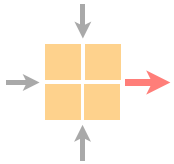

Taming the transient while reconfiguring BGP

Tibor Schneider*, Roland Schmid*,
Stefano Vissicchio‡, Laurent Vanbever*

SIGCOMM 2023, September 11, 2023

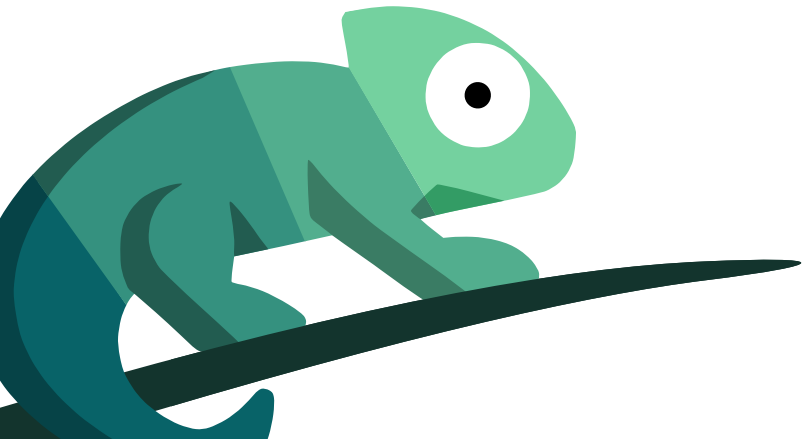

* ETH zürich
‡ UCL

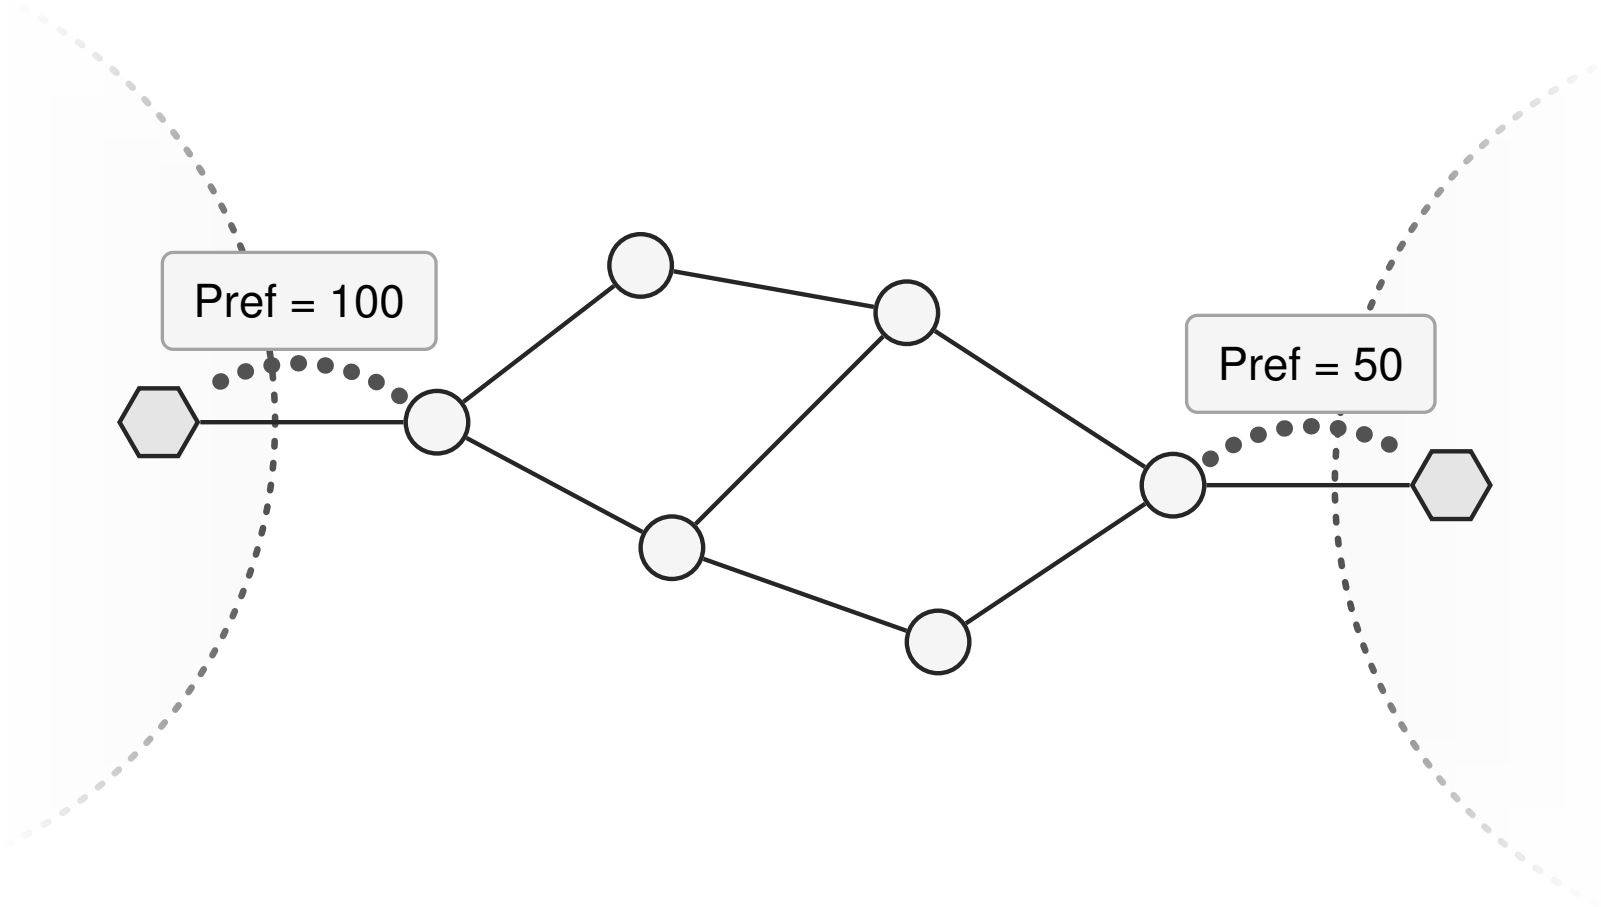
Pref = 100
Pref = 50

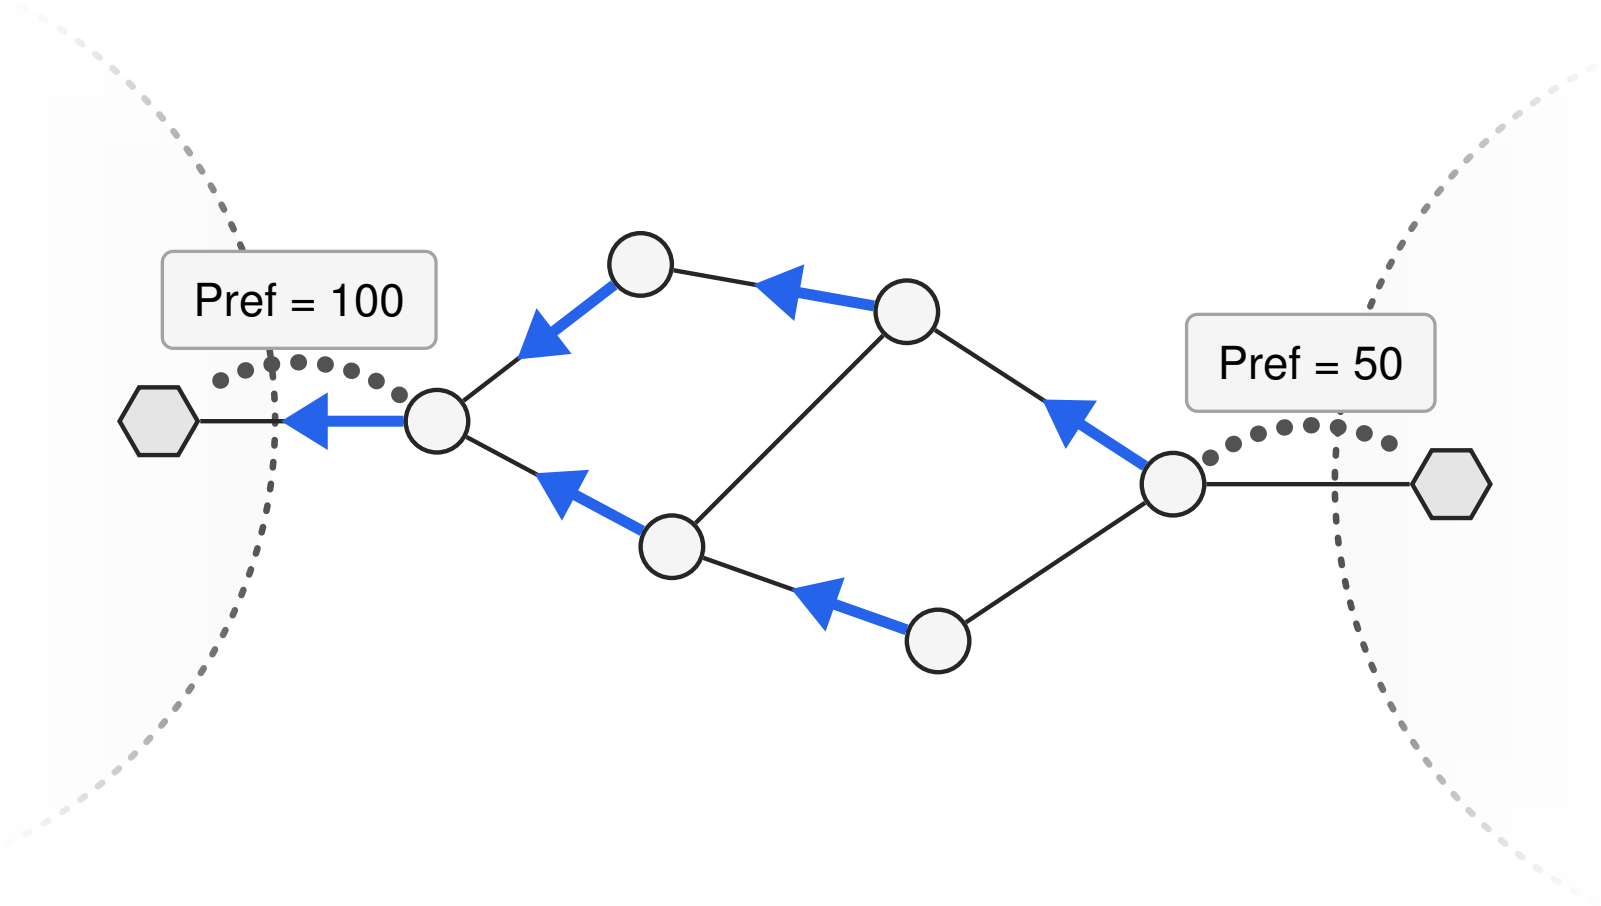
Pref = 100
Pref = 50

Remove the peering session
with the network on the left.

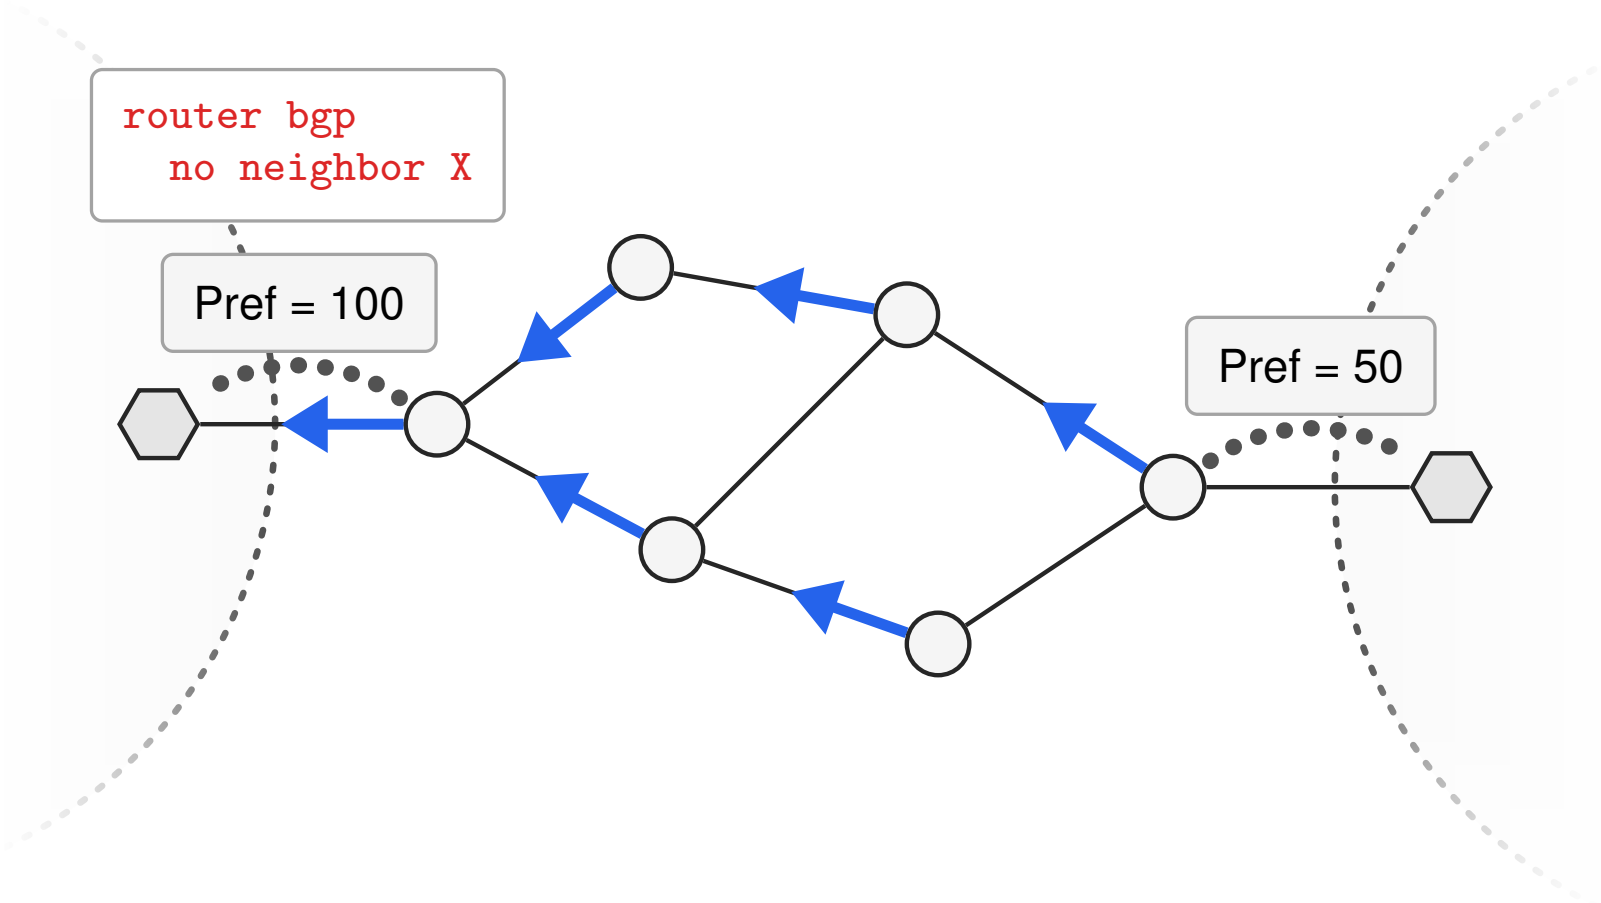

Naively reconfiguring the network causes significant packet loss.

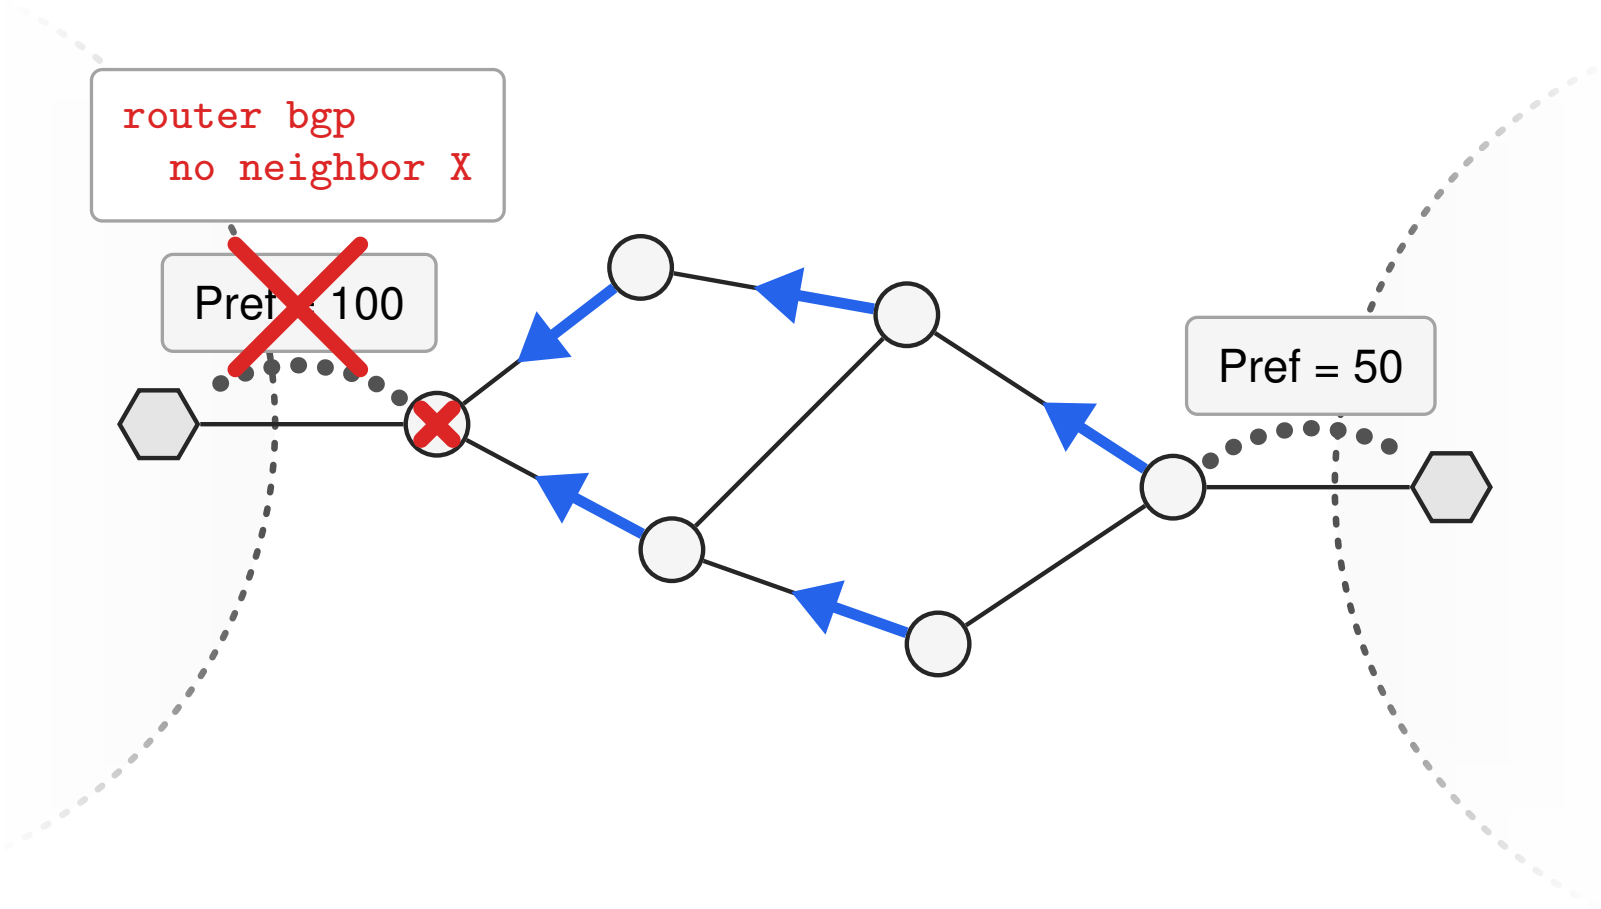

Best practice cannot give guarantees during convergence.

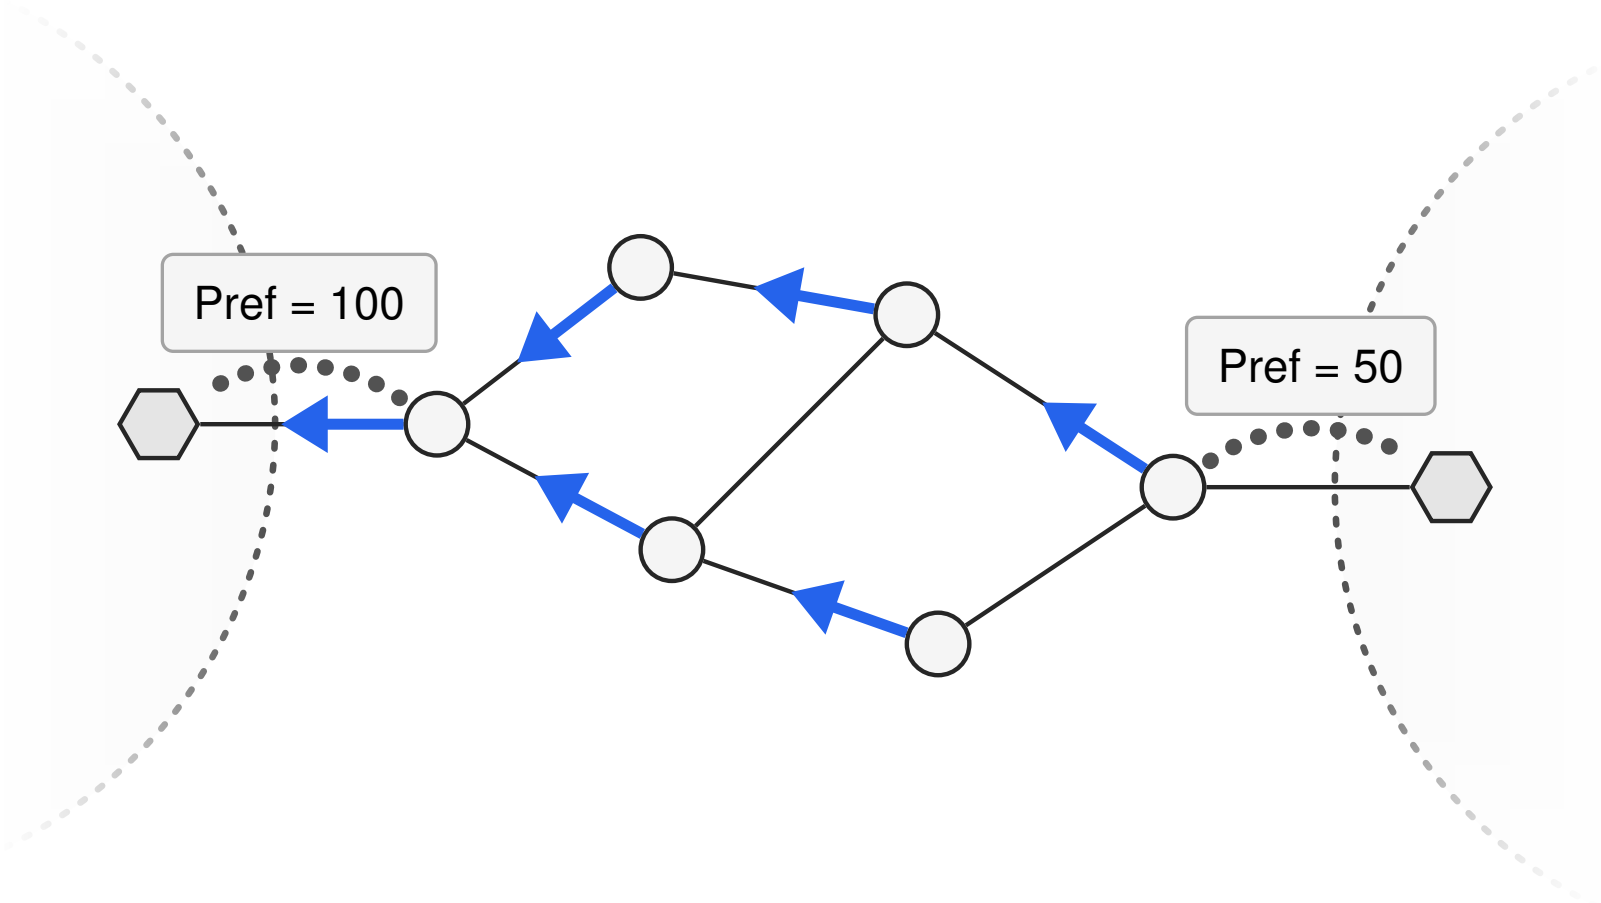

Best practice cannot give guarantees during convergence.

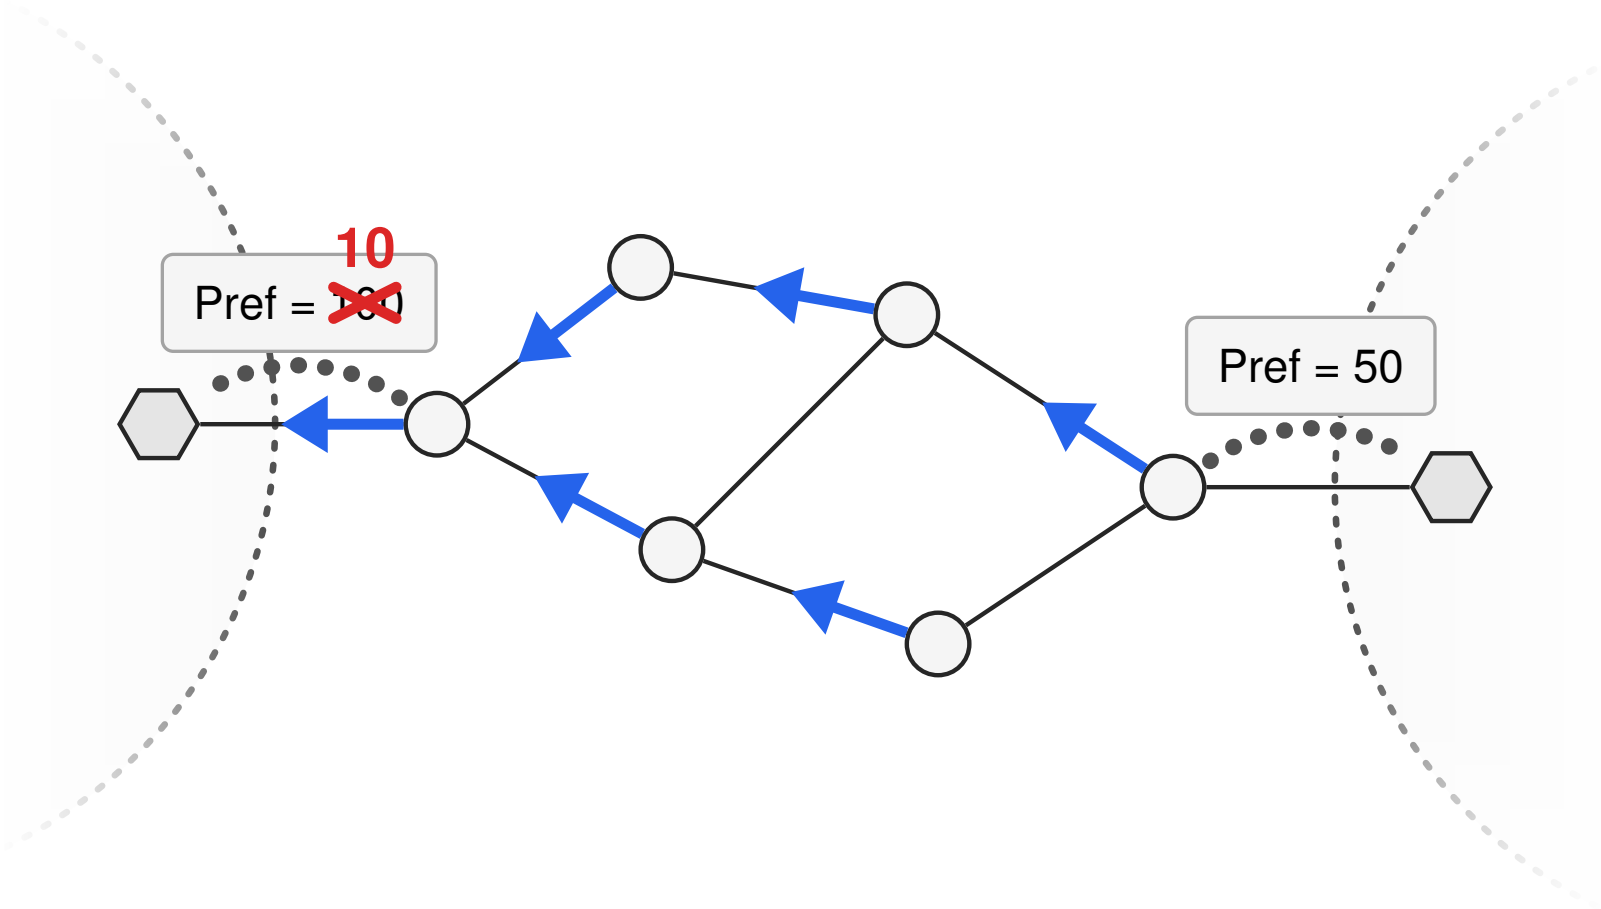

Best practice cannot give guarantees during convergence.

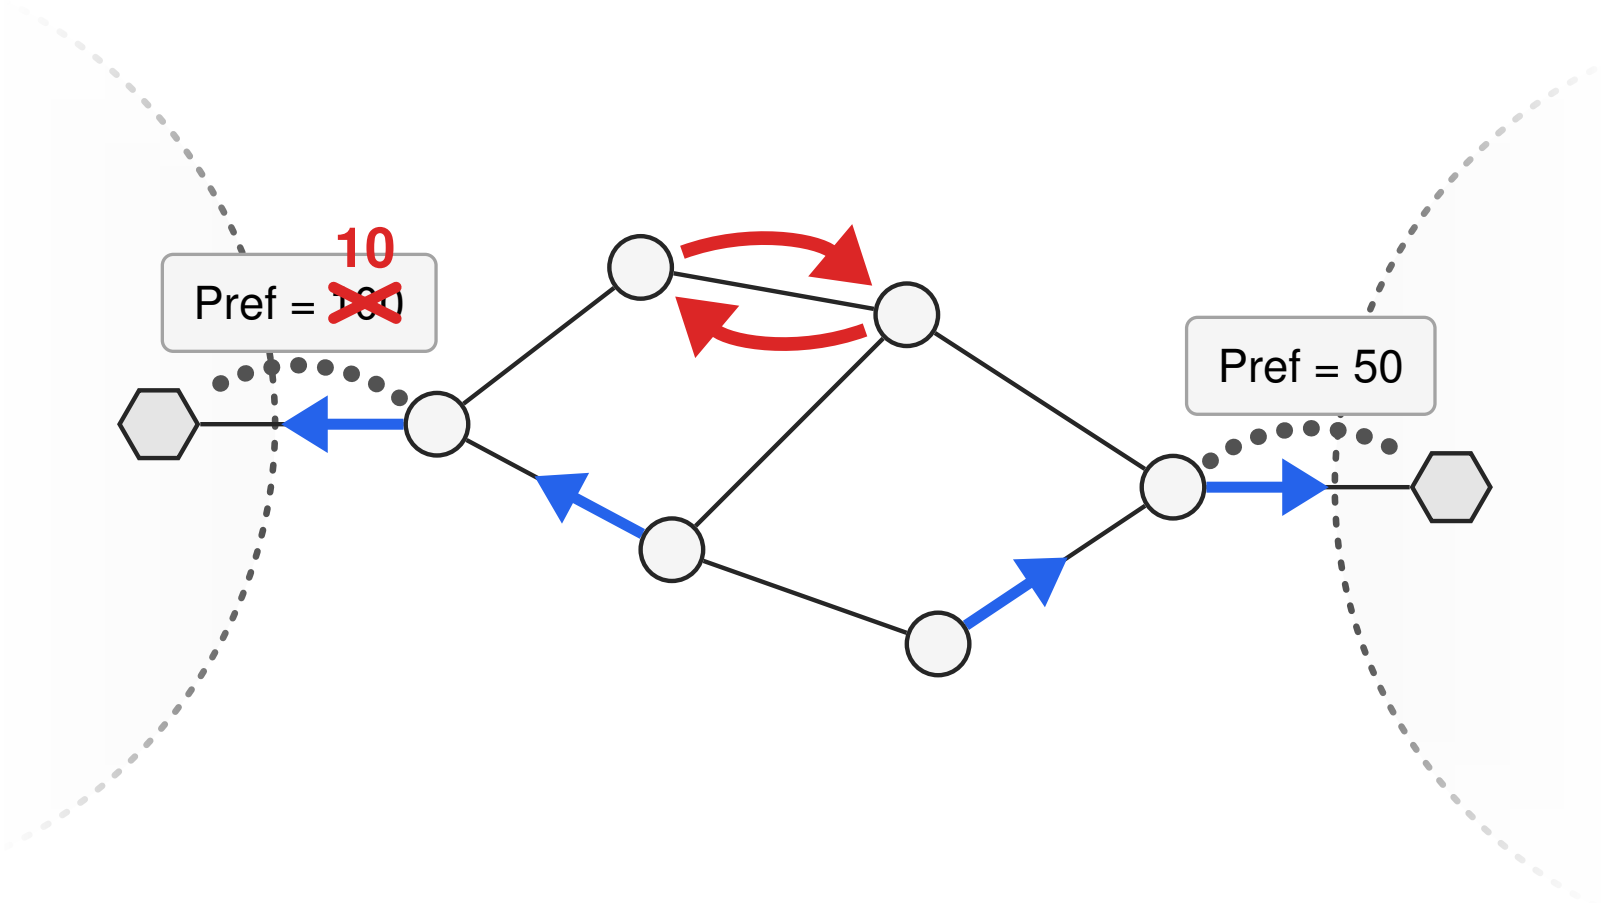

Decouple the control-plane from the data-plane to guarantee reachability.[1]

[1] S. Vissicchio et al. “Improving network agility with seamless BGP reconfigurations”. *IEEE/ACM TNet*. 2012

Decouple the control-plane from the data-plane to guarantee reachability.[1]

1. **Introduce the final configuration**
   while the initial configuration still drives the data-plane.

[1] S. Vissicchio et al. “Improving network agility with seamless BGP reconfigurations”. *IEEE/ACM TNet*. 2012

Decouple the control-plane from the data-plane to guarantee reachability.[1]

1. **Introduce the final configuration**
   while the initial configuration still drives the data-plane.
2. **Wait** until the final configuration is fully converged.

[1] S. Vissicchio et al. “Improving network agility with seamless BGP reconfigurations”. *IEEE/ACM TNet*. 2012

Decouple the control-plane from the data-plane to guarantee reachability.[1]

1. **Introduce the final configuration**
   while the initial configuration still drives the data-plane.
2. **Wait** until the final configuration is fully converged.
3. “Activate” the final configuration **one router at a time**.

[1] S. Vissicchio et al. “Improving network agility with seamless BGP reconfigurations”. *IEEE/ACM TNet*. 2012

Decouple the control-plane from the data-plane to guarantee reachability.[1]

1. **Introduce the final configuration**
   while the initial configuration still drives the data-plane.
2. **Wait** until the final configuration is fully converged.
3. "Activate" the final configuration **one router at a time**.

→ safe but resource hungry.

[1] S. Vissicchio et al. "Improving network agility with seamless BGP reconfigurations". *IEEE/ACM TNet*. 2012

Naive idea: use static routes to decouple the control-plane from the data-plane.

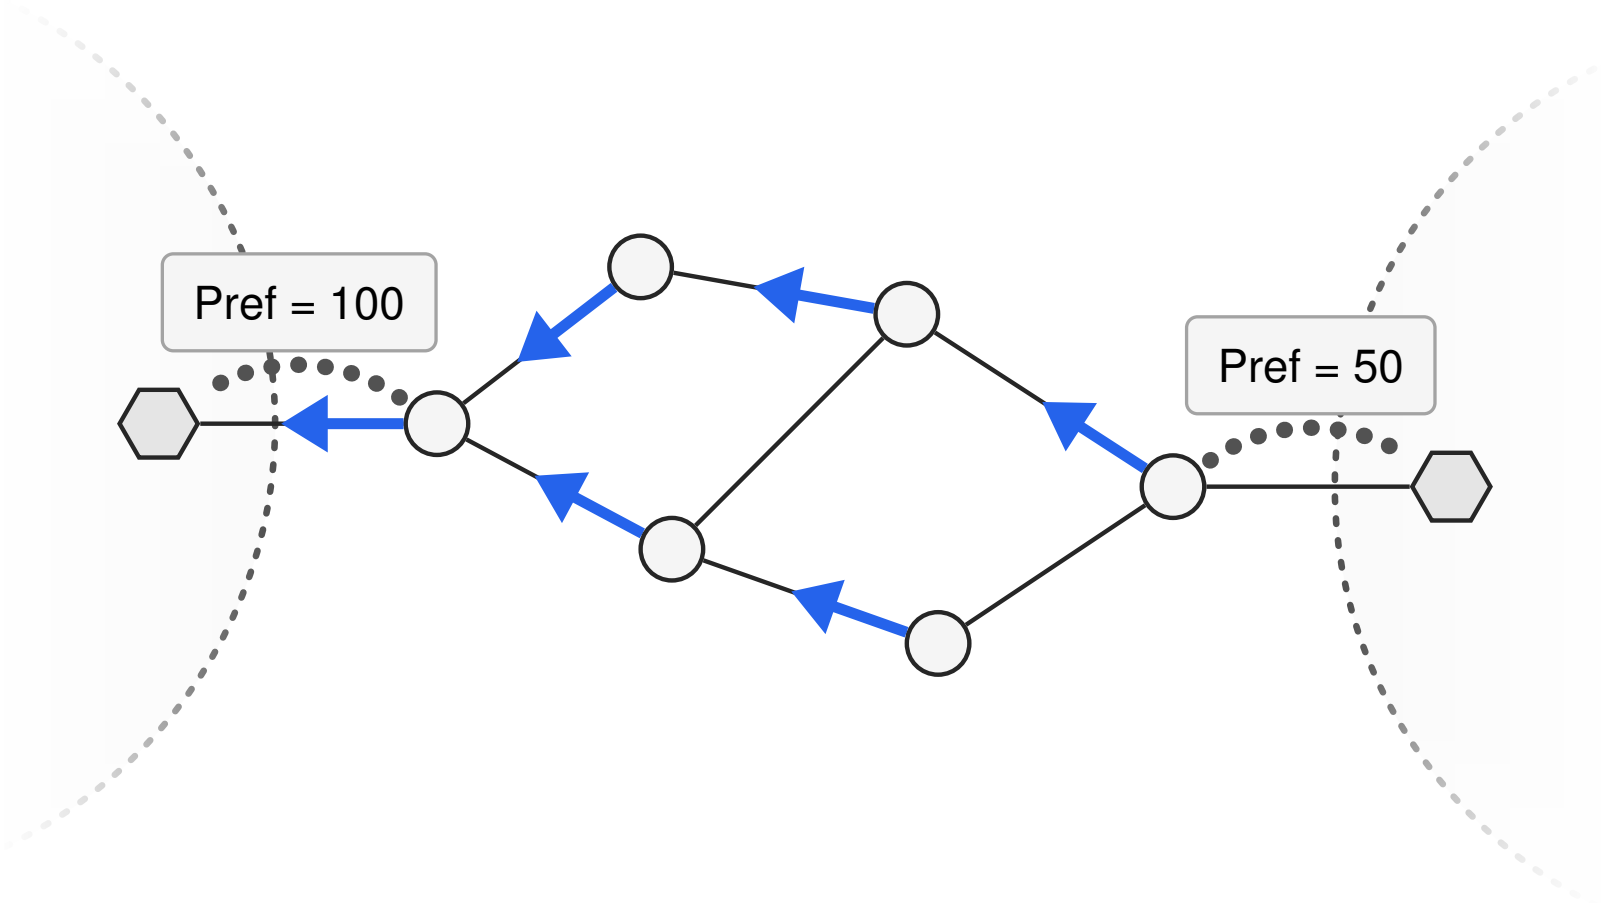

Naive idea: use static routes to decouple the control-plane from the data-plane.

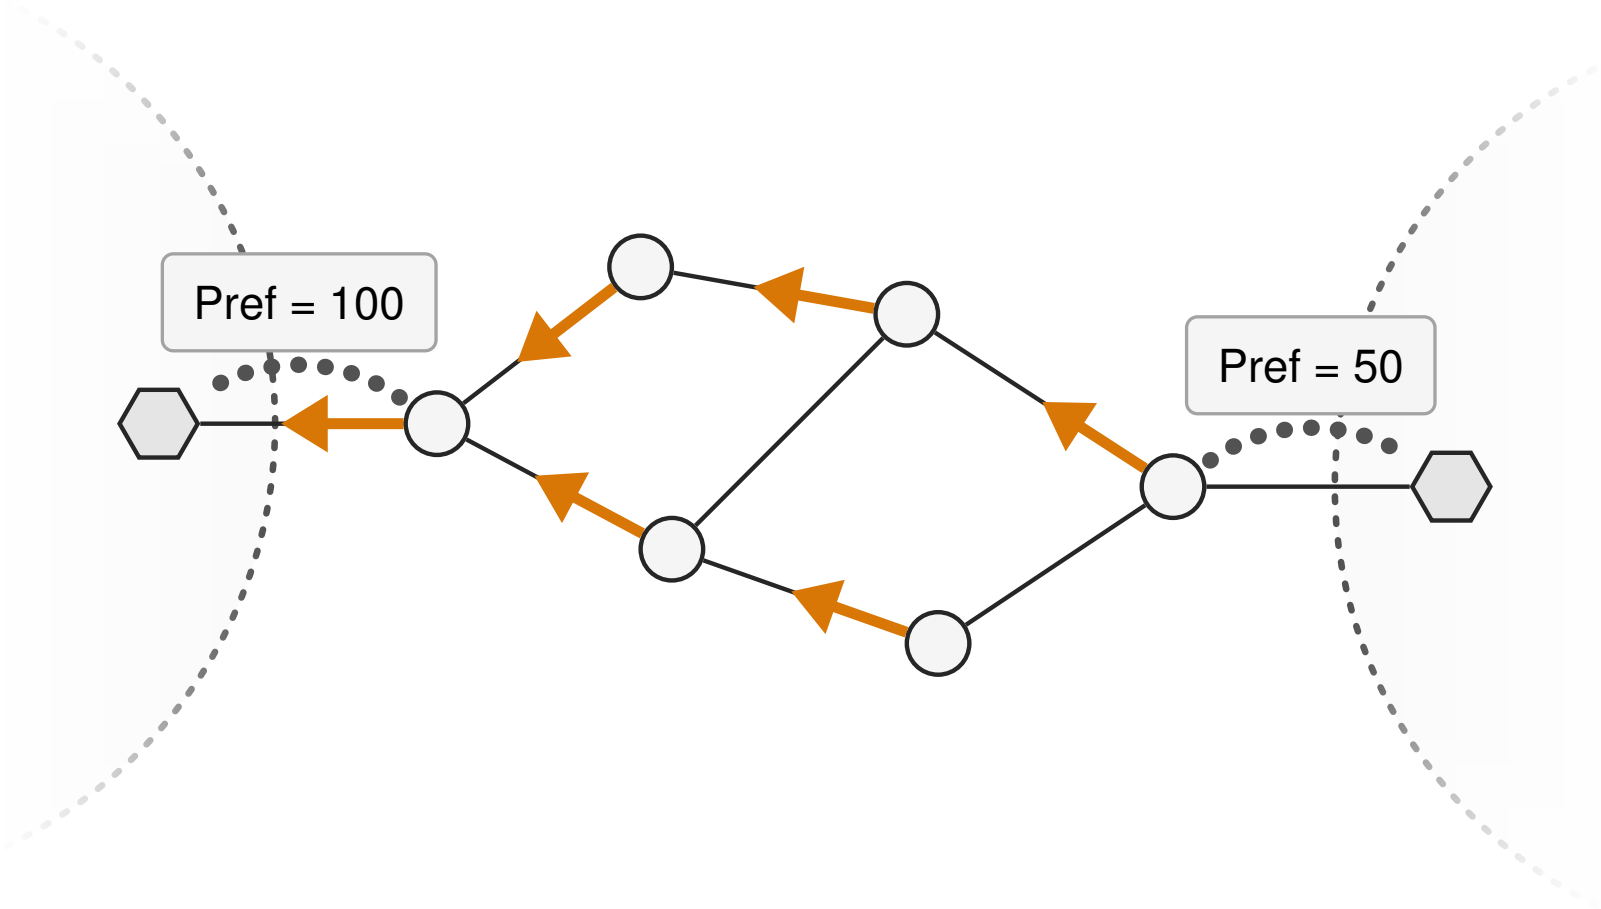

Naive idea: use static routes to decouple the control-plane from the data-plane.

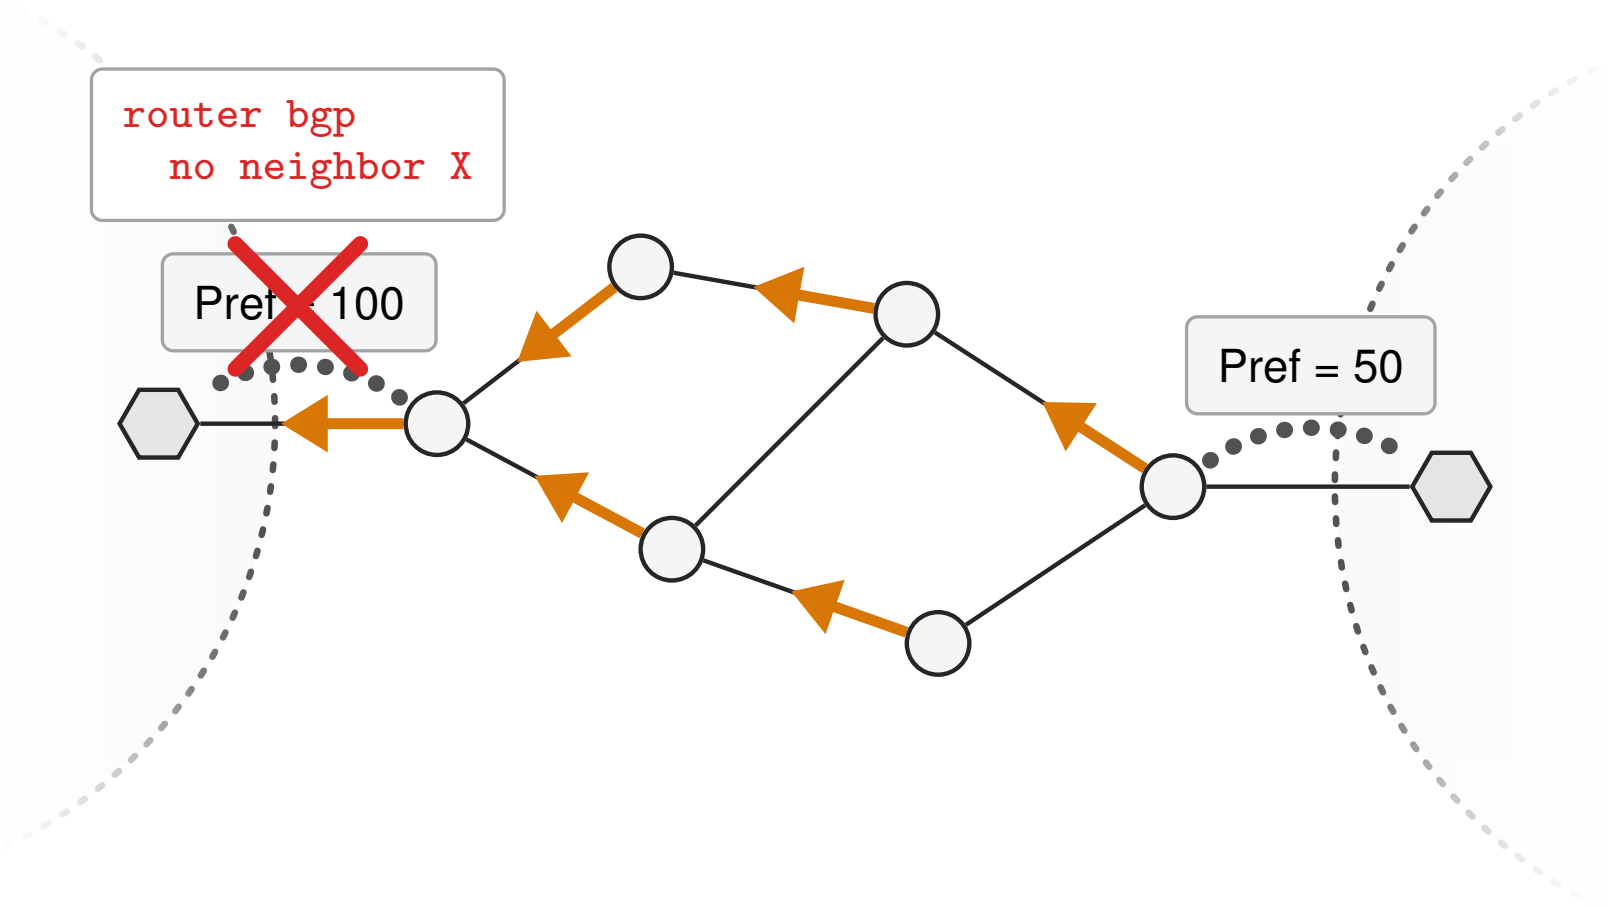

Naive idea: use static routes to decouple the control-plane from the data-plane.

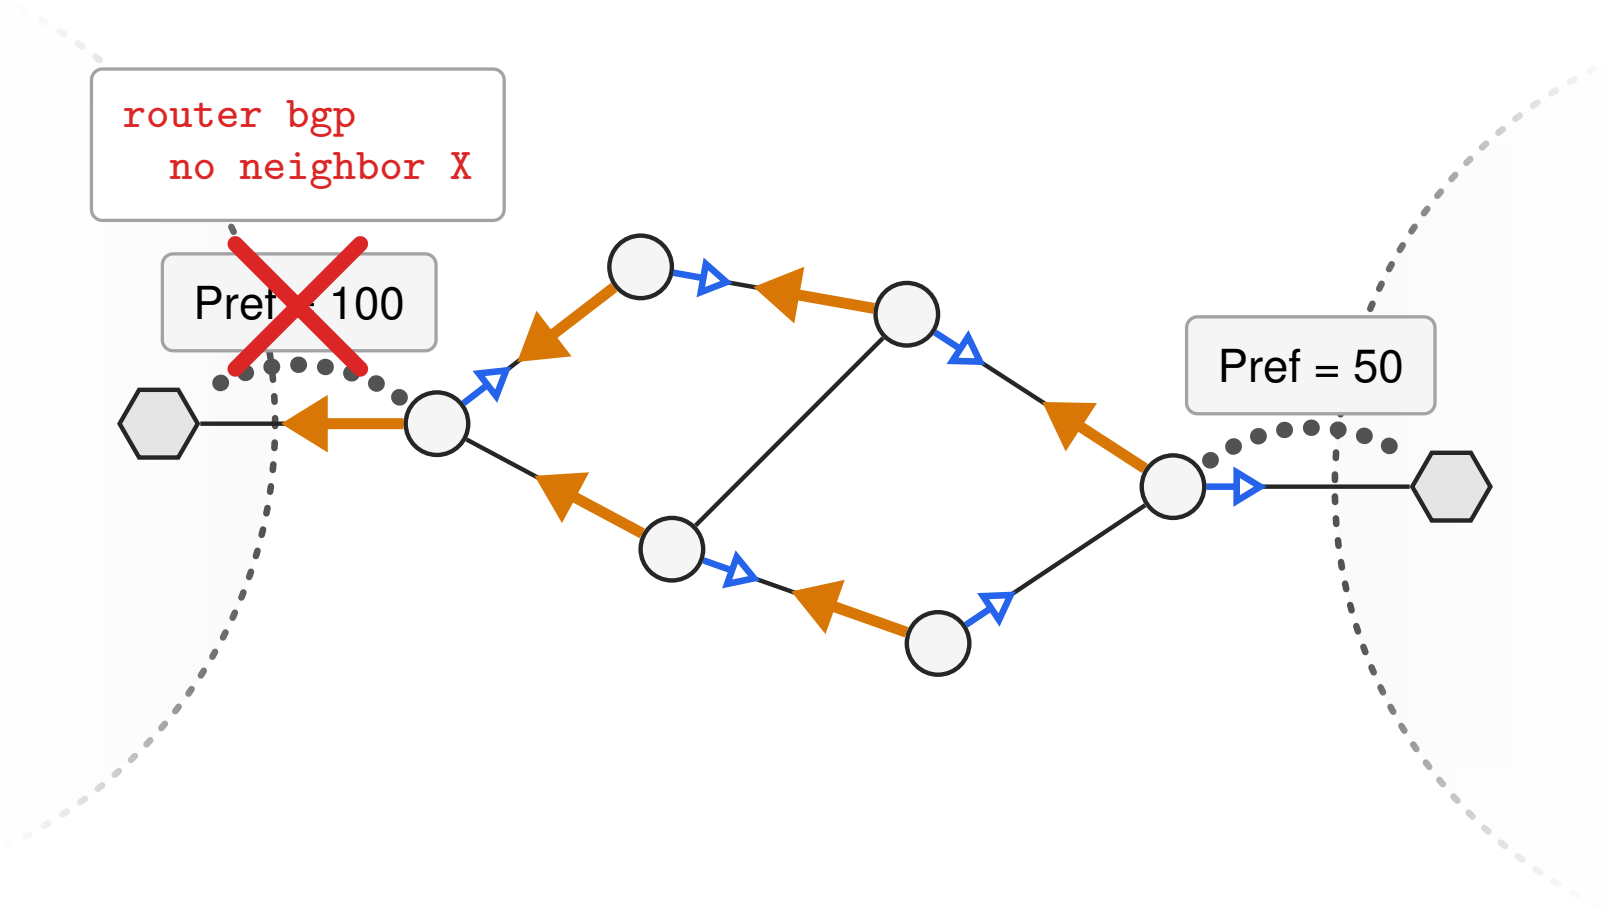

Naive idea: use static routes to decouple the control-plane from the data-plane.

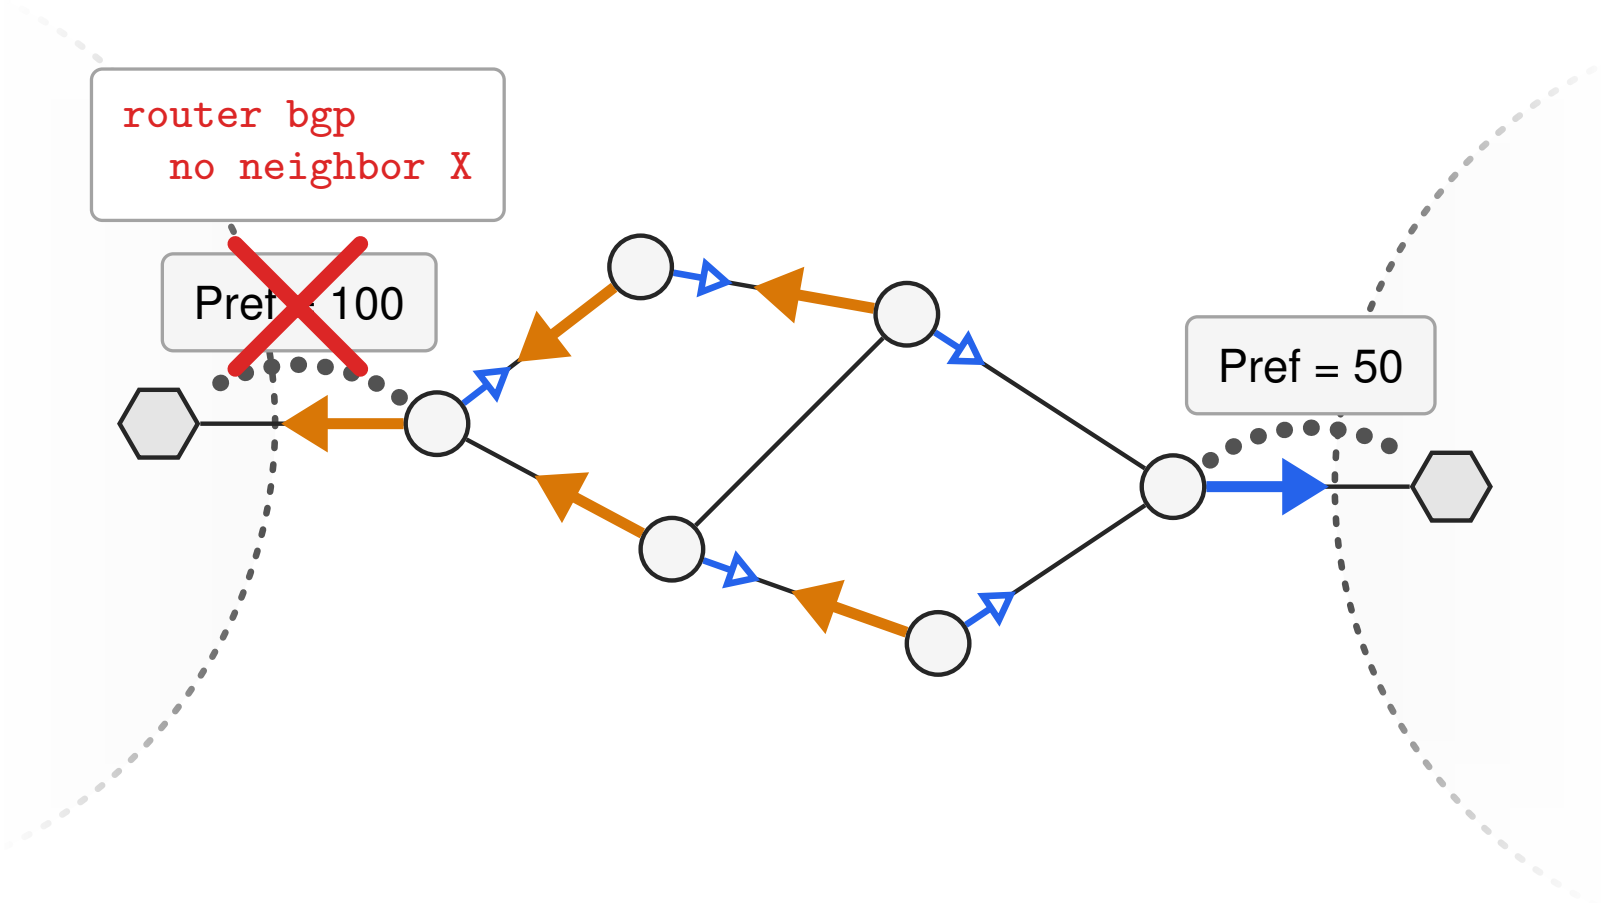

Naive idea: use static routes to decouple the control-plane from the data-plane.

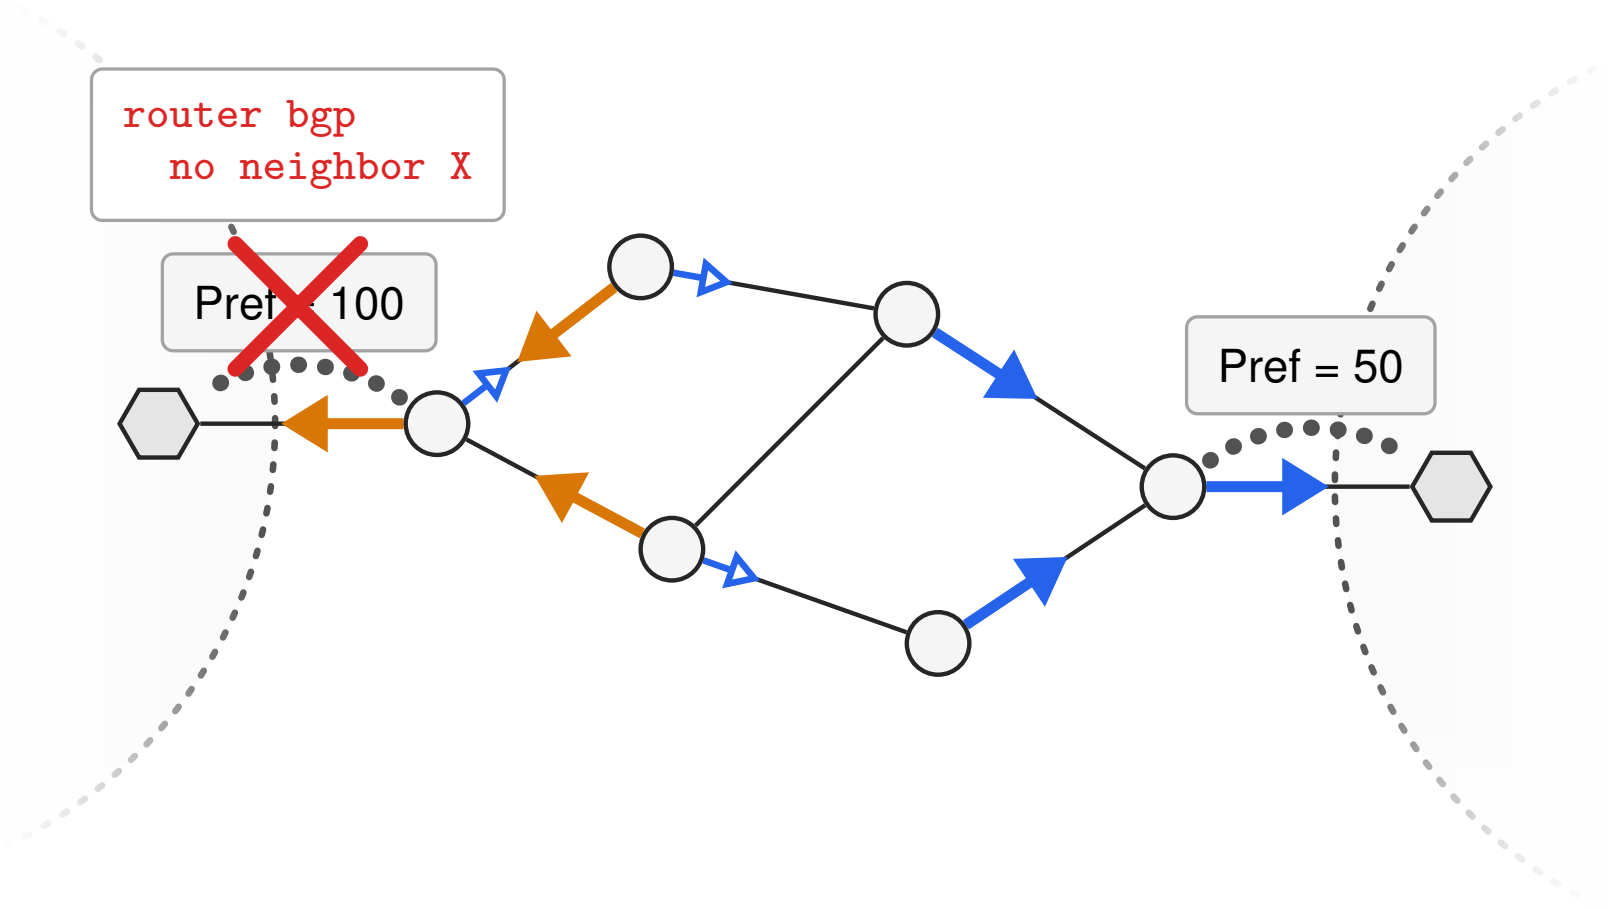

Naive idea: use static routes to decouple the control-plane from the data-plane.

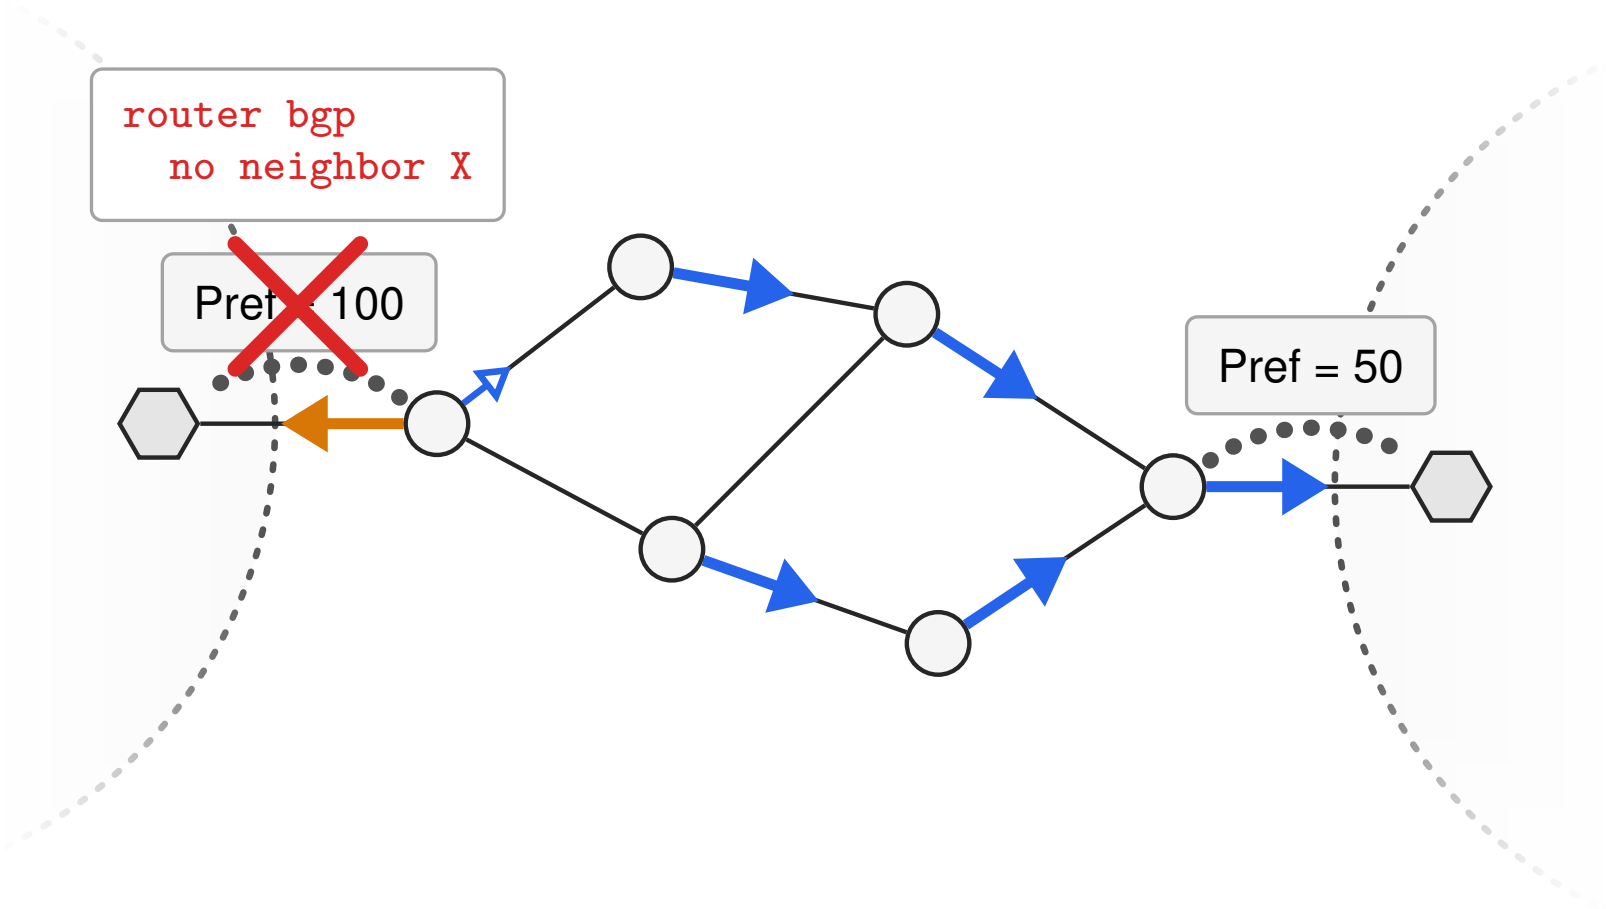

However, static routes cannot react to failures.

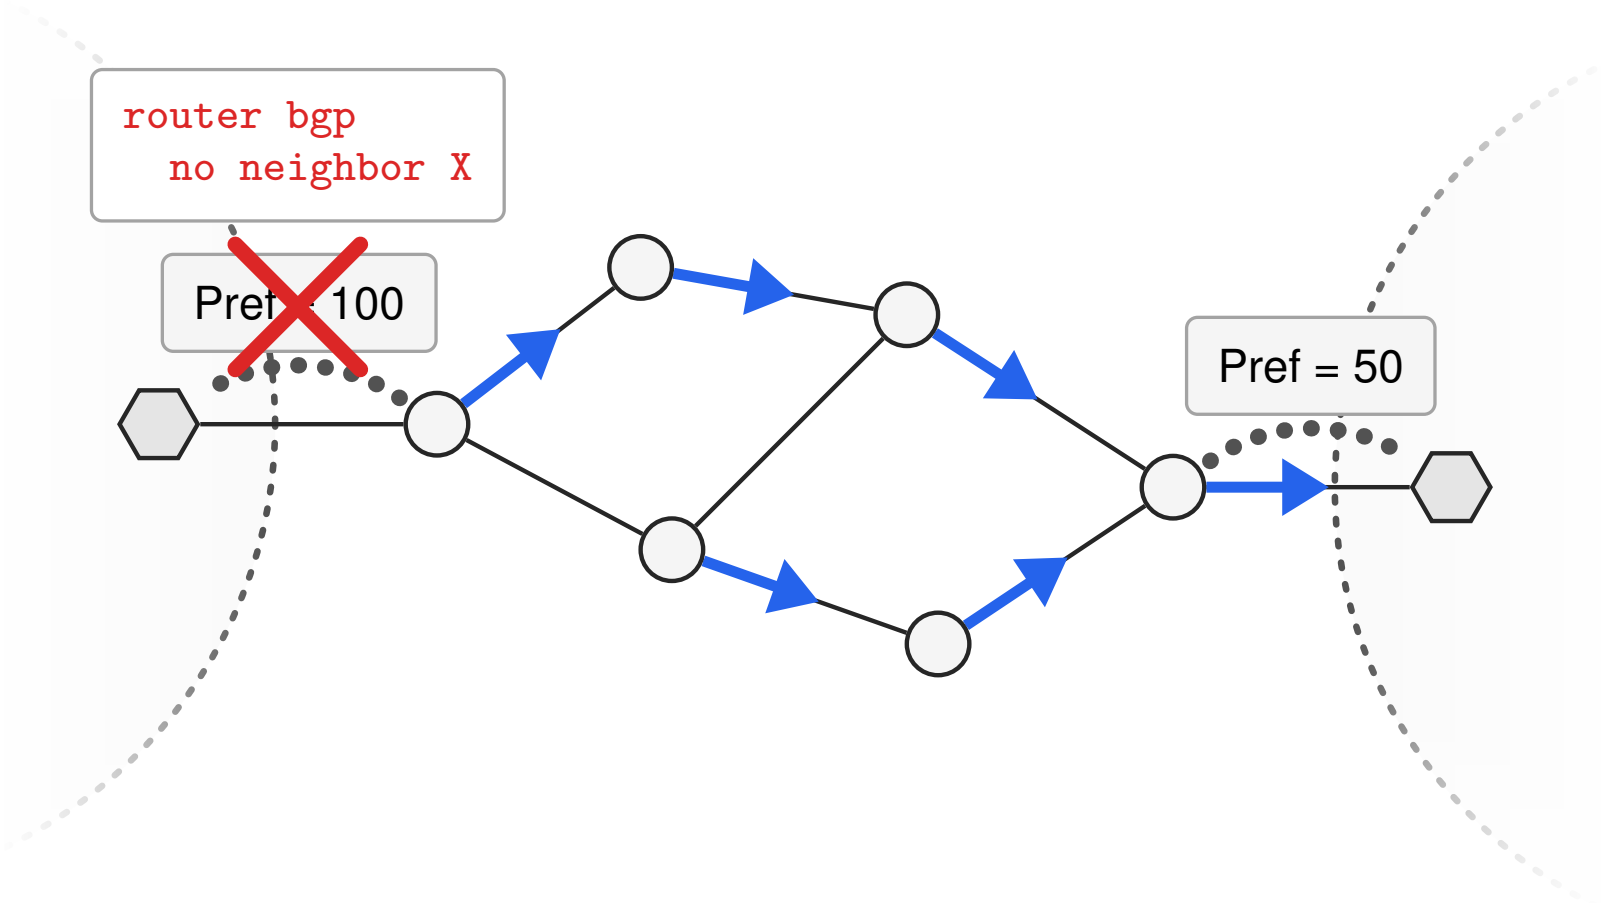

How to reconfigure BGP

How to reconfigure BGP

- **efficiently** while being **reactive**,

How to reconfigure BGP

- **efficiently** while being **reactive**,
- **safely**,

How to reconfigure BGP

- **efficiently** while being **reactive**,
- **safely**, and
- **today**?

How to reconfigure BGP

- **efficiently** while being **reactive**,
- **safely**, and
- **today**?

We **keep** the control- and data-plane **coupled**.

How to reconfigure BGP

- **efficiently** while being **reactive**, → We **keep** the control- and data-plane **coupled**.
- **safely**, and → We **control** the BGP convergence.
- **today**?

How to reconfigure BGP

- **efficiently** while being **reactive**, → We **keep** the control- and data-plane **coupled**.
- **safely**, and → We **control** the BGP convergence.
- **today**? → We only change route preferences.

We control the convergence
of each router **individually**.

We identify three sufficient properties:

We control the convergence
of each router **individually**.

We identify three sufficient properties:

1. The router selects the initial route.

We control the convergence
of each router **individually**.

We identify three sufficient properties:

1. The router selects the initial route.
2. The router receives the final route.

We control the convergence
of each router **individually**.

We identify three sufficient properties:

1. The router selects the initial route.
2. The router receives the final route.
3. Updating the router does not cause any other router to update.

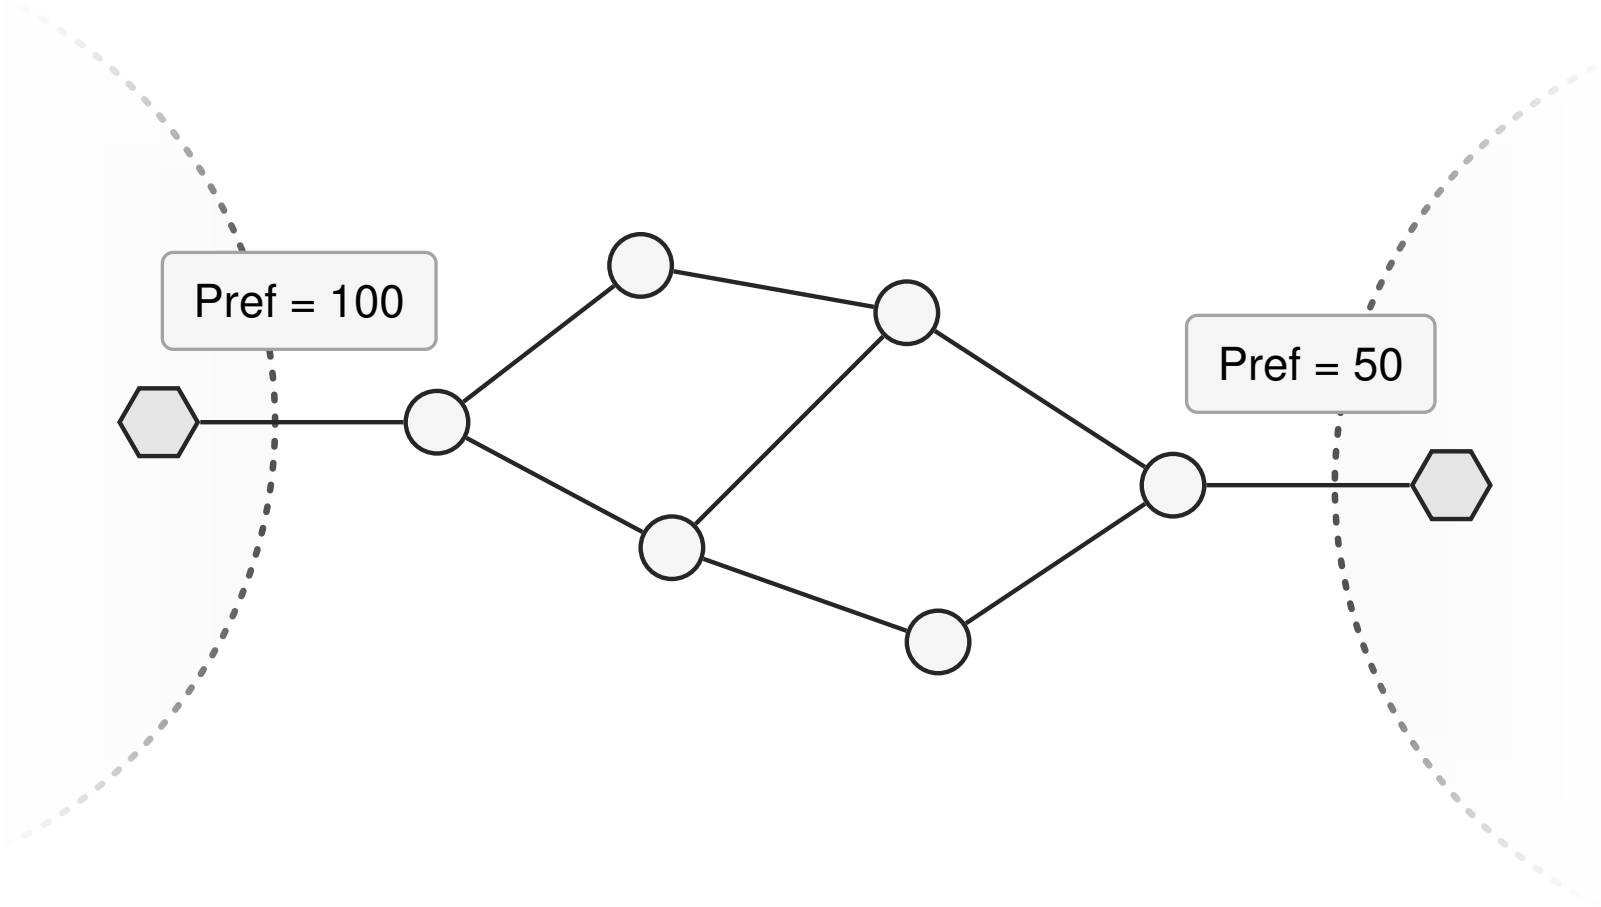
Pref = 100
Pref = 50

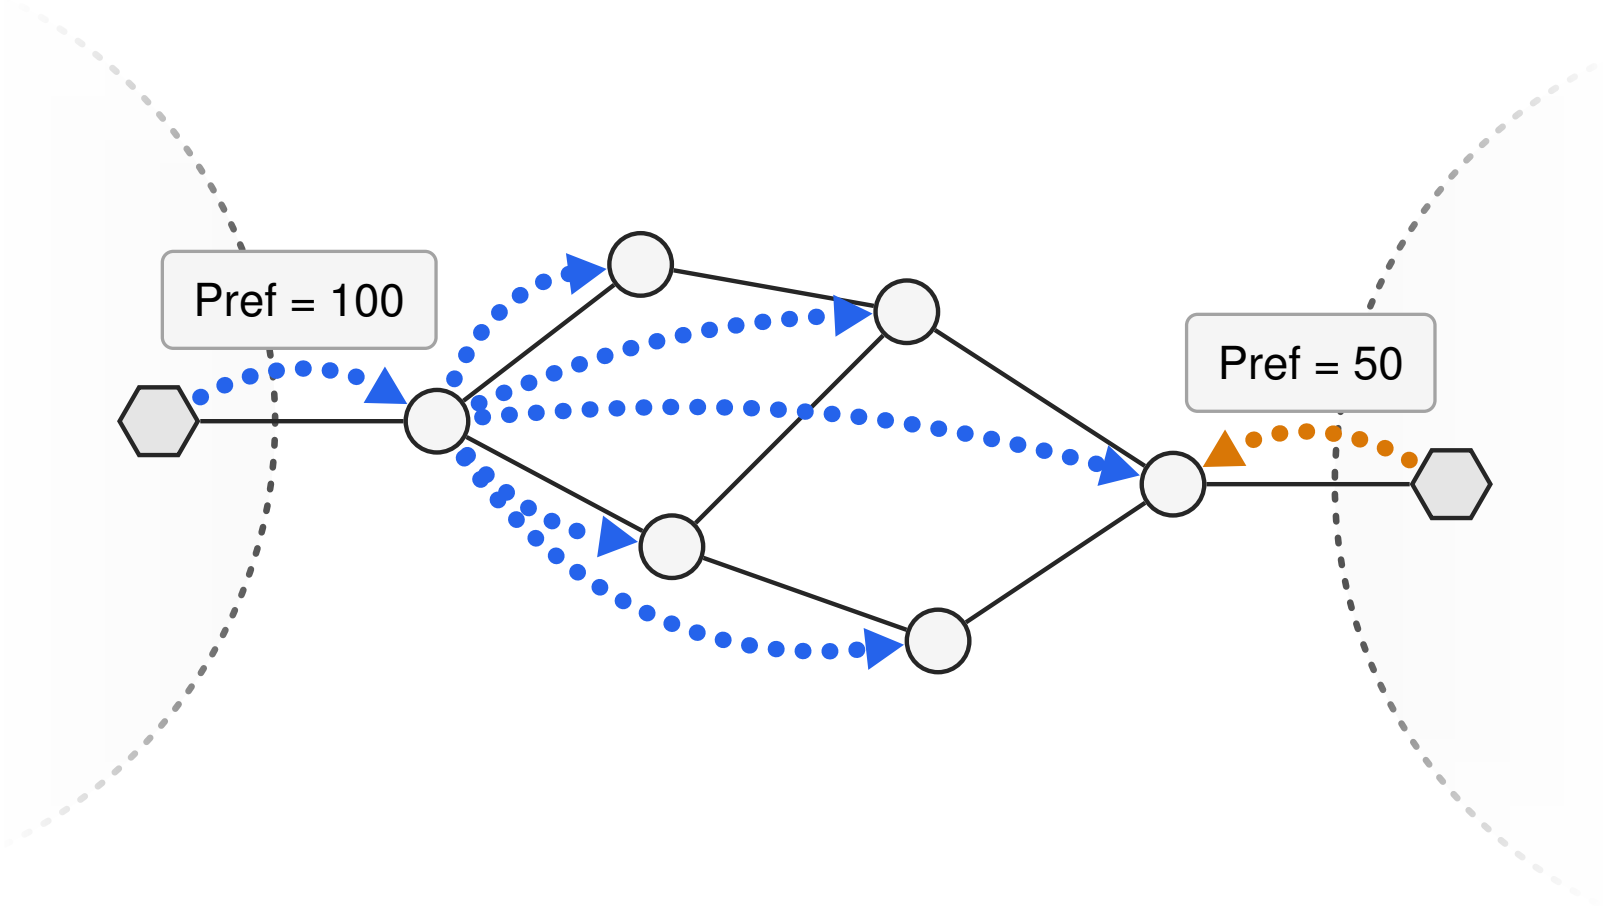
Pref = 100
Pref = 50

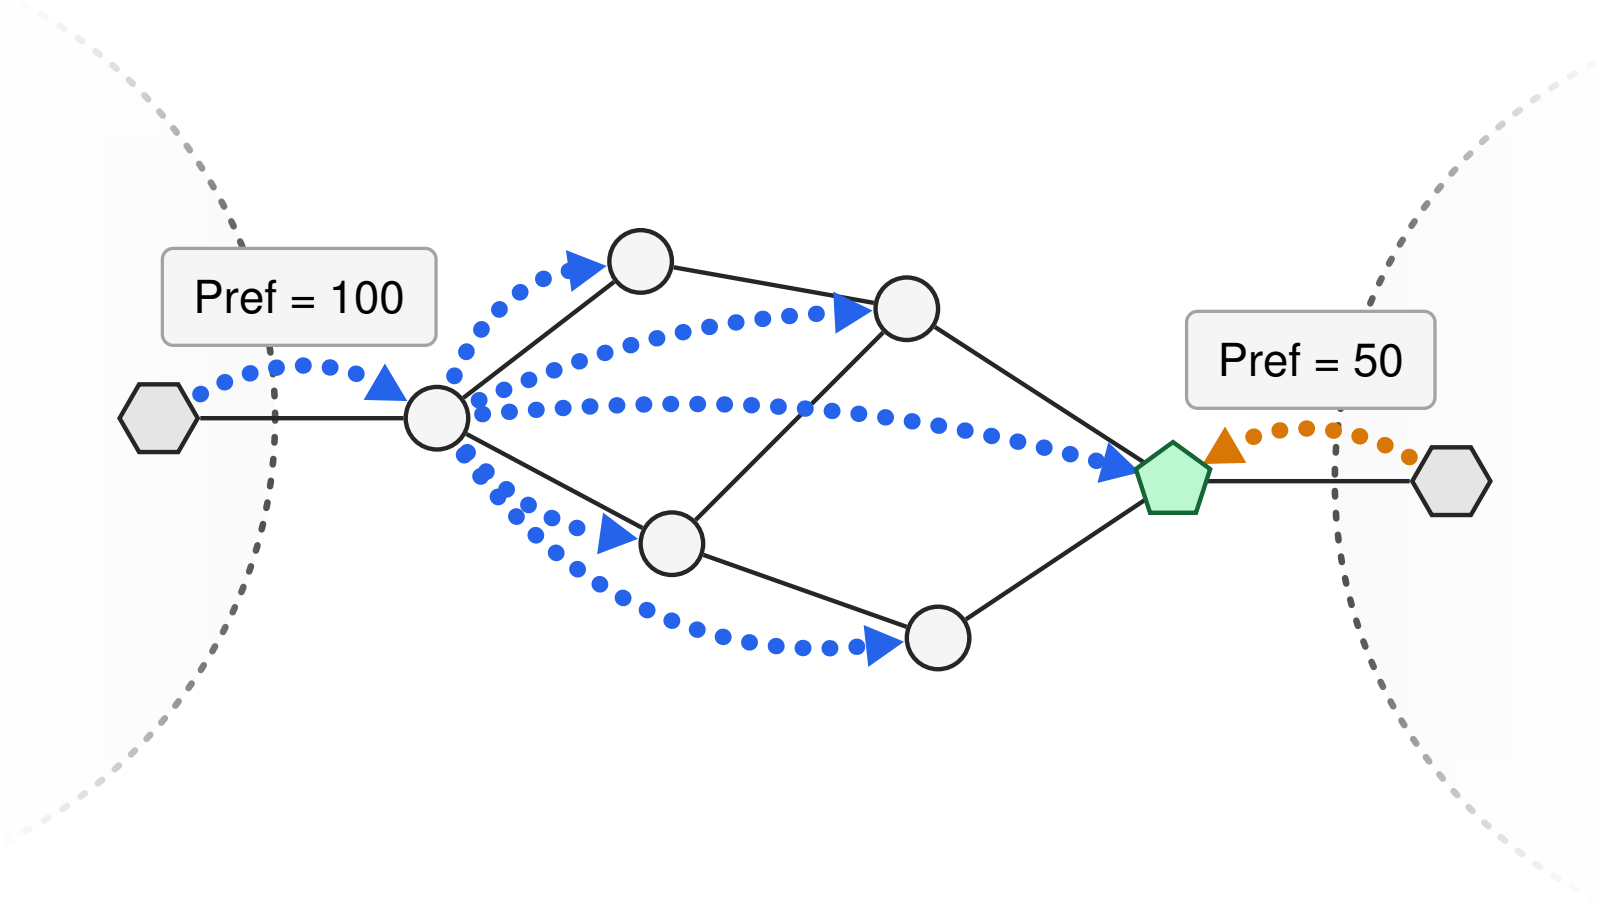
Pref = 100
Pref = 50

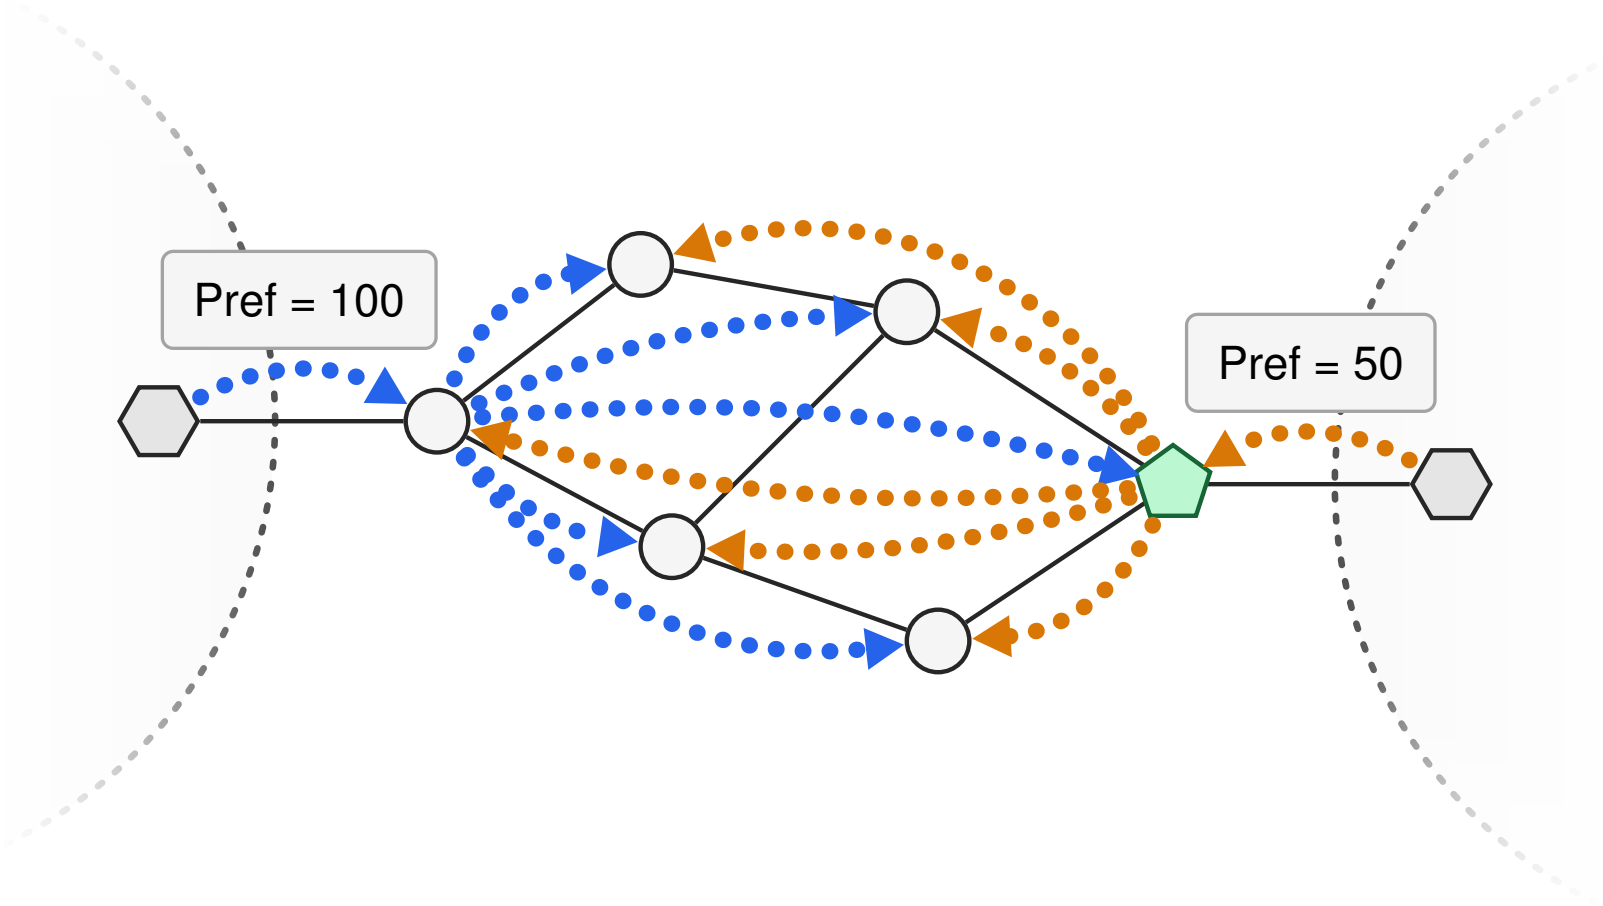
Pref = 100
Pref = 50

How can we deal with limited route visibility?

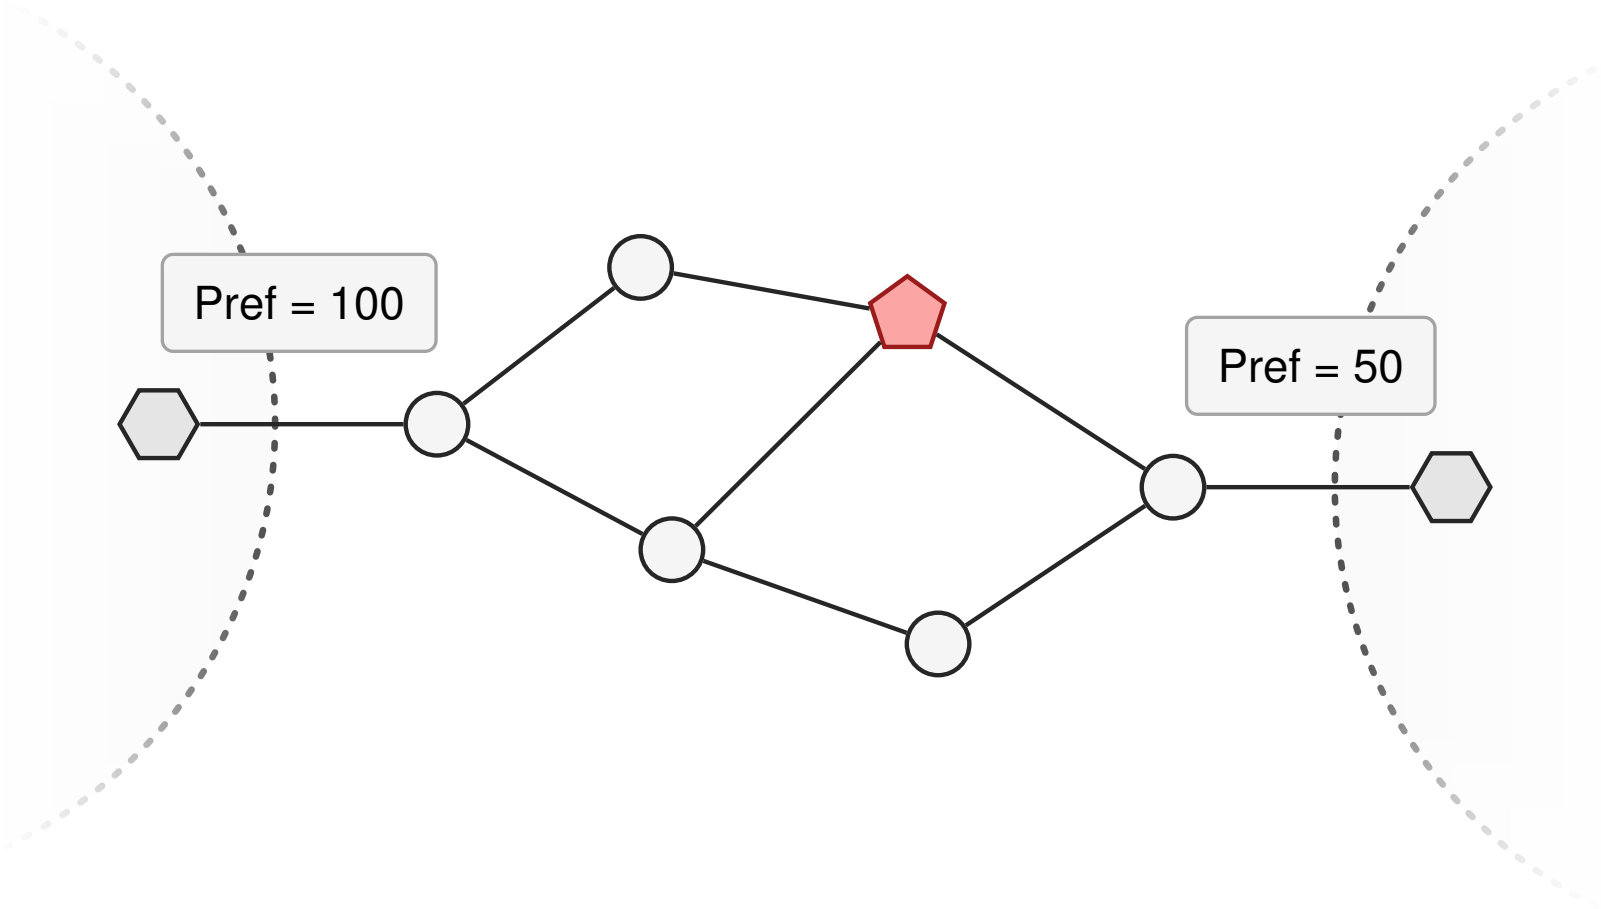

How can we deal with limited route visibility?

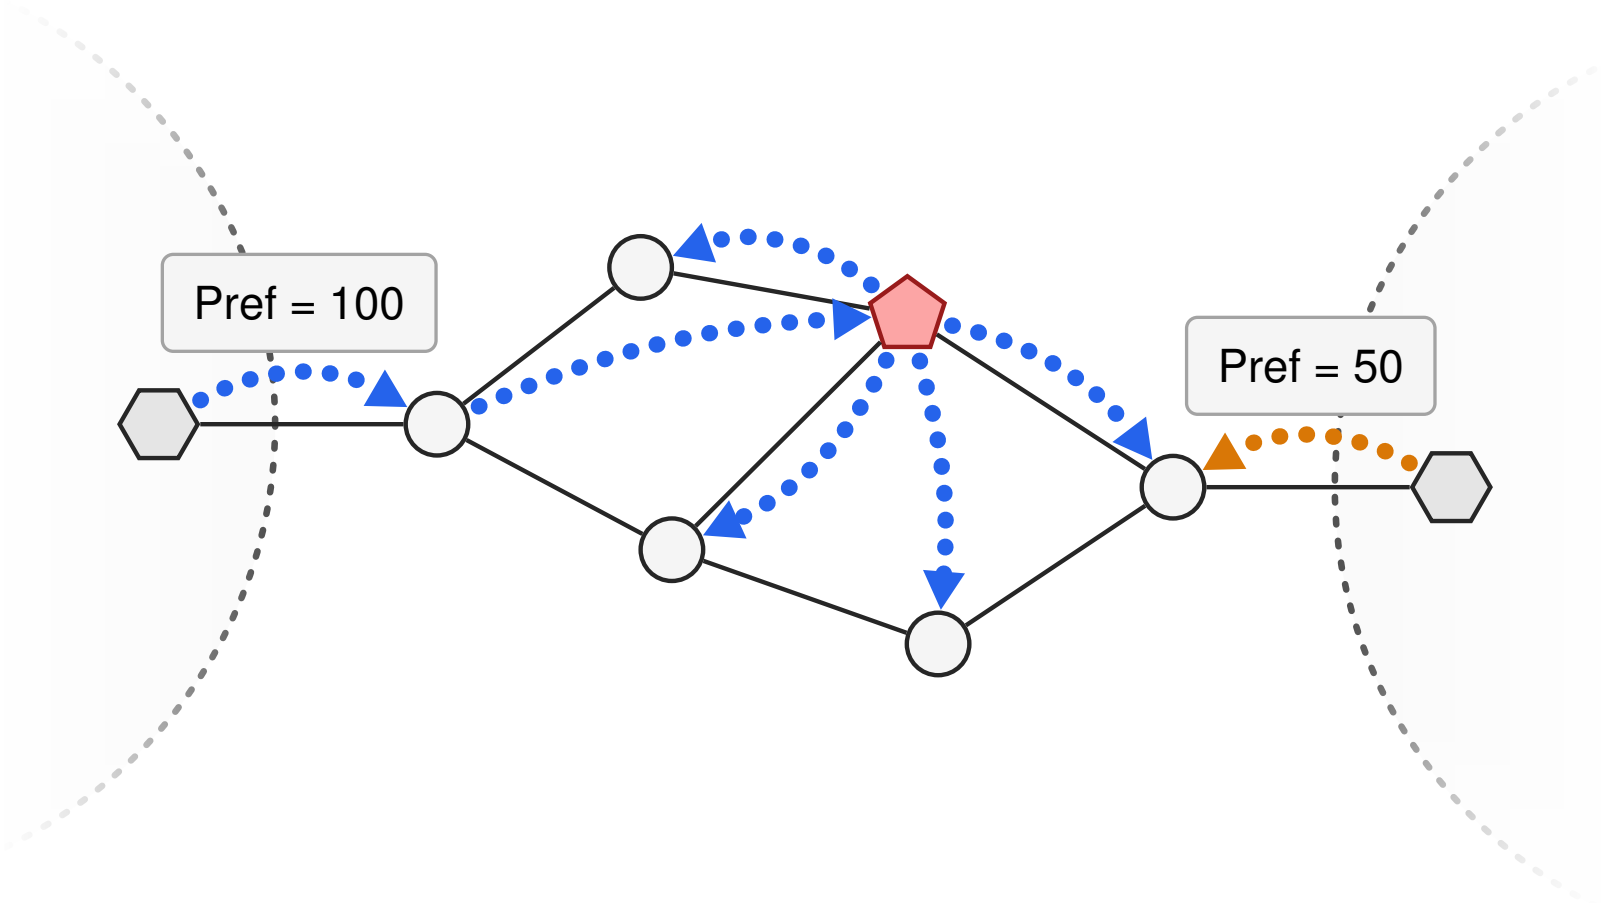

How can we deal with limited route visibility?

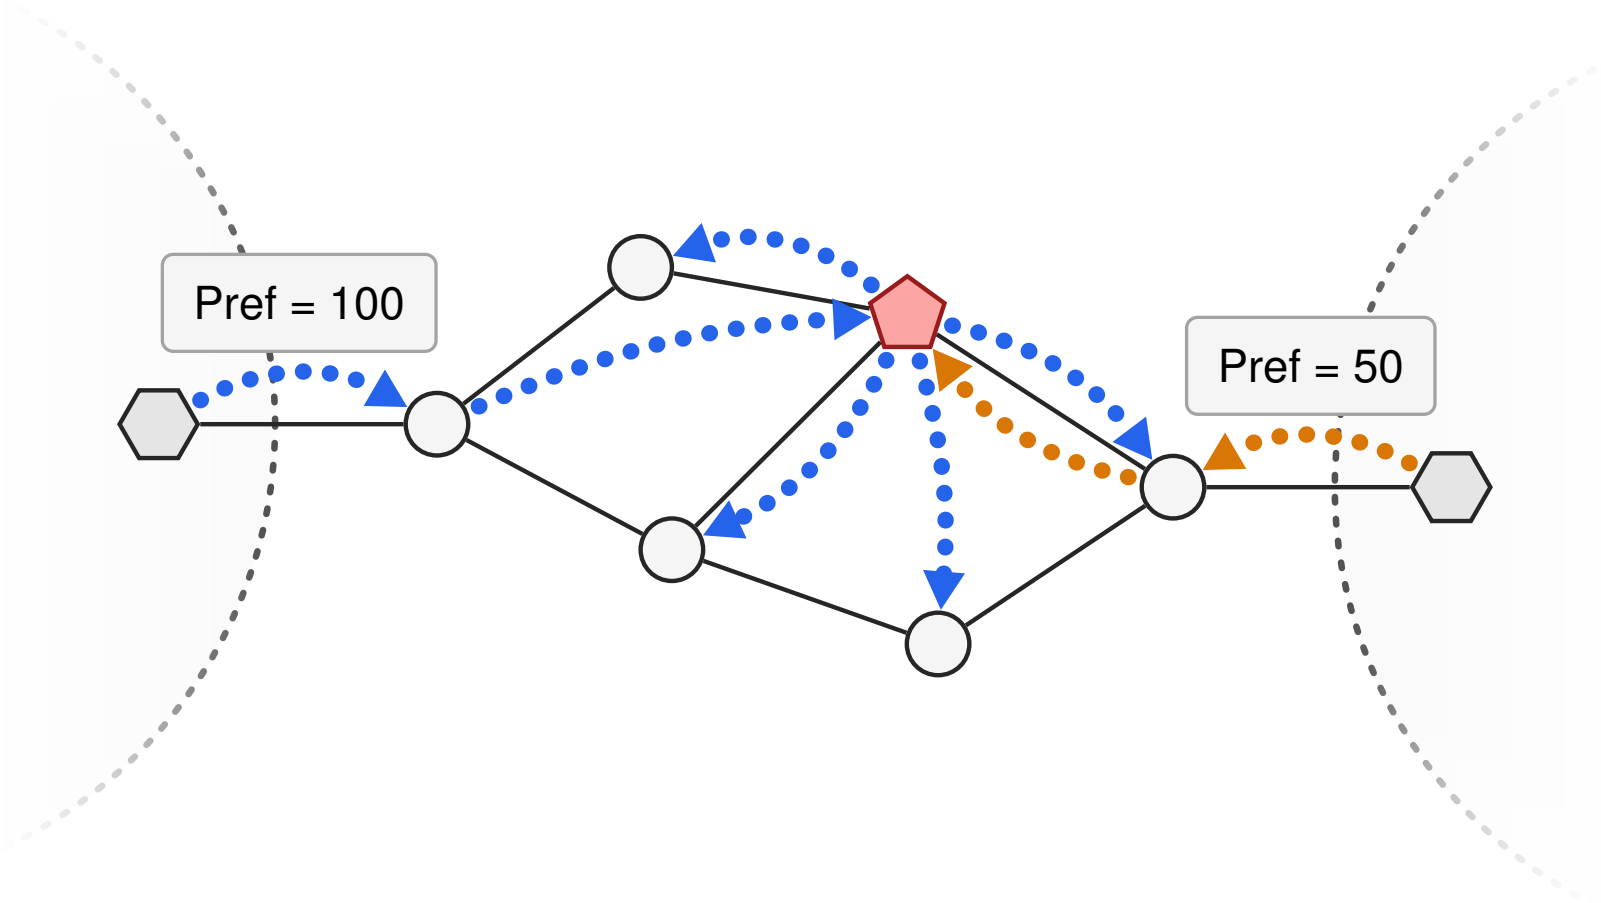

We deal with limited route visibility by
temporarily adapting the propagation graph.

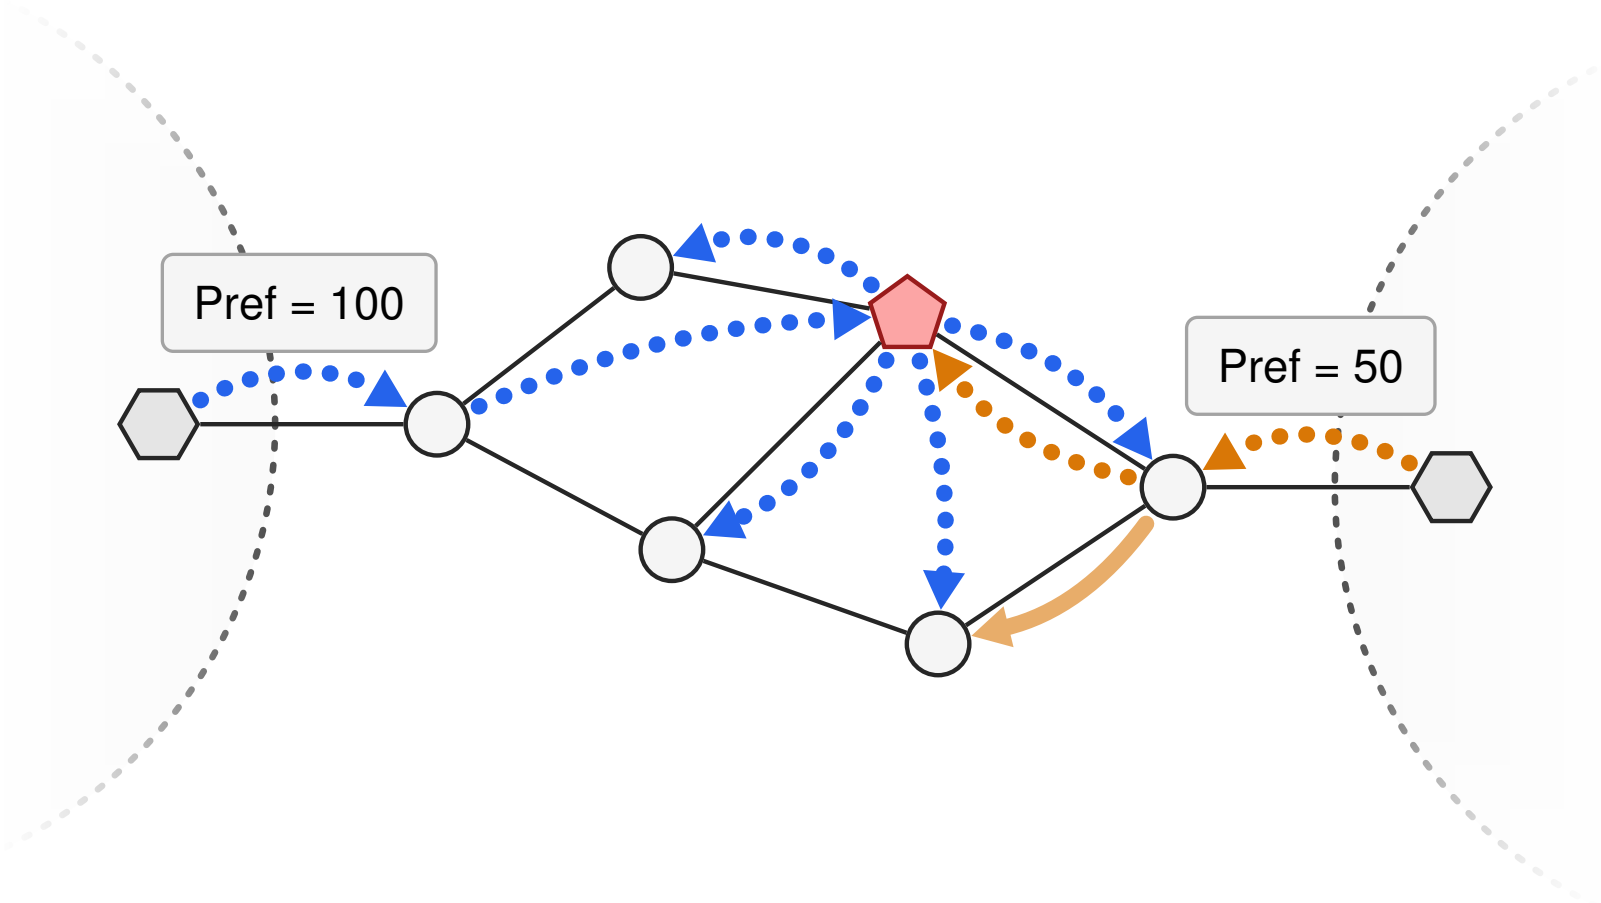

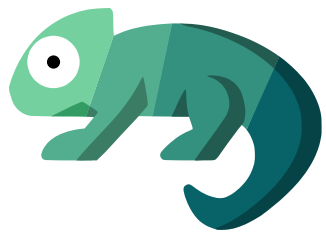

initial and final
configurations

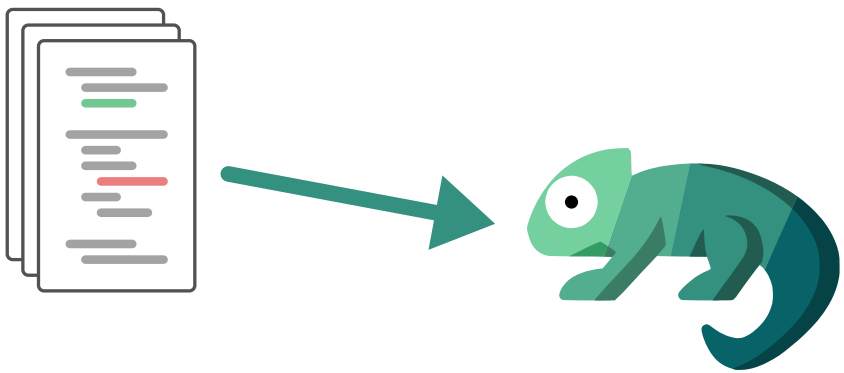

initial and final configurations

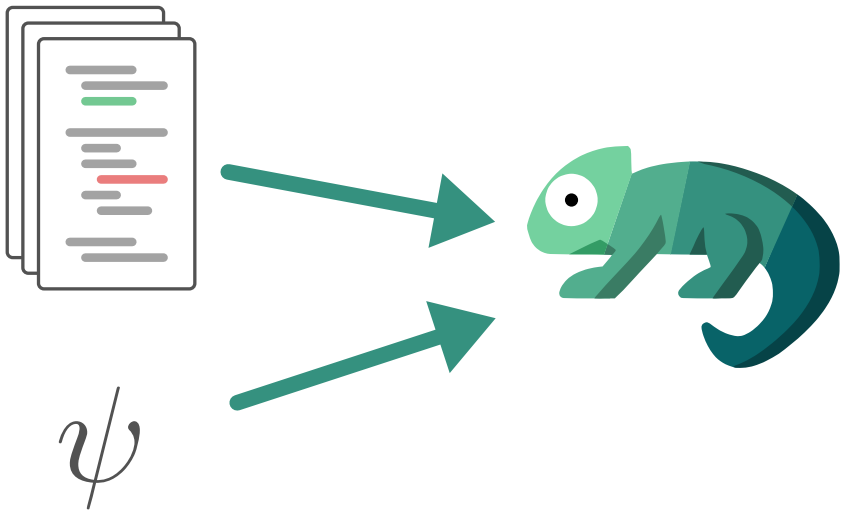

ψ

specification

initial and final configurations

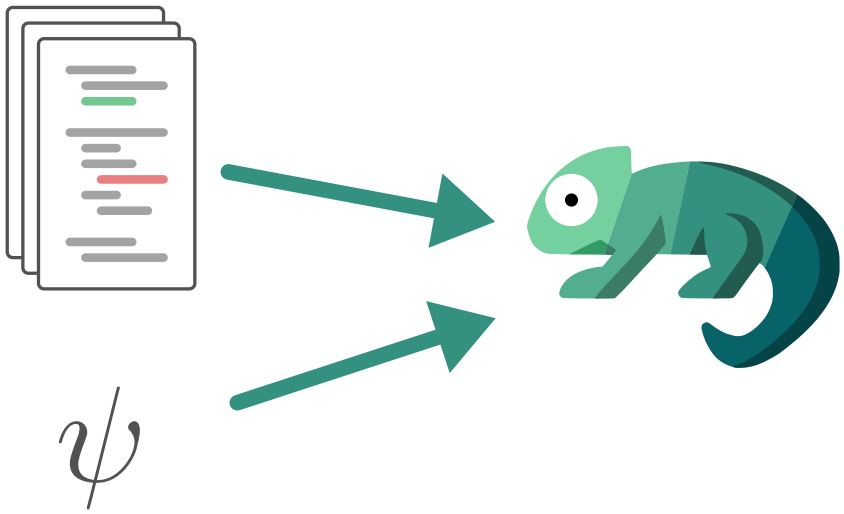

ψ

specification

- Follow our three rules

initial and final configurations

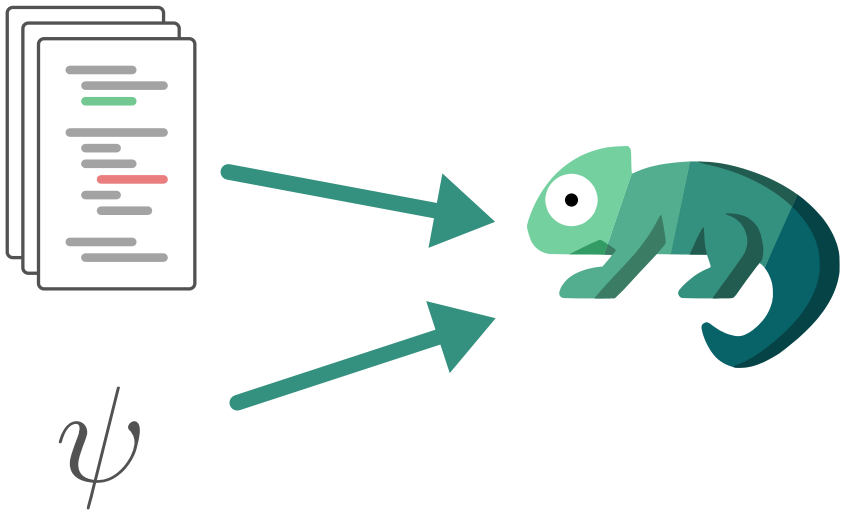

ψ

specification

- Follow our three rules
- Satisfy the specification

initial and final configurations

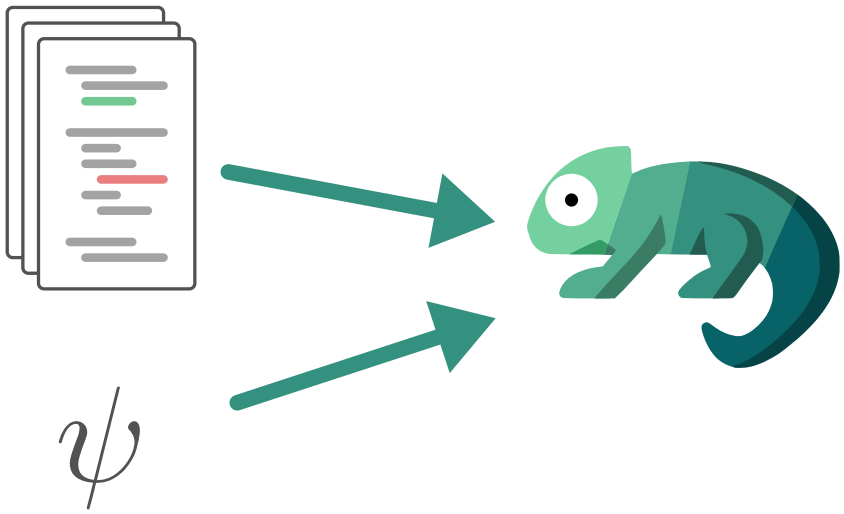

ψ

specification

- Follow our three rules
- Satisfy the specification
- Minimize reconfiguration time

initial and final configurations

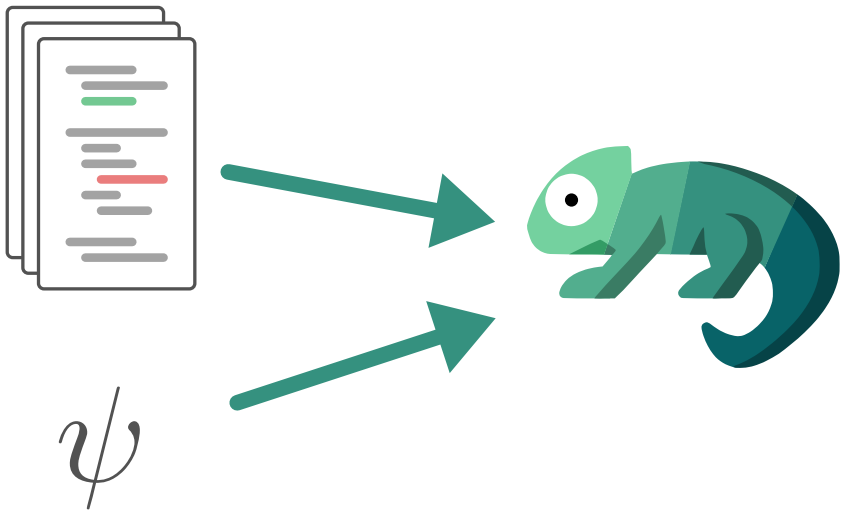

ψ

specification

- Follow our three rules
- Satisfy the specification
- Minimize reconfiguration time
- Minimize temp. sessions

initial and final configurations

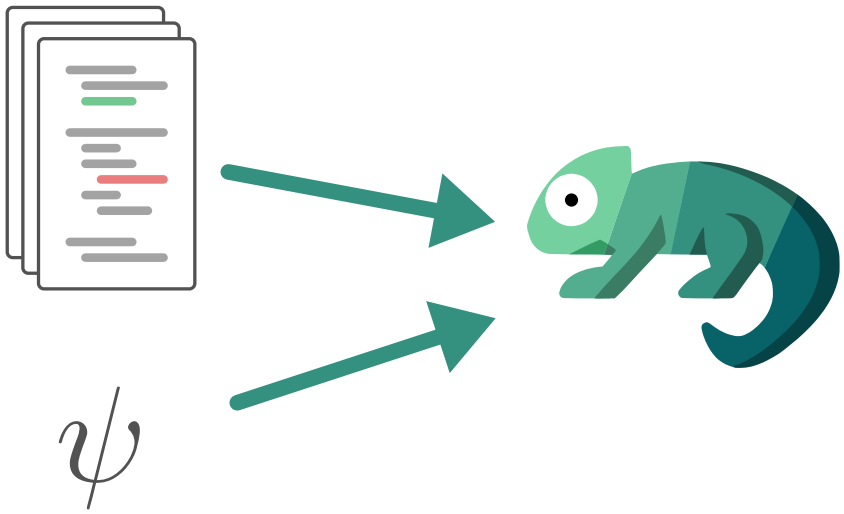

ψ

specification

- Follow our three rules
- Satisfy the specification
- Minimize reconfiguration time
- Minimize temp. sessions

→ **Optimization problem**

Chameleon issues temporary BGP commands.

initial and final configurations

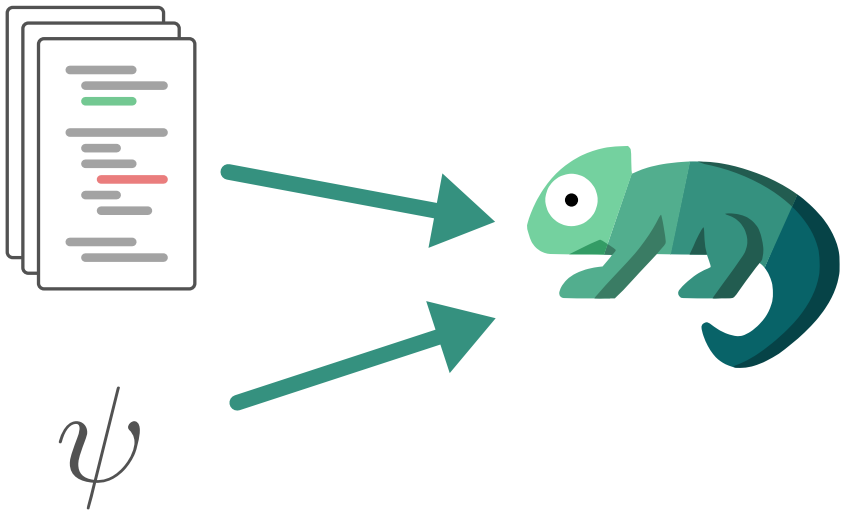

ψ

specification

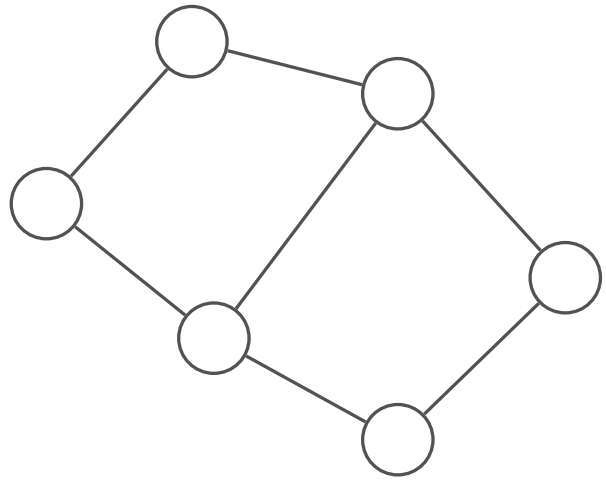

Chameleon issues temporary BGP commands.

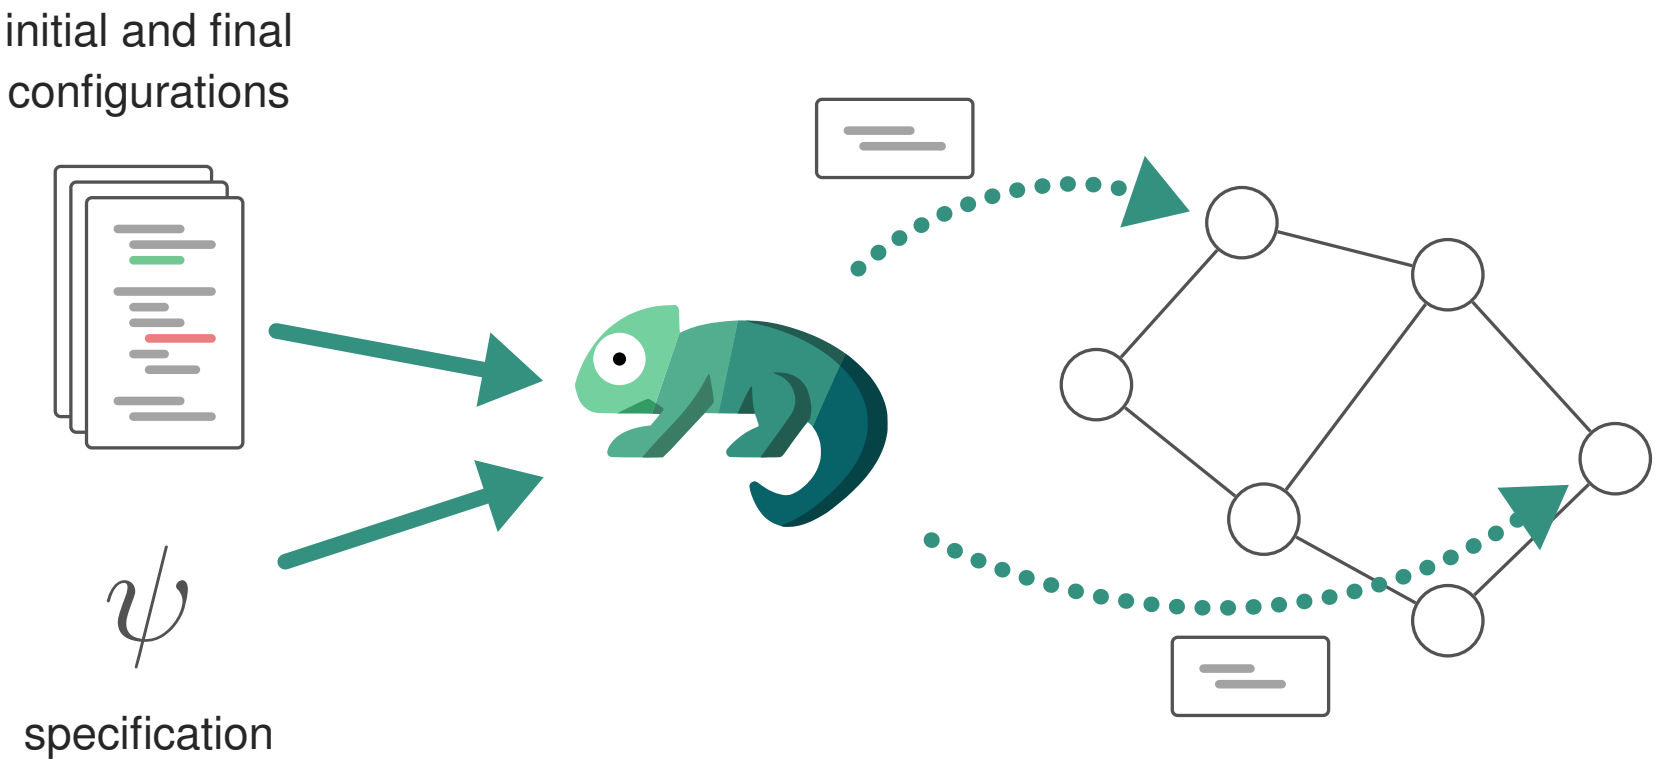

Chameleon issues temporary BGP commands and monitors the convergence progress.

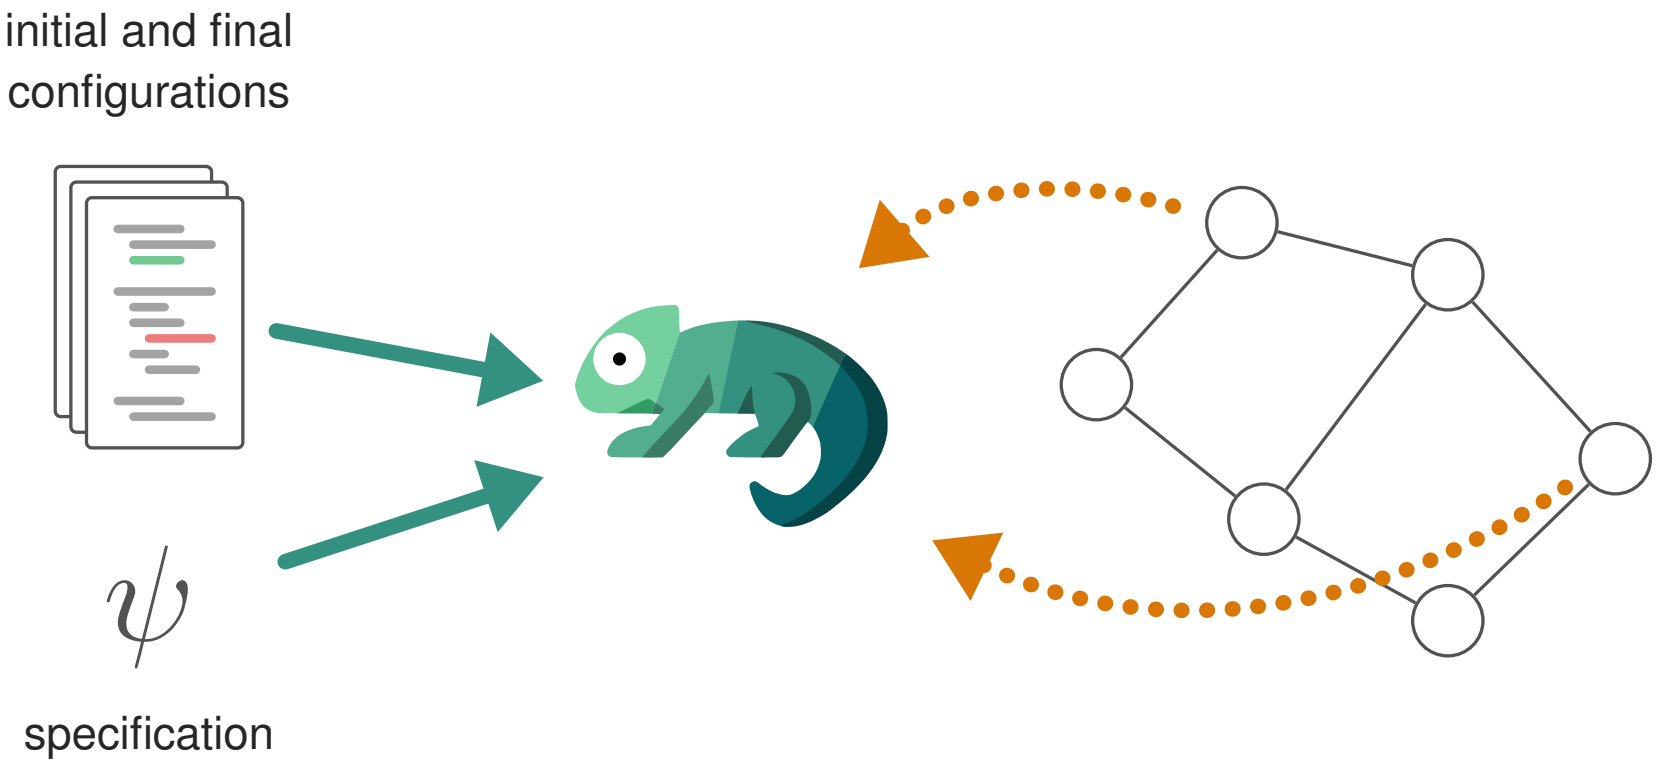

Chameleon issues temporary BGP commands and monitors the convergence progress.

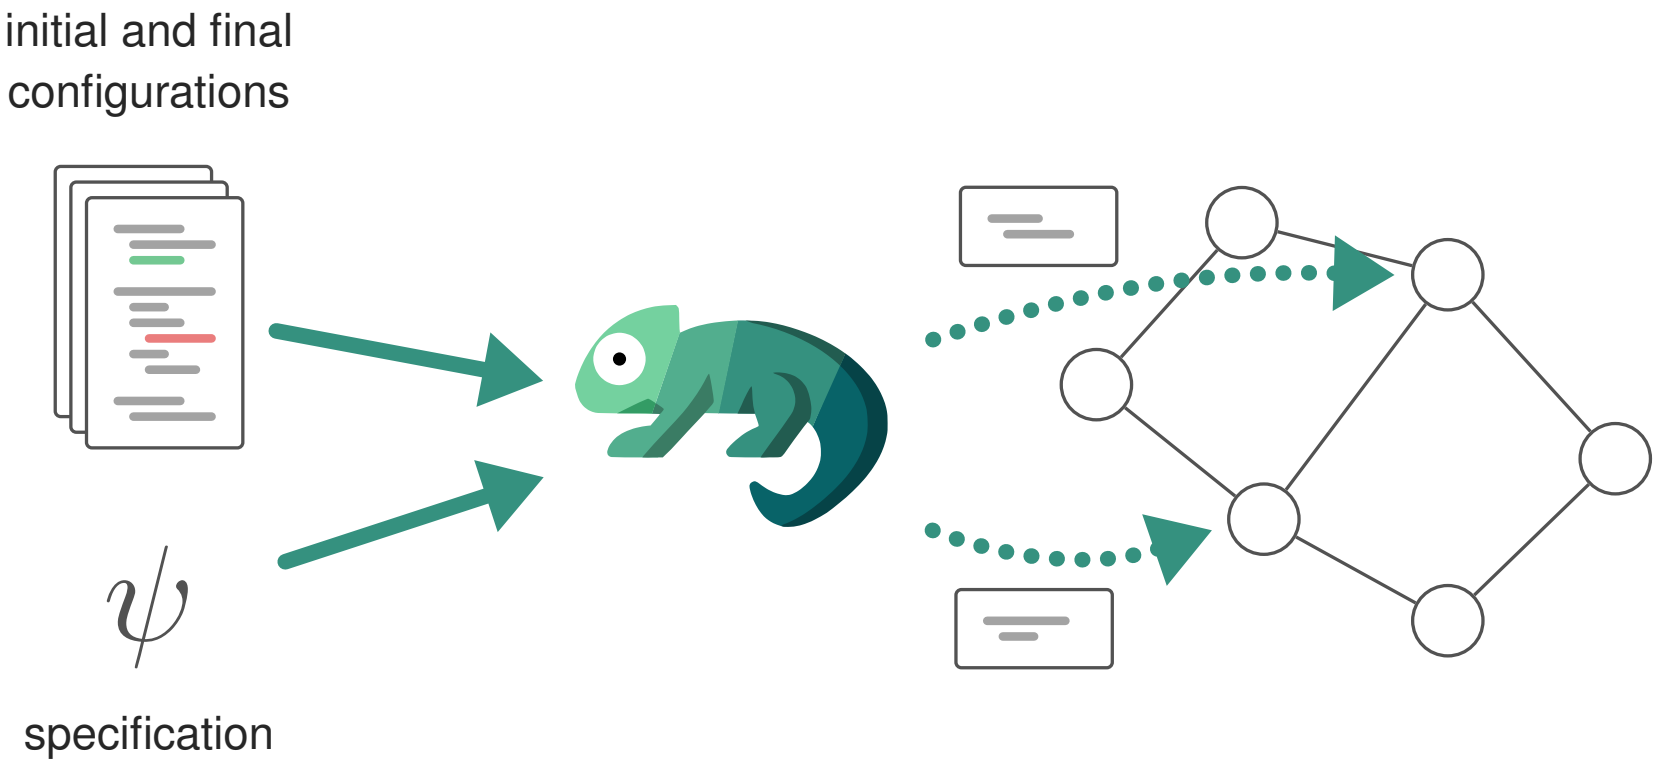

Chameleon schedules and performs reconfiguration within minutes.

Chameleon schedules and performs reconfiguration within minutes.

- Scheduling time is less than five minutes.

Chameleon schedules and performs reconfiguration within minutes.

- Scheduling time is less than five minutes.
- Reconfiguration time is less than five minutes.

Taming the transient while reconfiguring BGP

Tibor Schneider, Roland Schmid,
Stefano Vissicchio, Laurent Vanbever

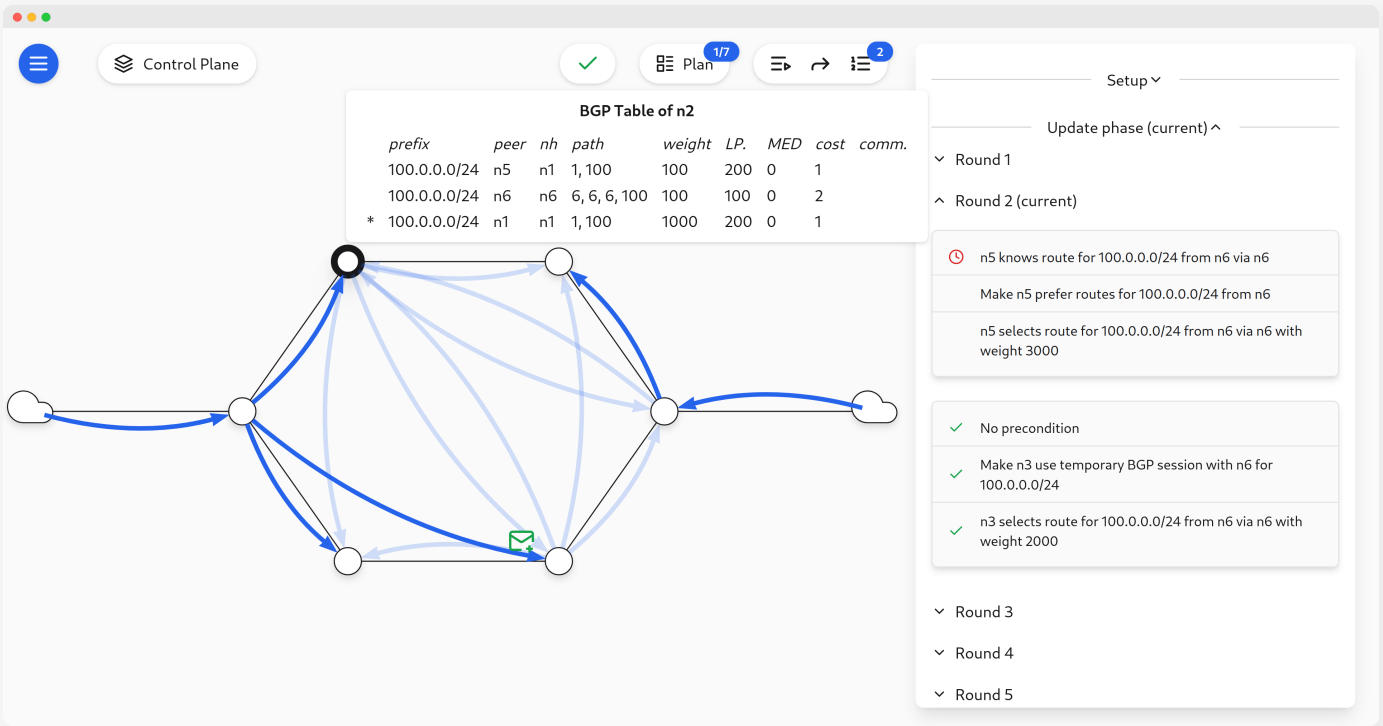

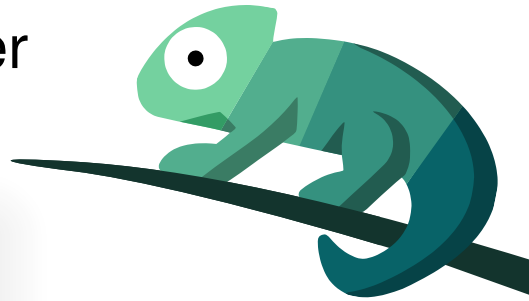

chameleon.ethz.ch