INTERNET GARBAGE STORAGE, ACCESS AND ANALYSIS

Alistair King
alistair@caida.org
CAIDA, UC San Diego

CAIDA www.caida.org

The Cooperative Association for Internet Data Analysis

- Independent analysis and research group
- Based at UC San Diego's San Diego Supercomputer Center
- Provide insights into Internet infrastructure, behavior, usage, and evolution
- Foster a collaborative environment in which data can be acquired, analyzed, and shared

CAIDA

Three primary areas of focus

RESEARCH + DATA SHARING + INFRASTRUCTURE

caida

RESEARCH

- **Topology Analysis**
  - Internet-scale router alias resolution
  - Comparing IPv4 and IPv6 topology
- **Security and Stability**
  - Large-scale Internet outages
  - Botnet activity
- **Internet Peering Analysis**
  - Inferring AS relationships
  - AS ranking
- **Modeling Complex Networks**
  - using hidden metric spaces
- **Interconnection Economics**
  - Modeling peering strategies
  - Transit pricing
- **Geolocation Analysis**
  - Comparing geolocation services
  - IP reputation vs. governance
- **Future Internet**
  - IPv6
  - Named Data Networking
- **Visualization**

caida

DATA SHARING caida.org/data/overview

Making data available to the community

- **Performance**
  - DNS root, gTLD RTT Data
- **Security**
  - Computer worms, backscatter, RSDoS attacks, Botnet scans
- **Topology**
  - AS Links, Prefix to AS, AS Rank, AS Relationships, IPv4 + IPv6 topology, Internet Topology Data Kit (ITDK)
- **Traffic**
  - Historical Telescope data, Live Telescope data, Anonymized Internet traces, Tier 1 packet traces
- **Meta-Data**
  - DatCat (datcat.org)

caida

INFRASTRUCTURE

Collecting the data

- **Archipelago (caida.org/projects/ark)**
  - Active measurement infrastructure
  - Supports ongoing topology measurement as well as customized experiments
- **Passive Trace Capture**
  - Captures two-way traffic on Tier 1 10GE backbone link
  - Shared anonymized headers only
- **UCSD Network Telescope**

caida

NETWORK TELESCOPE

TELESCOPE

caida.org/projects/network_telescope

The UCSD Network Telescope

- Passive traffic monitoring system
- Globally routed, lightly utilized **/8 network** (1/256 of the entire IPv4 address space)
- **24/7 full packet traces**
- **Archive of pcap data back to 2003 (sampled data prior to 2008)**
  - > 200TB currently
  - Stable growth by ~4TB per month

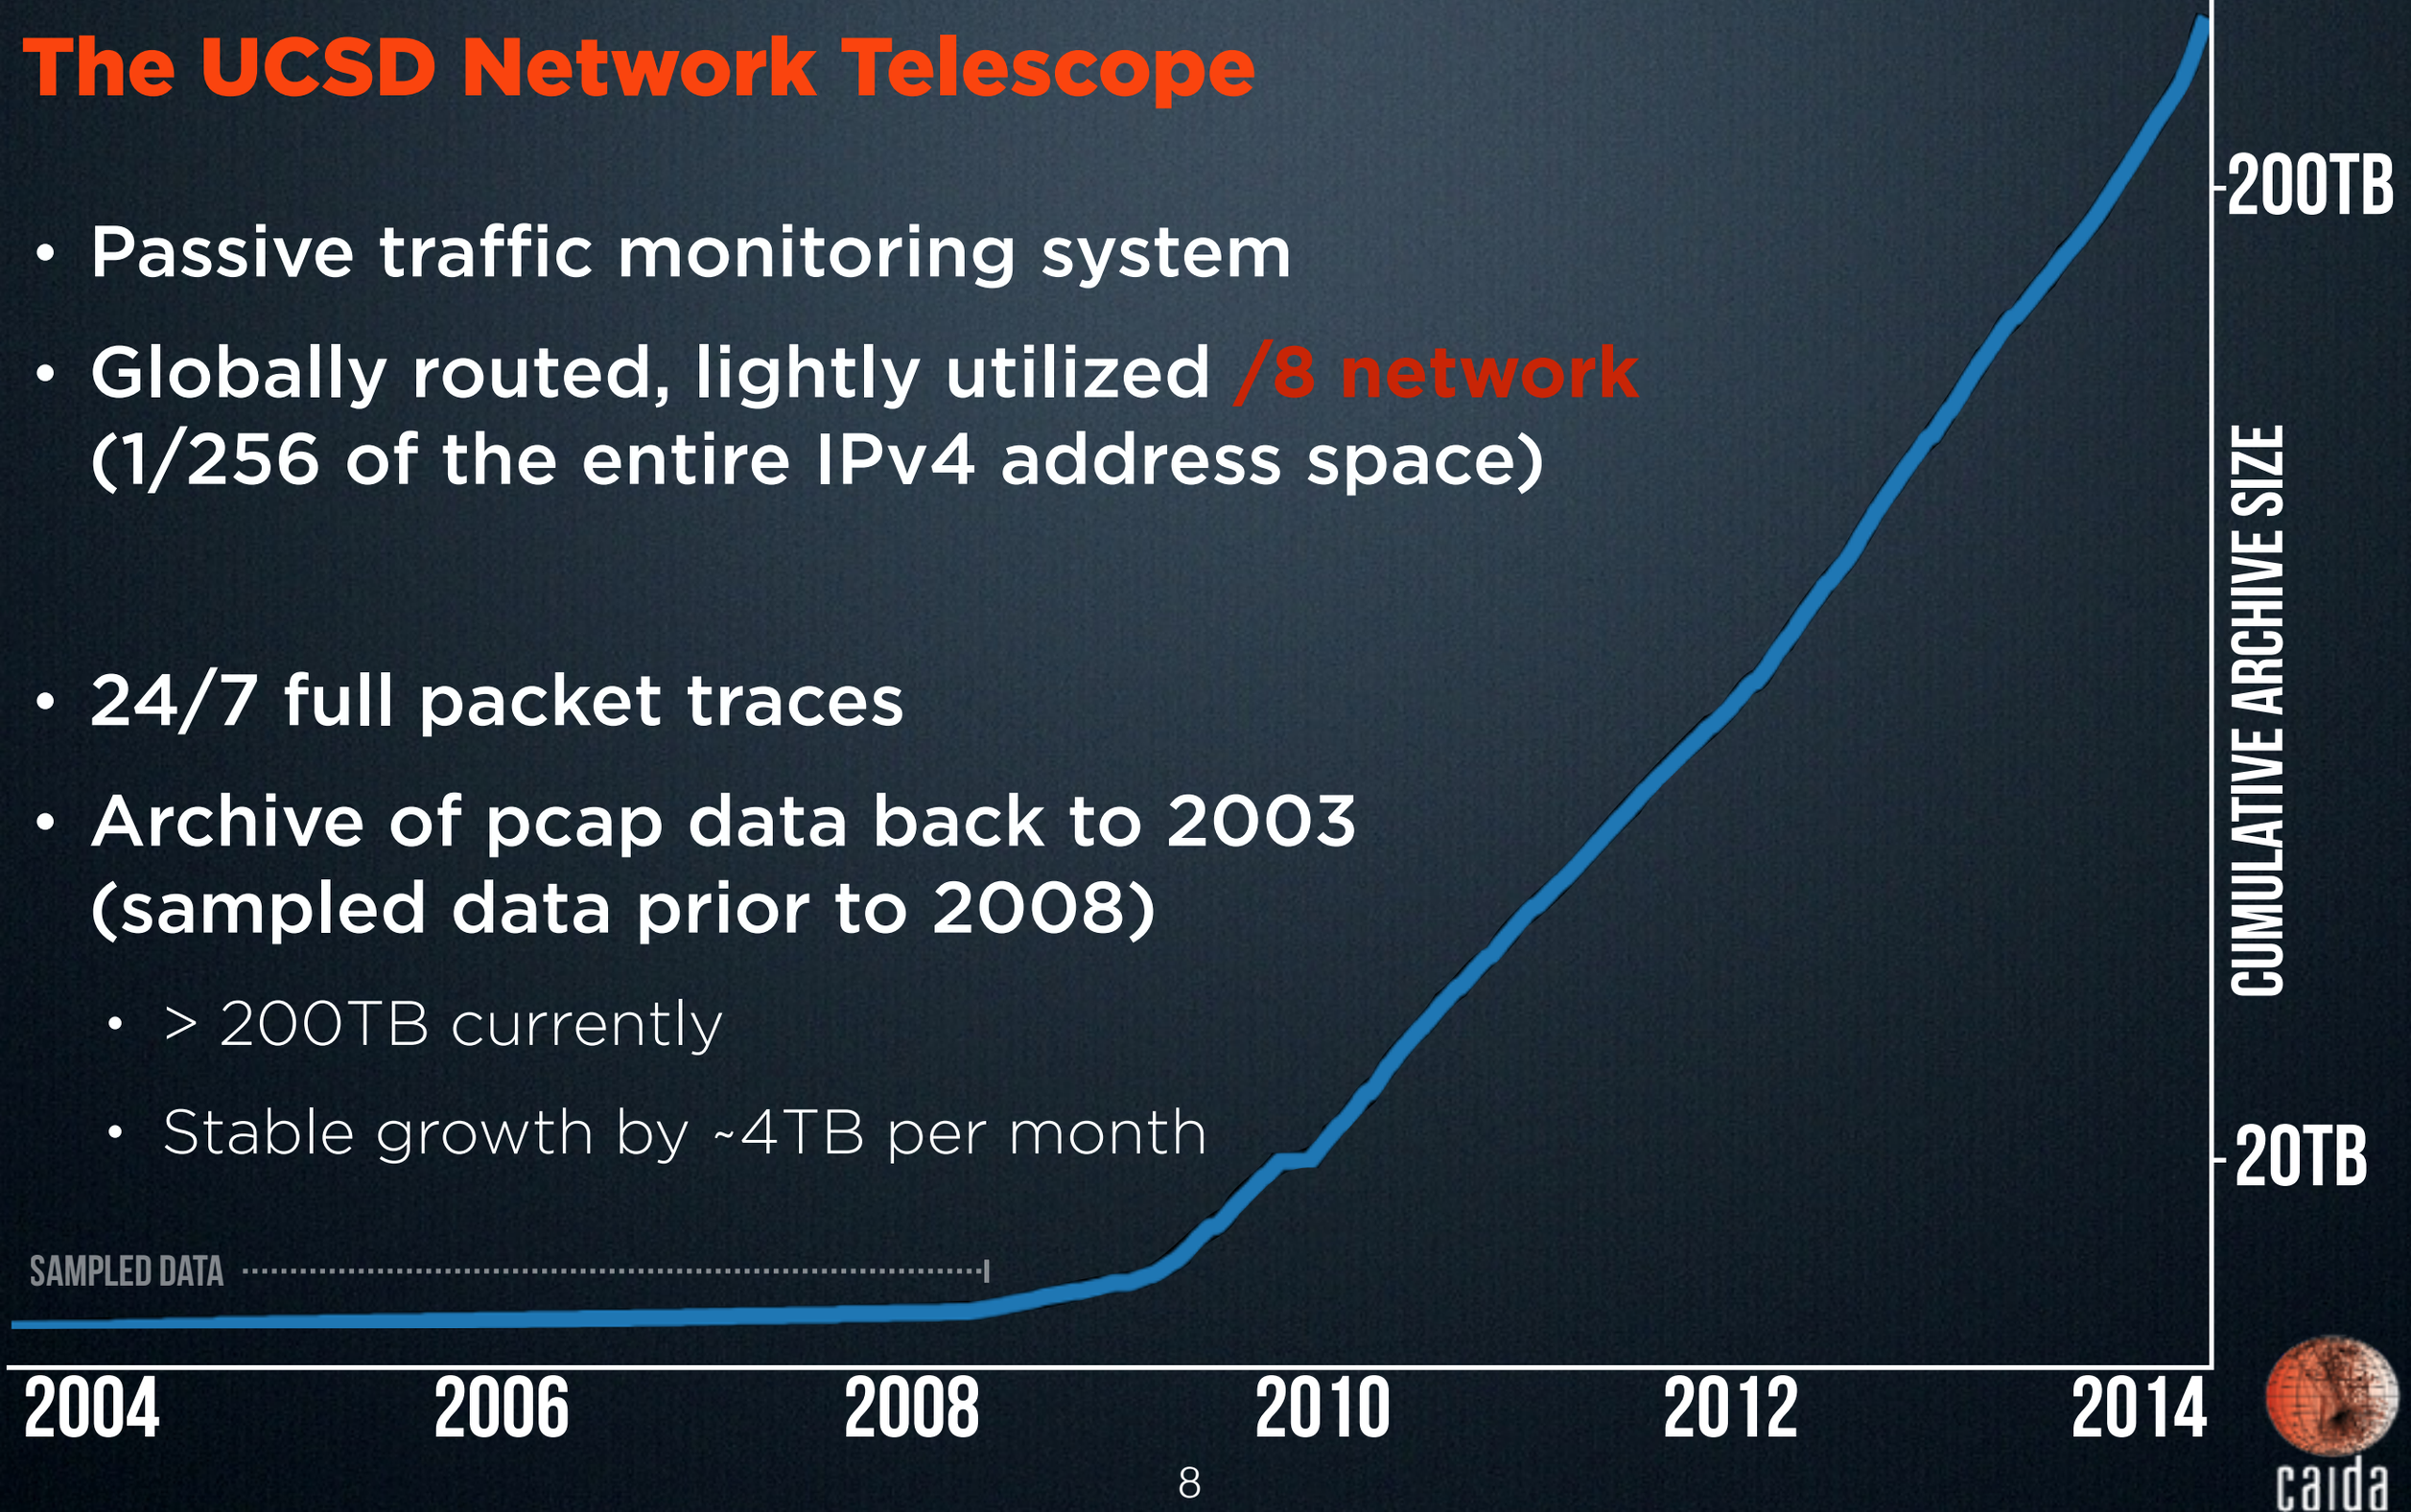

caida

HOW DOES IT WORK?

Who would send traffic to an unused network?

- Malware attempting to propagate
- Backscatter from spoofed DoS attacks
- Misconfigurations
- Network scans
- …

caida

WHAT IS IT GOOD FOR?

And what can this data be used to study?

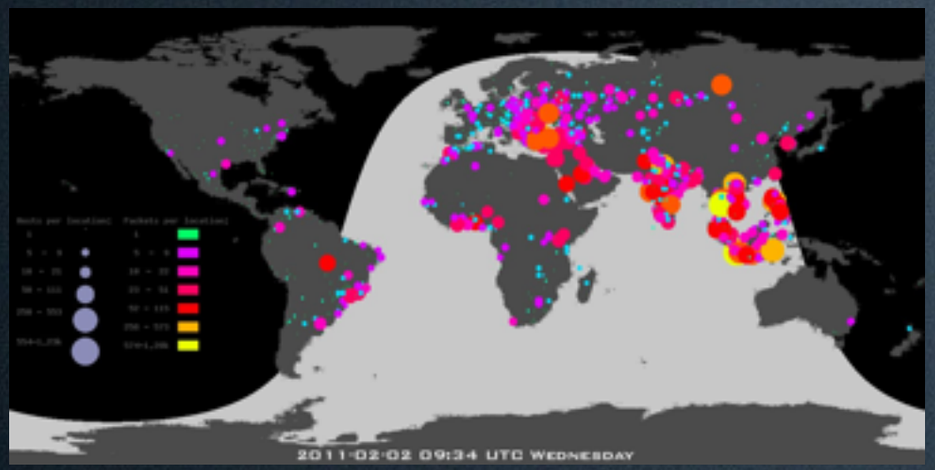

Malware Phenomena

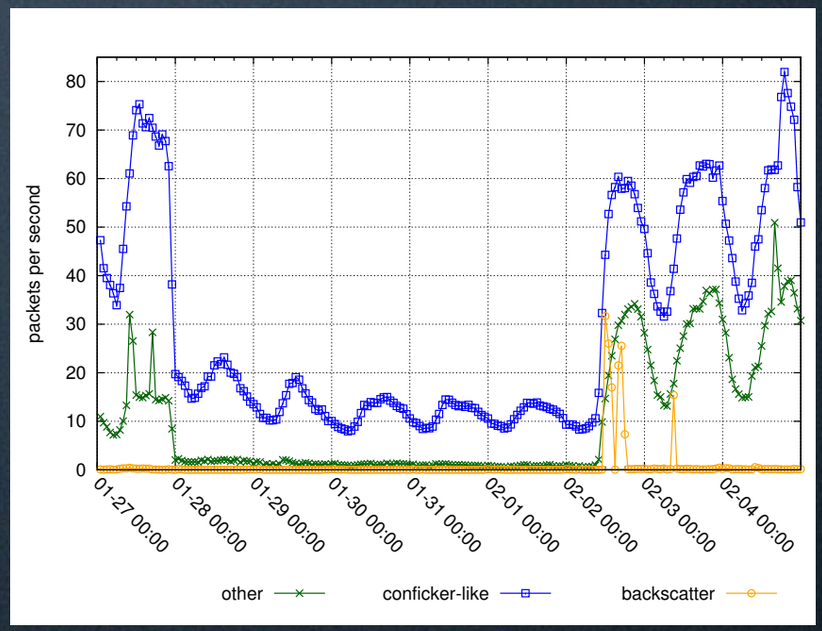

Connectivity Disruptions

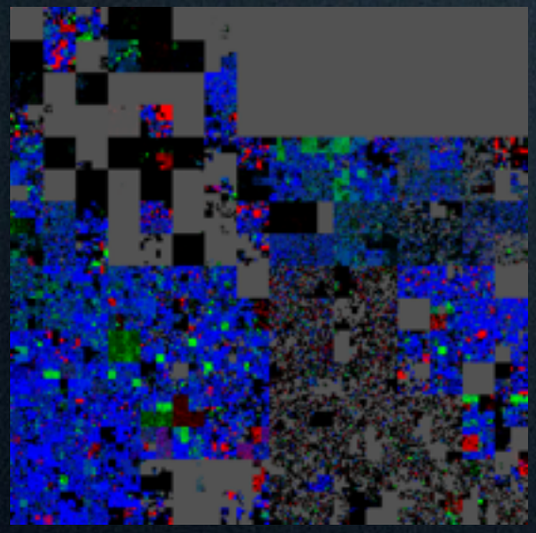

IPv4 address space usage

... and much more

caida

DATA HIERARCHY

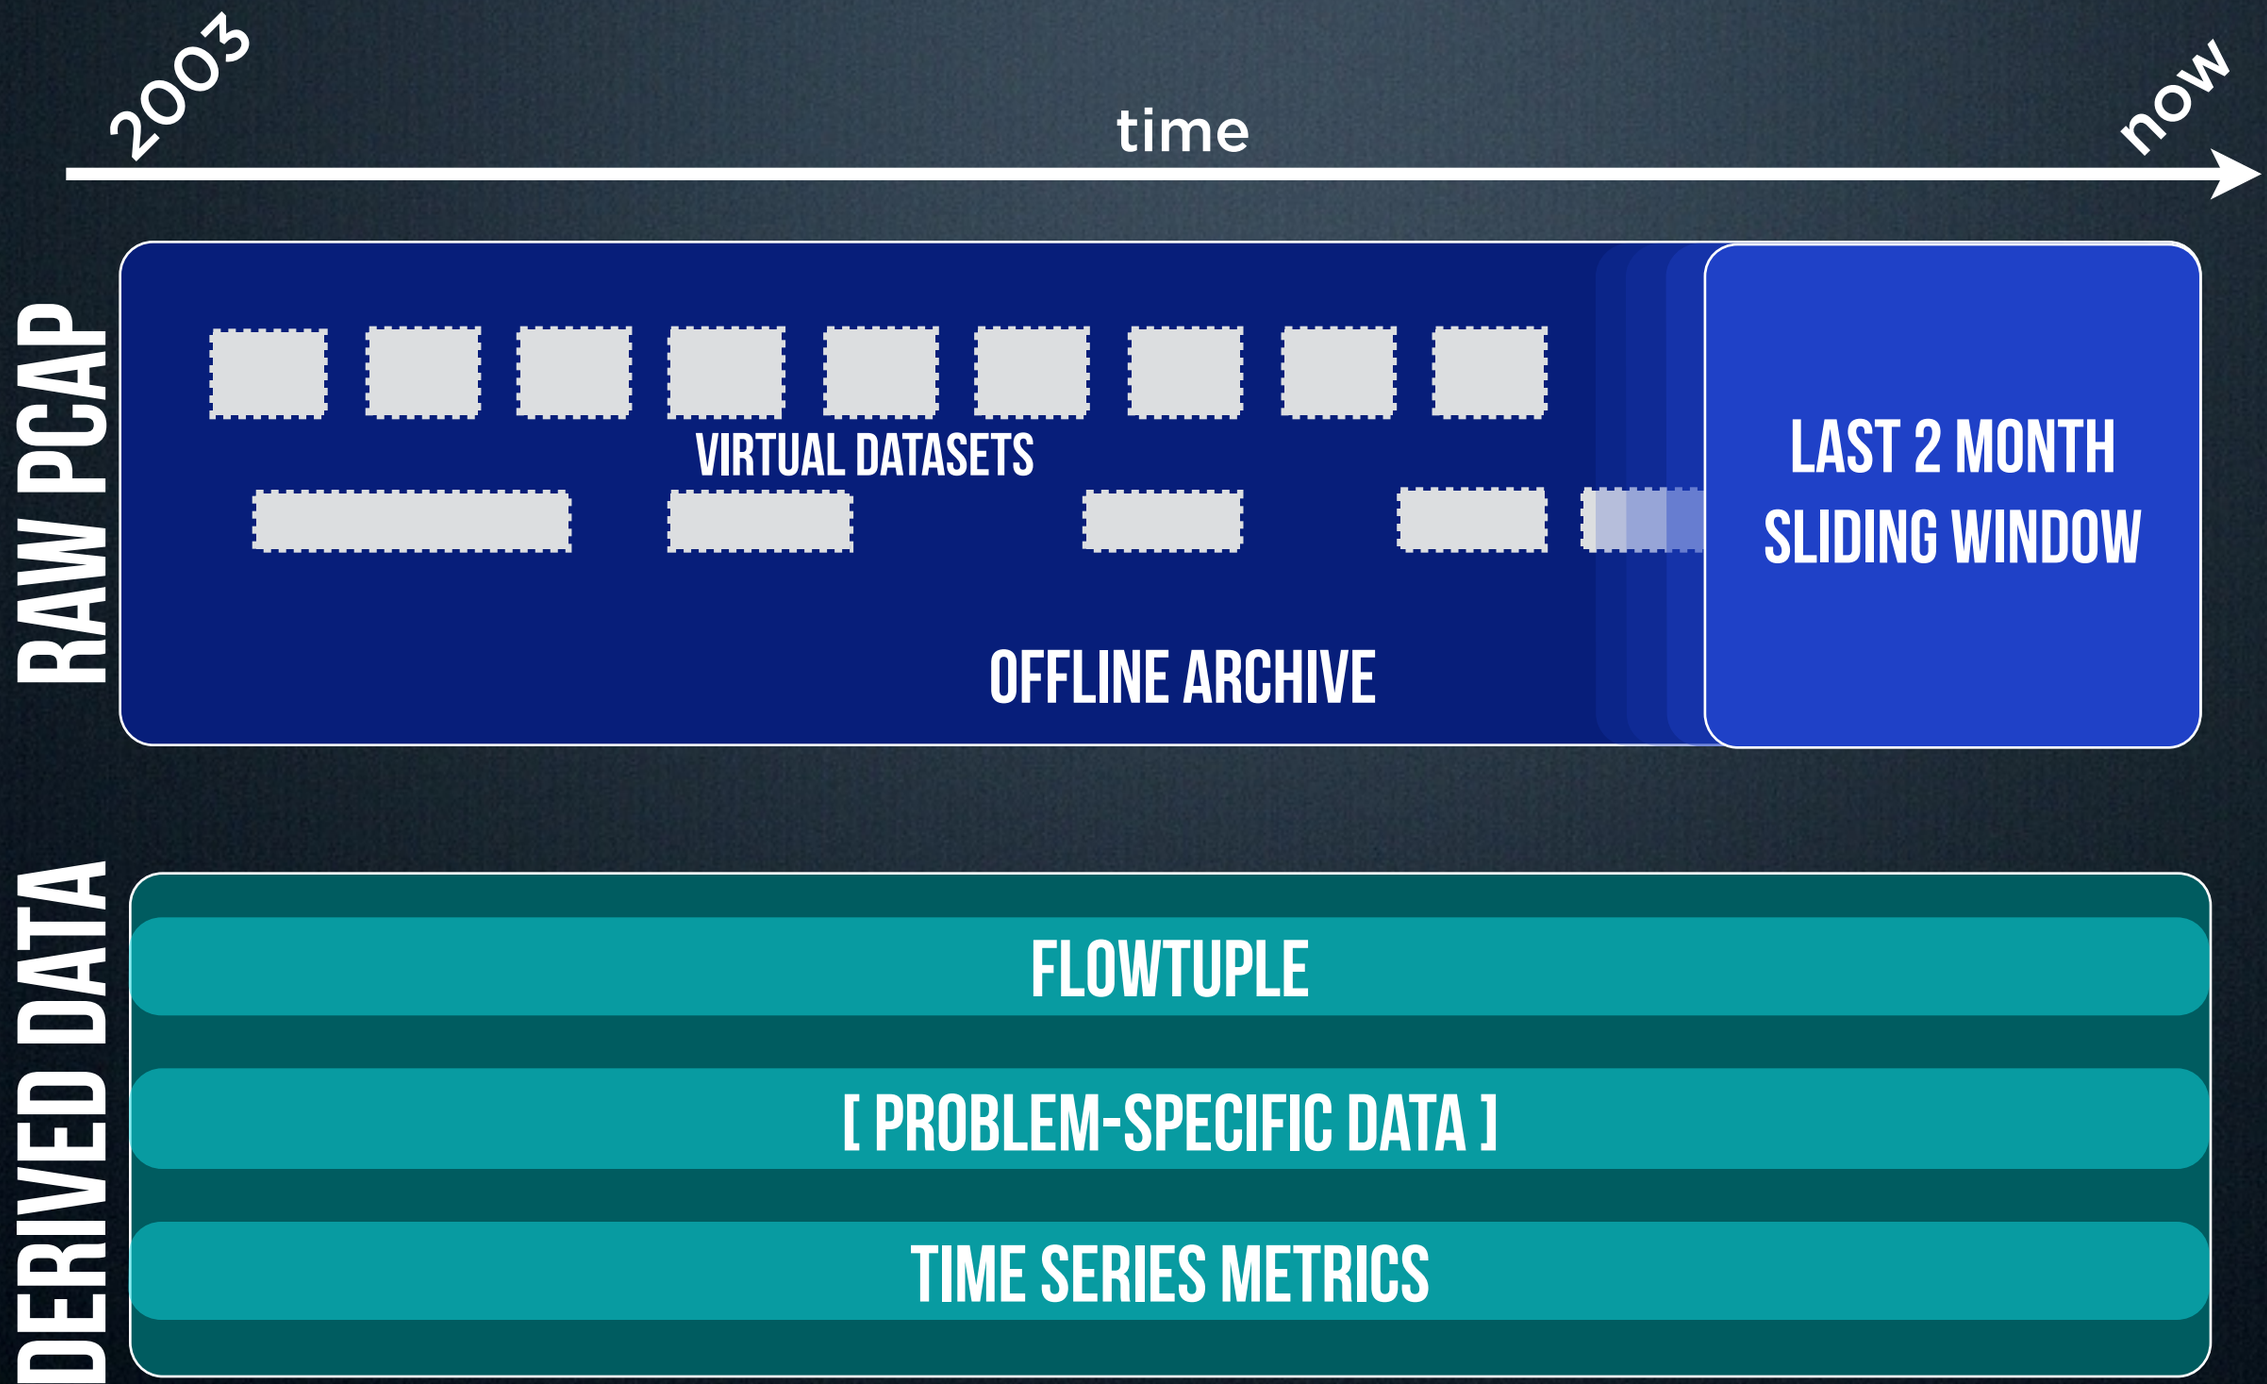

caida

DATA HIERARCHY

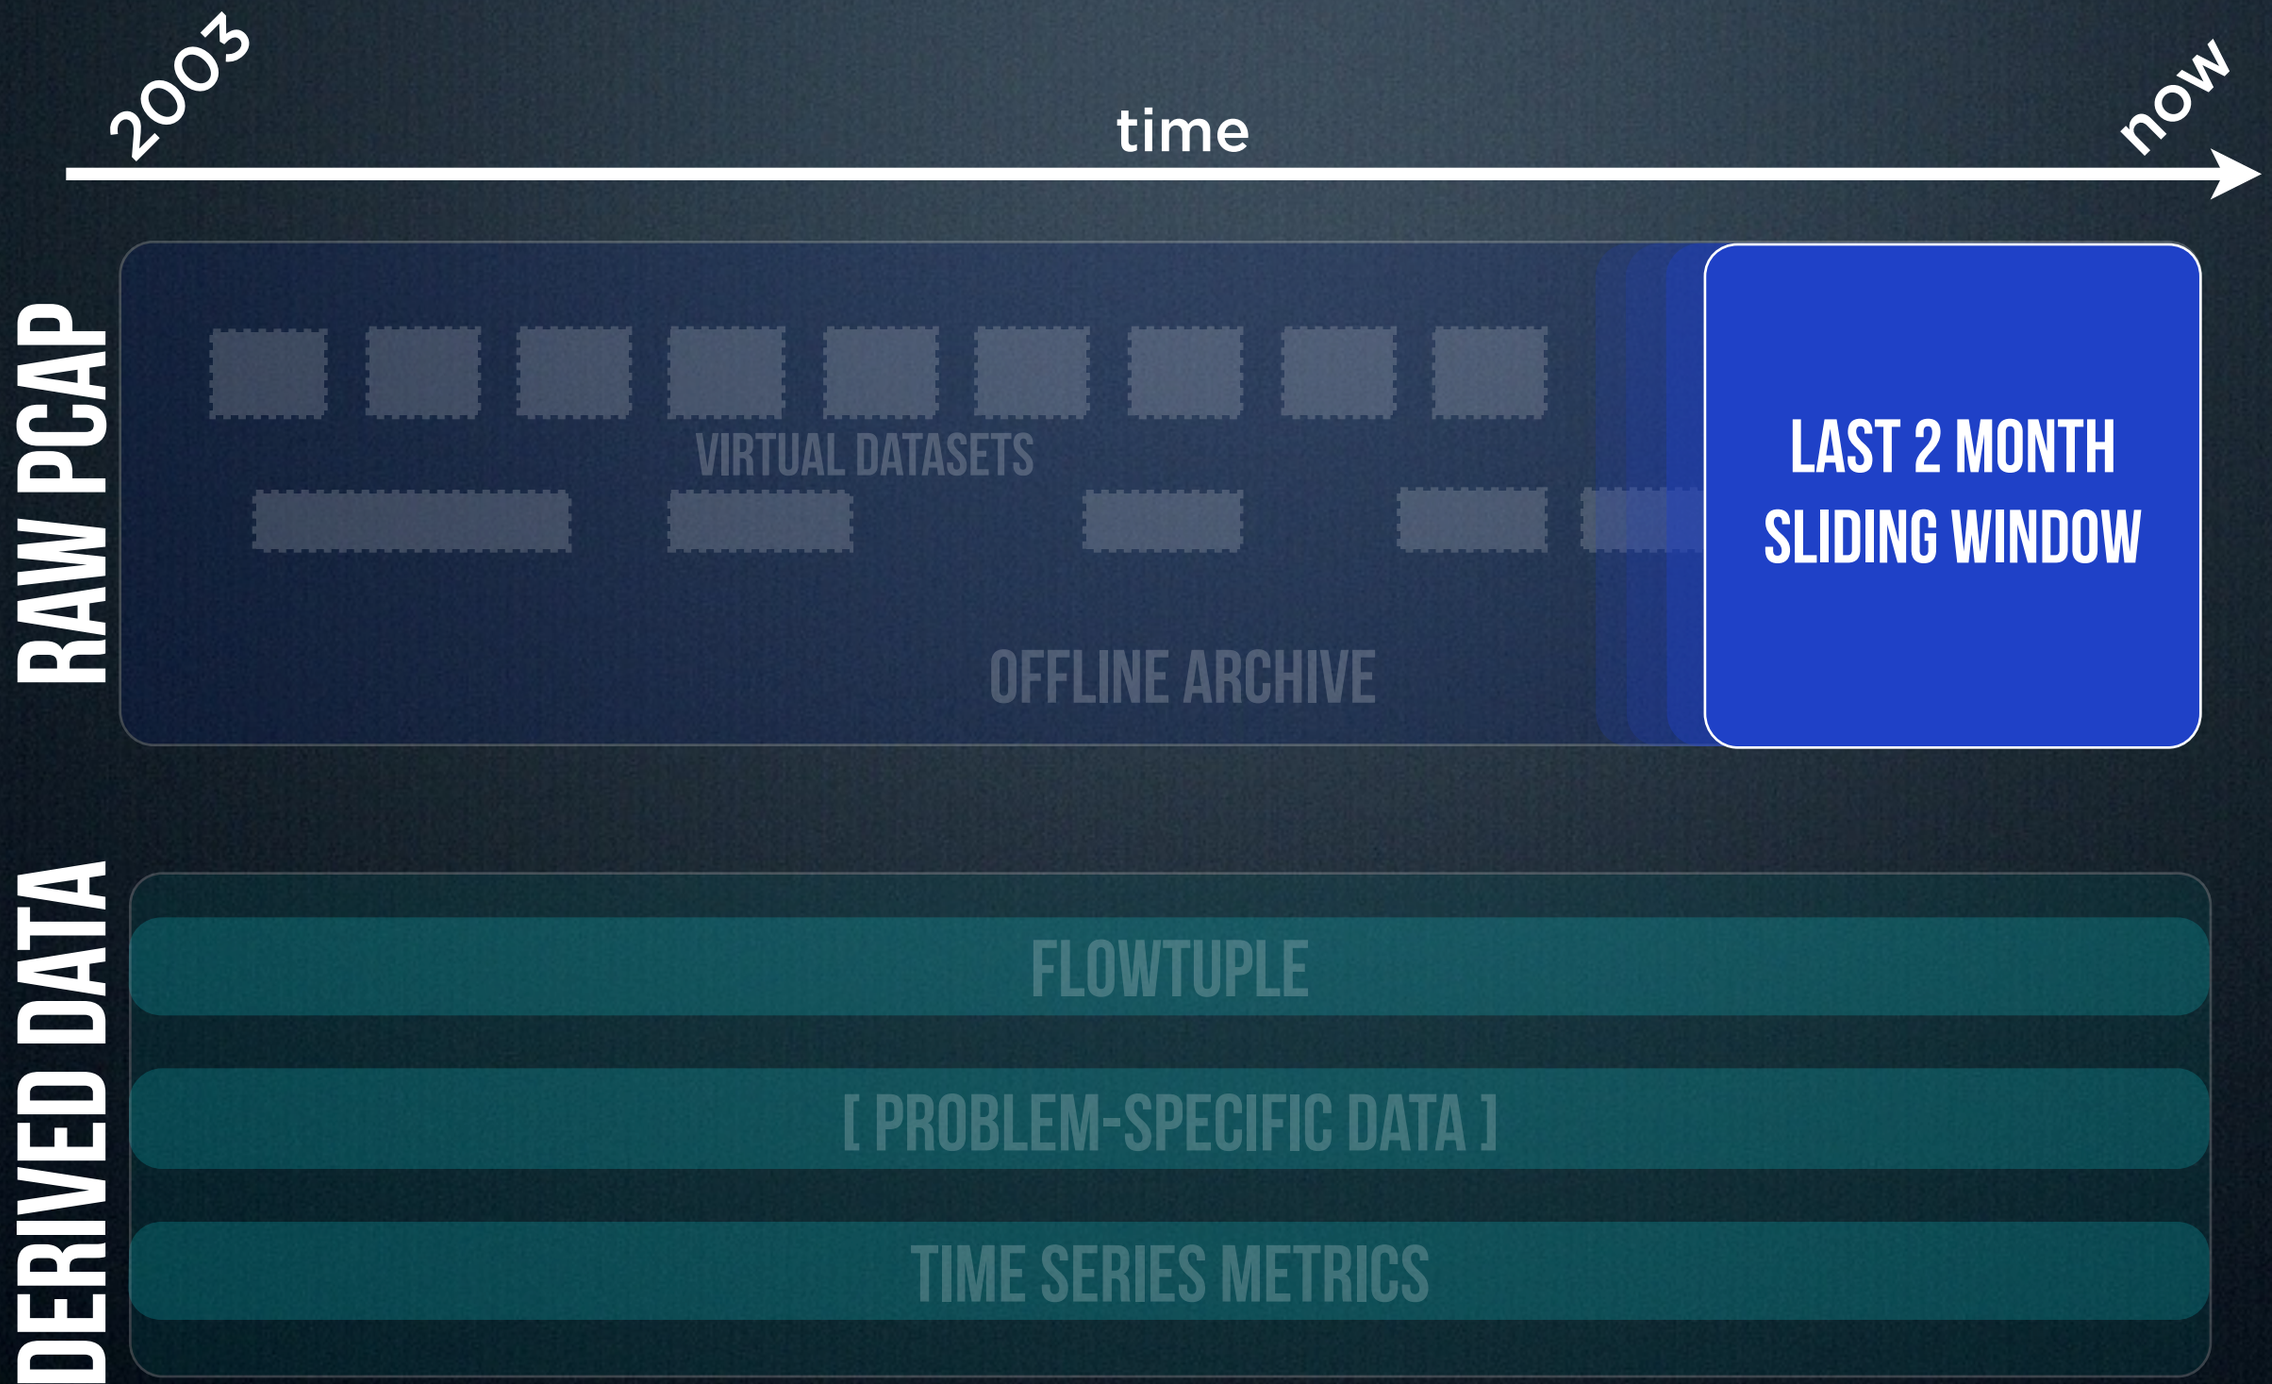

caida

DATA HIERARCHY

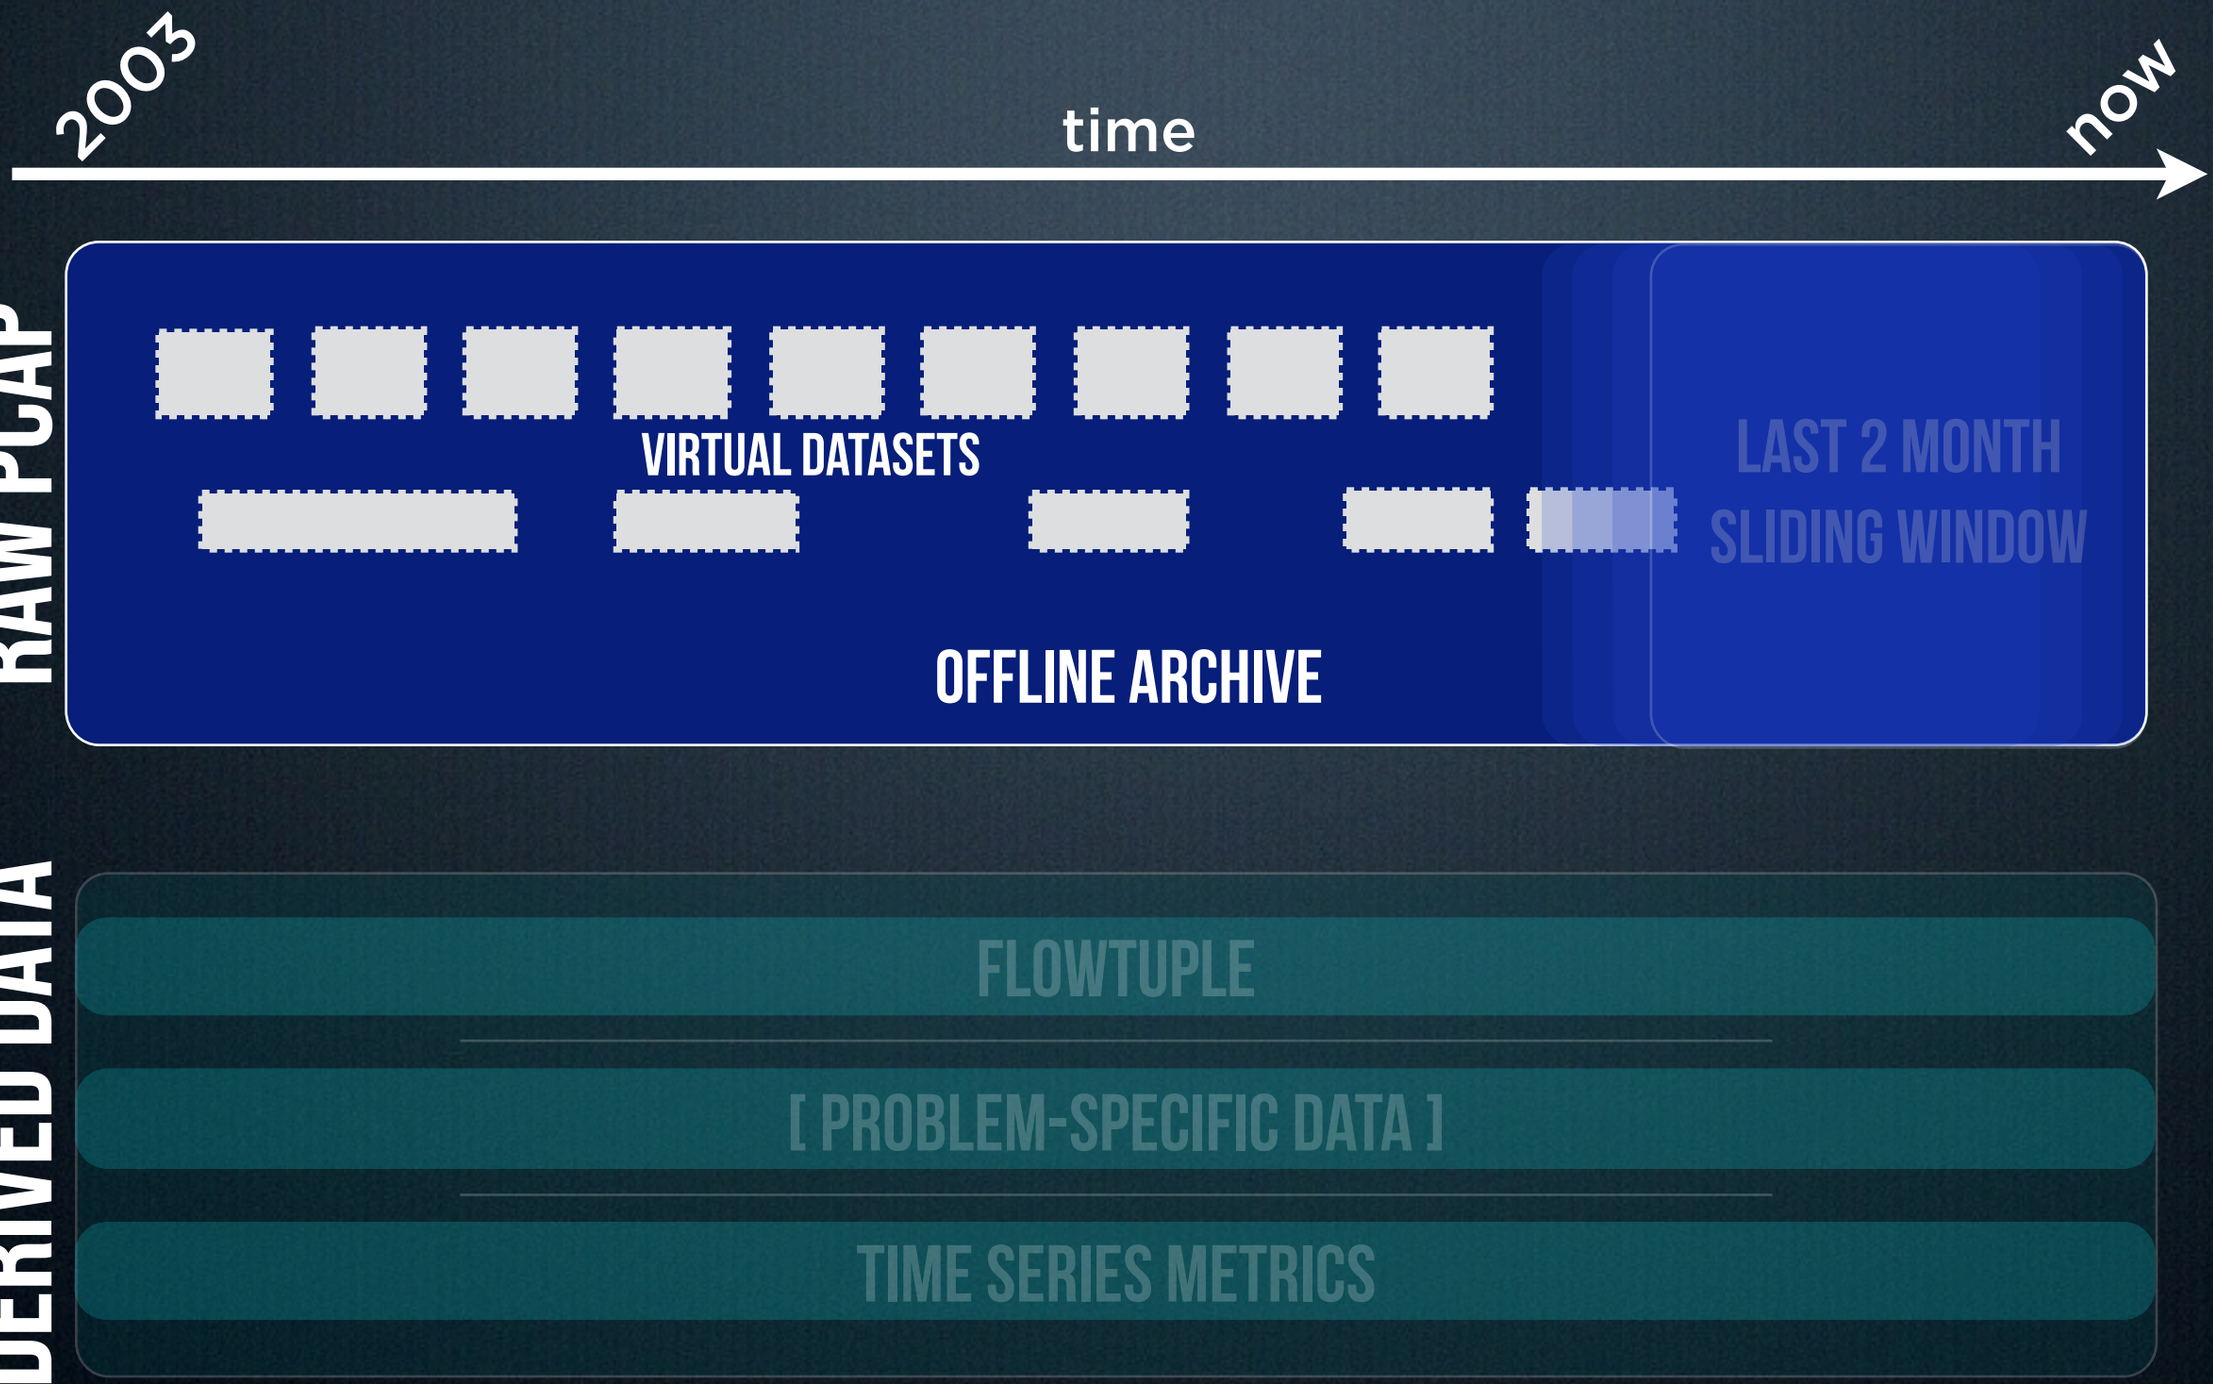

caida

DATA HIERARCHY

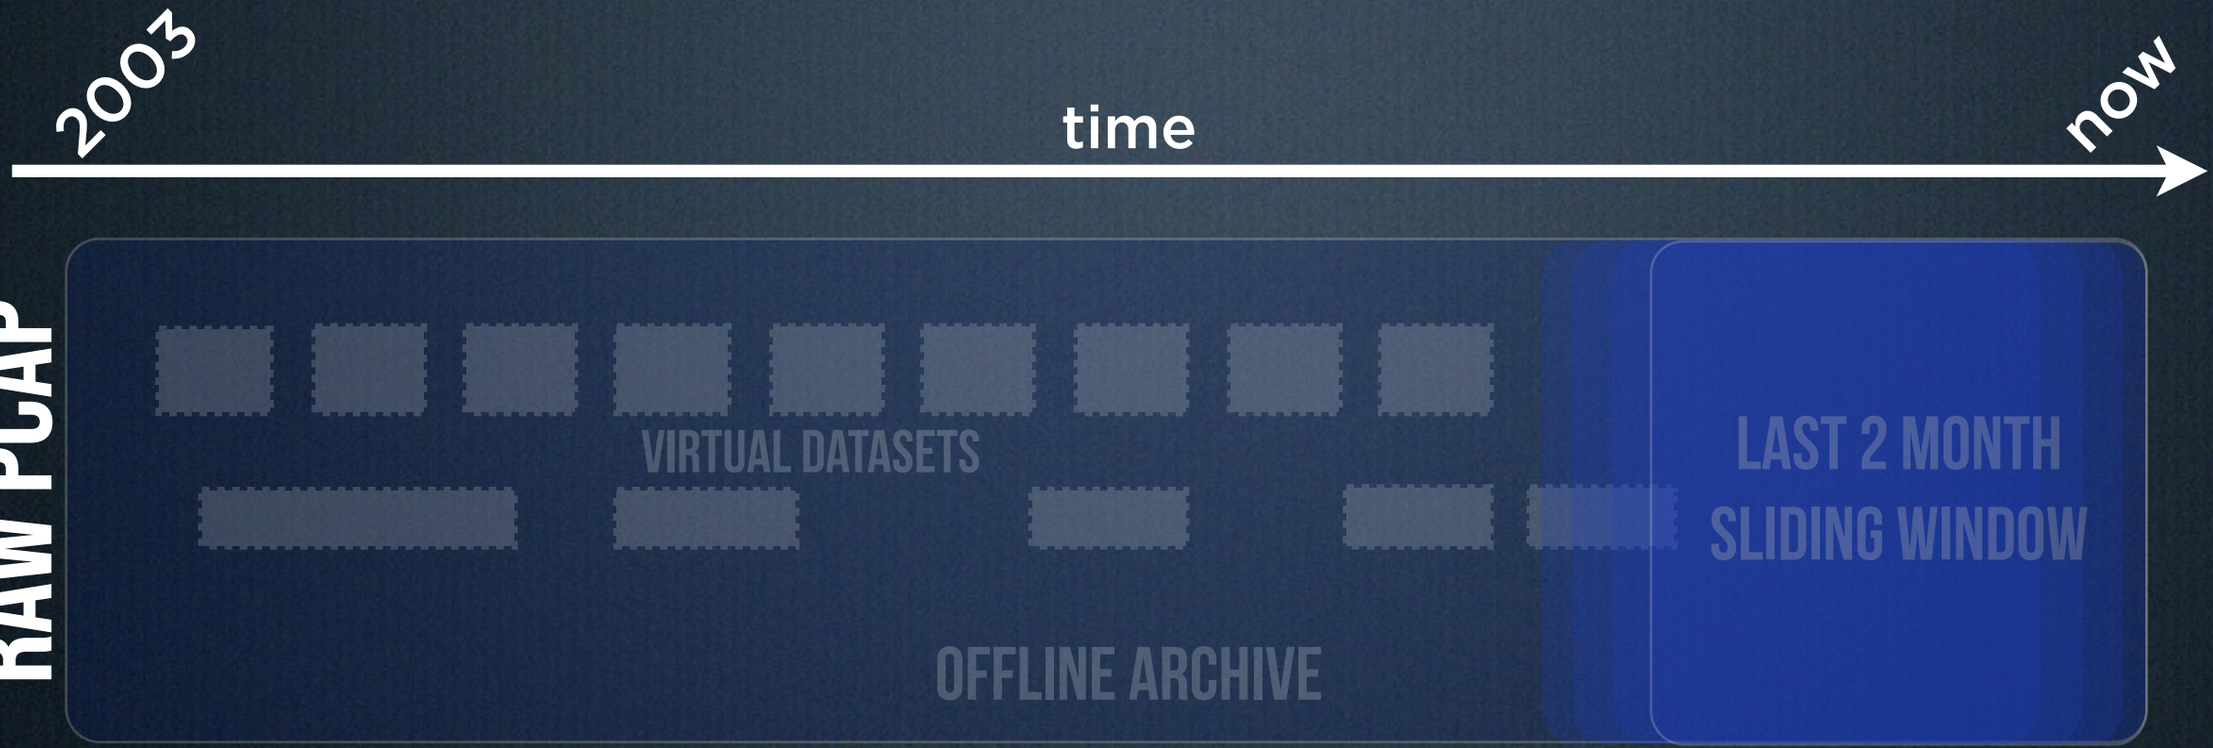

DERIVED DATA

- FLOWTUPLE
- [PROBLEM-SPECIFIC DATA]
- TIME SERIES METRICS

caida

DATA HIERARCHY

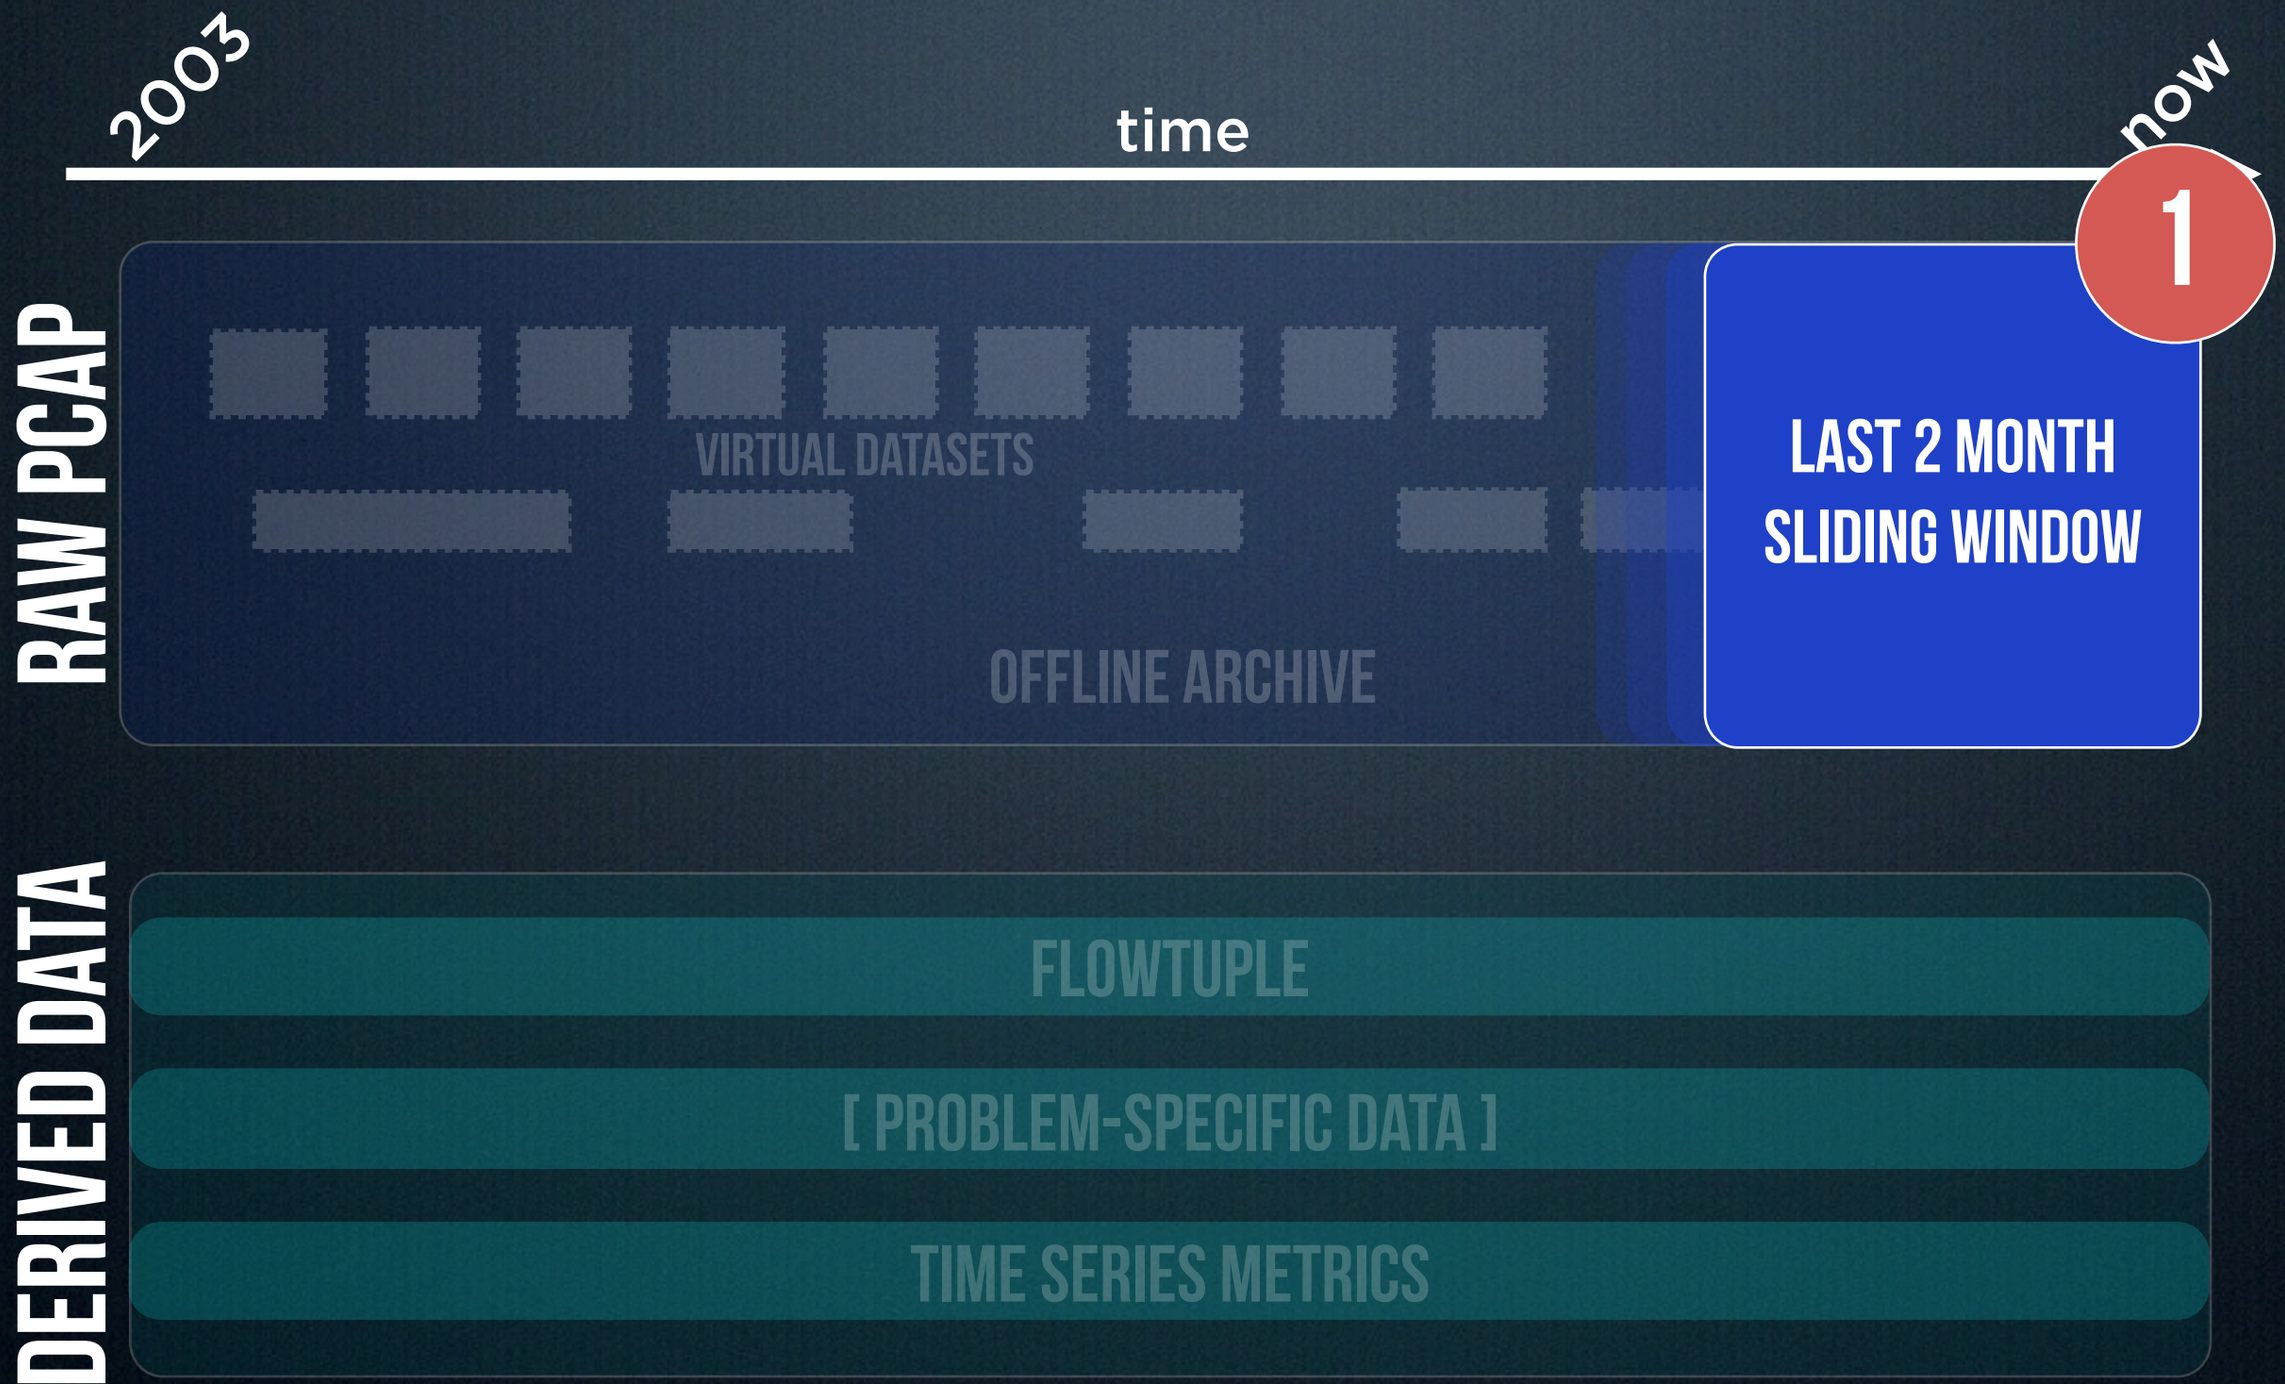

caida

DATA FLOW

Process overview: from collection to use.

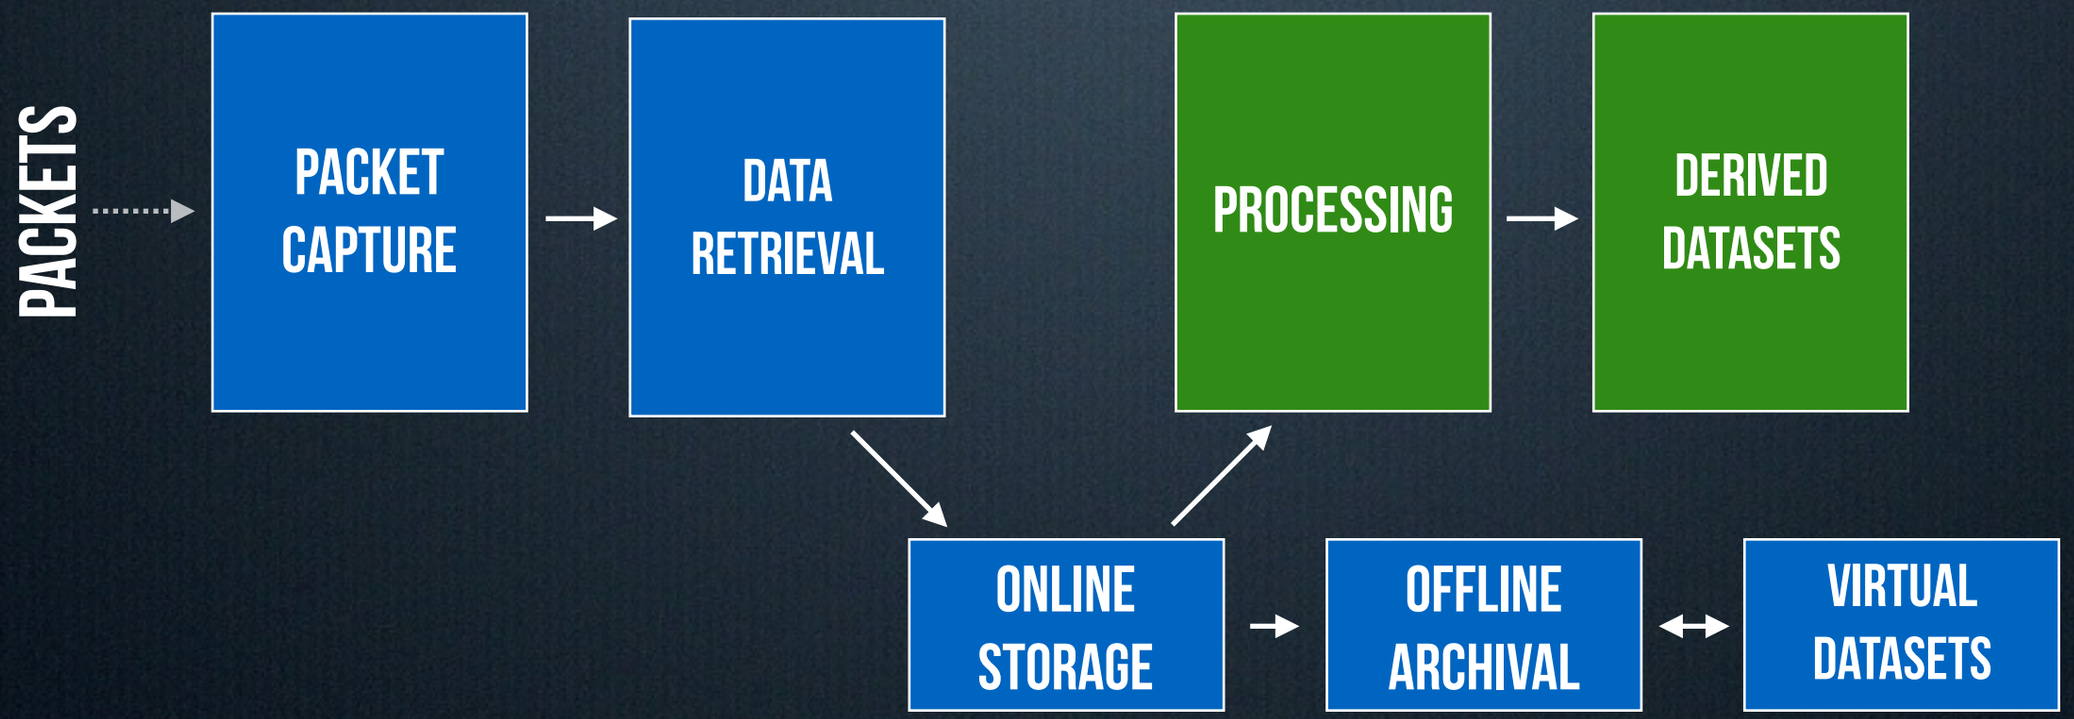

caida

DATA COLLECTION

Packet volume is not the challenge; velocity is

- X.0.0.0/8 is routed to us over a 1GE link
- ~40 Mbps, but…
- ~50k packets per second, mostly empty, with…
- Massive spikes in packet rate (> 100k pps)

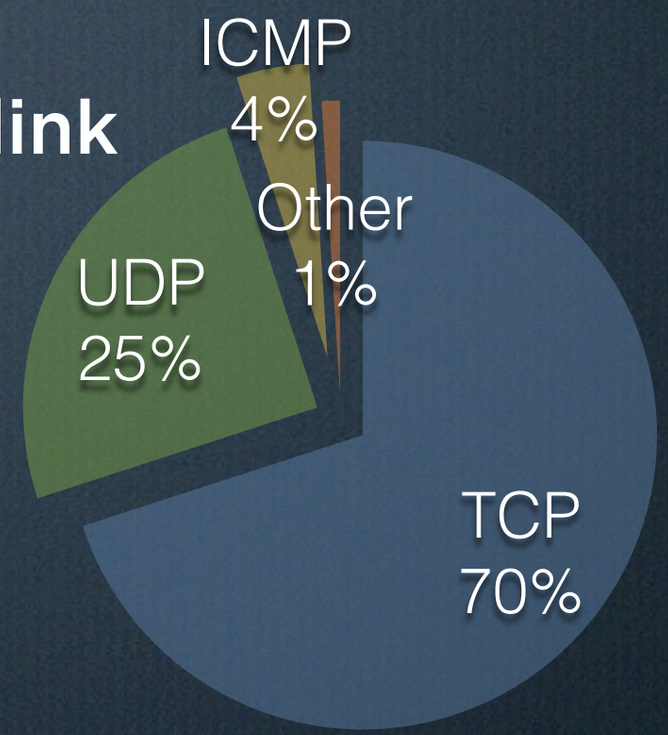

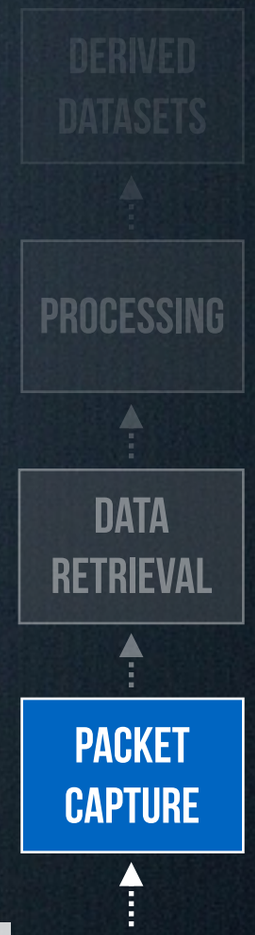

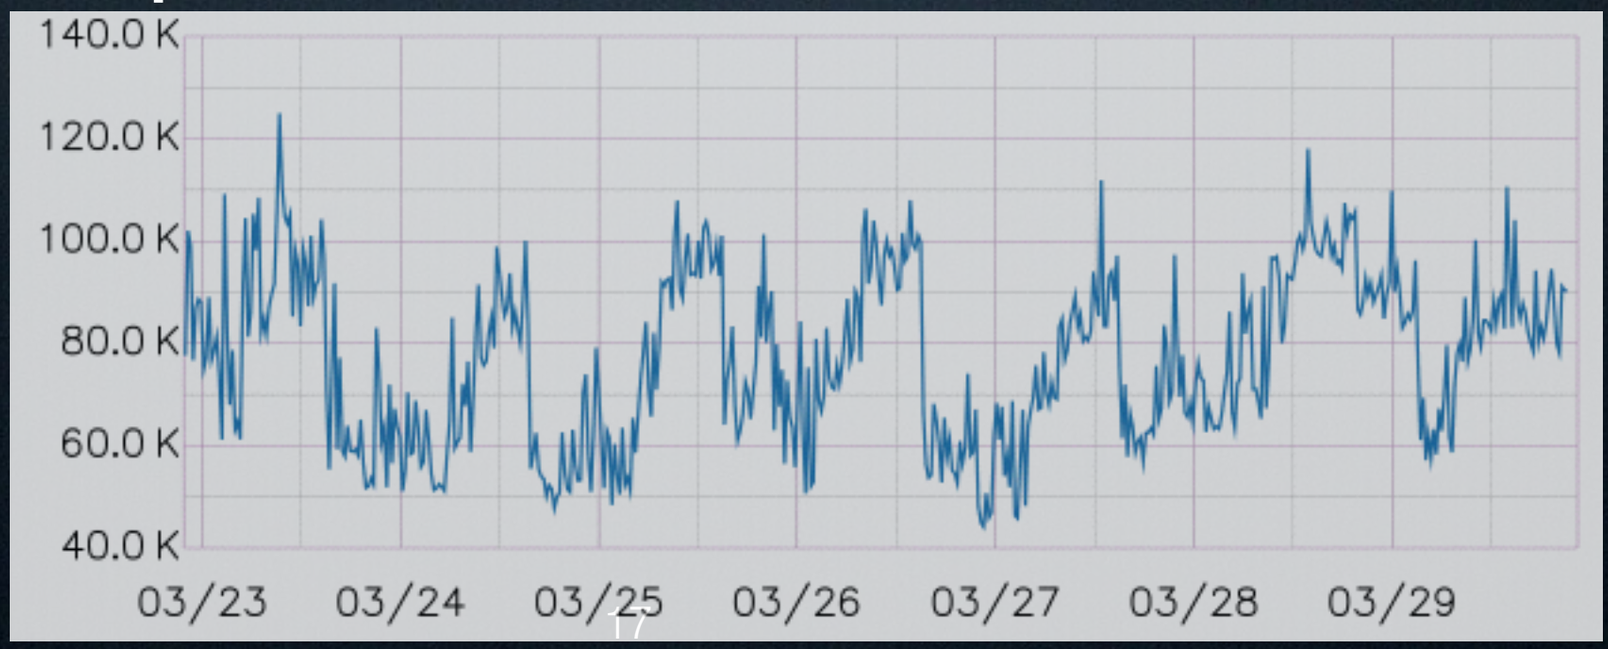

caida

DATA COLLECTION

Writing the packets to disk

- **FreeBSD server with commodity hardware**
  - Co-located with the telescope router
- **Customized pcap capture software**
  - Based on wdcap from WAND (research.wand.net.nz/software/wdcap.php)
  - Allows for on-the-fly gzip compression
- **< 0.0005% packet loss, but we know there is loss upstream**
- **Rotates pcap files every hour**

caida

DATA COLLECTION

Getting the hourly pcap files

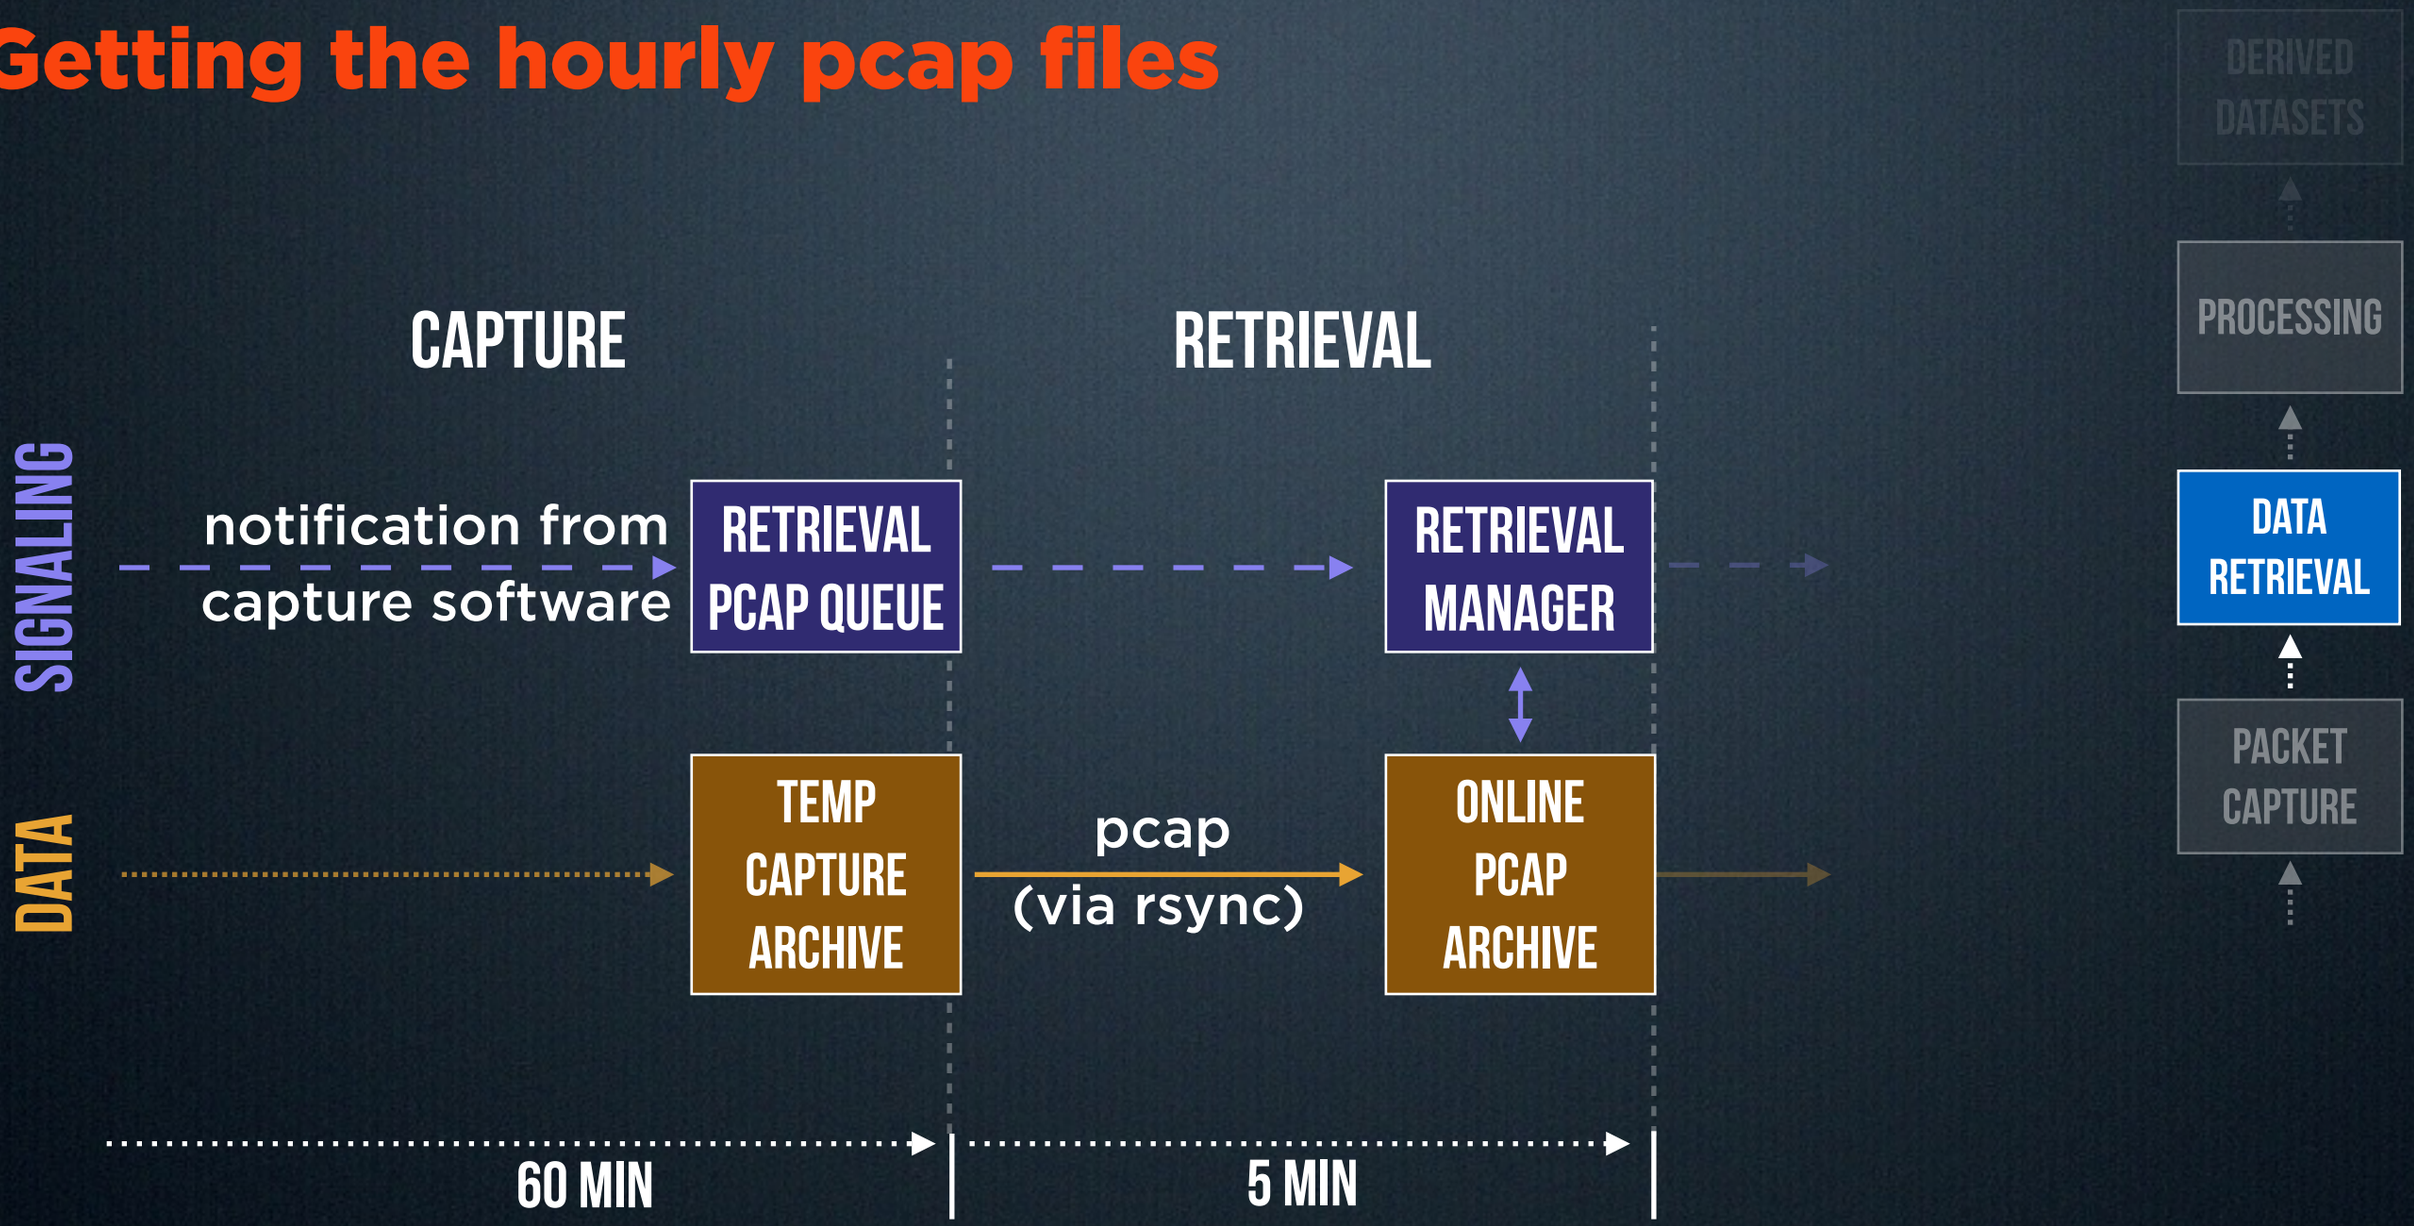

caida

DATA PROCESSING

Software for large-scale analysis of traffic traces

- Interval-driven packet processing software (Corsaro)
- Easily extensible, cooperative plugin architecture
- Per-packet operations include:
  - IP Geolocation
  - IP to ASN lookups
  - CryptoPAn Anonymization
  - Meta-data based packet filtering
- **Per-interval operations:**
  - Compute per-interval statistics
  - Write out aggregated data

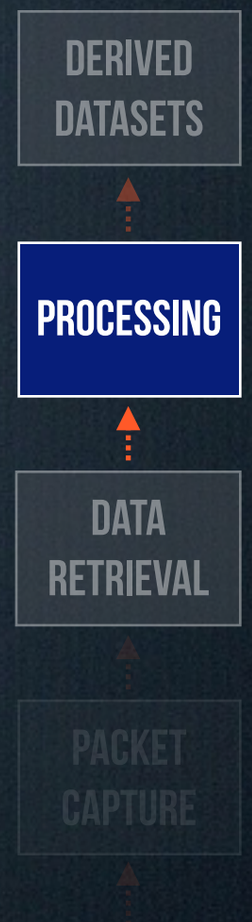

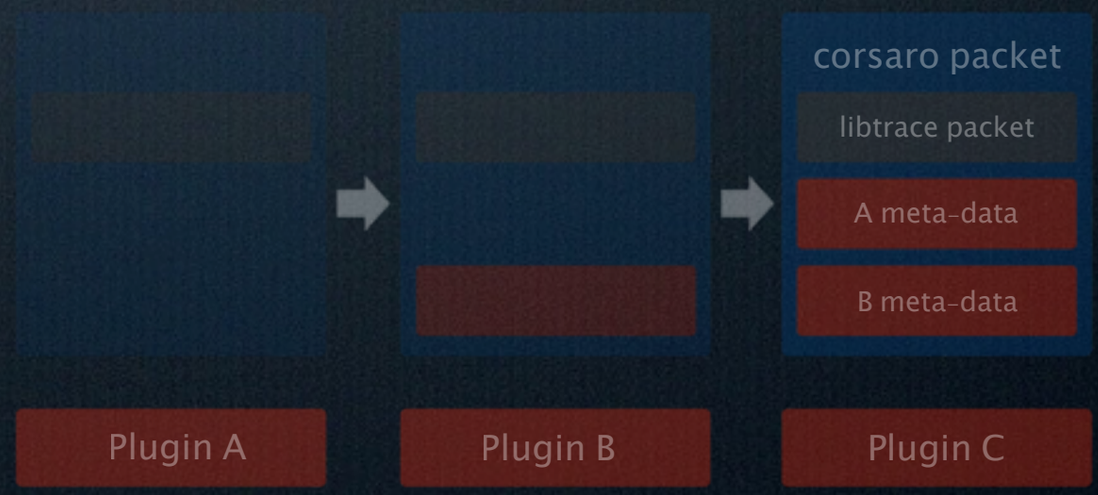

Corsaro
http://www.caida.org/tools/measurement/corsaro

caida

DATA PROCESSING

Looking at every packet

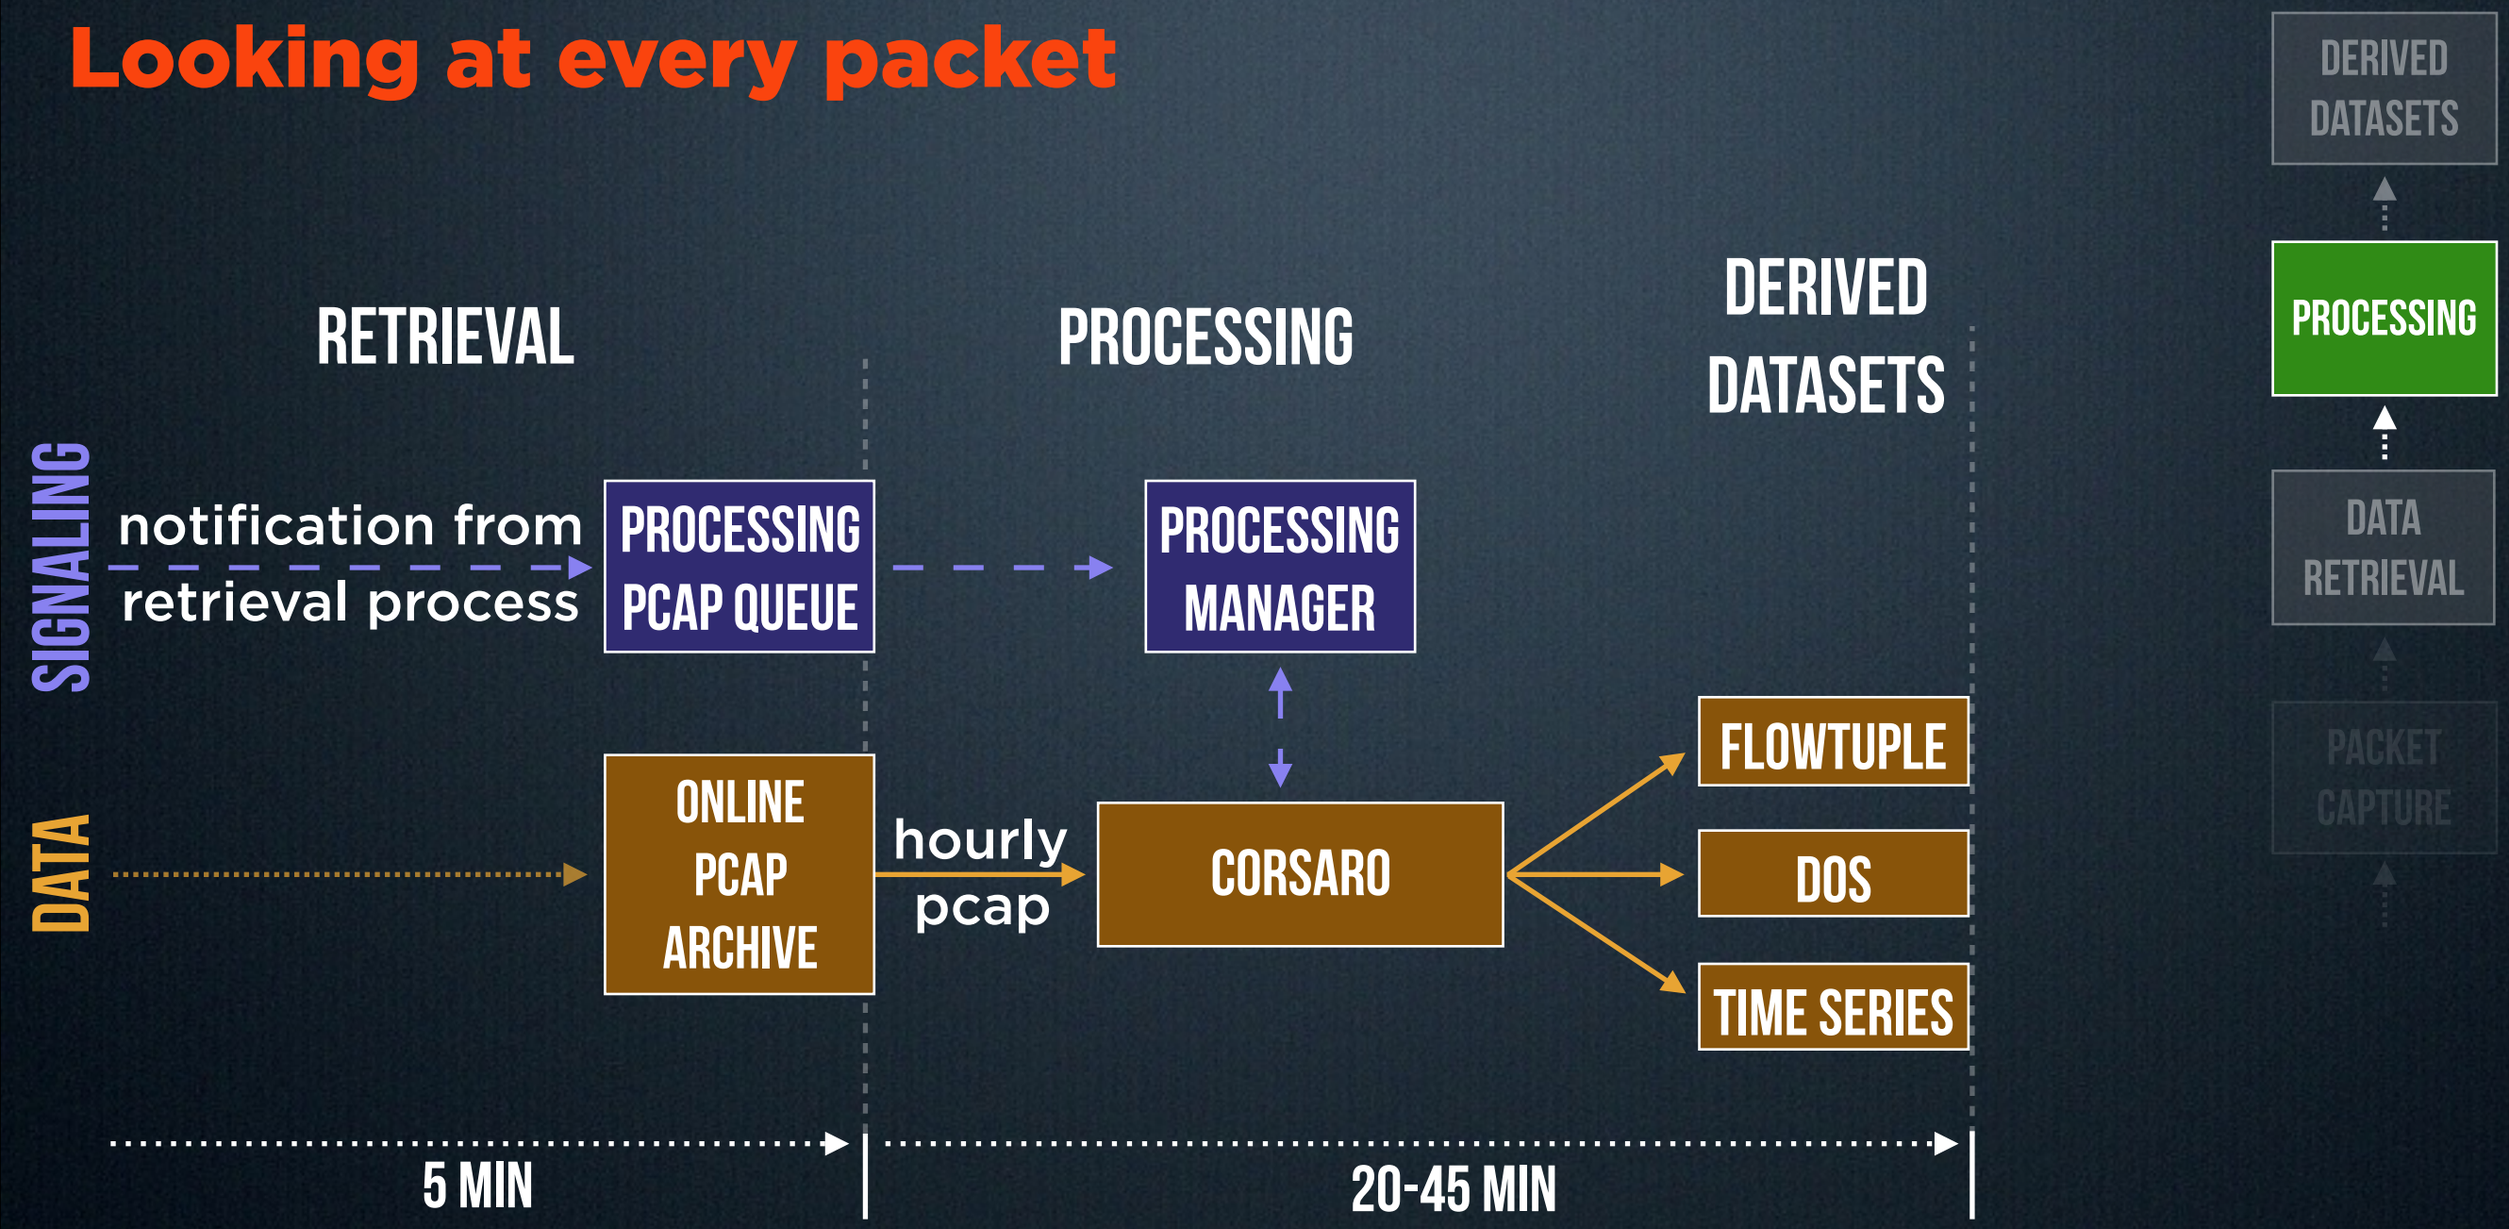

caida

DATA PROCESSING

Future plans

(REAL TIME)

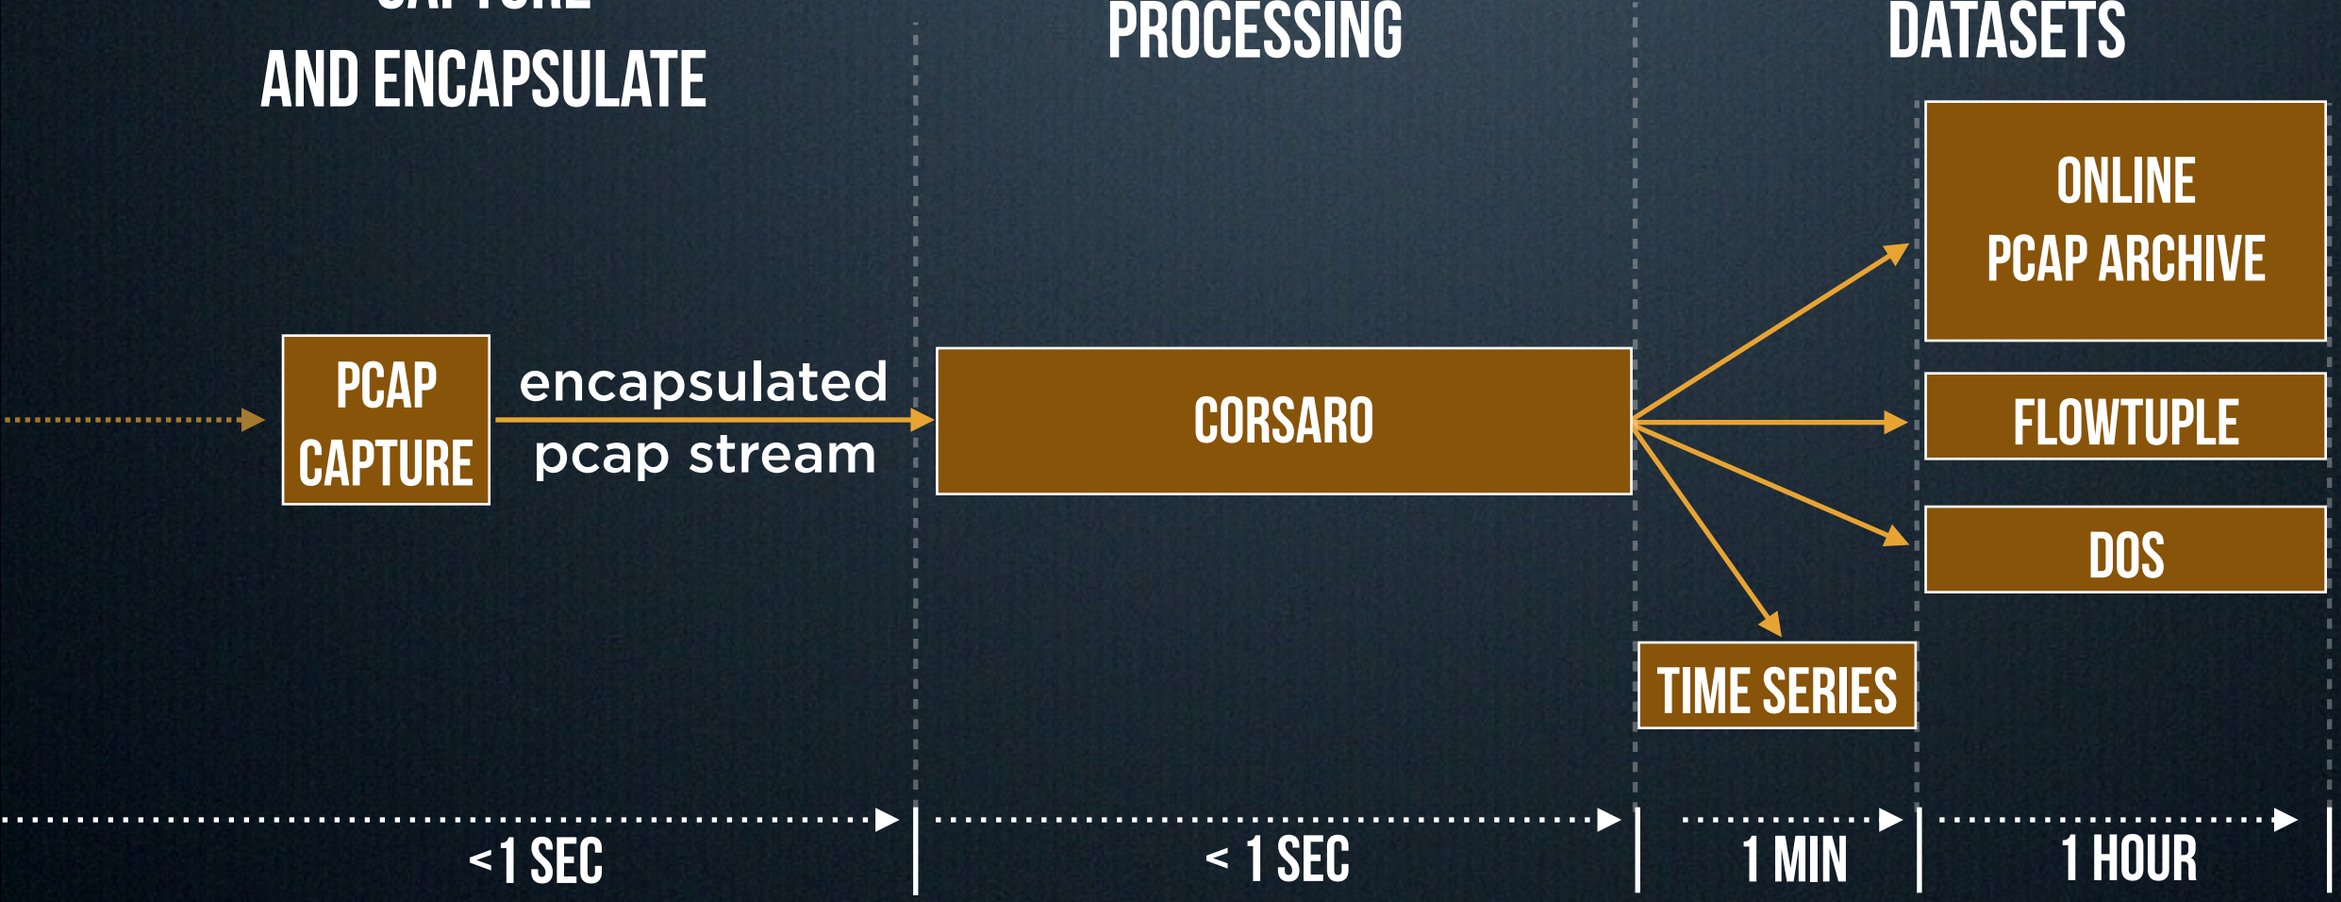

caida

DATA HIERARCHY

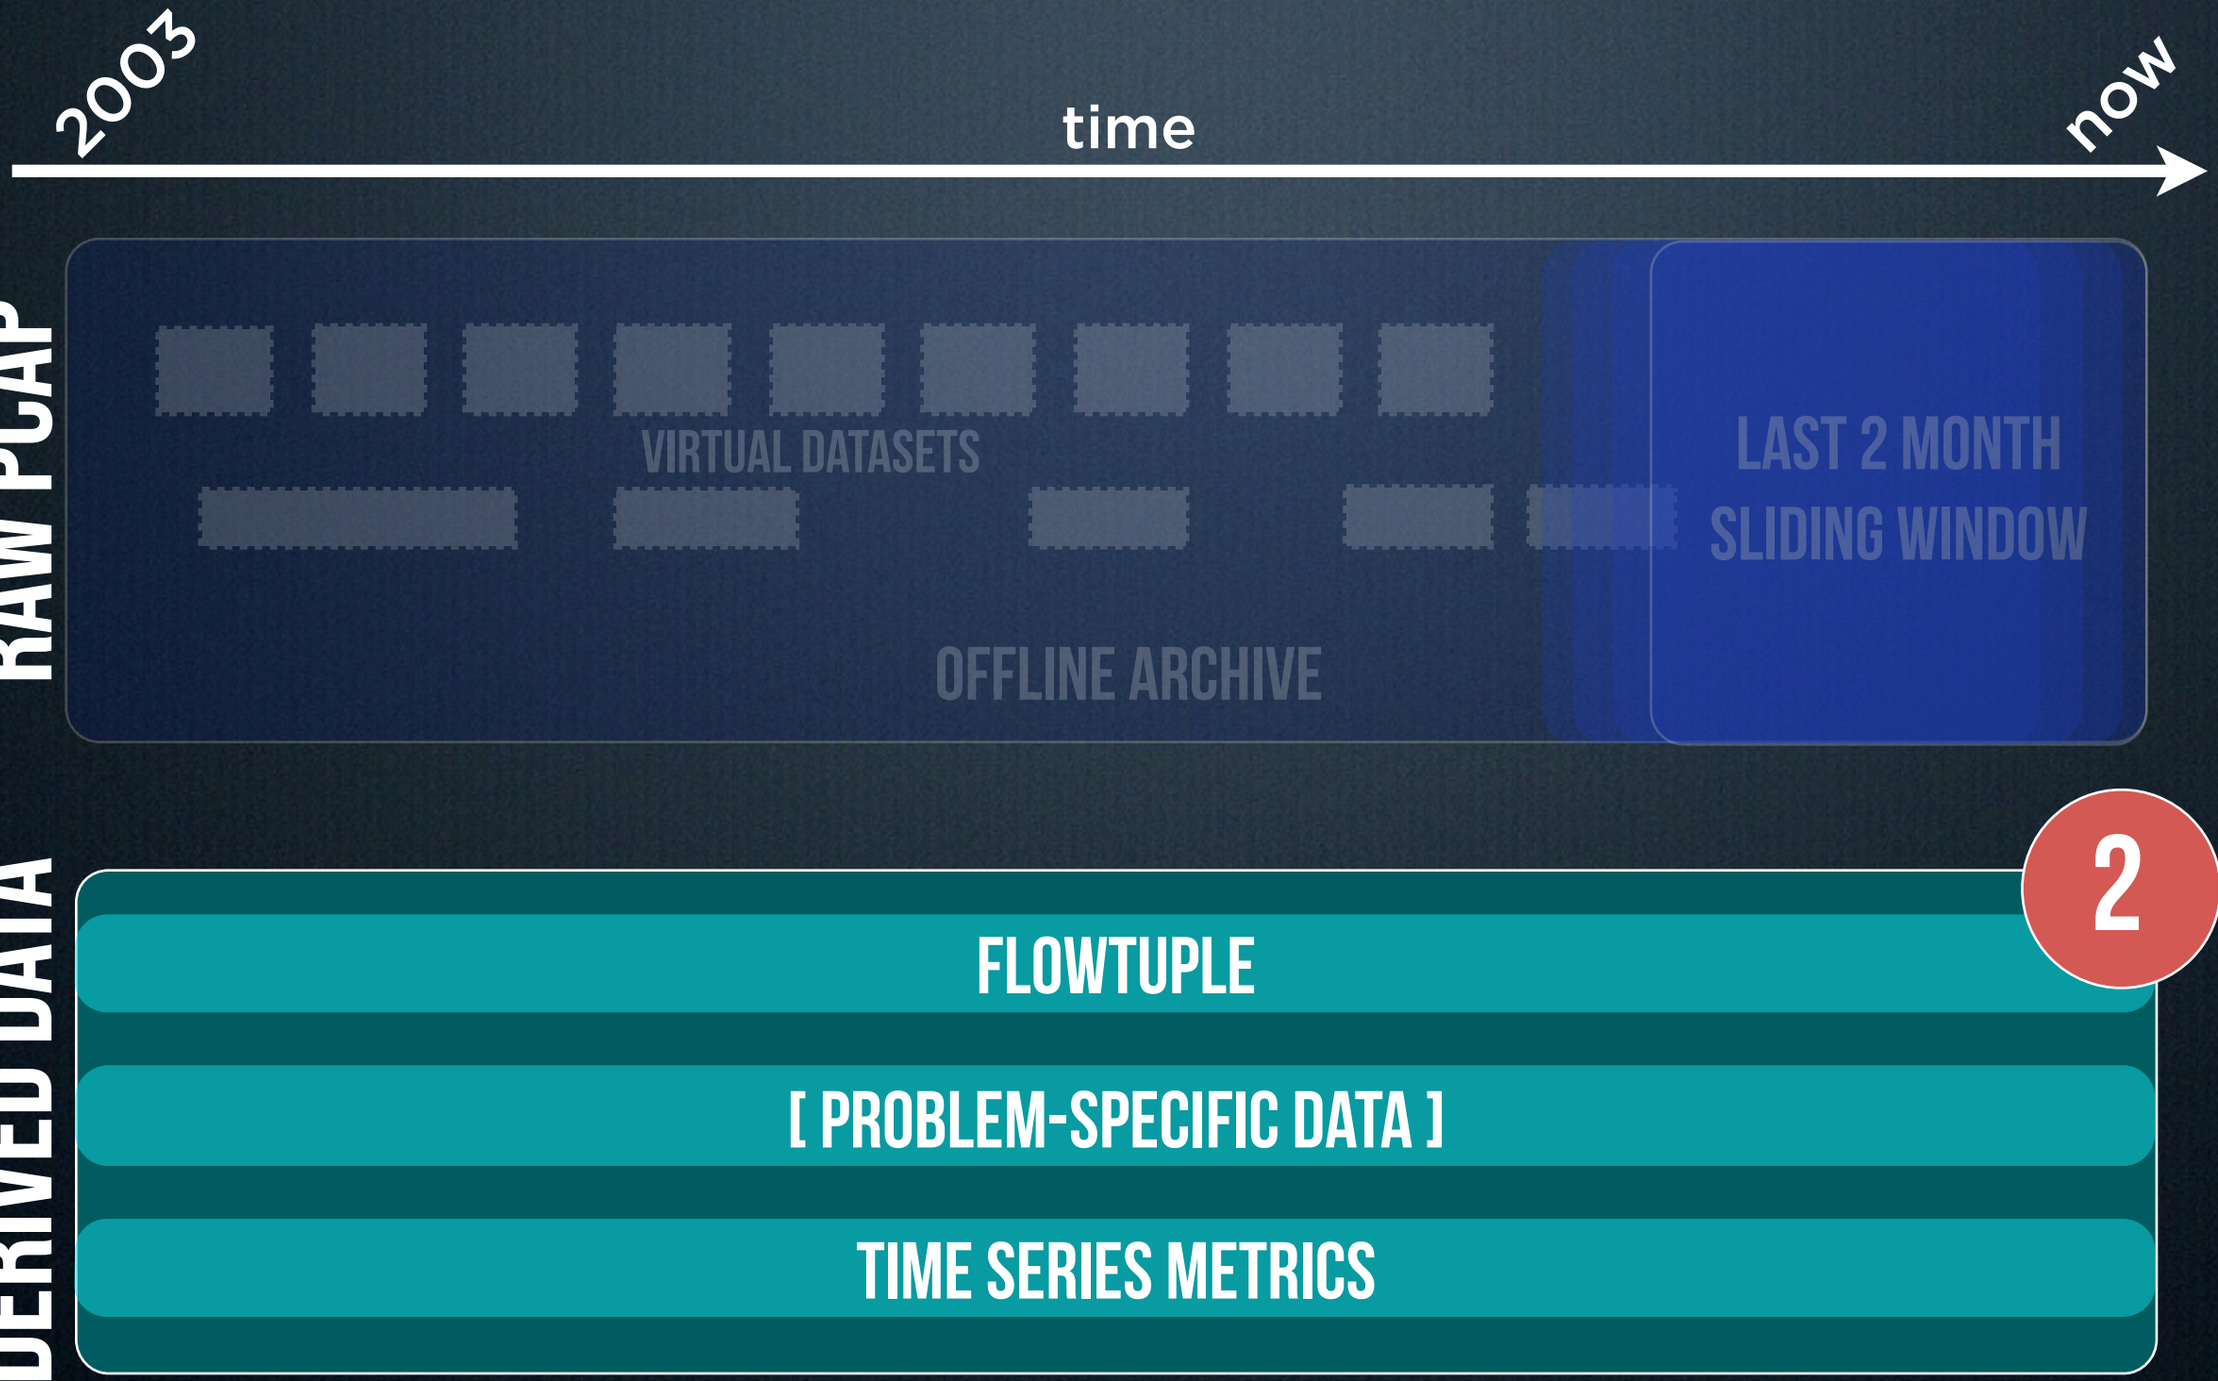

caida

FLOWTUPLE

Aggregates packets into 8-field flows

- Supports most telescope analysis use cases
- Src IP, Dest IP, Src Port, Dest Port, Protocol, TTL, TCP Flags, IP Len
- **Per-minute precision**
- **Serialized in efficient binary format**
  - ~10% size of compressed pcap
  - Easier to share
- **Corsaro reads and writes FlowTuple for easy analysis**

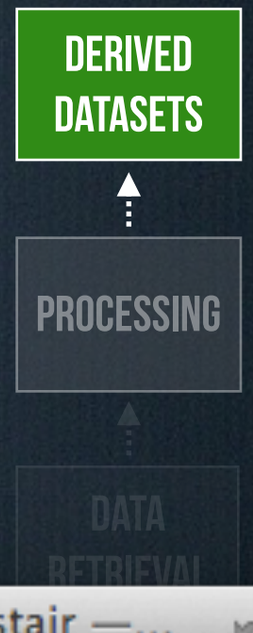

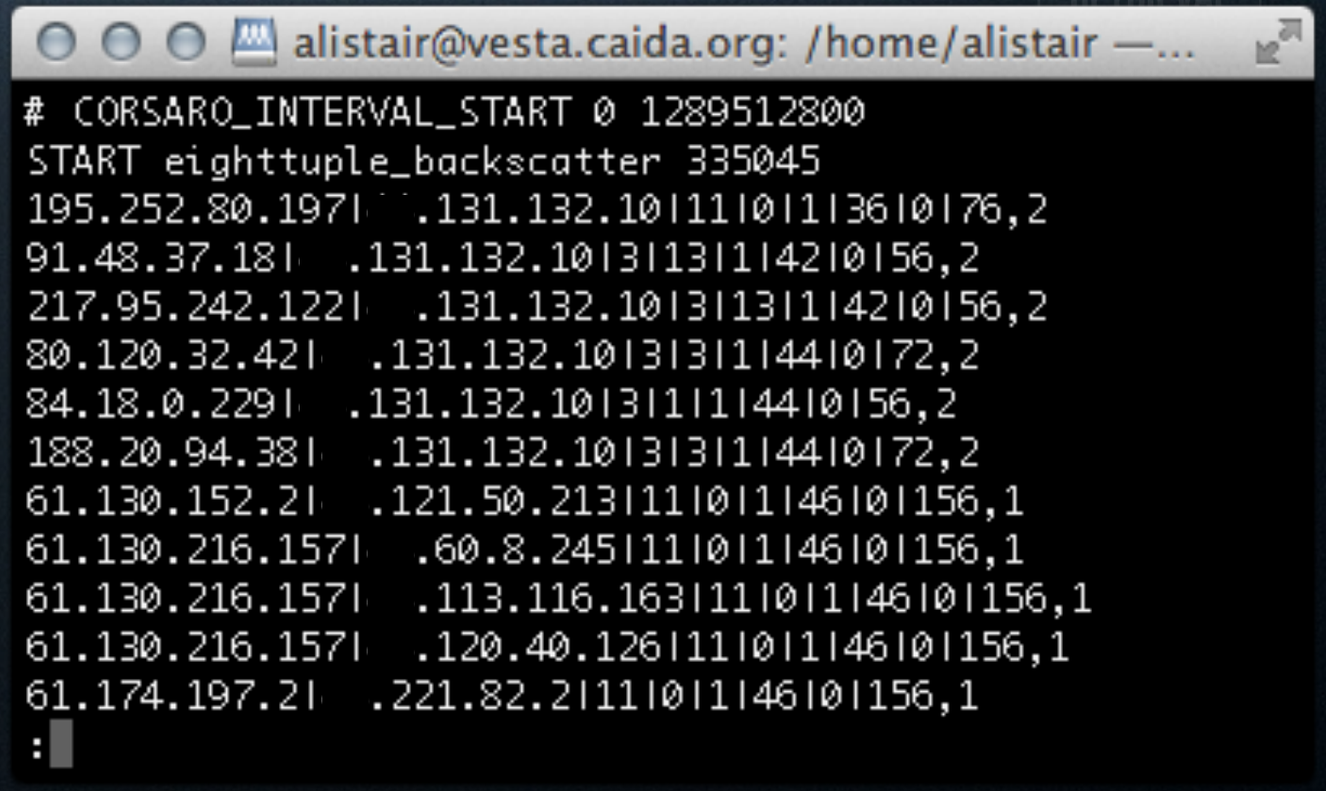

caida

TIME SERIES METRICS

Aggregates all packets along a single dimension

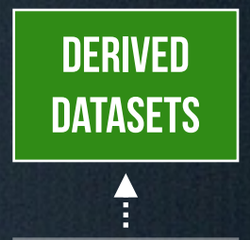

- E.g. # unique source IPs from Syria per minute
- We extract:
  - Packet count
  - Bytes
  - Unique Source IPs
  - Unique Destination IPs
- for:
  - Geolocation (country & region)
  - ASN
  - Protocol
  - Port (TCP/UDP, Src/Dest)
  - + more to come!

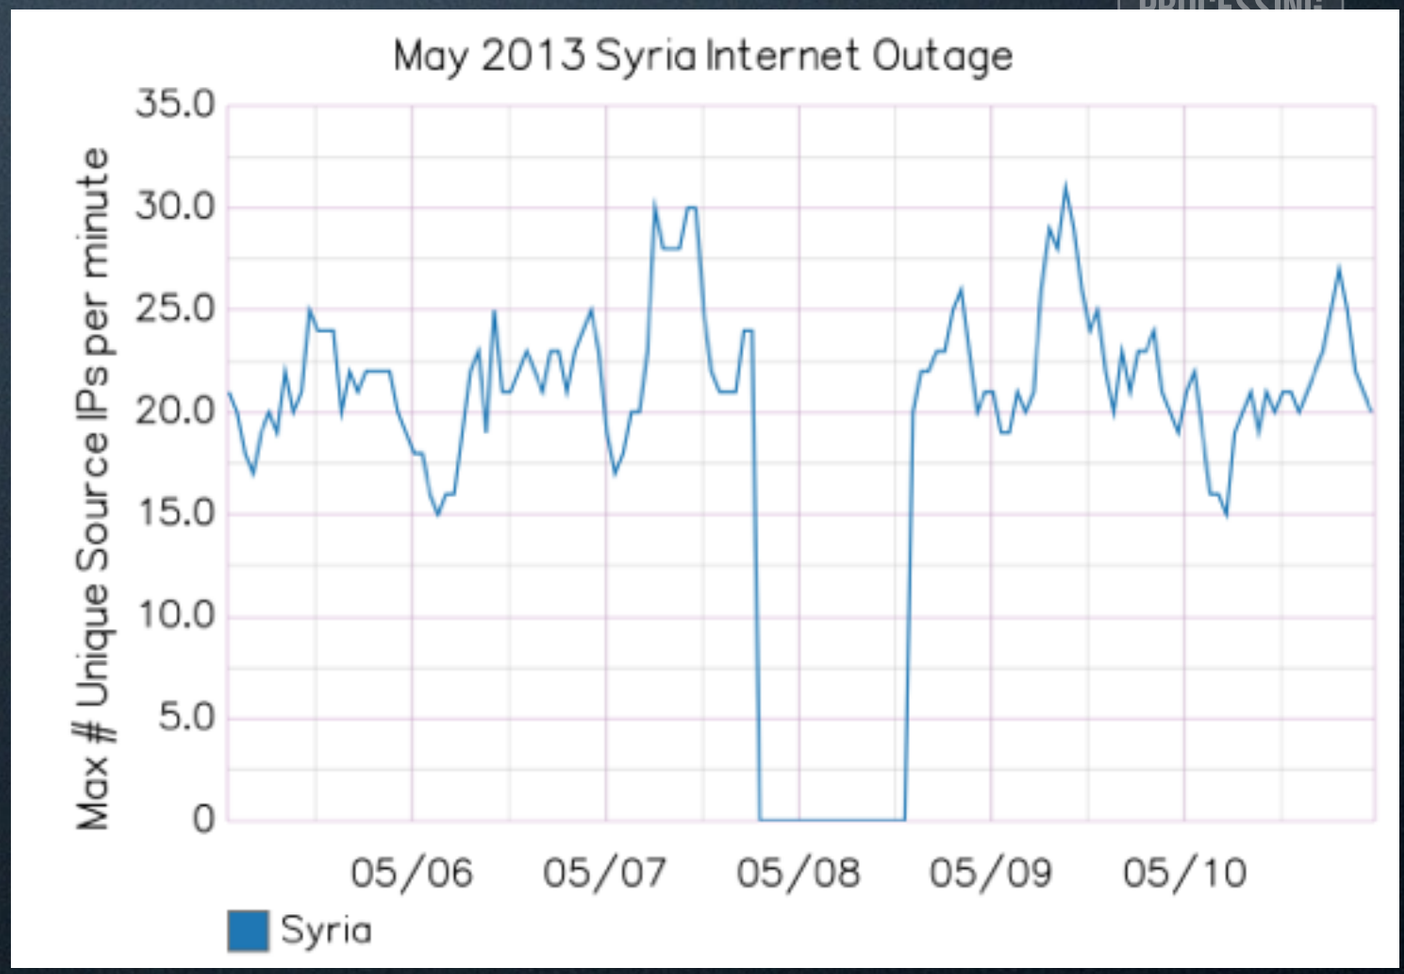

caida

TIME SERIES METRICS

Storing the metrics

- Corsaro generates >130k metrics per minute

```
darknet.ucsd-nt.geo.maxmind.--.--.ip_len          445246   1396047600
darknet.ucsd-nt.geo.maxmind.AS.AP.uniq_src_ip     31       1396047600
darknet.ucsd-nt.geo.maxmind.AS.AP.uniq_dst_ip     170      1396047600
darknet.ucsd-nt.geo.maxmind.AS.AP.pkt_cnt         176      1396047600
```

- **Inserted into RRD-style databases** (github.com/graphite-project/carbon)
- **Same performance issues as RRDtool:**
  - Metric hierarchy is encoded in filesystem
  - Every metric is a file (100k metrics => 100k files)
    - Huge IOPS requirements
- **Stored on dedicated DB server with 4TB SSD**

caida

DATA HIERARCHY

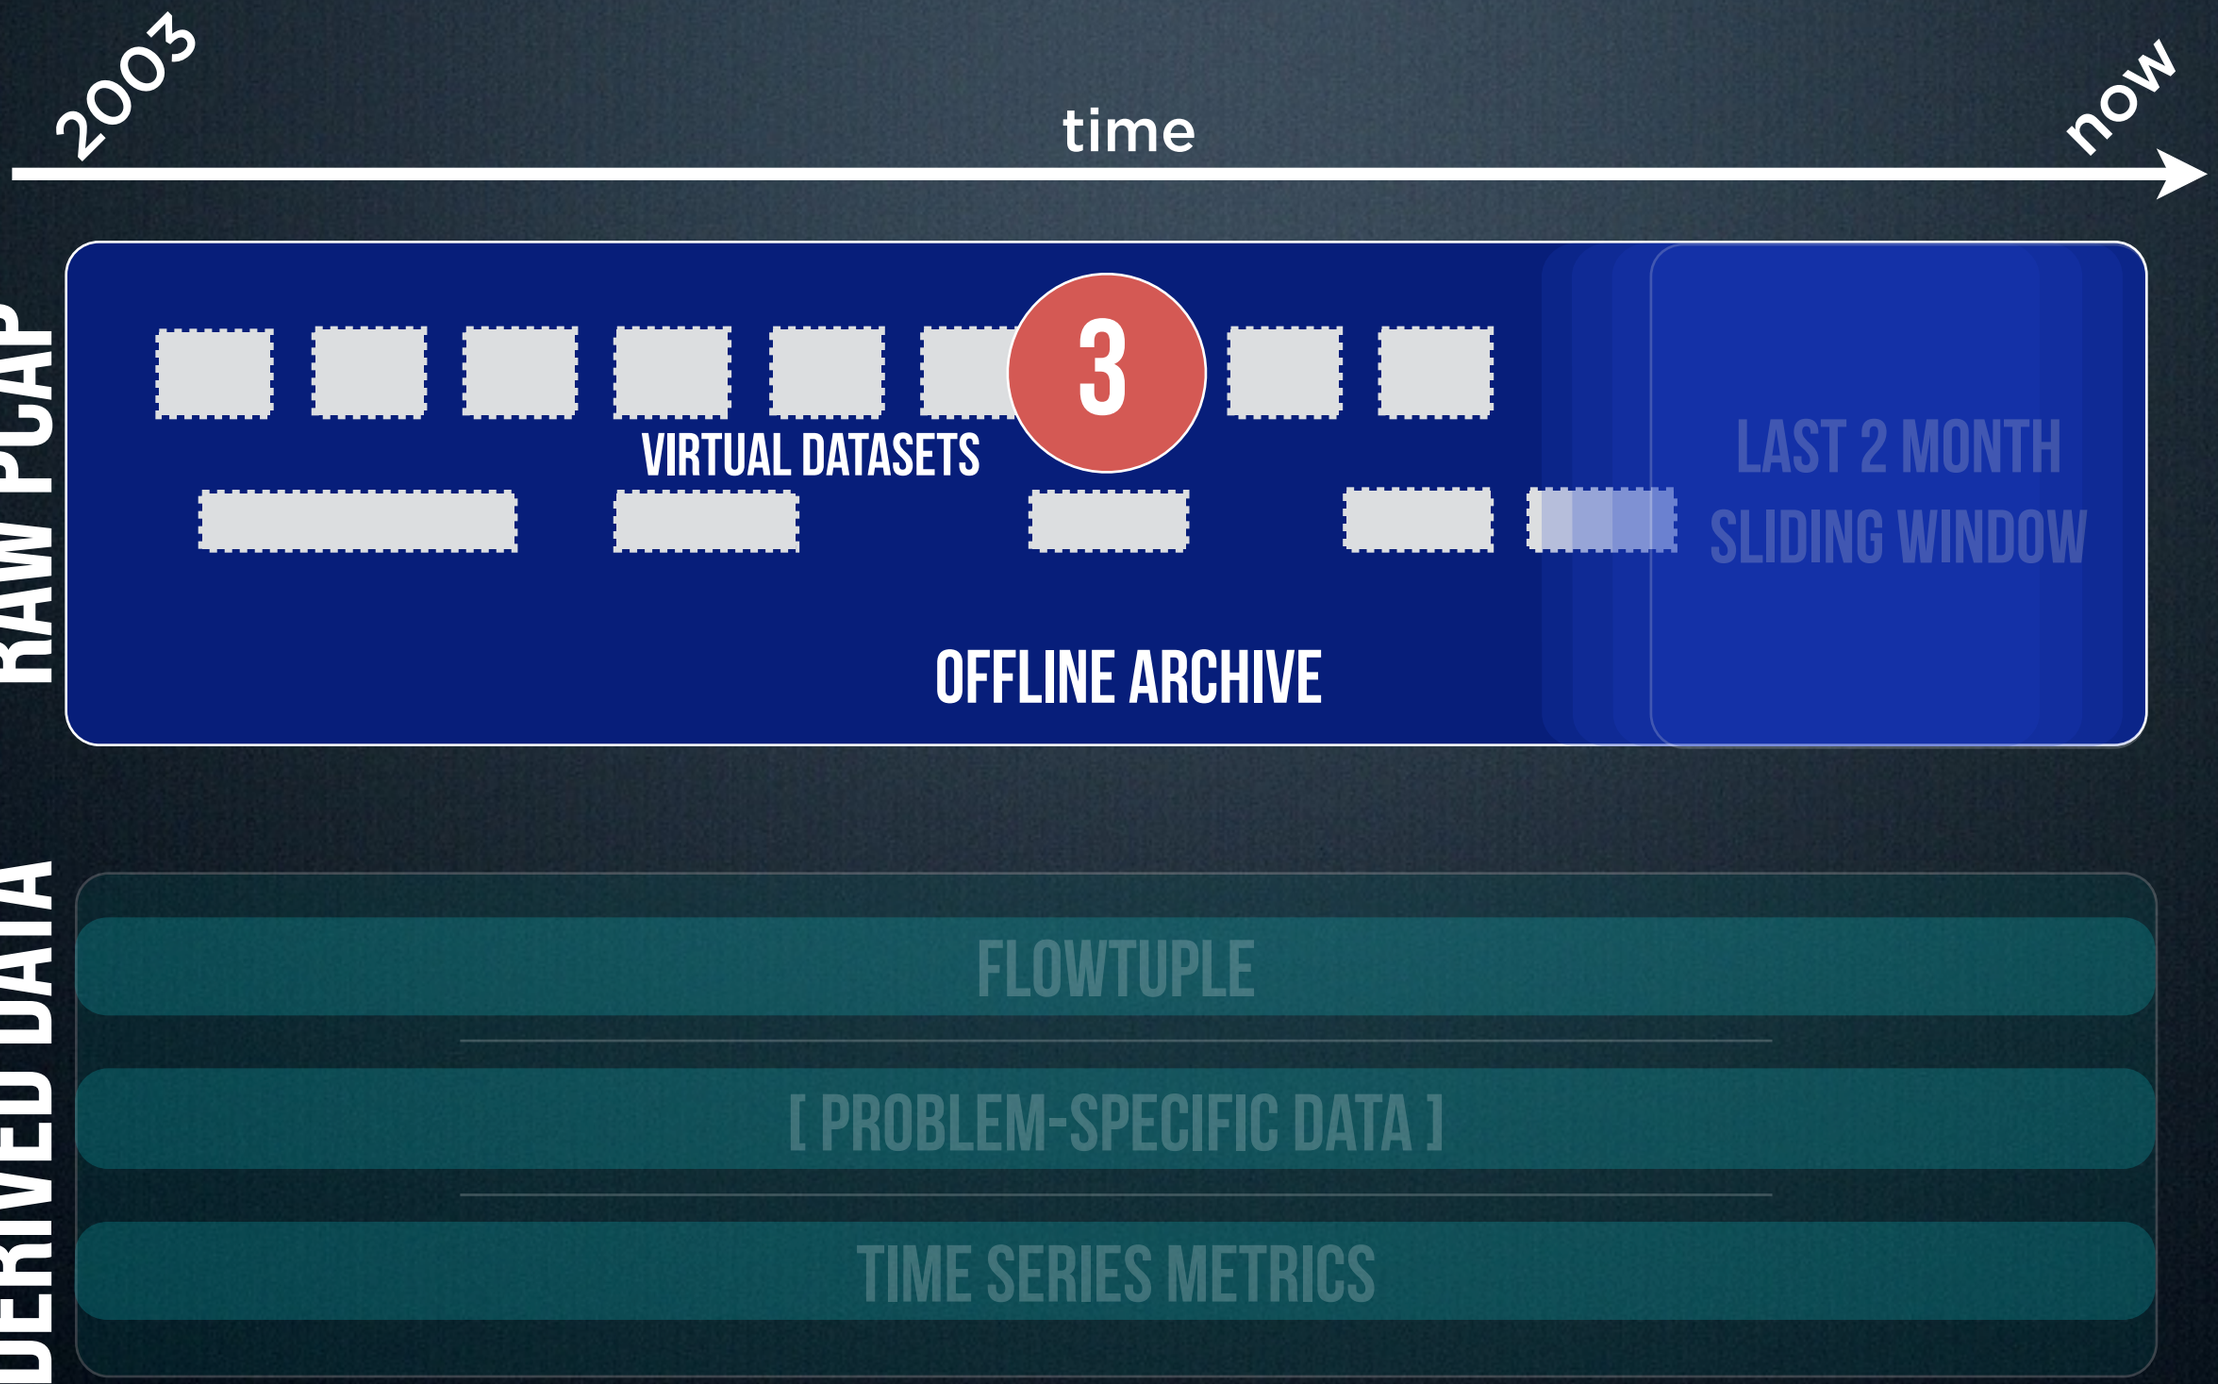

caida

VIRTUAL DATASETS

Structured access to archived data

- Archive data is offsite as raw hourly pcap
- Collections of **pointers** to pcap files for an event
- For example:
  - Routine samples (monthly, quarterly)
  - Published datasets
  - Data under investigation
- **Automated dataset restoration**
  - Restores a month of data in < 1 day

alistair@thor.caida.org: /data/telescope/meta/r

```
---
data_types: pcap
date_created: 2014-03-20
maintainer:
  email: alistair@caida.org
  username: alistair
name: 'UCSD-NT quarterly sample: 2014-1'
owner:
  email: alistair@caida.org
  username: alistair
pcap_ranges:
- - 2014-02-26
  - 2014-03-05
```

csd-nt.samples.quarterly.2014-1.dsd.yaml

USING THE DATA

AGILE TRAFFIC DATA ANALYSIS

Powerful pipeline for interactive time series exploration

CORSARO

Suite for large-scale analysis of (telescope) traffic traces

+

CHARTHOUSE

Interactive time series exploration

caida

CHARTHOUSE charthouse.caida.org

Web-based interactive time series exploration

- **Built on Graphite** (graphite.readthedocs.org)
- **Supports dynamic charting**
- **Limited support for geographic maps**

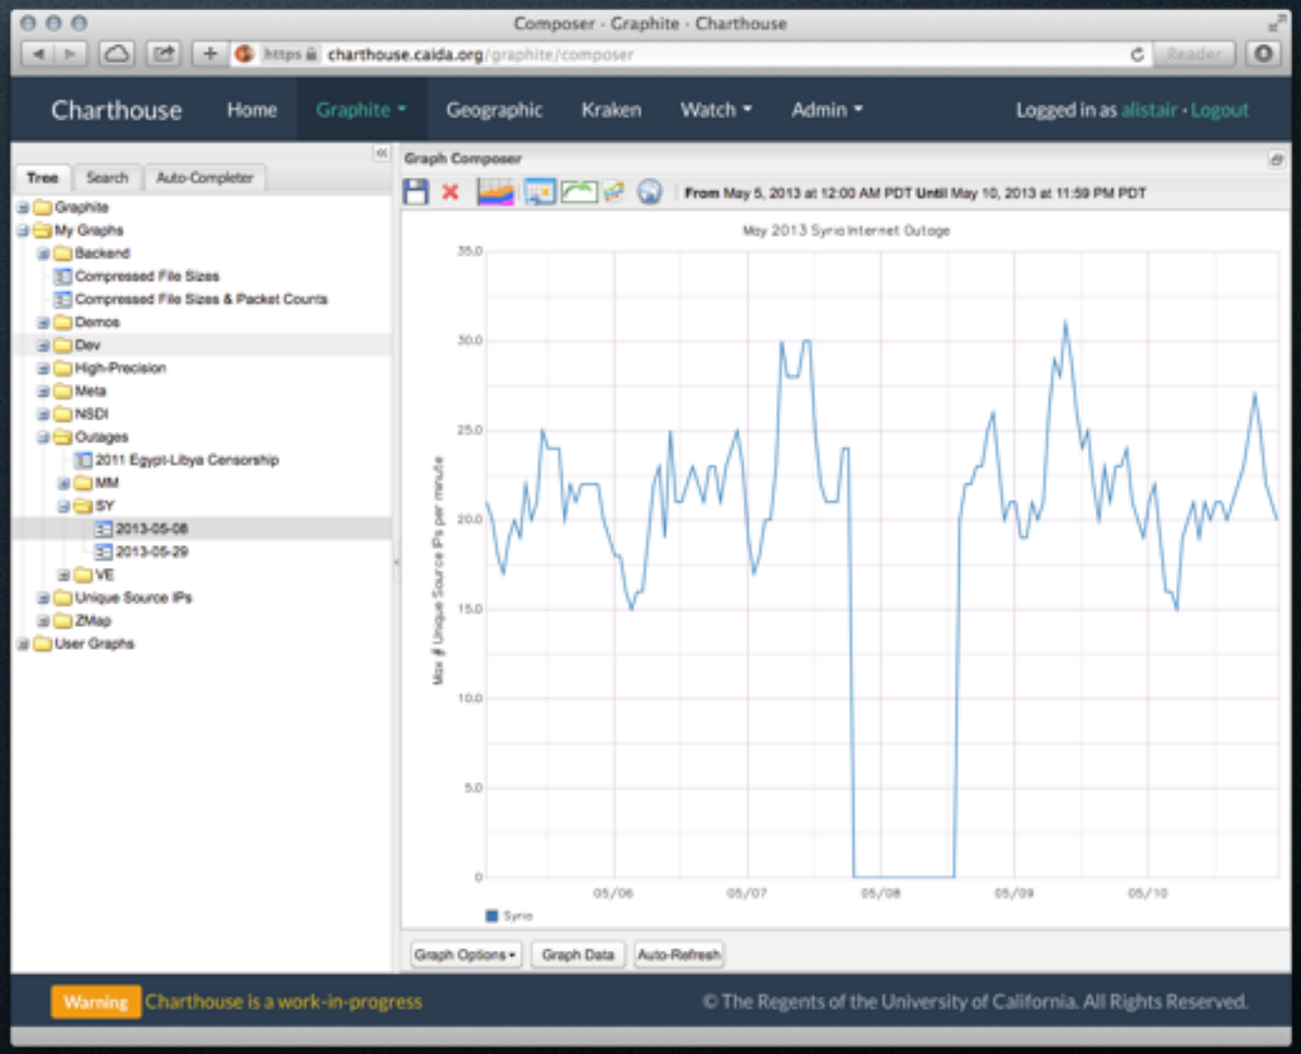

caida

CHARTHOUSE DEMO

TECH CHALLENGES, POLICY & FUNDING

TECHNICAL CHALLENGES

With data capture, storage, processing and sharing

- **Limited disk storage**
  - We almost deleted all historical pcap
- **Packet loss**
  - Our upstream router is dropping ~1%!
- **Disk IOPS limitations**
  - RRD-style DBs are tough on disks
- **Getting closer to realtime analysis**
  - Sub-second latency
- **Failure recovery**
  - Lots of moving parts to reassemble
- **Researchers…**

caida

PHILOSOPHIES

Words we live by

- **Reuse components whenever possible**
  - There is already code to do 70% of what you need. Hack the rest.
  - Resist the urge to engineer everything yourself
- **Performance is critical**
  - but only once it becomes critical
- **Disk is cheap, but not free**
- **Document returns on investment**
  - It helps justify ongoing funding
  - e.g. CAIDA maintains a list of non-CAIDA papers using CAIDA data (caida.org/data/publications)

caida

PRIVACY CONCERNS

Balance sharing against privacy concerns

- **Telescope data is potentially sensitive**
  - Source IPs may refer to vulnerable machines
  - Payload may include personally-identifiable information
- **Acceptable Use Policies written with legal assistance**
  - Privacy-Sensitive Sharing Framework (PS2) (caida.org/publications/papers/2010/dialing_privacy_utility)
  - Supports various levels of sensitivity of data
    - Code-to-data model for near-real-time traffic data
    - Enforced resource limits on compute machines

PROVIDING ACCESS

Promoting community use of the data

- **Collection and sharing of data is one of CAIDA's core objectives**
  - Support scientific analysis of Internet traffic, topology, routing, performance, and cybersecurity events
  - Data-sharing is not a trivial component of research effort
- **NSF now requires a Data management plan**
  - https://www.nsf.gov/eng/general/dmp.jsp
  - Must include plans for how to manage, share, sustain data
  - Enable reproducibility of science

FUNDING caida.org/funding

Sustainable measurement infrastructure

- Data provides opportunity for research, therefore funding
- Some research grants cover data collection costs indirectly
  - Funding agencies generally prefer funding research to measurement, data collection, curation, or sharing
- **As the Internet becomes increasingly critical infrastructure, recognition grows that we understand it too little**
- **But we are far from having a "Bureau of Internet Statistics" or "Internet Census Bureau"**

caida

OUR FUNDING SOURCES caida.org/funding

Mostly U.S. gov, some .com/.net

- **Department of Homeland Security S&T Directorate (HSARPA)**
  - PREDICT: funds CAIDA data curation, management, sharing
  - Cybersecurity Research: funds infrastructure, data analysis
- **NSF Computing Research Infrastructure (CRI) funding**
  - Support telescope infrastructure:
    - enhance tools for analysis and visualization
    - enable real-time sharing
    - community development
- **CAIDA commercial sponsors who want empirical research**

caida

SUSTAINABLE STEWARDSHIP

An eye toward the future

- Develop and maintain infrastructure as efficiently as possible
- Focus on problems relevant to many stakeholders
- **Build and support community**
- **Document uses of data (ours and others)**
- **Stay aware of funding opportunities**
  - e.g., NSF DIBBs (Data Infrastructure Building Blocks) program
    - Encourages development of robust and shared data-centric cyber-infrastructure capabilities

caida

FUTURE WORK

Watch this space

- **Realtime analysis**
  - Sub-second latency between capture and analysis
  - Automated large-scale outage detection
- **New Time Series DB**
  - Drastically improved performance
  - Under development by CAIDA
- **Charthouse Improvements**
  - Interactive geographic mapping
- **Pending CRI funding request**
  - E-RAID (Environment for Rapid Analysis of Internet Data)
  - Integrates multiple measurement types (Active, Routing, Telescope, …)

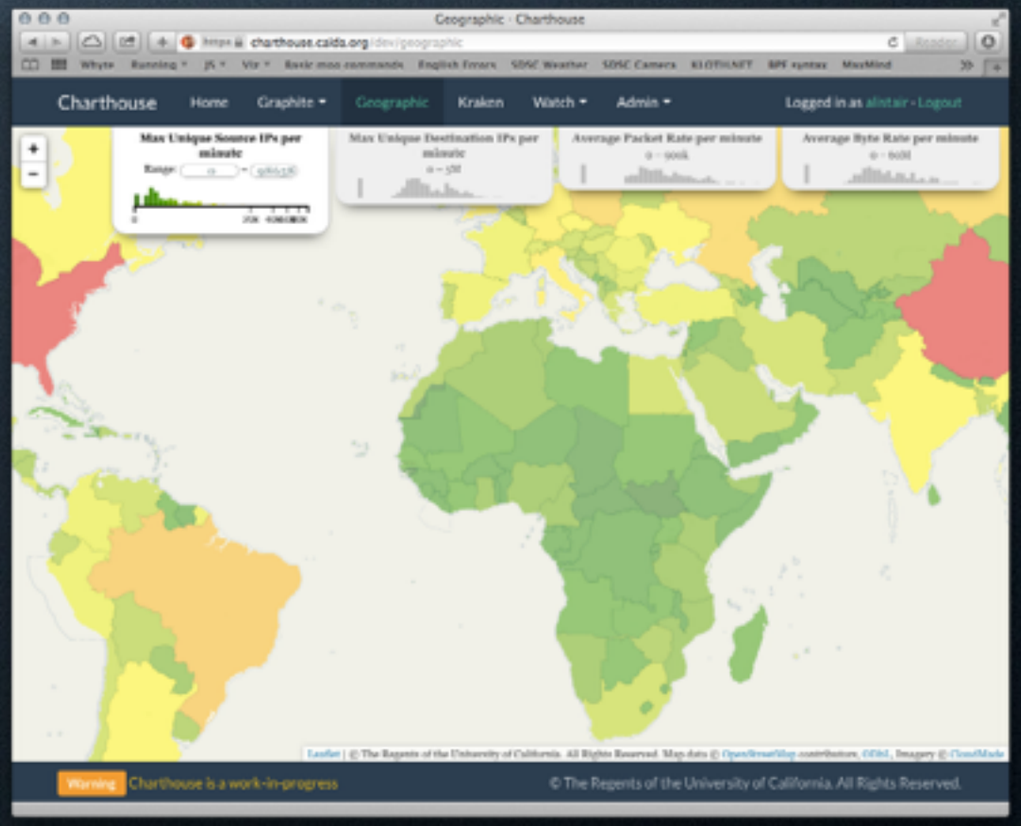

caida

QUESTIONS?