The Cable Guy - September 2003

74 out of 88 rated this helpful – Rate this topic

Default Gateway Behavior for Windows TCP/IP

By The Cable Guy

TCP/IP hosts can use the following methods to reach remote destinations:

- **Store a host-specific route to each remote destination.** This is obviously not practical or possible, as the routing table might have to contains thousands, or in the case of the Internet, millions of routes. The host routing table would have to change as new addresses are added or removed.
- **Store a network route to each remote subnet.** Although more possible, this is also not practical, as the routing table would still have to contain possibly hundreds, or in the case of the Internet, tens or thousands of routes. The host routing table would have to change as new subnets are added or removed.
- **Store a single default route that effectively summarizes all of the locations that are not located on the local subnet.** This is possible and practical. Only a single route is needed and does not need to change for nodes or subnets that are added or removed from the network.

By using a default route, the knowledge of the topology of the network and the set of reachable destinations is offloaded to the routers, rather than being a responsibility of the sending host. The advantage to this method is ease of configuration. The disadvantage is that the host can send traffic destined to unreachable addresses. When this happens, however, a router in the path to the destination informs the sending host with an ICMP Destination Unreachable-Host Unreachable message.

The default gateway setting, which creates the default route in the IP routing table, is a critical part of the configuration of a TCP/IP host. The role of the default gateway is to provide the next-hop IP address and interface for all destinations that are not located on its subnet. Without a default gateway, communication with remote destination is not possible, unless additional routes are added to the IP routing table.

Default Gateway Configuration

You can configure a default gateway on a computer running Microsoft Windows XP or Windows Server 2003 in the following ways:

- When obtaining an IP address configuration using DHCP, the default gateway becomes the value of the first IP address in the Router DHCP option, which is configured on the DHCP server to specify an ordered list of one or more default gateways.
- When obtaining an IP address configuration using the user-configured alternate configuration, the default gateway is the IP address specified in the **Default gateway** field on the **Alternate Configuration** tab for the properties of the Internet Protocol (TCP/IP) component in Network Connections. You can specify only a single default gateway.
- When manually specifying an IP address configuration, the default gateway is the IP address typed in the **Default gateway** field on the **General** tab for the properties of the Internet Protocol (TCP/IP). To specify multiple default gateways, you must add them from the **IP Settings** tab in the advanced properties of the Internet Protocol (TCP/IP).

When obtaining the IP address configuration using Automatic Private IP Addressing (APIPA), a default gateway is not configured. APIPA is only useful for a single subnet.

The configuration of a default gateway creates a default route in the IP routing table. The default route has a destination of 0.0.0.0 with a subnet mask of 0.0.0.0. In network prefix notation, the default route is 0.0.0.0/0, which is sometimes abbreviated to 0/0. The next-hop address, also known as the **Gateway** address in the display of the **route print** command, is set to the IP address of the default gateway. The next-hop interface is the interface assigned the IP address in the **Interface** column in the display of the **route print** command.

Based on the route determination process, the default route matches all destinations. If there is no other closer matching route for the destination, the default route is used to determine the next-hop address and interface. Default route traffic is traffic destined to a remote network that is forwarded to the default gateway (rather than traffic destined for the default gateway's IP address).

For more information about the route determination process, see Understanding the IP Routing Table, the December 2001 Cable Guy article.

Default route metric

TCP/IP for Windows XP and Windows Server 2003 by default automatically calculates a metric for the default route that is based on the speed of the adapter to which the default gateway is configured. For example, for a 100 megabit per second (Mbps) Ethernet adapter, the default route metric is set to 20. For a 10 Mbps Ethernet adapter, the default route metric is set to 30.

To override this behavior for DHCP-assigned default gateways, use the Default Router Metric Base Microsoft-specific DHCP option. To override this behavior for manually configured default gateways, clear the **Automatic metric** check box on the **TCP/IP Gateway Address** dialog box for the configured default gateways on the IP Settings tab in the advanced properties of the Internet Protocol (TCP/IP). The **TCP/IP Gateway Address** dialog box is shown in the following figure.

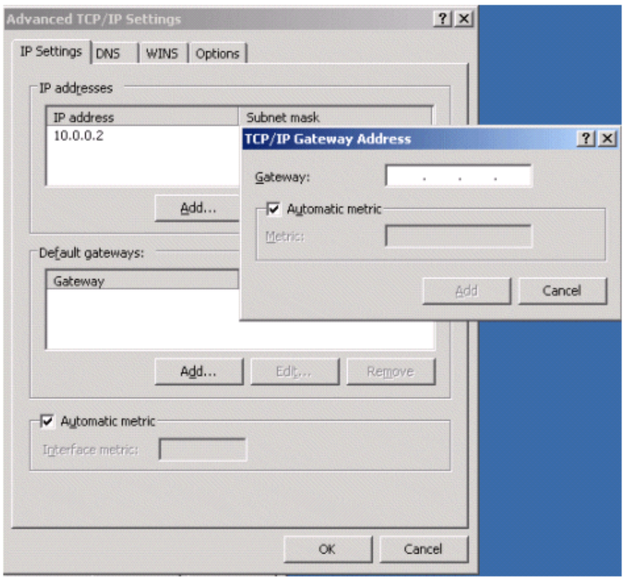

IP address: 10.0.0.2

Gateway: . . .

☑ Automatic metric

☑ Automatic metric

If your browser does not support inline frames, click here to view on a separate page.

Configuring Multiple Gateways

If you have multiple interfaces and you configure a default gateway for each interface, the default metric determination that is based on the speed of the interface causes your fastest interface to be used for default gateway traffic. This might be desirable in some configurations in which the computer is connected to the same network. For example, if you have a 100 Mbps Ethernet adapter and a 10 Mbps Ethernet adapter connected to the same organization intranet, you would want the default gateway traffic to be sent using the 100 Mbps adapter.

However, this default behavior might be a problem when the computer is connected to two or more disjoint networks; networks that do not provide symmetric reachability at the Network layer. Symmetric reachability exists when packets can be sent to and received from an arbitrary destination. For example, the Ping tool tests for symmetric reachability.

Examples of disjoint networks are the following:

- Networks that have no Network layer connectivity, such as an organization intranet and a test lab that has no IP router forwarding packets between them. A computer can be connected to both networks, but if there are no routes to reach both networks and the computer connecting them is not forwarding packets, the two networks are disjoint.
- A privately addressed intranet that has a routed connection to the Internet. In this case, there is asymmetric or one-way reachability; intranet hosts can send packets to Internet hosts from private IP addresses, but the return traffic cannot be delivered because routes for the private address space do not exist in the routing infrastructure of the Internet.

Connectivity to disjoint networks is important when organizations use the following:

- Either a proxy server, such as Microsoft Internet Security and Acceleration (ISA) Server 2000, or a Network Address Translator (NAT) to connect their private intranets to the Internet. In either case, the address space of the intranet is not directly accessible to Internet hosts, regardless of whether the organization is using private or public addressing. Intranet hosts can access Internet locations indirectly through proxy or translation, but Internet hosts cannot access arbitrary intranet locations directly. Hence, there is no symmetric reachability. This is a common configuration for organizations offering Internet connectivity to their employees.
- A virtual private network (VPN) server to allow remote users or remote sites to connect to a private intranet over the Internet. Although the VPN server is connected to both the Internet and a private intranet and is acting as a router, the configuration of packet filters on the Internet interface prevents it from accepting any traffic that is not VPN-based traffic. Internet hosts cannot directly reach intranet locations without an authenticated VPN connection.

Because the TCP/IP protocol only uses a single default route in the routing table at any one time for default route traffic, default gateways configured on multiple interfaces that are connected to disjoint networks can produce undesirable results.

For the examples of the ISA or VPN server, the default route traffic is either forwarded to the Internet or the intranet, but not both. From the ISA or VPN server, either all the locations on the Internet are reachable or all the locations on the intranet are reachable, but not both at the same time. However, IAS or VPN servers require simultaneous symmetric reachability for all the locations on both the Internet and the intranet to operate properly.

When default gateways are configured on multiple interfaces, the default route that is chosen for current use is based on the following:

- When there are multiple default routes in the routing table with the different metrics, TCP/IP for Windows XP and Windows Server 2003 chooses the default route with the lowest metric. If the adapters are of different speeds, then the adapter with the higher speed by default has the lower metric and is used to forward default route traffic.
- When there are multiple default routes in the routing table with the lowest metric, TCP/IP for Windows XP and Windows Server 2003 uses the default route corresponding to the adapter that is the highest in the binding order. You can view and modify the binding order from the **Adapters and Bindings** tab in the **Advanced Settings** dialog box for Network Connections, as shown in the following figure.

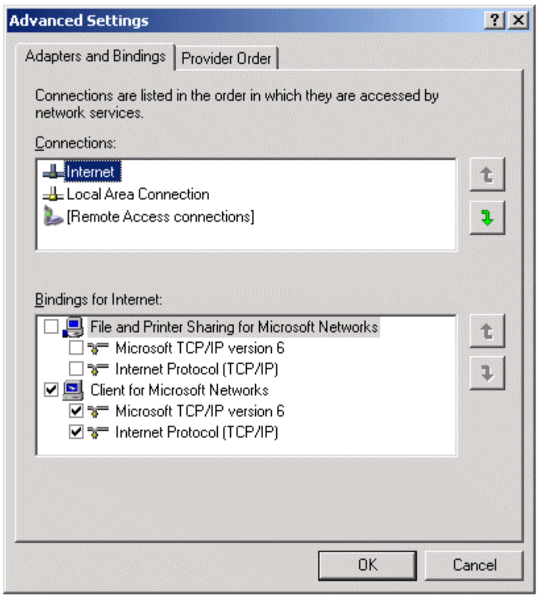

Advanced Settings

Connections are listed in the order in which they are accessed by network services.

Connections:

- Internet
- Local Area Connection
- [Remote Access connections]

Bindings for Internet:

- ☐ File and Printer Sharing for Microsoft Networks
  - ☐ Microsoft TCP/IP version 6
  - ☐ Internet Protocol (TCP/IP)
- ☑ Client for Microsoft Networks
  - ☑ Microsoft TCP/IP version 6
  - ☑ Internet Protocol (TCP/IP)

If your browser does not support inline frames, click here to view on a separate page.

To prevent the problem of disjoint network unreachability, you must do the following on the ISA or VPN server:

- Configure a default gateway on the interface that is connected to the network with the largest number of routes. In most configurations of disjoint networks, the network with the largest number of routes is the Internet.
- Do not configure a default gateway on any other interface. Instead use static routes or dynamic routing protocols to add the routes that summarize the addresses of the other disjoint networks to the local IP routing table.

For example, an ISA server is connected to the Internet and a private intranet. The private intranet uses the private IP address space. To configure this server so that all locations on both disjoint networks are reachable from the ISA server, you would do the following on the ISA server:

- Configure a default gateway on the adapter connected to the Internet. This creates a default route that points to the Internet, making all Internet locations reachable.
- Add the 10.0.0.0/8, 172.16.0.0/12, and 192.168.0.0/16 routes using the intranet-connected adapter as persistent static routes with the Route tool. This creates the routes that summarize all the addresses of the private intranet, making all intranet locations reachable.

In this example, static routes are added. It is also possible to configure the ISA server as a Routing Information Protocol (RIP) or Open Shortest Path First (OSPF) dynamic router so that rather than summarizing the entire private IP address space, subnet-specific routes are dynamically added and removed from the IP routing table based on the current intranet routing topology. To use RIP or OSPF, configure the Routing and Remote Access service.

Dead Gateway Detection

Dead gateway detection is used by the TCP component of Windows TCP/IP to detect the failure of the default gateway and to adjust the IP routing table to use the next default gateway when there are multiple default gateways configured.

When a TCP segment for a TCP connection forwarded via the default gateway is retransmitted three times (by default), dead gateway detection changes the Route Cache Entry (RCE) for that remote IP address to use the next default gateway in the list as its next-hop address. An RCE is an entry in the routing cache, which stores the next-hop IP address for a destination address.

When one fourth of the TCP connections routed through the default gateway have had their RCEs adjusted to the next default gateway, dead gateway detection informs IP to change the computer's default gateway to the one that the adjusted connections are now using. If TCP connections continue to fail, dead gateway detection attempts to use the next default gateway in the list, returning to the first default gateway after cycling through the entire list.

Dead gateway detection monitors only TCP traffic. If connectivity fails for other types of traffic, the default gateway is not switched. Dead gateway detection can cause the default gateway configuration to change when a remote router fails. Remote routers in the path between the host and the destination that fail might also cause TCP connections forwarded along that path to fail and for the host to switch its default gateway. Because dead gateway detection relies on an end-to-end protocol (such as TCP), a host can switch its default gateway even when the current default gateway is fully operational.

For More Information

For more information about this topic, consult the following resources:

- Microsoft Windows Server 2003 TCP/IP Implementation Details
- Understanding the IP Routing Table (December 2001 Cable Guy article)

For a list of all **The Cable Guy** articles, click here.

Did you find this helpful? ○ Yes ○ No