Indefensible Neighbors

Joel Jaeggli
Fastly

Intro

If it feels like I’ve been harping on this for awhile, It’s because I have.

RFC 6583

Draft-jaeggli-v6ops-indefensible-nd

All goes back to RFC 3756

14 years later what to do we have?

Problem

Since the definition of Interface IDs and accompanying subnet sizes in RFC 1885, the potential has existed for the forwarding and control plane resources of a router to be greatly exceeded by locally or remotely triggered attempts to discover connected neighbors.

Well understood by the time of RFC 3756

https://en.wikipedia.org/wiki/Where%27s_Wally%3F

We rediscover this problem time and again.

ARP

ND

IGMP

MLD

MSDP

This is not a new problem.

IPv4 has it.

Appletalk has it.

We know that building large or chatty a broadcast domains is an expensive proposition.

3rd Century BCE Greek Citadels have mitigations for it.

Various Approaches

Rate limit discovery under duress - RFC 6583

Make subnets link local - RFC 7404

Make subnets very very small - RFC 7608

Throw out ND

- Bind prefixes to loopback and route to it
- Address registration (6Lowpan)

Firewalls

The solution is generally routing...

It is supposed; that we like discovery, because it is flexible and doesn’t require state until it is used.

Bottleneck the problem so that it’s small enough to contain.

Take the load off the control-plane.

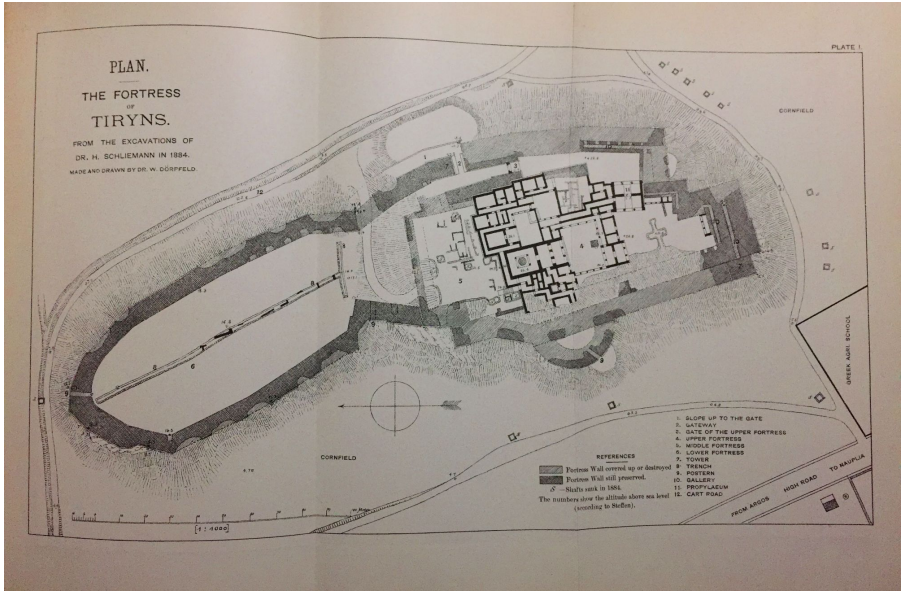

Citadel at Tiryns

Why discuss now?

Obvious that Datacenter / Network operators adopt various approaches to this problem.

Intersects with addressing architecture in funny ways.

6man - discussion “Re: SAILing LAPs”

- draft-bourbaki-6man-classless-ipv6
- Draft-carpenter-6man-lap

We’re Doomed?

If you’re building provider -> customer edges

Or

CPE

Did the IETF doom you to mediating this exposure with stateful firewalls?

Policing the control-plane is means reducing the efficacy / reliability of your discovery mechanism.

Playing with the Addressing architecture

Prefix / PD per host gives you something you can defend.

Binding a /64 to a loopback means you can respond to or discard as you wish

Null route on asic based hardware is as efficient as anything you can do anywhere to discard packets.

Playing with the addressing architecture

Longer than /64 IIDs violates RFC 2460 with respect to SLAAC

You can certainly employ longer prefixes if you are prepared to assign them via stateful DHCPv6.

What did we do?

Link local only subnets.

Subnets are locally significant, loopback addressing is globally significant.

All globally significant assigned addresses are applied to loopback, routed.

Data centers may use SLAAC / DHCP for bootstrapping but other tools can take over from there.

```
lo        Link encap:Local Loopback
          inet addr:127.0.0.1  Mask:255.0.0.0
          inet6 addr: 2a04:4e40:8050:0:2ae:6322:c8ee:506/128
Scope:Global
          inet6 addr: 2a04:4e47:3::/48 Scope:Global
          inet6 addr: 2a04:4e40:8010:0:2ae:6322:c8ee:506/128
Scope:Global
          inet6 addr: 2a04:4e40:8030:0:2ae:6322:c8ee:506/128
Scope:Global
          inet6 addr: 2a04:4e40:8010::/128 Scope:Global
          inet6 addr: ::1/128 Scope:Host
          inet6 addr: 2a04:4e47:2::/48 Scope:Global
          inet6 addr: 2a04:4e40:8030::/128 Scope:Global
          inet6 addr: 2a04:4e40:8050::/128 Scope:Global
          inet6 addr: 2a04:4e42::/32 Scope:Global
          inet6 addr: 2a04:4e40:80f0:0:2ae:6322:c8ee:506/128
Scope:Global
          inet6 addr: 2a04:4e40:8040:0:2ae:6322:c8ee:506/128
Scope:Global
          inet6 addr: 2a04:4e47:1::/48 Scope:Global
          inet6 addr: 2620:12a:8000::/44 Scope:Global
          inet6 addr: 2a04:4e40:8000:0:2ae:6322:c8ee:506/128
Scope:Global
          inet6 addr: 2a04:4e40:8000::/128 Scope:Global
          inet6 addr: 2a04:4e40:80f0::/128 Scope:Global
          inet6 addr: 2a04:4e40:8020:0:2ae:6322:c8ee:506/128
Scope:Global
          inet6 addr: 2a04:4e40:8020::/128 Scope:Global
          inet6 addr: 2a04:4e47::/48 Scope:Global
          inet6 addr: 2a04:4e40:8040::/128 Scope:Global
```

Getting to know your neighbors...

Neighbors have certain properties.

They share common resources such as subnets.

They can be subject to admissions controls.

They may share a common administrative domains, or at least have mutuality beneficial interests.

They mutually accumulate state.

The internet as a whole shares few of these properties.

Won't You Be My Neighbor? - 2018