Network Monitoring and Analysis in a Cluster Environment

One aspect of Sodalite is to monitor the usage of the resources on a Cluster and to use that information to adjust applications and their deployment to improve their performance.

Network Monitoring

Numerous tools exist to monitor critical resources in a computer cluster. One of the critical resources in a computer cluster is its network connectivity and performance. In this post we discuss Skydive – a network analysis tool.

What is Skydive?

Skydive (http://skydive.network) is a real-time network topology and protocols analyzer that provides detailed network topology and performance information. Skydive agents collect topology information and flows and forward them to a central agent for further analysis.

Network topology support of Skydive allows one to view relationships between hosts (both physical and virtual), containers (network namespace), and network entities (device, bridge, veth, tun, macvlan, etc). Topology agent probes exist for the following environments and tools: Docker, Ethtool, Libvirt, LLDP, Lxd, NetLINK, NetNS, Neutron, OVSDB, Opencontrail, runC, socketinfo, VPP.

Skydive is extensible and allows the development of probes for new kinds of environments. Probes already exist for Kubernetes, Istio, NSM, OVN. For example, the Kubernetes probe provides information on clusters, namespaces, nodes, pods, containers, services, network policies, volumes, and deployments.

Skydive also provides support to track specific network flows of various types including: sFlow, AFPacket, PCAP, DPDK, eBPF, and OpenvSwitch port mirroring.

The extensible Skydive framework allows a developer to implement probes for new kinds of topology entities and flow types.

Detailed documentation of Skydive can be found at http://skydive.network/documentation.
Additional insights can be gleaned by looking through source code, especially the configuration possibilities defined in the skydive.yml.default file.
See https://github.com/skydive-project/skydive/blob/master/etc/skydive.yml.default.

What is Skydive good for?

Various networking tools provide information on bandwidth and response time for network interfaces. Tools also exist to track (sniff) packets that flow through a network interface based on various filter criteria. Skydive builds upon these capabilities to provide information about end-to-end flows. For example, based on the standard network 5-tuple (protocol type, source address, destination address, source port, target port) in packet headers, Skydive can give detailed information about network flows on specific connections (port to port). This information is invaluable in helping an administrator understand how the network is being utilized and to discover anomalies in the network.

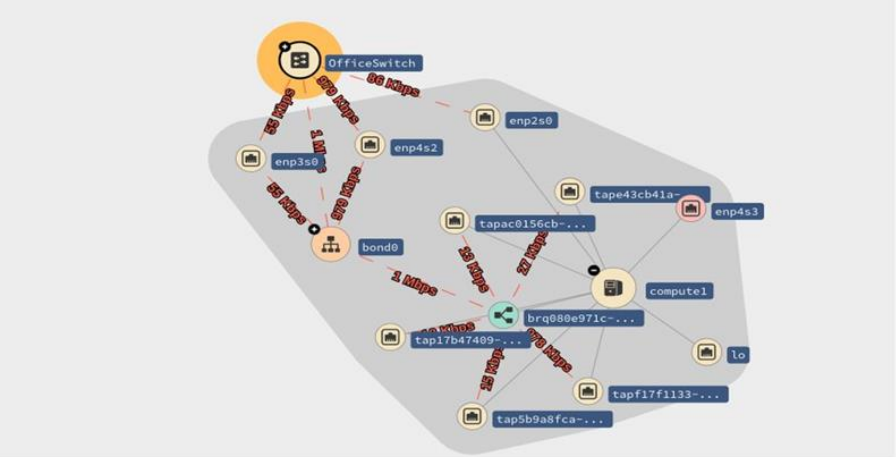

Let's examine several use cases.

Detect failing resources

Starting from the topology visualization of Kubernetes objects in the Skydive GUI, an administrator can see up/down status per resource (e.g. pod, pvc, pv), and can traverse from a failed pod down to the storage resource causing the error.

Network troubleshooting and debugging

Suppose a networking engineer needs to understand why one of the multi-service Kubernetes environments behaves strangely. He needs to explore, in real-time and in the right context, the network topology, bandwidth and latencies, and use a protocol analyzer to understand the flows. He then performs a fix and wants to monitor the system to assure that the abnormal network behavior was resolved.

Network operational analytics

Suppose an administrator needs to suggest enhancements to a company's large-scale Kubernetes environment. He needs to find out which areas of the Kubernetes environment are bottle-necks at what times. He then needs to drill-down to the pods level and check over time the communication behavior to better advise the set of app/pods to enhance.

Performance testing and profiling

Suppose a testing engineer needs to test the performance of new network equipment in the infrastructure level. He needs to inject traffic from multiple layers of the network: application level, Kubernetes level, and infrastructure level. Then, he monitors the behavior of the network to assure that the new equipment is working as expected.

Security coverage and compliance testing

Suppose a security engineer needs to validate that security rules in the network policy are working as expected. He needs a way to inject TCP and UDP traffic with pre-defined ports from and against various locations in the network topology to validate pass/fail conditions and create a compliance report.

All the above scenarios can be accomplished using Skydive features of topology discovery, flow analysis, packet injection, etc.

A Skydive network flow application - Security Advisor

IBM recently integrated a Security Advisor capability on top of Skydive. The Security Advisor filters the flow data obtained from Skydive, performs a data transformation, and saves the information to an object store in JSON format. This is the first in a series of Skydive data exporters planned to be released. Details describing how to use the Security Advisor are available on the Skydive blog page: http://skydive.network/blog. The Security Advisor looks at the flows reported by Skydive associated with the subnets defined for a cluster. A particular case of interest is when a network flow originates outside the cluster and connects to an entity inside the cluster. The Security Advisor flags these flows as "ingress" flows, and these may be further scrutinized by an administrator to determine whether there is a security issue. This feature is central to the notion of a security group for a Virtual Private Cloud (VPC). A separate subnet is defined for the VPC, and any traffic originating from outside the VPC is flagged. Even if there is a firewall to prevent the improper traffic from penetrating the subnets associated with the VPC, it is important to be aware of attempts by outsiders to connect to entities inside the VPC.

Additional details on how to use the Skydive data exporter, flow logs, and other features are available on the Skydive blog page http://skydive.network/blog.