


BGP Configuration from the IRR

Cengiz Alaettinoglu
Packet Design Inc.
cengiz@packetdesign.com

<https://www.isi.edu/ra/gpa/brainin/>
<http://www.isi.edu/ra/BAToolBo/>

BGP Configuration

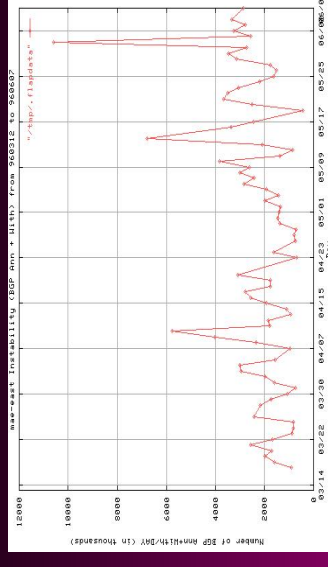


Too many routers
Too detailed, large & tedious
Consistency
Heavy consequences of mistakes


Mistakes don't happen!

BGP → RIP → BGP injection
128/7 leak
bogon 0/0, 10/8 leaks
daily, someone's leaking someone else's prefix

Consequences



Internet Routing Registry




Policy & contact information

Route: 128.9.0.0/16
descr: ISI-NET
origin: AS226
notify: Prue@isi.edu
mnt-by: LN-MAINT-MCI
changed: Prue@isi.edu 950420
source: CW

Internet Routing Registry

person: Walt Prue
address: USC/Information Sciences Institute
 4676 Admiralty Way Suite 1000
 Marina del Rey, California
 USA
phone: +1 310 822 1511 x89191
fax-no: +1 310 823 6714
e-mail: Prue@isi.edu
nic-hdl: WP8
notify: Prue@isi.edu
mnt-by: LN-MAINT-MCI
changed: Prue@isi.edu 20000222
source: CW

BGP Configuration from IRR




RPSL: Abstract, high level, per-as policies
IRR: Benefit from others' data & delegation
RiConfig: Details/redund aspects automated

Outline


- Motivation
- Routing Policy Definitions
- RPSL & RConfig
- Basic classes and notations
- Set classes
- import/export policy
- Structured Policy
- Other Features & Wrap-up

Internet's Routing




Interior routing is metric based
Inside MCI, shortest exit to GTE
OSPF, IS-IS, ...
Exterior routing is policy based
can ISI use GTE to reach MIT?
BGP

Inter-AS Topology




AS Relationships




Customer can use its resources
Provider prefer one provider over another
Peer can use its resources to reach customers

Typical routing policies implement these

Routing Policy



Not to/from a customer



Routing Policy

InterAS Routing

AS Level Peering



AS2 originates 128.9.0.0/16

AS2 exports 128.9.0.0/16 to AS1

AS1 imports 128.9.0.0/16 from AS2

BGP

Each route carries *path attributes*

Destination address prefixes

AS path

Originator AS

List of communities (flags)

Metrics: med, pref

Routing Policy

who are the peers

what routes are
originated
imported
exported
preferred

what to do when no route
what routes to aggregate

Outline

Motivation
Routing Policy Definitions
RPSL & RtConfig
Basic classes and notations
Set classes
import/export policy
Structured Policy
Other Features & Wrap-up

Routing Policy Specification Language (RPSL)

Object oriented language

route, autonomous system, router, contact and set objects
Defines the syntax/semantics/format of data in IRR

Vendor independent

Extensible

IETF Proposed Standard (RFC 2622)


Based on RIPE-181

RtConfig

Router configuration tool

cisco, juniper, bay, gated, rsd

user-controlled level of features/publishing



RPSL Classes

route
autonomous system
router
person, role, maintainer
set classes

Object Representation

person: **Randy Bush** Attribute name
address: RGnet NOC Attribute value
5147 Crystal Springs Drive NE
10361 NE Sasquatch
Bainbridge Island, WE 98110 Comment
USA
phone: +1 206 780 0431
fax-no: +1 206 780 0653
e-mail: randy@psg.com
nic-hdl: RB366 continuation
remarks: This object is automatically
converted from RIPE181
mnt-by: RGNET-MAINT-MCI
changed: randy@psg.com 19970614
source: MCI

Person Class

```
person: Randy Bush
address: RGnet NOC
        5147 Crystal Springs Drive NE
        10361 NE Sasquatch
        Bainbridge Island, WE 98110
        USA
phone: +1 206 780 0431 # day time
fax-no: +1 206 780 0653
e-mail: randy@psg.com
nic-hdl: RB366
remarks: This object is automatically
        converted from RIPE181
mnt-by: RGNET-MAINT-MCI
changed: randy@psg.com 19970614
source: MCI
```

Person class
attributes

Common
attributes
Maintenance

Common Attributes

descr: short free text description of the object
remarks: free text comment attribute
tech-c: Technical contact nic handles
admin-c: Administrative contact nic handles
notify: emails to send notification of changes
mnt-by: maintainer authorized to do changes
changed: <email> <date>
source: registry

Role Class

```
role: RIPE NCC Operations
address: Singel 258
        1016 AB Amsterdam
        The Netherlands
phone: +31 20 535 4444
fax-no: +31 20 545 4445
e-mail: ops@ripe.net
admin-c: CO19-RIPE
tech-c: RW488-RIPE
tech-c: JLSD1-RIPE
nic-hdl: OPS4-RIPE
notify: ops@ripe.net
changed: roderik@ripe.net 19970926
source: RIPE
```

Mntner Class

```
mntner: MAINT-RGNET
descr: RGnet RADB maintainer
admin-c: RB366
tech-c: RB366
upd-to: rw@rg.net
mnt-nfy: randy@psg.com
auth: PGPKEY-23F5CE3
mnt-by: MAINT-RGNET
changed: randy@psg.com 19970804
source: RADB
```

Defines access control for other objects in database.

Auth Attribute

```
* auth: PGPKEY-23F5CE3
* auth: CRYPT-PW lz1A7/JnflTI
* auth: MAIL-FROM cengiz@isi.edu
* auth: MAIL-FROM .*@canet.ca
* auth: NONE
```

Route Class

```
Route: 156.36.0.0/16
origin: AS2914
descr: my routes
mnt-by: MAINT-RGNET
tech-c: RB366
changed: randy@psg.com 19960829
source: RADB
```

Policy information

Route **156.36.0.0/16** is originated by **AS2914**

Notations

AS Numbers **AS 2914**
Address prefixes **156.36.0.0/16**
Route set names **RS-VERIO**
AS set names **AS-VERIO**

Words

Words can have **-** or **_** in the middle

RGNET-MAINT-MCI

Can have digits

RGNET-MAINT-MCI_1

Case insensitive

rgnet-MaInT-MCI

RtConfig Steps

Register objects
Create template, input file for RtConfig
Invoke RtConfig

RtConfig Templates

```
version 10.3
!
interface Ethernet0/0
ip address 198.32.4.1 255.255.255.0
!
! other usual cisco command should be put around here
!
router bgp 3561
neighbor 198.32.4.25 remote-as 4551
!
@RtConfig set cisco_map_name = "mymap"
@RtConfig set cisco_access_list_no = 500
@RtConfig import AS3561 1.1.1.1 AS114 1.1.1.2
! other usual cisco command should be put around here
end
```


RtConfig Invocation

```
% RtConfig < template > config-file
% RtConfig -config junos < template > config-file
% RtConfig -f file < template > config-file
```

access_list

```
./RtConfig
RtConfig> @RtConfig access_list filter as226
!
no access-list 100
access-list 100 permit ip 128.97.0.0 0.0.0.0 255.255.0.0 0.0.0.0
access-list 100 permit ip 128.99.0.0 0.0.0.0 255.255.0.0 0.0.0.0
access-list 100 permit ip 128.125.0.0 0.0.0.0 255.255.0.0 0.0.0.0
access-list 100 permit ip 128.149.0.0 0.0.0.0 255.255.0.0 0.0.0.0
access-list 100 permit ip 129.4.0.0 0.0.0.0 255.255.0.0 0.0.0.0
access-list 100 permit ip 129.193.0.0 0.0.0.0 255.255.0.0 0.0.0.0
access-list 100 permit ip 130.152.0.0 0.0.0.0 255.255.0.0 0.0.0.0
access-list 100 permit ip 130.154.0.0 0.0.0.0 255.255.0.0 0.0.0.0
```

Prefix_list

```
./RtConfig -cisco_use_prefix_lists
RtConfig> @RtConfig access_list filter as226
no ip prefix-list pl100
ip prefix-list pl100 permit 128.97.0.0/16
ip prefix-list pl100 permit 128.99.0.0/16
ip prefix-list pl100 permit 128.125.0.0/16
ip prefix-list pl100 permit 128.149.0.0/16
ip prefix-list pl100 permit 129.4.0.0/16
ip prefix-list pl100 permit 129.193.0.0/16
ip prefix-list pl100 permit 130.152.0.0/16
```

Juniper

```
./RtConfig -config junos
RtConfig> @RtConfig access_list filter as226
policy-statement prefix-list-100 {
  term prefixes {
    from {
      route-filter 128.97.0.0/16 exact accept;
      route-filter 128.99.0.0/16 exact accept;
      route-filter 128.125.0.0/16 exact accept;
      route-filter 128.149.0.0/16 exact accept;
      route-filter 129.4.0.0/16 exact accept;
    }
  }
}
```

Gated

```
% ./RtConfig -config gated
RtConfig> @RtConfig access_list filter as226
128.97.0.0 masklen 16 exact;
128.99.0.0 masklen 16 exact;
128.125.0.0 masklen 16 exact;
128.149.0.0 masklen 16 exact;
129.4.0.0 masklen 16 exact;
```

Bay/Nortel

```
% ./RtConfig -config bcc
RtConfig> @RtConfig access_list filter as226
network address 128.97.0.0 mask 255.255.0.0 match Exact
network address 128.99.0.0 mask 255.255.0.0 match Exact
network address 128.125.0.0 mask 255.255.0.0 match Exact
network address 128.149.0.0 mask 255.255.0.0 match Exact
network address 129.4.0.0 mask 255.255.0.0 match Exact
```

Other Formats

```
@RtConfig printPrefixes "my format %p/%l;\n" filter
as226
my format 128.97.0.0/16;
my format 128.99.0.0/16;
my format 128.125.0.0/16;
my format 128.149.0.0/16;
my format 129.4.0.0/16;
```

Other Options

```
@RtConfig printPrefixes "%p/%l %L %k %K\n"
filter AS226
198.178.204.0/24 8 255.255.255.0 0.0.0.255
198.178.205.0/24 8 255.255.255.0 0.0.0.255
199.249.180.0/24 8 255.255.255.0 0.0.0.255
199.249.181.0/24 8 255.255.255.0 0.0.0.255
```

%p prefix | **%l length** | **%L 32-length** | **%k mask** | **%K dont-care**

printPrefixRanges

```
@RtConfig printPrefixRanges "%p/%l from %n to %m\n"  
filter AS226
```

```
198.178.204.0/23 from 24 to 24
```

```
199.249.180.0/23 from 24 to 24
```

Operators

```
RtConfig> @RtConfig access_list filter as226 and as52
```

```
no ip prefix-list pl100
```

```
ip prefix-list pl100 permit 128.97.0.0/16
```

```
ip prefix-list pl100 permit 131.179.0.0/16
```

```
ip prefix-list pl100 permit 149.142.0.0/16
```

```
ip prefix-list pl100 permit 164.67.0.0/16
```

```
ip prefix-list pl100 deny 0.0.0.0/0 le 32
```

Roe

The screenshot shows the RTConfig interface with a list of prefix ranges and a detailed view of the 'demo 52' entry.

File	Show	Selection	Configure
128.97.0.0/16	---	NI:hd	demo:RS226 demo:RS52
128.59.0.0/16	---	NI:hd	demo:RS226 demo:RS226
128.125.0.0/16	---	NI:hd	demo:RS226 demo:RS226
128.149.0.0/16	---	NI:hd	demo:RS226 demo:RS226
128.4.0.0/16	---	NI:hd	demo:RS226 demo:RS226
128.133.0.0/16	---	NI:hd	demo:RS226 demo:RS226
130.152.0.0/16	---	NI:hd	demo:RS226 demo:RS226
130.154.0.0/16	---	NI:hd	demo:RS226 demo:RS226
130.221.0.0/16	---	NI:hd	demo:RS226 demo:RS226
131.173.0.0/16	---	NI:hd	demo:RS226 demo:RS226

demo 52

route: 128.97.0.0/16
descr: University of California
Los Angeles
Postler Hall and Mathematical Sciences Buil
Los Angeles
CA 90024, USA

origin: AS226
member-of: RS-COMM.INSFNET
mnt-by: MAINT-AS226
changed: P. S. 19951011
priority: 100
integrity: no-auth # 136 20000610 22:04:00 +00:00
source: demo

Buttons: Add Template, Delete Template, Update Template, Schedule, Update IRR, Pending Replies: 0

Roe

The screenshot shows the RTConfig interface with a list of prefix ranges and a detailed view of the 'demo 52' entry.

File	Show	Selection	Configure
128.97.0.0/16	---	NI:hd	demo:RS226 demo:RS52
130.154.0.0/16	---	NI:hd	demo:RS226 demo:RS226
131.179.0.0/16	---	NI:hd	demo:RS226 demo:RS226
149.142.0.0/16	---	NI:hd	demo:RS226 demo:RS226
157.127.0.0/16	---	NI:hd	demo:RS226 demo:RS1740
157.249.0.0/16	---	NI:hd	demo:RS226 demo:RS226
164.67.0.0/16	---	NI:hd	demo:RS226 demo:RS226
164.157.0.0/16	---	NI:hd	demo:RS226 demo:RS226
198.178.204.0/24	---	NI:hd	demo:RS226 demo:RS226
199.249.180.0/24	---	NI:hd	demo:RS226 demo:RS226
199.249.180.0/23	---	NI:hd	demo:RS226 demo:RS226

demo 52

route: 128.97.0.0/16
descr: University of California
Los Angeles
Postler Hall and Mathematical Sciences Buil
Los Angeles
CA 90024, USA

origin: AS226
member-of: RS-COMM.INSFNET
mnt-by: MAINT-AS226
changed: P. S. 19951011
priority: 100
integrity: no-auth # 136 20000610 22:04:00 +00:00
source: demo

Buttons: Add Template, Delete Template, Update Template, Schedule, Update IRR, Pending Replies: 0

Operators

```
RtConfig> @RtConfig access_list filter as226 and as52
no ip prefix-list pl100
ip prefix-list pl100 permit 128.97.0.0/16
ip prefix-list pl100 permit 131.179.0.0/16
ip prefix-list pl100 permit 149.142.0.0/16
ip prefix-list pl100 permit 164.67.0.0/16
ip prefix-list pl100 deny 0.0.0.0/0 le 32
```

Operators

```
@RtConfig access_list filter as226 and not as52
no ip prefix-list pl100
ip prefix-list pl100 permit 128.99.0.0/16
ip prefix-list pl100 permit 128.125.0.0/16
ip prefix-list pl100 permit 128.149.0.0/16
ip prefix-list pl100 permit 129.4.0.0/16
```

Peval

```
./peval as226
({205.147.0.0/20, 204.140.128.0/17, 204.80.138.0/24,
204.80.96.0/19, 204.57.0.0/21, 199.249.182.0/24,
199.249.180.0/24, 199.249.181.0/24, 199.249.179.0/24,
199.182.86.0/24, 199.182.81.0/24, 199.165.231.0/24,
199.120.212.0/24, 198.186.216.0/24, 198.178.206.0/24,
198.178.204.0/24, 198.178.205.0/24, 198.178.203.0/24,
198.147.64.0/18, 198.32.0.0/23, 192.251.121.0/24,
192.88.8.0/24, 192.45.0.0/16, 192.27.235.0/24,
192.12.44.0/24, 192.12.19.0/24, 192.5.133.0/24,
192.5.14.0/24, 192.5.10.0/24, 165.192.0.0/16, 164.67.0.0/16,
163.40.0.0/16, 157.242.0.0/16, 157.127.0.0/16,
155.157.0.0/16, 149.142.0.0/16, 144.215.0.0/16,
143.127.0.0/16, 140.171.0.0/16, 137.78.0.0/15, 134.70.0.0/16,
134.4.0.0/16, 131.215.0.0/16, 131.179.0.0/16, 130.221.0.0/16,
130.154.0.0/16, 130.152.0.0/16, 129.193.0.0/16, 129.4.0.0/16,
128.149.0.0/16, 128.125.0.0/16, 128.97.0.0/16})
```

Outline

- Motivation
- Routing Policy Definitions
- RPSL & RtConfig
- Basic classes and notations
- Set classes
- import/export policy
- Structured Policy
- Other Features & Wrap-up

route-set Class

```
route-set: rs-foo
members: 128.9.0.0/16, 128.9.0.0/24,
         128.8.0.0/16
descr: some address prefixes
mnt-by: MAINT-RGNET
tech-c: RB366
changed: randy@psg.com 19960829
source: RADE
```

```
route-set: rs-bar
members: 128.7.0.0/16, rs-foo
```

Route-set

```
route-set: RS-BCMI2
descr: routes via BCM to be announced to I2
members: 128.249.0.0/16, 192.31.88.0/24,
         192.147.26.0/24
admin-c: JCY
tech-c: SM346
mnt-by: MAINT-AS302
changed: smace@intt.org 20000213
source: demo
```

Route-set

```
RtConfig> @RtConfig access_list filter RS-BCMI2
no ip prefix-list p1100
ip prefix-list p1100 permit 128.249.0.0/16
ip prefix-list p1100 permit 192.31.88.0/24
ip prefix-list p1100 permit 192.147.26.0/24
ip prefix-list p1100 deny 0.0.0.0/0 le 32
```

Indirect Members

```
route-set: RS-ANS-IGP_ONLY
descr: ANS IGP aggregates
mbrs-by-ref: any

route: 207.25.17.0/24
origin: AS1675
member-of: RS-ANS-IGP_ONLY
mnt-by: MNT-ANS

route: 192.157.69.0/24
origin: AS1675
member-of: RS-ANS-IGP_ONLY
mnt-by: MNT-ANS
```

Restricted Indirect Members

```
route-set: RS-ANS-IGP_ONLY
descr: ANS IGP aggregates
mbrs-by-ref: MNT-ANS, MNT-CENGIZ

route: 207.25.17.0/24
origin: AS1675
member-of: RS-ANS-IGP_ONLY
mnt-by: MNT-ANS

route: 192.157.69.0/24
origin: AS1675
member-of: RS-ANS-IGP_ONLY
mnt-by: MNT-ANS
```

Direct & Indirect Members

```
route-set: RS-ANS-IGP_ONLY
descr: ANS IGP aggregates
members: 207.25.17.0/24, 207.25.16.0/24,
207.25.20.0/24
mbrs-by-ref: MNT-ANS

route: 207.25.17.0/24
origin: AS1675
member-of: RS-ANS-IGP_ONLY
mnt-by: MNT-ANS

route: 192.157.69.0/24
origin: AS1675
member-of: RS-ANS-IGP_ONLY
mnt-by: MNT-ANS
```

More Specific Operators

```
route-set: rs-martians
descr: most ASes do not import these routes
members: 0.0.0.0/0^32, 127.0.0.0/8^+,
10.0.0.0/8^+, 172.16.0.0/20^+,
192.168.0.0/16^+, 192.0.2.0/24^+,
128.0.0.0/16^+, 191.255.0.0/16^+,
192.0.0.0/24^+, 223.255.255.0/24^+,
224.0.0.0/3^+, 0.0.0.0/0^26-32
```

Inclusive more specifics: ^+

Exclusive more specifics: ^-

Length n more specifics: ^n

Length n-m more specifics: ^n-m

Route-set Namespaces

```
Route-set: AS4763:RS-ROUTES:AS681
descr: prefix filter for AS681
members: 130.216.0.0/16, 130.217.0.0/16,
132.181.0.0/16, 138.75.0.0/16, 139.80.0.0/16,
140.200.0.0/16, 156.62.0.0/16, 192.73.21.0/24
tech-c: JA39
mnt-by: MAINT-TELSTRA-NZ
changed: jabley@patho.gen.nz 19991118
source: RADB
```

as-set Class

```
as-set: AS-SESQUI-STUB
descr: Single Homed Sesquinet Customer ASs
members: AS1832, AS2712, AS302, AS3526, AS8
tech-c: SB98
mnt-by: MAINT-AS114
source: RADE
```

Same flexibility as route-set class

Indirect as-sets

```
as-set: as-aads-mlpa
descr: MLPA participants at the AADS NAP
mbrs-by-ref: any
admin-c: Andrew Schmidt
tech-c: Mark Cnota
notify: mlpa-participants@aads.net
mnt-by: MAINT-RSPEER
changed: auto-mlpa@aads.net 19971123
source: RADE

aut-num: AS4550
member-of: as-aads-mlpa

aut-num: AS683
member-of: as-aads-mlpa
```

As-sets

```
as-set: AS-GOODNET
descr: ASs routed through GoodNet
members: AS5696, AS1808, AS1932, AS2900, AS3111,
AS3365, AS3393, AS3844, AS3901, AS4314, ... AS-
ACESRESEARCH, AS-ALPHA, AS-GST, AS-DERU, AS-INQUO
admin-c: IP Admin DW970
tech-c: IP Admin DW970
notify: ipadmin@winstar.net
mnt-by: MAINT-AS5696
changed: darin@good.net 19990731
source: demo
```

as-set

```
@RtConfig access_list filter as-goodnet
no ip prefix-list pl100 permit 12.10.231.0/24
ip prefix-list pl100 permit 63.140.0.0/16
ip prefix-list pl100 permit 64.16.64.0/20
ip prefix-list pl100 permit 64.16.75.0/24
ip prefix-list pl100 permit 64.16.128.0/20
ip prefix-list pl100 permit 64.16.128.0/24
ip prefix-list pl100 permit 64.16.138.0/23
...
```

Outline

- Motivation
- Routing Policy Definitions
- RPSL & RConfig
- Basic classes and notations
- Set classes
- import/export policy**
- Structured Policy
- Other Features & Wrap-up

aut-num Class


```
aut-num: AS4591
as-name: Syra-NET
import: from AS4590
       action pref=1;
       accept AS4590
export: to AS4590
       announce AS4591
default: to AS4590
        action pref=1;
        networks {140.222.0.0/16}
admin-c: Warren Lavallee
tech-c: Warren Lavallee
mnt-by: MAINT-AS4591
changed: warren@Syra.NET 19950522
source: RADB
```

Policy Styles

RPSL allows import/export policies based on:

- prefix
- AS Path
- community
- prefix-length
- other/future attributes thru its dictionary

Prefix Based



```
aut-num: AS1
export: to AS2
       announce {128.9.0.0/16, 128.8.0.0/16}

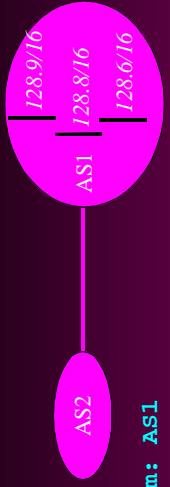
aut-num: AS2
import: from AS1
       accept {128.9.0.0/16, 128.8.0.0/16}
```


RtConfig: A tool to configure routers

```
% RtConfig
RtConfig> @RtConfig export AS1 1.1.1.1 AS2 2.2.2.2
no access-list 1
access-list 1 permit ip 128.8.0.0 0.0.0.0 255.255.0.0 0.0.0.0
access-list 1 permit ip 128.9.0.0 0.0.0.0 255.255.0.0 0.0.0.0
access-list 1 deny ip 0.0.0.0 255.255.255.255 0.0.0.0 255.255.255.255
!
no route-map foo
route-map foo permit 1
match ip address 1
!
!
```

Aut-num for AS2 is not needed!!

Cumbersome?



```
aut-num: AS1
export: to AS2
      announce {128.9.0.0/16, 128.8.0.0/16,
                128.6.0.0/16}

aut-num: AS2
import: from AS1
      accept  {128.9.0.0/16, 128.8.0.0/16,
              128.6.0.0/16}
```


Per route-set

```
route-set: rs-red
members:
  128.6.0.0/16, 128.9.0.0/16,
  128.8.0.0/16

aut-num: AS1
export: to AS2 announce rs-red

aut-num: AS2
import: from AS1 accept rs-red
```

Per Origin AS



```
route: 128.9.0.0/16 route: 128.8.0.0/16
origin: AS1 origin: AS1

aut-num: AS1
export: to AS2 announce AS1


aut-num: AS2
import: from AS1 accept AS1

AS1 == {128.9.0.0/16, 128.8.0.0/16}
```

RtConfig

```
% RtConfig
RtConfig> @RtConfig export AS1 1.1.1.1 AS2 2.2.2.2
no access-list 1
access-list 1 permit ip 128.8.0.0 0.0.0.0 255.255.0.0 0.0.0.0
access-list 1 permit ip 128.9.0.0 0.0.0.0 255.255.0.0 0.0.0.0
access-list 1 deny ip 0.0.0.0 255.255.255.255 0.0.0.0 255.255.255.255
!
no route-map foo
route-map foo permit 1
match ip address 1
!
```


Customers



```
aut-num: AS1
export: to AS2 announce AS1 OR AS3 OR ... AS6

aut-num: AS2
import: from AS1 accept AS1 OR AS3 OR ... AS6
```

Customers in a set




```
as-set: AS1:AS-Customers
members: AS1, AS3, AS4, AS5, AS6

aut-num: AS1
export: to AS2 announce AS1:AS-Customers

aut-num: AS2
import: from AS1 accept AS1:AS-Customers
```

Customers of Customers




```
Aut-num: AS1
export: to AS2 announce AS1:AS-Customers

as-set: AS1:AS-Customers
members: AS1, AS3, AS4, AS5, AS6:AS-Customers


as-set: AS6:AS-Customers
members: AS6, AS7, AS8
```

More Customers?



```
aut-num: AS2
import: from AS1 accept AS1:AS-Customers
import: from AS3 accept AS3:AS-Customers
import: from AS4 accept AS4:AS-Customers
```


Shortcut: PeerAS



```
as-set: AS2:AS-Customers
members: AS1, AS3, AS4

aut-num: AS2
import: from AS2:AS-Customers
accept PeerAS:AS-Customers
```

Preference/Cost



```
aut-num: AS4
import: from AS1 action pref = 10; accept ANY
import: from AS3 action pref = 15; accept ANY
```

Smaller the number, higher the preference!

Actions


```
import: from ... action XXX; accept ...
export: to ... action XXX; announce ...

med = 0;
med = igp_cost;

community.append(NO_EXPORT, 10250, 3561:90);
community.delete(NO_EXPORT);

aspath.prepend(AS1, AS1, AS1);
```


Action



```
aut-num: AS4
export: to AS1 announce AS4
export: to AS3 action aspath.prepend(AS4);
        announce AS4
```


What would happen if `aspath.prepend(AS4,AS4,AS4,AS4)`?

Peering Choice



```
aut-num: AS1
import: from AS2 accept AS2
```

Peering Choice




```
aut-num: AS1
import: from AS2
        action pref = 10;
        accept AS2
import: from AS2 1.1.1.2 at 1.1.1.1
        action pref = 5;
        accept AS2
```

Logical/Set Operators

```
Aut-num: AS1
import: from AS1
        accept (AS1 OR rs-red) AND NOT {0.0.0.0/0}
```

```
AS1
rs-red
== {128.9.0.0/16, 128.8.0.0/16}
== {128.6.0.0/16, 128.9.0.0/16}
AS1 OR rs-red == {128.6.0.0/16, 128.8.0.0/16, 128.9.0.0/16}
AS1 AND rs-red == {128.9.0.0/16}
AS1 AND NOT rs-red == {128.8.0.0/16}
```

AS Path Based



AS paths that start in AS1 and end in AS8:


`<^AS1 .* AS8$>`

No prefix filters here!

AS Path Regular Expressions

<code>AS1</code>	AS1 only
<code>.</code>	Any AS
<code>as-foo</code>	Any member of as-foo
<code>X*</code>	0 or more occurrence of X
<code>X+</code>	1 or more occurrence of X
<code>X?</code>	0 or 1 occurrence of X
<code>^</code>	beginning of path
<code>\$</code>	end of path
<code>X Y</code>	X or Y
<code>X Y</code>	X followed by Y


AS Paths into AS1's Customers



`<^AS1 * AS1:AS-Customers * $>` matches:

AS1
 AS1 AS3
 AS1 AS4
 AS1 AS5 AS6
 AS1 AS1 AS5 AS5 AS6

AS Path Based import/export



`import: from AS1 accept <^AS1 .* AS8$>`

`import: from AS1 accept <^AS1 AS1:AS-Customers*$>`

RtConfig

```
RtConfig> @RtConfig import AS2 1.1.1.1 AS1 2.2.2.2
!
no ip as-path access-list 1
ip as-path access-list 1 permit ^_1(([_0-9]+))*_8$
!
no route-map foo
route-map foo permit 1
match as-path 1
!
router bgp 2
neighbor 2.2.2.2 route-map foo in
```

community Based



AS4 wants AS3561 to prefer AS1 path


AS3561 prefers routes with

- no community
- with community 3561:90
- with community 3561:80
- with community 3561:70

AS3561's Policies

```
aut-num: AS3561
import: from AS-ANY
      action pref = 30;
      accept community(3561:70)
import: from AS-ANY
      action pref = 20;
      accept community(3561:80)
import: from AS-ANY
      action pref = 10;
      accept community(3561:90)
import: from AS-ANY
      action pref = 0;
      accept ANY
```

AS4's Policies



aut-num: AS4

export: to AS1 action community.={3561:90};

to AS3 action community.={3561:80};

announce AS4

RtConfig

```
aut-num: AS1
export: to AS2
announce AS2764 AND NOT { 0.0.0.0/0 }
      AND <^AS2764*$>
      AND NOT community(2764:1)
```

RtConfig

```
RtConfig> @RtConfig export AS1 1.1.1.1 AS2 2.2.2.2
access-list 1 permit ip 210.8.248.0 0.0.0.0 255.255.248.0 0.0.0.0
...
access-list 1 permit ip 210.9.2.0 0.0.0.0 255.255.254.0 0.0.0.0
access-list 1 deny ip 0.0.0.0 255.255.255.255 0.0.0.0 255.255.255.255
!
no ip as-path access-list 1
ip as-path access-list 1 permit ^(_2764)*$
!
ip bgp-community new-format
no community-list 1
ip community-list 1 deny 2764:1
ip community-list 1 permit internet
!
```

RtConfig (cont)

```
no route-map foo
route-map foo permit 1
  match as-path 1
  match community 1
  match ip address 1
!
router bgp 1
neighbor 0.0.0.0 route-map foo out
```

Prefix Length Based

```
aut-num: AS1
import: from as-any
      accept ANY AND NOT { 0.0.0.0/0^19-32 }
```

Outline

- Motivation
- Routing Policy Definitions
- RPSL & RConfig
- Basic classes and notations
- Set classes
- import/export policy
- Structured Policy**
- Other Features & Wrap-up

Accept Good Routes

```
from AS-ANY
  accept ANY and not {0.0.0.0/0};
```

Better yet:

```
from AS-ANY
  accept ANY and not RS-MARTIANS;
```

```
route-set: RS-MARTIANS
members: 0.0.0.0/0, 127.0.0.0/8^+, 10.0.0.0/8^+,
172.16.0.0/20^+, 192.168.0.0/16^+, 192.0.2.0/24^+,
128.0.0.0/16^+, 191.255.0.0/16^+, 192.0.0.0/24^+,
223.255.255.0/24^+, 224.0.0.0/3^+, 0.0.0.0/0^26-32
```

Set Preference Using Communities

```
from AS-ANY action pref=20;
  accept community(1:20);
from AS-ANY action pref=10;
  accept community(1:10);
from AS-ANY action pref=0;
  accept any;
```

Filter (at least) the customers

```
from AS2 accept AS2;
from AS3 accept AS3 or AS4;
from AS5 accept AS5:AS-Customers;
...
```


Combine: Structured Policy

```
import: {  
  from AS-ANY  
  accept ANY and not RS-MARTIANS;  
} refine {  
  from AS-ANY action pref=10;  
  accept community(1:10);  
  from AS-ANY action pref=20;  
  accept community(1:20);  
  from AS-ANY action pref=0;  
  accept any;  
}  
  from AS2 accept AS2;  
  from AS3 accept AS3 or AS4;  
  from AS5 accept AS5:AS-Customers;  
}
```

Structured Policy

```
aut-num: AS3561  
import: {  
  from AS-ANY action pref = 30;  
  accept community(3561:70);  
  from AS-ANY action pref = 20;  
  accept community(3561:80);  
} refine {  
  from AS1 accept AS1:AS-Customers;  
} except {  
  from AS2 accept AS2;  
  from AS3 accept AS3;  
}
```

AS1:AS-Customers contains AS2 and AS3

Outline

- Motivation
- Routing Policy Definitions
- RPSL & RConfig
- Basic classes and notations
- Set classes
- import/export policy
- Structured Policy
- Other Features & Wrap-up

Static Routes

```
route: 128.7.0.0/16  
origin: AS1  
inject: at 7.7.7.1  
  action next-hop = 7.7.7.2; cost = 10;  
  upon static  
inject: at 7.7.7.1  
  action next-hop = 7.7.7.3; cost = 20;  
  upon static
```

Aggregation

```
route: 128.8.0.0/15
origin: ASI
components: {128.8.0.0/15^-}
aggr-mtd: outbound AS-ANY
inject: at 1.1.1.1 action dpa = 100;
inject: at 1.1.1.2 action dpa = 110;
```

Inet-rtr

```
inet-rtr: aads.OneCall.Net
local-as: AS6402
ifaddr: 207.7.19.254 masklen 24
ifaddr: 207.112.228.254 masklen 24
peer: BGP4 198.32.130.52 asno(AS6583)
peer: BGP4 198.32.130.49 asno(AS5696)
admin-c: RI69
tech-c: RI69
remarks: AADS router for the OneCall NetWork
mnt-by: MAINT-AS6402
```

RtConfig

```
@RtConfig configureRouter aads.OneCall.Net

Generates import/export for:

peer: BGP4 198.32.130.52 asno(AS6583)
peer: BGP4 198.32.130.49 asno(AS5696)
```

ConfigureRouter

```
RtConfig> @RtConfig configureRouter aads.onecall.net
...
ip prefix-list pl102 permit 216.246.0.0/17
ip prefix-list pl102 permit 216.246.0.0/18
ip prefix-list pl102 deny 0.0.0.0/0 le 32
route-map MyMap_6402_15 permit 1
match ip address prefix-list pl102
set local-preference 980
router bgp 6402
neighbor 198.32.130.49 route-map MyMap_6402_15 in
```

RtConfig

```
@RtConfig import AS1 1.1.1.1 AS2 2.2.2.2.2
@RtConfig export AS1 1.1.1.1 AS2 2.2.2.2.2
@RtConfig configureRouter c56-11.t3.ans.net
@RtConfig networks AS1
@RtConfig default AS1 AS2
@RtConfig pkt_filter "eth0" AS1 1.1.1.1 AS2 2.2.2.2.2
@RtConfig outbound_pkt_filter "eth0" AS1 1.1.1.1 AS2 2.2.2
@RtConfig static2bgp AS1 1.1.1.1
@RTConfig printPrefixes "format" filter <filter>
```

Configuration formats:

cisco, juniper, gated, rsd, bay/nortel
lucent coming soon

Extensible thru dictionary

```
dictionary: RPSL
rp-attribute: pref # smaller values are preferred
operator=(integer[0, 65535])
rp-attribute: med # BGP multi_exit_discriminator
operator=(integer[0, 65535])
typedef: community_elm union
integer[1, 4294967200],
enum[internet, no_export, no_advertise]
rp-attribute: community # BGP community attribute
operator.=(list of community_elm)
operator()(community_elm, ...)
```

Pointers

<http://www.isi.edu/ra/rps/training>

Tutorial slides

RFC2622 (RPSL)

RFC2650 (RPSL Applications)

<http://www.isi.edu/ra/RAToolSet>

Mailing lists

rps@isi.edu

radasect@isi.edu

Cengiz Alaettinoglu <cengiz@packetdesign.com>