BGP Vortex

Update Message Floods Can Create Internet Instabilities

Felix Stöger, Henry Birge-Lee, Giacomo Giuliari, Jordi Subirà-Nieto, Adrian Perrig

ETH zürich PRINCETON UNIVERSITY MystenLabs

Internet Requires Consistent and Stable Routing

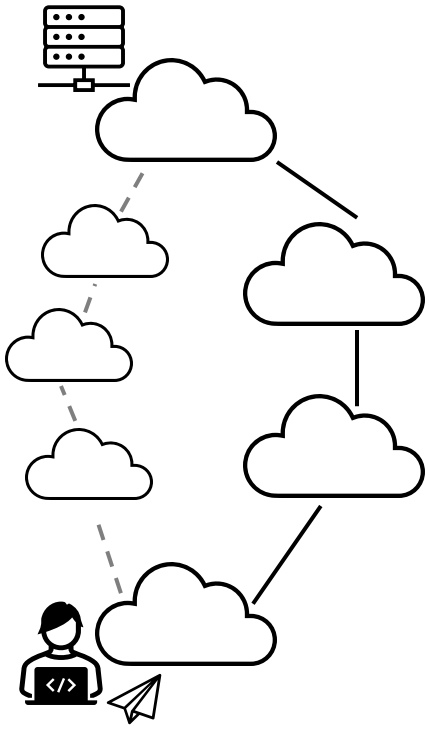

Internet Requires Consistent and Stable Routing

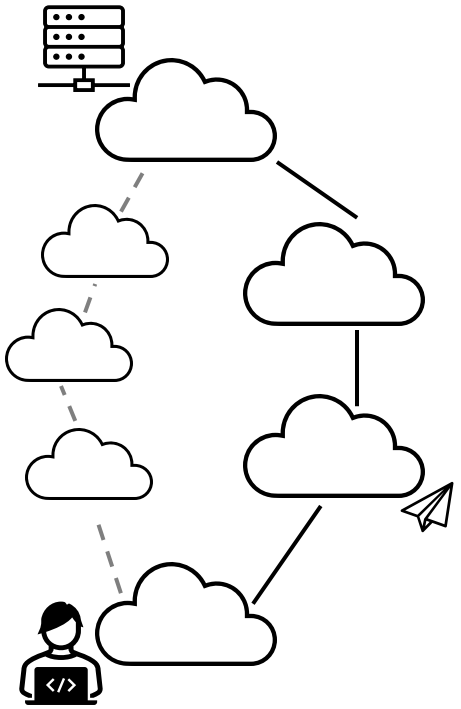

Internet Requires Consistent and Stable Routing

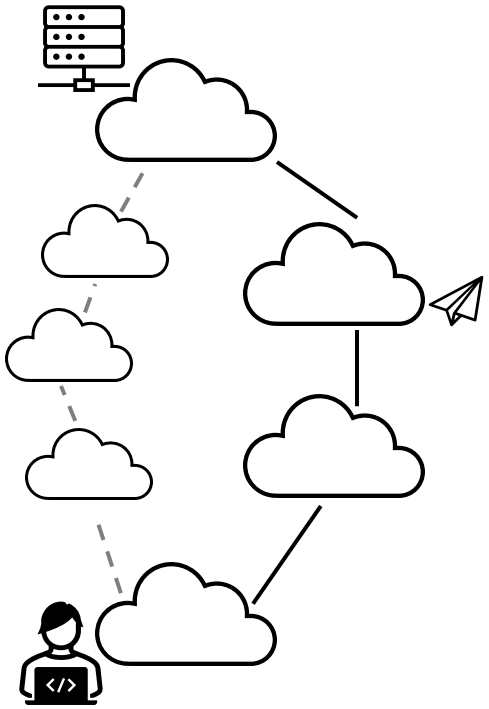

Internet Requires Consistent and Stable Routing

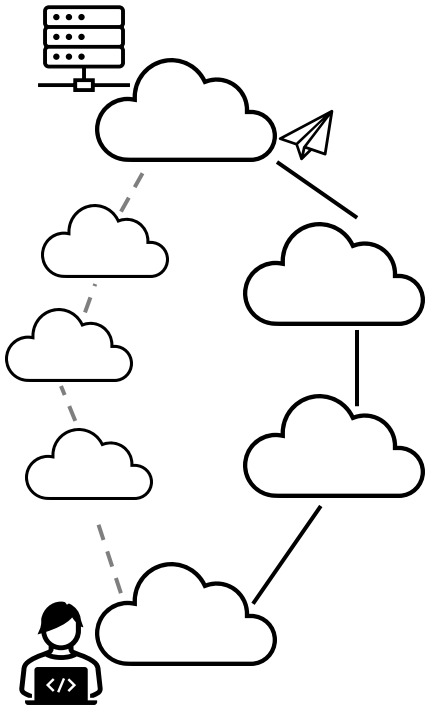

Internet Requires Consistent and Stable Routing

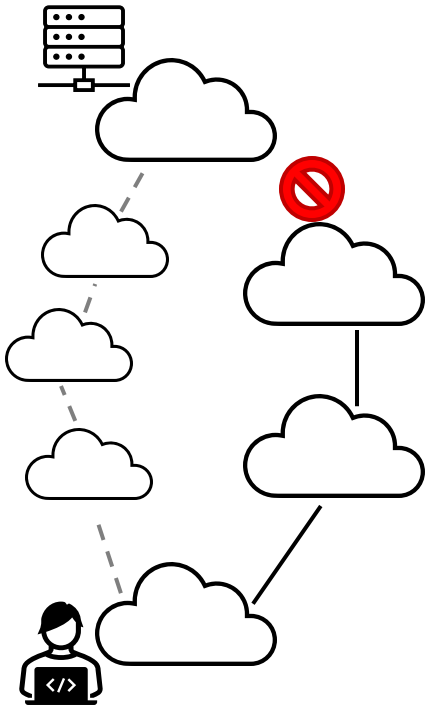

Internet Requires Consistent and Stable Routing

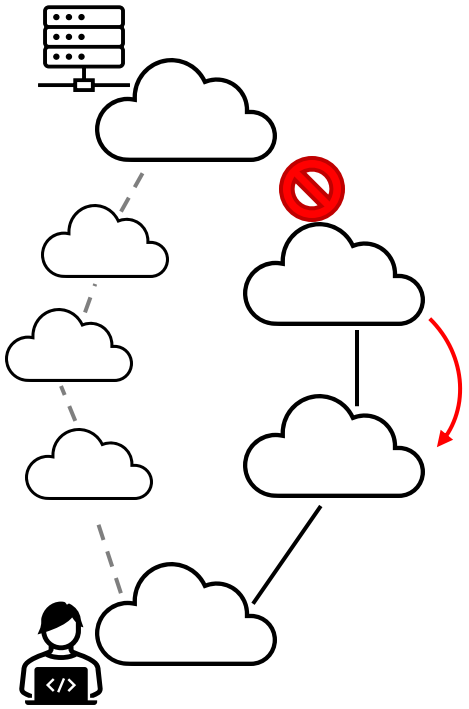

Internet Requires Consistent and Stable Routing

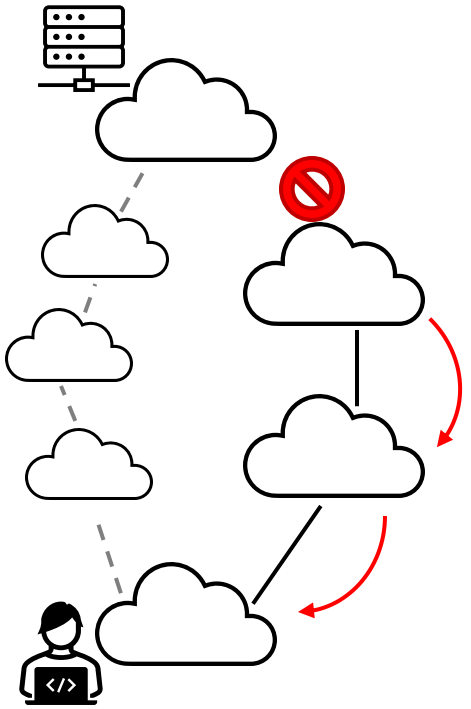

Internet Requires Consistent and Stable Routing

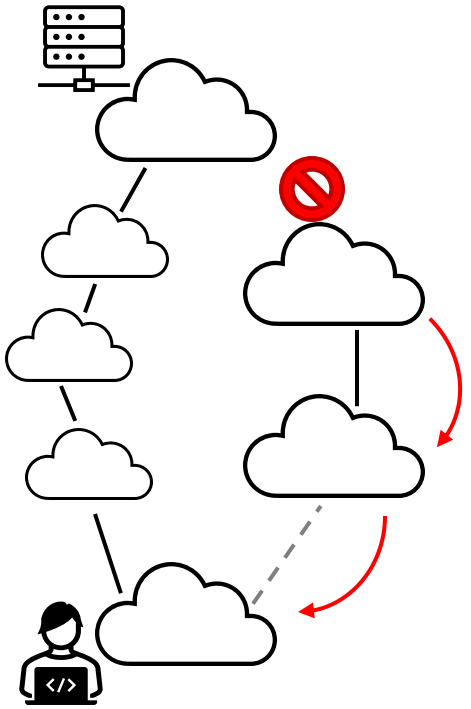

Internet Requires Consistent and Stable Routing

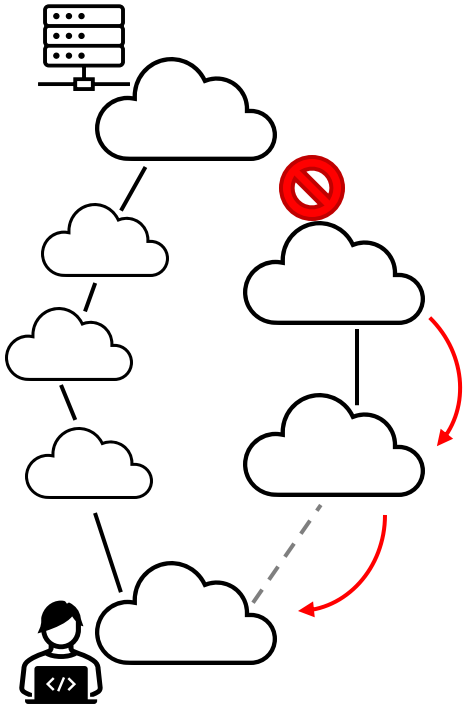

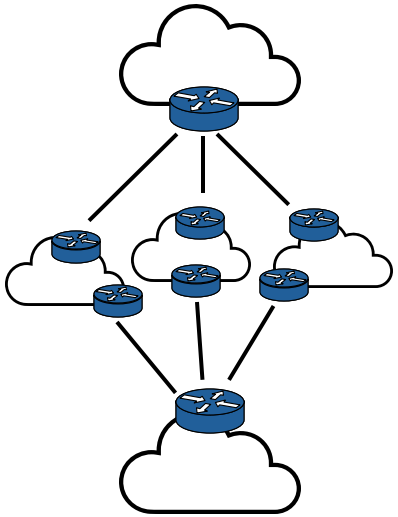

Internet Requires Consistent and Stable Routing

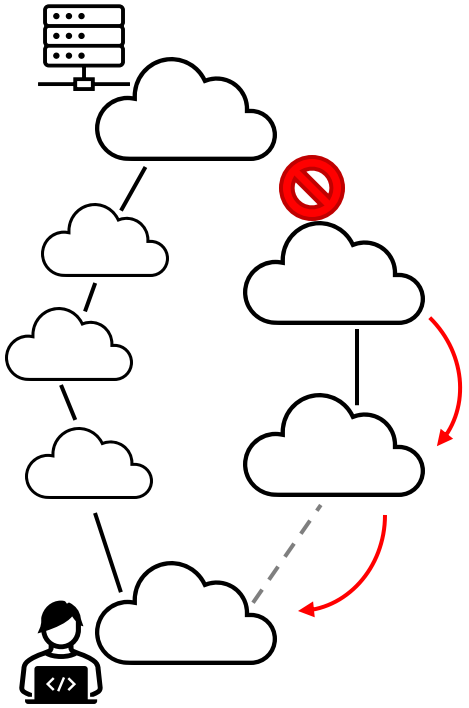

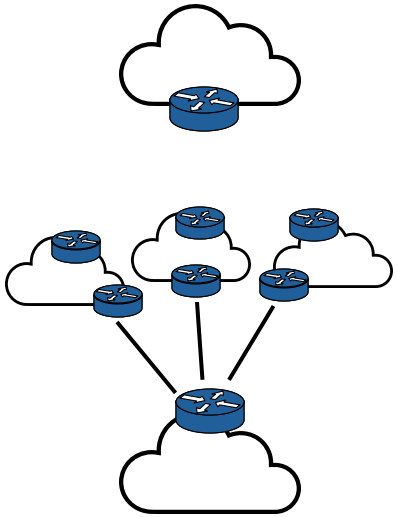

Internet Requires Consistent and Stable Routing

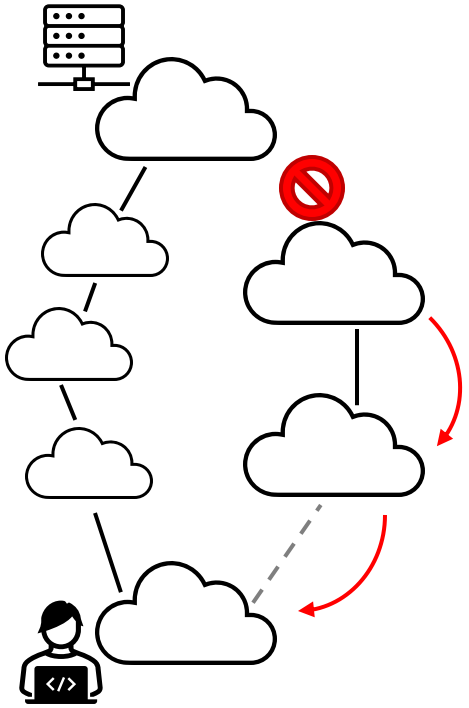

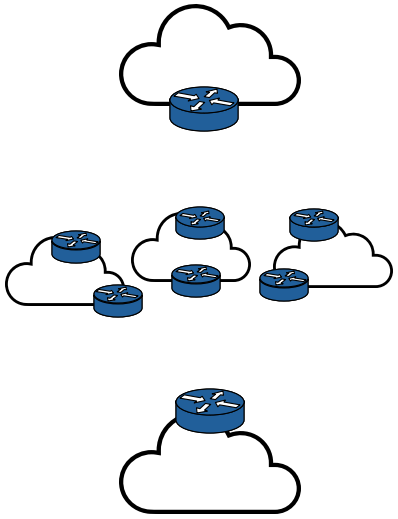

Internet Requires Consistent and Stable Routing

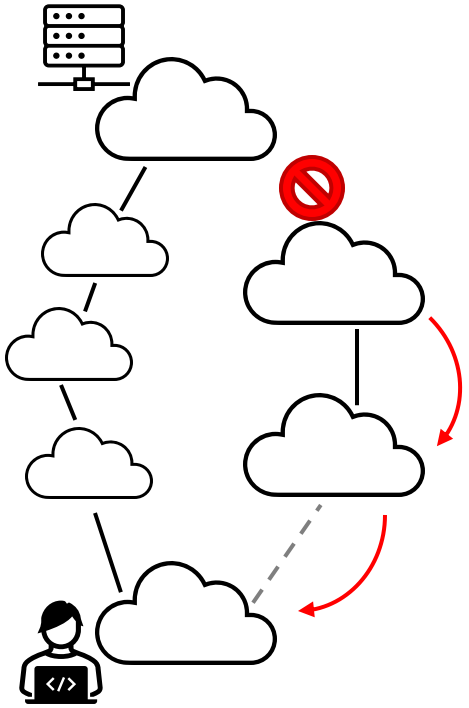

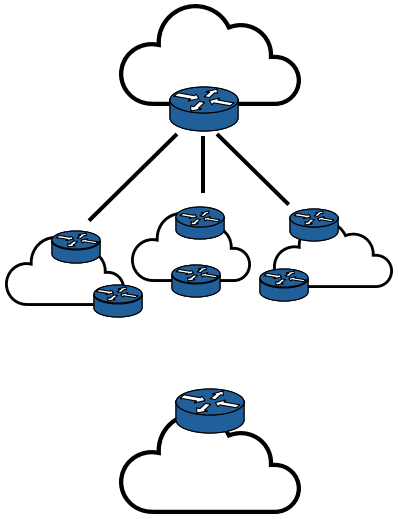

Internet Requires Consistent and Stable Routing

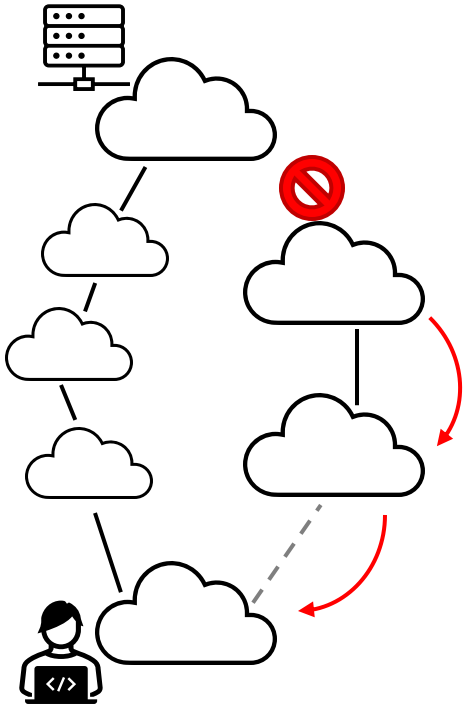

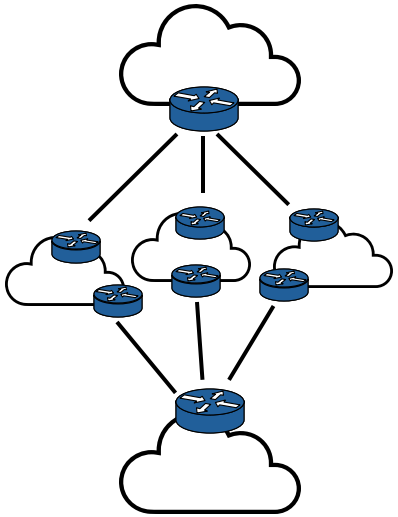

Internet Requires Consistent and Stable Routing

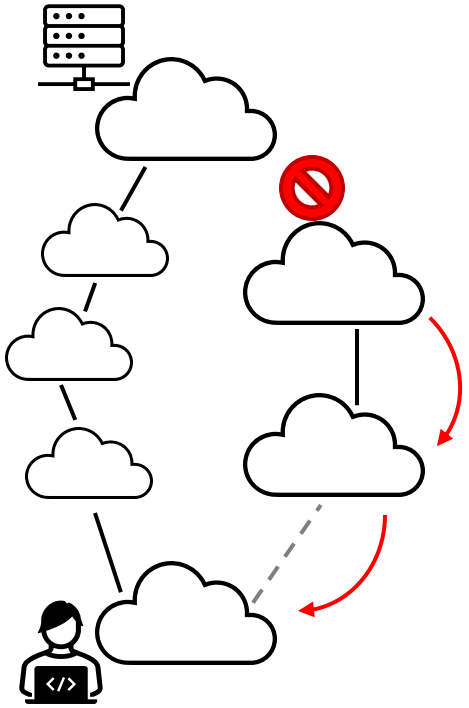

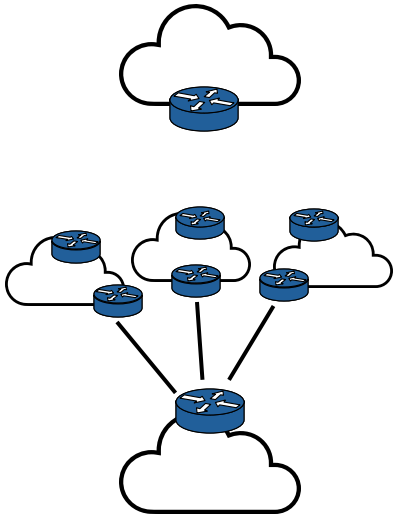

Internet Requires Consistent and Stable Routing

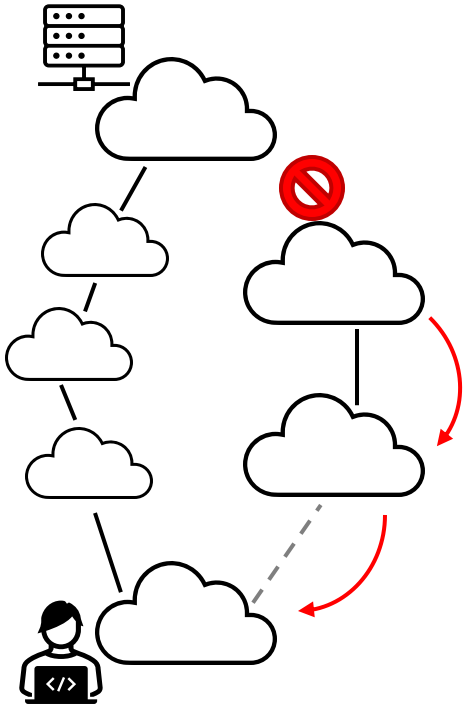

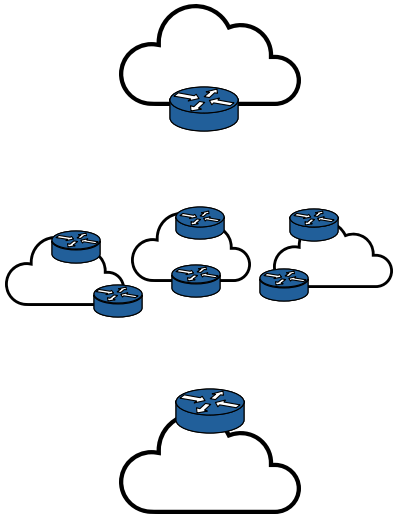

Internet Requires Consistent and Stable Routing

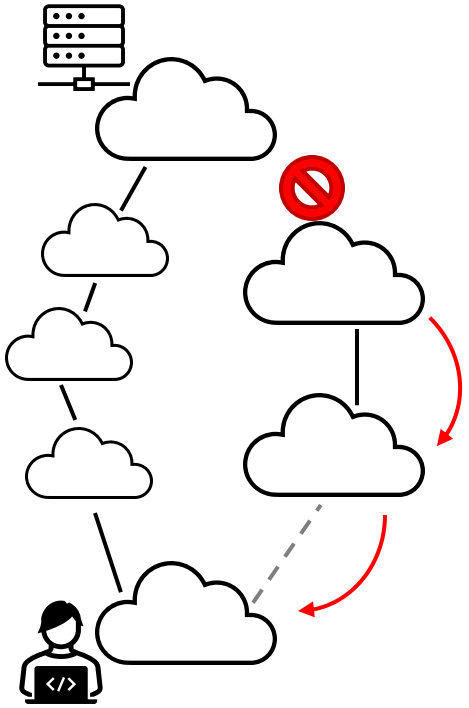

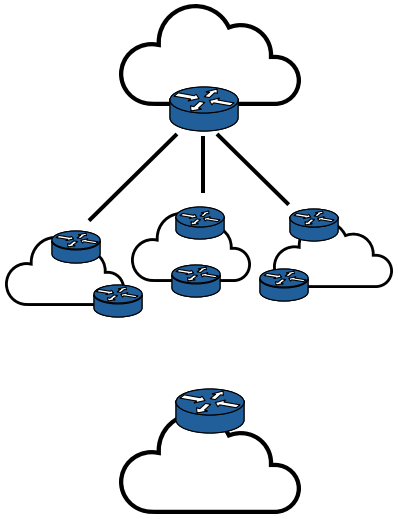

Internet Requires Consistent and Stable Routing

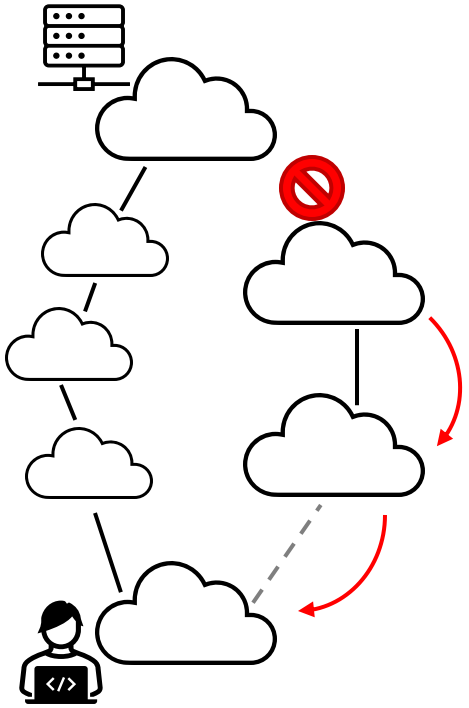

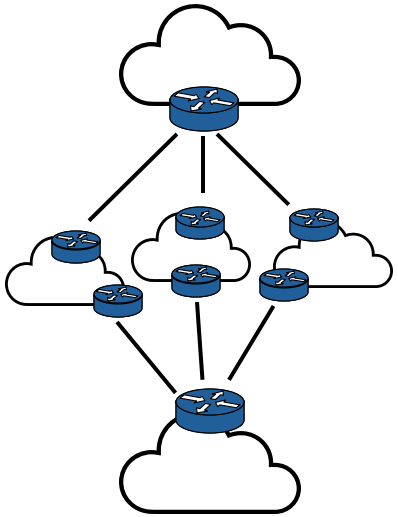

Internet Requires Consistent and Stable Routing

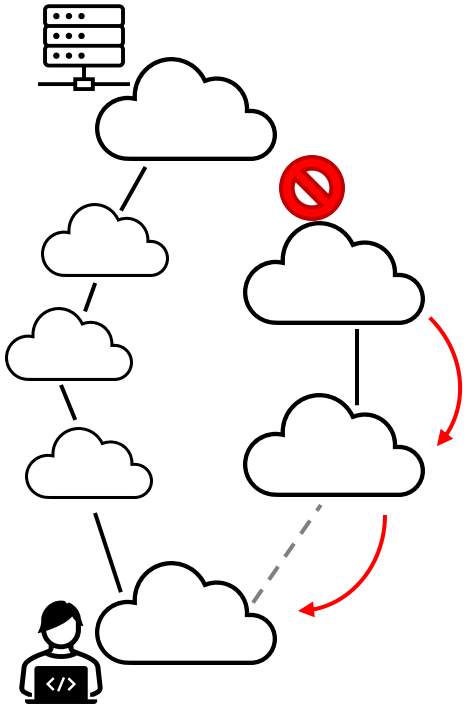

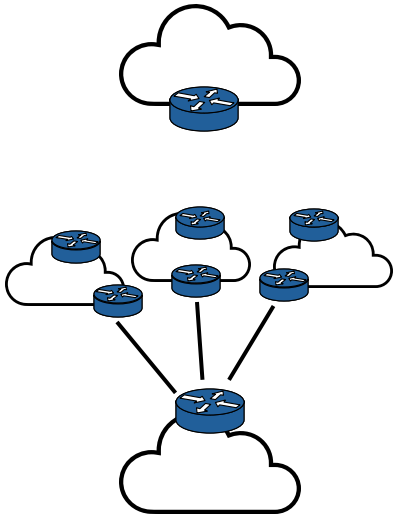

Internet Requires Consistent and Stable Routing

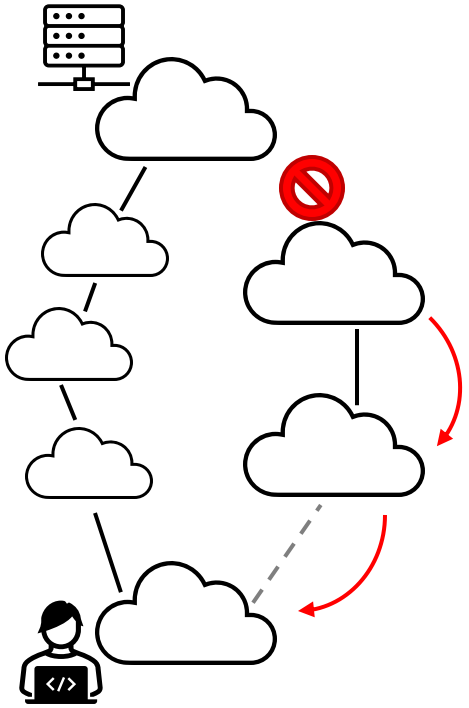

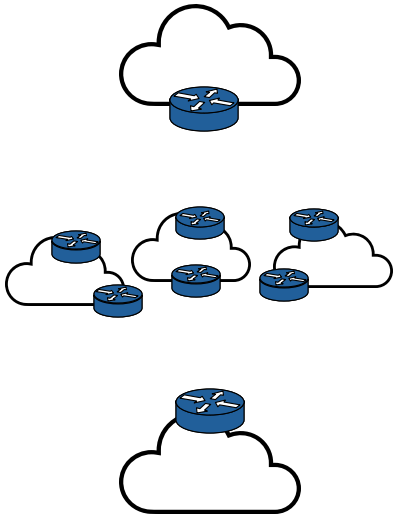

Internet Requires Consistent and Stable Routing

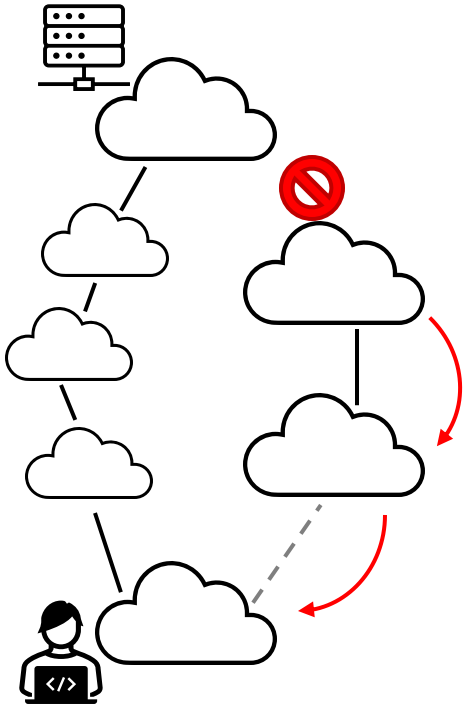

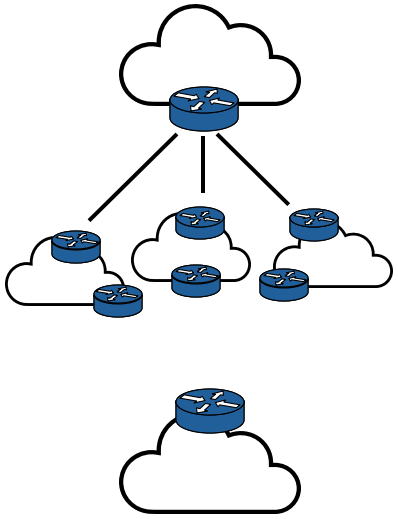

Internet Requires Consistent and Stable Routing

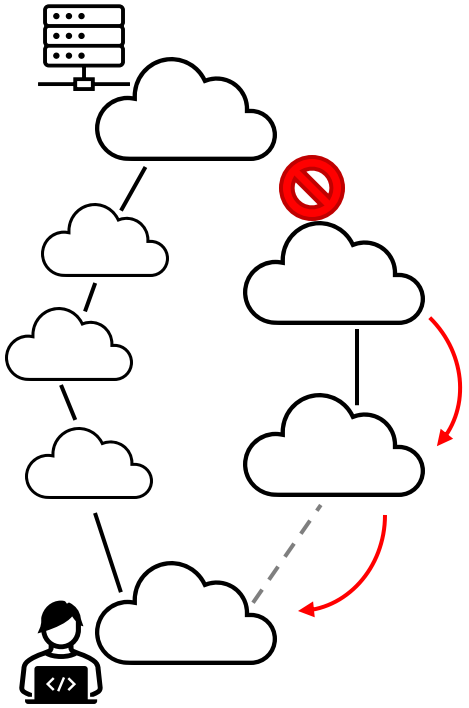

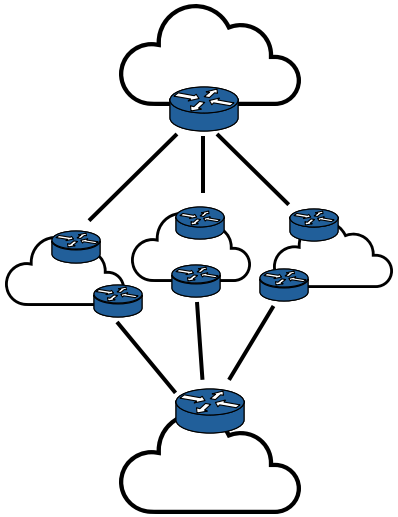

Internet Requires Consistent and Stable Routing

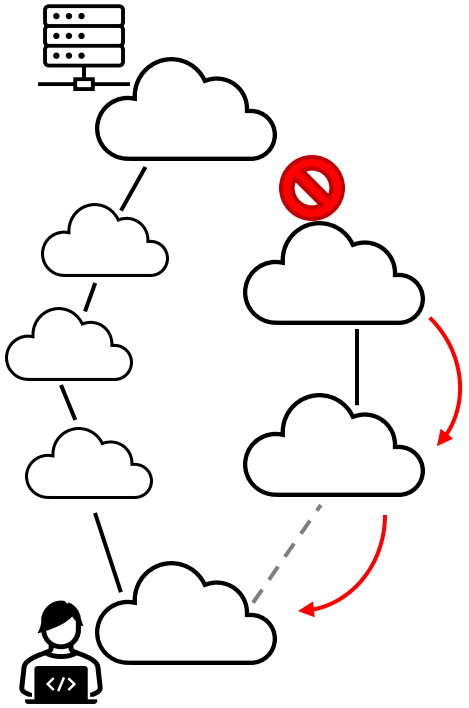

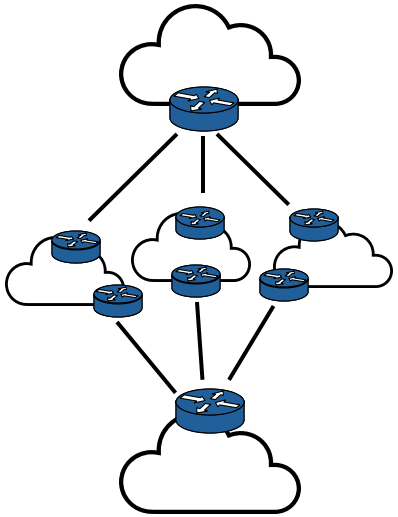

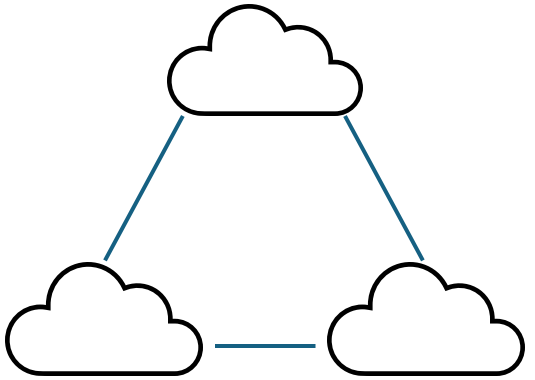

Internet Requires Consistent and Stable Routing

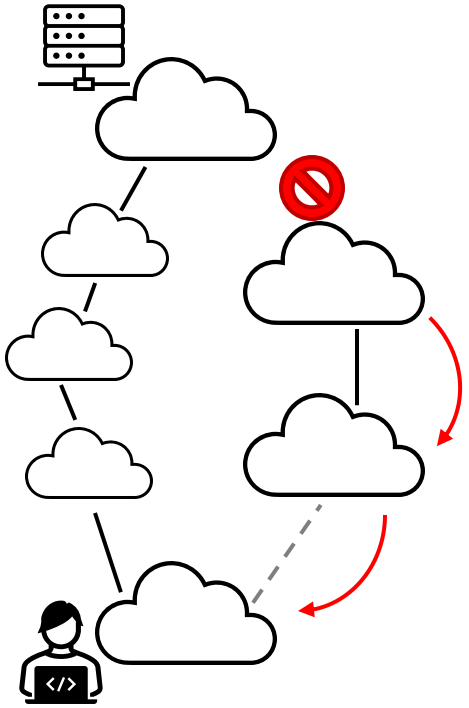

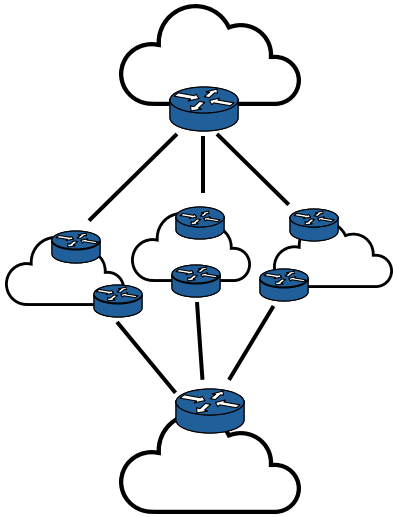

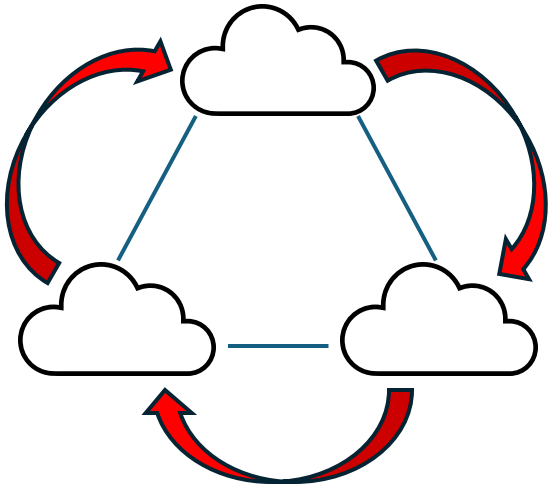

Internet Requires Consistent and Stable Routing

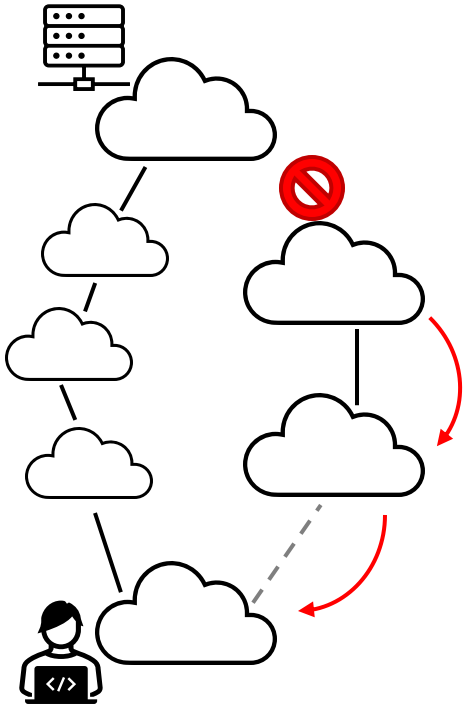

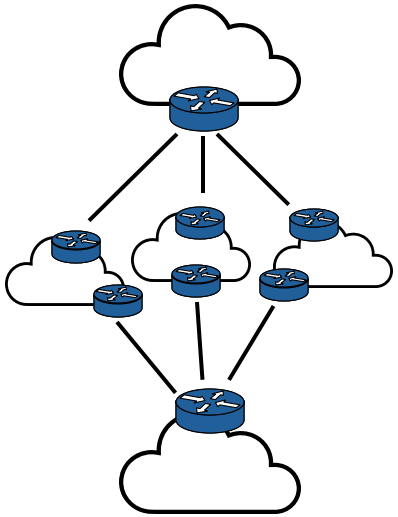

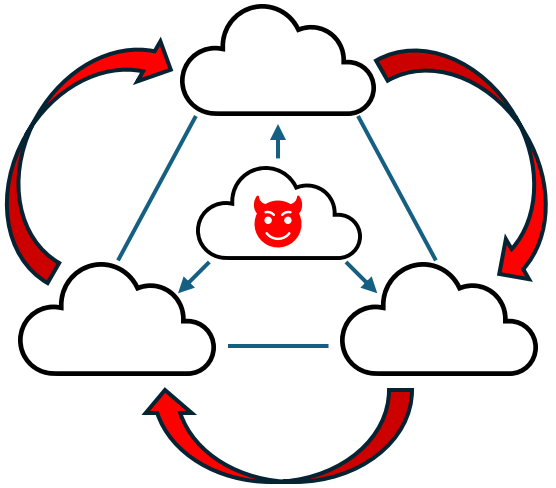

The Border Gateway Protocol (BGP)

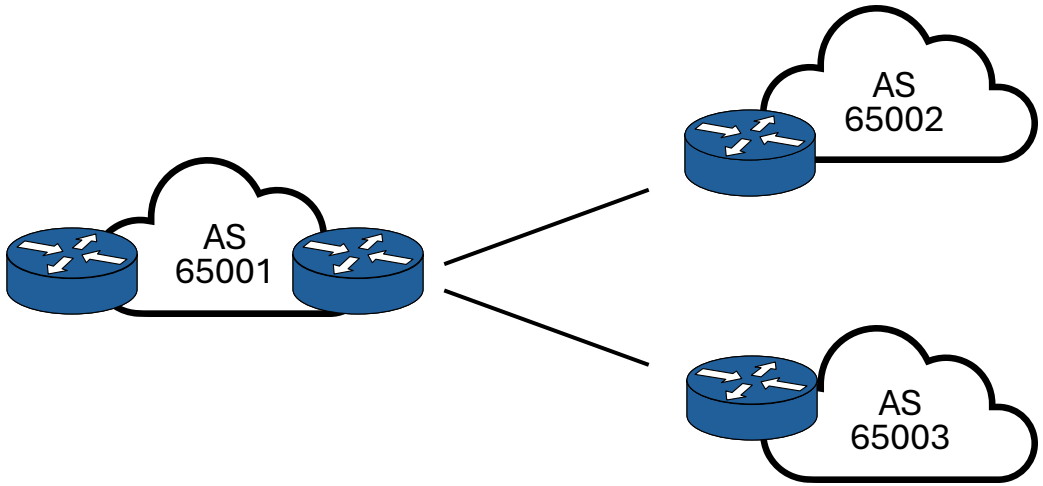

The Border Gateway Protocol (BGP)

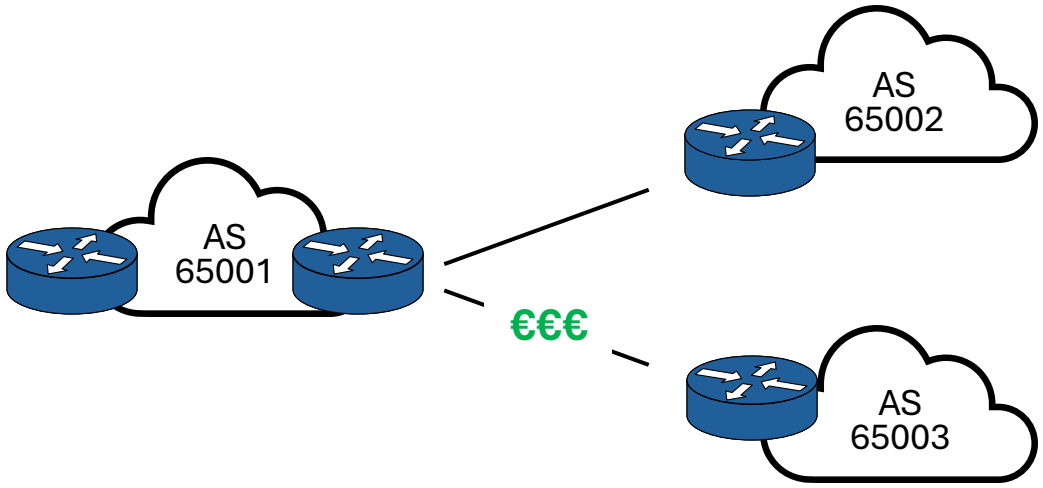

The Border Gateway Protocol (BGP)

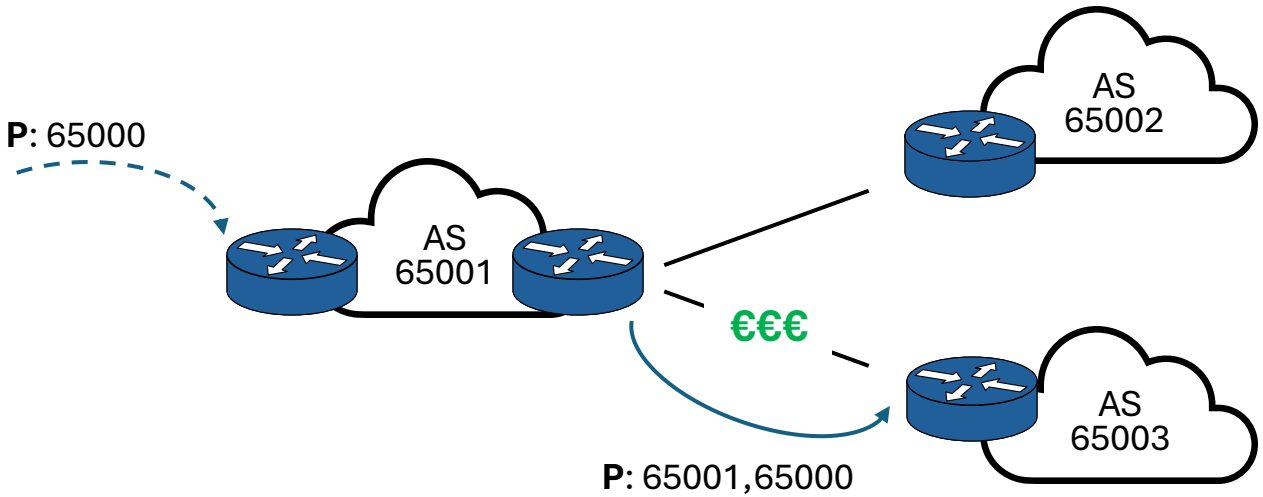

Does BGP Converge?

Theorem 5.1 (Simplified): A BGP system converges, if ASes prefer routes from customers over those received by peers or providers.(Gao-Rexford, 2001)

Does BGP Converge?

Theorem 5.1 (Simplified): A BGP system converges, if ASes prefer routes from customers over those received by peers or providers.(Gao-Rexford, 2001)

BGP AS Relationships and Routing Policies

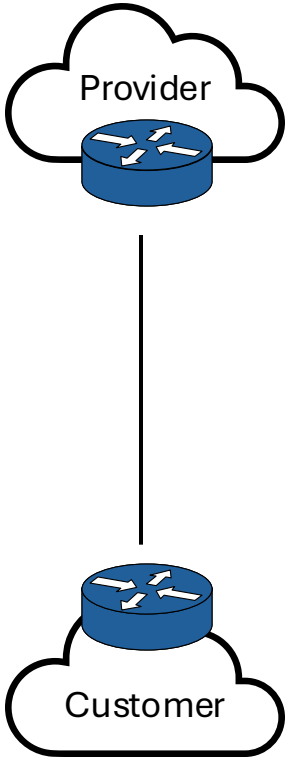

BGP AS Relationships and Routing Policies

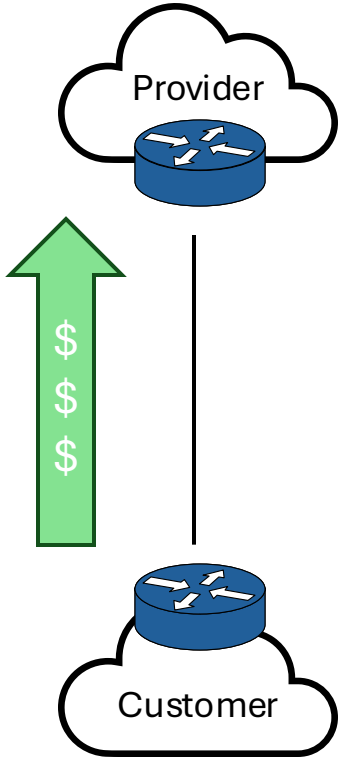

BGP AS Relationships and Routing Policies

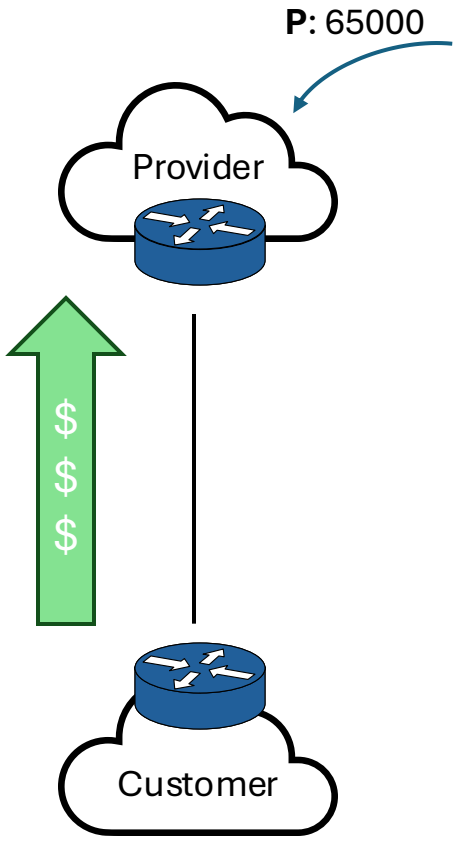

BGP AS Relationships and Routing Policies

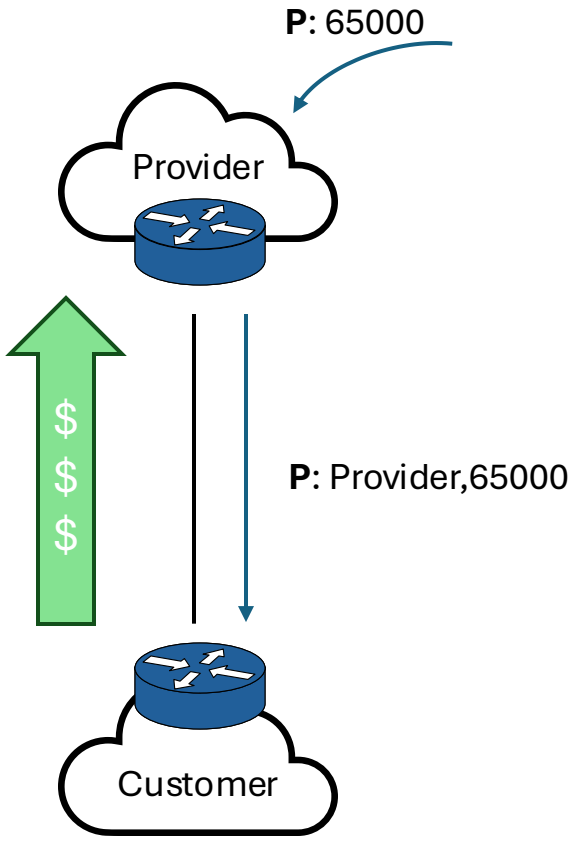

BGP AS Relationships and Routing Policies

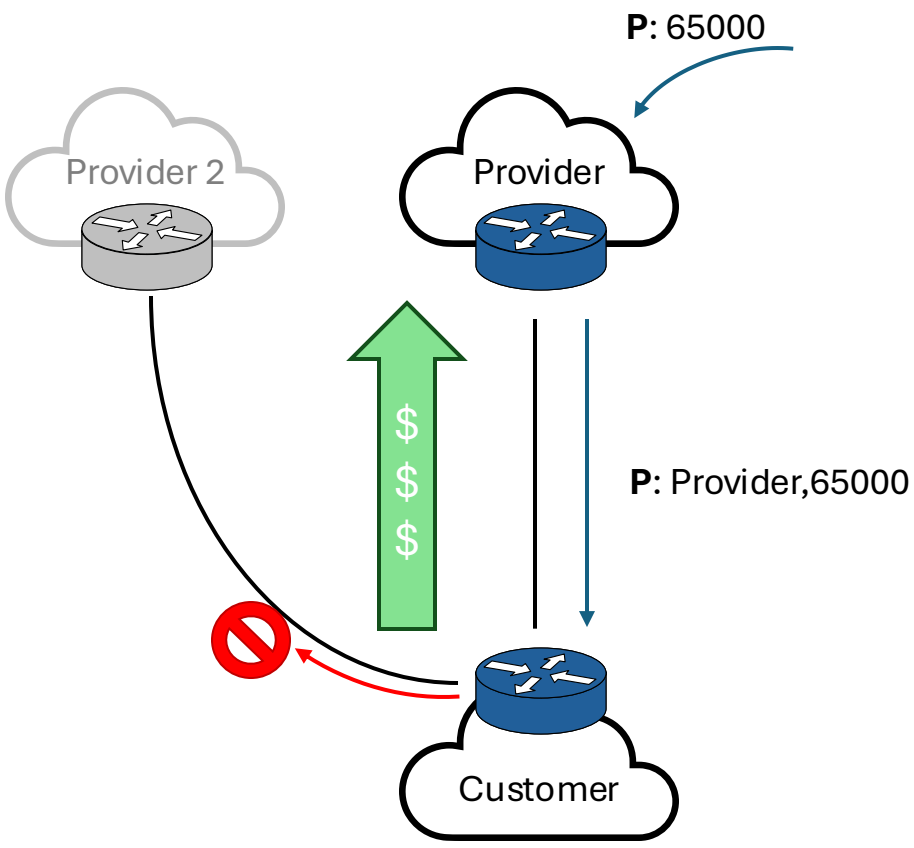

BGP AS Relationships and Routing Policies

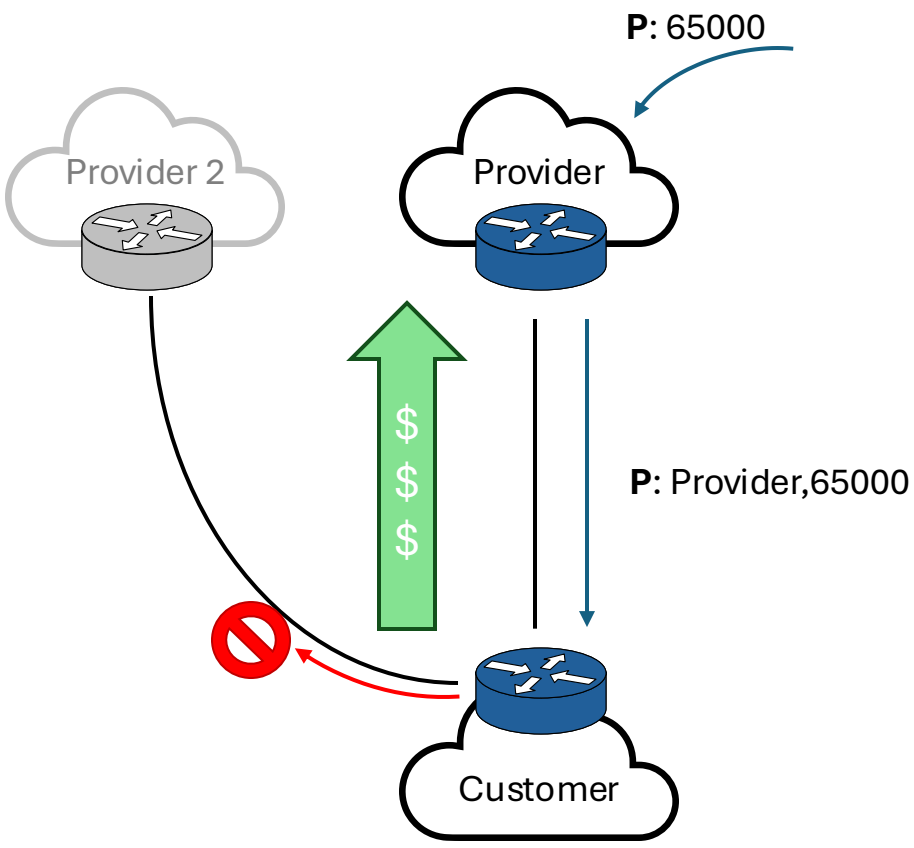

BGP AS Relationships and Routing Policies

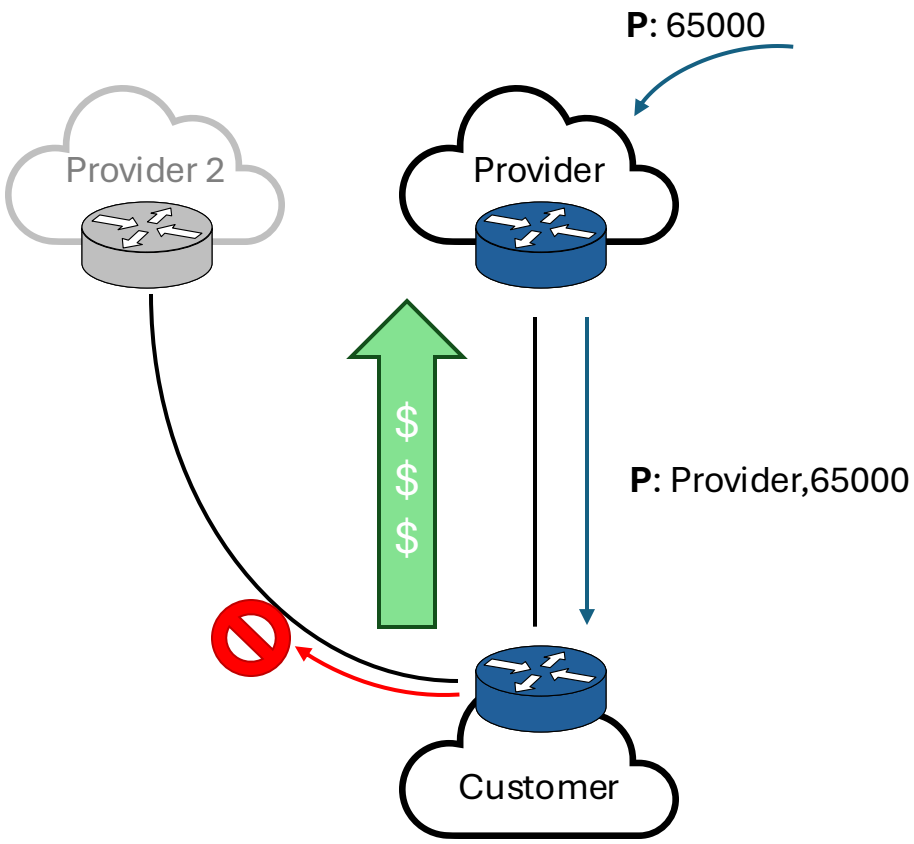

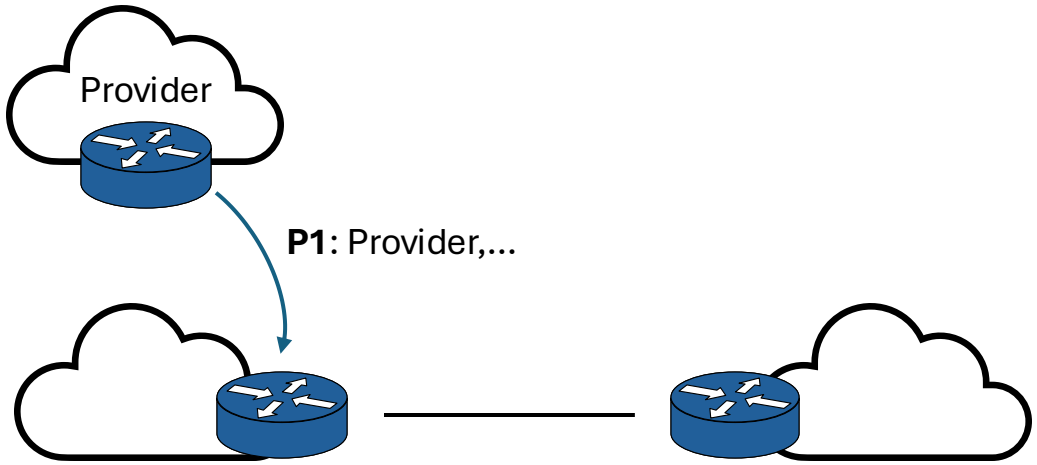

BGP AS Relationships and Routing Policies

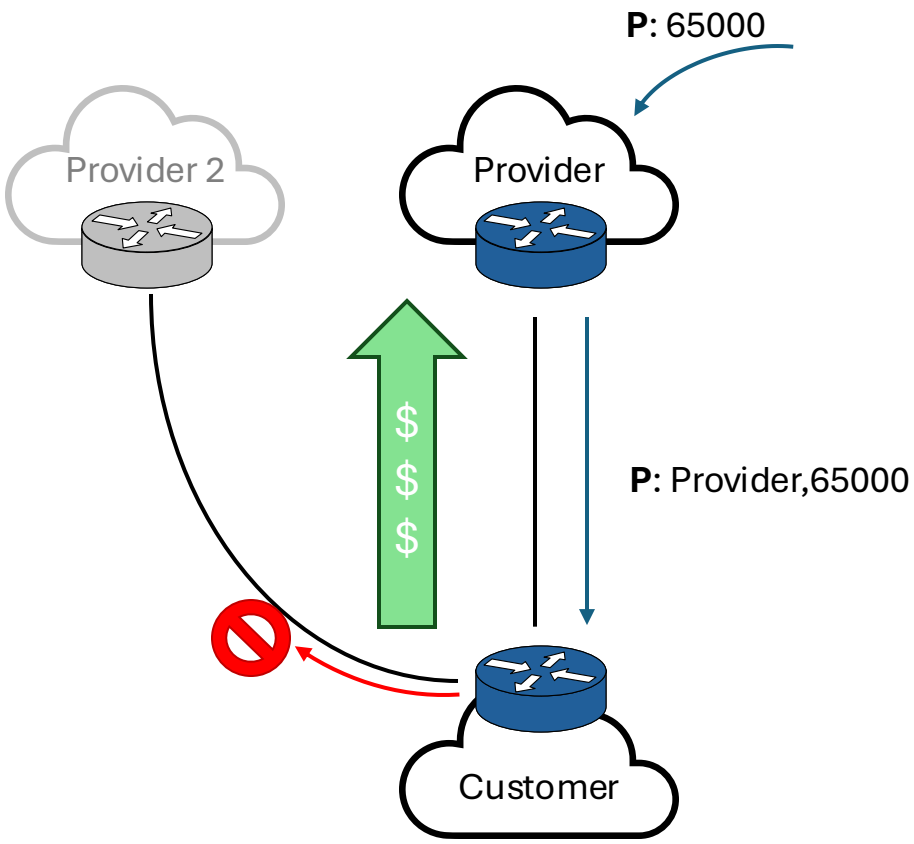

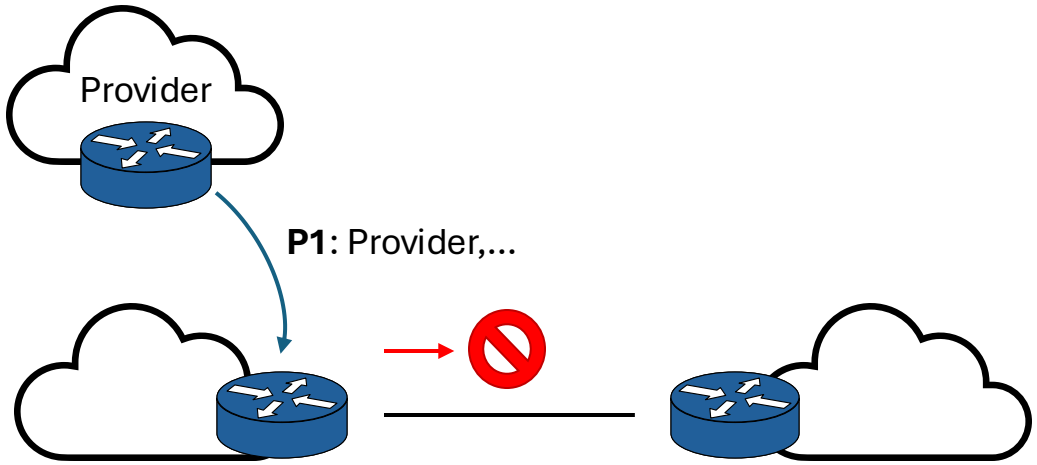

BGP AS Relationships and Routing Policies

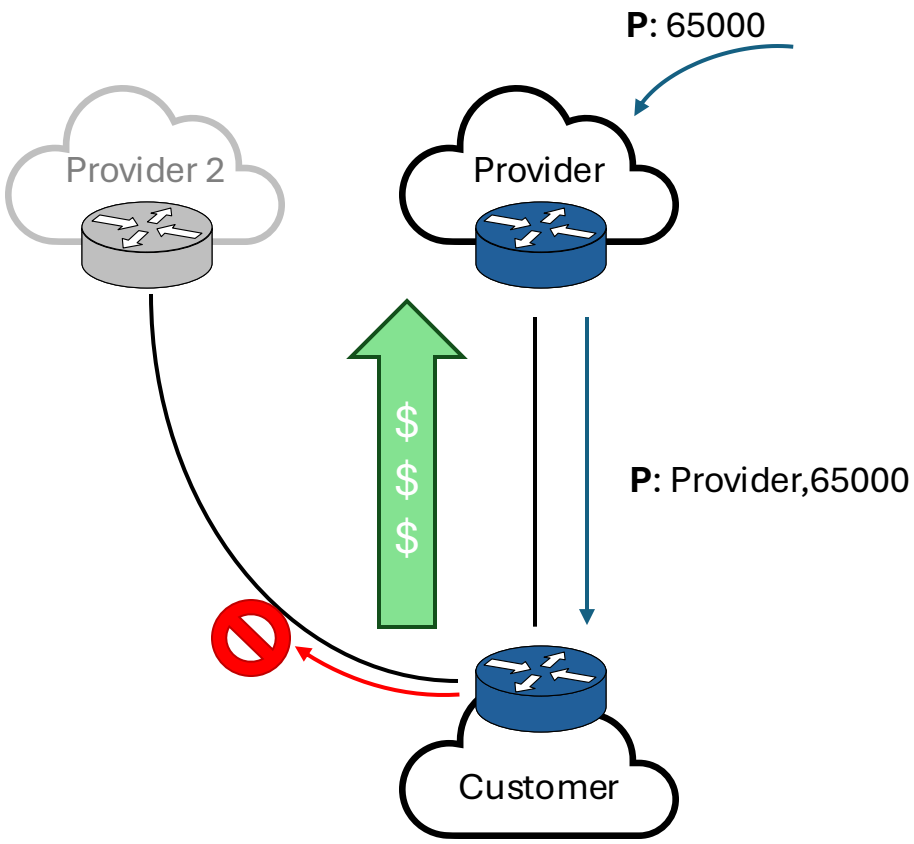

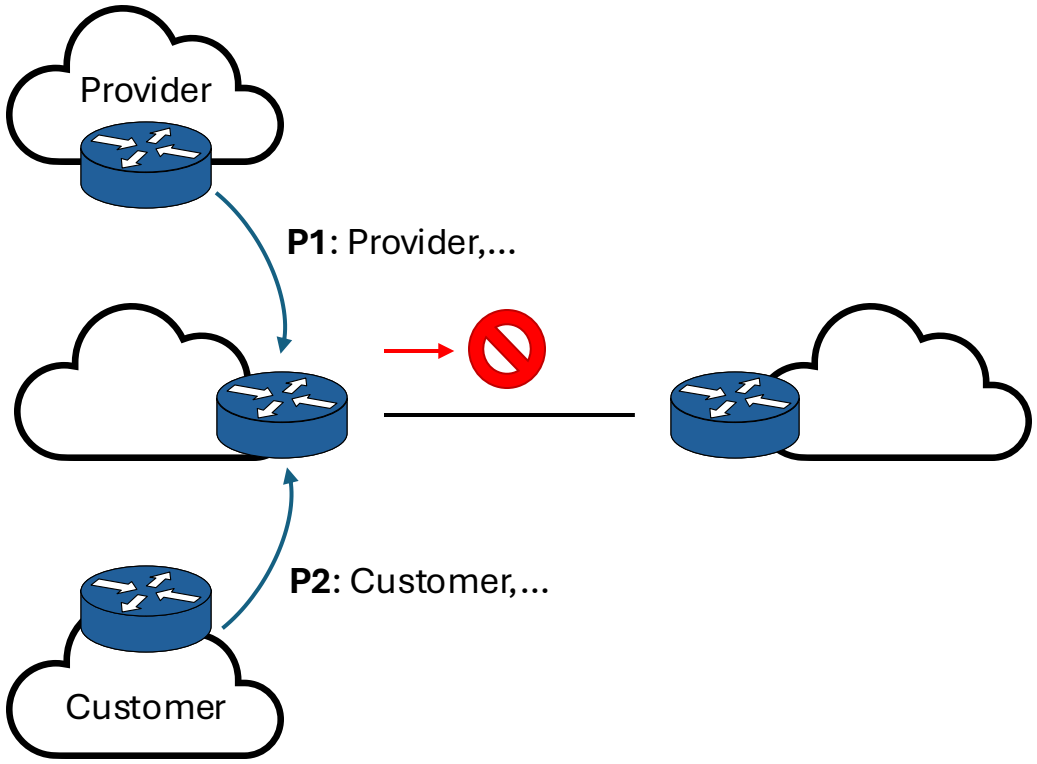

BGP AS Relationships and Routing Policies

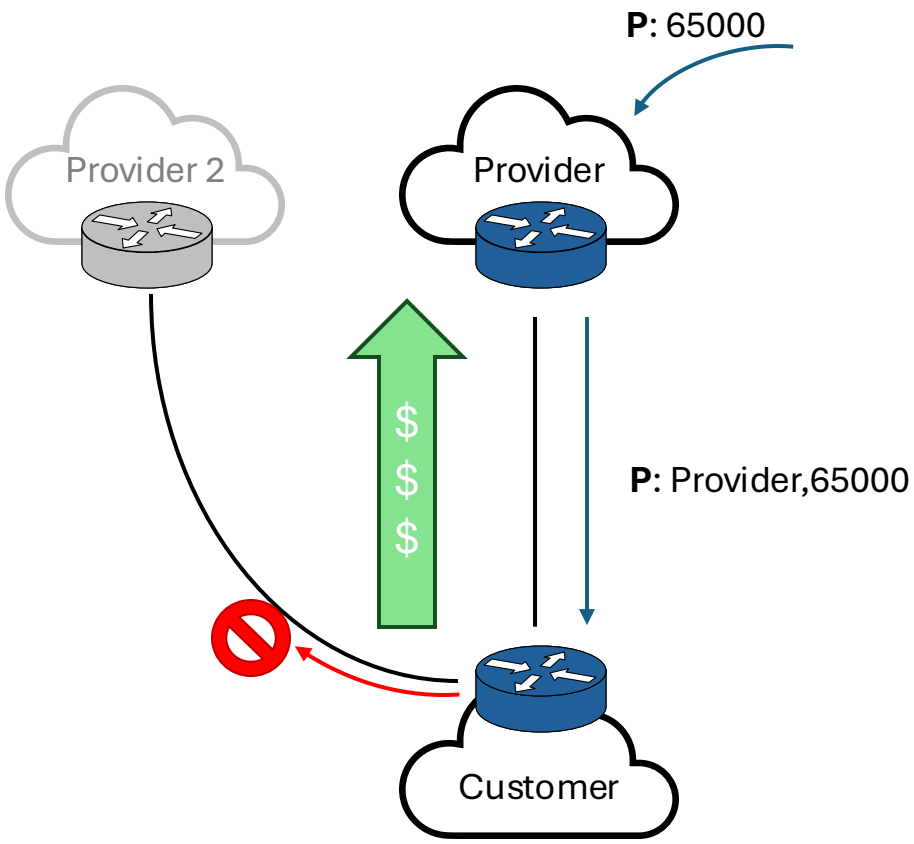

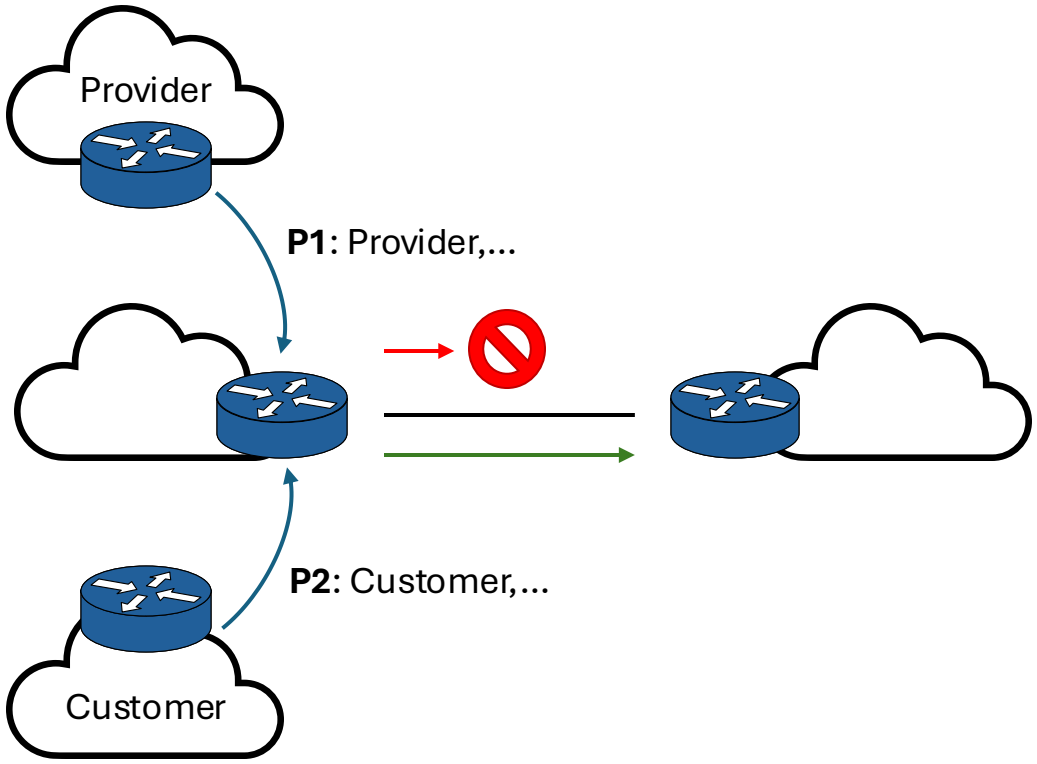

BGP AS Relationships and Routing Policies

Advertisement Forwarded To

Advertisement Received From	Customer	Peer	Provider
Customer	Advertise	Advertise	Advertise
Peer	Advertise	Drop	Drop
Provider	Advertise	Drop	Drop

Preference

Does BGP Converge?

Theorem 5.1 (Simplified): A BGP system converges, if ASes prefer routes from customers over those received by peers or providers.(Gao-Rexford, 2001)

Does BGP Converge?

Theorem 5.1 (Simplified): A BGP system converges, if ASes prefer routes from customers over those received by peers or providers. (Gao-Rexford, 2001)

ASSUMPTION VIOLATED

BGP Communities Can Violate Gao Rexford Rules

BGP Communities (RFC1997)

- Encode additional information
- Enhance Traffic Engineering
- Standard and custom communities
  - **Local Preference lowering (RFC1998)**
  - Do not forward route to specific AS

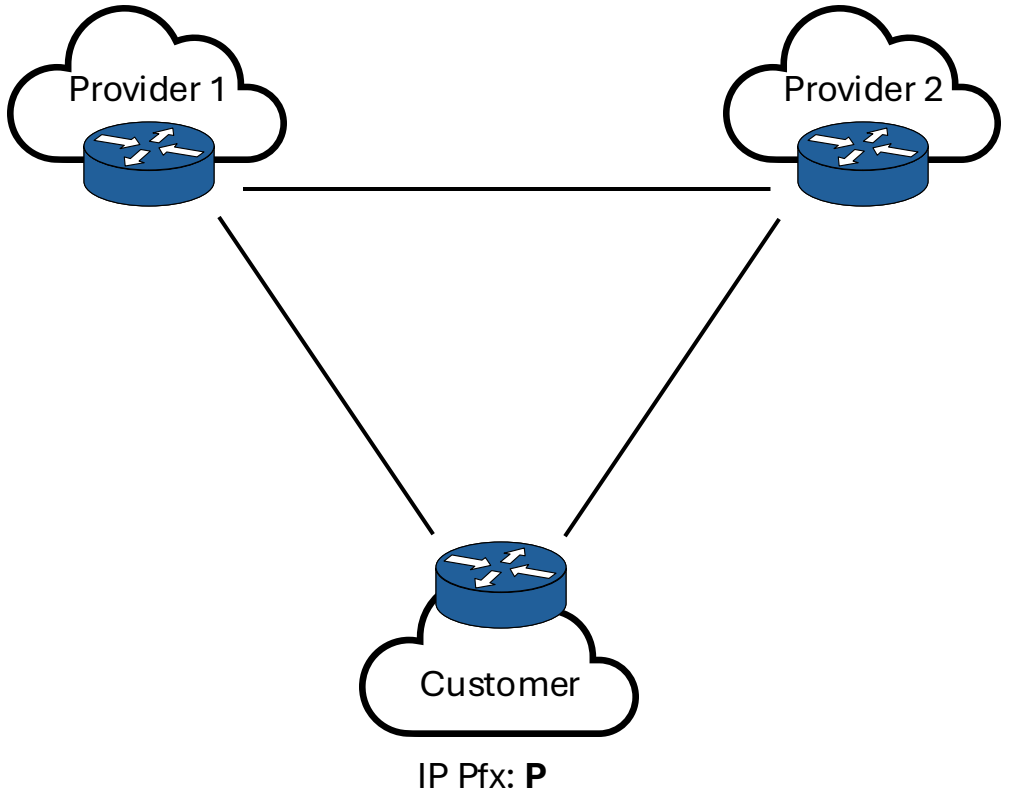

BGP Communities Can Violate Gao Rexford Rules

BGP Communities (RFC1997)

- Encode additional information
- Enhance Traffic Engineering
- Standard and custom communities
  - **Local Preference lowering (RFC1998)**
  - Do not forward route to specific AS

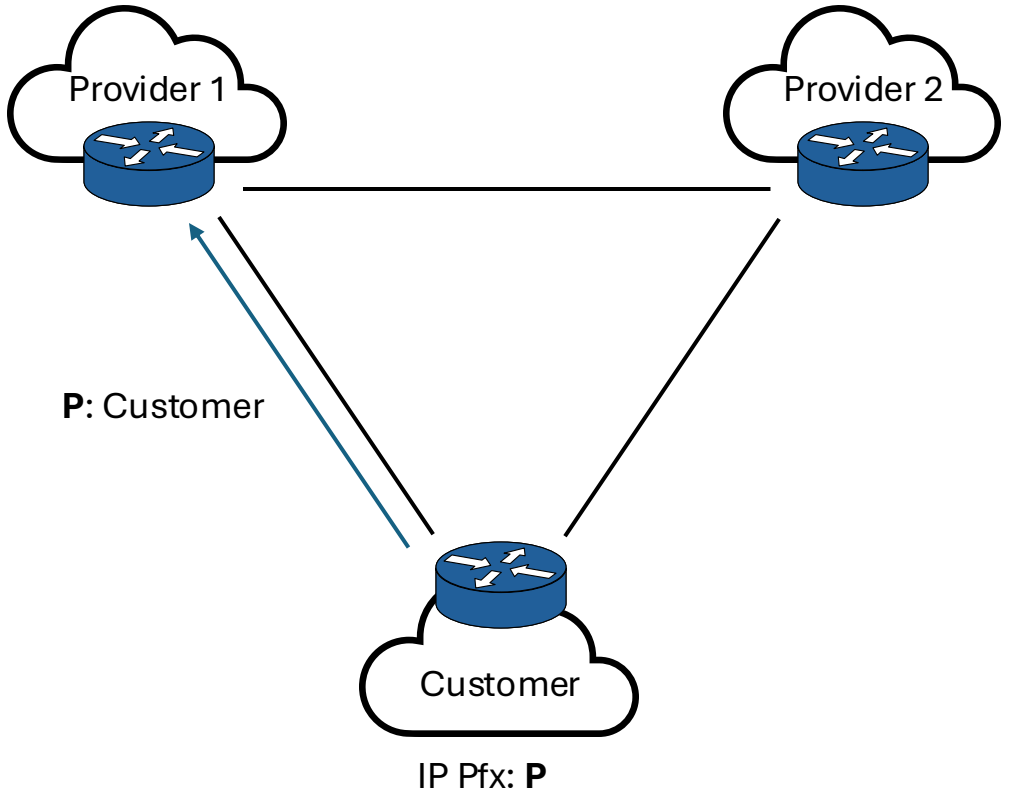

BGP Communities Can Violate Gao Rexford Rules

BGP Communities (RFC1997)

- Encode additional information
- Enhance Traffic Engineering
- Standard and custom communities
  - **Local Preference lowering (RFC1998)**
  - Do not forward route to specific AS

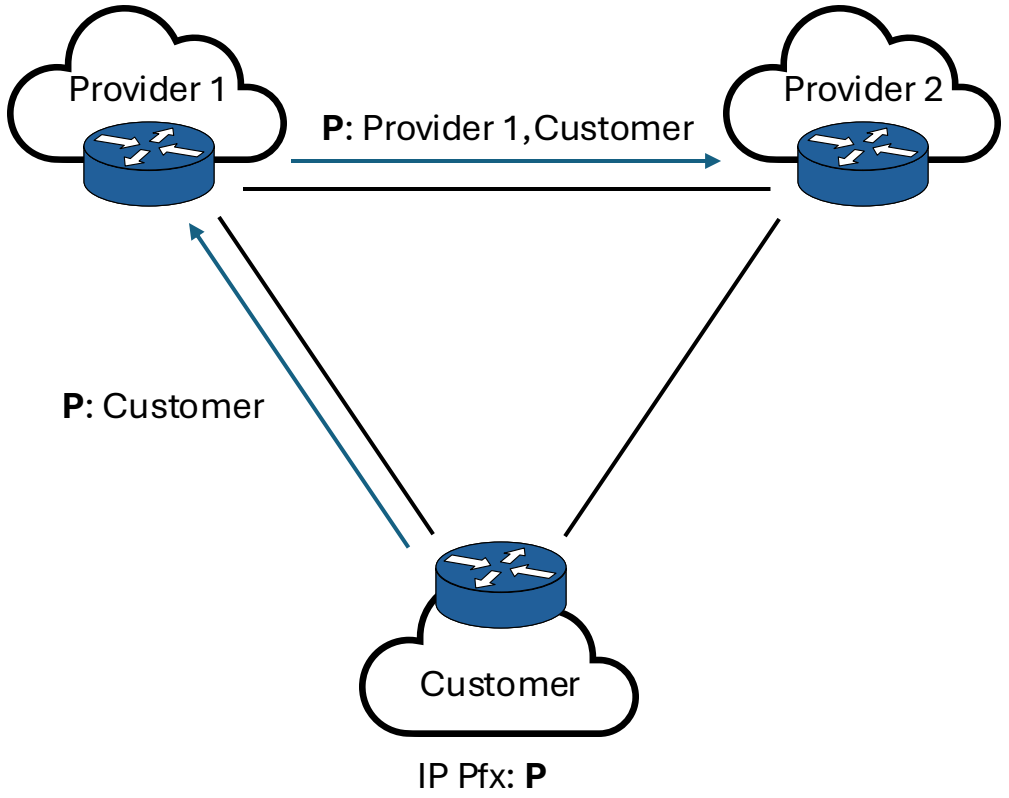

BGP Communities Can Violate Gao Rexford Rules

BGP Communities (RFC1997)

- Encode additional information
- Enhance Traffic Engineering
- Standard and custom communities
  - **Local Preference lowering (RFC1998)**
  - Do not forward route to specific AS

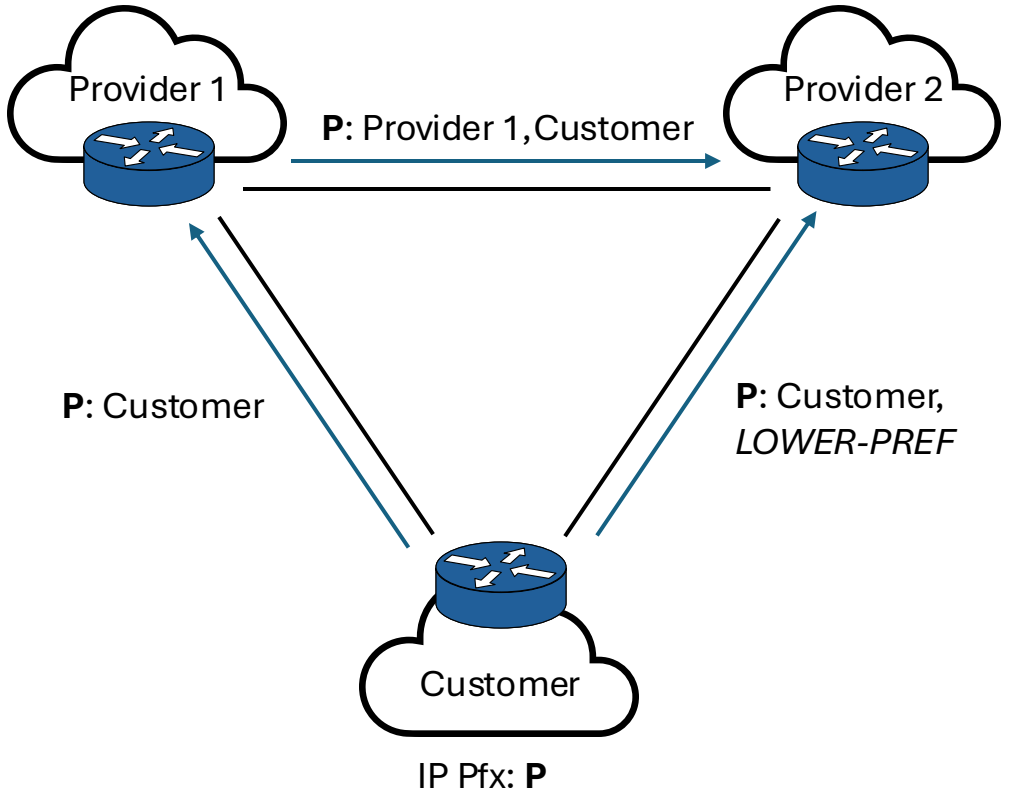

BGP Communities Can Violate Gao Rexford Rules

BGP Communities (RFC1997)

- Encode additional information
- Enhance Traffic Engineering
- Standard and custom communities
  - **Local Preference lowering (RFC1998)**
  - Do not forward route to specific AS

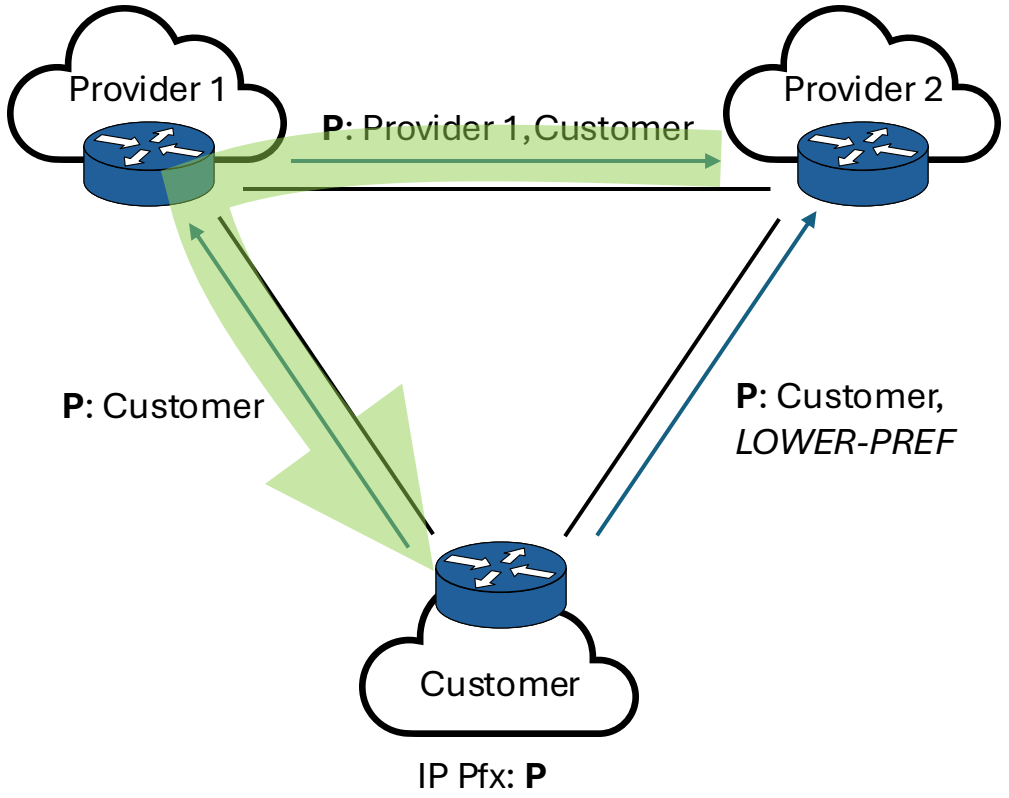

The BGP Vortex

Exploiting BGP Communities to Induce Persistent Oscillations

BGP Vortex Basics

Topology: 3 peers, 1 customer

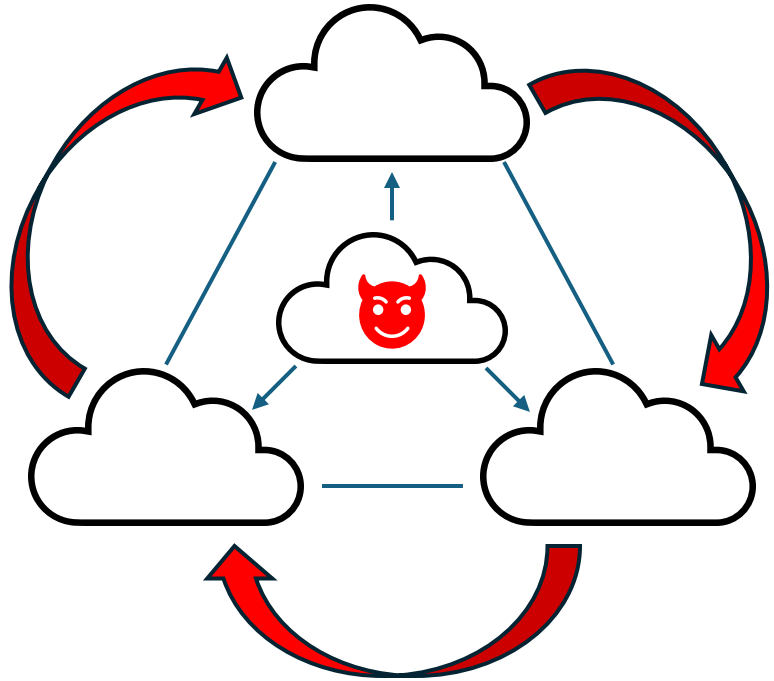

BGP Vortex Basics

Topology: 3 peers, 1 customer

Supported Communities:

- Local Preference lowering (RFC1998)
- Do not forward route to specific AS

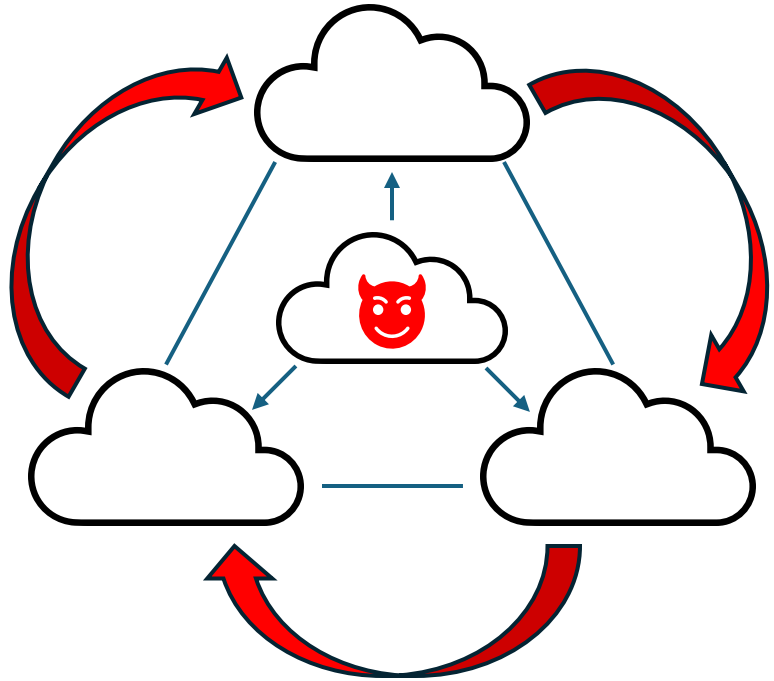

BGP Vortex Basics

Topology: 3 peers, 1 customer

Supported Communities:

- Local Preference lowering (RFC1998)
- Do not forward route to specific AS

Adversary Capabilities:

- Controls BGP speaker in customer AS
- Send legitimate, well-formed BGP UPDATEs

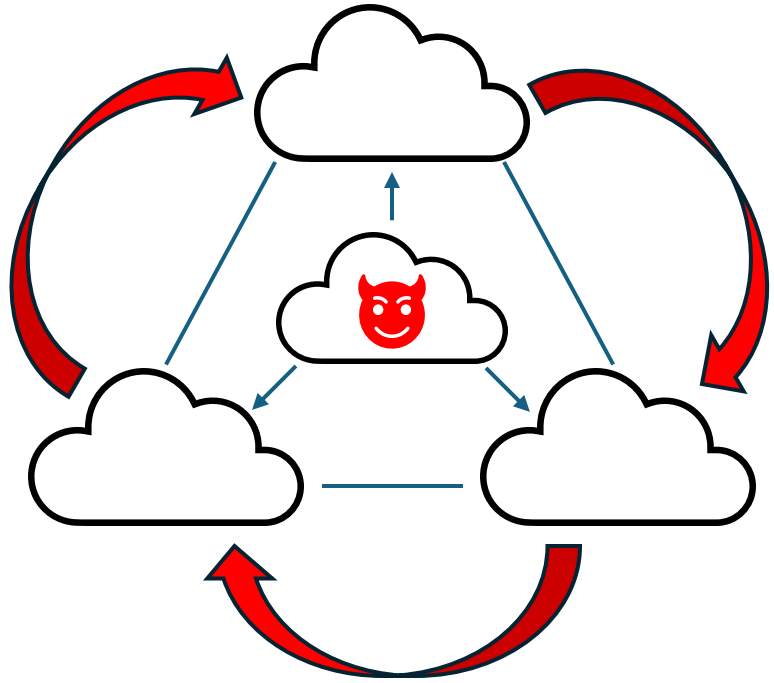

BGP Vortex: Inducing Persistent Oscillations

	Customer	Peer	Provider
Customer	Advertise	Advertise	Advertise
Peer	Advertise	Drop	Drop
Provider	Advertise	Drop	Drop

BGP Vortex: Inducing Persistent Oscillations

	Customer	Peer	Provider
Customer	Advertise	Advertise	Advertise
Peer	Advertise	Drop	Drop
Provider	Advertise	Drop	Drop

BGP Vortex: Inducing Persistent Oscillations

	Customer	Peer	Provider
Customer	Advertise	Advertise	Advertise
Peer	Advertise	Drop	Drop
Provider	Advertise	Drop	Drop

BGP Vortex: Inducing Persistent Oscillations

	Customer	Peer	Provider
Customer	Advertise	Advertise	Advertise
Peer	Advertise	Drop	Drop
Provider	Advertise	Drop	Drop

Preference

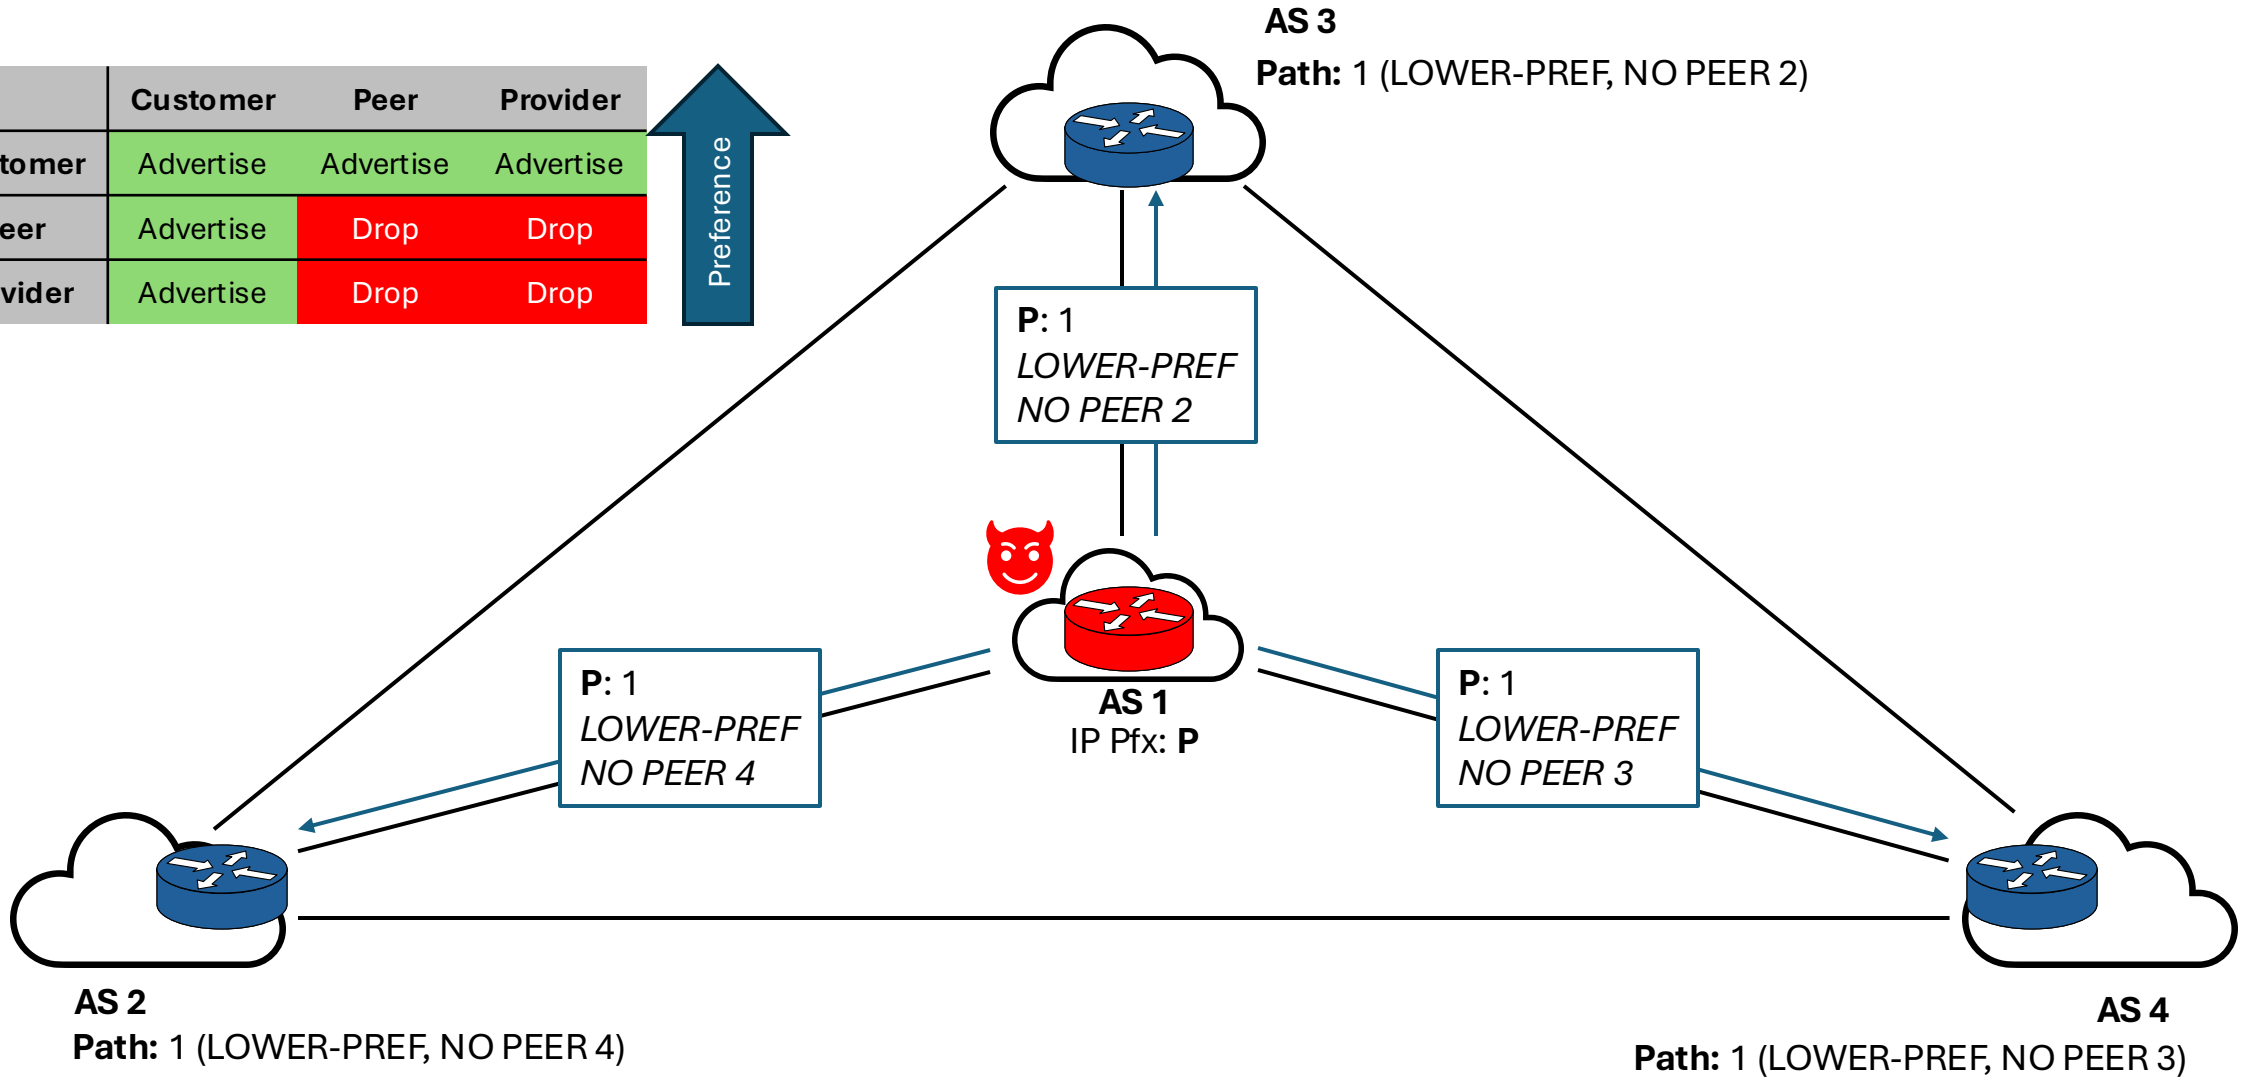

BGP Vortex: Inducing Persistent Oscillations

	Customer	Peer	Provider
Customer	Advertise	Advertise	Advertise
Peer	Advertise	Drop	Drop
Provider	Advertise	Drop	Drop

Preference

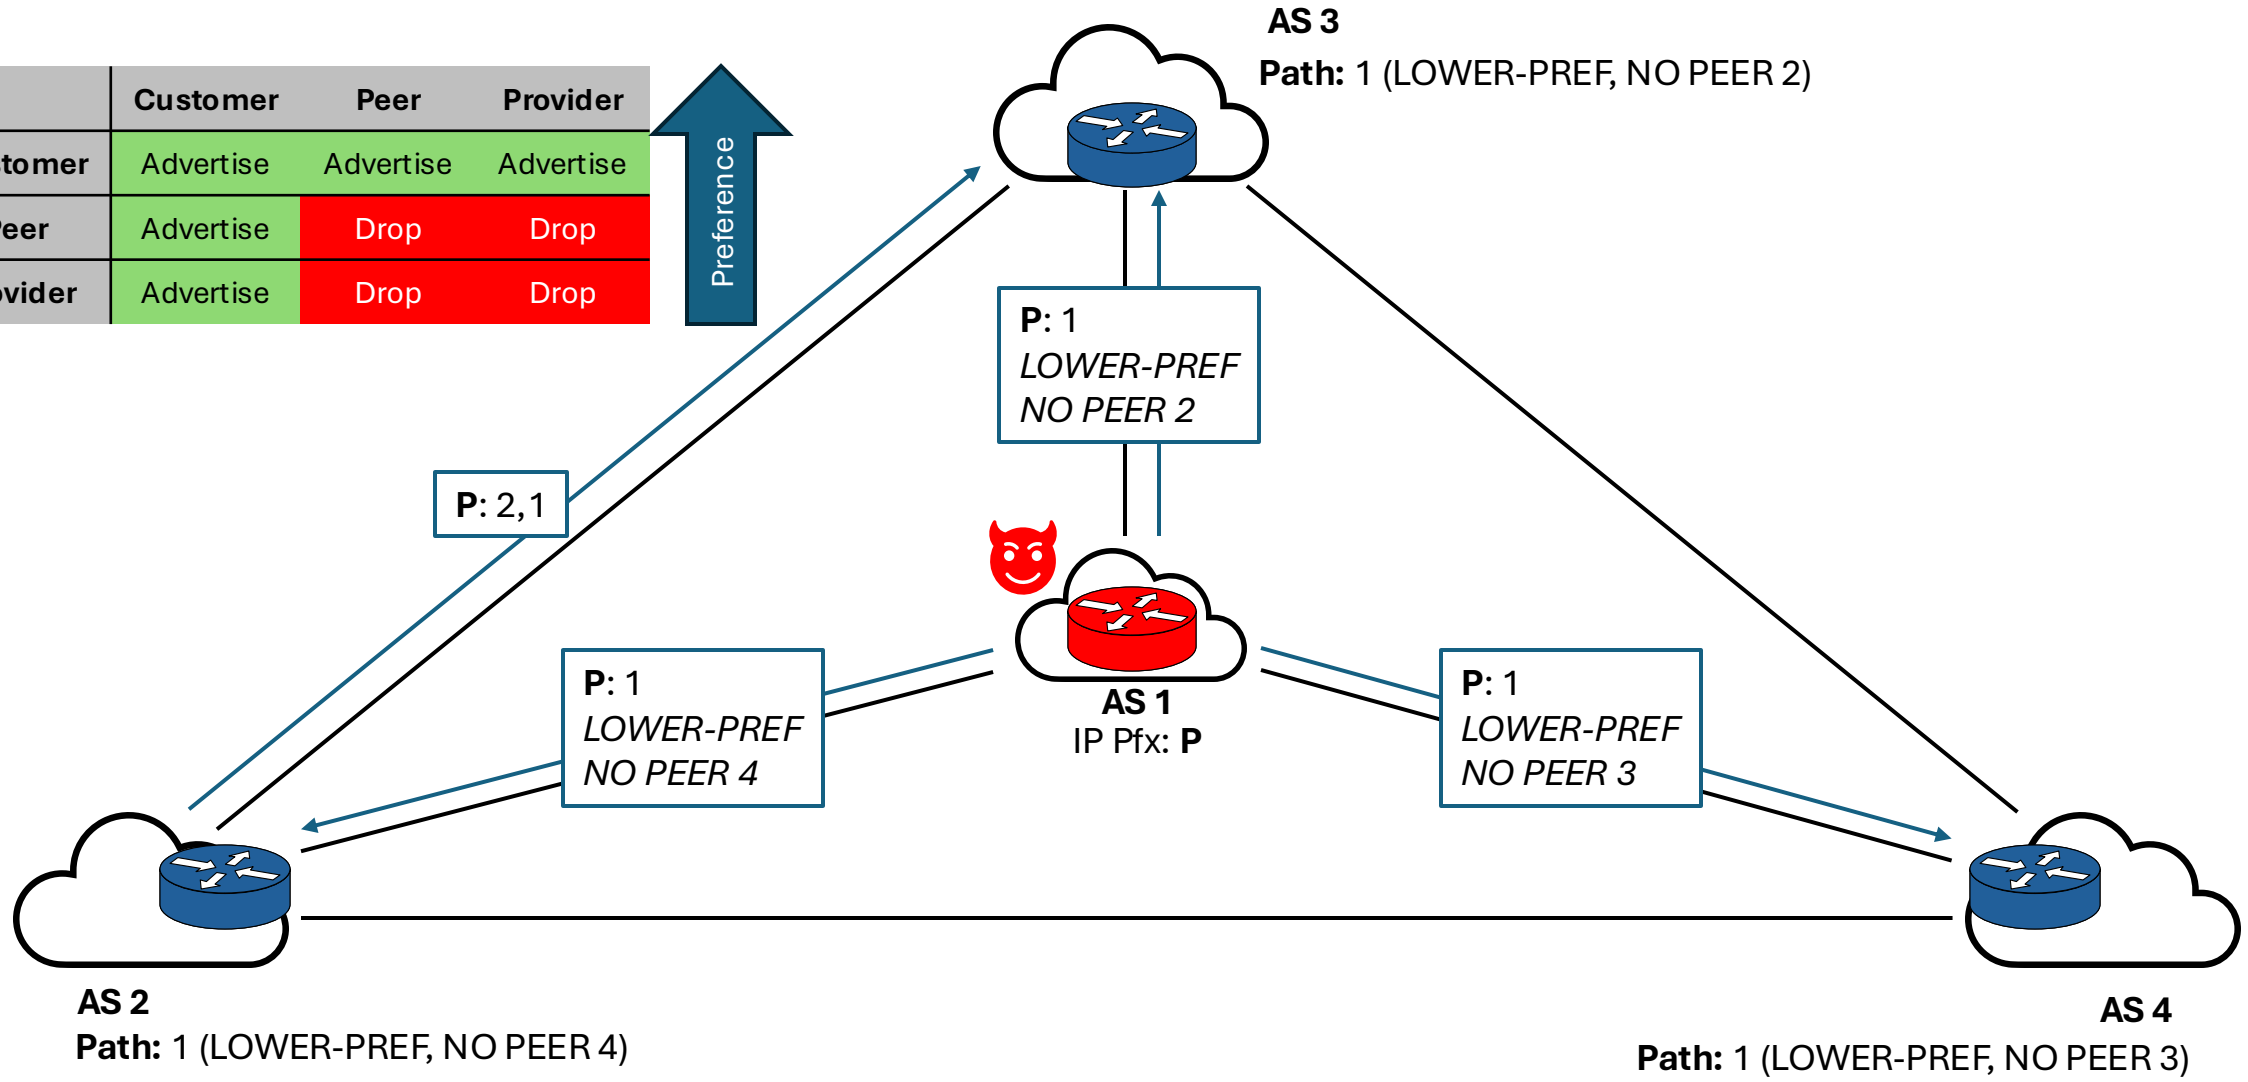

BGP Vortex: Inducing Persistent Oscillations

	Customer	Peer	Provider
Customer	Advertise	Advertise	Advertise
Peer	Advertise	Drop	Drop
Provider	Advertise	Drop	Drop

Preference

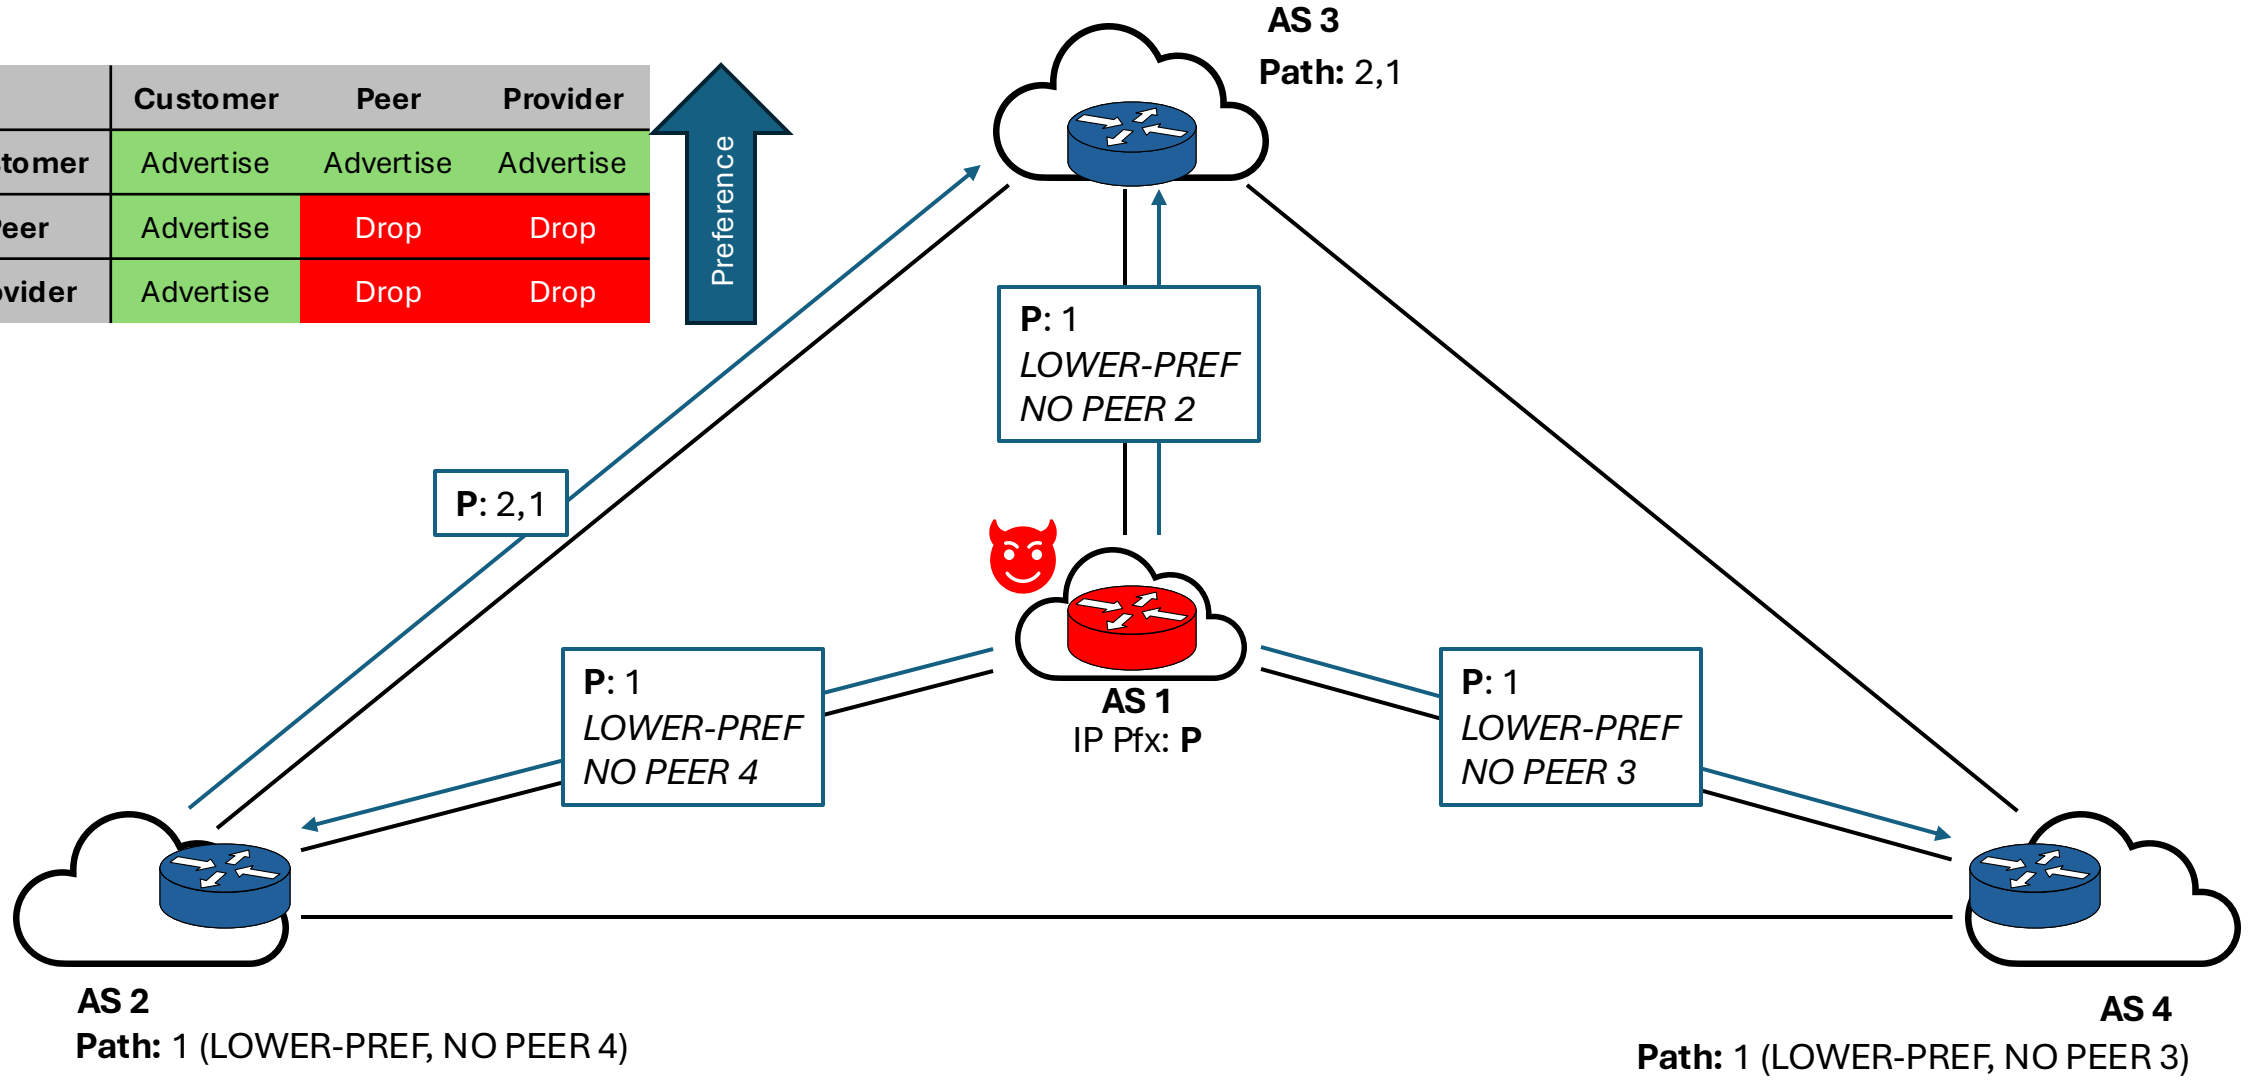

BGP Vortex: Inducing Persistent Oscillations

	Customer	Peer	Provider
Customer	Advertise	Advertise	Advertise
Peer	Advertise	Drop	Drop
Provider	Advertise	Drop	Drop

Preference

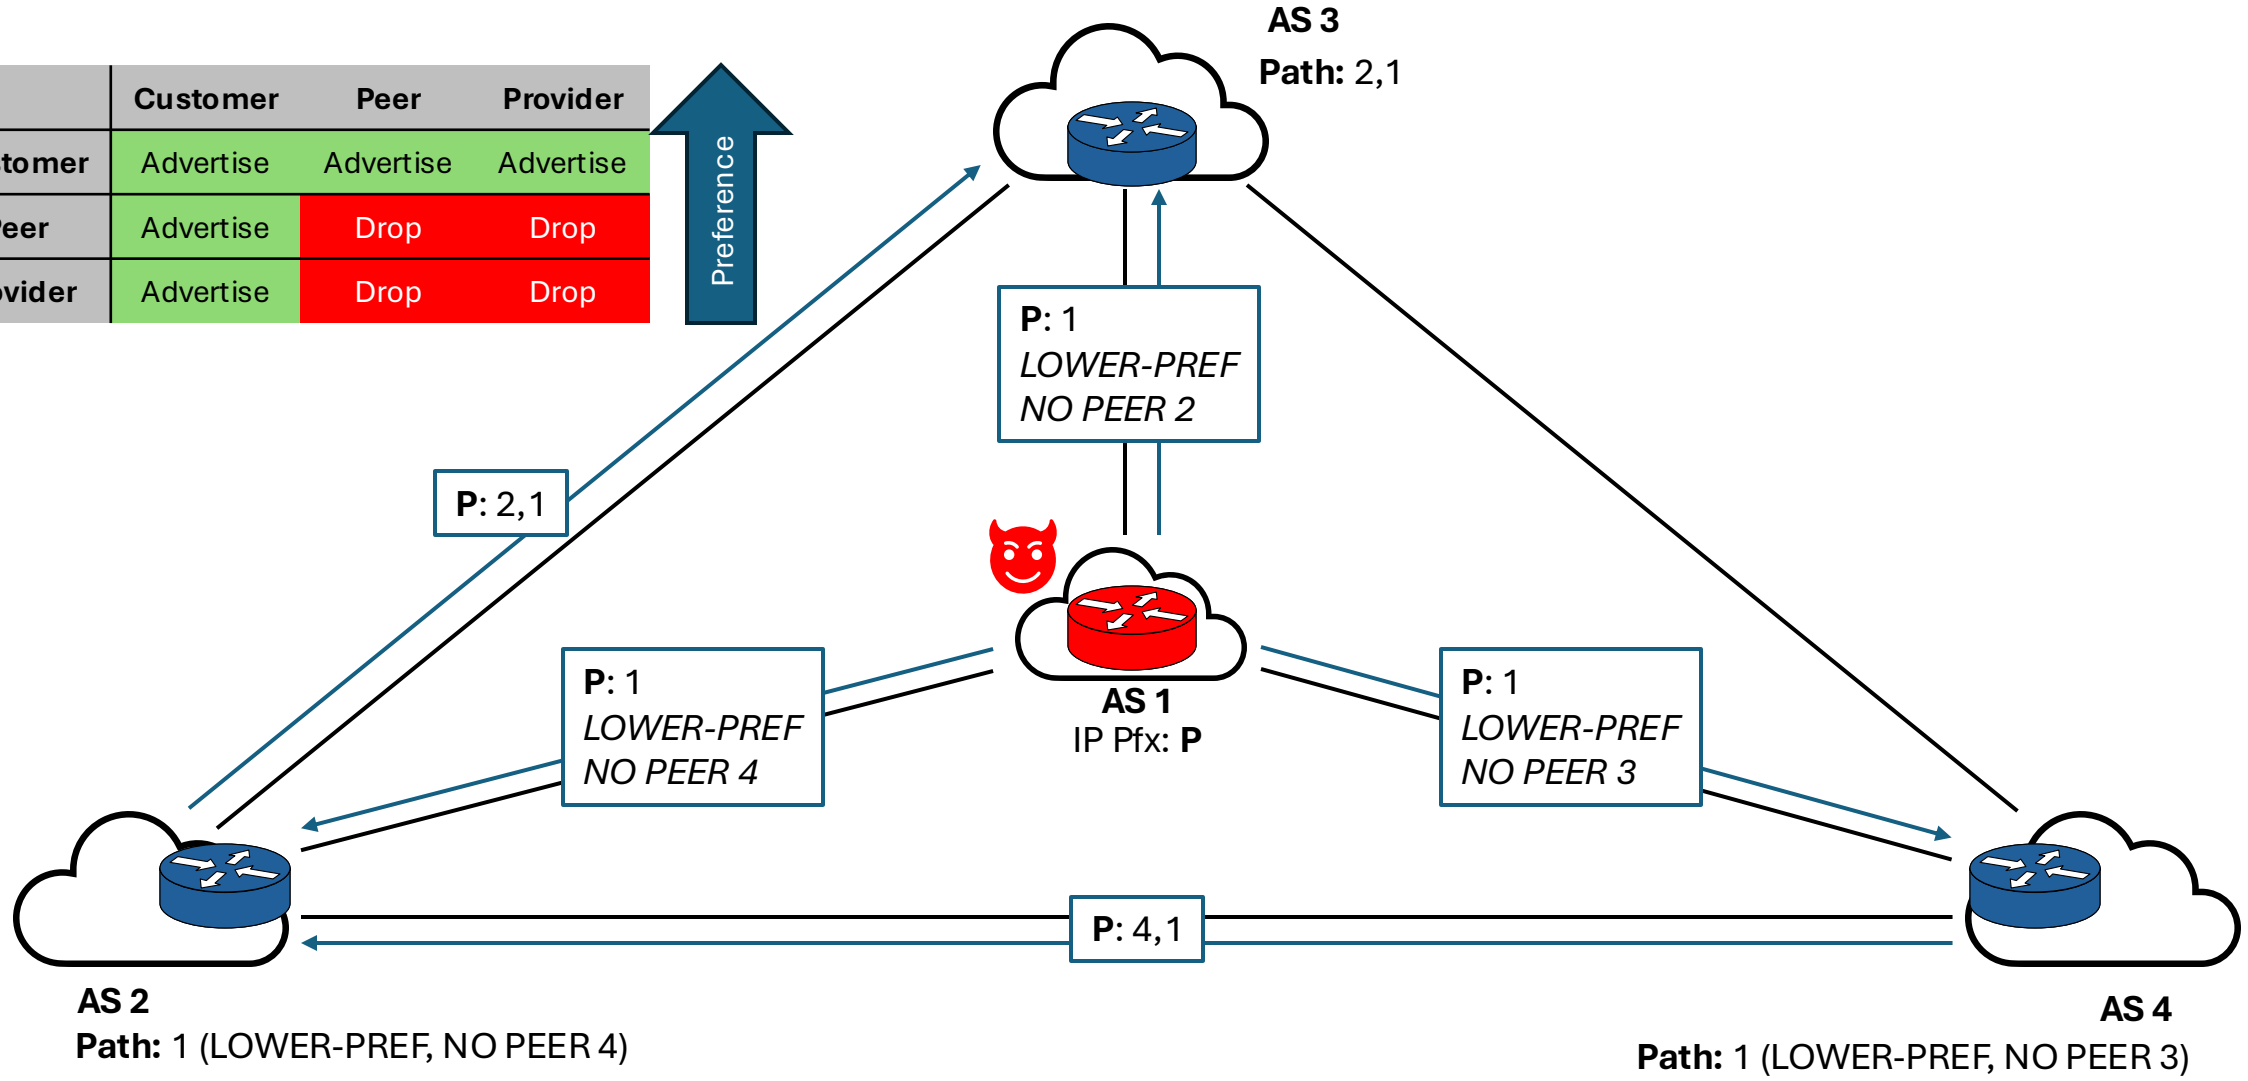

BGP Vortex: Inducing Persistent Oscillations

	Customer	Peer	Provider
Customer	Advertise	Advertise	Advertise
Peer	Advertise	Drop	Drop
Provider	Advertise	Drop	Drop

Preference

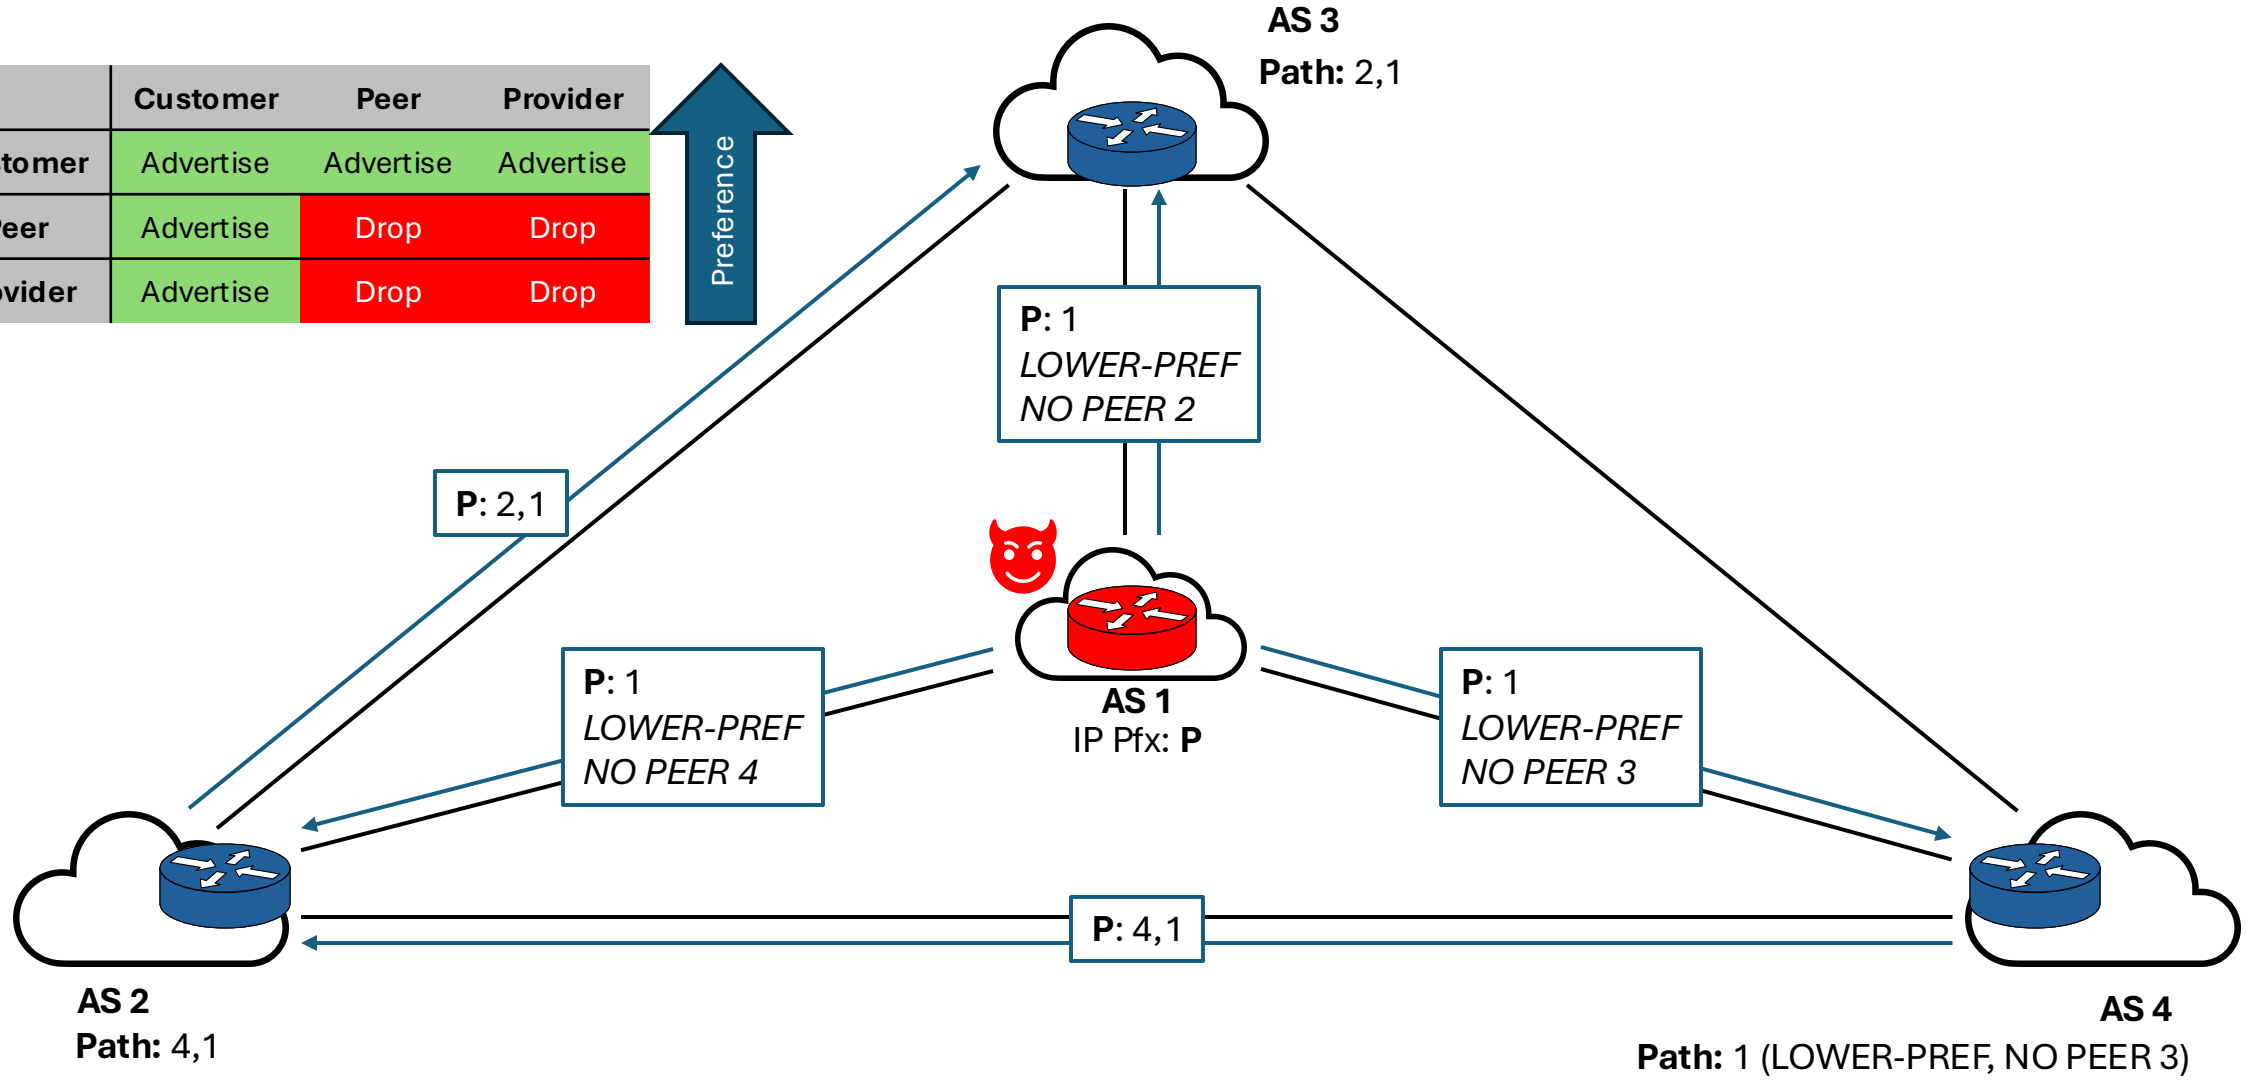

BGP Vortex: Inducing Persistent Oscillations

	Customer	Peer	Provider
Customer	Advertise	Advertise	Advertise
Peer	Advertise	Drop	Drop
Provider	Advertise	Drop	Drop

Preference

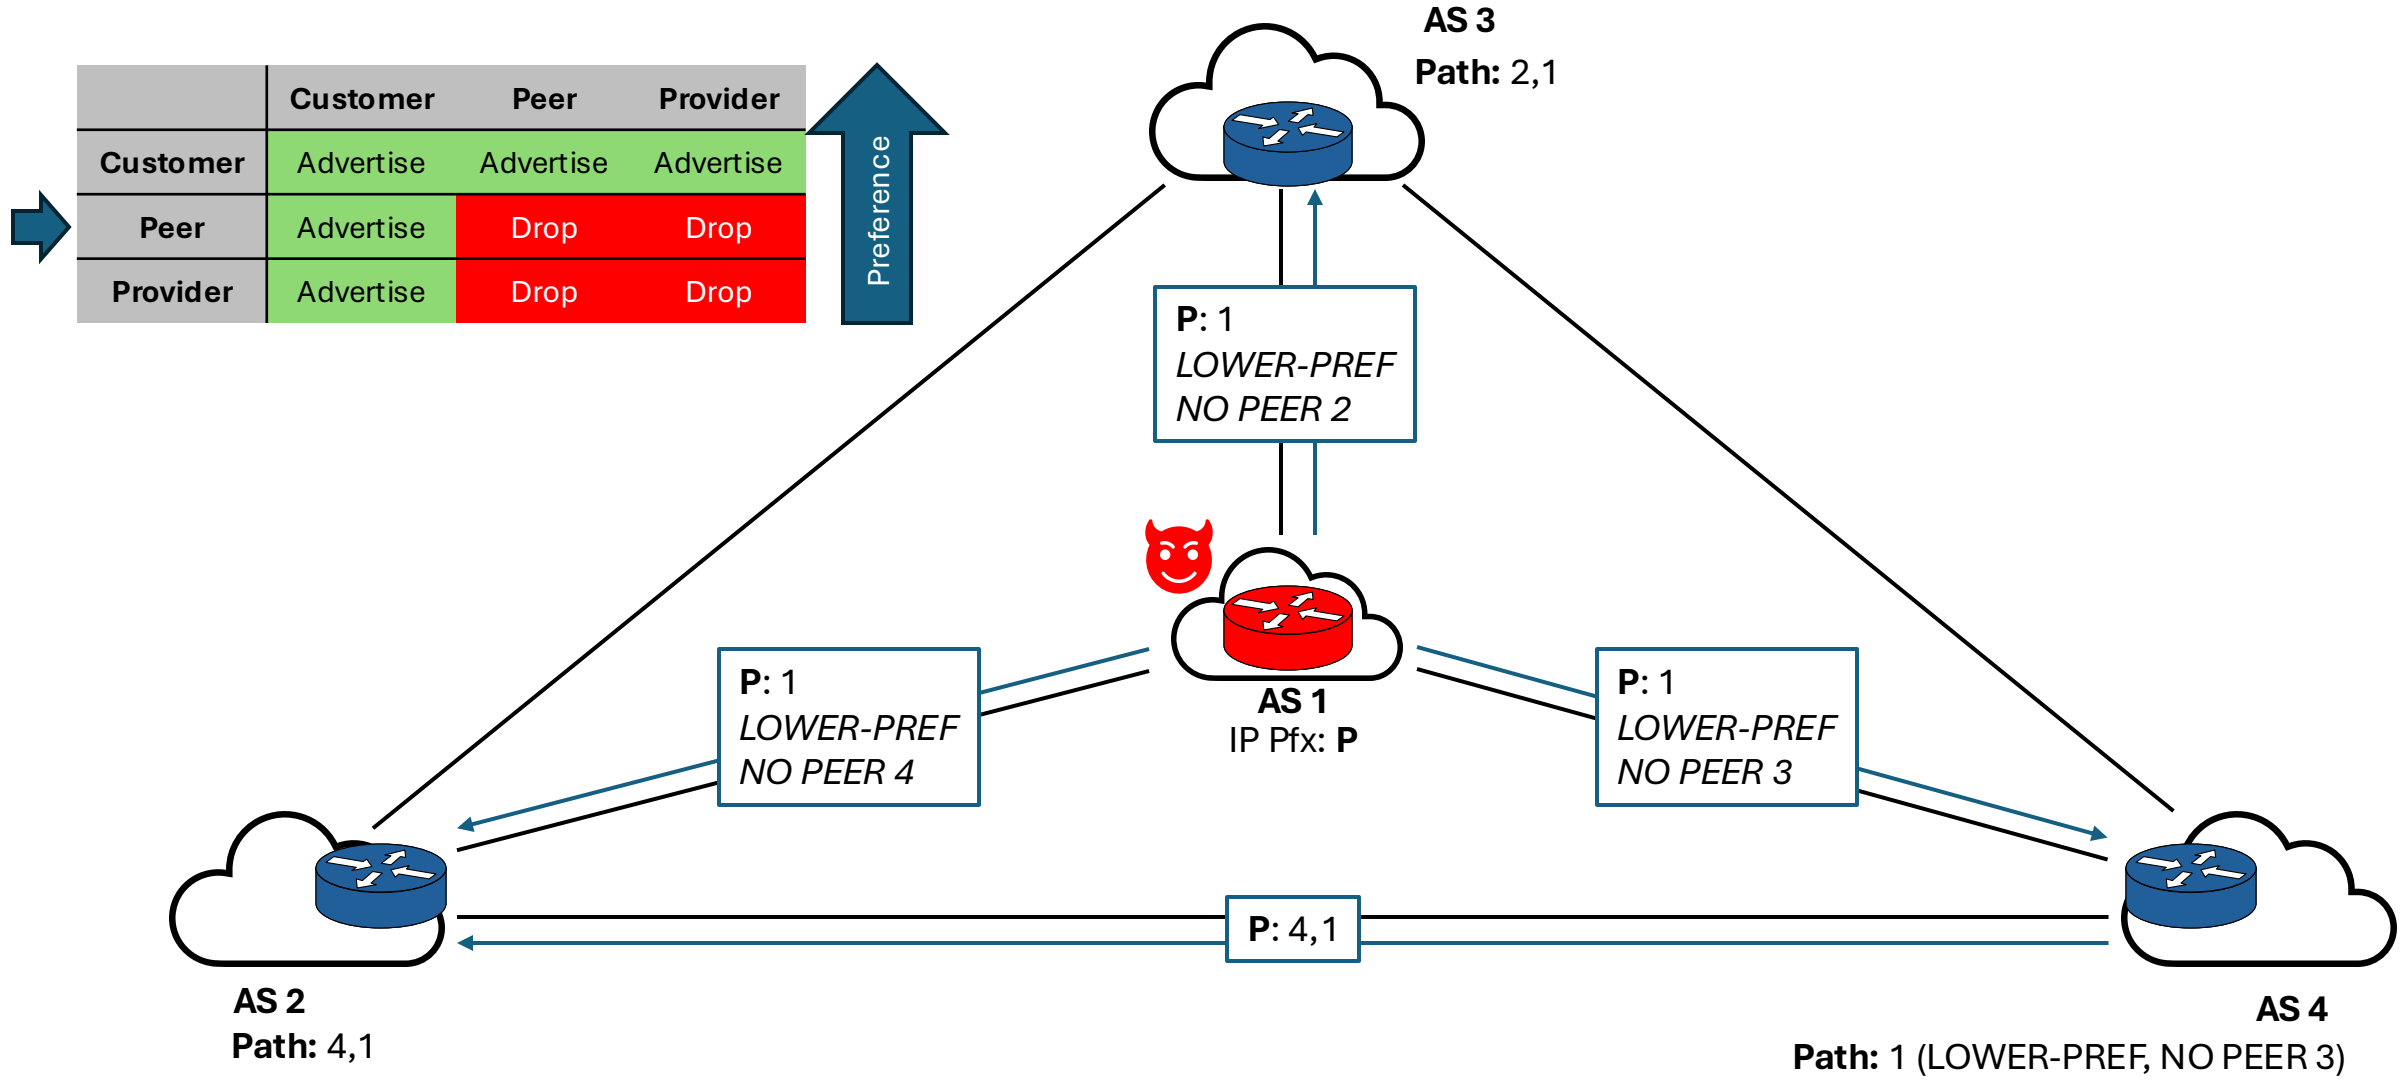

BGP Vortex: Inducing Persistent Oscillations

	Customer	Peer	Provider
Customer	Advertise	Advertise	Advertise
Peer	Advertise	Drop	Drop
Provider	Advertise	Drop	Drop

Preference

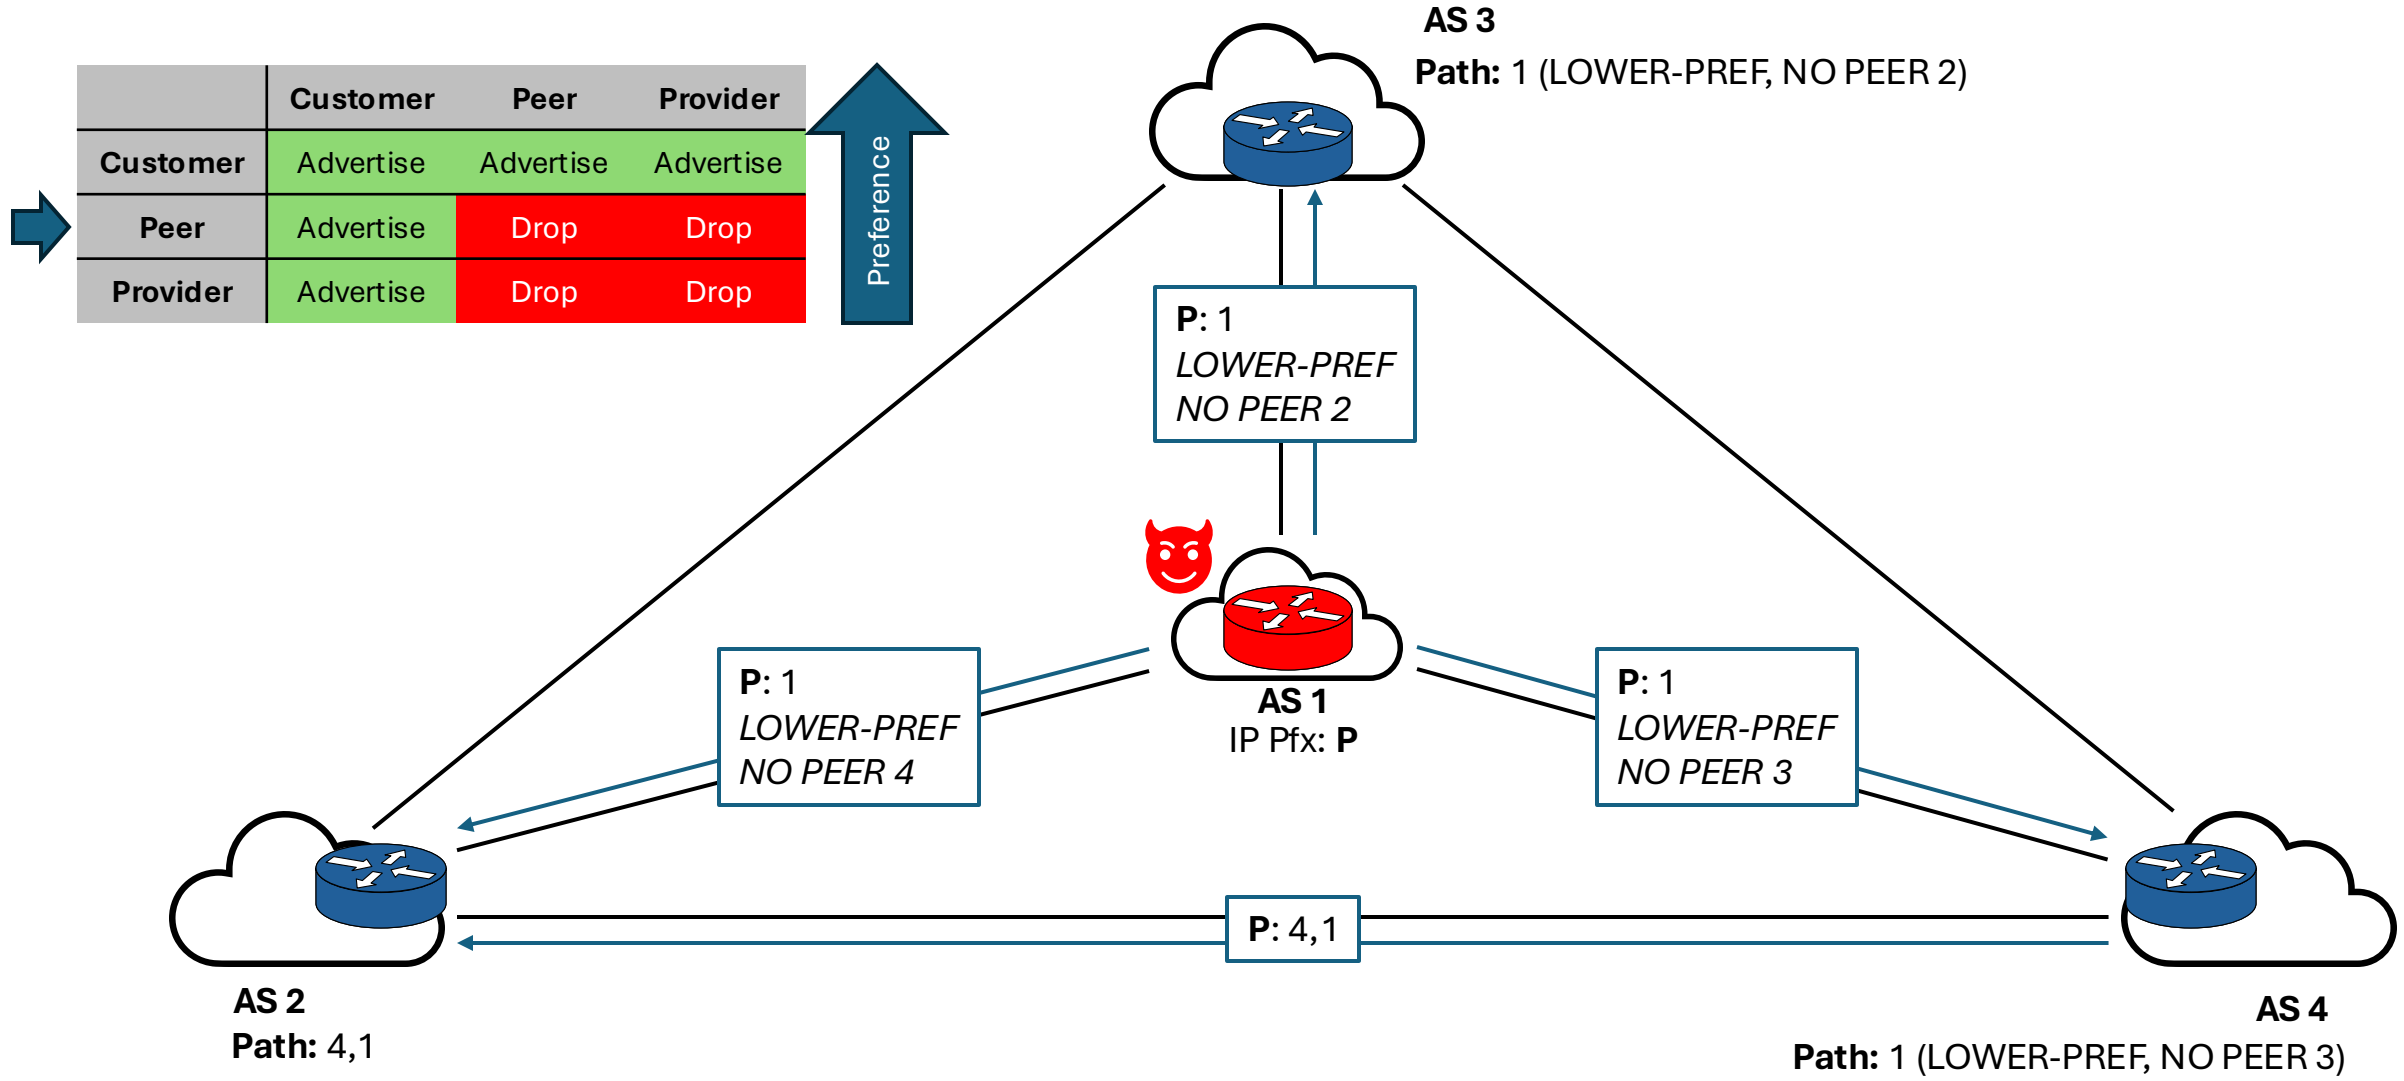

BGP Vortex: Inducing Persistent Oscillations

	Customer	Peer	Provider
Customer	Advertise	Advertise	Advertise
Peer	Advertise	Drop	Drop
Provider	Advertise	Drop	Drop

Preference

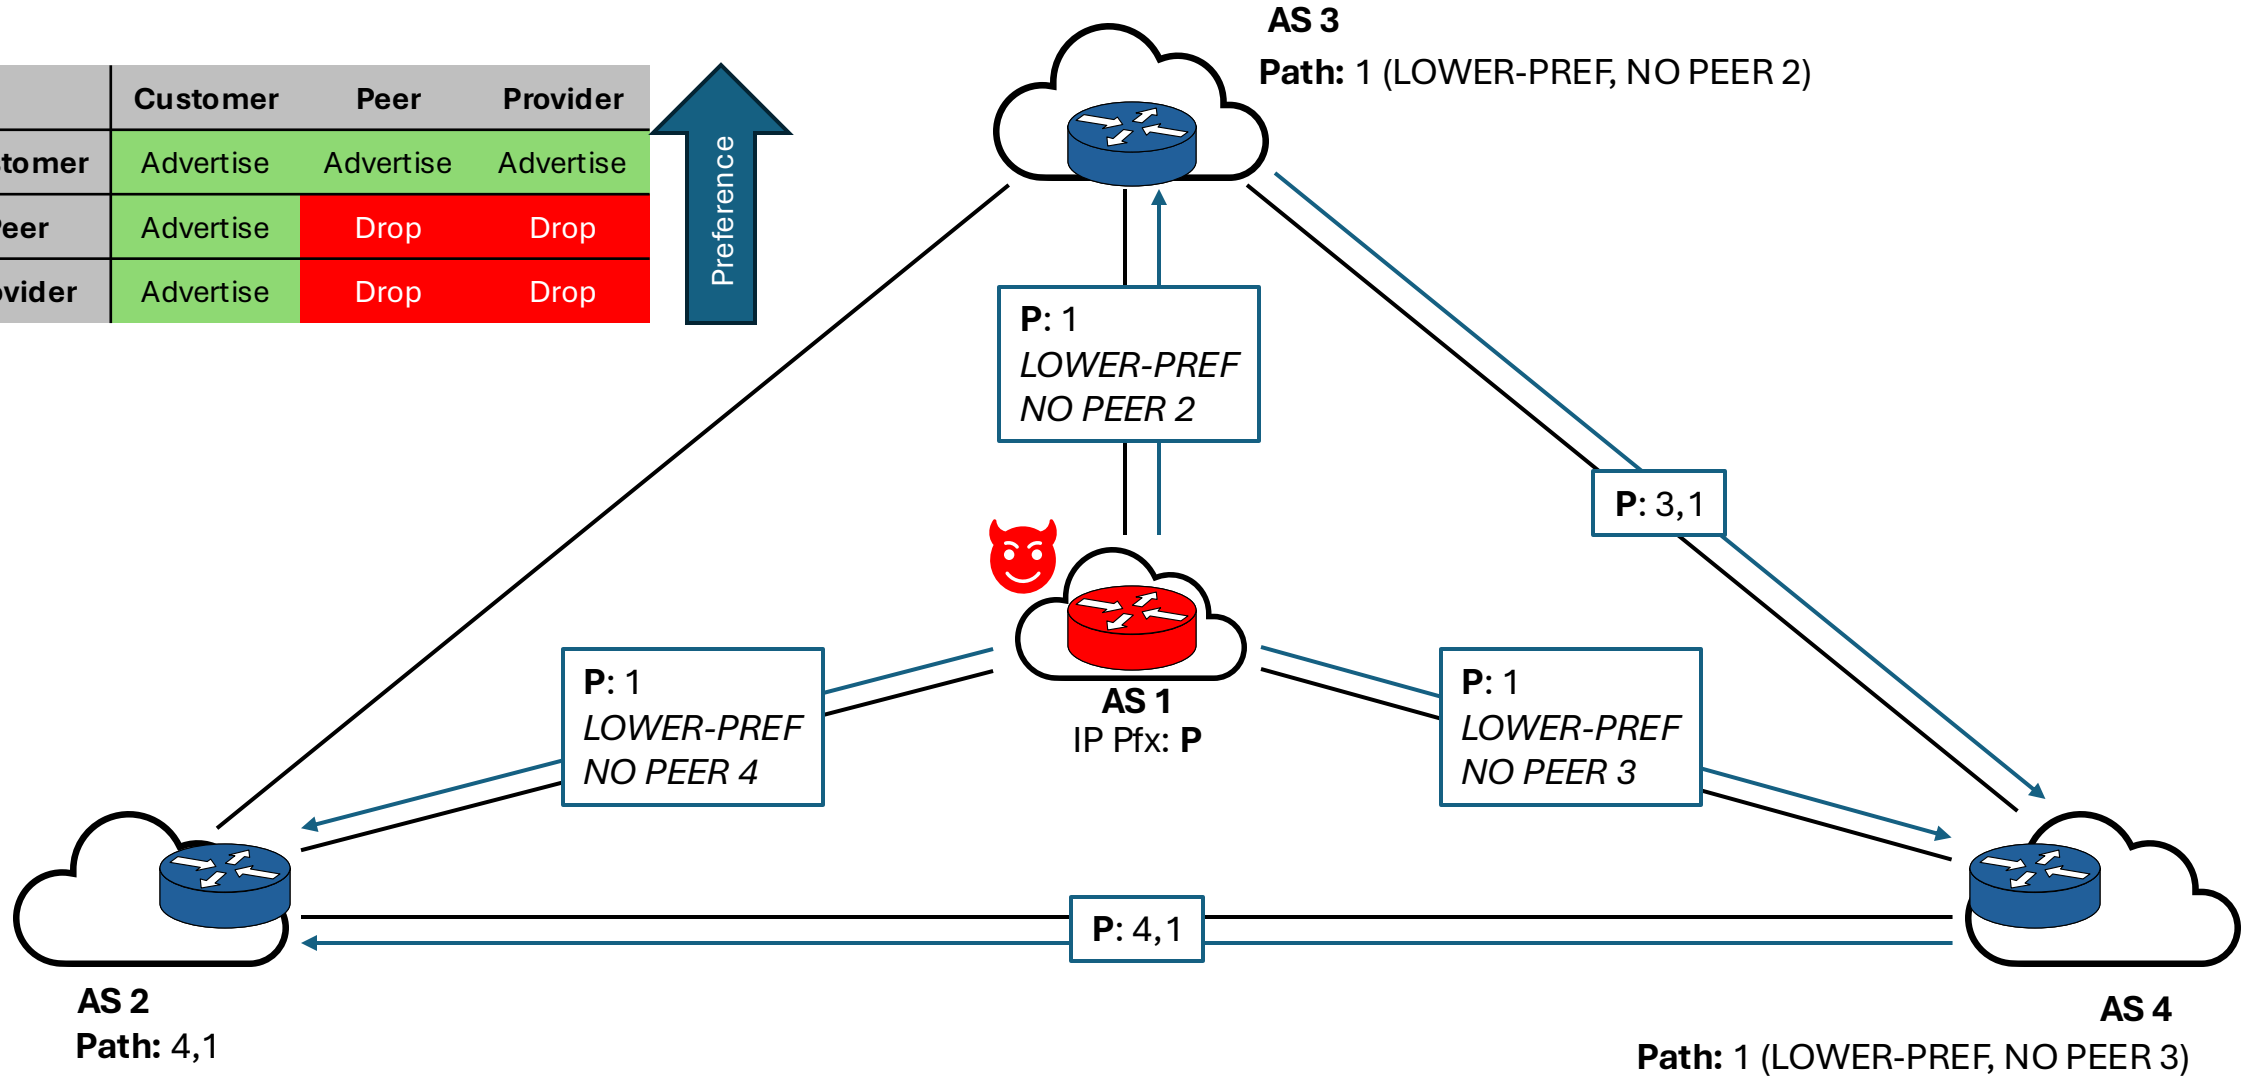

BGP Vortex: Inducing Persistent Oscillations

	Customer	Peer	Provider
Customer	Advertise	Advertise	Advertise
Peer	Advertise	Drop	Drop
Provider	Advertise	Drop	Drop

Preference

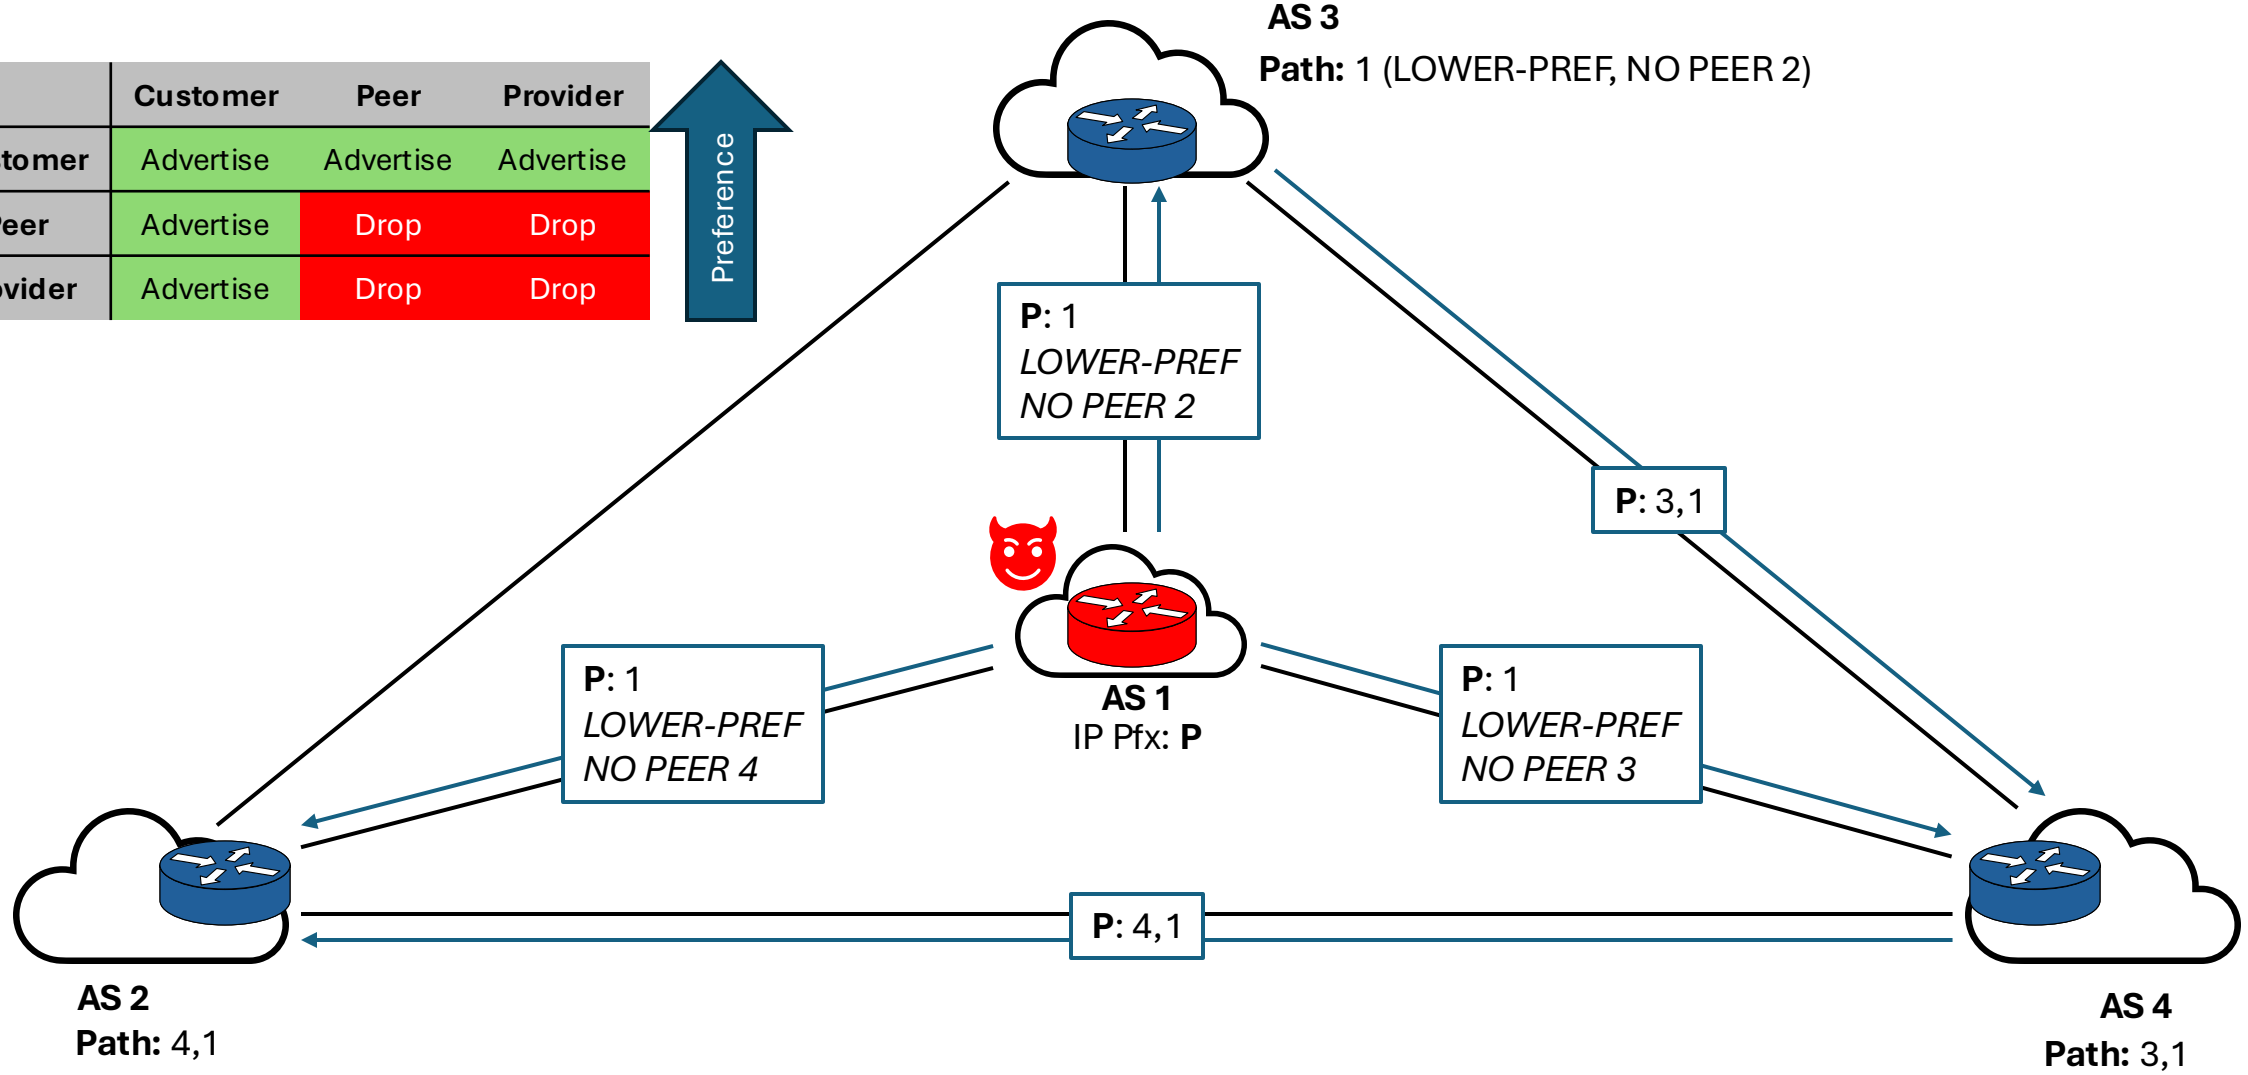

BGP Vortex: Inducing Persistent Oscillations

	Customer	Peer	Provider
Customer	Advertise	Advertise	Advertise
Peer	Advertise	Drop	Drop
Provider	Advertise	Drop	Drop

Preference

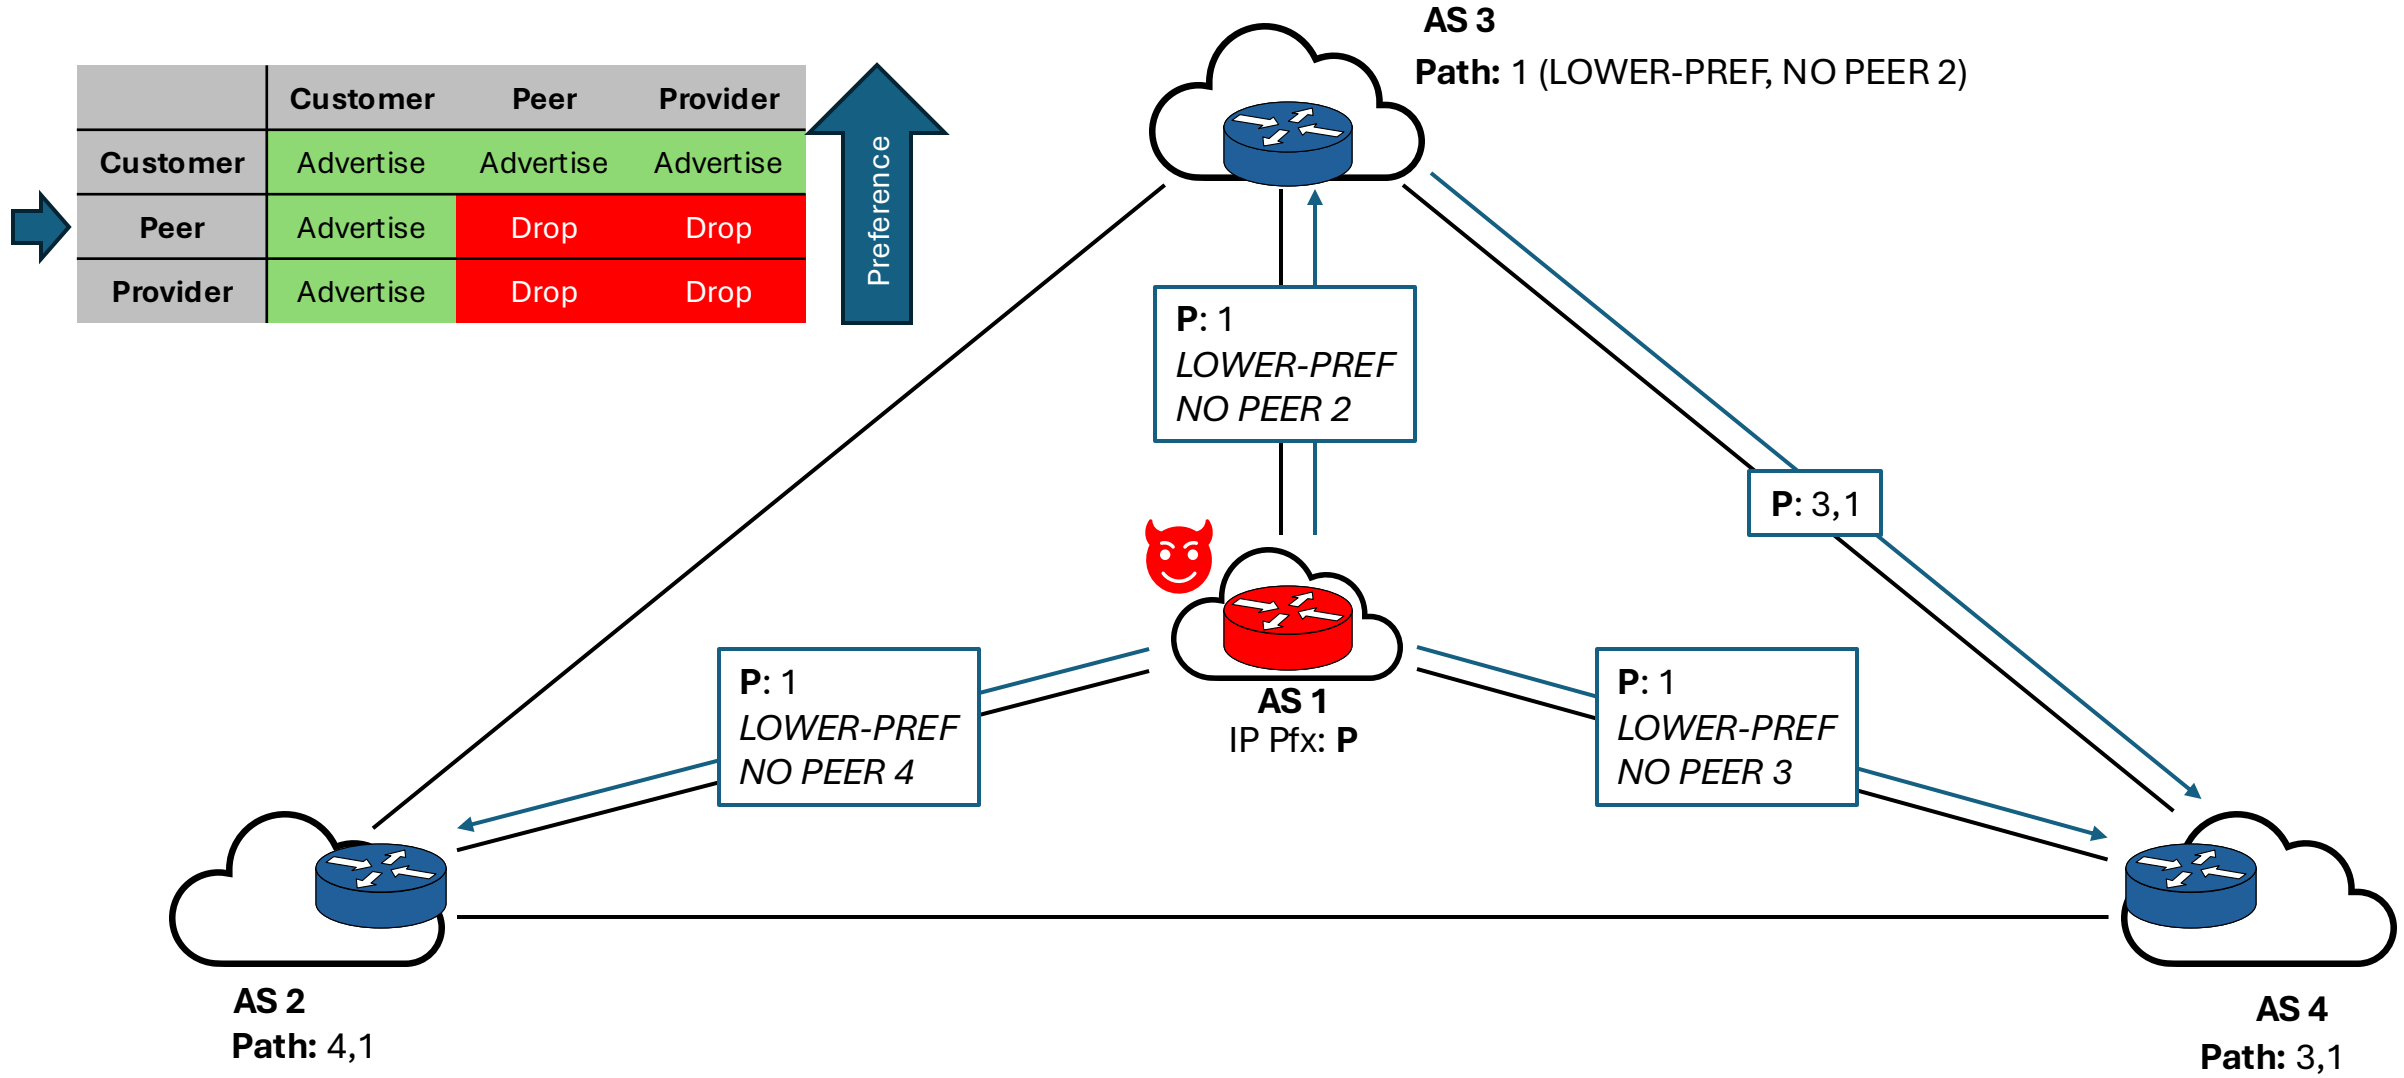

BGP Vortex: Inducing Persistent Oscillations

	Customer	Peer	Provider
Customer	Advertise	Advertise	Advertise
Peer	Advertise	Drop	Drop
Provider	Advertise	Drop	Drop

Preference

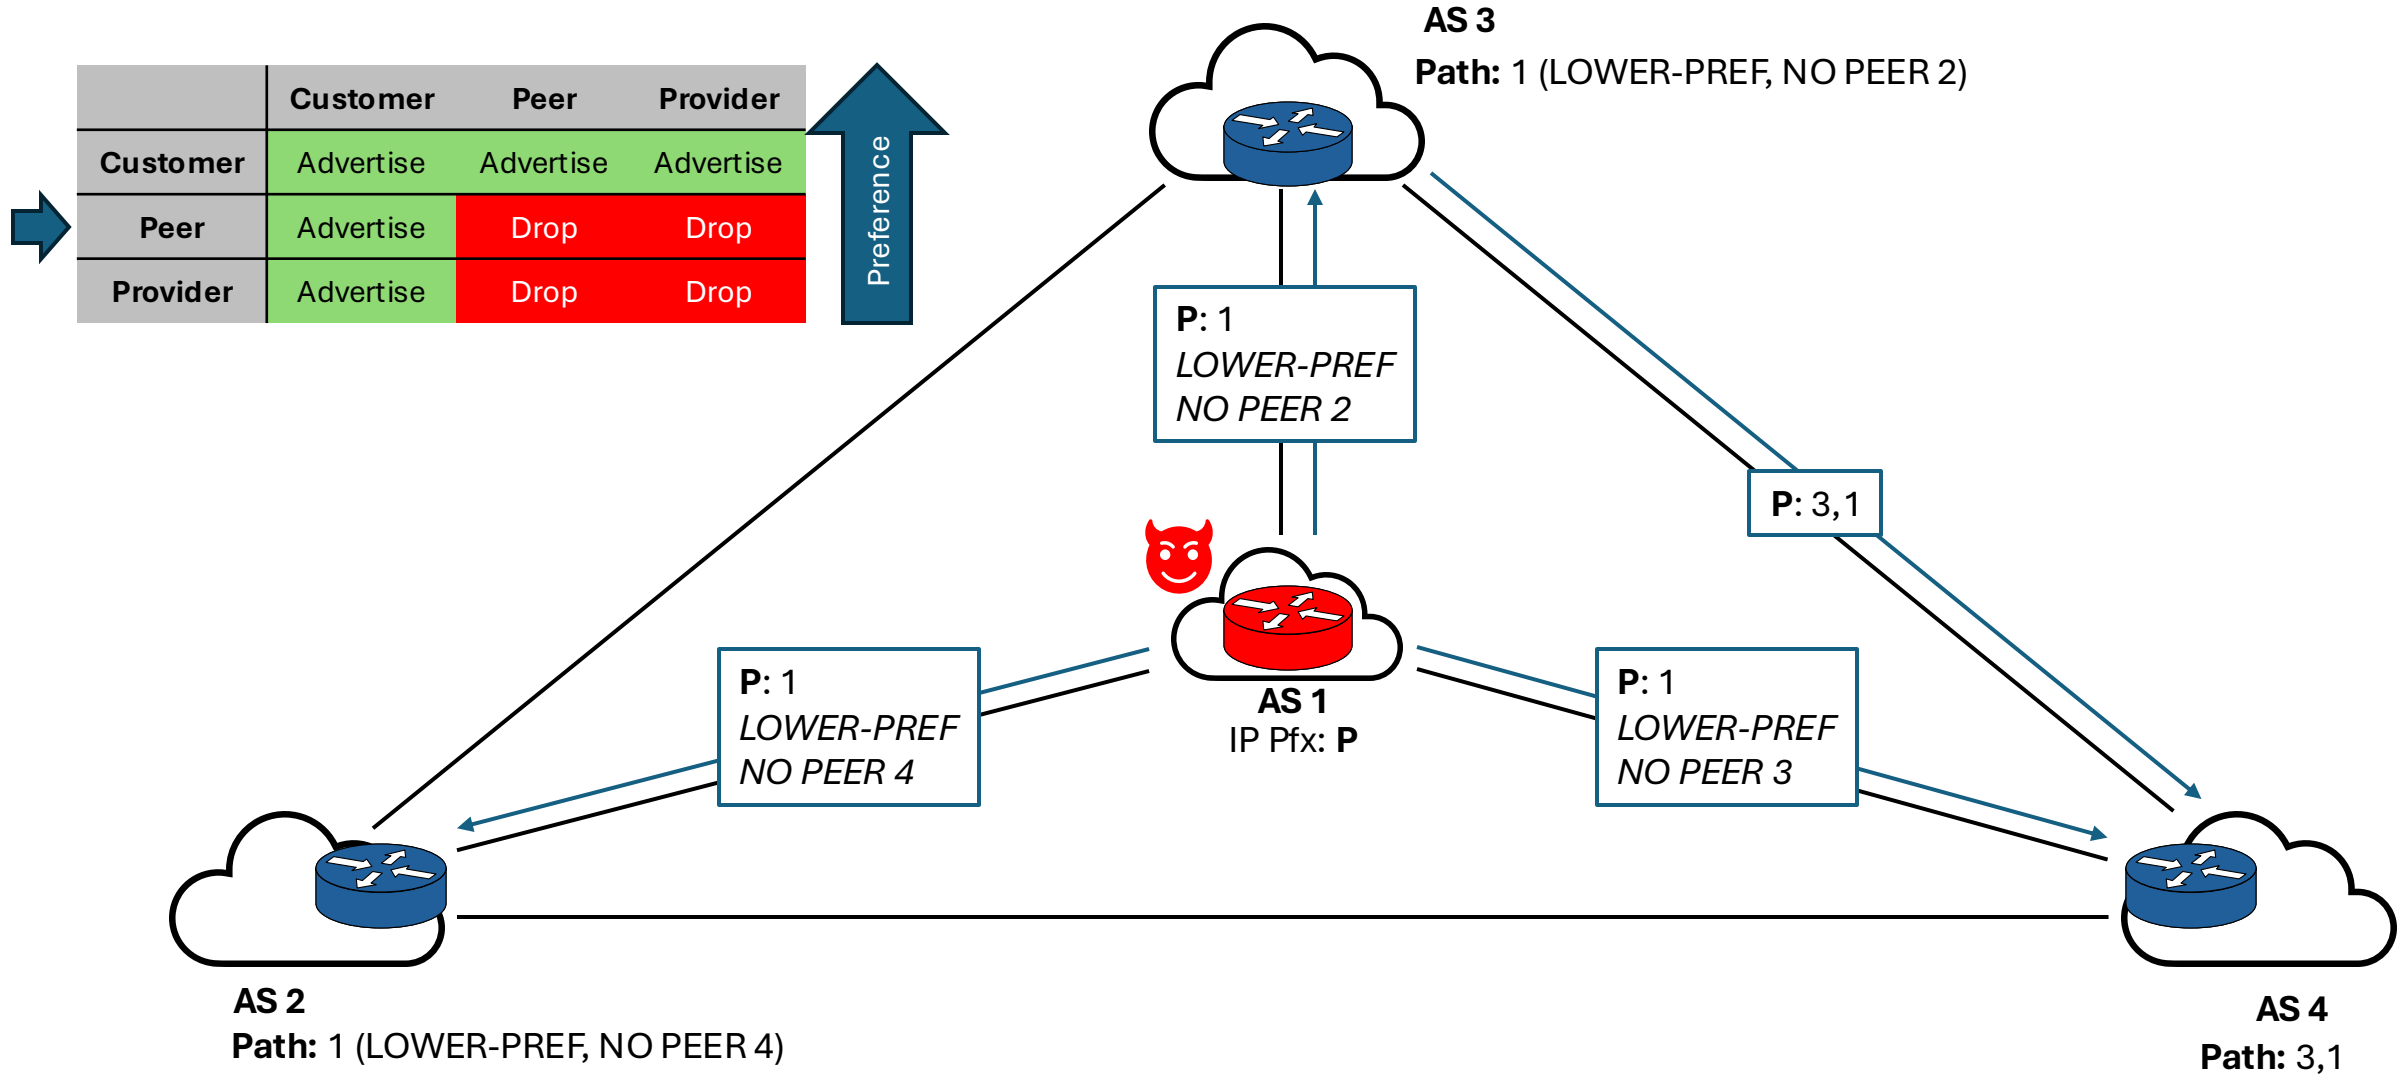

BGP Vortex: Inducing Persistent Oscillations

	Customer	Peer	Provider
Customer	Advertise	Advertise	Advertise
Peer	Advertise	Drop	Drop
Provider	Advertise	Drop	Drop

Preference

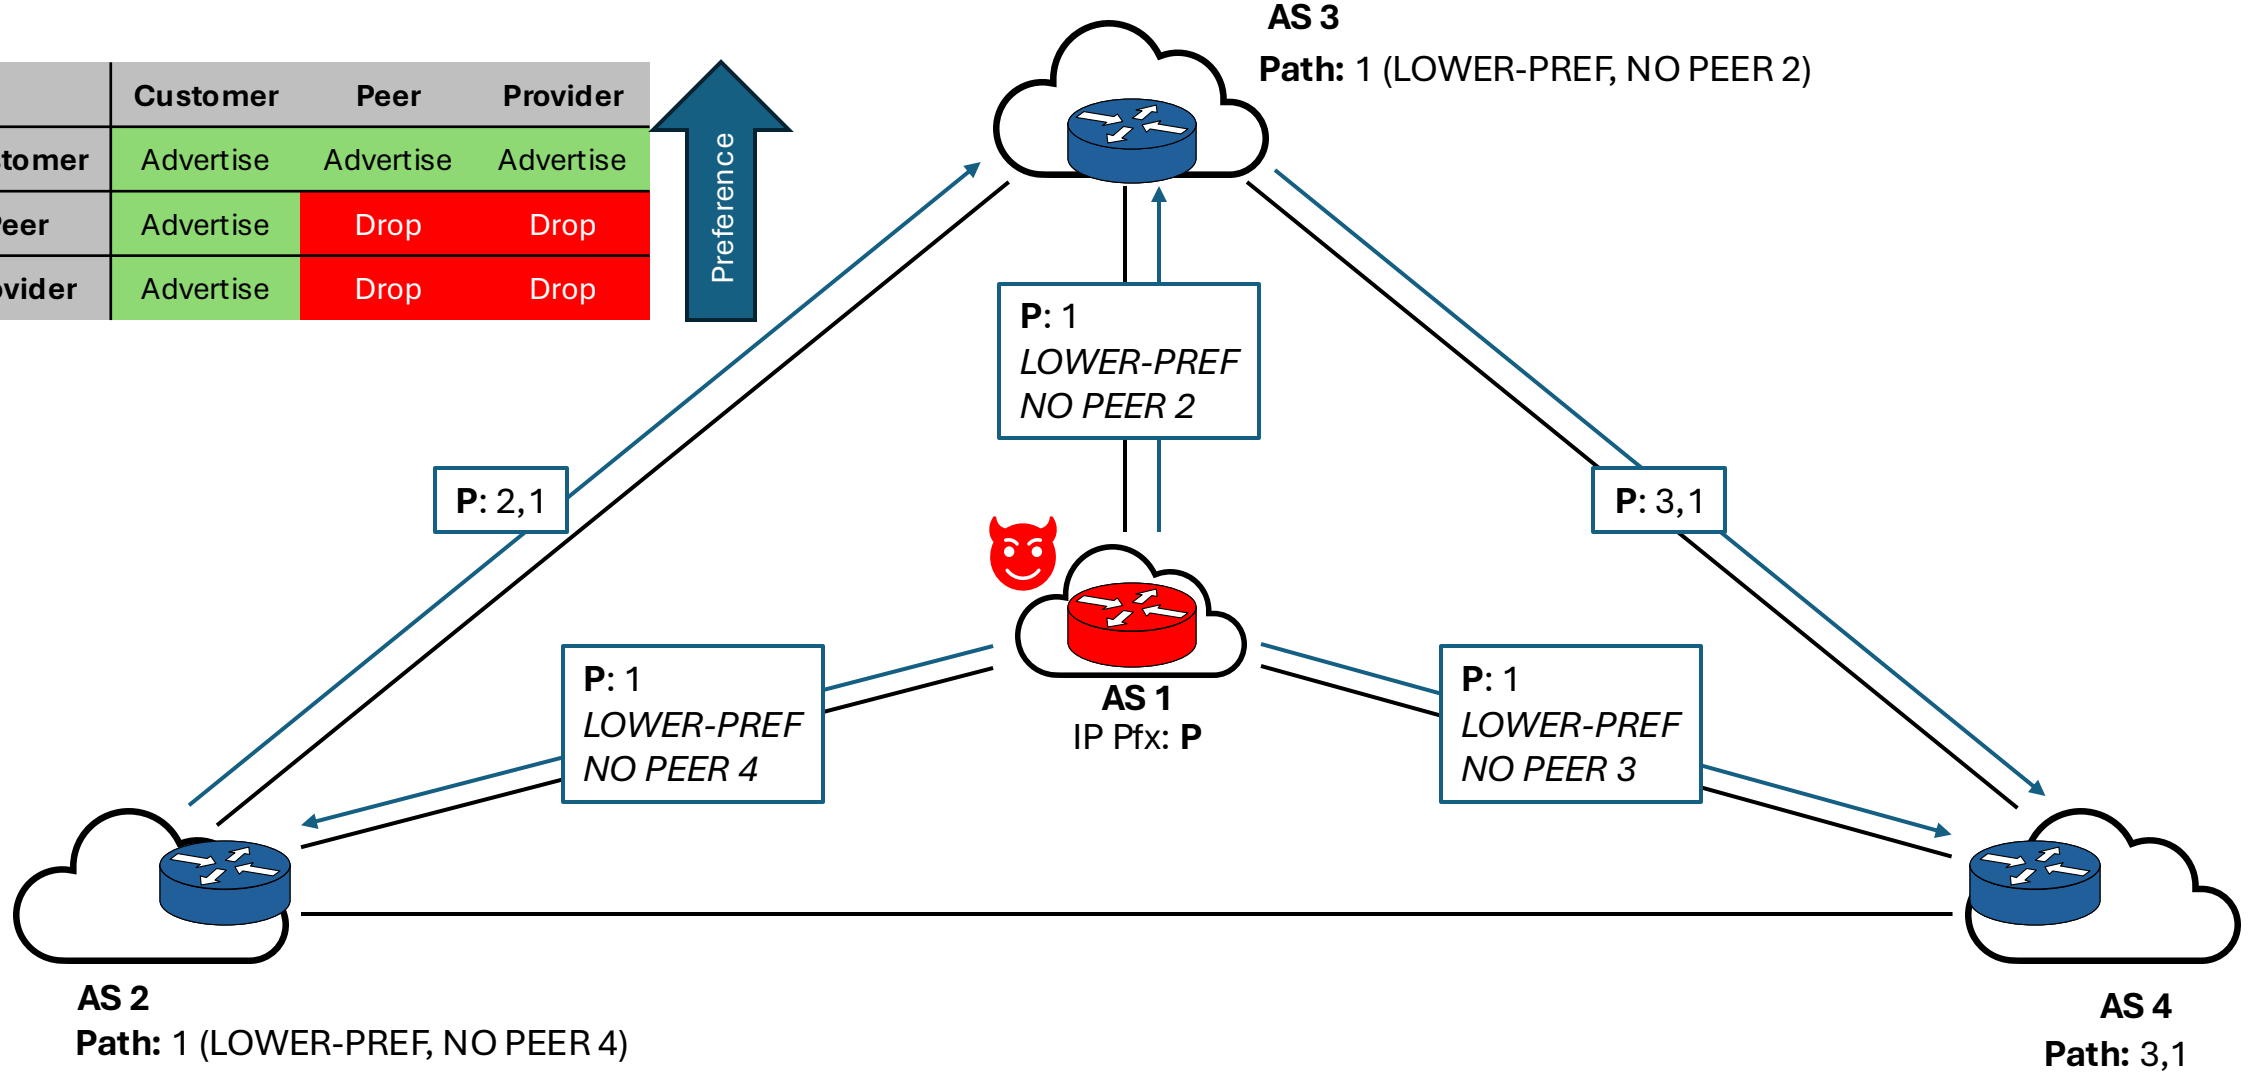

BGP Vortex: Inducing Persistent Oscillations

	Customer	Peer	Provider
Customer	Advertise	Advertise	Advertise
Peer	Advertise	Drop	Drop
Provider	Advertise	Drop	Drop

Preference

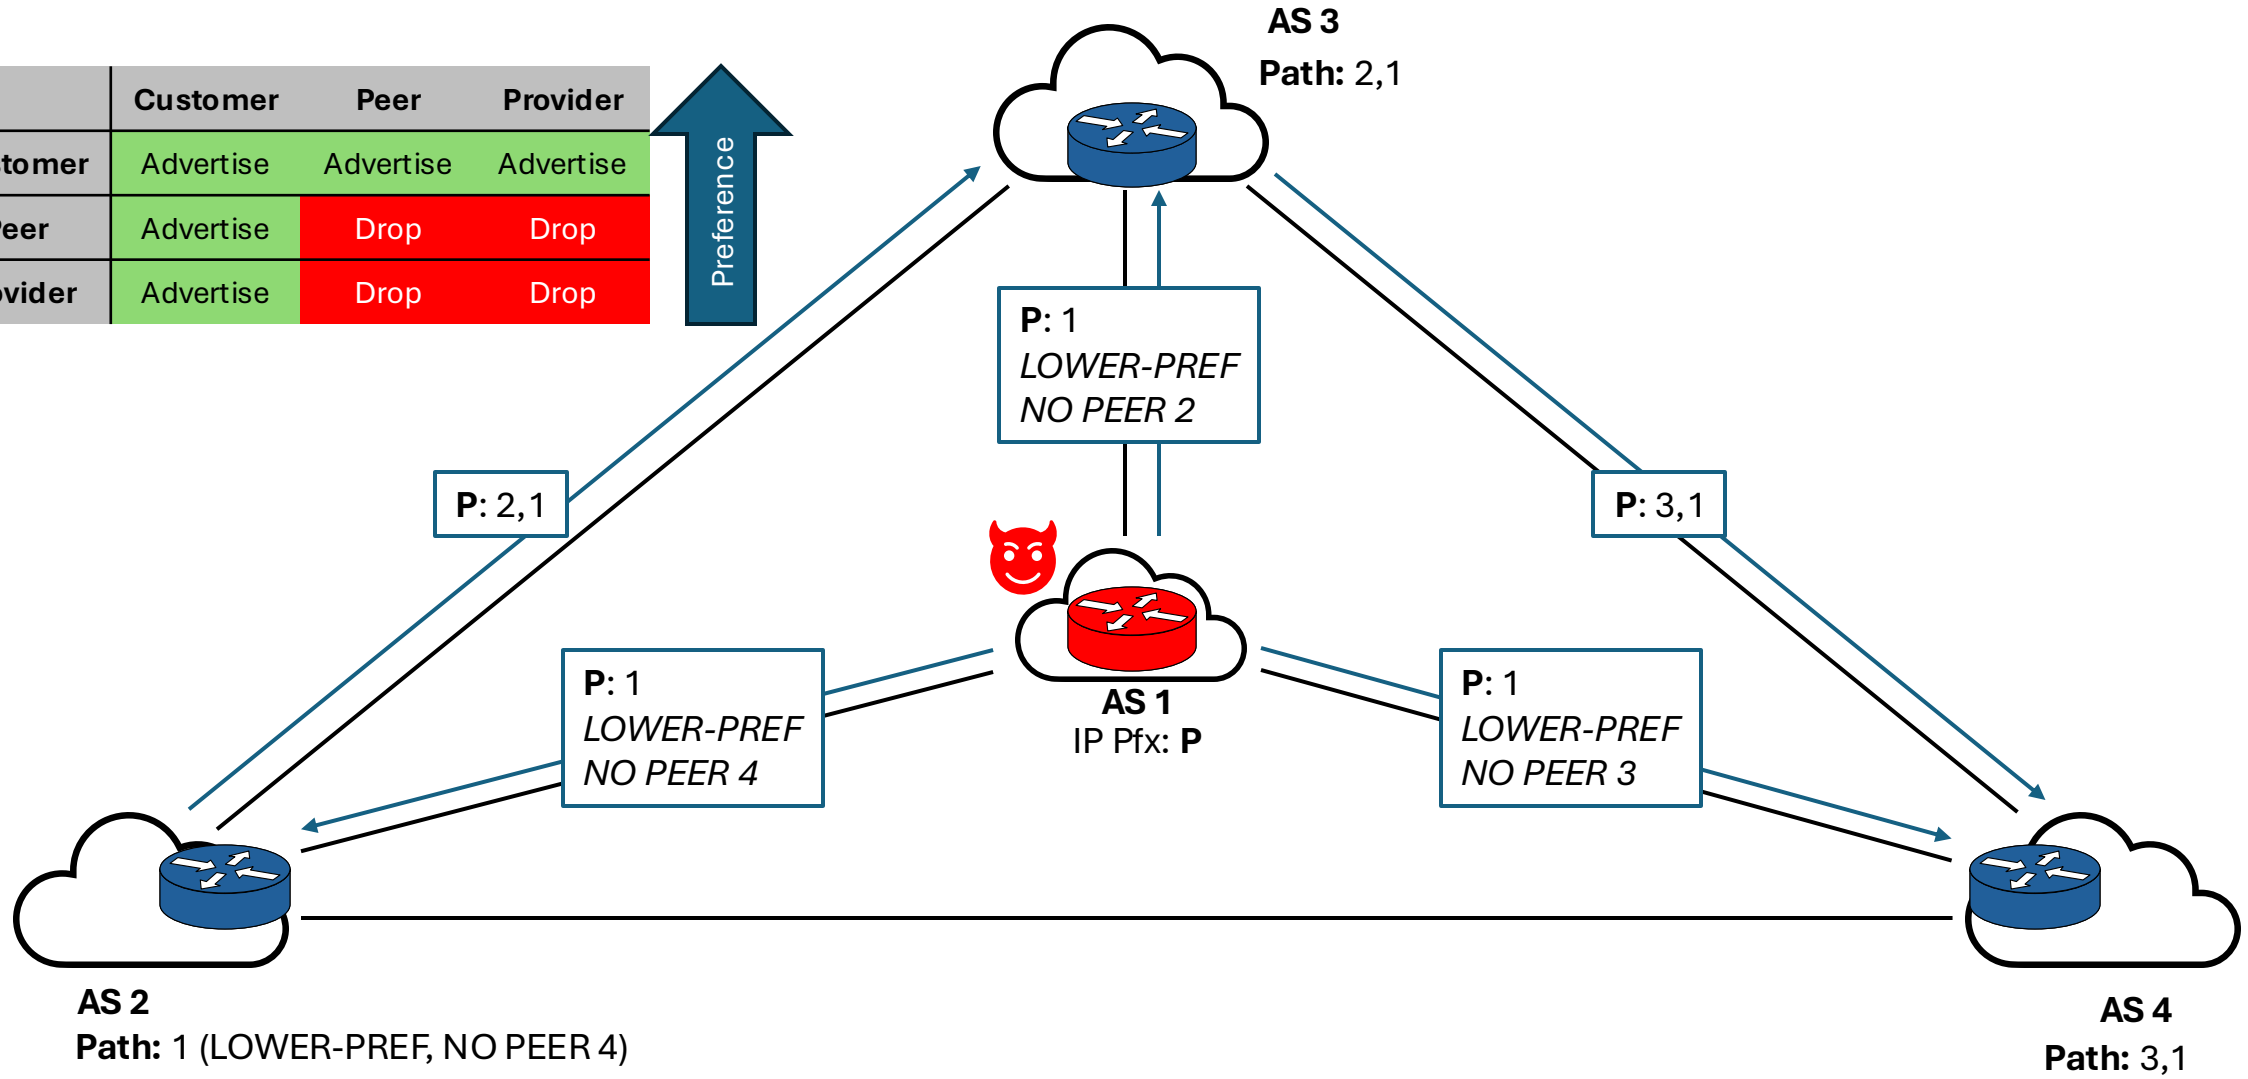

BGP Vortex: Inducing Persistent Oscillations

	Customer	Peer	Provider
Customer	Advertise	Advertise	Advertise
Peer	Advertise	Drop	Drop
Provider	Advertise	Drop	Drop

Preference

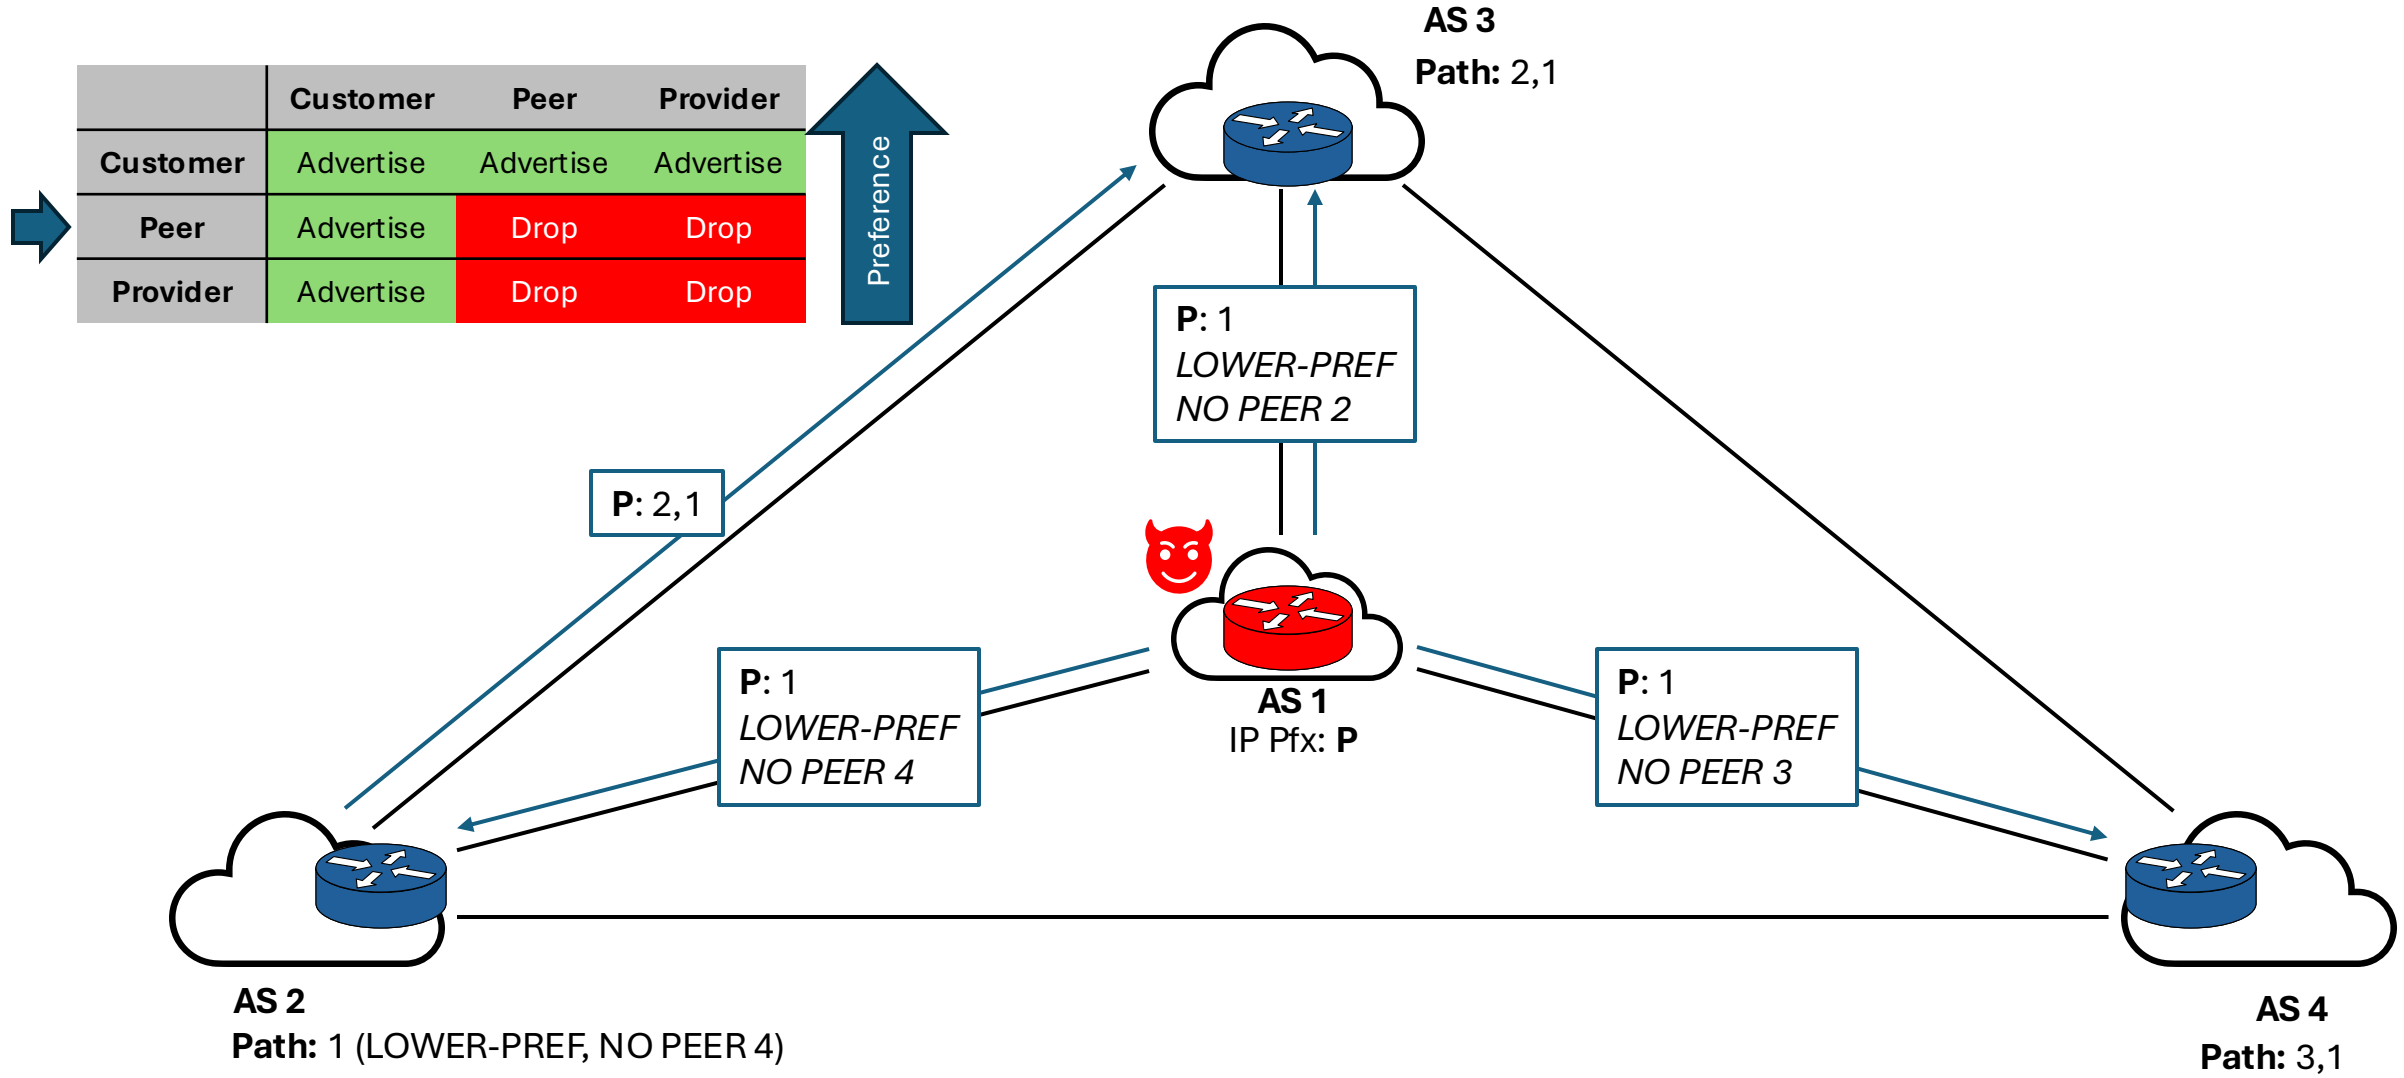

BGP Vortex Causes Collateral Damage

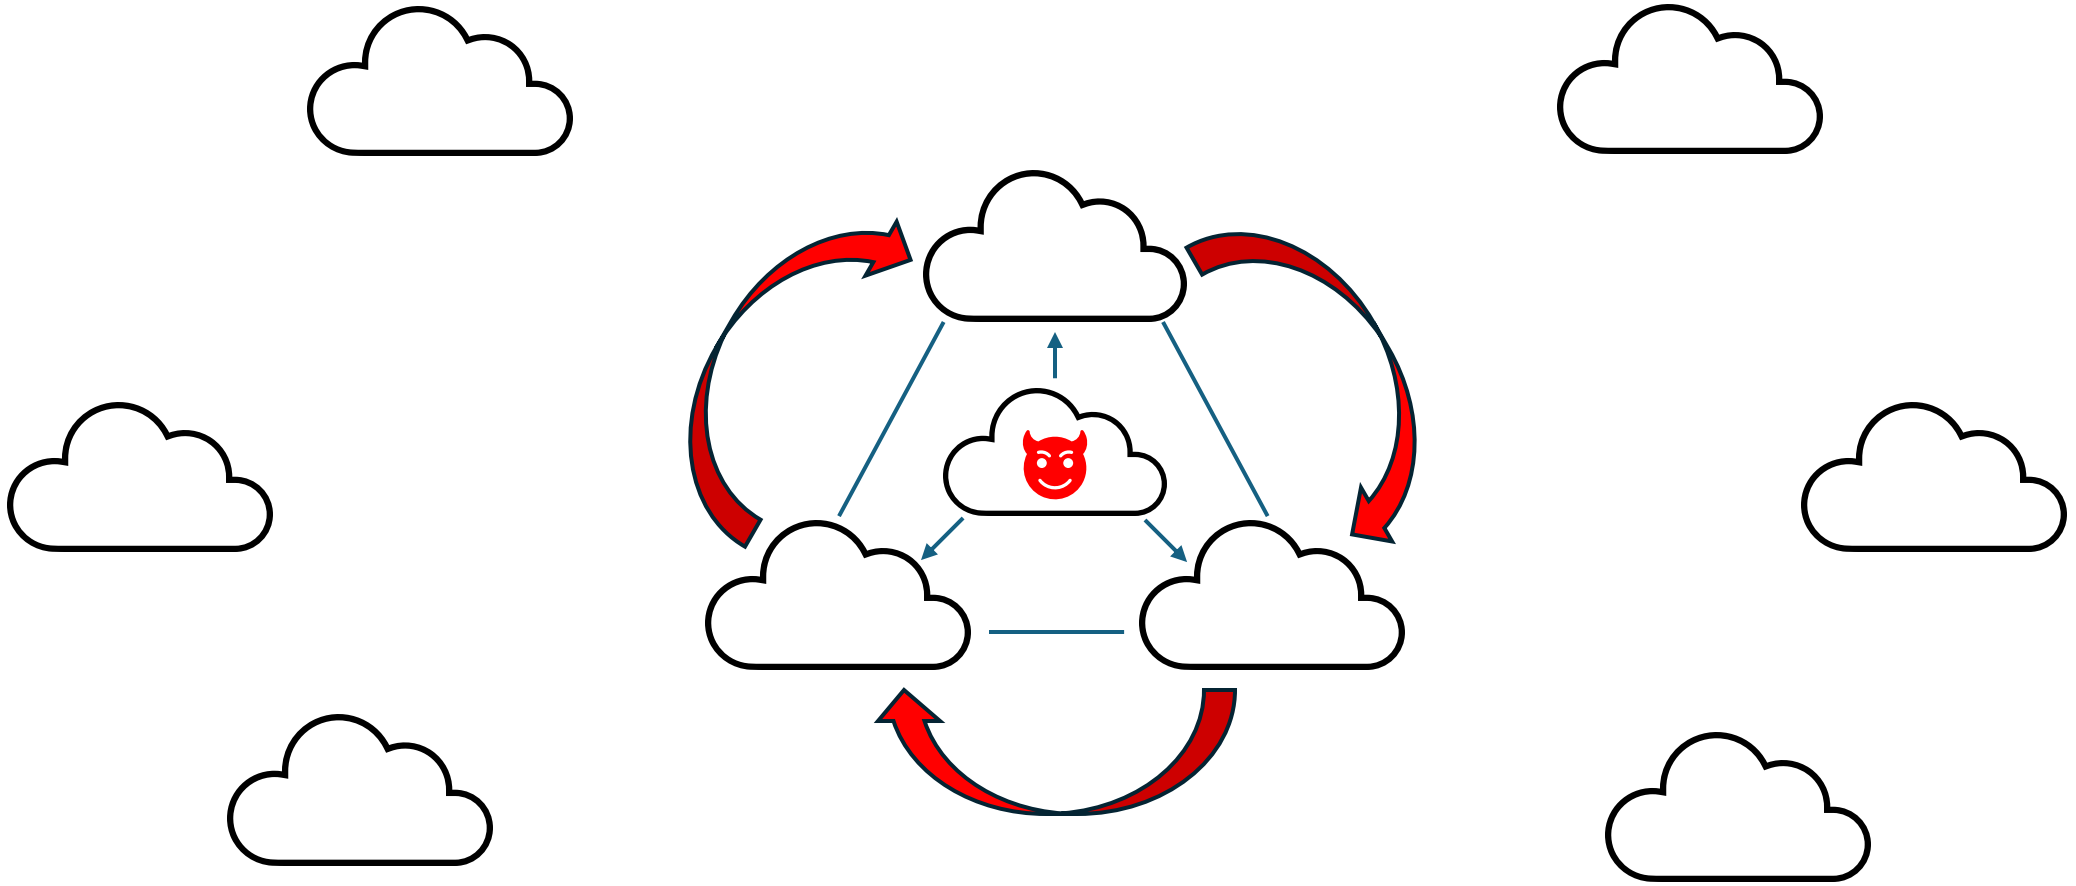

BGP Vortex Causes Collateral Damage

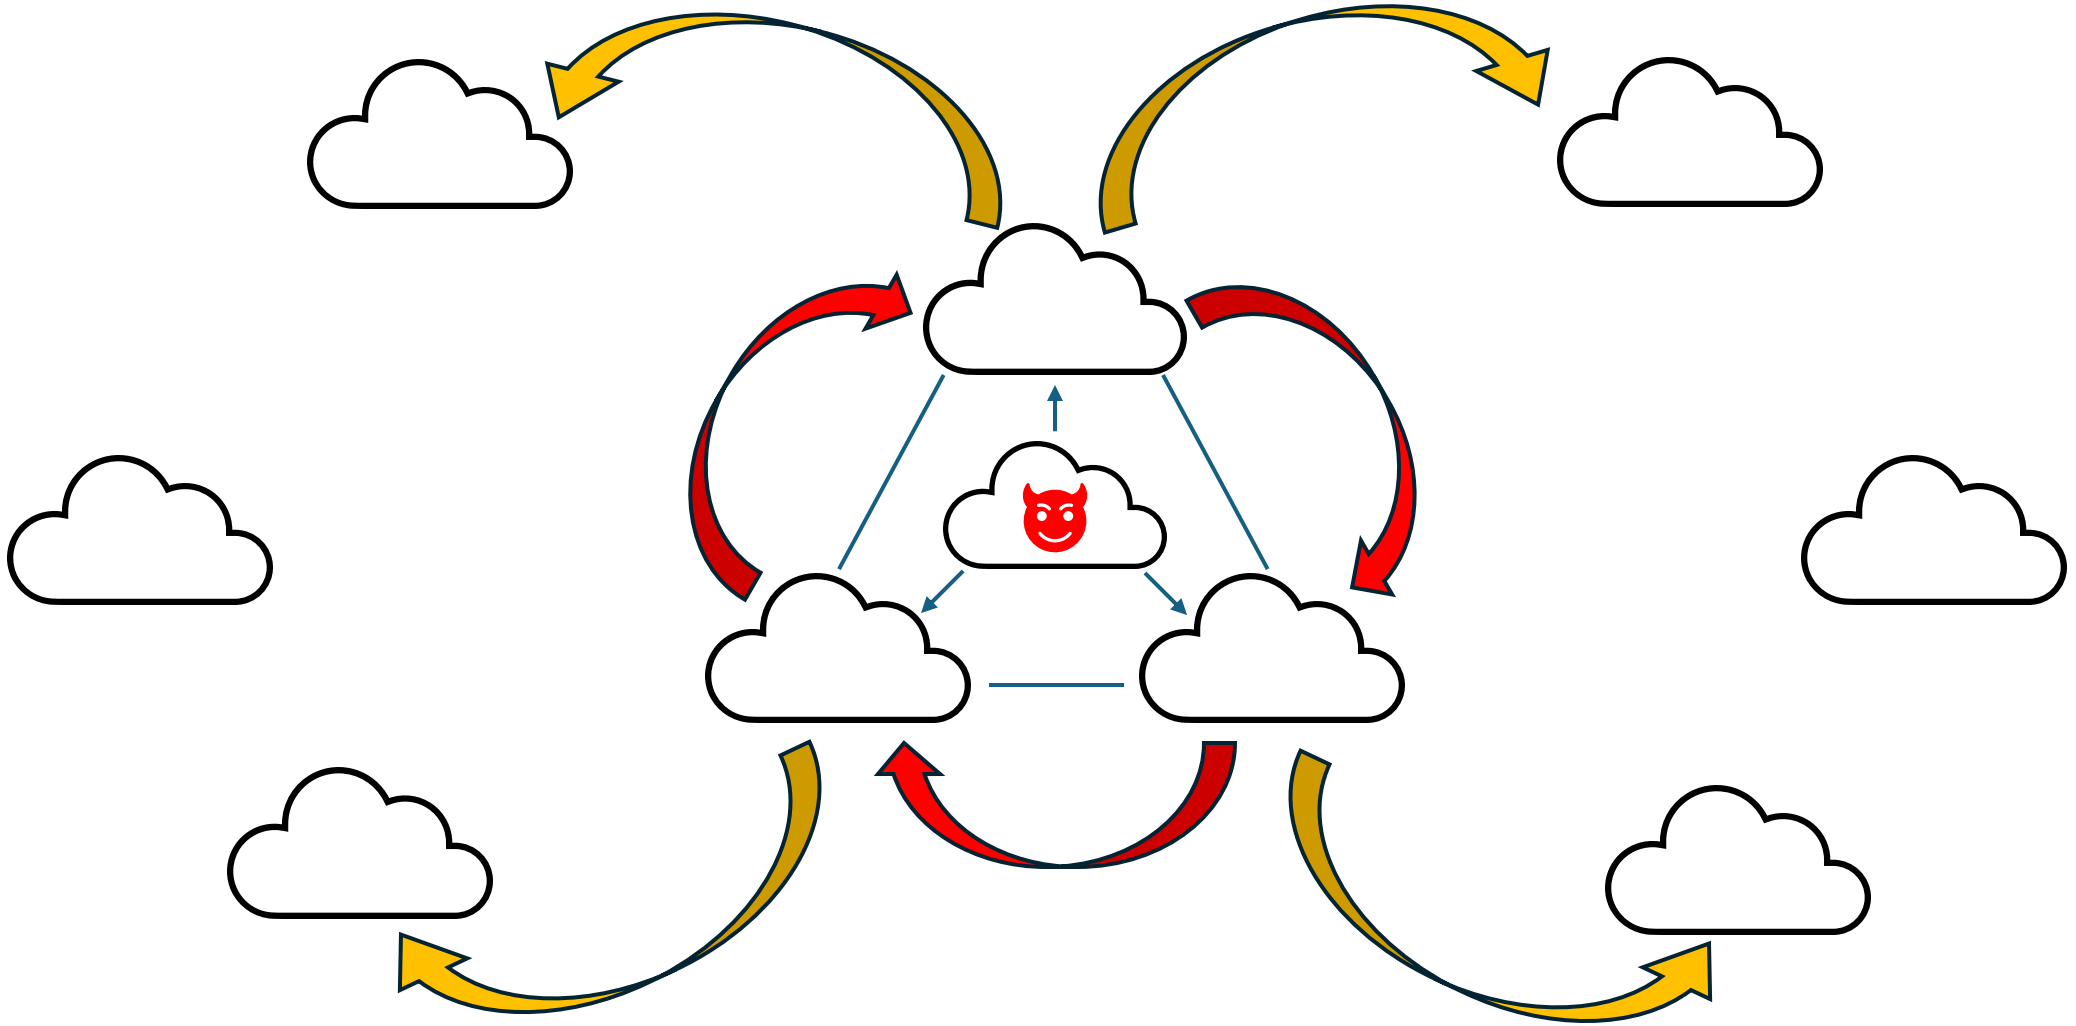

BGP Vortex Causes Collateral Damage

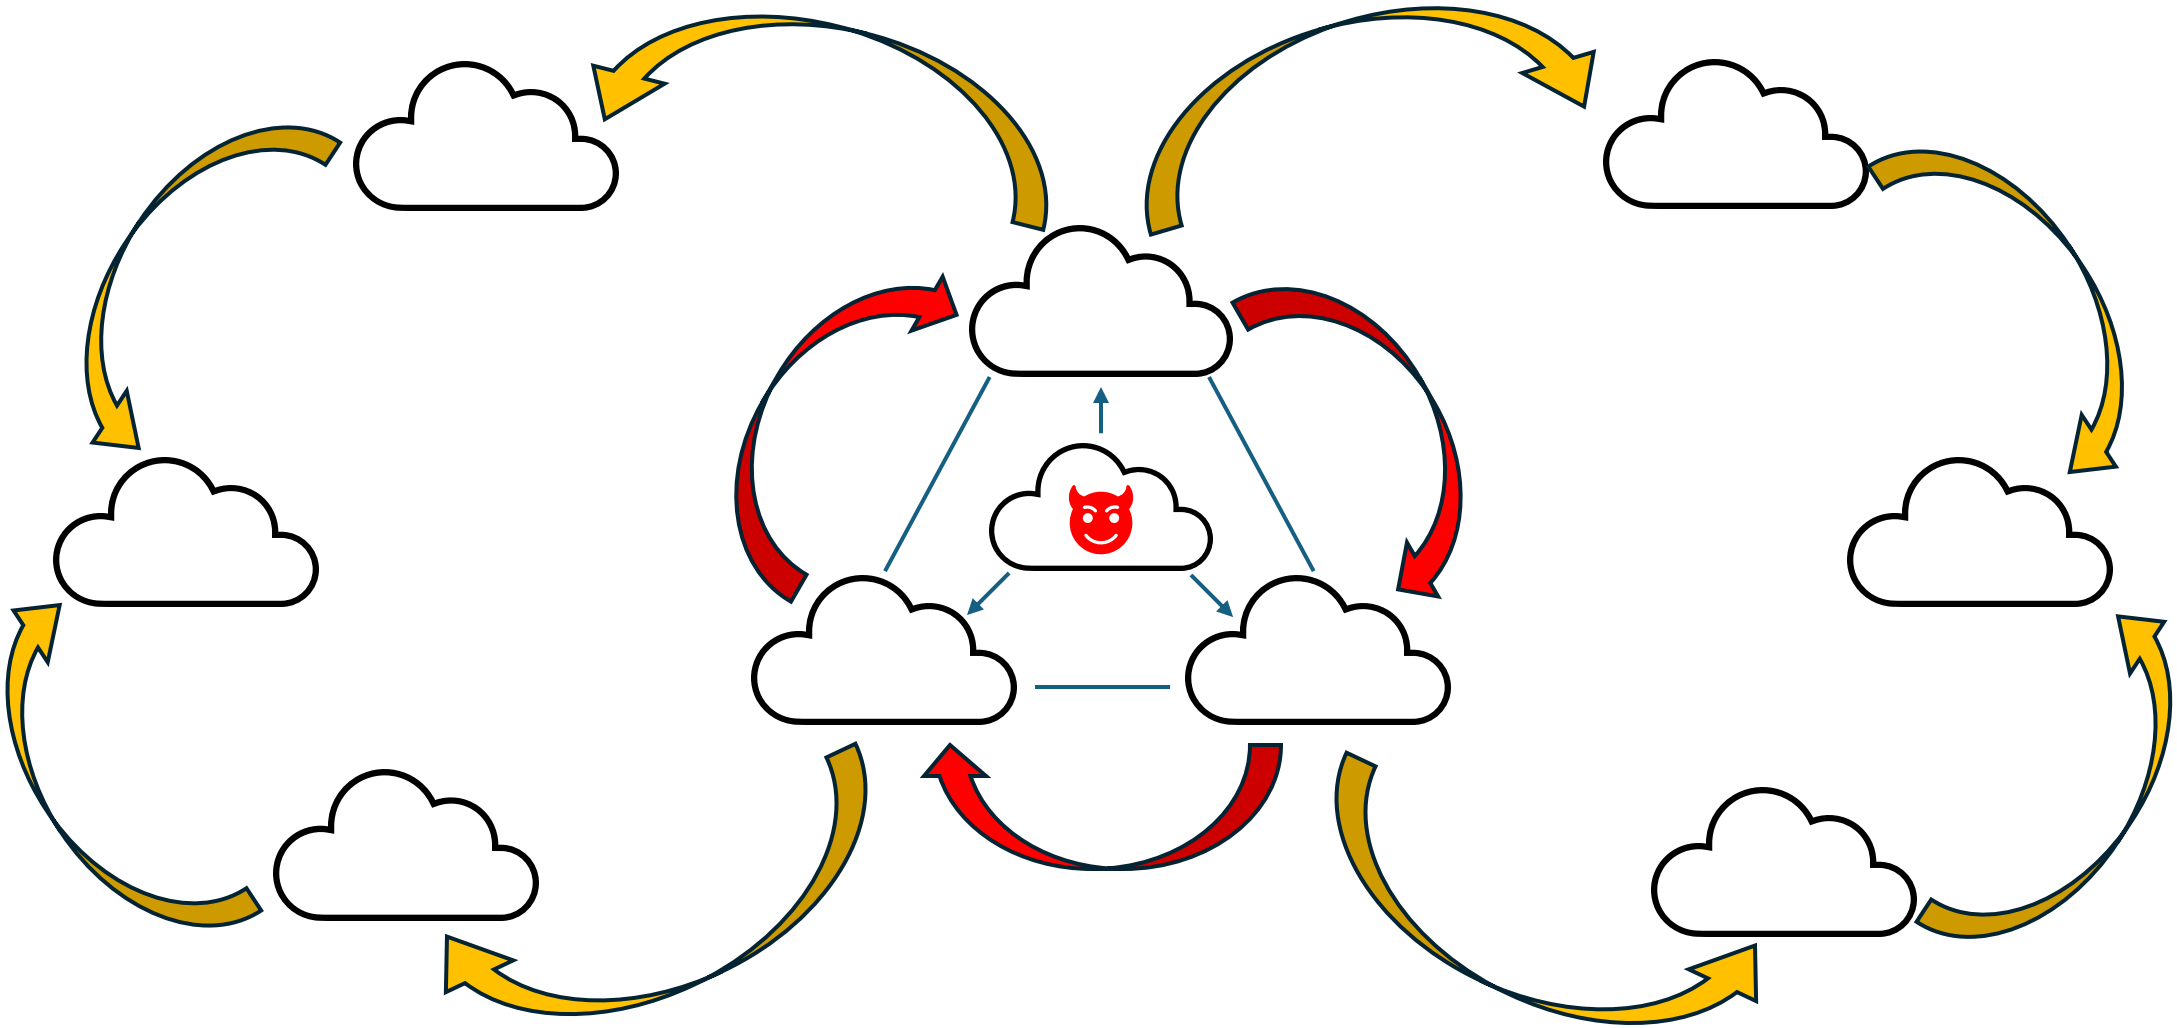

Measurements

Show the Prevalence of BGP Vortices in the Internet

Statically Identifying BGP Vortices

Statically Identifying BGP Vortices

=

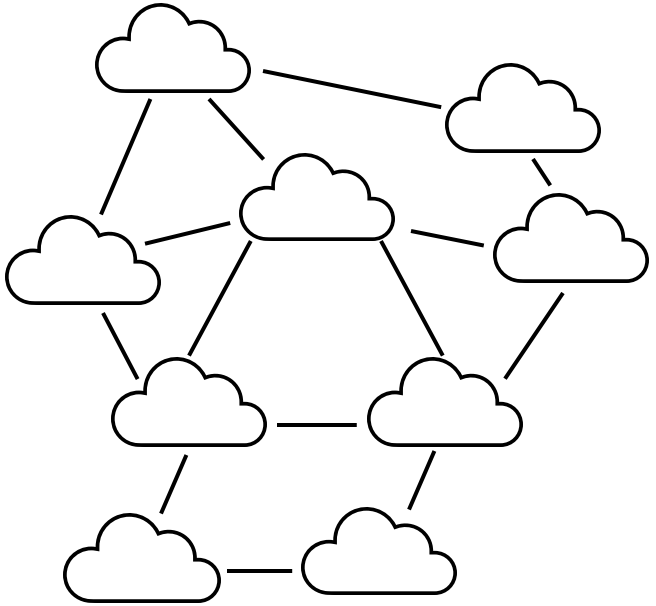

Statically Identifying BGP Vortices

+

ASN(Name)	LowerPref	NoExportSelect	NoExportAll
3356 (Level3)	Yes	Yes	Yes
1299 (Telia)	Yes	Yes*	Yes
174 (Cogent)	Yes	Yes**	Yes
2914 (NTT)	Yes	Yes	Yes
3257 (GTT)	Yes	Yes	Yes
6762 (Sparkle)	Yes	No	Yes***
6939 (Hurricane)	No	No	No
6453 (TATA)	Yes	Yes	Yes
3491 (PCCW)	Yes	Yes*	Yes
6461 (Zero)	Yes	Yes*	Yes
1273 (Vodafone)	Yes	Yes*	Yes
3549 (Level3)	Yes	Yes*	Yes
9002 (RETN)	Yes	Yes	Yes
12956 (Telefonica)	unknown	unknown	unknown
4637 (Telstra)	No	No	No
209 (CenturyLink)	Yes	Yes*	Yes
7473 (SINGTEL)	unknown	unknown	unknown
12389 (Rostelecom)	Yes	Yes*	Yes
20485 (TransTeleCom)	No	Yes*	Yes
3320 (Deutsche)	Yes	Yes	Yes
701 (MCI)	Yes	No	Yes
7018 (AT&T)	Yes	No	Yes
7922 (Comcast)	Yes	Yes	Yes
5511 (Orange)	Yes	Yes*	Yes
8359 (MTS)	No	Yes*	Yes
3216 (Vimpelcom)	Yes	Yes*	Yes
2828 (MCI)	Yes	Yes*	Yes
31133 (MegaFon)	Yes	Yes*	Yes
286 (KPN)	Yes	Yes	Yes
20764 (RASCOM)	Yes	Yes*	Yes

=

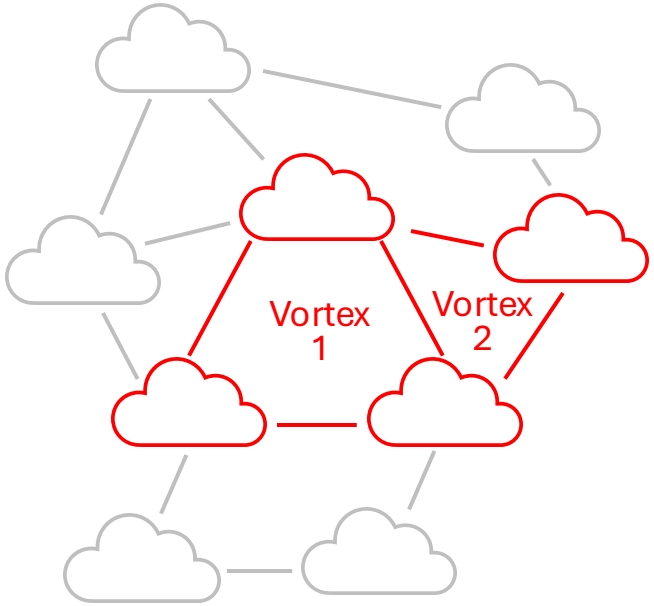

Statically Identifying BGP Vortices

21/30

ASes Support Both Vulnerable BGP Communities

20/21

ASes Contained in BGP Vortices

340

Viable BGP Vortices (30.4% of Peering Triangles)

58

BGP Vortices in Median

Collateral Damage Caused by BGP Vortices

>96%

of ASes are in the Customer Cone of at Least One BGP Vortex

801*

Route Updates Received by an AS in the Customer Cone in the Median Per Period

only a single prefix circulating in each BGP Vortex

6.6-40

Periods per Second of a BGP Vortex.

5281-32040**

BGP Updates per second

2.3 BGP Updates per second during normal operations

Experiments

Testing the Impact of BGP Vortices on Industry-Standard Router Software

Experiment 1: Test Router Susceptibility to BGP Vortex

Constructed BGP Vortex using 3 BGP implementations:

- Cisco IOS-XR (XRv9000)
- FRR 9.0.1
- BIRD 2.0.8

Results:

- All router implementations susceptible
- BIRD fastest circulation, FRR slowest

Experiment 2: Delayed Network Convergence

Goal: Measure delay of legitimate BGP Updates by BGP Vortex

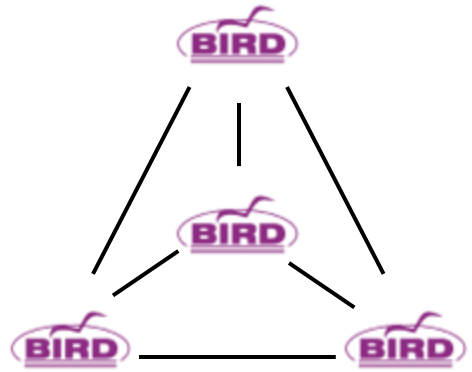

Experiment 2: Delayed Network Convergence

Goal: Measure delay of legitimate BGP Updates by BGP Vortex

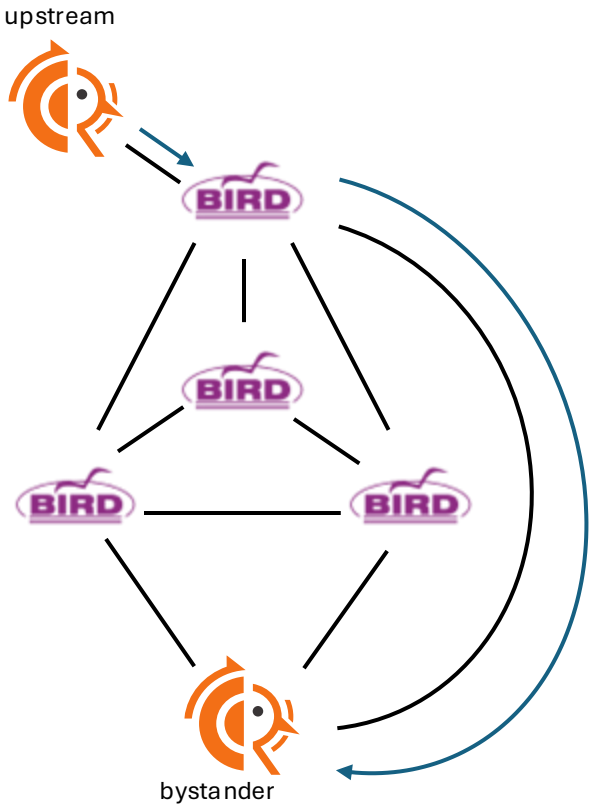

Experiment 2: Delayed Network Convergence

Goal: Measure delay of legitimate BGP Updates by BGP Vortex

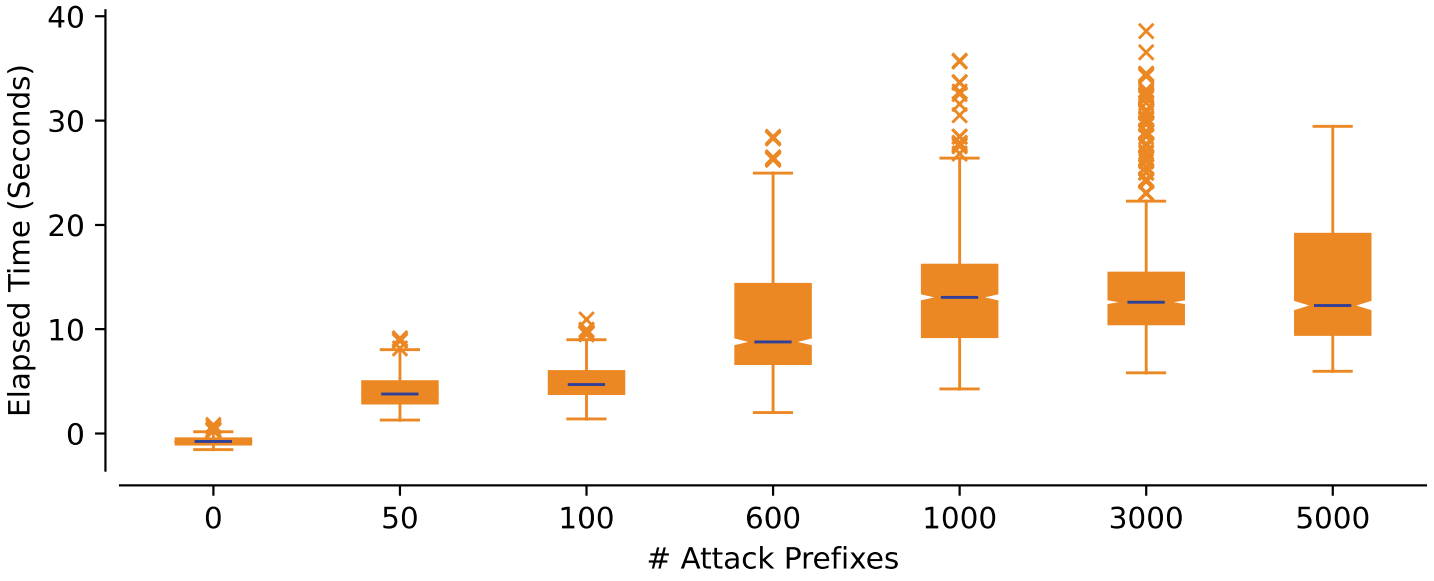

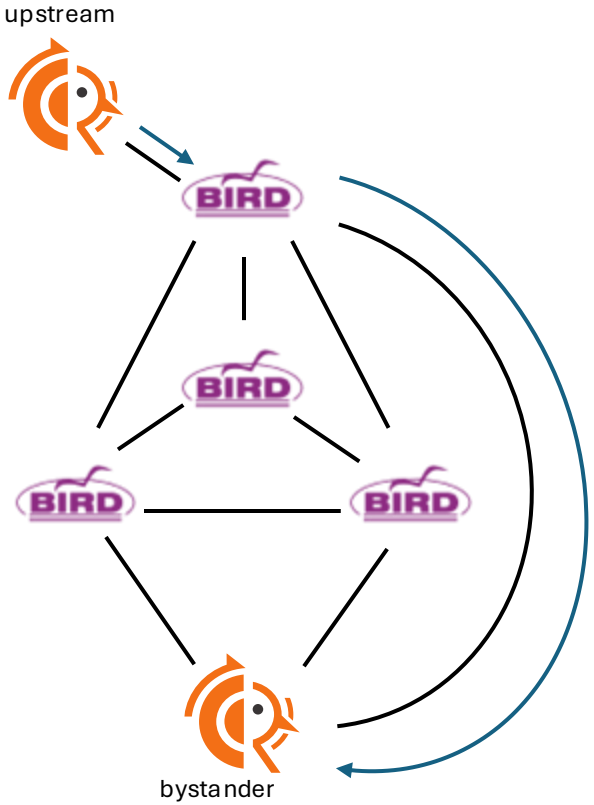

Experiment 3: Data Plane Outage

Goal: Measure duration in data plane unavailability

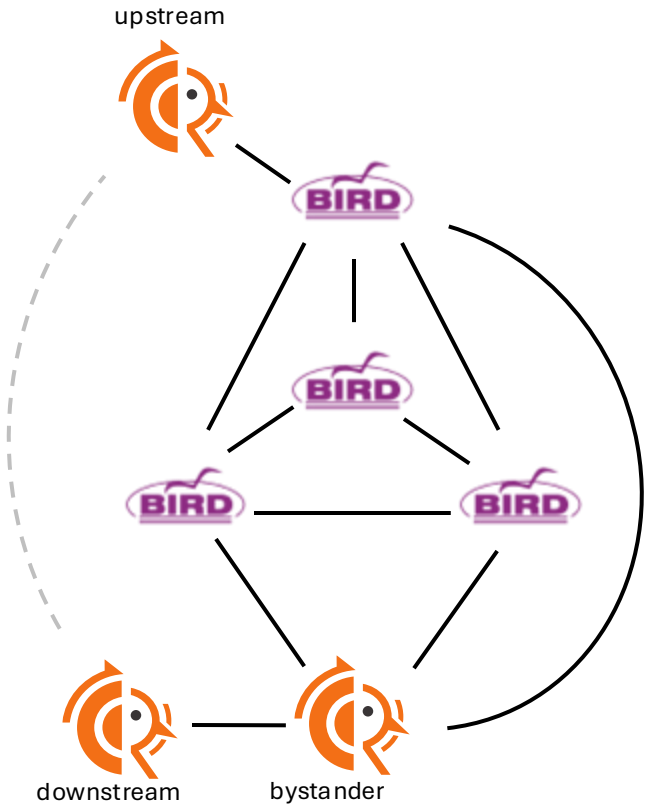

Experiment 3: Data Plane Outage

Goal: Measure duration in data plane unavailability

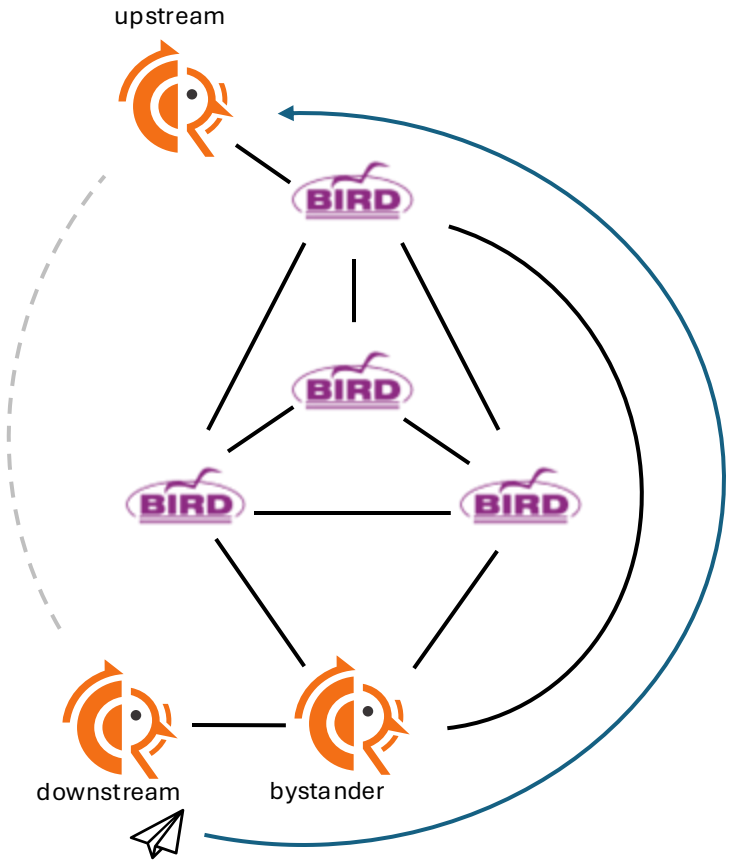

Experiment 3: Data Plane Outage

Goal: Measure duration in data plane unavailability

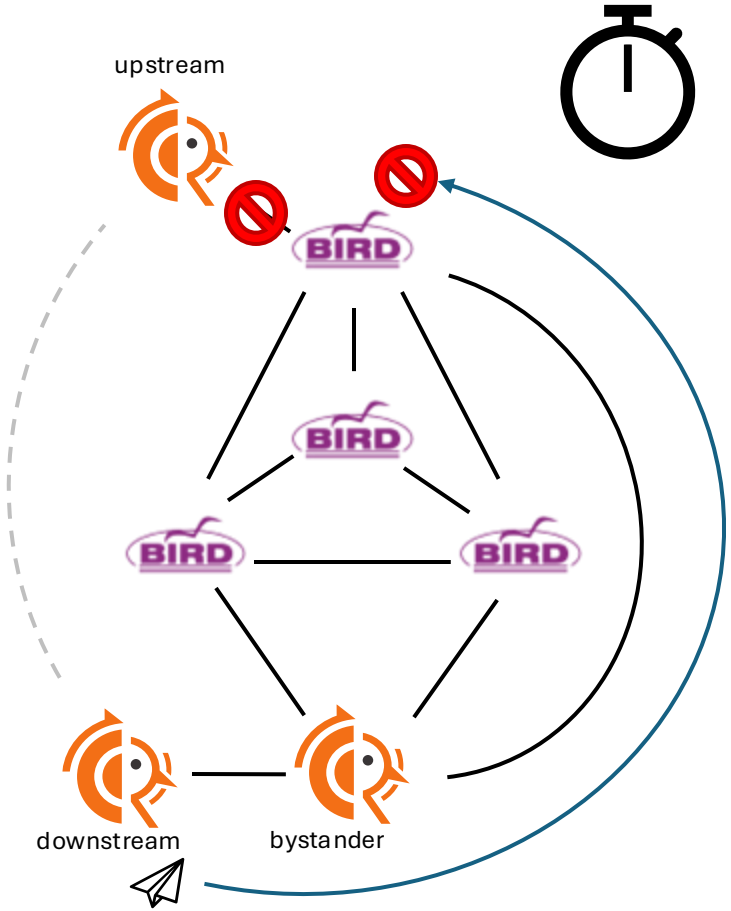

Experiment 3: Data Plane Outage

Goal: Measure duration in data plane unavailability

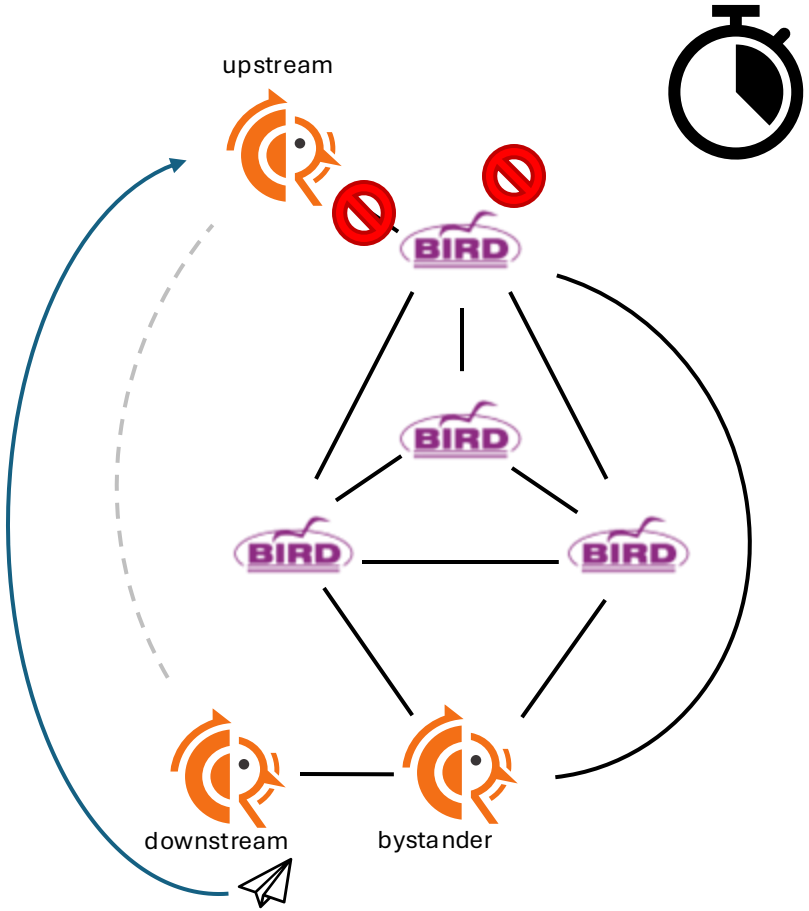

Experiment 3: Data Plane Outage

Goal: Measure duration in data plane unavailability

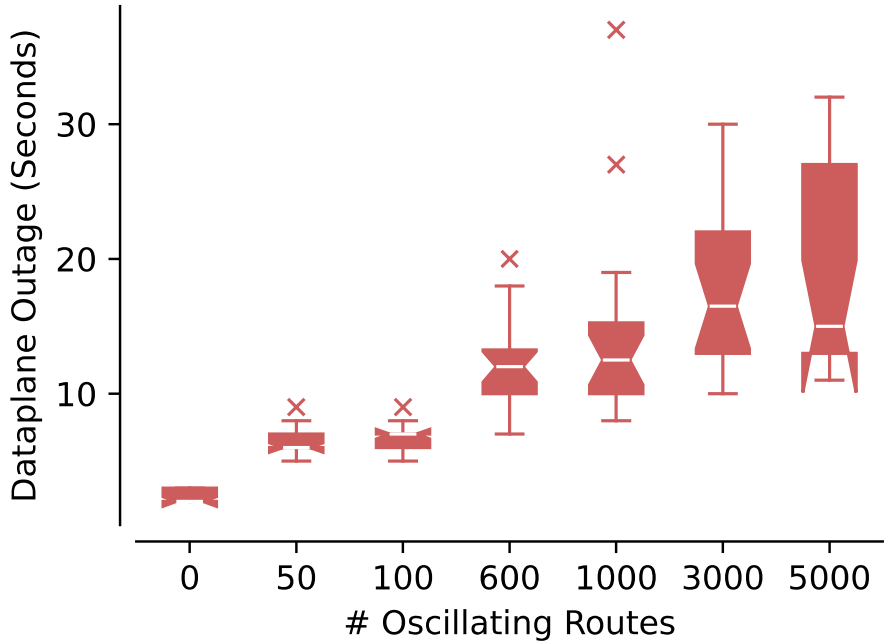

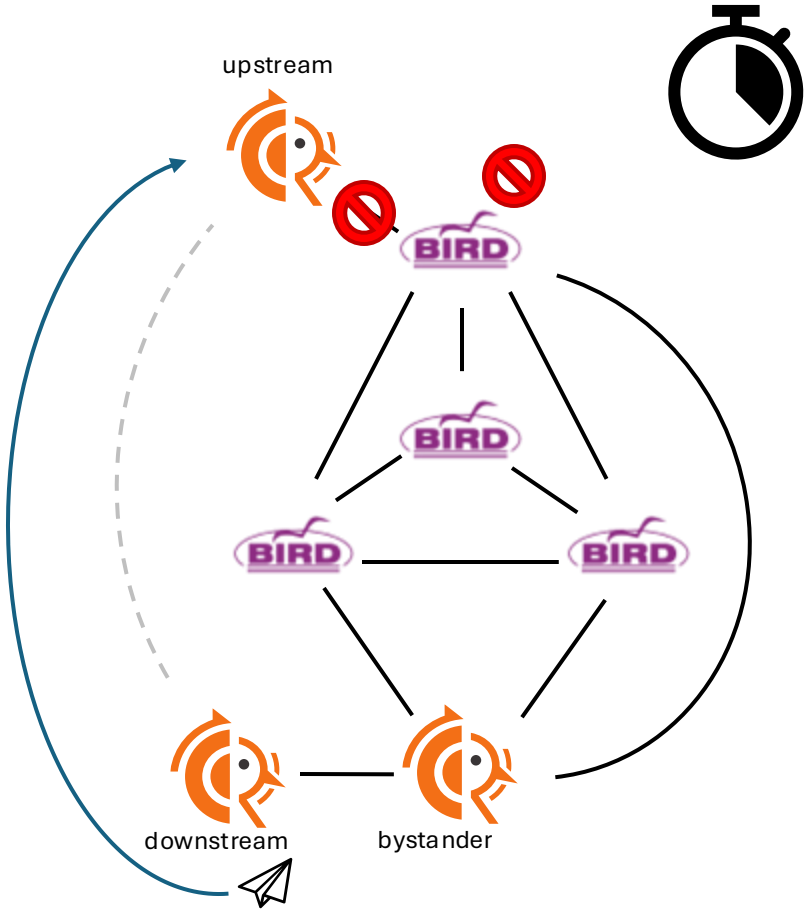

Mitigations

Safe BGP Communities

Recall: BGP converges if ASes route according to the Gao-Rexford conditions:

Advertisement Forwarded To

Advertisement Received From	Customer	Peer	Provider
Customer	Advertise	Advertise	Advertise
Peer	Advertise	Drop	Drop
Provider	Advertise	Drop	Drop

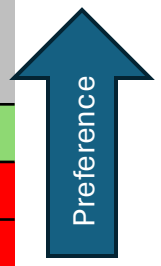

Safe BGP Communities

Recall: BGP converges if ASes route according to the Gao-Rexford conditions:

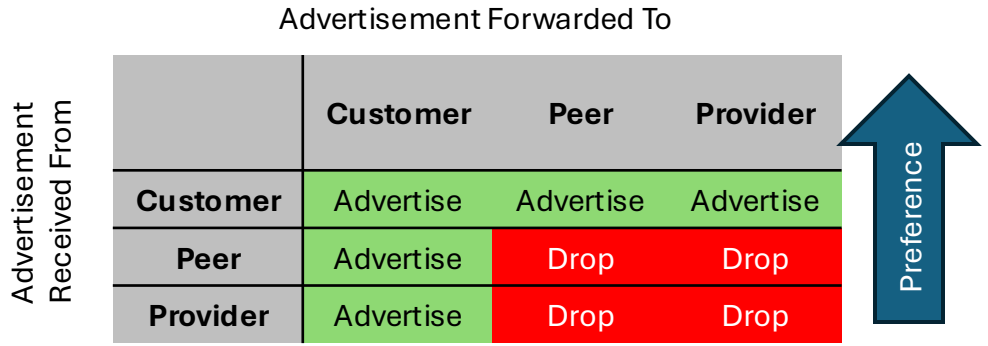

Idea: Disallow BGP communities violating Gao-Rexford conditions

Safe BGP Communities

Recall: BGP converges if ASes route according to the Gao-Rexford conditions:

Advertisement Forwarded To

Advertisement Received From	Customer	Peer	Provider
Customer	Advertise	Advertise	Advertise
Peer	Advertise	Drop	Drop
Provider	Advertise	Drop	Drop

Preference

Idea: Disallow BGP communities violating Gao-Rexford conditions

Restricting Lower-Pref BGP community:

Safe BGP Communities

Recall: BGP converges if ASes route according to the Gao-Rexford conditions:

Advertisement Forwarded To

Advertisement Received From	Customer	Peer	Provider
Customer	Advertise	Advertise	Advertise
Peer	Advertise	Drop	Drop
Provider	Advertise	Drop	Drop

Preference

Idea: Disallow BGP communities violating Gao-Rexford conditions

Restricting Lower-Pref BGP community:

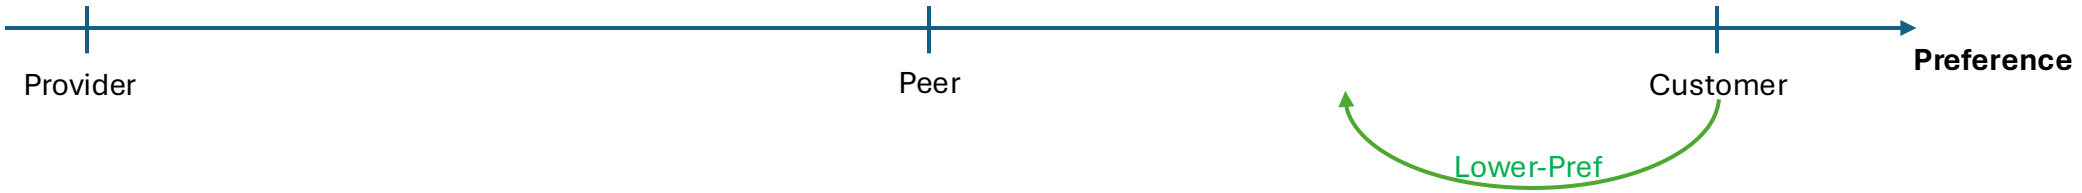

Safe BGP Communities

Recall: BGP converges if ASes route according to the Gao-Rexford conditions:

Advertisement Forwarded To

Advertisement Received From	Customer	Peer	Provider
Customer	Advertise	Advertise	Advertise
Peer	Advertise	Drop	Drop
Provider	Advertise	Drop	Drop

Preference

Idea: Disallow BGP communities violating Gao-Rexford conditions

Restricting Lower-Pref BGP community:

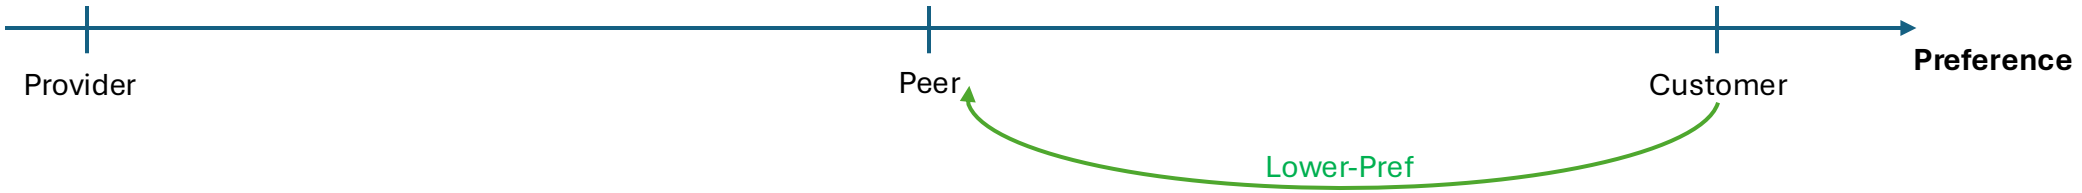

Safe BGP Communities

Recall: BGP converges if ASes route according to the Gao-Rexford conditions:

Advertisement Forwarded To

Advertisement Received From	Customer	Peer	Provider
Customer	Advertise	Advertise	Advertise
Peer	Advertise	Drop	Drop
Provider	Advertise	Drop	Drop

Preference

Idea: Disallow BGP communities violating Gao-Rexford conditions

Restricting Lower-Pref BGP community:

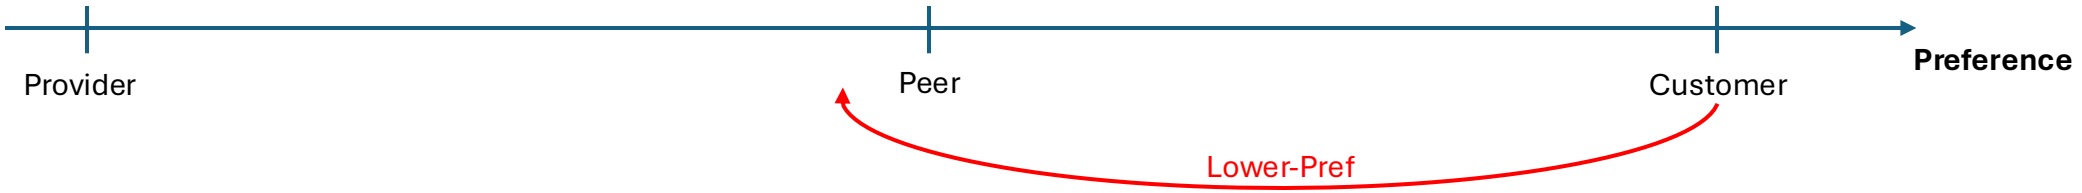

Conclusion

- **BGP Vortices** are an induced state of persistent route instability among 3 benign ASes

Conclusion

- **BGP Vortices** are an induced state of persistent route instability among 3 benign ASes
- A Vortex can be triggered using only **three BGP Update messages**

Conclusion

- **BGP Vortices** are an induced state of persistent route instability among 3 benign ASes
- A Vortex can be triggered using only **three BGP Update messages**
- BGP Vortices can generate **thousands of BGP Updates per second**

Conclusion

- **BGP Vortices** are an induced state of persistent route instability among 3 benign ASes
- A Vortex can be triggered using only **three BGP Update messages**
- BGP Vortices can generate **thousands of BGP Updates per second**
- **20/30** largest ASes are susceptible to BGP Vortices
- **96%** of all ASes in the customer cone of at least one BGP Vortex

Conclusion

- **BGP Vortices** are an induced state of persistent route instability among 3 benign ASes
- A Vortex can be triggered using only **three BGP Update messages**
- BGP Vortices can generate **thousands of BGP Updates per second**
- **20/30** largest ASes are susceptible to BGP Vortices
- **96%** of all ASes in the customer cone of at least one BGP Vortex
- Indication that long-standing assumptions on BGP stability may not hold

Conclusion

Thank you for your attention!

Questions?

felix.stoeger@inf.ethz.ch

- **BGP Vortices** are an induced state of persistent route instability among 3 benign ASes
- A Vortex can be triggered using only **three BGP Update messages**
- BGP Vortices can generate **thousands of BGP Updates per second**
- **20/30** largest ASes are susceptible to BGP Vortices
- **96%** of all ASes in the customer cone of at least one BGP Vortex
- Indication that long-standing assumptions on BGP stability may not hold