


# High Availability at Heroku

Mark McGranaghan



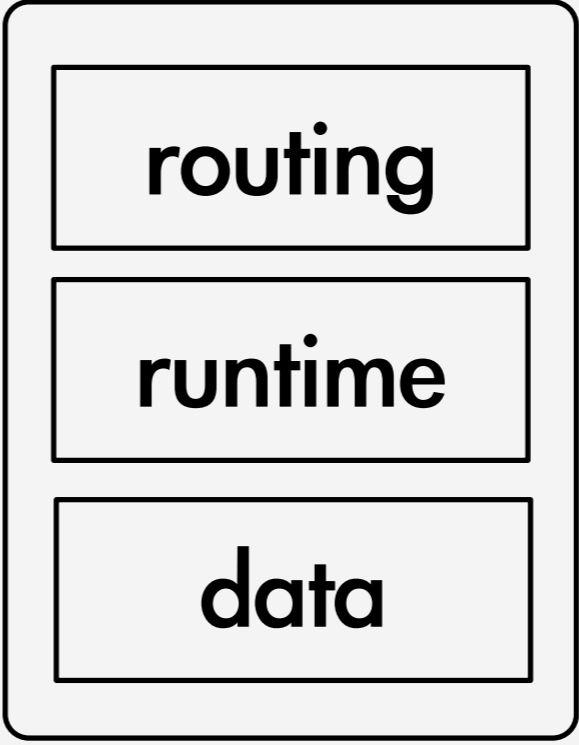



heroku

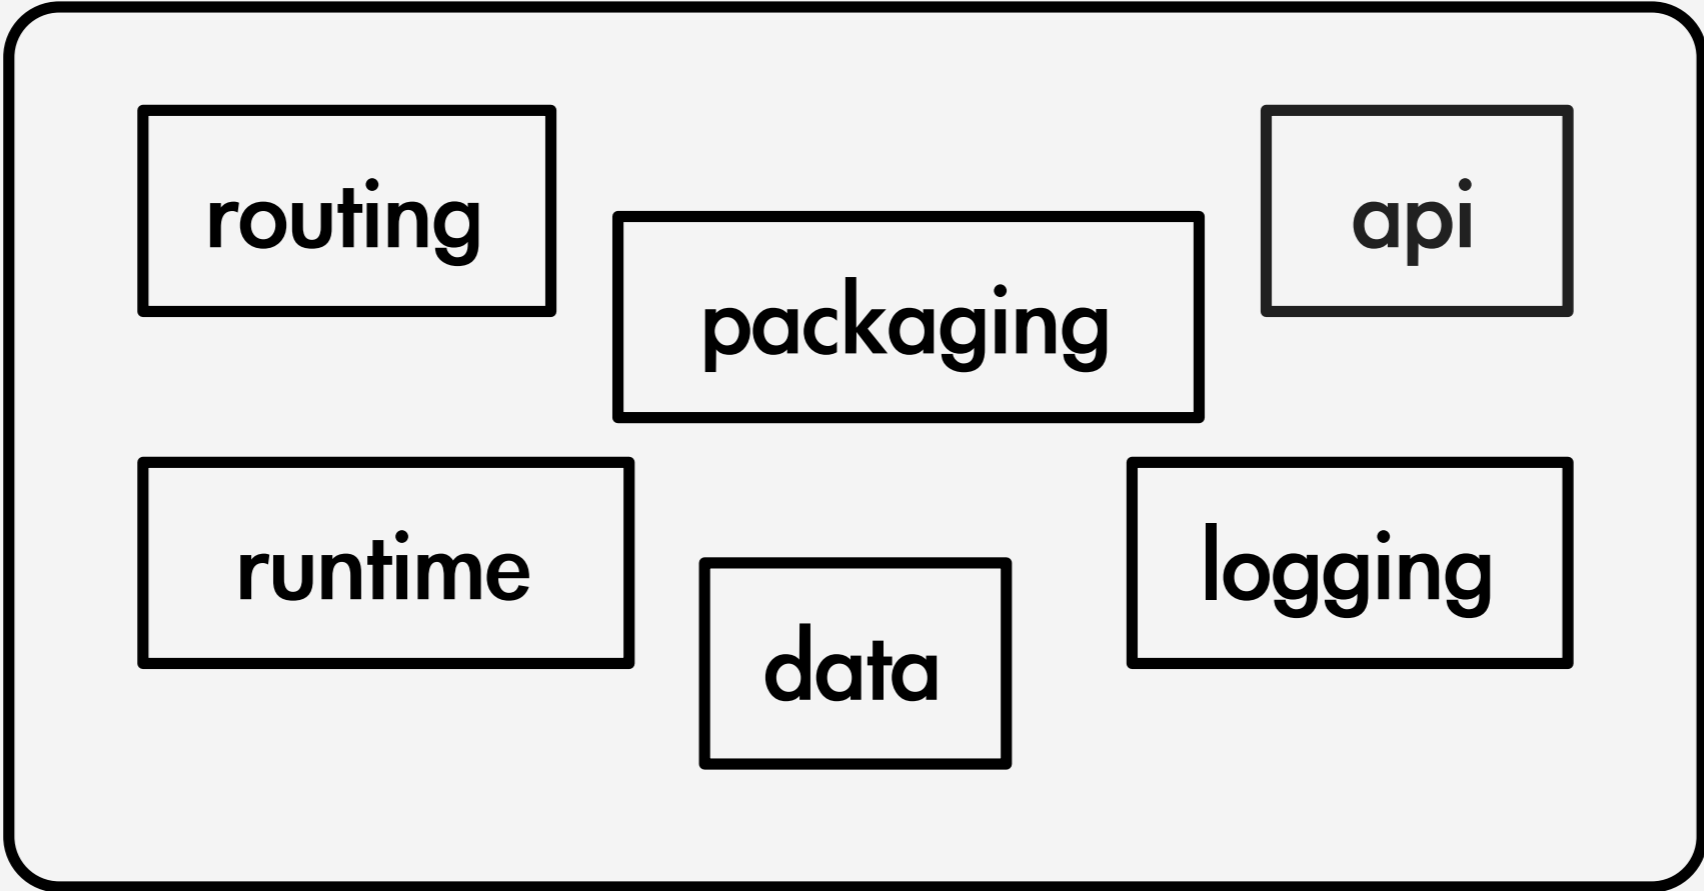
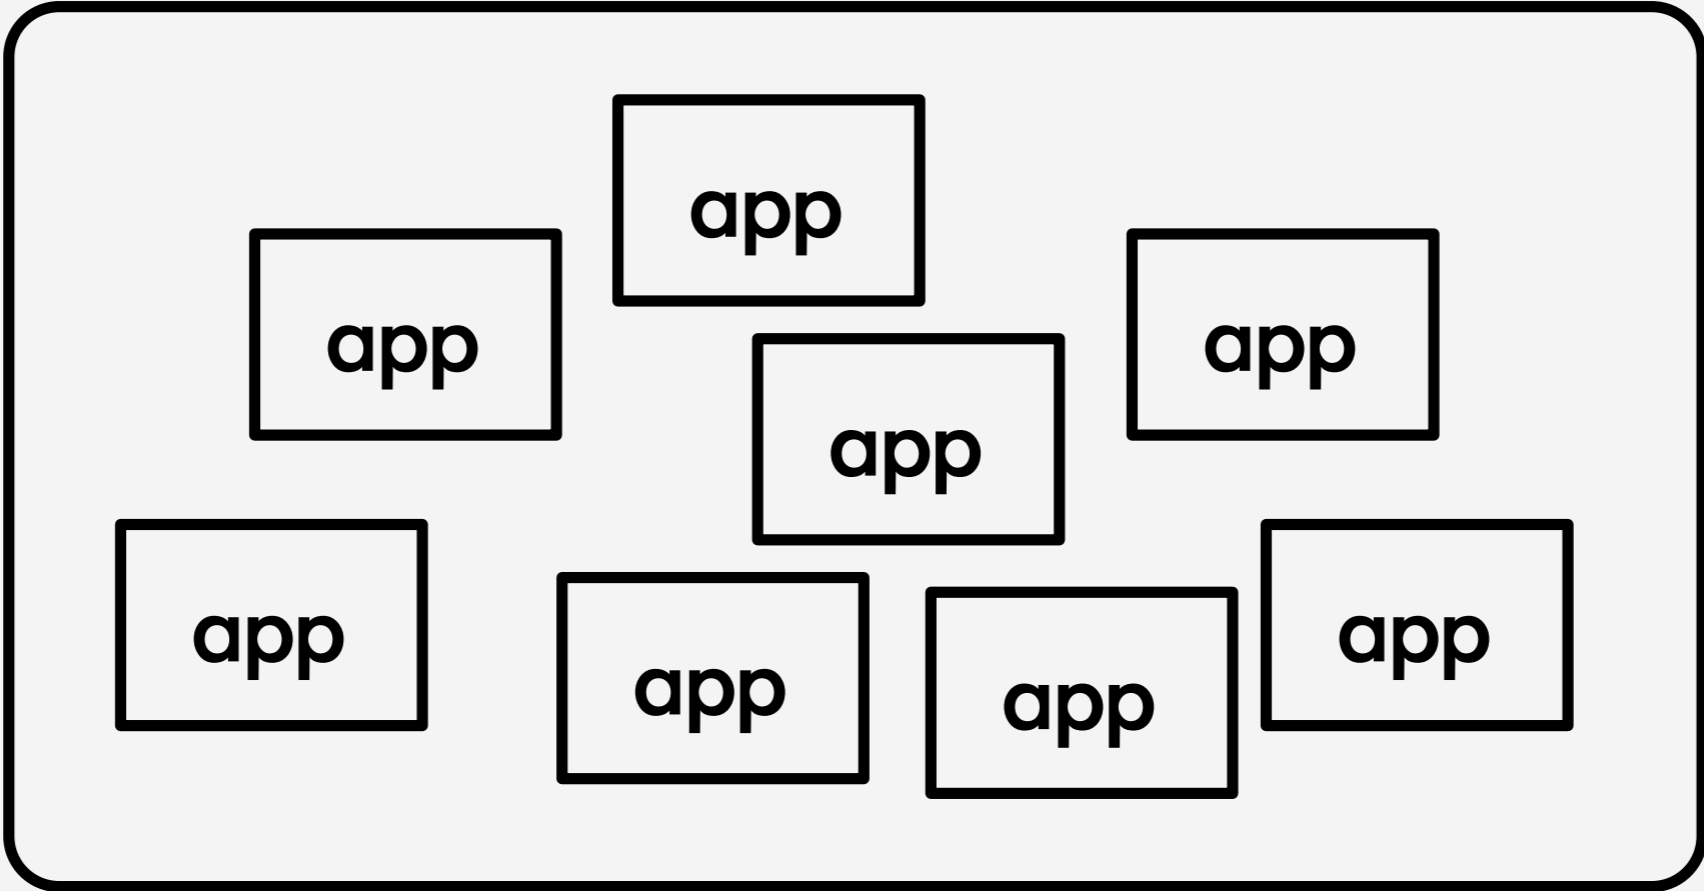
**app**



**routing**


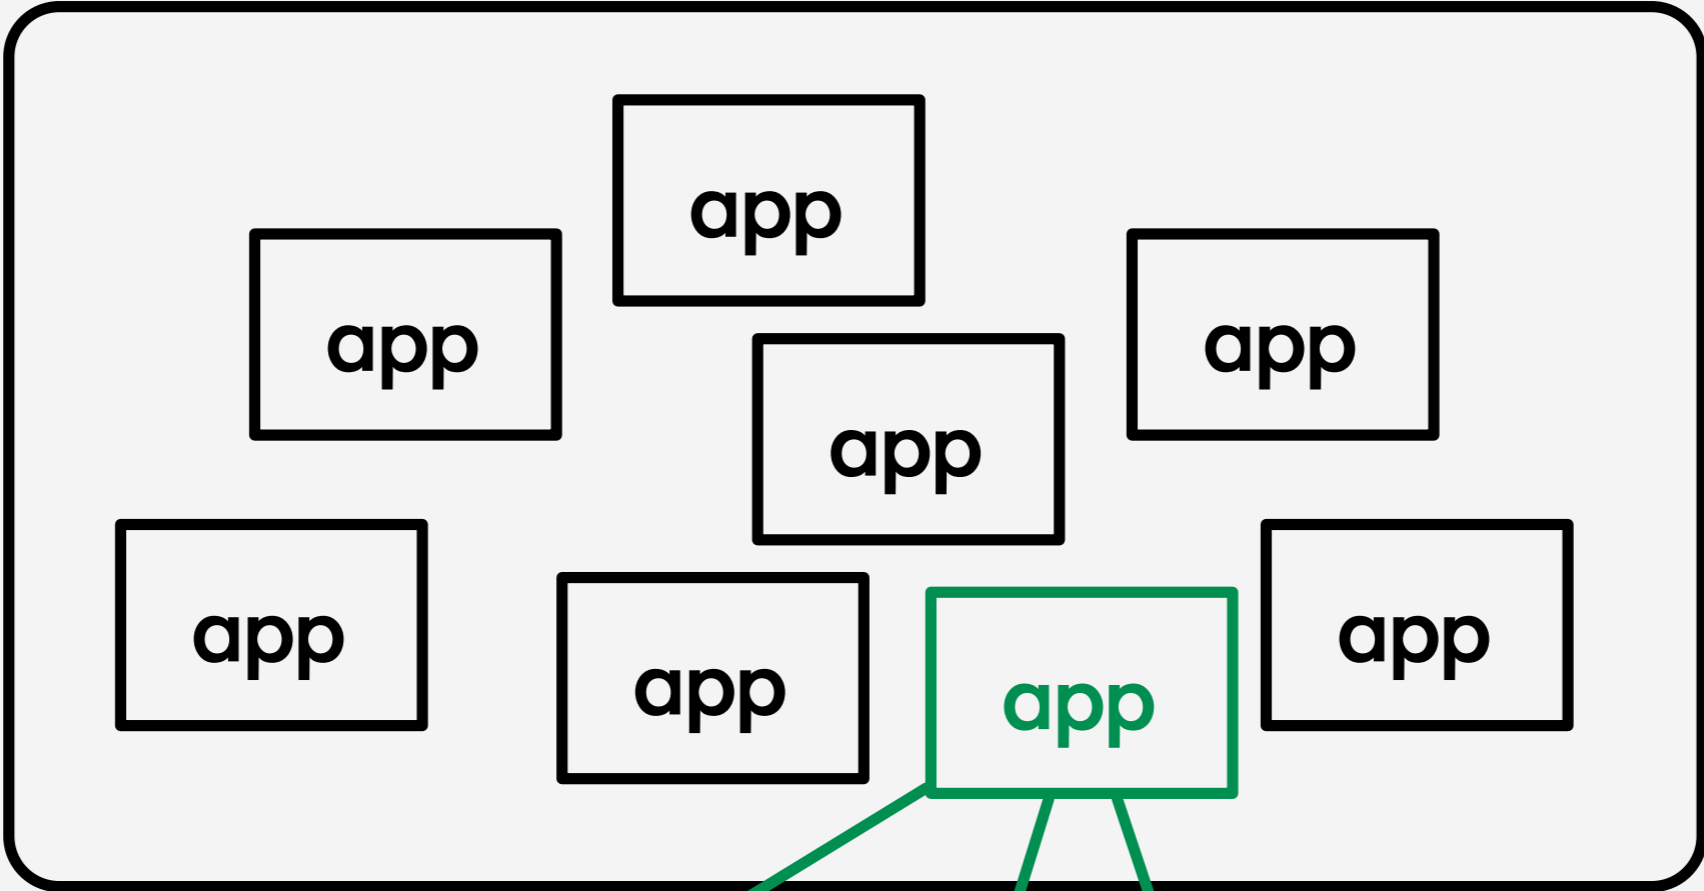
**runtime**

**data**

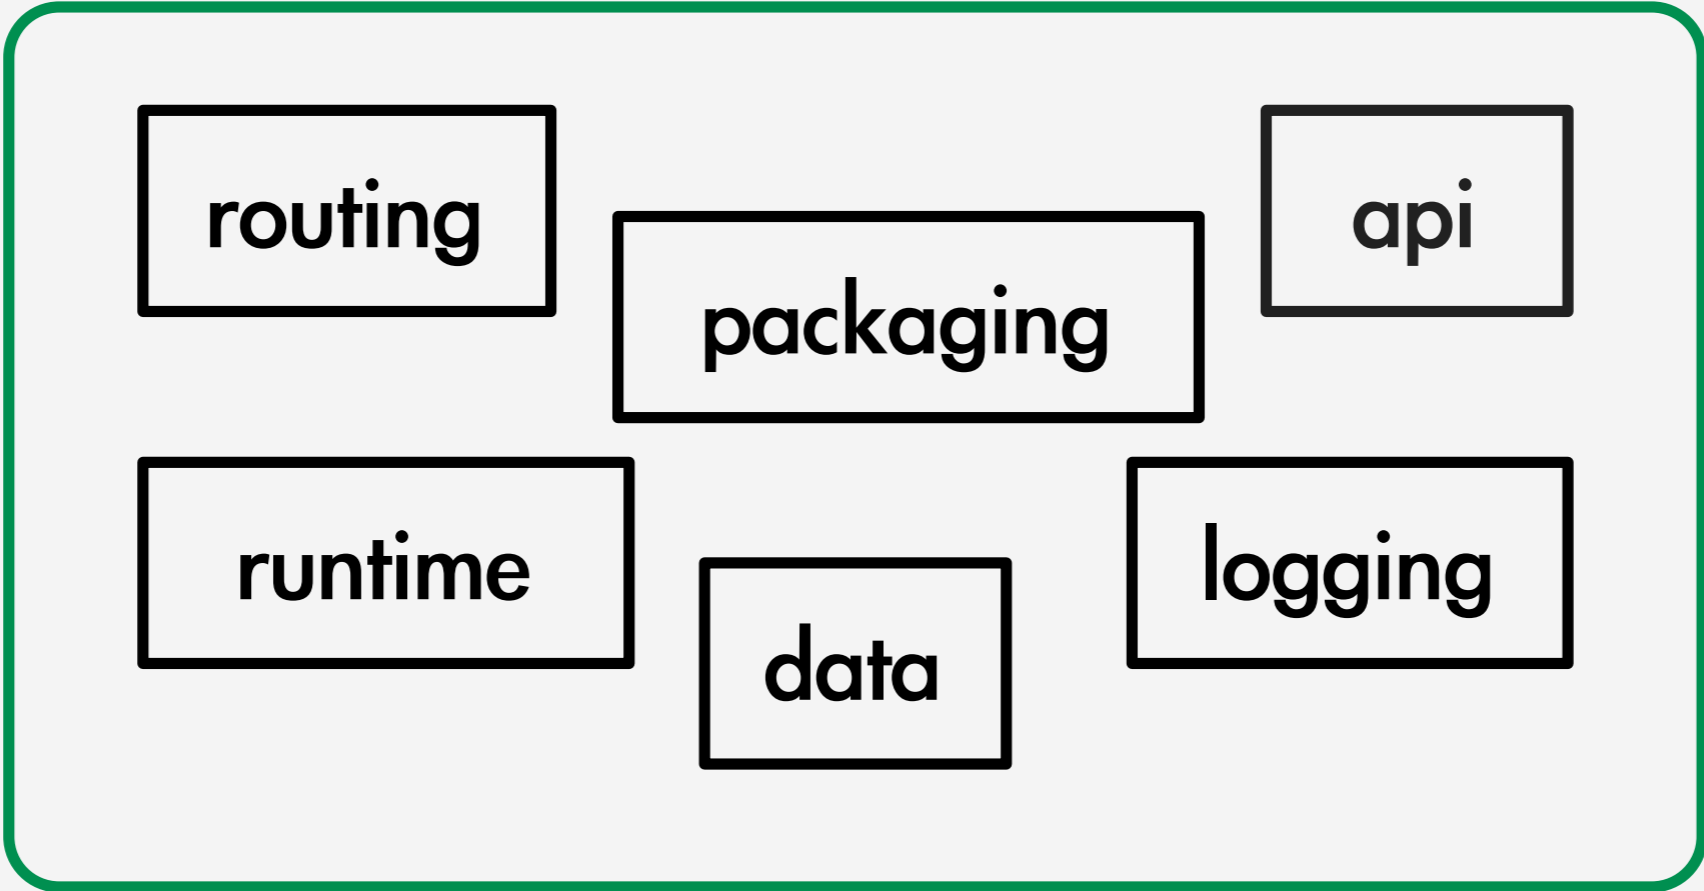
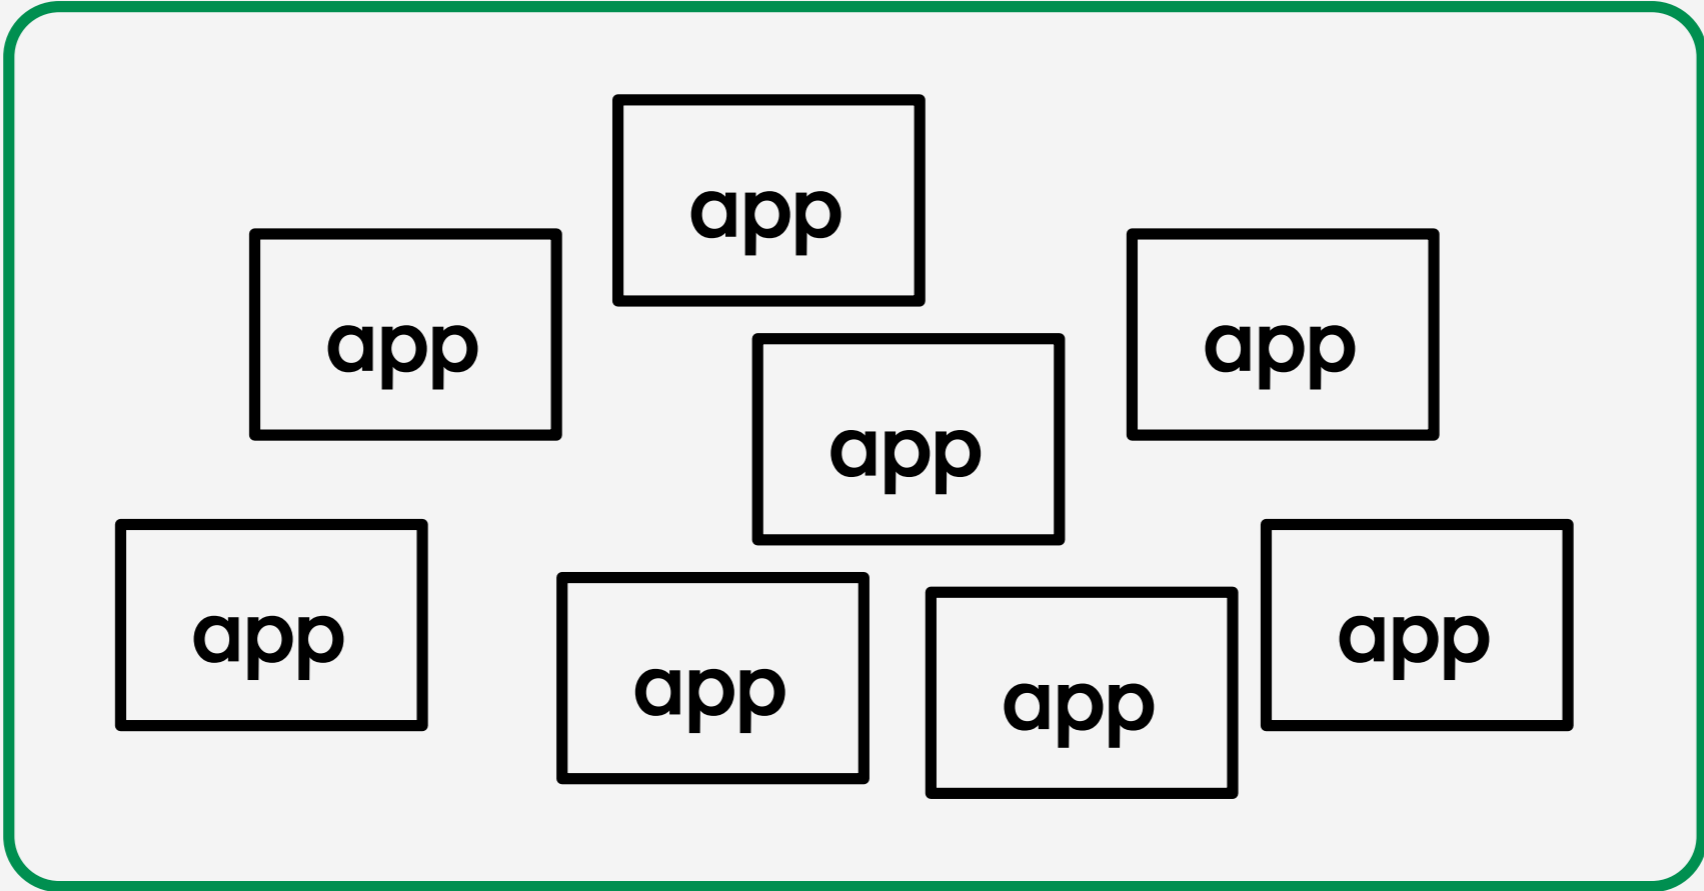












# Scale & Scope

**$O(1,000)$  instances**

**$O(1,000,000)$  apps**

**Success & Failure**

# Architecture Execution

**Architecture**

Execution


**routing**

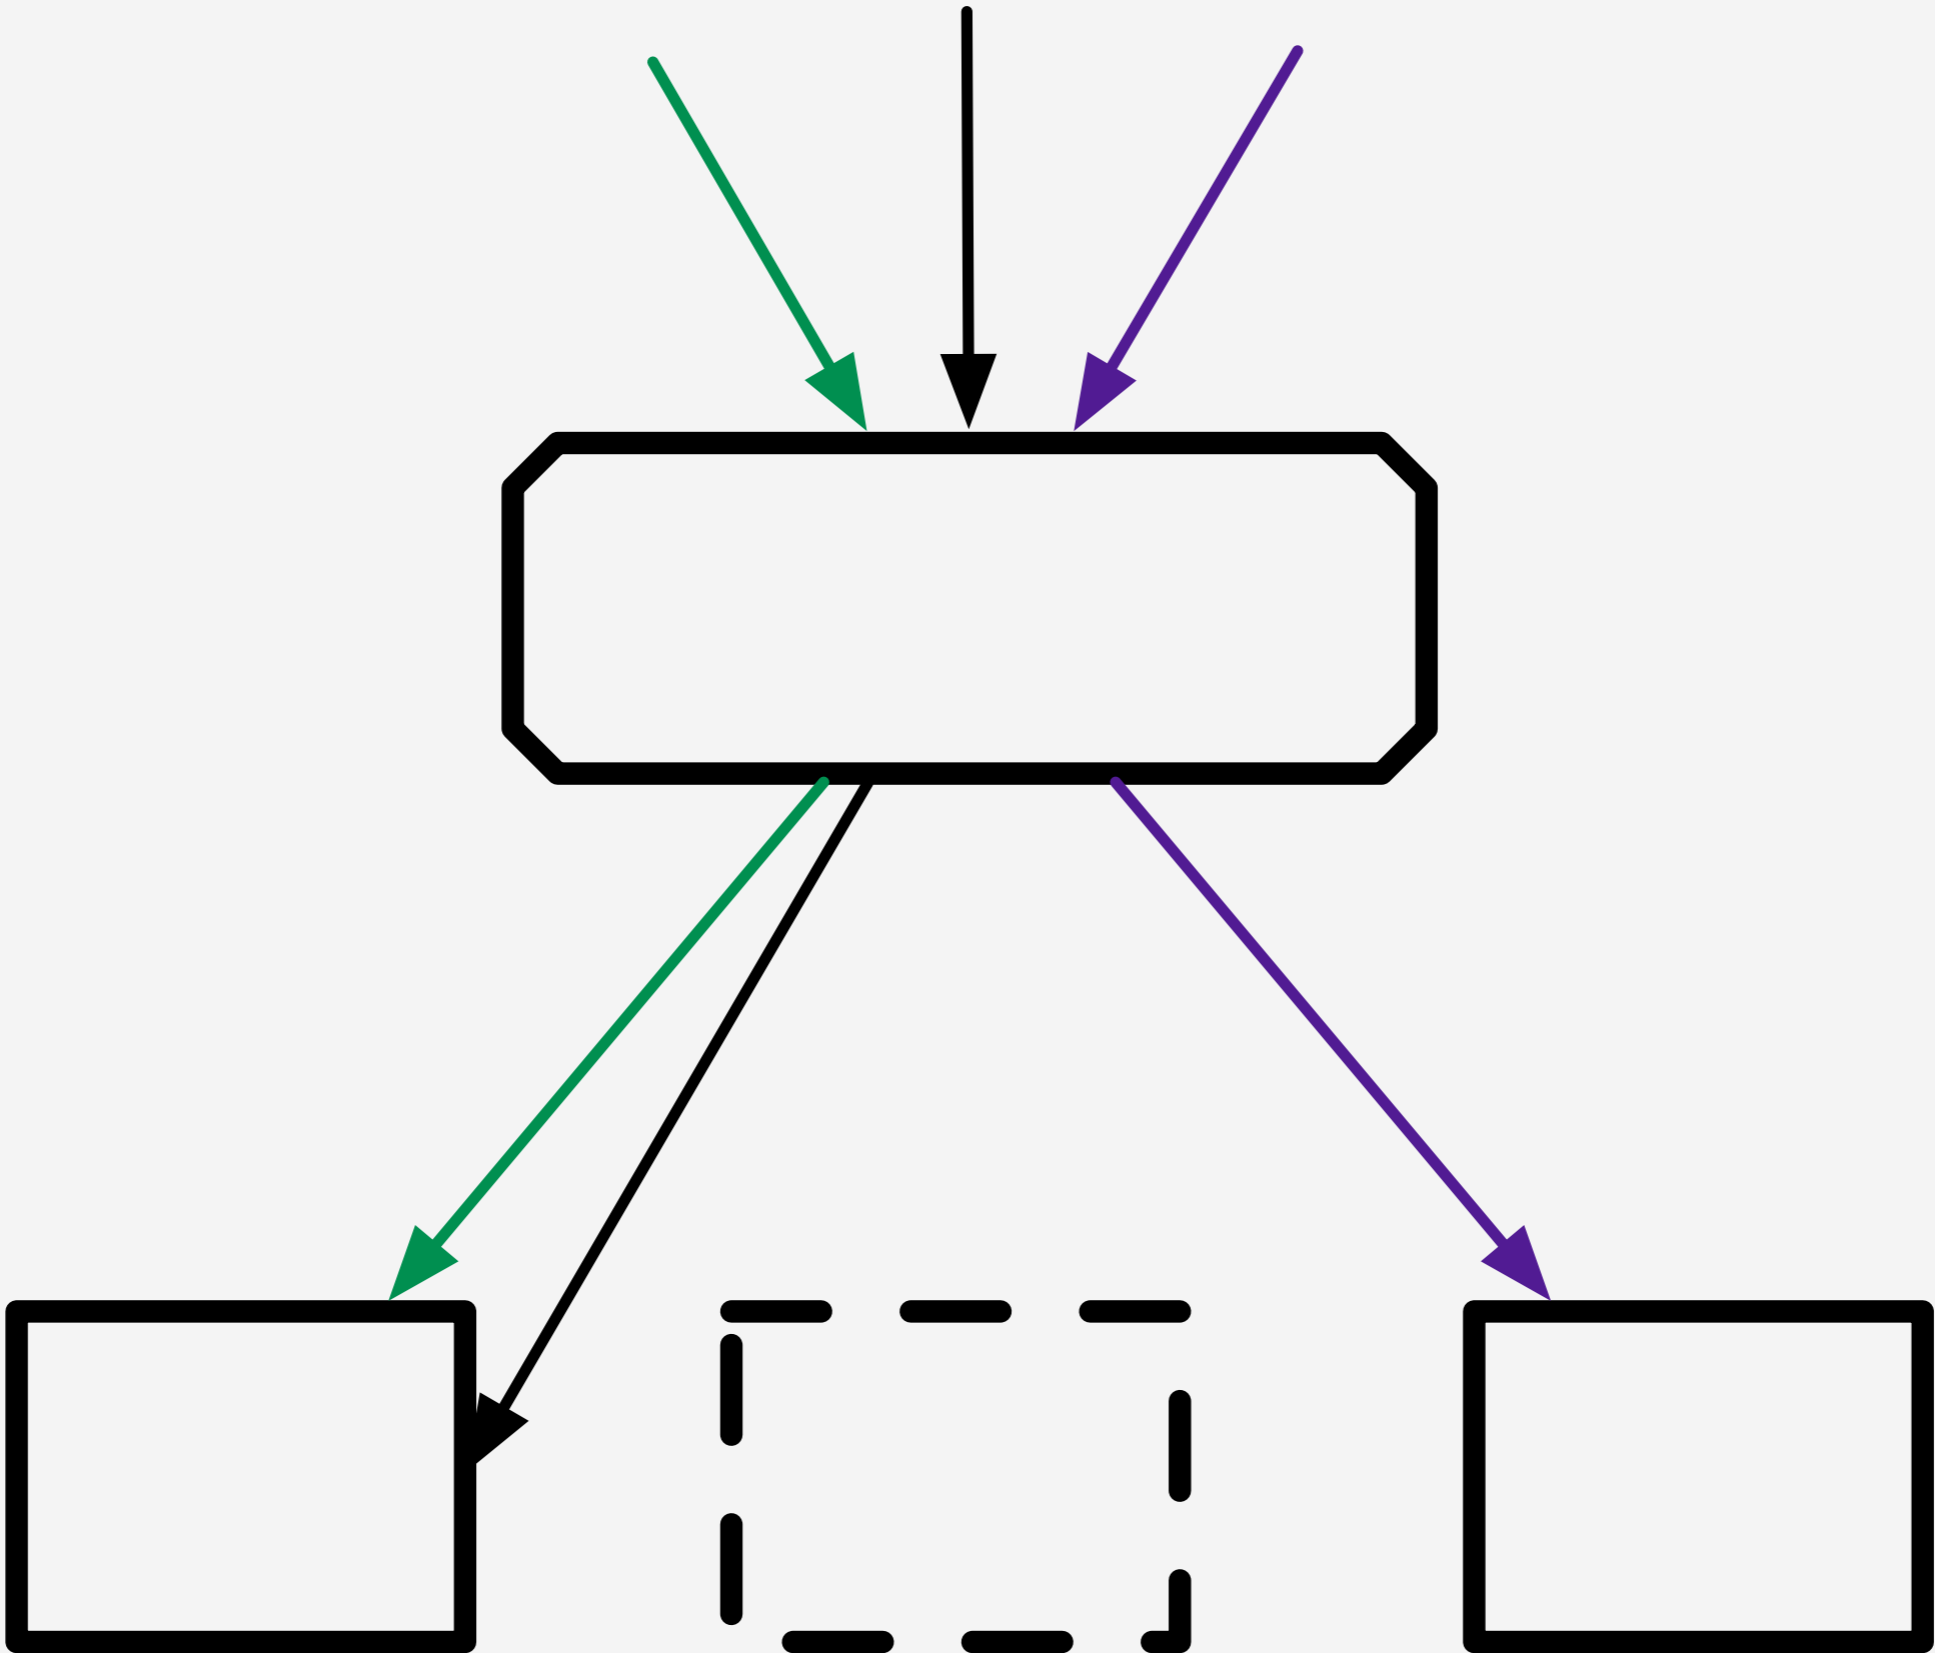
**runtime**



**data**


# Platform-Enabled HA Routing














**Crashes**

**&**

**Supervision**




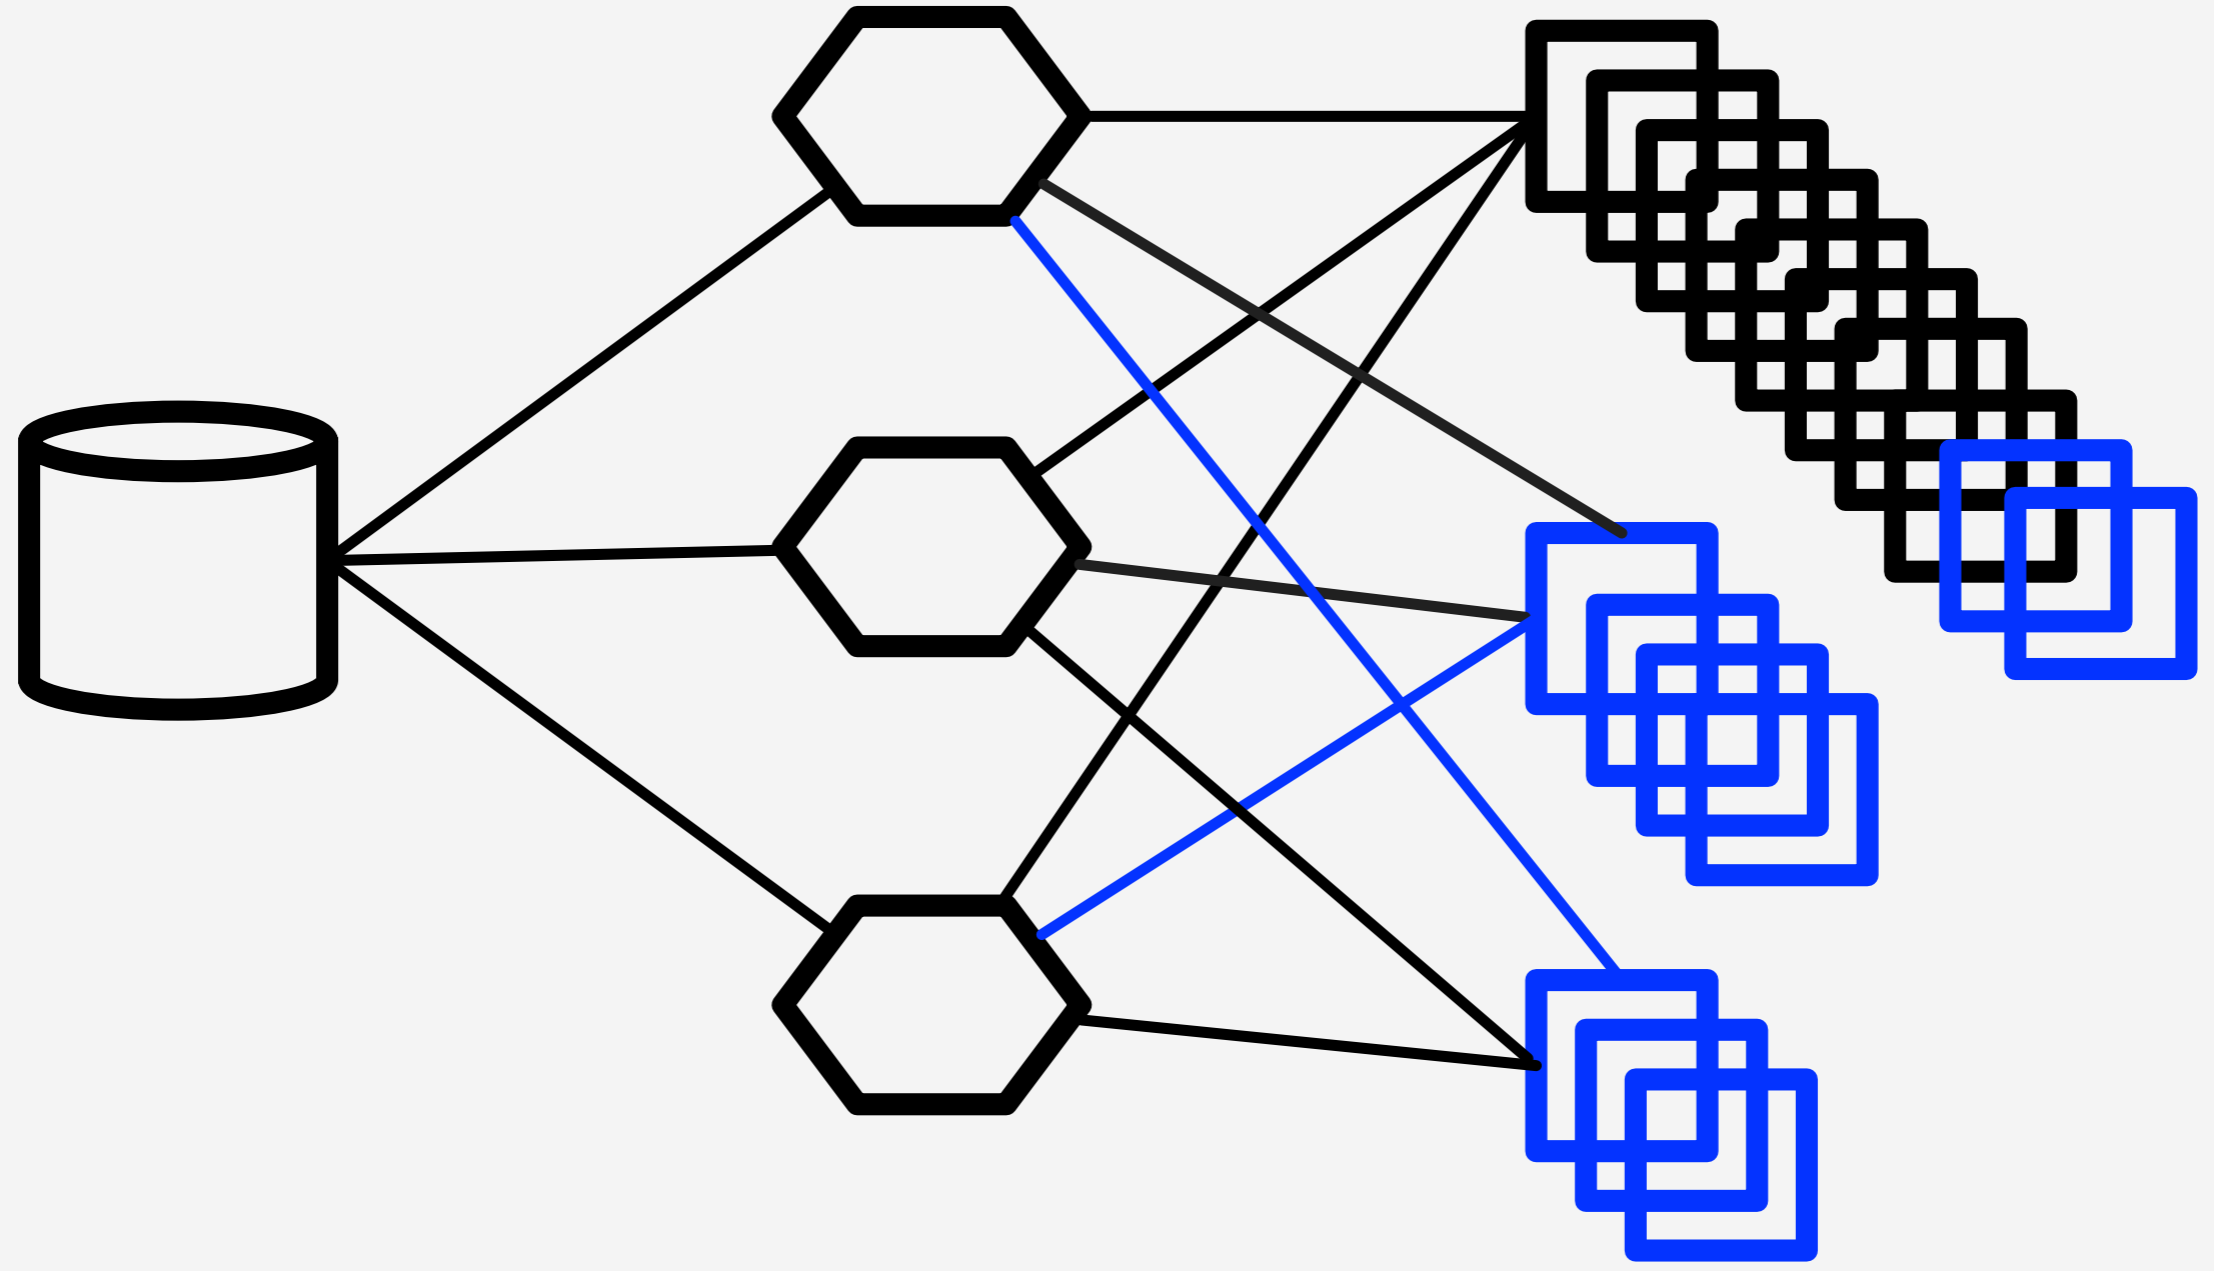







**Crashes as the  
only code path**






**Crashes as a hot  
code path**






# Error Kernel






**Layered design**




**Message  
passing...**



```
{slug: “https://aws...”,  
  cmd: “java ...”,  
  env: {“JAVA_OPTS”: ...,  
        “DATABASE_URL”: ...,  
        “SESSION_SECRET”: ...}}}
```


**...of narrow,  
versioned values**

```
{slug: "https://s3...",  
  cmd: "java -cp ...",  
  env: {"JAVA_OPTS": ...,  
        "DATABASE_URL": ...,  
        "SESSION_SECRET": ...},  
  flag: "extra_cpu"}
```





**No Stoppopping the  
World**



**Load balancing**


**Supervision**

**Crash-only**

**Error kernels**

**Layered design**

**Message-passing**



**ERLANG**






# Erlang

Designed for  
granular failure







# Distributed Systems


Defined as  
granular failure


# Brokered Queueing








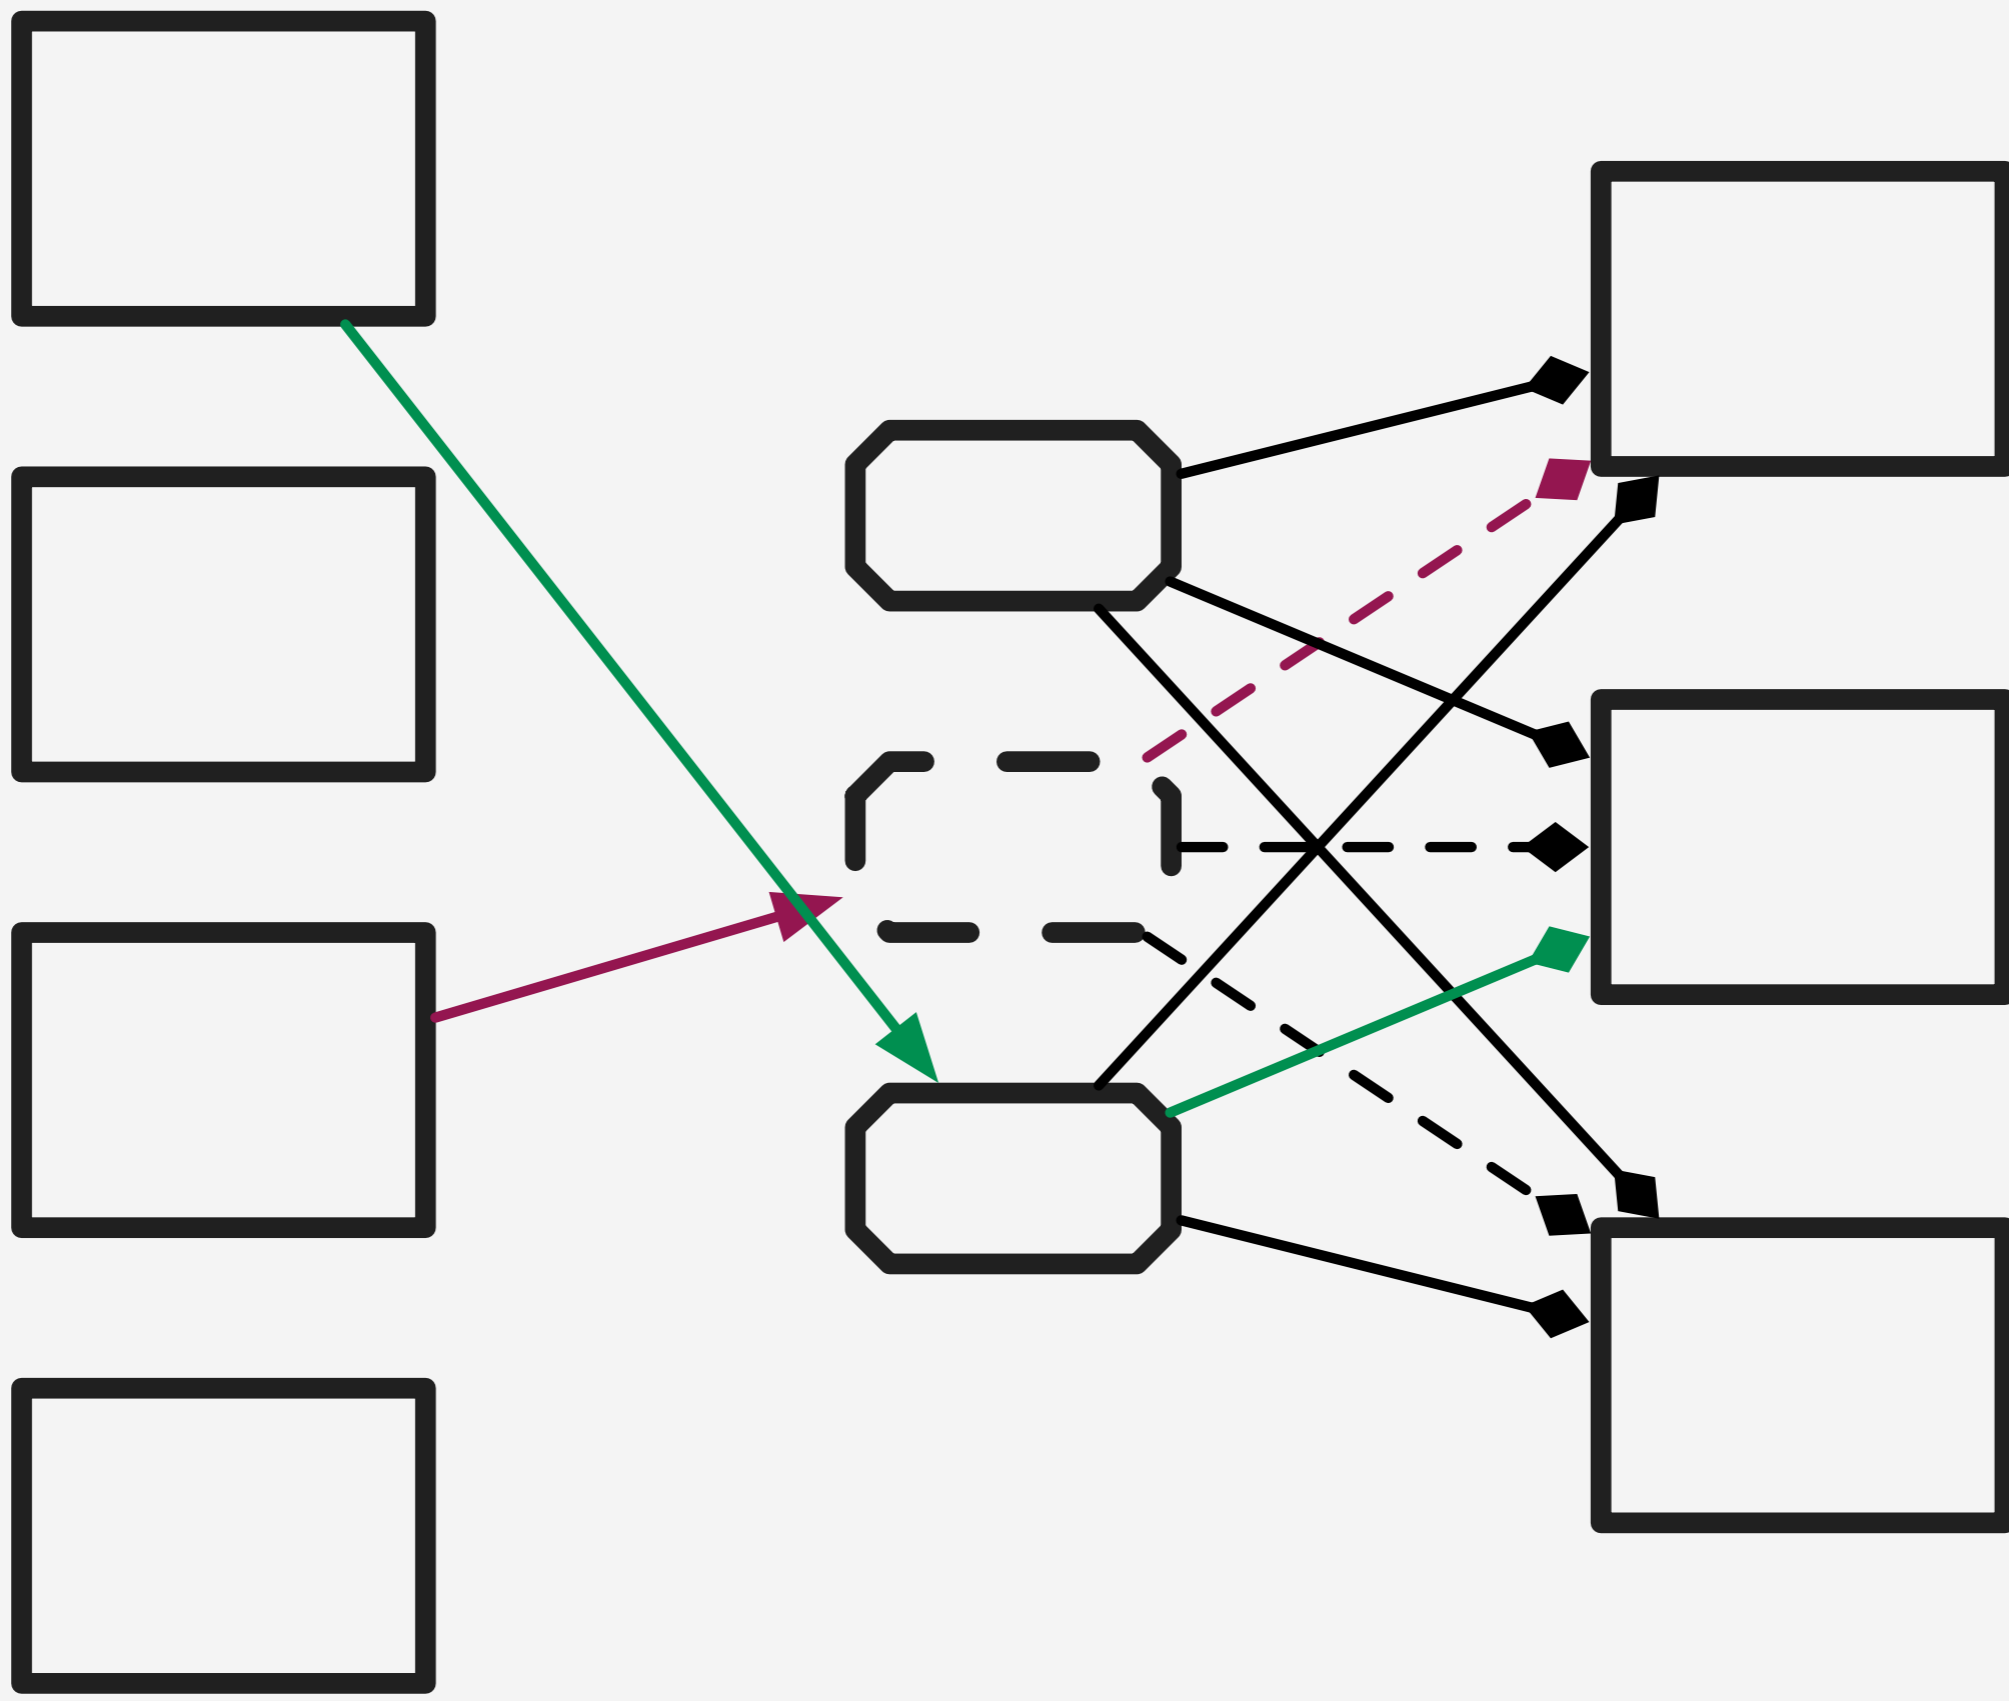





**Publish one /**

**Subscribe many**











# Distributed call graphs








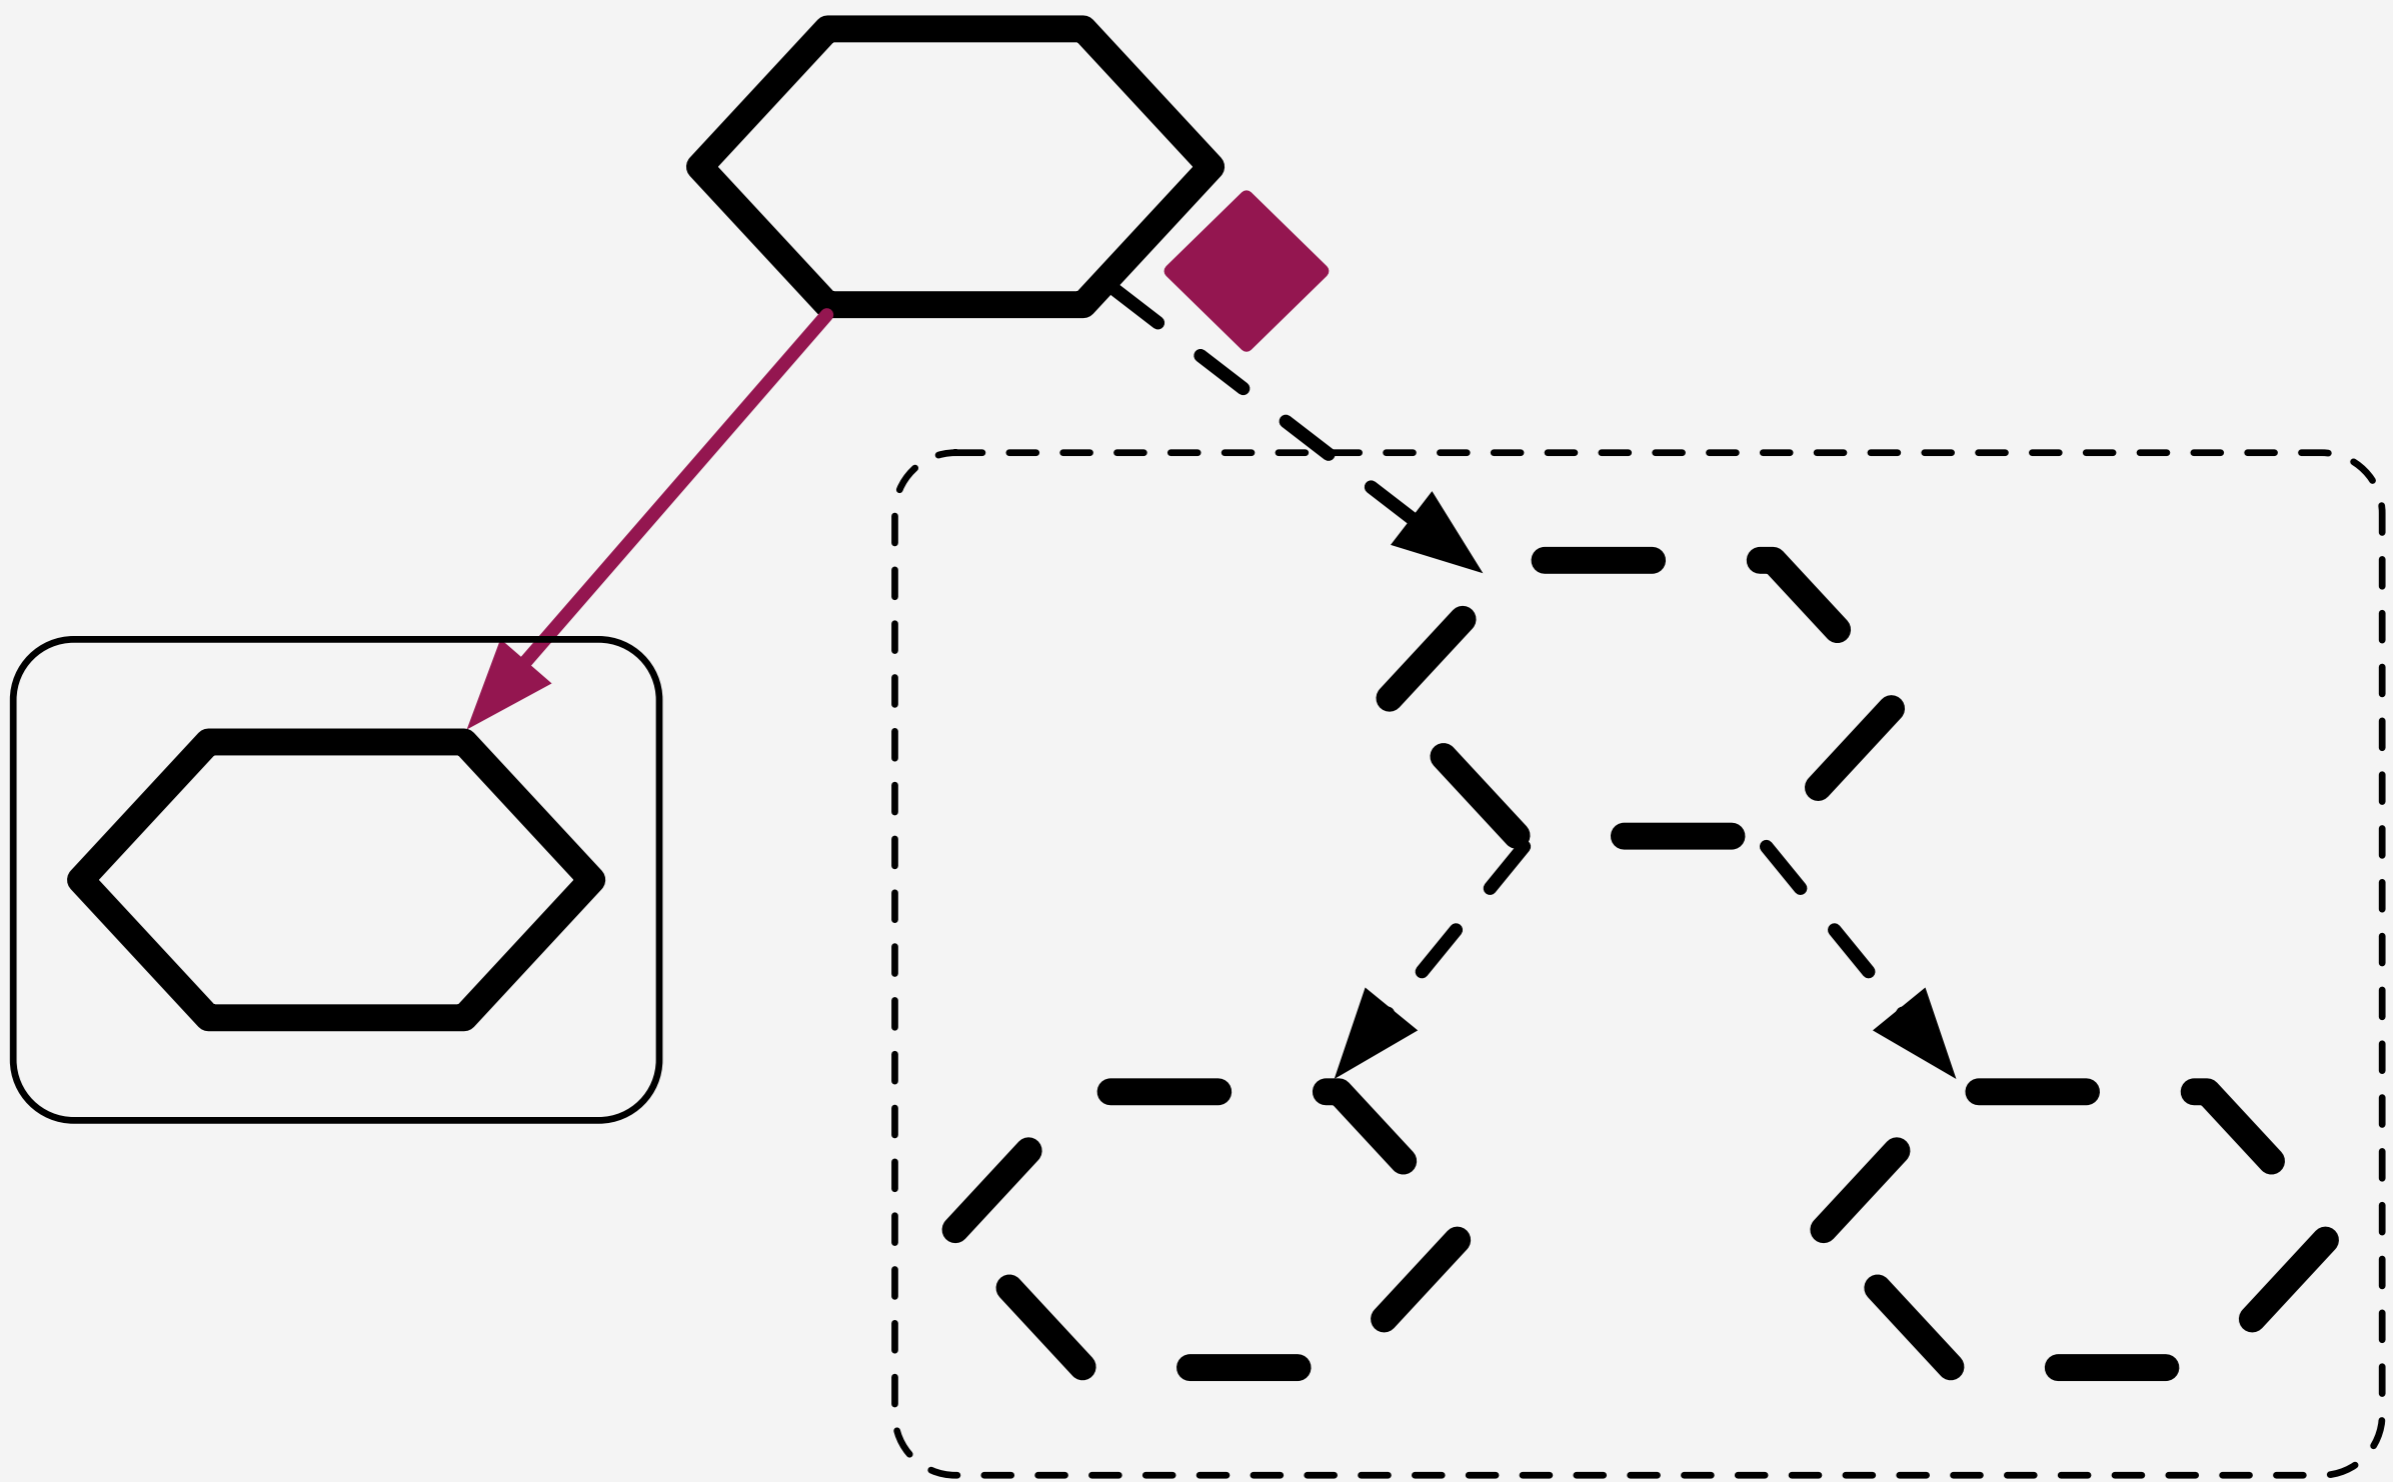





**Read call graph**  
**Partial failure**



**Write call graph  
de-synchronizing**







**Architecture**

Execution

# Architecture

# Execution

“...we deployed a code change...

...introduced a...problem...

...visible under unusual...conditions...

...engineers noticed a deviation...

...began to escalate...

...system...entered into a feedback loop...



...engineers...deactivated the feedback...

# **Evolving Socio-Technical Systems**

**routing**

**runtime**

**data**



**Availability >> Architecture**



**Failed deploys**

**Bad visibility**

**Cascading feedback**

# **Evolving Socio-Technical Systems**

**Failed deploys**

**Bad visibility**

**Cascading feedback**

**Deploy tooling**  
**Visibility services**  
**Feedback controls**

bin/ship

bin/ship \

--component api \

--version v408

**Incremental deploys**

**t+0**

**t+1**

**t+2**



**t+0**

**t+1**

**t+2**


# Incremental rollouts

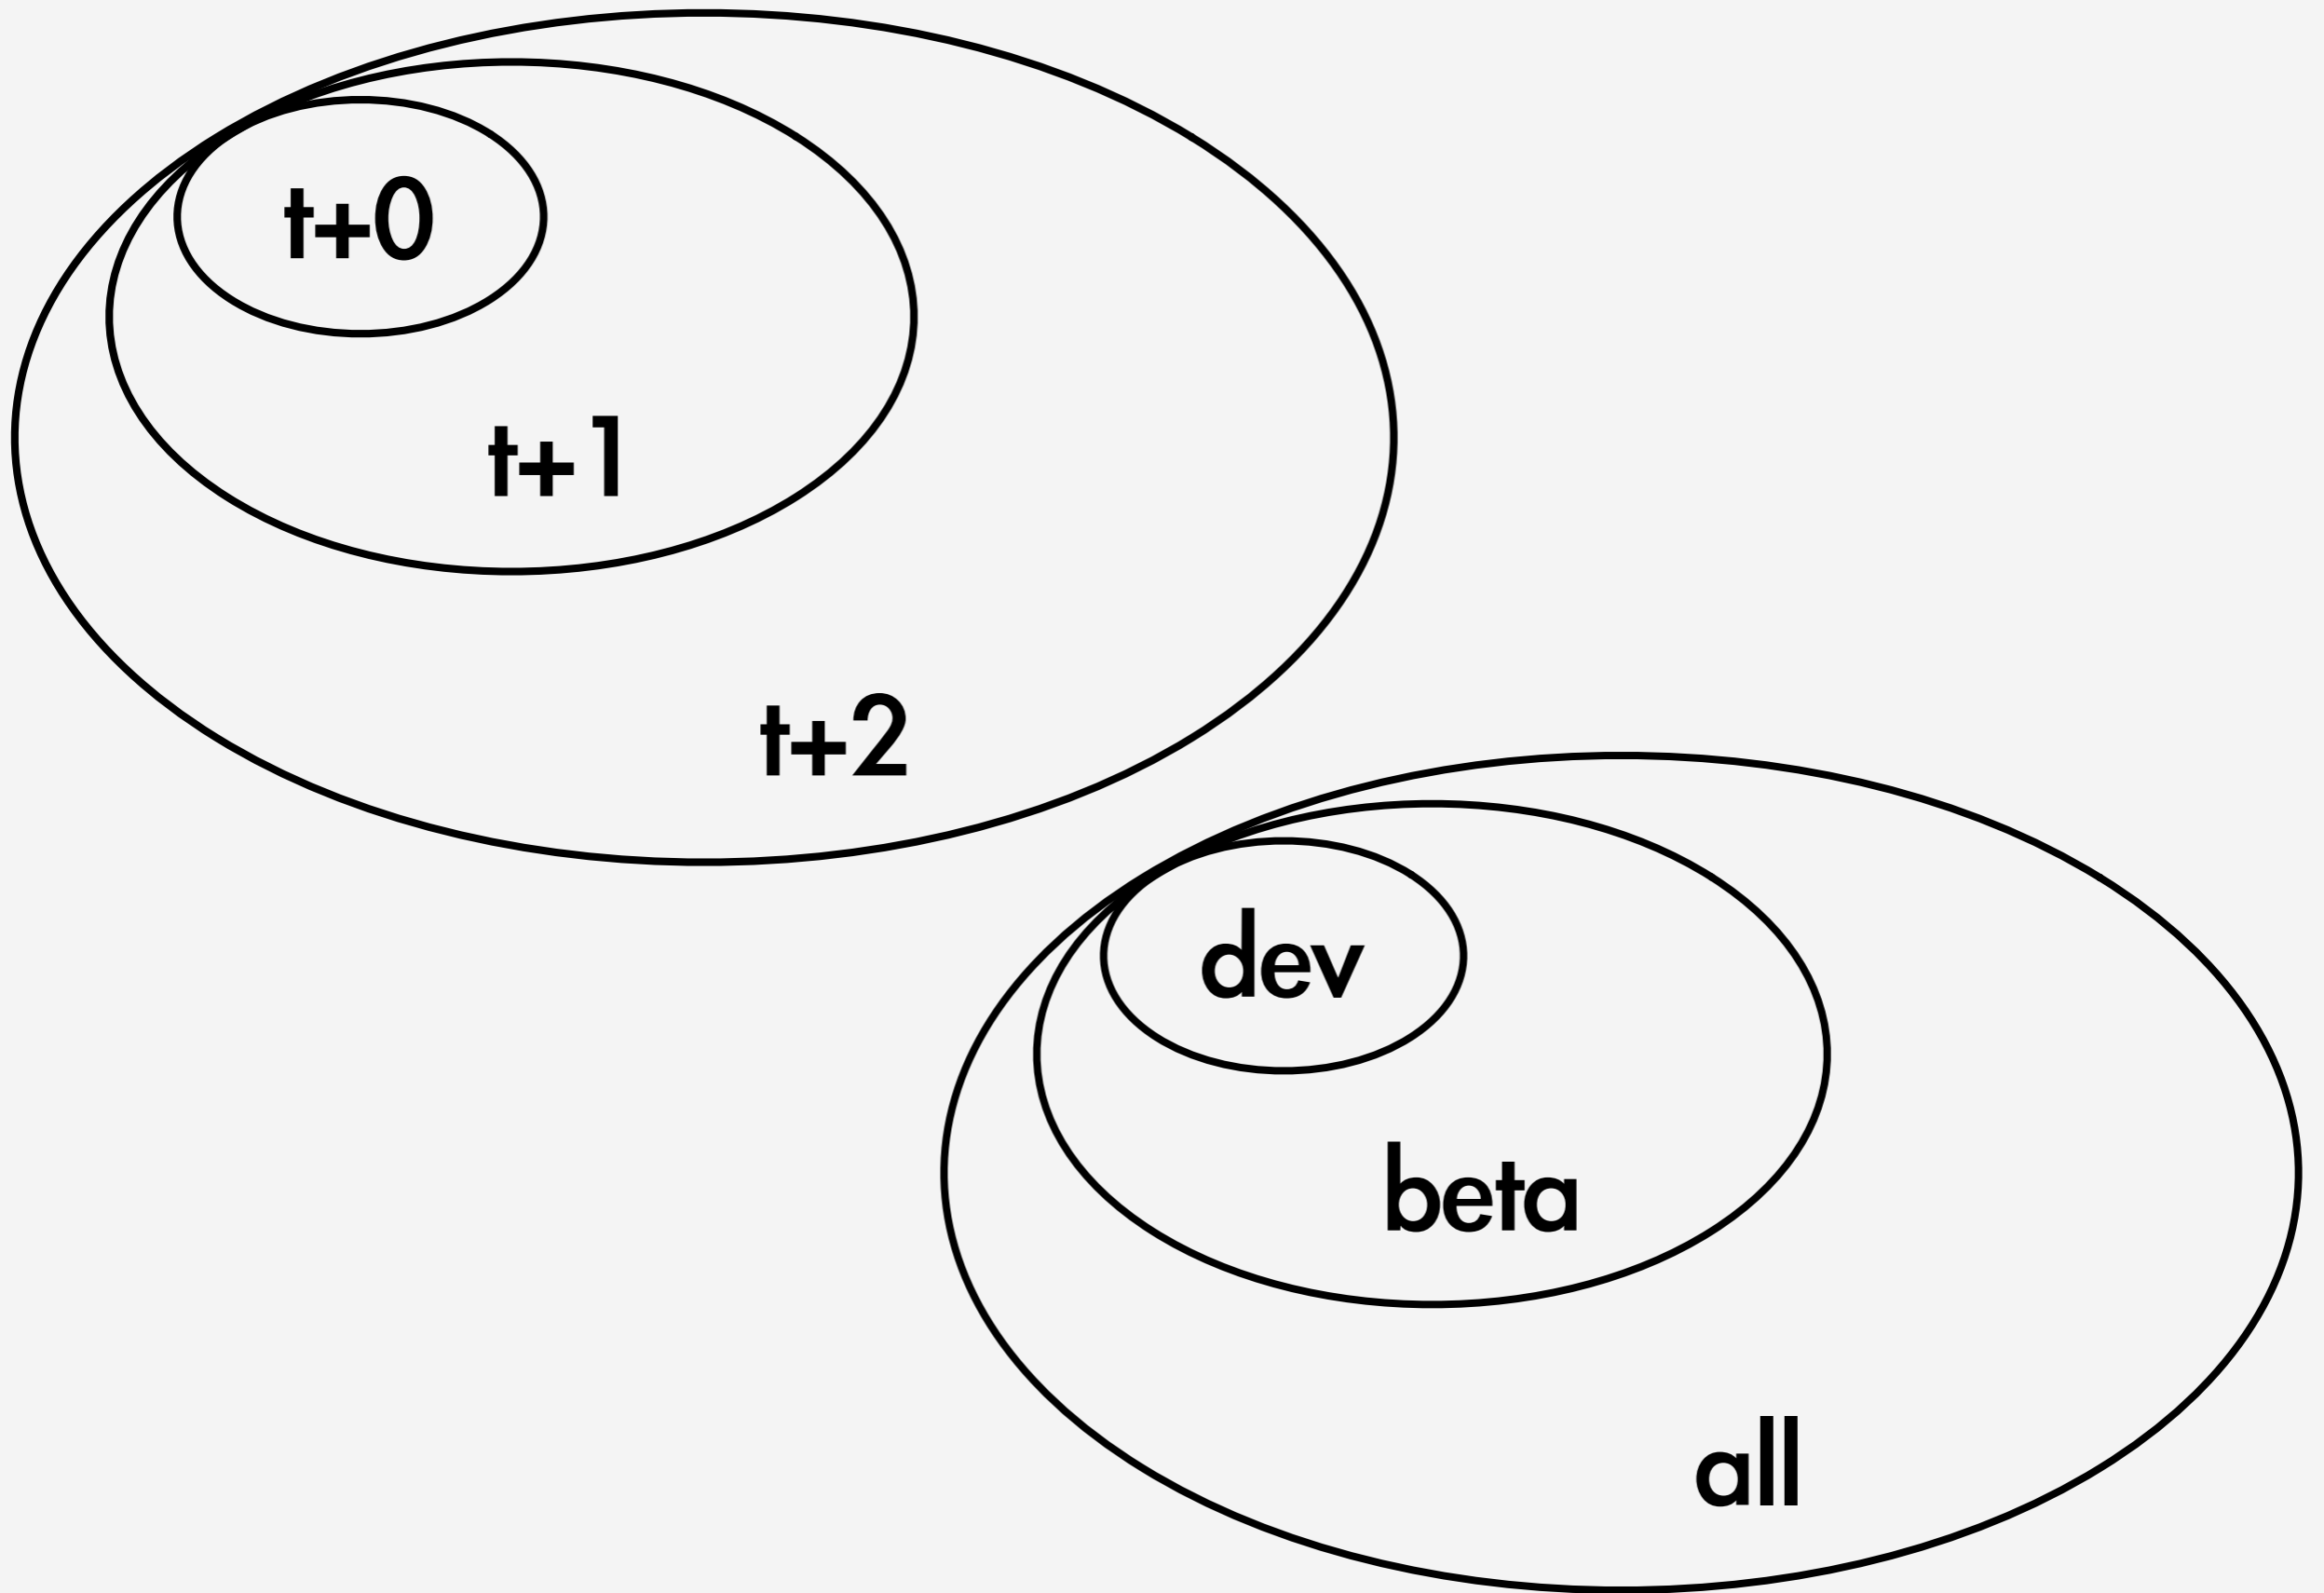
prep\_launch  
- launch\_without\_lxc  
+ launch-with\_lxc  
monitor\_launch

```
if flag_on?("lxc")
    launch_with_lxc
else
    launch_without_lxc
```

```
app.flag_on("lxc")
```

```
app.flag_off("lxc")
```





**Real-time visibility**





**Service-level assertions**




```
assert(index > 0)
```

```
assert(index > 0)
```

```
objects[index]
```

```
assert(p99_latency < 50)
```

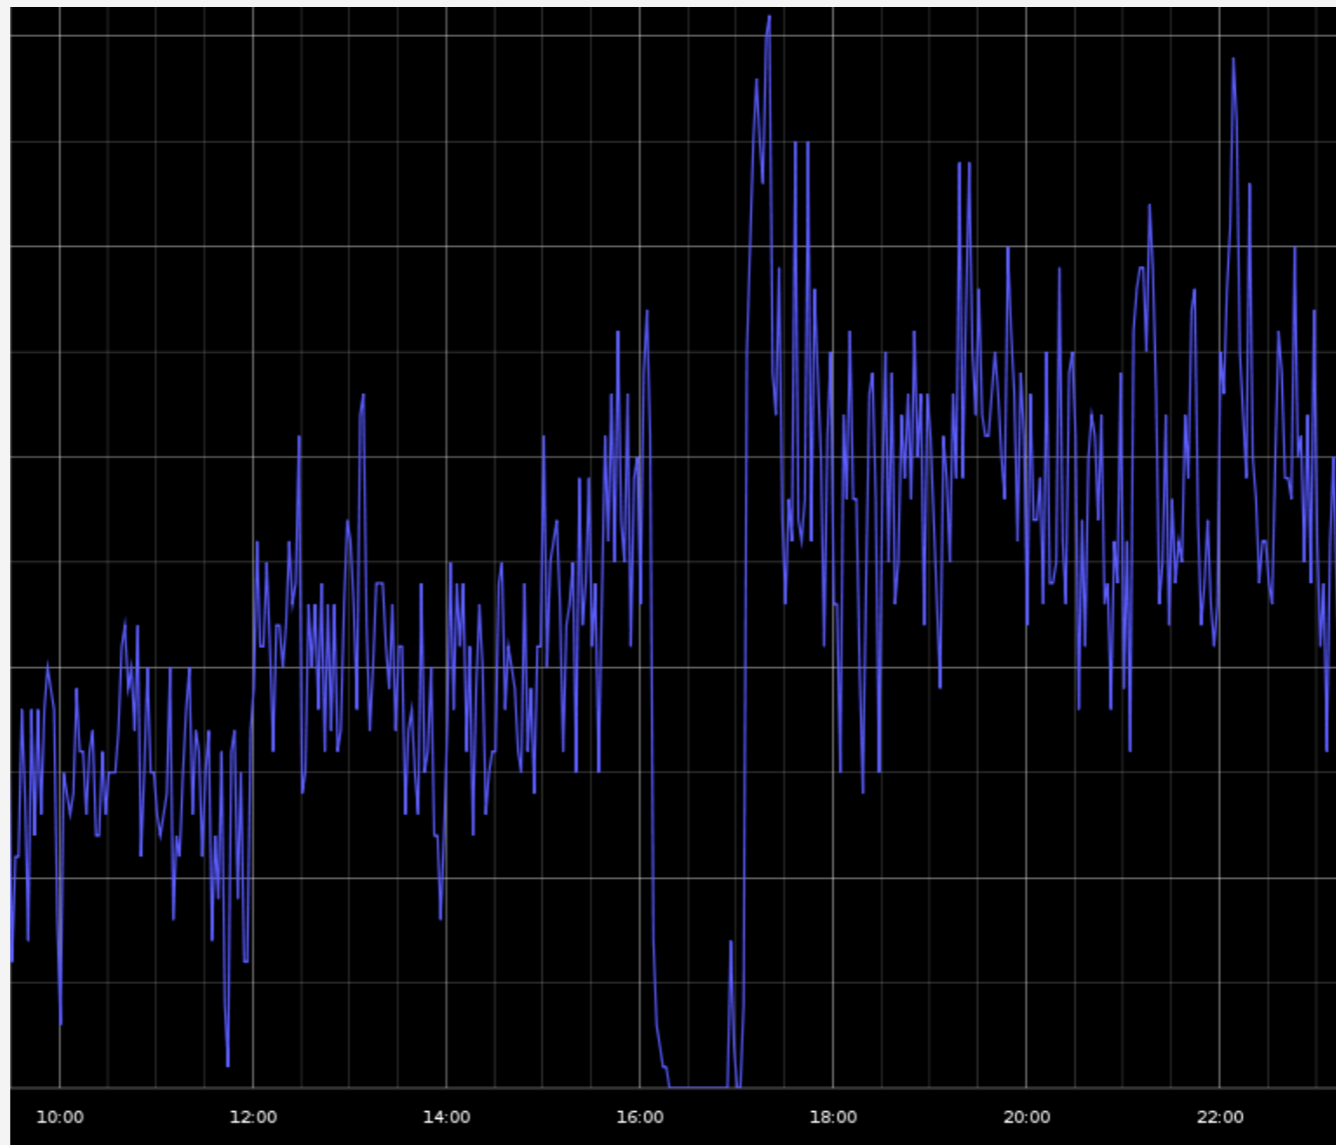
```
assert(p99_latency < 50)
```




```
assert(active_cons > 10)
```




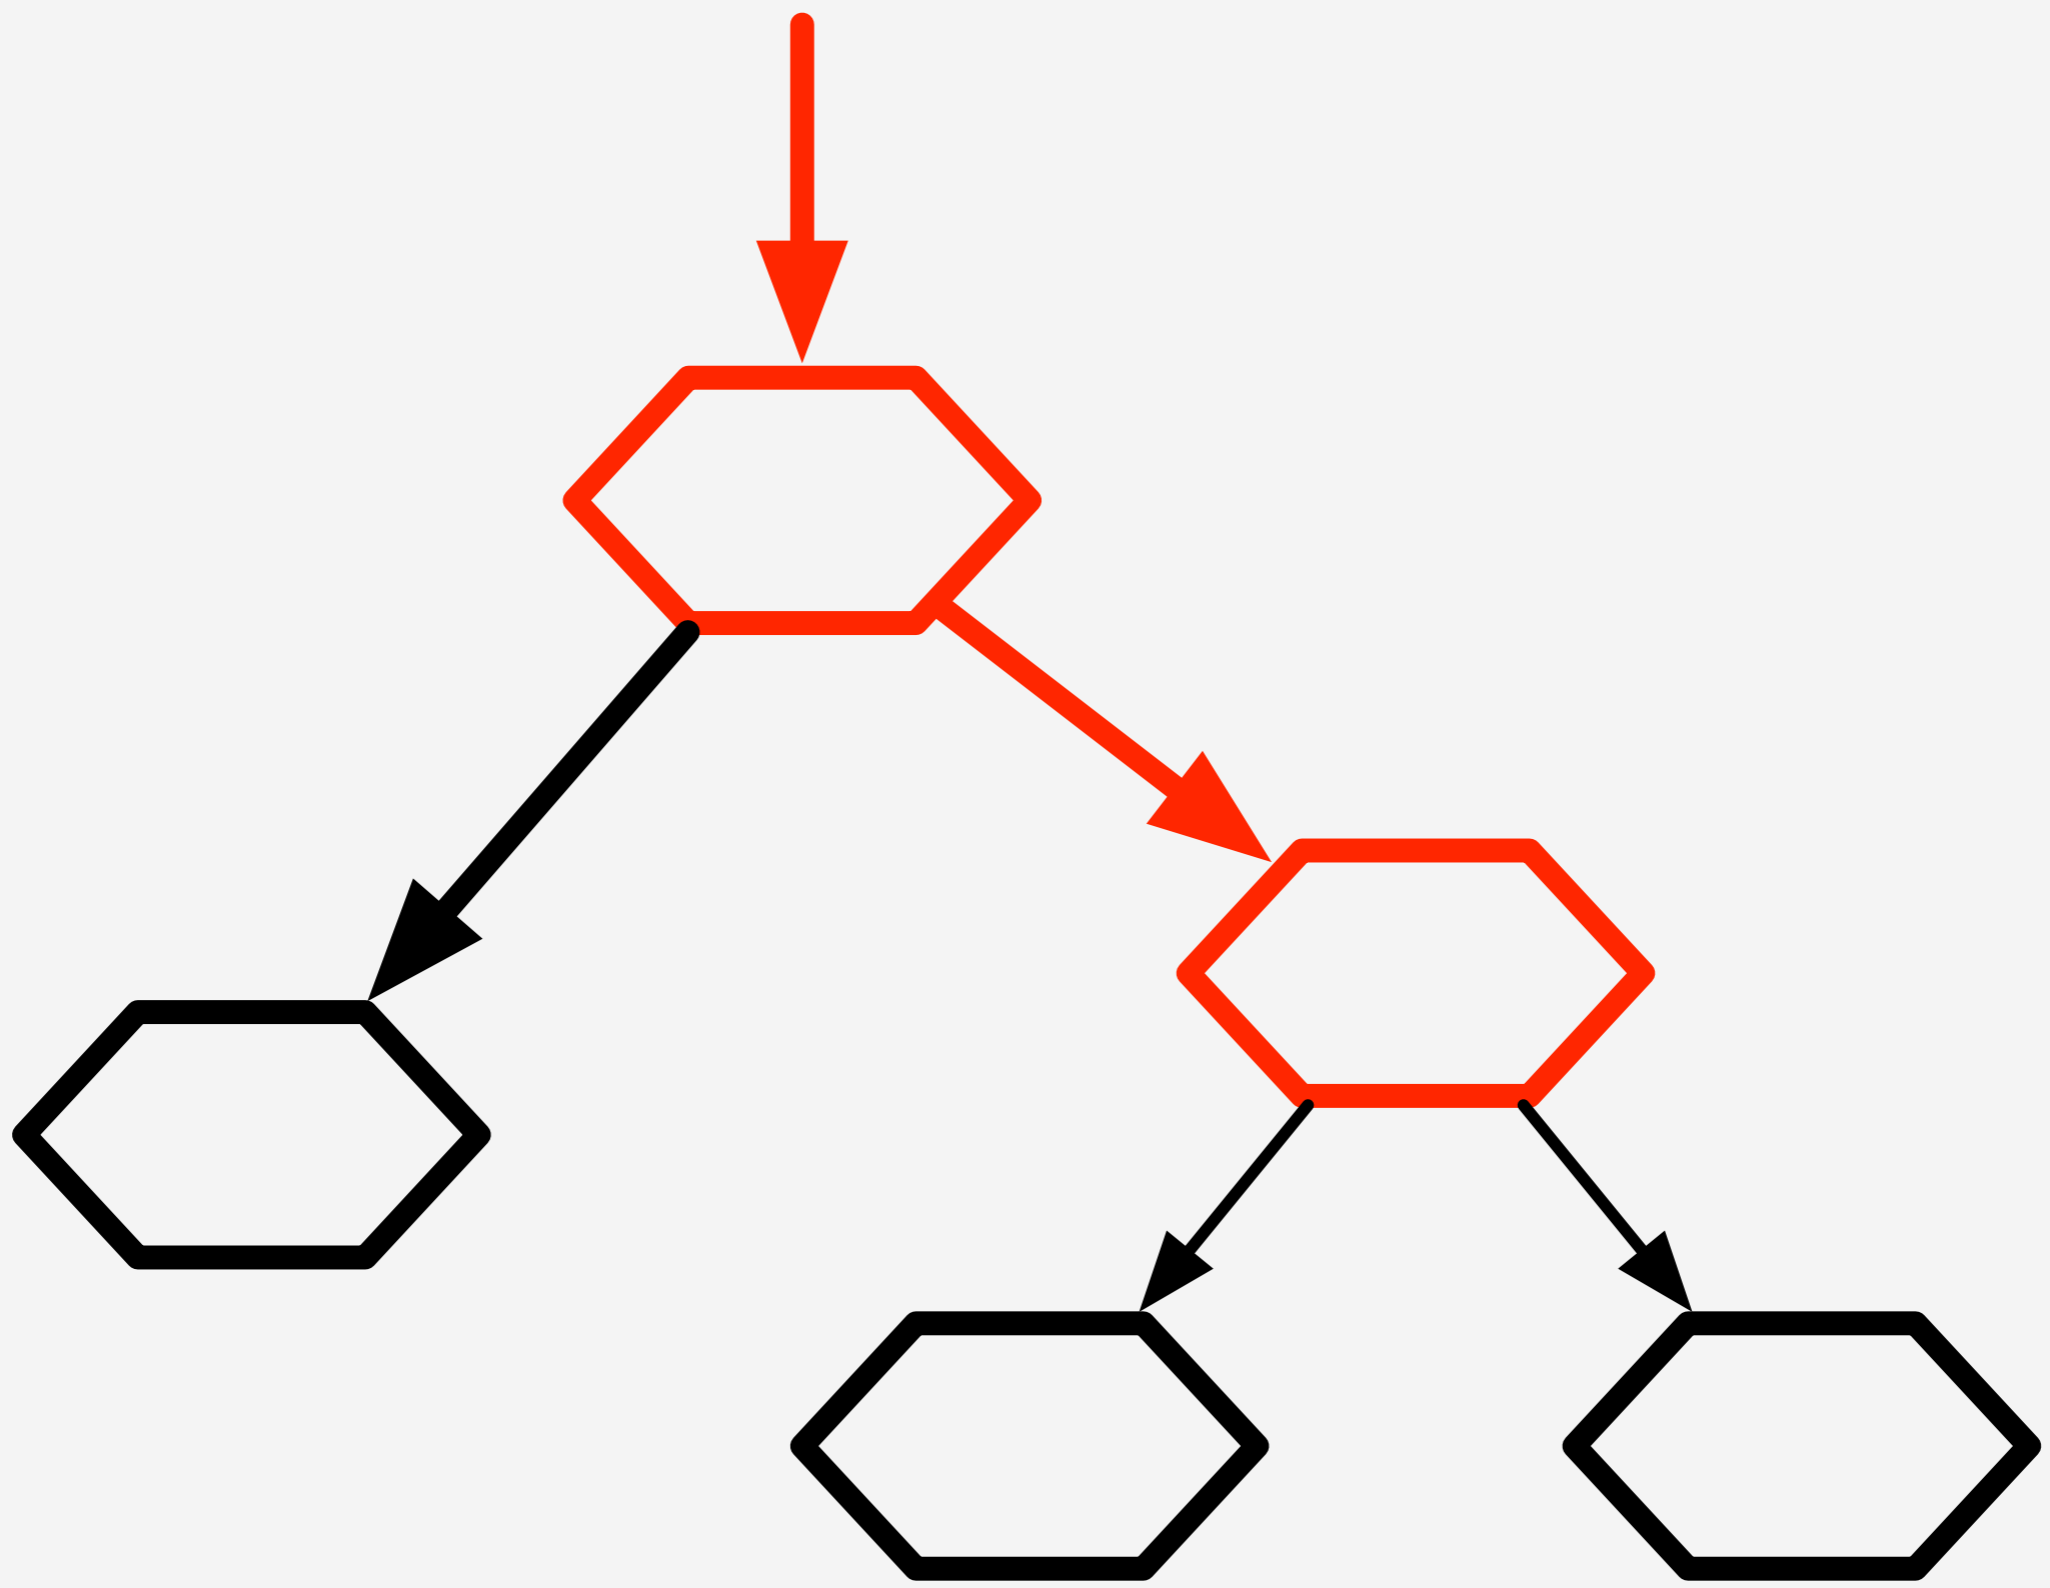
```
assert(active_cons > 10)
```

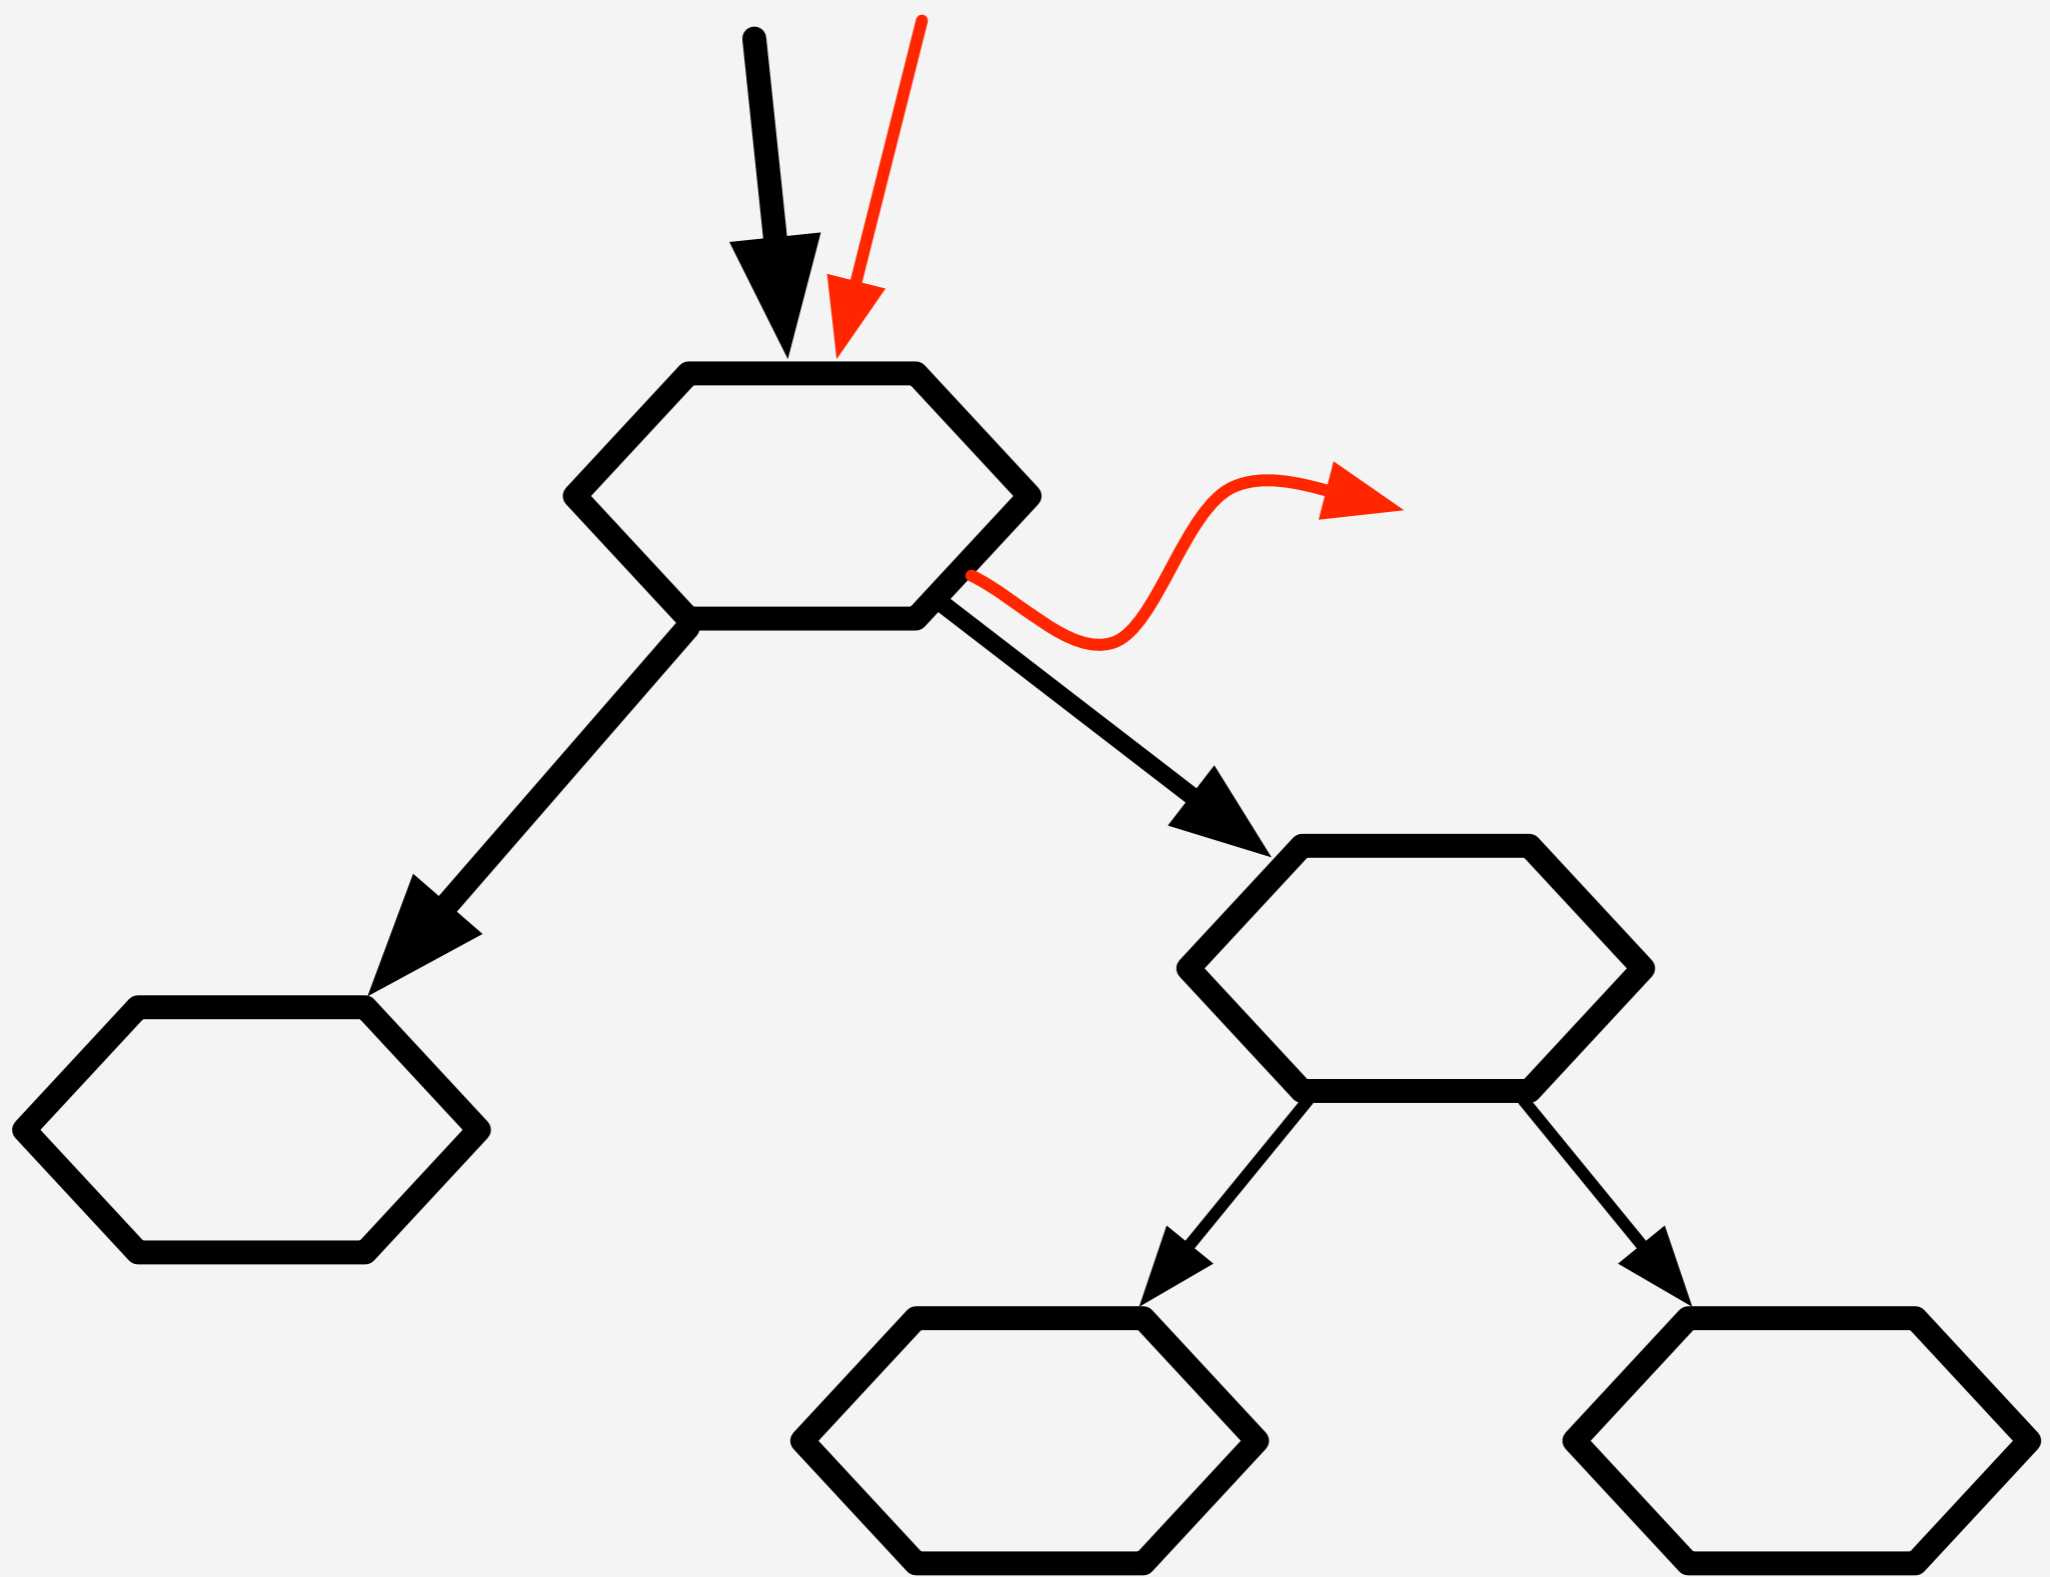


# Flow control & Backpressure









```
limit("publish", 10) do  
  Bus.publish(msg)  
end
```

```
limit("publish", prate) do  
  publish(msg)  
end
```


```
echo 0 >/etc/rates/publish
```



# **Evolving Socio-Technical Systems**

# Architecture

# Execution



heroku

**Thanks**

**Questions?**