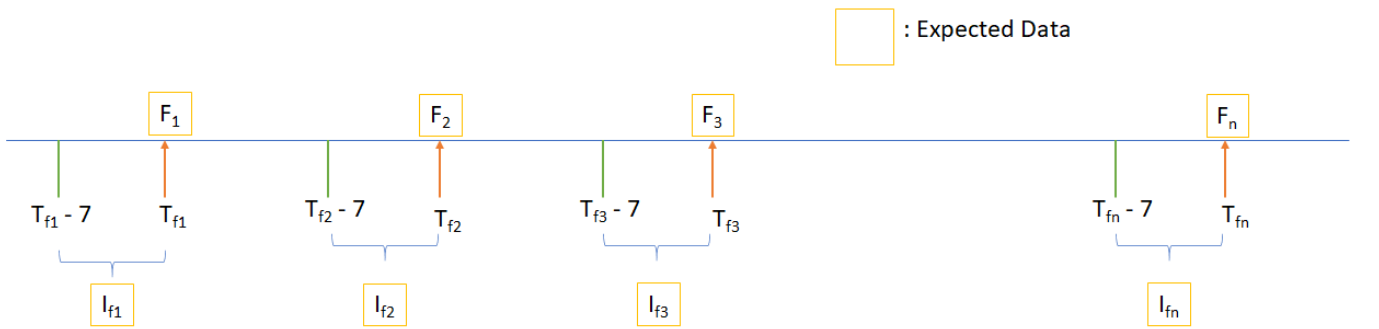


Failure Prediction using AI/ML in NFV Environments

Data



Failure Type	Failure parameter	<div><div><div>F_n</div></div><div>Failure Event</div></div>	<div><div><div>I_f</div></div><div>Infrastructure Metrics</div></div>	Comments								
Links	Link Down. Link removed	<div>Virtual Switch link failure</div> <div>Hardware Failure</div> <div>Interface Down</div> <div><table><tr><th>dhcp-agent.log</th><th>neutron-dhcp-agent</th></tr><tr><td>l3-agent.log</td><td>neutron-l3-agent</td></tr><tr><td>linuxbridge-agent.log</td><td>neutron-linuxbridge-agent</td></tr><tr><td>openvswitch-agent.log</td><td>neutron-openvswitch-agent</td></tr></table></div> <div>(Ref: https://docs.openstack.org/ocata/config-reference/networking/logs.html)</div>	dhcp-agent.log	neutron-dhcp-agent	l3-agent.log	neutron-l3-agent	linuxbridge-agent.log	neutron-linuxbridge-agent	openvswitch-agent.log	neutron-openvswitch-agent	<div>Network interface status,</div> <div>High packet drop,</div> <div>low throughput,</div> <div>excessive latency or jitter</div> <div>crc-statistics, fabric-link-failure, link-flap, transceiver-power-low</div>	
dhcp-agent.log	neutron-dhcp-agent											
l3-agent.log	neutron-l3-agent											
linuxbridge-agent.log	neutron-linuxbridge-agent											
openvswitch-agent.log	neutron-openvswitch-agent											
VM	<div>Deployment/Start Failures:</div> <div>1. Failed to start*</div> <div>2. Failed to boot*</div> <div>Post-Deployment /Start failures:</div> <div>1. Shutdown</div> <div>2. Crash</div> <div>3. Hang*</div> <div>4. Panic</div>	<div>nova-compute.log</div> <div>nova-api.log</div> <div>nova-scheduler.log</div> <div>libvirt.log</div> <div>qemu/\$vm.log</div> <div>neutron-server.log</div> <div>glance/cinder -</div> <div>flavor</div> <div>Node and Core-mapping</div>	<div>cpu: per-core utilization</div> <div>memory</div> <div>Interfaces statistics - sent, recv, drops</div> <div>Disk Read/Write</div>	<div>If possible, Infrastructure metrics and syslogs from within the VM should be collected.</div> <div>Deployment/Start failures can be the first step.</div>								
Container	<div>Deployment/Start Failures:</div> <div>1. Failed to start*</div> <div>2. Failed to boot*</div> <div>Post-Deployment /Start failures:</div> <div>1. Shutdown</div> <div>2. Crash</div> <div>3. Hang</div> <div>4. Panic</div>	<div><ul style="list-style-type: none">OS layer – syslog, boot.log, kern.log etc.Kubernetes Layer – container Logs (/var/log/containers)OpenStack Layer – OpenStack service Logs</div>	<div>cpu: per-core utilization</div> <div>memory</div> <div>Interfaces statistics - sent, recv, drops</div> <div>Disk Read/Write</div>									

Node	<p>A node failure (hardware failure, OS crash, etc)</p> <p>A) node network connectivity failure</p> <p>B) nova service failure</p> <p>C) Failure of other OpenStack services</p>	<p><code>/var/log/nova/nova-compute.log</code> (To ensure that it has successfully connected to the AMQP server Ref: https://docs.openstack.org/operations-guide/ops-maintenance-compute.html)</p> <table><thead><tr><th>Cloud controller</th><th>nova-*</th><th><code>/var/log/nova</code></th></tr></thead><tbody><tr><td>Cloud controller</td><td>glance-*</td><td><code>/var/log/glance</code></td></tr><tr><td>Cloud controller</td><td>cinder-*</td><td><code>/var/log/cinder</code></td></tr><tr><td>Cloud controller</td><td>keystone-*</td><td><code>/var/log/keystone</code></td></tr><tr><td>Cloud controller</td><td>neutron-*</td><td><code>/var/log/neutron</code></td></tr><tr><td>Cloud controller</td><td>horizon</td><td><code>/var/log/apache2/</code></td></tr><tr><td>All nodes</td><td>misc (swift, dnsmasq)</td><td><code>/var/log/syslog</code></td></tr><tr><td>Compute nodes</td><td>libvirt</td><td><code>/var/log/libvirt</code> <code>/libvirt.log</code></td></tr><tr><td>Compute nodes</td><td>Console (boot up messages) for VM instances:</td><td><code>/var/lib/nova</code> <code>/instances/instance- <instance id>/console.log</code></td></tr><tr><td>Block Storage nodes</td><td>cinder-volume</td><td><code>/var/log/cinder</code> <code>/cinder-volume.log</code></td></tr></tbody></table> <p>(Ref: https://docs.openstack.org/operations-guide/ops-logging.html)</p> <p>A) node network connectivity failure</p> <ol style="list-style-type: none">1. management network2. VMs communication network3. storage network <p>B) nova service failure (e.g., process crashed) -- detected and restarted by a local watchdog process</p> <ol style="list-style-type: none">1. compute2. volume3. network4. scheduler5. api. <p>C) Failure of other OpenStack services -- N/A, assuming redundant /highly available configuration</p> <ol style="list-style-type: none">1. Glance2. Keystone	Cloud controller	nova-*	<code>/var/log/nova</code>	Cloud controller	glance-*	<code>/var/log/glance</code>	Cloud controller	cinder-*	<code>/var/log/cinder</code>	Cloud controller	keystone-*	<code>/var/log/keystone</code>	Cloud controller	neutron-*	<code>/var/log/neutron</code>	Cloud controller	horizon	<code>/var/log/apache2/</code>	All nodes	misc (swift, dnsmasq)	<code>/var/log/syslog</code>	Compute nodes	libvirt	<code>/var/log/libvirt</code> <code>/libvirt.log</code>	Compute nodes	Console (boot up messages) for VM instances:	<code>/var/lib/nova</code> <code>/instances/instance- <instance id>/console.log</code>	Block Storage nodes	cinder-volume	<code>/var/log/cinder</code> <code>/cinder-volume.log</code>	<p>Interfaces statistics - sent, recv, drops</p> <p>Hypervisor Metrics, Nova Server Metrics, Tenant Metrics, Message Queue Metrics</p> <p>Keystone and Glance Metrics</p>	
Cloud controller	nova-*	<code>/var/log/nova</code>																																
Cloud controller	glance-*	<code>/var/log/glance</code>																																
Cloud controller	cinder-*	<code>/var/log/cinder</code>																																
Cloud controller	keystone-*	<code>/var/log/keystone</code>																																
Cloud controller	neutron-*	<code>/var/log/neutron</code>																																
Cloud controller	horizon	<code>/var/log/apache2/</code>																																
All nodes	misc (swift, dnsmasq)	<code>/var/log/syslog</code>																																
Compute nodes	libvirt	<code>/var/log/libvirt</code> <code>/libvirt.log</code>																																
Compute nodes	Console (boot up messages) for VM instances:	<code>/var/lib/nova</code> <code>/instances/instance- <instance id>/console.log</code>																																
Block Storage nodes	cinder-volume	<code>/var/log/cinder</code> <code>/cinder-volume.log</code>																																
Application	Crash/Connectivity /Non-Functional	Application Log i.e. If it is Apache then logs of Apache (<code>/var/log/apache2</code>)	Packet Drops, Latency, Throughput, Saturation, Resource Usage	Deploy Collectd within the application and collect both application logs and infrastructure metrics																														
Middleware Services																																		

Models

We have taken three types of models and in those models we have considered **Failure Prediction** problem and the remaining types are given as:

1. Event Correlation
2. Anomaly Detection
3. Failure Prediction

We are focusing on Failure Prediction of Node, Application, VM, Service, Container and Links. Our aim is to predict the failures before they happen so that user can take necessary actions regarding those failures. So, to implement Failure Prediction models we are developing our models using Classical Neural Networks techniques i.e. **RNN & LSTM**.

Gaps

From the perspective of Telco after doing a literature survey we found most of the work has done on VM and Applications. There are less work has done for Node Failures, Link Failures, Middleware Services and also there is a lack of Publicly available datasets for these failures. Majority of researchers have used **ARIMA & RNN** so to improve the performance of the prediction model we can do some experiments with **Generative Adversarial Networks (GAN)**, **Graphical Neural Network (GNN)**. Also, in our literature survey we found that majority of the publicly available data does not contain time stamp. So to make the future predictions we will need of Time Series data.

Enhancements