Reducing the impact of IXP maintenance

Will Hargrave // LONAP

will@lonap.net

LONAP
London Access Point

What is an IXP?

- A switched fabric for interconnecting networks
- BGP is the control plane for network operators to signal reachability
- What happens to production traffic when an IXP operator does maintenance?

A typical IXP

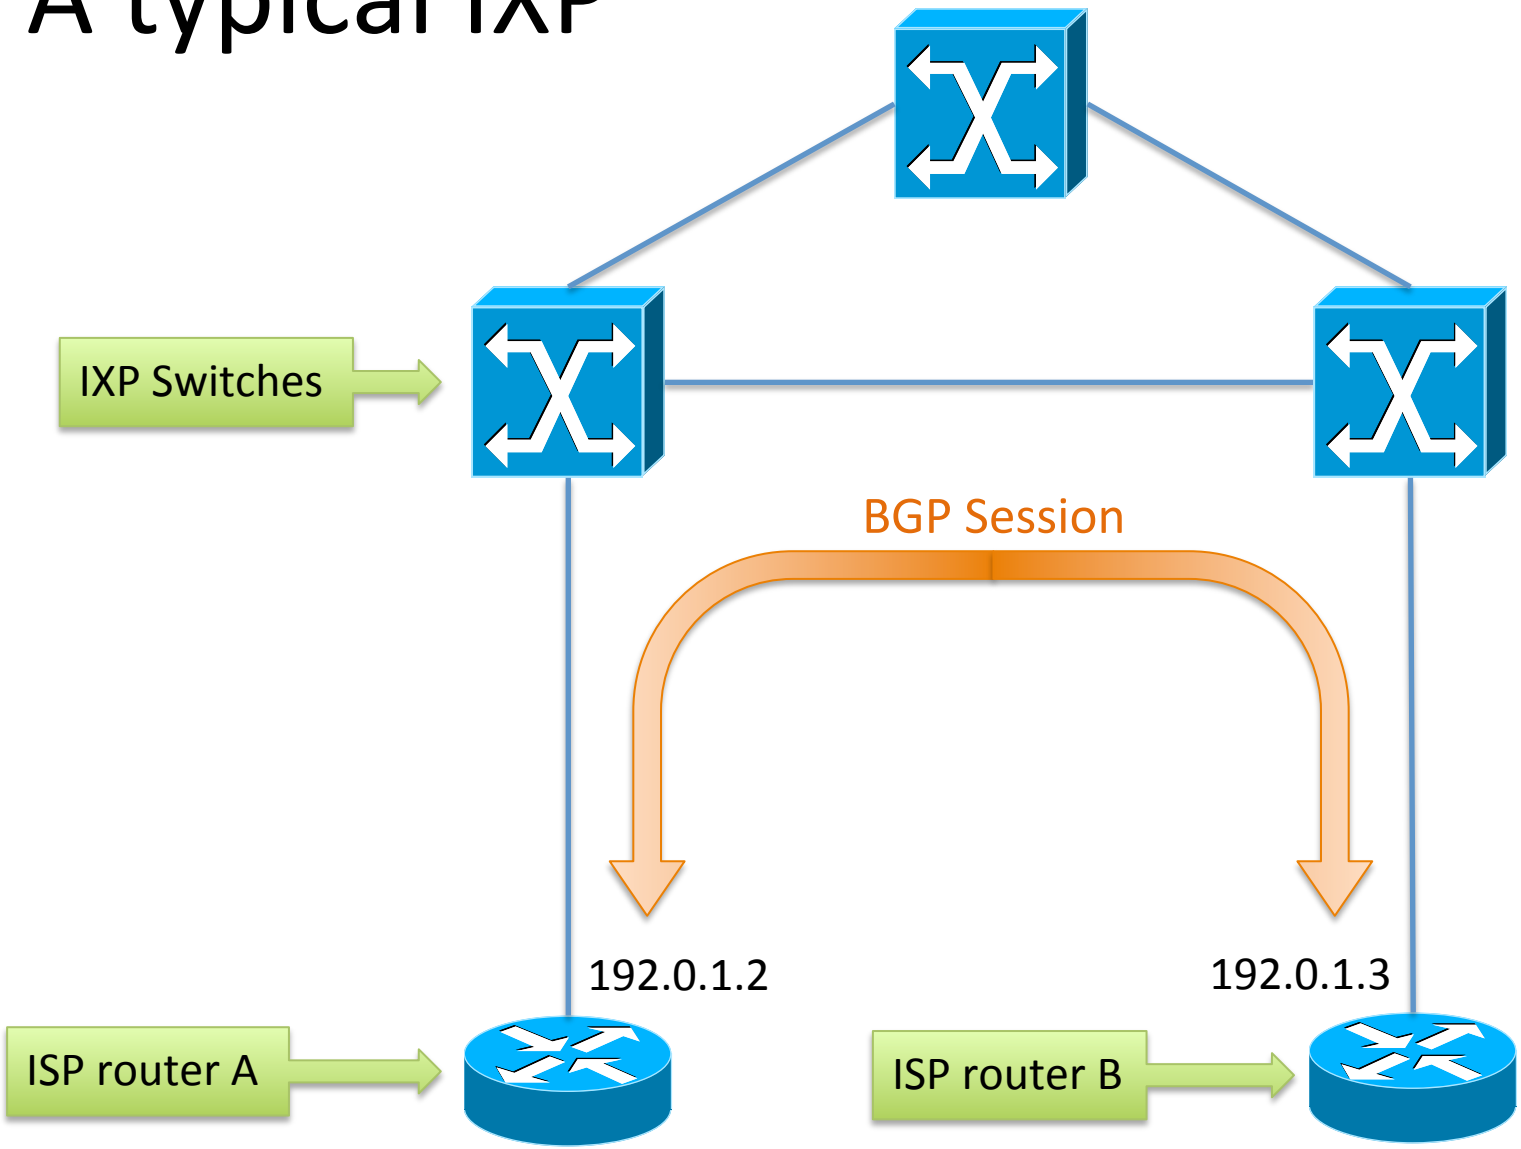

Maintenance

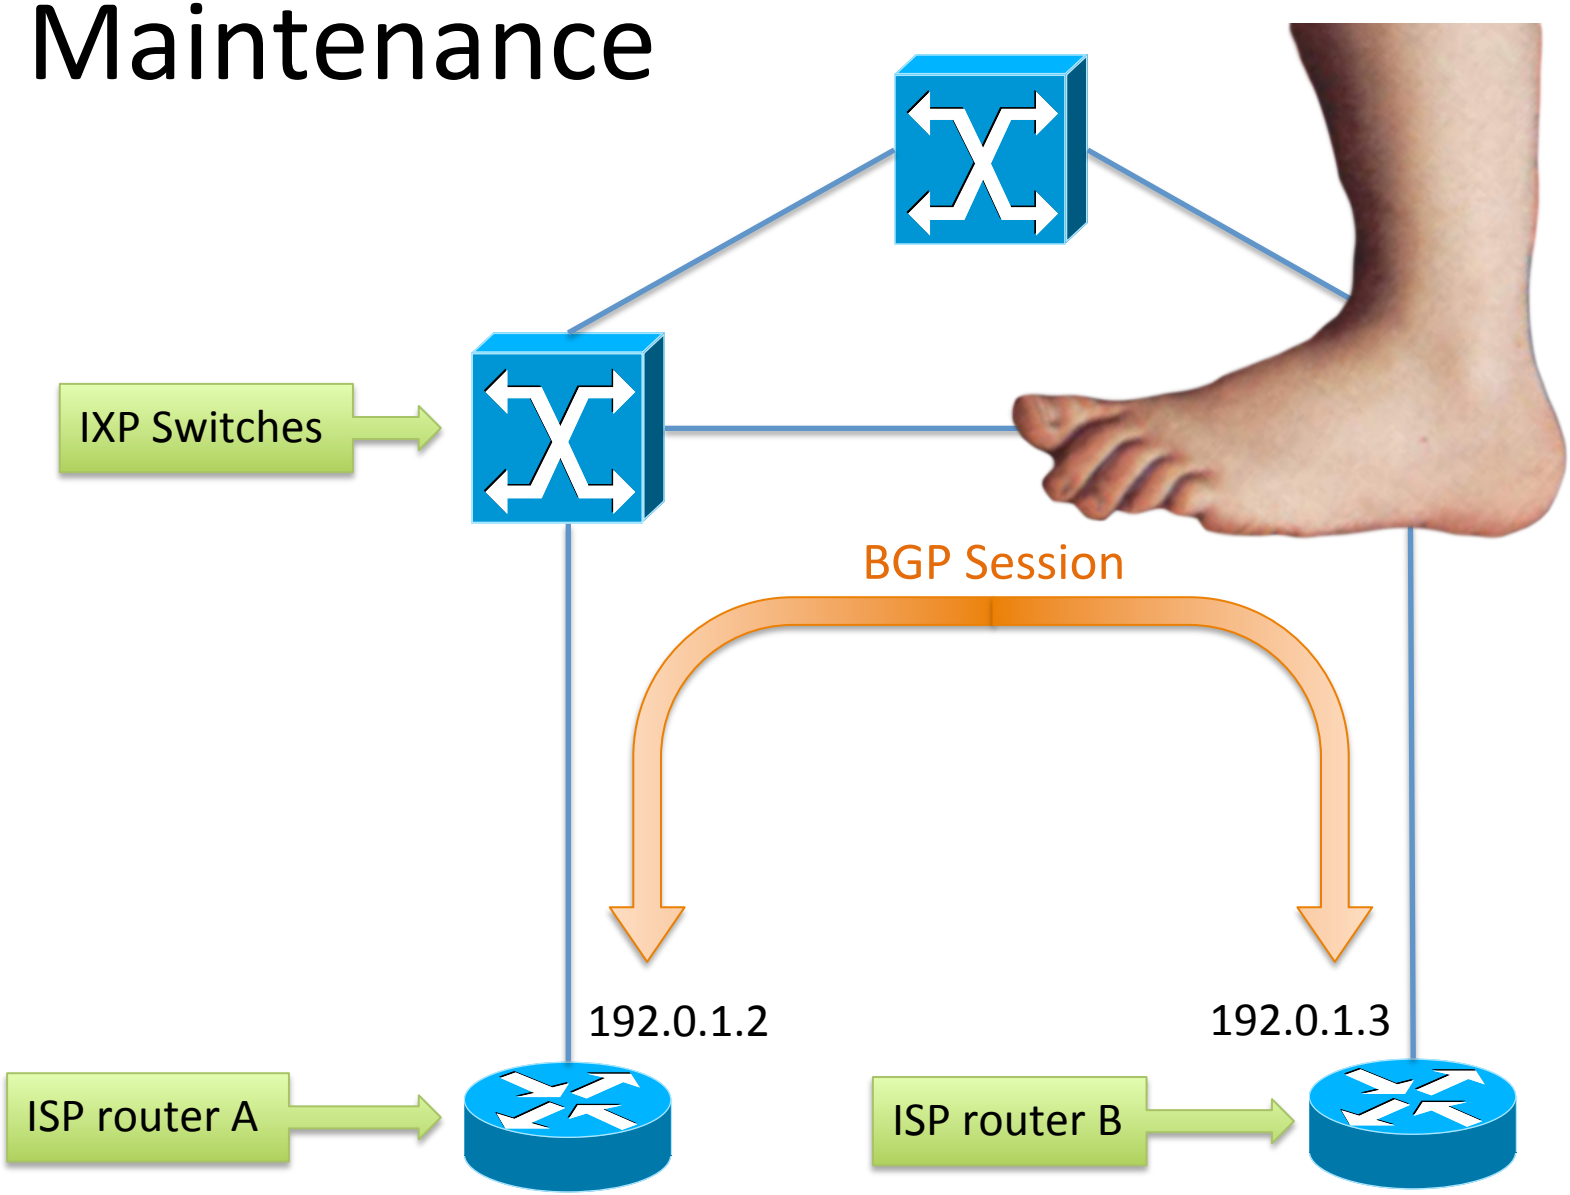

Waiting for BGP holdtimers…

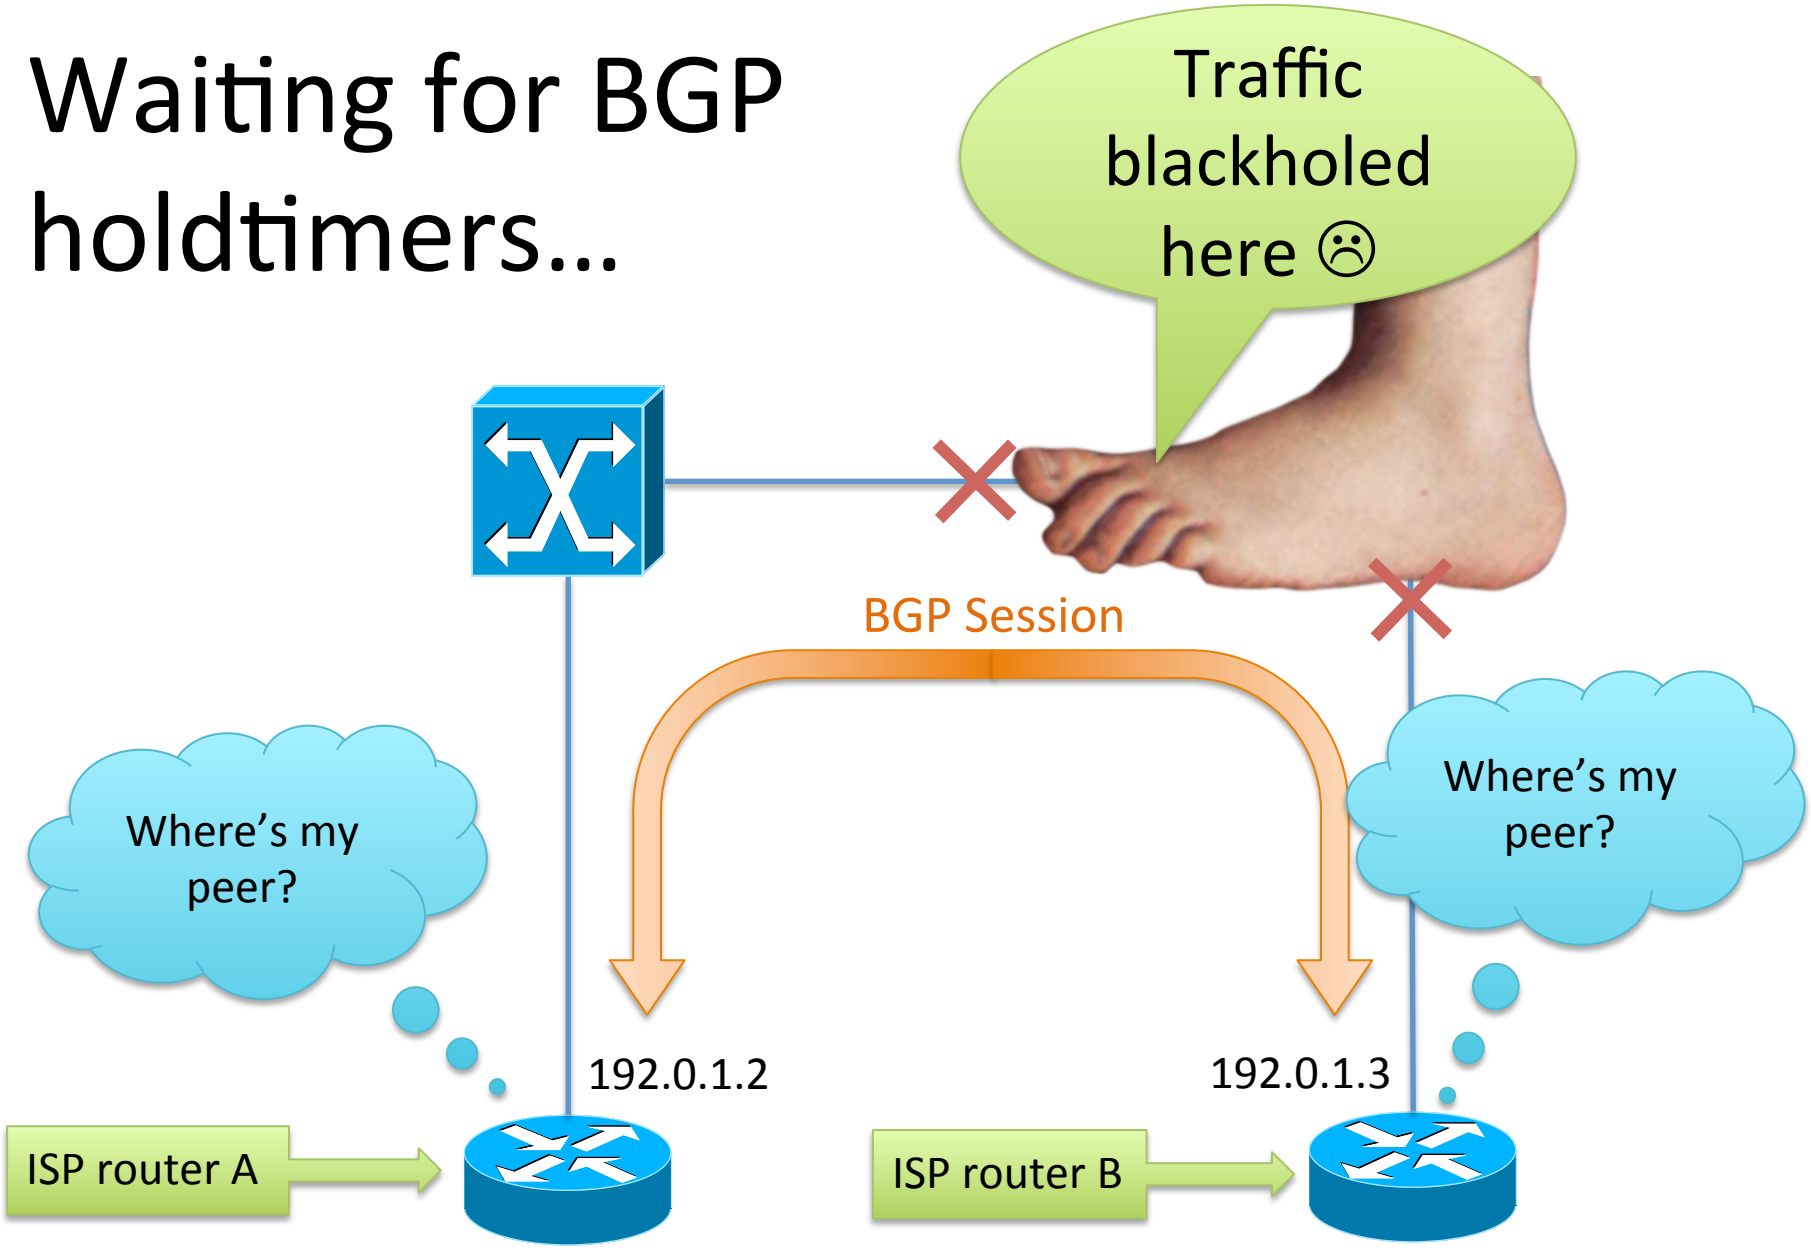

30 to 180 seconds later….

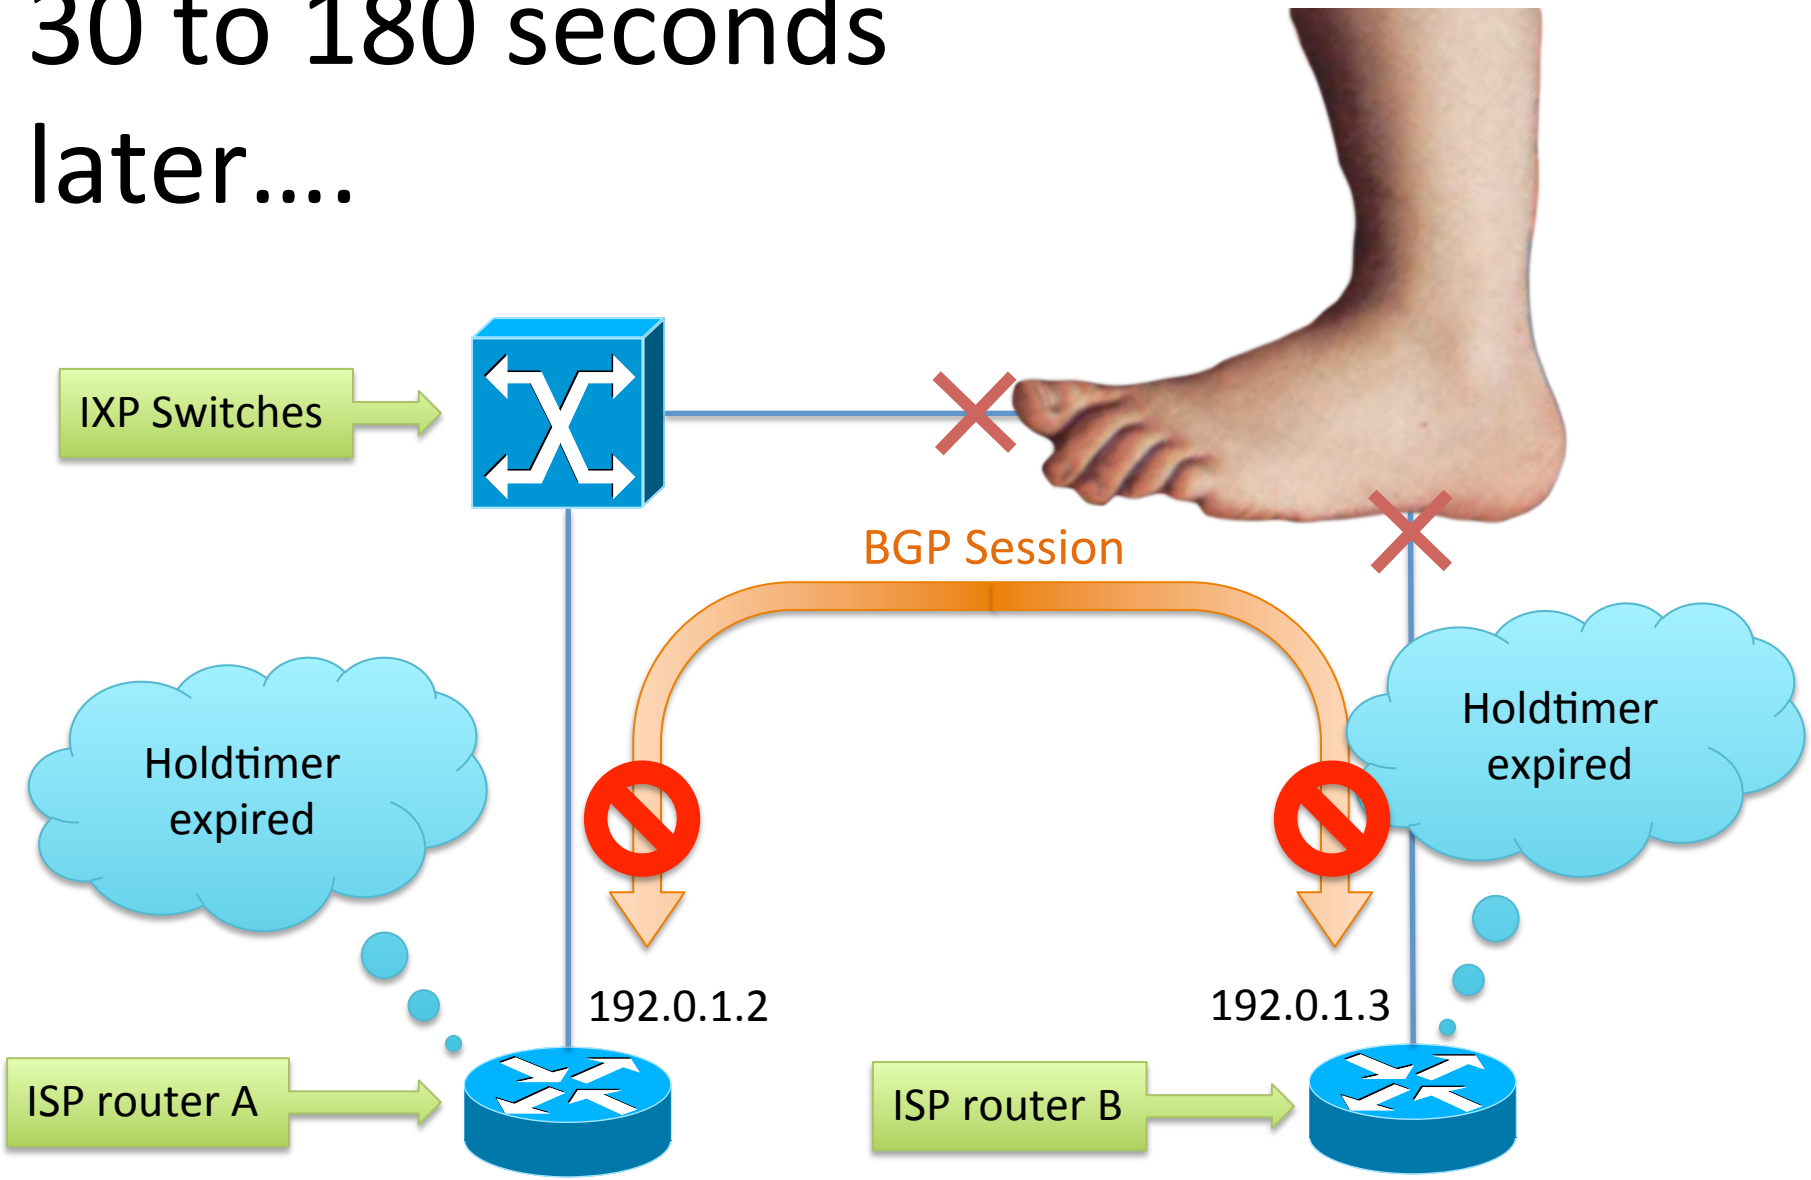

Improving the experience

- In many cases, IXP switch maintenance causes 90seconds+ blackholing of production traffic
- We are doing things the wrong way round!
- **Solution:**
  1. At the start of the maintenance window, tear down the control plane!
  2. Wait for traffic to diminish (3-5 minutes)
  3. Now do your maintenance
- How to tear down the control plane?
  Answer: L4 ACLs on IXP port!

L4 BGP ACLs on IXP

```
entry DenyBGPv4e {
  if {
    source-address 5.57.80.0/22;
    protocol tcp;
    source-port 179;
  } then {
    deny;
  }
}

entry DenyBGPv4i {
  if {
    source-address 5.57.80.0/22;
    protocol tcp;
    destination-port 179;
  } then {
    deny;
  }
}

entry DenyBGPv6e {
```

Your IXP subnet

Block in both directions, otherwise sessions will re-establish

Obviously repeat this for IPv6 too

Conclusions

- We tested this during several LONAP maintenances with good success.
- Subsequently, tested at INEX in late 2013.
- Traffic removal in < 3mins
- Restore traffic when you are ready
  - even after multiple reboots – think microcode upgrades, mistakes

End

- This benign technique makes a better internet
- Time for a BCOP, RFC, other document??

- Questions? Comments?

Will Hargrave
will@lonap.net

LONAP
London Access Point