

RED HAT
SUMMIT

TECHNOLOGICAL PIVOT POINTS LIGHTNING TALKS

Jon Masters
Steven Carter
Veer Muchandi & Tom Corcoran
Tom Coughlan


May 2, 2017

The logo for Red Hat Summit, featuring the words "RED HAT" in a smaller font above "SUMMIT" in a larger, bold font, both in white on a red, downward-pointing arrow-shaped background.

RED HAT
SUMMIT


Modernizing HPC for Cloud

Steven Carter
Chief Cloud Architect, Public Sector
2 May 2017

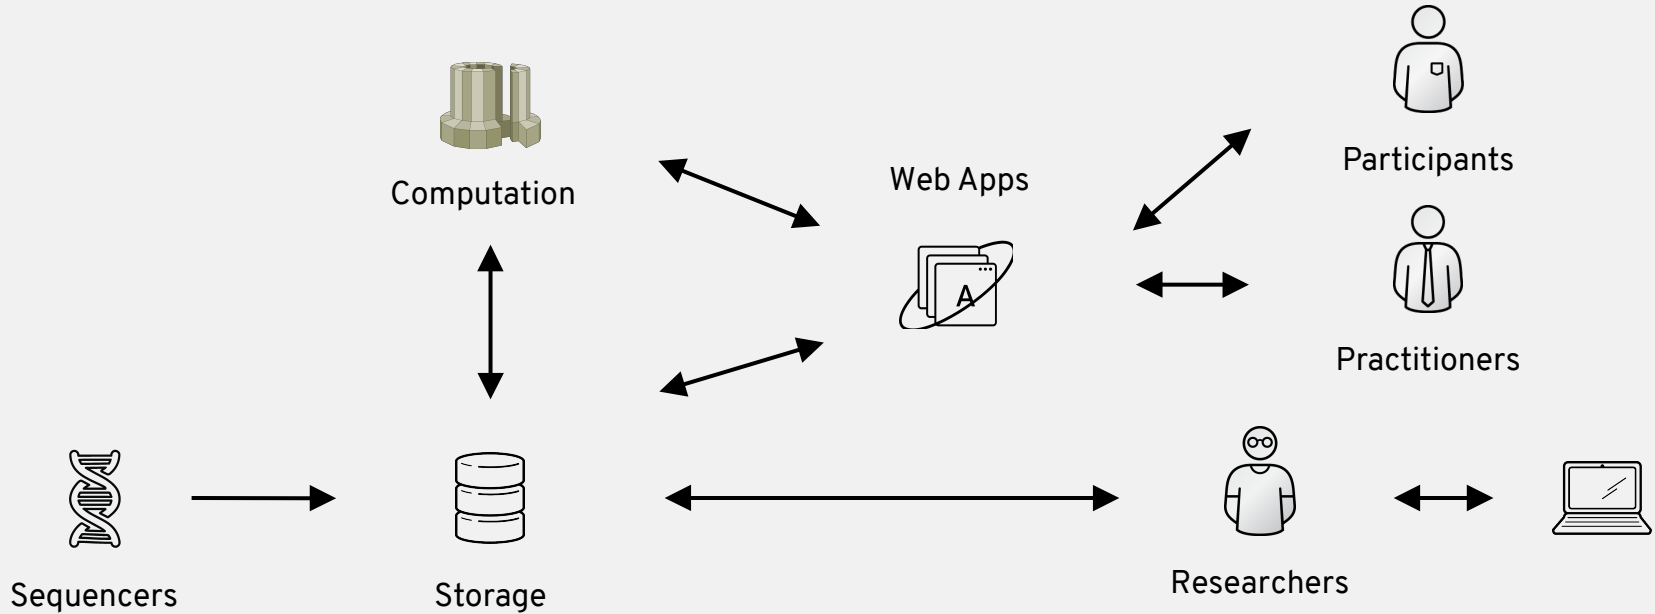


Wall Clock


VS



Results



Science is Multidisciplinary




Supercomputer


VS



Cloud



But, not just a Supercomputer in the
cloud...



Posix Filesystem

VS




Object Storage




Code

VS




Containers

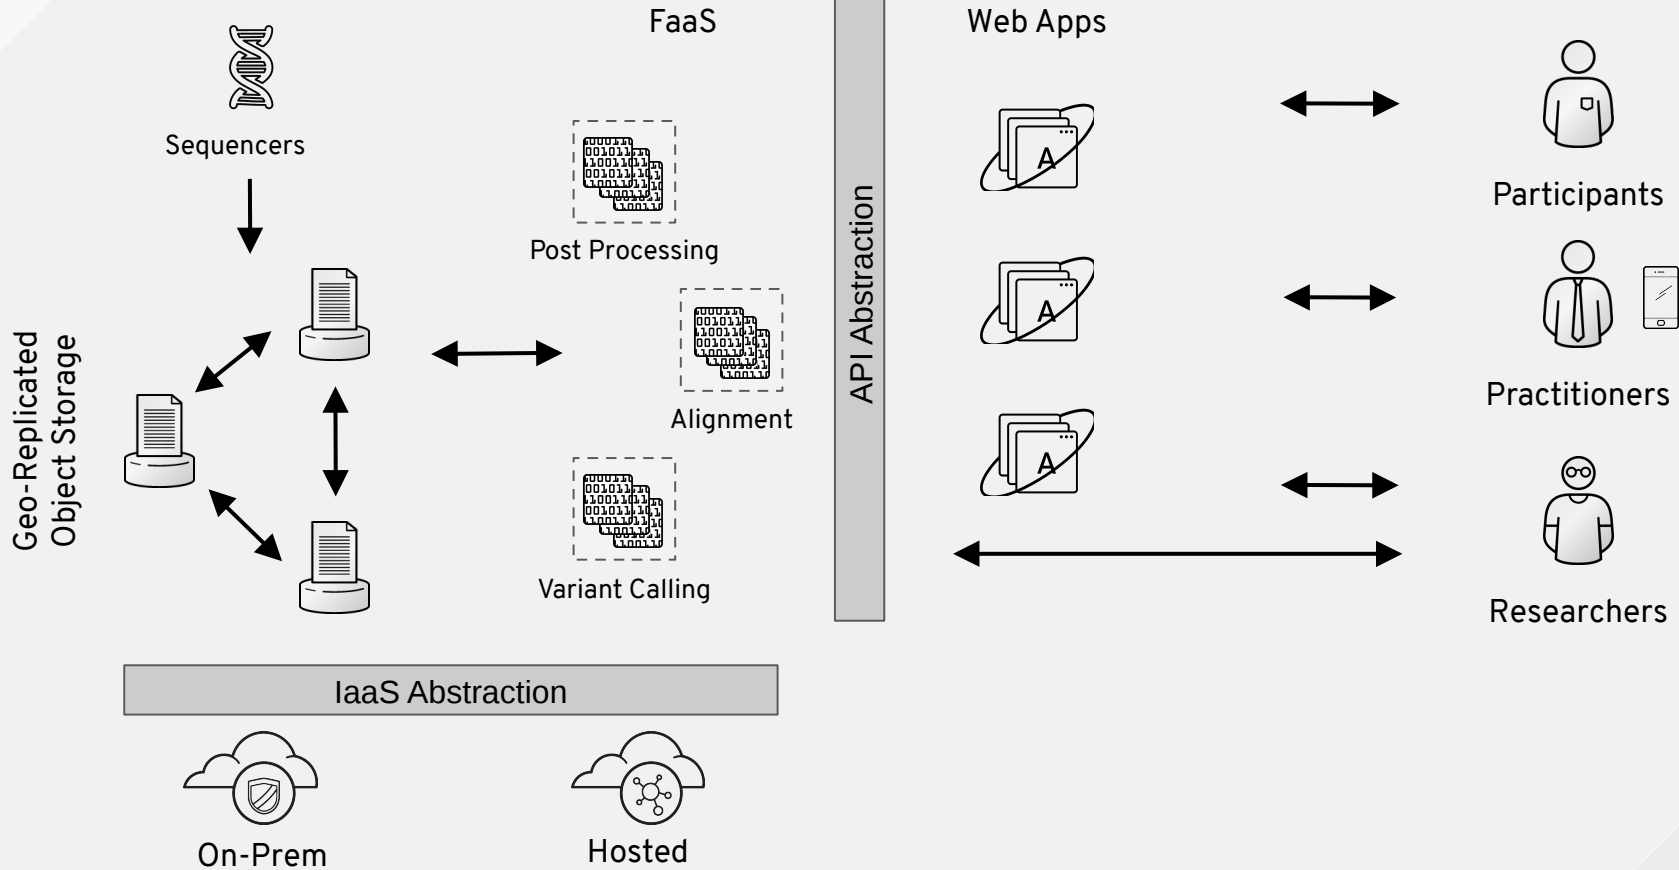


Monolithic

VS



MicroServices



Circuit Breaker using 3Scale

Dealing with Microservice failures


Tom Corcoran, Sr. Solution Architect, Red Hat

Veer Muchandi, Principal Architect, Red Hat


05/02/2017

Problem


Client calling service



Service Failure



Service failure effect multiplied



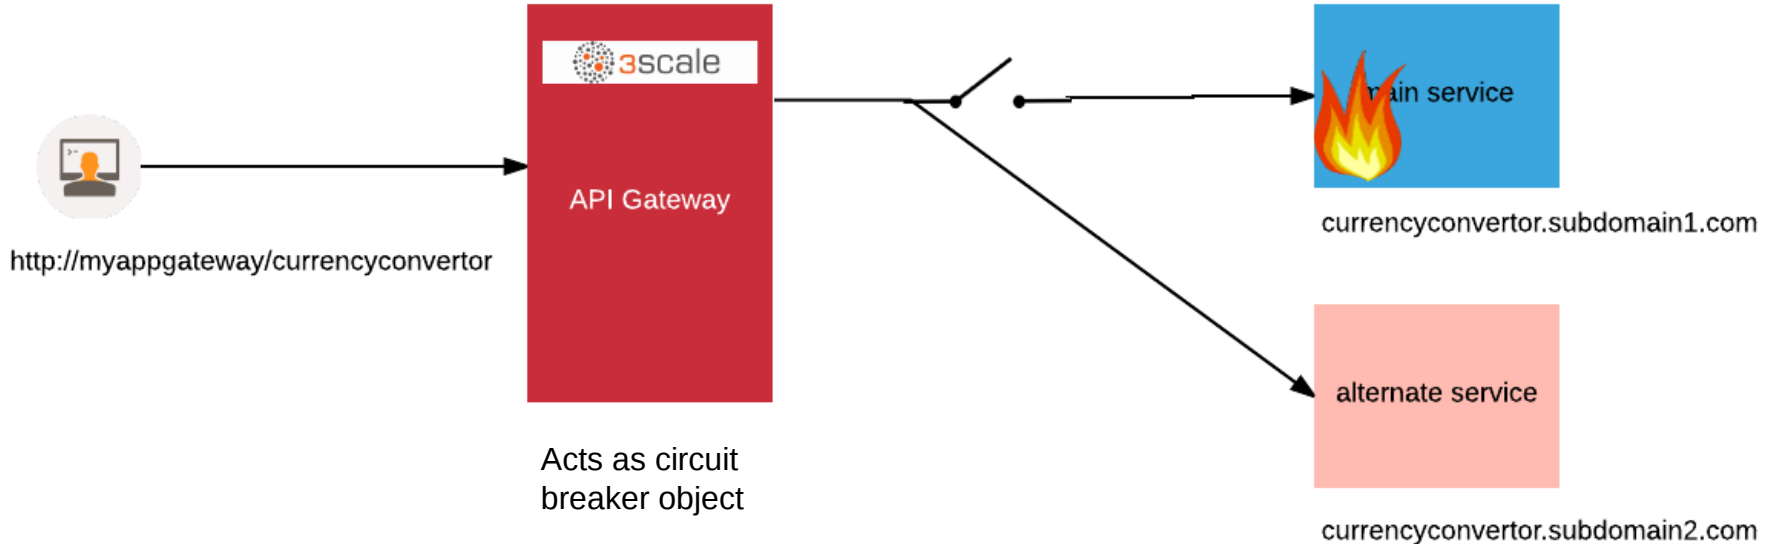
Circuit Breaker Implementation

using 3Scale

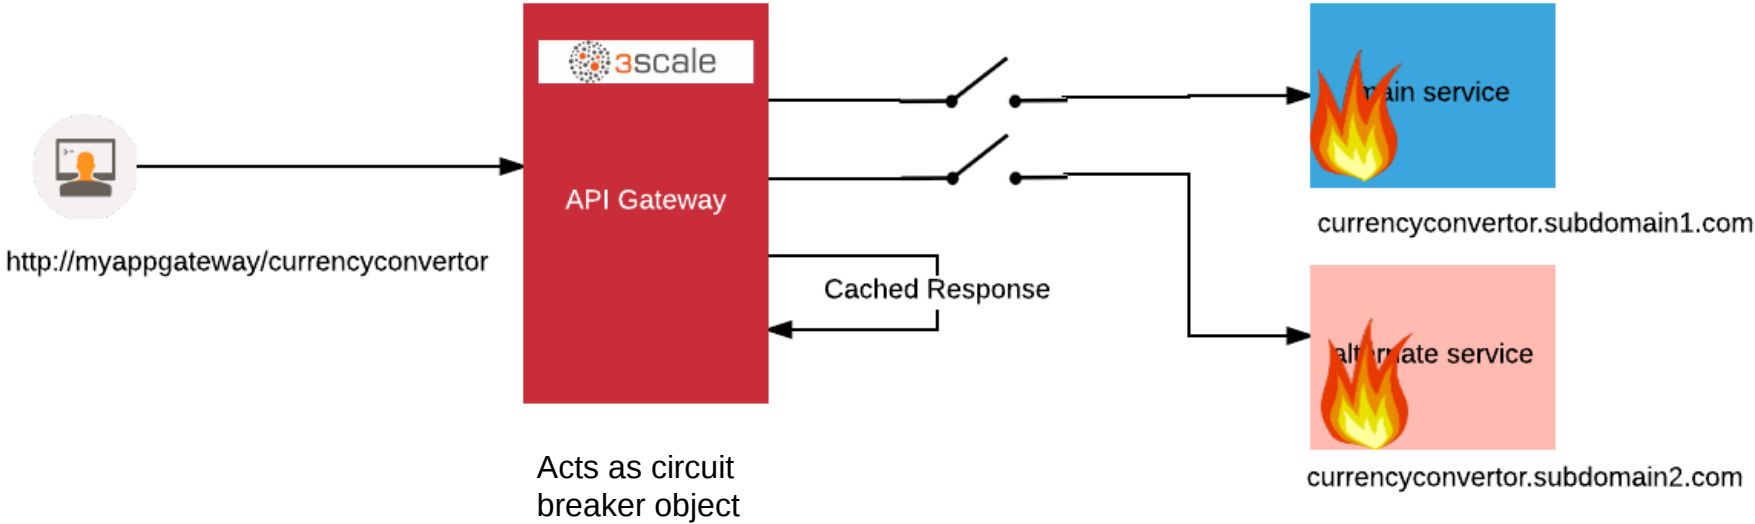
Normal Operation - 3 Scale forwards requests to the main service



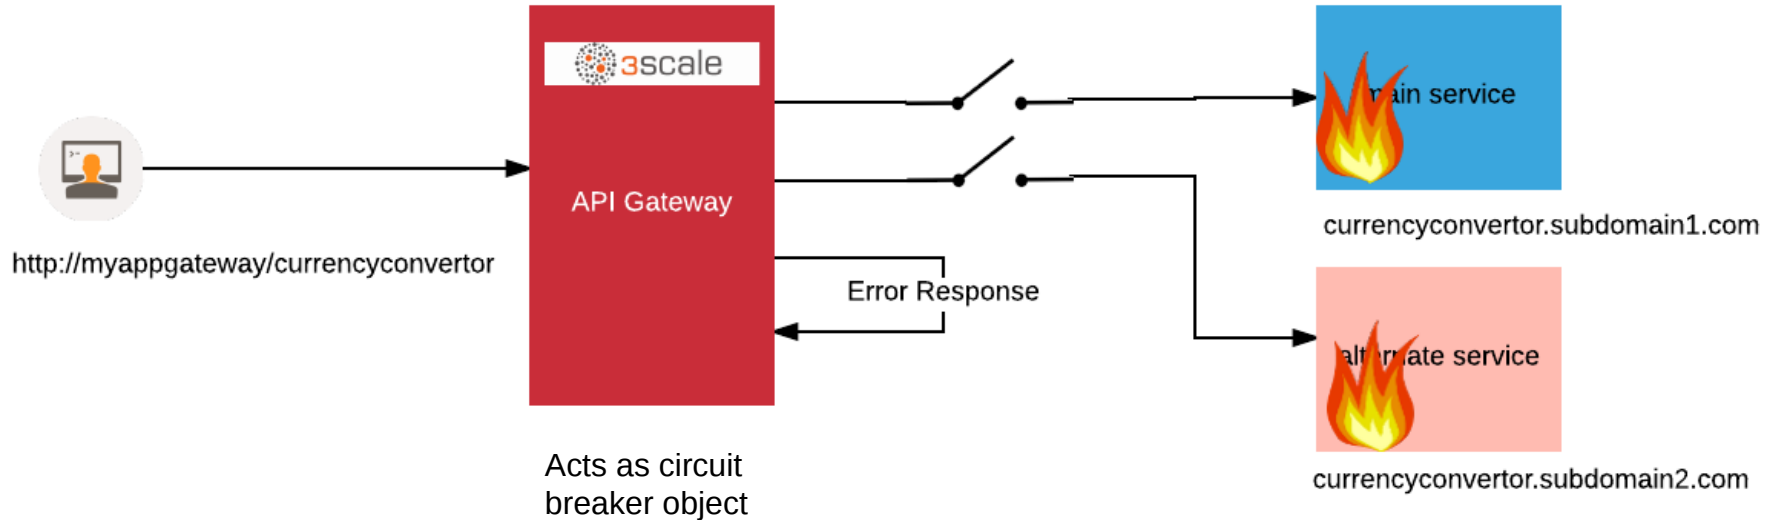
Fail over - Circuit broken to the main service and requests forwarded to the alternate service



Graceful handling on Alternative service failure with cached response



Graceful handling of Alternative service failure with Error response



Advantages of using 3Scale

Circuit breaker becomes an infrastructure service

- Works with simple configuration changes in the `nginx.conf`
- No code intrusion, No annotations in business logic

Underlying technology does not matter

Language Neutral

Demo

SUPPORT – SLIDE - Circuit Breaker – on Nginx

Nginx uses

- Upstreams – for API Back Ends

```
upstream backend_sms_API {  
    ip_hash;  
    server subdomain1.sms-backend.com;  
    server subdomain2.sms-backend.com;  
}
```

- Locations – to capture incoming API requests and proxy pass them

```
location /sms {  
  
    proxy_pass http://backend_sms_API;  
  
}
```

SUPPORT – SLIDE - Circuit Breaker – on Nginx

Back end failover – fine grained control

Override default selection of Next Upstream.

```
location /sms {  
    proxy_next_upstream    error timeout invalid_header http_504 http_503;  
    proxy_connect_timeout  2;  
    proxy_pass http://backend_sms_API;  
}
```


SUPPORT – SLIDE - Circuit Breaker – on Nginx

Failure Mitigation

Slow Start

```
upstream backend_sms_API {  
    slow_start=30s;  
  
    server subdomain1.sms-backend.com;  
    server subdomain2.sms-backend.com;  
}
```

Rate Limiting

```
http {  
    limit_req_zone $server_addr zone=moderateReqs:1m rate=100r/s;  
    ...  
    server {  
        ...  
        limit_req zone=moderateReqs burst=150;  
        limit_req_status 503;  
        ...  
    }  
}
```

SUPPORT – SLIDE - Circuit Breaker – on Nginx

Last Resort Options – no back end available

Response caching.

```
proxy_cache_path /app/cache levels=1:2 keys_zone=oauth_cache:10m max_size=10m inactive=15s

location /v1/users {
    ....
    proxy_cache_valid      200 30s;
    proxy_cache_use_stale  error timeout invalid_header updating
                          http_500 http_502 http_503 http_504;
}
```

Error Page.

```
location / {
    error_page 503 = @fallback;
}

location @fallback {
    proxy_pass http://alternative-backend;
}
```

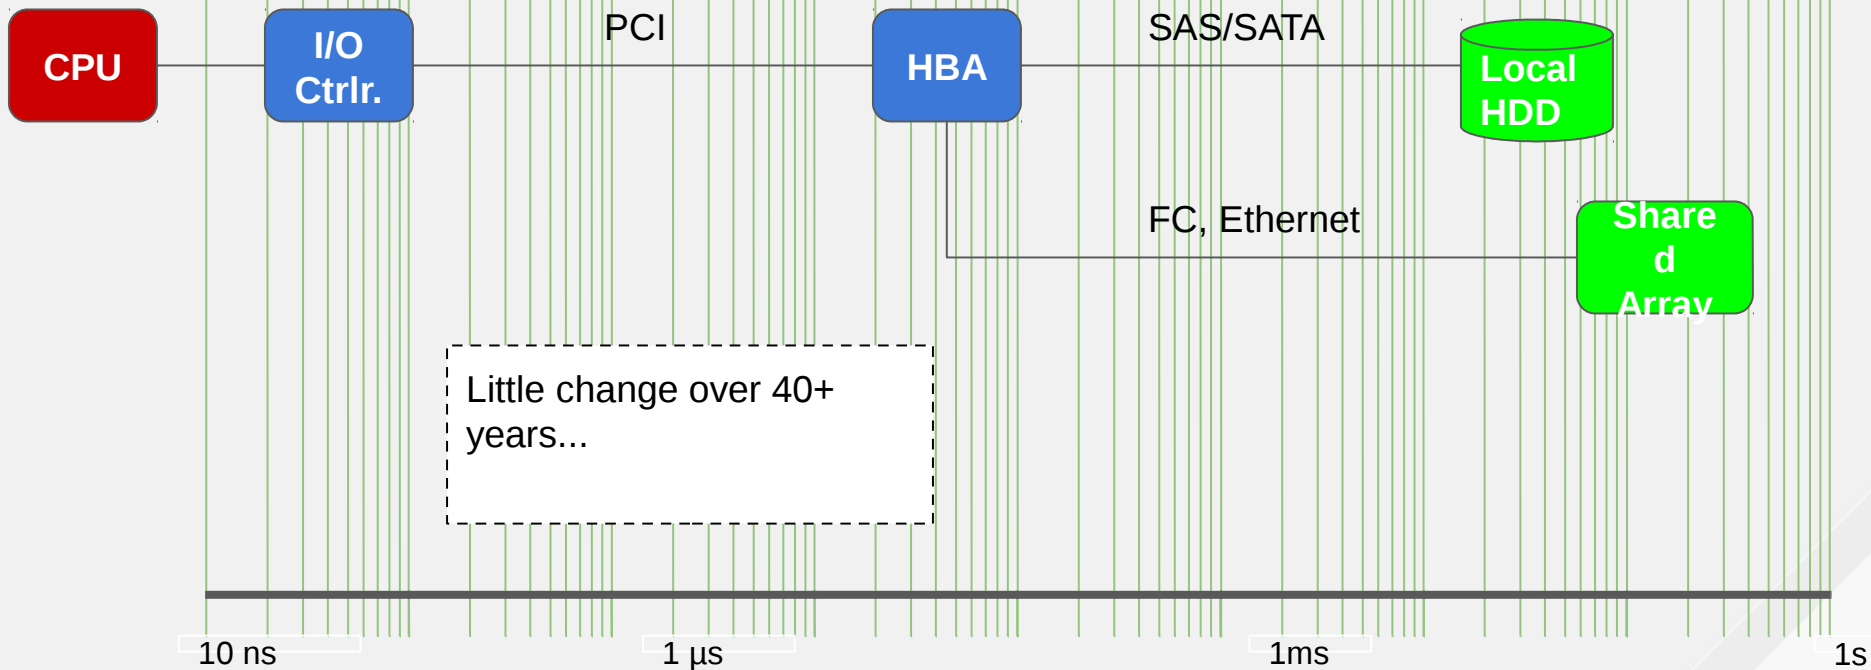
The logo for Red Hat Summit, featuring the words "RED HAT" in a smaller font above "SUMMIT" in a larger, bold font, all contained within a white speech bubble shape.

RED HAT
SUMMIT

Fast as lightning : NVMe and NVDIMM's Impact on your data center

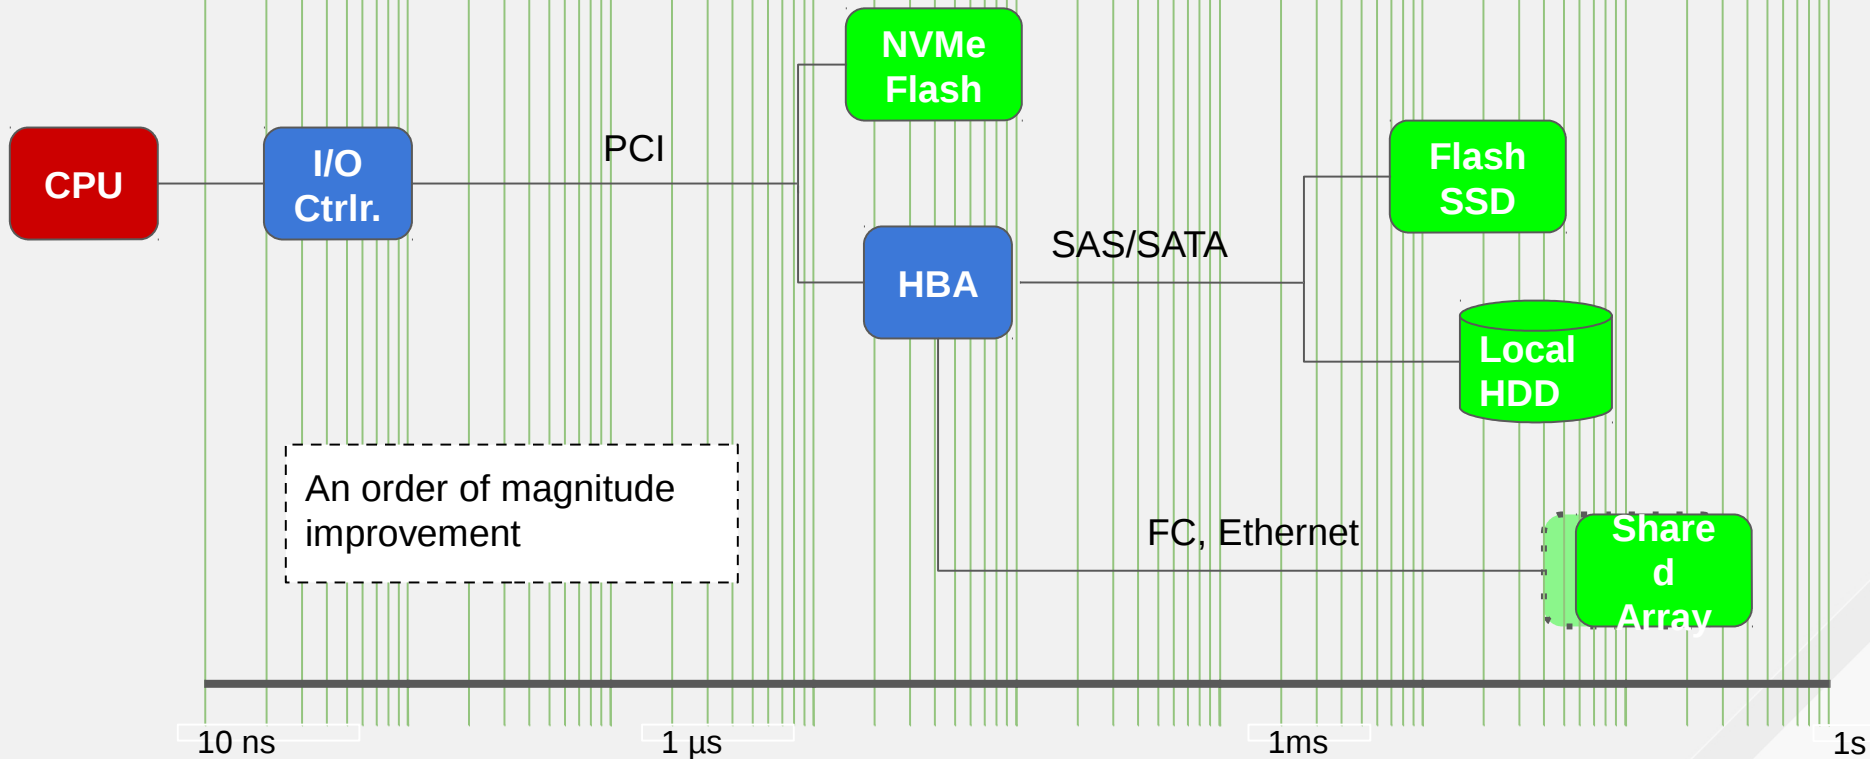
Tom Coughlan
Sr. Engineering Manager
Filesystem and Storage Team, Red Hat
May 2, 2017

I/O Latency



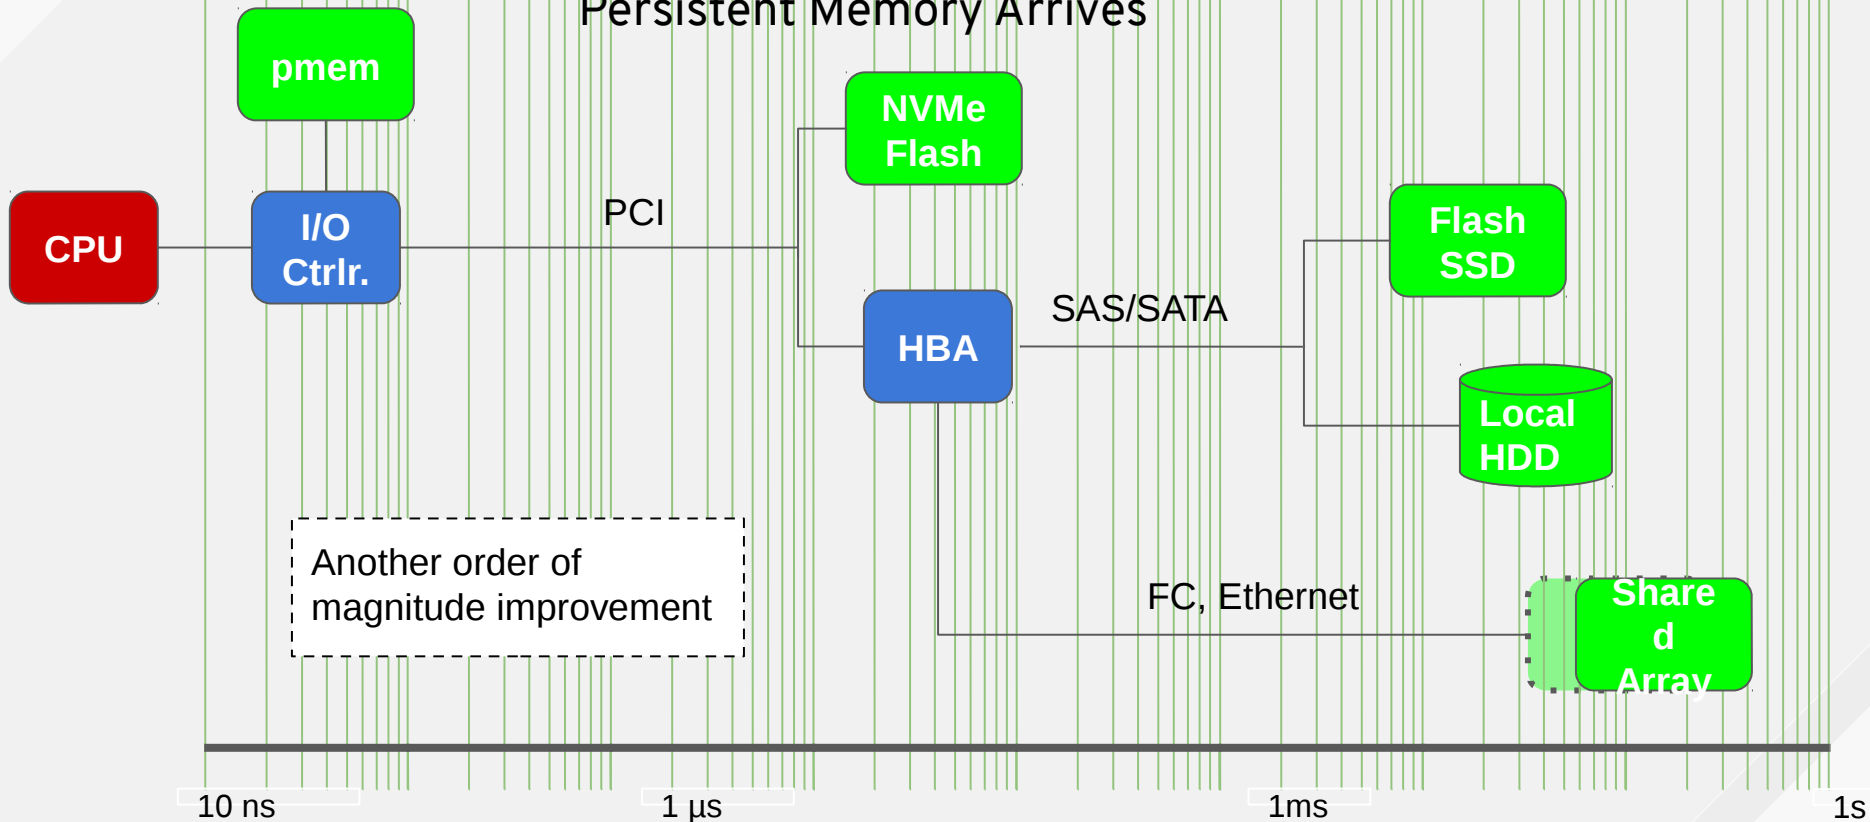
I/O Latency

Flash and NVMe arrive.



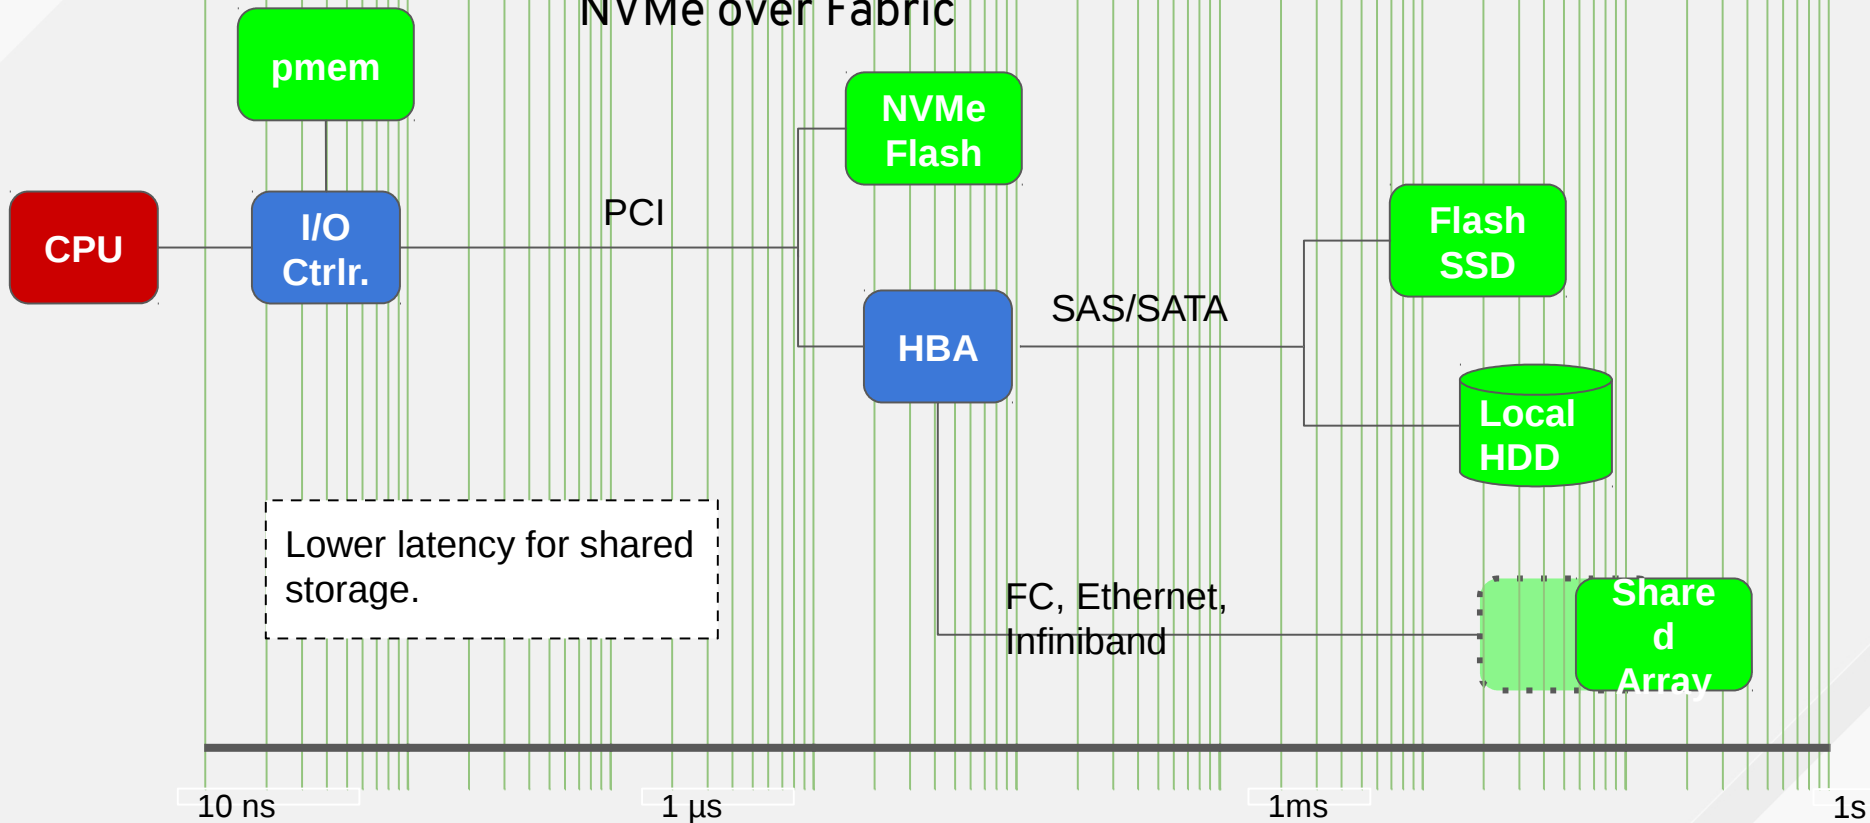
I/O Latency

Persistent Memory Arrives

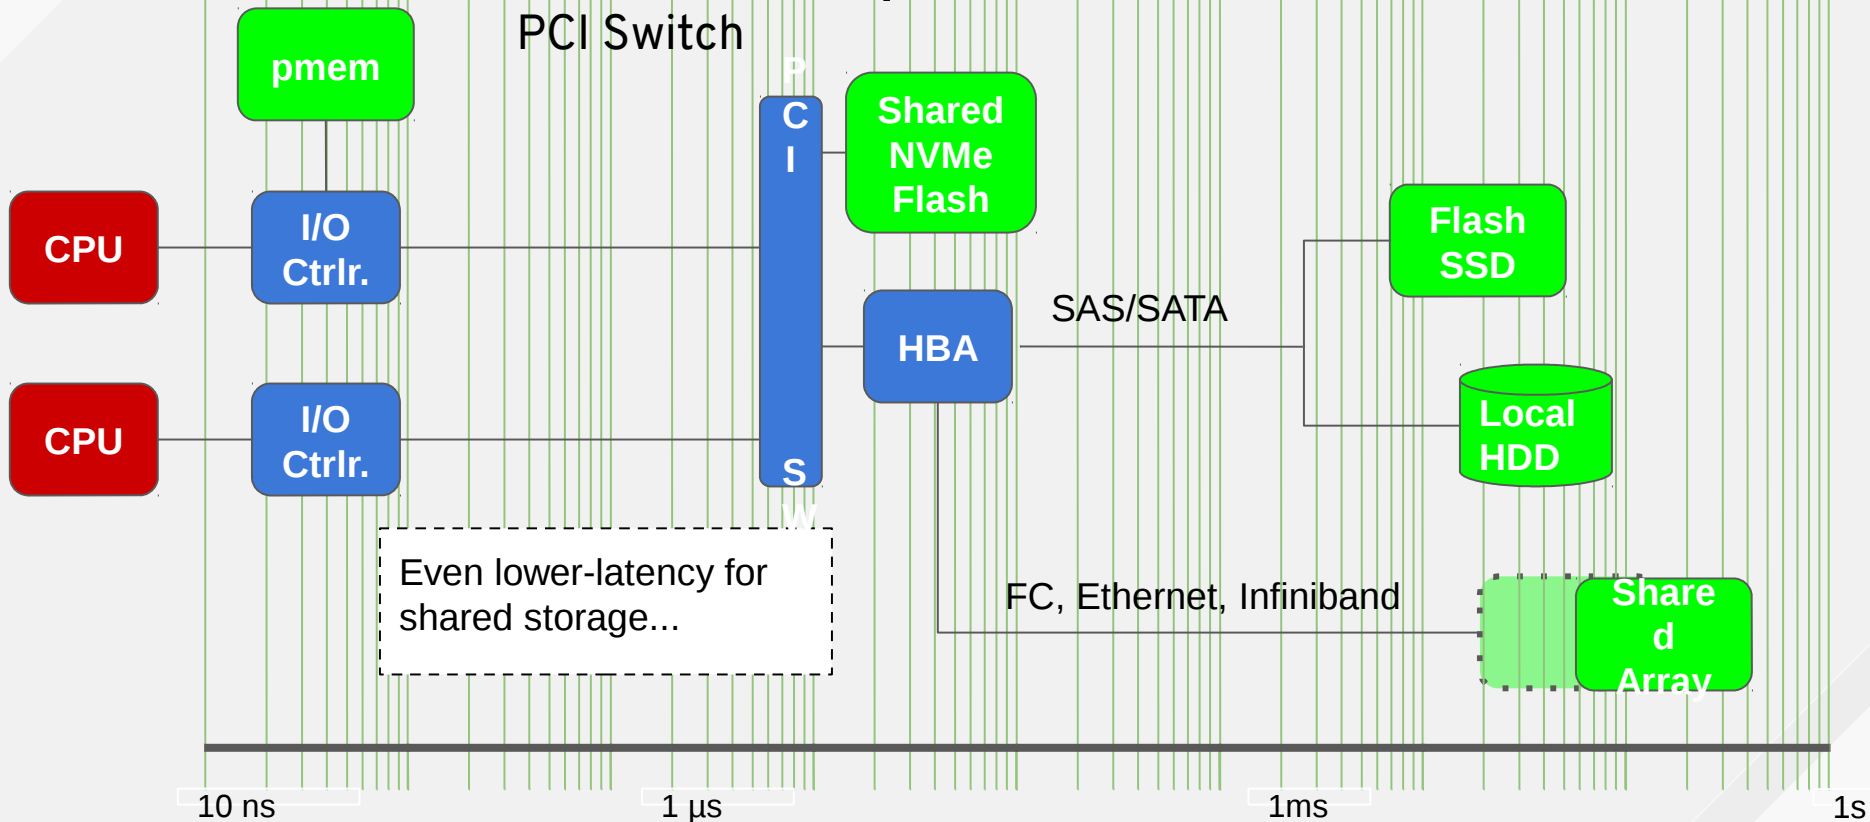


I/O Latency

NVMe over Fabric



I/O Latency



RHEL Status

- Persistent Memory
 - Block-mode access supported in 7.3
 - Byte-mode direct access (DAX) is Tech. Preview in 7.3 (and 7.4)
 - Uses mmap on XFS or ext4
 - HPE results, using NVDIMMs for...
 - Oracle OLTP redologs
 - <https://www.hpe.com/h20195/v2/GetPDF.aspx/4AA6-6008ENW.pdf>
 - PostgreSQL Write-Ahead Logging
 - <http://h20195.www2.hpe.com/V2/GetDocument.aspx?docname=4AA6-8881ENW>
- NVMeF - try it in RHEL 7.4 Beta

Additional sessions of interest...

- “On the path to persistent memory”
BOF B113761
Tuesday, 10:15 to 11 am (oops...) Room 158
Linda Knippers (HPE) and Jeff Moyer (Red Hat)
- “Bring your performance and scale problems to the experts”
Session 101558
Tuesday 5:30 PM - 7:00 PM Room 205C
- “Utilizing Persistent Memory to Improve DB Performance and Reduce Costs”
Session S111008
Wednesday at 4:30 pm Room 105
Karen Dorhamer (HPE)

RED HAT
SUMMIT

THANK YOU



plus.google.com/+RedHat



facebook.com/redhatinc



linkedin.com/company/red-hat



twitter.com/RedHatNews



youtube.com/user/RedHatVideos

The logo consists of a red speech bubble shape pointing downwards, containing the text "RED HAT" in a smaller font above "SUMMIT" in a larger, bold font.

RED HAT
SUMMIT

**LEARN. NETWORK.
EXPERIENCE
OPEN SOURCE.**