December 18, 2018

SAFE UNROUTED ADDRESSING IN A CLOUD OR DATACENTER NETWORK

Kyle Mestery

Ian Wells

SAFE UNROUTED ADDRESSING IN A CLOUD OR DATACENTER NETWORK

AUTHORS:
Kyle Mestery
Ian Wells

ABSTRACT

Techniques are described for multi-tenant isolation in cloud and datacenter networks using normal Layer 3 (L3) routing rather than overlay networks. Basically, the wider routed network is divided into three components: public and corporate addresses, private addresses, and fully reachable addresses (i.e., addresses that can route to both the public and corporate and the private addresses).

DETAILED DESCRIPTION

In a typical cloud network environment composed of a combination of physical servers, virtual machines, containers, and server-less functions, the need for addressing is a requirement usually placed on the cloud management system. This handles addressing the underlying physical infrastructure. However, tenants themselves have addressing requirements. This results in virtual networks being created to handle these tenants, and allow for the overlapping of Internet Protocol (IP) addresses inside. Typically the addressing the tenant receives is of little consequence, and in fact the default for a tenant network is often used by most tenants. This means most tenants have the same overlapping IP space. Accordingly, there exist tenant isolation protocols to handle this overlapping address space.

Similarly, when building a cloud infrastructure, the cloud components have a need to communicate with each other internally. One method of deployment is to use either Internet Engineering Task Force (IETF) Request for Comments (RFC) 1918 space or Unique Local Addresses (ULAs) to provide for this, but in both cases the cloud can be on a corporate network where IETF RFC 1918 networking is already in use. It is desirable for the address to be within the same routing domain (e.g., Virtual Routing and Forwarding (VRF), etc.) so it must be unique, but need not be widely routable (e.g., a split horizon where that address is routed within the infrastructure but not outside).

For both cases, a mechanism is provided to find a set of addresses that can be used for this purpose and are unique within the space of reachable addresses the domain might wish to connect to.

In particular, overlay networks are removed as a form of isolation and instead normal Layer 3 (L3) routing is used. Normal L3 routing (both IPv4 and IPv6) is used to isolate tenant traffic when the underlying infrastructure is providing cloud networking. Tenants can have either IPv4 or IPv6 addresses, or both. The infrastructure itself will handle delegating ranges of addresses when new components come onboard.

In one example, the infrastructure makes a request of the wider network for the addresses to use. The addresses provided are guaranteed to be unique within (a) the domain they are being delegated to and (b) the wider network of addresses, so that a routable address in the domain can be assured that it can go north to a Wide Area Network (WAN) or corporate network address, or south to a delegated address, since no address is duplicated in either of its horizons.

However, the delegated addresses do not provide routability in general because there is no route that will return traffic to them. They are intended for internal use within the domain. This means that they can be re-used multiple times, e.g. for multiple cloud deployments in the infrastructure case or for multiple tenants within the cloud in the workload case.

Since tenants typically only want a network with addressing, and not a network with specific addresses, this is more optimal than making them choose and using overlay networks to isolate traffic. The same goes for cloud infrastructure deployment. The cloud can use a set of addresses it is assured will not affect its routability. In both cases, the person in charge does not have an overall view of addresses used in the network and cannot make a good choice for itself.

Basically, the wider routed network is divided into three components: public and corporate addresses, private addresses, and fully reachable addresses (i.e., addresses that can route to both the public and corporate and the private addresses). For this to work, the private addresses must not overlap with the public and corporate addresses, but may overlap with other domains of private addresses also attached to the network. A system in the corporate network or a function on the nexthop router provides an address range that

may be assigned to tenant workloads or infrastructure components. This range may be repeatedly used in different subdomains within the network matching the above properties. For instance, multiple tenants may be given 10.0.0.0/24, since they do not interact with each other, or a cloud being deployed may be given 10.0.0.0/24 for its compute hosts, since the compute hosts need only be reachable from the control hosts. Bastion hosts or control hosts can then be given fully reachable addresses from the corporate network address space using conventional IP Address Management (IPAM) methods (e.g., a big spreadsheet). The system providing private addresses may re-use the private address space assigned for the purpose over and over again, providing it does not provide the same address range more than once to the same private domain.

It may also place routing policies on the nexthop router as addresses are assigned either to prevent packets with those source addresses being forwarded into the network or noting and reporting such attempts (which would indicate misconfigurations within the private space). Further, it may be used to configure Network Address Translation (NAT) should the private addresses be permitted to reach into the public space (the private and public spaces do not overlap).

This differs from conventional IPAM and IPv6 prefix delegation. IPAM in the conventional sense is used to manage addresses with full routability. Prefix delegation similarly provides fully routable addresses. This is providing addresses which are only consumable within the domain, but avoids the use of VRFs by assuring that they do not overlap with routable addresses, simplifying software deployment.

Figure 1 below illustrates uniqueness requirements of example address spaces.

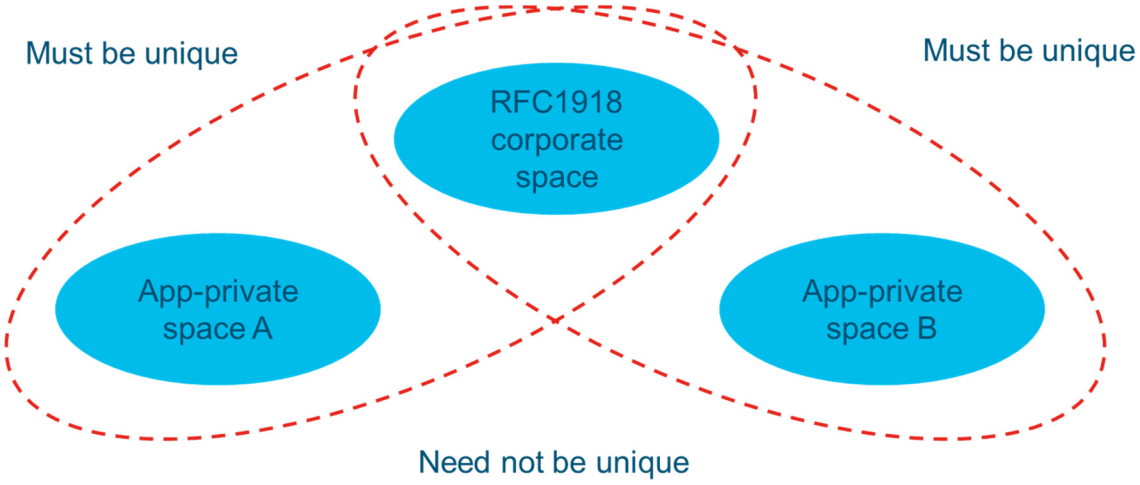

Figure 1

Figure 2 below illustrates relevant properties of the example address spaces.

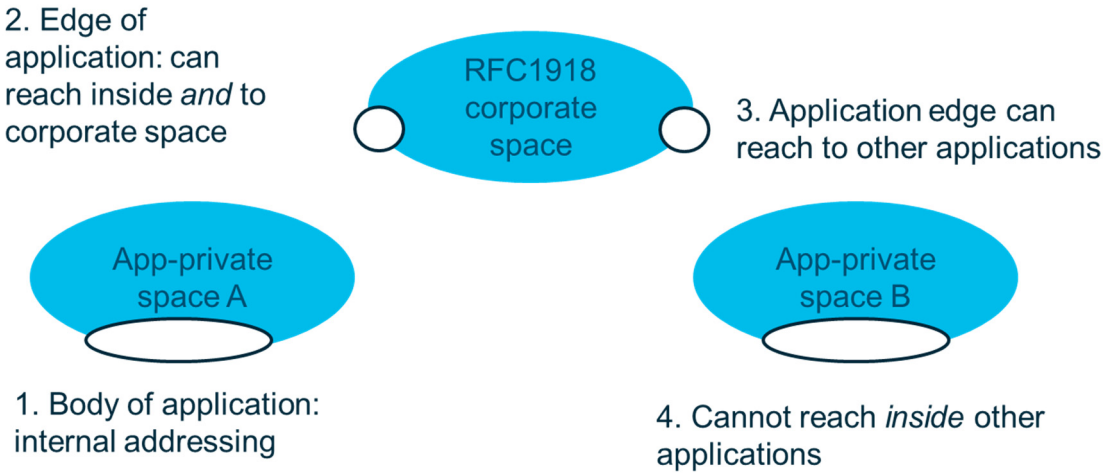

Figure 2

Figure 3 below illustrates how the same address range can be given to multiple applications because the applications need not reach inside each other.

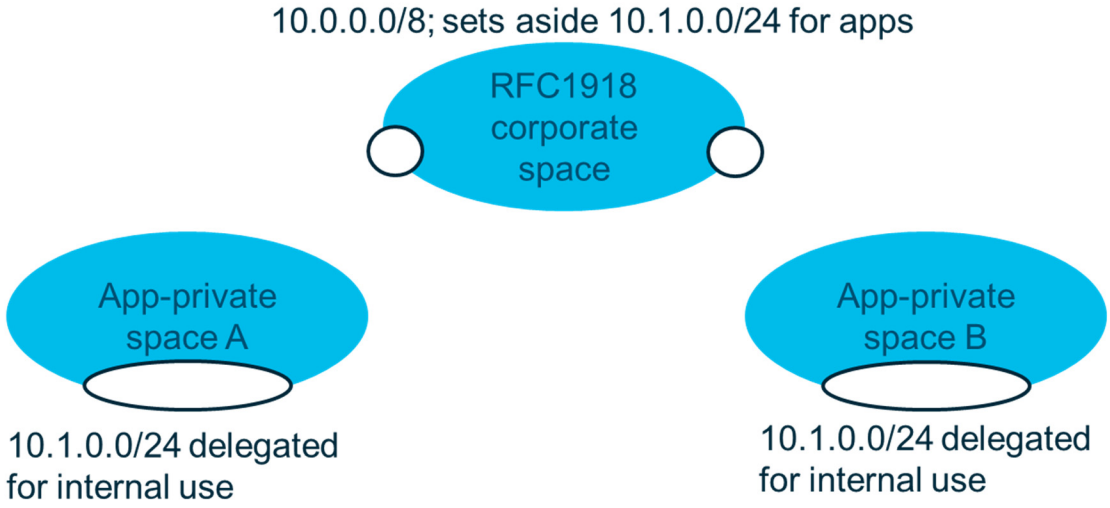

Figure 3

In summary, techniques are described for multi-tenant isolation in cloud and datacenter networks using normal L3 routing rather than overlay networks. Basically, the wider routed network is divided into three components: public and corporate addresses, private addresses, and fully reachable addresses (i.e., addresses that can route to both the public and corporate and the private addresses).