neat

Transport Evolution on top of the BSD's

[tj]
tj@enoti.me

NEAT is funded by the European Union's Horizon 2020 research and innovation programme under grant agreement no. 644334.

Horizon 2020
European Union funding
for Research & Innovation

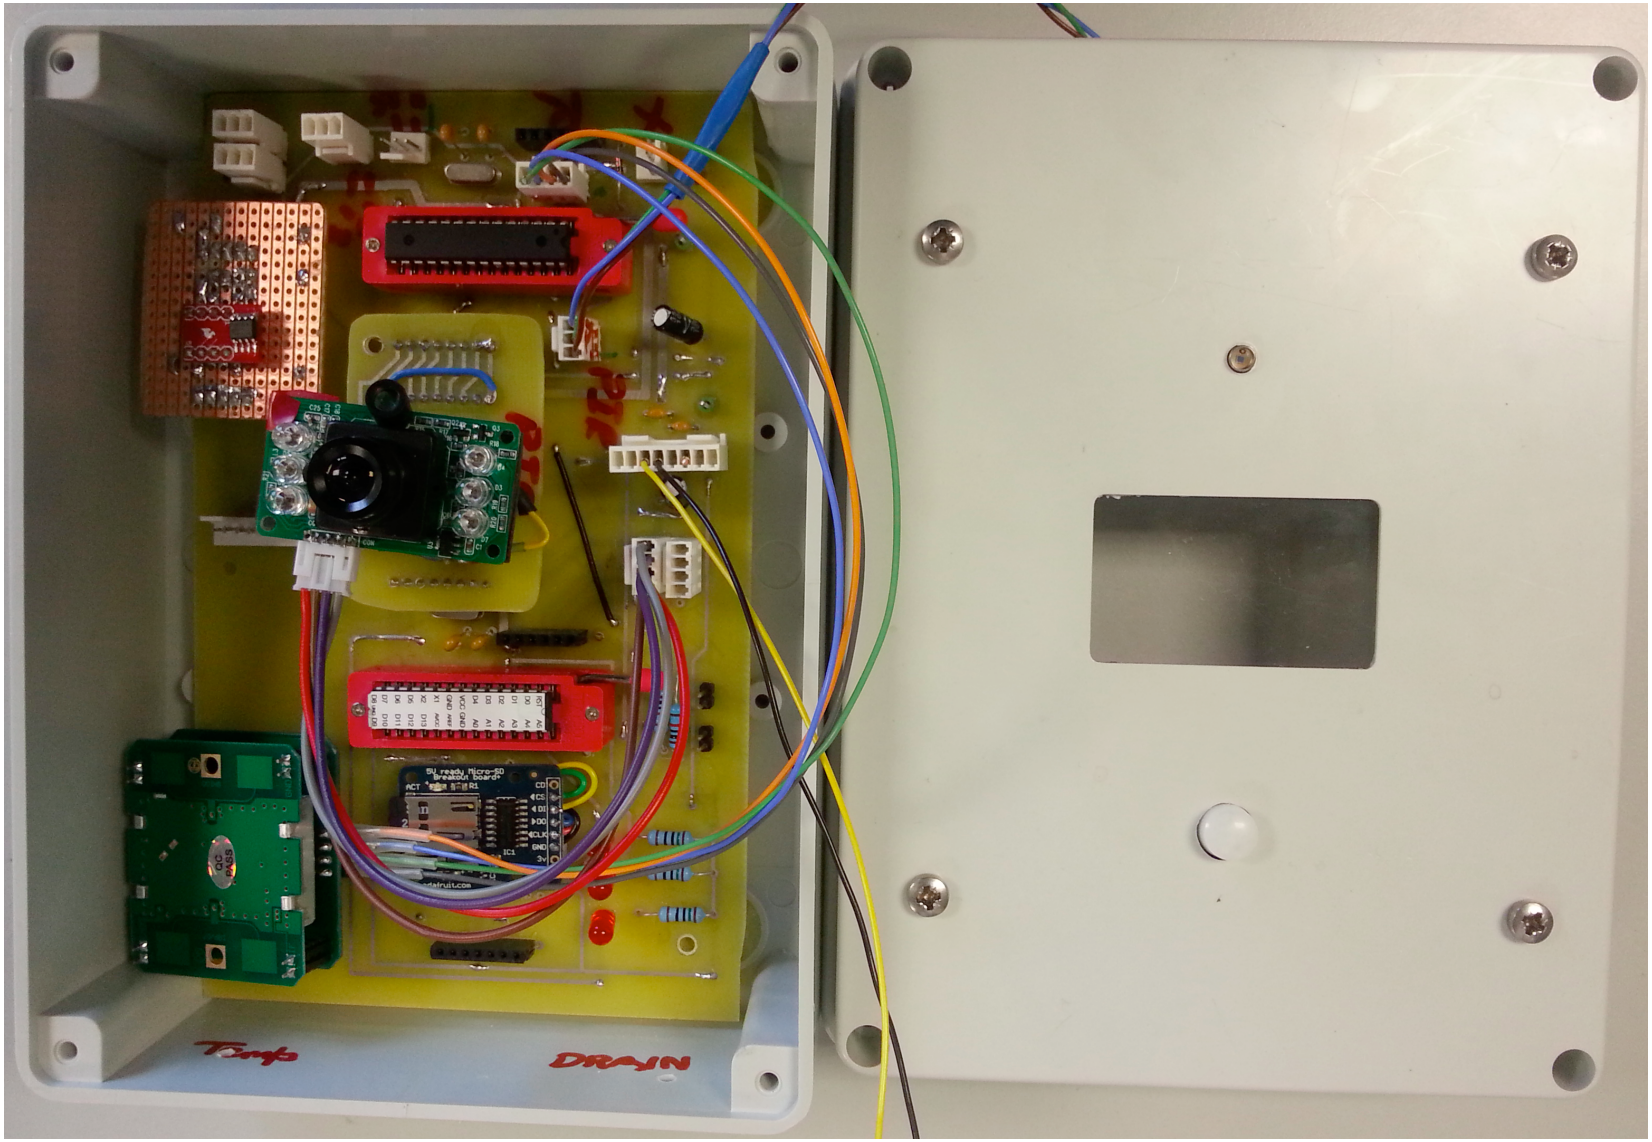
DRAIN

TCP Congestion Control

- Slow Start Phase
- Steady State
- Multiplicative Reduction to Loss

HTTP Video Workloads

- Bursty
- Terrible Reaction to loss
  - ‘Confused, Timid and Unstable’
- Very long connection life time
- Lots of packets in flight

NewCWV

- TCP Adaptation for Rate Limited Traffic
- Improved window validation
- FreeBSD Implementation here:
  - https://bugs.freebsd.org/bugzilla/show_bug.cgi?id=191520

TCP Congestion Window

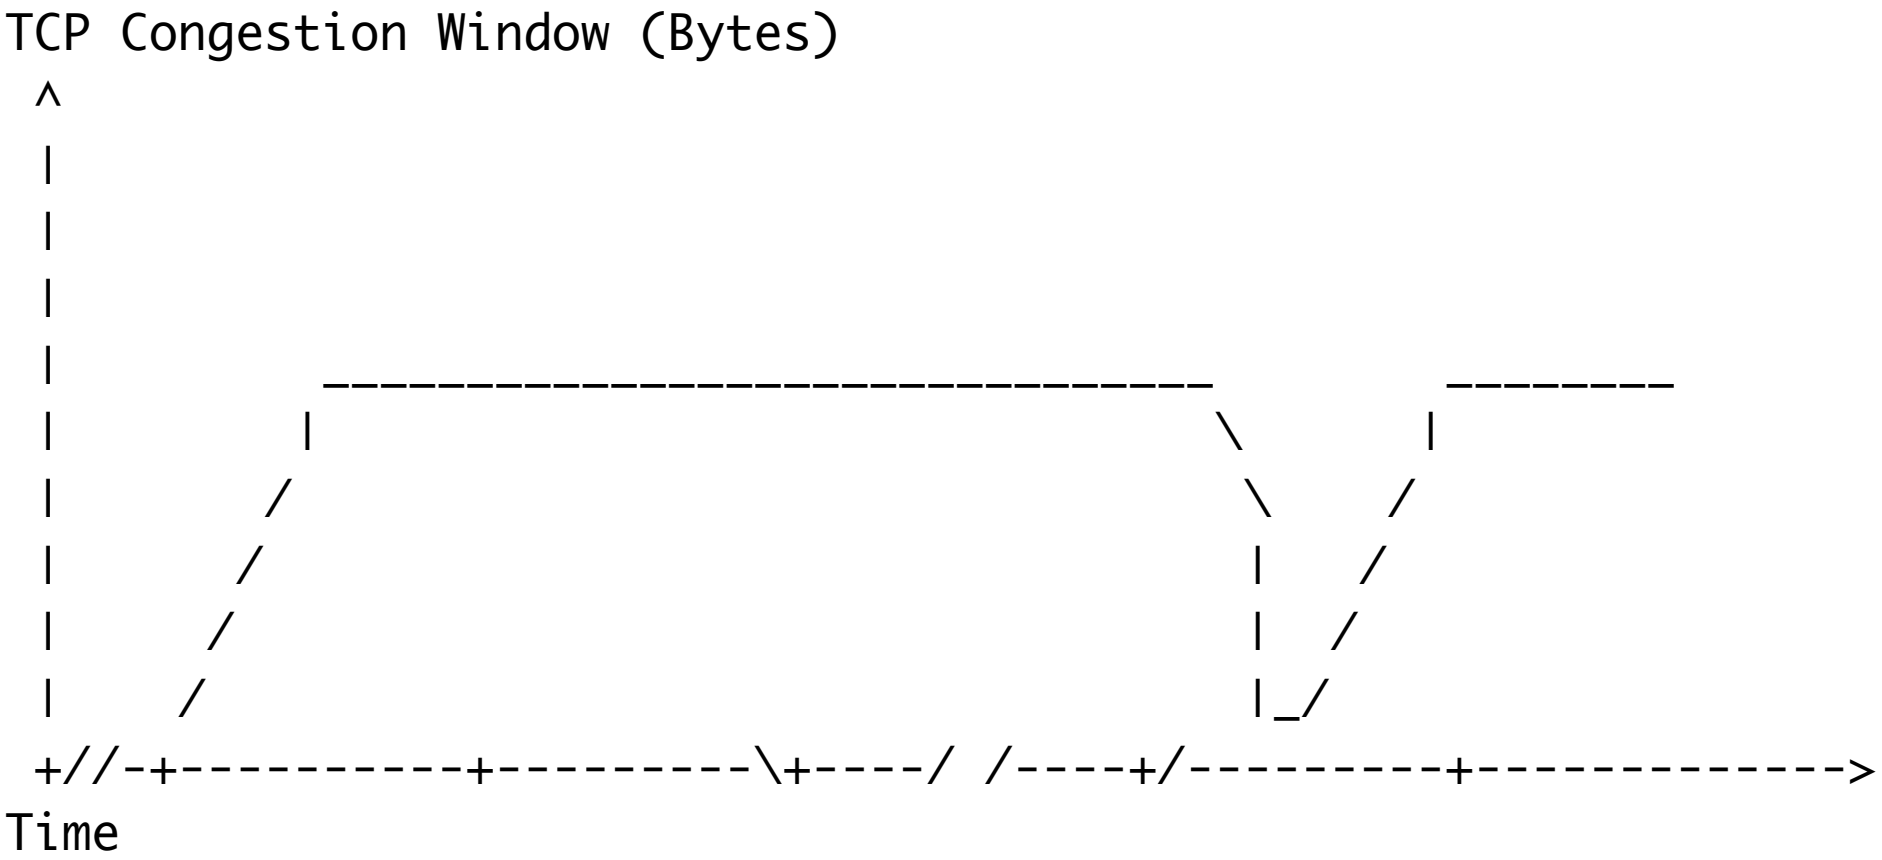

Pipe ack

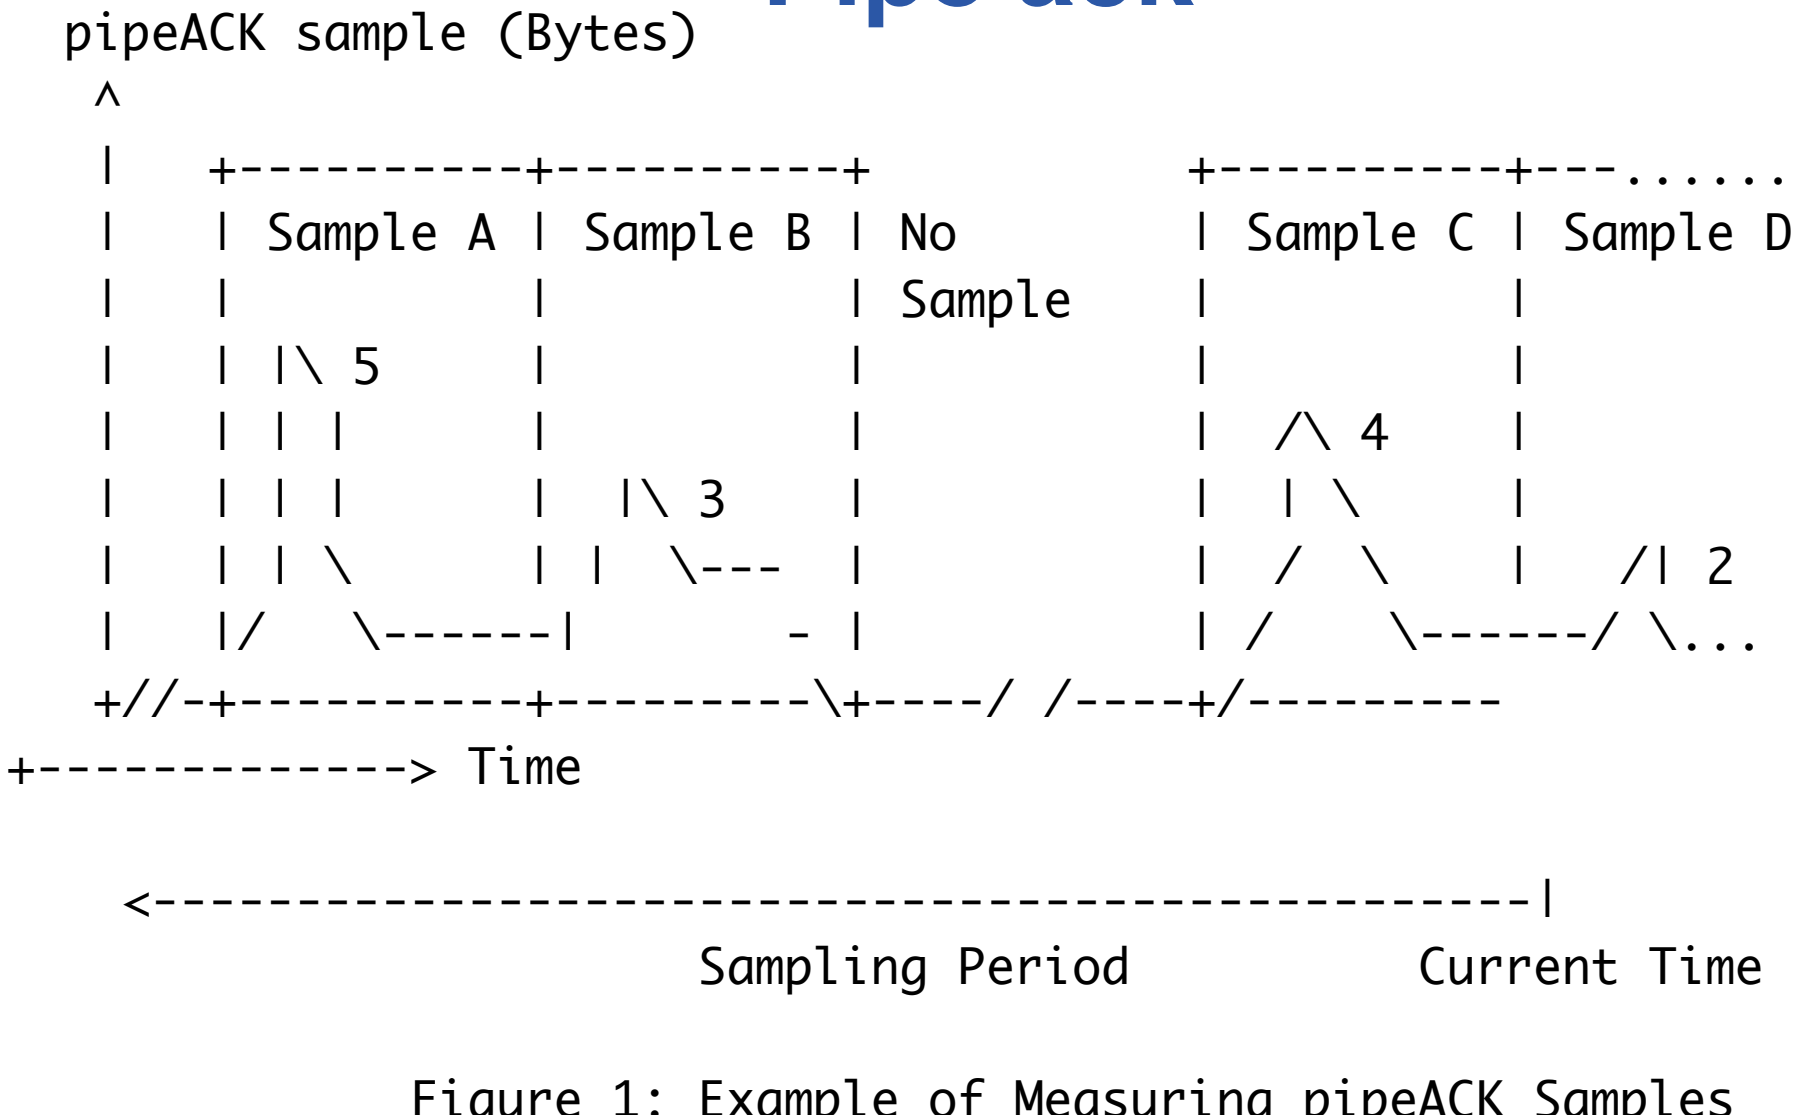

Figure 1: Example of Measuring pipeACK Samples

NewCWV Status

Hiren Panchasara 2016-12-22 20:22:46 UTC Comment 5

This work is supposed to land from Netflix in coming months.

Internet Protocol Datagram

RFC791

Source

Destination

Version ☐ *If other than version 4, attach form RFC 2460.*

Type of Service
- ☐ high reliability
- ☐ high throughput
- ☐ low delay

Protocol
- ☐ TCP
- ☐ UDP
- ☐ Other _____

Precedence
- ☐ Routine
- ☐ Priority
- ☐ Immediate
- ☐ Flash
- ☐ Flash Override
- ☐ CRITIC/ECP
- ☐ Internetwork Control
- ☐ Network Control

Fragmentation **Offset**

Transport layer use only

- ☐ more to follow
- ☐ do not fragment
- ☐ this bit intentionally left blank

Identifier _____

Length **Header Length**

Time to Live **Options**

Do not write in this space.

Header Checksum

Data

Print legibly and press hard. You are making up to 255 copies.

for more info, check IPv4 specifications at http://www.ietf.org/rfc/rfc0791.txt

The Naive View of the Network

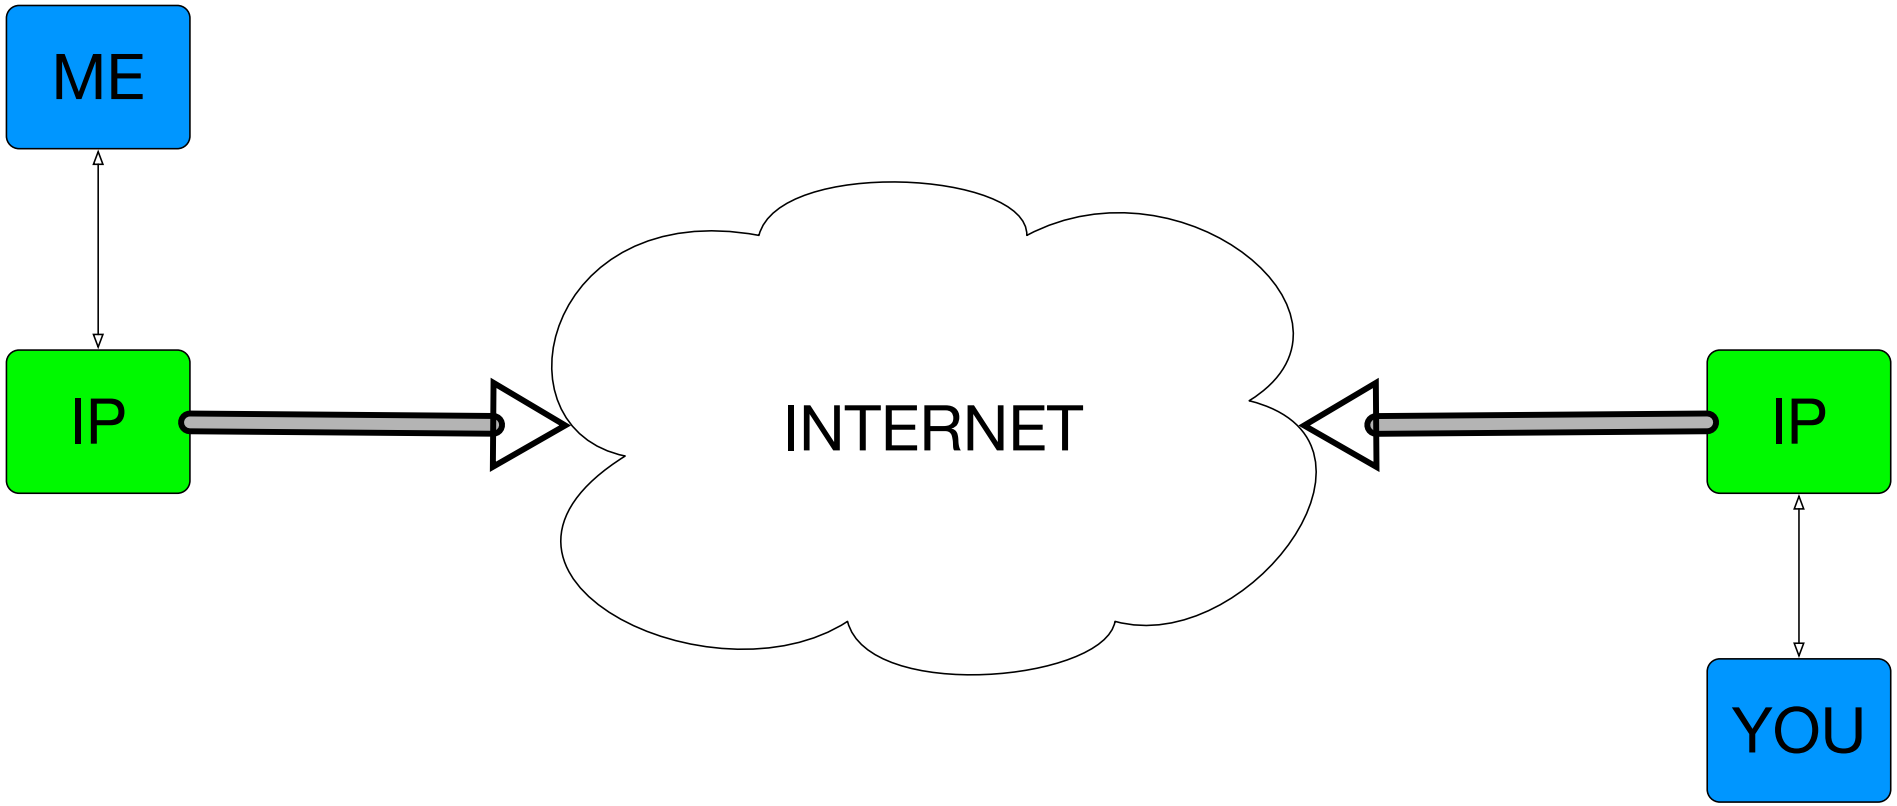

The Reality of the Network

Middleboxes drop non http traffic
Firewalls Block protocols
Security devices inject ‘bad’ data
Proxies break TLS

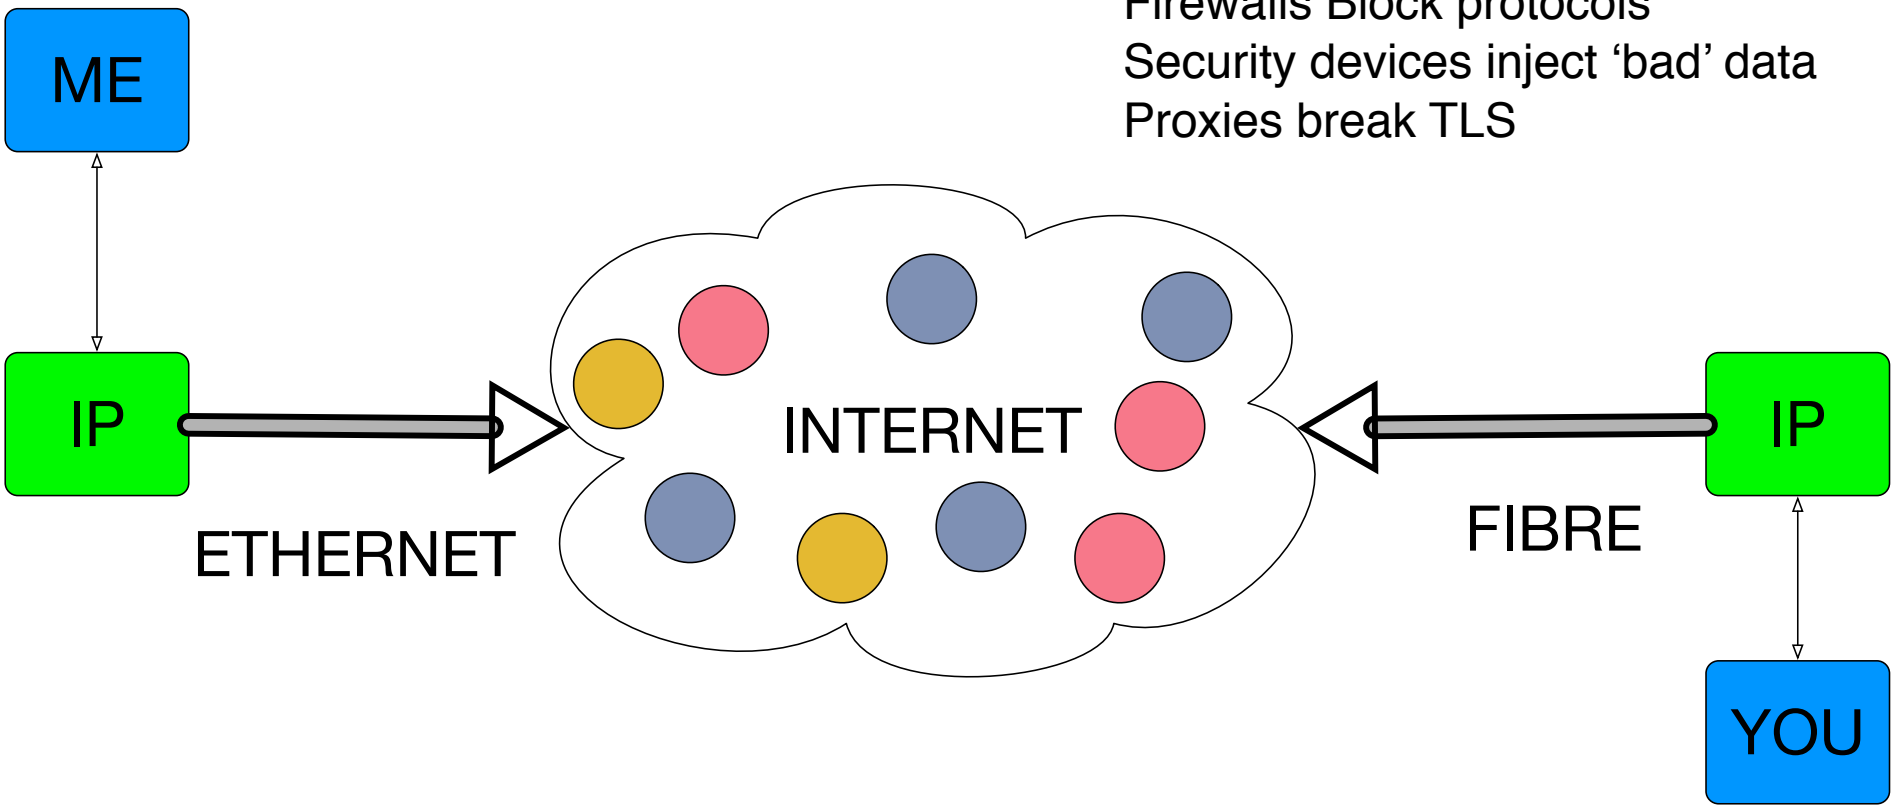

Networks vary a ton

Middleboxes drop non http traffic
Firewalls Block protocols
Security devices inject ‘bad’ data
Proxies break TLS

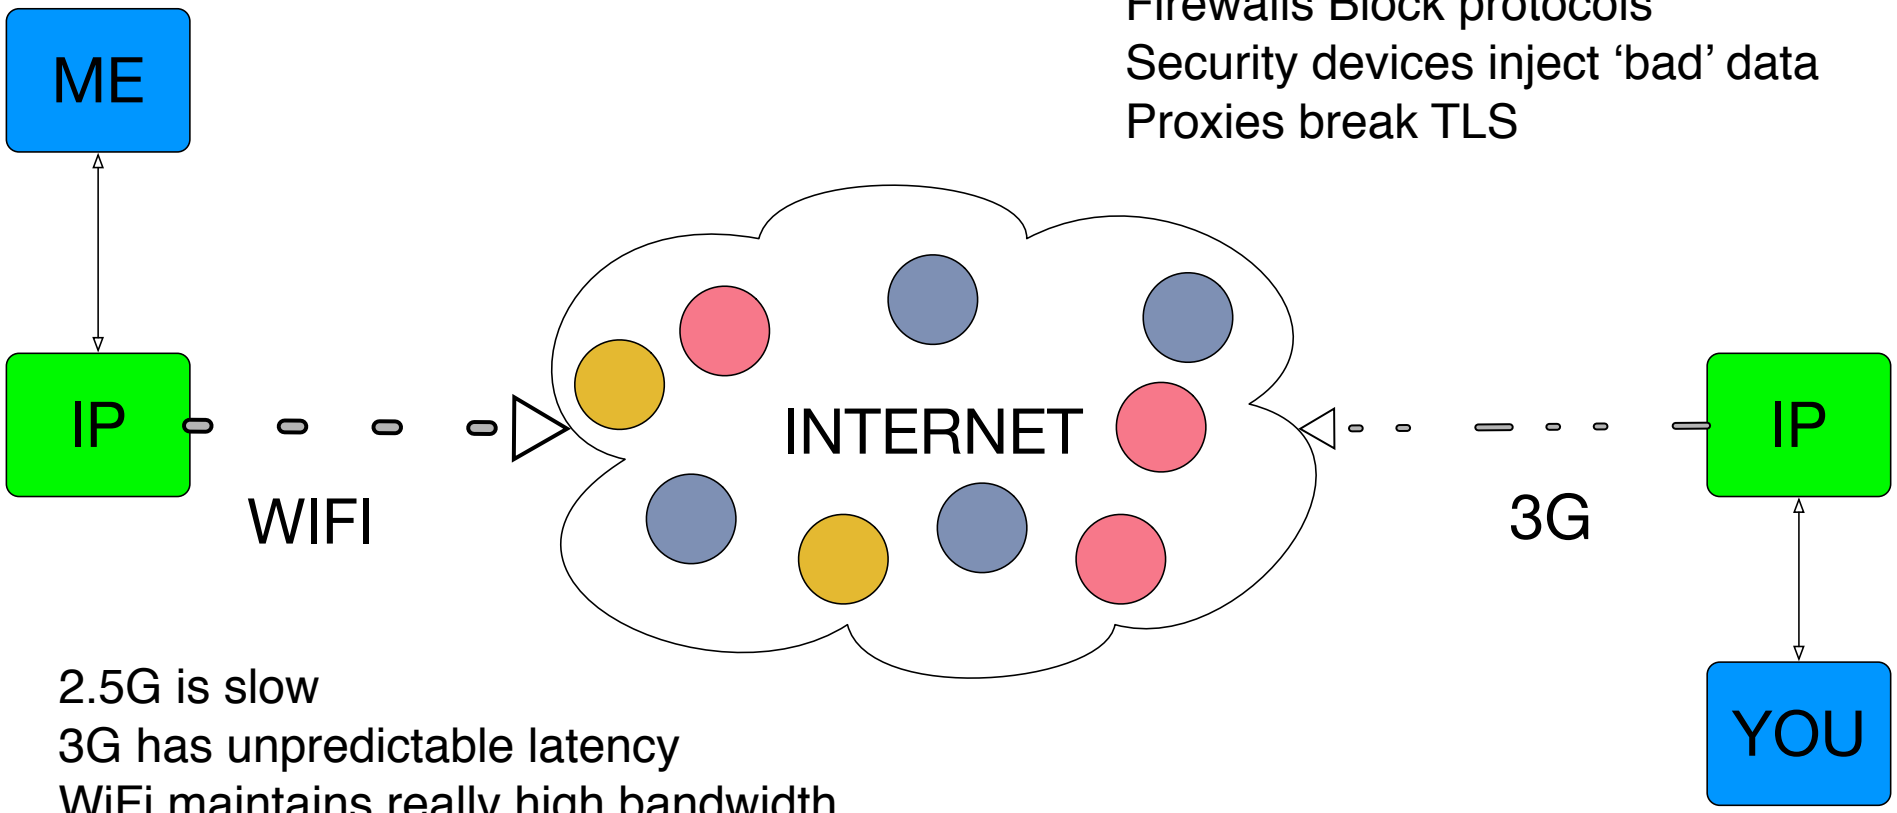

2.5G is slow
3G has unpredictable latency
WiFi maintains really high bandwidth
4G has high, but variability in bandwidth

The socket API has ossified

```
getaddrinfo();   // Look up host
socket();        // Create a socket

setsockopt();    // Configure the socket
getsockopt();    // Check parameters

connect();       // Start connection

send();
recv();
```


Do it in userspace

- usersctp
- QUIC
  - draft-ietf-quic-transport-01
- GUE
  - draft-ietf-nvo3-gue-05
- TCP over UDP

The NEAT System

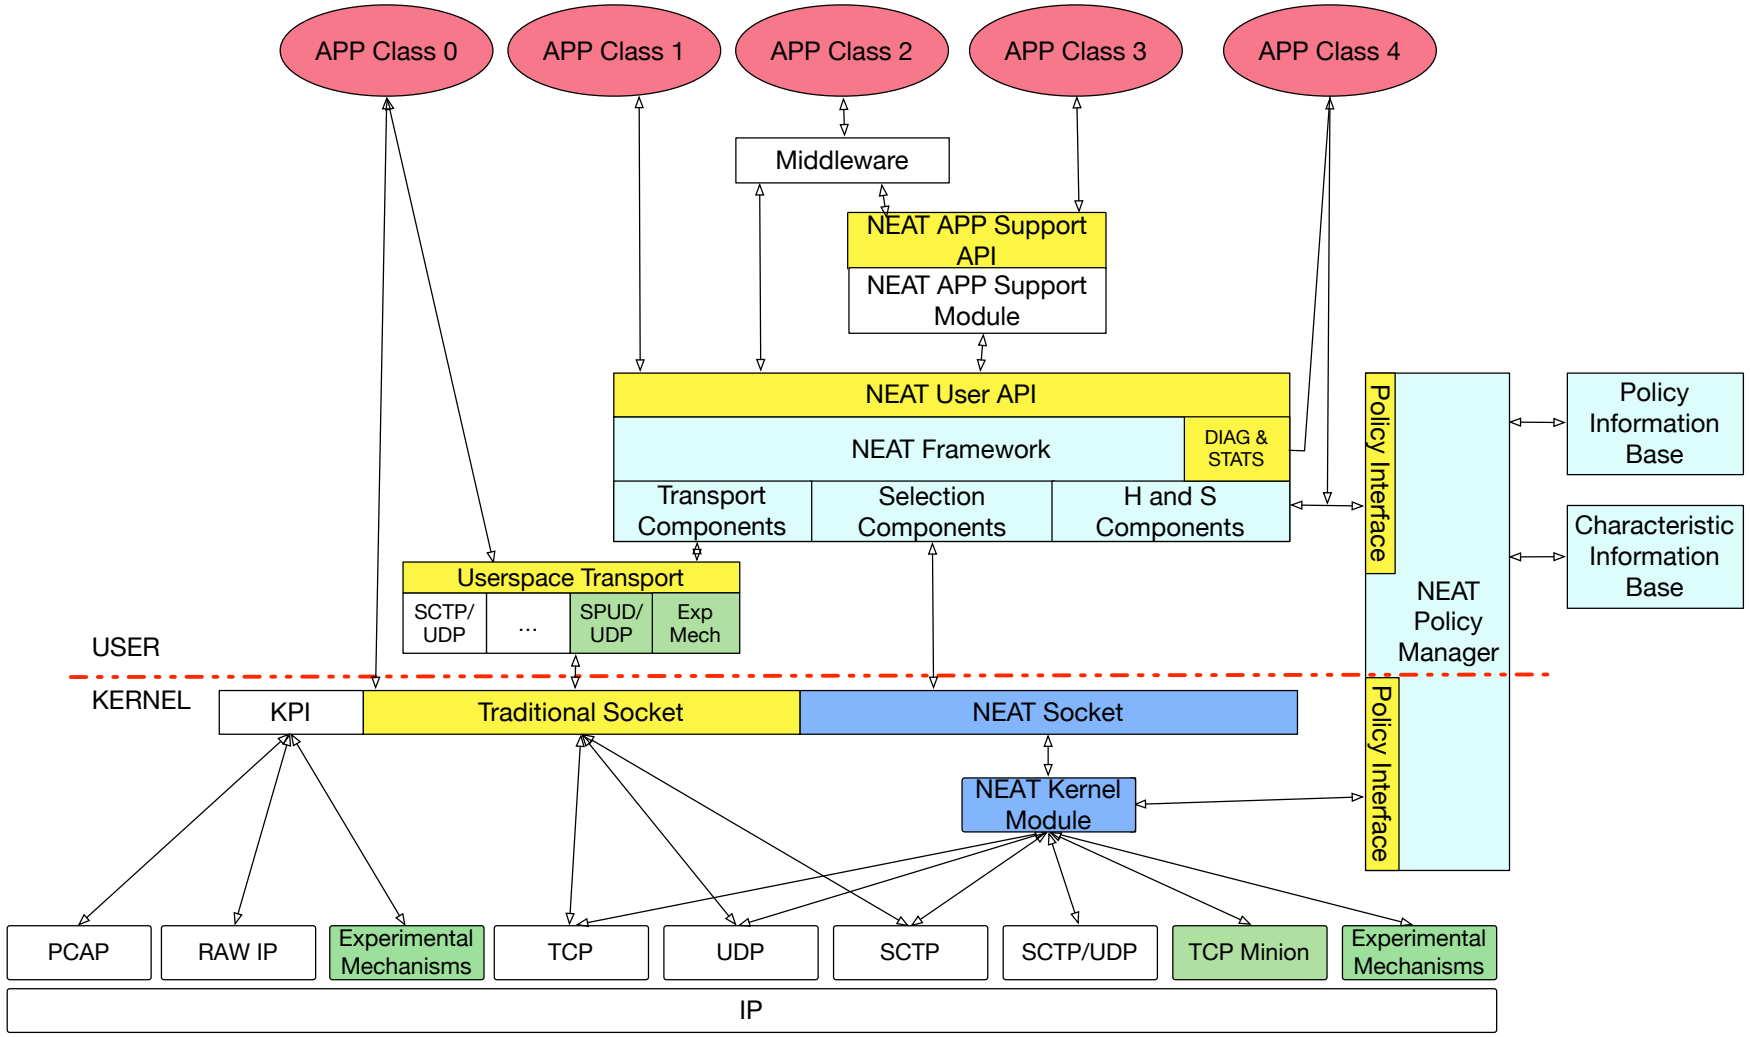

NEAT Application

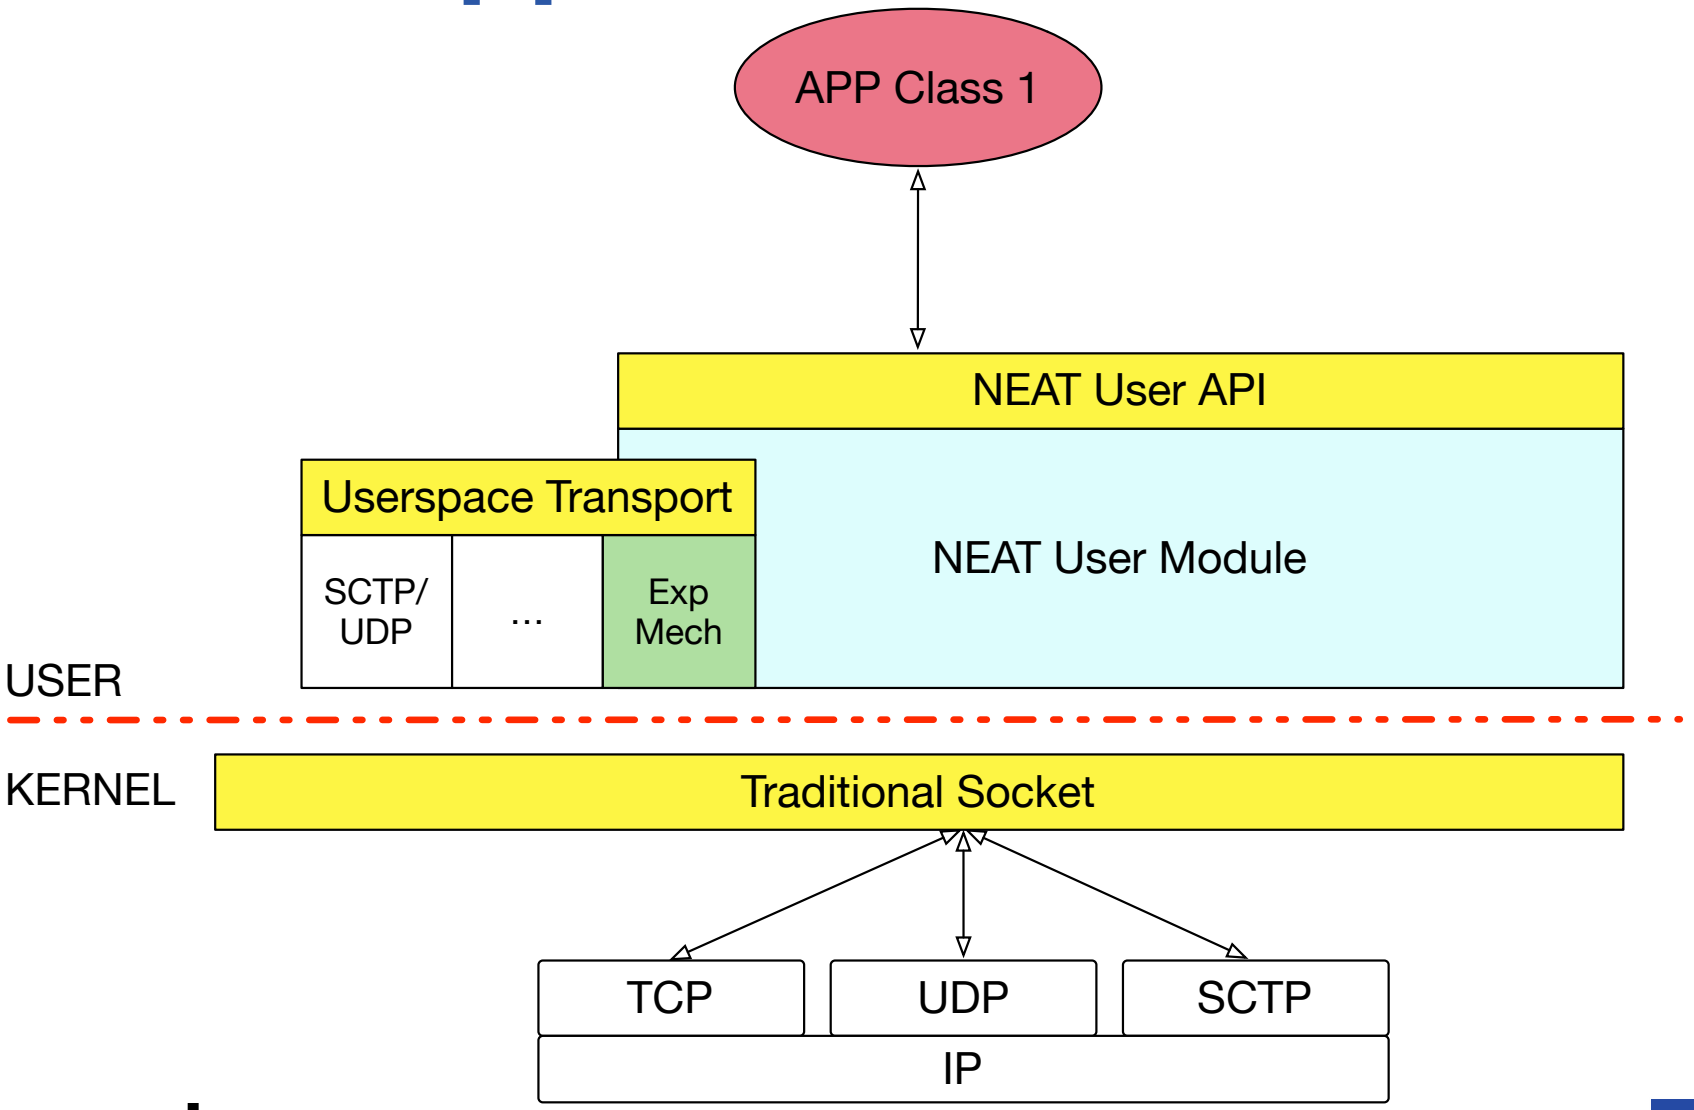

The NEAT User Module

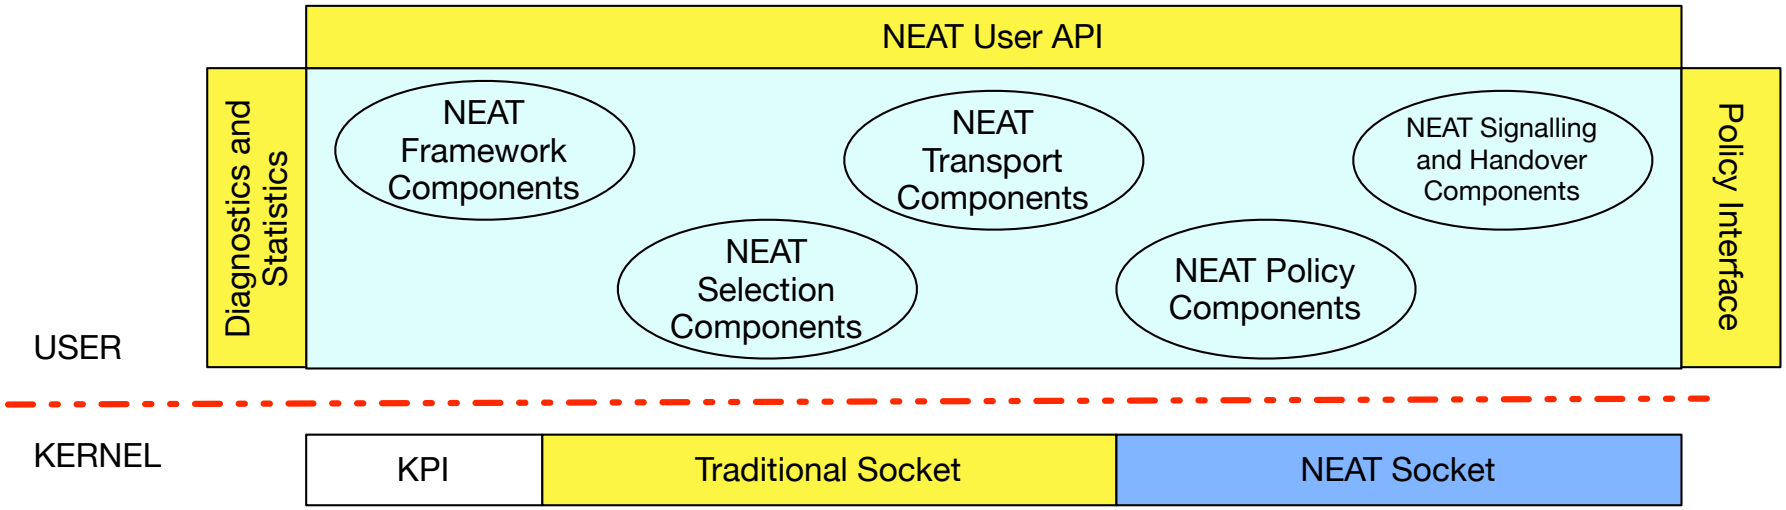

The NEAT User Module

NEAT Flow Endpoint

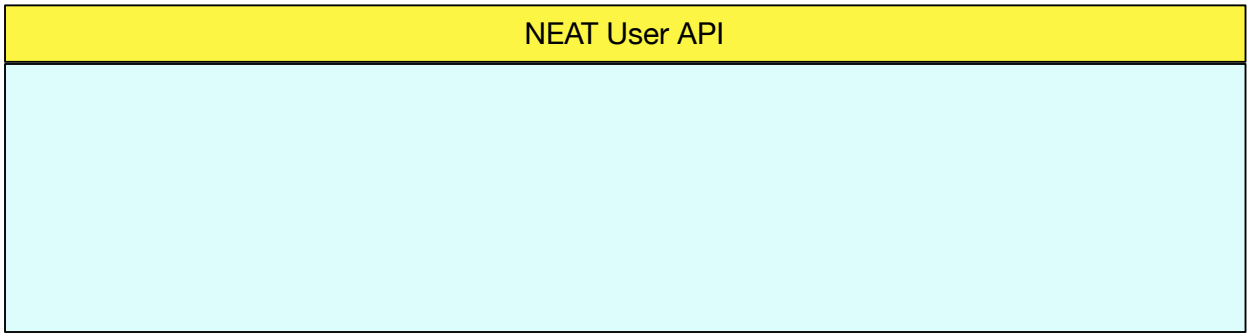

The NEAT User Module

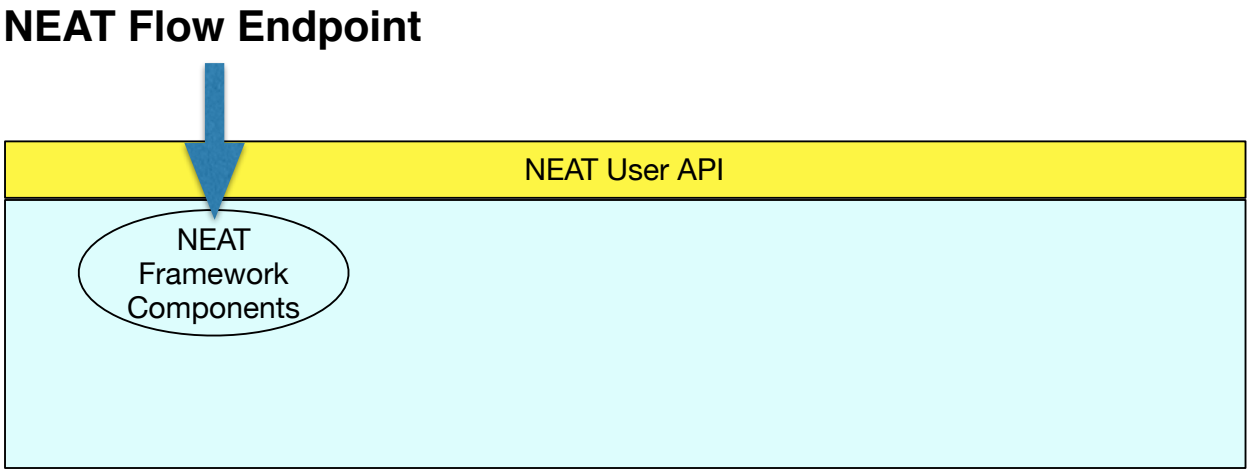

5 Groups of components:

The NEAT User Module

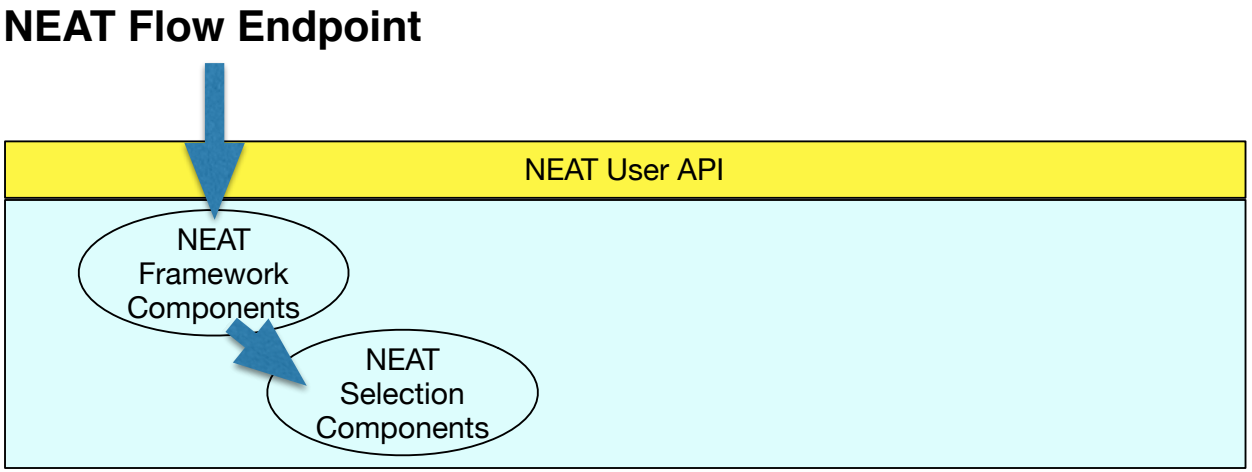

5 Groups of components:

- NEAT Framework Component: API, Logic

The NEAT User Module

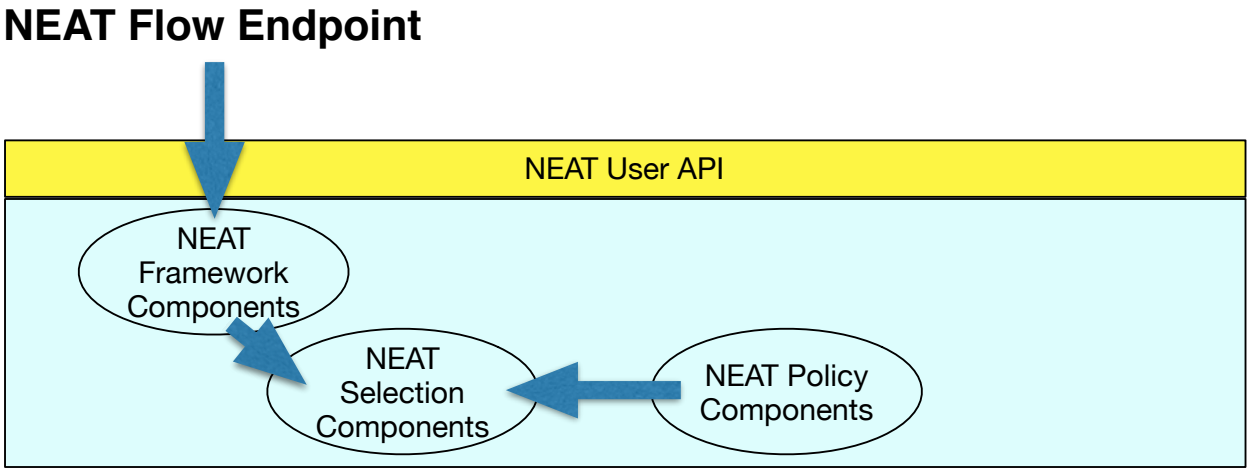

5 Groups of components:

- NEAT Framework Component: API, Logic
- NEAT Selection Components: Choose candidates

The NEAT User Module

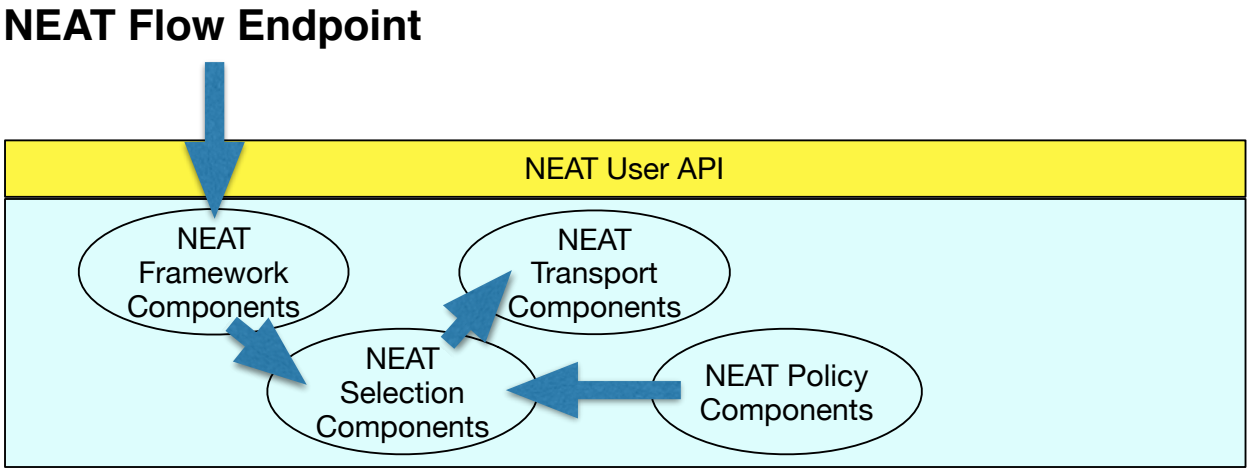

5 Groups of components:

- NEAT Framework Component: API, Logic
- NEAT Selection Components: Choose candidates
- NEAT Policy Components: Policy and Characteristics

The NEAT User Module

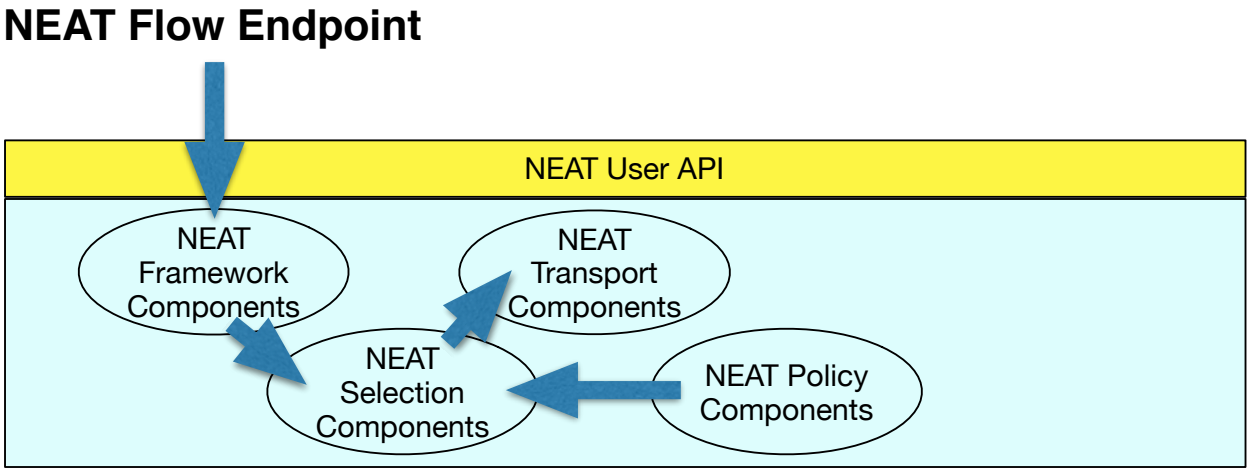

5 Groups of components:

- NEAT Framework Component: API, Logic
- NEAT Selection Components: Choose candidates
- NEAT Policy Components: Policy and Characteristics
- NEAT Transport Components: Instantiate transports

The NEAT User Module

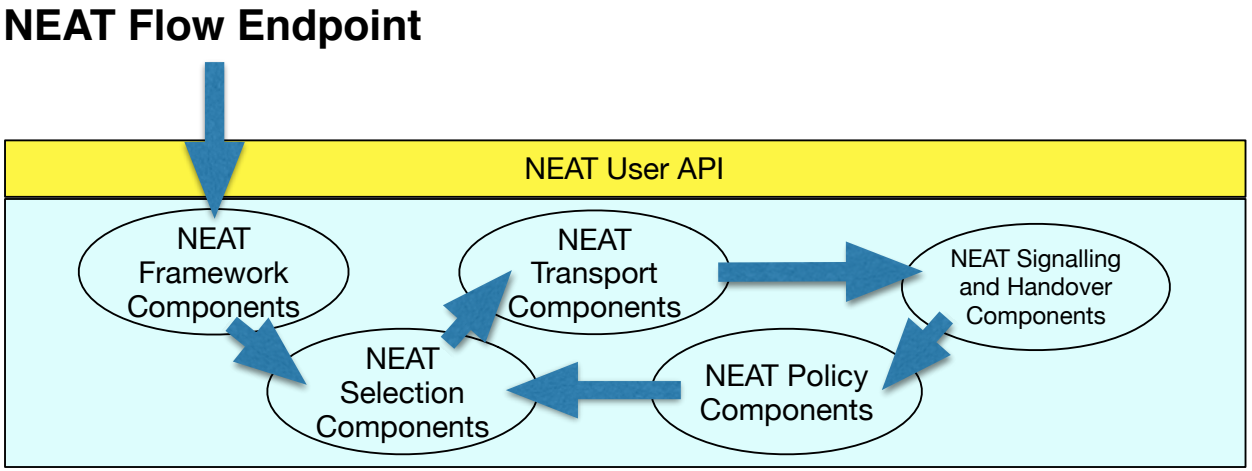

5 Groups of components:

- NEAT Framework Component: API, Logic
- NEAT Selection Components: Choose candidates
- NEAT Policy Components: Policy and Characteristics
- NEAT Transport Components: Instantiate transports
- NEAT Handover and Signalling Components

The NEAT User Module

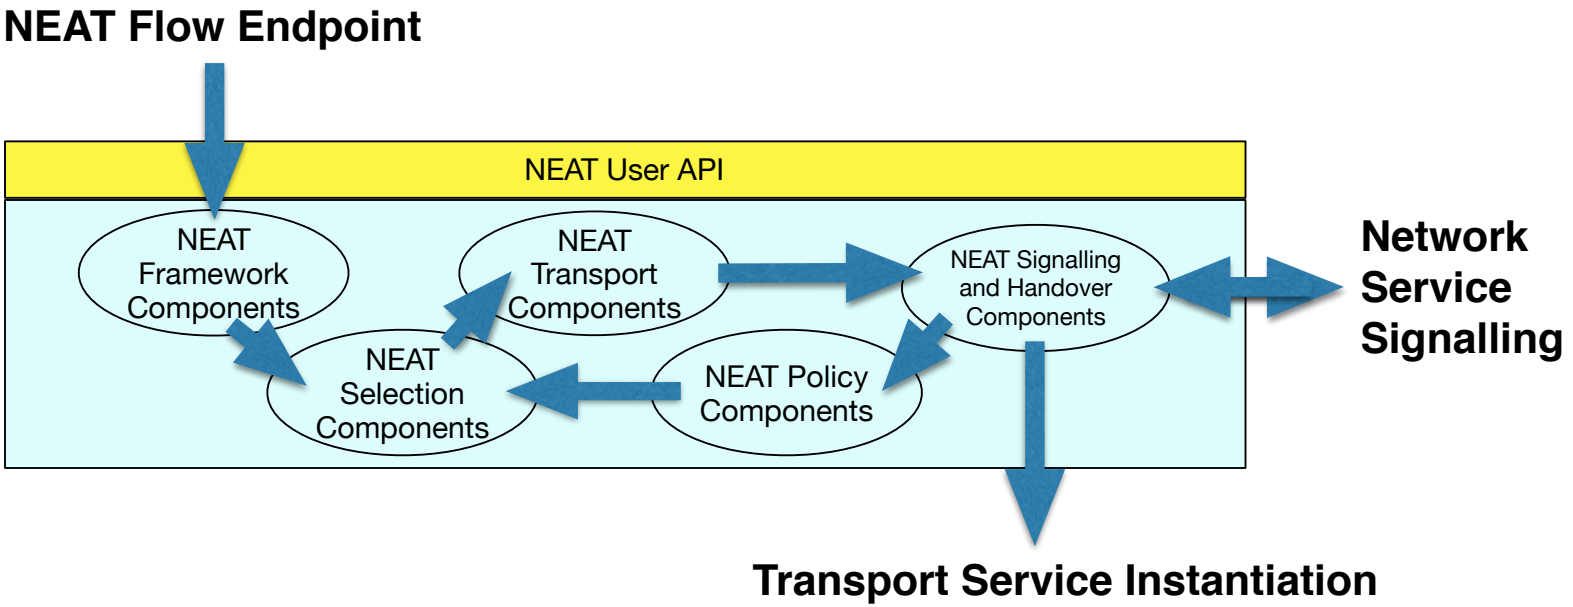

5 Groups of components:

- NEAT Framework Component: API, Logic
- NEAT Selection Components: Choose candidates
- NEAT Policy Components: Policy and Characteristics
- NEAT Transport Components: Instantiate transports
- NEAT Handover and Signalling Components

NEAT Application

```
static struct neat_flow_operations ops;
static struct neat_ctx *ctx = NULL;
static struct neat_flow *flow = NULL;

ctx = neat_init_ctx()
flow = neat_new_flow(ctx)

prop = NEAT_PROPERTY_UDP_REQUIRED | NEAT_PROPERTY_IPV6_REQUIRED;
neat_set_property(ctx, flow, &prop)

ops.on_writable = on_writable;
ops.on_readable = on_readable;
ops.on_error = on_error;

neat_set_operations(ctx, flow, &ops)
neat_open(ctx, flow, argv[argc - 2], argv[argc - 1])

neat_start_event_loop(ctx, NEAT_RUN_DEFAULT);
```

NEAT Application

```
static neat_error_code
on_writable(struct neat_flow_operations *opCB)
{
    neat_write(opCB->ctx, opCB->flow, buf)
    return NEAT_OK;
}

static neat_error_code
on_readable(struct neat_flow_operations *opCB)
{
    neat_read(opCB->ctx, opCB->flow, buf)
    return NEAT_OK;
}
```

https://github.com/NEAT-project/neat/blob/master/examples/client.c

Porting Apps

- Firefox
- rsync

buildbot

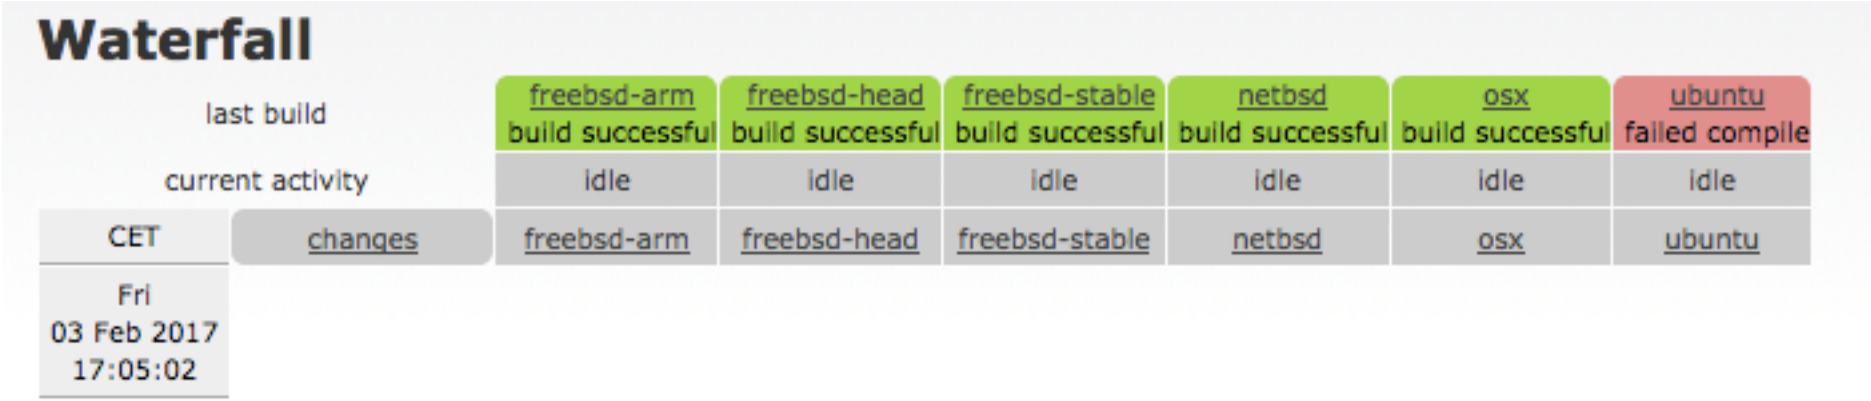

Waterfall

		freebsd-arm	freebsd-head	freebsd-stable	netbsd	osx	ubuntu
last build		build successful	build successful	build successful	build successful	build successful	failed compile
current activity		idle	idle	idle	idle	idle	idle
CET	changes	freebsd-arm	freebsd-head	freebsd-stable	netbsd	osx	ubuntu
Fri 03 Feb 2017 17:05:02							

neat

https://www.neat-project.org

https://github.com/neat-project/neat

NEAT is funded by the European Union's Horizon 2020 research and innovation programme under grant agreement no. 644334.

Horizon 2020
European Union funding
for Research & Innovation