Taming Tenancy, Cost and Architecture at Collibra

Through OpenTelemetry and our Telemetry Backbone

Alex Van Boxel

Almost 30 years in the sector

Mostly as Software Engineer

Web - 3D - Middleware - Mobile - Big Data

More recent as Architect

Data - SRE - Infrastructure

Community

Apache Beam contributor

OpenTelemetry Collector contributor

Collibra

Principal Systems Architect

Collibra

A data intelligence platform powered by active metadata

AI Governance

Data Catalog

Data Governance

Data Lineage

Data Notebook

Data Privacy

Data Quality & Observability

Protect

"How much does X cost?"

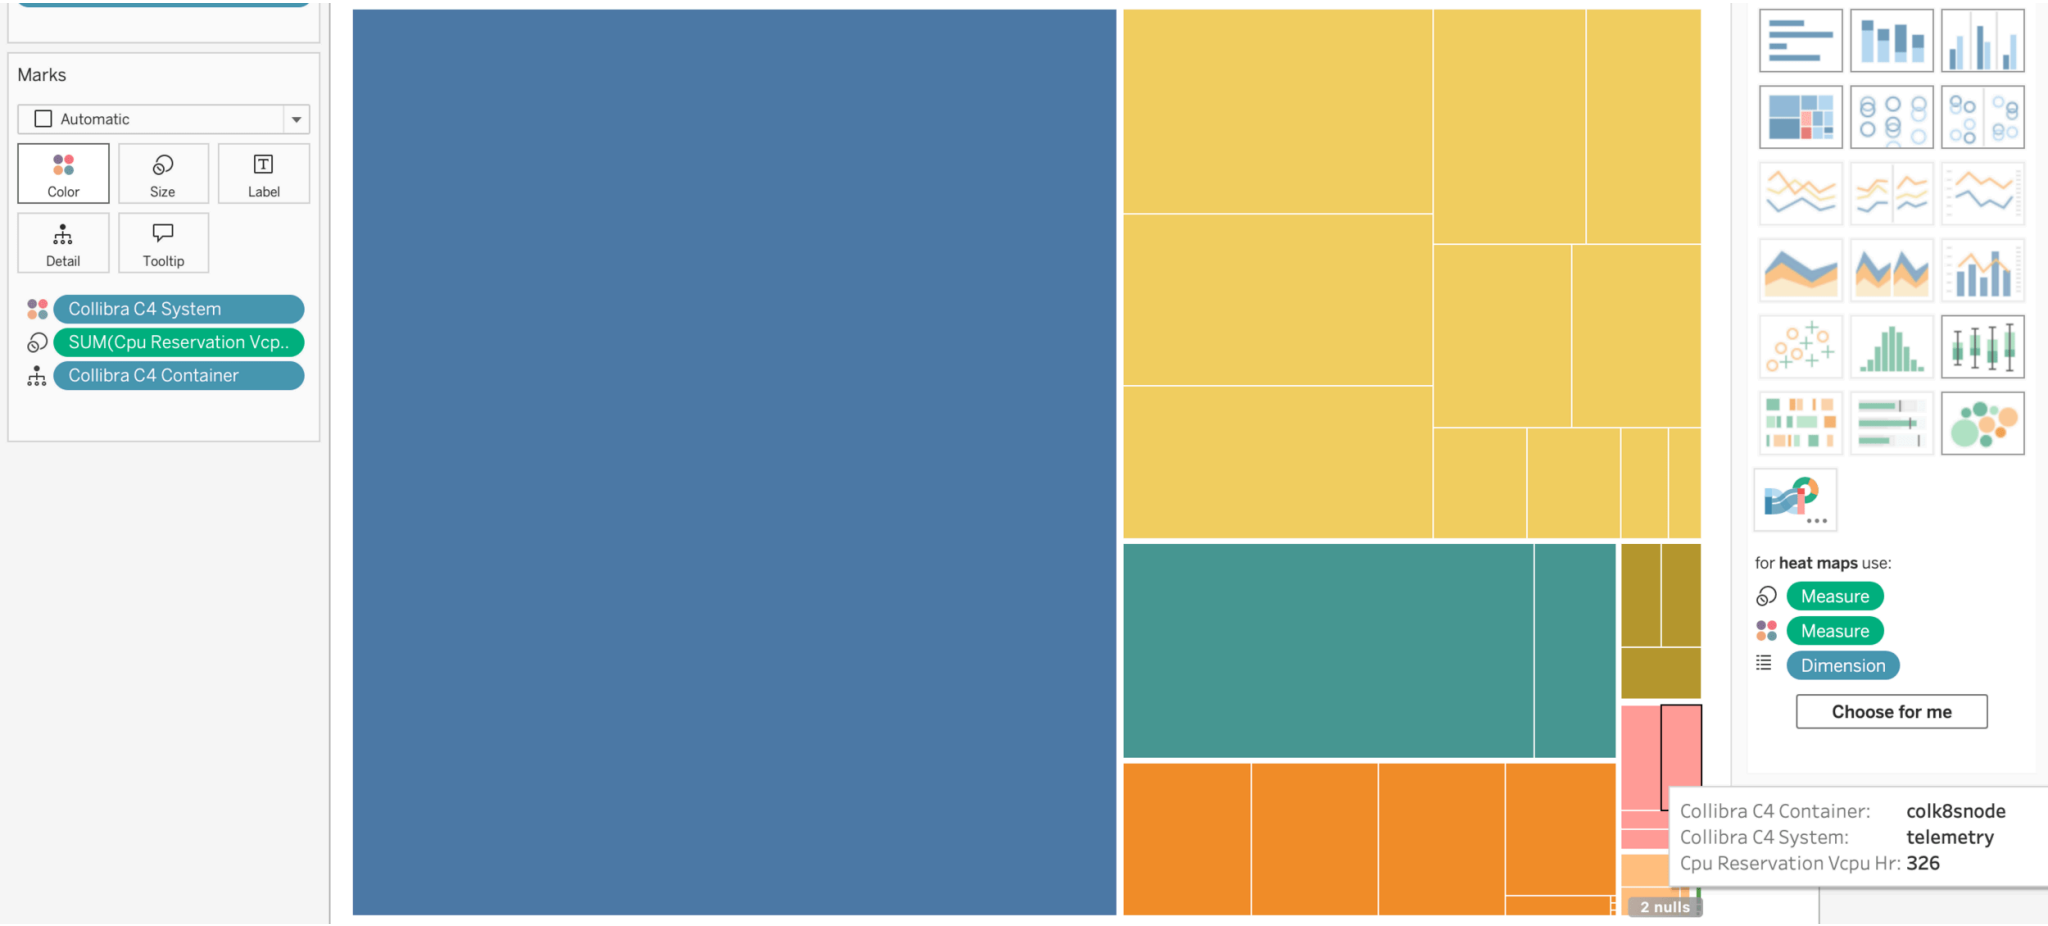

"Compare the Usage and Reservation?"

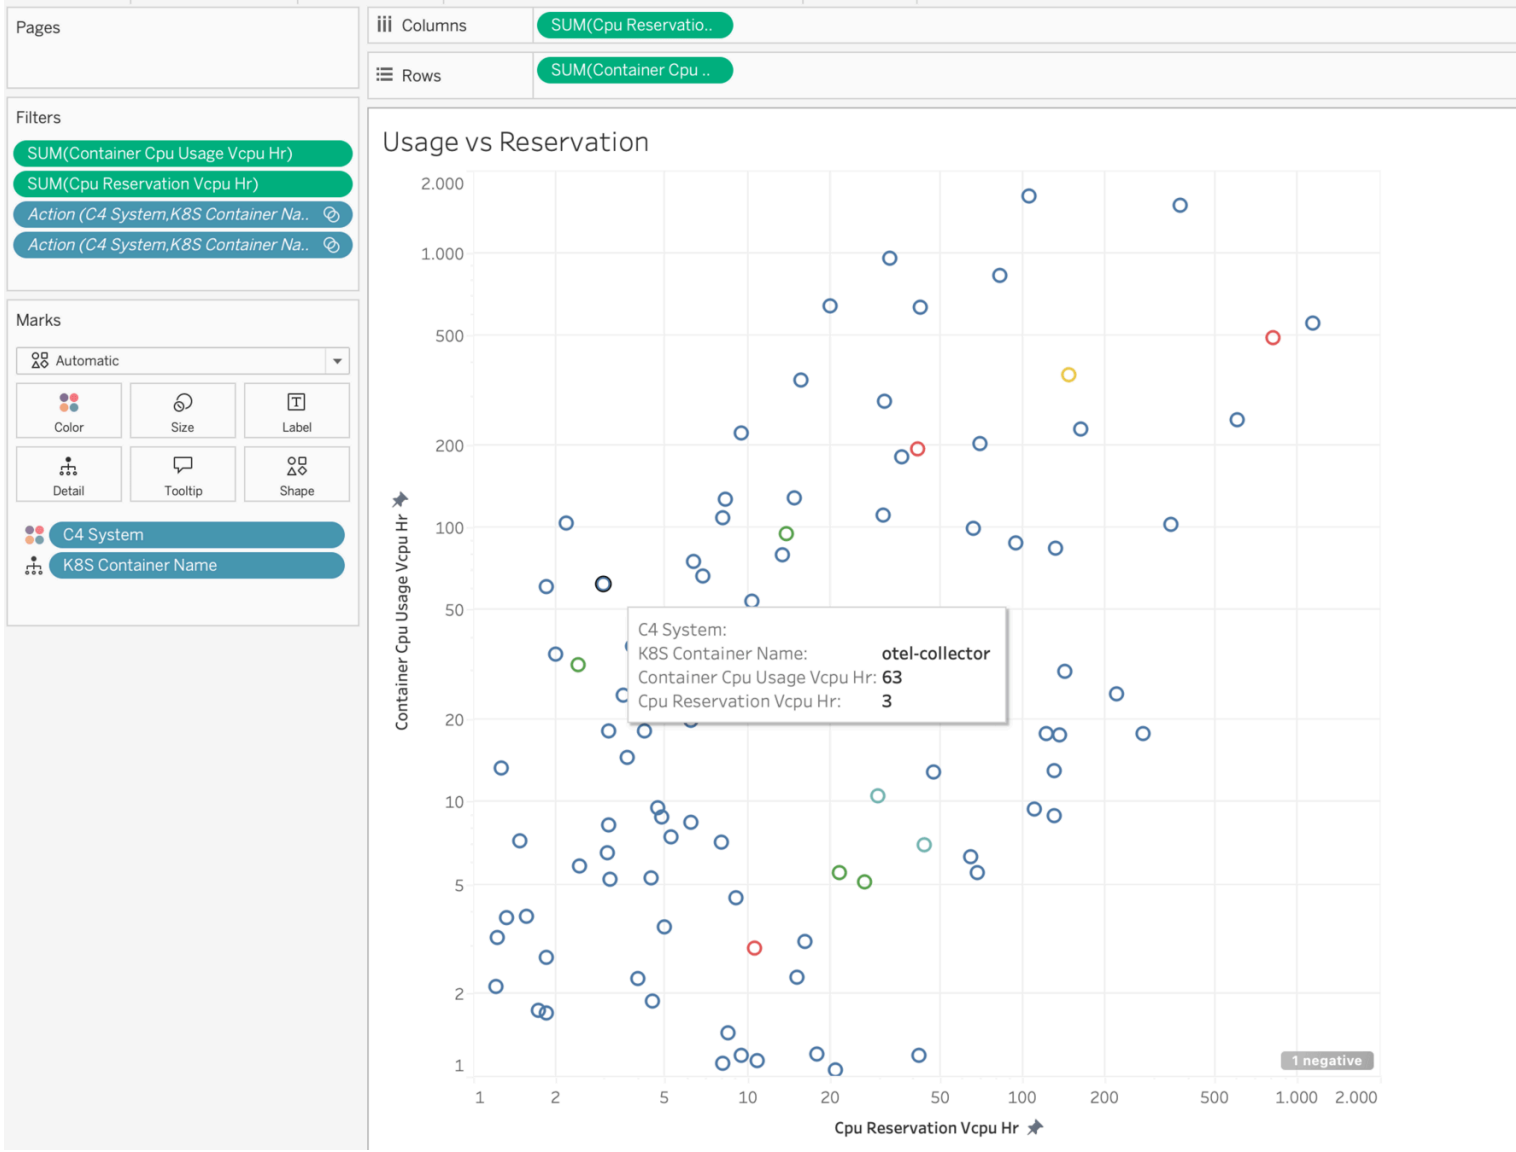


Trend analysis per tenant

Architecture
and Collection at the Edge

Collibra Architecture

- **On-prem heritage:** single deployable monolith → hosted SaaS on VMs (single-tenant isolation for free)
- **Kubernetes shift:** microservices for team velocity & polyglot (Python now dominant in AI)
- **The pivot:** "container per service per tenant" was unsustainable → shared multi-tenant on K8s

Shared multi-tenancy saves cost — but makes it harder to figure out the cost per tenant… how can we solve this?

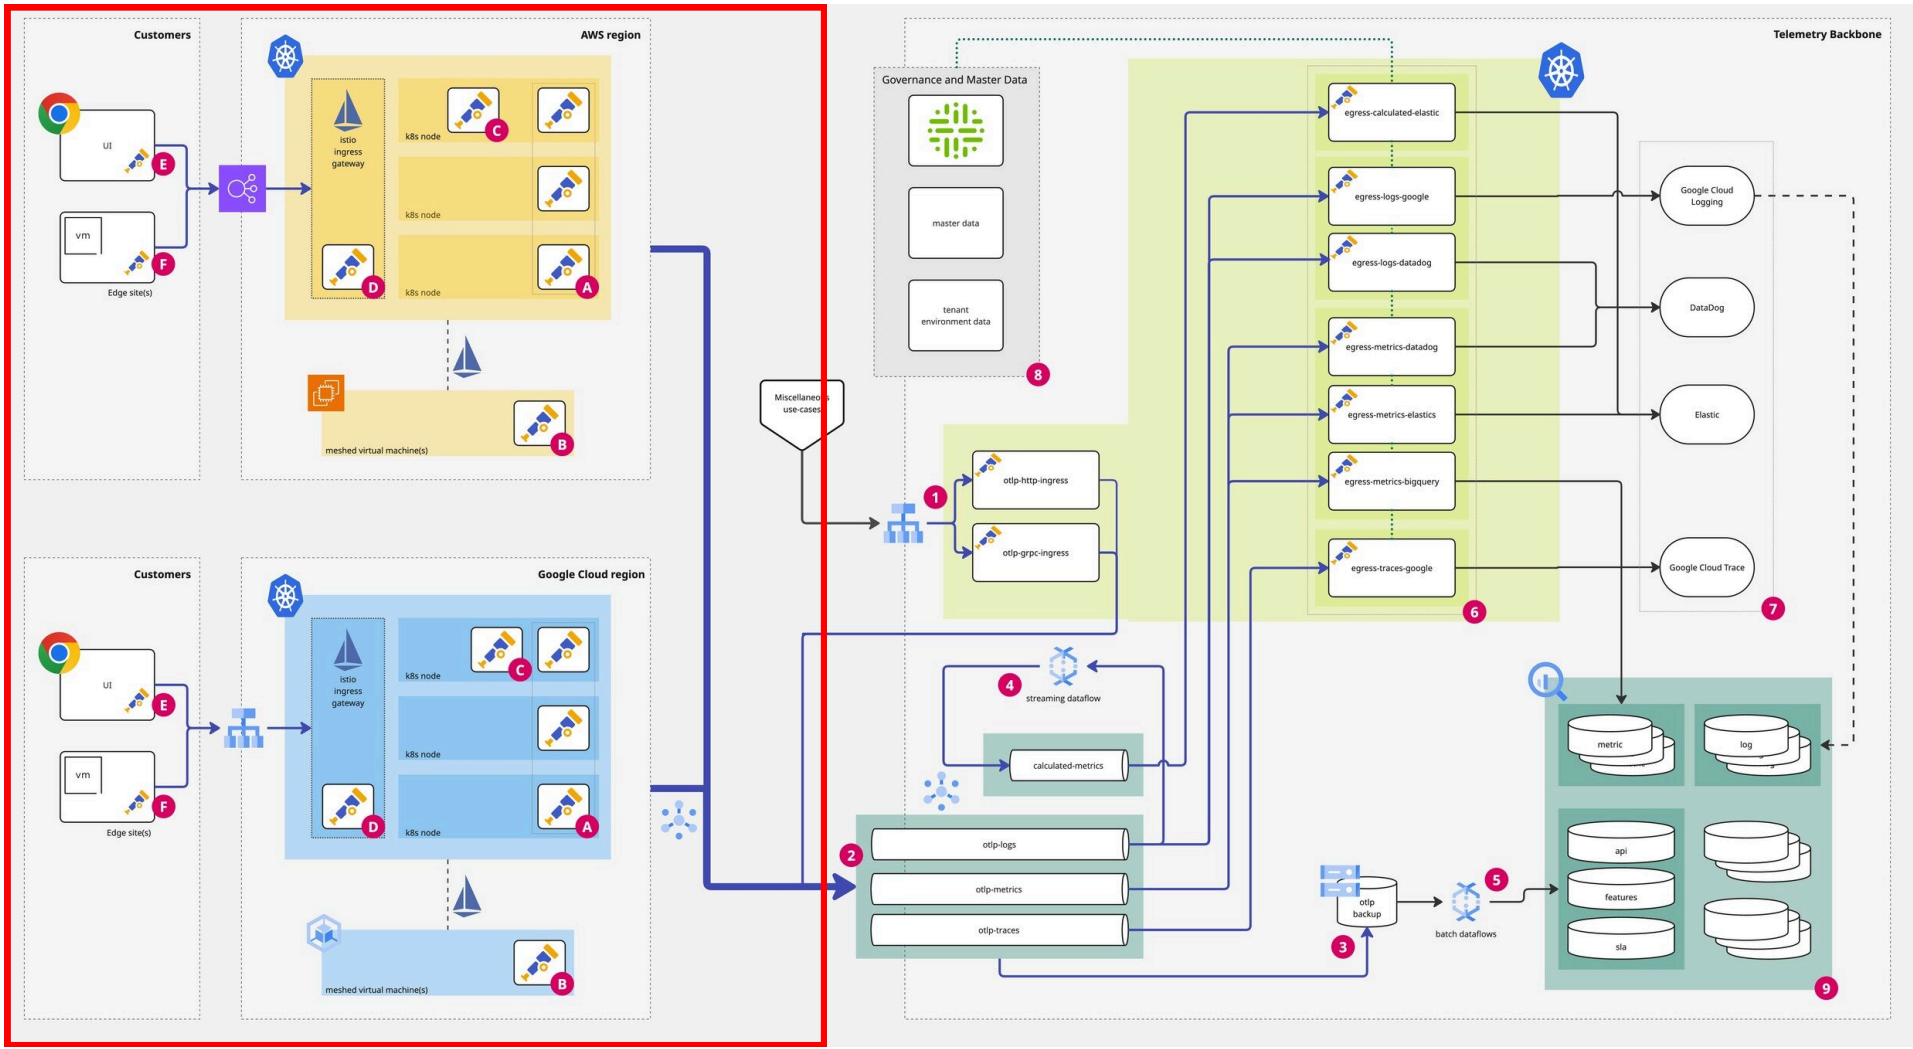
Customers
AWS region
UI
vm
Edge site(s)
istio ingress gateway
k8s node
k8s node
k8s node
meshed virtual machine(s)
Customers
Google Cloud region
UI
vm
Edge site(s)
istio ingress gateway
k8s node
k8s node
k8s node
meshed virtual machine(s)
Governance and Master Data
master data
tenant environment data
8
Miscellaneous use-cases
otlp-http-ingress
otlp-grpc-ingress
1
streaming dataflow
4
calculated-metrics
otlp-logs
otlp-metrics
otlp-traces
2
otlp backup
3
batch dataflows
5
Telemetry Backbone
egress-calculated-elastic
egress-logs-google
egress-logs-datadog
egress-metrics-datadog
egress-metrics-elastics
egress-metrics-bigquery
egress-traces-google
6
Google Cloud Logging
DataDog
Elastic
Google Cloud Trace
7
metric
log
api
features
sla
9

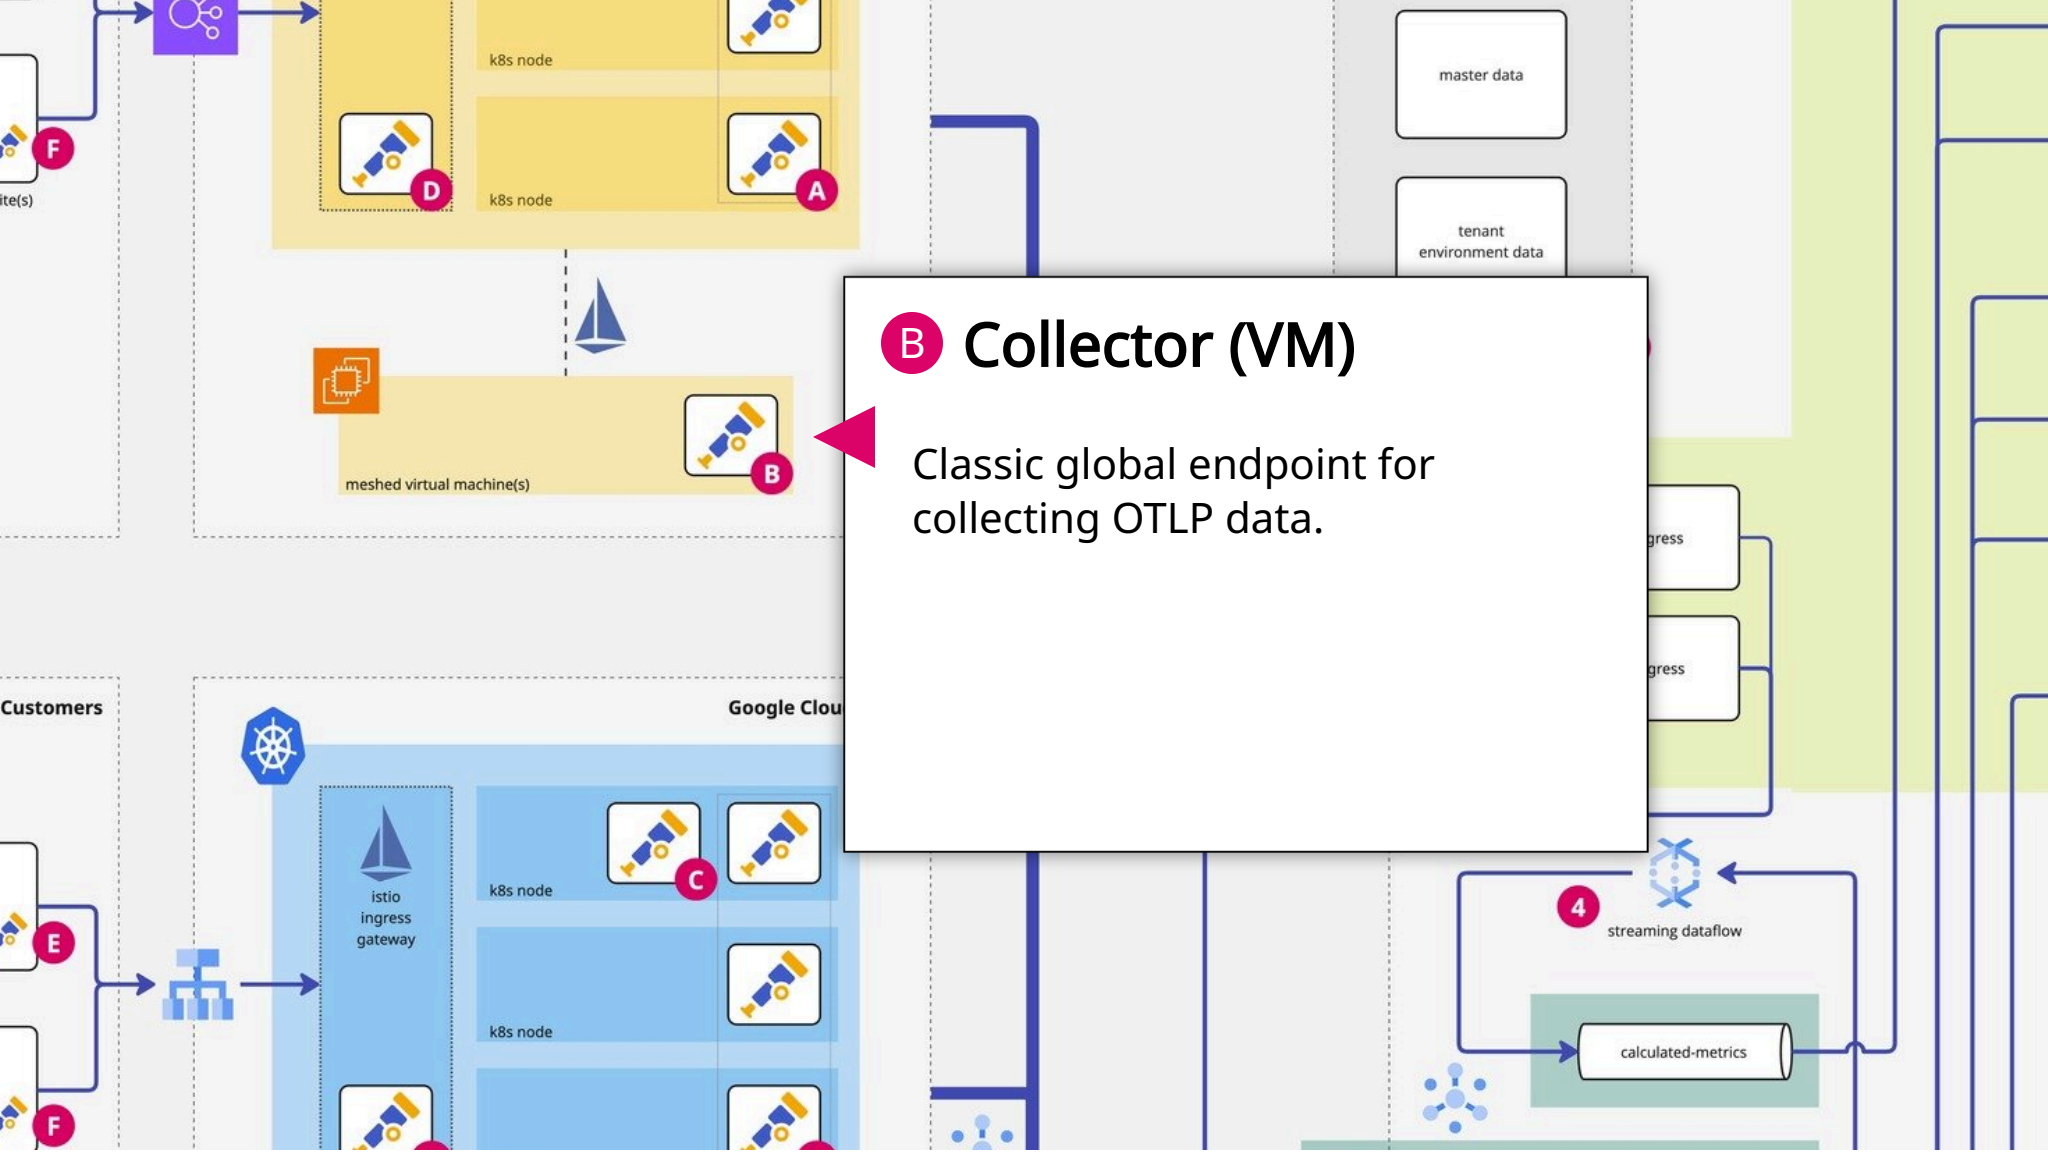
B Collector (VM)
Classic global endpoint for collecting OTLP data.
k8s node
k8s node
meshed virtual machine(s)
master data
tenant environment data
Customers
Google Clou
istio ingress gateway
k8s node
k8s node
4
streaming dataflow
calculated-metrics

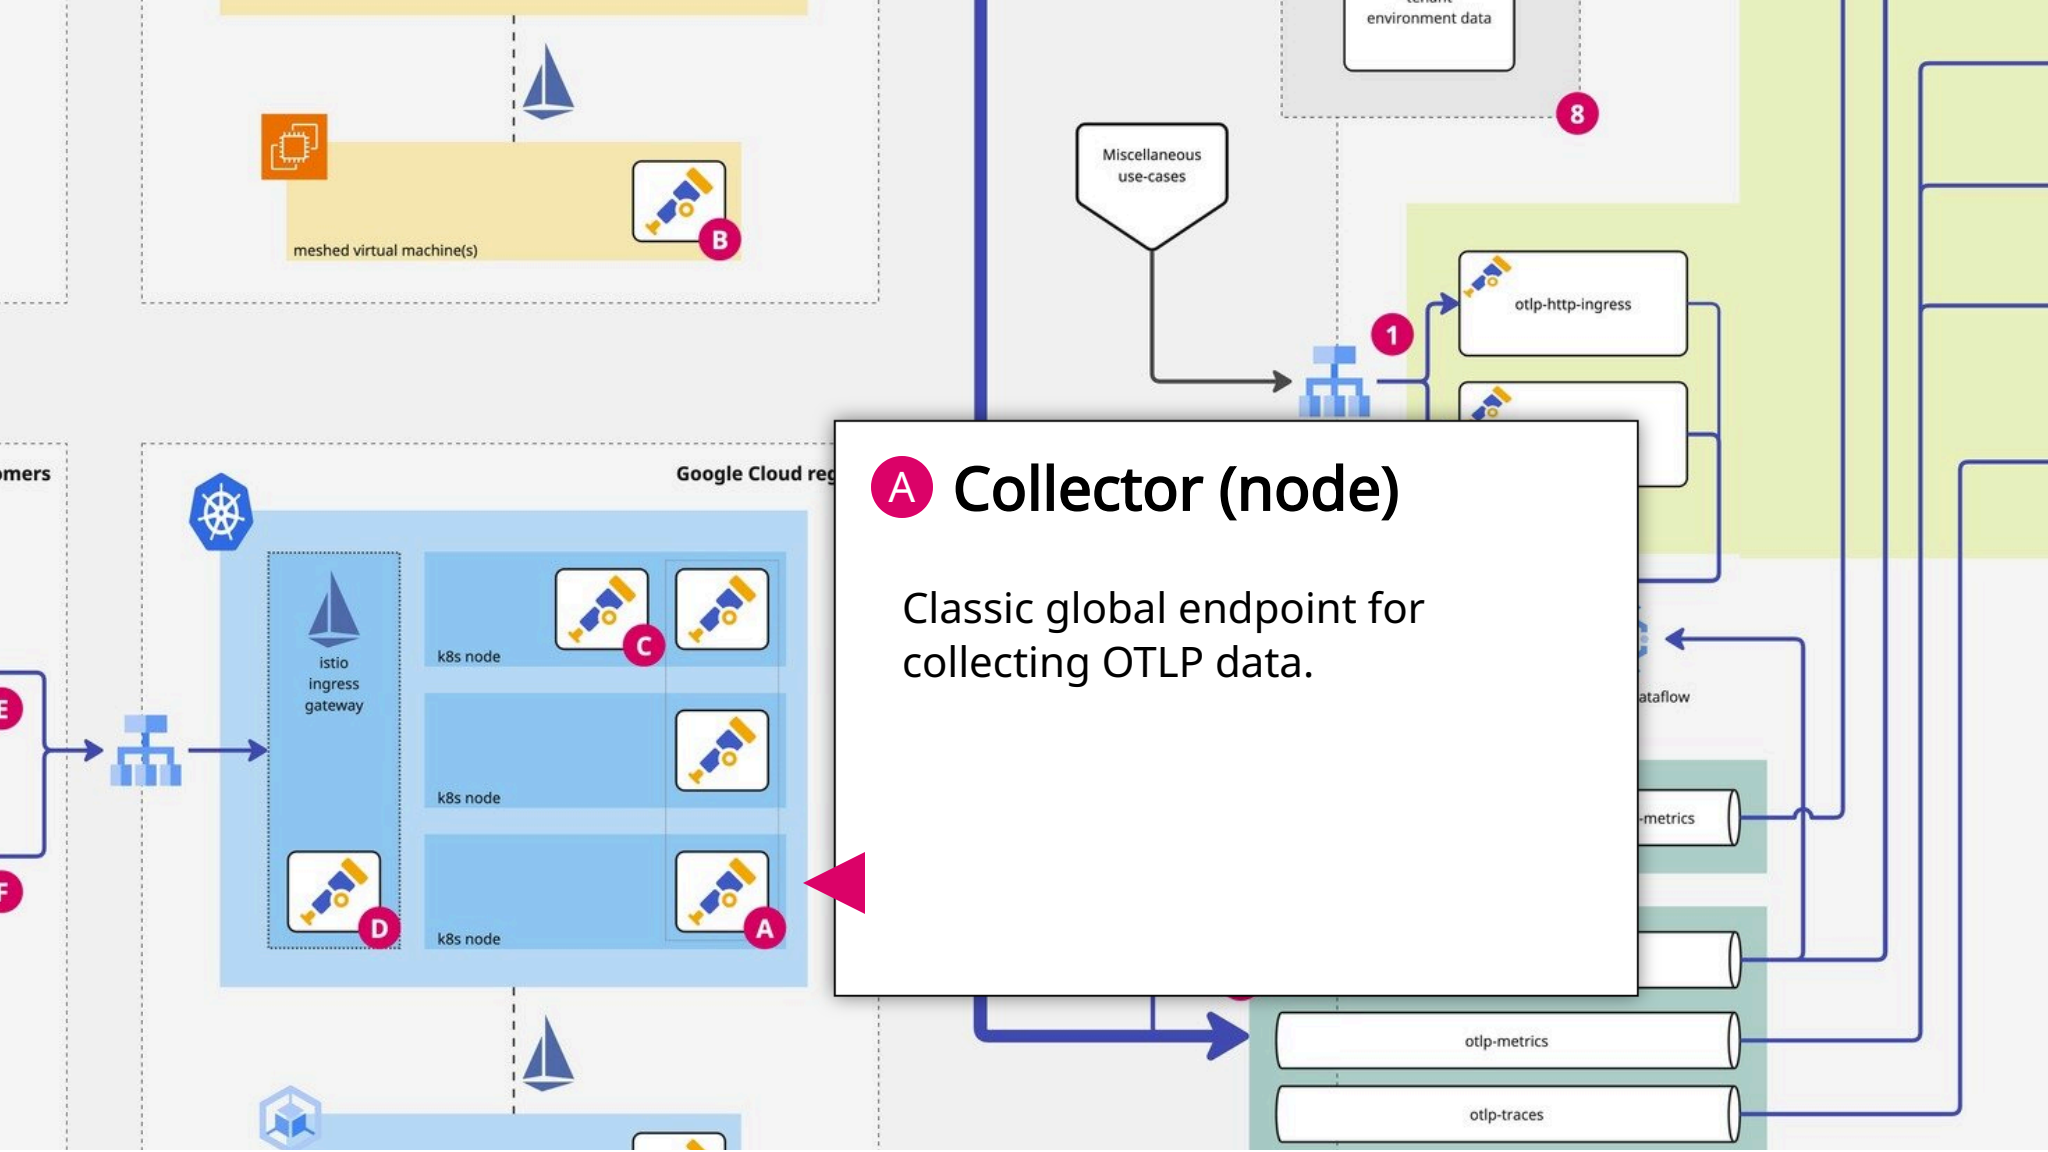
environment data
8
Miscellaneous use-cases
otlp-http-ingress
1
meshed virtual machine(s)
B
Google Cloud reg
istio ingress gateway
k8s node
C
k8s node
k8s node
D
A
ataflow
-metrics
otlp-metrics
otlp-traces
A Collector (node)
Classic global endpoint for collecting OTLP data.

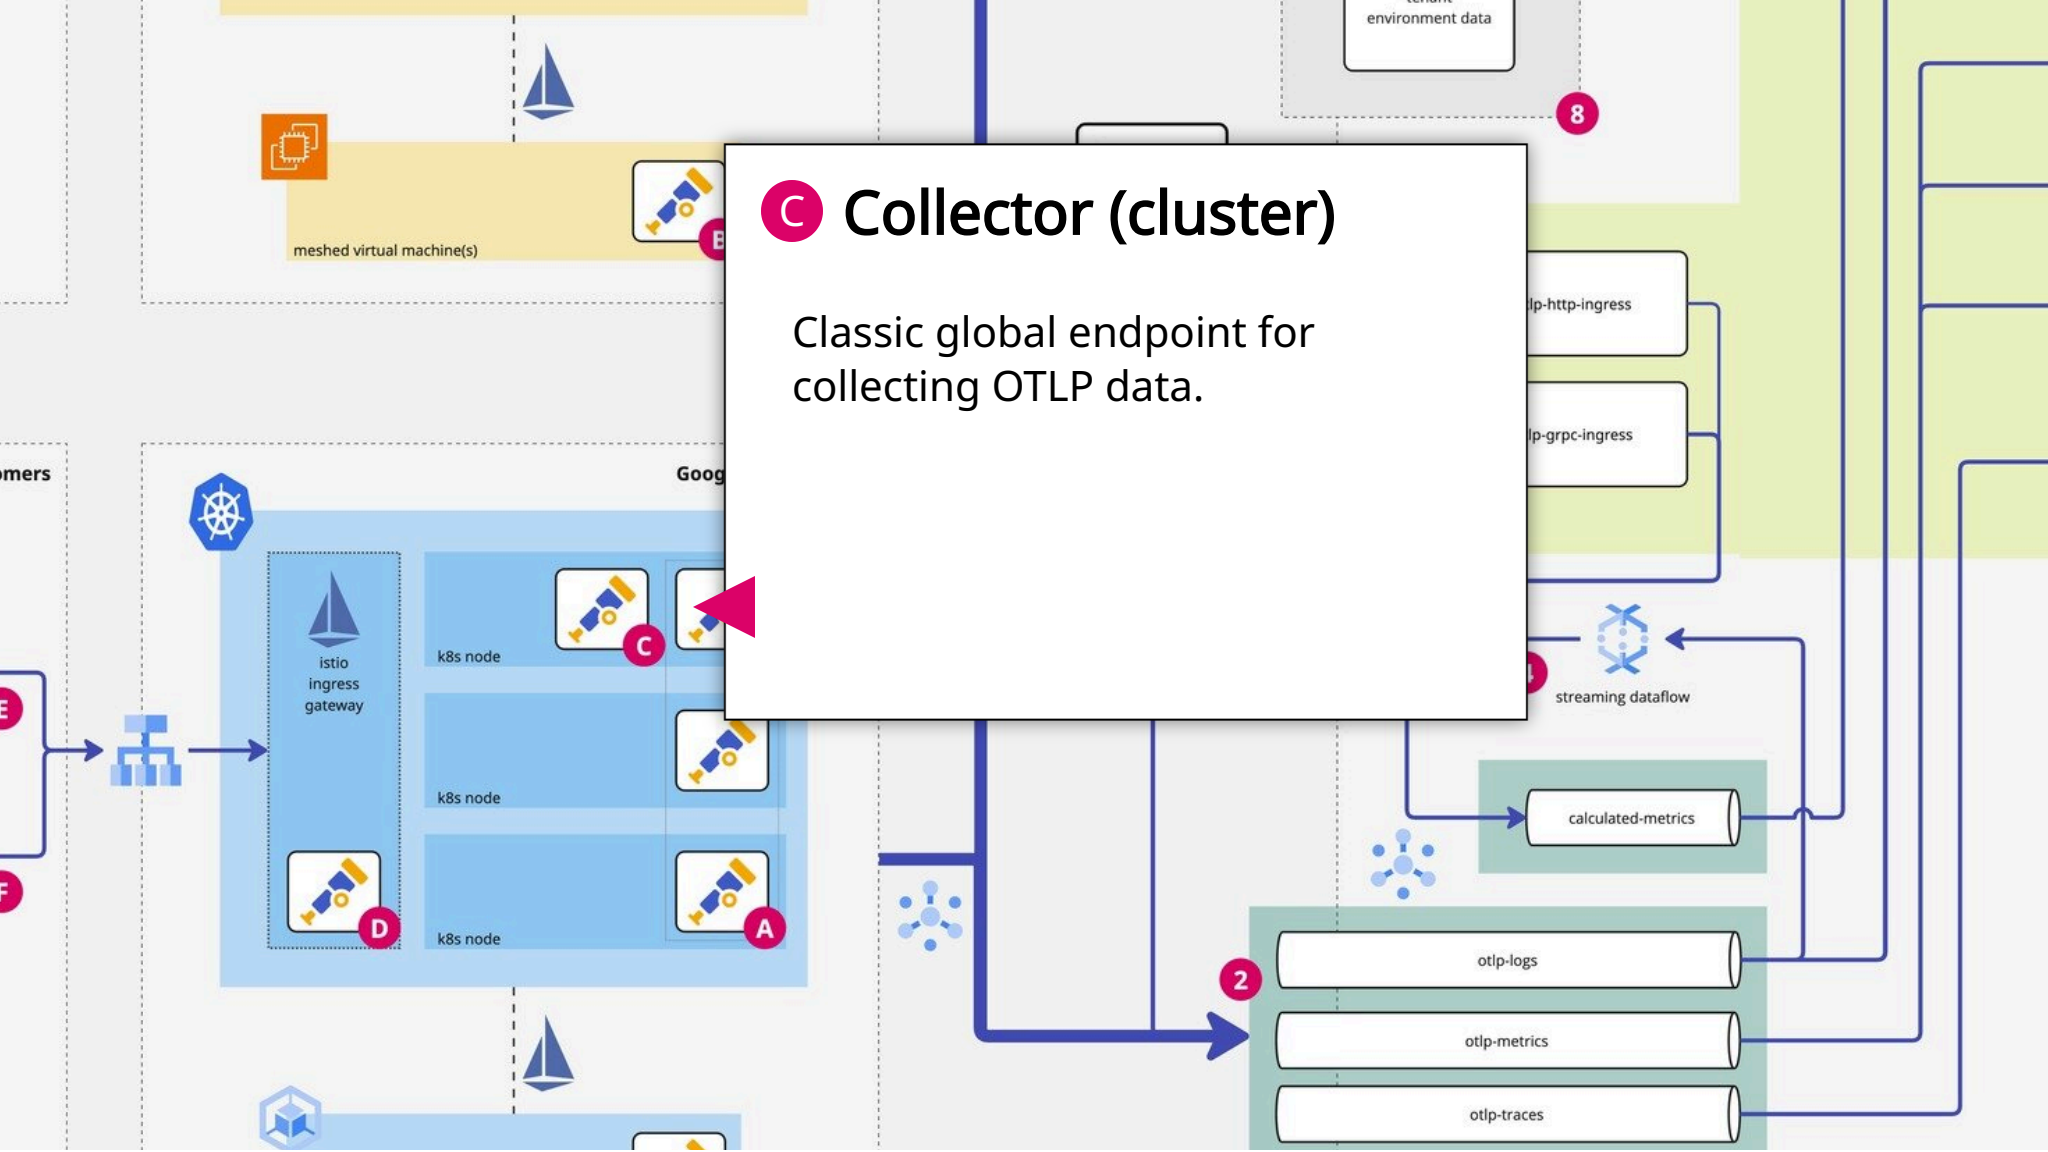
environment data
8
meshed virtual machine(s)
B
lp-http-ingress
lp-grpc-ingress
C Collector (cluster)
Classic global endpoint for collecting OTLP data.
omers
Goog
istio ingress gateway
k8s node
C
k8s node
k8s node
D
A
streaming dataflow
calculated-metrics
2
otlp-logs
otlp-metrics
otlp-traces

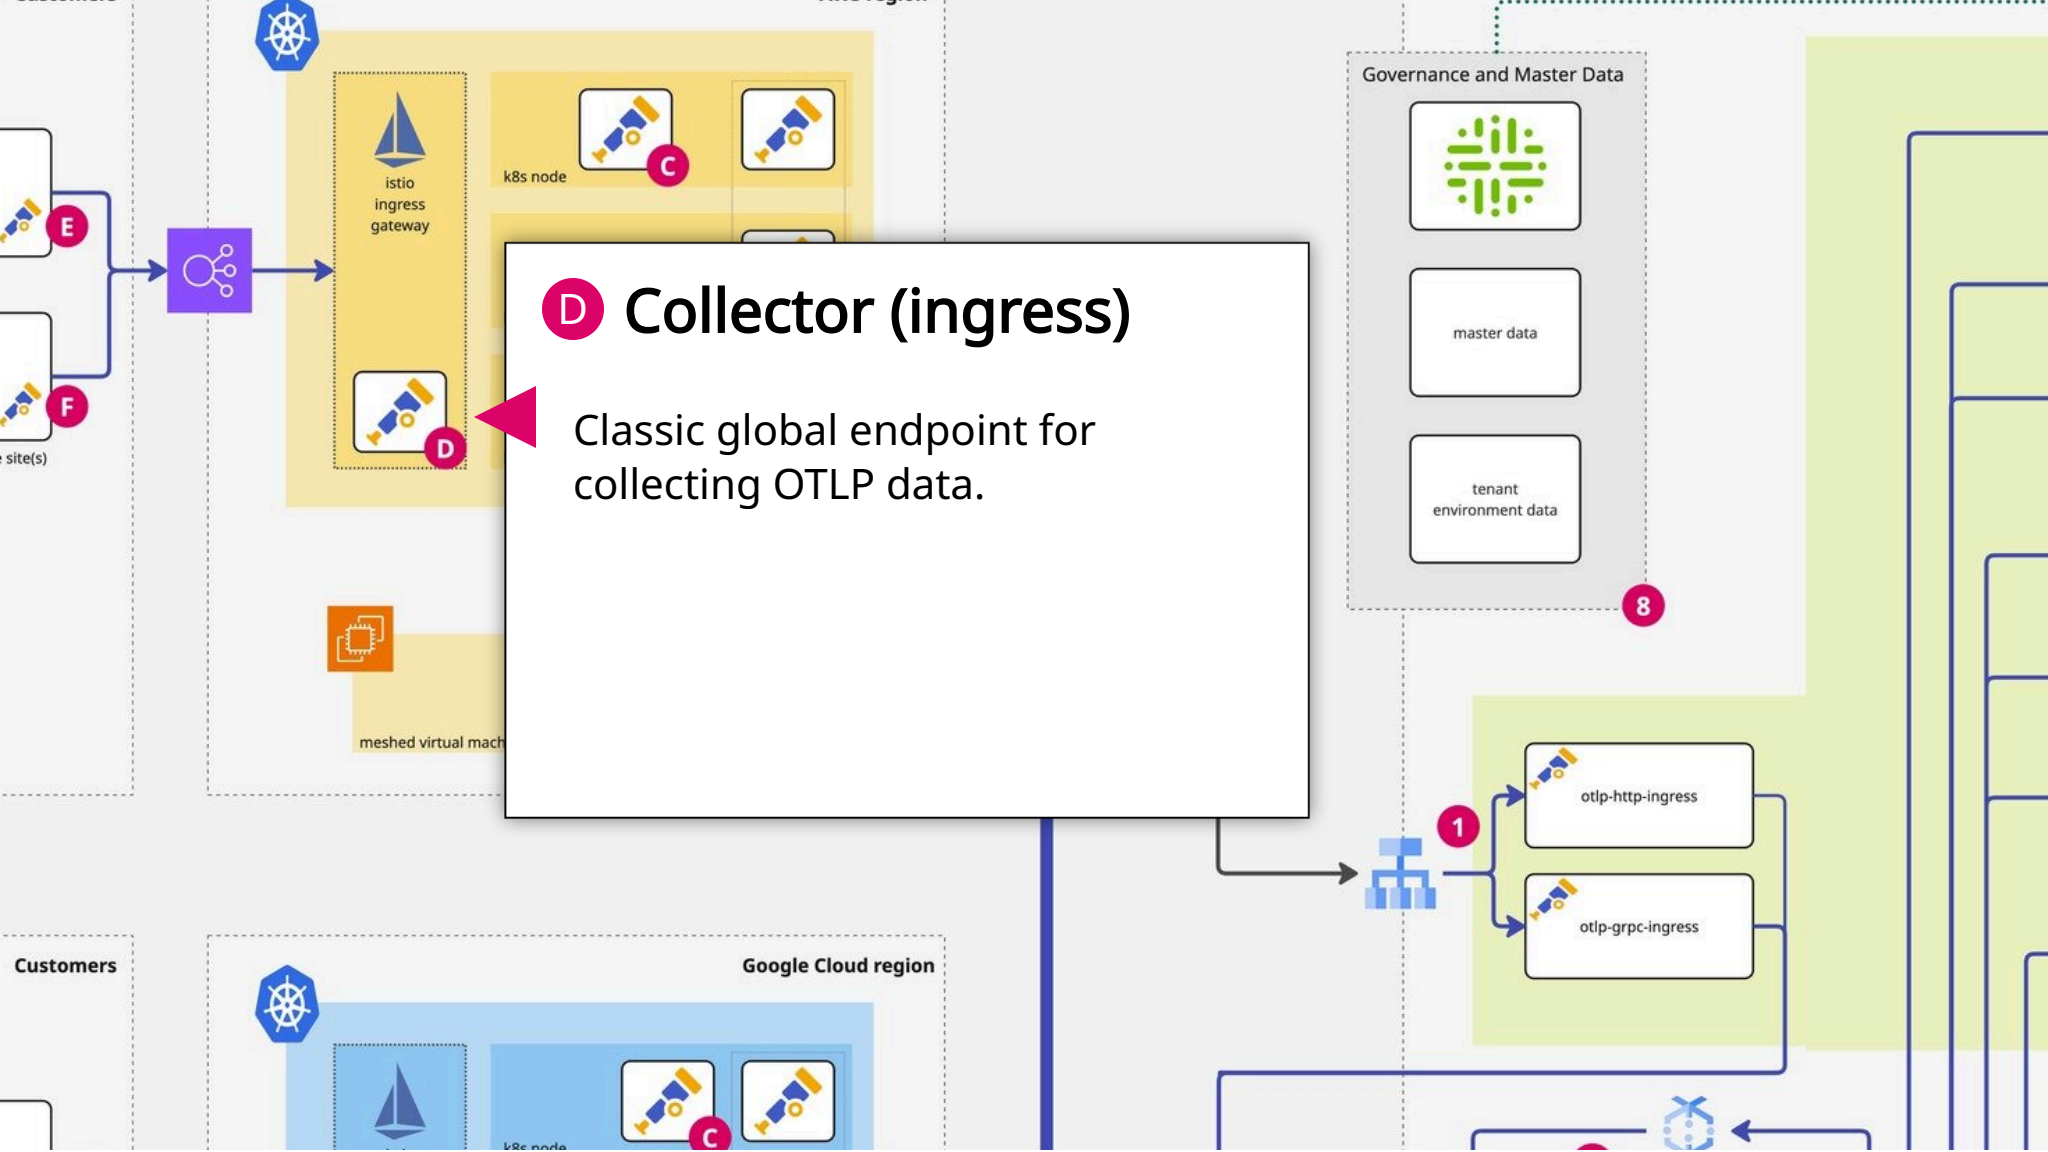
D Collector (ingress)
Classic global endpoint for collecting OTLP data.
Governance and Master Data
master data
tenant environment data
8
istio ingress gateway
k8s node
C
D
E
F
site(s)
meshed virtual mach
otlp-http-ingress
otlp-grpc-ingress
1
Customers
Google Cloud region

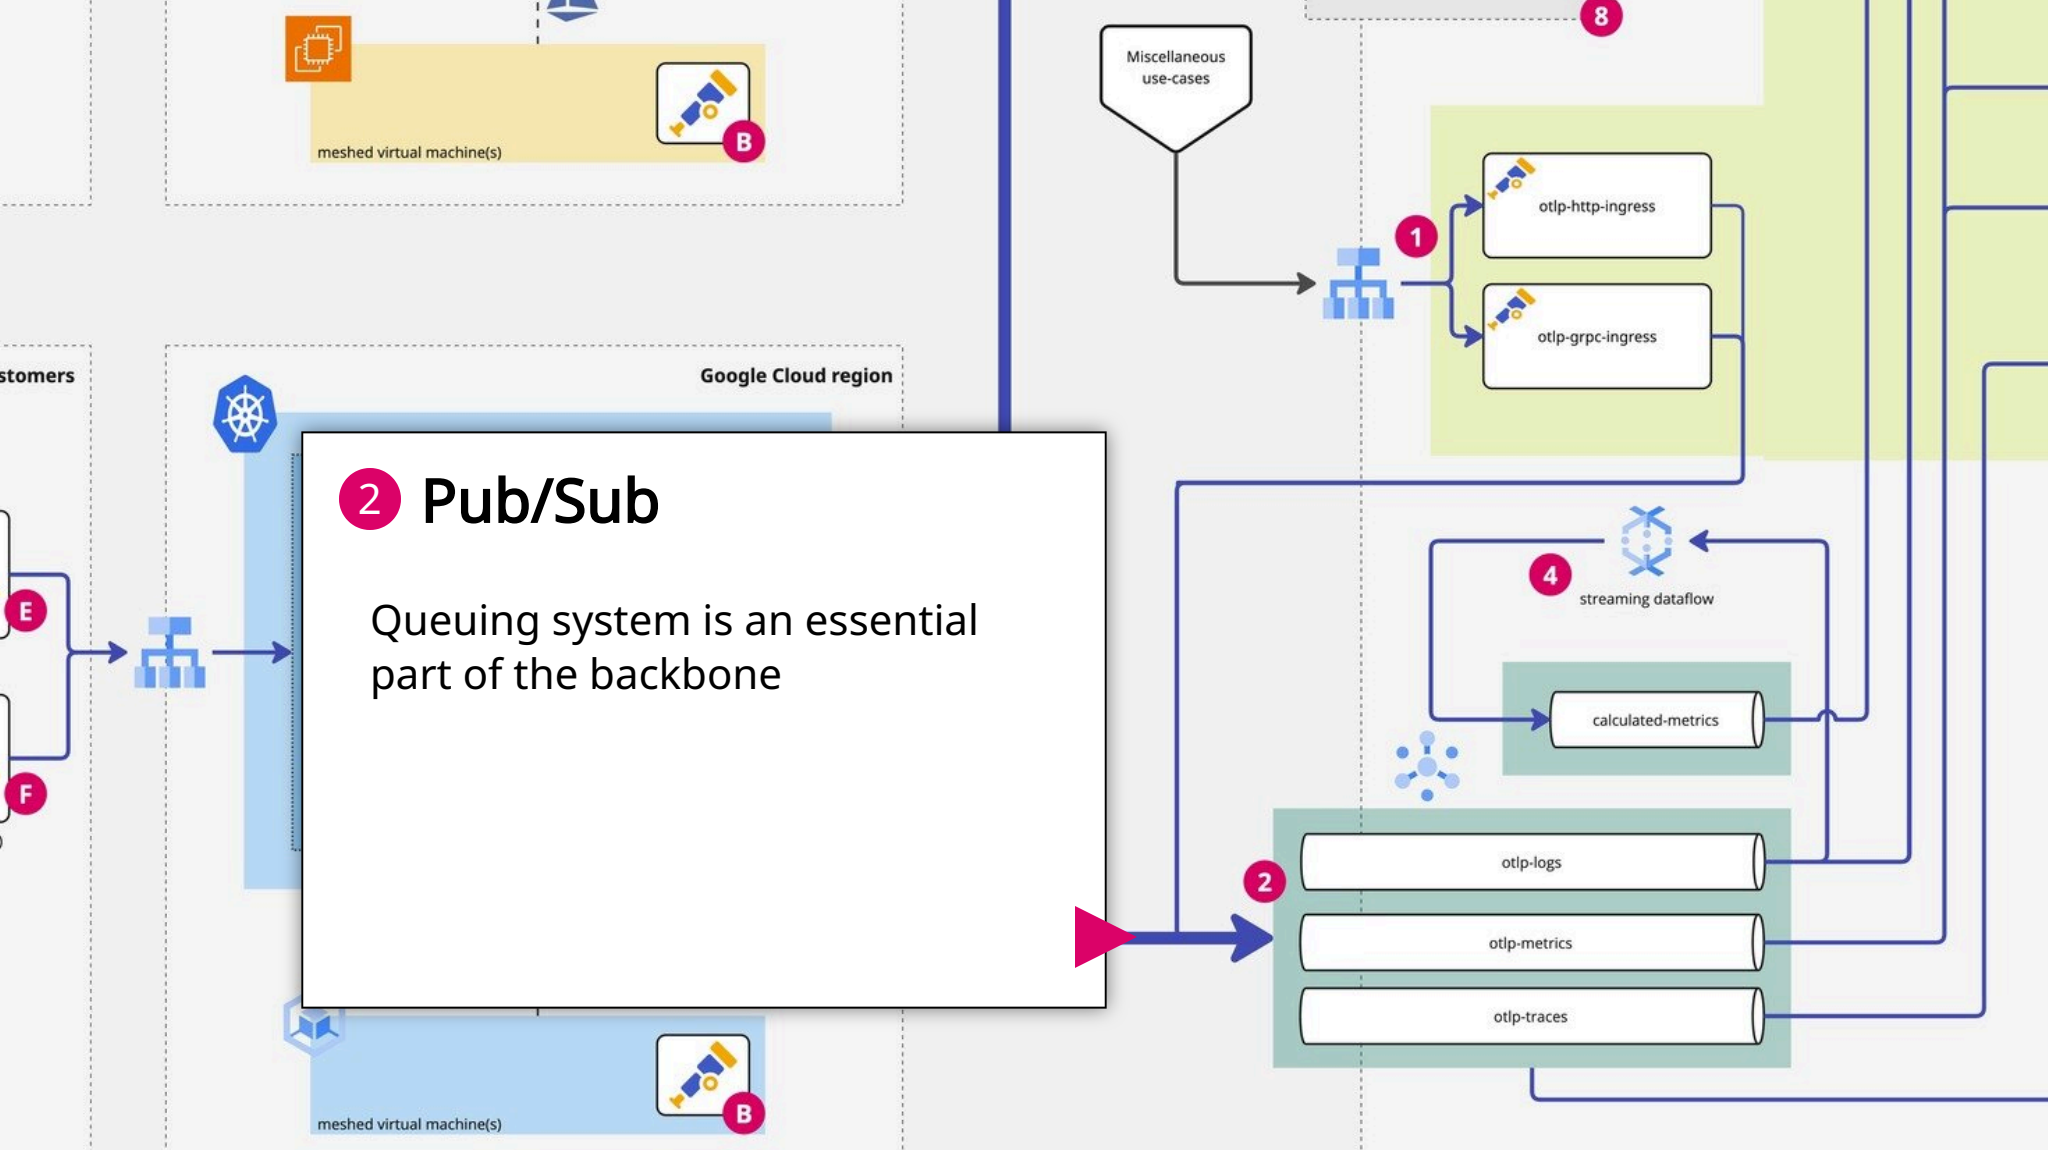
Miscellaneous use-cases
otlp-http-ingress
otlp-grpc-ingress
streaming dataflow
calculated-metrics
otlp-logs
otlp-metrics
otlp-traces
meshed virtual machine(s)
meshed virtual machine(s)
Google Cloud region
2 Pub/Sub
Queuing system is an essential part of the backbone

Golden Signals

OpenTelemetry Attributes

Gold Attr. #1: Tenant

`collibra.tenant.environment_id`

- VMs Resource Attributes - Configurated at the collector
- Pod Resource Attributes - For single tenant pods
- Multi-tenant Pod - Signal Attributes

Gold Attr. #1: Tenant

`collibra.tenant.environment_id`

- VMs Resource Attributes - Configurated at the collector
- Pod Resource Attributes - For single tenant pods
- Multi-tenant Pod - Signal Attributes

```
{
  "event_name": "workflow:started",
  "tenant_environment_id": "...",
  "asset_id": "..."
}
```

CSTE - Collibra Structured Telemetry **Event**: Events are our golden signal

Gold Attr. #1: Tenant

`collibra.tenant.environment_id`

- VMs Resource Attributes - Configurated at the collector
- Pod Resource Attributes - For single tenant pods
- Multi-tenant Pod - Signal Attributes

Multi-tenant service? Dev's responsibility to add signals in code, eg. Mapped Diagnostic Context

```
{
  "event_name": "workflow:started",
  "tenant_environment_id": "...",
  "asset_id": "..."
}
```

```
MDC.put("tenant_environment_id",
        ctx.getTenantEnvironmentId());
try {
  // all logs in this thread
} finally {
  MDC.clear();
}
```

Gold Attr. #2: Architecture

https://c4model.com/ - The C4 model is an easy to learn, developer friendly approach to software architecture diagramming (by Simon Brown)

- **System** - logical product capability
- **Container** - service, logical database, topic, module
- **Deployment Node** - where it runs (can nest)

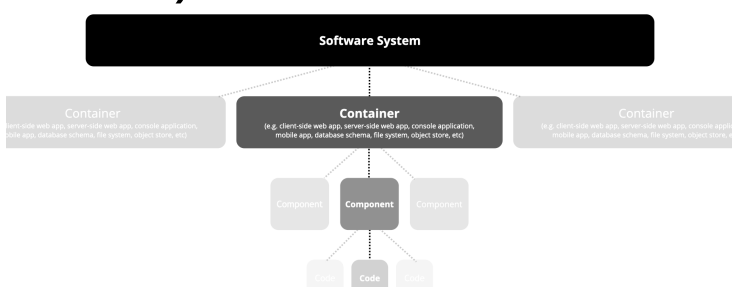

A **software system** is made up of one or more **containers** (applications and data stores), each of which contains one or more **components**, which in turn are implemented by one or more **code** elements (classes, interfaces, objects, functions, etc).

Gold Attr. #2: Architecture

https://c4model.com/ - The C4 model is an easy to learn, developer friendly approach to software architecture diagramming (by Simon Brown)

- **System** - logical product capability
- **Container** - service, logical database, topic, module
- **Deployment Node** - where it runs (can nest)

```
collibra.c4.system
collibra.c4.container
collibra.c4.deployment
```

Gold Attr. #2: Architecture

- **Pod** Resource Attributes - Easy with 1:1 mapping

```
labels:
  c4.collibra.com/system: telemetry
  c4.collibra.com/container: colkyverno
```

```
collibra.c4.system: telemetry
collibra.c4.container: colkyverno
```

- **Modular Monoliths** Signal Attributes - It's not only out single tenant core, but also k8s jobs

Modular Monoliths (it becomes the resposability for devs

Telemetry Backbone

Enrichment and Routing

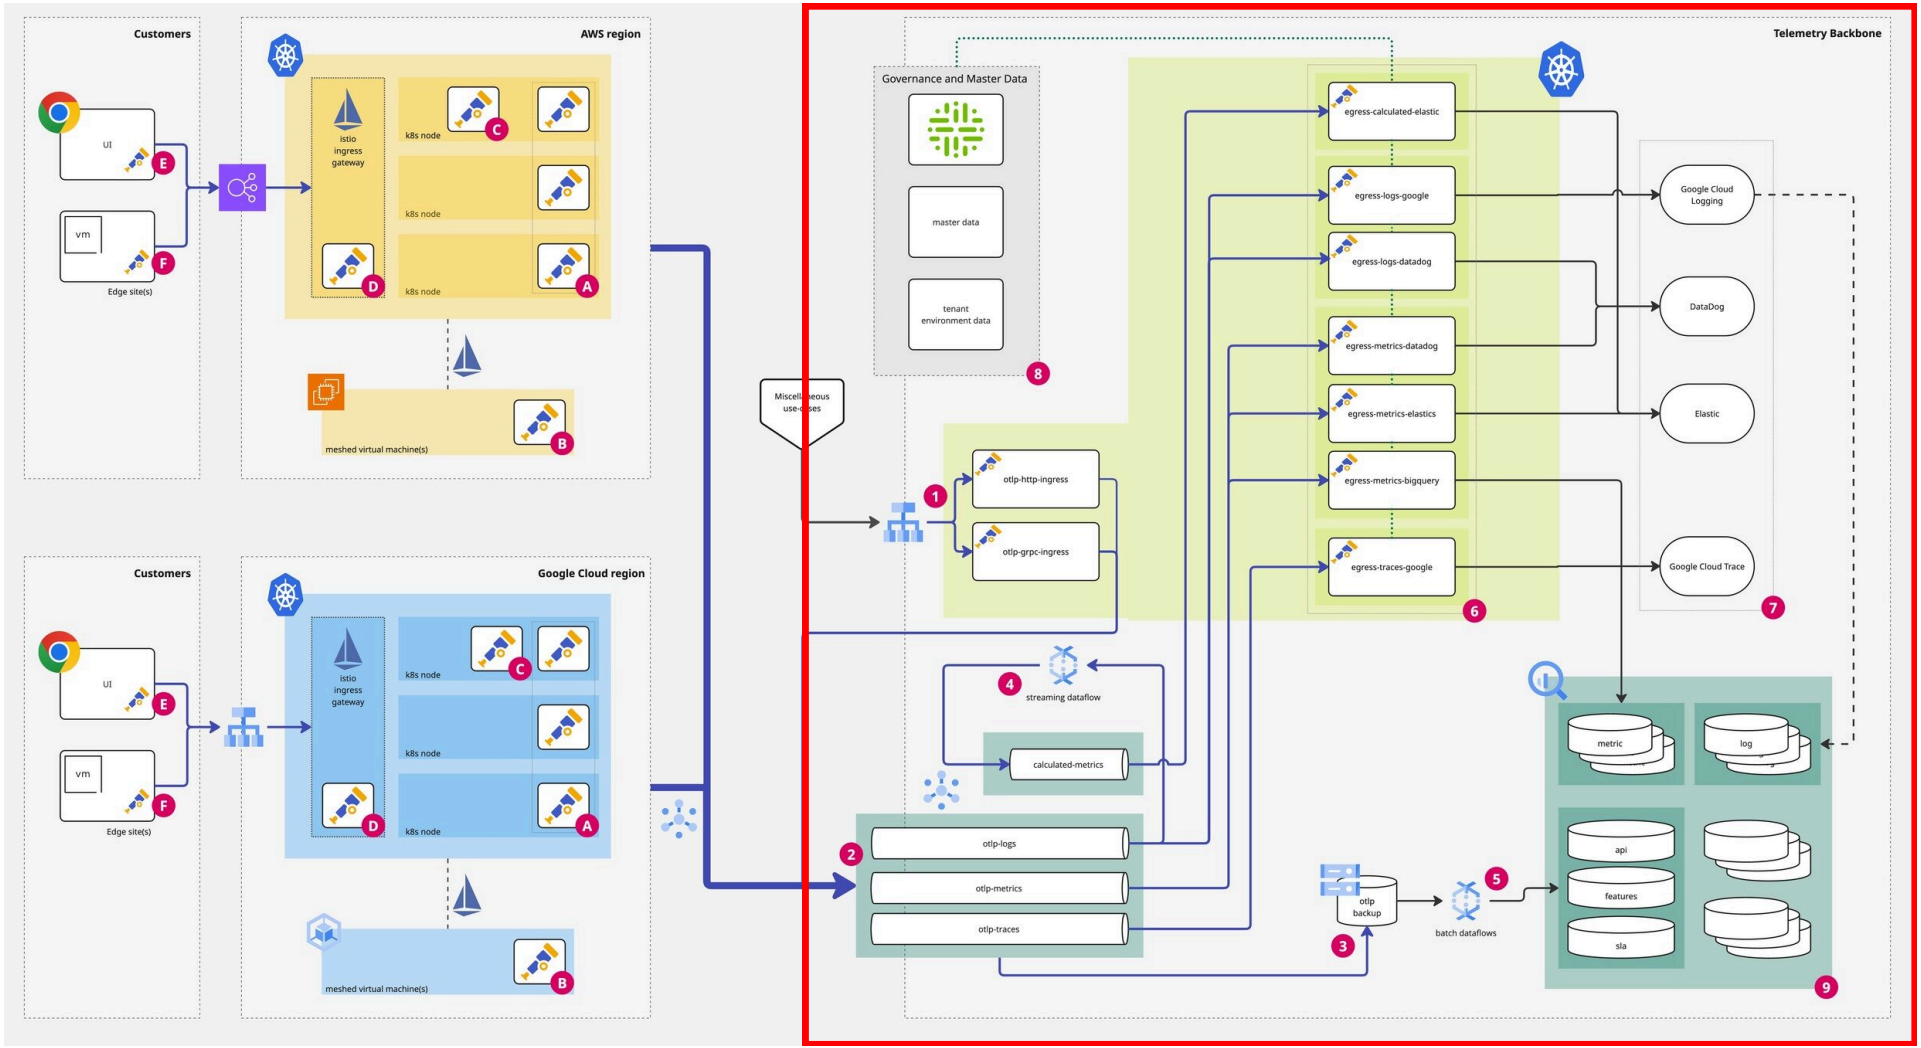
Customers
AWS region
UI
vm
Edge site(s)
istio ingress gateway
k8s node
k8s node
k8s node
meshed virtual machine(s)
Customers
Google Cloud region
UI
vm
Edge site(s)
istio ingress gateway
k8s node
k8s node
k8s node
meshed virtual machine(s)
Telemetry Backbone
Governance and Master Data
master data
tenant environment data
Miscellaneous use-cases
otlp-http-ingress
otlp-grpc-ingress
streaming dataflow
calculated-metrics
otlp-logs
otlp-metrics
otlp-traces
otlp backup
batch dataflows
egress-calculated-elastic
egress-logs-google
egress-logs-datadog
egress-metrics-datadog
egress-metrics-elastics
egress-metrics-bigquery
egress-traces-google
Google Cloud Logging
DataDog
Elastic
Google Cloud Trace
metric
log
api
features
sla

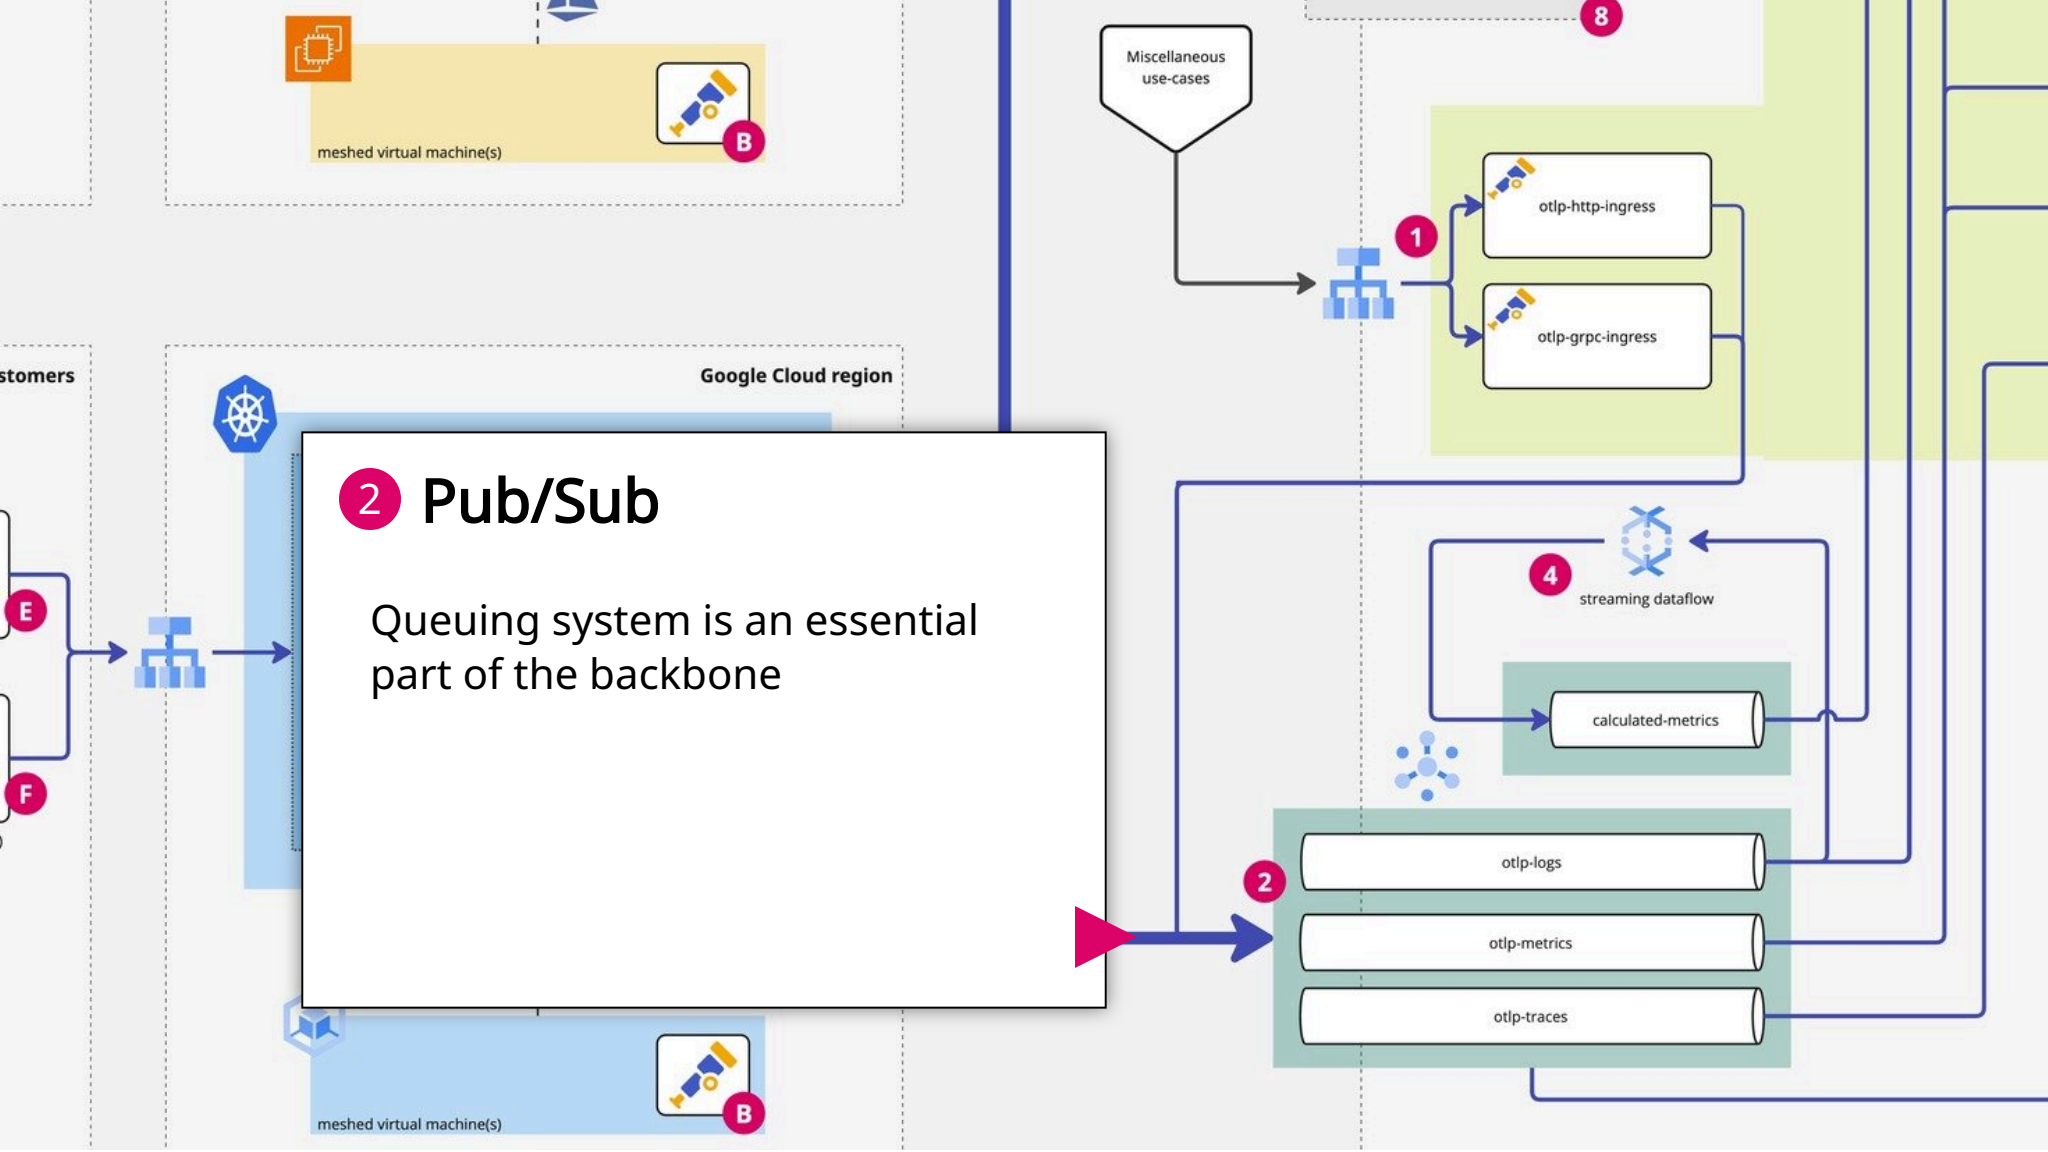
Miscellaneous use-cases
otlp-http-ingress
otlp-grpc-ingress
streaming dataflow
calculated-metrics
otlp-logs
otlp-metrics
otlp-traces
meshed virtual machine(s)
meshed virtual machine(s)
Google Cloud region
2 Pub/Sub
Queuing system is an essential part of the backbone

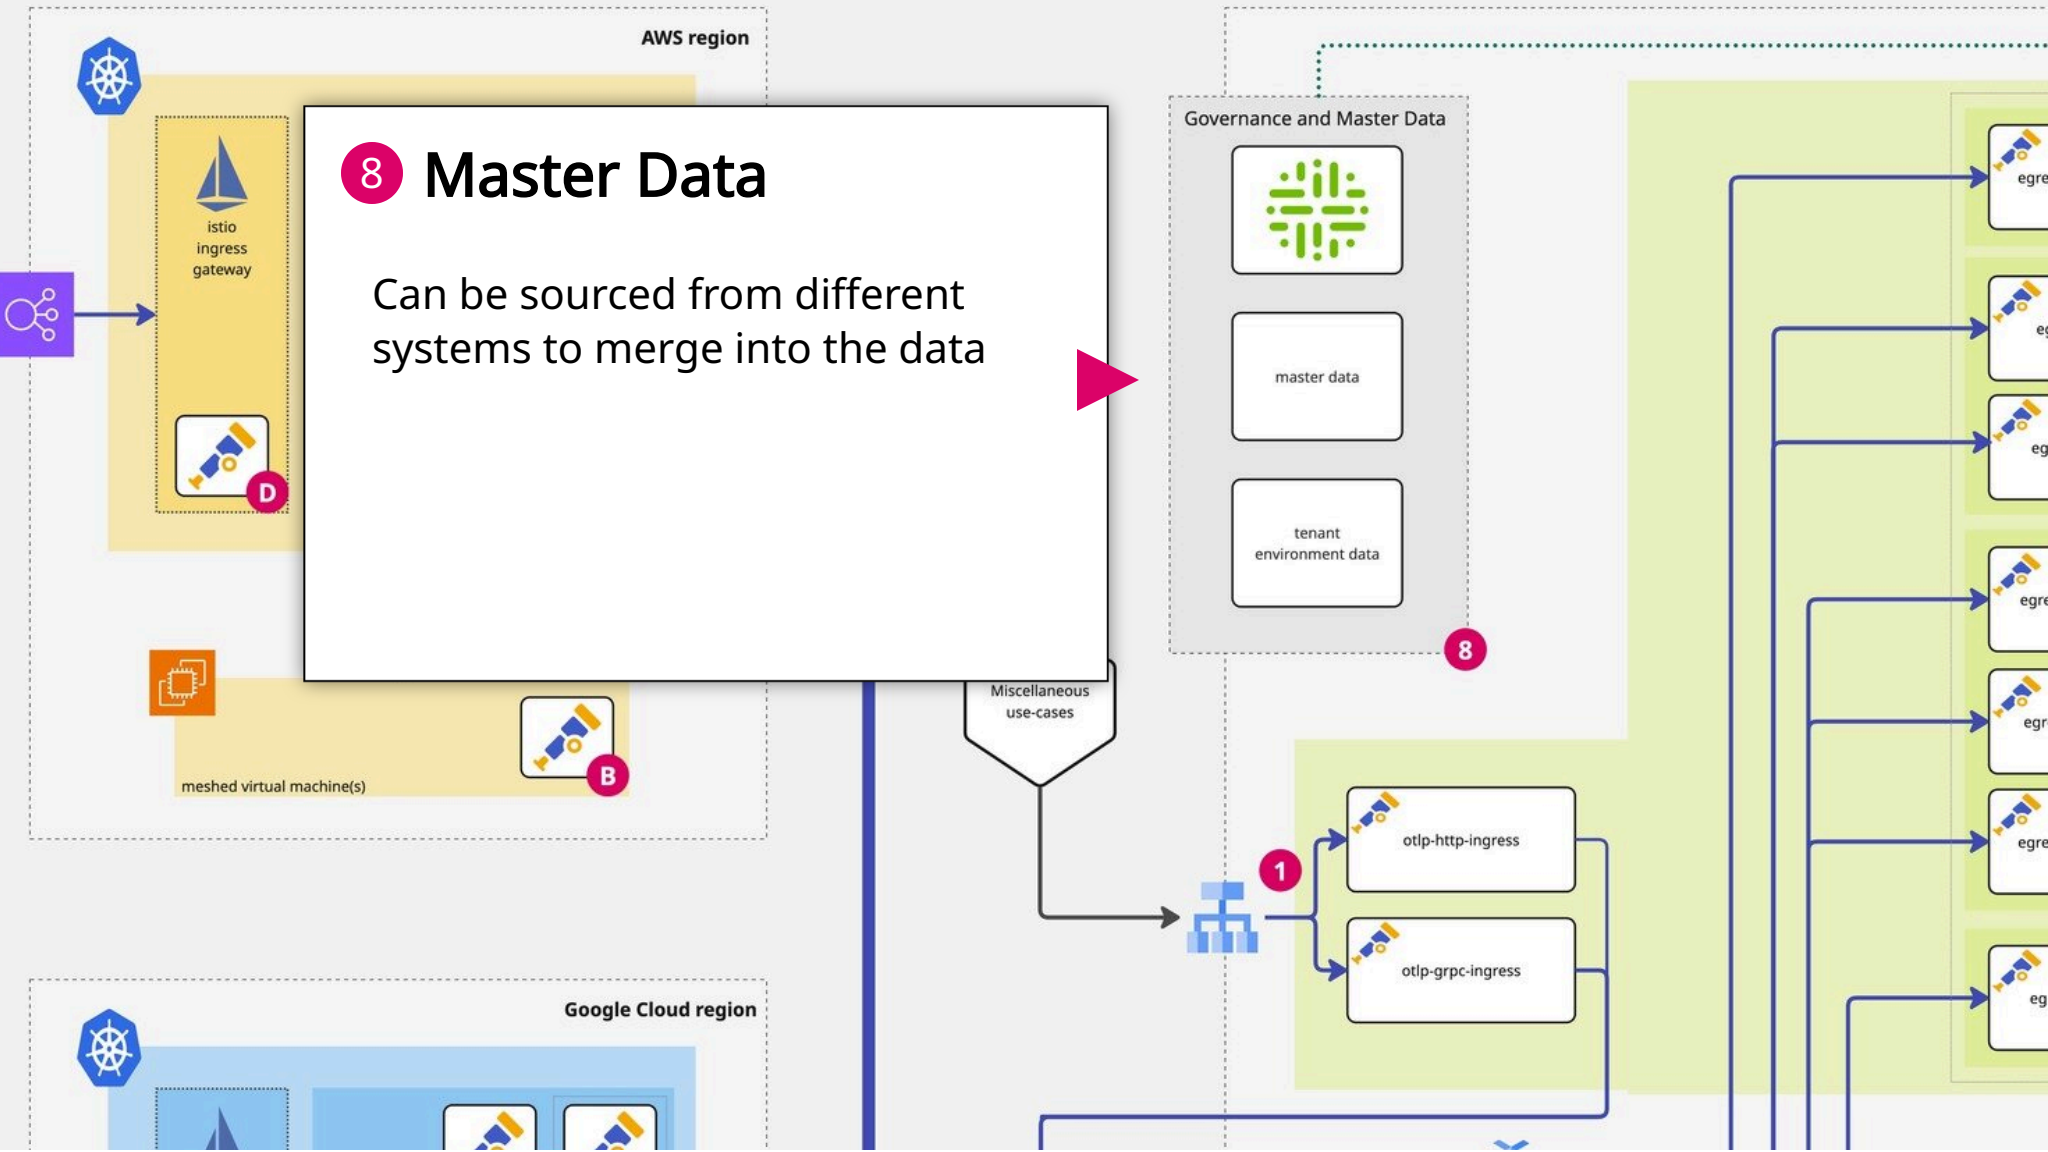

8 Master Data

Can be sourced from different systems to merge into the data

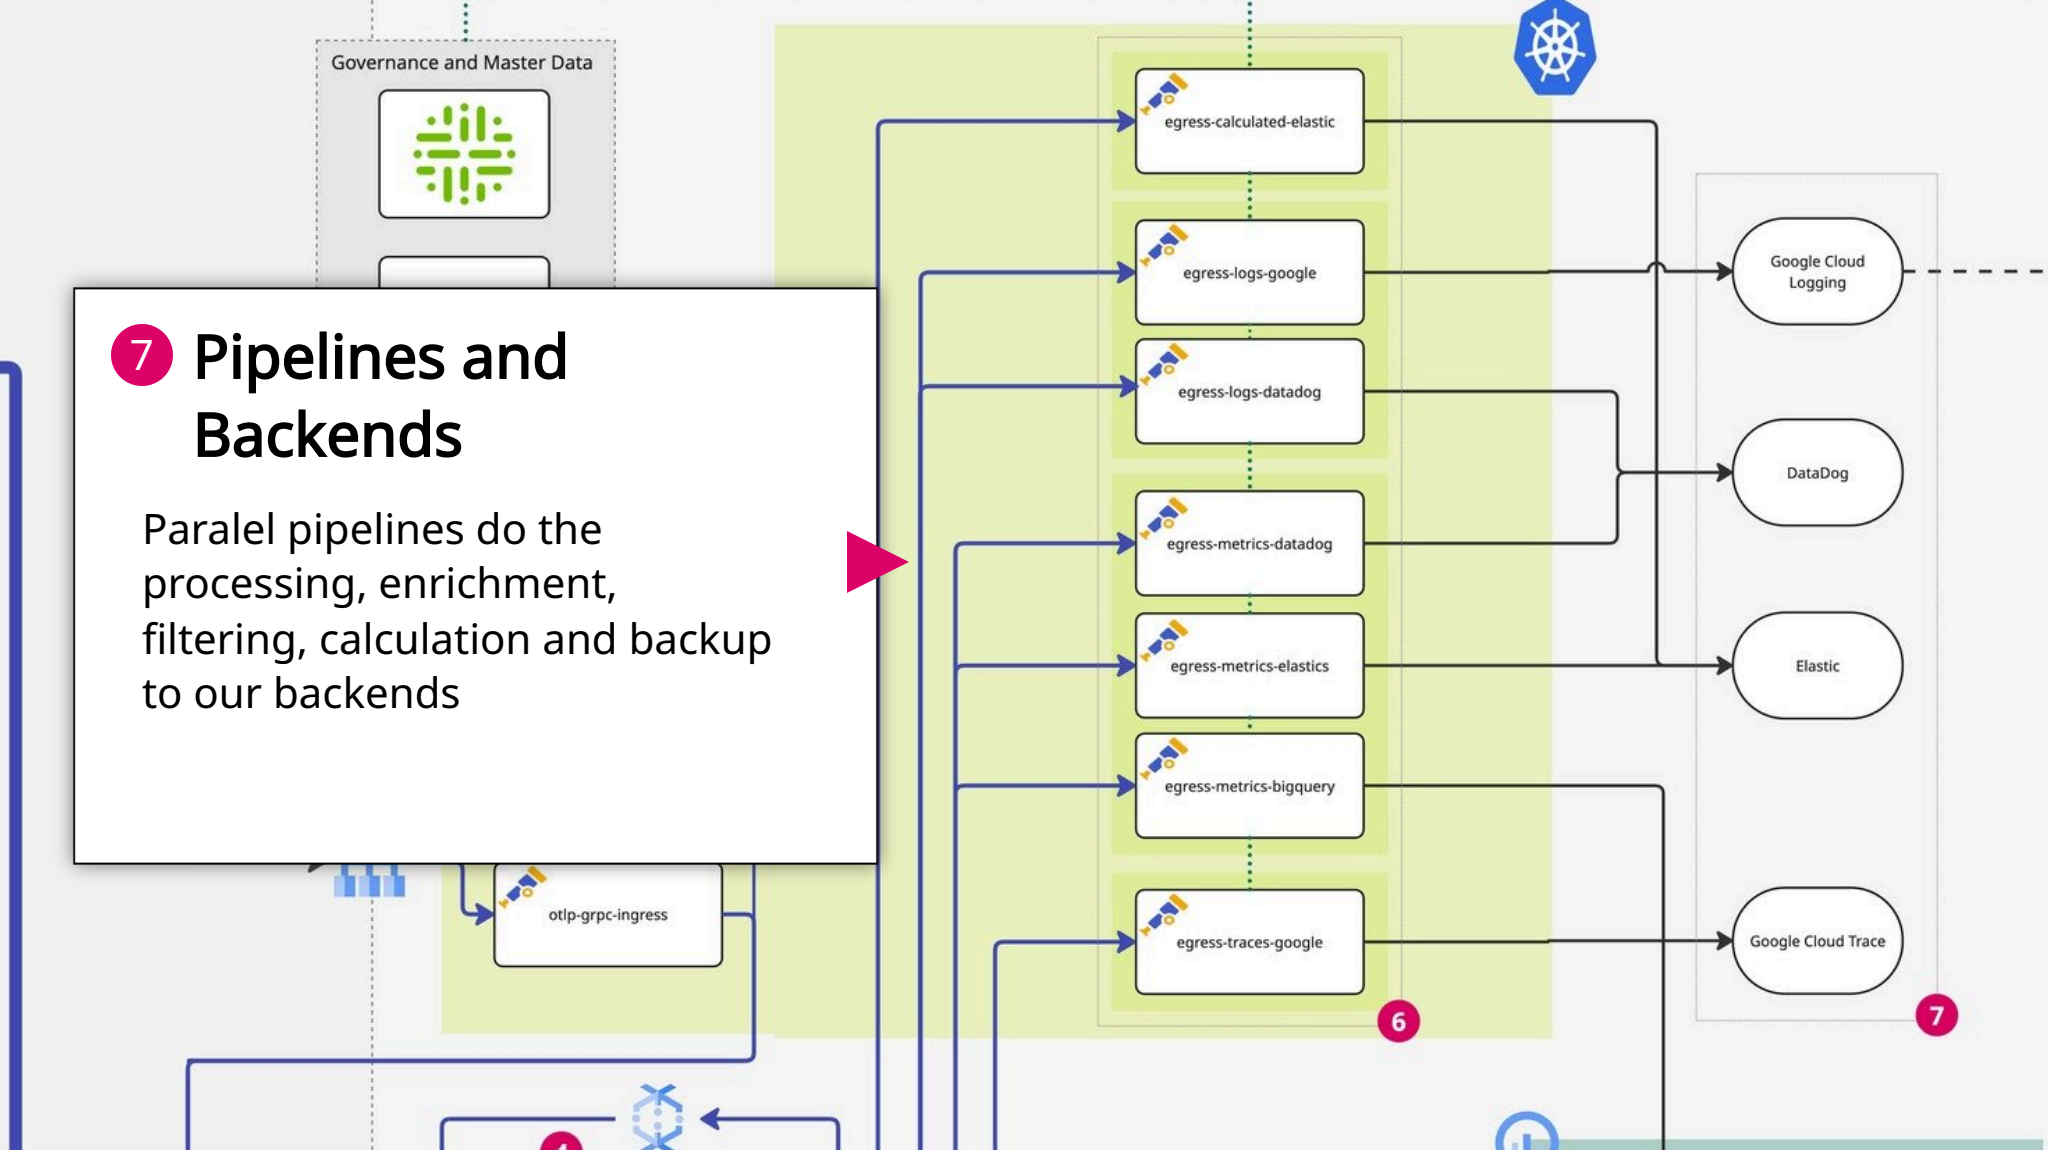

7 Pipelines and Backends

Paralel pipelines do the processing, enrichment, filtering, calculation and backup to our backends

In-Flight Enrichment from Master Data

- **JSON field promotion**: body fields → signal OpenTelemetry attributes
- **Master data lookup**: keyed by collibra.tenant.environment_id

Devs don't need to know contract terms or support levels — they just log the tenant environment ID, and the backbone dynamically infers and injects the rest.

OpenAPI Reverse-Mapping

"We don't measure URLs. We measure contracts."

- **URL cardinality** explosion bloats metric DBs and costs
- Reverse-map Istio URL + method → OpenAPI operationId
- Low-cardinality, semantic endpoint stream → automated SLOs across all microservices
- Also: aggressively drop runtime spam & infra-sweep noise before vendors see it

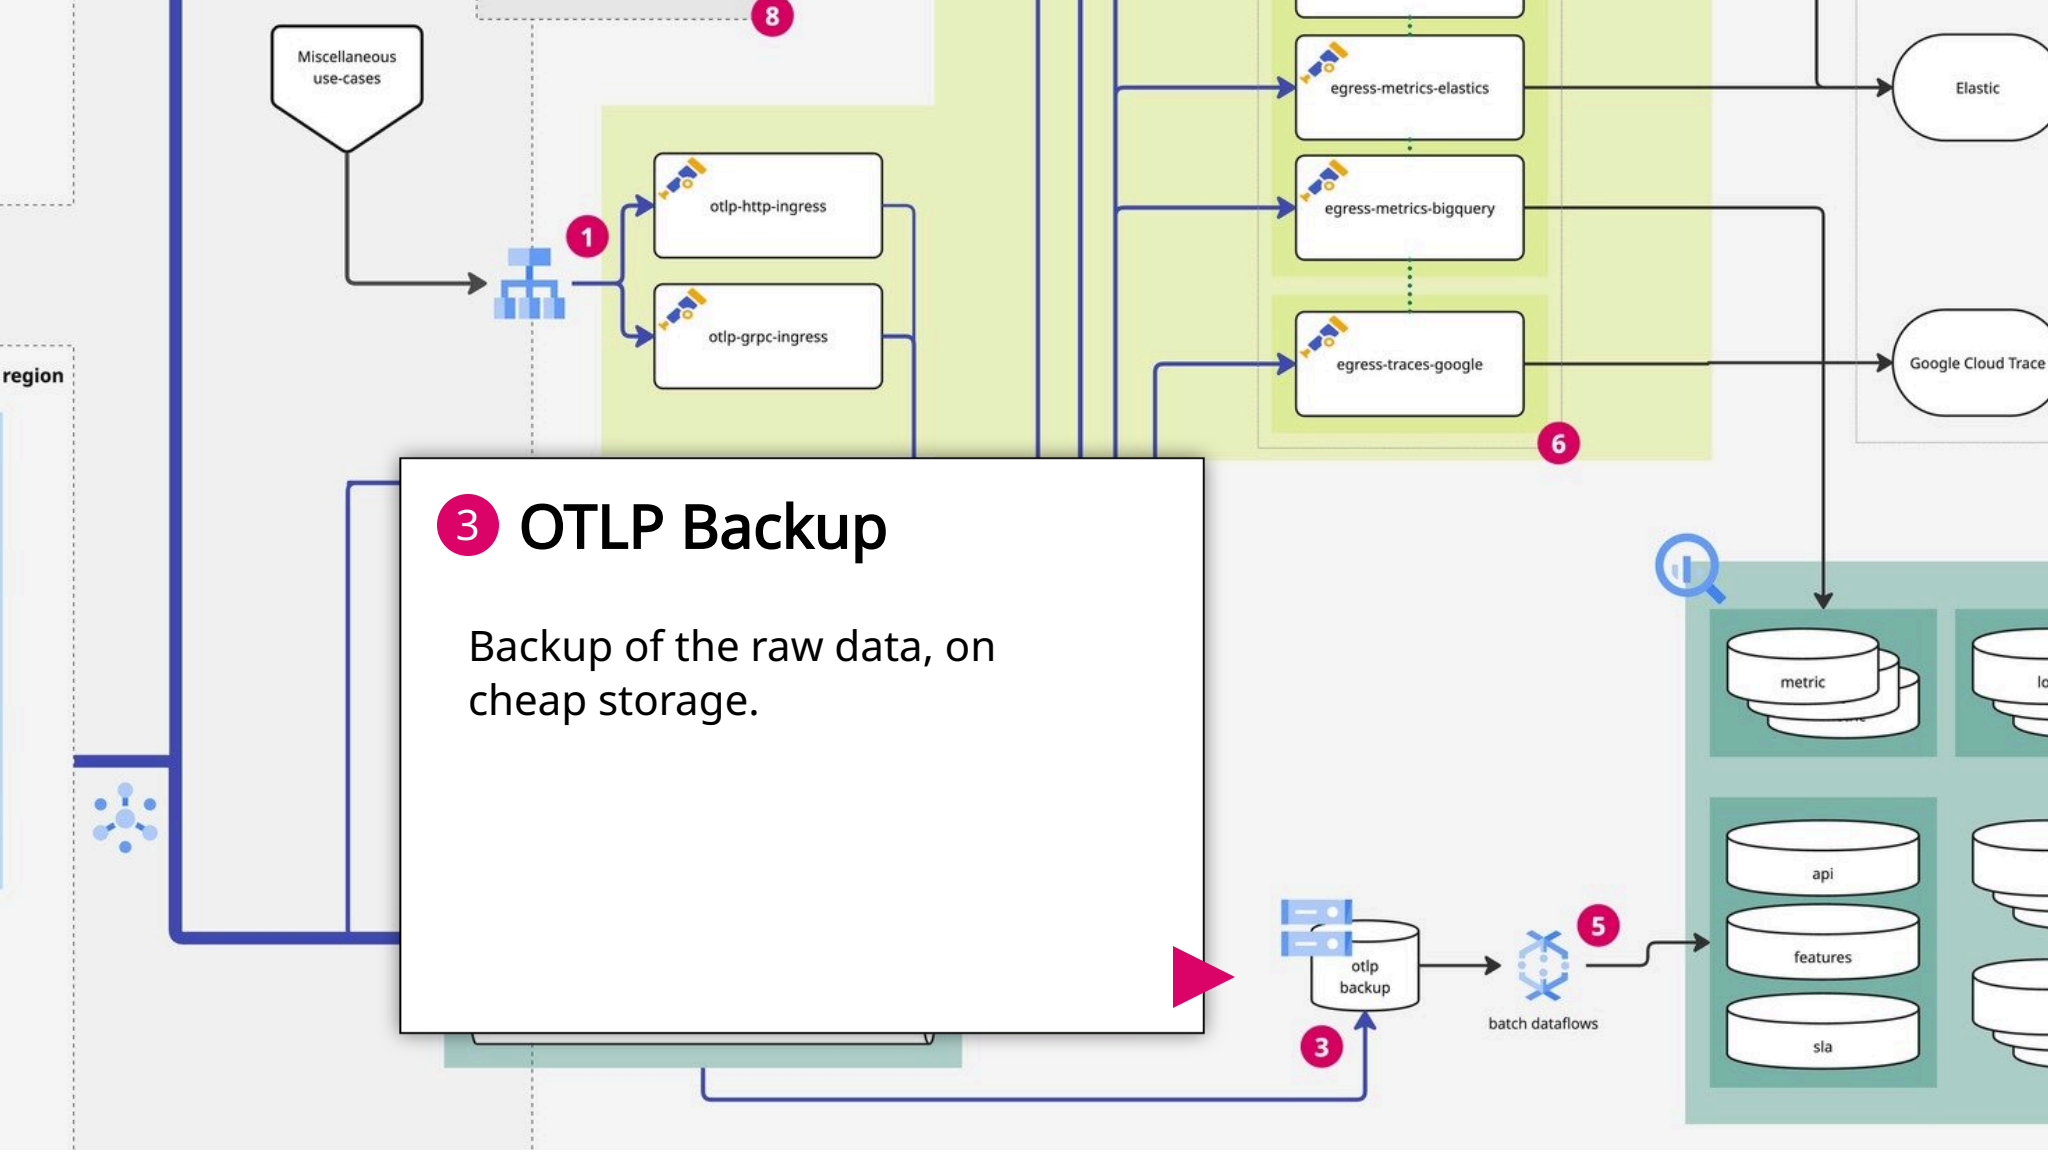
Miscellaneous use-cases
8
otlp-http-ingress
otlp-grpc-ingress
1
egress-metrics-elastics
egress-metrics-bigquery
egress-traces-google
6
Elastic
Google Cloud Trace
region
3 OTLP Backup
Backup of the raw data, on cheap storage.
metric
api
features
sla
otlp backup
3
batch dataflows
5

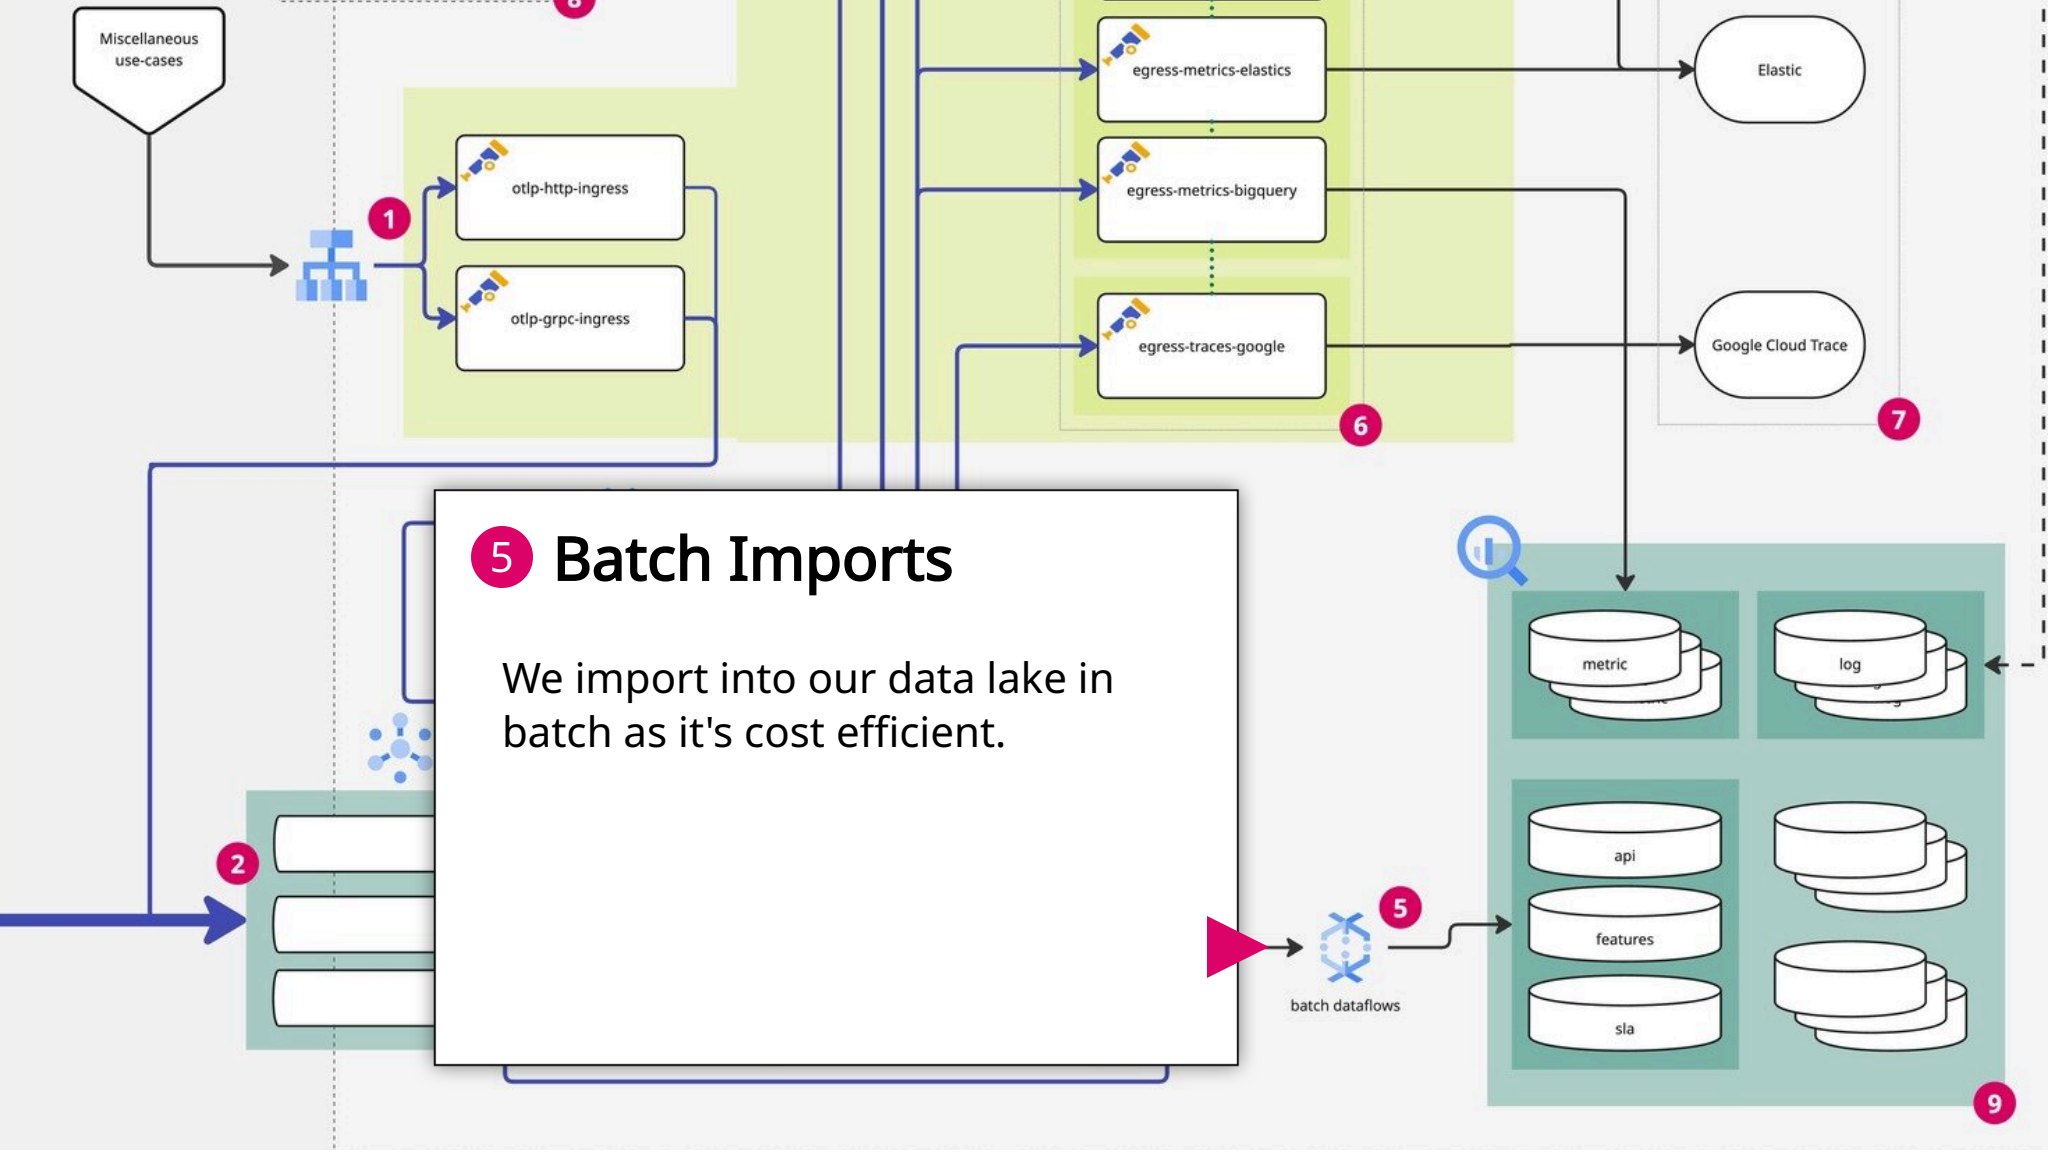

5 Batch Imports

We import into our data lake in batch as it's cost efficient.

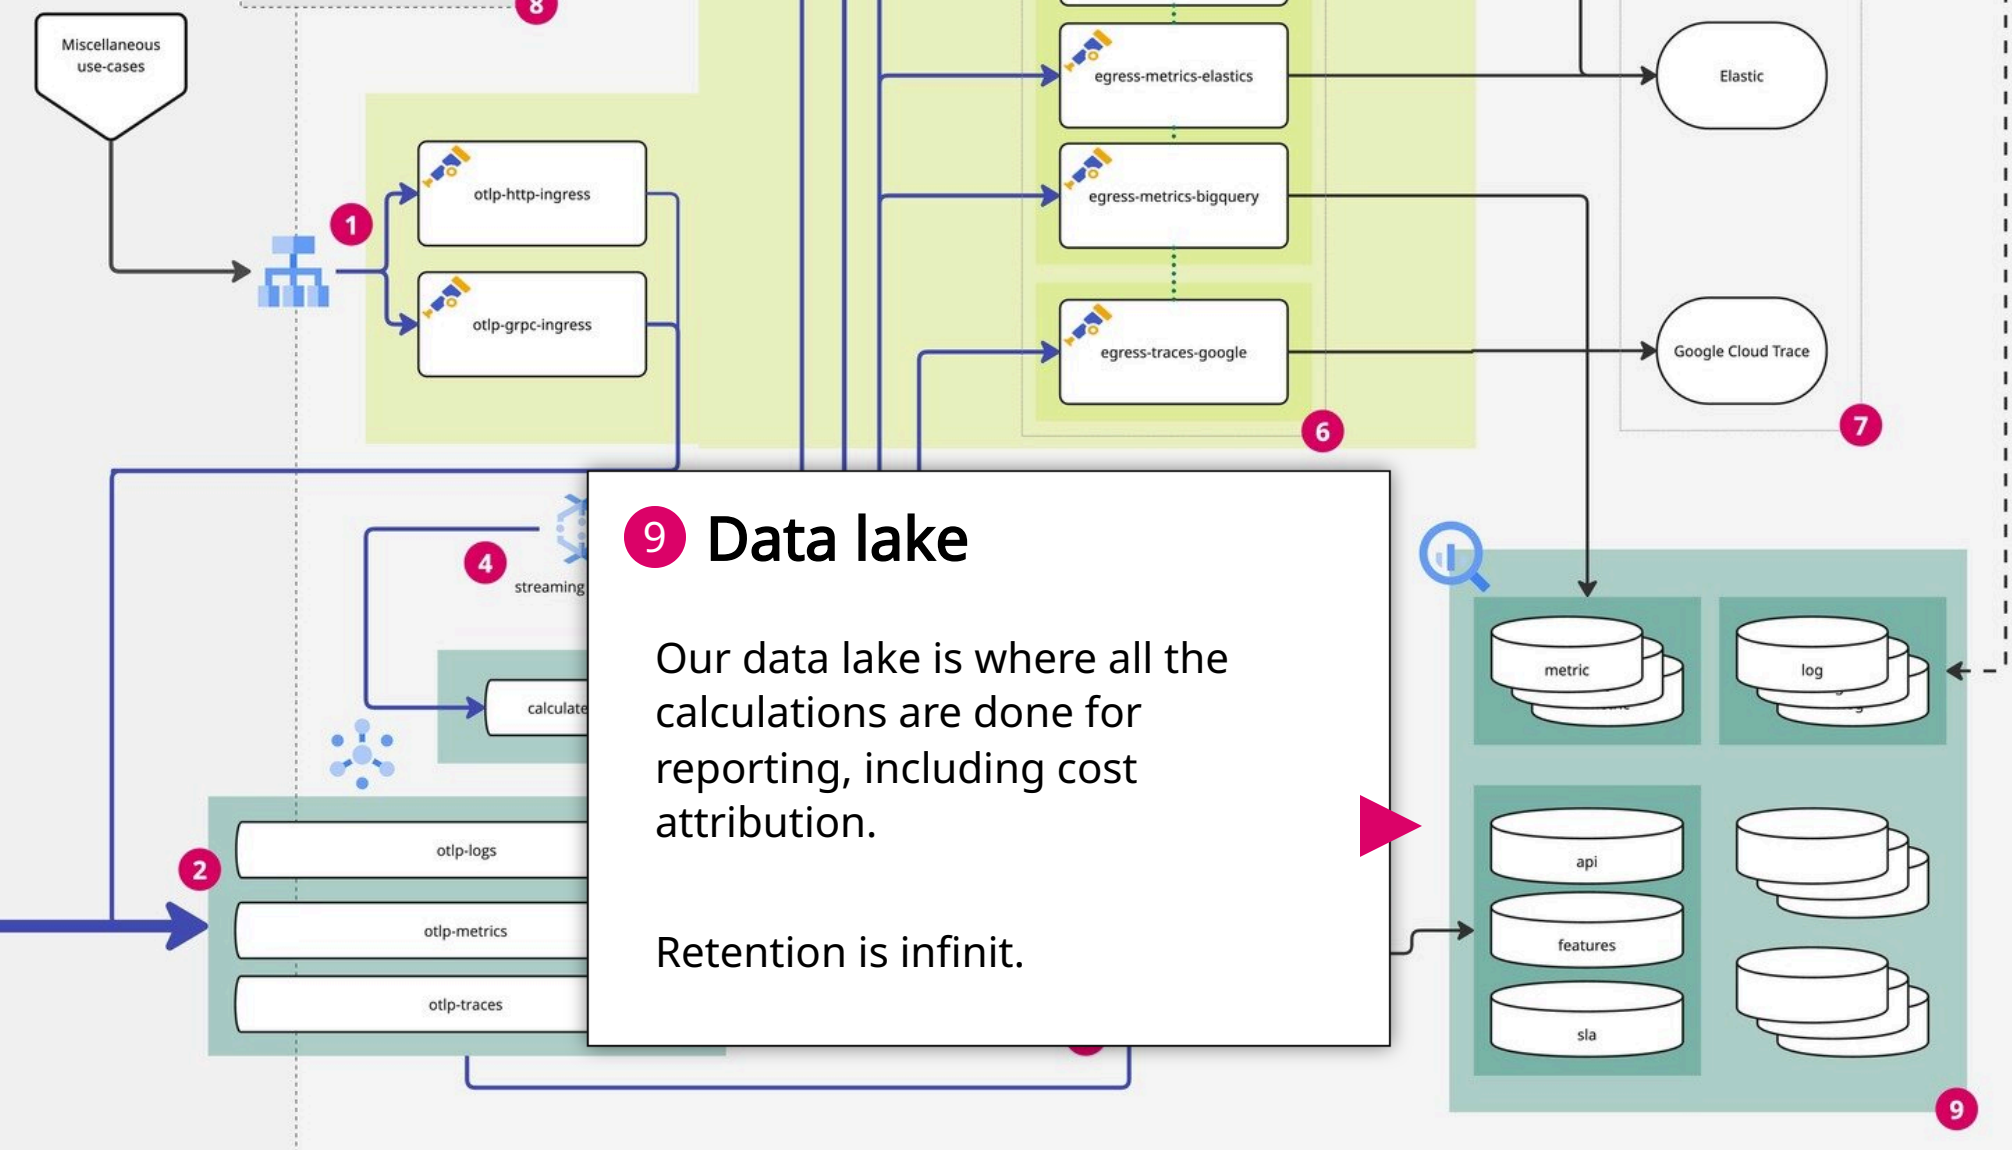

9 Data lake

Our data lake is where all the calculations are done for reporting, including cost attribution.

Retention is infinit.

Cost Attribution - Closing the Loop

- **Telemetry volume cost** — aggregate signal volume per C4 system × tenant
- **Compute cost slicing** — CPU / mem / disk / network by tenant and C4
- **C4-aware provisioning** — Collibra Infra CRDs carry C4 metadata; cloud billing maps to logical owner

Open problem: defensible "virtual dollar" formula for cross-team chargebacks.

Semantic Conventions
and wiring

Semantic Conventions

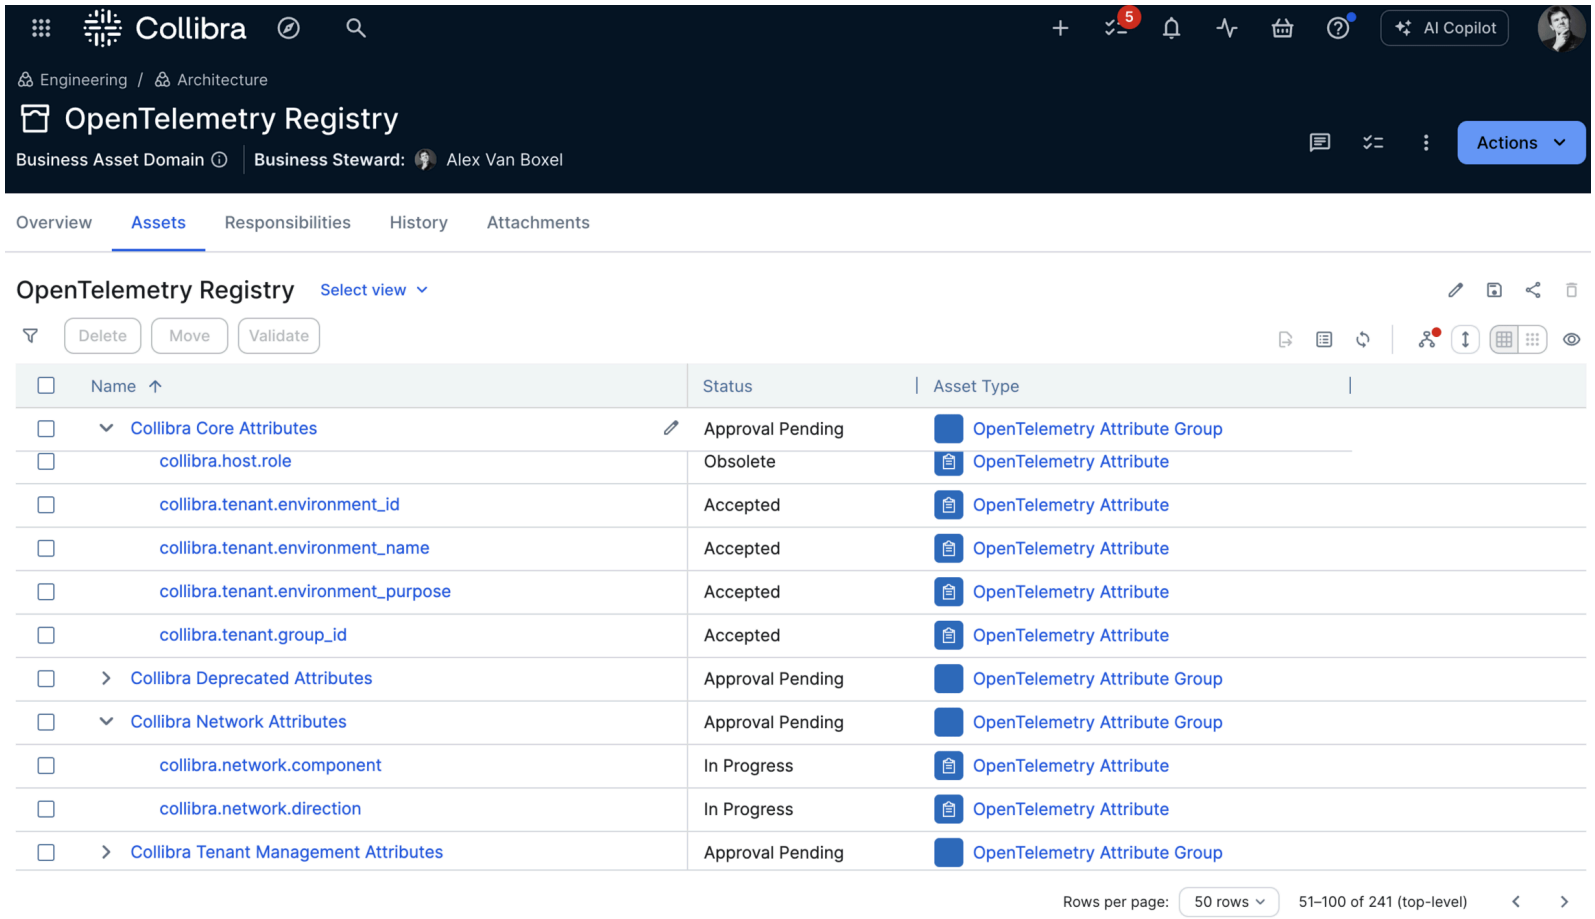

Semantic Conventions

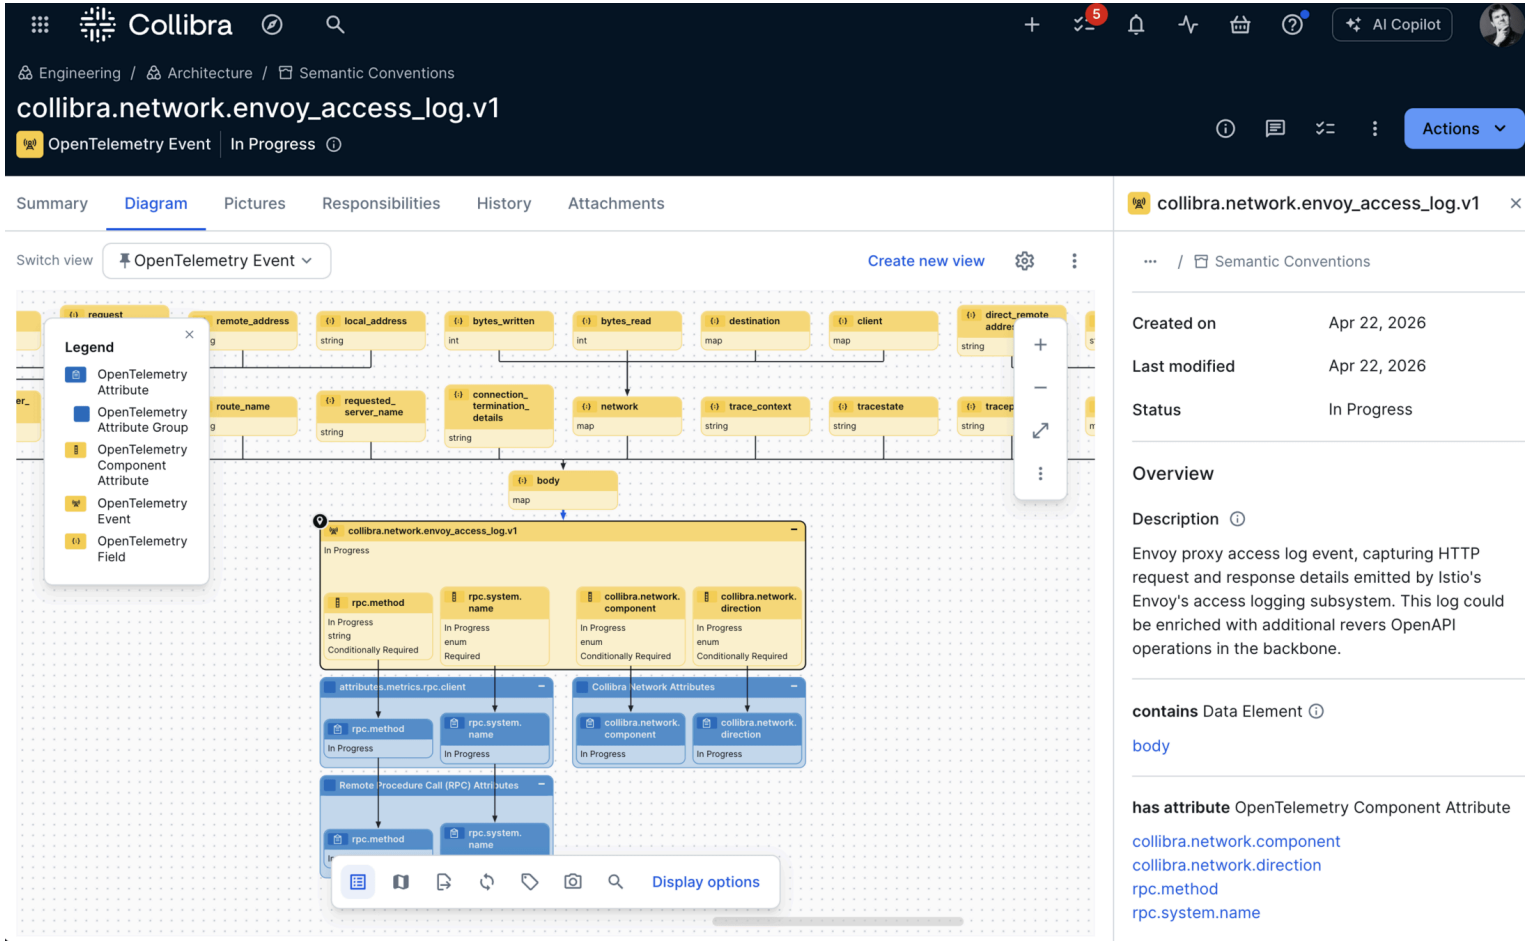

Wiring - SemConv + Weaver

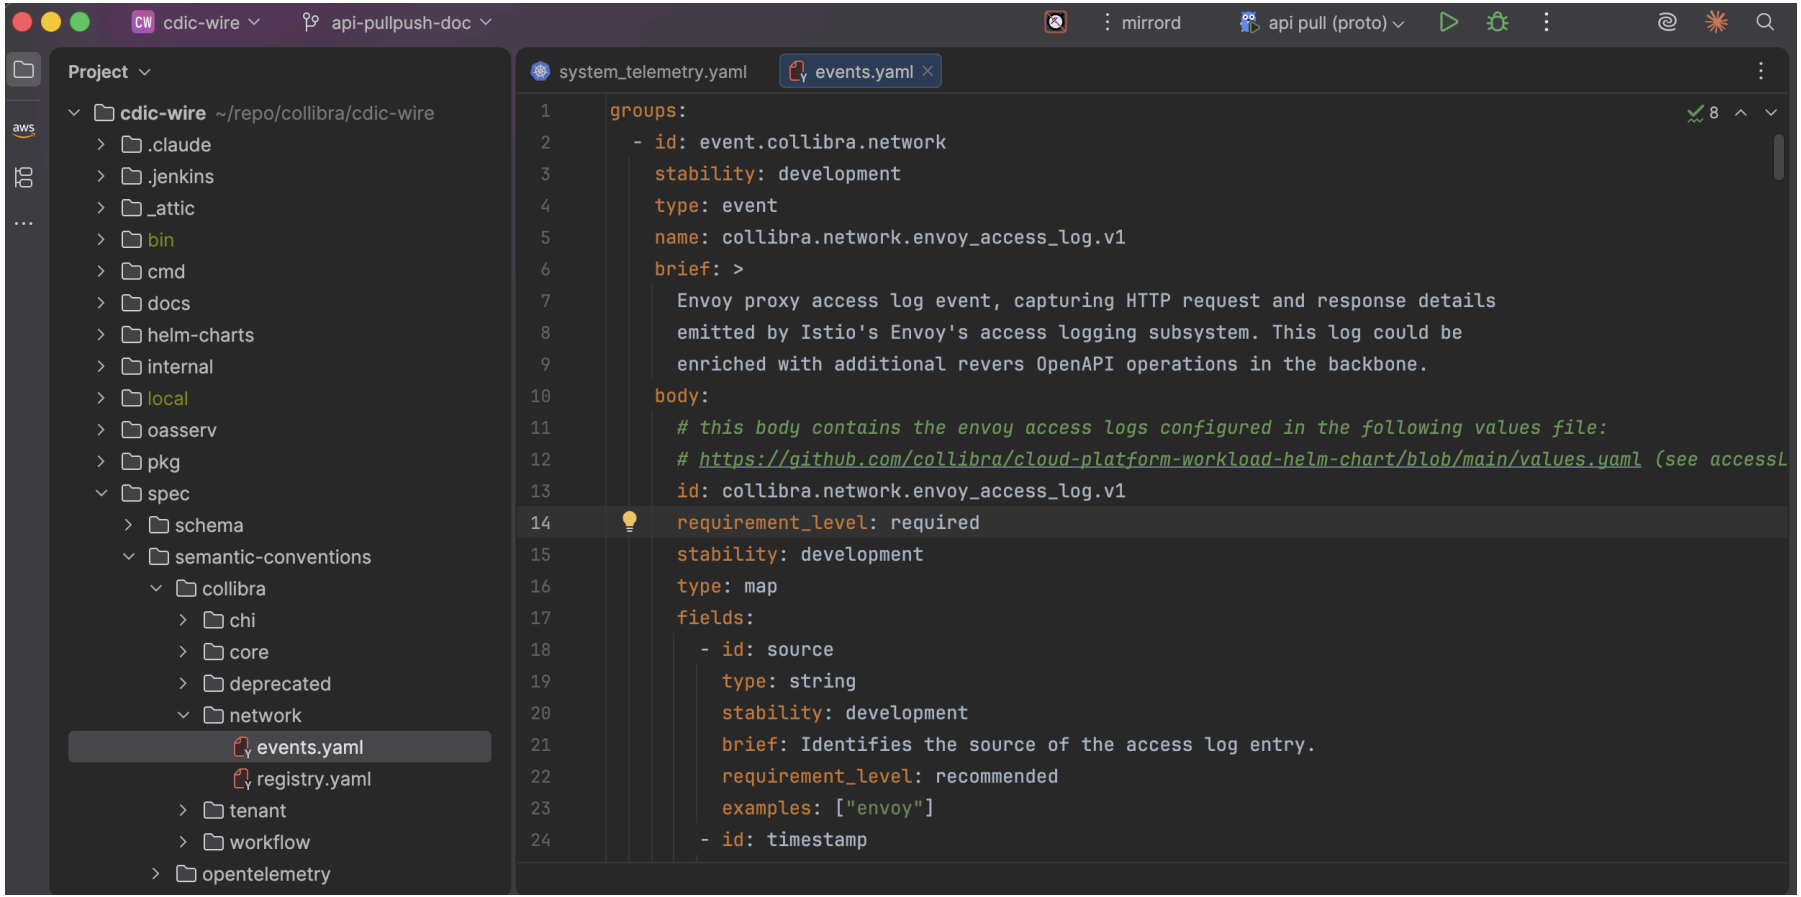

Wiring - More YAML

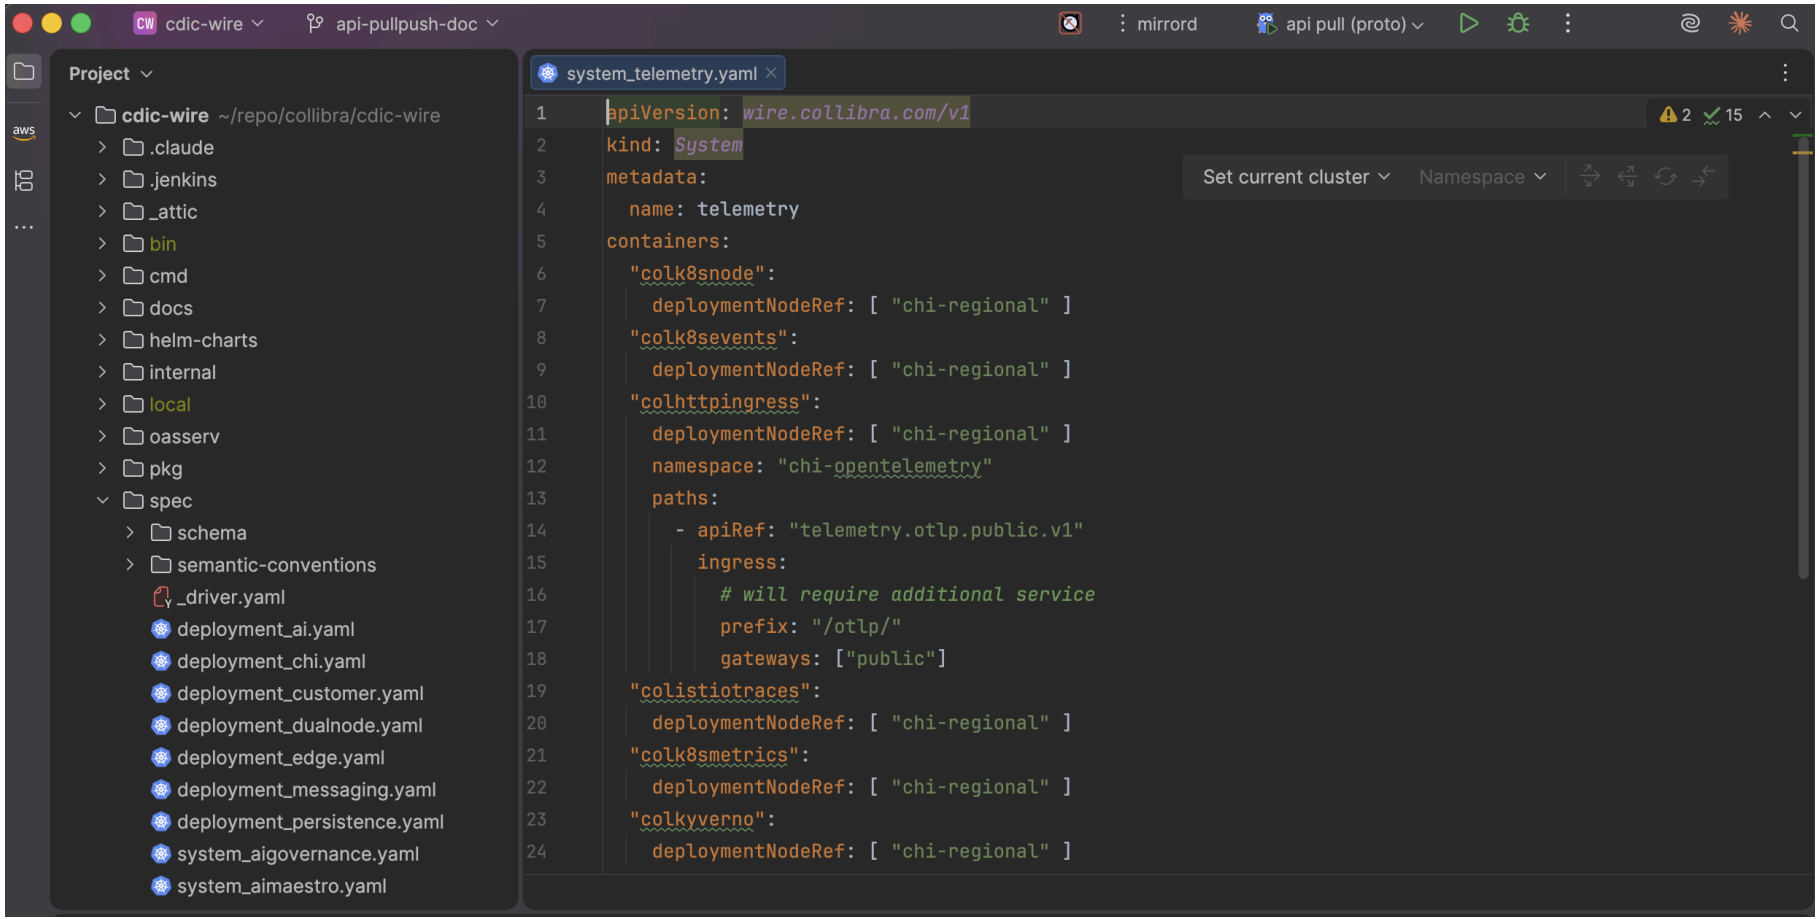

Future
and takeaway

OpAMP — Pushing Control to the Collection Edge

- Bandwidth problem: backbone filtering saves on vendors, but raw telemetry still costs WAN egress
- OpAMP: dynamic management & configuration of the entire collector fleet
- Adaptive edge sampling: normal tenant → aggressive sampling; incident → dial up fidelity at source

Key Takeaways

① Golden attributes on day one

Define tenancy and architecture dimensions before you split into microservices, not after.

② Decouple with a backbone

Buffer-first ingestion (Pub/Sub) + centralized enrichment unlocks both ops and FinOps / BI.

③ Invest in semantic contracts

They structure your signals today and become the foundation for AI diagnostic agents tomorrow.