



routing with

NLNETLABS

SURF Internet Measurements and Analysis Workshop 5 April 2023
Luuk Hendriks & Jasper den Hertog

WHAT DO WE DO?


CORE INTERNET

Development



Open Source Software

Standards (IETF)



Research

Governance related advisory

OPEN SOURCE SOFTWARE

soon-to-be super fast authoritative
DNS resolver



RPKI validator
("Relying Party Software")



multi-purpose
DNS resolver



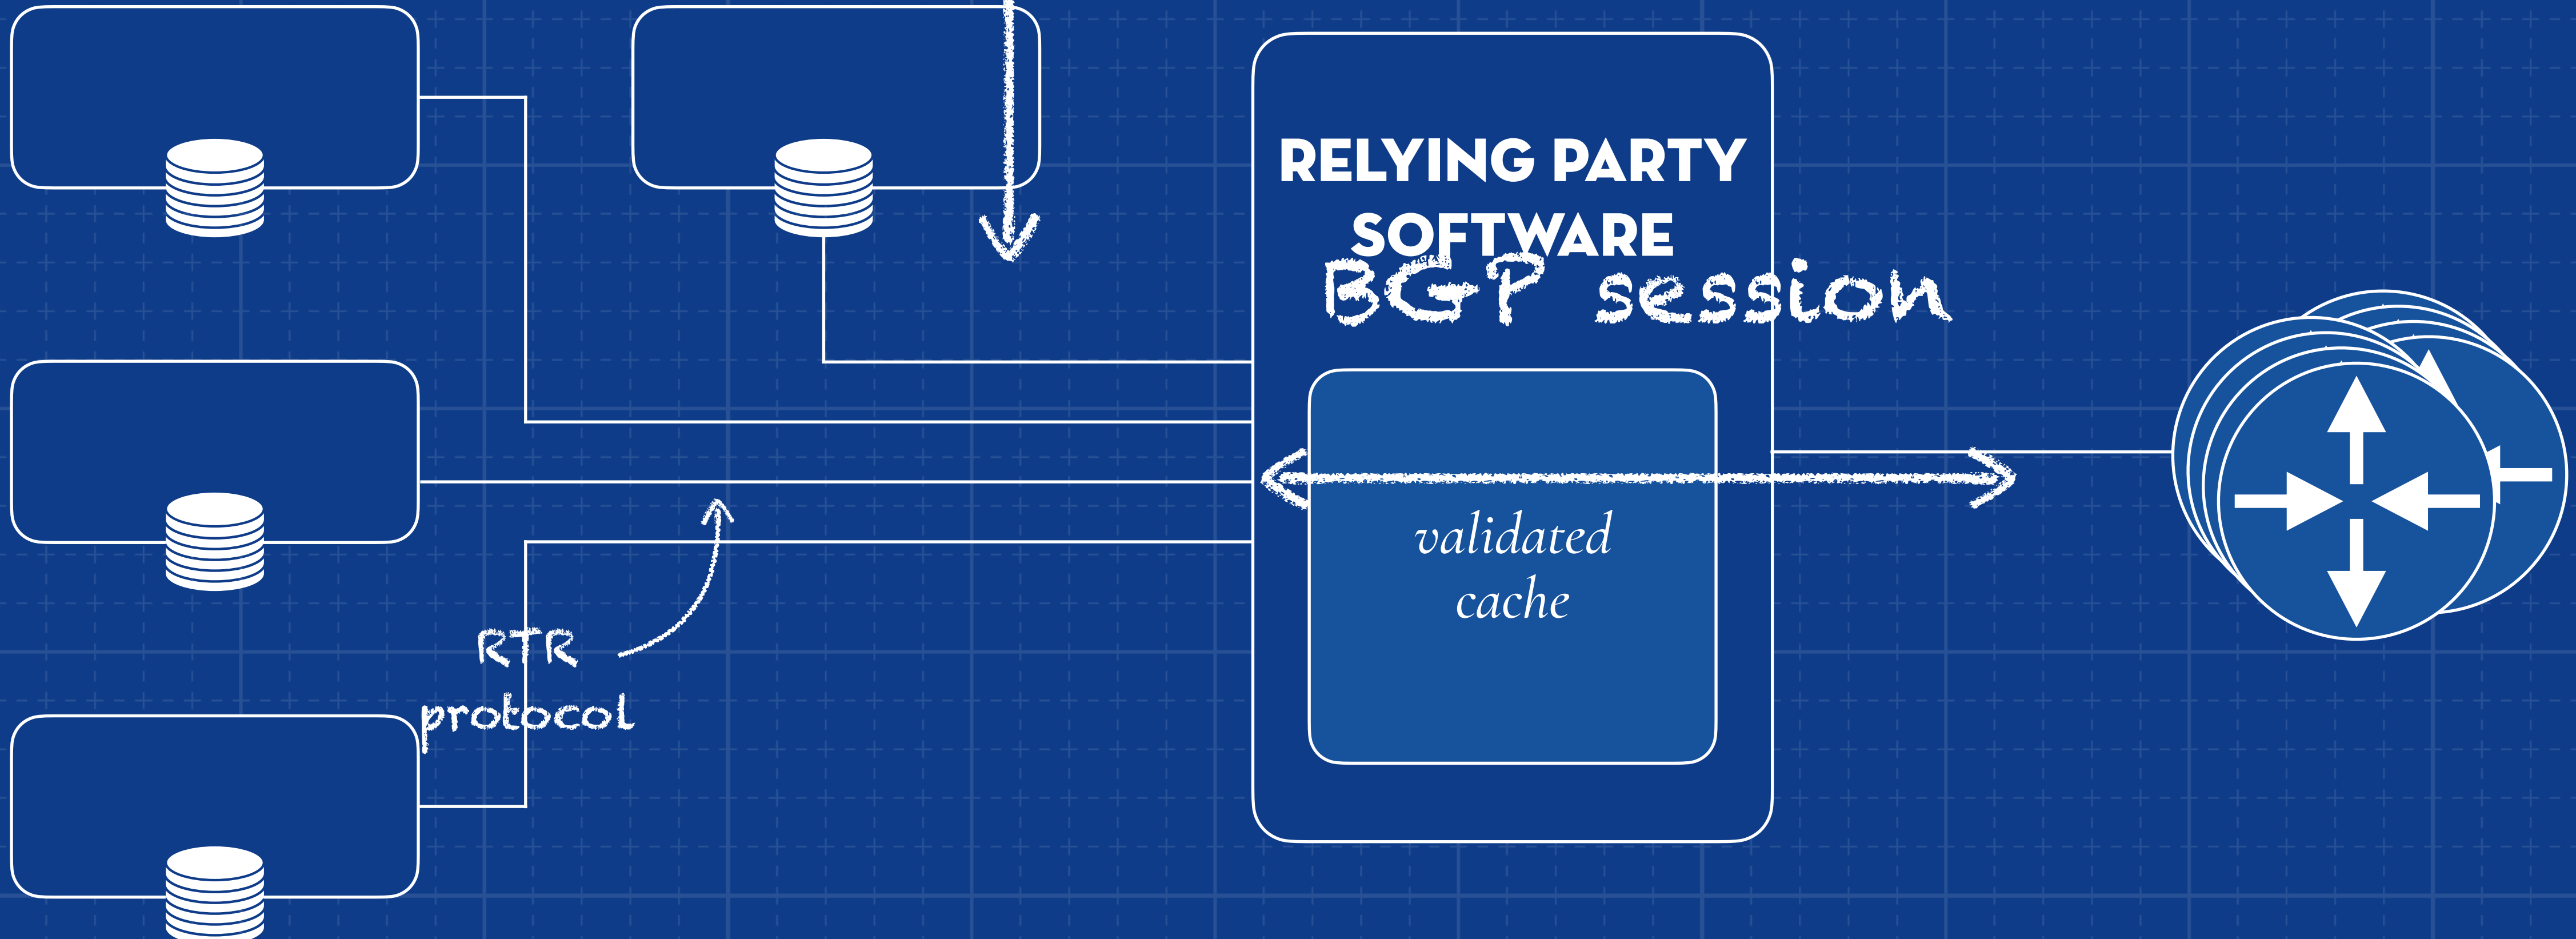
all singing and dancing
delegated RPKI



memory-safety
strongly, statically typed
raw performance (no garbage collection)
composability (dependencies, modules)

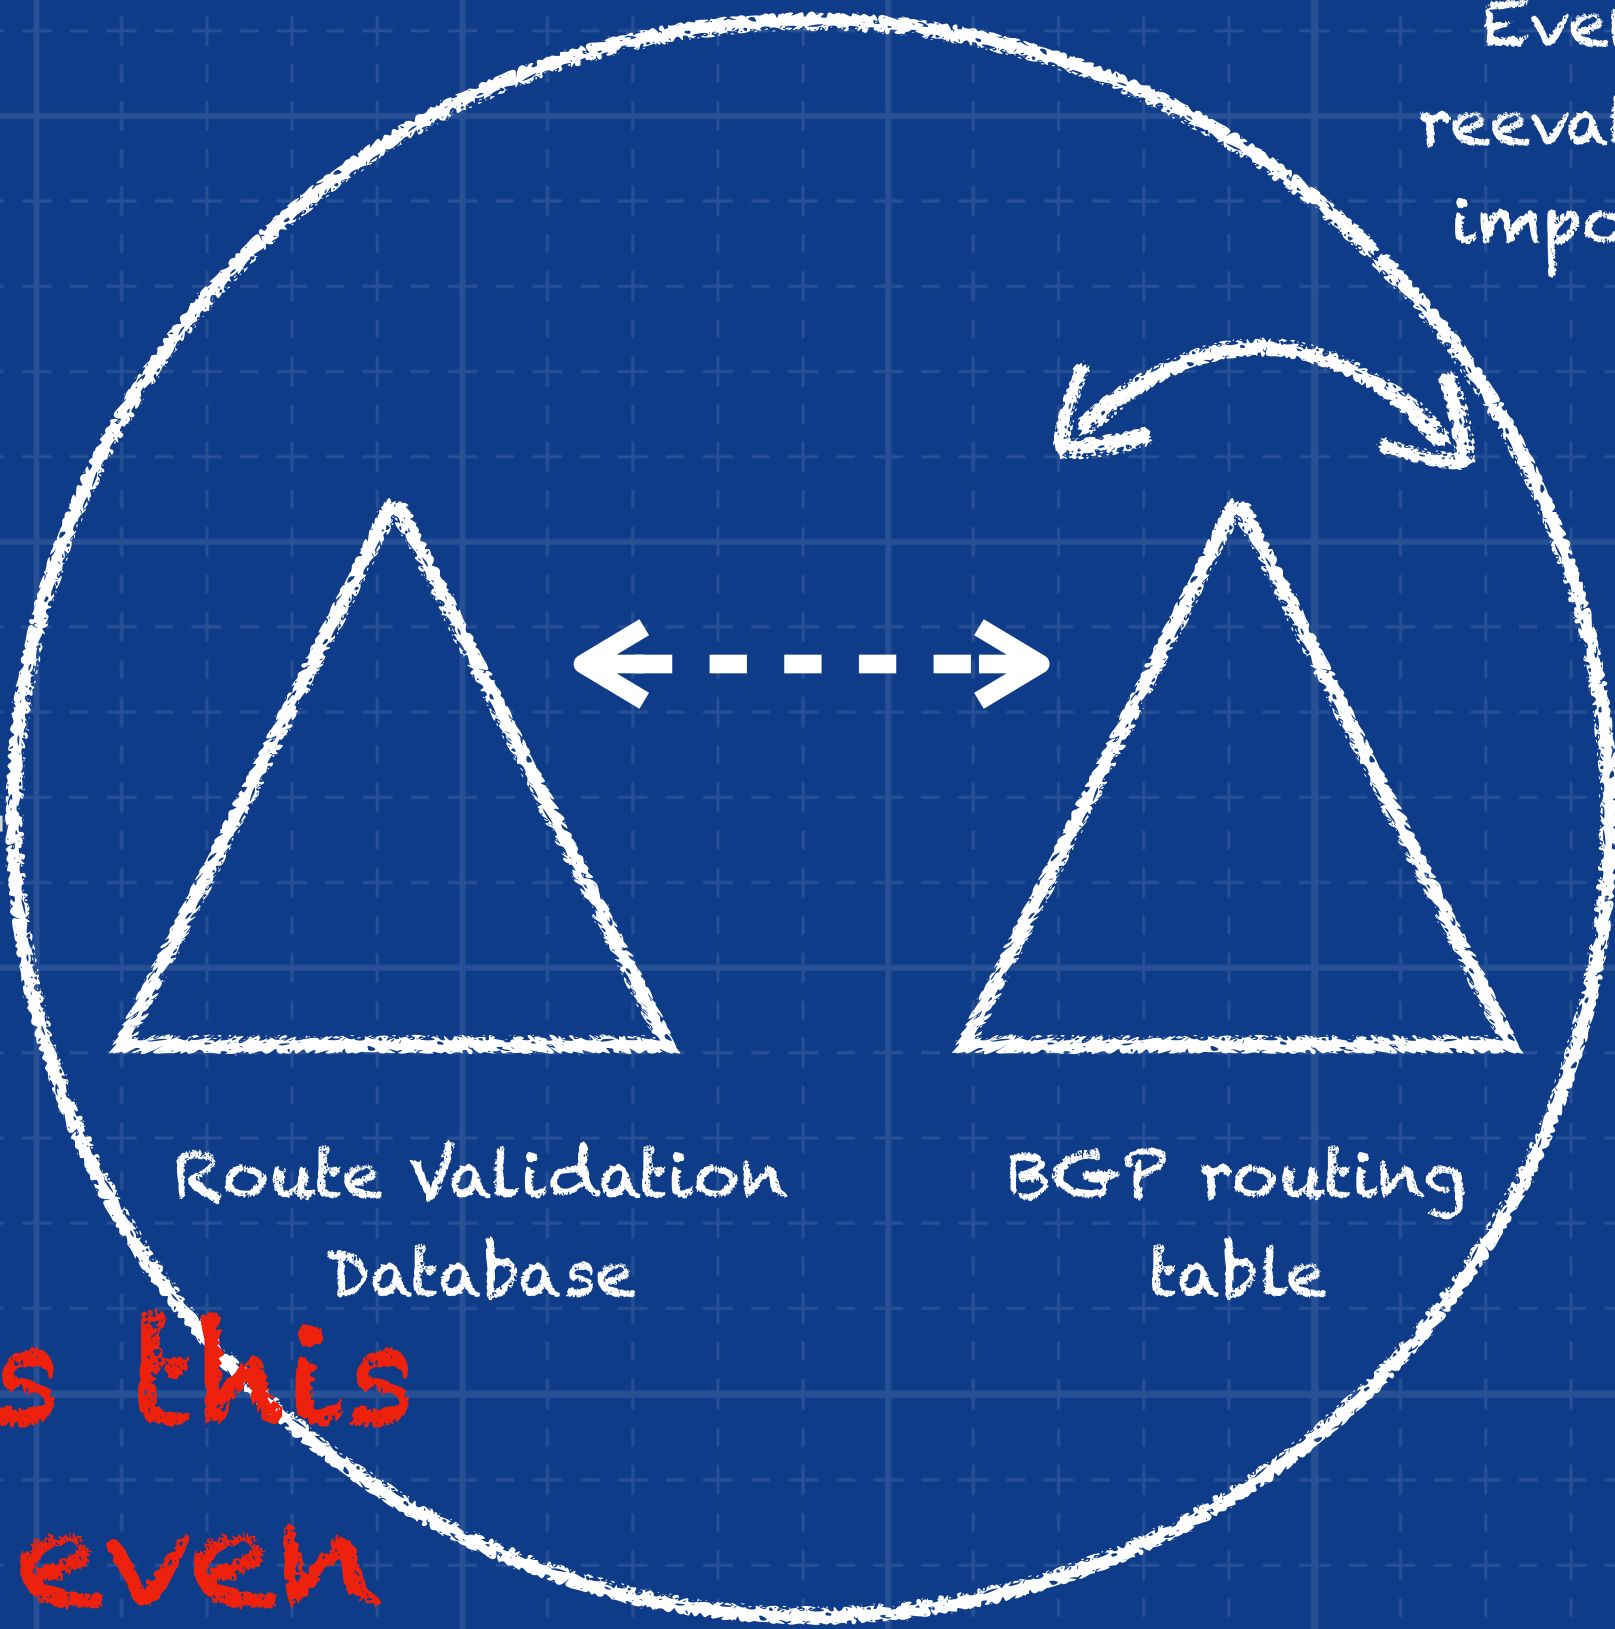
ROUTING
OBSERVABILITY

prefix filters



out of sync

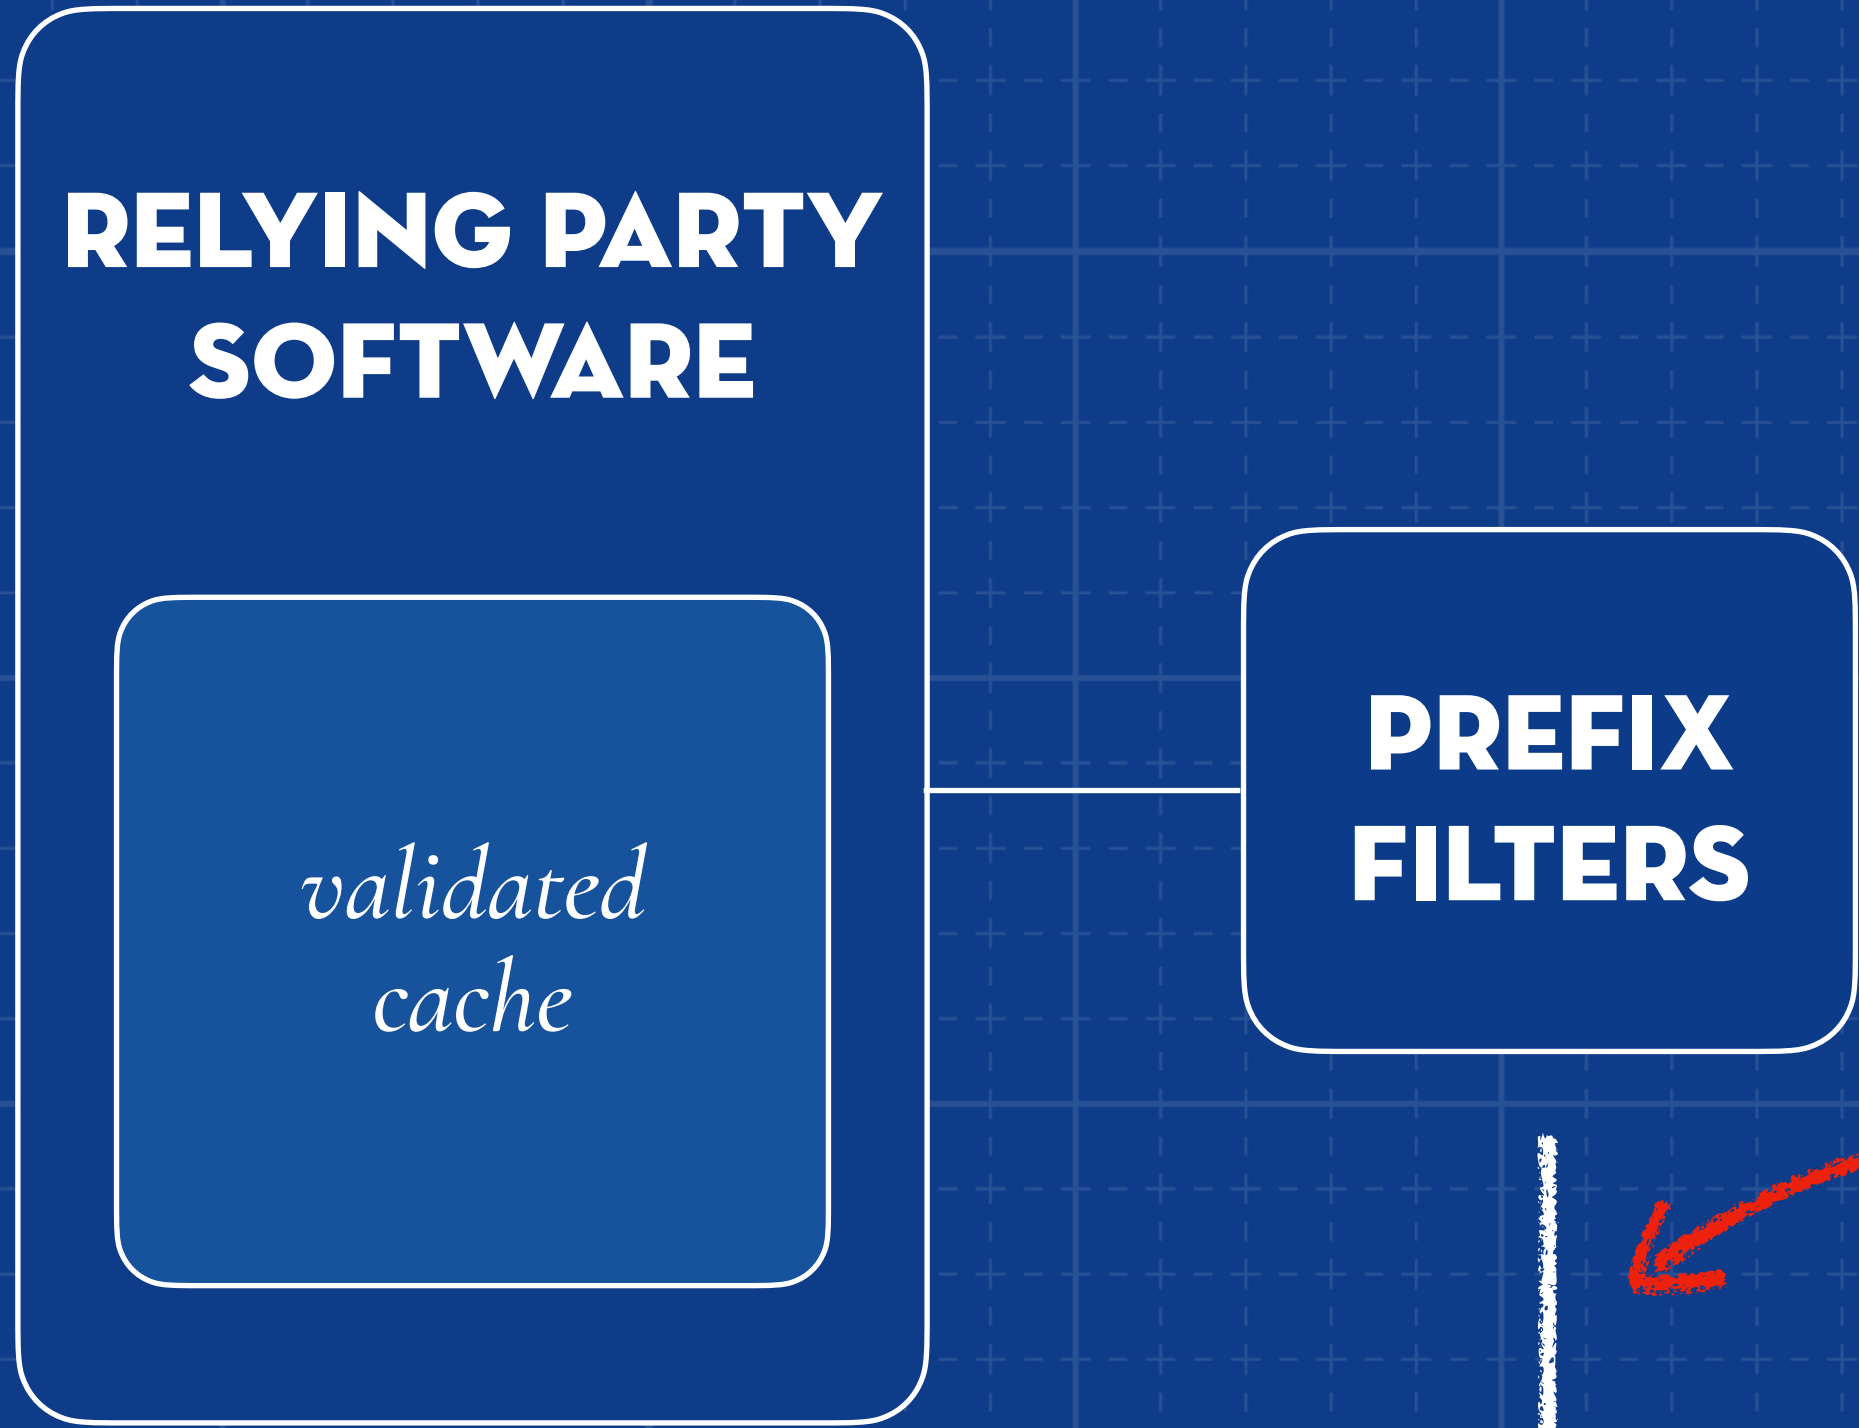
prefix filters



Event-based reevaluation of import policy


what is this thing even doing to my routes??!?

A SOLUTION

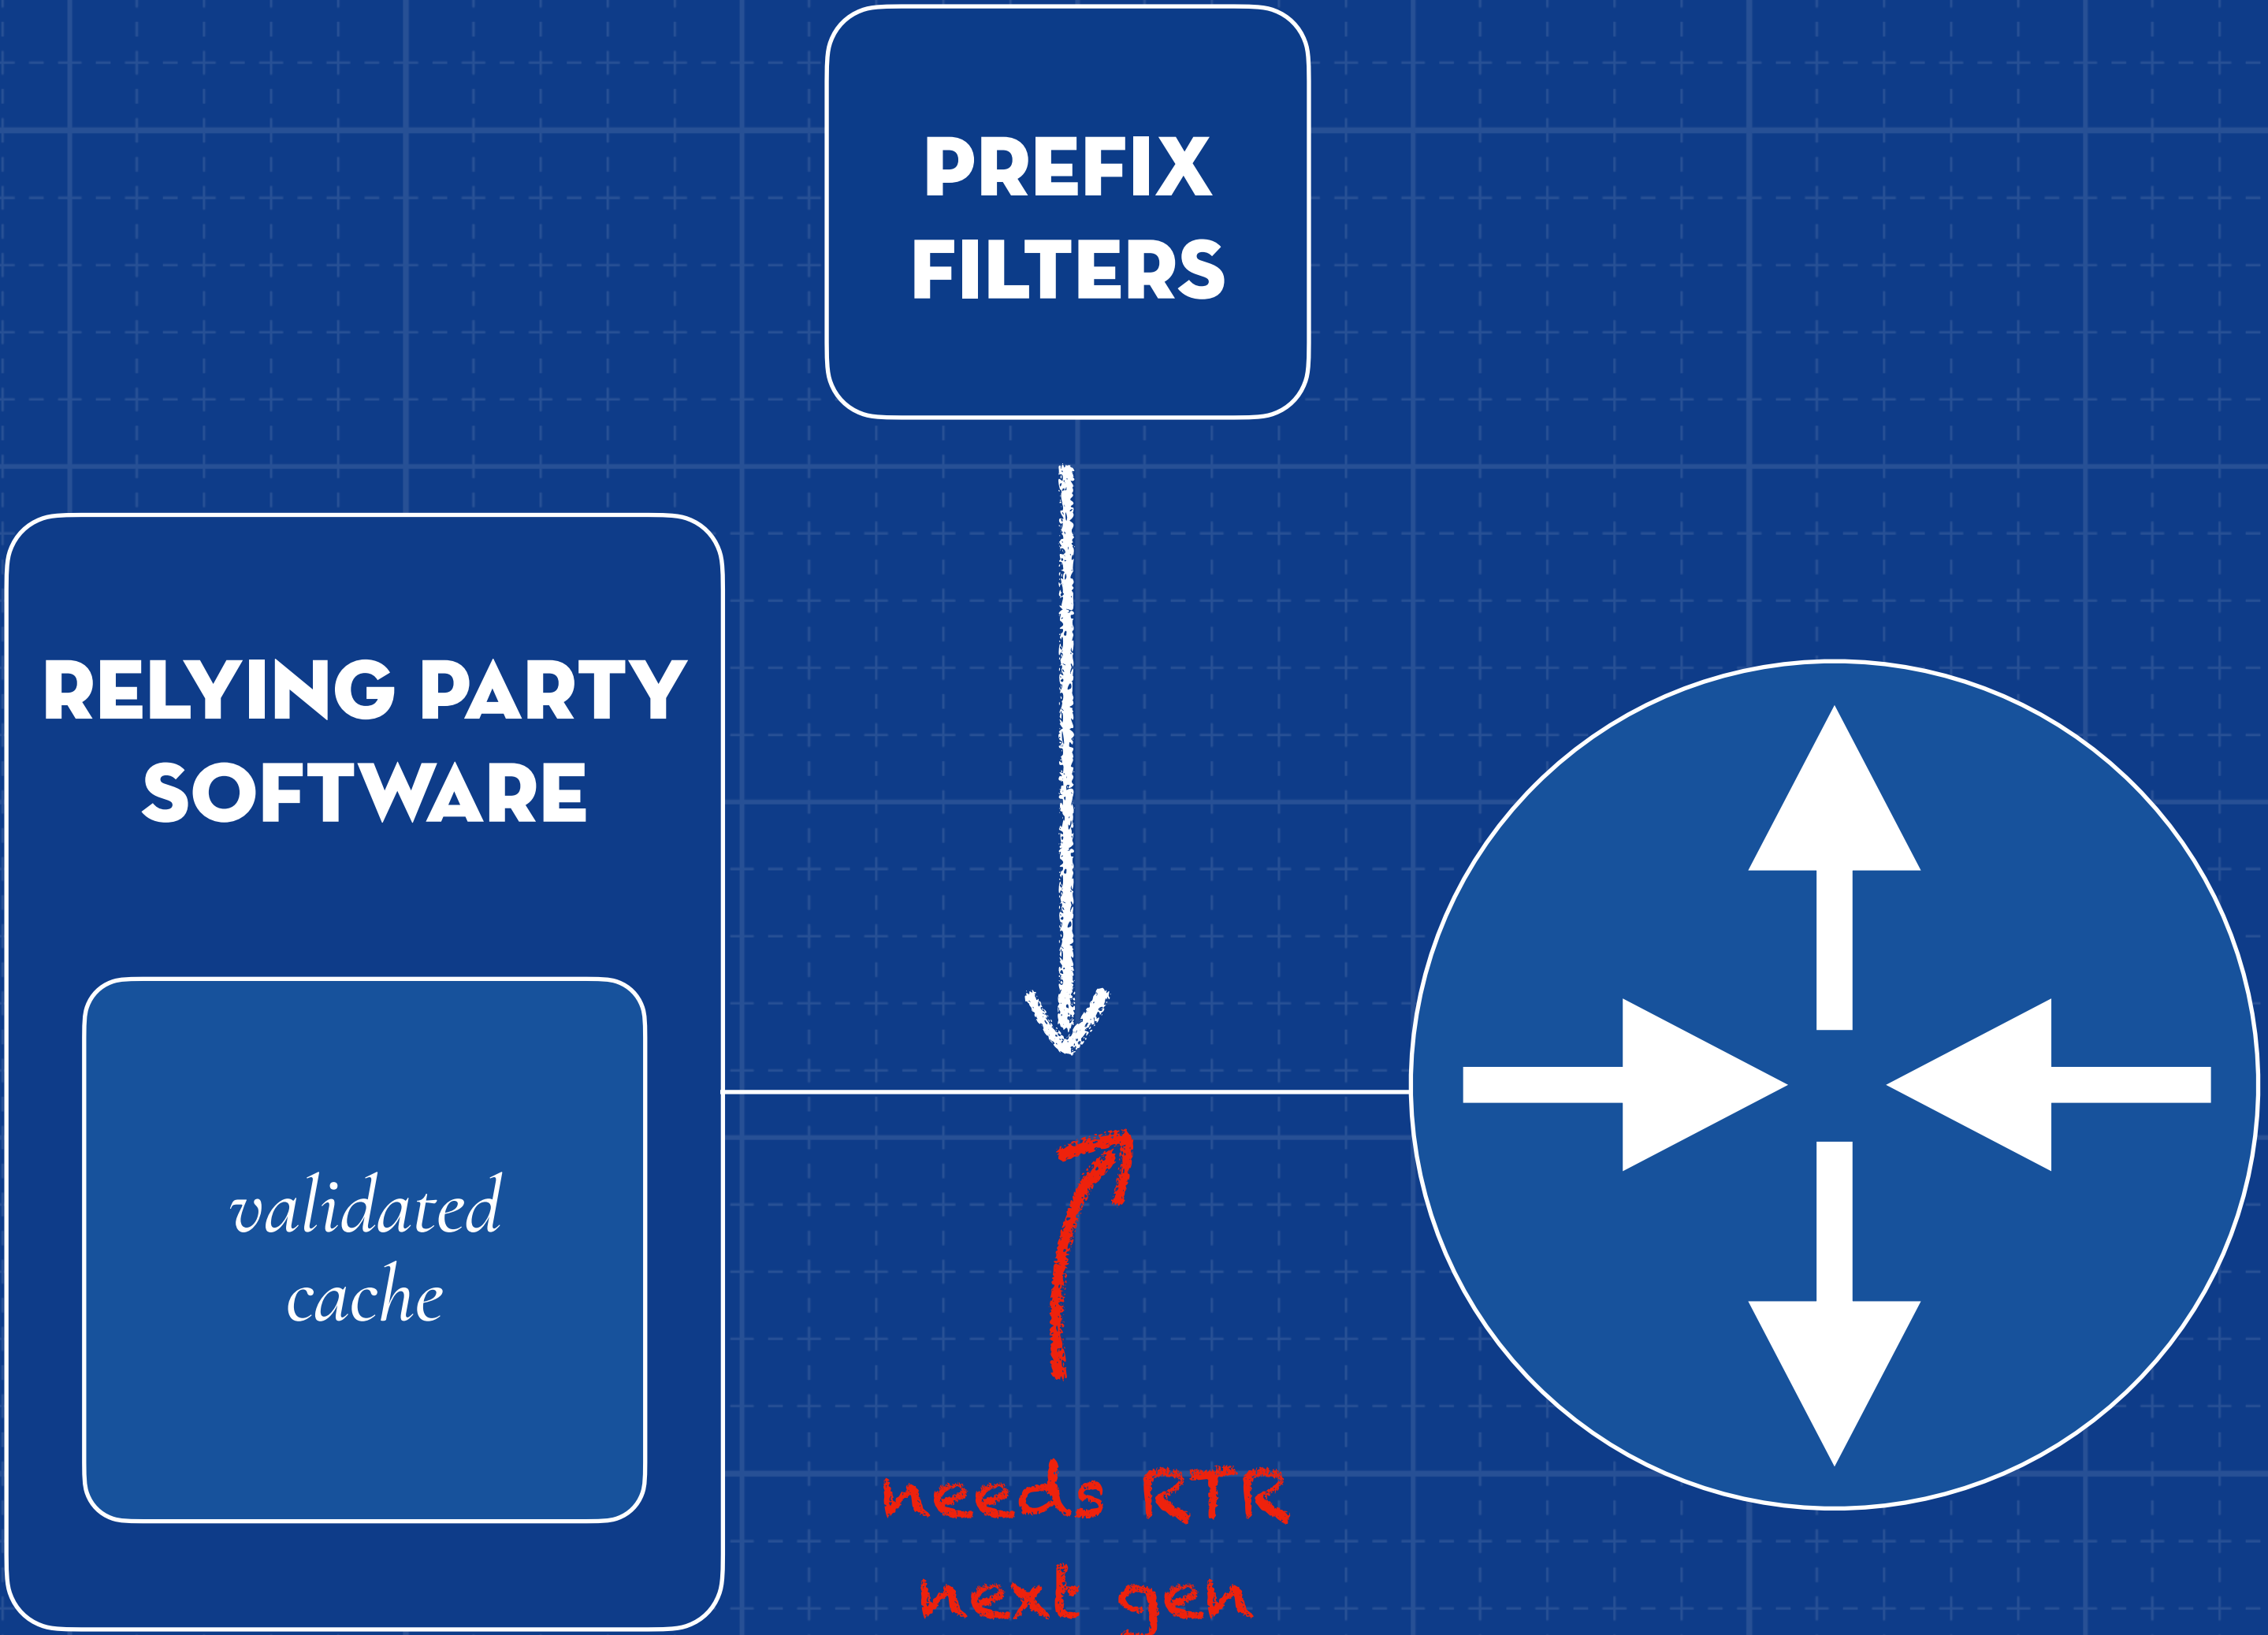


lots of processing & slow update cycle

A red curved arrow points from this text back to the 'PREFIX FILTERS' box.



A BETTER SOLUTION



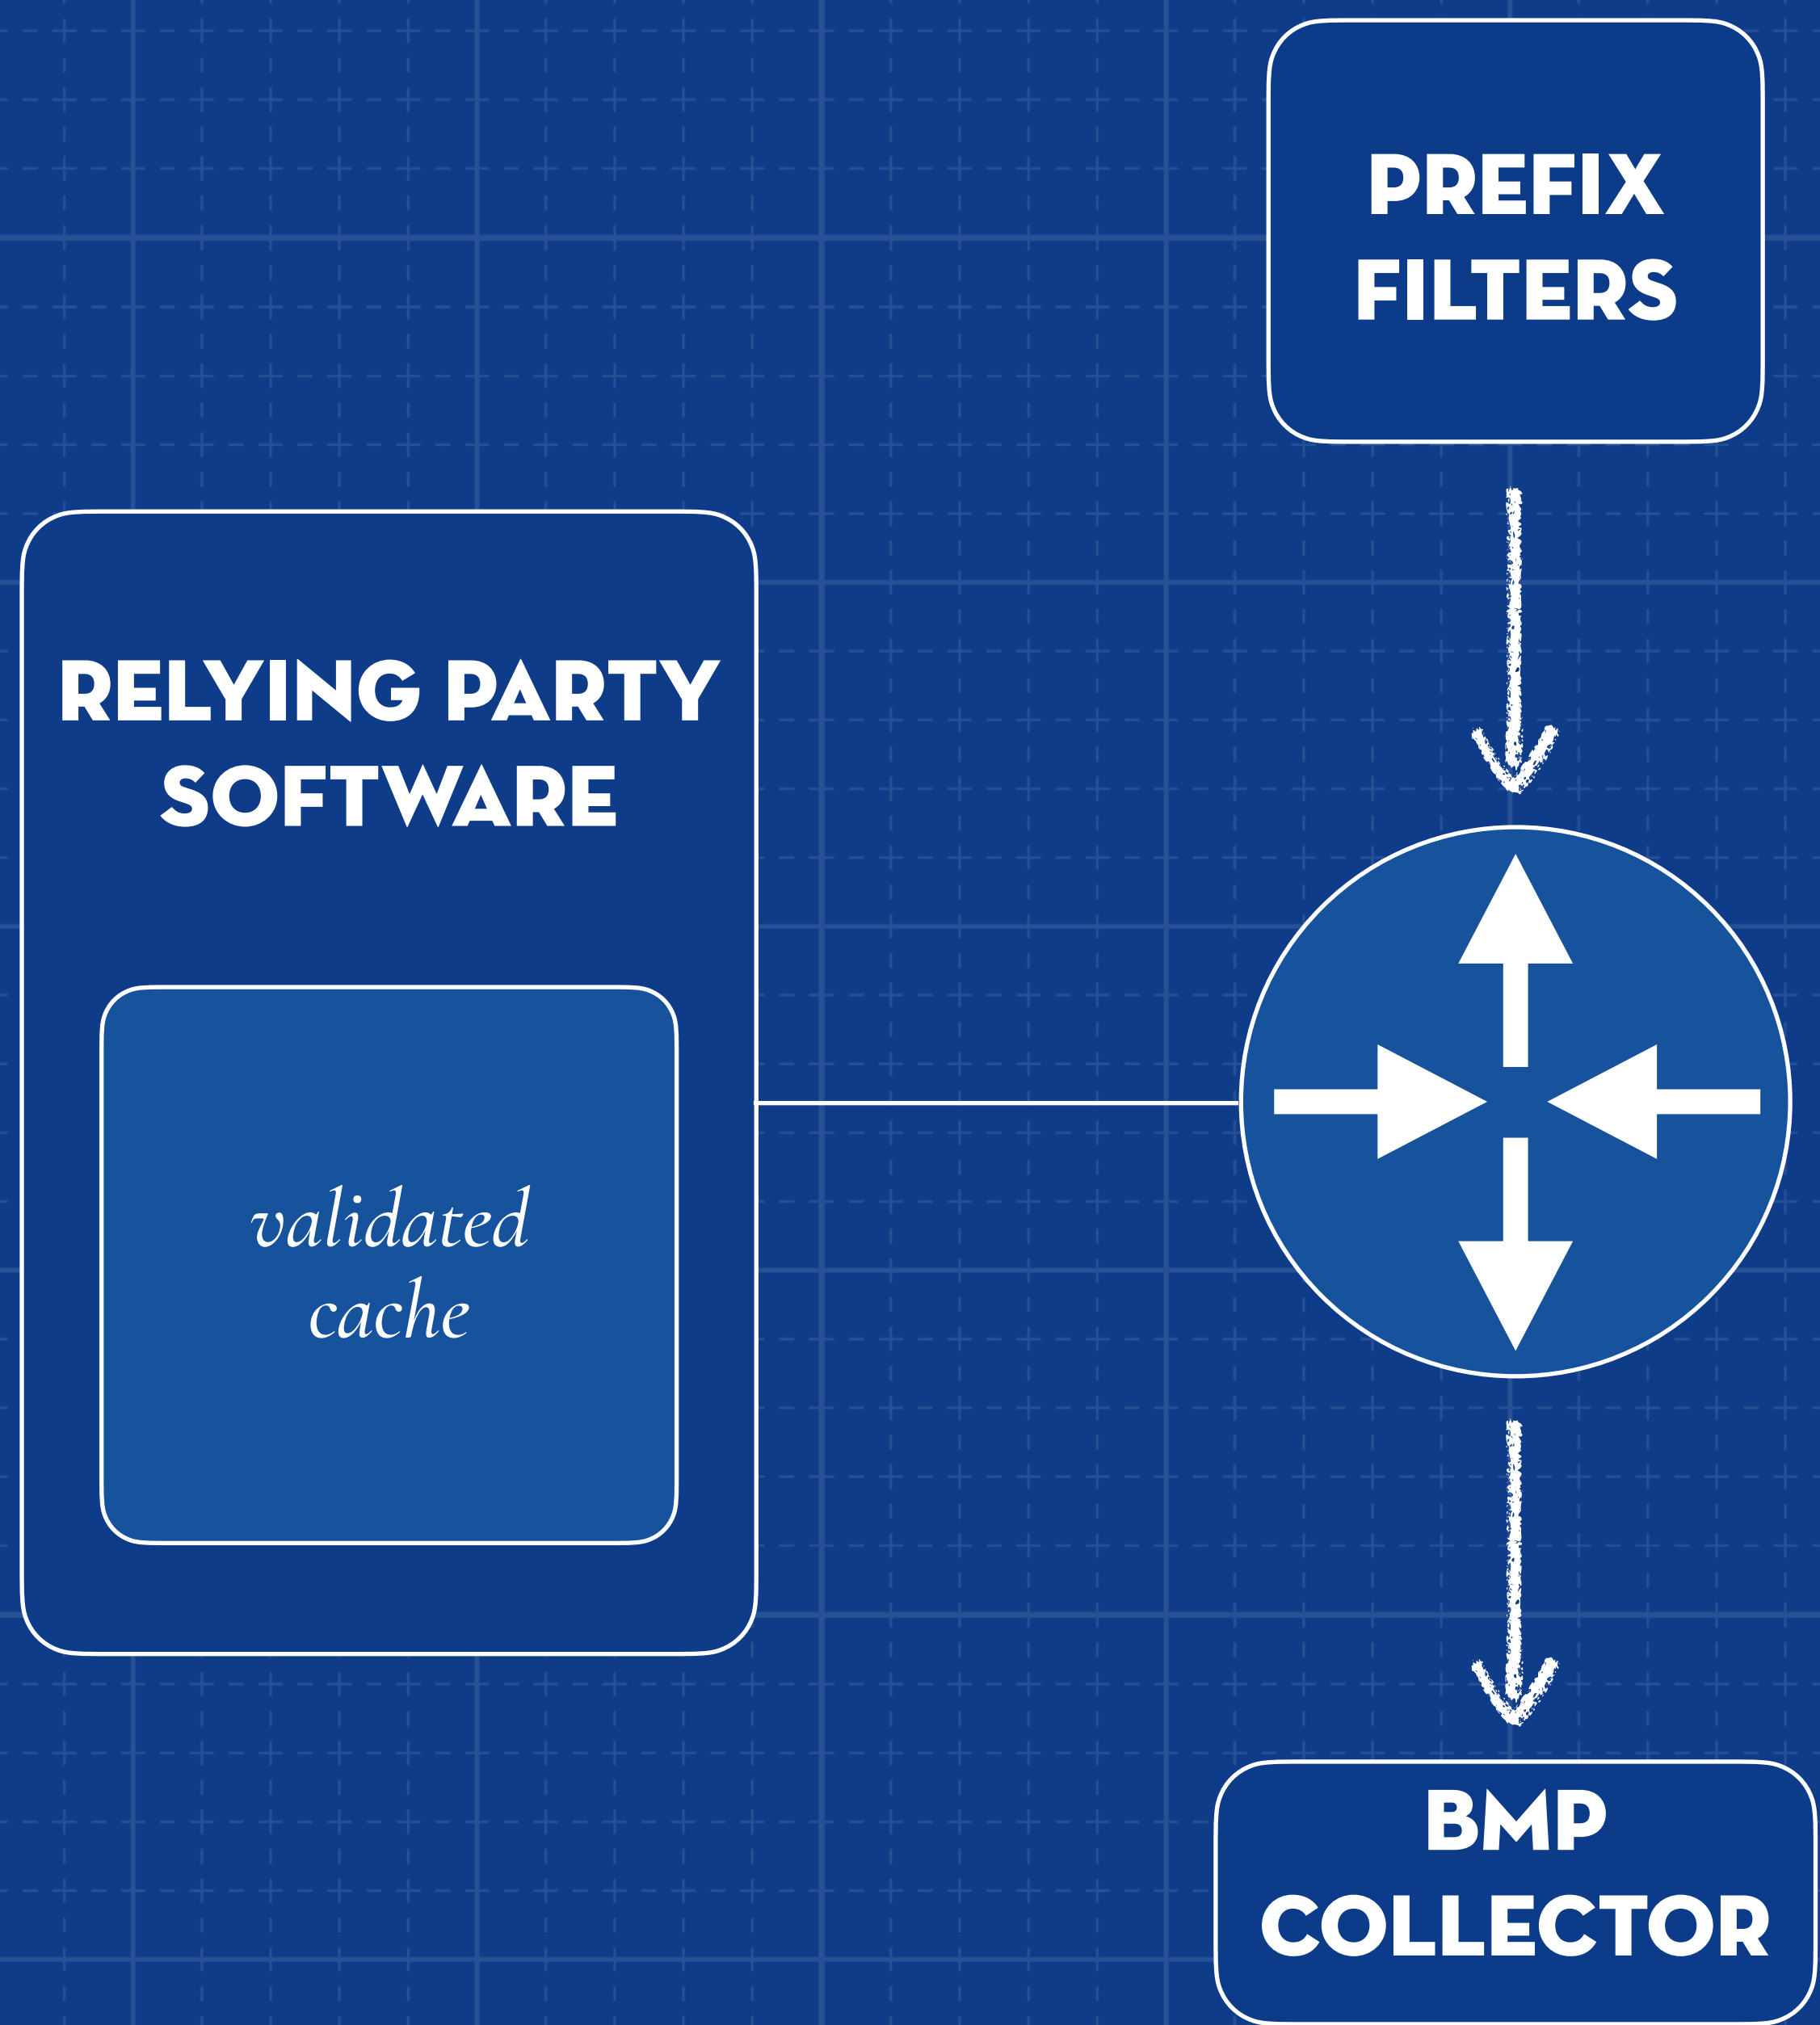
**PREFIX
FILTERS**

**RELYING PARTY
SOFTWARE**

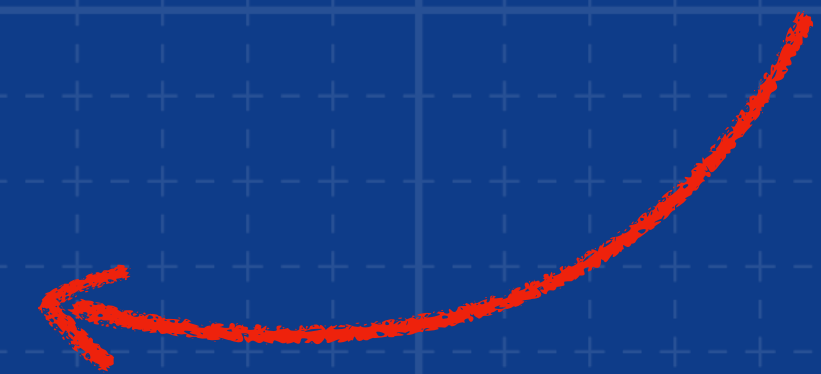
*validated
cache*

*needs RTR
next gen*

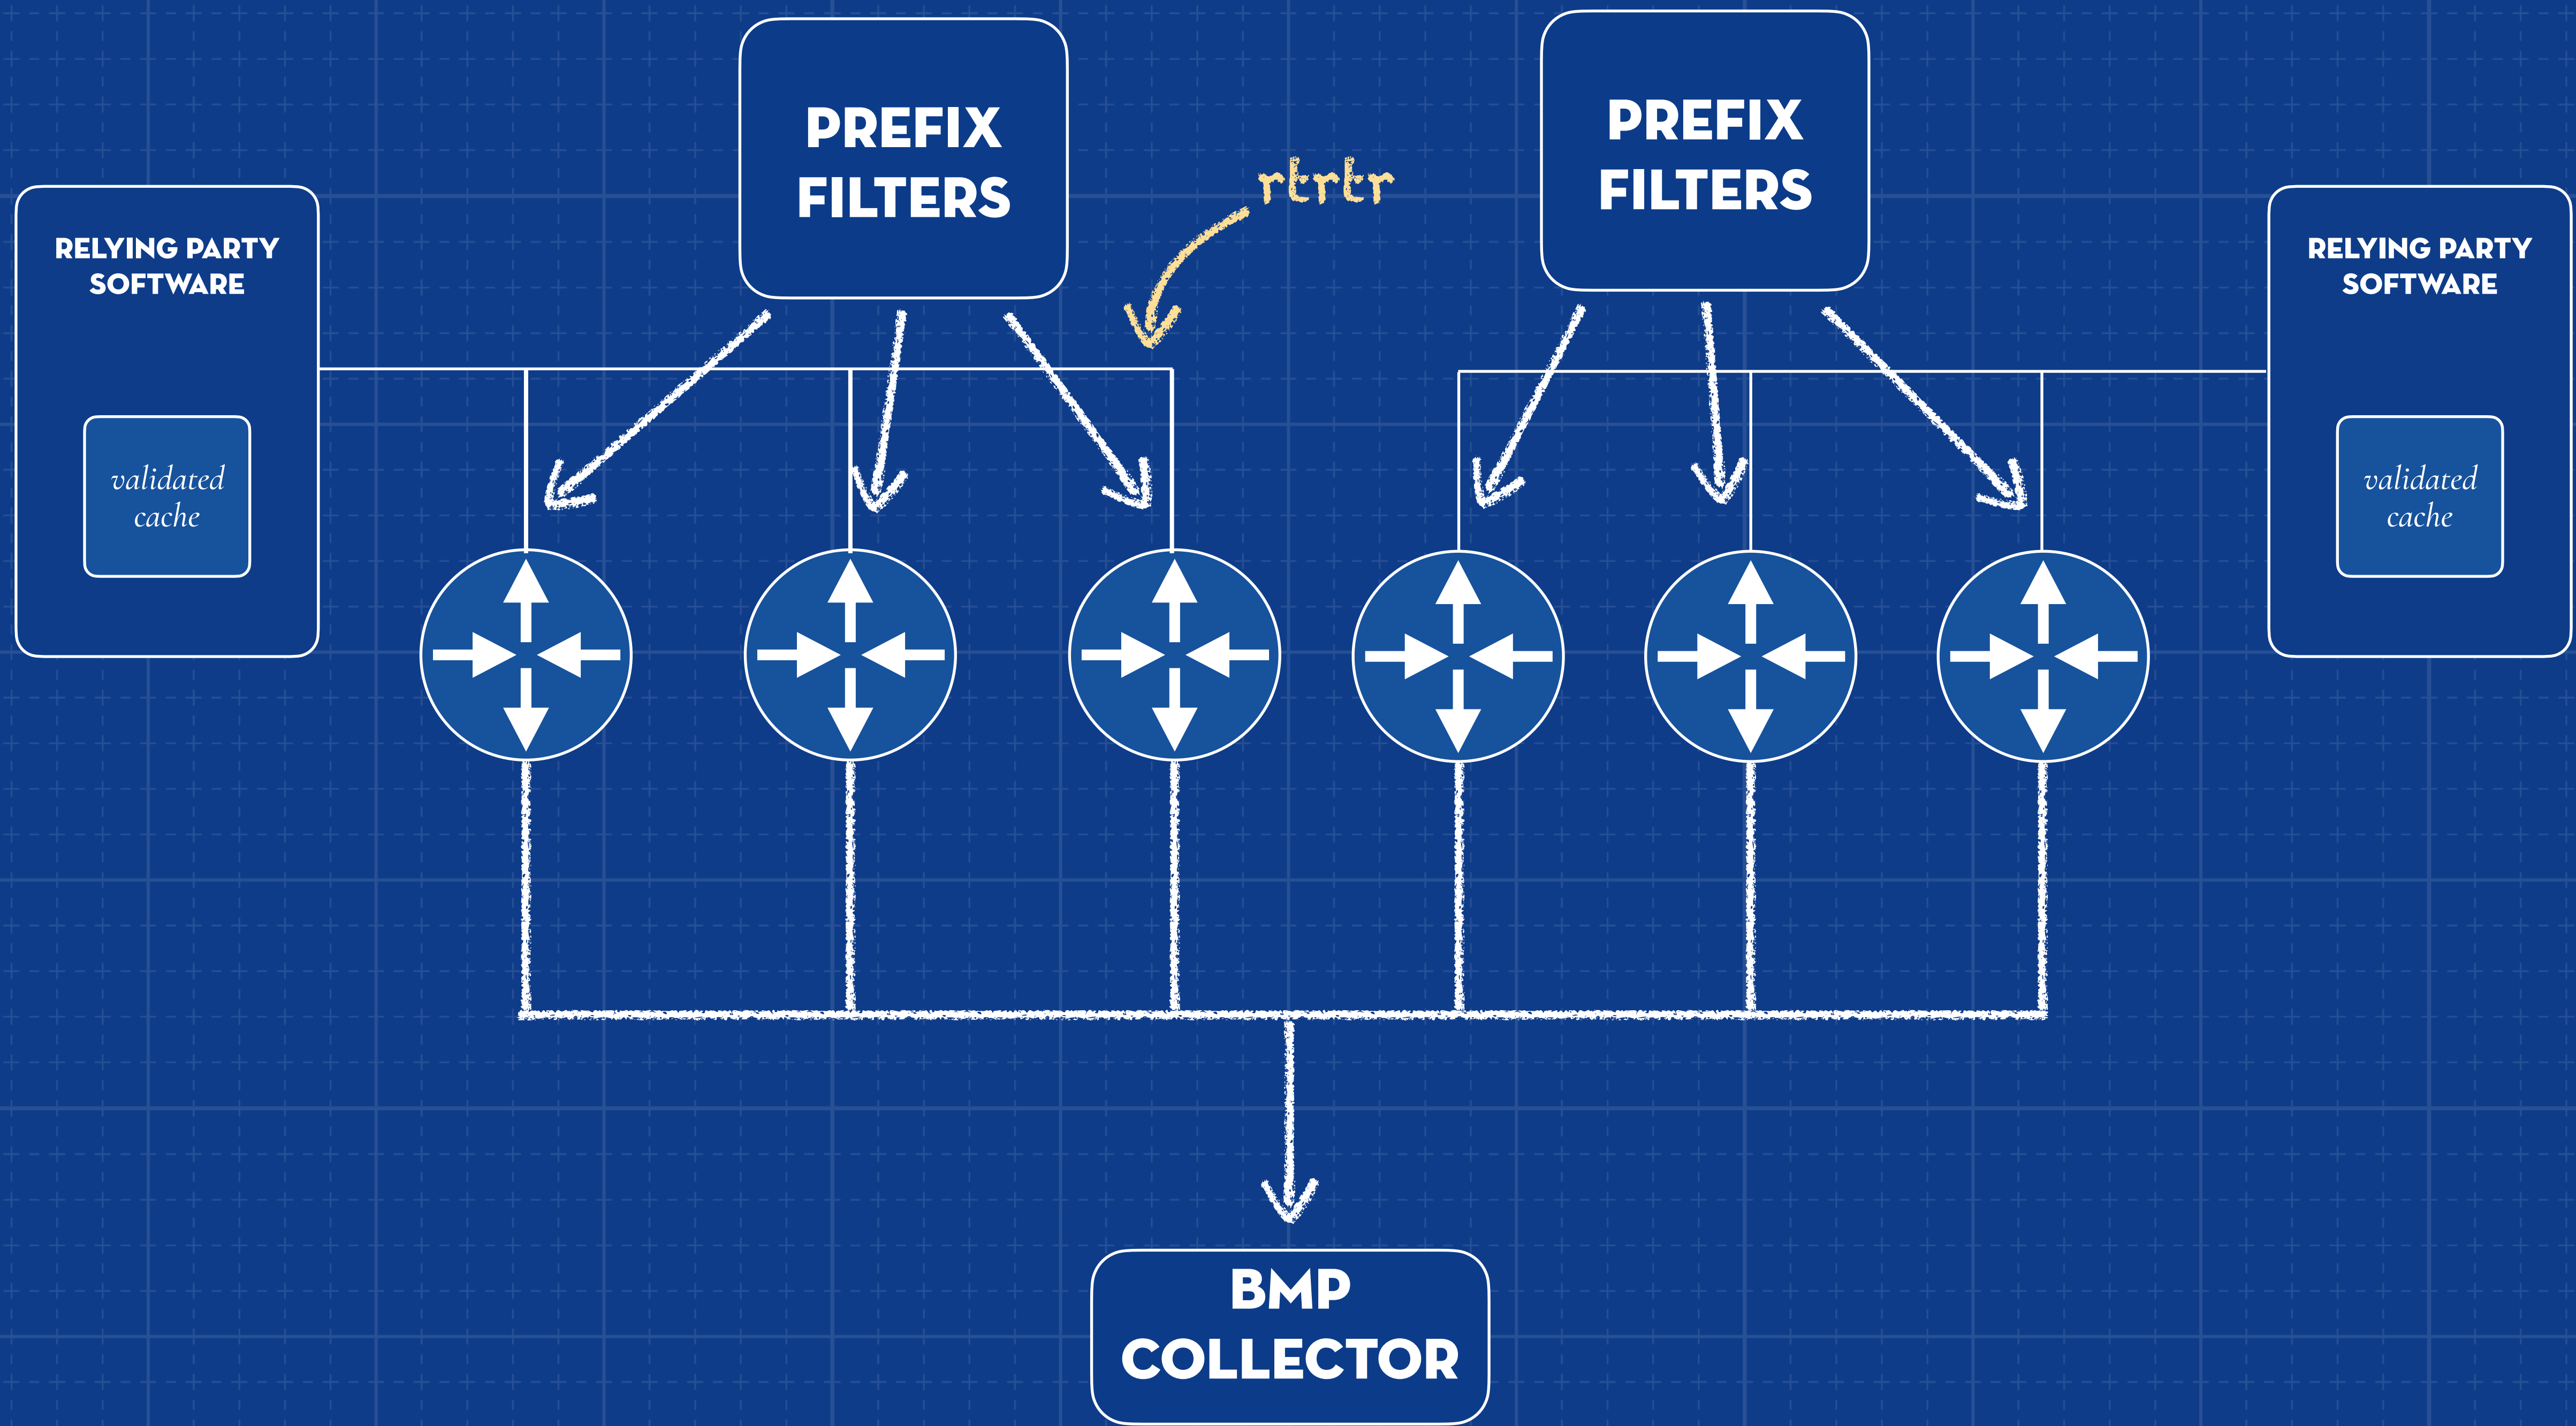
ANOTHER SOLUTION



needs a high-performance, scalable, real-time streaming, stateful collector







ANATOMY OF A COLLECTOR

COMPOSABLE RUNTIMES

combine into single runtime

multiple runtimes connected
by RTR-inspired protocol
(codename "rotoro")

builds on our RTRTR library

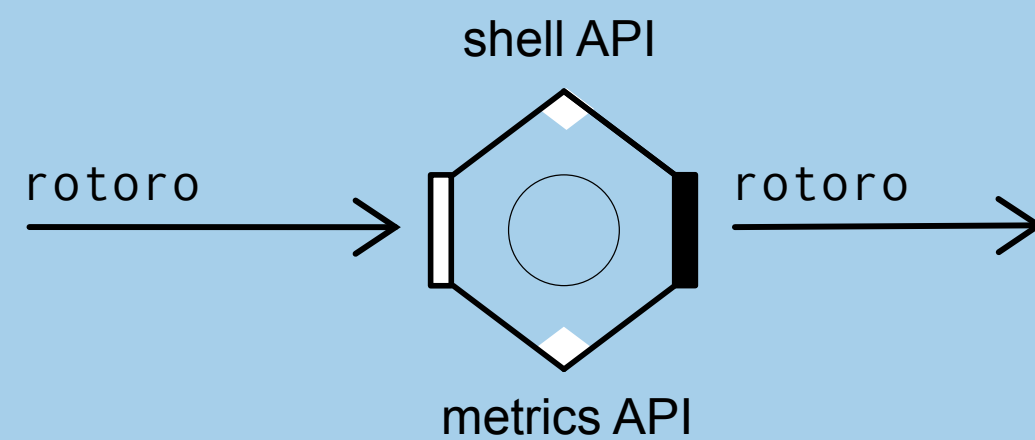
not only BMP, also BGP

many output formats,
streams, files, pub/sub

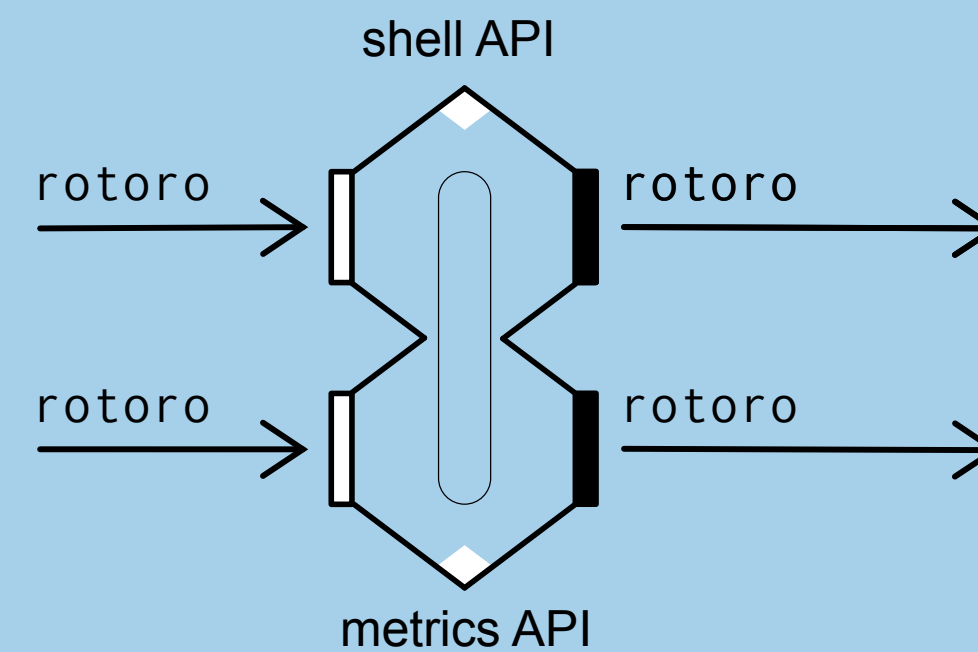
RUNTIMES

STORAGE

single input

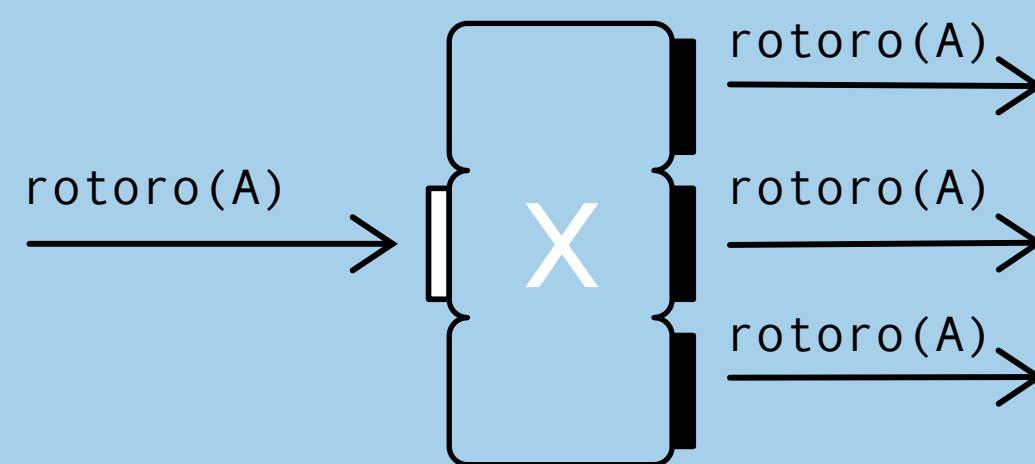


multiple input

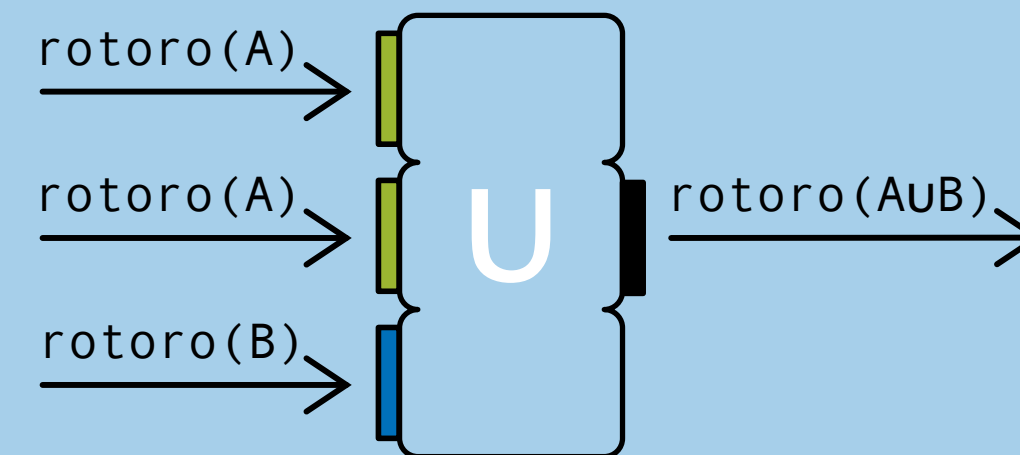


CHANNEL OPERATORS

broadcast




merge




TRANSFORMERS

input transformers



output transformers




TOWARDS ROUTE SERVICES

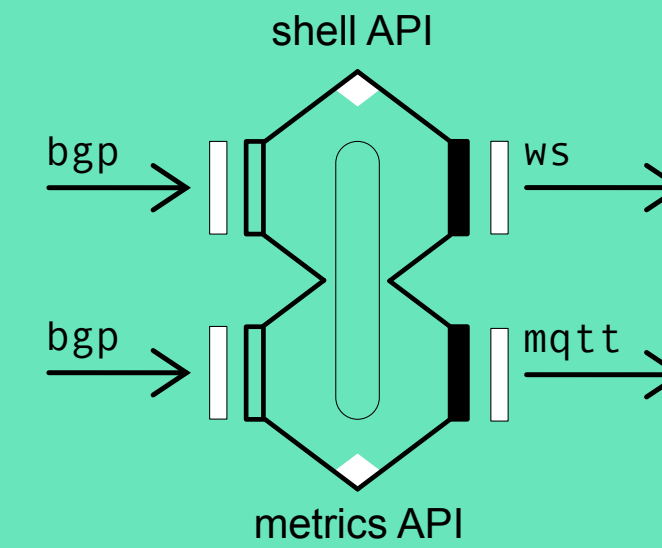
route collector
route monitor
(offline looking glass)
route server
route reflector

EXAMPLE CONFIGURATIONS

single runtime/single transformed output




single runtime/multiple input/two transformers

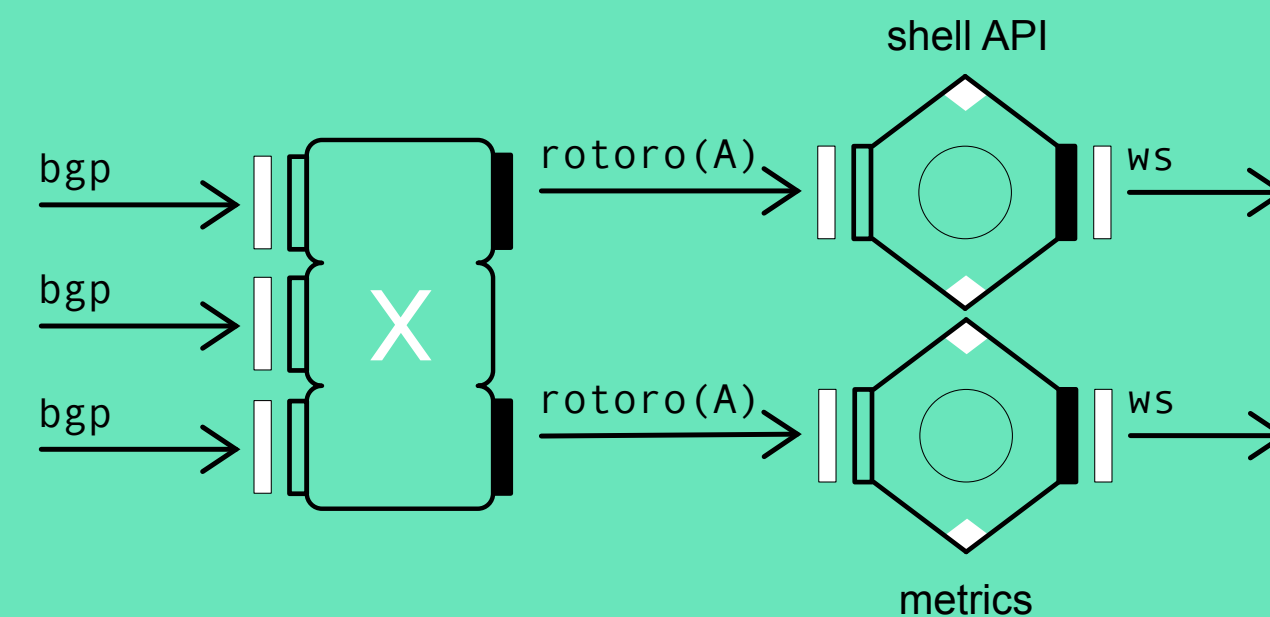


two runtimes

single input/two transformed outputs



redundant ingress



SINGLE-THREADED PREFIX STORE

STORE

(prefix, meta-data)
pairs with hierarchy

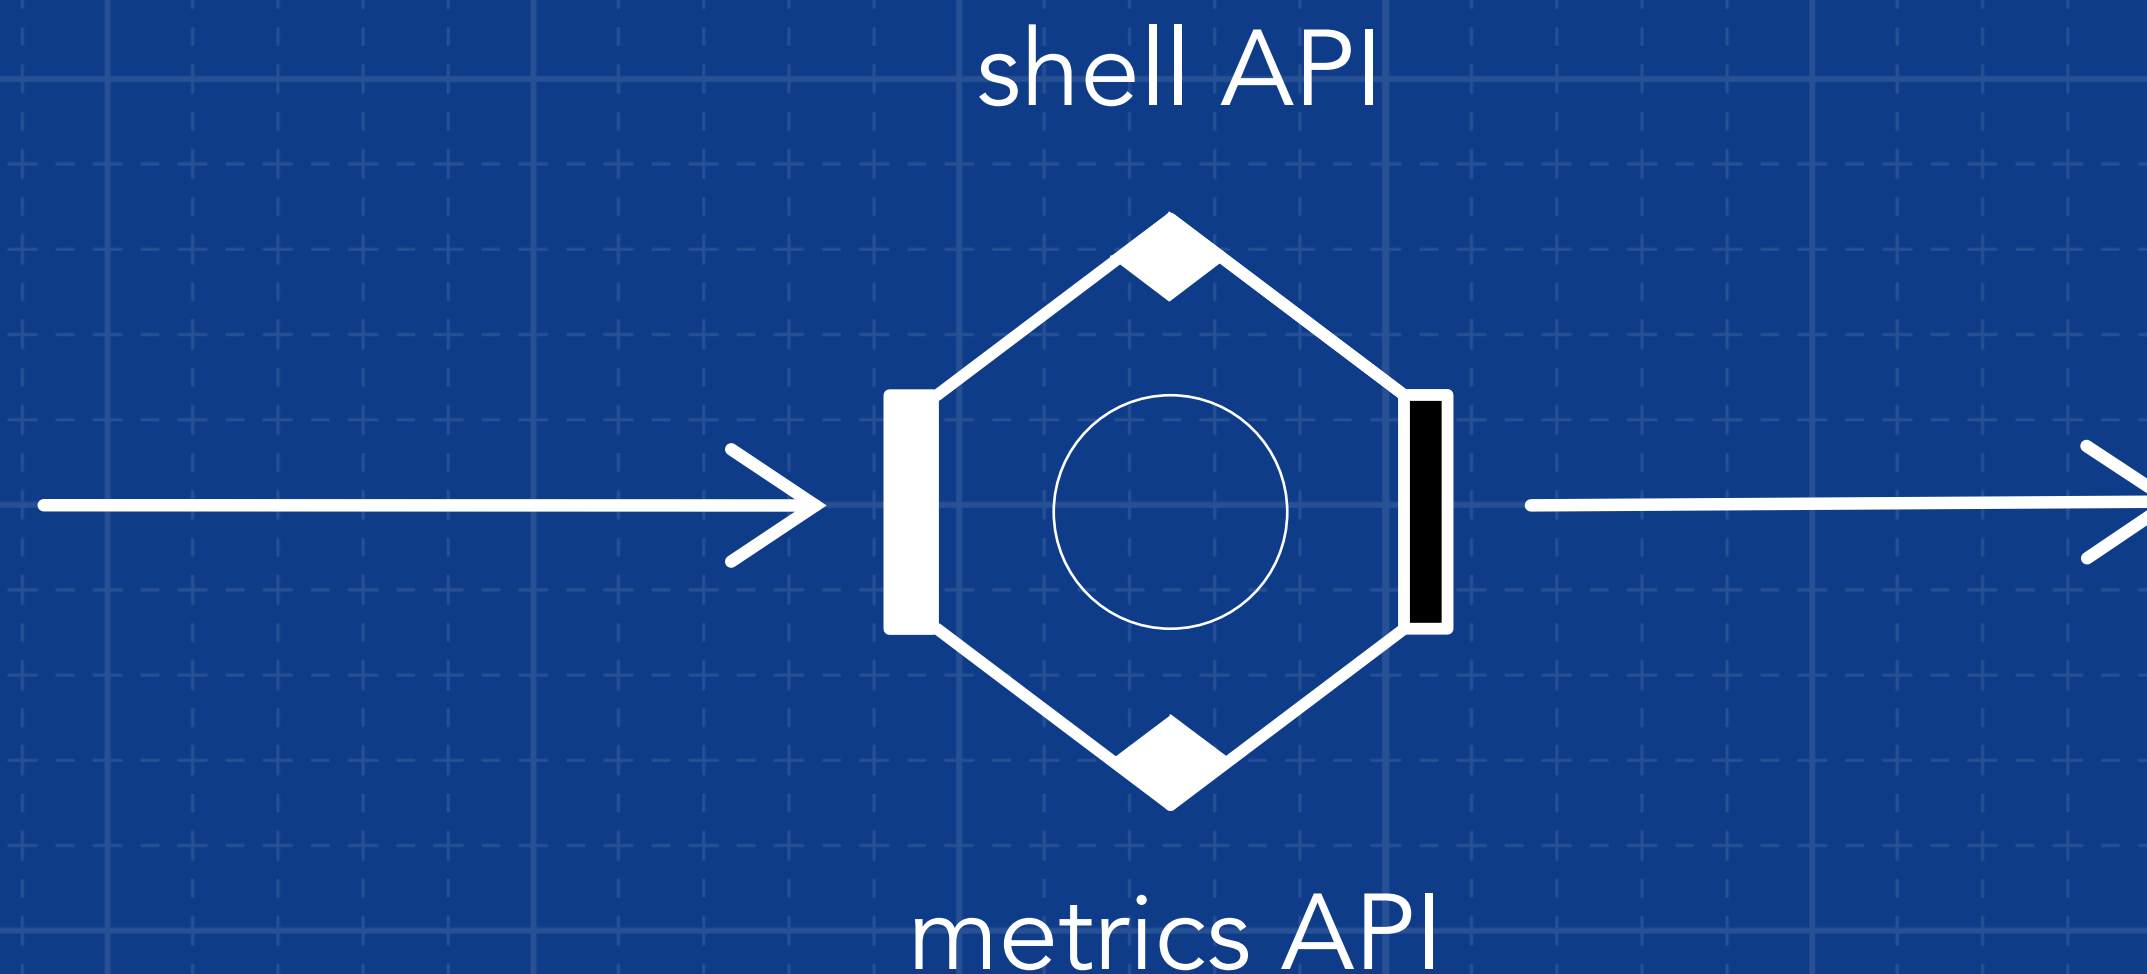
INPUTS

single input/
single type

OUTPUTS

single output/
single type

can be a stand-alone runtime



MULTI-THREADED PREFIX STORE

STORE

(prefix, meta-data)
pairs with hierarchy

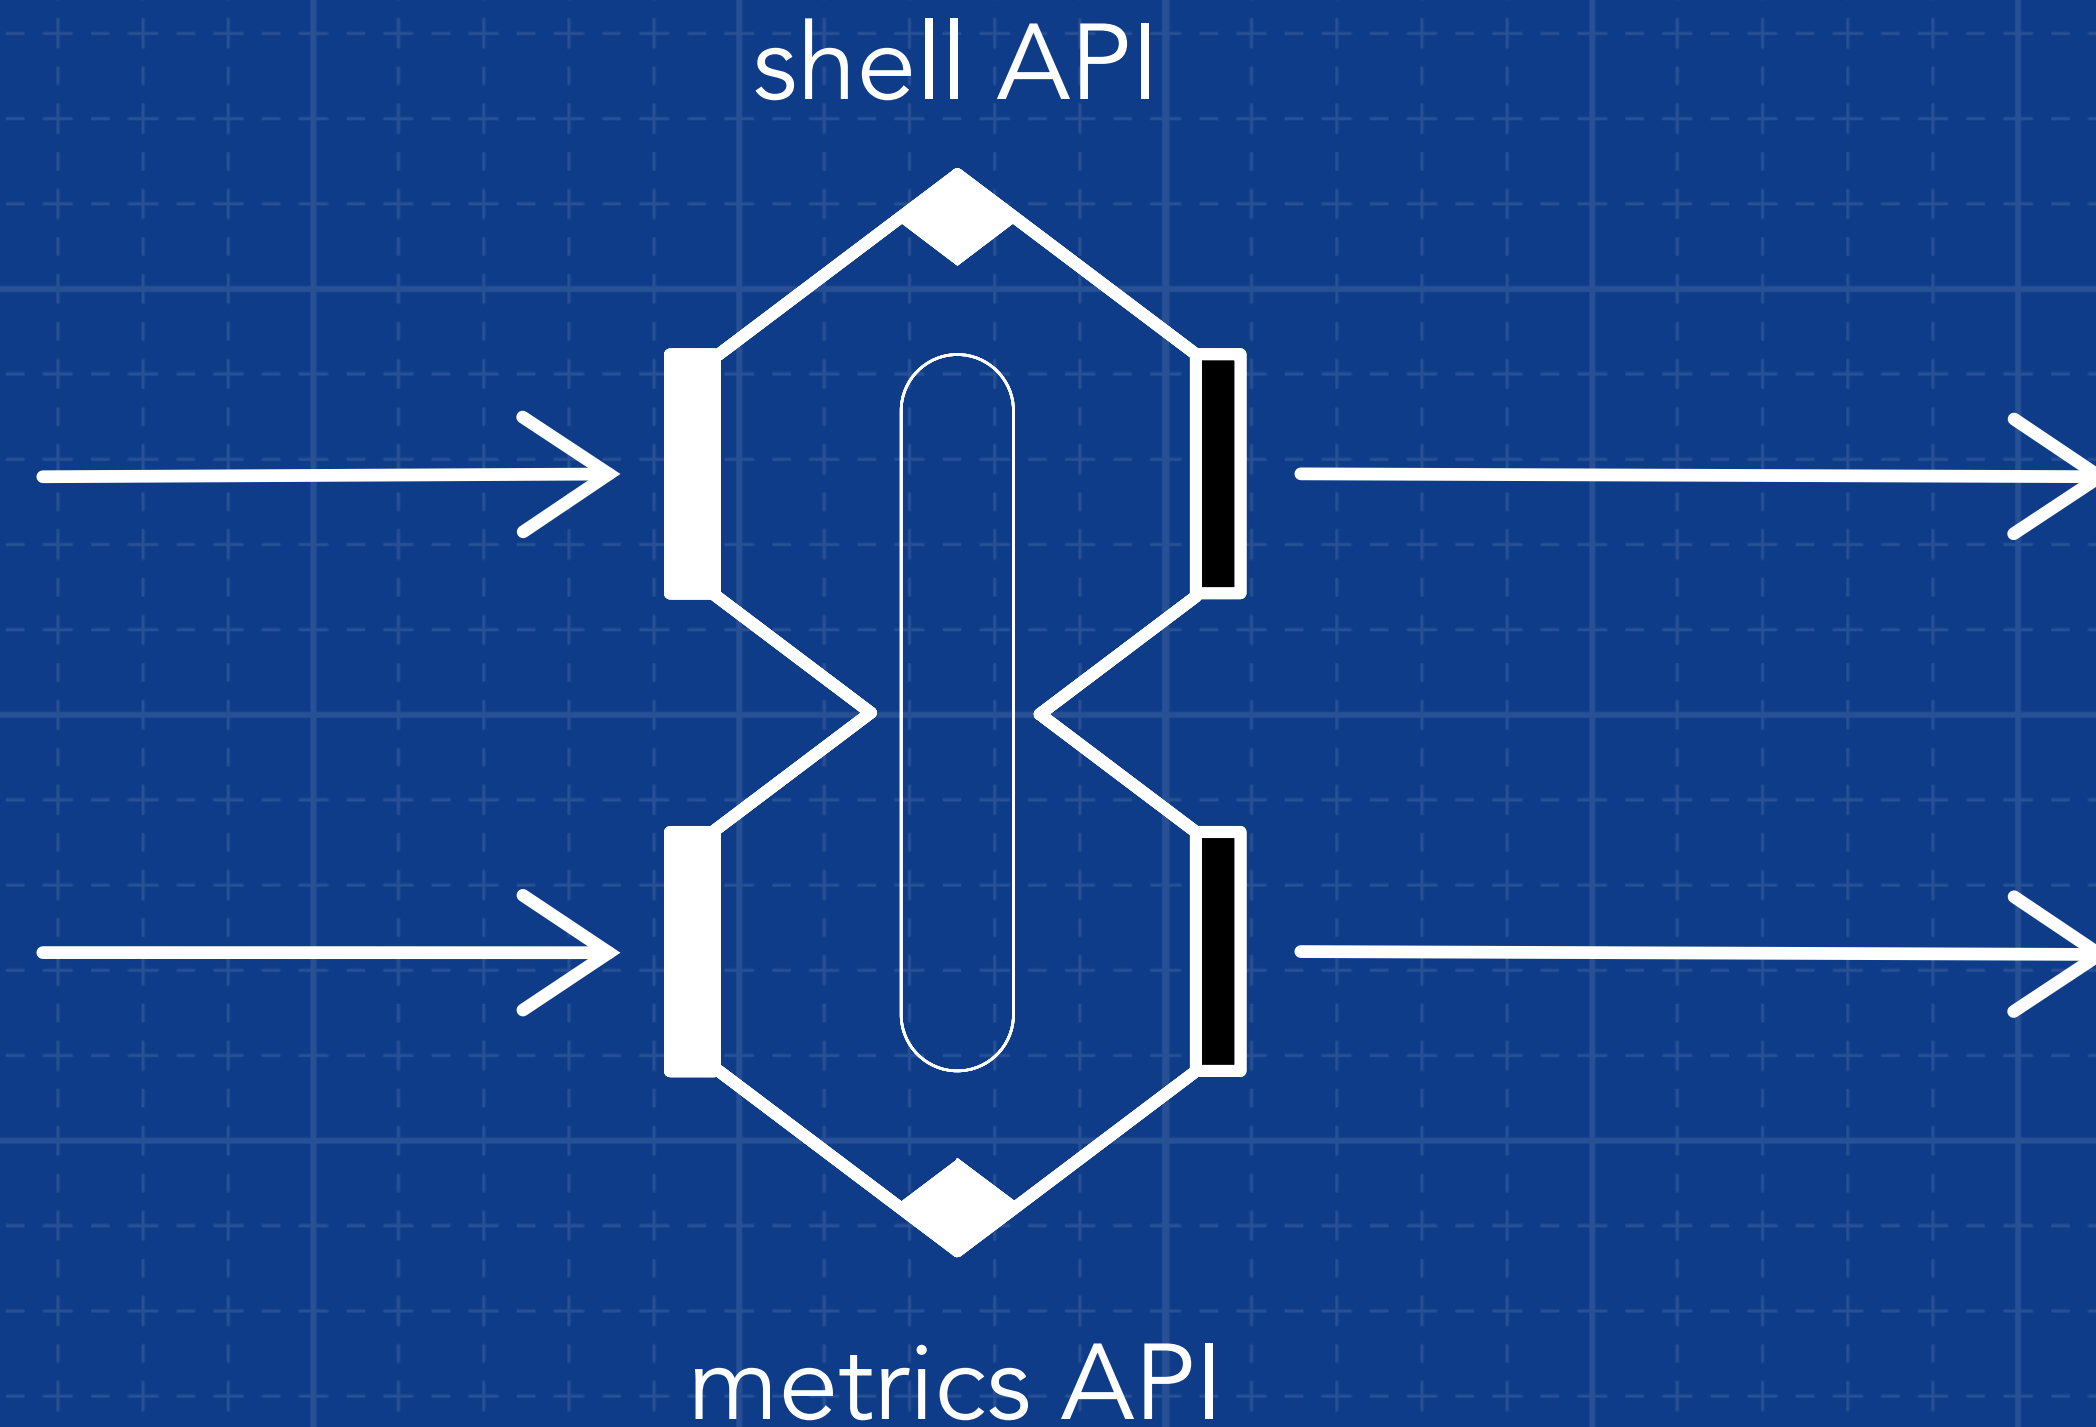
INPUTS

multiple input/
single type

OUTPUTS

multiple output/
single type

can be a stand-alone runtime



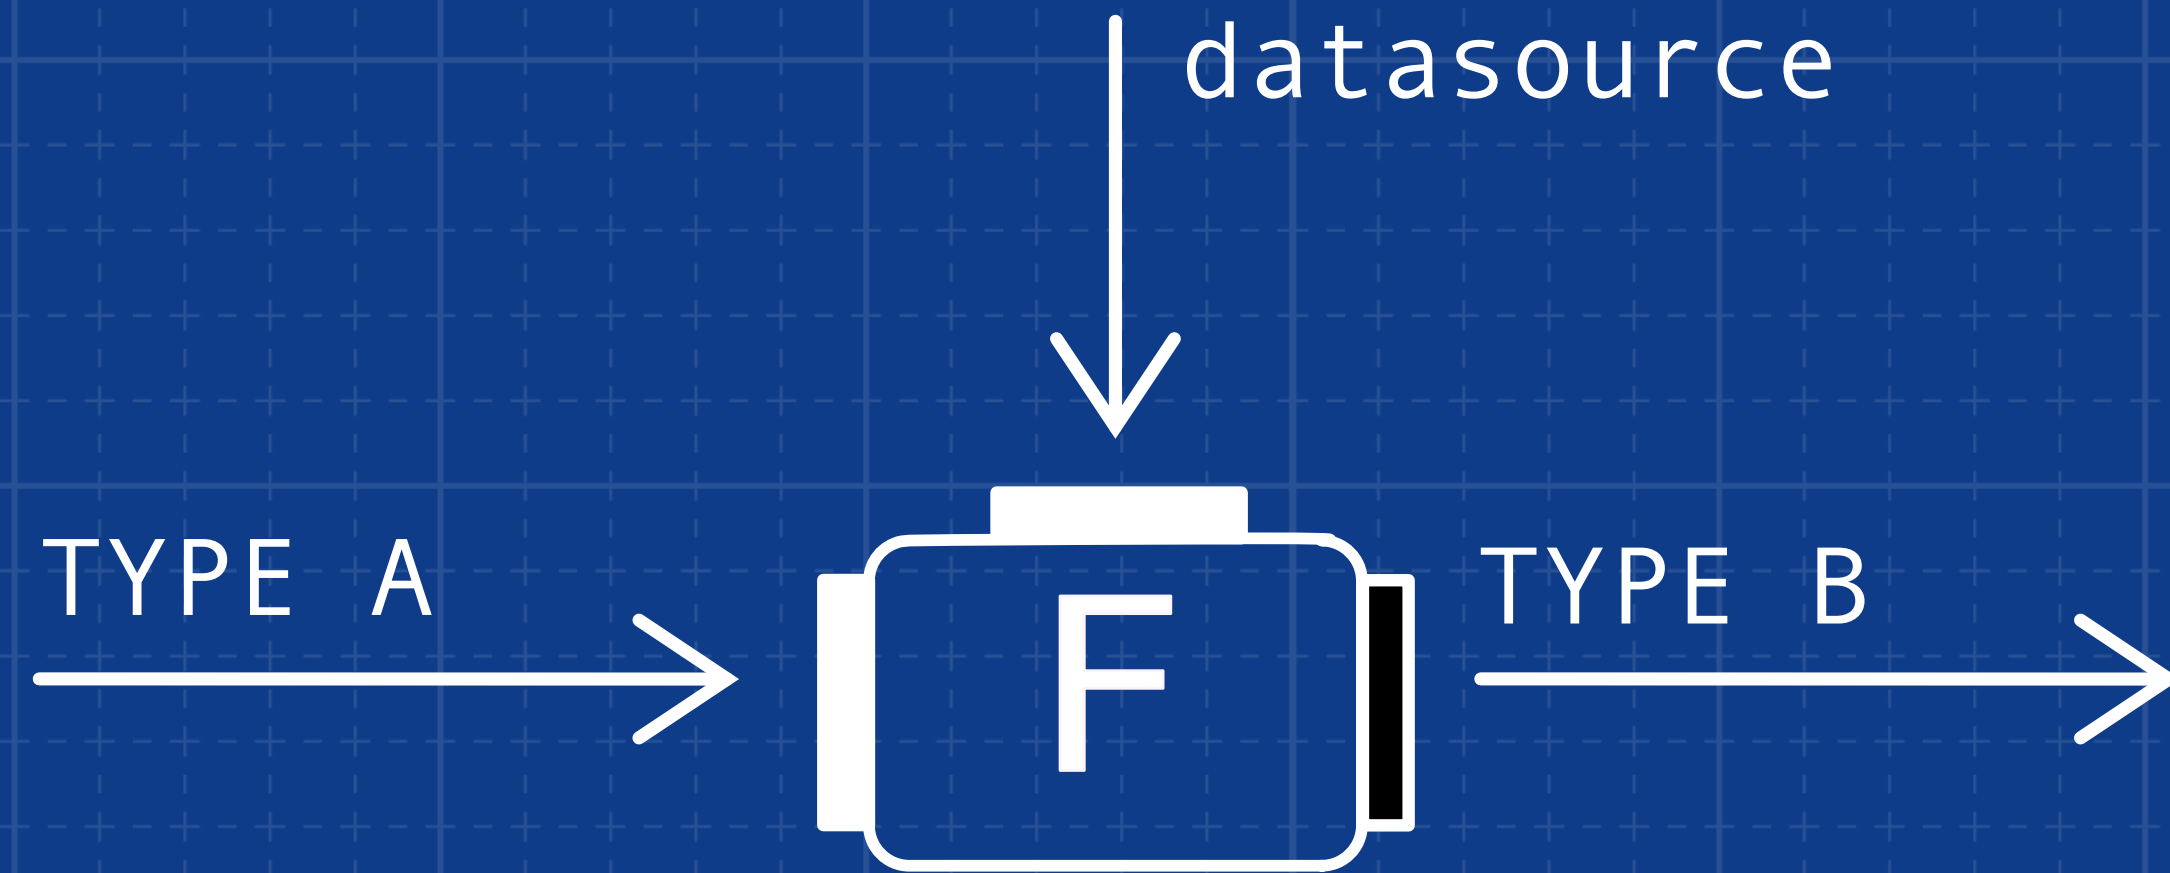
PROGRAMMABLE FILTERMAP - 1

INPUTS

a single typed data stream

OUTPUTS

a single typed data stream

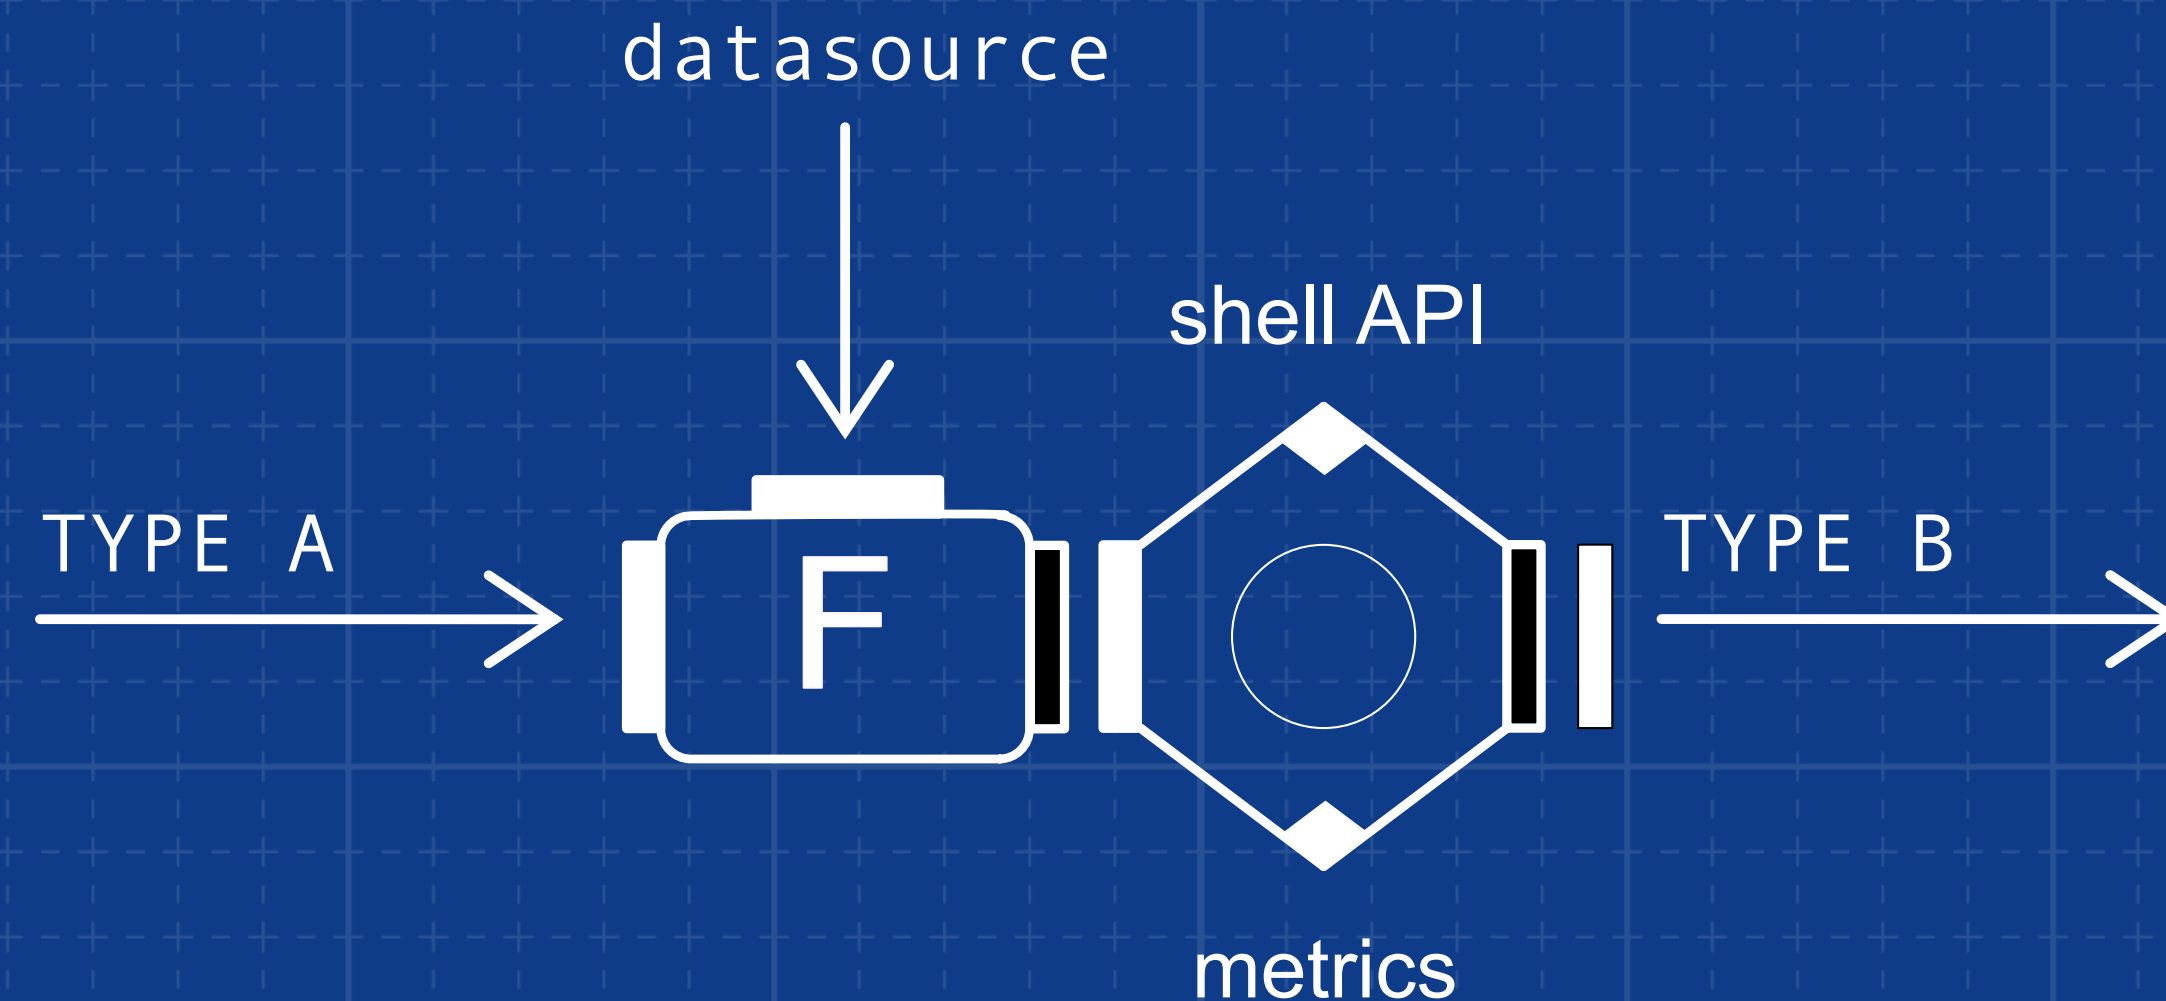


lives in a filter stack in front of a RIB

PROGRAMMABLE FILTERMAP - 2

PROCESSING

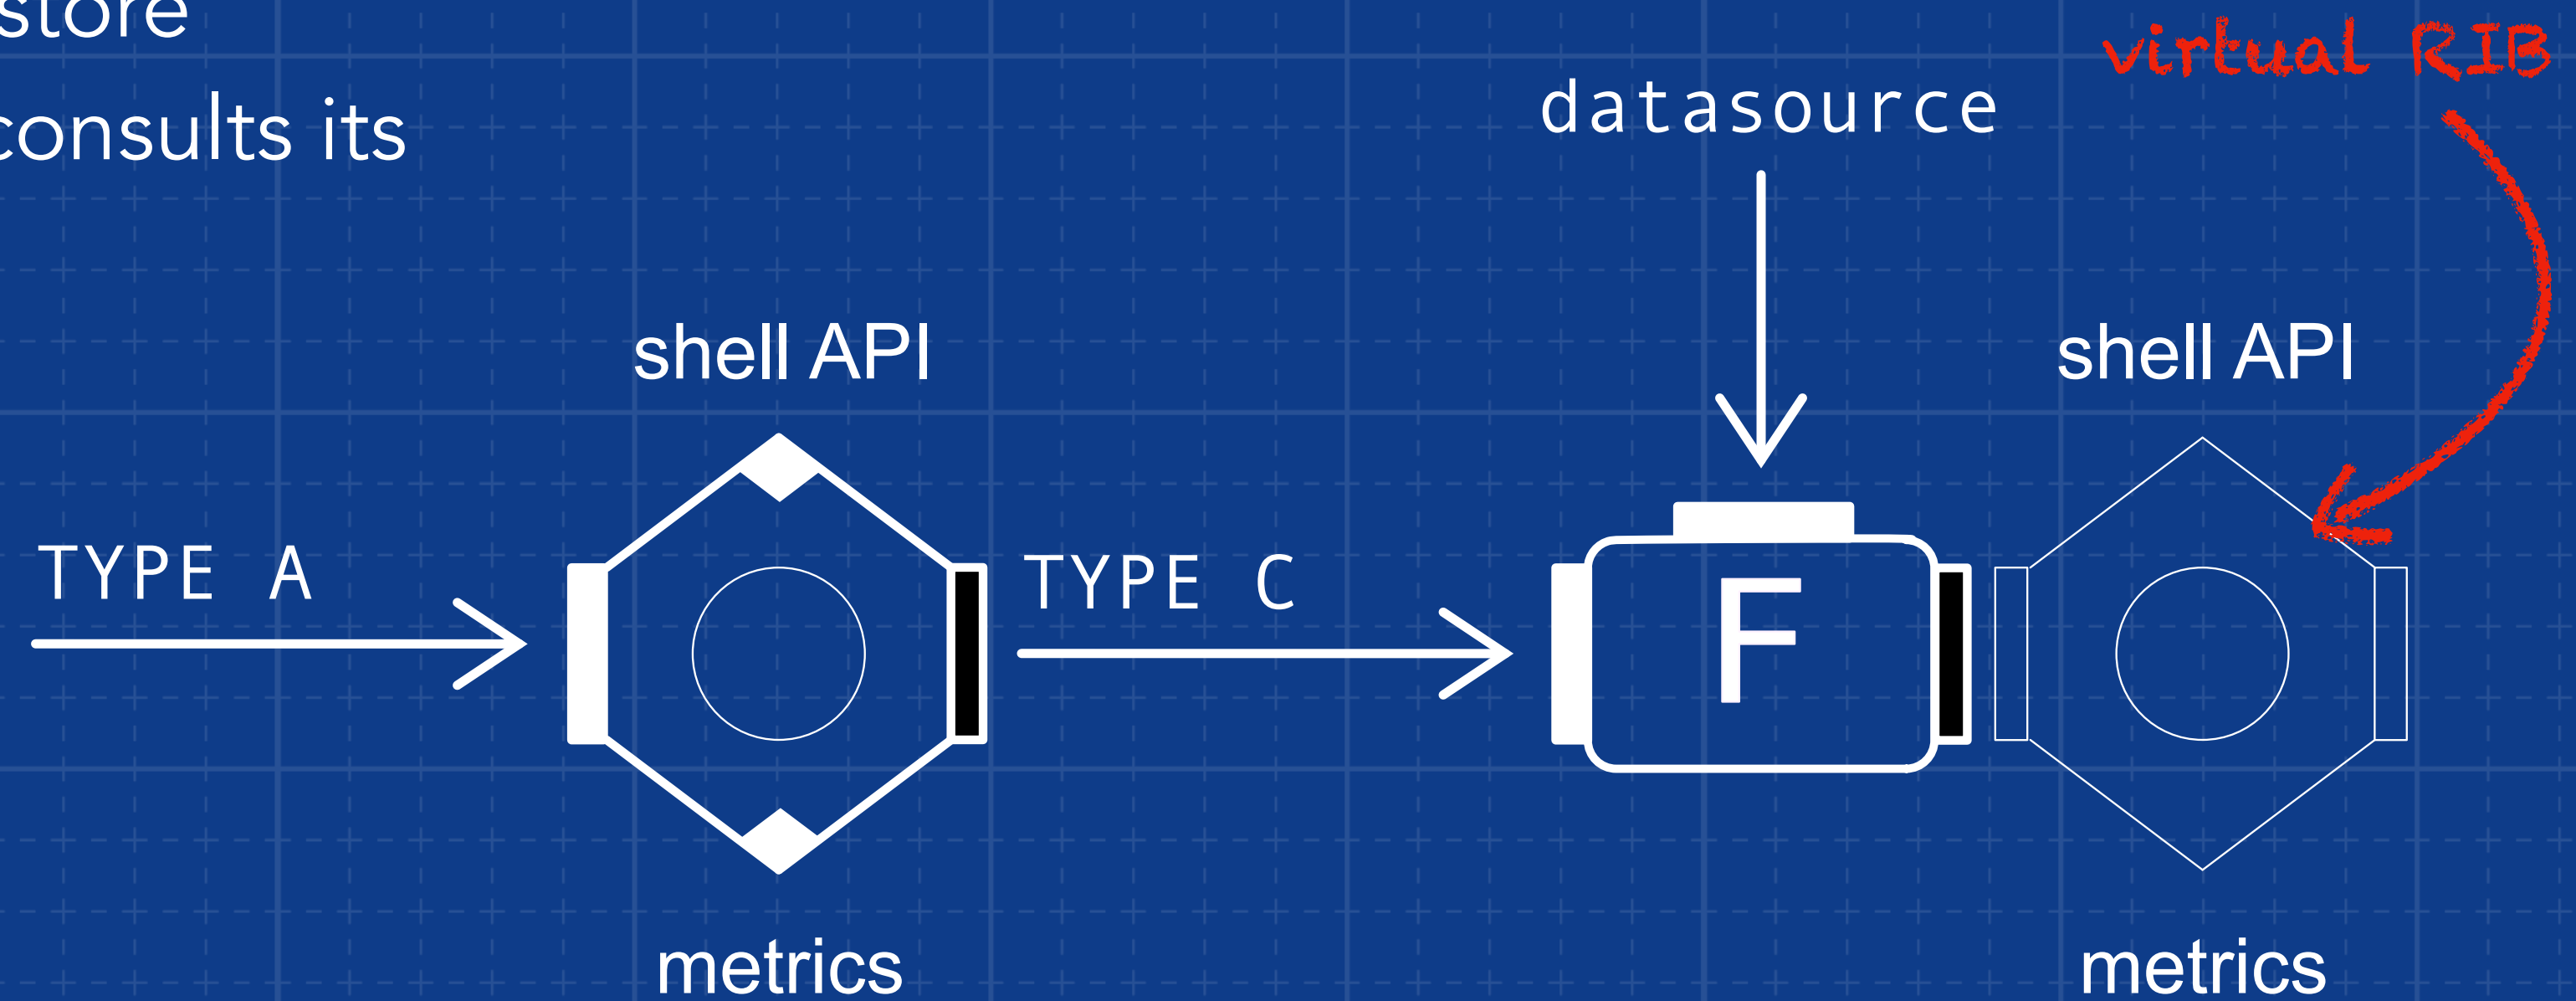
- ROTO filter/config language
- accept/reject
- transform input stream
- realtime per packet processing



lives in a filter stack in front of a RIB

VIRTUAL RIB

same functionality as a normal RIB,
but doesn't actually store
anything, instead it consults its
western RIB



lives in a filter stack in front of a RIB

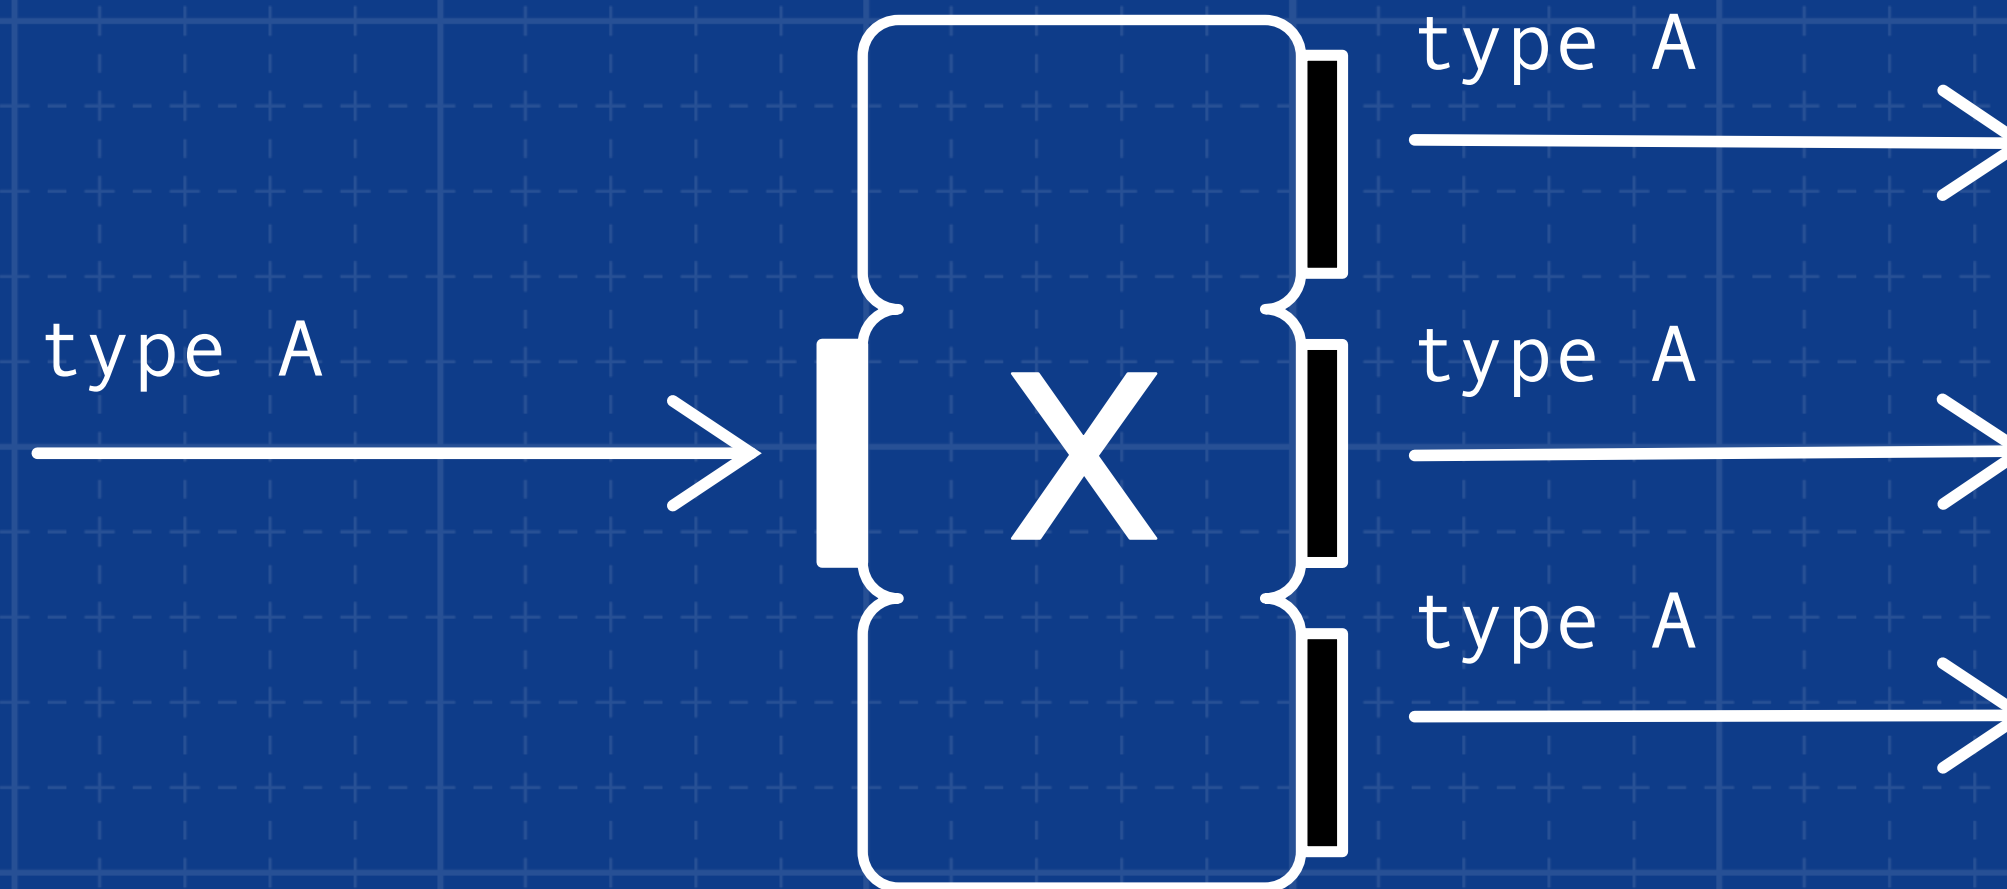
BROADCAST STREAM

INPUTS

single input/
single type

OUTPUTS

multiple outputs/
single type



can be a stand-alone runtime

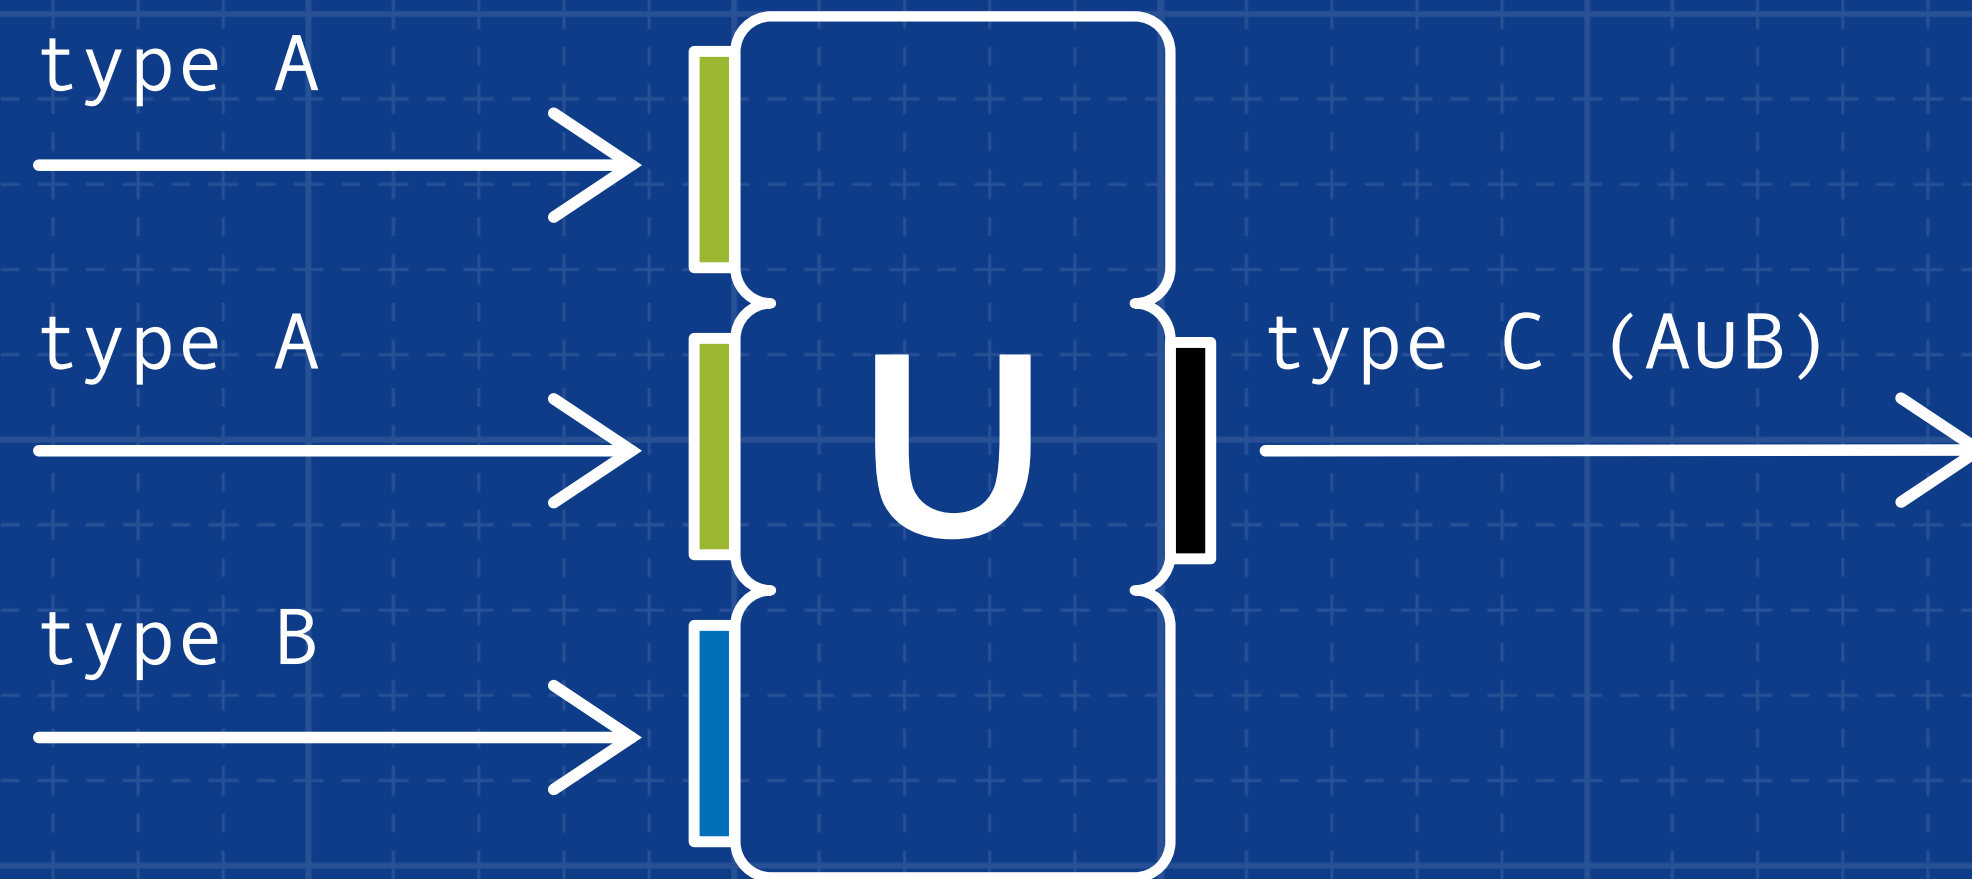
MERGE STREAMS

INPUTS

multiple inputs/
multiple types

OUTPUTS

single output/
single type



can be a stand-alone runtime

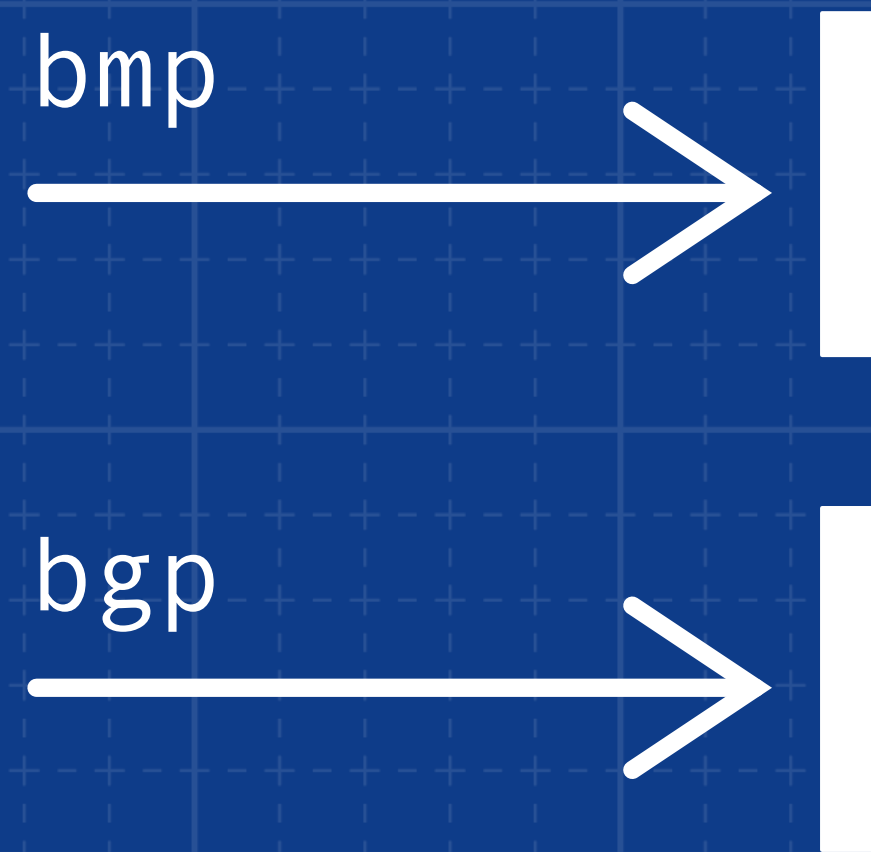
CONNECTORS: INPUTS

INPUTS

single input/
single types

OUTPUTS

single output/
single type



attached to a runtime

CONNECTORS: OUTPUTS

INPUTS

single input/
single type


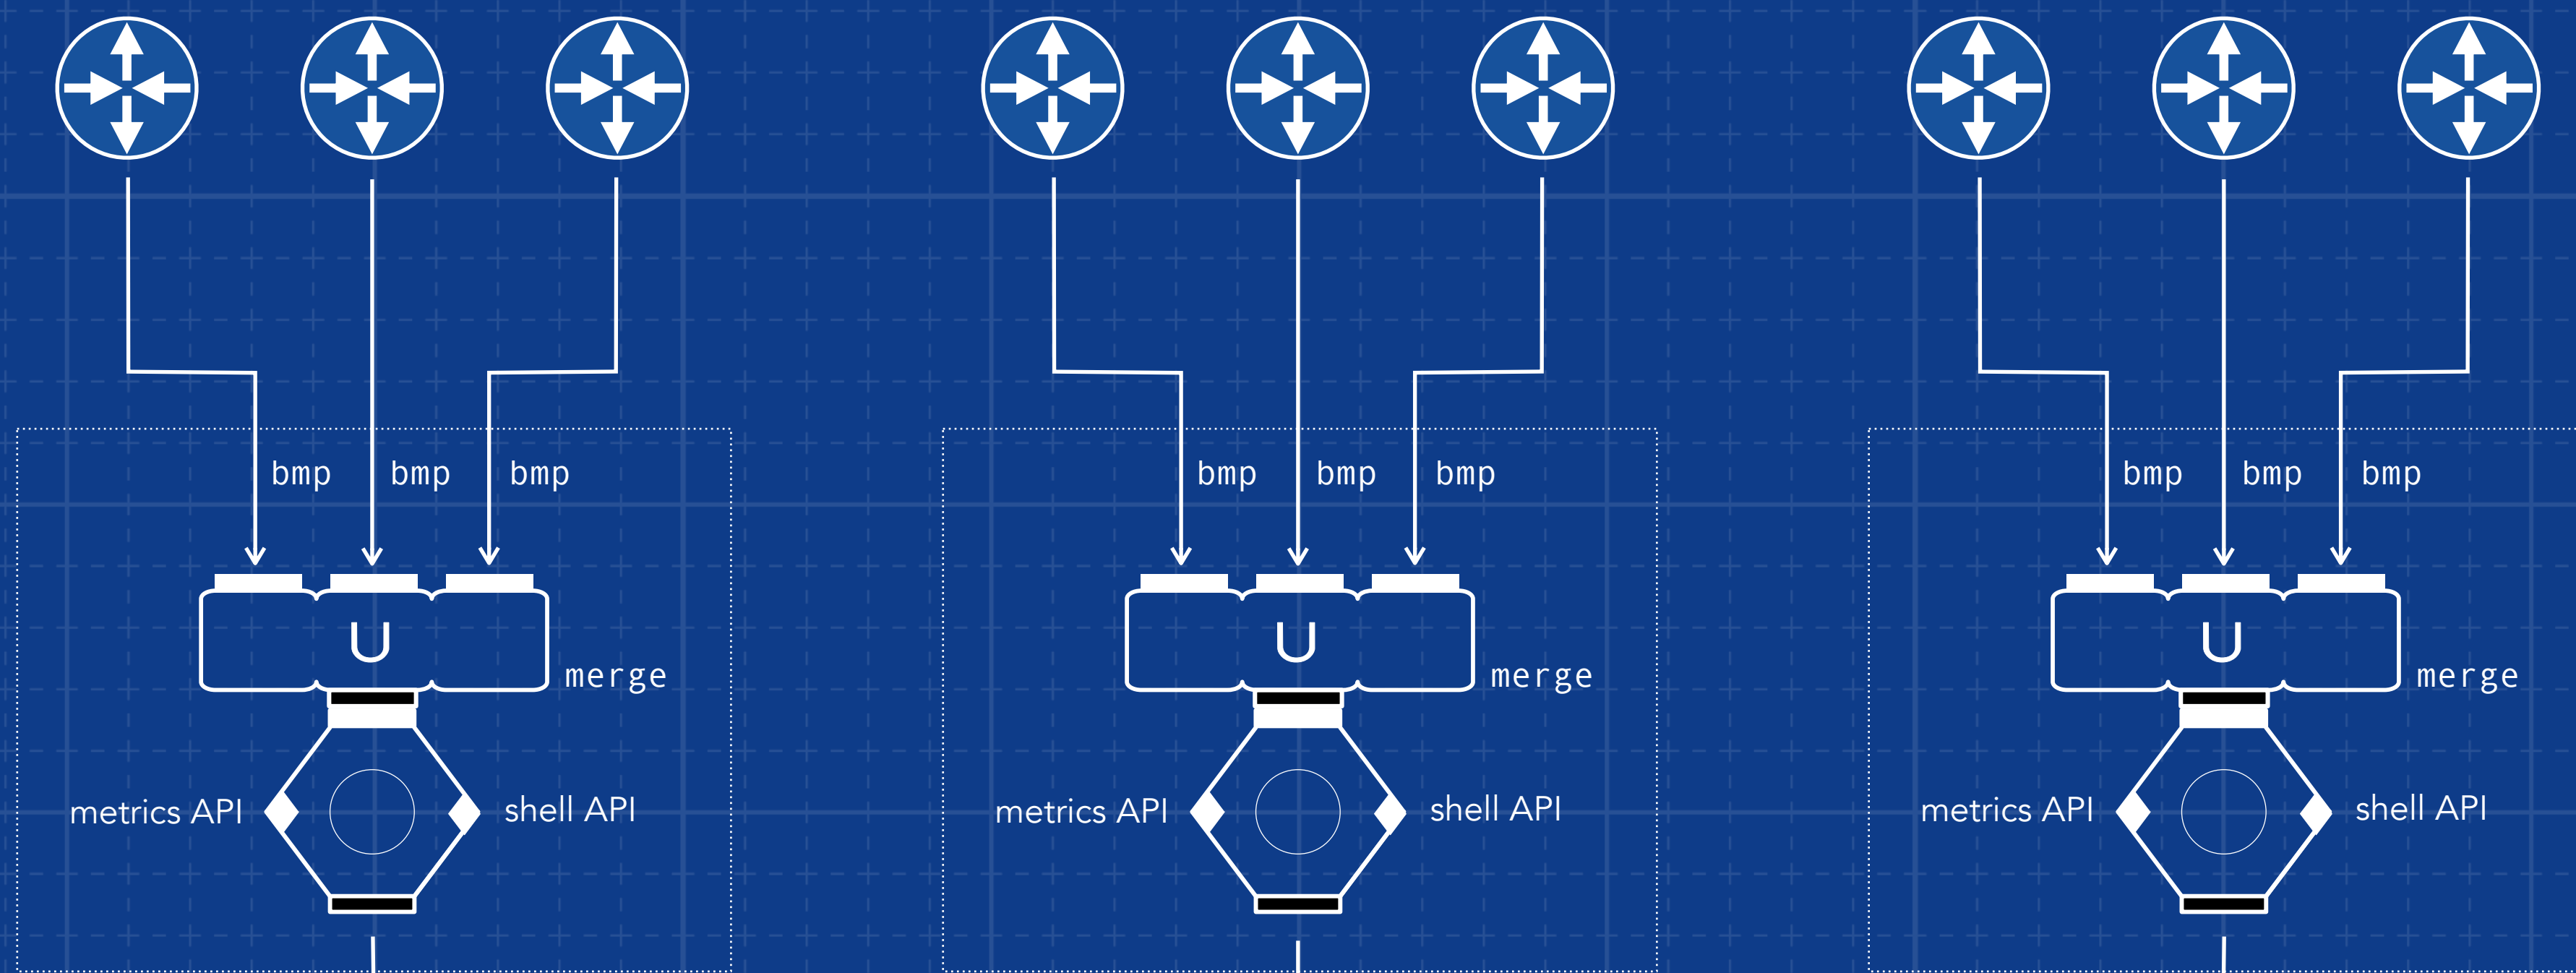
OUTPUTS

single output/
single type






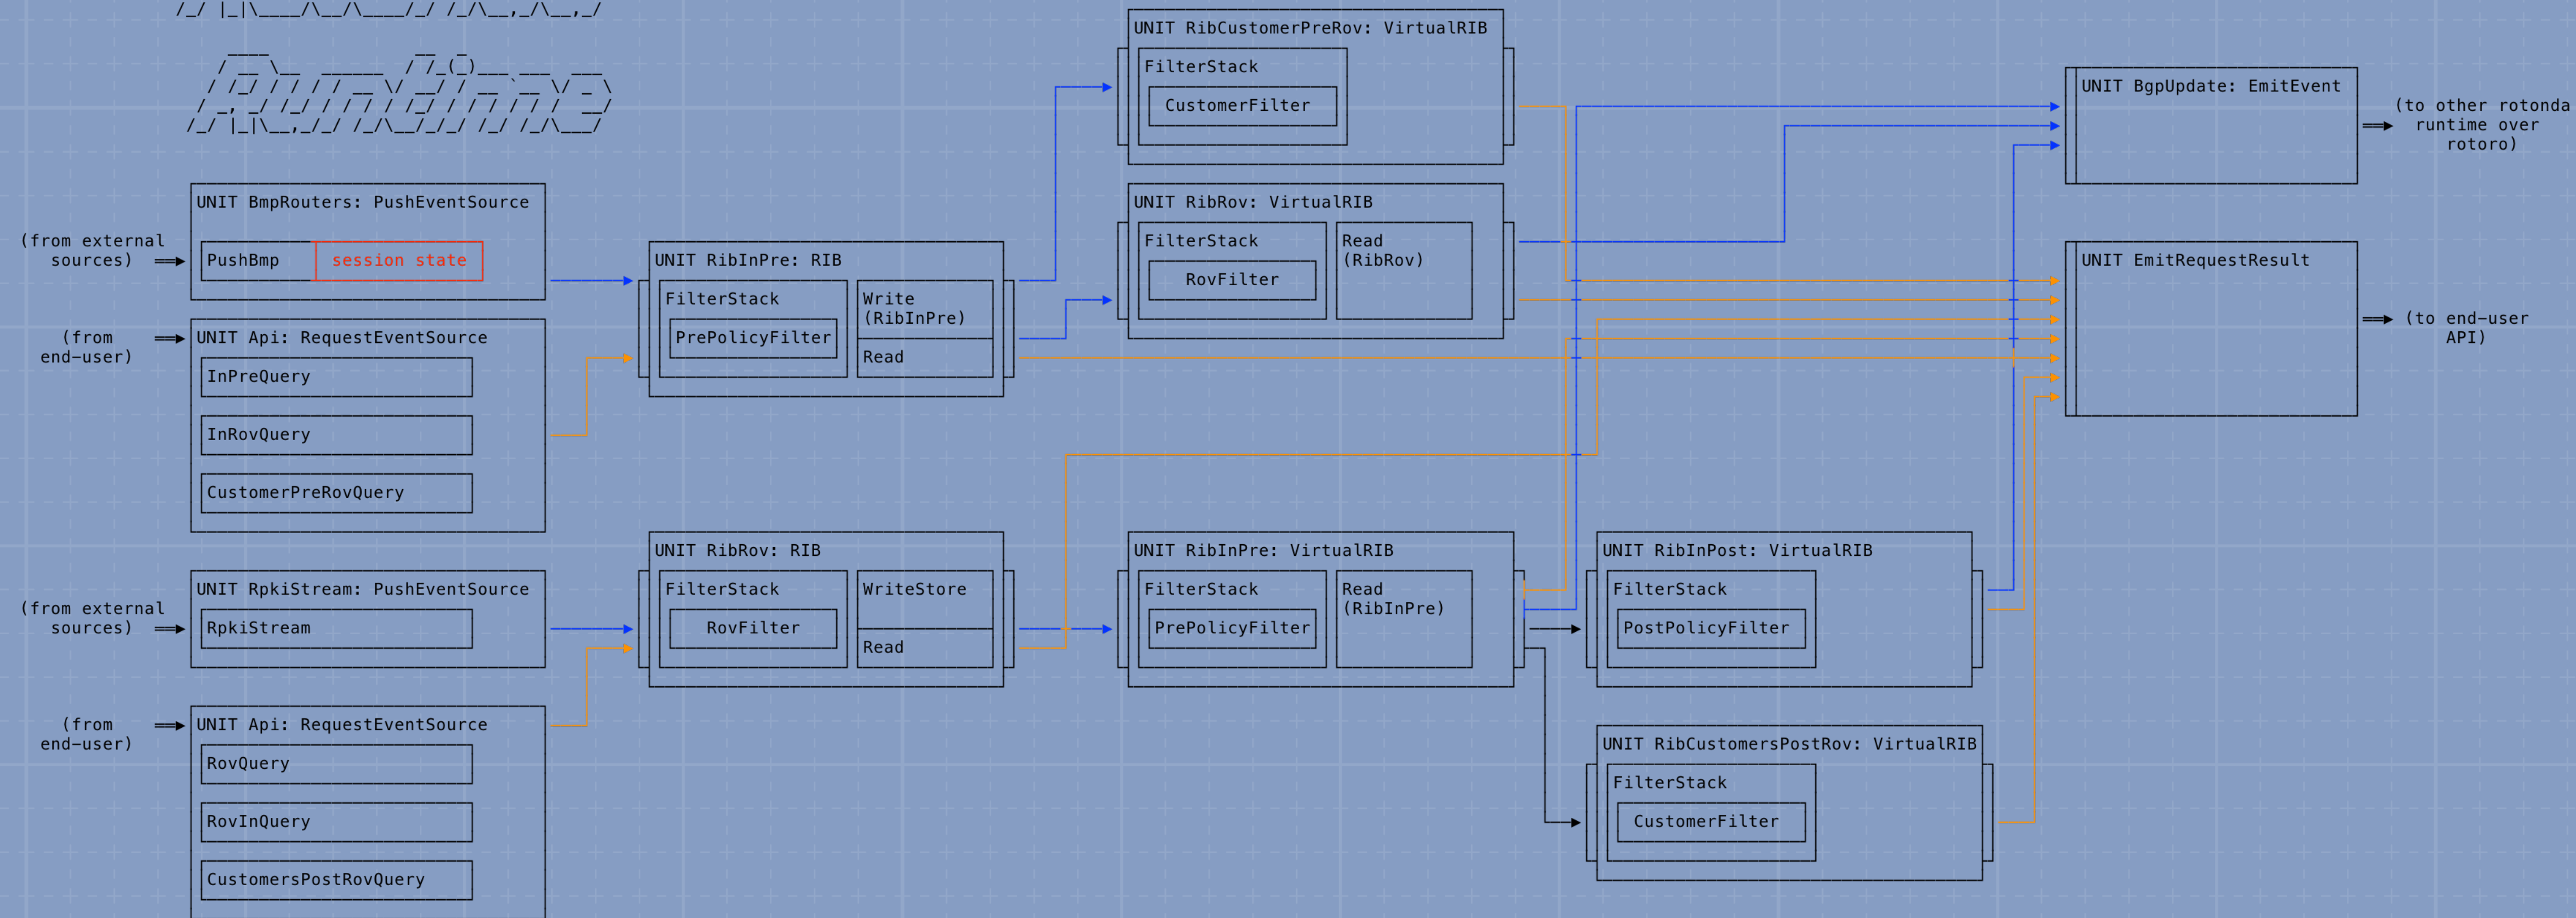
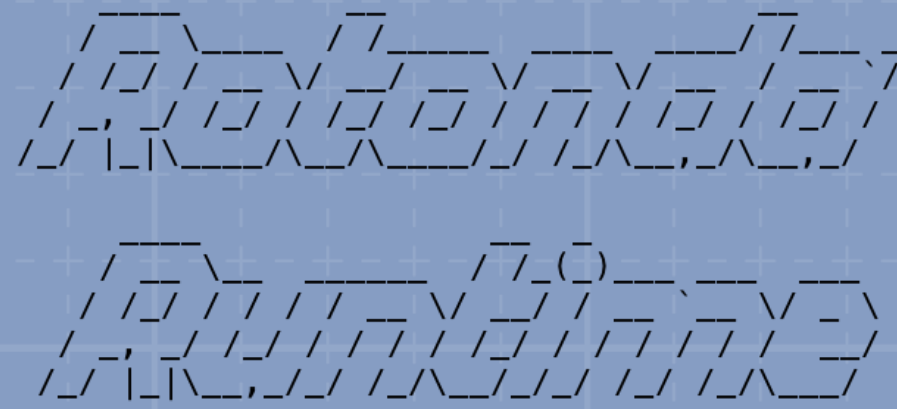
attached to a runtime

WORK IN PROGRESS

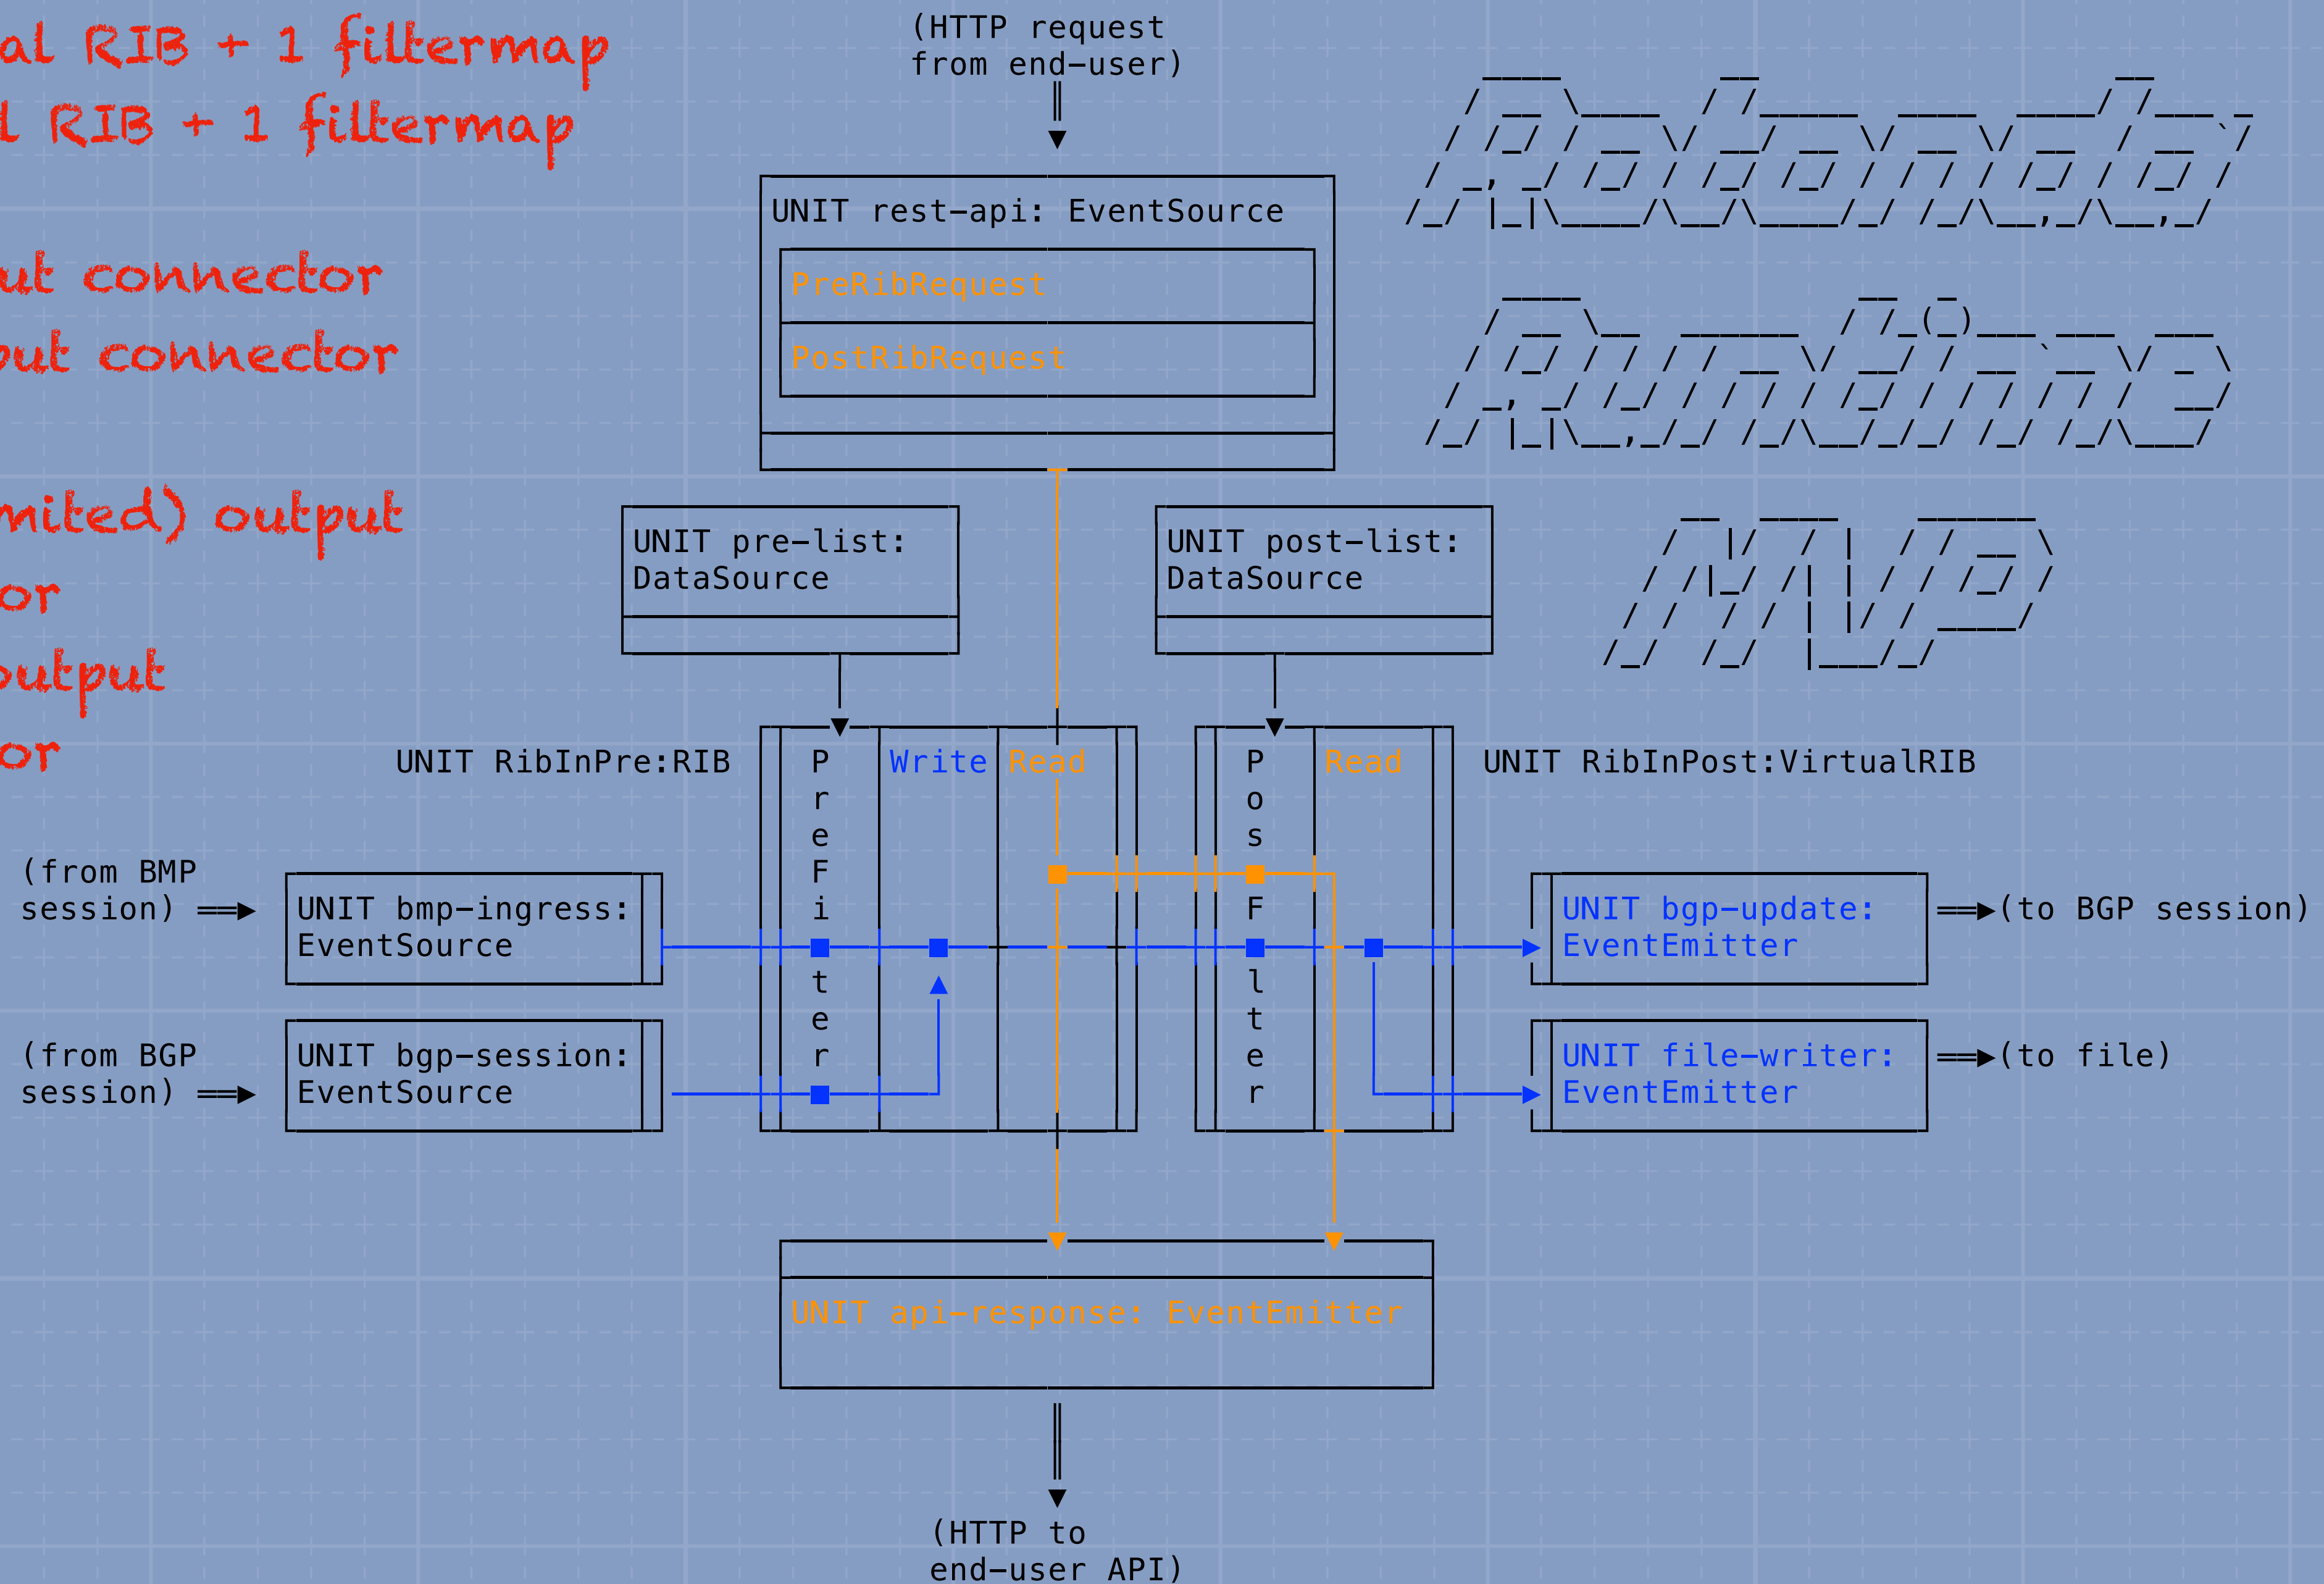


stream output

-  stand-alone runtime
-  single-threaded store
-  multi-threaded store



- 1 physical RIB + 1 filtermap
- 1 virtual RIB + 1 filtermap
- BMP input connector
- BGP input connector
- BGP (limited) output connector
- to file output connector



THANK YOU

BMP AS DATA INGRESS

“BMP provides access to the Adj-RIB-In of a peer on an ongoing basis (...). From a high level, BMP can be thought of as the result of multiplexing together the messages received on the various monitored BGP sessions.”




AS1



AS2



AS3



AS4



AS1



AS2



AS3



AS4




10.1.2.0/24
AS_PATH: AS4




AS1

10.1.2.0/24
AS_PATH: AS3 AS4




AS2

10.1.2.0/24
AS_PATH: AS3 AS4

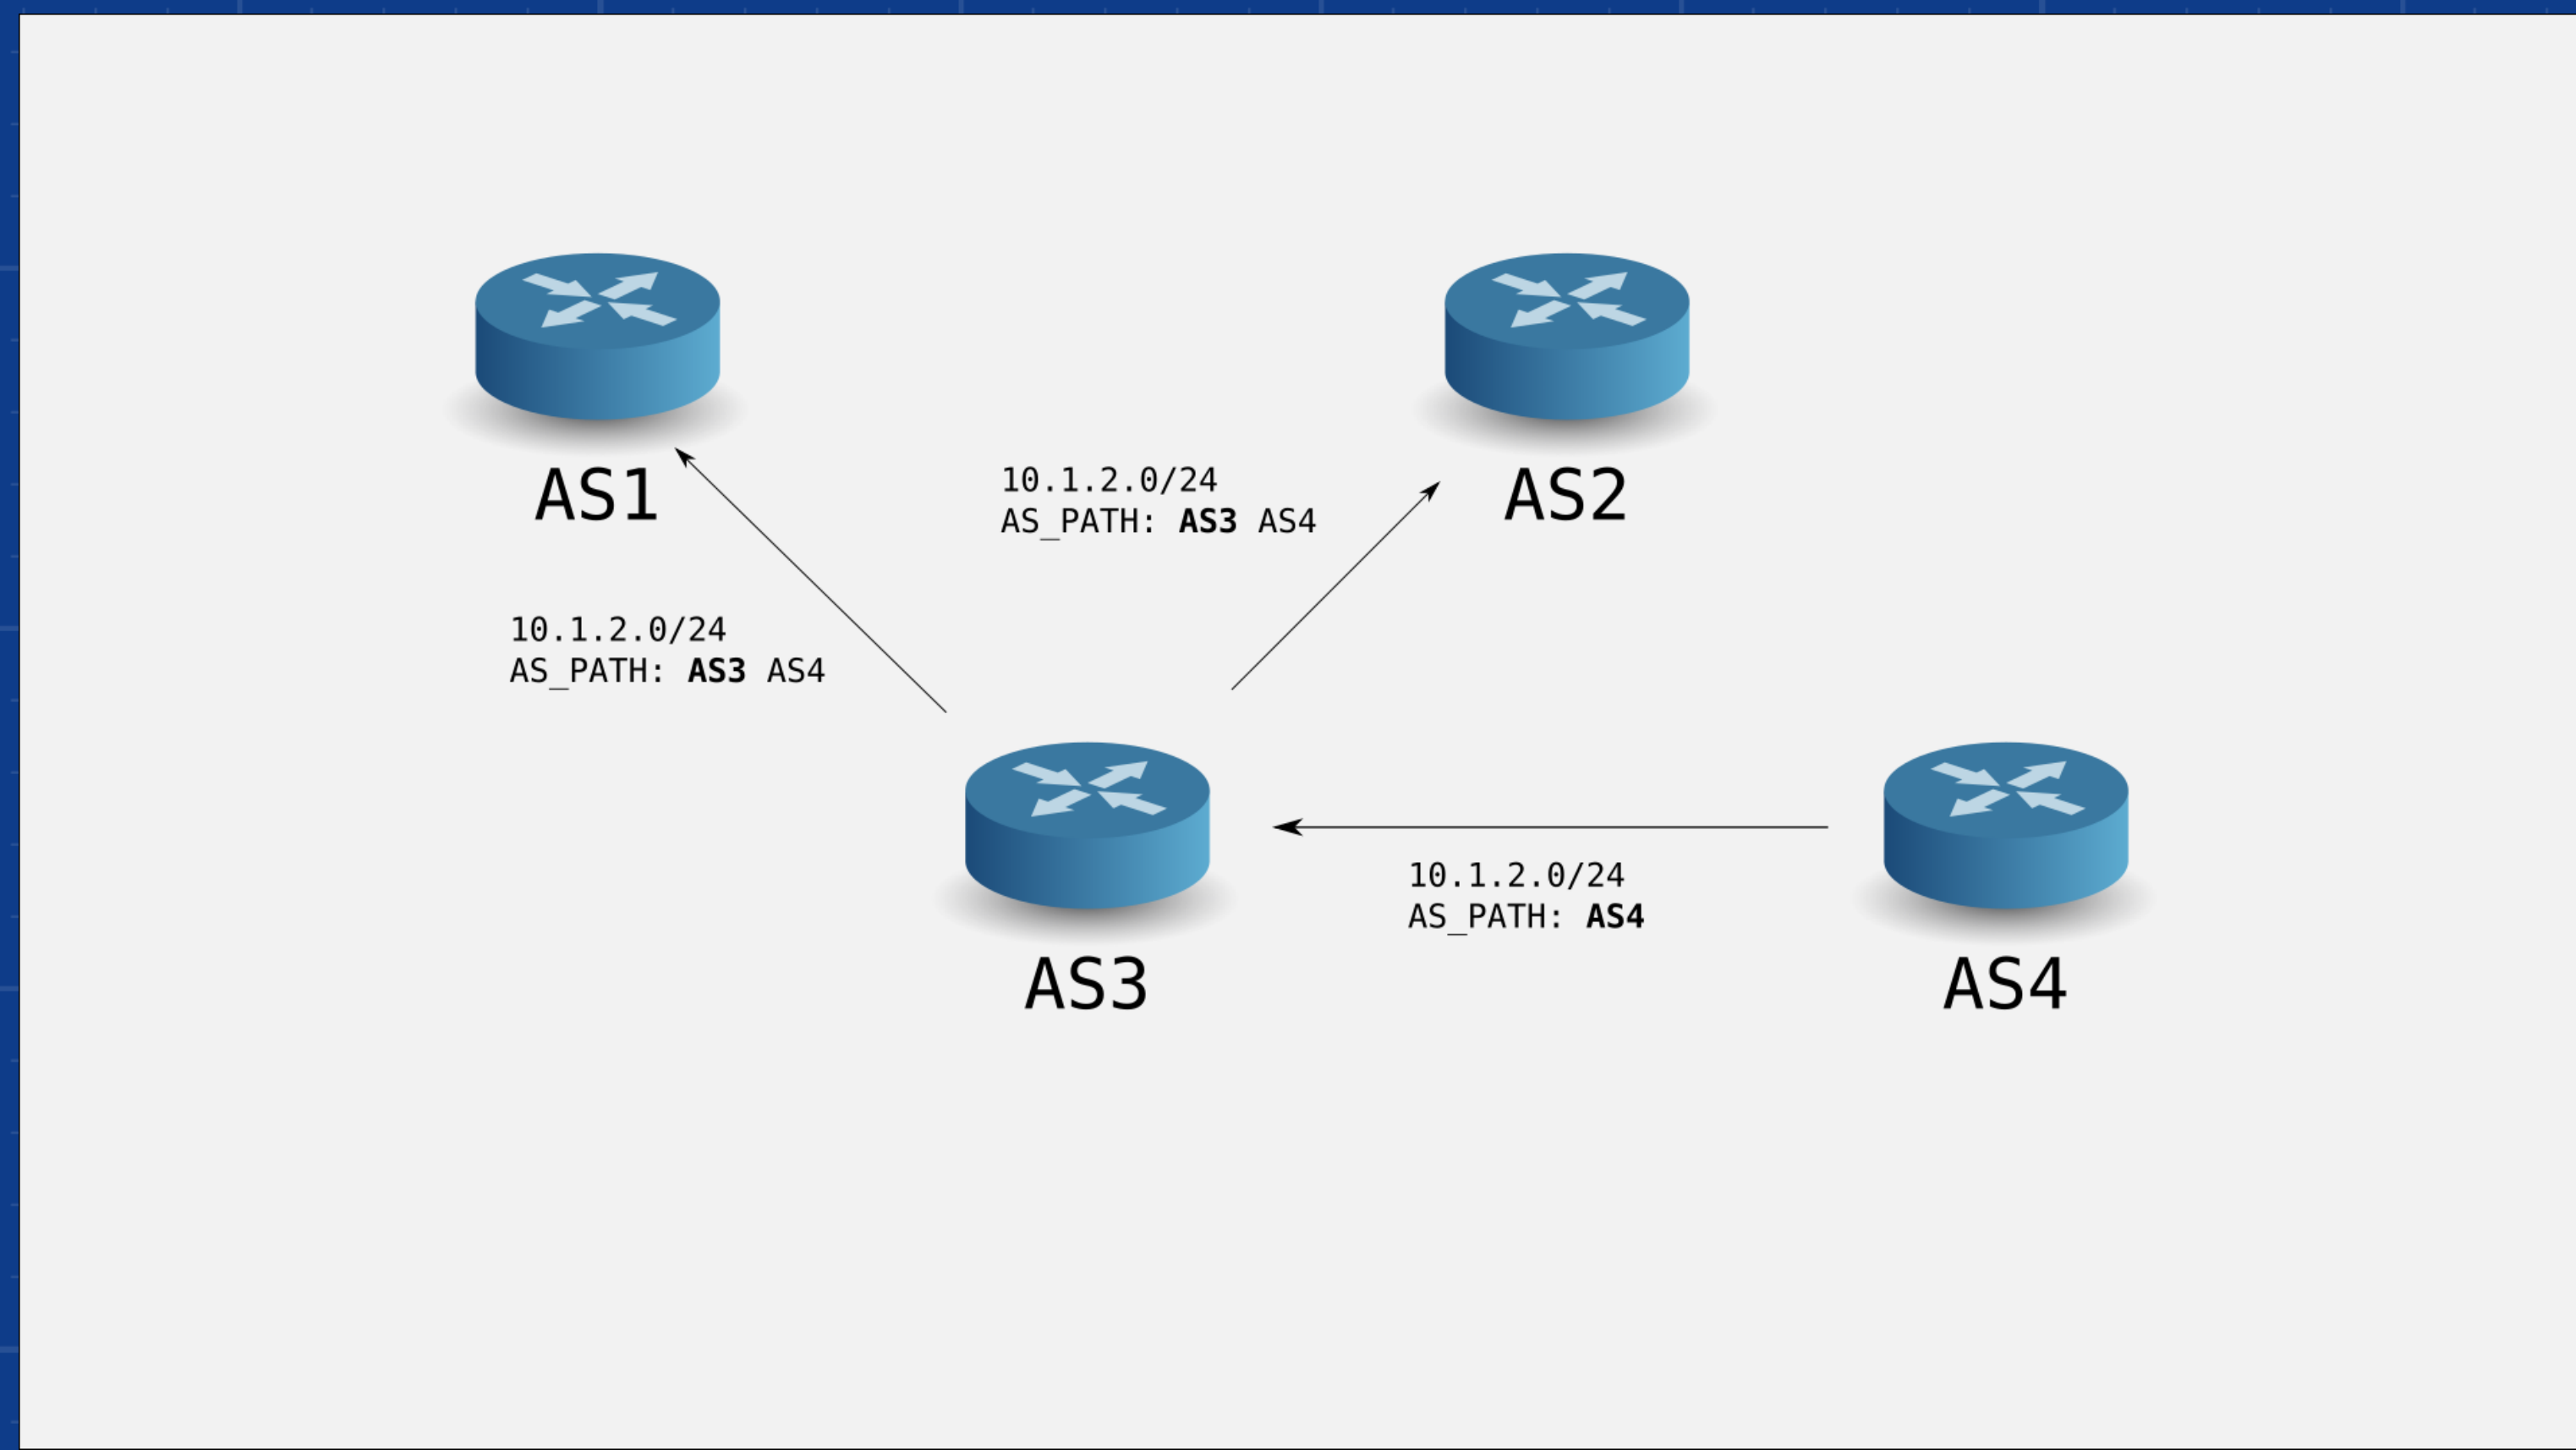


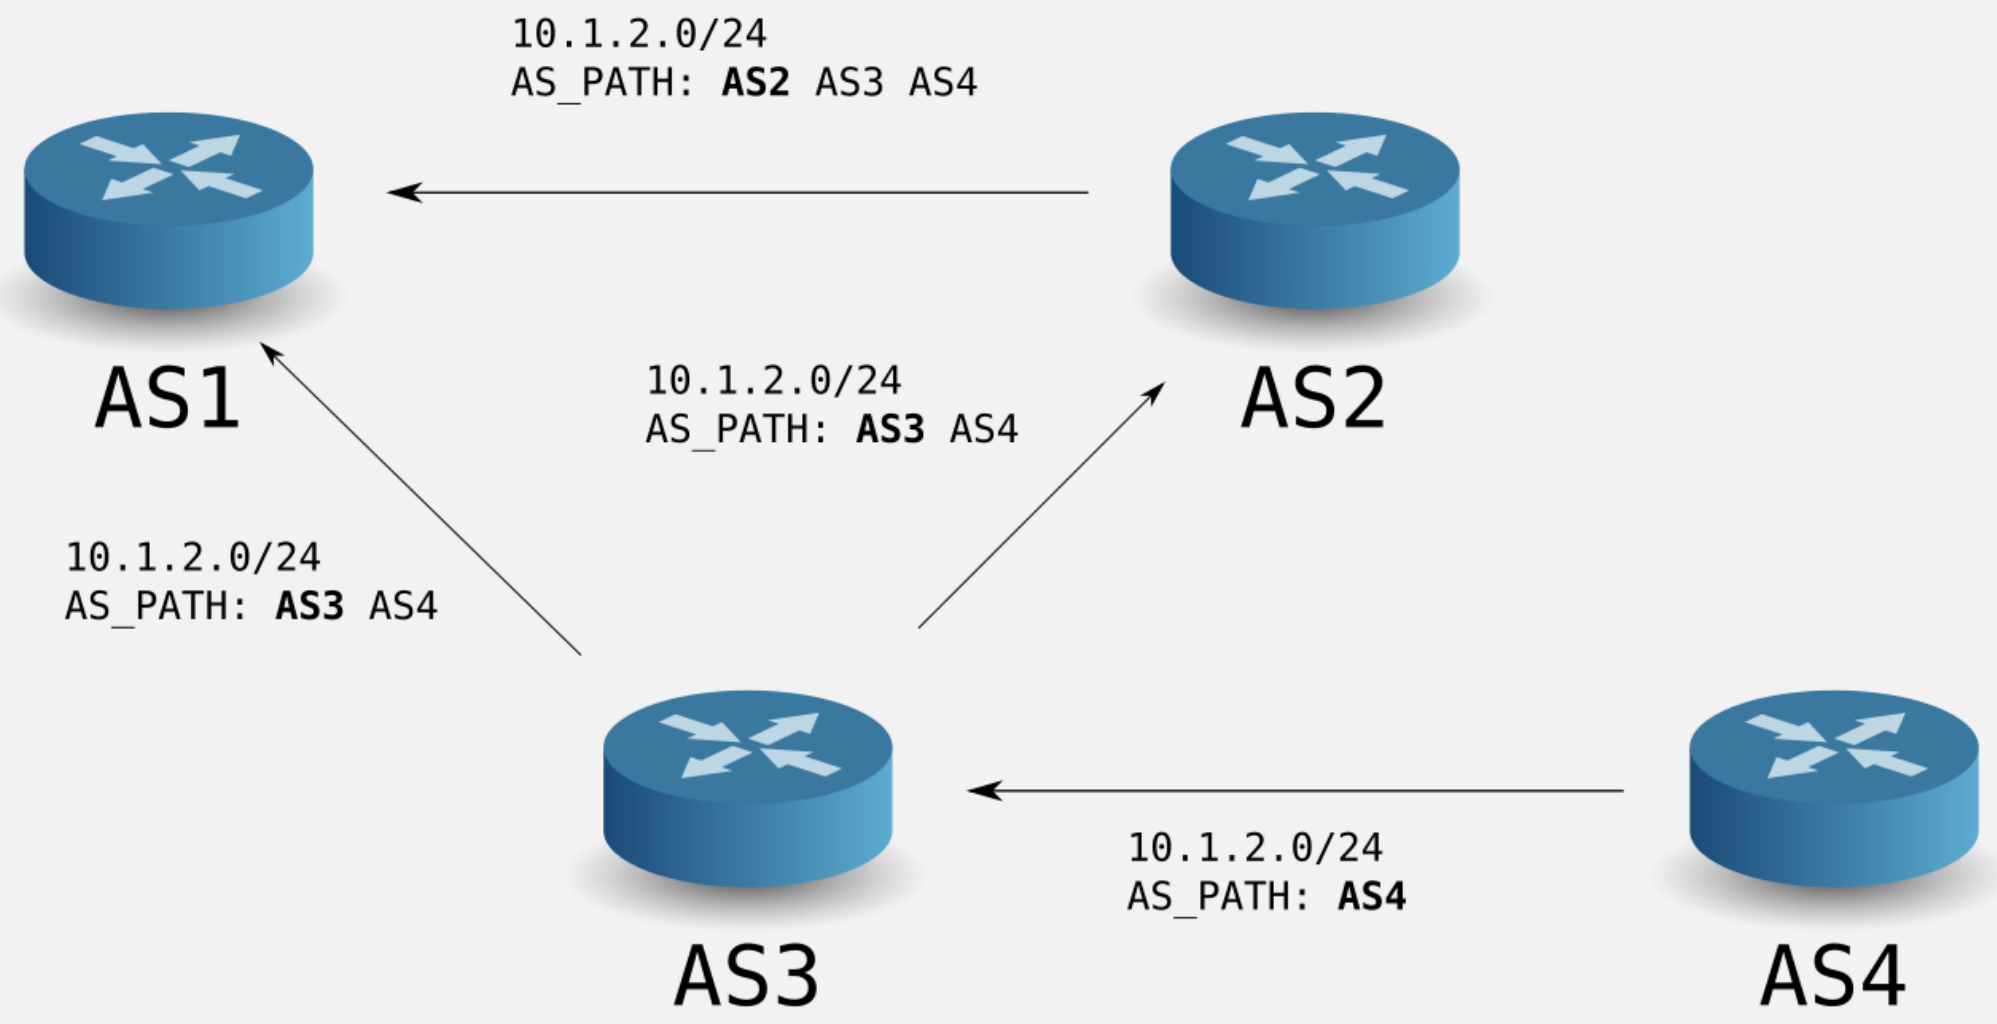
AS3




AS4

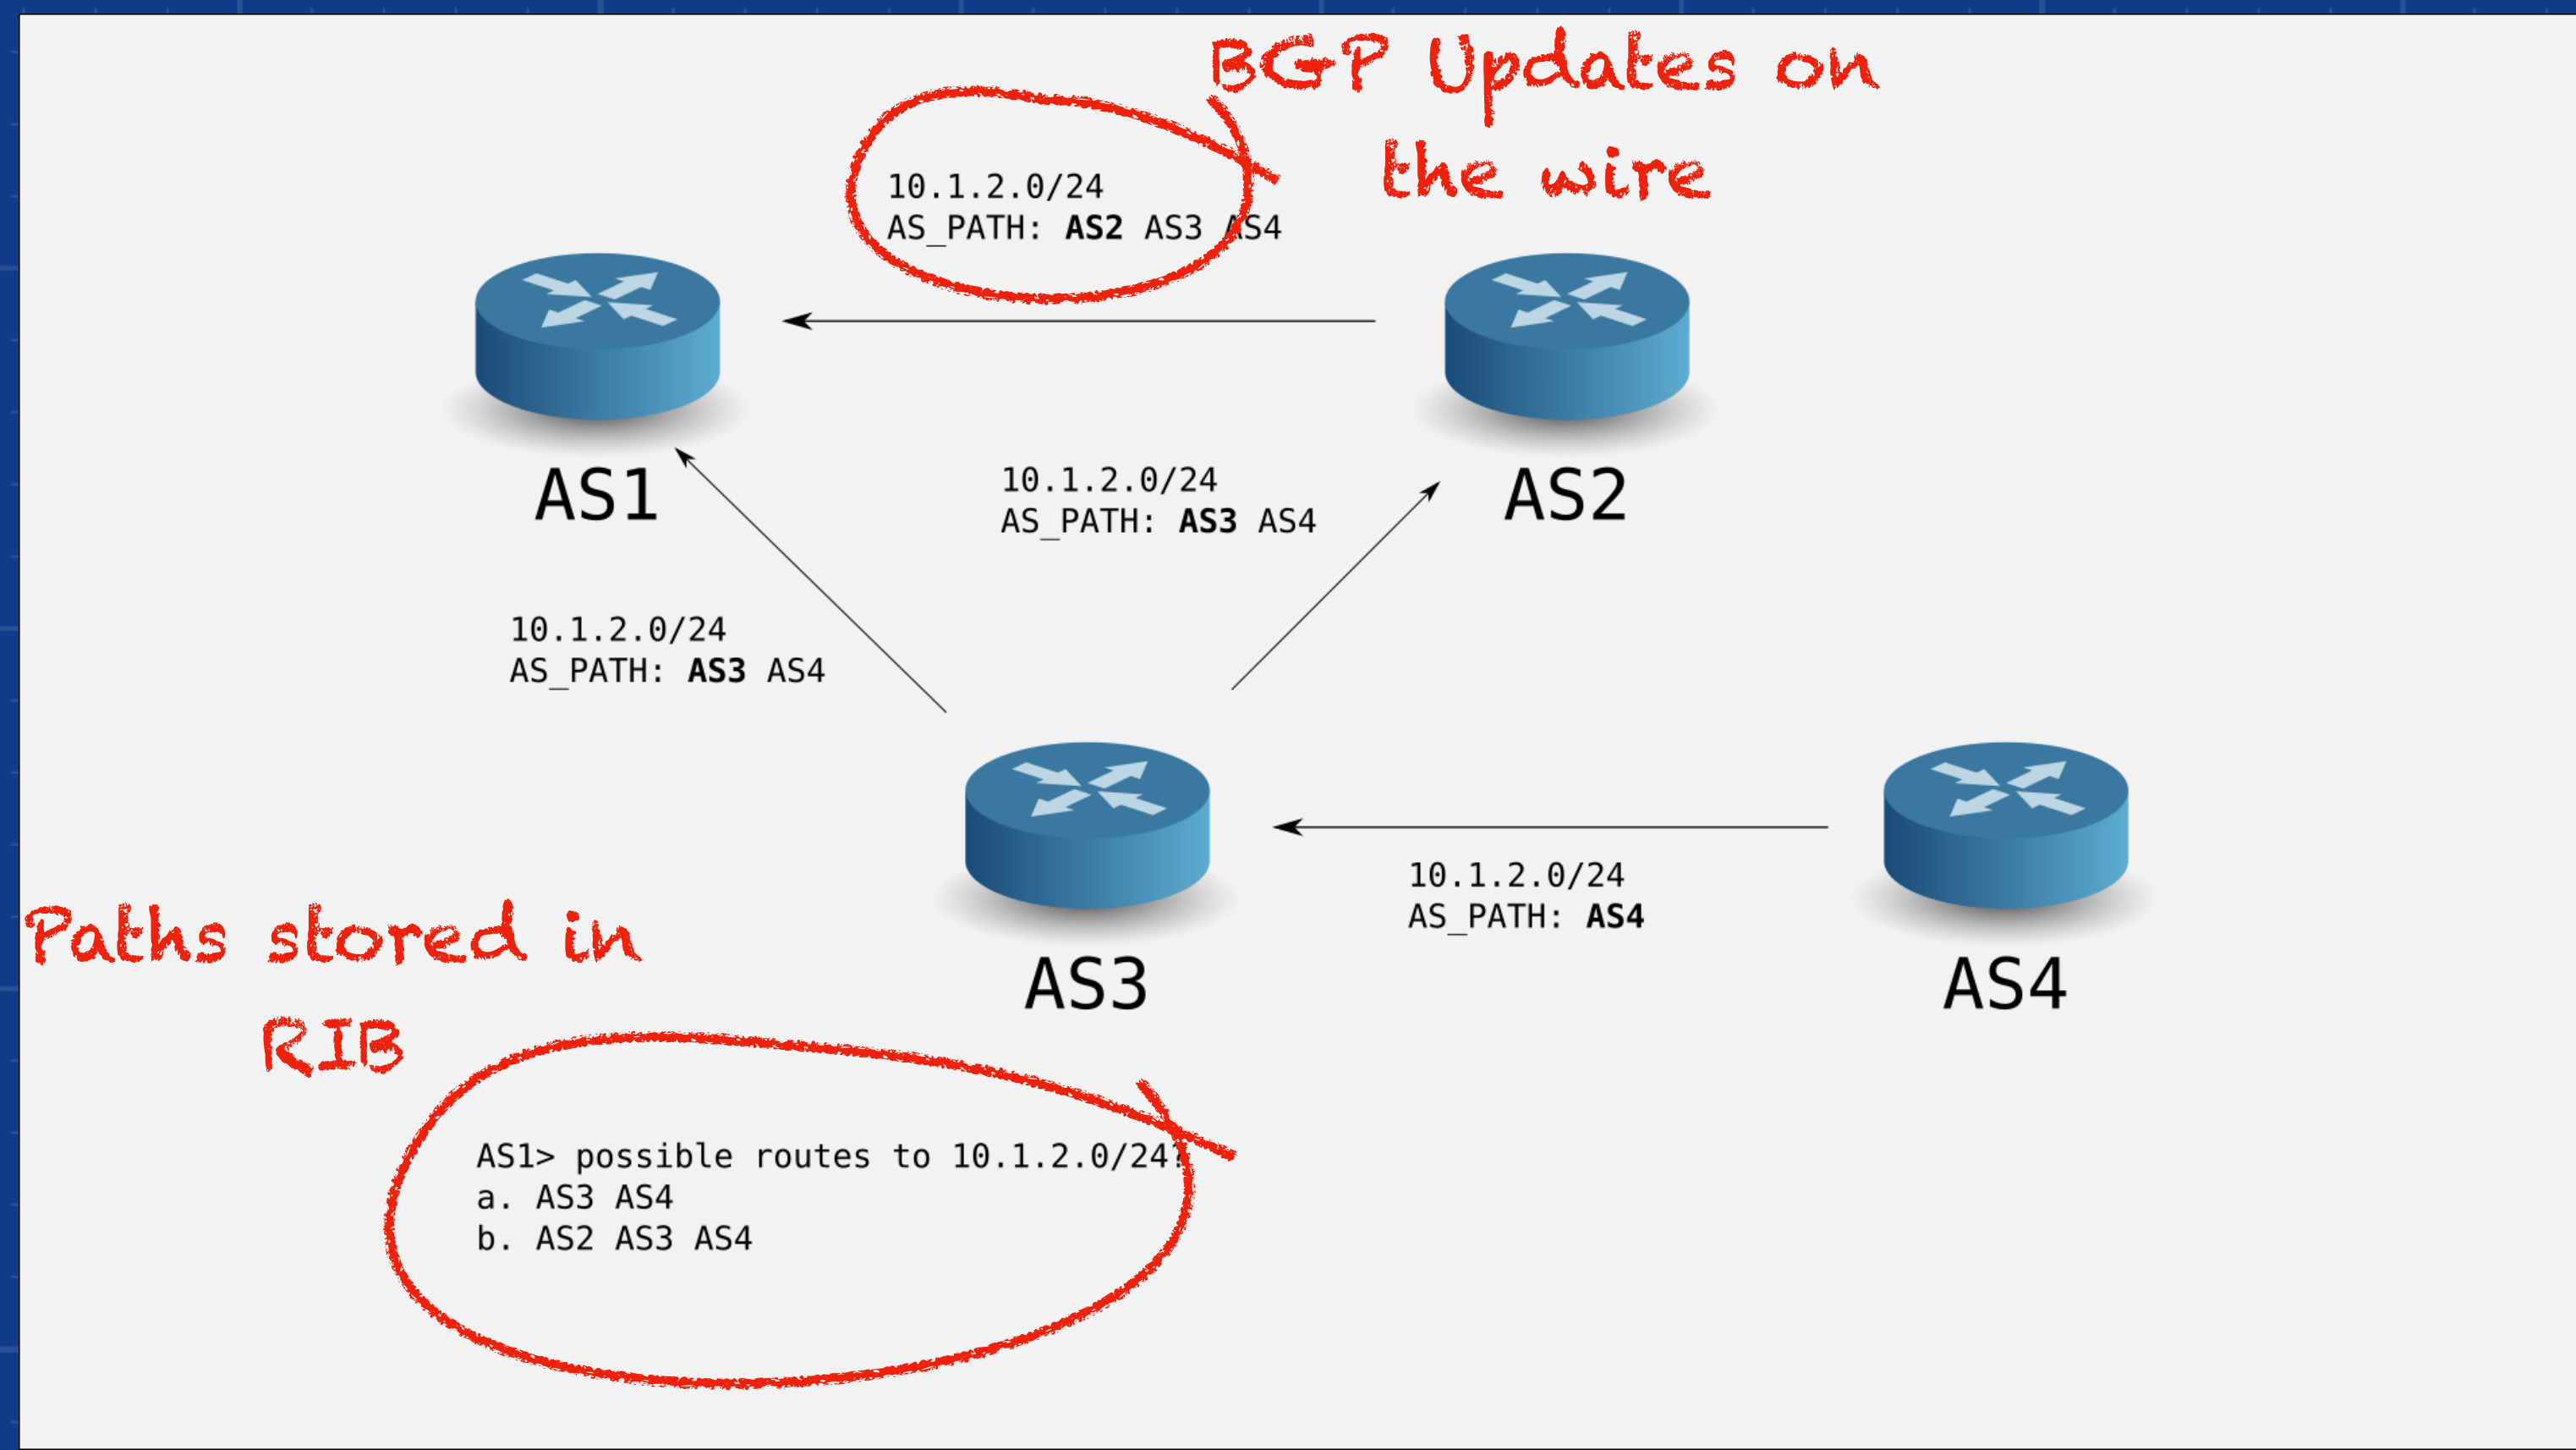
10.1.2.0/24
AS_PATH: AS4







AS1> possible routes to 10.1.2.0/24?
a. AS3 AS4
b. AS2 AS3 AS4



GETTING DATA FROM YOUR ROUTERS

PRE BMP

- Scraping the CLI
- Configure a peering session with the router, receive all routes via BGP

WITH BMP

- Uni-directional stream of data over a single TCP connection from router to BMP station ('Collector')
- Re-using BGP wire-formats, 'encapsulating' BGP UPDATEs with Per-Peer Headers
- Additional information available: Statistics Reports with various counters, session related configuration




AS1



AS2



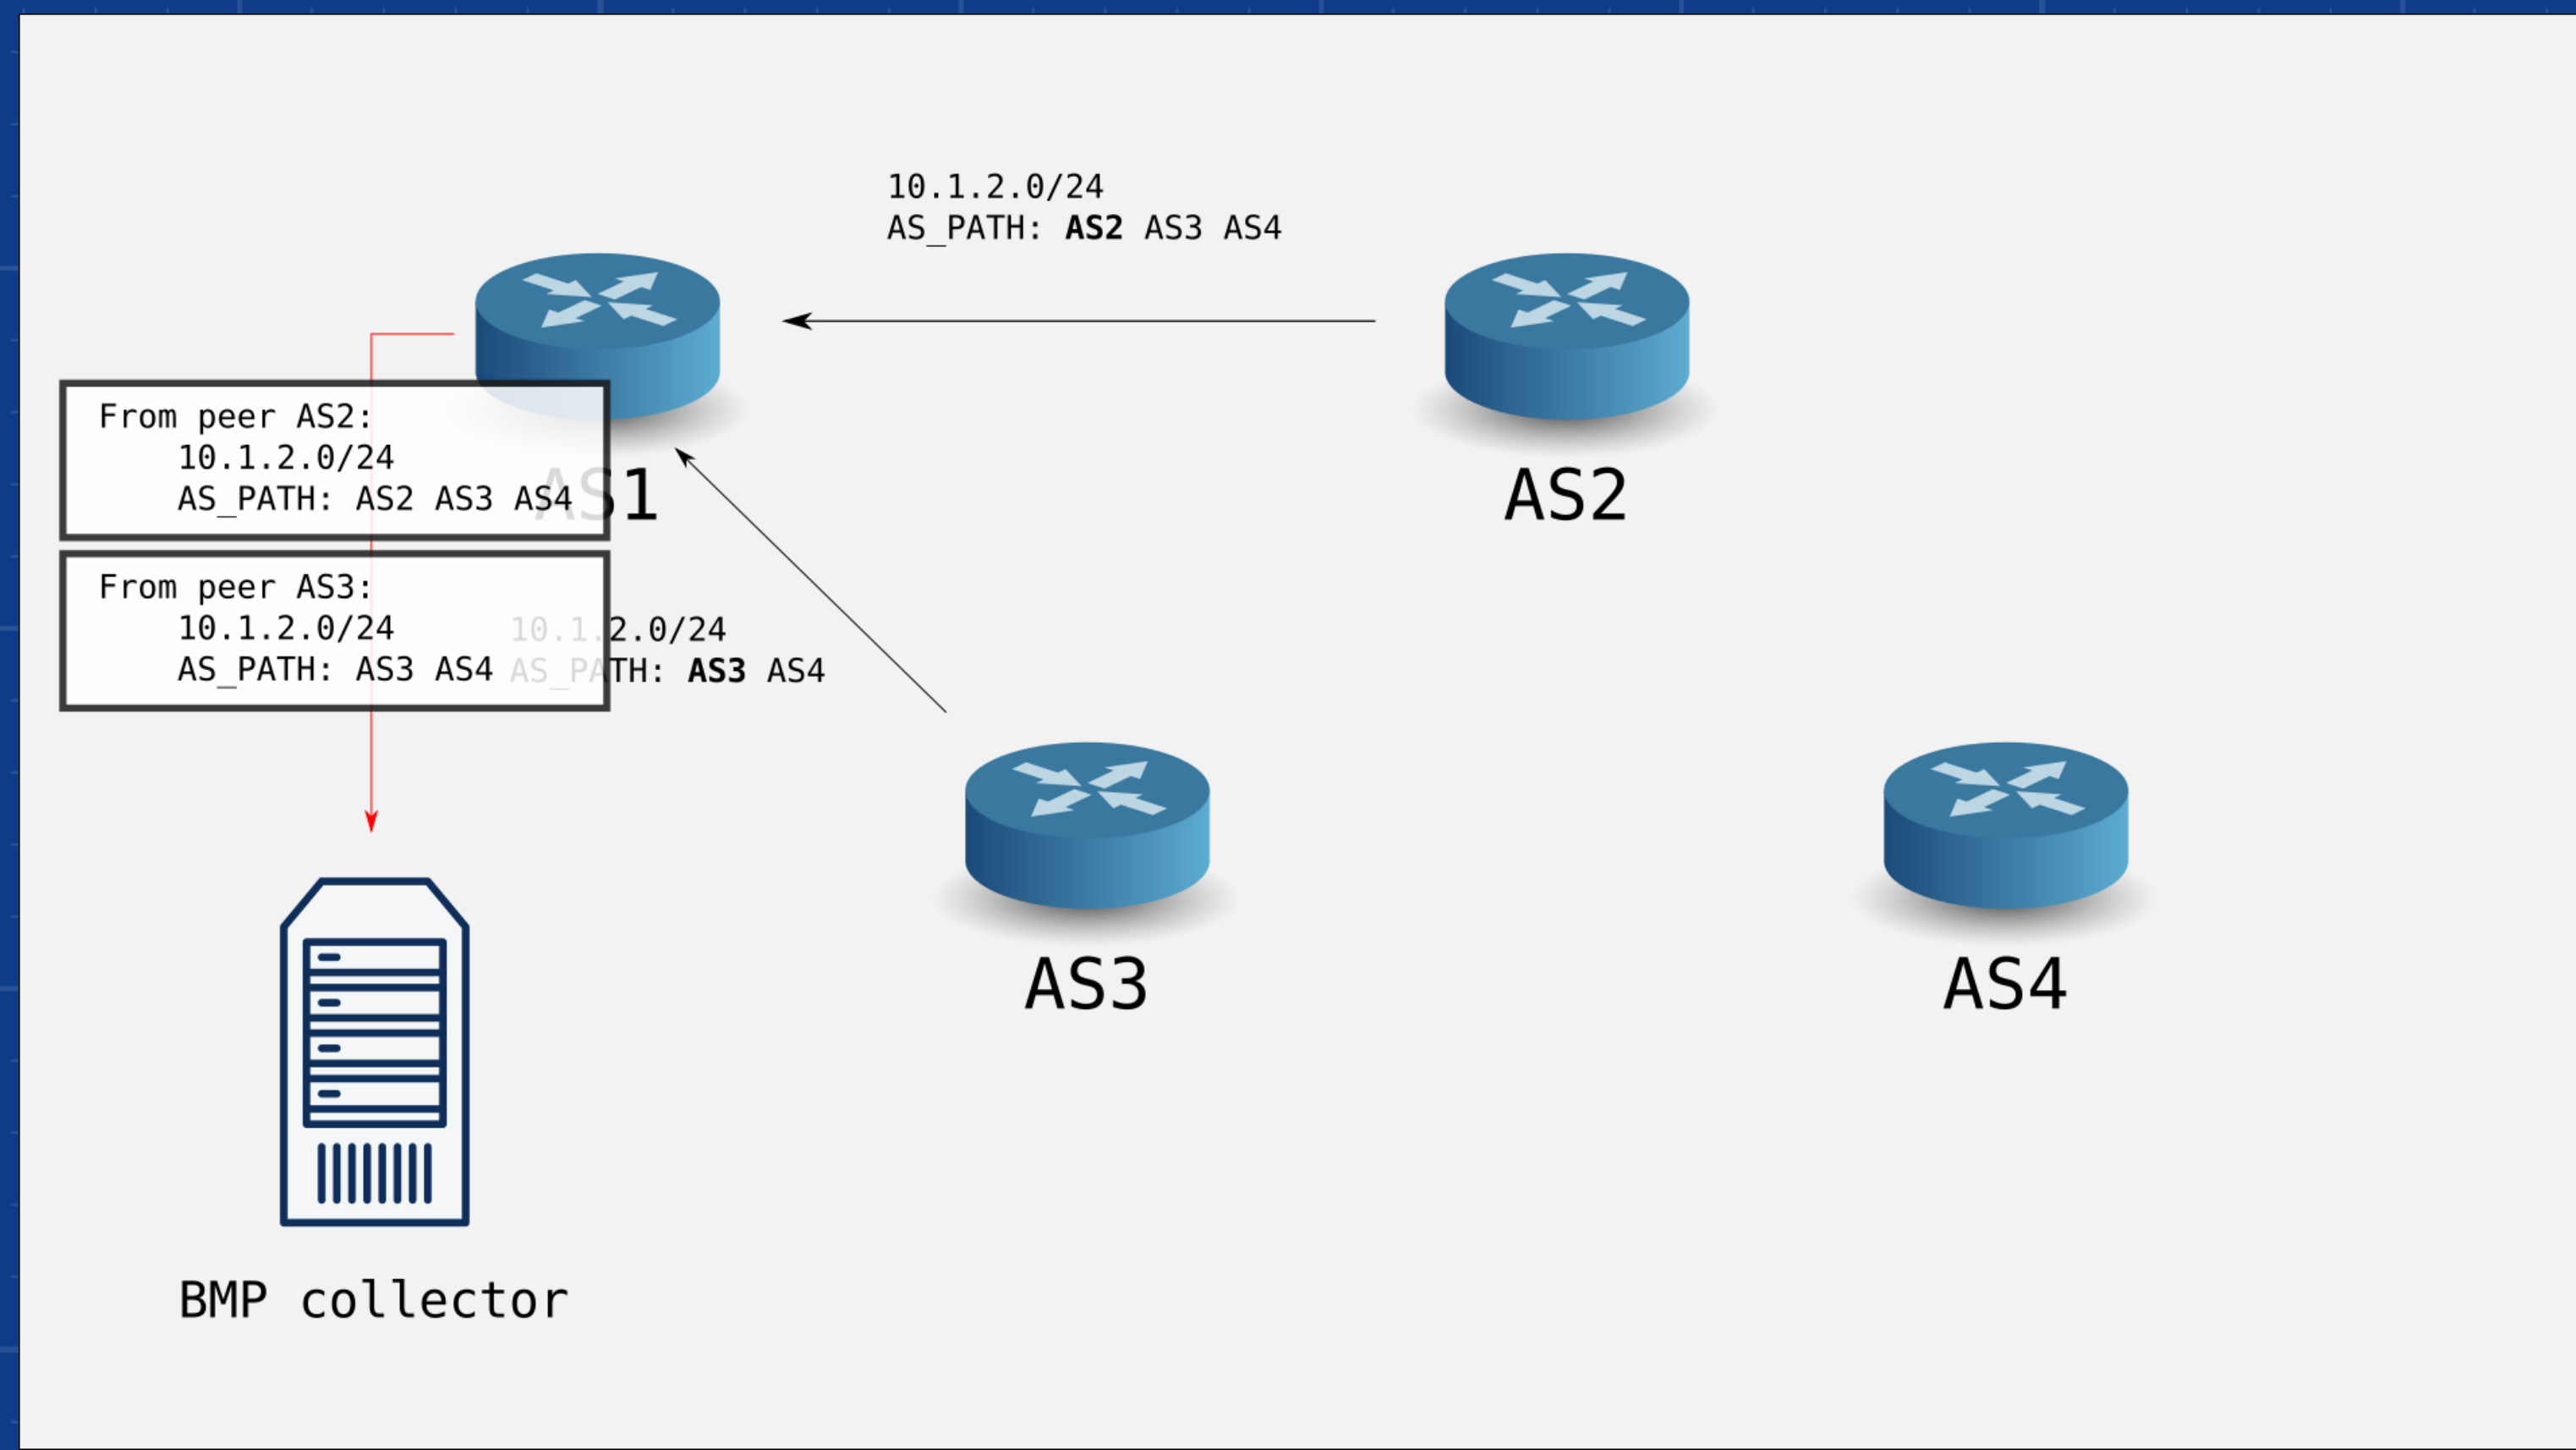
AS3

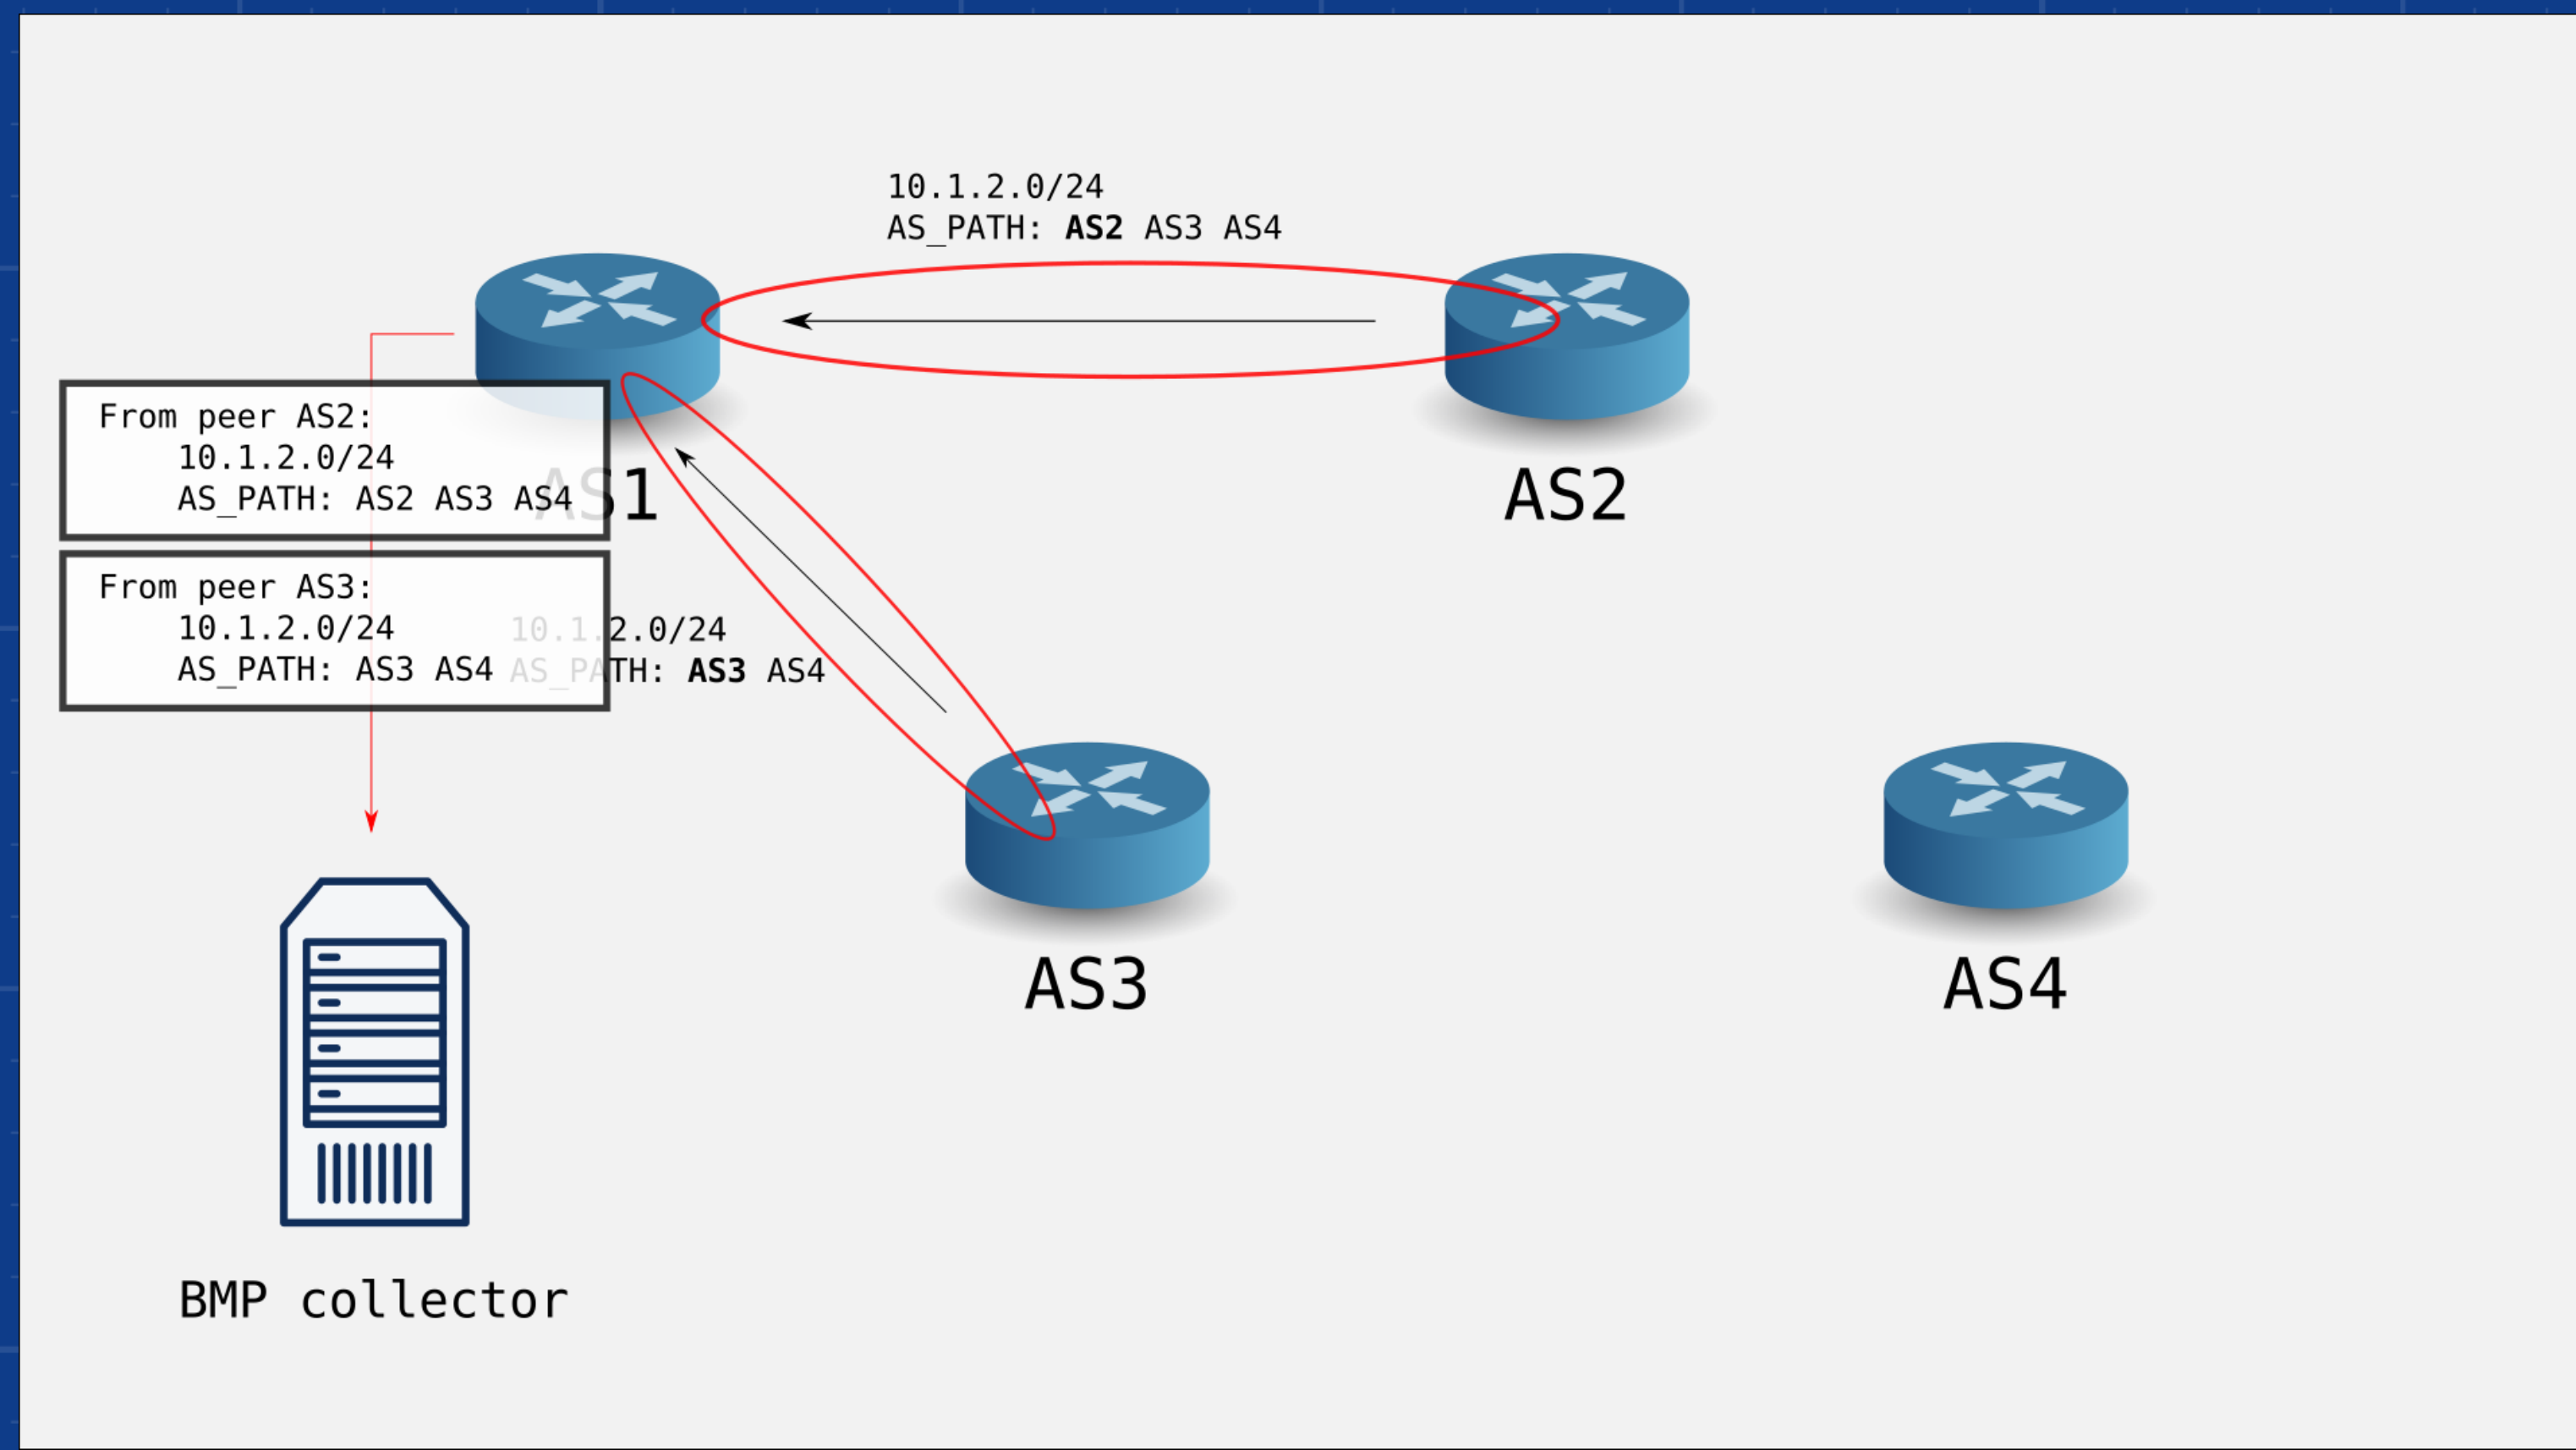


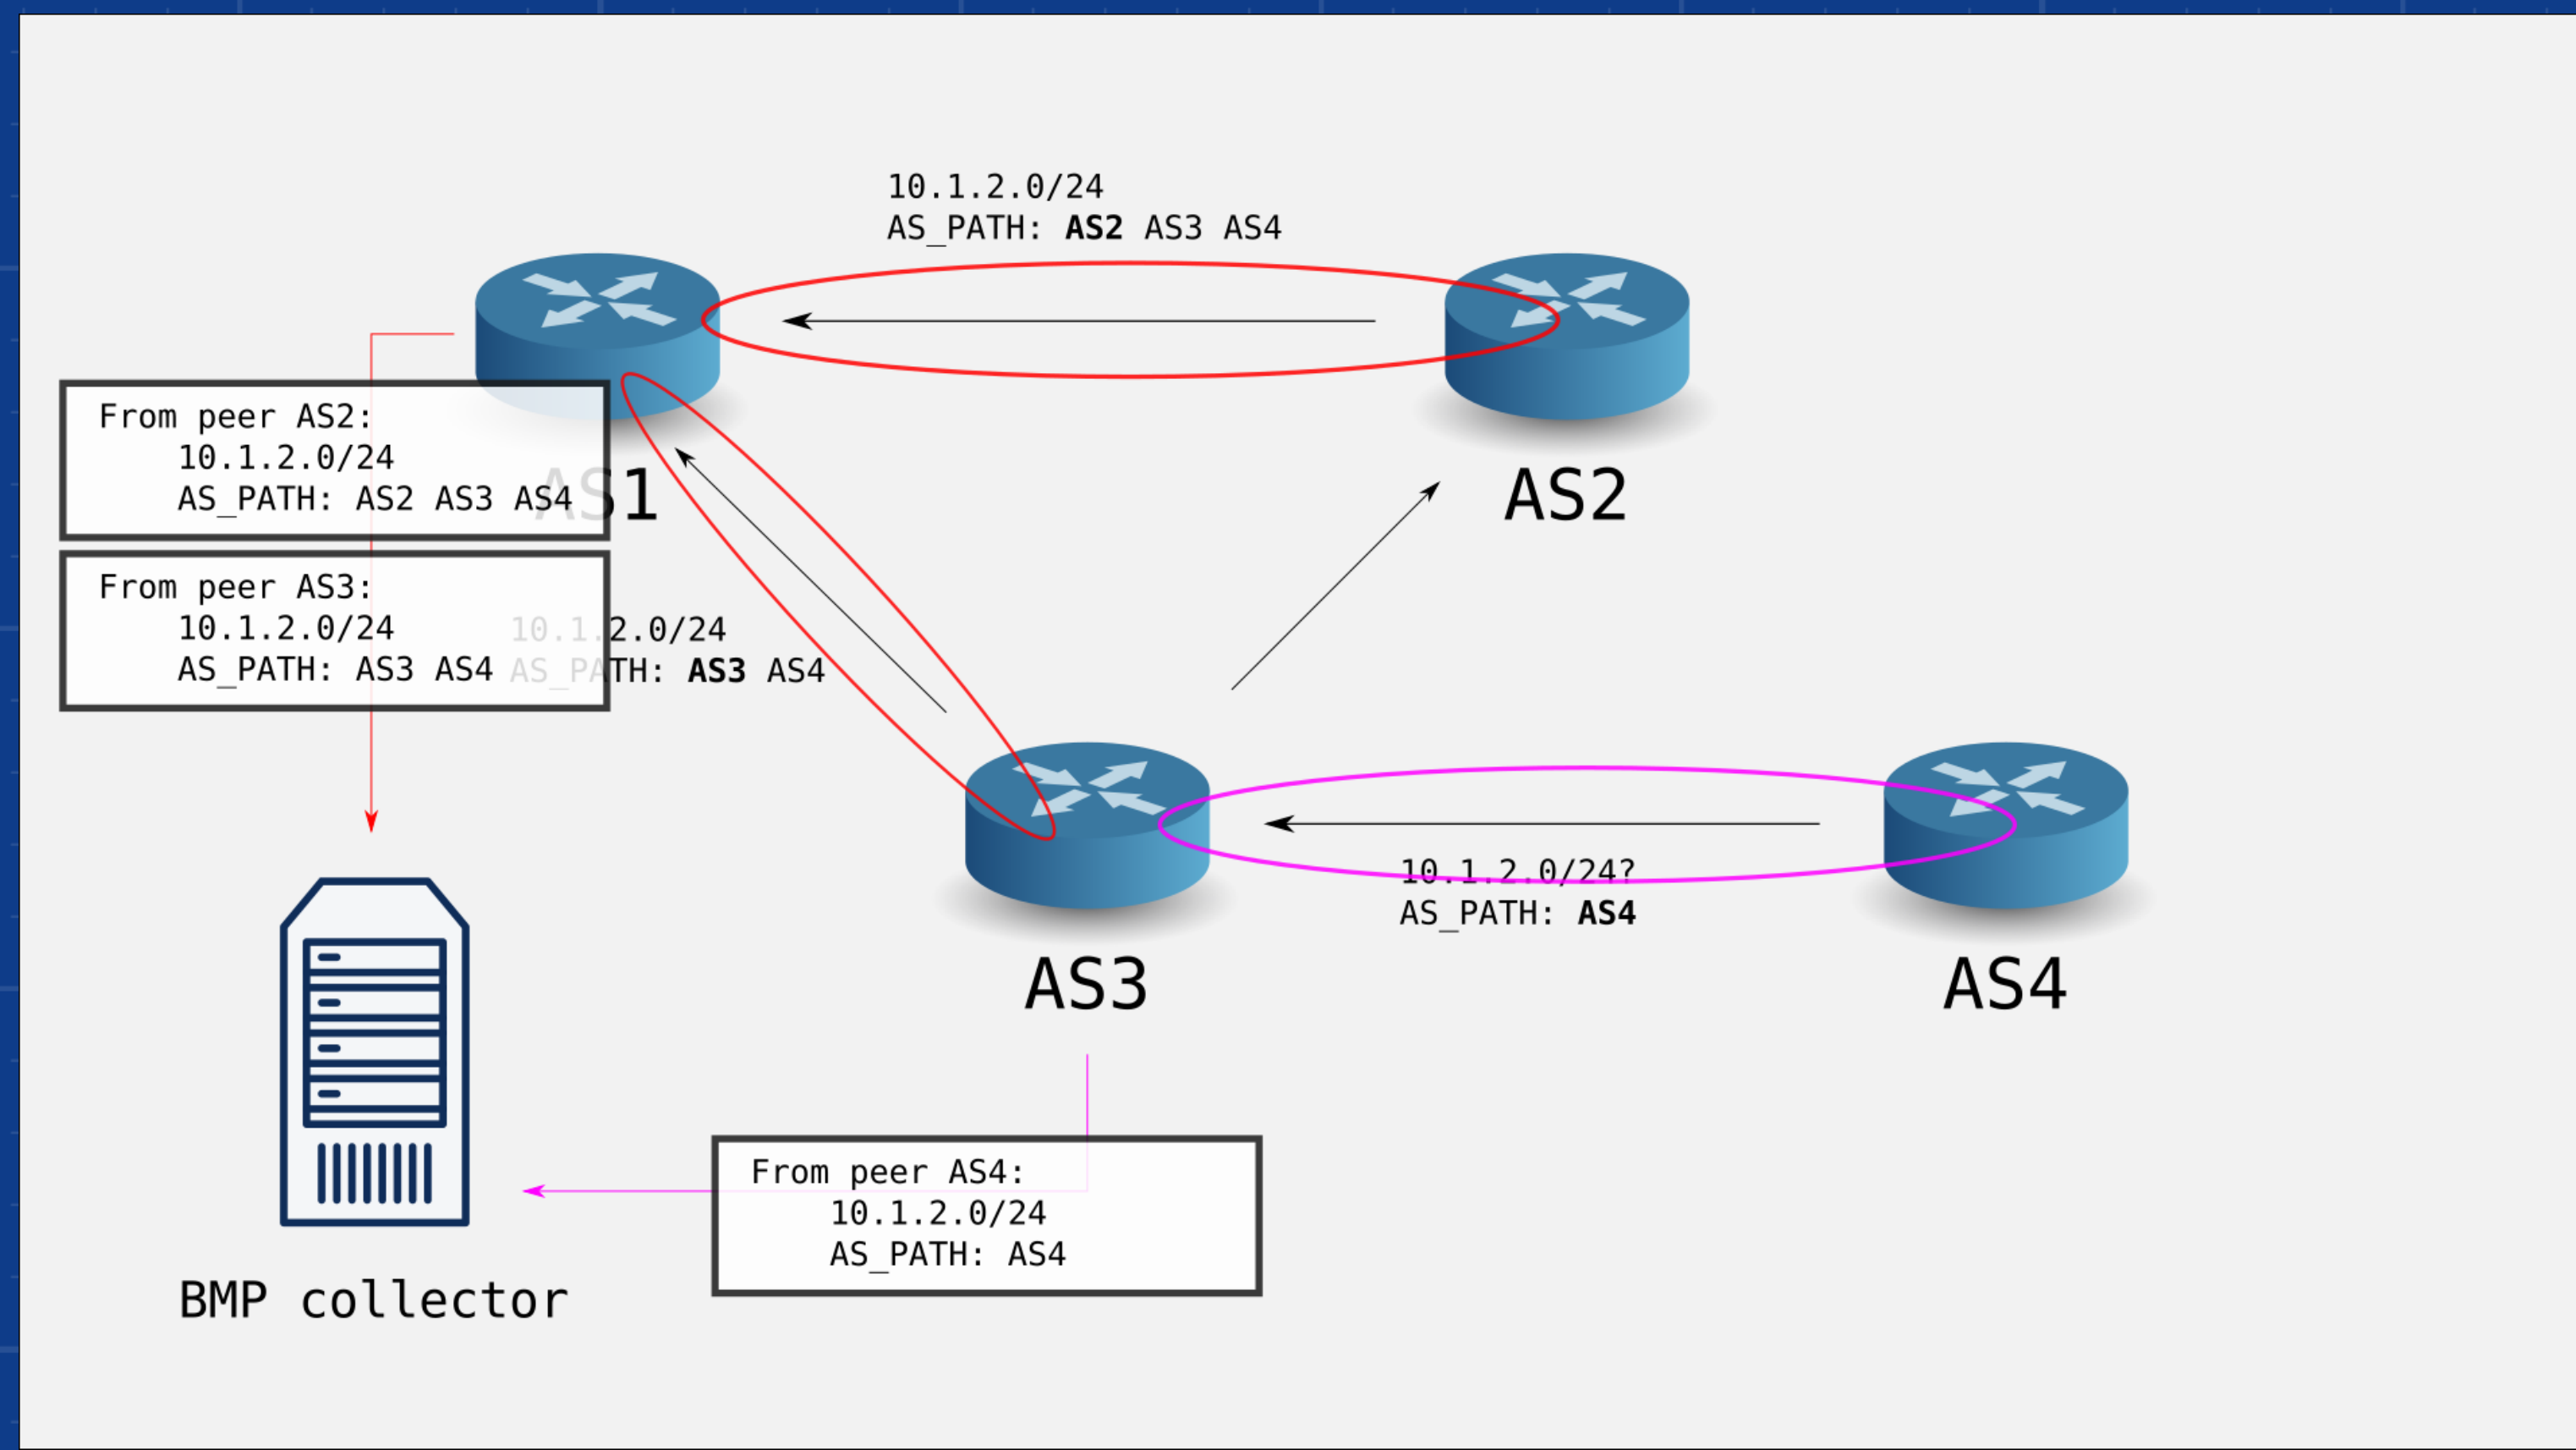
AS4



BMP collector







**WHAT
HAPPENED
BEFORE
"STEADY STATE"?**

WITH BMP

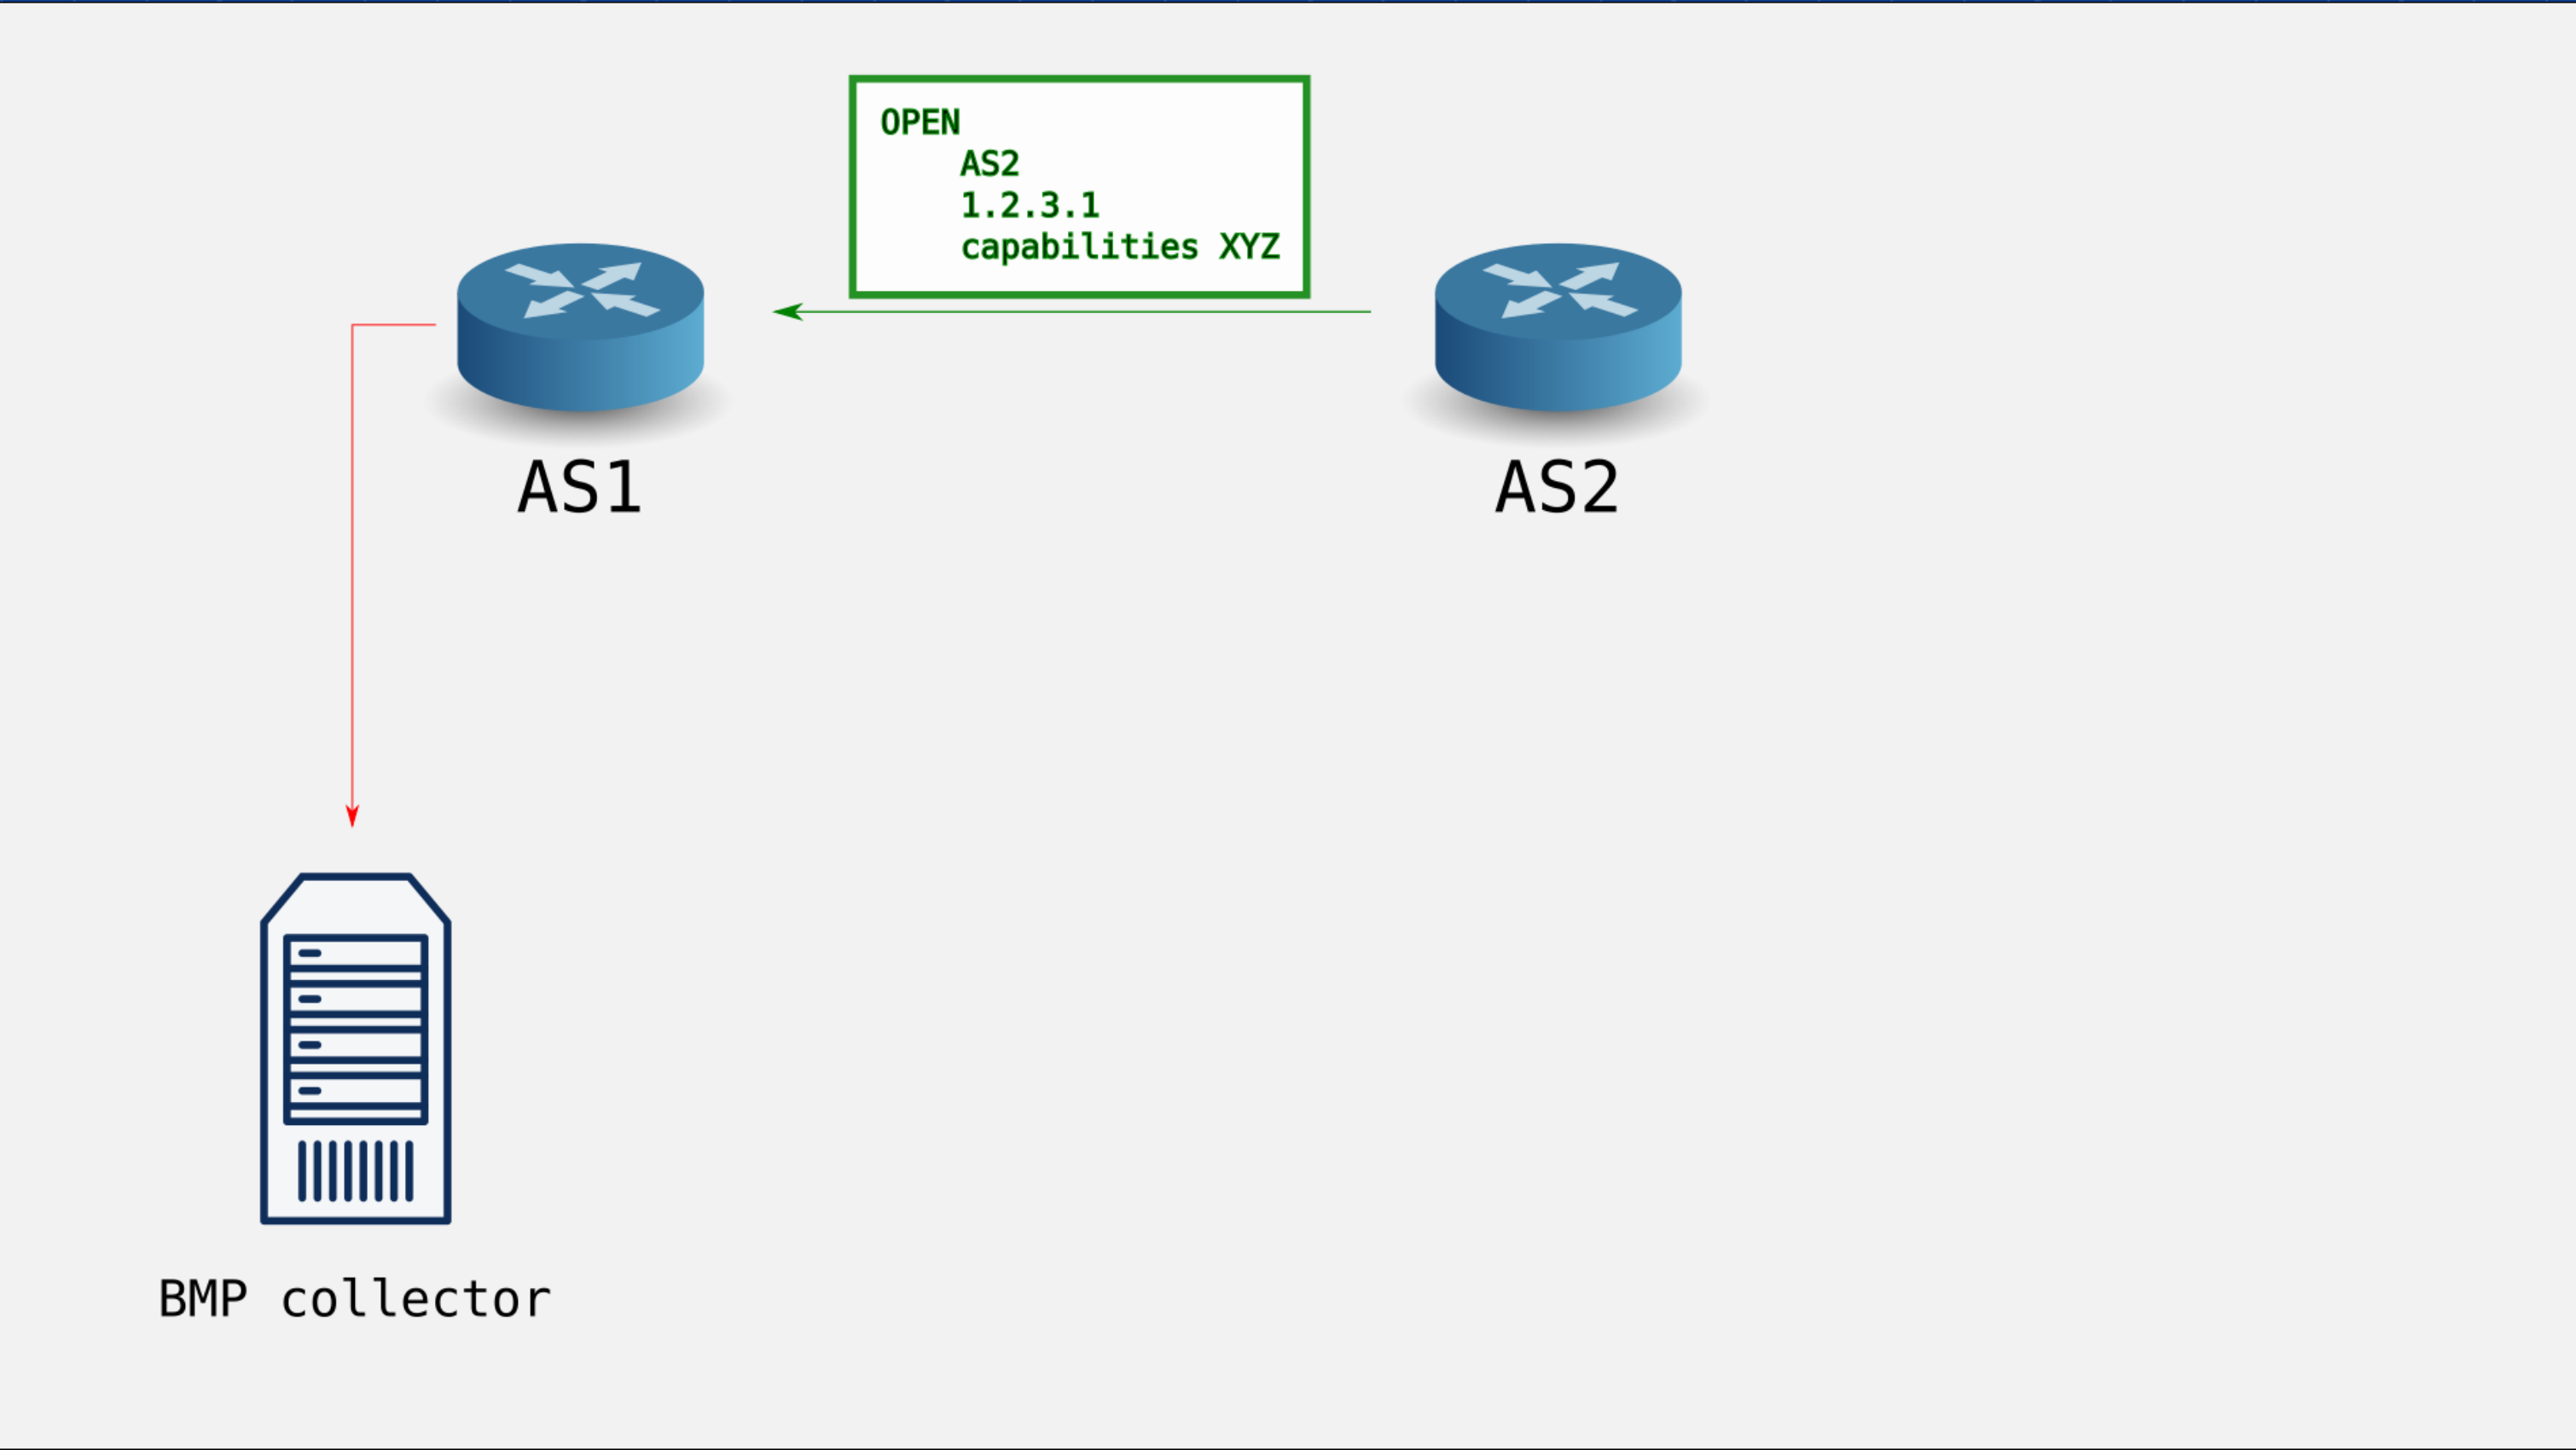
- What about the BGP UPDATEs exchanged before our BMP session was established?
- What happens when a new BGP session between our monitored router and a peer is established?

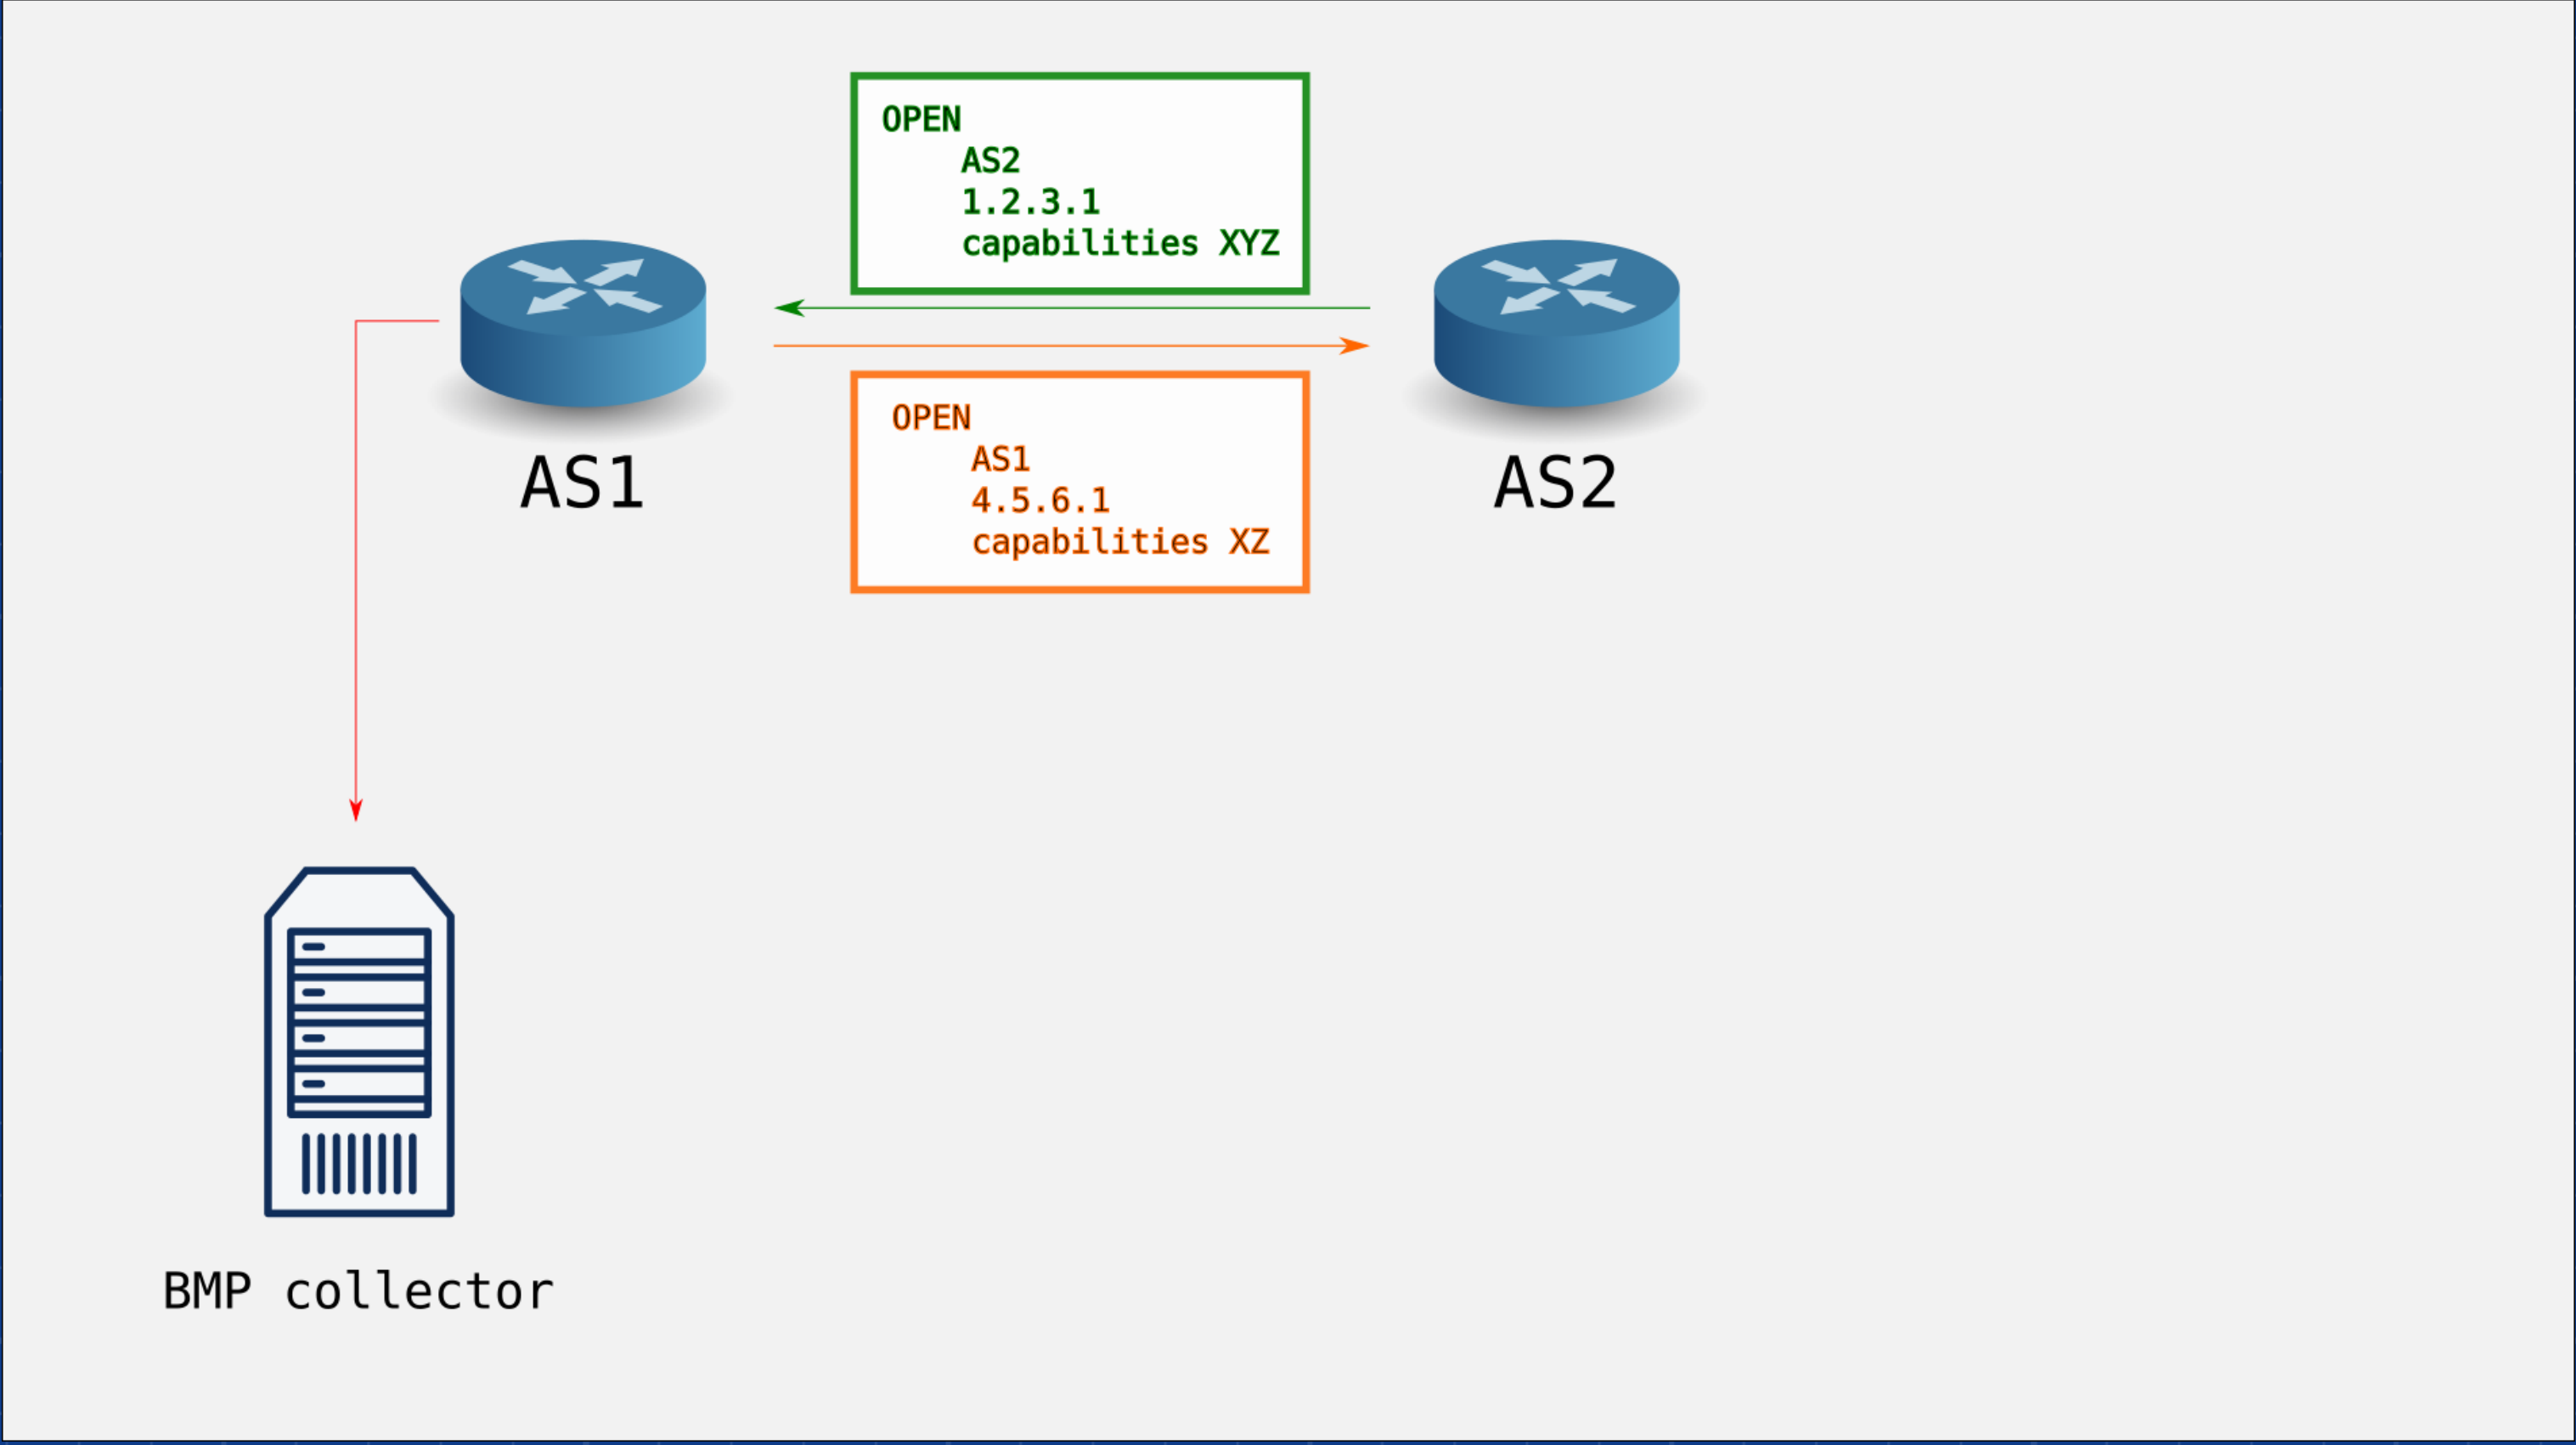
WITH BMP

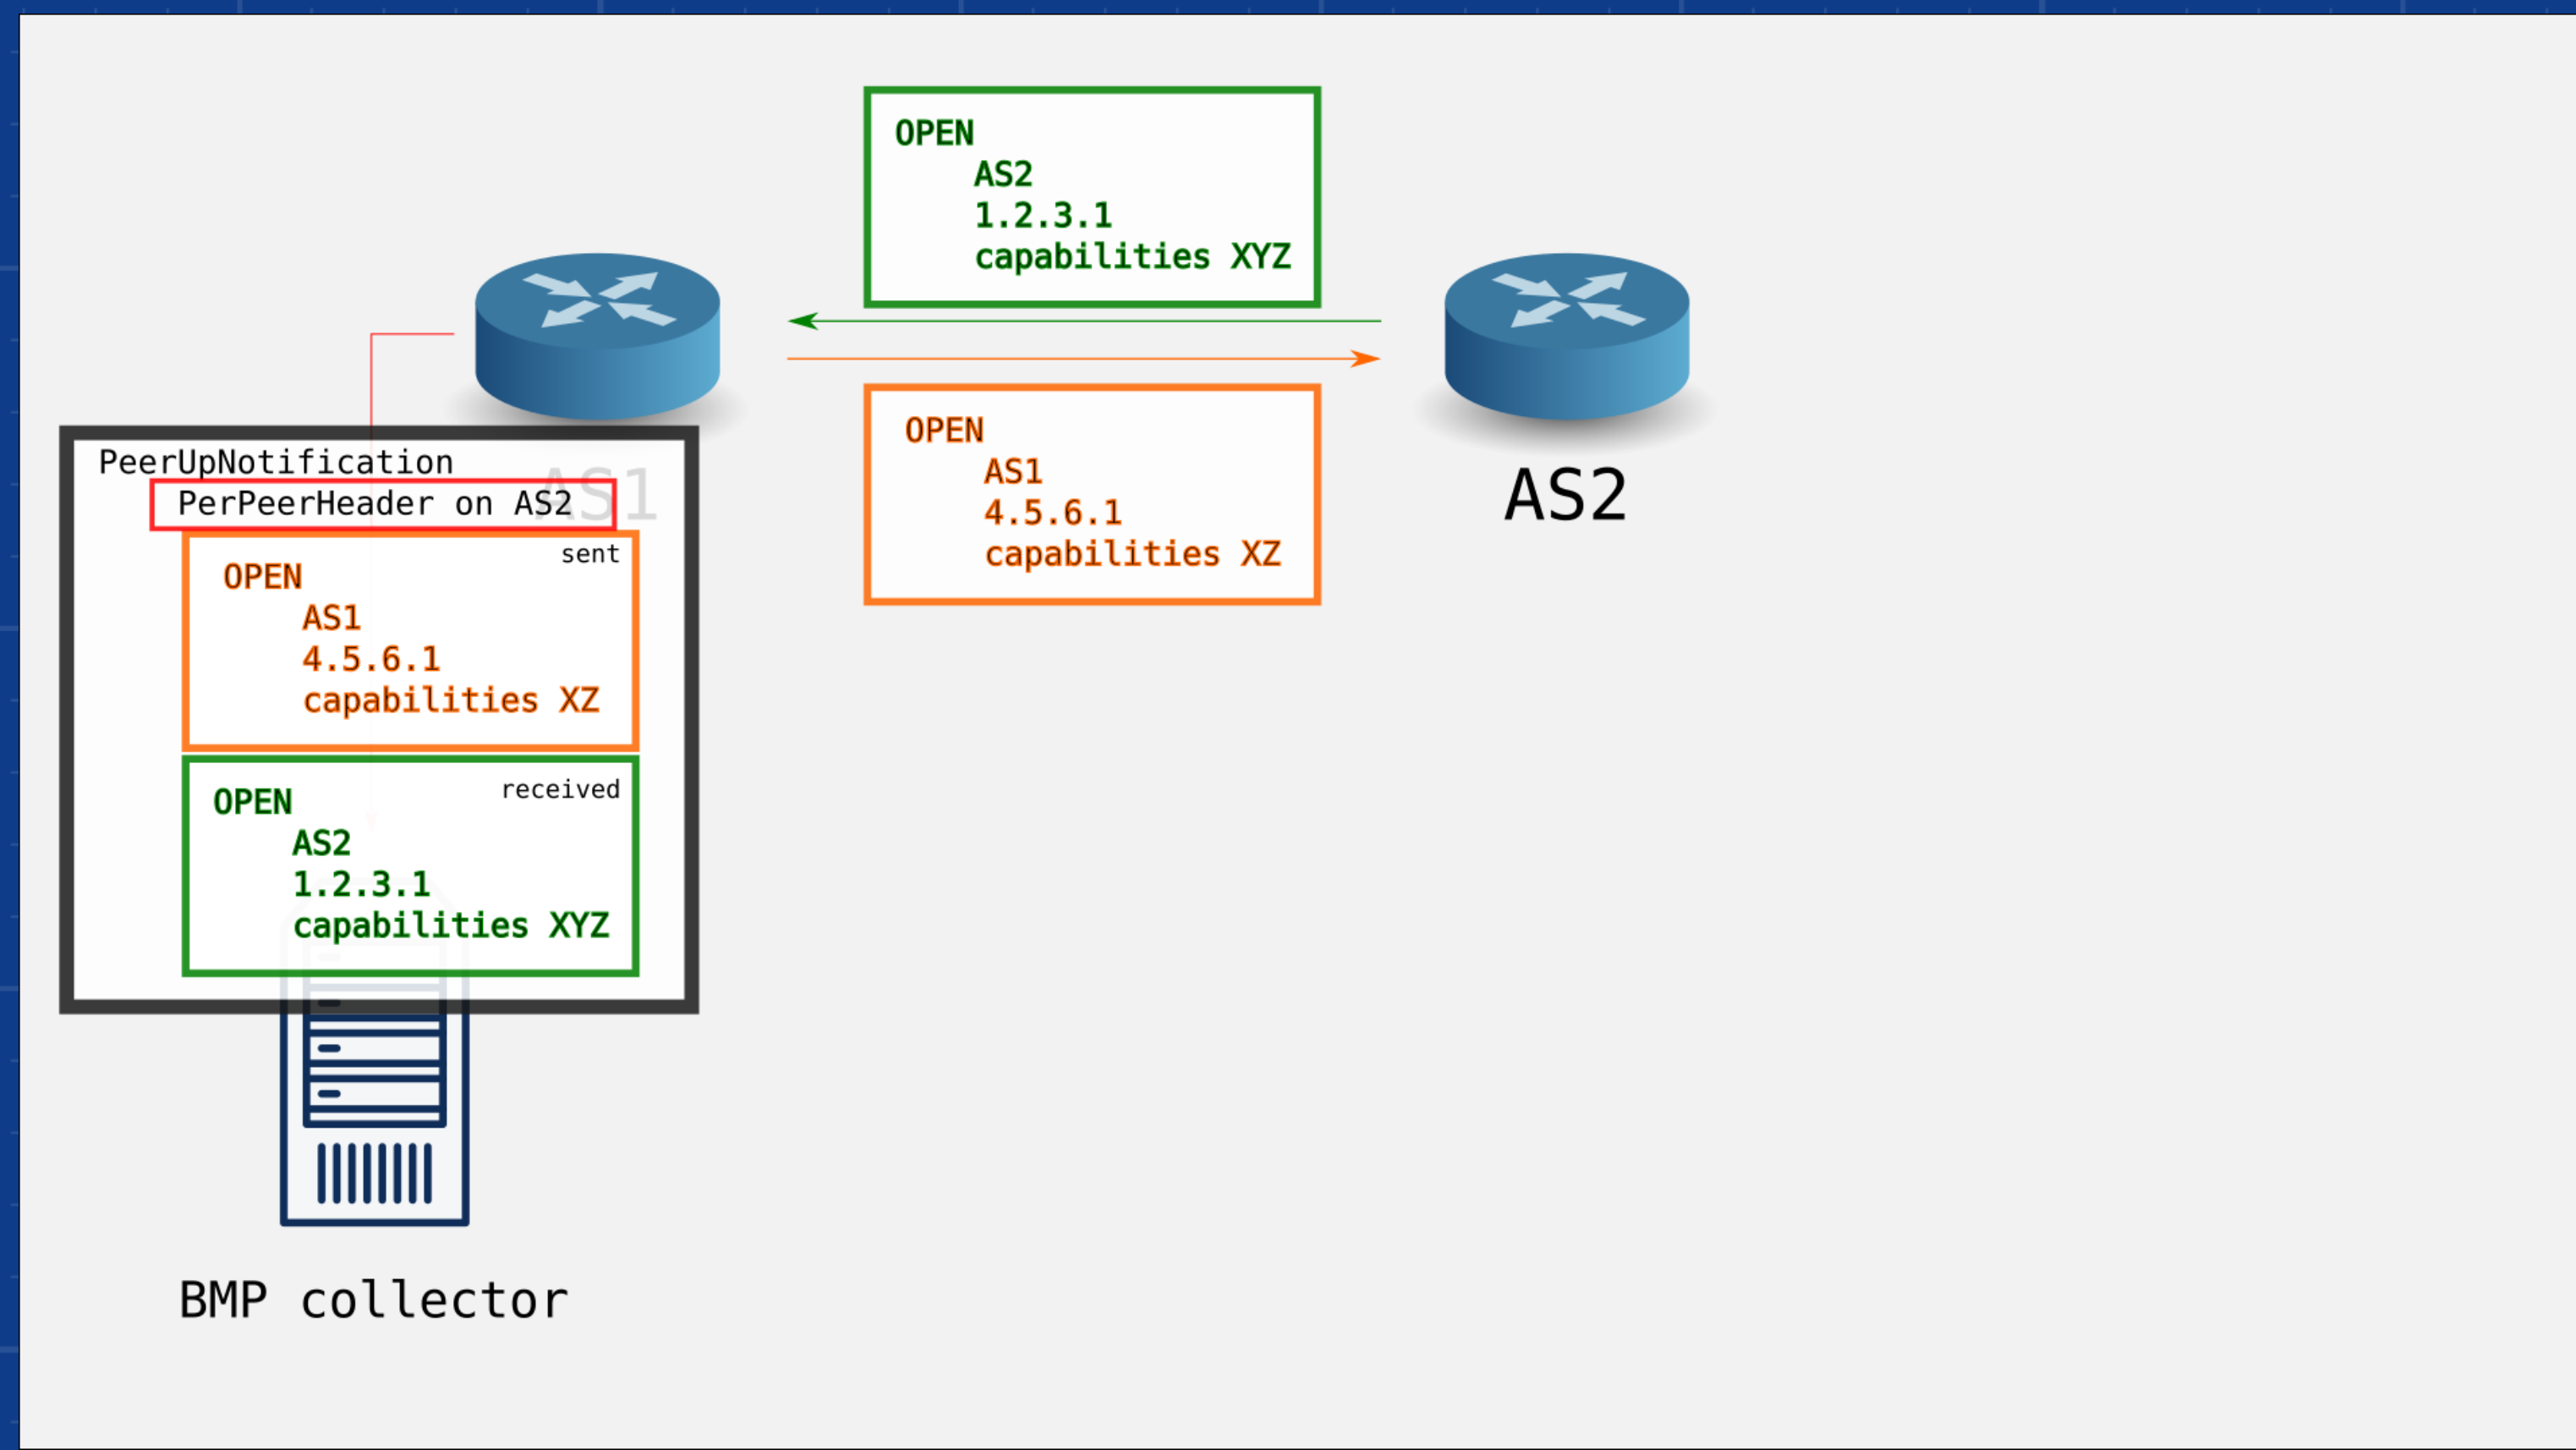
All contents of the monitored RIB are exported to the collector using the same BMP messages as when a new BGP sessions comes up:

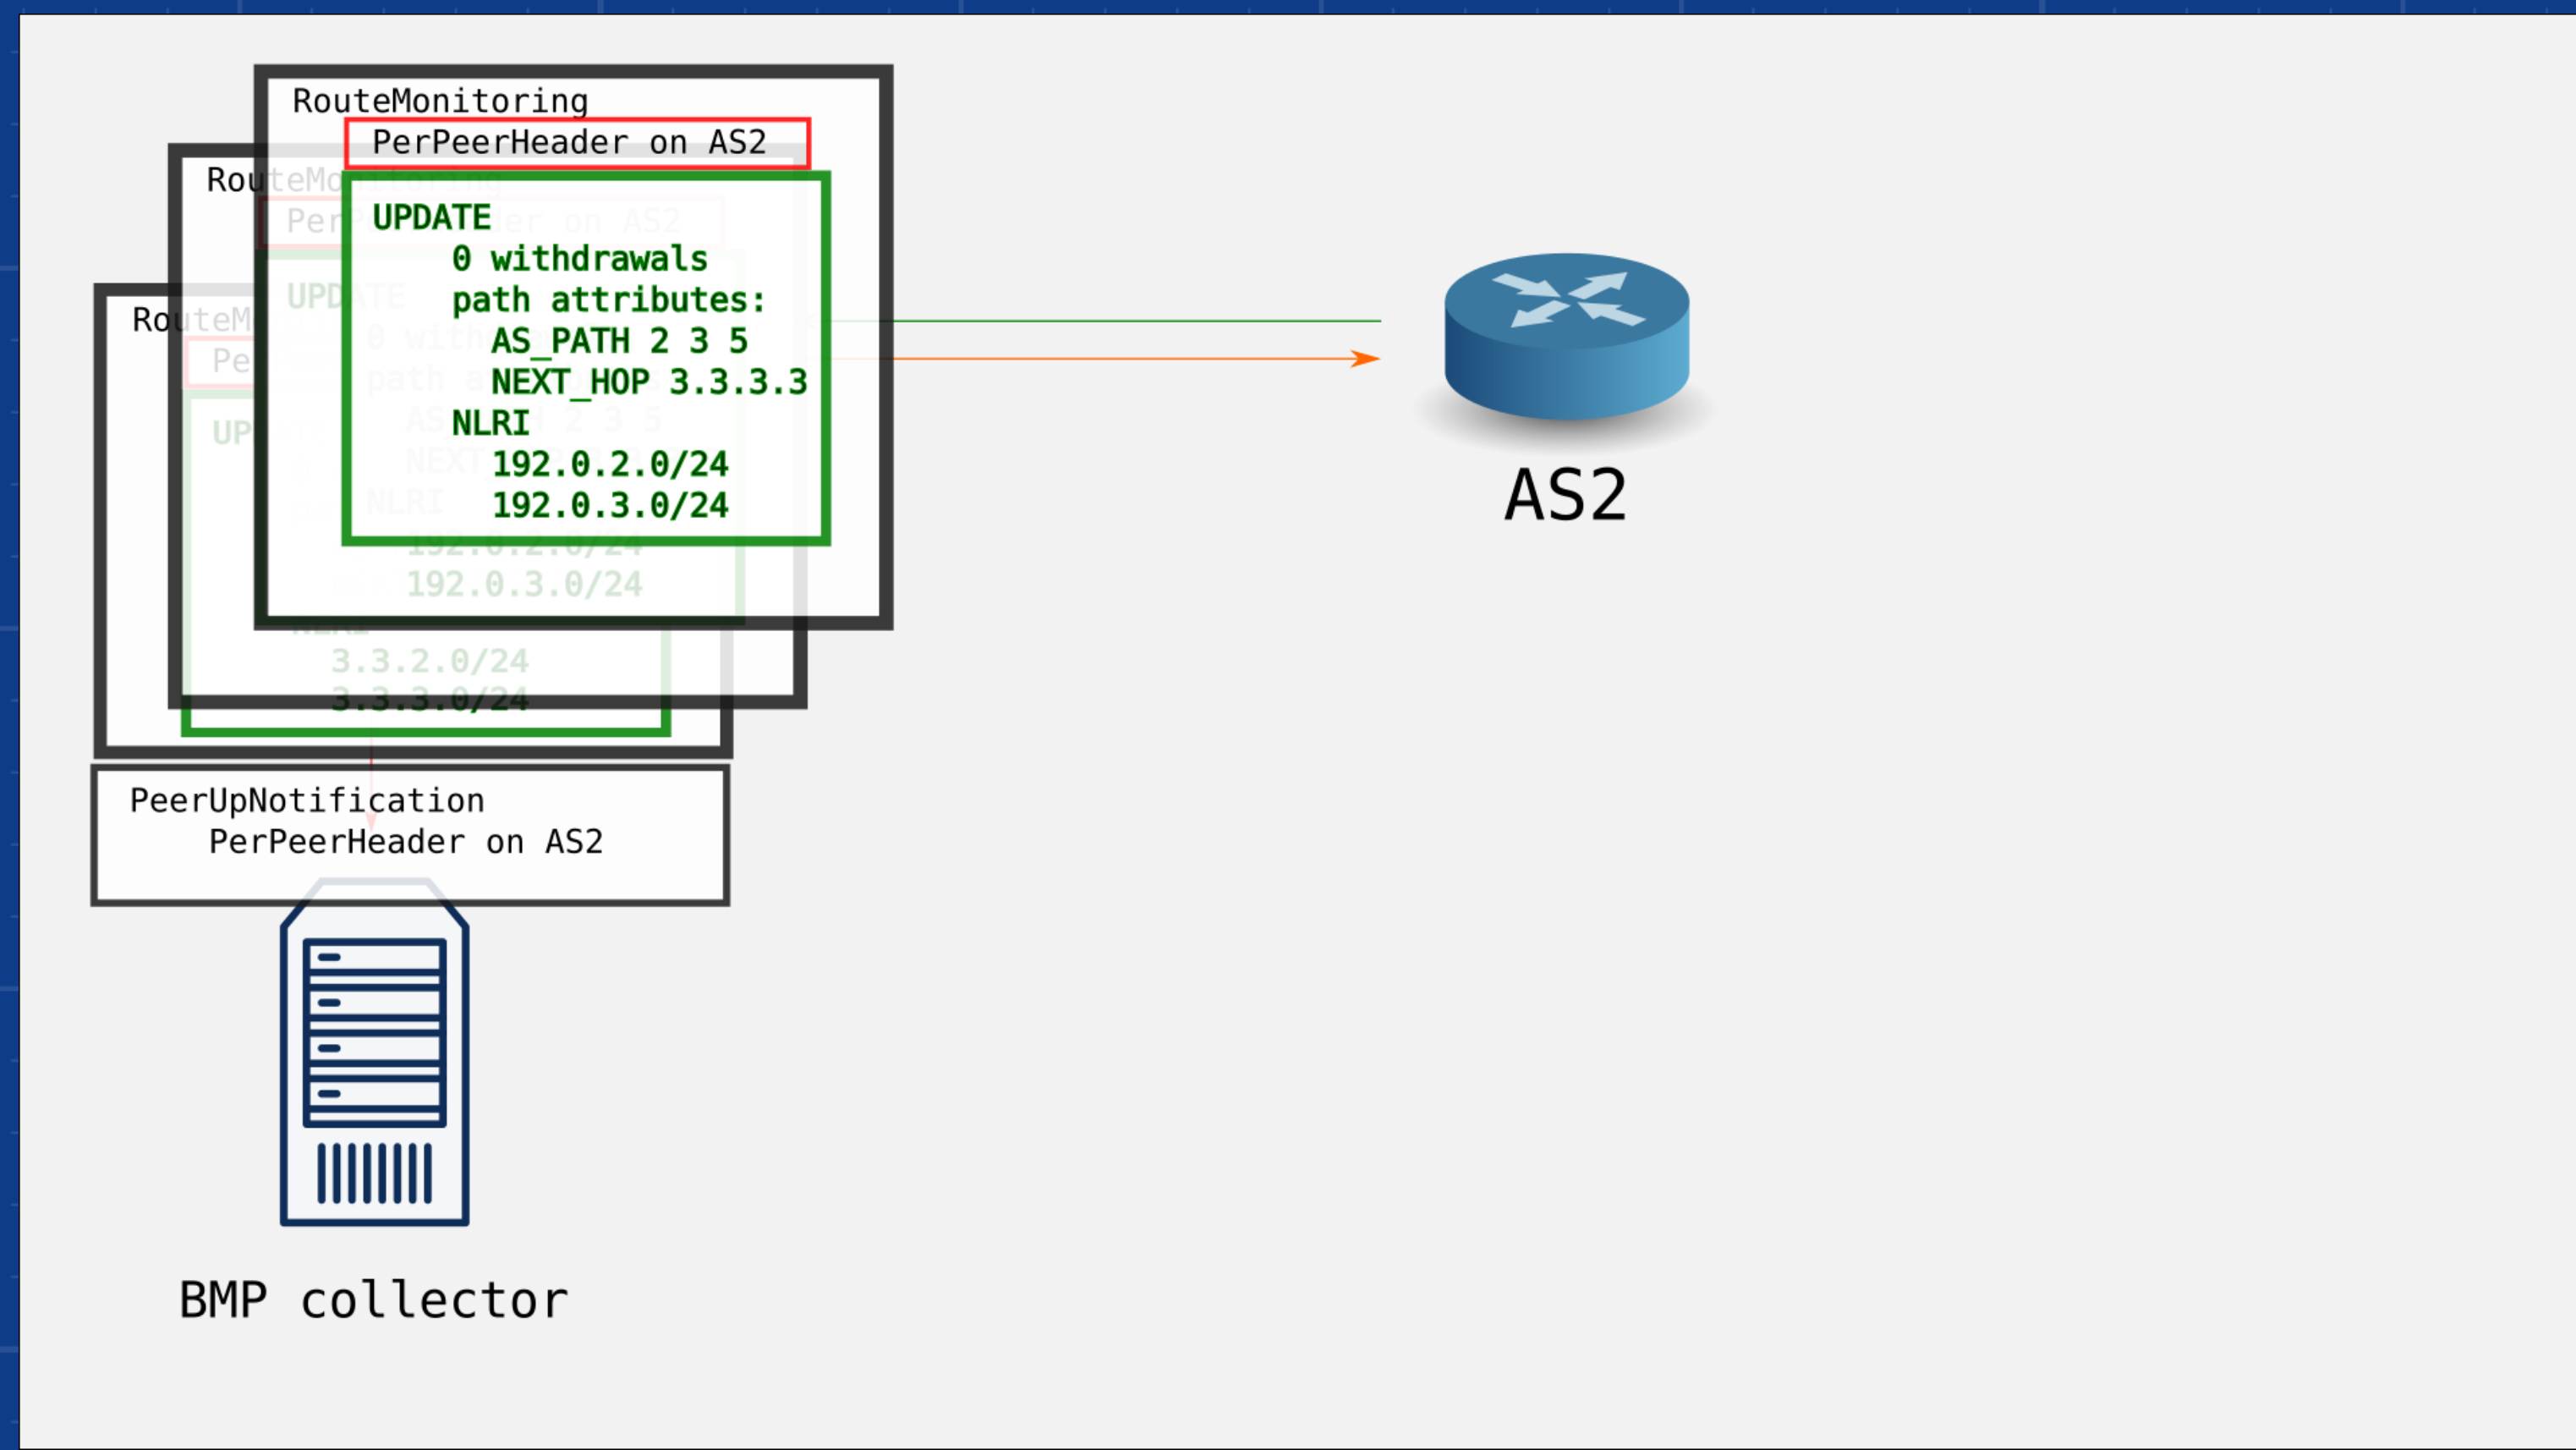
```
PEERUPNOTIFICATION,  
[ROUTEMONITORINGMESSAGE, ..],  
RM-WITH-EOR,  
STEADY-STATE
```











**OUR EXPERIENCES
SO FAR**

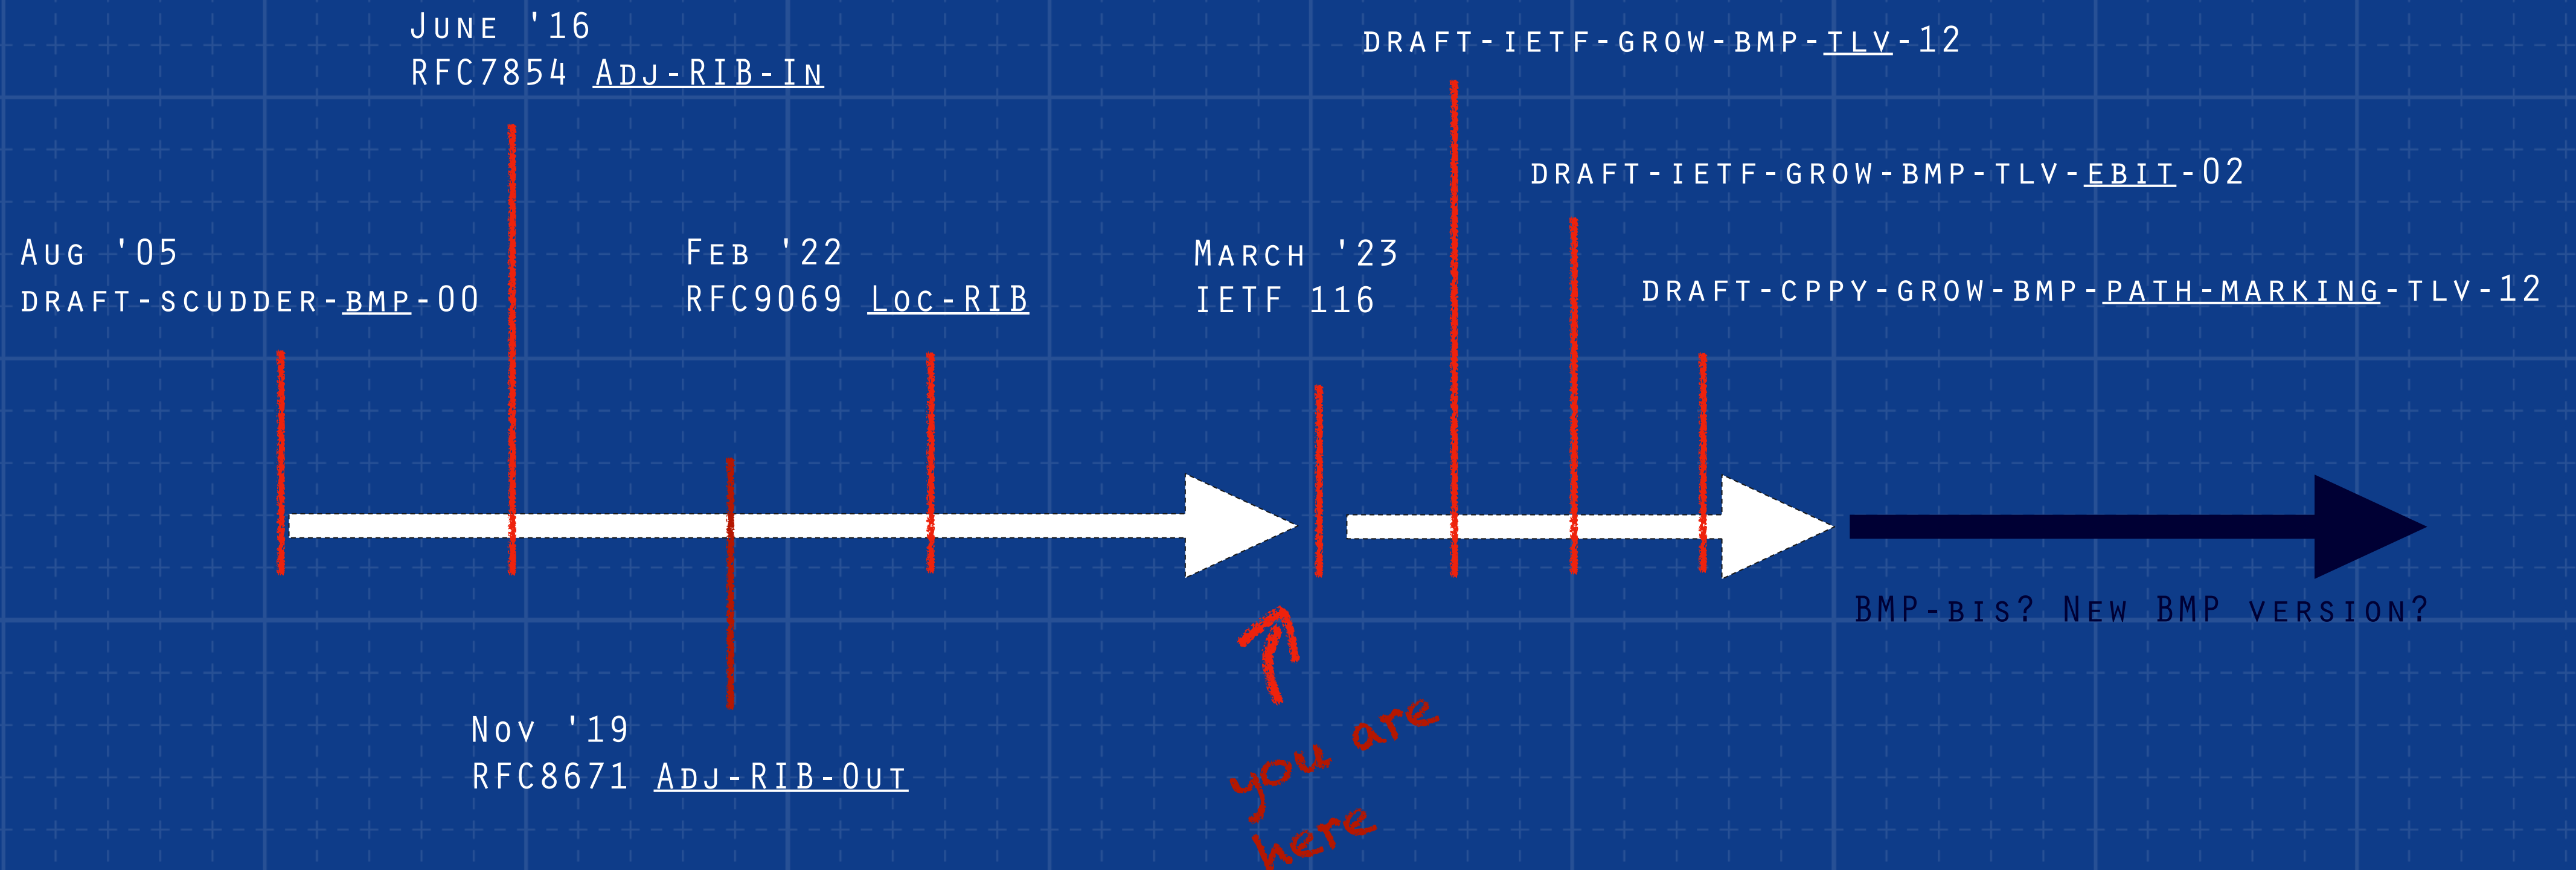
THE GOOD STUFF


- Monitoring ~250 production routers from a Tier 1 provider, three different vendors, it works.
- Easy (re-)synchronisation of monitored RIBs without tears
- Re-use of BGP wire-formats speeds up development

THE COULD-BE-NICER STUFF

- Not all parsing can be done stateless, information from PeerUpNotifications needed to process RouteMonitoring Messages. (inherited from BGP)
- EoR signalling is messy and incomplete, or even non-existent in some vendors
- Inconsistencies in (headers in) BMP messages, and textually in standards documents

EVOLUTION OF BMP STANDARD DOCUMENTS





AS1> possible routes to 10.1.2.0/24?

a. AS3 AS4 → Path marked "Best path"

b. AS2 AS3 AS4 → Path marked "Backup" + "Not installed" or Path marked "ROV Invalid"

THANK YOU