

*High-speed traffic encryption on x86_64
with Snabb*

FOSDEM 20

Max Rottenkolber

max@mr.gy

@eugeneia_

whoami



Max Rottenkolber <max@mr.gy>

Open source hacker, working on Snabb since 2014

Consulting on software networking (in userspace),
protocols, optimization...



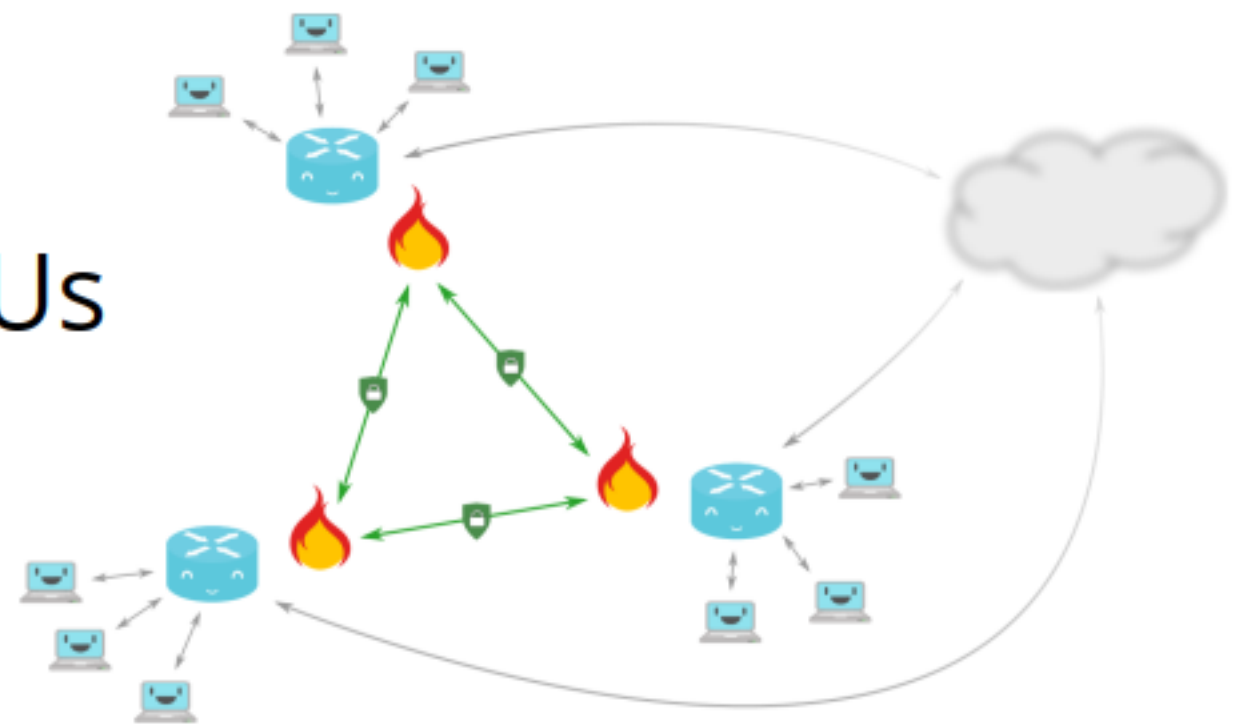
**inter—
stellar**

Vita

Vita is a high-performance site-to-site VPN gateway

Fully open source (and hackable!)

Runs on generic x86_64 server CPUs




Vita 

Based on  **snabb**

Written in a high-level language (Lua)


RAPTORJIT

Made possible by  **nl-net**
FOUNDATION



snabb

```
while not link.empty(input) do
  local p = link.receive(input)

  if ipv4_ttl(p) > 0 then
    link.transmit(output, p)
  else
    link.transmit(time_exceeded, p)
  end
end

end
```

Vita

~3 Mpps per core on a modern CPU (duplex)

...or ~5 Gbps of IMIX traffic per core

Medium-term goal: 100 Gbps at 60 byte packets on a generic x86 server

How?

In Snabb-land we like to write software that is both fast and simple

...and we don't like vendor lock-in

No QuickAssist, crypto cards... Only x86_64!

How?

For crunching numbers (encryption): AES-NI, AVX2
(optimized AES-GCM implementation written
in DynASM)

```
function ghash_mul(Dst, gh, hk, t1, t2, t3)
| vpcmqlqdq xmm(t1), xmm(gh), xmm(hk), 0x11
| vpcmqlqdq xmm(t2), xmm(gh), xmm(hk), 0x00
| vpcmqlqdq xmm(t3), xmm(gh), xmm(hk), 0x01
| vpcmqlqdq xmm(gh), xmm(gh), xmm(hk), 0x10
| vpxor xmm(gh), xmm(gh), xmm(t3)
...
```


How?

For route lookups (longest prefix match):
Optimized Poptrie implementation (again, DynASM)

```
function lookup (Dst, Poptrie, keysize)
    if Poptrie.direct_pointing then
        -- v = extract(key, 0, Poptrie.s)
        local direct_mask = bit.lshift(1ULL, Poptrie.s) - 1
        -- v = band(key, direct_mask)
        | mov v_dw, dword [key]
        | and v, direct_mask
```

....

How?

RaptorJIT + FFI

(simple and fast implementation of IPsec ESP)

```
esp_head = ffi.typeof[[
    struct {
        uint32_t spi;
        uint32_t seq_no;
    } __attribute__((packed))
]]
```

```
esp_tail = ffi.typeof[[
    struct {
        uint8_t pad_length;
        uint8_t next_header;
    } __attribute__((packed))
]]
```

How?

DSL for match-action pipeline, based on pcap-filter(7)
language (code generation at runtime)

```
pf_match.compile([[match {  
  ip dst host %s and icmp => icmp4  
  ip dst host %s => protocol4_unreachable  
  ip => forward4  
  arp => arp  
  otherwise => reject_ethertype  
}]]):format(conf.node_ip4, conf.node_ip4))
```

How?

Problem: can not parallelize SA

Every packet on an SA gets a unique sequence number

Synchronization problem if spread across cores

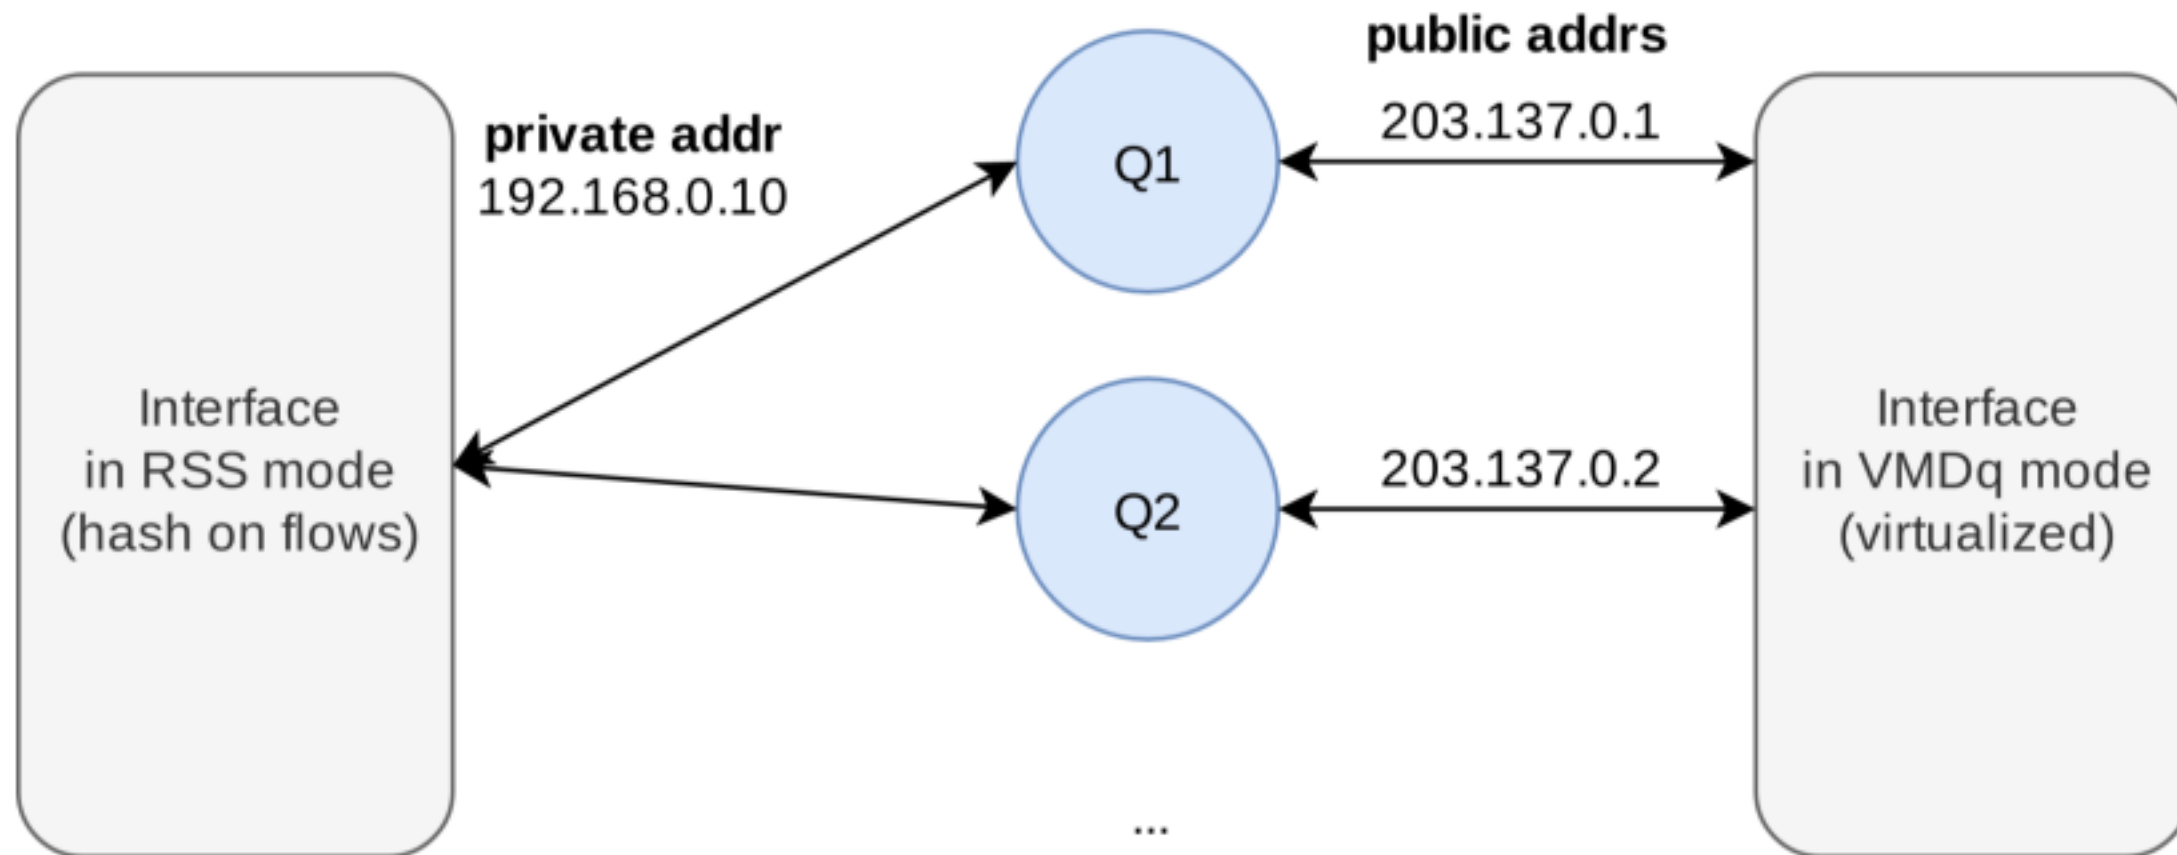
How?

Solution: scale out (multiple SAs per route)

RSS on private interface: distribute onto SAs

VMDq on public interface: aggregate SAs

How?



Network Drivers

The Snabb way: simple drivers written in Lua

New: Snabb drivers for XDP, Intel AVF

XDP?

Immediate goal: make Vita easily deployable in cloud

XDP: - a bit heavy to setup, currently some limitations
- working with kernel upstream looks promising (kudos to Björn Töpel)

AKE (authenticated key exchange)

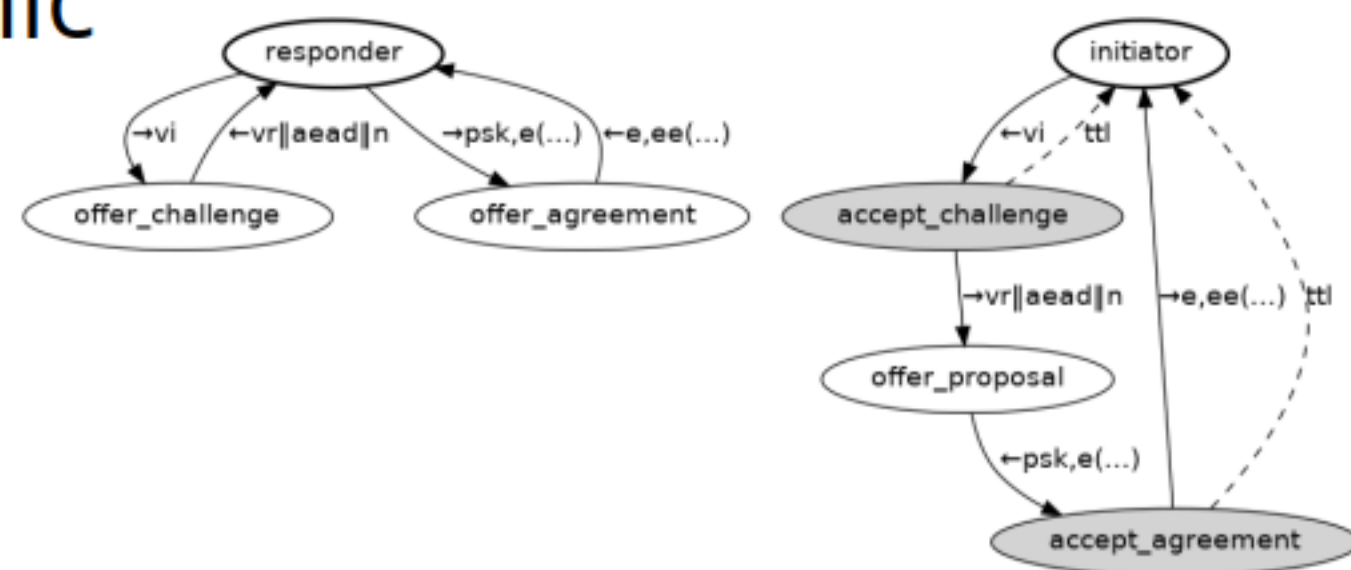
Cycle SAs often and without losing packets
(perfect forward secrecy)

Low-throughput, but largest attack surface

AKE (authenticated key exchange)

Simple PSK based protocol
(based on the Noise protocol framework)

Using minimal set of cryptographic primitives, constructions and protocol implemented in Lua



AKE (authenticated key exchange)

Alternatively: IKEv2 via StrongSwan

SWITCH engineer Alexander Gall developed
StrongSwan plugin+interop with Snabb

Configuration & operation

Based on a YANG model

...includes runtime statistics

```
module vita-esp-gateway {  
    ...  
}
```

Configuration & operation

Query/update configuration via RPC

Query runtime statistics via RPC

```
$ snabb config get-state vita /gateway-state/private-interface
```

Thanks!

Get involved:

github.com/snabbco/snabb

github.com/inters/vita

Get support and consulting:

<https://inters.co>

Email me:

max@mr.gy

Gritty details on my blog:

<https://mr.gy/blog>

**inter—
stellar**