Cassandra Secure-by-default

Ayushi Singh
ApacheCon North America, 2022

Hello! I am

- Ayushi Singh
- Sr. Software Engineer
- Netflix (~2 years)
- Online Datastores
- Ex-PayPal (4 years)

Acknowledgement

Sumanth Pasupuleti
co-authored this feature

Objective

How we have used Cassandra’s pluggable nature to integrate our own authentication and authorization.

Agenda

- Introduction
- Motivation
- Solution
- Performance
- Benefits
- Issues
- Demo

Authentication

Authorization

Motivation

- Managing passwords is not very secure.
- Add temporary node or remove a node
- Less visibility in roles and permissions.
- Incompatibility with unified authentication system

Components

Metatron authenticator

- Cryptographic identity
- Metatron Secrets
- Metatron TLS

Gandalf authorizer

- Framework/Service for authorization policies
- Gandalf agent
- Gandalf portal

Components

Spinnaker

- Open-source, multi-cloud continuous delivery platform
- Provides application management and deployment

Fast properties

- Dynamic properties
- Various scope of fast properties.

Solution

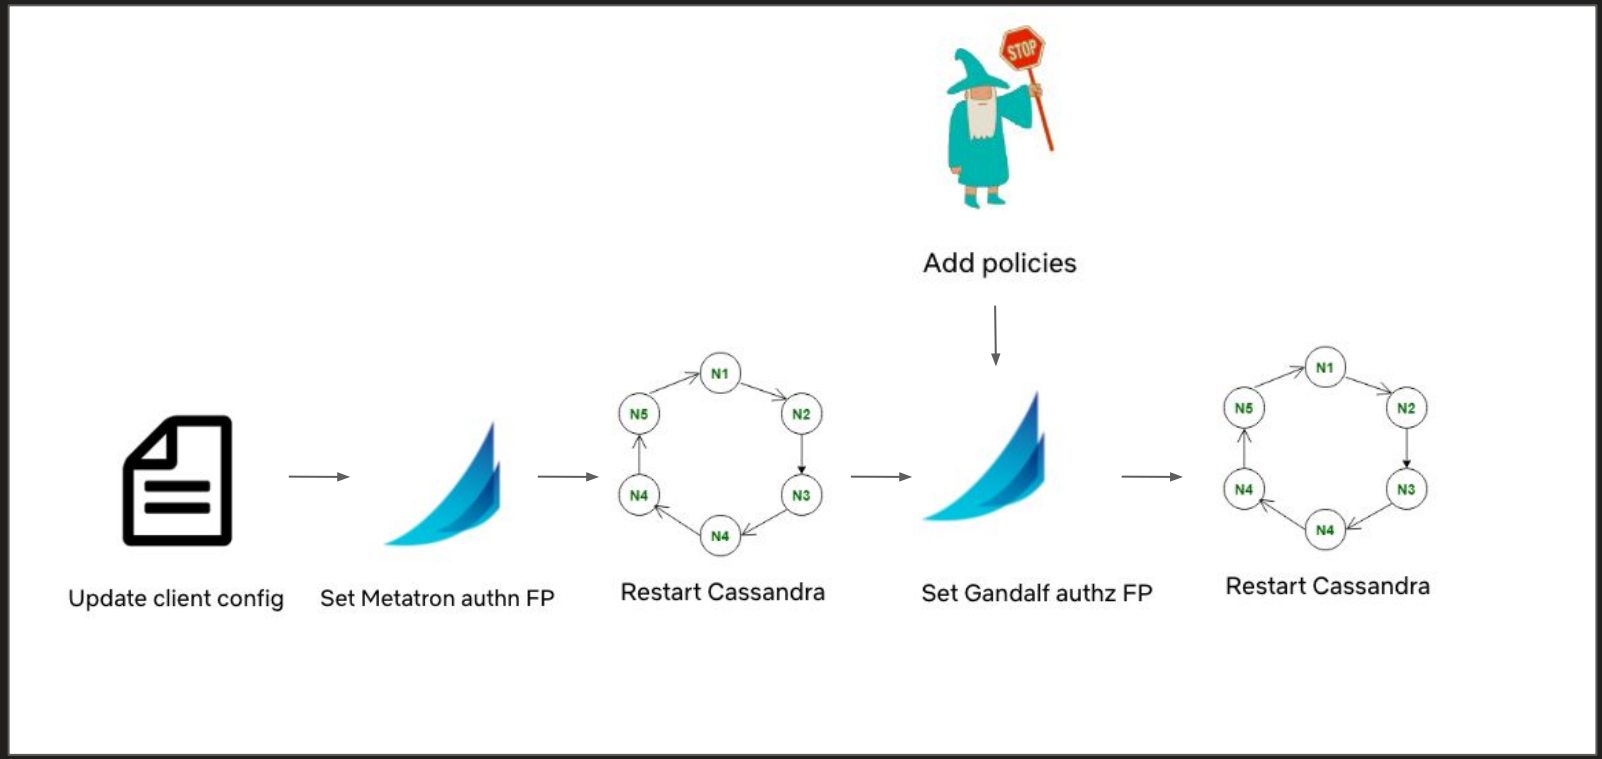
STOP
Add policies
N1
N2
N3
N4
N5
Update client config
Set Metatron authn FP
Restart Cassandra
Set Gandalf authz FP
Restart Cassandra

Fast properties to Cassandra

Authentication

- Implements IAuthenticator class of Cassandra
- Certificates get downloaded on instance boot
- SASL based

Authentication

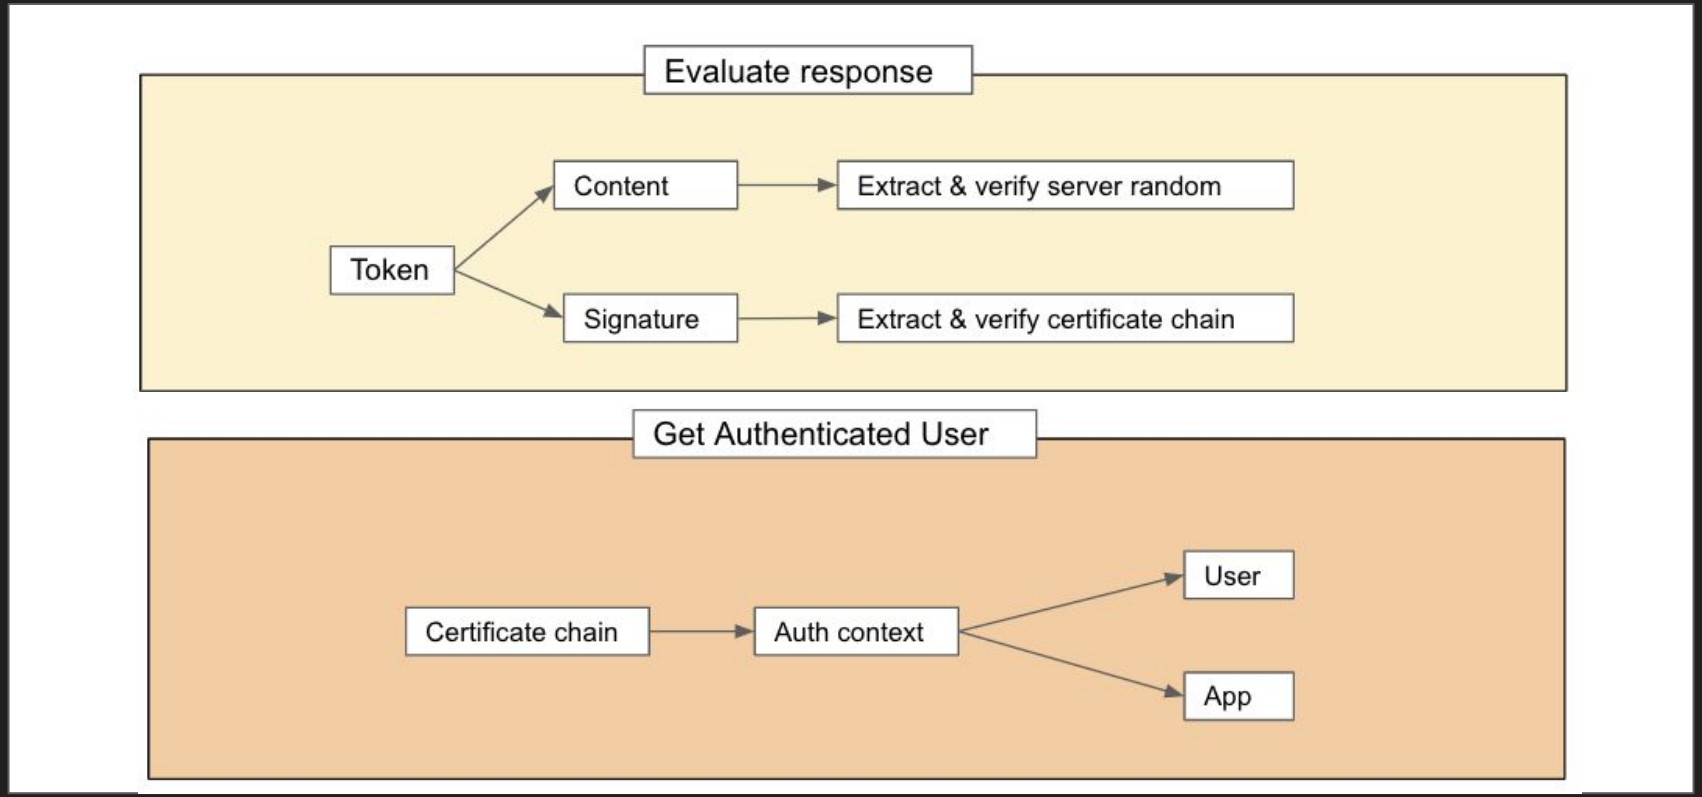

Authorization

- Implements IAuthorizer class of Cassandra
- Gandalf agent is present on every node
- Policy added to Gandalf portal
- Policy name: <App name>-<resource>-<type>-<env>
- Narrow-broad approach
  - Keyspace
  - Data

Authorization

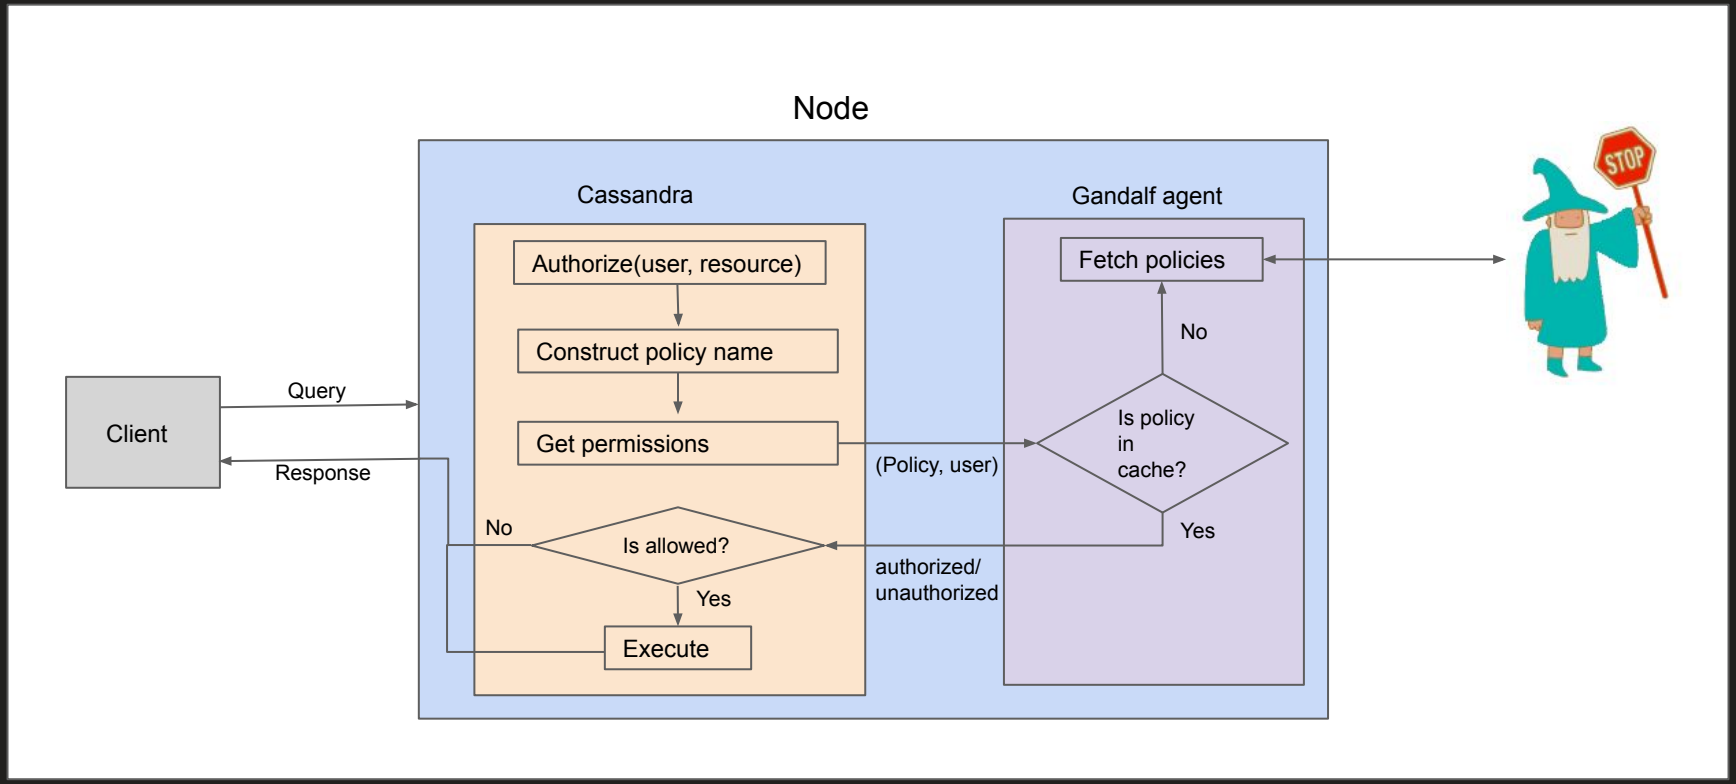

Cache Update

- Periodically query the Distributor for updates.
- Patch of changes
- Dynamic updates.
- Application can actively refresh the cache

Performance

- Gandalf agent
  - Max: ~0.11ms Avg: ~0.09 ms

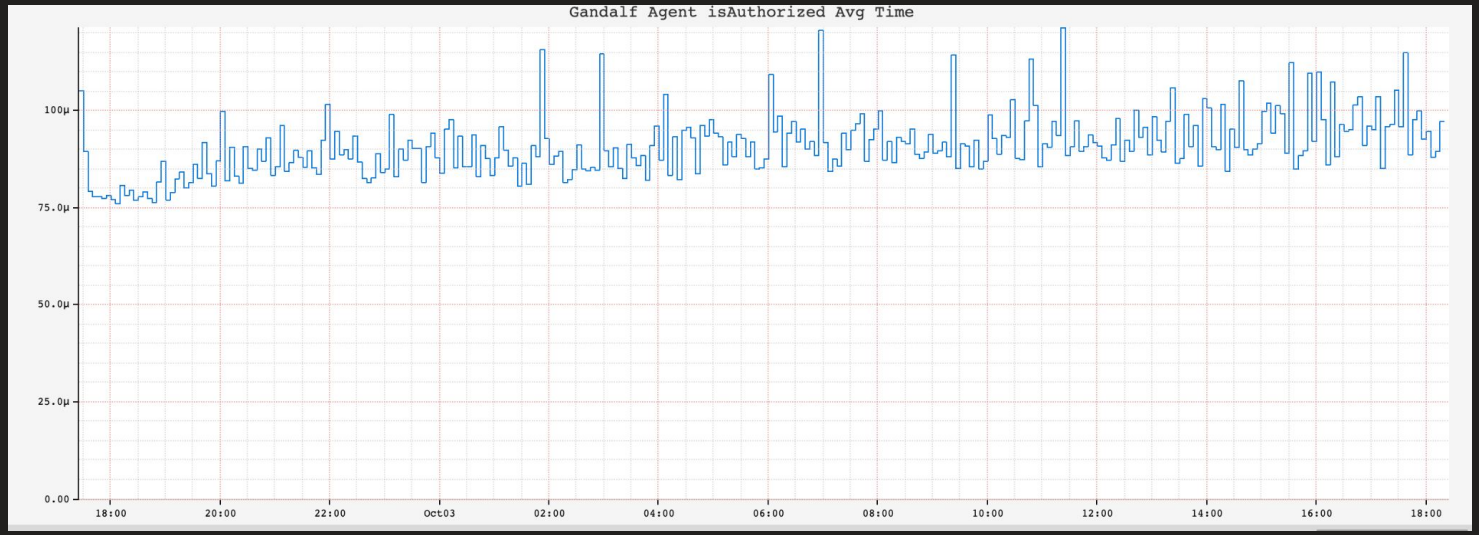

Performance

Cassandra

	Without authn/authz	Traditional authn/authz	New authn/authz
Read	14.8 ms	15.7 ms	15.3 ms
Write	2.23 ms	2.28 ms	2.29 ms

Performance

- Propagating policies to Gandalf can take time.

Benefits

- Hot reload of server certificates
- Rolling restart, no downtime
- Easy to operate
- High visibility
- Modular design
- No need to maintain fork
- E2E integration with Netflix ecosystem

Issues!!

- Increase in authz errors on Gandalf
- Each query verified for read/write/schema
- Unable to switch authenticators without downtime

Future work

- Transition mode
- Optimize authorization interface

Demo

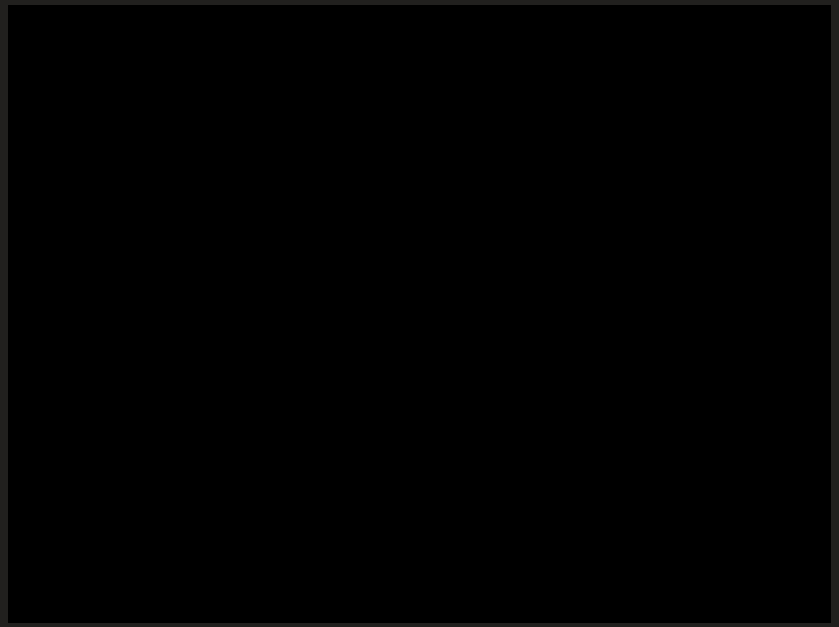

Thank you