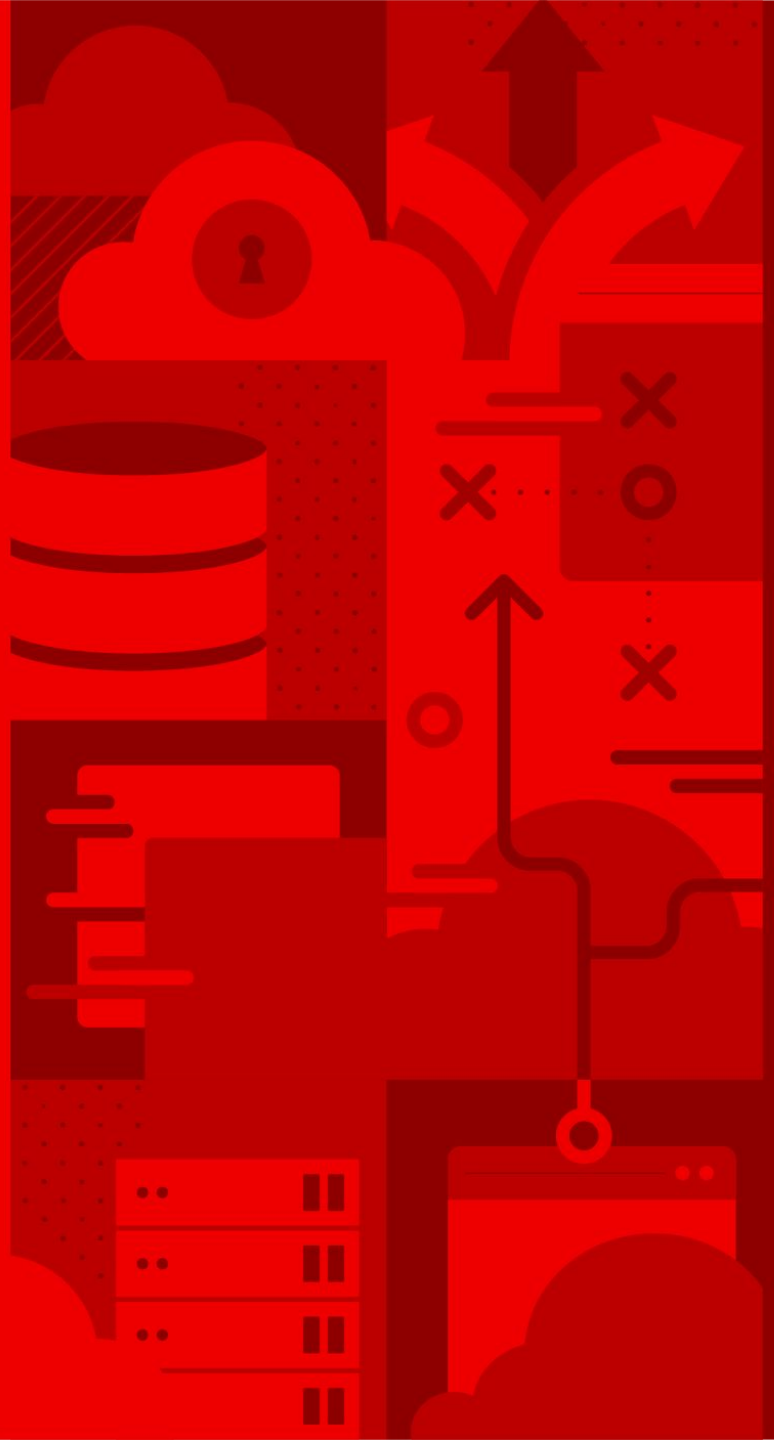

Automating RCA in OVS/OVN-Based Deployments Using AI/ML

Gurpreet Singh

Principal Product Manager (Red Hat)

Point of observations (topology view)

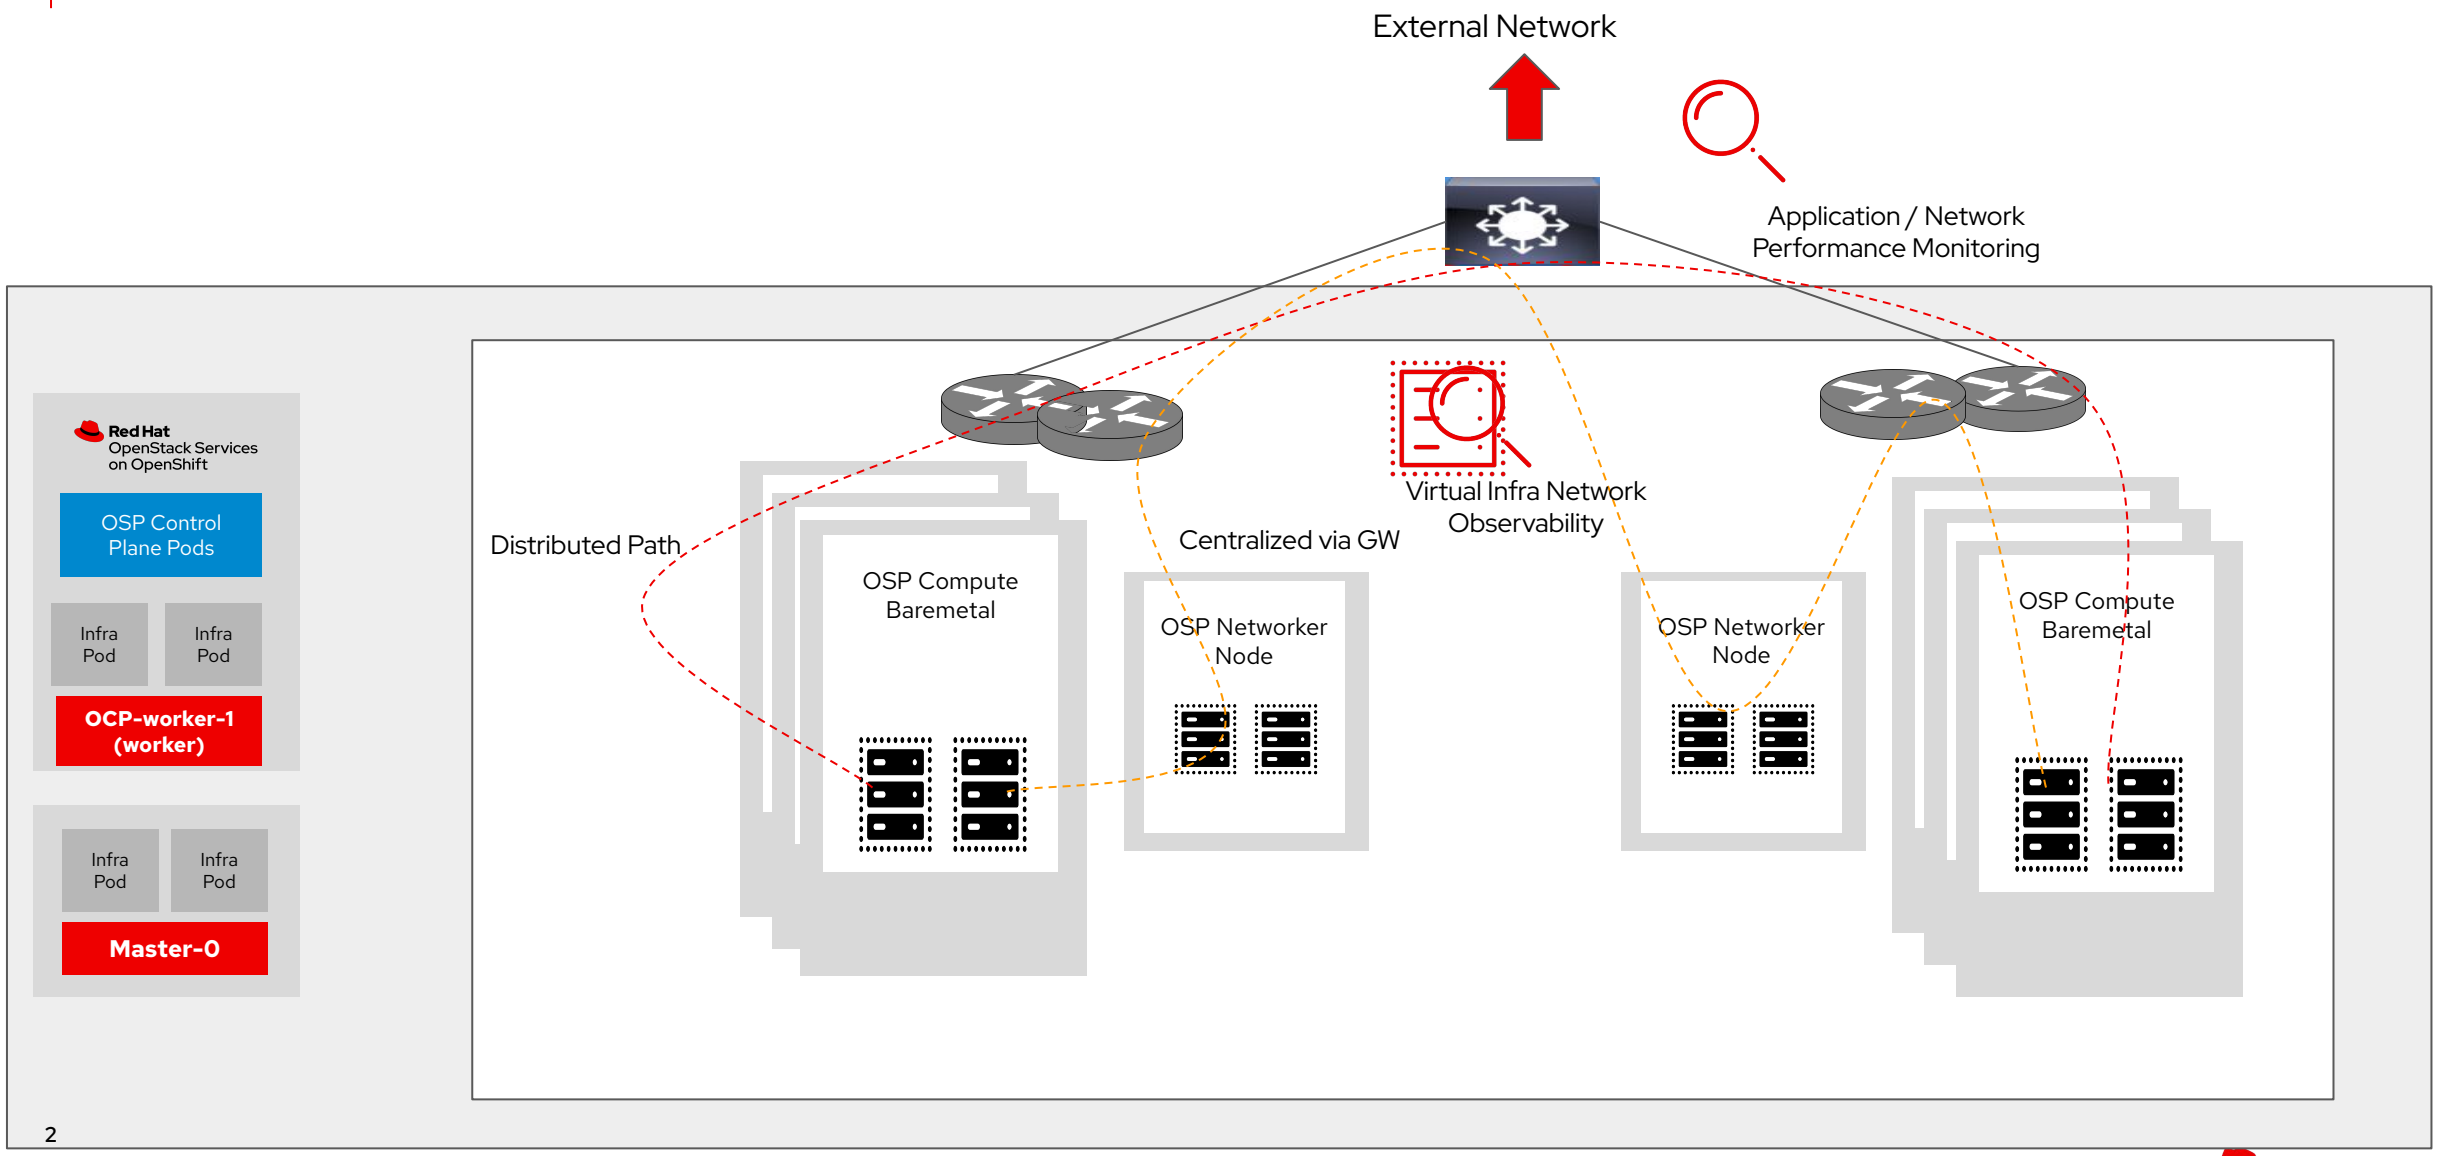

Network Observability Metrics

Metrics / KPIs

- Control plane
  - OVN DB
- Data Plane
  - OVS, OVN (interface/port level metrics)
  - OVS-DPDK
  - System metrics
  - Resource Utilization

Root Cause / Point of Contention

- NIC, OVS/OVN
- Buffer oversubscription, queue lengths
- CPU & Memory allocation

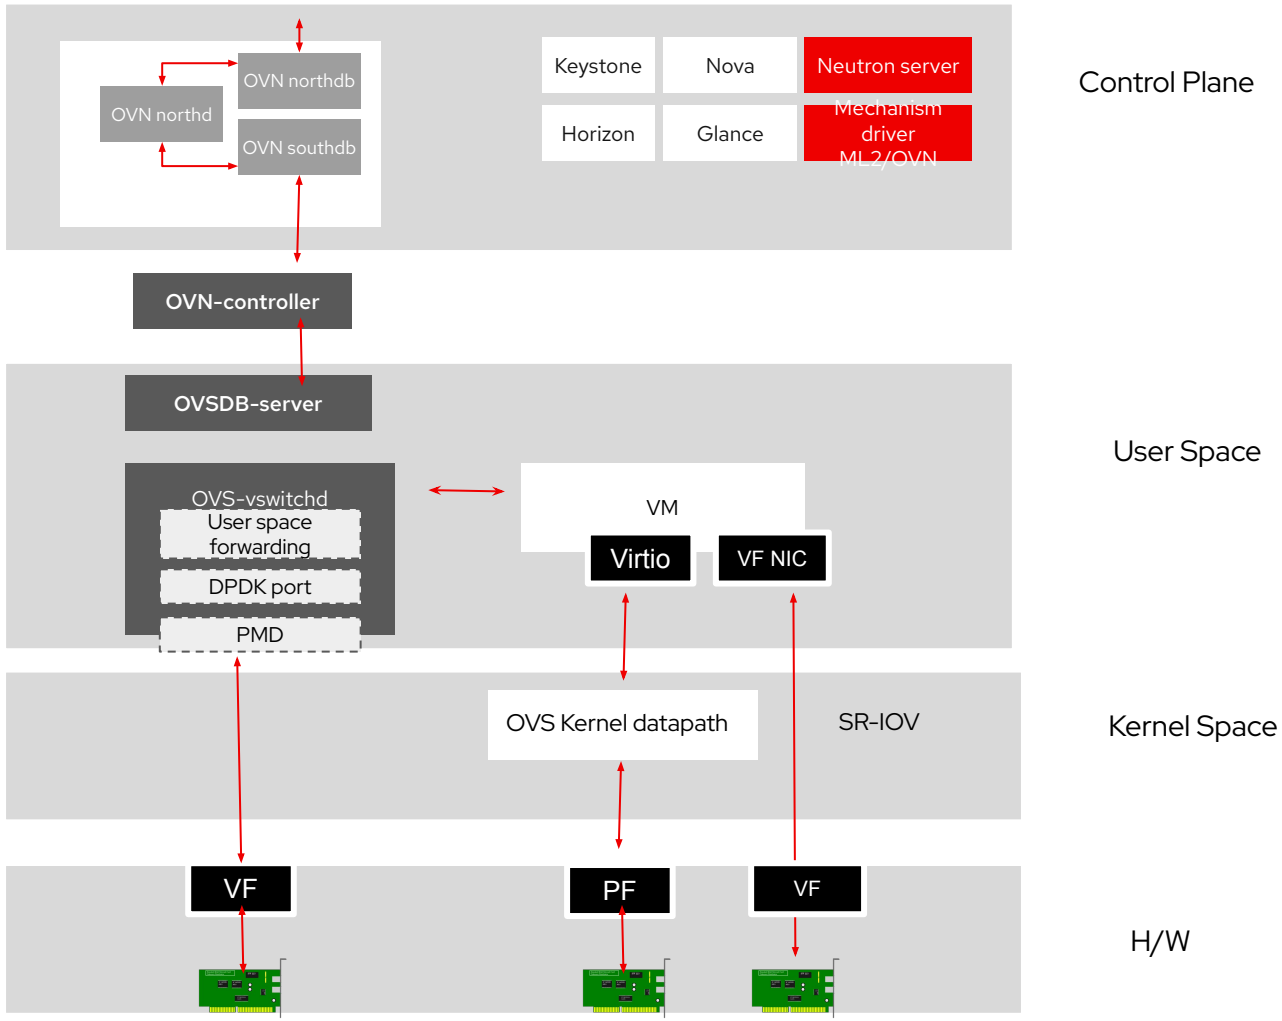

Fine Tuning Performance

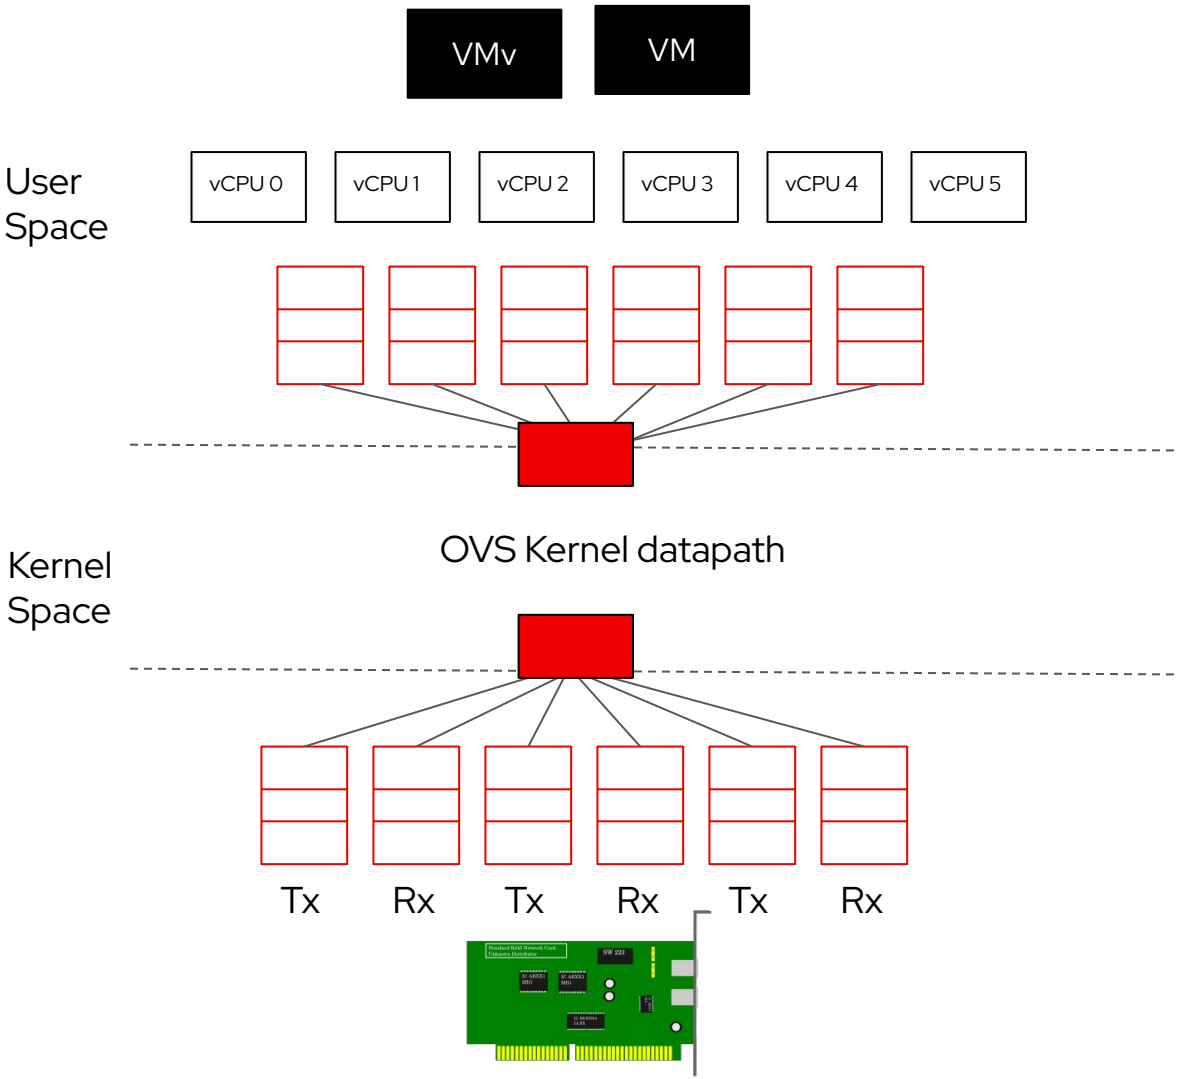

- vCPU core, memory allocation
  - Resource Optimization
  - Performance isolation
    - CPU Partitioning
    - Isolcpus
- Non-voluntary cpu context switches on OVS PMD and QEMU vCPU threads
- IRQ balance
- vCPU to queue mapping
  - Hashing or 1:1
- Tx/Rx queue length
- Throughput / Latency trade-off
- Acceleration technologies

Traditional Root Cause Analysis

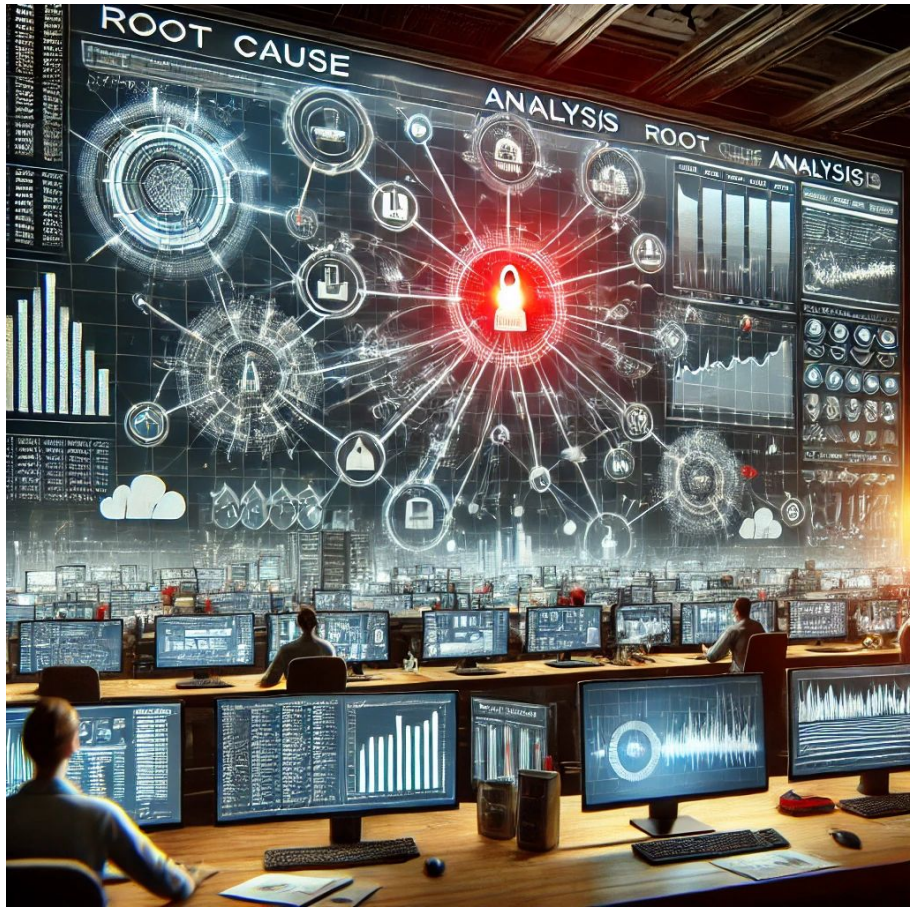

- Largely manual with reliance on monitoring dashboards
- Multiple entities and iterations
- Sometimes can take weeks or months to resolve
- Reactive approach impacting SLAs and service quality

Overprovision capacity significantly

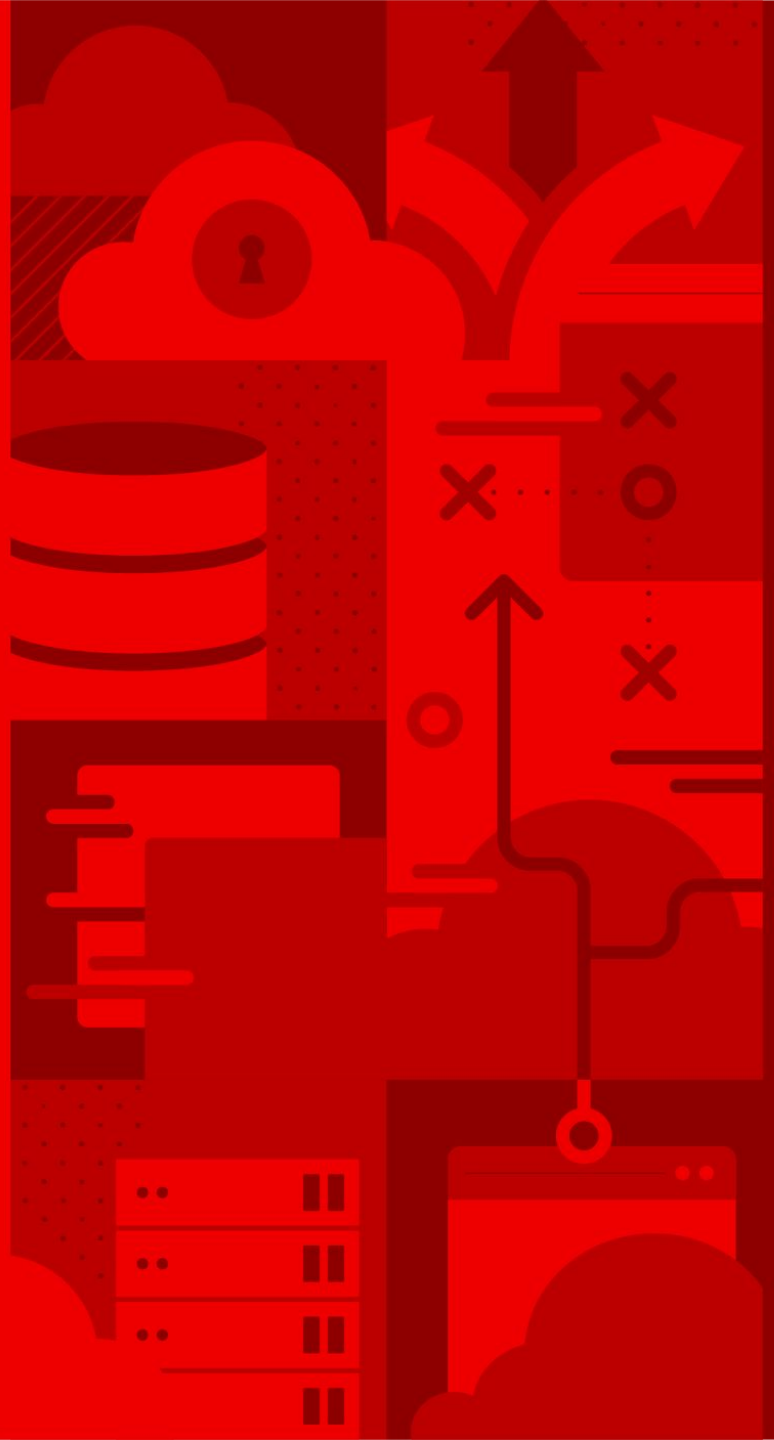

Anomaly Detection

Red Hat

Analytics

Anomaly detection

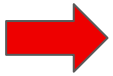

Rule-based

- Established patterns, signatures and thresholds

- Complex and difficult to scale
- Lack of novel anomalies

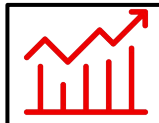

Statistical

- Statistical methods to measure deviation from normal

- Diversity of workload, traffic and usage patterns
- Long convergence times

Machine Learning

- Automation
- Correlation across metric sets
- Proactive vs Reactive
- Scalability

- Insufficient market adoption
- Interpretability

Self Organizing Maps (SOM) based approach

- Joint analysis of infrastructure resource utilization (infra metrics) & application performance metrics (VNF metrics)
- Real data from DC
- Somoclu open source SOM implementation
- 99% accuracy

Challenges

- High processing time, large data sets
- Scale challenges

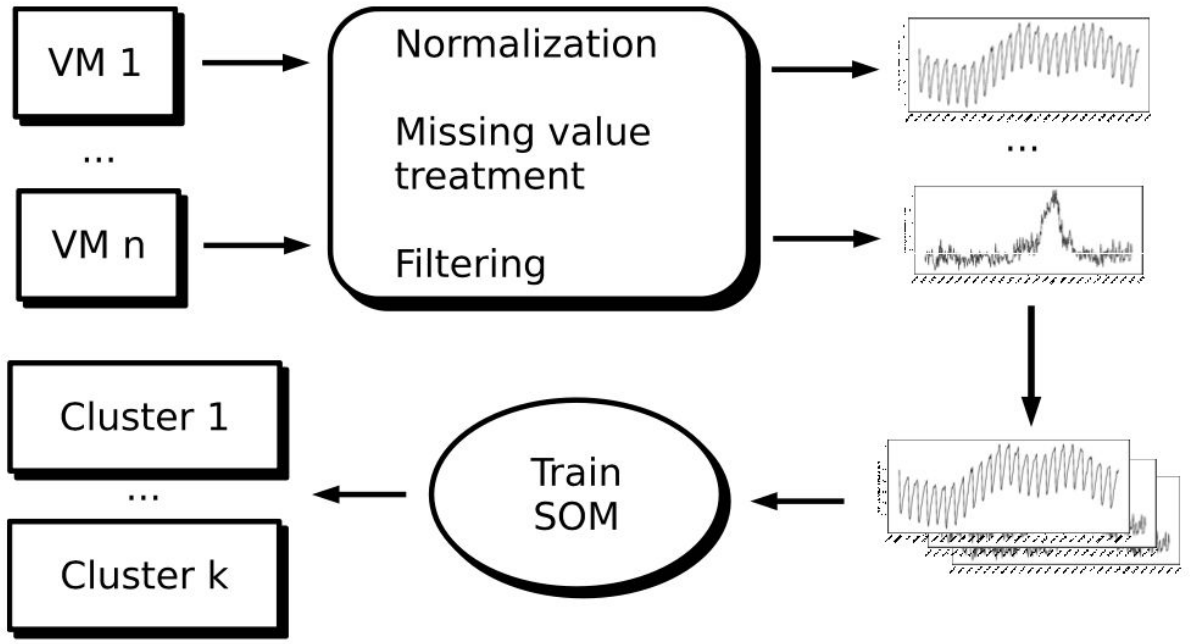

Reference: Giacomo Lanciano, Antonio Ritacco, Tommaso Cucinotta, Marco Vannucci, Antonino Artale, Luca Basili, Enrica Sposato, and Joao Barata. 2020. SOM-based behavioral analysis for virtualized network functions. In Proceedings of the 35th Annual ACM Symposium on Applied Computing (SAC '20). Association for Computing Machinery, New York, NY, USA, 1204–1206. https://doi.org/10.1145/3341105.3374110

Model Training

Simulated environment

- Simulate workload to mimic real-world scenarios
- Replicate resource allocation and configuration

Fault Injection

- CPU stress
- Network stress (delay)
- Memory stress
- Tx/Rx queue oversubscription
- Limit to service degradation and not service outage

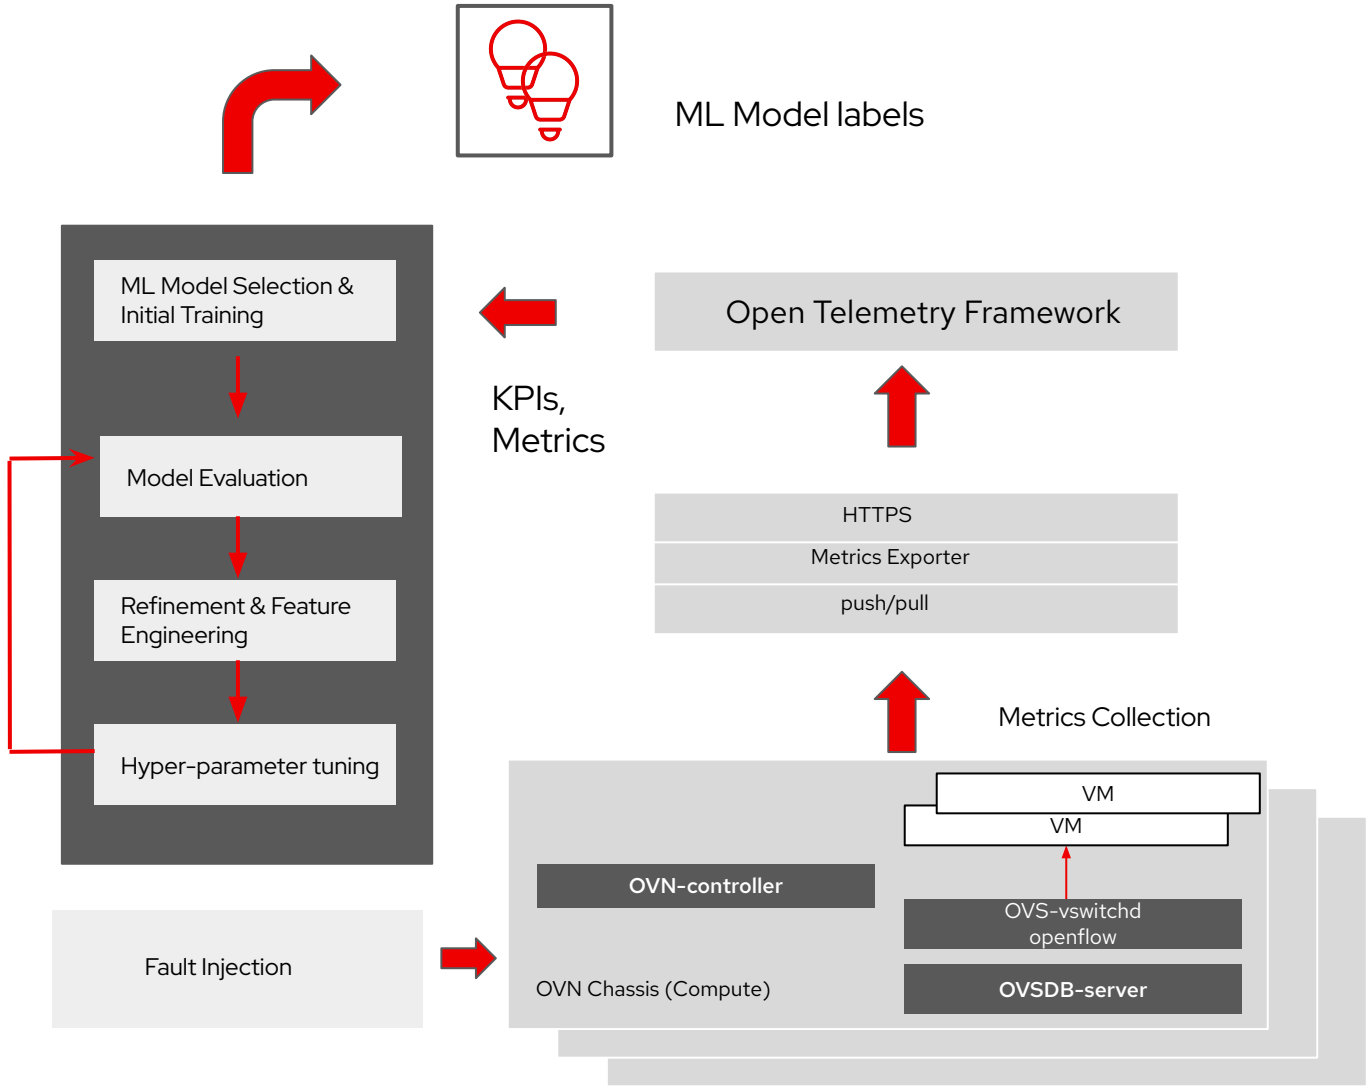

OSP Sandbox for pre-training ML models

Closed Loop Automation

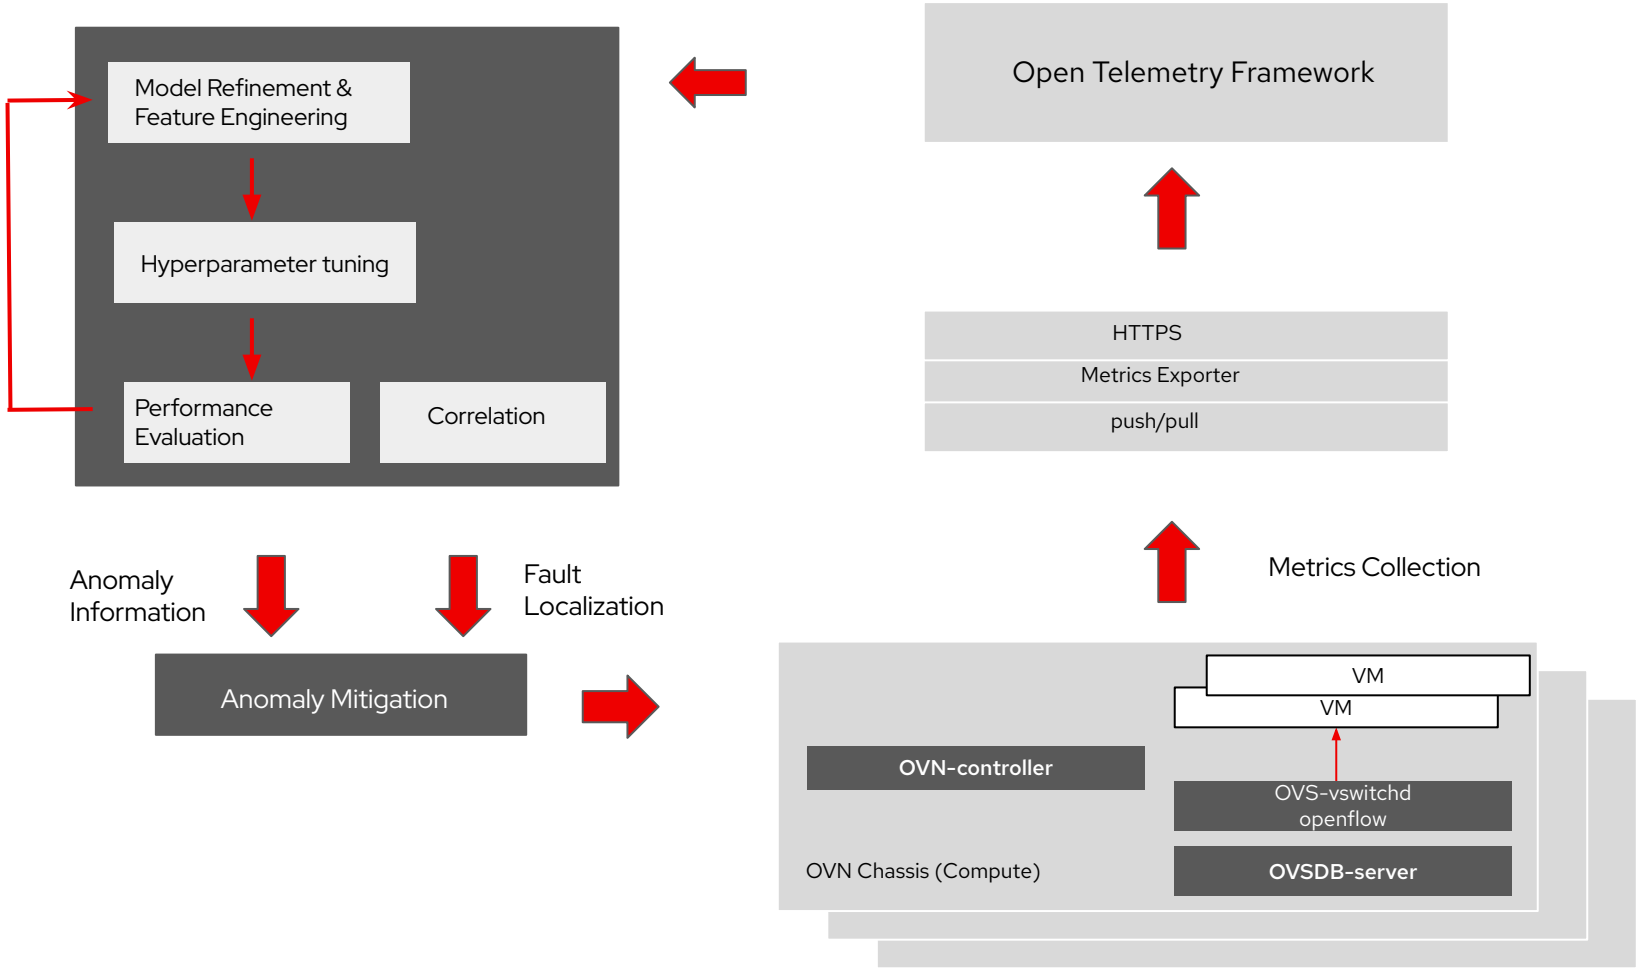

Future Work

- Standardized observability framework such as OpenTelemetry
- Open-source tools, sandbox that enables fault injection in simulated environments for training AI/ML models for fault/anomaly detection
- Open-source test VNFs to enable VNF centric telemetry data collection and test anomaly mitigation

Useful Links & References

[1] Grout: DPDK based Graph router

[2] Somoclu: Open Source SOM Implementation

[3] Open Telemetry

[4] Zehra S, Faseeha U, Syed HJ, Samad F, Ibrahim AO, Abulfaraj AW, Nagmeldin W. Machine Learning-Based Anomaly Detection in NFV: A Comprehensive Survey. Sensors (Basel). 2023 Jun 5;23(11):5340. doi: 10.3390/s23115340. PMID: 37300067; PMCID: PMC10256098.

Thank you

Red Hat is the world's leading provider of enterprise open source software solutions. Award-winning support, training, and consulting services make Red Hat a trusted adviser to the Fortune 500.

linkedin.com/company/red-hat

youtube.com/user/RedHatVideos

facebook.com/redhatinc

twitter.com/RedHat

Red Hat