Quick Note 20

Configuring a GRE tunnel over an IPSec tunnel and using BGP to propagate routing information

(GRE over IPSec with BGP)

Contents

1 INTRODUCTION

1.1 Outline

This document describes how to configure a GRE tunnel within an IPSec tunnel to secure communications between routers. The GRE tunnel proveds a point-to-point link between the routers that can be used by BGP as well as transferring regular data.

If BGP is not used, but Static Routes are, Please refer to the AN74 “How to configure a GRE over IPsec Tunnel between Digi TransPort WR Routers”.

The scenario considered in this document s the following:

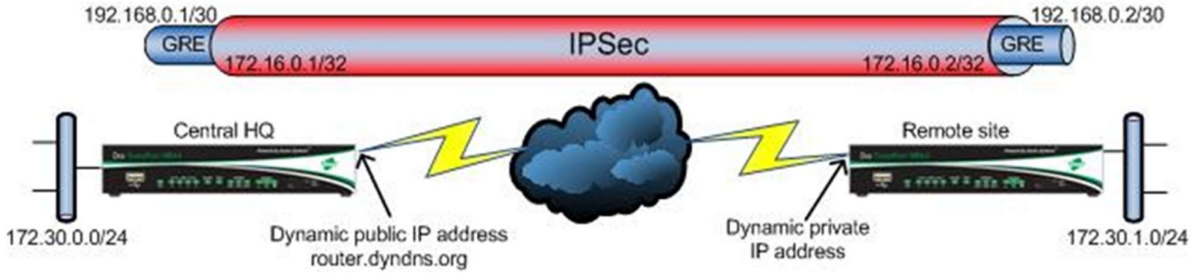

An IPSec tunnel is setup to ensure secure communications between the Central HQ and the Remote site. A GRE tunnel is configured to run through the IPSec tunnel to allow point to point communication between the 2 sites. This is used when a process such as a routing protocol needs point to point communication between 2 sites and a point to point link such as leased line is not available.

Both routers have been configured with internet connectivity, the Central HQ router uses ADSL with a dynamic public IP address but uses the DynDNS service so it can always be reached at router.dyndns.org; the Remote site router has a cellular link and is allocated a private IP address by the mobile operator. LAN segments are attached on Eth0.

1.2 Assumptions

This guide has been written for technically competent personnel who are able to configure a standard IPSec tunnel between 2 TransPort WR routers and are familiar with the use of routing protocols.

Configuration: This guide assumes that the routers have been configured already with internet access.

This application note applies to;

Models shown: Central HQ router, Digi Transport DR64 router with ADSL. Remote site router, WR44 with a cellular link running. Both routers are running firmware version 5081.

Other Compatible Models: All Digi Transport products.

Firmware versions: 4905 or later.

1.3 Corrections

Requests for corrections or amendments to this application note are welcome and should be addressed to: tech.support@digi.com.

Requests for new application notes can be sent to the same address.

1.4 Version

Version Number	Status
1.0	Published
2.0	Re-written and updated
2.1	Update for new GUI
3.0	Update New WEB GUI and branding. Overall fixes and reference to new doc for GRE with Static Route

2 CONFIGURE IPSEC VPN

2.1 Configure IKE

On both routers, browsing in the WEB GUI to the IKE section and configure as follows:

Central HQ (IPSec responder):

CONFIGURATION - NETWORK > VIRTUAL PRIVATE NETWORKING (VPN) > IPSEC > IKE > IKE RESPONDER

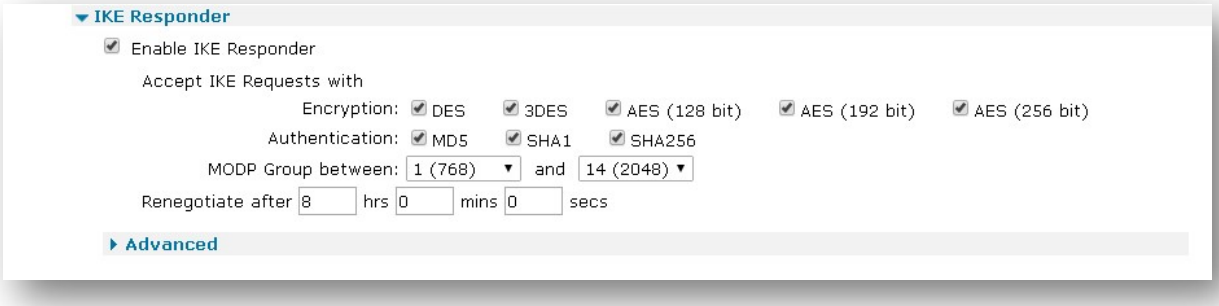

Remote site (Initiator):

CONFIGURATION - NETWORK > VIRTUAL PRIVATE NETWORKING (VPN) > IPSEC > IKE > IKE 0

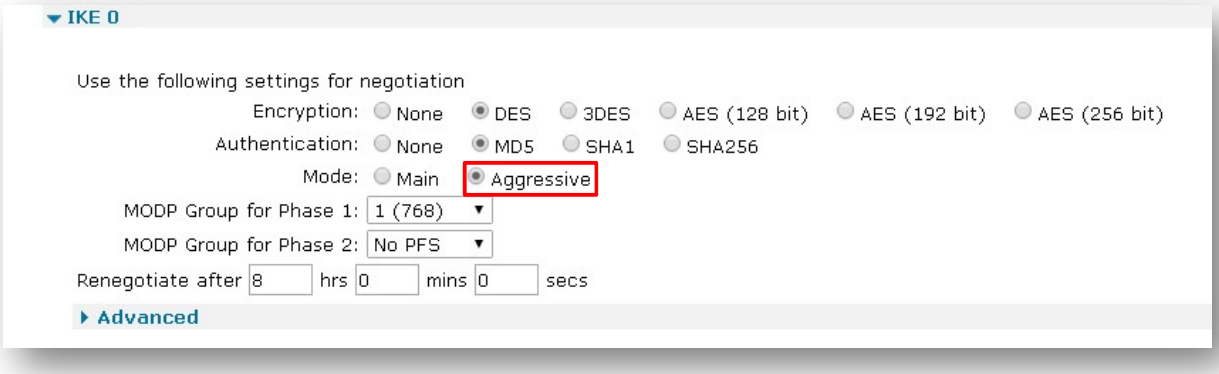

The IKE configuration is default on both routers except for enabling aggressive mode on the Remote site IPSec initiator.

2.2 Configure IPSec

On both routers, configure the IPSec tunnel as follows:

Central HQ (IPSec responder):

CONFIGURATION - NETWORK > VIRTUAL PRIVATE NETWORKING (VPN) > IPSEC > IPSEC TUNNELS > IPSEC 0

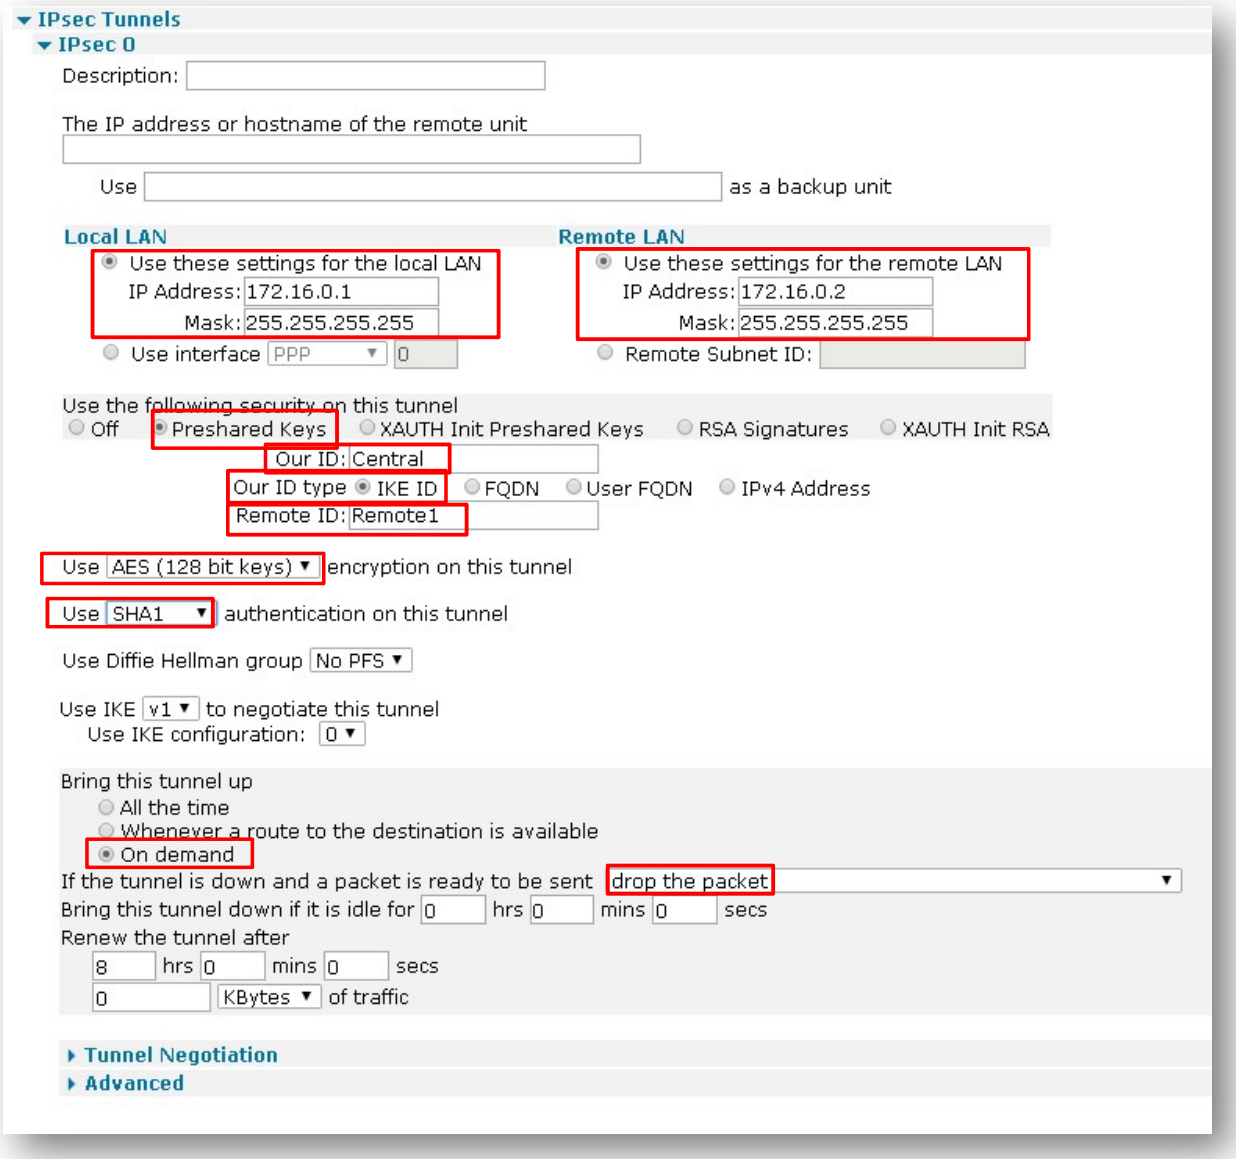

Remote site (Initiator):

CONFIGURATION - NETWORK > VIRTUAL PRIVATE NETWORKING (VPN) > IPSEC > IPSEC TUNNELS > IPSEC 0

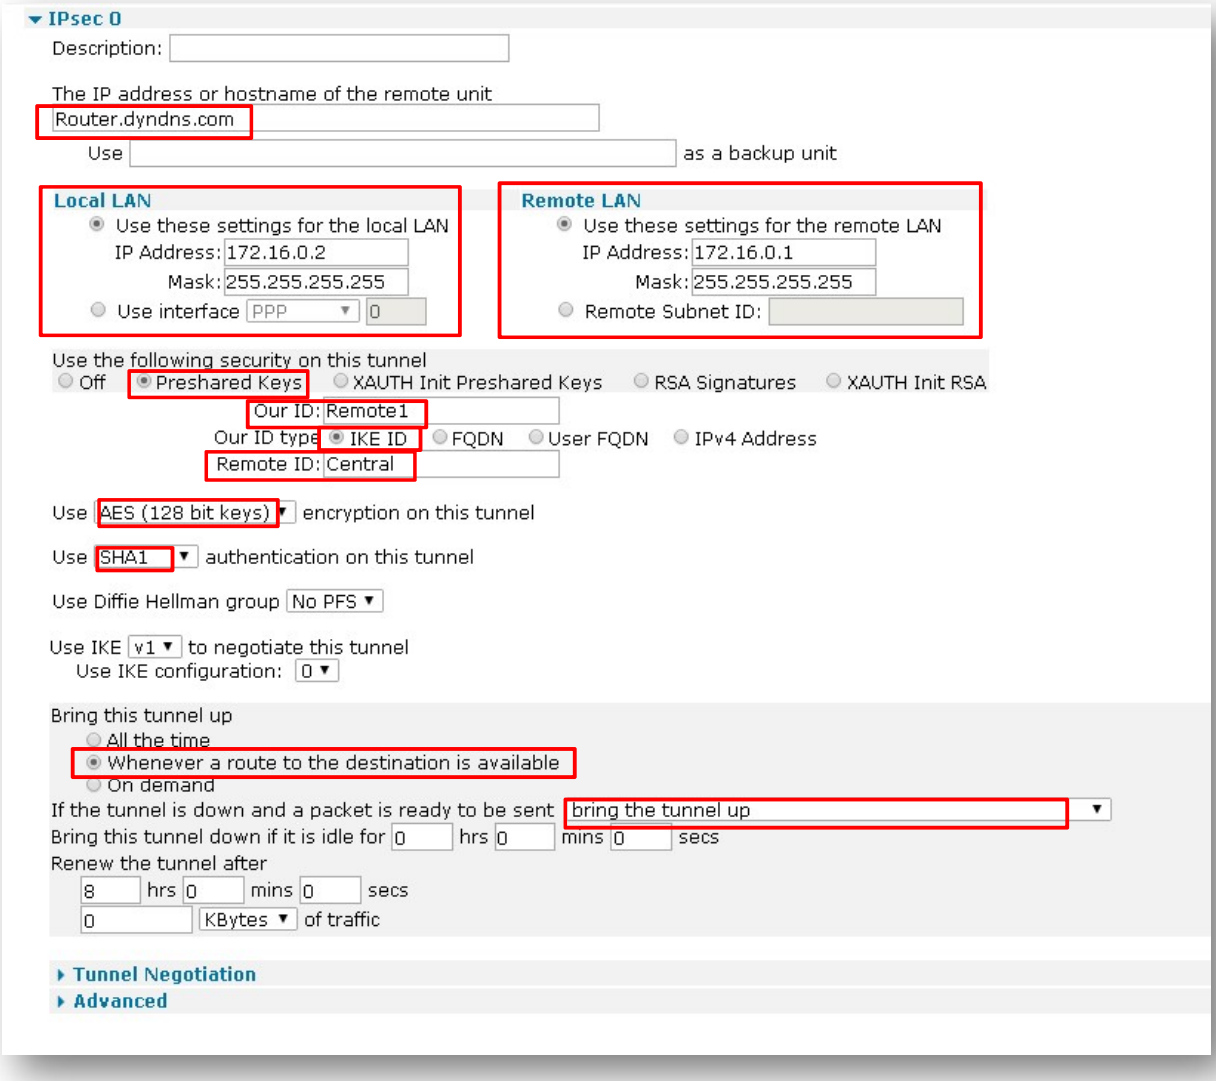

This Eroute config is exactly the same as a regular IPSec tunnel except for the following fields:

Local subnet IP address, Local subnet mask, Remote subnet IP address, Remote subnet mask

These fields are configured with a host IP address that does not actually exist (use an unused IP address from an unused subnet, it doesn't matter what is used). These are the end points of the IPSec tunnel. In this example 172.16.0.1 is used on the Central HQ router and 172.16.0.2 is used on the Remote site router, both with the subnet mask 255.255.255.255

2.3 Configure Pre-Shared Key

The PSK is configured as in a regular IPsec Tunnel, using the Users section.

Central HQ (IPSec responder):

CONFIGURATION - SECURITY > USERS > USER 10 - 14 > USER 10

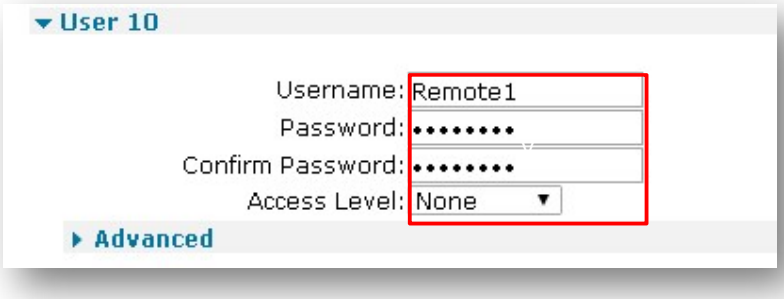

Remote site (Initiator):

CONFIGURATION - SECURITY > USERS > USER 10 - 14 > USER 10

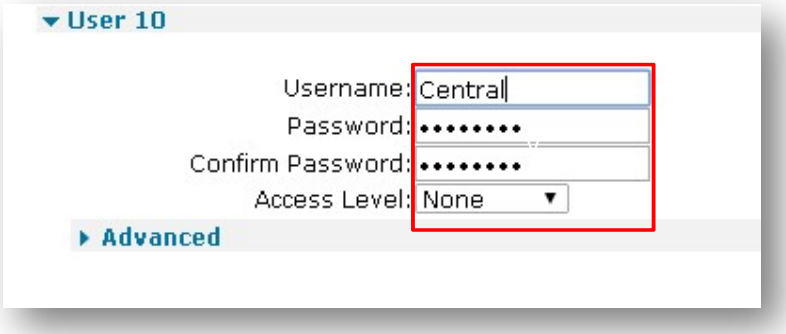

The pre-shared key is configured as shown, the name is the ID that the other router sends as its 'Our ID' from the eroute parameters. The Password needs to match on both routers as this is the shared key. The Access level should be none, as this user does not need access to the router administration interfaces.

3 CONFIGURE GRE TUNNELS

Central HQ (IPSec responder):

CONFIGURATION - NETWORK > INTERFACES > GRE > TUNNEL 0

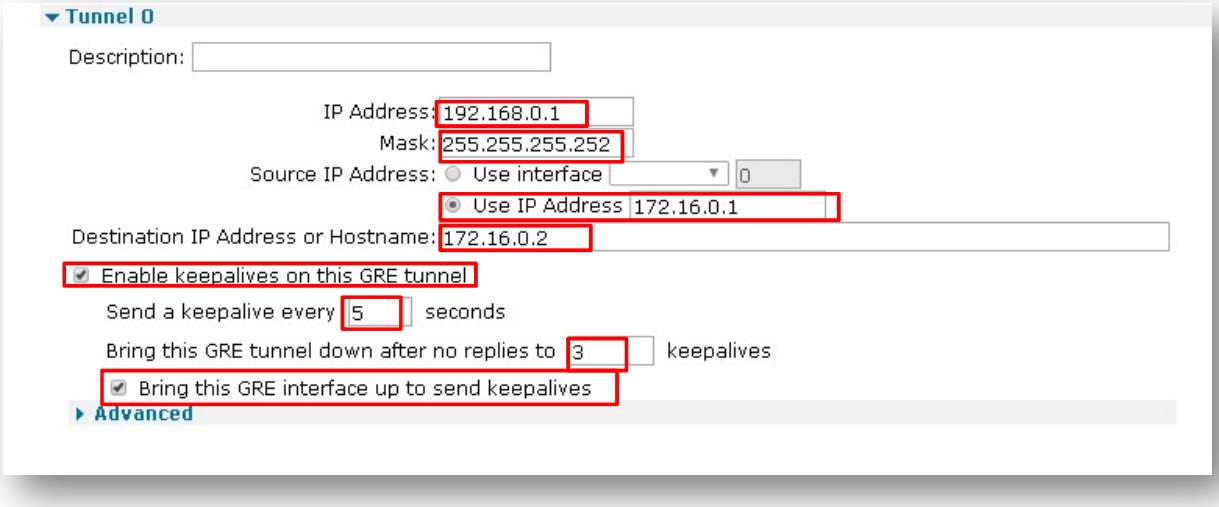

Remote site (Initiator):

CONFIGURATION - NETWORK > INTERFACES > GRE > TUNNEL 0

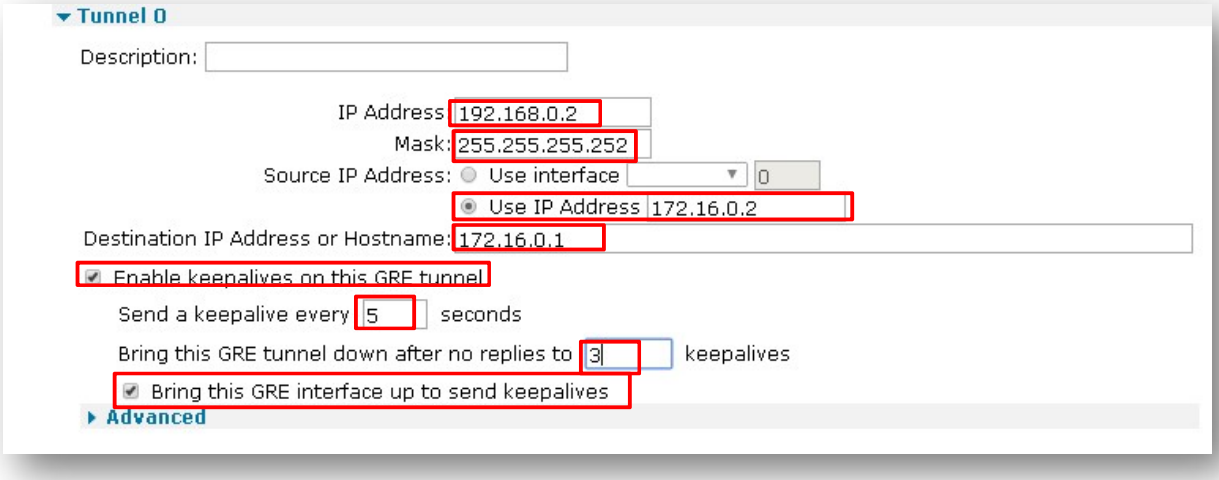

he GRE tunnel is configured as a point to point connection using the 192.168.0.0/30 subnet. Note the usage of the previously configured addresses 172.16.0.1 and 172.16.0.2 from within the Eroute settings, these are the source and destination IP addresses of the IPSec tunnel that GRE will tunnel through.

4 CONFIGURING BGP

Each router will need a bgp.conf file creating, this is a plain text file created using notepad. The file contains the parameters that BGP will use.

The *#macros* section does not need to be used but can contain be used to define parameters such as hello intervals that will be used across all sites.

The *#global configuration* section is where the main BGP configuration is defined.

An example bgp.conf file contains:

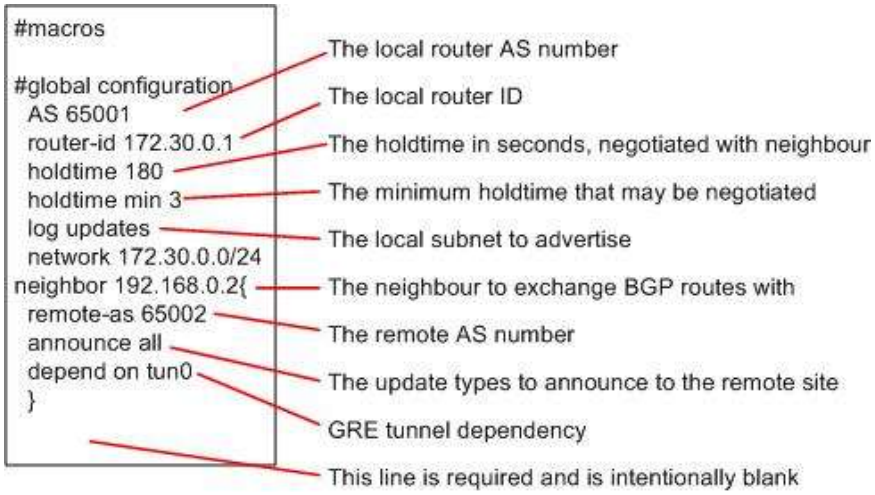

AS numbers

The AS numbers should be configured in the private AS range 64512-65535.

If the local and remote AS is the same then IBGP is inferred and when scaling up routes will only be exchanged between directly connected neighbours.

If the local and remote AS are different then EBGP is inferred and when scaling up routes will be exchanged between directly connected neighbours and all other neighbours connected to the central router.

Neighbor

The IP address defined in the statement “neighbor 192.168.0.2” is the IP address assigned to the remote end point of the GRE tunnel.

NOTE:

The blank line at the end of the bgp.conf, after the final “}” IS required, otherwise BGP will not start. Be aware of the American spelling of “neighbor”.

The options for the bgp.conf file are explained fully at the following web site:
http://www.openbsd.org/cgi-bin/man.cgi?query=bgpd.conf

4.1 Create the bgp.conf text files

Using notepad create a file with the following contents for each router. The files can be named anything you like, but we recommend something like bgp.conf so it is obvious what the file is.

In this example, the files are named bgpc.conf for the central router & bgpr.conf for the remote router.

Central HQ (IPSec responder)	Remote site (Initiator)

#macros
 	
#global configuration
AS 65001
router-id 172.30.0.1
holdtime 180
holdtime min 3
log updates
network 172.30.0.0/24
neighbor 192.168.0.2{
remote-as 65002
announce all
depend on tun0
}
[blank line]
	#macros
	
	#global configuration
	AS 65002
	router-id 172.30.1.1
	holdtime 180
	holdtime min 3
	log updates
	network 172.30.1.0/24
	neighbor 192.168.0.1{
	remote-as 65001
	announce all
	depend on tun0
	}
	[blank line]

These bgp configuration text files need to be FTP uploaded onto the respective routers.

4.2 Enable BGP

Once the configuration files are uploaded into the routers, BGP needs to be enabled and the BGP file associated. The configuration wll be the same on both central and remote router:

Central HQ (IPSec responder) & Remote site (Initiator)

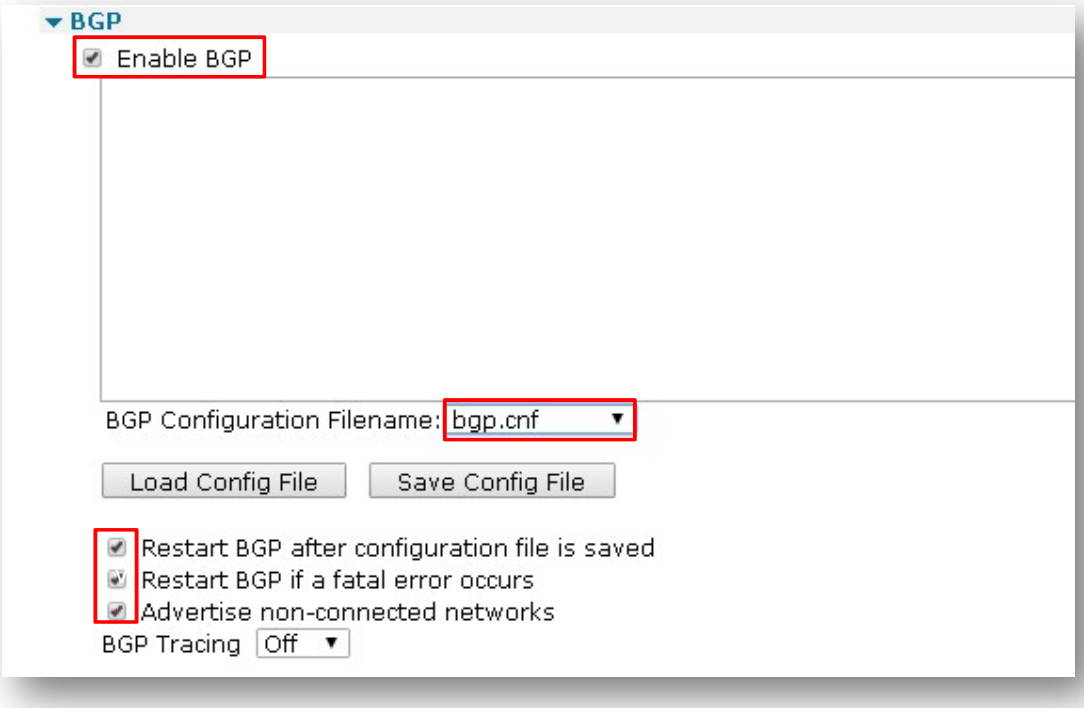

NOTE: Be sure that the two Restarts option highlighted are ticked.

4.3 Save your config changes to profile 0

ADMINISTRATION - SAVE CONFIGURATION

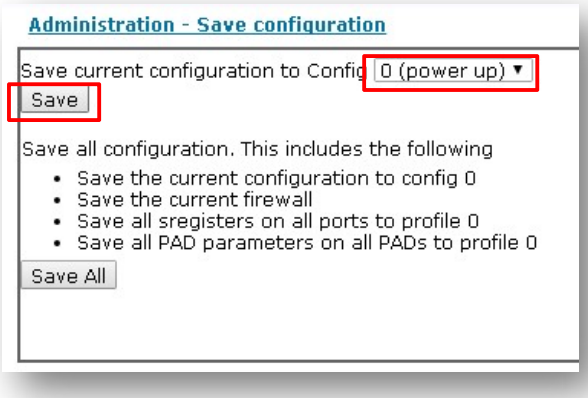

5 TESTING

5.1 Check the routing tables

Check the Routing Tables on both routers with the “route print” command:

Central HQ (IPSec Responder):

```
route print

       Destination          Gateway   Metric   Protocol  Idx Interface  Status
-------------------------------------------------------------------------------
        10.1.0.0/16        10.1.51.2      1      Local     -       ETH 0   UP
   95.154.209.28/32    217.34.133.29      2     Static     0       ETH 3   UP
      172.30.0.0/24       172.30.0.1      1      Local     -       ETH 1   UP
      172.30.1.0/24      192.168.0.2     20       EBGP     -       TUN 0   UP
     192.168.0.0/30      192.168.0.1      1      Local     -       TUN 0   UP
   217.34.133.16/28    217.34.133.21      1      Local     -       ETH 3   UP

         0.0.0.0/0     217.34.133.29      2     Static     3   ETH 3      UP
OK
```

Remote site (IPSec Initiator):

```
route print

       Destination          Gateway   Metric   Protocol  Idx Interface  Status
-------------------------------------------------------------------------------
        10.1.0.0/16        10.1.51.4      1      Local     -       ETH 0   UP
      172.30.0.0/24      192.168.0.1      0       EBGP     -       TUN 0   UP
      172.30.1.0/24       172.30.1.1      1      Local     -       ETH 1   UP
     192.168.0.0/30      192.168.0.2      1      Local     -       TUN 0   UP

         0.0.0.0/0       10.1.2.100      1     Static     0   ETH 0      UP

OK
```

Viewing the routing table shows:

The local LAN segment and the interface it is configured on.

The GRE tunnel /30 subnet. The gateway address is the remote GRE IP address.

The remote LAN subnet. This will be routed to via TUN 0, its gateway will be the remote IP address of the GRE tunnel and the protocol will be EBGP (or IBGP if the same AS numbers were used).

5.2 Test connectivity

An easy test to check connectivity is to ping from each router to the ETH port of the other one:

Central HQ (IPSec responder)

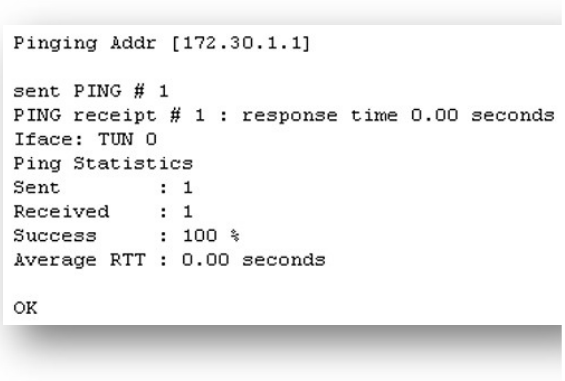

Remote site (Initiator)

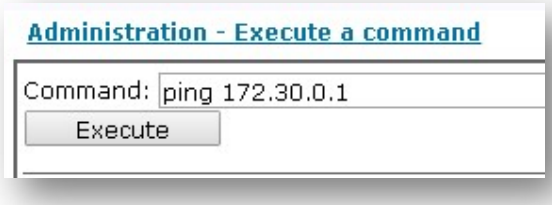

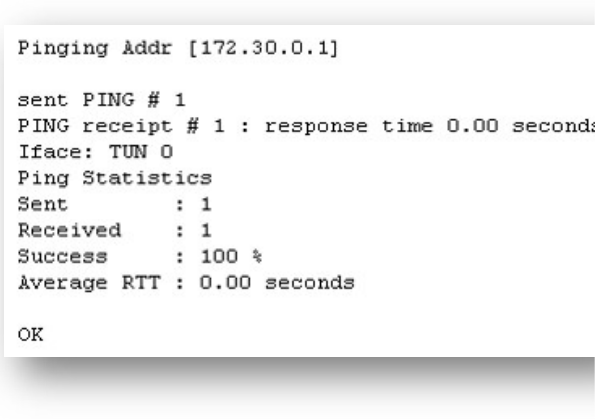

Note: Although this guide is written using ADSL and Cellular connectivity the testing was done using Ethernet as the WAN connectivity, this is why the ping response time is 0.00 seconds.

6 SCALING UP - ADDING MORE SITES

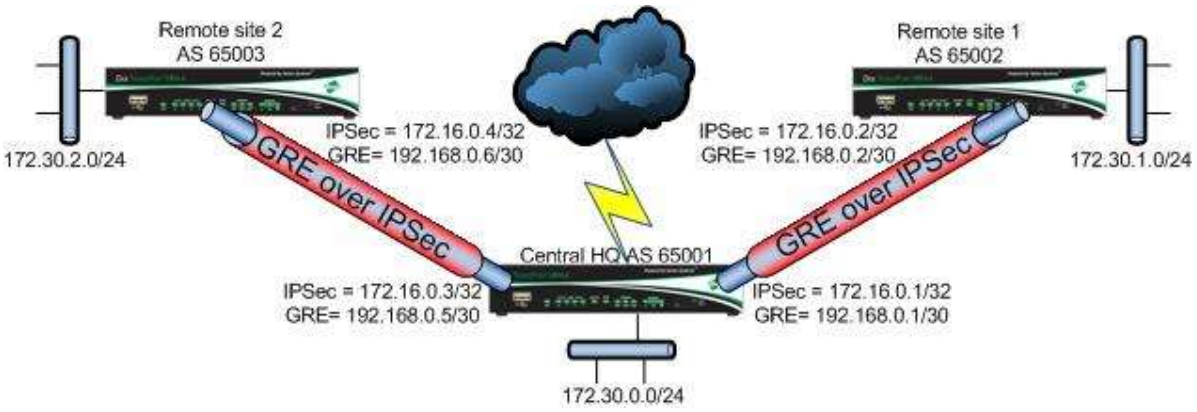

This scenario can be scaled up to add more connected sites to the Central HQ router. To add another site, create an IPSec/GRE tunnel between the Central HQ router and the new site router. The next tunnel on the Central HQ router to the new site will be Tun1 with local IP address 192.168.0.5/30 and remote IP address 192.168.0.6/30.

Note the use of the new command “set nexthop self” in the BGP configuration file, this is only used on the Central HQ router to enable routing between sites, the command will set the Central HQ router as the next hop when advertising updates about remote networks.

bgp.conf from Site 2 router	bgp.conf from Central HQ router	bgp.conf from Site 1 router

#macros
# global configuration 	
AS 65003
router-id 172.30.2.1
holdtime 180
holdtime min 3
log updates
network 172.30.2.0/24
neighbor 192.168.0.5{
remote-as 65001
announce all
depend on tun0
}
[blank line]
	#macros
	# global configuration
	AS 65001
	router-id 172.30.0.1
	holdtime 180
	holdtime min 3
	log updates
	network 172.30.0.0/24
	neighbor 192.168.0.2{
	remote-as 65002
	announce all
	set nexthop self
	depend on tun0
	}
	neighbor 192.168.0.6{
	remote-as 65003
	announce all
	set nexthop self
	depend on tun1
	}
	[blank line]
		#macros
		# global configuration
		AS 65002
		router-id 172.30.1.1
		holdtime 180
		holdtime min 3
		log updates
		network 172.30.1.0/24
		neighbor 192.168.0.1{
		remote-as 65001
		announce all
		depend on tun0
		}
		[blank line]

7 BASIC TROUBLESHOOTING

In order to to do a basic troubleshoot on this configuration, do the following steps:

- Make sure the IPSec tunnel is up. Execute “sastat” from the CLI.
- Check the GRE tunnel is up. Execute “tunstat 0” or “tunstat 1” from the CLI.
- Disable the firewall if it is enabled.
- Enable Debug. From CLI, “debug 0” if using a serial connection, “debug t” if using telnet. Then issue the command “bgp 0 debug 3” for high level debug output.
- Stop and start BGP from the CLI and make sure BGP is starting correctly from the output. To stop BGP “bgp 0 enable off”, to restart BGP “bgp 0 enable on”.

The output should be similar to the following:

```
bgp 0 enable 1
OK

Start BGP and enabled
AS 65001
router-id 10.1.0.2
holdtime 180
holdtime min 3
fib-update yes
log updates

network 10.1.0.0/16


neighbor 192.168.0.2 {
        remote-as 65002
        announce all
        enforce neighbor-as yes
        depend on "tun0"
        announce IPv4 unicast
        softreconfig in yes
        softreconfig out yes
}


startup
route decision engine ready
session engine init
listening on 0.0.0.0
listening on 0.0.0.0
session engine init done
neighbor 192.168.0.2: state change None -> Idle, reason: None
neighbor 192.168.0.2: state change Idle -> Connect, reason: Start
neighbor 192.168.0.2: state change Connect -> OpenSent, reason: Connection opened
neighbor 192.168.0.2: state change OpenSent -> OpenConfirm, reason: OPEN message received
neighbor 192.168.0.2: state change OpenConfirm -> Established, reason: KEEPALIVE message received
neighbor 192.168.0.2 (AS65002) update 172.16.51.0/24  via 192.168.0.2
nexthop 192.168.0.2 now valid: directly connected: via 192.168.0.1
```

BGP is confirmed to be enabled, but not yet running and routing processes.

The bgp.conf file is read and displayed in the debug output.

Startup, shows that the bgp.conf was read correctly and BGP is now started.
The route decision engine starts, the session engine initialises & BGP listens for BGP updates from other routers.

Neighbourships are then created and BGP update messages are displayed.

The BGP process can be further debugged using the bgpctl command. The usage of bgpctl is documented at the following web site: http://www.openbsd.org/cgi-bin/man.cgi?query=bgpctl

Useful (abbreviated) commands are:

bgpctl sh nei	Shows neighbours and stats
bgpctl sh fib	Show the forwarding information base
bgpctl sh rib	Shows the routing information base
bgpctl sh sum	Shows a summary of neighbours, AS's & uptime
bgpctl sh ip bgp det	Shows details information about BGP neighbours