# Floodlight Tracing

**Tulio A. Ribeiro**
*Fernando M. V. Ramos, Alysson Bessani*

University of Lisbon, Faculty of Sciences.
LaSIGE - Large-Scale Informatic Systems Laboratory.
Portugal.

NavTalk 03/11/2016

Ciências ULisboa | Informática

LASIGE
Large-Scale Informatics Systems Laboratory

# Outline

- SDN – Context
- (Strong) Consistency Matters
- Proposed Controller Architecture
- Ravana x Floodlight Tracing
- Floodlight Pipeline Tracing
- Como Evitar Sobrecarga
- Ganho esperado

# SDN - Context

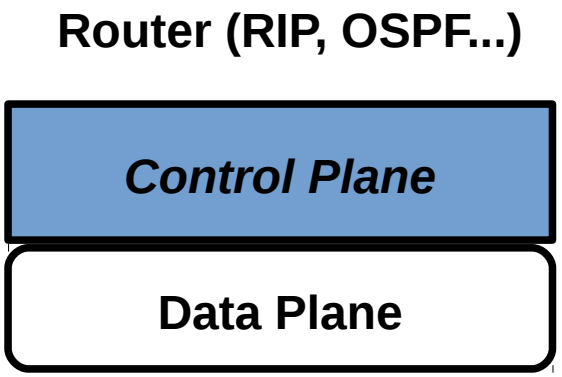


**Traditional Network Device.**

# SDN - Context

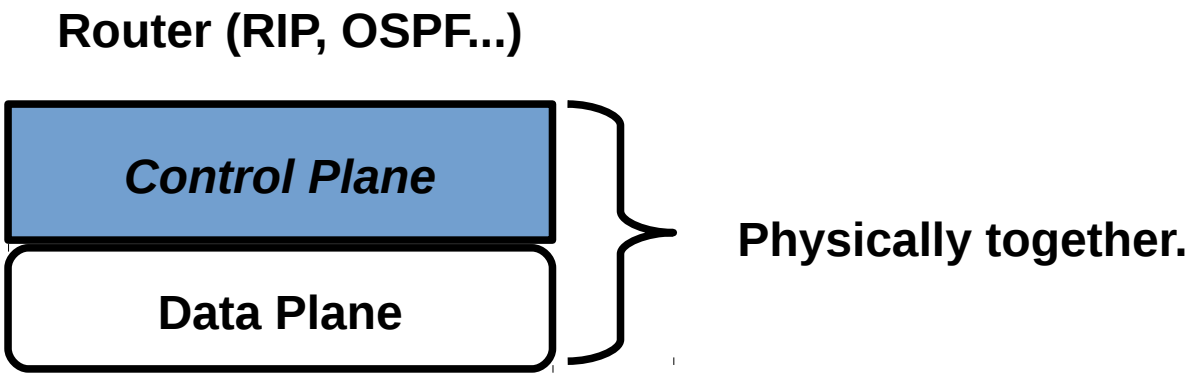


**Traditional Network Device.**

# SDN - Context

**Separation of concerns.**

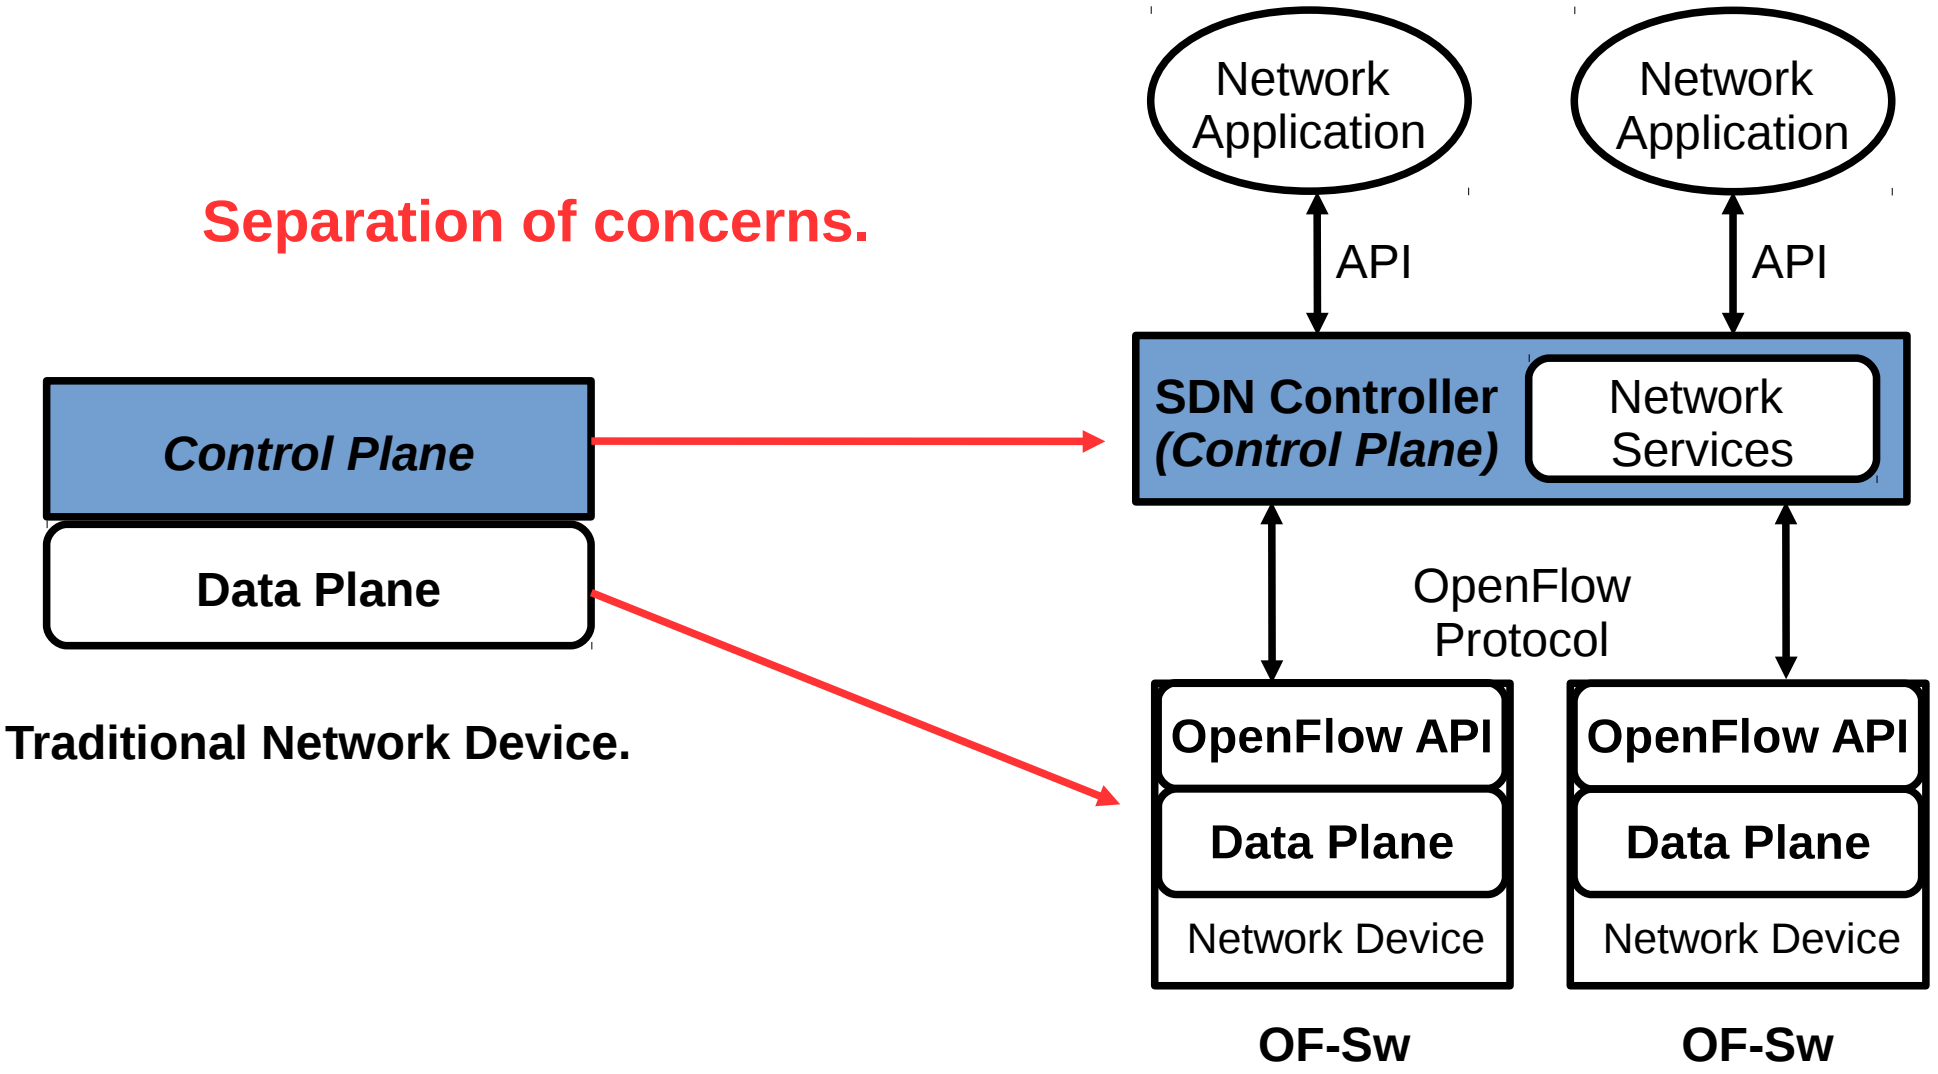


**Traditional Network Device.**

# SDN - Context

**Application Layer**

Traffic Engineering.

**Control Layer**

Device Manager, Routing.

**Infrastructure Layer**

Forwarding.

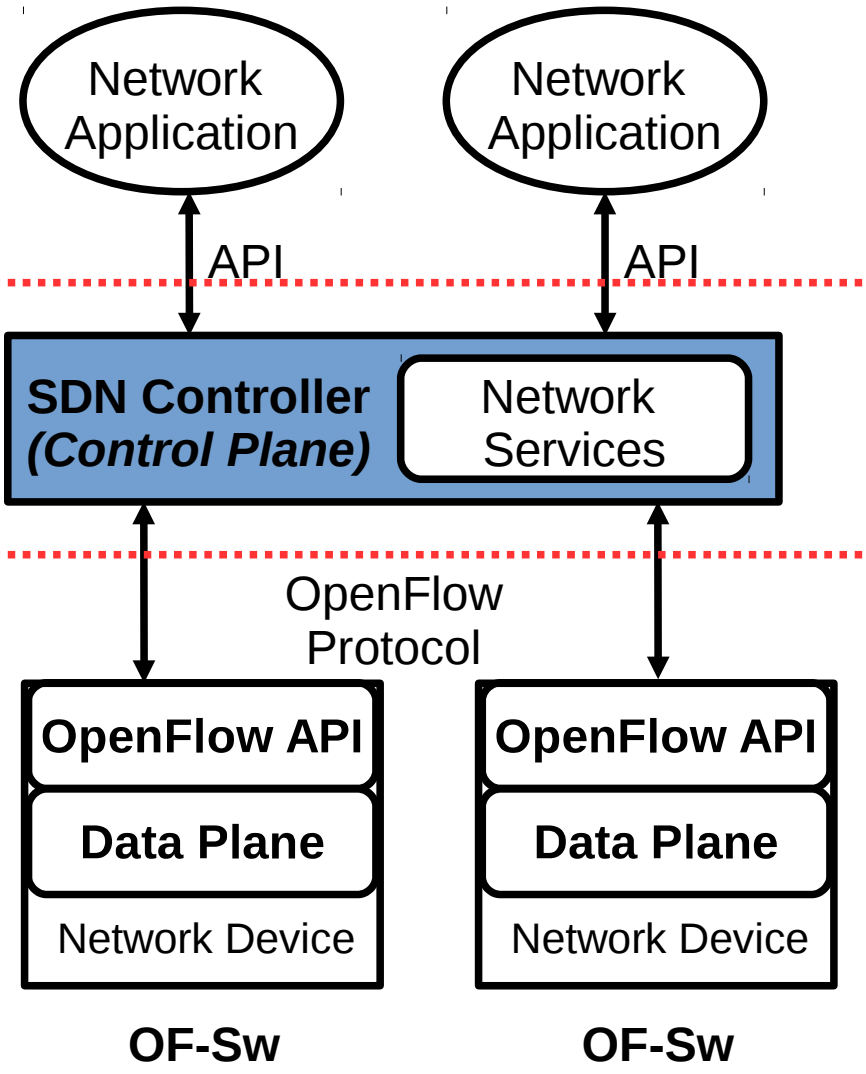

# SDN - Context

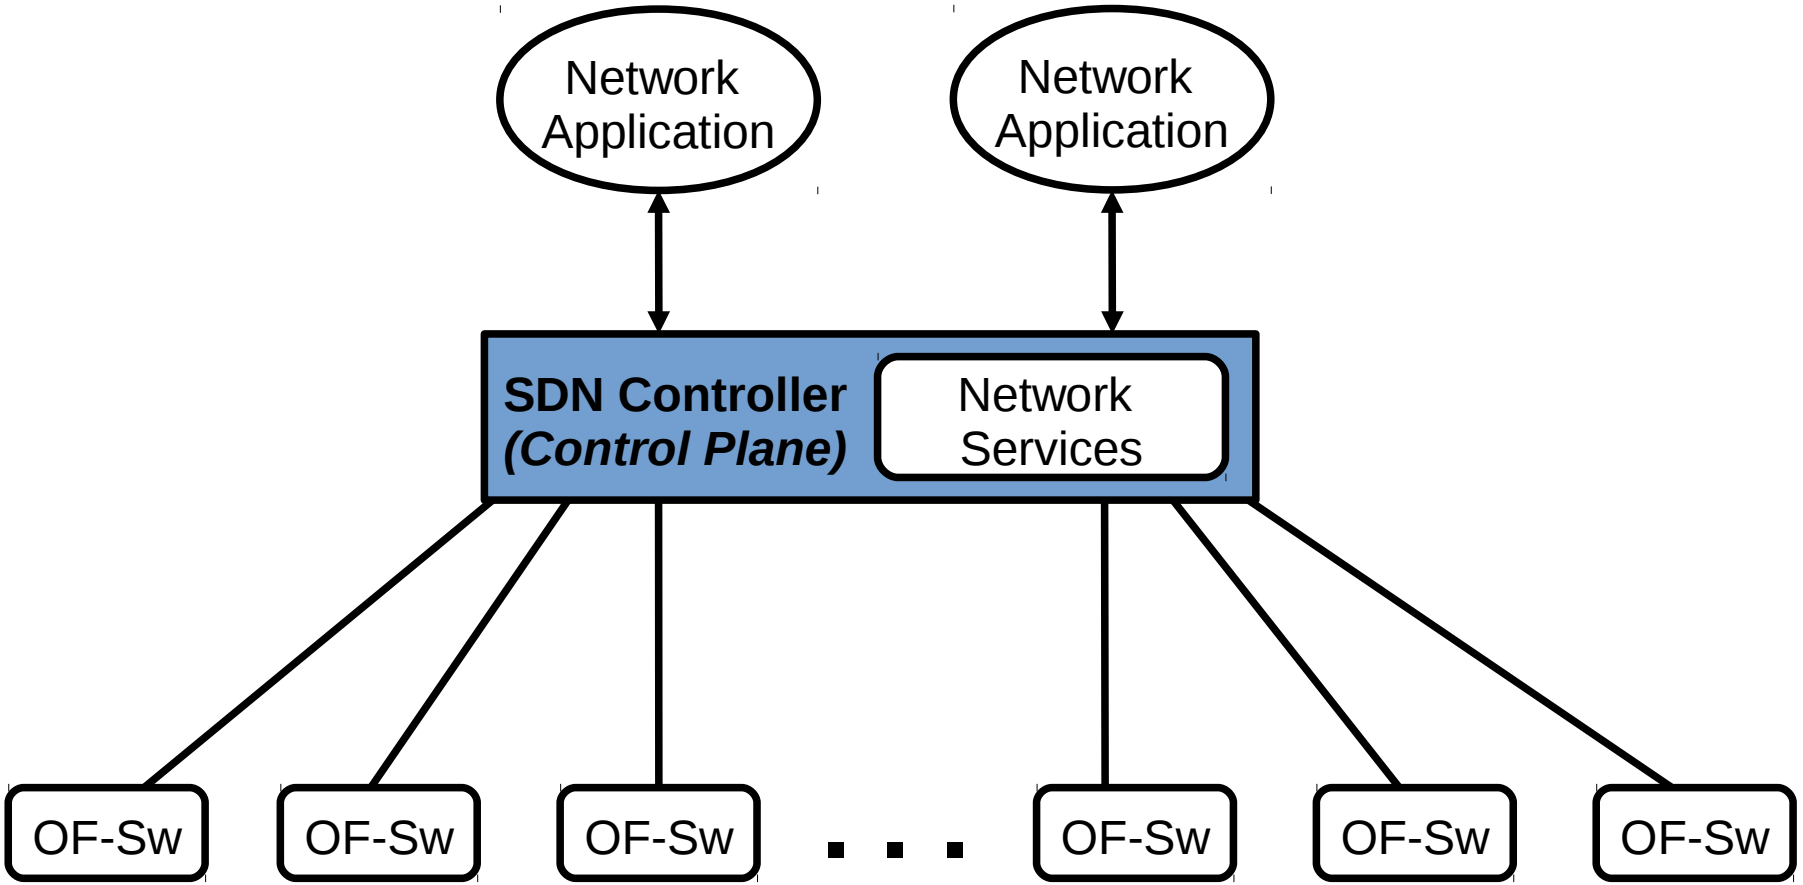
Network Application
Network Application
SDN Controller (Control Plane)
Network Services
OF-Sw
OF-Sw
OF-Sw
OF-Sw
OF-Sw
OF-Sw

# SDN - Context

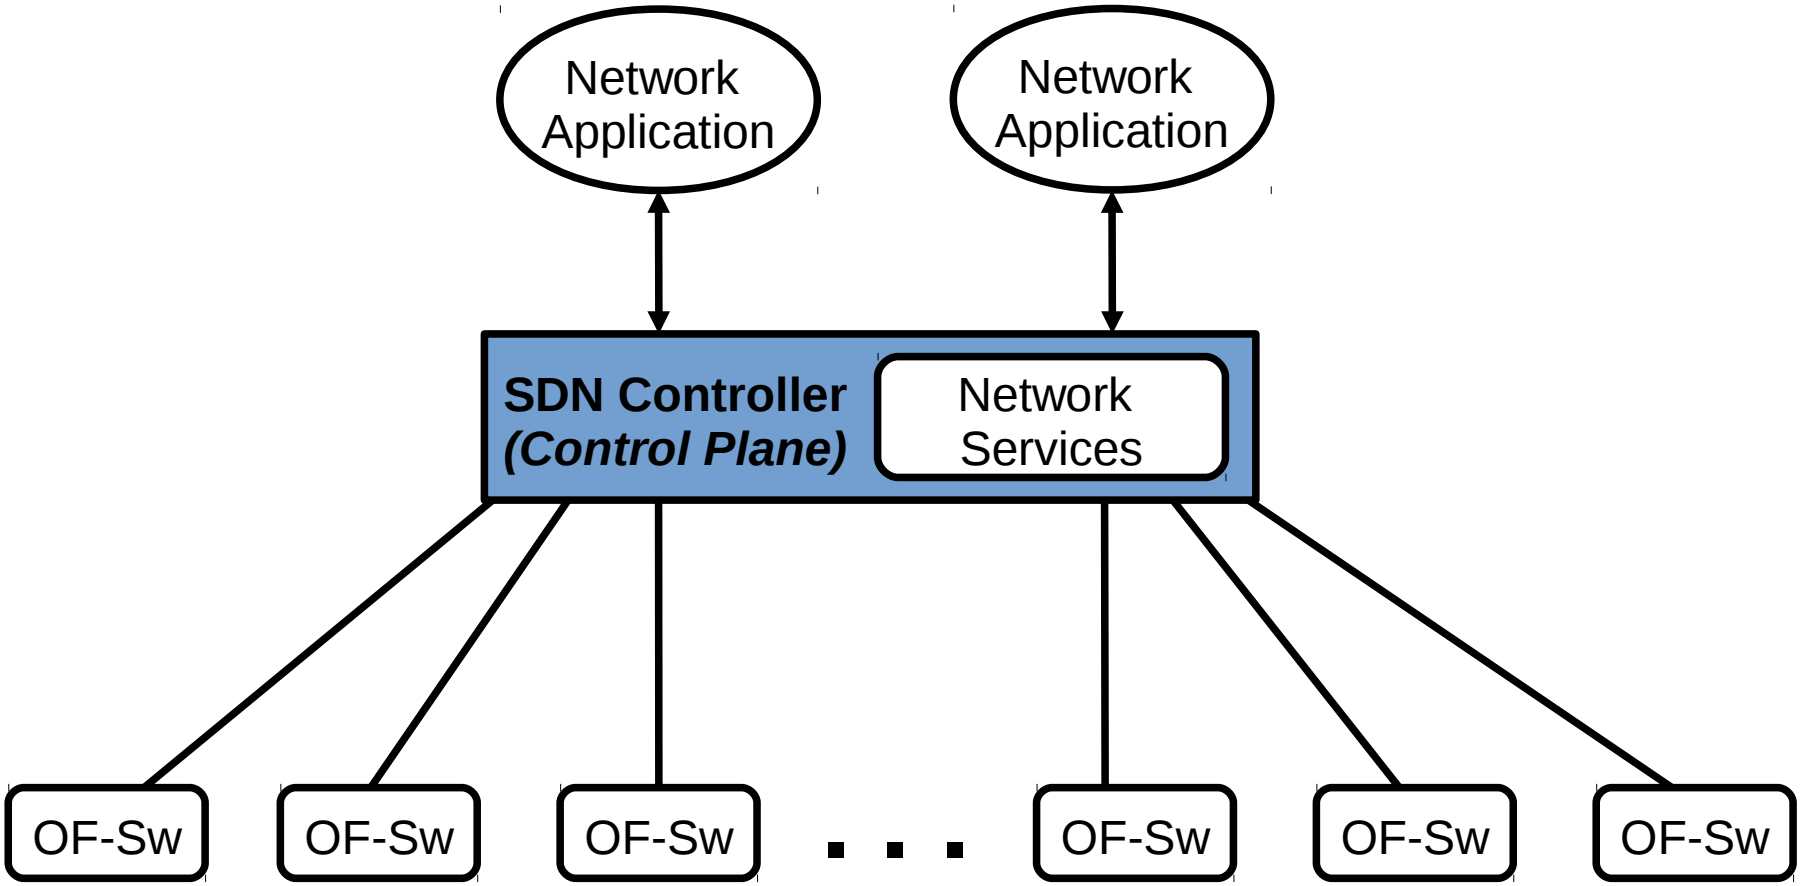


**It does not scale.**
**&**
**Single point of failure.**

# SDN - Context

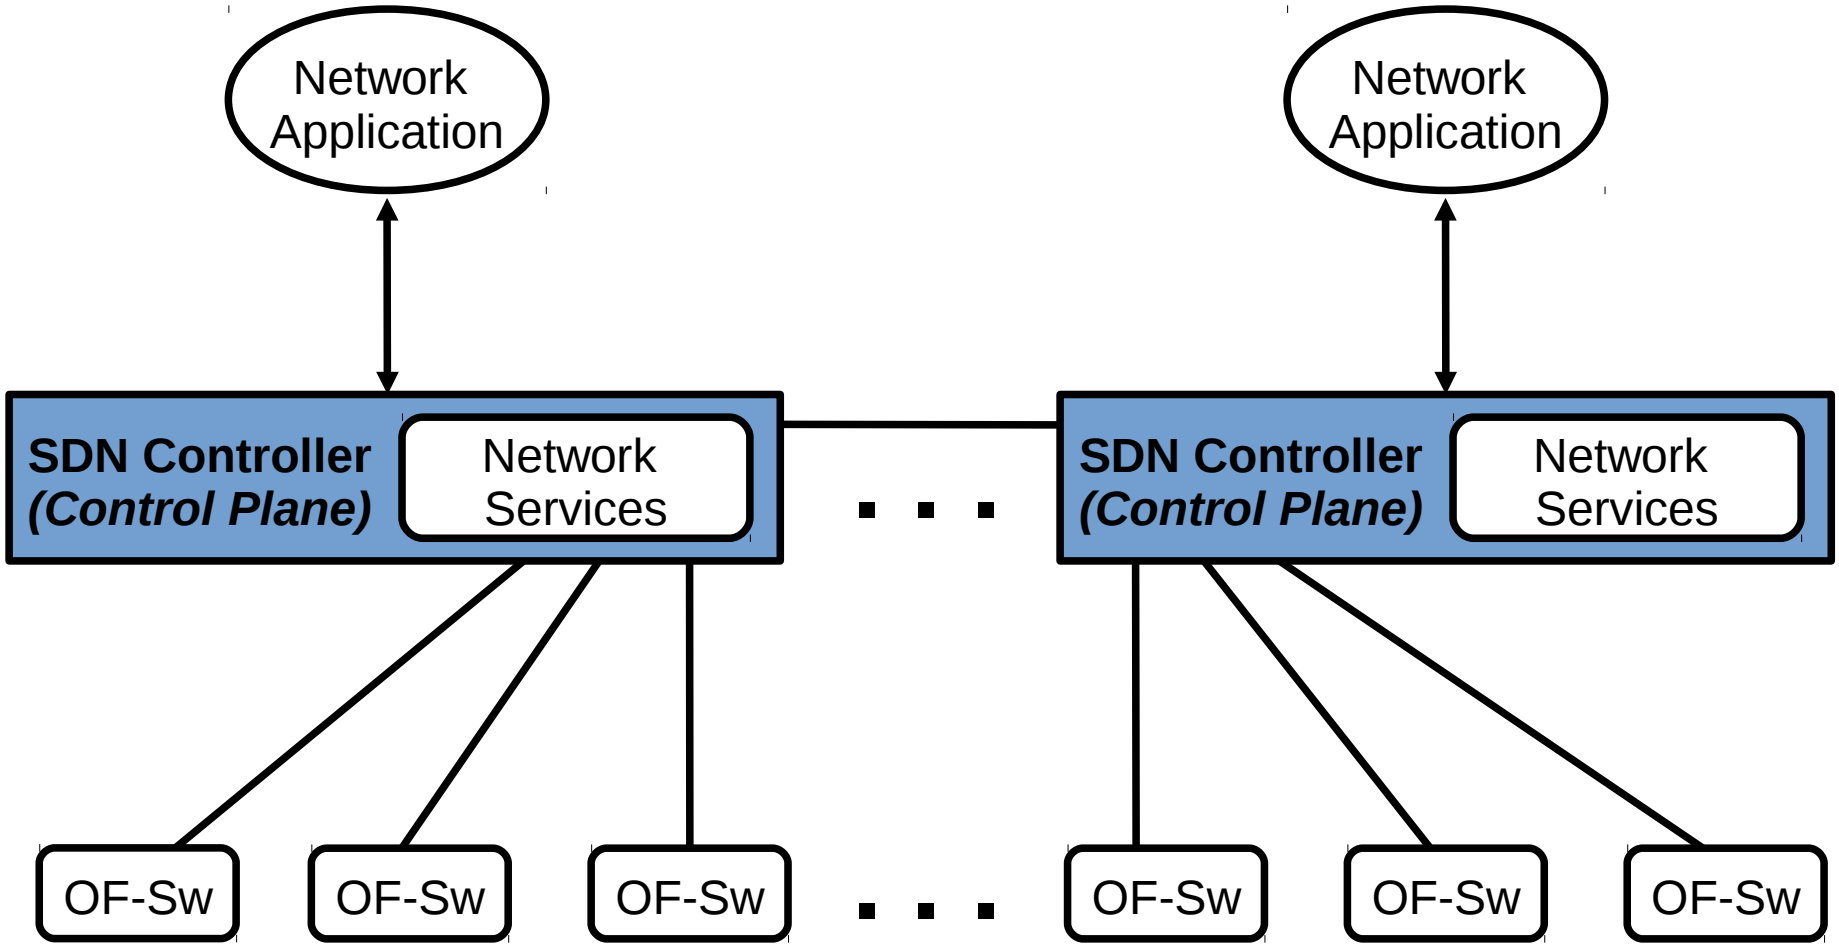


**Scales better, but...**

# SDN - Context

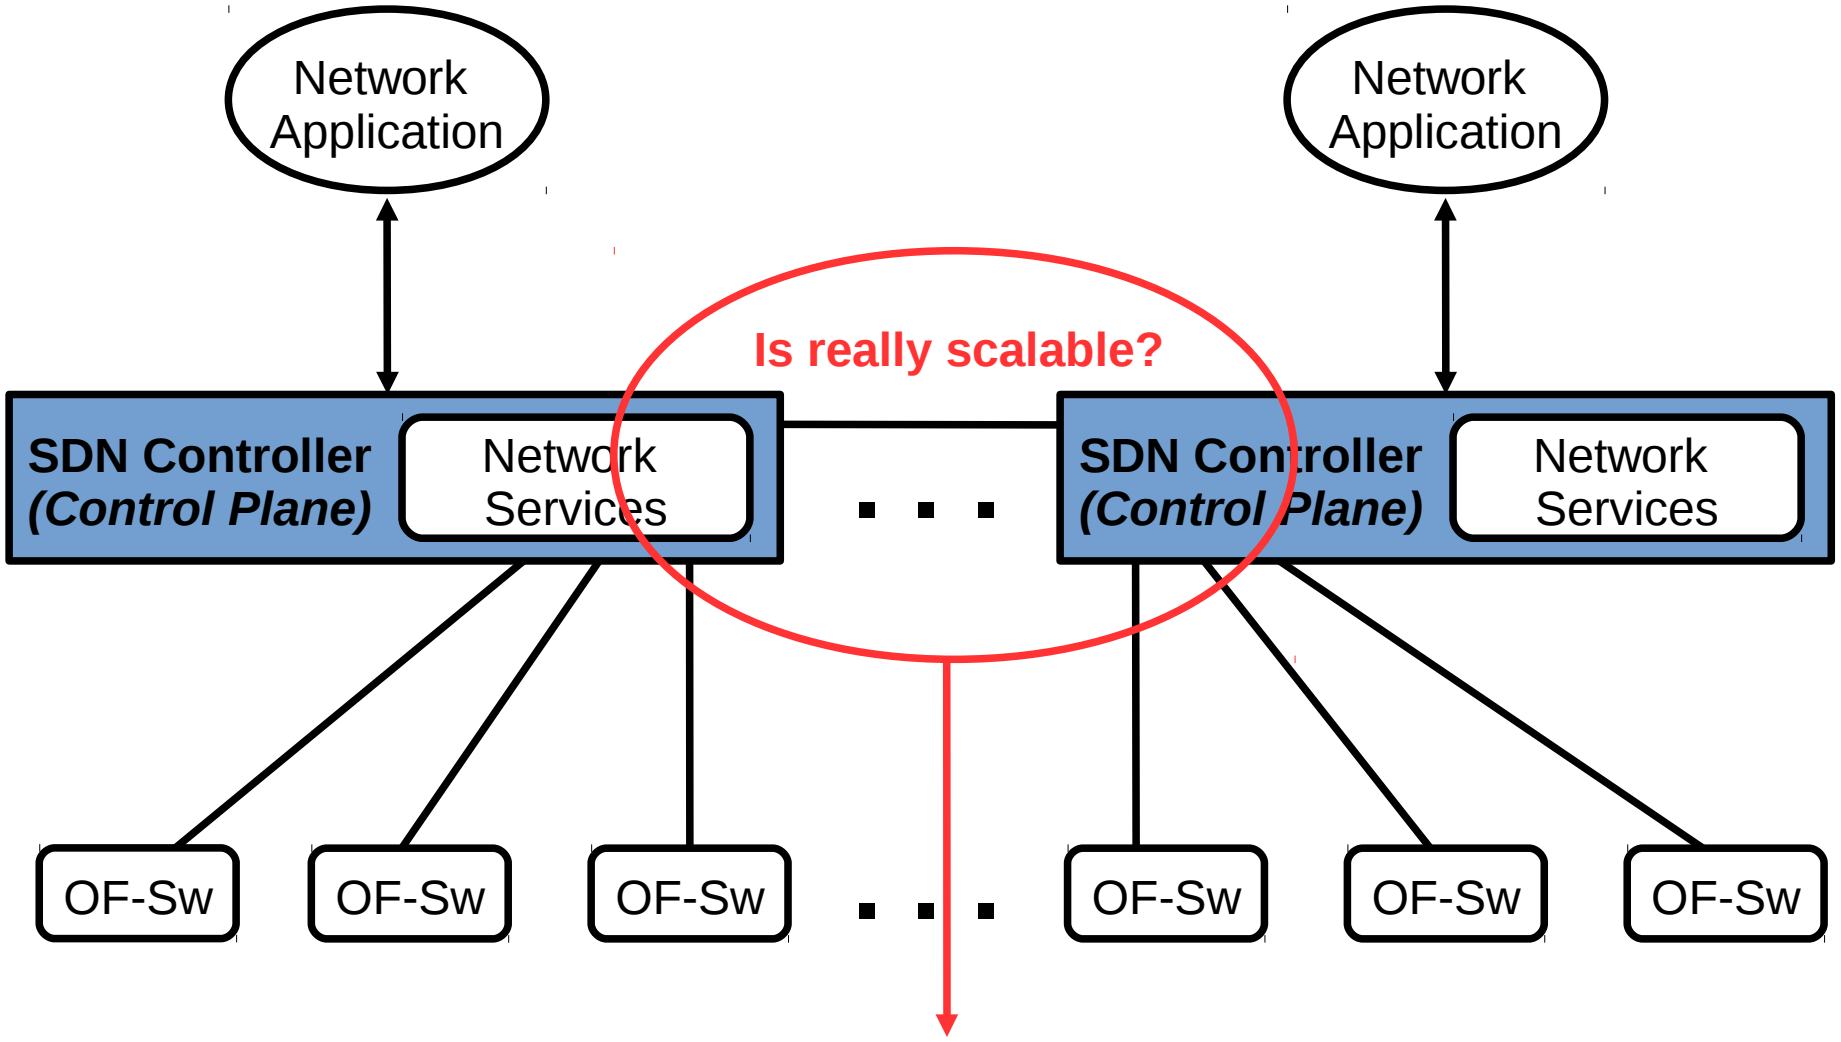


Synchonized?, != network view?, eventually consistency?, consensus algorithms?

# SDN - Context

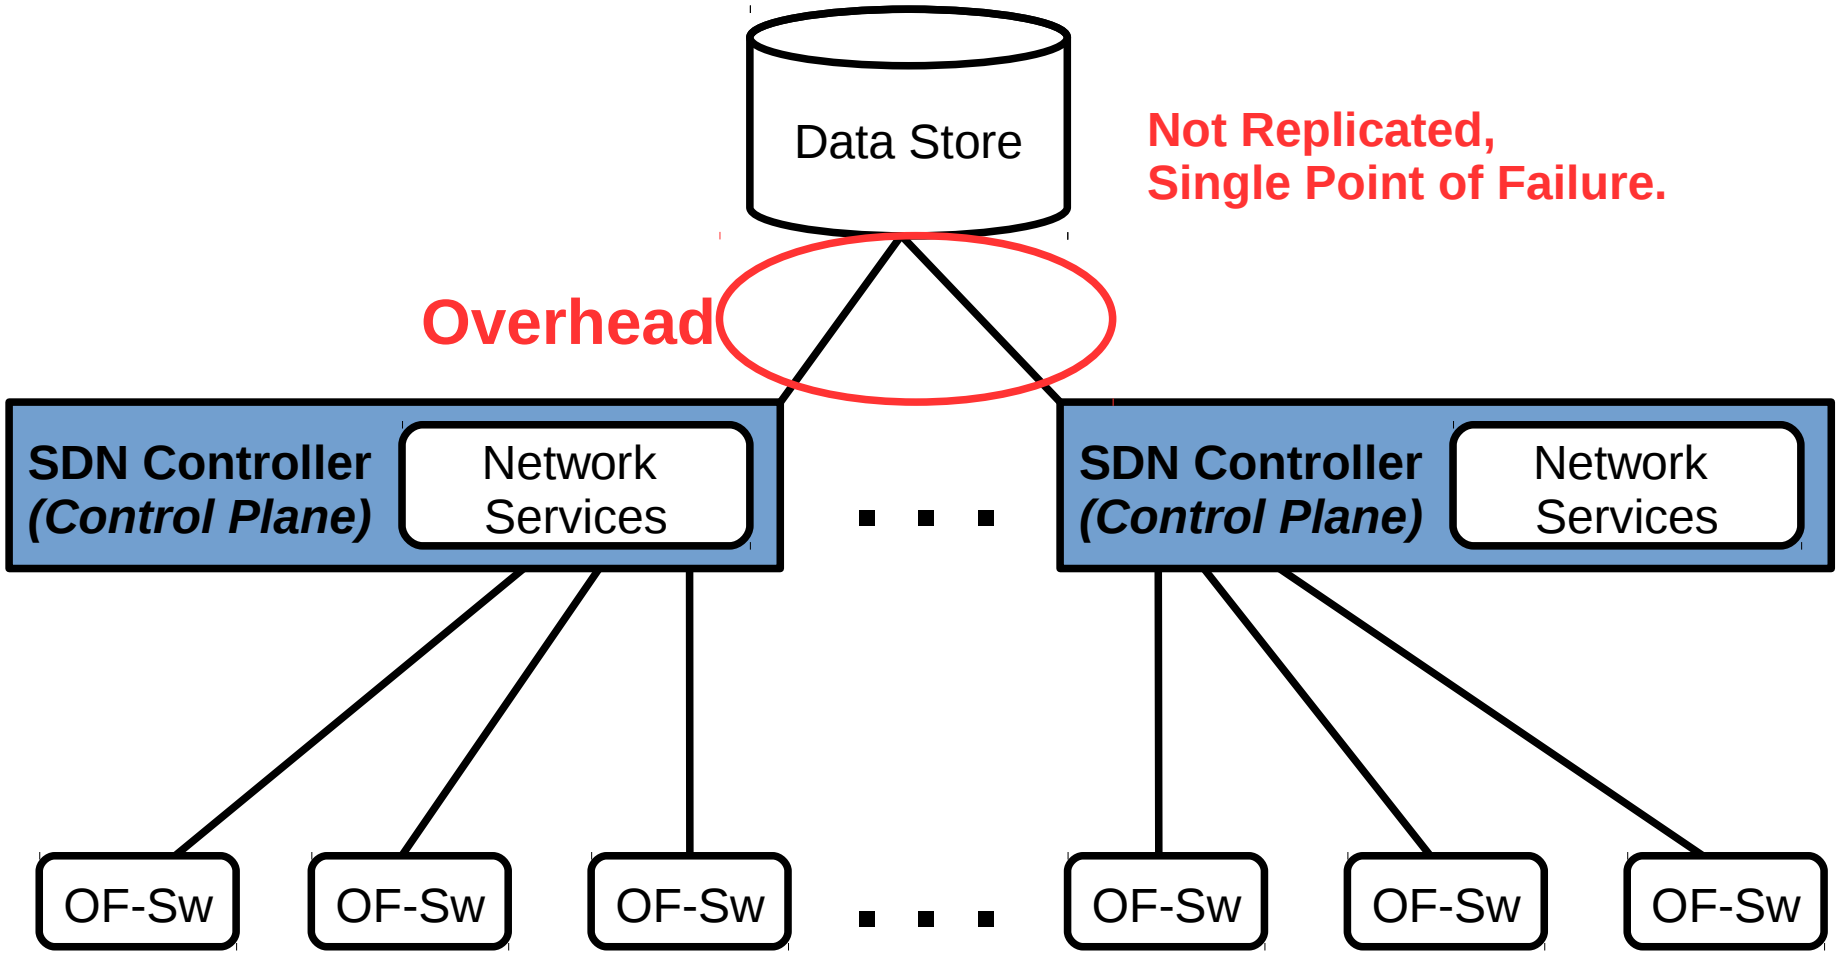
Data Store
Not Replicated,
Single Point of Failure.
Overhead
SDN Controller
(Control Plane)
Network
Services
SDN Controller
(Control Plane)
Network
Services
OF-Sw
OF-Sw
OF-Sw
OF-Sw
OF-Sw
OF-Sw

# SDN - Context

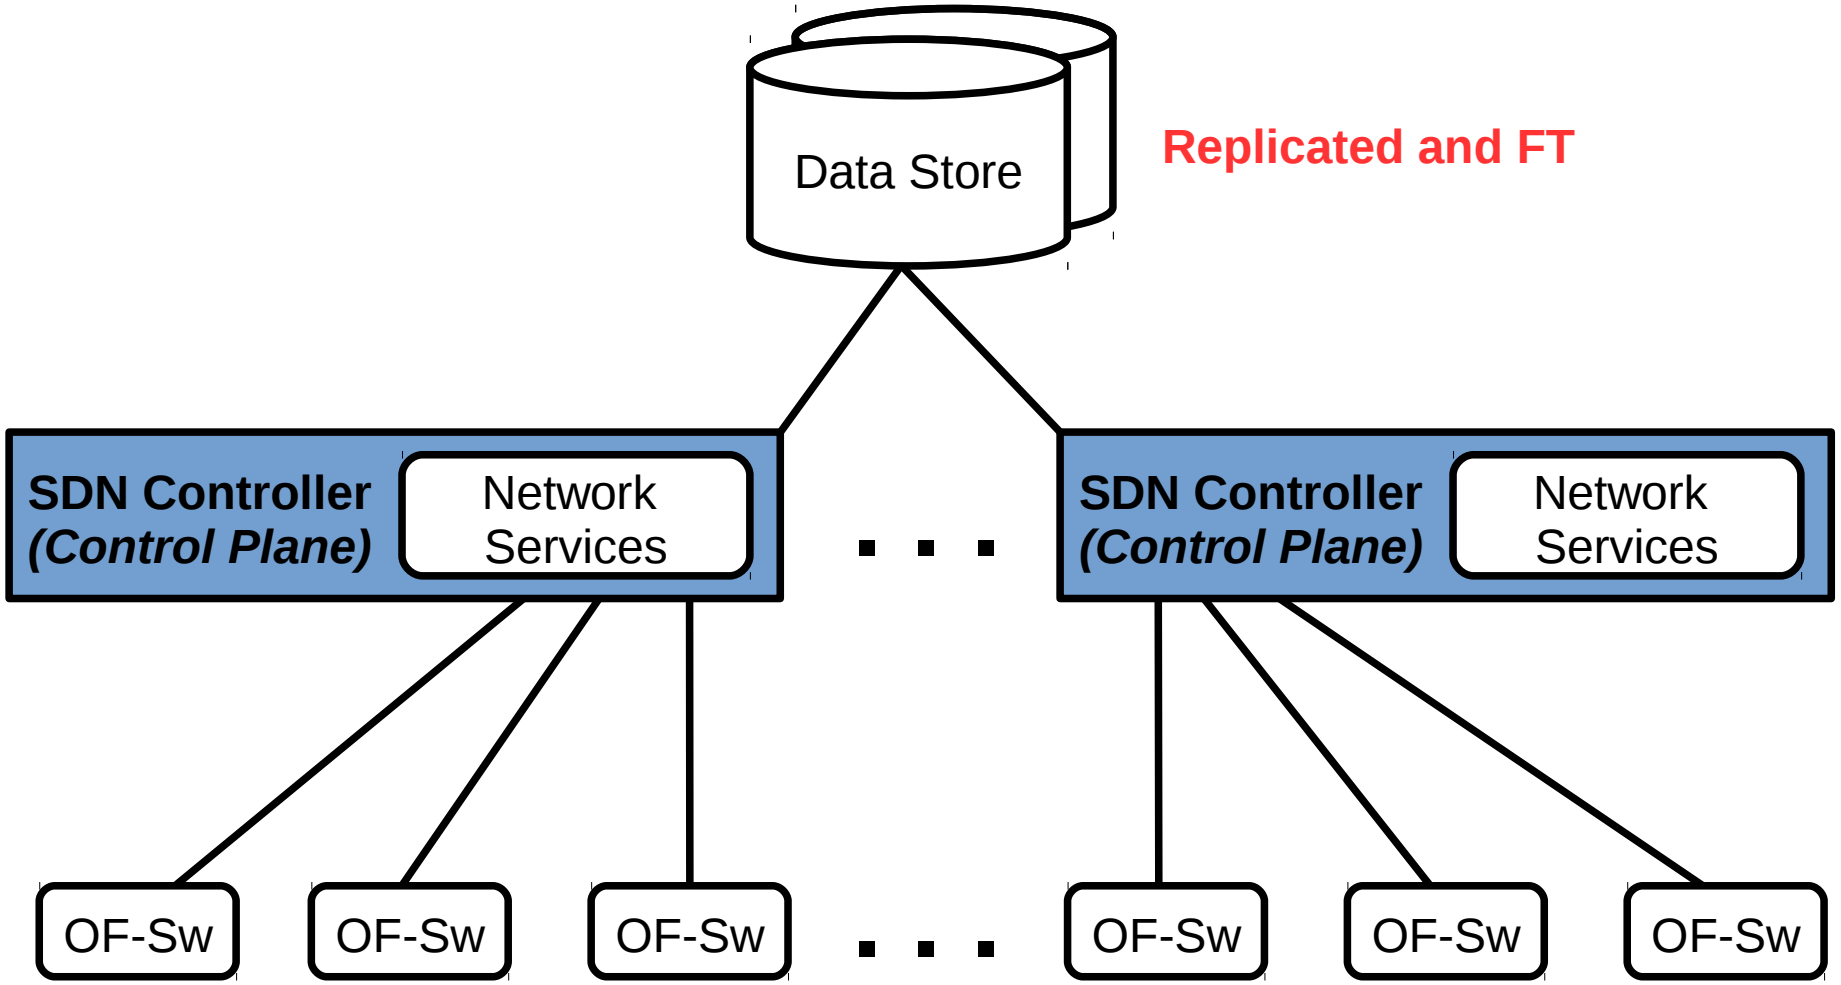


Eventually consistency scales out but problems might appear

# Outline

- SDN – Context
- (Strong) Consistency Matters
- Proposed Controller Architecture
- Ravana x Floodlight Tracing
- Floodlight Pipeline Tracing
- Como Evitar Sobrecarga
- Ganho esperado

# (Strong) Consistency Matters

- Why?

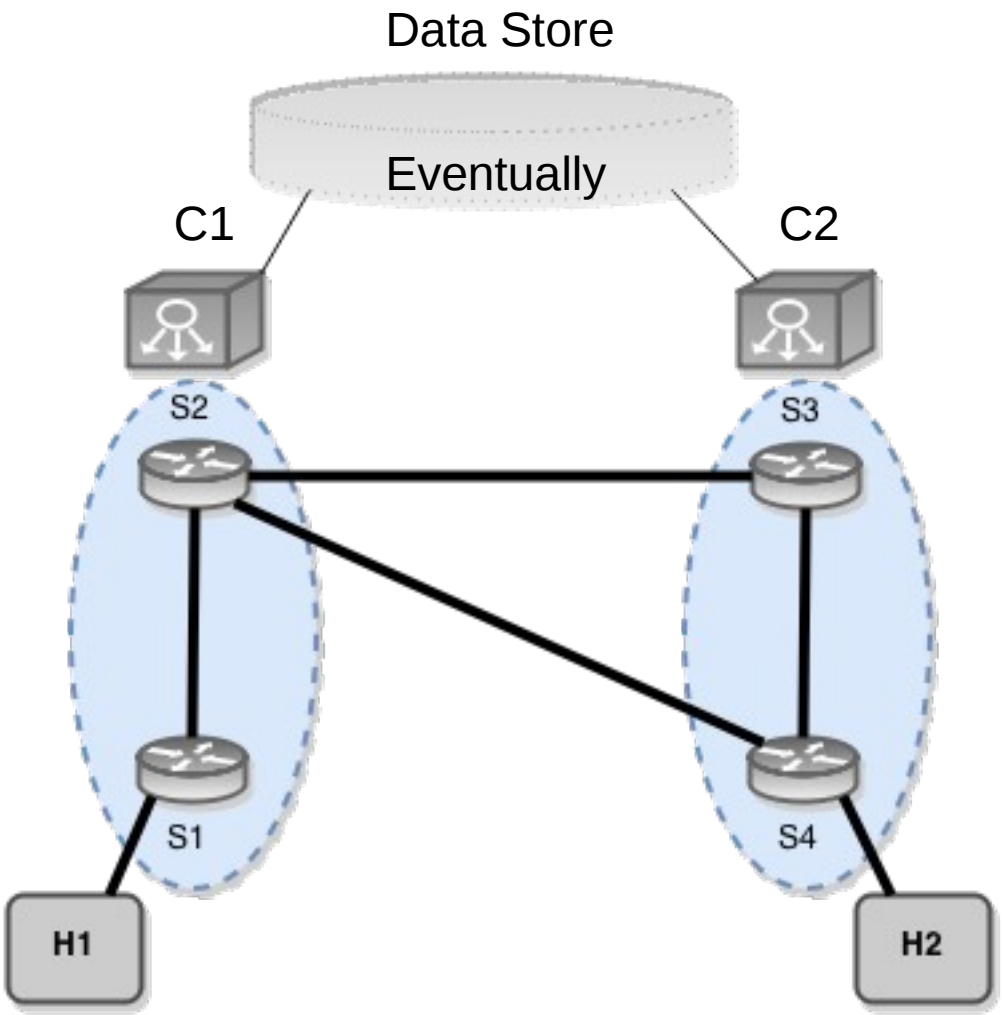

# (Strong) Consistency Matters

Initial flow / path

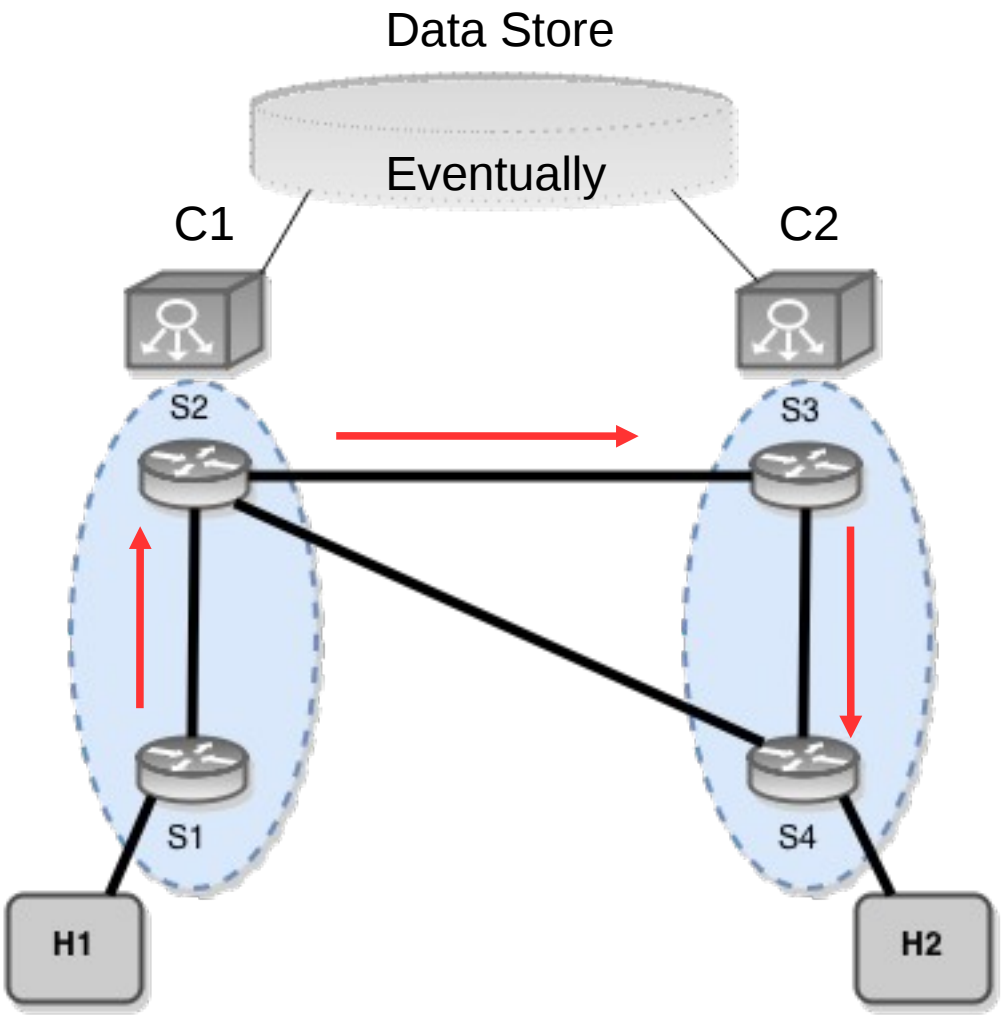

# (Strong) Consistency Matters

Initial flow / path

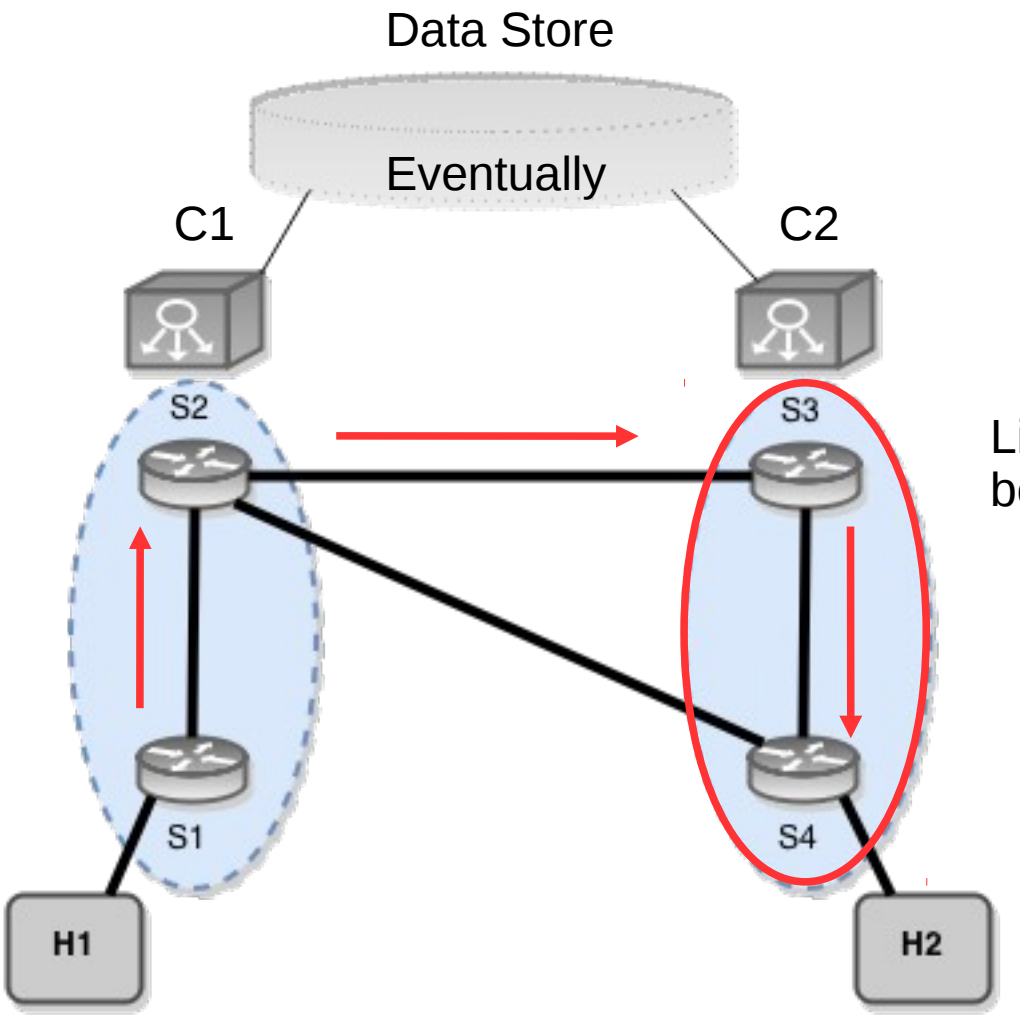


Link between S3-S4 is becoming congested.

# (Strong) Consistency Matters

Initial flow / path

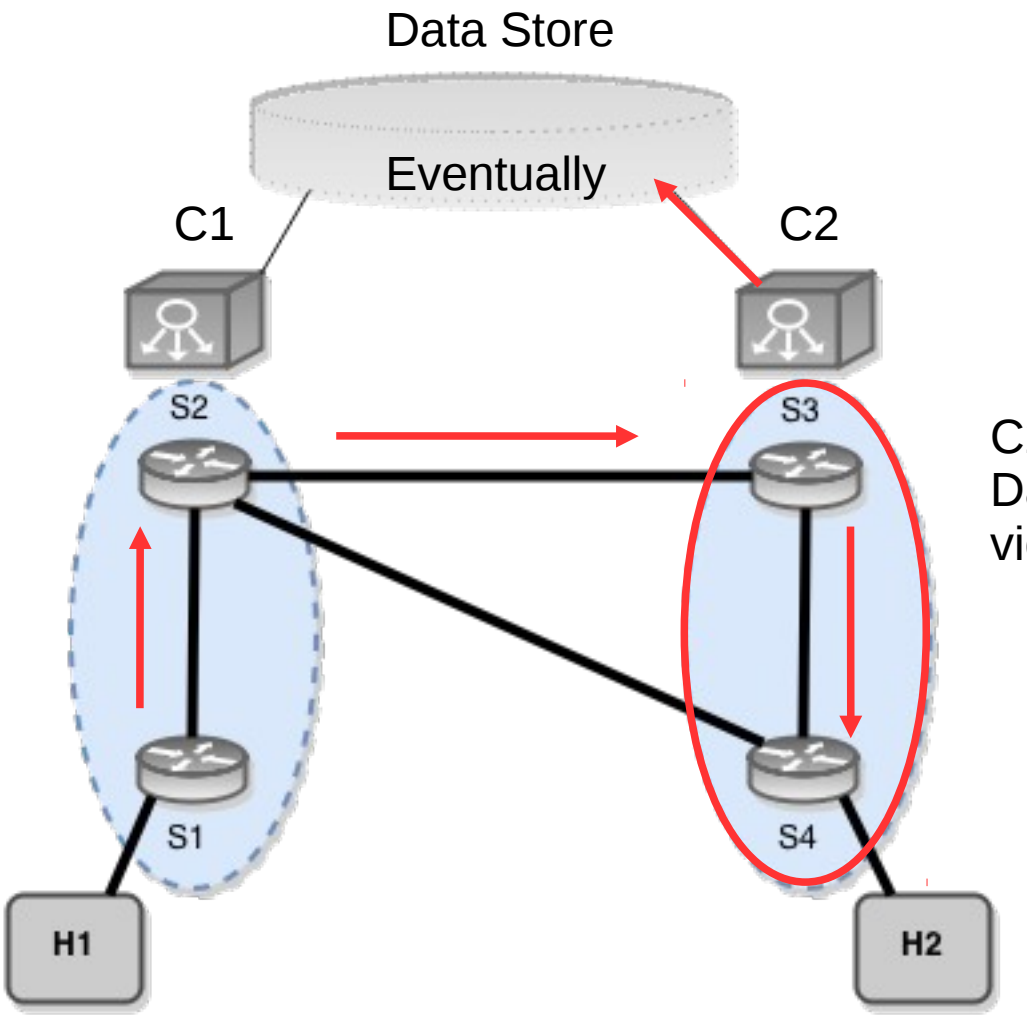


C2 writes changes into Data Store (network view).

# (Strong) Consistency Matters

Initial flow / path

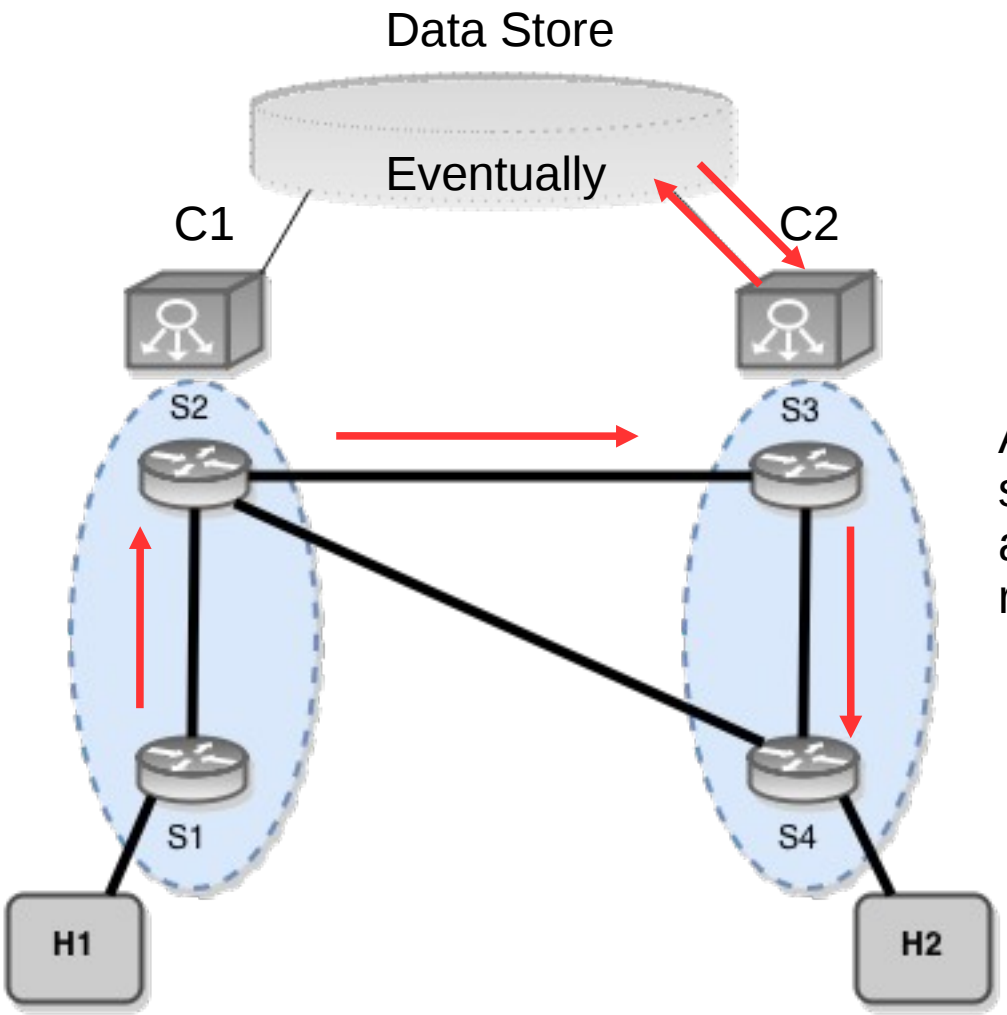


An eventually data store will reply as soon as it process the request.

# (Strong) Consistency Matters

Initial flow / path

After change

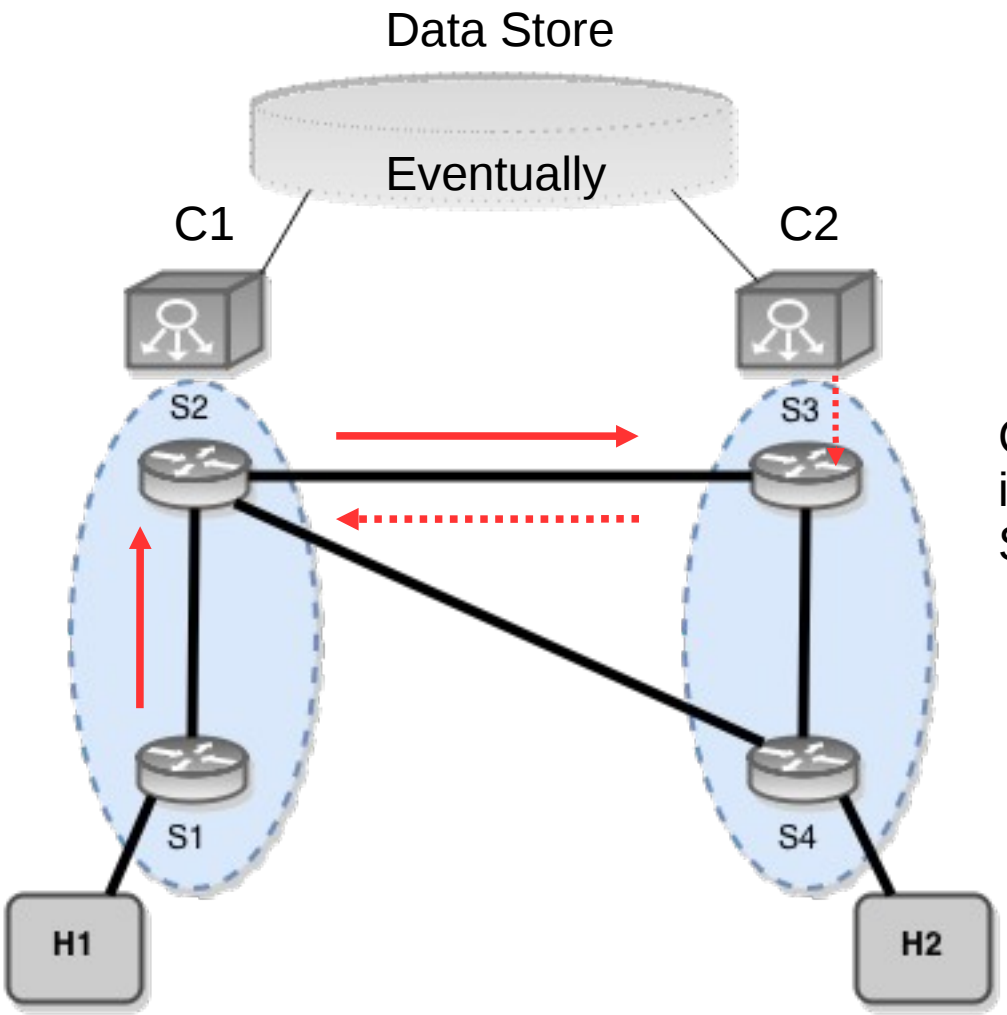


C2 will immediately install the new rules at S3.

# (Strong) Consistency Matters

- Consistency loop scenario.

Initial flow / path

After change

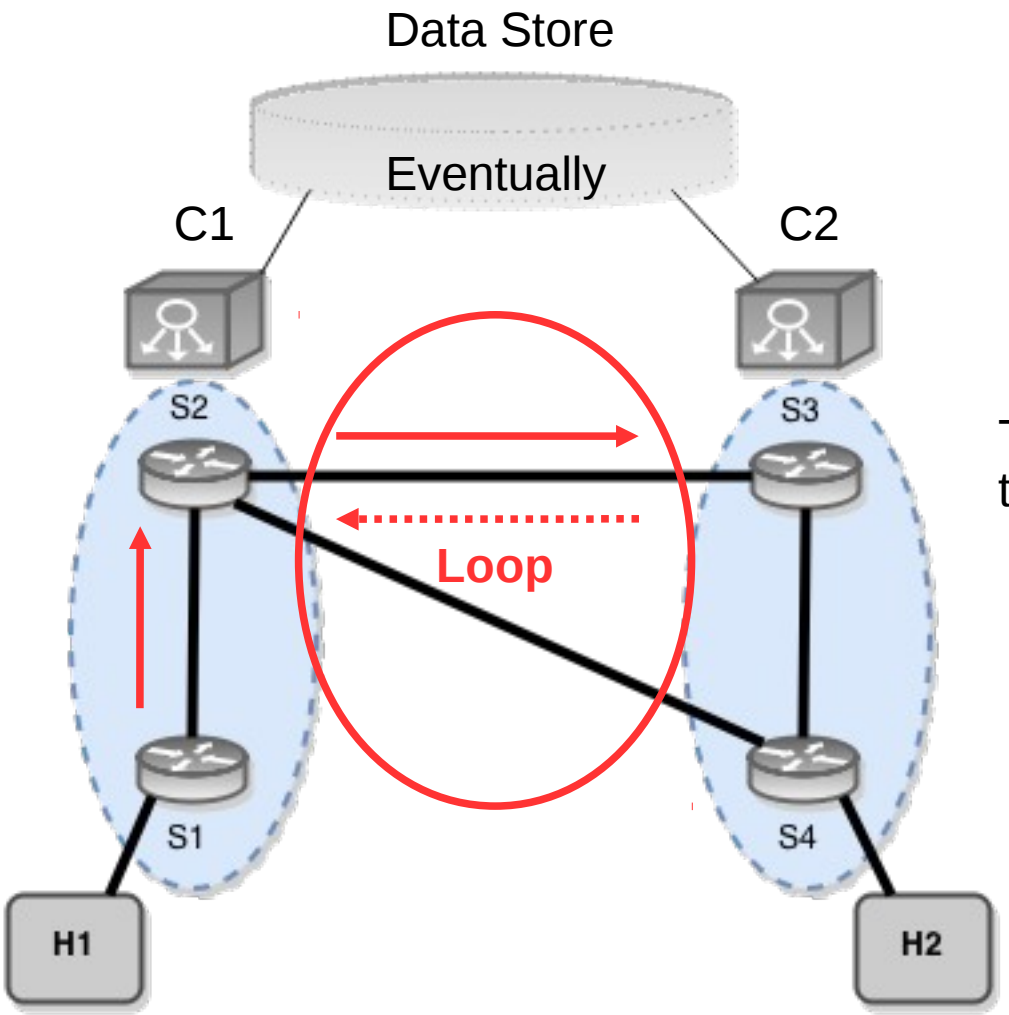


This will create a transient loop.

# (Strong) Consistency Matters

Initial flow / path

After change

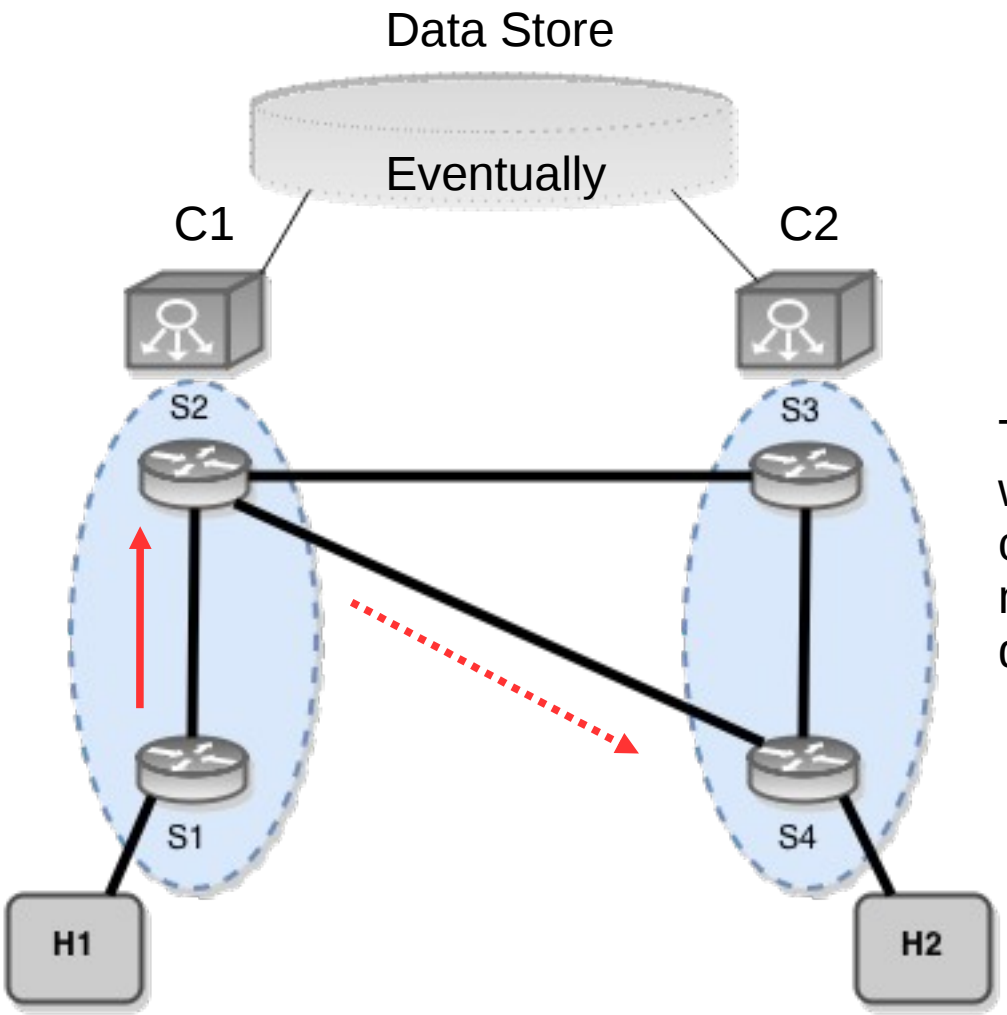


The network anomaly will eventually be corrected when the network state in the data store converges.

# (Strong) Consistency Matters

Initial flow / path

After change

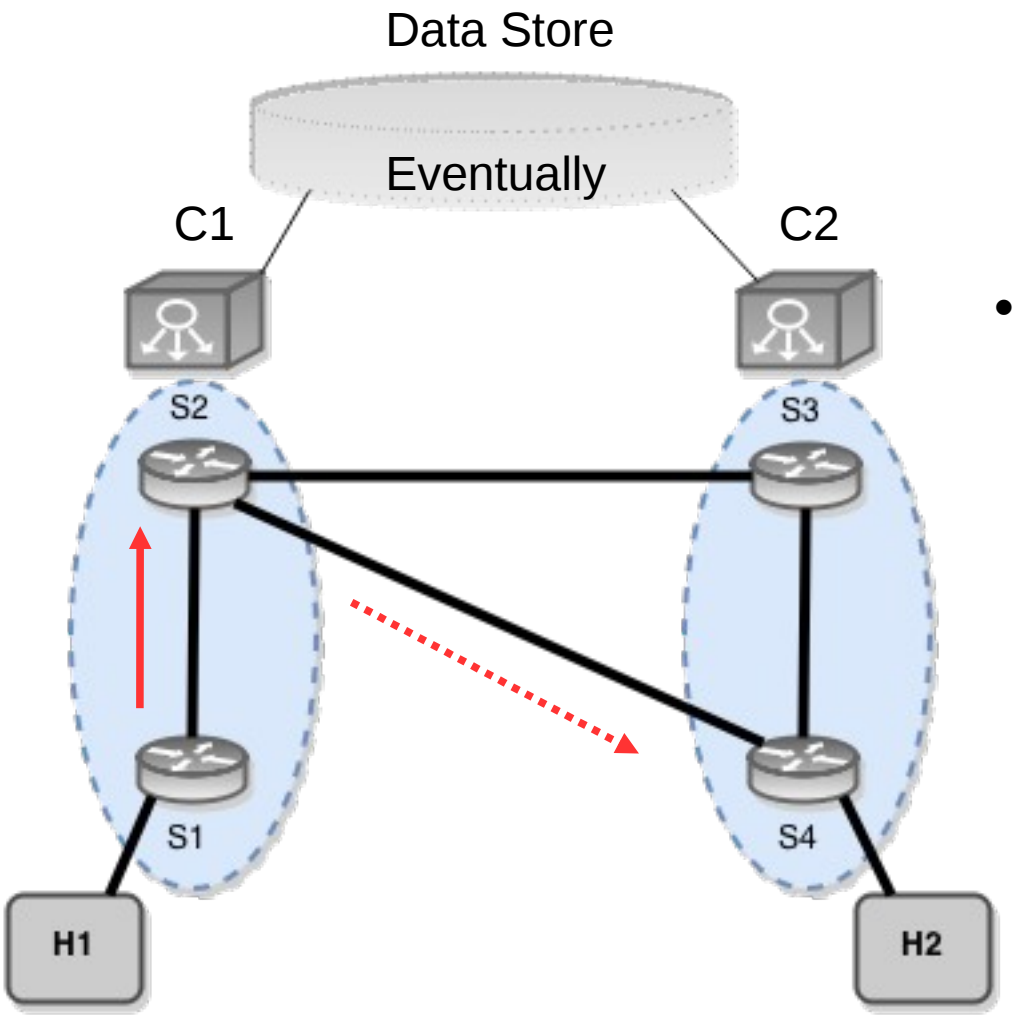


- Consequences:
  - Hiccup in VoIP call.
  - Lost server connection.
  - Security breaches.

# Outline

- SDN – Context
- (Strong) Consistency Matters
- Proposed Controller Architecture
- Ravana x Floodlight Tracing
- Floodlight Pipeline Tracing
- Como Evitar Sobrecarga
- Ganho esperado

# Proposed Controller Architecture

- Controllers of different network domains coordinate their actions using a logically centralized (consistent and fault-tolerant) data store.

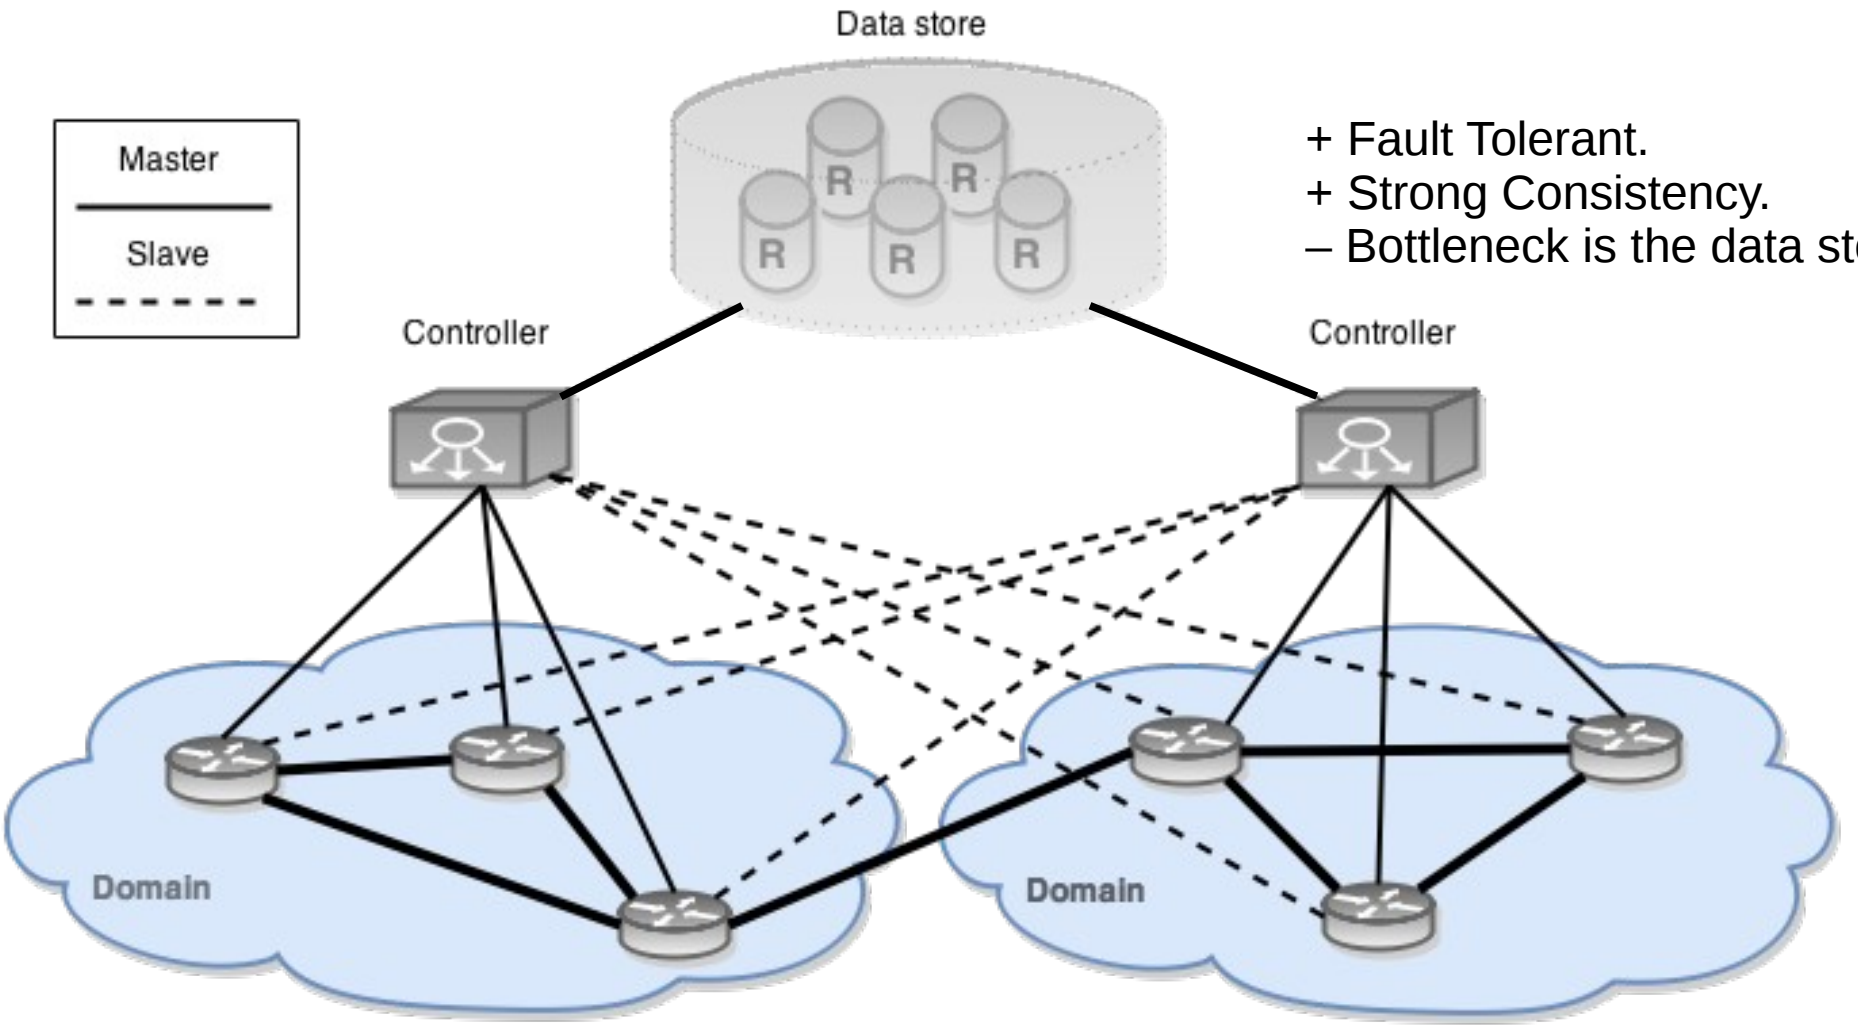


+ Fault Tolerant.
+ Strong Consistency.
– Bottleneck is the data store.

# Proposed Controller Architecture

- Controllers of different network domains coordinate their actions using a logically centralized (consistent and fault-tolerant) data store.

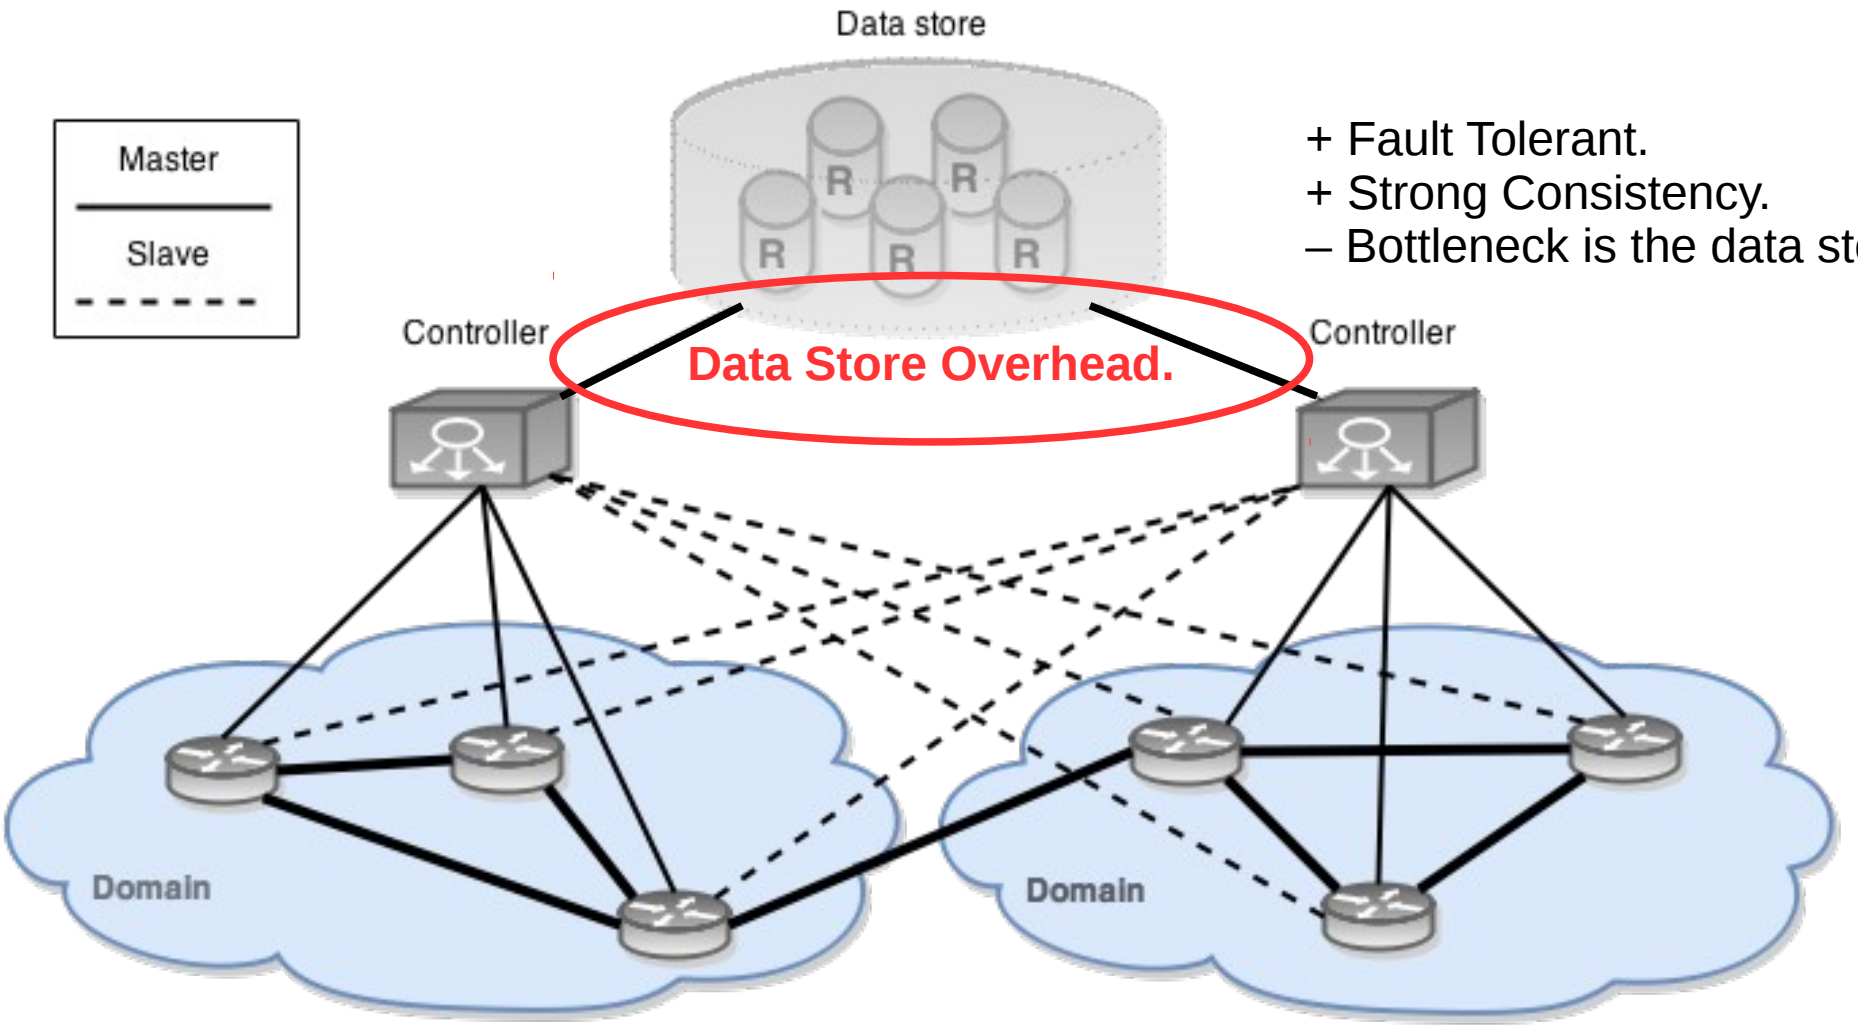


+ Fault Tolerant.
+ Strong Consistency.
– Bottleneck is the data store.

# Outline

- SDN – Context
- (Strong) Consistency Matters
- Proposed Controller Architecture
- Ravana x Floodlight Tracing
- Floodlight Pipeline Tracing
- Como Evitar Sobrecarga
- Ganho esperado

# Ravana x Floodlight tracing

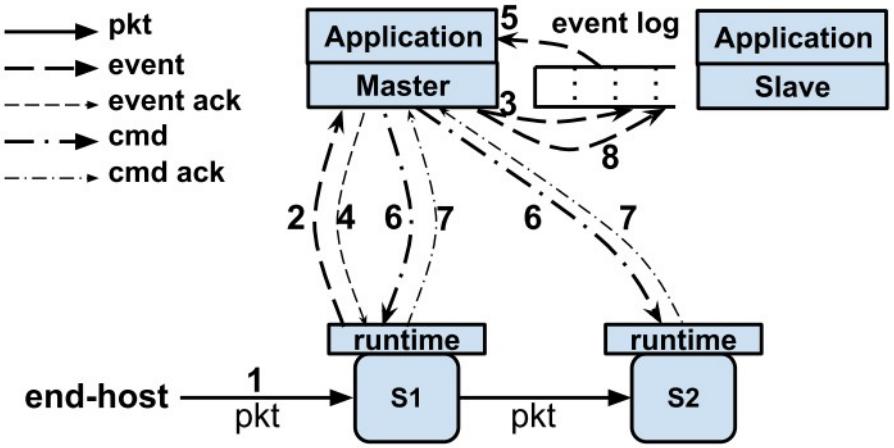


Figure 4: Steps for processing a packet in Ravana.

Floodlight tracing

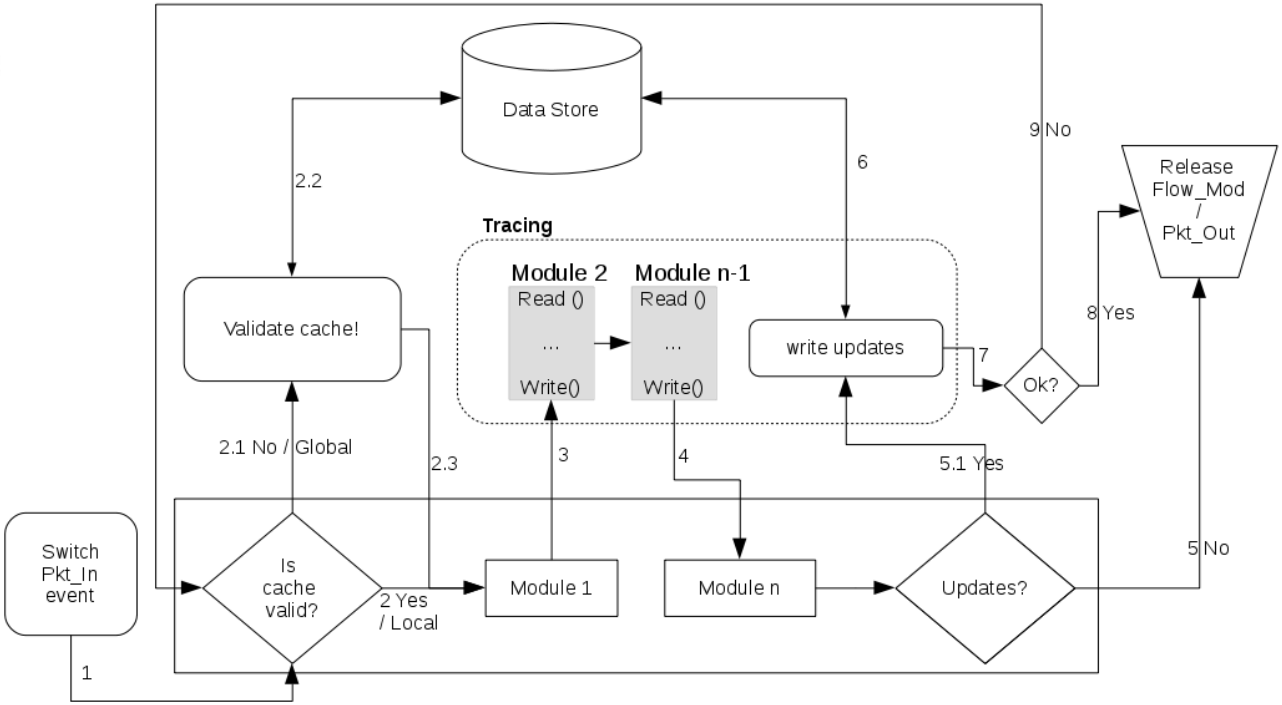

# Outline

- SDN – Context
- (Strong) Consistency Matters
- Proposed Controller Architecture
- Ravana x Floodlight Tracing
- Floodlight Pipeline Tracing
- Como Evitar Sobrecarga
- Ganho esperado

# Floodlight Pipeline Tracing

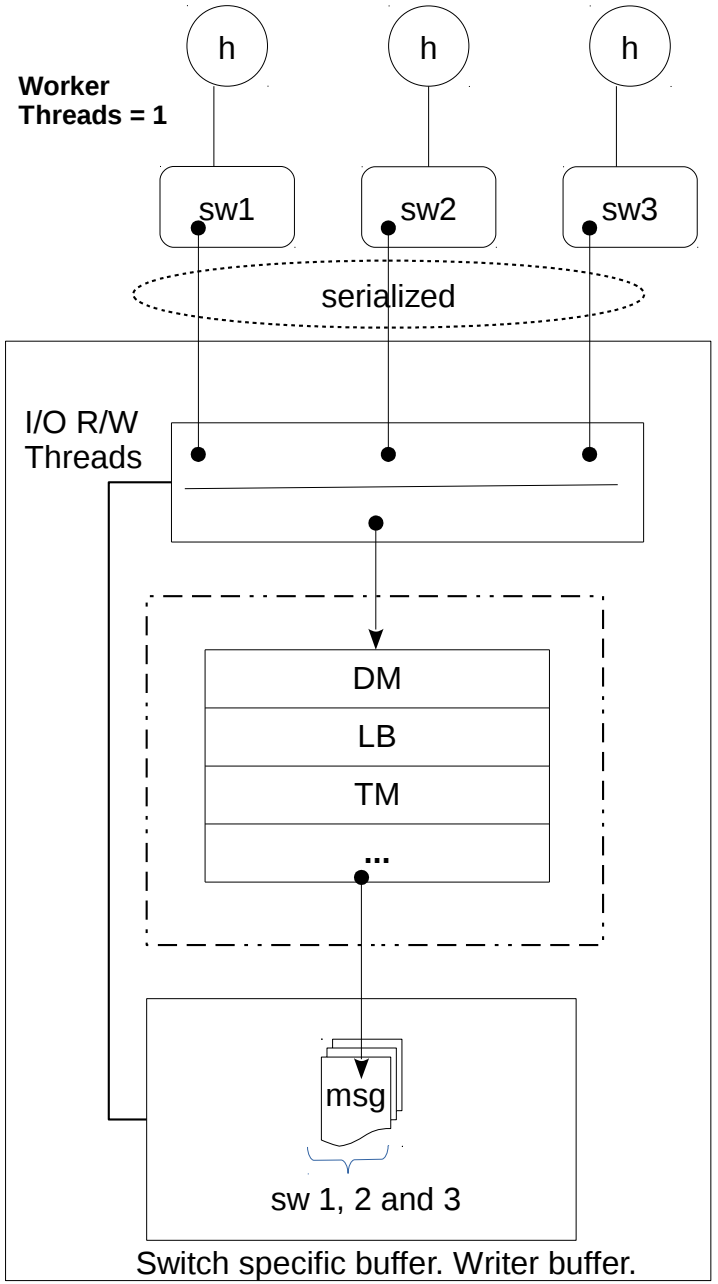

# Floodlight Pipeline Tracing

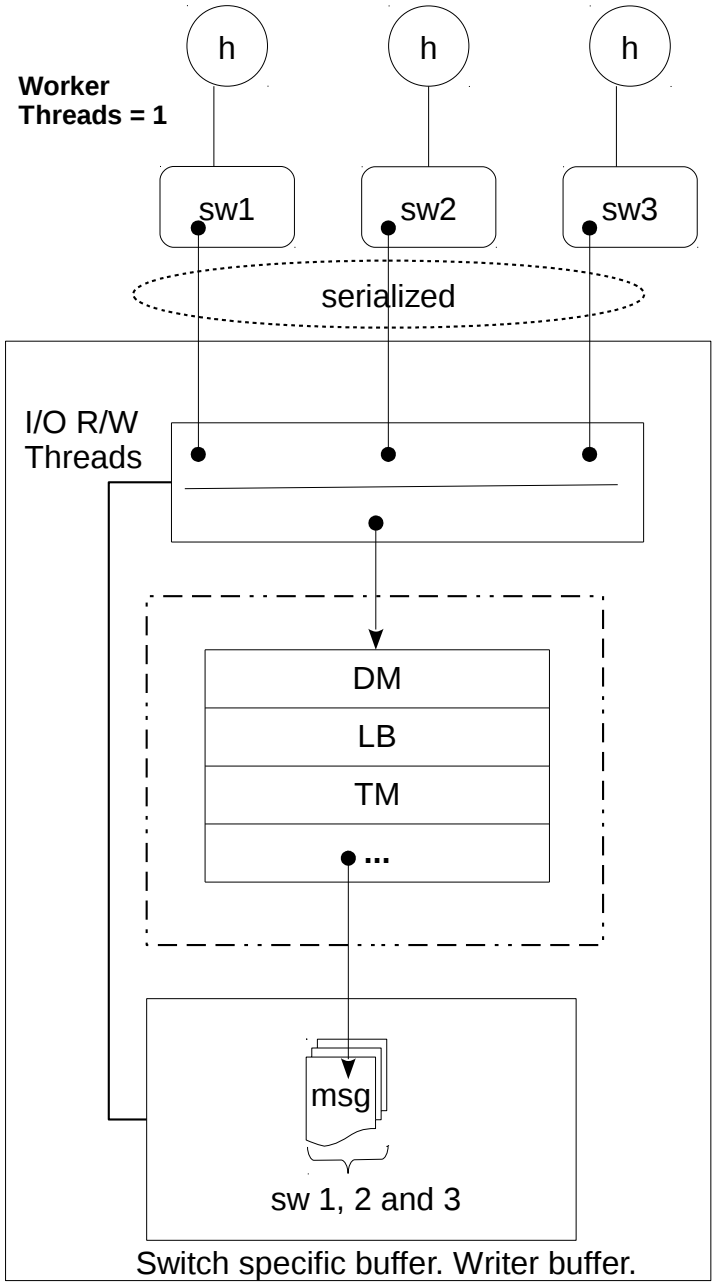


## Definição Estrutura de Dados

### Sem tracing

```
ConcurrentHashMap<Long, Device> deviceMap;//original
```

### Com tracing

```
ConcurrentHashMapTracing<Long, Device> deviceMap;
```

## Inicialização Estrutura de Dados

### Sem tracing

```
deviceMap = new ConcurrentHashMap<Long, Device>();//original
```

### Com tracing

```
/**
* @param 1: Data structure (HashMap, ArrayList...)
* @param 2: Unique name on system (there is no conflict resolution).
* @param 3: Scope: GLOBAL / LOCAL.
*/
deviceMap = tarService.createConcurrentHashMapTracing(new
    ConcurrentHashMap<Long, Device>(), "deviceMap", Scope.GLOBAL);
```

# Floodlight Pipeline Tracing

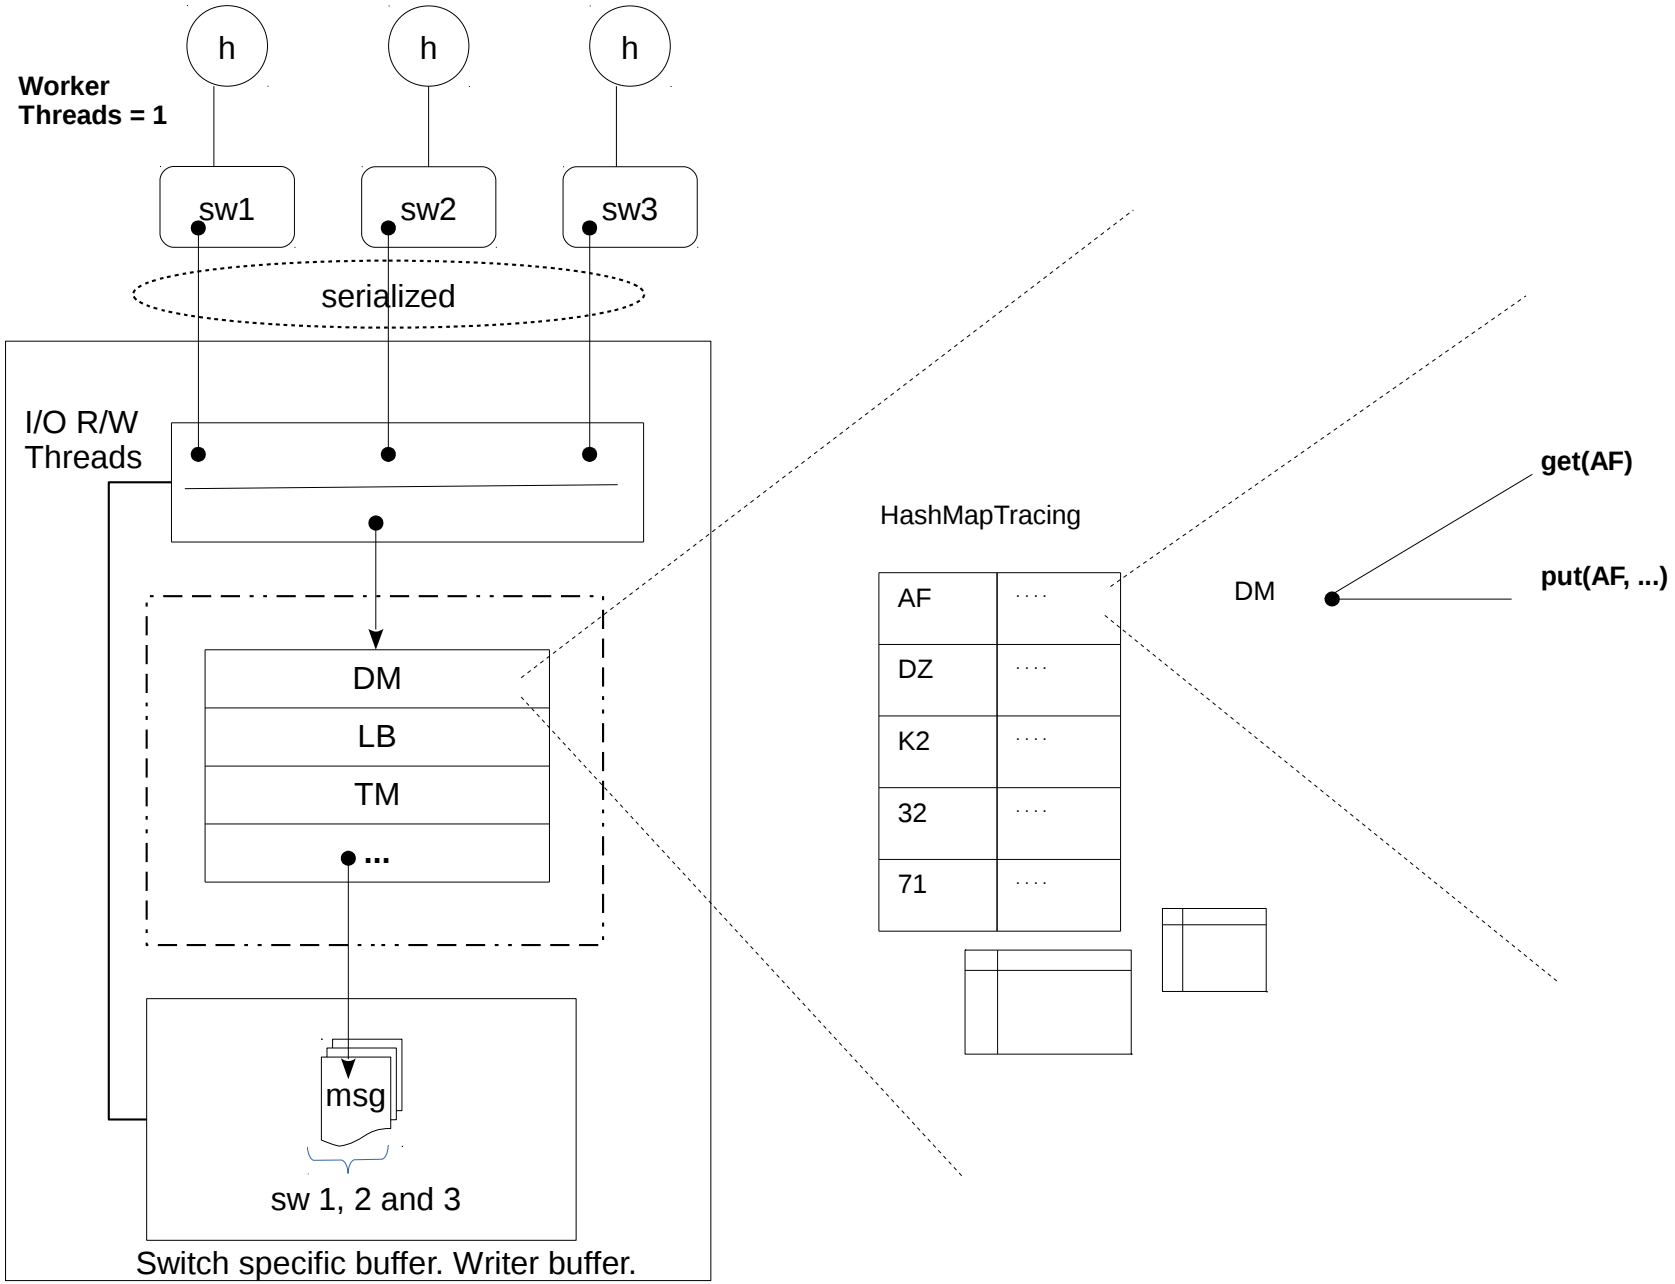

# Floodlight Pipeline Tracing

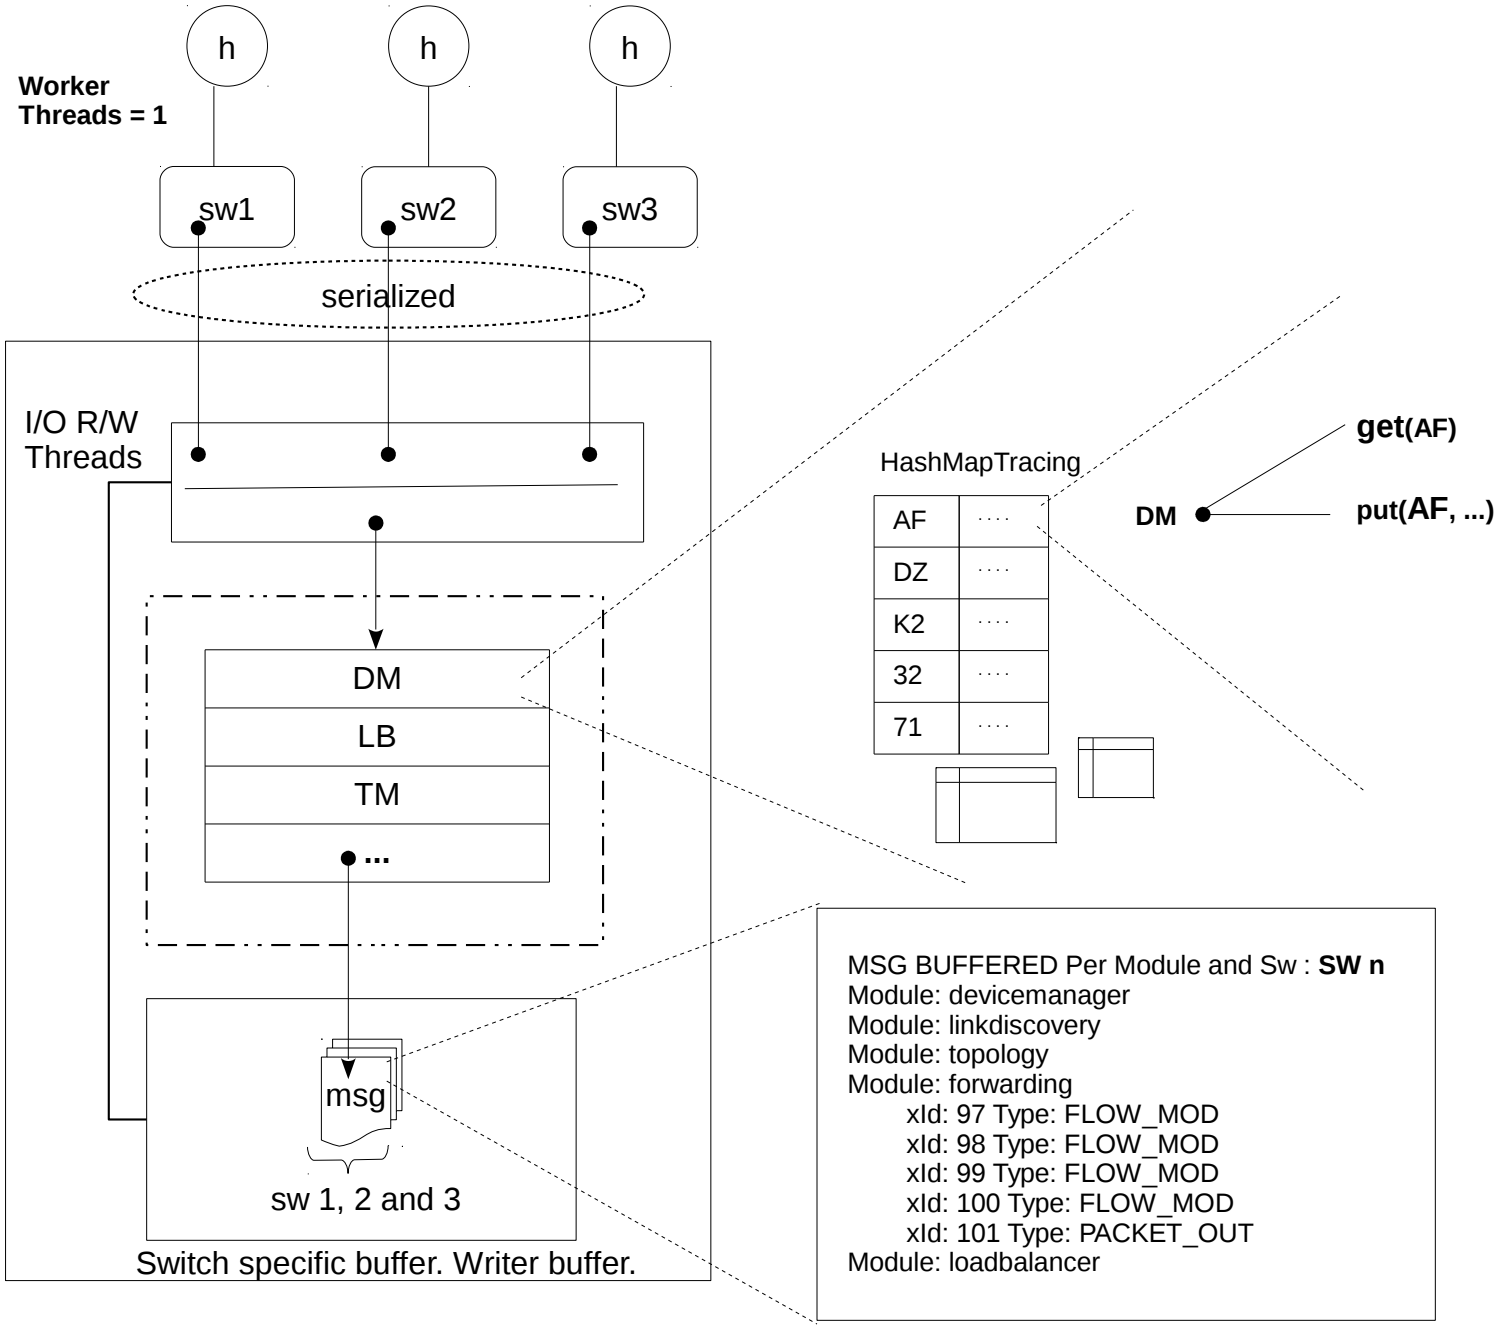

# Outline

- SDN – Context
- (Strong) Consistency Matters
- Proposed Controller Architecture
- Ravana x Floodlight Tracing
- Floodlight Pipeline Tracing
- Como Evitar Sobrecarga
- Ganho esperado

# Como Evitar Sobrecarga

Ping: host 1 → host 2

Memory Changes Per PacketIn/Switch: **00:00:00:00:00:00:00:01**, ThreadID: 62:
# DM:deviceMap
---: **put(K,V)**
---: 6
---: Device [deviceKey=6, entityClass=DefaultEntityClass,
MAC=00:00:00:00:01:10, IPv4s=[], IPv6s=[], APs=[SwitchPort [switchDPID=00:00:00:00:00:00:00:01, port=1, errorStatus=null]]]
---: LOCAL
---: 0

**DISPATCHING MESSAGE BUFFER** threadId: 62
Module: devicemanager
Module: linkdiscovery
Module: topology
Module: forwarding
xId: 157 Type: PACKET_OUT
Module: loadbalancer

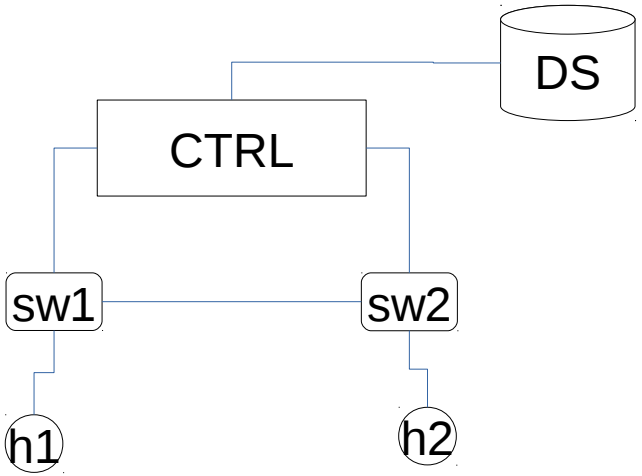


Memory Changes Per PacketIn/Switch: **00:00:00:00:00:00:00:02**, ThreadID: 62:
# DM:deviceMap
---: **get(K)**
---: 6
---: LOCAL
---: 0

**DISPATCHING MESSAGE BUFFER** threadId: 62
Module: devicemanager
Module: linkdiscovery
Module: topology
Module: forwarding
xId: 158 Type: PACKET_OUT
Module: loadbalancer

# Como Evitar Sobrecarga

Retorno do ping: host 2 → host 1

Memory Changes Per PacketIn/Switch: **00:00:00:00:00:00:00:02**, ThreadID: 62:
# DM:deviceMap
---: **put(K,V)**
---: 7
---: Device [deviceKey=7, entityClass=DefaultEntityClass, MAC=00:00:00:00:02:20, IPv4s=[], IPv6s=[], APs=[SwitchPort [switchDPID=00:00:00:00:00:00:00:02, port=1, errorStatus=null]]]
---: LOCAL
---: 0
# DM:deviceMap
---: **get(K)**
---: 6
---: LOCAL
---: 0

**DISPATCHING MESSAGE BUFFER** threadId: 62
Module: devicemanager
Module: linkdiscovery
Module: topology
Module: forwarding
    xId: 160 Type: FLOW_MOD
    xId: 161 Type: FLOW_MOD
    xId: 162 Type: PACKET_OUT
Module: loadbalancer

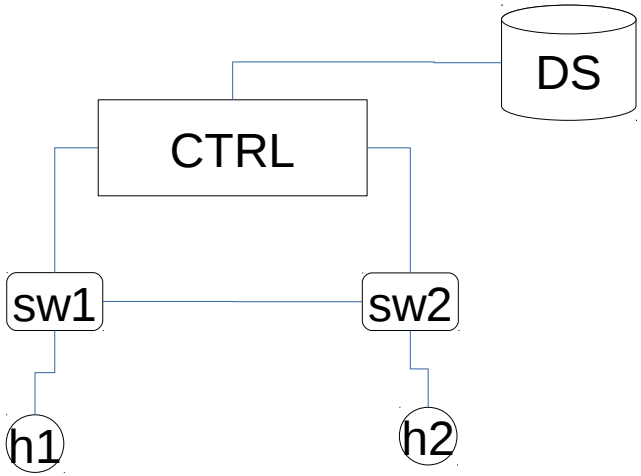

# Como Evitar Sobrecarga

## Segundo ping: host 1 → host 2

Memory Changes Per PacketIn/Switch: **00:00:00:00:00:00:00:01**, ThreadID: 62:
# DM:deviceMap
---: **get(K)**
---: 6
---: LOCAL
---: 0
# DM:deviceMap
---: **get(K)**
---: 7
---: LOCAL
---: 0

DISPATCHING MESSAGE BUFFER threadId: 62
Module: devicemanager
Module: linkdiscovery
Module: topology
Module: forwarding
        xId: 639 Type: FLOW_MOD
        xId: 640 Type: FLOW_MOD
        xId: 641 Type: PACKET_OUT
Module: loadbalancer

Memory Changes Per PacketIn/Switch: **00:00:00:00:00:00:00:02**, ThreadID: 62:
# DM:deviceMap
---: **get(K)**
---: 6
---: LOCAL
---: 0
# DM:deviceMap
---: **get(K)**
---: 7
---: LOCAL
---: 0
DISPATCHING MESSAGE BUFFER threadId: 62
Module: devicemanager
Module: linkdiscovery
Module: topology
Module: forwarding
        xId: 642 Type: PACKET_OUT
Module: loadbalancer

# Outline

- SDN – Context
- (Strong) Consistency Matters
- Proposed Controller Architecture
- Ravana x Floodlight Tracing
- Floodlight Pipeline Tracing
- Como Evitar Sobrecarga
- Ganho esperado

# Ganho esperado

Data Store throughput: 100 Kops/s

Controlador throughput: 100 Kops/s

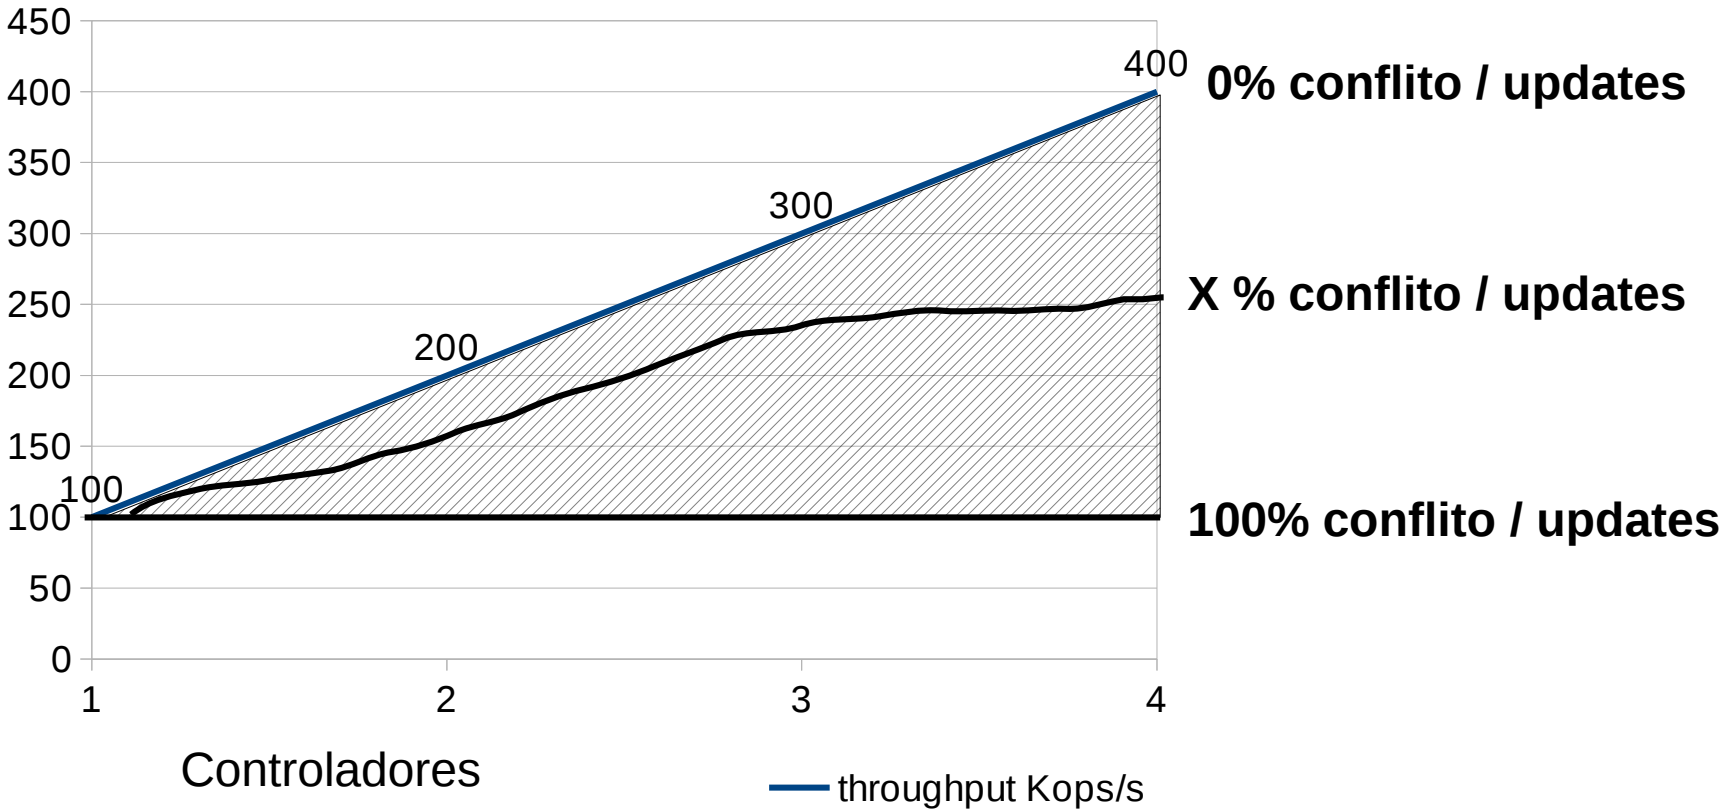


**0% conflito / updates**

**X % conflito / updates**

**100% conflito / updates**

# Obrigado.

- LaSIGE: http://www.lasige.di.fc.ul.pt/

- Tulio Ribeiro: http://www.lasige.di.fc.ul.pt/user/422

**Questões / Dúvidas ?**

**FIM**