

---

# 2007 Day In The Life Mistakes We Don't Want To Repeat

Duane Wessels  
The Measurement Factory/CAIDA

WIDE+CAIDA Workshop #9  
January 20, 2008

---

# OARC Disk Space

## Problem:

- The OARC server did not have enough disk space to store all the data it should have received.
- FYI: DITL-2007 is about 800 GB, mostly from 4 roots.

## Solution:

- More disk space
- New file server

---

# Bandwidth Bottleneck or Local Disk Shortage

Problem:

- Some collector sites cannot upload pcap files fast enough and run out of local disk space.

Solution:

- Use an intermediate storage node?
- Operated by whom?

---

## Senders Can't List or Delete

Problem:

- Once data has been uploaded, senders cannot list, read, or remove the file on the OARC server.

Solution:

- There was talk of providing access via sftp?

---

# Partial Uploads

Problem:

- Partial uploads were hard to detect because the receiver did not use temporary file names.

Solution:

- This bug has been fixed. Files will be given their final name only if the upload succeeds.

---

## No Metadata

Problem:

- OARC upload interface does not provide a way to also send metadata. Senders must write down or otherwise remember metadata details until data is indexed in DatCat.

Solution:

---

## Varying pcap Size and Boundaries

Problem:

- Senders use different time limits on pcap files (1 hour, 5 minutes, variable)
- File start/stop times based on when the program started.

Solution:

- Use *dnscap* with one hour time limit.
- New file will be started at the top of each hour.

---

# Clock Skew

Problem:

- Queries containing timestamps uncovered clock skew issues.

Solution:

- Send skew-detecting queries to all server unicast addresses?
- What about load-balanced nodes?
- Convince operators to make their NTPs readable?



---

# Truncated Packets

Problem:

- Some folks forgot '-s 0' on *tcpdump* or used '-s 1500'.

Solution:

- Use *dnscap* instead.

---

## Extra Data

Problem:

- Some pcap files contained “extra data” – i.e., packets from nearby DNS servers using the same infrastructure.

Solution:

- Perhaps we should normalize pcap files after they have been received.

---

## VLAN tags

Problem:

- Some pcap files were captured from VLAN-enabled networks and create a minor nuisance for pcap readers and bpf filters.

Solution:

- Perhaps we should normalize pcap files after they have been received.

The End