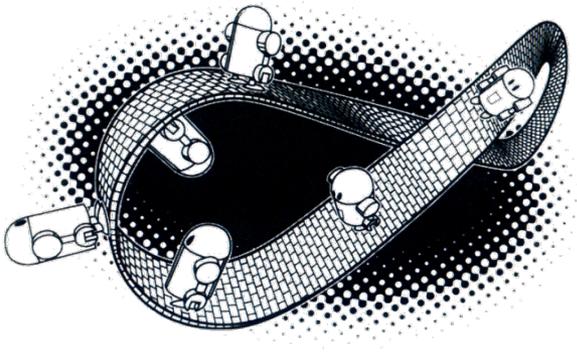


# SDX-Based Flexibility or Internet Correctness? Pick Two!



Rüdiger Birkner, Arpit Gupta  
Nick Feamster, Laurent Vanbever

SOSR'17  
April, 4<sup>th</sup> 2017



**ETH** zürich

# Traditional interdomain routing is limited

BGP does **not allow...**

...fine-grained routing;

...multiple paths;

...direct control over forwarding.

# Traditional interdomain routing is limited

BGP does **not allow...**

...fine-grained routing;

...multiple paths;

...direct control over forwarding.

All of it **is possible** with **SDX**

# SDXes are gaining traction



**Nick Feamster**  
@feamster



Exciting! Bryan Larish @NSAGov discusses our joint #SDX project - [goo.gl/X2QJal](http://goo.gl/X2QJal) and SDN in the wide-area -

## ► CARRIER SDN / SDN TECHNOLOGY

### French Exchange Sees Classic SDN Benefits



France is steeped in the classics, so it's appropriate that a small French Internet exchange is seeing the classic benefits of SDN.

TOUIX serves Toulouse, in the Midi-Pyrénées region in the south of France. It's a small exchange, with about ten interconnected service providers, three PoPs and less than 1 Gbit/s traffic. But it's got the same problems as any exchange of any size and is looking to SDN for the same benefits, including innovation, manageability, enabling new business models and vendor

all size makes it a good candidate for innovation, Marc Bruyère, co-ctaneutral, a nonprofit service provider with a seat on the board of Light Reading. Tetaneutral has 1,000 customers (which it calls and provides WiFi and DSL connectivity in areas called "zones white spaces -- without connectivity.

Toulouse



## On the feasibility of converting AMS-IX to an Industrial-Scale Software Defined Internet Exchange Point

Thursday Keynote - SDN in Enterprises  
ing Summits, Inc. (ONS) represents the industry's y non-profit conference that brings together the g ecosystem to ...

UNIVERSITY OF AMSTERDAM  
*System and Network Engineering*  
MSc Research Project

### Authors

Siem Hermans  
[siem.hermans@os3.nl](mailto:siem.hermans@os3.nl)  
Jeroen Schutruf  
[jeroen.schutruf@os3.nl](mailto:jeroen.schutruf@os3.nl)

July 29, 2016

### Supervis

Arien  
[arien.vijn@ams-ix](mailto:arien.vijn@ams-ix)  
Joris Claas  
[joris.claassen@ams-ix](mailto:joris.claassen@ams-ix)

## Project CARDIGAN An SDN Controlled Exchange Fabric



Scott Whyte  
[subyte@google.com](mailto:subyte@google.com)

## CODE DROP! ISDX PROJECT TO OPEN SOURCE SDN

Posted on March 18, 2016 - By Dan Pitt, Executive Director, ONF

With the Open Networking Summit taking place this week in Santa Clara, open source SDN projects are the talk of the town. We'd like to talk to you about one such project that we're particularly proud to be involved with – ISDX, an industrial-scale Software Defined Internet Exchange point.

This week, Princeton University took center stage at ONS to announce the contribution of ISDX code to ONF's open source community, OpenSourceSDN.org. Lead by Arpit Gupta, a team of researchers has been working for the last three years to develop mechanisms that could fundamentally change how independently operated networks interconnect and exchange traffic. Over the past year, the Princeton team has worked closely with partner enterprise networks to harden the code and test it on hardware switches, including the low-cost Quanta LY2. A recent breakthrough with the project has enabled the developers to allow the scaling challenge of using SDN controllers to manage high-volume traffic and the applying of flow-based policy beyond what conventional routing can offer. The project is now at the stage where it can be shared with the broader open source community.

ONF Open Networking Fndn  
@openflow



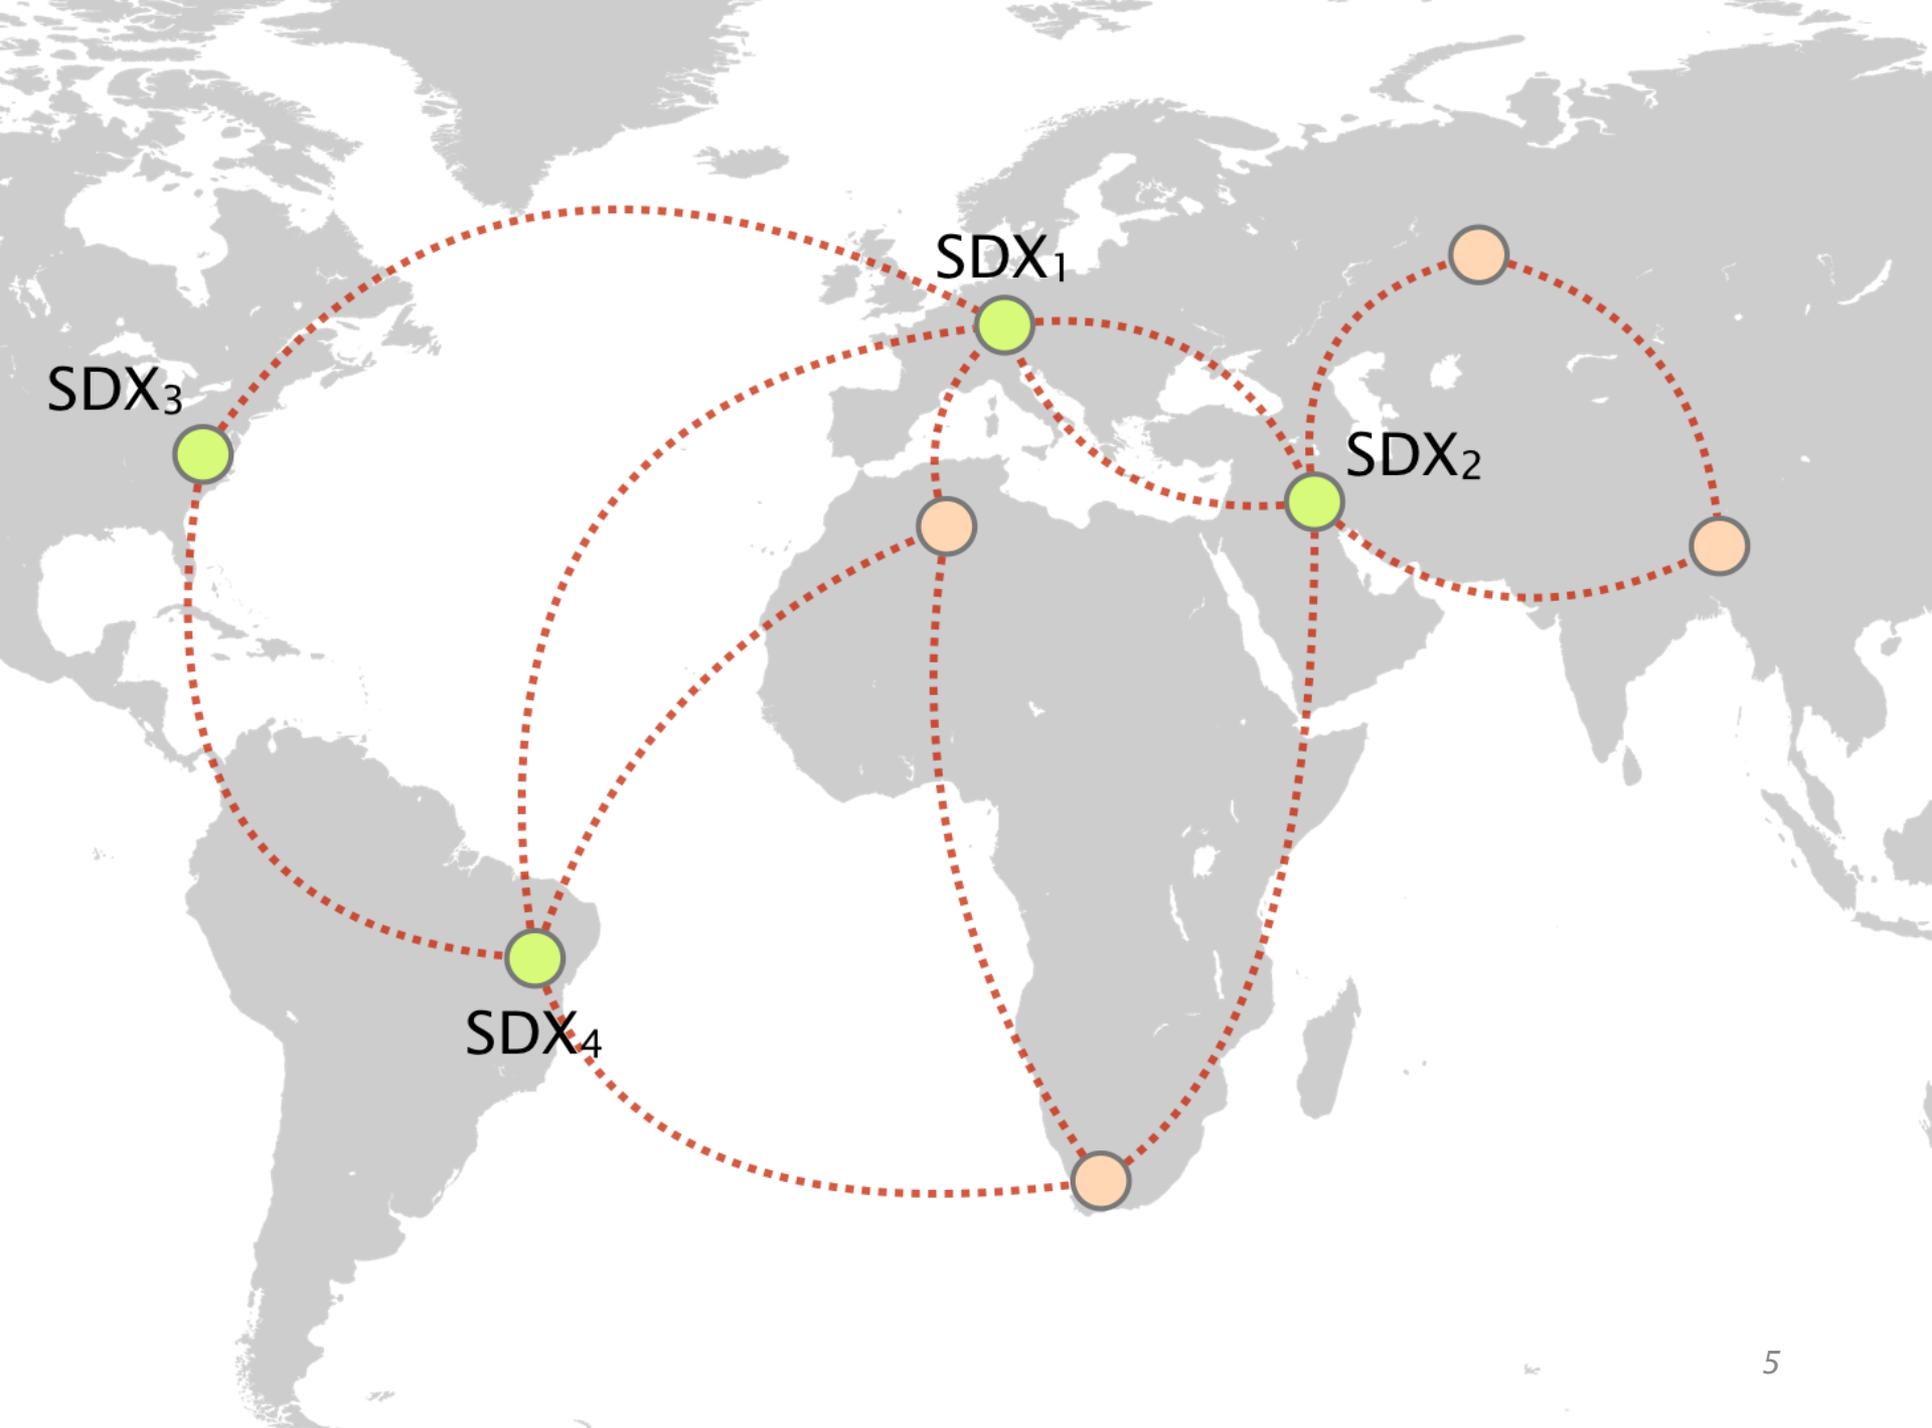
.@Bauer\_Rick, @glexqsd, & @feamster give a closer look at #ISDX in our latest @SDxCentral webinar. Watch it here

## ONF & ISDX Webinar: Software-Defined Internet Exchange Points & Solving Internet Routing Problems

Video: ONF & ISDX Webinar - SDN, Internet Exchange Points

ONF & ISDX Webinar featuring Princeton University present an exciting Webinar on software-defined Internet exchange point (ISDX). Watch the

[sdxcentral.com](http://sdxcentral.com)

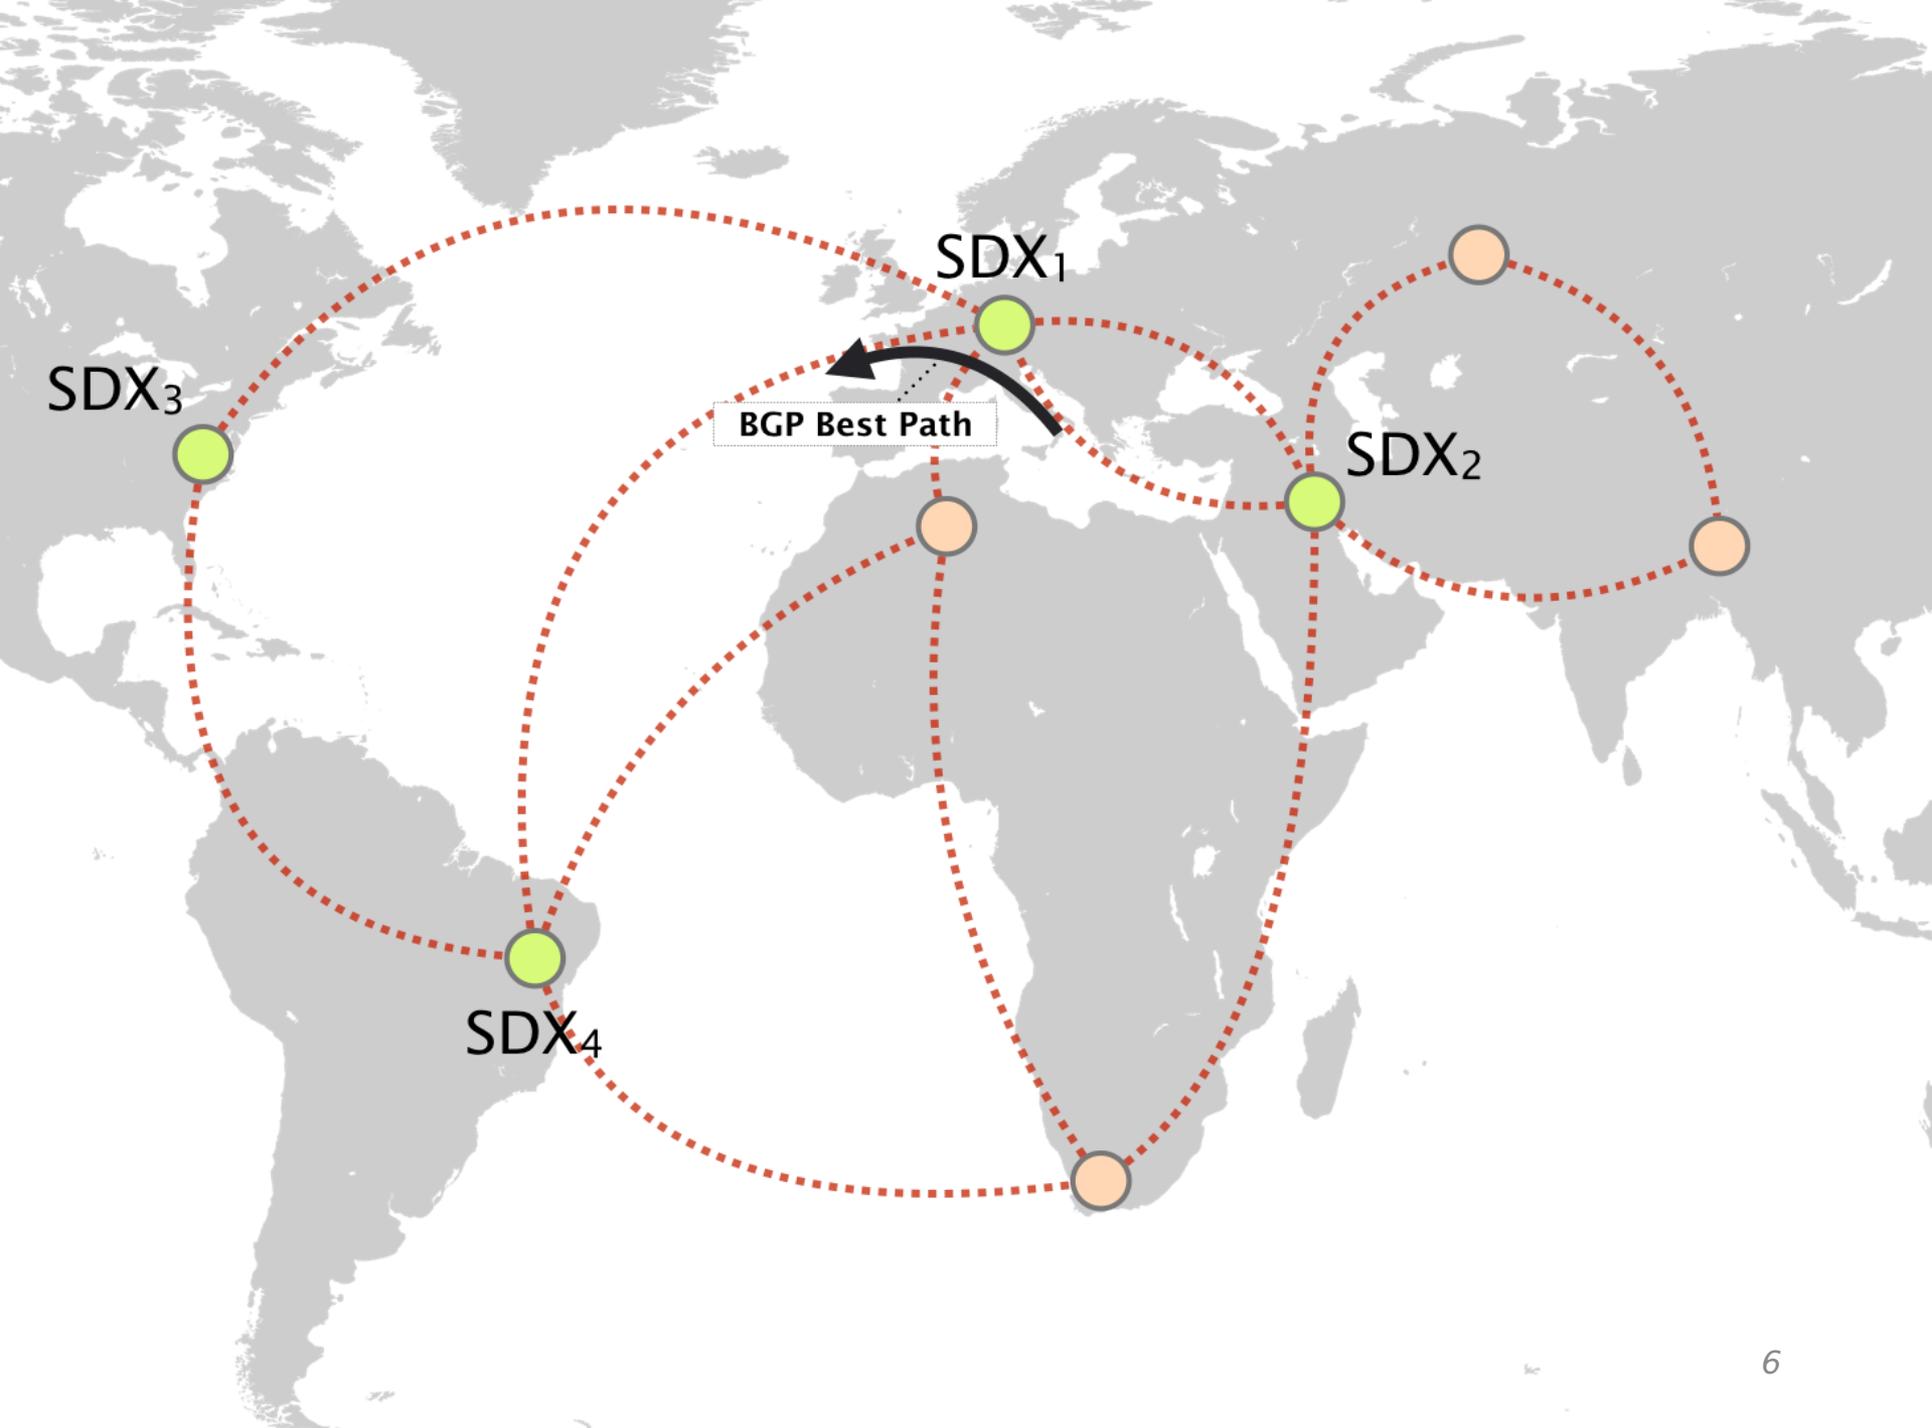


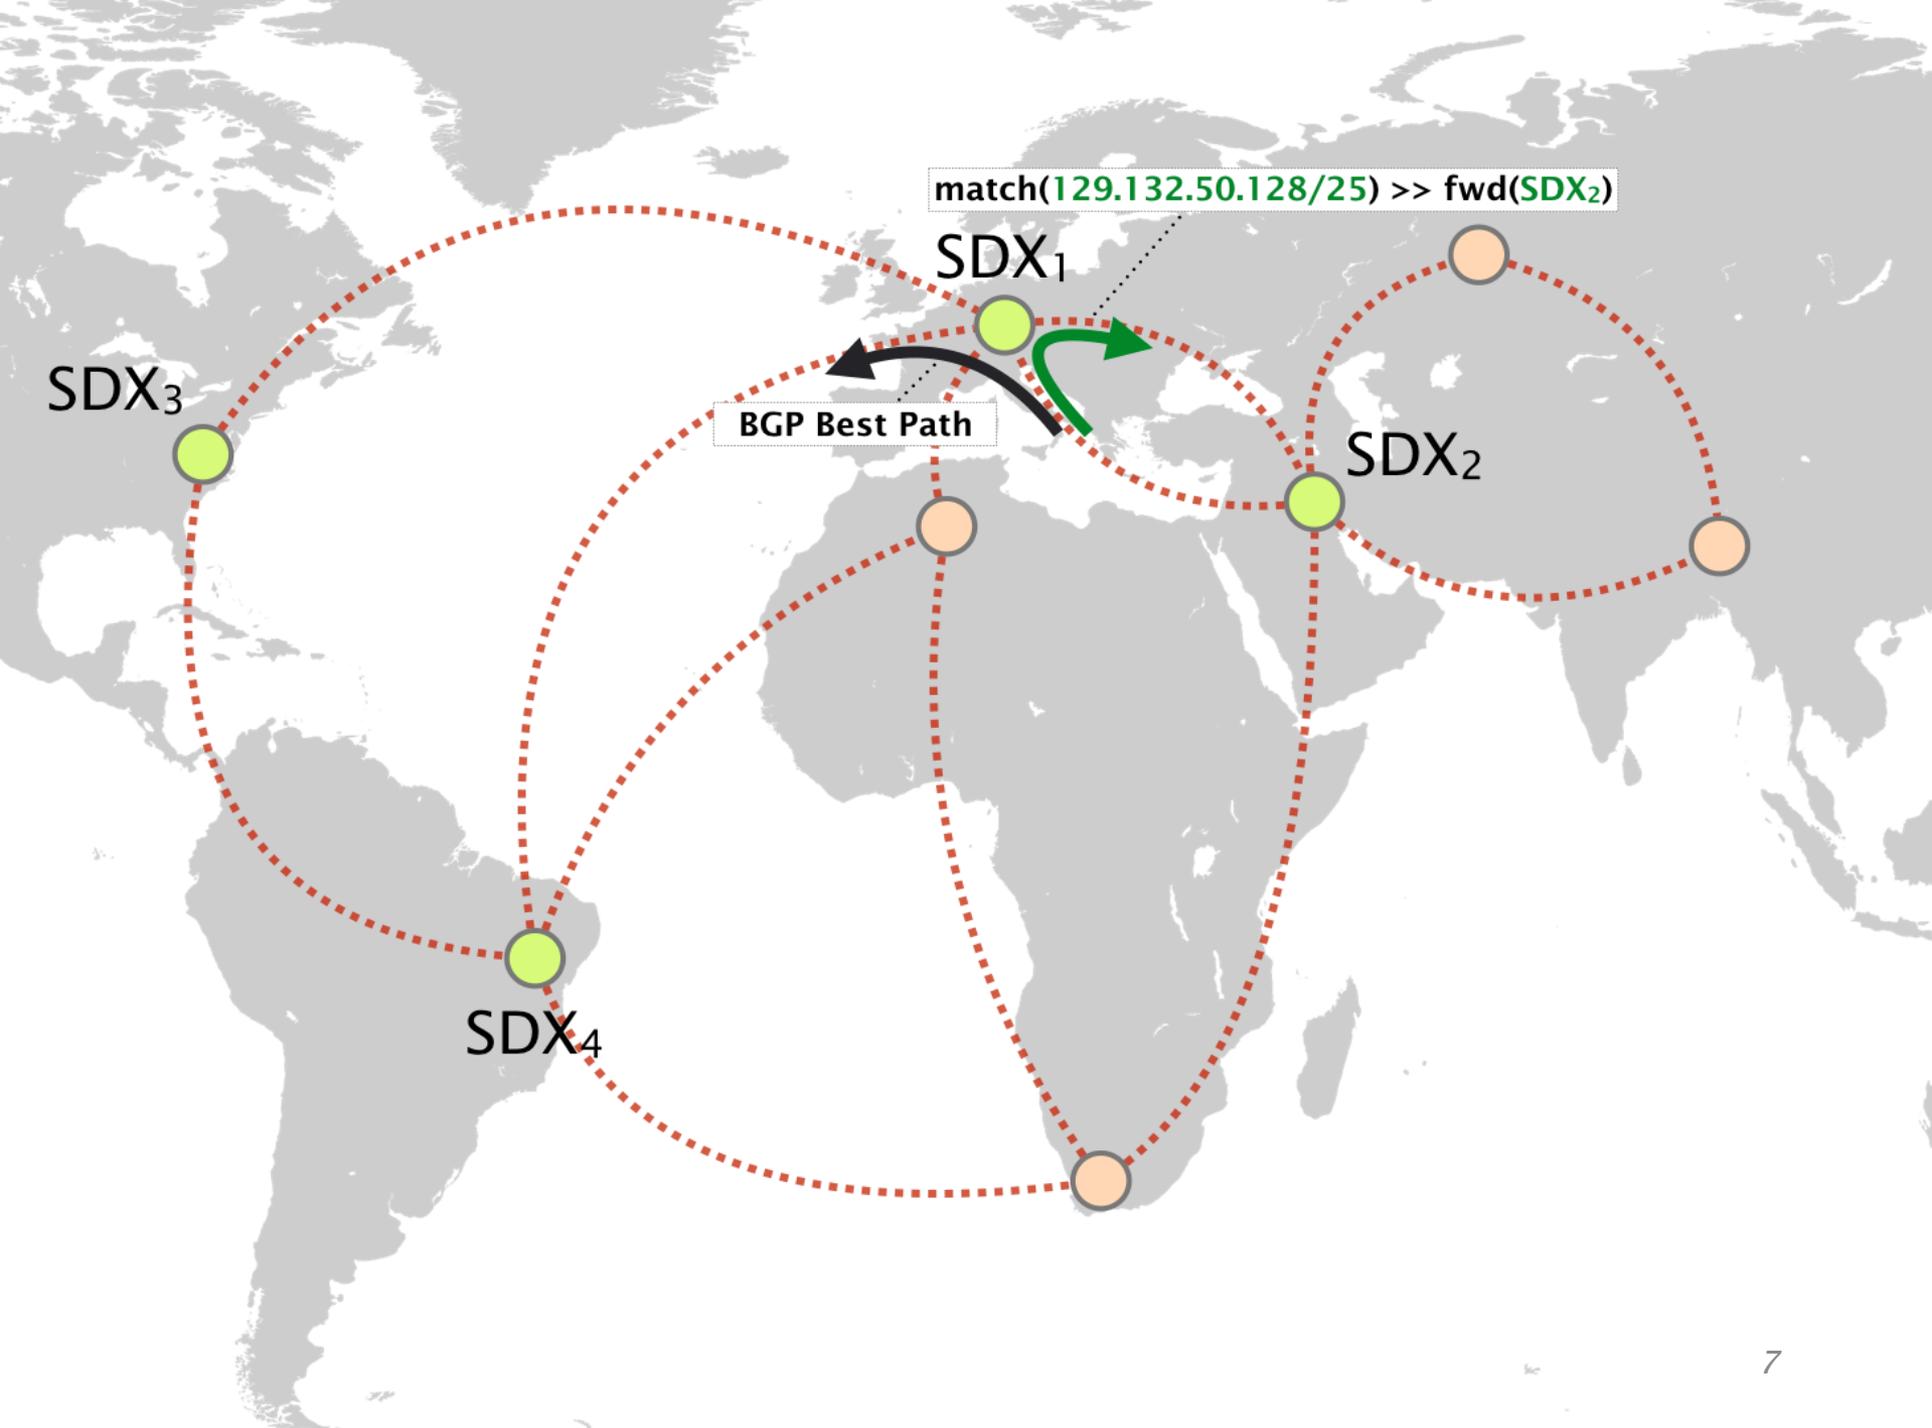
SDX<sub>3</sub>

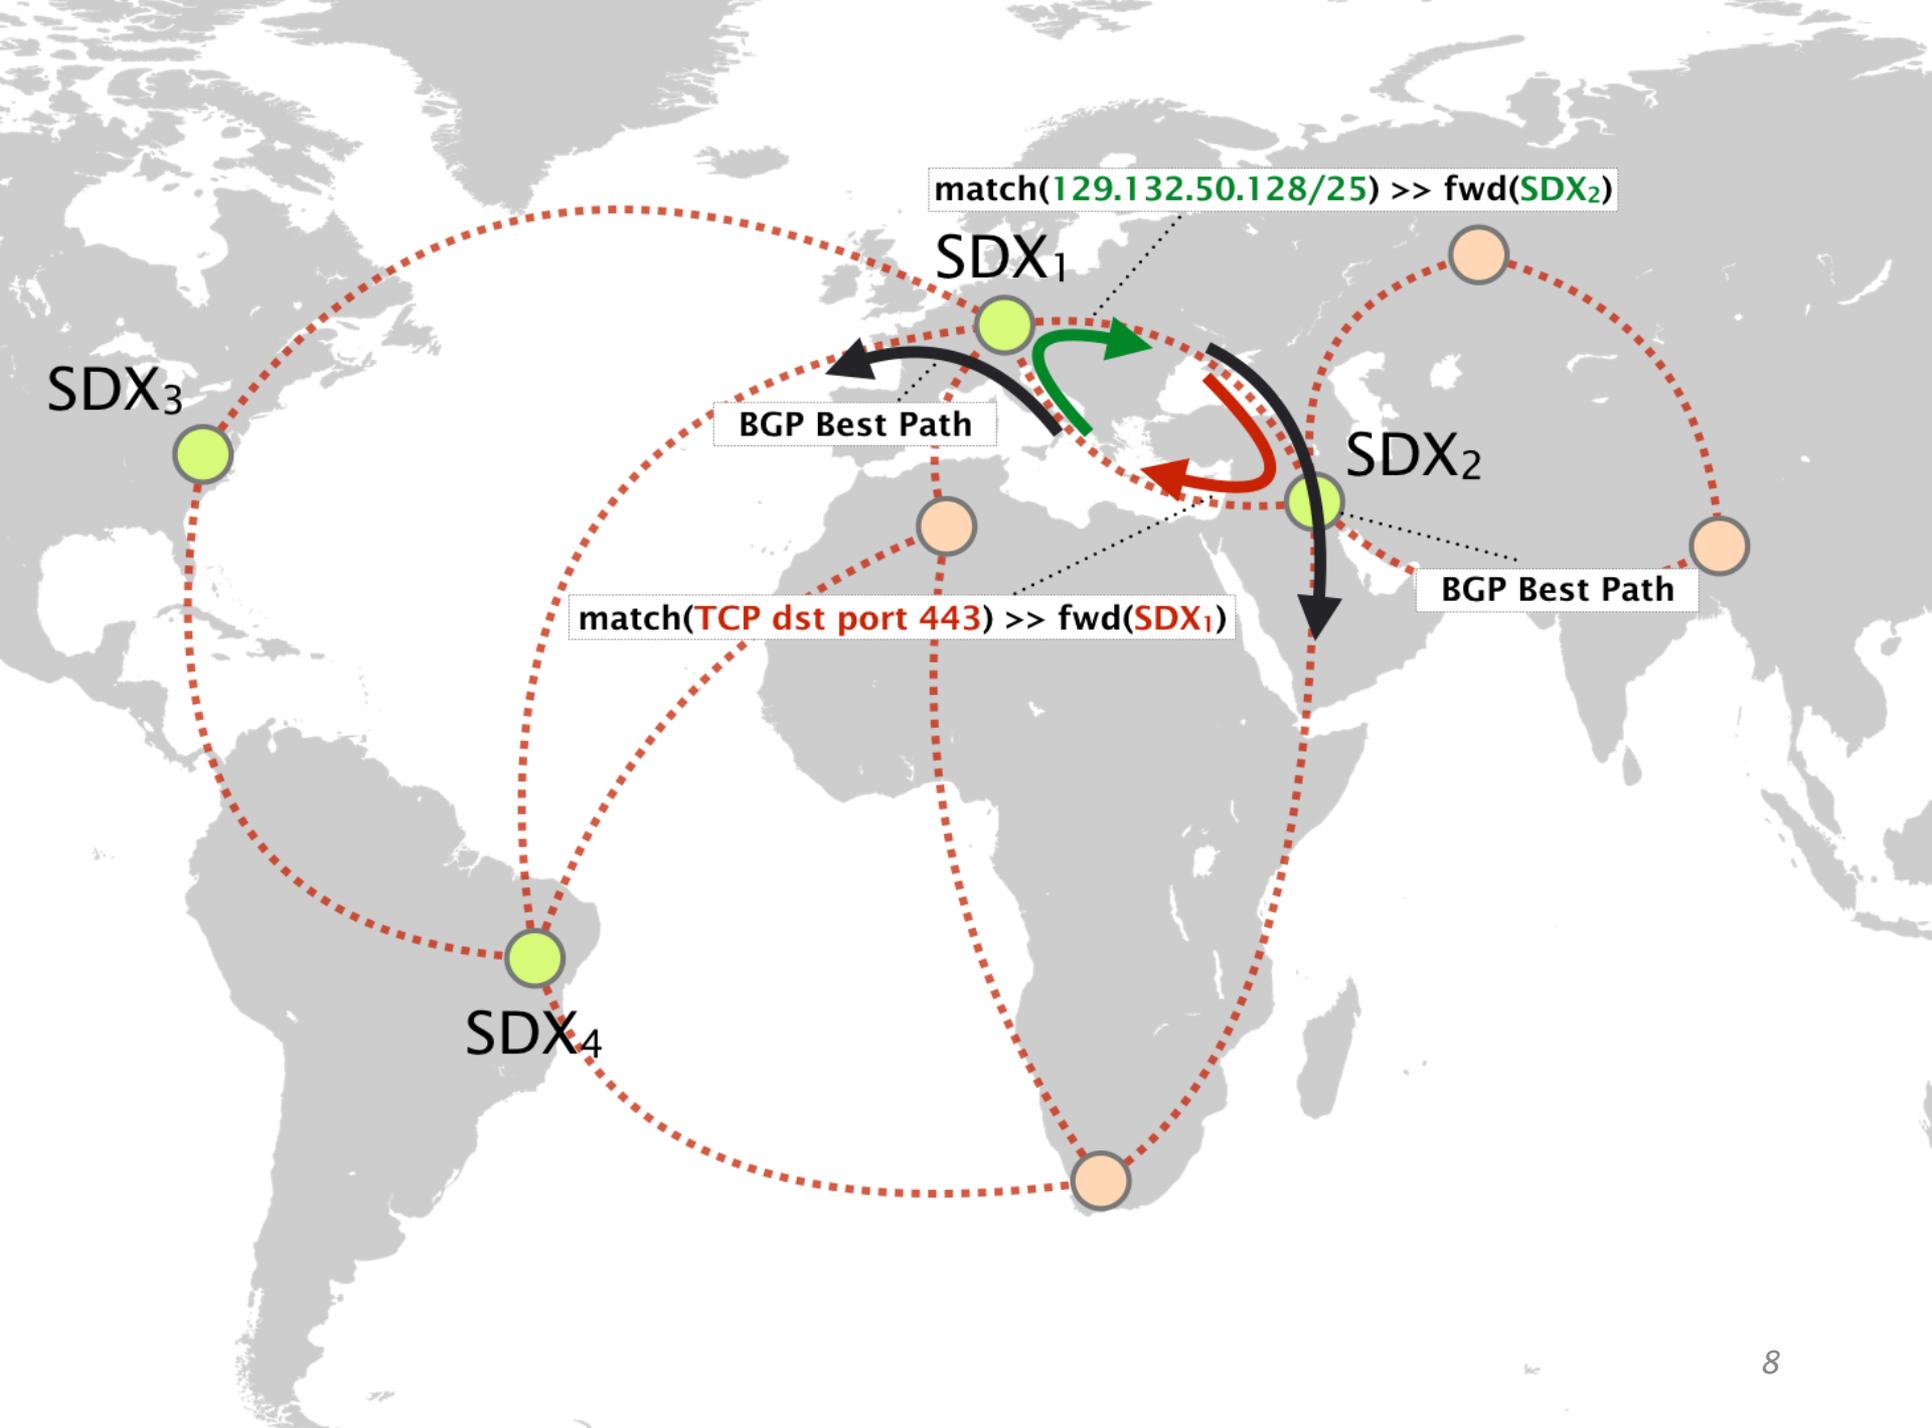
SDX<sub>1</sub>

SDX<sub>2</sub>

SDX<sub>4</sub>







match(129.132.50.128/25) >> fwd(SDX<sub>2</sub>)

SDX<sub>1</sub>

SDX<sub>3</sub>

BGP Best Path

SDX<sub>2</sub>

BGP Best Path

match(TCP dst port 443) >> fwd(SDX<sub>1</sub>)

SDX<sub>4</sub>

```
match(129.132.50.128/25) >> fwd(SDX2)
```

## Uncoordinated policy installations cause forwarding loops

SDX<sub>1</sub>

SDX<sub>1</sub> Best Path

SDX<sub>2</sub>

match(

**Detection is hard**

as probing is not feasible.

SDX<sub>4</sub>

```
match(129.132.50.128/25) >> fwd(SDX2)
```

## Uncoordinated policy installations cause forwarding loops

SDX<sub>1</sub>

SDX<sub>2</sub>

SDX<sub>4</sub>

Detection is hard  
as probing is not feasible.

**Loop prevention is useless**  
as the control plane is not  
aware of remote policies.

```
match(129.132.50.128/25) >> fwd(SDX2)
```

## Uncoordinated policy installations cause forwarding loops

SDX<sub>1</sub>

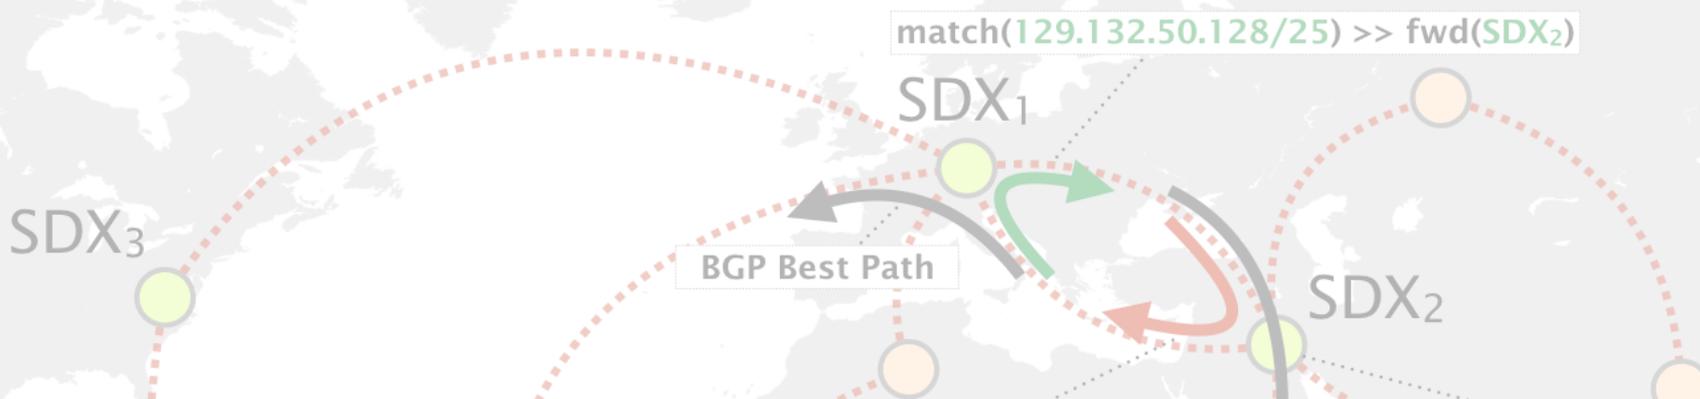
SDX<sub>2</sub>

SDX<sub>4</sub>

Detection is hard  
as probing is not feasible.

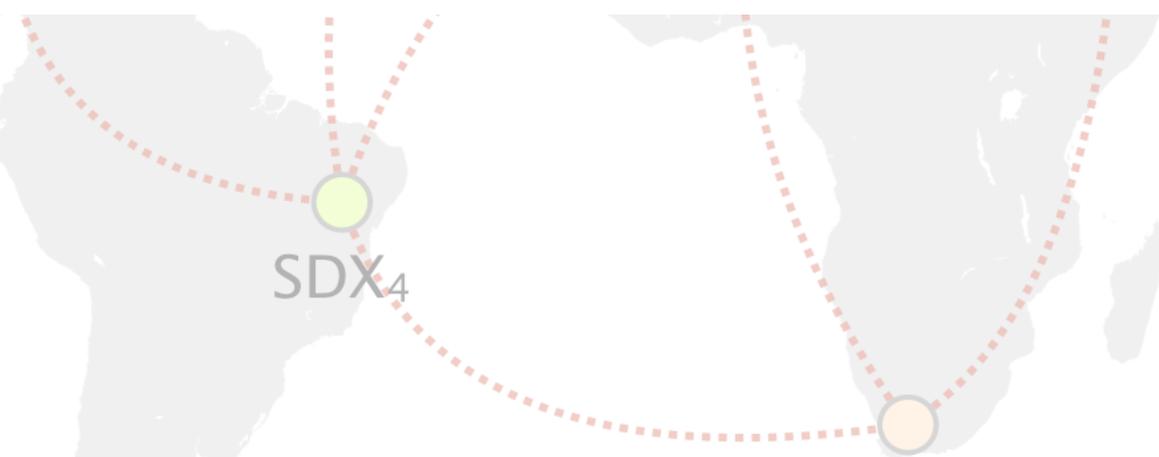
Loop prevention is useless  
as the control plane is not  
aware of remote policies.

**To guarantee correctness,**  
we need to exchange state.



## Solution: SIDR

# Safe Interdomain Deflection based Routing



# SIDR faces three challenges

# SIDR faces three challenges

Privacy

SDX participants are reluctant to share the exact policies.

# SIDR faces three challenges

Privacy

SDX participants are reluctant to share the exact policies.

Scalability

Exchanging state about all deflections gets out of hand.

# SIDR faces three challenges

Privacy

SDX participants are reluctant to share the exact policies.

Scalability

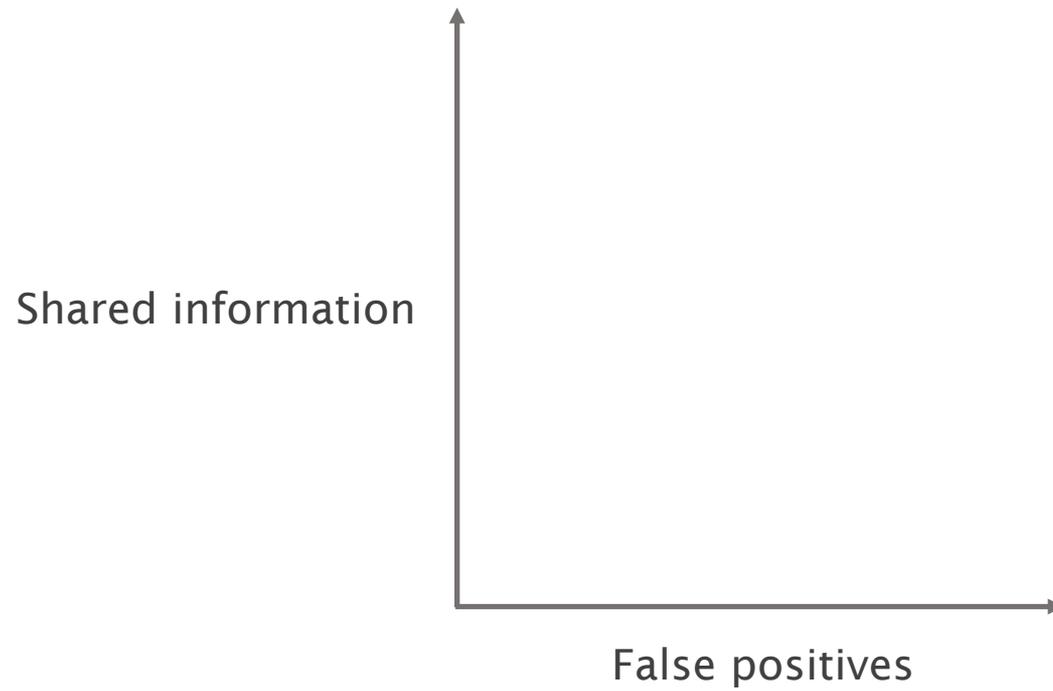
Exchanging state about all deflections gets out of hand.

Flexibility

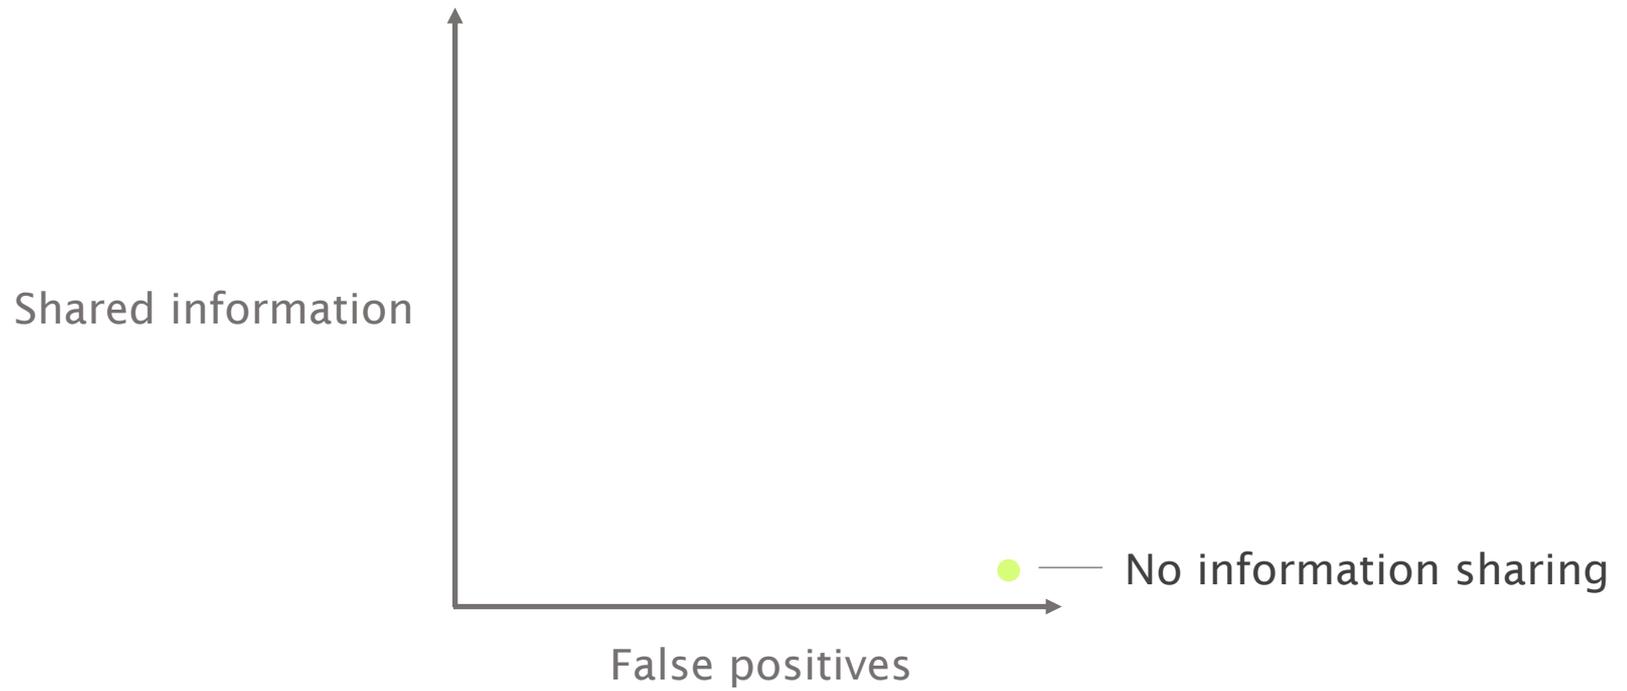
Installing as many safe policies as possible.

# Trade-off

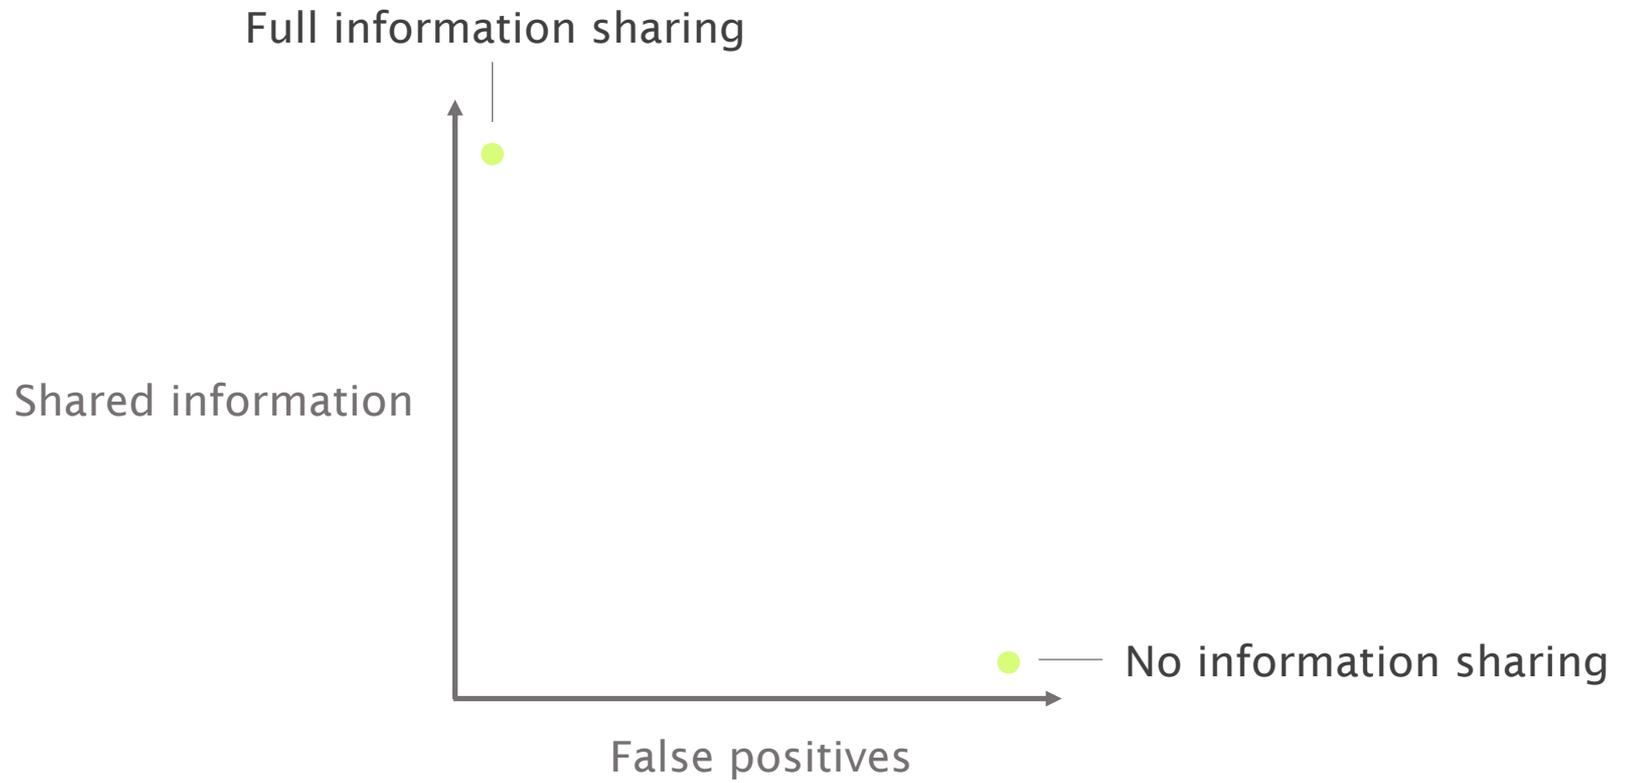
# Trade-off



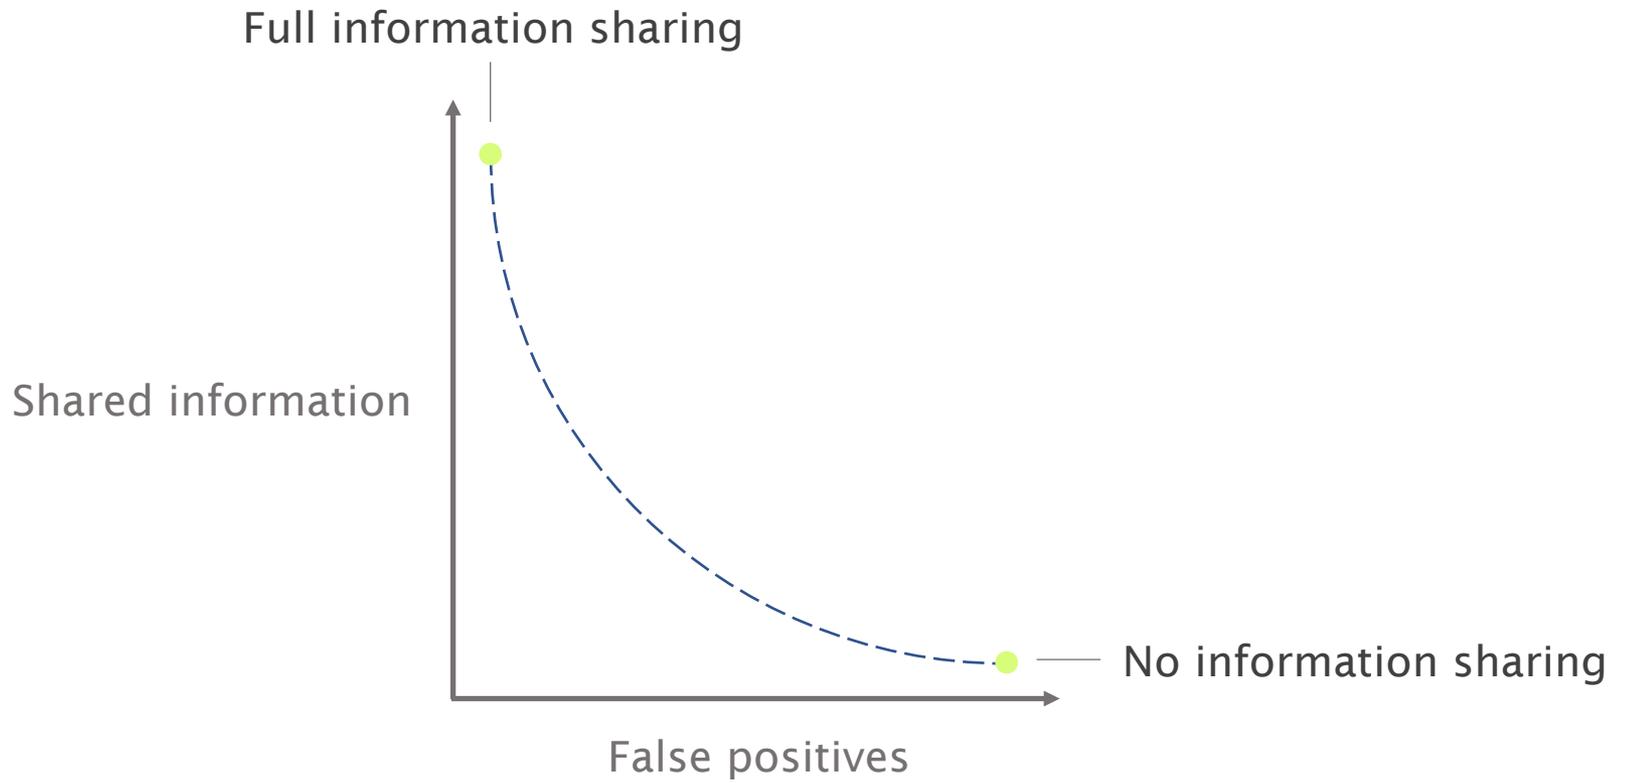
# Trade-off



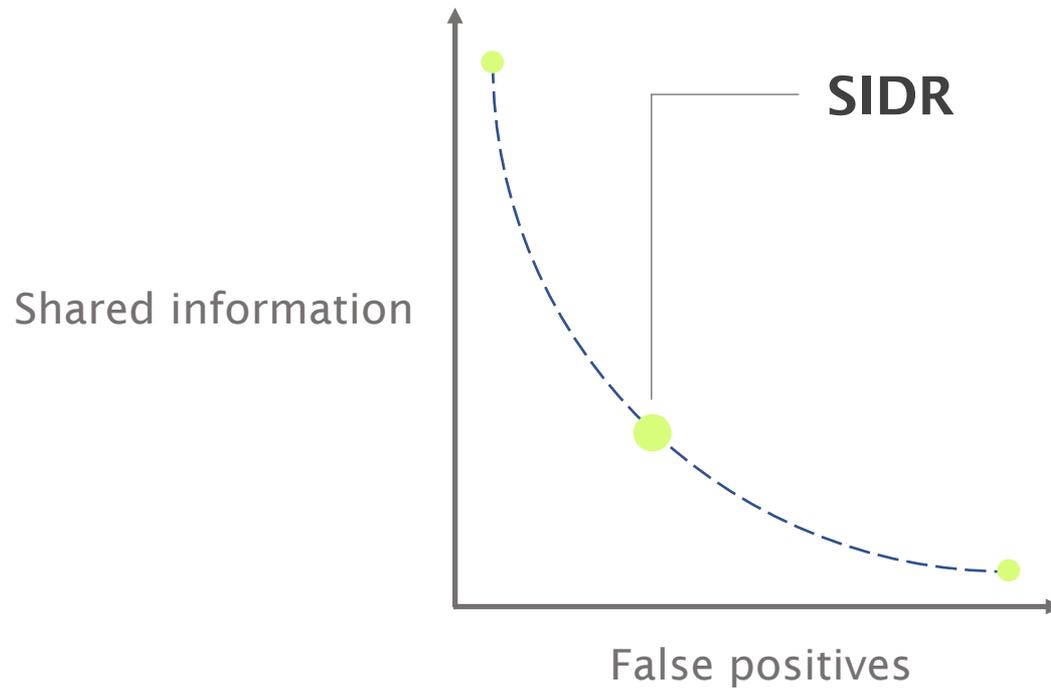
# Trade-off



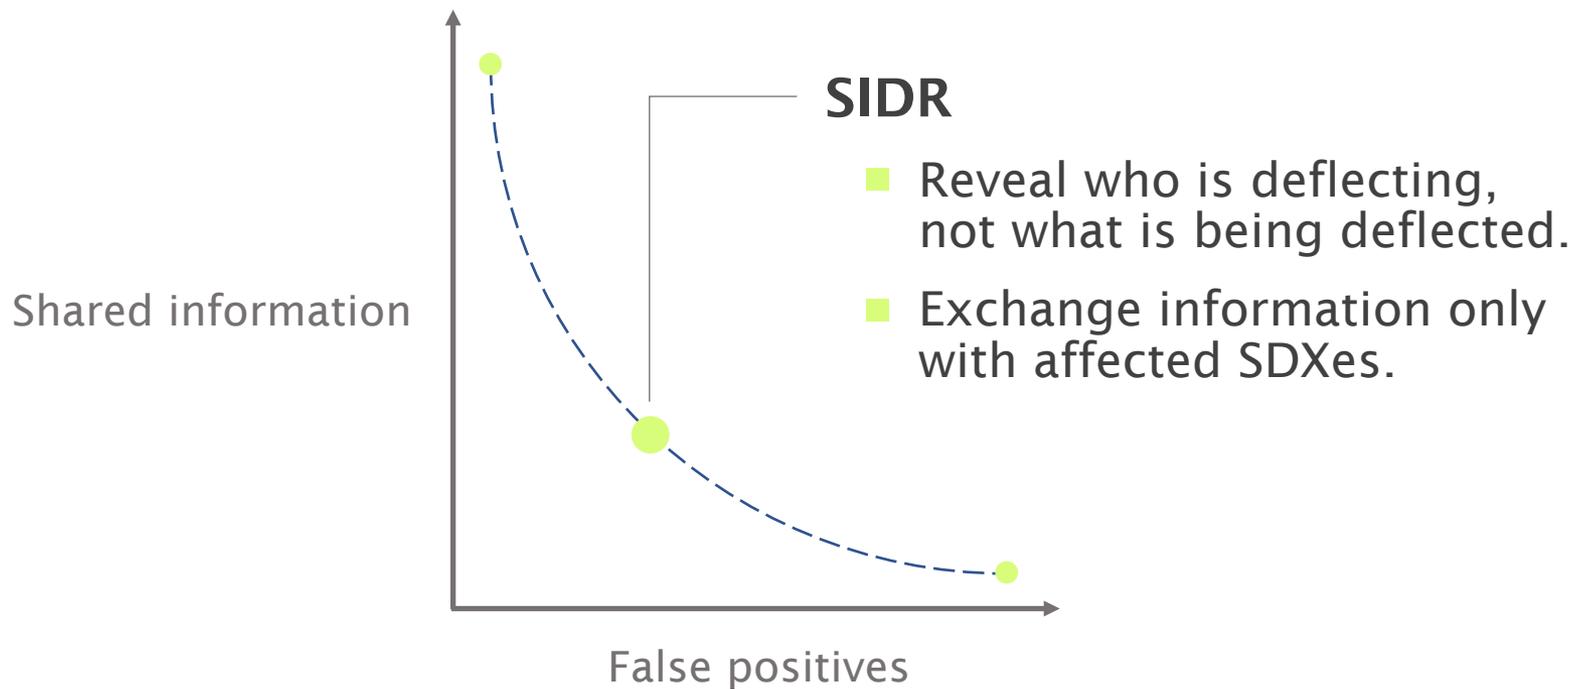
# Trade-off



# Trade-off



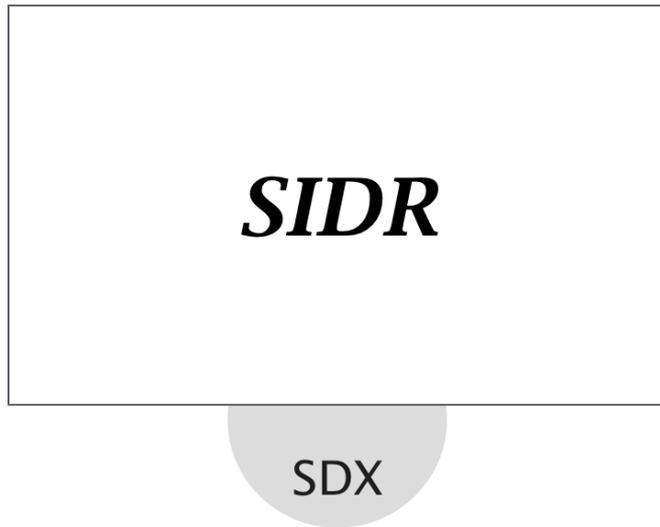
# Trade-off



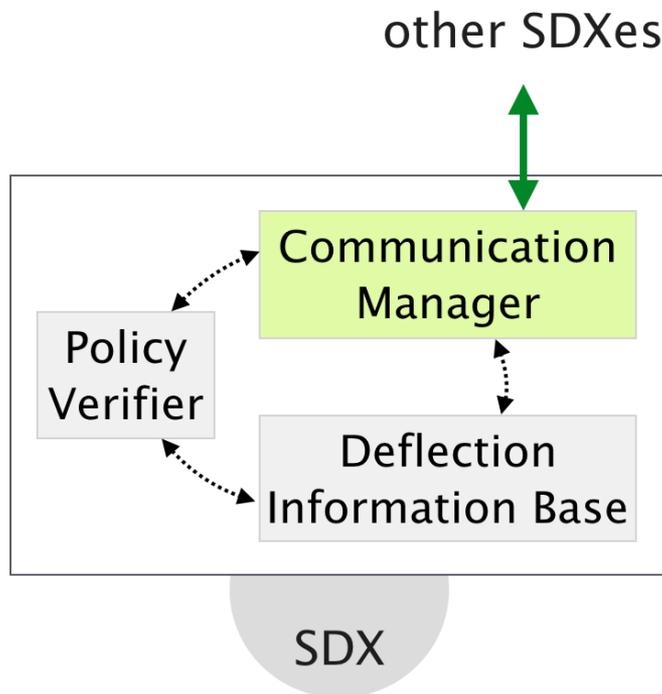
# SIDR coordination framework



# SIDR coordination framework



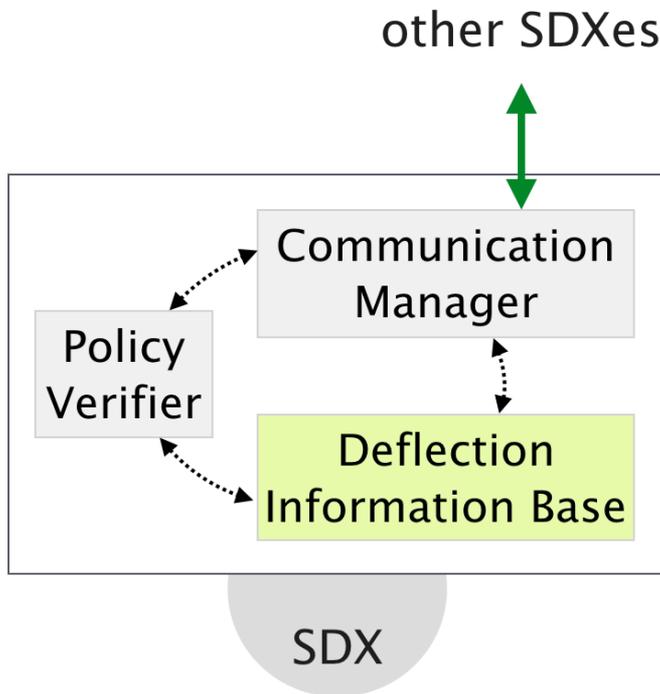
# SIDR coordination framework



Communication Manager

Learn about remote deflections.

# SIDR coordination framework



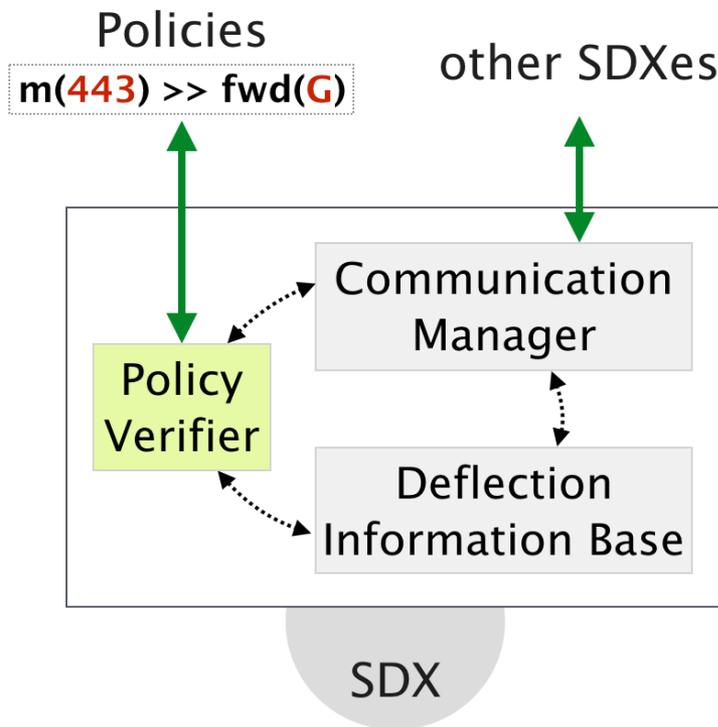
Communication Manager

Learn about remote deflections.

Deflection Information Base

Keep track of remote deflections.

# SIDR coordination framework



Communication Manager

Learn about remote deflections.

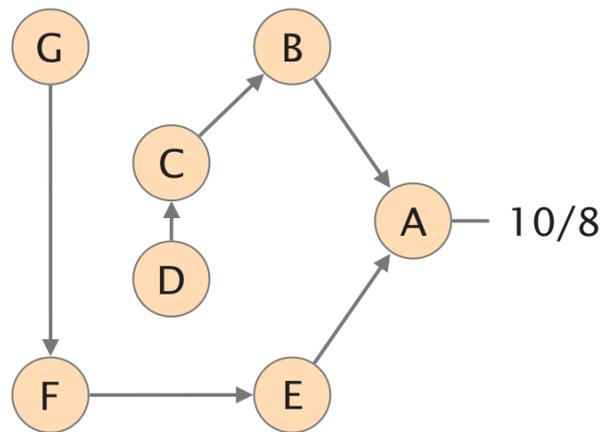
Deflection Information Base

Keep track of remote deflections.

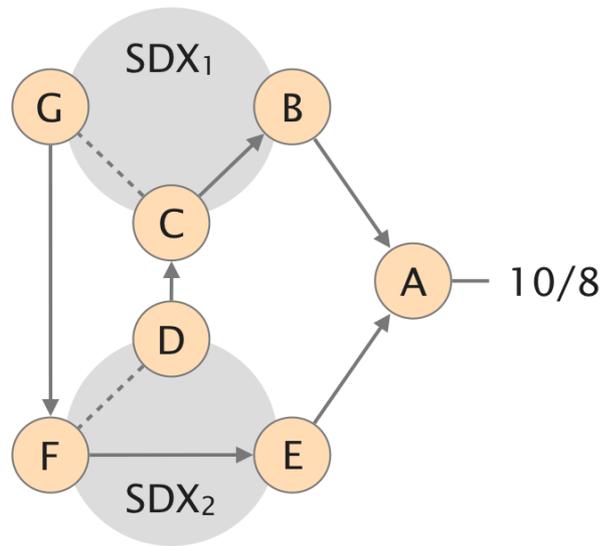
Policy Verifier

Consult BGP best path and DIB to check policies.

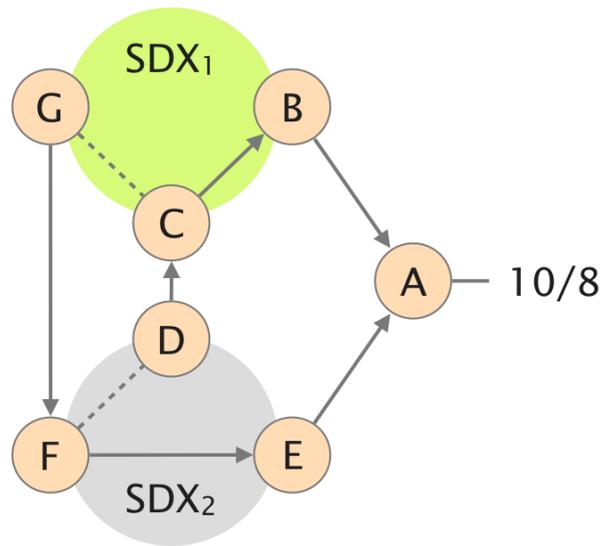
# Safe Inter-domain Deflection Routing



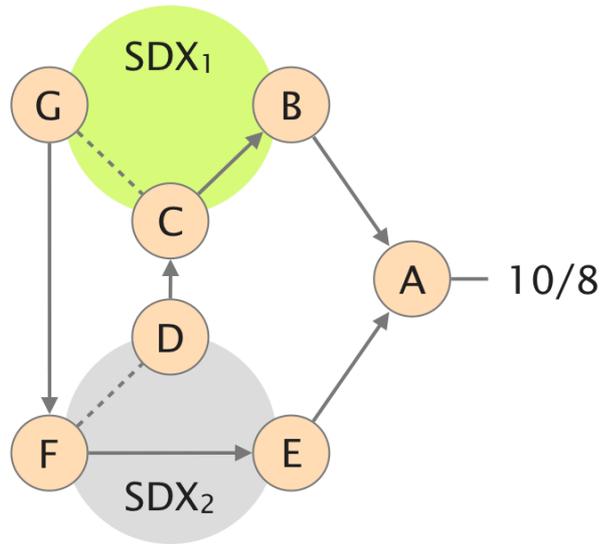
# Safe Inter-domain Deflection Routing



# Safe Inter-domain Deflection Routing



# Safe Inter-domain Deflection Routing



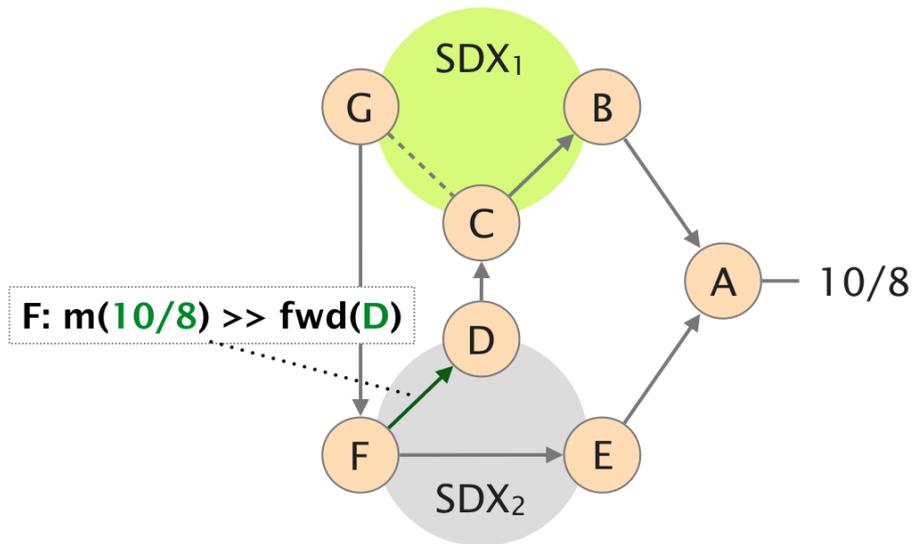
*Deflection DB @ SDX<sub>1</sub>*

	Prefix	Deflection Set
...	...	...

*BGP Paths @ SDX<sub>1</sub>*

	Prefix	AS Path
C	10/8	[G, F, E, A]
...	...	...

# Safe Inter-domain Deflection Routing



*Deflection DB @ SDX<sub>1</sub>*

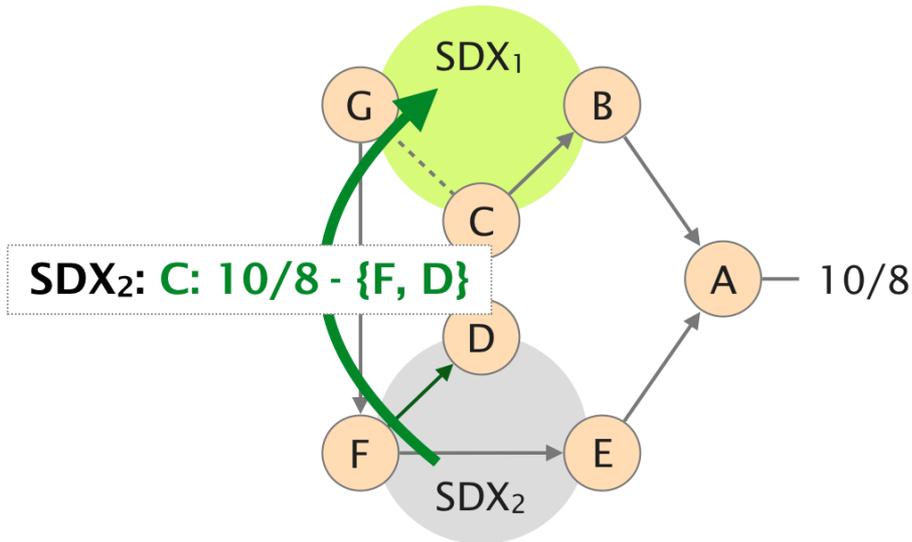
	Prefix	Deflection Set
...	...	...

*BGP Paths @ SDX<sub>1</sub>*

	Prefix	AS Path
C	10/8	[G, F, E, A]
...	...	...

# Safe Inter-domain Deflection Routing

## (1) Policy notification



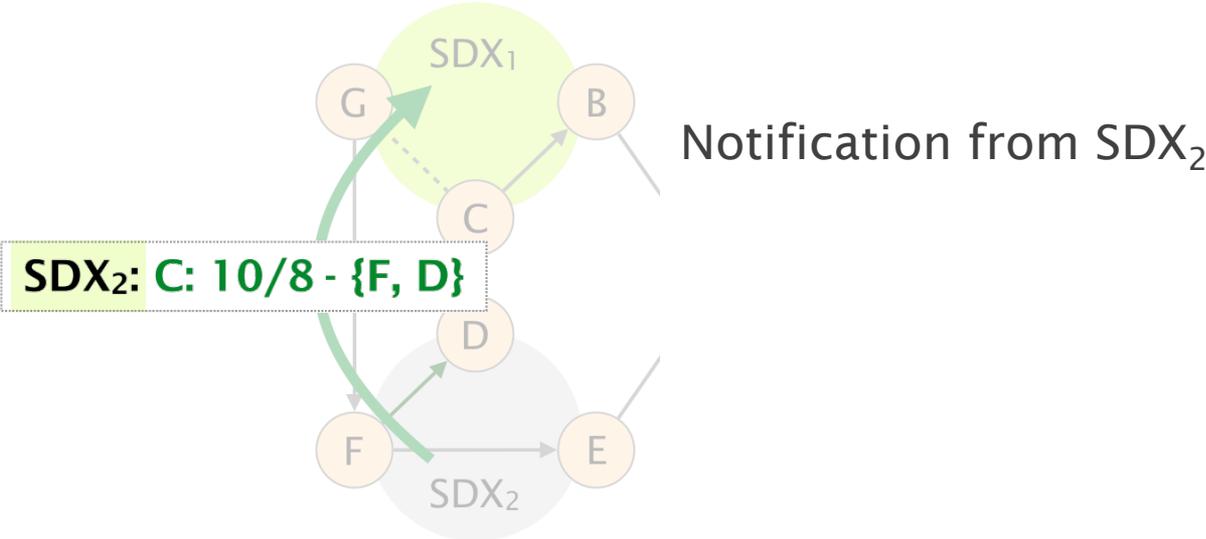
## Deflection DB @ SDX<sub>1</sub>

	Prefix	Deflection Set
...	...	...

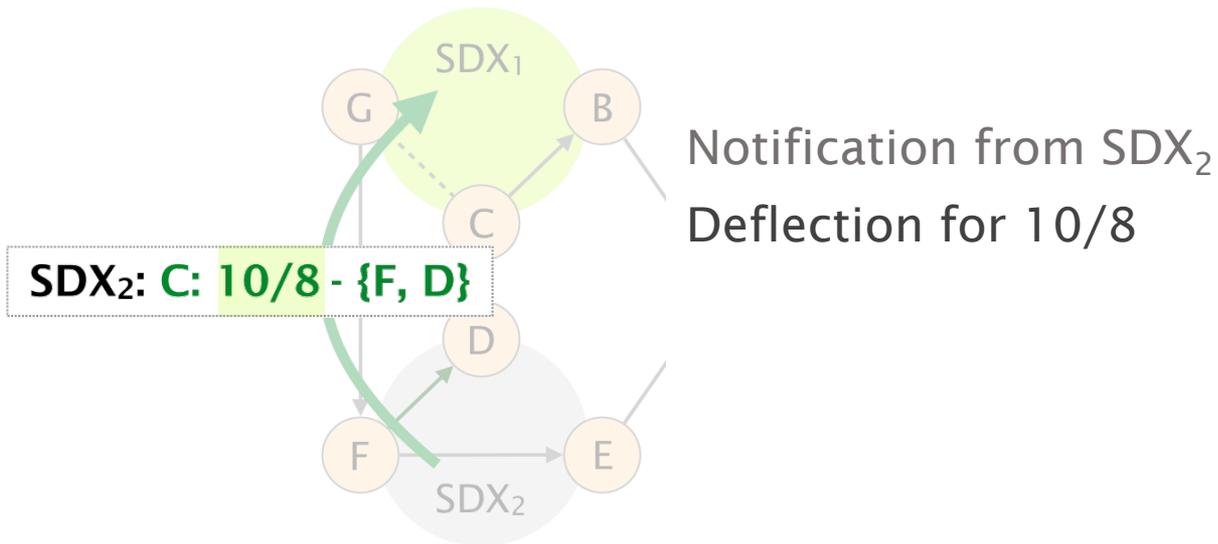
## BGP Paths @ SDX<sub>1</sub>

	Prefix	AS Path
C	10/8	[G, F, E, A]
...	...	...

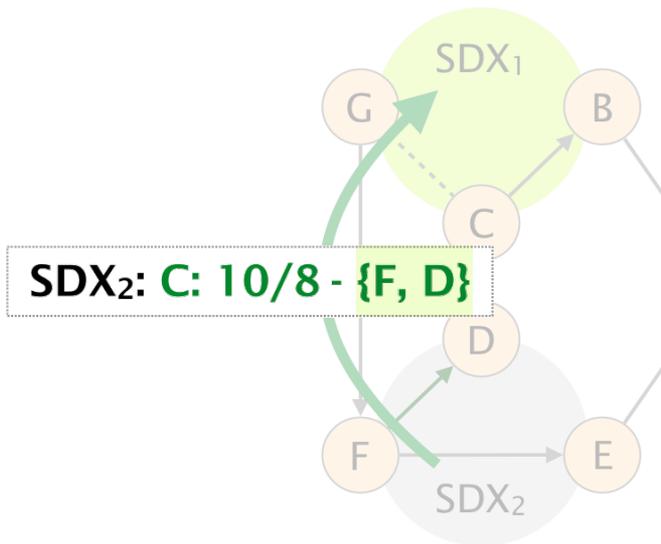
# Safe Inter-domain Deflection Routing



# Safe Inter-domain Deflection Routing



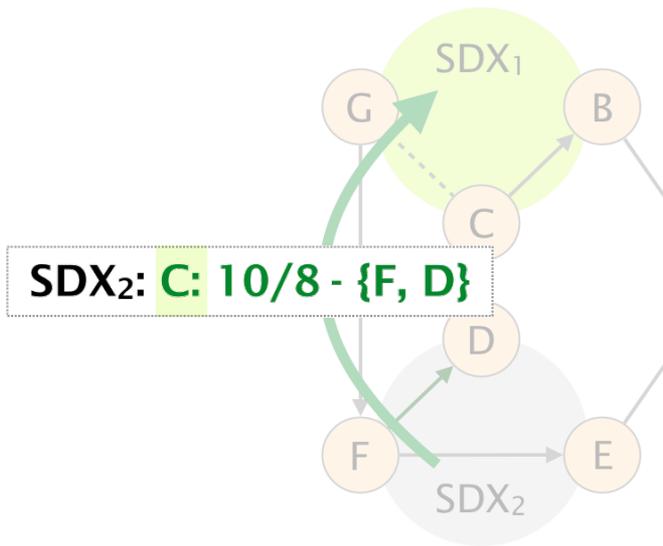
# Safe Inter-domain Deflection Routing



Notification from SDX<sub>2</sub>

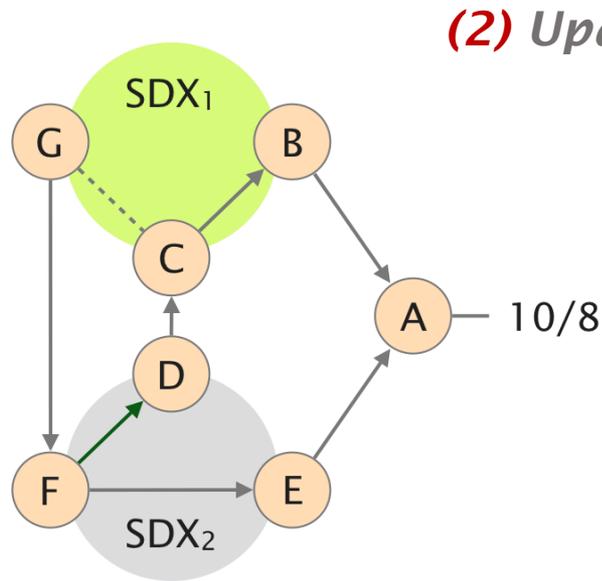
Deflection for 10/8 between F and D

# Safe Inter-domain Deflection Routing



Notification from SDX<sub>2</sub>  
Deflection for 10/8 between F and D  
SDX<sub>1</sub> is affected through C

# Safe Inter-domain Deflection Routing



*Deflection DB @ SDX<sub>1</sub>*

	Prefix	Deflection Set
C	10/8	{F, D}
...	...	...

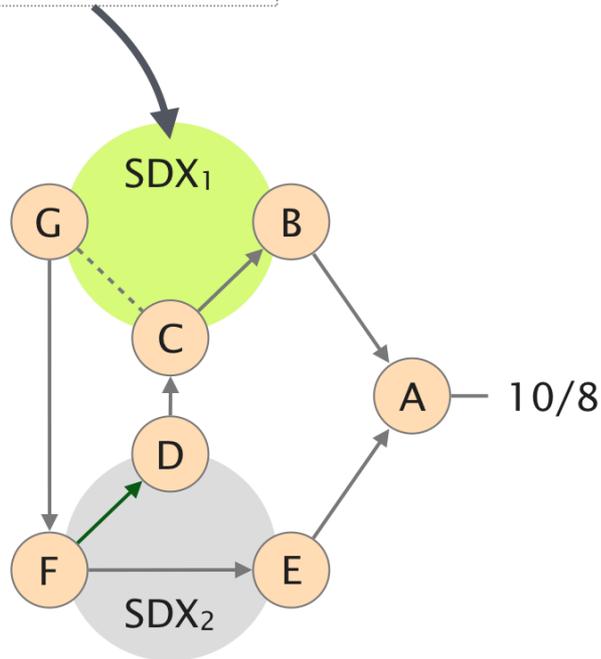
*BGP Paths @ SDX<sub>1</sub>*

	Prefix	AS Path
C	10/8	[G, F, E, A]
...	...	...

# Safe Inter-domain Deflection Routing

## (3) Policy activation request

C: m(443) >> fwd(G)



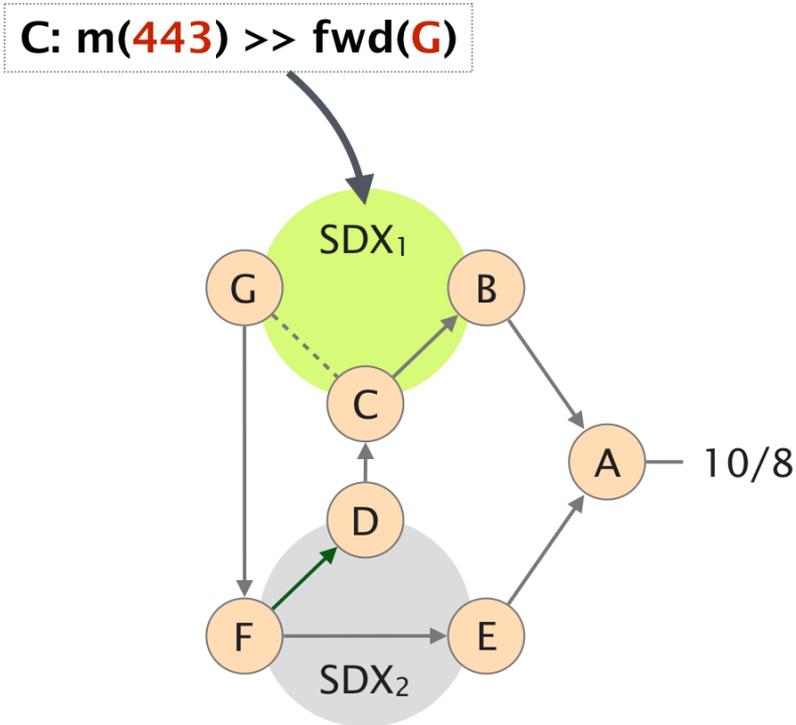
*Deflection DB @ SDX<sub>1</sub>*

	Prefix	Deflection Set
C	10/8	{F, D}
...	...	...

*BGP Paths @ SDX<sub>1</sub>*

	Prefix	AS Path
C	10/8	[G, F, E, A]
...	...	...

# Safe Inter-domain Deflection Routing



*Deflection DB @ SDX<sub>1</sub>*

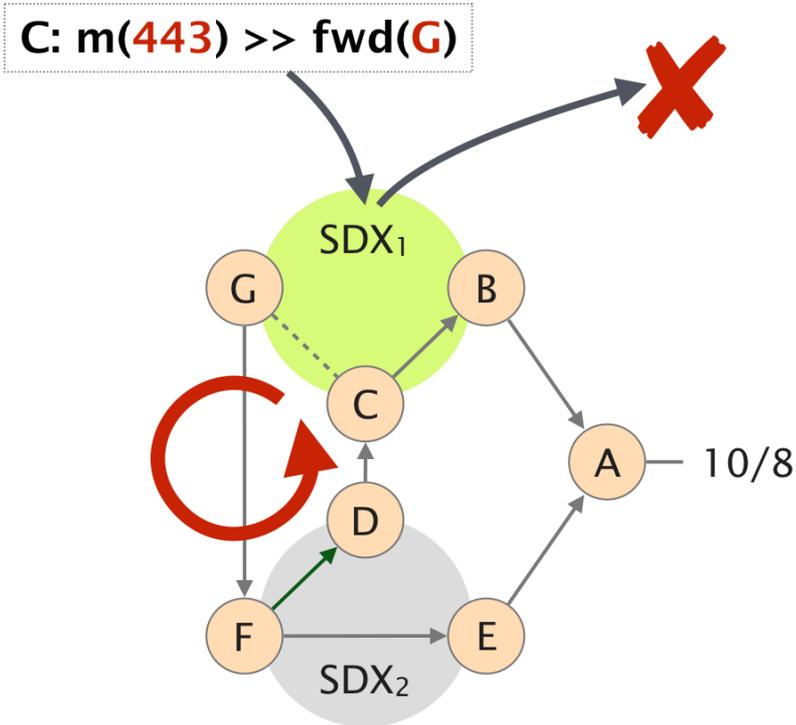
	Prefix	Deflection Set
C	10/8	{F, D}
...	...	...

*BGP Paths @ SDX<sub>1</sub>*

	Prefix	AS Path
C	10/8	[G, F, E, A]
...	...	...

# Safe Inter-domain Deflection Routing

(4) *Reject*



*Deflection DB @ SDX<sub>1</sub>*

	Prefix	Deflection Set
C	10/8	{F, D}
...	...	...

*BGP Best Path @ SDX<sub>1</sub>*

	Prefix	AS Path
C	10/8	[G, F, E, A]
...	...	...

# SIDR works in practice

# SIDR works in practice

Dataset

CAIDA AS graph (50k nodes, 200k edges)

Combined IXP dataset (421 IXPs)

# SIDR works in practice

Dataset	CAIDA AS graph (50k nodes, 200k edges) Combined IXP dataset (421 IXPs)
Methodology	Augmented with IXP links (1M edges) Compute BGP paths to 1000 destinations Generate policies using iSDX's model Install one policy after another

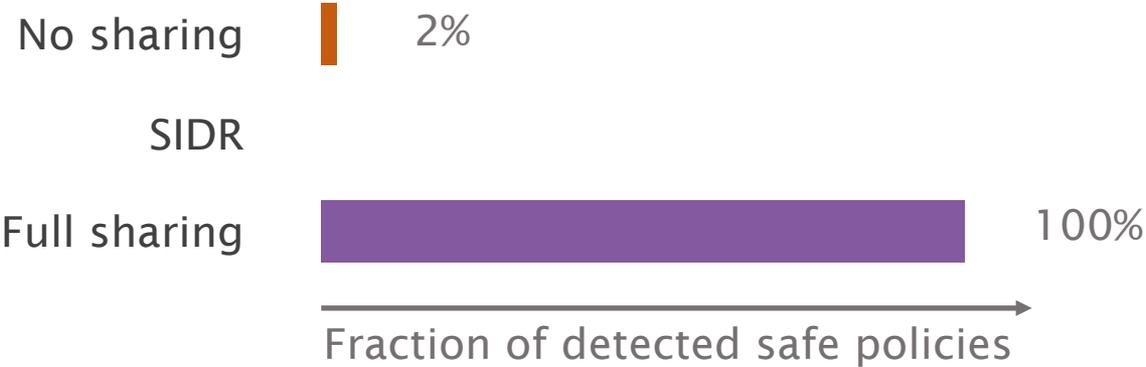
# SIDR is able to activate a majority of the safe policies



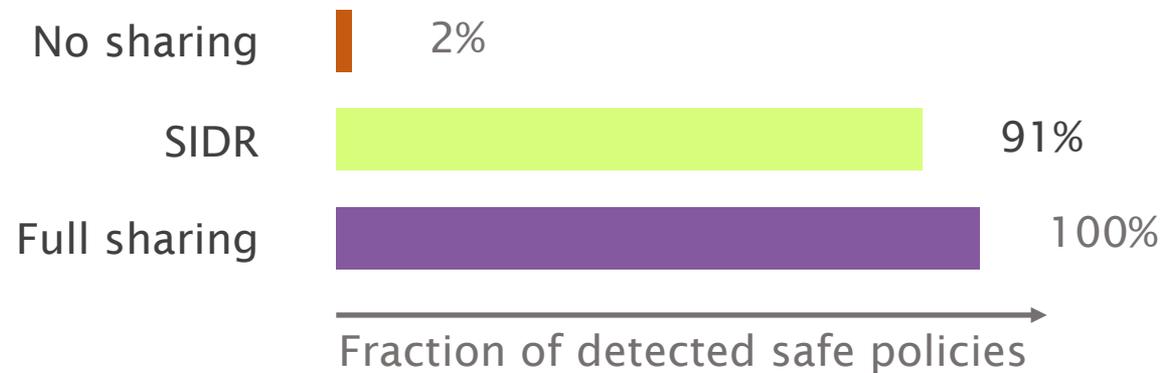
# SIDR is able to activate a majority of the safe policies



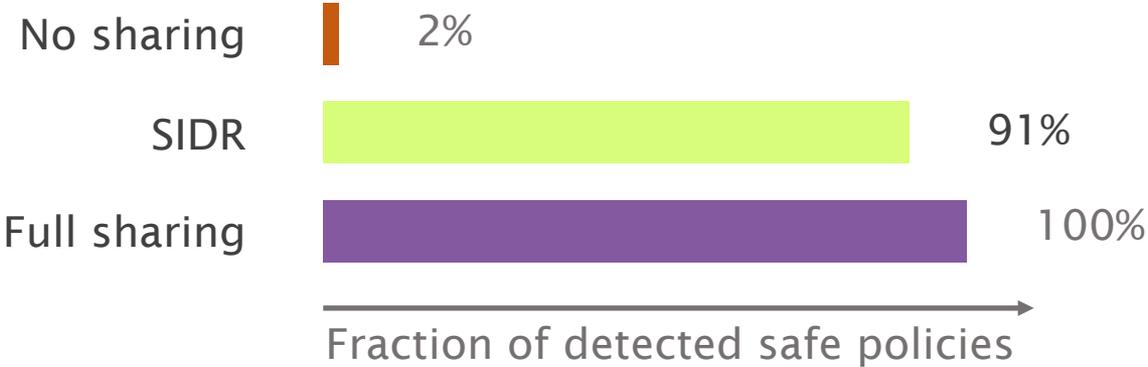
# SIDR is able to activate a majority of the safe policies



# SIDR is able to activate a majority of the safe policies



# SIDR is able to activate a majority of the safe policies

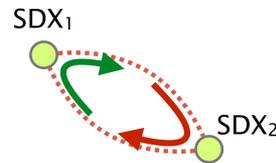


No false negatives, but false positives

SDX-Based Flexibility or Internet  
Correctness? Pick Two!

# SDX-Based Flexibility or Internet Correctness? Pick Two!

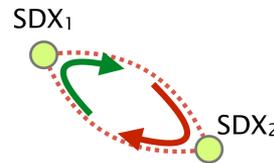
Part I



Multiple SDXes can create loops

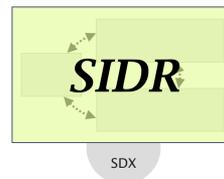
# SDX-Based Flexibility or Internet Correctness? Pick Two!

Part I



Multiple SDXes can create loops

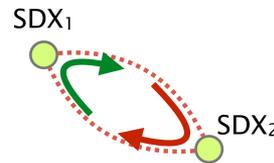
Part II



SIDR prevents these loops

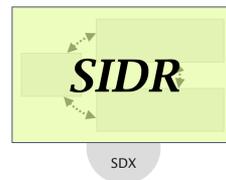
# SDX-Based Flexibility or Internet Correctness? Pick Two!

Part I



Multiple SDXes can create loops

Part II



SIDR prevents these loops

Part III



while maintaining flexibility.