


From 100k/s
To
2 million/s



Siqi Deng
OTG Monitoring Team
siqideng@lbl.gov
May 8, 2023

Operations Monitoring & Notification Infrastructure



- TSDB
 - Compatible with Prometheus
 - Stores all metrics
-
- Log aggregation system
 - Prometheus ecosystem
 - Stores all logs



Phase 1:

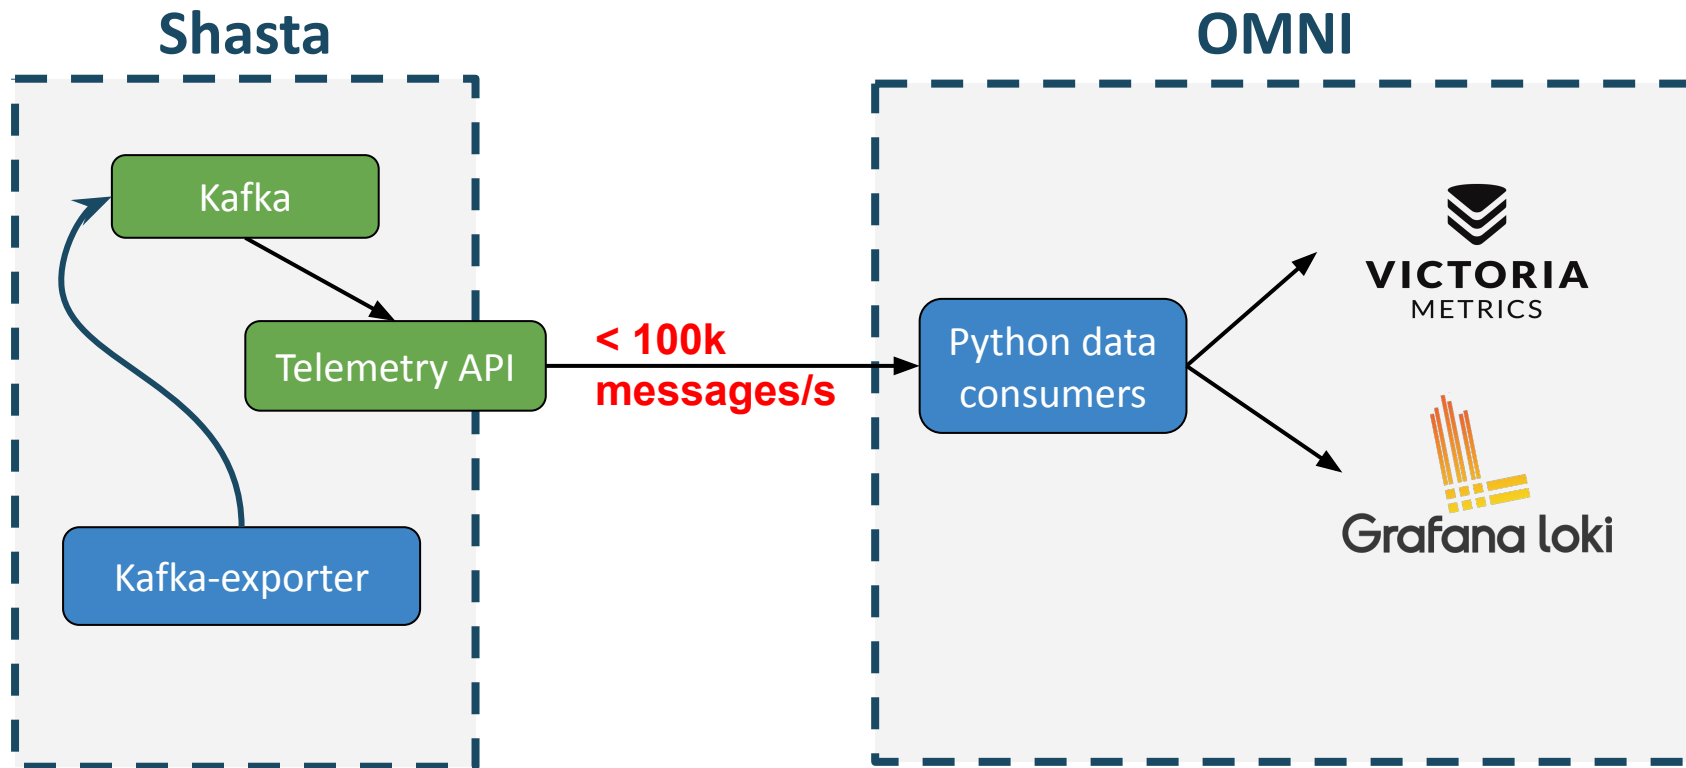
1536 compute nodes with
6,000+ NVIDIA A100 GPU
100+ non-compute nodes

Phase 2:

3,000+ CPU-only compute
nodes



Data Pipeline(2021~Early 2023)



Kafka-exporter(Non-JMX exporter)



- Data production rate
- Data consumption rate
- Lags

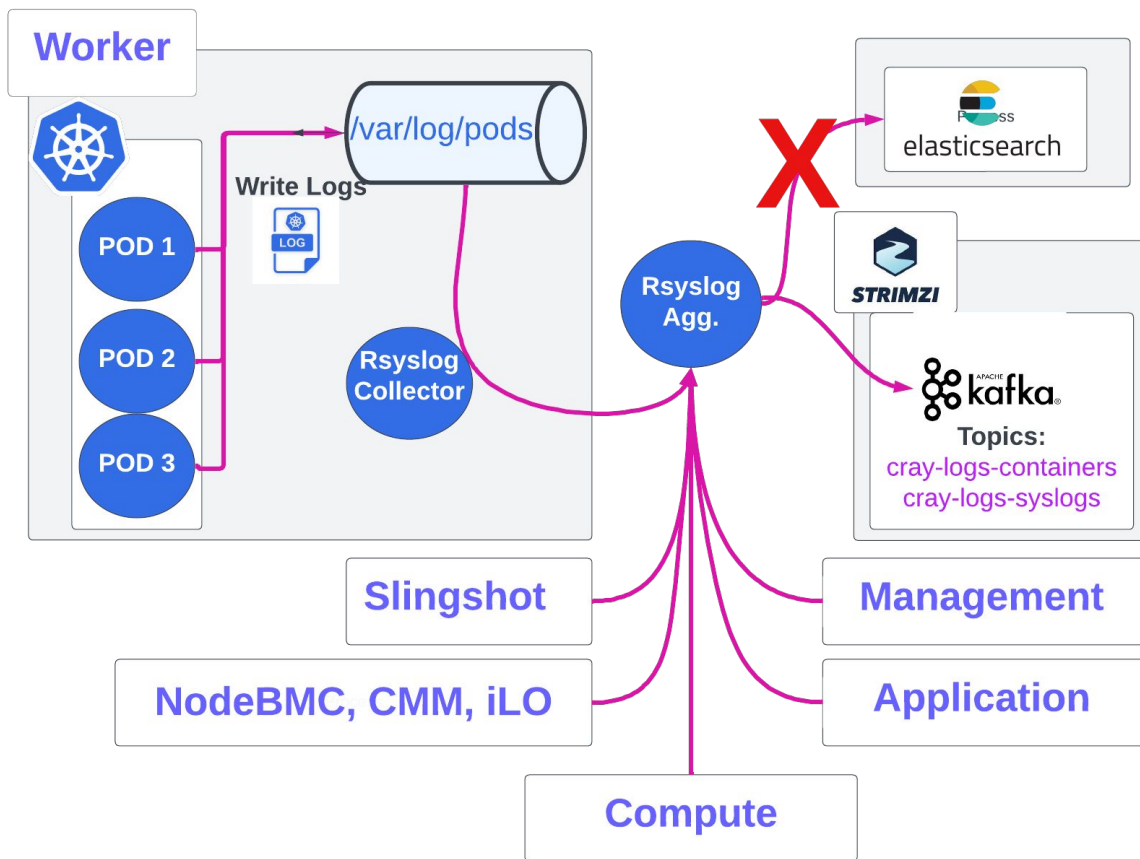
Data Production Rate by Topics



An Inverse Relationship



Bottlenecks

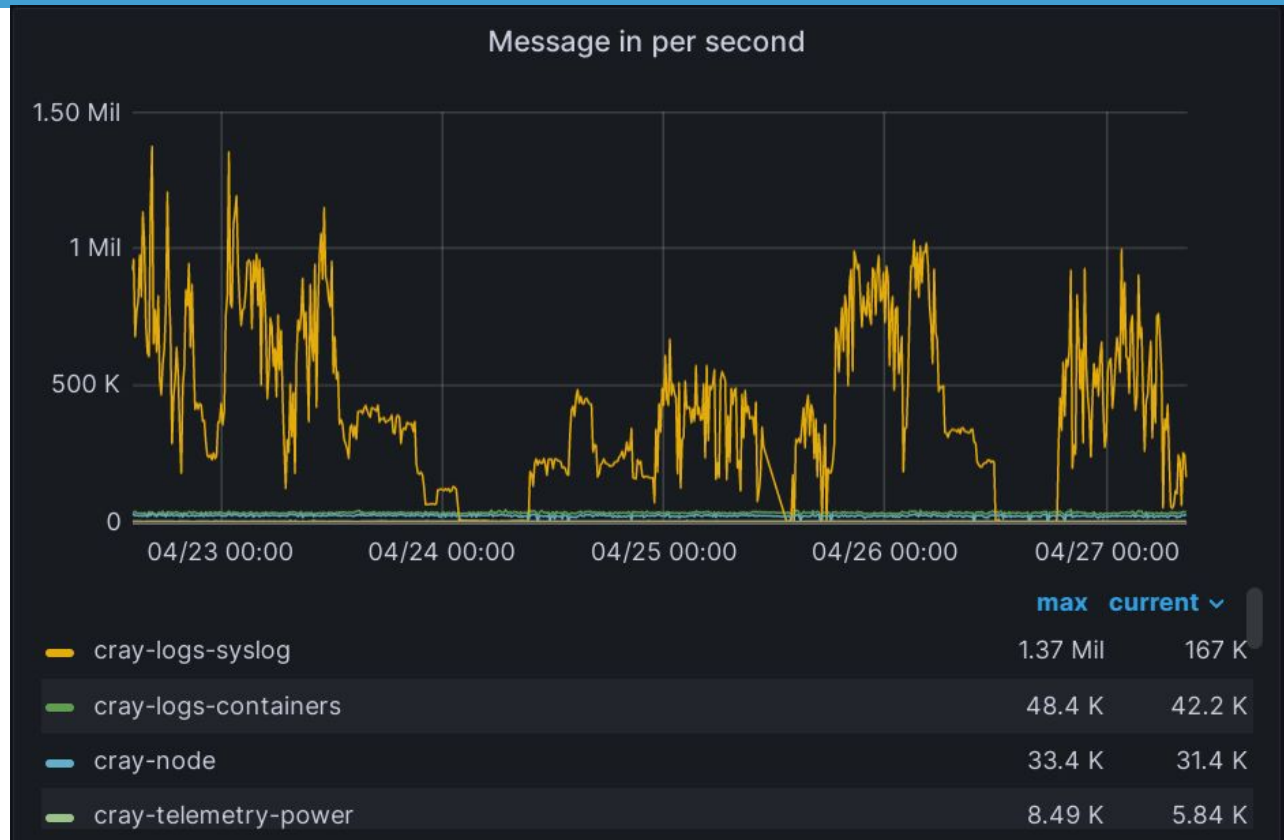


After Scaling Up Rsyslog & Kafka

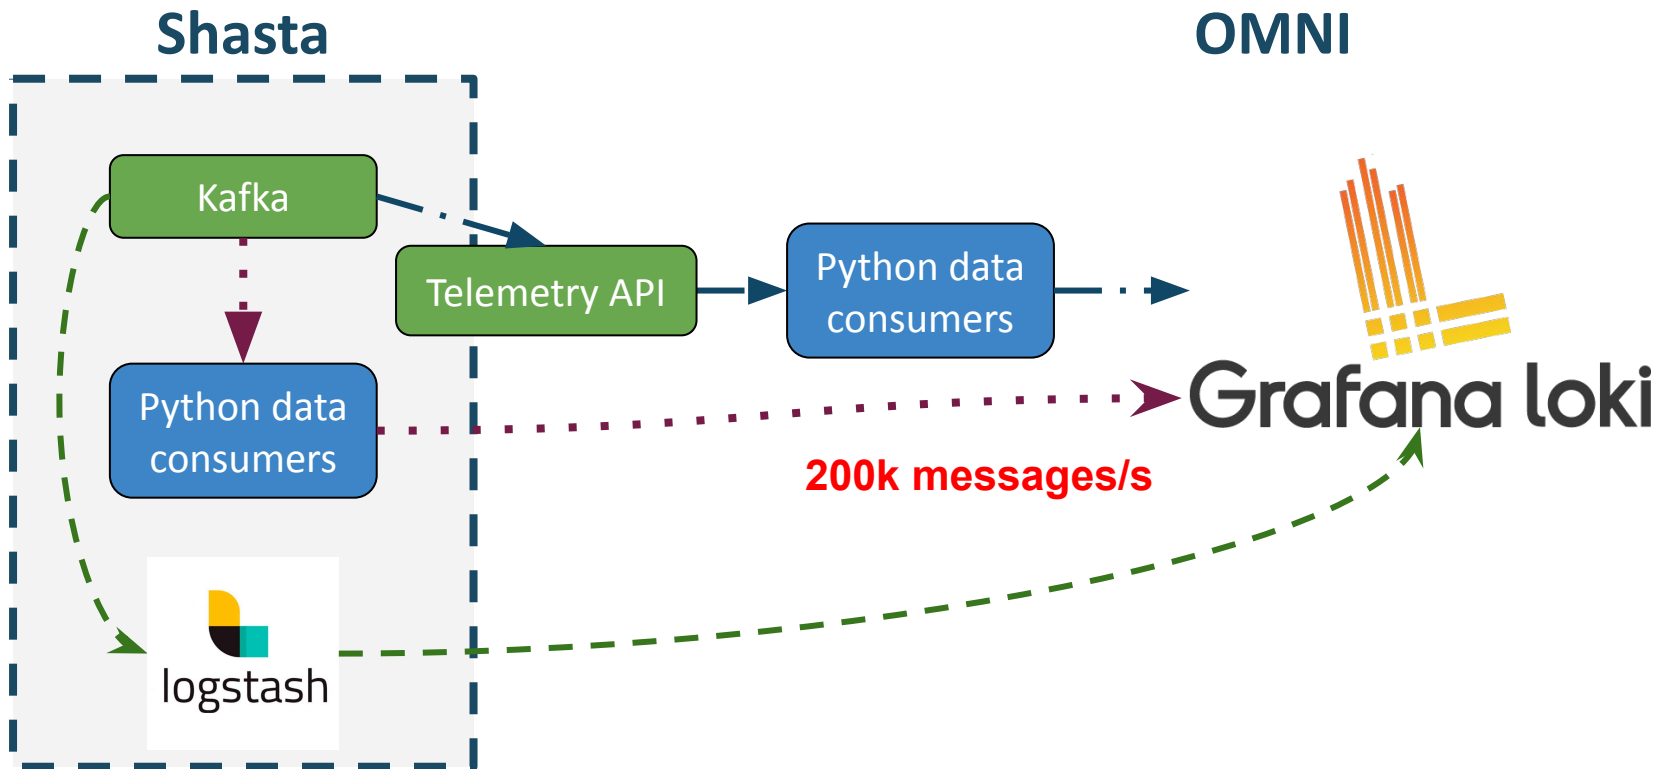


Over 1 million messages per second

Syslog: 90%+



Solutions to send logs to Loki



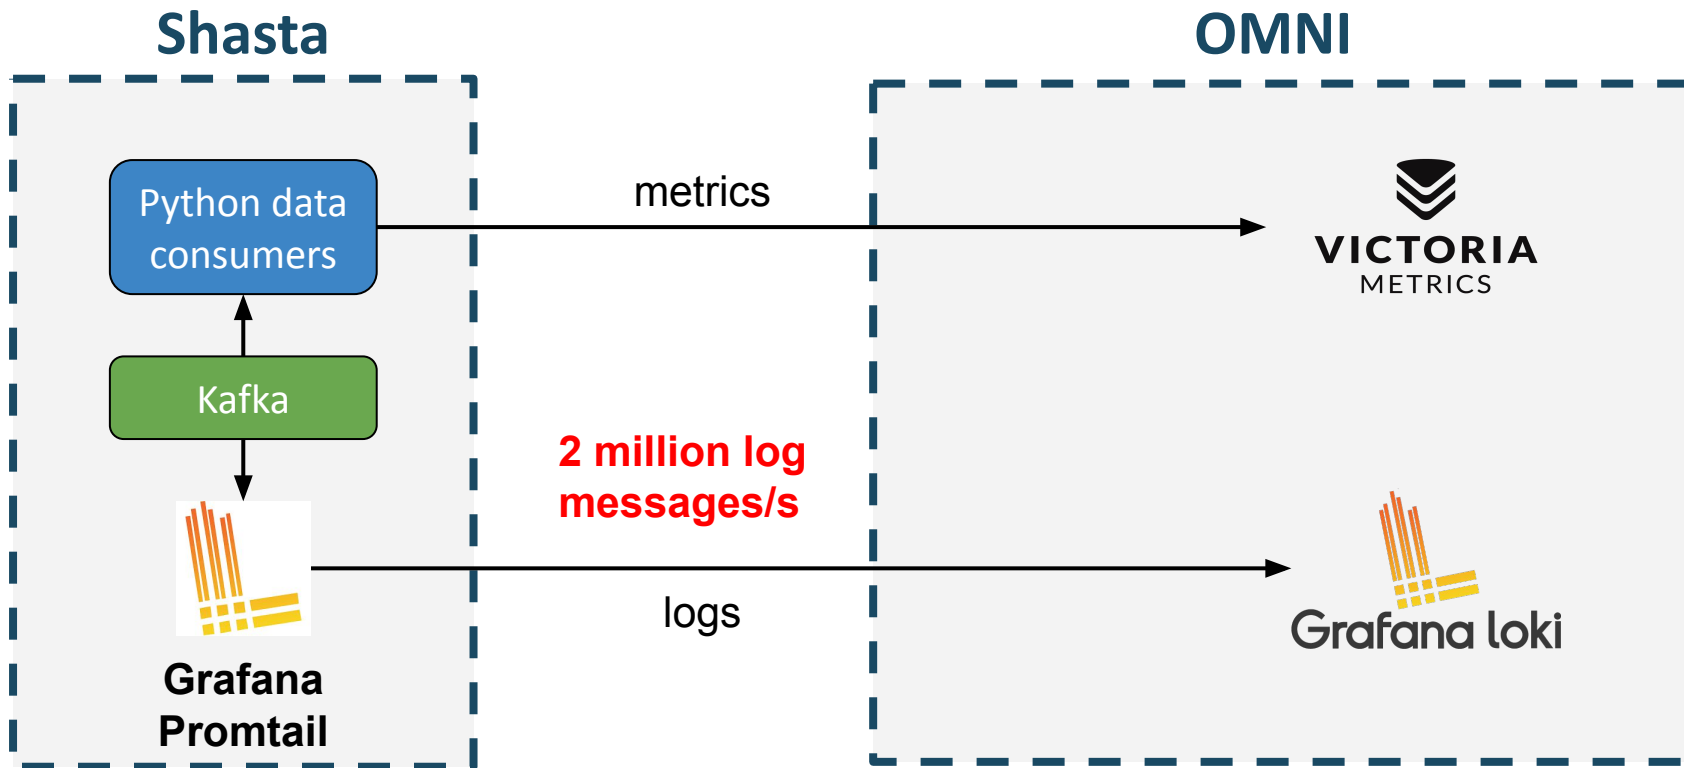
Data Pipeline: 2 million logs/s



2 CPU cores +
150MB Mem
per pod x 64~128
pods



Data Pipeline: 100k/s >> 2m/s



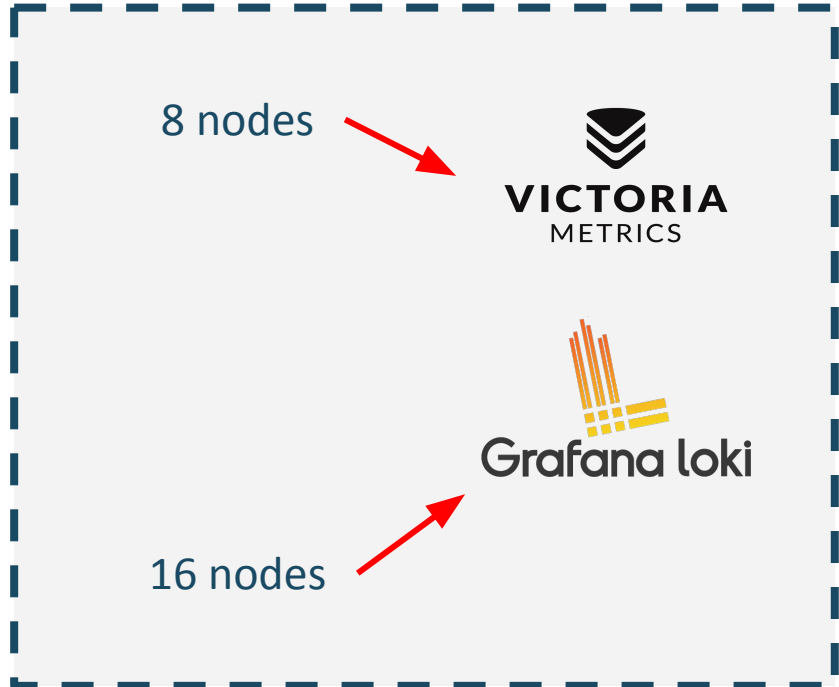
Hardware Upgrade:

- 8-node >> 8-node + 16-node
- 1 gateway >> 4 gateways
- Upcoming: higher bandwidth & more memory

Software Upgrade:

- Better scalability
- Better data compression ratio
- Better search performance

OMNI

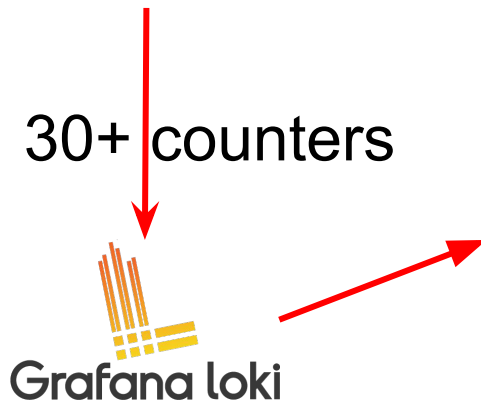


New Use Case: Cassini Monitoring



2023-03-30 12:12:01 nersc-nodeprolog|6324876|Telemetry:cxii2:

```
pct_mst_hit_on_som:119,pct_sct_timeouts:1,pct_spt_timeouts:10,pct_tct_timeouts:0,pct_bad_seq_nacks:1494,pct_no_mst_nacks:0,pct_no_tct_nacks:0,pct_no_trs_nacks:0,pct_trs_rsp_nack_drops:0,pct_trs_replay_pend_drops:0,hni_rx_paused_0:545743258,hni_rx_paused_1:317616923,hni_rx_paused_2:0,hni_rx_paused_3:30545,hni_rx_paused_4:0,hni_rx_paused_5:882208505,hni_rx_paused_6:0,hni_rx_paused_7:0,hni_rx_paused_std:1128451822,hni_tx_paused_0:8055899,hni_tx_paused_1:33258947,hni_tx_paused_2:34996804,hni_tx_paused_3:33258947,hni_tx_paused_4:6527124,hni_tx_paused_5:0,hni_tx_paused_6:6527124,hni_tx_paused_7:33258947,parbs_tarb_pi_non_posted_blocked_cnt:418236214,parbs_tarb_pi_non_posted_pkts:2197636770,parbs_tarb_pi_posted_blocked_cnt:4960363957,parbs_tarb_pi_posted_pkts:2537599855
```

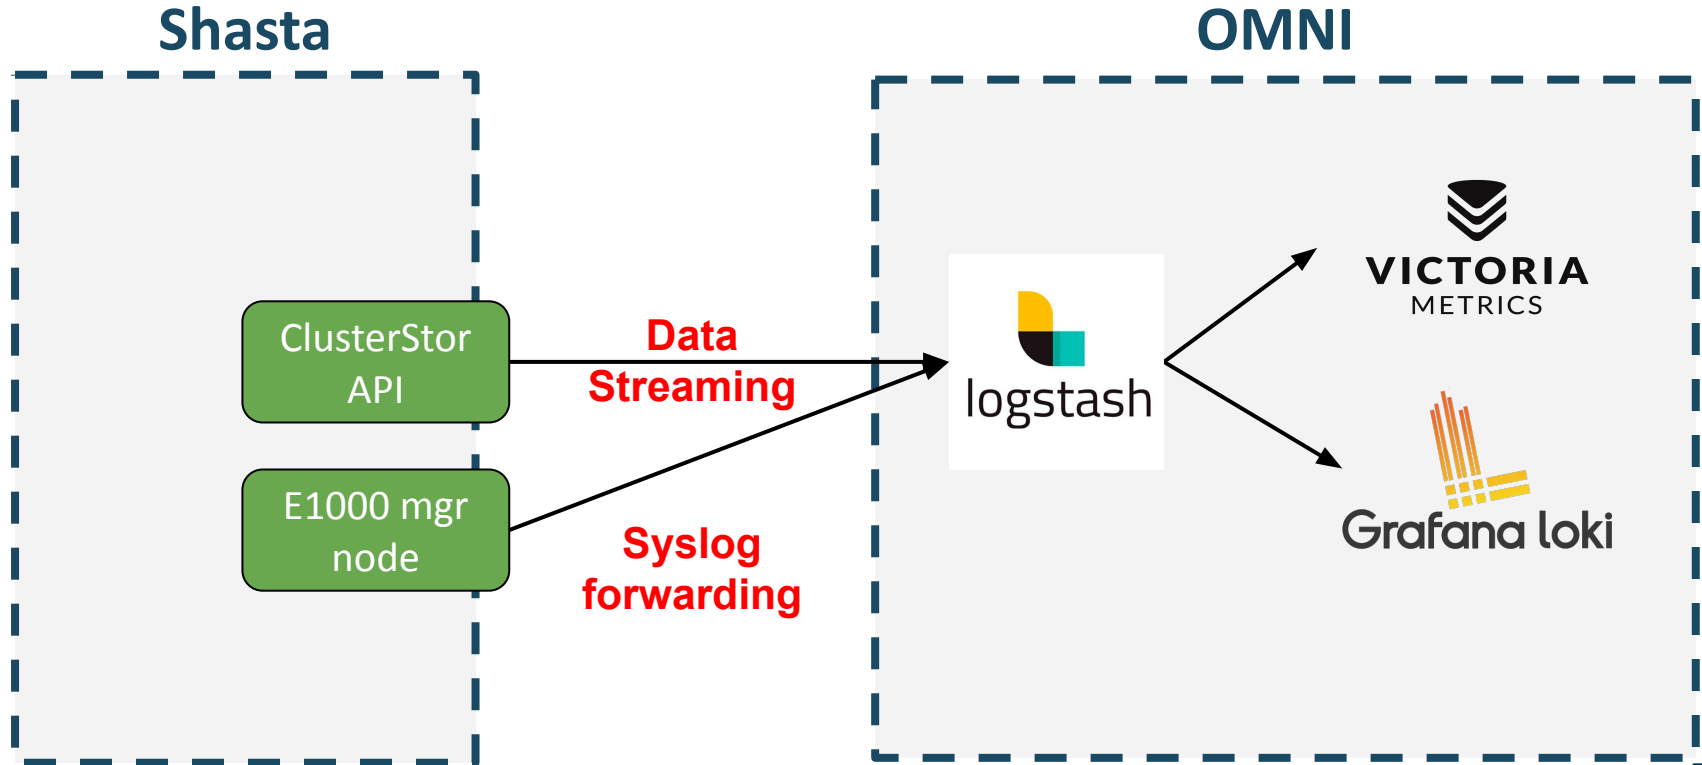


Lessons Learned

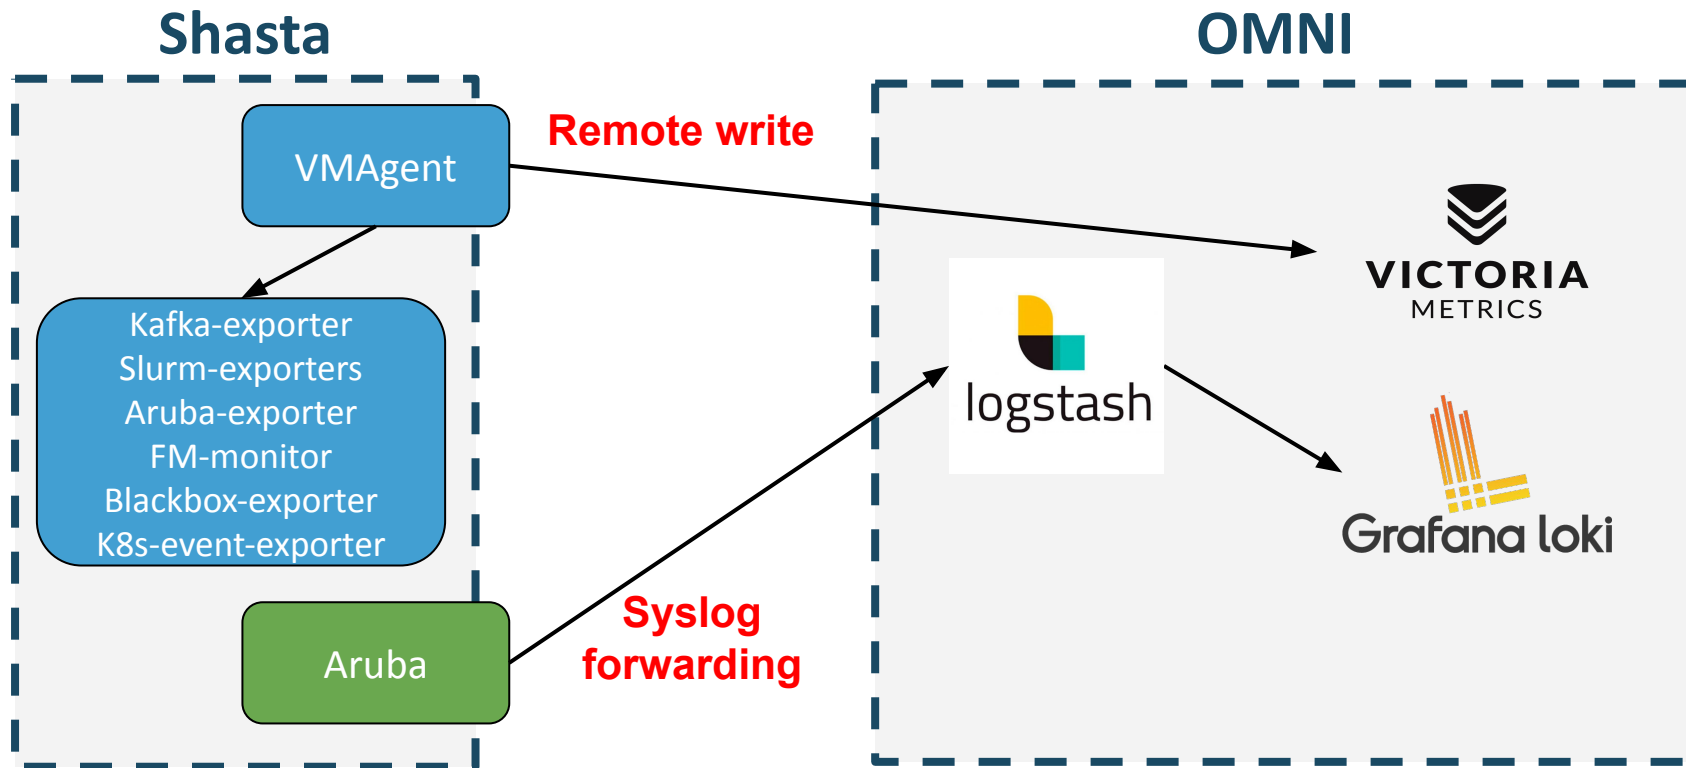



- the default configuration is not suitable for everyone
- additional monitoring is required

ClusterStor Monitoring Data Pipeline



Other Monitoring Data Pipelines





Questions?