Scamper Development

Matthew Luckie

mjl@wand.net.nz

Last year’s scamper work

- Inferring and Debugging PMTUD Failures
  - Paper at IMC 2005
- Datalink access
  - Bypass IP stack for writing some packets
    - Avoid IP stack fragmentation
    - Precise control over what appears on wire
  - Fixed bug in FreeBSD for writing to DLT_NULL BPF devices (6.0-RELEASE): PR kern/82157
  - Did away with any code that manipulated the routing table

Last year’s scamper work

- TCP Traceroute (IPv4)
- Darryl Veitch’s TSC clock
- File Format Work
- SSL version of skdatad

Future Development

- SSL version of skcollectd
- Implement Doubletree Algorithm
- BGP
- Annotate data files with IP to DNS name at the time of capture
- Rocketfuel
- Pin-pointing Routing Changes

SSL version of skcollectd

- Currently, skitter uses kerberos 4.
- FreeBSD 5 onwards ships with kerberos 5.
- Replace authentication / encryption with SSL which is more widely deployed and understood.
  - Some of this work has been done

The Doubletree Algorithm

- Donnet / Raoult / Friedman / Crovella
  - (sigmetrics 2005)
- Goal is to reduce the number of times an interface is visited during large-scale topology discovery
- Intra-monitor redundancy
  - Many interfaces are visited repetitively from the same source
- Inter-monitor redundancy
  - Many interfaces are visited repetitively from the distributed monitors

The Doubletree Algorithm

- Avoid probing the destination and working backwards (as in Mercator)
  - a mesh of hosts probing the same address simultaneously might appear as a DDoS
- Challenge then becomes determining where (*h*) to begin probing
  - Selecting an appropriate value of *h*
  - Current suggestion is to determine the hop distance where most new interfaces are discovered

BGP

- Obtain routing table and updates by peering with a local router.
- Annotate data files with BGP data at the time of capture
- Re-probe path changes signaled by BGP update
- Use AS path to guide Doubletree, discussed as future work in sigmetrics 2005 paper.

DNS

- Annotate data files with IP to Name information
- Not really a focus until Doubletree is used, since goal is to reduce network load
- Leads to Rocketfuel…

Rocketfuel

- Spring, Mahajan, Wetherall (SIGCOMM 2002)
- Detailed ISP maps
  - They used 300 traceroute web servers…
  - … don't have required skitter monitor deployment for this (20 or so monitors) but it is worth thinking about
  - there are other contributions to use
    - BGP to reduce probe count (RouteViews)
    - DNS data to determine purpose, location of routers
    - Alias resolution

Pin-pointing Routing Changes

- Teixeria, Rexford. SIGCOMM NetTs 2004
- Straw-man proposal for inter-ISP cooperation through an “Omni”
- Use BGP + Local AS policy + IGP
- Not sure this fits with other proposed work, but it is something I’m interested in.

Feedback Invited