

# UNIMON

Efficient Monitoring of Unikernel-Based VNFs

Will Fantom

[<w.fantom@lancaster.ac.uk>](mailto:w.fantom@lancaster.ac.uk)

PhD Student

Lancaster University




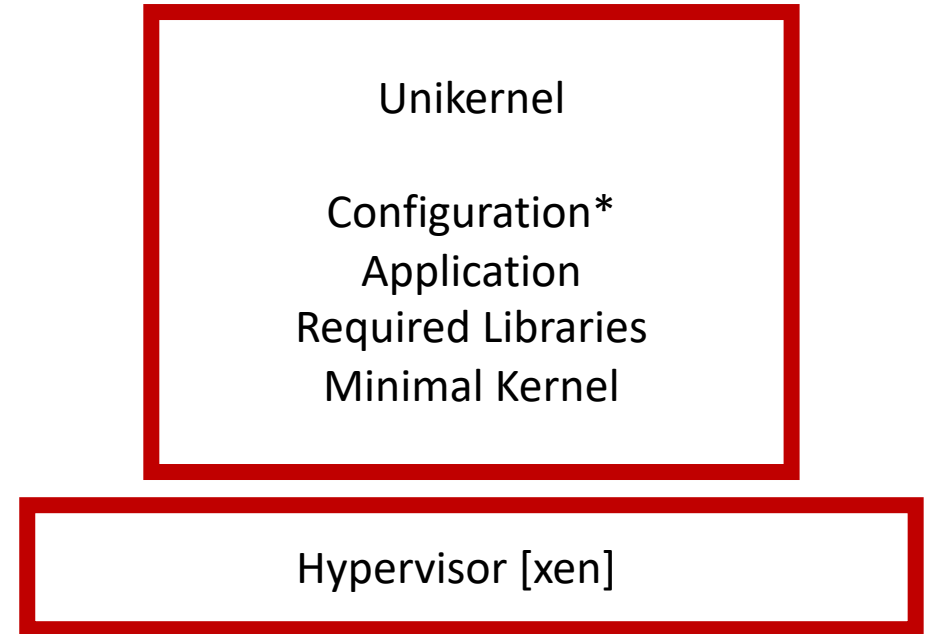
# Unikernels

## Pros:

- Small Image Size
- No Context Switching
- Small Attack Surface

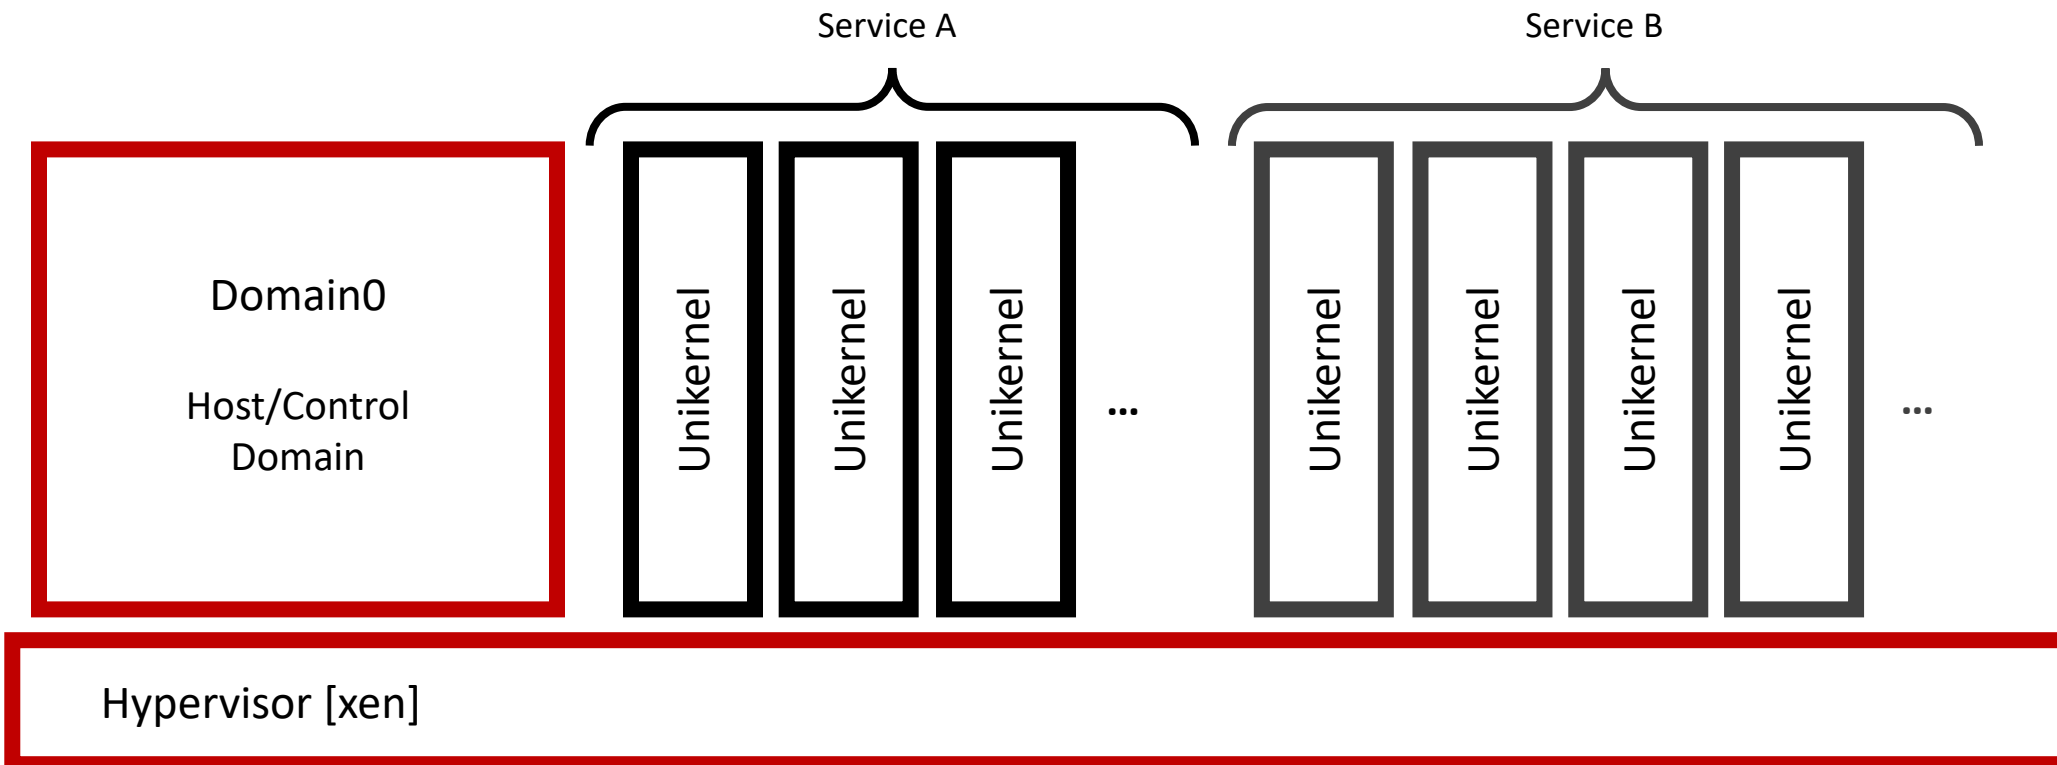
## Cons:

- Difficult to Develop
- Harder to Manage



# Scaling

- Low resource cost of a single Unikernel allows for vast horizontal scaling



(NFV+DevOps)= 

- The marriage of microservices and CI/CD as seen in SaaS
- Unikernels can  $\cong$  Microservices for NFV in the context of CI/CD
- Can use the DevOps methodology with Unikernels

# Challenges

- Poll-based packet I/O uses all free cycles to poll
  - 'Black-box' Unikernel always appears to use 100% CPU
- Scaling decisions made by IM/MANO platform
  - Hardware statistics w/o context can be invalid
  - Limited interface options [no ssh]
- Single-core + cooperative scheduling [with ClickOS<sup>[1]</sup> example]
  - Guarantees in-pipeline monitoring to have impact
- Difficult development environment
  - Leads to issues to be discovered post deployment

# UNIMON Goals

- Gain contextual data regarding the state of a Unikernel based VNF
- Minimize the impact of in-pipeline VNF monitoring
- Offer flexibility to both developers and operators
- Push stats out of the VM with minimal overhead

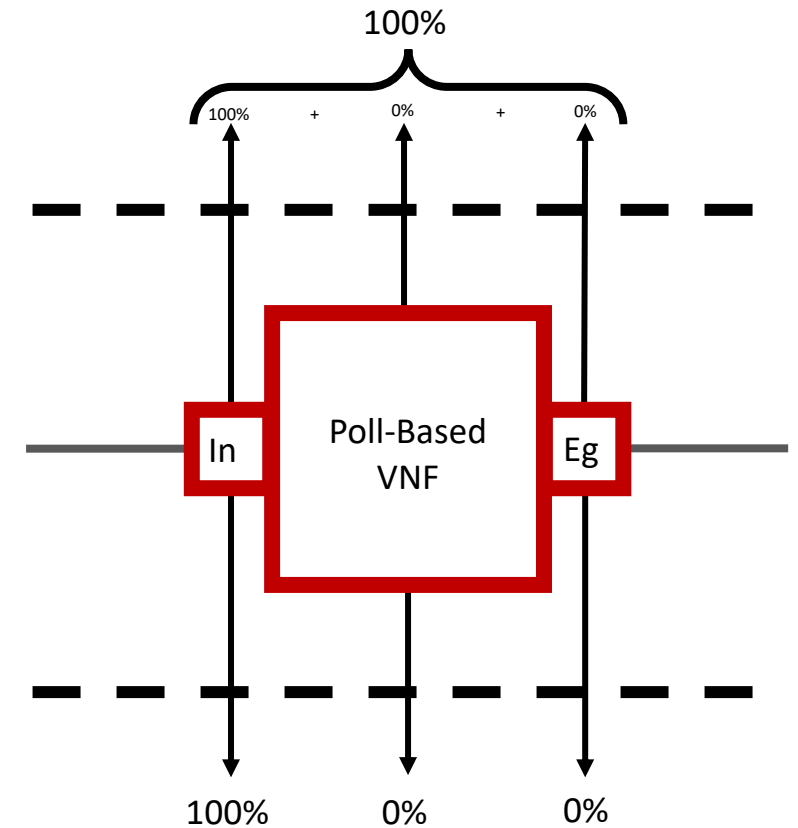
# UNIMON:

## Getting contextual data

- Deep external monitoring not suited for live VNFs. Uniprof<sup>[2]</sup> is an example and requires pausing VMs and only provides CPU usage when combined with unstripped binary.
- UNIMON opts for internal:
  - Can be used in live deployment scenarios
  - Comes inside the binary
  - Offers contextual data

NAME	STATE	CPU(sec)	CPU(%)	MEM(k)	MEM(%)
clickos	[m -----r	29	100.0	524288	1.6

CPU: As seen by IM or MANO platform

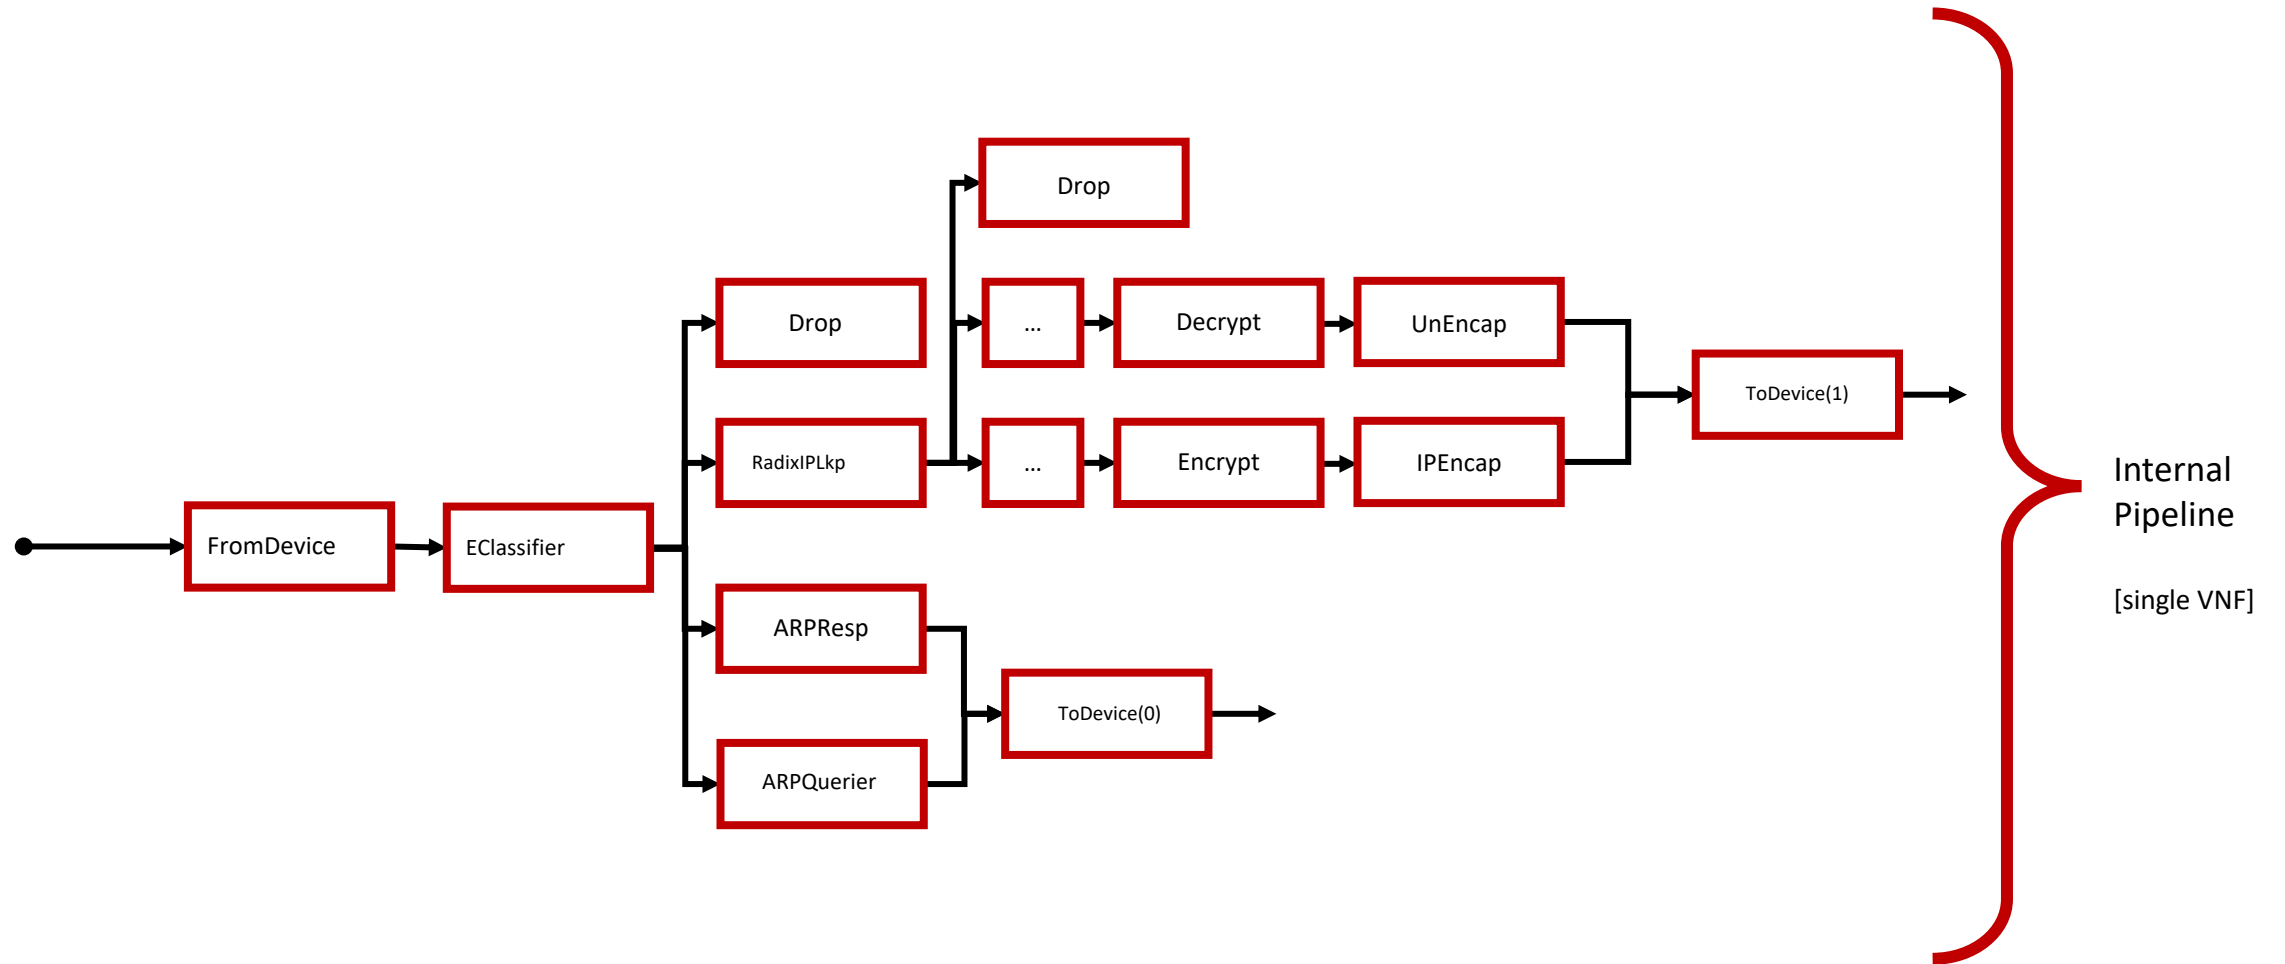


CPU: Split Values

[2] Schmidt, Florian. (2017). uniprof: A Unikernel Stack Profiler.

# UNIMON:

## ClickOS IPsec Element Pipeline





# UNIMON:

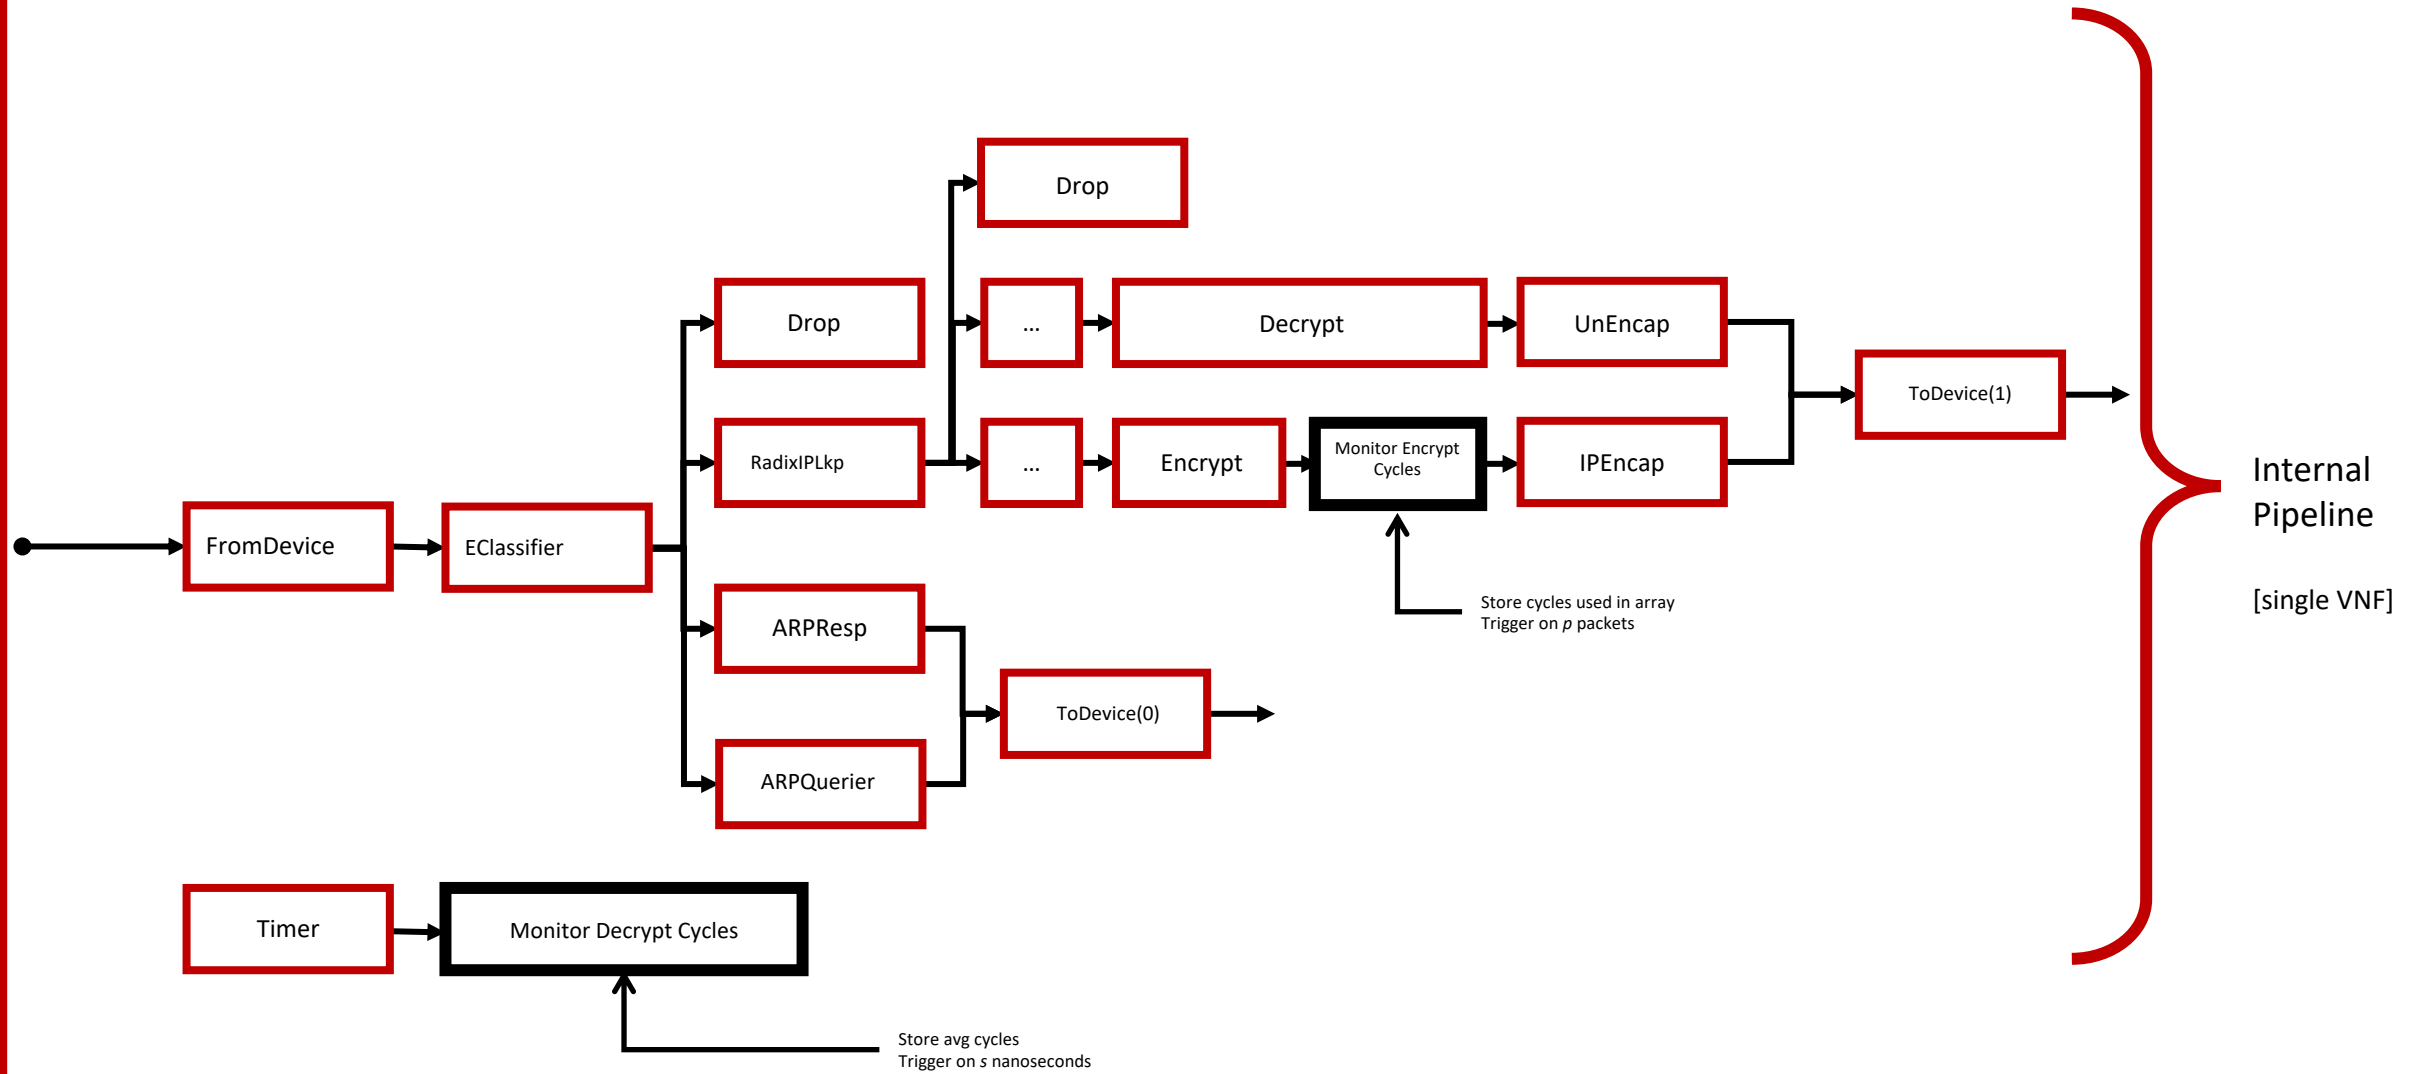
## In-pipeline Monitoring

Attempts to reduce performance impact:

- Don't monitor the whole system, only suspect 'elements'
- Provide flexibility to both developers and operators of Unikernel VNFs
  - Per monitored element poll frequency
  - Separated frequency for pushing out collected stats
  - Wide range of values that can be monitored whilst keeping context

# UNIMON:

## Pipeline with Monitoring



# UNIMON:

## Zero-Copy Stats Transfer

```
Unimon Shared Struct
    uint16_t tmc;
tagged_mon_t tagged_monitor[tmc];
mon_data_t mon_data[tmc];
```

```
Tagged Monitor Struct
    uint8_t status;
char element_name[32];
char handler_name[32];
    uint32_t data_size;
    uint32_t data_offset;
```

```
Monitor Data Struct
    uint64_t data[n];
```

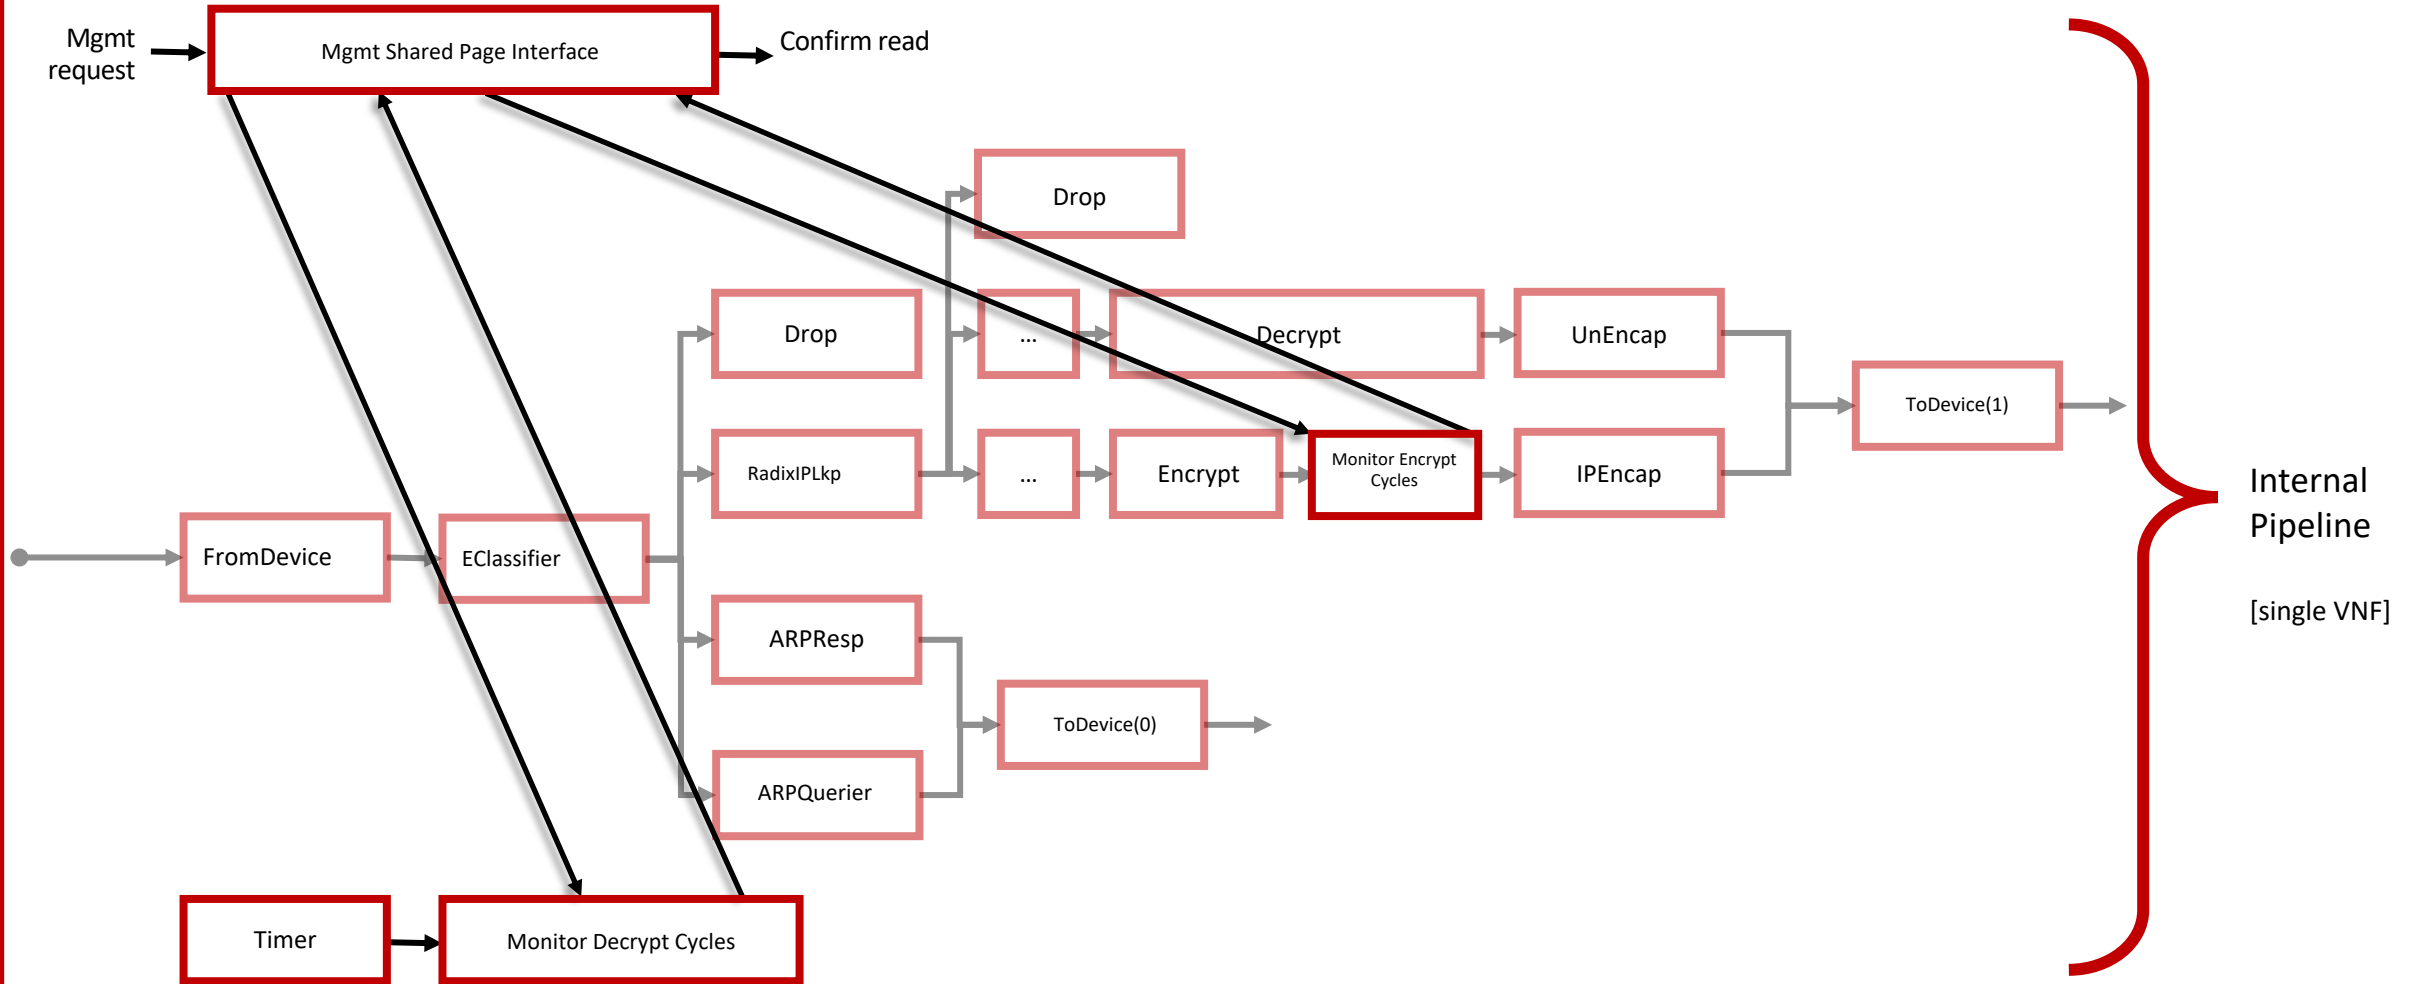
### Optimizations to the flow of monitoring data:

- Separate frequencies for monitoring frequency and data push outs (leaving memory decisions up to the dev)
- Configurable average vs. time series array
- Using shared pages between the host domain and the VNF to allow for zero copy stats transfer (XenSocket<sup>[3]</sup>)

[3] Zhang, Xiaolan & McIntosh, Suzanne & Rohatgi, Pankaj & Linwood Griffin, John. (2007). XenSocket: A High-Throughput Interdomain Transport for Virtual Machines.

# UNIMON:

## Exporting Data



# Testing

## Does UNIMON Work?

```
Avg: 145709907
Avg: 145782630
Avg: 145484148
Avg: 145828617
Avg: 145506150
Avg: 145876485
Avg: 145538511
Avg: 145621572
Avg: 145526115
Avg: 145962381
Avg: 145654845
Avg: 145565616
Avg: 145606869
Avg: 146035269
Avg: 145009377
Avg: 146110041
Avg: 145739064
Avg: 145672473
Avg: 145351104
Avg: 145590792
Avg: 145843362
Avg: 145747425
Avg: 145534869
Avg: 145567332
Avg: 145568208
Avg: 145533060
Avg: 145500000
```

~/projects/clickos-monitoring/clickos-elements(master\*) » sudo pkt-gen -i vale420:tx -f tx -l 64 -b 128

Output when monitoring CPU cycles of an ‘fromdevice element’ every 0.1seconds

```
Avg: 0
Avg: 0
Avg: 0
Avg: 0
Avg: 0
Avg: 0
Avg: 0
Avg: 0
Avg: 0
Avg: 0
Avg: 0
Avg: 0
Avg: 0
Avg: 0
Avg: 0
Avg: 0
Avg: 0
Avg: 0
Avg: 0
Avg: 0
Avg: 0
Avg: 0
Avg: 0
Avg: 0
Avg: 0
Avg: 0
Avg: 0
Avg: 0
Avg: 0
Avg: 0
Avg: 0
Avg: 0
Avg: 0
Avg: 0
Avg: 0
Avg: 0
Avg: 0
```

~/projects/clickos-monitoring/clickos-elements(master\*) » sudo pkt-gen -i vale420:tx -f tx -l 64 -b 128

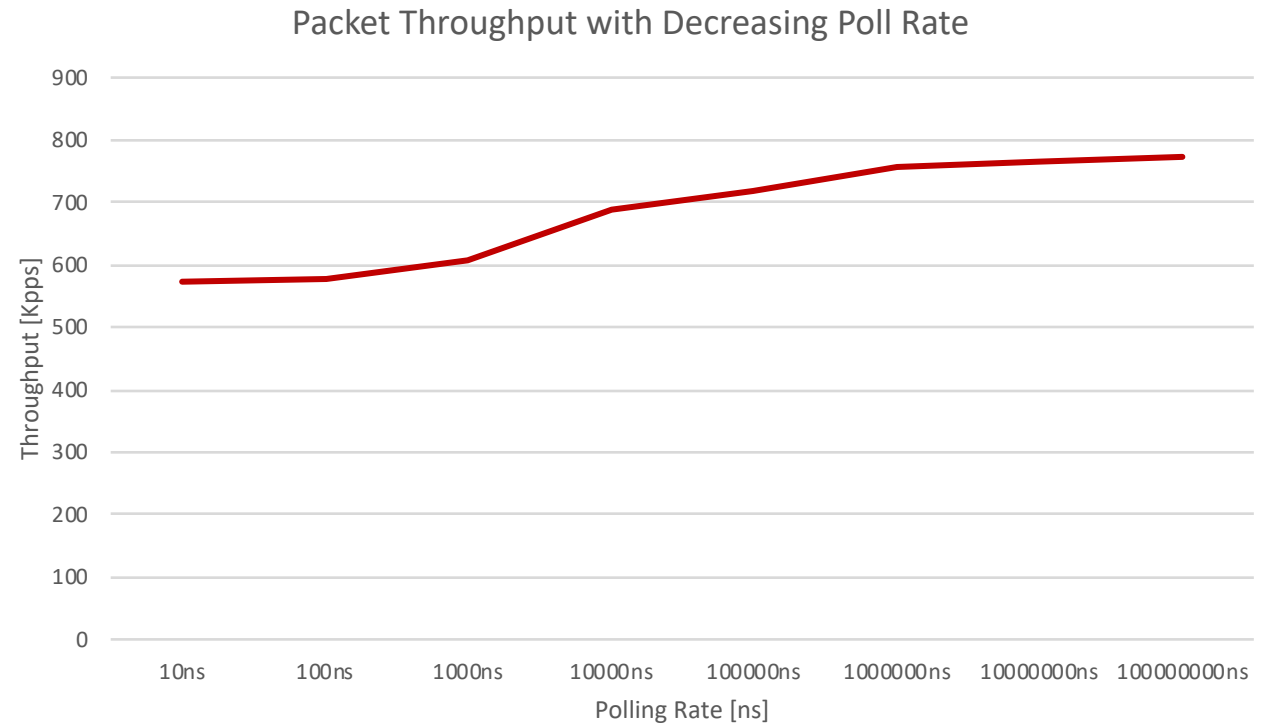
Output when monitoring CPU cycles of an ‘ethermirror element’ every 0.1seconds

# Testing

What's the Impact?

What does this show?

- 3.4% overhead when polling 1000x per second
- 27% overhead when polling every 10 nanoseconds



# Future Work

- Enable host domain stats aggregation  
(Fractal presentation tomorrow covers how we could do a bit of this)
- Support live reconfiguration of VNF monitoring
- Batch monitoring of multiple elements
- Budget focused adaption

# Thanks for Listening

## Any Questions?

Very different UNIMON...



### Unimon

**Unimon** is a Mythical Animal Digimon. It is a composite Digimon that has both the horn of the **Unicorn**, the legendary holy beast, and the wings of the **Pegasus**. With the gigantic wings growing from its back, it instantly runs about the world of the **Computer Network**, stabbing enemies with the sharp horn extending from its brow. A wild(?) Unimon has a rough temperament like a restive horse, but once tamed, it can be handled as if it was the tamer's own limbs.<sup>[2]</sup>

#### Attacks

- **Aerial Attack**<sup>[3]</sup> (*Holy Shot*): Spews a **Qigong** shot from its gigantic mouth.
- **Aerial Gallop**<sup>[4]</sup> (*'corn Thrust*): Stabs the enemy with the horn on its head.
- **Uni-Gallop**: Kicks with its hind legs.

