

# High availability services

---

Elliott Sprehn (Google, Inc.)

[esprehn@gmail.com](mailto:esprehn@gmail.com)

<http://www.elliottsprehn.com/>

# Scalable ~~High availability~~ services

---


Elliott Sprehn (Google, Inc.)

[esprehn@gmail.com](mailto:esprehn@gmail.com)

<http://www.elliotsprehn.com/+/>

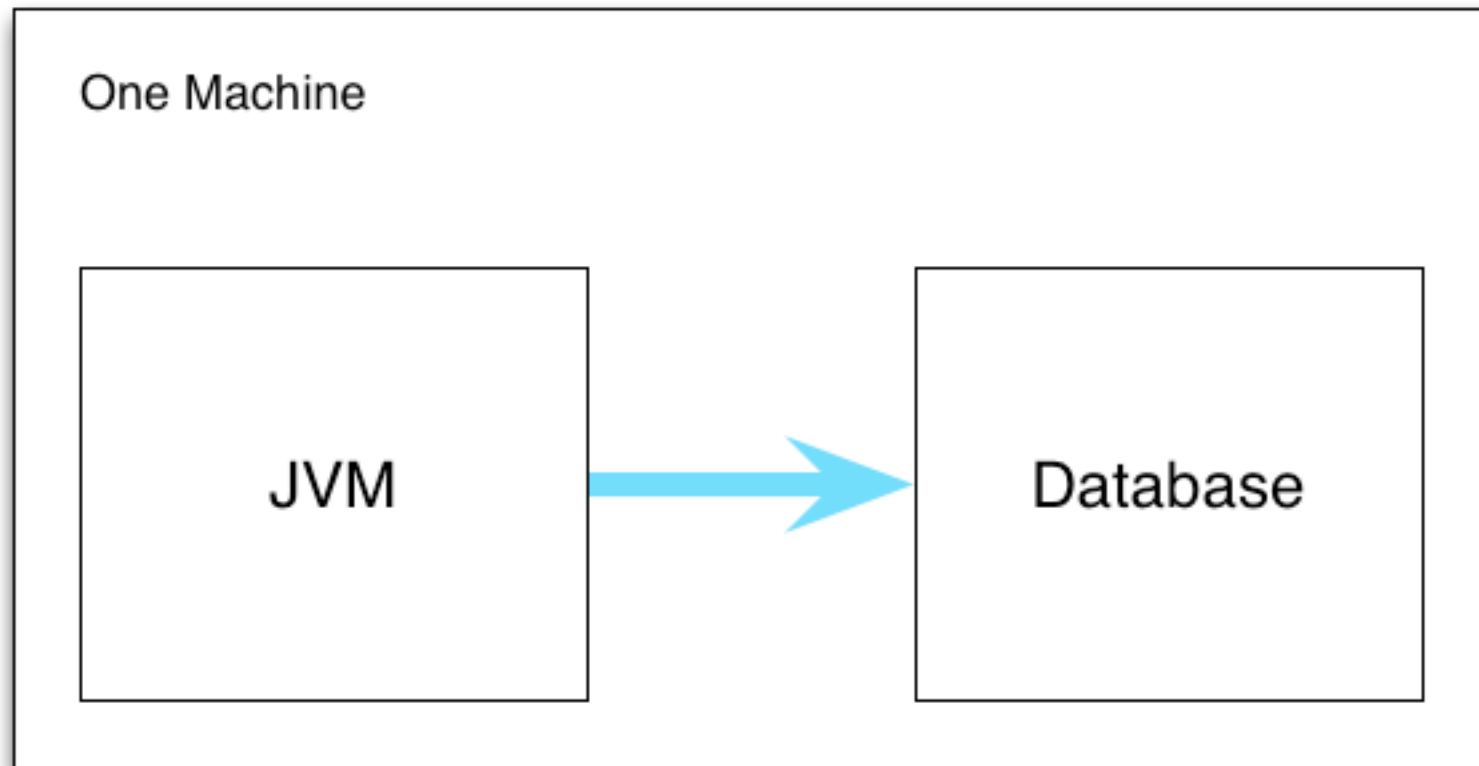
# Evolution of an application

---



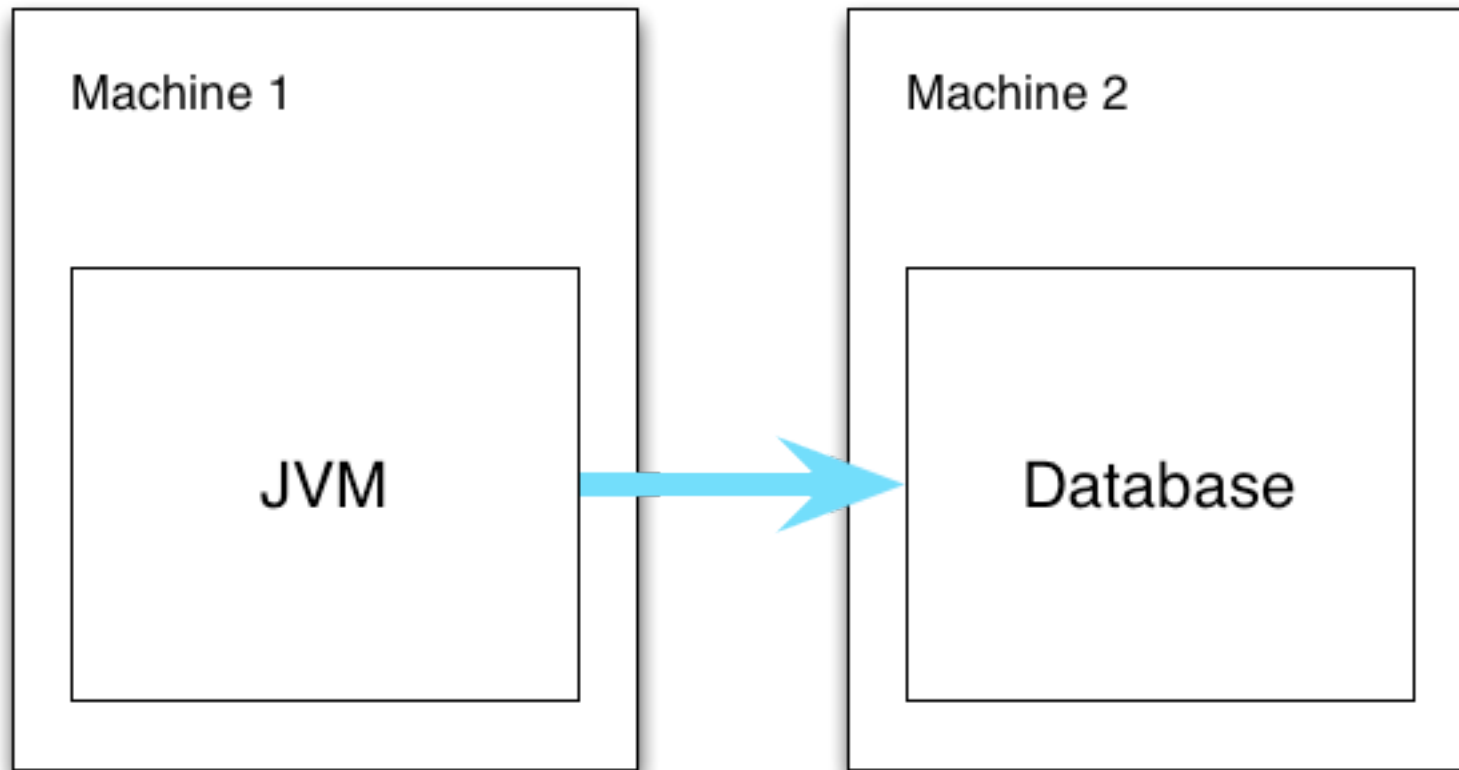
# Evolution of an application

---



# Evolution of an application

---



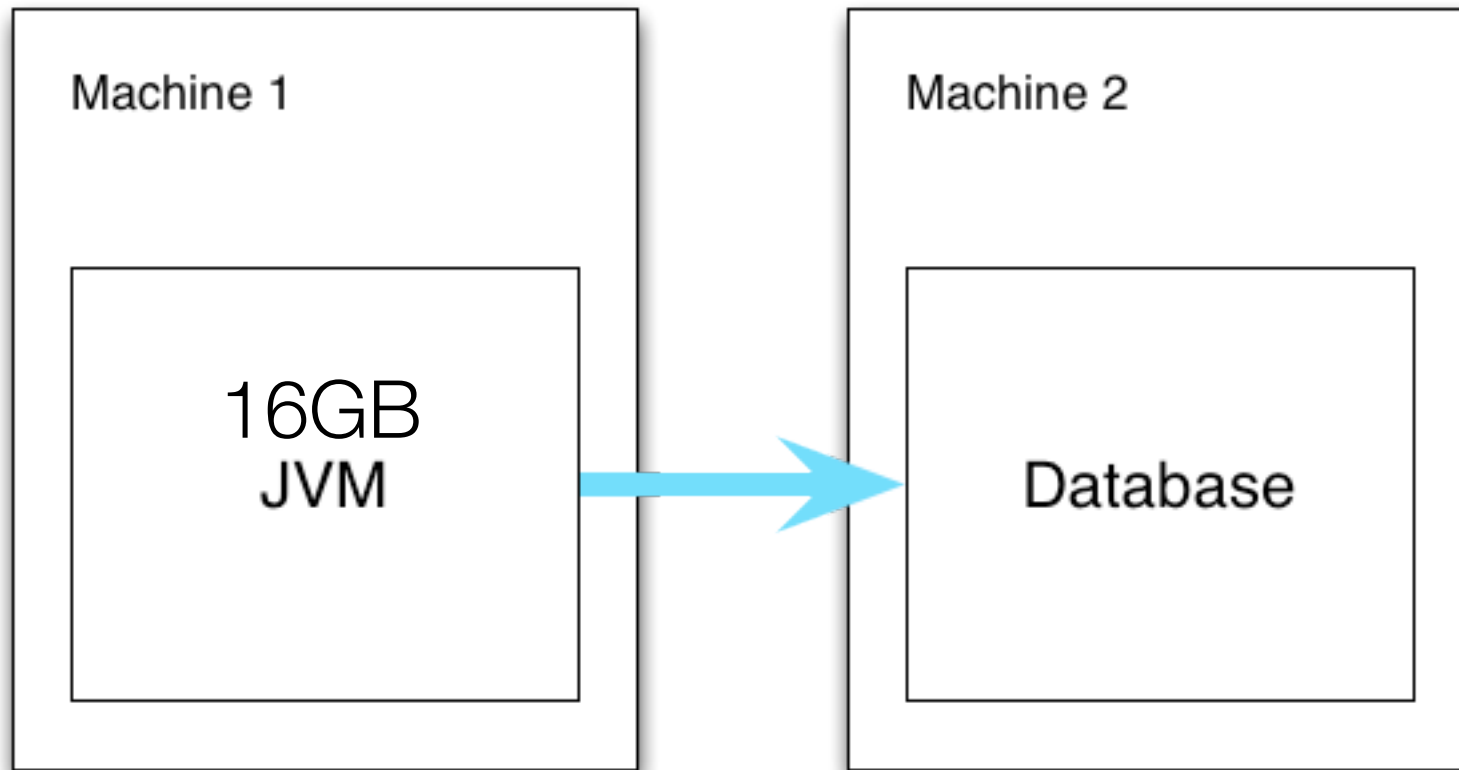
# Evolution of an application

---

Add more RAM.


# Evolution of an application

---



# Evolution of an application

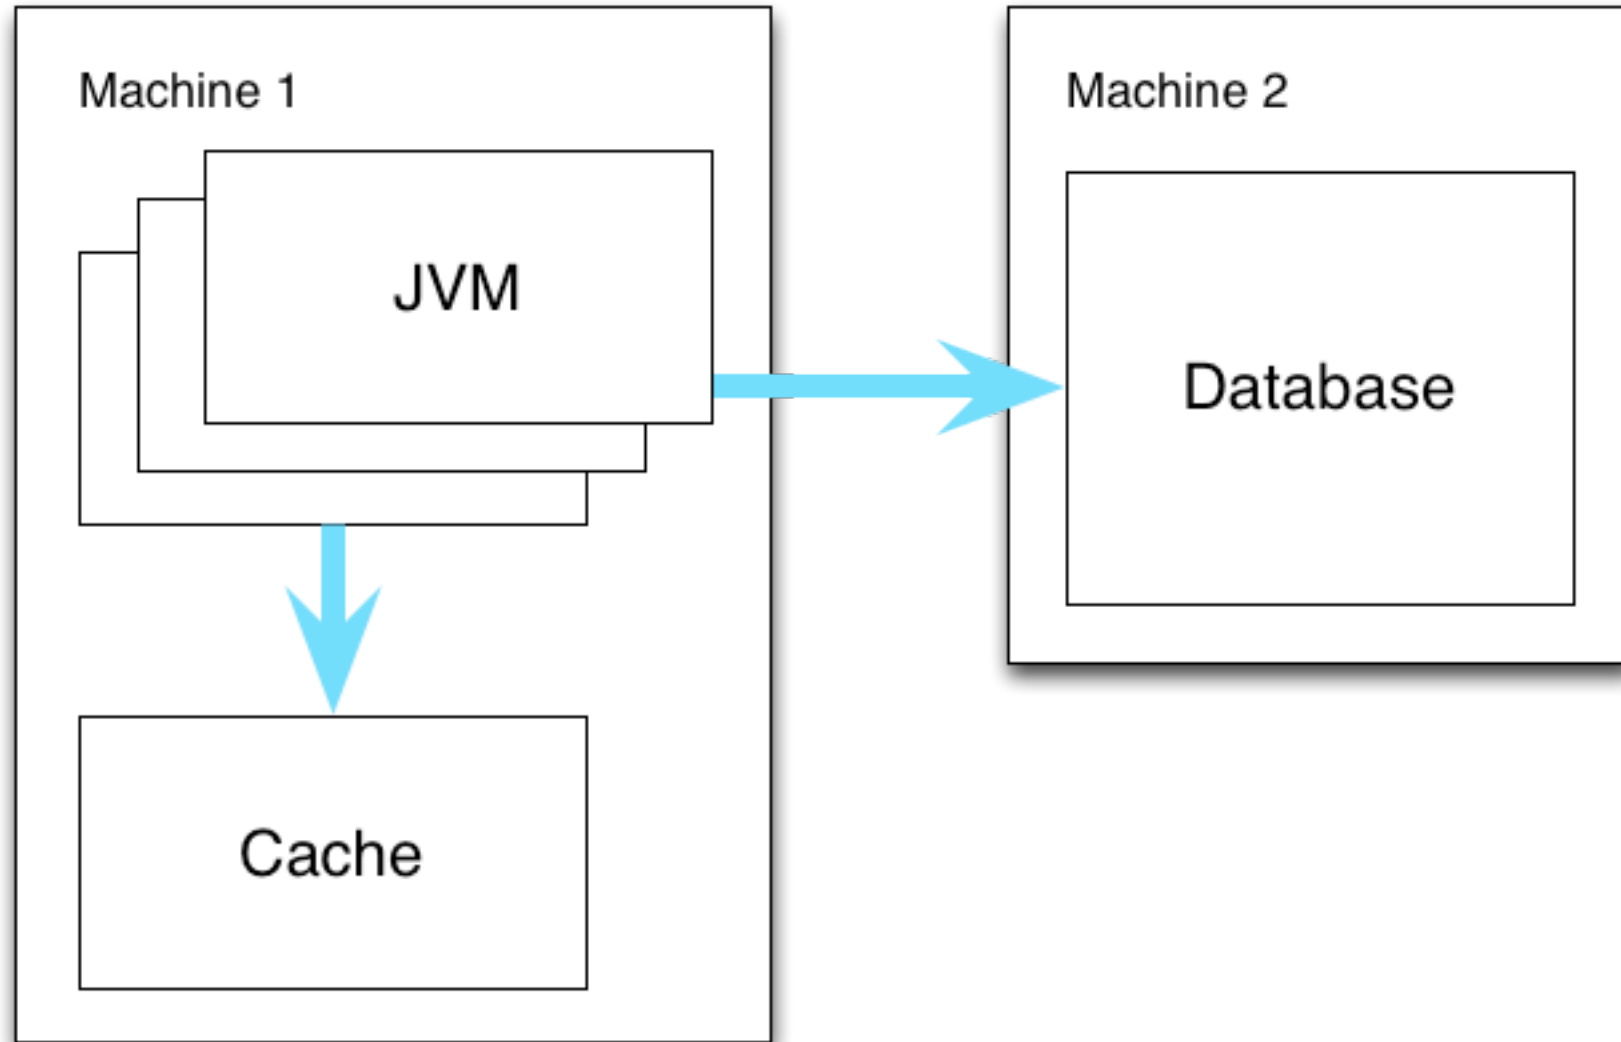
---






# Evolution of an application

---



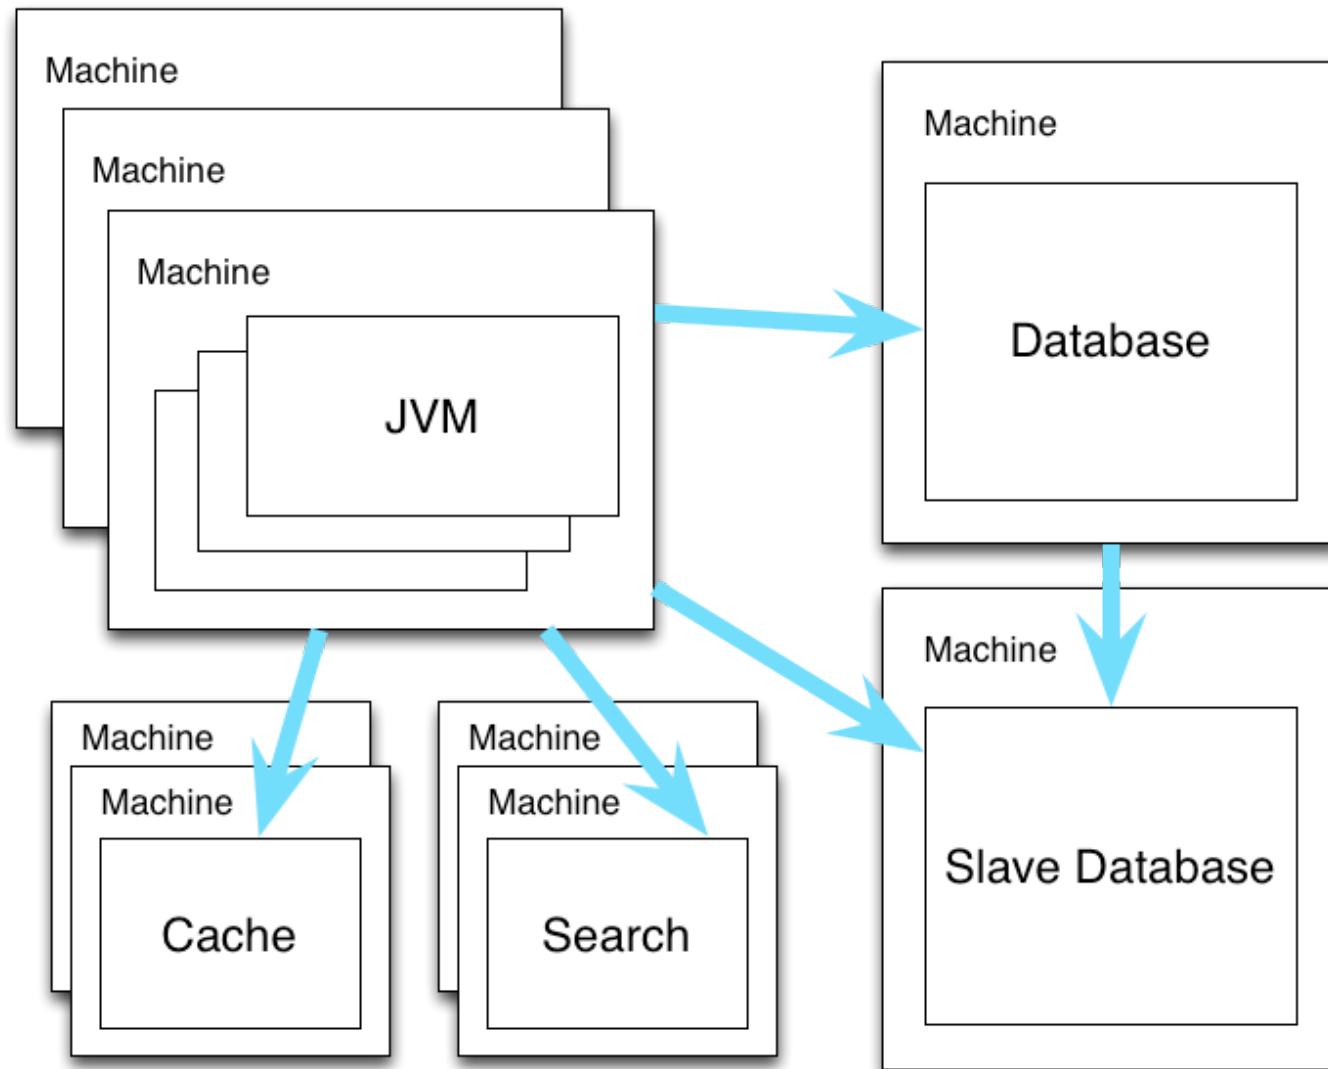
# Evolution of an application

---



# Evolution of an application

---



# Evolution of an application

---

One big binary.

## Change your mindset

---

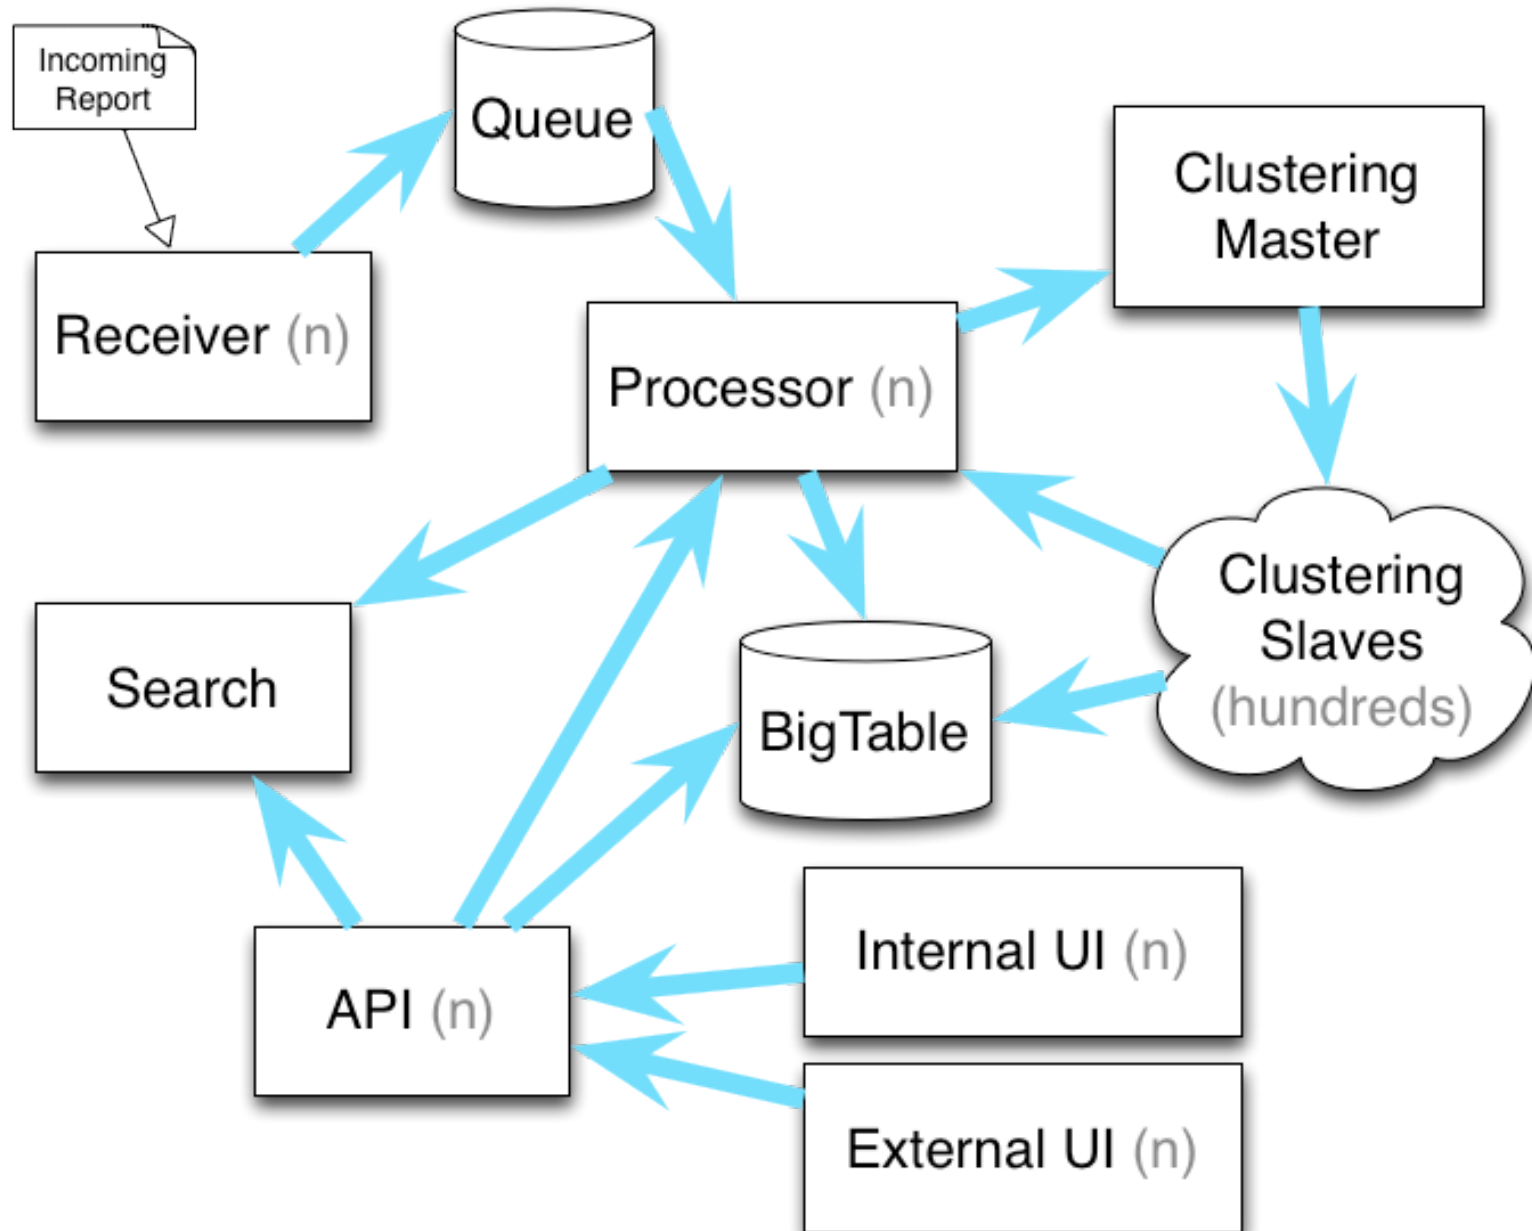
- Many small binaries
- Many small JVMs
- Connected with RPC

# Google Feedback

---

- 4 data centers
- 8+ types of jobs
- Processing queue

# Google Feedback



# Issues

---

- Delay in processing queue
- Single points of failure
- Master election



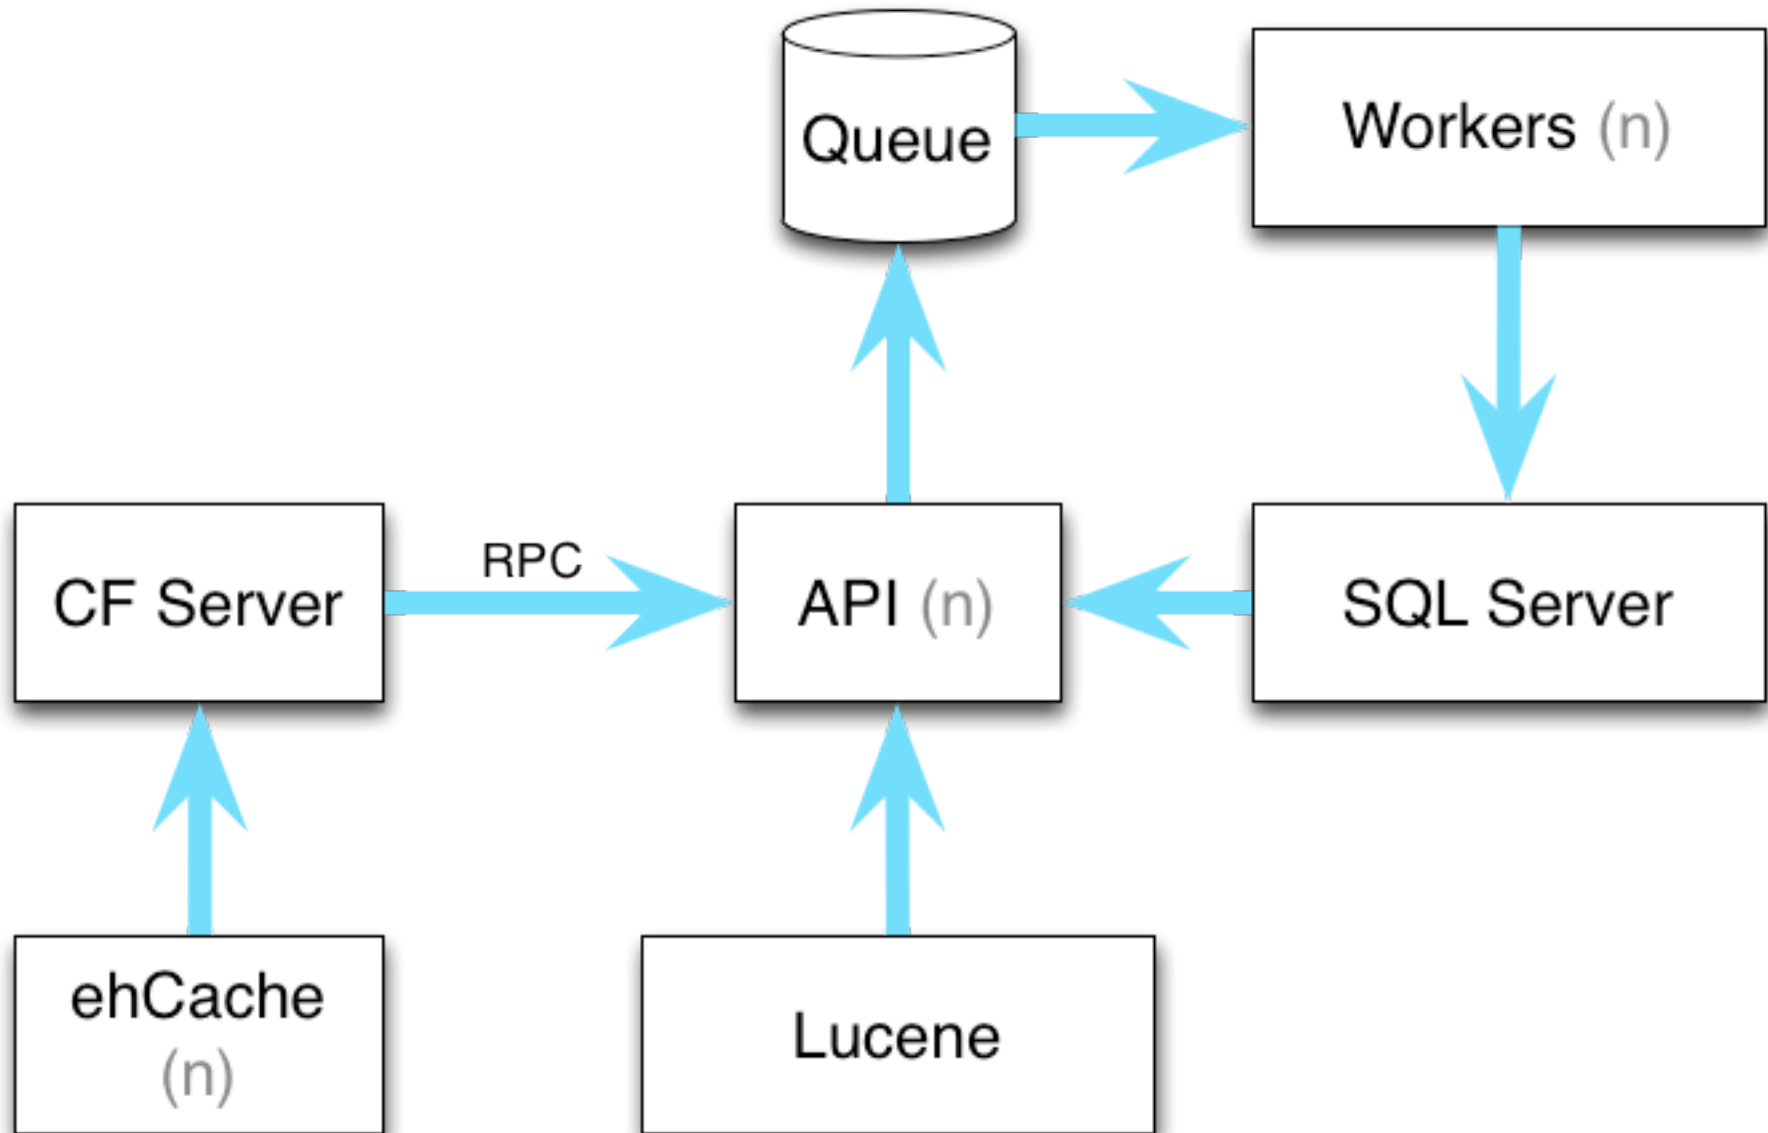
# Scalable ColdFusion

---

- Many small binaries
- Many small JVMs
- Connected with RPC

# Scalable ColdFusion

---



## Purpose built jobs

---

- Async processing queue
- Many small JVMs
- End request as fast as possible

# Kinds of jobs

---

- Virus scanning
- Image processing
- Data transforms
- Index updates

# Pros

---

- Reliability
- Scalability
- Performance (distributed latency)

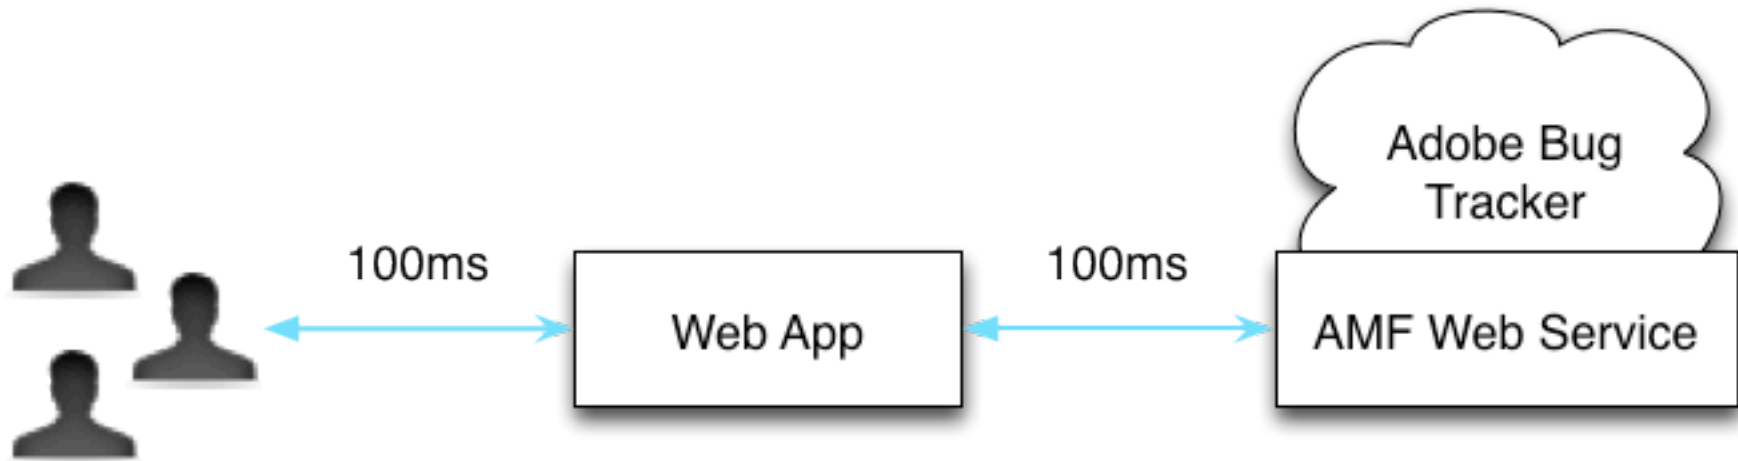
## Cons

---

- Complicated push
- Complicated failures
- Relaxed consistency

# CFBug Mirror

---



# FreshBot

---

- Search engine spider
- 3 modes
  - Full
  - Recent
  - Freshness

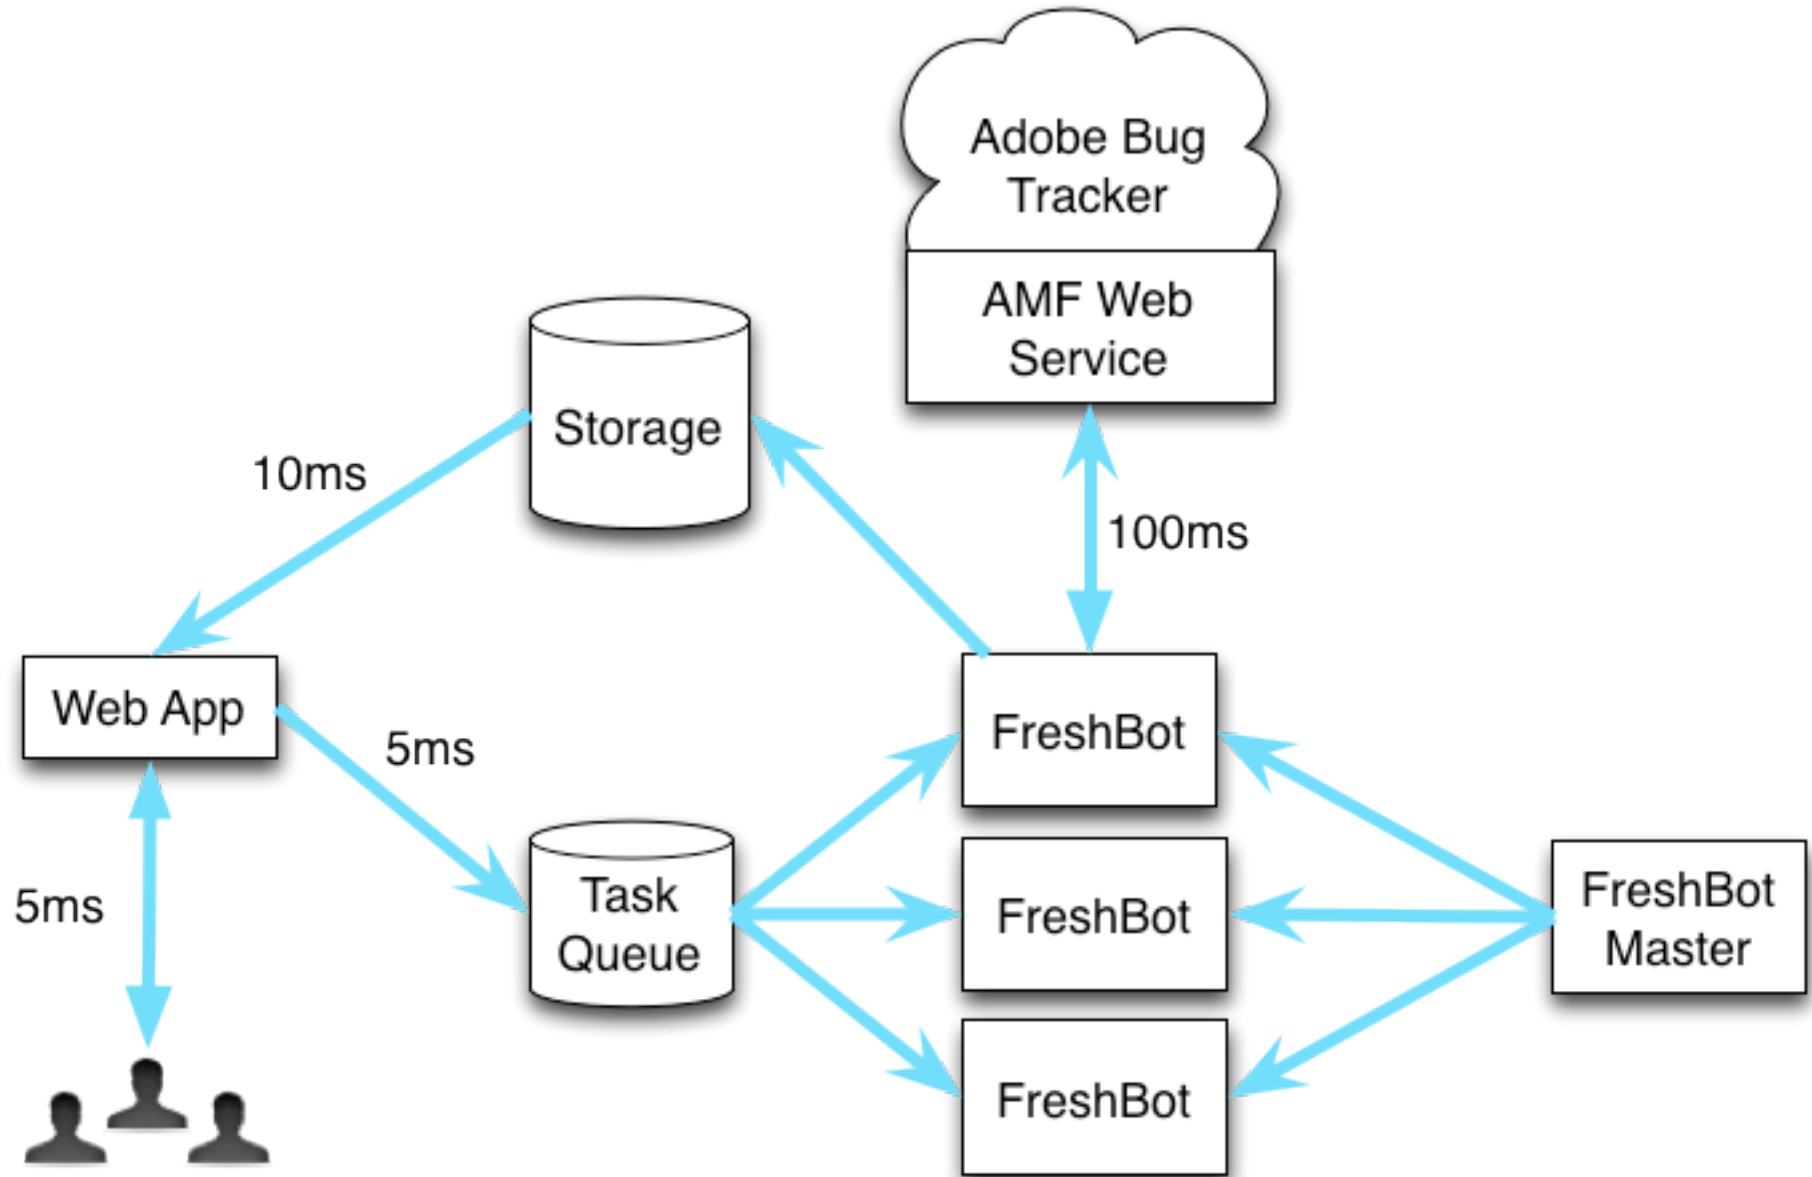


# FreshBot indexing

---

- Full
  - 1-2 hours (once a day)
- Recent
  - 200ms (every n minutes)
- Freshness
  - 100ms (item) + 100ms (page)

# CFBug Mirror (Freshbot)



vs 1 JVM and many threads (and RAM)

---

- Bad failures modes
- Slow restarts

vs 1 JVM and work queue in memory

---

- Bad failures modes
- Consistency

## Benefits

---

Use the right  
language.

## Pro tip: Rolling updates

---

- Avoid downtime
- Never show  
**“We’re down for maintenance”**

Pro tip: Add lots of logging

---

Debugging  
distributed  
systems is hard.

# Questions

---

