

Observations of a Layer 2 Clos Fat-tree

IEEE 802.1-24-0014-00-ICne

Sai Prakash Chittampalli <sai.chittampalli@colorado.edu>

Ruchi Dhamnani <Ruchi.Dhamnani@colorado.edu>

Prateek Kumar <prku2113@colorado.edu>

Aditya Srivastava <Aditya.Srivastava@colorado.edu>

Weiqiang Wang <weiqiang.wang@colorado.edu>

EthAirNet Associates

University of Colorado at Boulder

14 March 2024

Related Contributions

- Data Center Collective Multicast using BARC-assigned Address Blocks

<https://ieee802.org/1/files/public/docs2024/cq-Marks-collective-multicast-0324-v00.pdf>

- Collective Communication in a Layer 2 Clos Fat-tree

IEEE 802.1-24-0012

https://mentor.ieee.org/802.1/documents?is_group=ICne&is_year=2024&is_dcn=0012

Implementation of Layer 2 Clos Fat-tree with Programmable Switches

IEEE 802.1-24-0013

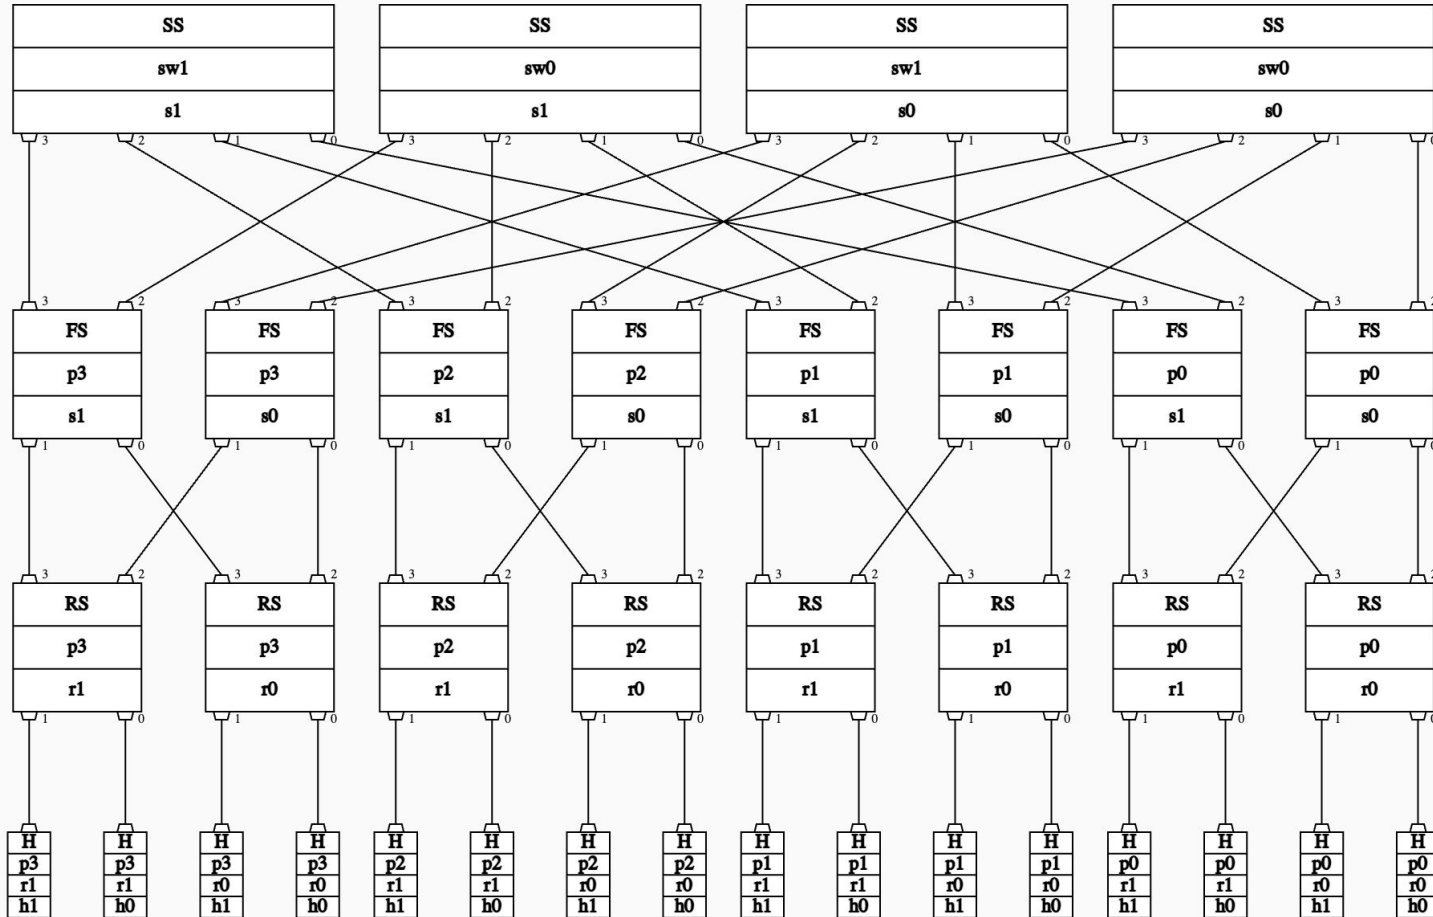
https://mentor.ieee.org/802.1/documents?is_group=ICne&is_year=2024&is_dcn=0013

Introduction

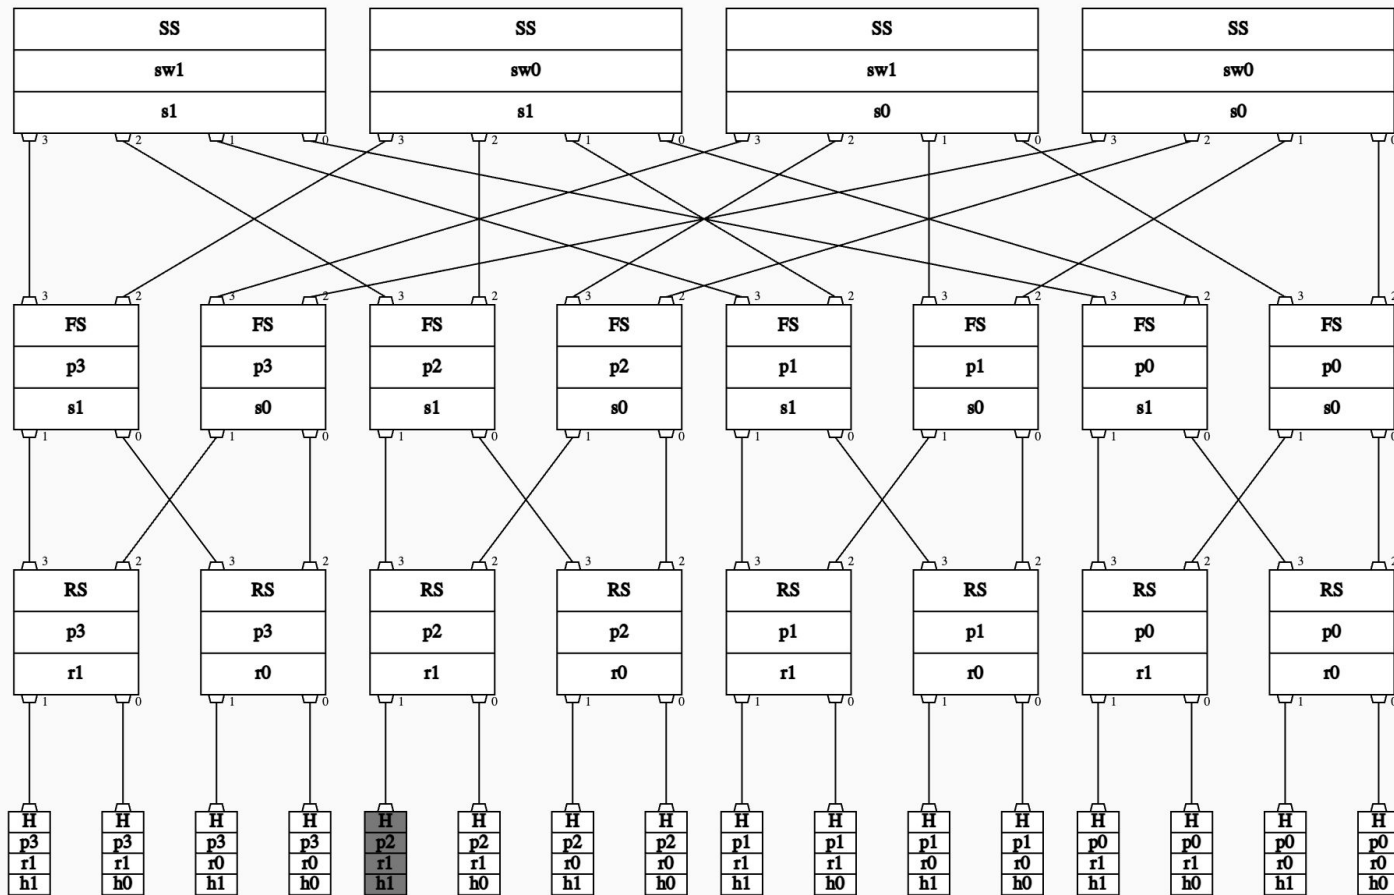
- The implementation of methods of communication in a Clos Fat-tree computing networks were described in “Implementation of Layer 2 Clos Fat-tree with Programmable Switches”
- This contribution presents simple results of that implementation, illustrating frame forwarding paths in the network.

BARC Frame Forwarding

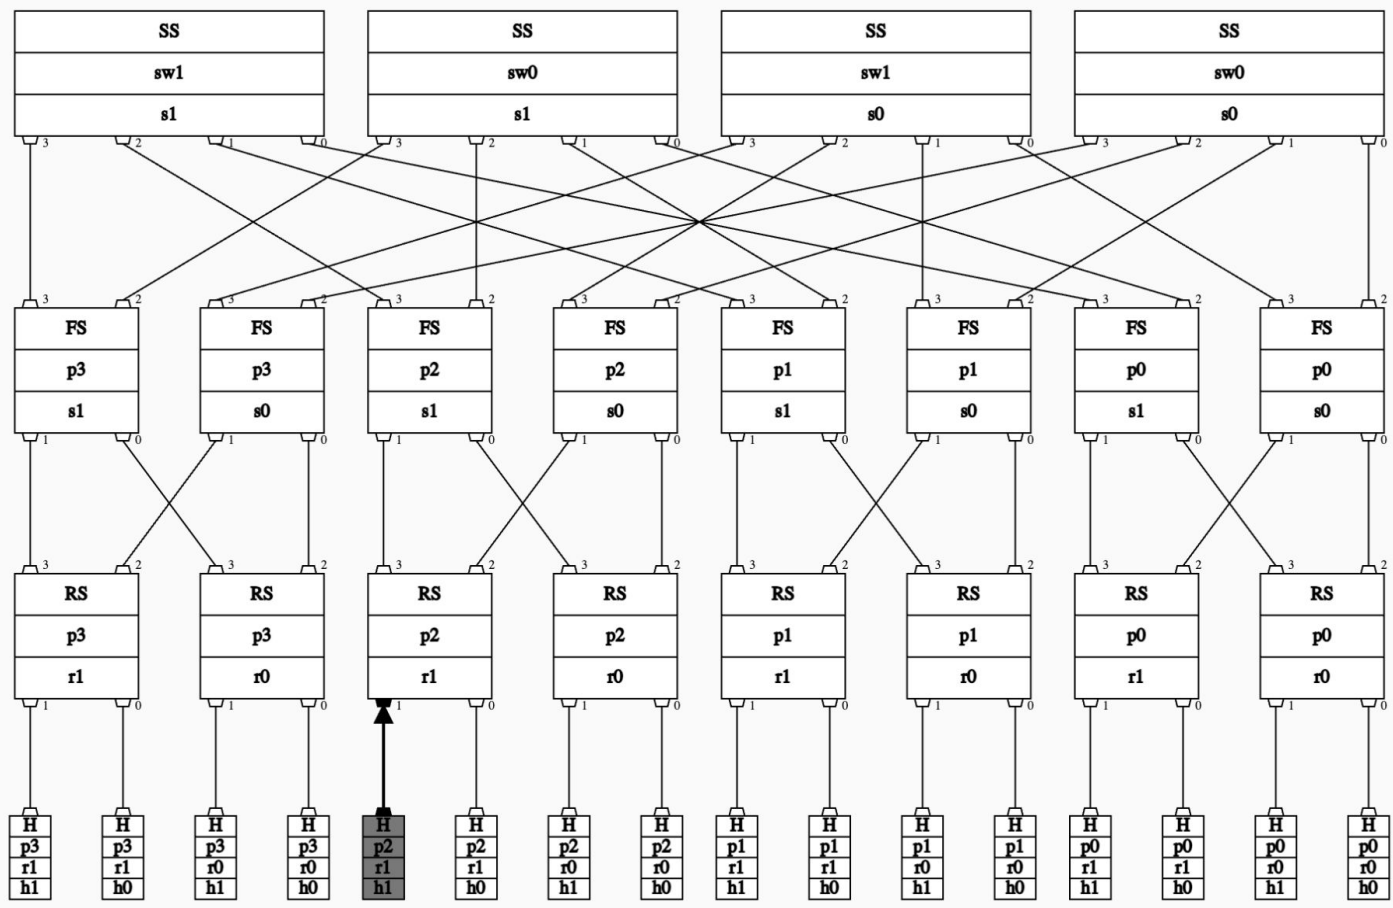
BARC FRAME FORWARDING



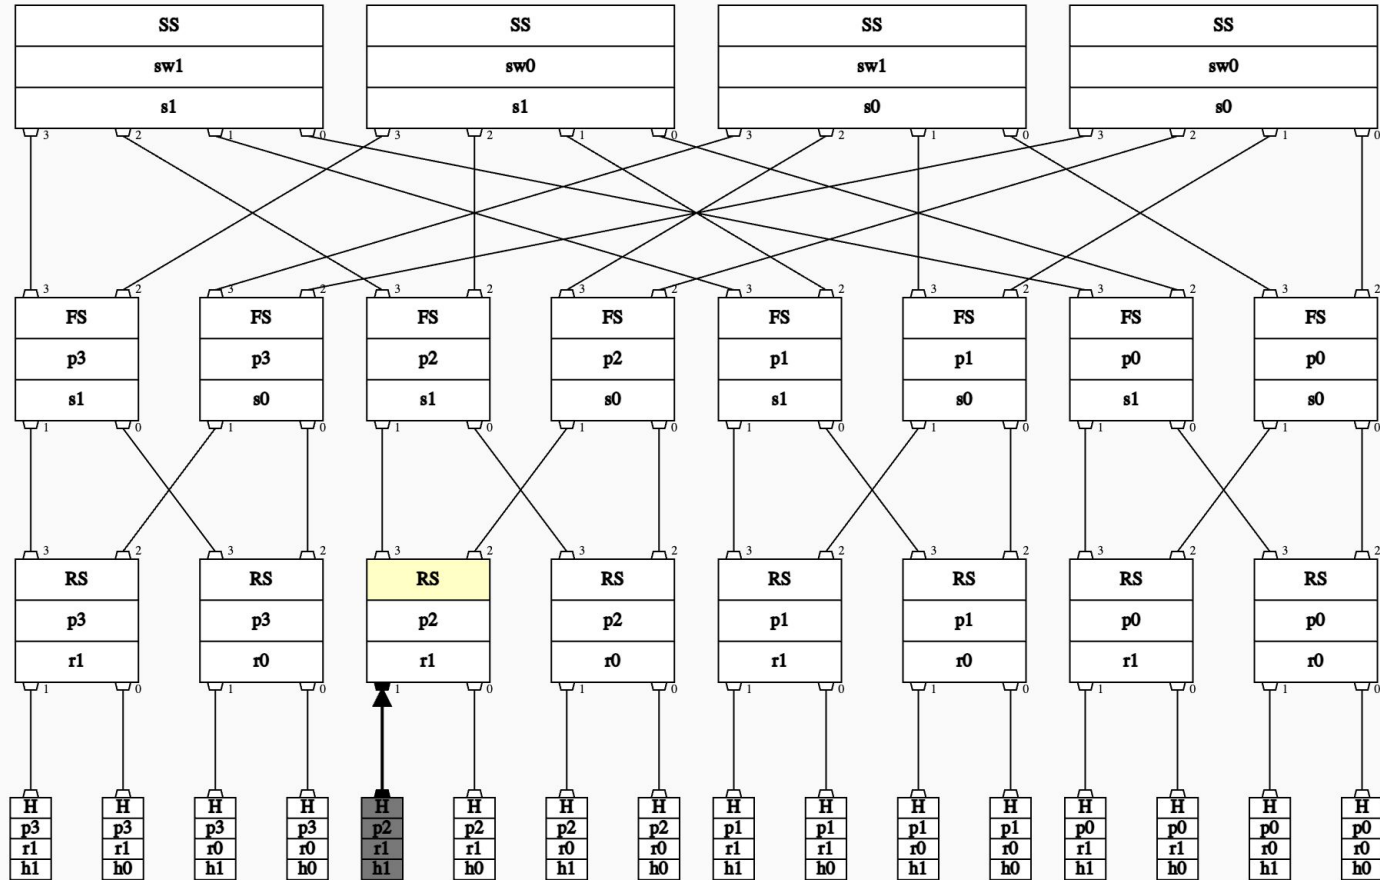
BARC FRAME FORWARDING



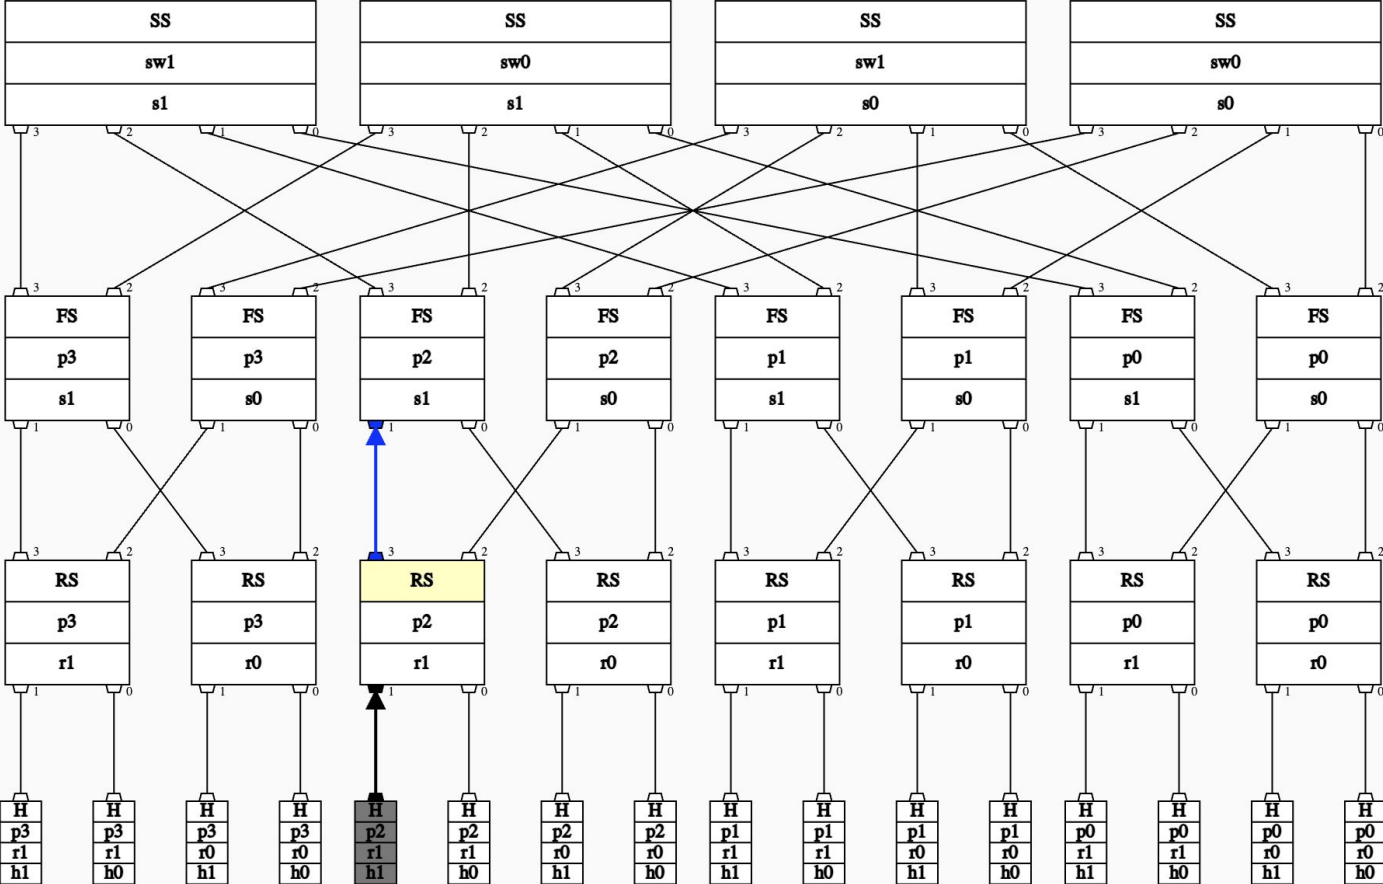
BARC FRAME FORWARDING



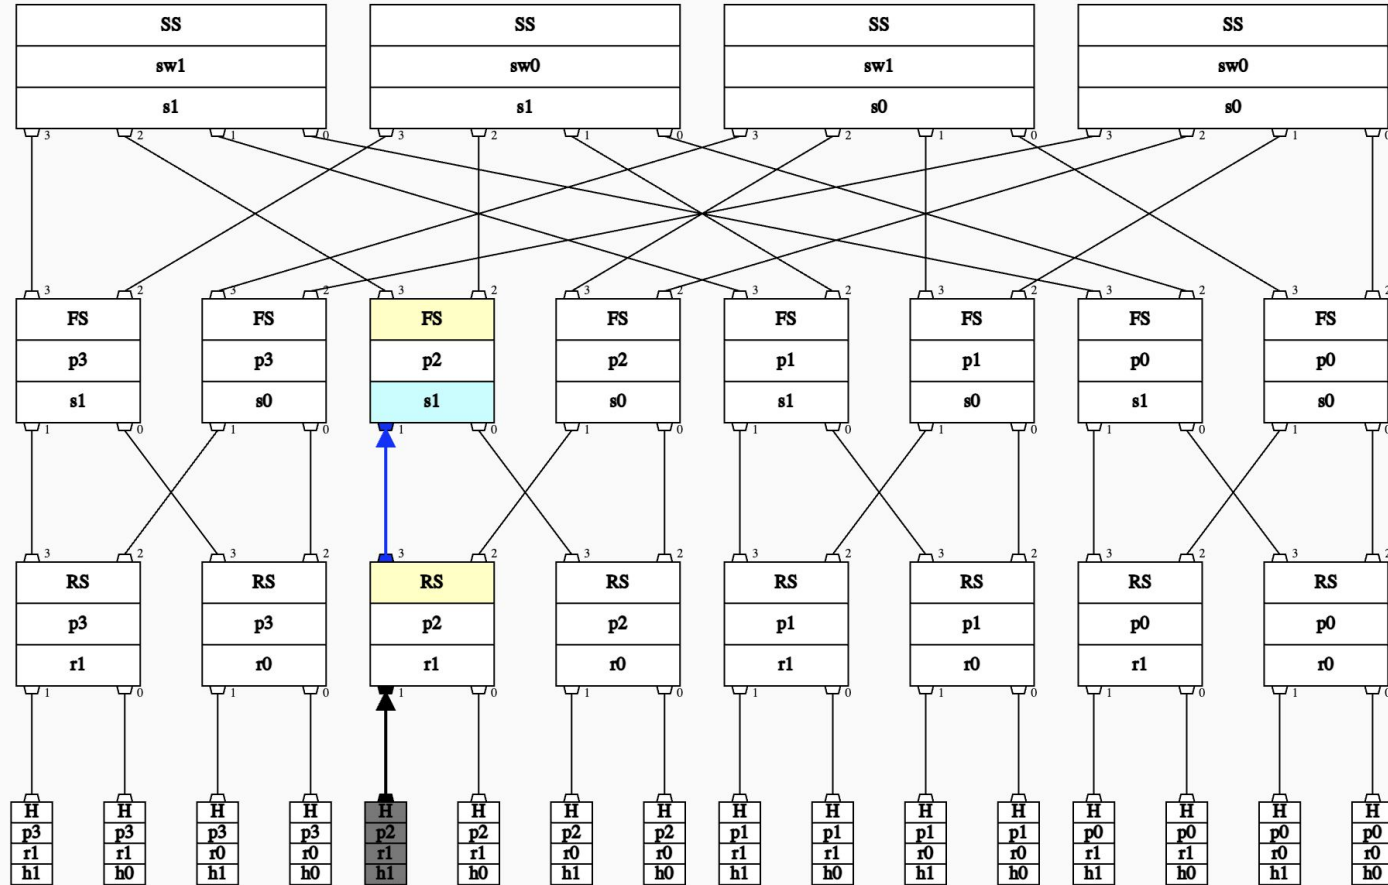
BARC FRAME FORWARDING



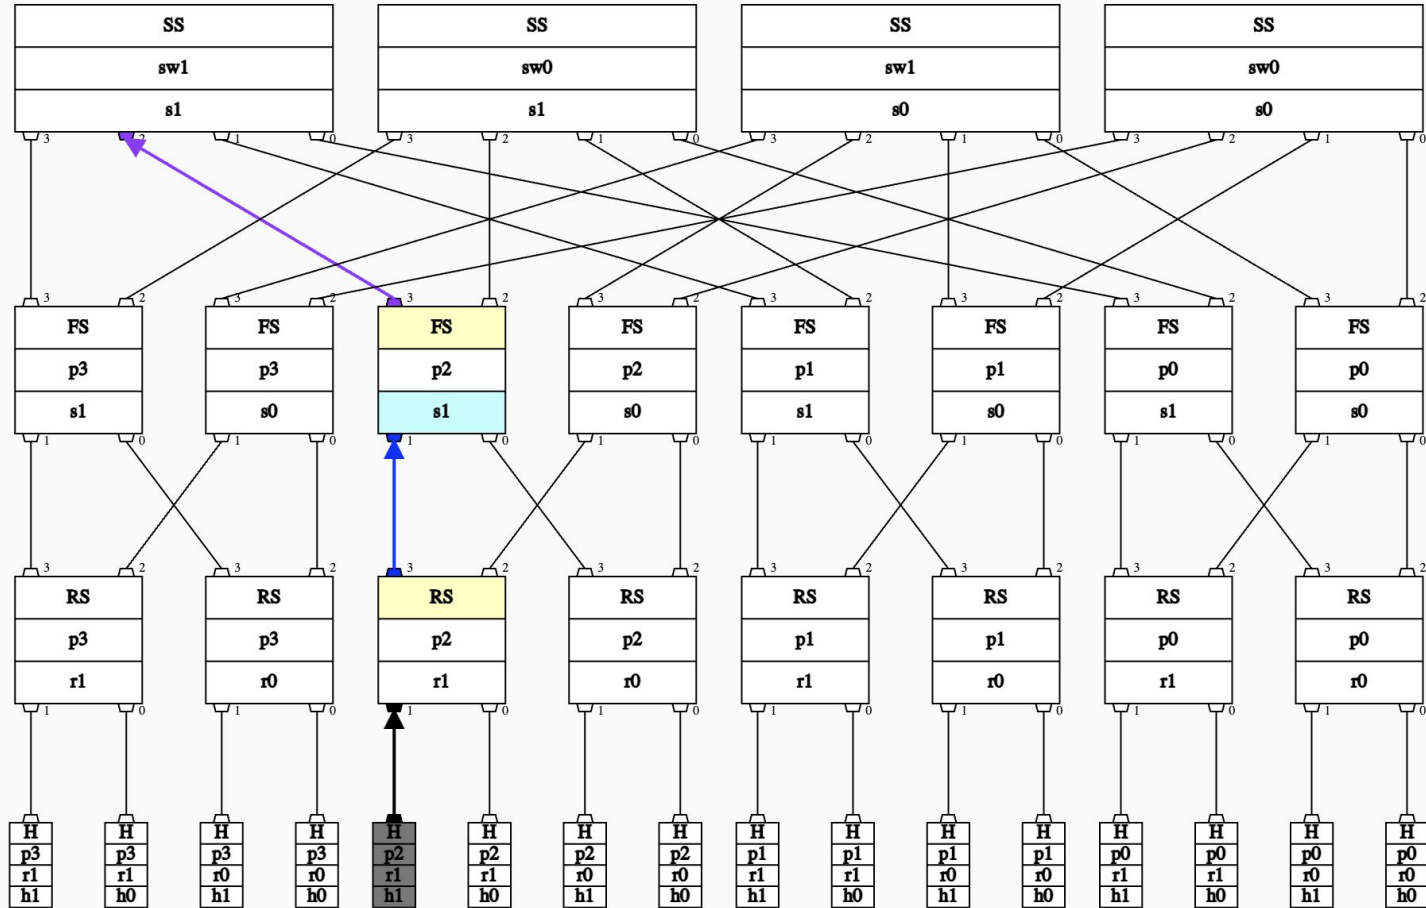
BARC FRAME FORWARDING



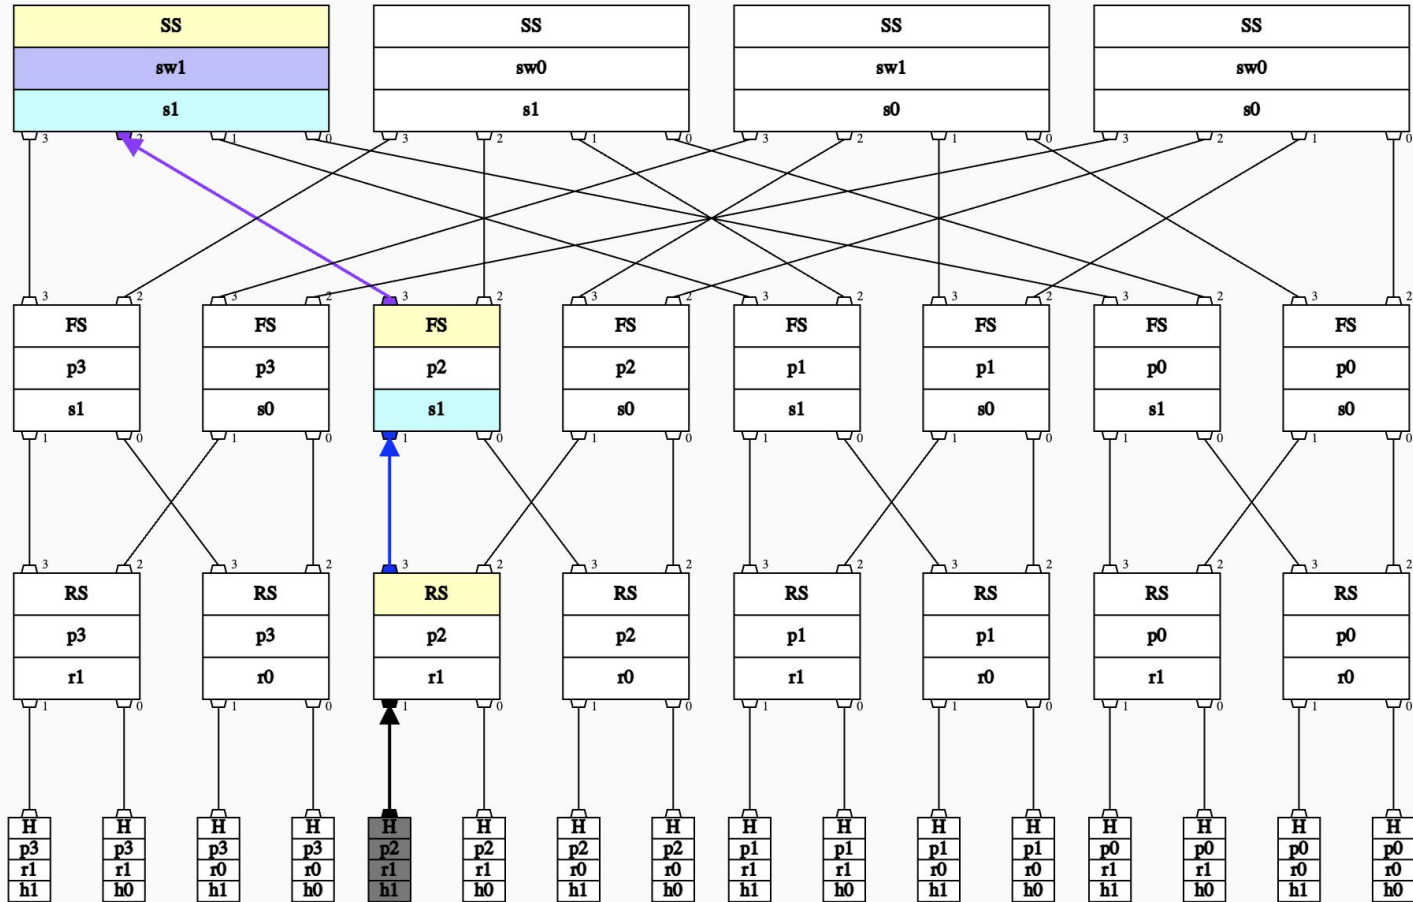
BARC FRAME FORWARDING



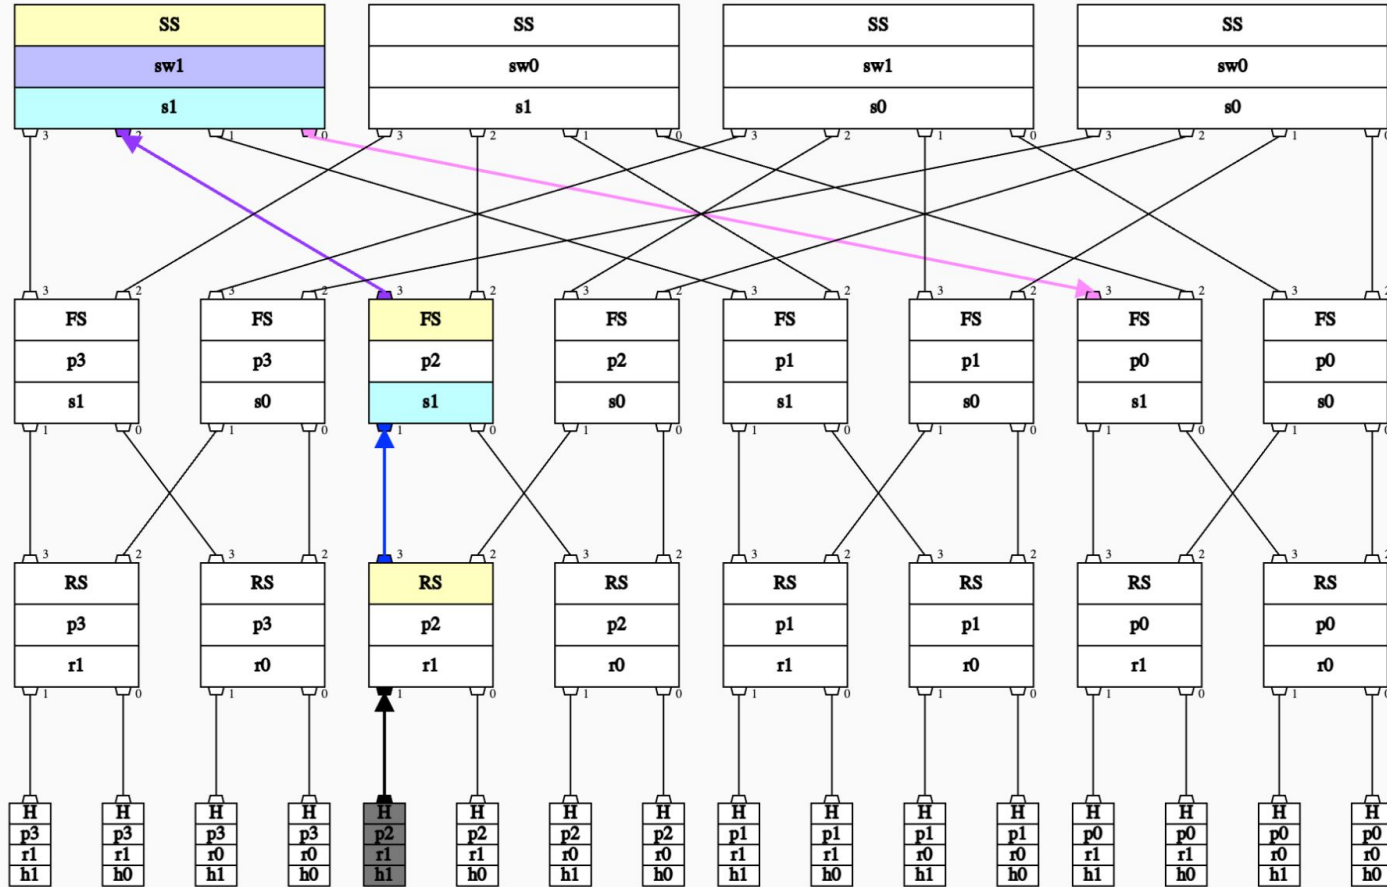
BARC FRAME FORWARDING



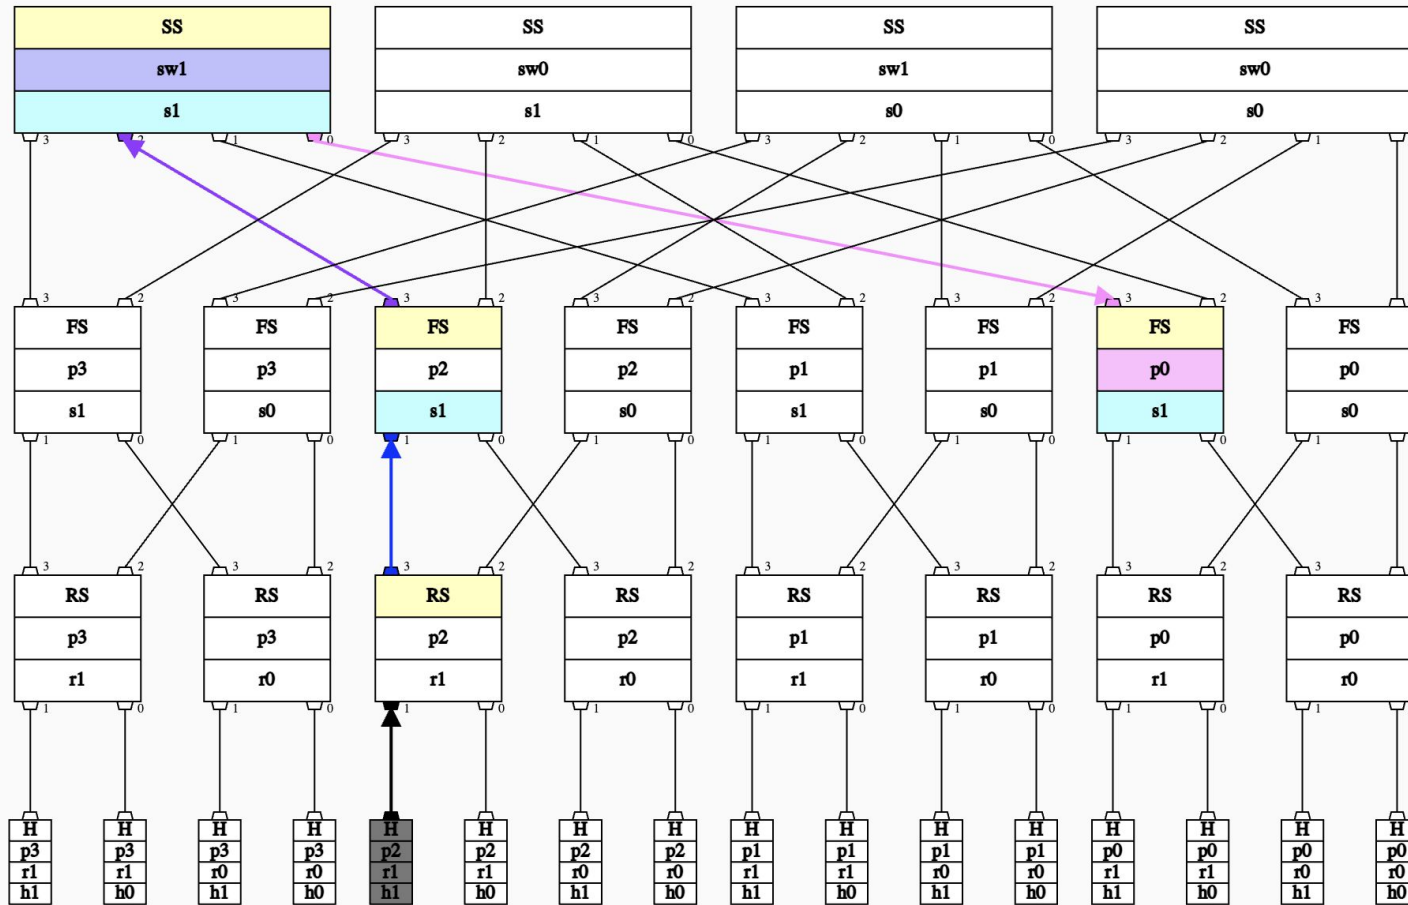
BARC FRAME FORWARDING



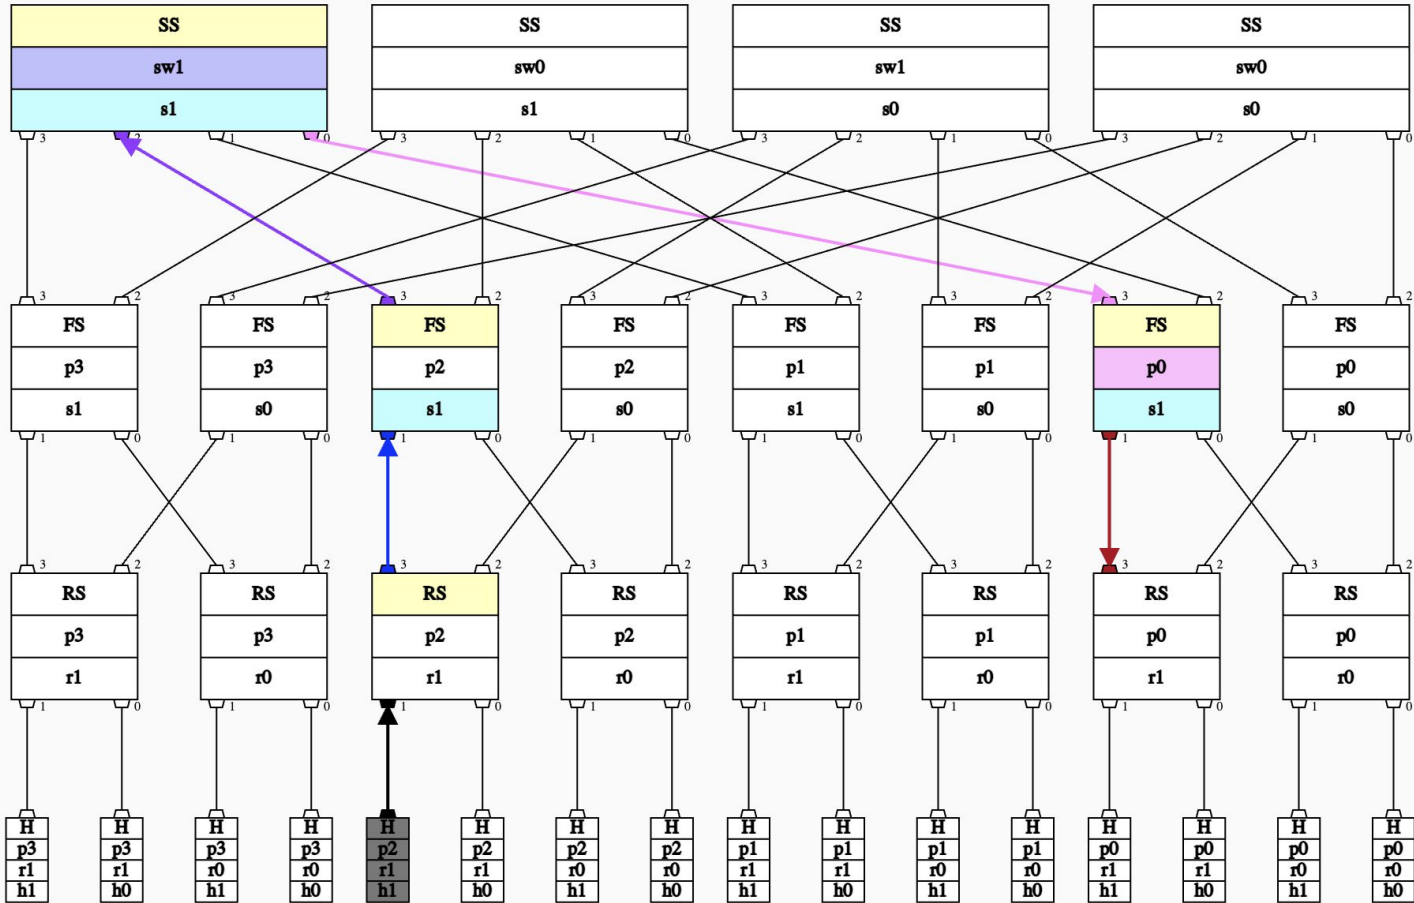
BARC FRAME FORWARDING



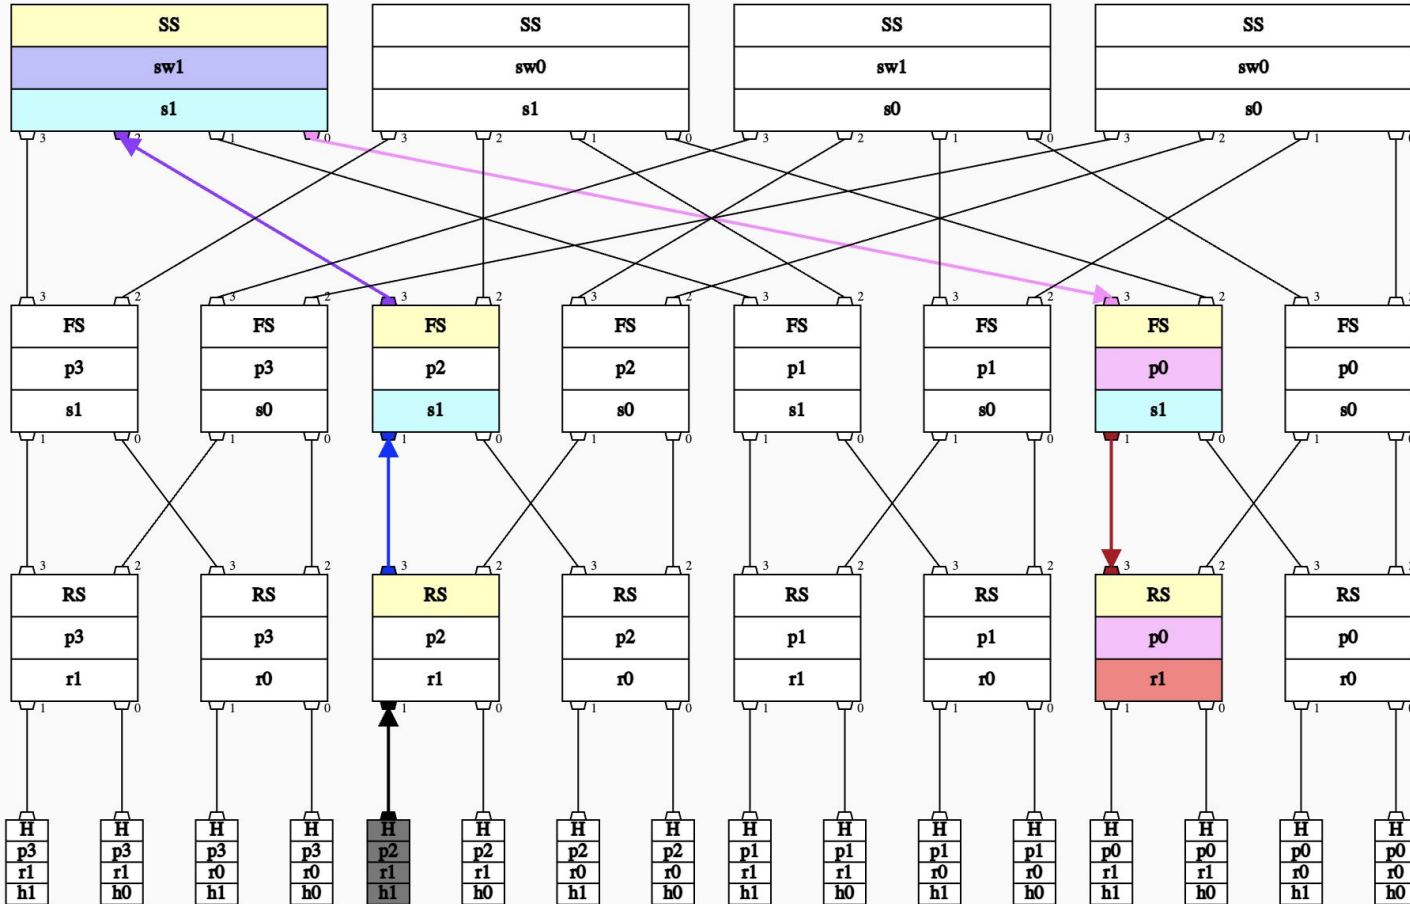
BARC FRAME FORWARDING



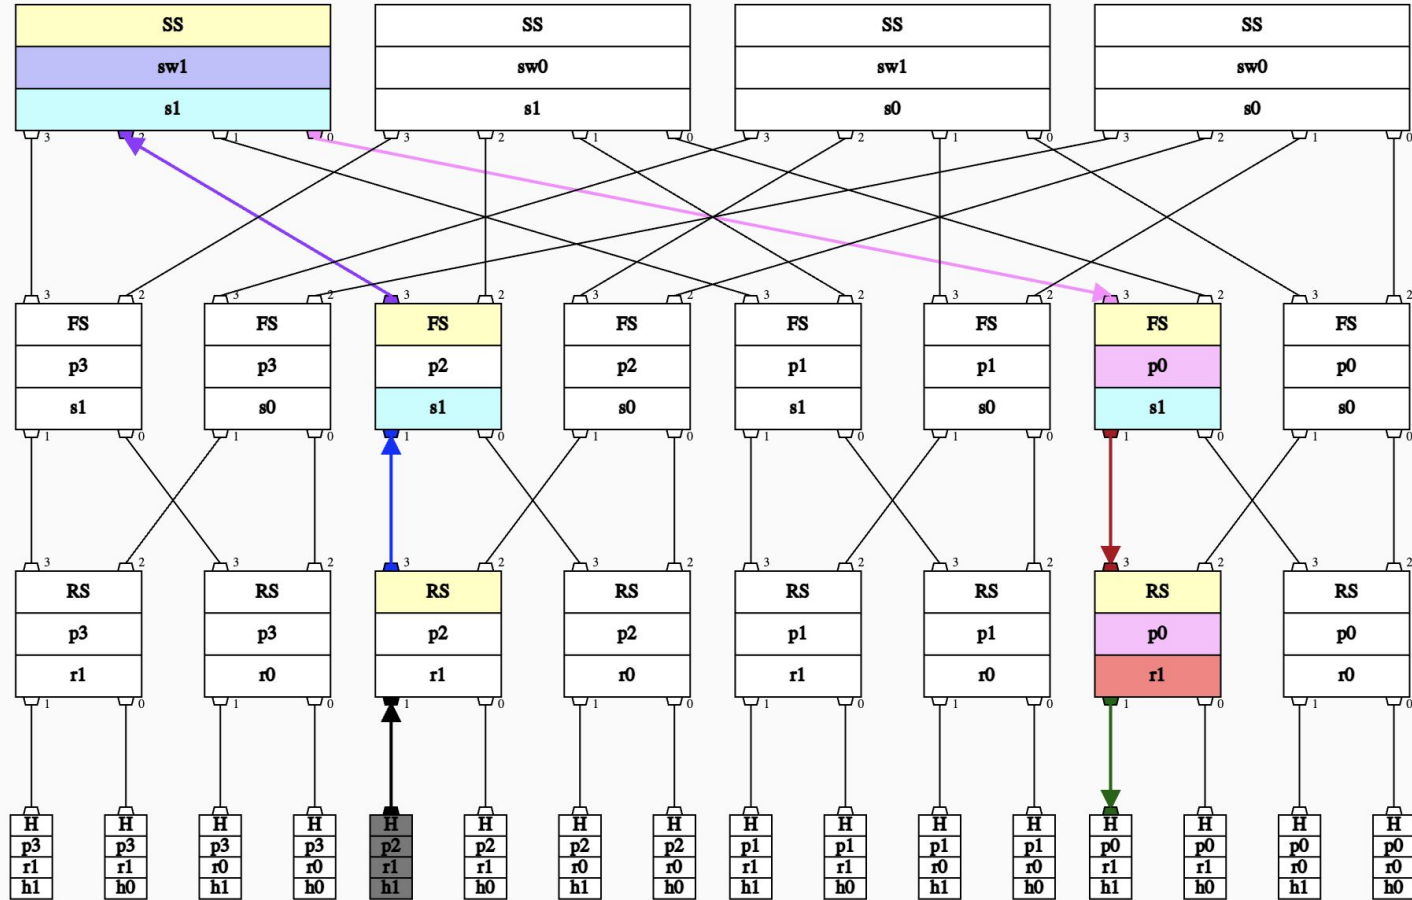
BARC FRAME FORWARDING



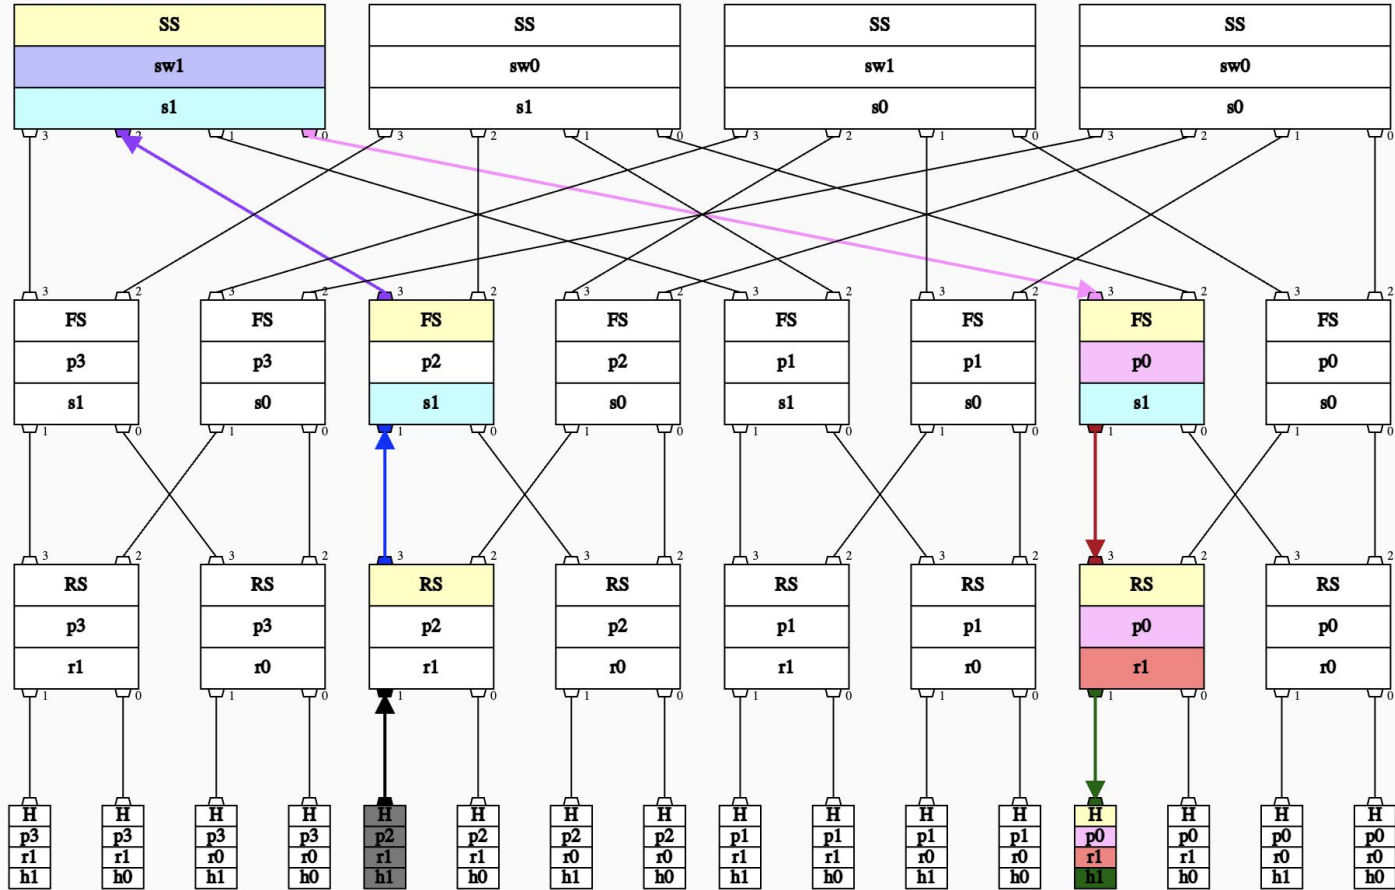
BARC FRAME FORWARDING



BARC FRAME FORWARDING

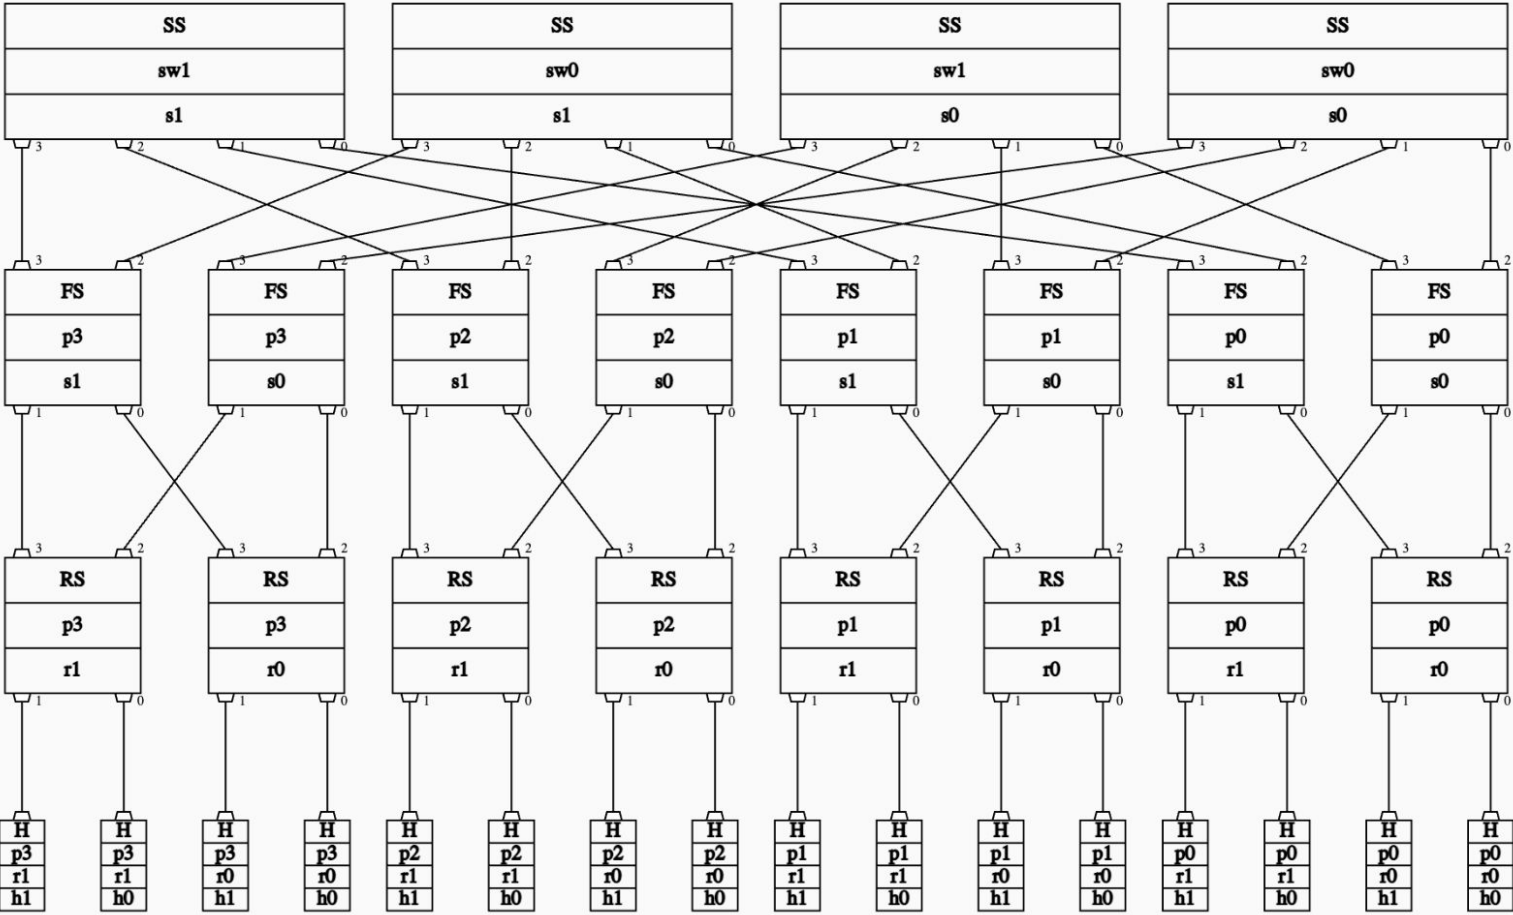


BARC FRAME FORWARDING



Unicast Frame Forwarding

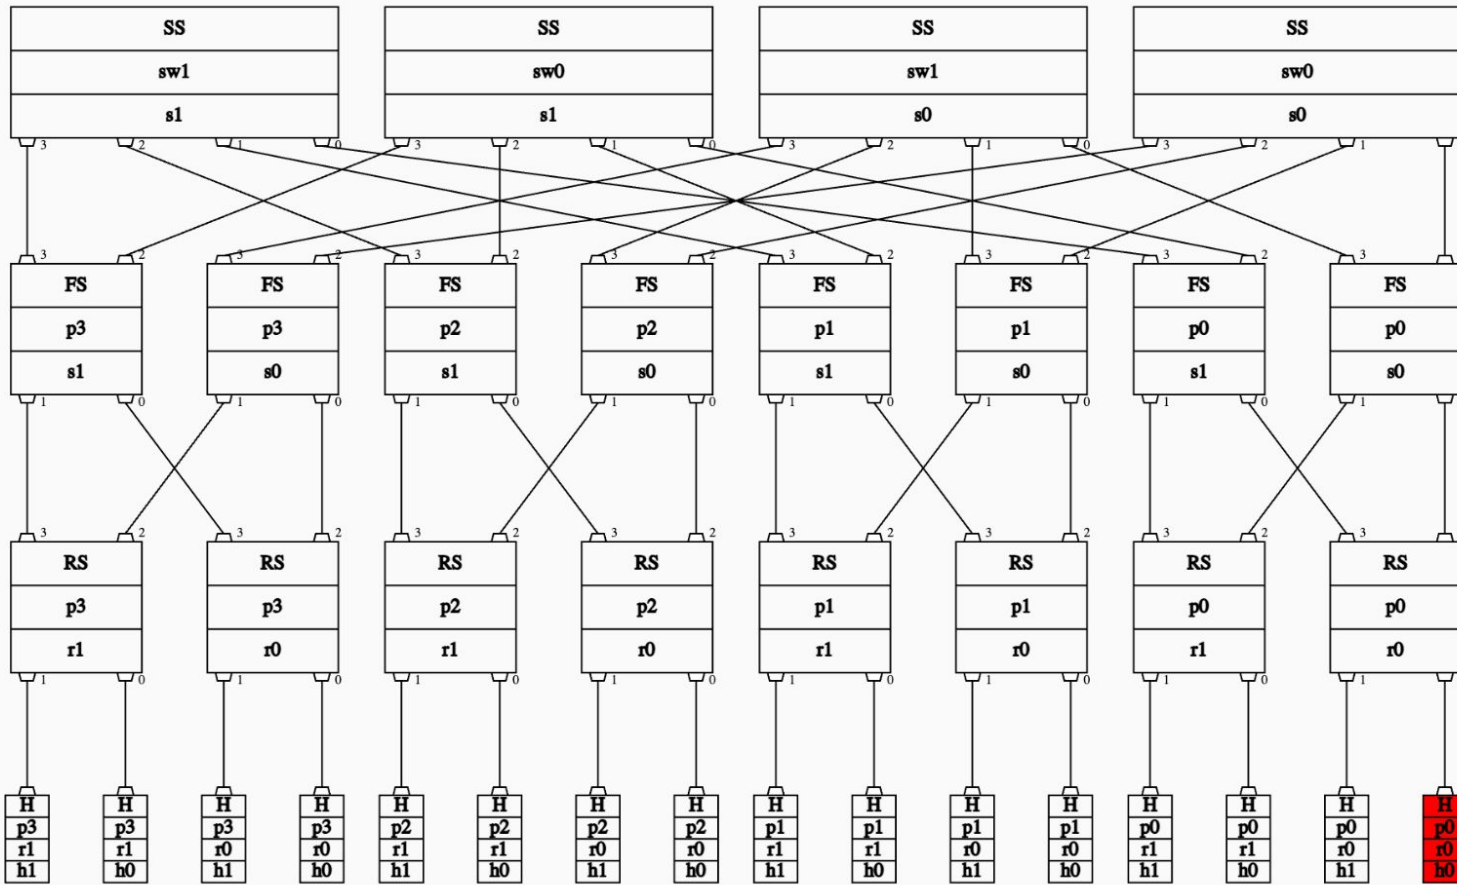
UNICAST FRAME FORWARDING



Destination Address

0xAE
Pod ID: 2
Rack ID: 1
Host ID: 0
Spine Switch ID: 0
Spine ID: 0

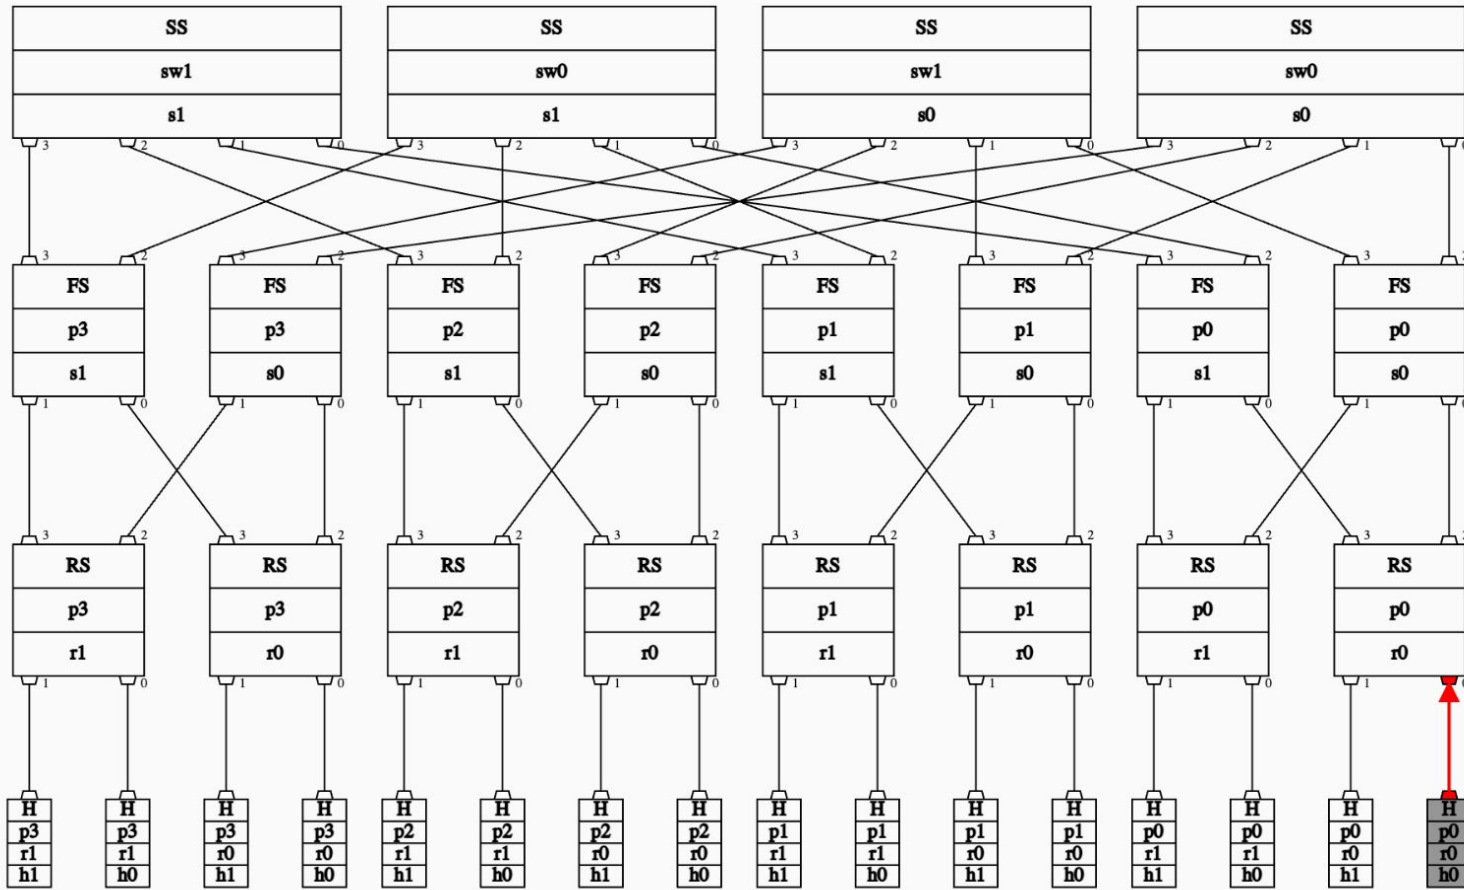
UNICAST FRAME FORWARDING



Destination Address

0xAE
Pod ID: 2
Rack ID: 1
Host ID: 0
Spine Switch ID: 0
Spine ID: 0

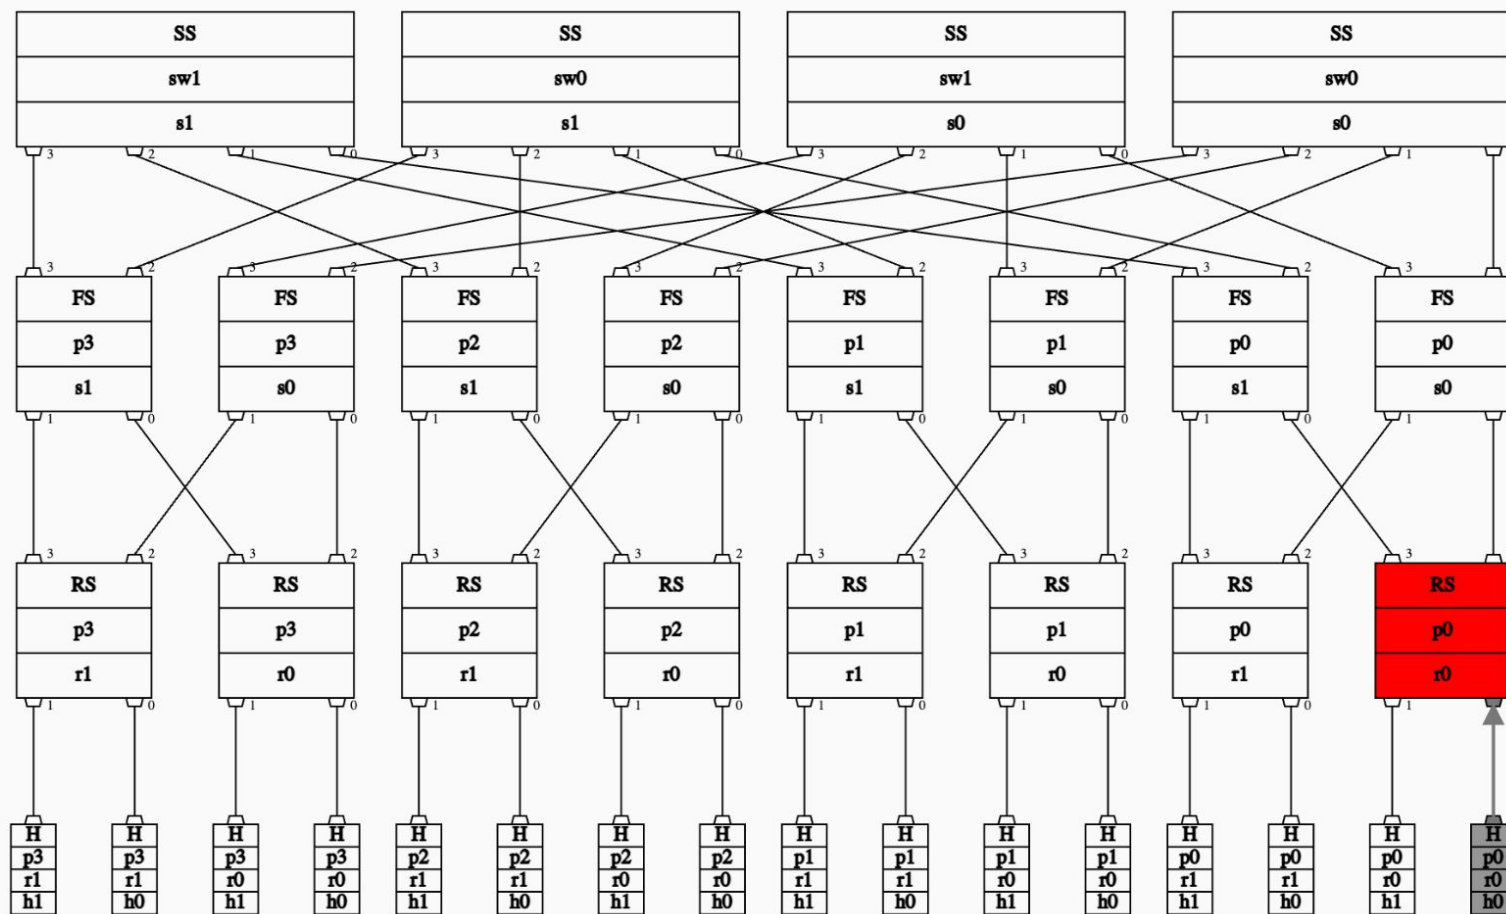
UNICAST FRAME FORWARDING



Destination Address

0xAE
Pod ID: 2
Rack ID: 1
Host ID: 0
Spine Switch ID: 0
Spine ID: 0

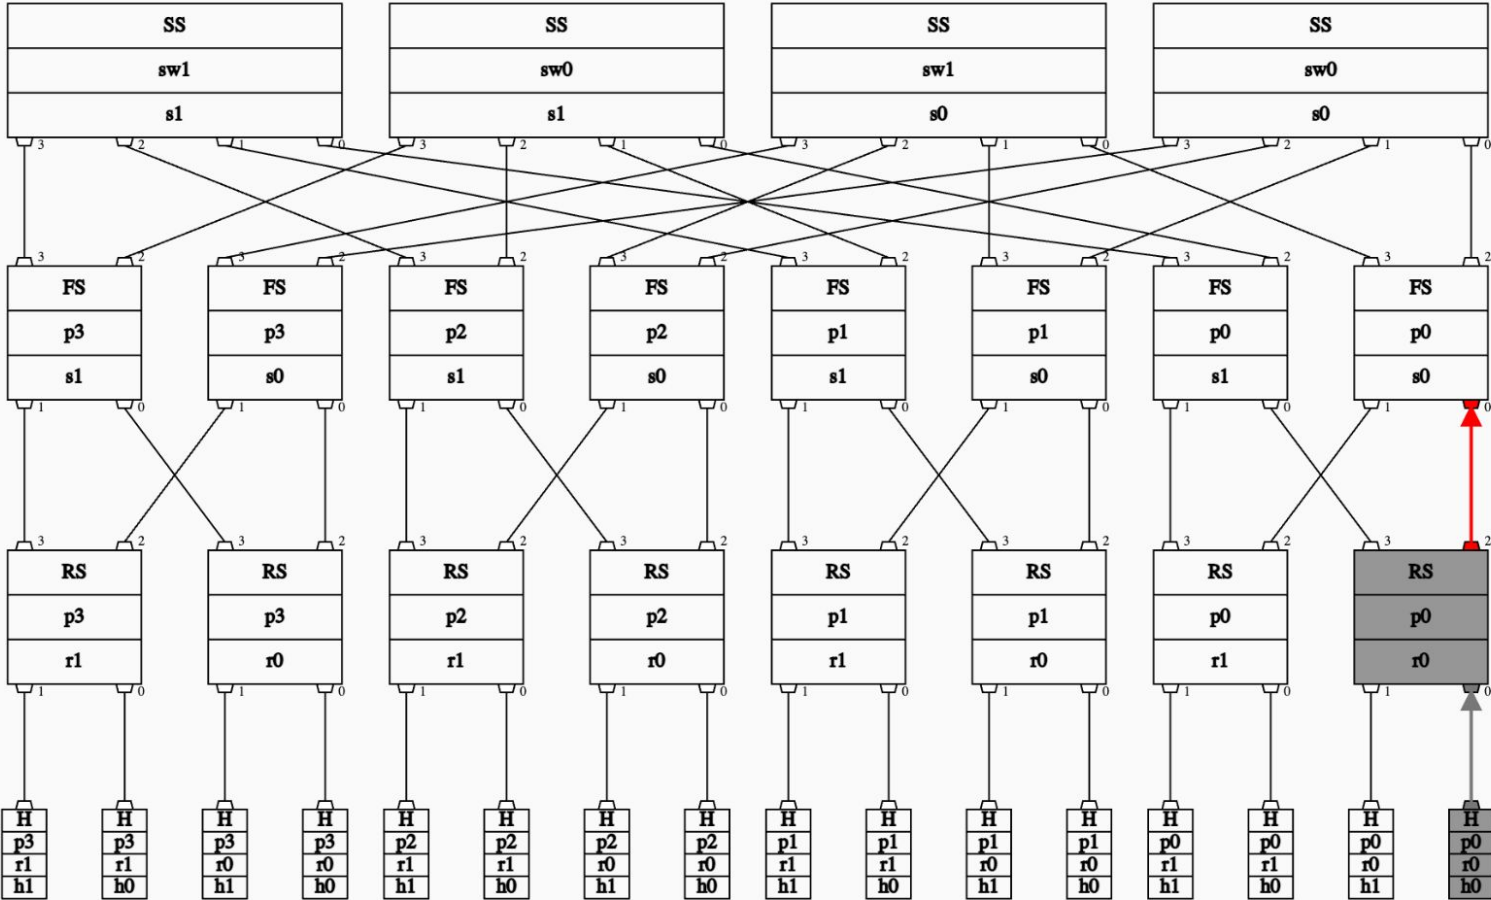
UNICAST FRAME FORWARDING



Destination Address

0xAE
Pod ID: 2
Rack ID: 1
Host ID: 0
Spine Switch ID: 0
Spine ID: 0

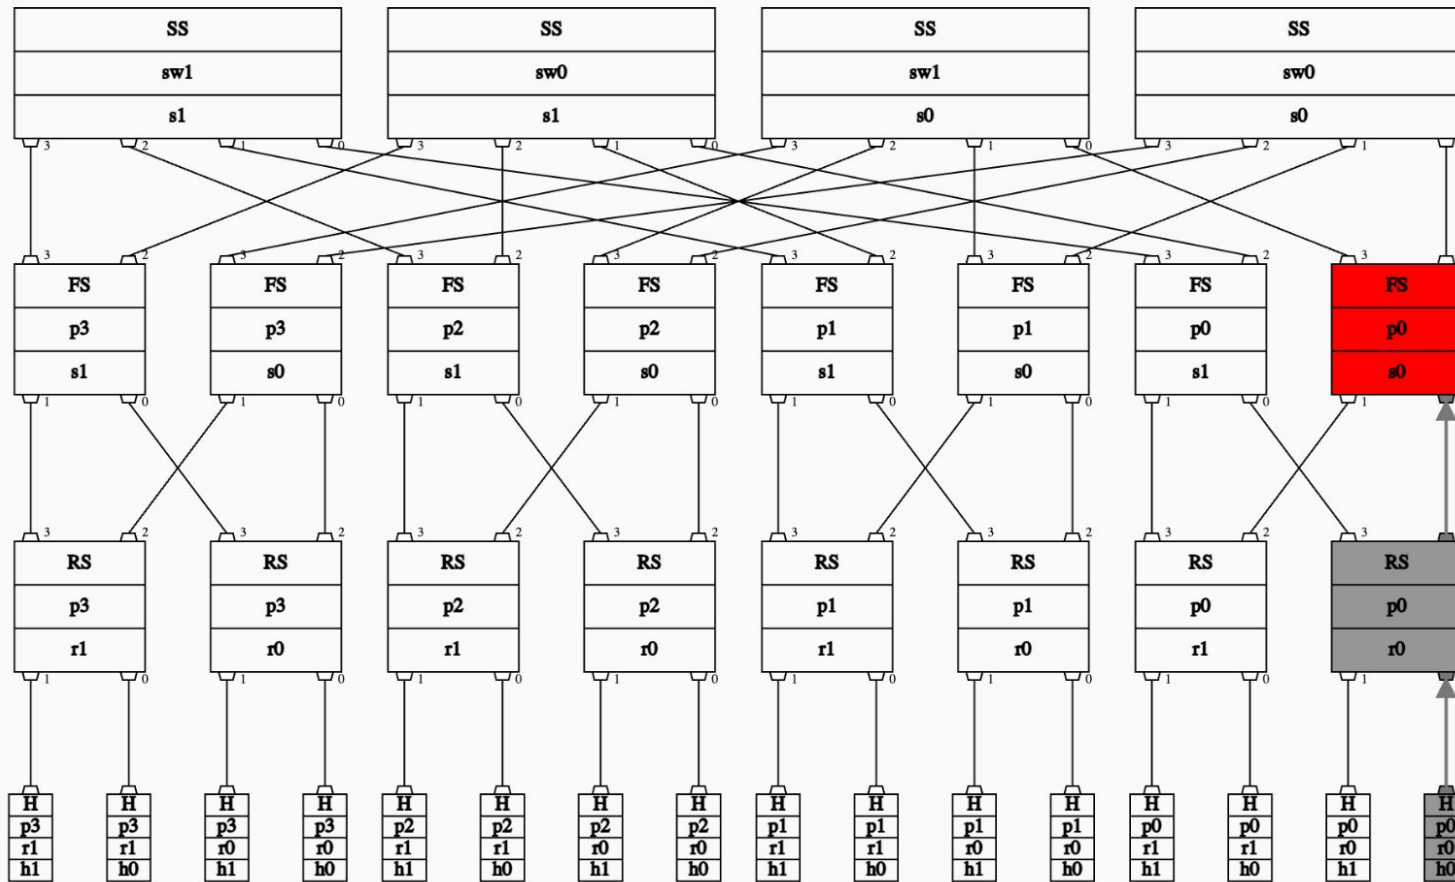
UNICAST FRAME FORWARDING



Destination Address

0xAE
Pod ID: 2
Rack ID: 1
Host ID: 0
Spine Switch ID: 0
Spine ID: 0

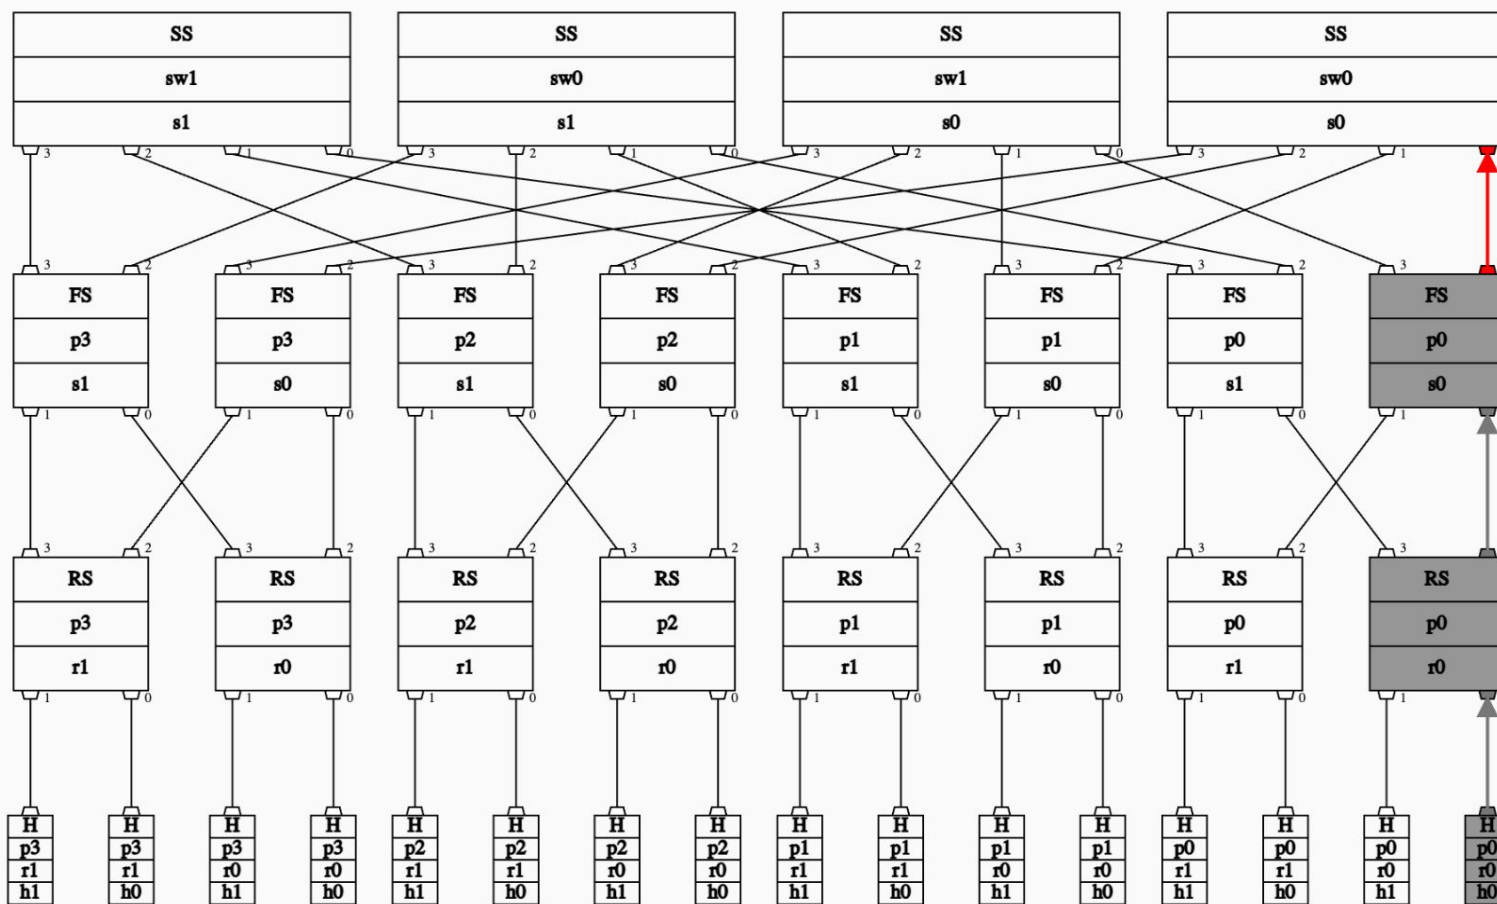
UNICAST FRAME FORWARDING



Destination Address

0xAE
Pod ID: 2
Rack ID: 1
Host ID: 0
Spine Switch ID: 0
Spine ID: 0

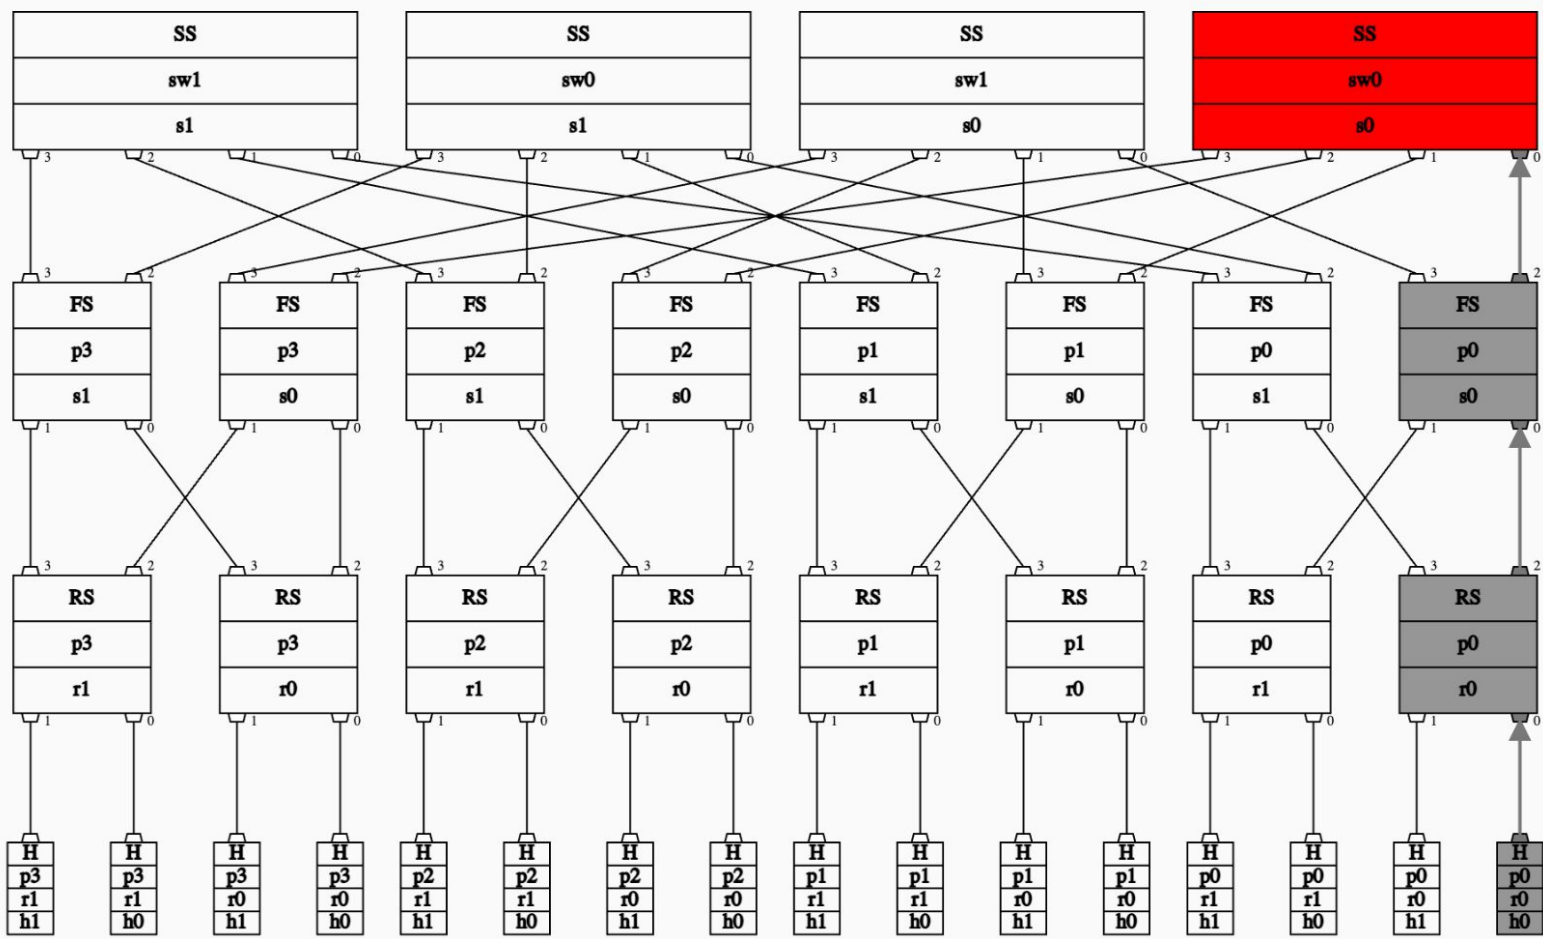
UNICAST FRAME FORWARDING



Destination Address

0xAE
Pod ID: 2
Rack ID: 1
Host ID: 0
Spine Switch ID: 0
Spine ID: 0

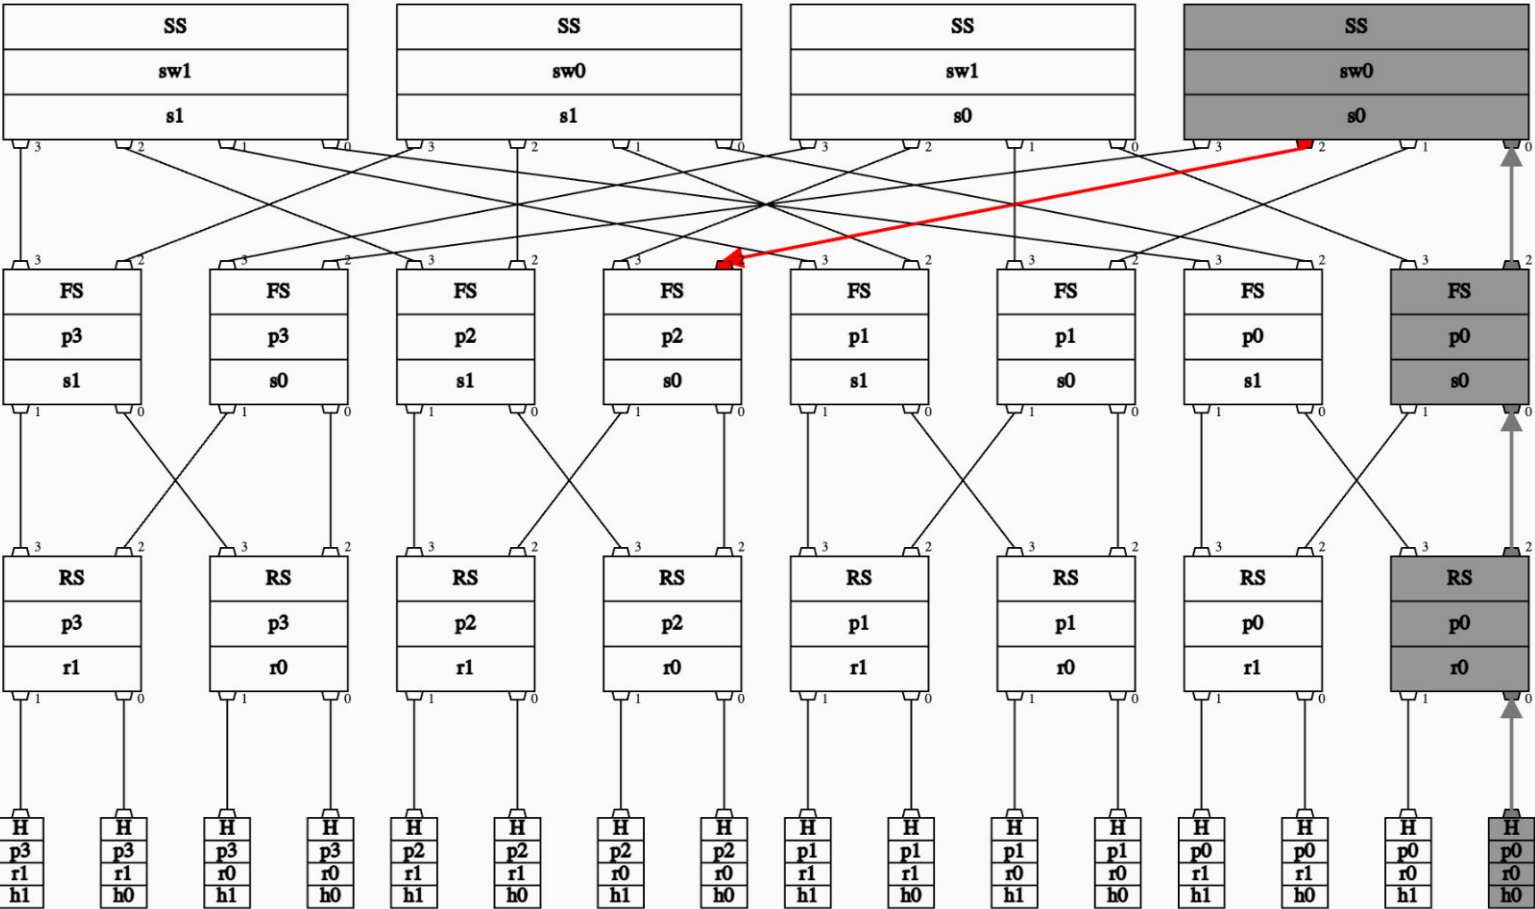
UNICAST FRAME FORWARDING



Destination Address

0xAE
Pod ID: 2
Rack ID: 1
Host ID: 0
Spine Switch ID: 0
Spine ID: 0

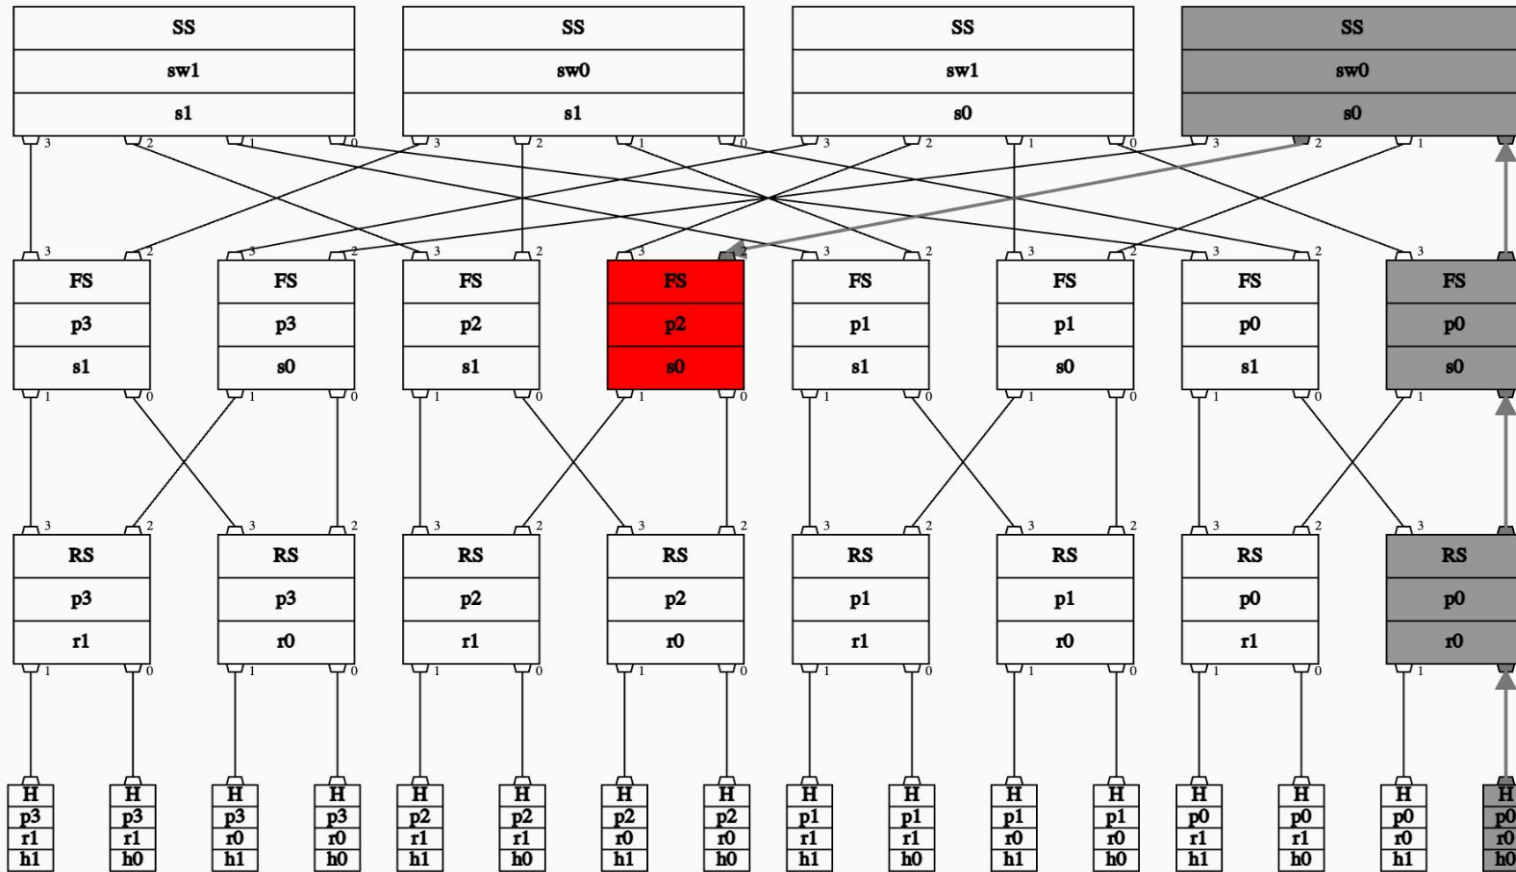
UNICAST FRAME FORWARDING



Destination Address

0xAE
Pod ID: 2
Rack ID: 1
Host ID: 0
Spine Switch ID: 0
Spine ID: 0

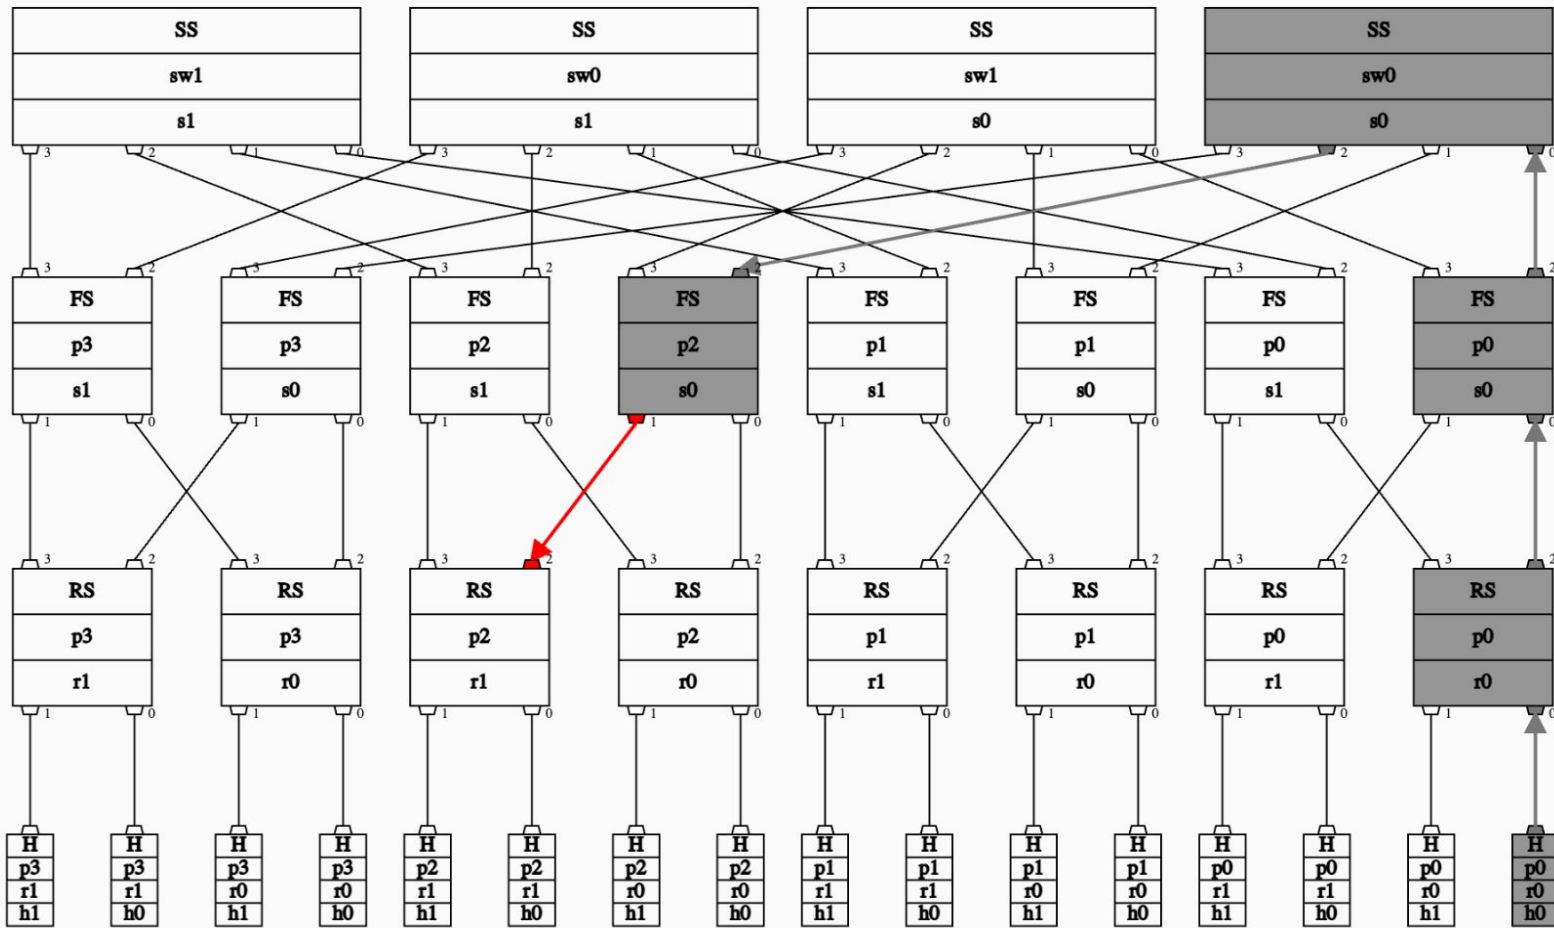
UNICAST FRAME FORWARDING



Destination Address

0xAE
Pod ID: 2
Rack ID: 1
Host ID: 0
Spine Switch ID: 0
Spine ID: 0

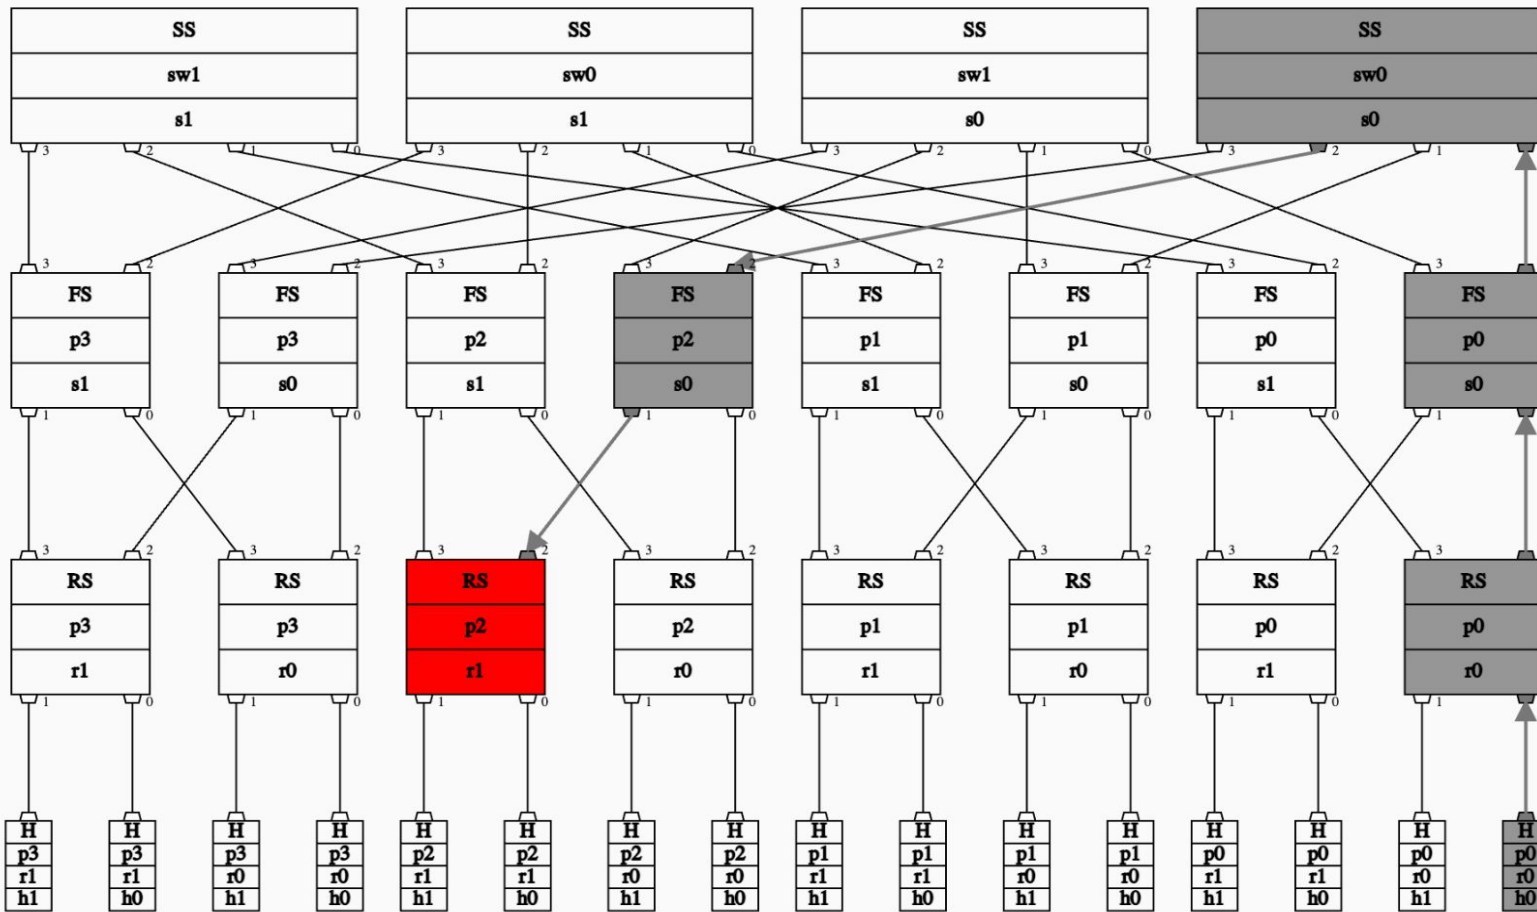
UNICAST FRAME FORWARDING



Destination Address

0xAE
Pod ID: 2
Rack ID: 1
Host ID: 0
Spine Switch ID: 0
Spine ID: 0

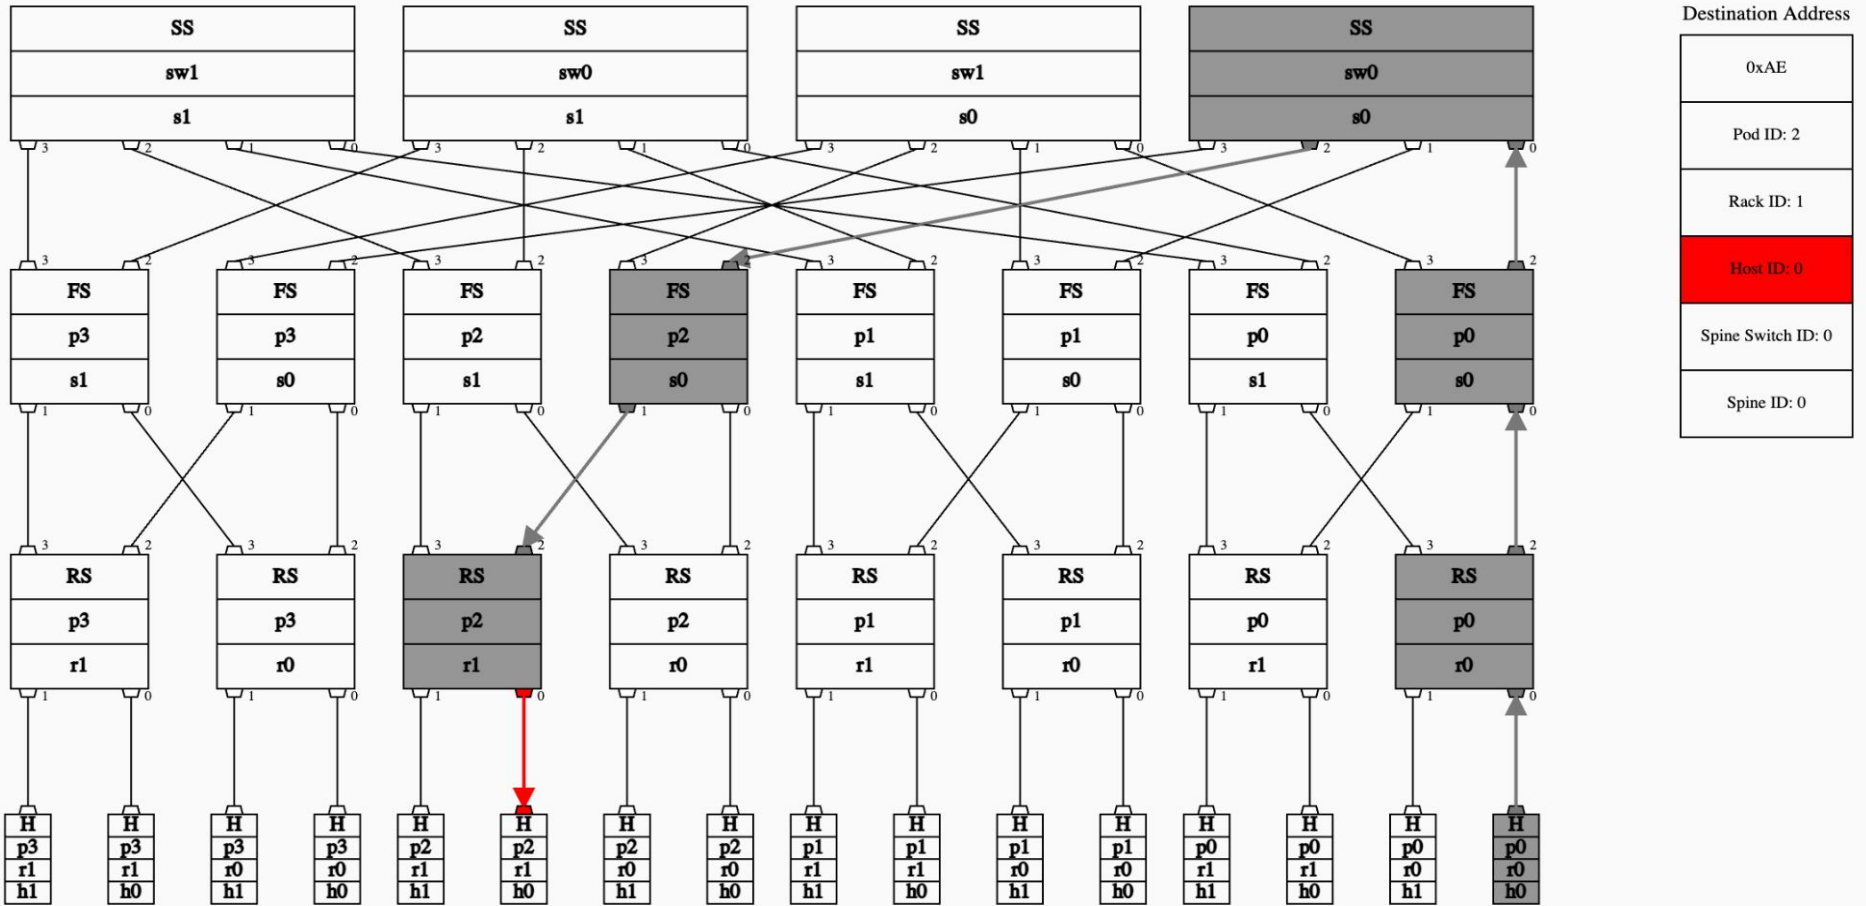
UNICAST FRAME FORWARDING



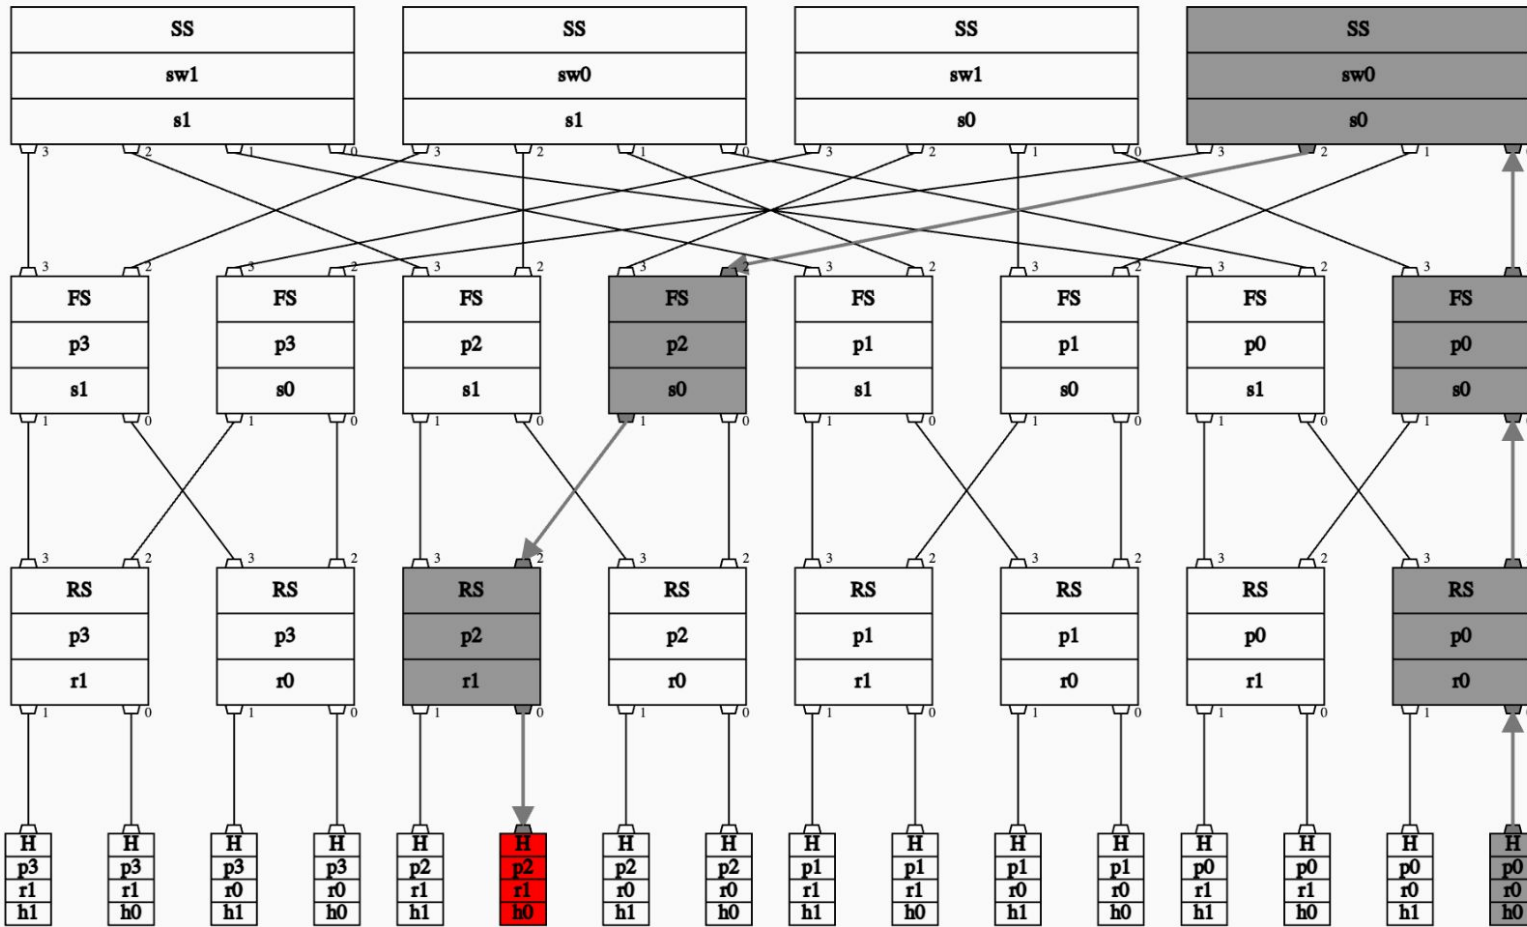
Destination Address

0xAE
Pod ID: 2
Rack ID: 1
Host ID: 0
Spine Switch ID: 0
Spine ID: 0

UNICAST FRAME FORWARDING



UNICAST FRAME FORWARDING

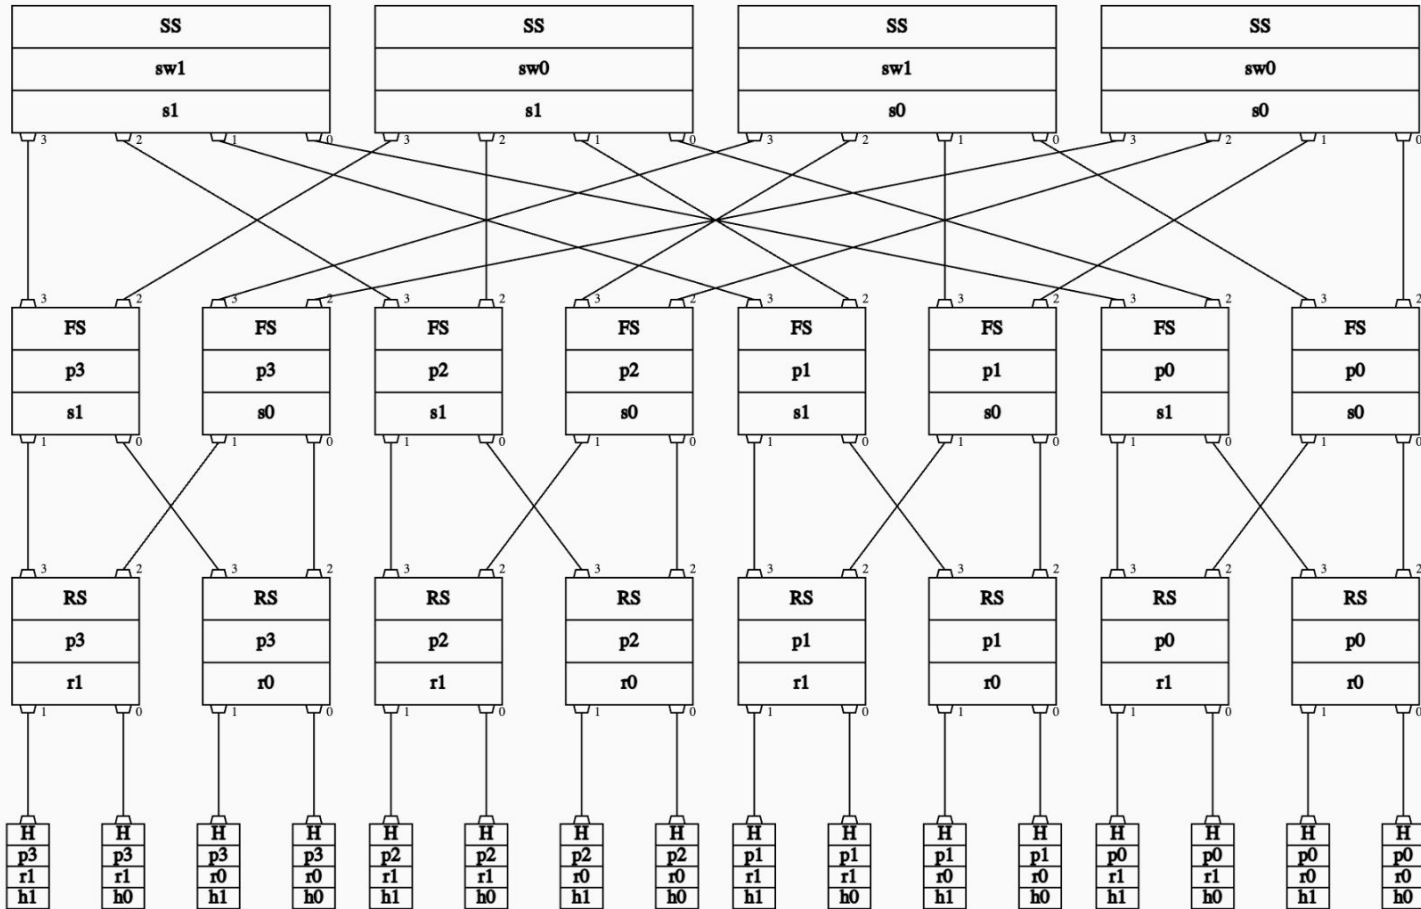


Destination Address

0xAE
Pod ID: 2
Rack ID: 1
Host ID: 0
Spine Switch ID: 0
Spine ID: 0

Multicast Join Forwarding

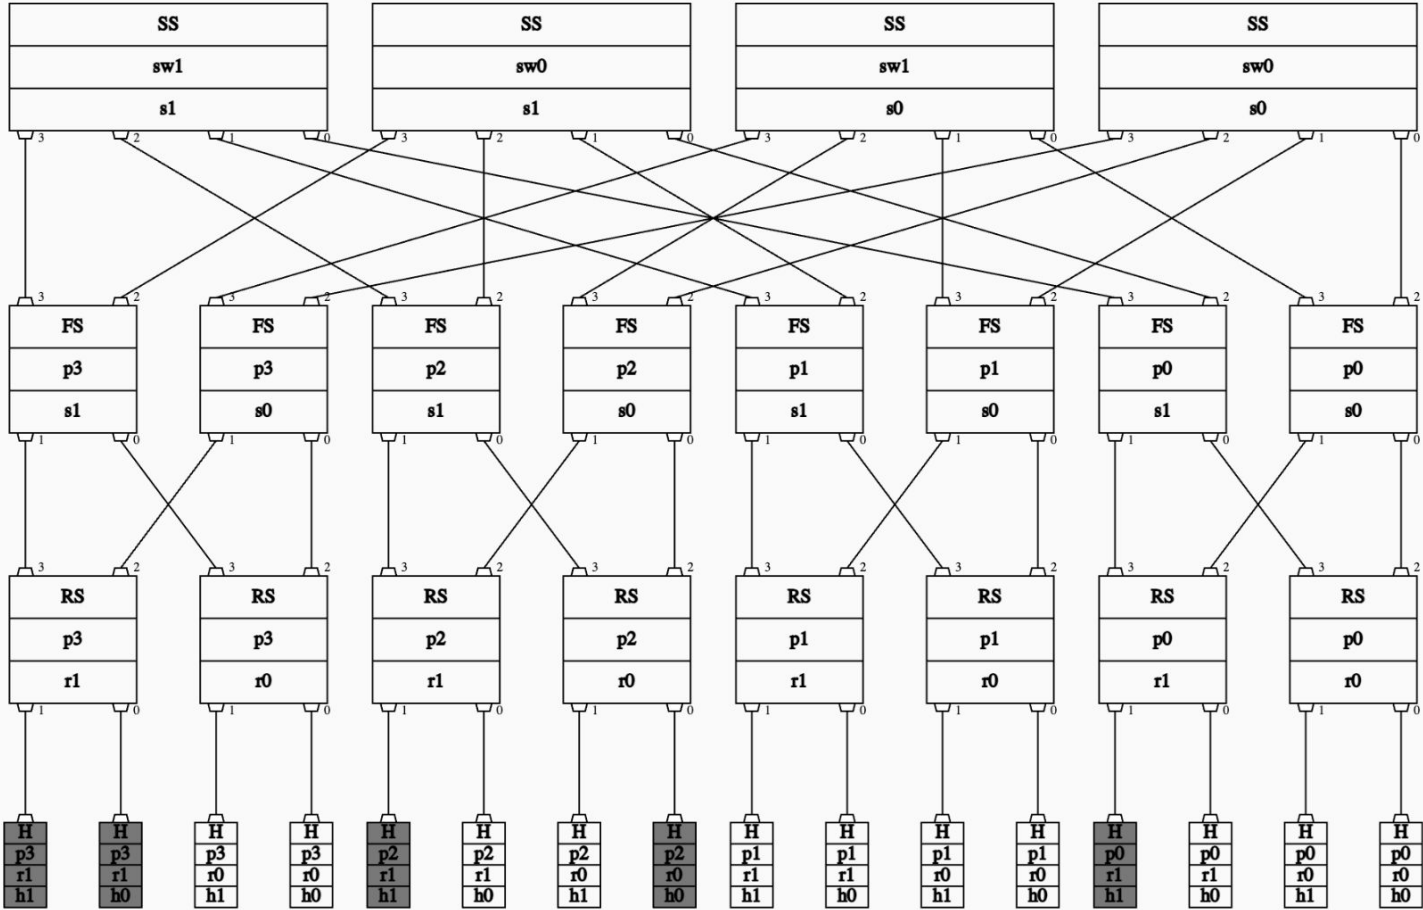
MULTICAST - JOIN



CA

0xBF
Spine Switch ID: 1
Spine ID: 0
CA identifier
CA identifier
CA identifier

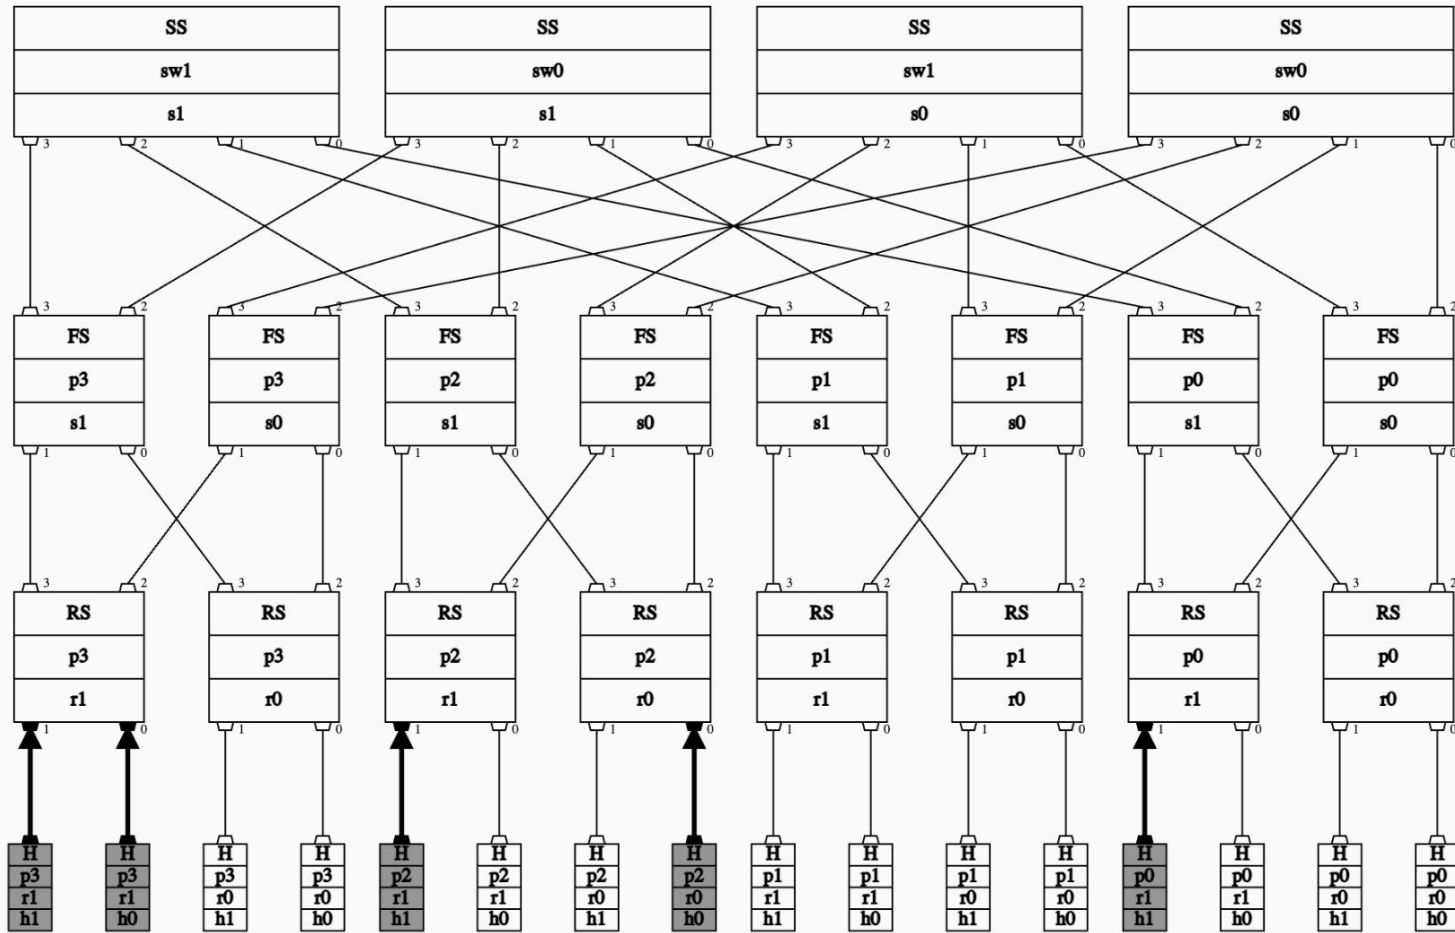
MULTICAST - JOIN



CA

0xBF
Spine Switch ID: 1
Spine ID: 0
CA identifier
CA identifier
CA identifier

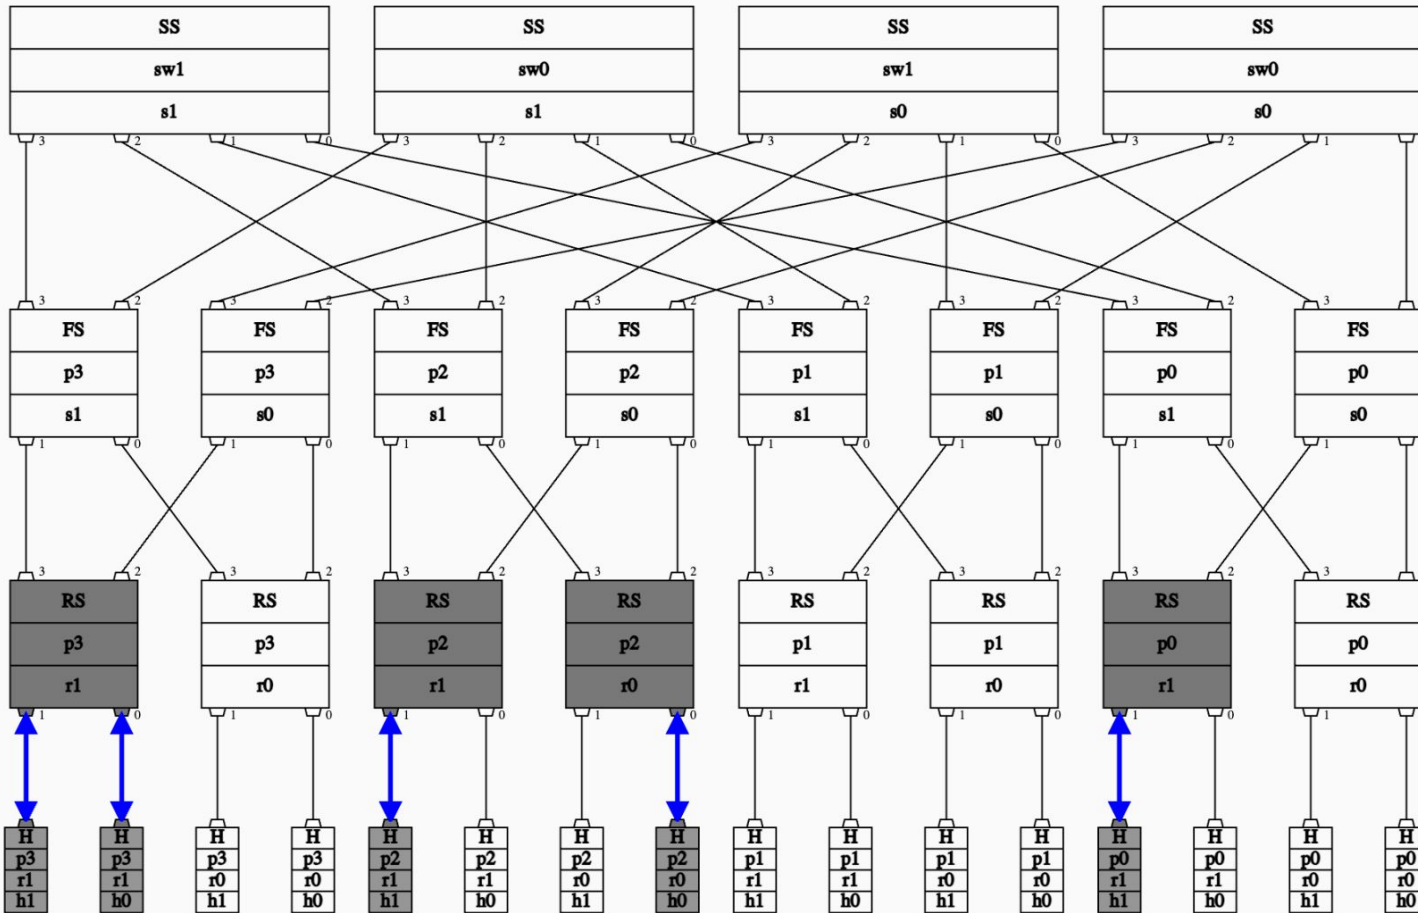
MULTICAST - JOIN



CA

0xBF
Spine Switch ID: 1
Spine ID: 0
CA identifier
CA identifier
CA identifier

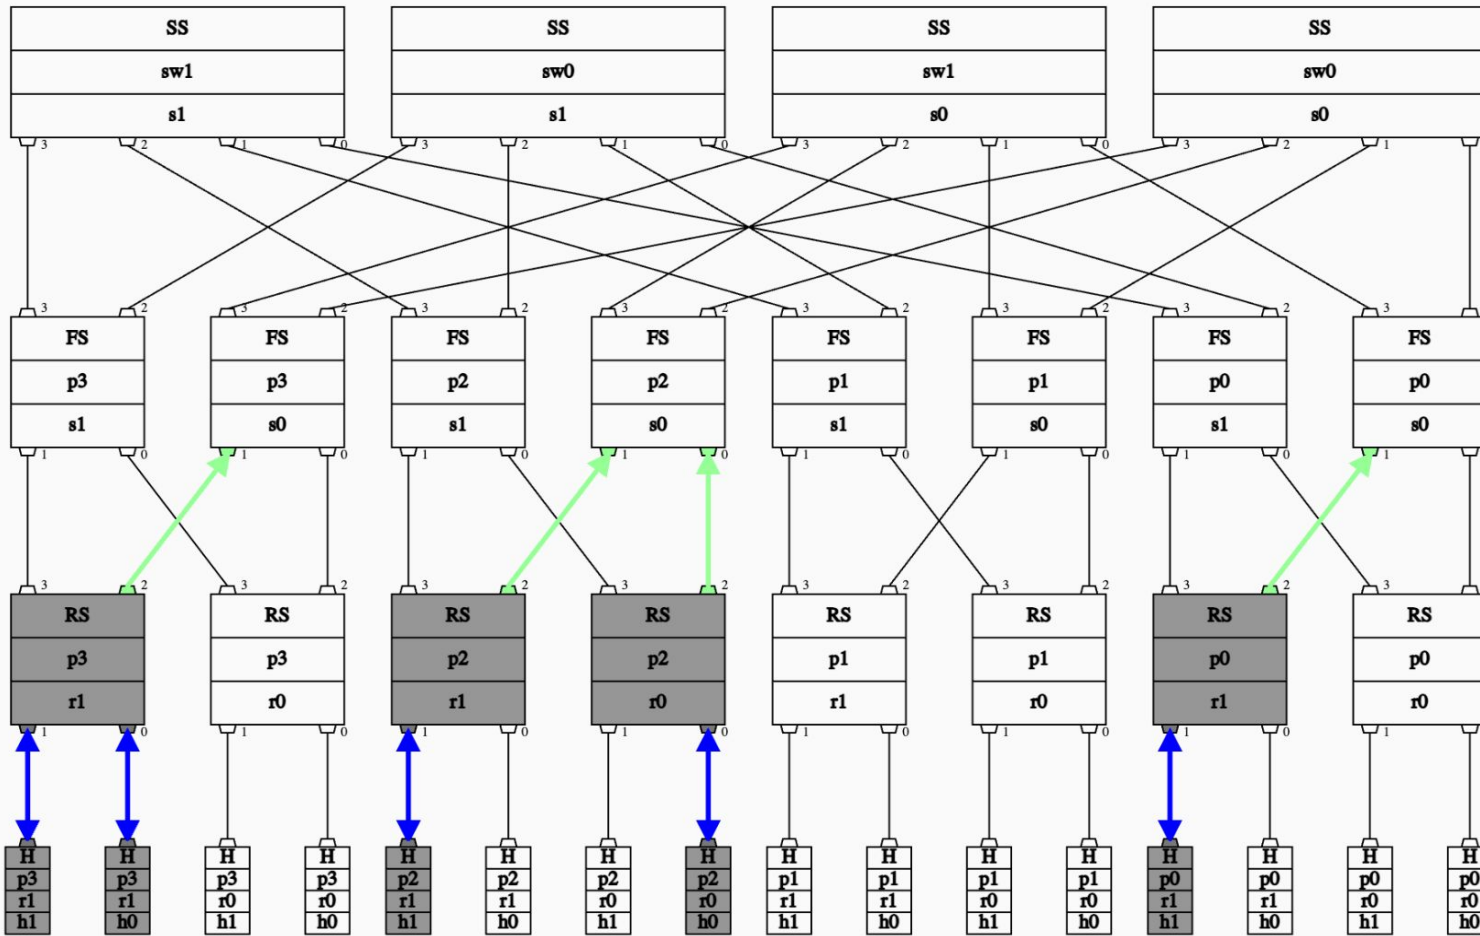
MULTICAST - JOIN



CA

0xBF
Spine Switch ID: 1
Spine ID: 0
CA identifier
CA identifier
CA identifier

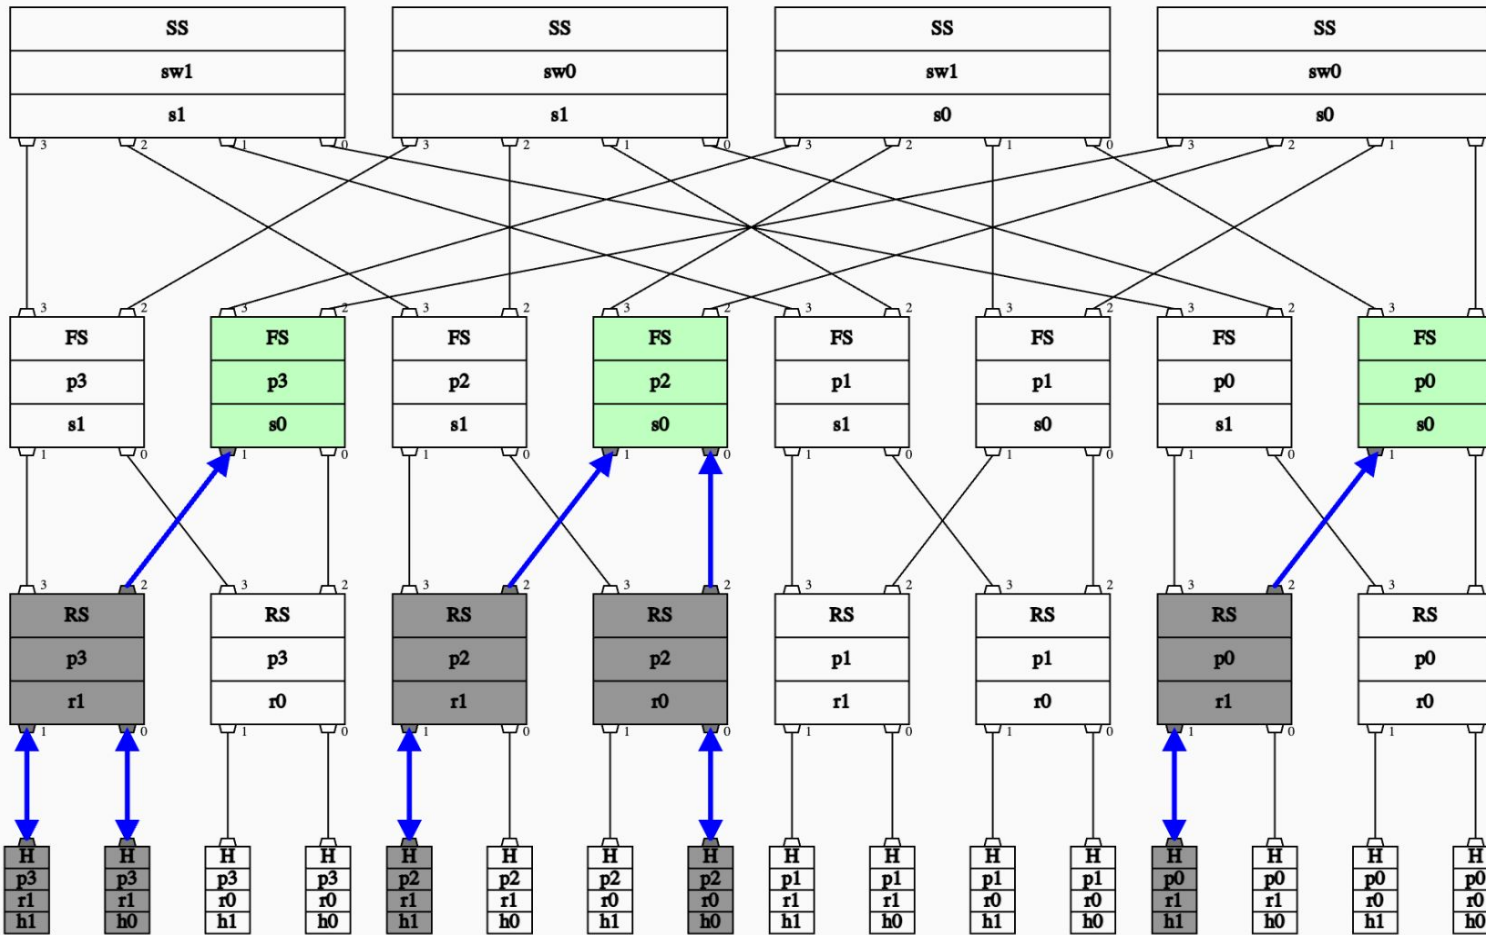
MULTICAST - JOIN



CA

0xBF
Spine Switch ID: 1
Spine ID: 0
CA identifier
CA identifier
CA identifier

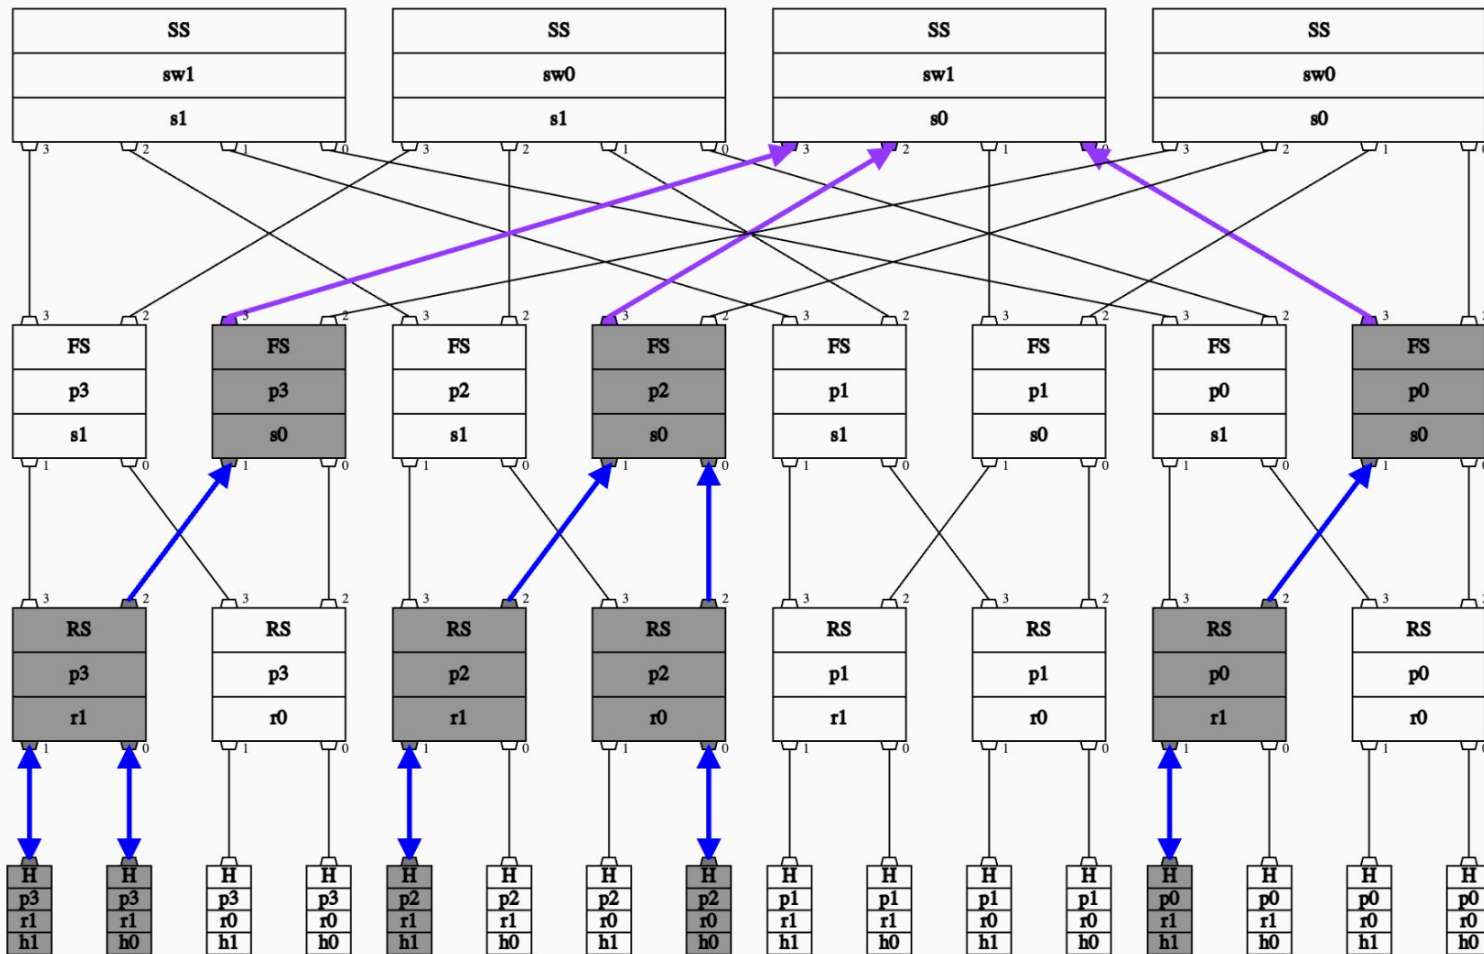
MULTICAST - JOIN



CA

0xBF
Spine Switch ID: 1
Spine ID: 0
CA identifier
CA identifier
CA identifier

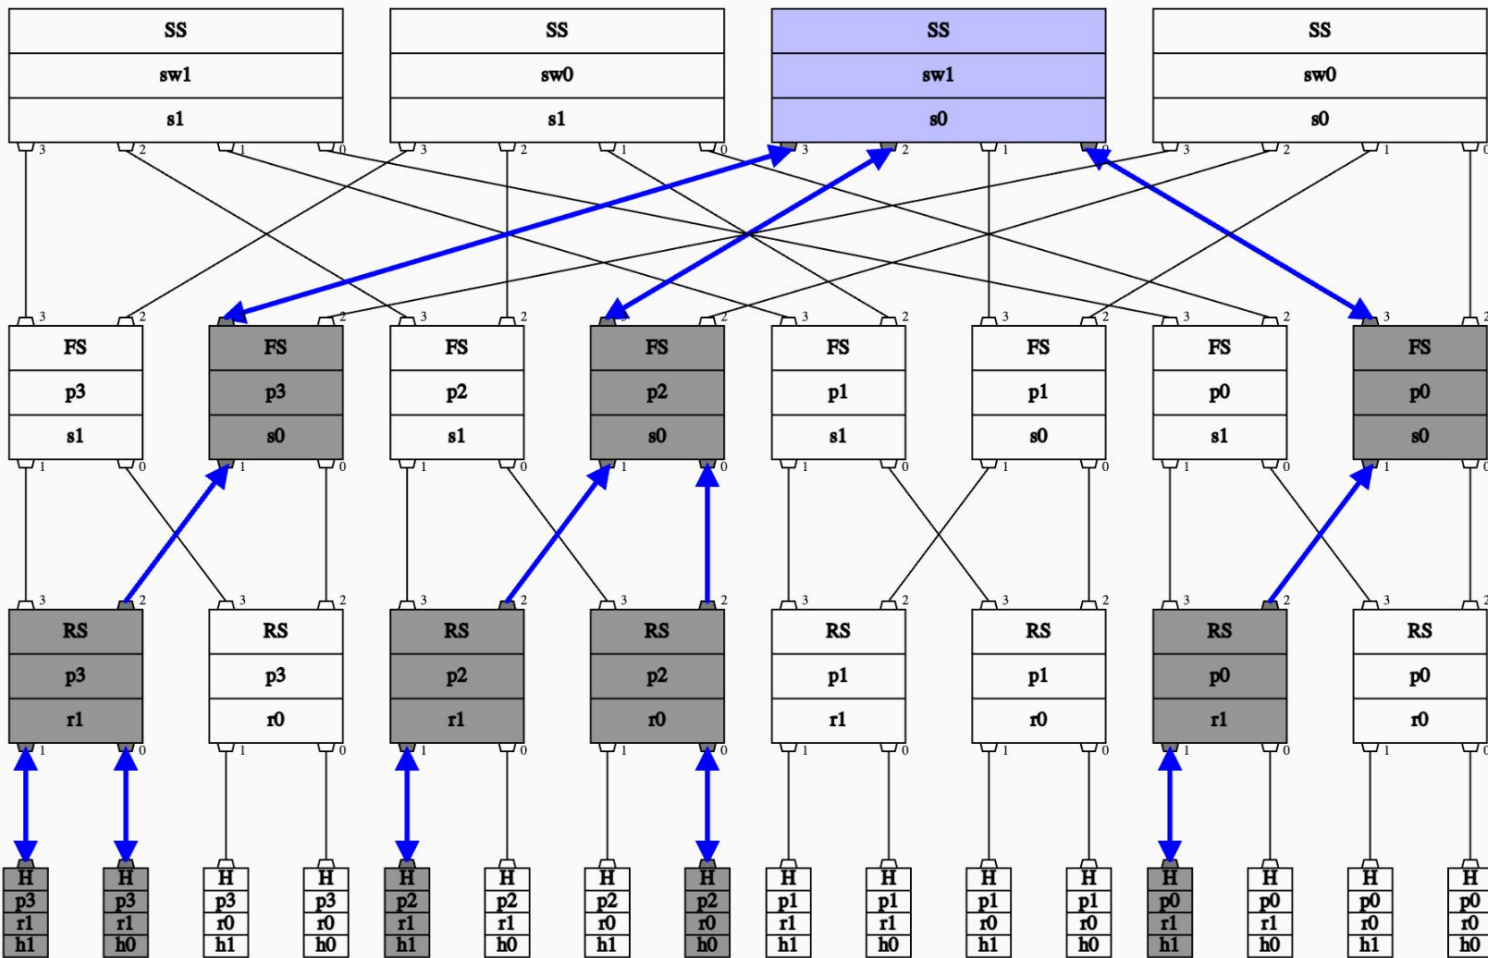
MULTICAST - JOIN



CA

0xBF
Spine Switch ID: 1
Spine ID: 0
CA identifier
CA identifier
CA identifier

MULTICAST - JOIN

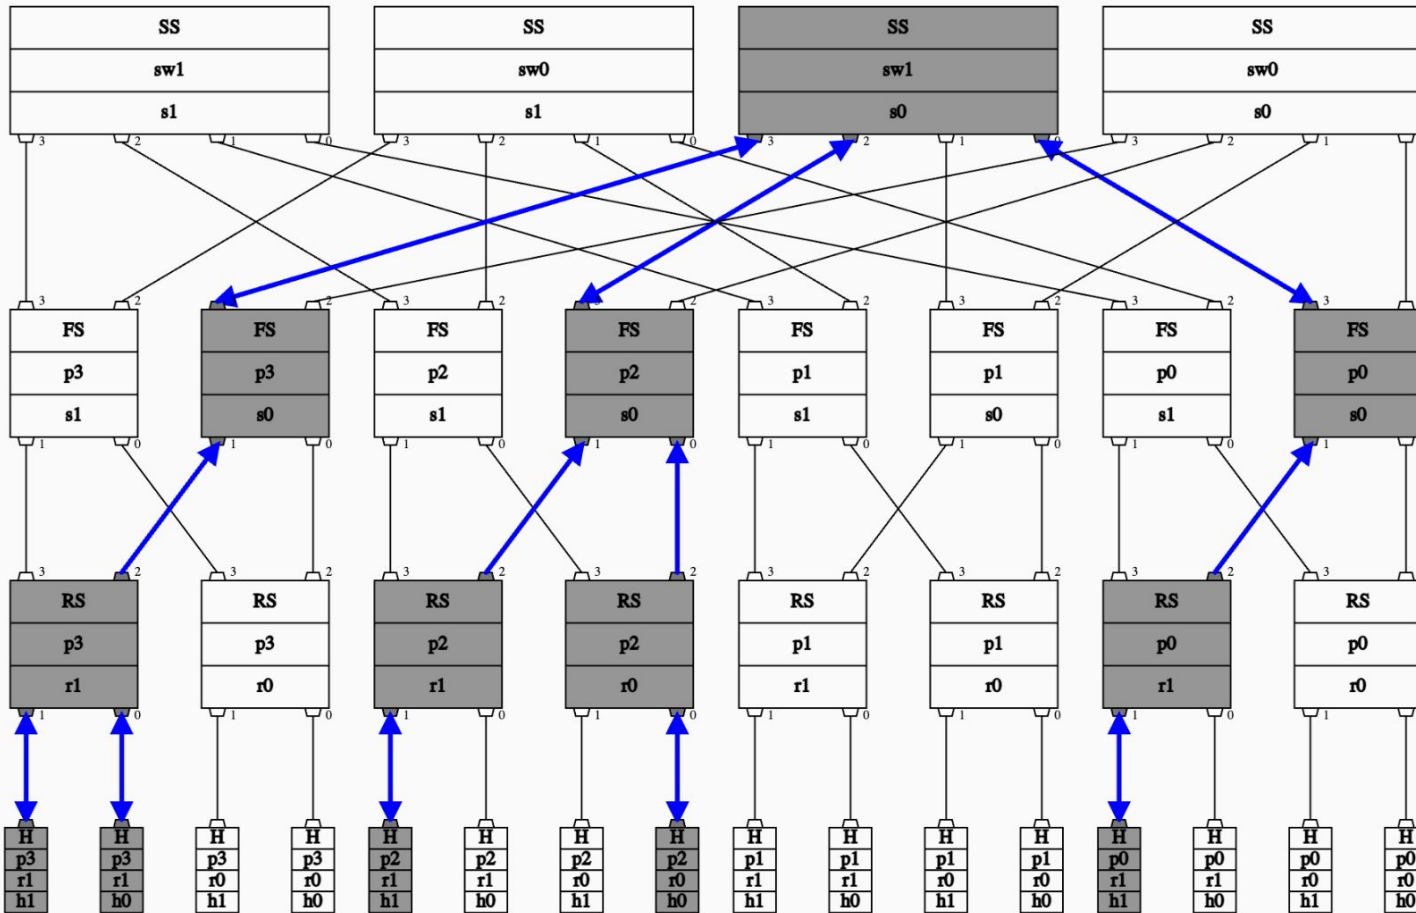


CA

0xBF
Spine Switch ID: 1
Spine ID: 0
CA identifier
CA identifier
CA identifier

Multicast Frame Forwarding

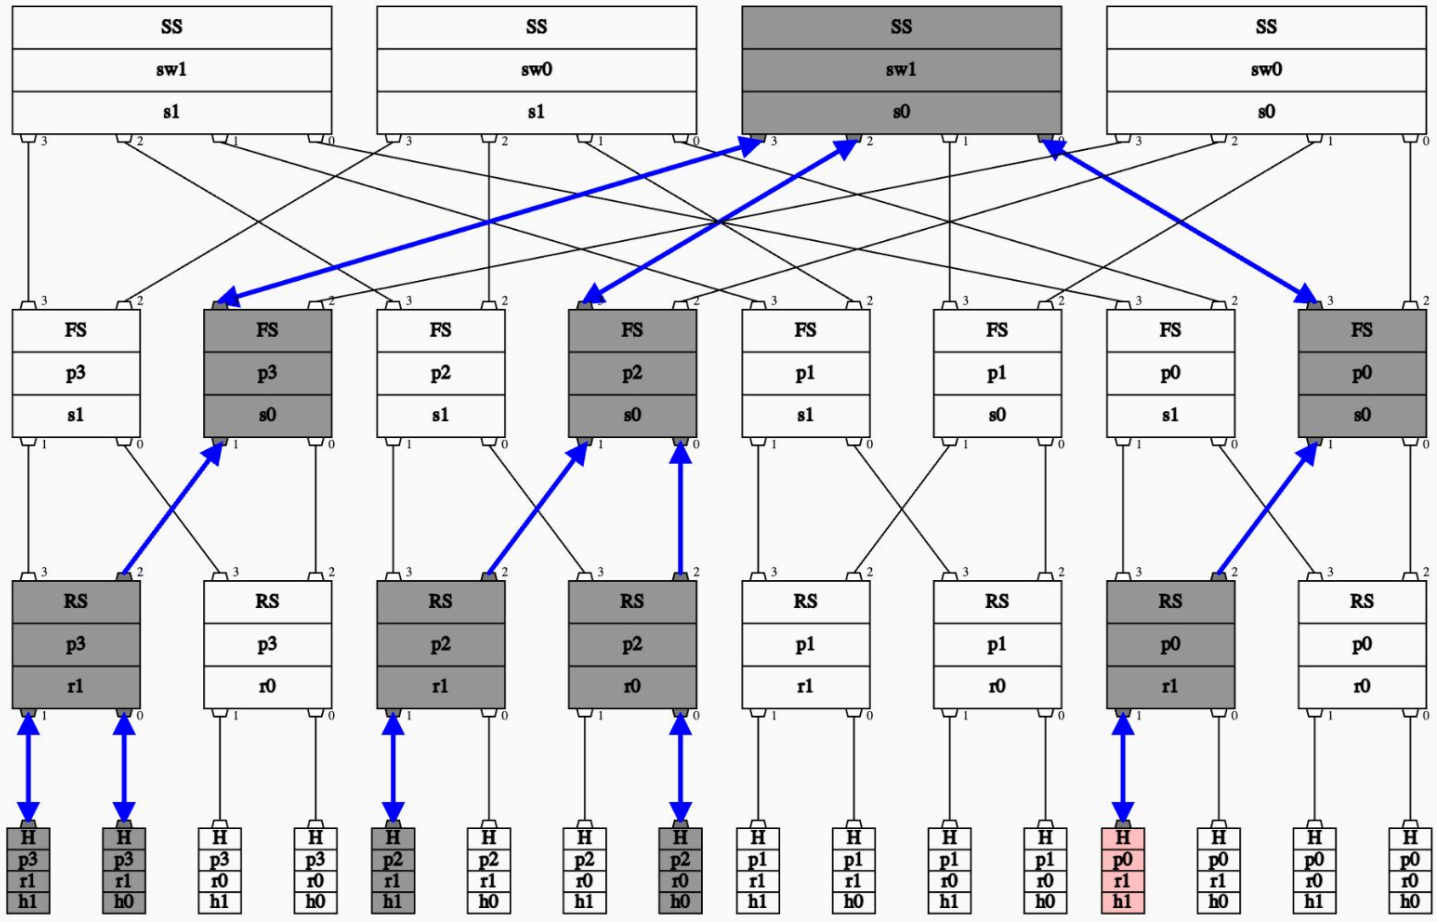
MULTICAST - DATA



CA

0xBF
Spine Switch ID: 1
Spine ID: 0
CA identifier
CA identifier
CA identifier

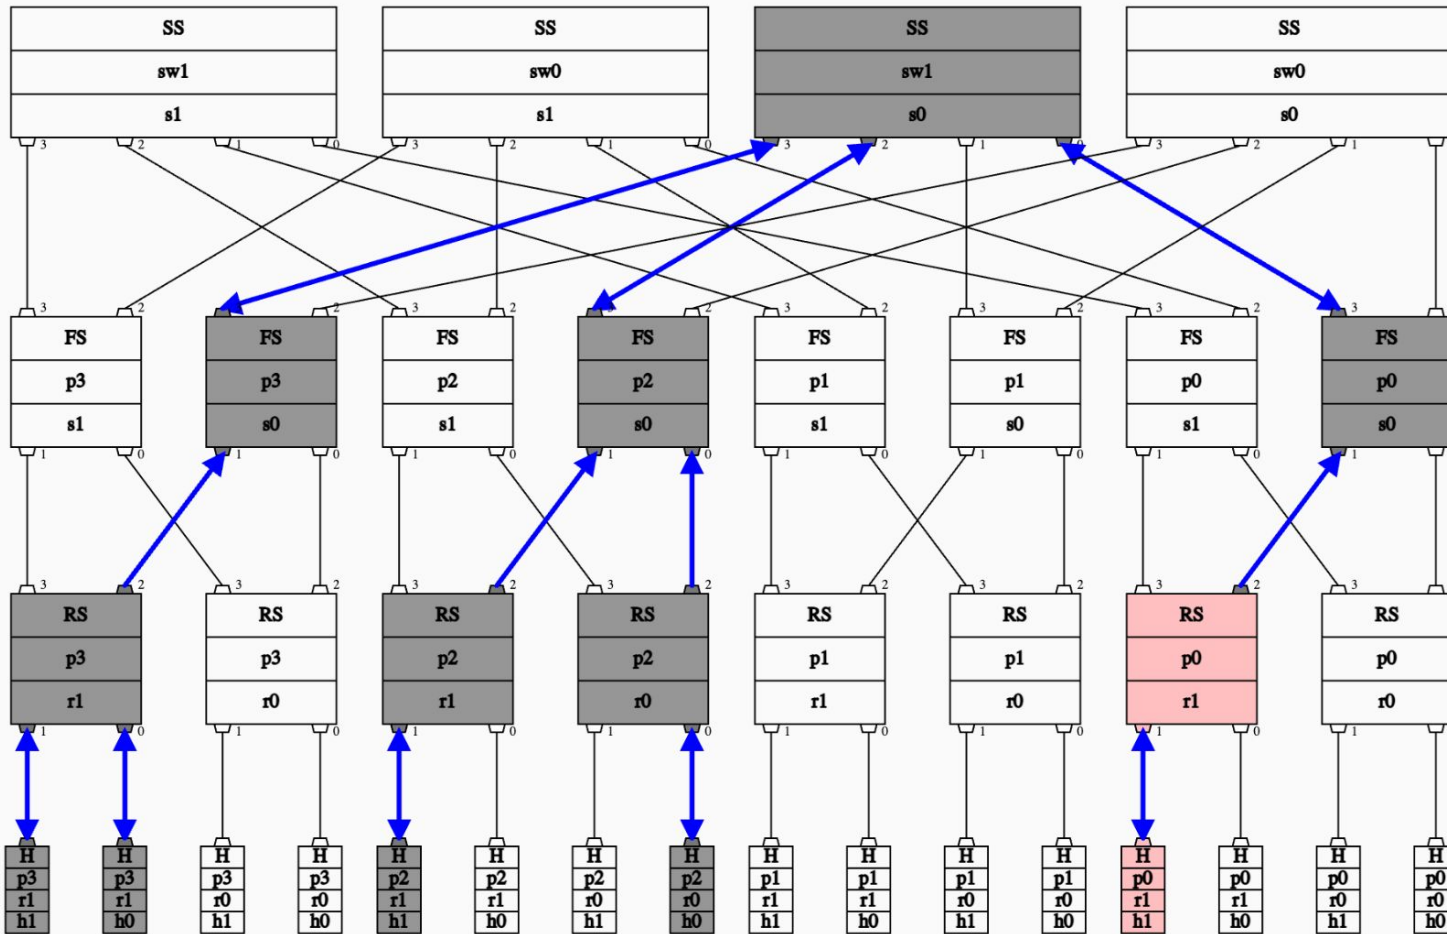
MULTICAST - DATA



CA

0xBF
Spine Switch ID: 1
Spine ID: 0
CA identifier
CA identifier
CA identifier

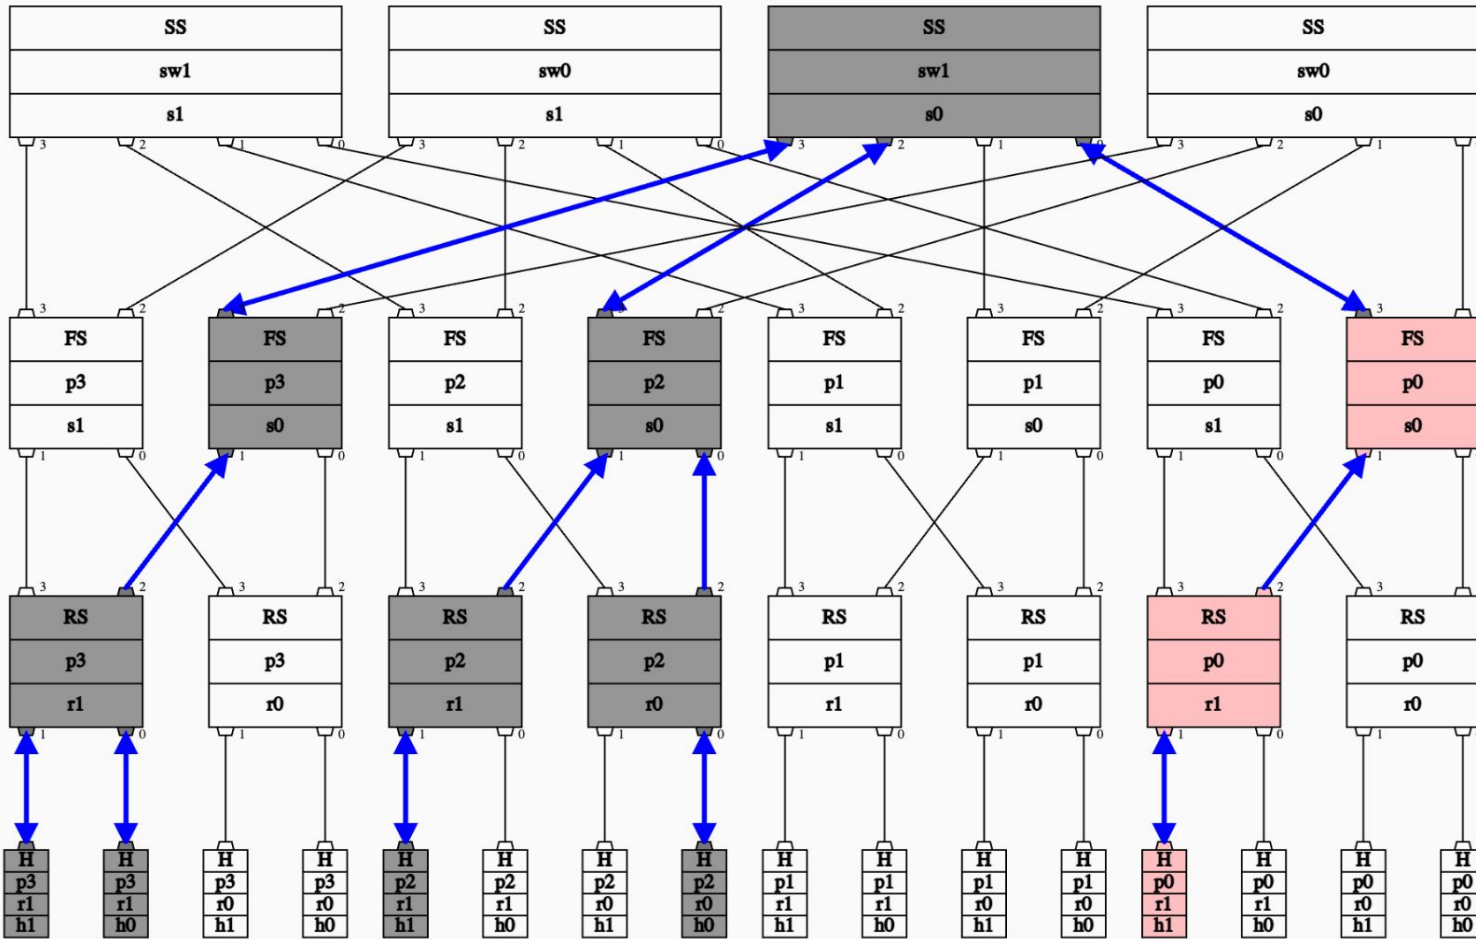
MULTICAST - DATA



CA

0xBF
Spine Switch ID: 1
Spine ID: 0
CA identifier
CA identifier
CA identifier

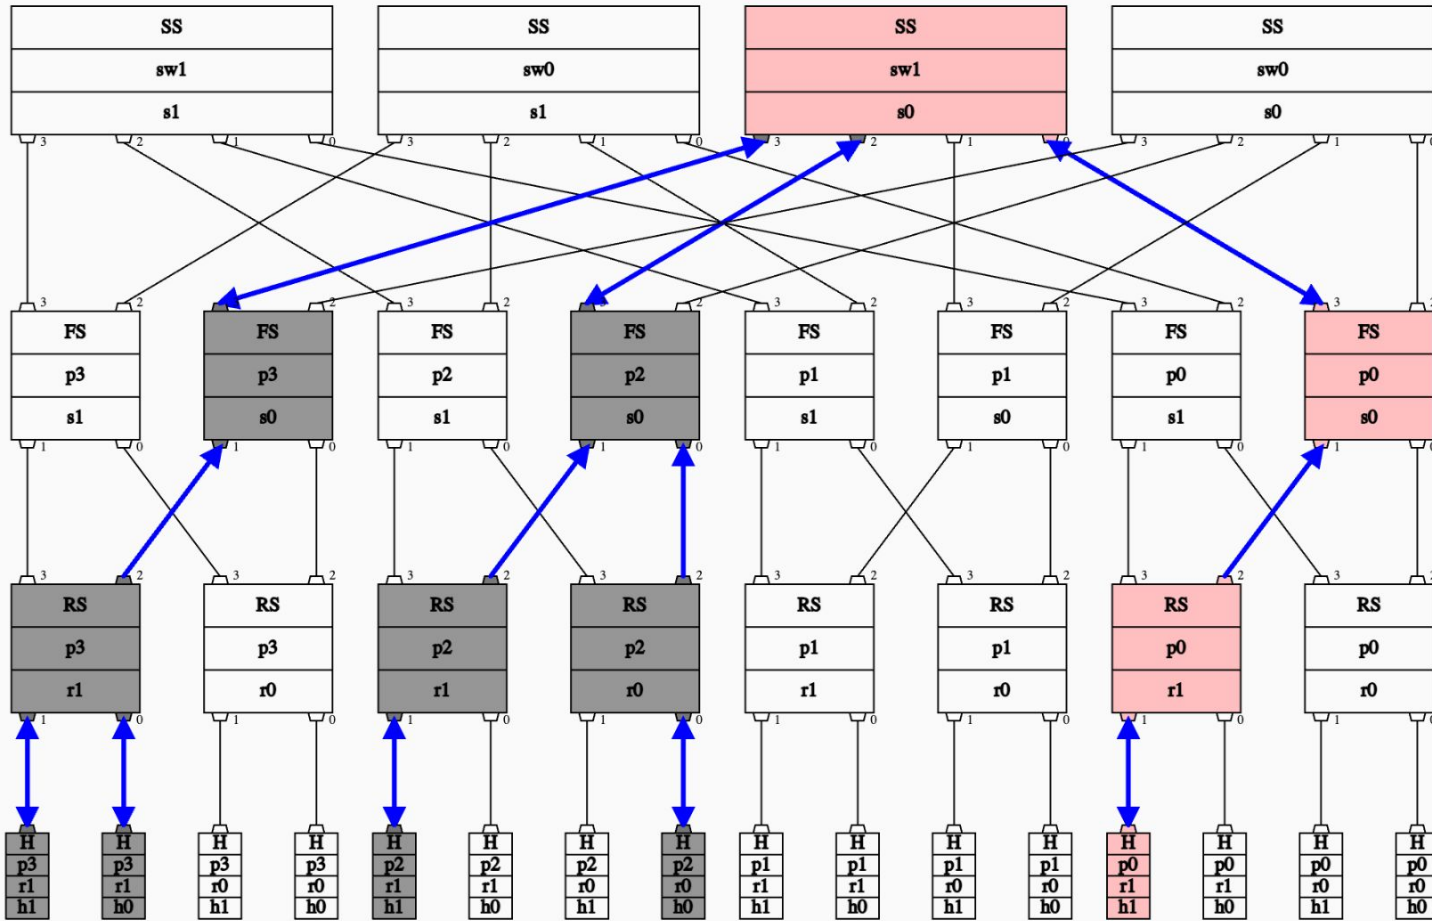
MULTICAST - DATA



CA

0xBF
Spine Switch ID: 1
Spine ID: 0
CA identifier
CA identifier
CA identifier

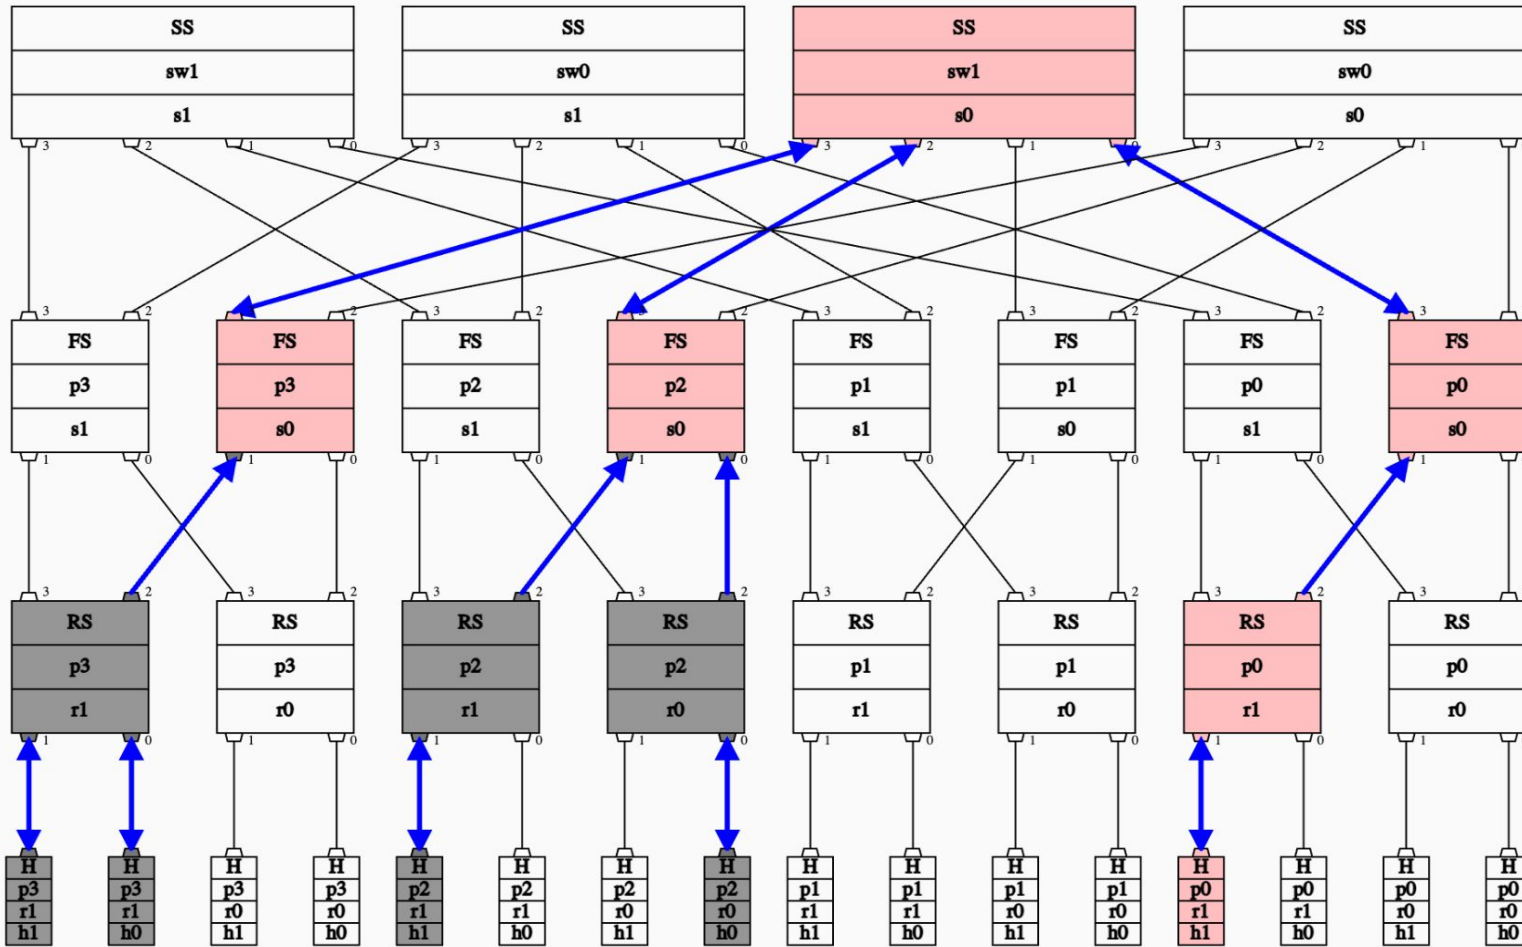
MULTICAST - DATA



CA

0xBF
Spine Switch ID: 1
Spine ID: 0
CA identifier
CA identifier
CA identifier

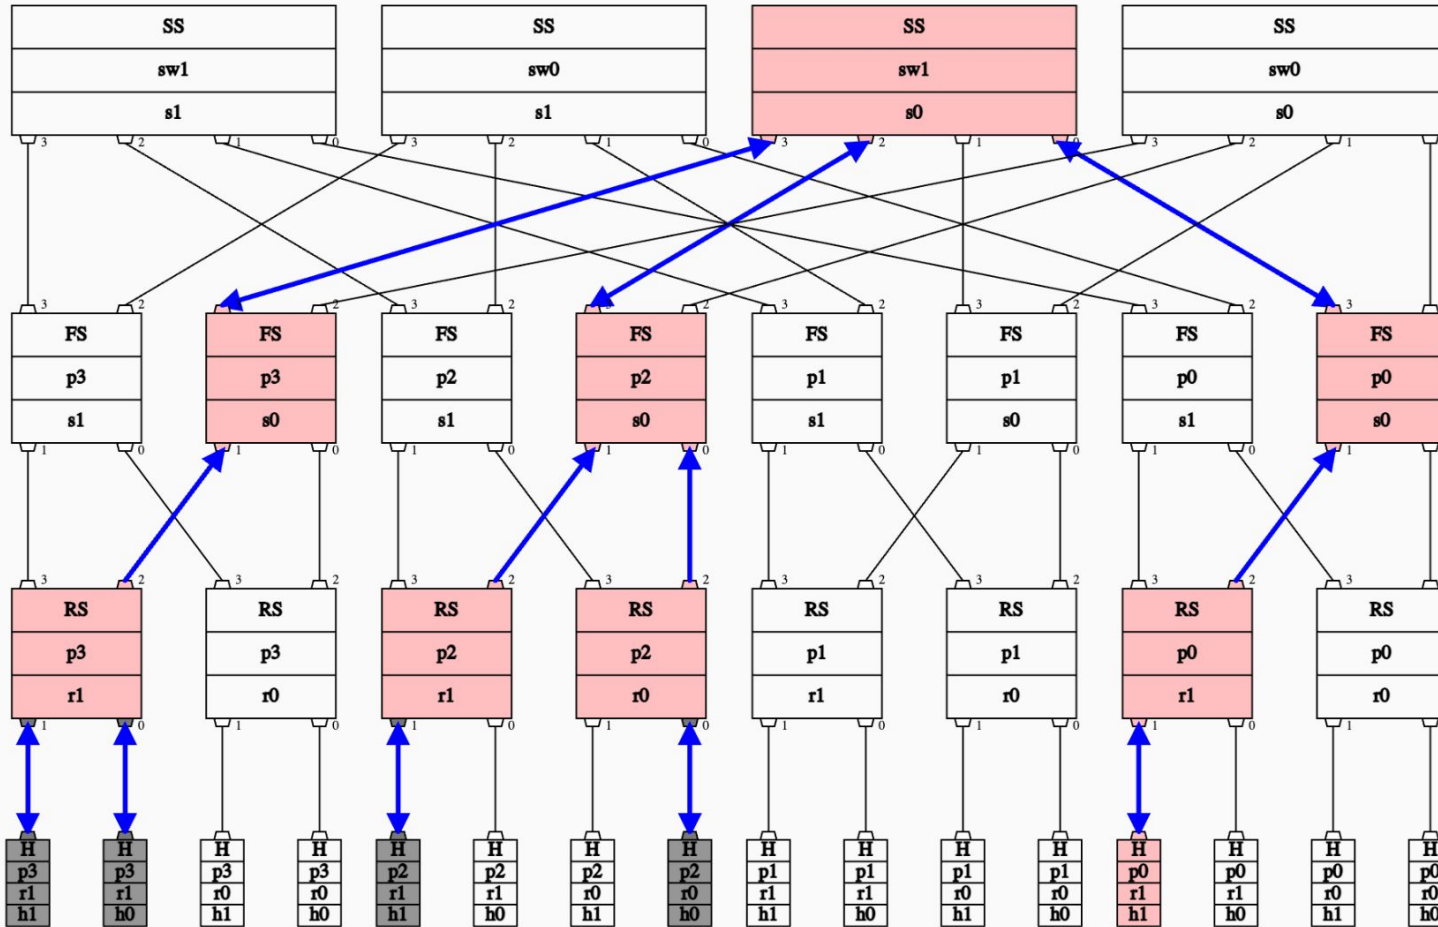
MULTICAST - DATA



CA

0xBF
Spine Switch ID: 1
Spine ID: 0
CA identifier
CA identifier
CA identifier

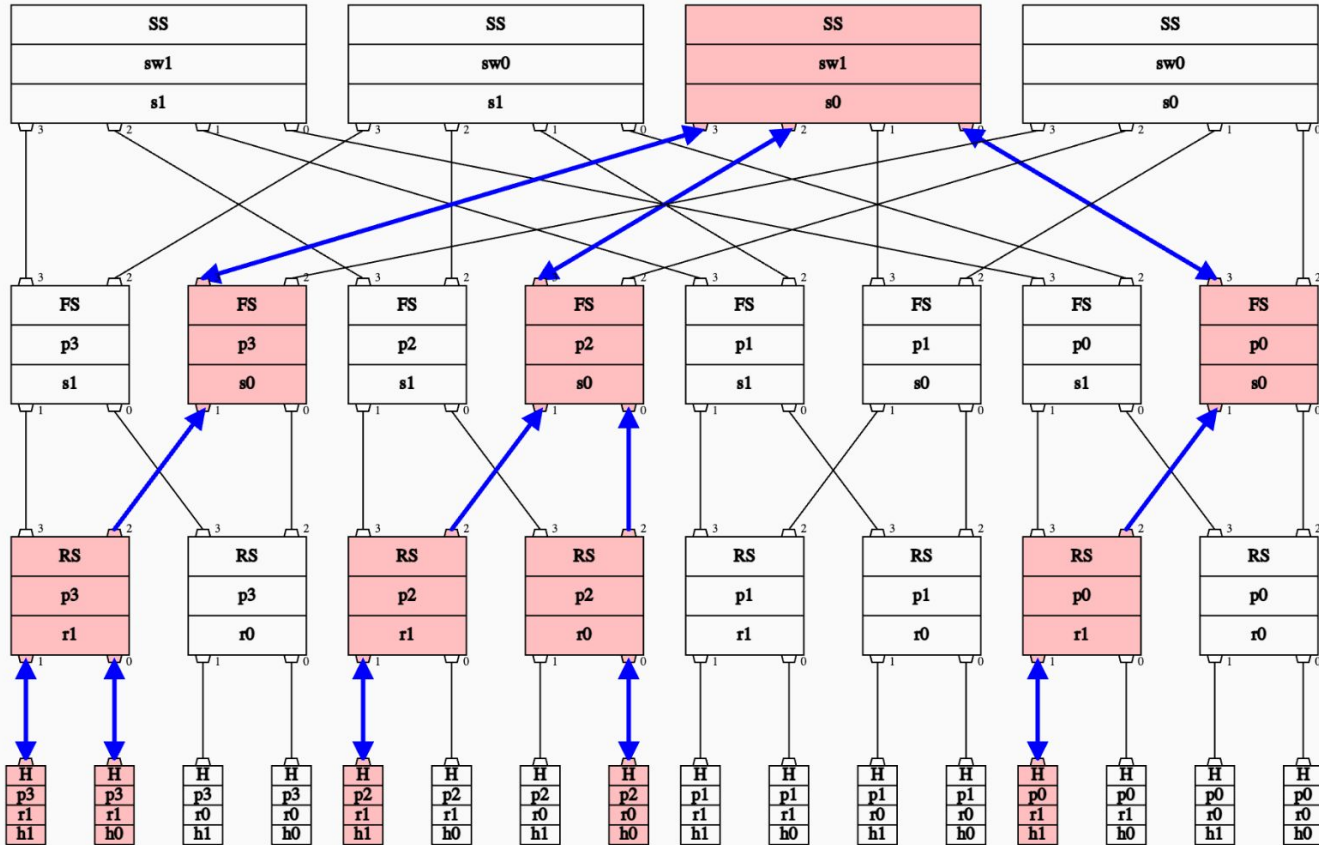
MULTICAST - DATA



CA

0xBF
Spine Switch ID: 1
Spine ID: 0
CA identifier
CA identifier
CA identifier

MULTICAST - DATA



CA

0xBF
Spine Switch ID: 1
Spine ID: 0
CA identifier
CA identifier
CA identifier

Summary

- The implementation BARC, unicast, and collective multicast in a Clos Fat-tree computing network operates as expected.