RomandIX
Year one.

Agenda

-Presenter Bio
-What is RomandIX
-Why another IXP
-What we did
-Lesson learned / Next steps
-Questions

Bio

-Nicolas Desir

Founder of Saitis (AS6893) an ISP in Lausanne, working with community networks for 20+ years, Datacenter builder, hardware tinkerer […]

-Will van Gulik

Running AS2613, Network Engineer at AS6893, Ripe Connect-WG co-chair, Ripe Atlas Ambassador, MANRS advocate, Dj, Writer […]

What is RomandIX

RomandIX is a Lausanne IXP founded as a non-profit organisation the 13th March 2017 by :

-AlpineDC – AS198385
-AS2613
-DFI Services – AS12333
-EDSI-Tech – AS202194
-Evok – AS203882
-Infraly – AS50963
-Openbusiness – AS49457
-Saitis Networks – AS6893
-Telcopack – AS206214

Currently 9 Active networks, 3 more connecting

Why another IXP

- Keep your traffic local
- Simpler interconnections (no full mesh)
- Additional services
- We were bored and didn’t have enough stuff to do
- The Fun ;)

What we did #1

-Build an association
-Get IPv4/6 space and an ASN from Ripe NCC
-Buy a switch and rack it (Arista 7050)
-Get two servers for our services in two different locations (AlpineDC room in Brain hosted by EDSI and in Provence hosted by Saitis)
-Create the PeeringDB entries / IRR Records
-Use virtualization, VMs, Install IXP Manager and …

What we did #2

… install the Route Servers, activate filtering day one, daily updates

```
./gen_prefix_filters_verbose.sh
[...]
AS2613: [IPv4: 1 total; 0 stale; 0 new; DB updated] [IPv6: 3 total; 0 stale; 0 new; DB updated]
PCH3856: [IPv4: 1 total; 0 stale; 0 new; DB updated] [IPv6: 1 total; 0 stale; 0 new; DB updated]
PCH42: [IPv4: 272 total; 0 stale; 0 new; DB updated] [IPv6: 237 total; 0 stale; 0 new; DB updated]
Saitis: [IPv4: 72 total; 0 stale; 0 new; DB updated] [IPv6: 69 total; 0 stale; 0 new; DB updated]
Telcopack: [IPv4: 35 total; 0 stale; 0 new; DB updated] [IPv6: 6 total; 0 stale; 0 new; DB updated]
ROIX: [IPv4: 0 total; 0 stale; 0 new; DB updated] [IPv6: 0 total; 0 stale; 0 new; DB updated]
   Customer not a RS client or IRRDB filtered. Prefixes, if any, wiped from database.
[...]
```

-Configure the switch with a Reseller VLAN mechanism, one VLAN per port

- But what’s the point, everyone can do that. Nico wanted to do more.

Let's hack a switch!

Broadcast/mac learning, ARP/NDP inspection, sponge: Is there another paradigm?

- We need something not too expensive/energy consuming
- NDP inspection can't be found today on available second hand hardware like ours
- Is there an alternative way?
- Can we also solve other problems:
    - broadcast when mac is forgotten?
    - ARP broadcast load?

L2 Part

-Separation first
-Each participant in a separate VLAN
-Let's give a try to the diretflow feature of Arista

```
flow to_peer_001_ipv4
    Table Trident ifp
    match destination mac 42:07:f8:b5:00:00 mask ff:ff:ff:ff:ff:f0
    match ethertype ip
    match vlan 512 mask 3584
    action output interface Et47
    action set vlan 512
 !
 flow to_peer_001_ipv6
    Table Trident ifp
    match destination mac 42:07:f8:b5:00:00 mask ff:ff:ff:ff:ff:f0
    match ethertype 34525
    match vlan 512 mask 3584
    action output interface Et47
    action set vlan 512
 !
```

L3 Part

It won't work, ARP and NDP need [b/m]cast !

-No inspecting/filtering: let's generate ARP and NDP response
-so easy access to the shell in EOS
-executable build on debian7/amd64 run on the switch
-each VLAN can be seen as a standard interface in the switch

ARP: the simple case

-arp-dictator : a home made simple executable
-static config file which IP/MAC mapping
-one per VLAN with rate-limiting

```
bash-4.3# pwd
/mnt/flash/bin
bash-4.3# cat v4tomac
185.1.93.1 42:07:F8:B5:00:01
185.1.93.2 42:07:F8:B5:00:02
185.1.93.3 42:07:F8:B5:00:11
185.1.93.4 42:07:F8:B5:00:04
185.1.93.5 42:07:F8:B5:00:05
[...]
```

NDP: the tricky case

- ndp-dictator
- must answer for the real and link-local addresses
- multicast
- no specific Ethernet protocol
    - directflow will forward unicast response
    - can be faked
    - we must filter
- let's deny inbound all ND except Neighbor Solicitation

```
ipv6 access-list ndonlyns
   20 deny icmpv6 any any 133
   30 deny icmpv6 any any 134
   40 deny icmpv6 any any 136
   50 deny icmpv6 any any 137
   60 permit ipv6 any any


interface Ethernet11
    switchport access vlan 11
    ipv6 access-group ndonlyns in
```

NDP: same idea as ARP

```
bash-4.3# cat v6tomac
# [syntax ipv6 address] [mac]
2001:7f8:b5::10:1 00:1b:21:99:a4:21
fe80::21b:21ff:fe99:a421 00:1b:21:99:a4:21
2001:7f8:b5::11:1 e4:8D:8C:A4:B7:4E
fe80::e68d:8cff:fea4:b74e e4:8D:8C:A4:B7:4E
[...]
```

- -local IPv6 adressed are made with MAC addresses
  - -will be automatic in the futur ;)

Caveats

-Broken ARP/NDP reachability information

  -to be addressed/implemented by dynamic checks and communication between daemons

  -BGP's responsibility

-AFAIK not so big and bad problem:

  -ARP can be also very slow

Mac Addressing

- Can be configured on more and more devices
- Some still cause trouble
- Exception are accepted
- Better to give the MAC address to the participant
    - insure the participant is able to configure it
    - participant will be able to do it again when choosing/changing equipment

https://kb.juniper.net/InfoCenter/index?page=content&id=KB28918 :

"As confirmed by the example above, the MAC address configured under the edit interfaces IRB unit <xyz> hierarchy is not effective for the inet6 family.
This has always been the case. There are no plans to include this support in the future."

MAC Addressing #2

- Got an OUI-36 for RomandIX:
    - http://standards-oui.ieee.org/oui36/oui36.txt

```
[...]
70-B3-D5                  (hex)          Association Romandix
AEB000-AEBFFF    (base 16)
        Association Romandix
                          rue de Sebeillon 9b
                 Lausanne  Vaud  1004
                 CH
[...]
```

- Arista works with mac mask
    - one directflow entry can accommodate several peers
- Participant get a /44 on their port
- So no excuse of possible conflicts...

Also an other political paradigm?

-Common IXP way
  -collaborative (democratic/free speech)
  -inspection/blocking (police)
  -quarantine (jail)

-This proposed way
  -separation (Divide),
  -and dictate (Rule)
  -suppress unwanted packets (kill)

Todo for L2/L3 part :

-code cleaning and optimization
-dynamic peer accessibility between daemon (e.g. via IPC)
-ring for maximum performance: PACKET_RX_RING/PACKET_TX_RING
-external code/methodology reviewing (YOU?)

Lessons learned / What's next
Conclusions

-Our switches run Linux, we used it, so could you

-Doing things non-profit takes longer

-We suck at web mastering

But :

-First Swiss member of MANRS for IXPs

-We enforce IRR based filtering on our RS

Todo :

-Reach ZRH (Colozüeri)

-Reach GVA (Cern - IXEu^H^H^HEquinix's GV1/2)

-More members (you ?), more traffic, rpki, röstigraben packet translation, …

-Our tools are available on https://www.romandix.ch/code/

-Contact us : info@romandix.ch

Questions ?

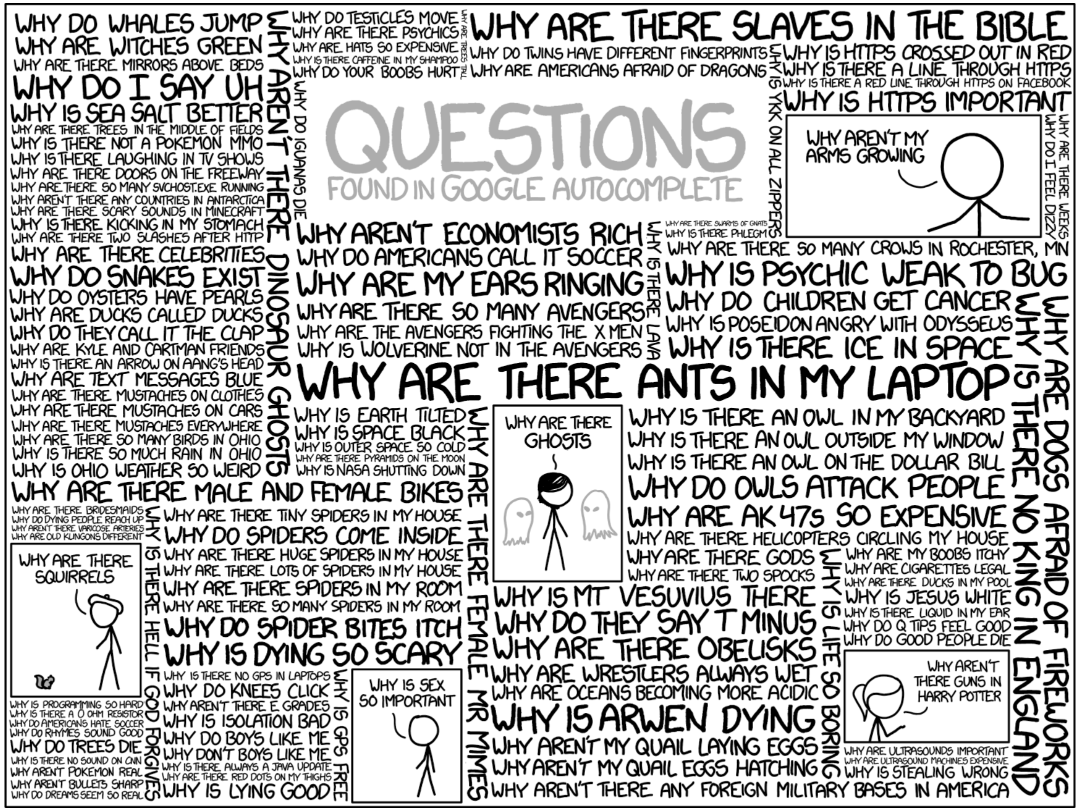
QUESTIONS
FOUND IN GOOGLE AUTOCOMPLETE
WHY DO WHALES JUMP
WHY ARE WITCHES GREEN
WHY ARE THERE MIRRORS ABOVE BEDS
WHY DO I SAY UH
WHY IS SEA SALT BETTER
WHY ARE THERE TREES IN THE MIDDLE OF FIELDS
WHY IS THERE NOT A POKEMON MMO
WHY IS THERE LAUGHING IN TV SHOWS
WHY ARE THERE DOORS ON THE FREEWAY
WHY ARE THERE SO MANY SVCHOST.EXE RUNNING
WHY AREN'T THERE ANY COUNTRIES IN ANTARCTICA
WHY ARE THERE SCARY SOUNDS IN MINECRAFT
WHY IS THERE KICKING IN MY STOMACH
WHY ARE THERE TWO SLASHES AFTER HTTP
WHY ARE THERE CELEBRITIES
WHY DO SNAKES EXIST
WHY DO OYSTERS HAVE PEARLS
WHY ARE DUCKS CALLED DUCKS
WHY DO THEY CALL IT THE CLAP
WHY ARE KYLE AND CARTMAN FRIENDS
WHY IS THERE AN ARROW ON AANG'S HEAD
WHY ARE TEXT MESSAGES BLUE
WHY ARE THERE MUSTACHES ON CLOTHES
WHY ARE THERE MUSTACHES ON CARS
WHY ARE THERE MUSTACHES EVERYWHERE
WHY ARE THERE SO MANY BIRDS IN OHIO
WHY IS THERE SO MUCH RAIN IN OHIO
WHY IS OHIO WEATHER SO WEIRD
WHY AREN'T THERE DINOSAUR GHOSTS
WHY DO IGUANAS DIE
WHY DO TESTICLES MOVE
WHY ARE THERE PSYCHICS
WHY ARE HATS SO EXPENSIVE
WHY IS THERE CAFFEINE IN MY SHAMPOO
WHY DO YOUR BOOBS HURT
WHY ARE TREES TALL
WHY ARE THERE SLAVES IN THE BIBLE
WHY DO TWINS HAVE DIFFERENT FINGERPRINTS
WHY ARE AMERICANS AFRAID OF DRAGONS
WHY IS YKK ON ALL ZIPPERS
WHY IS HTTPS CROSSED OUT IN RED
WHY IS THERE A LINE THROUGH HTTPS
WHY IS THERE A RED LINE THROUGH HTTPS ON FACEBOOK
WHY IS HTTPS IMPORTANT
WHY AREN'T MY ARMS GROWING
WHY ARE THERE WEEKS
WHY DO I FEEL DIZZY
WHY AREN'T ECONOMISTS RICH
WHY DO AMERICANS CALL IT SOCCER
WHY ARE MY EARS RINGING
WHY ARE THERE SO MANY AVENGERS
WHY ARE THE AVENGERS FIGHTING THE X MEN
WHY IS WOLVERINE NOT IN THE AVENGERS
WHY IS THERE LAVA
WHY ARE THERE SWARMS OF GNATS
WHY IS THERE PHLEGM
WHY ARE THERE SO MANY CROWS IN ROCHESTER, MN
WHY IS PSYCHIC WEAK TO BUG
WHY DO CHILDREN GET CANCER
WHY IS POSEIDON ANGRY WITH ODYSSEUS
WHY IS THERE ICE IN SPACE
WHY ARE THERE ANTS IN MY LAPTOP
WHY IS EARTH TILTED
WHY IS SPACE BLACK
WHY IS OUTER SPACE SO COLD
WHY ARE THERE PYRAMIDS ON THE MOON
WHY IS NASA SHUTTING DOWN
WHY ARE THERE MALE AND FEMALE BIKES
WHY ARE THERE GHOSTS
WHY ARE THERE FEMALE MR MIMES
WHY IS THERE AN OWL IN MY BACKYARD
WHY IS THERE AN OWL OUTSIDE MY WINDOW
WHY IS THERE AN OWL ON THE DOLLAR BILL
WHY DO OWLS ATTACK PEOPLE
WHY ARE AK 47s SO EXPENSIVE
WHY ARE THERE HELICOPTERS CIRCLING MY HOUSE
WHY ARE THERE GODS
WHY ARE THERE TWO SPOCKS
WHY IS MT VESUVIUS THERE
WHY DO THEY SAY T MINUS
WHY ARE THERE OBELISKS
WHY ARE WRESTLERS ALWAYS WET
WHY ARE OCEANS BECOMING MORE ACIDIC
WHY IS ARWEN DYING
WHY AREN'T MY QUAIL LAYING EGGS
WHY AREN'T MY QUAIL EGGS HATCHING
WHY AREN'T THERE ANY FOREIGN MILITARY BASES IN AMERICA
WHY IS LIFE SO BORING
WHY ARE MY BOOBS ITCHY
WHY ARE CIGARETTES LEGAL
WHY ARE THERE DUCKS IN MY POOL
WHY IS JESUS WHITE
WHY IS THERE LIQUID IN MY EAR
WHY DO Q TIPS FEEL GOOD
WHY DO GOOD PEOPLE DIE
WHY AREN'T THERE GUNS IN HARRY POTTER
WHY ARE ULTRASOUNDS IMPORTANT
WHY ARE ULTRASOUND MACHINES EXPENSIVE
WHY IS STEALING WRONG
WHY IS THERE NO KING IN ENGLAND
WHY ARE DOGS AFRAID OF FIREWORKS
WHY ARE THERE BRIDESMAIDS
WHY DO DYING PEOPLE REACH UP
WHY AREN'T THERE VARICOSE ARTERIES
WHY ARE OLD KLINGONS DIFFERENT
WHY ARE THERE SQUIRRELS
WHY IS PROGRAMMING SO HARD
WHY IS THERE A 0 OHM RESISTOR
WHY DO AMERICANS HATE SOCCER
WHY DO RHYMES SOUND GOOD
WHY DO TREES DIE
WHY IS THERE NO SOUND ON CNN
WHY AREN'T POKEMON REAL
WHY AREN'T BULLETS SHARP
WHY DO DREAMS SEEM SO REAL
WHY IS THERE HELL IF GOD FORGIVES
WHY ARE THERE TINY SPIDERS IN MY HOUSE
WHY DO SPIDERS COME INSIDE
WHY ARE THERE HUGE SPIDERS IN MY HOUSE
WHY ARE THERE LOTS OF SPIDERS IN MY HOUSE
WHY ARE THERE SPIDERS IN MY ROOM
WHY ARE THERE SO MANY SPIDERS IN MY ROOM
WHY DO SPIDER BITES ITCH
WHY IS DYING SO SCARY
WHY IS THERE NO GPS IN LAPTOPS
WHY DO KNEES CLICK
WHY AREN'T THERE E GRADES
WHY IS ISOLATION BAD
WHY DO BOYS LIKE ME
WHY DON'T BOYS LIKE ME
WHY IS THERE ALWAYS A JAVA UPDATE
WHY ARE THERE RED DOTS ON MY THIGHS
WHY IS LYING GOOD
WHY IS GPS FREE
WHY IS SEX SO IMPORTANT

-Thanks to xkcd.com for the image

-Thanks to Sébastien Keller, Florian Hibler and Romain Aviolat

-https://Romandix.ch (com/net/org)

-https://peeringdb.com/ix/1984

-https://www.manrs.org/ixps/