The SpeedTouch™ and Virtual Private Networks

Author: Jan Wuyts

Date: March 2006

Version: v1.0

Abstract: This application note provides technical information on Virtual Private Networks (VPN) and how this relates to the SpeedTouch™. In the introduction a brief background on VPNs is presented. The subsequent sections give detailed information on user scenarios and how to configure a SpeedTouch™. The last section explains services that can be enabled on a VPN.

Applicability: This application note applies to

- All SpeedTouch™ (Wireless) Business DSL Routers Release R5.4 and higher.

Updates: THOMSON continuously develops new solutions, but is also committed to improve its existing products.

For more information on THOMSON's latest technological innovations, documents and software releases, visit us at:

http://www.speedtouch.com

CONTENTS

1 INTRODUCTION

Introduction

Most often the use of DSL-based routers is limited to well-known Internet connectivity. However, more and more we see new implementations popping up. One of these implementations is the use of DSL connectivity for implementing Virtual Private Networking.

The Virtual Private Network (VPN) allows different trusted sites to be interconnected transparently through an untrusted provider network.

In this document

This document covers:

2 INTRODUCTION TO VIRTUAL PRIVATE NETWORKS

What is a VPN

A VPN is a network to interconnect different related sites on top of an existing network. The sites are trusted, as only trusted people and authenticated devices have access to these parts of the network. The underlying network is untrusted, as people who are not related and unauthenticated devices have access to it.

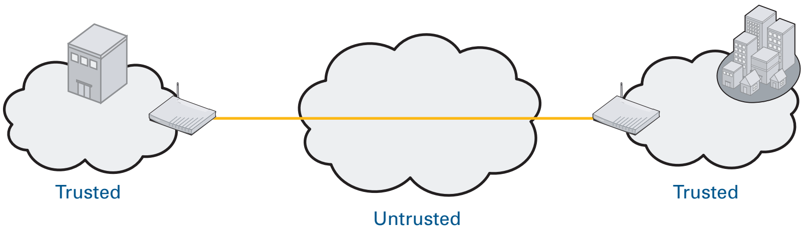

Kinds of VPNs

Basically, 2 kinds of VPNs exist:

- VPNs offered by a service provider.

  This type of network is created by means of switching techniques. Some examples of switching techniques are Asynchronous Transfer Mode (ATM), Frame Relay, Multi-Protocol Label Switching (MPLS), ...

  Basically, it is sufficient to set up a connection to the provider. The technology used within the provider's transport network is irrelevant for our VPN. The IP addresses that are used within the VPN are unknown and unreachable to the untrusted network.

- Self-made VPNs over a public network.

  This type of network is created by means of tunnelling techniques. Some examples of tunnelling techniques are IPSec, Point-to-Point Tunnelling Protocol (PPTP), Layer 2 Tunnelling Protocol (L2TP), ... An additional tunnel is placed on top of the Internet connectivity. In case of L2TP the tunnel is created on a Layer 2 protocol, p.e. Ethernet. VPN traffic is encrypted before it is sent to the untrusted network. The trusted sites share a key to decrypt the VPN traffic.

2.1 VPN Topologies

Introduction

Depending on the number of sites and how these sites are interconnected VPNs can exist in many different topologies. The preferred topology is determined by:

- The technology used to create the VPN.
- The company policy.

  For example, IPSec supports full mesh topology, but company policy may prefer a hub-and-spoke approach.

Following topologies will be discussed:

- " Point-to-Point" on page 5
- " Hub-and-Spoke" on page 6
- " Star" on page 7
- " Mesh" on page 8

Point-to-Point

The most simple VPN topology to achieve is a point-to-point connection between two sites. Commonly service providers will call this a "leased line" service.

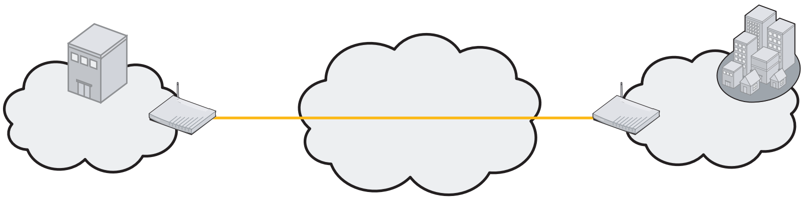

Typical VPN technologies to achieve a point-to-point VPN are ATM and Frame Relay.

Hub-and-Spoke

A hub-and-spoke VPN can be achieved by making a number of point-to-point connections from branch offices to a central office. The hub commonly has a specific function in the VPN, such as VPN access concentration.

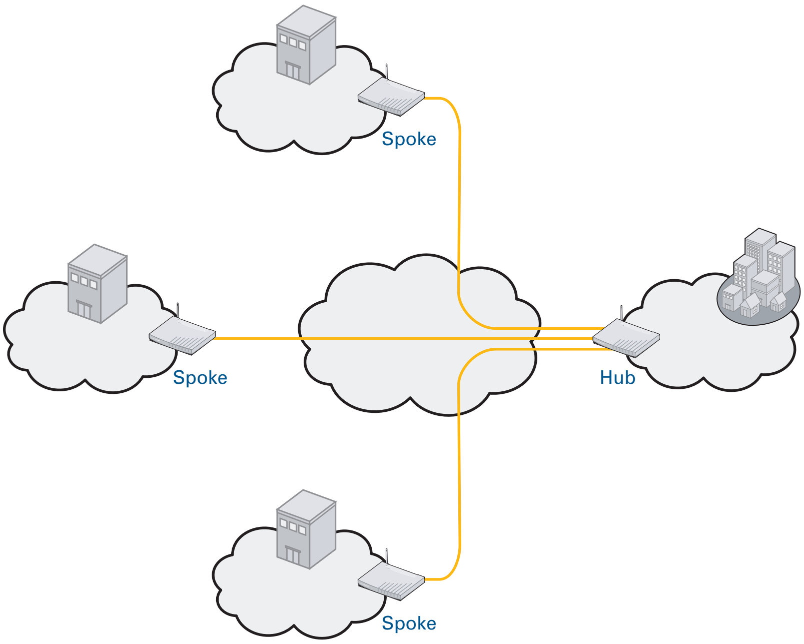

The hub-and-spoke topology is typically based on Frame Relay. This technology requires that all traffic passes a central node, that is part of the VPN, to relay frames to their destination.

Another use of a hub and spoke VPN is a company consisting of:

- A central office hosting all servers, a firewall and an internet gateway.
- Branch offices only have clients that need connectivity to the servers in the central office. Company policy dictates that all internet traffic has to go via the firewall in the central office.

Star

The star topology is achieved when each site has 1 connection to the provider network. It is the provider's task to make sure all the sites of the VPN are interconnected.

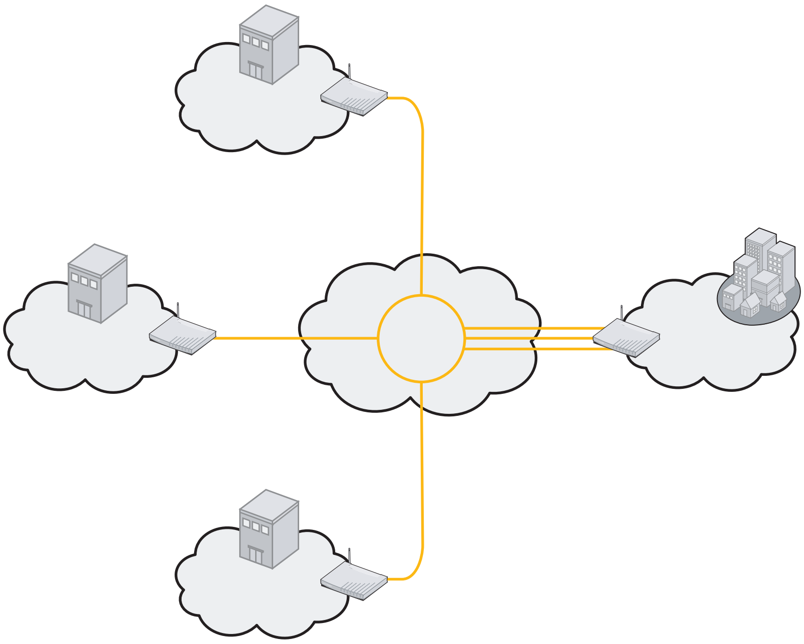

A typical technology to achieve the star topology is MPLS.

Mesh

In a fully meshed VPN, each site has a direct point-to-point connection to all other sites.

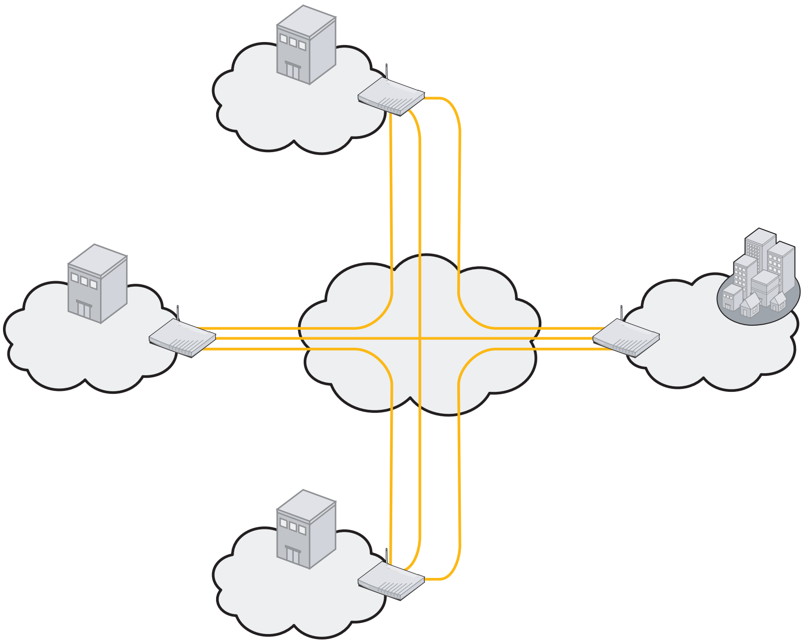

As leased lines are usually very expensive, a fully meshed VPN will be based on IPSec tunnels.

It is possible to have a partially meshed VPN. In this case, only the heavily used connections will be provisioned by a point-to-point connection in addition to a backbone, interconnecting all the sites to the central office.

2.2 VPN Technologies

Introduction

To set up a VPN, an underlying physical network is needed. We distinguish between

- ***provider provisioned VPNs*** where the physical network is tailored to the VPN needs of the customer
- ***self-made VPNs*** where the customer uses an existing IP network to establish secure communication.

Typical technologies in provider provisioned VPNs are:

- ATM
- Frame Relay
- MPLS

Self-made VPNs are commonly based on IPSec. An IPSec tunnel can be configured on top of:

- The Internet
- A Provider provisioned VPN

2.2.1 Provider Provisioned VPN

ATM

ATM is a communication method based on connections. A number of virtual connections can coexist on a link between ATM switches. By linking virtual connections in the switches, it is possible to establish end to end virtual connections through the ATM network.

In general, ATM cells are used as a transport method for higher layer protocols such as IP. As such, edge nodes encapsulate IP packets in ATM cells for transmission and decapsulate ATM cells to retrieve the original IP packets.

Frame Relay

Frame relay is a communication method based on packet switching. Packets are switched based on a small Frame Relay header.

All packets are sent to a central node that decides where to forward the packet. Its decision is based solely on the Frame Relay header.

MPLS

MPLS is a high performance method for forwarding packets. Routers at the edge of the MPLS network add very simple labels (also known as tags) to the packets. MPLS core switches can perform forwarding with very limited overhead. Additionally, MPLS has extended traffic engineering features which enable extended Class of Service (CoS) capabilities.

An MPLS network logically consists of 3 components:

- Provider core switches: The main task of these devices is forwarding packets
- Provider edge routers: They define the edge of the MPLS cloud. The main task of these devices is adding correct tags to the packets.
- Customer edge routers: These are the routers in the customer premises, that are linked to the MPLS cloud. These devices are not MPLS aware. This is where the SpeedTouch™ is placed.

2.2.2 Self-made VPN

IPSec tunnel over the public Internet

An IPSec tunnel is set up on top of the public Internet. The purpose is to share a secret key between the end points of the tunnel to encrypt and decrypt data. Untrusted nodes can not join the communication stream through the IPSec tunnel.

The SpeedTouch™608 and SpeedTouch™620 (Wireless) Business DSL Routers feature an embedded VPN IPSec software module, offering VPN IPSec capabilities.

For more information see the “SpeedTouch™ IPSec Quick Start Guide”.

IPSec tunnel on Top of Provider Provisioned VPN

Although a provider provisioned VPN ensures interconnectivity, privacy and security, it is possible to add IPSec tunnels on top of the VPN. This is important for those customers who want to add a higher level of security that they can control themselves.

PPTP/L2TP

PPTP and L2TP are tunnelling protocols that are based on PPP. While an IPSec tunnel features device authentication, PPTP and L2TP feature user authentication. PPTP and L2TP are implemented mainly by Microsoft to achieve a secure VPN connection from a Windows PC to a central office site. This will automatically lead to a hub-and-spoke topology.

3 THE SPEEDTOUCH™ IN A VPN

The SpeedTouch™ as Customer Edge Router in a Provider Provisioned VPN

Suppose we have to connect a branch office to the central office of our company, where the main servers, such as file server, mail server and Domain Name Server (DNS) are located. The network provider has given us access to do so by means of a VPN.

The VPN can be based on one of the following technologies:

- ATM

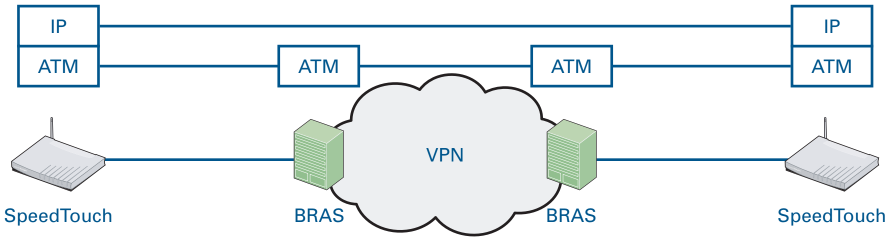

- Frame Relay

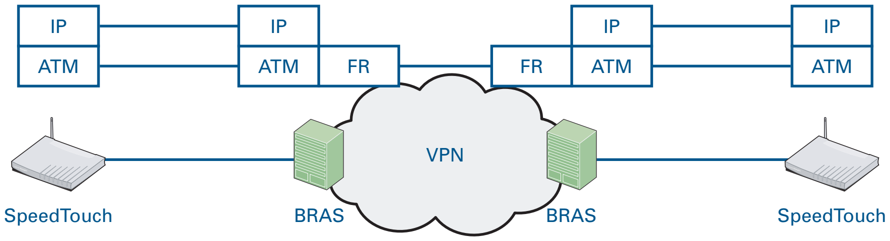

- MPLS

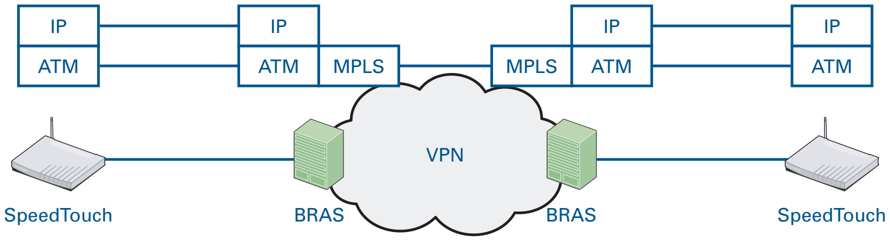

The technology used within the provider provisioned VPN is irrelevant to the SpeedTouch™.

To enable communication between the computers in the branch office and the servers in the main office, we have to follow the following procedure:

1. Create an IP over ATM (IPoA) interface on the SpeedTouch™. Another type of connection is also possible.

   In case the VPN is based on ATM technology the IPoA interface will be a point-to-point connection to the gateway of the main office. In case of MPLS the IPoA interface will be a point-to-point connection to the provider edge router of the MPLS network. The SpeedTouch™ does not need to be aware of the VPN technology.

2. Add a private IP address to the Local Area Network (LAN) interface, to be used as VPN gateway.

   All devices on the LAN will receive private IP addresses within the range of the VPN gateway, configured manually or via Dynamic Host Configuration Protocol (DHCP). Every site needs a number of unique IP addresses.

3 Configure Routing Information Protocol (RIP) on the SpeedTouch™ to ensure routes are exchanged between the main office and the branch office.

If both the main office and the branch office have Internet connectivity, the default route may be omitted from the RIP messages.

The SpeedTouch™ End-Point of an IPSec Tunnel

Suppose both the branch office and the main office of our company have Internet connectivity and we want to exchange confidential information. Using the default connection, we can communicate, but since the information is not encrypted, anyone can join in the communication stream.

To ensure confidential information is kept confidential, the information needs to be encrypted. This can be done by setting up an IPSec tunnel between the gateway in the main office and the gateway in the branch office. An IPSec tunnel can be constructed on top of

- The public Internet

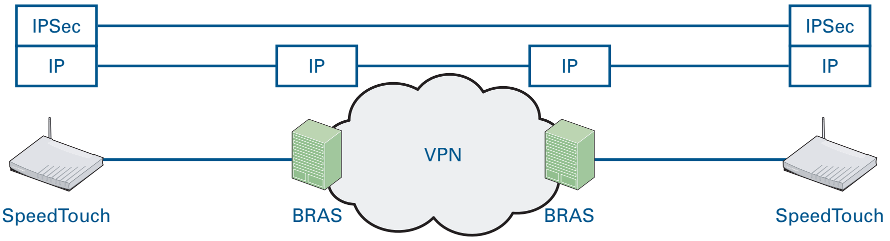

- A provider provisioned VPN

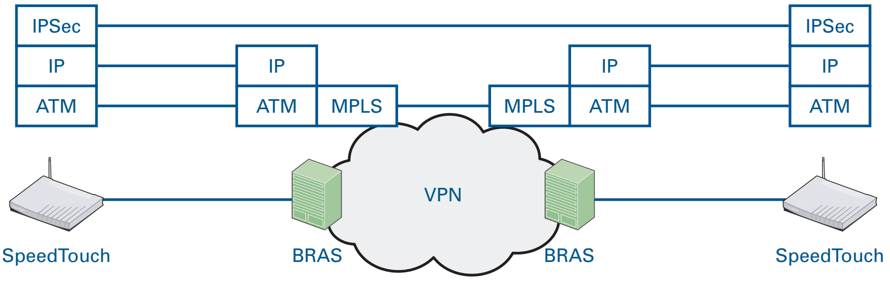

The SpeedTouch™608 and SpeedTouch™620 (Wireless) Business DSL Routers feature an embedded VPN IPSec software module. For more information, see the “SpeedTouch™ IPSec Quick Start Guide” and “SpeedTouch™ IPSec Configuration Guide”.

4 CONFIGURING THE SPEEDTOUCH™ AS VPN ROUTER

Introduction

In this chapter, all features of the SpeedTouch™ necessary to configure a VPN are addressed. As an illustration, the following network example is created.

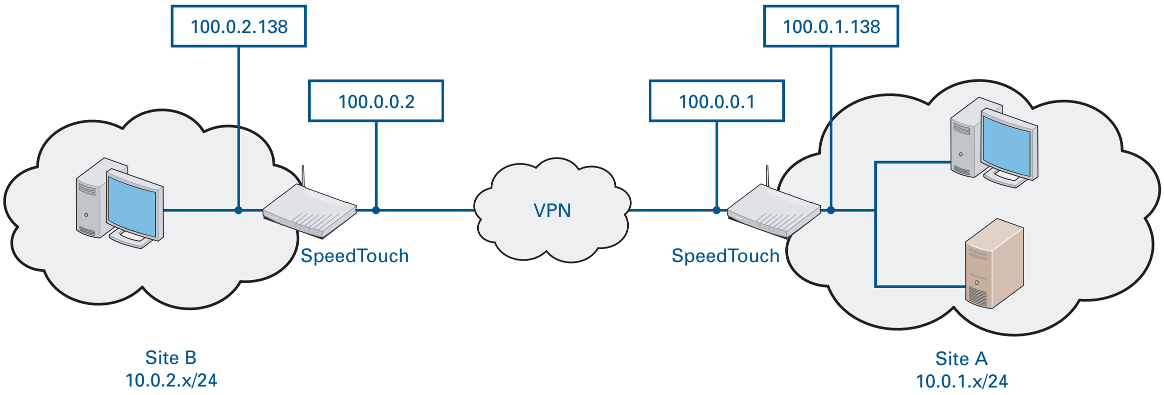

This example will serve in “5 Services Running on Top of a VPN” on page 21 as a basis to support services or extra features.

General procedure

The SpeedTouch™ Business DSL Routers are equipped with all the features to connect to a VPN. When configuring the SpeedTouch™, the following aspects have to be taken into account:

- “4.1 Local Addressing” on page 14
- “4.2 VPN Connection” on page 15
- “4.3 Routing over a VPN” on page 17

4.1 Local Addressing

Introduction

All devices used in the VPN need a unique IP address. To make sure that all IP addresses are unique, it is advised that each site has its own subnet.

As shown below the following subnets are defined:

- Hosts on site A have an IP address in subnet 10.0.1.x/24.
- Hosts on site B have an IP address in subnet 10.0.2.x/24.

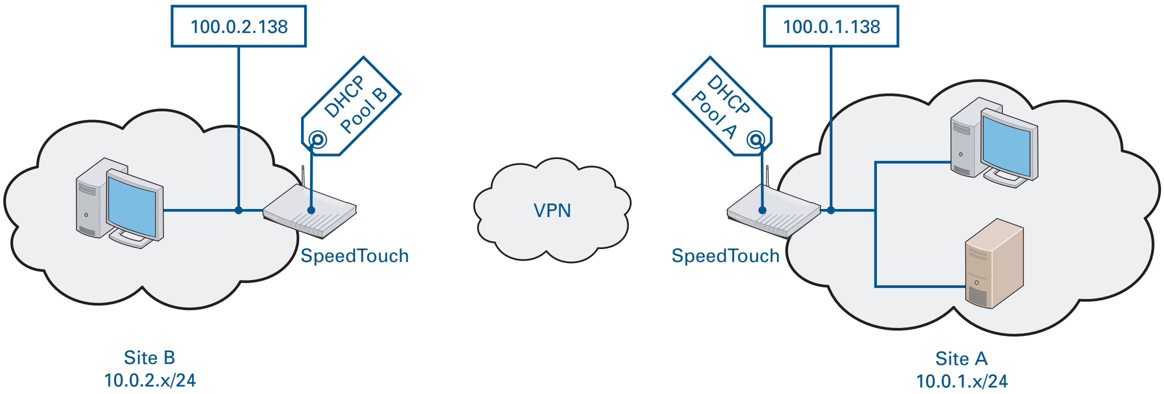

Addresses can be assigned

- Manually.
- Automatically by DHCP.

Local gateway

To make sure the hosts on site A will be able to communicate with the hosts on site B once the VPN is set up, the SpeedTouch™ needs to be configured as a gateway for the defined subnets.

Execute the following CLI command on the SpeedTouch™ at site A to configure its gateway address:

```
{ST-A} :ip ipadd intf=eth0 addr=10.0.1.138 netmask=24
```

Execute the following CLI command on the SpeedTouch™ at site B to configure its gateway address:

```
{ST-B} :ip ipadd intf=eth0 addr=10.0.2.138 netmask=24
```

To make the distinction between CLI commands on the different SpeedTouch™ Business DSL Routers, the CLI commands on the SpeedTouch™ at site A is preceded by **`{ST-A}`** and the CLI commands on the SpeedTouch™ at site B is preceded by **`{ST-B}`**.

DHCP server

The SpeedTouch™ DSL Routers feature an embedded DHCP server, offering the capabilities to distribute IP addresses on the local network.

Execute the following CLI commands on the SpeedTouch™ at site A to configure its DHCP server:

```
{ST-A} :dhcp server lease flush
{ST-A} :dhcp server pool flush
{ST-A} :dhcp server pool add name=poolA
{ST-A} :dhcp server pool config name=poolA intf=eth0 poolstart=10.0.1.1 poolend=10.0.1.64
netmask=24 gateway=10.0.1.138
```

Execute the following CLI commands on the SpeedTouch™ at site B to configure its DHCP server:

```
{ST-B} :dhcp server lease flush
{ST-B} :dhcp server pool flush
{ST-B} :dhcp server pool add name=poolB
{STB} :dhcp server pool config name=poolB intf=eth0 poolstart=10.0.2.1 poolend=10.0.2.64
netmask=24 gateway=10.0.2.138
```

The default configuration of the SpeedTouch™ features a default DHCP pool. This pool is not unique throughout the VPN we are creating. Therefore, we need to remove the default address pool from the DHCP server. This is why we need to execute the two flush commands before configuring the DHCP pool.

4.2 VPN Connection

Introduction

The connections on the SpeedTouch™ are always ATM based. Although the VPN can be based on other technologies like Frame Relay or MPLS, the connection to the VPN will always be ATM based. The most basic connection on top of ATM is the IP over ATM (IPoA) connection.

- In case of a VPN based on MPLS or Frame Relay the IPoA connection is terminated at the VPN access node.
- In case of an ATM VPN the connection is terminated by the SpeedTouch™ at the other site.

We will focus on an ATM VPN connection from the SpeedTouch™ at site A to the SpeedTouch™ at site B as shown below:

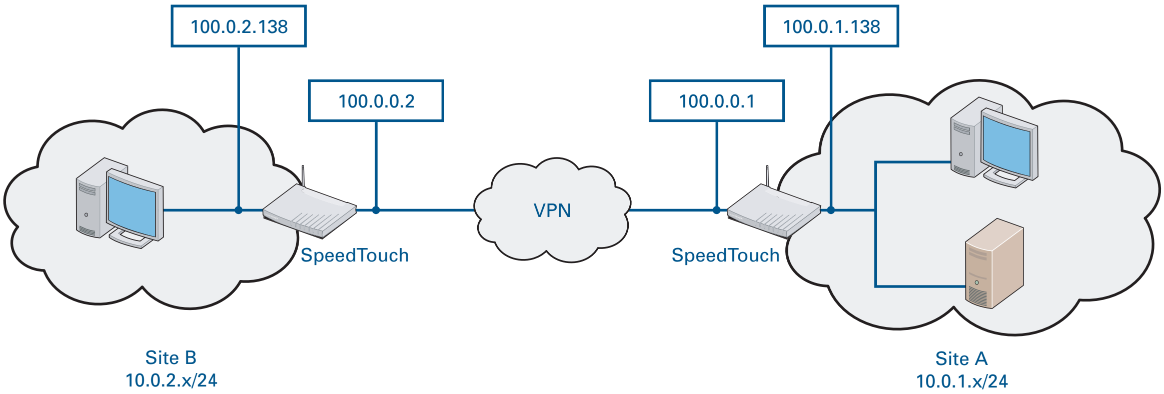

Creating an IP over ATM interface at both sites

Proceed as follows to create an IPoA interface on the SpeedTouch™ at site A:

1 Create an ATM phonebook entry. The assumption is made that a connection is available on VPI/VCI 8/50 for the VPN interface.

```
{ST-A} :atm phonebook add name=pvc_8_50 addr=8.50
```

2 Create an ATM interface. The assumption is made that ATM encapsulation method VCMUX is used.

```
{ST-A} :atm ifadd inft=atm_vpn
{ST-A} :atm ifconfig intf=atm_vpn dest=pvc_8_50 encaps=vcmux ulp=ip
{ST-A} :atm ifattach intf=atm_vpn
```

3 Create an IP interface. The assumption is made that the IP address of the peer is 100.0.0.2.

```
{ST-A} :ip ifadd intf=ip_vpn_A
{ST-A} :ip ifconfig intf=ip_vpn_A dest=atm_vpn
{ST-A} :ip ipadd intf=ip_vpn_A addr=100.0.0.1 pointopoint=100.0.0.2 addroute=enabled
{ST-A} :ip ifattach intf=ip_vpn_A
```

Follow the same procedure to configure the IPoA interface on the SpeedTouch™ at site B:

```
{ST-B} :atm phonebook add name=pvc_8_50 addr=8.50
{ST-B} :atm ifadd inft=atm_vpn
{ST-B} :atm ifconfig intf=atm_vpn dest=pvc_8_50 encaps=vcmux ulp=ip
{ST-B} :atm ifattach intf=atm_vpn
{ST-B} :ip ifadd intf=ip_vpn_B
{ST-B} :ip ifconfig intf=ip_vpn_B dest=atm_vpn
{ST-B} :ip ipadd intf=ip_vpn_B addr=100.0.0.2 pointopoint=100.0.0.1 addroute=enabled
{ST-B} :ip ifattach intf=ip_vpn_B
```

The settings of the IPoA interface on the SpeedTouch™ at site B correspond to the settings of the SpeedTouch™ at site A. If we connect to an access node of a provided provisioned VPN, you need the following data from the provider:

- ATM PVC to use
- Encapsulation method
- IP address reserved for the SpeedTouch™
- IP address of the VPN access node

4.3 Routing over a VPN

Introduction

To allow communication between hosts at site A and hosts at site B, it is necessary to populate the routing tables of the gateways at these sites. These routing tables can be populated

- " Manual configuration of the IP routing table" on page 17
- " Automatic update of the IP routing table via RIP in the SpeedTouch™608 WL" on page 17

Manual configuration of the IP routing table

To configure the routing table manually, you have to add all the routes as required by your configuration. Execute the following CLI commands on the SpeedTouch™ DSL Routers at sites A and B to populate the routing tables for the example VPN:

```
{ST-A} :ip rtadd dst=100.0.0.2/32 intf=ip_vpn_A
{ST-A} :ip rtadd dst=10.0.2.0/24 gateway=100.0.0.2
{ST-B} :ip rtadd dst=100.0.0.1/32 intf=ip_vpn_B
{ST-B} :ip rtadd dst=10.0.1.0/24 gateway=100.0.0.1
```

The configured IP routes can be consulted by means of the `:ip rtlist` CLI command.

Automatic update of the IP routing table via RIP in the SpeedTouch™608 WL

Routing Information Protocol (RIP) is used between routers to communicate which networks can be reached. A Router that receives a RIP message will update its routing table if the destination network is not yet known or if the new route has a better metric. In the SpeedTouch™608 WL by default the firewall is configured to allow incoming RIP messages. The SpeedTouch™608 WL participates in the RIP protocol in a passive mode, i.e. it receives and processes RIP messages from other routes, but it does not progagate routing information.

Proceed as follows to enable RIP on the SpeedTouch™608 WL:

1 Configure RIP

Execute the following CLI command to enable the RIP engine on the SpeedTouch™ at site A. The RIP version depends on the actual situation. In this scenario RIP version 2 is used.

```
{ST-A} :grp rip config state=enabled version=rip_v2
```

2 Enable RIP on an interface

Execute the following CLI command to enable the reception of RIP messages on the ip_vpn interface. To avoid misconfiguration, you can force the RIP version on this specific interface, overruling the global RIP settings.

```
{ST-A} :grp rip ifconfig intf=ip_vpn_A rip=enabled rxversion=rip_v2
```

Execute the following CLI command to see the status of the RIP engine of the SpeedTouch™ at site A, as well as the routes that are currently stored in the routing table:

```
{ST-A} :grp rip show
```

Below is the generated output of site A:

- The first part of the output is the RIP engine status.

```
 RIP routing protocol config dump
------------------------------------
        RIP daemon is enabled
        Global RIP queries received : 0
        Global RIP route changes : 25
        Default version : send rip_v2, receive rip_v2
        Default redistribution metric is 1
        Sending routing table updates every 30 seconds with +/-5%
        Route timeout after 180 seconds
        Route garbage collect after 120 seconds
        Import of connected routes is enabled
        Import of kernel routes is enabled
        Import of static routes is enabled
        Import of default kernel route is enabled
        Import of default static route is enabled
        Export of RIP routes is enabled
        Export of default RIP route is enabled
        Transmission of default RIP route is enabled
```

- The second part of the output consists of the RIP routes.

```
RIP routing table dump
--------------------------
        Codes : K - Kernel, C - connected, S - Static, R - RIP, * - FIB route
   Network            Next Hop         Metric  From            Flags
----------------------------------------------------------------------
C  10.0.0.0/24                         1                       <>  *
R  169.254.0.0/16     100.0.0.2        2       100.0.0.2       <>  *
R  172.16.1.0/24      100.0.0.2        2       100.0.0.2       <>  *
C  192.168.1.0/24                      1                       <>  *
C  100.0.0.1/32                        1                       <>  *
K  100.0.0.2/32       100.0.0.1        1                       <>  *
C  10.0.1.0/24                         1                       <>  *
R  10.0.2.0/24        100.0.0.2        2       100.0.0.2       <>  *
```

Automatic update of the IP routing table via RIP in the SpeedTouch™620

Routing Information Protocol (RIP) is used between routers to communicate which networks can be reached. A Router that receives a RIP message will update its routing table if the destination network is not yet known or if the new route has a better metric. In addition it can propagate routing information to its neighbours. The SpeedTouch™620 is capable to receive and transmit routing information. By default the firewall is configured to allow both incoming and outgoing RIP messages.

Proceed as follows to enable RIP on the SpeedTouch™620:

1 Configure RIP

Execute the following CLI command to enable the RIP engine on the SpeedTouch™ at site A. The RIP version depends on the actual situation. In this scenario RIP version 2 is used.

```
{ST-A} :router rip config state=enabled
```

2 Enable RIP on an interface

Execute the following CLI command to enable the reception of RIP messages on the ip_vpn interface. To avoid misconfiguration, you can force the RIP version on this specific interface, overruling the global RIP settings.

```
{ST-A} :router rip ifconfig intf=ip_vpn_A ripin=enabled ripout=enabled version=v2
```

Execute the following CLI command to see the status of the RIP engine of the SpeedTouch™ at site A, as well as the routes that are currently stored in the routing table:

```
{ST-A} :router rip list
```

Below is the generated output of site A:

```
 RIP daemon specific information
  -------------------------------
        RIP admin state                enabled
        Update Time                    30 +/- Random(1..5)s
        Expire Time                    180 s
        Garbage Time                   120 s
        Flash Update Time              Random(1..5)s
        Split Horizon                  simple
        Default Metric                 1
        Recognition RIPv2 Host Routes  disabled
        ECMP                           disabled
        RIPV1 MBZ Check                enabled
        # RIP Routes per Destination   1
        # Updates on Terminate         4

                     Interface Admin   Version   Recv  Send  Metin   Metout  PrimAuth  SecAuth
--------------------------------------------------------------------------------------------
---
                   ISDN_backup down    v2        on    on    1       0       none      none
           ISDN_backup_trigger down    v2        on    on    1       0       none      none
                      Internet up      v2        on    on    1       0       none      none
                          dmz1 down    v2        on    on    1       0       none      none
                        guest1 down    v2        on    on    1       0       none      none
                          lan1 down    v2        on    on    1       0       none      none
                          wan1 down    v2        on    on    1       0       none      none

  RIP unicast neighbor list
  -------------------------
        No unicast neighbors configured

  RIP trusted neighbor list
  -------------------------
        All onlink routers are trusted
```

To see the routes that are currently stored in the RIP routing table, execute the following command:

```
{ST-A} :router rip rtlist
```

Using RIP filters in the SpeedTouch™620

In case there are routes or destinations that are only available to the local site, make use of the feature called 'route filtering'. This functionality allows to define specific routes that must be filtered out of the routing protocol communication. For example, if subnet 172.16.1.0/24 at site B is a local network that is only available for site B, proceed as follows to prevent propagation of this destination to site A:

1 Create an access list:

```
{ST-B} :router policy accesslist add name=dontprop
```

2 Create a rule in the access list that denies the propagation of the route 172.16.1.0/24:

```
{ST-B} :router policy accesslist addrule name=dontprop seqnbr=10 action=deny
addr=172.16.1.0/24
```

3 Use the access list in the RIP policy:

```
{ST-B} :router rip policy distlistout intf=ip_vpn_B protocol=any accesslist=dontprop
```

If you combine VPN connectivity with Internet access in a single SpeedTouch™620, you want to prevent that the default route to the Internet is advertized inside the VPN community. You can achieve this with the VPN filtering as described above. Make sure to use exactly the route 0.0.0.0/0, and do not make use of the parameters ge (greater or equal) and le (less than or equal) for the netmask. Otherwise, you will filter out more networks than you intend to!

5 SERVICES RUNNING ON TOP OF A VPN

Introduction

Basically all relevant services for business communication can run on top of a VPN. In this section some of these services are highlighted.

5.1 DHCP Relay

Introduction

As mentioned in “4 Configuring the SpeedTouch™ as VPN Router” on page 13 IP addresses of the LAN clients can be configured manually or dynamically. Dynamic configuration will make use of a DHCP server. The SpeedTouch™ can be configured as a:

- “ DHCP relay to an external DHCP server” on page 21
- “ DHCP Server relaying DHCP requests to a SpeedTouch™ at another site” on page 22

DHCP relay to an external DHCP server

If an external DHCP server is used, all SpeedTouch™ devices are configured as DHCP relay:

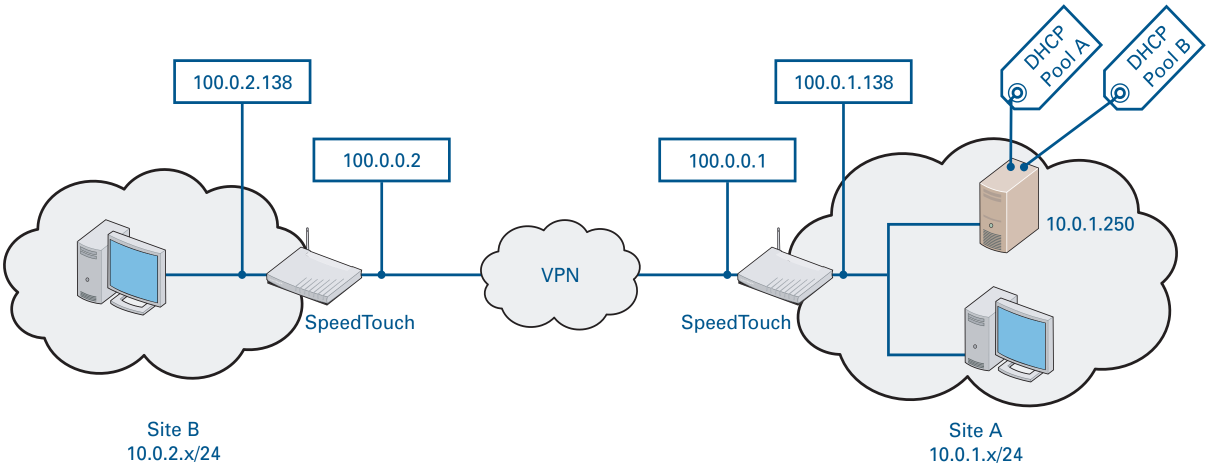

Using an External DHCP Server

The SpeedTouch™ DHCP service is by design connected to the DHCP relay agent. By default, the DHCP server is enabled to distribute IP addresses on the local network. If an external DHCP server is used, the internal DHCP server has to be disabled.

Execute the following CLI command to disable the internal DHCP server of the SpeedTouch™ at site A:

```
{ST-A} :dhcp server config state=disabled
```

Execute the following CLI command on the SpeedTouch™ at site B to relay DHCP requests to the DHCP server at site A:

```
{ST-B} :dhcp server config state=disabled
{ST-B} :dhcp relay add addr=10.0.1.250 intf=eth0 giaddr=10.0.2.138
```

The DHCP server has to support the giaddr feature to select the pool based on the gateway IP address.

DHCP Server relaying DHCP requests to a SpeedTouch™ at another site

The SpeedTouch™ at site A can act as DHCP server:

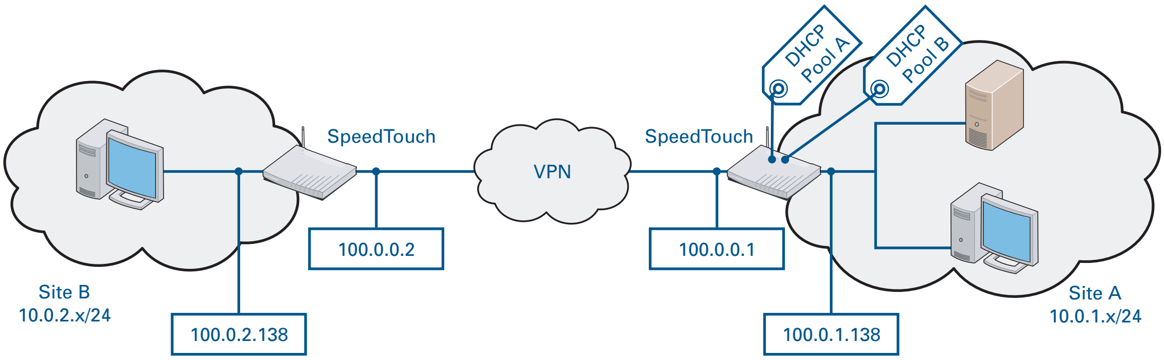

By default only packets coming from the LAN are relayed to the internal DHCP server.

Execute the following CLI commands to enable relaying of DHCP request coming from the ip_vpn_A interface:

```
{ST-A} :dhcp relay ifconfig intf=ip_vpn_A relay=on
{ST-A} :dhcp relay add addr=127.0.0.1 intf=ip_vpn_A
```

Execute the following CLI commands to relay local DHCP requests at site B to the SpeedTouch™ at the site A:

```
{ST-B} :dhcp server config state disabled
{ST-B} :dhcp relay add addr=100.0.0.1 intf=eth0 giaddr=10.0.2.138
```

The assumption is made that the SpeedTouch™ at site A is the DHCP server. The internal DHCP server is configured with a local address pool (**`poolA`**) and a remote address pool (**`poolB`**).

Execute the following CLI commands to configure the SpeedTouch™ at site A with a local address pool and a remote address pool:

```
{ST-A} :dhcp server lease flush
{ST-A} :dhcp server pool flush
{ST-A} :dhcp server pool add name=poolA
{ST-A} :dhcp server pool config name=poolA intf=eth0 poolstart=10.0.1.1 poolend=10.0.1.64
netmask=24 gateway=10.0.1.138
{ST-A} :dhcp server pool add name=poolB
{ST-A} :dhcp server pool config name=poolB intf=ip_vpn poolstart=10.0.2.1 poolend=10.0.2.64
netmask=24 gateway=10.0.2.138
```

5.2 Internet connectivity

Company policy to provide Internet connectivity

In the following sections, two different approaches are described to provide Internet access to the VPN members. Both approaches have their pros and cons. The selection depends on the company policy.

! The advertisement of routes via the RIP protocol is possible only with the SpeedTouch™620. So, this chapter is applicable to the SpeedTouch™620 only.

One Site has Internet Connectivity

In this approach, Internet connectivity for the VPN members is provided via a single site, most probably the central headquarters of the company. The rationale for this strategy may be the use of a central firewall for the whole company. As a drawback, all Internet traffic passes via the central node of the VPN, which puts a higher traffic load on this node,

In this network topology, the default route to the Internet has to be distributed to the other sites:

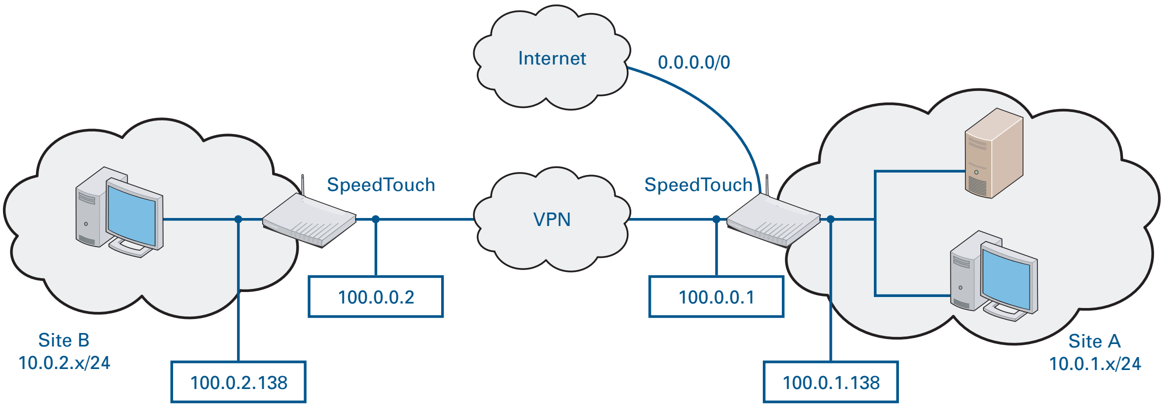

This can be achieved by configuring the RIP engine accordingly. By default, the default route to the Internet is communicated by RIP.

If in your current configuration the communication of the default route is disabled, you should enable it. Execute the following CLI commands to enable the advertisement of the default route via RIP:

```
{ST-A} :router policy accesslist add name=defrtprop
{ST-A} :router policy accesslist addrule name=defrtprop seqnbr=1 action=permit addr=0.0.0.0/0
{ST-A} :router rip policy distlistout intf=ip_vpn_A protocol=any accesslist=defrtprop
```

All Sites have Internet Connectivity

A company may decide to provide local Internet access at every VPN site. The advantage of this strategy is to reduce the traffic that passes via the central node.

In this network topology, each site has its own default route to the Internet, and we do not want to distribute the default route via RIP:

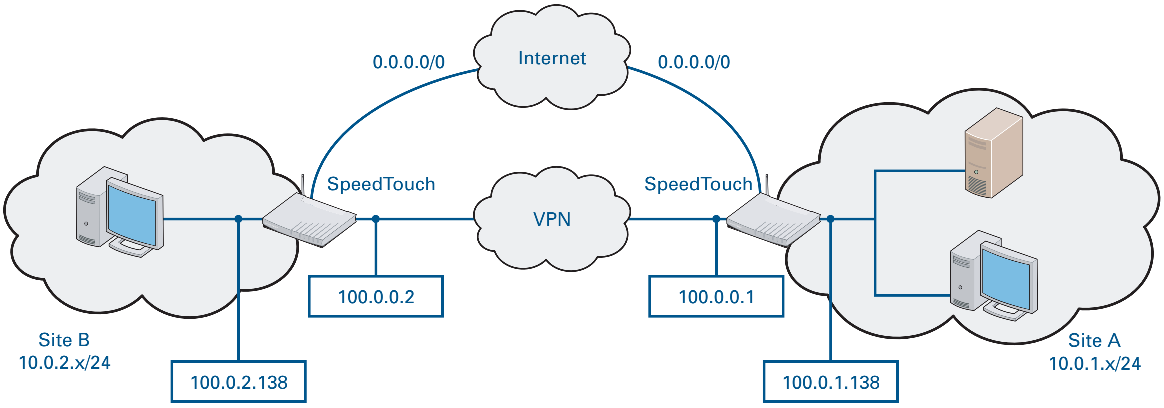

In each of SpeedTouch™620, execute the following CLI commands to avoid communication of the default route to the other VPN sites via RIP:

```
{ST-A} :router policy accesslist add name=defrtdontprop
{ST-A} :router policy accesslist addrule name=defrtdontprop seqnbr=1 action=deny
addr=0.0.0.0/0
{ST-A} :router rip policy distlistout intf=ip_vpn_A protocol=any accesslist=defrtprop

{ST-B} :router policy accesslist add name=defrtdontprop
{ST-B} :router policy accesslist addrule name=defrtdontprop seqnbr=1 action=deny
addr=0.0.0.0/0
{ST-B} :router rip policy distlistout intf=ip_vpn_A protocol=any accesslist=defrtprop
```

If an IPSec tunnel is used to set up the VPN, local Internet connectivity is available by default at all sites. However, you can force Internet access through the main site (site A). In this case, transmission of the default route has to be enabled at the main site (site A), but disabled at the remote site (site B).

5.3 Domain Name Server (DNS)

Introduction

The DNS service allows you to contact a server using an understandable name instead of an IP address. The client contacts the locally configured DNS server to request the IP address for this name.

Usually, the DNS service and the DHCP service run on the same machine. So, if the SpeedTouch™ acts as DHCP server, by default it announces its own IP address as primary DNS server

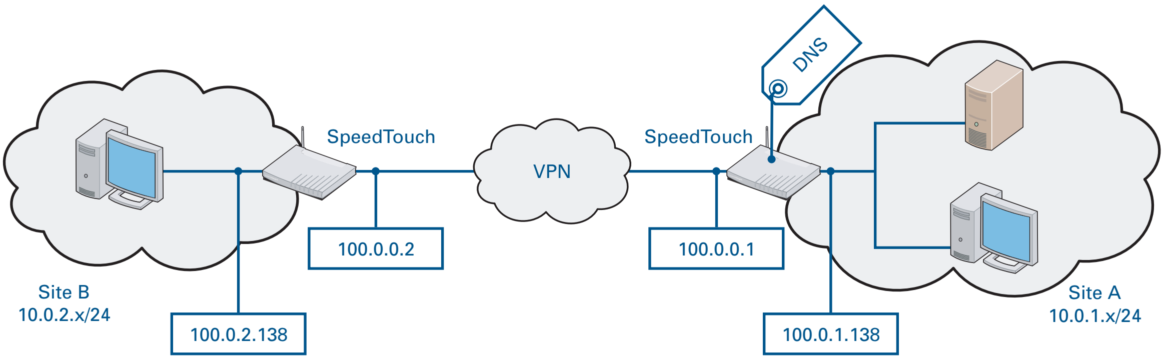

In case the SpeedTouch™ relays all DHCP requests to an external server, an external server will act as primary DNS server as well:

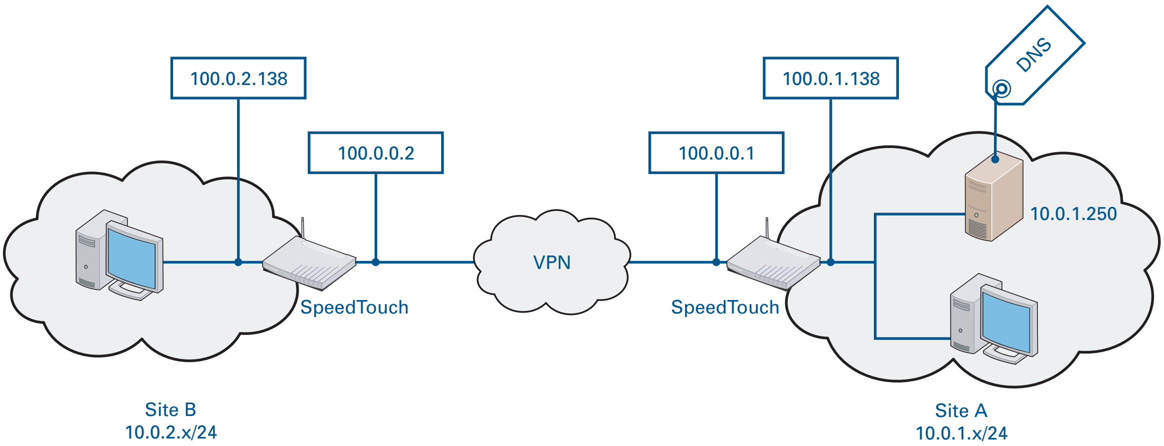

Enabling DNS Service on WAN

By default the SpeedTouch™ only responds to DNS requests from the local network. In order to allow the SpeedTouch™ to respond to DNS requests from the WAN, the DNS service has to be adapted.

Execute the following CLI command to enable the DNS service on the WAN interface of the SpeedTouch™ at site A:

```
{ST-A} :service system ifadd name=DNS-S group=wan
```

DNS Server Configuration

The domain that is controlled by the SpeedTouch™ can be given a specific name, chosen by the customer.

Execute the following CLI command to set the domain name for the SpeedTouch™ at site A to myvpn:

```
{ST-A} :dns server config domain=myvpn
```

DNS Server Host List

The CLI command group **`dns server host`** allows to configure understandable host names on the local network. In general, hosts configured by DHCP have been introduced in the DNS table by the Host Manager. However, sometimes the default names are not very descriptive.

Execute the following CLI command to show the known host names on the SpeedTouch™ at site A:

```
{ST-A} :dns server host list
Address          Hostname                        TTL (s)   Creator id
10.0.2.1       * laptop_1                              0            1
10.0.1.200     * desktop_1                             0            1
```

It is possible to add more understandable names to the DNS table.

Execute the following CLI commands to add hostnames dummy__1 and dummy__2 to the SpeedTouch™ at site A:

```
{ST-A} :dns server host add name=dummy__1 addr=10.0.1.200 ttl=10
{ST-A} :dns server host add name=dummy__2 addr=10.0.2.1 ttl=10
```

The host list is appended with the added hosts.

```
{ST-A} :dns server host list
Address          Hostname                        TTL (s)   Creator id
10.0.2.1         dummy__2                             10            0
10.0.1.200       dummy__1                             10            0
10.0.2.1       * laptop_1                              0            1
10.0.1.200     * desktop_1                             0            1
```

The asterisk preceding the hostname indicated that the DNS entry has been added automatically (via Host Manager).

A VPN with different DNS zones

It is possible that different zones are defined within a VPN. In this case, each zone will have its DNS server. In our example, 3 DNS zones have been defined, Antwerp, Brussels and Internet:

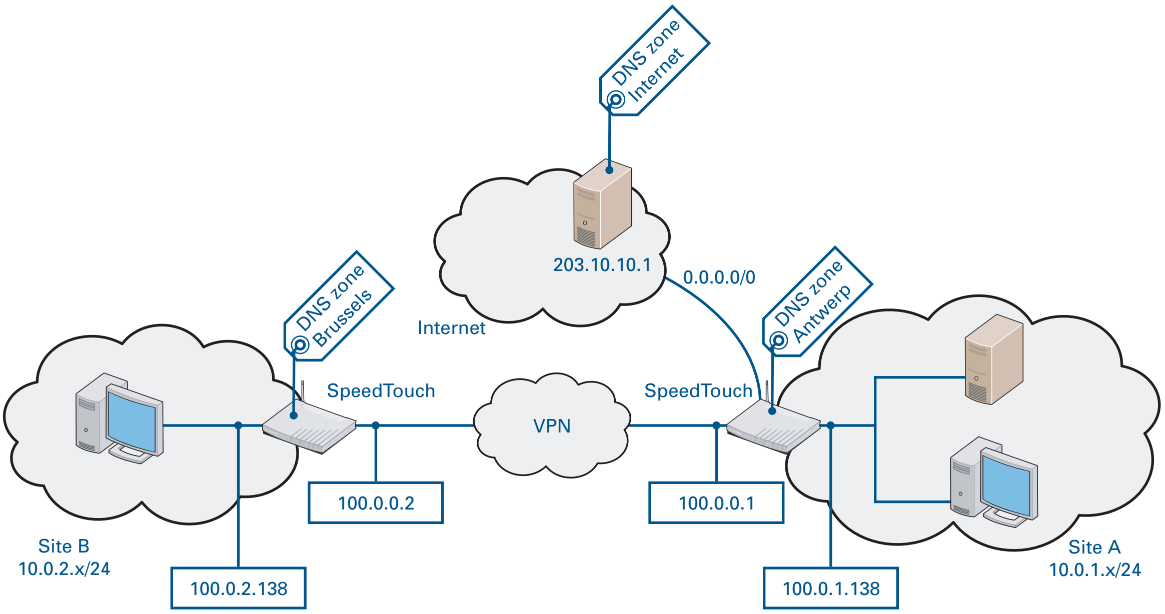

If the requested hostname is unknown to the DNS server in zone Brussels, it forwards the DNS request to the DNS server in zone Antwerp.

Execute the following CLI command to configure the DNS route.

```
{ST-B} :dns server route add dns=100.0.0.1
```

If the DNS server at site A (zone Antwerp) does not know the requested hostname, there are 2 options:

- The hostname belongs to zone Brussels.
- The hostname belongs to zone Internet (not Brussels and not Antwerp)

Execute the following CLI commands to configure the DNS routes.

```
{ST-A} :dns server route add dns=100.0.0.2 domain=brussels.mycompany
{ST-A} :dns server route add dns=203.10.10.1
```

There is no one-to-one correspondence between zones and sites.

5.4 Encrypting Traffic on the VPN

By default, the traffic is isolated in the customer VPN. Basically, this is what a VPN is all about: building virtual, separated networks on top of an existing shared network. A provider network may host a lot of VPNs simultaneously. By design, these VPNs are totally invisible for each other.

In some user scenarios the customer still wants a higher level of security to protect his communication. For these highly secured VPNs the SpeedTouch™608 and SpeedTouch™620 (Wireless) Business DSL Routers feature an IPSec software module encrypting all data passing over the VPN.

For more information see the "The SpeedTouch™ IPSec Quick Start Guide".

This only applies to provider provisioned VPNs.

5.5 Troubleshooting Tools

Introduction

The CLI commands **`ping`** and **`itraceroute`** are available in each command groutp. These commands allow the customer to check the status of his connections at any time.

For more information see the "SpeedTouch™ Operator's Guide".

Ping

Ping is a very basic IP tool, that lets you verify whether a specified IP host is reachable. Loosely, ping means "to get the attention of" or "to check the presence of" another IP host.

The SpeedTouch™ contains an extended version of ping, offering extra options that can be very interesting while debugging. For example the ability to define the size of the packets can be of great help when facing problems with MTU sizes, re-segmentation, etc.

In the example network, you could send a ping from the SpeedTouch™ at site A to a host at site B by executing the following CLI command:

```
{ST-A} :ping addr=10.0.2.1
9 bytes from 10.0.2.1: icmp_id=12, icmp_sec=0
```

The reply from the queried host is displayed in the CLI. If there is no feedback on the CLI, this means there is no reply.

Additional to the destination IP address, the SpeedTouch™ ping command can be used with optional parameters, allowing to modify the ping sequence.

For more information see the "SpeedTouch™ Operator's Guide".

Traceroute

Traceroute is a utility that records the route between you and a specified destination.

This is done by means of limiting the time to live (ttl). Traceroute starts by sending a packet with ttl=1. A router that receives a packet with ttl=1 is not allowed to forward this packet. Instead it sends a "ttl expired" message to the sender. Traceroute then increases the ttl, so the next packet arrives one hop further. This procedure continues until the destination is reached.

In our example network, you could find the path followed for an IP packet travelling from the SpeedTouch™ at site A towards a host at site B by executing the following CLI command:

```
{ST-A} :traceroute addr=10.0.2.1
ttl=1 100.0.0.2        10816 us        10813 us        9970 us
ttl=2 10.0.2.1         11327 us        11330 us        10853 us
```

The hops identified by the traceroute tool are displayed, with their IP information and round-trip time for each try.

Additional to the destination IP address, the SpeedTouch™ traceroute command can be used with optional parameters, allowing to modify the traceroute sequence.

For more information see the "SpeedTouch™ Operator's Guide".

5.6 SLA monitoring tools

Service Level Agreement monitoring

If you have a Service Level Agreement with your ISP, you probably like to monitor the performance of the network service by yourself. The SpeedTouch™620 offer you some basic means to automate continuous monitoring of network paths, and get feedback on the results. The repetitive network tests are based on the troubleshooting tools **`ping`** and **`traceroute`**. The test results can be monitored either via the CLI, the web interface or via an SNMP monitoring tool.

Typically you use one SpeedTouch™ to monitor the quality of all the connections within the VPN.

Setting up an SLA test

The procedure to set up a repetitive SLA test based either on **`ping`** or **`traceroute`** is similar.

Proceed as follows to set up an SLA test based on **`ping`**:

1 Add a new test:

```
:sla ping add name=pingtest1 addr=10.0.2.1
```

2 Set the parameters:

```
:sla ping modify test=pingtest1 size=200 count=5 frequency=120
```

This command sets the repetition period of the test to 120 seconds. Each test consists of five ping messages (the probe), the length of each message is 200 octets. For detailed information on the parameters, see the "SpeedTouch™ CLI Guide".

3 Start the repetitive test:

```
:sla ping start test=pingtest1
```

Viewing the test results

In the CLI, type the following command to get the history of an SLA ping test:

```
:sla ping hist test=test1 owner=modem
```

When an **`sla ping`** test is configured on the SpeedTouch™ it is configured by default with **`owner=modem`**

The result looks like this :

```
Index  Rtt[us]               Status    RC                Timestamp
    1      844        resp received     0 23/01/06 12:34:11.006124
    2      683        resp received     0 23/01/06 12:36:12.024748
    3      703        resp received     0 23/01/06 12:38:13.554346
    4      675        resp received     0 23/01/06 12:40:14.573502
```

To see similar results in the web interface, browse to:

Expert Mode -> SpeedTouch -> SLA ->History

The details of the tests and the last results can be displayed in the CLI with the **`sla ping list`** command.

```
:sla ping list
test1_sla_ping : [owner = modem] dest = 10.0.2.1
        size = 200 timeout[s] = 3 count = 5
        datafill =
        frequency[s] = 120 maxrows = 50
        trapflag = testFailure
        probefailfilter = 1 testfailfilter = 1
        type = IcmpEcho storagetype = nonVolatile
        descr =
        srcaddr = 0.0.0.0
        intf = none bypassrt = no dsfield = 0

        result Info
        status = in progress
        minrtt[us] = 4609 maxrtt[us] = 5539
        avgrtt[us] = 5135 rttsumofsqr[ms] = 107
        reponses = 5 sentprobes = 5
        lastgoodresponse = 01/02/05 16:26:52.689245
```

SNMP traps

The SpeedTouch™ allows you to generate traps as an outcome of an SLA test. The traps can be sent to an SNMP server for further processing.

Three types of traps can be generated:

- pingProbeFailed: when a specified number of consecutive probes is not successful, i.e. a number of consecutive individual pings within one test fails
- pingTestFailed: when a specified number of pings within a single test is not successful.
- pingTestCompleted: at the end of every test. When enabled, this trap is sent with the same repitition period as the SLA test.

Reported objects

For each of the traps, the following objects are reported:

- pingCtlTargetAddressType
- pingCtlTargetAddress
- pingResultsOperStatus
- pingResultsIpTargetAddressType
- pingResultsIpTargetAddress
- pingResultsMinRtt
- pingResultsMaxRtt
- pingResultsAverageRtt
- pingResultsProbeResponses
- pingResultsSentProbes
- pingResultsRttSumOfSquares
- pingResultsLastGoodProbe

Sending SNMP traps to an SNMP server

Proceed as follows to send traps to an SNMP server.

1 Enable the required SpeedTouch™ system services:

```
:service system modify name=SNMPV3_AGENT state=enabled
:service system modify name=SNMPV3_TRAPS state=enabled
```

2 Specify an SNMP server as destination for the traps:

```
:snmp target add name=SNMPserver addr=192.168.1.64 taglist=V1TrapTag params=V1Params
```

3 Specify which traps you want to generate. In the following example, all traps are activated:

```
:sla ping modify test=test1 trap=probeFailure+testFailure+testCompletion
```

5.7 Remote Management

Remote Access

Users who can access the SpeedTouch™ from the WAN (for example by means of telnet) have remote access. By default, users with the following roles have remote access:

- root (the super user)
- SuperUser (access from anywhere to any service)
- TechnicalSupport (access from WAN to any service, typical user profile for remote access)
- WAN_Admin (access from WAN to a limited number of services)

In general, a specific subnet will be created to allow remote access for technical support. Examples of technical support are:

- Retrieve past **`sla ping`** or **`sla traceroute`** events to investigate the situation in the past or start **`ip debug ping`** or **`ip debug traceroute`** commands to study the current situation.
- Configure the SpeedTouch™ according to the wishes and requirements of the customer.

Simple Network Management Protocol (SNMP)

The SpeedTouch™ supports a number of SNMP Management Information Bases (MIBs). By default, the SNMP agent is activated on the LAN. In the example network, the customer wishes to manage his devices from an SNMP server at site A. Therefore, the SNMP agent embedded in the SpeedTouch™ at site B has to be activated on the WAN.

To enable SNMP communication from the WAN on the SpeedTouch™ at site B, execute the following CLI command:

```
{ST-B} :service system ifadd name=SNMP_AGENT group=wan
```

Then, the SNMP application of the remote operator can be used to browse the MIBs of the SpeedTouch™, for example the SLA MIB.

5.8 ISDN Fall-Back Connectivity

Introduction

Some business customers need a reliable connection between branch offices and main office, even if the DSL service drops out. In this section, it is assumed that only the main office is connected to the public network via a SpeedTouch™ equipped with an ISDN fall-back interface.

In this case, the SpeedTouch™ Integrated ISDN modem fall-back solution can guarantee that connectivity is maintained. The switch-over requires that routes are communicated between the gateways over the ISDN connection if the DSL service drops out. The switch-over can be realised by means of RIP or link sensing.

As an alternative, you can configure the ISDN interface as a dial-on-demand link. In this case, the activation of the ISDN is triggered by outgoing traffic from the main office. This scenario is described in “Fall-back WAN Connectivity with the Integrated ISDN Modem”. In order to guarantee outgoing traffic from the central office, the SLA PING test can be used. This traffic is able to trigger the activation of the ISDN fall-back interface. This is described in the following sections.

Configuring the ISDN fall-back interface for dial-on-demand

The configuration procedure for an operational ISDN fall-back interface is described in “Fall-back WAN Connectivity with the Integrated ISDN Modem”.

Type the following commands to set the ISDN fall-back interface to dial-on-demand mode:

```
:ppp ifdetach intf=ISDN_fallback
:ppp ifconfig intf=ISDN_fallback demanddial=enabled
:ppp rtadd int=ISDN_fallback dst=0.0.0.0/0 metric 10
:ppp ifattach intf=ISDN_fallback
```

Configuring an SLA ping test that can trigger the ISDN fall-back interface

Proceed as follows to configure an SLA ping test that is able to trigger the ISDN fall-back interface in case the DSL interface fails:

1 Add a new test. Select a destination at one of your remote offices:

```
:sla ping add name=pingtest1 addr=10.0.2.1
```

2 Set the parameters:

```
:sla ping modify test=pingtest1 intf=none
```

Setting the interface to “none” means that no specific interfaces is specified. This implies that the ping messages can be sent via any of the available interface, according to the routing rules in the routing table. If the DSL interface is not available, the fall-back route will be triggered.

3 Start the repetitive test:

```
:sla ping start test=pingtest1
```

Configuring an SLA ping test that will not trigger the ISDN fall-back interface

If, on the contrary, you want to use the SLA ping test exclusively to monitor your DSL line and you do not want this test to trigger the ISDN fall-back interface, proceed as follows.

In this case you have to define a specific interface to launch the SLA ping test, and you have to bypass the routing tables. This is done in the following way.

```
:sla ping modify test=pingtest1 intf=Internet bypassrt=enabled
```

Visit us at:

www.speedtouch.com

Acknowledgements

All Colleagues for sharing their knowledge.

Coordinates

THOMSON Telecom Belgium
Prins Boudewijnlaan 47
B-2650 Edegem
Belgium
E-mail: documentation.speedtouch@thomson.net

speed**touch**™

Copyright

A **THOMSON** BRAND