



EPI Framework Demo

Tim Müller (CCI Group, UvA)

t.muller@uva.nl

UMC Utrecht - 23 October 2023

Schedule

- 14:30 - 15:05: **Hello, world! in Brane** (*guided hands-on*)
- 15:05 - 15:25: **EPIF in the PoC** (*presentation*)
- 15:25 - 15:30: **Questions, thoughts, evaluation, ...**

Hello, world! in Brane (guided hands-on)

- Write your first **Hello, world!-package!**
- See the steps at <https://wiki.enablingpersonalizedinterventions.nl/user-guide>
 - Bottom-left, scroll down to “**35. Tutorials**”, then “**35.2.1. Hands-on session: Hello, world!**”
 - Or see: <https://tinyurl.com/umc-utrecht-demo>
- I’ll go through it on the board!

EPIF in the PoC

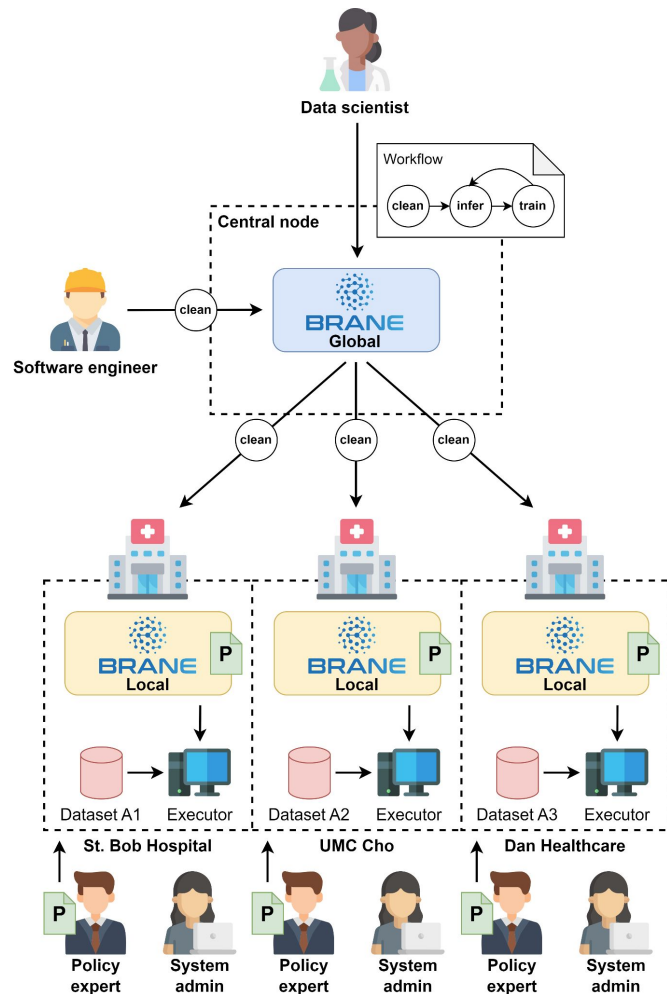
Or: The EPIF admin-side



I. Proof-of-Concept (PoC)

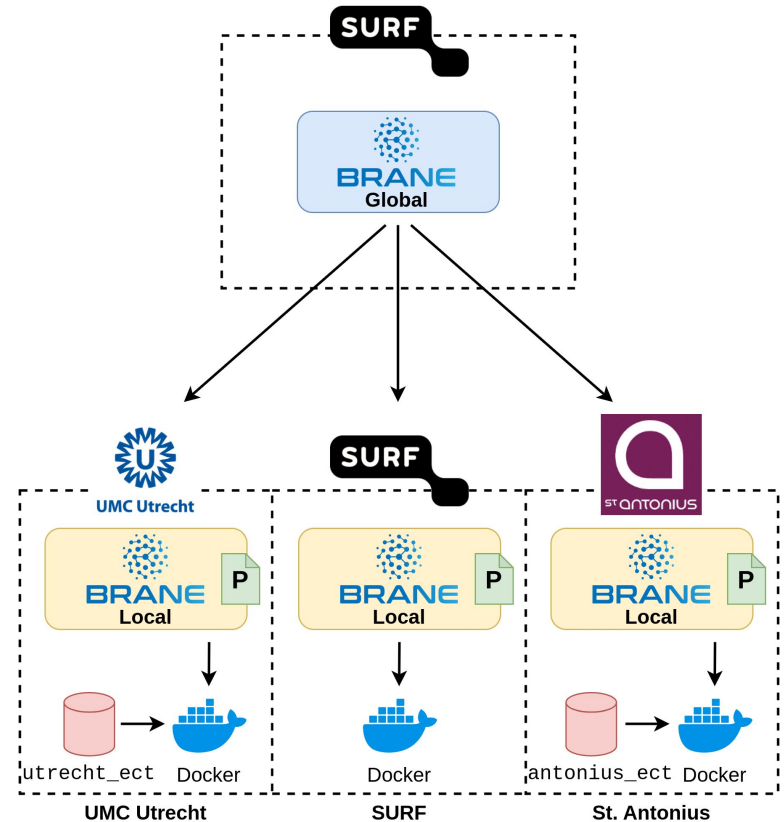
Where we left off...

- The **EPI Framework** is a:
“Federated workflow execution engine”
- We’ve discussed using the framework
 - Data scientist
 - Software engineer
 - Policy expert
- Now: **PoC-specifics as a system admin**



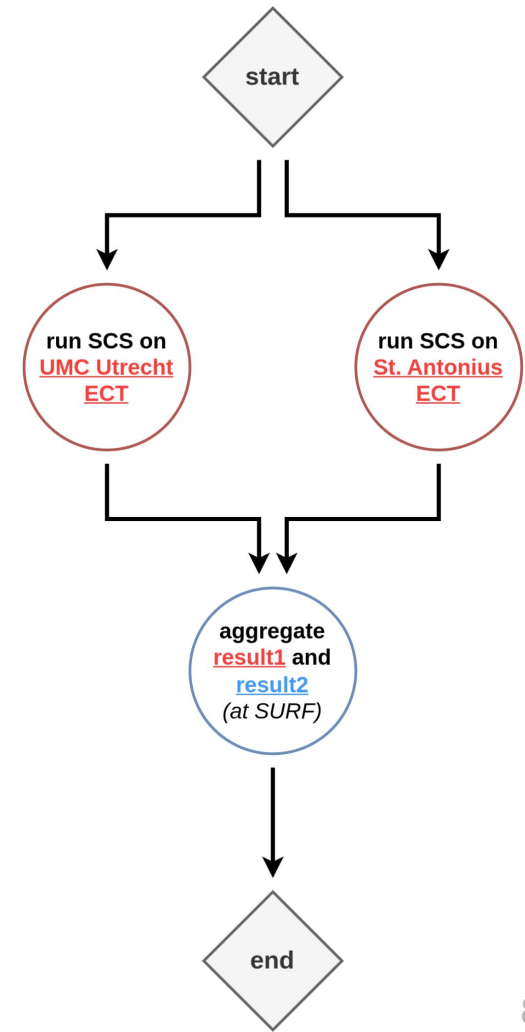
PoC - The EPIF-perspective

- One **central node**
 - Hosted by SURF
- Three **worker nodes**
 - SURF (aggregation)
 - St. Antonius, UMC Utrecht (local compute)
- Two **datasets**
 - umc_utrecht_ect
 - st_antoniuss_ect
- Two **use-cases**
 - Rosanne's use-case (stratified confidence)
 - Saba's use-case (synthetic data)



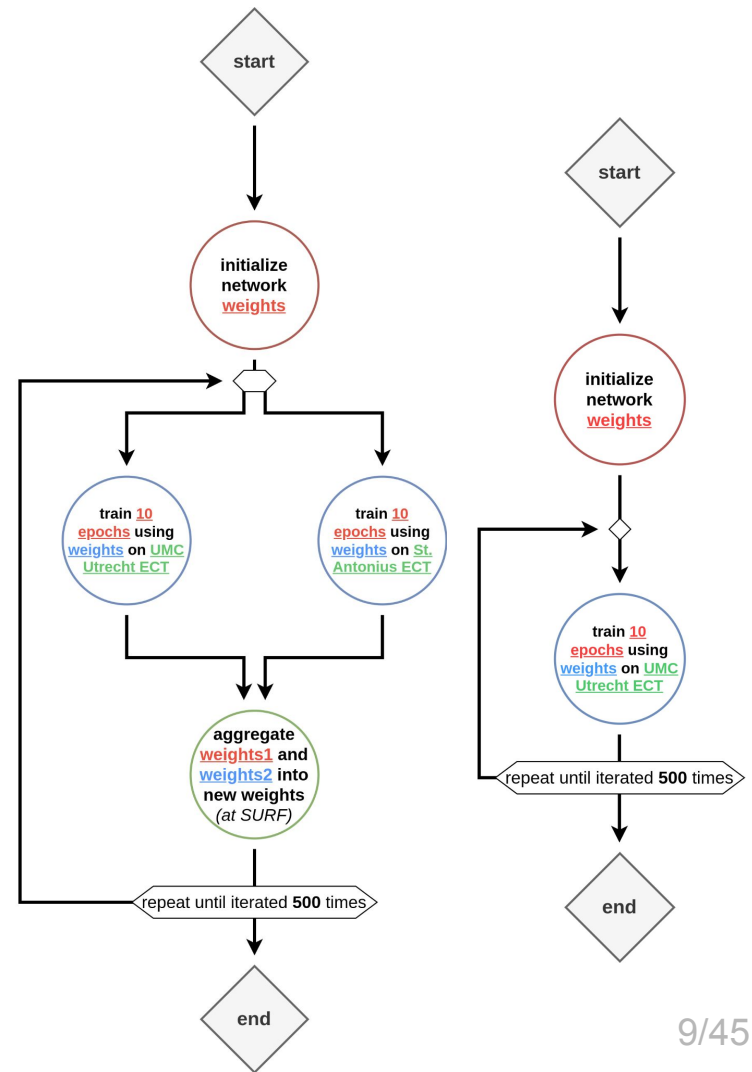
PoC - Rosanne's use-case

- **Stratified Confidence Sequence (SCS) analysis**
- **Federated analysis**
 - **Compute SCS locally** (UMC Utrecht, St. Antonius)
 - Send to **Trusted Third-Party (TTP)** (SURF)
 - Aggregate into **global result**
- **Result: single value (number)**





PoC - Saba's use-case

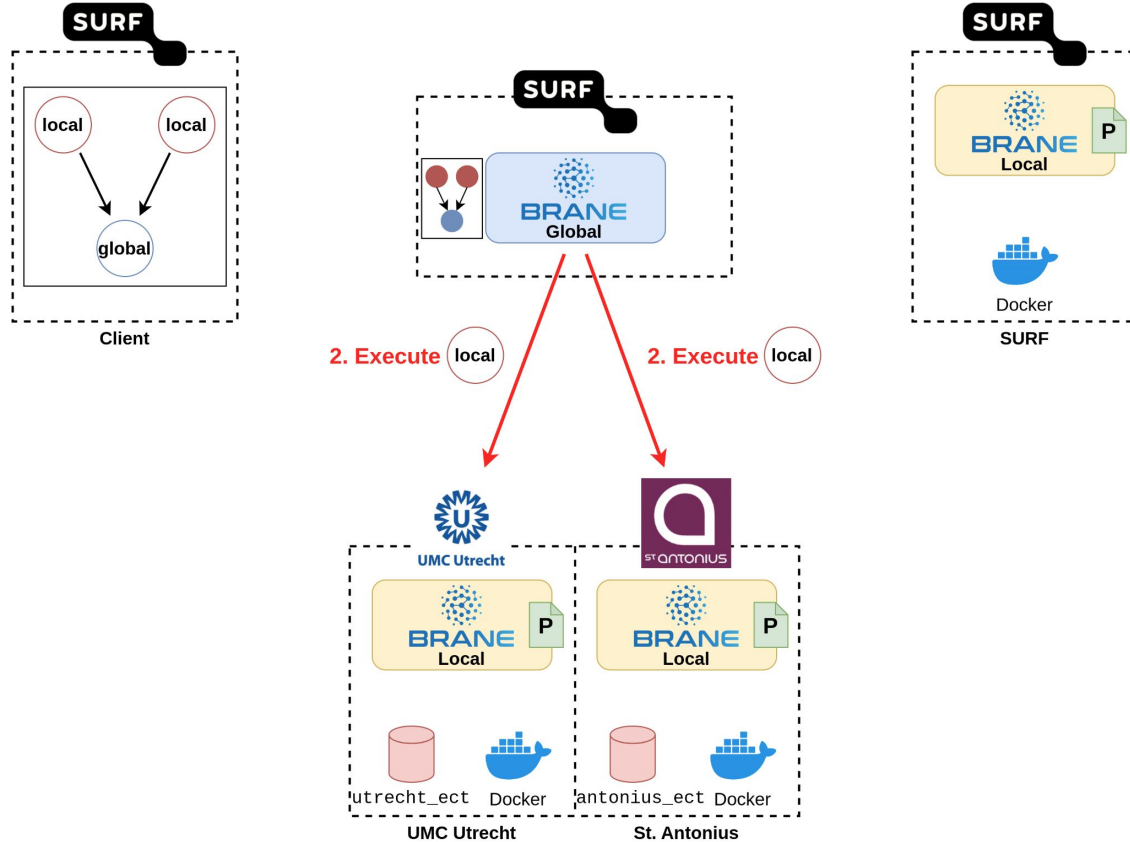
- Training a **synthetic generation algorithm**
 - First train as **federated algorithm**
 - Then generate new set from **central algorithm**
- Training simple validation **neural network**
 - Once as **federated algorithm** on raw data
 - Twice as **centralised algorithm** on half data
 - Once on **synthetic data** (also centralised)
- **Result: various trained NN models (weights)**
 - Different hyperparameters (number of iterations, hidden layers)
 - Compare using a test set (20% of data)






PoC in action - Rosanne's use-case



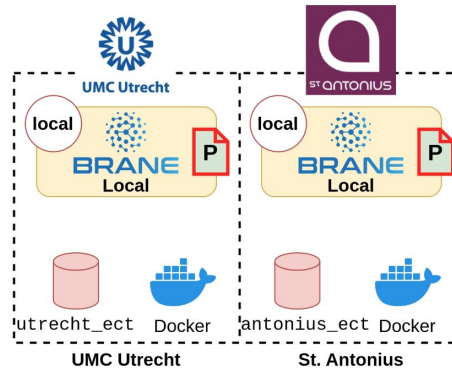
PoC in action - Rosanne's use-case



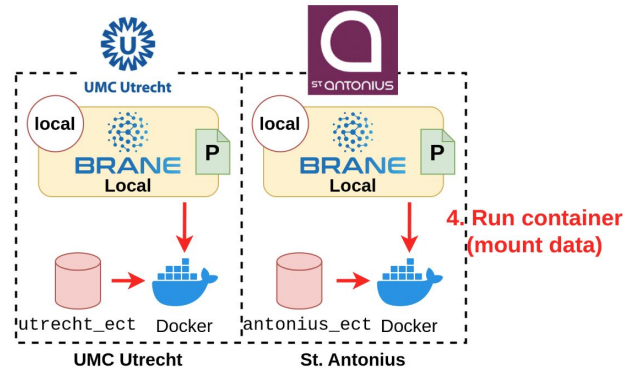



PoC in action - Rosanne's use-case



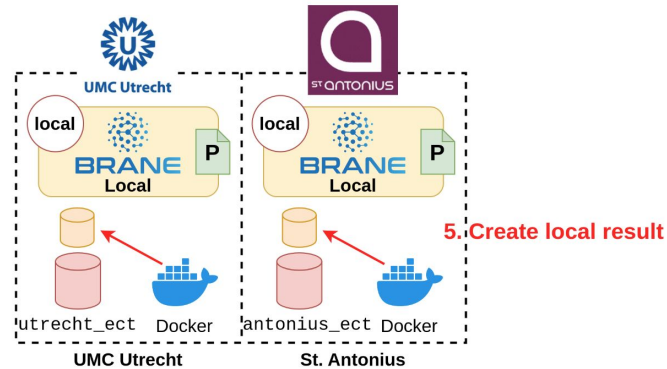



3. Consult policy



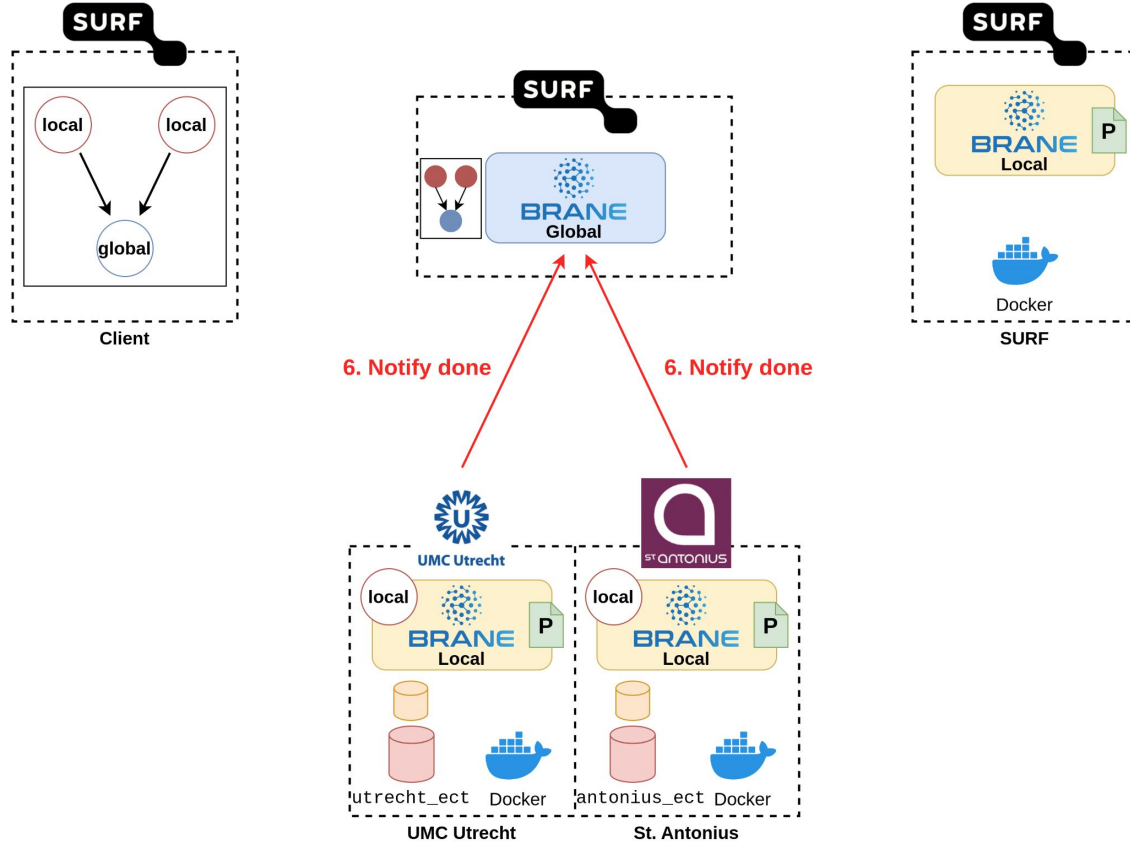
PoC in action - Rosanne's use-case





PoC in action - Rosanne's use-case





PoC in action - Rosanne's use-case





PoC in action - Rosanne's use-case



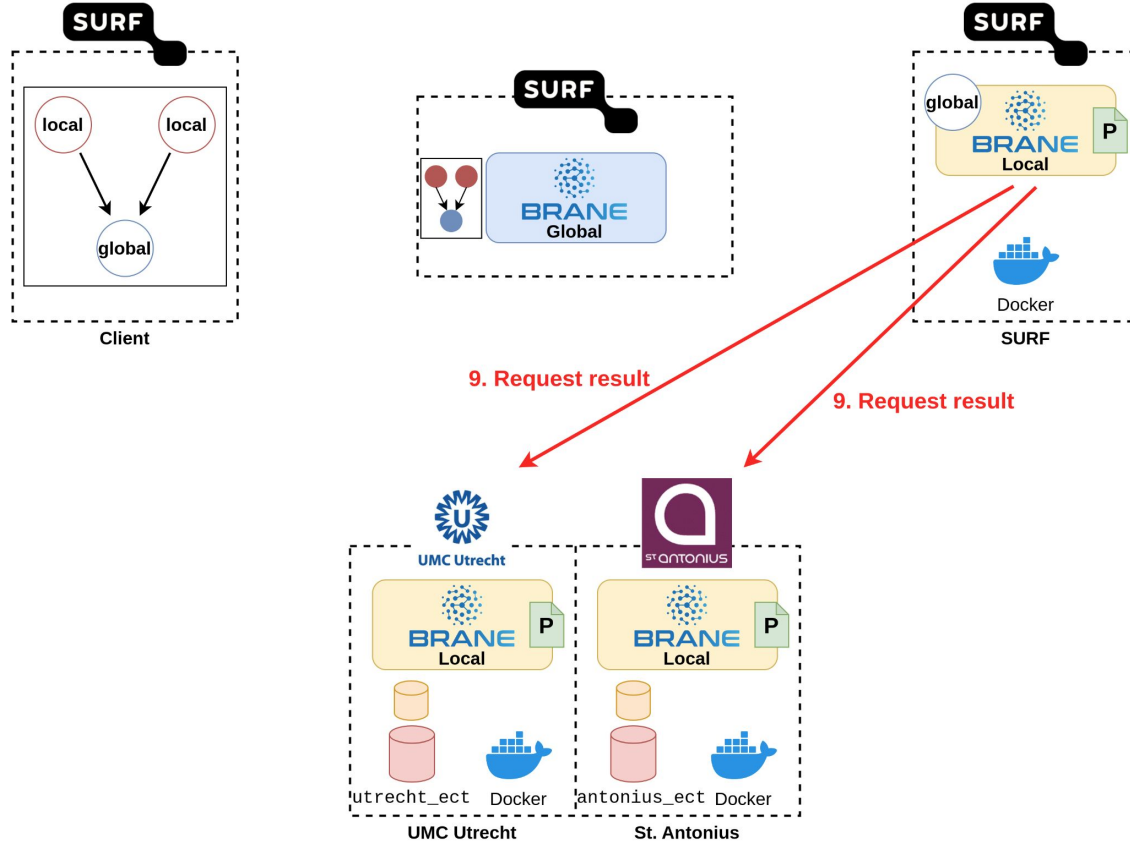
PoC in action - Rosanne's use-case






8. Consult policy



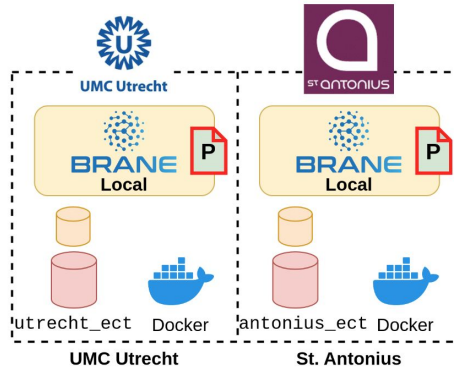
PoC in action - Rosanne's use-case



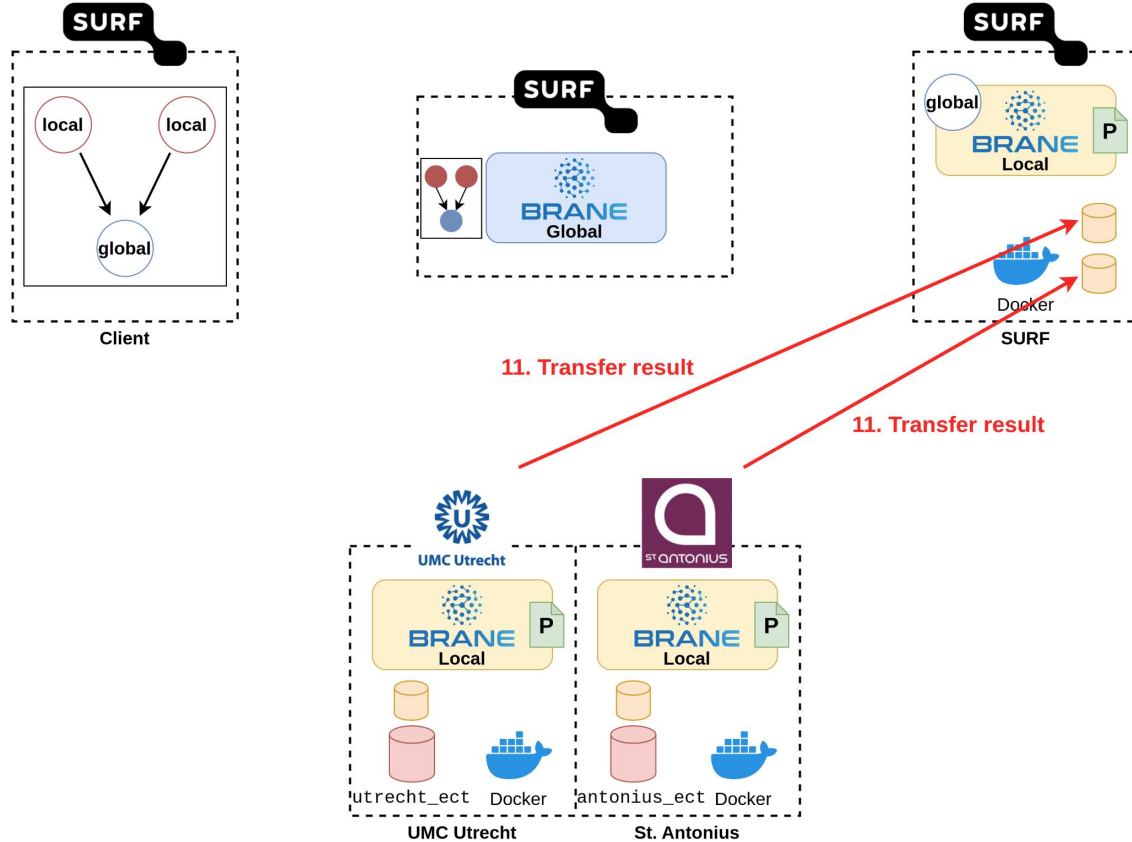
PoC in action - Rosanne's use-case






10. Consult policy




PoC in action - Rosanne's use-case





PoC in action - Rosanne's use-case





12. Run container (mount data)





PoC in action - Rosanne's use-case





13. Create global result





PoC in action - Rosanne's use-case





PoC in action - Rosanne's use-case




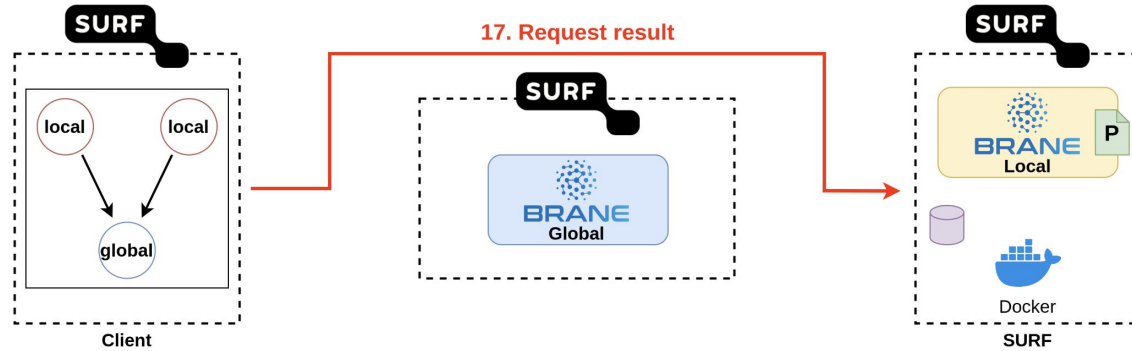
PoC in action - Rosanne's use-case





16. Remove intermediate results





PoC in action - Rosanne's use-case




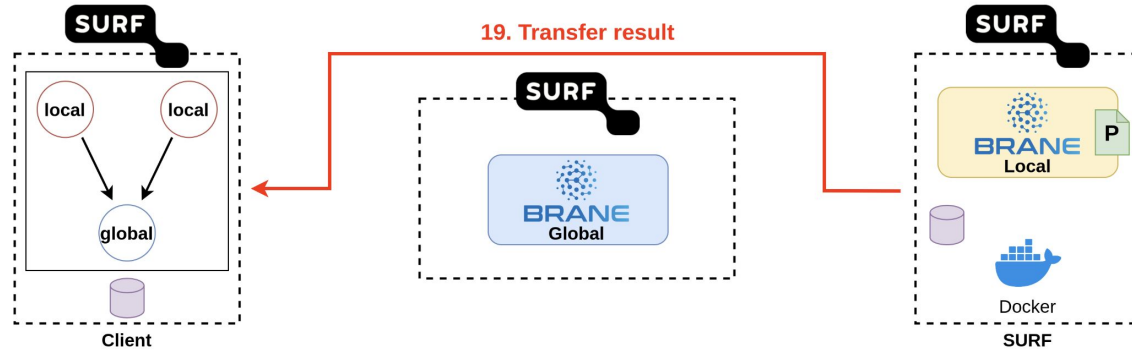
PoC in action - Rosanne's use-case






18. Consult policy




PoC in action - Rosanne's use-case




PoC in action - Rosanne's use-case



20. Done!!

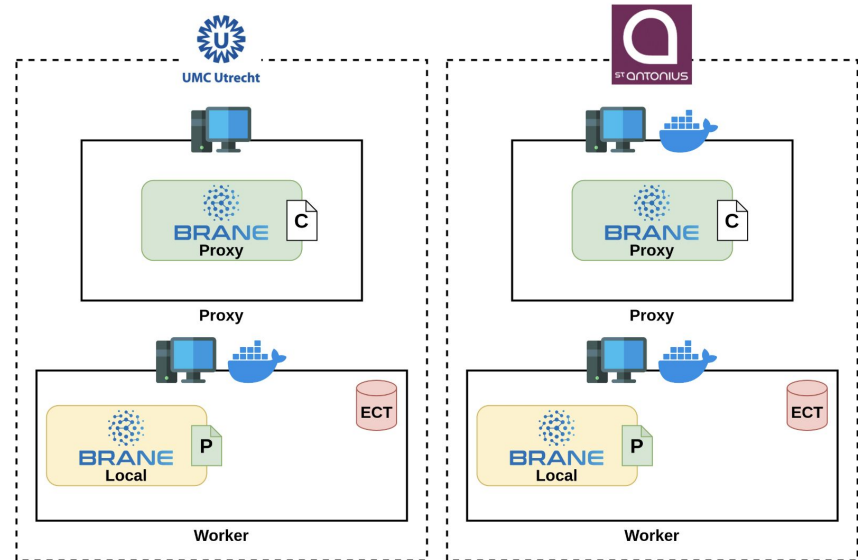
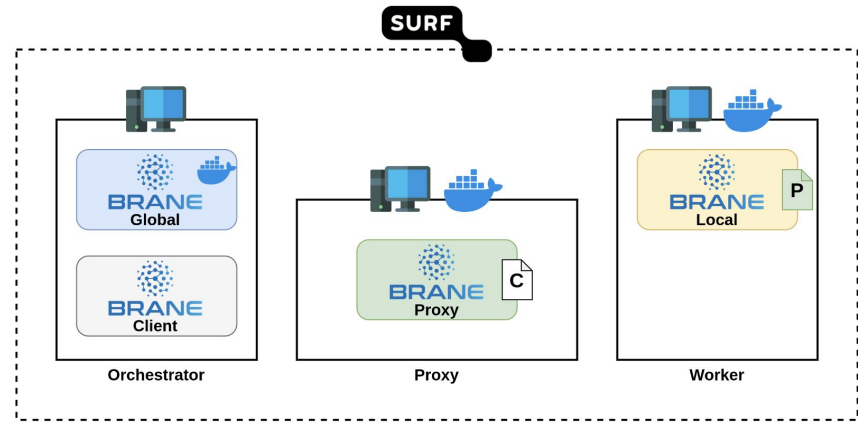




II. PoC Setup

Getting more accurate

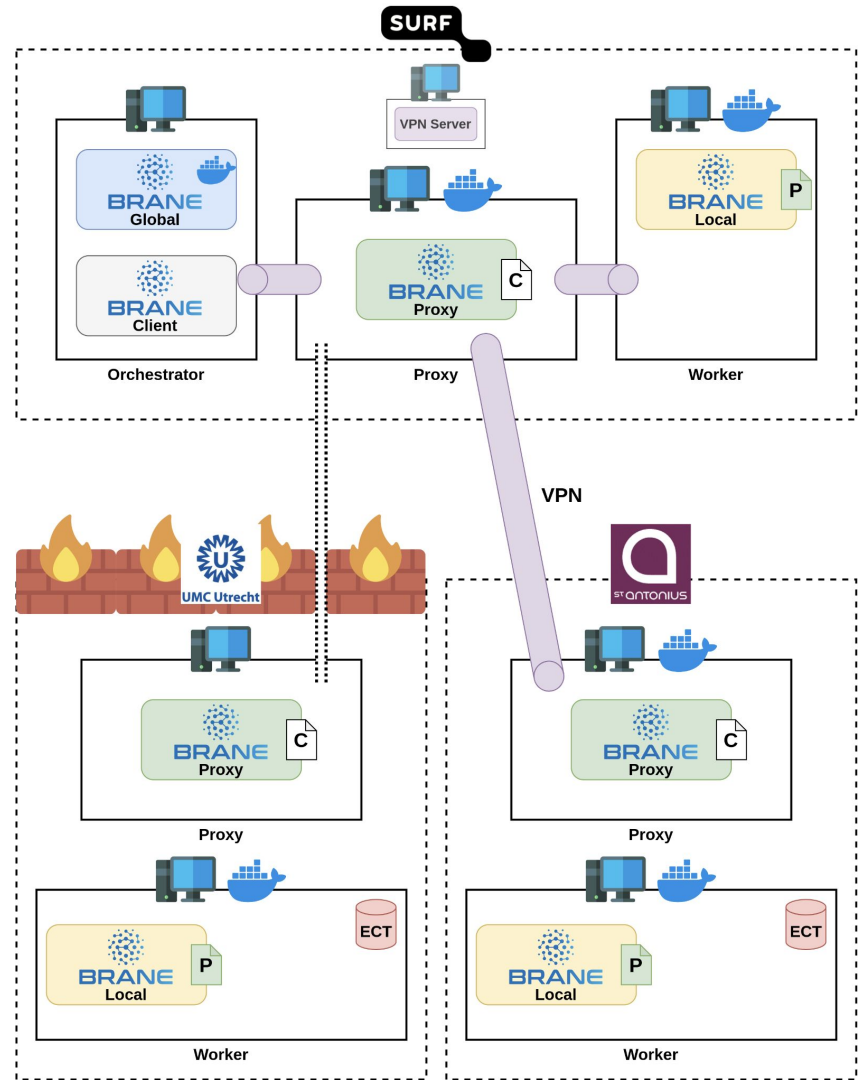
- Let's get detailed!
- **Multiple VMs** per domain
- **Proxy nodes**
 - Channel communication
 - "Gateways" for nodes
- Third VMs unused
 - No time to add Jamila's framework



Getting more secure

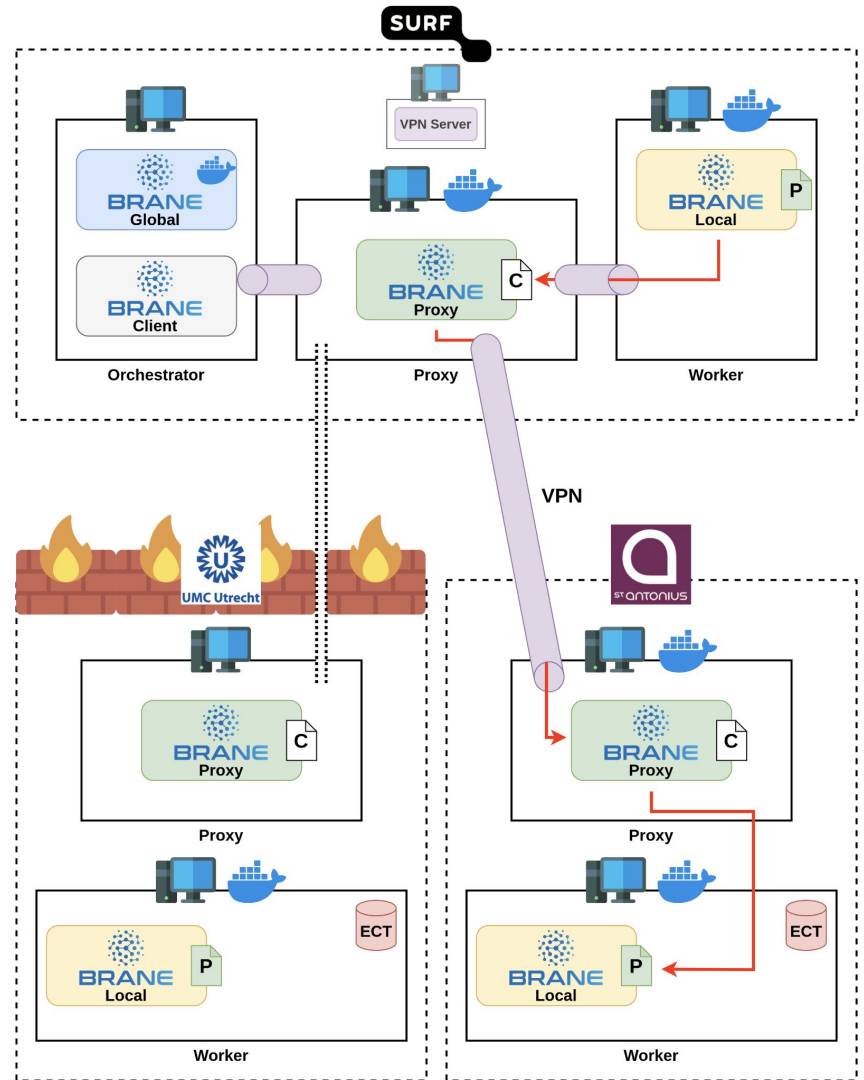
- Attempted to create realistic network
- **VPN St. Antonius / SURF**
 - strongSwan VPN¹ (IPsec)
- **UMC Utrecht firewall** restrictions
 - Only proxy nodes are allowed to talk

¹ <https://www.strongswan.org/>



Getting more secure

- To adhere to security, we need to define **specific network routes**
 - Specific hops
 - Specific interfaces





III. Configuring Brane nodes

node.yml

- Defines **node context**
 - Defines **node kind** (central, worker, proxy, ...)
 - Defines **other config locations**
 - Defines **ports**
 - Defines **container names**
 - ...
- Comparable to
~/kube/config.yaml

```
12
13
14 hostnames:
15   central: 145.38.187.47
16   surf: 145.38.187.47
17   umc_utrecht: 143.121.240.12
18   st_antonius: 194.13.118.4
19 node: !proxy
20 paths:
21   certs: /home/muller/brane/config/certs
22   proxy: /home/muller/brane/config/proxy.yml
23 services:
24   prx:
25     name: brane-prx-proxy
26     address: http://brane-prx-proxy:1080
27     bind: 0.0.0.0:1080
28     external_address: http://proxy:1080
29
```

```
12
13
14 hostnames:
15   central: 145.38.187.47
16   surf: 145.38.187.47
17   umc_utrecht: 143.121.240.12
18   st_antonius: 194.13.118.4
19 node: !worker
20 name: umc_utrecht
21 paths:
22   certs: /home/muller/brane/config/certs
23   packages: /home/muller/brane/packages
24   backend: /home/muller/brane/config/backend.y
25   policies: /home/muller/brane/config/policies
26   proxy: null
27   data: /home/muller/brane/data
28   results: /home/muller/brane/results
29   temp_data: /home/muller/brane/temp_data
30   temp_results: /home/muller/brane/temp_result
31 services:
32   prx: !external
33     address: http://umc_utrecht:1080
34   reg:
35     name: brane-reg-umc_utrecht
36     address: https://brane-reg-umc_utrecht:108
37     bind: 0.0.0.0:1080
38     external_address: https://umc_utrecht:1080
39   job:
```



node.yml - specific interfaces

- We use `hostnames` to **customize addressing**
 - Node-local contents of `/etc/hosts`
- **Different nodes talk to different interfaces**
 - ...while sending the same hostname around
- Hacky, but it works!

```
12
13
14 hostnames:
15   central: 10.0.42.98
16   surf: 10.0.42.98
17   umc_utrecht: 143.121.240.12
18   st_antonius: 194.13.119.49
19 node: !worker
20 name: st_antonius
```



```
12
13
14 hostnames:
15   central: 145.38.187.47
16   surf: 145.38.187.47
17   umc_utrecht: 143.121.240.12
18   st_antonius: 194.13.118.4
19 node: !proxy
20 paths:
21   certs: /home/muller/brane/config/certs
22   proxy: /home/muller/brane/config/proxy.yml
23 services:
24   prx:
25     name: brane-prx-proxy
26     address: http://brane-prx-proxy:1080
27     bind: 0.0.0.0:1080
28     external_address: http://proxy:1080
29
```

```
12
13
14 hostnames:
15   central: 145.38.187.47
16   surf: 145.38.187.47
17   umc_utrecht: 143.121.240.12
18   st_antonius: 194.13.118.4
19 node: !worker
20 name: umc_utrecht
21 paths:
22   certs: /home/muller/brane/config/certs
23   packages: /home/muller/brane/packages
24   backend: /home/muller/brane/config/backend.y
25   policies: /home/muller/brane/config/policies
26   proxy: null
27   data: /home/muller/brane/data
28   results: /home/muller/brane/results
29   temp_data: /home/muller/brane/temp_data
30   temp_results: /home/muller/brane/temp_result
31 services:
32   prx: !external
33     address: http://umc_utrecht:1080
34   reg:
35     name: brane-reg-umc_utrecht
36     address: https://brane-reg-umc_utrecht:1080
37     bind: 0.0.0.0:1080
38     external_address: https://umc_utrecht:1080
39   job:
```

proxy.yml

- **Routes** network traffic
- **Authenticates** clients
 - Only clients presenting signed client certificate
- Routes through BFCs

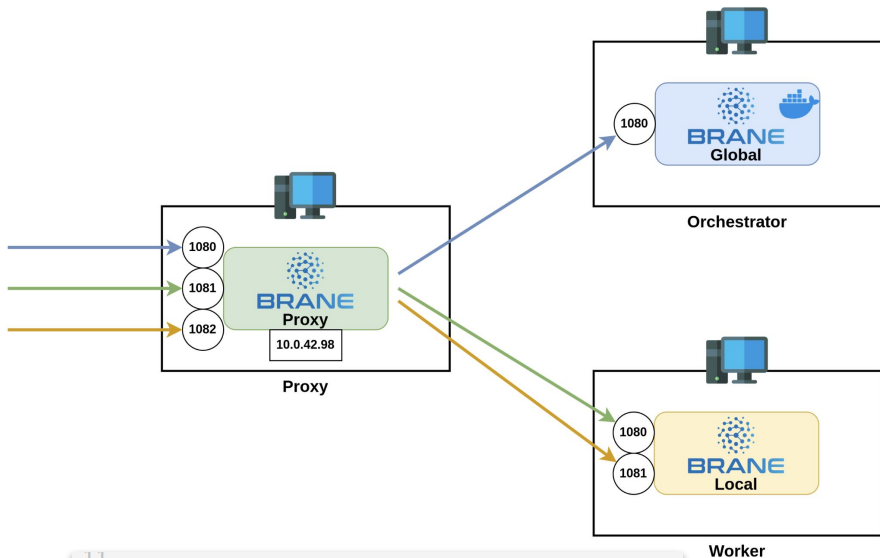
```
[lut_99@gamelinux config]$ branectl generate proxy
Successfully generated ./proxy.yml
[lut_99@gamelinux config]$ _
```

```
11
12
13   outgoing_range:
14     start: 1082
15     end: 1085
16   incoming:
17     1080: https://143.121.240.11:1080
18     1081: https://143.121.240.11:1081
19   forward: null
20
```

proxy.yml - specific hops

- Route incoming traffic through **single IP**
 - Essentially NATs using the `incoming-table`
- Outgoing traffic already **routed through proxy** by default

```
[lut_99@gamelinux config]$ branectl generate proxy
Successfully generated ./proxy.yml
[lut_99@gamelinux config]$ _
```



```
11
12
13   outgoing_range:
14     start: 1082
15     end: 1085
16   incoming:
17     1080: https://143.121.240.11:1080
18     1081: https://143.121.240.11:1081
19   forward: null
20
```

policies.yml



- Defines **policies!**
 - Implemented as simple rule-based rules
- One set defines **which container to execute**
- One set defines **who can access which dataset**
 - Identification based on **client-side certificates**

```
[lut_99@gameLinux worker]$ branectl packages hash ~/.local/share/brane/packages_2.0.0/epi_rosanne/1.0.0/image.tar
QS43h4ycr/PdYZTtwUAKw0c68qkEZiz9oDwCo0kMdgGE=
[lut_99@gameLinux worker]$ _
```

```
15
16 users:
17 - policy: deny
18   user: surf
19   data: umc_utrecht_ect
20 - policy: deny
21   user: surf
22   data: umc_utrecht_ect_train
23 - policy: deny
24   user: surf
25   data: umc_utrecht_ect_test
26 - policy: allow_user_all
27   user: surf
28 - policy: deny
29   user: client
30   data: umc_utrecht_ect
31 - policy: deny
32   user: client
33   data: umc_utrecht_ect_train
34 - policy: deny
35   user: client
36   data: umc_utrecht_ect_test
37 - policy: allow_user_all
38   user: client
39 - policy: deny_all
40 containers:
41 - policy: allow
42   name: Rosanne's use-case container
43   hash: "QS43h4ycr/PdYZTtwUAKw0c68qkEZiz9oDwCo0kMdgGE="
44 - policy: allow
45   name: Saba train synthetic container
46   hash: "HdyBrlwCpl9ltQDl7u9Ac/vexBBInRgoWSuHR09C"
47 - policy: deny_all
48
```

Certificates

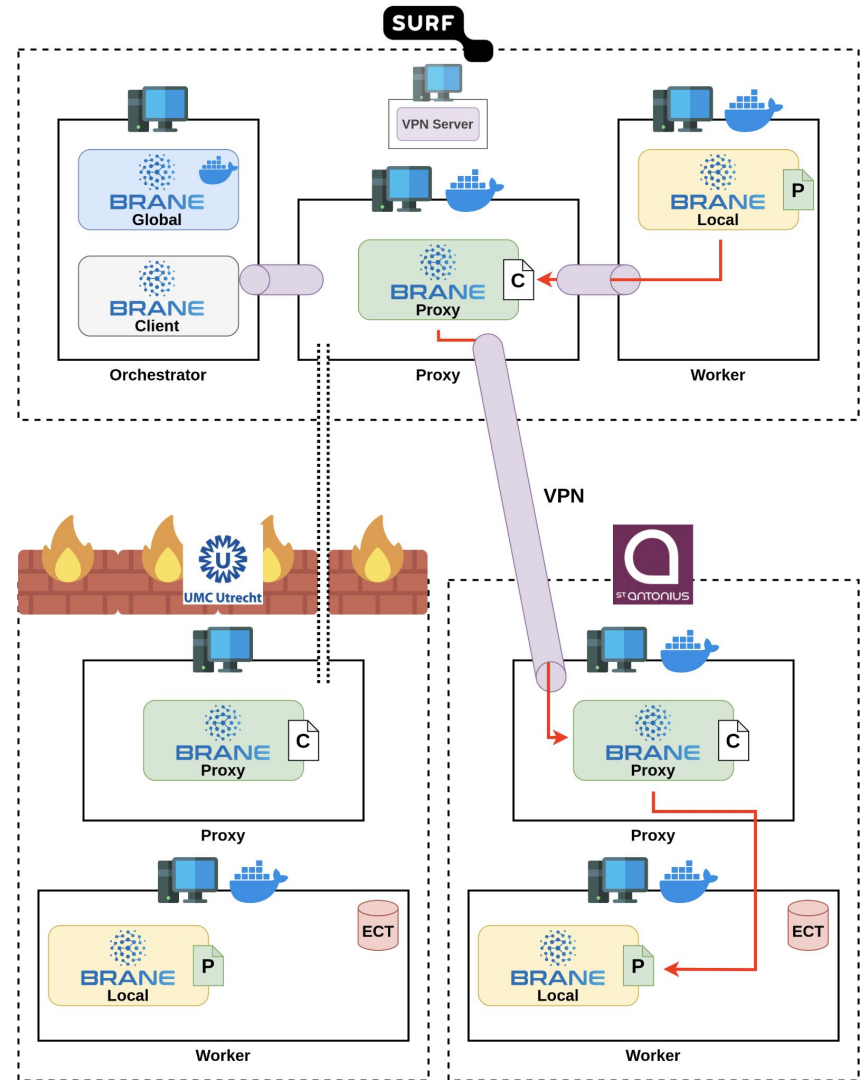
- **Authenticates clients**
 - So that policy may authorize them
- Used to **encrypt traffic** (without BFCs, that is)
- Required:
 - Root certificate (per node)
 - Server certificate (per node)
 - Client certificate (per node, per client)

```
[lut_99@gameLinux certs]$ branectl generate certs server umc_utrecht
Downloading http://github.com/cloudflare/cfssl/releases/download/v1.6.3/cfssl\_1.6.3\_linux\_amd64...
> Checksum 16b42bfc592dc4d0ba1e51304f466cae7257edec13743384caf4106195ab6047 OK
Successfully generated server certificates for domain umc_utrecht
[lut_99@gameLinux certs]$ ls
ca.csr ca-key.pem ca.pem server.csr server-key.pem server.pem
[lut_99@gameLinux certs]$ _
```

```
[lut_99@gameLinux surf]$ branectl generate certs client surf -c ../../ca.pem -k ../../ca-key.pem
Successfully generated client certificates for domain surf
[lut_99@gameLinux surf]$ _
```


Takeaways

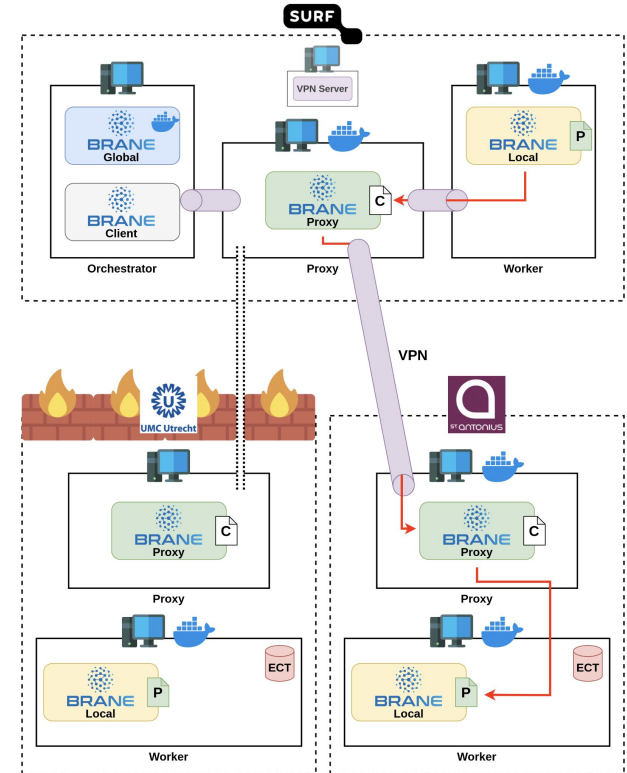
- Brane allows configuring **specific network routes**
 - Specific hops ✓
 - Specific interfaces ✓
- Policies defined through **rule-based YAML file**
- **Certificates** used for authentication/encryption



IV. Conclusion

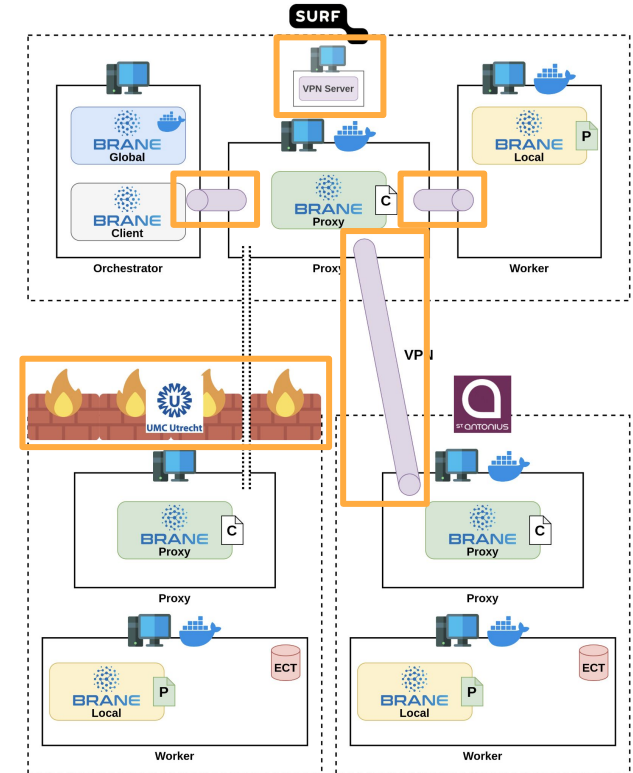
EPIF in the PoC

- Setup between **SURF**, **UMC Utrecht** and **St. Antonius**
- **Two use-cases**
 - Rosanne's stratified confidence sequence analysis
 - Saba's synthetic data generation
- **Realistic network security** (hopefully)
 - Brane supports required routing requirements
 - ...also *because* of the Proof-of-Concept



What if... BFC Framework

- BFC Framework can add in security as **Virtualized Network Functions**
 - Spawn as Docker container
 - Route traffic through container
- Can interact with **policy**
 - e.g., “Only share with St. Bob Hospital if they are trusted and setup a VPN with us”
- Only useful for **inter-domain** networking





BRANE

Tim Müller (t.muller@uva.nl)

<https://enablingpersonalizedinterventions.nl>

<https://github.com/epi-project/brane>

<https://wiki.enablingpersonalizedinterventions.nl>



The icons (not logos) in this presentation are from: Freepik, juicy_fish, Ultimatearm, Vector Valley