

# CISCO BGP router configuration for Internet Exchange Connection

by [Thomas Mangin](#)

Posted 02/06/2003

Updated 28/07/2003

## Prelude

In order to minimise the need for experimentation on your network when connecting to an Internet exchange, here is an example of an external BGP router which will provide you with an overall view of a real-world system.

Please feel free to [report](http://thomas.mangin.me.uk/email.html) any possible improvements, inaccuracies or faults contained in this document.

## Information

The article is based on the following virtual network. The routers are owned by an ISP with its own AS number. The IGP is EIGRP as it is the protocol I am most familiar with (it could be OSPF or IS-IS).

### Servers

10.0.0.1	Primary caching DNS
10.1.0.1	Secondary caching DNS
10.2.3.200	SNMP Monitoring station
10.2.3.201	SYSLOG logging server
10.2.3.202	NTP server
10.2.3.205	Netflow monitoring server

### IP Range

10.0.0.0/16	AS 65200
10.1.0.0/16	AS 65200
10.2.3.0/24	AS 65200

### The internet exchange

The exchange Information is running a dual ring topology with a separate /23 for each

172.16.0.0/23	AS 65400
172.16.2.0/23	AS 65400

# Router configuration

First a generic router configuration.

! Help telnet connection

**service nagle**  
**no service pad**

! Deal with dead connections gracefully

**service tcp-keepalives-in**  
**service tcp-keepalives-out**

! Logging information structure

**service timestamps debug datetime msec localtime show-timezone**  
**service timestamps log datetime msec localtime show-timezone**

**service slave-log**

**service password-encryption**

! Remove all useless services

**no service compress-config**  
**no service udp-small-servers**  
**no service tcp-small-servers**

**no service config**

**no service dhcp**

**no ip bootp server**

**no ip finger**

**no ip identd**

! Only use if you are really concerned about the router physical security

! no service password-recovery

! Router name

**hostname my\_router**

! Log in memory and not to console

**logging buffered 16384 errors**  
**no logging console**

! create user (for telnet and console login)

**username *the\_login\_name* password 0 *the\_password***

! set a password for the privileged mode

**enable secret 0 the\_enable\_password**

! Set time for UK

**clock timezone GMT 0**

**clock summer-time BST recurring**

! Allow use of all subnet

**ip subnet-zero**

! Do not allow packet to specify their own route

**no ip source-route**

! Enable Cisco Express Forwarding technology

**ip cef**

! Do not do any host lookup but configure it should we want it

**no ip domain-lookup**

! But make sure it is configured, should we want it

**ip domain-list [isp.net.uk](http://isp.net.uk)**

**ip domain-list .**

**ip domain-name [isp.net.uk](http://isp.net.uk)**

**ip name-server 10.0.0.1**

**ip name-server 10.1.0.1**

! Should be the default ...

**ip classless**

**ip routing**

! Should we not know a route send it to the IGP router which have a route for 10.0.1.0

**ip default-network 10.0.1.0**

! Do not allow http management

**no ip http server**

! Log all the information to a remote syslog server

**logging trap debugging**

**logging facility local6**

**logging source-interface Loopback0**

**logging 10.2.3.201**

! Only the snmp server IP

**access-list 20 permit 10.2.3.200**

! Match ICMP traffic

**access-list 110 permit icmp any any echo**

**access-list 110 permit icmp any any echo-reply**

! Allow snmp monitoring from the snmp server only

**snmp-server community snmp\_community\_password RO 20**

**snmp-server host 10.2.3.200 snmp\_community\_password**

! Welcome banner when telneting to the router

**banner login ^C**

\*\*\*\*\*

## **NOTICE TO USERS**

**This equipment is for authorized use only. Users (authorized or unauthorized) have no explicit or implicit expectation of privacy.**

**Any or all uses of this system and all files on this system may be intercepted, monitored, recorded, copied, audited, inspected, and disclosed to authorized site and law enforcement personnel.**

**By using this system, the user consents to such interception, monitoring, recording, copying, auditing, inspection, and disclosure at the discretion of authorized site.**

**Unauthorized or improper use of this system may result in administrative disciplinary action and civil and criminal penalties. By continuing to use this system you indicate your awareness of and consent to these terms and conditions of use.**

**LOG OFF NOW if you do not agree to the conditions stated in this warning.**

**ISP - noc@isp.net.uk - Phone number : 00 44 .....**

\*\*\*\*\*

**^C**

! Protect our router asking for username and password and then enable password

**line con 0**

**login local**

**line aux 0**

**login local**

**transport input all**

**transport output none**

**line vty 0 4**

! One hour timeout is not very secure but much more practical ..

**exec-timeout 60 0**

**login local**

! Keep a higher command history

**history size 256**

! Keep the router time correct

**ntp server 10.2.3.202**

! Undocumented command to improve the speed at which BGP routes are learned

**spd headroom 1000**

! Loopback interfaced used to source the IP the the packet

**interface Loopback0**

**ip address 10.2.3.14 255.255.255.255**

! See [http://www.cisco.com/en/US/products/sw/iosswrel/ps5187/products\\_command\\_reference\\_chapter09186a008017d163.html](http://www.cisco.com/en/US/products/sw/iosswrel/ps5187/products_command_reference_chapter09186a008017d163.html)

**no ip redirects**

**no ip unreachable**

**no ip proxy-arp**

**no ip directed-broadcast**

**no ip unreachable**

**interface FastEthernet0/0**

**description "ISP Backbone"**

**ip address 10.0.0.254 255.255.255.0**

**no ip unreachable**

**no ip proxy-arp**

**no ip directed-broadcast**

! Allow netflow monitoring of the interface

**ip route-cache flow**

! Force speed and duplex (always a good idea)

**speed 100**

**full-duplex**

! Should you need to use this interface as hosts' gateway

! ip route-cache same-interface

**interface FastEthernet1/0**

**description "Primary Exchange Connection"**

**ip address 172.16.0.100 255.255.254.0**

! We do not want to send or received spoofed packets

**ip access-group network\_isp\_in in**

**ip access-group network\_isp\_out out**

! See cisco site for more details

**no ip redirects**

**no ip unreachable**

**no ip proxy-arp**

**no ip directed-broadcast**

**no ip mask-reply**

**ip verify unicast reverse-path**

! Limit the amount of ICMP traffic (DDOS protection)

**rate-limit input access-group 110 2048000 8000 8000 \**

**conform-action transmit exceed-action drop**

! Allow netflow monitoring of the interface

**ip route-cache flow**

! Force speed and duplex (always a good idea)

**speed 100**

**full-duplex**

! You do not want to have CDP running to the exchange

! Read your exchange MoU :-)

**no cdp enable**

! Improve the speed at which we learn BGP routes

**hold-queue 1500 in**

! Whatever your IGP, Here EIGRP

**router eigrp 65200**

! Let in IGP static and connected routes

**redistribute connected**

**redistribute static**

! Do not send or receive IGP packet to the exchange

**passive-interface FastEthernet1/0**

! Our network

**network 10.2.2.0**

! Do not advertise the default route

**distribute-list 10 out**

! Read more about EIGRP :-)

**no auto-summary**

**eigrp log-neighbor-changes**

! Here we are with BGP !!  
**ip bgp-community new-format**

**router bgp 65200**

! We want to advertise routes even if not in our IGP  
**no synchronization**

! See <http://www.ripe.net/ripe/docs/ripe-229.html>  
**no bgp fast-external-falover**

! Keep track of BGP neighbour flaps  
**bgp log-neighbor-changes**

! Use route dampening  
! If route are flapping, we should remove them  
**bgp dampening route-map graded-flap-dampening**

! Load balance across the two best routes  
**maximum-paths 2**

! If you see any network withing the 10.0.0.0/15, aggregate and announce the /15  
! It mean that one of your iBGP routers must have a network statement  
! otherwise you will not announce any routes to the exchange. ie:  
! Network 10.0.1.0  
**aggregate-address 10.0.0.0 255.254.0.0 summary-only**

! Define a internal BGP router profile  
**neighbor peer-ibgp peer-group**  
**neighbor peer-ibgp remote-as 65200**  
! Make sure the connection will never flap is we have more than one  
! interface/route to the peer.

**neighbor peer-ibgp update-source Loopback0**  
**neighbor peer-ibgp version 4**

! Propagate community information  
**neighbor peer-ibgp send-community**  
! Enables the storage of received route updates, regardless of the inbound policy  
**neighbor peer-ibgp soft-reconfiguration inbound**

! We like those neighbor and we want to use them  
**neighbor peer-nap-high peer-group**  
**neighbor peer-nap-high description "Peering (High Preference)"**  
**neighbor peer-nap-high version 4**

**neighbor peer-nap-high next-hop-self**  
**neighbor peer-nap-high send-community**  
**neighbor peer-nap-high soft-reconfiguration inbound**

! This allow to prefer a router over another to send traffic in or out.

**neighbor peer-nap-high route-map peer-nap-high-in in**  
**neighbor peer-nap-high route-map peer-nap-out out**

! If we are learning more than 100 route kill the peering relation (route leak)

**neighbor peer-nap-high maximum-prefix 100**

! We do not like those peer as much as the previous ones

```
neighbor peer-nap-low peer-group
neighbor peer-nap-low description "Peering (Low Preference)"
neighbor peer-nap-low version 4
neighbor peer-nap-low next-hop-self
neighbor peer-nap-low send-community
neighbor peer-nap-low soft-reconfiguration inbound
neighbor peer-nap-low route-map peer-nap-low-in in
neighbor peer-nap-low route-map peer-nap-low-out out
neighbor peer-nap-low maximum-prefix 100
```

! We do not want to learn bogon-network or our own IP range from the internet,  
! <http://www.cymru.com/Documents/bogon-dd.html>

```
neighbor transit-london prefix-list bogon-external in
! And we do not want to announce bogon-networks neither
neighbor transit-london prefix-list bogon-internal out
```

! This is the full list of our iBGP routers

```
neighbor 10.2.3.10 peer-group peer-ibgp
neighbor 10.2.3.11 peer-group peer-ibgp
```

! Peering connection with another ISP (we like a lot)

```
neighbor 172.16.0.10 remote-as 65300
neighbor 172.16.0.10 peer-group peer-nap-high
neighbor 172.16.0.10 description "ISP One Primary"
neighbor 172.16.0.10 maximum-prefix 1000
```

```
neighbor 172.16.2.10 remote-as 65300
neighbor 172.16.2.10 peer-group peer-nap-high
neighbor 172.16.2.10 description "ISP One Secondary"
neighbor 172.16.2.10 maximum-prefix 1000
```

! Peering connection with another ISP (we like less)

```
neighbor 172.16.0.20 remote-as 65310
neighbor 172.16.0.20 peer-group peer-nap-low
neighbor 172.16.0.20 description "ISP Two Primary"
```

```
neighbor 172.16.2.20 remote-as 65310
neighbor 172.16.2.20 peer-group peer-nap-low
neighbor 172.16.2.20 description "ISP Two Secondary"
```

! We are not here to aggregate the internet :-)  
**no auto-summary**

! We do not want to learn invalid information from our iBGP peers  
! This can be done using the distribute-list with the ibgp peer group as well  
**distribute-list prefix bogon-external out FastEthernet0/0**  
**distribute-list prefix bogon-internal in FastEthernet0/0**



```
! Export netflow information for traffic analysis
ip flow-export source Loopback0
ip flow-export version 5 peer-as
ip flow-export destination 10.2.3.205 2055
```

```
! Our own AS
ip as-path access-list 60 permit ^$
```

```
! All ASes
ip as-path access-list 50 permit .*
```

```
! Our own AS and the one of the customer we provide transit for.
ip as-path access-list 70 permit ^$
ip as-path access-list 70 permit ^65350_
ip as-path access-list 70 permit ^65360_
```

```
! First transit customer AS
ip as-path access-list 100 permit ^65350_
```

```
! Second transit customer
ip as-path access-list 105 permit ^65360_
```

```
! Never used this, should be off if not for the cisco IPv4 bug...
ip pim bidir-enable
```

```
! Deny all reserved AS, Allow all the valid ones
ip as-path access-list 10 deny _(6451[2-9]|645[2-9][0-9]|64[6-9][0-9][0-9]|
65[0-4][0-9][0-9]|655[0-2][0-9]|6553[0-5])_
ip as-path access-list 10 permit .*
```

! We should never see those range from our external BGP peers

```
ip prefix-list bogon-external seq 1 deny 0.0.0.0/0
ip prefix-list bogon-external seq 10 deny 0.0.0.0/7 le 32
ip prefix-list bogon-external seq 20 deny 2.0.0.0/8 le 32
ip prefix-list bogon-external seq 30 deny 5.0.0.0/8 le 32
ip prefix-list bogon-external seq 40 deny 7.0.0.0/8 le 32
ip prefix-list bogon-external seq 50 deny 10.0.0.0/8 le 32
ip prefix-list bogon-external seq 60 deny 23.0.0.0/8 le 32
ip prefix-list bogon-external seq 70 deny 27.0.0.0/8 le 32
ip prefix-list bogon-external seq 80 deny 31.0.0.0/8 le 32
ip prefix-list bogon-external seq 90 deny 36.0.0.0/7 le 32
ip prefix-list bogon-external seq 100 deny 39.0.0.0/8 le 32
ip prefix-list bogon-external seq 110 deny 41.0.0.0/8 le 32
ip prefix-list bogon-external seq 120 deny 42.0.0.0/8 le 32
ip prefix-list bogon-external seq 130 deny 49.0.0.0/8 le 32
ip prefix-list bogon-external seq 140 deny 50.0.0.0/8 le 32
ip prefix-list bogon-external seq 150 deny 58.0.0.0/7 le 32
ip prefix-list bogon-external seq 160 deny 60.0.0.0/8 le 32
ip prefix-list bogon-external seq 170 deny 70.0.0.0/7 le 32
ip prefix-list bogon-external seq 180 deny 72.0.0.0/5 le 32
ip prefix-list bogon-external seq 190 deny 83.0.0.0/8 le 32
ip prefix-list bogon-external seq 200 deny 84.0.0.0/6 le 32
ip prefix-list bogon-external seq 210 deny 88.0.0.0/5 le 32
ip prefix-list bogon-external seq 220 deny 96.0.0.0/3 le 32
ip prefix-list bogon-external seq 230 deny 169.254.0.0/16 le 32
ip prefix-list bogon-external seq 240 deny 172.16.0.0/12 le 32
ip prefix-list bogon-external seq 250 deny 192.0.2.0/24 le 32
ip prefix-list bogon-external seq 260 deny 192.168.0.0/16 le 32
ip prefix-list bogon-external seq 270 deny 197.0.0.0/8 le 32
ip prefix-list bogon-external seq 280 deny 198.18.0.0/15 le 32
ip prefix-list bogon-external seq 290 deny 201.0.0.0/8 le 32
ip prefix-list bogon-external seq 300 deny 222.0.0.0/7 le 32
ip prefix-list bogon-external seq 310 deny 224.0.0.0/3 le 32
ip prefix-list bogon-external seq 500 deny 159.101.0.0/16 le 32
```

! We do not want to learn about our own network from someone else

```
ip prefix-list bogon-external seq 510 deny 10.0.0.0/16 le 32
ip prefix-list bogon-external seq 520 deny 10.1.0.0/16 le 32
```

! Do not accept any route with a netmask smaller than /24

```
ip prefix-list bogon-external seq 900 permit 0.0.0.0/0 le 24
```

! We should never see those networks from inside our network

```
ip prefix-list bogon-internal seq 1 deny 0.0.0.0/0
ip prefix-list bogon-internal seq 10 deny 0.0.0.0/7 le 32
ip prefix-list bogon-internal seq 20 deny 2.0.0.0/8 le 32
ip prefix-list bogon-internal seq 30 deny 5.0.0.0/8 le 32
ip prefix-list bogon-internal seq 40 deny 7.0.0.0/8 le 32
ip prefix-list bogon-internal seq 50 deny 10.0.0.0/8 le 32
ip prefix-list bogon-internal seq 60 deny 23.0.0.0/8 le 32
ip prefix-list bogon-internal seq 70 deny 27.0.0.0/8 le 32
ip prefix-list bogon-internal seq 80 deny 31.0.0.0/8 le 32
ip prefix-list bogon-internal seq 90 deny 36.0.0.0/7 le 32
ip prefix-list bogon-internal seq 100 deny 39.0.0.0/8 le 32
ip prefix-list bogon-internal seq 110 deny 41.0.0.0/8 le 32
ip prefix-list bogon-internal seq 120 deny 42.0.0.0/8 le 32
ip prefix-list bogon-internal seq 130 deny 49.0.0.0/8 le 32
ip prefix-list bogon-internal seq 140 deny 50.0.0.0/8 le 32
ip prefix-list bogon-internal seq 150 deny 58.0.0.0/7 le 32
ip prefix-list bogon-internal seq 160 deny 60.0.0.0/8 le 32
ip prefix-list bogon-internal seq 170 deny 70.0.0.0/7 le 32
ip prefix-list bogon-internal seq 180 deny 72.0.0.0/5 le 32
ip prefix-list bogon-internal seq 190 deny 83.0.0.0/8 le 32
ip prefix-list bogon-internal seq 200 deny 84.0.0.0/6 le 32
ip prefix-list bogon-internal seq 210 deny 88.0.0.0/5 le 32
ip prefix-list bogon-internal seq 220 deny 96.0.0.0/3 le 32
ip prefix-list bogon-internal seq 230 deny 169.254.0.0/16 le 32
ip prefix-list bogon-internal seq 240 deny 172.16.0.0/12 le 32
ip prefix-list bogon-internal seq 250 deny 192.0.2.0/24 le 32
ip prefix-list bogon-internal seq 260 deny 192.168.0.0/16 le 32
ip prefix-list bogon-internal seq 270 deny 197.0.0.0/8 le 32
ip prefix-list bogon-internal seq 280 deny 198.18.0.0/15 le 32
ip prefix-list bogon-internal seq 290 deny 201.0.0.0/8 le 32
ip prefix-list bogon-internal seq 300 deny 222.0.0.0/7 le 32
ip prefix-list bogon-internal seq 310 deny 224.0.0.0/3 le 32

ip prefix-list bogon-internal seq 900 permit 0.0.0.0/0 le 24
```

! Those network contain root DNS servers,

! we do not want to ever loose them from our routing table

**ip prefix-list golden-networks description "root DNS server networks"**

**ip prefix-list golden-networks seq 100 permit 198.41.0.0/24**

**ip prefix-list golden-networks seq 105 permit 128.9.0.0/16**

**ip prefix-list golden-networks seq 110 permit 192.33.4.0/24**

**ip prefix-list golden-networks seq 115 permit 128.8.0.0/16**

**ip prefix-list golden-networks seq 120 permit 192.203.230.0/24**

**ip prefix-list golden-networks seq 125 permit 192.5.5.0/24**

**ip prefix-list golden-networks seq 130 permit 192.112.36.0/24**

**ip prefix-list golden-networks seq 135 permit 128.63.0.0/16**

**ip prefix-list golden-networks seq 140 permit 192.36.148.0/24**

**ip prefix-list golden-networks seq 145 permit 192.58.128.0/24**

**ip prefix-list golden-networks seq 150 permit 193.0.14.0/24**

**ip prefix-list golden-networks seq 155 permit 198.32.64.0/24**

**ip prefix-list golden-networks seq 160 permit 202.12.27.0/24**

**ip prefix-list golden-networks seq 165 permit 192.5.6.0/24**

**ip prefix-list golden-networks seq 170 permit 192.33.14.0/24**

**ip prefix-list golden-networks seq 175 permit 192.26.92.0/24**

**ip prefix-list golden-networks seq 180 permit 192.31.80.0/24**

**ip prefix-list golden-networks seq 185 permit 192.12.94.0/24**

**ip prefix-list golden-networks seq 190 permit 192.35.51.0/24**

**ip prefix-list golden-networks seq 195 permit 192.42.93.0/24**

**ip prefix-list golden-networks seq 200 permit 192.54.112.0/24**

**ip prefix-list golden-networks seq 205 permit 192.43.172.0/24**

**ip prefix-list golden-networks seq 210 permit 192.48.79.0/24**

**ip prefix-list golden-networks seq 215 permit 192.52.178.0/24**

**ip prefix-list golden-networks seq 220 permit 192.41.162.0/24**

**ip prefix-list golden-networks seq 225 permit 192.55.83.0/24**

! It does what is written on the tin

```
ip prefix-list max22-23 description Apply to /22 and /23 prefixes  
ip prefix-list max22-23 seq 5 permit 0.0.0.0/0 ge 22 le 23
```

```
ip prefix-list min24 description Apply to /24 and longer prefixes  
ip prefix-list min24 seq 5 permit 0.0.0.0/0 ge 24
```

! Match our RIPE allocated ranges

```
ip access-list standard network_isp_in  
deny 10.2.3.0 0.0.0.255 log  
deny 10.0.0.0 0.0.255.255  
deny 10.1.0.0 0.0.255.255  
permit any
```

! Match our RIPE allocated ranges and the one of our transit customers

```
ip access-list standard network_isp_out  
permit 10.2.3.0 0.0.0.255 log  
permit 10.0.0.0 0.0.255.255  
permit 10.1.0.0 0.0.255.255  
permit 192.168.0.0 0.0.0.255  
permit 192.168.1.0 0.0.0.255  
deny any
```

! All but default route

```
access-list 10 deny 0.0.0.0  
access-list 10 permit any
```

! Default route only

```
access-list 11 permit 0.0.0.0  
access-list 11 deny any
```

! Accept our syslog server only

```
access-list 20 permit 10.2.3.200
```

! Match ICMP traffic

```
access-list 110 permit icmp any any echo  
access-list 110 permit icmp any any echo-reply
```

**route-map peer-nap-high-in permit 10**

! Tell all IBGP routers to send us the traffic for those routes

**set ip next-hop peer-address**

! Indicate to IBGP routers that those route should be preferred (150 better than 145)

**set local-preference 150**

! Mark those route with a community

**set community 65200:65400**

**route-map peer-nap-low-in permit 10**

**set ip next-hop peer-address**

**set local-preference 145**

**set community 65200:65400**

**route-map peer-nap-out permit 10**

! only announce our and our transit customer route

! we are peering not offering a free transit service

**match as-path 70**

**set community 65200:65200**

! Never dampening networks with root servers

**route-map graded-flap-dampening deny 10**

**match ip address prefix-list golden-networks**

! Apply those settings (see RIPE 229 documents)

**route-map graded-flap-dampening permit 20**

**match ip address prefix-list min24**

**set dampening 30 820 3000 60**

**route-map graded-flap-dampening permit 30**

**match ip address prefix-list max22-23**

**set dampening 15 750 3000 45**

**route-map graded-flap-dampening permit 40**

**set dampening 10 1500 3000 30**

# Links

## BGP related links:

- [ISP IOS essentials](http://www.alliancedatacom.com/isp/IOS_Essentials_2-9.pdf), To be read  
[http://www.alliancedatacom.com/isp/IOS\\_Essentials\\_2-9.pdf](http://www.alliancedatacom.com/isp/IOS_Essentials_2-9.pdf)
- [ISP Security essential](http://www.cisco.com/public/cons/isp/security/B_Intro_v2-3.pdf), as well  
[http://www.cisco.com/public/cons/isp/security/B\\_Intro\\_v2-3.pdf](http://www.cisco.com/public/cons/isp/security/B_Intro_v2-3.pdf)
- [Secure BGP Template](http://www.cymru.com/Documents/secure-bgp-template.html), Advice on how to configure a cisco router  
<http://www.cymru.com/Documents/secure-bgp-template.html>
- [BGP Information](http://www.cisco.com/cgi-bin/Support/browse/psp_view.pl?p=Internetworking:BGP), The BGP section of the CISCO web site  
[http://www.cisco.com/cgi-bin/Support/browse/psp\\_view.pl?p=Internetworking:BGP](http://www.cisco.com/cgi-bin/Support/browse/psp_view.pl?p=Internetworking:BGP)
- [Nanag 25 Presentation](http://www.nanog.org/mtg-0206/ppt/philip/), BGP for Internet Service Providers  
<http://www.nanog.org/mtg-0206/ppt/philip/>
- [Bogon networks](http://www.cymru.com/Documents/bogon-dd.html), The list of routes you should never find in BGP  
<http://www.cymru.com/Documents/bogon-dd.html>
- [Golden Networks](http://www.golden-networks.net/), List of the root DNS servers  
<http://www.golden-networks.net/>
- [RIPE 229](http://www.ripe.net/ripe/docs/ripe-229.html), Explain route dampening  
<http://www.ripe.net/ripe/docs/ripe-229.html>
- [Undocumented IOS command](http://www.elemental.net/~lf/undoc/), List of non-completing command  
<http://www.elemental.net/~lf/undoc/>
- [Undocumented cisco IOS commands](http://www.i-n-t.de/ccie/ios_commands.html)  
[http://www.i-n-t.de/ccie/ios\\_commands.html](http://www.i-n-t.de/ccie/ios_commands.html)
- [RIPE documents](http://www.ripe.net/ripe/docs/alltitle.html), All available RIPE document  
<http://www.ripe.net/ripe/docs/alltitle.html>
- [More Resources](http://macross.dynodns.net/idr/), A good list of BGP related links  
<http://macross.dynodns.net/idr/>
- [RIPE whois](http://www.ripe.net/ripenncc/pub-services/db/whois/whois.html)  
<http://www.ripe.net/ripenncc/pub-services/db/whois/whois.html>

## Side notes

To see the routes learned from the exchange:  
***show ip bgp community 65200:65400***

To see how many spoofed packets have been stopped by the access list:  
***show ip access-lists network\_osp\_in | inc deny***  
***show ip access-lists network\_osp\_out | inc deny***

If you are using serial connections, use "no keepalive" on the interface.

You can find more information and other templates at  
<http://thomas.mangin.me.uk/networking.html>

## To do

- Fix all possible faults
- Limit telnet to loopback interface
- add scheduler allocate command for 7200+ router
- Add the RIPE object associated with this configuration
  - RIPE preference : lower is more favorable
  - Local Preference : higher is more favorable
- You name it ..