



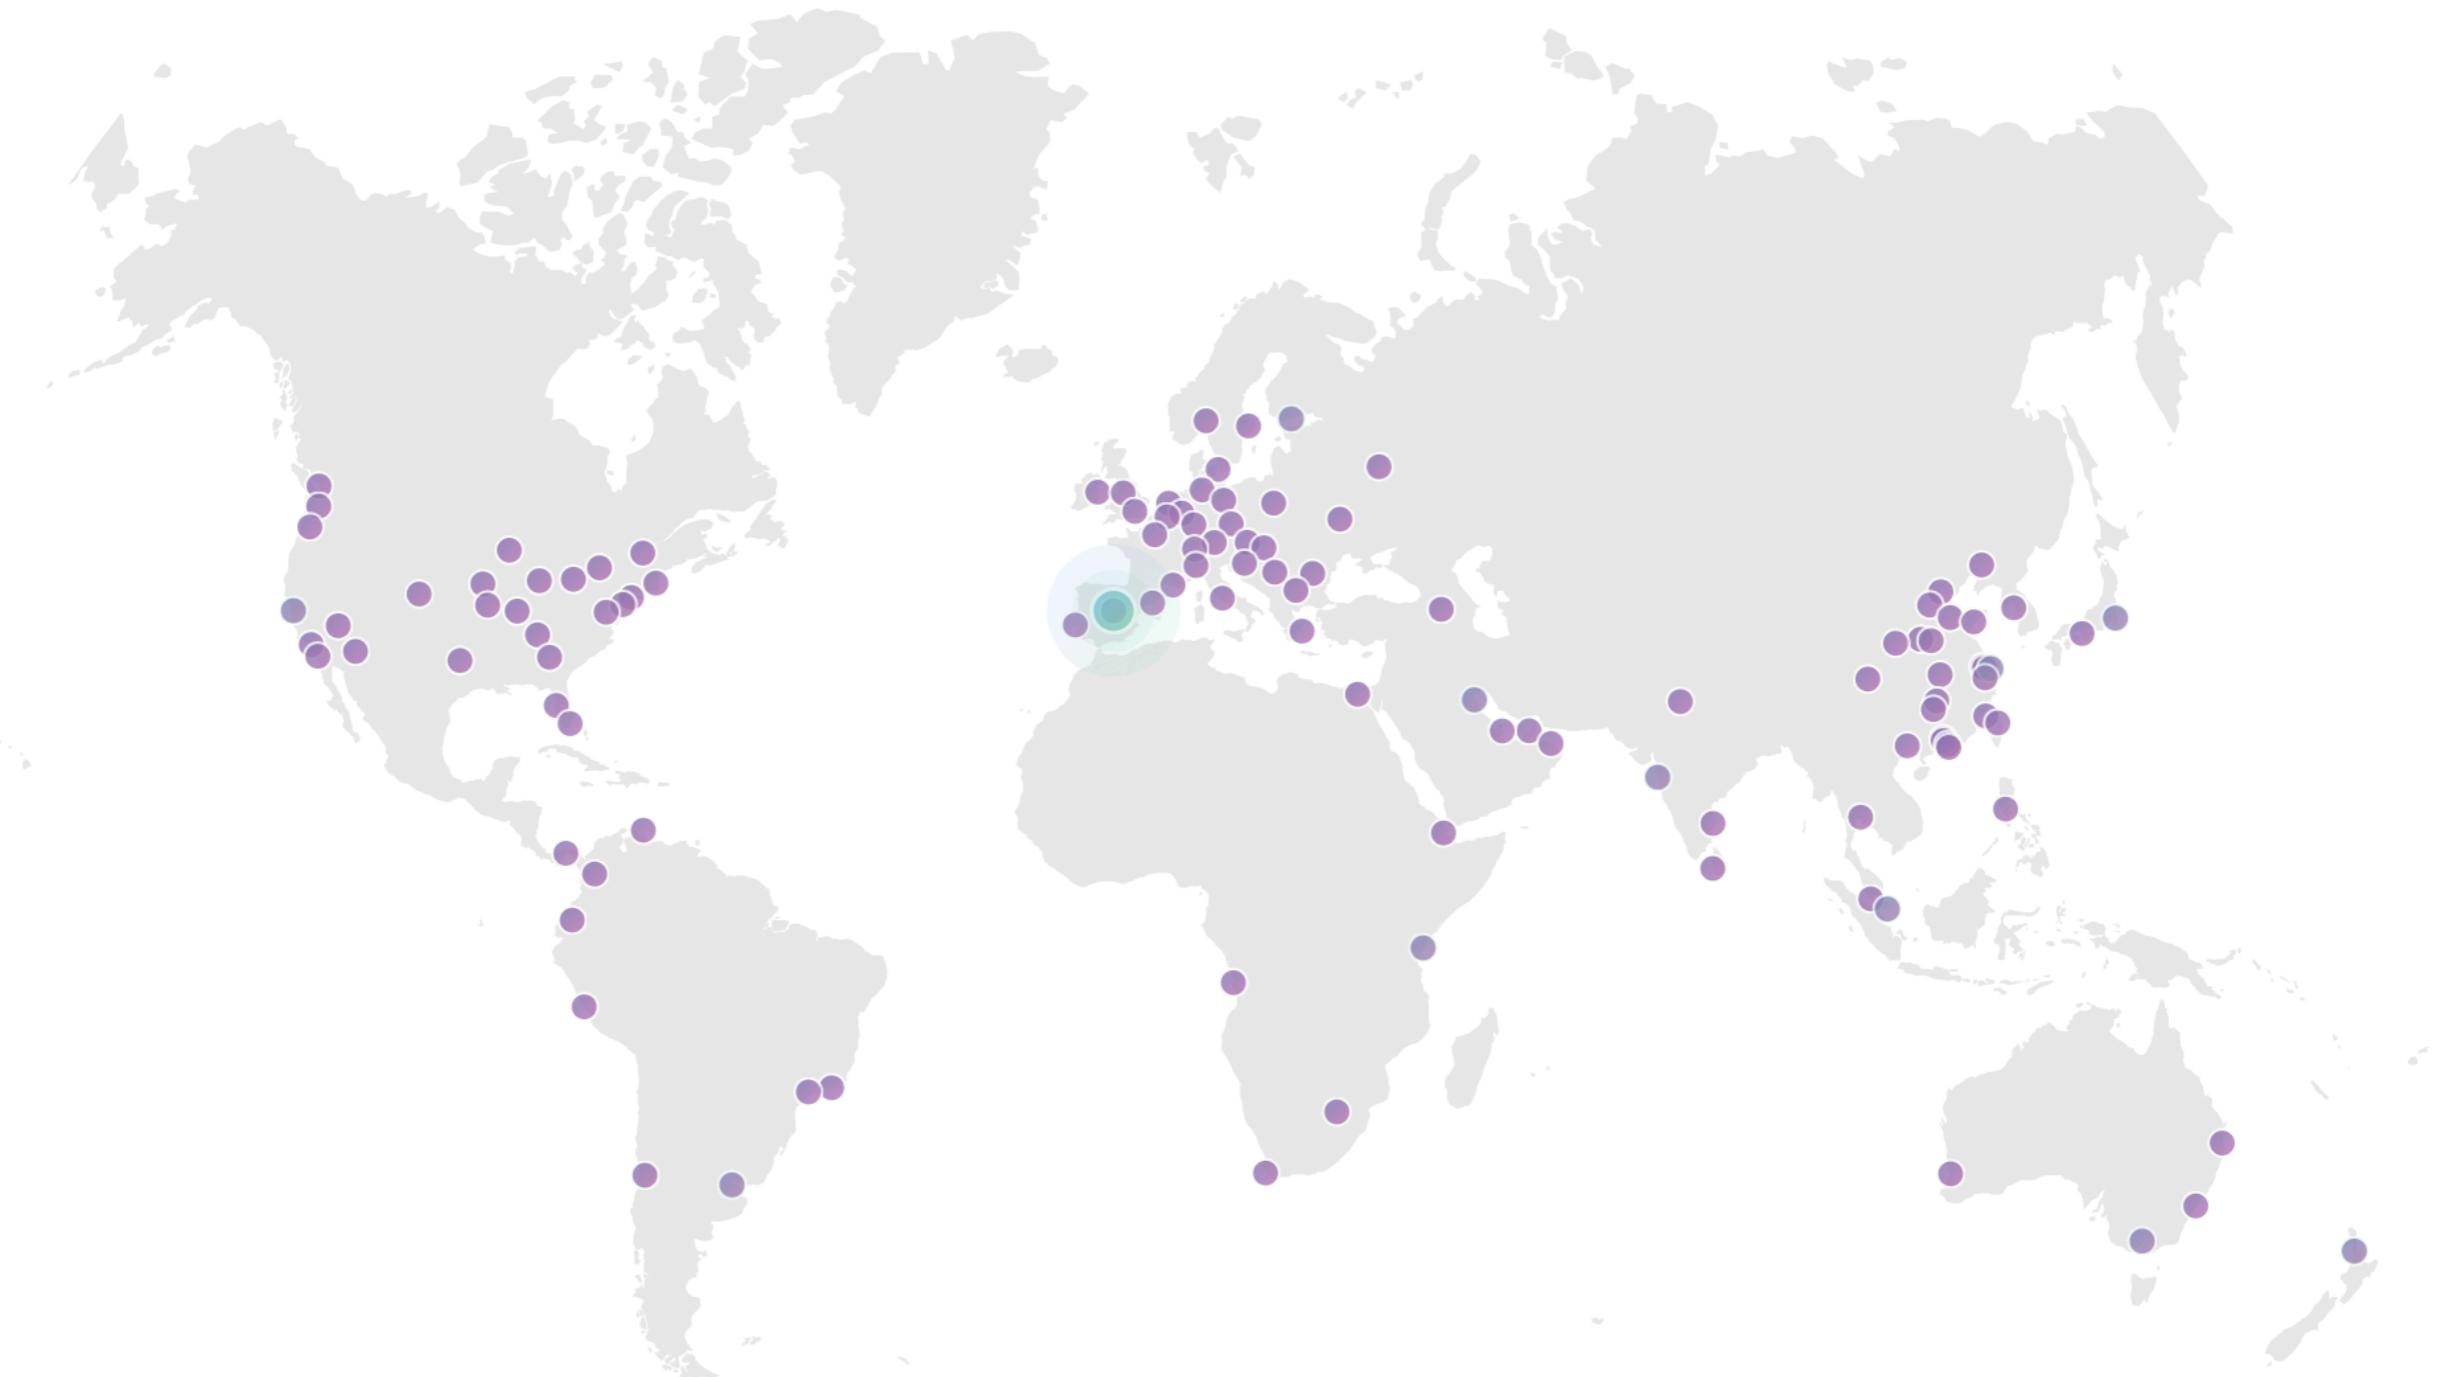
Gatelogic FRP framework

...200 lines of Python I still regret...

Marek Majkowski

@majek04

Who we are?

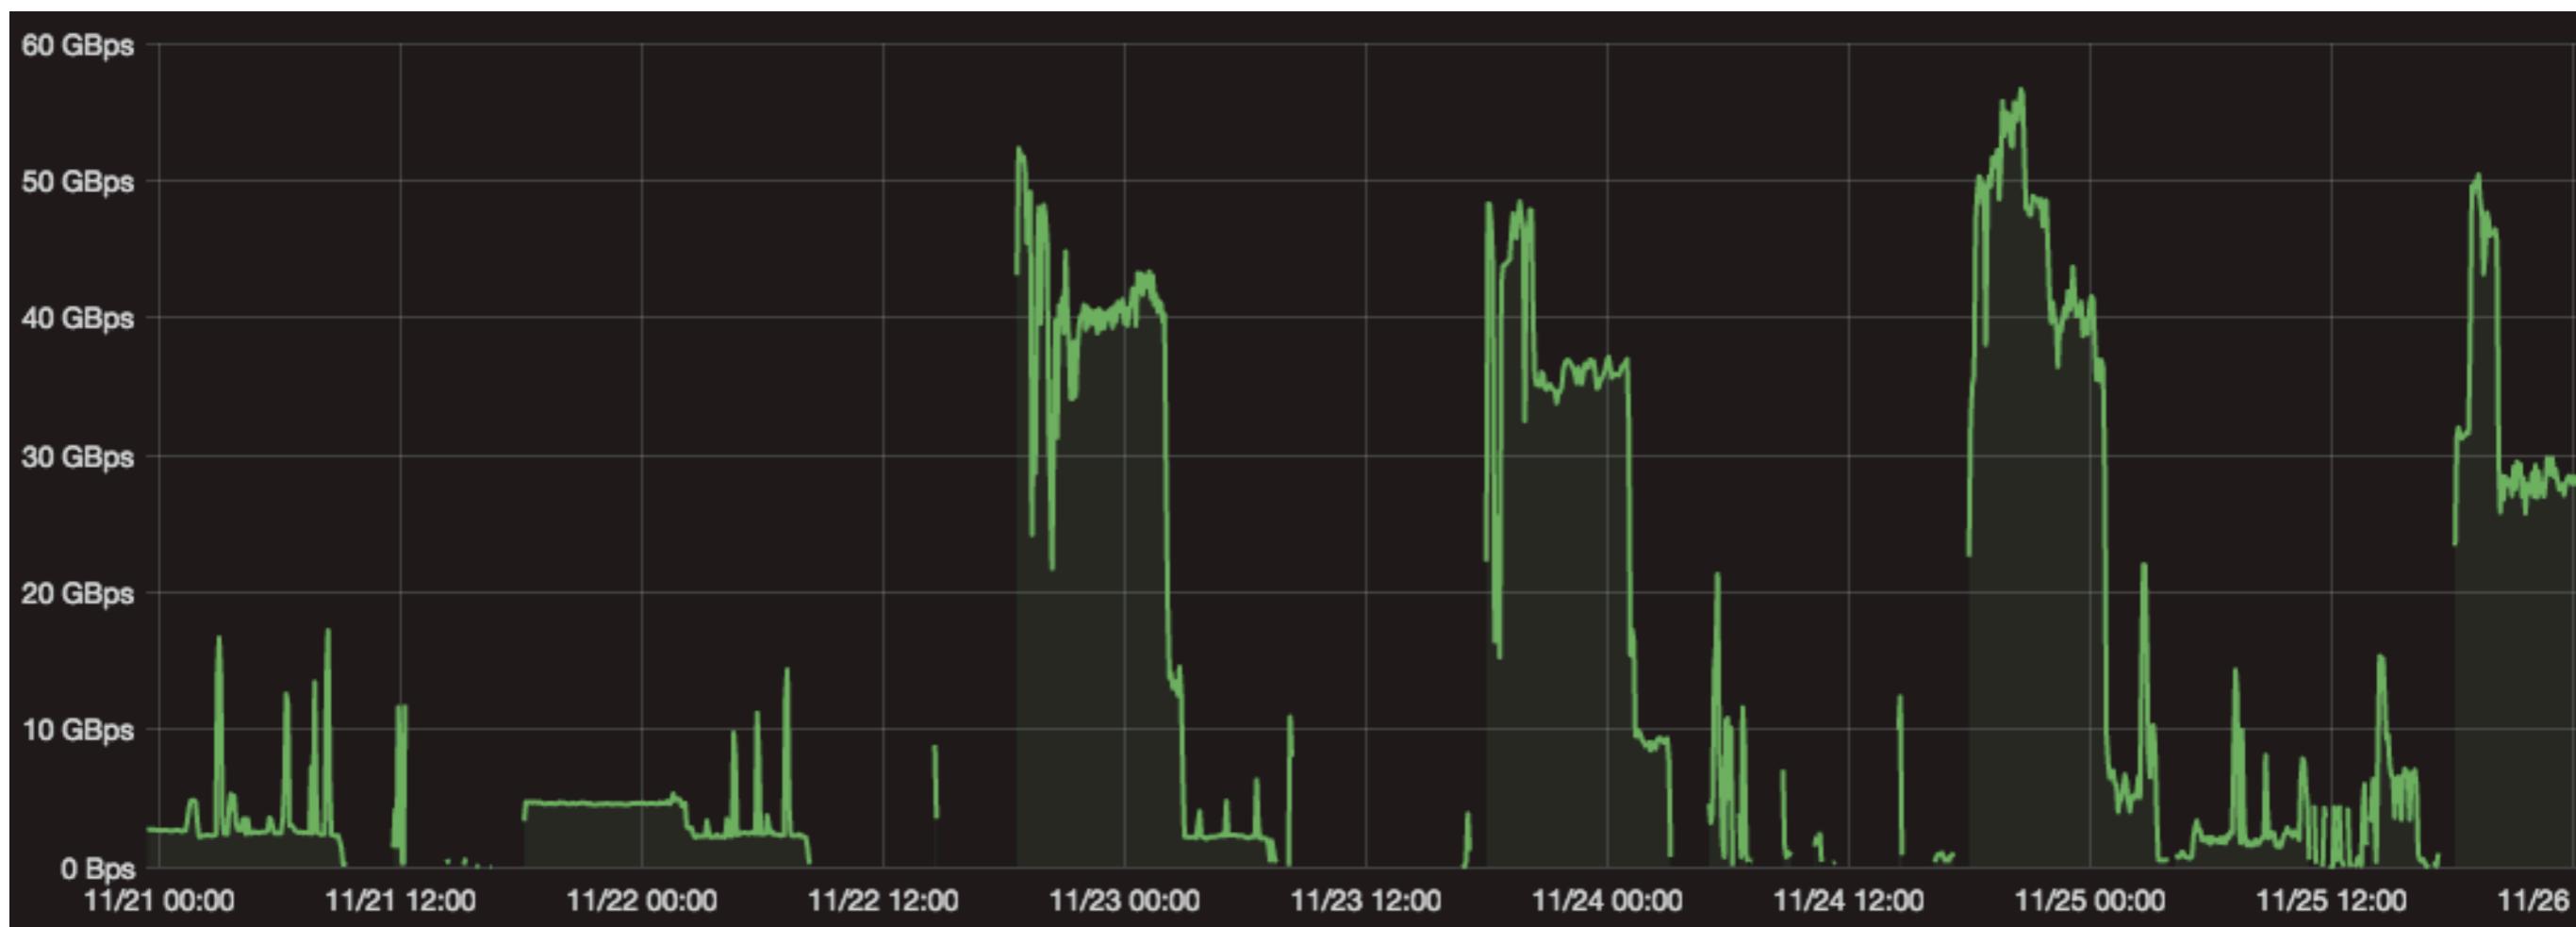


Reverse proxy

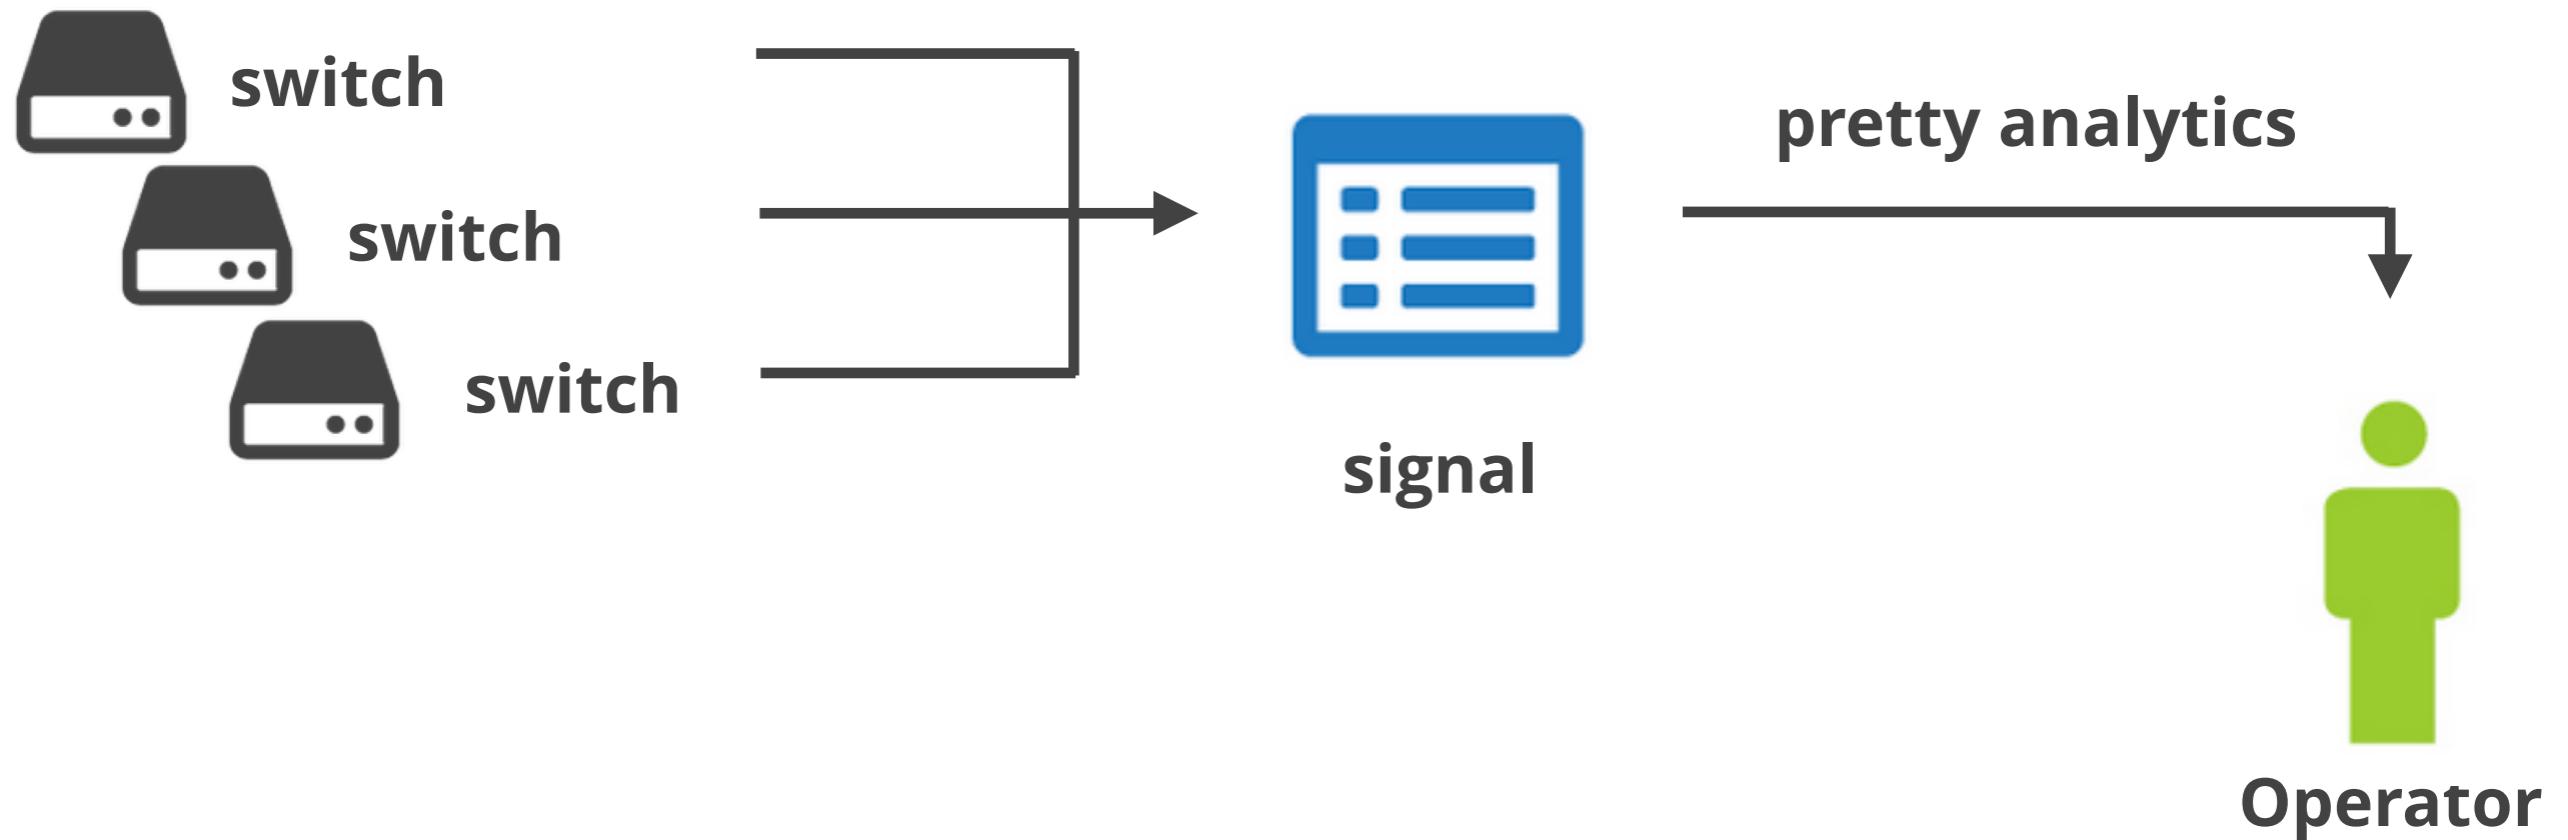


- Optimizations
- Caching
- DDoS protection
- Security

Attacked



Signal



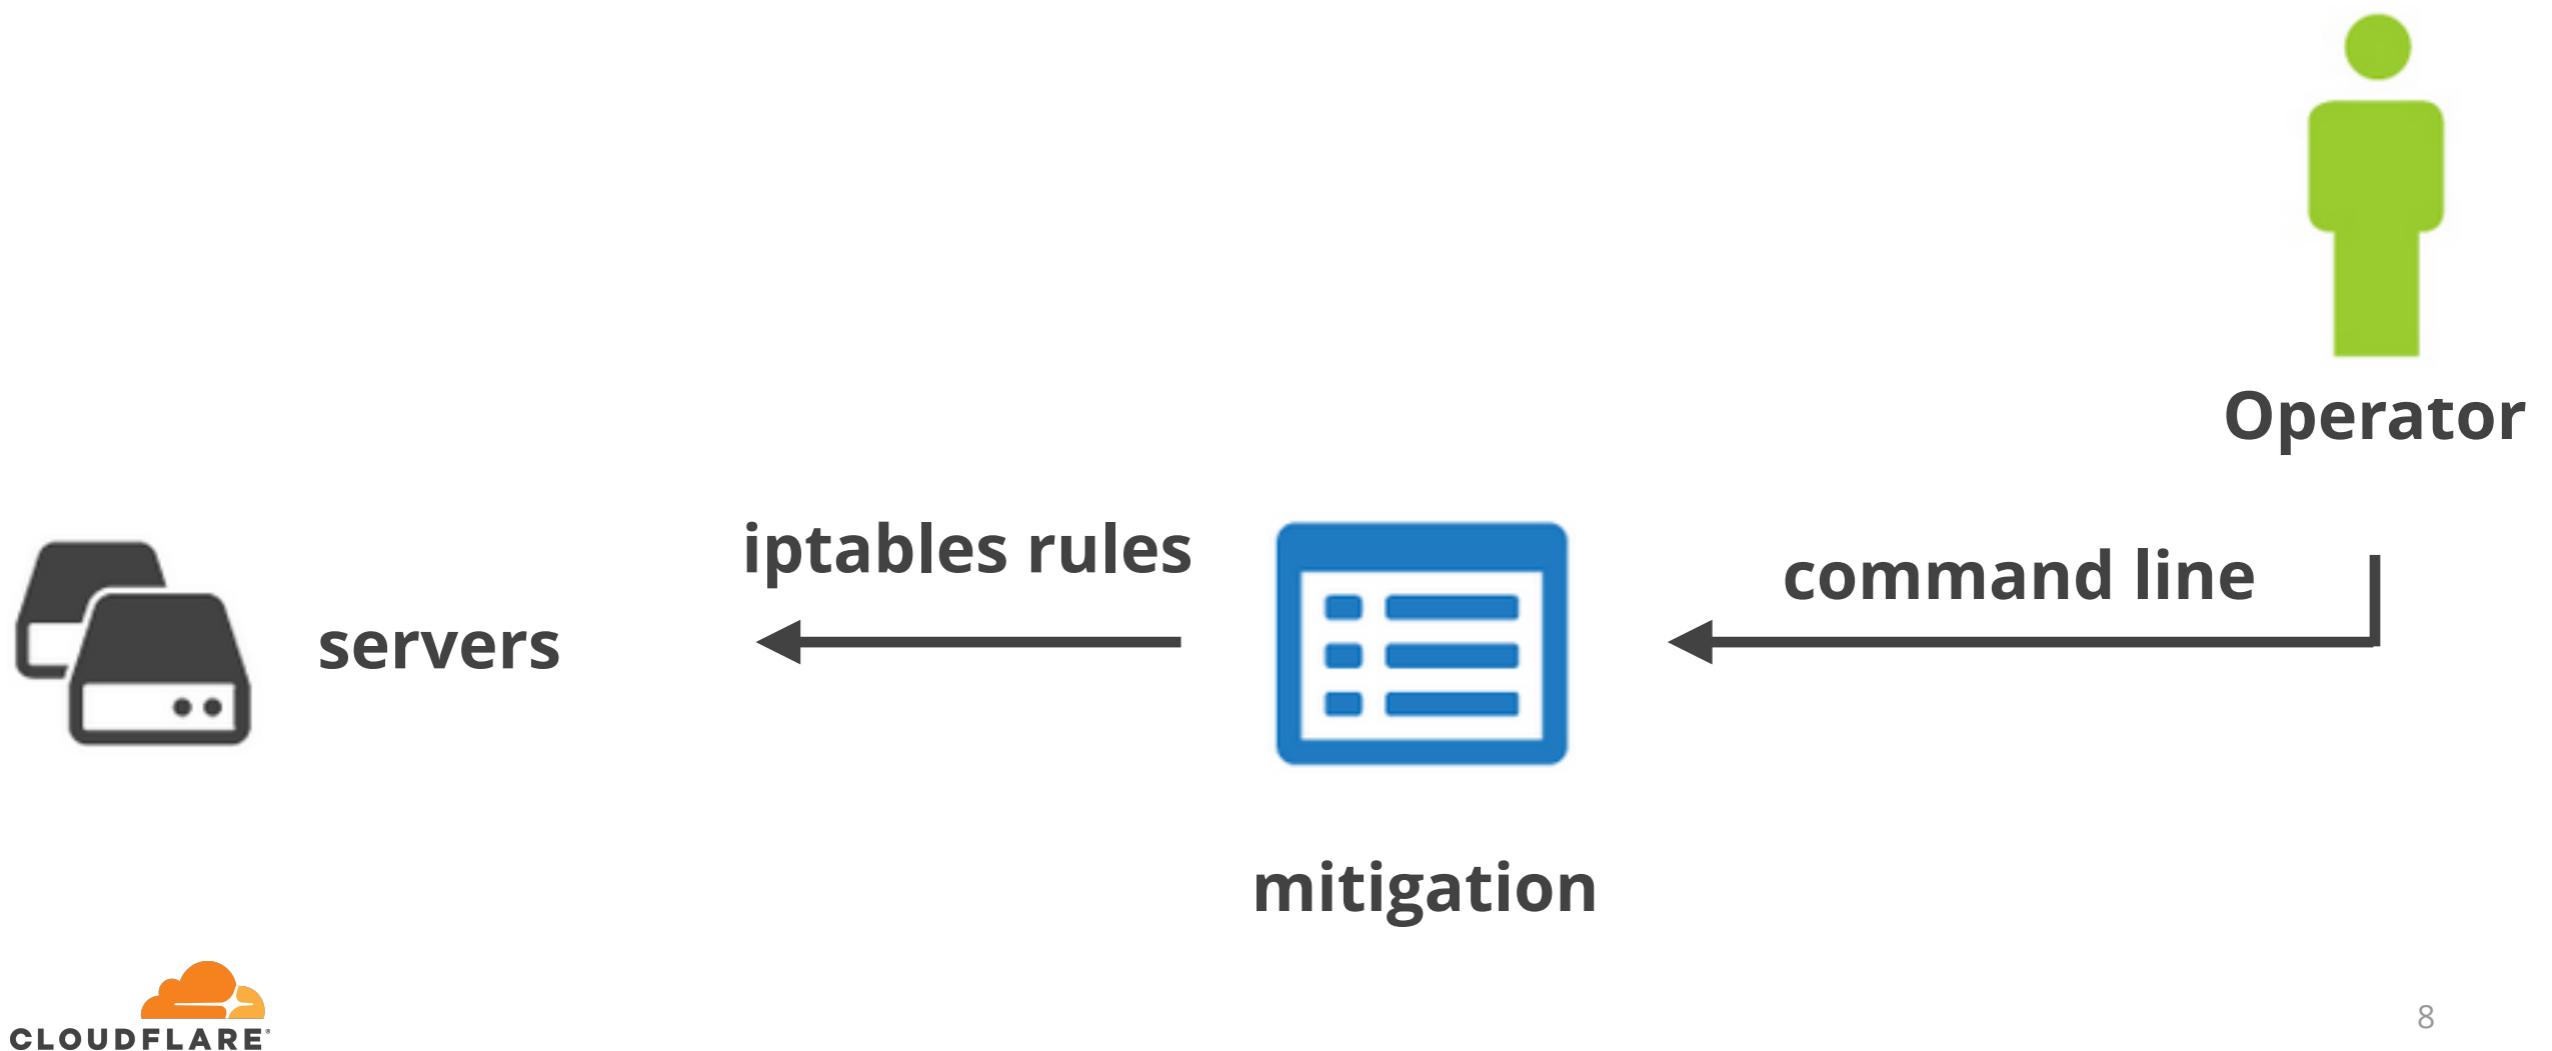
Signal

```
INCOMING/UDP  samples=11836 pps={total= 65.176M, accounted= 65.174M}
```

Mpps	rule	hint	
1.160	dnsbpf	-- --ip=173.245.59.	www.████████.com
3.953	dnsbpf	-- --ip=173.245.59.	www.████████.com
3.892	dnsbpf	-- --ip=173.245.58.	████████.com
4.169	dnsbpf	-- --ip=173.245.58.	████████.com
3.750	dnsbpf	-- --ip=173.245.58.	www.████████.com
4.098	dnsbpf	-- --ip=173.245.59.	████████.com
0.965	dnsbpf	-- --ip=173.245.59.	████████.com
3.910	dnsbpf	-- --ip=173.245.59.	████████.com
0.973	dnsbpf	-- --ip=173.245.59.	www.████████.com
3.811	dnsbpf	-- --ip=173.245.58.	www.████████.com
3.770	dnsbpf	-- --ip=173.245.59.	████████.com
3.879	dnsbpf	-- --ip=173.245.58.	████████.com
3.857	dnsbpf	-- --ip=173.245.58.	████████.com
3.975	dnsbpf	-- --ip=173.245.58.	www.████████.com
0.545	dnsbpf	-- --ip=████████.	
3.406	dnsbpf	-- --ip=████████.	
0.315	dnsbpf	-- --ip=173.245.58.	
0.415	dnsbpf	-- --ip=173.245.59.	*.████████.com

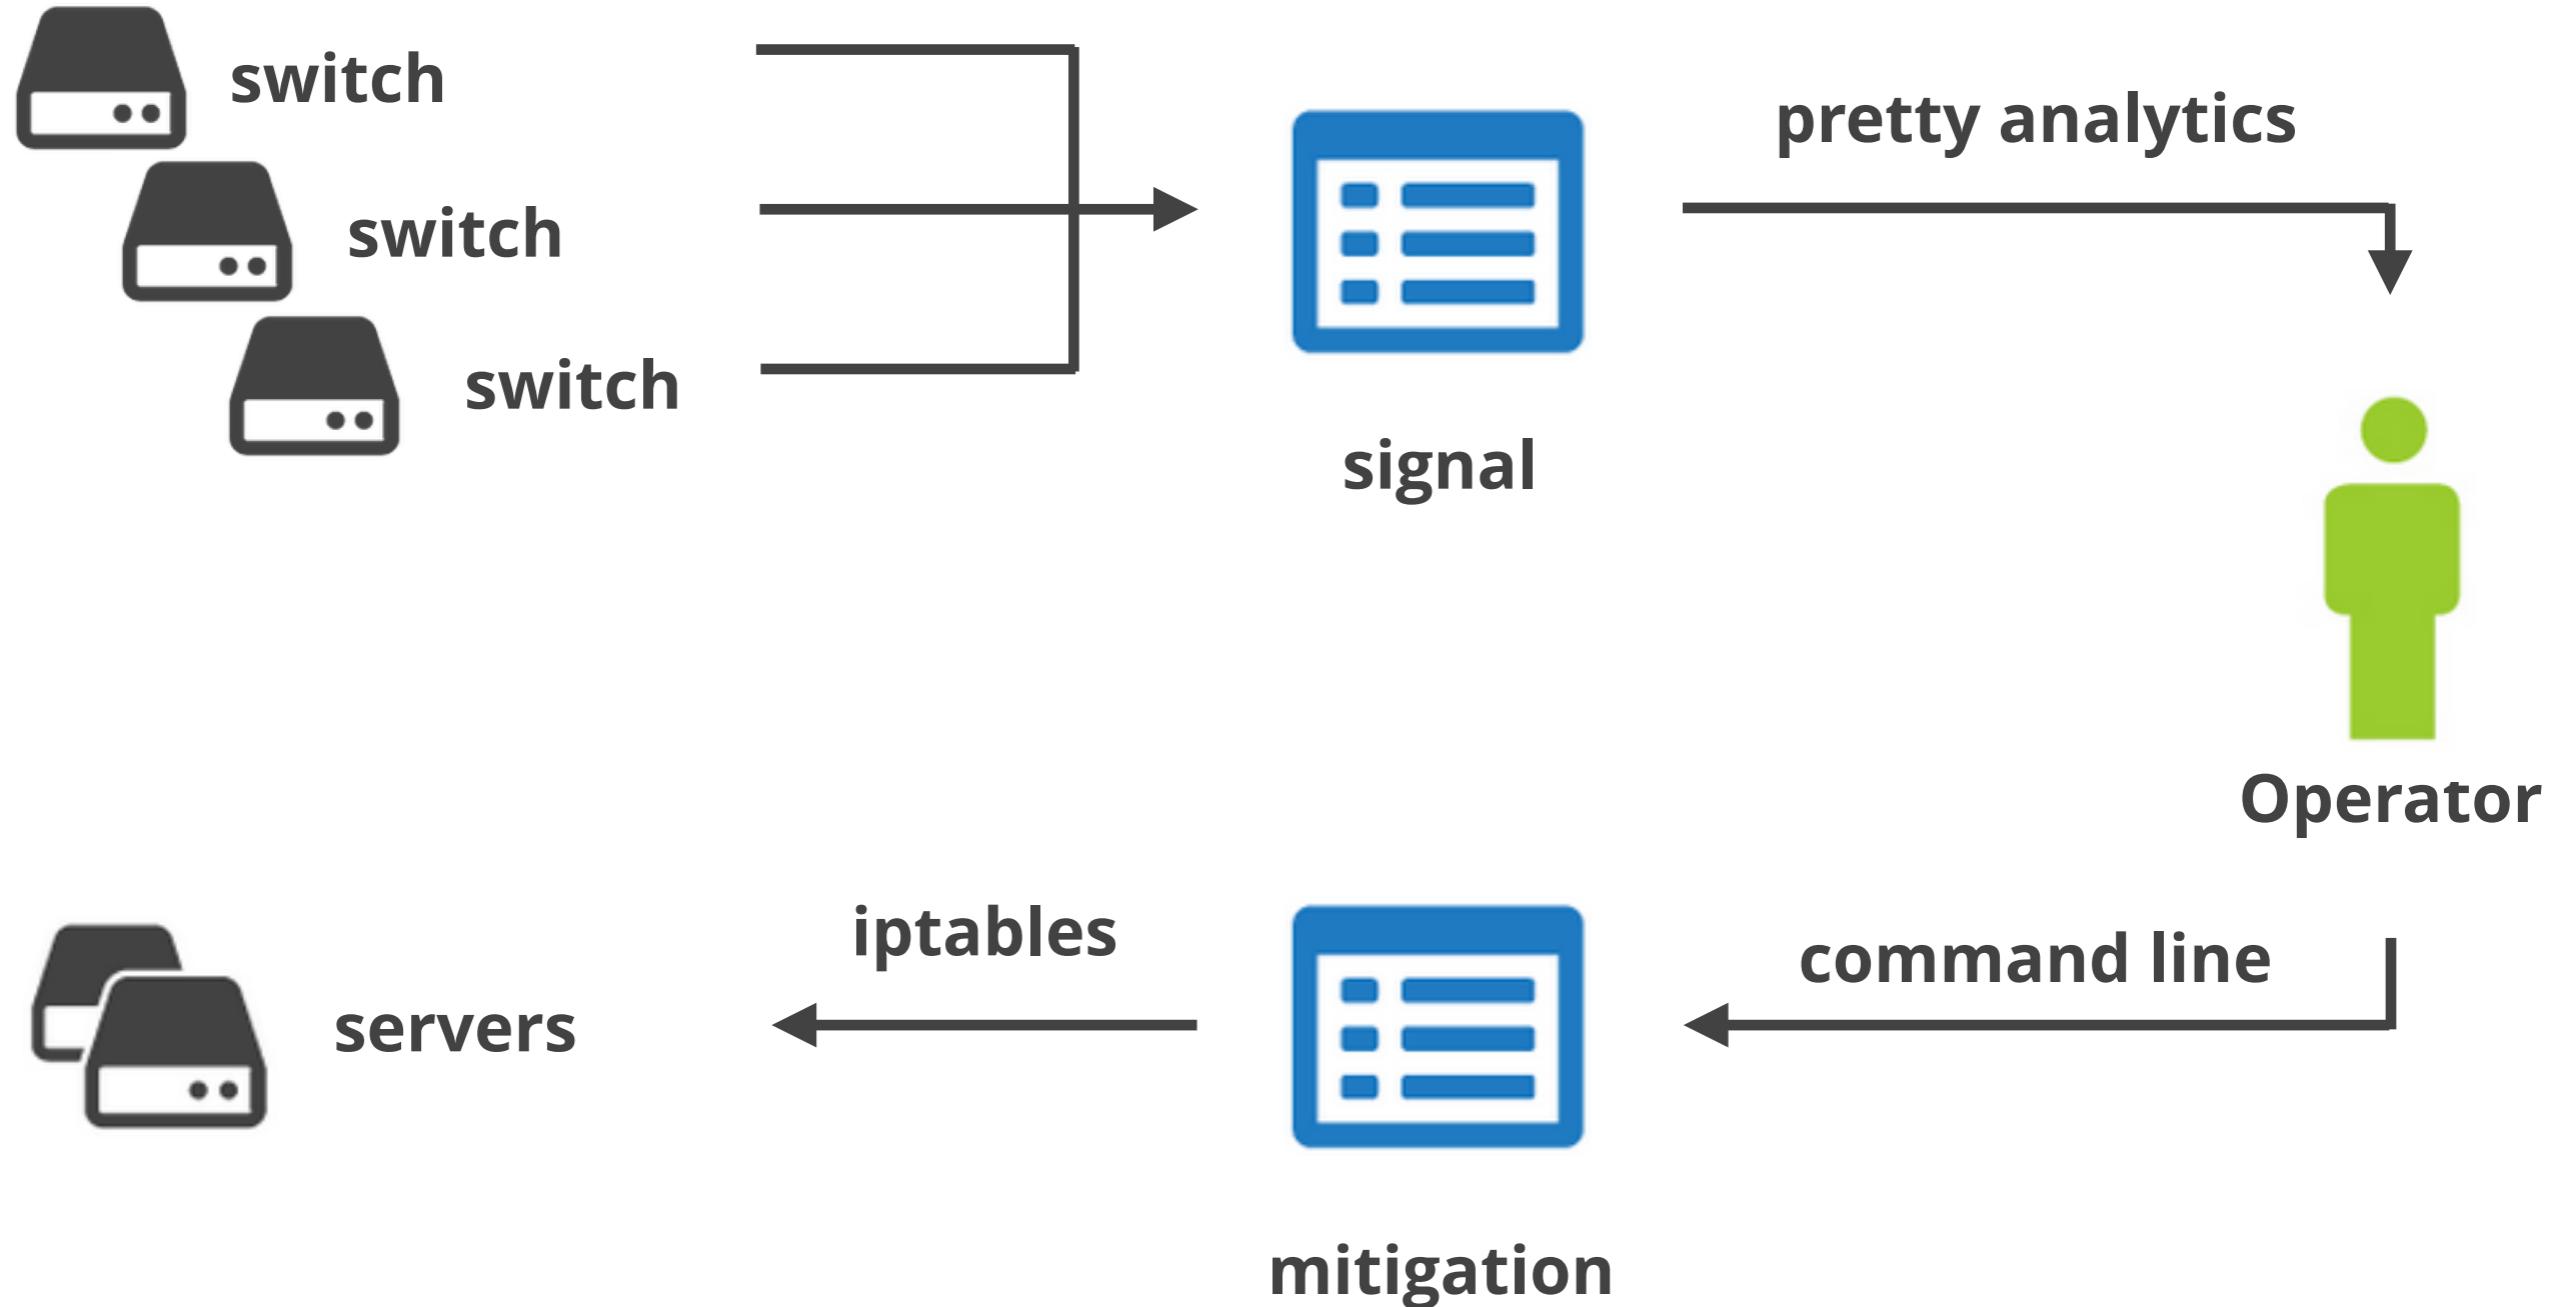
WHAT?

Mitigation



Mitigation

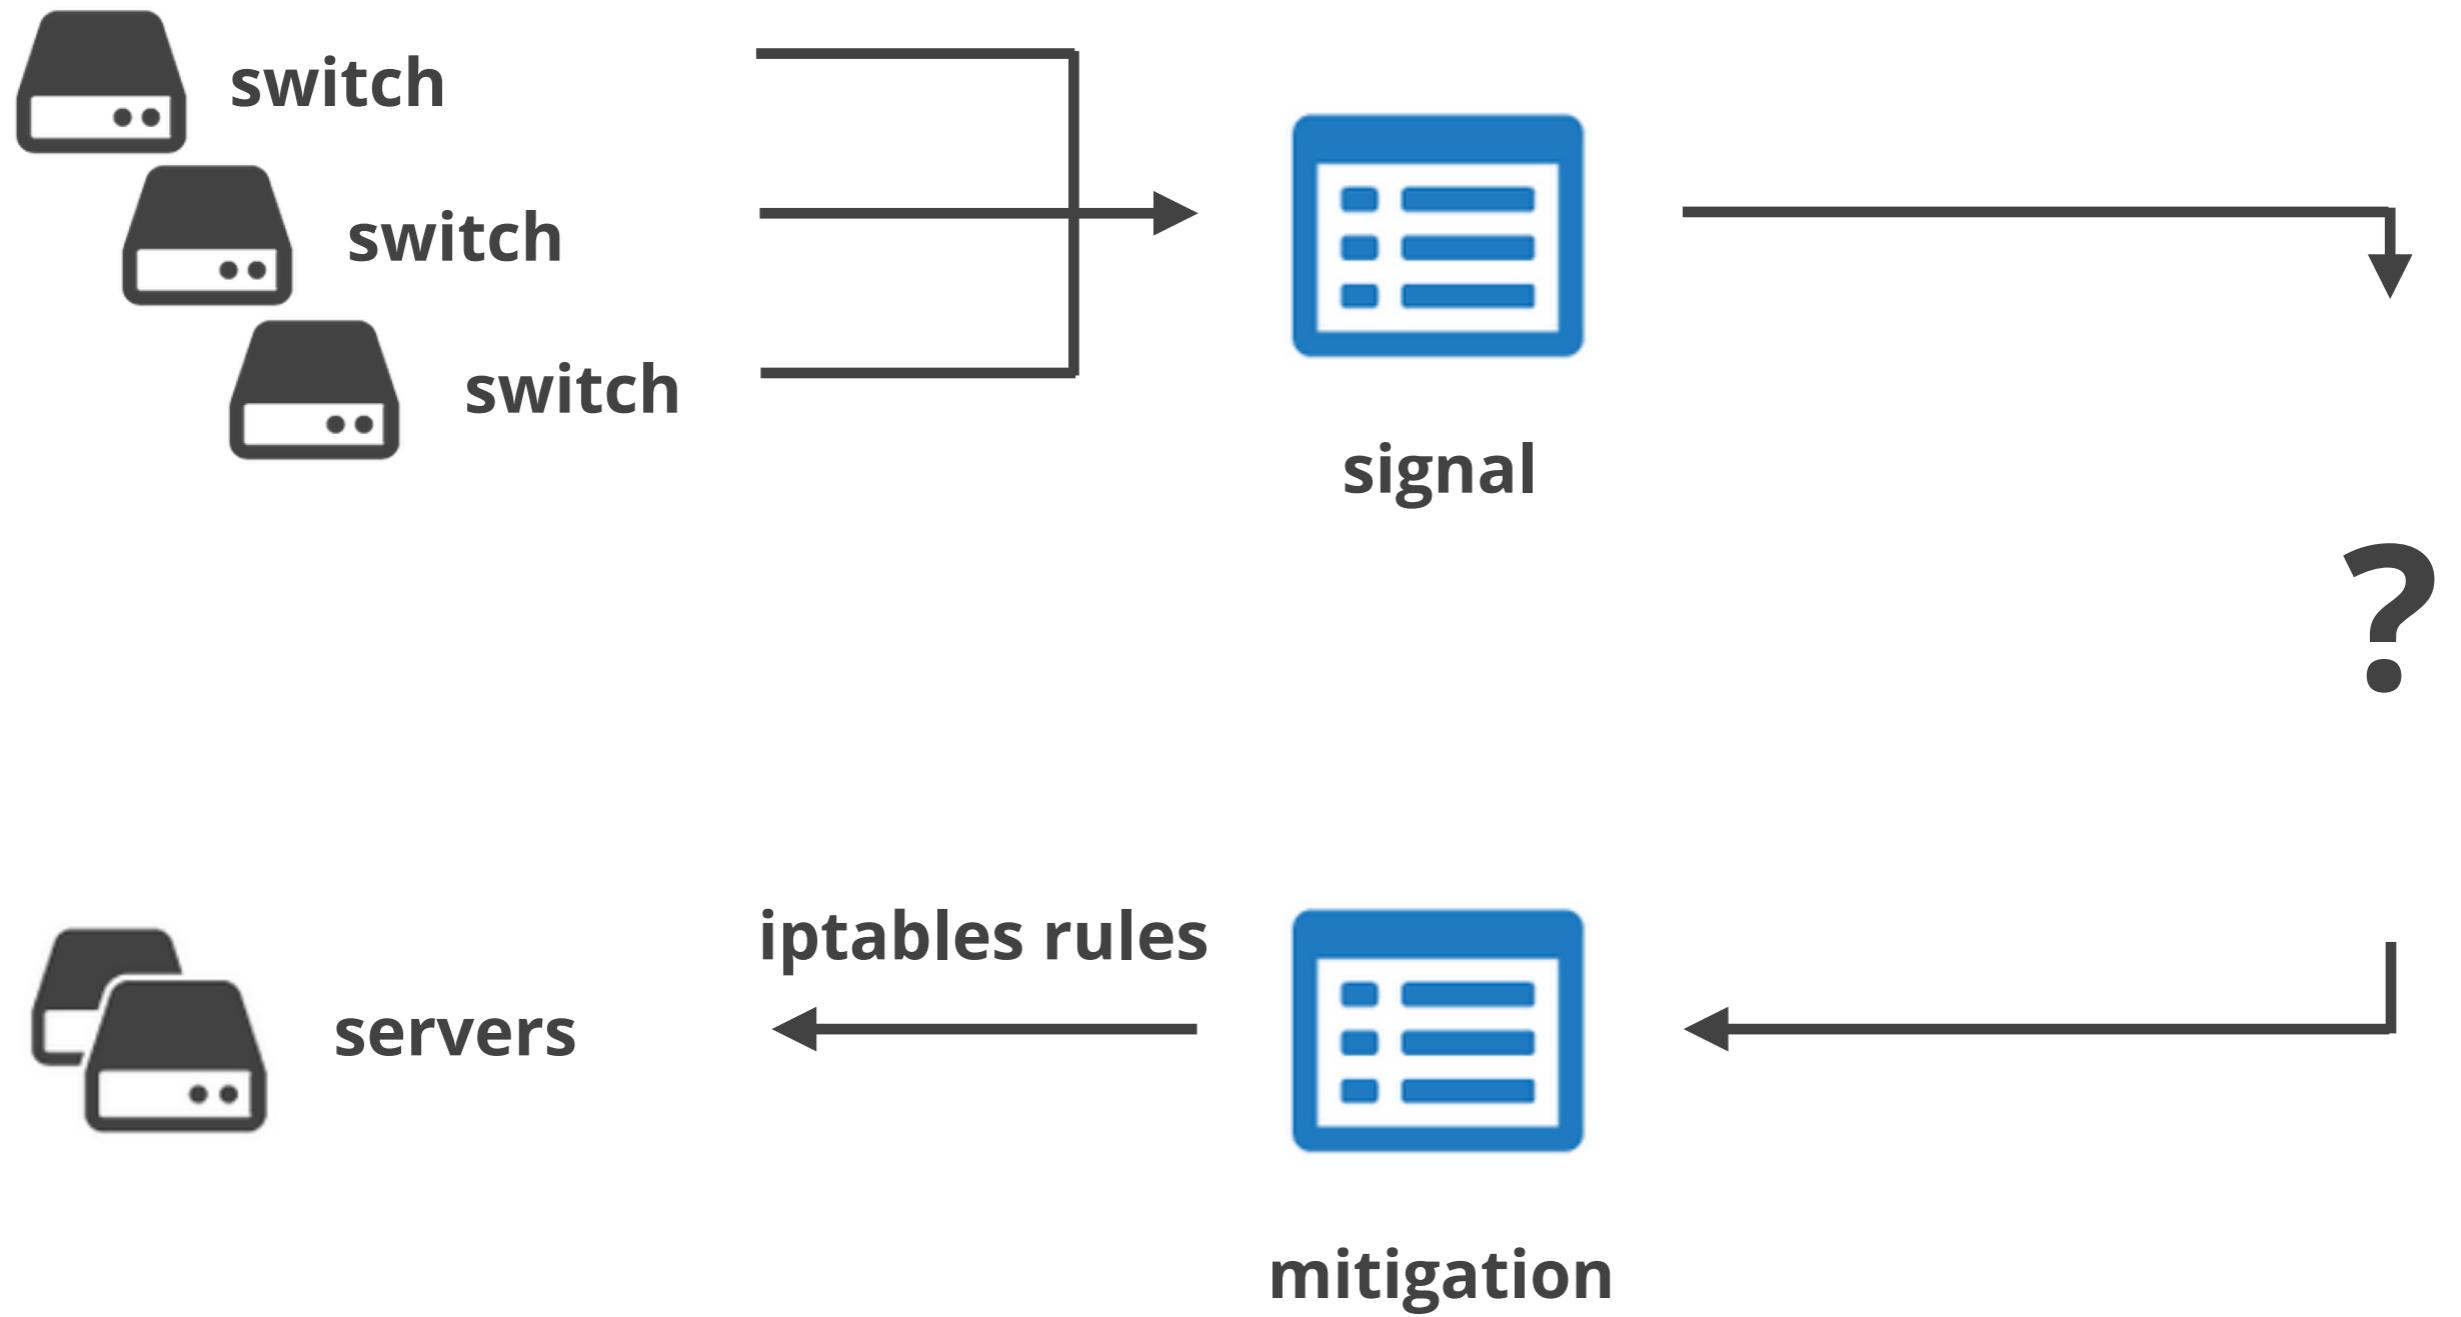
```
$ gatekeeper dnsbpf list
[ ] Imported: 0 removed: 0 trusted: 25
--expiry=365d -- --ip=173.245.██████████ www.████████.com
--expiry=7d -- --ip=██████████ ██████████ m.████████.com
--expiry=7d -- --ip=██████████ ██████████ m.████████.com
--expiry=7d -- --ip=██████████ ██████████ www.████████.com
--expiry=7d -- --ip=██████████ ██████████ www.████████.com
--ip=173.245.██████████ ██████████ www.████████.com
--ip=173.245.██████████ *.*www.████████.com
--ip=173.245.██████████ ██████████ www.████████.com
--ip=██████████ ██████████ ██████████ .com
--ip=██████████ ██████████ ██████████ .com
--ip=██████████ *.*.████████.com *.*.████████.com *.*www.████████.com
--ip=██████████ *.*www.████████.com *.*.████████.com
--ip=██████████ *.*www.████████.com *.*www.████████.com
--ip=██████████ ██████████ ██████████ ,████████,████████
--ip=██████████ *.*.████████.com *.*.████████.com *.*www.████████.com
--ip=██████████ *.*.████████.com *.*.████████.com
```



Copy-pasta



"Business logic"



12 months later...

```
--ip=1.2.3.4 example.com
```



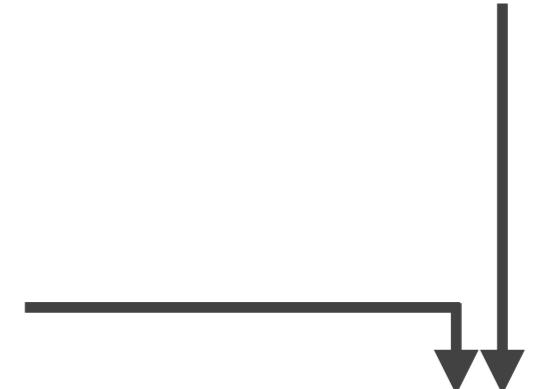
Business logic

```
--ip=1.2.3.4 example.com --qps=100
```



example.com = FREE | PAID

```
--ip=1.2.3.4 example.com
```



Business logic

```
--ip=1.2.3.4 example.com --qps=500
```



example.com subdomains:
(www, ns1, ns2)

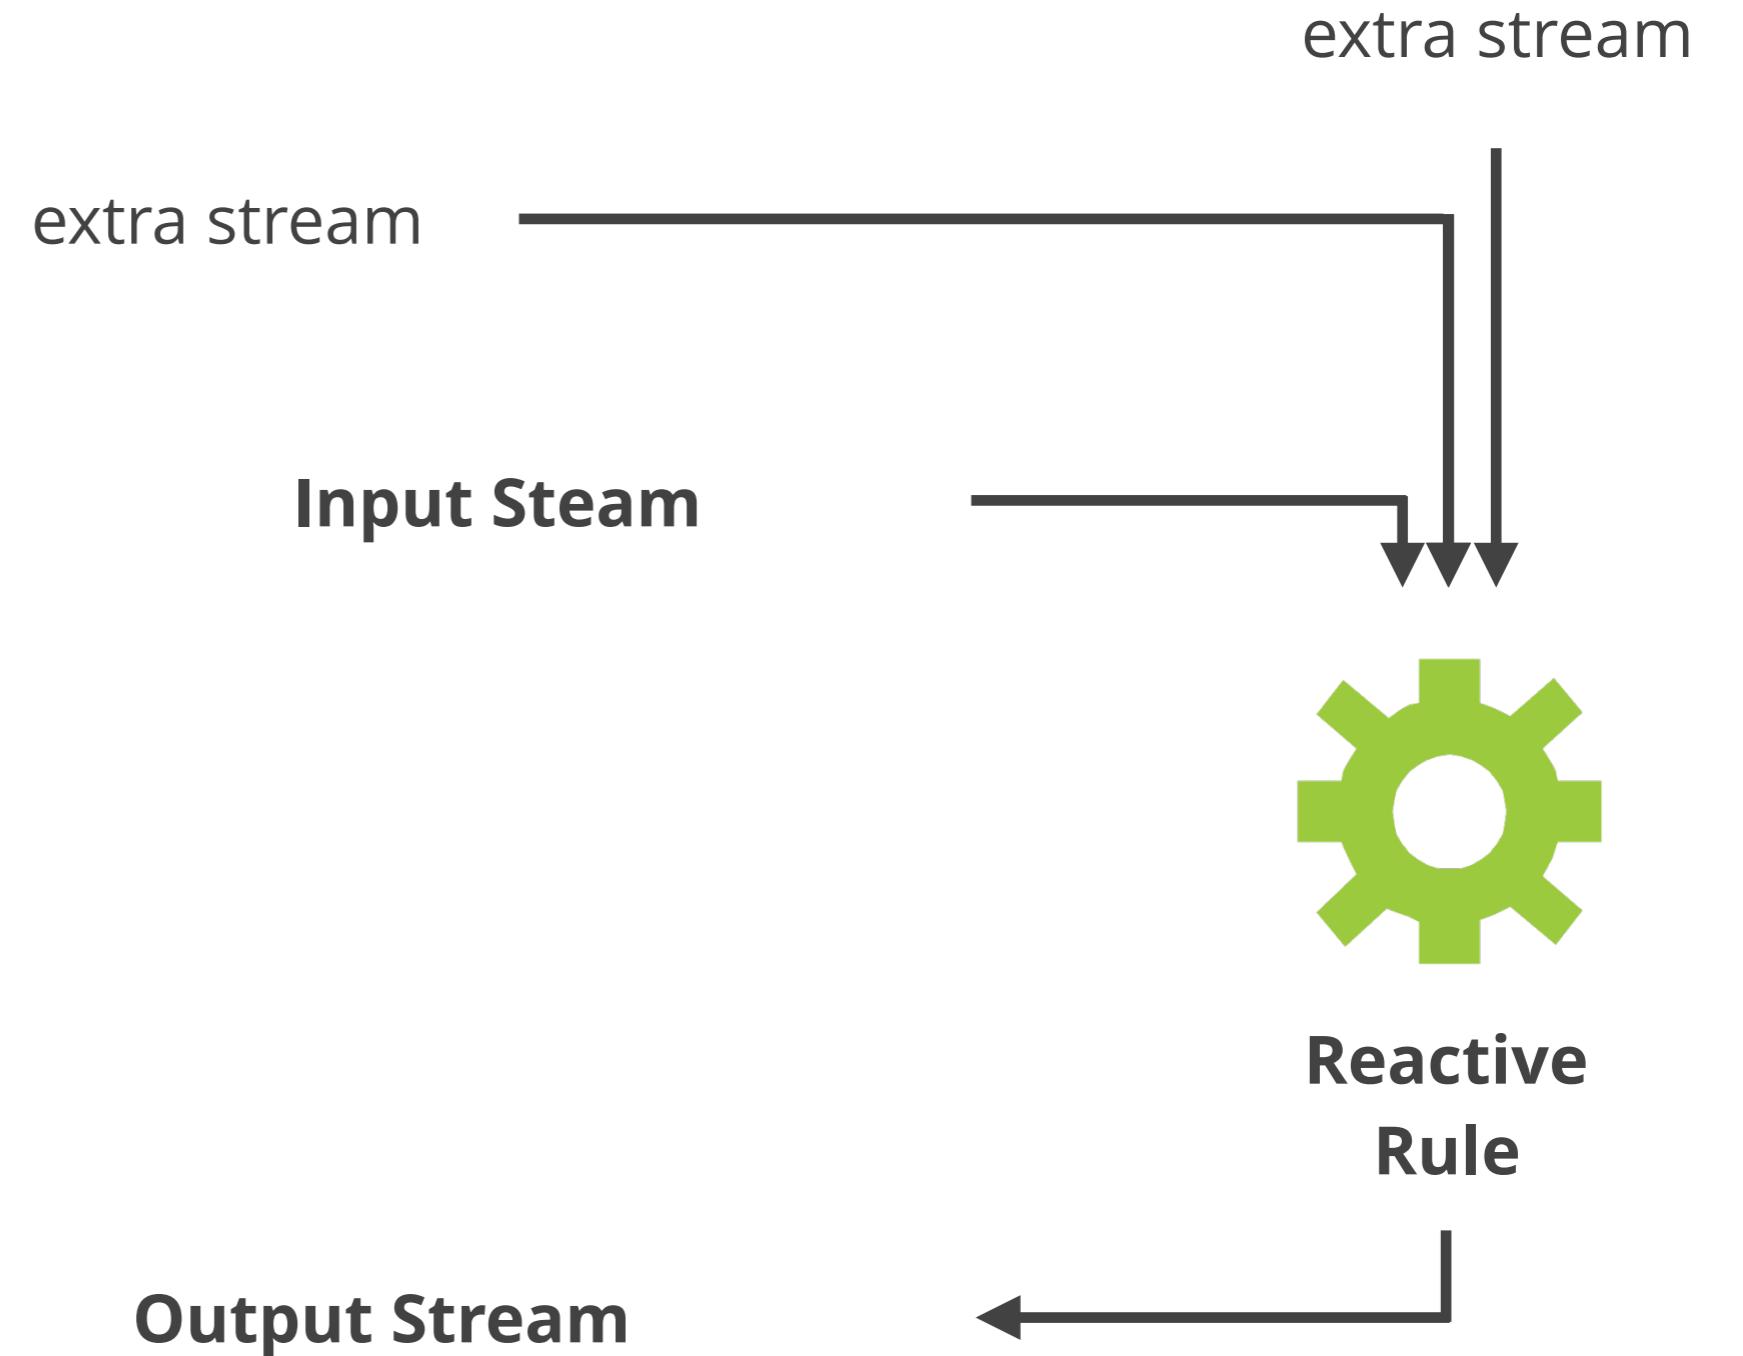
example.com = FREE | PAID

--ip=1.2.3.4 example.com



Business logic

--ip=1.2.3.4 example.com --except www,n1,ns2 --qps=500



Reactive rule

```
def dns_mitigation(attack, plan, subdomains, toggles):
    domain = attack['domain']

    if toggles['all_mitigations_disabled']: return

    qps = 100
    if plan[domain] == 'business':
        qps = 500

    mitigation =
        attack['description'] + \
        ' --qps=%s' % qps + \
        ' --except=%s'.join(subdomains[domain])

    return mitigation
```

Subscriptions

```
def dns_mitigation(attack, plan, subdomains, toggles):
    domain = attack['domain']

    if toggles['all_mitigations_disabled']: return

    qps = 100
    if plan[domain] == 'business':
        qps = 500

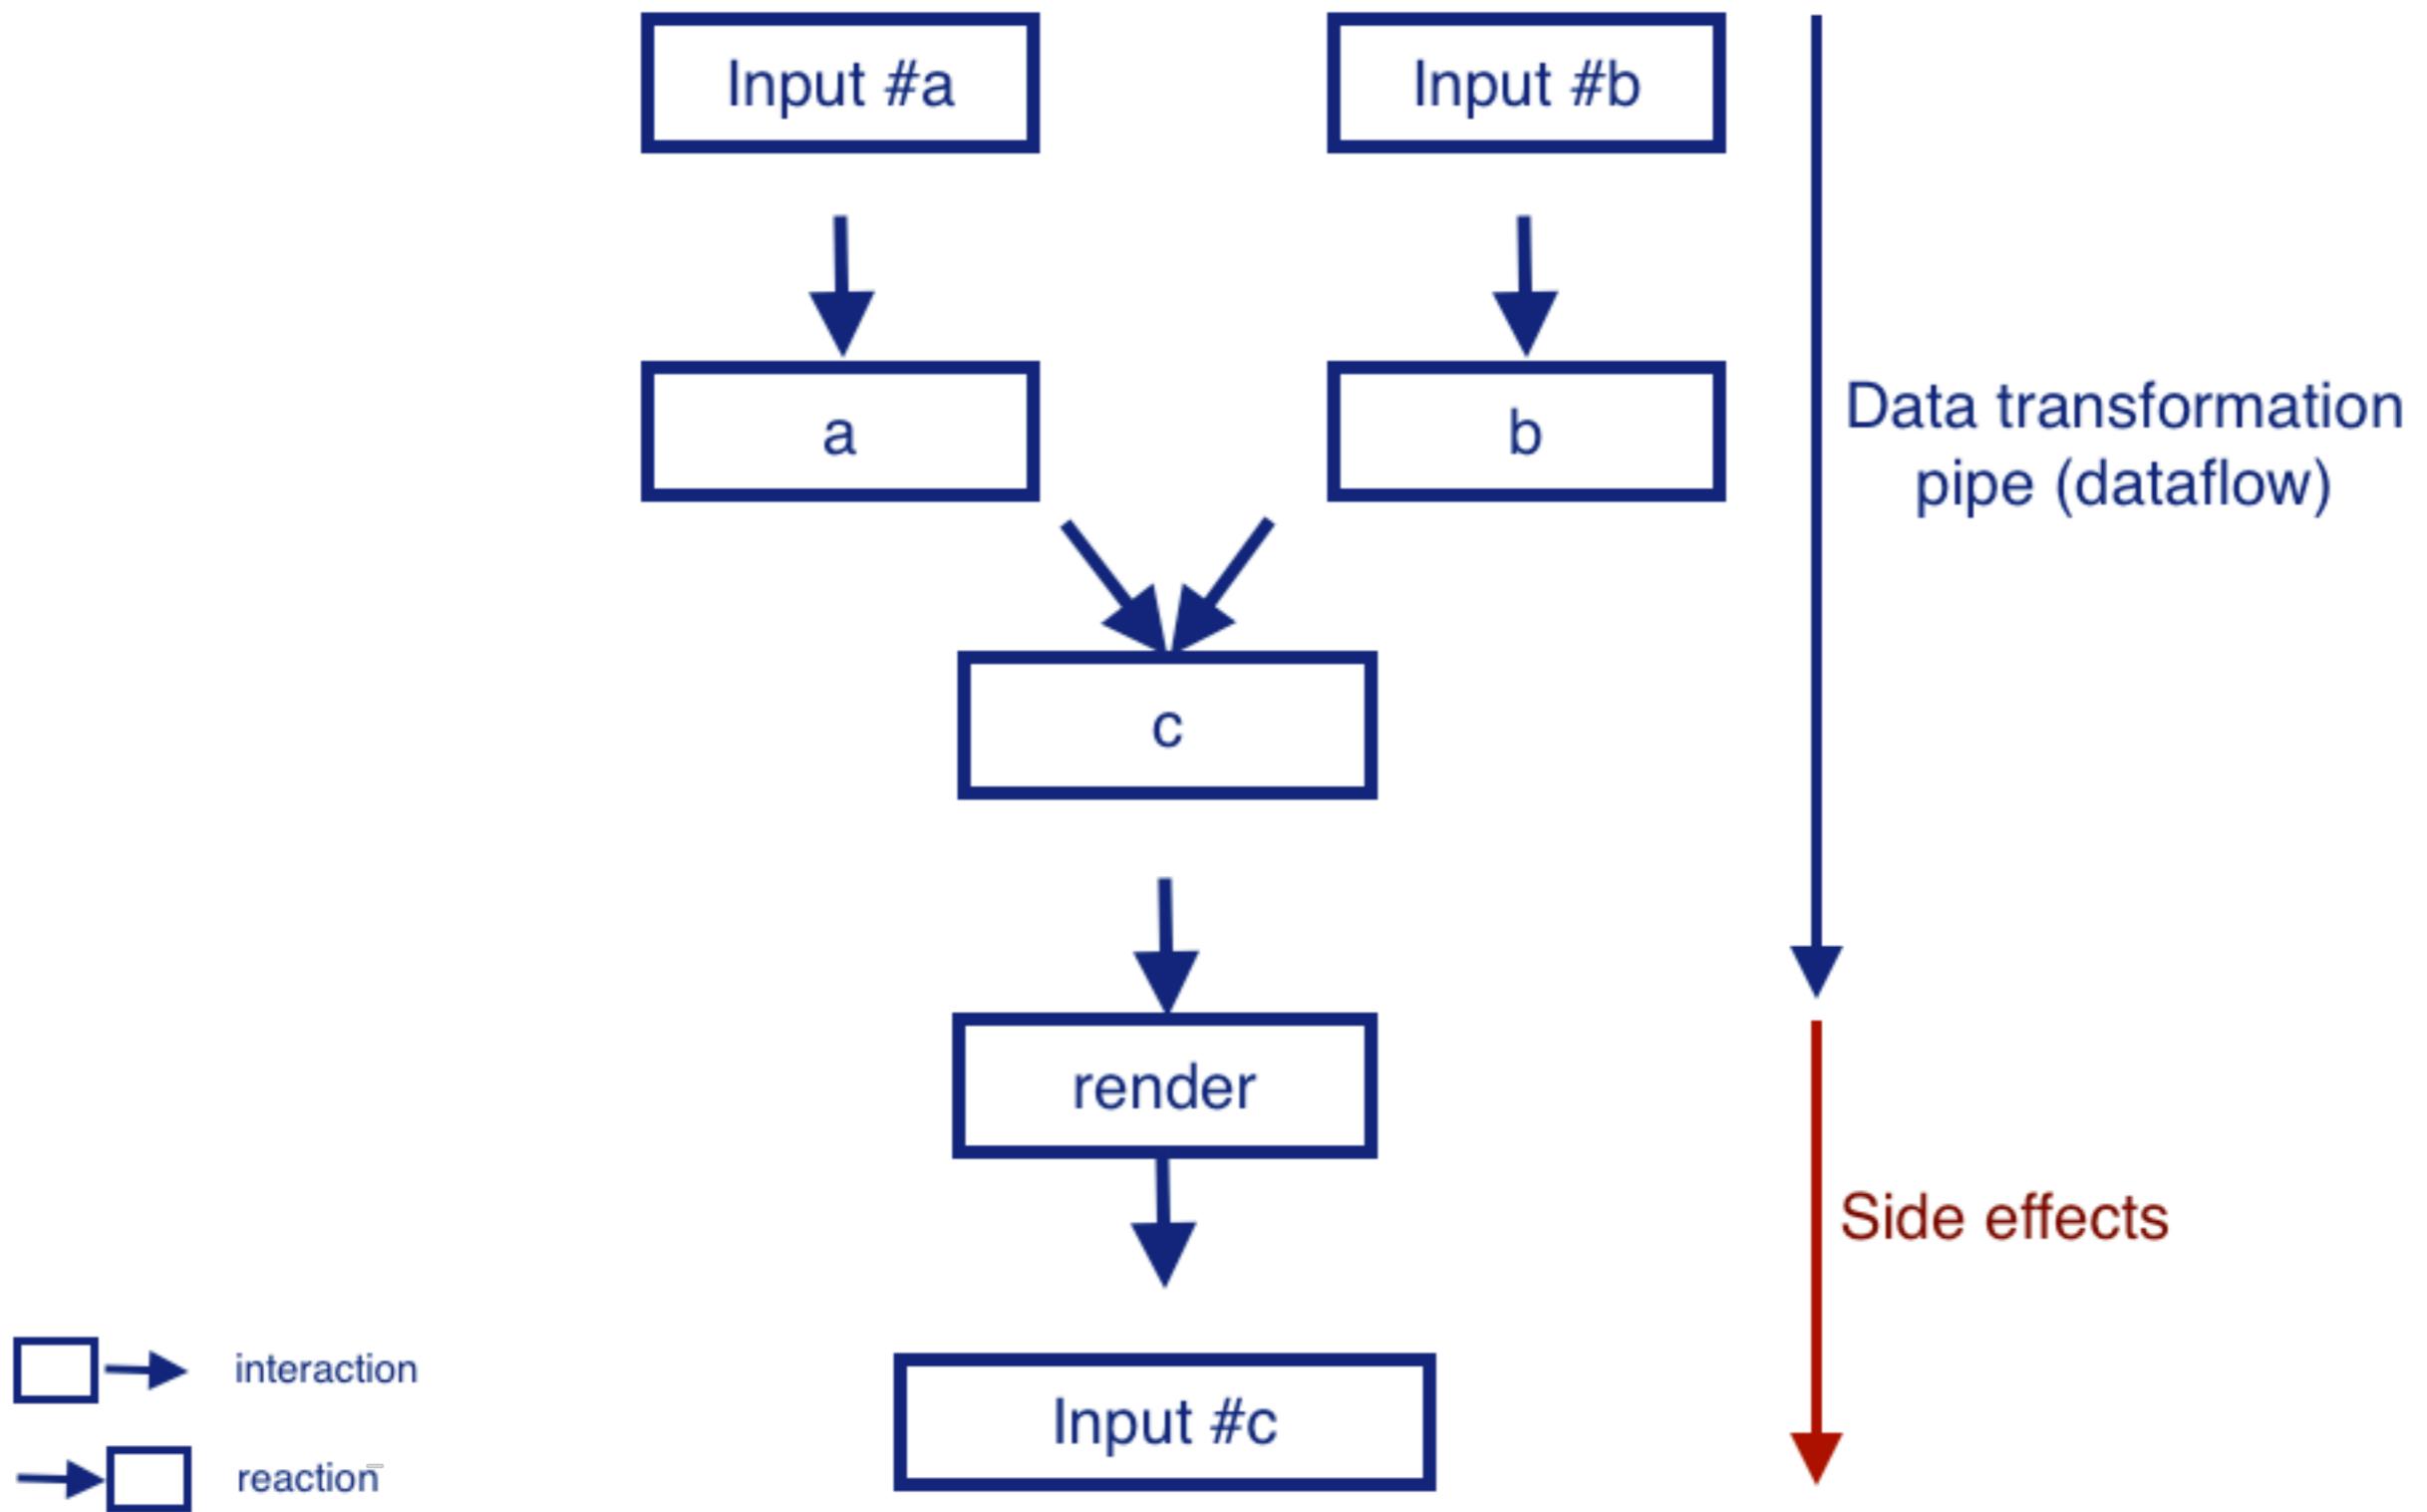
    mitigation =
        attack['description'] + \
        ' --qps=%s' % qps + \
        ' --except=%s'.join(subdomains[domain])

    return mitigation
```

Business logic

- Hard problem!
- Multiple DB lookups
- Wait for operator confirmation
- Critical path

Functional reactive programming

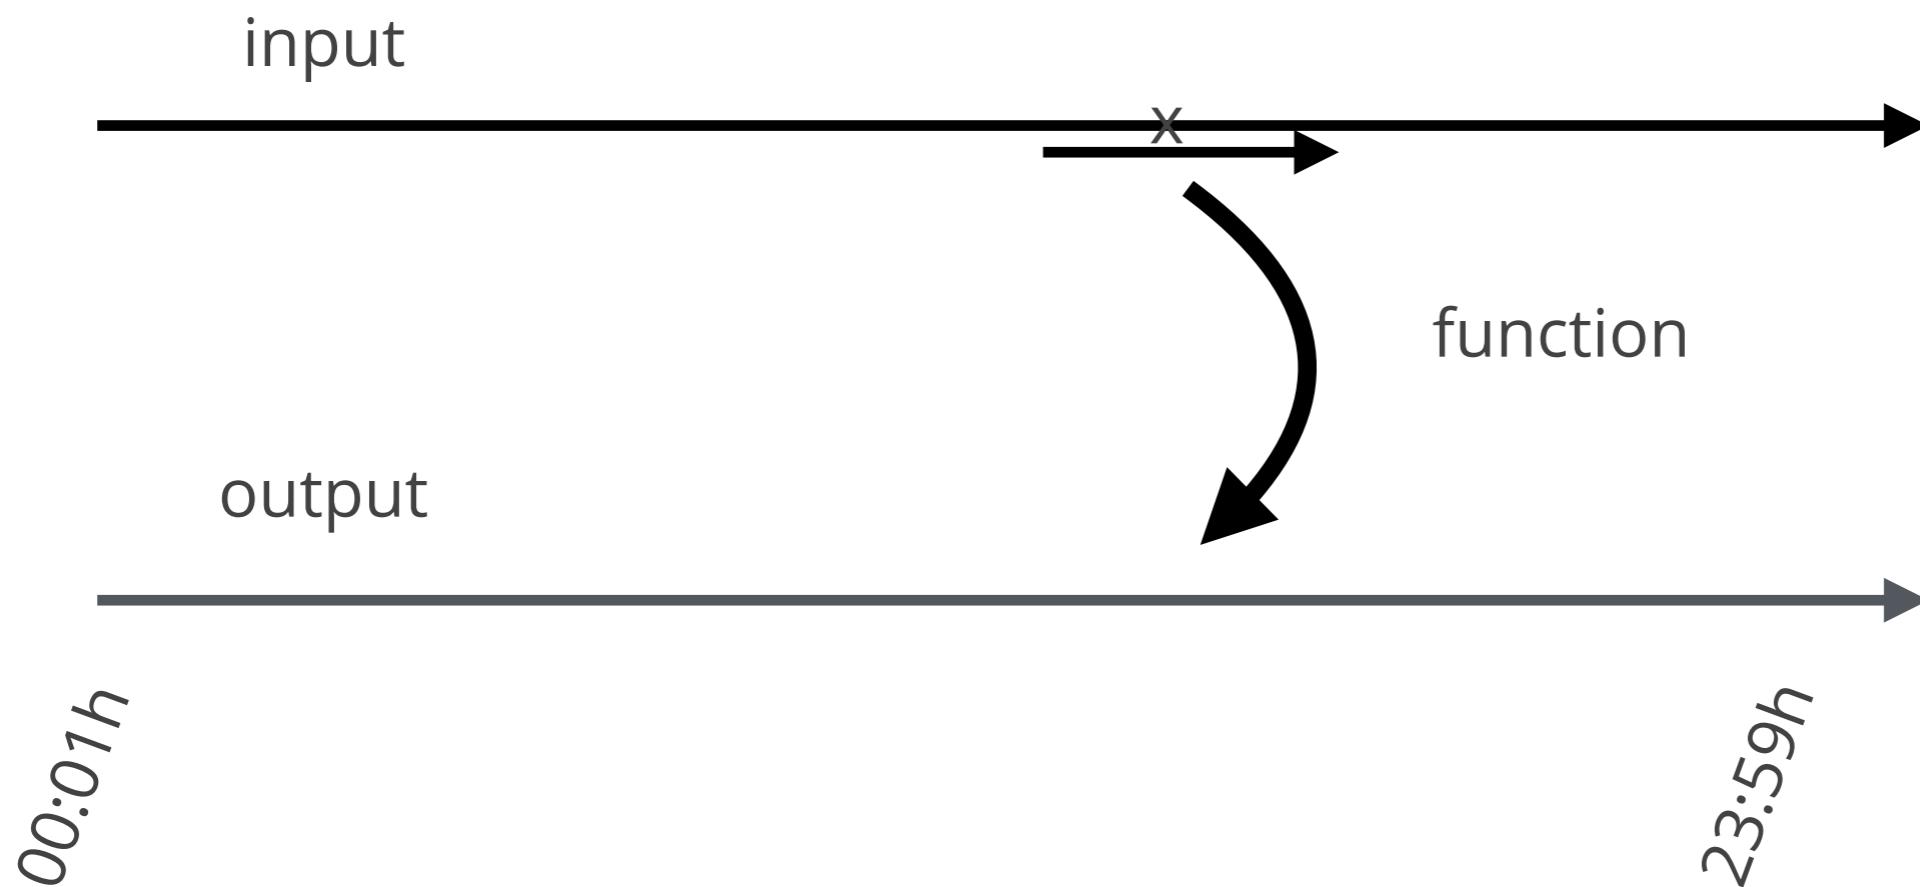


models - Excel

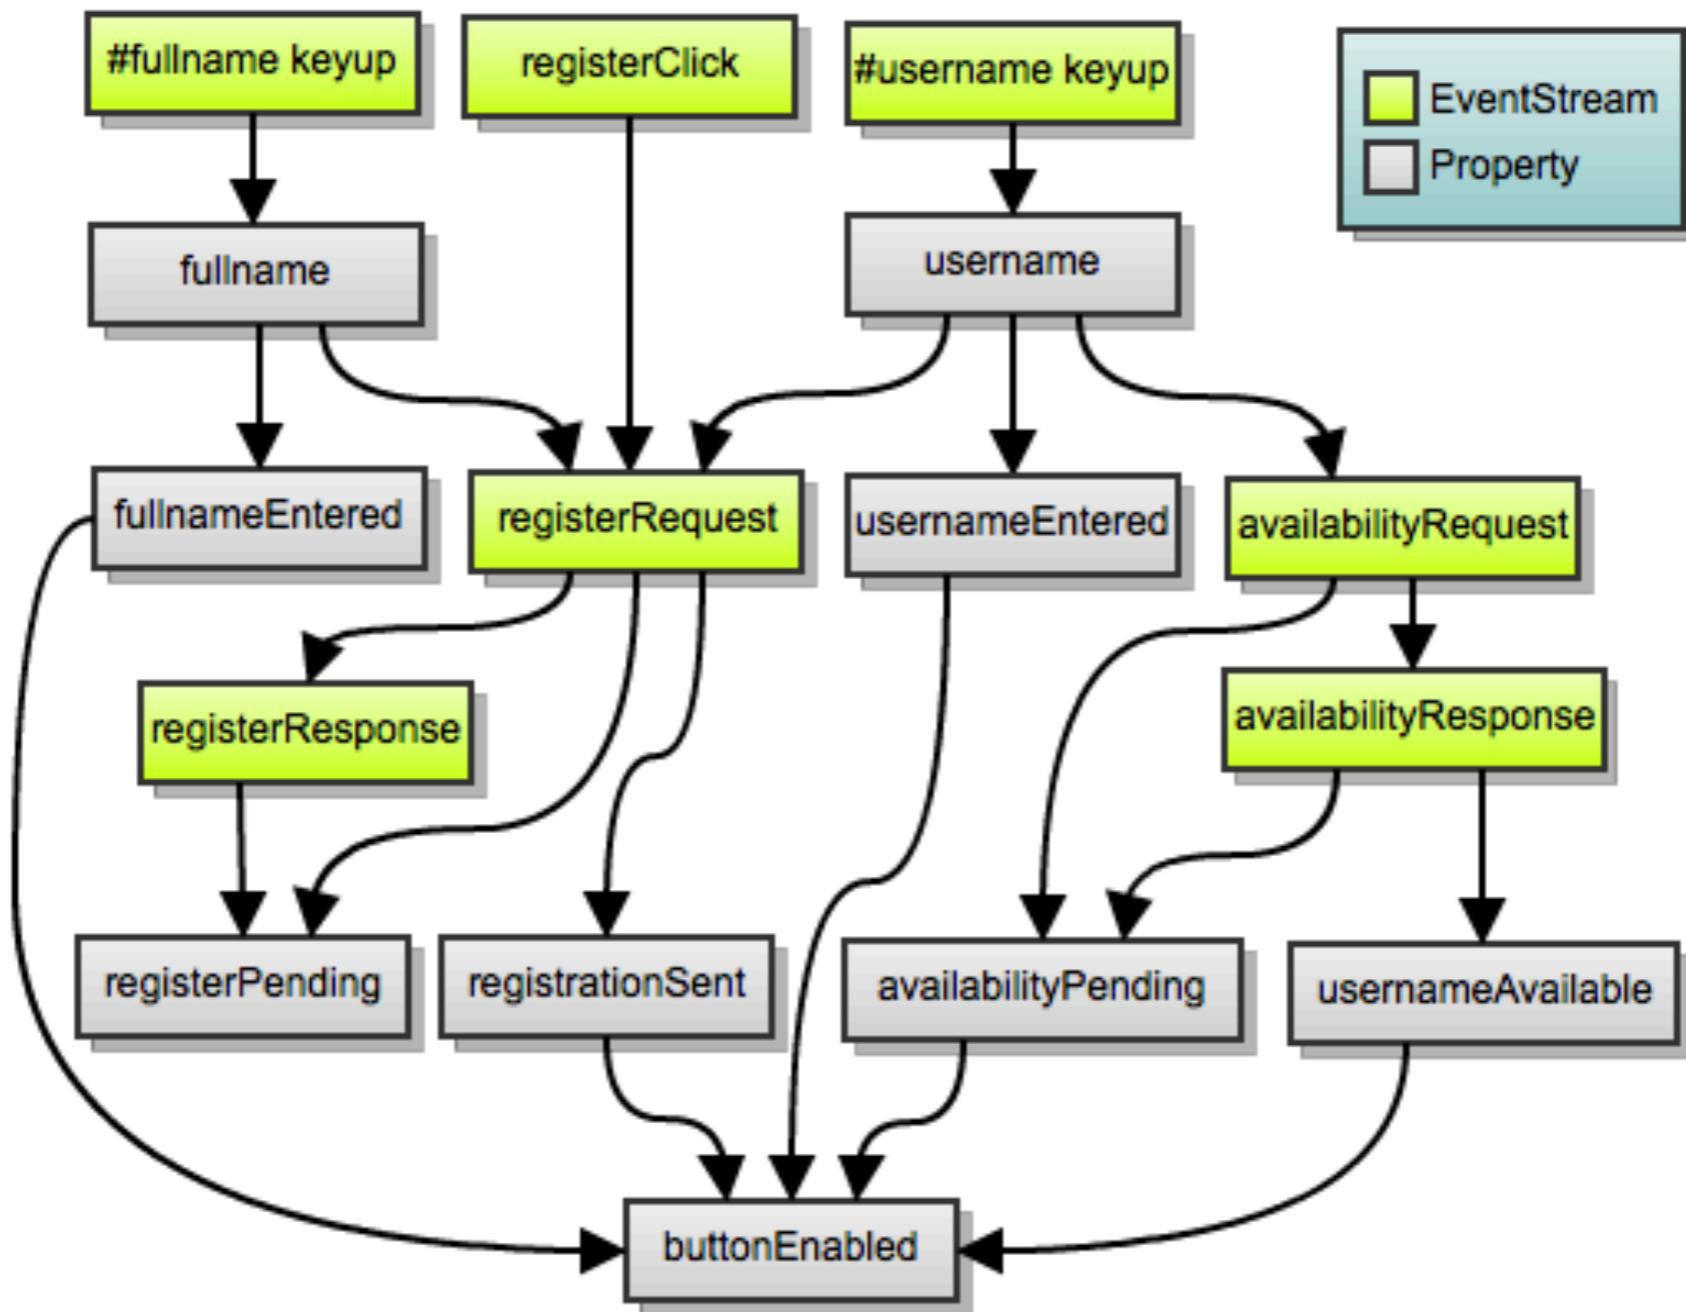
The screenshot shows a Microsoft Excel spreadsheet titled "PowerBI_Test_Data.xlsx - Excel". The window title bar also displays "Mark Kaelin". The ribbon menu includes File, Home, Insert, Page Layout, Formulas, Data, Review, View, and a search bar. The status bar at the bottom indicates "Ready" and shows zoom levels from 100% to 400%.

	A	B	C	D	E	F	G	H	I	J
1	Stock Name	Symbol	Shares	Purchase Price	Cost Basis	Current Price	Market Value	Gain/Loss	Dividend/share	Annual Yield
2	Apple	AAPL	100	\$90.00	\$9,000.00	\$144.13	\$14,413.27	\$14,269.14	\$2.28	1.58%
3	Microsoft	MSFT	200	\$32.00	\$6,400.00	\$65.57	\$13,114.14	\$13,048.57	\$1.56	2.38%
4	Salesforce	CRM	150	\$25.00	\$3,750.00	\$82.57	\$12,385.50	\$12,302.93	\$0.00	0.00%
5	Oracle	ORCL	250	\$50.00	\$12,500.00	\$44.56	\$11,138.75	\$11,094.20	\$0.64	1.44%
6	Hewlett Packard Enterprise	HPE	500	\$18.00	\$9,000.00	\$17.69	\$8,842.50	\$8,824.82	\$0.26	1.47%
7	Alphabet	GOOG	100	\$225.00	\$22,500.00	\$833.36	\$83,336.00	\$82,502.64	\$0.00	0.00%
8	Intel	INTC	200	\$22.00	\$4,400.00	\$36.07	\$7,213.00	\$7,176.94	\$1.09	3.02%
9	Cisco	CSCO	225	\$18.00	\$4,050.00	\$33.24	\$7,478.78	\$7,445.54	\$1.16	3.49%
10	Qualcomm	QCOM	185	\$65.00	\$12,025.00	\$56.48	\$10,447.88	\$10,391.40	\$2.12	3.75%
11	Amazon	AMZN	50	\$800.00	\$40,000.00	\$897.64	\$44,882.00	\$43,984.36	\$0.00	0.00%
12	Redhat	RHT	100	\$95.00	\$9,500.00	\$86.26	\$8,626.00	\$8,539.74	\$0.00	0.00%
13	Facebook	FB	1000	\$17.00	\$17,000.00	\$141.64	\$141,640.00	\$141,498.36	\$0.00	0.00%
14	Twitter	TWTR	500	\$45.00	\$22,500.00	\$14.61	\$7,302.55	\$7,287.94	\$0.00	0.00%
15										

models - Materialized data



models - Signals



Pure FRP is useless

- Weird language - (ELM anyone?)
- Fixed signal flow
- Strictly no side-effects

Dirty FRP is awesome

- Weird language
 - **Python**
- Fixed signal flow
 - **Attacks come and go, but patterns fixed**
- Strictly no side-effects
 - **Dynamic "subscriptions", but idempotent**

Prior art - Trellis

The screenshot shows a web page for the PEAK Python Enterprise Application Kit. The header features a large logo with the word "PEAK" in bold black letters, where the letter "A" contains a stylized mountain peak. Below the logo, the text "PYTHON ENTERPRISE APPLICATION KIT" is visible. To the right of the logo, the word "Trellis" is written in red. A navigation bar below the header includes links for "The PEAK Developers' Center" (which is highlighted in red), "FrontPage", "RecentChanges", "TitleIndex", "WordIndex", "SiteNavigation", and "HelpContents". To the right of the navigation bar are icons for User Preferences, XML, email, and search. The main content area has a yellow background and displays the text: "Event-Driven Programming The Easy Way, with `peak.events.trellis`". Below this, a note in parentheses states: "(NOTE: As of 0.7a1, many new features have been added to the Trellis API, and some old ones have been deprecated. If you are upgrading from an older version, please see the [porting guide](#) for details.)". Further down, there is a statement: "Whether it's an application server or a desktop application, any sufficiently complex system is event-driven -- and that usually means callbacks." At the bottom, another note says: "Unfortunately, explicit callback management is to event-driven programming what explicit memory management is to most other kinds of programming: a tedious hassle and a significant source of unnecessary bugs."

Software Transactional Memory (STM) And Observers

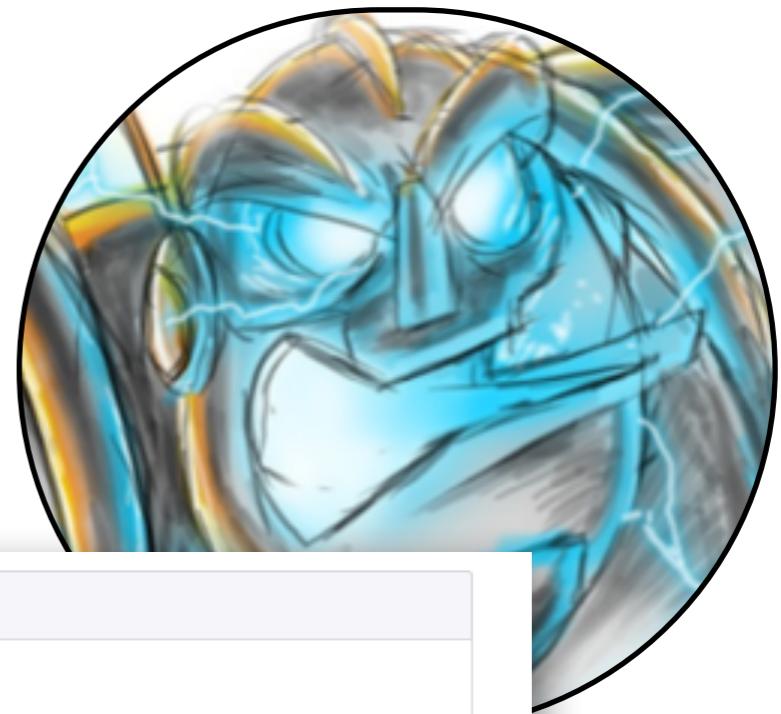
The Trellis is built on a simple Python STM (Software Transactional Memory) and "Observer Pattern" implementation. This document specifies how that implementation works, and tests it.

You should read this document if you plan to implement your own Trellis cell types or other custom data structures, or if you just want to know how things work "under the hood".

Table of Contents

- [STM History](#)
 - [Commit/Abort Notices](#)
 - [Undo, Rollback and Save Points](#)
 - [Commit Callbacks](#)
 - [Logged Setattr](#)
- [The Observer Framework](#)
 - [Subjects, Listeners, and Links](#)
 - [Subjects](#)
 - [Listeners](#)
 - [Controllers](#)
 - [The Singleton Controller](#)
 - [Creating Custom Cell Types \(TBD\)](#)

Gatelogic!



README.md

Gatelogic - somewhat reactive programming framework

Gatelogic is a functional reactive programming / software transactional memory inspired framework.

In classic FRP programming model, the signal graph is defined on the start and doesn't change during the lifetime of a program. For our usage we needed something more dynamic. We also wanted to do express the computable code in an easy to understand Python, as opposed to a classic composition of pure functions.

The idea is to write pretty simple, side-effect free functions that describe some business logic, while allowing them to dynamically "subscribe" arbitrary signal sources.

For example, a basic "action" code could look like:

```
def action(row):
    if confirm_hub.get('all_mitigations_enabled').value != 'True':
        return None
```

This will actually "subscribe" to `confirm_hub/all_mitigations_enabled` input signal. Note, that this subscription is *declared* in the code, as opposed to predefined as input. Change of that input signal value will trigger recomputation of all dependent action functions.

<https://github.com/cloudflare/gatelogic>

Gatelogic

- Input - ReadableHub
 - update(full_data)
- Processing - ComputableHub
 - maintain(key, function)
 - unmaintain
- Subscriptions - QueryHub
 - update(full_data)

Input data - a dict

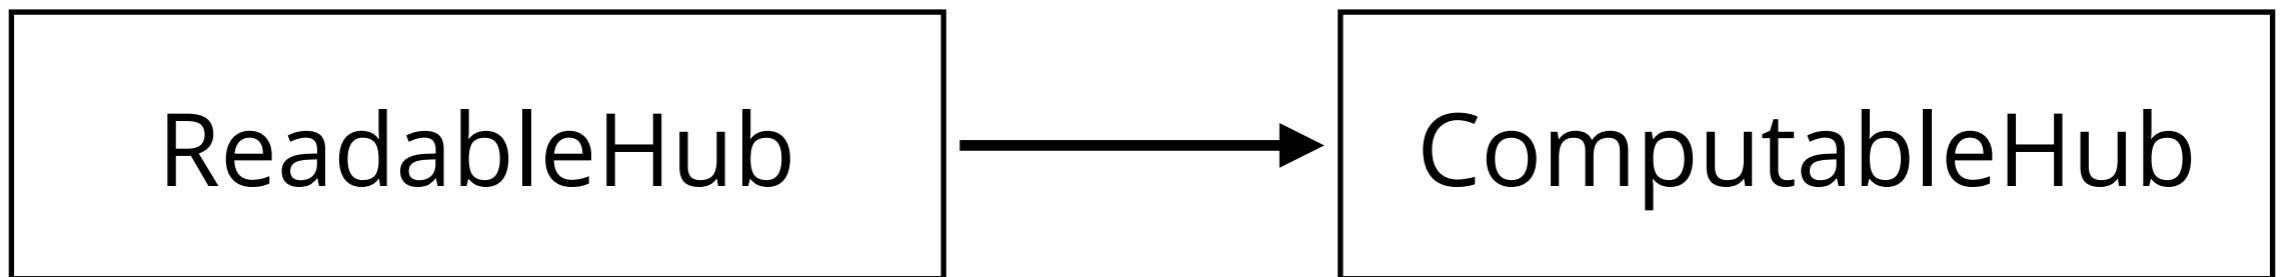
```
{  
    '00001': {'ip': '1.2.3.4', 'port': 80, 'domain': 'bar.com'},  
    '00002': {'ip': '1.2.3.5', 'port': 80, 'domain': 'foo.com'},  
    ...  
}
```

update()
→

ReadableHub

```
update

{
  'attack1': 'example.com',
  ...
}
```



```
def on_hook(_, kind, k, row):
    if kind == 'add':
        mitigations.maintain(k, action, row)

    if kind == 'delete':
        mitigations.unmaintain(k)

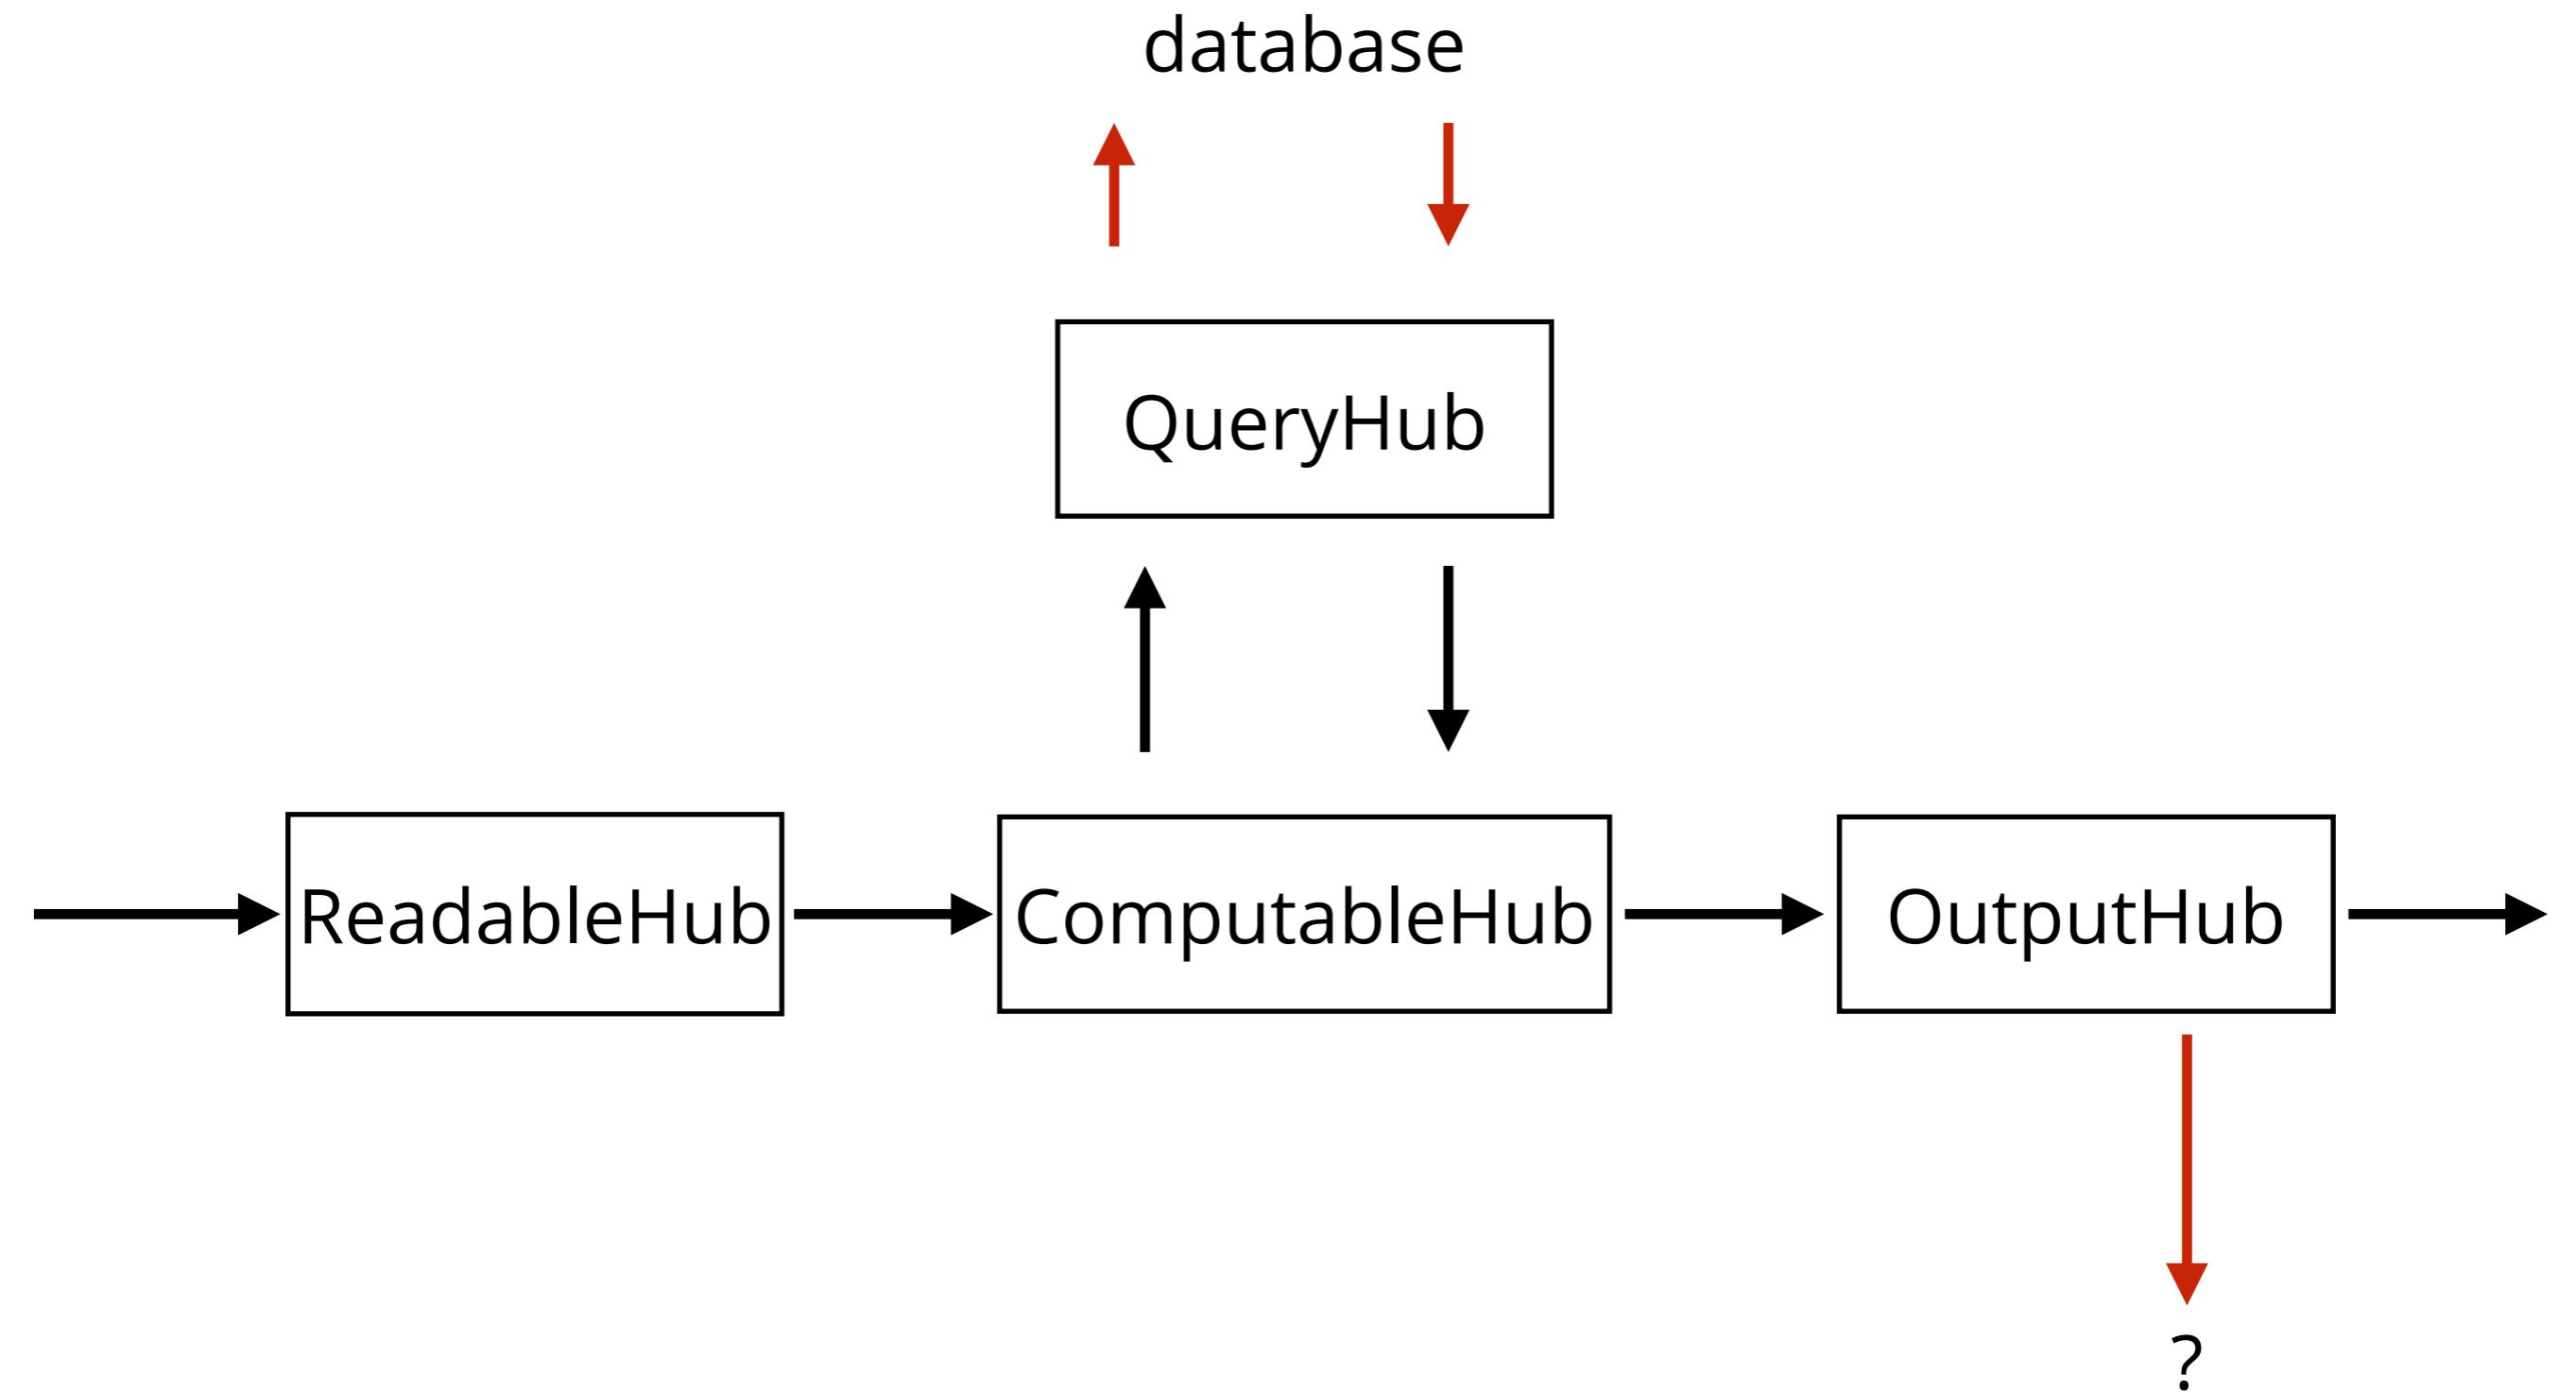
subscribe(readable, on_hook)

def action(row):
    return None
```

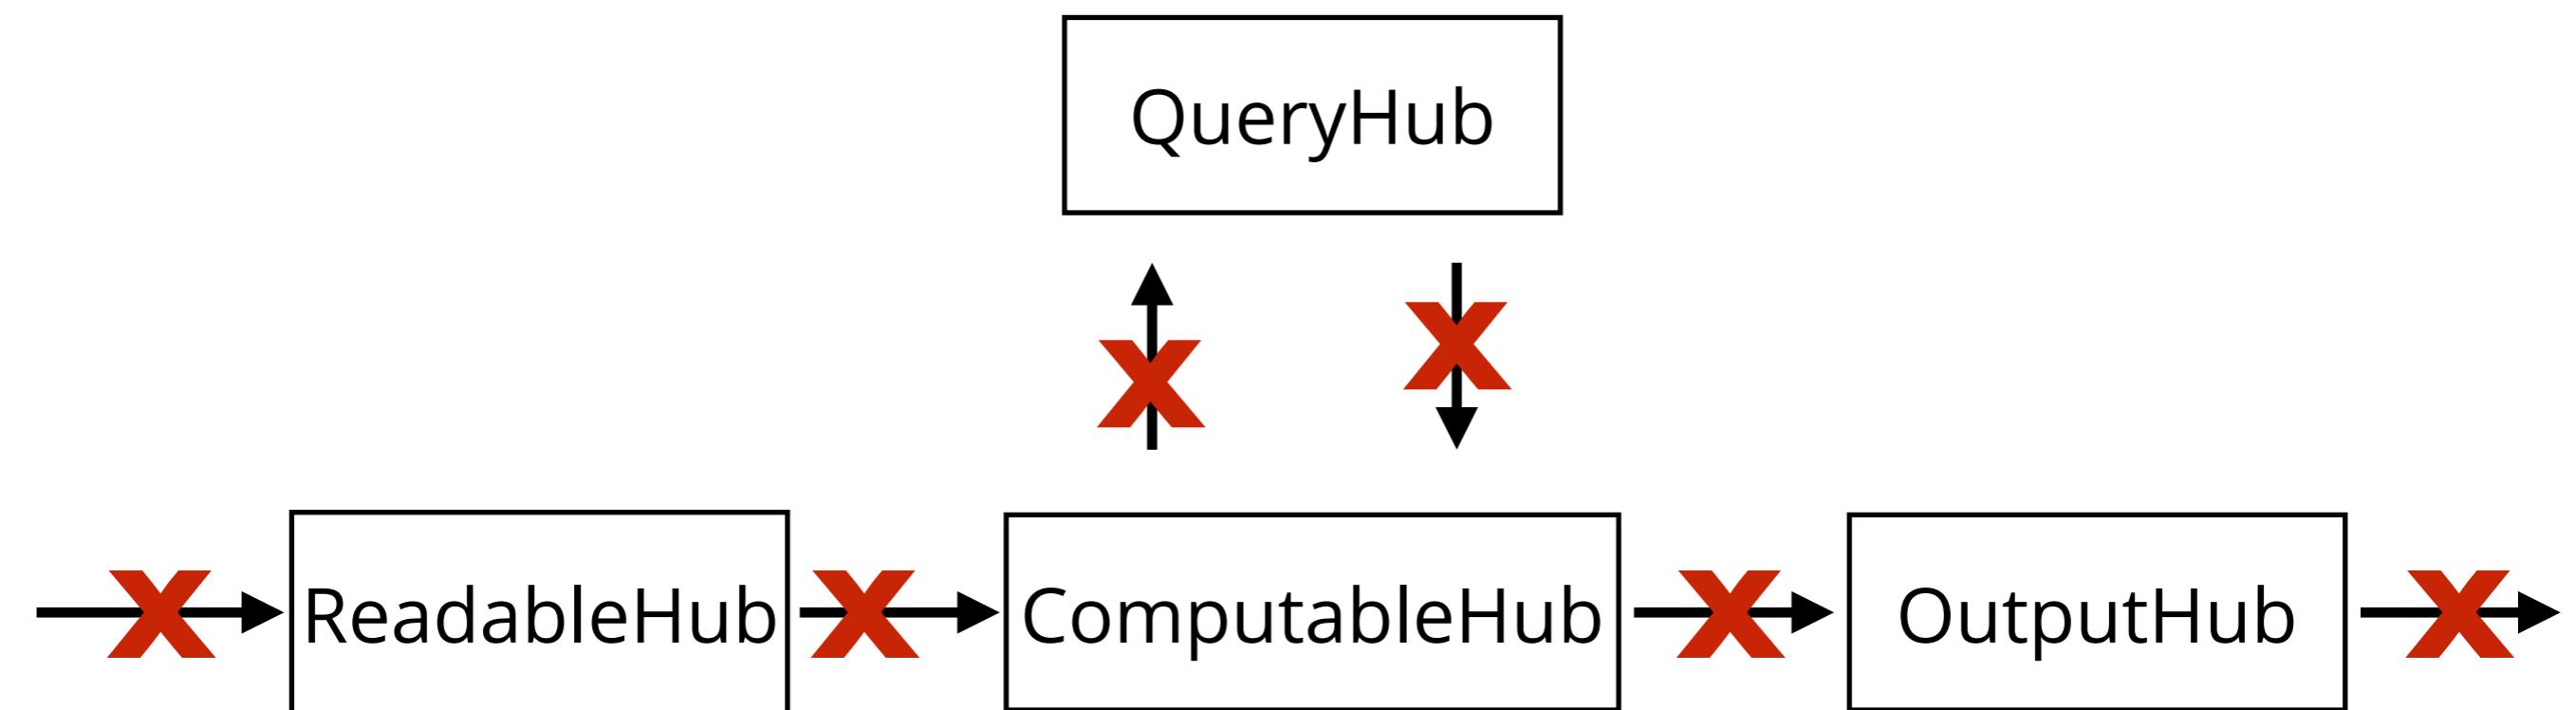
ComputableHub

OutputHub

?



Materialized



```
def action(row, plan_hub, subdomain_hub, toggle_hub):
    domain = row.value
    if not domain:
        return None

    if toggle_hub.get('all_mitigations_disabled').value != 'True':
        return None

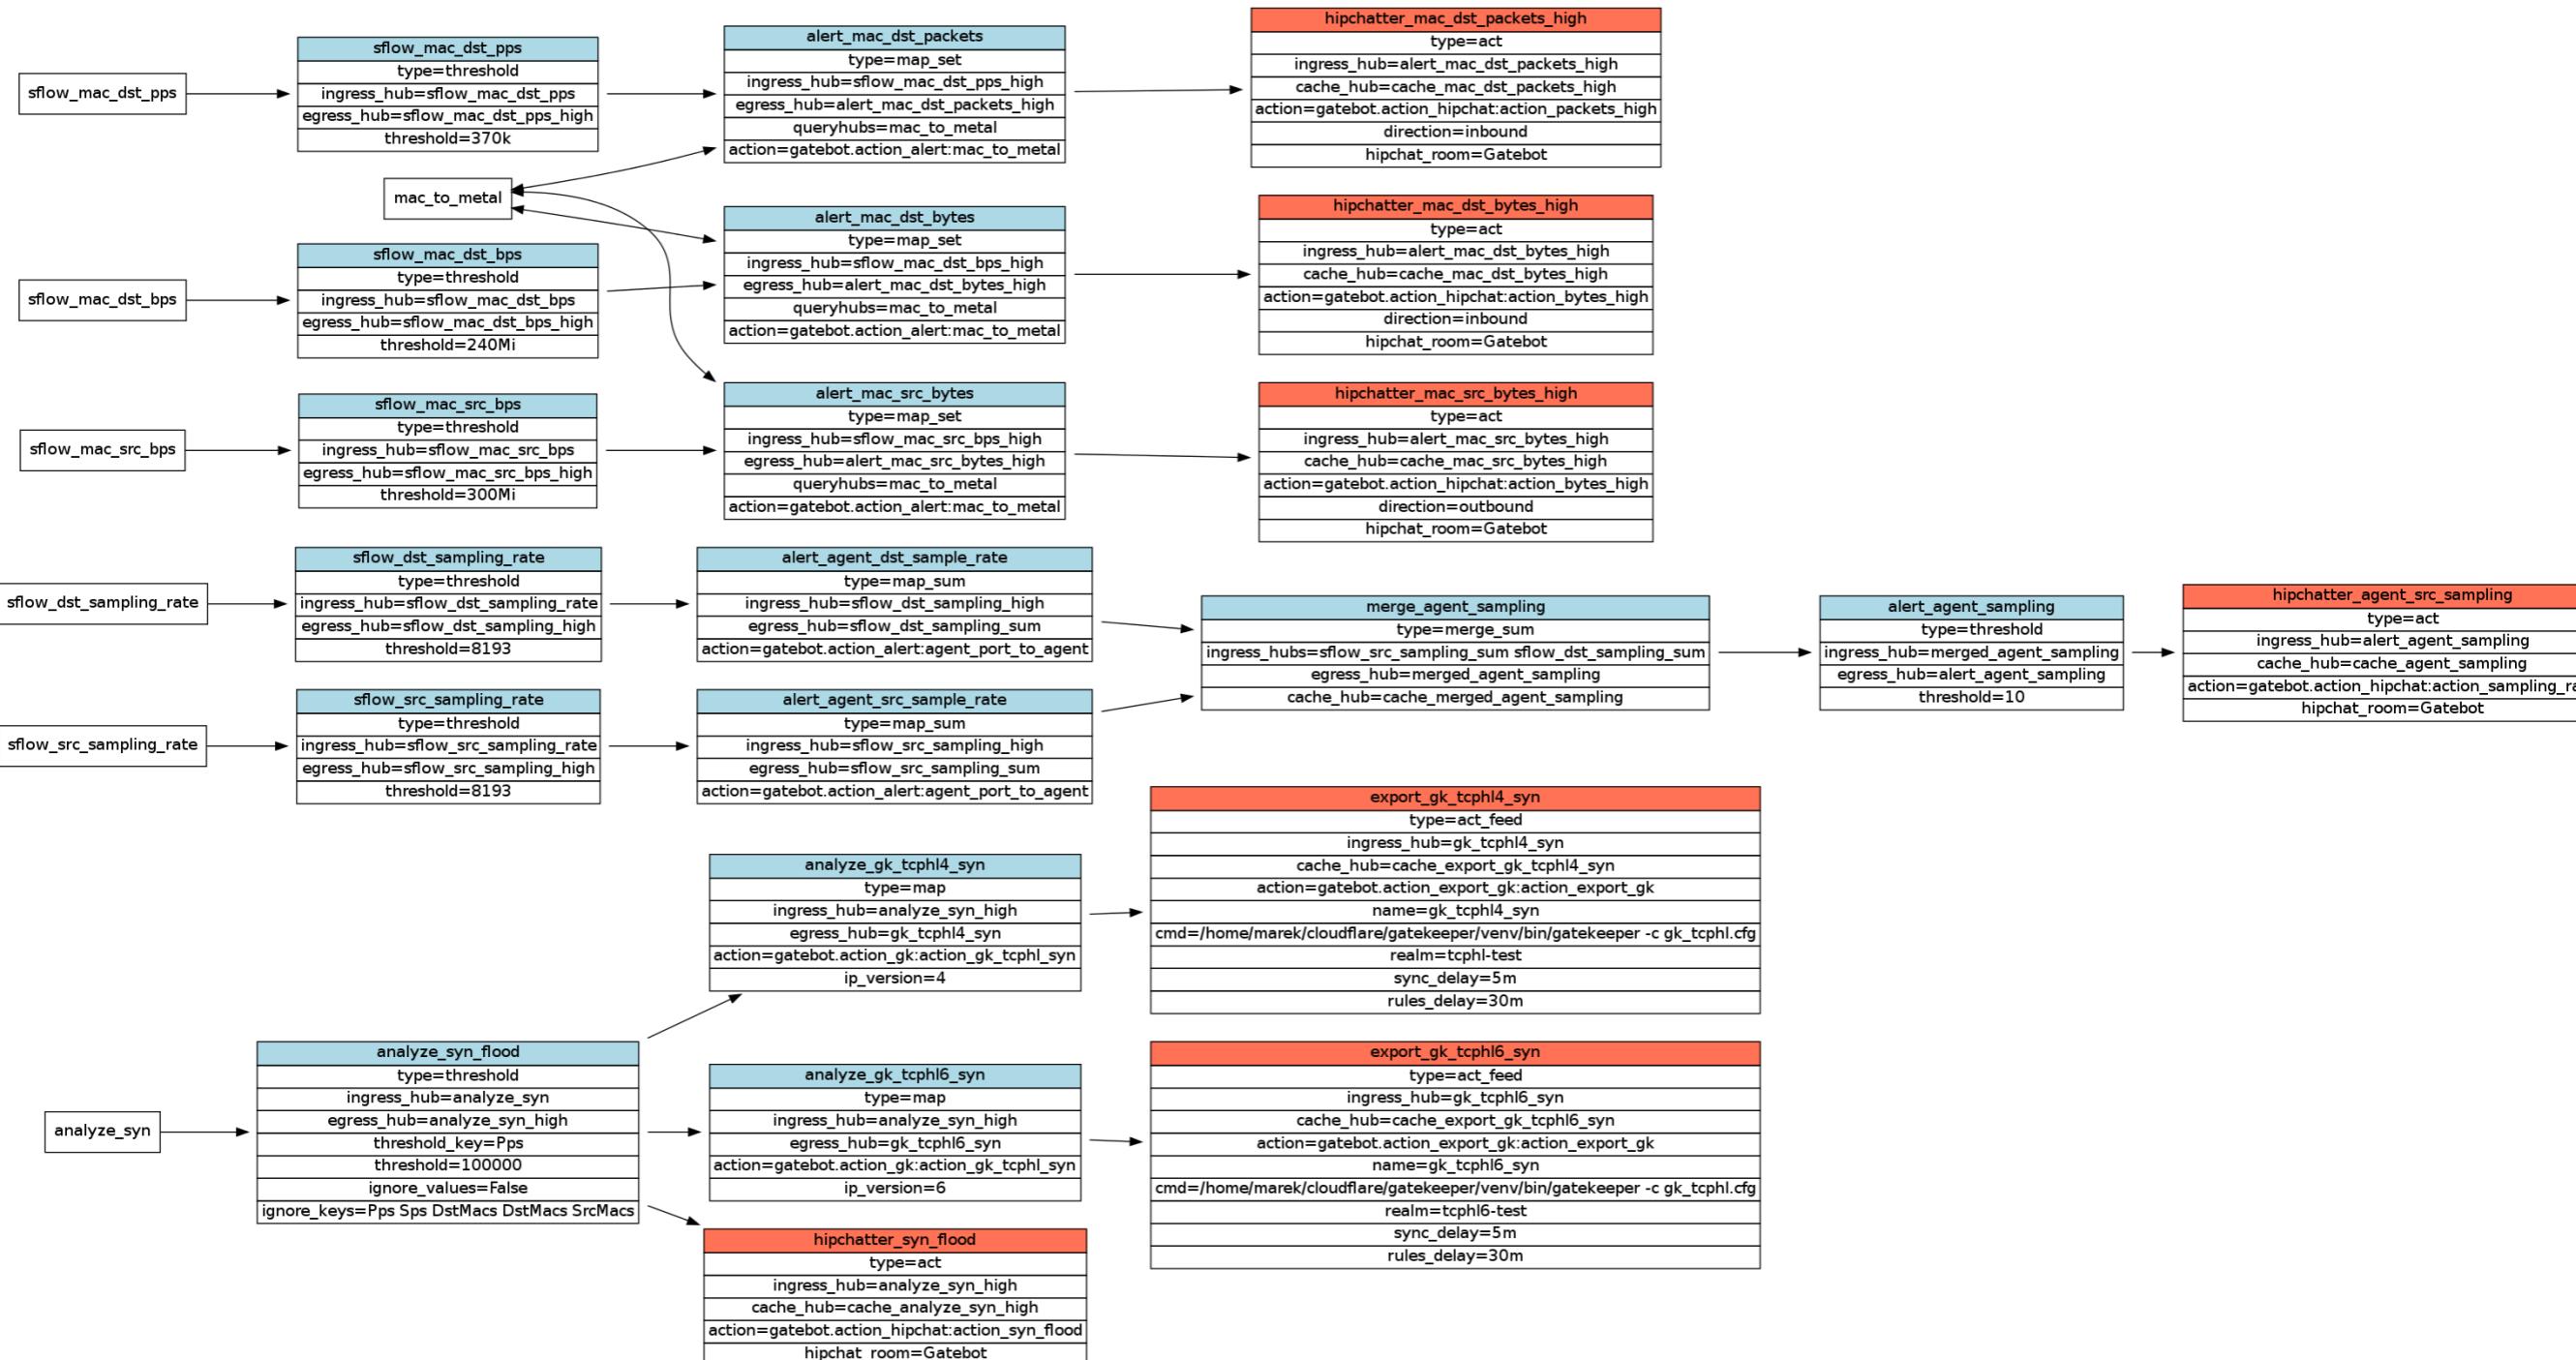
    qps = 100
    if plan_hub.get(domain).value in ('business', 'b'):
        qps = 500

    sd = (subdomain_hub.get(domain).value or '').split(' ')
    mitigation = \
        domain + \
        ' --qps=%s' % qps + \
        ' '.join('--except=%s' % s for s in sd)

    return mitigation
```

It works!

- Solid foundation!
- Composable!
- Scalable
- Maintainable
- But:
 - no event loop
 - lacks higher-order abstractions



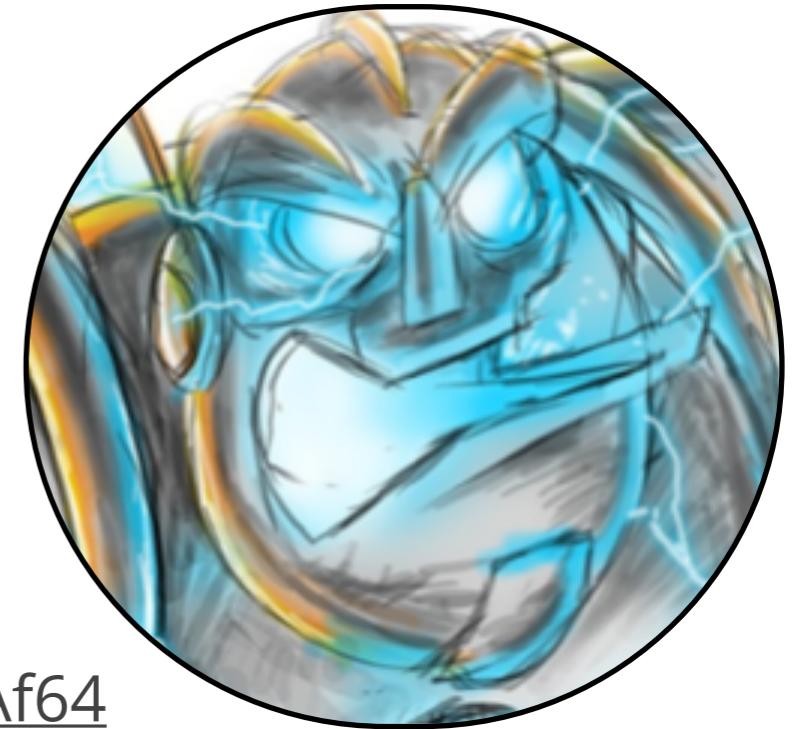
204 loc

```
marek@ubuntu-saucy:~/cloudflare/gatelogic/gatelogic$ cloc *py
 3 text files.
 3 unique files.
 0 files ignored.

http://cloc.sourceforge.net v 1.60  T=0.01 s (240.8 files/s, 23197.8 lines/s)
-----
      Language          files        blank      comment      code
-----
      Python              3            61           24         204
-----
      SUM:                3            61           24         204
-----
```

Thanks!

- FRP is great
 - <http://www.flapjax-lang.org/>
 - <https://www.youtube.com/watch?v=mEvo6TVAf64>
 - <https://www.youtube.com/watch?v=Agu6jipKfYw>



<https://github.com/cloudflare/gatelogic>

marek@cloudflare.com

@majek04