The story of Greendale

Turbinia: Automation of forensic processing in the cloud

Why are WE here?

Thomas Chopitea   @tomchop_

Aaron Peterson   @aarontpeterson

DFIR @ Google

- We write code, we use it to hunt bad guys

- **dfTimewolf** / **Turbinia** core devs

- Try to automate ourselves out of a job

Why are YOU here?

- You'll learn about the Cloud part of our forensics toolkit
    - It's all **Free and Open Source Software**

- You'll see how these tools fit together through a fictional **scenario**

We'll focus on:

- dfTimewolf
- Turbinia

- Plaso
- Timesketch

😅

Log2timeline / plaso

- Recursively parses **everything** in your filesystem and extracts timestamp information

- Builds a **system timeline** from this information

timesketch

- Forensic **timeline visualization** tool

- Plays well with **plaso**

- Multi-user, multi-case, multi-timeline

Log2timeline / dfTimewolf

- CLI utility - **the Glue** between different tools

- **Modules** (e.g. collectors, processors, exporters)

- **Recipes** (directions on how to chain Modules)

Log2timeline / dfTimewolf

Recipe name

Recipe parameters

```
$ dftimewolf grr_artifact_hosts tomchop.greendale.edu 12345 --sketch_id 666
```

Target GRR host

Comment

Existing sketch ID (optional)

tomchop.greendale.edu

/tmp/C.123

/tmp/C.123.plaso

GRR

Plaso

Time sketch

http://timesketch.greendale.edu/sketch/666
Timeline description: 12345

```python
contents = {
    'name': 'local_plaso',
    'short_description': _short_description,
    'modules': [{
        'name': 'FilesystemCollector',
        'args': {
            'paths': '@paths',
        },
    }, {
        'name': 'LocalPlasoProcessor',
        'args': {
            'timezone': @timezone,
        },
    }, {
        'name': 'TimesketchExporter',
        'args': {
            'endpoint': '@ts_endpoint',
            'username': '@ts_username',
            'password': '@ts_password',
            'incident_id': '@incident_id',
            'sketch_id': '@sketch_id',
        }
    }]
}
```

Turbinia

- **Open-source framework** for deploying, managing and running forensic workloads

- **Automate common tools** like Plaso, bulk_extractor, strings, etc) in cloud environments

- **Parallel processing** whenever possible

*"Grab this piece of **evidence**, run **plaso** on it, and **dump** results in a cloud bucket"*

Other details

- Written in Python

- PoC written in 2015 by @jberggren and @coryaltheide

- Rewritten starting in 2017

- We're good at logos!

Turbinia Installation Types

- **Cloud**
 - Storage, processing, metadata 100% on GCP Cloud

- **Hybrid**
  - Workers run on local machines with shared storage
  - Only metadata is sent to the Cloud
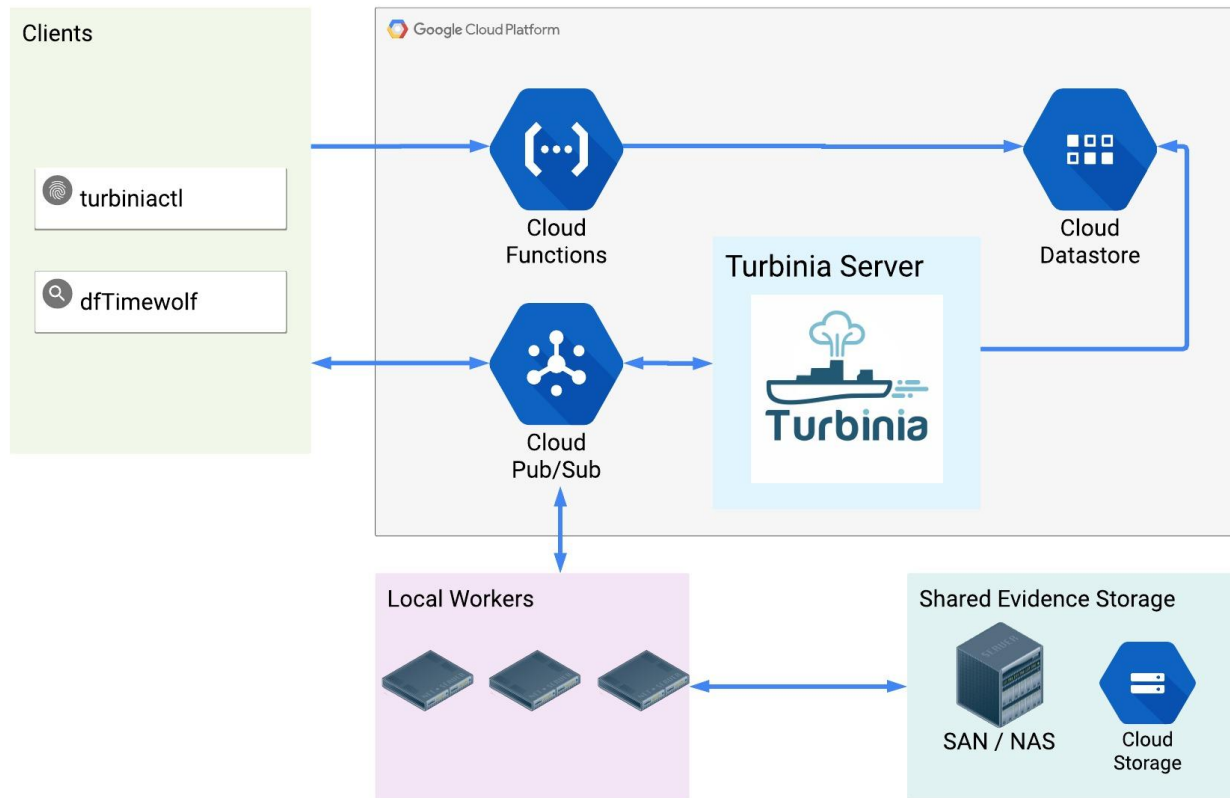 - All processed data stays local

- **Local**
  - No cloud dependencies
  - Uses Celery / Kombu / Redis
  - Contributed by Facebook (Eric Zinnikas, ericz.com)

Installation Types Pros/Cons

	😄 Pros	😰 Cons
Cloud	• No infrastructure management	• Evidence may need to be uploaded
Hybrid	• Shifts costs • No server management • Data stays local	• Local machine management
Local	• No cloud dependencies • Data stays local	• Local machine management • Local service management (Celery, Kombu, Redis)

What is Evidence?

- Evidence can be anything we want to process
    - E.g. RawDisk, GoogleCloudDisk, PlasoFile, etc
- Definitions in Python
- Tasks can generate new Evidence, which may be re-processed
- Evidence as seen by Client/Server is just metadata
- Actual data stored in shared storage or Google Cloud Storage

Pre/Post-Processors

- Pre-processors make Evidence available to Tasks

  - Mounting images and attaching cloud disks, etc.

  - CloudPersistentDisk → RawDisk

- Post-processors clean-up

- Evidence can be "stacked"

  - GoogleCloudDiskRawEmbedded Evidence

  - Pre-processor for outer Cloud Disk attaches outer disk

  - Pre-processor for RawDisk mounts inner raw disk

Output Manager

- Some Evidence types are "copyable"

  - PlasoFile, PlasoCSVFile, TextFile, etc

- Copyable Evidence can be automatically pulled from storage

  - Google Cloud Storage

  - Copyable generated Evidence can also be copied back

- Non-copyable Evidence requires shared storage

A typical Turbinia workflow

- Client sends processing request to server

- Server schedules Tasks from Jobs that can process that Evidence

- Workers from the pool run Tasks to process the Evidence
  a. Tasks read Evidence from shared storage or copied from cloud storage
  b. Task runner pre-processes the Evidence
  c. Task does actual processing
  d. Task generates new Evidence objects (e.g. RawDisk → PlasoFile)
  e. Tasks return this new Evidence to the Server to be processed

Job Graph
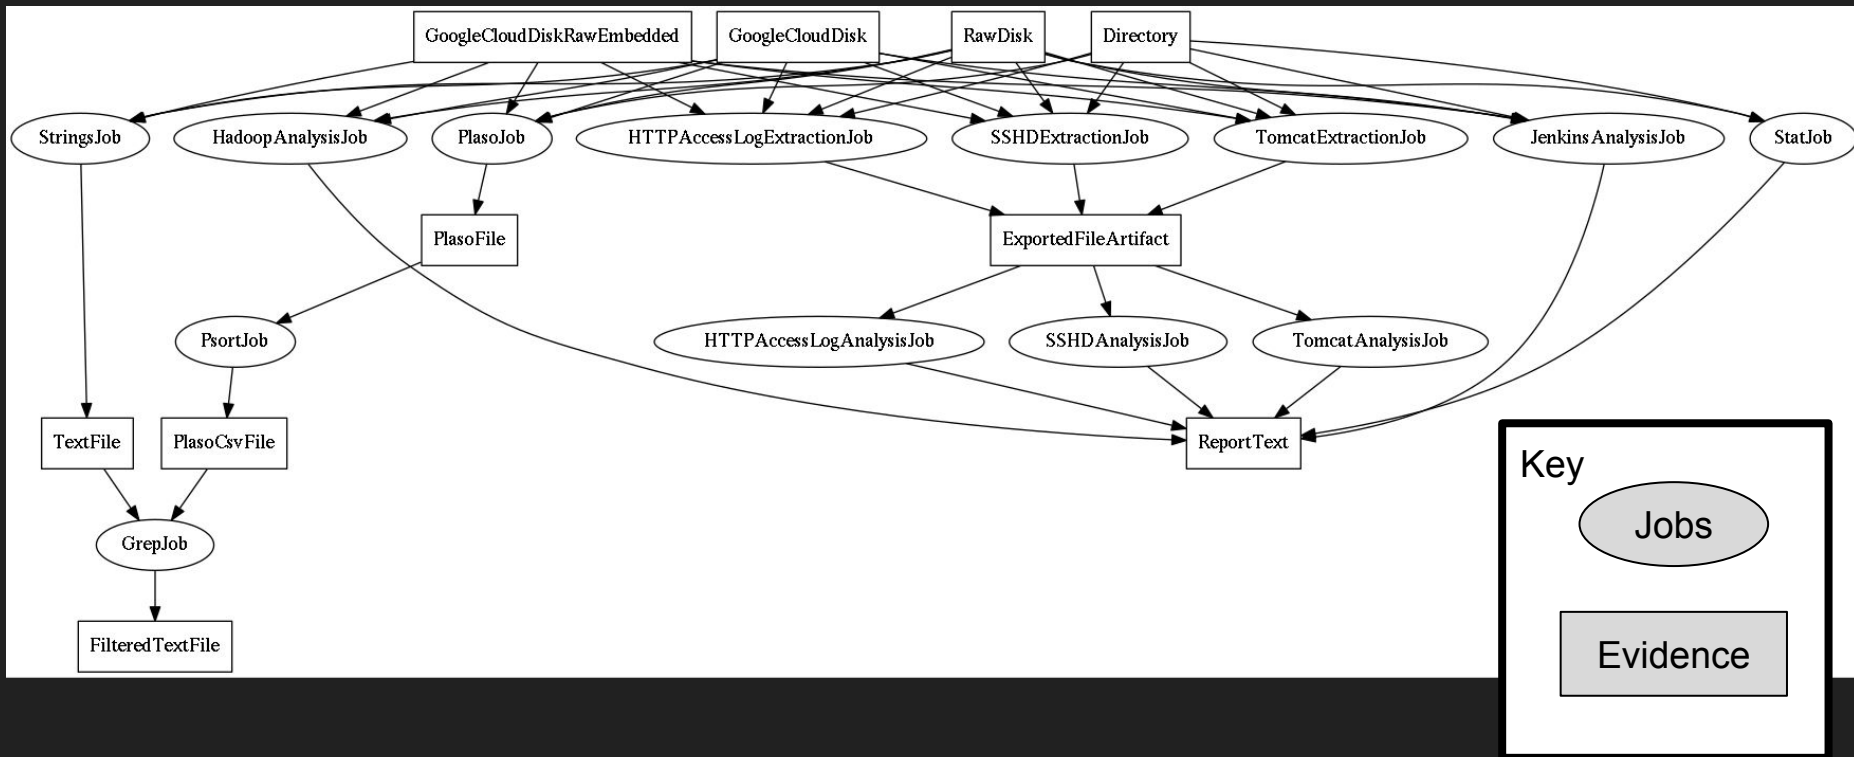
Turbinia analysis modules

Analysis modules make sense of forensic evidence.

- Look in logs for **successful exploitation** of {Wordpress,Tomcat,Jenkins}.

- Highlight **insecure** {SSH,Redis,MongoDB} **configuration** files

- Python code in **tasks** holds the analysis logic.

- Extracts Artifacts and then analyzes them

CREATING NEW TASKS IS EASY

- Simple execution tasks can be 10-15 lines of actual code

- Documentation at docs/developing-new-tasks.md

```python
output_evidence = TextFile()
base_name = os.path.basename(evidence.local_path)
output_file_path = os.path.join(
    self.output_dir, '{0:s}.ascii'.format(base_name))
output_evidence.local_path = output_file_path

cmd = 'strings -a -t d {0:s} > {1:s}'.format(
    evidence.local_path, output_file_path)
result.log('Running strings as [{0:s}]'.format(cmd))
self.execute(
    cmd, result, new_evidence=[output_evidence], close=True, shell=True)
```

Turbinia Scope

- Orchestration happens externally
    - dfTimewolf
- Intentionally limited privs
- Push evidence instead of pull

Turbinia Recent Updates

- Several new analysis tasks

- dfTimewolf integration published

- Python 3 support

- New Job manager

- Other clean-up and fixes: Task locking, Evidence and processor refactoring, better Task error handling, libcloudforensics fixes, etc.

- Initial structure for terraform support

Turbinia Next Steps

- Disk Volume Enumeration

- Encrypted disk support

- More Tasks in general (they're easy to write!)

- Cloud bootstrapping automation

- Reporting

- Recipes

- Killing off Python 2

Big Picture

- Hunting: **GRR**

- Gathering: **dfTimewolf**

- Processing: **Turbinia**
    - Via: Plaso, libyal, TSK, etc

- Analysis: **Timesketch**

The scenario

DISCLAIMER

None of what I'm about to talk about is true

(except for the demos)

The victim

Greendale Poly - the most famous **fictitious** university

Everyone's on semester break when... someone gets a tip.

Suspicious domain reported by admin:

grendale.xyz

Greendale just migrated to the cloud...

Honing in on the INITIAL tip...

- Typosquatting on **grendale.xyz**

- Looks **targeted**... Let's look for related artifacts

- Let's see what our cloud forensics options are...

Dem0 (dfTimeWolf gcp_forensics)

DemO (Turbinia)

Forensics in the cloud

dfTimewolf

plz copy disk(s)

greendale-student-{0..n}

*disk ID
123-copy*

*plz forensicate
disk 123-copy*

Turbinia

timeline.plaso

*Grab your
plaso file from
gs://123.plaso*

Time
sketch

Dem0 (dfTimewolf with Turbinia)

Timesketch

A new host was compromised!

Netflow shows connections between the compromised GCP host and a web server... running WordPress.

Let's see if Turbinia's Wordpress analysis module can help.

Dem0 (Turbinia analysis plugins, WP compromise)

Disaster averted!

- Payload was a keylogger; no traces of lateral movement found.
  - Plus, Greendale uses 2FA tokens for all sensitive access

- Attacker's objective was likely to disrupt the launch of Greendale's new PhD program in AC flow study.

What else can these tools do?

- GRR
  - Some host timelining, run custom Python scripts
- Plaso
  - Focus processing on specific user-selected artifacts
- dfTimewolf
  - Chain any system with an API into your workflow
- Timesketch
  - Histogram and heatmap view to view data differently, graphs
- Turbinia
  - Repetitive, parallelizable tasks

Key takeaways

Tools that you might have a place in **your** ecosystem

Used daily by IR teams at Google

Contributions are **encouraged**

Apache 2 license

Where to find us

Slack channel

https://open-source-dfir.slack.com

http://join-open-source-dfir-slack.herokuapp.com/

GRR

github.com/google/grr

Plaso

github.com/log2timeline/plaso

dfTimewolf

github.com/log2timeline/dftimewolf

Turbinia

github.com/google/turbinia

github.com/google/timesketch

The End!