

High Availability for VPNs

# ALTERNATIVE PATH

www.scs.hu

IPSec prevents many of the clever tricks high-availability products employ. We'll show you a solution that provides transparent backup for IPSec connections. **BY JOHANNES HUBERTZ**

**S**ystem administrators often want a network connection system that switches transparently to a backup if the primary connection goes down. But if you use a VPN with IPSec to protect your traffic en route through the Internet, the backup line needs some special attention.

The reason for this attention is that IPSec [1] [2] requires consistent IP addresses at the endpoints of a tunnel, so when the network switches to a different tunnel, the IP addresses must switch to the new endpoints or else existing connections will be terminated. The Border Gateway Protocol (BGP [3]) offers a reliable means of maintaining a highly-available pool of IP addresses with a number of providers. Unfortunately, provider service agreements often prevent admins from using BGP for an existing Internet connection.

As a workaround, many admins do without automation, and if worst comes to worst, switch manually from the standard line to the backup line in a process

that often involves physically patching the interfaces. Not exactly state-of-the-art. It would be preferable for the network devices to detect a line failure and switch automatically. The ideal system would also automatically generate the settings for both endpoints by referring to a central configuration.

For firewalls and IPsec gateways, the central configuration is state-of-the-art technology. In Linux, SSPE (Simple Security Policy Editor, [4]) handles this. However, the HA solution introduced in this article is not SSPE capable as yet.

## Linux-HA

The Linux-HA project [5] is dedicated to high availability solutions for Linux servers. This software allows admins to set up a highly-available VPN (that doesn't require BGP) that will switch from standard to backup operations more quickly. To implement this solution, you need two independent parallel tunnels; you'll only use one of these tunnels at any given time. Each tunnel in each net-

work has its own endpoint, which serves as a gateway to the local network. Linux-HA implements automatic reconfiguration of IP addresses to support this setup.

The HA nodes both have an individual address and a shared address. The shared address is only used by one of the machines at any given time. The mechanism is designed for operating (Web, email...) servers with a shared address. The service runs on both machines and listens to each IP address, however, requests only arrive via the shared IP. This allows Linux-HA to assign the external address to the second machine in case of emergency, ideally without users noticing that the first machine is down.

A serial cable provides the heartbeat for the two machines. The heartbeat is an important part of Linux-HA, as the computers involved use it to check the availability of their peers. If one computer fails, the partner computer adopts the IP address belonging to the first com-

puter. It generates ARP broadcasts using the shared IP and its own MAC address on the LAN (and thus performs a kind of legal ARP-Spoofing). Additionally, the machine enables an interface alias. As soon as the first node gets back online, the heartbeat protocol ensures that Server 2 will disable the alias interface again, while Server 1 enables its interface and ARP broadcasts on the LAN.

## Two Linux-HA Installations

The scenario shown in Figure 1 displays the customer on the left, and the outsourcing company on the right. Customer side (left, *gw-aa*) and service provider side IPsec gateways (right, *gw-ba*) use ESP (Encapsulating Security Payload, an IPsec protocol) to send

packets across the tunnel. The replacement tunnel (2) is configured on the two backup routers at the top (*gw-ab*, *gw-bb*) in a similar way to Tunnel 1.

Linux-HA is running between *gw-aa* and *gw-ab* and between *gw-ba* and *gw-bb*. The heartbeat uses the serial port, which is not required for any other purpose and is thus independent of the network, IPsec, and IPtables. Both HA installations work separately and independently of one another. In normal operations, *gw-aa* and *gw-ba* have local router IPs (10.1.255.254 left and 10.31.0.254 right). In case of an outage, these addresses migrate to *gw-ab* and *gw-bb*. To allow admins to explicitly connect to the computers in an HA pool, both additionally have static addresses

on their respective LANs; this is 10.31.0.252 for *gw-bb* for example.

## Scripting an HA-VPN

The default gateways *gw-aa* and *gw-ba* run a shell script like the one shown in Listing 1. The script is launched by an *inittab* entry and monitors the accessibility of the other side, independently of the IPsec tunnel. To do so, the script sends short UDP packets to the echo port. Echo ping is a standard component for most Linux distributions and can use UDP if you set the right options. This approach avoids issues with over zealous firewalls that drop any ICMP messages and thus block normal pings.

As long as the connection between *gw-aa* and *gw-ba* is up, the *FAIL* counter

### Listing 1: HA VPN Supervisor

```

01 #!/bin/bash
02 # HA VPN Supervisor on gw-aa
03
04 # The other end of the tunnel
05 TARGET="gw-ba"
06
07 # Number of seconds between
08 # pings
09 TIMEOUT=1
10 # Wait MAXFAIL * TIMEOUT,
11 # before enabling backout
12 MAXFAIL=5
13 # Wait HYSTERE * TIMEOUT after
14 # end of outage
15 # before the script switches
16 # back to normal operations
17 HYSTERE=180
18 # Assumption: No error
19 # condition on start
20 FAIL=0
21
22 VERBOSE=""
23 ACTION_FAIL_START="/root/bin/
24 HA-VPN-action-script start"
25 ACTION_OK_AGAIN="/root/bin/
26 HA-VPN-action-script stop"
27
28 PING=/usr/bin/echoping
29 LOG="/usr/bin/logger -t
30 HA-VPN"
31
32 math () {
33     eval echo "\${($*)}"
34 }
35
36 echo "`date +%Y%m%d%H%M%S`
37 `basename $0` starting" | $LOG
38
39 while :
40 do
41     VAL=`$PING ${VERBOSE} -u -t
42     $TIMEOUT -s 5 ${TARGET} 2>&1`
43     ERROR=$?
44     if [ $ERROR -gt 0 ] ; then
45         echo "$DAT $ERROR $FAIL
46         $VAL" | $LOG
47         # Timeout occurred
48         if [ $FAIL -lt 0 ] ; then
49             # Another error during
50             # the recovery phase
51             FAIL=`math $MAXFAIL + 1`
52         fi
53         if [ $FAIL -eq $MAXFAIL ]
54         ; then
55             # Start backout
56             :
57             FAIL=`math $FAIL + 1`
58             echo "$DAT starting
59             backup now: ${ACTION_FAIL_
60             START}" | $LOG
61             ${ACTION_FAIL_START}
62         else
63             if [ $FAIL -lt
64             $MAXFAIL ] ; then
65                 FAIL=`math $FAIL + 1`
66             fi
67         fi
68     fi
69     fi
70     fi
71     fi
72     fi
73     fi
74     fi
75     fi
76     #echo "$DAT $ERROR $FAIL $VAL"
77     | $LOG
78     sleep $TIMEOUT
79 done
80 # never happens:
81 exit 0

```

will always be zero. If the ping fails, the script increments the *FAIL* variable in line 53, as long as the value is below *MAXFAIL*. If the next ping is ok, the script sets *FAIL* back to zero (Line 62). If the value reaches *MAXFAIL*, line 50 calls *ACTION\_FAIL\_START* (Listing 2 with the *start* parameter). *ACTION\_FAIL\_START* disables the local heartbeat, causing the backup gateway to automatically adopt the local router IP.

### Waiting for Normal Operations

The infinite loop keeps running, waits for the line to go back up, and waits for echo ping to return a normal value of 0. When the first response arrives after an outage, the line may not be completely stable, so the script waits before reinstating normal operations. It sets *FAIL* to the negative value of *HYSTERE* (Line 59) and increments *FAIL* for each successful ping (line 67). If another error occurs, Line 43 sets the *FAIL* variable to a value greater than *MAXFAIL* – this keeps the system using the backup.

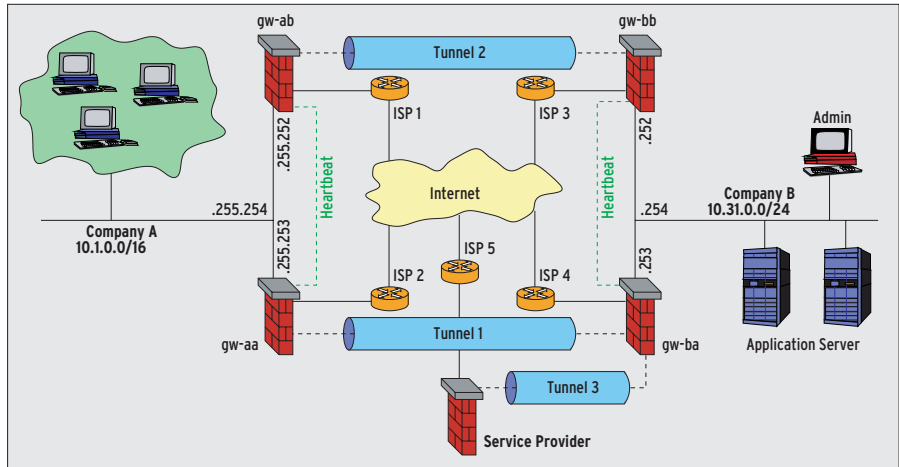


Figure 1: A customer network is attached to an outsourcing service provider via a VPN. This connection uses two alternative paths; if required, Tunnel 2 takes over the role of Tunnel 1.

Normal operations begin when *FAIL* gets back to zero and the script calls *ACTION\_OK\_AGAIN* in Line 71 (Listing 2 with *stop* parameter). The program also re-enables the heartbeat and thus reinstates the local router IP address.

This approach avoids routing flipflops, where gateways oscillate between tunnels. Internet connections are fraught

with short outages, which are resolved after a short period. Three minutes is a good value to ensure stable operations.

Past experience shows that, if the heartbeat at one end of the tunnel downs the gateway, the other end follows suit after a delay of one second at the most.

### Satisfaction Guaranteed

This solution has been online since 2004 and has repeatedly demonstrated that users do not even notice provider-side router outages. In case of total failure, the normal retry mechanisms used by TCP/IP stacks on the application servers and client PCs can easily bridge the ten seconds before the backup solution cuts in. And the kind of disruption that accompanied outages of this kind previously is now a thing of the past. ■

#### Listing 2: HA-VPN Action Script

```

01 #!/bin/bash
02 # HA-VPN-Aktionsskript
03
04 #VERBOSE=-v
05 VERBOSE=""
06
07 NAME=`basename $0`
08 LOG="/usr/bin/logger -t
09   HA-VPN"
10
11 ACTION_FAIL_START="/etc/init.d
12 /heartbeat stop"
13
14 ACTION_OK_AGAIN="/etc/init.d/h
15 eartbeat start"
16
17 case $PARAMETER in
18   start) $LOG
19     ${ACTION_FAIL_START} ;
20
21     ${ACTION_FAIL_START} ;;
22   stop) $LOG
23     ${ACTION_OK_AGAIN} ;
24
25     ${ACTION_OK_AGAIN}
26 ;;
27 esac
28
29 exit 0

```

#### INFO

- [1] RFC 2401: <http://www.ietf.org/rfc/rfc2401.txt>
- [2] Freeswan: <http://www.freeswan.ca/code/super-freeswan/>
- [3] RFC 1745: <http://www.ietf.org/rfc/rfc1745.txt>
- [4] SSPE: <http://sspe.sourceforge.net>
- [5] Linux-HA: <http://www.linux-ha.org>

#### THE AUTHOR

Johannes Hubertz has been a Linux fan since 0.99.2, mainly because Linux allows him to get down to bit level if he wants to. On a professional level, Johannes has had fun managing a few dozen machines for one of Europe's biggest IT enterprises since 1996.