


# Certificate Transparency

Supporting Critical Internet Infrastructure


[certificate.transparency.dev](https://certificate.transparency.dev)

# Brought to you by:




**HTTP**

**HTTPS**




Percentage of Web Pages Loaded by Firefox Using HTTPS  
Source: Mozilla





Chrome page loads over HTTPS (with TLS)  
Source: Google



**95-99%**

Chrome browsing time over HTTPS (with TLS)  
Source: Google






**TLS**






ISRG Root X1



R3



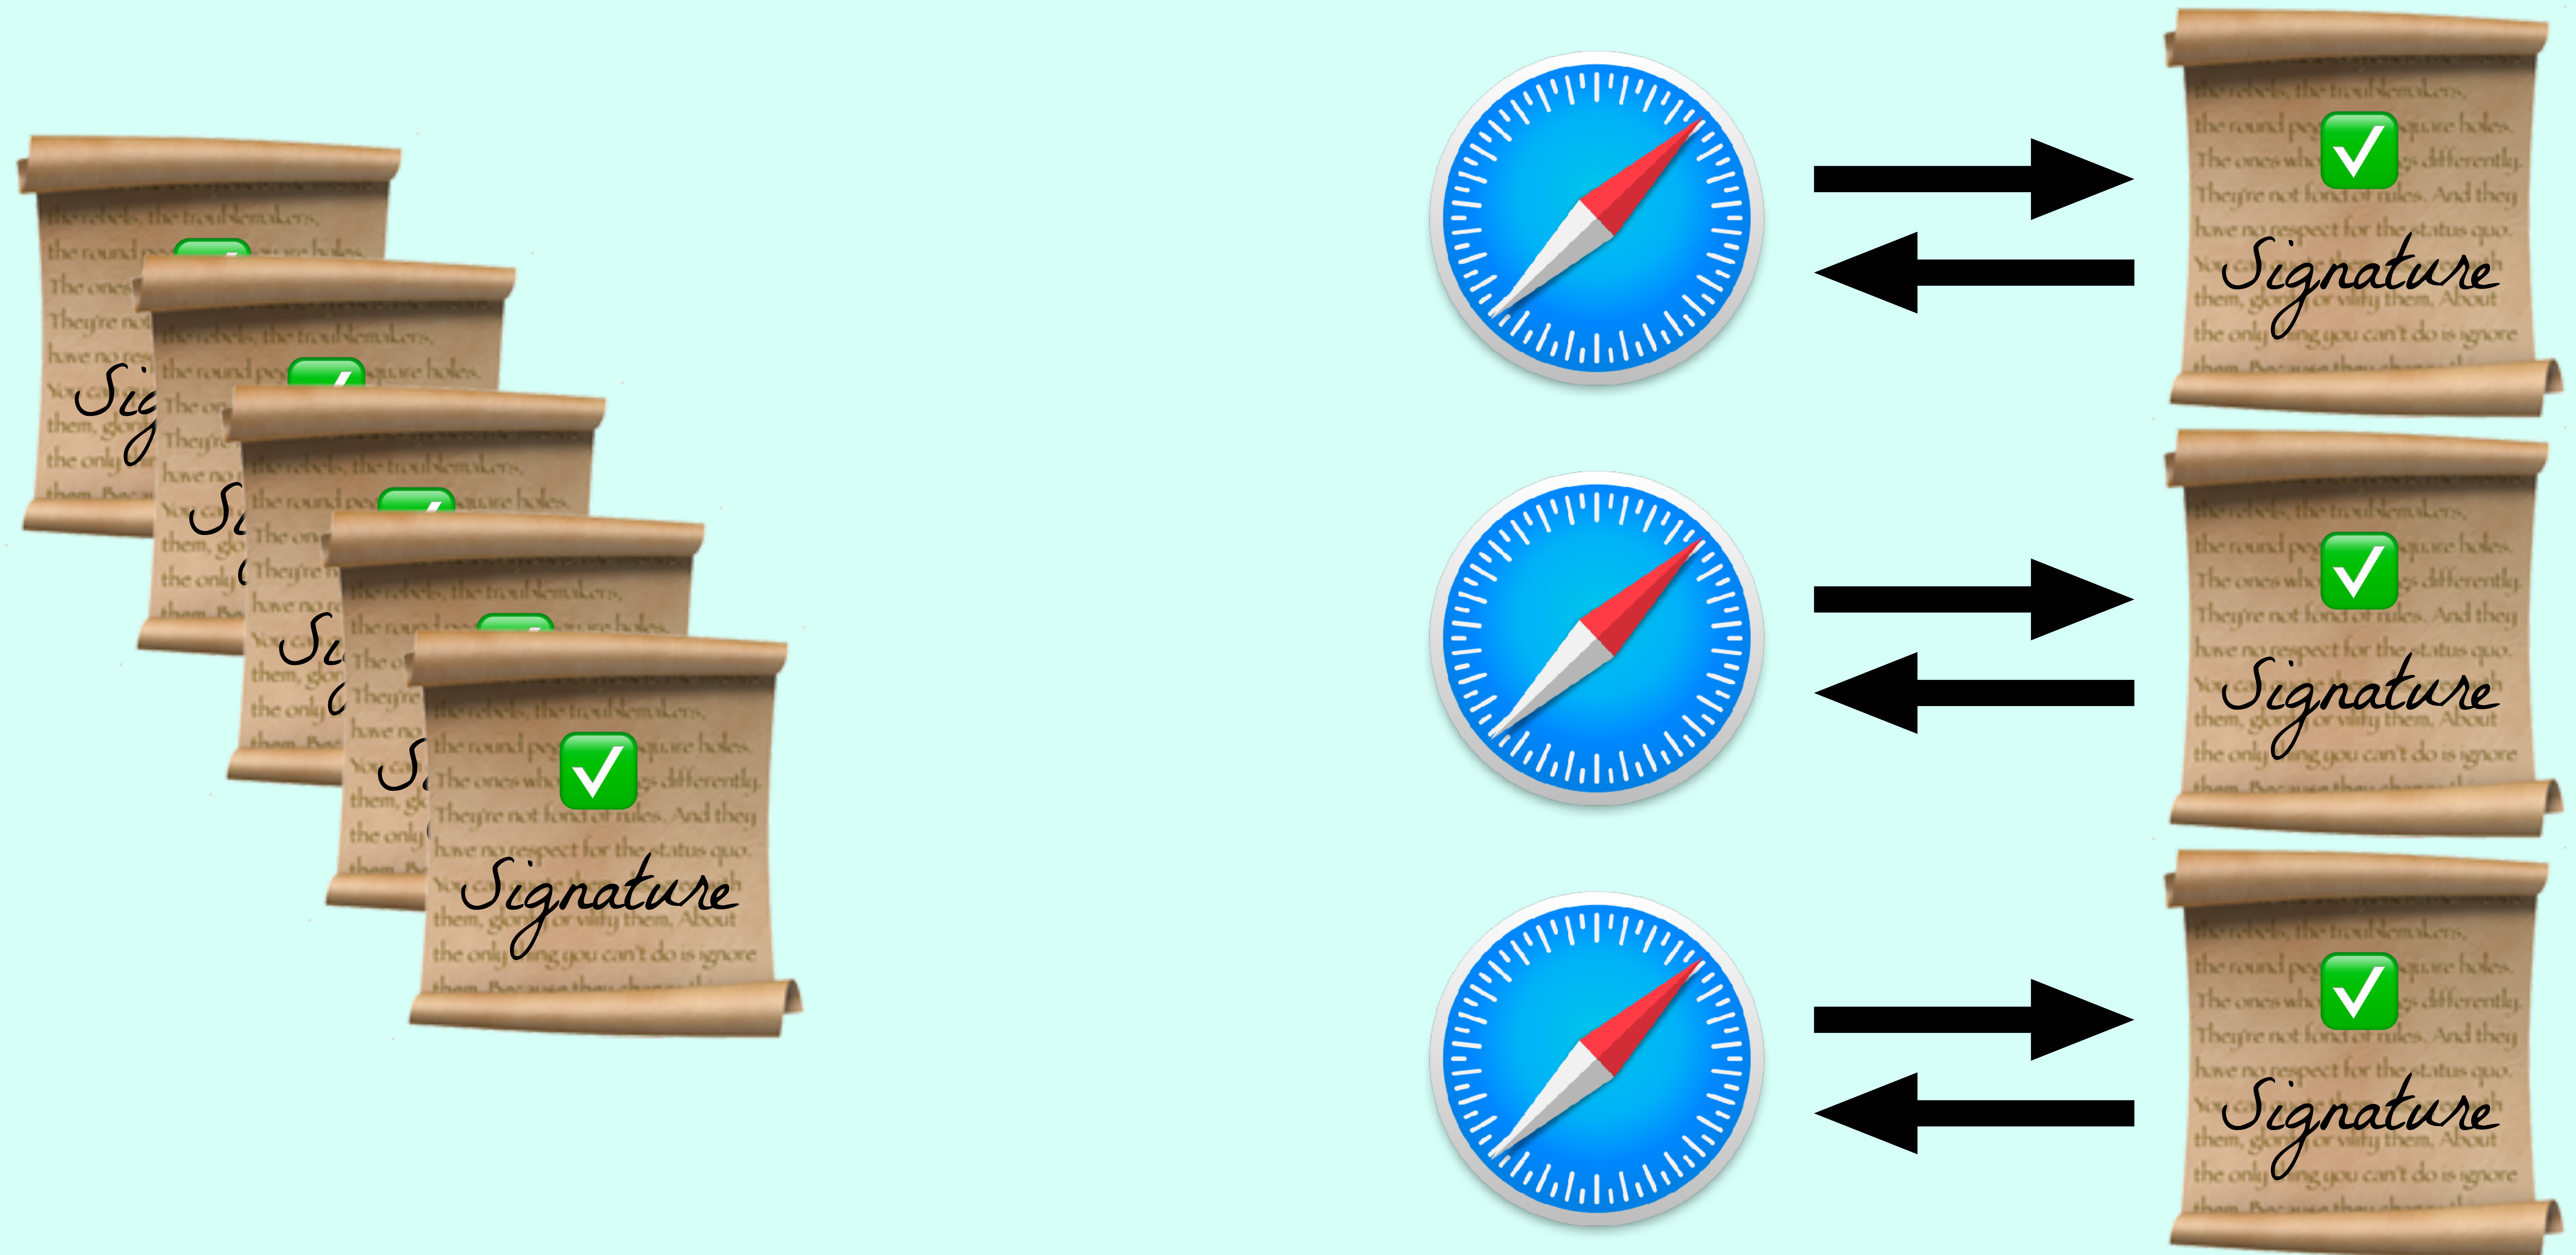
swinog.ch



# WebPKI




# RPKI vs WebPKI






# WebPKI Size

- Over 4.8 **B**illion Certificates stored
- Over 0.5 Billion Active / Unexpired
- Over 250,000 new certificates per hour



Currently Active / Unexpired

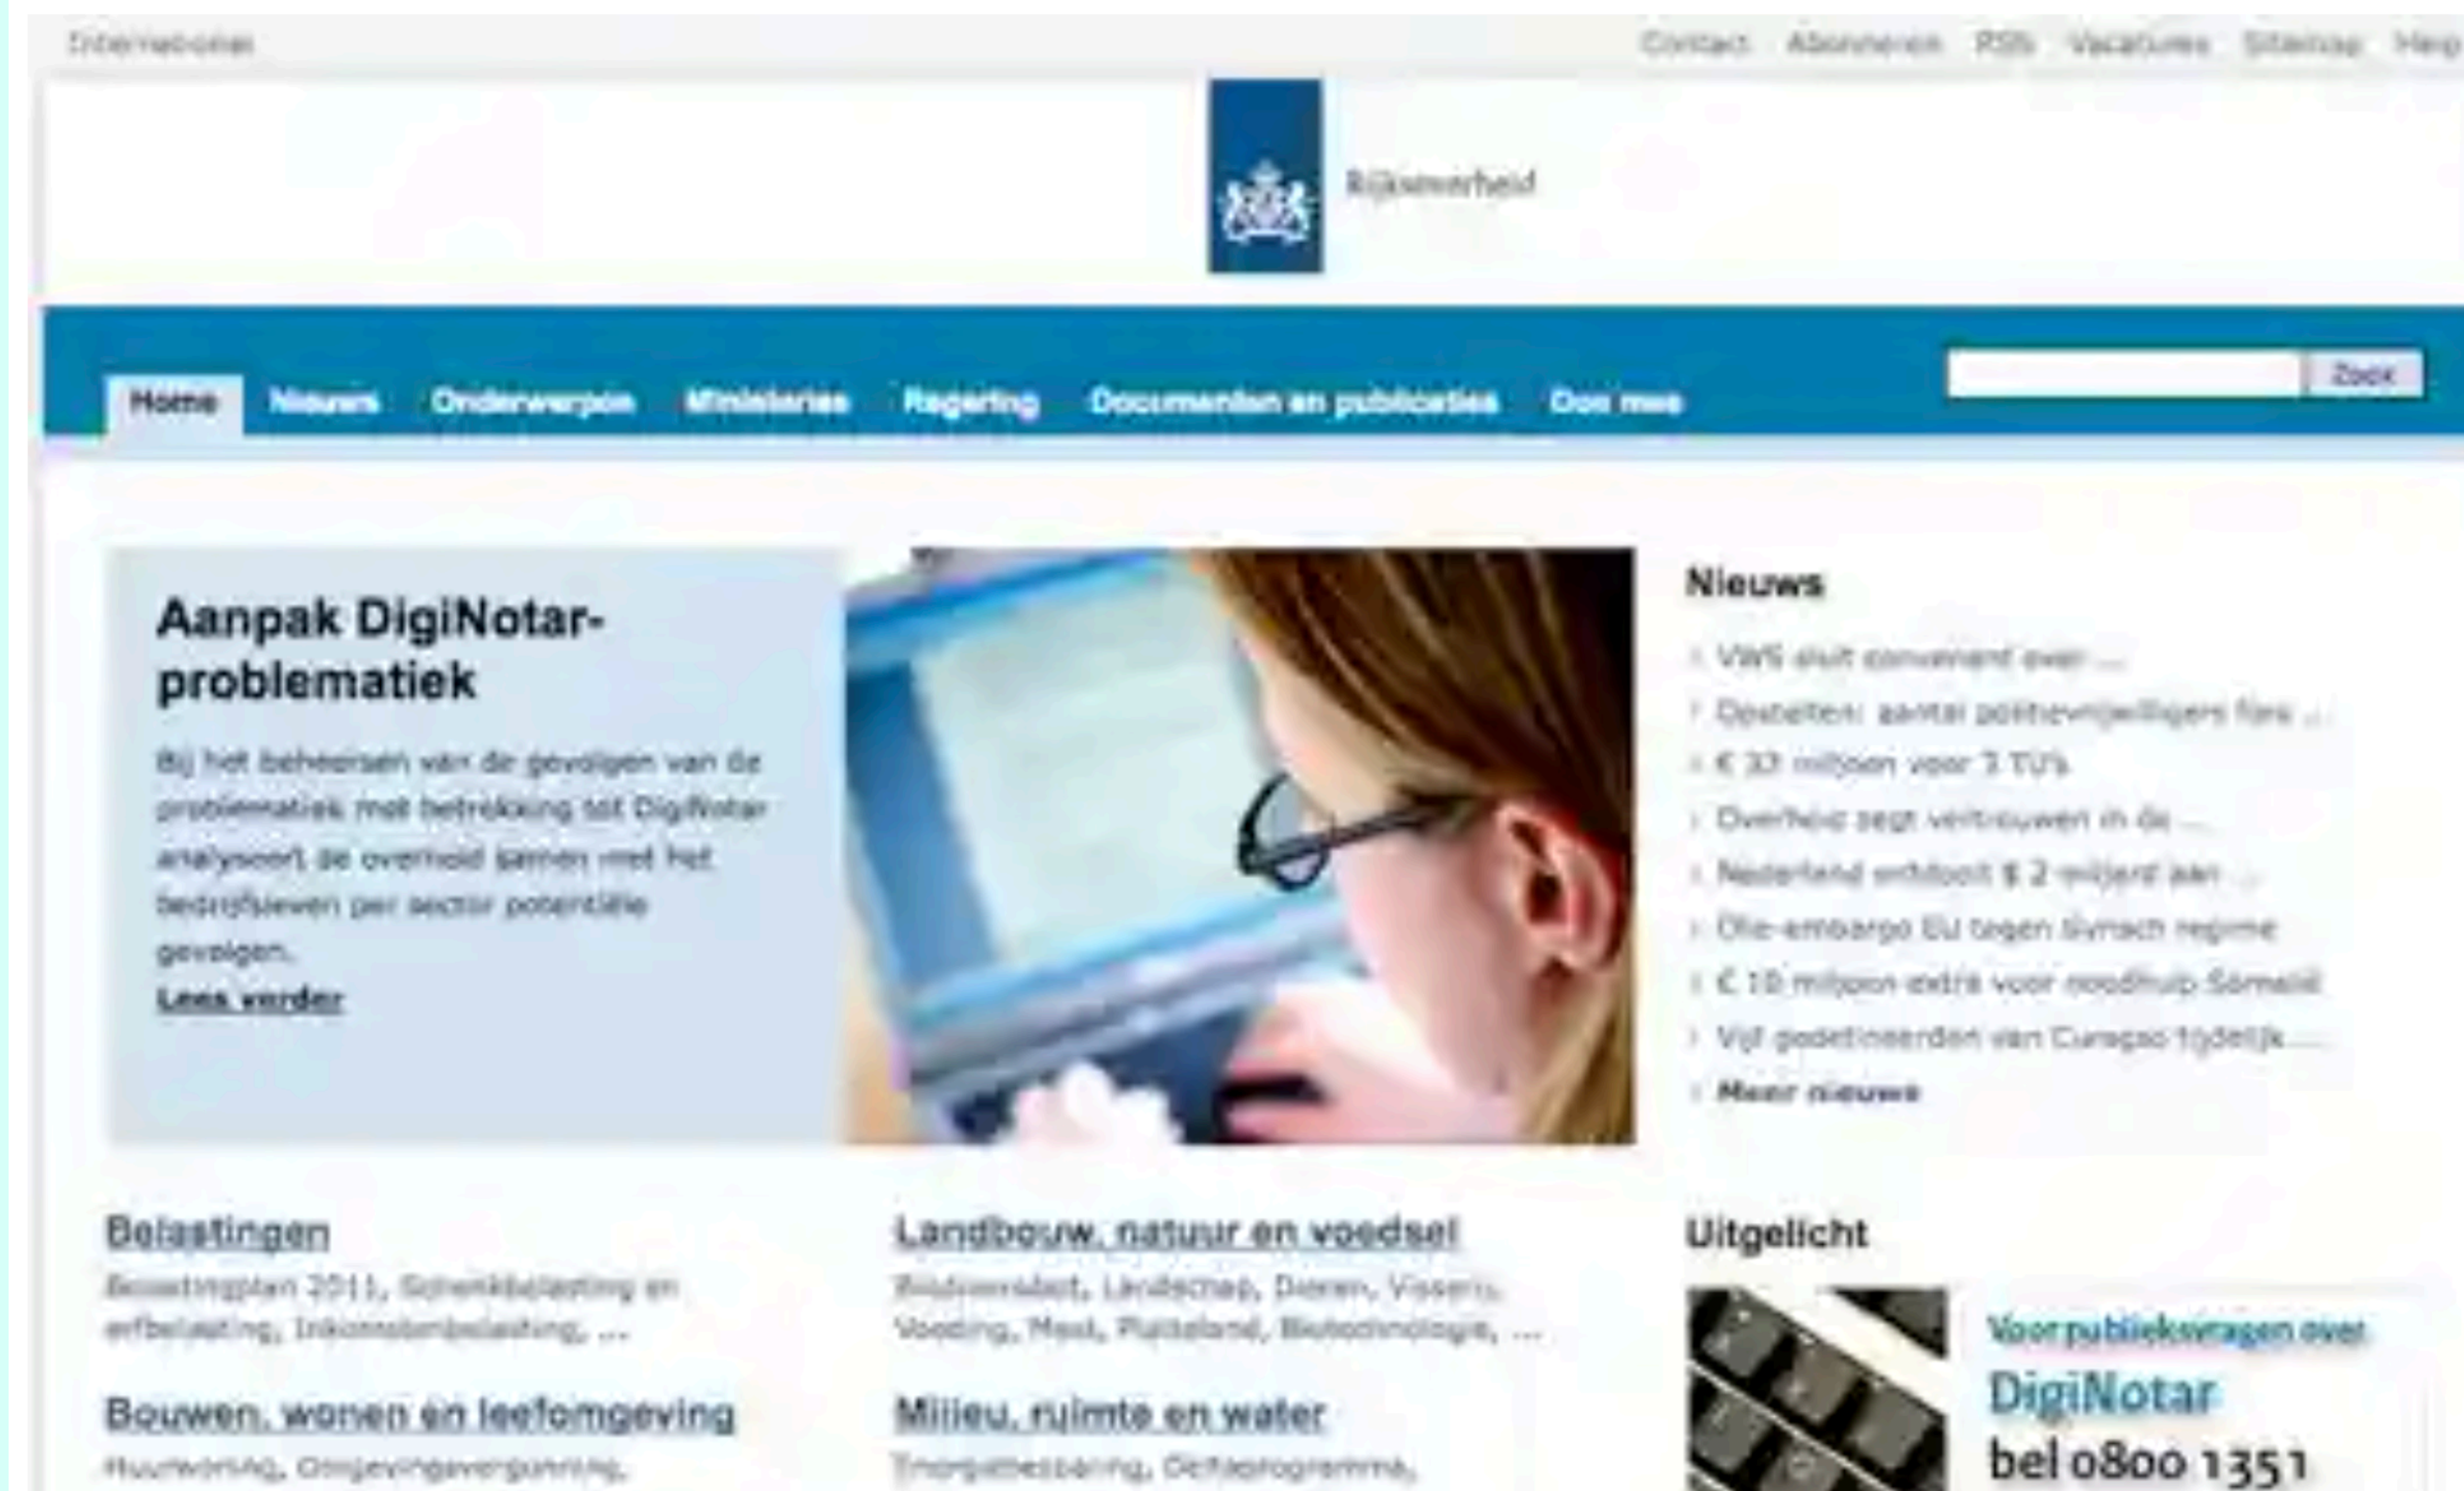


Total collected



# DigiNotar SSL certificate hack amounts to cyberwar, says expert

**Dutch government revokes certificates used for all its secure online transactions, while CIA, Google, Microsoft and others affected by hack called 'worse than Stuxnet'**




📷 The Dutch government has revoked all trust in digital certificates issued by DigiNotar



# Certificate Transparency

CT Log








## Your connection is not private

Attackers might be trying to steal your information from **fincen.gov** (for example, passwords, messages, or credit cards). [Learn more](#)

NET::ERR\_CERTIFICATE\_TRANSPARENCY\_REQUIRED

Advanced

Back to safety



digicert®



SECTIGO®





**Phil, what does it take to  
operate one of these?**



# Operating CT Logs

Wins and Woes


# Who am I?

- [phil@letsencrypt.org](mailto:phil@letsencrypt.org)
- [github.com/pgporada](https://github.com/pgporada)
- [linkedin.com/in/philporada](https://linkedin.com/in/philporada)



# What is Let's Encrypt?

Let's Encrypt Certificates Issued Per Day





# Blogs we've written about CT

## 2019

Introducing Oak, a Free and Open Certificate Transparency Log

<https://letsencrypt.org/2019/05/15/introducing-oak-ct-log.html>

How Let's Encrypt Runs CT Logs



<https://letsencrypt.org/2019/11/20/how-le-runs-ct-logs.html>

## 2022

Nurturing Continued Growth of Our Oak CT Log

<https://letsencrypt.org/2022/05/19/nurturing-ct-log-growth.html>

# Logs? Shards? Oh my!



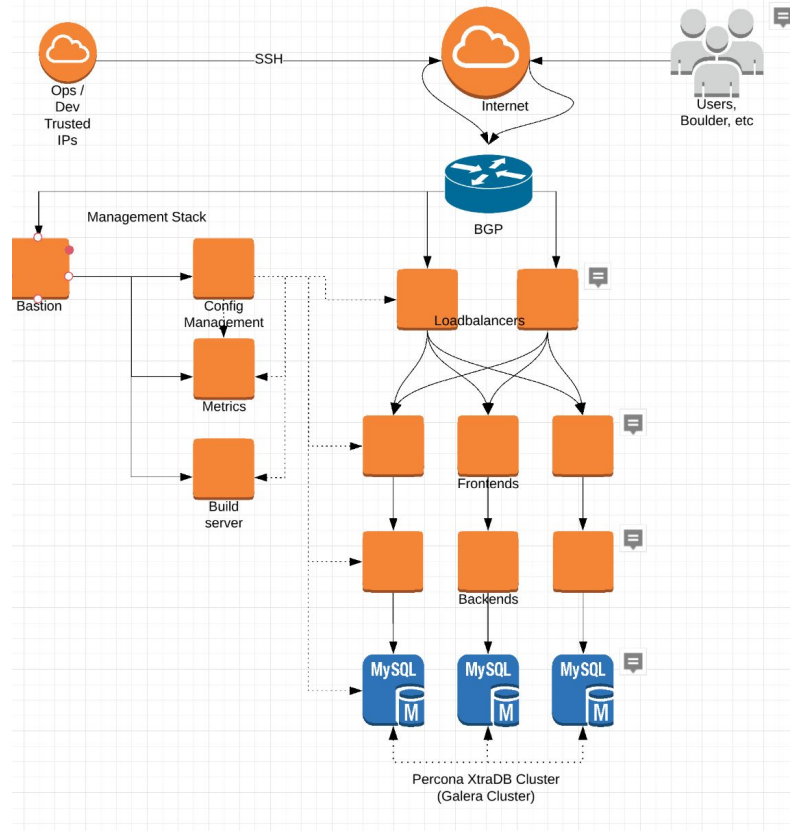
# How does a CT log benefit a CA?

## Embedded SCTs

Log ID	B7:3E:FB:24:DF:9C:4D:BA:75:F2:39:C5:BA:58:F4:6C:5D:FC:42:CF:7A:9F:35:C4:9E:1D:0...
Name	Let's Encrypt Oak 2023
Signature Algorithm	SHA-256 ECDSA
Version	1
Timestamp	Thu, 01 Jun 2023 23:20:23 GMT
Log ID	7A:32:8C:54:D8:B7:2D:B6:20:EA:38:E0:52:1E:E9:84:16:70:32:13:85:4D:3B:D2:2B:C1:3...
Name	Cloudflare "Nimbus2023"
Signature Algorithm	SHA-256 ECDSA
Version	1
Timestamp	Thu, 01 Jun 2023 23:20:23 GMT

# Initial Failed Architectures


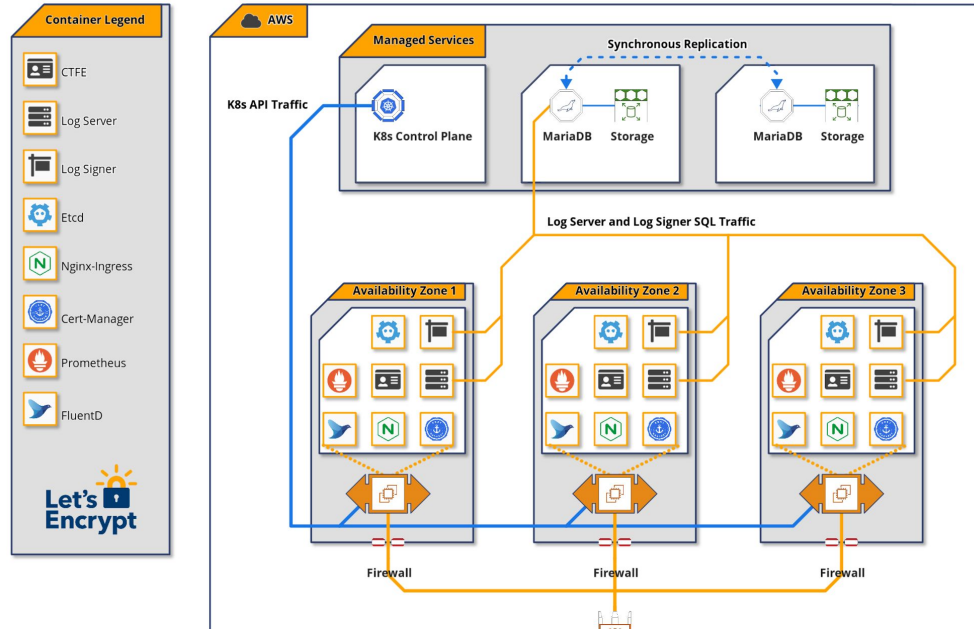
## CT Architecture



```
1 galera:
2 # Path to Galera library
3 wsrep_provider: /usr/lib64/galera3/libgalera_smm.so
4
5 # Cluster connection URL contains IPs of nodes
6 # If no IP is found, this implies that a new cluster needs to be created,
7 # in order to do that you need to bootstrap this node
8 wsrep_cluster_address: gcomm://10.88.162.21,10.88.162.23,10.88.162.29
9
10 # In order for Galera to work correctly binlog format should be ROW
11 binlog_format: ROW
12 default_storage_engine: InnoDB
13 wsrep_slave_threads: 8
14 wsrep_log_conflicts: NO_VAL
15
16 # This changes how InnoDB autoincrement locks are managed and is a requirement for Galera
17 innodb_autoinc_lock_mode: 2
18
19 wsrep_node_address: {{ salt['grains.get']('ipv4')[0] }}
20 wsrep_cluster_name: birch-ct
21 wsrep_on: ON
22 pxc_strict_mode: ENFORCING
23 wsrep_sst_method: xtrabackup-v2
24
25 # Encrypted as sstuser:GENERATEDPASSWORD
26 wsrep_sst_auth: |
27     -----BEGIN PGP MESSAGE-----
28     -----END PGP MESSAGE-----
29 skip_external_locking: NO_VAL
30 wsrep_retry_autocommit: 5
31 wsrep_sst_donor: birch-db-99
```




# Current Functioning Architecture




```
$ kubectl get pods -n oak-2023
```

NAME	READY	STATUS	RESTARTS	AGE
oak-2023-etcd-operator-etcd-operator-etcd-operator-77968f6j66vs	1/1	Running	0	42m
prometheus-mysqld-exporter-deployment-65c58775f7-sp2kn	1/1	Running	2 (147d ago)	351d
trillian-ctfe-deployment-6fc9cff9d6-8f2nq	1/1	Running	0	15md
trillian-ctfe-deployment-6fc9cff9d6-f4rjw	1/1	Running	0	42m
trillian-etcd-cluster-jrfm7b4nqp	1/1	Running	0	351d
trillian-etcd-cluster-n8fjkgpht8	1/1	Running	0	41m
trillian-etcd-cluster-pld2rzj8q8	1/1	Running	0	90d
trillian-logserver-deployment-785d5c444d-b5kd8	1/1	Running	0	21d
trillian-logserver-deployment-785d5c444d-qpj9v	1/1	Running	0	21d
trillian-logsigner-deployment-d68cc6bf7-fnmvl	1/1	Running	0	21d
trillian-logsigner-deployment-d68cc6bf7-ruc3f5	1/1	Running	0	21d

# Tree Growth Rate




# Request rate by endpoint



# A Fun Incident

```
"log":"W0808 19:22:35.000186      1 tree_storage.go:81] Failed to set strict mode on mysql db: commands out of sync. Did you run multiple statements at once?",
```



Error reproduction: [github.com/koshatul/go-mysql-sync-issue](https://github.com/koshatul/go-mysql-sync-issue)

# Cost of running our CT logs

## Human

- 1 - 2x SREs spending ~3 months worth of time over the course of a year

## Compute

- Compute nodes are basically commodity hardware.
- Storage and RAM for compute nodes is also negligible. Just enough to run applications.

## Database


- This will be your pain point. Horizontal scaling continues to be the solution, and is straightforward to apply
- Faster storage will give better log performance to a point when data is read from disk.



# Lessons Learned and Takeaways

- Have a testing log so you don't prematurely ruin your production log.
- Logs are ephemeral. When your log fails, root cause why and build a new better log with the lessons learned.
- ```
sql> SET GLOBAL event_scheduler=ON;  
sql> CREATE EVENT analyze_subtree ON SCHEDULE EVERY 1 DAY DO ANALYZE TABLE Subtree;
```
- Separation of concerns: run each binary in a different container, VM, or physical host. You're after reliability.
- The log\_signers (sequencers) perform an etcd election to determine which cluster member will communicate with the database. Make an alert if more than 1 cluster member has mastership for a particular shard because it will indicate a split brain scenario and cause an incident. We've been there.
- Build a rate limiting story to protect your log.
- We don't run database backups for CT logs.
- More automation is better than less. That's why we put our logs in Kubernetes.

**Cool!**






digicert®



**SECTIGO®**

Google



**We need more CT Logs!**



# Setting up a CT Log

- Apple & Google have guidelines & requirements
- You can apply like this for consideration
- You can join the community mailing list



# You're not alone!

Use the mailing list, or feel free to contact [morectlogs@daknob.net](mailto:morectlogs@daknob.net) for discussion, help, insights, advice, or anything we can do at any phase of the project :)