

2016-04-20 Workbench load balancer

Notes from discussion with Nebula team about possible options for the [Cluster Loadbalancer](#).

Overview

The basic requirement is to allow NDS Labs Workbench users to securely access workbench services, which include both HTTP and TCP services. This means implementing TLS for all services. Services run as Pods in Kubernetes. Via Kubernetes network model, each service is assigned a unique IP. Internally, services can run on well-known ports (e.g. 80, 8080, 1247, etc).

Options

We discussed three different solutions for routing traffic from a load balancer to services running in Kubernetes.

1. Path: Loadbalancer serves as re-writing proxy. HTTP requests are in the form `labs.nds.org/project/service` and routed to the correct Pod by monitoring changes to etcd
2. Port: Loadbalancer exposes non-HTTP services as ports and routes to the correct Pod by monitoring changes to etcd.
3. IP Blocks
 - a. Nebula team allocates a fixed block of IPs
 - b. Each project is assigned an IP address and CNAME (e.g., `project.labs.nds.org`)

Other notes

- NCSA and <https://letsencrypt.org/> provide free SSL certs. Letsencrypt will do so on the fly – via API
- Certs available from NCSA by sending request to help+security@ncsa.illinois.edu (1-day response)
- DNS:
 - there are no name restrictions if not in `illinois.edu` domain
 - Neteng handles reverse lookup
 - Send DNS requests to help+neteng@ncsa.illinois.edu
 - Nebula doesn't support DNS today, but hopes to in the next 6 months
- Allocated 1 IP for management (Workbench GUI – probably need 1 IP for production 1 for staging.
- For each project:
 - Allocate 1 IP
 - Send CNAME request to neteng (1 business day)
- How many projects do we anticipate
 - Nebula team prepared to allocate 128 IPs for NDS

General questions

- Q. What is nebula storage?
 - Instance storage = ephemeral storage
 - Volume storage = Cinder, iSCSI
- Q. When and why Swift
 - Offering Swift storage in addition to Cinder
 - Swift is object based with stupid S3
 - When do you want to use an object store — additional metadata
 - Make multiple copies, etc.
 - Currently working on storage:
 - Network performance
 - Move Glance to shared backend
 - Size up Cinder (volumes big enough)
 - Build out Swift
 - Growing cluster — they have a bunch of storage nodes that have ungodly
- Q. Is there a practical limit on the number of volumes
- Q. Running production on Nebula
 - No, we shouldn't be – not yet. You're restricted to running on a single compute node
 - For example, `compute3` is problematic. Or not yet
 - When the backend is in a parallel, and you can migrate, as we do in VMware
 - Then you can run a production service
 - No UPS on Nebula
 - It runs at NCSA — power outages happen
 - So do Kerberos and LDAP (do have UPS)
 - Nebula running since September
 - Currently have one zone, verified that they can run multiple zones
 - Elastic guys? For elastic compute?
 - LSST only uses compute resources at night.
- Q. Where do production services run ?
 - VSphere/VMWare
 - They are hosted for projects

- For example, license servers for HPC, Jenkins for BlueWaters
 - LDAP services, Webservices for NCSA
 - Runs here and NPCF (BlueWaters)
- OpenStack user group meetings