

Internet Measurement and Monitoring

Vinayak Hegde

Structure

- Part 1 – Basics of Measurement (35 min)
- Part 2 – Basic Statistics primer (35 min)
- Part 3 – Measurement Case Studies (20 min)
- Part 4 – Overview of Tools (20 min)
- Q&A – (10 min)

Why are measurement needed ?

- Capacity planning and network design
- Finding anomalies and fault detection
- Defining a baseline for policy / pricing
- Measuring adoptions of technology
- Mapping the Internet
- Academic research
- Measuring QoS and SLAs

What can you measure ?

- Latency
- Throughput
- Connectivity
- Periodicity

Types of Measurements

- Active Measurements

The active approach relies on the capability to inject test packets into the network and follow them and measuring service obtained from the network/application.

Types of Measurements

- Passive Measurements

The passive approach uses devices to watch the traffic as it passes by and collect data. Often they do not collect every data point but sample data

Active Measurements

- Pros
 - More “objective” since you can control some parts the measurement environment
 - Easier to emulate scenarios by scheduling, mimicing traffic patterns
 - Better control over sampling
- Cons
 - Measurement could modify the test environment
 - Increases network traffic

Passive Measurements

- Pros
 - Measures real traffic
 - Extremely valuable in network-debugging
 - Does not create extra traffic
- Cons
 - Can lead to processing lot of data. Proper sampling is crucial.
 - Can add extra devices to monitor live network
 - Privacy & Security issues

Software

- Remote monitoring (RMON)
- SNMP
- Netflow
- RIPE Atlas
- M-lab

Active Measurements

- One-way Measurements (OWAMP)
 - RFC 4656
- Two-way Measurements (TWAMP)
 - RFC 6038
- TCP Throughput Testing
 - RFC 6349
- Loss Episode Metrics
 - RFC 6534

Passive Measurements

- IPPM Draft -
<http://datatracker.ietf.org/doc/draft-morton-ippm-act>
- IPPM Draft -
<http://datatracker.ietf.org/doc/draft-zheng-ippm-frame>
-

Challenges and Considerations

- Setting up the test environment
- Understanding Traffic patterns
- Removing “white noise”
- Understanding layers underneath
- Sampling correctly

IETF WGs

- IPPM
- BMWG
- LMAP
- PMOL directorate

Stats 101 for Measurements

Basic terminology

- Distribution
- Mean
- Mode
- Median
- Variance
- Standard Deviation
- Population
- Sampling

Distributions

- Normal Distributions
- Poission Distribution
- Binomial Distribution
- Bimodal Distribution
- Bernoulli Distribution
- Lognormal Distribution
- Zipf's law

Sampling

- Process of Sampling
- Types of Sampling
 - Simple Sampling
 - Stratified Sampling
 - Systematic Sampling
 - Cluster Sampling
 - Opportunity Sampling





More concepts

- Sampling Bias
- Outliers
- Margin of Error
- Sample Size
- Statistical significance
- A/B Testing
- Correlation
- Percentiles

Gotchas & Traps

- Correlation is not causation
- Visualise your distribution (Anscombe Quartet)
- Confounding Variables

Anscombe's Quartet



Case Studies

IPv6 Adoption Metrics

- <http://www.worldipv6launch.org/measurements/>

Reachability of Anycast DNS K-Root Servers

- <https://www.ripe.net/publications/docs/ripe-393>

Latency Analysis

- Cover one example from cable cut
 - <http://research.dyn.com/2008/01/mediterranean-cable-break/>
- Cover one example of blocking
 - <https://labs.ripe.net/Members/emileaben/a-ripe-atlas-view-of-internet-meddling-in-turkey>

Tools Overview (demo)

RIPE Atlas

RIPE Stats

Measurement Labs

Q & A

- Thank you