Introduction to IP Addressing and Routing in MCP 1.0

Description

Describes how private and public IP's in an MCP 1.0 location are assigned, what their routing characteristics are, and the Public IP Address Usage Policy. Only IPv4 addresses are supported in MCP 1.0 locations.

NOTE: This article applies only to MCP 1.0 locations. For details on MCP 2.0 behavior, see Introduction to IP Addressing in MCP 2.0

Content / Solution:

Public IPv4 Addresses (Public Internet Routable)

In MCP 1.0, each Cloud Network comes with a small block of public IPv4 addresses already assigned to it. Users can view their "default" IP block or add additional public IP address blocks as described in How to Manage Public IPv4 Addresses on a Network or Network Domain

Clients are not charged for the default public IP addresses deployed with a Cloud Network, however additional public IP addresses added to the network do incur usage charges. For more details, see How are Public IP Addresses Billed for on the Cloud

Regardless of pricing, customers are allowed to deploy IP's onto a Cloud Network only up to a maximum ratio of three public IP addresses per deployed Cloud Server. So if a customer has deployed a total of 10 Cloud Servers on a Cloud Network, they are allowed to deploy only a maximum total of 30 public IP addresses on that Cloud Network.

Most data center locations assign IP addresses in blocks of 2 IP addresses. However, a few older data center locations have a mix of block sizes. In the Admin UI, customers cannot choose the size of the block - the smallest size available is always added. The Cloud v0.9 API does allow a request size through the *Reserve Public IP Block with Size* function, but that size may not be honored. The affected data centers with a mix of block sizes are:

- **US-East (NA1), US-West (NA3), Amsterdam (EU1):** Cloud Networks and IP blocks use a mix of 2 and 8 IP addresses
- **Sydney (AU1):** Cloud Networks and IP blocks use a mix of 4 and 8 IP addresses
- **All other Locations**: Cloud Networks and IP blocks use a size of 2 IP addresses

Note that in MCP 1.0 locations only, Server to VIP connectivity on a public IP address is disabled - meaning that by default, you cannot communicate to a public IPv4 VIP from Cloud Servers located on the same Cloud Network. This behavior can be adjusted as described in How to Manage Public IPv4 Addresses on a Network or Network Domain. However, as described in the article, there is a limit on the number of IP blocks on a given Cloud Network that can have this function enabled.

Private IPv4 Addresses (Routable Both Within and Between Cloud Networks)

Each Cloud Network is assigned a single /24 block of 255 private IP addresses, although only 244 private IP addresses are available for use. The system reserves 10.x.x.0 through 10.x.x.10 and 10.x.x.255 for its own use. This space is used to support backups, multicast, and other Cloud features.

The entire IPv4 RFC 1918 range is routable between different Cloud Networks, including Cloud Networks in different data center location, different Geographic Regions. **This is true even of Cloud Networks controlled by different CloudControl organizations, meaning that it is critical that users take into account this routing behavior from a security perspective when defining ACL rules, particularly since the "default" ACL rules on Cloud Networks allow such private IP traffic.**

In addition, when communicating between Public MCP 1.0 locations, this functionality allows us to automatically encrypt traffic between Cloud Network deployments in geographically dispersed locations. Traffic with Private Cloud locations is encrypted but not accelerated. This behavior is described in more detail in How does Outbound Traffic Work on a Cloud Network in a MCP 1.0 location.

In an MCP 1.0 location, each Cloud Server has a single NIC and can only be deployed onto a single Cloud Network. Users can choose the private IPv4 address of the Cloud Network at deployment or CloudControl can assign it from the available IP addresses registered with the system. If changes are made inside the Guest OS to these assignments, users should update the Cloud system record as described in How to Notify CloudControl of a Change to the Private IP Address of a Server in a MCP 1.0 Data Center.

It is not recommended or supported, but it is possible to deploy additional private IP addresses to a server as described in How to Add Additional IP Addresses to a Red Hat / CentOS Linux Server. However, the customer takes all responsibility for potential private IP address conflicts that may occur as a result of using such an approach.

Clients can make the private IP addresses of these MCP 1.0 Cloud Servers publicly available to the Internet using either:

- The NAT function described at How to Create a NAT Rule on a Network or Network Domain. This function directly maps all ports on a public IP address to the same ports on single private IP address.
- The VIP function described at What is a VIP in a MCP 1.0 Data Center Location allows a customer to map a specific port on a public IP address to a port (potentially a different port) on one or more private IP addresses. This allows customers to do both port mapping and load balancing functions.

IP Uniqueness and Exclusive Reservation of IP Addresses

IP Uniqueness Policy *(Changes Effective October 19, 2017)*

CloudControl tracks which private IPv4 addresses are on Cloud Servers on the Cloud Network. Prior to October 19th, the system prevented those addresses from being "re-used" and assigned to other Cloud Servers. Effective October 19th, the system no longer enforces uniqueness constraints on the use of private IPv4 addresses. Therefore, users can assign the same IPv4 address to different Cloud Servers on the Cloud Network. Such behavior has the potential to cause IP conflicts when multiple servers broadcast the same IP address on a Cloud Network. Non-unique address assignments are required in a variety of scenarios but those scenarios usually require configurations not supported within MCP 1.0.

A few notes and restrictions on the uniqueness policy

1. The system will still require a unique IPv4 private address when deploying a Cloud Server. However, once the server has been deployed, users can change the IPv4 addresses as described in How to Notify CloudControl of a Change to the Private IP Address of a Server in a MCP 1.0 Data Center and uniqueness is no longer enforced.
2. To help avoid and diagnose IP conflict issues, the system now allows you to "disconnect" a NIC while leaving a server powered on, allowing you to remove connectivity to the VLAN while maintaining the server in a powered on state so that you can access the Guest OS to diagnose issues. For details, see Introduction to Cloud Server NICs in MCP 1.0.
3. The UI will notify of potential IP conflicts based on its record of the IPv4 NIC address assignments if the following conditions are met:
    a. Two (or more) devices share the same IP Address
    b. The devices are powered on
    c. The NICs are connected

IP Exclusive Reservations

The system does support "exclusive reservation" of private IPv4 addresses in MCP 1.0, but the functionality is currently available only using the API. For details on the API function, see the current documentation at API 2 - Documentation Downloads

1. "Exclusively reserving" an IPv4 address in MCP 1.0 now prevents the use of the IP address going forward from the time of the reservation. This means that once an IP address is exclusively reserved, the system will not allow that IP address be used when deploying Cloud Servers or notifying the system of an IP address change.
    a. If this address is uniquely used at the time the IP is designated as exclusively reserved, it means it will stay uniquely assigned thereafter.
    b. However, an exclusively reserved address doesn't ensure uniqueness as the IP address may have been used on multiple Cloud Servers prior to its designation as exclusively reserved.
2. When a Cloud Server is deleted, if the associated IP address is exclusively reserved and it is the only server known to CloudControl to be using the private IPv4 address, the exclusive reservation is removed along with the server.
3. Cloud Networks cannot be deleted if an IP address is exclusively reserved on them. The reservation must be removed via API in order to allow deletion of the Cloud Network. If you do not know how to use the API and are prevented from deleting a Cloud Network due to this issue, contact Support.

IPv6 Addresses (Not Supported)

IPv6 addresses are not supported or routable in MCP 1.0 locations.

Related Articles

- How are Public IP Addresses Billed for on the Cloud
- How to Notify CloudControl of a Change to the Private IP Address of a Server in a MCP 1.0 Data Center
- Introduction to IP Addressing and Routing in MCP 1.0
- How to Manage Public IPv4 Addresses on a Network or Network Domain
- Import Function Support for Microsoft Windows Operating Systems