CS294-4: Network Support for Reliable Adaptive Distributed Systems

Reliability attributes of networks (26 Jan 2005)

Lecturer: Randy H. Katz *Scribes: David Chu*

1 Reliability relationships

Last class we outlined relationships among systems reliability metrics. We discussed what these mean in the context of networks. It may be more meaningful to measure networks from the host's perspective rather than the router's. DQ offers a particularly insightful tradeoff between quality and quantity of responses.

$$Avail = \frac{MTBF}{MTBF + MTTR} \tag{1}$$

$$Uptime = \frac{MTBF - MTTR}{MTBF} \tag{2}$$

$$Yield = \frac{Queries\ completed}{Queries\ offered} \tag{3}$$

$$Harvest = \frac{Data\ avail}{Complete\ data} \tag{4}$$

$$DQ = Data\ per\ query \times Query\ per\ second = Constant \tag{5}$$

2 Workloads

Let us define performability as performance under failure. This relates to Vahdat's red skies, perfect storm taxonomy. We then discussed the three dimensions:

- performance throughput, latency, data quality
- workload and faults
- cost

We discussed what "data quality" means at the network level. Quality is typically service-specific.

We could obtain source data-workloads from:

- synthetically generated workloads
- Department network. EECS Network Engineer is Fred Archibald.
- DETER project.
- Vern Paxson and LBL data.

Related to the workload, what is the Faultload: low likelihood, mid to high severity events. Symptoms of faults include: router software failures, router buffer overflows, router high CPU utilization. Causes of faults are: packets floods where each packet is the minimum allowable size; packet floods where each packet is a random destination. These causes stress the slow path of router's packet processing. Another common cause is configuration errors.

3 Paper discussion

A. Vahdat, J. Chase, M. Dahlin The Perfect Storm: Reliability Benchmarking for Global-Scale Services 2003

Good characterization of workload taxonomies. Metaphor of clear days, red skies and perfect storms helps convey message. We may be less concerned with investigating effects of particular exploits on performance. How do we deal with the 100-year event? Severe network events seem to be quite frequent (1/month).

4 Quality of the network

Humans perceive quality in discrete units. It might be modeled as a step function. For perceiving "bad" vs. "good" quality, this is a binary decision function.

Network protocols are designed for stability. Does this necessarily sacrifice the protocol's time to convergence when a change does occur?

Keynote (http://www.keynote.com/) offers network measurements from disparate sites around the globe. What can we learn from this distributed view?

Root cause analysis for service failures may be visualized as a graph. Could we do something similar for networks? For example, could we trace the correlated effects a DNS outage would have? How can we provide this type of visibility?