Akamai DNS

Providing Authoritative Answers to the World's Queries

Kyle Schomp, Onkar Bhardwaj, Eymen Kurdoglu, Mashooq Muhaimen, Ramesh K. Sitaraman

end-user

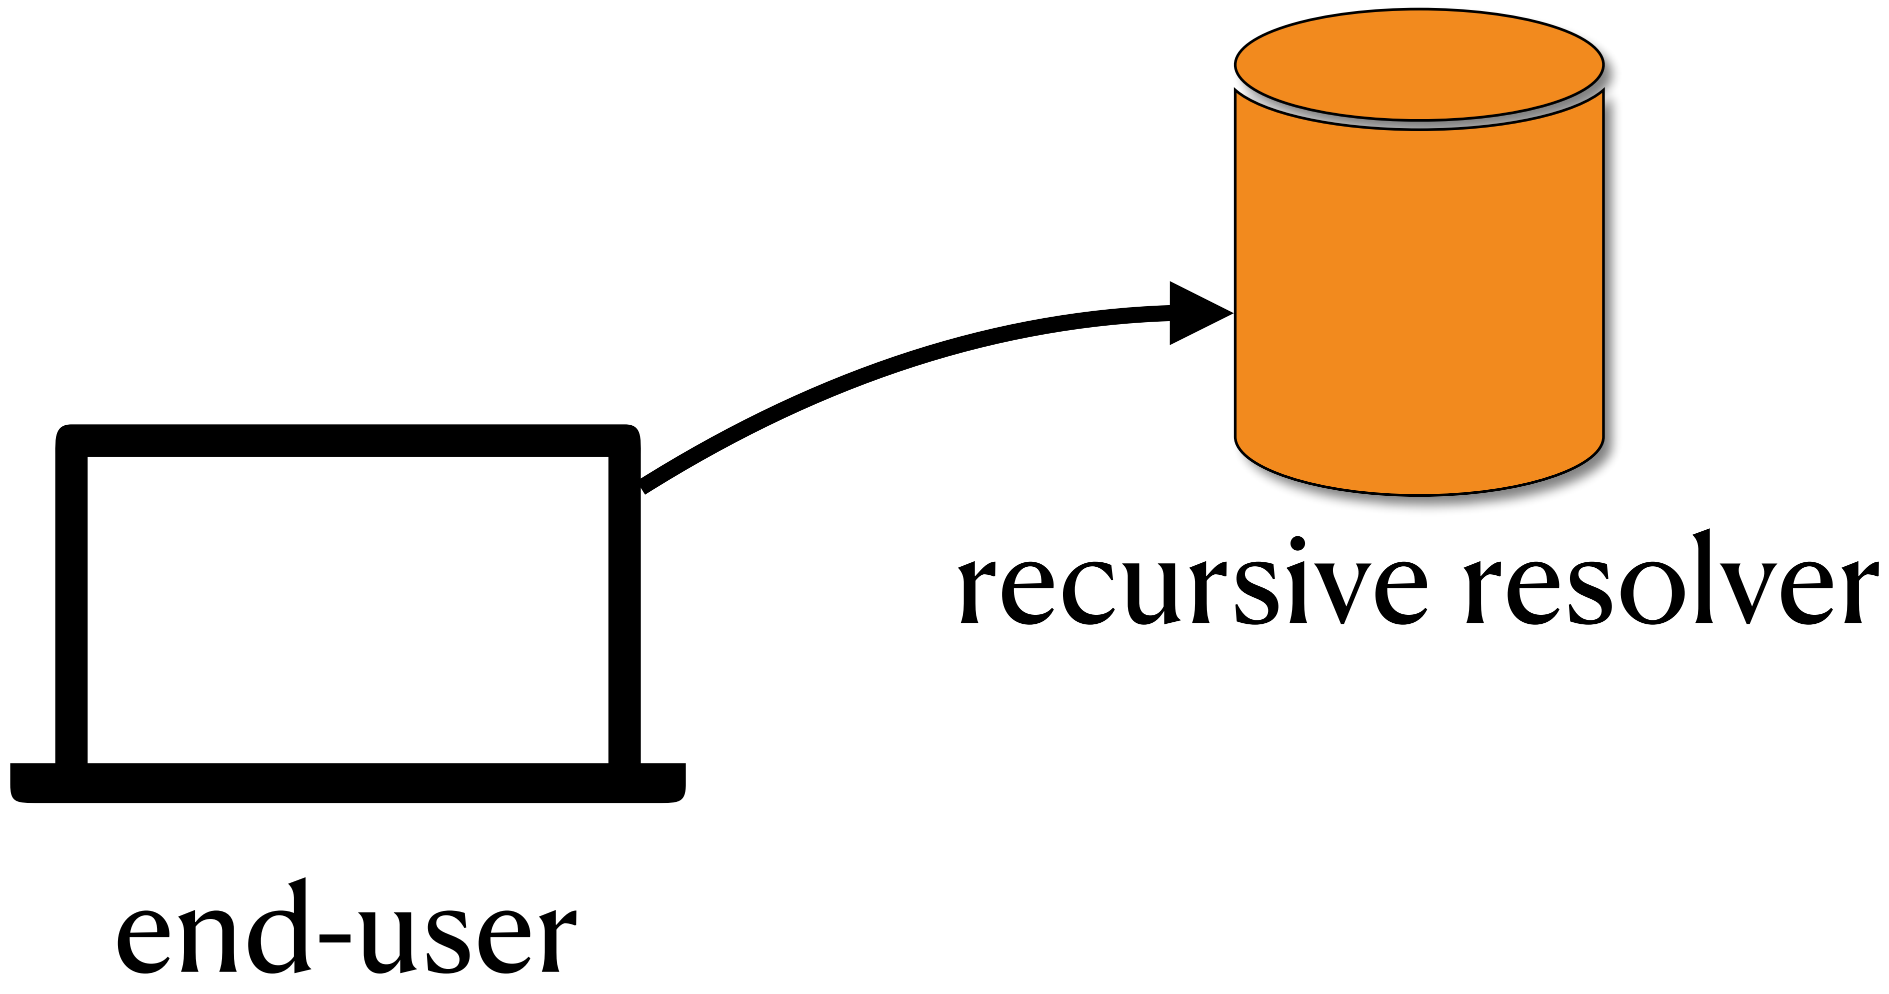
recursive resolver
end-user

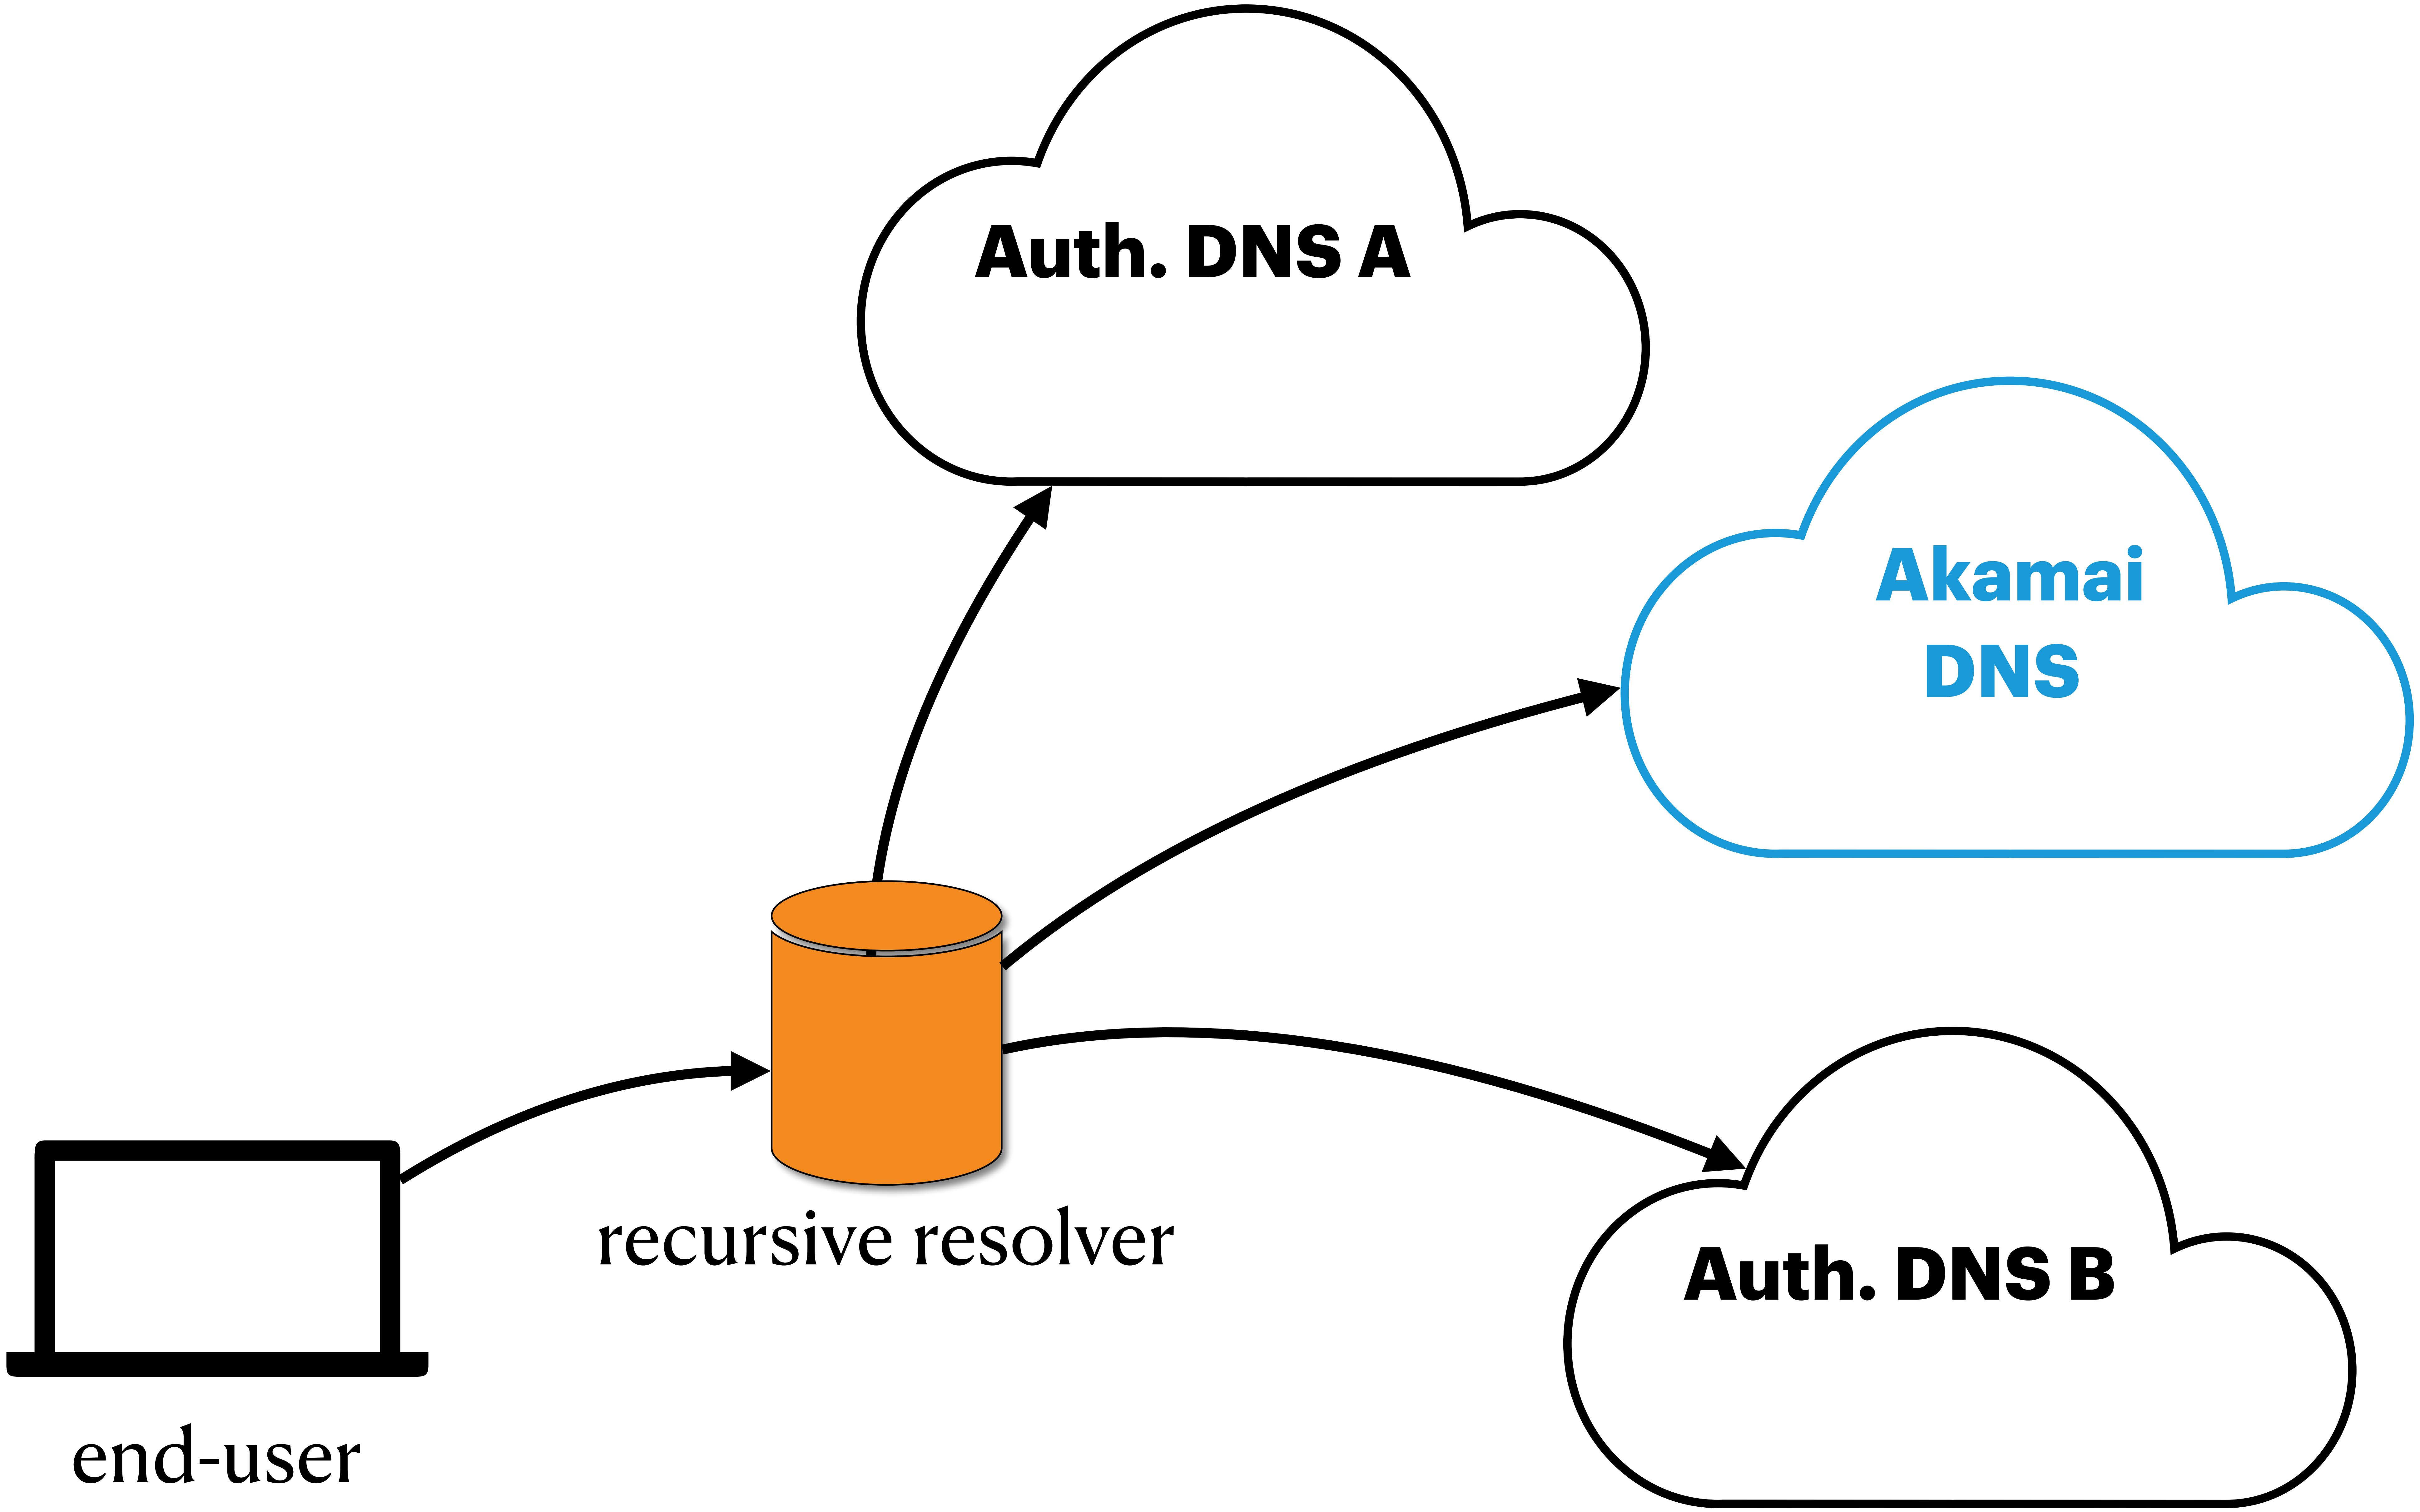
Auth. DNS A
Akamai DNS
recursive resolver
Auth. DNS B
end-user

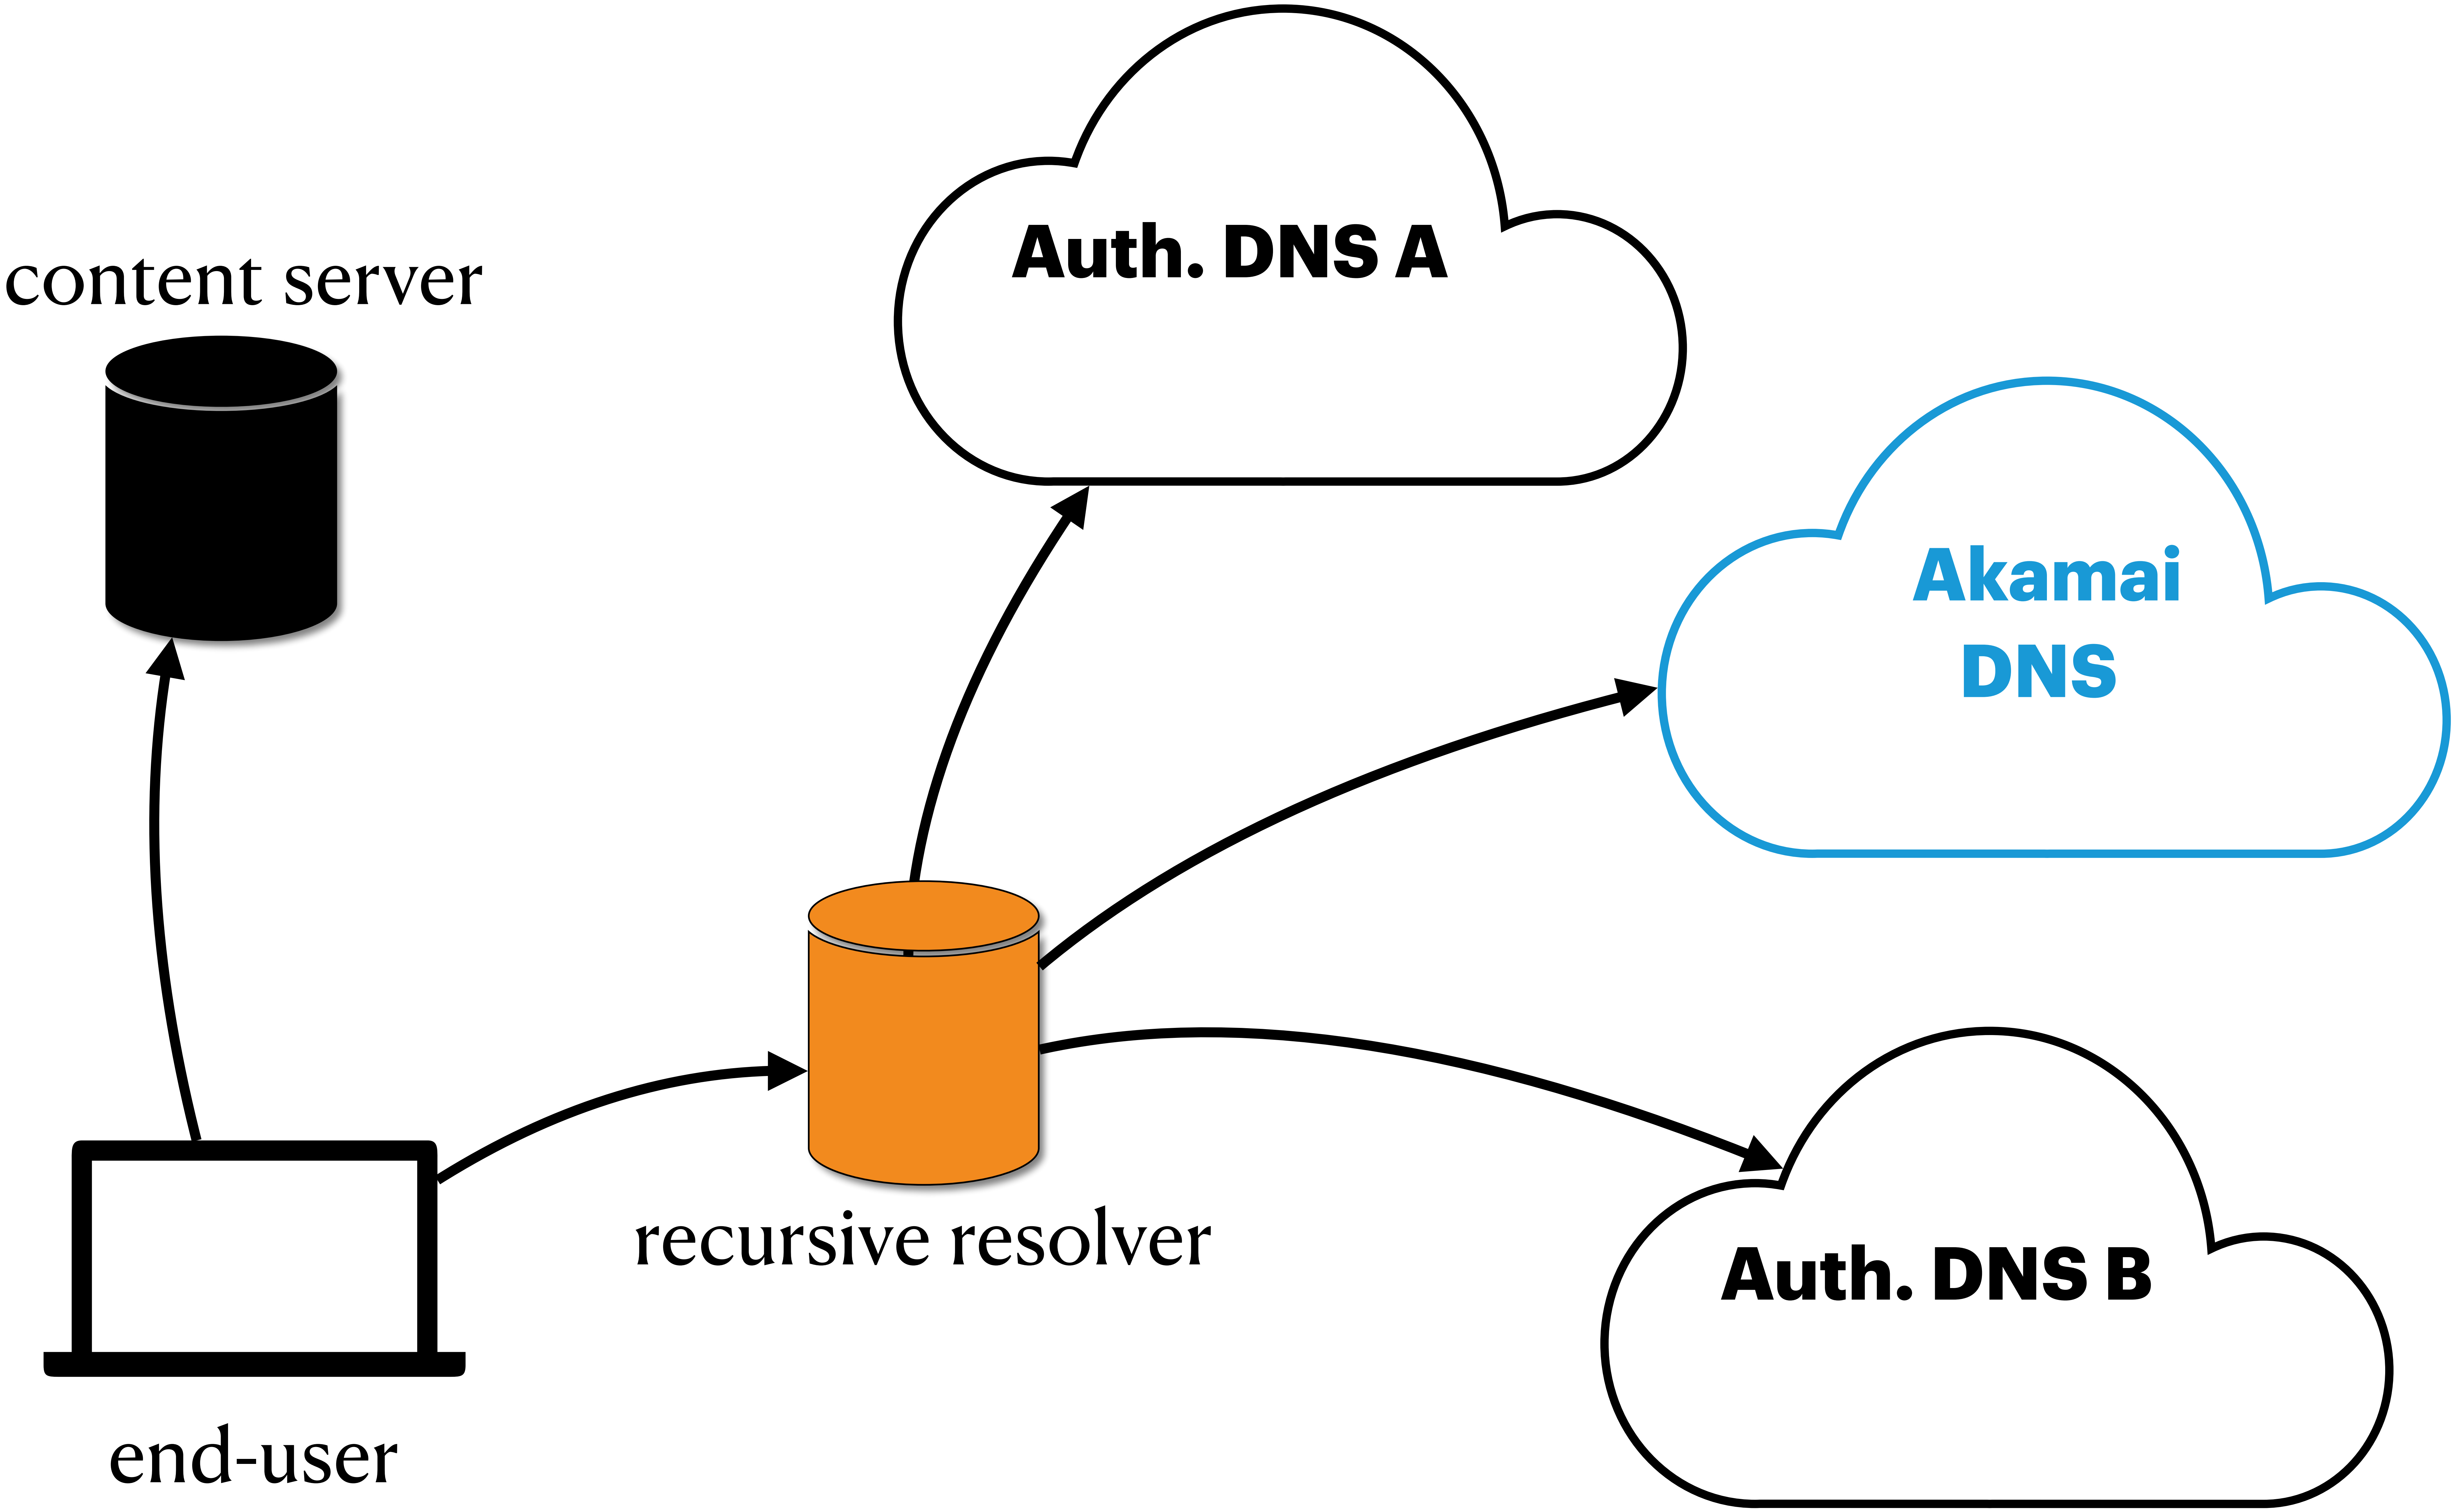
content server
Auth. DNS A
Akamai DNS
recursive resolver
end-user
Auth. DNS B

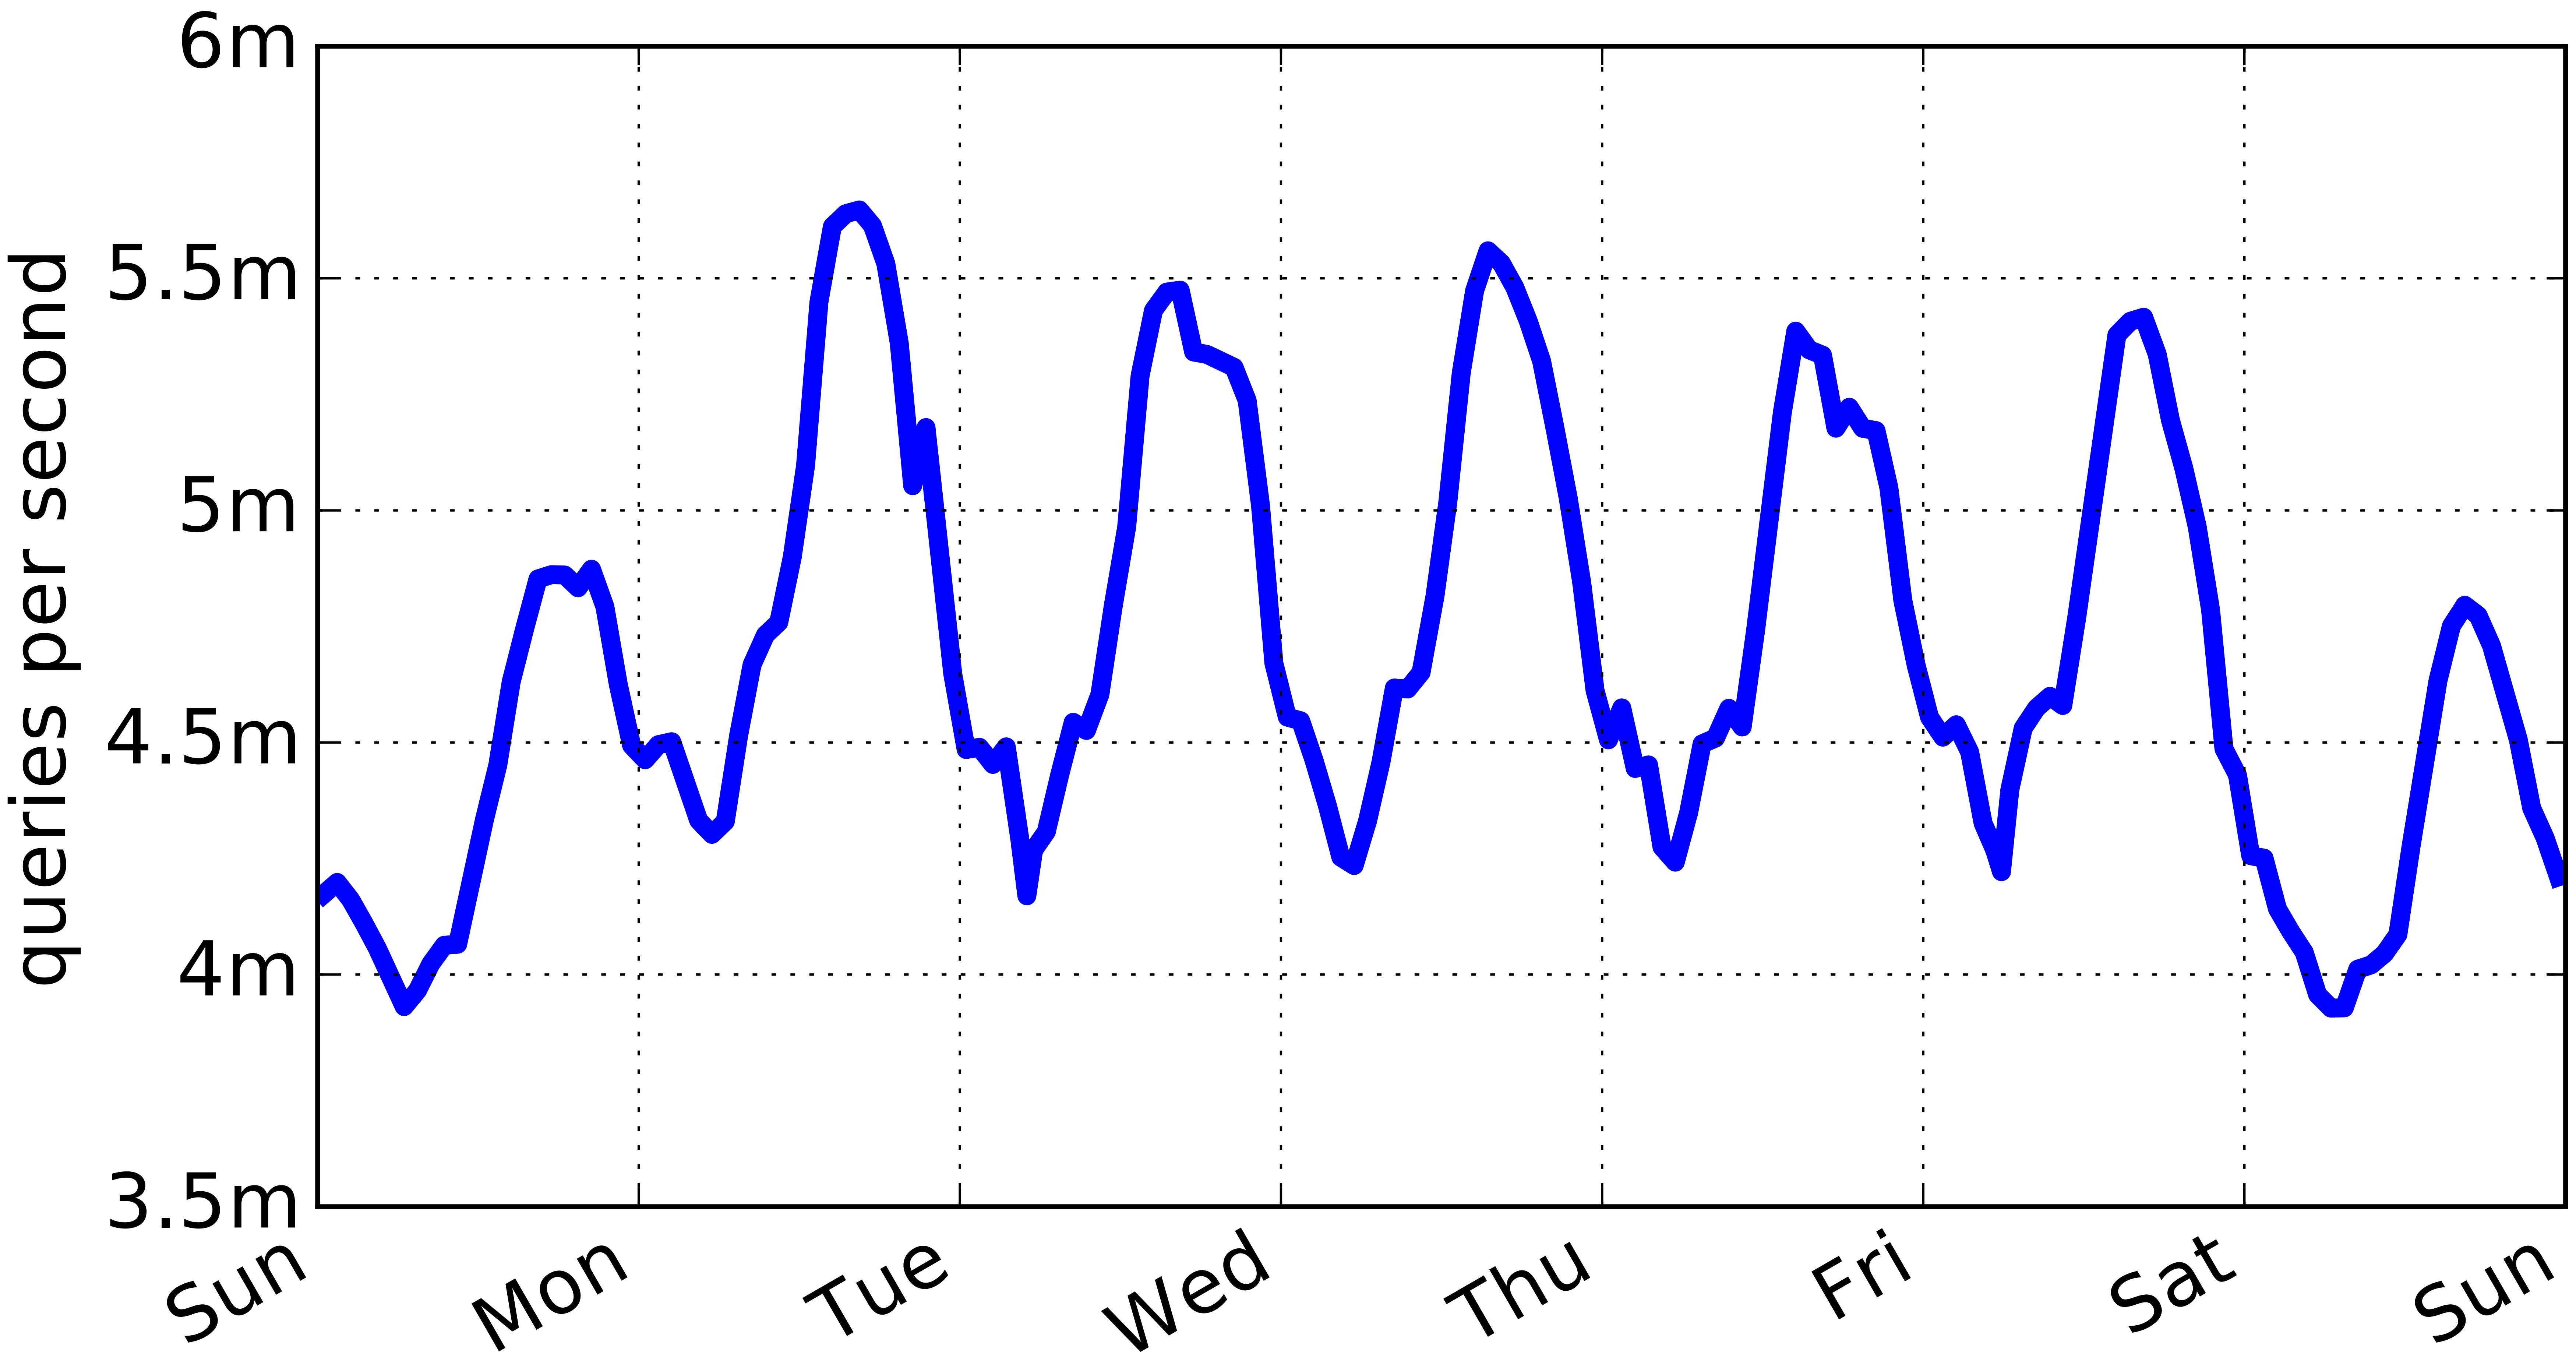
queries per second
6m
5.5m
5m
4.5m
4m
3.5m
Sun
Mon
Tue
Wed
Thu
Fri
Sat
Sun

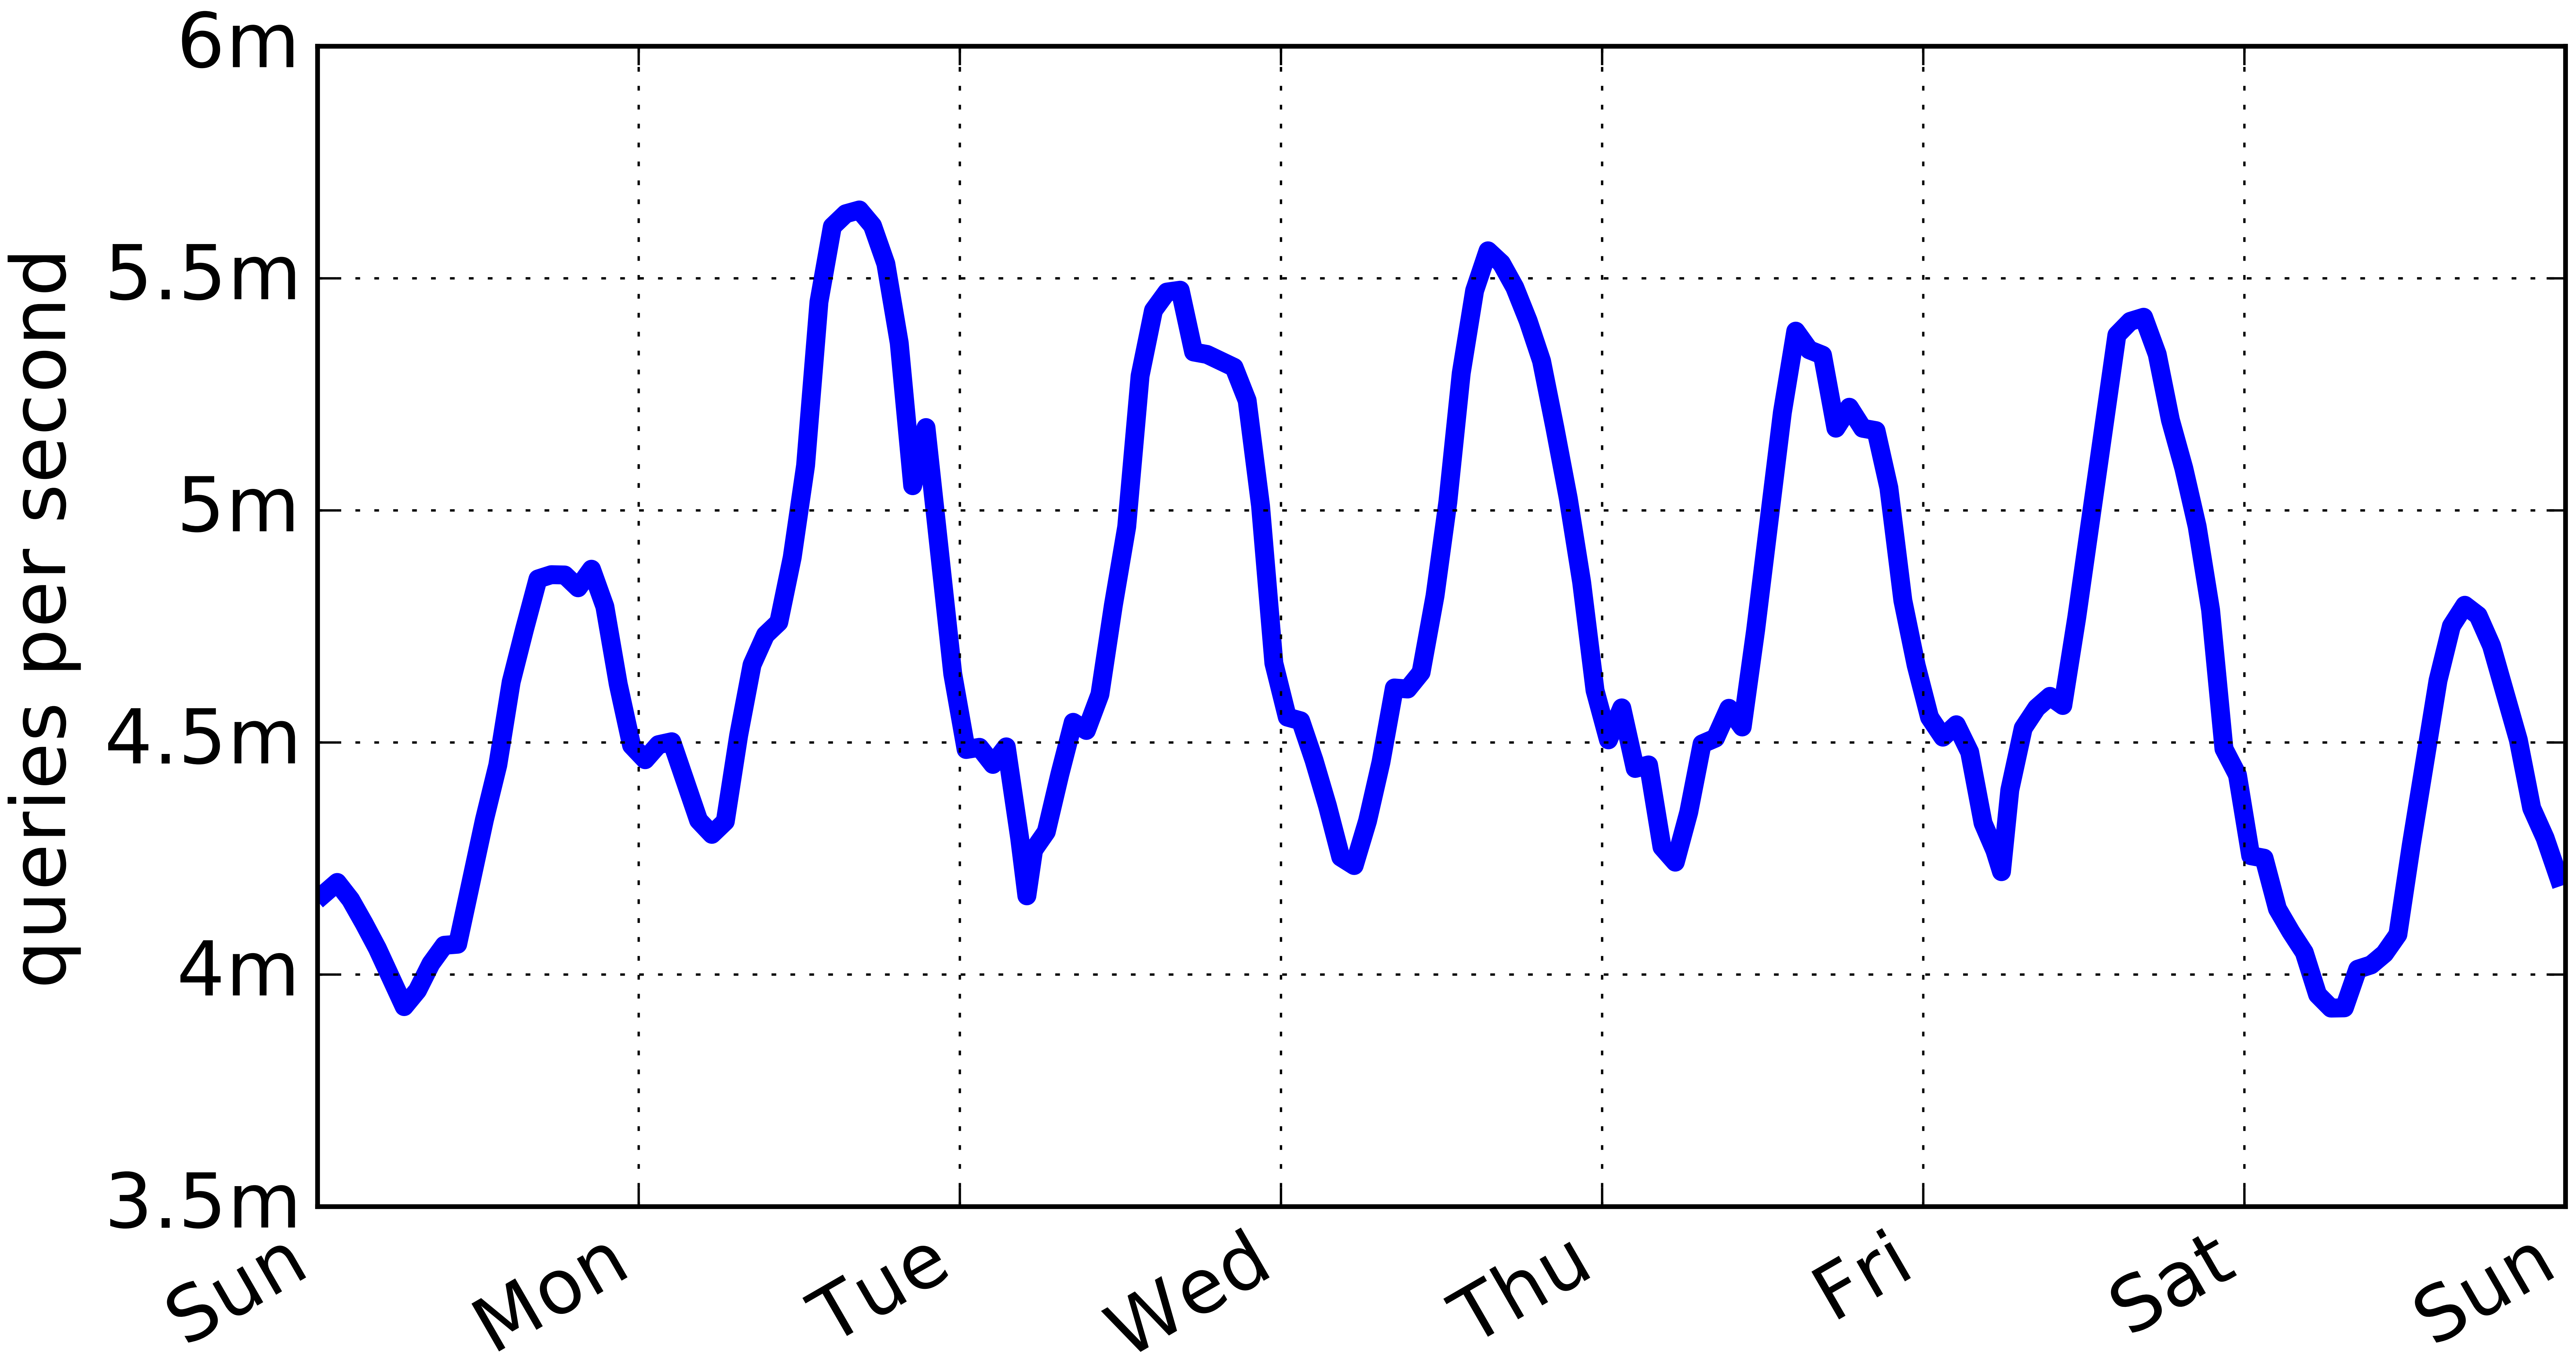

Akamai DNS is the starting point for a significant fraction of the world's Internet interactions and has a critical role in the Internet ecosystem

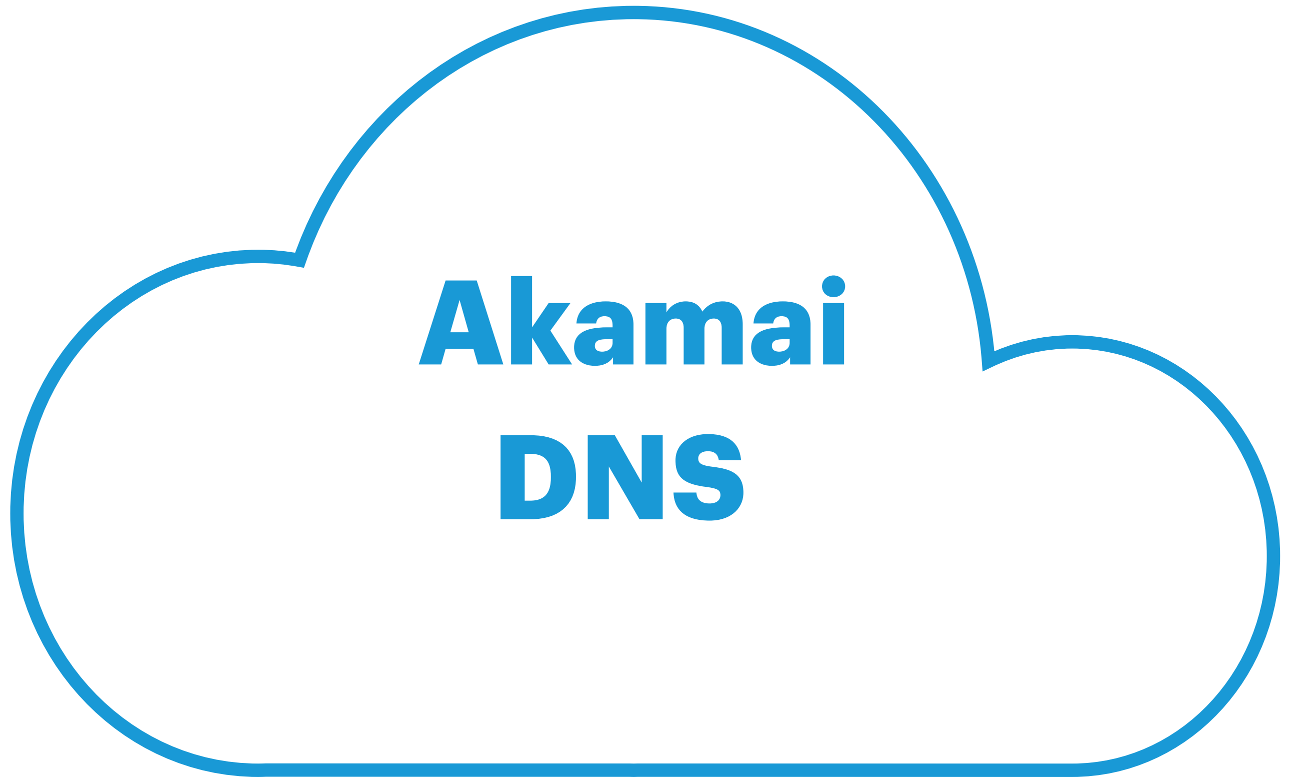

Akamai DNS

Deployed Points of Presence (PoPs) distributed around the world

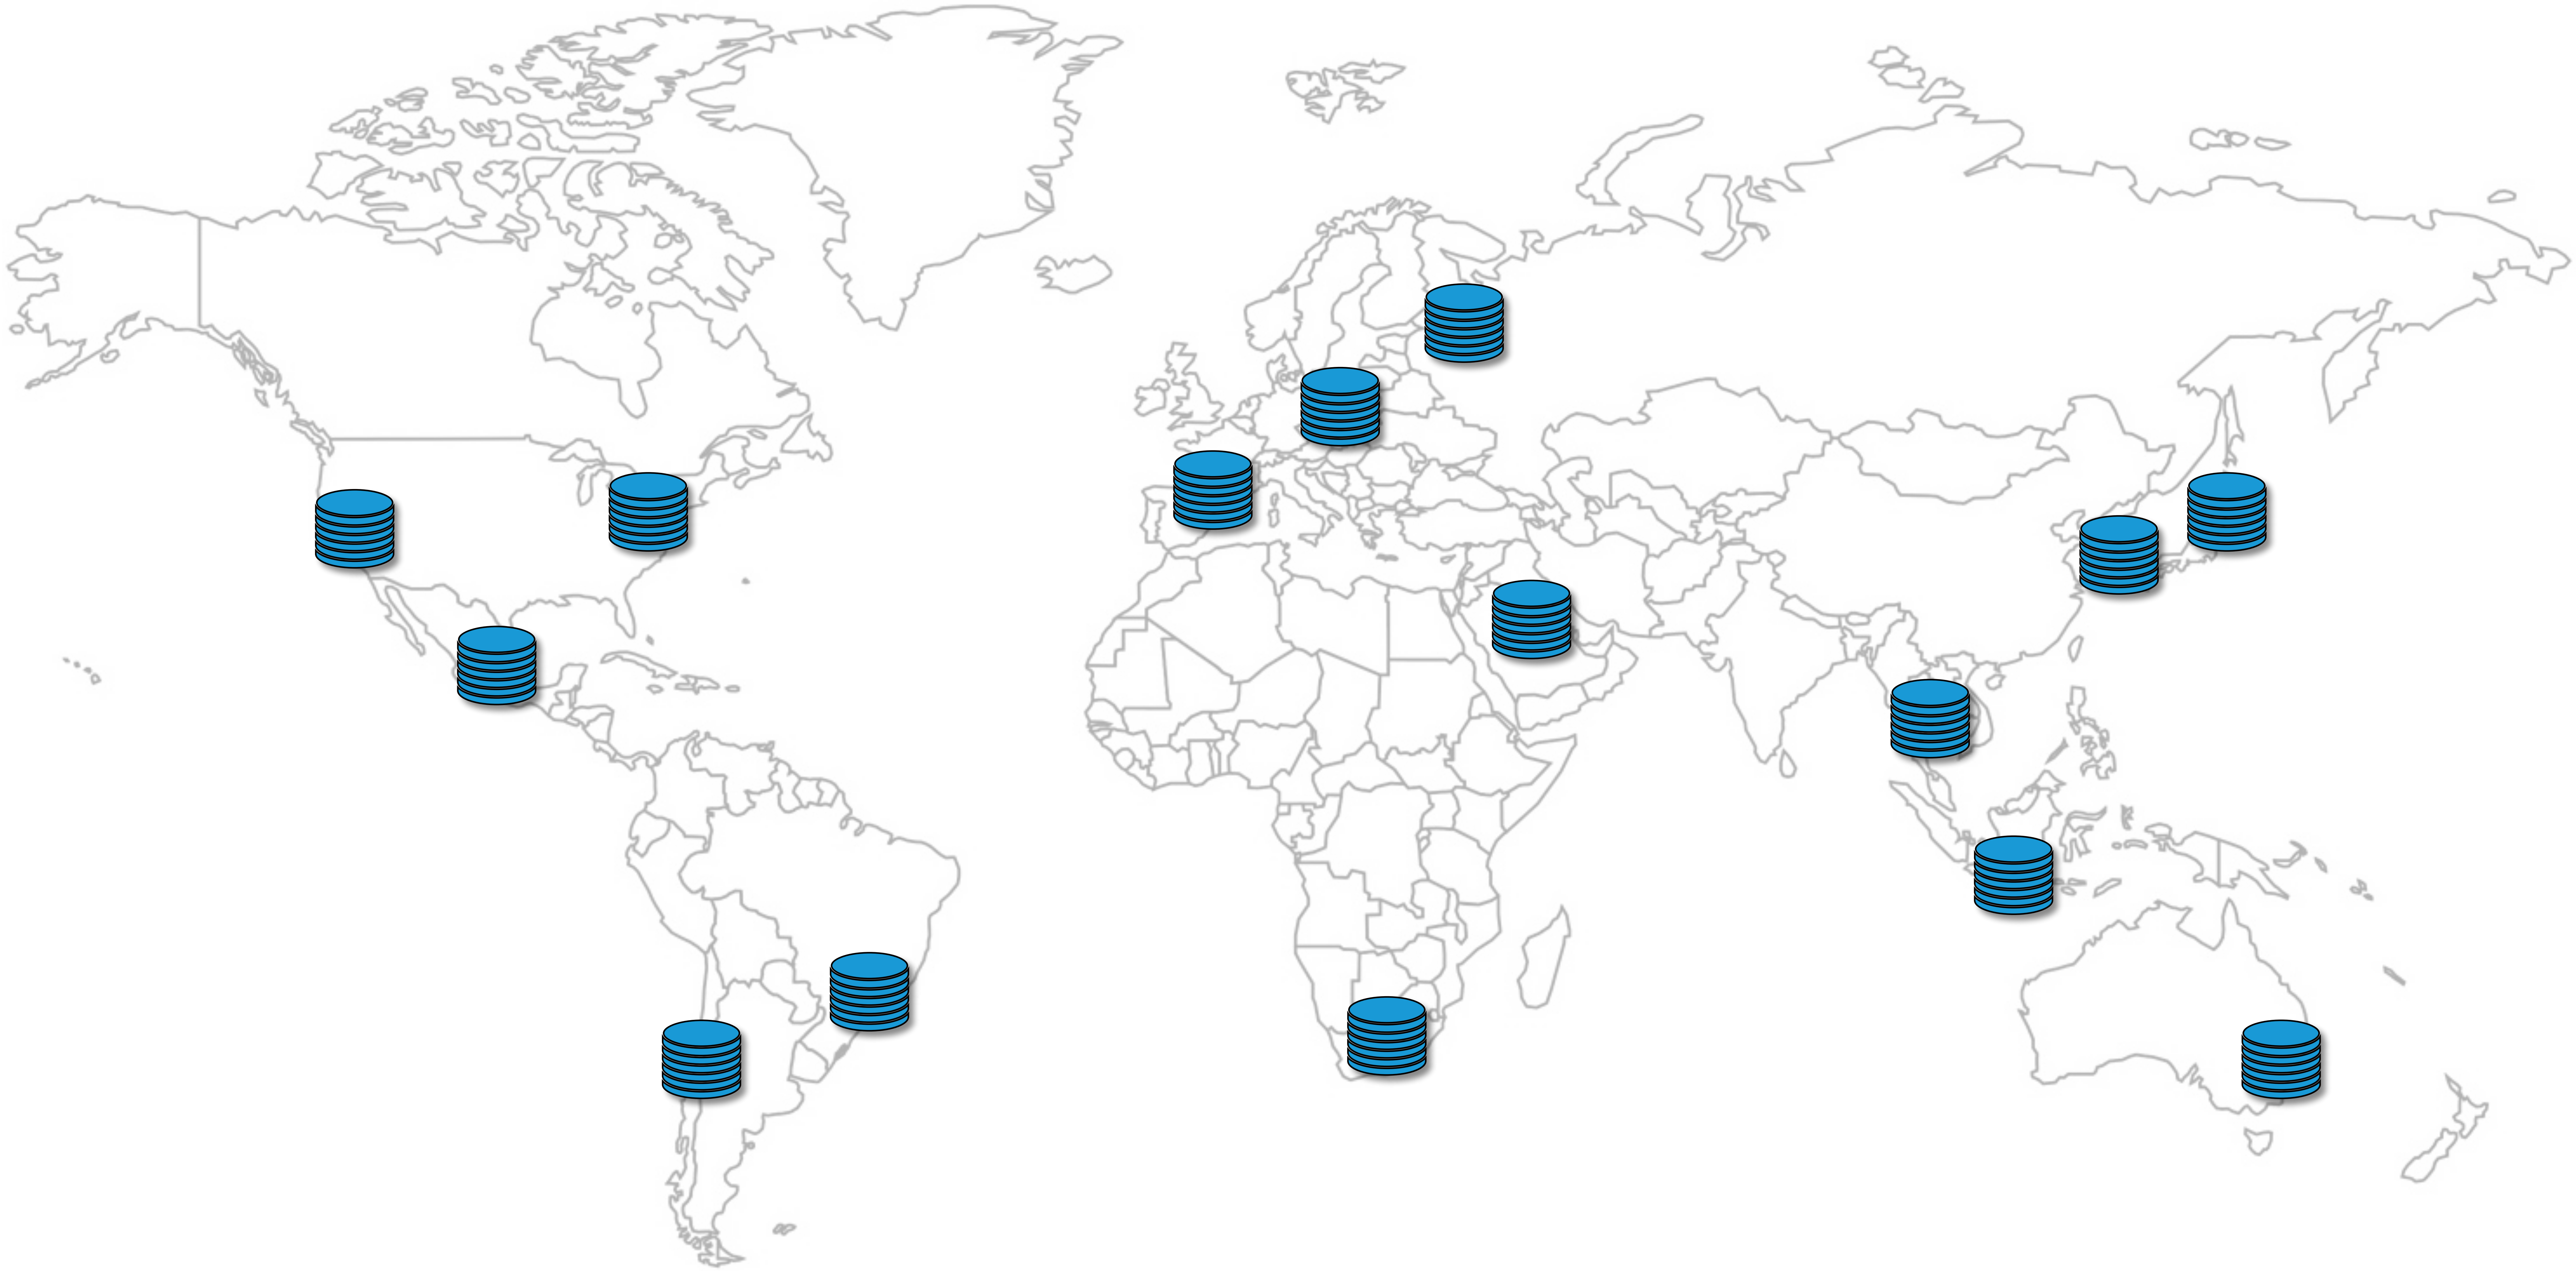

Deployed Points of Presence (PoPs) distributed around the world

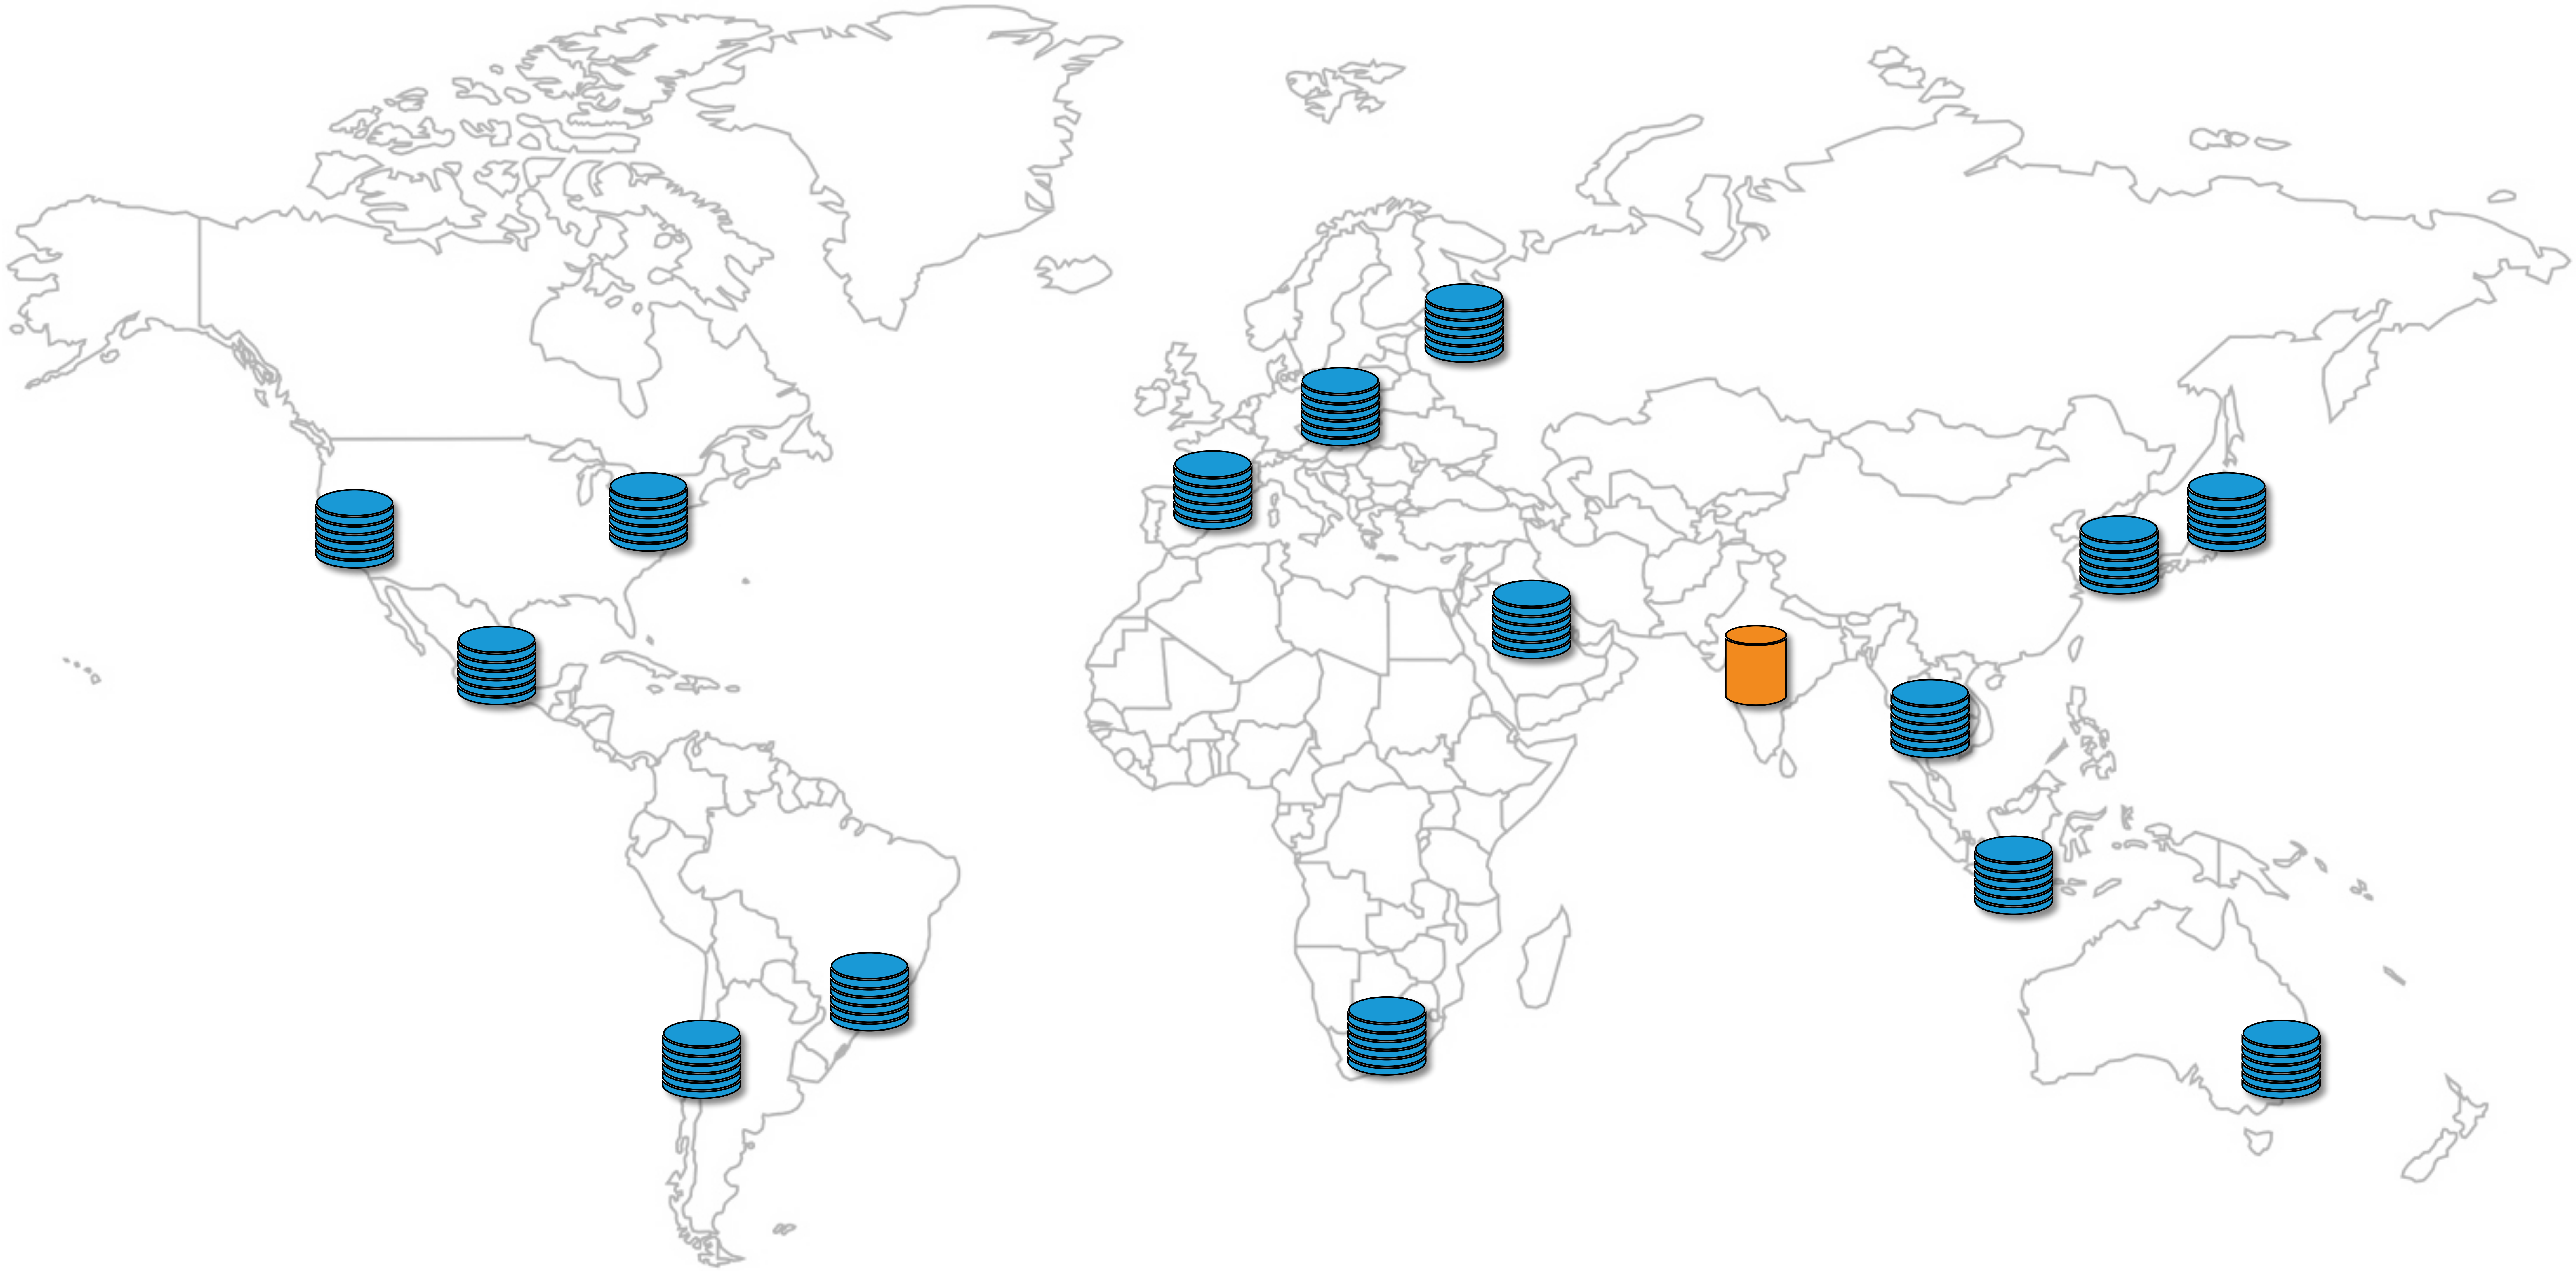

Deployed Points of Presence (PoPs) distributed around the world

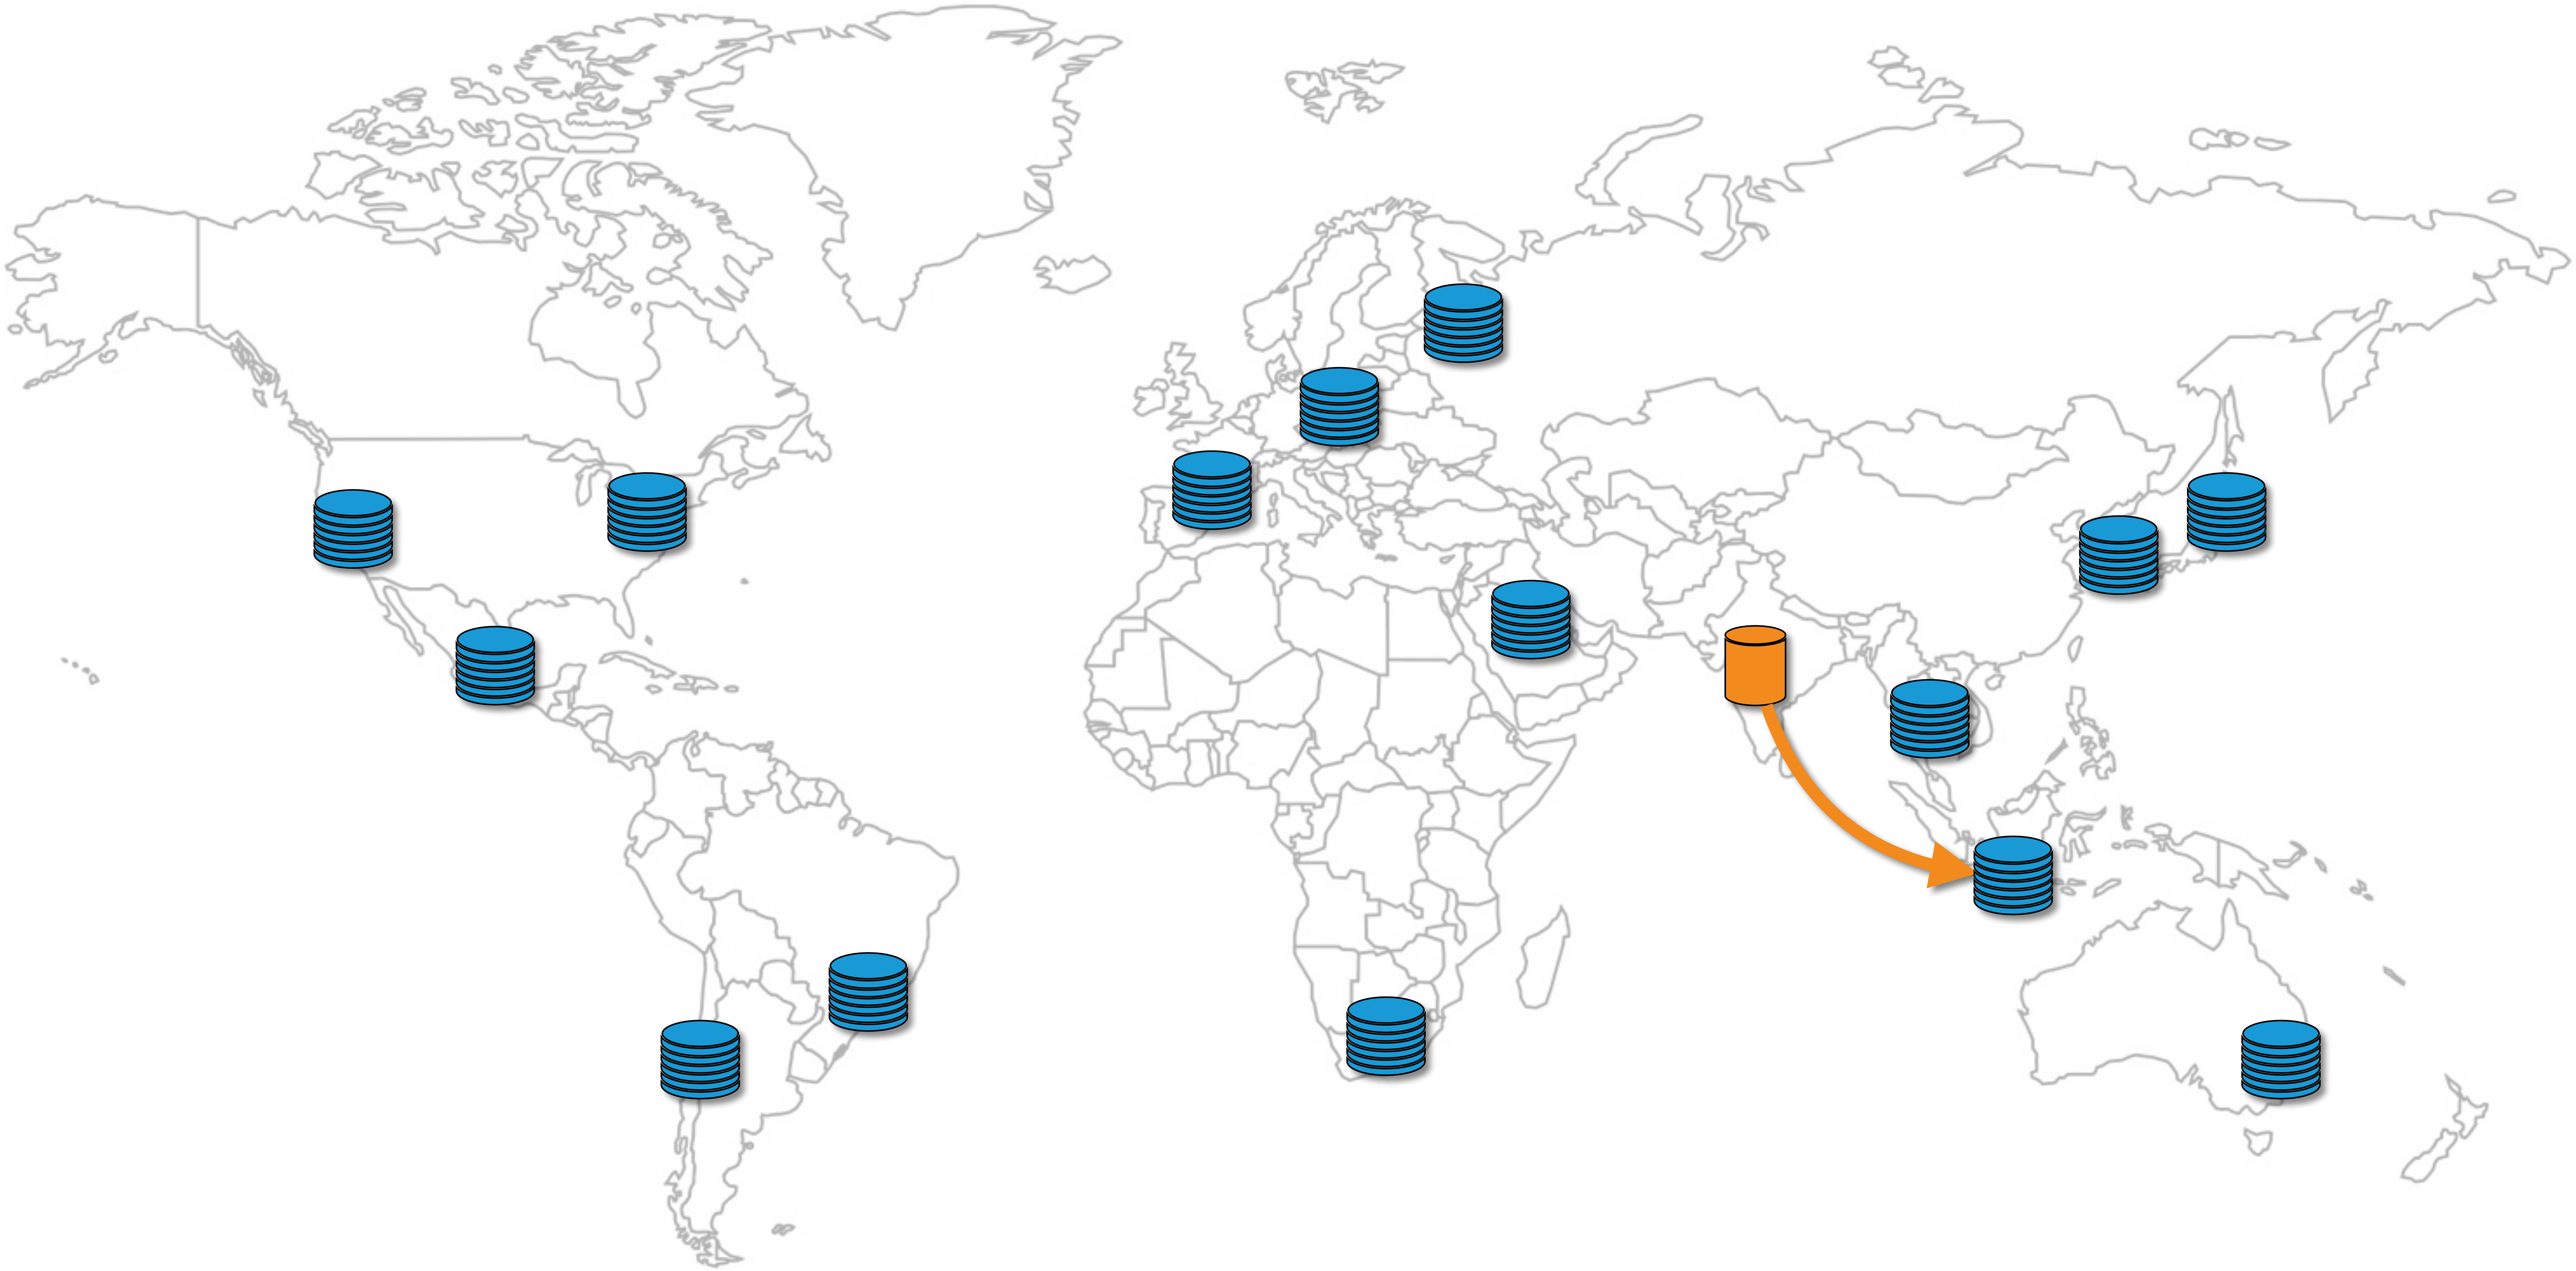

IP Anycast routes Recursive Resolvers to a PoP

One PoP

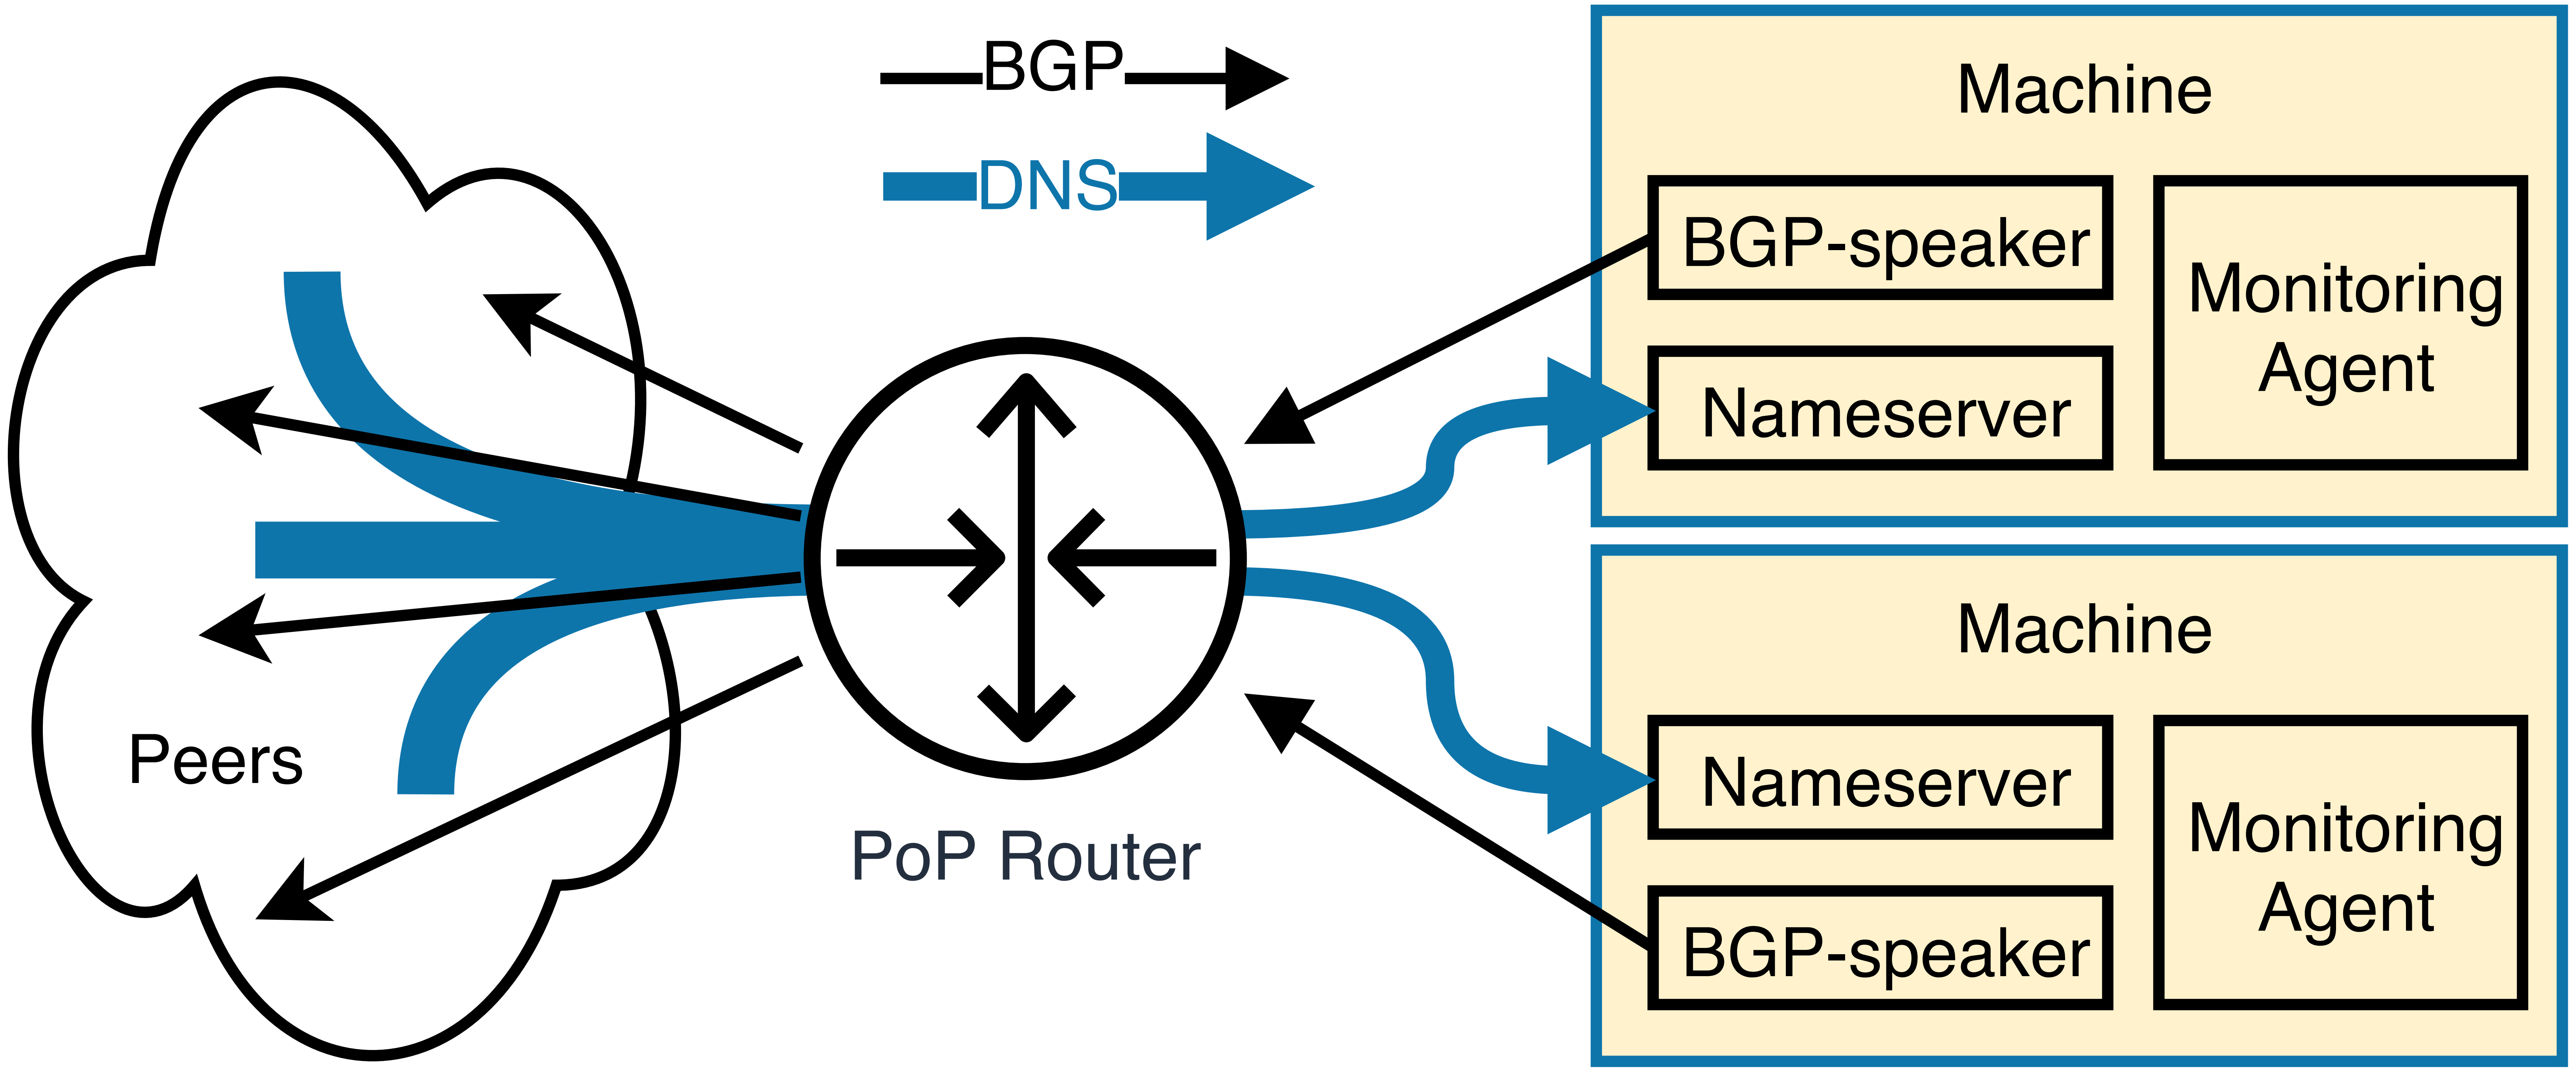
BGP
DNS
Machine
BGP-speaker
Nameserver
Monitoring Agent
Machine
Nameserver
BGP-speaker
Monitoring Agent
Peers
PoP Router

One PoP

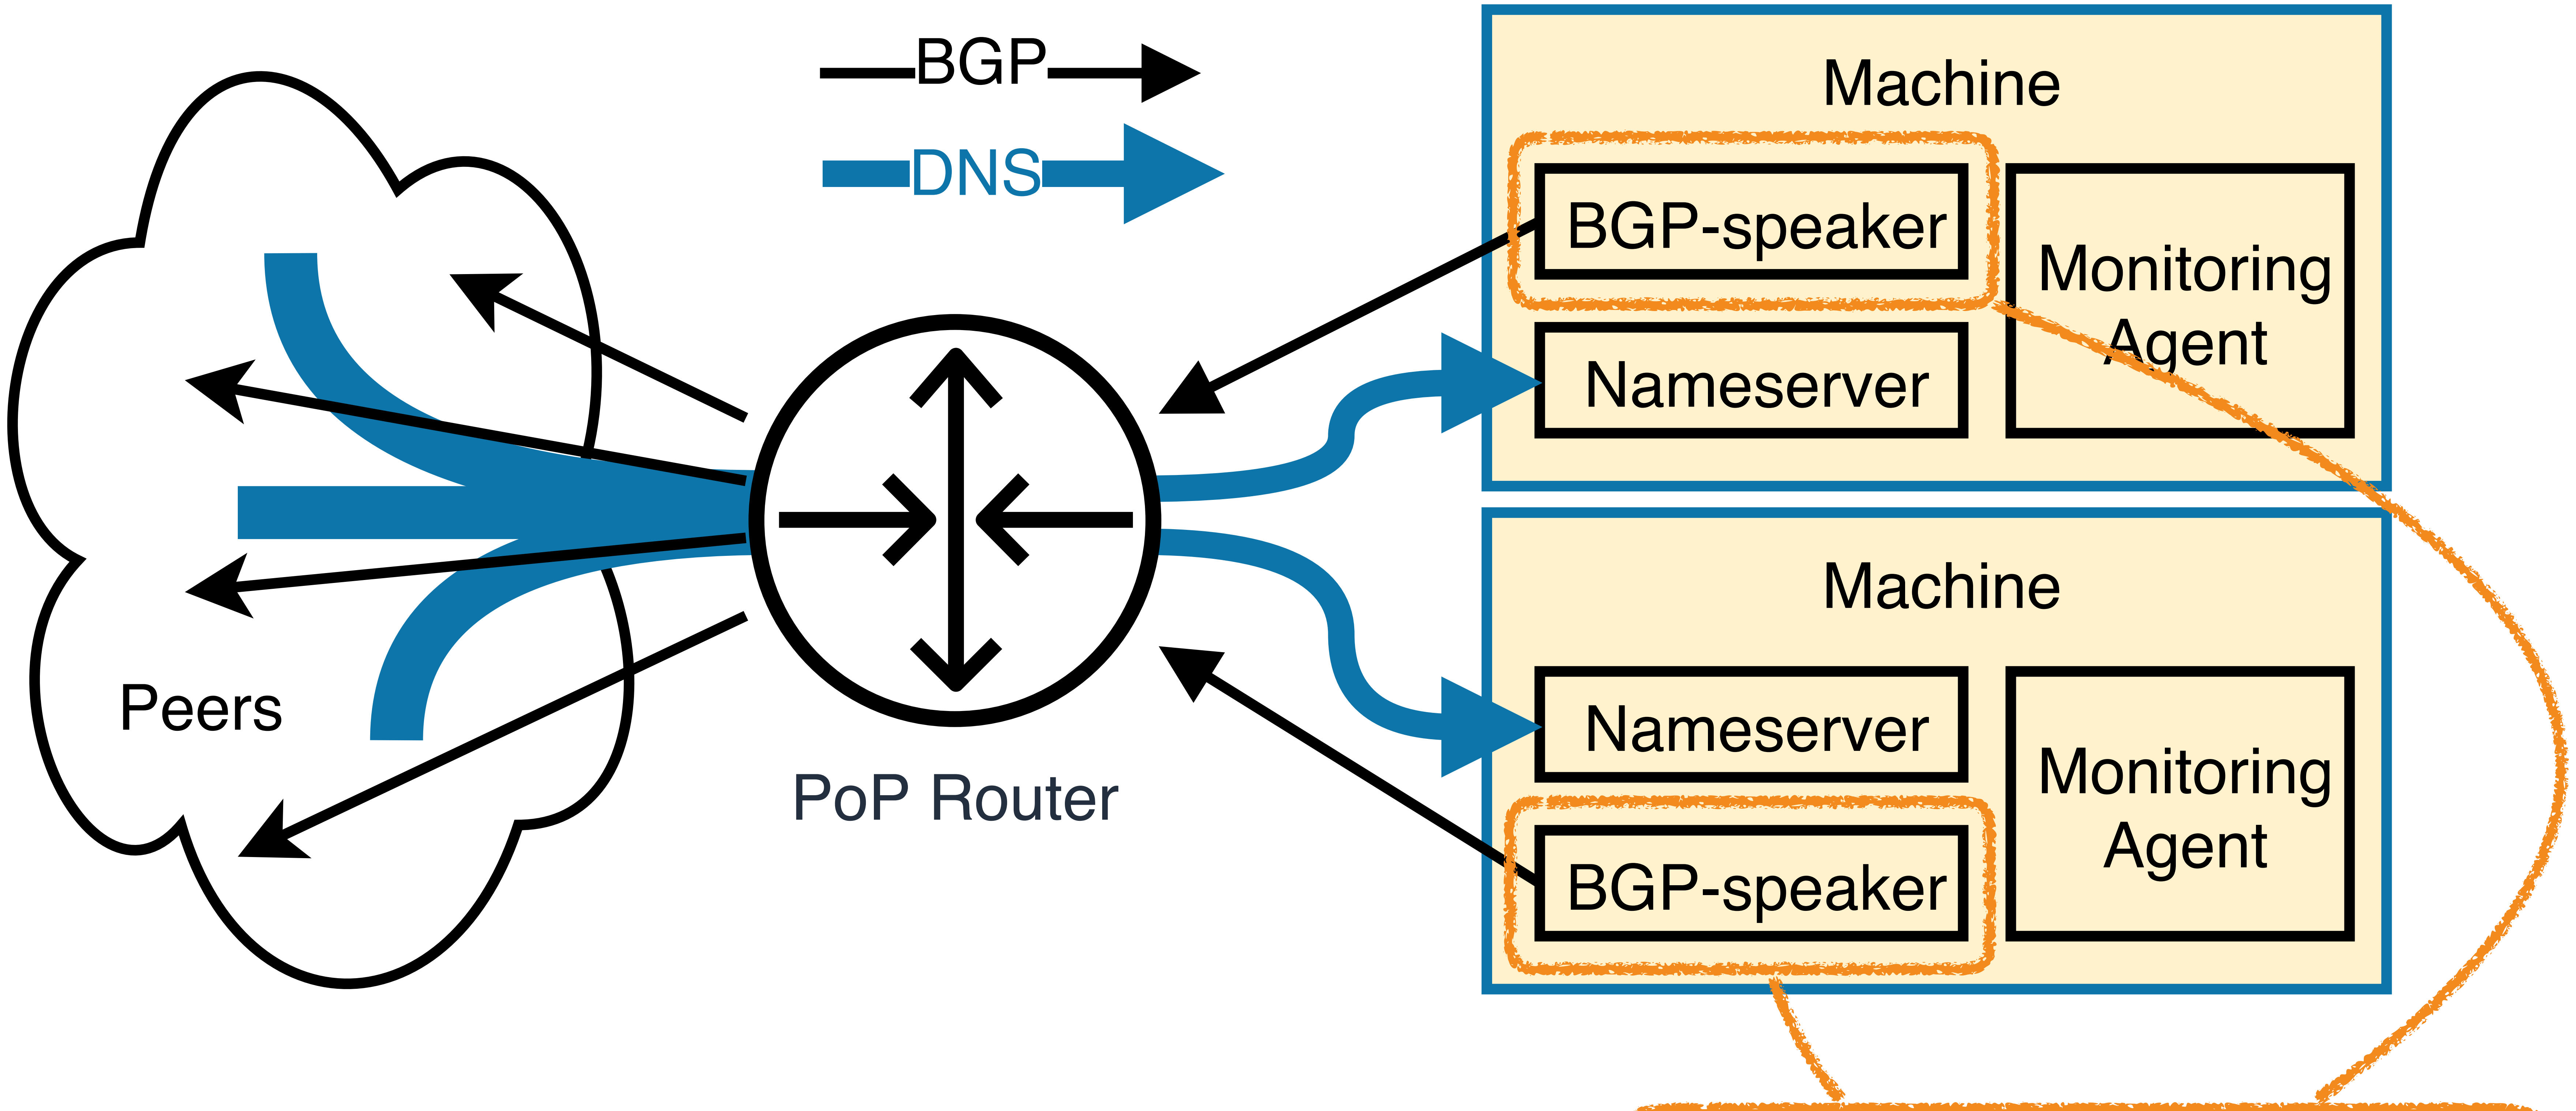

advertise anycast prefixes to router...

One PoP

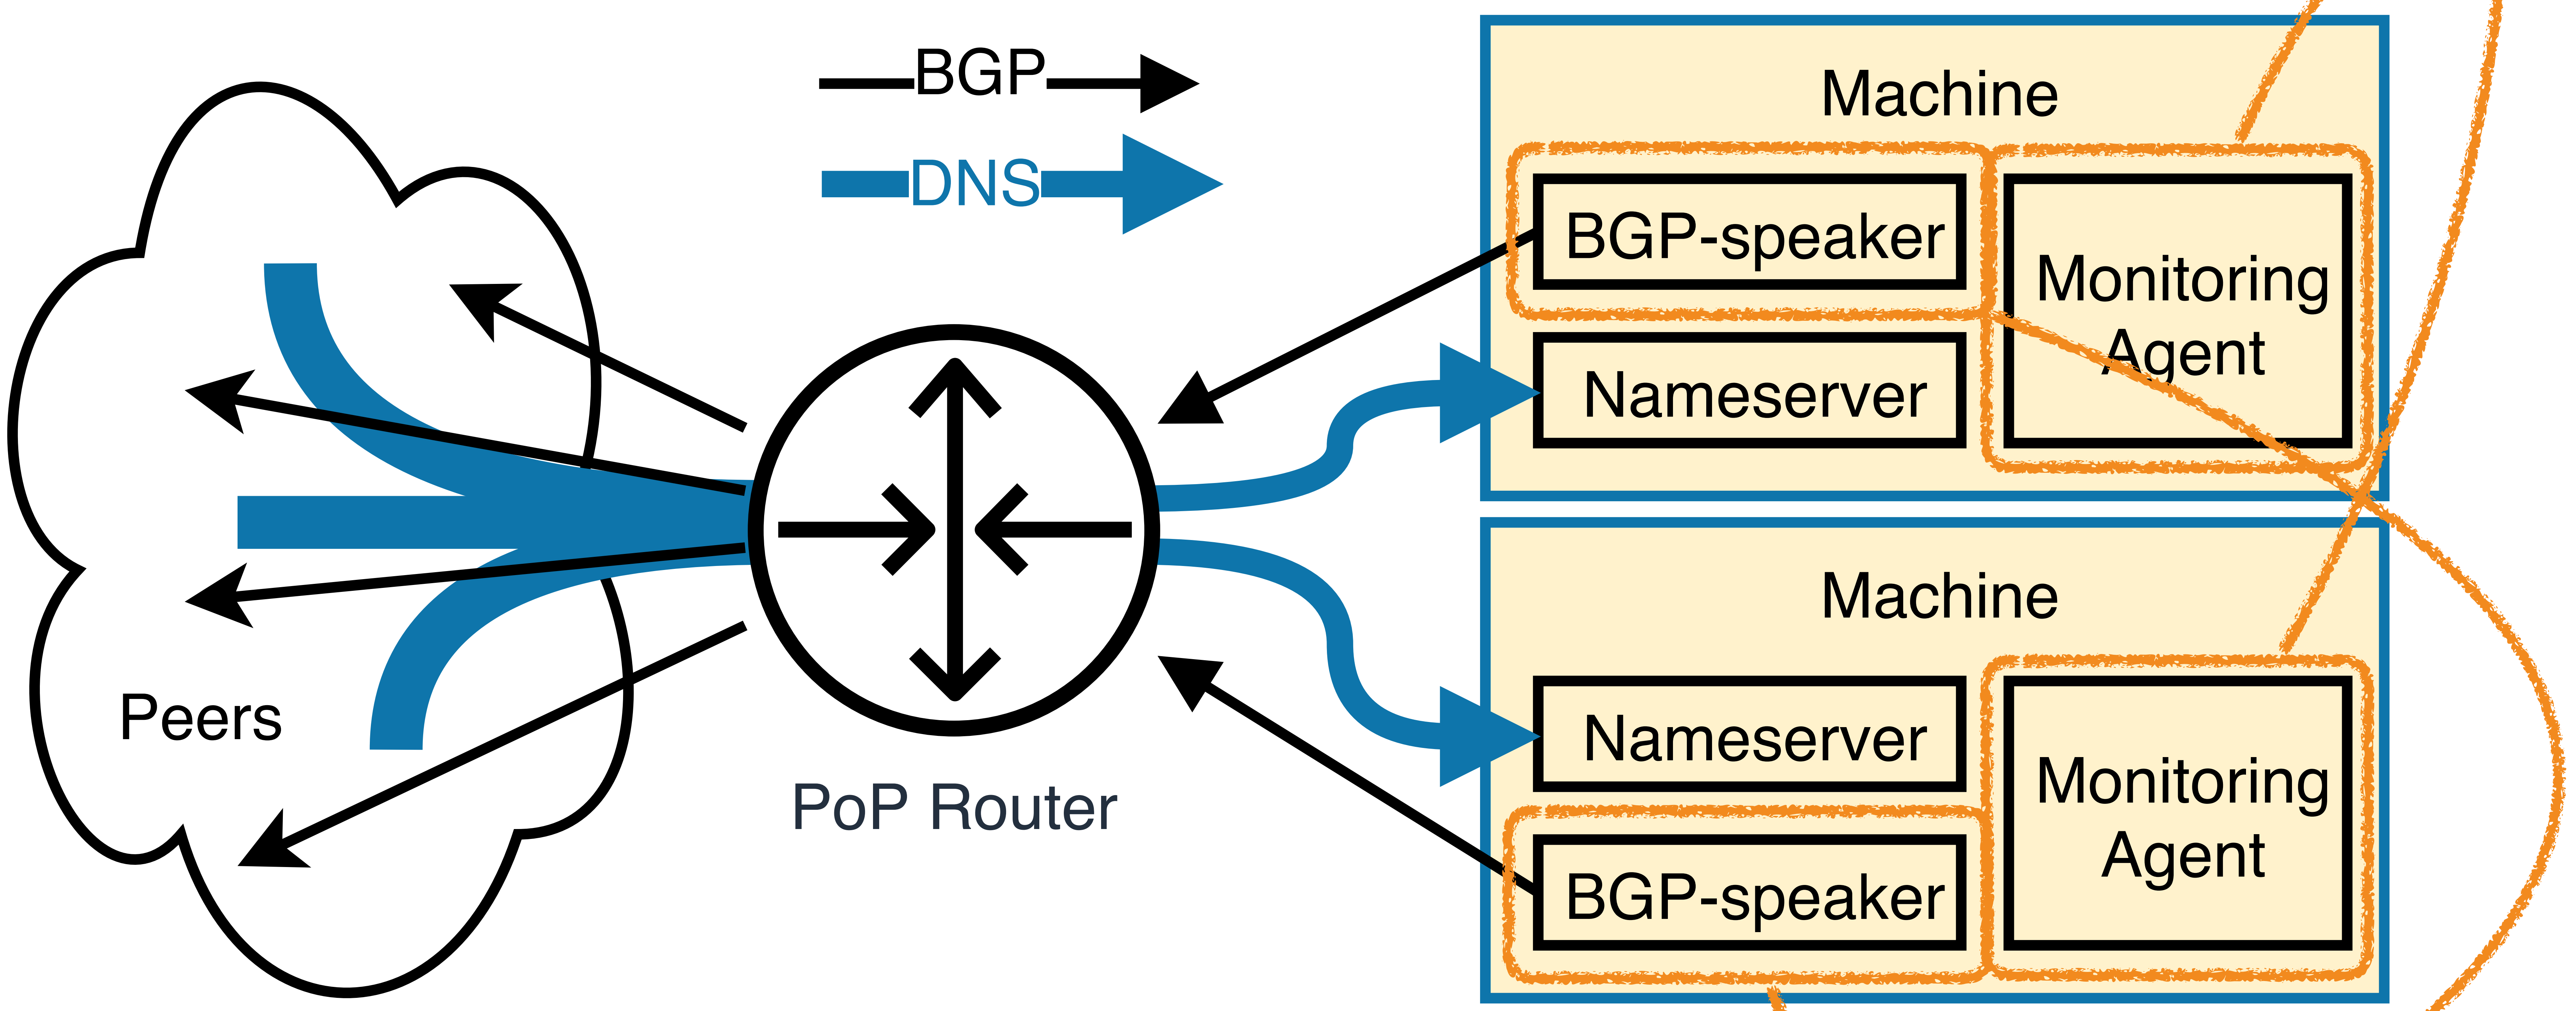

...if nameserver is healthy

advertise anycast prefixes to router...

One PoP

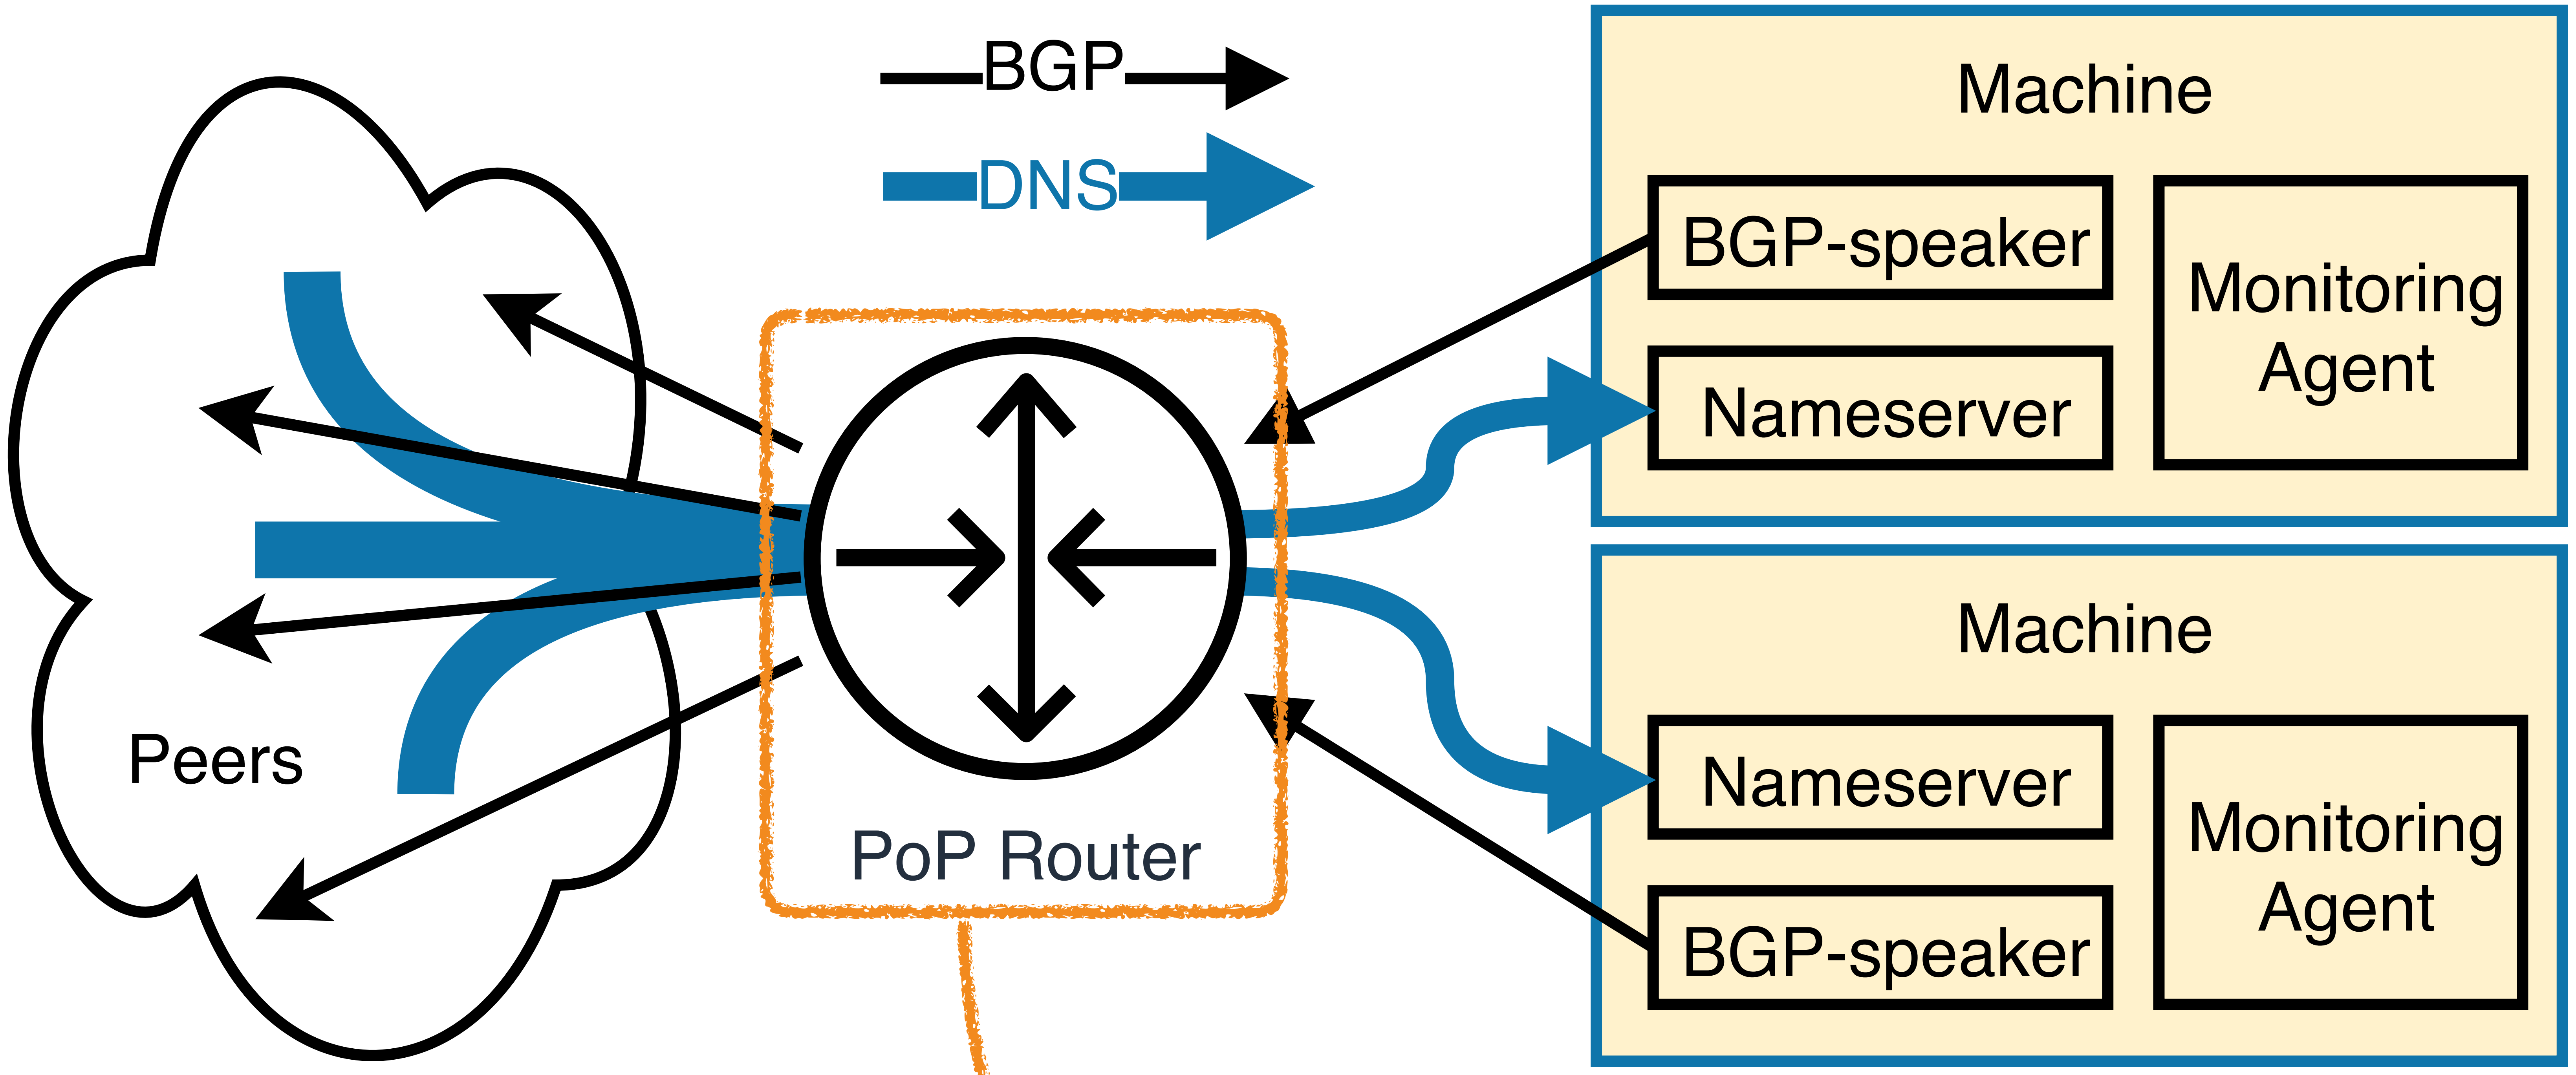

router advertises anycast prefixes to N peers

DNS traffic from peers spread using Equal-Cost-MultiPath (ECMP)

One PoP

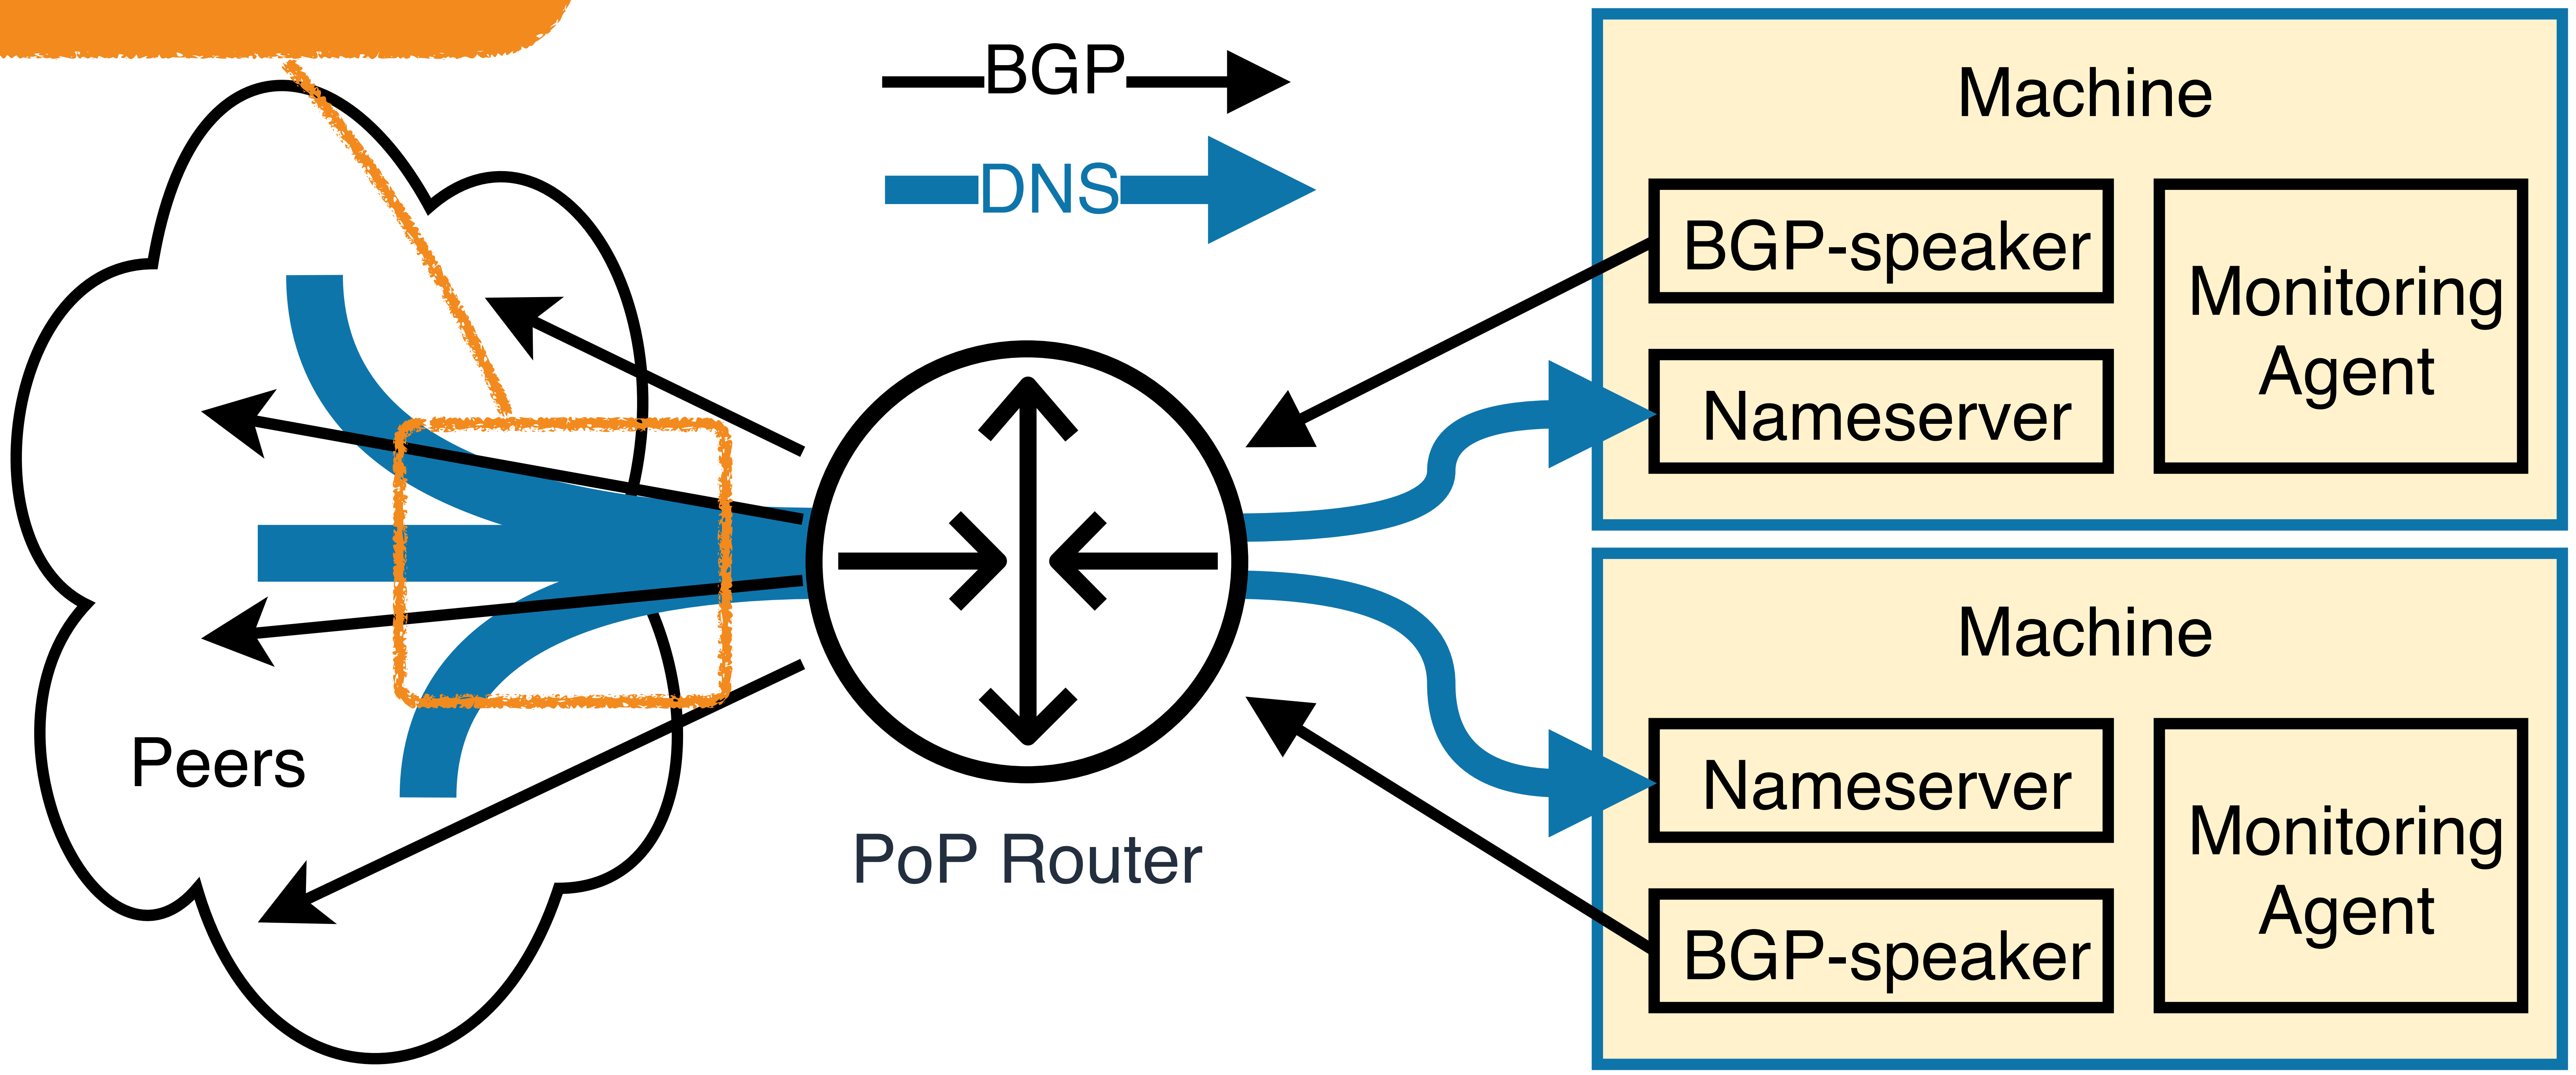

Akamai DNS

Failure Resiliency

Attack Resiliency

Sources of Failure

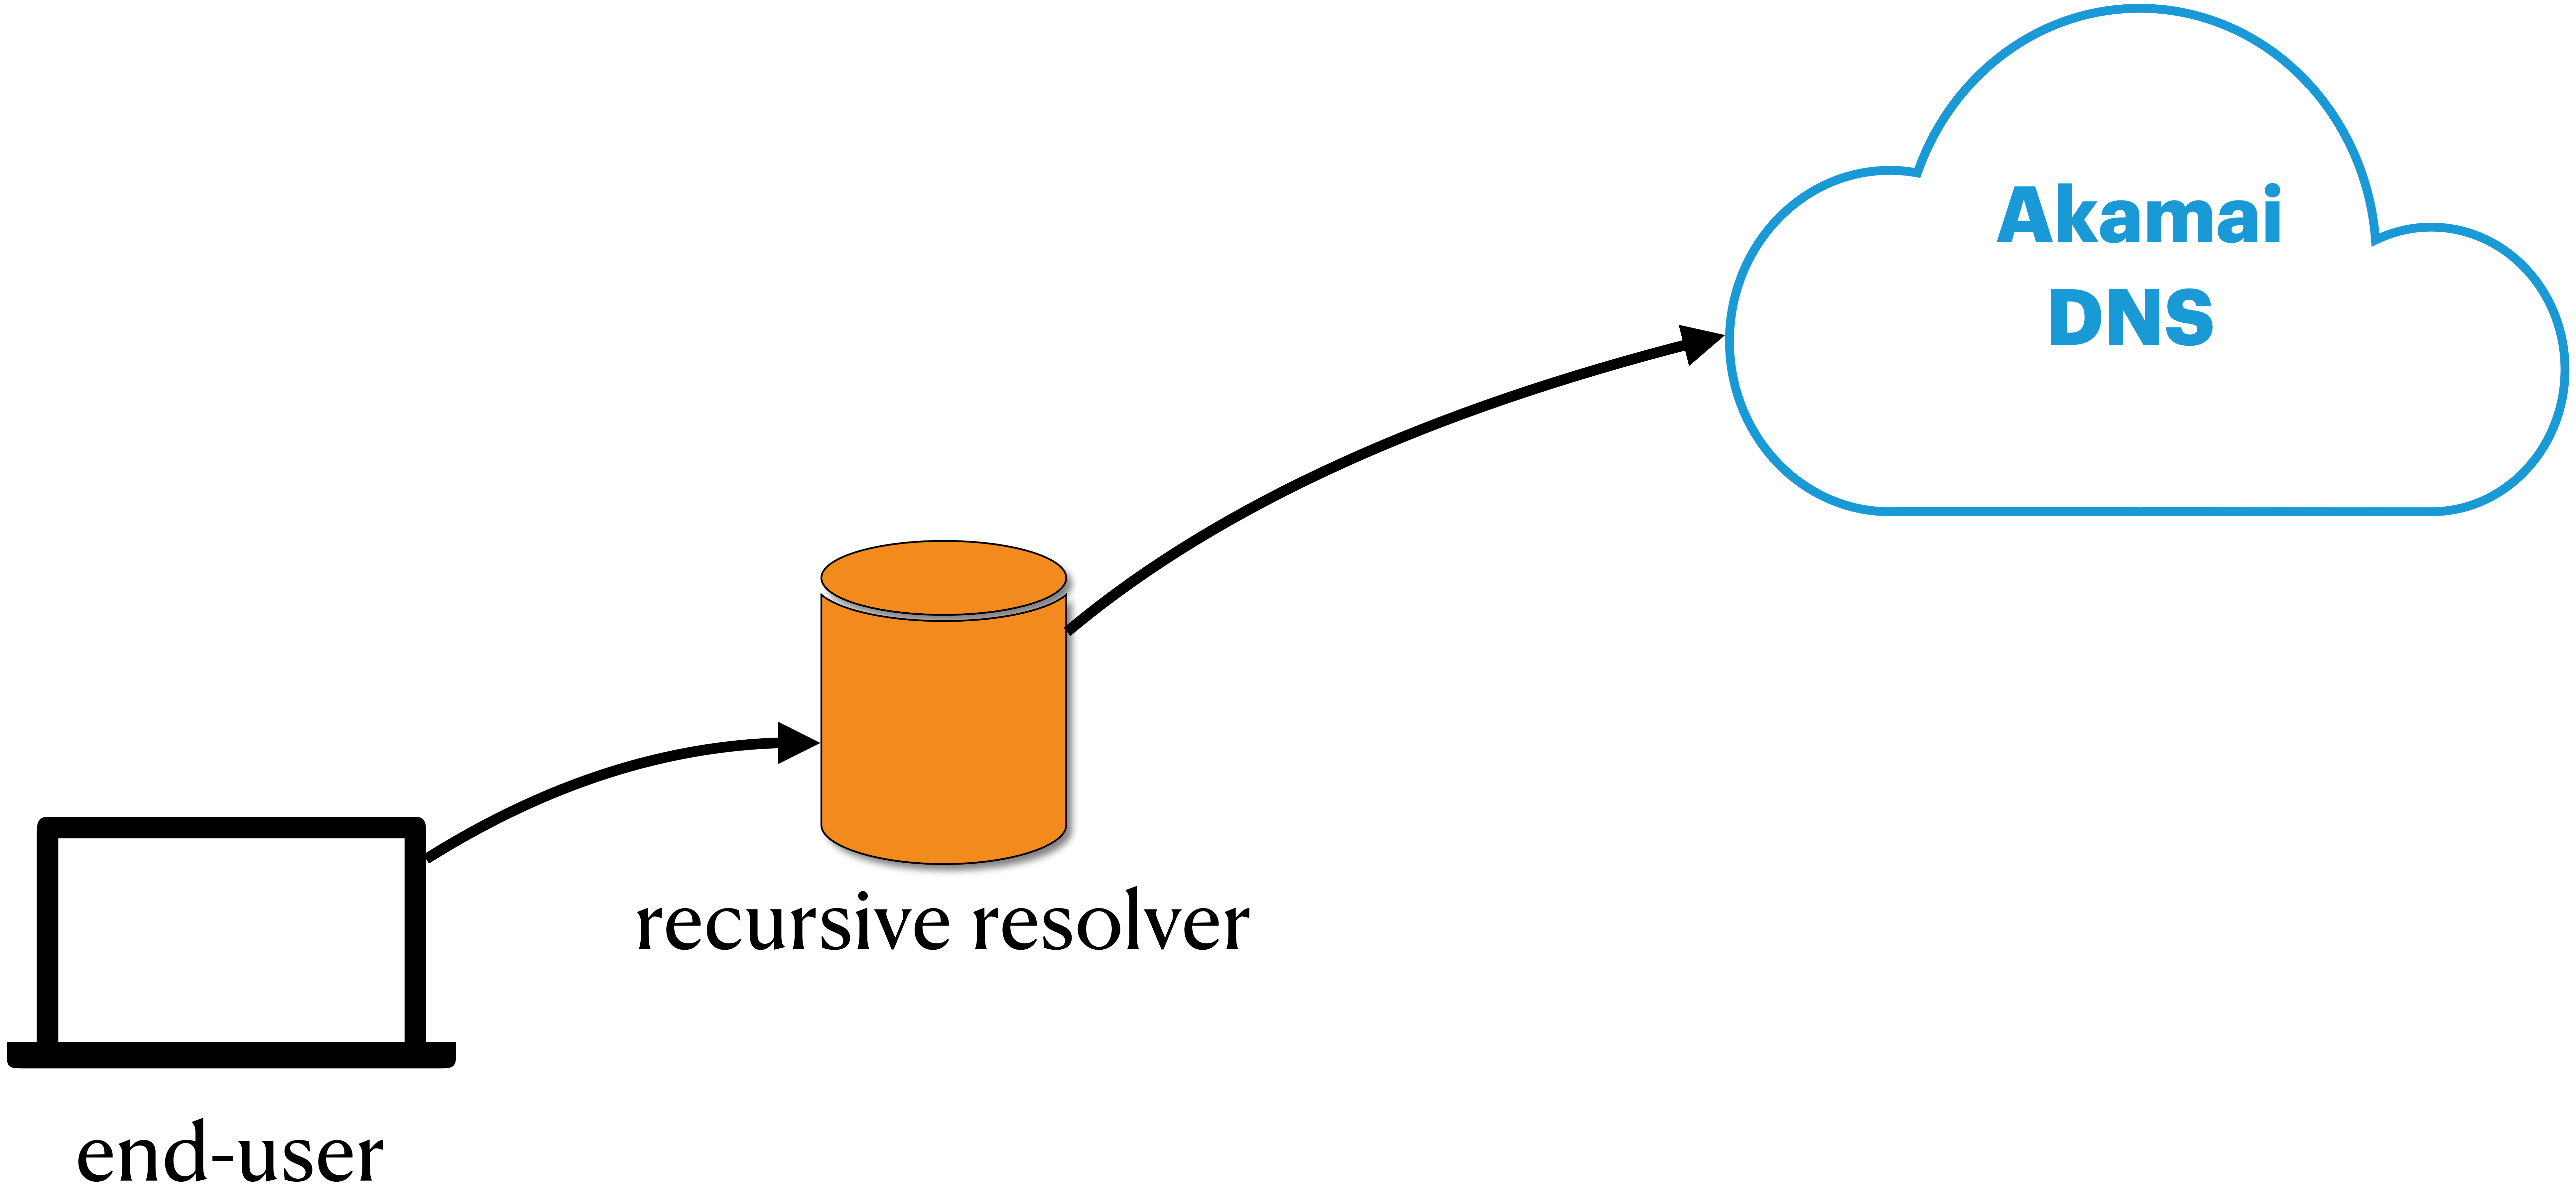

Sources of Failure

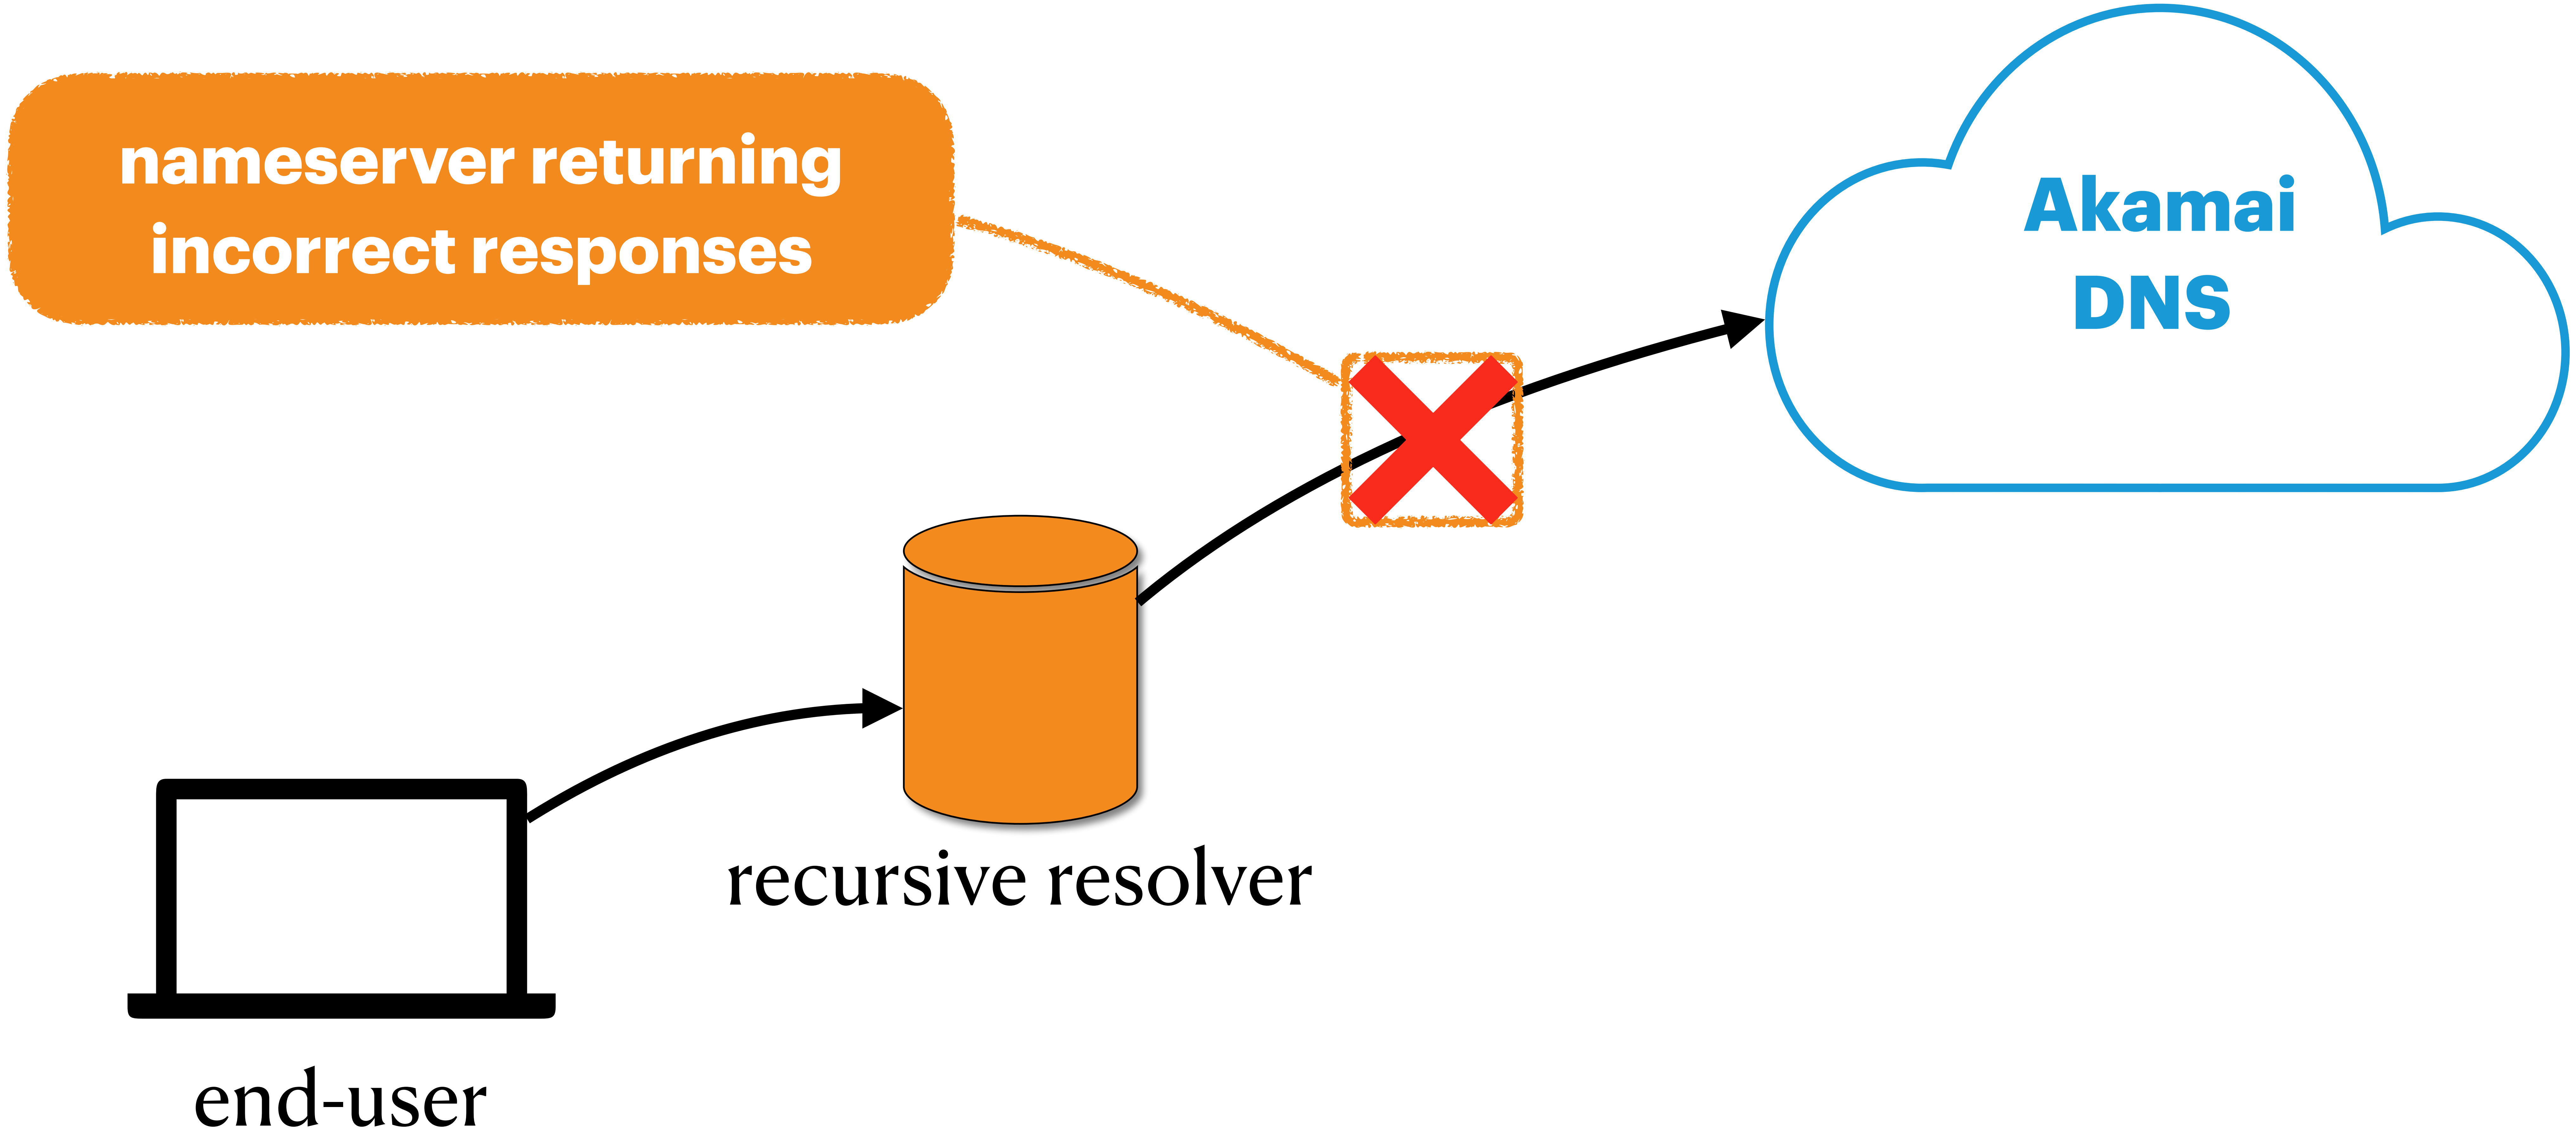

nameserver returning incorrect responses

Akamai DNS

recursive resolver

end-user

Incorrect Response Mitigation

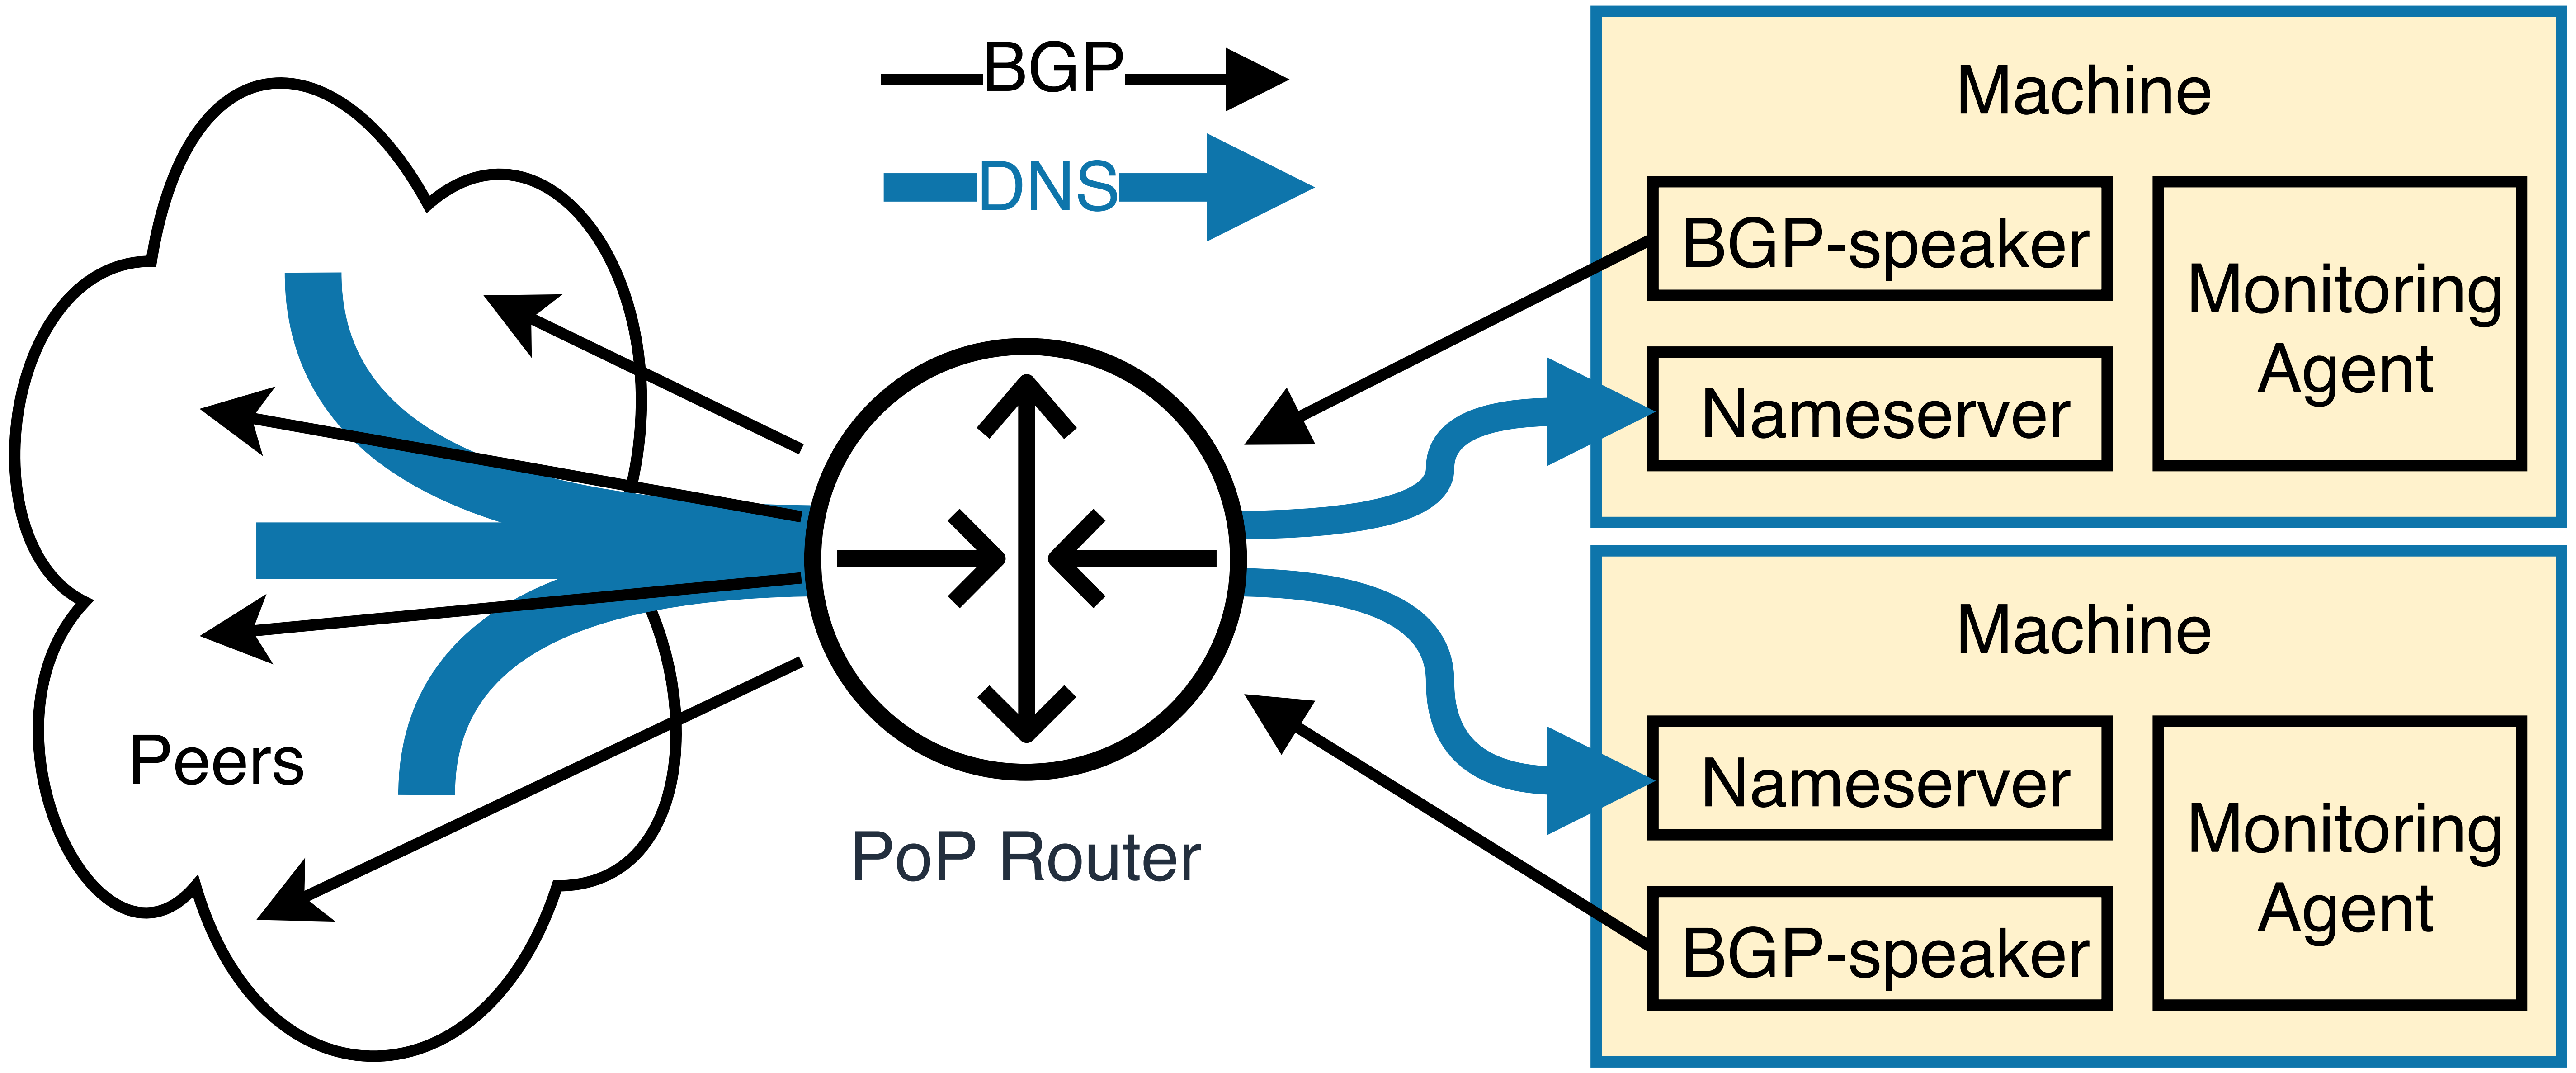

Incorrect Response Mitigation

1. monitoring agent detects issue

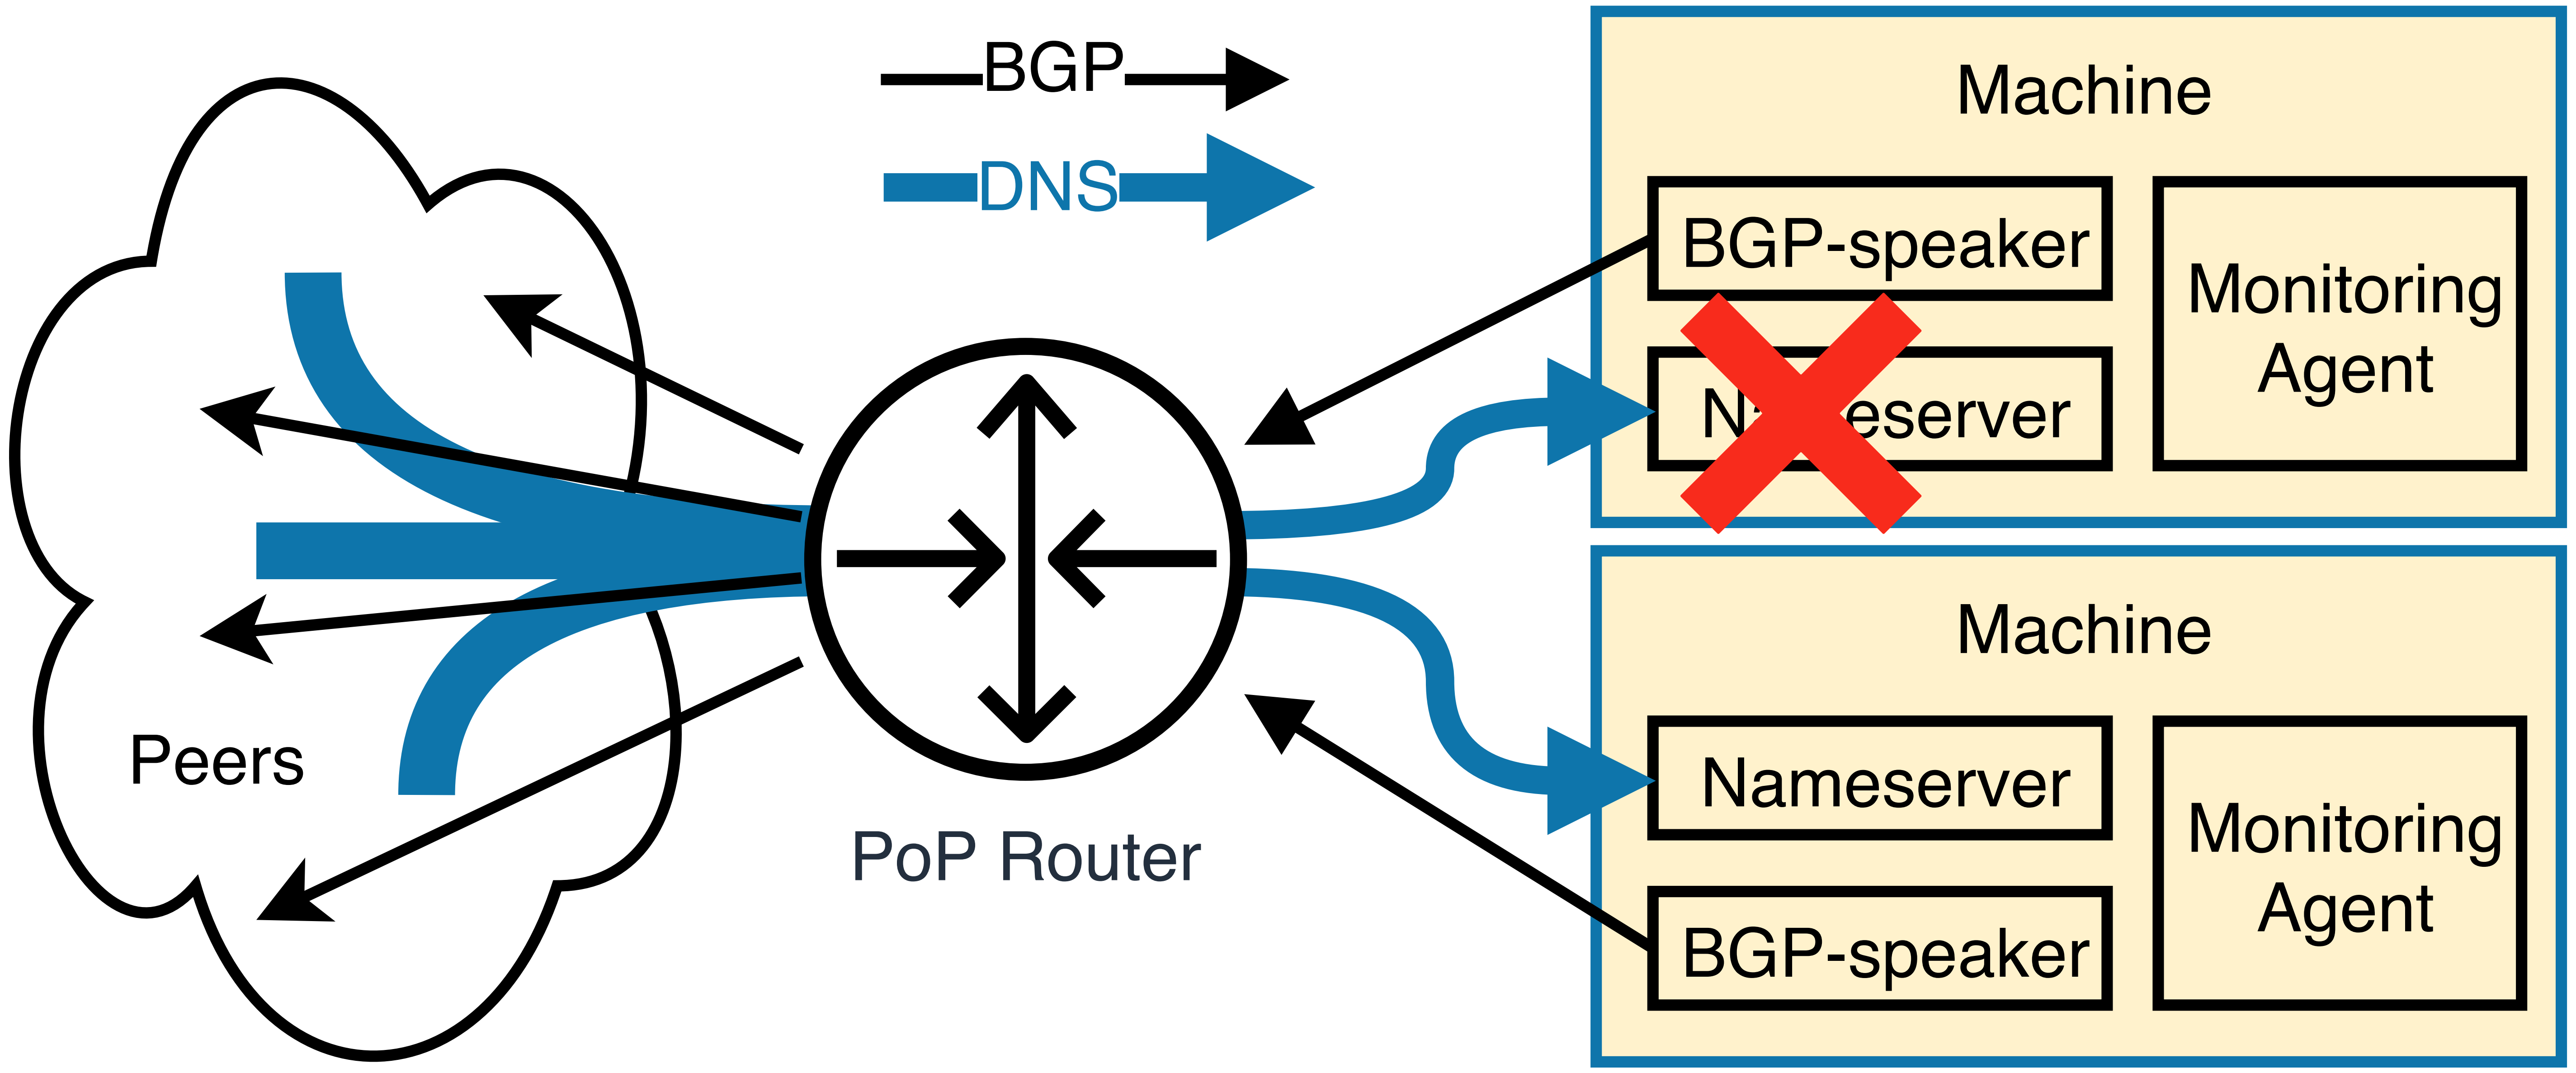

Incorrect Response Mitigation

1. monitoring agent detects issue
2. withdraws BGP advertisement

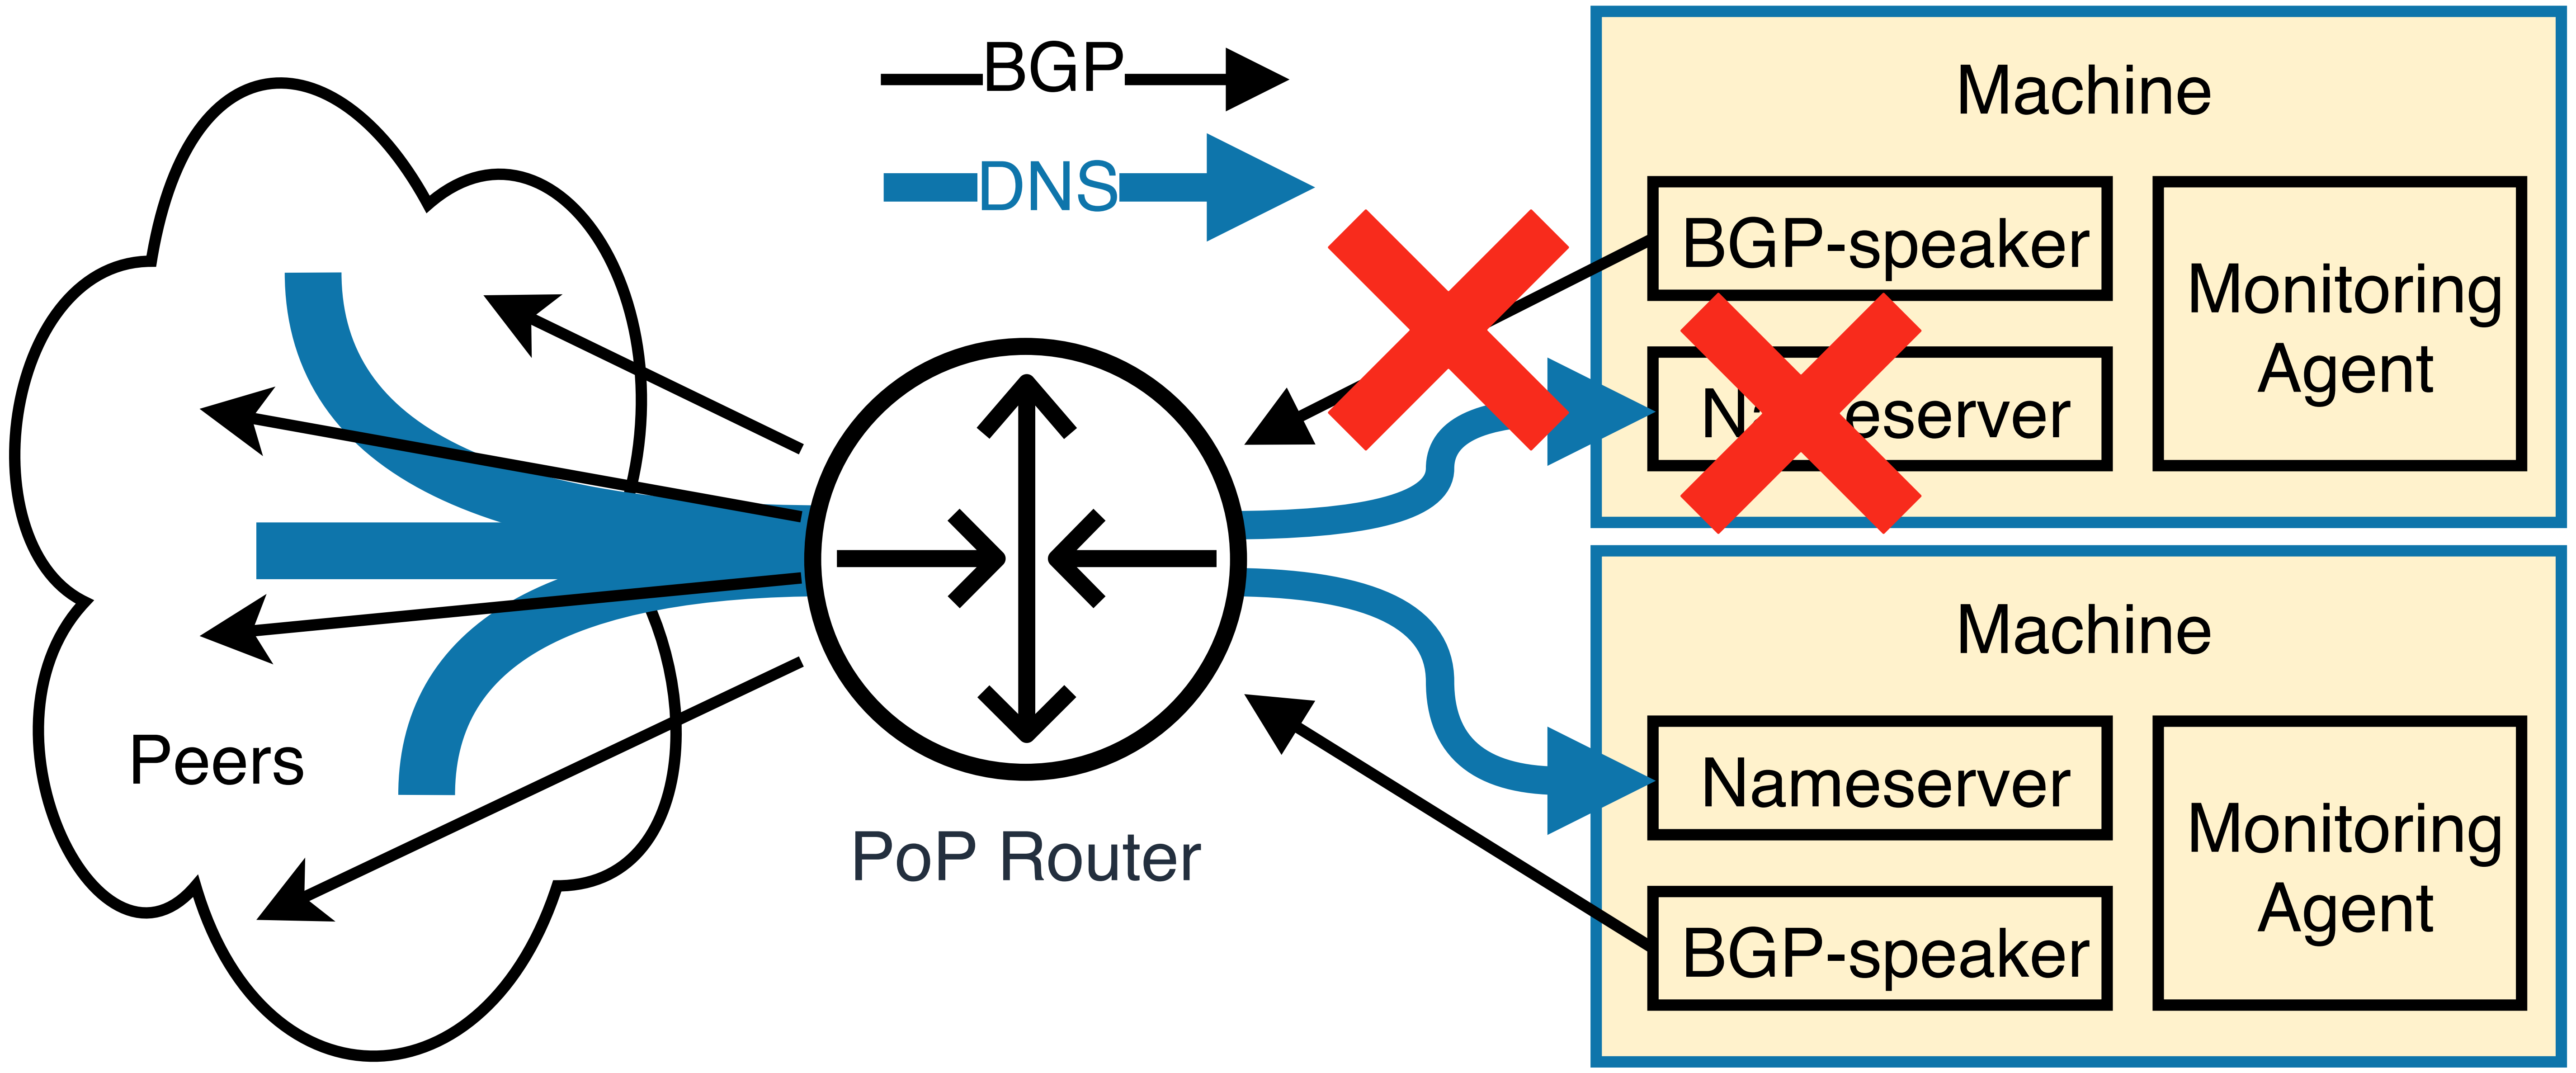

Incorrect Response Mitigation

1. monitoring agent detects issue
2. withdraws BGP advertisement
3. router forwards traffic to other nameservers in PoP

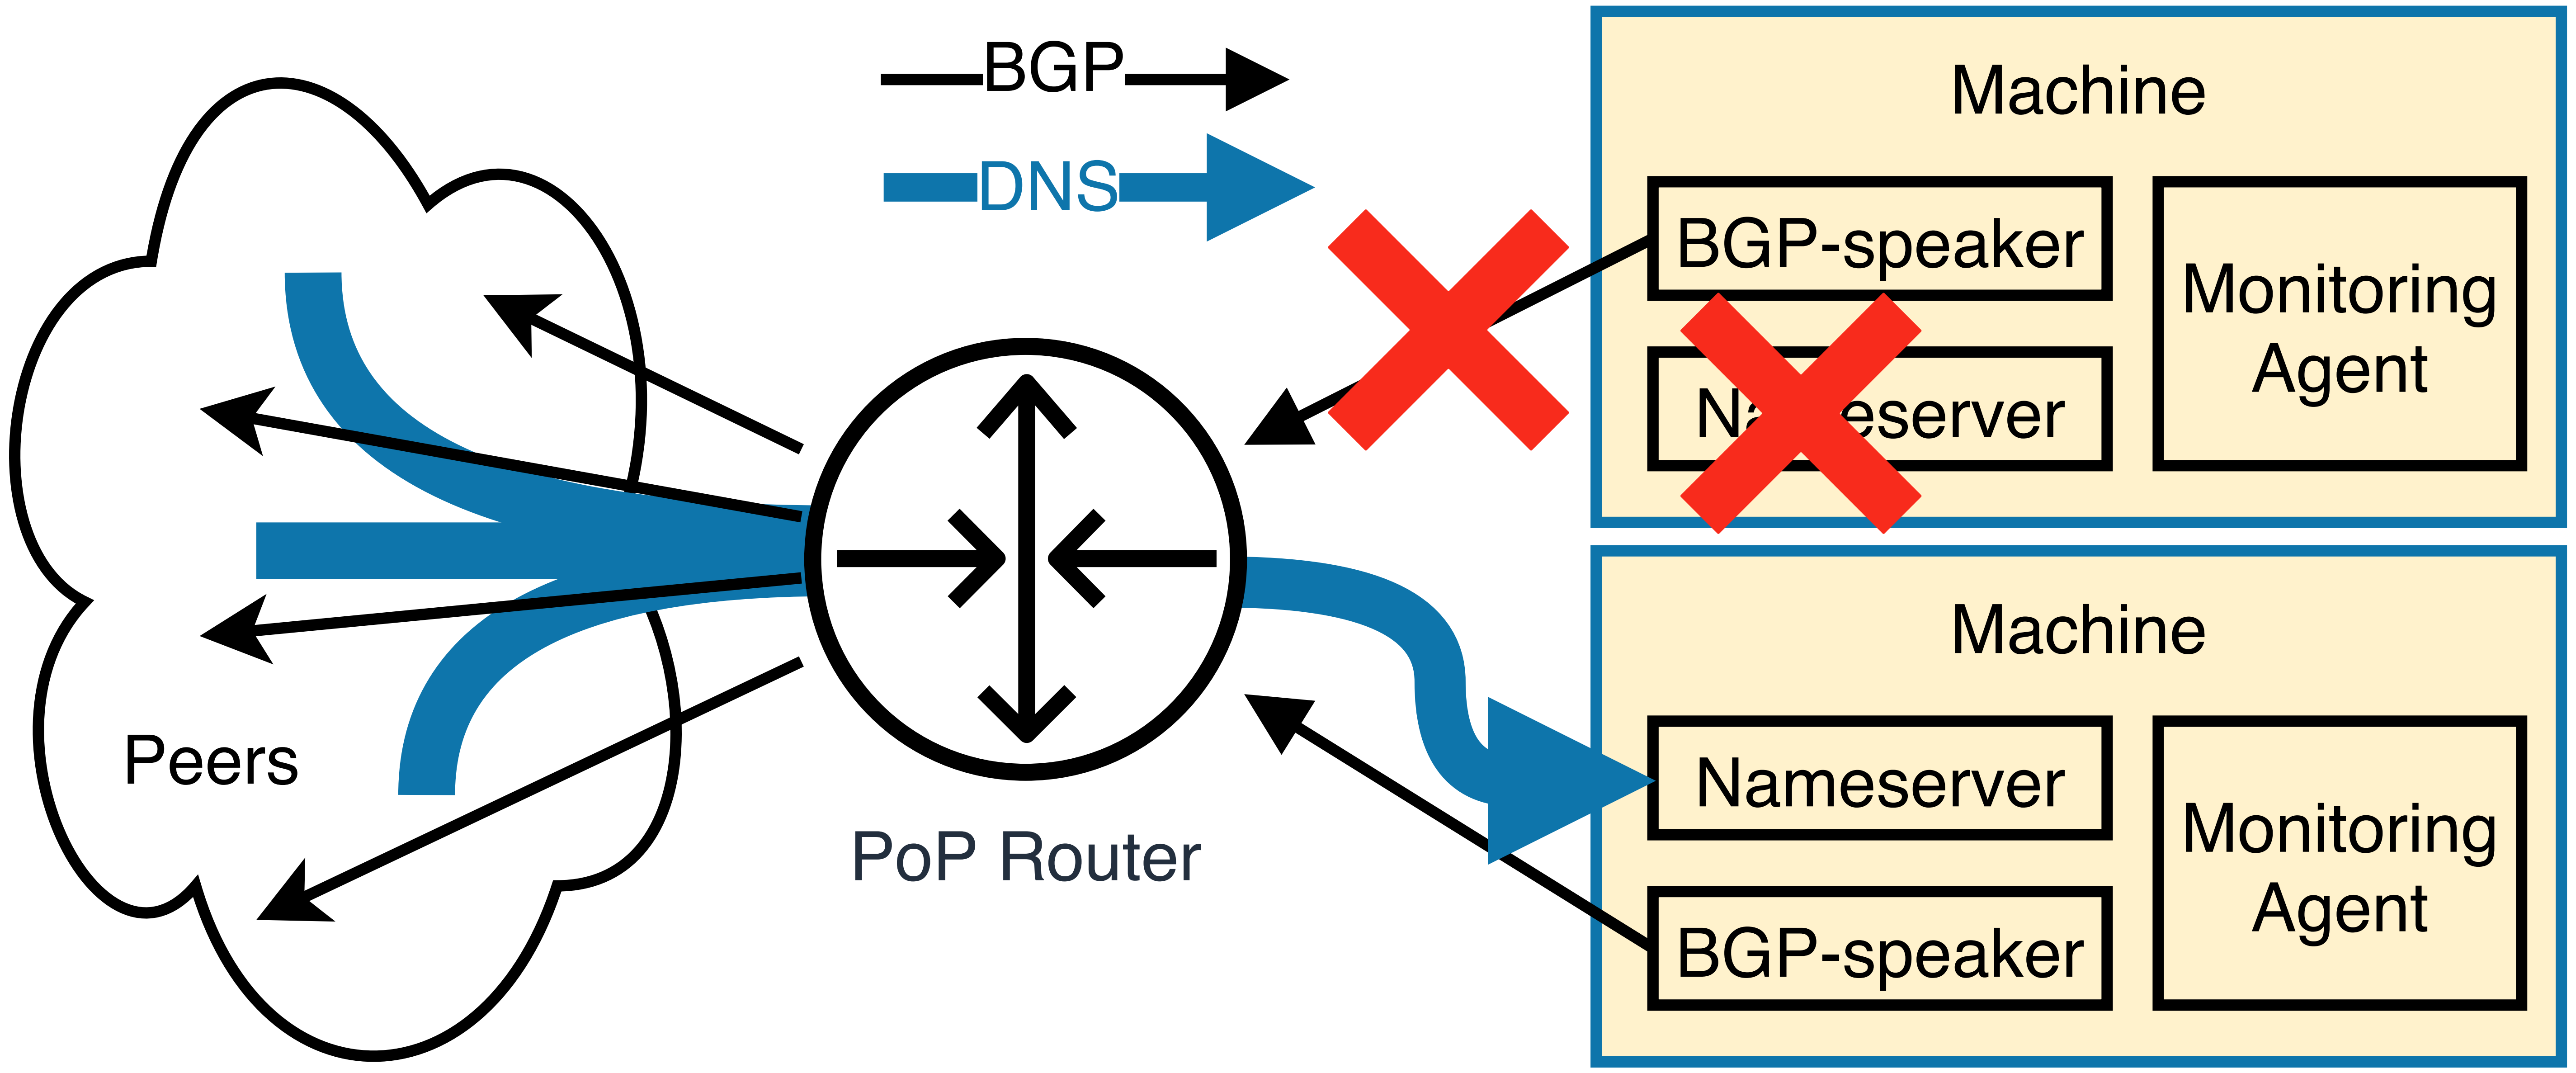

Incorrect Response Mitigation

1. monitoring agent detects issue
2. withdraws BGP advertisement
3. router forwards traffic to other nameservers in PoP
4. if all nameservers withdraw advertisement, router withdraws advertisement from peers

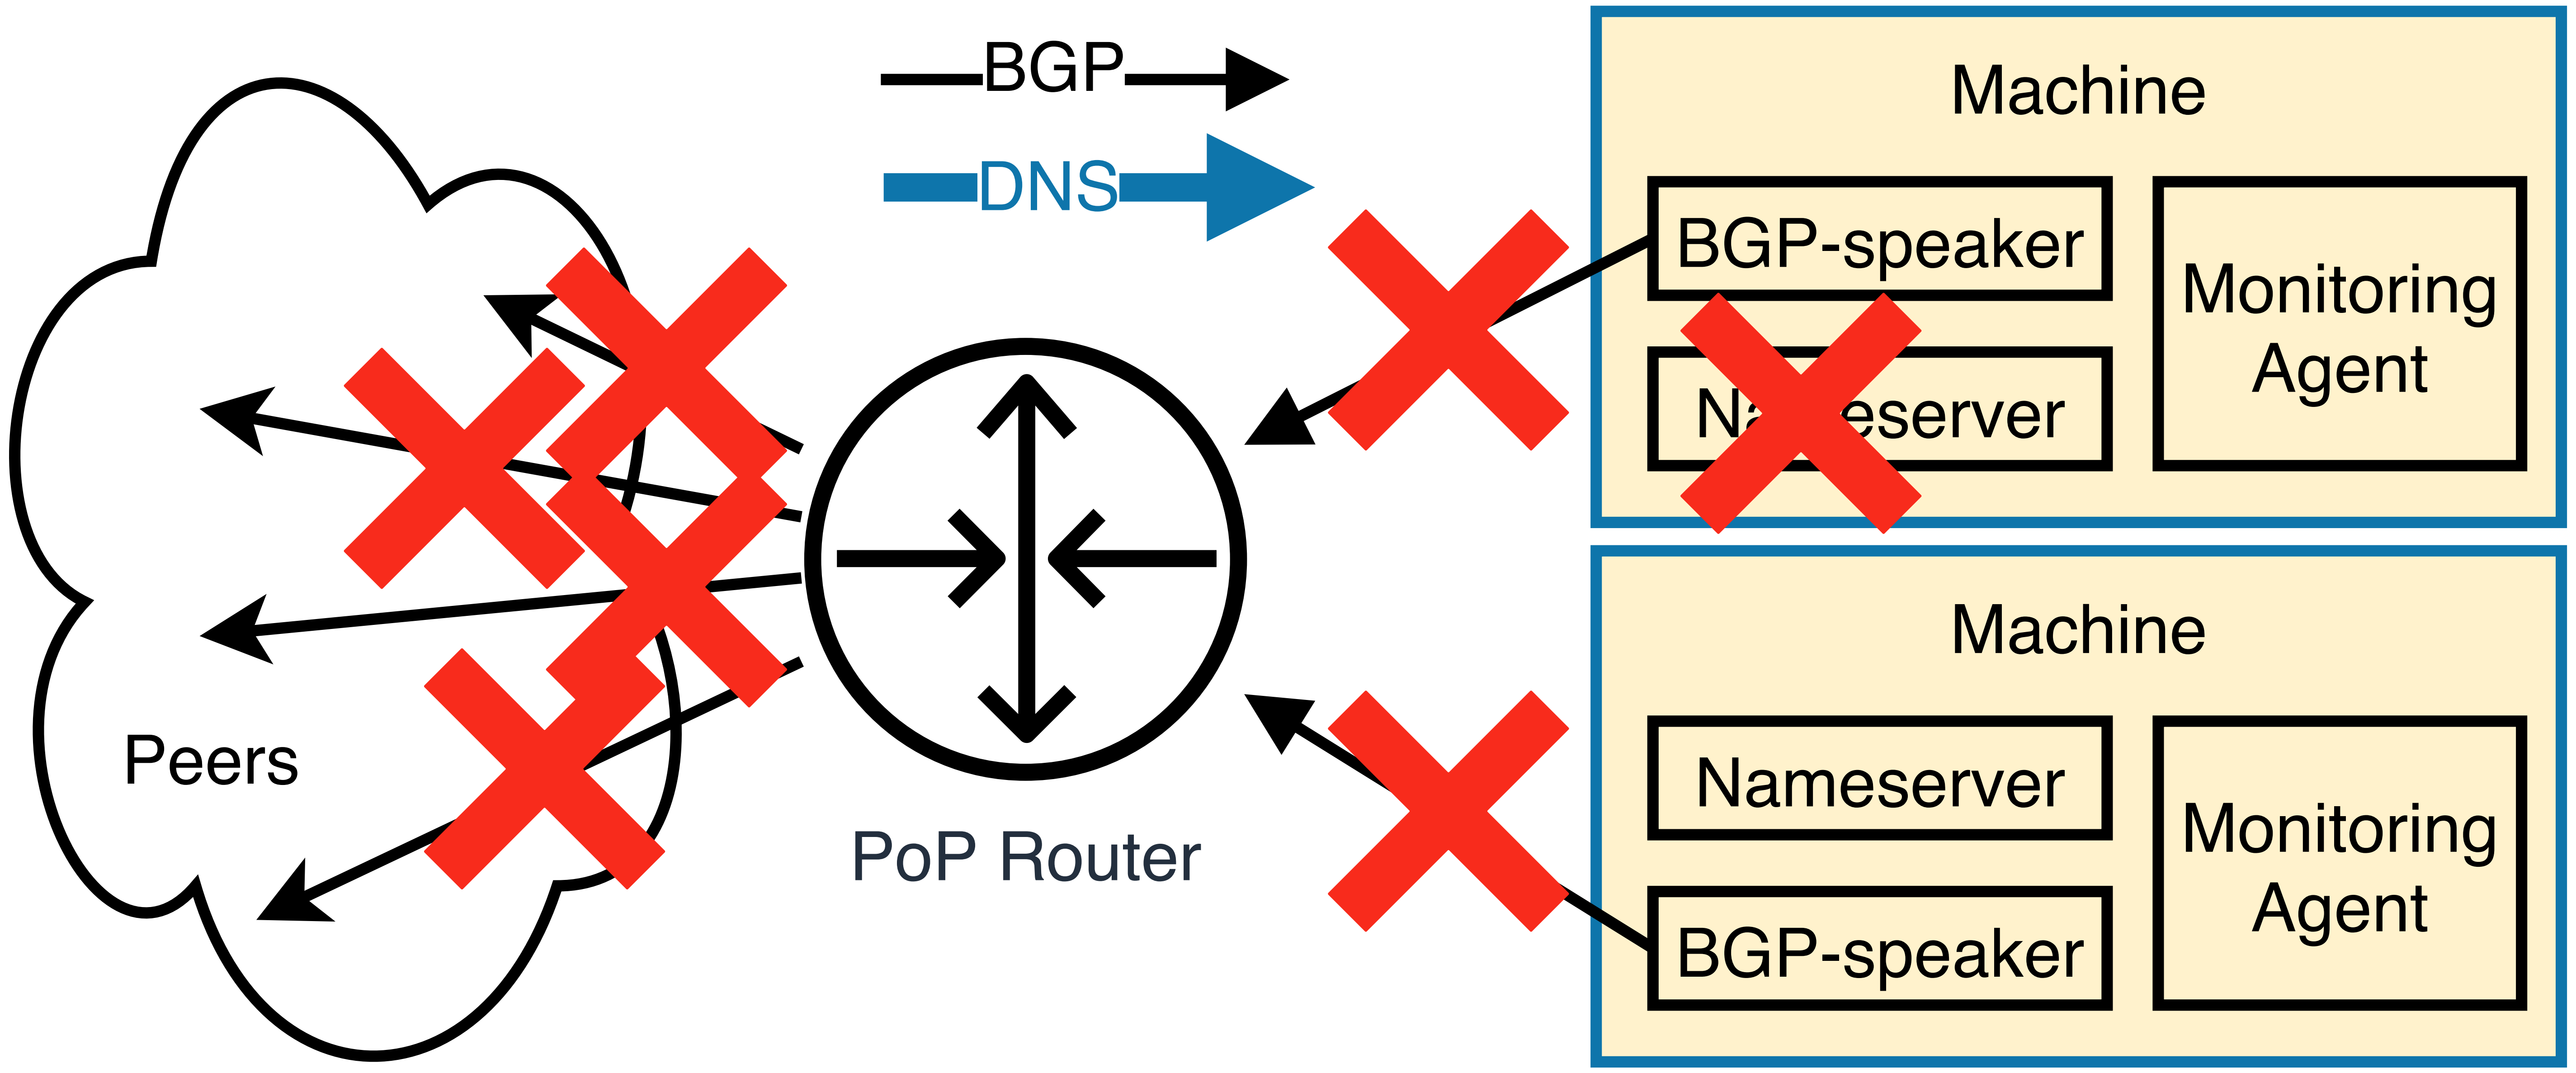

Sources of Failure

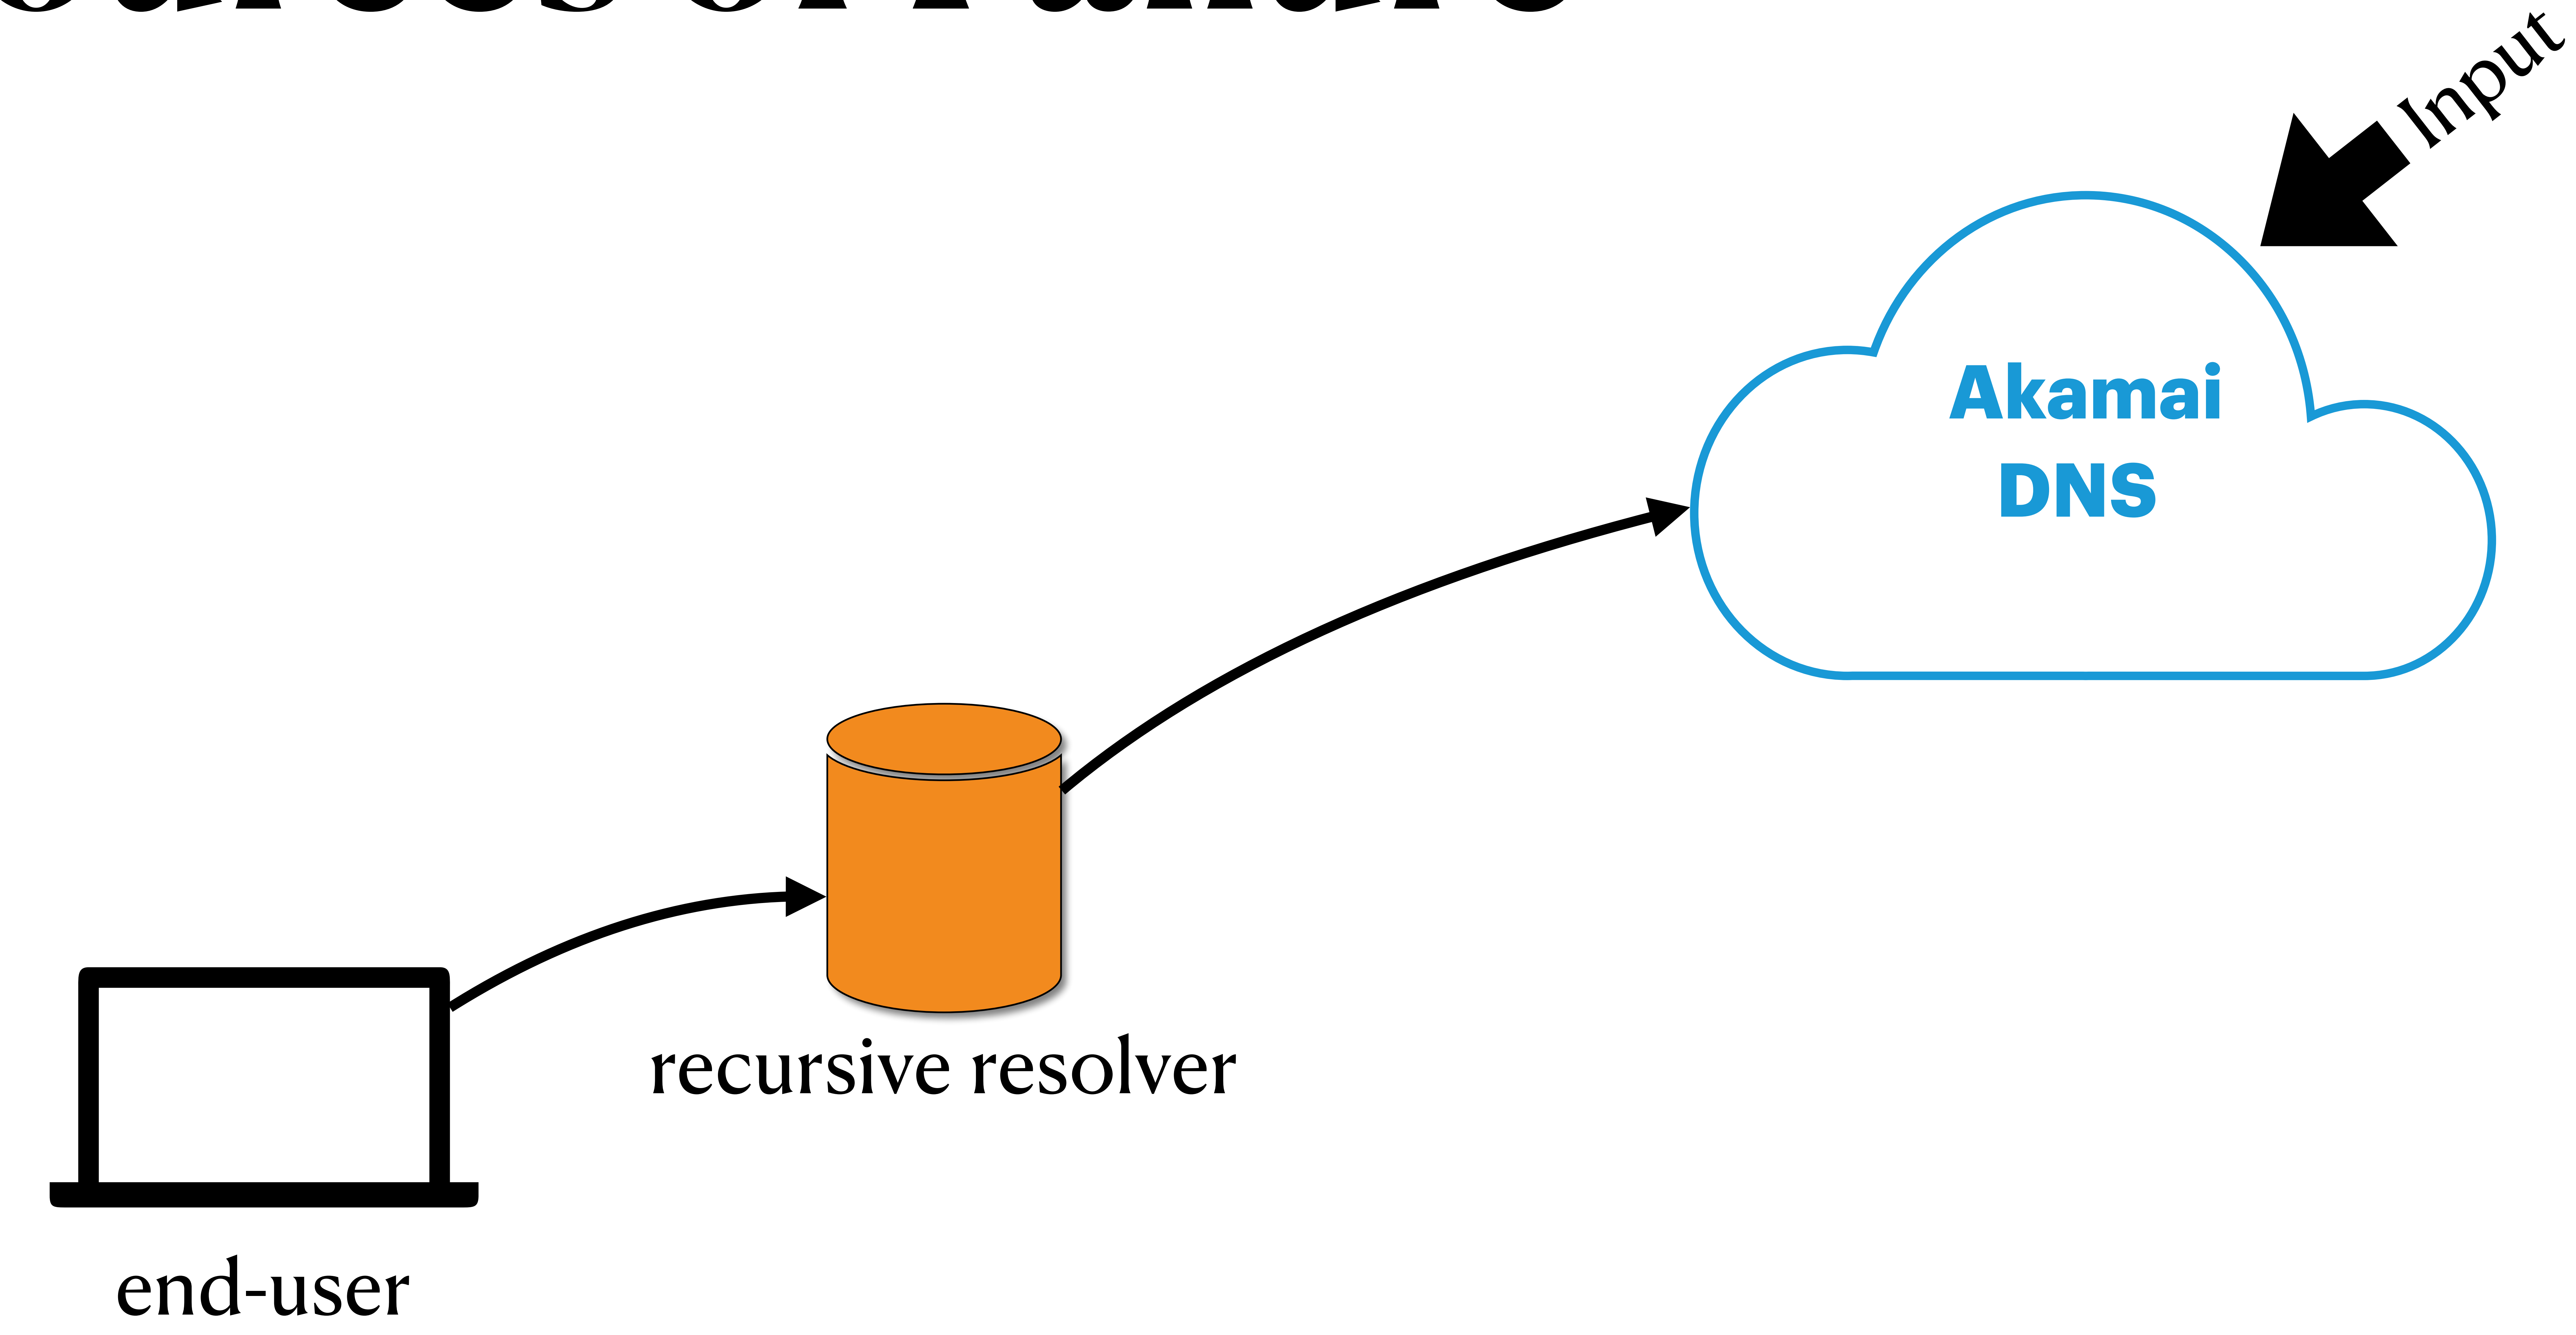

Sources of Failure

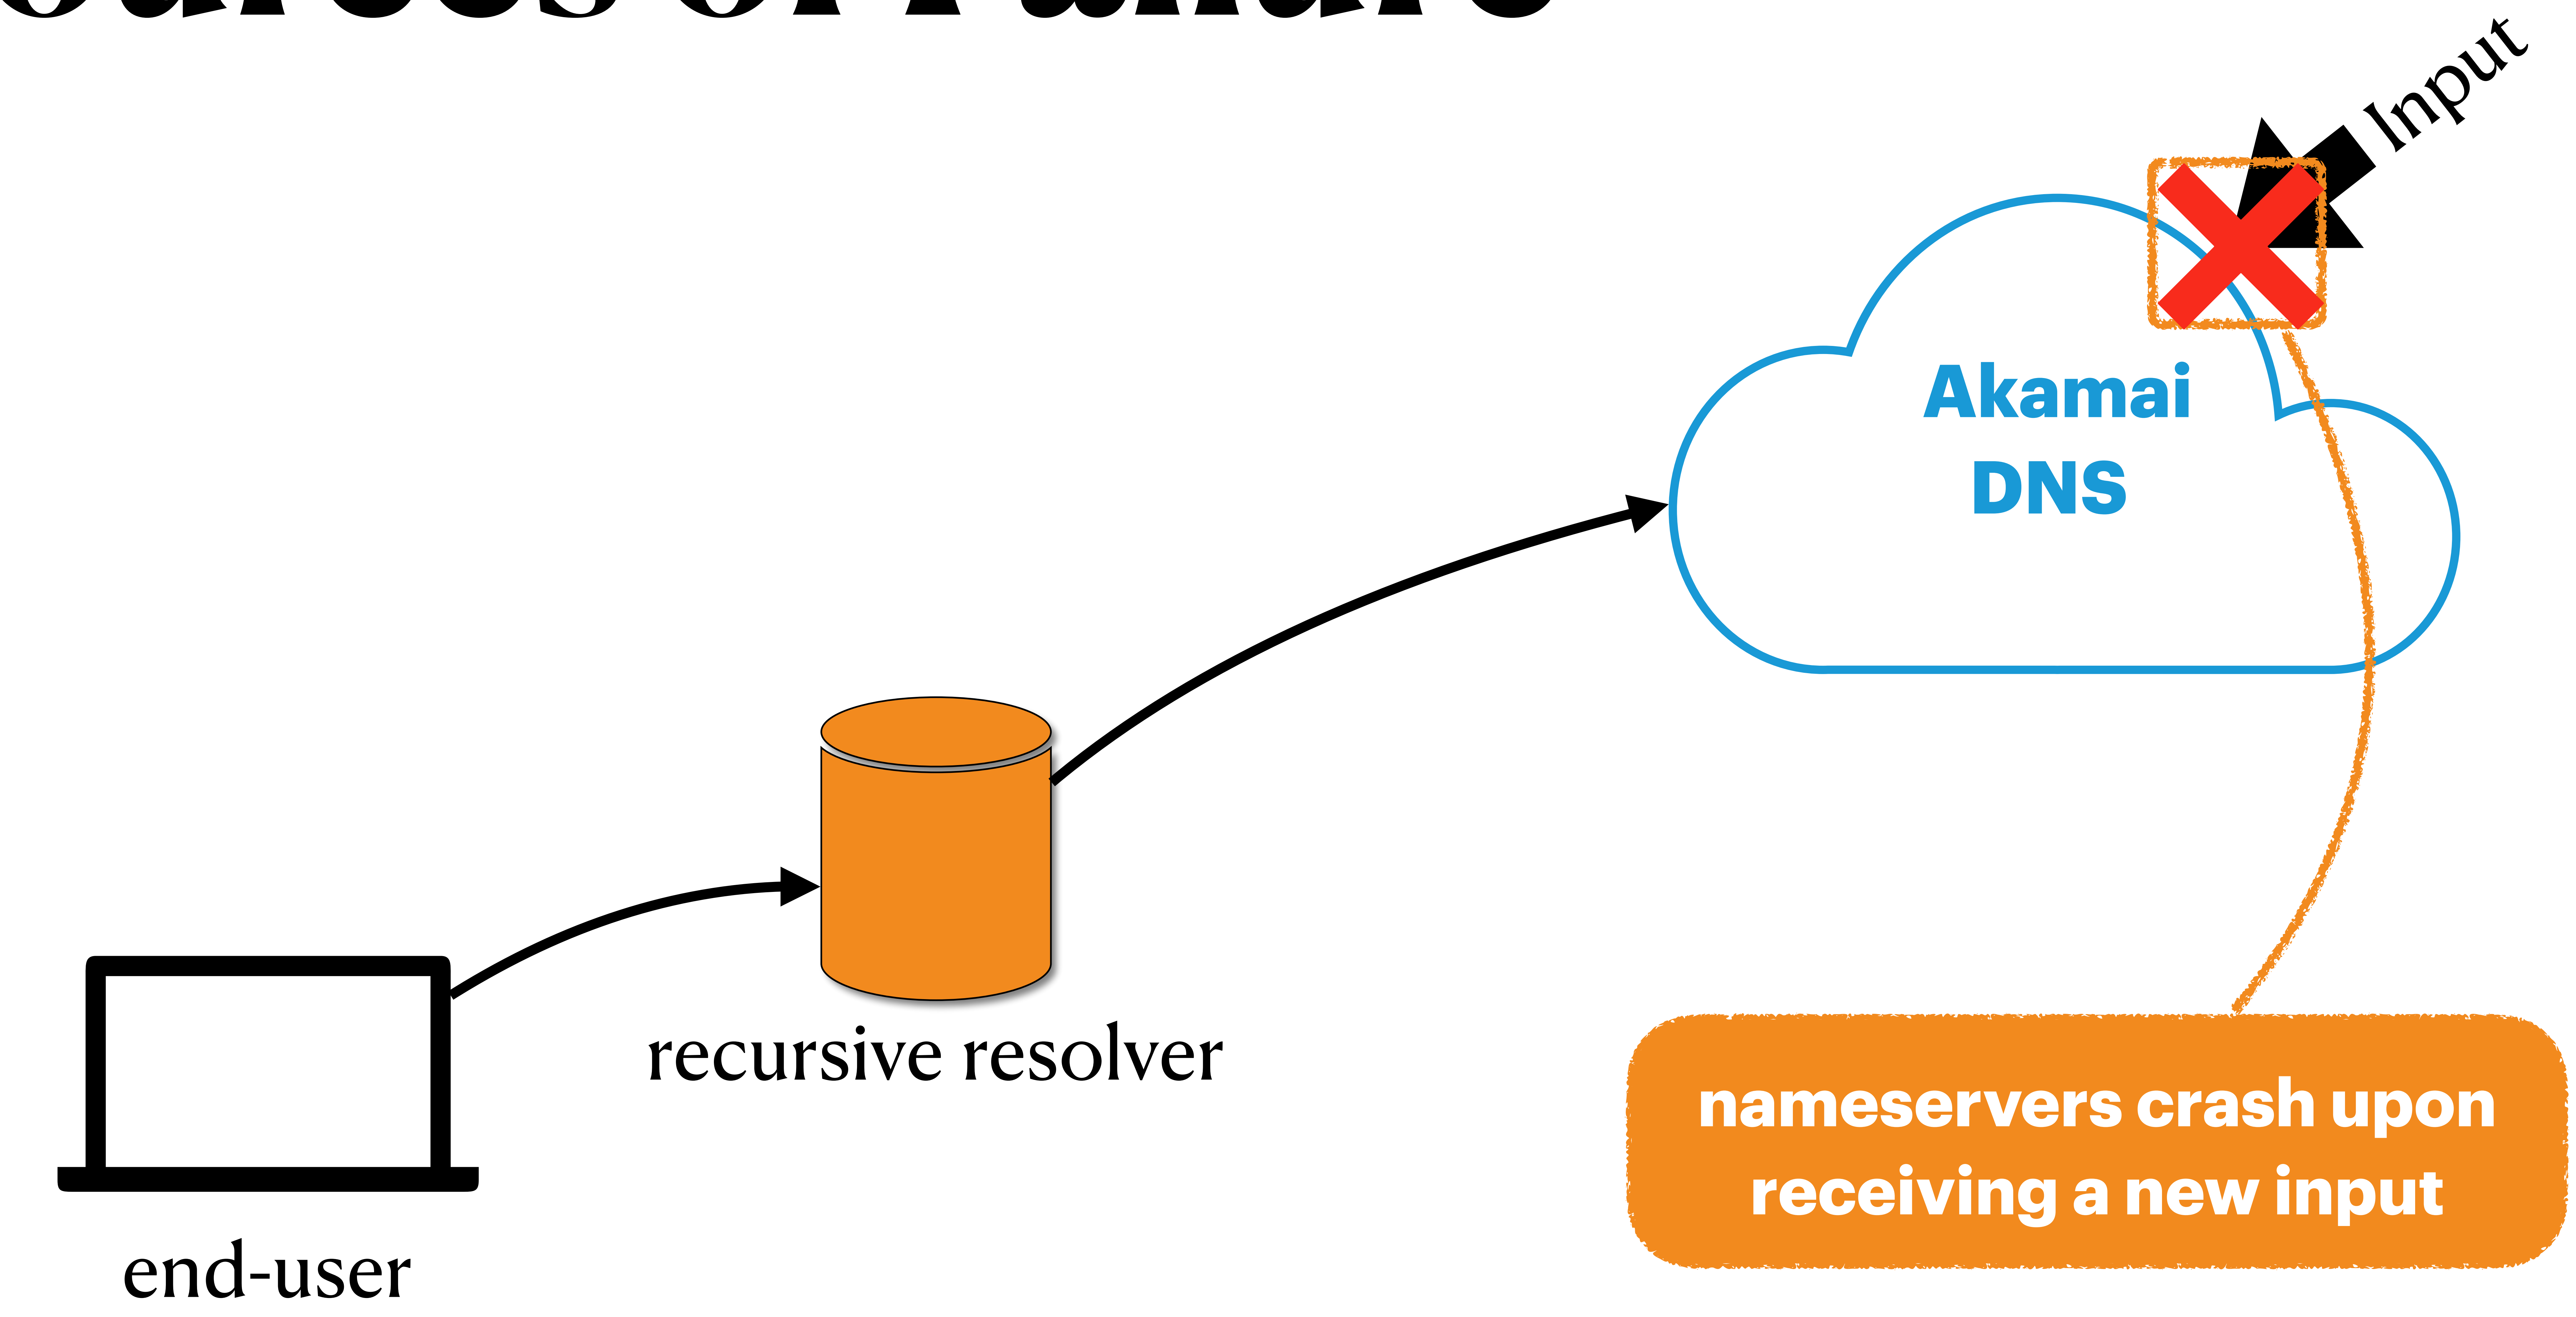

nameservers crash upon receiving a new input

Input-induced Crash Mitigation

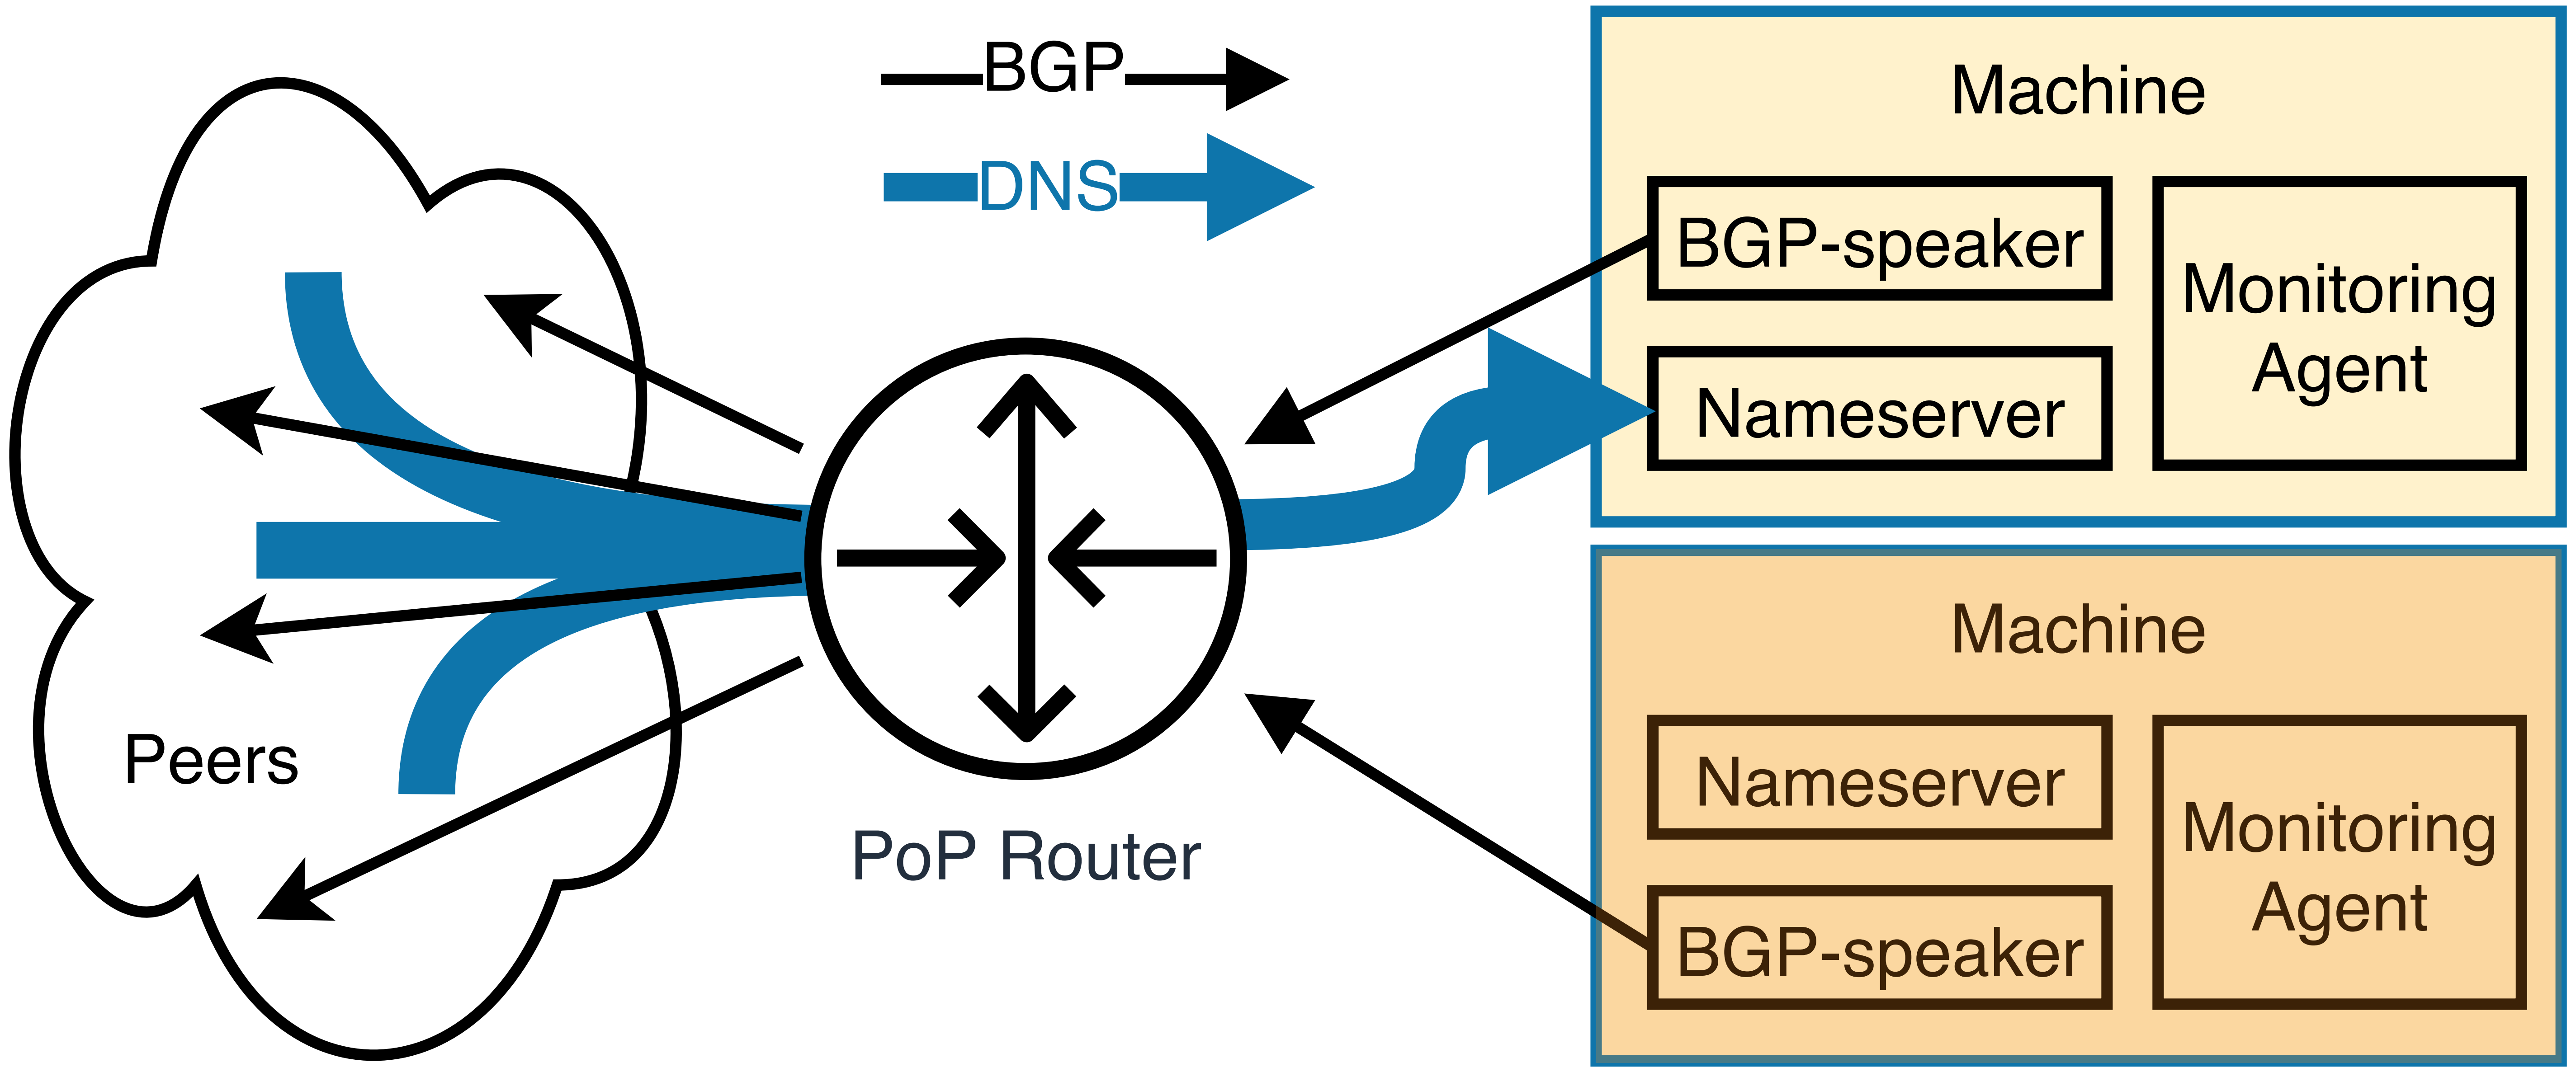

Input-induced Crash Mitigation

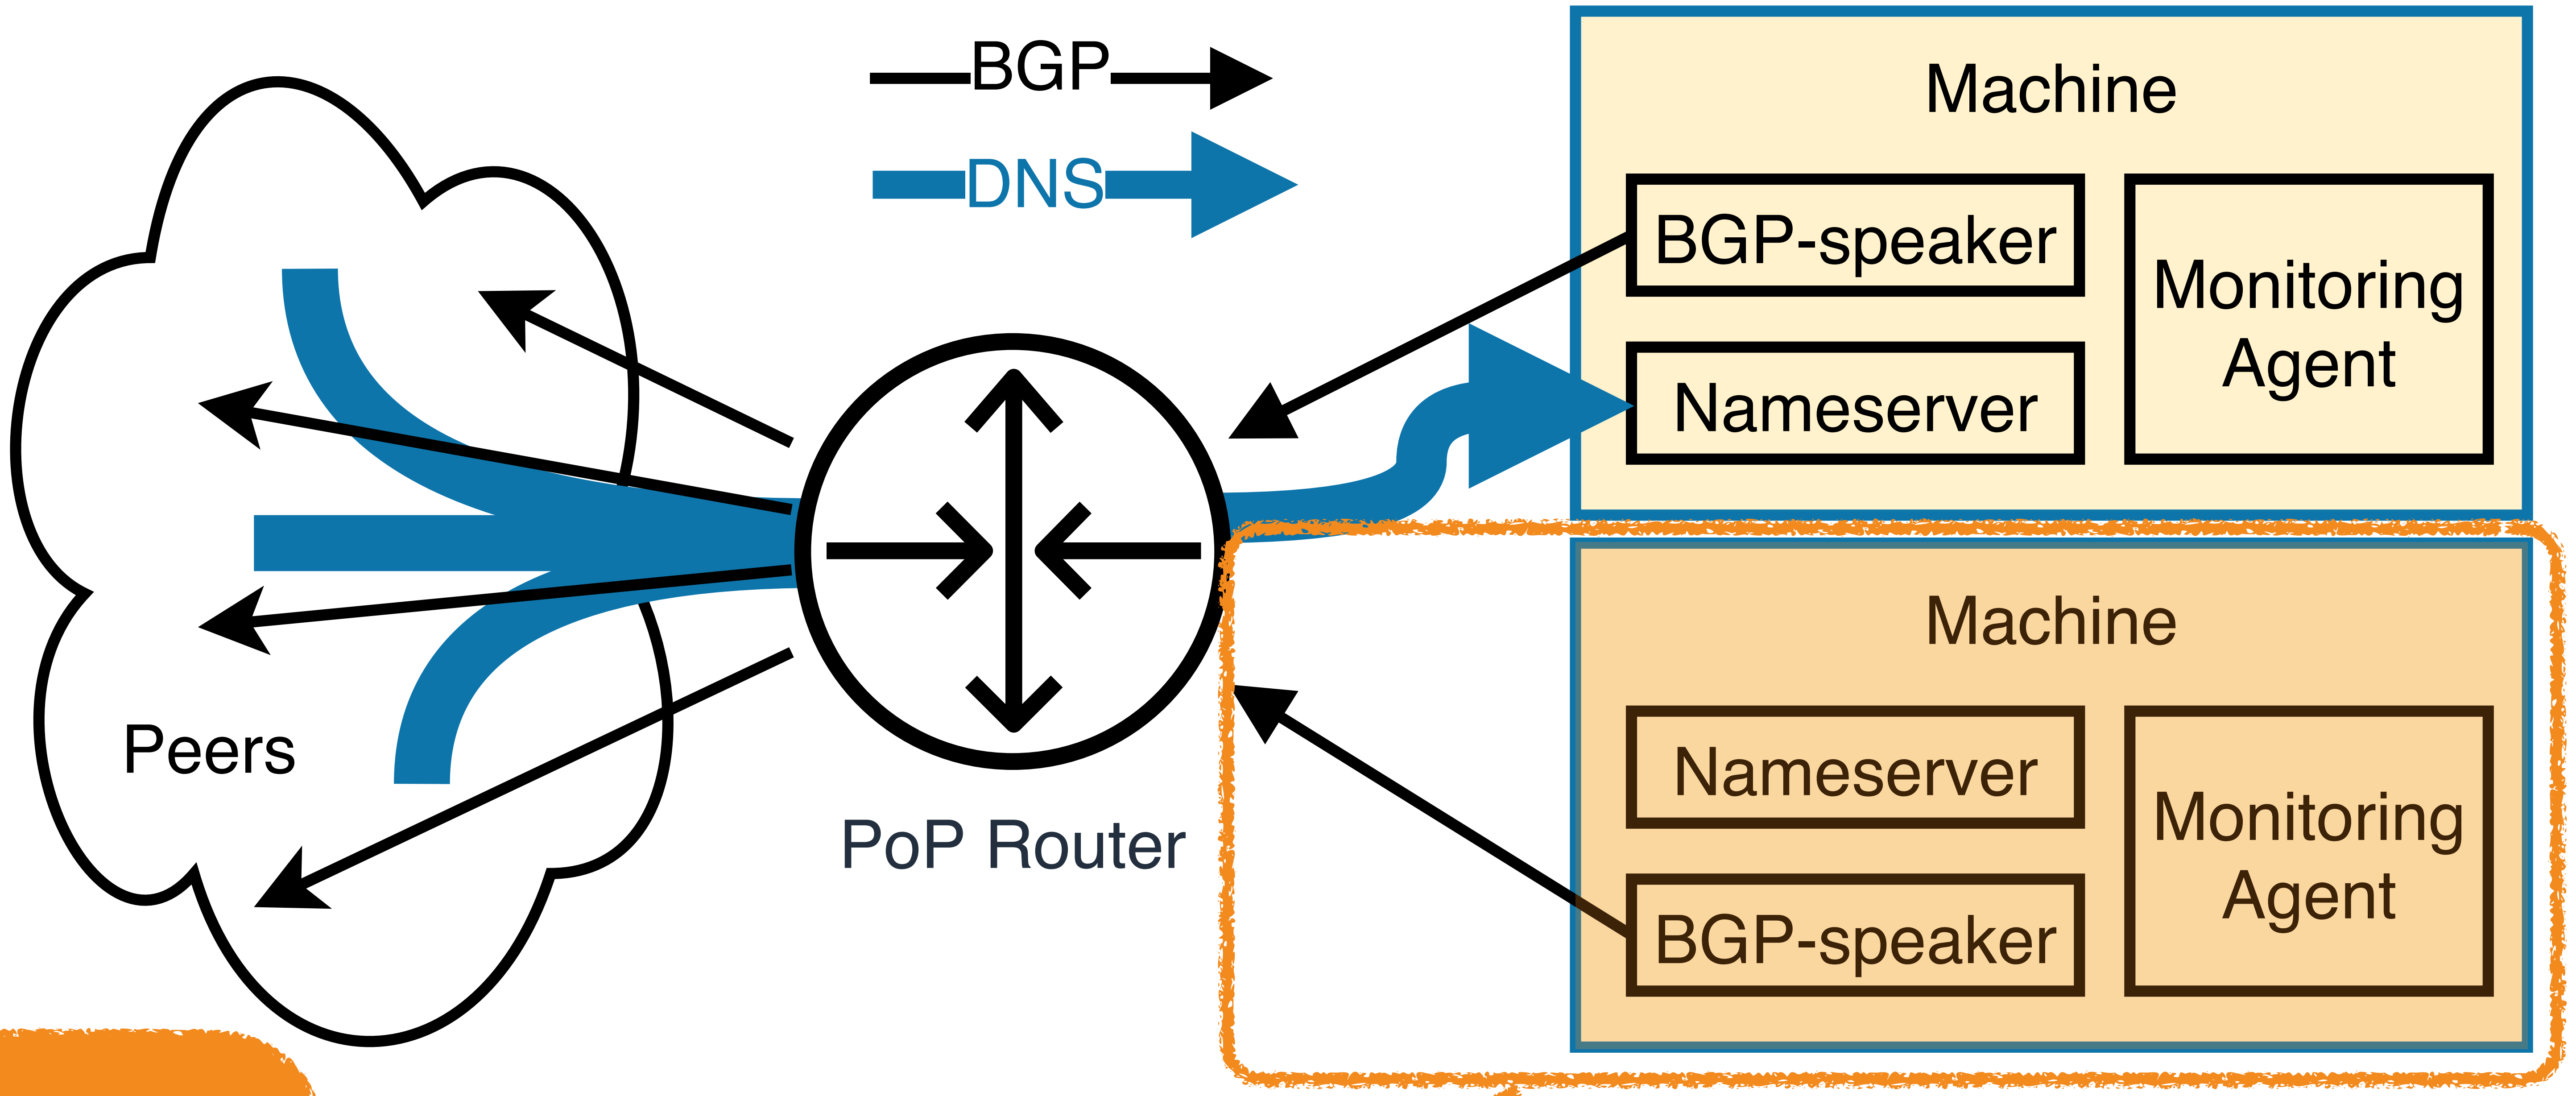

input-delayed nameservers receive input with artificially imposed delayed

Input-induced Crash Mitigation

1. nameserver crashes

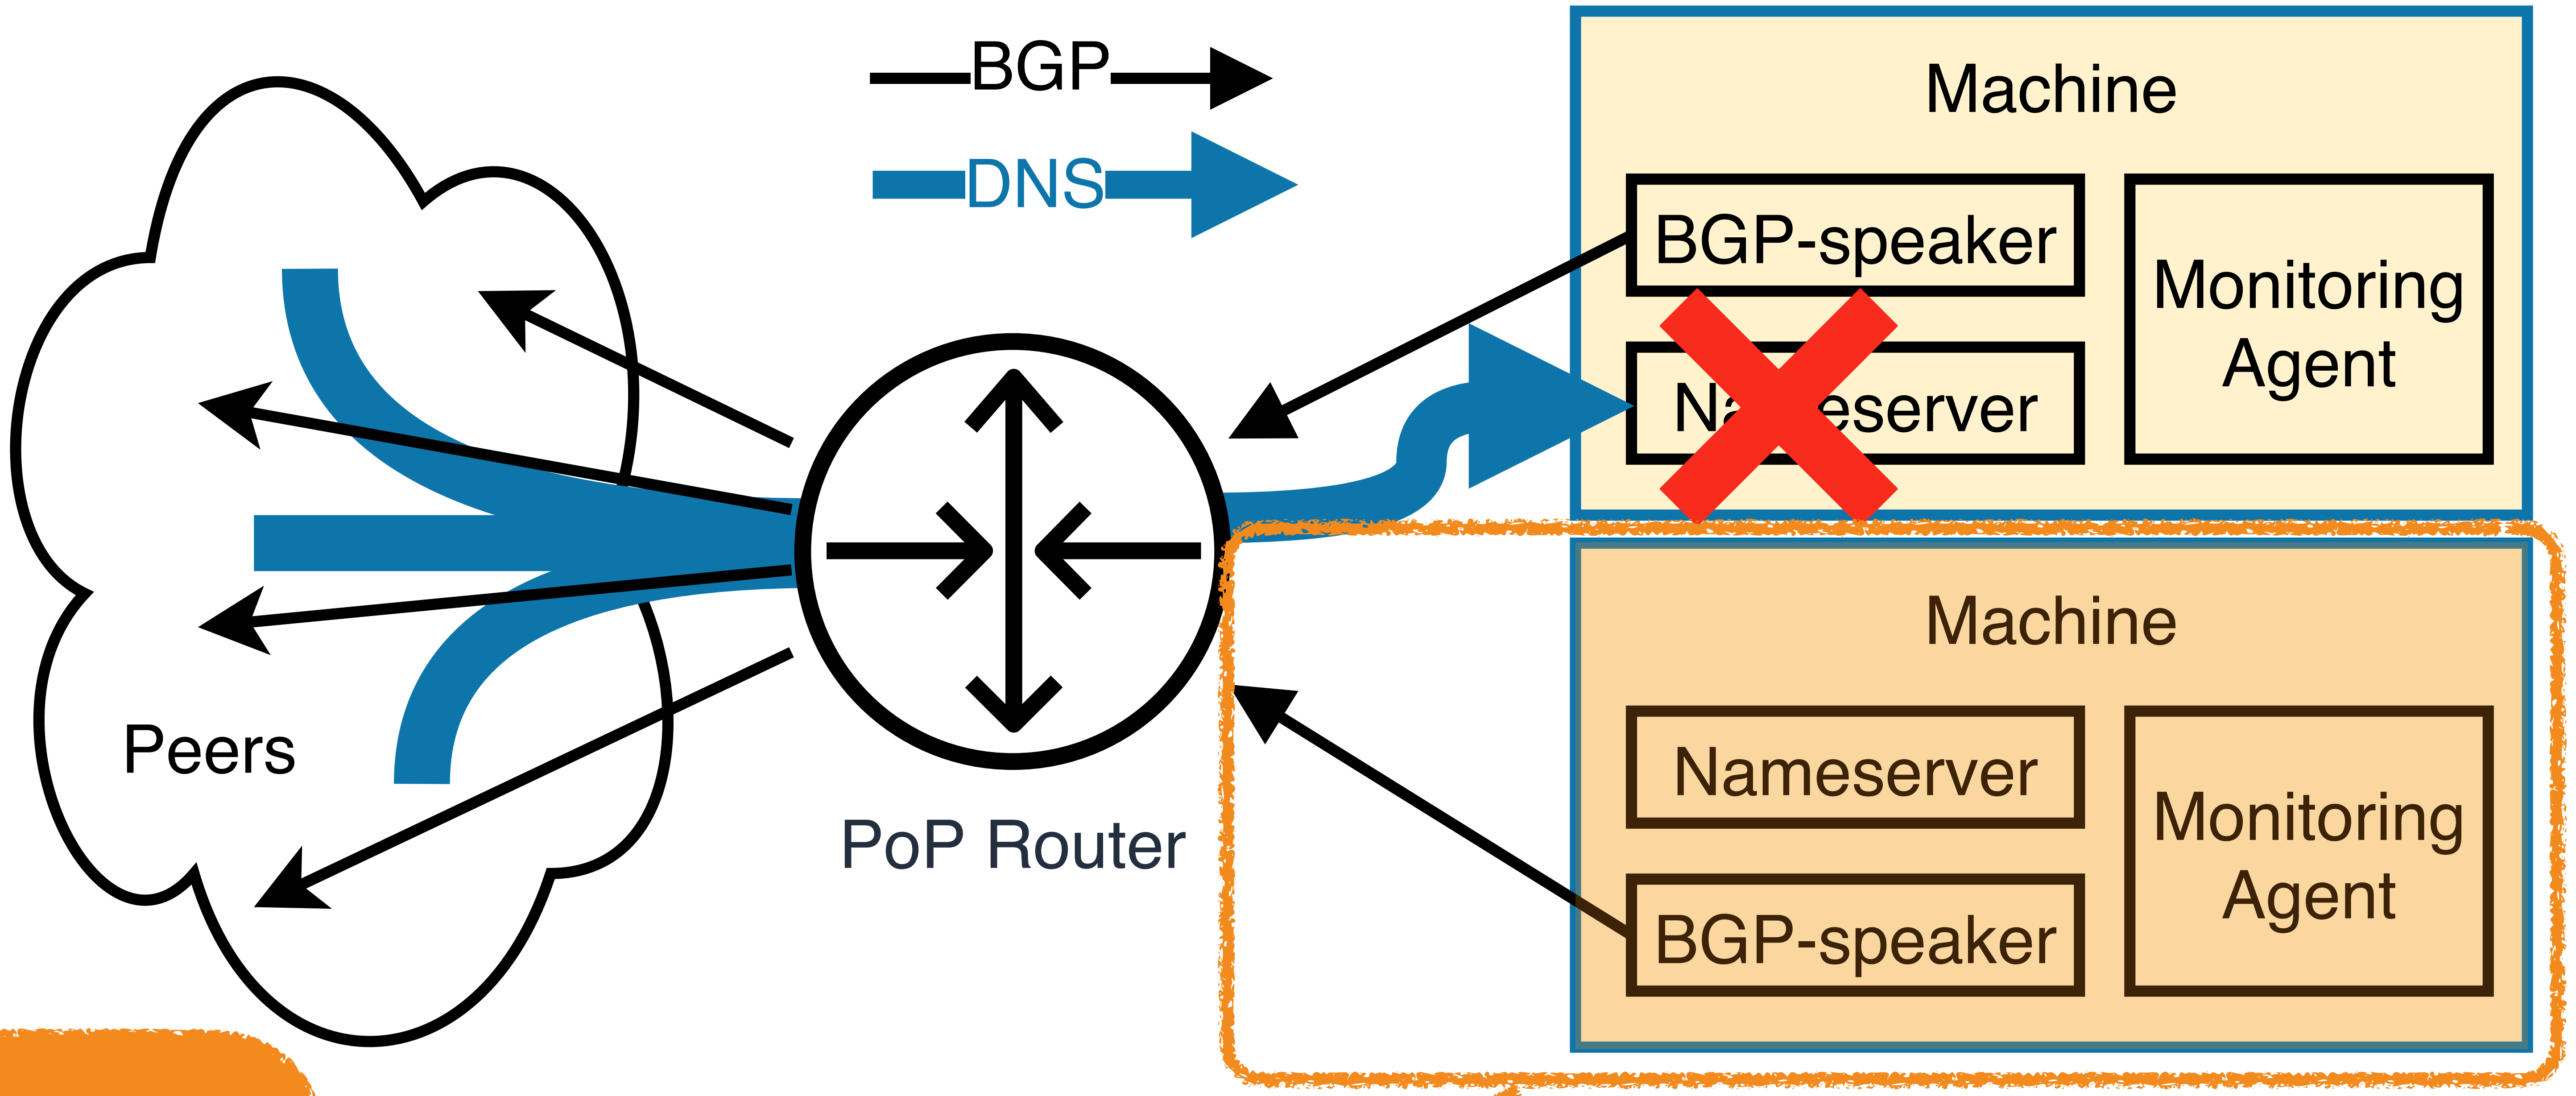

input-delayed nameservers receive input with artificially imposed delayed

Input-induced Crash Mitigation

1. nameserver crashes
2. monitoring agent withdraws BGP advertisement

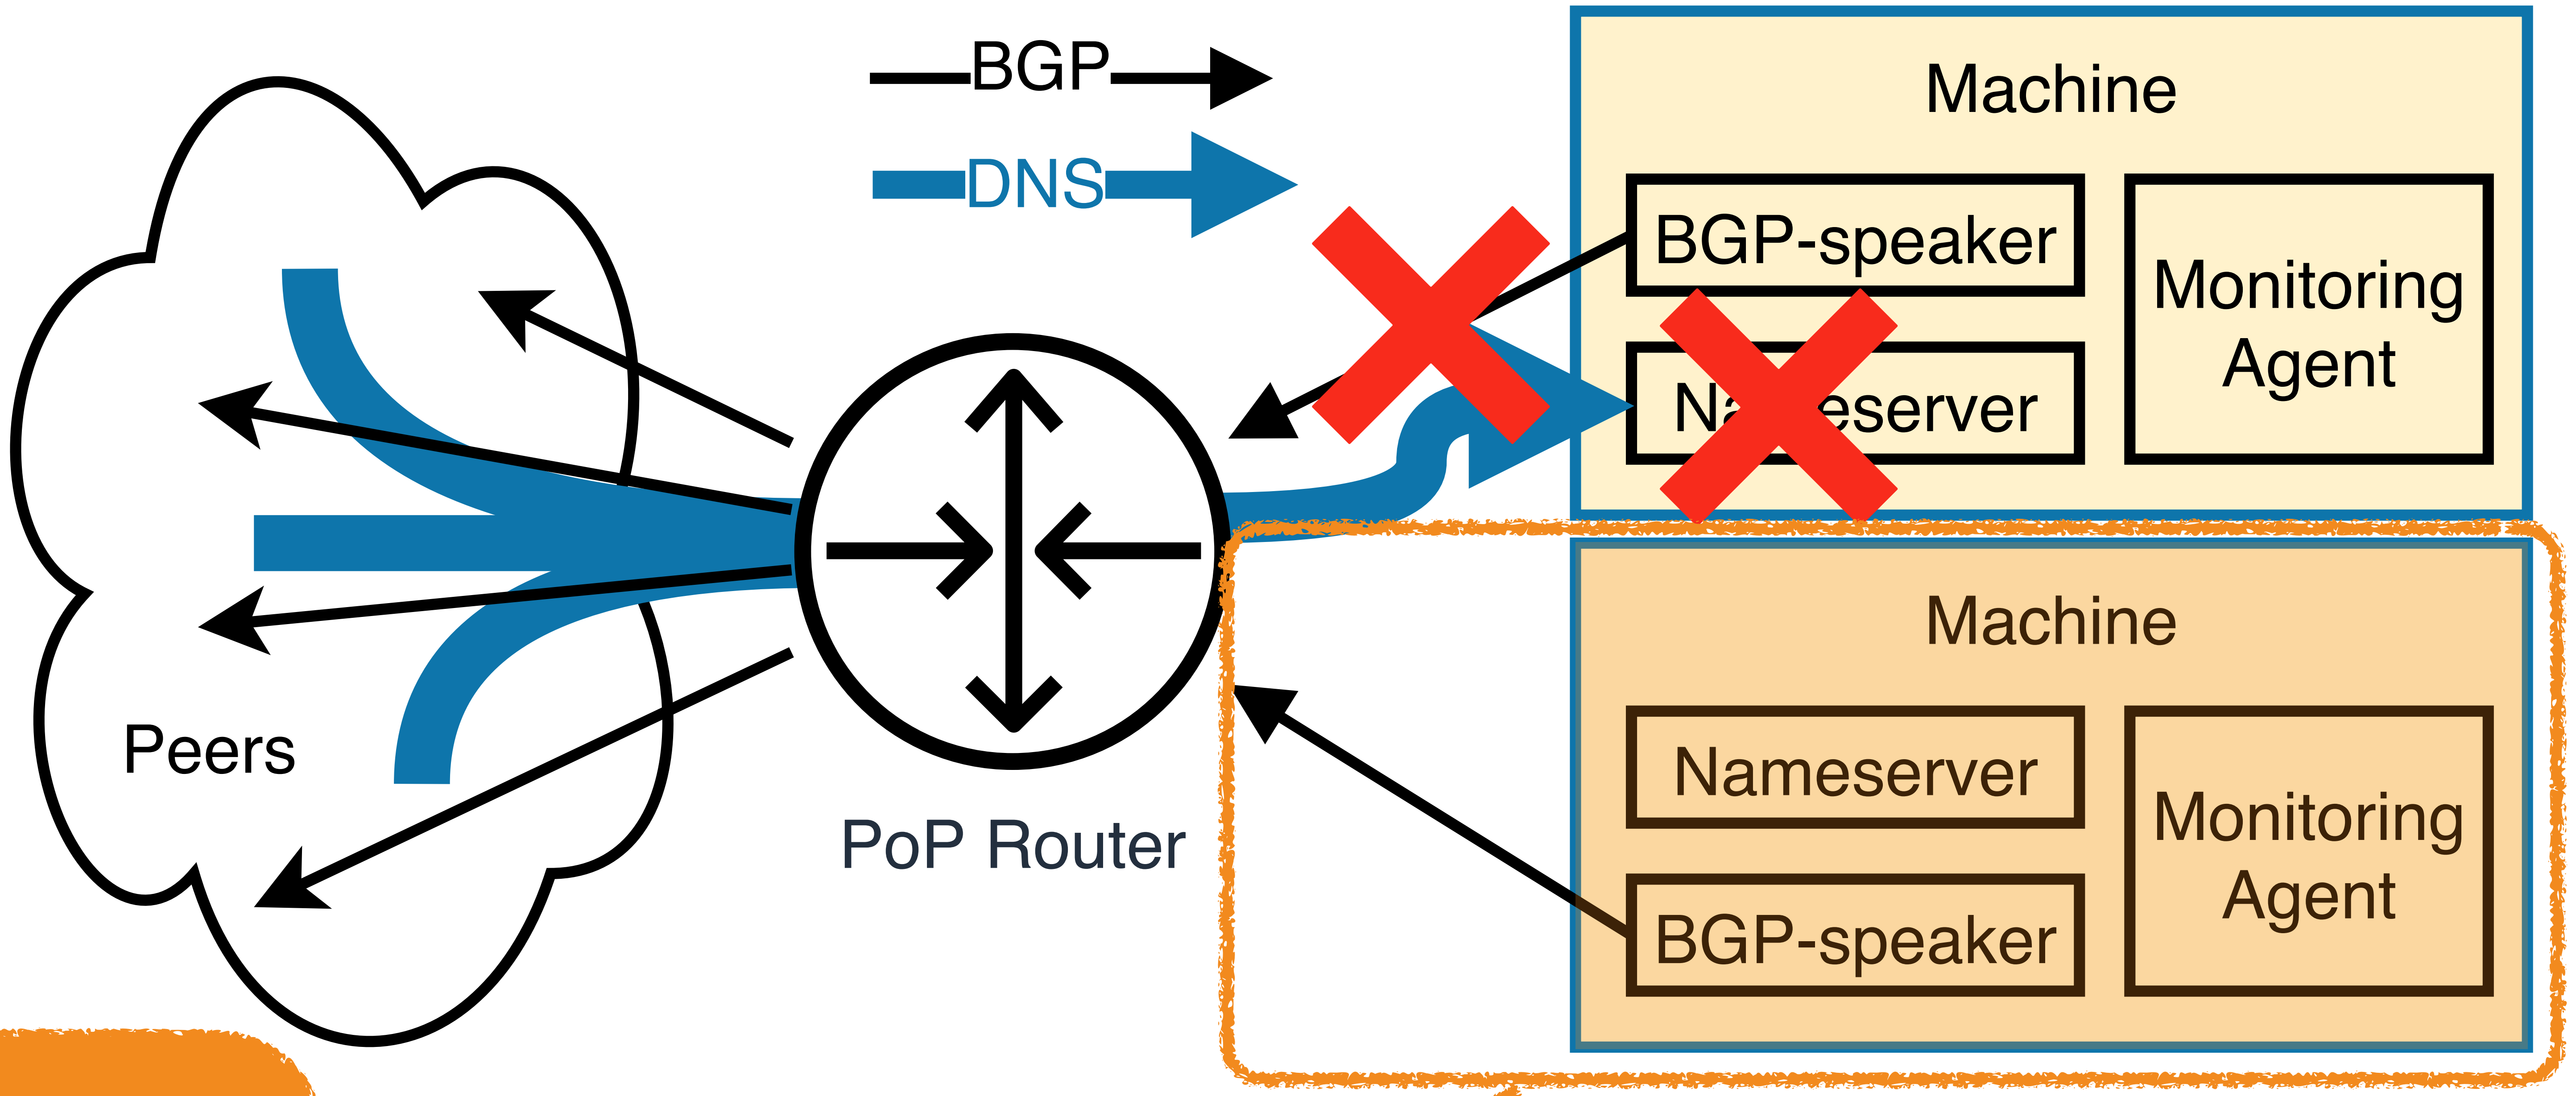

input-delayed nameservers receive input with artificially imposed delayed

Input-induced Crash Mitigation

1. nameserver crashes
2. monitoring agent withdraws BGP advertisement
3. router forwards traffic to input-delayed nameserver

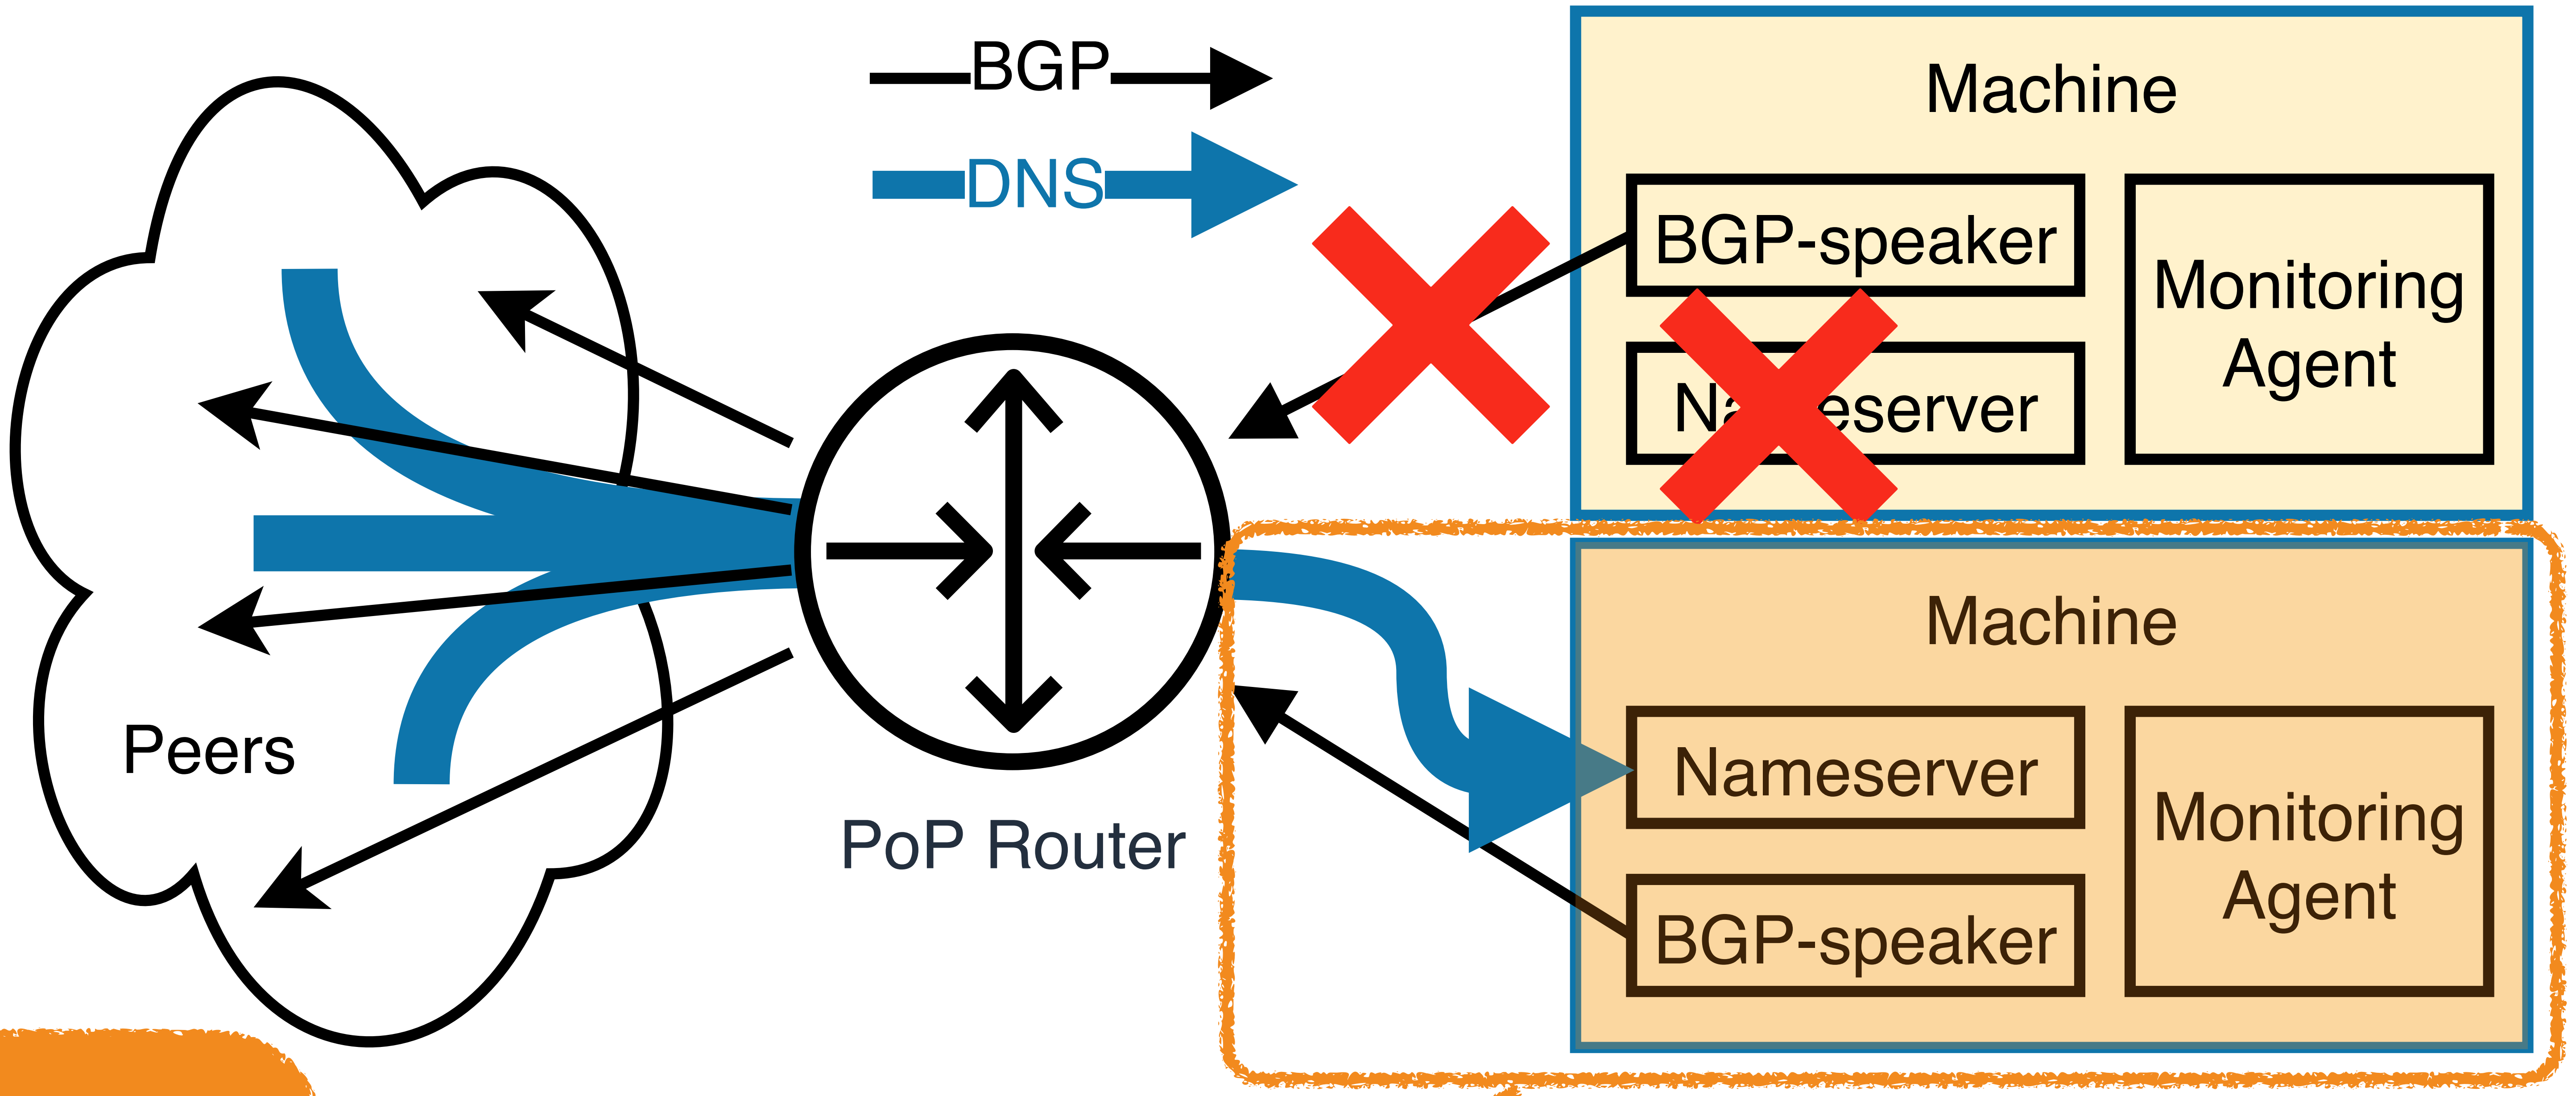

input-delayed nameservers receive input with artificially imposed delayed

Akamai DNS

Failure Resiliency

Attack Resiliency

Anycast prefixes (A B C D E F G ...) advertised from different PoPs

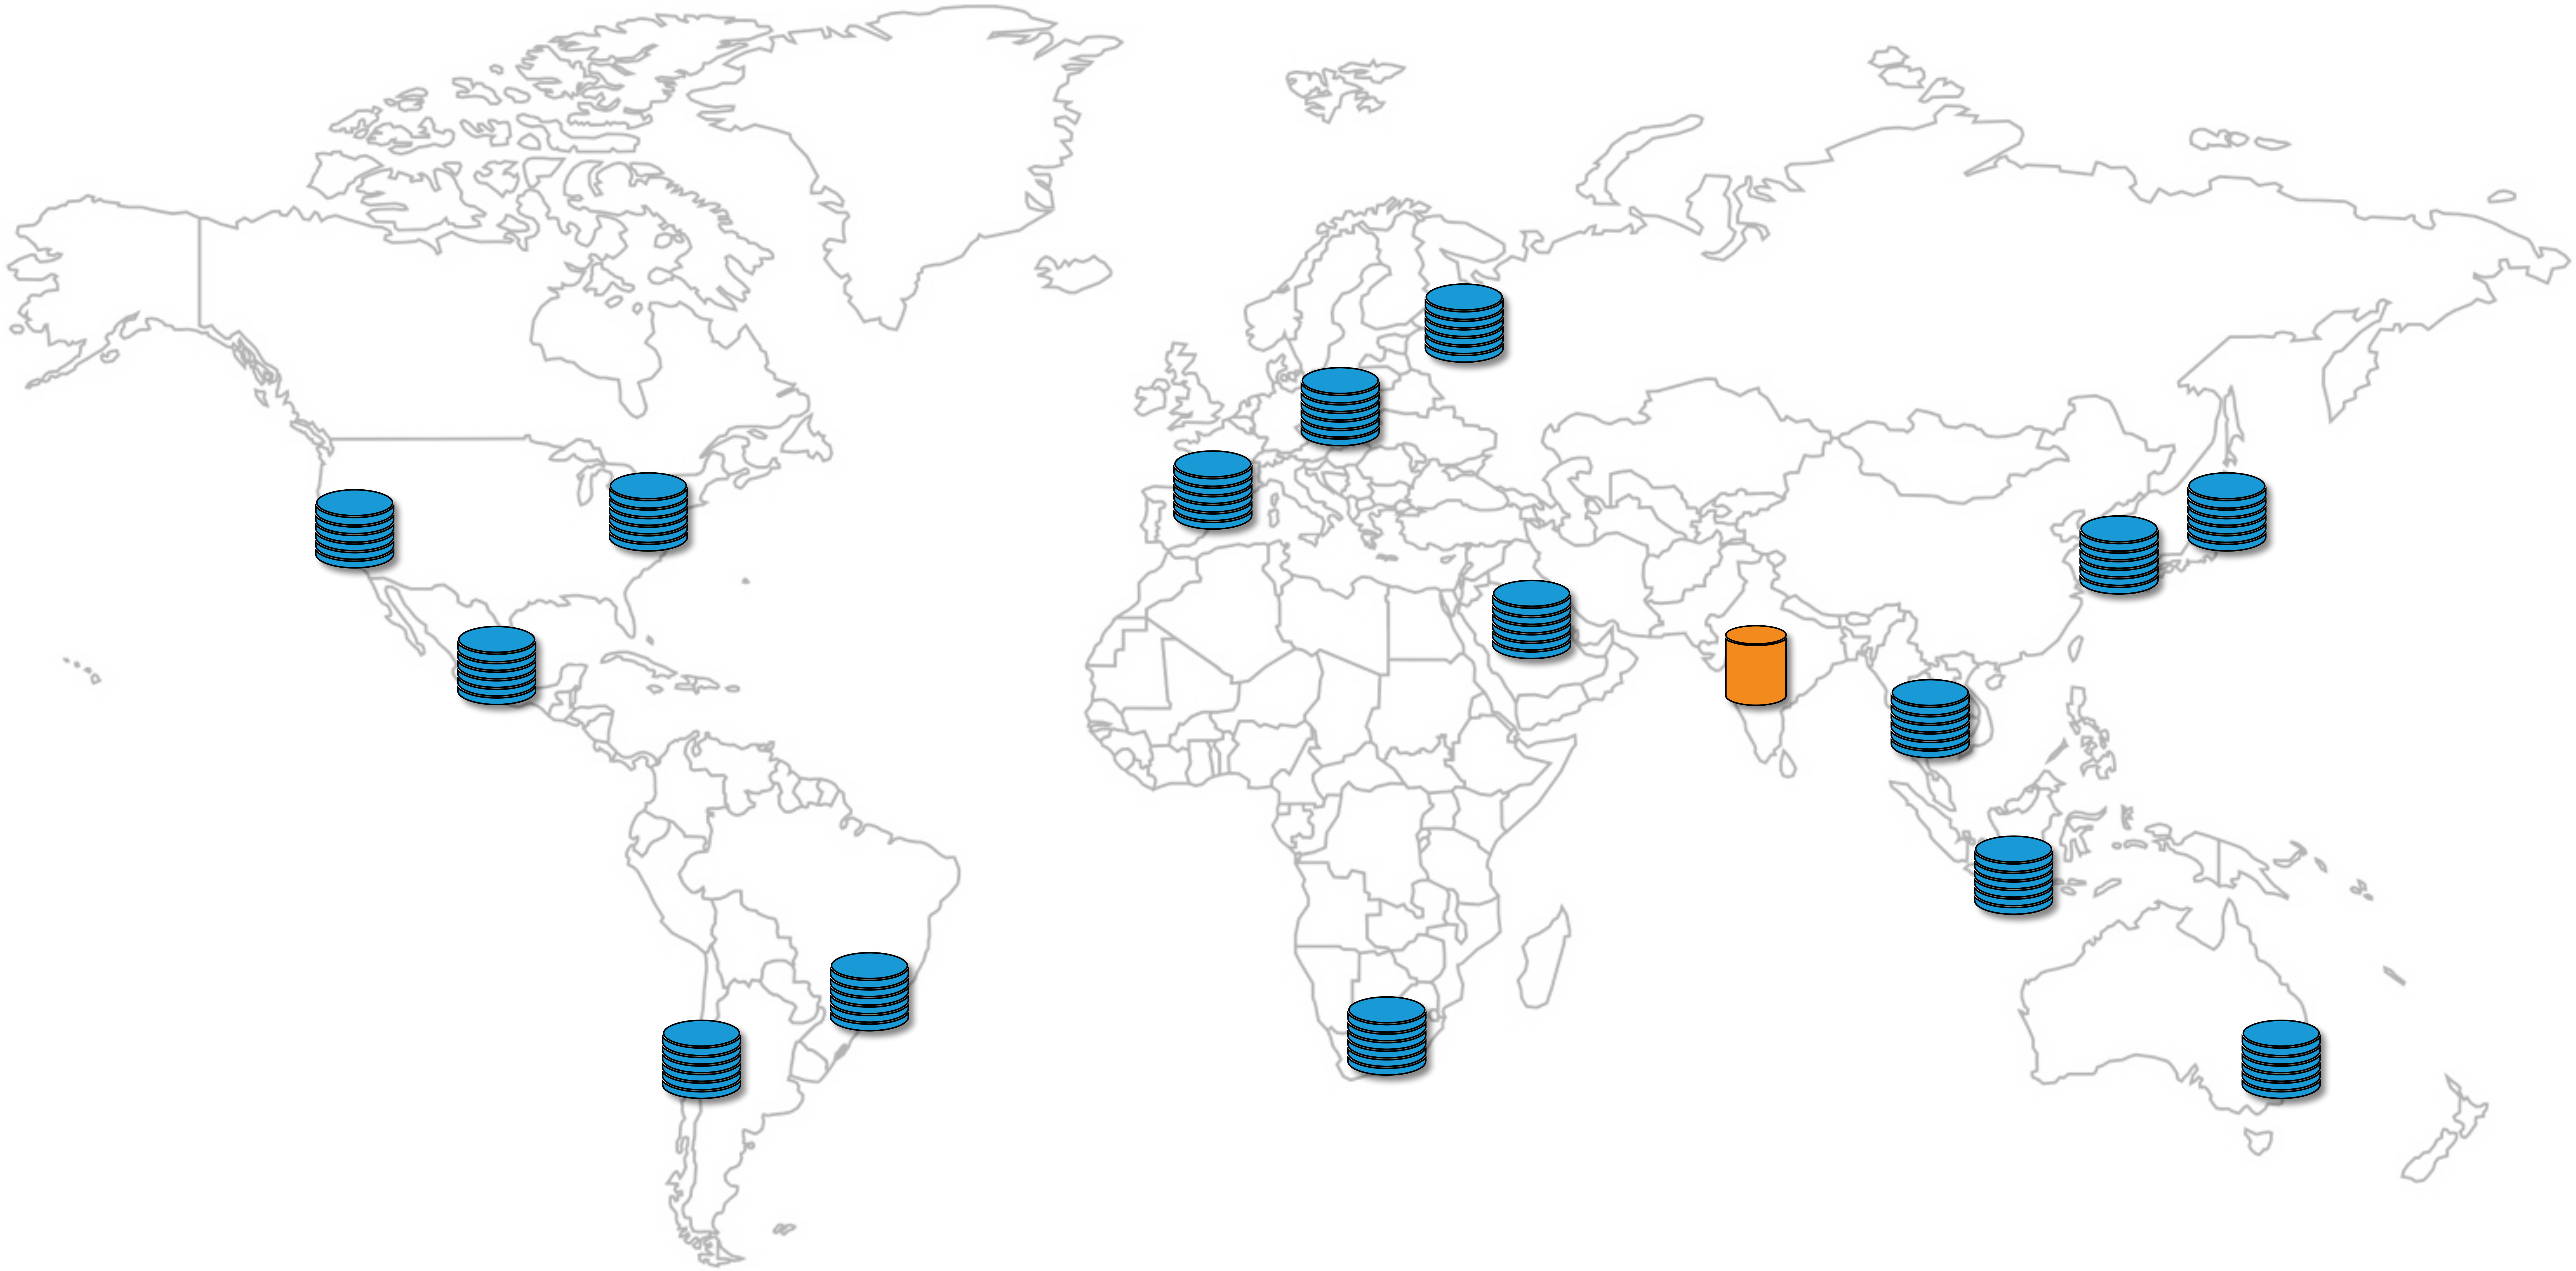

Anycast prefixes (A B C D E F G ...) advertised from different PoPs

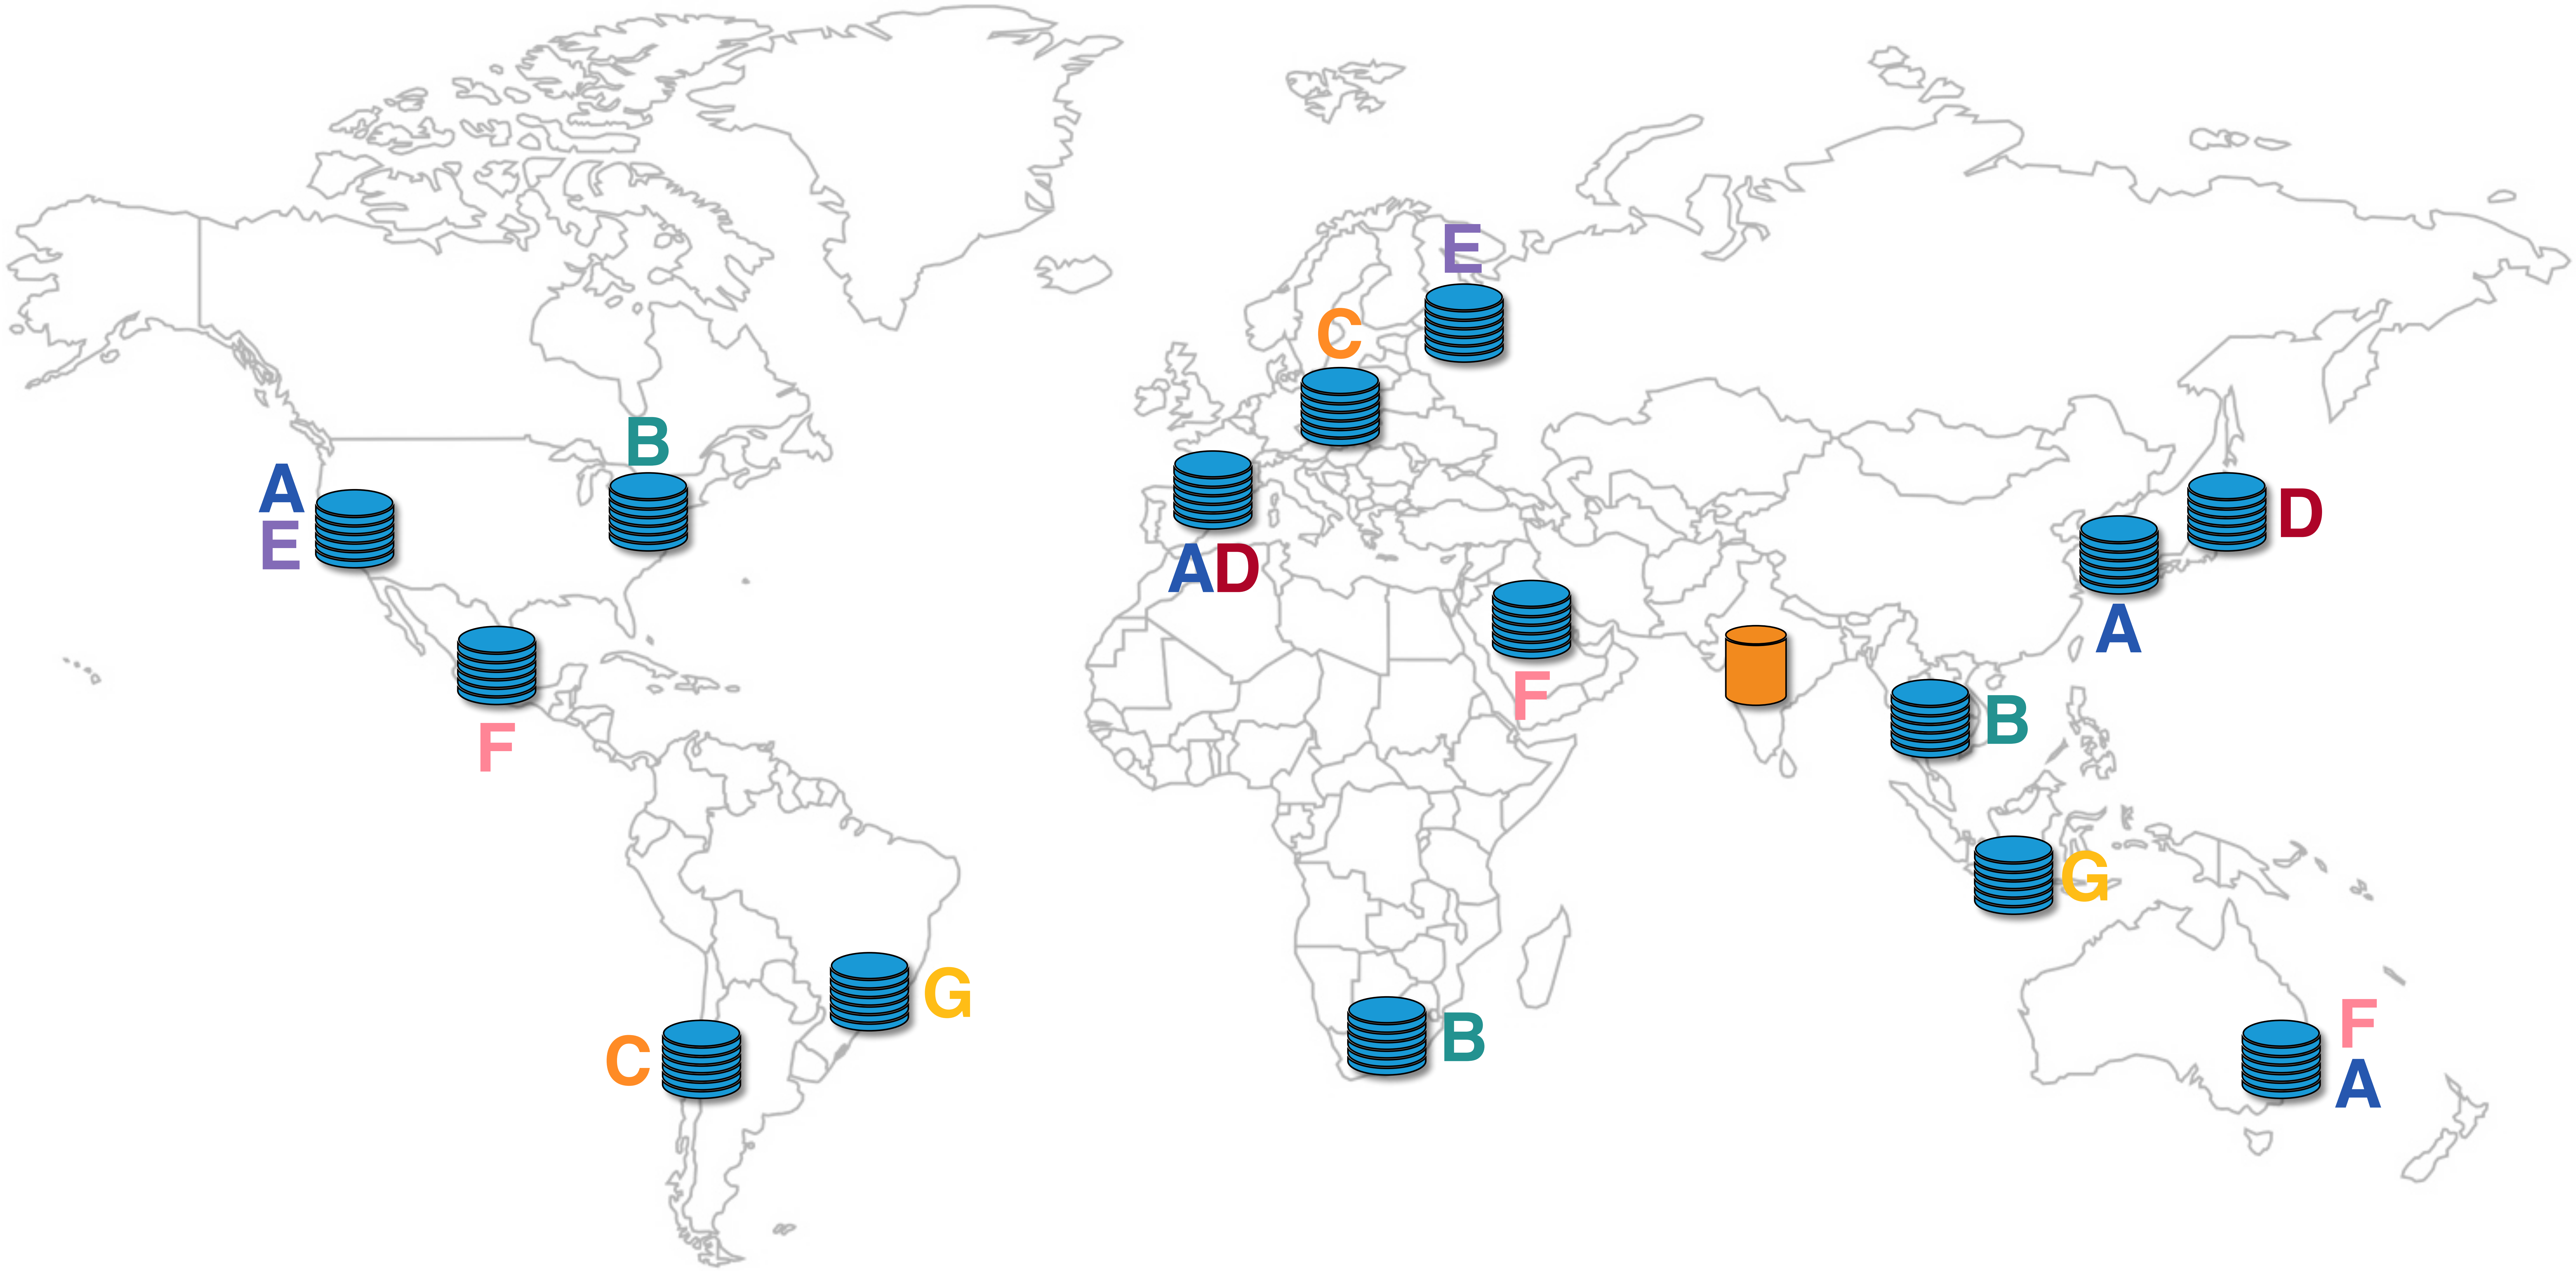

Anycast prefixes (A B C D E F G ...) advertised from different PoPs

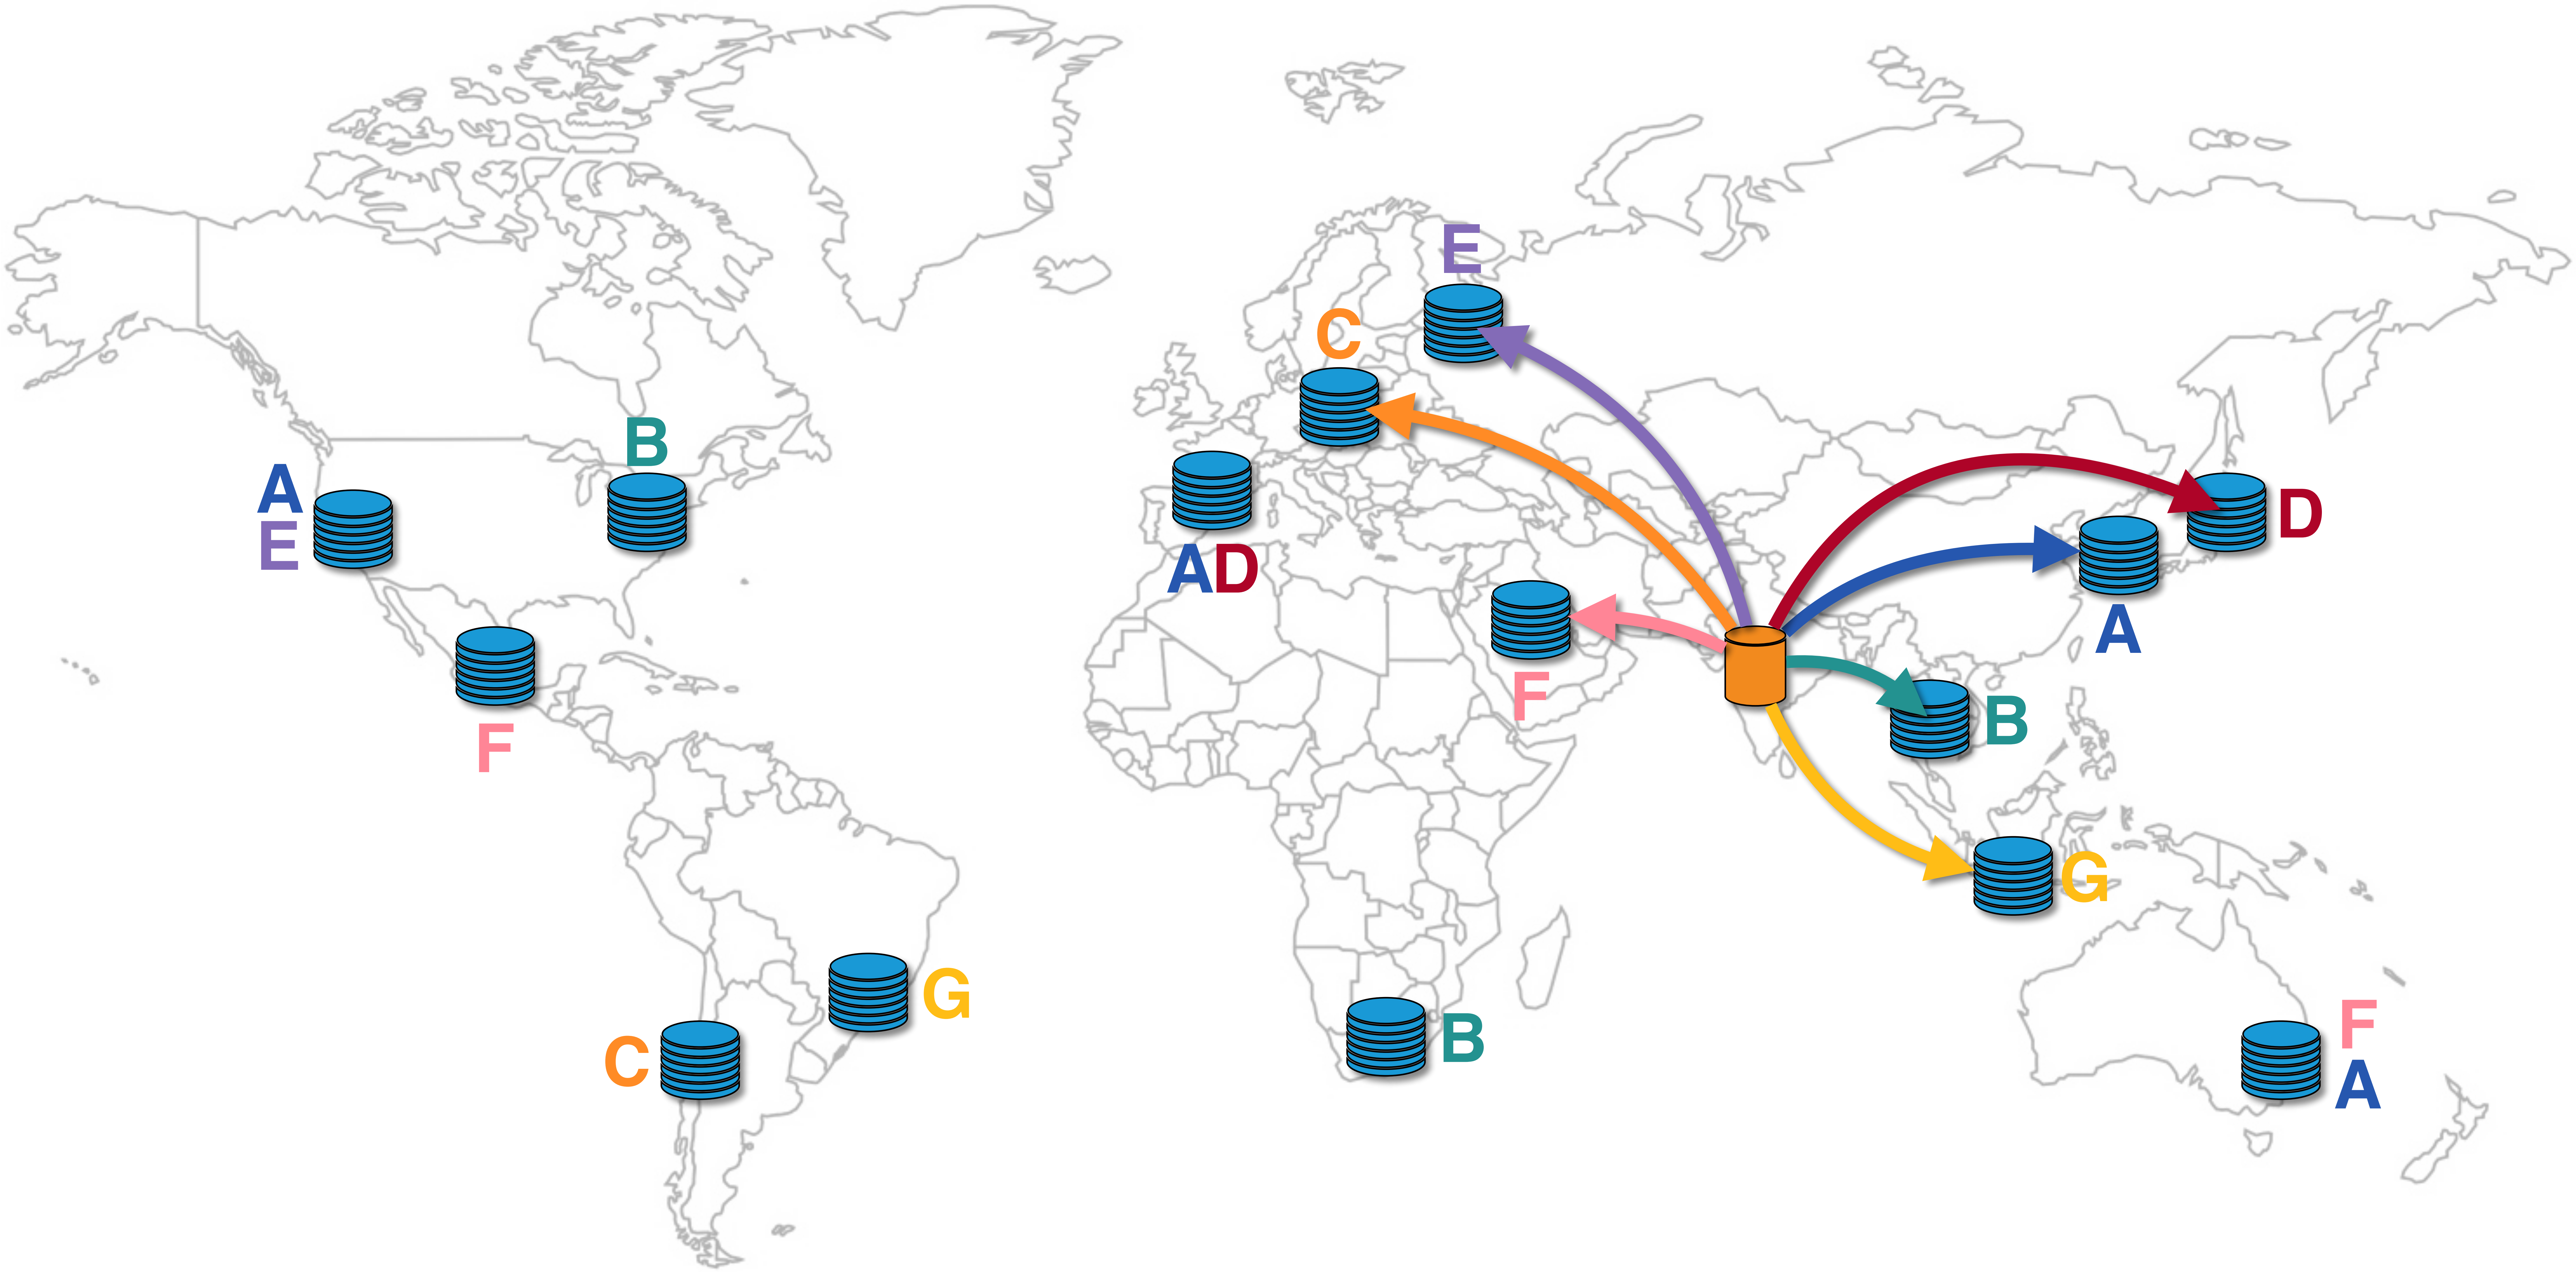

Anycast prefixes (A B C D E F G ...) advertised from different PoPs

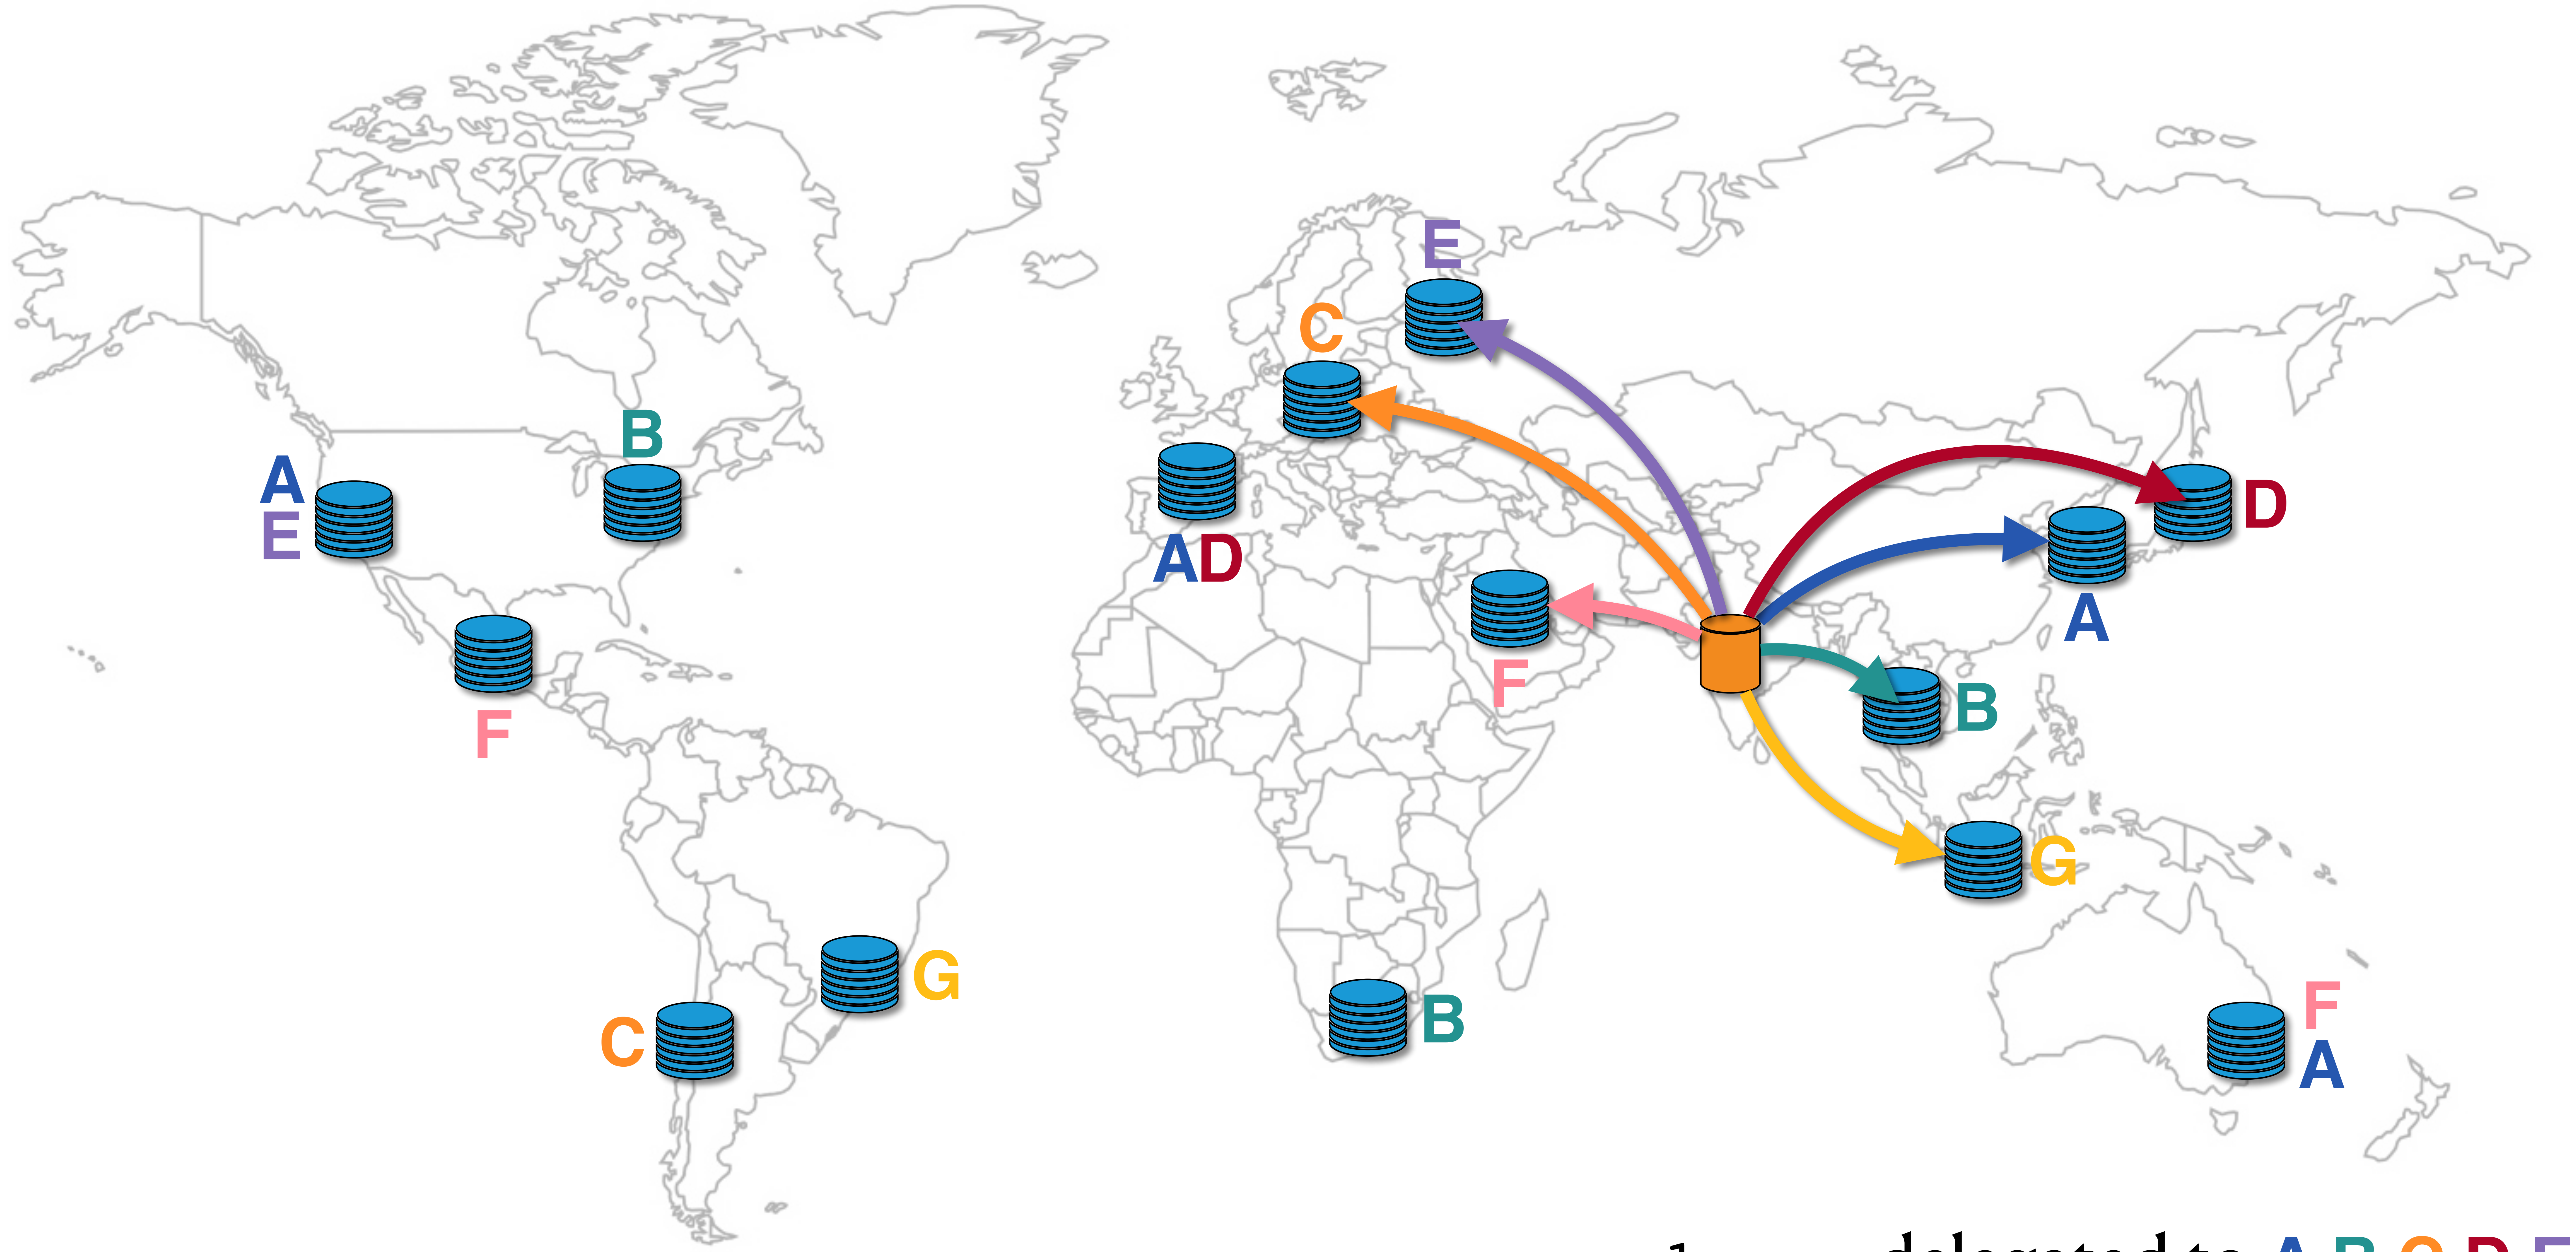

ex1.com delegated to A B C D E F
ex2.com delegated to A B C D E G

Anycast prefixes (A B C D E F G ...) advertised from different PoPs

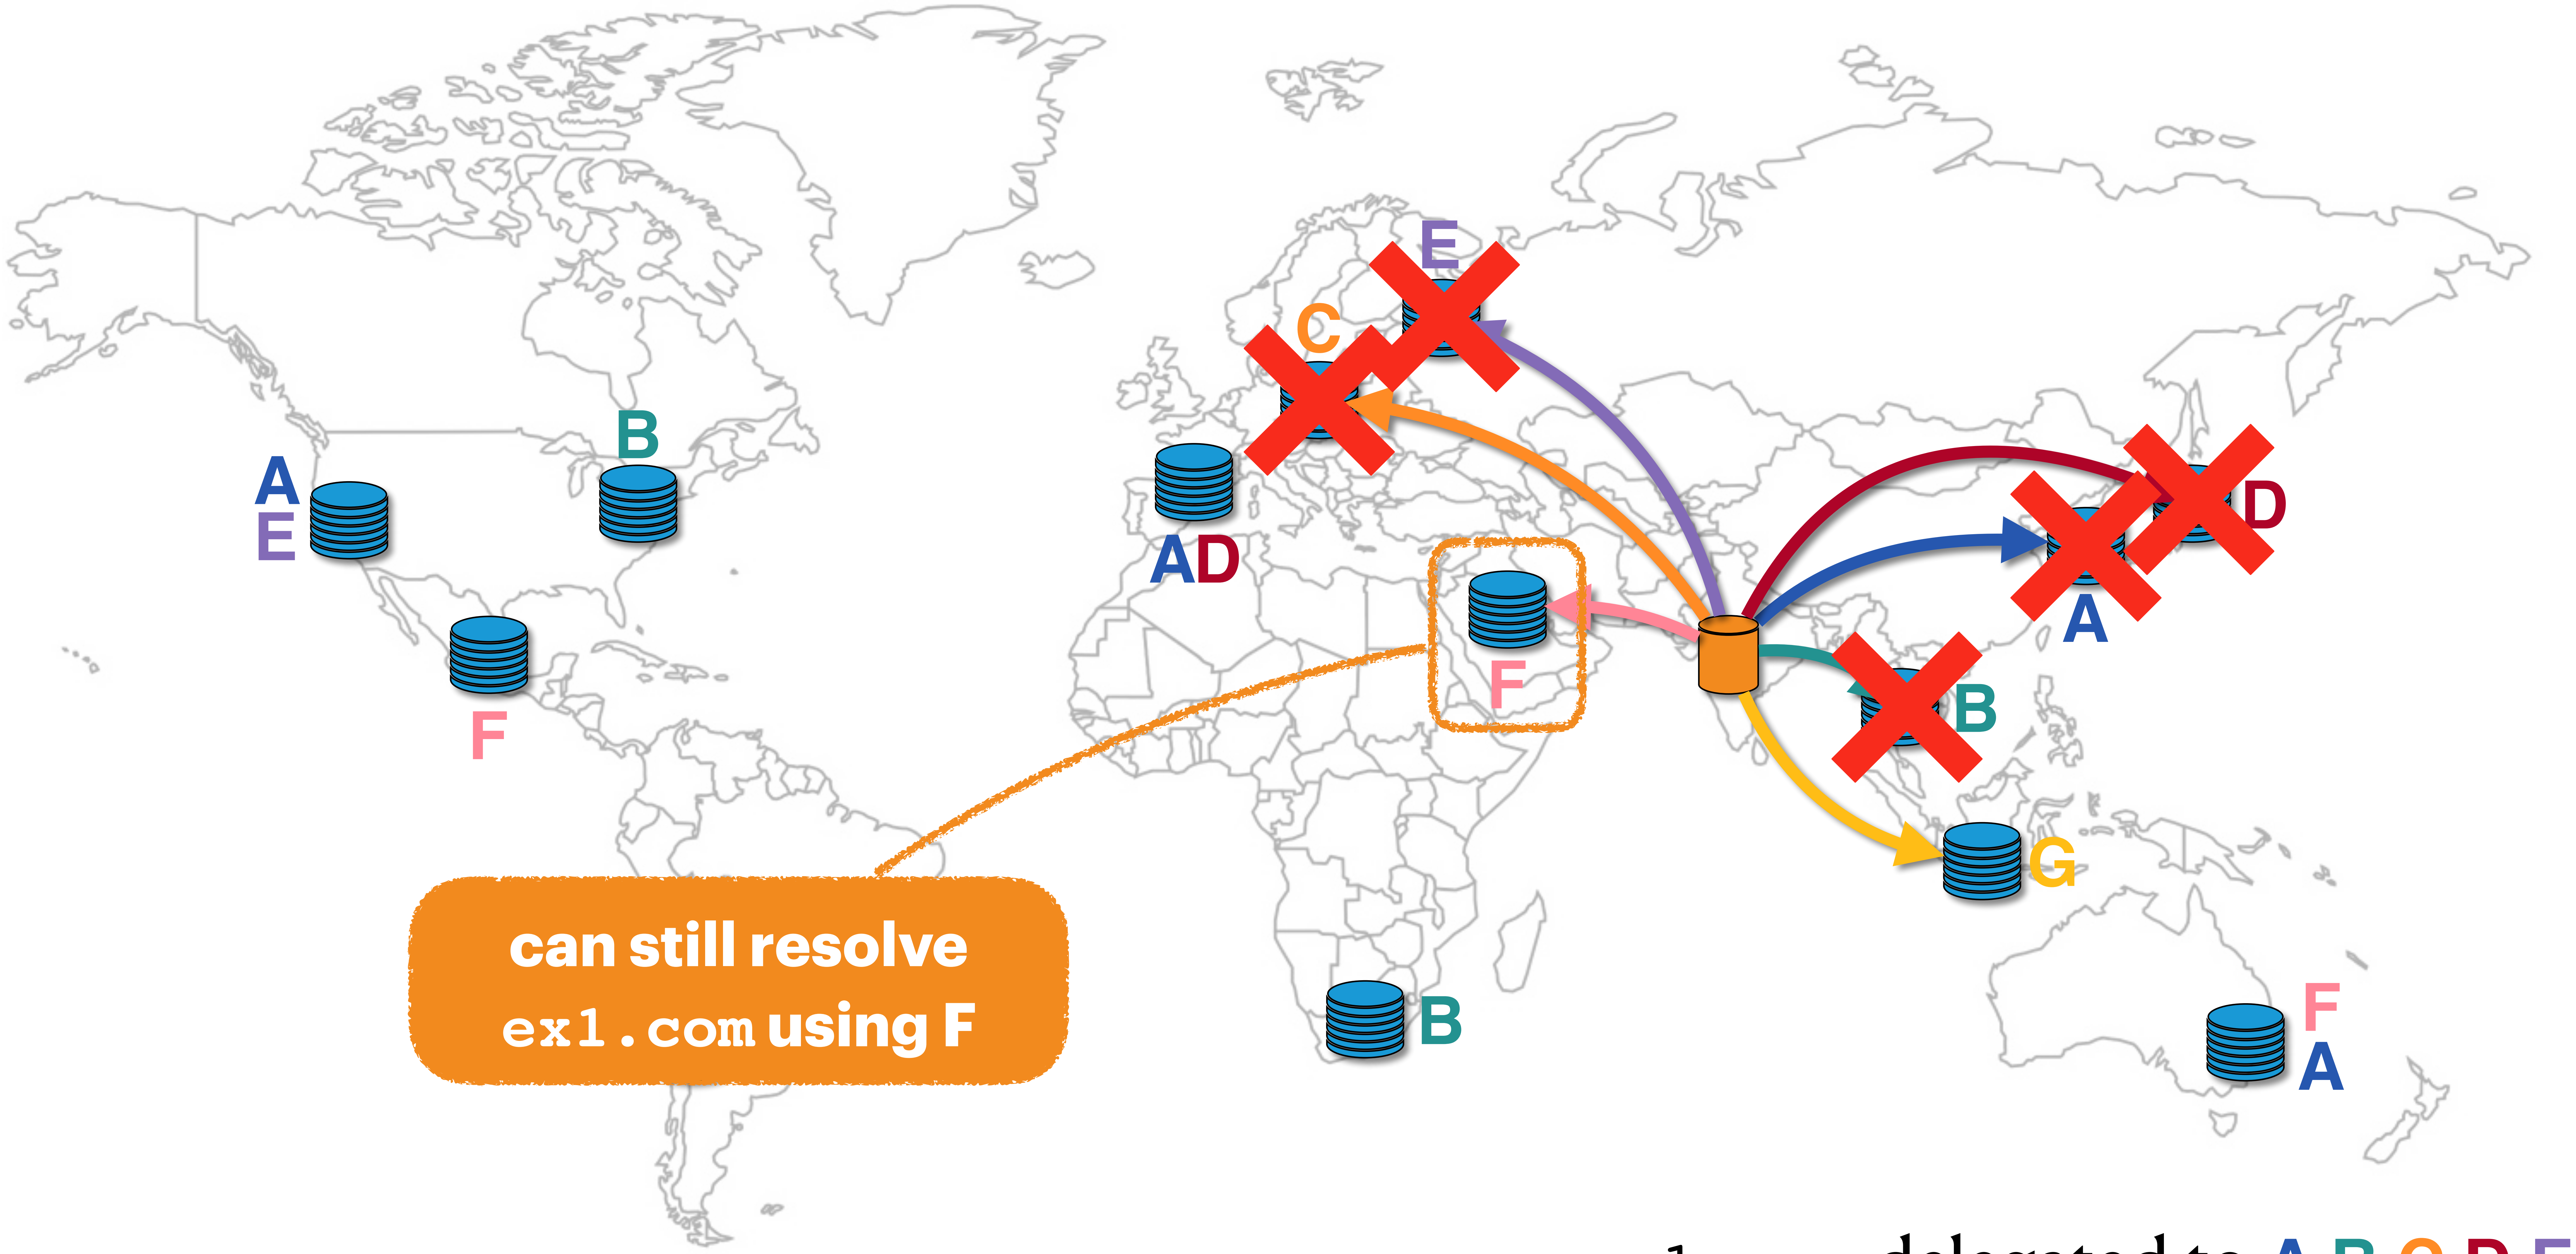

can still resolve
`ex1.com` using F

`ex1.com` delegated to A B C D E F
`ex2.com` delegated to A B C D E G

Anycast prefixes (A B C D E F G ...) advertised from different PoPs

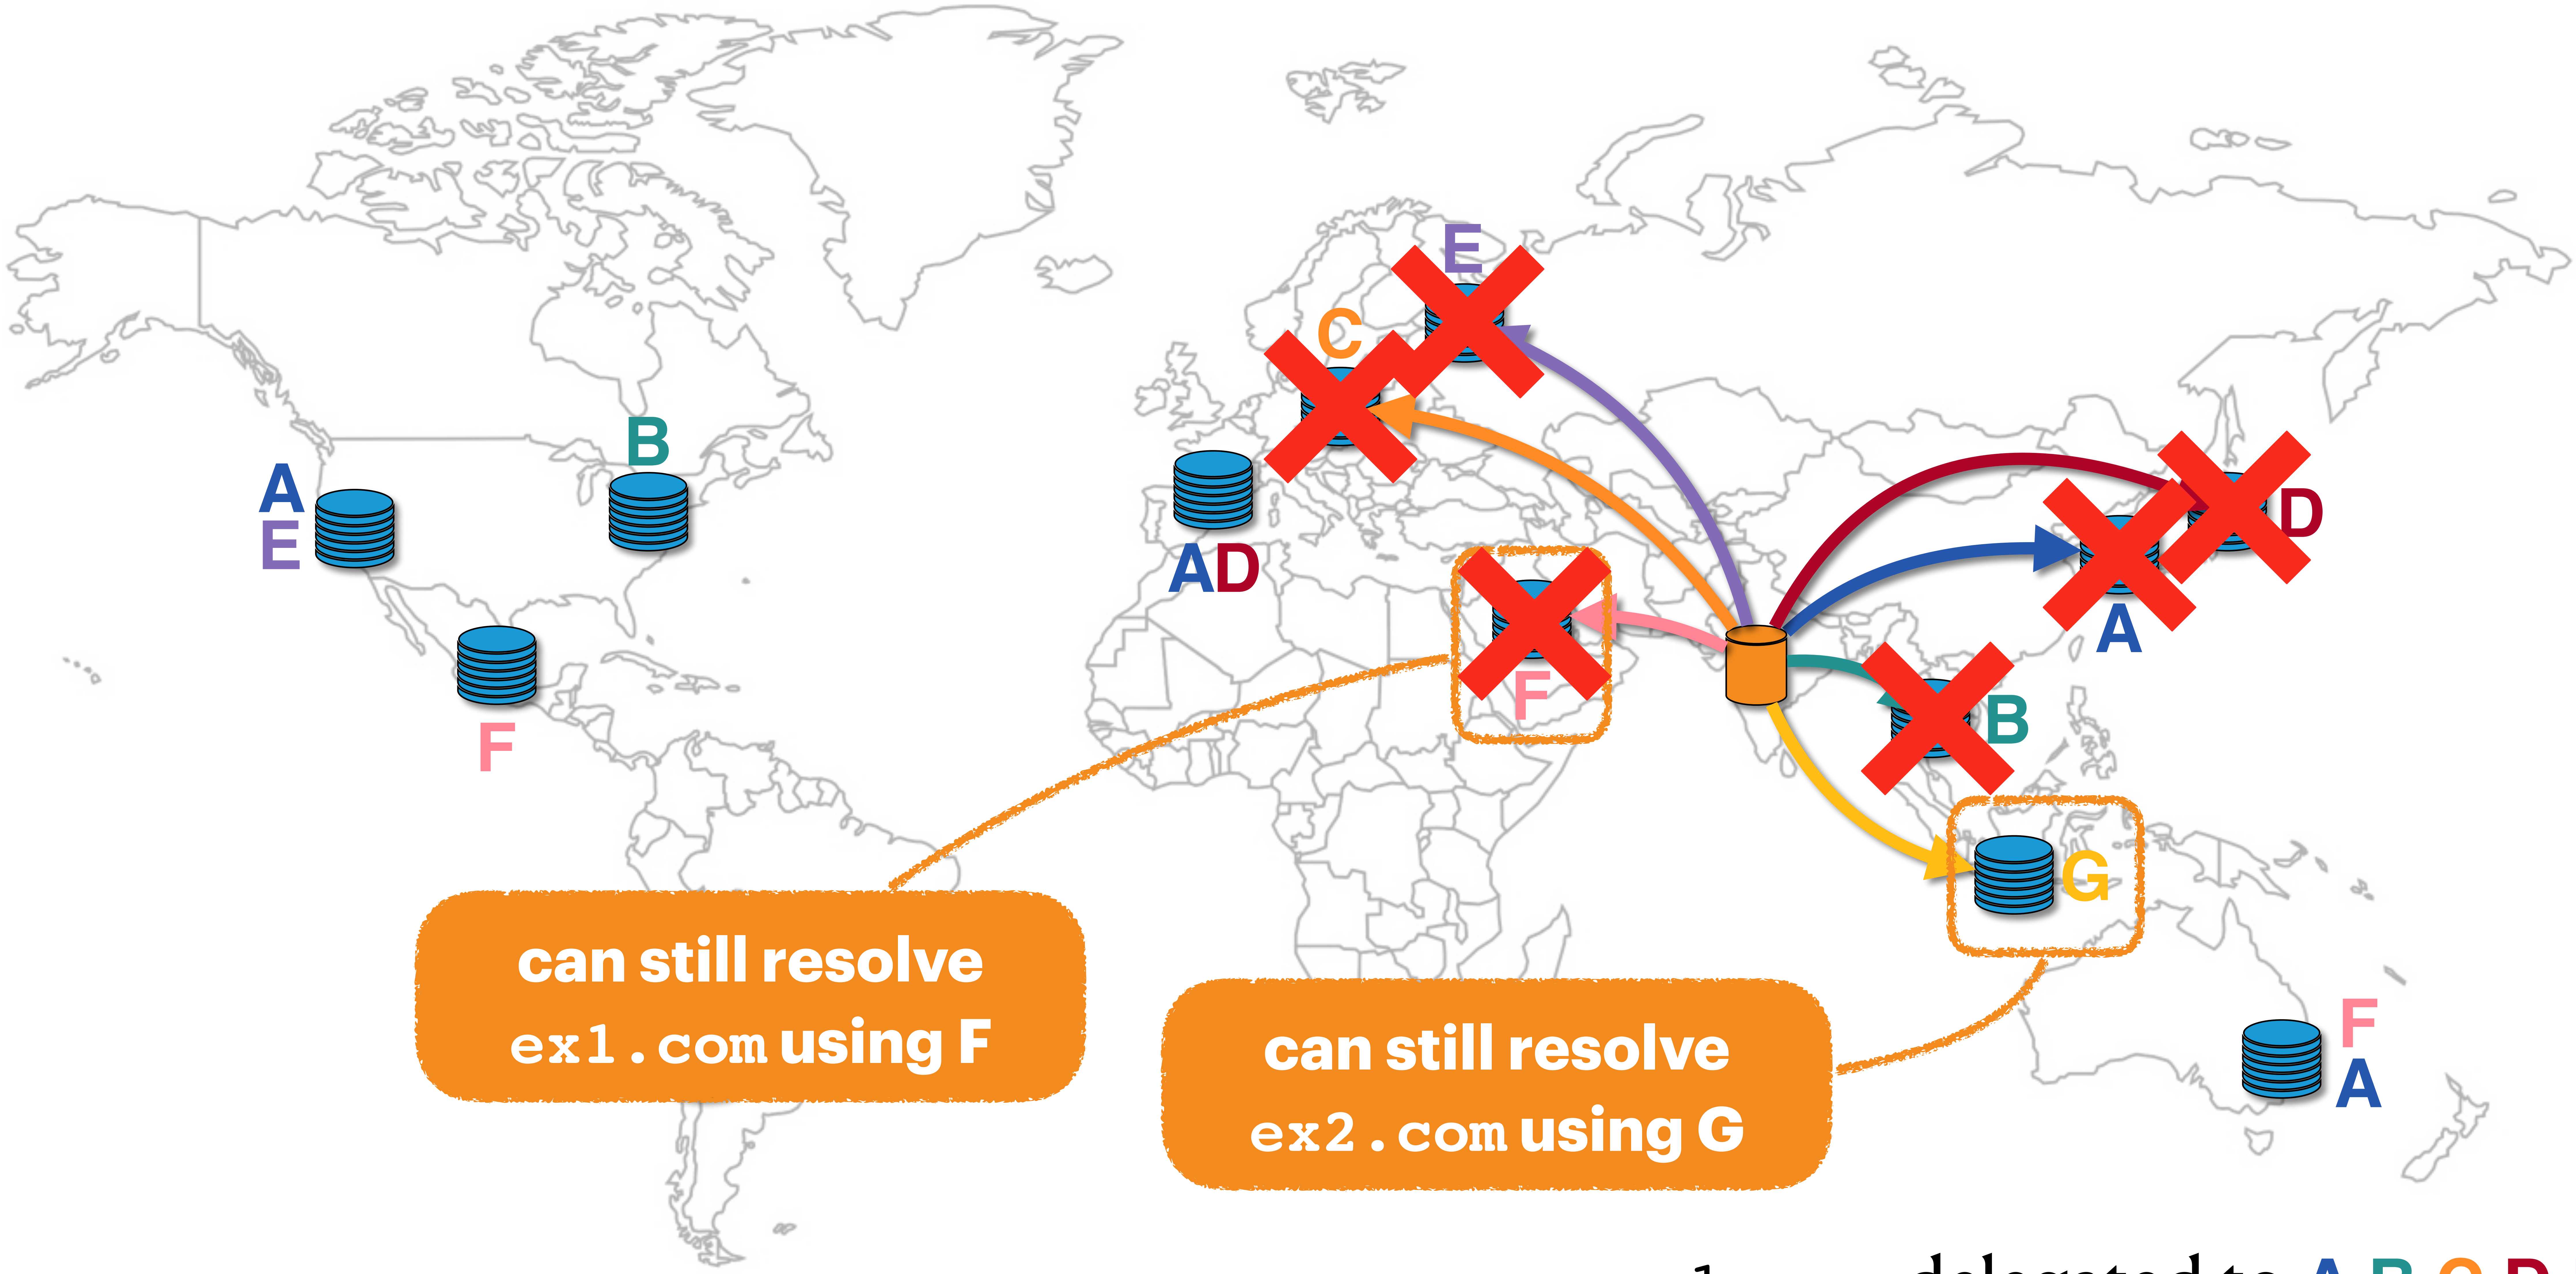

can still resolve
`ex1.com` using F

can still resolve
`ex2.com` using G

`ex1.com` delegated to A B C D E F
`ex2.com` delegated to A B C D E G

Automated Mitigations

Authoritative nameservers prioritize answering legitimate queries over suspicious ones

Automated Mitigations

Authoritative nameservers prioritize answering legitimate queries over suspicious ones

Query Scoring

1. each query passes through multiple filters
2. each filter adds a penalty
3. query added to queue according to total penalty
4. large penalty queries dropped outright

Automated Mitigations

Authoritative nameservers prioritize answering legitimate queries over suspicious ones

Query Scoring

1. each query passes through multiple filters
2. each filter adds a penalty
3. query added to queue according to total penalty
4. large penalty queries dropped outright

Query Processing

1. queues read in order of increasing penalty
2. low penalty queries preferred
3. high penalty queries potentially discard

Testbed demonstration of filtering

Fixed rate of legitimate traffic, increasing rate of attack traffic

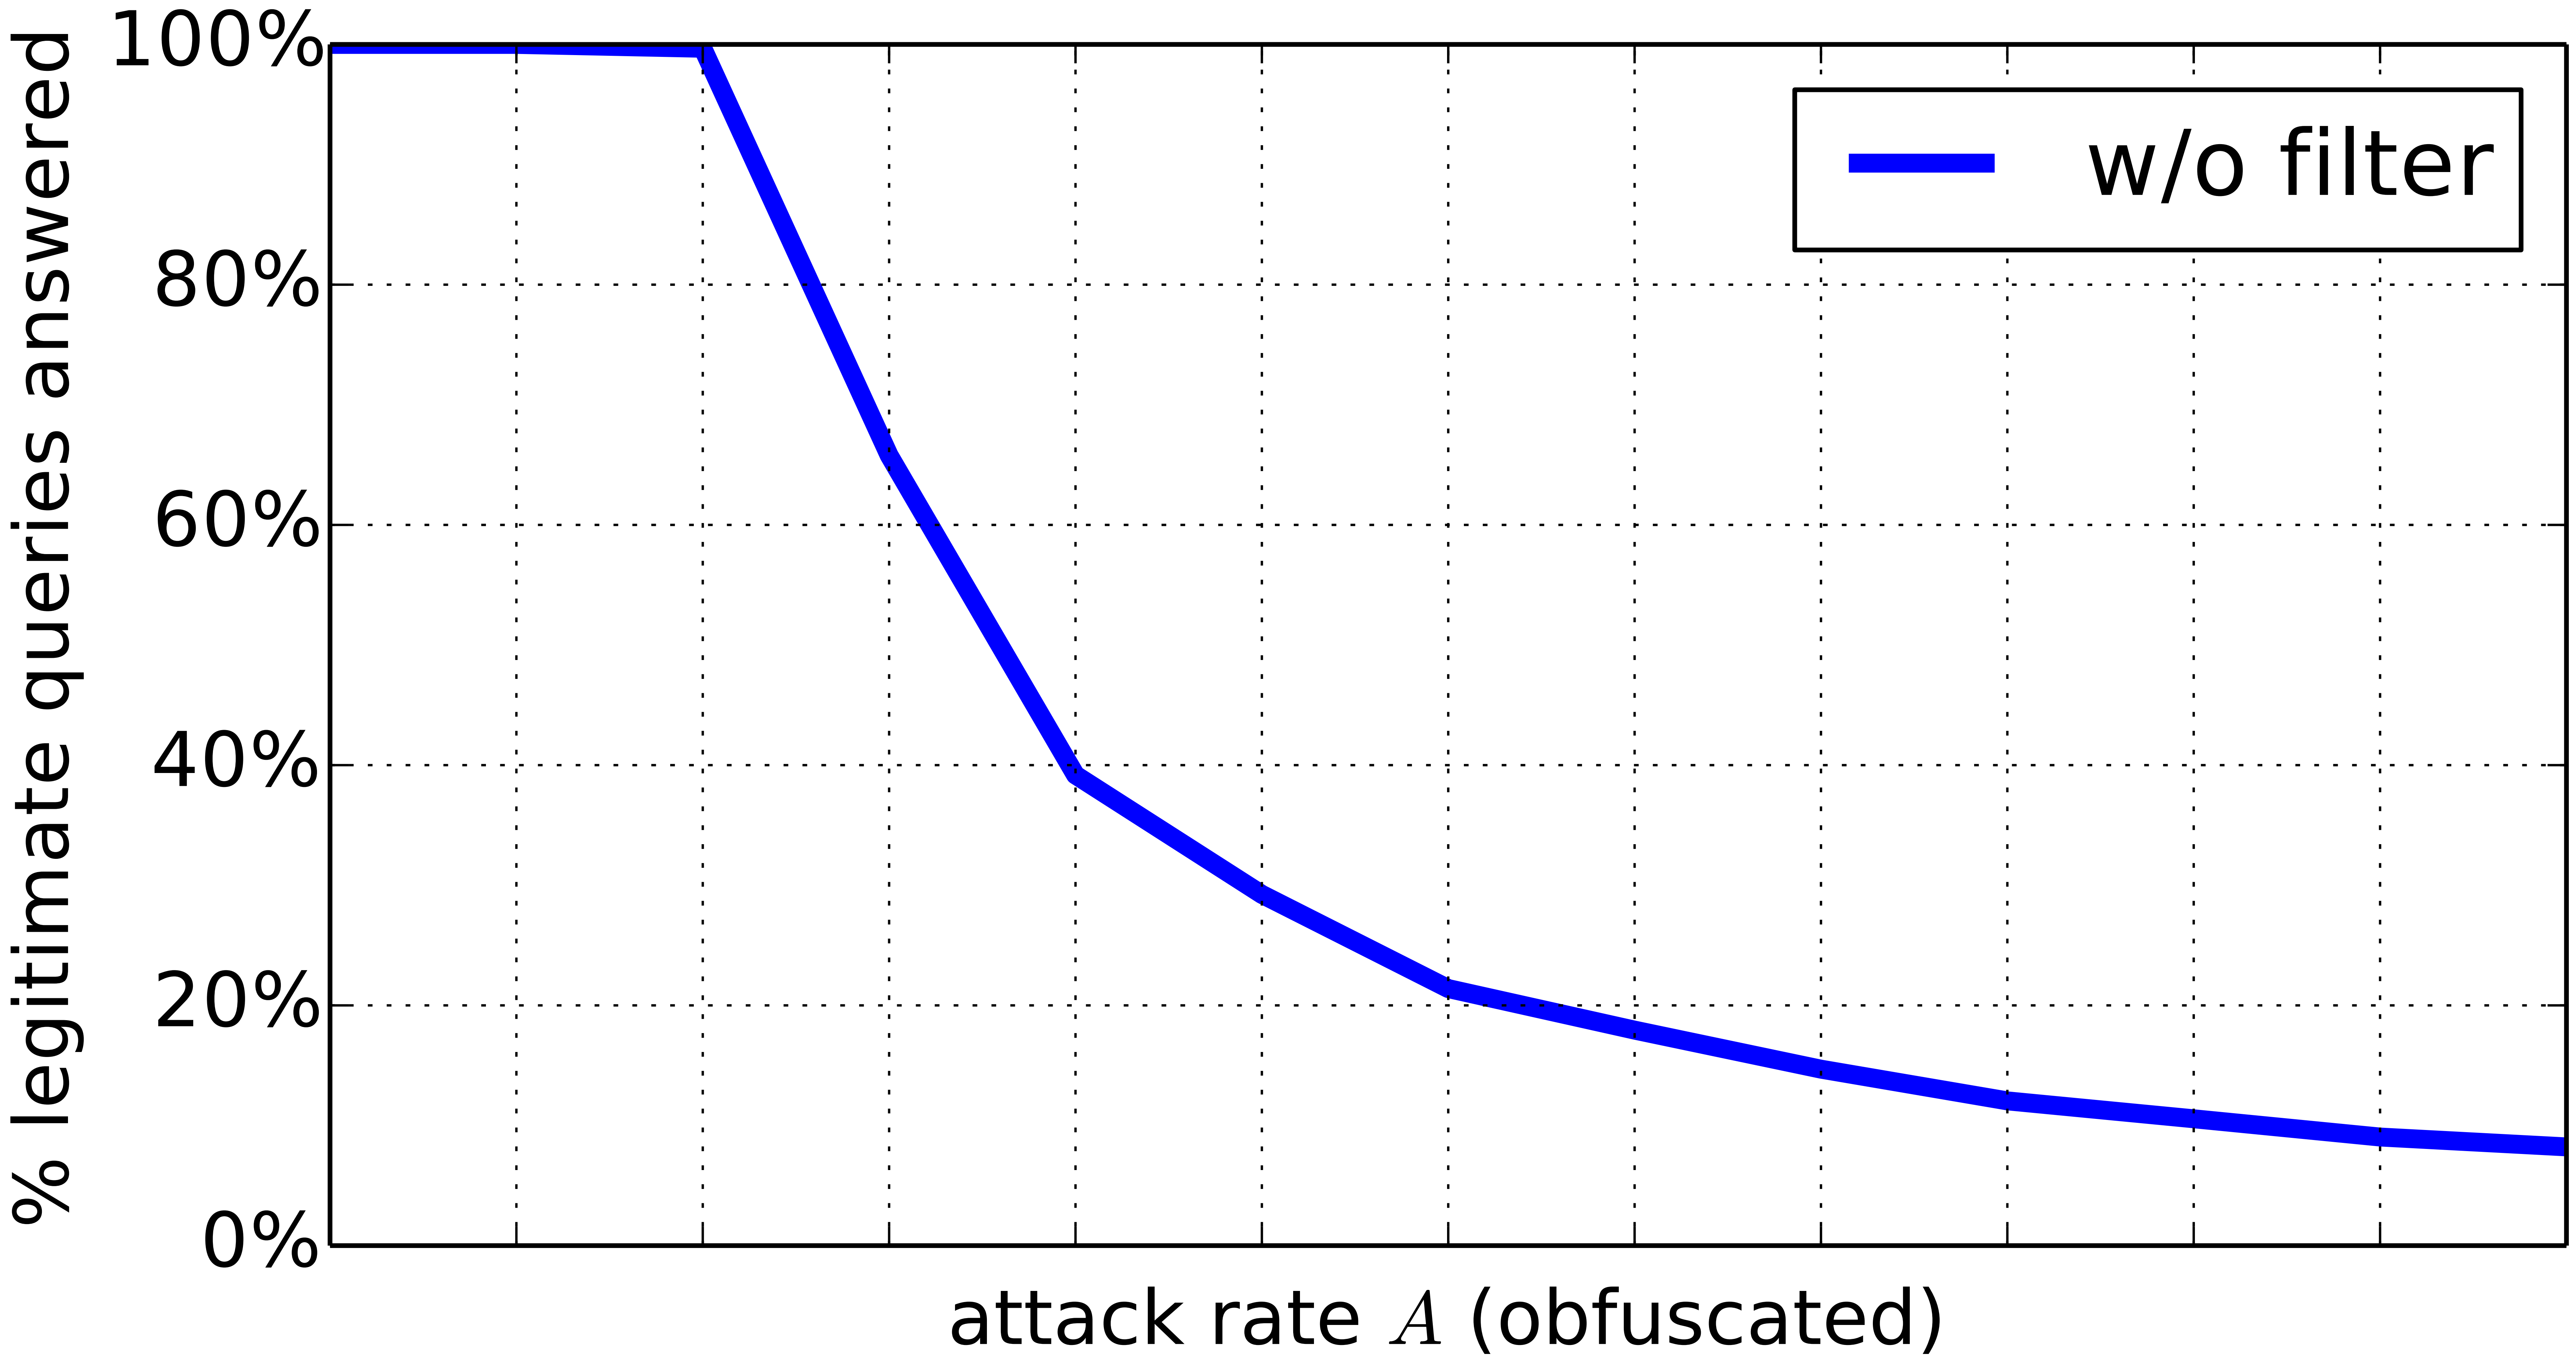

Testbed demonstration of filtering

Fixed rate of legitimate traffic, increasing rate of attack traffic

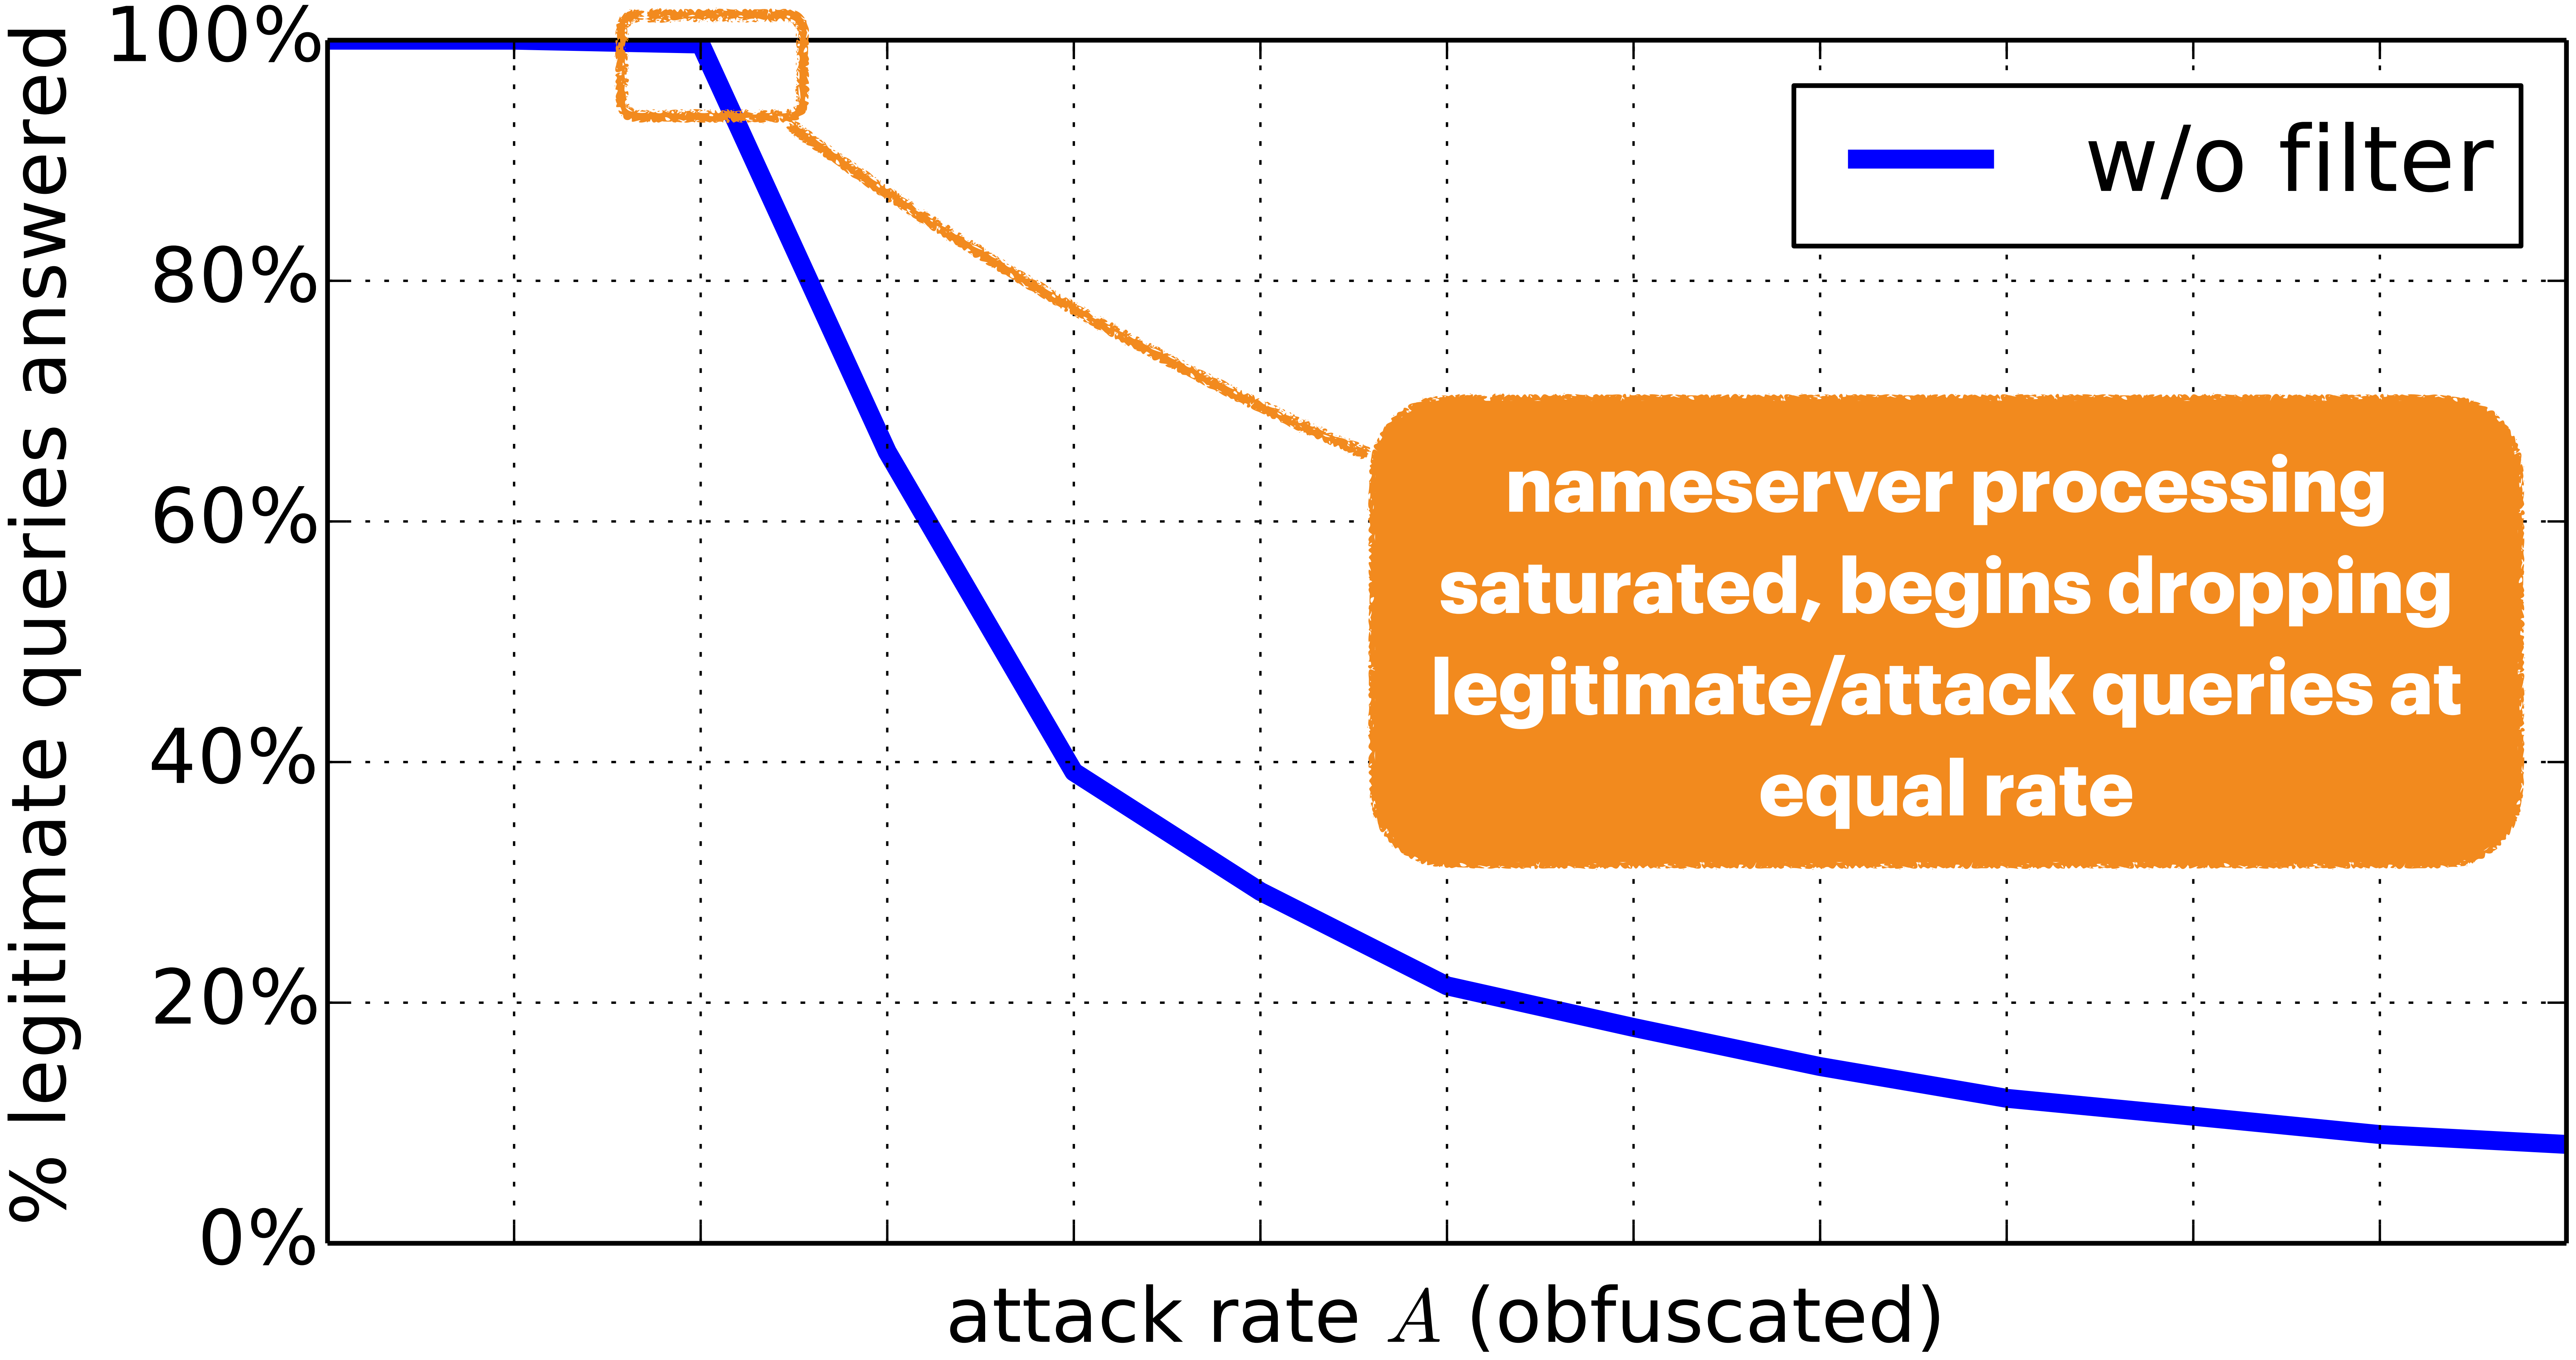

nameserver processing saturated, begins dropping legitimate/attack queries at equal rate

Testbed demonstration of filtering

Fixed rate of legitimate traffic, increasing rate of attack traffic

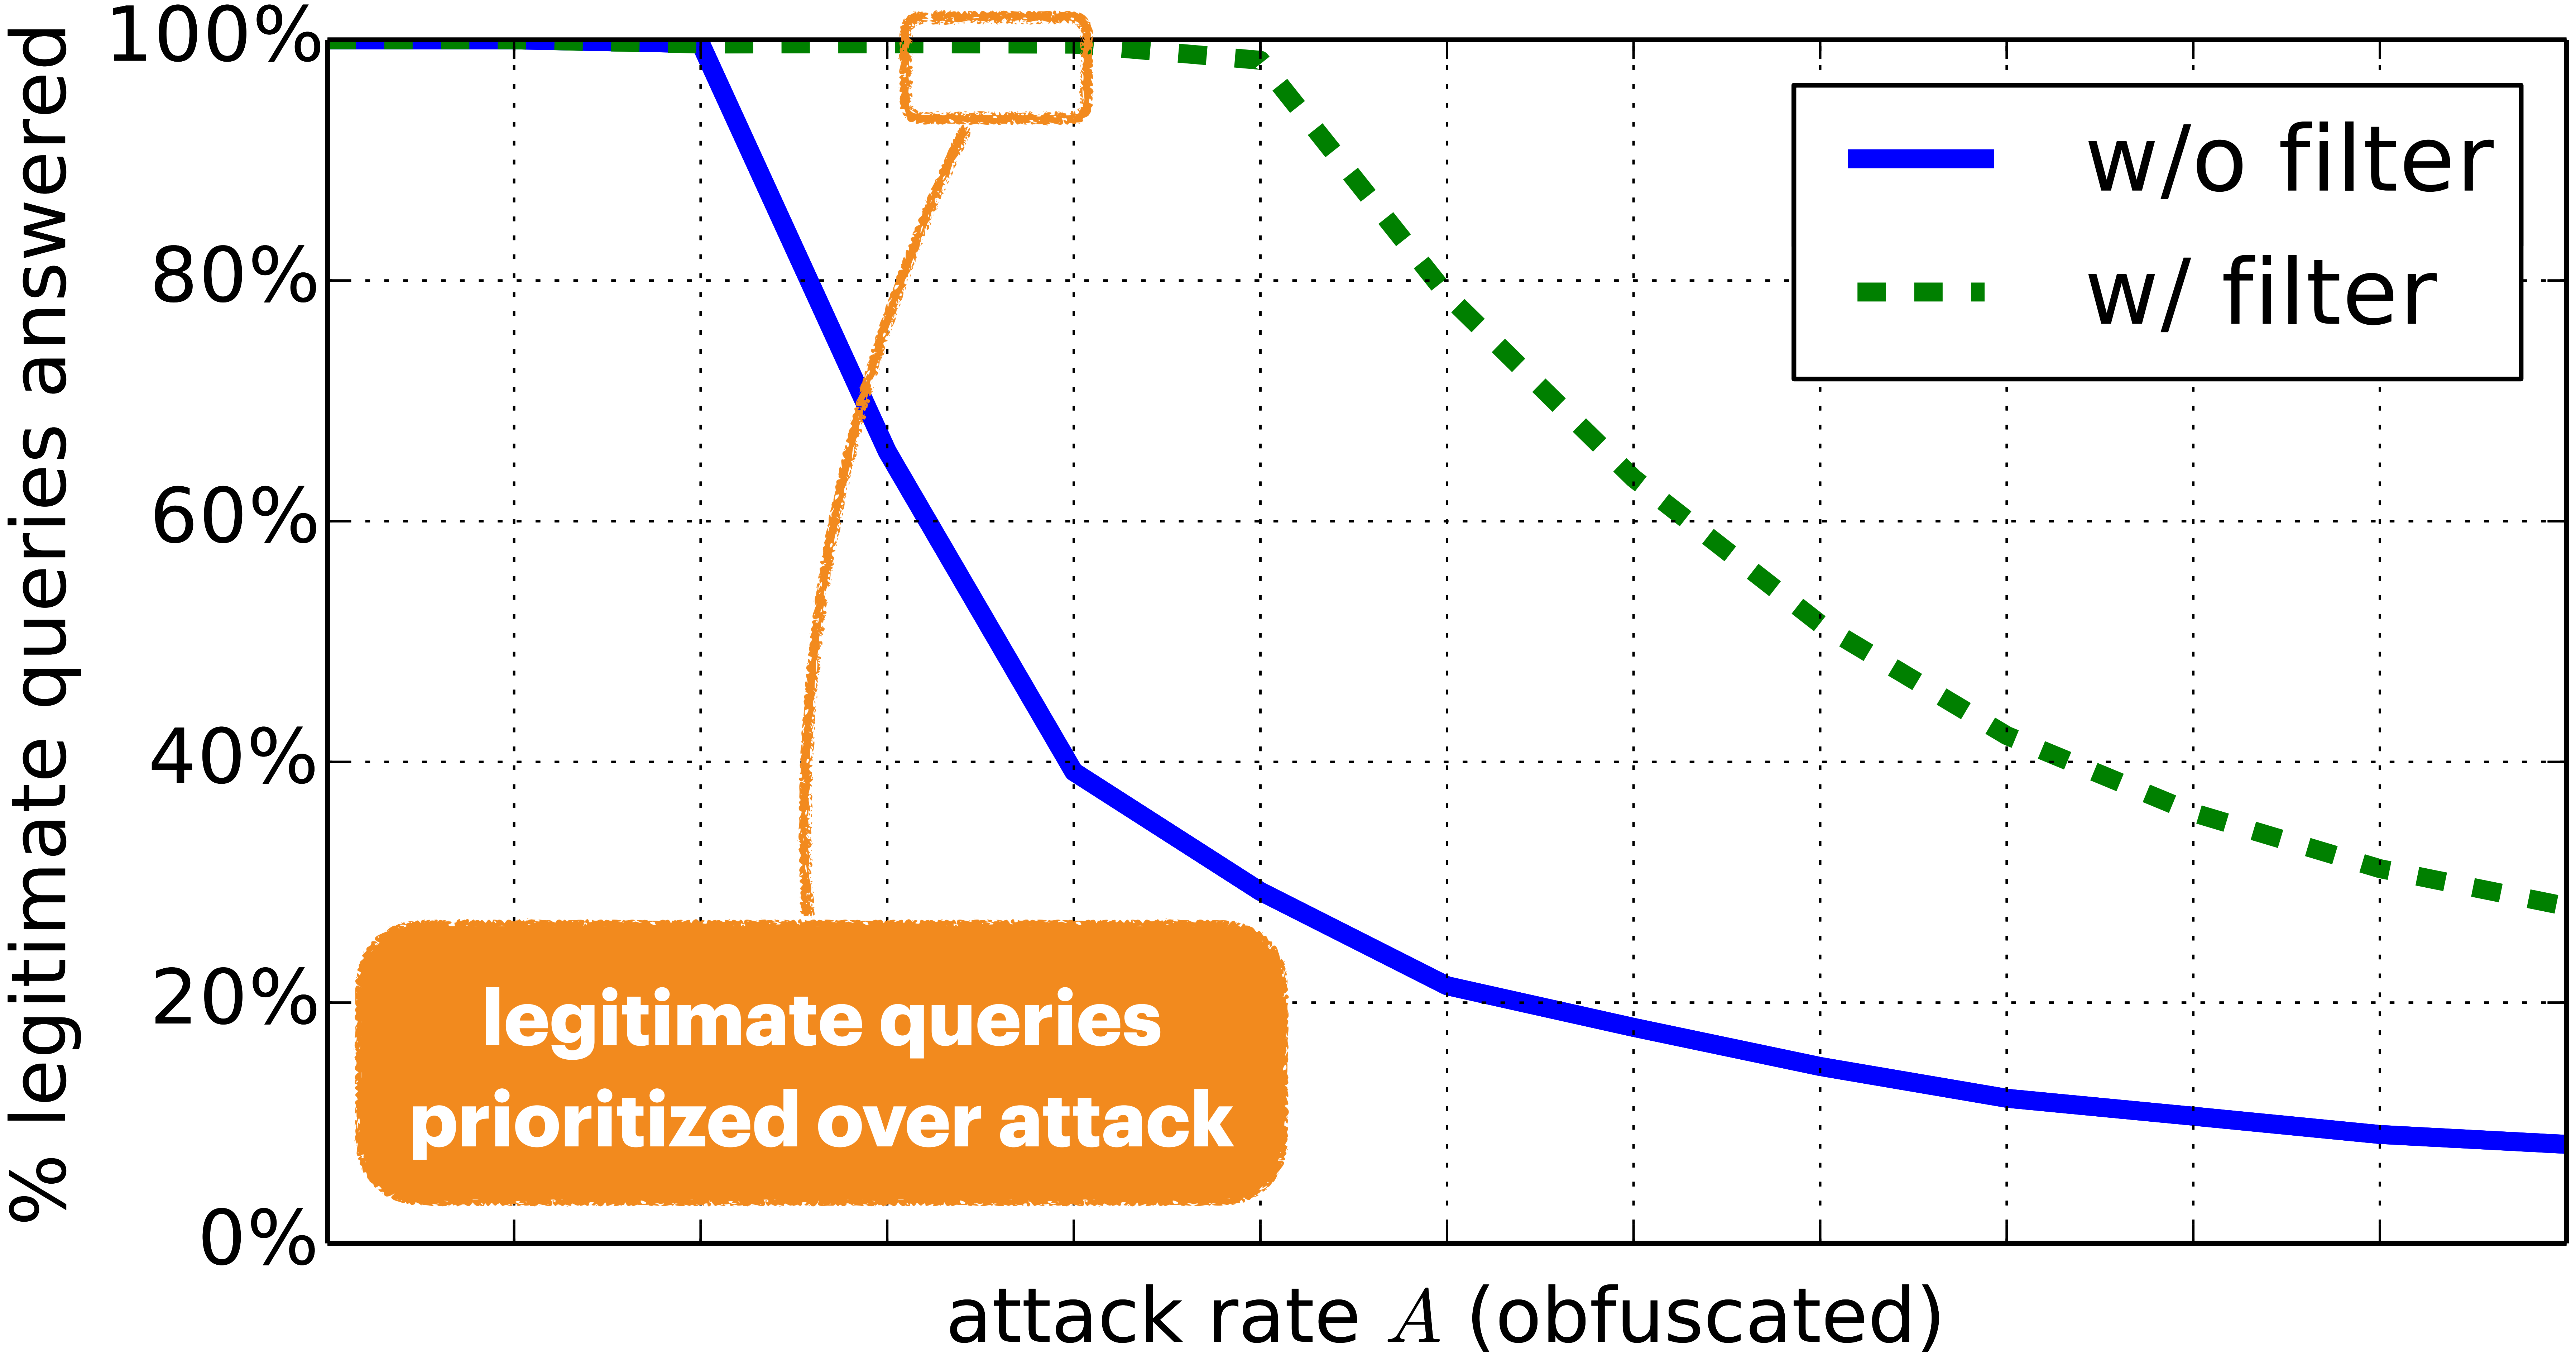

legitimate queries prioritized over attack

Testbed demonstration of filtering

Fixed rate of legitimate traffic, increasing rate of attack traffic

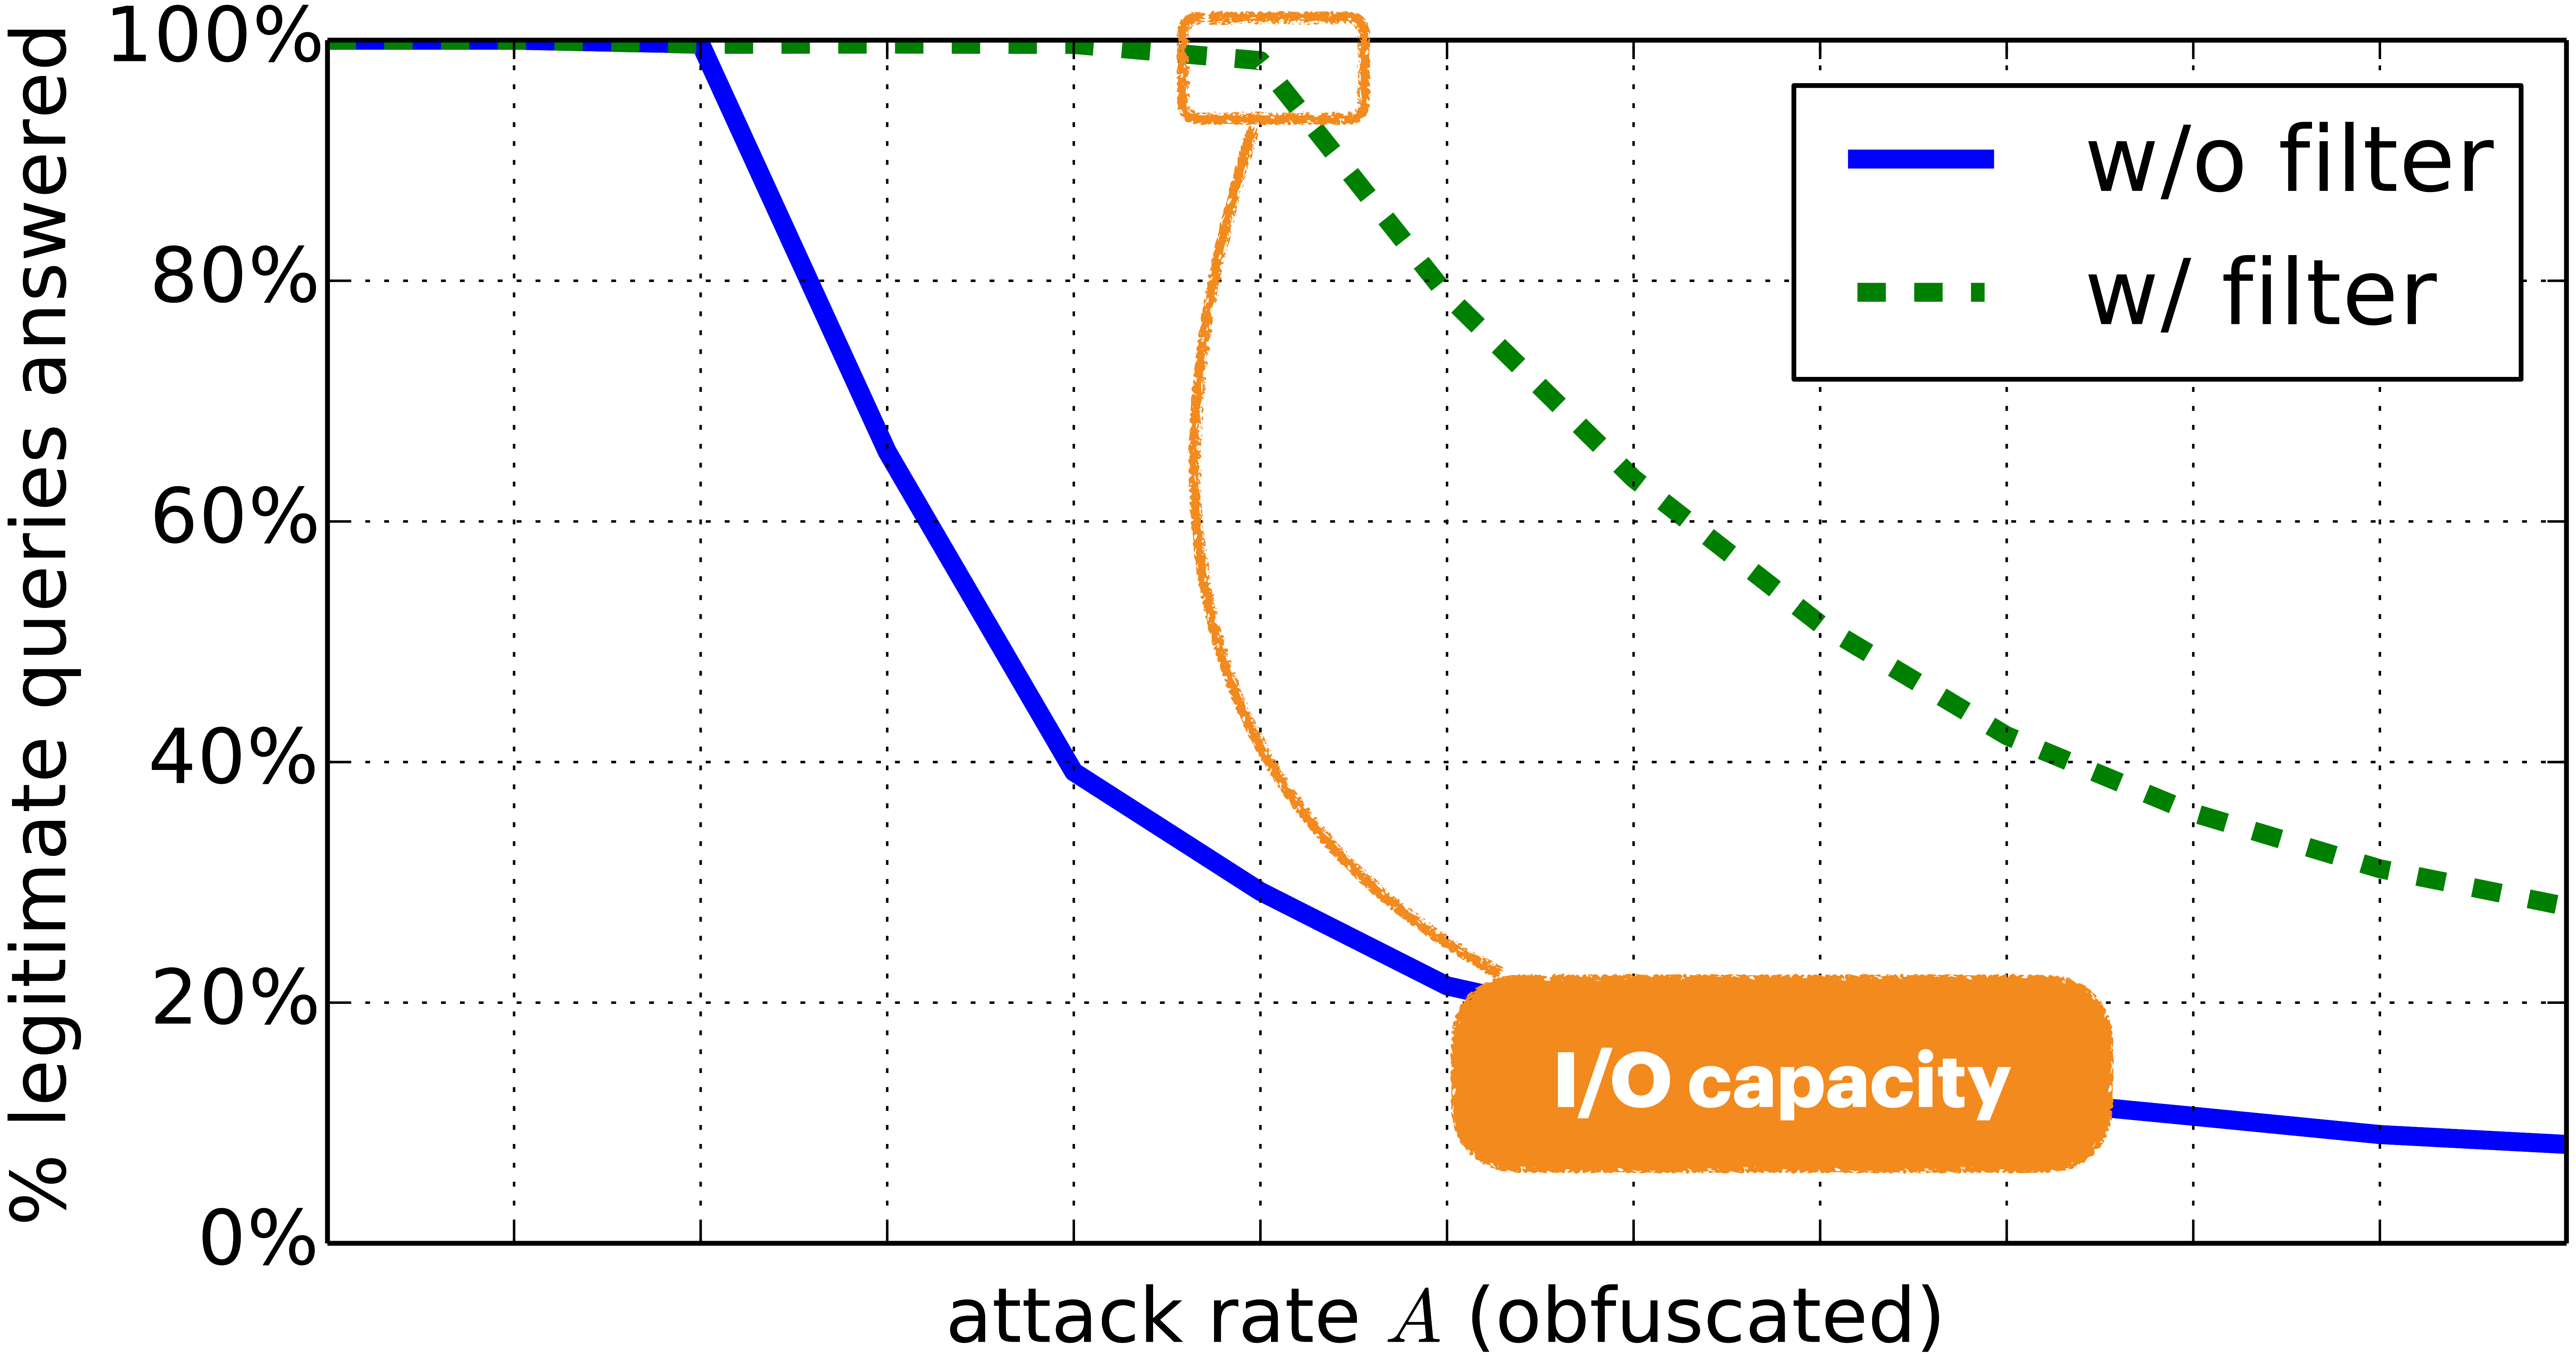

I/O capacity

In the worst case,

an attack is indistinguishable from legitimate traffic.

So that Akamai DNS is always available, we build contingencies for even extremely unlikely but high impact scenarios.

In the worst case,

an attack is indistinguishable from legitimate traffic.

So that Akamai DNS is always available, we build contingencies for even extremely unlikely but high impact scenarios.

1. overprovision bandwidth in peering links
2. and compute in nameservers
3. compartmentalize infrastructure to minimize collateral damage

Conclusion

We've presented design principles and experiential insights gleaned over two decades of architecting, deploying, and operating Akamai DNS.

We've shown how the architecture provides:

1. Failure Resiliency
2. Attack Resiliency

please read our paper for more!