



Assigning Network and IP Addresses

This appendix contains the following sections:

- [Ethernet Interfaces, on page 1](#)
- [Selecting IP Addresses and Netmasks, on page 1](#)
- [Strategies for Connecting Your Content Security Appliance , on page 3](#)

Ethernet Interfaces

Appliances have up to four Ethernet interfaces located on the rear panel of the system, depending on the configuration (whether or not you have the optional optical network interface). They are labeled:

- Management
- Data1
- Data2
- Data3
- Data4

Selecting IP Addresses and Netmasks

When you configure the network, the appliance must be able to select a unique interface to send an outgoing packet. This requirement drives some of the decisions regarding IP address and netmask selection for the Ethernet interfaces. The rule is that only one interface can be on a single network (as determined through the applications of netmasks to the IP addresses of the interfaces).

An IP address identifies a physical interface on any given network. A physical Ethernet interface can have more than one IP address for which it accepts packets. An Ethernet interface that has more than one IP address can send packets over that interface with any one of the IP addresses as the source address in the packet. This property is used in implementing Virtual Gateway technology.

The purpose of a netmask is to divide an IP address into a network address and a host address. The network address can be thought of as the network part (the bits matching the netmask) of the IP address. The host address is the remaining bits of the IP address. The number of bits in a four octet address that are significant are sometimes expressed in Classless Inter-Domain Routing (CIDR) style. This is a slash followed by the number of bits (1-32).

A netmask can be expressed in this way by simply counting the ones in binary, so 255.255.255.0 becomes “ /24 ” and 255.255.240.0 becomes “ /20. ”

Sample Interface Configurations

This section shows sample interface configurations based on some typical networks. The example uses two interfaces called Int1 and Int2. In the case of the appliance, these interface names can represent any two interfaces out of the three interfaces (Management, Data1, Data2).

Network 1:

Separate interfaces must appear to be on separate networks.

Interface	IP Address	Netmask	Net Address
Int1	192.168.1.10	255.255.255.0	192.168.1.0/24
Int2	192.168.0.10	255.255.255.0	192.168.0.0/24

Data addressed to 192.168.1.X (where X is any number from 1 through 255, except for your own address, 10 in this case) go out on Int1. Anything addressed to 192.168.0.X goes out on Int2. Any packet headed for some other address not in these formats, most likely out on a WAN or the Internet, is sent to the default gateway, which must be on one of these networks. The default gateway then forwards the packet on.

Network 2:

The network addresses (network parts of the IP addresses) of two different interfaces cannot be the same.

Ethernet Interface	IP Address	Netmask	Net Address
Int1	192.168.1.10	255.255.0.0	192.168.0.0/16
Int2	192.168.0.10	255.255.0.0	192.168.0.0/16

This situation presents a conflict in that two different Ethernet interfaces have the same network address. If a packet from the content security appliance is sent to 192.168.1.11, there is no way to decide which Ethernet interface should be used to deliver the packet. If the two Ethernet interfaces are connected to two separate physical networks, the packet may be delivered to the incorrect network and never find its destination. The appliance does not allow you to configure your network with conflicts.

You can connect two Ethernet interfaces to the same physical network, but you must construct IP addresses and netmasks to allow the appliance to select a unique delivery interface.

IP Addresses, Interfaces, and Routing

When you select an interface on which to perform a command or function in the GUI or CLI that allows you to select an interface (for example, upgrading AsyncOS or configuring DNS), routing (your default gateway) takes precedence over your selection.

For example, suppose that you have a appliance with the three network interfaces configured, each on a different network segment (assume all /24):

Ethernet	IP
Management	192.19.0.100
Data1	192.19.1.100

Ethernet	IP
Data2	192.19.2.100

And your default gateway is 192.19.0.1.

Now, if you perform an AsyncOS upgrade (or other command or function that allows you to select an interface) and you select the IP that is on Data1 (192.19.1.100), you would expect all the TCP traffic to occur over the Data1 Ethernet interface. However, instead the traffic goes out of the interface that is set as your default gateway, in this case Management, but is stamped with the source address of the IP on Data1.

Summary

The content security appliance must always be able to identify a unique interface over which a packet can be delivered. To make this decision, the appliance uses a combination of the packet's destination IP address, and the network and IP address settings of its Ethernet interfaces. The following table summarizes the preceding examples:

	Same Network	Different Network
Same Physical Interface	Allowed	Allowed
Different Physical Interface	Not allowed	Allowed

Strategies for Connecting Your Content Security Appliance

Keep the following in mind when connecting your appliance :

- Administrative traffic (CLI, web interface, log delivery) is usually little compared to email traffic.
- If two Ethernet interfaces are connected to the same network switch, but end up talking to a single interface on another host downstream, or are connected to a network hub where all data are echoed to all ports, no advantage is gained by using two interfaces.
- SMTP conversations over an interface operating at 1000Base-T are slightly faster than conversations over the same interfaces operating at 100Base-T, but only under ideal conditions.
- There is no point in optimizing connections to your network if there is a bottleneck in some other part of your delivery network. Bottlenecks most often occur in the connection to the Internet and further upstream at your connectivity provider.

The number of interfaces that you choose to connect and how you address them should be dictated by the complexity of your underlying network. It is not necessary to connect multiple interfaces if your network topology or data volumes do not call for it. It is also possible to keep the connection simple at first as you familiarize yourself with the gateway and then increase the connectivity as volume and network topology require it.

