Measurement in ISP Backbones
Capacity Planning and SLA Monitoring

NANOG 26 - October 2002

Tony Tauber

Genuity Network Architecture

GENUITY

Different approaches meeting in the middle

- Analytical
  - Models
  - Numbers from vendors and industry literature
    - **MTBF (Mean Time Between Failures)**
    - **Fiber cuts/distance**
  - Simplifying assumptions
    - **Number of nodes reduced**
  - Rules of thumb
- Empirical
  - Based on observation vs. theory
  - Measured in the real world
    - **Interface counters (SNMP)**
    - **Flow Statistics**

Capacity Planning – Inputs and Methodology

- Derive Traffic Matrix
  - Topology
  - IGP Costs
    - **Run SPF offline**
  - Gather flow stats
  - Reconcile with interface stats
    - **Conservation of flows (in == out)**
  - Model onto network
    - **Find best-fit matrix**
- Try "what-if" scenarios
  - Failure analysis
    - **Mostly single failures with inheritance**
      - **This box fails, then this card fails, then these links fail**
    - **Multiple (complex) failures add another dimension to the matrix**
      - **Probably not that necessary given low likelihood**

Capacity Planning – stuff to make your head hurt

- Look at empirical trends
- Incorporate forecasts from Sales department (!)
  - New demands aren't point-to-point but point-to-multipoint
- Consider largish exogenous factors
  - Changes in exit capacity (eg. interconnection)
  - Behavior of large sources and sinks
    - **Very big customers may shift their traffic among providers**
    - **The law of large numbers doesn't apply in these parts**
  - Construction delays
    - **Also what intervals to expect?**
      - **If the average buildout takes 3 months what does the occasional 12 month interval do to expectations?**
      - **How to model economic ramifications of building too fast vs. too slow?**
  - New applications (eg. peer-to-peer file sharing, data replication)

Capacity Planning Challenges

- Sampling
  - Grab a monotonically increasing counter at two different times and infer the rate to be the difference in the counter value divided by the length of the time interval
  - Detailed information on actual traffic arrival rates is lost
- Bursts
  - Instantaneous offered load exceeds instantaneous capacity
  - Queuing (jitter) or drops (loss)
  - Data (Internet) traffic is very variable and appears to be self-similar
    - **When traffic is aggregated, peaks don't smooth out so much**
      - **Fat pipes need commensurately big overhead to handle big bursts**
    - **Time of day, week (and year?) effects are pronounced**
- Difficult to calculate latent (true) demand
  - Will removing one bottleneck just reveal another?

Capacity Planning Challenges, too

- How well does the approach scale?
  - How many stats (and how much data) from how many devices?
- Get most bang for the buck in the data we choose to gather
  - Inasmuch as we have to limit ourselves
- Implementation problems with vendors
  - MIB counters
    - **not compliant**
    - **missing**
    - **broken**
  - Ditto for Flow Stats
- Backend systems which drive collections and gather topology and routing hints are key to making this work
  - Historical descriptive information must be kept along with stats measures to make sense of them

Service Level Assurances (SLAs)

- A marketing tool
  - 100% availability guaranteed? For real?
  - SLAs involve these four things
    - **What's guaranteed to the customer**
    - **Expectations set with customer**
    - **What the network is designed to achieve**
    - **How the network actually performs**
- Hard to compare different companies' SLAs
  - Far from being apples to apples
- All about balancing risks
  - For really, really important things, a guarantee isn't enough
  - Like your fire insurance, you don't really want to collect that refund
  - Would you trust your life to IP?
    - **Well, if I could set the ToS bits...**

SLA Measurements

- Some metrics people (and possibly applications) care about
  - Availability, latency, jitter, packet loss
- Do our measurements catch extremely transient events?
  - That's where the bursts and bad things happen
  - Internal study using two probes with same period but out of phase
    - **Introduced episodes of loss**
    - **One or the other might catch the episode**
    - **Reduced the period (time between) probe packets**
    - **When probe period reached ~1sec, both would generally reflect the loss**

SLA Functional Challenges

- What *CAN* be measured serves a proxy for what may be of most interest to a particular user
  - Too many possible paths and behaviors for any potential unicast conversation
  - Point-to-point SLAs on a connectionless network?
    - **Measure by sending test traffic on same data path**
    - **More hardware (and expense) or integrated in current boxes?**
- Can we correlate measured data with known events?
  - Customer caused outages, planned maintenance
  - Hard to say what may have happened a week and a half ago to cause a particular blip in monitoring traffic data
    - **Much less a customer-reported incident after the fact**

SLA Big Picture Challenges

- Education, education, education (for staff and customers)
  - What are these numbers telling us?
- Hard (if not impossible) to engineer to a desired quantitative behavior given the changing demands on the overall system
  - We may not live to see the promised land
- Will people pay for premium performance or the promise thereof and if so, how much?
- Economic analysis of SLAs?
  - What does all of this buy us or our customers?

Acknowledgements and Credit where it's due!

BBN/Genuity Network Analysis folks who've demonstrated particular patience in describing their art to me!

Doug Browning
Steve Chiu
Susan Coombs
Carl Holmberg
Kagan Gokbayrak
Irvin Schick
Fred Serr
Alex Shulman
Liyi Dai
David McGuigan

Thanks!