

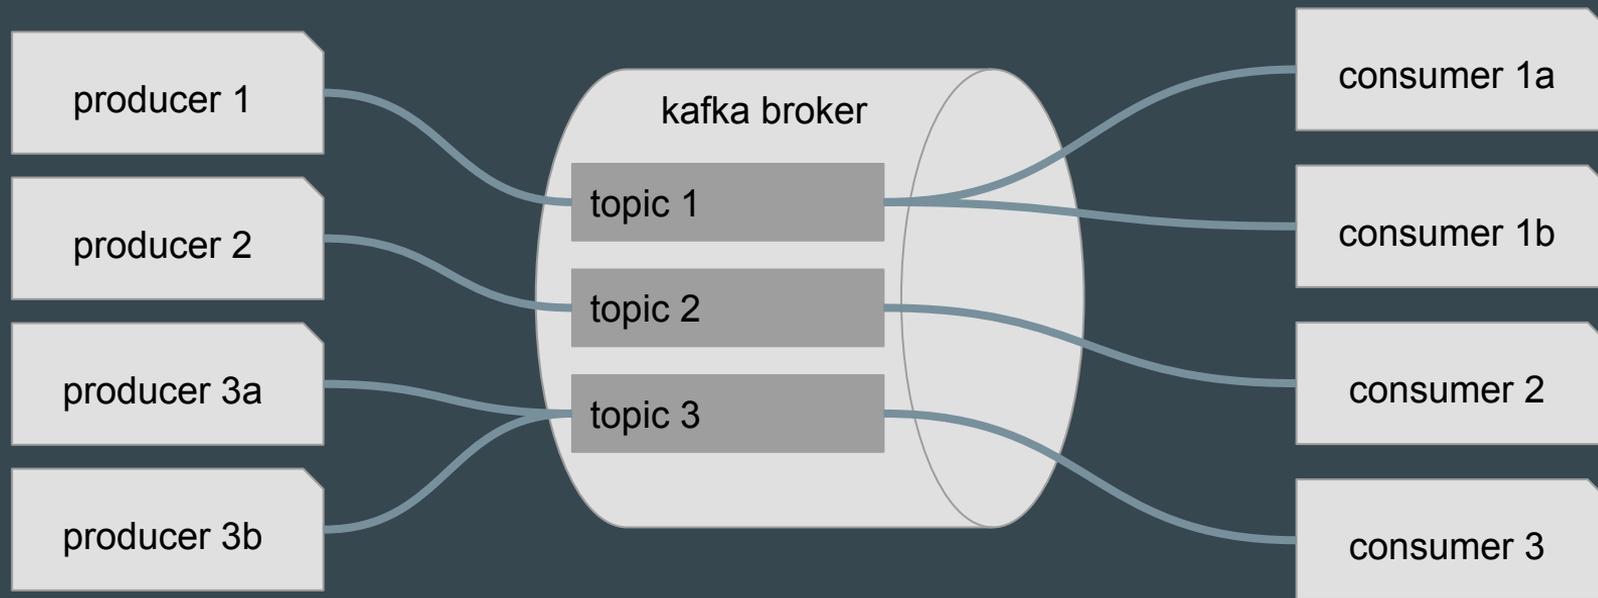
kafka and InfluxDB Grafana Update

...

28 Oct 2015

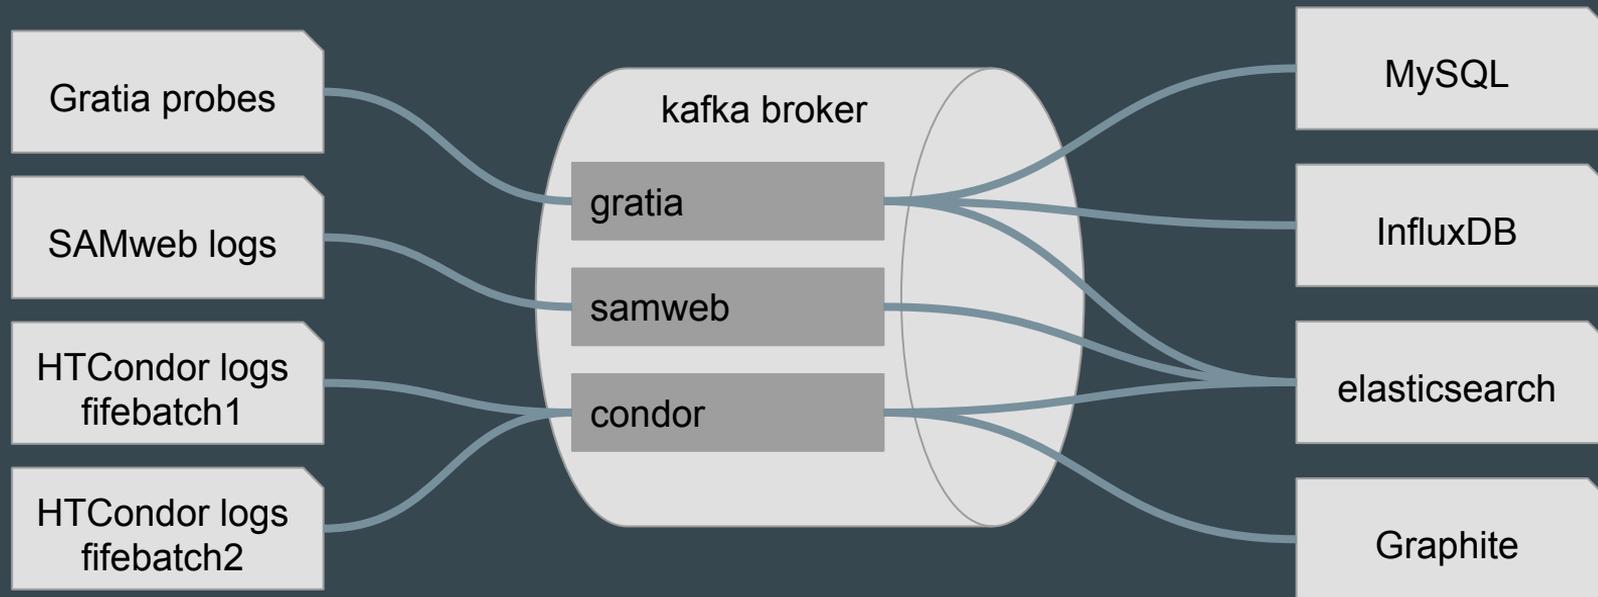
kafka

kafka (<http://kafka.apache.org>) is a distributed, scalable, pub/sub message broker



kafka

could act as a central hub for all monitoring and accounting data



InfluxDB

InfluxDB (<https://influxdb.com/>) is a time-series database like graphite, with some differences:

- Data is stored in LevelDB key-value store
- Data is stored by measurement, with arbitrary string tags and values
- SQL-like query language
- Aggregations are done via “continuous queries” that summarize data from one measurement into another
- Concerns: relatively young (circa 2014ish), data size versus RRD/whisper, performance

Grafana 2.5 released today!

- UI improvements
- Elasticsearch data source
- Cloudwatch data source

Running on Fifemon dev (<https://fermicloud079.fnal.gov/monitor>) now!