

CellSDN: Taking control of cellular core networks



Laurent Vanbever

Princeton University

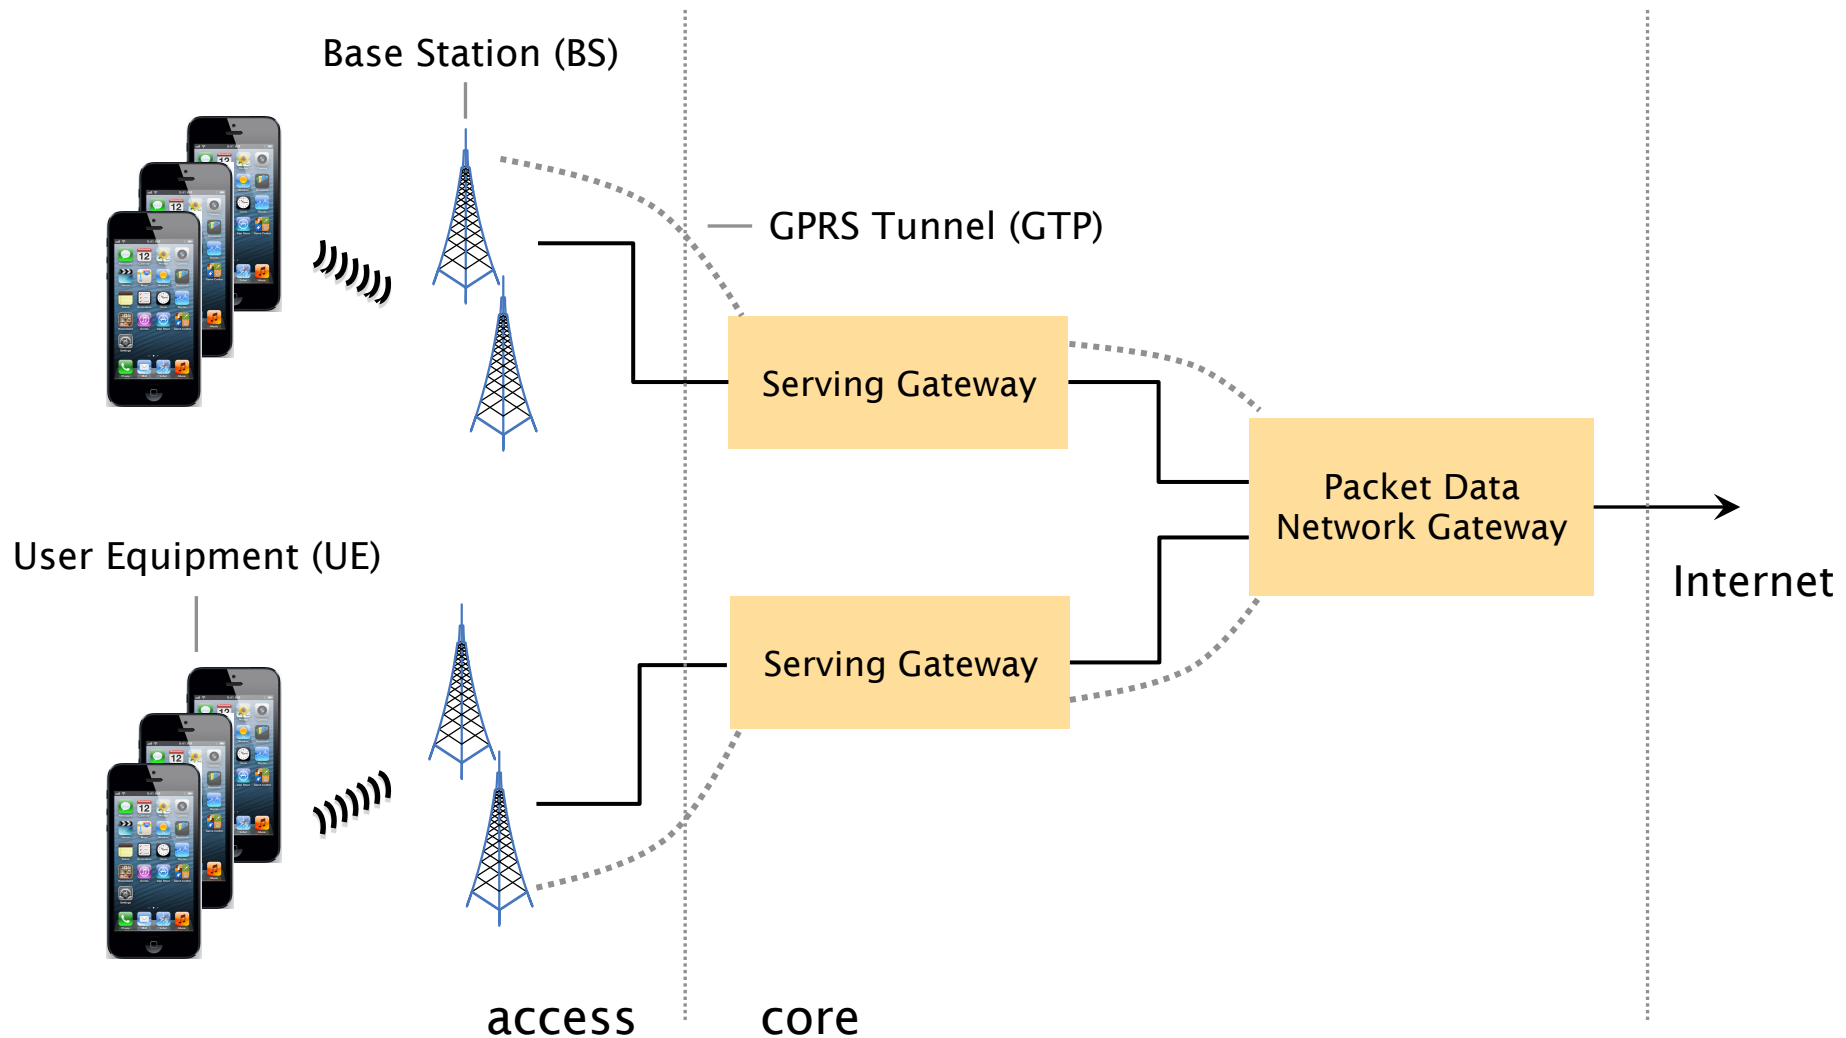
IBM Watson Research Center

Tue 9 April 2013

Joint work with Xin Jin, Li Erran Li, and Jennifer Rexford

Cellular core networks
are not **flexible**

Cellular core networks are not **flexible**



Cellular core networks are not **flexible**

Most functionalities are implemented at the PGW:

- Stateful firewall, DPI, lawfull intercept
- Application optimization
- Paging
- Monitoring and Billing

Packet Data
Network Gateway

Cellular core networks are not **flexible**

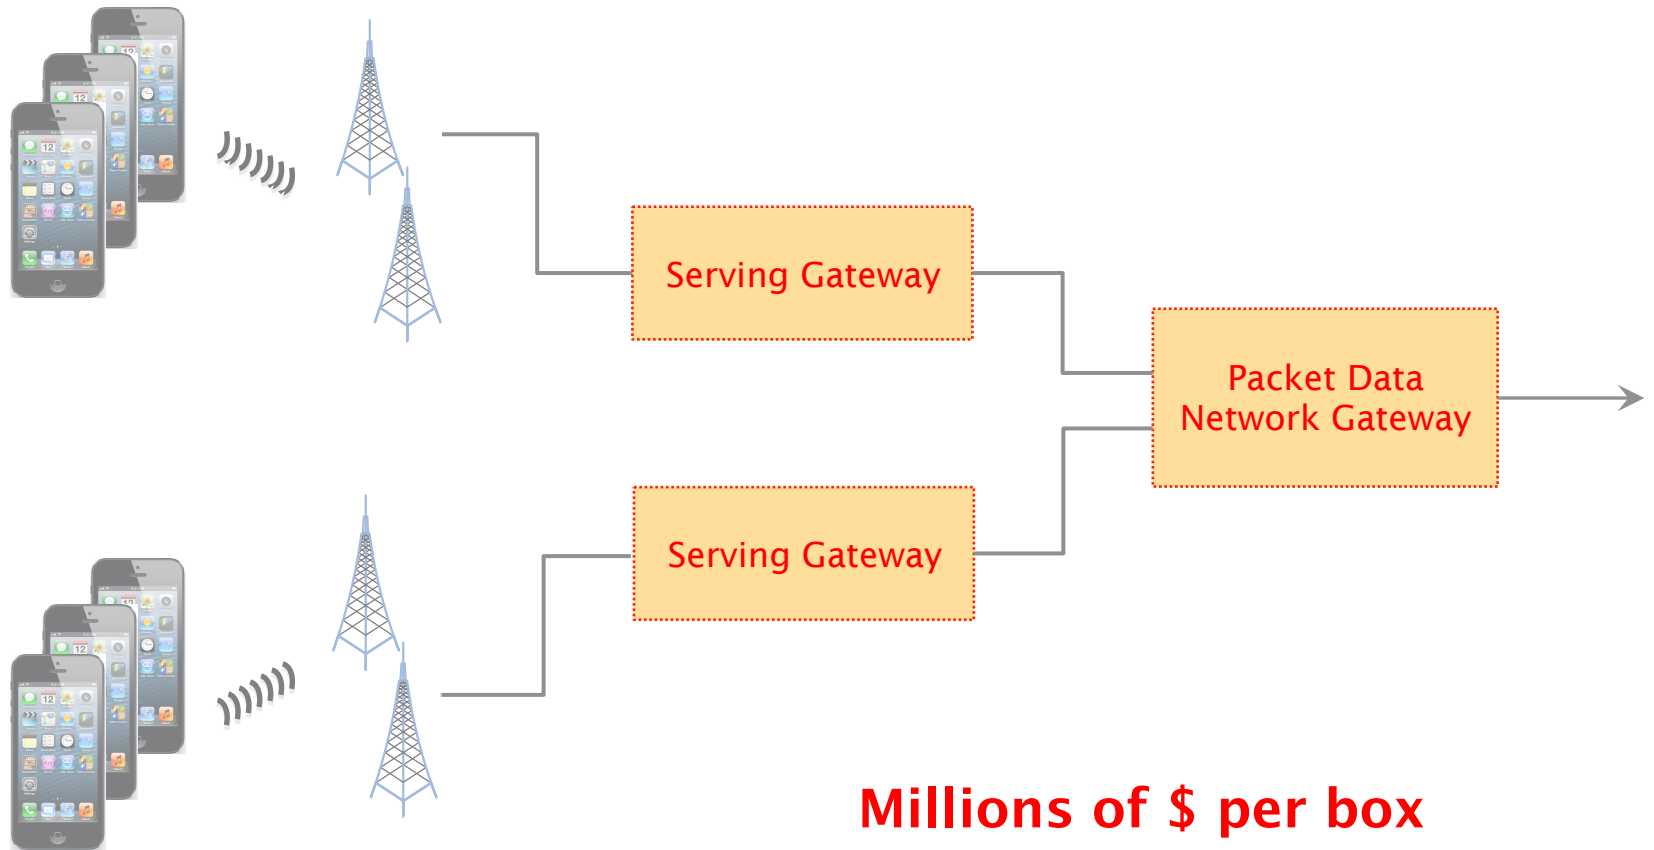
Most functionalities are implemented at the PGW:

- Stateful firewall, DPI, lawful intercept
- Application optimization
- Paging
- Monitoring and Billing

Packet Data
Network Gateway

- **Hard to mix-and-match**
- **Slow innovation**

Cellular core networks are not flexible **and costly!**



CellSDN enables
flexible and cost-effective networks

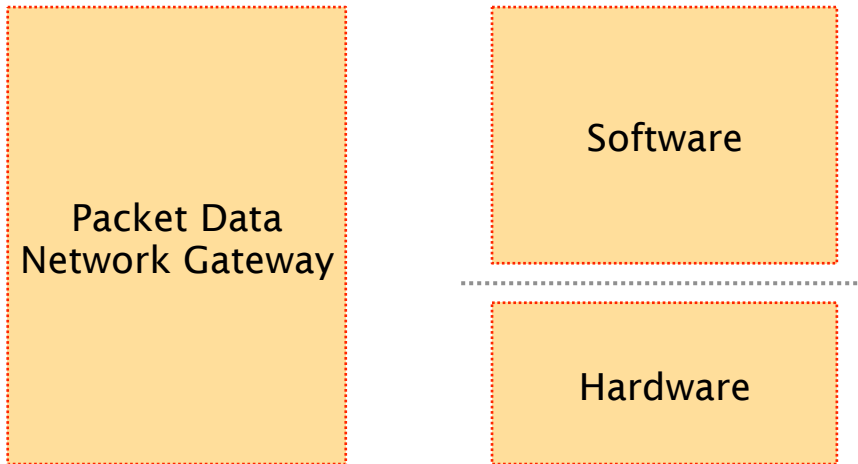
CellSDN enables flexible and cost-effective networks

Key idea: separate Software from Hardware

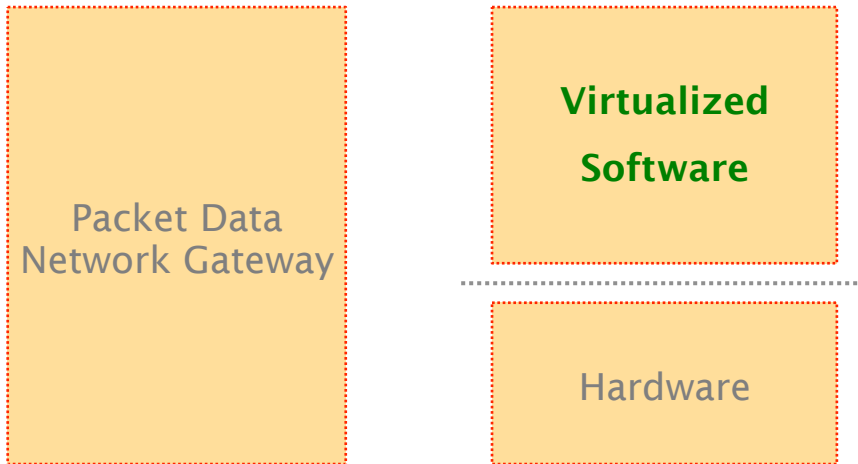


Packet Data
Network Gateway

CellSDN enables flexible and cost-effective networks




CellSDN enables flexible and cost-effective networks



CellSDN enables **flexible** and cost-effective networks



CellSDN enables flexible and cost-effective networks



Virtualized
Software

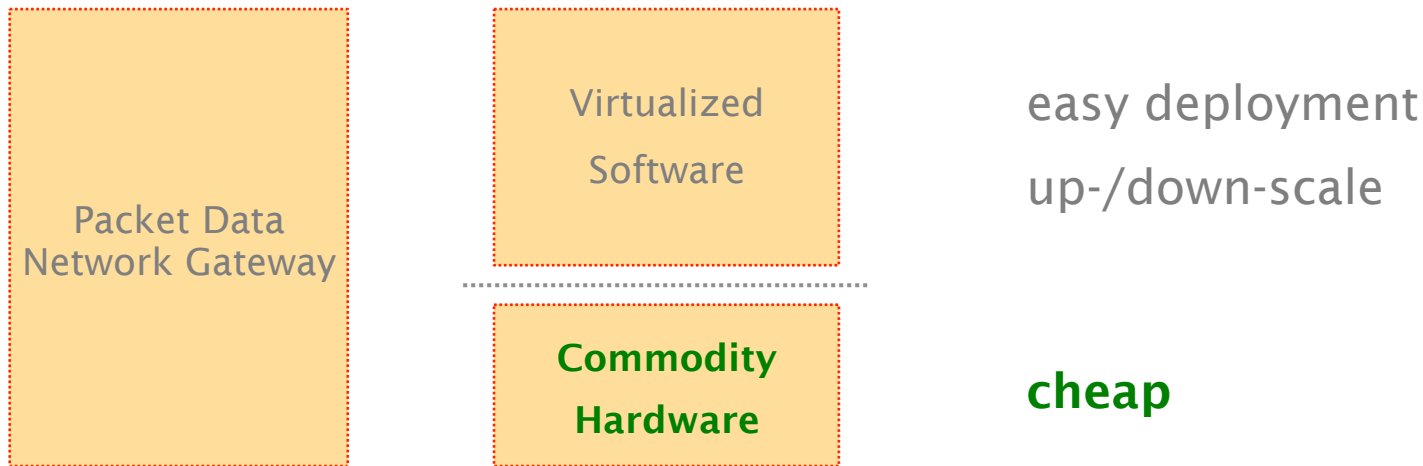
Virtualized Software

**Commodity
Hardware**

Commodity Hardware

easy deployment
up-/down-scale

CellSDN enables flexible and **cost-effective** networks



To enable CellSDN flexibility, the challenge is **scalability**

Cellular policies are fine-grained, based on

- customer-attributes
smartphone model, OS, billing plan, ...
- applications
video, web, voice, ...
- a combination of both
video from iPhone 5 gold subscribers, ...

To enable CellSDN flexibility, the challenge is **scalability**

Cellular policies are fine-grained, based on

- customer-attributes
smartphone model, OS, billing plan, ...
- applications
video, web, voice, ...
- a combination of both
video from iPhone 5 gold subscribers, ...

can lead to *million* of different flows

To enable CellSDN flexibility,
the challenge is **scalability**

Challenges

Data-plane

Forwarding table size

Control-plane

Overhead

Slow reactivity

To enable CellSDN flexibility,
the challenge is **scalability**

Challenges

Solutions

Data-plane

Forwarding table size

Hierarchical tagging

Control-plane

Overhead

Slow reactivity

To enable CellSDN flexibility,
the challenge is **scalability**

Challenges

Solutions

Data-plane

Forwarding table size

Hierarchical tagging

Control-plane

Overhead

Slow reactivity

Hierarchical controller

CellSDN: Taking control of cellular core networks



Architecture

software-defined network

Scaling the data-plane

multi-dimensional tagging

Scaling the control-plane

tasks delegation

CellSDN: Taking control of cellular core networks



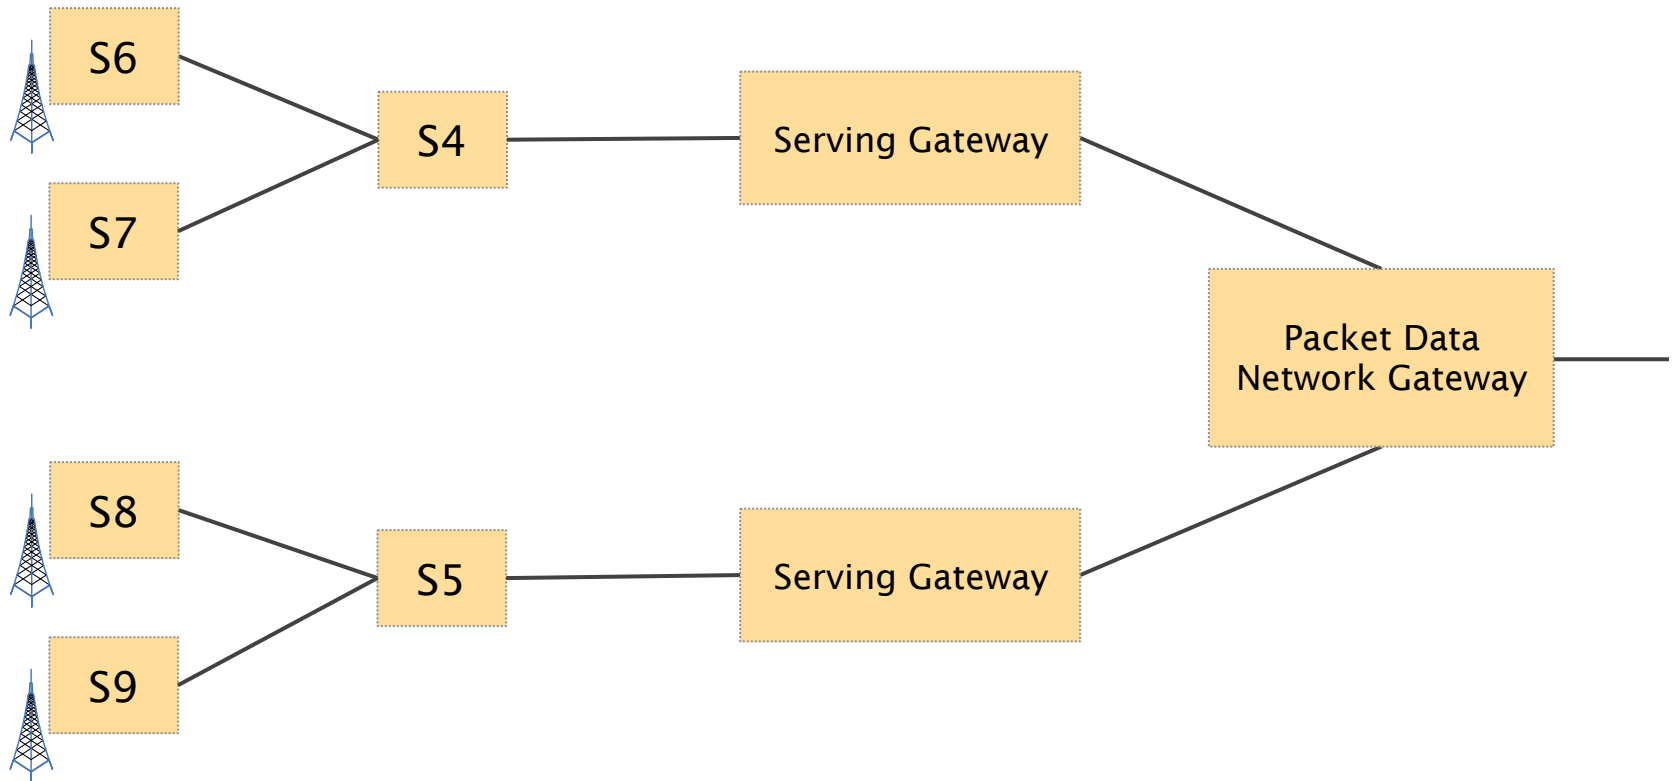
1 Architecture

software-defined network

Scaling the data-plane
multi-dimensional tagging

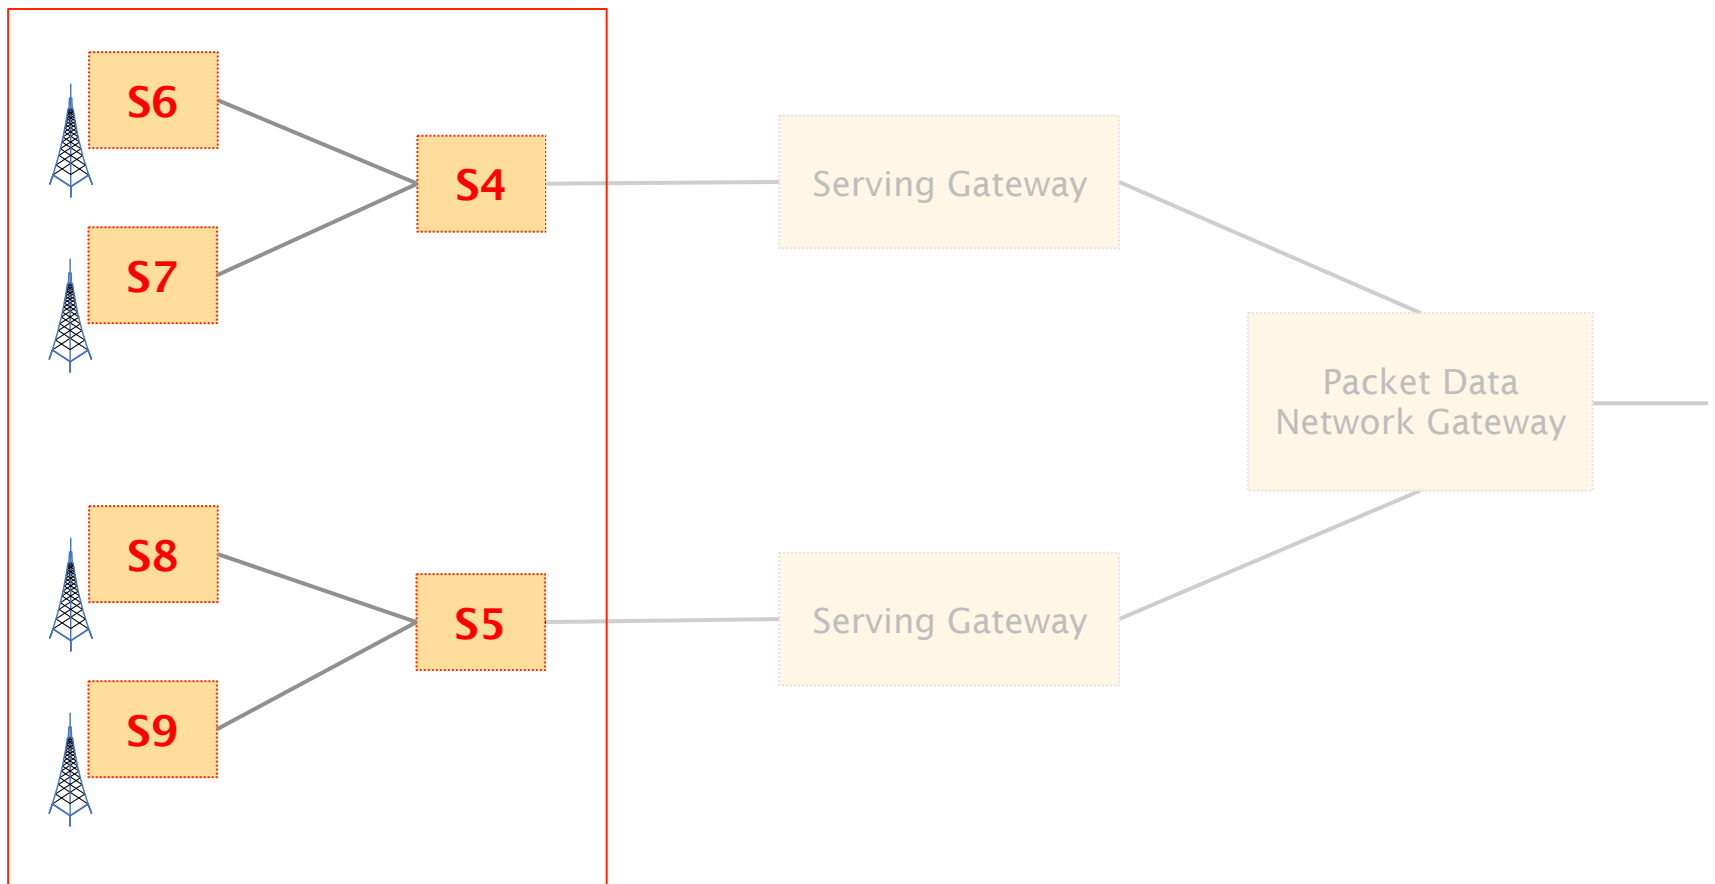
Scaling the control-plane
tasks delegation

CellSDN relies on commodity, SDN-enabled switches

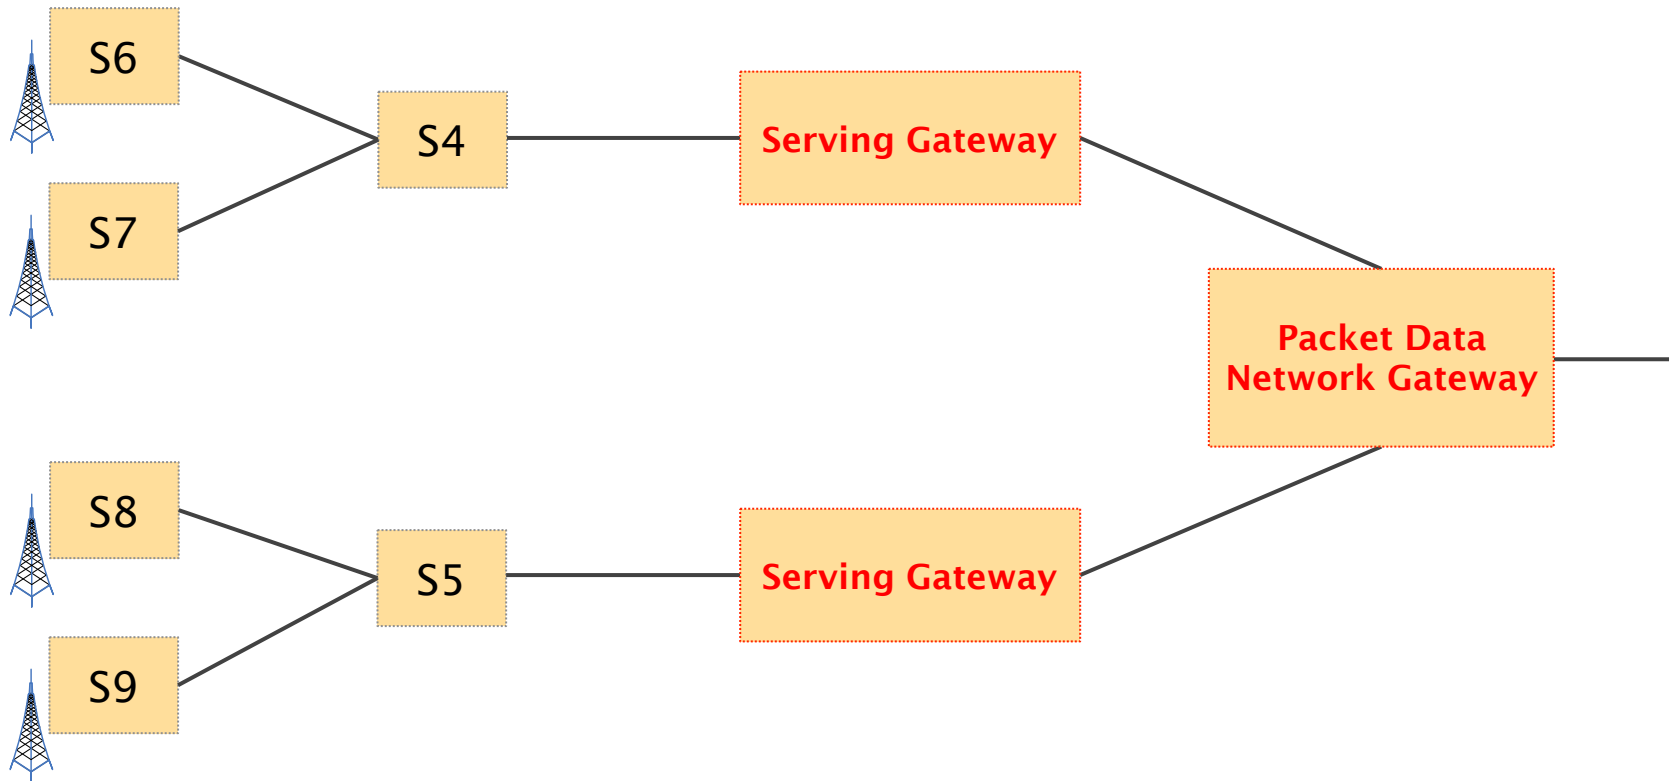


CellSDN relies on commodity, SDN-enabled switches


OpenFlow switches




CellSDN also relies on commodity SDN switches *for specialized appliances*




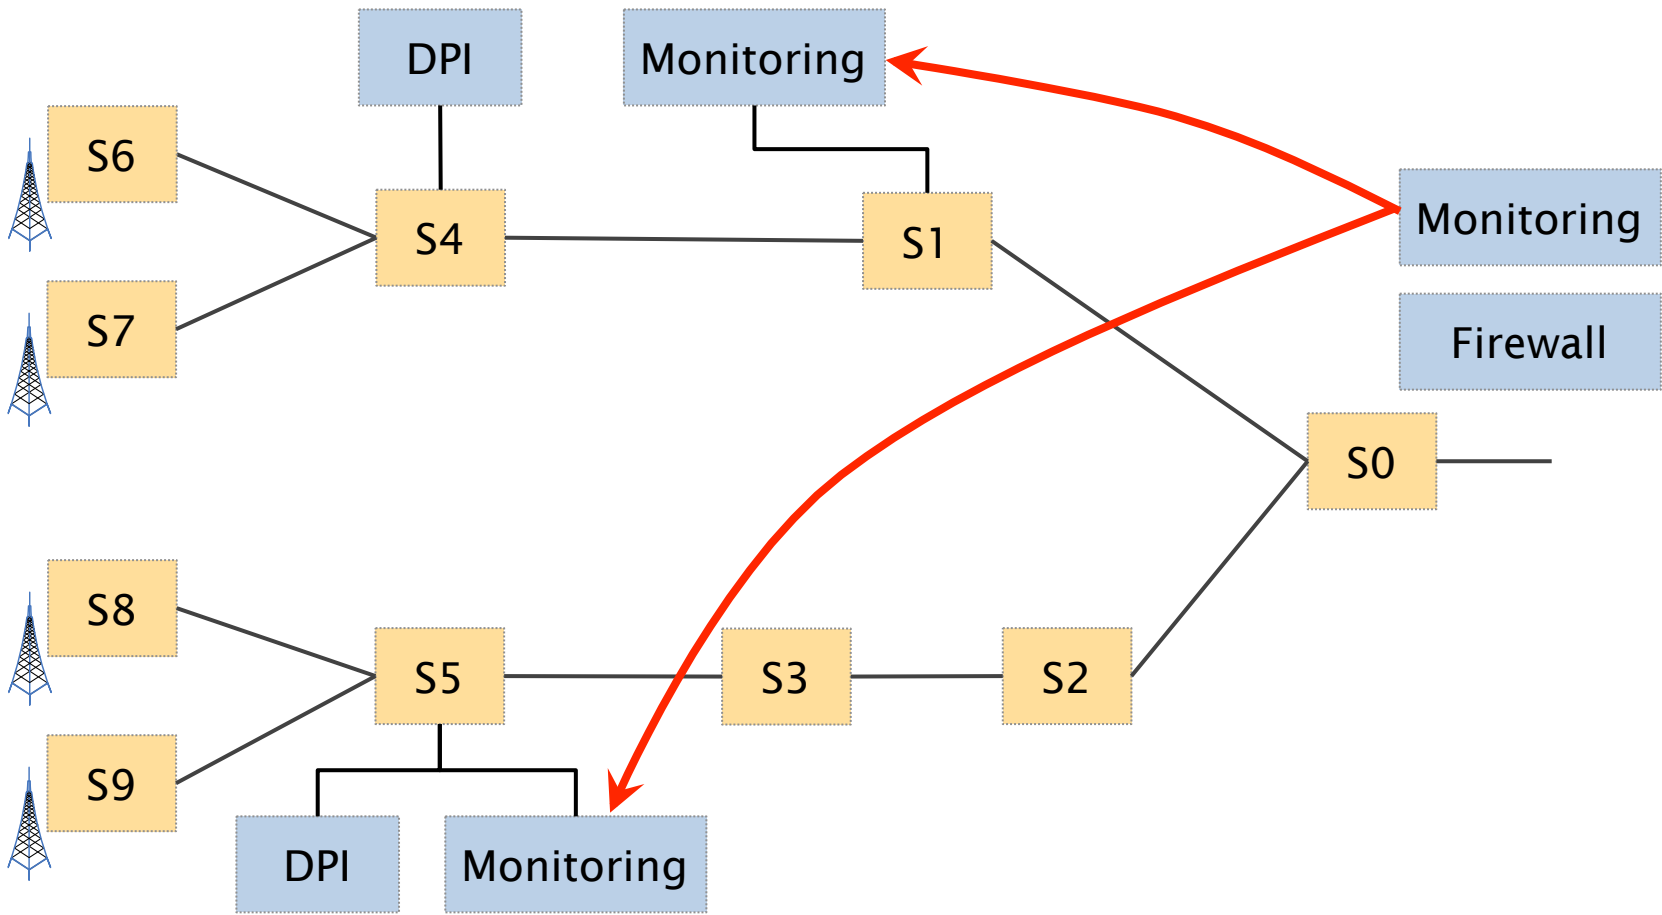
CellSDN also relies on commodity SDN switches *for specialized appliances*

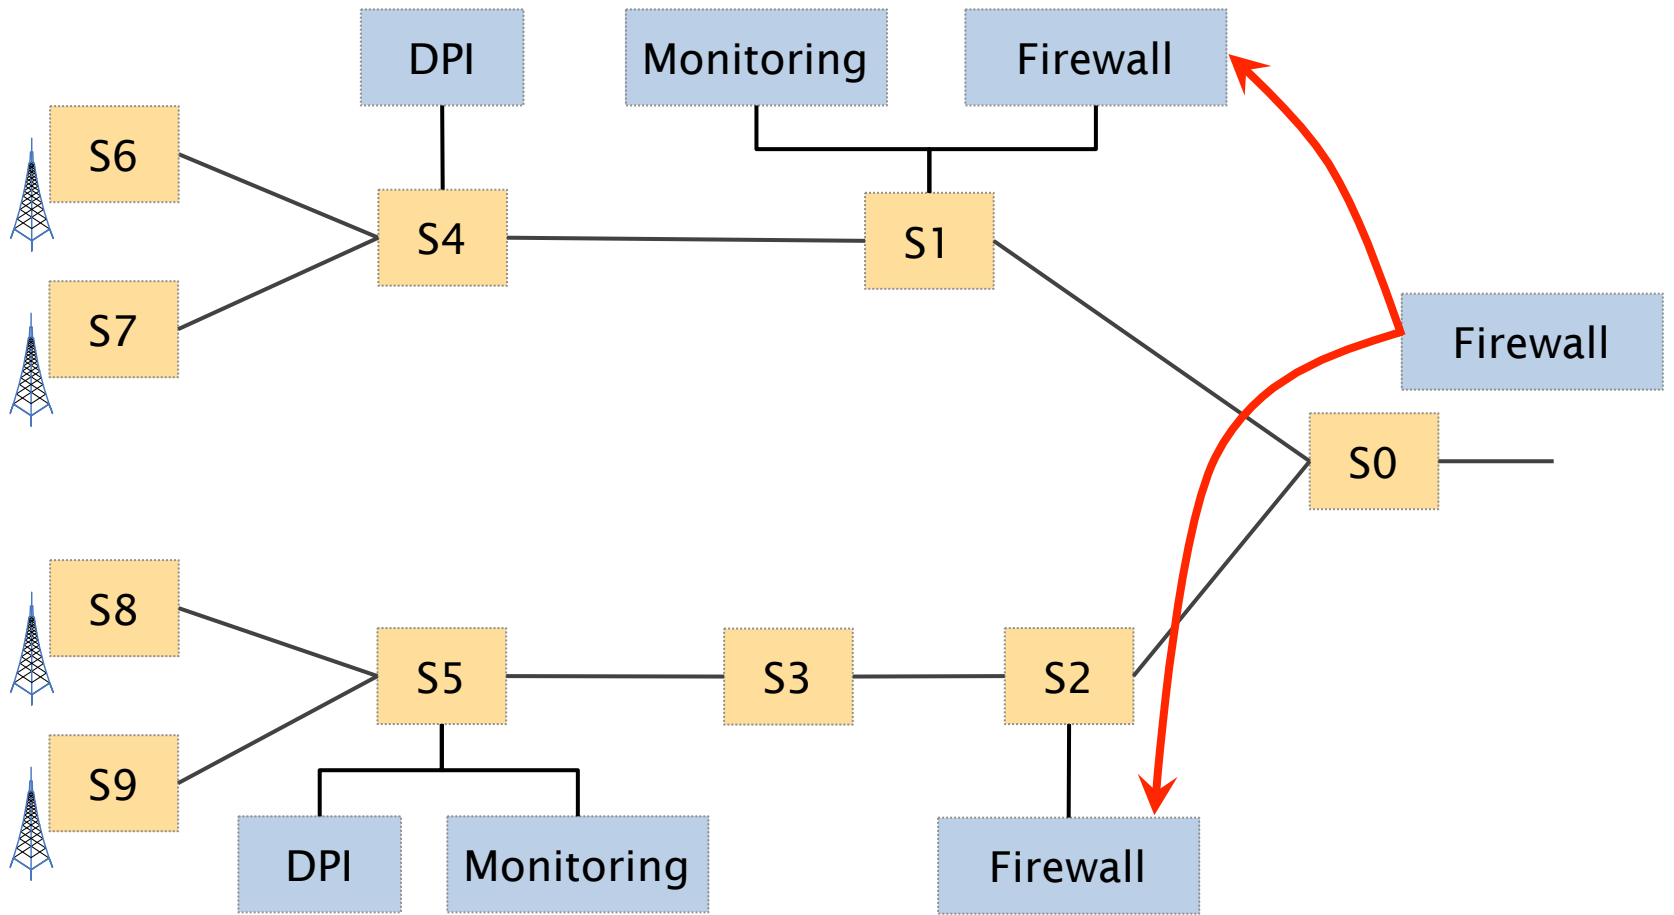


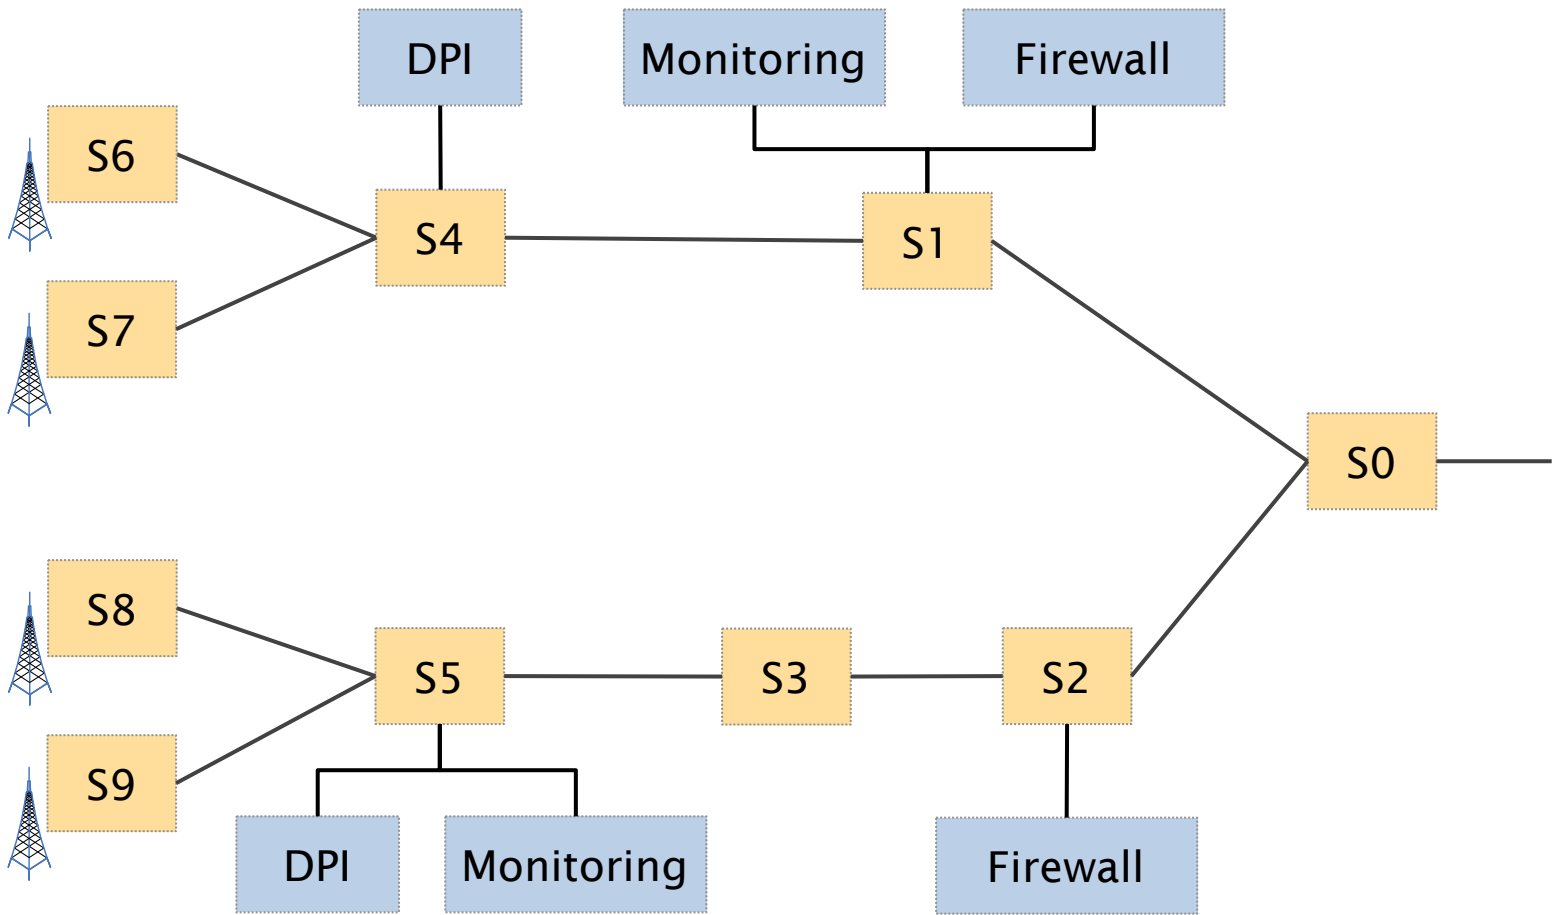
CellSDN distributes network processing











CellSDN route traffic based on high-level service policies

A service policy is composed of a

predicate

A boolean expression, e.g. on subscriber attributes, applications attributes, cell properties ...

action

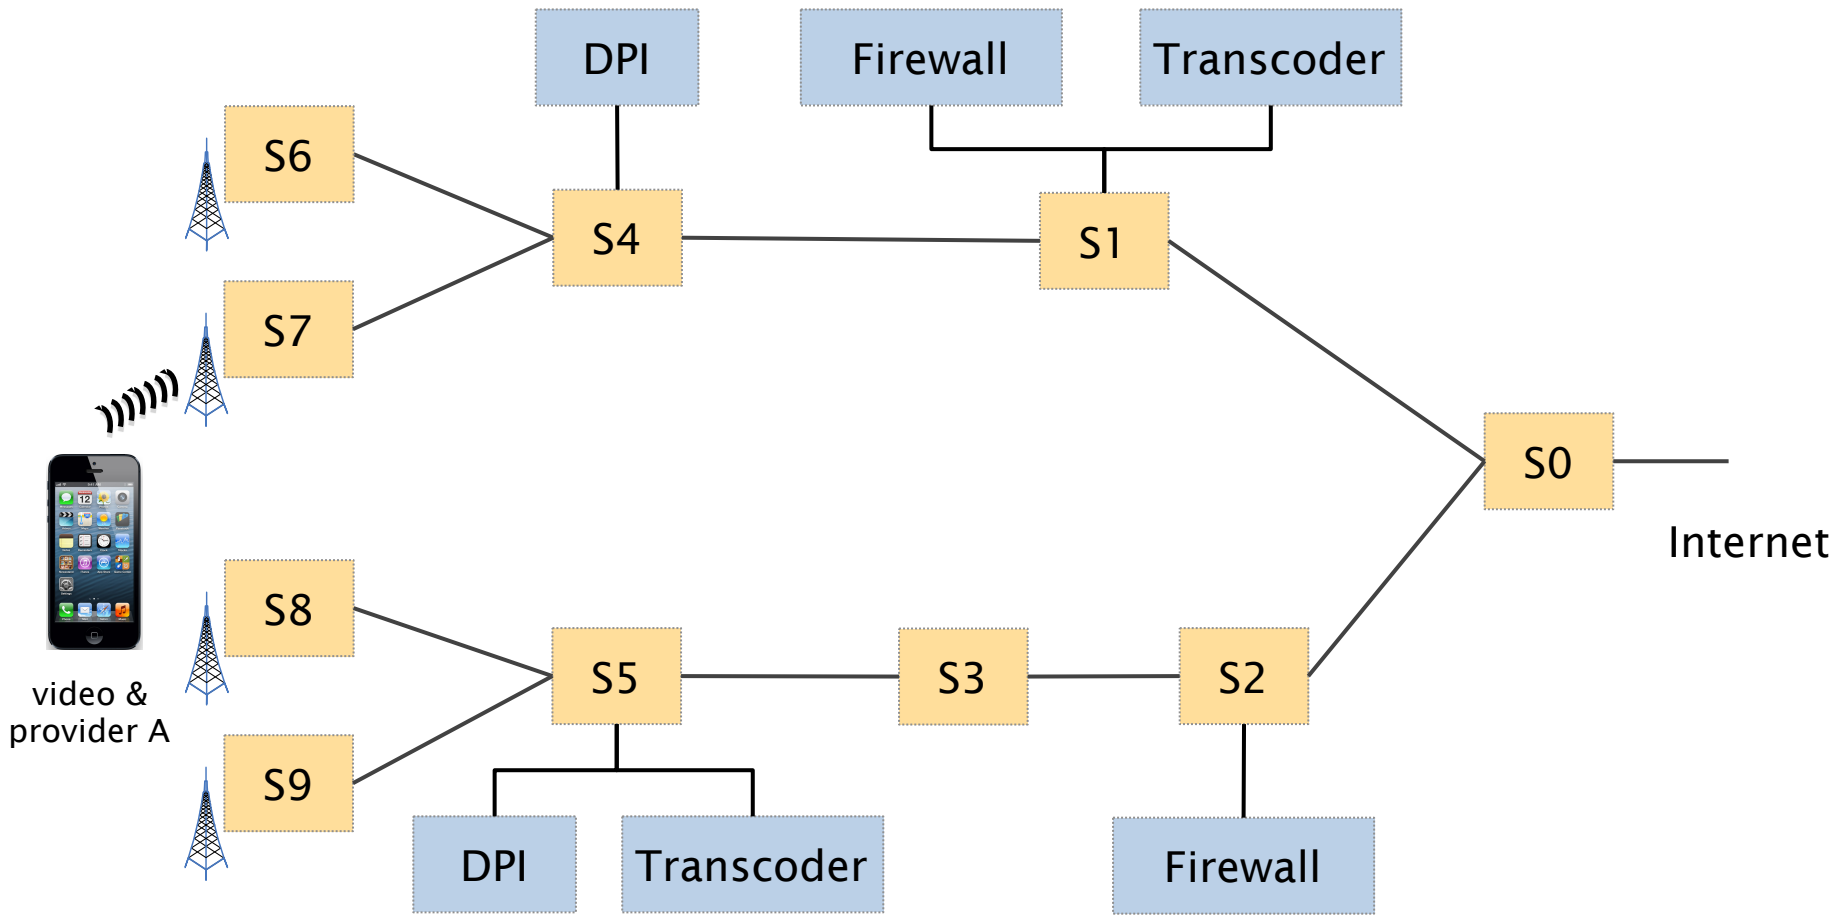
A list of middleboxes, access-control specifications (i.e., allow/deny) ...

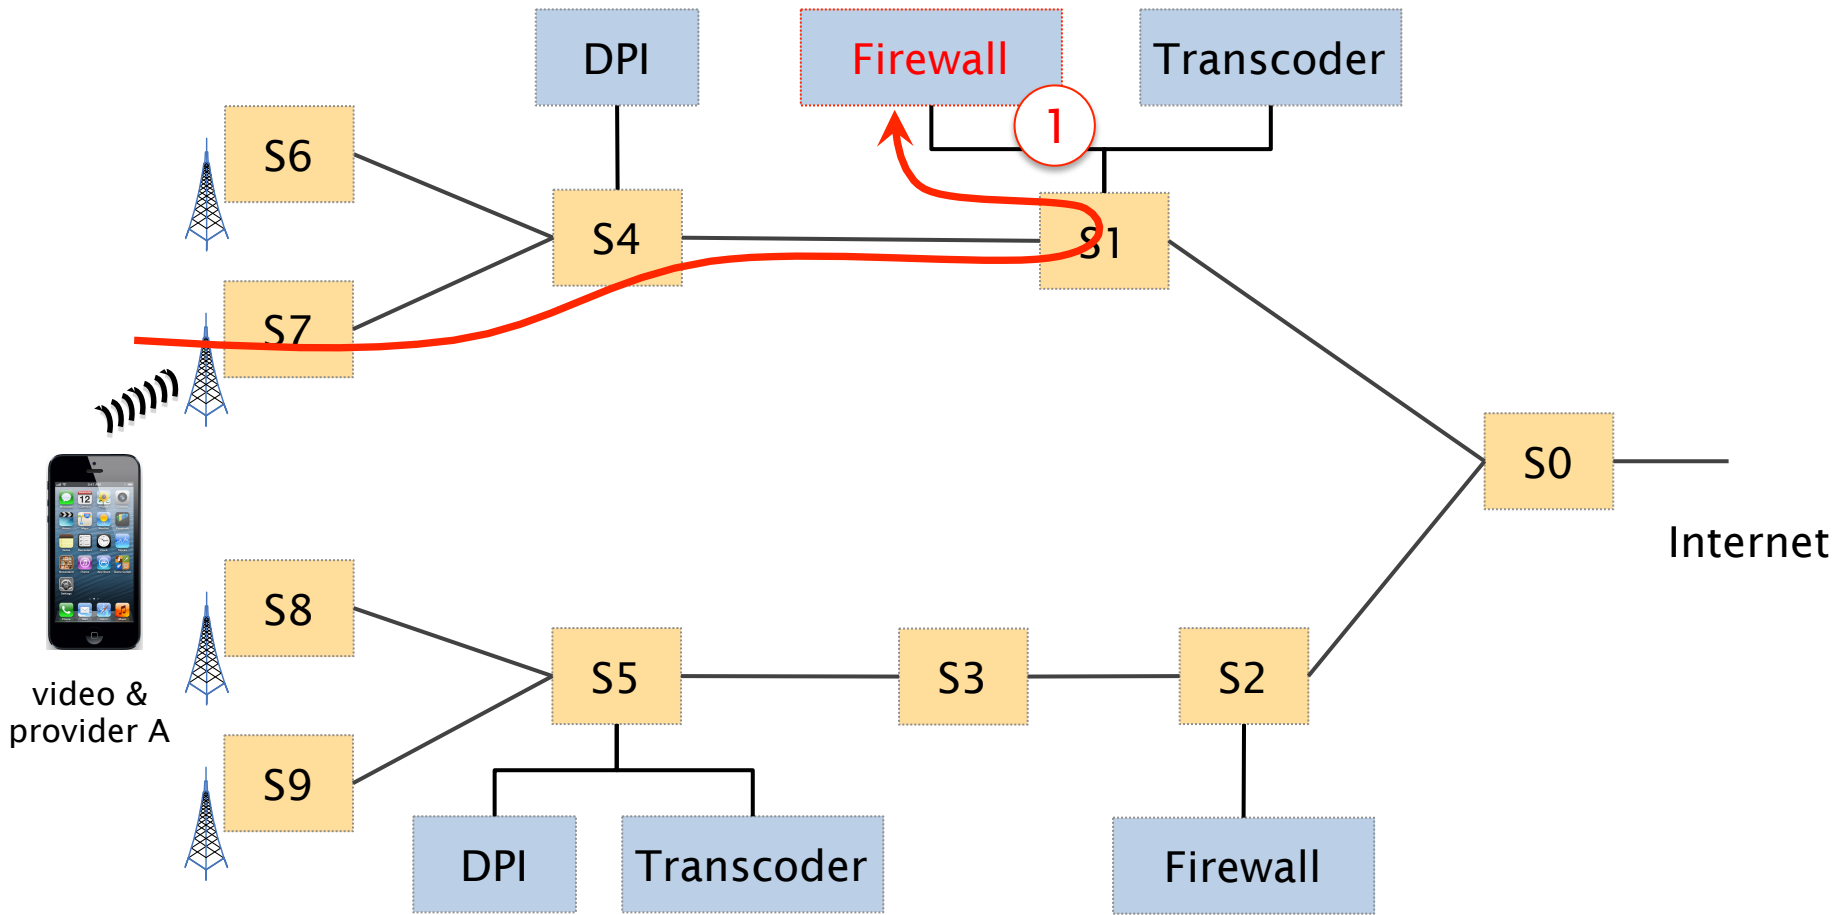
priority

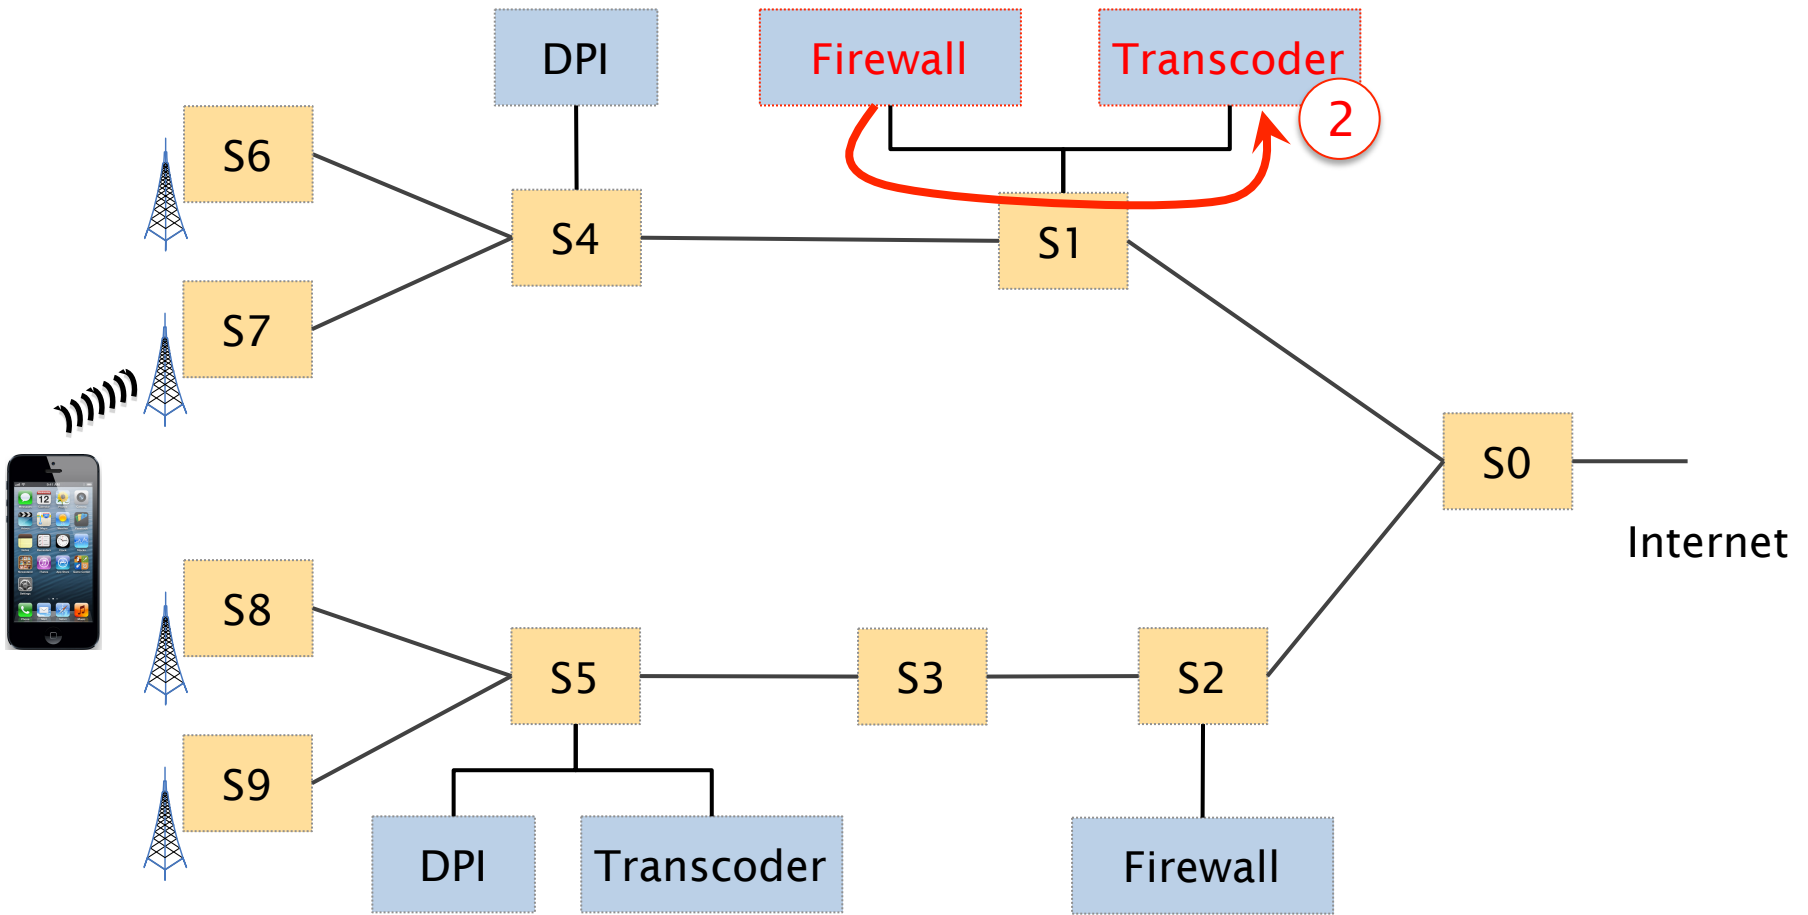
An integer. To disambiguate overlapping policies

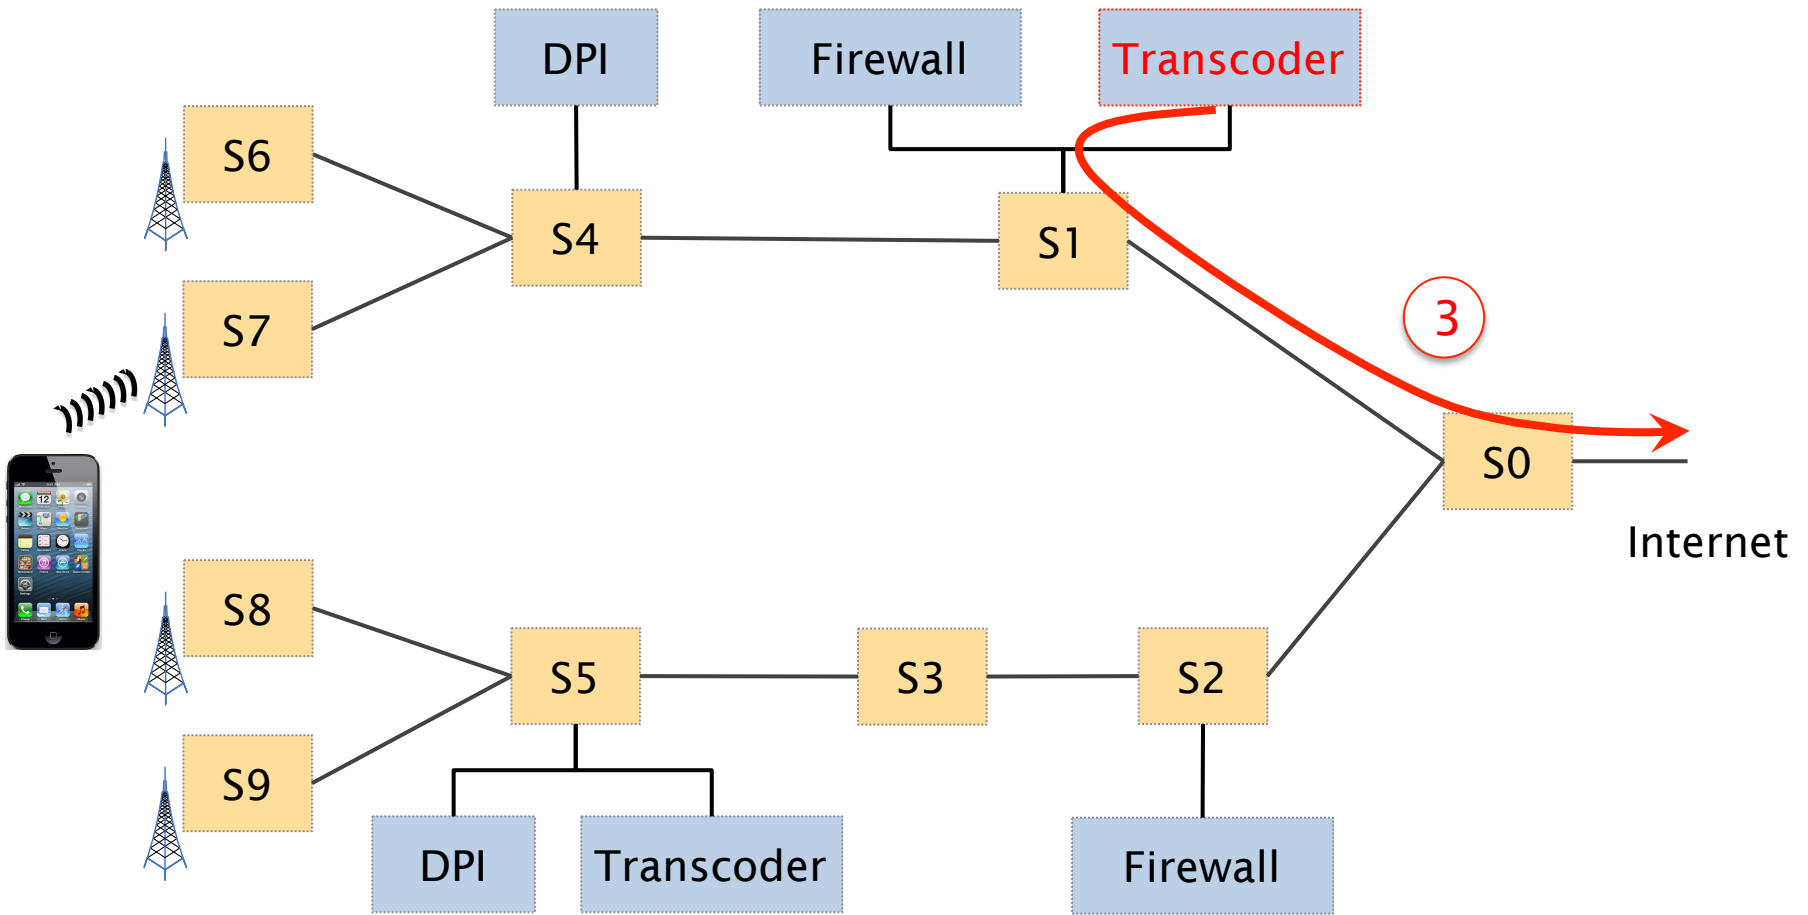
Example of service policy

priority	predicate	action
1.	provider is <i>B</i>	Firewall
2.	provider is not A	Drop
3.	traffic is video and plan is Silver	Firewall → Transcoder
4.	*	Firewall

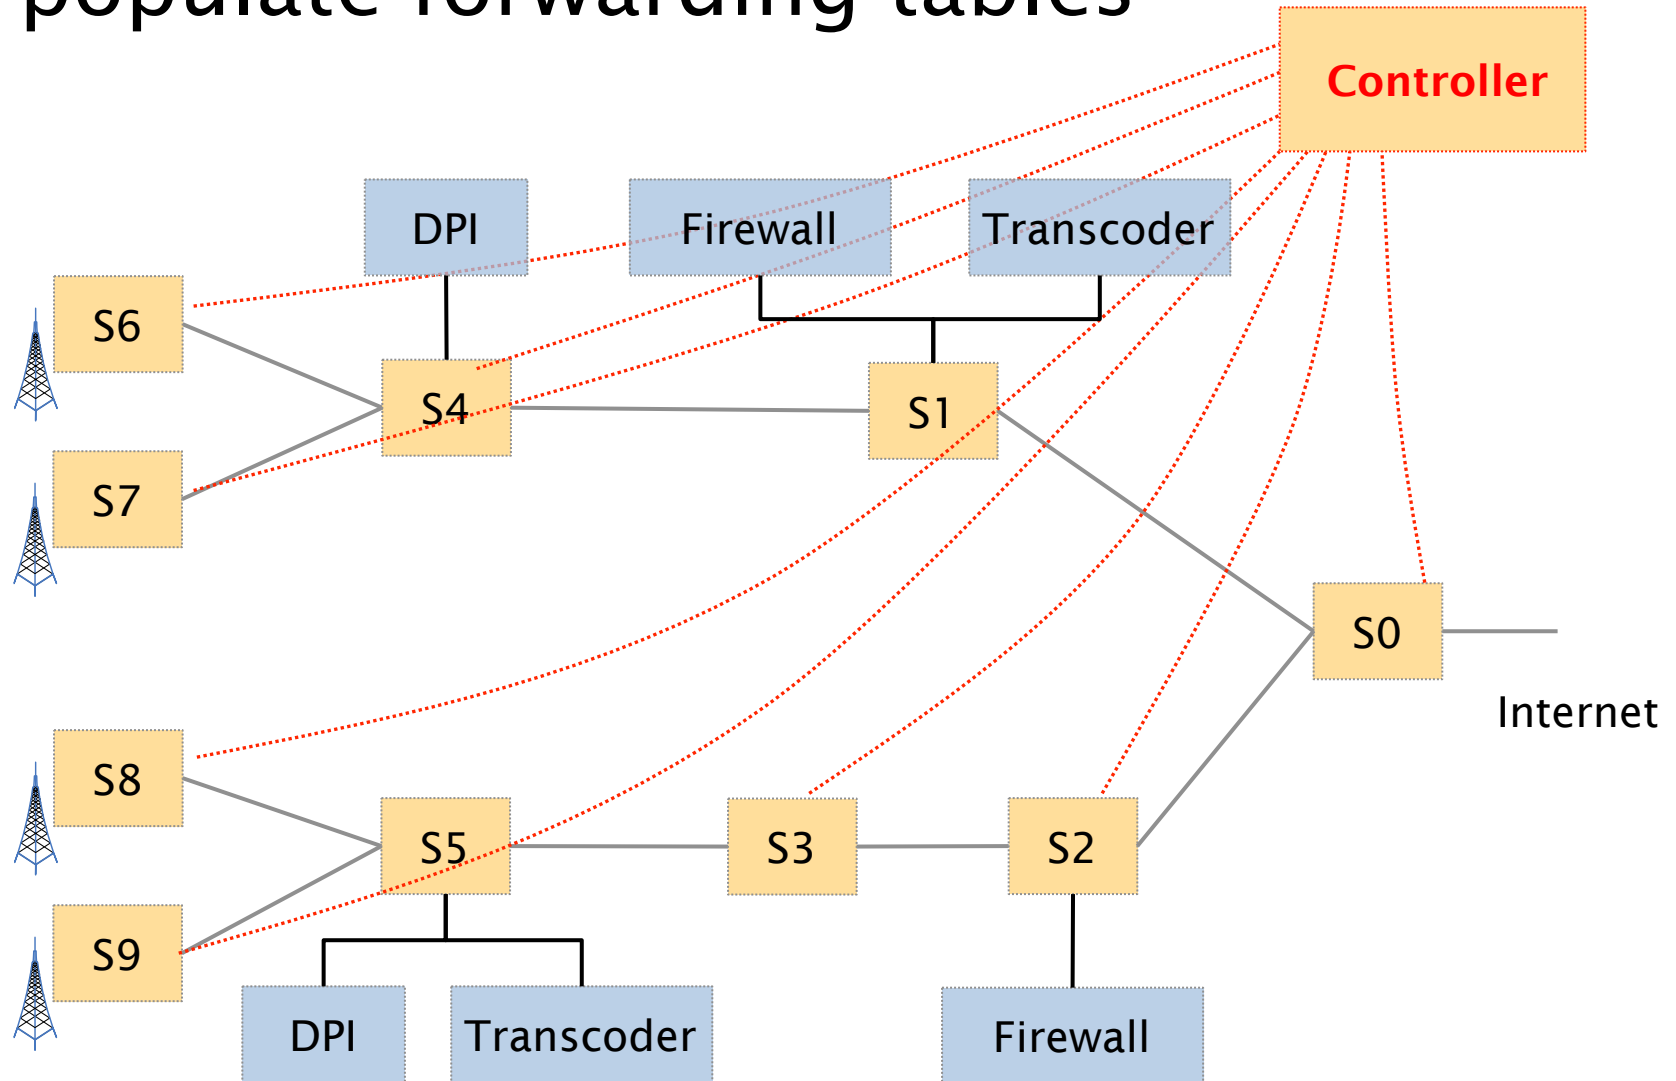








CellSDN uses a centralized controller to populate forwarding tables



CellSDN: Taking control of cellular core networks



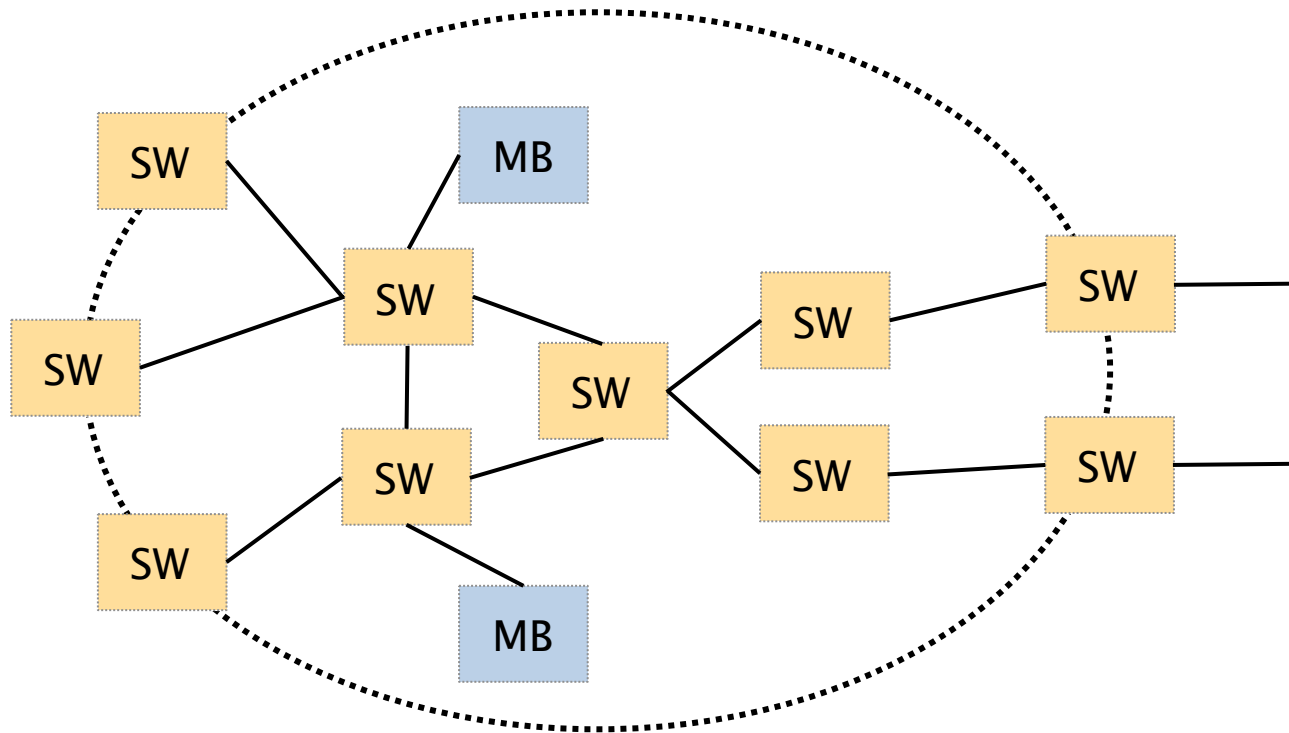
Architecture

software-defined network


2 **Scaling the data-plane**
multi-dimensional tagging

Scaling the control-plane
tasks delegation

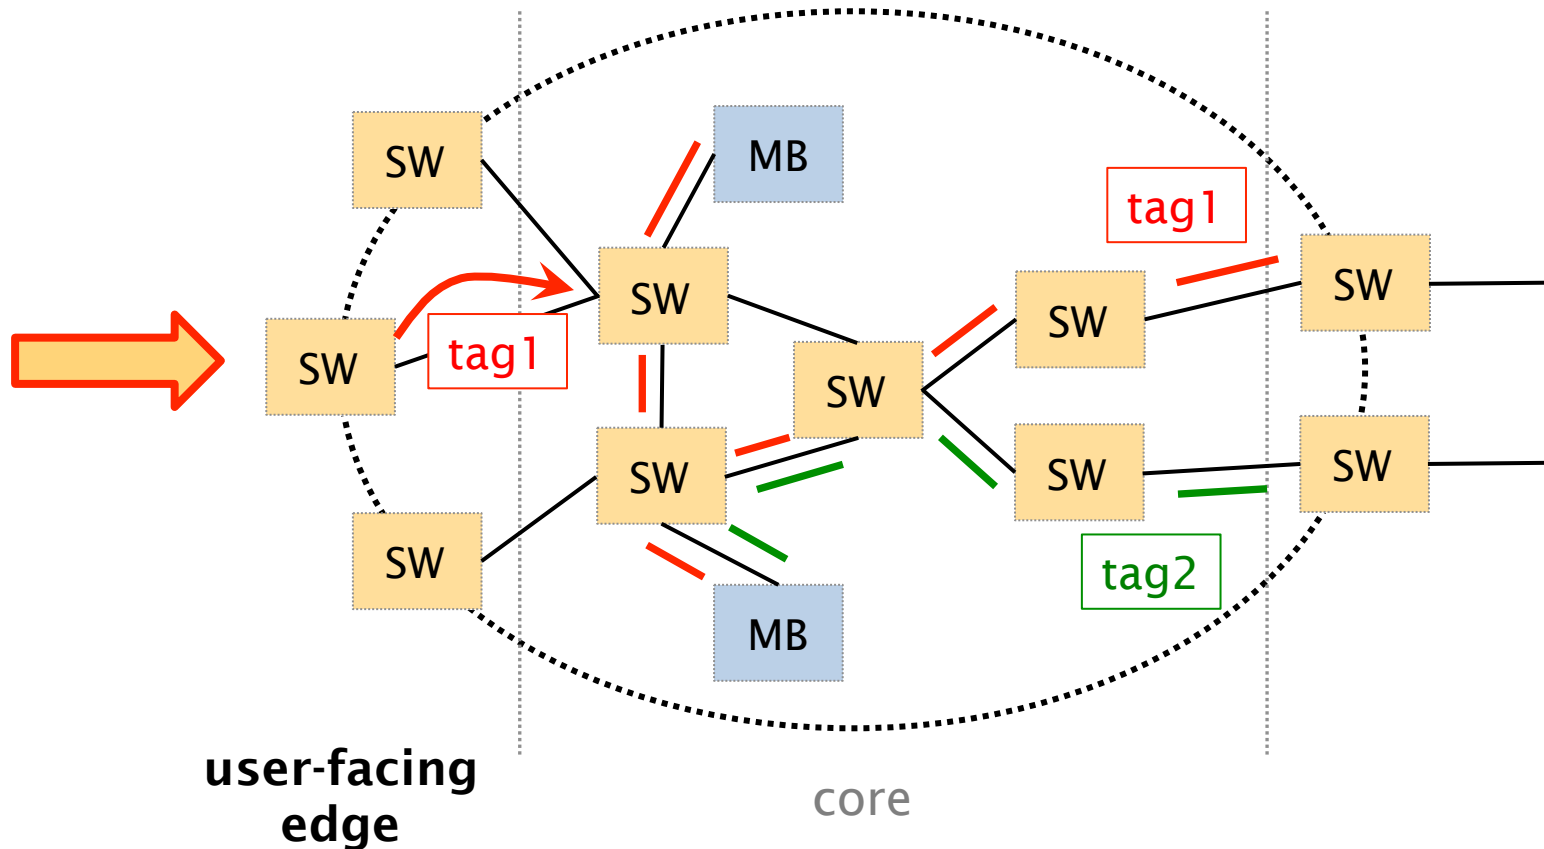
CellSDN relies on tagging to implement forwarding policies



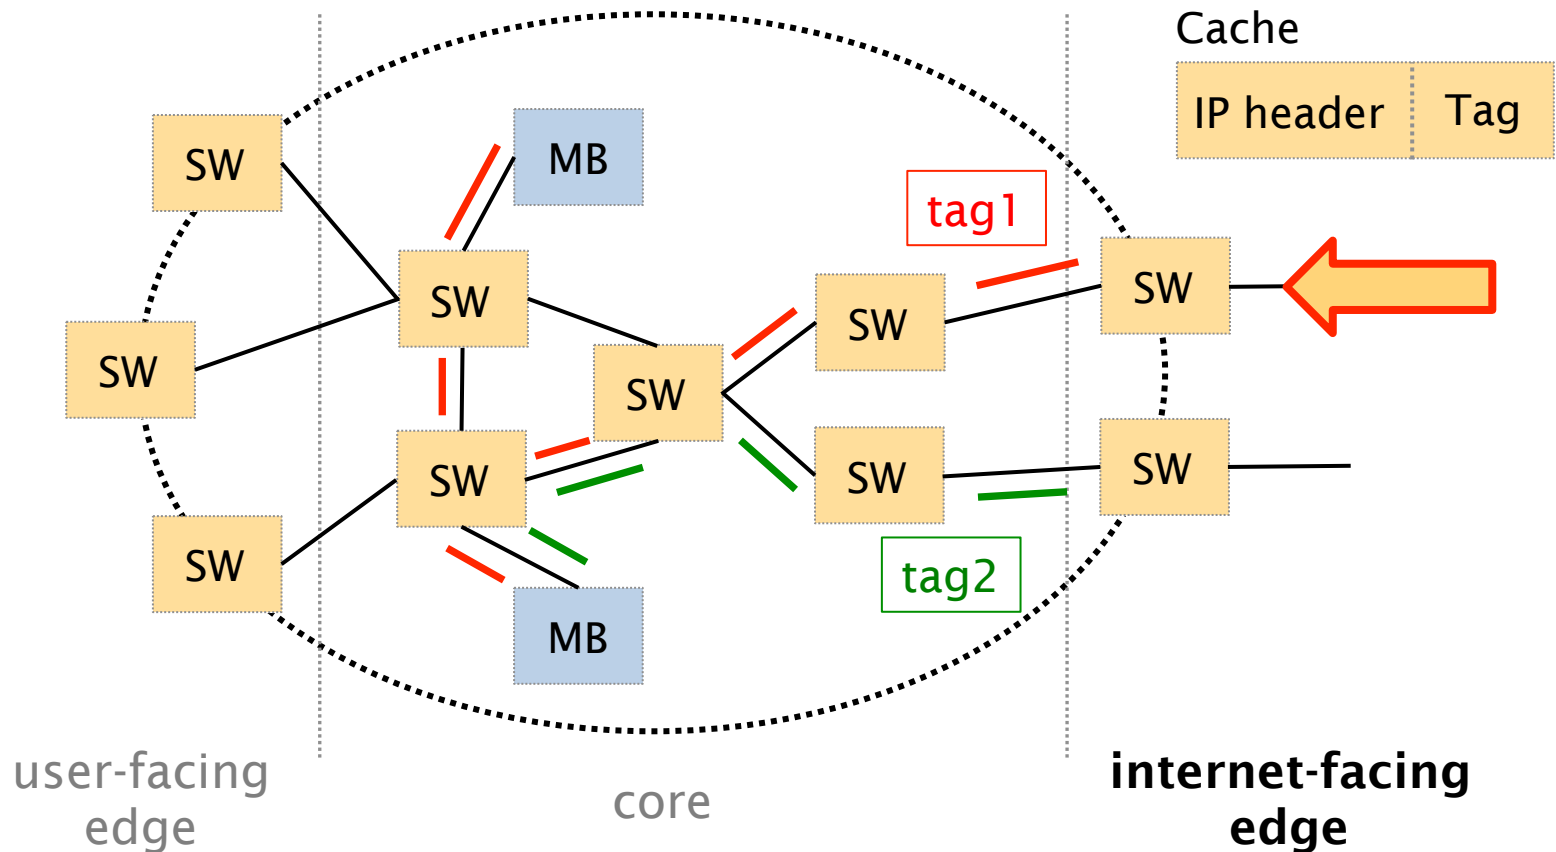
CellSDN installs stable policy path in the core



CellSDN classifies and tags traffic at the user-facing edge



CellSDN caches tags at the Internet-facing edge for the return traffic



If each path gets one tag,
millions of tag are needed

$$\begin{array}{rcccl} 1000 & * & 1000 & = & \mathbf{1 \text{ million}} \\ \# \text{ base station} & & \frac{\# \text{ policy paths}}{\text{base station}} & & \mathbf{\# \text{ paths}} \end{array}$$

If each path gets one tag, millions of tag are needed

$$\begin{array}{ccccccc} 1000 & * & 1000 & = & \mathbf{1 \text{ million}} \\ \# \text{ base station} & & \frac{\# \text{ policy paths}}{\text{base station}} & & \# \text{ paths} \end{array}$$

At best, switches support a few tens of thousands entries

[Stephen, Conext12]

CellSDN relies on multi-dimensional tag

CellSDN tags are composed of three parts:

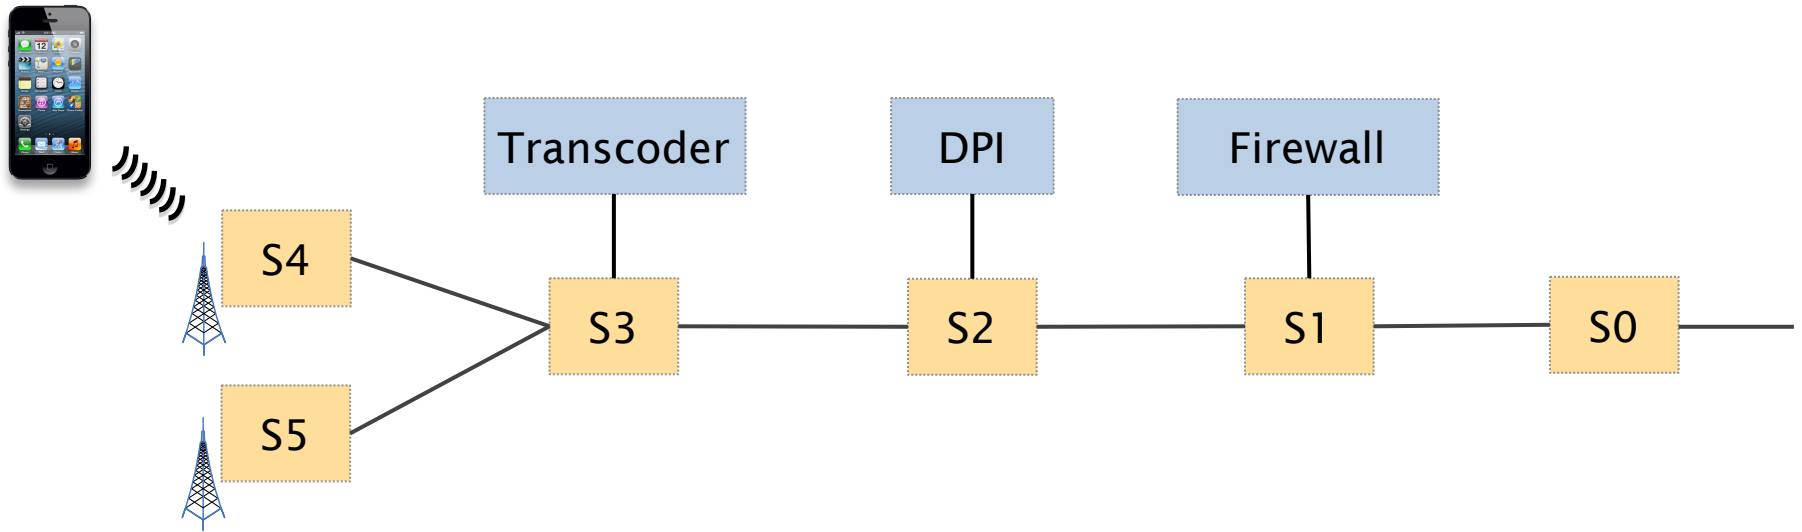
- Policy
- Location
- User Equipment Identifier

switches can selectively match on any part

CellSDN tags are composed of three parts:


- **Policy**
- Location
- User Equipment Identifier

Video Traffic + Gold Membership = Transcoder || Firewall



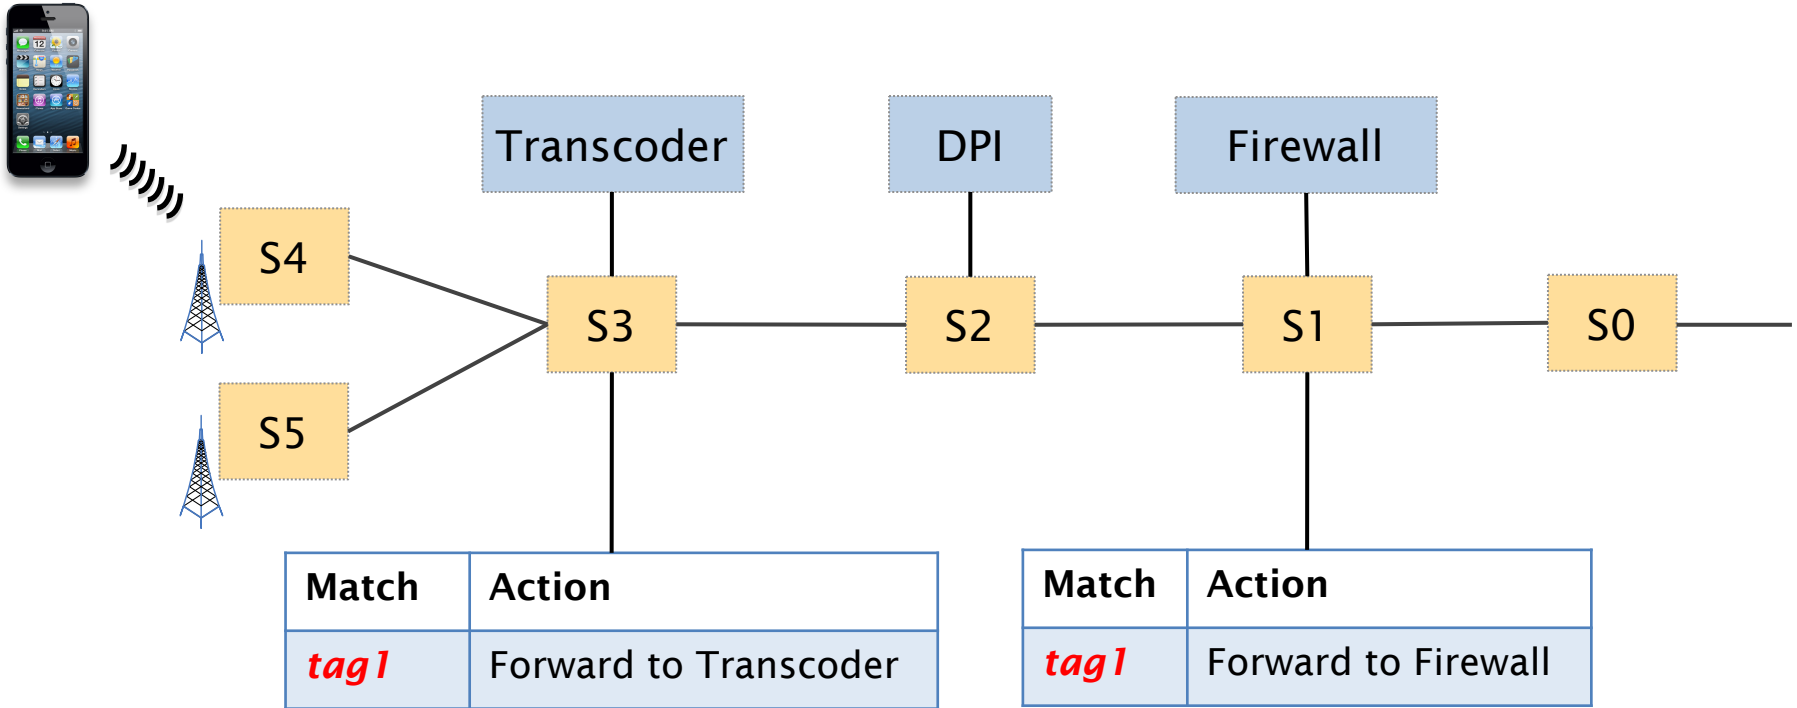
Video Traffic + Gold Membership = **Transcoder || Firewall**

tag1

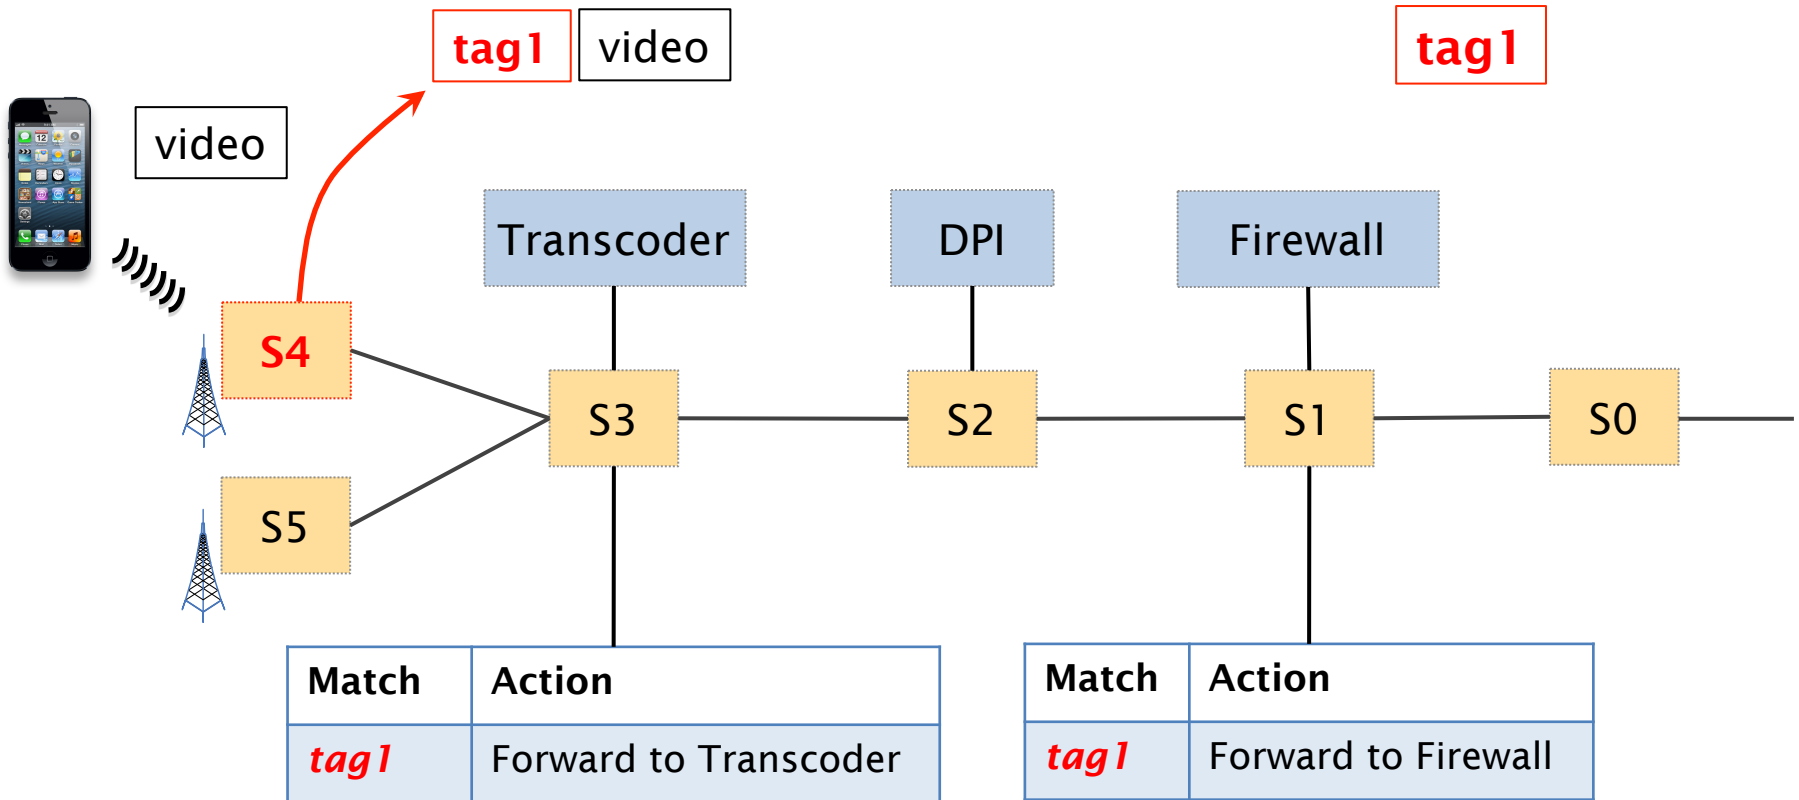


Video Traffic + Gold Membership = **Transcoder || Firewall**

tag1

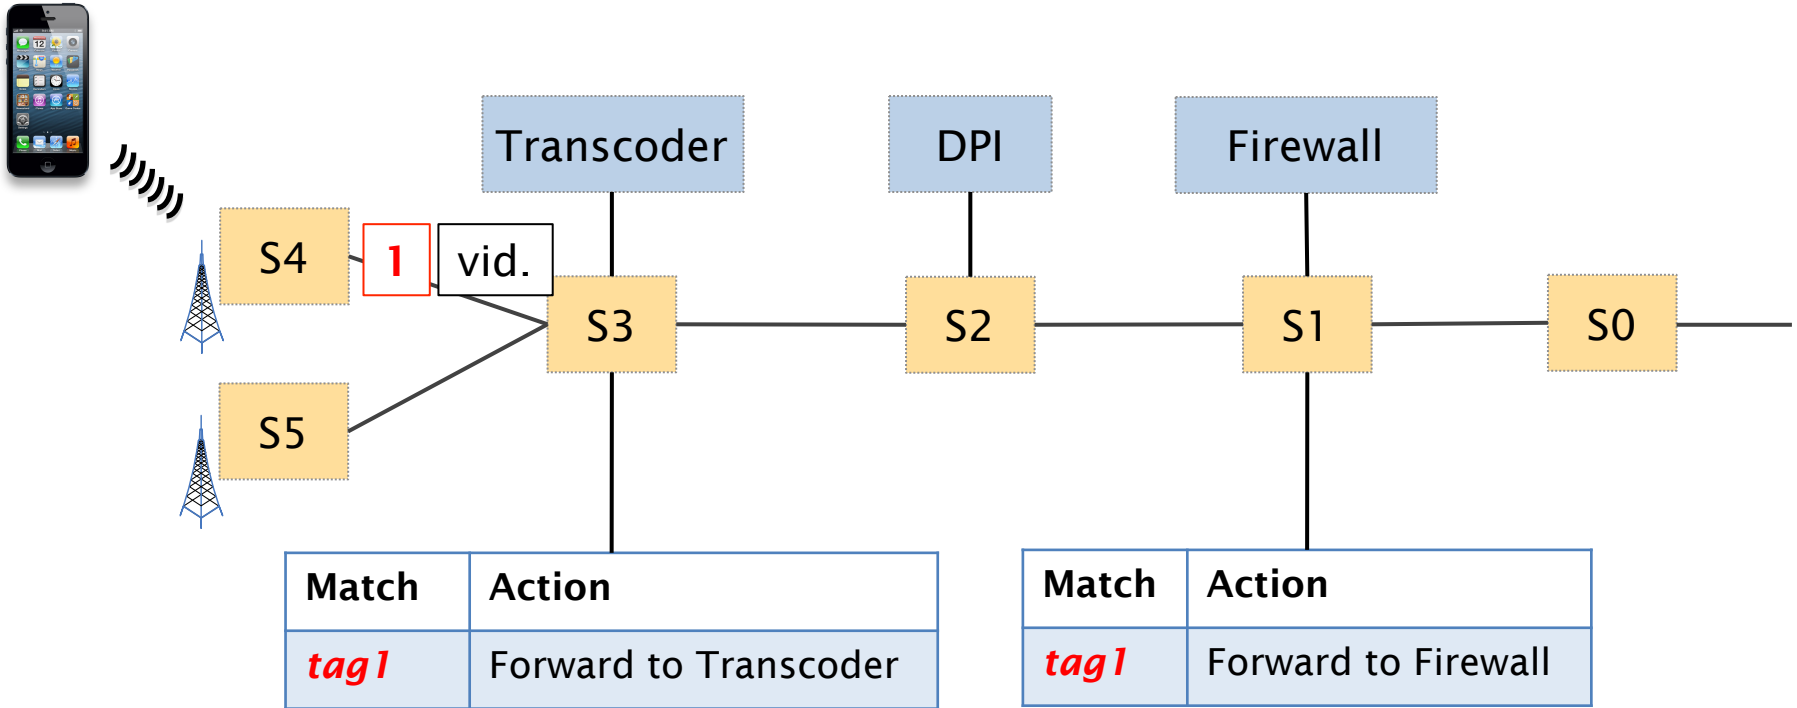


Video Traffic + Gold Membership = **Transcoder || Firewall**




Video Traffic + Gold Membership = **Transcoder || Firewall**

tag1

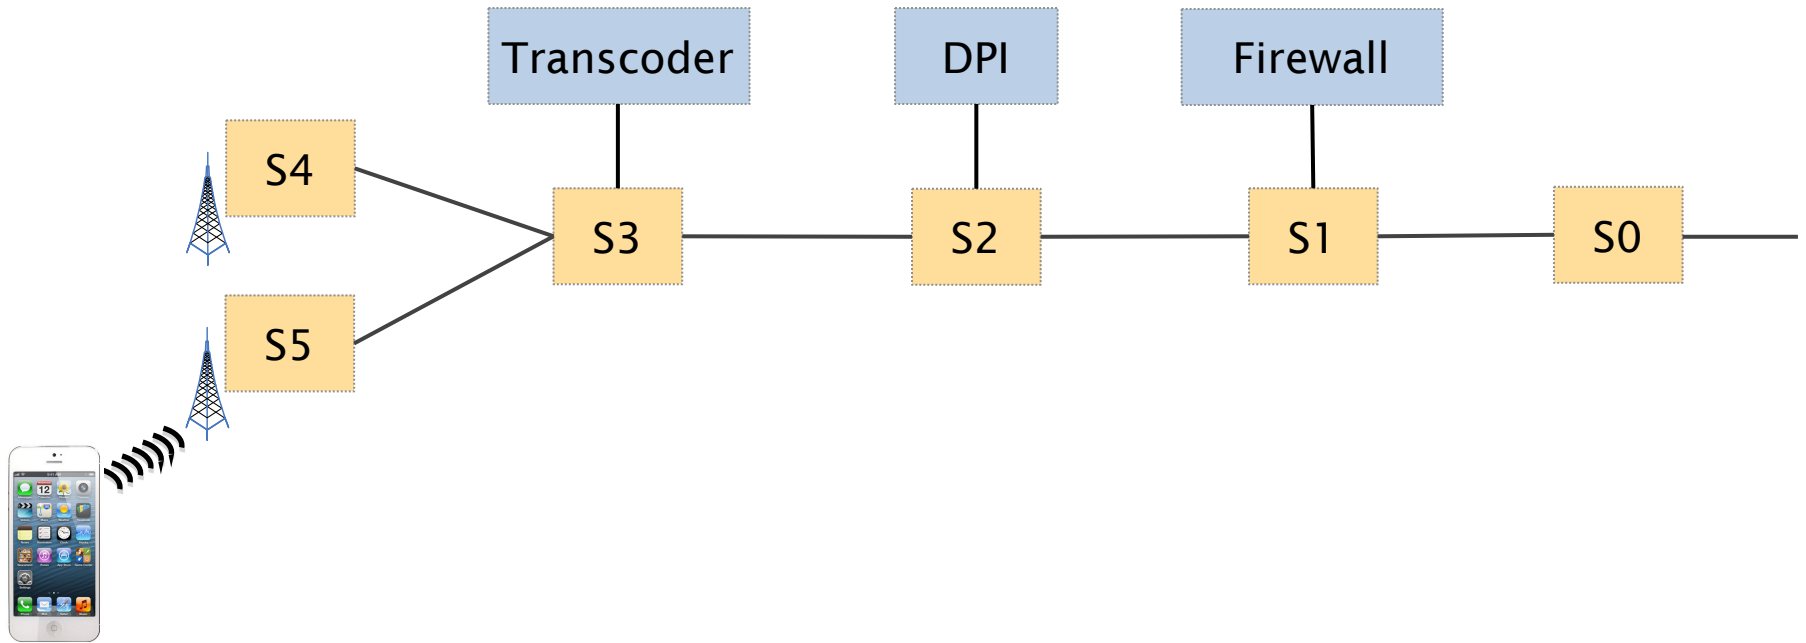


Web Traffic + Feature Phone = DPI || Firewall



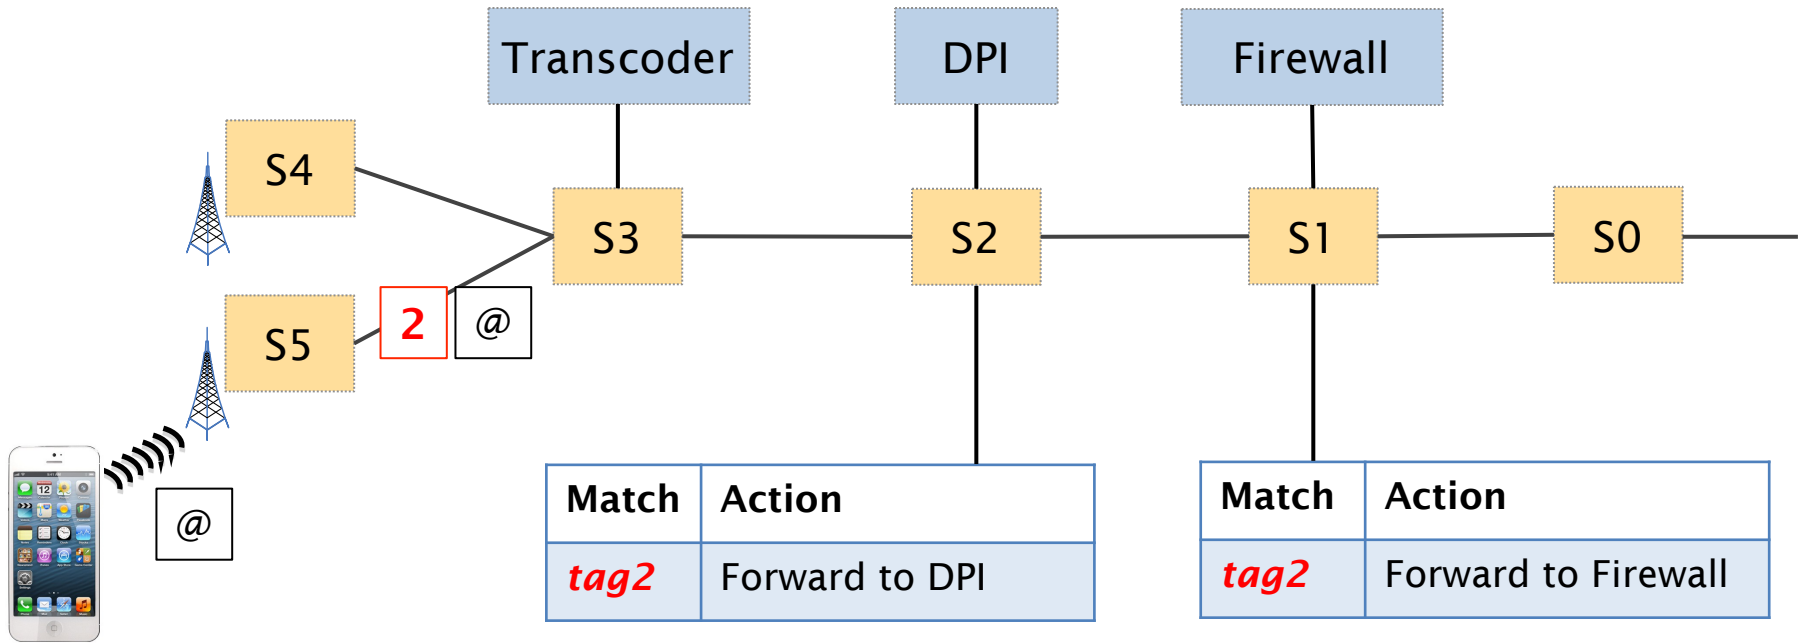
Web Traffic + Feature Phone = DPI || Firewall

tag2



Web Traffic + Feature Phone = DPI || Firewall

tag2




CellSDN tags are composed of three parts:


- Policy
- **Location**
- User Equipment Identifier

In CellSDN, each base station is associated with an IP prefix

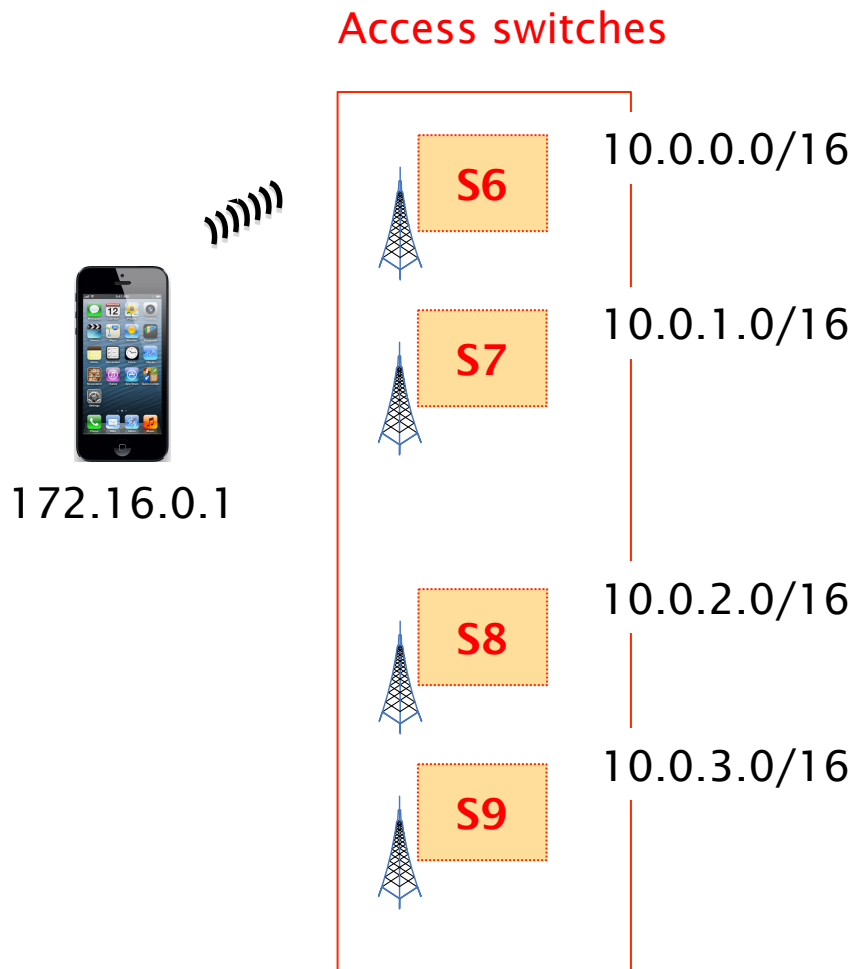
Access switches




In CellSDN, each base station is associated with an IP prefix




User Equipment receive a unique IP address



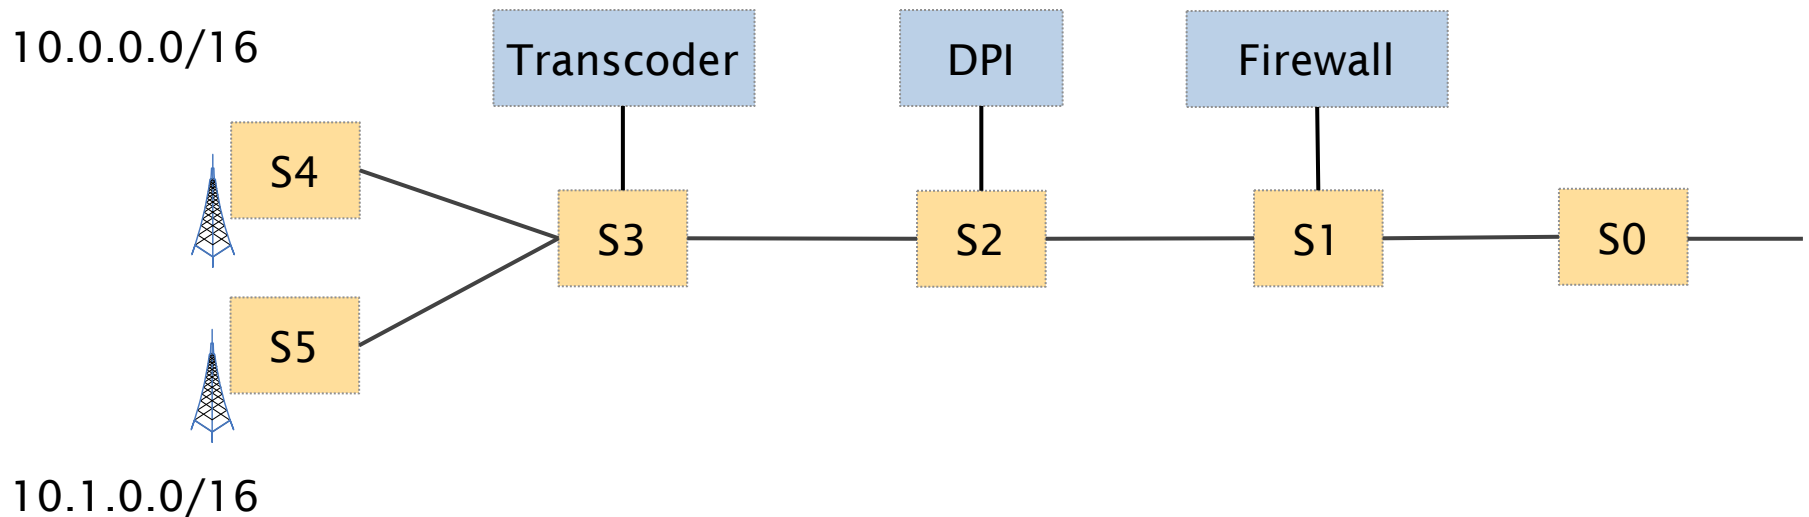
Access switches rewrite the UE IP into a location-dependent address



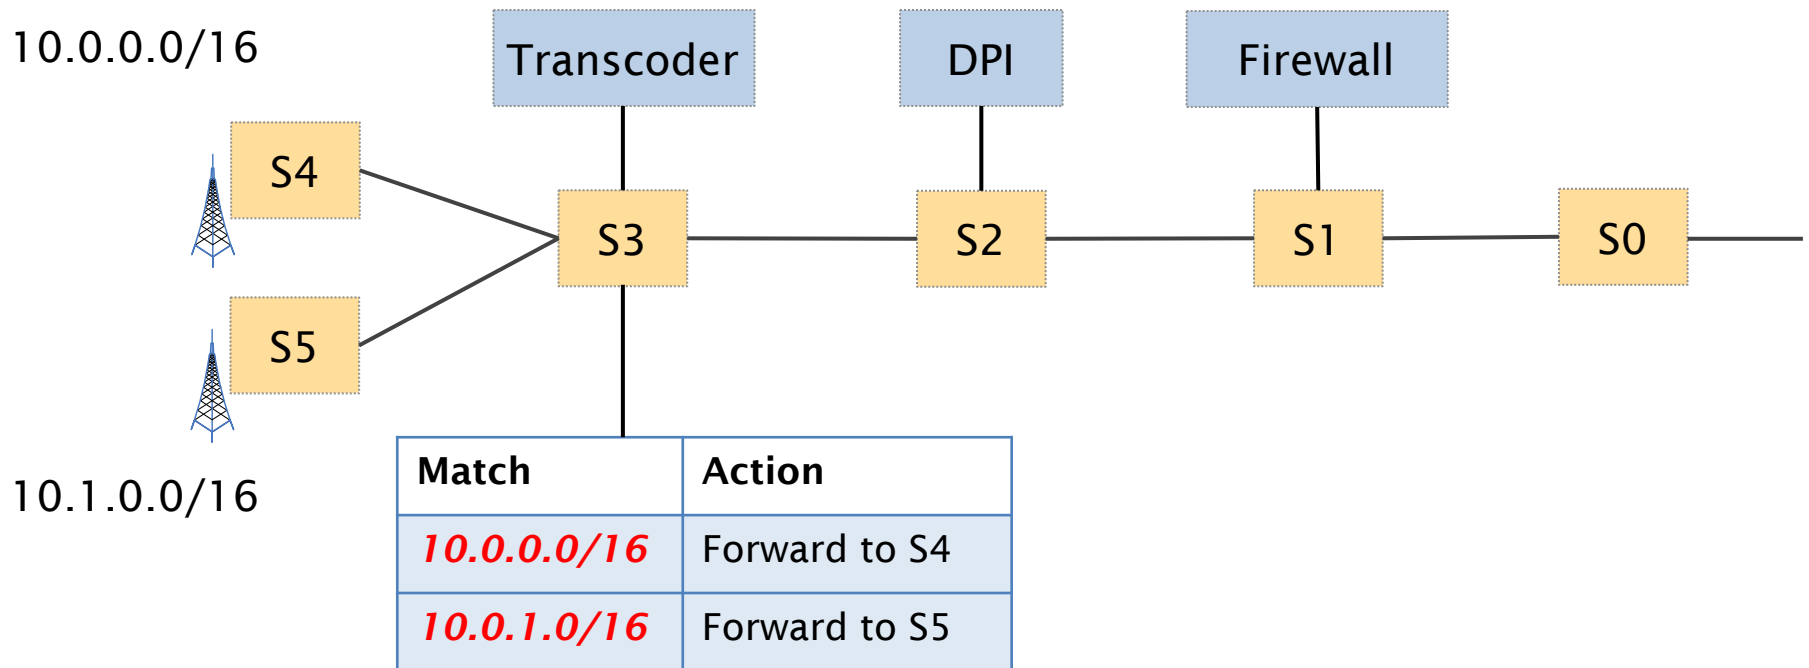
CellSDN can route based on Base Station prefixes



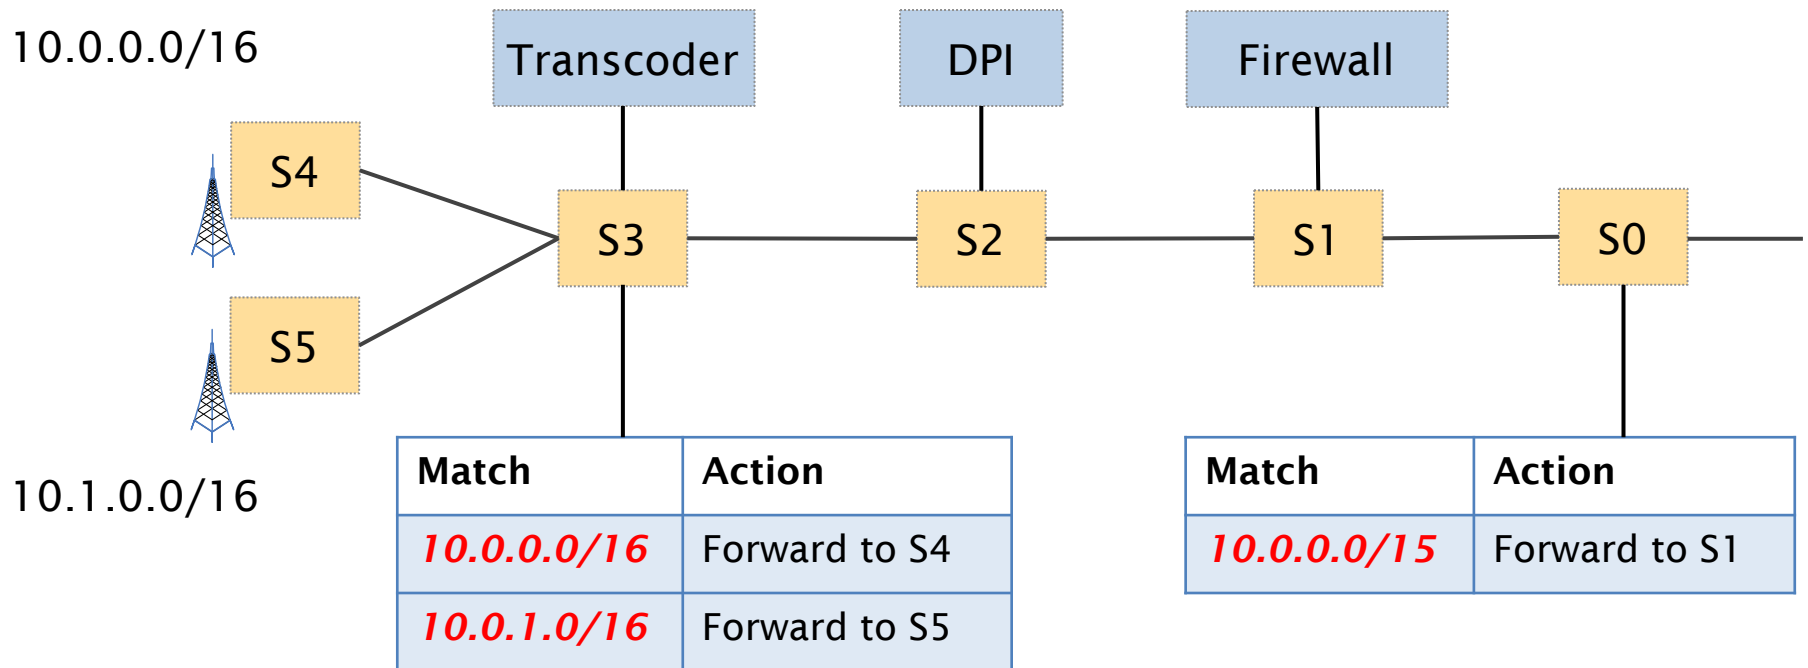
CellSDN can route based on Base Station prefixes



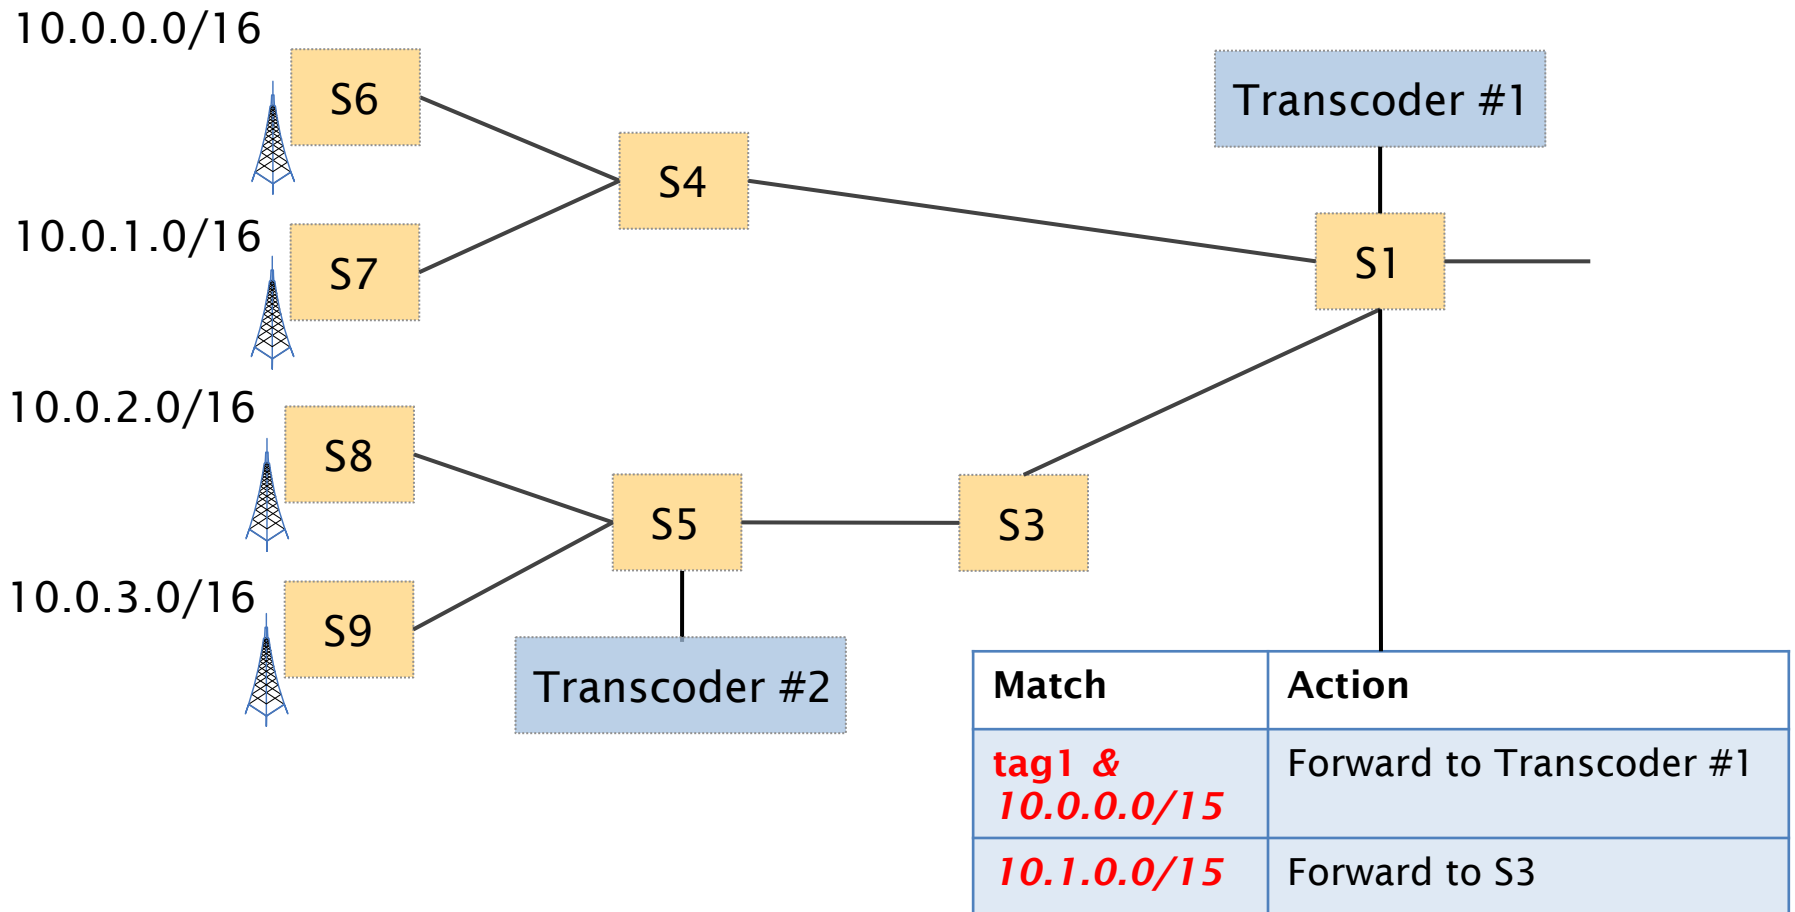
CellSDN can route based on Base Station prefixes



CellSDN automatically aggregates adjacent prefixes



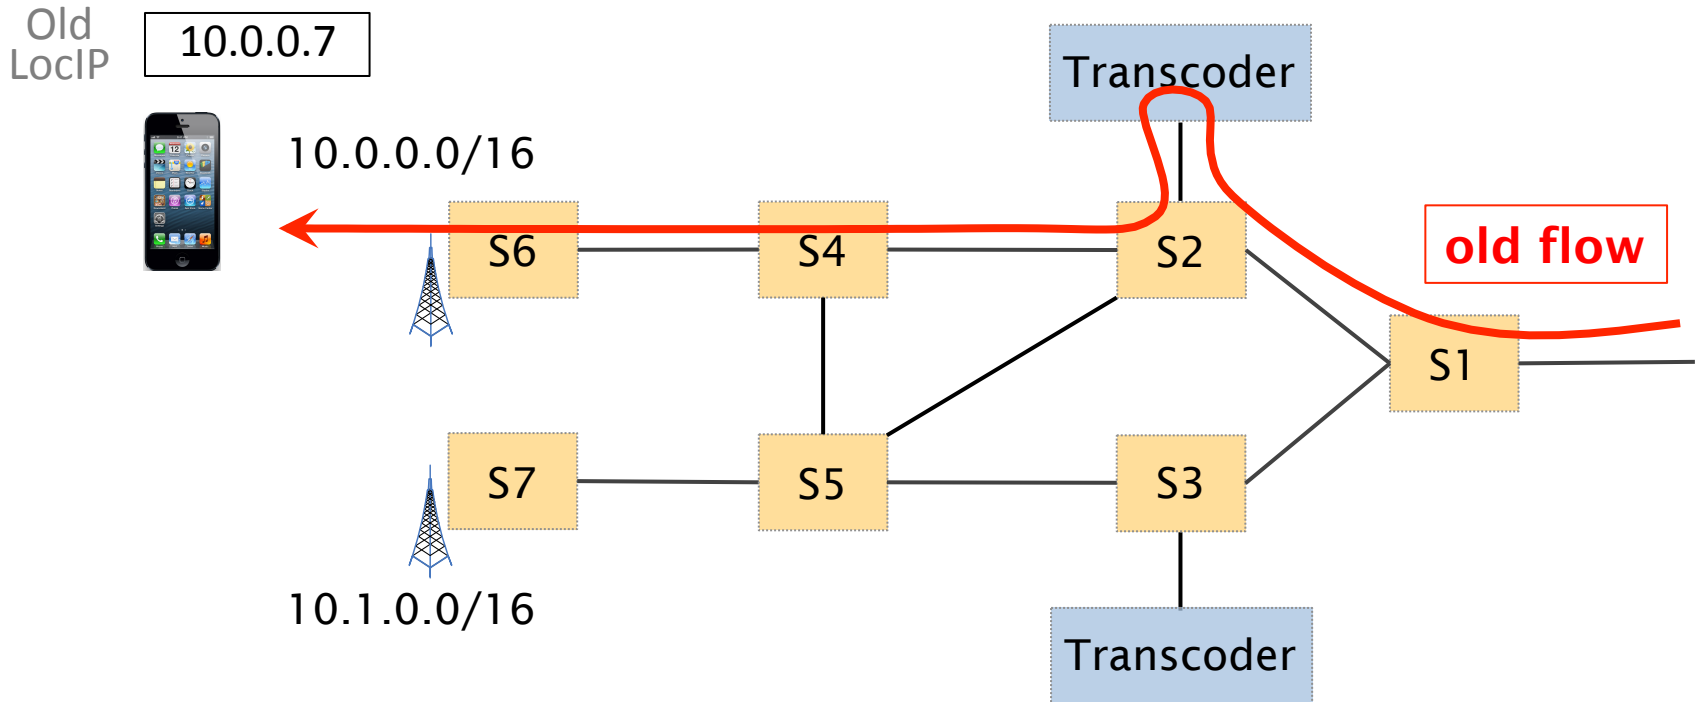
CellSDN can selectively match on tag and BS ID for load-balancing



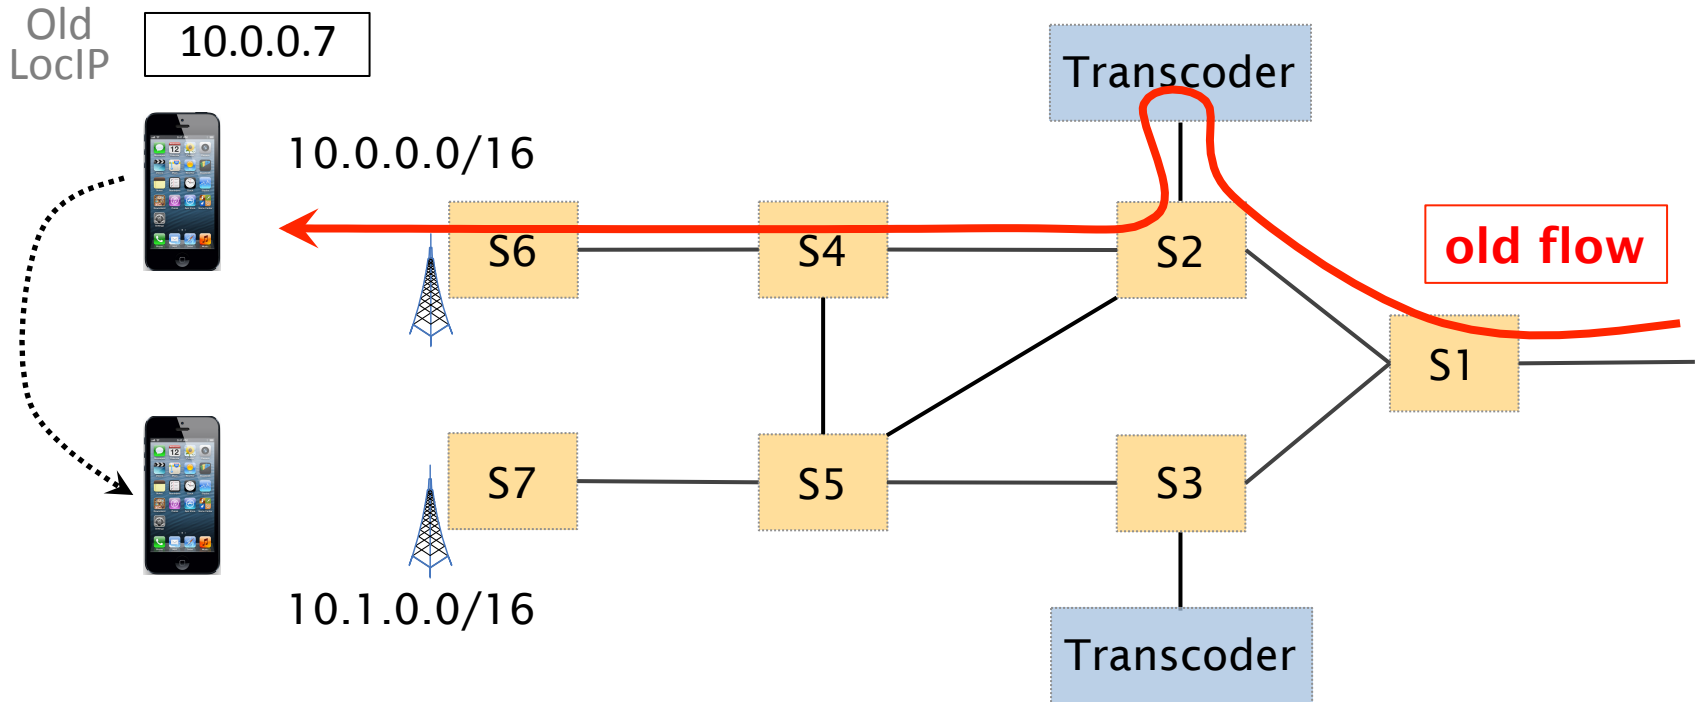
CellSDN tags are composed of three parts:

- Policy
- Location
- **User Equipment Identifier**

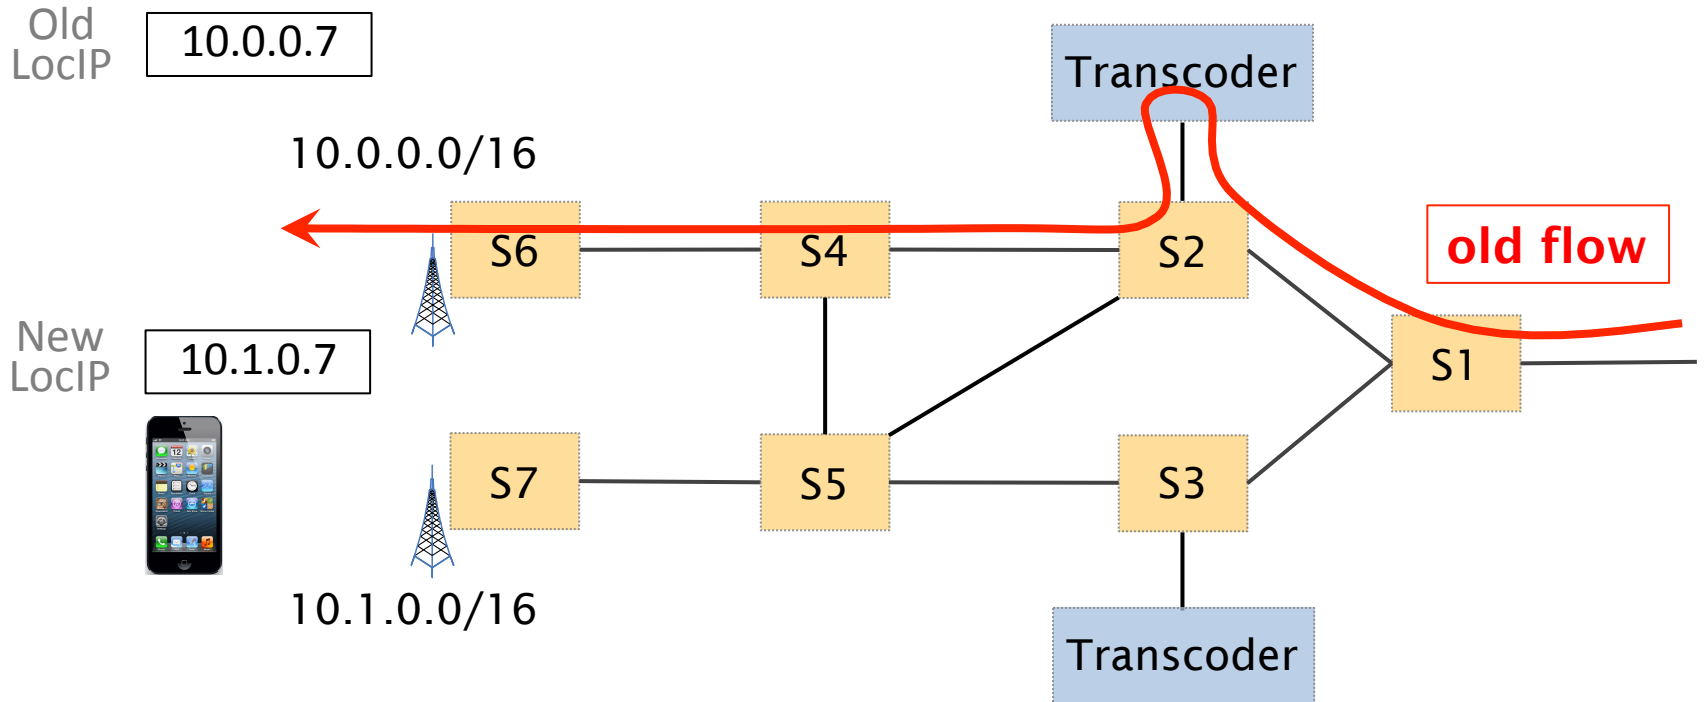
CellSDN UE tag enables mobility



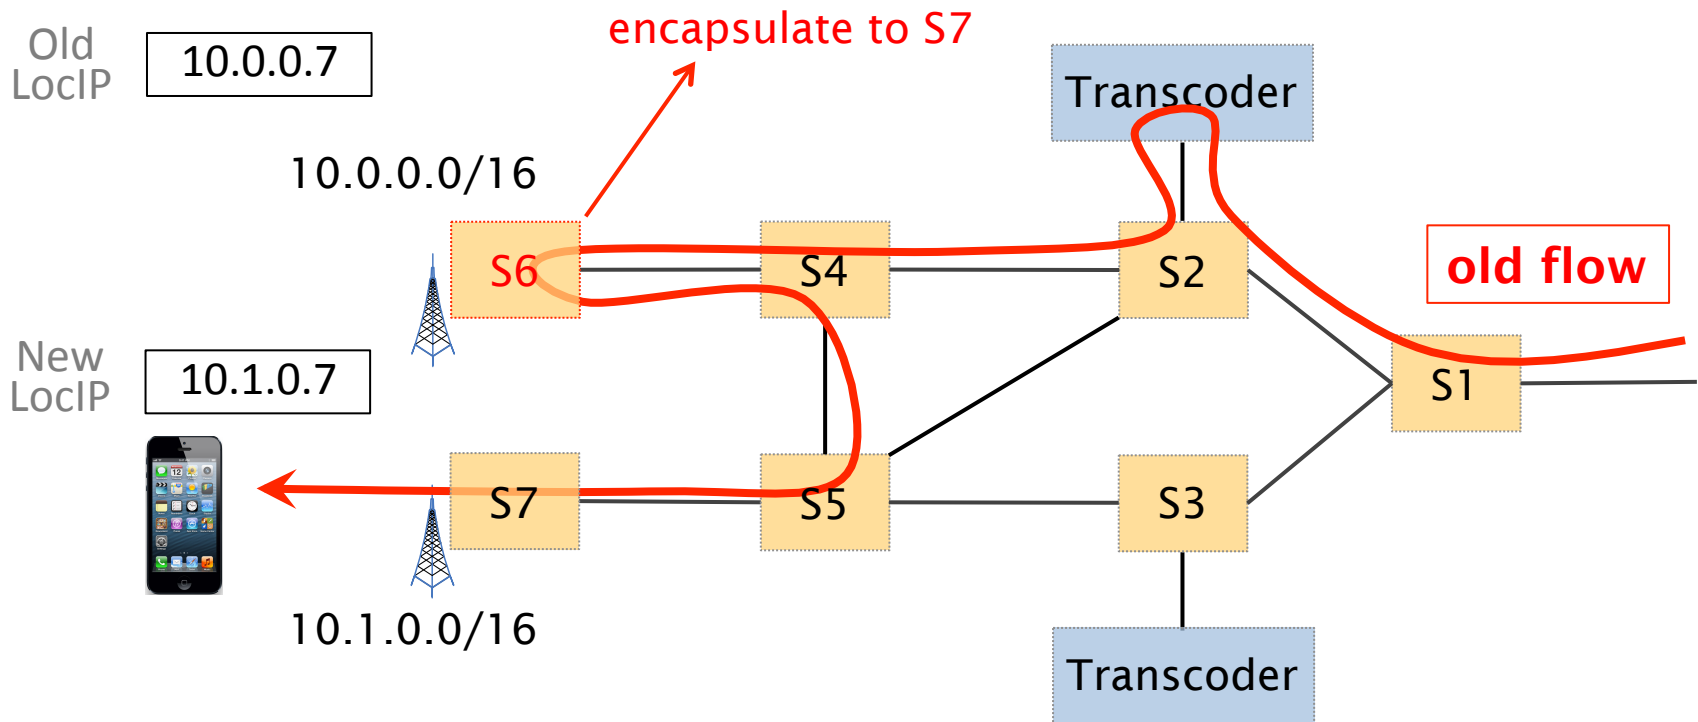
When an user moves, she receives a new LocIP



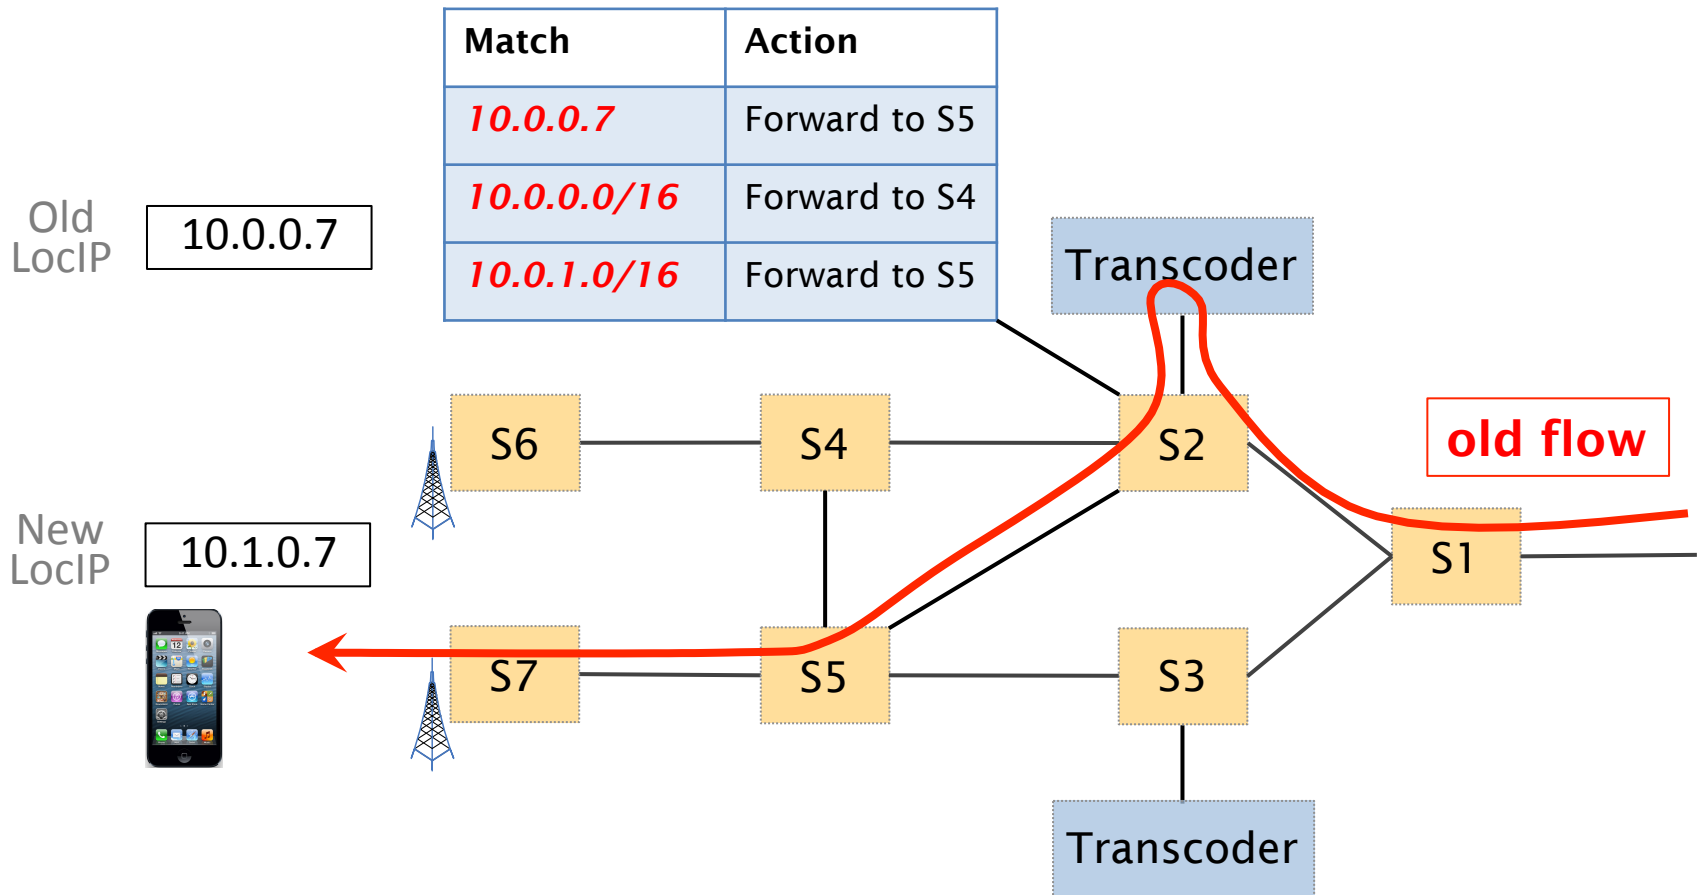
When an user moves, she receives a new LocIP



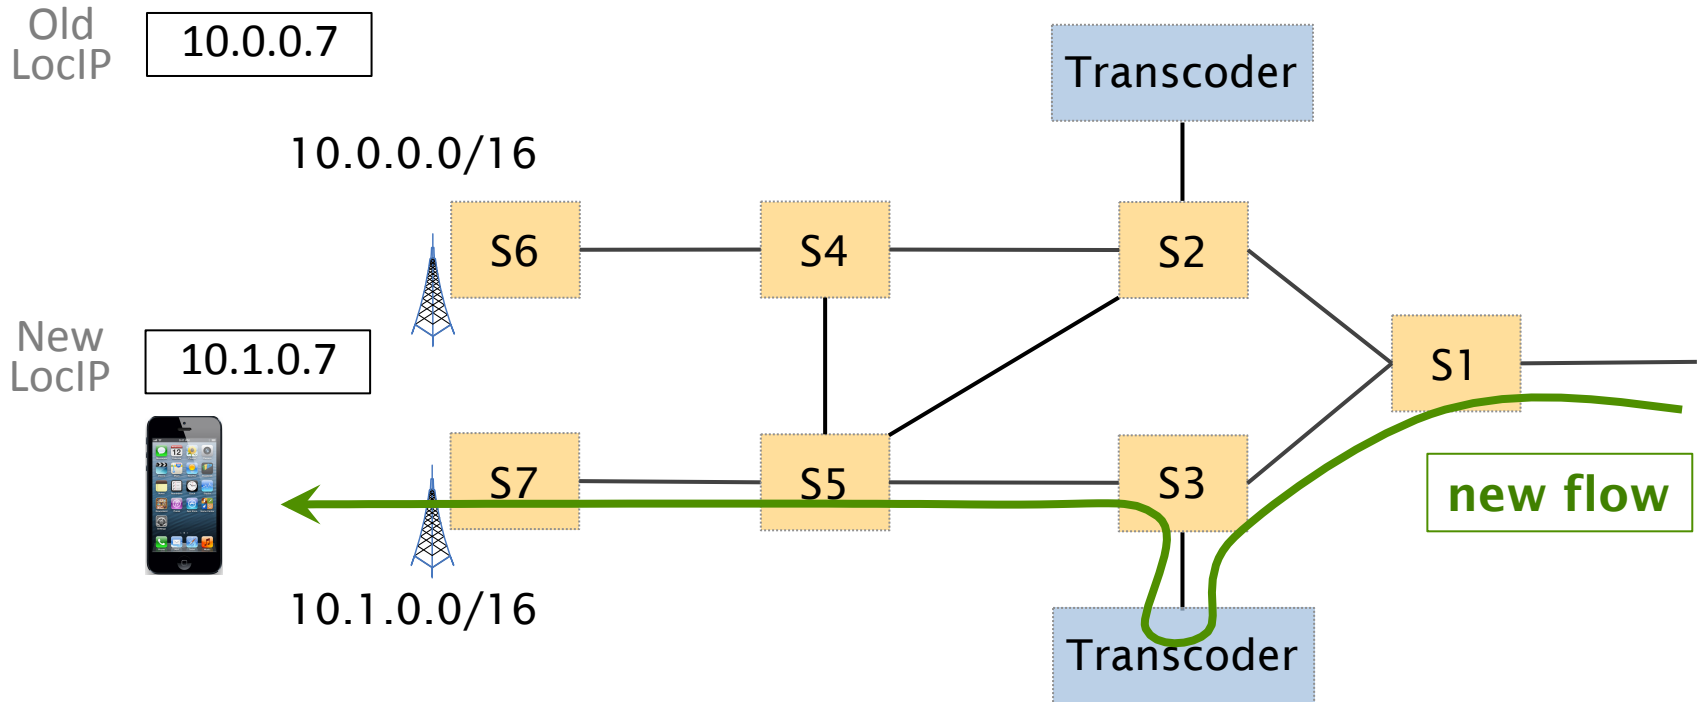
Old flows reach the previous base station, ensuring policy consistency



To avoid triangle routing, CellSDN can install shortcut path after the last MB



New flows automatically flow along new policy paths



CellSDN tags are composed of three parts:

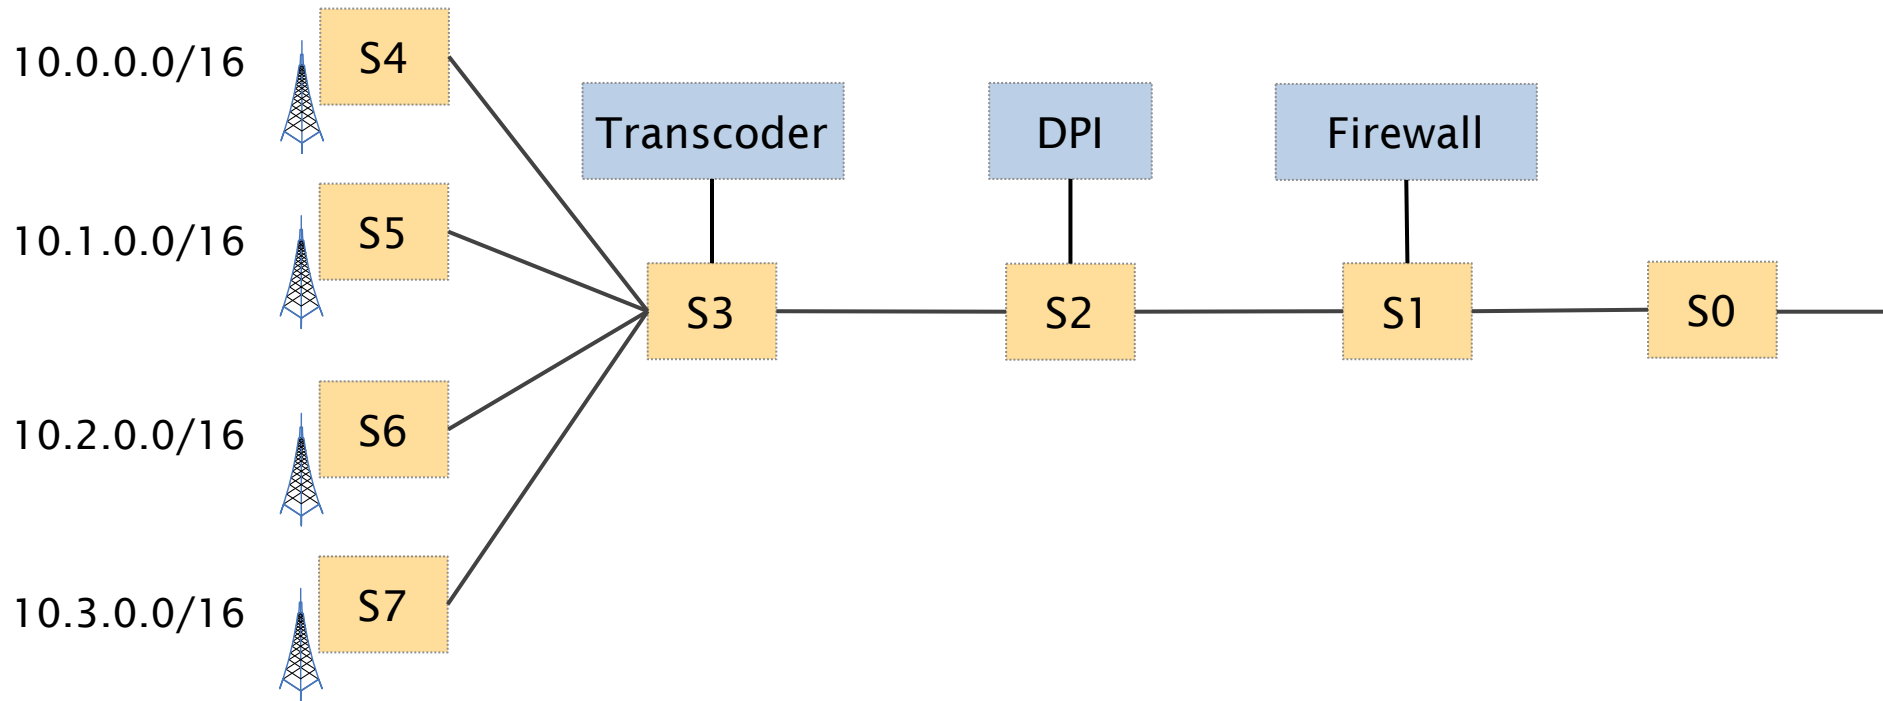
- Policy
- Location
- User Equipment Identifier

How does CellSDN compute these tags?

CellSDN minimizes the forwarding tables size by reusing tags

Given a service policy-path,

1. Compute the actual path P used in the network
2. For each candidate tag t used on the path P ,
 Compute the # of new rules needed if t is used
3. Select the candidate tag minimizing the # of new rules,
 Create new tag if none available
4. Install the forwarding entries in the network



policy path

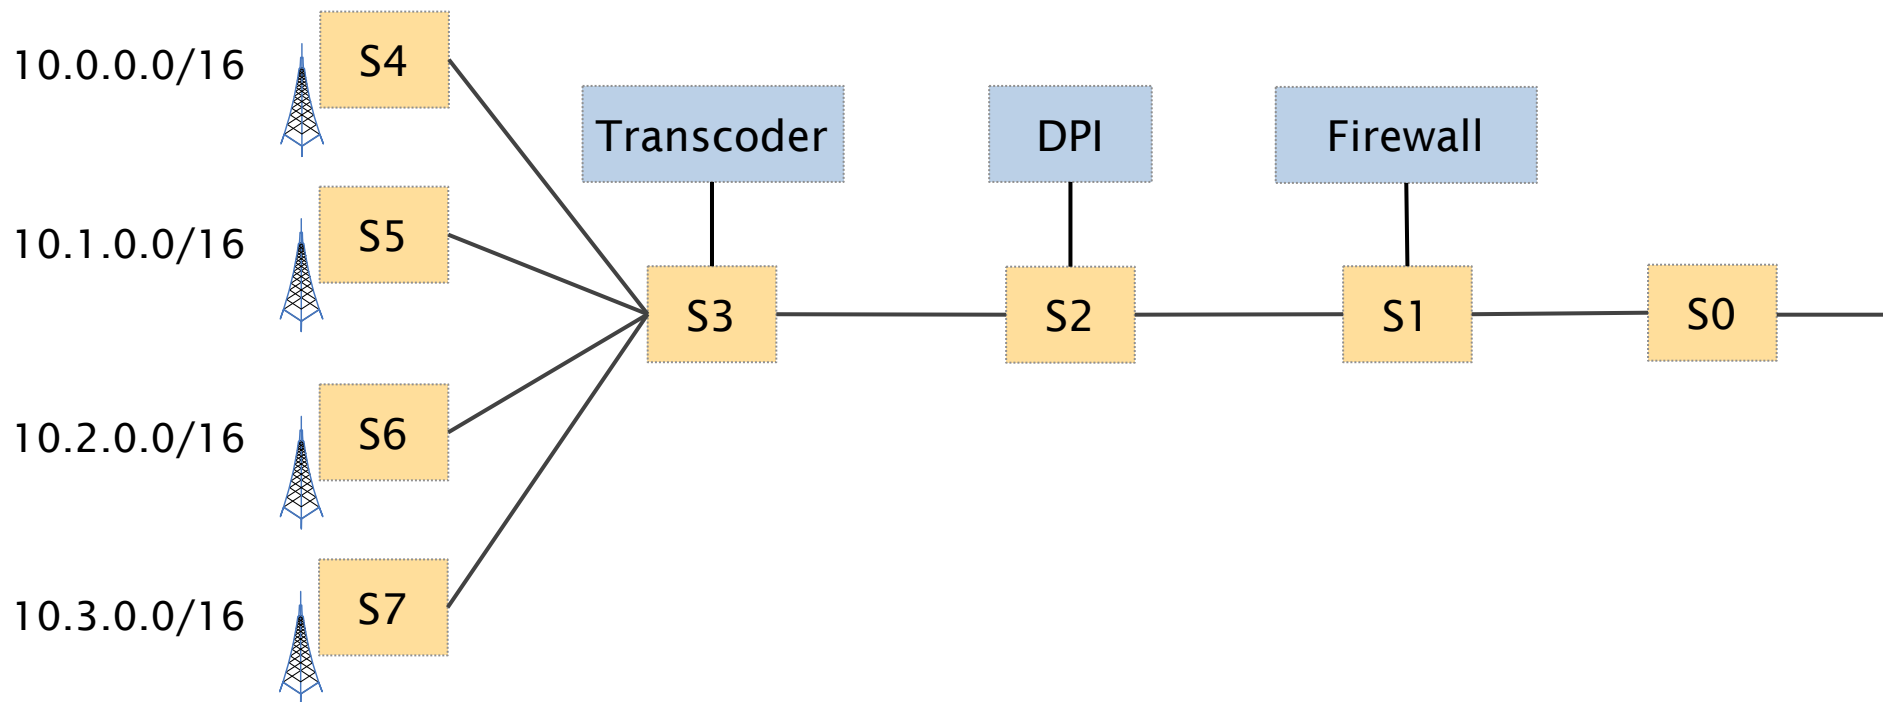
S4



DPI



Firewall



policy path

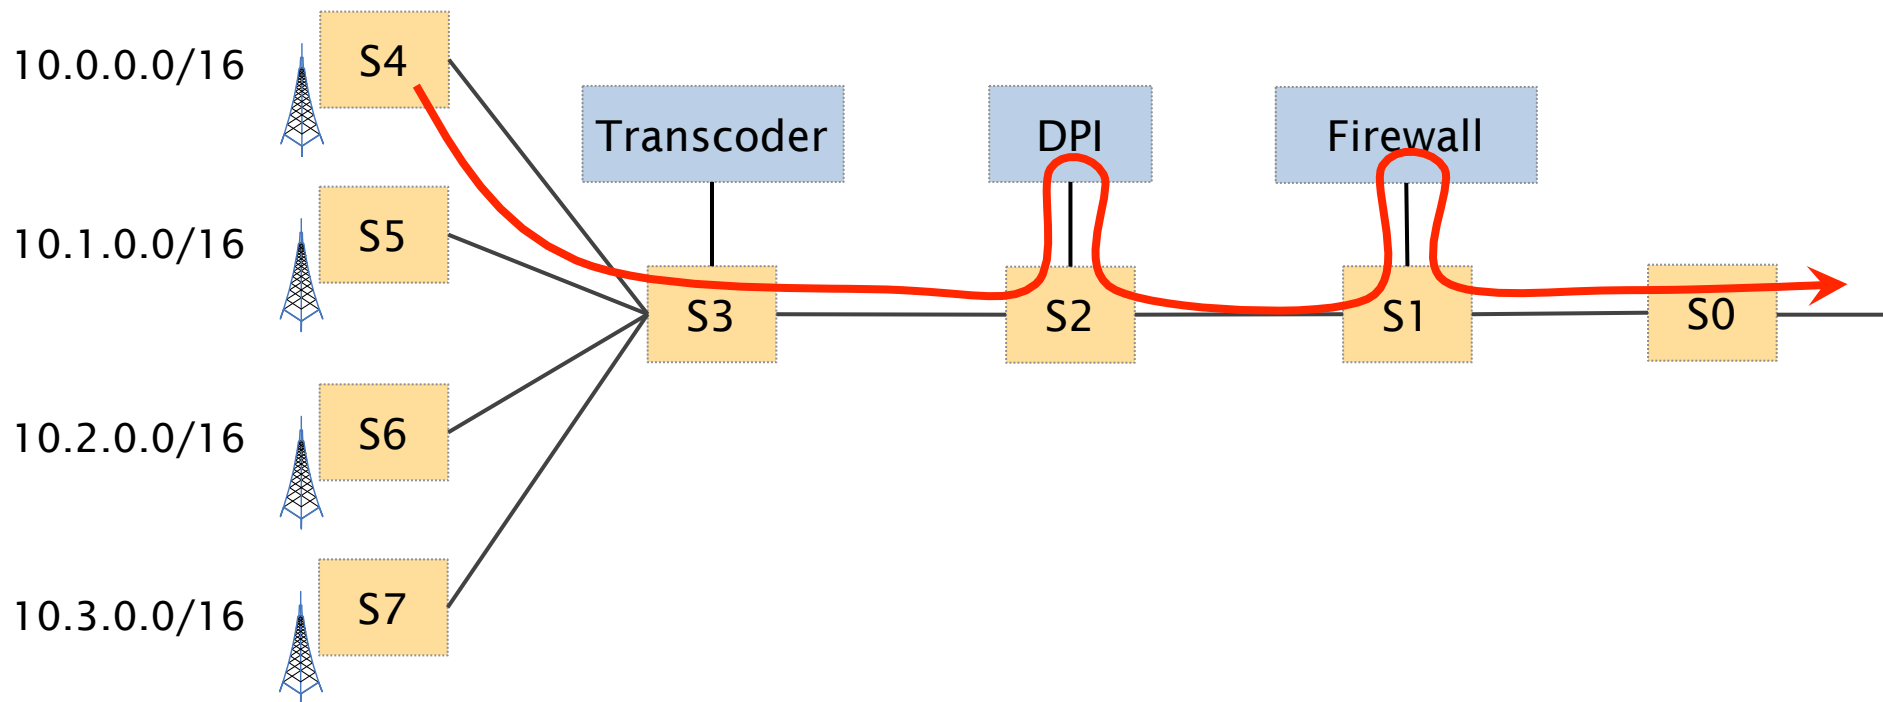
S4



DPI

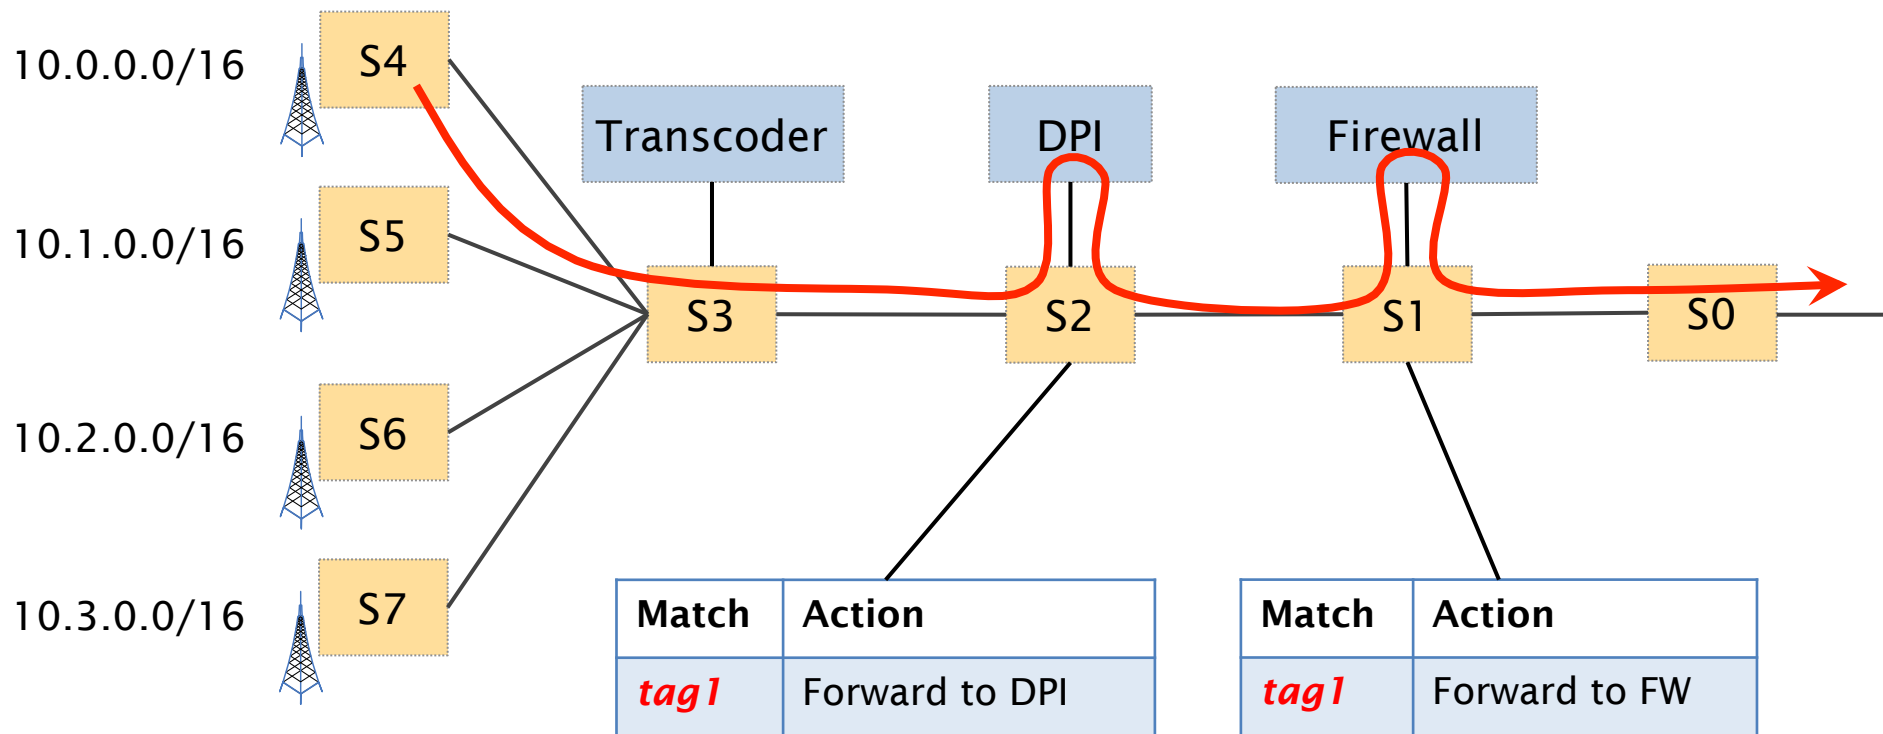


Firewall



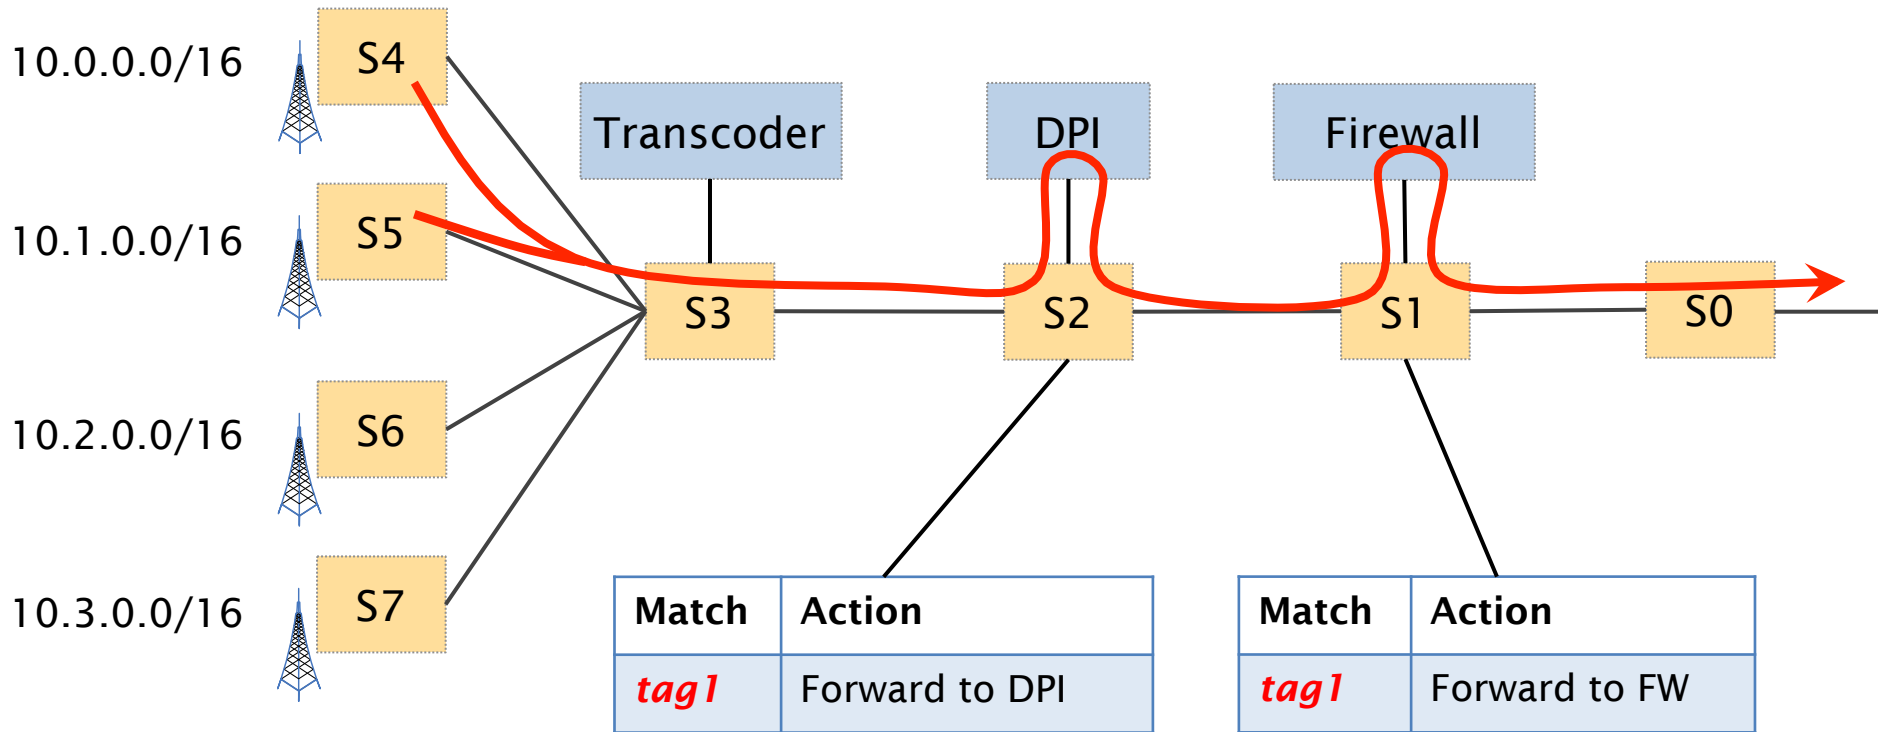
policy path

S4 → DPI → Firewall



policy path

S5 → DPI → Firewall

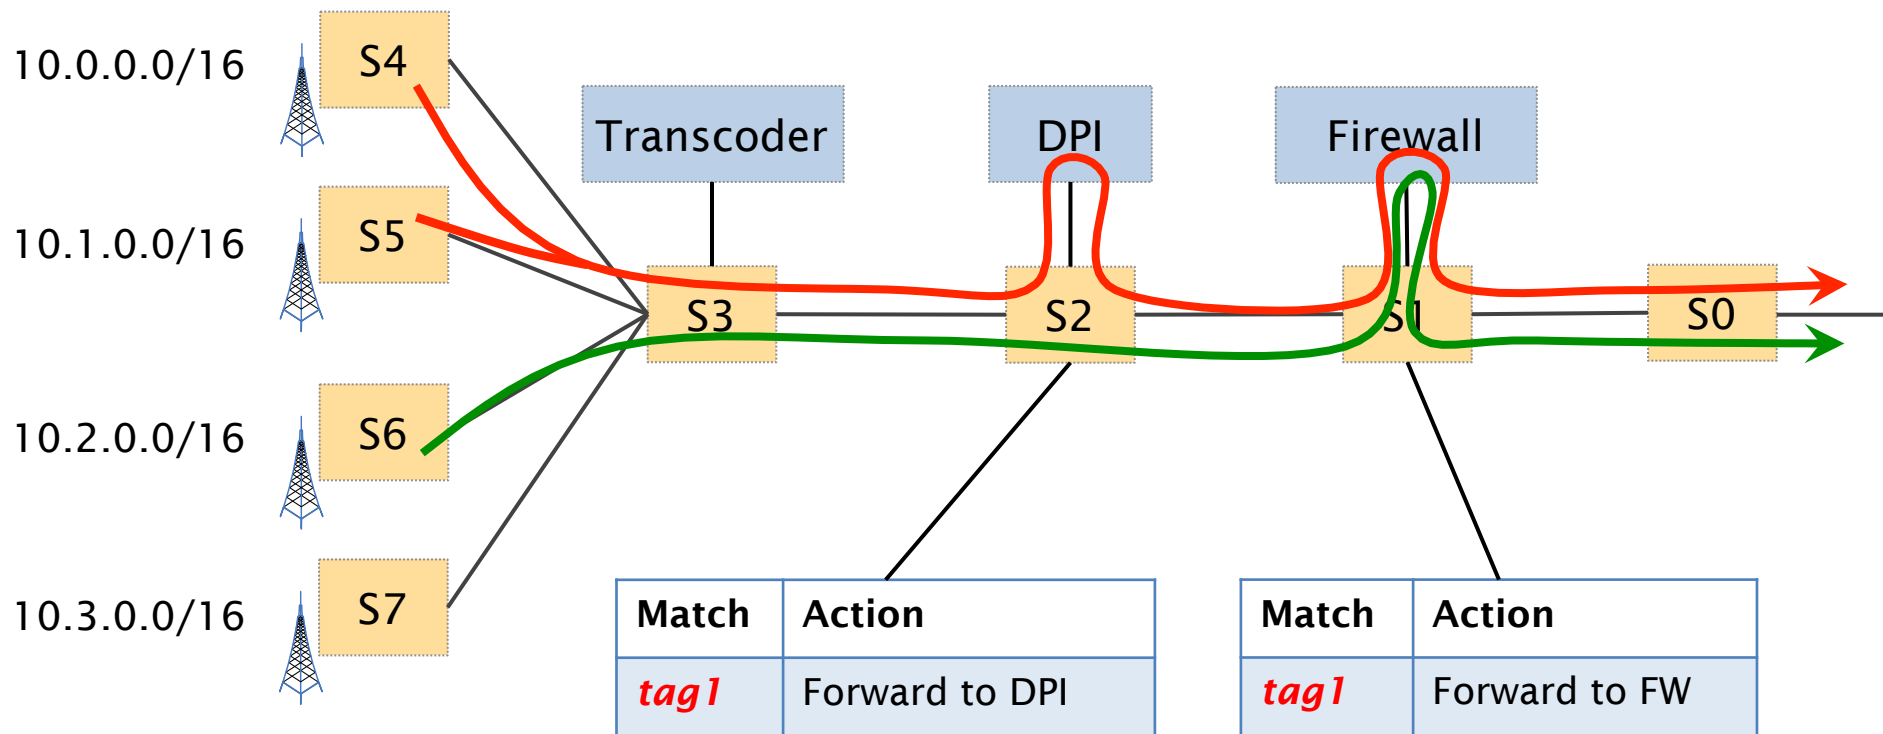


policy path

S6



Firewall




policy path

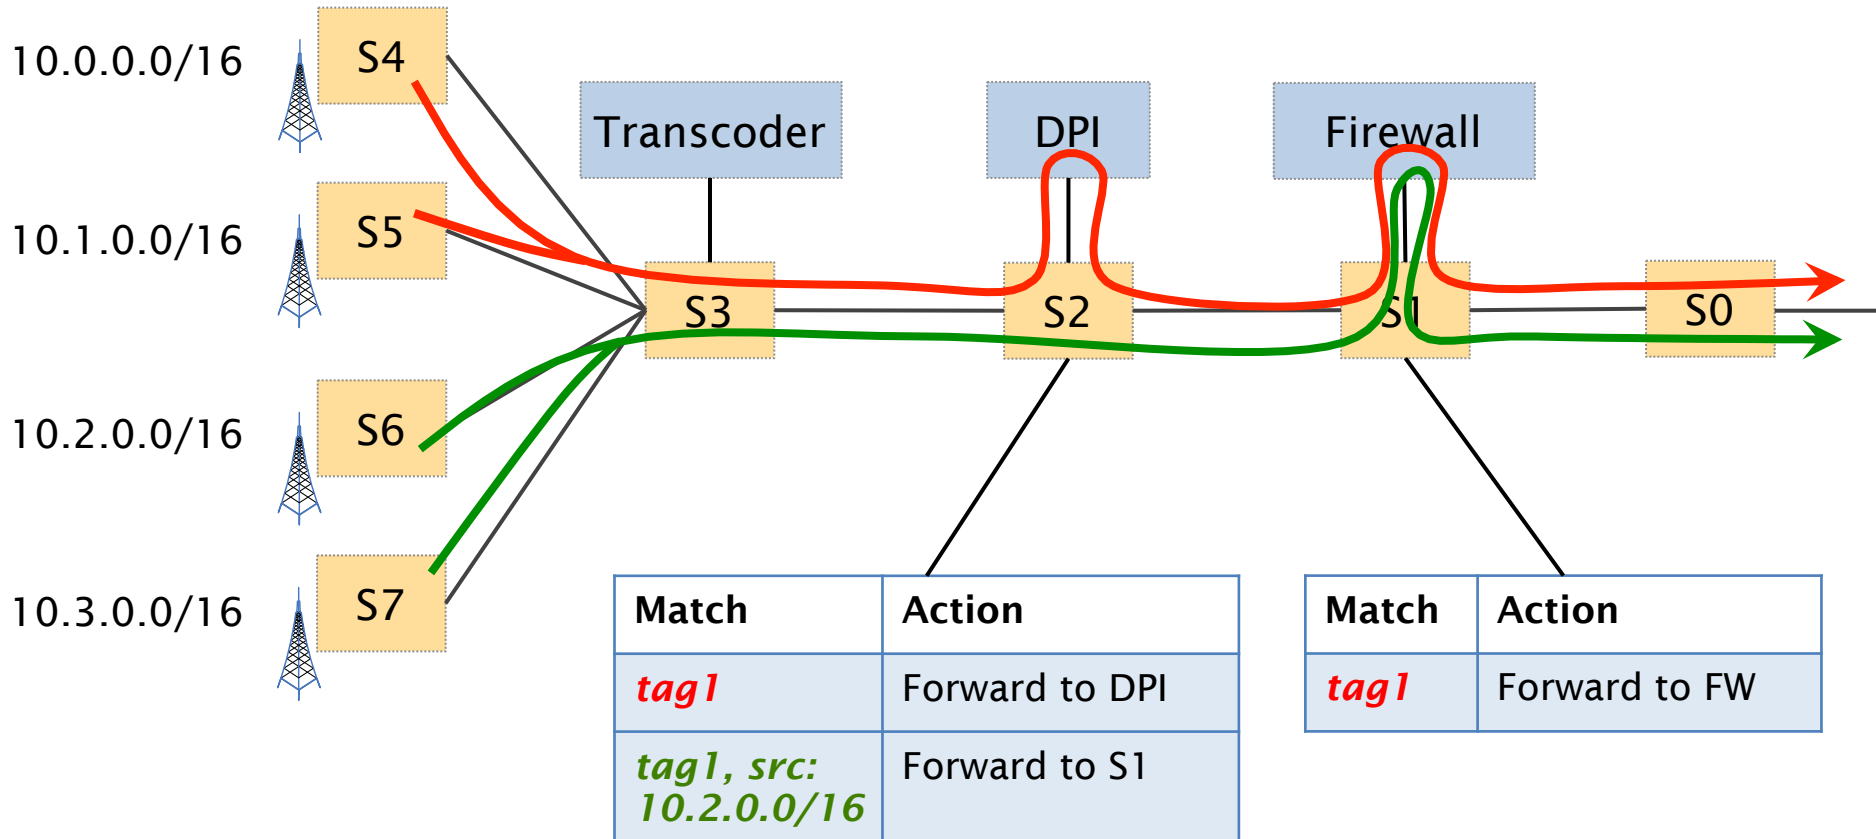
S6



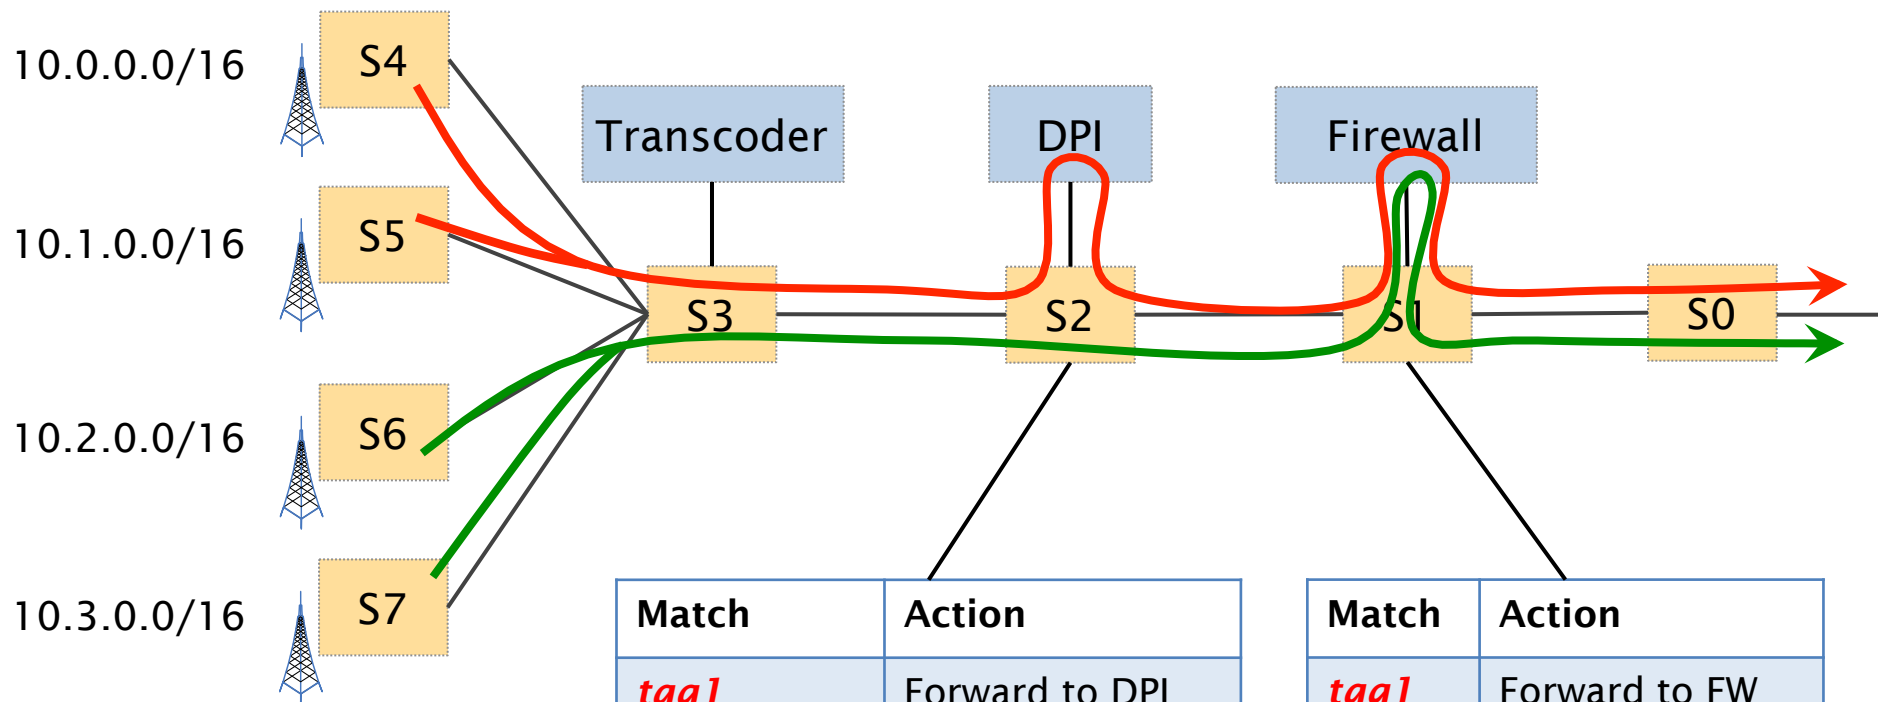
Firewall



policy path **S7**  **Firewall**



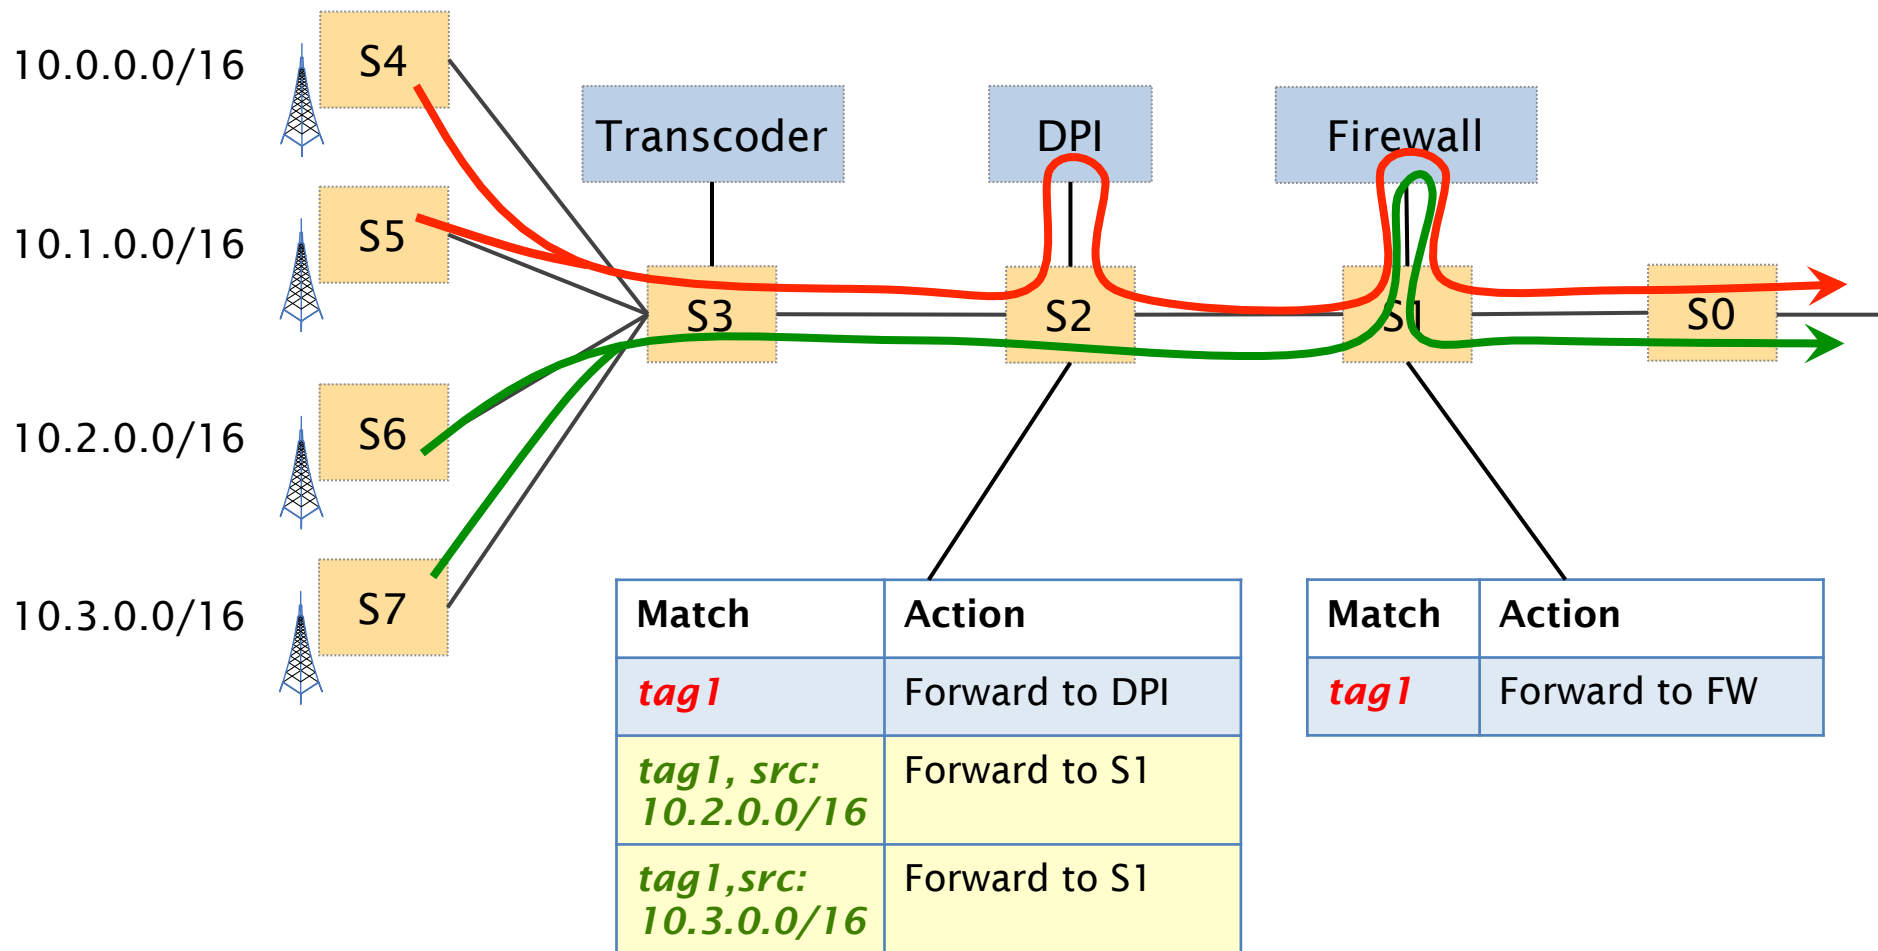
policy path **S7** \longrightarrow Firewall



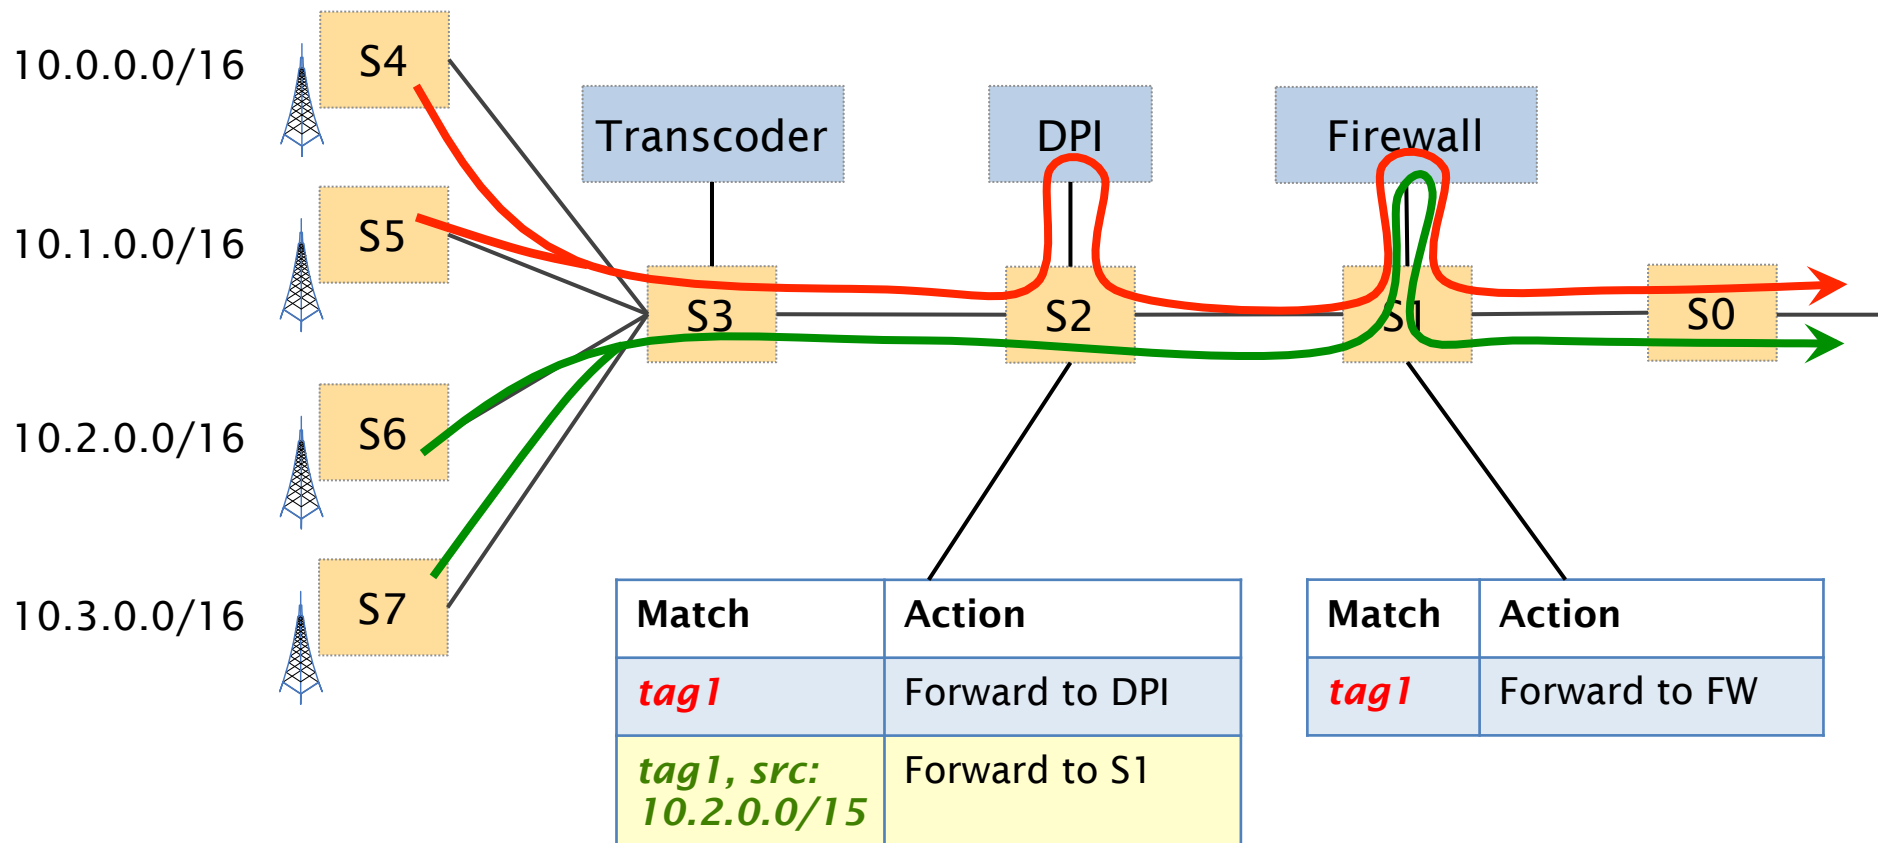
Match	Action
<i>tag1</i>	Forward to DPI
<i>tag1, src: 10.2.0.0/16</i>	Forward to S1
<i>tag1, src: 10.3.0.0/16</i>	Forward to S1

Match	Action
<i>tag1</i>	Forward to FW

policy path **S7** \longrightarrow **Firewall**



policy path **S7**  **Firewall**



CellSDN dataplane can support a large number of service policies


Simulation

- FatTree topology
- 128 switches && 1280 base stations
- 8 types of middleboxes
- 5 middleboxes long service policy


Evaluation

table size in function of the # of policies

Switches table size increase linearly wrt to the # of policy paths

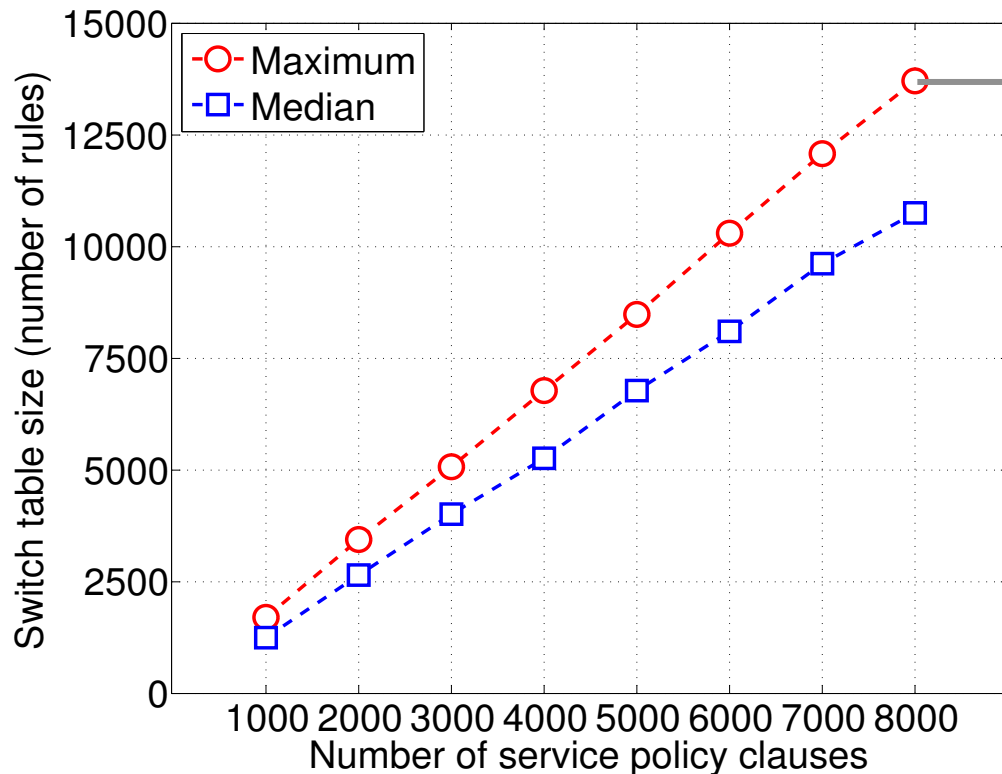


Switches table size increase linearly wrt to the # of policy paths




Only 5k entries are required to support 4k policy paths

Switches table size increase linearly wrt to the # of policy paths



In the worst case, 8k policy paths require 13.6k entries


Switches table size increase linearly wrt to the # of policy paths



In the worst case, 8k policy paths require 13.6k entries

Today, operators require a few hundreds policies

Switches table size increase linearly wrt to the # of policy paths



CellSDN **works** on commodity switches

CellSDN: Taking control of cellular core networks



Architecture


software-defined network

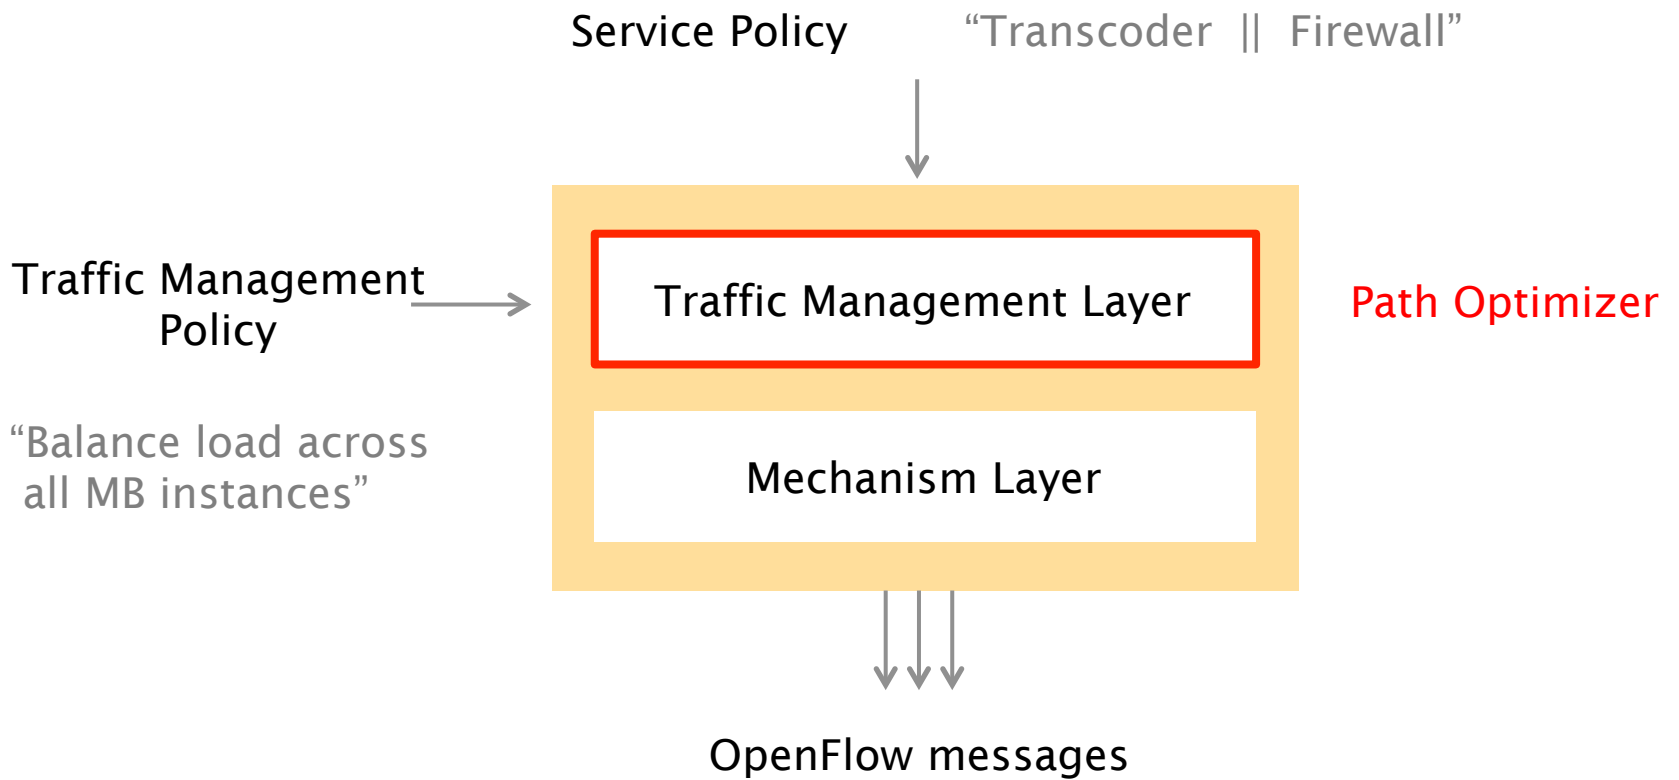
Scaling the data-plane

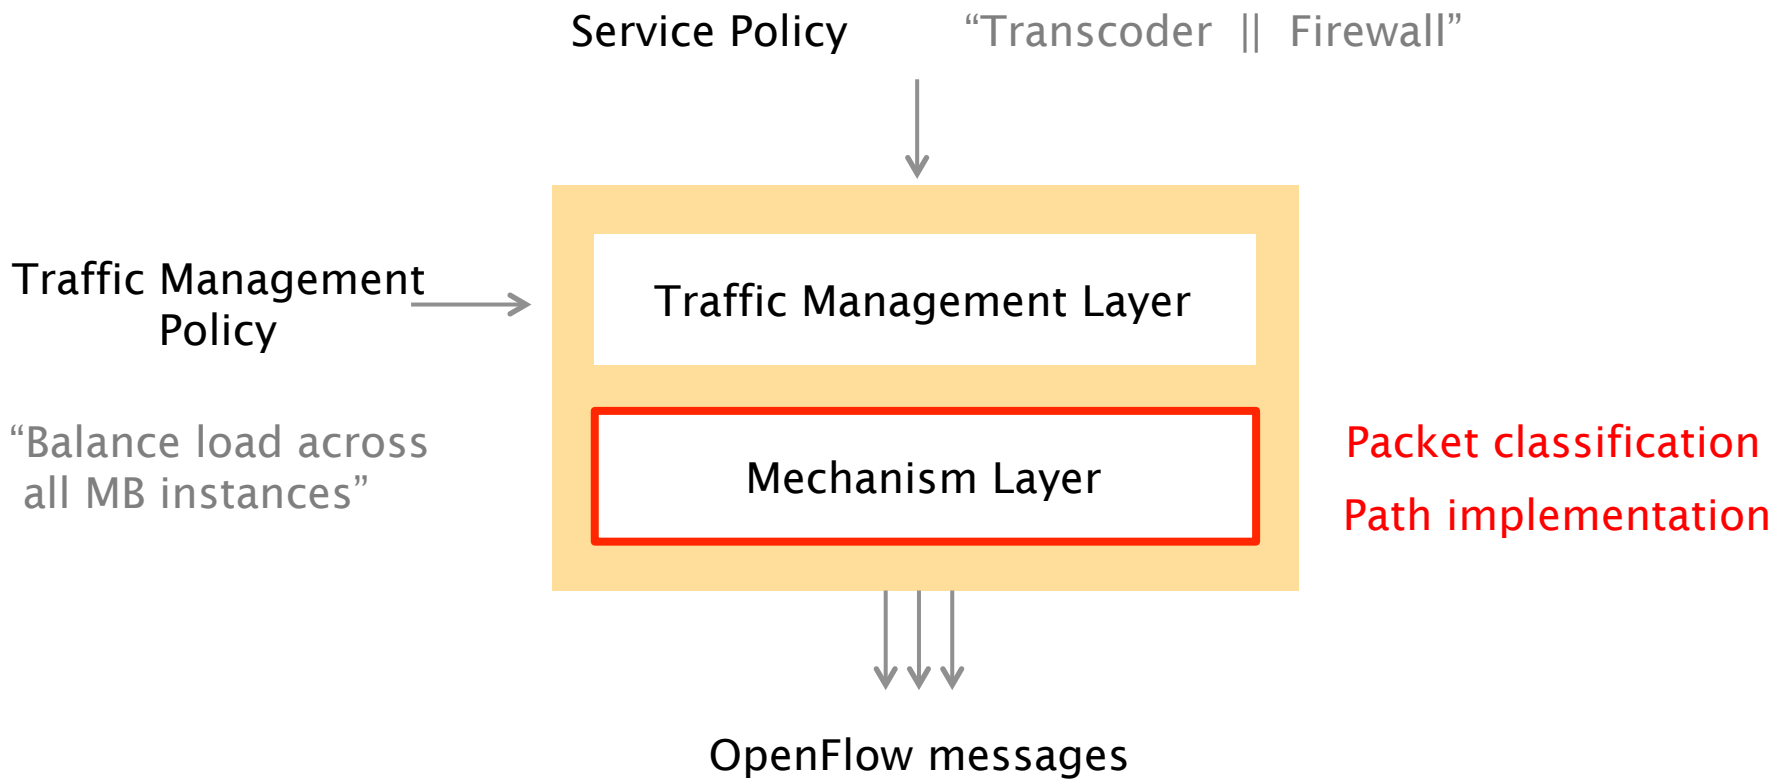
multi-dimensional tagging

3 **Scaling the control-plane
tasks delegation**

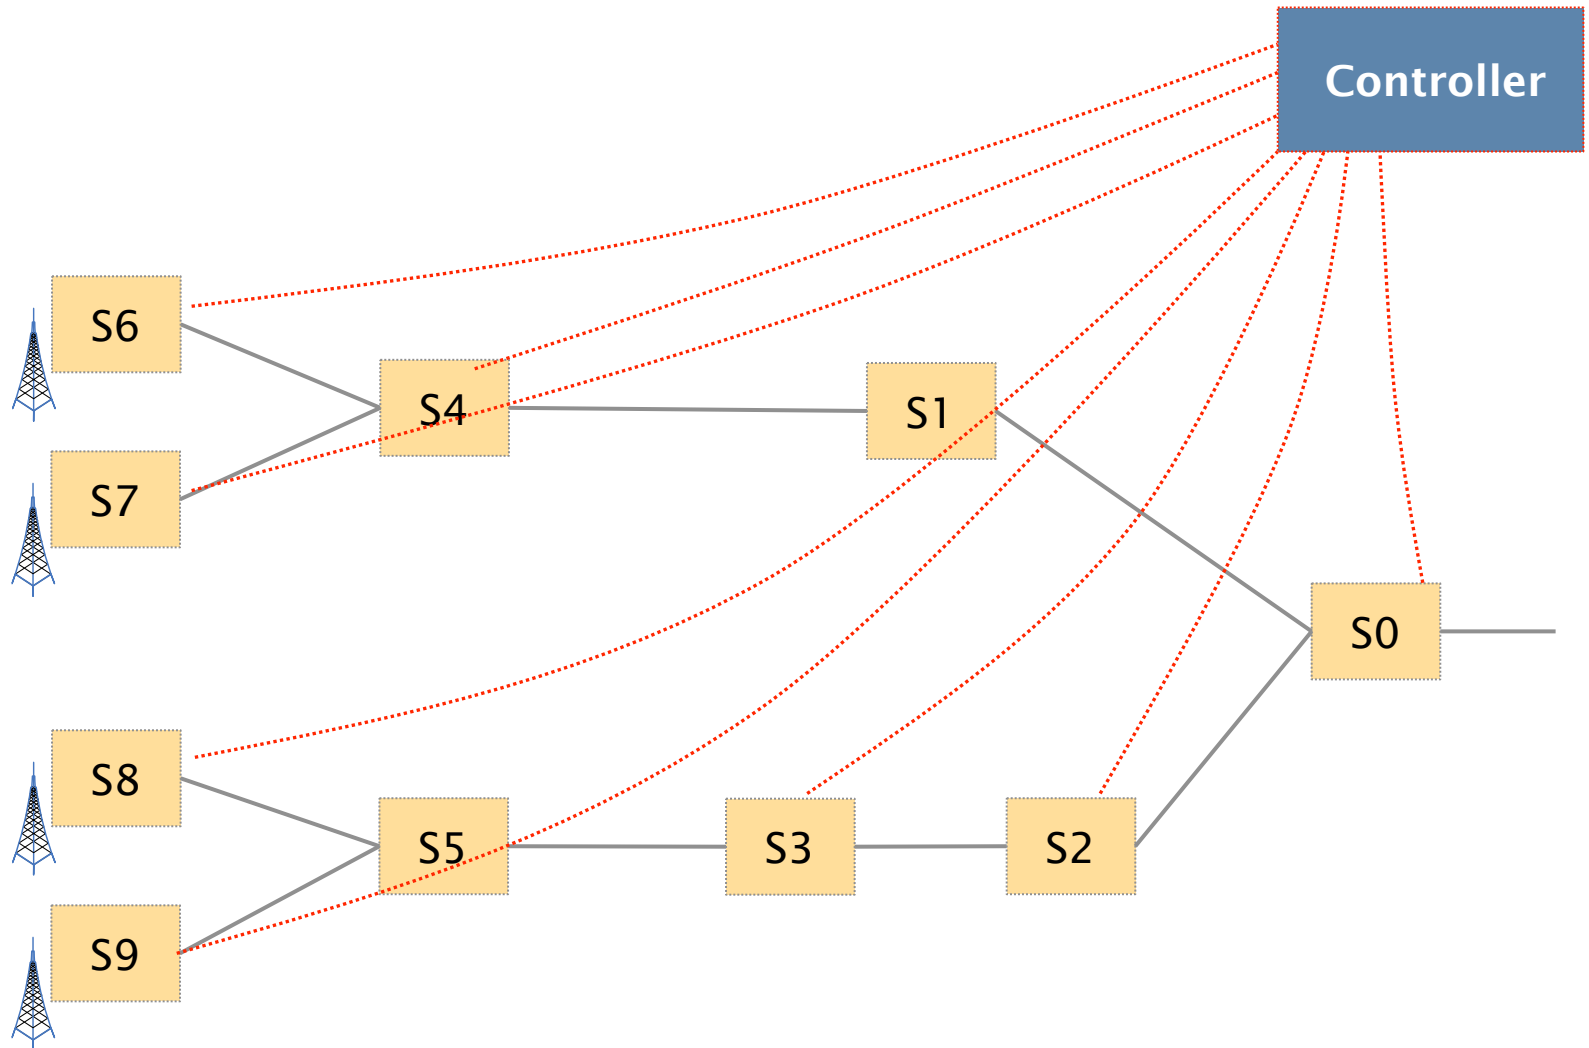
CellSDN controller separates traffic management from rules installation



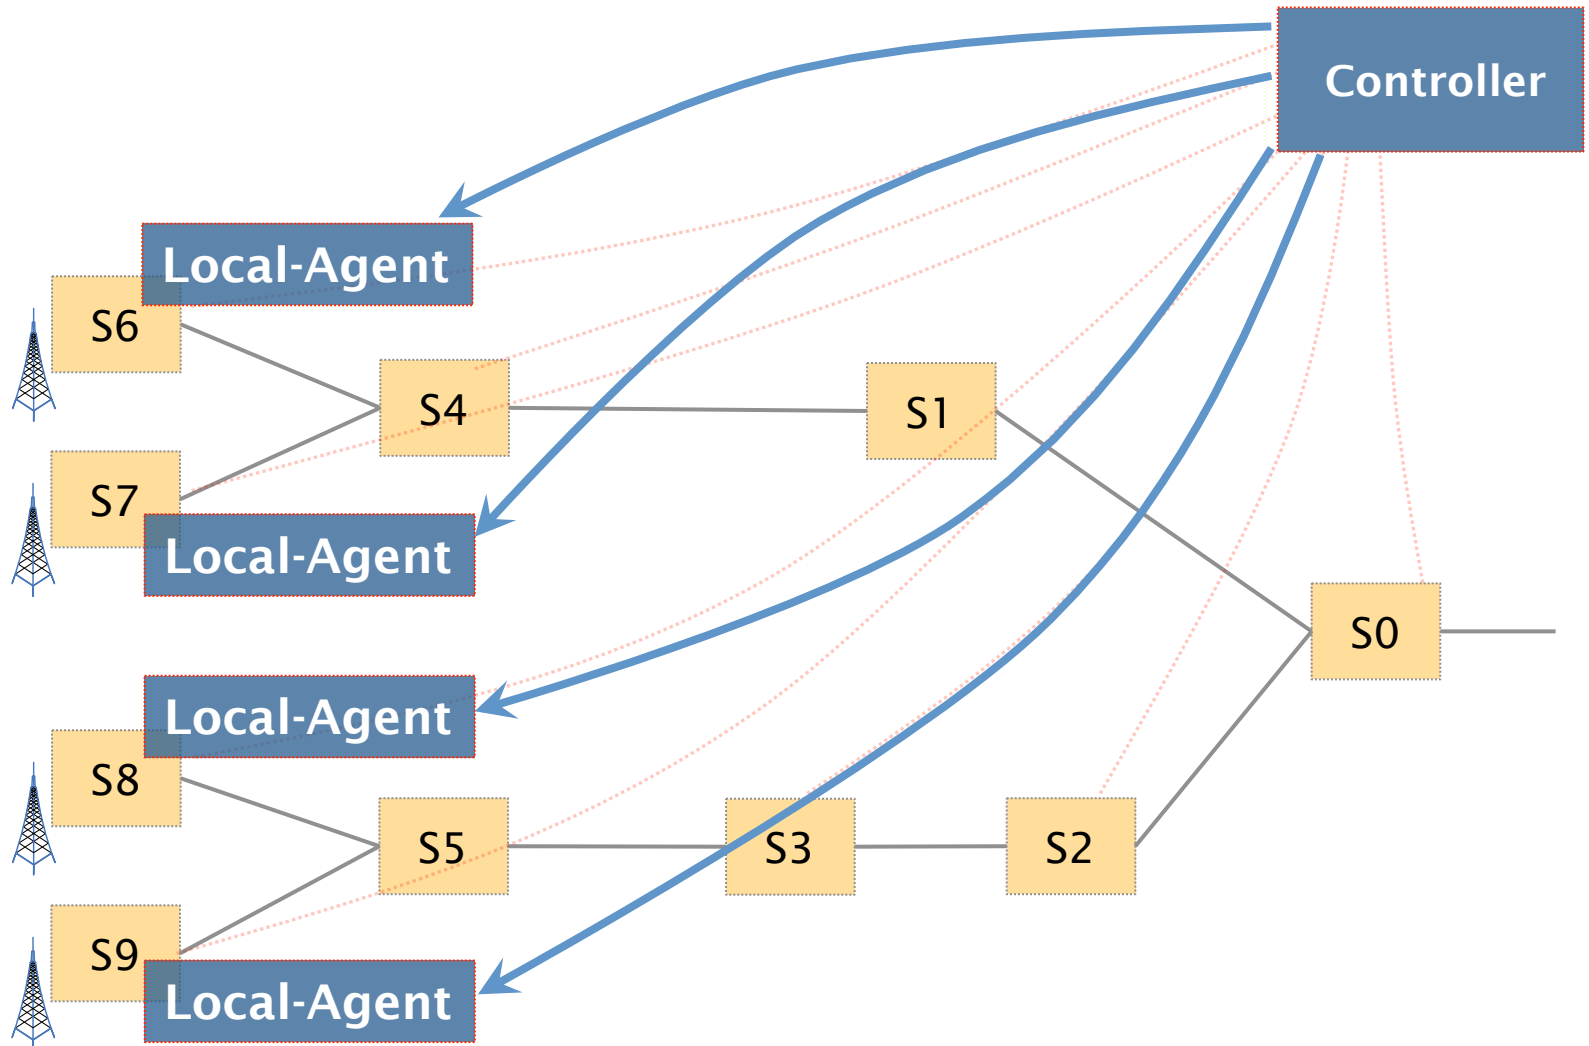




To scale, CellSDN uses a hierarchical controller



Most of the tasks are delegated to local-agents



Local-agents act as cache, reducing the load on the main controller

Local agents handle *locally* most frequent events

- cache a list of packet classifiers
- contact central controller upon cache miss
- tag flows

The central controller deals with less frequent, but more complex events

Local agents handle *locally* most frequent events

- cache a list of packet classifiers
- contact central controller upon cache miss
- tag flows

Central controller *globally* handle less frequent events

- UE arrival, handoff
- topology changes
- dynamic policies

CellSDN control-plane can handle the load of large cellular networks


Dataset

- 1 week of traces from a large LTE network
- 1500 base stations
- 1 million users


Evaluation

- # of events per second
- # of active users per second

The number of events going to the main controller is small



The number of events going to the main controller is small



99th percentile


214/s

280/s



Total \ll 1000

The number of events going to the main controller is small



Today's controllers can process tens of thousands events/s

99th percentile

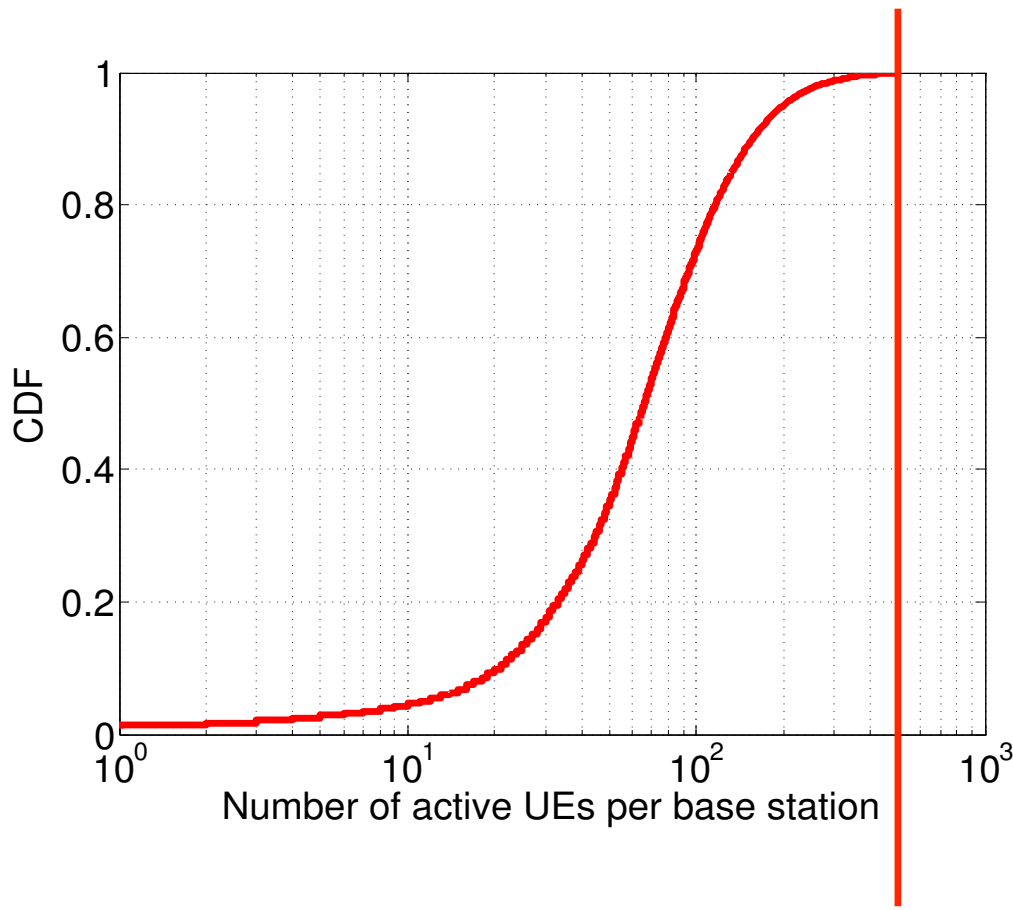
214/s

280/s



Total \ll 1000

The number of flows handled by the local-agent is small




99th percentile

514/s



~ 5k flows

The number of flows handled by the local-agent is small



Today's controllers can process tens of thousands events/s

99th percentile

514/s

→ ~ 5k flows

CellSDN: Taking control of cellular core networks



Architecture

software-defined network

Scaling the data-plane

multi-dimensional tagging

Scaling the control-plane

tasks delegation

CellSDN enables flexible and cost-effective cellular networks

CellSDN supports flexible **fine-grained** policies

CellSDN achieves **scalability** with

- Multi-dimensional aggregation
- Asymmetric edge design
- Hierarchical controller

CellSDN: Taking control of cellular core networks



Laurent Vanbever

www.vanbever.eu

IBM Watson Research Center

Tue 9 April 2013