

Introducing Hedgehog

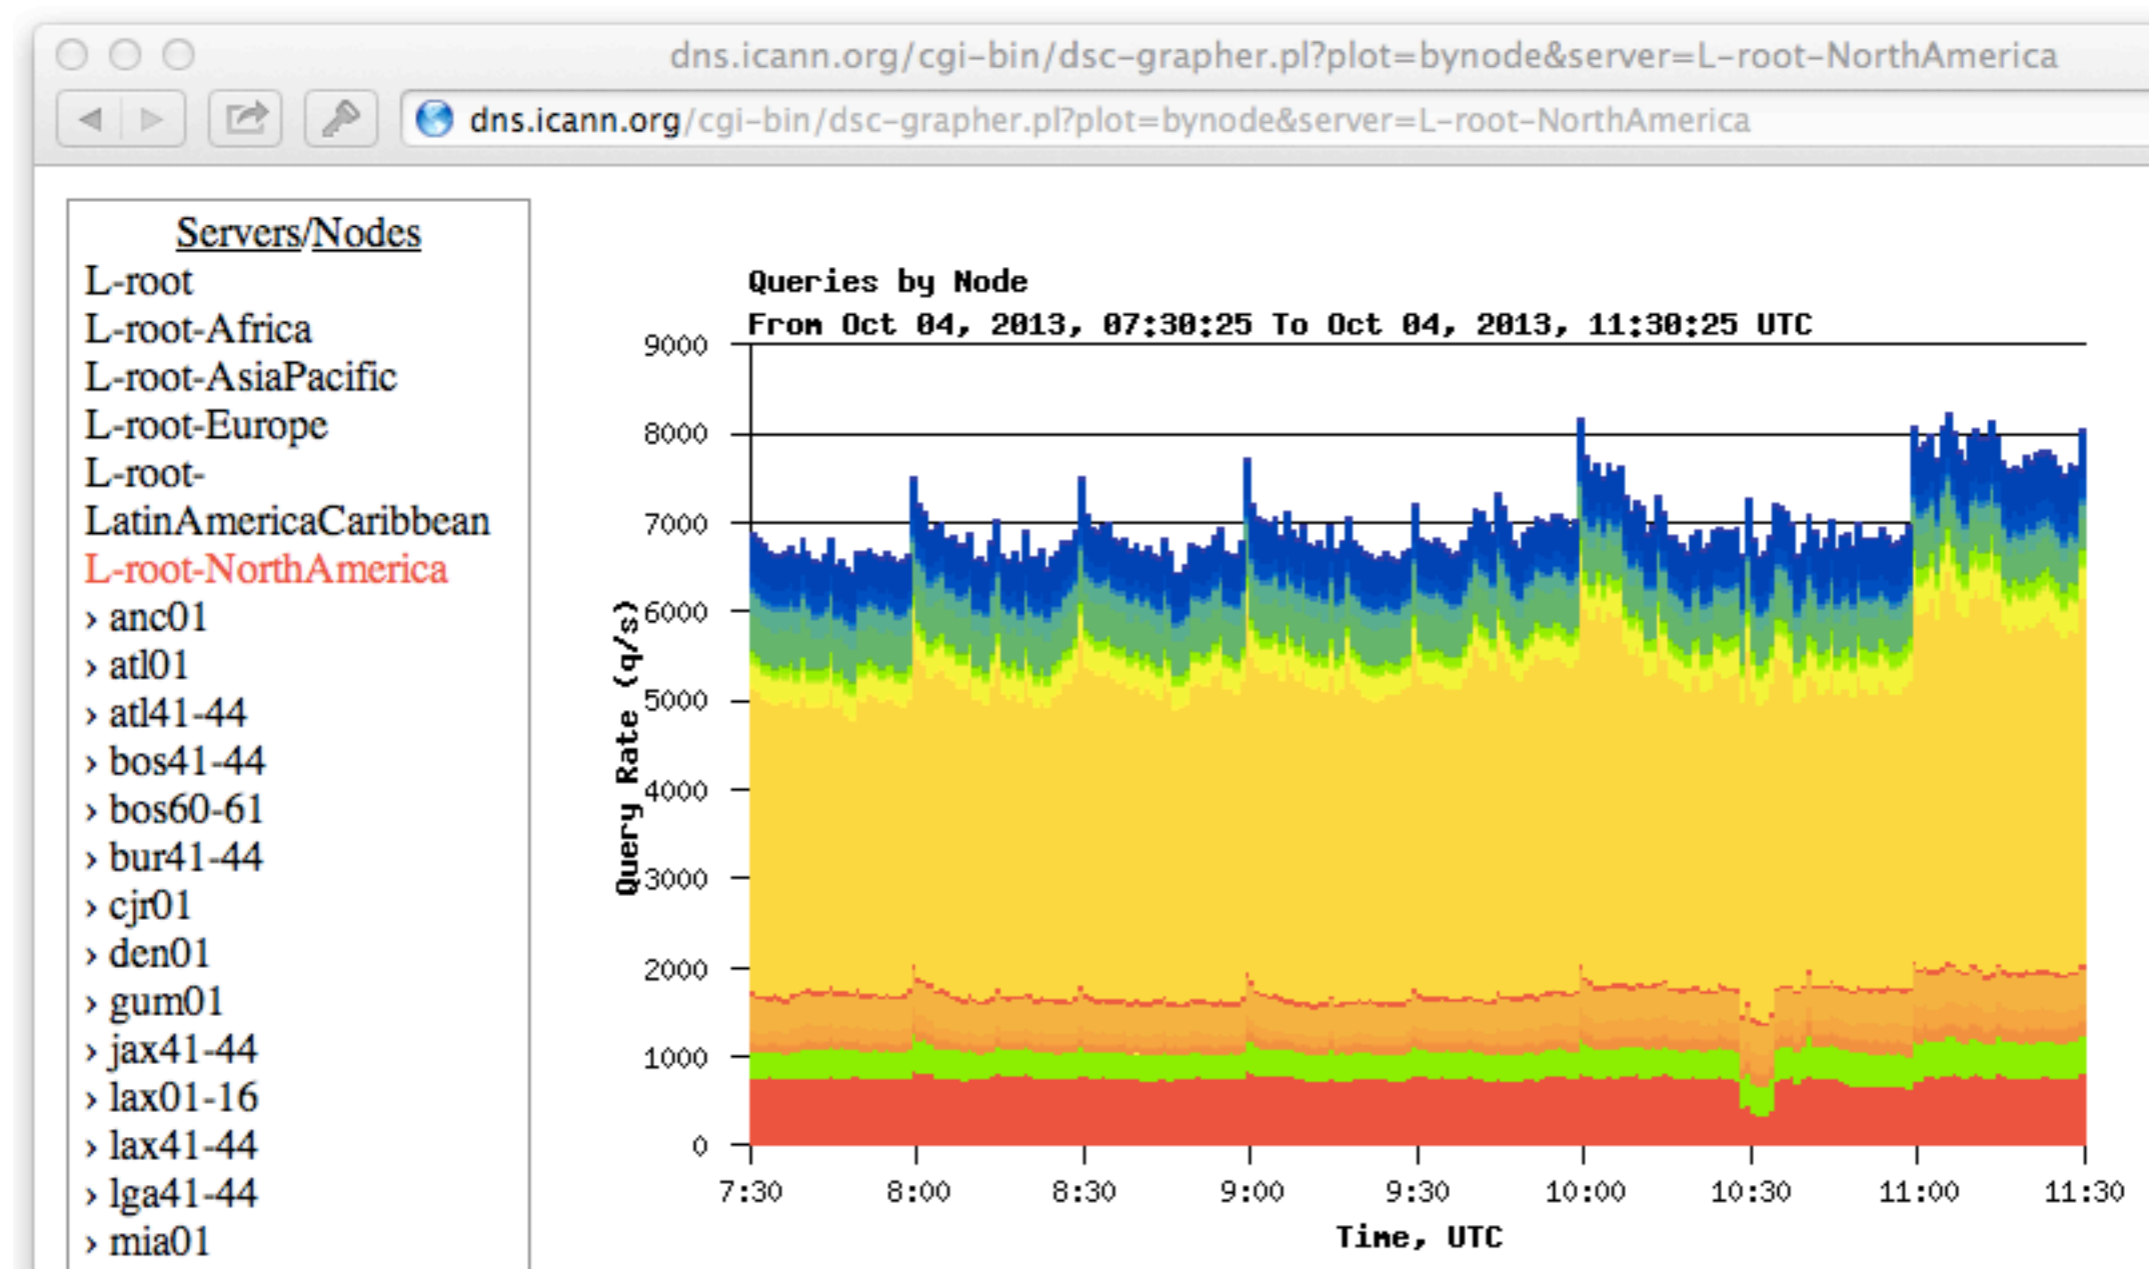
RIPE 67 DNS Working Group

Athens, October 2013

Dave Knight



Hedgehog is a replacement for the DSC presenter



Why replace DSC?

- L-root has used DSC for many years
- Recently L-root has grown, a lot
- DSC didn't cope with the increased scale
 - ▶ Presenting data from hundreds of sources
 - ▶ Huge amount of data to move around

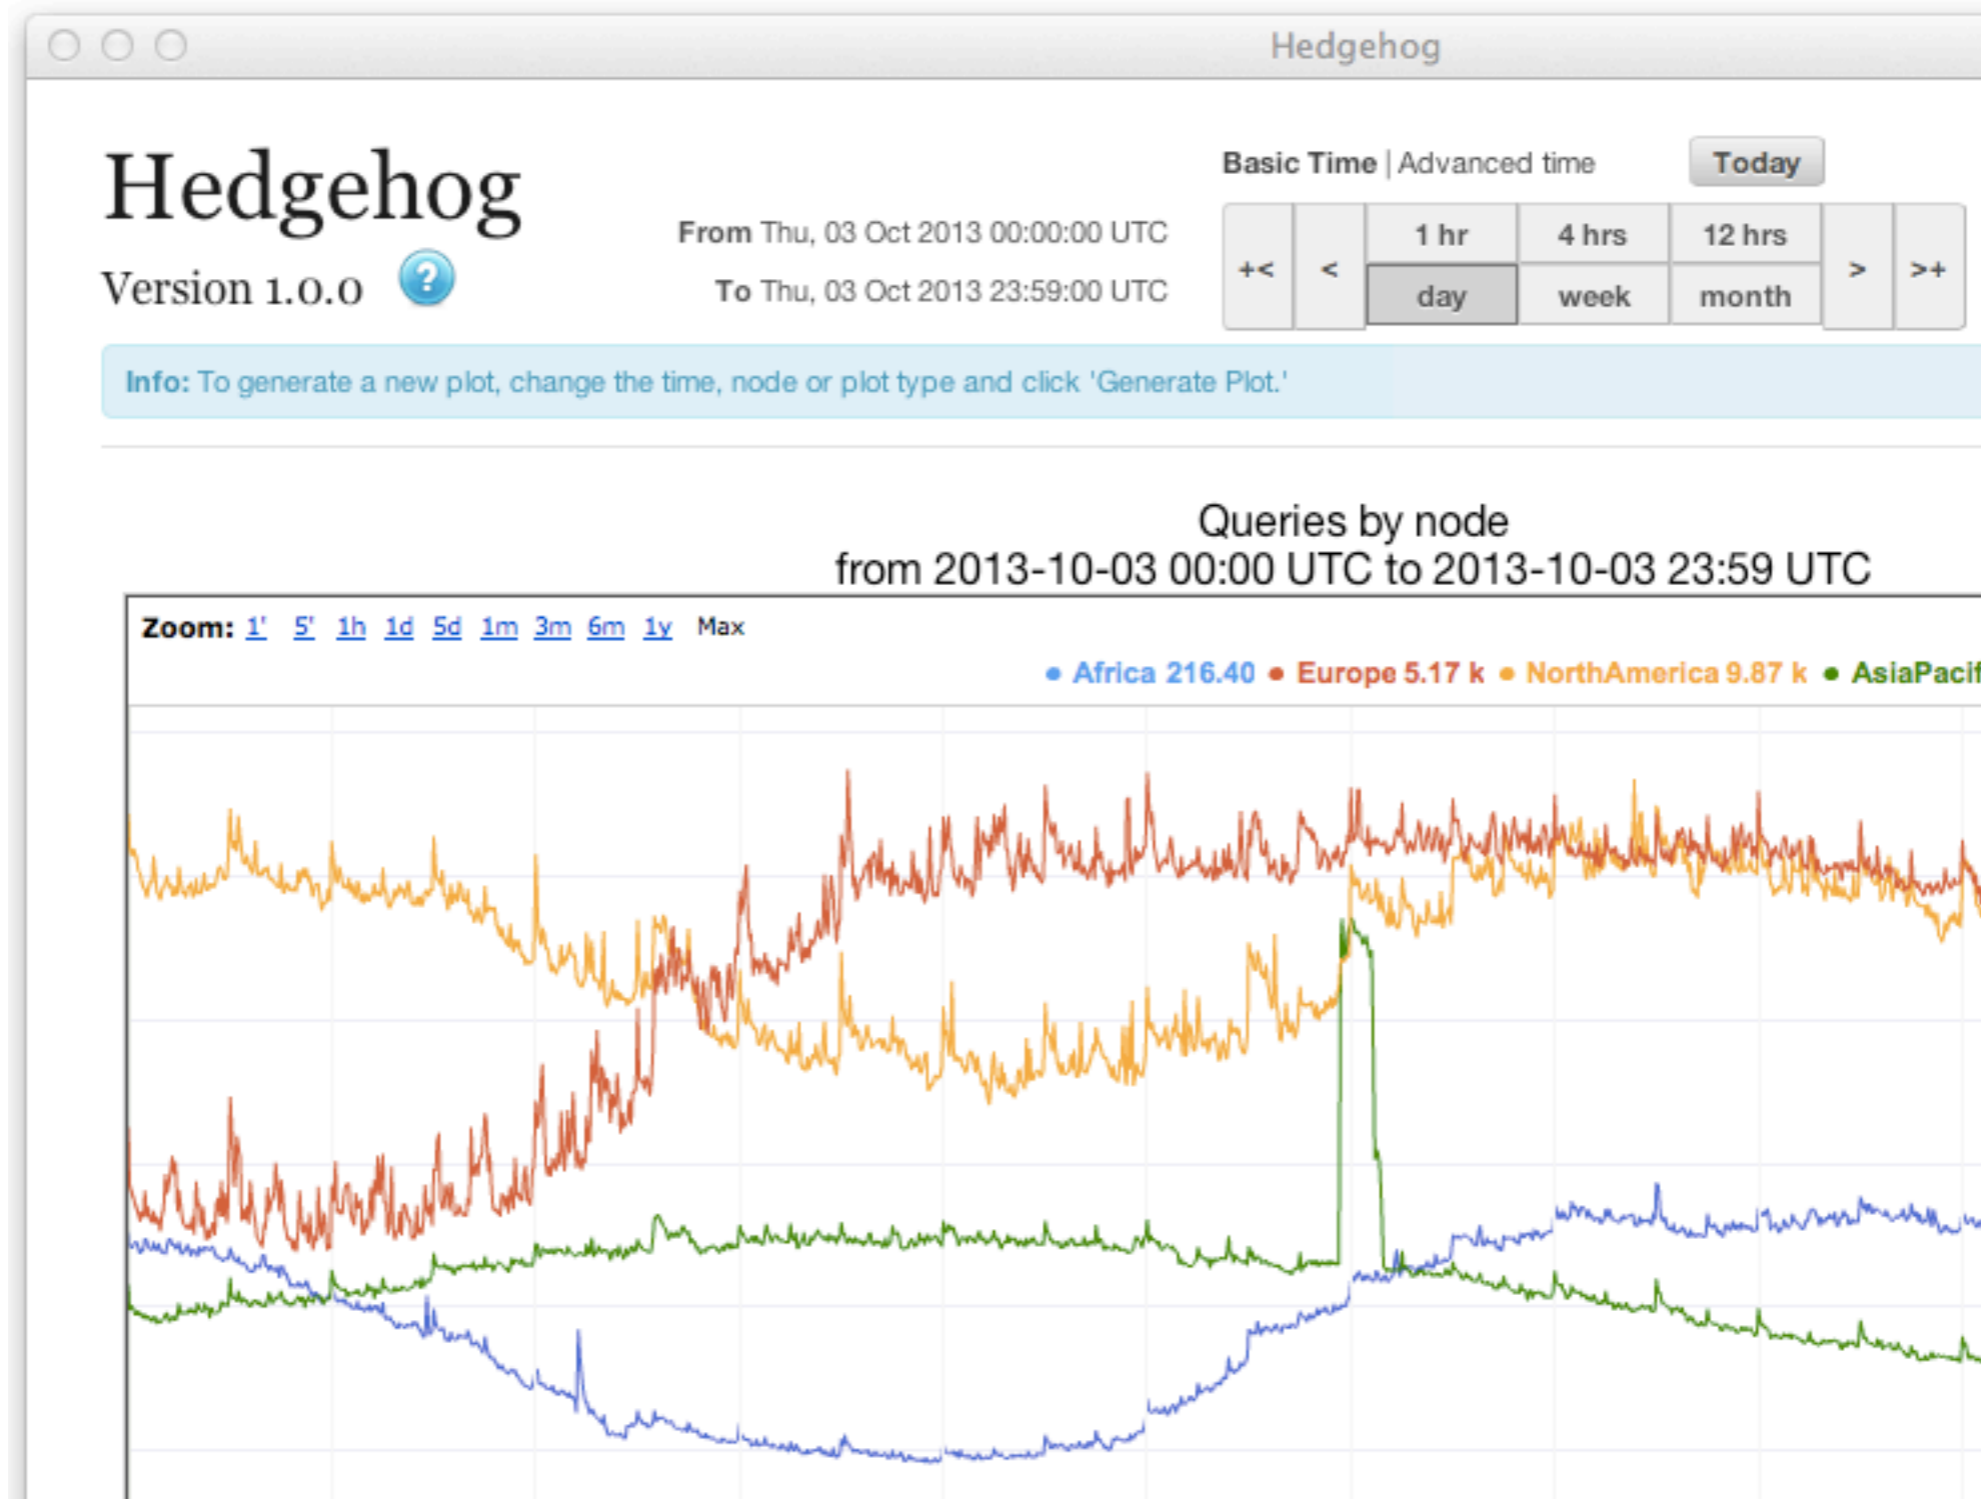
Why not an existing alternative?

- Ultimately we want more than a simple DSC replacement
 - ▶ Want to better understand the health of the service
 - ▶ Deeper inspection of queries
- Didn't seem to be a lot of active development in other projects
- The size of the L-root deployment presents unique challenges

Implementation

- Uses the DSC collector
 - ▶ XML data is uploaded using Apache and mod_dav_fs secured with SSL certificates
- Data importer written in C++ is a drop in replacement for DSC's dsc-extractor.pl
 - ▶ Importer adds the ability to XML or DAT to a Postgres DB whose tables are structured to be as similar as possible to the DAT files
- Web interface uses R, R-Apache, HTML5, javascript and jquery
 - ▶ R is a statistical language which will allow statistical analysis of the data

Live demo!



1.0.0

- Can import DSC XML, or DAT files
- Postgres DB
- Web GUI has interactive plots, and also static PNG plots for reporting
- Basic and Fine Grain time controls
- Built using R stats package for extensibility
- Web GUI optimized for large datasets
- Scales to >300 nodes
- Support for remote GUI access to DB
- Automated caching of daily `all node` plots

1.x

- Support for RSSAC plots and daily reports
- Improvements to node selection widget
- Cron job to cache week/month plots
- Ability to toggle keys
- Geo map by node
- Display metadata for each plot
 - ▶ Detailed description
 - ▶ Date cached
 - ▶ Plot URL
 - ▶ etc

2.0

- Alternative node selection in a directory check box style listing on the left hand side
- Graphical time picker for Advanced Time mode
- Option to export raw data
- Statistical analysis of data (attempt to measure health)

Future work

- More interesting data displays than line and bar charts
 - ▶ Heat maps
 - ▶ Word clouds
 - ▶ etc
- Ability to drill down into PCAP data
- Scaling improvements
 - ▶ Move data processing from the presenter out to the nodes
 - ▶ New data upload mechanism which is lighter than XML

Availability

Hedgehog for L-root is online at:

<http://hedgehog.dns.icann.org/>

ICANN intends to release Hedgehog as Free/Open Source Software. Licensing details are yet to be determined.

Questions?

terry.manderson@icann.org