ESnet
ENERGY SCIENCES NETWORK

Minding our MANRS

Dale W. Carder

ESnet Network Engineering

dwcarder@es.net

Internet2 TechEx '23

2023-09-21

U.S. DEPARTMENT OF ENERGY
Office of Science

BERKELEY LAB

ESnet Customer Provisioning

- All customer prefixes are checked against RIR database
  - We do not accept a "Letter of Authorization"
- Prefixes entered into our database
- Provisioning is fully orchestrated
- BCP38 ACL's applied on interfaces
- explicit BGP import filters
- All customers MUST have route/route6 IRR objects and get included into ESnet's as-set

How do 3rd party networks find our prefixes?

```
aut-num:    AS293
as-name:    ESNET
descr:      Energy Sciences Network
export:     to AS-ANY   announce AS293:AS-ESNET
admin-c:    ESnet Network Engineering Group
tech-c:     ESnet Network Engineering Group
notify:     hostmaster@es.net
mnt-by:     MAINT-ESNET
changed:    dwcarder@es.net 20230915  #18:51:23Z
source:     RADB
```

peeringdb.com

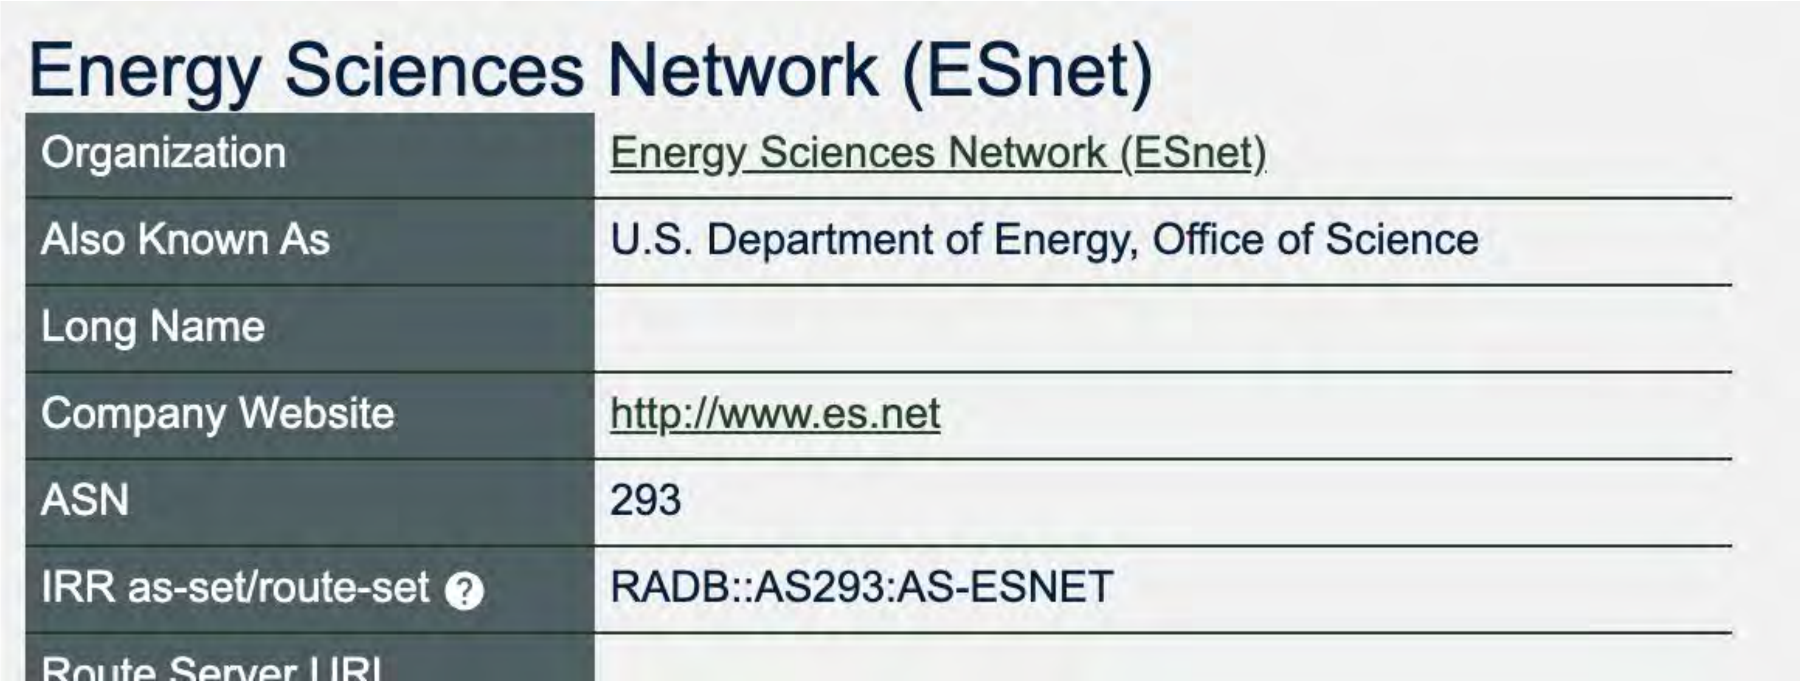

Energy Sciences Network (ESnet)

Organization	Energy Sciences Network (ESnet)
Also Known As	U.S. Department of Energy, Office of Science
Long Name	
Company Website	http://www.es.net
ASN	293
IRR as-set/route-set ?	RADB::AS293:AS-ESNET
Route Server URL	

```
as-set:     AS293:AS-ESNET
descr:      AS Cone of ESnet
members:    AS16, AS43, AS44, AS45, AS50, AS68, AS291,
AS292, AS293, AS377, AS513:AS-CERN-NREN, AS683, AS2640,
AS2936, AS2937, AS3152, AS3380, AS3425, AS3428, AS3424,
AS3431, AS3443, AS3445, AS3562, AS3671, AS3970, AS6406,
AS10702, AS11678, AS14702, AS16411, AS32982, AS36288,
AS46846, AS54297, AS62555, AS63331, AS396098, AS398900,
AS400066, AS-PEERING-TESTBED
tech-c:     ESnet Network Engineering Group
mnt-by:     MAINT-ESNET
changed:    dwcarder@es.net 20230911  #16:15:51Z
source:     RADB
```

```
> bgpq4 -A -6 AS293:AS-ESNET | wc -l
      80

> bgpq4 -A -4 AS293:AS-ESNET | wc -l
     271
```

bgpq4 can automagically create filters

```
> bgpq4 -A -6 AS-ESNET
no ipv6 prefix-list NN
ipv6 prefix-list NN permit 2001:400::/32
ipv6 prefix-list NN permit 2001:67c:2c4::/48
ipv6 prefix-list NN permit 2001:7fb:fd02::/48
```

... and so on

json format:

```
> bgpq4 -A -6 -j AS-ESNET
{ "NN": [
   { "prefix": "2001:400::\/32", "exact": true },
   { "prefix": "2001:67c:2c4::\/48", "exact": true },
   { "prefix": "2001:7fb:fd02::\/48", "exact": true },
   ... and so on
```

Recently published our peering policy

https://www.es.net/engineering-services/the-network/peering-connections/

- Peer must have a publicly assigned Autonomous System Number (ASN) from a Regional Internet Registry (RIR).
- All prefixes announced must be publicly routable and properly registered with the corresponding RIR.
- Prefixes will be exchanged over BGP.
- Peer must maintain an AS-SET and keep up to date entries in the Internet Routing Registry (IRR) system for all prefixes announced.
  - Prefixes not registered in the IRR system will not be accepted.
  - Prefixes that are RPKI invalid will not be accepted.
- Peer must maintain an up to date PeeringDB entry, including a 24x7 NOC Contact, AS-SET, and prefix limits.
- Unless specifically agreed upon beforehand, peers are expected to peer in all locations where mutually present and announce a consistent set of prefixes at all locations.
- Peer must adhere to MANRS industry standards for route security, including BCP38 filtering of its customer cone.

Peer Network Provisioning

- AS Number, peer type, other bgp intent entered into database
- Orchestration process
  - Looks in PeeringDB for IRR object
  - calls bgpq4 to get prefixes (returned as json)
    - rpki-invalid prefixes are dropped
  - provisions explicit BGP import policy & prefix accept lists
- **Any prefix you send us not in the list is dropped!**

ESnet

In our Internet DFZ VRF (mix of commercial & R&E peers):

- 127 peer networks (non-customer)
- 110 have entries in PeeringDB

- What if a PeeringDB as-set entry doesn't exist?
  - we can statically configure the IRR object in our database
    - 8 networks only define their as-set in their aut-num object
    - 2 we figured out by inverse query of mnt-by records
  - or, just do lookup of RIR data (only works for stub asn's, though)
    - 2 networks fall into this category

ESnet

THE "WALL OF SHAME"

Only 5 ESnet peers have no discernable IRR as-set object:

- NASA
- TWAREN
- CUDI
- TRANSPAC
- INTERNET2 <------ you are here

Why doesn't Internet2 have an as-set?

- Maybe nobody notices traffic taking commercial paths?
  - *possibly*, see Steve's talk about routing intent from Weds

- Maybe I2 doesn't care about routing security?
  - *highly unlikely*

- Maybe I2 doesn't know who their customers are?
  - *highly unlikely*

ESnet

Maybe there's a lot of legacy stuff and it's just very very hard to make an as-set for Internet2?

- bazillion asn's
- connector networks
- regional networks
- state networks
- campuses
- k12's
- and so on

Maybe there's a lot of legacy stuff and it's just very very hard to make an as-set for Internet2?

FALSE

ESnet

Maybe there's a lot of legacy stuff and it's just very very hard to make an as-set for Internet2?

FALSE

Proof: I created one!

I created an as-set for Internet2!

```
as-set:     AS293:AS-FROM-I2-TO-ESNET
members:    AS-CALREN2
members:    AS-CARNE
members:    AS-DARTMOUTH
members:    AS-FLRnet-Aggregate
members:    AS-FRGP
members:    AS-GPN-PEERS
members:    AS-ICCN
members:    AS-INGIG
members:    AS-KINBER
members:    AS-LONI-members
   ... and on & on
```

Other recent example issues in R&E

- Backup paths and other special arrangements complicate things
  - GEANT sending ESnet extra R&E reachability
    - Some of it turned out to be a mistake
    - Some of it is intentional
    - ESnet, of course, drops all of these prefixes as unauthorized
  - NORDUnet's IRR object potentially missing entries
    - ESnet, of course, drops all of these prefixes as unauthorized

- These issues cause asymmetry & complicate troubleshooting

It's time to get real about routing security

- R&E Networks do lag behind the commercial sector
    - even I2PX has an IRR object: AS11164:AS-ALL

  – lack of prefix filtering (GEANT)

  – lack of ROA's

- With some motivation we can reap the benefits.

extra slides if we have additional time

Things ESnet still needs to improve #1

- Update prefix lists more regularly
  - Right now, we update filters ~quartery or as-needed
  - Need to do this way more regularly, like nightly
  - Waiting on a better way to safely run bulk automation jobs
- AS-Path filtering from customers
  - While we do filter prefixes, make sure they can only use their AS #'s
- Tier-1 as-path filtering from peers
  - Best practice: filter out paths where a "tier-1" is in the AS-PATH
  - ESnet does peers with quite a few of these, so this will take us some extra effort to add to our automation and not drop traffic

Things ESnet still needs to improve #2

- RPKI drop-invalid
  - validators are running, but need effort to add drop-invalid policy via provisioning automation
- Sign ROA's for our IPv4 prefixes
  - just submitted updated LRSA 9/20!!!
- Announce / sinkhole address space we use internal to DOE but does not appear on the external Internet
  - Some IP space looks "not in use" and attractive to steal
  - Maybe sign AS 0 ROA first, but it may be more interesting to capture traffic.