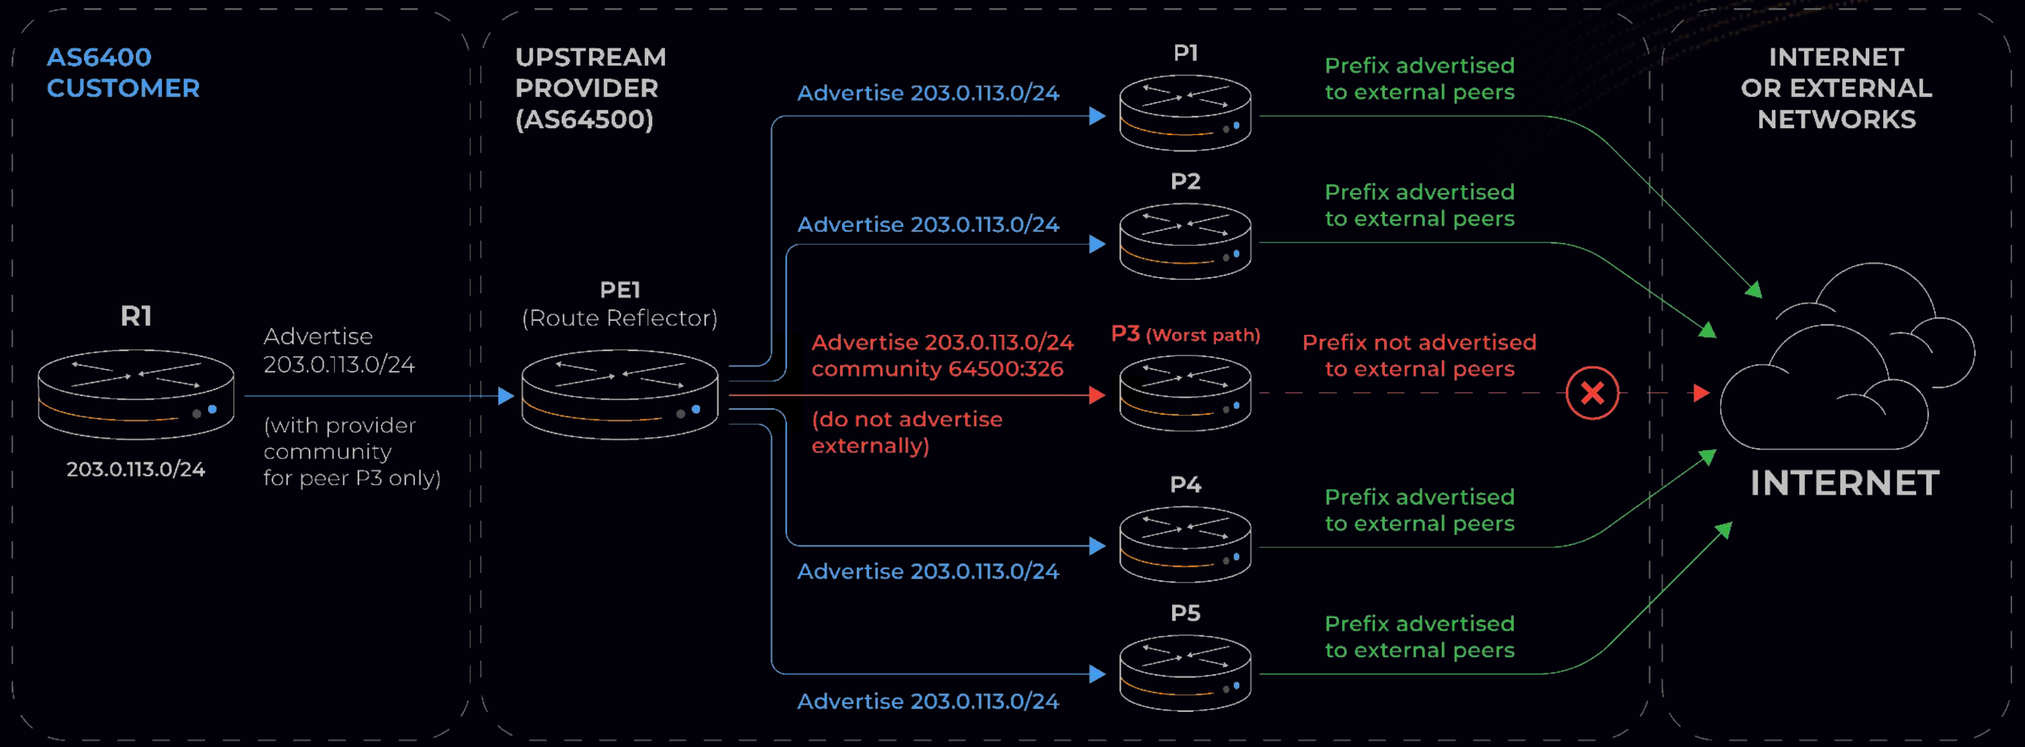
Shifting Gears - From Cost to Performance



Manual Detection of Inbound Path Degradation

Diagnostic Layer	Tools & Commands	Target Metrics / Red Flags
1. External Inbound Probing	• RIPE Atlas / Tier-1 Looking Glasses • mtr --tcp -P 443 <IP> • Public Cloud Instances (AWS/GCP)	• High packet loss on incoming hops • Latency spikes along shorter AS_PATH paths
2. Edge-Router Diagnostics	• TWAMP (RFC 5357) • Cisco IP SLA / Juniper RPM • show ip sla statistics	• Hardware-timestamped RTT variations • Packet loss (>1%) toward upstream ISP PE interface
3. Passive Flow & TCP Health	• NetFlow v9 / IPFIX / sFlow • tcpdump / Wireshark • Flow Collectors	• Spikes in TCP Retransmissions & Duplicate ACKs • Degradation in TCP SYN-to-ACK handshake deltas
4. Interface Physical Layer	• Digital Optical Monitoring (DOM) • FIFO Buffer & Queue Counters • show controllers optics	• CRC / Frame errors (dirty fiber or failing SFPs) • Micro-burst FIFO overruns and queue drops

Deflecting a poor performing path inside a single provider

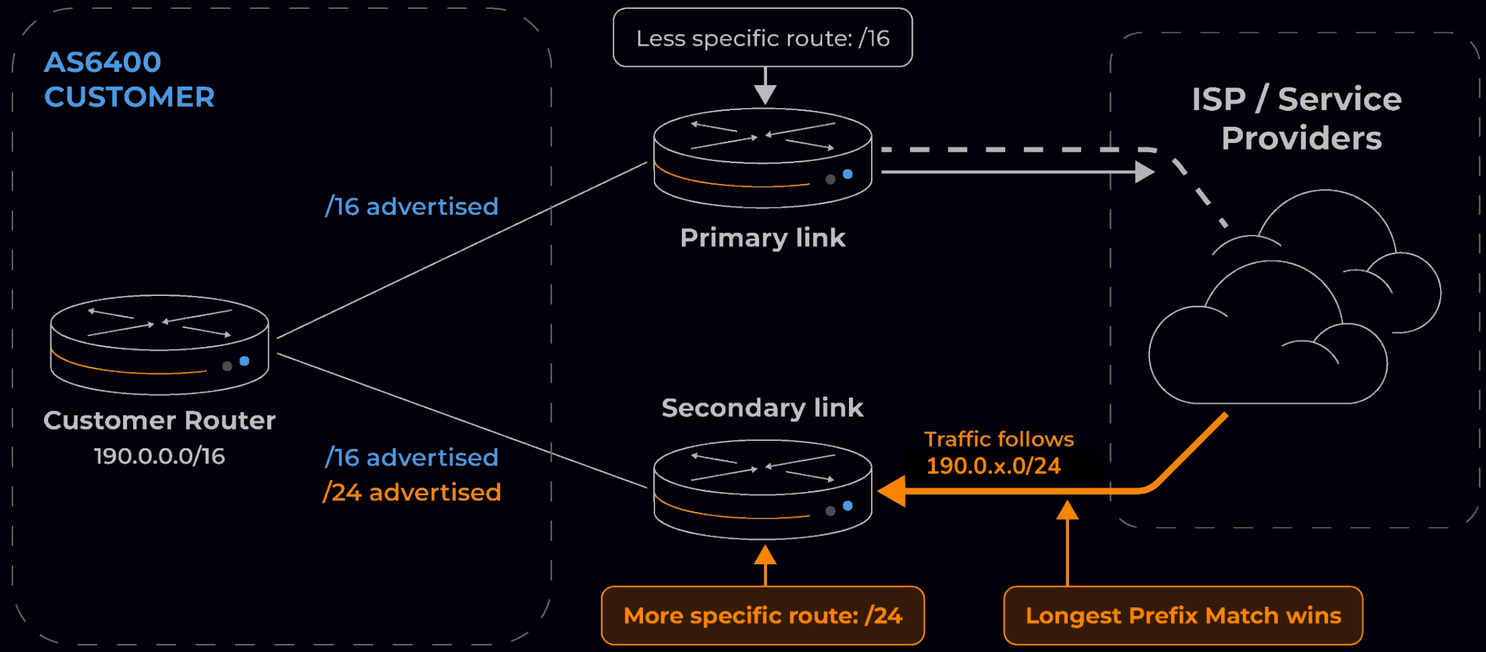


Result: inbound traffic is deflected from the worst-performing peer inside the provider.

The provider does not advertise the prefix to the selected external peer based on the received BGP community, while other peers continue to receive and advertise it.

The community value and action (e.g. no-advertise, prepend, lower local-preference) are provider-specific. Only the selected peer is affected; the prefix remains advertised to other peers and continues to reach the Internet through them.

The Ultimate Trump Card - Longest Prefix Match





Summary & Best Practices

- For Bandwidth Commit Control: Don't rely solely on AS_PATH prepending. Talk to your ISPs, get their BGP Community guides, and use Local Preference tagging.
- For Multi-Homed Single ISPs: Use MED to dictate your preferred entry points.
- For Performance Optimization: Automate! BGP was built for reachability, not performance. If inbound performance is critical to your revenue, look into the automated traffic manipulation platform that make these BGP attribute changes dynamically based on telemetry and active probing.

Questions & Answers

