

Whats New with DSC (since last time)

Duane Wessels
The Measurement Factory/CAIDA

DNS-OARC
June 4, 2008

The December 2007 Backlog

- In late December 2007 the OARC file server failed in some serious way(s) that kept it offline for a couple of weeks.
- This caused a big backlog on DSC collectors. Some ran out of local disk space.
- A backlog of this duration creates millions of small files in a single directory, which makes filesystems and people that run */bin/lis* unhappy.
- So this was a good reason and time to implement some long-overdue changes...

minfree_bytes

- The collector's *dsc.conf* now has a *minfree_bytes* setting.
- If non-zero, DSC will not write to disk if free space is less than *minfree_bytes*.
 - If this happens, DSC data will be lost until free space recovers.
- Typical setting might be 5 GBytes or so.

Single XML blob

- Previously, DSC wrote each dataset (every minute) to a separate XML file.
- \approx 20 datasets means 20 new XML files per minute.
- Now DSC writes a single, larger XML blob each minute.
- Collector is backward-compatible and accepts dataset files either way.

Date-Based Upload Directories

- Previously, the collector's *upload-prep.sh* script placed all to-be-uploaded XML files into one directory.
- Now it uses date-based subdirectories.
- For example:
`~dsc/run/sfo2/upload/collect.oarc.isc.org/2008-05-28/1211933520.dscdata.xml`
- This, combined with single XML blob, means that there won't be more than 1440 files per directory in case of a backlog.
- The old *upload-prep.sh* script was re-implemented in Perl, which adds Perl and Proc::PID::File as a dependency on the collector.

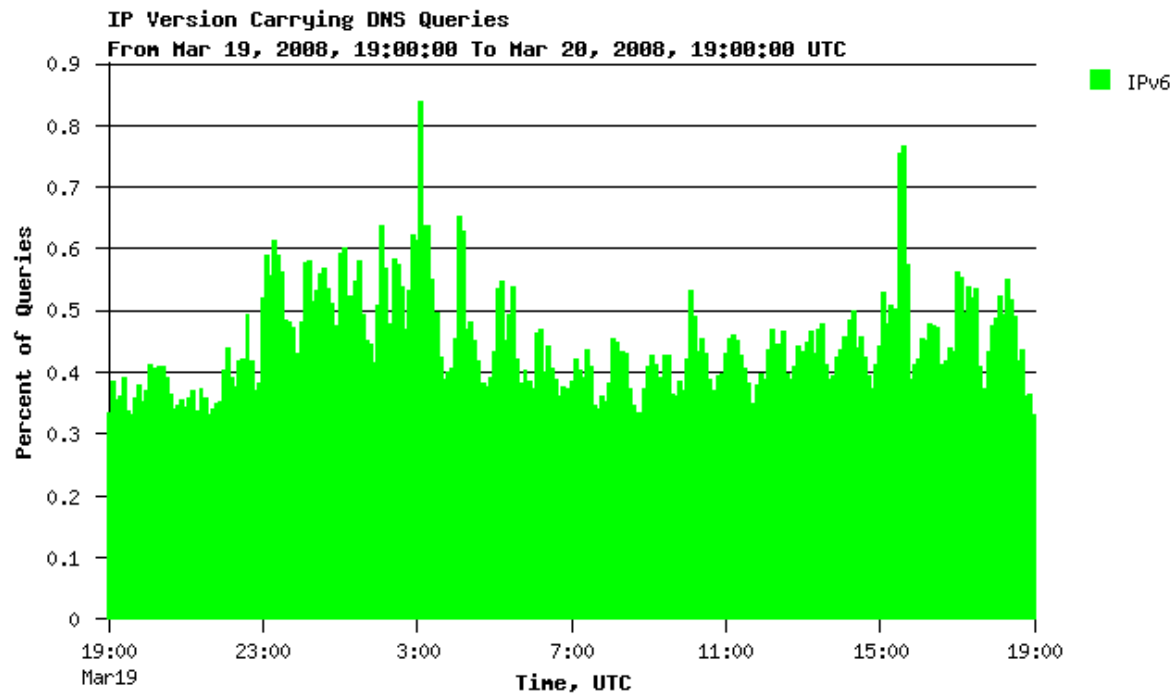
No More Forking

- The DSC collector process no longer forks at the end of each 60-second collection period.
- This change was necessary to correctly handle TCP sessions that span collection periods.

IP version of DNS message

- A new indexer reports whether DNS messages came in on IPv4 or IPv6.
- For example:

```
dataset dns_ip_version_vs_qtype dns IPVersion:dns_ip_version Qtype:qtype queries-only;
```




NCAP

- DSC can now read NCAP streams.
- This work was done to help visualize ISC's SIE data.
- But SIE uses a store-and-forward (rather than streaming) delivery technique, which confuses DSC.
- Additional changes will be necessary to fully support SIE.


hide_nodes

- New presenter feature allows you to “hide” the list of nodes within a server.

Normal



Hidden



./configure

- The collector code now has a ./configure script which should improve portability.
- Please report any portability issues so that they can be resolved.

Please Upgrade!

- If you are running a DSC collector for OARC, please make some time to upgrade to the latest version.
- The upgrade will help both you and OARC in case there is another backlog.
- Free hand-holding available.

The End

<http://dns.measurement-factory.com/tools/dsc/>