ARIN
American Registry for Internet Numbers

aflow – a Netflow Packet Analyzer

Dan Ardelean, ARIN, 2004

Quick Netflow Overview

- **developed at Cisco starting in 1996**
- **primary network accounting technology**
- **Netflow Version 9 emerging as an IETF standard**
  - work in IPFIX
  - various goals:
    - Accountability
    - Network monitoring
    - DDOS and DOS real time detection
- **What is a flow ?**
  - defined by several fields in a packet - together they form a primary key

aflow Overview

- **aflow ? what does it mean ?**
  - aflow = ARIN Flow
  - a tool for analyzing Netflows
- **Why aflow ?**
  - we need to:
    - understand network services' and applications' behavior
    - improve network performance
    - early detect and understand network events
  - the answer:
    - a tool for analyzing the network (netflows)
- **download from:**
  - http://www.aflow.org

aflow Overview

- **aflow major goals:**
  - ability to easily compile and install
  - powerful configuration
  - flow analysis based on filters and classification
- **aflow main tasks:**
  - collect Netflow UDP packets
  - classify each flow
  - store statistical information

aflow philosophy

- **defines the following constructs:**
  - datapoints
    - containers that store data over time
    - stored in memory as arrays of floating point numbers
    - characteristics:
      - datapoint type
      - datapoint filter
        - a logic infix expression (tcpdump or pf style)
        - used to classify a flow

aflow philosophy

- collections
  - group datapoints together
  - define common characteristics:
    - buffering (how many data generations we keep in memory before dumping to disk)
    - database file information
    - sample rate – the distance between two consecutive instances of a datapoint
    - common filter (in the works)

aflow config – classification definition

```
collection example {
        file “database_name.rrd”;
        filter dst 192.168.1.1;
        datapoint x {
                type bytes;
                filter proto 6 and srcport 22;
        }
        datapoint y {
                type packets;
                filter proto 6 and srcport 25;
        }
}
```

aflow config – graph definition

```
graph g1 {
        file “~/www/ex.cgi”
        variable var1 {
                collection example; datapoint x;
                color “green”;
                legend “ssh”;
                form “area”;
        }
        variable var2 {
                collection example; datapoint y;
                color “red”;
                legend “smtp”;
                form “line”;
        }
}
```

aflow example

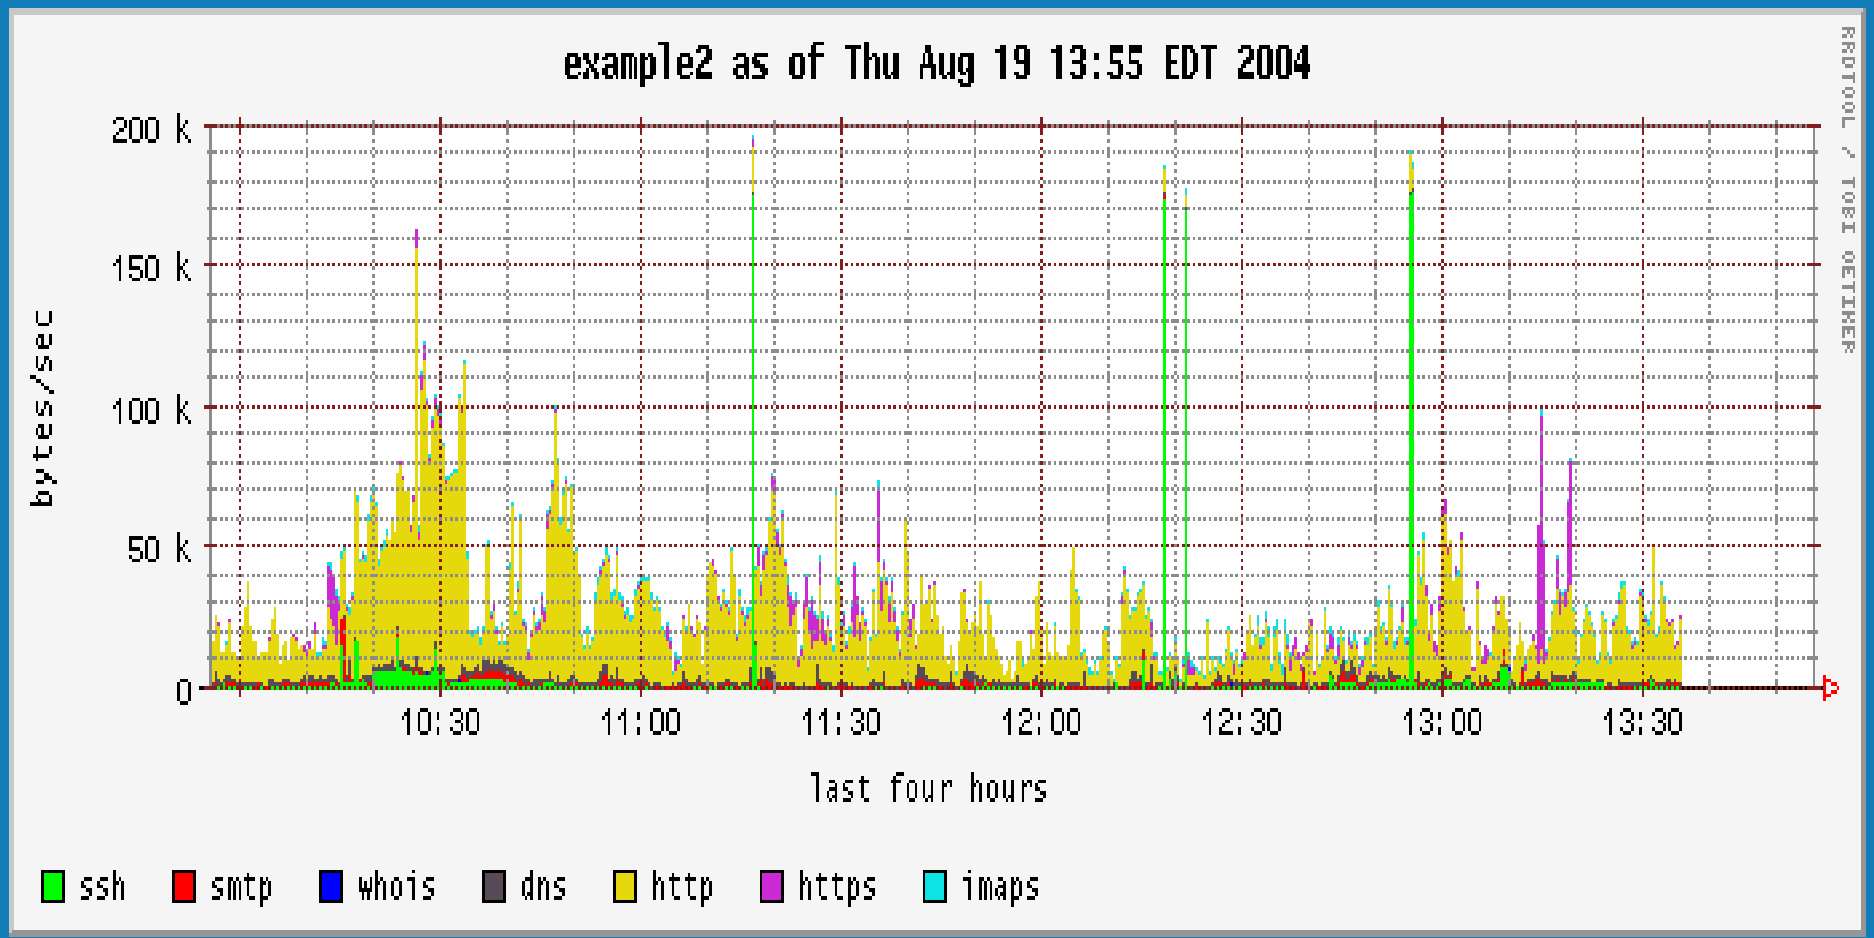

aflow features

- **real-time:**
  - capture
  - classification
  - processing
- **simple, but powerful config language:**
  - filter capability
  - organize data in collections and datapoints
- **a platform for network analysis**
  - based on extendable modules

aflow plans

- near real-time traffic information client/server
- build module for traffic pattern learning and real time alerting (based on learning filters)
- build a raw packet collector for aflow
- build a Web GUI for generating the configuration file
- extend support for all Netflow versions

So what? Don't we have cflowd and flowscan?

- **installation:**
  - aflow: one project, autoconf/automake based
  - cflowd/flowscan: collection/analysis separate (multiple dependencies)
- **functionality:**
  - aflow: flexible configuration language
  - flowscan: aggregates
- **data collection:**
  - aflow: data not captured in the classification is thrown away
  - cflowd: dumps netflow data in files that are examined at a later point in time

Summary

- **aflow:**
  - captures Netflow UDP packets
  - classifies the information in collections and datapoints based on infix filter expressions
  - generates cgis for displaying data
  - a platform for Netflow analysis
- **Download location**
  - http://www.aflow.org
- **Discussion & Announcements List:**
  - aflow@arin.net
  - “subscribe aflow” in body to majordomo@arin.net

aflow

Thank you !!!

http://www.aflow.org