

A close-up photograph of a dandelion seed head, showing the intricate structure of the seeds and their fine, hair-like pappus. The image is partially obscured by a green rectangular overlay.

# FLEXALGO

---

Shraddha Hegde and Ron Bonica

May 2020

JUNIPER  
NETWORKS

Engineering  
Simplicity

# OVERVIEW

---

- In Segment Routing (SR), a prefix segment follows the least-cost path from its ingress to its egress
  - An Intradomain Gateway Protocol (IGP) calculates the least-cost path
- A network operator can
  - Configure multiple prefix segments ending on the same node
  - Use Flexalgo to influence how the IGP calculates the least-cost path for each prefix segment
- Therefore, each prefix segment can traverse a unique set of links

# HOW IT WORKS ( IGP MAGIC!! )

---

- All nodes advertise information regarding links to which they are attached
  - Metrics (IGP, TE, latency-based)
  - Administrative groups (colors)
- Selected nodes advertise Flexible Algorithm Definitions (FAD)
  - Identifier, Metric type, Calculation Type, Constraints
- A segment endpoint advertises multiple prefix segments
  - Associating each with a FAD
- Each node calculates multiple least-cost paths to the segment endpoint
  - Once for each FAD / prefix segment
- Each least-cost path can traverse a unique set of links

# POPULAR USE-CASES

---



- Low latency / high bandwidth paths
- Path Diversity



## Low Latency / High Bandwidth Paths

# POLICY

- All flows follow the lowest latency path available
  - In this network, latency is a function of circuit length
- However, high bandwidth flows must avoid 10G links



# LINK ADVERTISEMENTS

---

| Link  | IGP Metric | TE Metric | Administrative Group |
|-------|------------|-----------|----------------------|
| R1-R2 | 400        | 400       | Blue                 |
| R1-R3 | 400        | 400       | Blue                 |
| R1-R4 | 300        | 300       | Red                  |
| R2-R4 | 400        | 400       | Blue                 |
| R3-R4 | 400        | 400       | Blue                 |

# FLEXALGO DEFINITIONS (FAD)

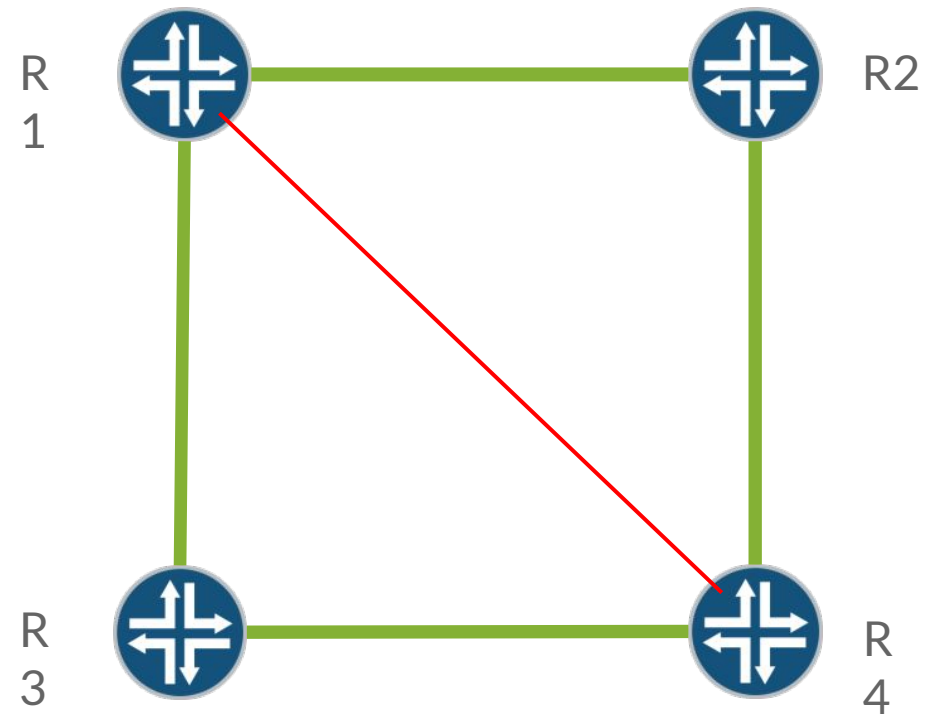
---

| FAD            | Metric Type | Calculation Type | Constraints |
|----------------|-------------|------------------|-------------|
| Low Latency    | IGP         | SPF              | Include all |
| High bandwidth | IGP         | SPF              | Exclude red |



# PULLING IT TOGETHER

- R4 advertises Segment A
  - Associates it with the low latency FAD
- R4 advertises Segment B
  - Associates it with the high bandwidth FAD
- R1 calculates the least-cost path to Segment A
  - Next Hop is R4
  - Because low latency FAD includes all links
- R1 calculates the least-cost path to Segment B
  - Next Hop is ECMP (either R2 or R3)
  - Because high bandwidth FAD excludes red links






# Path Diversity

# POLICY

- Red flows traverse red links
  - And no others
- Orange flows prefer red links
  - But can fail over to blue links
- Blue flows traverse blue links
  - And no others
- Yellow flows prefer blue links
  - But can fail over to red links



# LINK ADVERTISEMENTS

---

| Link  | IGP Metric | TE Metric | Administrative Group |
|-------|------------|-----------|----------------------|
| R1-R2 | 200        | 400       | Red                  |
| R1-R3 | 400        | 200       | Blue                 |
| R2-R4 | 200        | 400       | Red                  |
| R3-R4 | 400        | 200       | Blue                 |


# FLEXALGO DEFINITIONS (FAD)

---

| FAD    | Metric Type | Calculation Type | Constraints  |
|--------|-------------|------------------|--------------|
| Red    | IGP         | SPF              | Exclude blue |
| Orange | IGP         | SPF              | Include all  |
| Blue   | TE          | SPF              | Exclude red  |
| Yellow | TE          | SPF              | Include all  |


# PULLING IT TOGETHER

- R4 advertises four prefix segments
  - Segment A associated with the red FAD
  - Segment B associated with the orange FAD
  - Segment C associated with the blue FAD
  - Segment D associated with the yellow FAD
- R1 calculates the least-cost path to R4 four times
  - Once for each FAD / prefix segment




# R1 ROUTES TO R4

- Via Prefix Segment A (red)
  - Next Hop is R2
  - No failover
  - Because red FAD excludes blue links
- Via Prefix Segment B (orange)
  - Next Hop is R2
    - Because orange FAD uses IGP metrics
    - Because IGP metrics are lower on red links
  - Failover is R3
    - Because orange FAD includes all links



# R1 ROUTES TO R4 (CONTINUED)

- Via Prefix Segment C (blue)
  - Next Hop is R3
  - No failover
  - Because blue FAD excludes red links
- Via Prefix Segment D (yellow)
  - Next Hop is R3
    - Because orange FAD uses TE metrics
    - Because TE metrics are lower on blue links
  - Failover is R2
    - Because yellow FAD includes all links







# Conclusion

# FLEXALGO IS POWERFUL

---

- Many networks require only course-grained TE
  - As in the use-cases described above
- Benefits of deploying Flexalgo into such networks
  - Each SR path is reduced to a single segment
  - No need to specify TE policy on a controller or on each segment egress node
  - Operational simplicity



THANK YOU

---

JUNIPER  
NETWORKS

Engineering  
Simplicity