

Environment

- OpenShift 3.9
- Microsoft Azure

Issu Summary

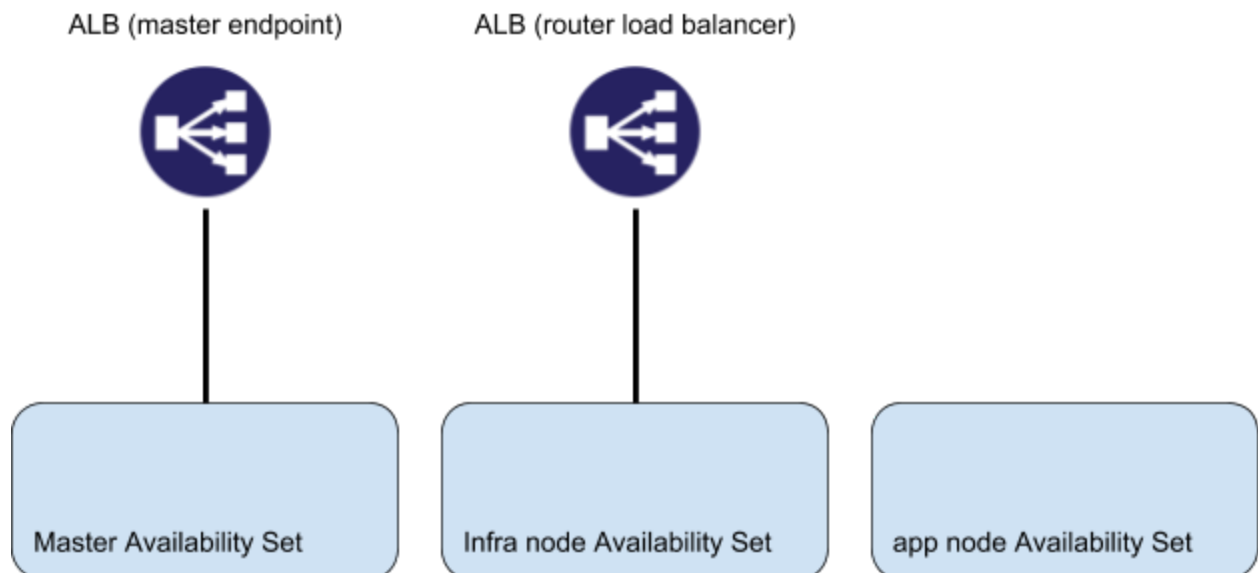
If OpenShift cluster has an Azure Load Balancer (ALB) as a router load balancer, OpenShift can't provision LoadBalancer Type service whose backend pod is running on infra nodes.

Issue Detail

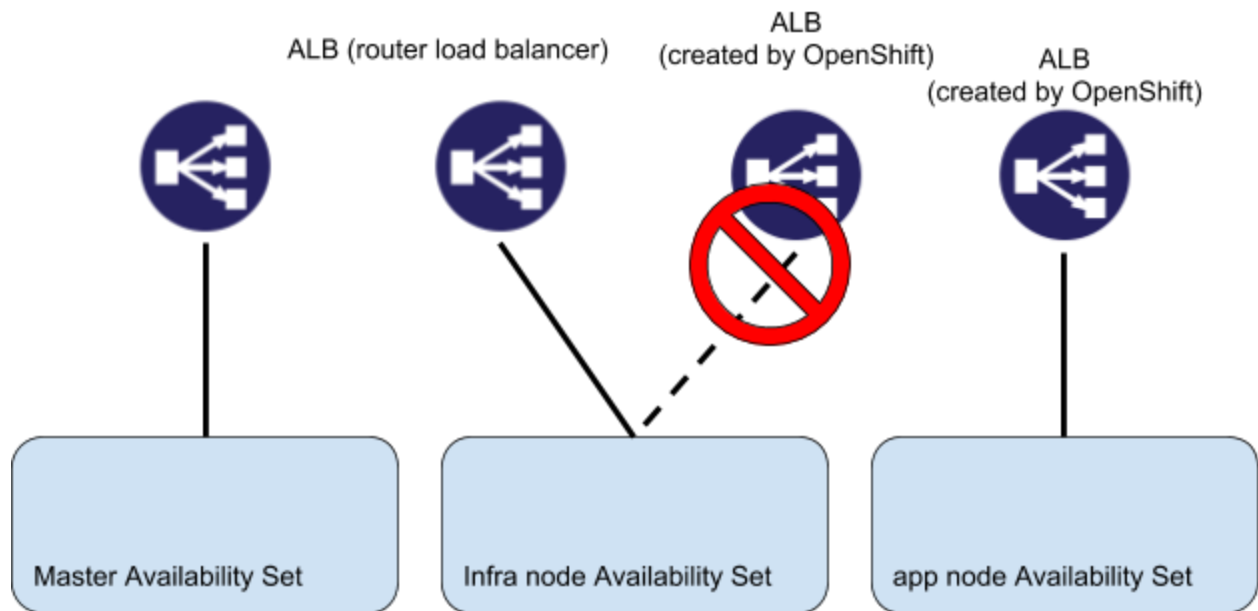
A customer wants to provision [LoadBalancer Type service \[1\]](#) in OpenShift on Azure. In Azure, kubernetes creates [Azure Load Balancer \(ALB\) Basic \[2\]](#) for LoadBalancer Type Service. ALB Basic has following limitations.

- [One ALB has only a single Availability Set as a backend pool \[3\]](#)
- One Availability Set can belong to at most one ALB

On the other hand, the customer can provision ALB as an OpenShift router load balancer. This is described in [OpenShift on Azure reference architecture \[4\]](#). In this case, an availability set for infra nodes has already belonged to an ALB.

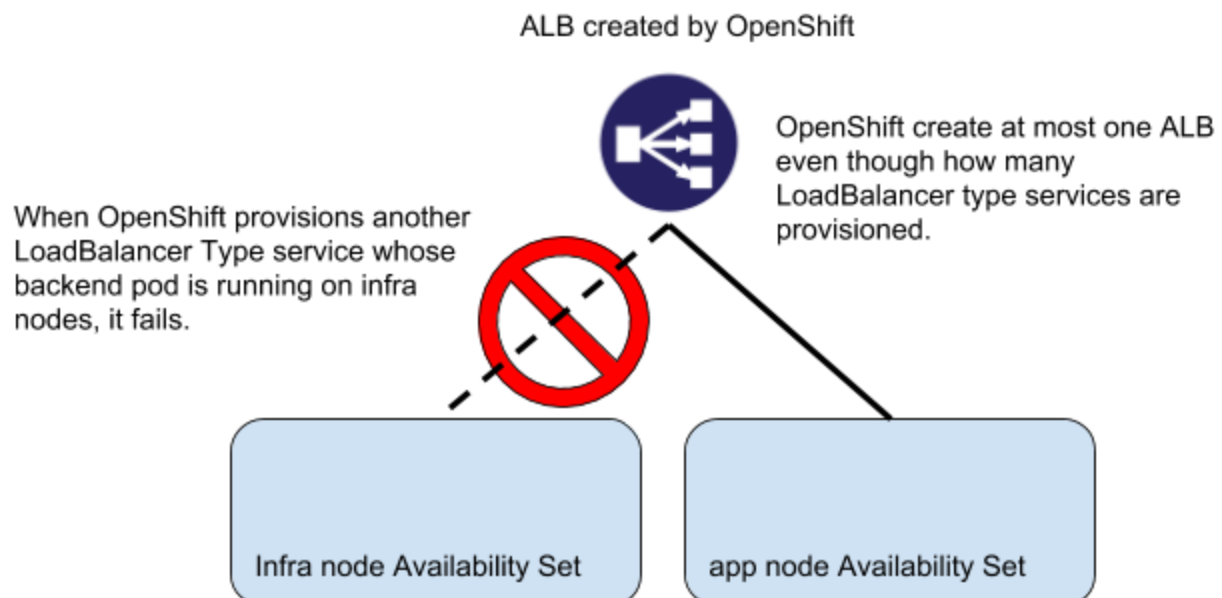


In this situation, if the customer creates the LoadBalancer Type service whose backend pod is running on one of infranodes, OpenShift tries to create an ALB whose backend pool is the availability set but it fails by the above ALB limitation.



Notes:

In OpenShift 3.7 or lower version, meaning kubernetes 1.7 or lower, OpenShift can't consider several Availability Set for LoadBalancer Type service. It means OpenShift fails to create two or more LoadBalancer Type services whose backend pods belong to application nodes and infra nodes.



In OpenShift 3.9 or later, LoadBalancer Type service considers which Availability Set should be used with "service.beta.kubernetes.io/azure-load-balancer-mode" annotation. However, this is undocumented annotation even in the kubernetes upstream document. I have made [a documentation request \[5\]](#).

Background - why the customer wants to use LoadBalancer Type service.

OpenShift default router supports only http/https/sni. If a customer wants to expose other ports such as HTTP 2.0 (e.g. gRPC) and MQTT, they have to route the traffic to nodes. It can be done manually without using LoadBalancer Type service, but it's too annoying. LoadBalancer Type service enables to expose these ports with OpenShift.

What's next?

At least, "service.beta.kubernetes.io/azure-load-balancer-mode" annotation should be documented and supported if we support LoadBalancer type service in OpenShift on Azure. In addition,

- [] Use Application Gateway or non-Azure managed Load Balancer for OpenShift router load balancer.
- [] Support ALB Standard SKU for kubernetes LoadBalancer Type service in Azure. (ALB Standard SKU)
- [] Fix kubernetes logic to use an existing ALB if existed instead of creating a new ALB.

[1] <https://kubernetes.io/docs/concepts/services-networking/service/#type-loadbalancer>

[2] <https://docs.microsoft.com/en-us/azure/load-balancer/load-balancer-overview>

[3]

<https://docs.microsoft.com/en-us/azure/load-balancer/load-balancer-standard-overview#why-use-standard-load-balancer>

[4]

https://access.redhat.com/documentation/en-us/reference_architectures/2018/html/deploying_and_managing_openshift_3.9_on_azure/red_hat_openshift_container_platform_instance_prerequisites#router_load_balancer

[5] <https://github.com/kubernetes/website/issues/8154>