NapAfrica Connection Guidelines

Assistance for networks connecting to NapAfrica

Monday, 08 February 2010

Version 0.1

Please email info@napafrica.com if you would like to speak to someone regarding a connection.

Author: Andy Davidson

Contents

Version History:

Version	Author	Date	Notes
0.1	Andy Davidson	8th Feb 2010	Initial Version

Emergency Information.

If you require urgent support from NapAfrica, please contact support@napafrica.com or telephone +2721 2000 230

Introduction

NapAfrica is a neutral peering point, also known as an Internet Exchange Point. Its function is to allow Internet Service Providers, Hosting companies, and other networks to exchange internet traffic between their organisations.

Connecting your network has many advantages, including :

- **Reduced transit costs** – you may find other networks on NapAfrica are willing to peer with you on a 'settlement free' basis, meaning that two networks give each other free access to each others' network for mutual advantage.
- **Improved Access speed** – a connection at 1Gbit to NapAfrica is a superfast connection to nearby networks, meaning your network and customers will enjoy much higher capacity to networks you peer with.
- **Fault Tolerance** – peering means that there are multiple points of entry and exit to your network, which helps you reduce the effect out outages on your internet connections, plus peering with local providers means that any problems on your international circuits will not break local service.
- **Reduced latency** – peering reduces latency between peers, so the quality of internet voice and video services will improve for your customers.

NapAfrica is a stable Internet Exchange platform, because we have built a set of technical guidelines which are designed to make the exchange as robust as possible. This document explains what these guidelines are, and offers sample configuration to help you meet these guidelines.

The configuration examples in this document will help you keep your network secure and reliable, so that you can enjoy the benefits of peering without risk.

NapAfrica is also a neutral facility, which means that each network connected to the exchange will be treated equally and fairly. We will run a reliable Internet Exchange point, for the equal benefit of all participants.

Connection Process

It is simple to connect to NapAfrica.

If you already have routers installed in Teraco then the port fee is waived. The connection procedure is as follows:

Process	Timeline
Complete the connection form at http://www.napafrica.net/ You will need your AS Number to hand to complete this form.	20 minutes
Order a Cross-Connect to the NapAfrica switch. This should be raised as a work request to your account manager.	1-2 working days
Connect the Cross-Connect cable to your router. Please leave the port shutdown until you have been assigned an IP address.	0 days
Test the connection by assigning a temporary address to your router in our quarantine VLAN. NapAfrica engineers will inspect the traffic from your router to ensure that the technical guidelines have been met. *Example configuration for your router is available later in this document.*	1 working day
Assign a Peering LAN IP address, which will be communicated to you by the NapAfrica engineer after you have passed Quarantine. You should then configure a BGP session to the **IXP Collector** which is used for debugging and testing.	20 minutes
If you have an open peering policy then you can configure a session to the **NapAfrica Route Servers**. This will give you automatic peering with all other networks connected to the Route Servers. We will invite you to do this when the connection process is complete.	20 minutes

If you do not have equipment present in the Teraco data centres, then a small port fee will apply. You should order a connection between your router's current location, and the NapAfrica switch. Because NapAfrica is neutral, then you are welcome to place an order with any fibre provider.

When the connection from your router to the NapAfrica switch has been installed, then the connection process is similar to the above – you will be first provisioned in the quarantine VLAN, at which time your connection to the exchange will be tested to ensure that it complies with the guidelines listed in this document, and then you will be assigned a production peering address.

Technical Guidelines

These guidelines are important. Everyone following the same rules will ensure that the Internet Exchange is reliable and safe.

It is simple to follow the guidelines in this document, simply take a look at the example configurations listed later in this document. If you wish to check your configuration before connecting, or at any time after connection, please get in touch with NapAfrica technical support, and we will be happy to advise.

- Connected ports must be Ethernet. 100Mbit ports should be configured with a fixed 100Mbit speed, and Full Duplex. Gigabit Ethernet ports can configured as auto-sensing or manual, please tell is what you prefer.
- Any intermediate devices between the NapAfrica switch and your router, for example an aggregation switch, a third party Ethernet transport provider, or media conversion devices must not emit any traffic towards the exchange.
- By default, participants will be assigned an access port on the public peering VLAN. Multiple VLAN ports are available, and we will use 802.1q VLAN tags (ethertype 0x8100) to signal which VLAN the frame is in. You should not send frames with VLANs that are tagged for VLANs which we have not configured on your port. A change of service from a single VLAN port to a port supporting 802.1q tagging will cause a small interruption to your service.
- We recommend that interfaces on an access port have 1500 byte MTUs. Ports with 802.1q VLAN tagging can be 1516 so that the containing frame does not need to be fragmented.
- **All frames from a single port, which are forwarded to the NapAfrica exchange must have the same source MAC-address**. If this is not configured, then port-security on the NapAfrica switch will shut down your peering port. This technique is the most important loop prevention technology used on NapAfrica, so exceptions are not permitted.
- On the public peering LAN, only ethertypes 0x0800 (IPv4), 0x08dd (IPv6) and 0x0806 (ARP) are permitted. If you wish to exchange other types of packets with connected participants, then a Closed User Group VLAN can be created between the participants who wish to exchange other types of traffic.
- Connected networks should disable Proxy ARP on the router interface connected to NapAfrica.
- Connected networks should disable non permitted link-local protocols on the router interface connected to NapAfrica. Illegal protocols include DHCP, Spanning Tree, DEC MOP, CDP, VTP and Layer 2 Keepalive packets. The only permitted link-local protocols are ARP and IPv6 Neighbour Discovery.
- Connected networks must only use BGP to exchange routing information. This explicitly means that other routing protocols such as OSPF, ISIS, EIGRP, and IPv6 Router Solicitation should never be used. **It is always forbidden to point any static routes at other exchange participants.**
- NapAfrica recommend against overloading your ports, and we will get in touch when traffic at 95th percentile exceeds 75% of the port capacity. When traffic reaches 50% of your port

capacity, then NapAfrica will discuss suitable upgrade options (higher port capacity or aggregated links)

- Networks using aggregated ports will follow 802.3ad specifications. The aggregated links must be of the same media type and link speed.
- All connected networks must BGP peer with the Internet Exchange's collector. This session allows NapAfrica to check the health of the Internet Exchange.
- Connected networks should not export the NapAfrica peering LAN address space to other networks without permission.

NapAfrica are always able to help with your peering port configurations.

If we notice that any forbidden traffic is reaching the exchange from your router, we will always contact you with advice in the first instance. However, if we see traffic from you that could damage the stability of the exchange, then we may be required to move your port from the peering LAN into the quarantine LAN whilst we diagnose the issue with you.

Interface Configuration Guides

NapAfrica are here to make peering easy! Our guidelines above translate into very simple configuration that you can deploy on your routers. Please look over our suggested configuration.

Cisco Router connected to NapAfrica

Here is an example typical Cisco router configuration for your interface. Use this configuration if you have a single VLAN port, and plug your router directly into the NapAfrica switch (this is the recommended connection method).

```
Interface GigabitEthernet0/1
      ip address x.x.x.x y.y.y.y
      description PEERING:: NapAfrica Internet Exchange
      no ip redirects
      no ip proxy-arp
      no cdp enable
      no ip directed-broadcast
      no mop enable
      no keepalive
```

If you run a multiple VLAN port, then you should use the same properties on the subinterface which is created to communicate with the Peering VLAN. You should use 802.1q encapsulation for these extra interfaces, e.g :

```
Interface GigabitEthernet 0/1.300
      encapsulation dot1Q 300
      ip address x.x.x.x y.y.y.y
```

Do not worry about Multiple VLAN ports – you will be setup as a single VLAN port user unless you specifically ask for multiple VLAN support, in order to keep your configuration simple, and reduce the risk of mistakes.

Cisco Router, connecting via an intermediate switch.

This configuration can be used if you are short of router ports, and need to use a switch to add ports to your edge. It is not the recommended configuration, because this setup can cause frames from the intermediate switch to leak to the exchange.

There is a very wide variety of Cisco devices which can act as intermediate switches. This configuration assumes you are running IOS on your switch. If a command in this sample configuration is rejected by the switch, then it may be safe to assume the feature is not available on your switch and the line is not needed.

Intermediate Switch:

```
Interface GigabitEthernet0/2
      description Link facing NapAfrica Internet Exchange
      switchport access vlan 300
      switchport mode access
      switchport nonegotiate
      no keepalive
      no udld enable
      no cdp enable
      no lldp receive
      no lldp transmit
      spanning-tree bpdufilter enable
end
vlan 300
      name NAPAFRICA


Interface GigabitEthernet0/1
      description Link facing my router
      switchport mode trunk
      switchport trunk allowed vlan 300
```

Router:

```
Interface GigabitEthernet0/0
      description: Link facing intermediate switch
      no ip address

Interface GigabitEthernet0/0.300
      description PEERING:: NapAfrica
      encapsulation dot1q 300
      ip address x.x.x.x y.y.y.y
```

Please configure the Cisco subinterface rules in the same way as the router port in the first example (i.e. disable proxy-arp, cdp, etc)

Juniper Router

The Juniper configuration is a little smaller, based on the router not running some of these protocols by default.

The recommendations we make, increase the logging depth for your interfaces, which will improve debugging in the future, and turn off prohibited protocols across the NapAfrica exchange.

```
interfaces {
    traceoptions {
        file interfaces.log size 10m files 20;
        flag change-events;
        flag config-states;
    }

    ge-0/0/0 {
        description "PEERING:: NapAfrica Internet Exchange Point”;
        unit 0 {
            family inet {
                no-redirects;
                address x.x.x.x/24;
            }
        }
    }
```

If you run a multiple vlan tag port, then the configuration for your interface would change to this:

```
    ge-0/0/0 {
        description "PEERING:: NapAfrica Internet Exchange Point”;
        vlan-tagging;
        unit 300 {
            description “PEERING:: NapAfrica Peering VLAN”
            vlan-id 300;
            family inet {
                no-redirects;
                address x.x.x.x/23;
            }
        }
    }
```

You then create additional units underneath the interface for each additional VLAN over the exchange.

BGP Session Configuration Guidelines

Direct Sessions

It is strictly forbidden to use any non-BGP method for exchanging prefixes over the exchange. Essentially this means :

- You must never point default or other static routes at any other exchange participant.
- You must disable OSPF, IS-IS, RIP, etc., on interfaces facing the internet exchange.

Configuring BGP sessions involves seeking the agreement to “peer” with other providers on the exchange. Typically, you ask a person responsible for peering at the other network, and you both configure a BGP session that connects each others’ routers over the internet exchange.

This BGP session should only exchange the network prefixes of your own network, and your customers, unless otherwise agreed, so you should consider filters which are similar to the example shown below:

```
router bgp aaaaa
     router-id b.b.b.b
     bgp log-neighbor-changes
     neighbor peer-napafrica peer-group
     neighbor peer-napafrica next-hop-self
     neighbor peer-napafrica soft-reconfiguration inbound
     neighbor peer-napafrica prefix-list bogons in
     neighbor peer-napafrica route-map napafrica-in in
     neighbor peer-napafrica route-map napafrica-out out
     neighbor x.x.x.1 remote-as yyyy
     neighbor x.x.x.1 description PEER: XYZisp
     neighbor x.x.x.1 peer-group peer-napafrica
     neighbor x.x.x.1 maximum-prefix 100 restart 15

ip prefix-list bogons seq 10 deny 10.0.0.0/8 le 32
ip prefix-list bogons seq 20 deny 172.16.0.0/12 le 32
ip prefix-list bogons seq 30 deny 192.168.0.0/16 le 32
ip prefix-list bogons seq 40 deny 169.254.0.0/16 le 32
ip prefix-list bogons seq 50 deny 192.0.2.0/24 le 32
ip prefix-list bogons seq 60 deny 224.0.0.0/4 le 32
ip prefix-list bogons seq 70 deny 240.0.0.0/4 le 32
ip prefix-list bogons seq 80 deny [myprefixes] le 32
ip prefix-list bogons seq 99 permit 0.0.0.0/0 le 24

ip prefix-list sendout seq 10 permit [myprefix]

route-map napafrica-out permit 10
     match ip address prefix-list sendout

route-map napafrica-in permit 10
     set local-preference 999
```

You should change the network statements and filters to suit the configuration which you have in place. If you have BGP downstream customers, then you should add them to your ‘sendout’ prefix list, or use communities to signal to your router which prefixes are eligible for sending to your peers. It is also good practice to build rules which test for a specific observed as adjacency.

Multilateral peering sessions

Multilateral peering allows you to set up a single peering session, over which you see many other networks.

The configuration is slightly different because the route-server does not participate in the forwarding of traffic, and the route-server also passes through the prefixes of many different ASNs in a single session.

You can use many advanced filtering techniques to signal the behaviour on the route-server, but in this example we will setup a simple session with no filtering of your prefixes. This is the right configuration for most networks.

```
router bgp aaaaa
     no bgp enforce-first-as
     neighbour peer-napafrica-rs peer-group
     neighbor peer-napafrica-rs next-hop-self
     neighbor peer-napafrica-rs soft-reconfiguration inbound
     neighbor peer-napafrica-rs prefix-list bogons in
     neighbor peer-napafrica-rs route-map napafrica-in in
     neighbor peer-napafrica-rs route-map napafrica-rs-out out
     neighbor peer-napafrica-rs remote-as rrrrr
     neighbor x.x.x.2 peer-group peer-napafrica-rs
     neighbor x.x.x.3 peer-group peer-napafrica-rs

route-map napafrica-rs-out permit 10
     match ip address prefix-list sendout
     set community zzz:Zzz
```

You will see that we send community zzz:zzz to the route server. This is a special signal that tells the route-server to send the prefix to all other participants. Note that there is a new peer-group. This allows you to shut down route-server sessions in the event of a problem without affecting your direct peering, and is also necessary because we use a different route-map for outgoing sessions.

When you connect to NapAfrica, the support team will contact you to discuss route-server peerings, and whether they are a suitable system