Intent Based Networking - the technology

NANOG™

Jeff Tantsura
Head of Networking Strategy @Apstra
Chair IETF Routing and RIFT working groups, IAB

Why IBN?

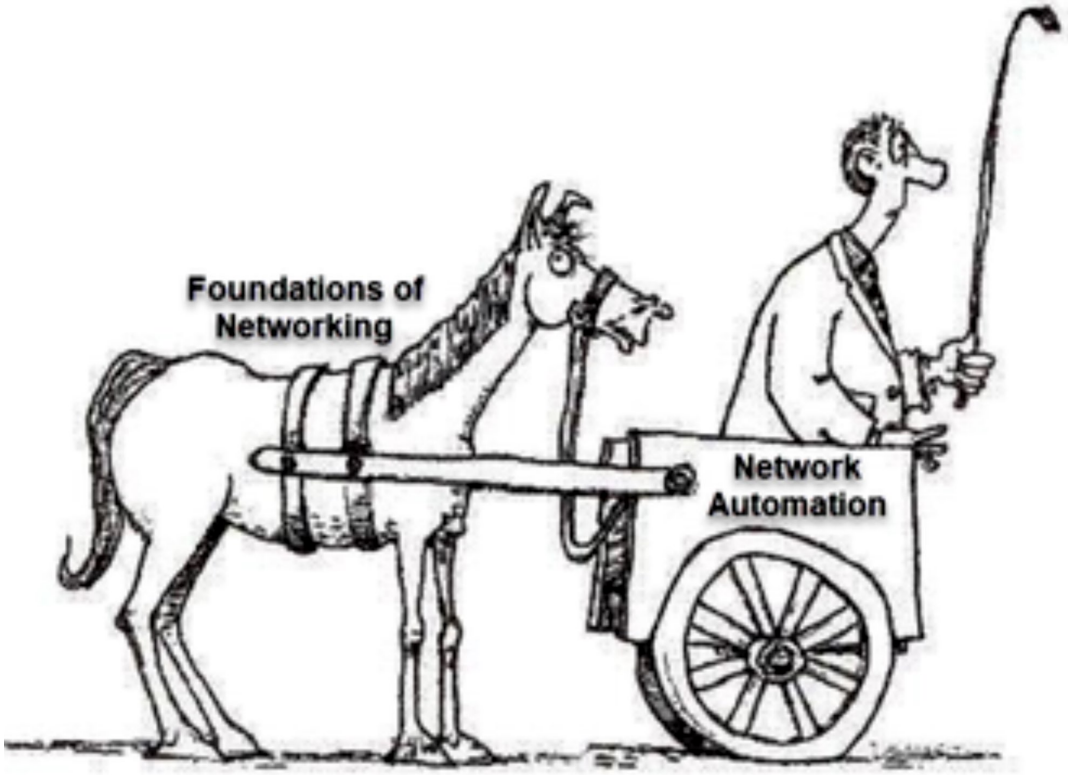

Source: Twitter

Why IBN?

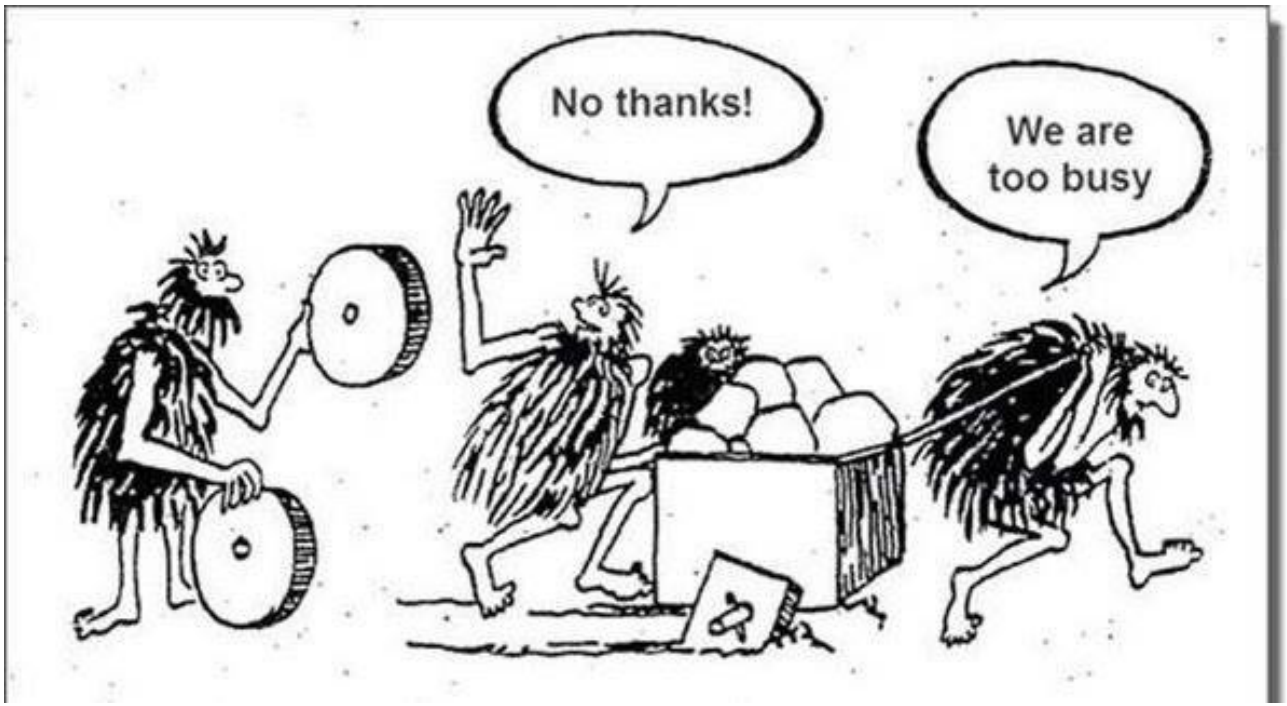

Source: Twitter

IBN Landscape

Source: Twitter

IBN Landscape

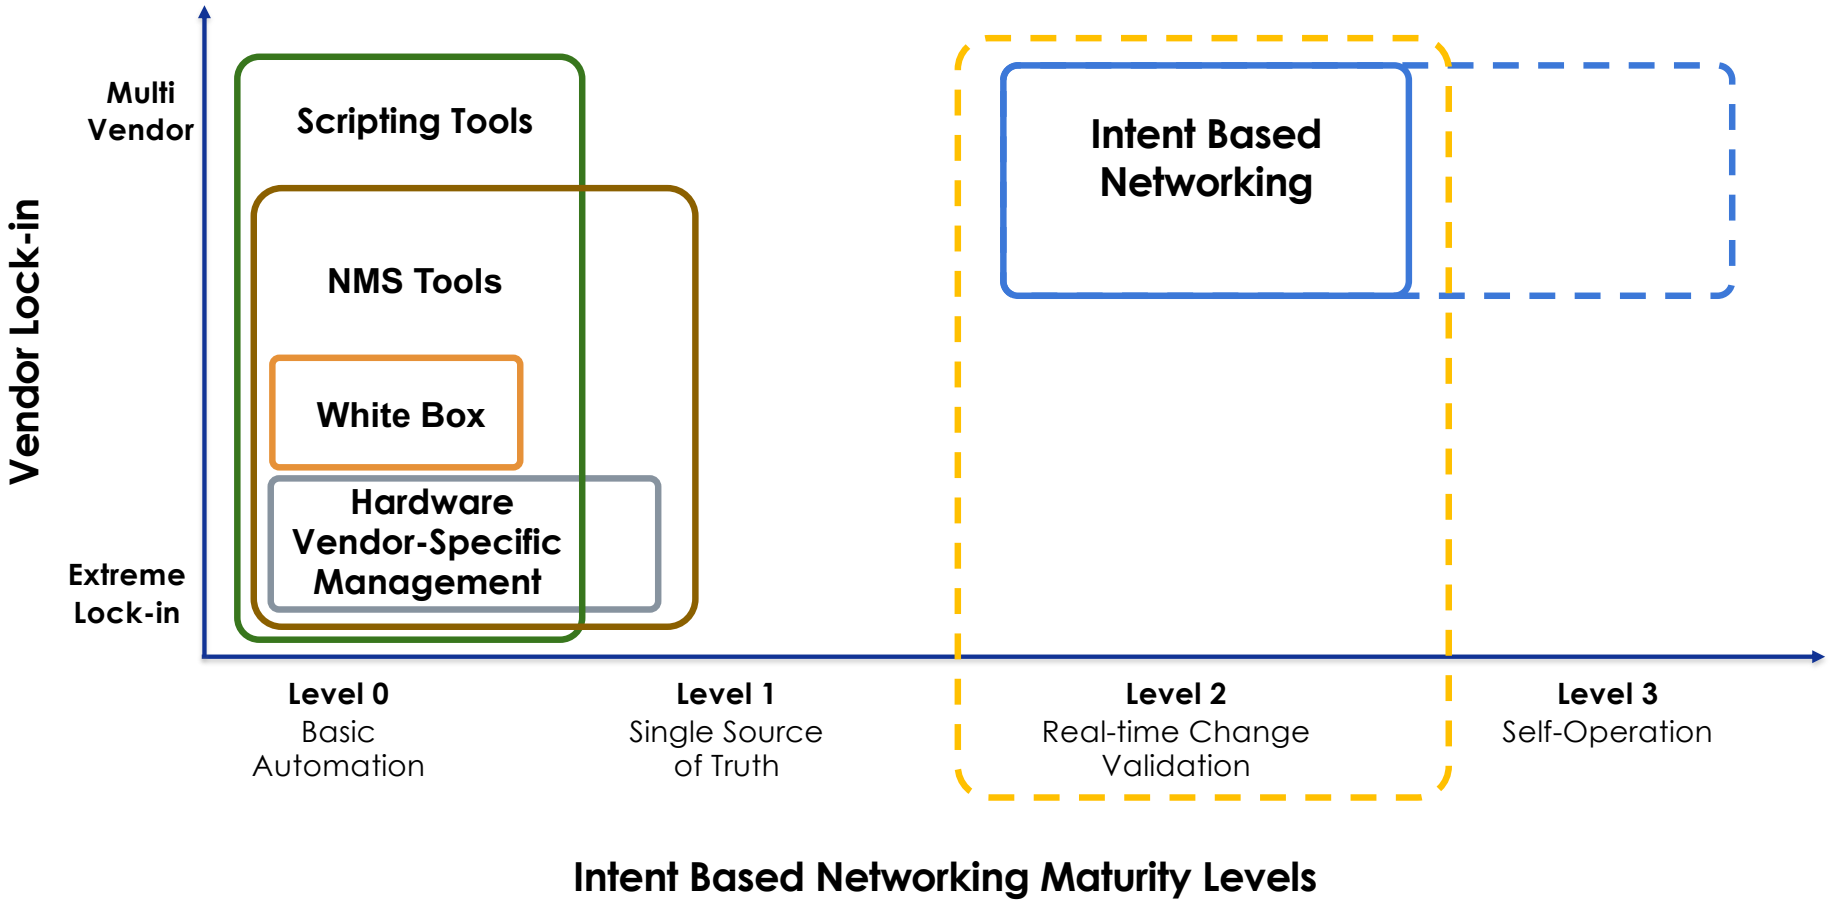

IBN standardization – just the beginning

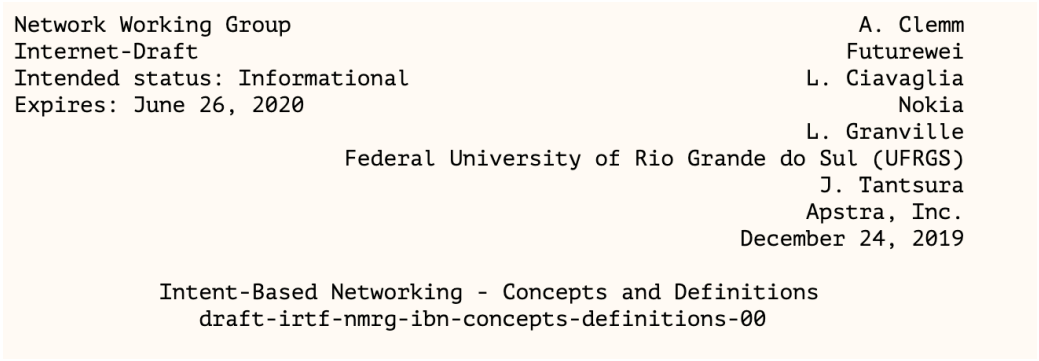

Network Working Group — A. Clemm
Internet-Draft — Futurewei
Intended status: Informational — L. Ciavaglia
Expires: June 26, 2020 — Nokia
L. Granville
Federal University of Rio Grande do Sul (UFRGS)
J. Tantsura
Apstra, Inc.
December 24, 2019

Intent-Based Networking - Concepts and Definitions
draft-irtf-nmrg-ibn-concepts-definitions-00

← Has recently been adopted as the rg document

Outlines number of fundamental IBNS building blocks and their relationship:

- SSoT: Single Source of Truth - A functional block in an IBN system that normalizes users intent and serves as the single source of data (normalized intended state) for every consumer.
- IBA: Intent Based Analytics - Analytics that are defined and derived from user' intent and used to validate the intended state.
- PDP: Policy Decision Point – part of intent definition, technology agnostic.
- PEP: Policy Enforcement Point – technology/device aware (e.g ACL or FW rule).

IBN – why DC is a good starting point?

Ability to reason about Intent is a fundamental property of an IBNS!

Complex systems fail in mysterious ways ;-)

CLOS topologies are extremely regular/uniform and mutate (expand) in a very predictable way. PoD structure provides clear boundaries.

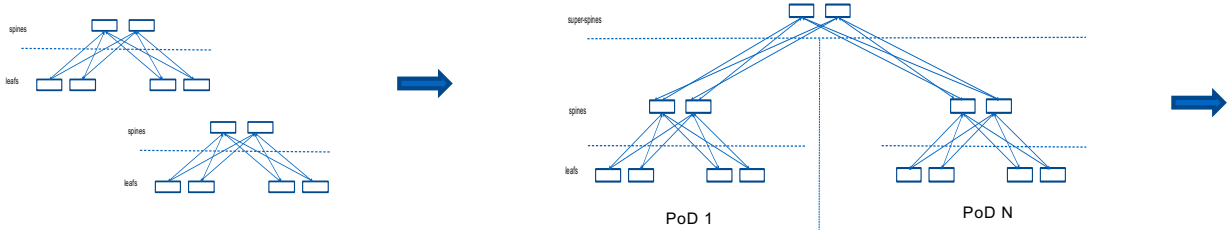

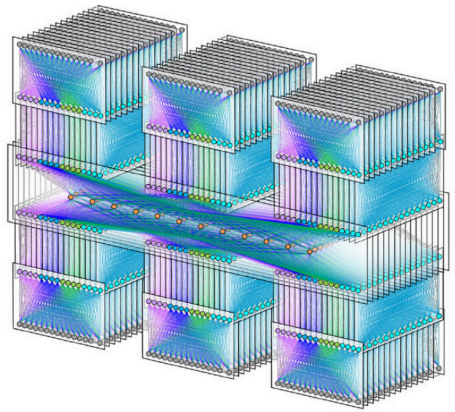

https://engineering.fb.com/data-center-engineering/f16-minipack/

IBN – why DC is a good starting point?

Telco WAN evolution

Day 0

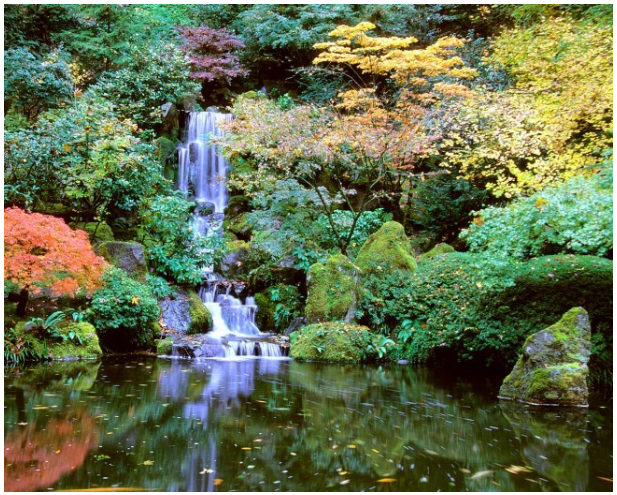

https://www.travelportland.com/article/portland-japanese-garden/

Day 2

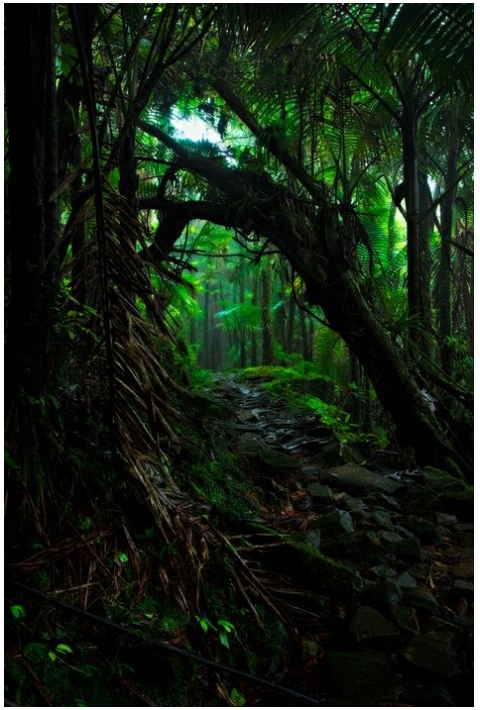

https://www.pinterest.com/pin/300756081335600951

IBN – why DC is a good starting point?

IBN in Telco WAN – is there a hope?

IETF TEAS ACTN framework is a step in the right direction

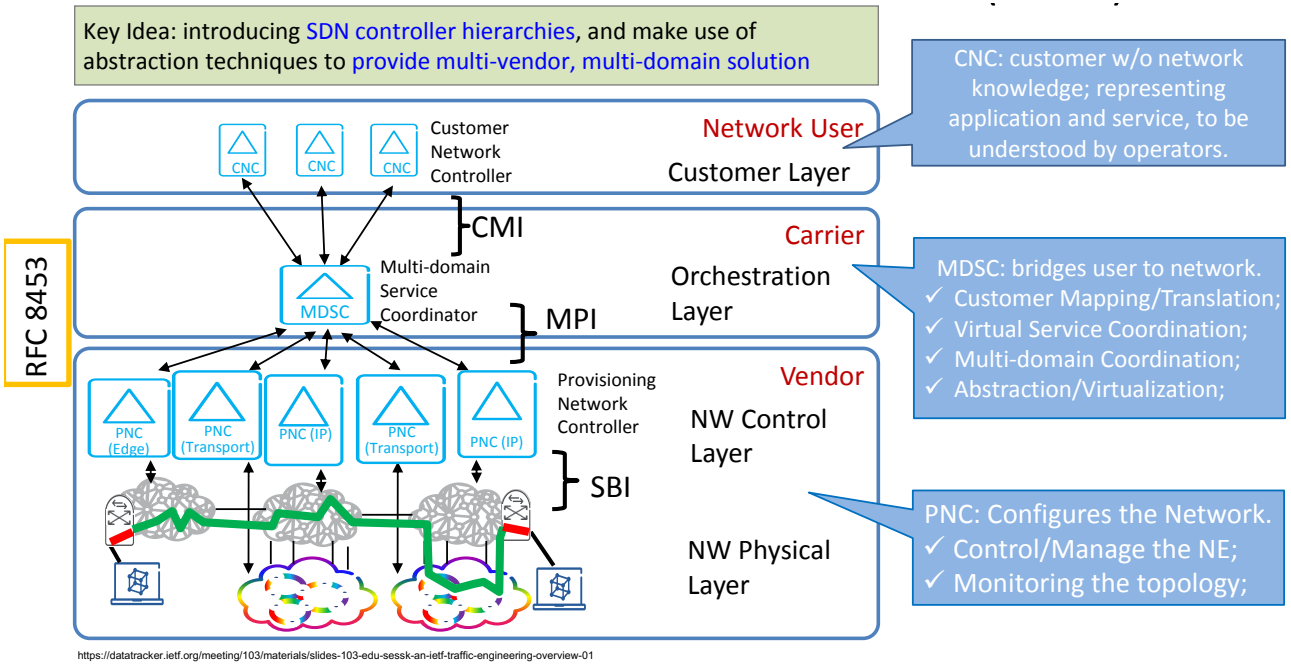

https://datatracker.ietf.org/meeting/103/materials/slides-103-edu-sessk-an-ietf-traffic-engineering-overview-01

IBN – why DC is a good starting point?

IBN in Telco WAN – is there a hope? What about 5G?

E2E Network Slicing is too complex, transport part of it is a perfect candidate for IBN

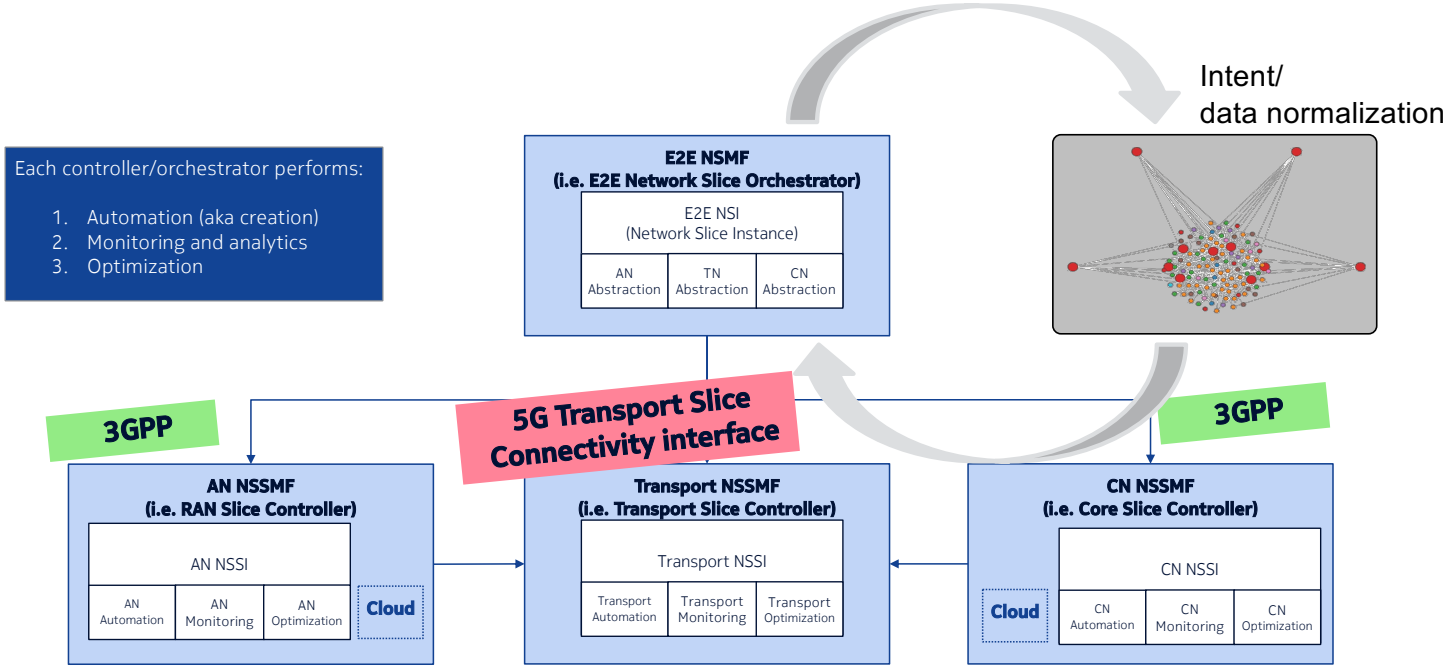

https://datatracker.ietf.org/meeting/105/materials/slides-105-teas-sessa-03-5g-transport-slice-connectivity-interface-00

IBN – why DC is a good starting point?

IBN in Telco WAN – 5G NS

E2E Network Slicing is too complex, transport part of it is a perfect candidate for IBN

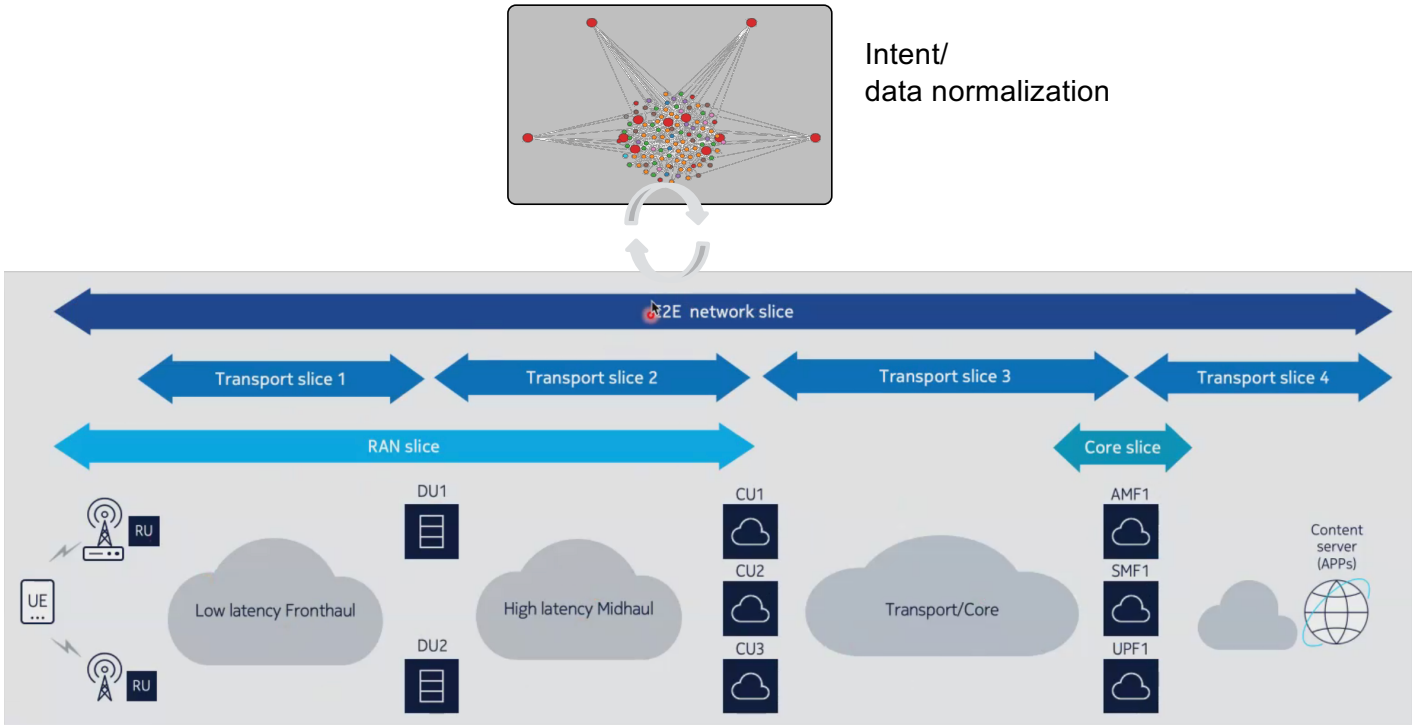

IBN Design Philosophy

Networks managed as a whole system, not individual components

Successful networks are defined by the outcomes produced by the whole system

Intent Based Networking is about "*what*" not "*how*"

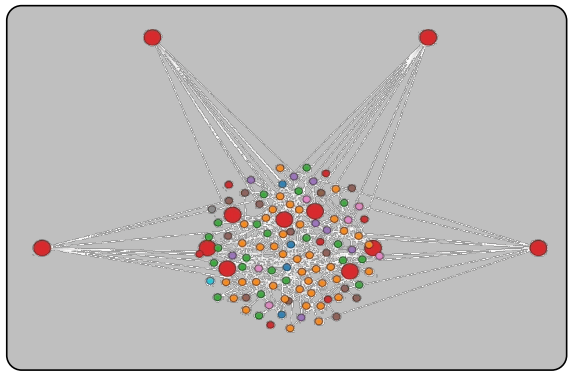

More details in: draft-irtf-nmrg-ibn-concepts-definitions

IBN life cycle

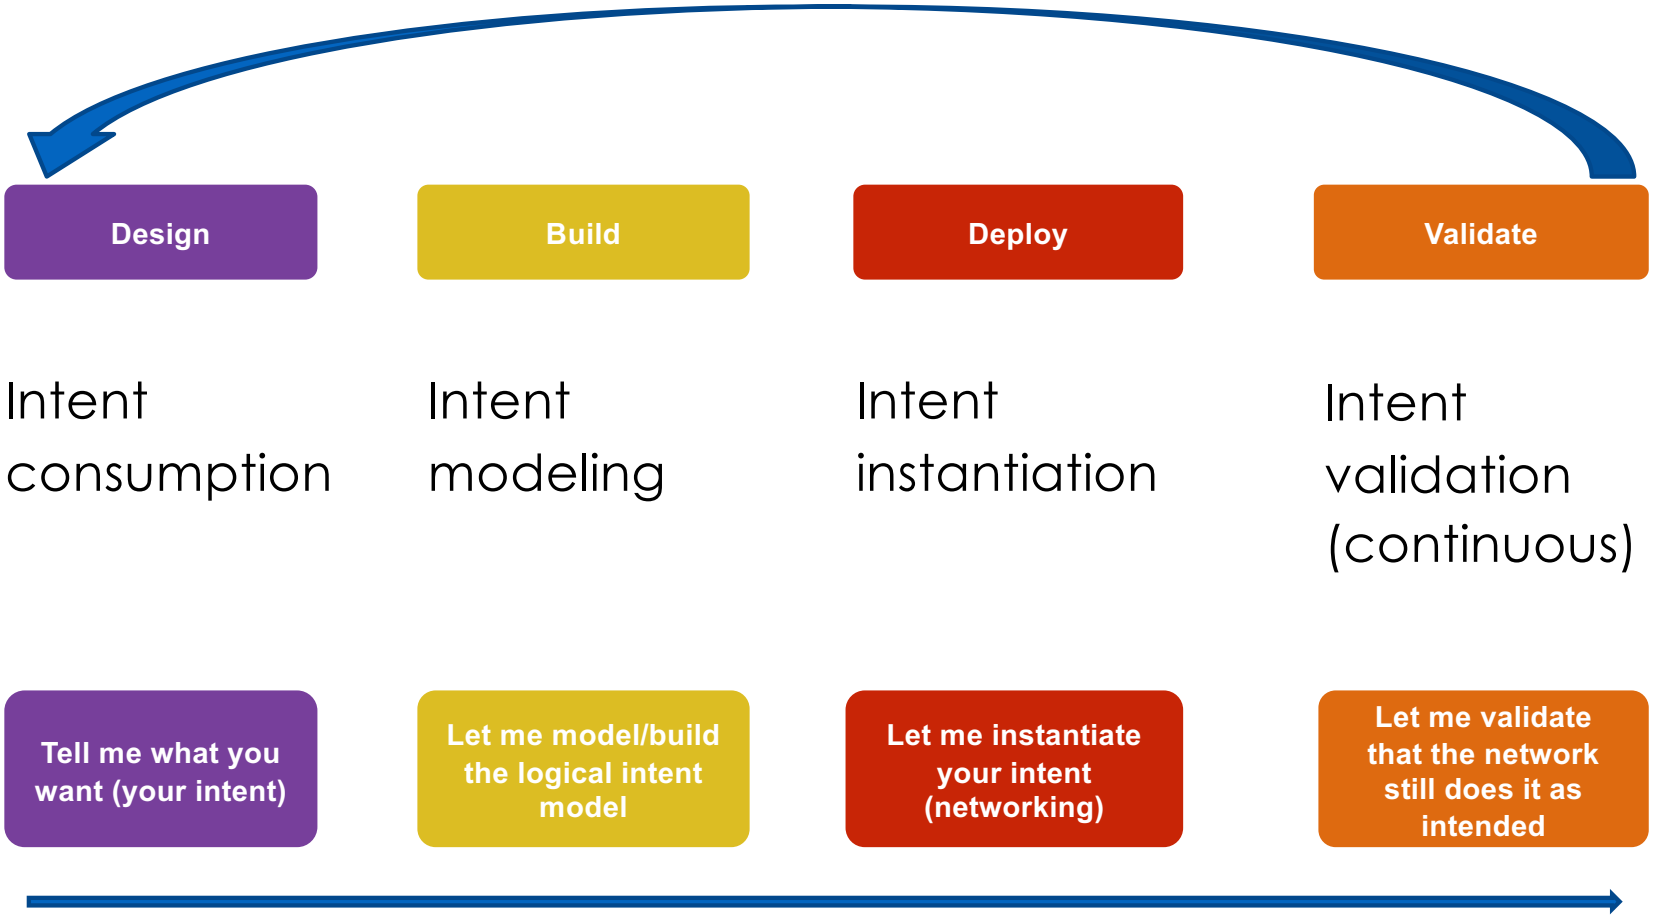

Day in the IBNS life

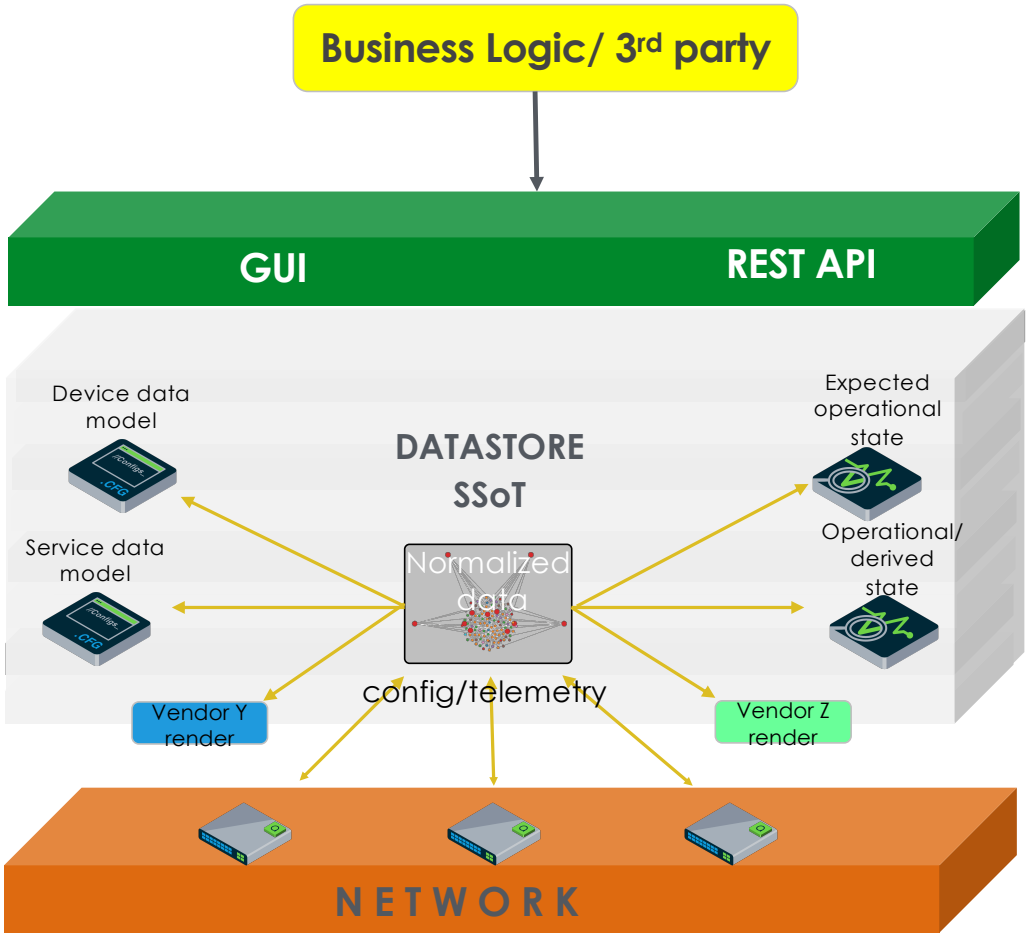

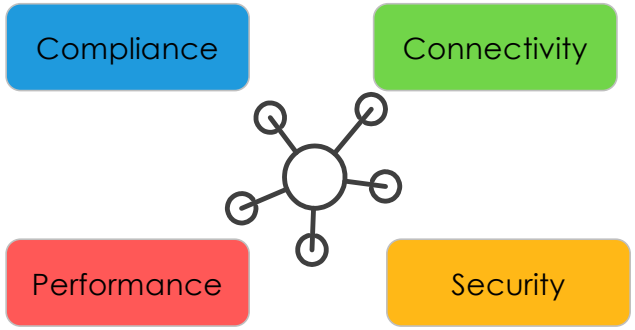

Programmable Network and Interfaces

Network

A network is programmable when the control and data planes provide an interface that allows the state of the network to be modified and monitored through a machine readable data-driven API's

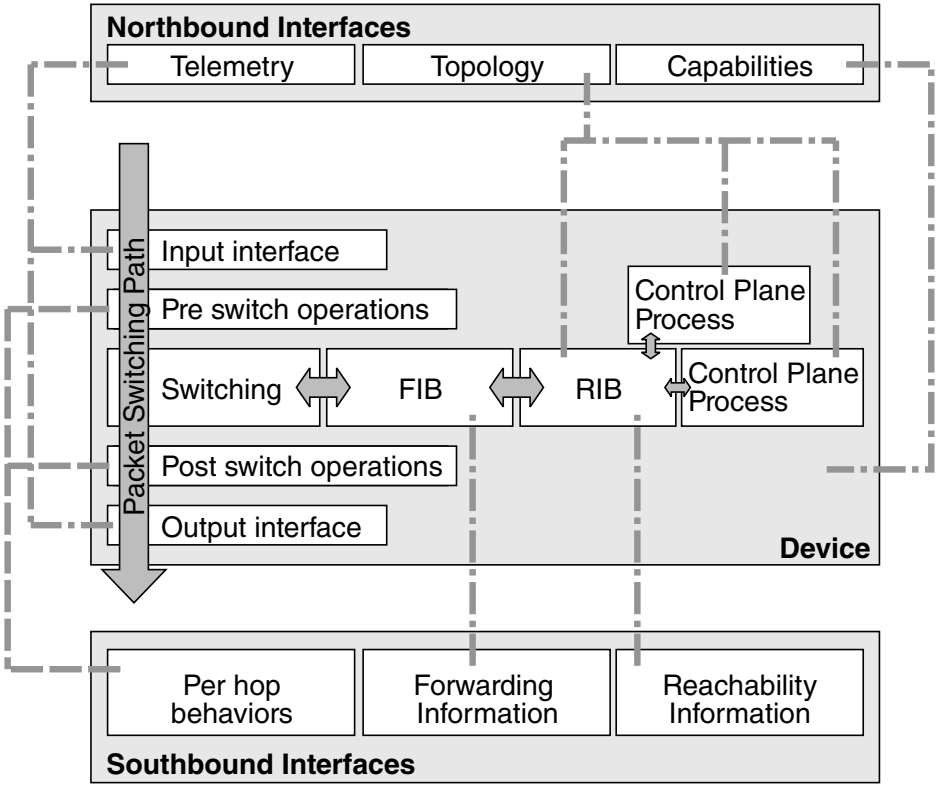

Source: Navigating Network Complexity: Book by Jeff Tantsura and Russ White

Programmable Network and Interfaces

Interfaces

Capabilities: Provides information what can be programmed or controlled. This included schema's and metadata, or rather information about how the control structures are organized, and how information is presented by network devices to the controller.

Inventory: Provides information about what devices are installed where in the network, potentially including any information about physical connections.

Topology: Provides information about the state of links connecting network devices. This includes all the artifacts of the topology described.

Telemetry: Includes operational state, counters, and other information about the current network state. This includes but not limited by: resources used/available, queue depths, delay, jitter, etc…

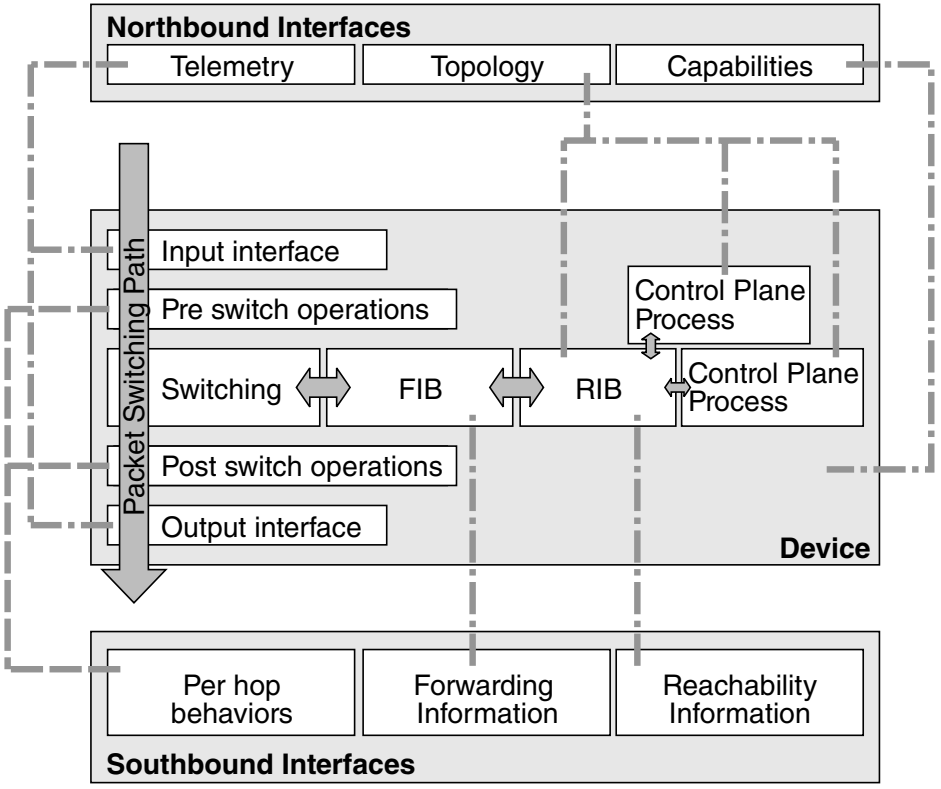

Source: Navigating Network Complexity: Book by Jeff Tantsura and Russ White

The need for data normalization

Source: food.com

➔ Fruitcake

The need for data normalization

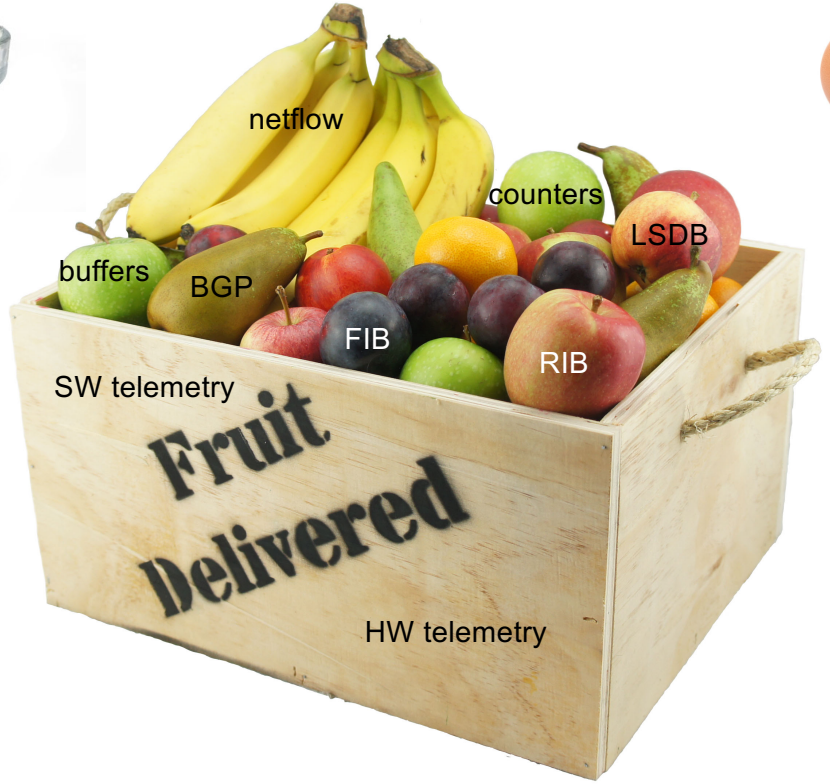

Source: Twitter

Architectural Goals of IBN

apstra

Problems to be solved:

- Composition/decomposition @scale
- Dealing with changes:
  - Planned change – can I achieve desired (future) state while preserving original intent (meeting SLO's)
  - Unplanned change – impact of the change, difference between intended and operational states, how to get to intended state (remediation/notification)

Architectural Goals of IBN

apstra

Problems to be solved:

- Closed loop validation:
  - continuously validate *outcomes* against the *intent* to ensure that the *composition* is working as intended
  - extract more knowledge by collecting less data thru IBA (Intent Based Analytics)
  - highly optimized SNR (signal to noise ratio) in analytics

Dealing With Scale?

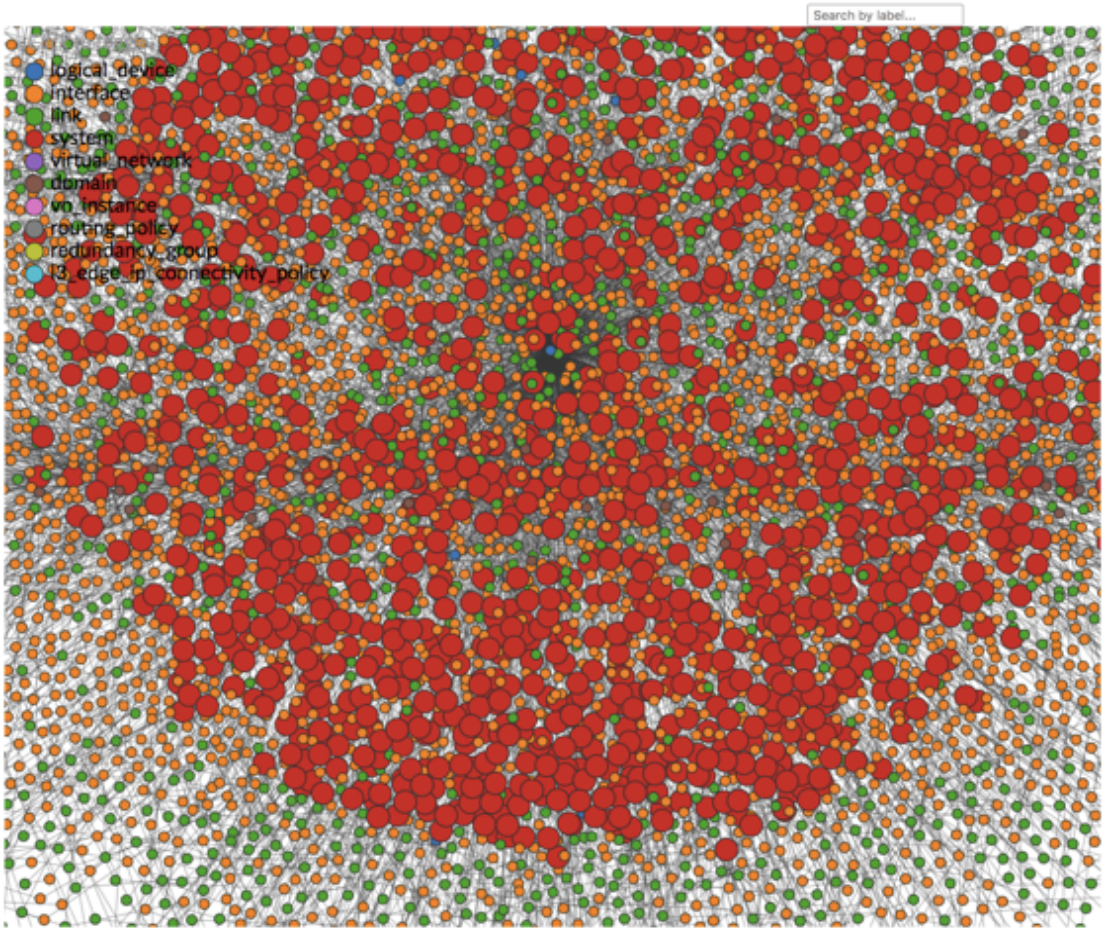

Composition

apstra

WIKIPEDIA
The Free Encyclopedia

Main page
Contents
Featured content
Current events
Random article
Donate to Wikipedia

Article Talk

Function composition (computer science)

From Wikipedia, the free encyclopedia

Not to be confused with object composition.

In computer science, **function composition** is an act or mechanism to combine simple functions to build more complicated ones. Like the usual composition of functions in mathematics, the result of each function is passed as the argument of the next, and the result of the last one is the result of the whole.

Programmers frequently apply functions to results of other functions, and almost all programming languages allow it. In some cases, the composition of functions is interesting as a function in its own right, to be used later. Such a function can always be defined but languages with first-class functions make it easier.

The ability to easily compose functions encourages factoring (breaking apart) functions for maintainability and code reuse. More generally, big systems might be built by composing whole programs.

Why model a graph?

- Networks are intuitively the connected set of **nodes** and **relationships**
- As network requirements **change** the model can be easily **extended**
- Efficiently run **queries** that were **not anticipated** at model design time

  Hint: you **will not** know all the queries at model definition time

Intent-> Graph composition

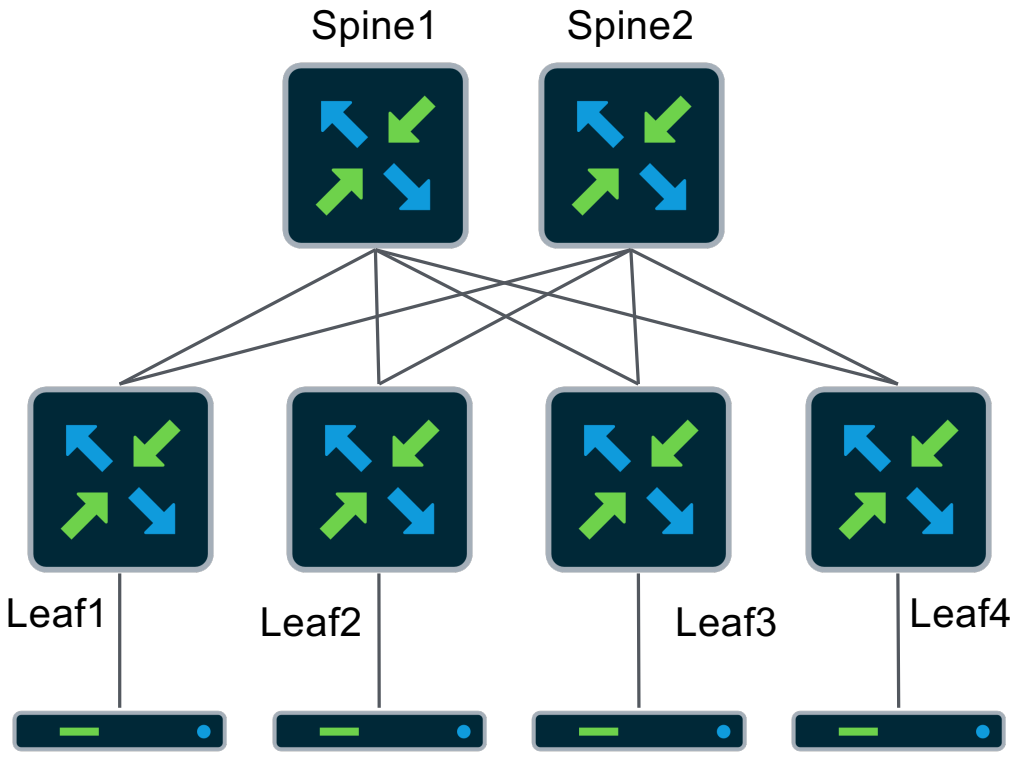

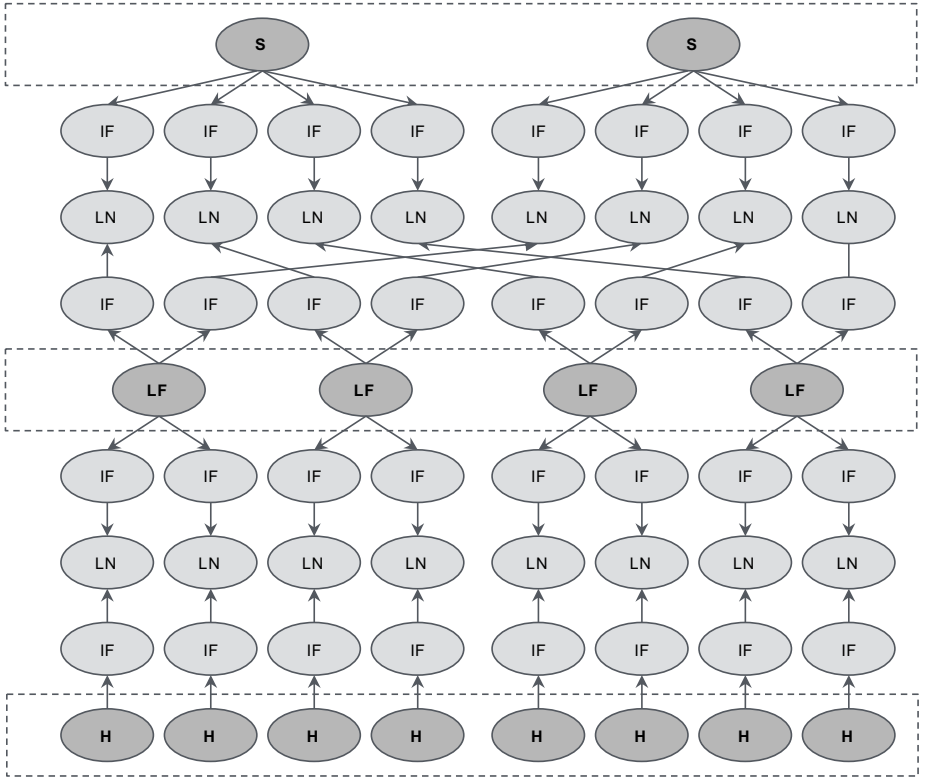

Function composition

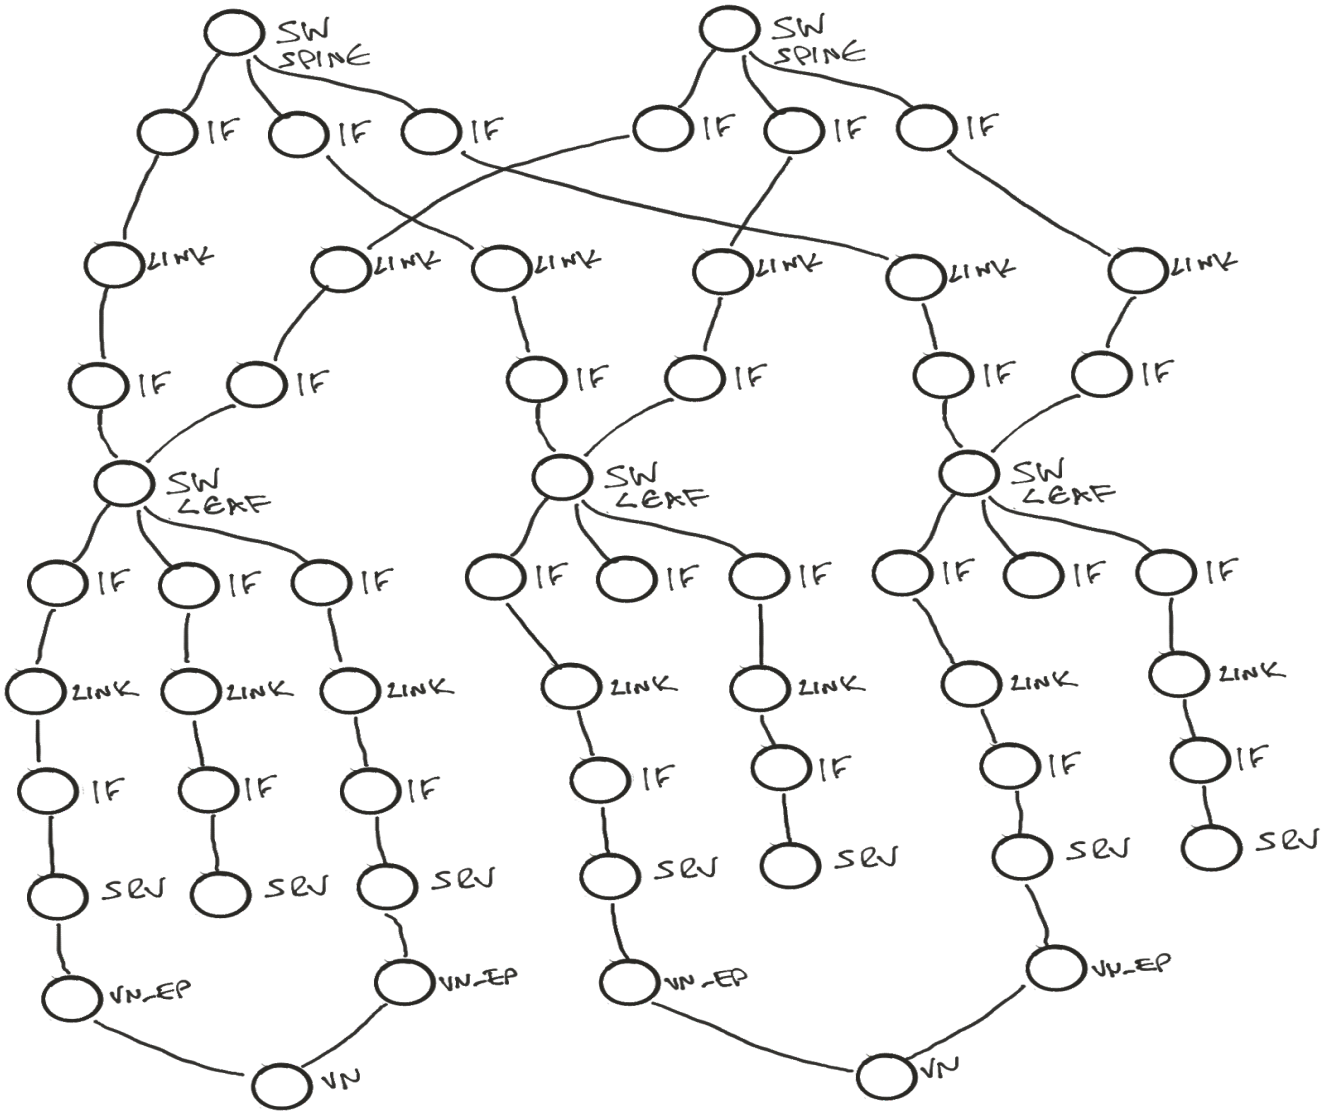
SW SPINE
SW SPINE
IF
LINK
SW LEAF
SRV
VN_EP
VN

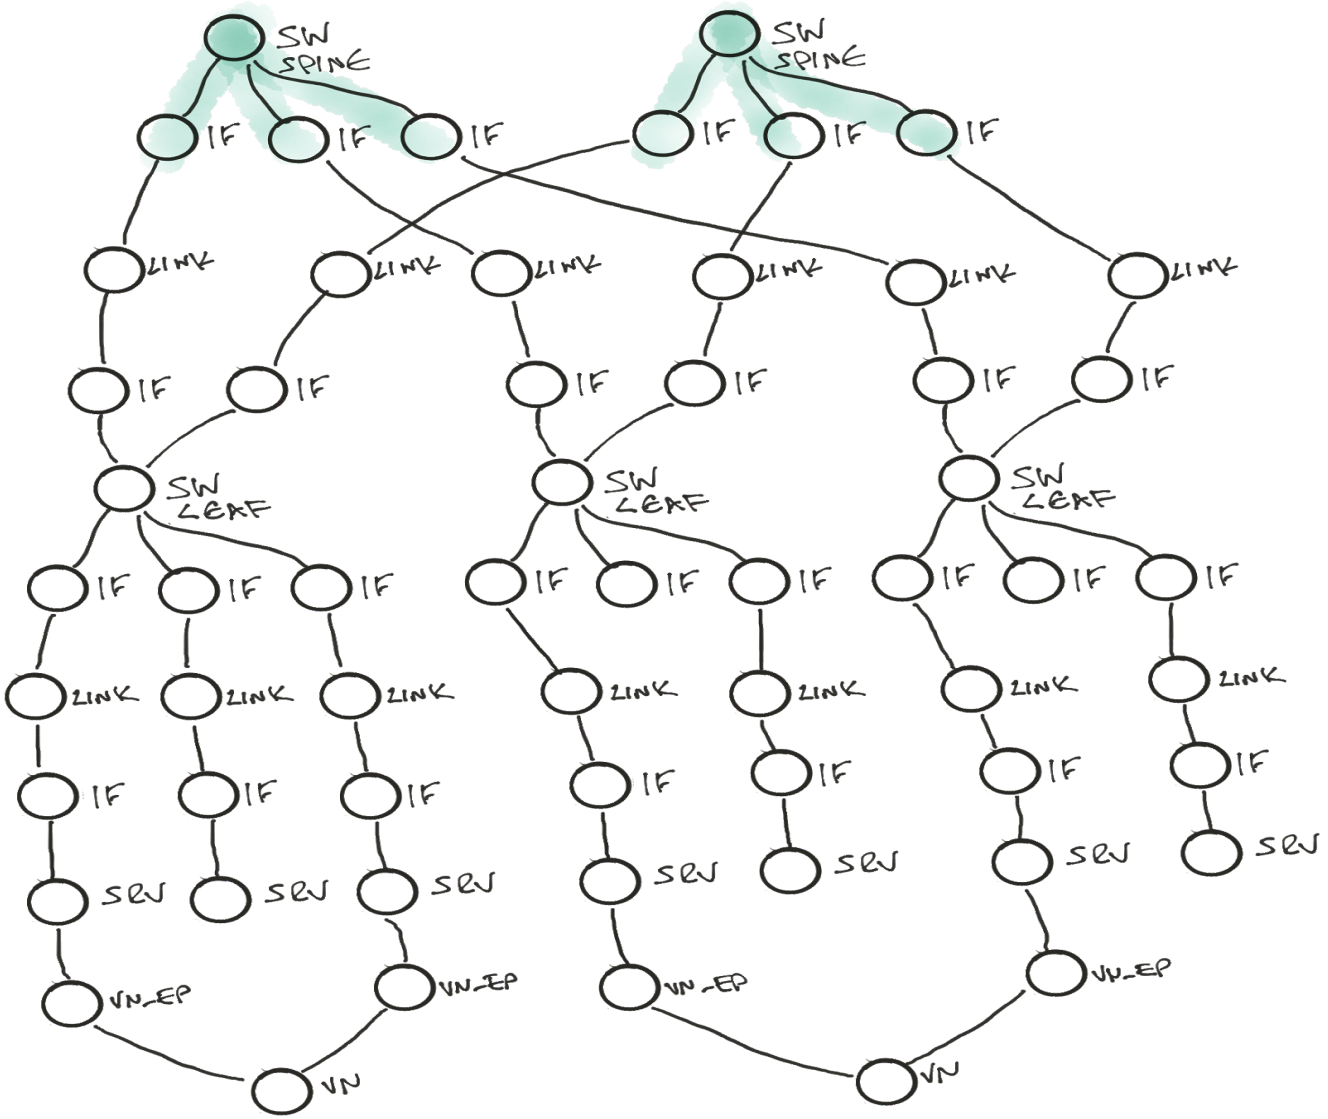

SPINE
VALIDATOR

apstra

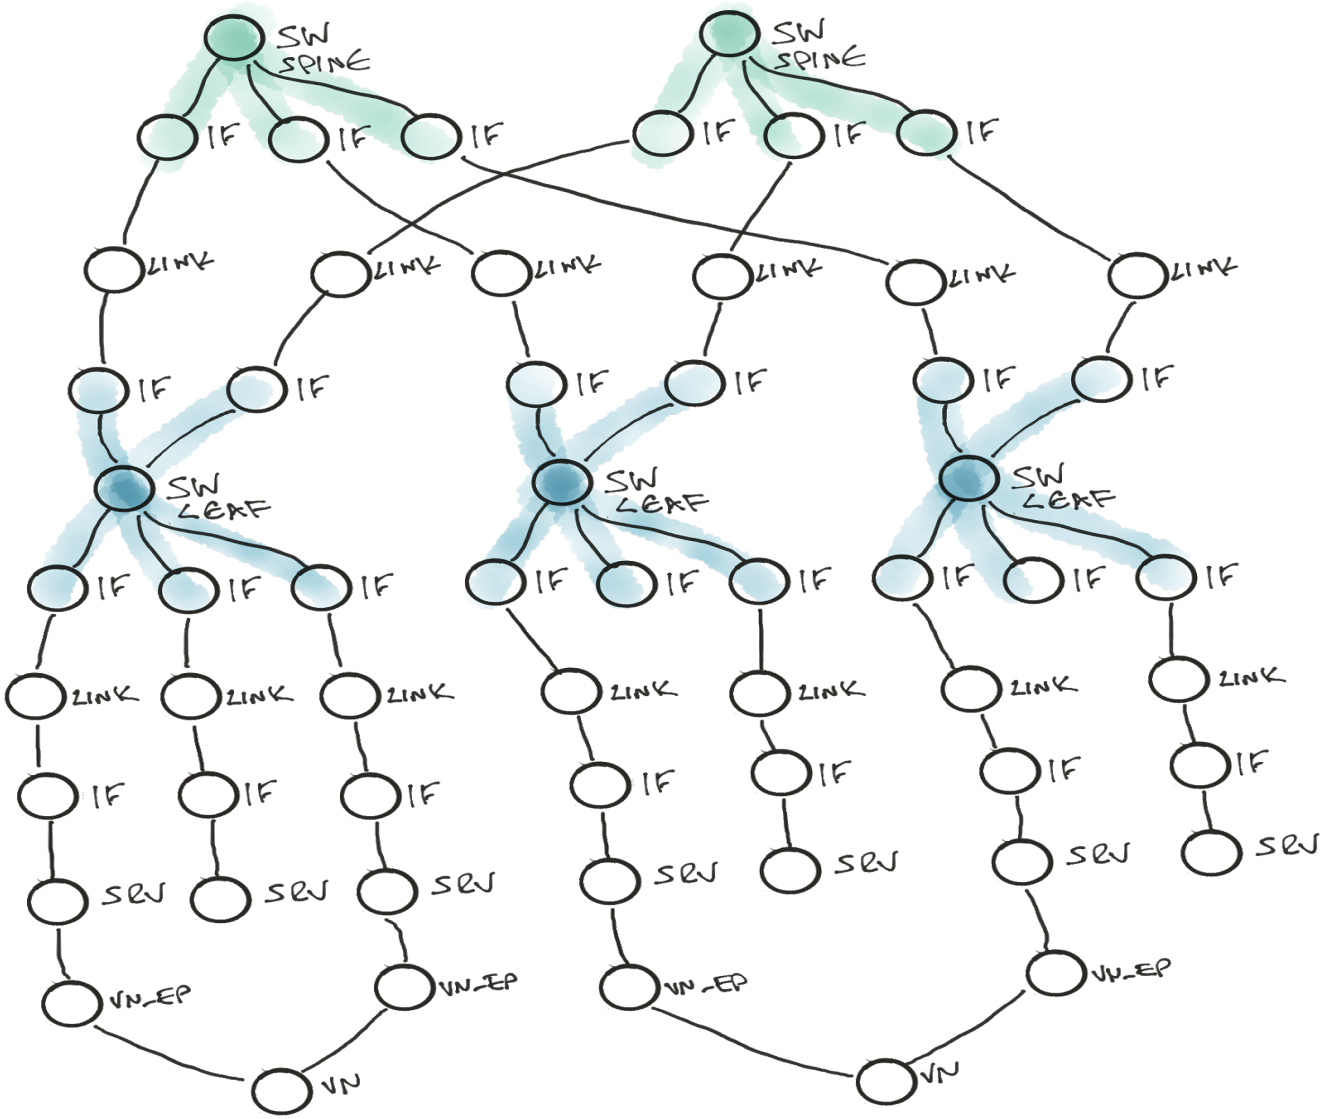

SPINE VALIDATOR

LEAF VALIDATOR

apstra

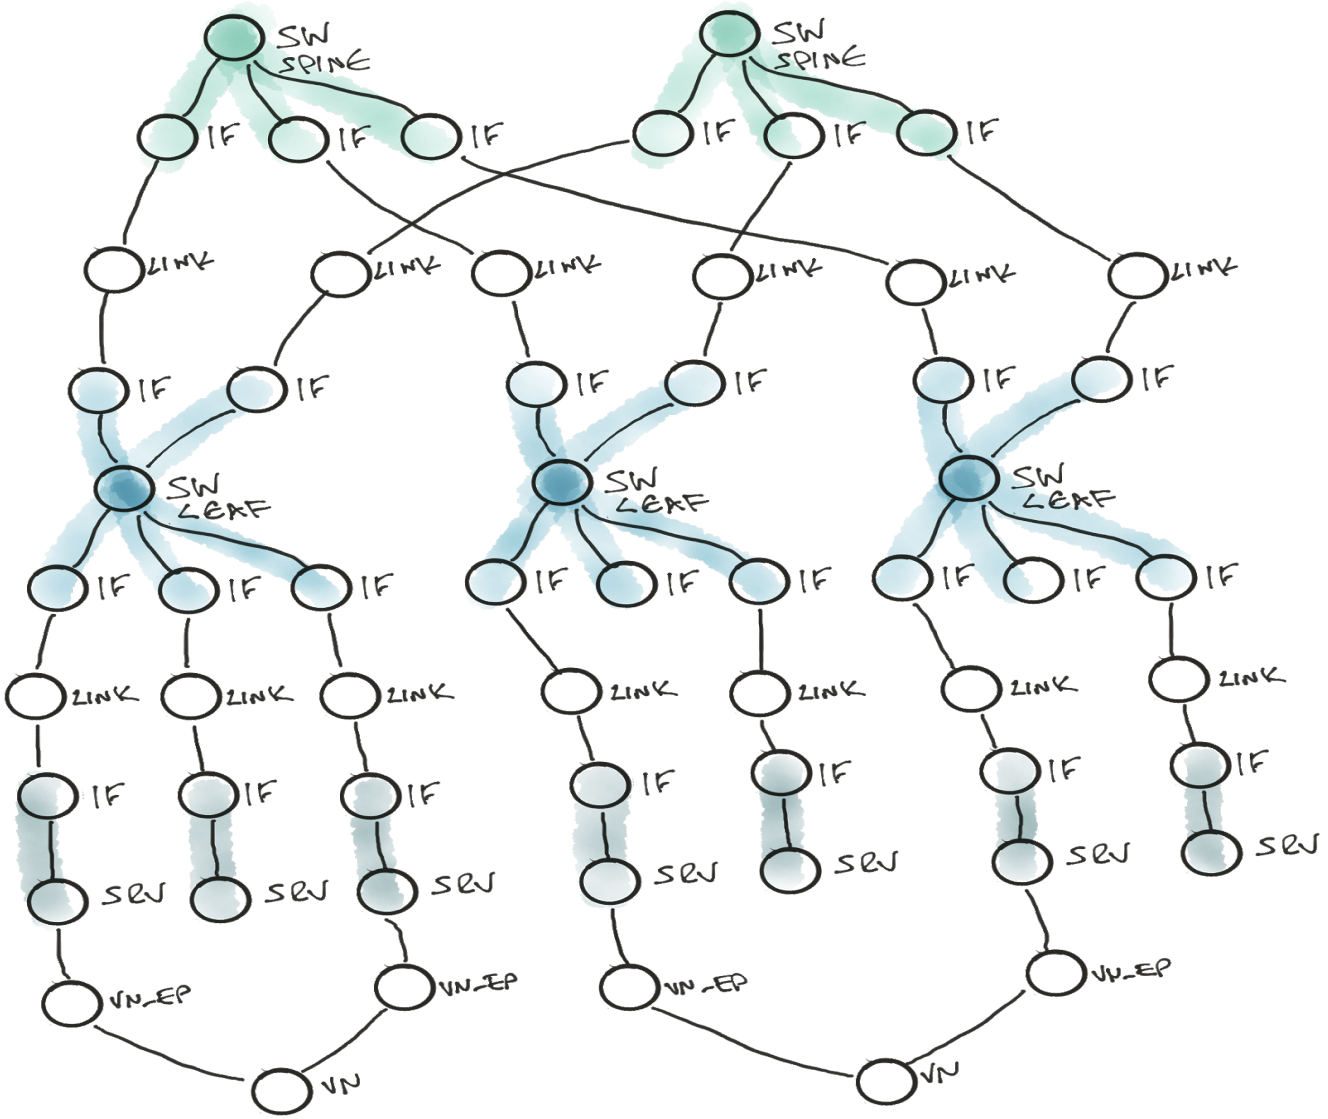

SPINE VALIDATOR

LEAF VALIDATOR

SERVER VALIDATOR

apstra

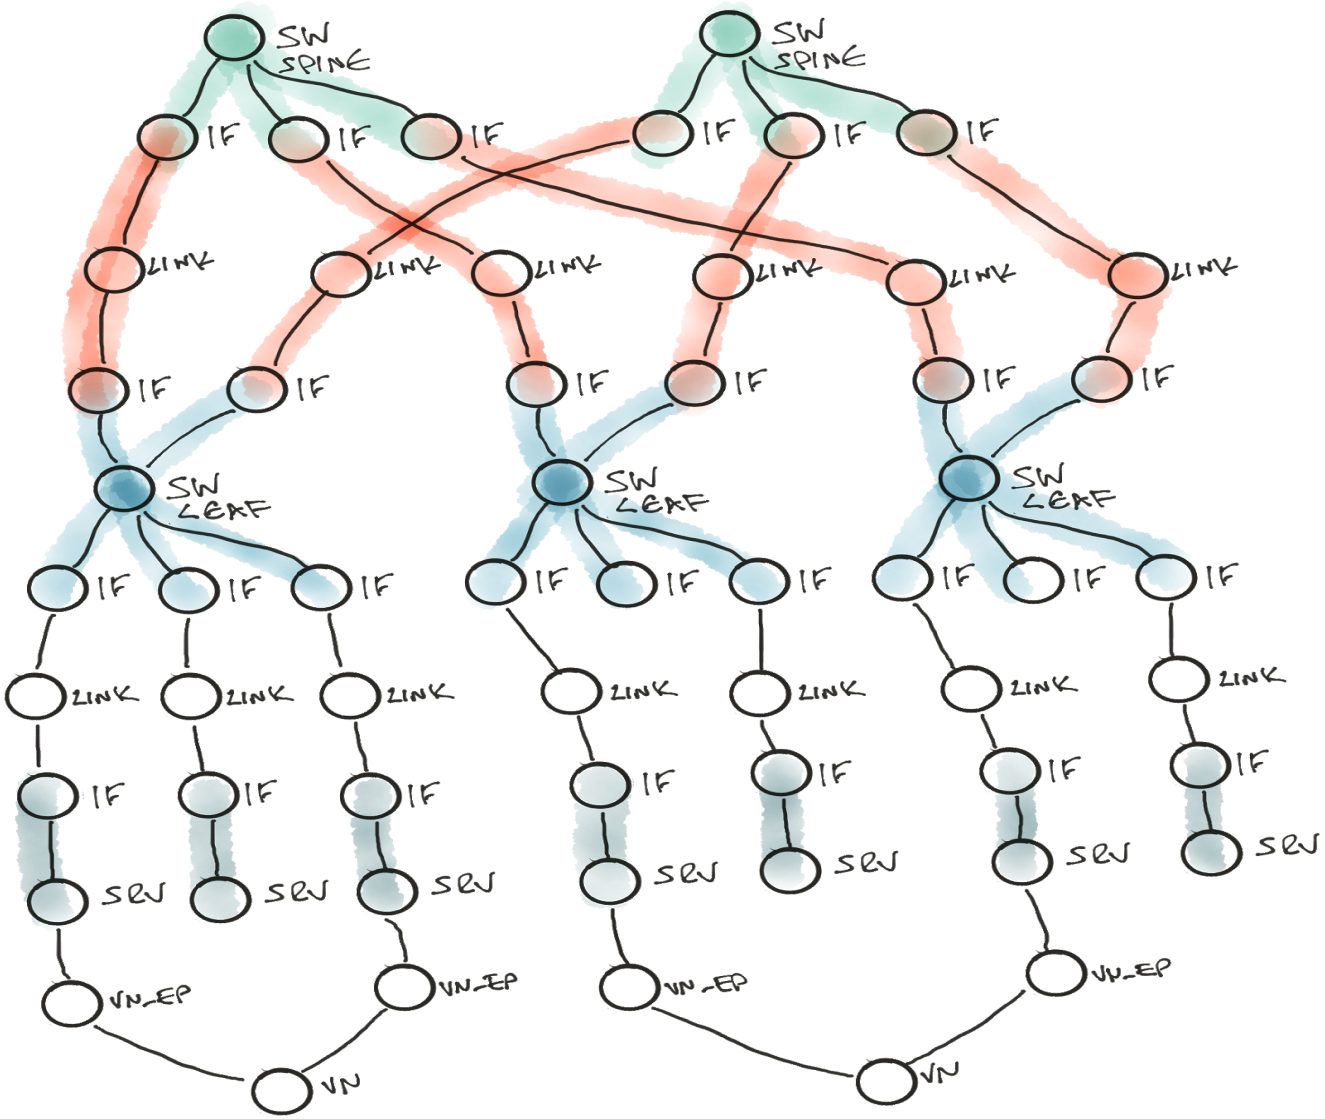

SPINE VALIDATOR

LEAF VALIDATOR

SERVER VALIDATOR

FABRIC LINK VALIDATOR

apstra

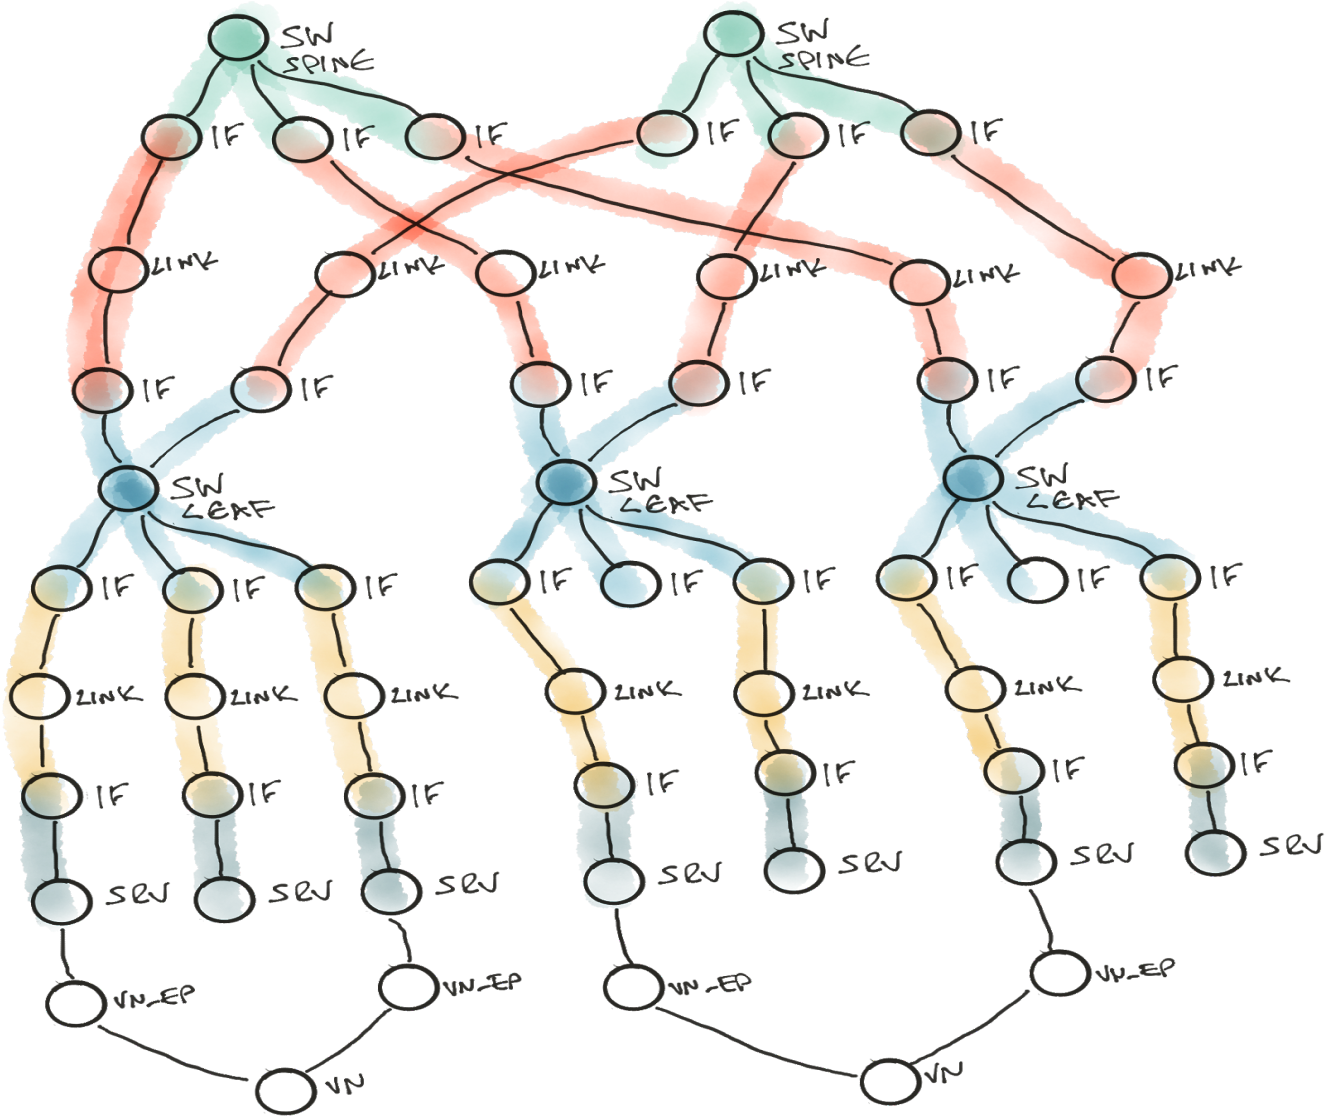

SPINE VALIDATOR

LEAF VALIDATOR

SERVER VALIDATOR

FABRIC LINK VALIDATOR

SERVER LINK VALIDATOR

apstra

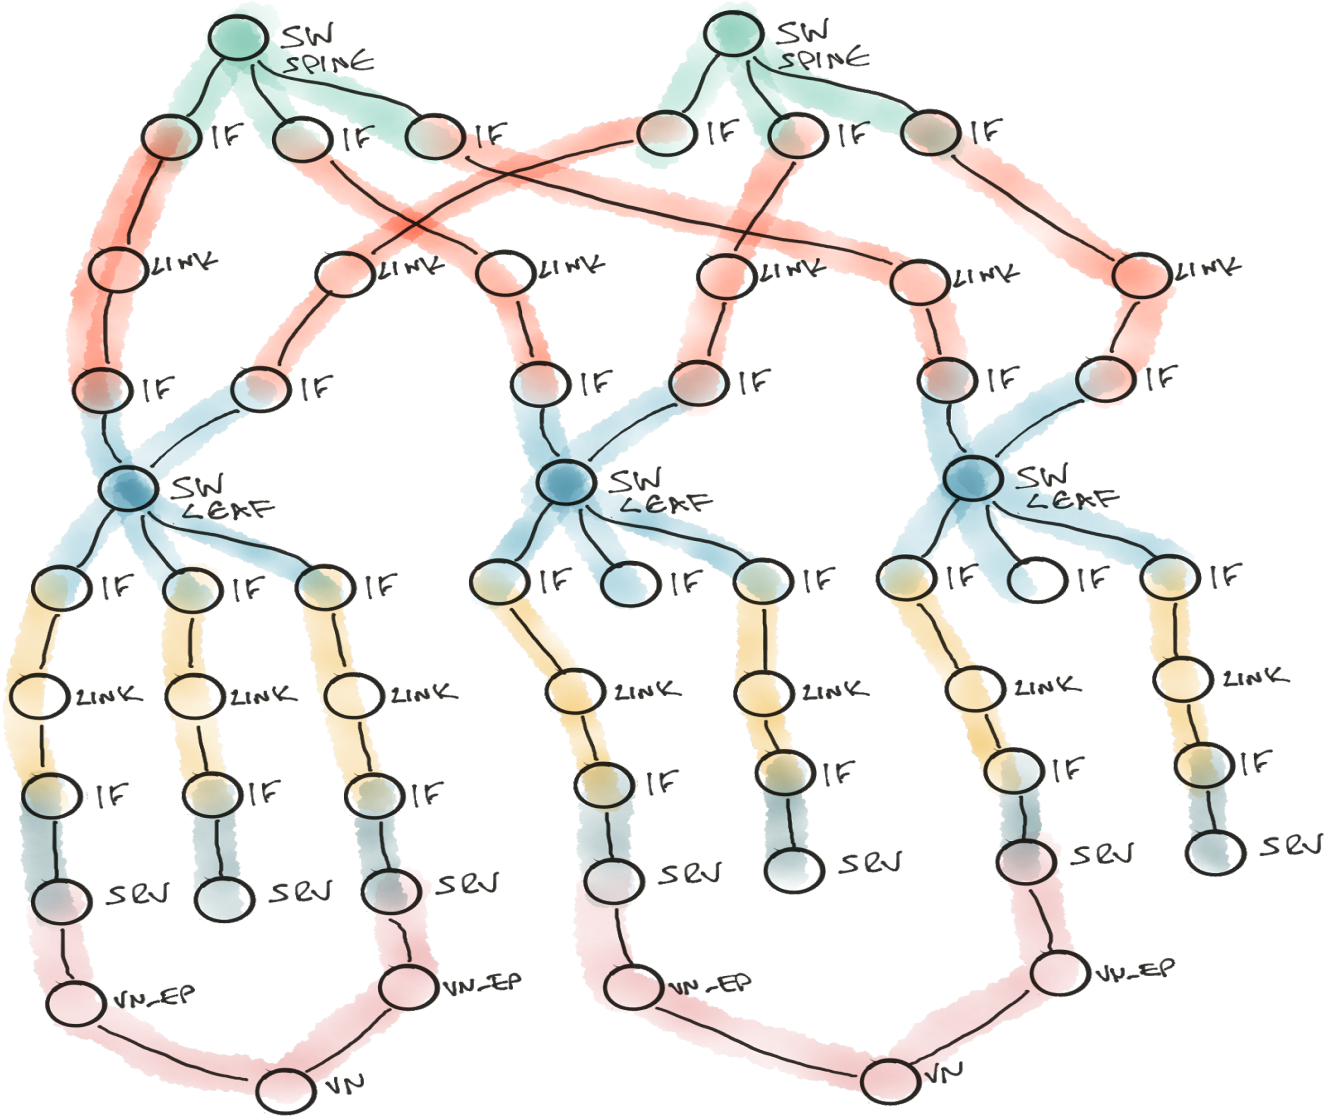

apstra

SPINE VALIDATOR

LEAF VALIDATOR

SERVER VALIDATOR

FABRIC LINK VALIDATOR

SERVER LINK VALIDATOR

VIRTUAL NETWORK VALIDATOR

Resulting Model

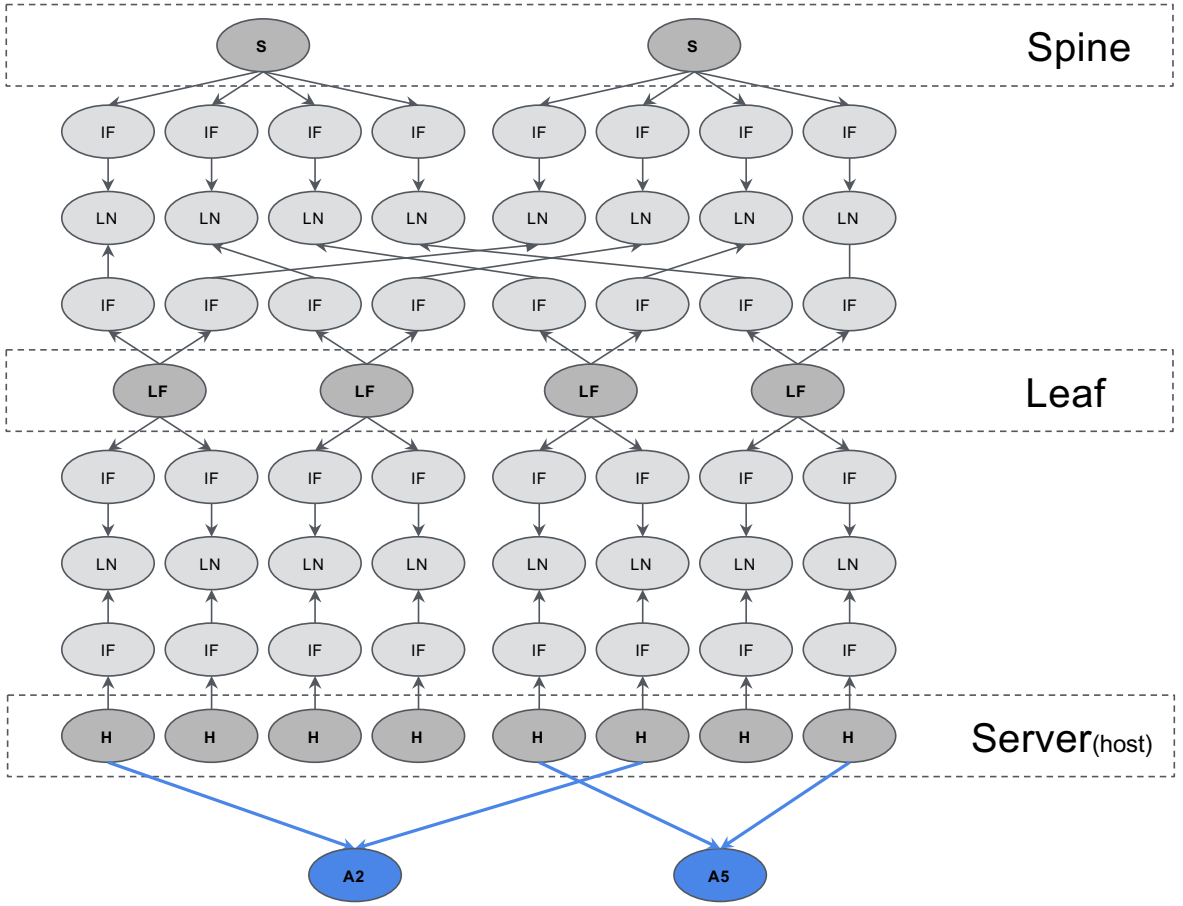

Query: Links that carry “A2” traffic

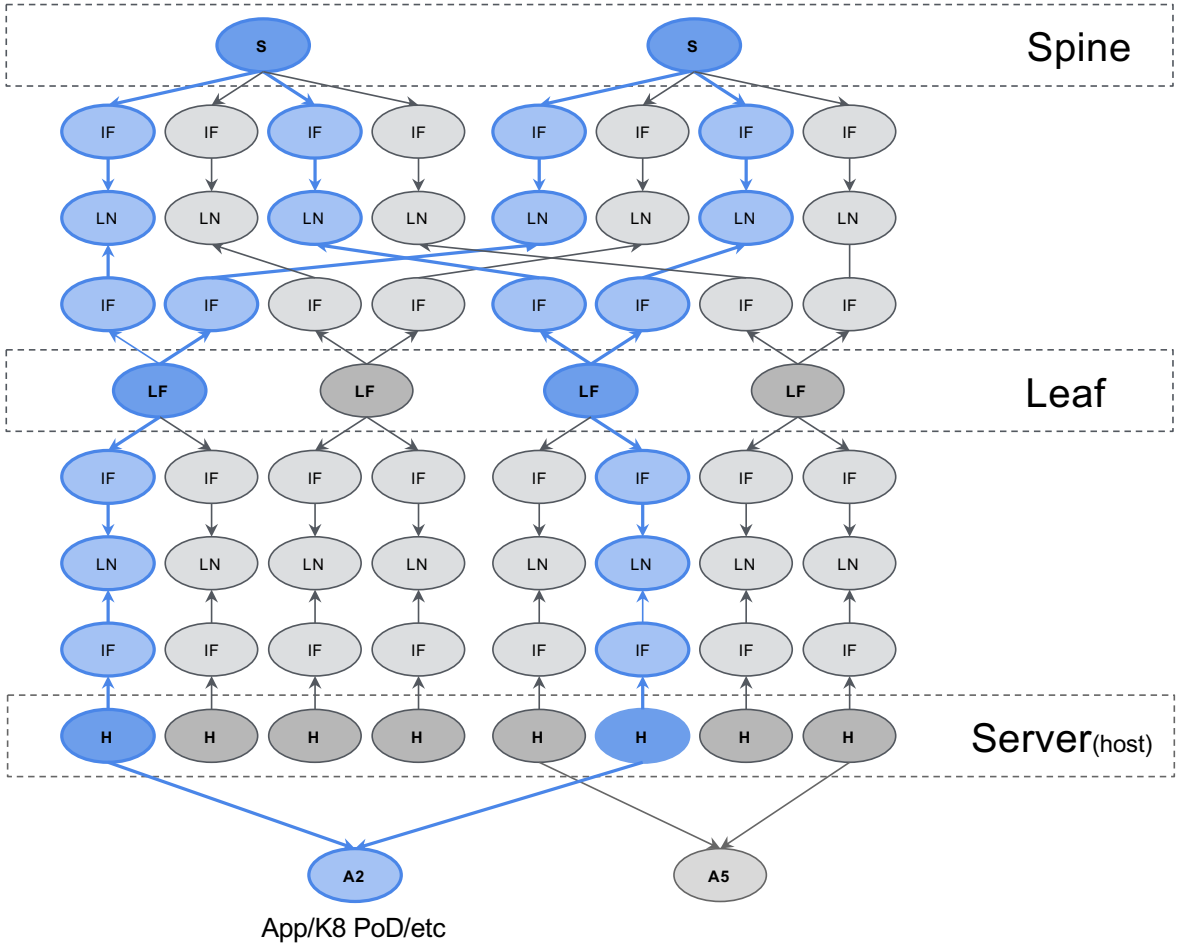

Decomposition

WIKIPEDIA
The Free Encyclopedia

Main page
Contents
Featured content
Current events
Random article
Donate to Wikipedia
Wikipedia store

Interaction

Help
About Wikipedia
Community portal
Recent changes

Decomposition (computer science)

From Wikipedia, the free encyclopedia

Decomposition in computer science, also known as **factoring**, is breaking a complex problem or system into parts that are easier to conceive, understand, program, and maintain.

Contents [hide]

DECOMPOSITION

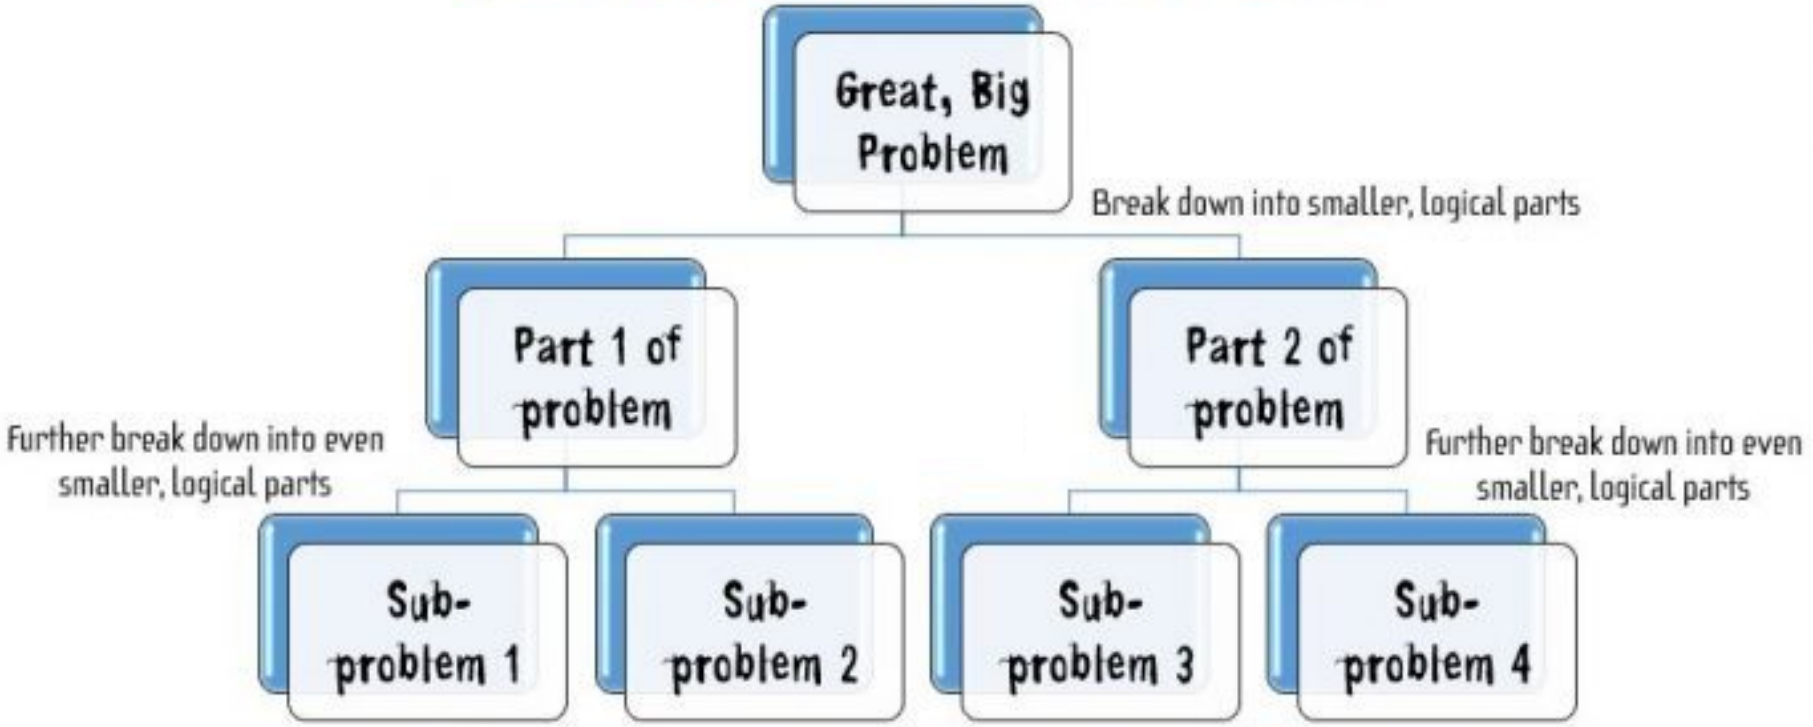

Decomposition: walking the graph

Query:

```
match(
    node("system", role="spine")
    .out()
    .node("interface")
    .out()
    .node("link")
    .in_()
    .node("interface")
    .in_()
    .node("system", role="leaf")
)
```

Execute Query

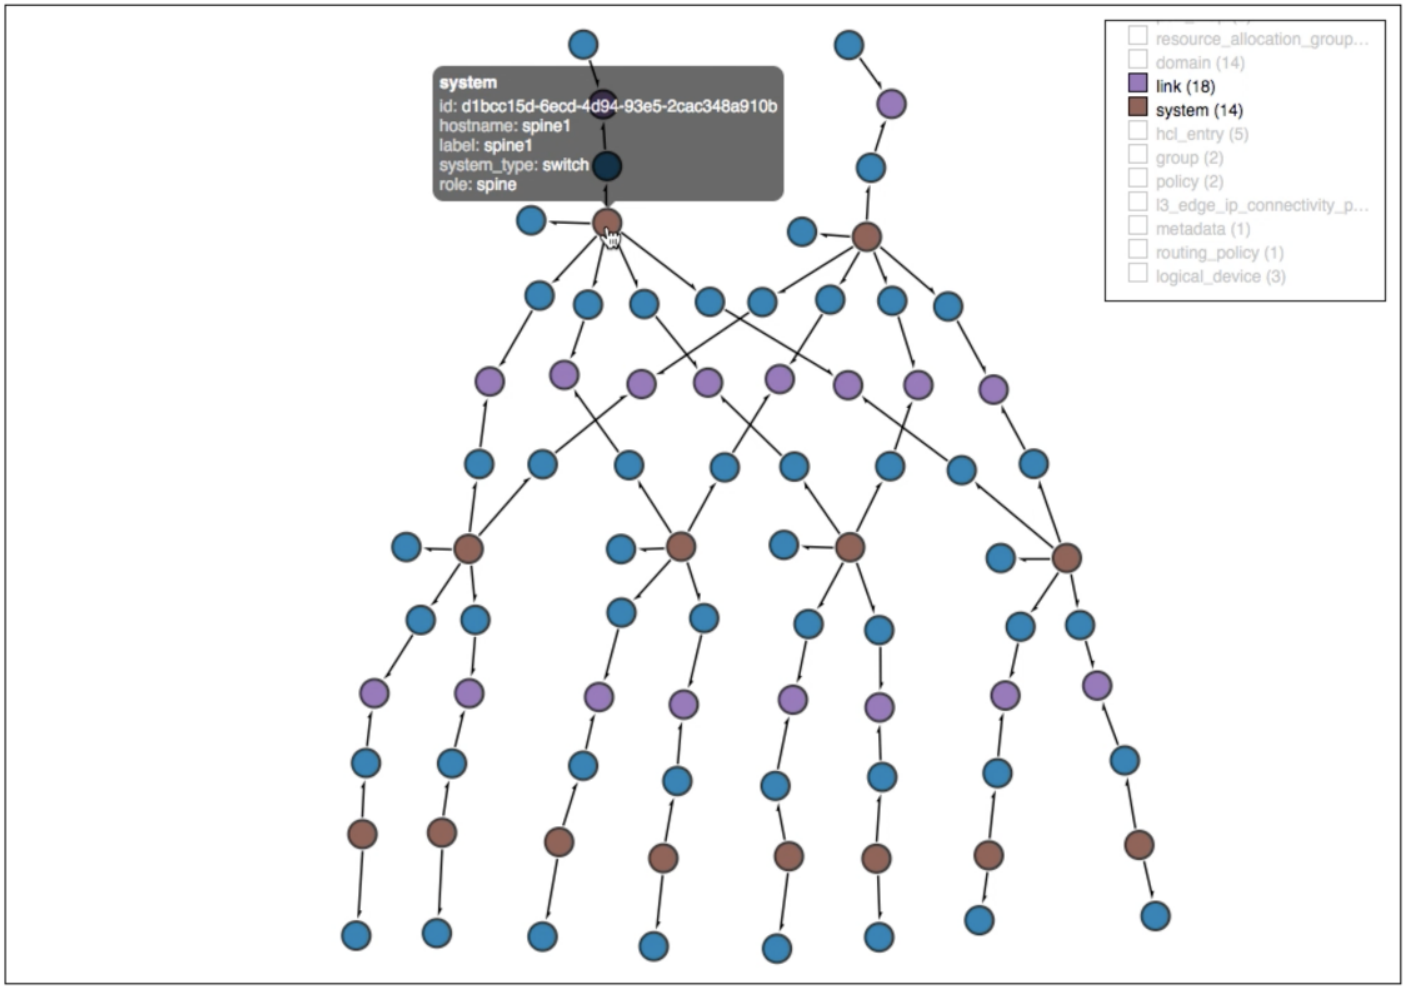

apstra

Query:

```
match(
    node("system", role="spine")
    .out()
    .node("interface")
    .out()
    .node("link")
    .in_()
    .node("interface")
    .in_()
    .node("system", role="leaf")
)
```

Execute Query

Close

Query

match(node("system", role="spine") .out()
.node("interface") .out() .node("link") .in_()
.node("interface") .in_() .node("system", role="leaf"))

Steps

start 0 1 2 3 4 5 6 7 8

<Start>

Paths (0)

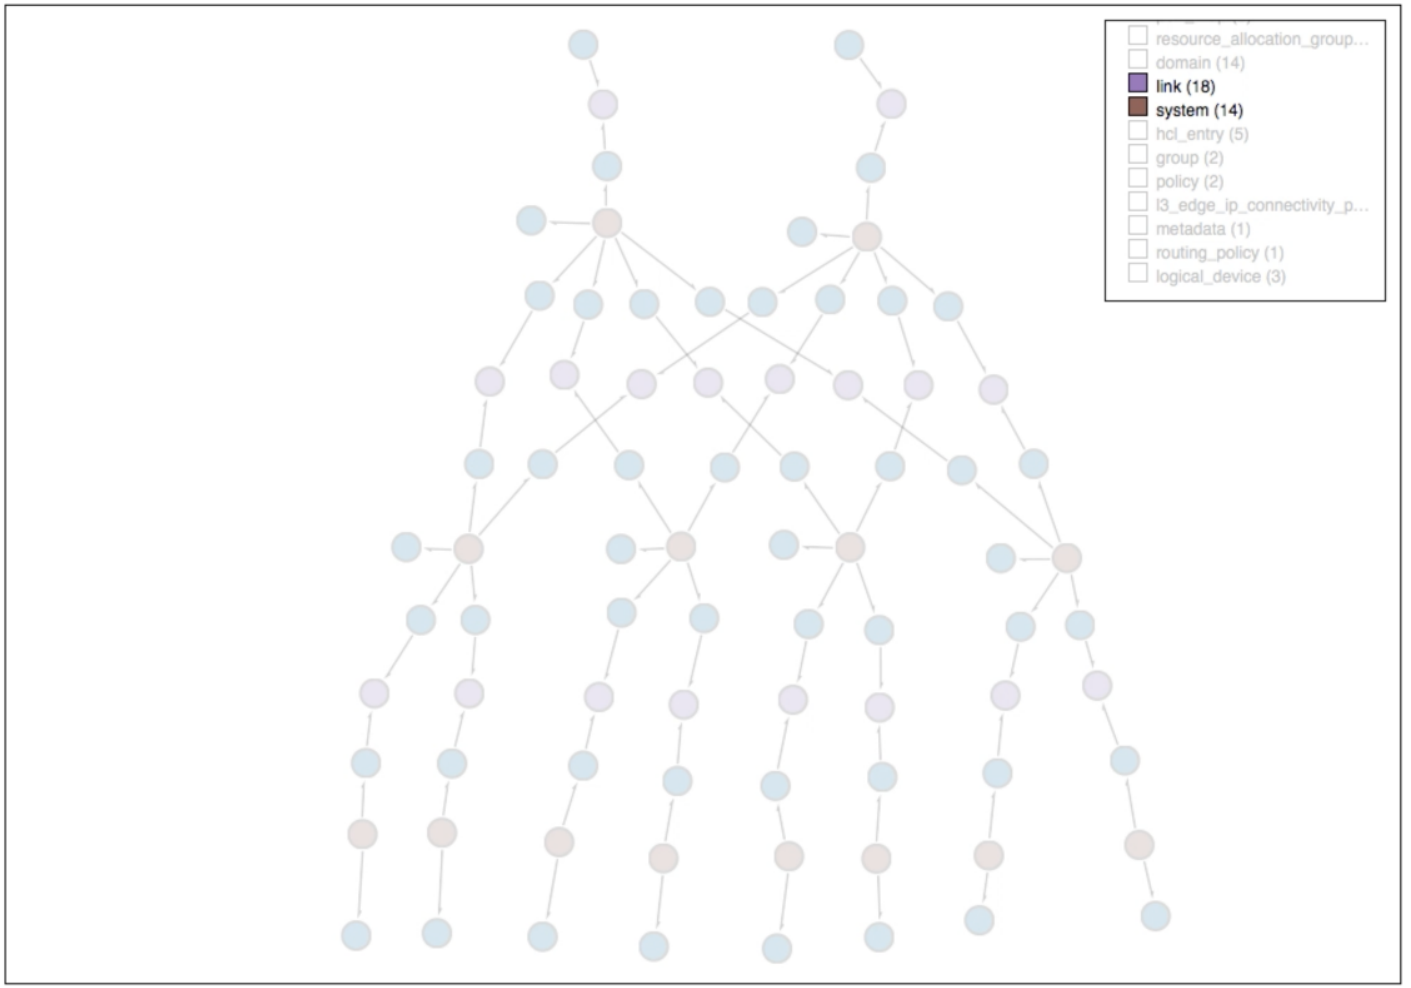

apstra

Query:

```
match(
    node("system", role="spine")
    .out()
    .node("interface")
    .out()
    .node("link")
    .in_()
    .node("interface")
    .in_()
    .node("system", role="leaf")
)
```

Execute Query

Close

Query

match(node("system", role="spine") .out()
.node("interface") .out() .node("link") .in_()
.node("interface") .in_() .node("system", role="leaf"))

Steps

start 0 1 2 3 4 5 6 7 8

<FindNodeAction type=system role=== spine>

Paths (2)

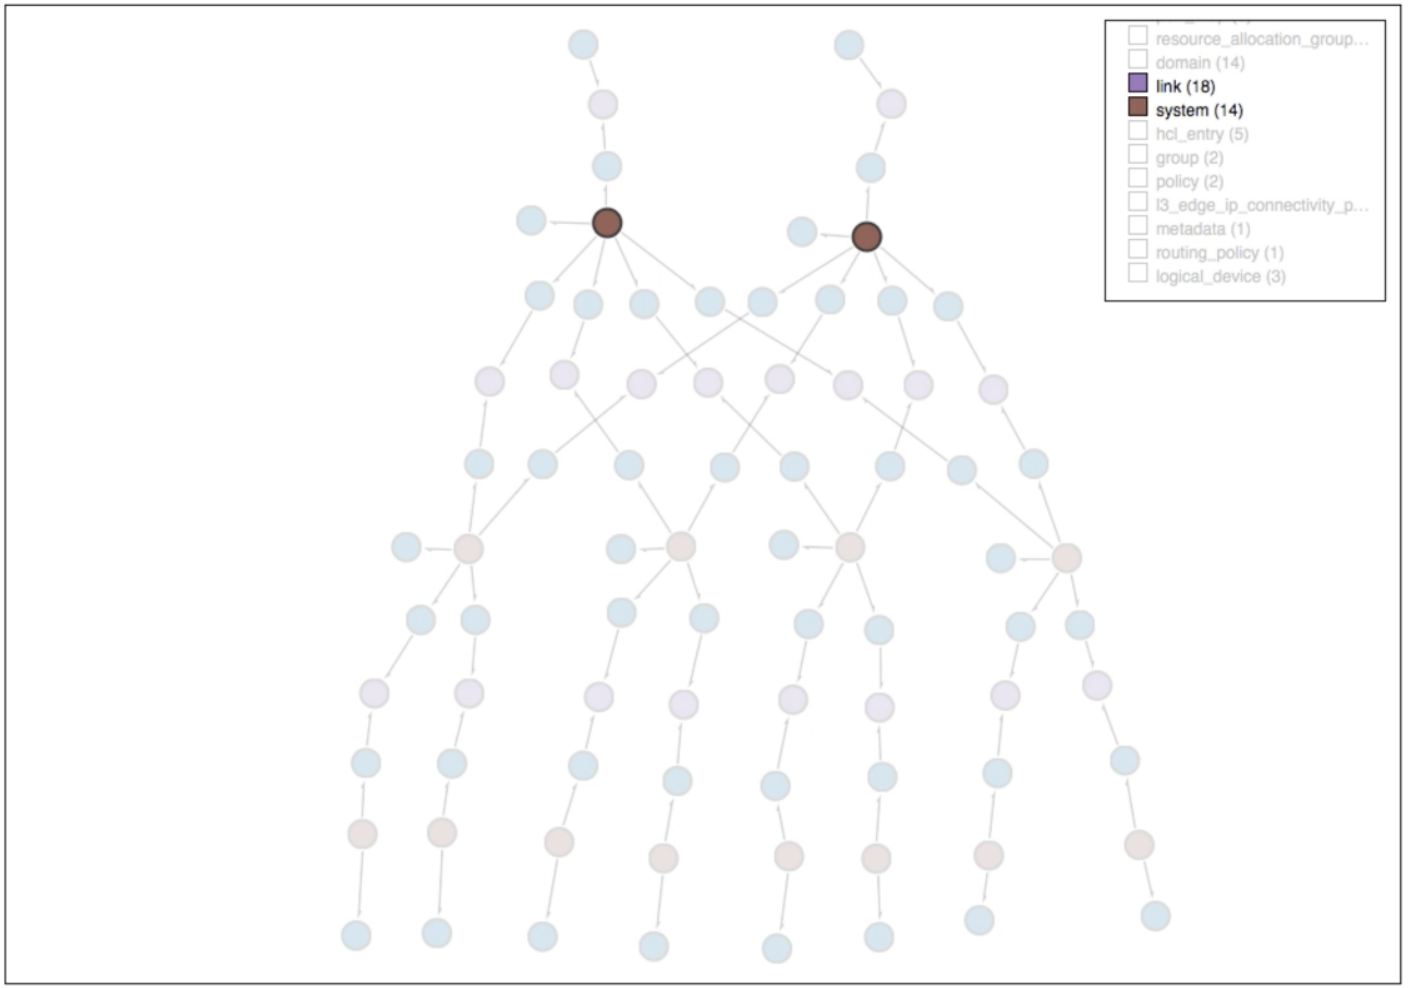

Query:

```
match(
    node("system", role="spine")
    .out()
    .node("interface")
    .out()
    .node("link")
    .in_()
    .node("interface")
    .in_()
    .node("system", role="leaf")
)
```

Execute Query

Close

Query

match(node("system", role="spine") .out()
.node("interface") .out() .node("link") .in_()
.node("interface") .in_() .node("system", role="leaf"))

Steps

start 0 1 2 3 4 5 6 7 8

<NodeOutRelationshipAction index=0>

Paths (14)

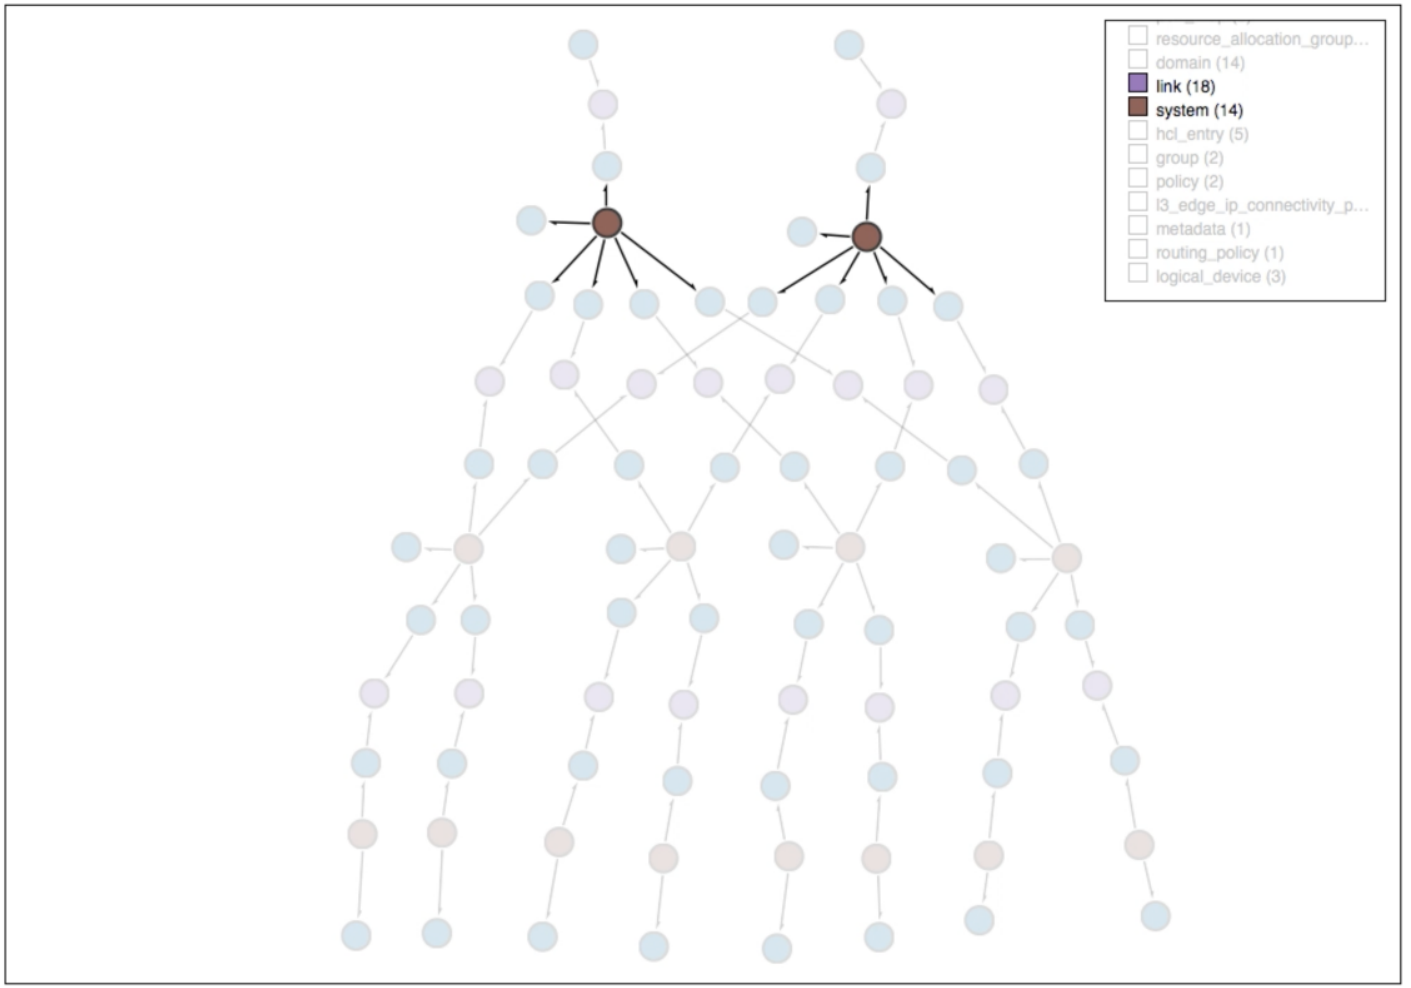

Query:

```
match(
    node("system", role="spine")
    .out()
    .node("interface")
    .out()
    .node("link")
    .in_()
    .node("interface")
    .in_()
    .node("system", role="leaf")
)
```

Execute Query

Close

Query

match(node("system", role="spine") .out()
.node("interface") .out() .node("link") .in_()
.node("interface") .in_() .node("system", role="leaf"))

Steps

start 0 1 2 3 4 5 6 7 8

<RelationshipTargetAction index=1 type=interface>

Paths (12)

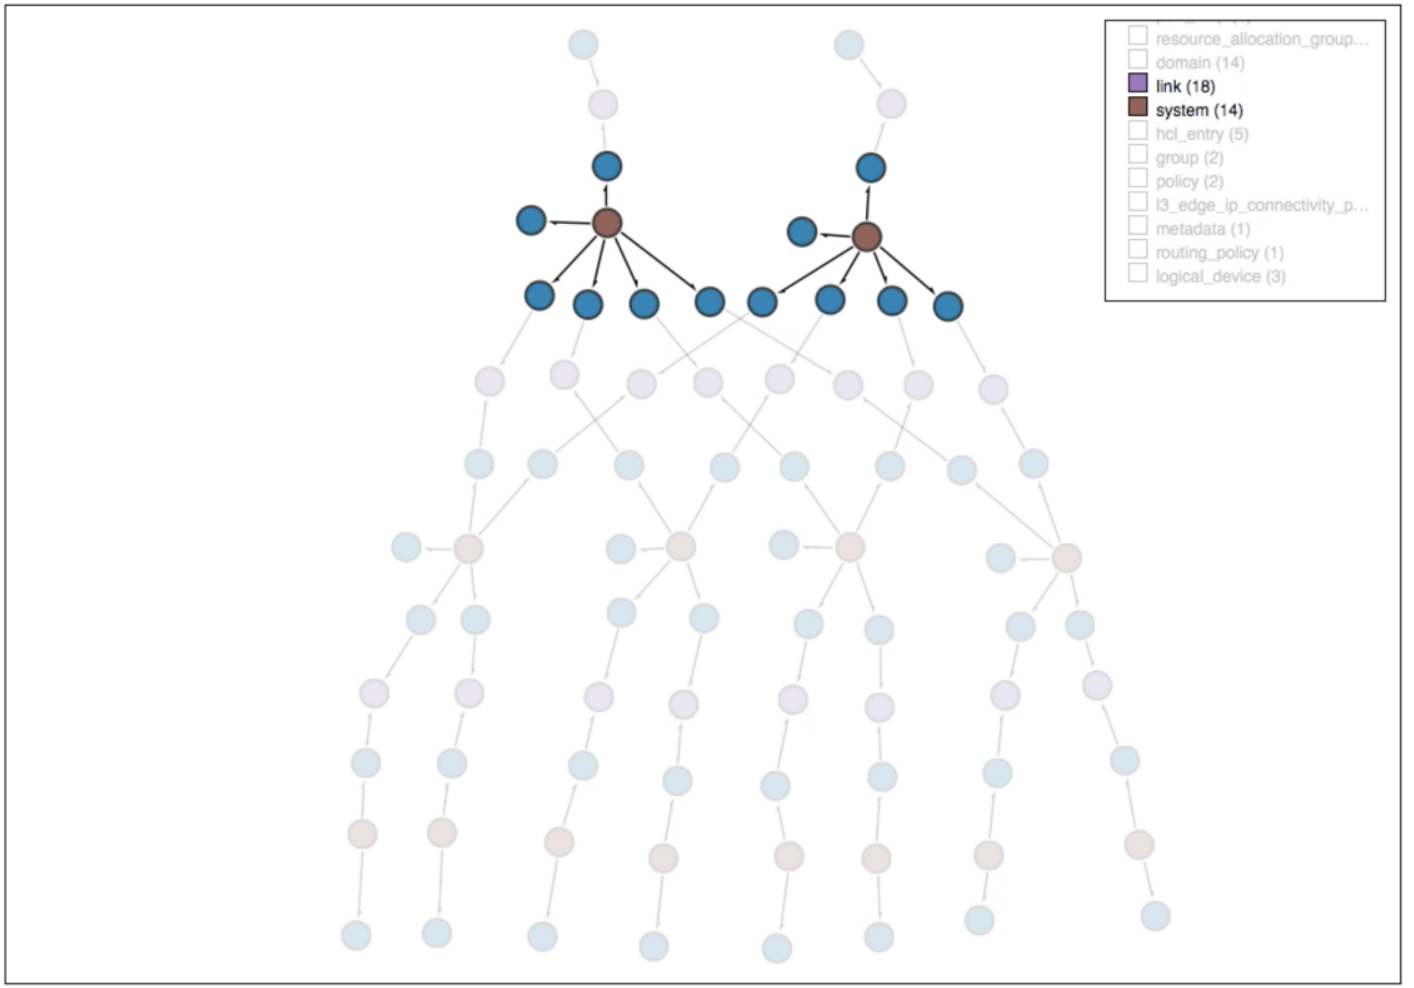

apstra

Query:

```
match(
    node("system", role="spine")
    .out()
    .node("interface")
    .out()
    .node("link")
    .in_()
    .node("interface")
    .in_()
    .node("system", role="leaf")
)
```

Execute Query

Close

Query

match(node("system", role="spine") .out()
.node("interface") .out() .node("link") .in_()
.node("interface") .in_() .node("system", role="leaf"))

Steps

start 0 1 2 3 4 5 6 7 8

<NodeOutRelationshipAction index=2>

Paths (10)

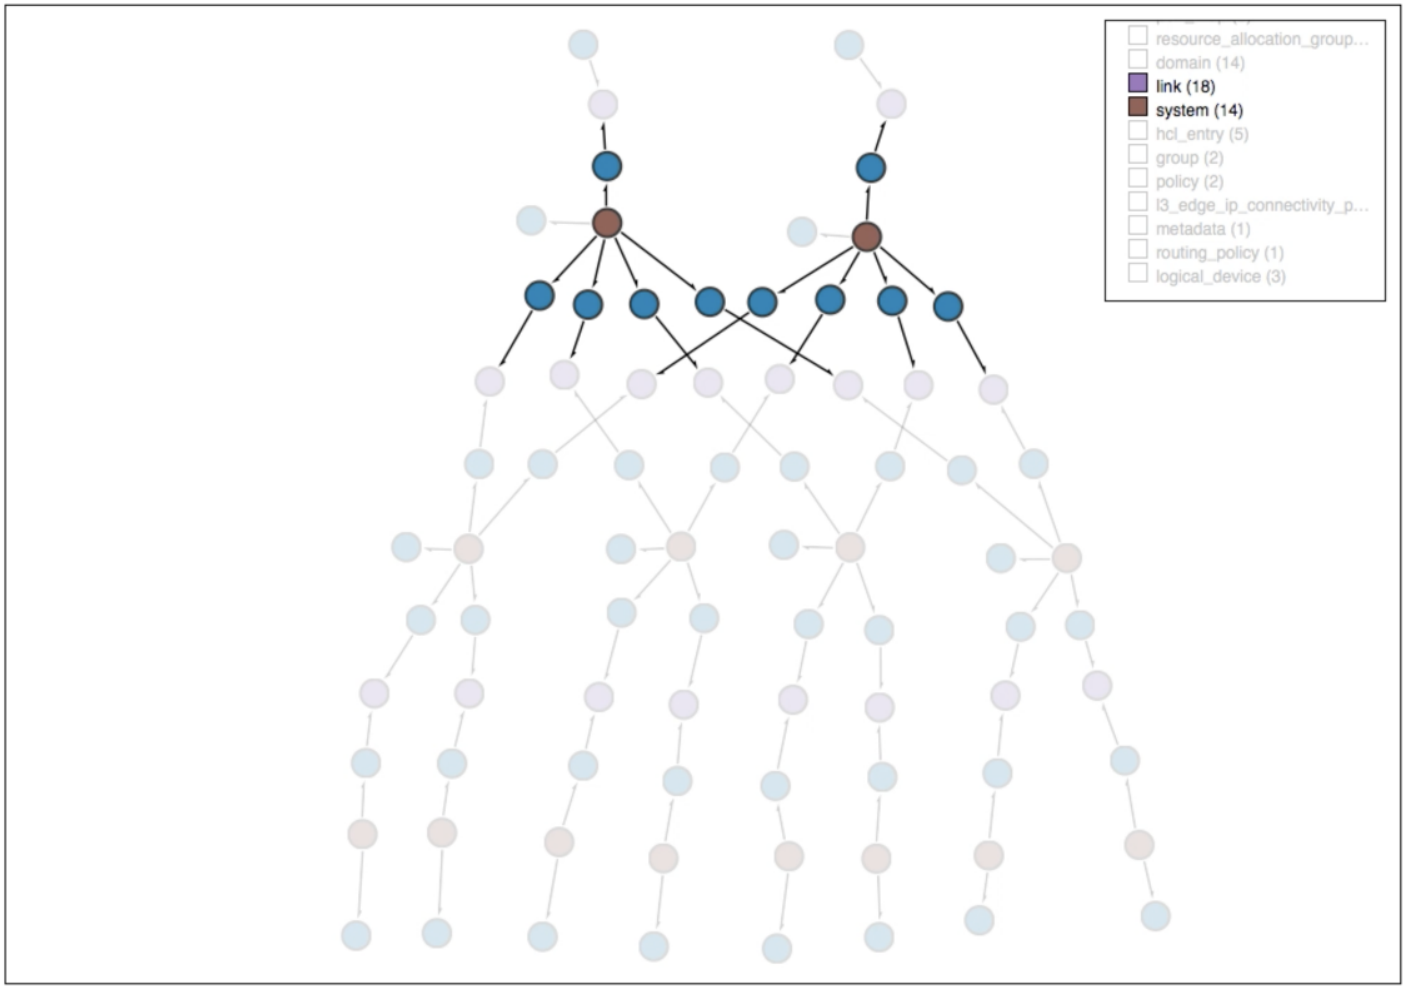

Query:

```
match(
    node("system", role="spine")
    .out()
    .node("interface")
    .out()
    .node("link")
    .in_()
    .node("interface")
    .in_()
    .node("system", role="leaf")
)
```

Execute Query

Close

Query

match(node("system", role="spine") .out()
.node("interface") .out() .node("link") .in_()
.node("interface") .in_() .node("system", role="leaf"))

Steps

start 0 1 2 3 4 5 6 7 8

<RelationshipTargetAction index=3 type=link>

Paths (10)

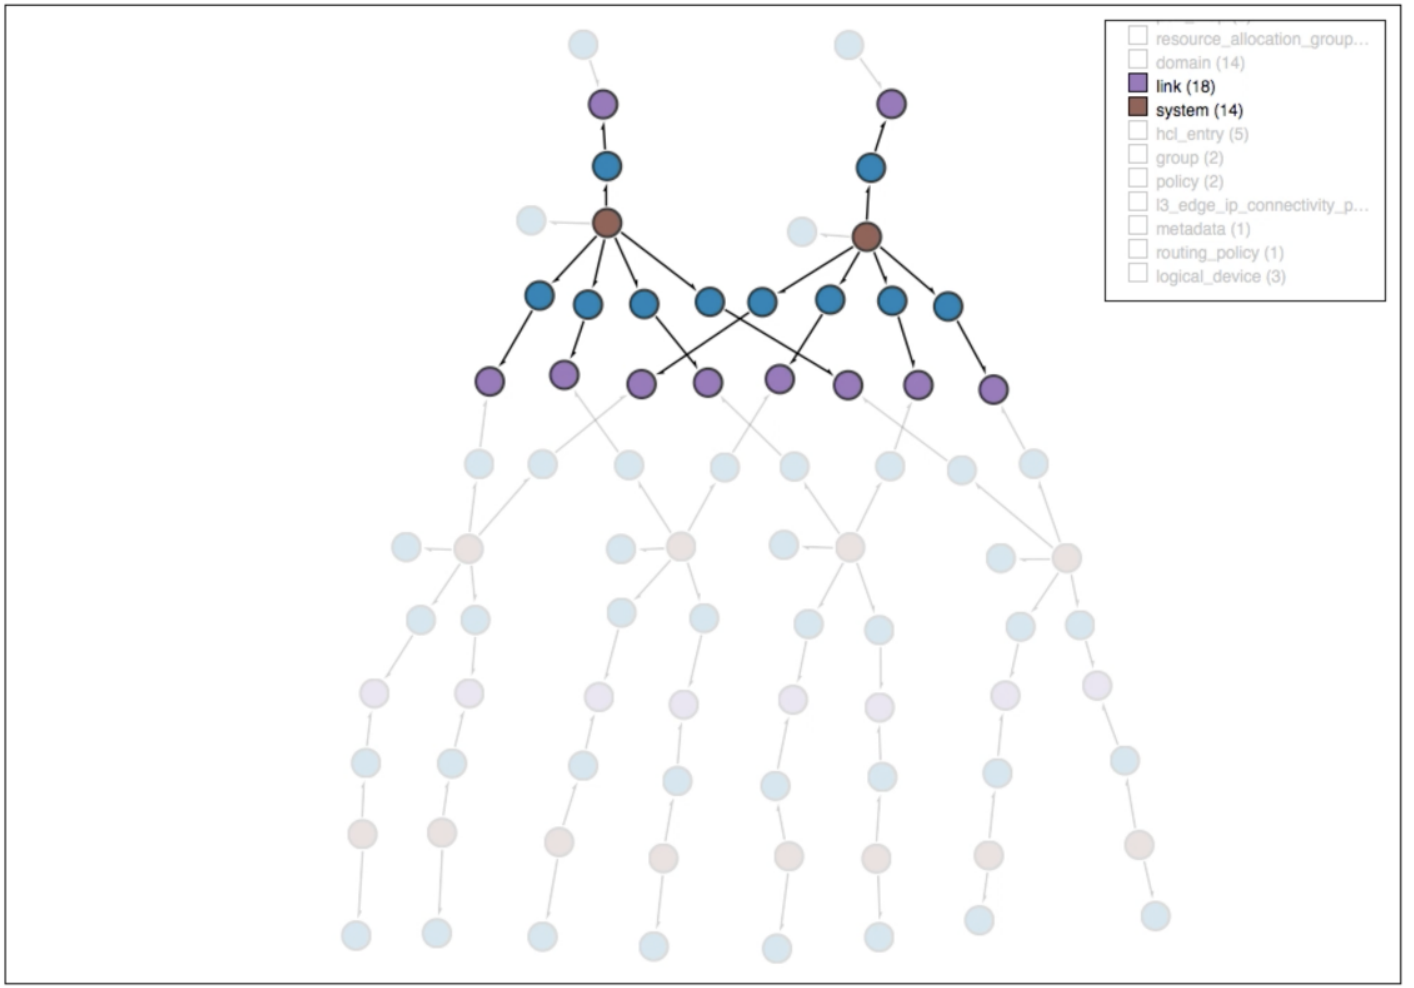

Query:

```
match(
    node("system", role="spine")
    .out()
    .node("interface")
    .out()
    .node("link")
    .in_()
    .node("interface")
    .in_()
    .node("system", role="leaf")
)
```

Execute Query

Close

Query

match(node("system", role="spine") .out()
.node("interface") .out() .node("link") .in_()
.node("interface") .in_() .node("system", role="leaf"))

Steps

start 0 1 2 3 4 5 6 7 8

<NodeInRelationshipAction index=4>

Paths (20)

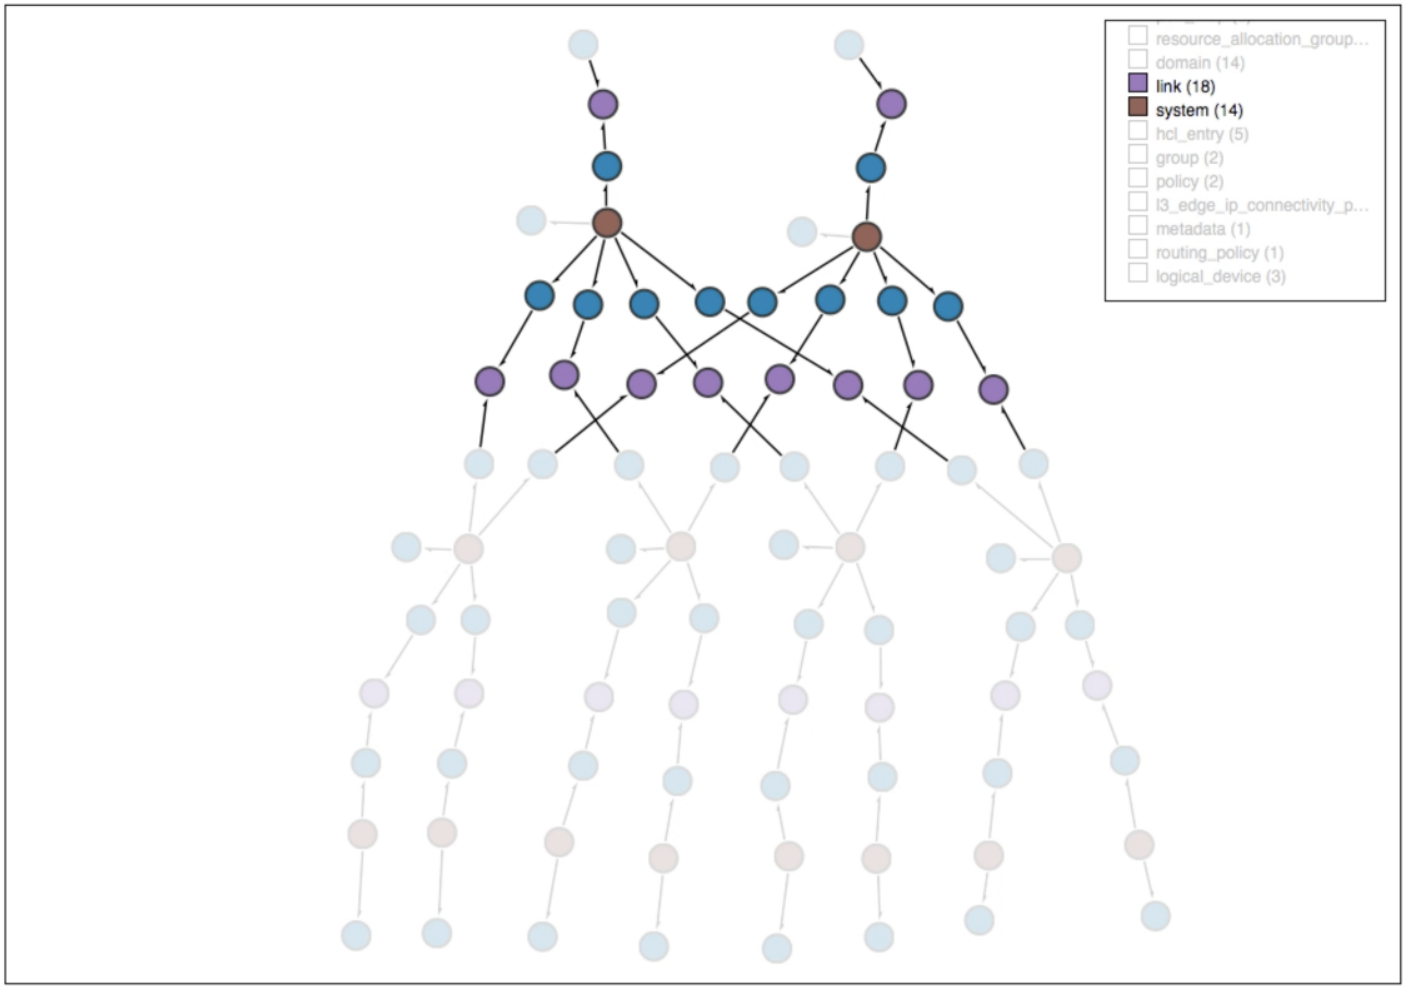

Query:

```
match(
    node("system", role="spine")
    .out()
    .node("interface")
    .out()
    .node("link")
    .in_()
    .node("interface")
    .in_()
    .node("system", role="leaf")
)
```

Execute Query

Close

Query

match(node("system", role="spine") .out()
.node("interface") .out() .node("link") .in_()
.node("interface") .in_() .node("system", role="leaf"))

Steps

start 0 1 2 3 4 5 6 7 8

<RelationshipSourceAction index=5 type=interface>

Paths (20)

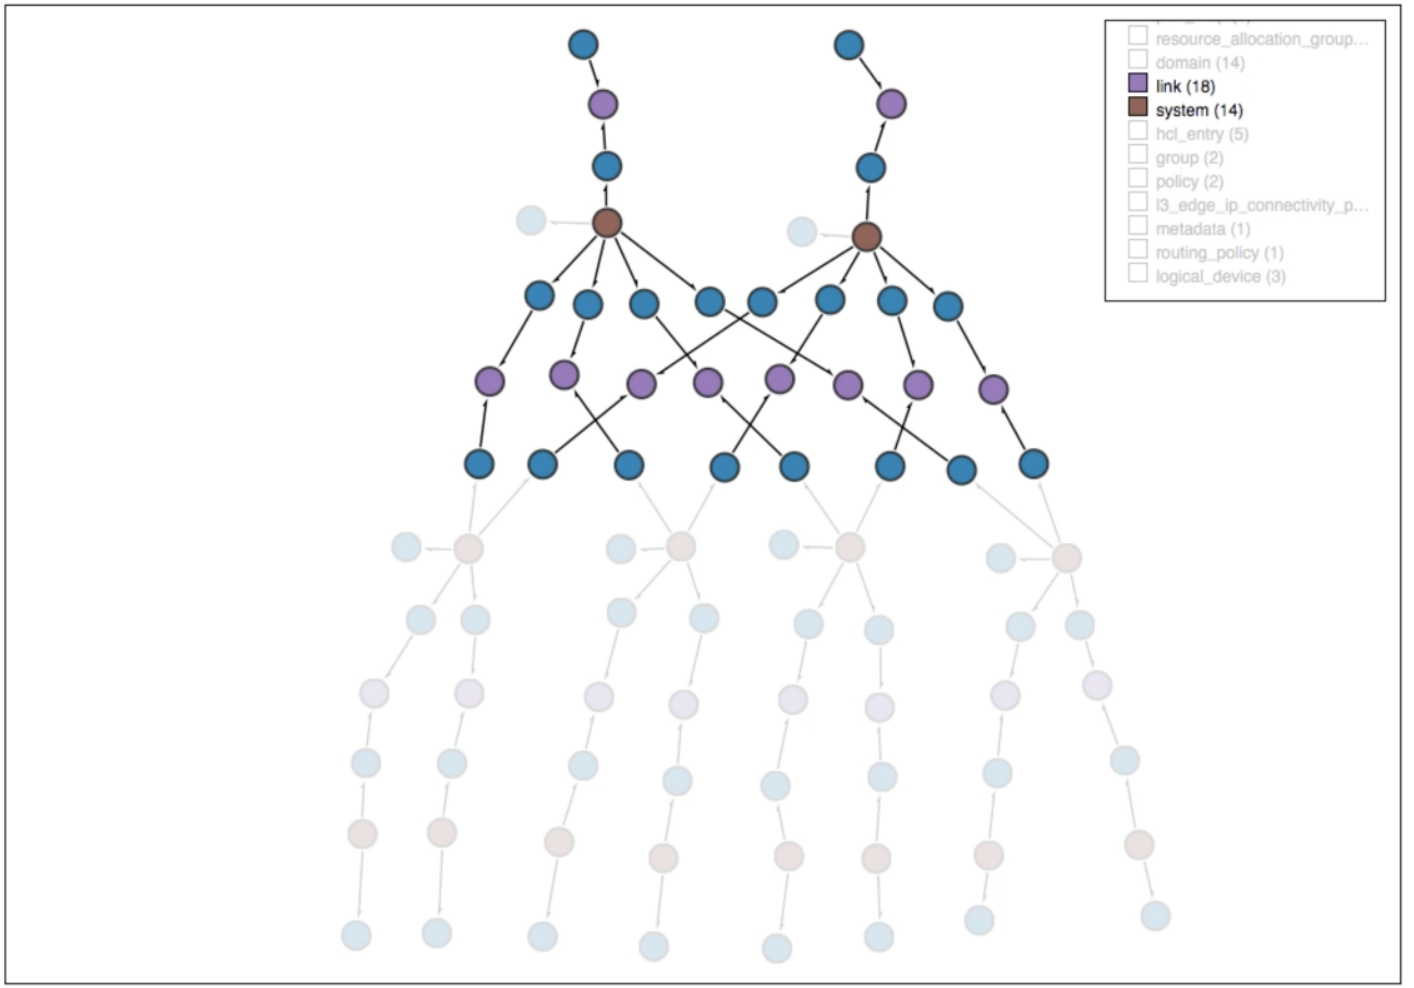

apstra

Query:

```
match(
    node("system", role="spine")
    .out()
    .node("interface")
    .out()
    .node("link")
    .in_()
    .node("interface")
    .in_()
    .node("system", role="leaf")
)
```

Execute Query

Close

Query

match(node("system", role="spine") .out()
.node("interface") .out() .node("link") .in_()
.node("interface") .in_() .node("system", role="leaf"))

Steps

start 0 1 2 3 4 5 6 7 8

<NodeInRelationshipAction index=6>

Paths (18)

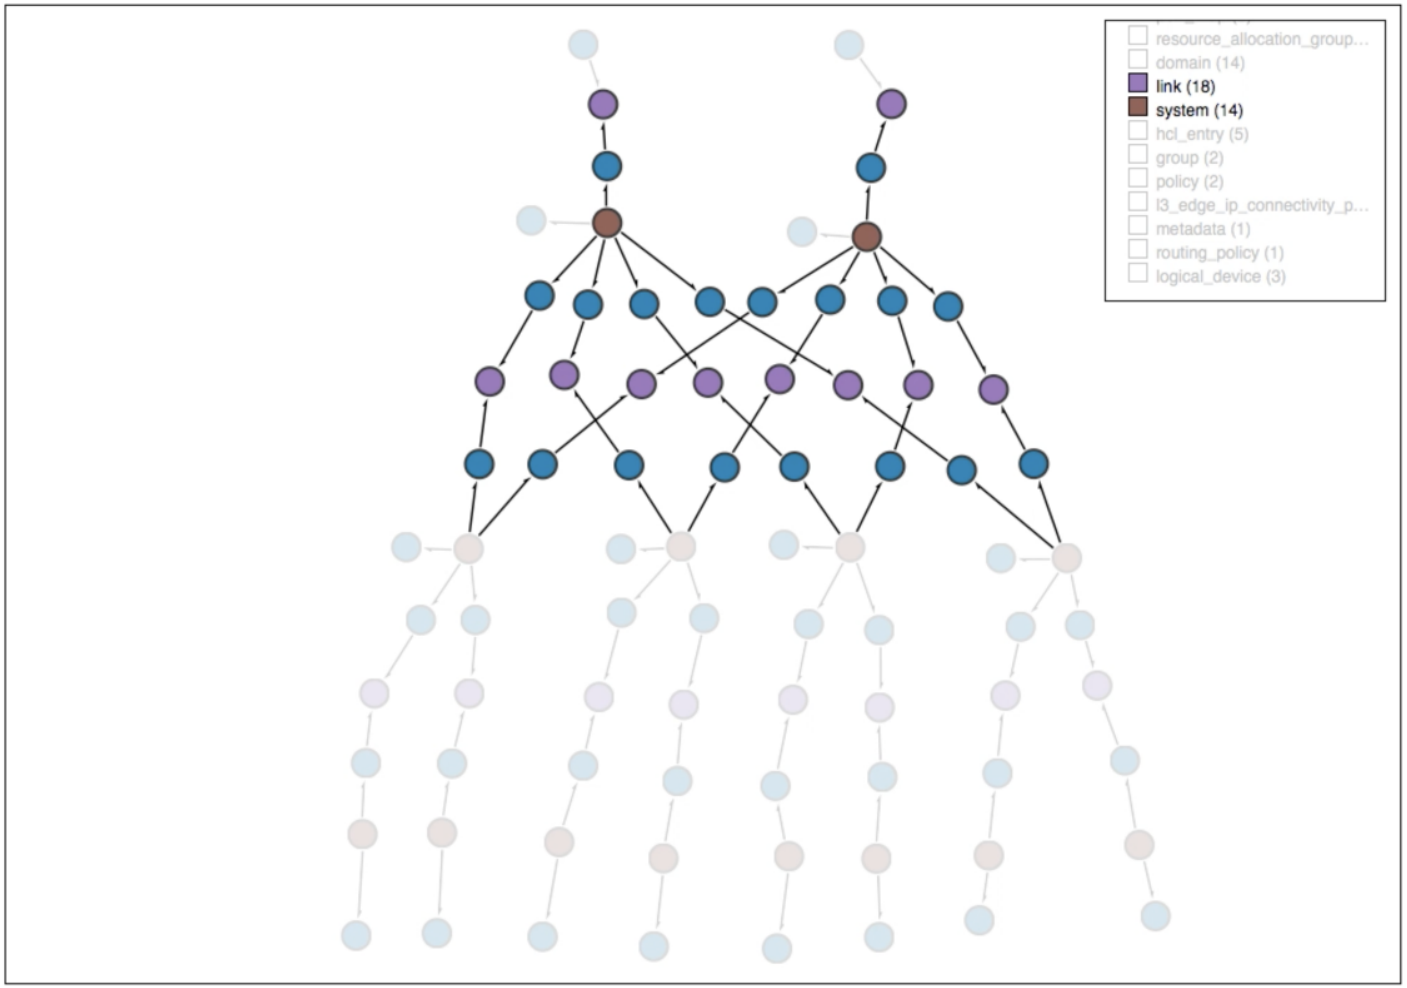

Query:

```
match(
    node("system", role="spine")
    .out()
    .node("interface")
    .out()
    .node("link")
    .in_()
    .node("interface")
    .in_()
    .node("system", role="leaf")
)
```

Execute Query

Close

Query

match(node("system", role="spine") .out()
.node("interface") .out() .node("link") .in_()
.node("interface") .in_() .node("system", role="leaf"))

Steps

start 0 1 2 3 4 5 6 7 8

<RelationshipSourceAction index=7 type=system role=== leaf>

Paths (8)

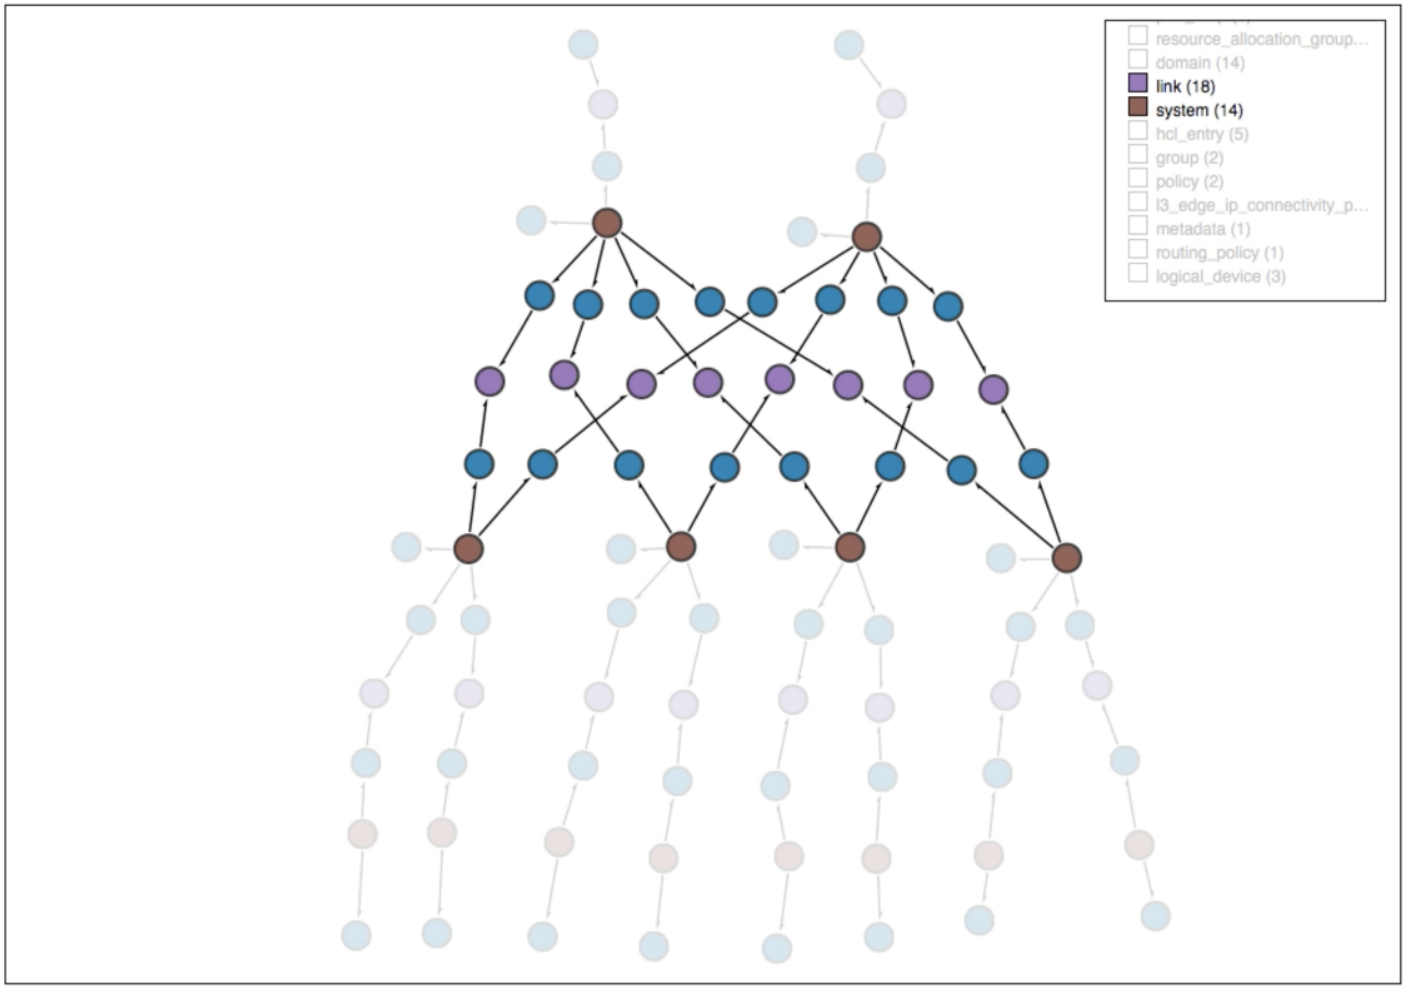

Query:

```
match(
    node("system", role="spine")
    .out()
    .node("interface")
    .out()
    .node("link")
    .in_()
    .node("interface")
    .in_()
    .node("system", role="leaf")
)
```

Execute Query

Close

Query

match(node("system", role="spine") .out()
.node("interface") .out() .node("link") .in_()
.node("interface") .in_() .node("system", role="leaf"))

Steps

start 0 1 2 3 4 5 6 7 8

<RelationshipSourceAction index=7 type=system role=== leaf>

Paths (8)

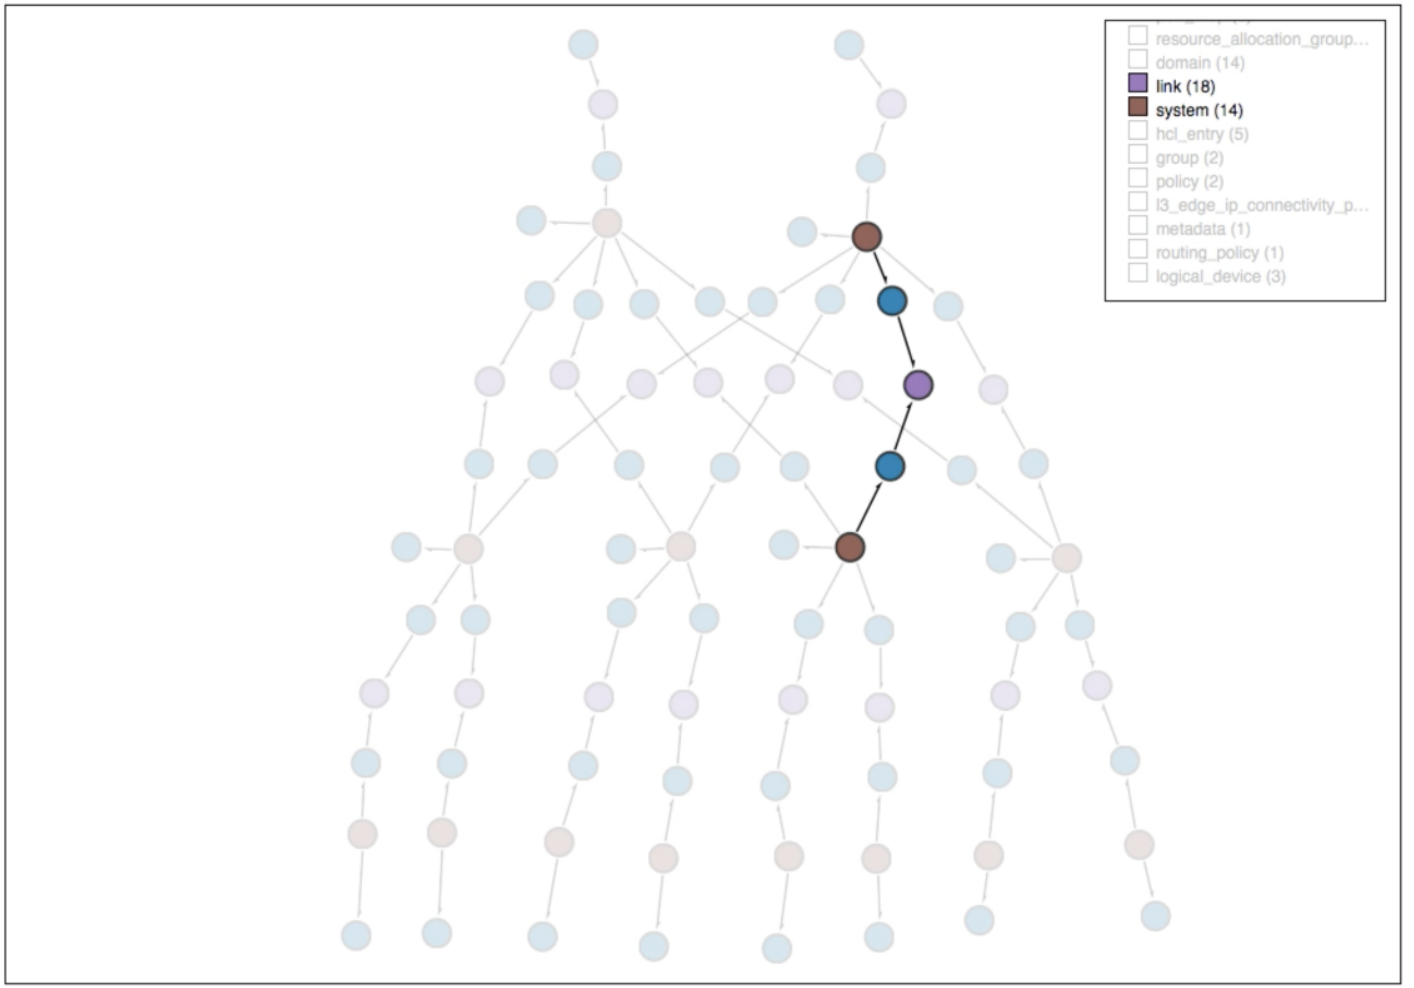

Query:

```
match(
    node("system", role="spine")
    .out()
    .node("interface")
    .out()
    .node("link")
    .in_()
    .node("interface")
    .in_()
    .node("system", role="leaf")
)
```

Execute Query

Close

Query

match(node("system", role="spine") .out()
.node("interface") .out() .node("link") .in_()
.node("interface") .in_() .node("system", role="leaf"))

Steps

start 0 1 2 3 4 5 6 7 8

<RelationshipSourceAction index=7 type=system role=== leaf>

Paths (8)

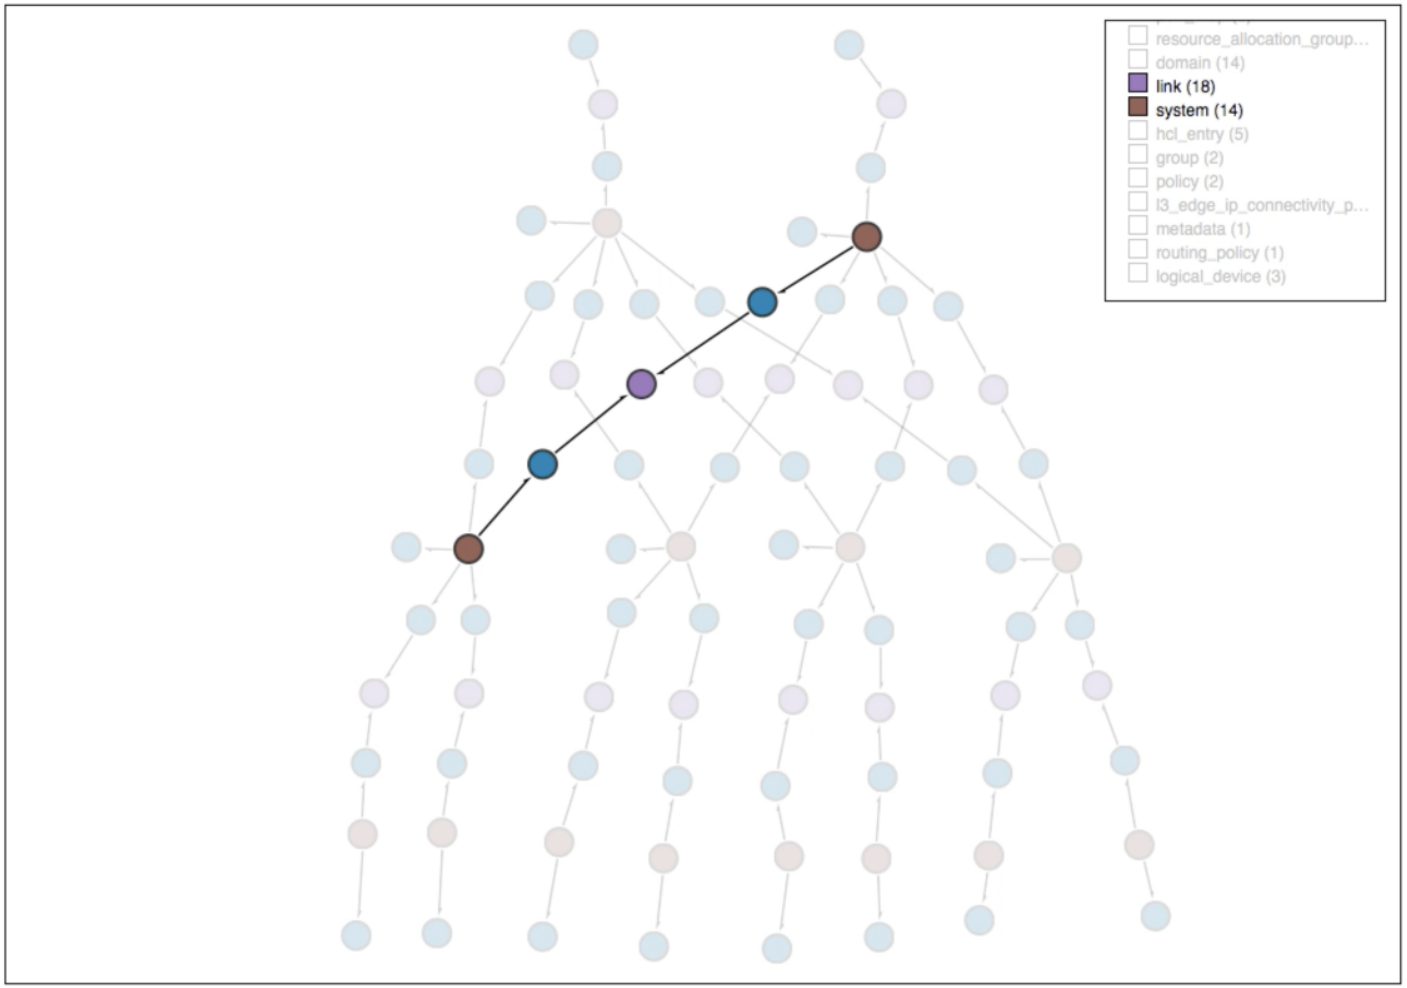

Query:

```
match(
    node("system", role="spine")
    .out()
    .node("interface")
    .out()
    .node("link")
    .in_()
    .node("interface")
    .in_()
    .node("system", role="leaf")
)
```

Execute Query

Close

Query

match(node("system", role="spine") .out()
.node("interface") .out() .node("link") .in_()
.node("interface") .in_() .node("system", role="leaf"))

Steps

start 0 1 2 3 4 5 6 7 8

<RelationshipSourceAction index=7 type=system role=== leaf>

Paths (8)

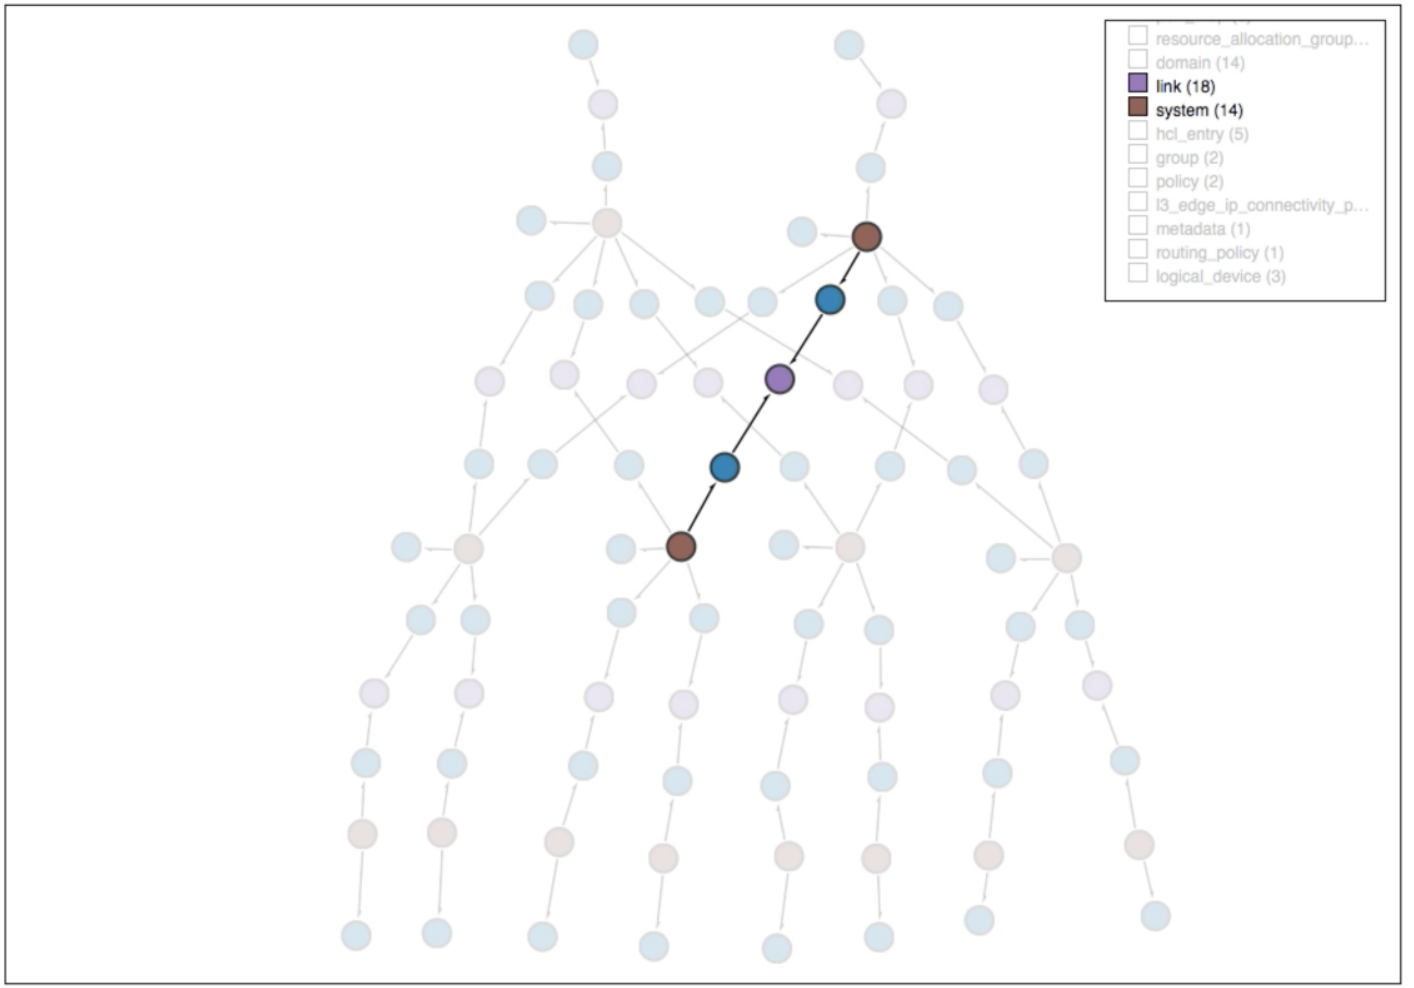

Query:

```
match(
    node("system", role="spine")
    .out()
    .node("interface")
    .out()
    .node("link")
    .in_()
    .node("interface")
    .in_()
    .node("system", role="leaf")
)
```

Execute Query

Close

Query

match(node("system", role="spine") .out()
.node("interface") .out() .node("link") .in_()
.node("interface") .in_() .node("system", role="leaf"))

Steps

start 0 1 2 3 4 5 6 7 8

<RelationshipSourceAction index=7 type=system role=== leaf>

Paths (8)

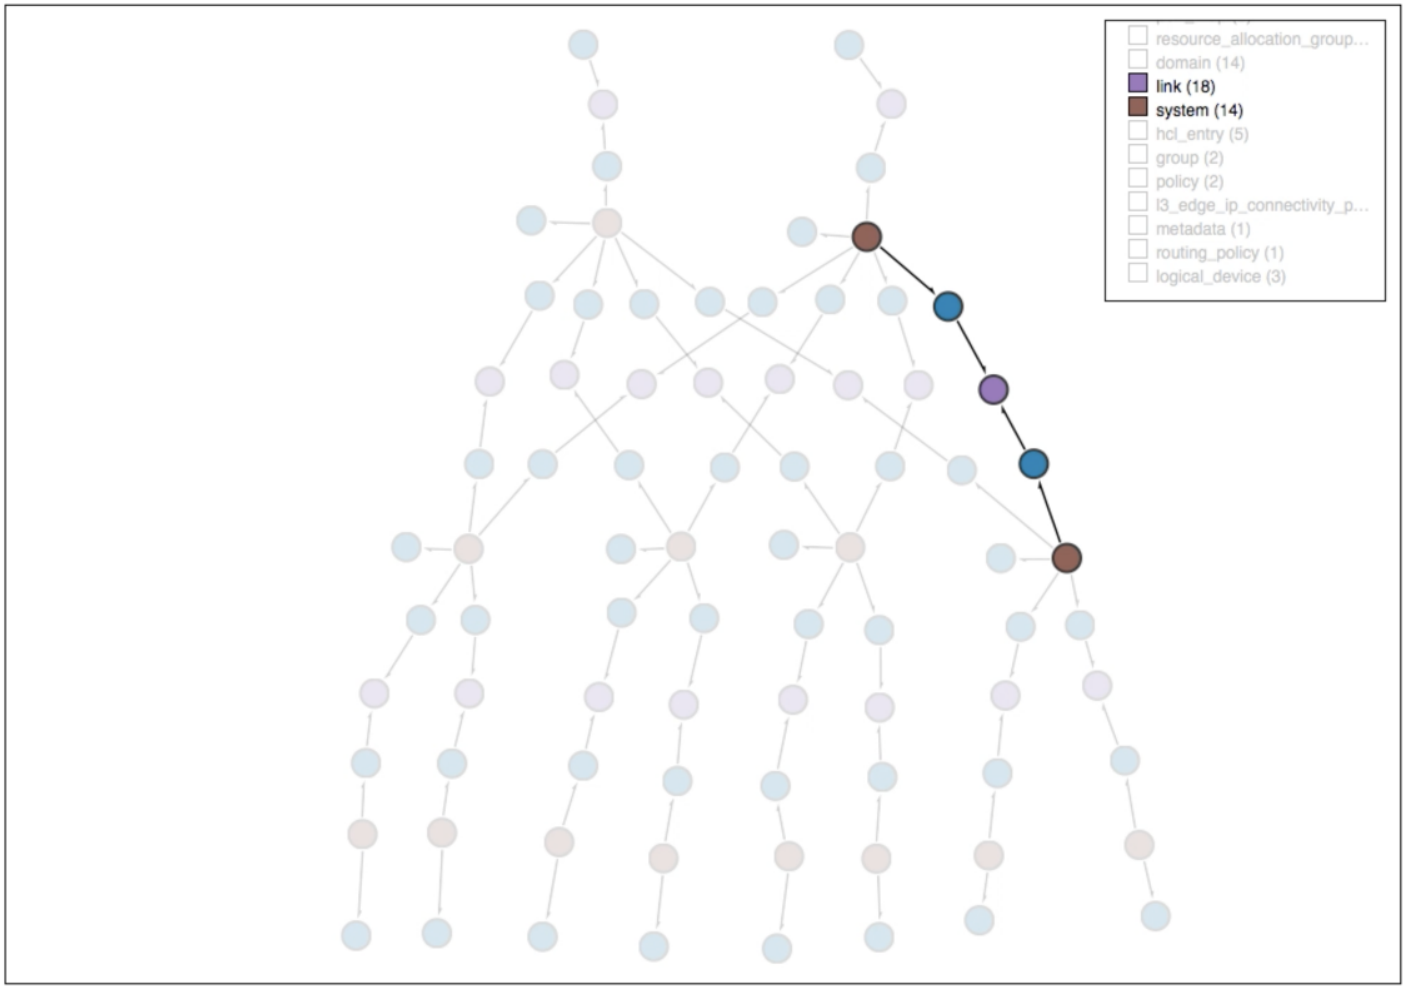

Query:

```
match(
    node("system", role="spine")
    .out()
    .node("interface")
    .out()
    .node("link")
    .in_()
    .node("interface")
    .in_()
    .node("system", role="leaf")
)
```

Execute Query

Close

Query

match(node("system", role="spine") .out()
.node("interface") .out() .node("link") .in_()
.node("interface") .in_() .node("system", role="leaf"))

Steps

start 0 1 2 3 4 5 6 7 8

<RelationshipSourceAction index=7 type=system role=== leaf>

Paths (8)

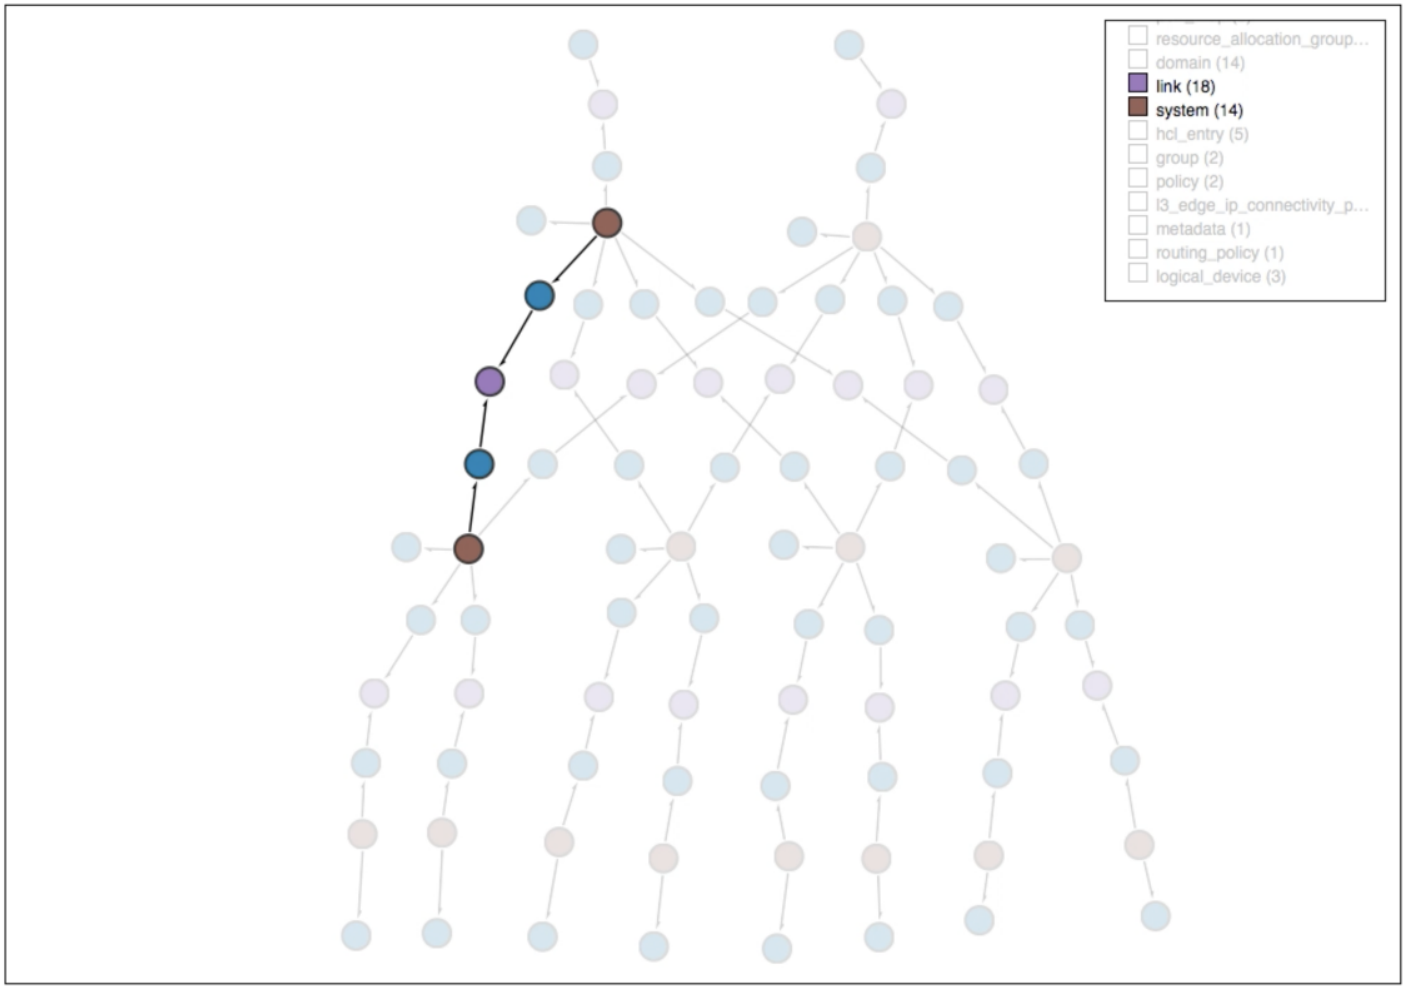

Query:

```
match(
    node("system", role="spine")
    .out()
    .node("interface")
    .out()
    .node("link")
    .in_()
    .node("interface")
    .in_()
    .node("system", role="leaf")
)
```

Execute Query

Close

Query

match(node("system", role="spine") .out()
.node("interface") .out() .node("link") .in_()
.node("interface") .in_() .node("system", role="leaf"))

Steps

start 0 1 2 3 4 5 6 7 8

<RelationshipSourceAction index=7 type=system role=== leaf>

Paths (8)

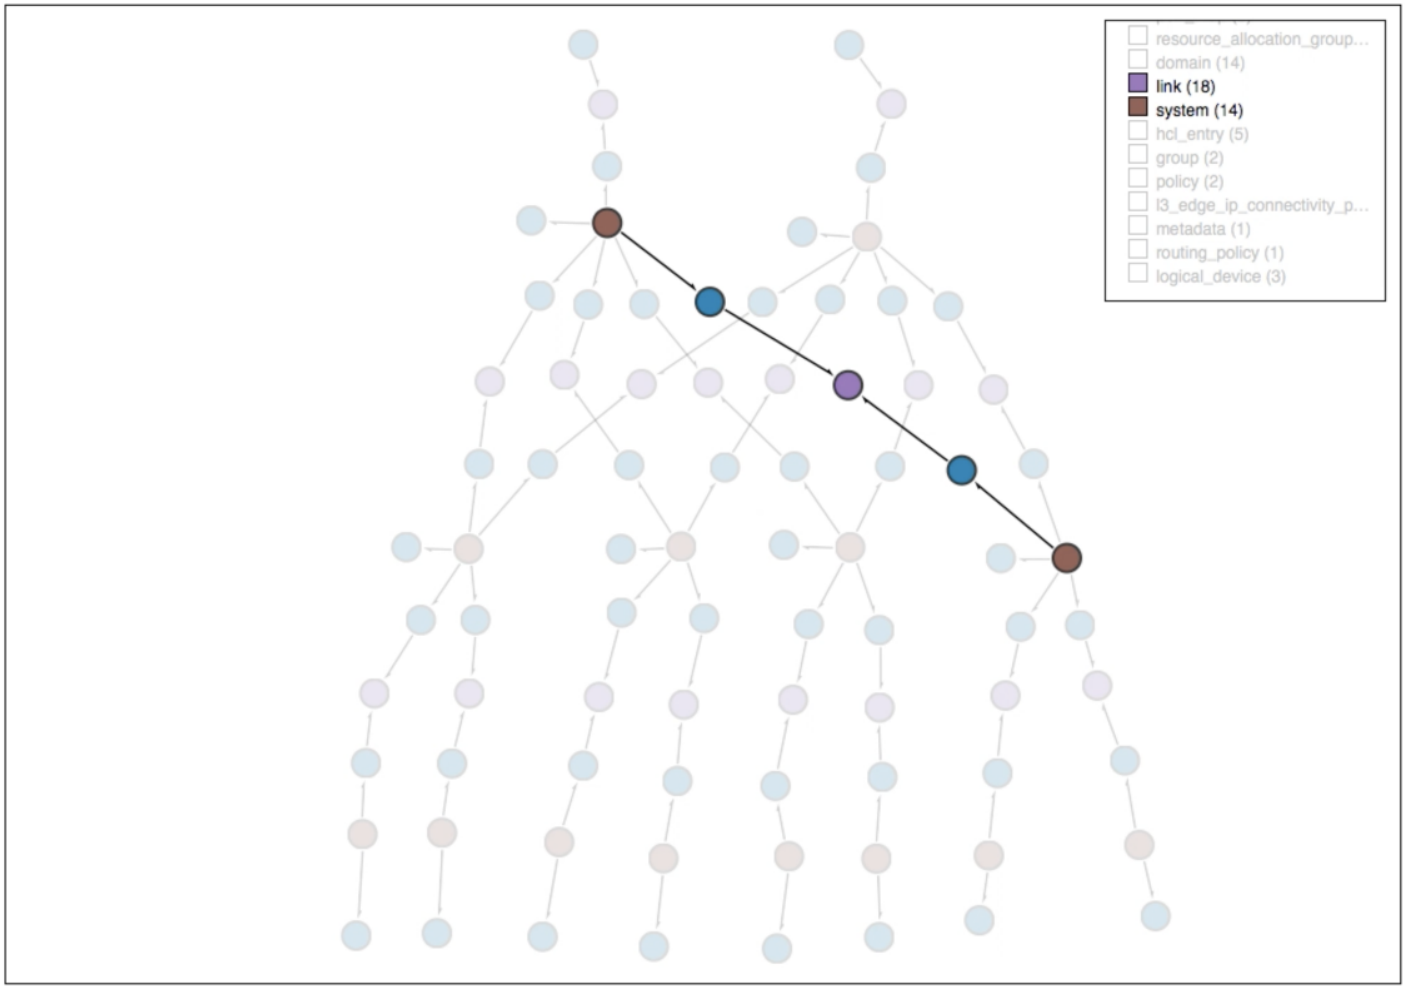

Query:

```
match(
    node("system", role="spine")
    .out()
    .node("interface")
    .out()
    .node("link")
    .in_()
    .node("interface")
    .in_()
    .node("system", role="leaf")
)
```

Execute Query

Close

Query

match(node("system", role="spine") .out()
.node("interface") .out() .node("link") .in_()
.node("interface") .in_() .node("system", role="leaf"))

Steps

start 0 1 2 3 4 5 6 7 8

<RelationshipSourceAction index=7 type=system role=== leaf>

Paths (8)

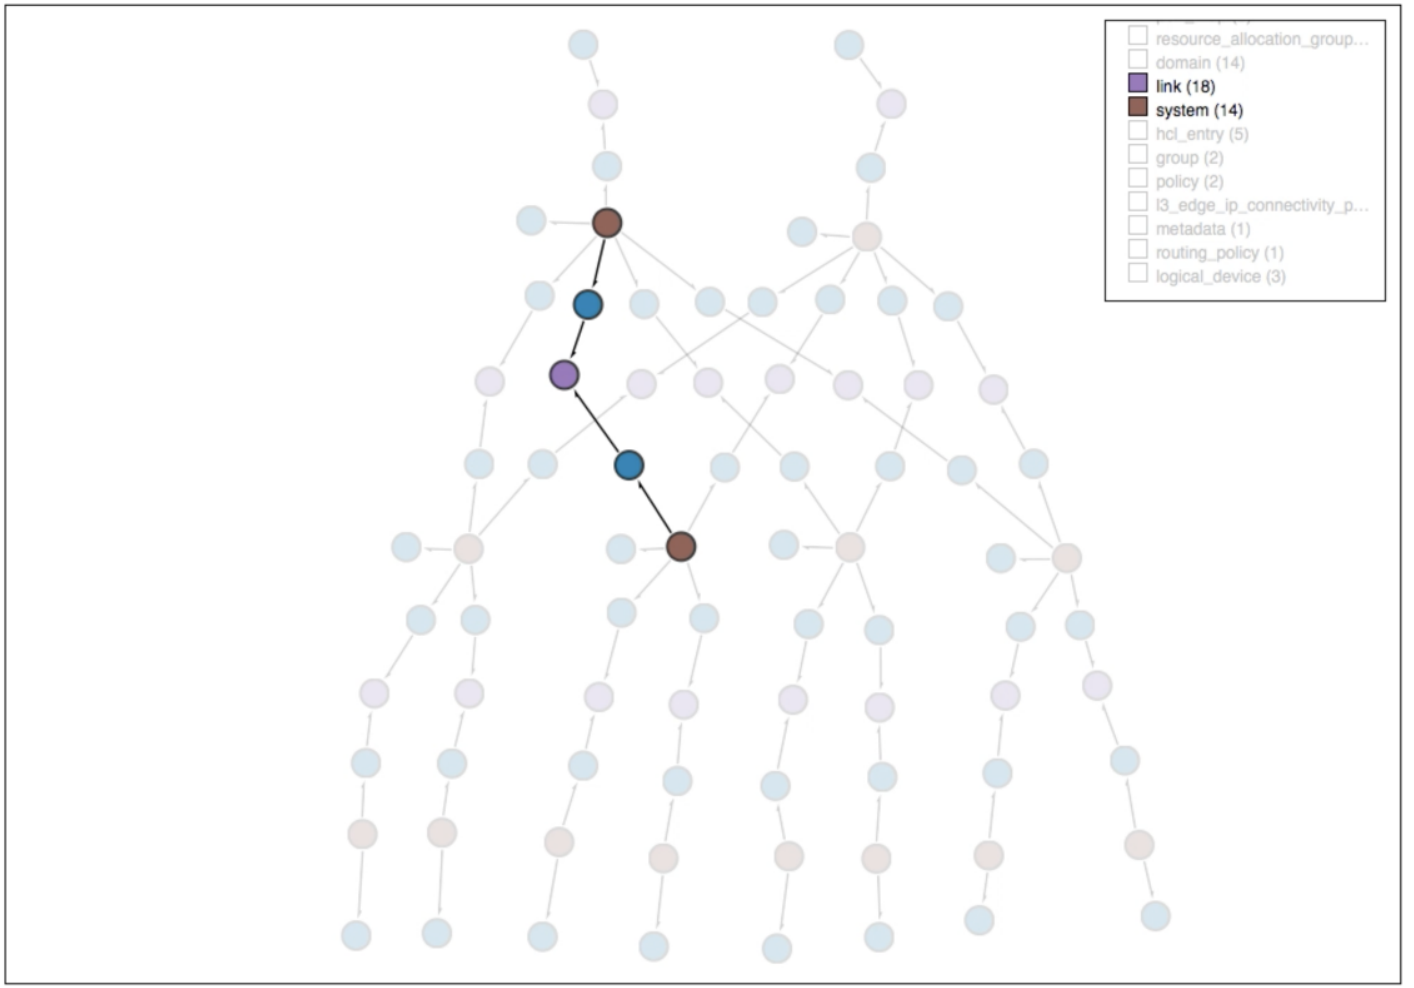

Query:

```
match(
    node("system", role="spine")
    .out()
    .node("interface")
    .out()
    .node("link")
    .in_()
    .node("interface")
    .in_()
    .node("system", role="leaf")
)
```

Execute Query

Close

Query

match(node("system", role="spine") .out()
.node("interface") .out() .node("link") .in_()
.node("interface") .in_() .node("system", role="leaf"))

Steps

start 0 1 2 3 4 5 6 7 8

<RelationshipSourceAction index=7 type=system role=== leaf>

Paths (8)

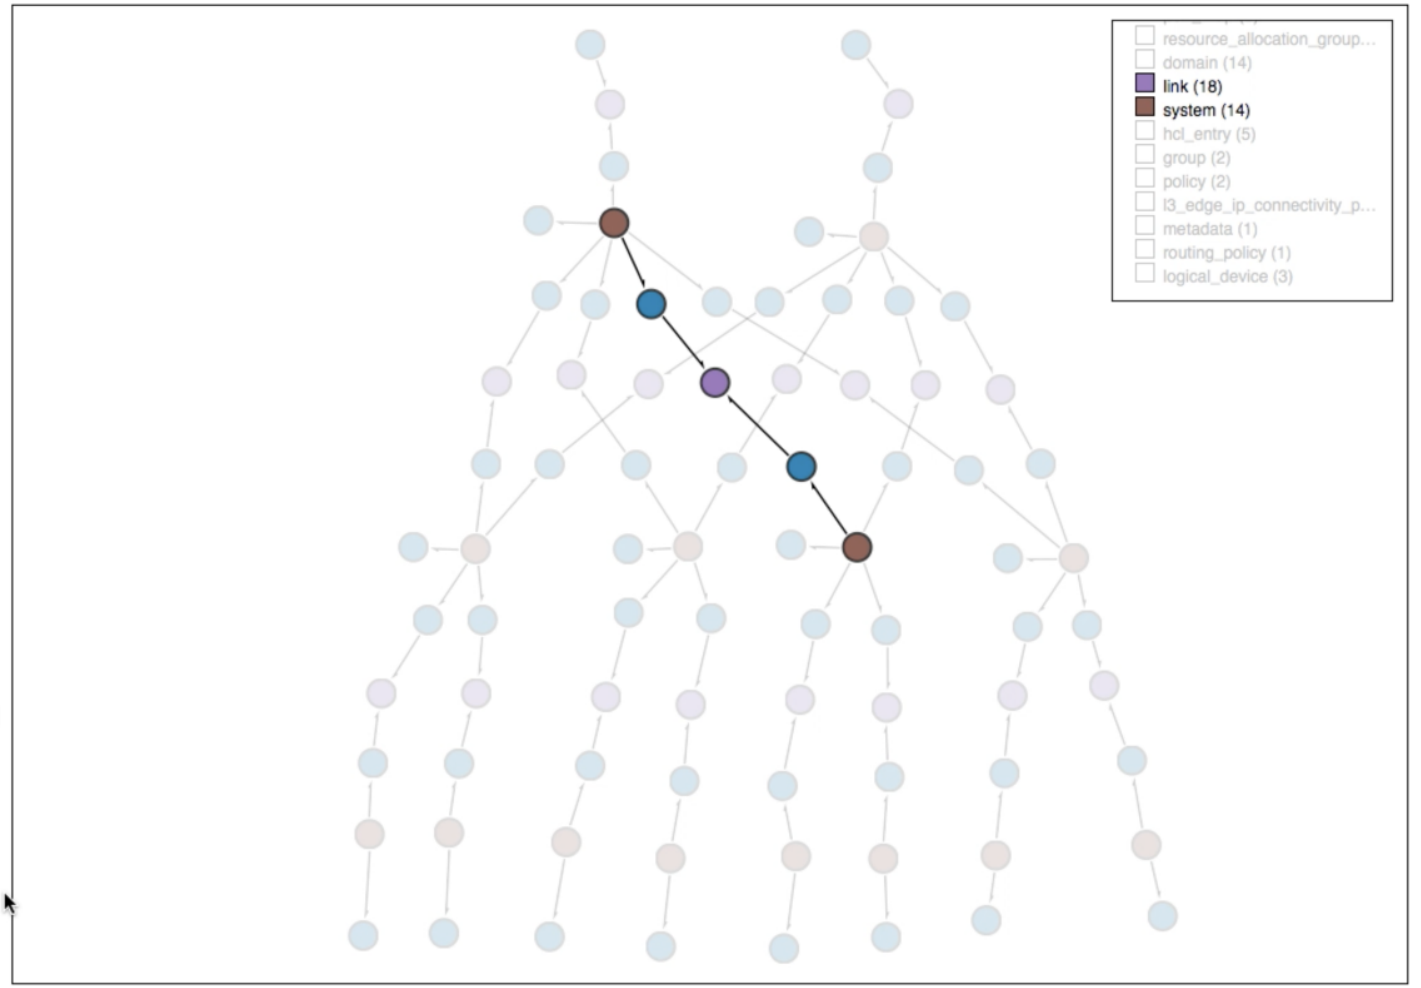

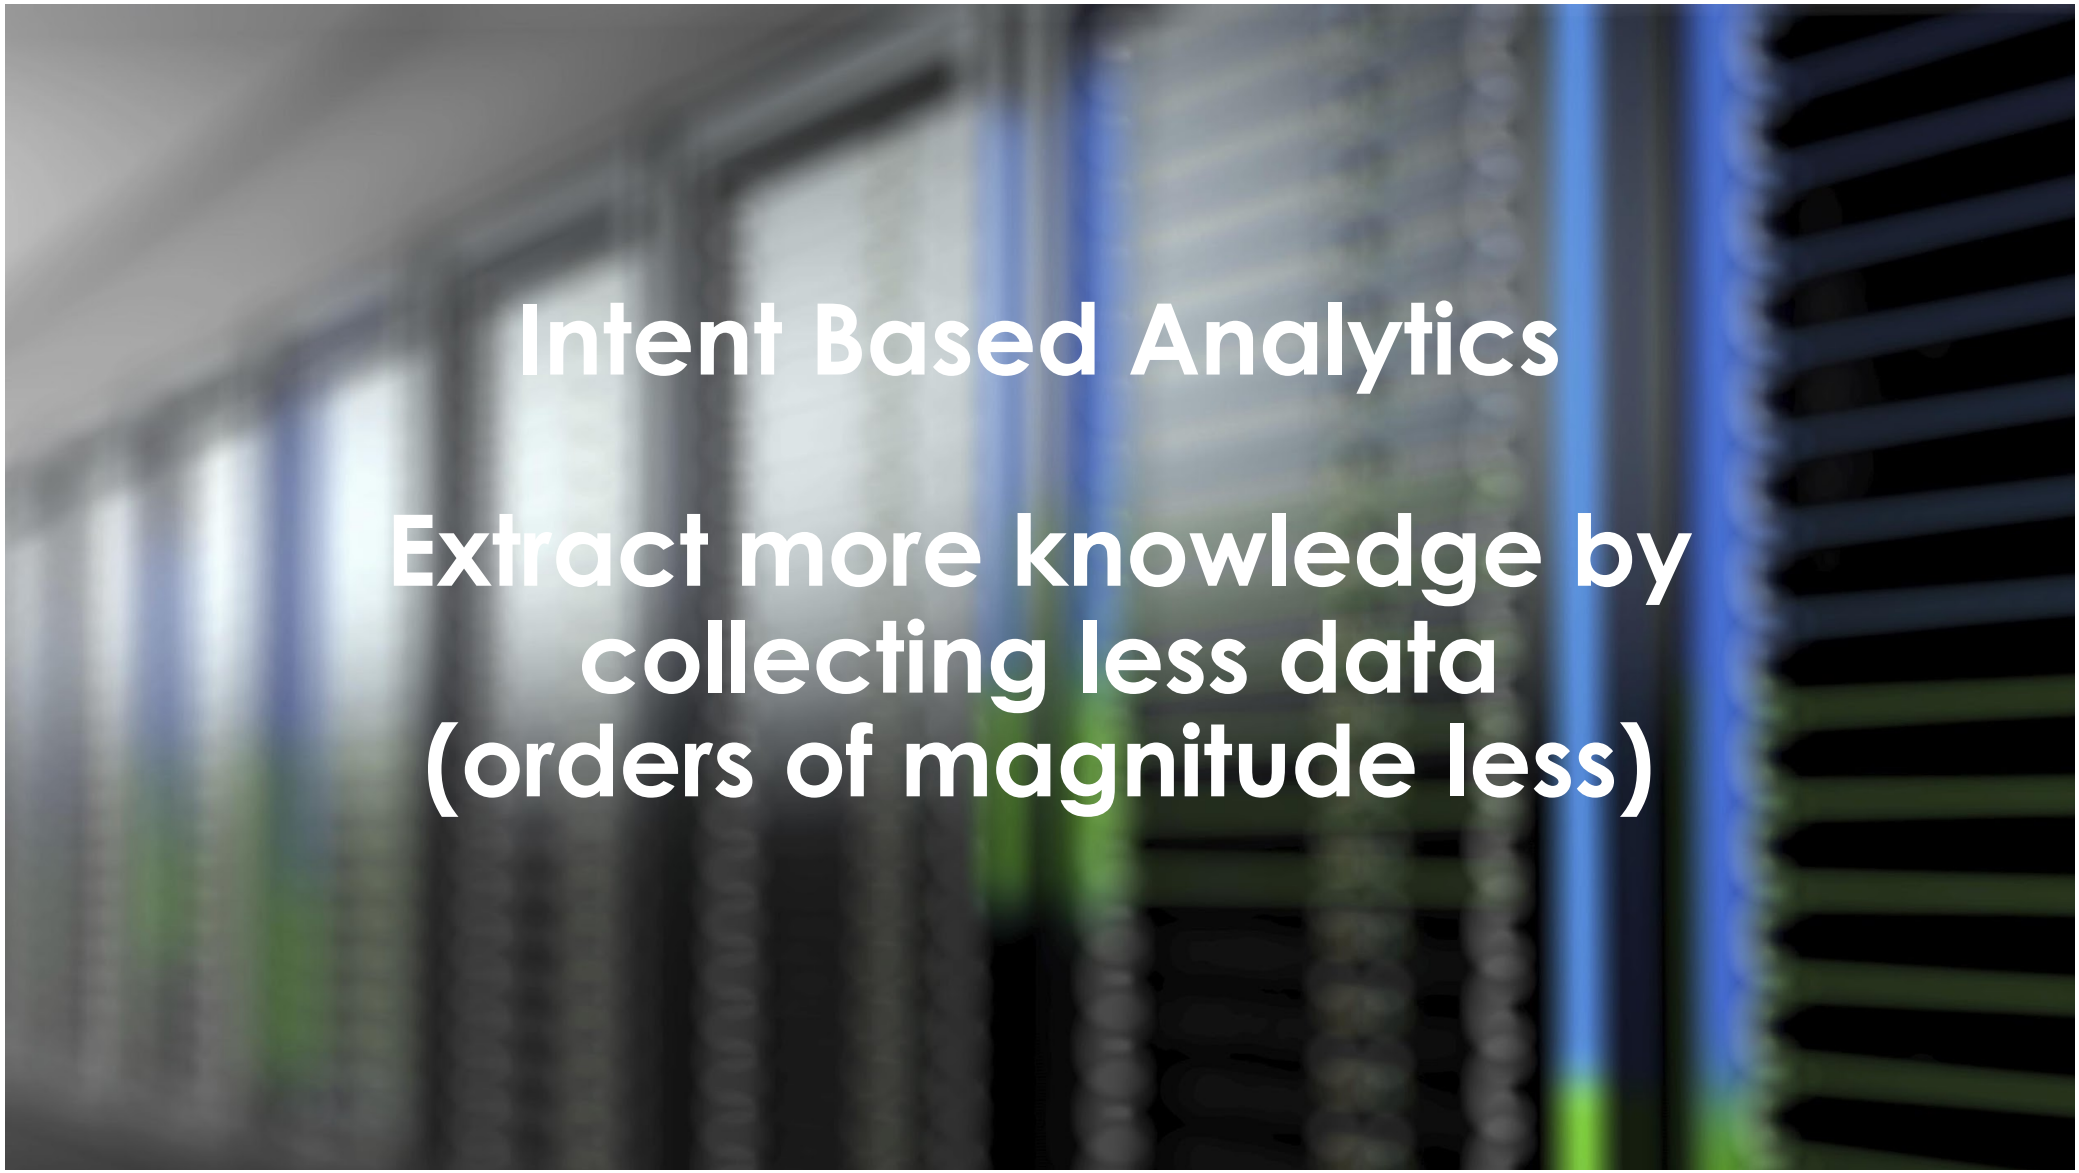
Intent Based Analytics
Extract more knowledge by collecting less data (orders of magnitude less)

Was I looking for something?

apstra

Gathering high def telemetry

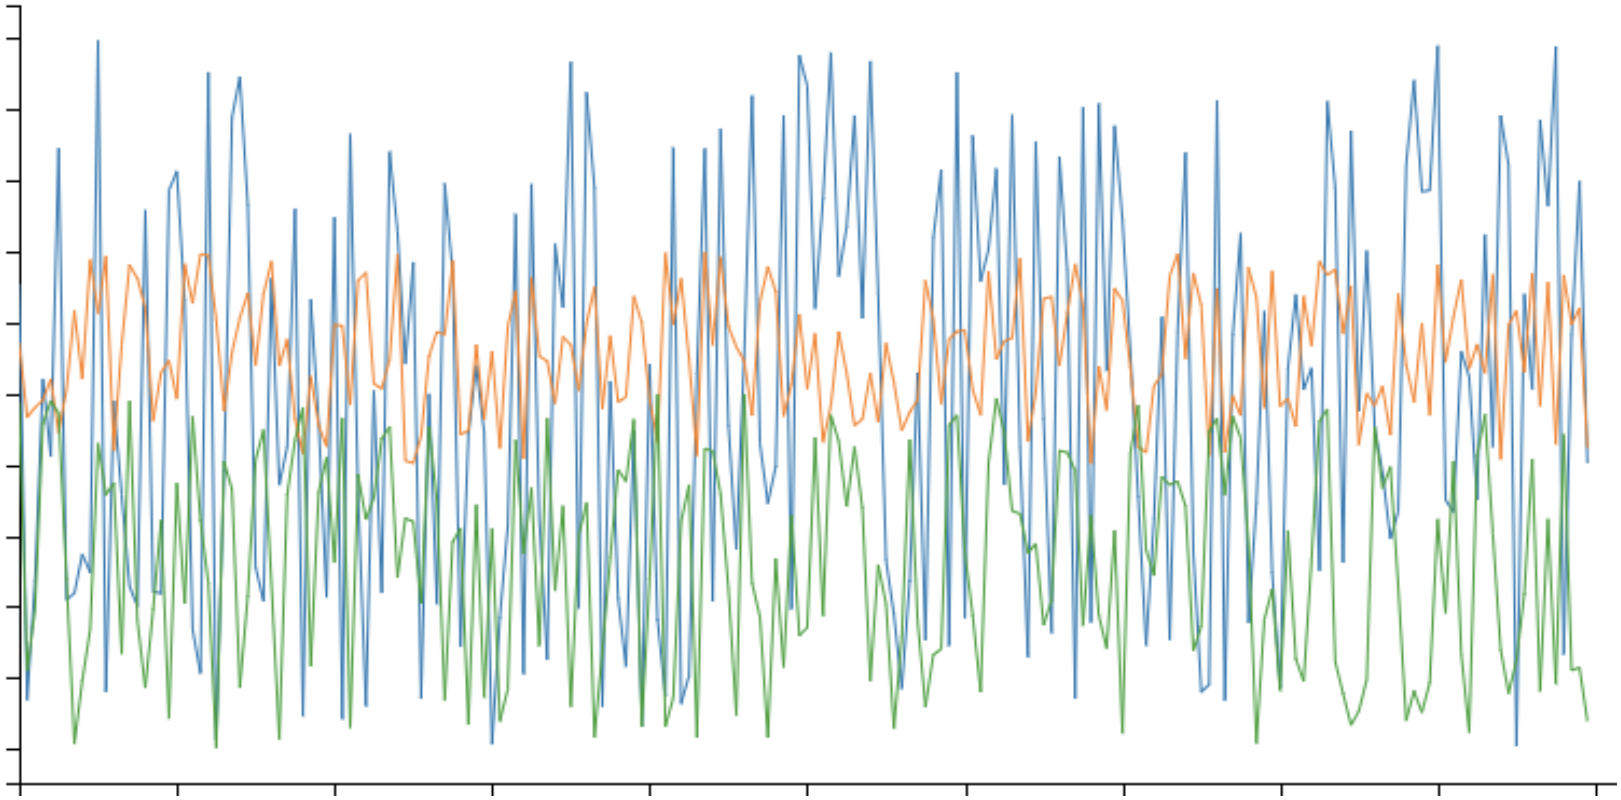

For all my leaf1 interfaces

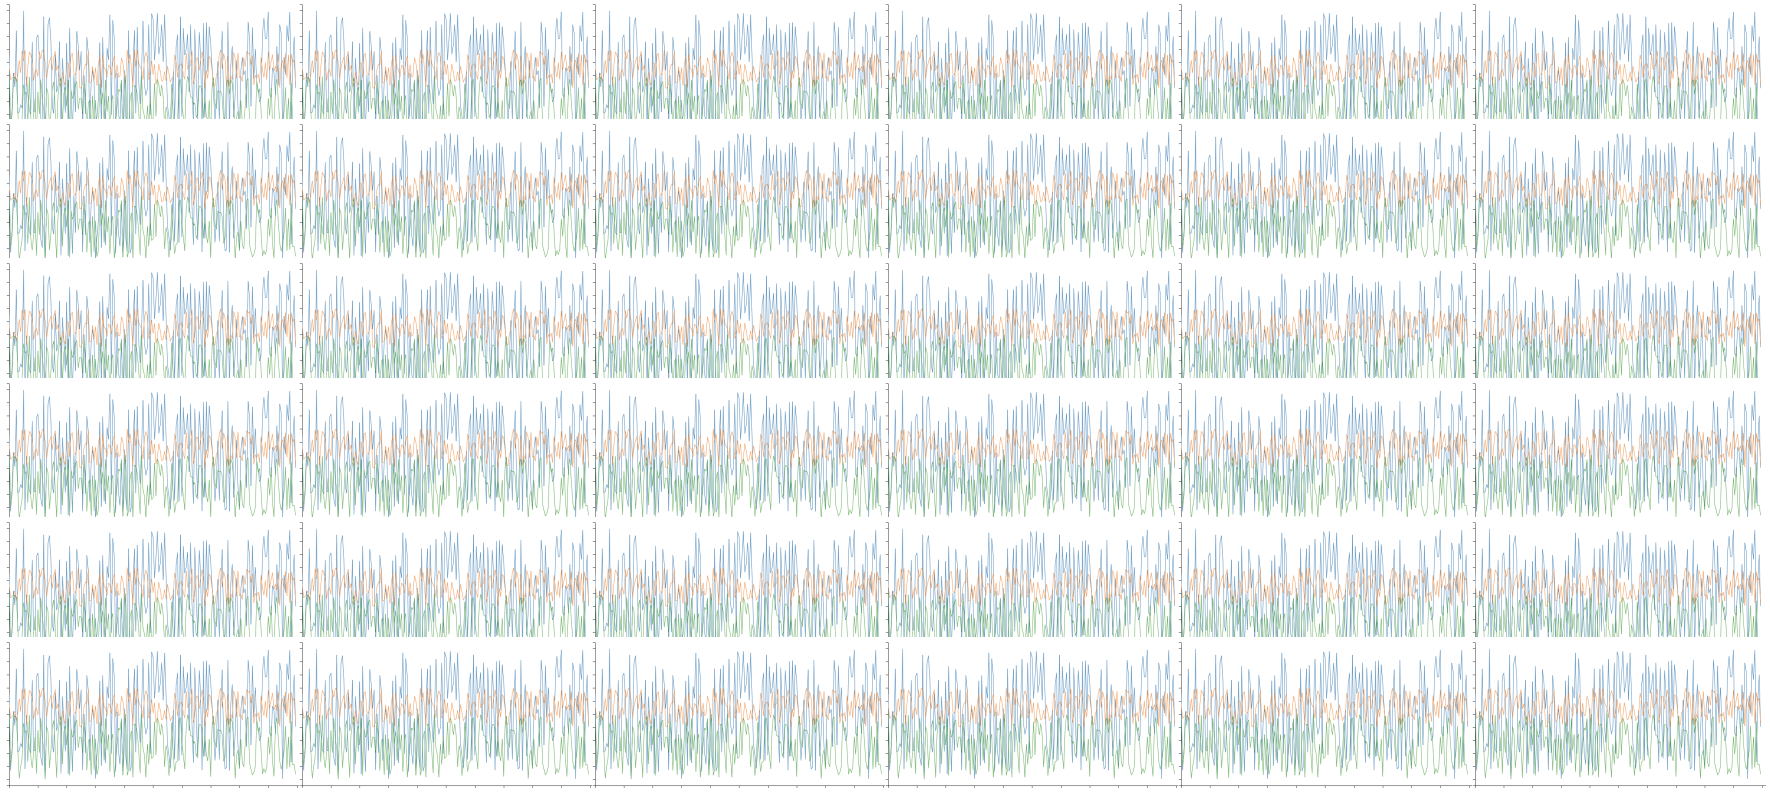

apstra

For all my leafs

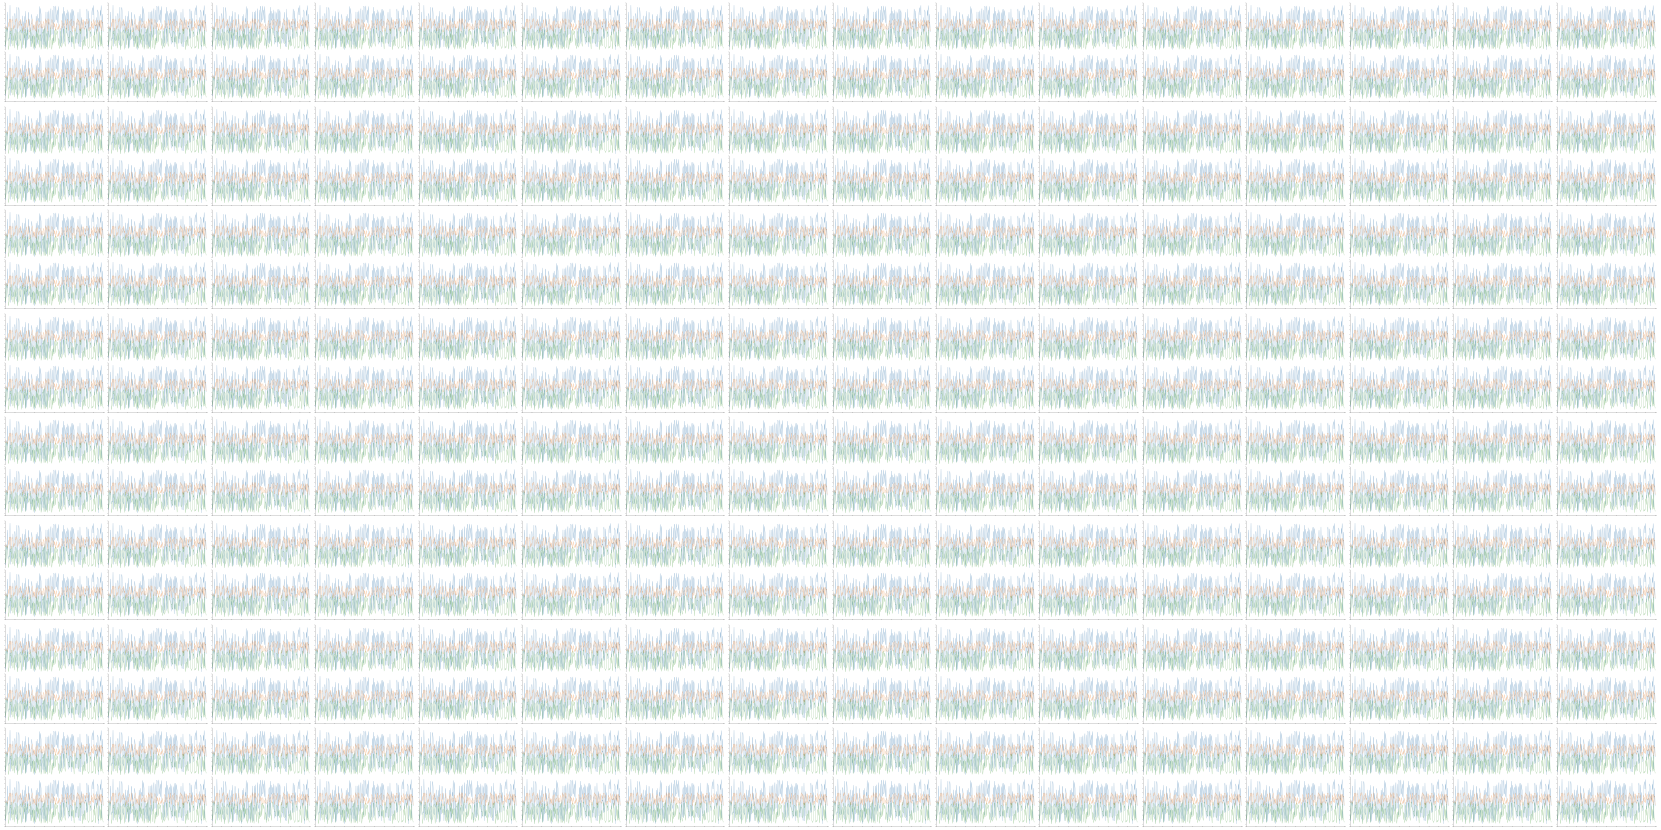

So that I have insight

IBA : ECMP fabric health (load sharing across fabric links)

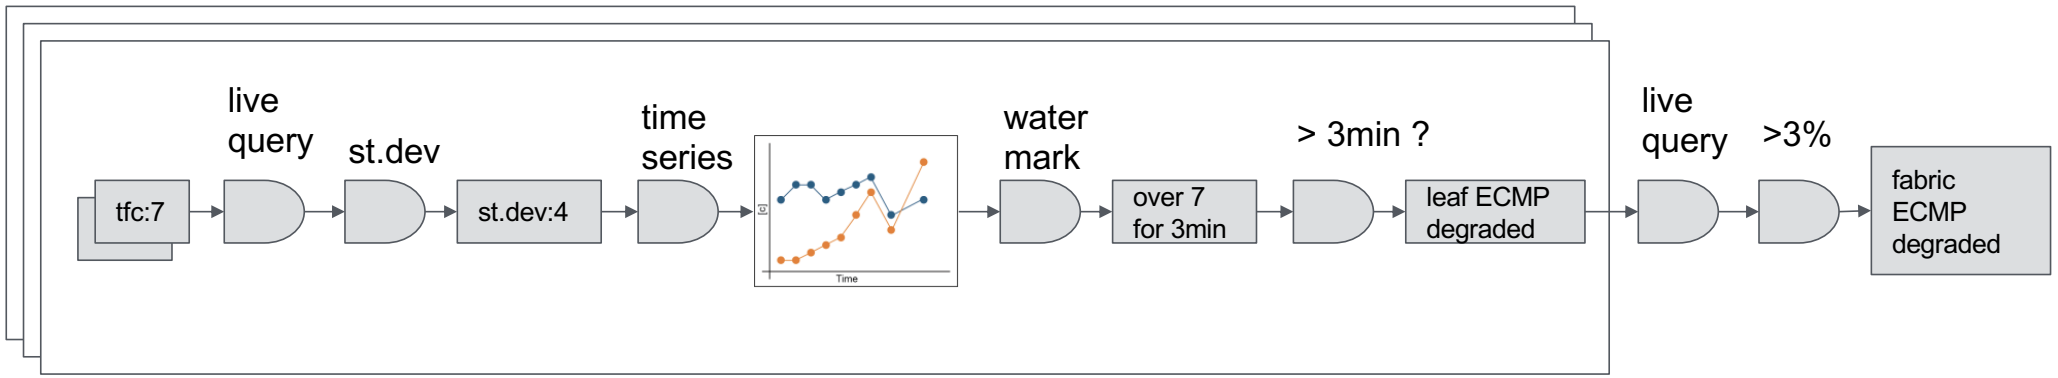

Would ML be helpful?

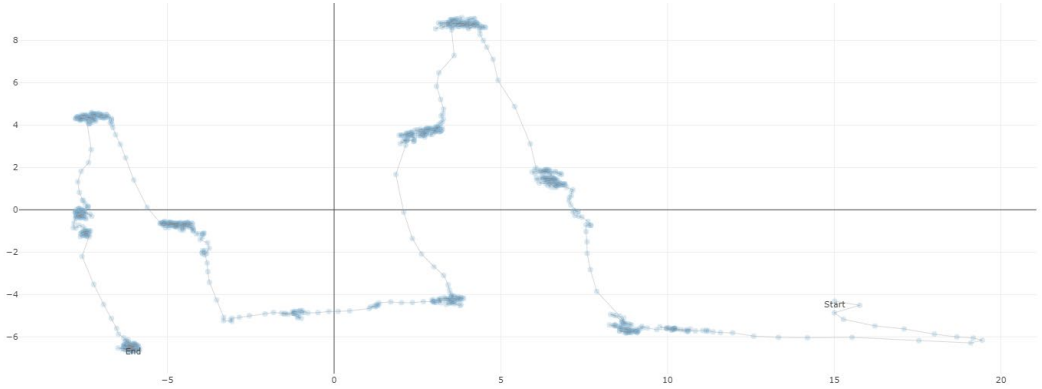

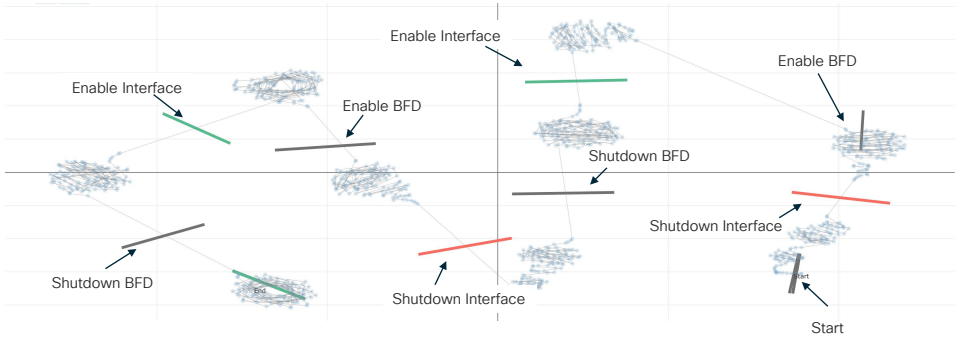

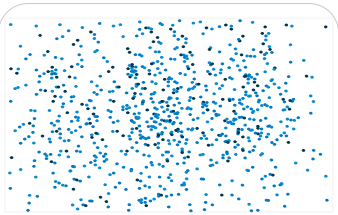

Holistic view:
Collect “all” counters all the time.

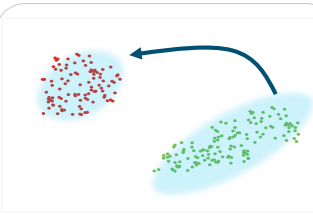

Macroscopic view:
Catch interesting state changes.
Dim.-Redux, Cluster.

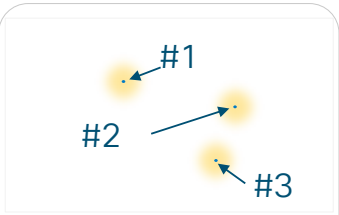

Microscopic view:
Choose counters which best describe the state change.

Source: BRKOPS-3825, Frank Brockners

IBA – context aware analytics

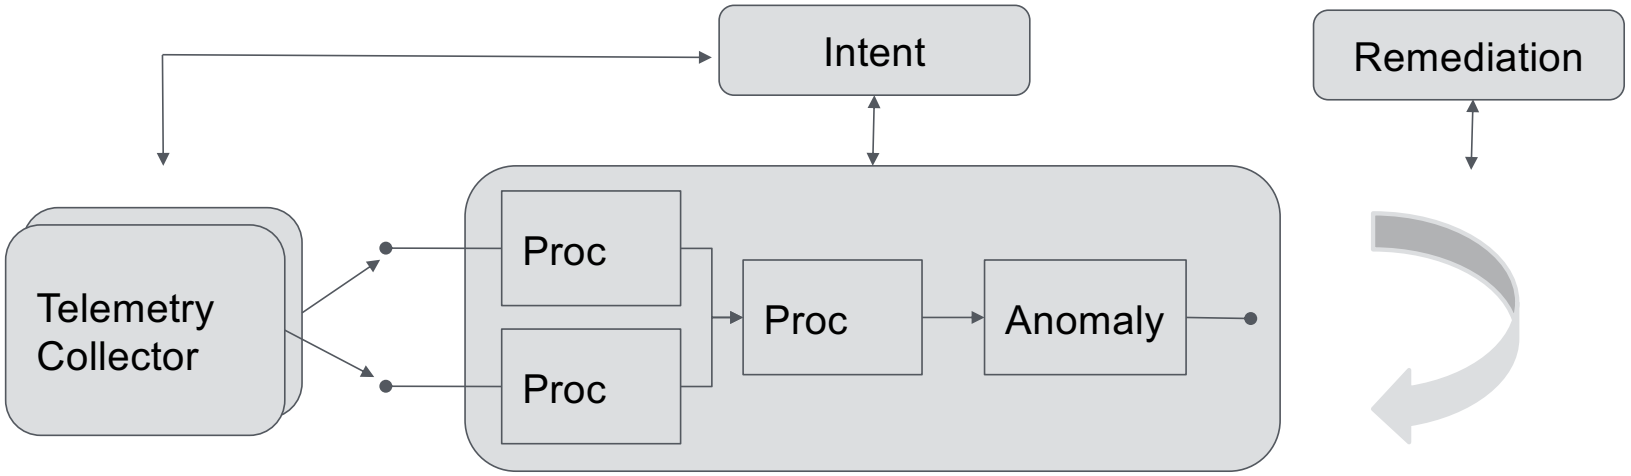

Declaratively specified, definition is de-coupled from instantiation

Once specified, is in constant sync with intent

Extracts knowledge out of the raw telemetry – context drives the content

New telemetry is “wired-in”

Conclusion

- Basic automation, while hot topic - is the first and easiest step in the IBN journey
- Single source of truth is mandatory for an IBN system to be able to reason about any change
- Day 2 operations @scale:
  - context aware continues validation
  - dealing with changes
  - configuration drift
  - remediation

  is the most complicated area of technologies to deal with!

Questions

Thank You!

www.apstra.com


@ApstraInc


https://www.linkedin.com/company/apstra


https://www.facebook.com/apstrainc/