Sasquatch for DFs

Frossie & Angelo

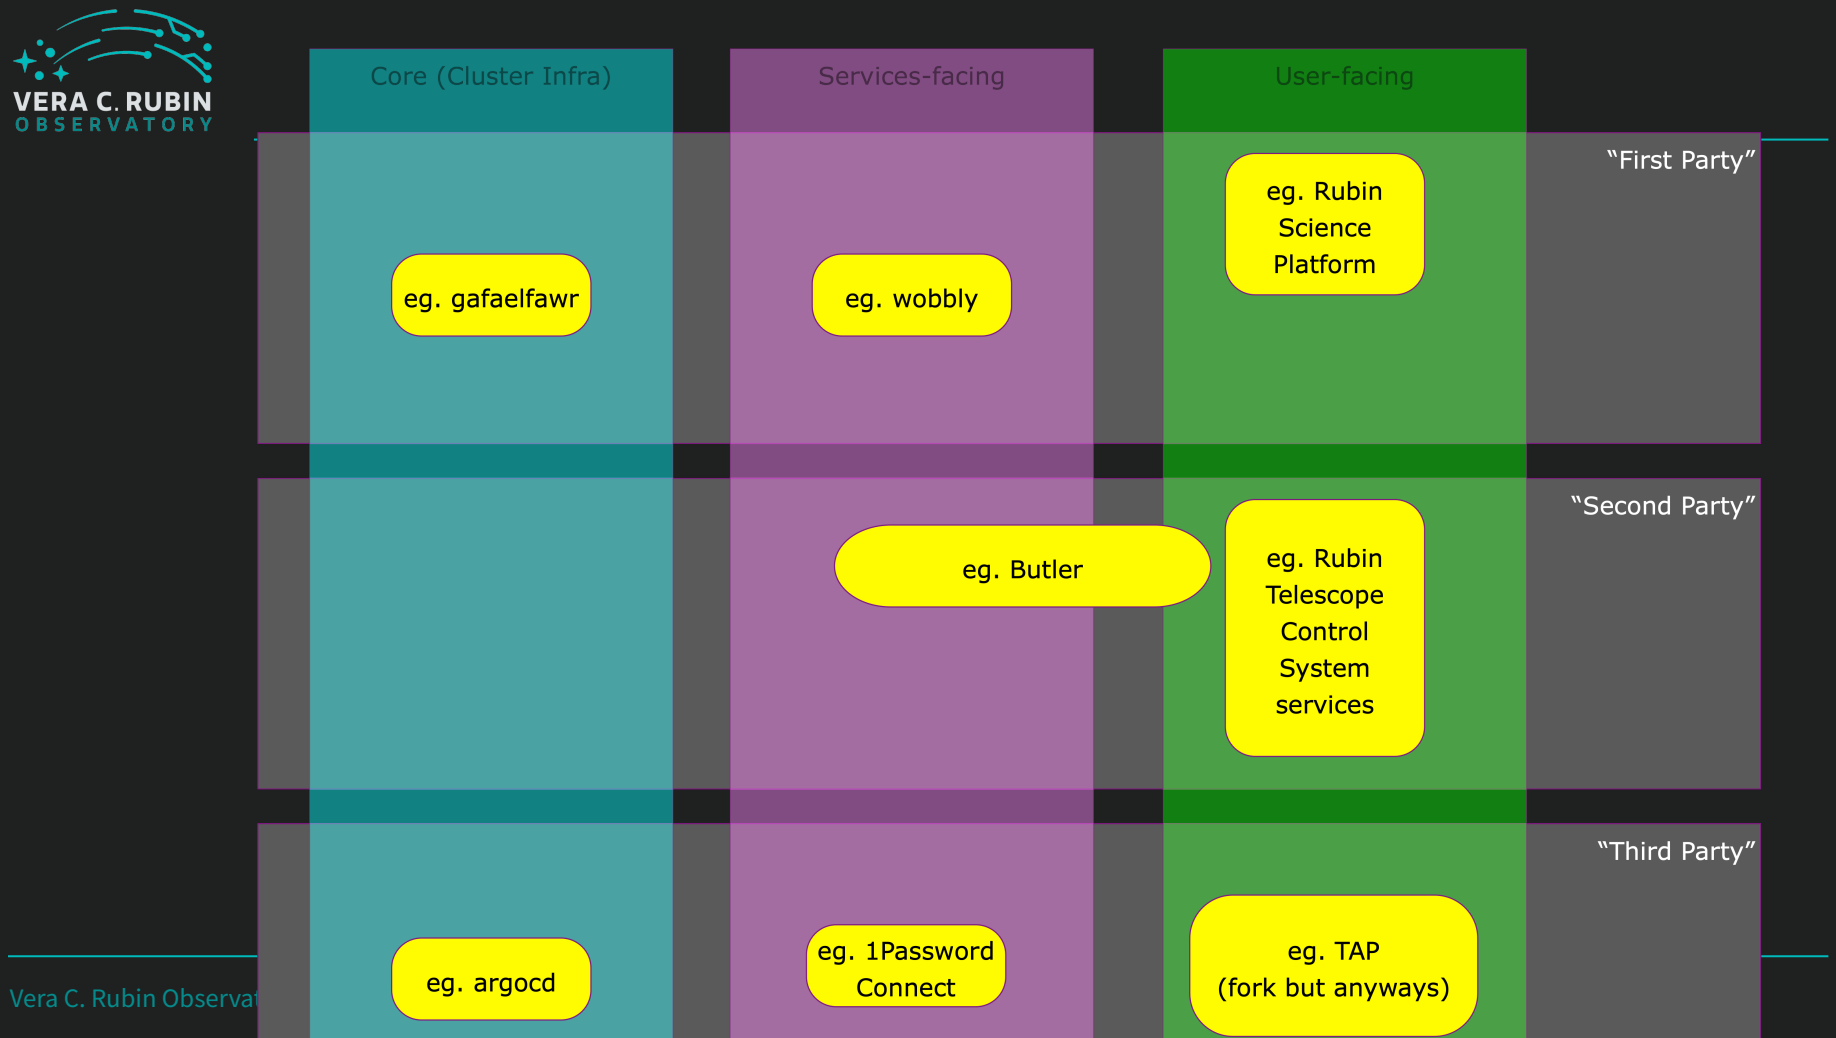
Core (Cluster Infra)
Services-facing
User-facing
"First Party"
eg. gafaelfawr
eg. wobbly
eg. Rubin Science Platform
"Second Party"
eg. Butler
eg. Rubin Telescope Control System services
"Third Party"
eg. argocd
eg. 1Password Connect
eg. TAP (fork but anyways)

VERA C. RUBIN OBSERVATORY

	Core (Cluster Infra)	Services-facing	User-facing
"First Party"	eg. gafaelfawr	eg. wobbly	eg. Rubin Science Platform
"Second Party"		g. Butler	Rubin escope ontrol ystem ices
"Third Party"	eg. argocd	eg. 1Password Connect	ork but any s)

DEFIES CLASSIFICATION

What is Sasquatch

- Sasquatch is SQuaRE's "first-party" low-latency time series data service
- Captures, stores, replicates, UI
- Completely reliant on third-party Kafka + InfluxDB (but a time series DB swap cannot be ruled out)
- Sometimes referred to as the EFD (Engineering Facilities Database) but we renamed it to Sasquatch because:
  - Confusion with the system it replaced
  - The smarts is in the Kafka/API layer
  - It's not just for ~~breakfast~~ the EFD anymore
  - EFD is not what most people think it is anyway
  - We now use EFD to mean specifically the data holdings from the summit components

Must Read Docs

- sqr-068.lsst.io (architecture)
- sasquatch.lsst.io (docs)
- dmtn-290.lsst.io (high level summary)

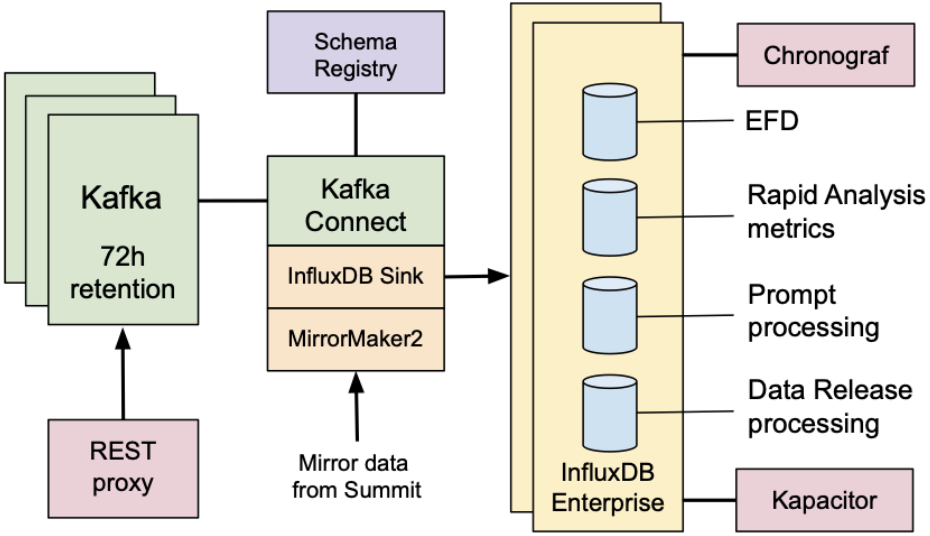

Figure 1. Sasquatch deployment at the US Data Facility.

Current Sasquatch holdings

- EFD (summit, USDF) - source: Summit Kafka brokers (SAL brokers, Camera vis SAL and own brokers), REST API (inc backpack)
- Pipeline Metrics (summit, USDF) - source: Butler put via REST proxy (also “telemetry gateway” functionality)
- Phalanx Metrics (Cloud) (Kafka)
- Could be more coming

Sasquatch on-prem - Things to Know 1/2

- Influx and Kafka both need fast (SSD) node-local storage
- Using taints/affinities to ensure dedicated nodes crucial for performance
- Weird shit can happen (maybe) if you pull the rug out of from under Influx (see compaction incident)
- Replication is “pull” with MirrorMaker2 (needs network *from* the replicator - eg USDF *to* the replicated (eg the summit) - so VPNs etc would be an issue.
- Highly prefer additional replication “taps” rather than accessing Kafka cross-cluster (do not create cluster inter-dependance)

Sasquatch on-prem - Things to Know 2/2

- EFD is massive and considered irreplaceable raw data
- The EFD “schema” is ridiculous but can’t get attention and downtime to fix it from the telescope folks
- People are always sending more (up to several 200Hz topics now), do not rely on sizing
- Influx Enterprise license effectively required for EFD (if included) for incremental backups if nothing else, also HA cluster setup and InfluxData support
- We are desperate to shift most users off Influx into ConsDB/Postgres aggregation
- Chronograf (Influx UI) is killing us, plan to switch to Grafana
- We will get rid of Segwardes when we have service discovery

Did I mention EFD is massive

Database growth during operations

Estimated increase of 30TB/year during survey operations

End of Year	Database size	RF	Total storage	Main drivers
2022	2TB	2		AuxTel
2023	15TB + SM	2	60 TB	M1M3, TMA, M2 + AuxTel
2024	**35 TB + SM**	**2**	**100 TB (*)**	**MainTel + AuxTel**
2025	65 TB + SM	2	160 TB (*)	Survey operations
2035	0.3-0.5PB + SM	2	0.7-1PB (*)	Survey operations

- Safety margin (SM) of 15TB + Replication factor (RF) of 2

However massive it is, people keep sending more stuff

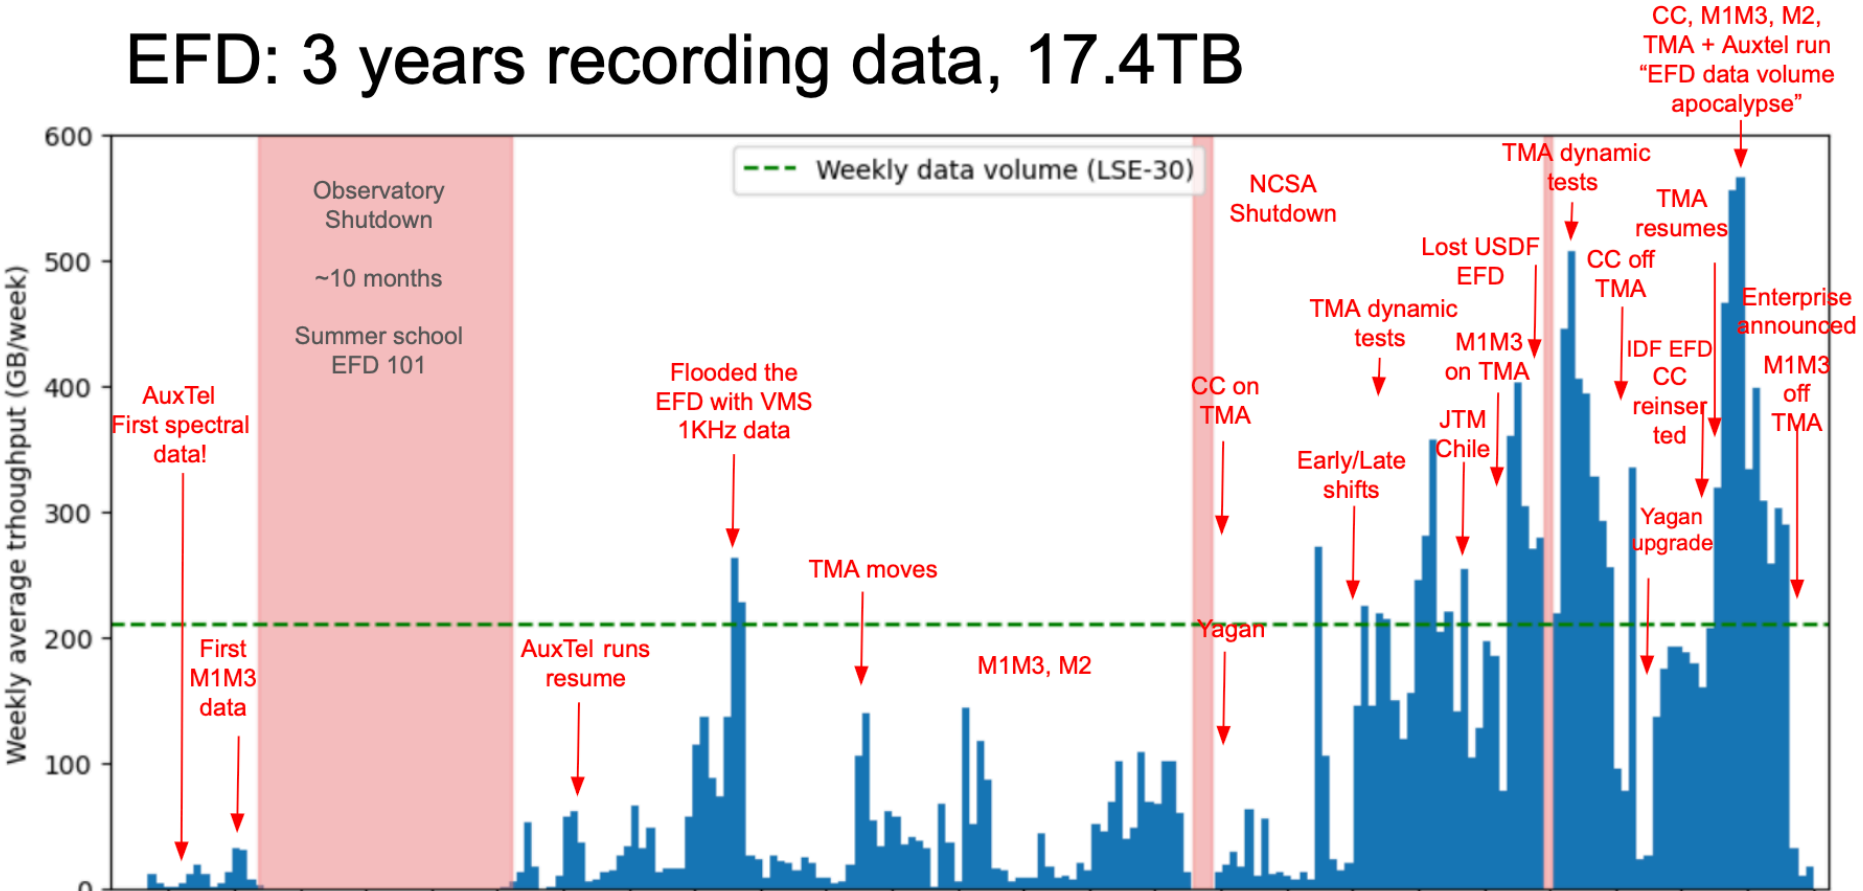

Should UKDF/FRDF run Sasquatch

- No
- Yes
- Maybe
- I don’t know

Should UKDF/FRDF run Sasquatch

- No (you should not have to care about the EFD)
- Yes (Sasquatch stores pipeline metrics)
- Maybe (You might be interested in the phalanx metrics)
- I don’t know (Are there plans to store eg. CM metrics or you can dispatch pipeline metrics to USDF directly)