Disaster Recovery Made Simple

Stephen Walsh
Web Systems Administrator
ACU National

Questnet 2007

Introduction

Who I Am

Define the Recovery

What services do your clients what?

Email
Web applications
Internet Access
Staff Services

What your clients really want

IM Messengers
webmail
flickr galleries
LOLcats

Define the Recovery

What do you need to provide those services

DNS
DHCP
Web servers & Applications
Connectivity

Connectivity

BGP can be used with HSRP to provide a “multi-homed” capacity.

Ie - * All outbound traffic sourced from hosts on network 192.168.21.0/24 and destined to the Internet must be routed through Router 1 to ISP-A. However, if that link fails or Router 1 fails, all outbound traffic must be rerouted through Router 2 to ISP-B (and then to the Internet) without manual intervention.

*All inbound traffic destined to an autonomous system, AS 100, from the Internet must be routed by way of R1. In the event that the link from ISP-A to Router 1 fails, the inbound traffic must automatically be rerouted through ISP-B to Router 2.

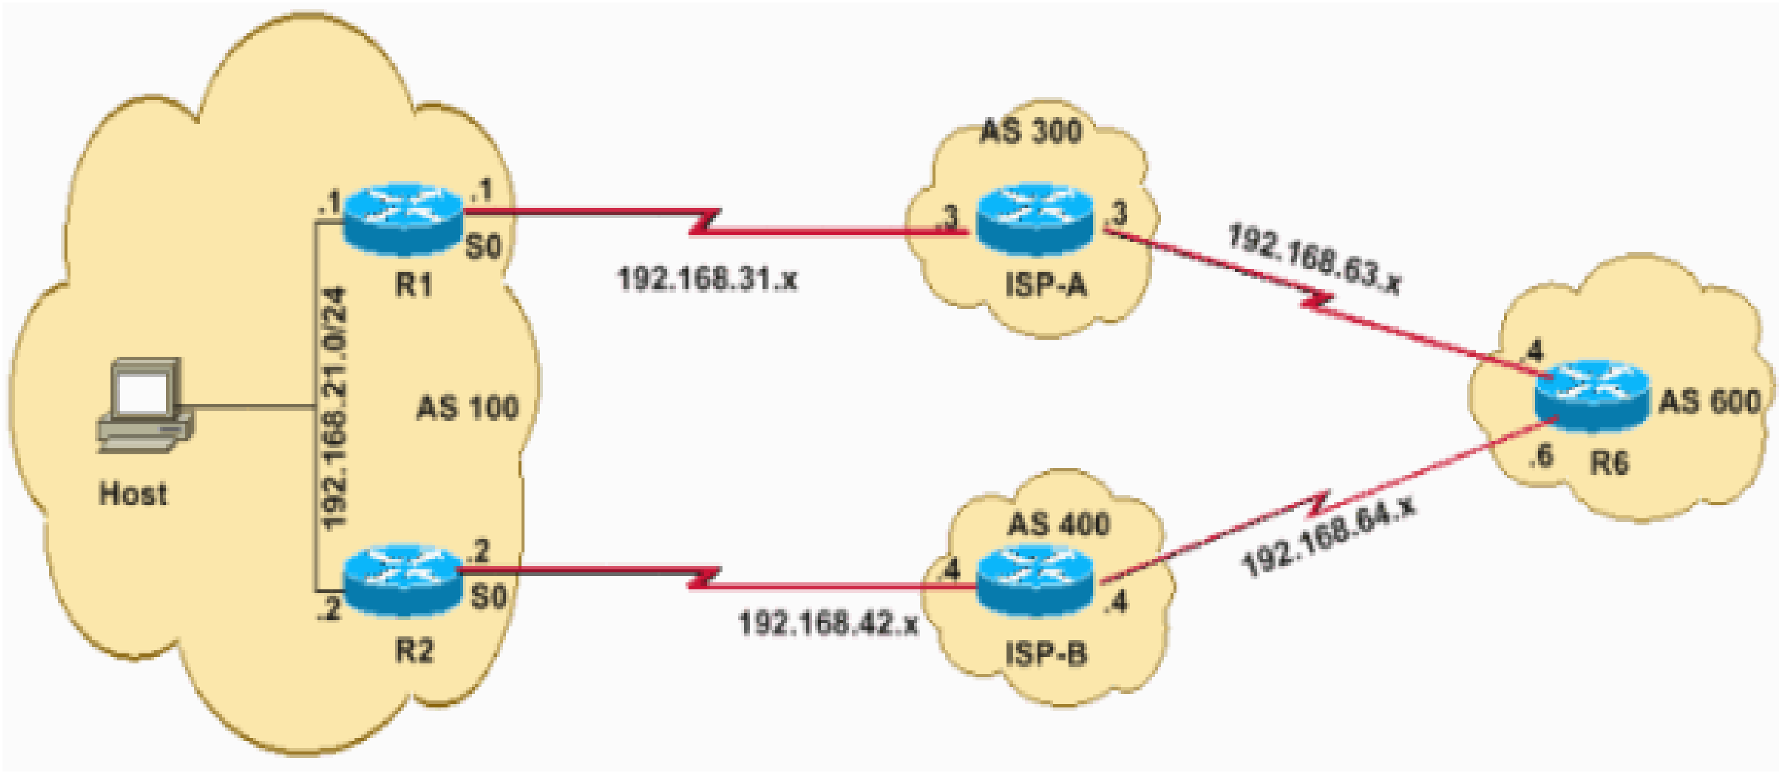
AS 300
.1
.1
.3
.3
S0
R1
192.168.31.x
ISP-A
192.168.63.x
192.168.21.0/24
.4
AS 100
AS 600
Host
.6
R6
AS 400
192.168.64.x
.2
.4
S0
.4
.2
192.168.42.x
R2
ISP-B

Router 1 Config

```
Current configuration
hostname Router1
!
interface serial 0
ip address 192.168.31.1 255.255.255.0
!
interface Ethernet1
  ip address 192.168.21.1 255.255.255.0
  standby 1 priority 105
  standby 1 preempt delay minimum 60
  standby 1 ip 192.168.21.10
  standby 1 track Serial0
!
!
router bgp 100
  no synchronization
  network 192.168.21.0
  neighbor 192.168.21.2 remote-as 100
  neighbor 192.168.21.2 next-hop-self
  neighbor 192.168.31.3 remote-as 300
  no auto-summary
!
```

Router 2 config

```
hostname  Router2
!
interface serial 0
ip address 192.168.42.2 255.255.255.0
!
interface Ethernet1
 ip address 192.168.21.2 255.255.255.0
 standby 1 priority 100
 standby 1 preempt
 standby 1 ip 192.168.21.10
!
!
```

```
router bgp 100
 no synchronization
 network 192.168.21.0
 neighbor 192.168.21.1 remote-as 100
 neighbor 192.168.21.1 next-hop-self
 neighbor 192.168.42.4 remote-as 400
 neighbor 192.168.42.4 route-map foo out
!--- It appends AS 100 to the BGP updates sent to AS 400
!--- in order to make it a backup for the ISP-A to R1 path.
 no auto-summary
!
access-list 1 permit 192.168.21.0
route-map foo permit 10
 match ip address 1
 set as-path prepend 100
end
```

Web server recovery with High Availability and Clustering

The High Availability (HA) tools runs a “heartbeat” client to determines when the requirements have been met for a new “master” node to take over the public IP.

A single Public IP is presented, but can change from host to host within the cluster as needed. This creates an effective “cloud” of private hosts, but with a single publically presented IP address.

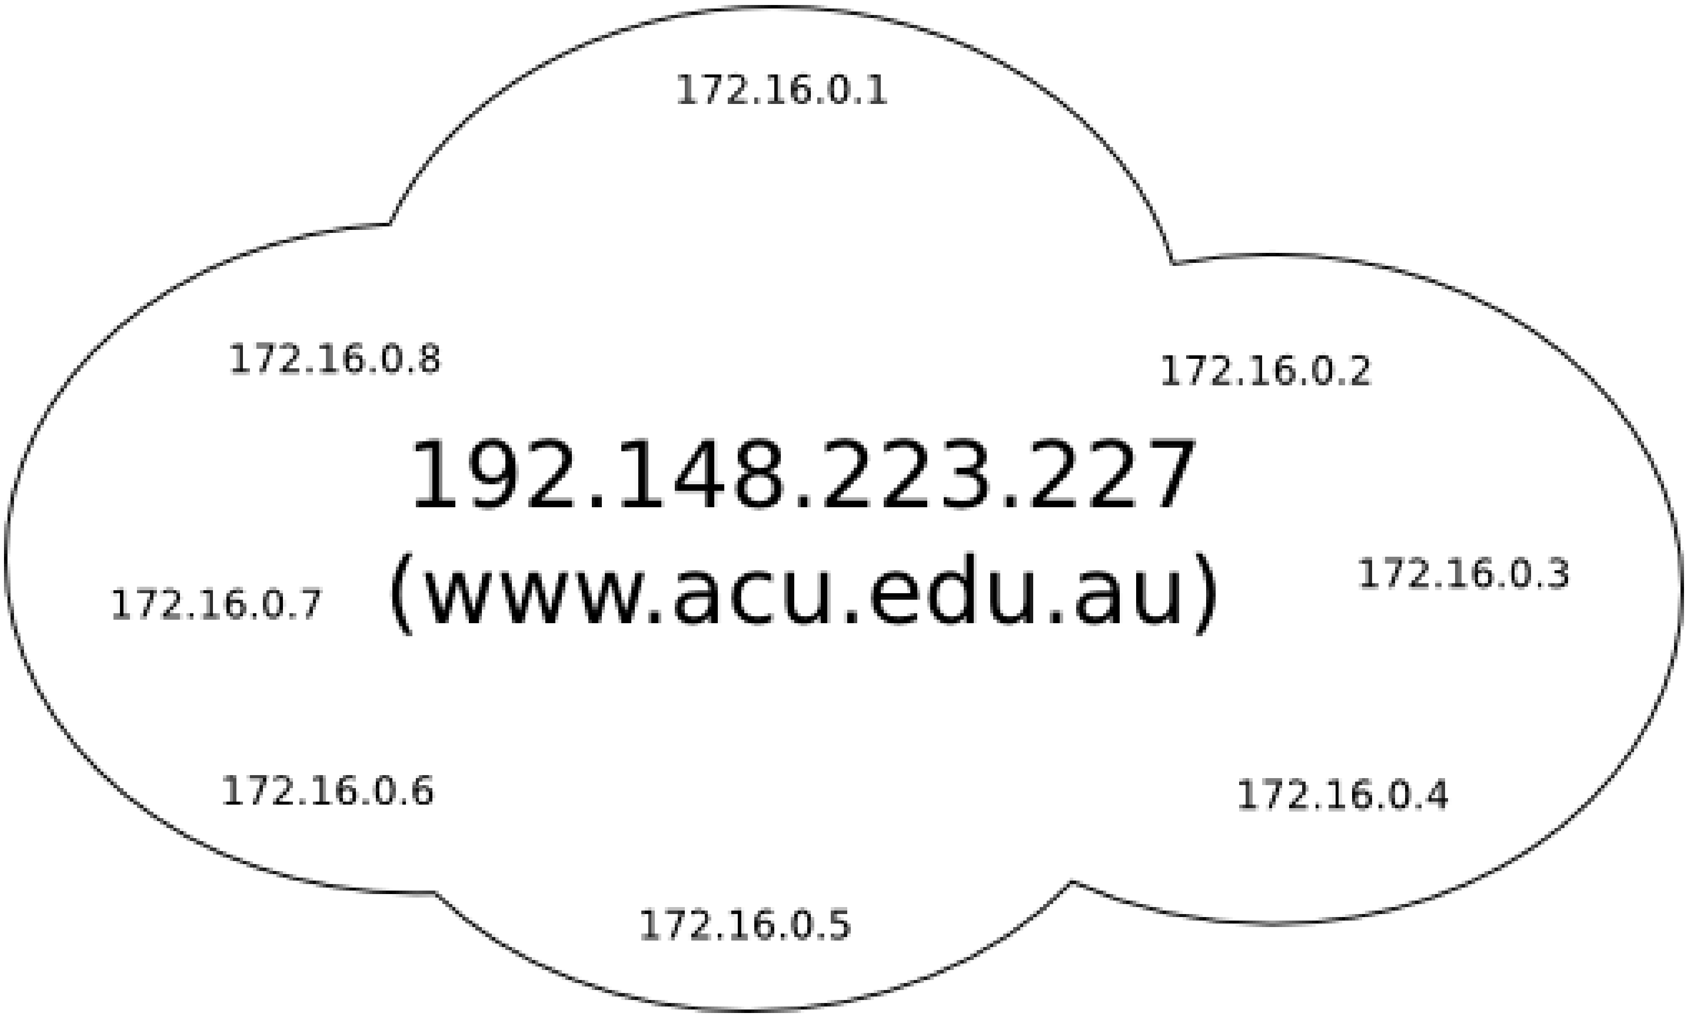
172.16.0.1
172.16.0.8
172.16.0.2
192.148.223.227
(www.acu.edu.au)
172.16.0.7
172.16.0.3
172.16.0.6
172.16.0.4
172.16.0.5

HA Configuration

Recovering from DHCP failure

Just bring another DHCP server online in its place, but the information about leases will be lost, forcing clients to acquire new addresses.

In that situation, some clients would have to break any existing network connections, and in some cases, local X-windows sessions would also break.

If you’re bored sometime, try changing the hostname of a *nix machine when running a live X desktop.

The recovery process can be amusing, and is a great way to put off finishing the house extension or the retaining wall.

DHCP

Alternatively, we could plan for a downtime by increasing lease times from 30 minutes to the better part of a day.

That would reduce—but not completely remove—the risk of any given client having its lease expire while the server is off-line, but any newly arriving client won’t get an address.

This is where the Fail-over capability in the ISC DHCP server comes in useful.

Failover needs “ip_helper” enabled on Cisco Routers, allowing DHCP requests to be sent to both hosts.

Configuration
Master (192.168.200.2)

```
failover peer "dhcp-failover" {
        primary; # declare this to be the primary server
        address 192.168.200.2;
        port 520;
        peer address 192.168.200.3;
        peer port 520;
        max-response-delay 30;
        max-unacked-updates 10;
        load balance max seconds 3;
        mclt 1800;
        split 128;
        }
..
        pool {
                failover peer "dhcp-failover";
                max-lease-time 1800; # 30 minutes
                range 192.168.200.100 192.168.200.254;
            }
```

Configuration
Slave (192.168.200.3)

```
failover peer "dhcp-failover" {
        secondary; # declare this to be the secondary server
        address 192.168.200.3;
        port 520;
        peer address 192.168.200.2;
        peer port 520;
        max-response-delay 30;
        max-unacked-updates 10;
        load balance max seconds 3;
        }

..
        pool {
            failover peer "dhcp-failover";
            max-lease-time 1800; # 30 minutes
            range 192.168.200.100 192.168.200.254;
            }
```

Failover DHCP notes

SE Linux prevents DHCP from binding to the required failover port. Create an appropriate context if using SE Linux.

NTP is Critical with failover, the logs <u>must</u> match for auditing purposes

When Primary fails, logs will show;
Jun 6 09:50:51 secondary dhcpd: failover peer dhcp-failover: I move from normal to communications-interrupted

When the Primary Recovers, Logs will Show;
Jun 6 09:51:37 secondary dhcpd: failover peer dhcp-failover: I move from communications-interrupted to normal

As the servers balance the pool, occasional log entries are written;
Jun 6 02:27:09 secondary dhcpd: pool 98e82b8 192.168.200.0/24 total 155 free 38 backup 37 lts 0