FIRST Amsterdam – April 2017

+bash-4.3$ echo 'PCAP cant scale’

PCAP cant scale

+bash-4.3$ echo 'PCAP cant scale'| sed 's/cant/does/’

PCAP does scale

+bash-4.3$ whoami

Erik Waher (Security Engineer @ Facebook)

Matt Moran (Security Engineer @ Facebook)

Incident Response Life Cycle

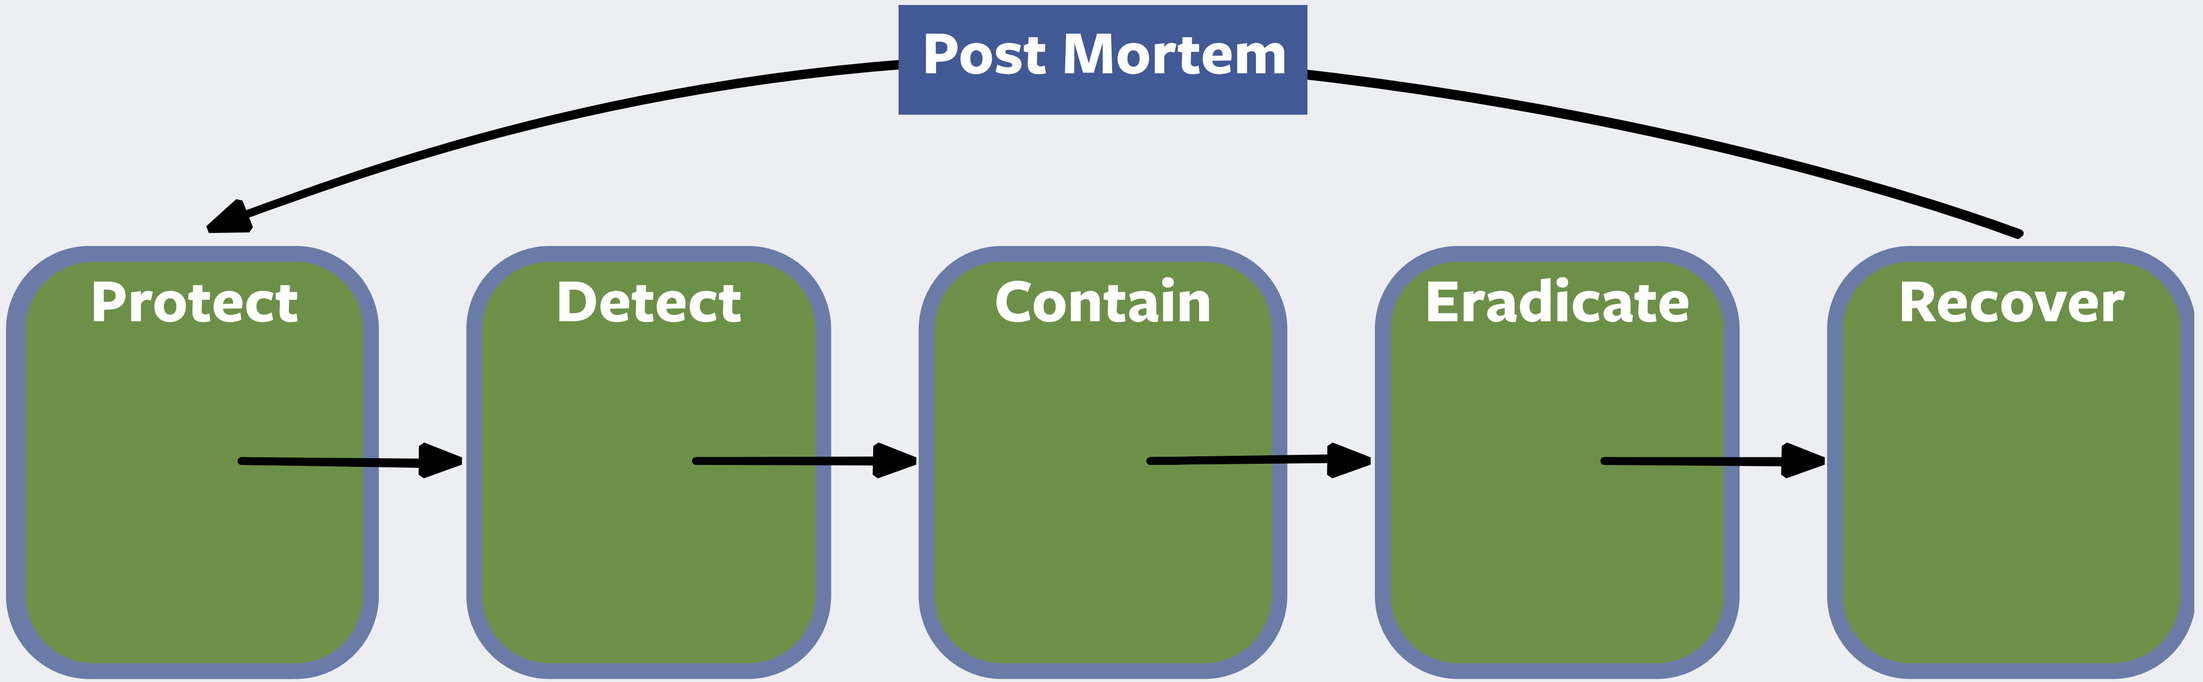

Detection Problem:
IR needs network visibility

Specifically they need indicators of compromise (IOC)

to track down badness…

IOC live in packet captures

Difficulty to Solve: High

- Physical space limitations
- Retention time
- High network throughput
- Commercial $olution$

Solution: build something better

We set out to build:

- Flexible storage size
- Low Cost
- High throughput
- Network accessible storage
- PCAP as a service

We built:

- 1PB – 3.6PB
- OCP hardware
- Speed - 44Gbps per host
- NFS backed by GlusterFS
- PCAP as a service
- Open platform anyone can build

Demo

		2.168.1.77	173.1	
09	13.662945	173.194.33.6	192.168.1.77	TCP
710	13.995895	Actionte_d8:a3:88	Msi_74:82:e6	ARP
711	13.995922	Msi_74:82:e6	Actionte_d8:a3:88	ARP
712	15.030559	fe80::bdca:e67b:5eb7:8	ff02::c	SSDP
13	15.058140	192.168.1.76	239.255.255.250	UDP
714	15.123002	192.168.1.74	239.255.255.250	UDP
715	17.628874	192.168.1.77	208.43.115.82	TCP
	711021	208.43.115.82	192.168.1.77	

Demo

PACKET CAPTURE [?]

Search | Results

From Time	To Time	Query	Cloudshark ID	Status
2017-04-21 10:21:00	2017-04-21 10:31:00	host 172.24.39.59 and host 64.62.174.41	dc4e5fe12349	Finished/OK

How does all that work?

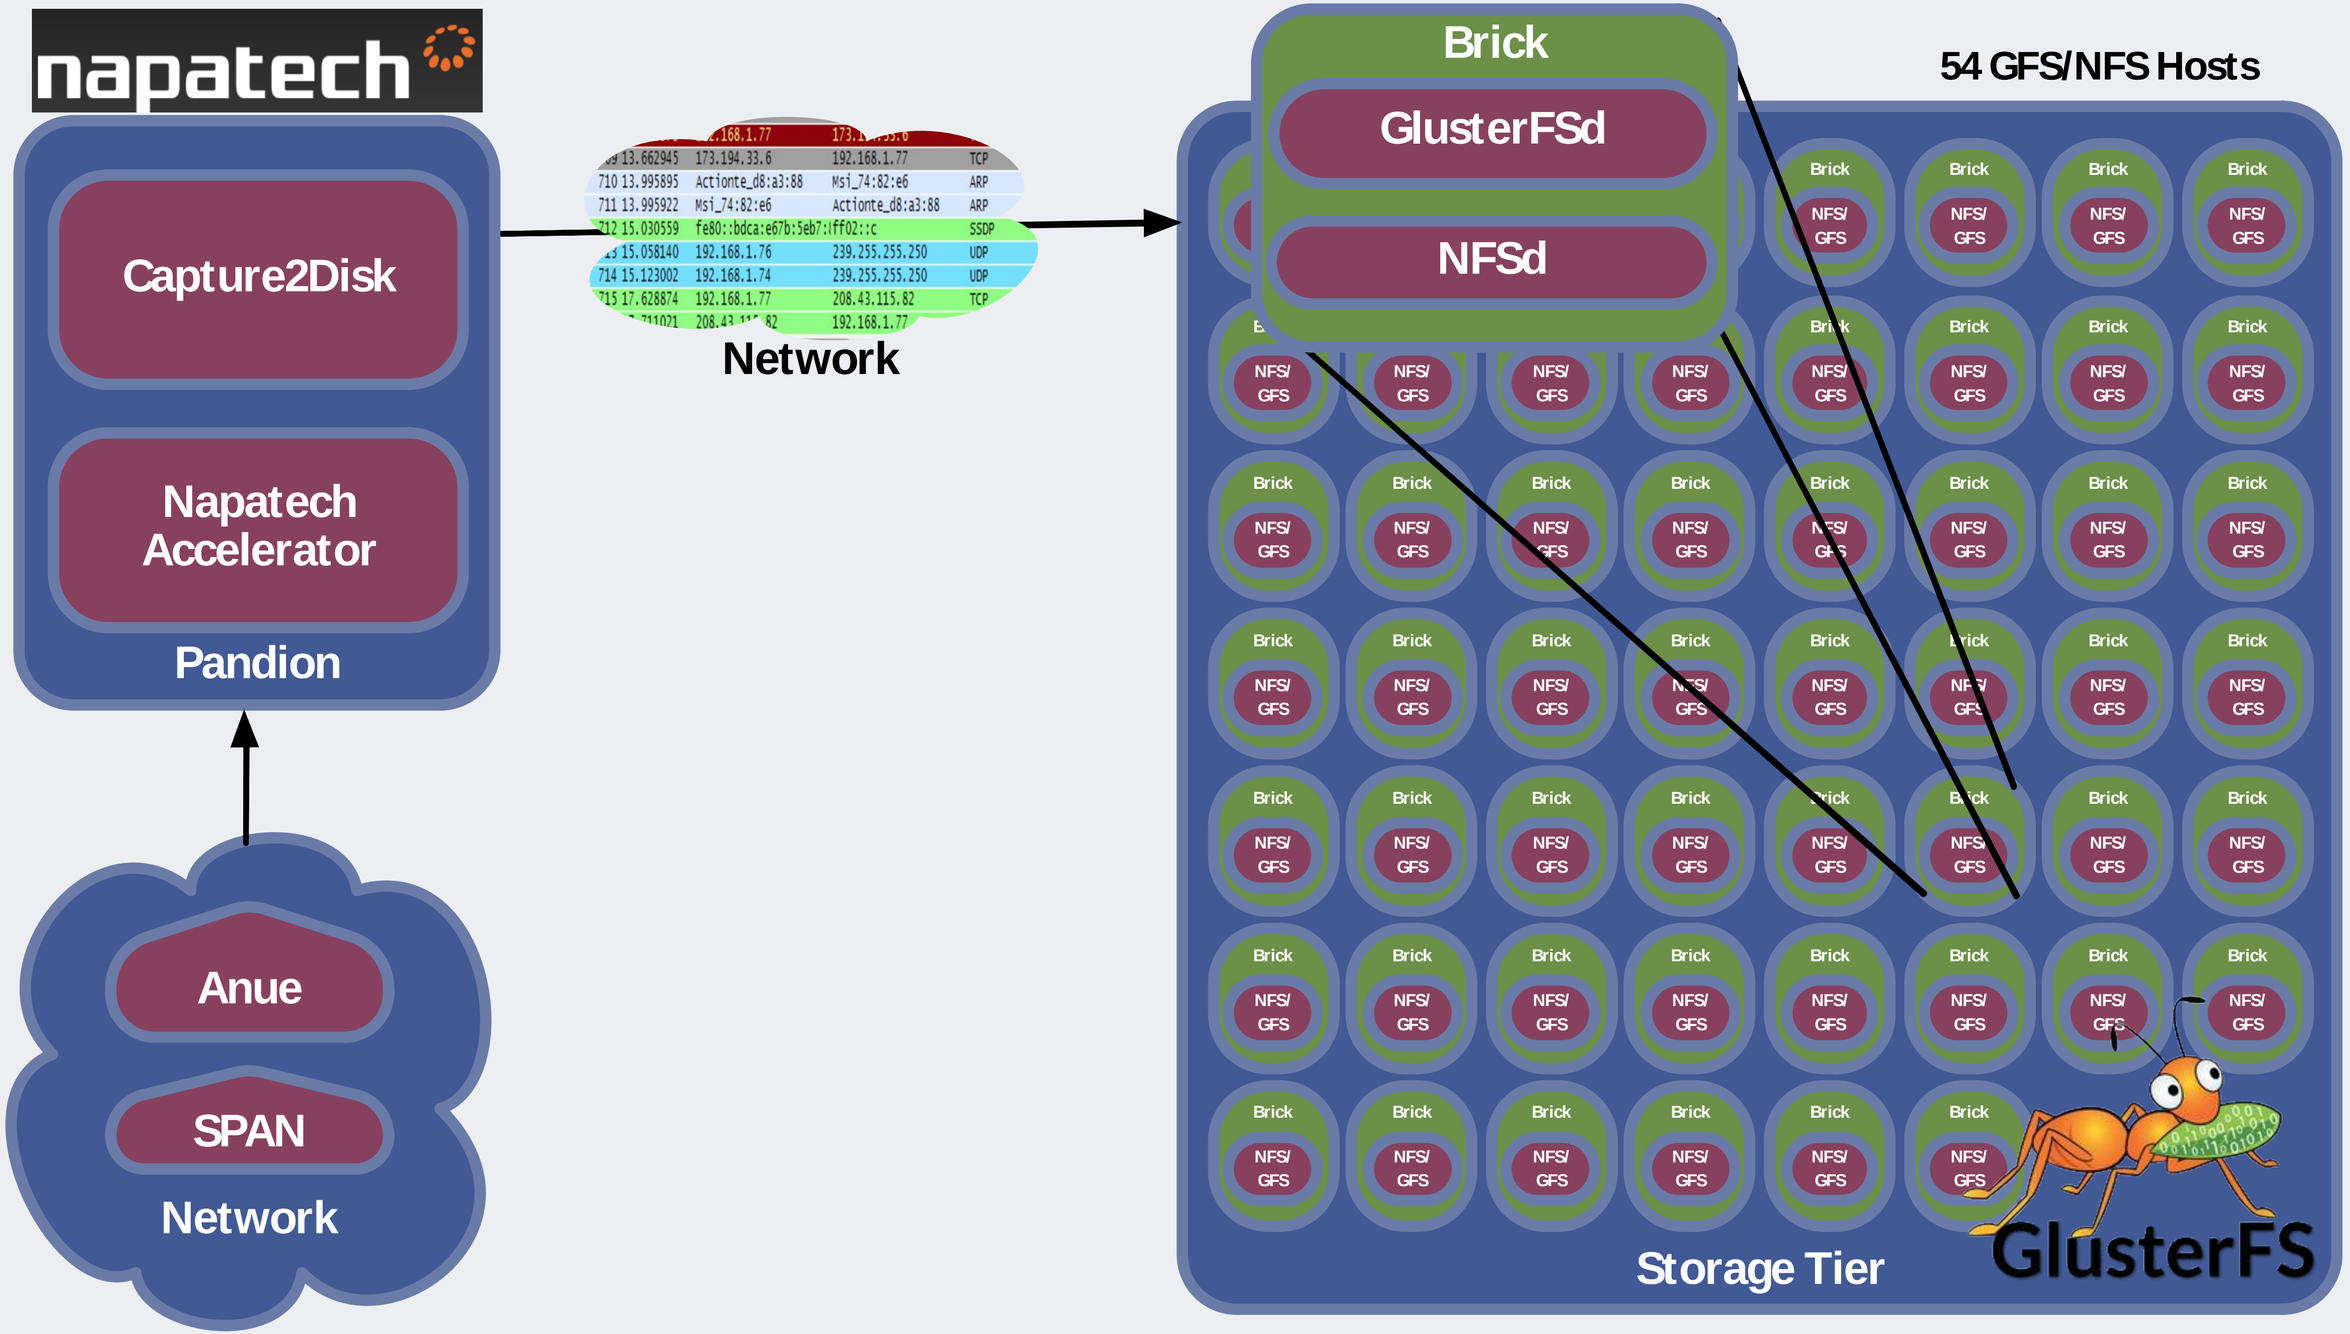
napatech
Capture2Disk
Napatech Accelerator
Pandion
Anue
SPAN
Network
Network
Brick
GlusterFSd
NFSd
54 GFS/NFS Hosts
Brick
NFS/ GFS
Storage Tier
GlusterFS

Network

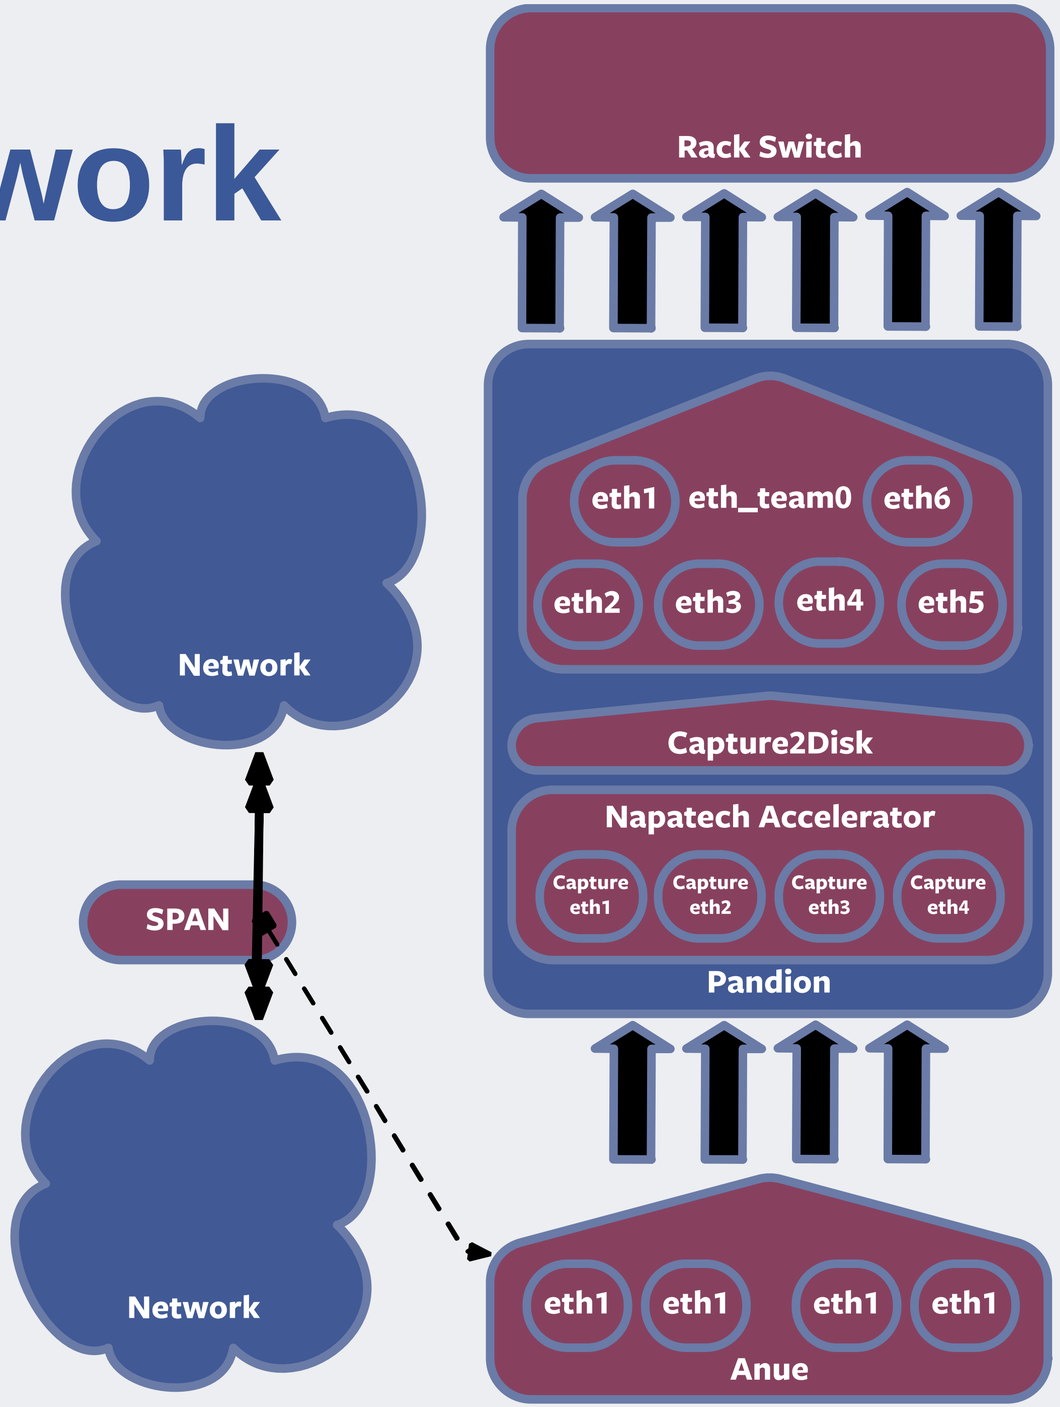

pcap-rack-switch1# sh port-channel load-balance

Port Channel Load-Balancing Configuration:
System: source-dest-ip

```
#ifcfg-team0
DEVICETYPE=Team
...
TEAM_CONFIG='{"runner": { "name": "loadbalance" },
 "tx_hash": [ "eth", "l3", "l4" ], "tx_balancer": { "name": "basic" } }'
...
#ifcfg-eno1
DEVICETYPE=TeamPort
...
TEAM_MASTER=team0
```

Load Balance Status

Enabled Status: **4 of 4 ports enabled**

Force Link Up: **Disabled**

Enabled Port Status

Combined Speed: **40G**

IPv6 Packets

Specify the fields to be used to load balance IPv6 packets (Ethertype 0x86DD).

- ☑ Source and destination IP addresses
- ☑ Next Header
- ☑ Source and destination L4 ports

PCAP Server

Napatech PandionFlex

- 4 SFP+ capture interfaces to Napatech accelerator
- PCAP writer/reader
- NIC Teaming - 6 interfaces (SFP+) for NFS traffic

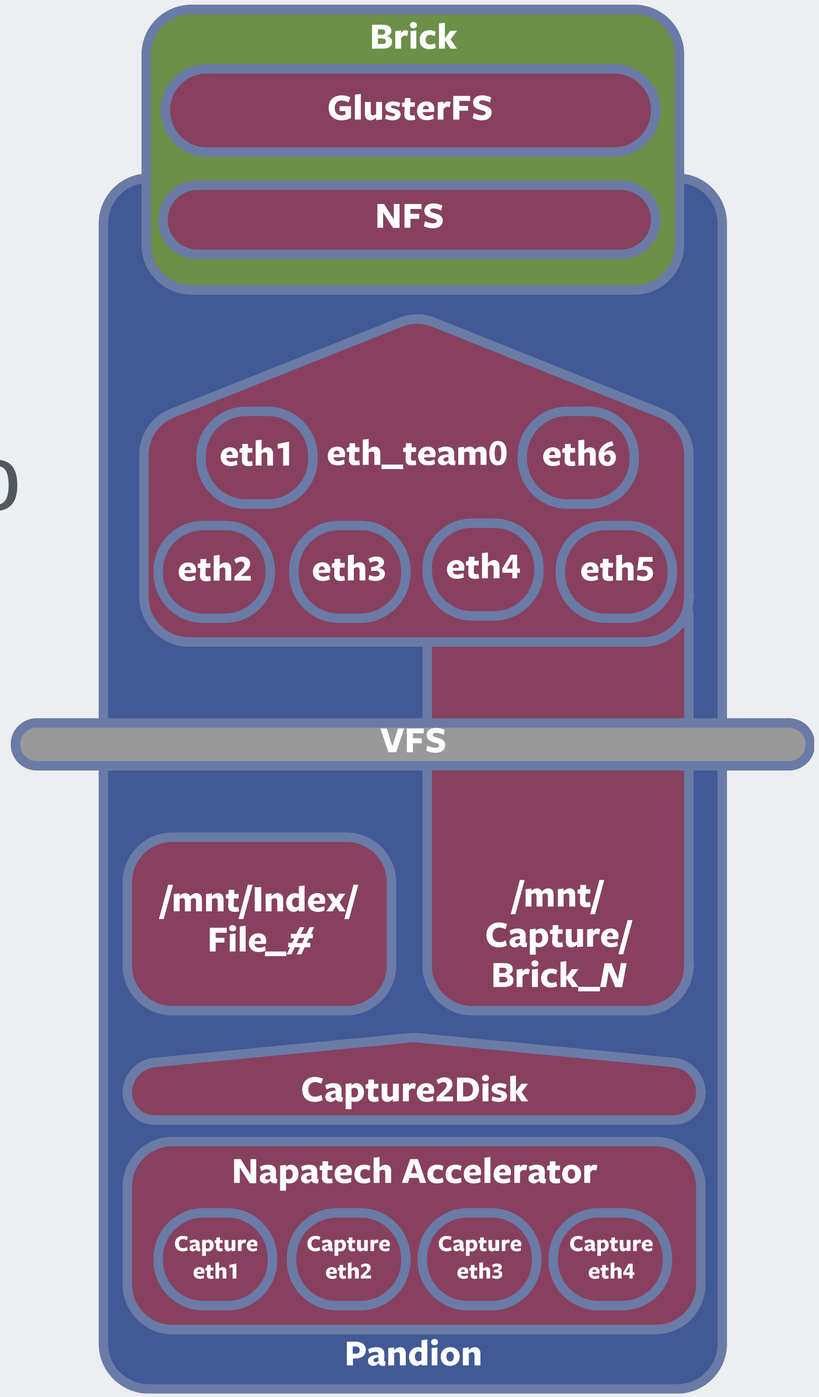

Storage Tier

GlusterFS & NFS on Open Compute Project Hardware

- Bricks run NFSd and GFSd
- Brick = 30 x 4TB HD ~100TB useable
- Storage Tier = 54 Bricks = 1.3PB useable

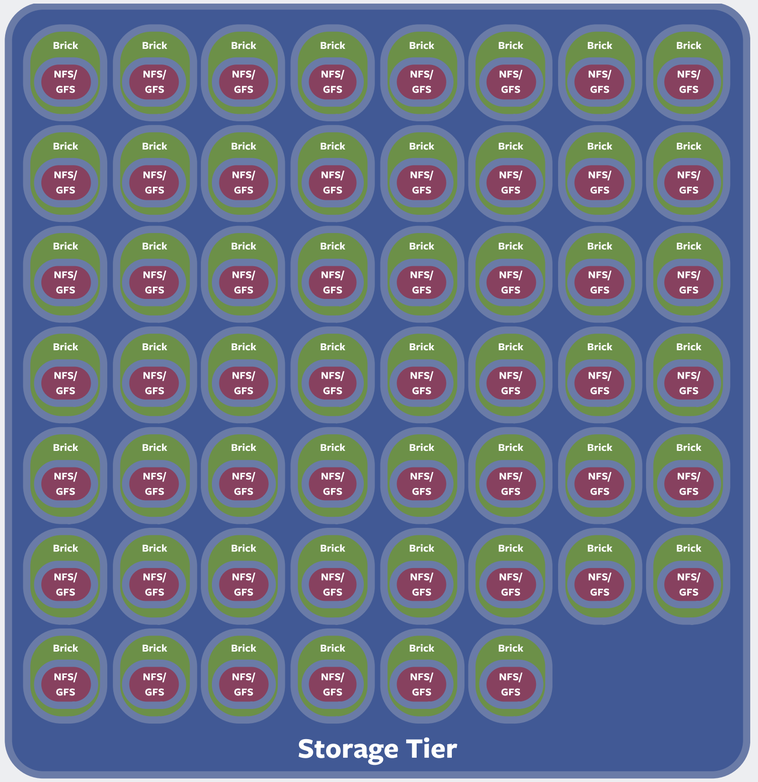

GlusterFS Introduction

High Level

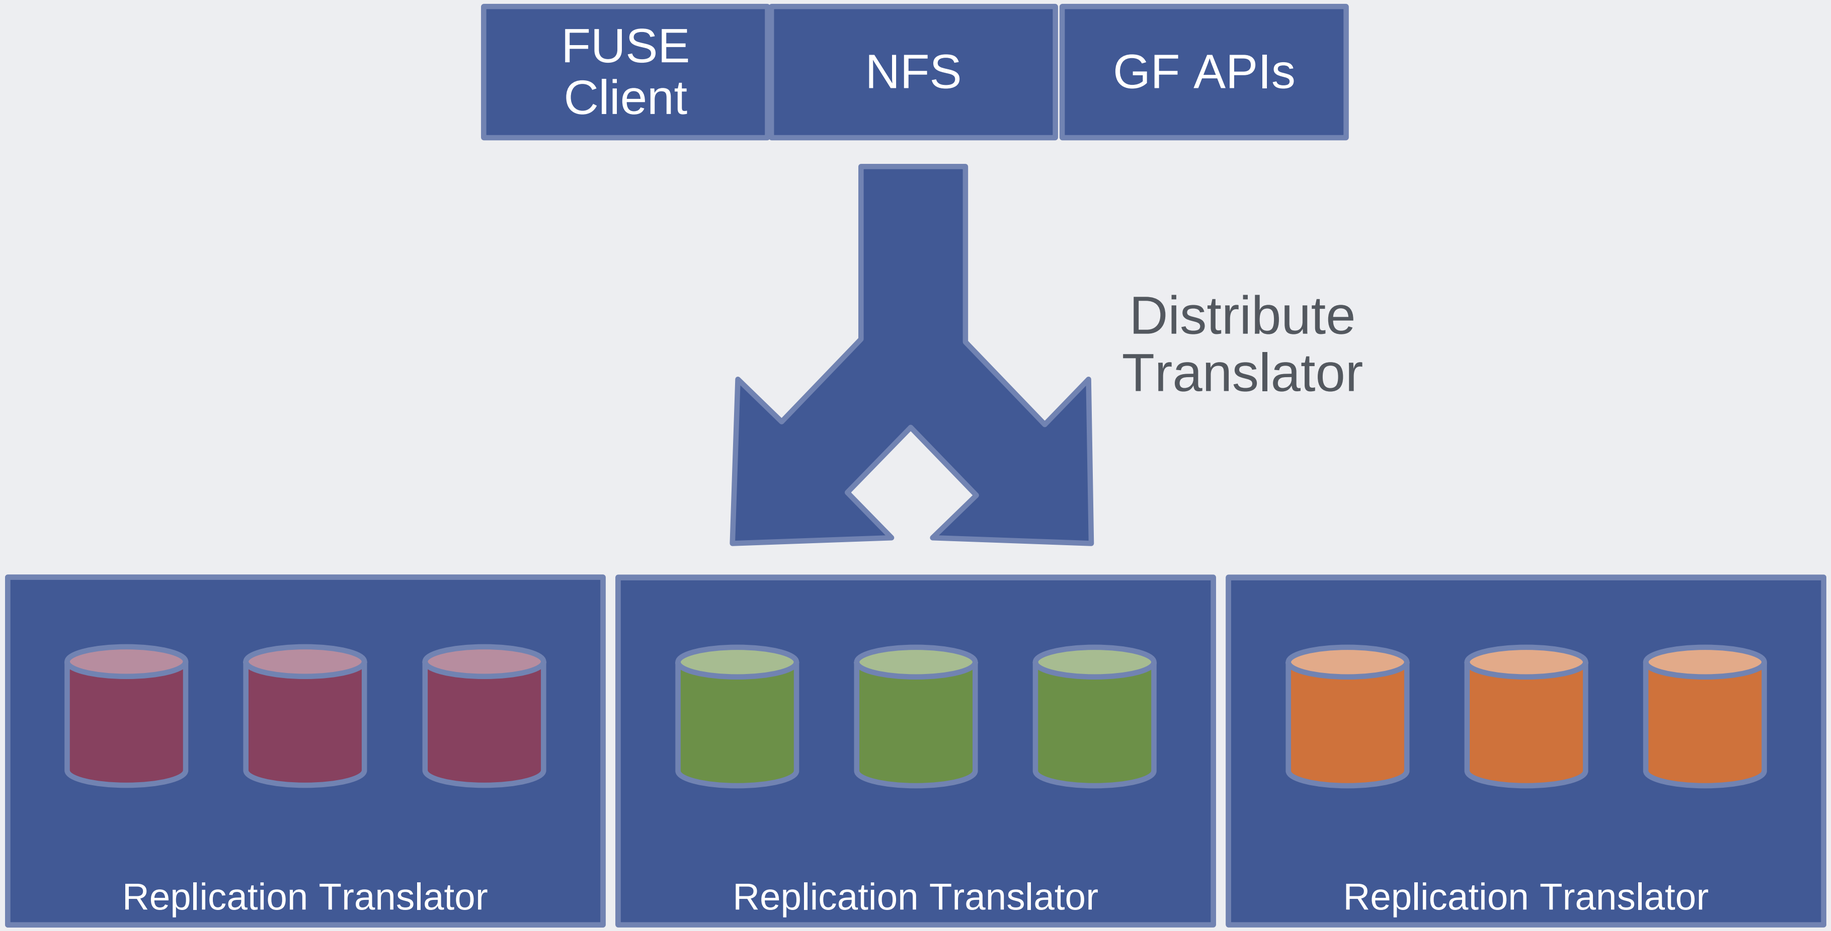

GlusterFS Introduction

Replicate Translator

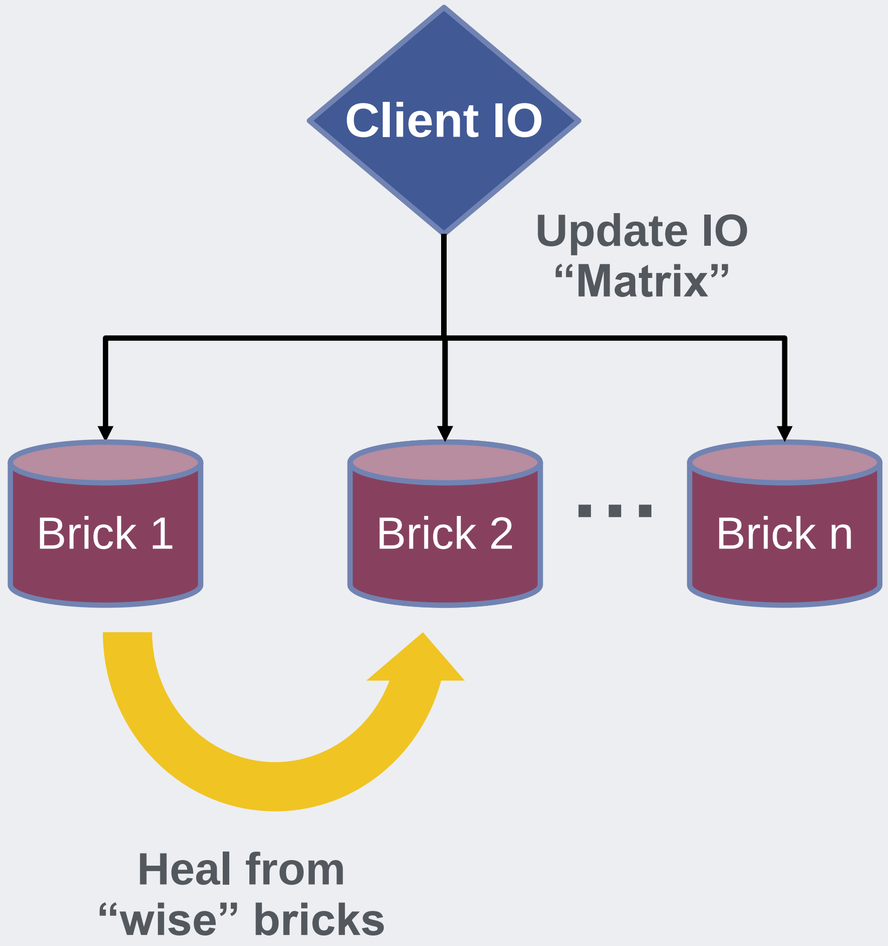

Why Napatech, Why Gluster? Why X?

- Didn’t want to create or support a pcap application
  - Existing Napatechs already in fleet, building features on accelerator cards
- Facebook GlusterFS active development branch https://github.com/gluster/glusterfs/tree/release-3.8-fb
- NFS = FUSE = easy to take existing products to network storage
  - Network storage provides flexibility to deploy anywhere

Things to watch out for

- Long Fat Network (LFN's)
  - Bandwidth Delay Product
- Multi-Homing pcap hosts
- Playing nice with your network
- File System best practices (directory structures, reads)
- NFS host failures
- PCAP buffers
  - Microbursts
- Fault Tolerance
  - Gluster Servers
  - PCAP Services
    - Reading
    - Writing

Next Generation Architecture

- Probably internal HDFS infra
  - MAP/Reduce function
    - Every time you touch a file, can you improve it?
      - Slices packets, run yara sigs over it? Generate new meta?
- Publisher/subscriber model
  - De-couples reader/writer
  - Readers can determine if file is ready to read without querying writer
  - Writer tasked with writing only

Thanks

- Questions?
- https://github.com/gluster/glusterfs/tree/release-3.8-fb
- edub@fb.com & mmoran@fb.com