Afilias SM

2nd Global Annual Symposium on DNS Security, Stability and Resiliency

James Galvin
Howard Eland

Data Data Everywhere

- How to manage the huge amounts of collected data
  - How raw is raw?
  - In front or behind routers and load balancers
- Today it is typical to collect full packet traces
  - How long do we keep it?
  - Where do we keep it?
- How is data moved?
- How much data is retained?
  - All of it?
  - Summary?
  - Analysis?

Here a slice, there a slice

- Will traditional measurement principles work?
  - Collect everything, keep as long as practical
  - What if I need “X” but can only get “% of X”?
    - Engineering constraints
    - Bandwidth constraints
- Will traditional statistical principles work?
  - Does sampling really work?
  - This can and should be tested empirically
- Health analysis: there should exist an operational advisory board to help understand the collected data

Unintended Consequences

- Zone size increases: DNSSEC could cause about a factor of 4 increase
- As the number of validating resolvers increases, what will be the amplification to normal operations?
- What does this mean to DDoS attacks?
  - How will this affect data collection?
  - How will this affect bandwidth health?
- Is “instantaneous” propagation still the right model for zone transfers?
  - What if you are at the end of a thin pipe?

Now you see it, now you don't

- DNS views are likely to become mainstream
- Views at the enterprise level are common place
- Governments are moving in the direction of filtering what the Internet looks like to their constituents
- Filtering is likely to be mandated or regulated
- All data (the complete zone) all the time may not be an option in the future