VirtIX

Virtual Internet Exchange

IPredator et al

virtix.st

What is the VirtIX?

- The idea of VirtIX is to build an ISP-level VPN provider for Tier 3
  – and to allow peering at a Virtual Internet Exchange using common BGP

What does the ISP-level VPN provider solve?

- Tier 1 and Tier 2 providers are in bed with the government,
  – and betray the general trust model the internet was built on
- The VirtIX enables Tier 3 ISPs to encrypt traffic exchanged with other ISPs,
  to peer independently of geographical location,
  and to keep traffic private from Tier 1 and 2
  ... and anyone else tapped in

Overcoming geographical restrictions

- Currently peerings are subject to geographical location –
  – limiting the direct peering between Tier 3 providers
- Most countries do not have regulation for ISP peerings,
  – providing room for creativity in times of dire need
- Virtual encrypted circuits allow for peering with one another without the need of being in the same location –
  – it allows for more direct, *closer* exchange of traffic

Encryption on Layer 3 and 4 vs upper layers

- So far, cryptography is only considered in the upper layers of the OSI model
  –
  – which poses a few problems
    - Scaling
    - Privacy considerations regarding meta data
    - Professionalism and ethics

Scaling problems

- Does not really scale *right* –

  – we are leaving encryption only up host to host communications
- In a world of servers, computers and laptops, smart phones, tablets, and networked fridges

  – a gazillion hosts are supposed to securely encrypt traffic between each other,
- All, while the 47k networks providing for all that traffic have basically not implemented any encryption at all

Meta data and the flow of the internet

- As we have no concept of encrypting traffic on the backbone level –
  – we leave the very flow of the internet exposed
- The flow is what all this Meta data fuzz is about,
  – our social graphs
- To understand the flow is to defeat anonymity, –
  – especially for the general public

Question of professionalism

- Taking care of the privacy of the flow should be considered due diligence for ISPs in a post-Snowden world
- We should be as ambitious when it comes to defend from illegitimate spying,
  - by someone preying on our networks, infrastructure and integrity,
  - just as we are about optimizing our networks for speed and reliability
- We seem to be trapped and helpless in this world of NSA and GCHQ and wiretapping submarines –
  - but we must not be
- It is time we start using some unconventional strategies and push back

We ¡3 infrastructure

- As a Tier 3 provider feeling responsible for users as well as infrastructure, we want to break free from the upper tiers –

  – and see if we can create an oasis of free and protected exchange for other smaller ISPs

Our thesis

- It is possible to build a public BGP overlay network for Tier 3 providers
  – for research purposes
- The negative effects of the overlay network can be managed
- It won't break the internet. We are past the BGP year2k ; 512k+ IPv4 routing table entries already happened
- The success of any such endeavor will depend on its tooling
  – the ability to integrate a session to the VirtIX into existing BGP peering infrastructure
- It needs to be compatible with the workflow at your average AS out there

What does the live internet look like?

Some statistics

- To see if we could backup our feeling we ran some statistics on our BGP routing table

  – trying to see how networks on the internet are distributed
- Are there enough "small" ISPs at all?
- Is there a use-case for interconnecting plenty of small ISPs?
- Our routine to find out:
  - grab full table BGP feed
  - merge redundantly announced prefixes
  - sort network sizes into CIDR classes
  - run some statistics

AS distribution by aggregated IP network size

80% out of 47k AS announce networks not bigger than 8192 IPs

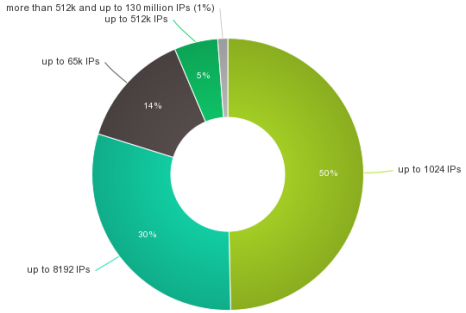

1% of AS announce networks between 512k and 130 million IPs

Aggregated IP address distribution by AS

1110 AS own more than 83% of the announced IP addresses

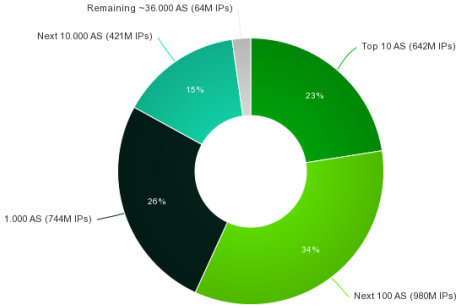

36k AS own only 2% of the anounced IPs

Ingredients

- You need to be your own Autonomous System (AS)
- You need to bring public IP space to the exchange
- You need to connect via OpenVPN to the exchange

OpenVPN setup

- OpenVPN topology subnet mode
- Scales well horizontally by adding more servers
- OpenVPN as it is the most trusted VPN protocol –
  – even with OpenSSL et al

Routing

- Requires multihop BGP sessions by default to allow extension with additional address space
- Redundant route reflector setup
  - default deny
  - reverse top-talker list, enable peering starting with lowest traffic peers for IPv4 –
    – IPv6 should not be an issue at all

Issues

- Scalability, higher latencies, reduced performance
- Route fragmentation, TCAM issues in widely used devices
- General net quality, shitty asymmetric routing
- Default setup susceptible to DDOS
  - might be worth to blacklist IX networks for rest of the internet
- Traffic cost
  - needs to be for free to get a sizable test crowd
- Needs more than one instance – lets fragment BGP :P

Where to take it from here

- We have an AS to use, around 40Gbits of bandwidth, a trusted location and can provide power and VPN connectivity
- We need feedback, contact to parties interested in testing, etc pp