Do We Need to Rethink Network Monitoring?

Kemal Šanjta
Customer Success

ksanjta@thousandeyes.com
@kemalsanjta

ThousandEyes

Troubleshooting Lifecycle

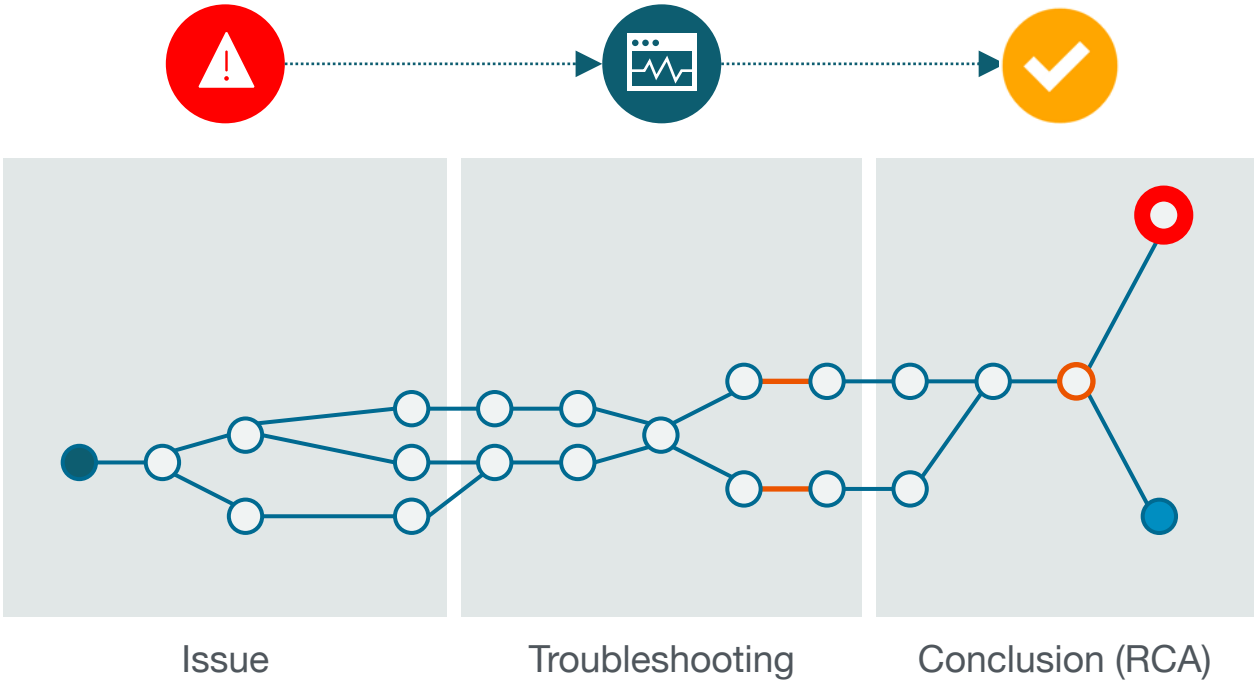

Issues with Troubleshooting Tool Set

Traceroute

- Fails to discover nodes
- Fails to discover links
- Reporting of false links

Ping

- Control plane reliance

Improvements to Troubleshooting Tool Set

Ping and traceroute as good as a starting point,
but we realized we need something more

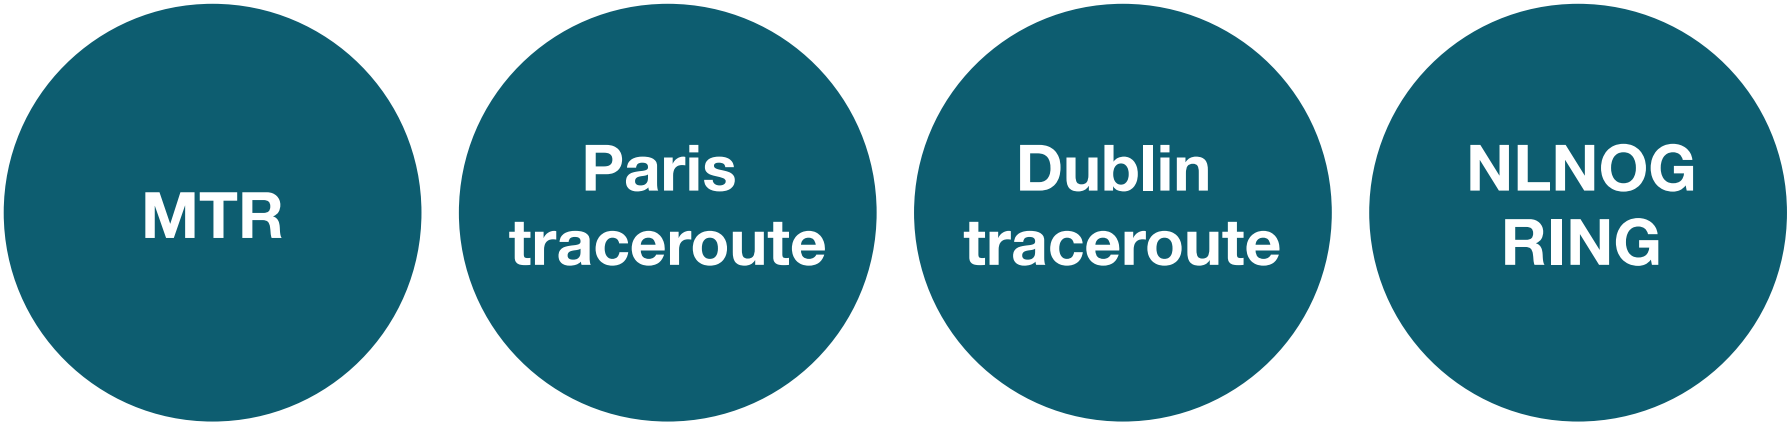

Various Sources for Alerting

- SYSLOG
- SNMP
- Streaming telemetry solutions
- Various “collections” (ssh logins, etc).

Can your control plane handle it?

What is the Problem?

TIME

Reactive nature of troubleshooting

Slow response

Service degradation

Unhappy customers

Is There Other Way?

Open Question

Is there any way to be proactive?

Advantage: Large scale data sets and machine learning (large companies)

AUTOMATION

We discovered…

- **Python**
  (and countless libraries)
- **Go Programming Language**
  (and its concurrency)
- **A few frameworks along the way like Ansible**

Once Automation Provided Results…

Are vendors telling
the full truth about performance
of their networks?

How Many Times Have You Heard?

- Linecards rebooting as a result of solar flares?
  (No root cause analysis)
- Counters for _exactly that_ issue are not user exposed?
- Counters exist, but you need to be a linecard level wizard to get them?
  (Involves knowing a good deal about architecture and silicon/ASIC type)
- Backplane was hit with this specifically crafted packet that took your fully redundant backplane down?
- Control plane cannot handle it?

Automation Gave Us A Product Called…

VENDOR DISTRUST

Active Network Monitoring

Challenges with Active Network Monitoring

- Large scale/enterprise networks moved to CLOS Fabric Designs
- Limiting the “blast radius”
- Smaller scale devices, in turn, suffer from smaller RIB/FIB sizes and weak Control planes

Are They Really Smaller Scale Devices?

- **Juniper PTX1000**
  24x100GbE, 72x40GbE, 288x10GbE = 2.88Tbps
- **Cisco NCS500 series**
  32x100GbE, 32x40GbE, 128x25GbE, 128x25GbE = 3.2Tbps
- **Arista 7170 series**
  32x100GbE, 64x50GbE, 32x40GbE, 128x25GbE, 130x10GbE = 6.4Tbps

Depends on the angle… Better to lose 2.8Tbps – 6.4Tbps capacity compared to fully loaded ASR 9022 taking down 160Tbps

Conceptual Solution

- Utilize data plane to measure experience
  (fundamental concept behind the Active Network Monitoring)
- Synthetic Traffic (UDP or TCP)

Practical Applications for the Solution

- Commercially available
- Open source solutions:
  - Matroschka prober
    (testing your networks with GRE and MPLS Tunnels)
  - **OpenNetNorad** (Facebook Open source solution—UDP based)

Backbone Related Challenges

Label switched networks (backbone networks) utilizing features like auto-bw are not that straightforward to implement active network monitoring on.

Potential Solution for Backbone Networks?

- Probe underlying IGP paths
- Control over IGP paths means same rules apply
- Best IGP path == Best MPLS path (often)
- ”Some” coverage is better than no coverage!

Did We Forget About Something?

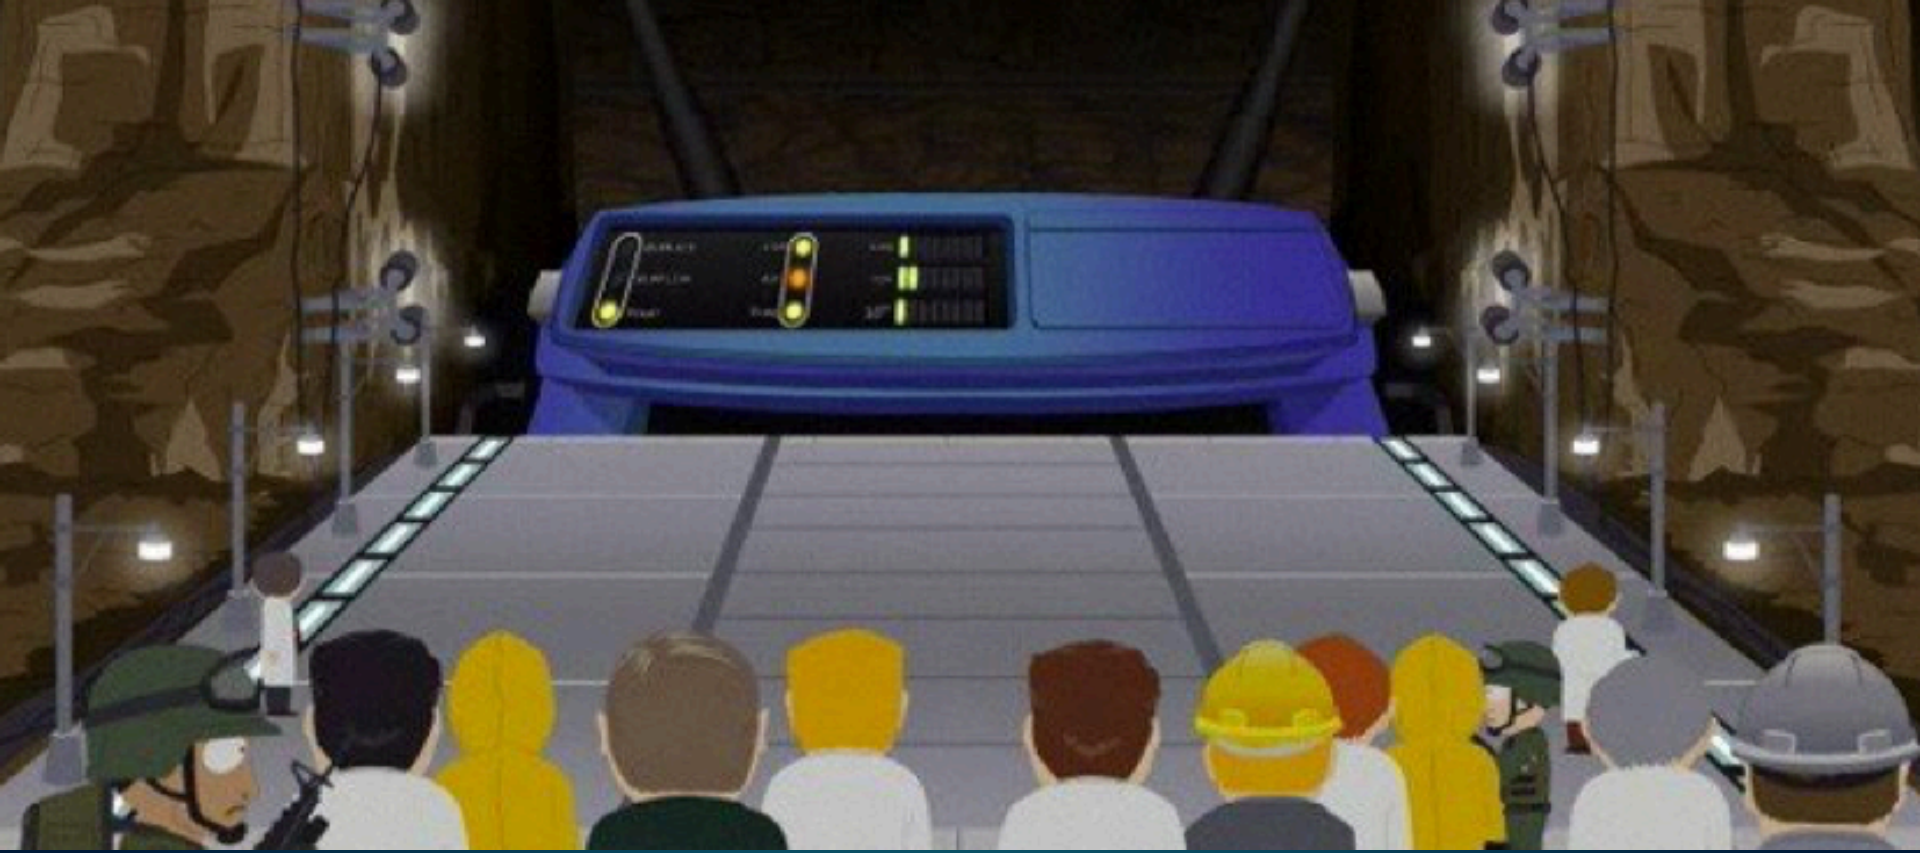

THE INTERNET

THE INTERNET

- Packet Loss
- Latency
- Jitter
- BGP
  (advertisements & withdrawals)
- Prefix hijacks

Solutions for Internet Monitoring

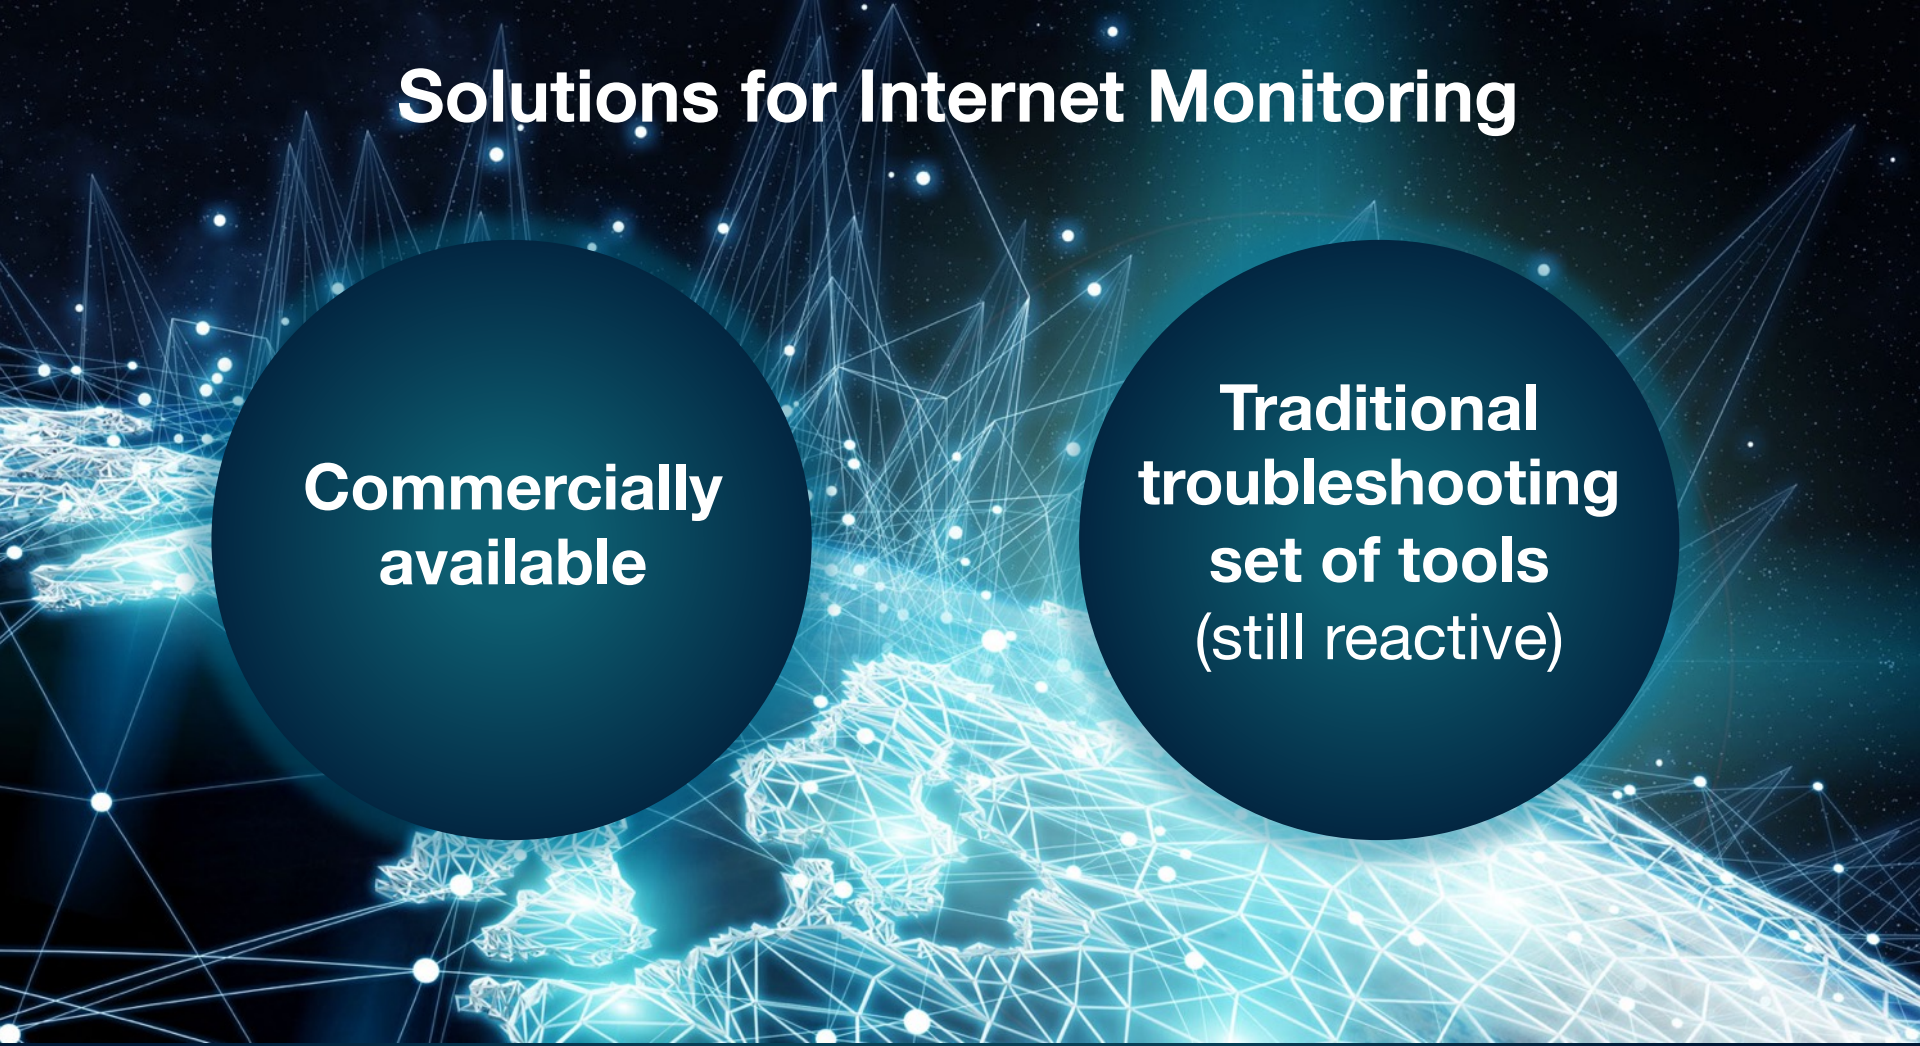

Conclusions

- Learn how to code
  (required skill to deploy and manage networks and market is moving towards it)
- Utilize research papers on data center and backbone design
  (do not repeat someone else's mistakes)
- Utilize both active and passive network monitoring

Conclusions

- Monitor performance of your internet paths as if life of your packets, and patience of your customers depends on it
- Don’t stop there – extend monitoring solutions to the services (know and monitor them and timely alert on issues)


ThousandEyes
Thrive in a connected world™