



CROWDSTRIKE

**A principled approach to monitoring data
infrastructure at scale**

WE ARE


- Eric Schow (Sr. Manager, Engineering)
- Praveen Yedidi (Software Engineer)



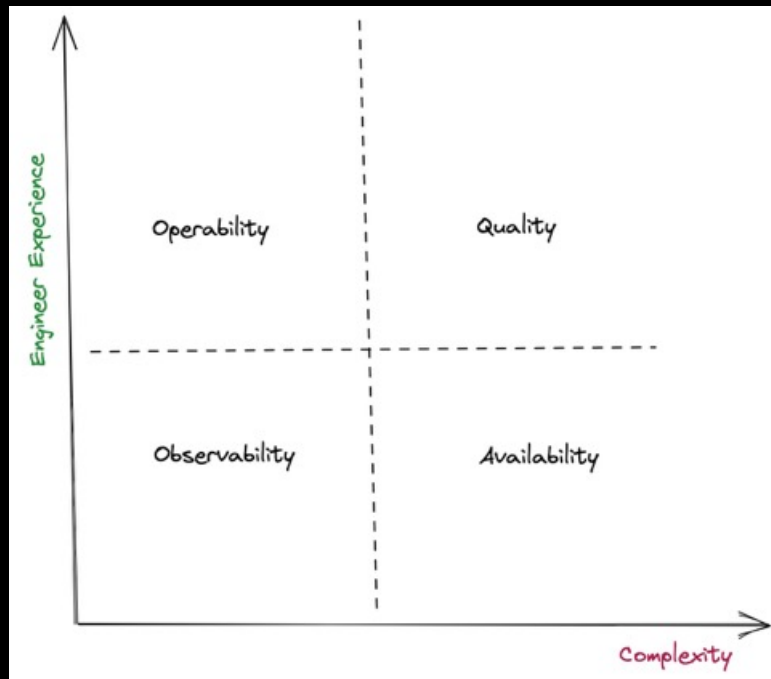
AGENDA

- Kafka at CrowdStrike
- Four Quadrant Model
 - Observability
 - Availability
 - Operability
 - Quality
- Wrap up

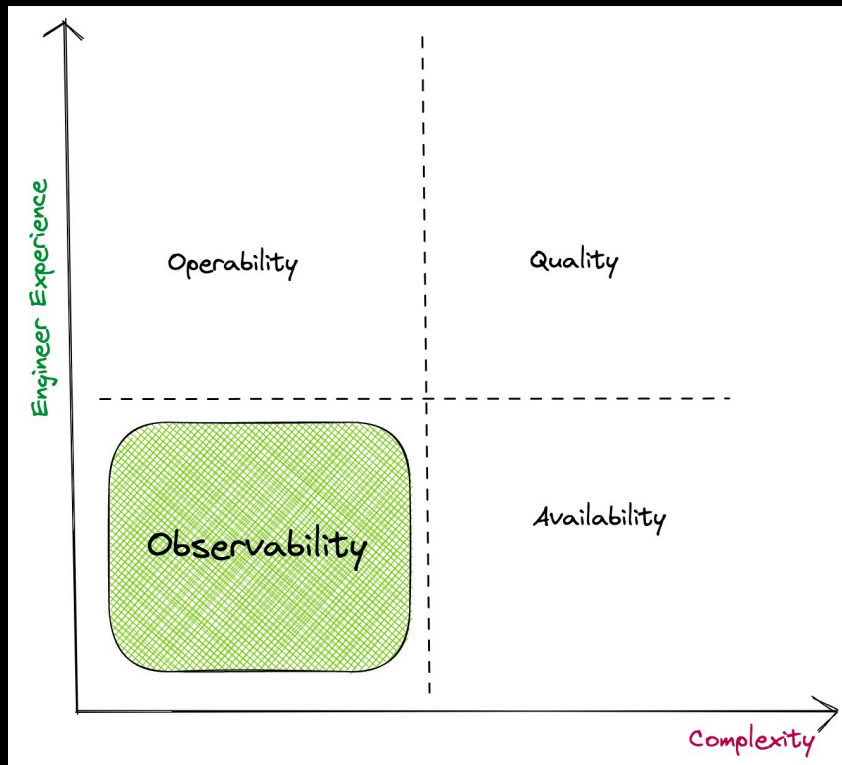
KAFKA @ CROWDSTRIKE




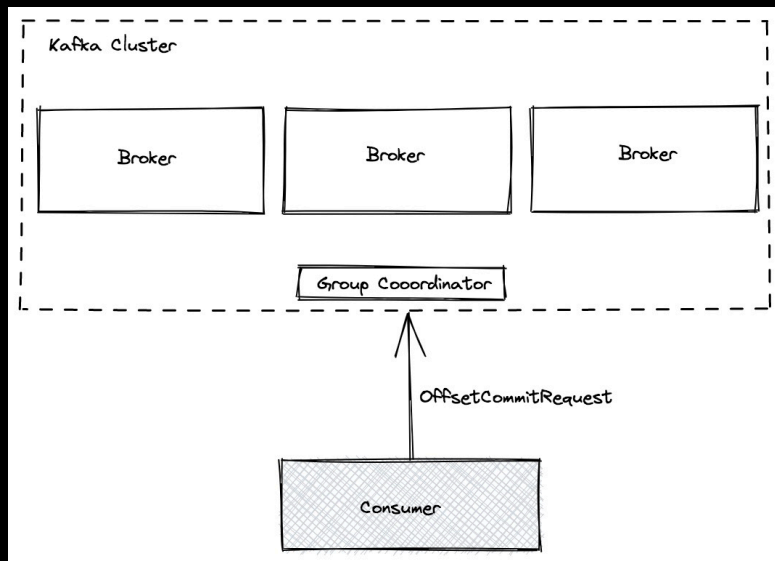
FOUR QUADRANT MODEL



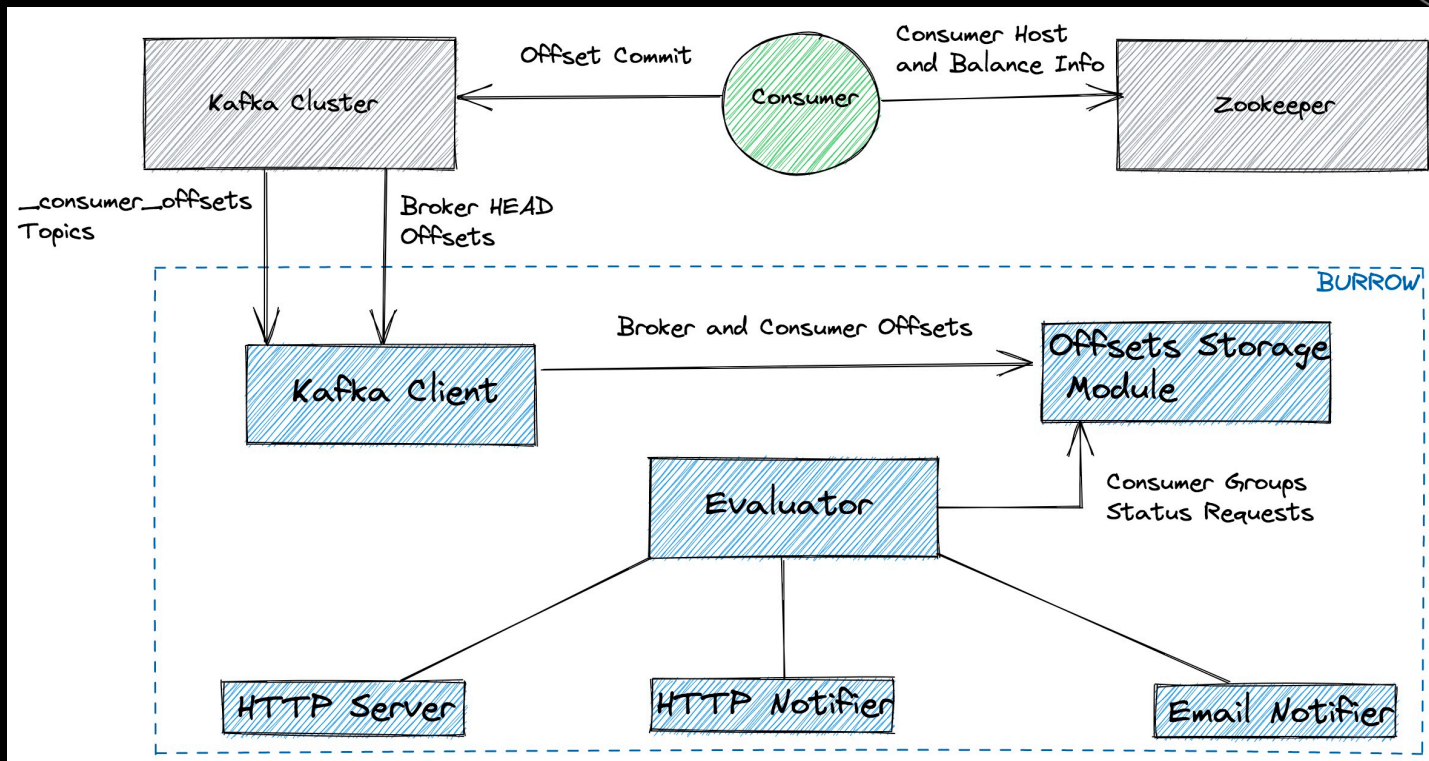
OBSERVABILITY



CONSUMER LAG MONITORING



MONITORING LAG WITH BURROW



STATUS & LAG FROM BURROW

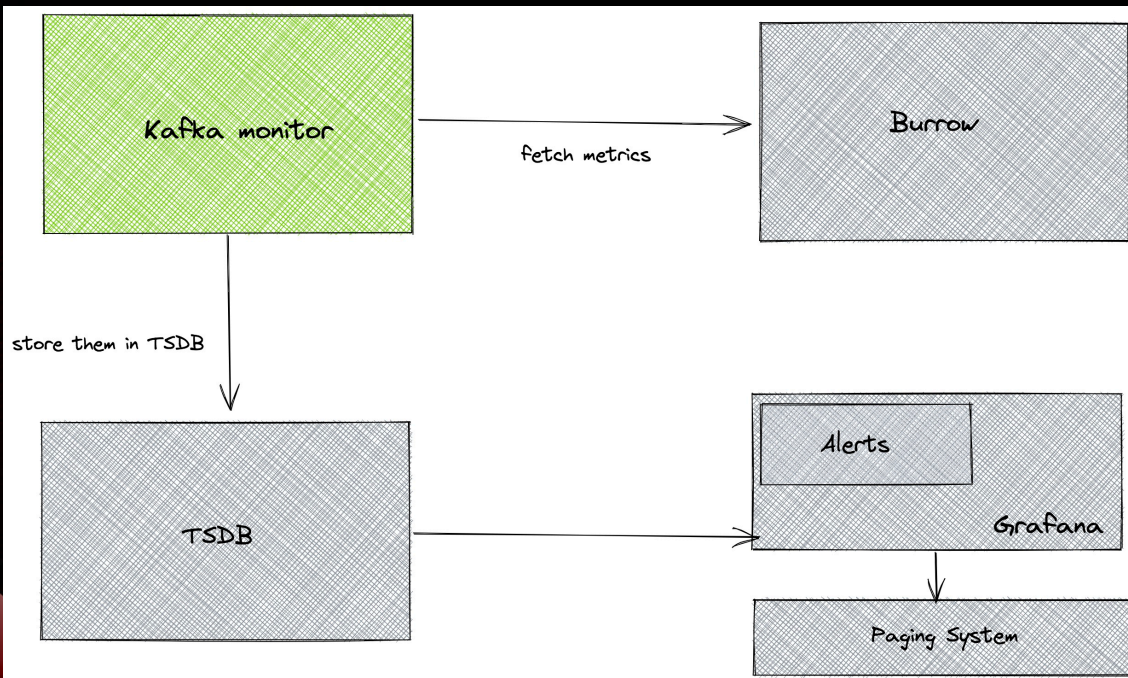
```
{
  "error": false,
  "message": "consumer status returned",
  "status": {
    "cluster": "dodo-red-main-s001-kafka-processing01",
    "group": "csm.datarouter",
    "status": "OK",
    "complete": 1,
    "partitions": [],
    "partition_count": 2,
    "maxlag": {
      "topic": "cs.csm.request.mutation",
      "partition": 0,
      "owner": "/10.X.X.X",
      "client_id": "csmdatarouter-594697b755-8ztrm",
      "status": "OK",
      "start": {
        "offset": 364152158,
        "timestamp": 1631759432556,
        "observedAt": 1631759489000,
        "lag": 0
      },
      "end": {
        "offset": 364159684,
        "timestamp": 1631759984555,
        "observedAt": 1631759993000,
        "lag": 0
      },
      "current_lag": 0,
      "complete": 1
    },
    "totallag": 0
  },
  "request": {
    "url": "/v3/kafka/dodo-red-main-s001-kafka-processing01/consumer/csm.datarouter/status",
    "host": "ip-10-X-X-X"
  }
}
```



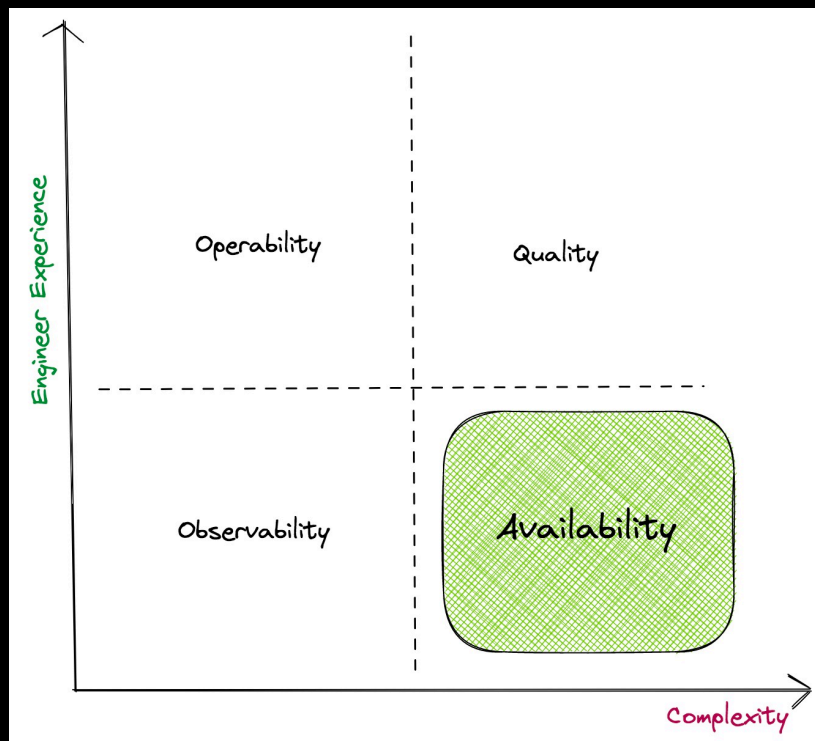
KAFKA MONITOR

Benefits:

- Historical trends
- Velocity calculations (ex: Mean Recovery time)

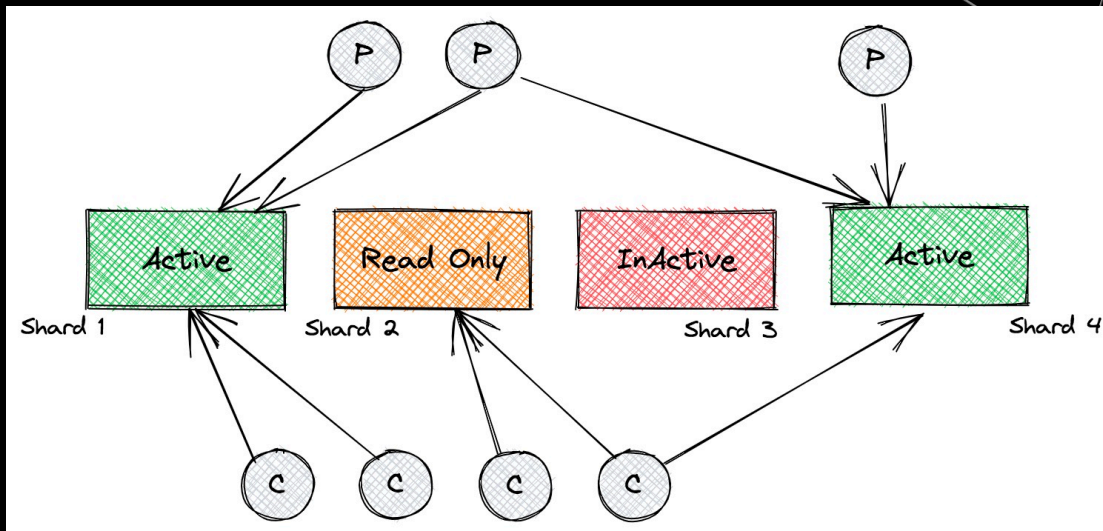


AVAILABILITY




SHARDING KAFKA FOR INCREASED AVAILABILITY & RELIABILITY

- Conceptual sharding of Kafka into shards via parallel clusters
- Nodes are configured to be active, read-only, and inactive states

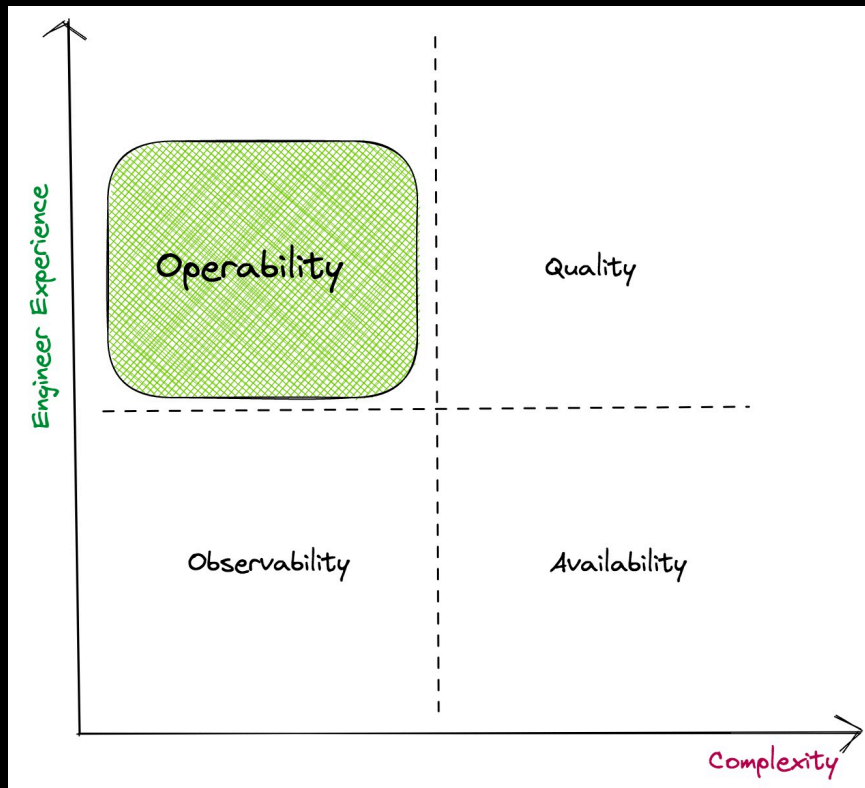


BENEFITS

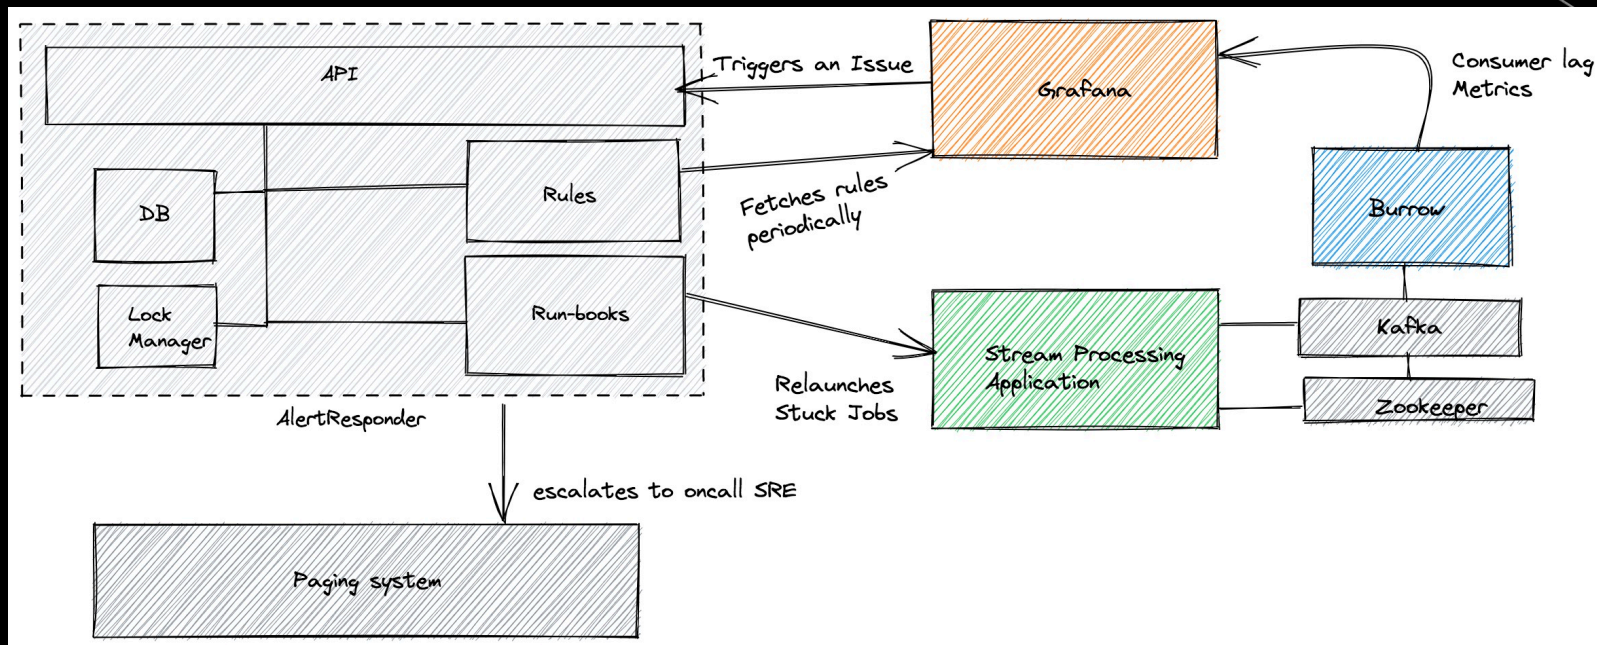
- Infinite Scaling
- Ease of maintenance
- Fault reliability



OPERABILITY

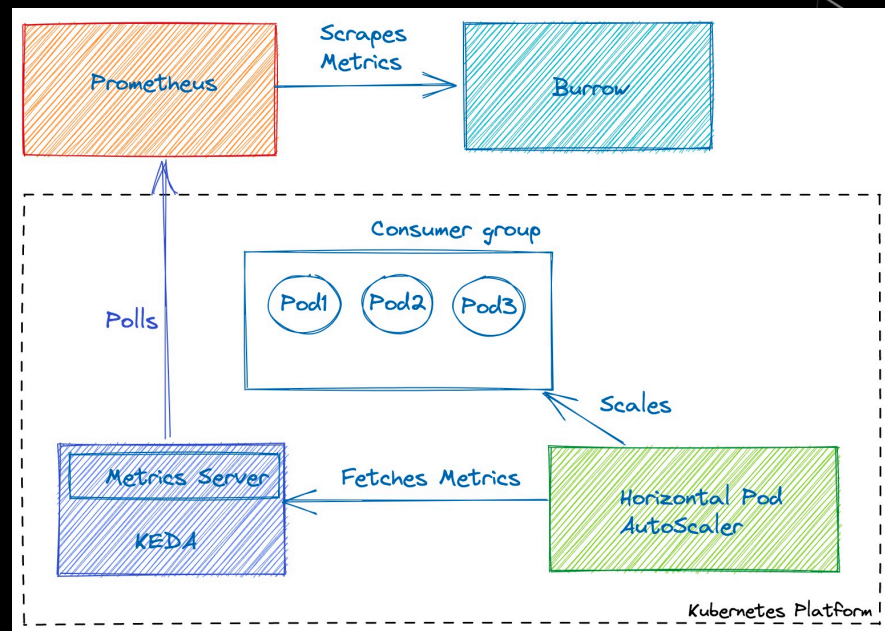


ALERT RESPONDER

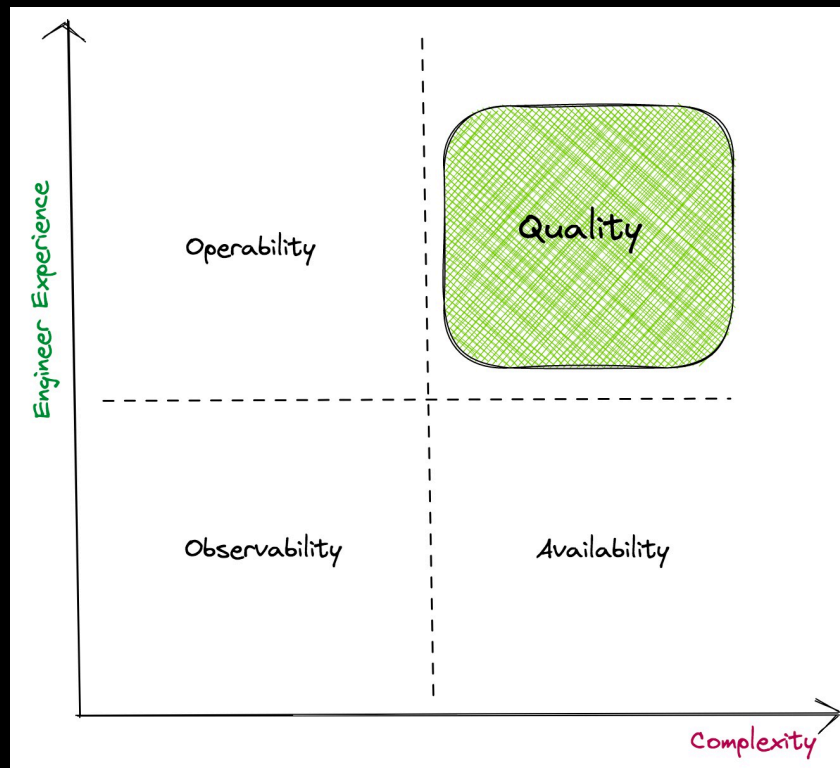


AUTOSCALING

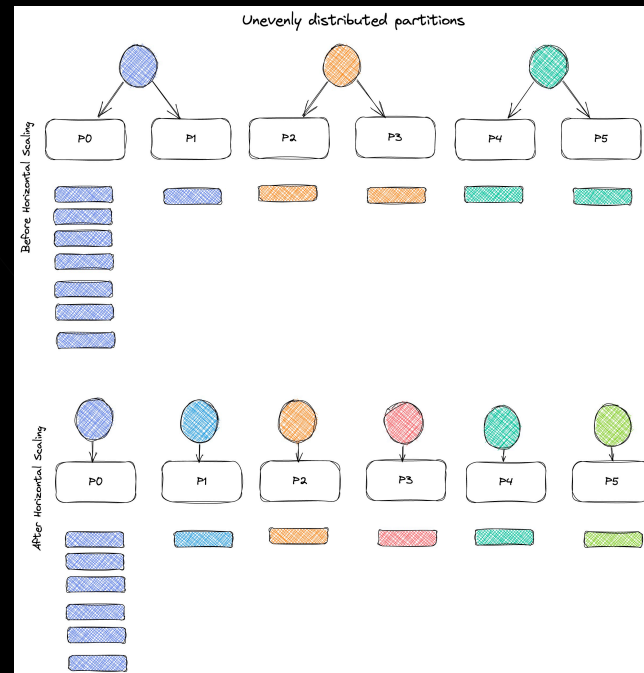
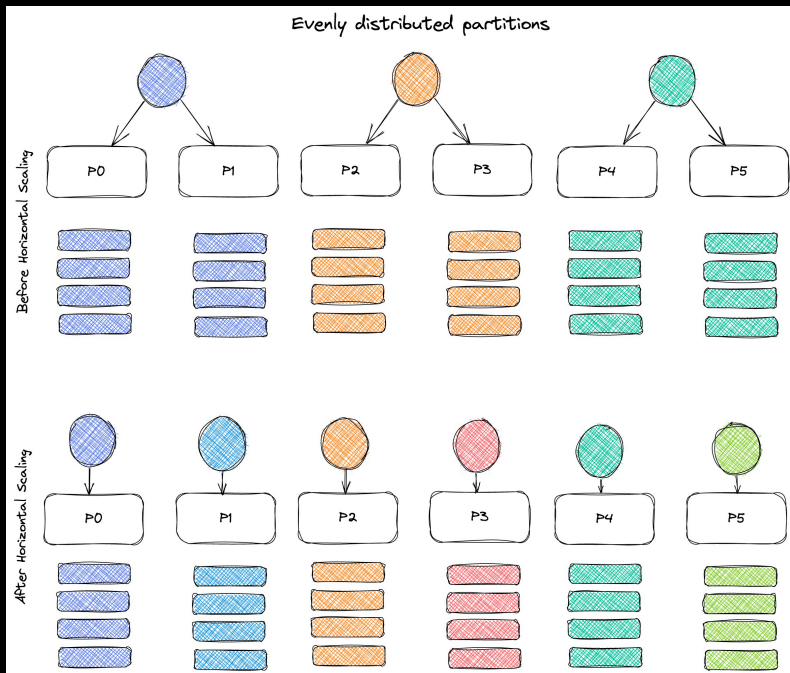
- Scheduled Scaling
- Scaling based on consumer lag



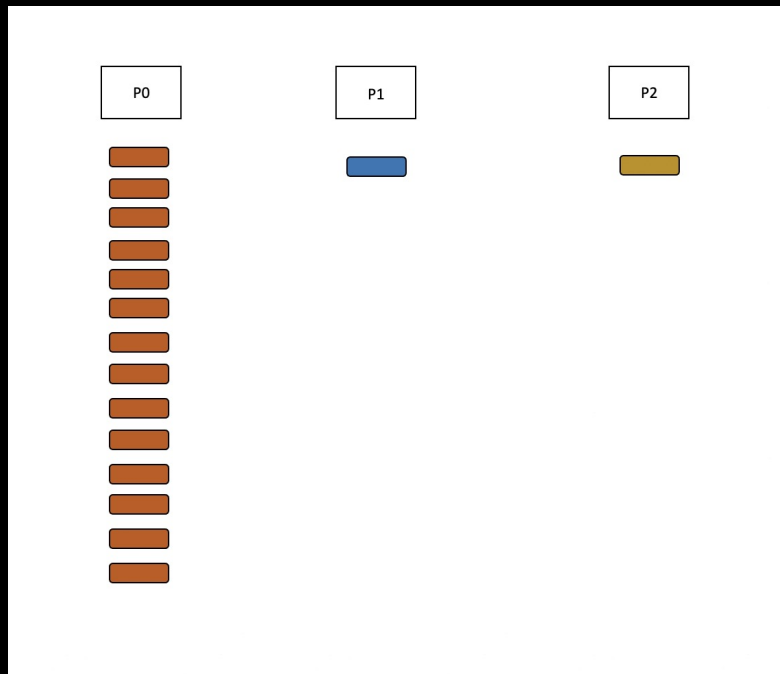
QUALITY



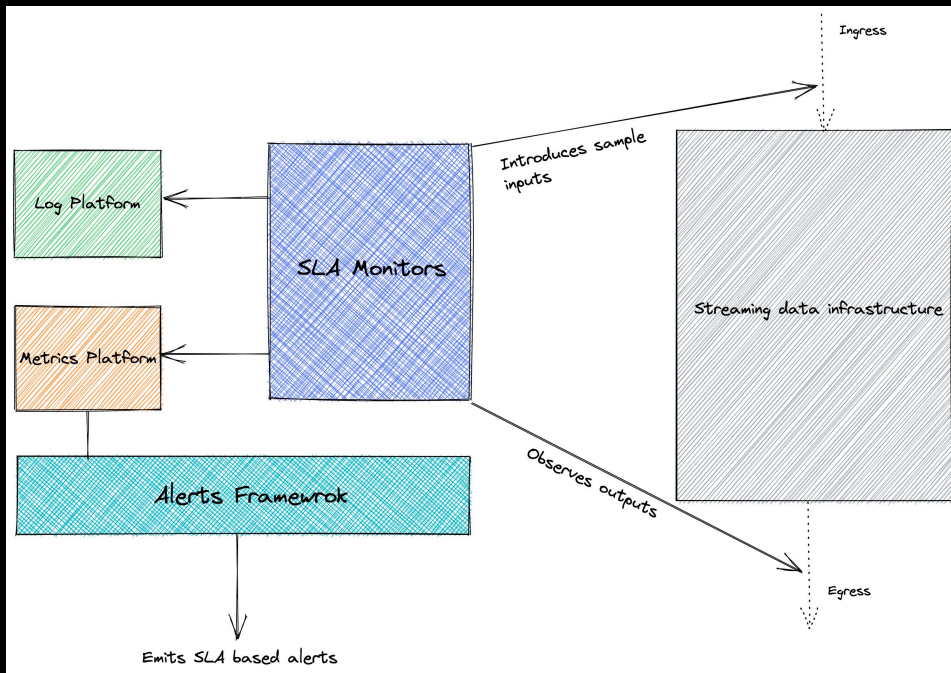
ADDRESSING UNEVEN PARTITION LAG IN KAFKA



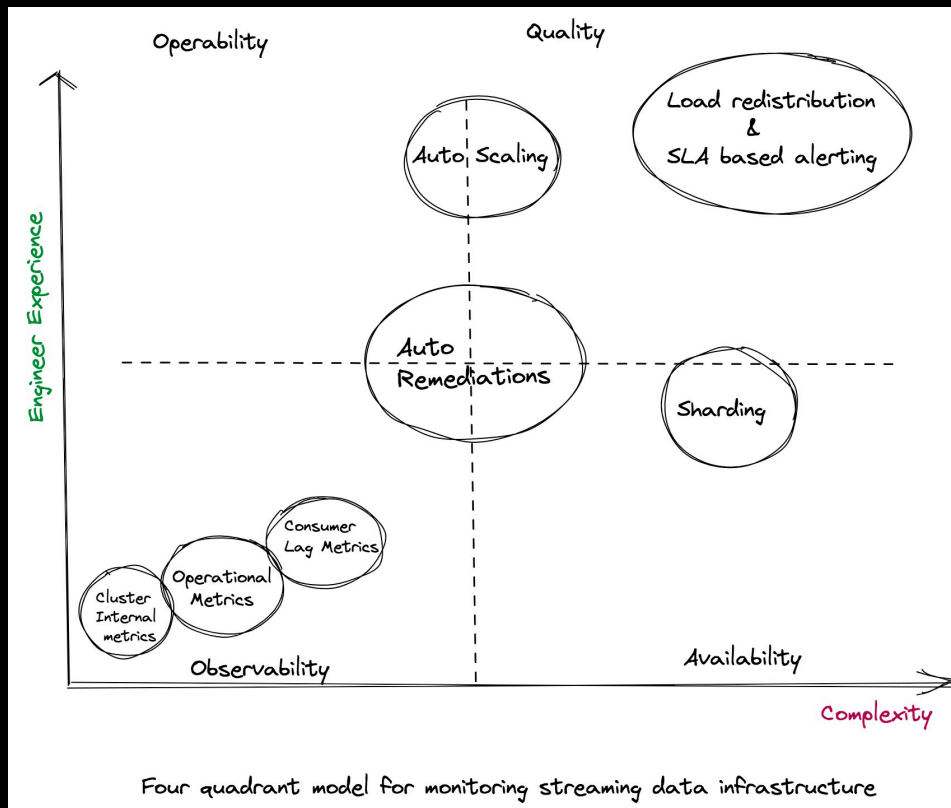
ADDRESSING UNEVEN PARTITION LAG IN KAFKA



SLA BASED MONITORING



WRAPPING UP



Thank You

Join us: <https://crowdstrike.com/careers>

Latest Blogs: <https://crowdstrike.com/blog>

