SUPERFLUIDITY

Demand-driven orchestration for 5G deployments

BEFORE: Administrators used to manage workloads

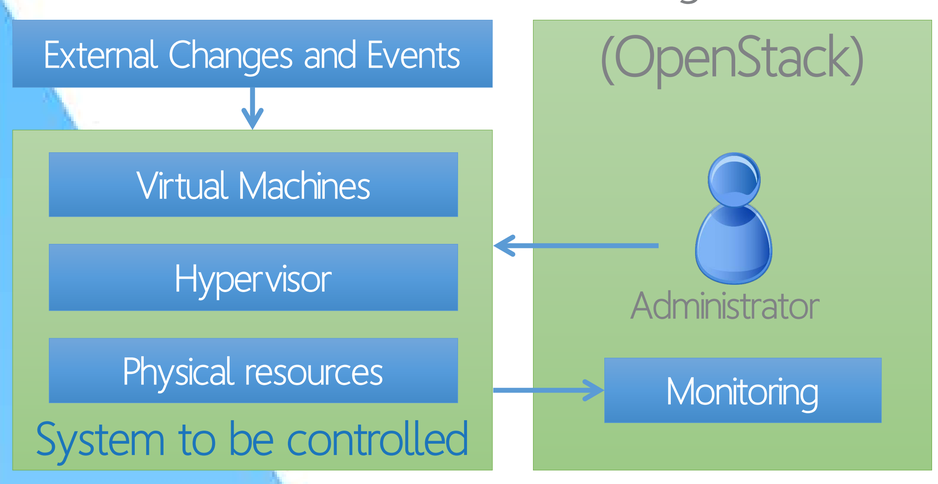

CHALLENGES

- Meeting low-latency and high-bandwidth, always on connections
  - Currently have long provision times and waste resources
  - Video-streaming services are pre-provisioned
- Number of devices in 5G landscape is growing exponentially
  - Expect 1000x volume of traffic of 3G/4G
- Ensure that KPIs / SLAs are monitored and enforced
- Performance of orchestration framework must meet demands

INNOVATIONS

- On-demand creation of resources to react to telemetry
- Automated, transparent service delivery mechanisms
- Model based approach for determining resource placement
- Offline low-level analysis of KPI metrics for optimisations
- Scale-independence, allowing addition of resources anywhere

Architected to work on heterogeneous hardware platforms.

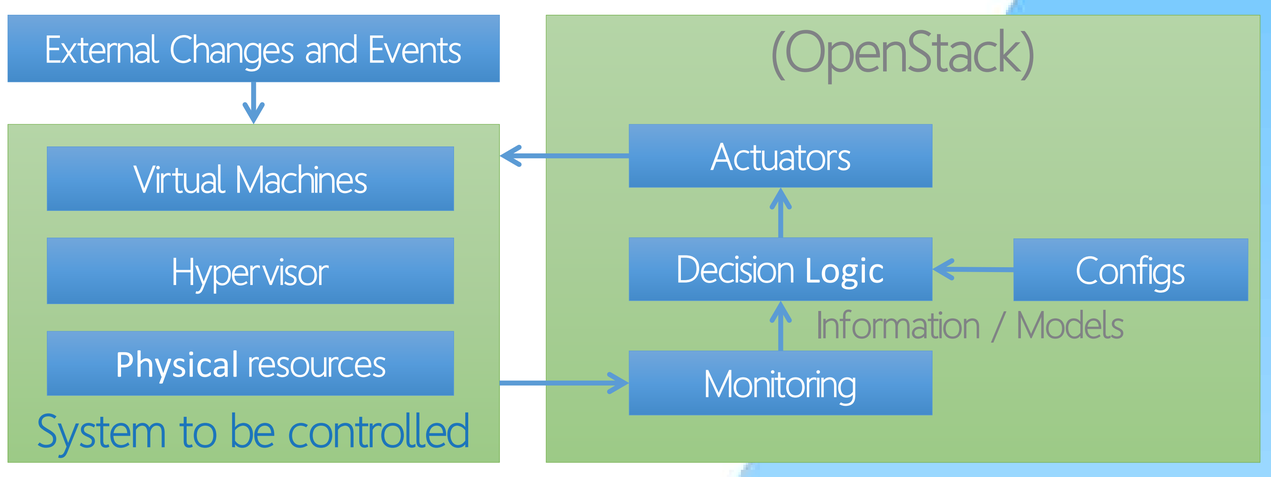

AFTER: 5G deployments, need (semi-)automated orchestration frameworks

SCENARIOS

a) KPI-driven Scaling in the Video Streaming Cloud

- Determination of Server Side Key Performance indicators based on Analytics and client side KPI's in offline experiments
- Monitoring of Key Performance Indicators – in demo, Ceilometer monitors loss rate and network throughput
- Critical values for KPIs trigger scaling
- Scaling of the virtual infrastructure – in demo, Nokia Cloudband's Mistral plus load balancer
- Technology improves Video Streaming and other applications with varying resource demands in the cloud

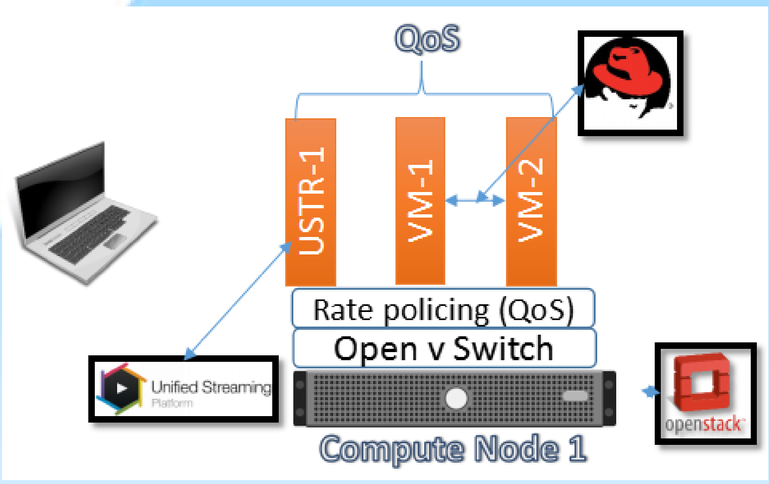

b) Noisy Neighbour mitigation in the Video Streaming Cloud

- VMs on same /nearby physical infrastructure can compete for network usage: "noisy neighbours"
- Rate Policing on top of Open vSwitch mitigates the problem & improves quality of video streaming
- Implemented and contributed to OpenStack

BENEFITS

- Automated mapping of service KPIs to platform metrics of most significant influence using Telemetry and Statistical modelling
- Workload Characterisation using TALE (Throughput/Anomalies/Latency/Entropy) Methodology Real-Time Telemetry to monitor a service

- Application driven on-the-go resource scaling
- Cost reduction and quality improvement for video streaming
- Noisy Neighbour interference mitigation by QoS Policing
- Useful for Video Streaming and other bandwidth savy applications
- Starting point for the superfluidity architecture that features on-the-go scaling of compute and network resources