

The Emergence of the Datacenter Developer

Tobi Knaup, Mesosphere · @superguenter



A Brief History of Operating Systems

1950's

Punchcards

No operating systems

Time Sharing

Computing as a utility



1960's

UNIX

Small, composable programs

Machine-independent
language (C)

Multitasking & Multiuser

Everything is a file

Shell



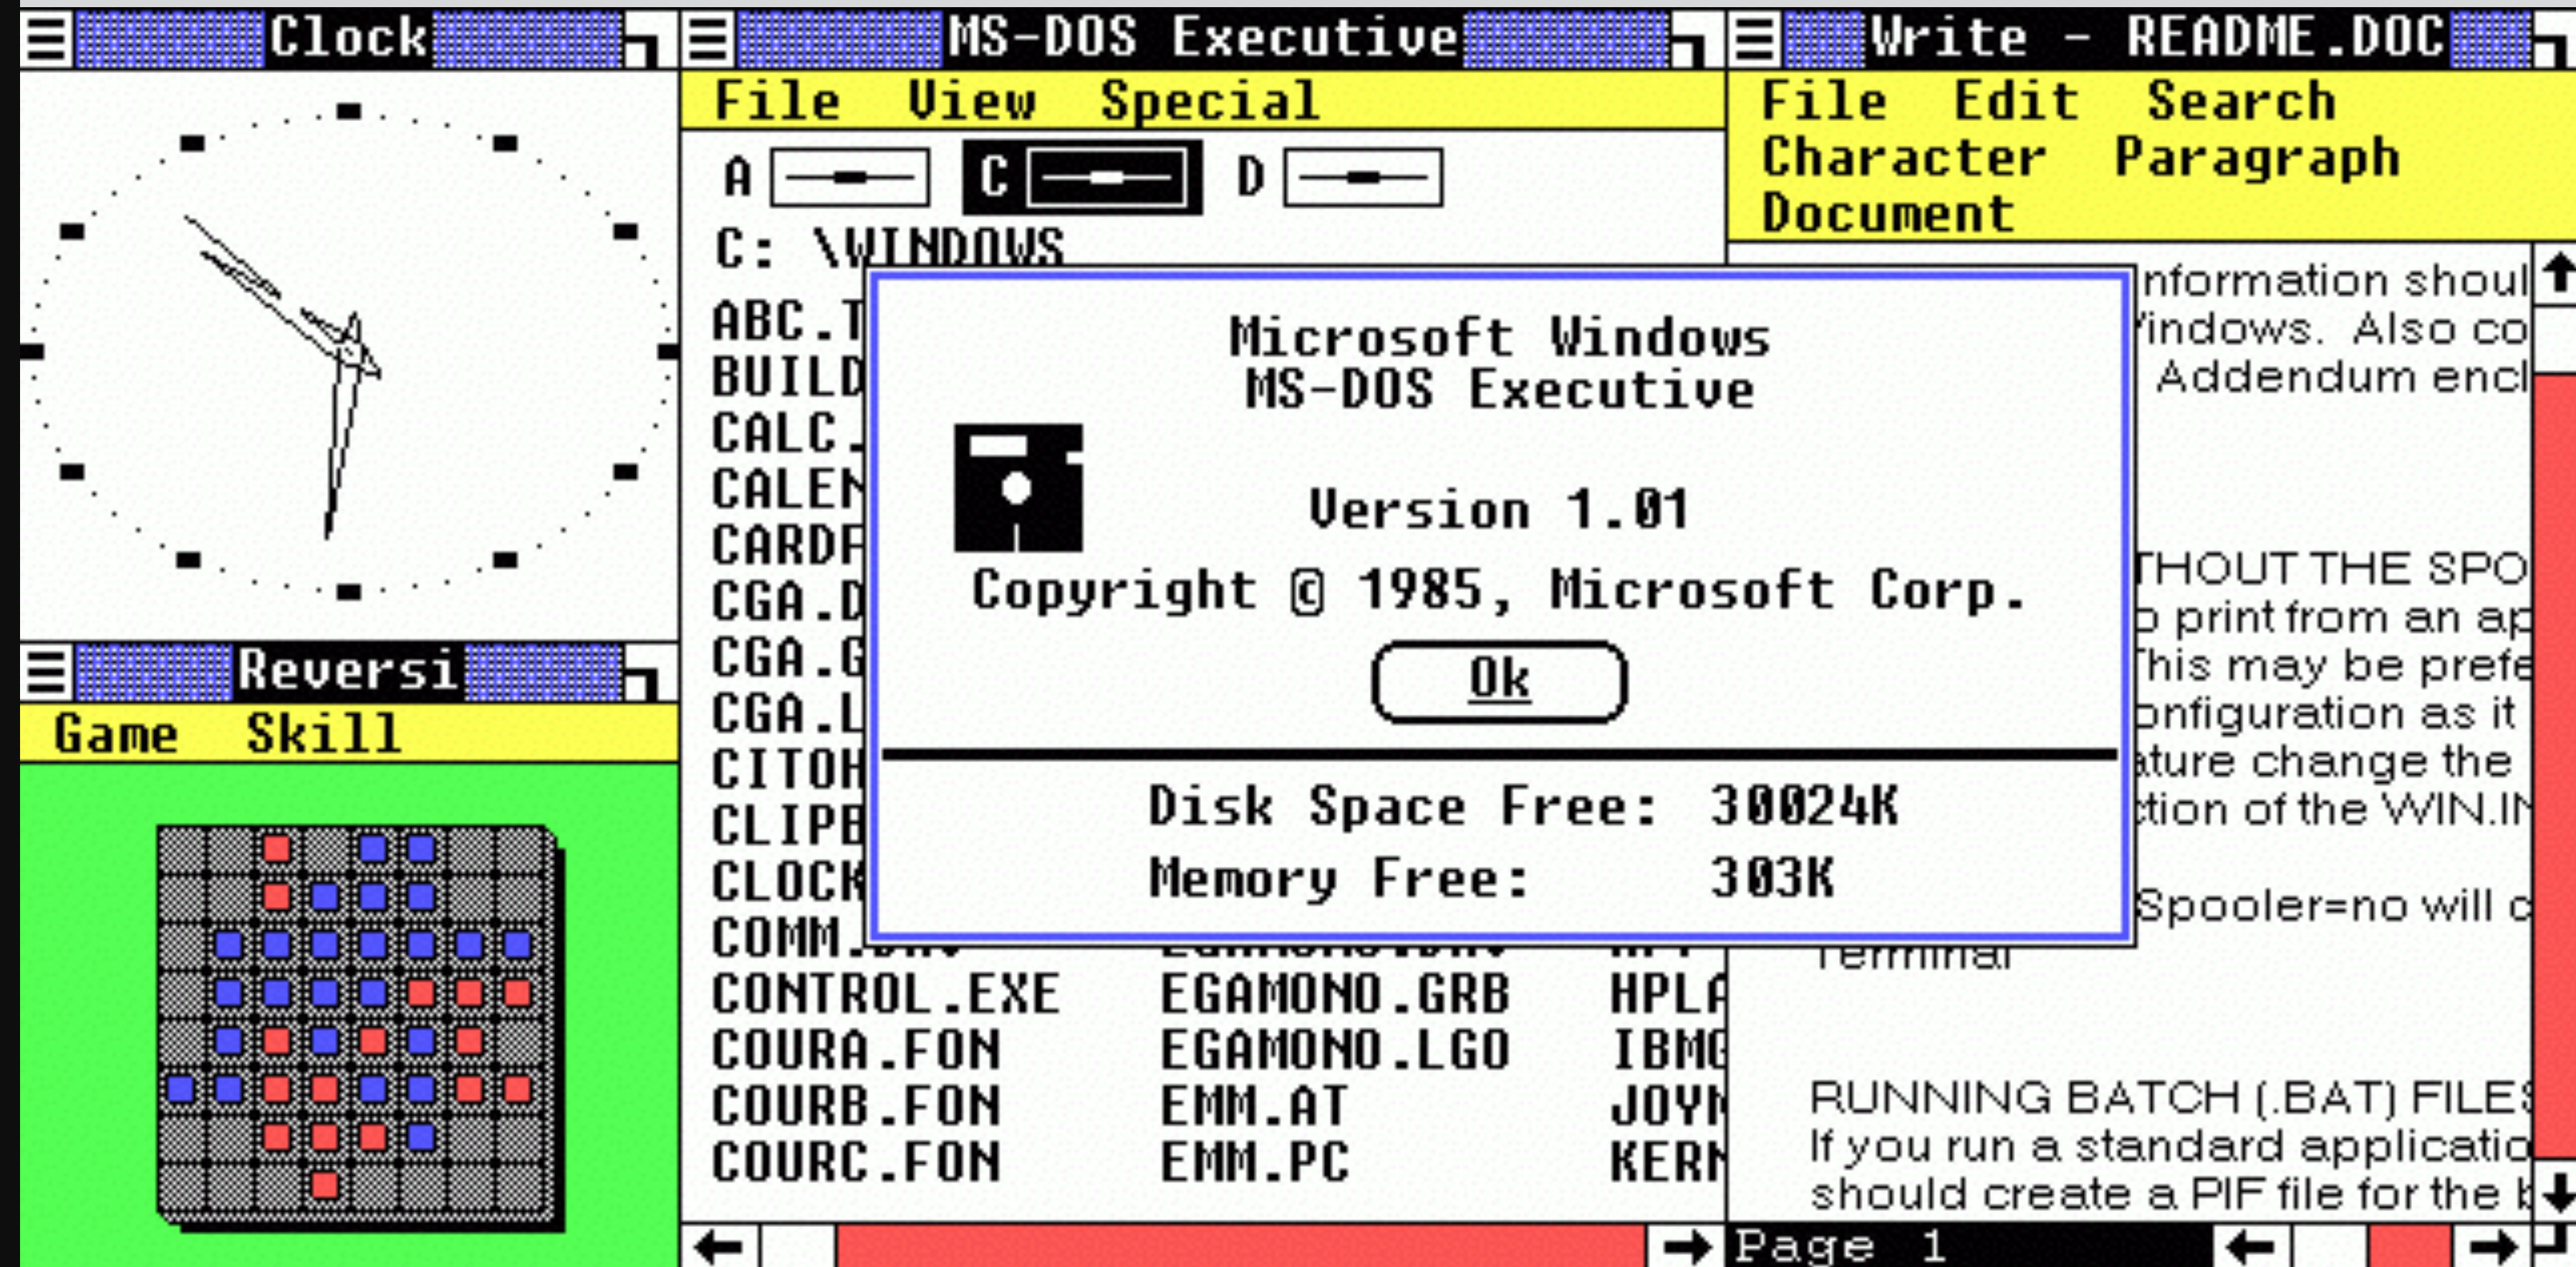
1980's
PC OS

Windows, Mac OS, Amiga

GUI

Mouse

Virtual Memory



2000's

Mobile OS

iOS, Android

Touchscreen

Camera

GPS, Accelerometer



2010's

Datacenter OS

Apache Mesos

Linux Containers


Cluster Scheduling

API for the datacenter



The sidebar shows the DCOS logo at the top, followed by the Datacenter Name '172.03.12.1'. Below are navigation links for 'Dashboard', 'Services' (highlighted), and 'Datacenter'. At the bottom, it indicates 'Mesosphere DCOS v.1.6' and an 'Open Command Line' button.

Services



38 Total Services

Filtering controls: All (38), Healthy (24), Sick (7), and a search filter box.

A grid of service cards. Each card includes a checkbox, a service icon, the service name, health status, task count, and version. The 'marathon-production' card is selected with a checkmark.

Service	Health	Tasks	Version
marathon-development	Healthy	37	Marathon v.1.0.2
marathon-production	Healthy	14	Marathon v.1.0.2
spark-financial-forecasin...	Sick	401	Spark v.1.2.1
kubernetes-p	Healthy	56	Kubernetes

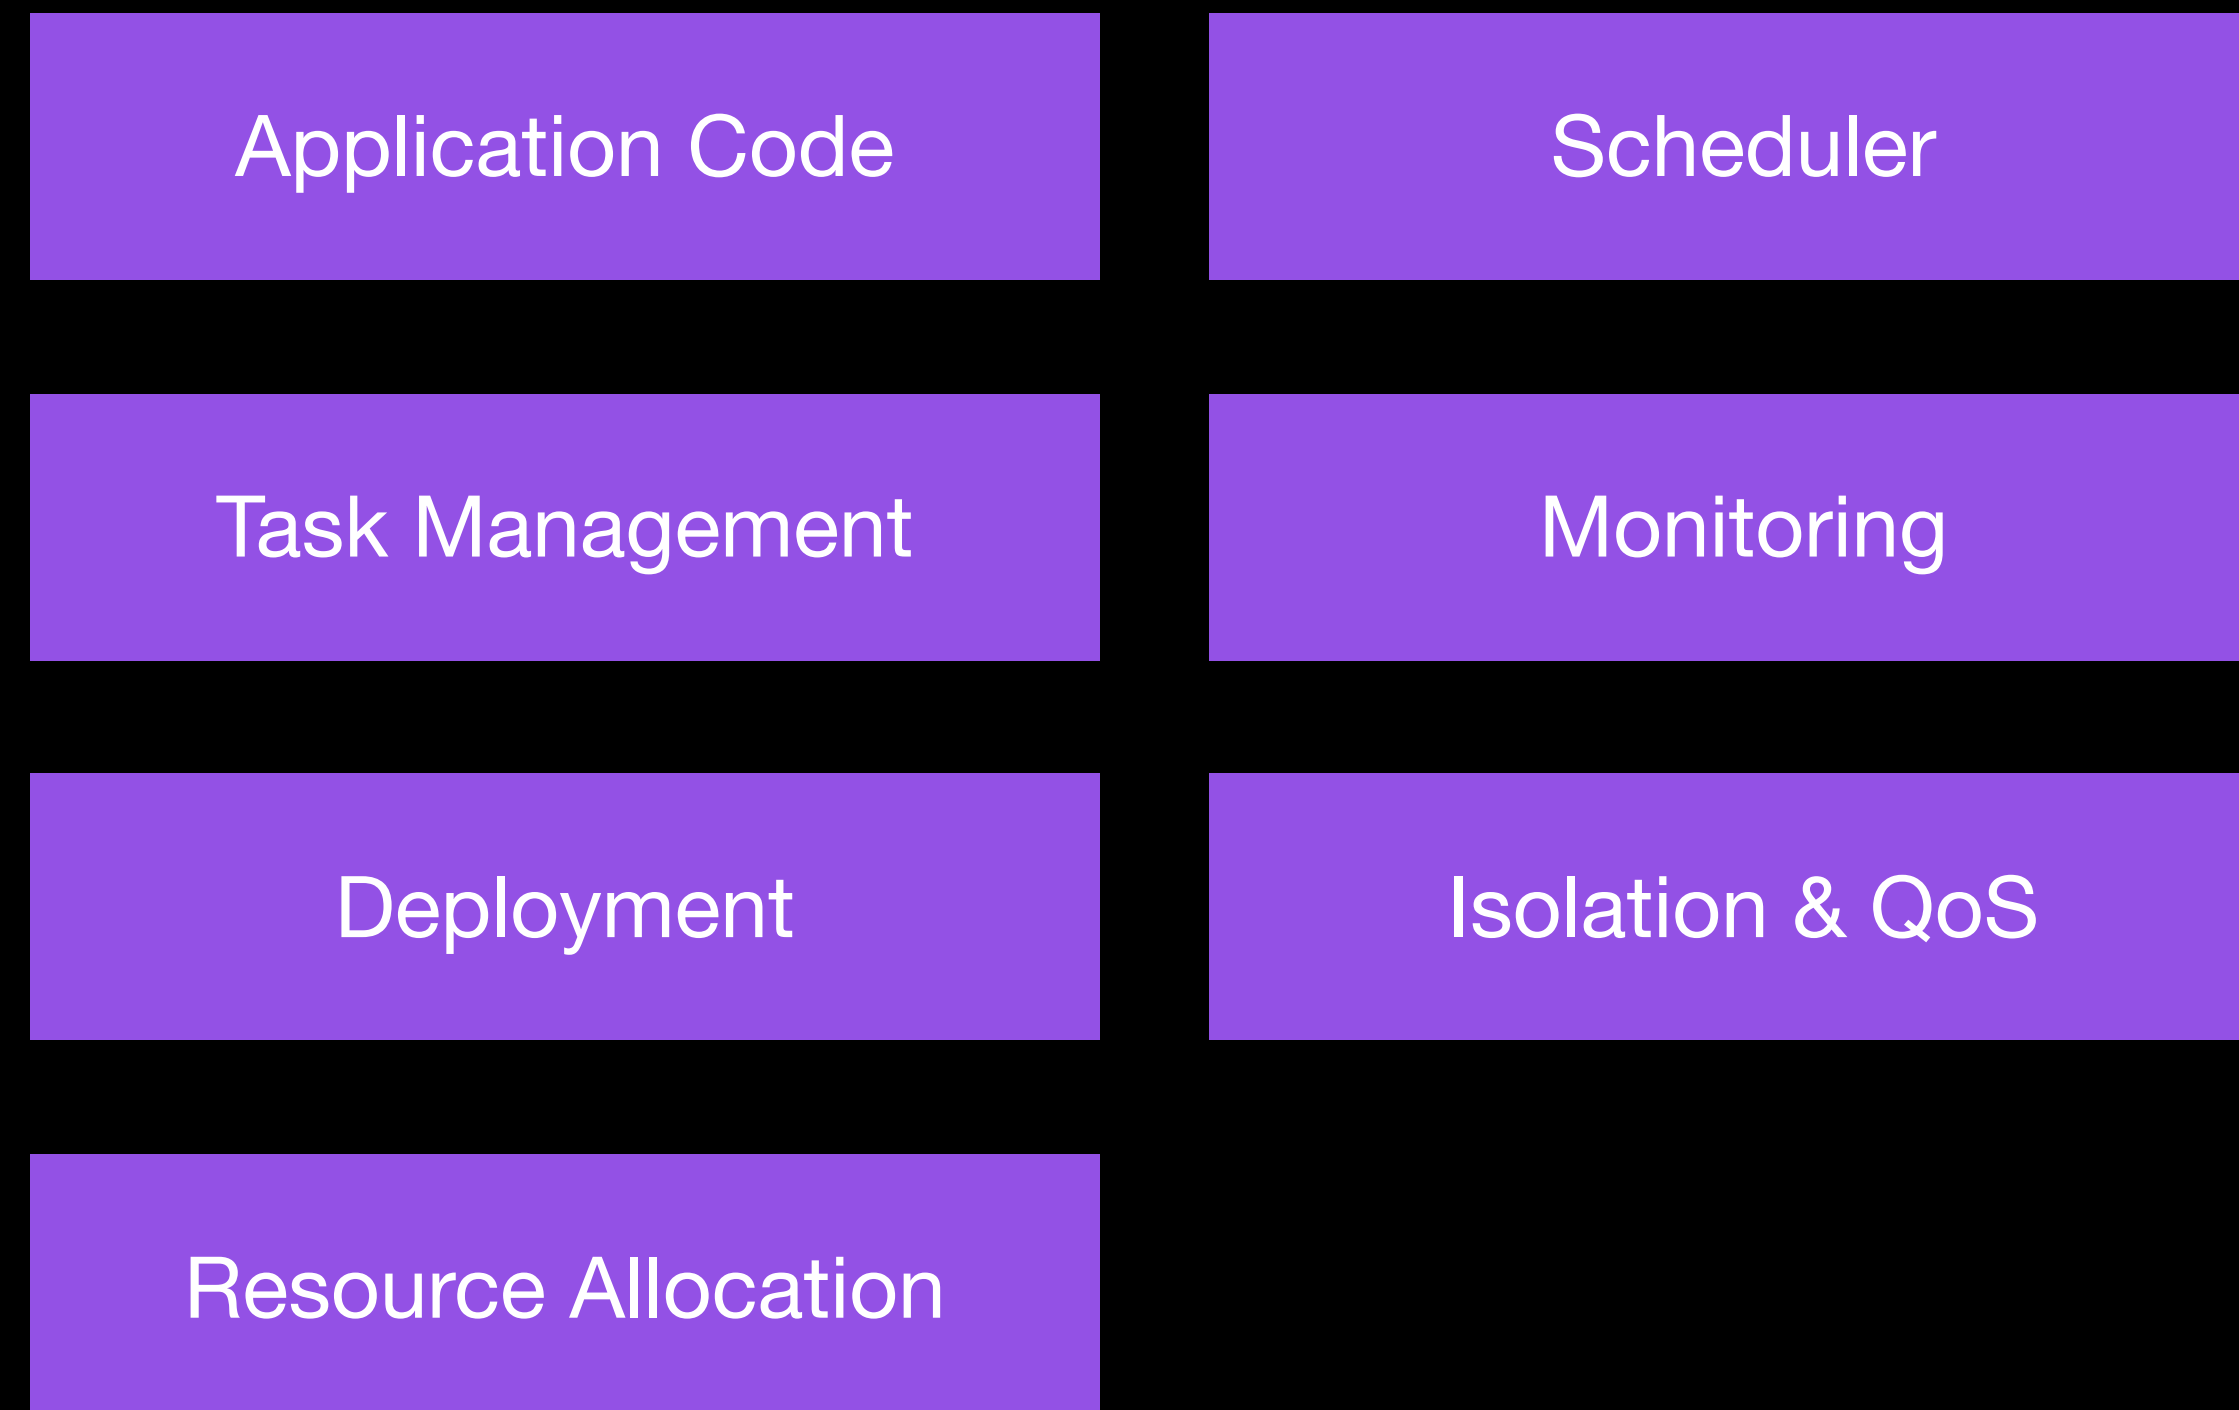
The Datacenter Developer

What is a Datacenter Developer?

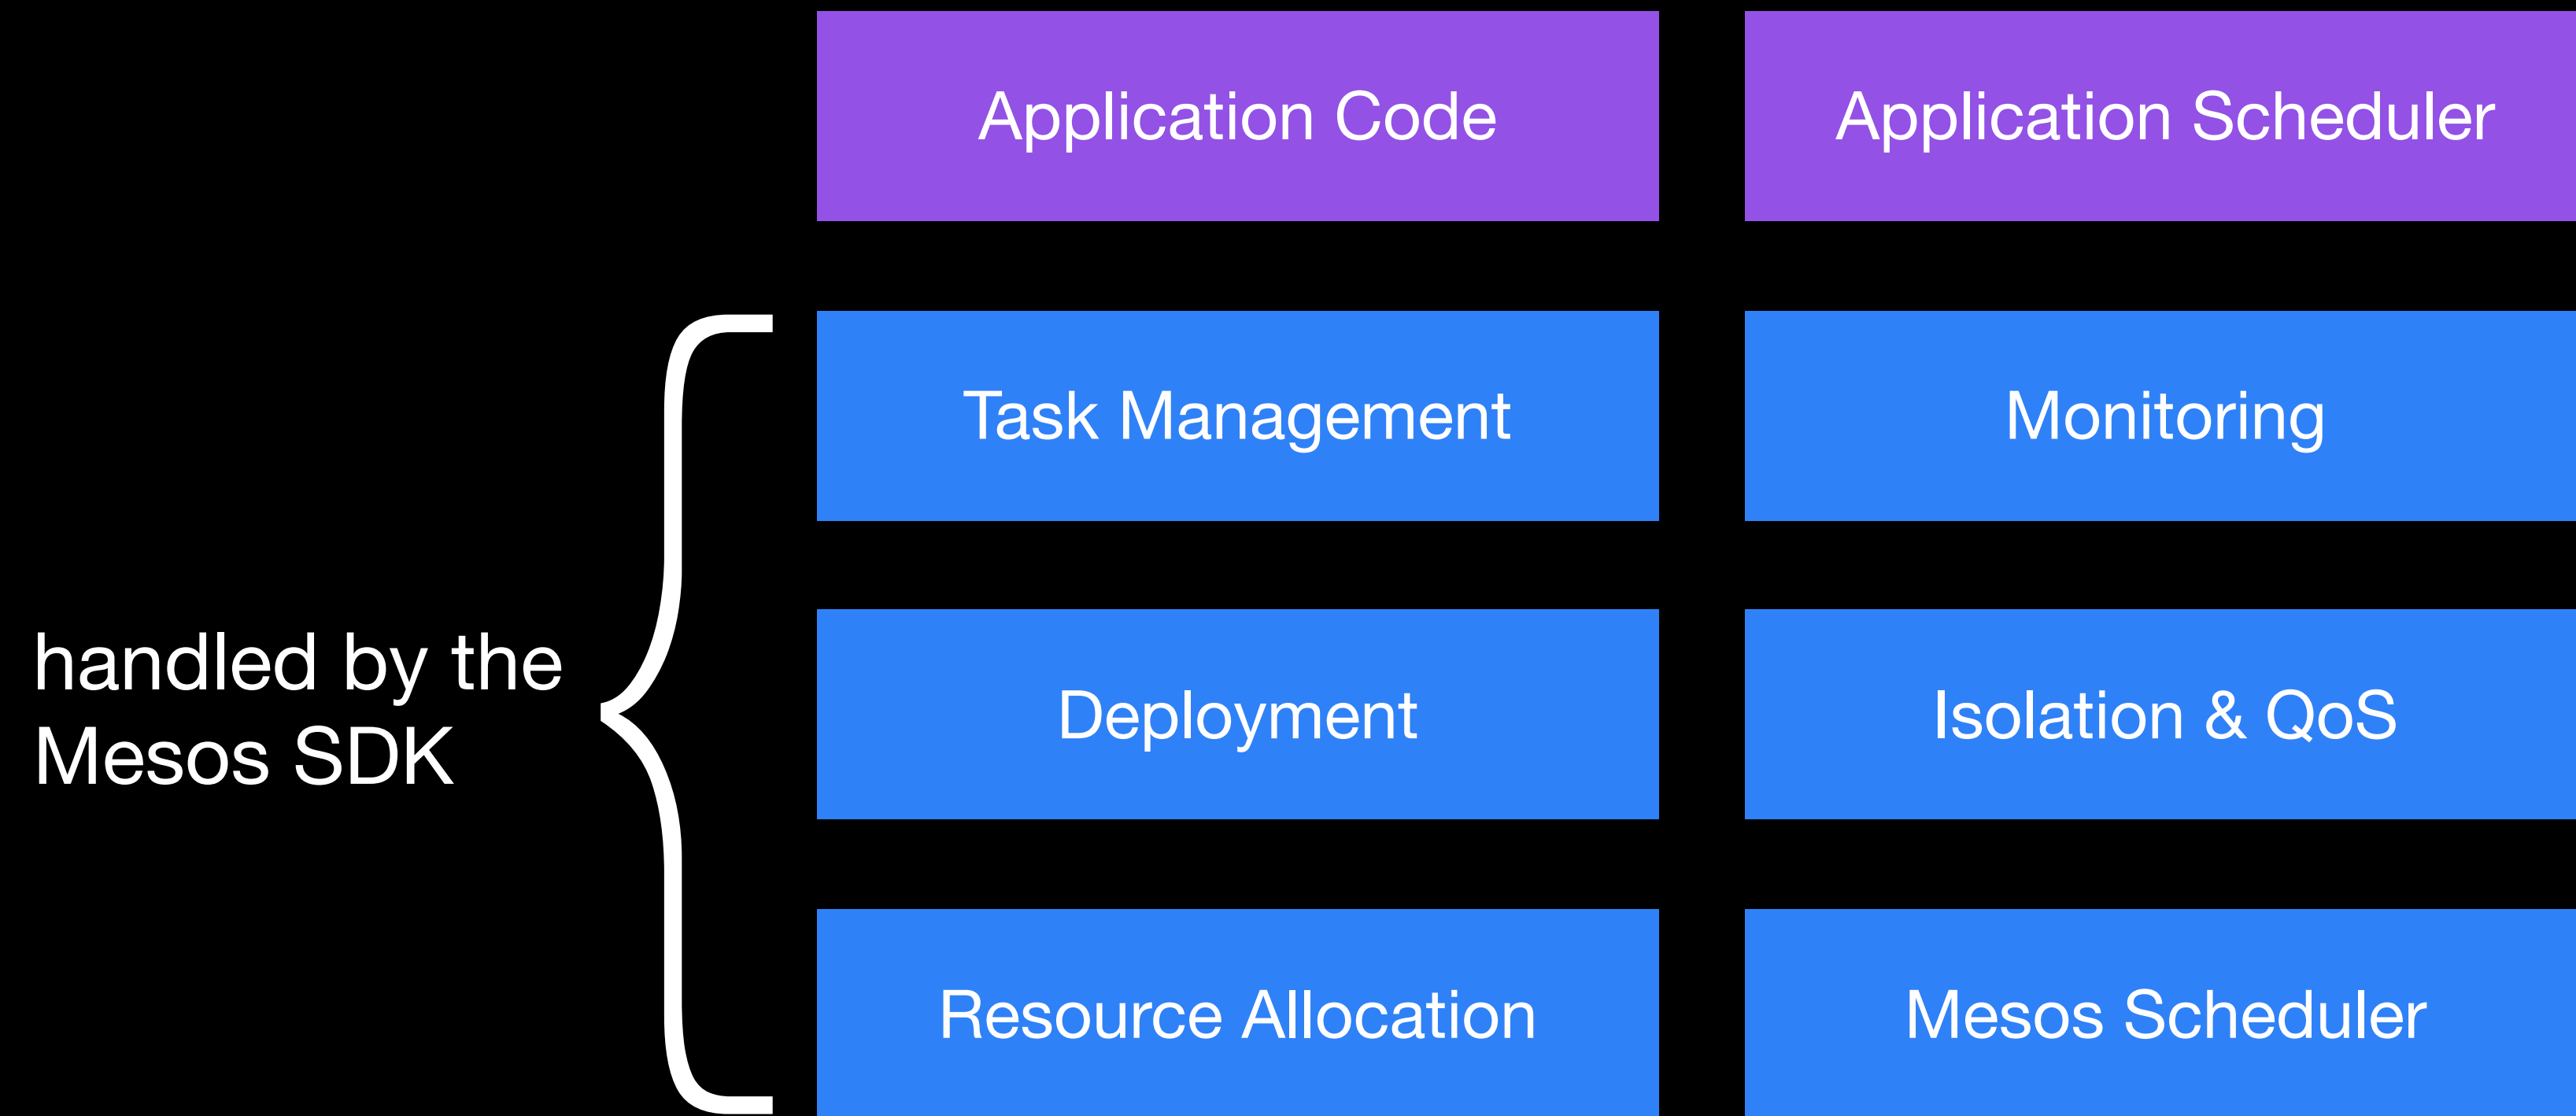
A developer who uses a datacenter SDK to build distributed systems that can dynamically leverage all the resources available in a datacenter.

Building Distributed Systems

Without Mesos



With Mesos



A Mesos Datacenter Developer...

... uses Containers to isolate tasks

Linux: process isolation

private address space, file descriptors, e.g stdout

Mesos: linux containers

Time share of CPU, private memory/pid/user namespace, chroot

Disk and network QoS

A Mesos Datacenter Developer...

... can populate cluster nodes dynamically

No need to provision new nodes or write recipes

Mesos fetches immutable artifacts via HTTP/HDFS or Docker

Can also be used to deploy config, models, ...

A Mesos Datacenter Developer...

... gets task lifecycle management

Mesos manages task distribution: starting, monitoring, killing, cleanup

Metadata

A Mesos Datacenter Developer...

... can easily respond to failures

Machine and network failure detection

Health checks: HTTP, TCP, or custom program

A Mesos Datacenter Developer...

... gets notified of failed tasks

```
TaskStatus {  
  task_id: "dispatch.45b70272-d88c-11e4-98ef-a689777343be",  
  state: TASK_FAILED,  
  source: SOURCE_SLAVE,  
  reason: REASON_COMMAND_EXECUTOR_FAILED,  
  message: "Command exited with status 1"  
}
```


A Mesos Datacenter Developer...

... gets notified when tasks run out of memory

```
TaskStatus {  
  task_id: "dispatch.45b70272-d88c-11e4-98ef-a689777343be",  
  state: TASK_FAILED,  
  source: SOURCE_SLAVE,  
  reason: REASON_MEMORY_LIMIT,  
  message: "Task exceeded its memory limit" // TODO: MESOS-2035  
}
```

A Mesos Datacenter Developer...

... gets notified of failing machines

```
TaskStatus {  
  task_id: "dispatch.45b70272-d88c-11e4-98ef-a689777343be",  
  state: TASK_LOST,  
  source: SOURCE_MASTER,  
  reason: REASON_SLAVE_DISCONNECTED,  
  message: "Slave disconnected"  
}
```

A Mesos Datacenter Developer...

... gets easy access to logs

```
$ mesos tail -f poseidon.*
```

```
==>poseidon.4729069a-d87d-11e4-8cee-f20d55470486:stdout<==
```

```
[2015-04-07 19:37:23,842] INFO 104.236.41.240 - - [07/Apr/2015:19:37:23 +0000] "GET /ping HTTP/1.1"
```

```
[2015-04-07 19:37:33,915] INFO 104.236.41.240 - - [07/Apr/2015:19:37:33 +0000] "GET /ping HTTP/1.1"
```

```
[2015-04-07 19:37:43,877] INFO 104.236.41.240 - - [07/Apr/2015:19:37:43 +0000] "GET /ping HTTP/1.1"
```

```
[2015-04-07 19:38:44,037] INFO 104.236.41.240 - - [07/Apr/2015:19:38:43 +0000] "GET /ping HTTP/1.1"
```

```
[2015-04-07 19:38:54,004] INFO 104.236.41.240 - - [07/Apr/2015:19:38:53 +0000] "GET /ping HTTP/1.1"
```

```
==>poseidon.15bb96b9-d63a-11e4-93c0-f6e64c94ec3c:stdout<==
```

```
[2015-04-07 19:37:23,905] INFO 104.236.41.240 - - [07/Apr/2015:19:37:23 +0000] "GET /ping HTTP/1.1"
```

```
[2015-04-07 19:37:33,891] INFO 104.236.41.240 - - [07/Apr/2015:19:37:33 +0000] "GET /ping HTTP/1.1"
```

```
[2015-04-07 19:37:43,924] INFO 104.236.41.240 - - [07/Apr/2015:19:37:43 +0000] "GET /ping HTTP/1.1"
```

```
[2015-04-07 19:38:44,022] INFO 104.236.41.240 - - [07/Apr/2015:19:38:43 +0000] "GET /ping HTTP/1.1"
```

```
[2015-04-07 19:38:54,042] INFO 104.236.41.240 - - [07/Apr/2015:19:38:53 +0000] "GET /ping HTTP/1.1"
```

A Mesos Datacenter Developer...

... gets powerful resource monitoring

```
$ curl http://localhost:5051/monitor/statistics.json
{
  "statistics": {
    "cpus_limit": 0.35,
    "cpus_nr_periods": 520883,
    "cpus_nr_throttled": 2163,
    "cpus_system_time_secs": 154.42,
    "cpus_throttled_time_secs": 145.96,
    "cpus_user_time_secs": 258.74,
    "mem_anon_bytes": 109137920,
    "mem_file_bytes": 30613504,
    "mem_limit_bytes": 167772160,
    "mem_mapped_file_bytes": 8192,
```

A Mesos Datacenter Developer...

... can run her apps next to the data

Data-heavy applications can run where the data is stored.

A Mesos Datacenter Developer...

... programs against a whole datacenter

Virtual memory: no longer program against physical addresses

Mesos: no longer program against machines, but the whole datacenter

Apps can allocate more resources on existing machines or populate new ones via an API

A Mesos Datacenter Developer...

... can build her own cluster scheduler easily

Linux: Completely Fair Scheduler (CFS) for process scheduling, maximizes utilization and performance

Mesos: Dominant Resource Fairness (DRF), two level scheduler

Provides fair resource allocation while allowing applications to implement their own specialized schedulers

A Mesos Datacenter Developer...

... has more time to focus on her application

Scheduler, Executor, Task, FrameworkMessage are high level APIs that abstract the plumbing.

Developers only need to know which resources they need, not where they are.

Developers can focus on their application instead of solving hard low level distributed systems problems.

Projects built on top of Mesos

Apache Projects ported to Mesos

Storm

Cassandra

Kafka

Myriad (YARN)

Other Projects ported to Mesos

Jenkins

Kubernetes

ElasticSearch

Projects based on Mesos

Apache Aurora

Apache Spark

ArangoDB

Condor

Crate.IO

Marathon

Chronos

An Example

Twitter

Anatomy of a Web App

Millions of users

Lots of data

Data products



Tobi Knaup
@superguenter

TWEETS **1,011** FOLLOWING **600** FOLLOWERS **1,314**

Who to follow · Refresh · View all

Peter Schuller @scode ×
[+ Follow](#)

Mats Linander @matslina ×
[+ Follow](#)

[Popular accounts](#) · [Find friends](#)

Trends · Change

- #NationalBeerDay
- Rand Paul
- HBO Now
- #WorldHealthDay
- The Bold Italic
- #BadChoiceFuneralSongs
- Deus Ex
- #ScotDebates
- #NCAACHampionship
- Arrested Development

What's happening?

[View 2 new Tweets](#)

Alex Williams @alexwilliams · 2m
Congrats to @derrickharris, formerly of @gigaom, who is joining @mesosphere as a senior research analyst.
[↩](#) [↻](#) [★](#) [⋮](#)

Zed @zedshaw · 6m
The Bold Italic folds [thebolditalic.com/articles/7229-...](#) Next up, The Comic Sans magazine about Duluth, and The Mono Arial about Topeka.
[↩](#) [↻](#) 1 [★](#) [⋮](#) [View summary](#)

Dennis Goedegebuure @TheNextCorner · 7m
Rare Sunrise Blood Moon This Weekend [bit.ly/1NFO3pg](#)
[↩](#) [↻](#) [★](#) [⋮](#) [View summary](#)

Interstate 80 Tahoe @i80chains · 7m
OPEN: Trucks Screened (14:32 04/07)
[↩](#) [↻](#) [★](#) [⋮](#)

Nicholas Weaver @lynxbat · 9m
I'm not smart enough for this room.
[↩](#) [↻](#) [★](#) [⋮](#)

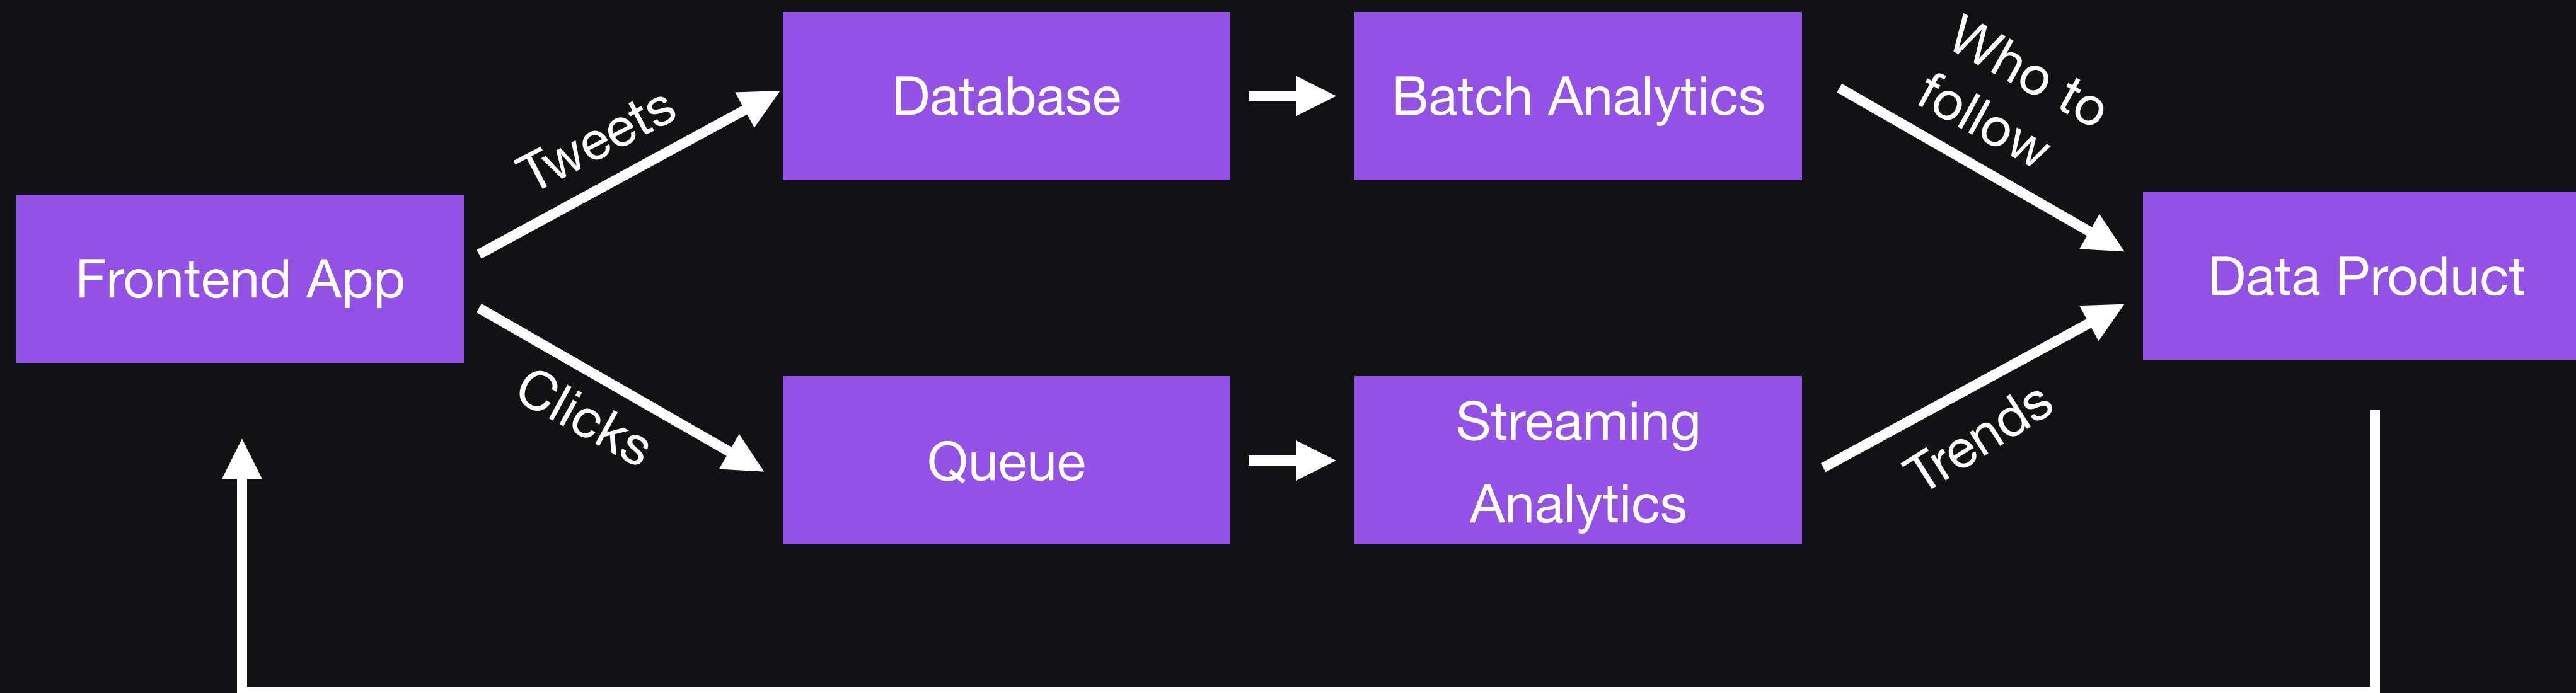
Zed retweeted

Sarah N. Emerson @SarahNEmerson · 23m
The dystopian FBI document used to keep spy technology hidden [motherboard.vice.com/read/the-dysto...](#)

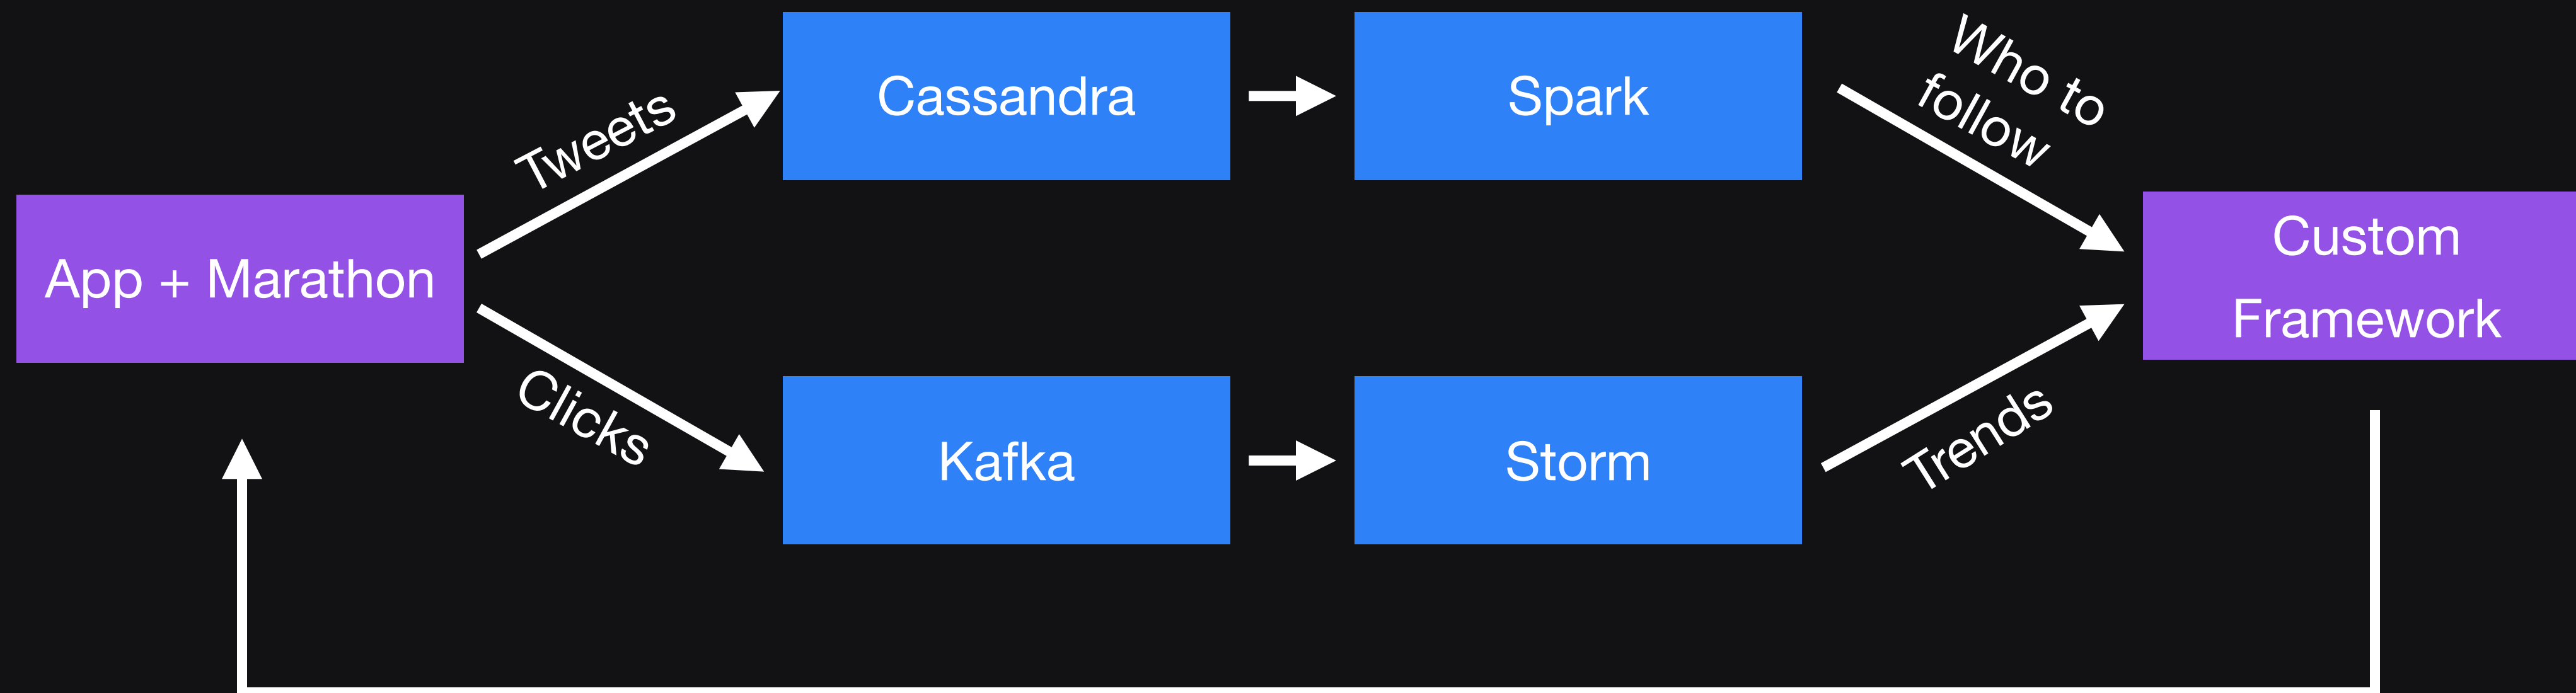
Police are not allowed to talk about the technology in press releases, in court documents, during judicial hearings, or during other public forums or proceedings."

- If a district prosecutor intends to use details about how the Stingray works during a court case, police must alert the FBI "in order to allow sufficient time for

Anatomy of a Web App



Anatomy of a Web App - entirely on Mesos



Become a Datacenter Developer!

Mesos Docs: <https://mesos.apache.org/documentation/latest/>

RENDERER: <https://github.com/mesosphere/RENDERER>

Go: <https://github.com/mesos/mesos-go>

Akka/Scala: <https://github.com/drexin/akka-mesos>

Clojure: <https://github.com/dgrnbrg/clj-mesos>