UNIVERSITEIT VAN AMSTERDAM

Multicast support on the AMS-IX platform

Attilla de Groot & Yuri Schaeffer

Research Project 2

1 July, 2009

AMS-IX

- One of the largest internet exchanges in the world
- Peak traffic at 675Gbit/s
- Broadcast Domain
- Separate vlan for multicast

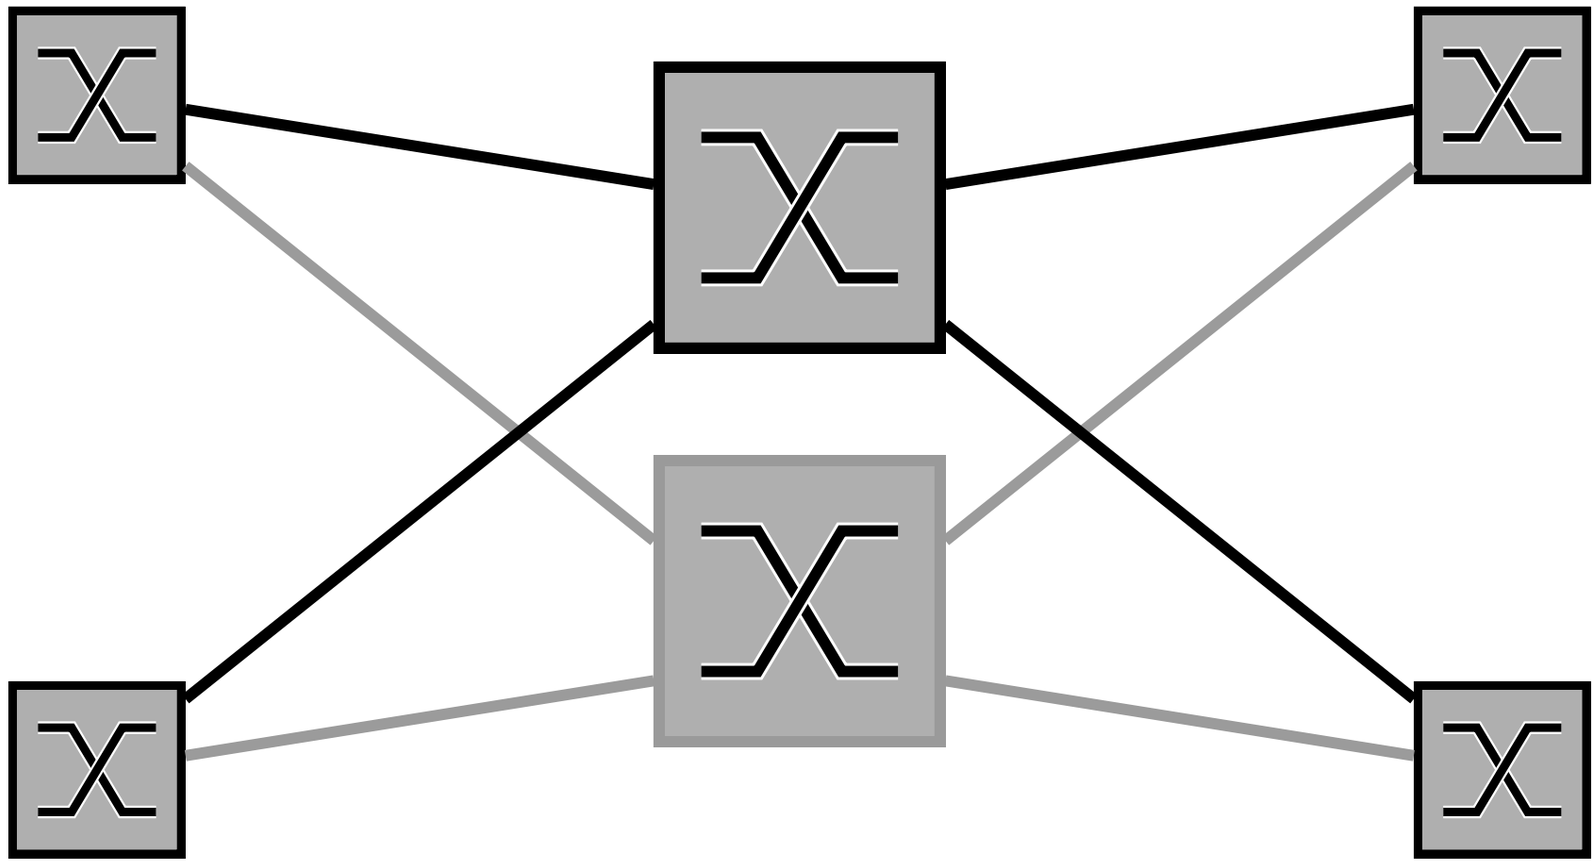

AMS-IXv4

- Current platform, maximum of scalability
- No 100Gbit/s available
- Implementation of MPLS/VPLS

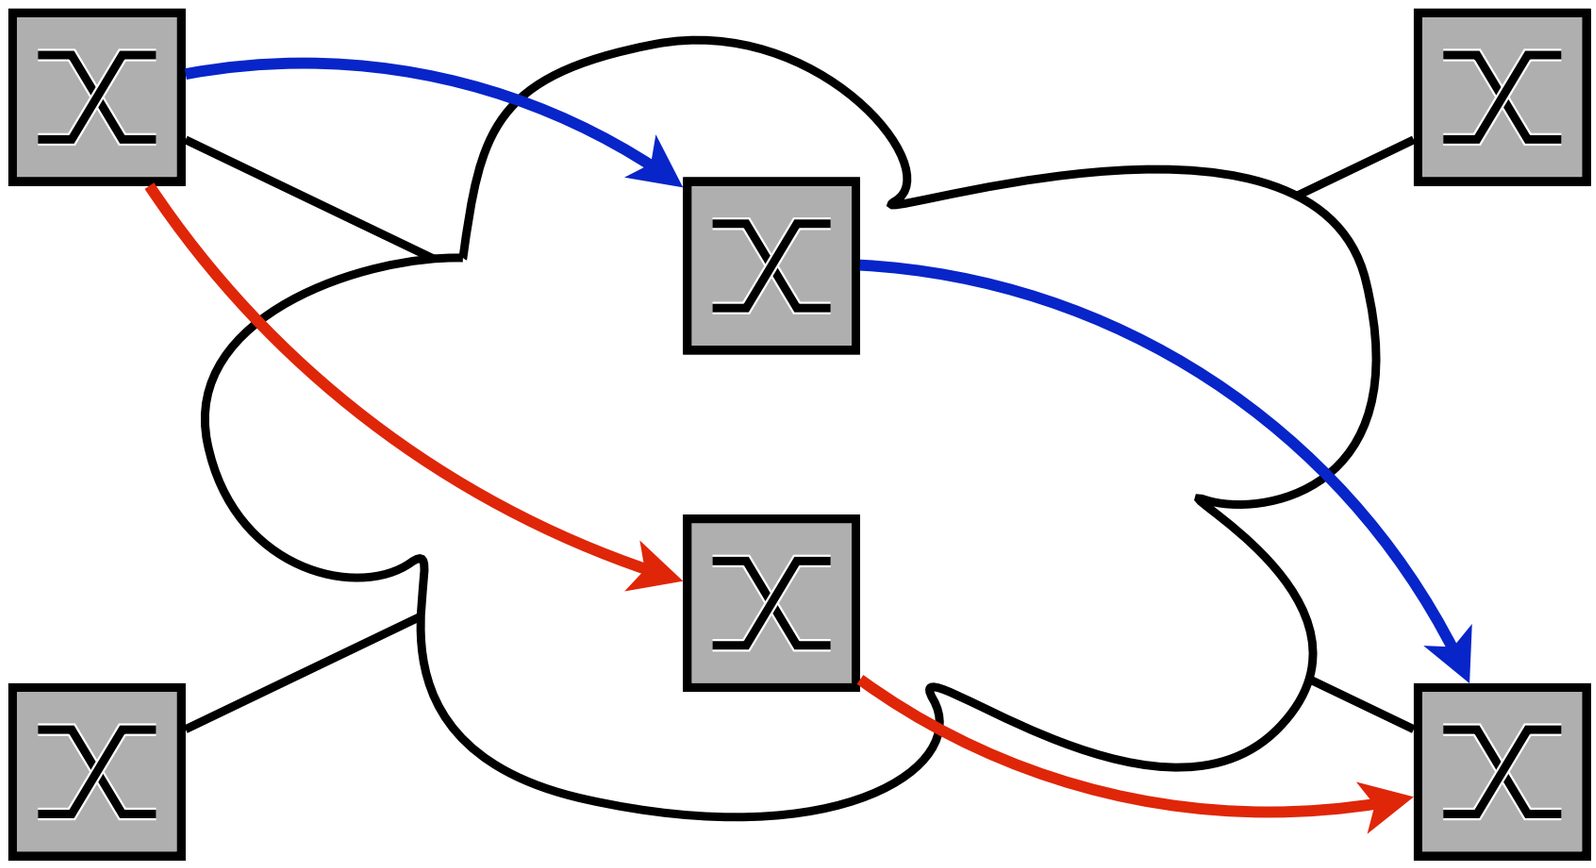

AMS-IXv4

- Current platform, maximum of scalability
- No 100Gbit/s available
- Implementation of MPLS/VPLS

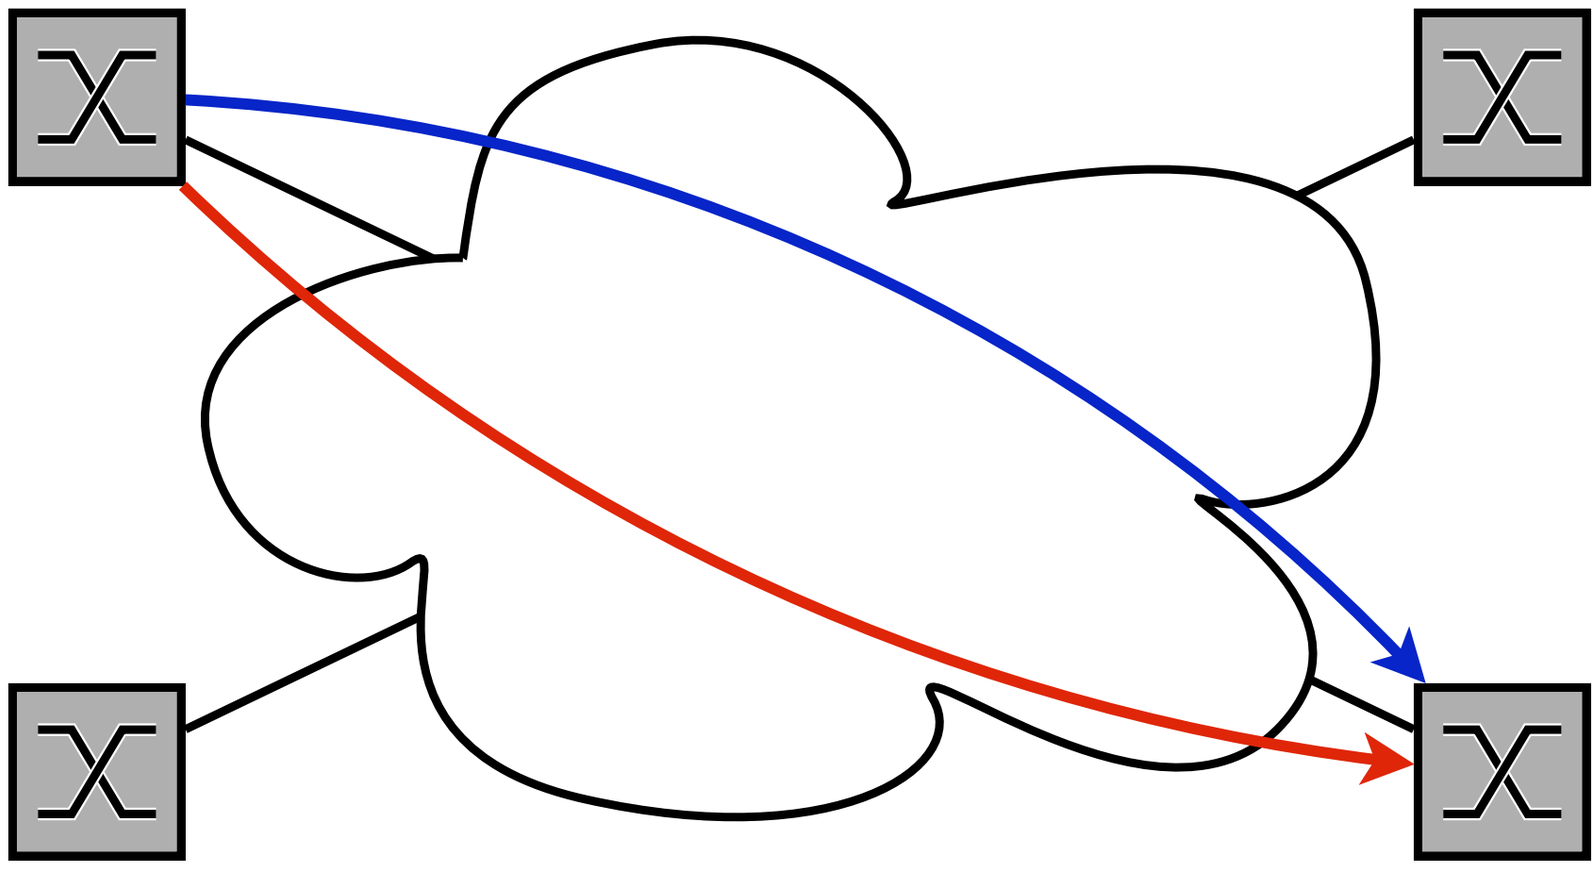

AMS-IXv4

- Current platform, maximum of scalability
- No 100Gbit/s available
- Implementation of MPLS/VPLS

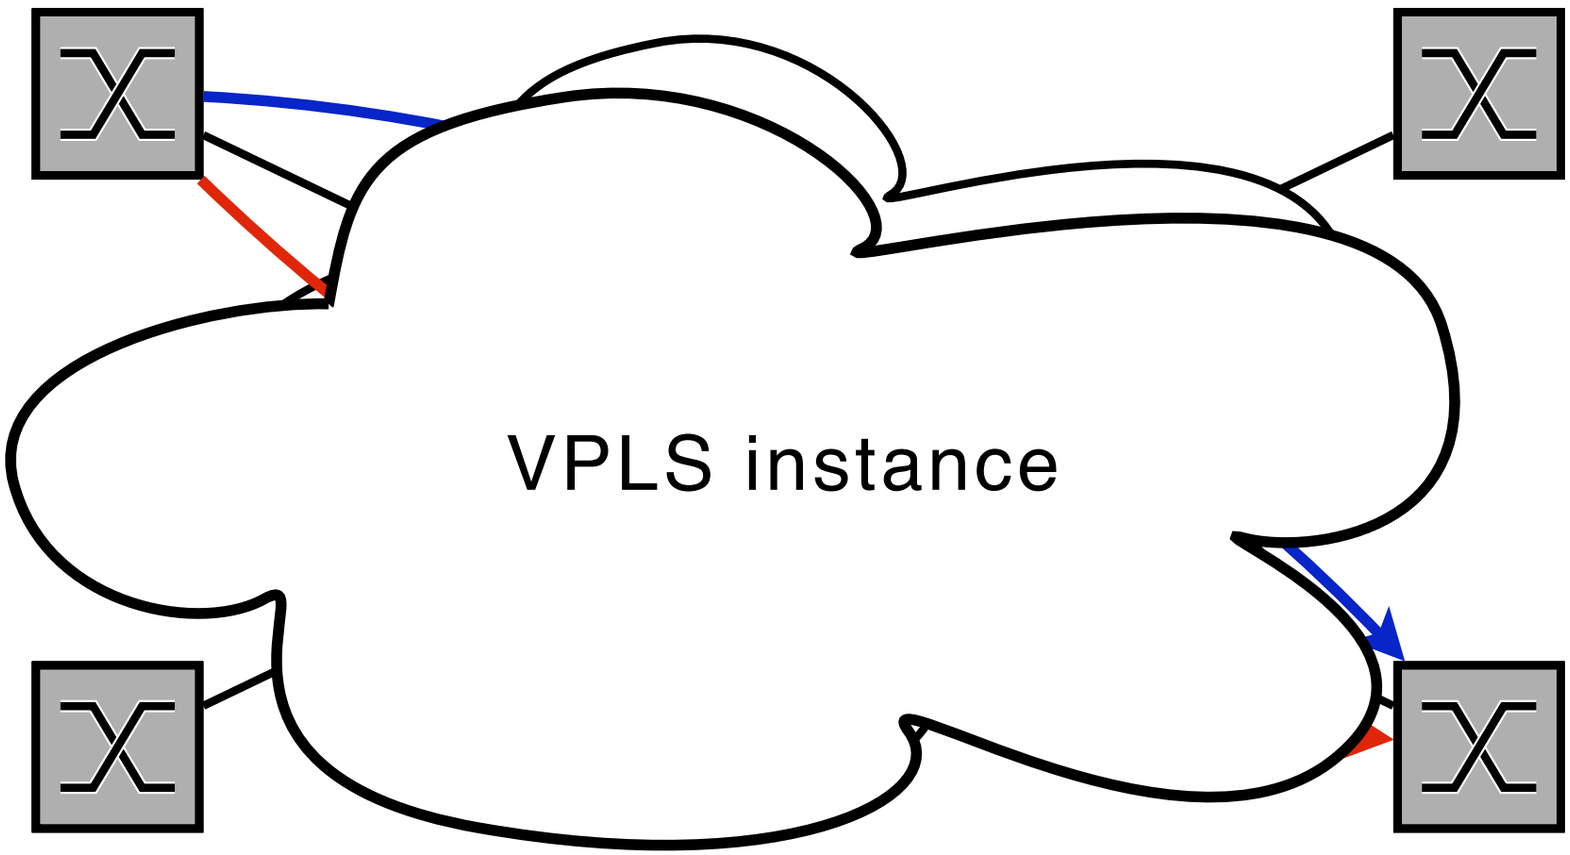

Research Question

How can multicast support be provided on a VPLS platform, such as is implemented on the AMS-IX, in an efficient way in respect to scalability, performance, and stability?

Multicast

Multicast Protocol	Usage
IGMP/MLDv2	Local
PIM-DM	Intra-domain
PIM-SM	**Intra/Inter-domain**
PIM-SSM	**Intra/Inter-domain**
PIM-BIDIR	**Intra/Inter-domain**
MSDP	**Inter-domain**

Routing Protocol	Usage
MOSPF	Intra-domain
DVMRP	Intra-domain
MBGP	**Inter-domain**

Multicast tree

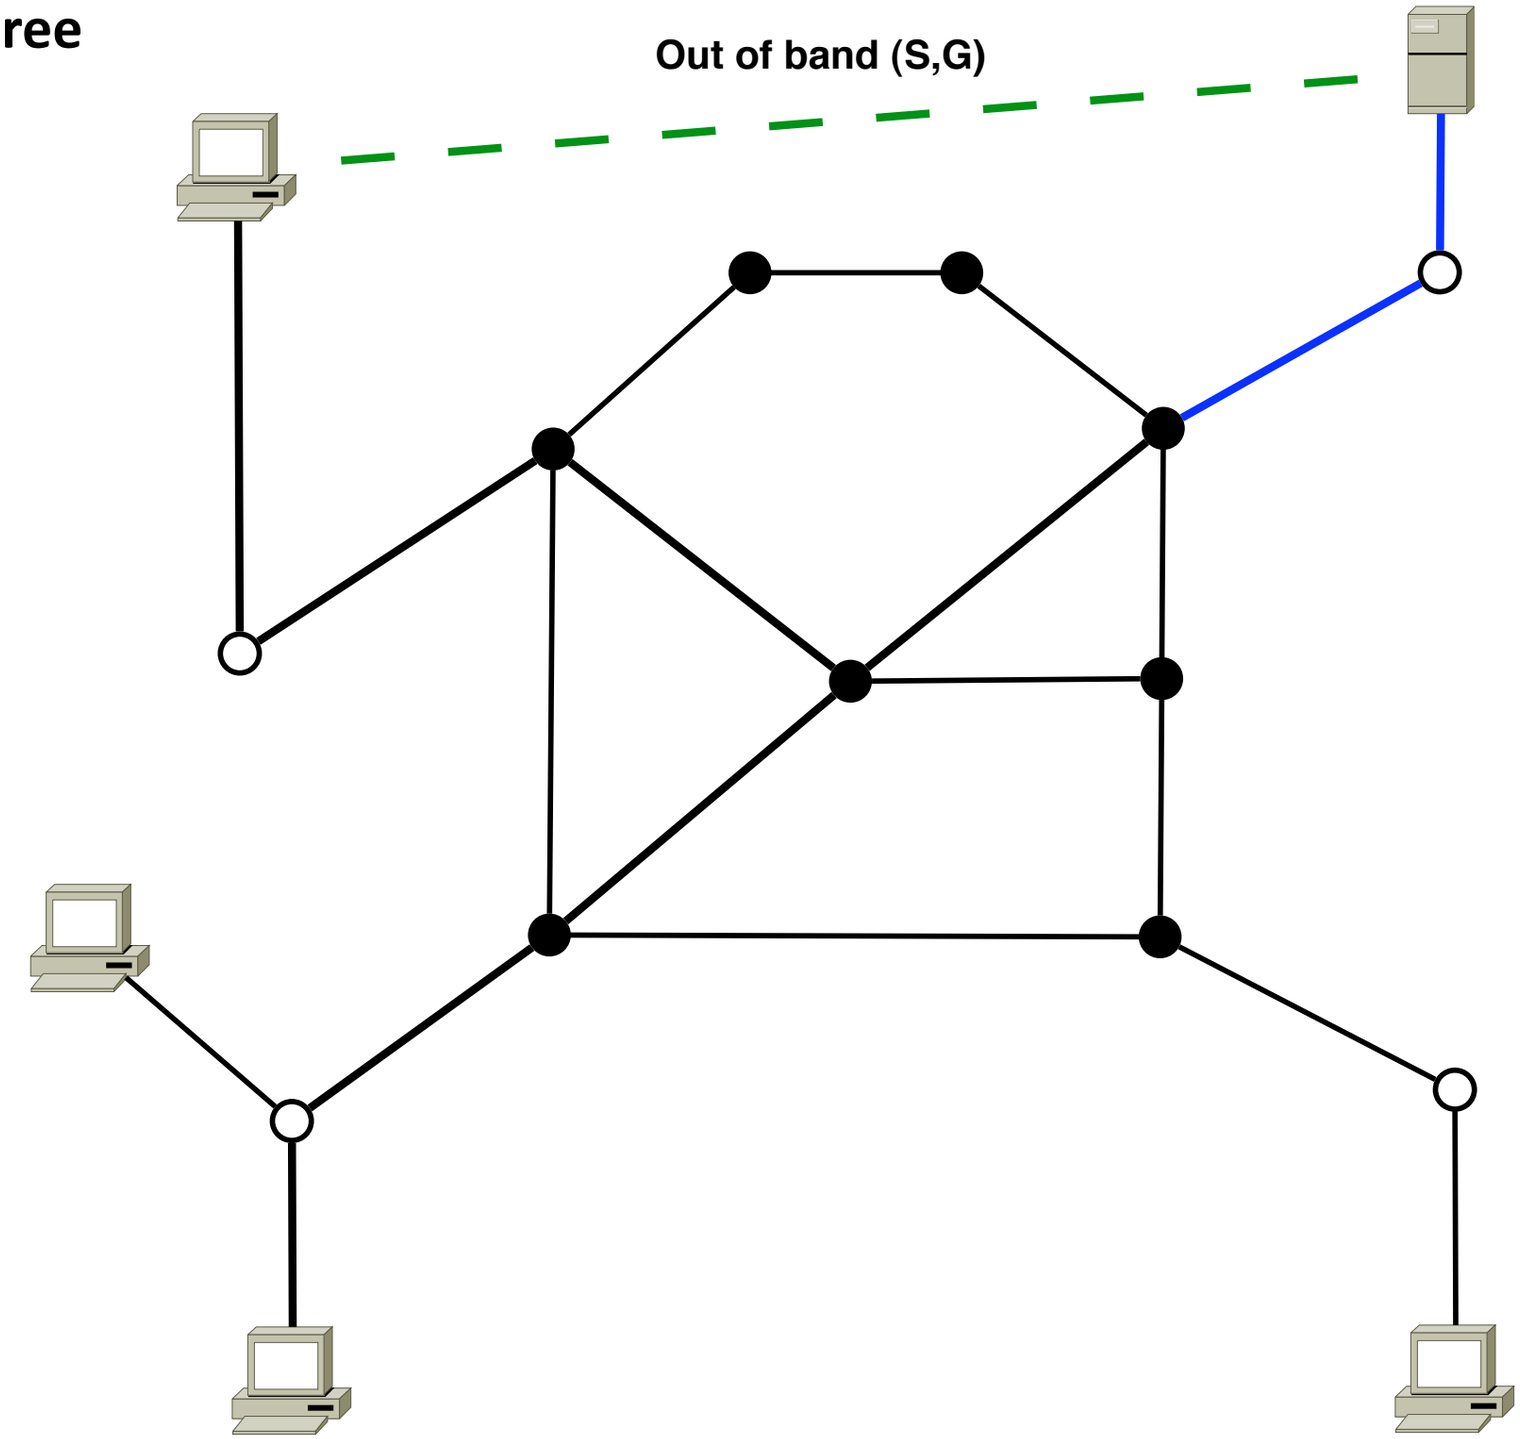

Multicast tree

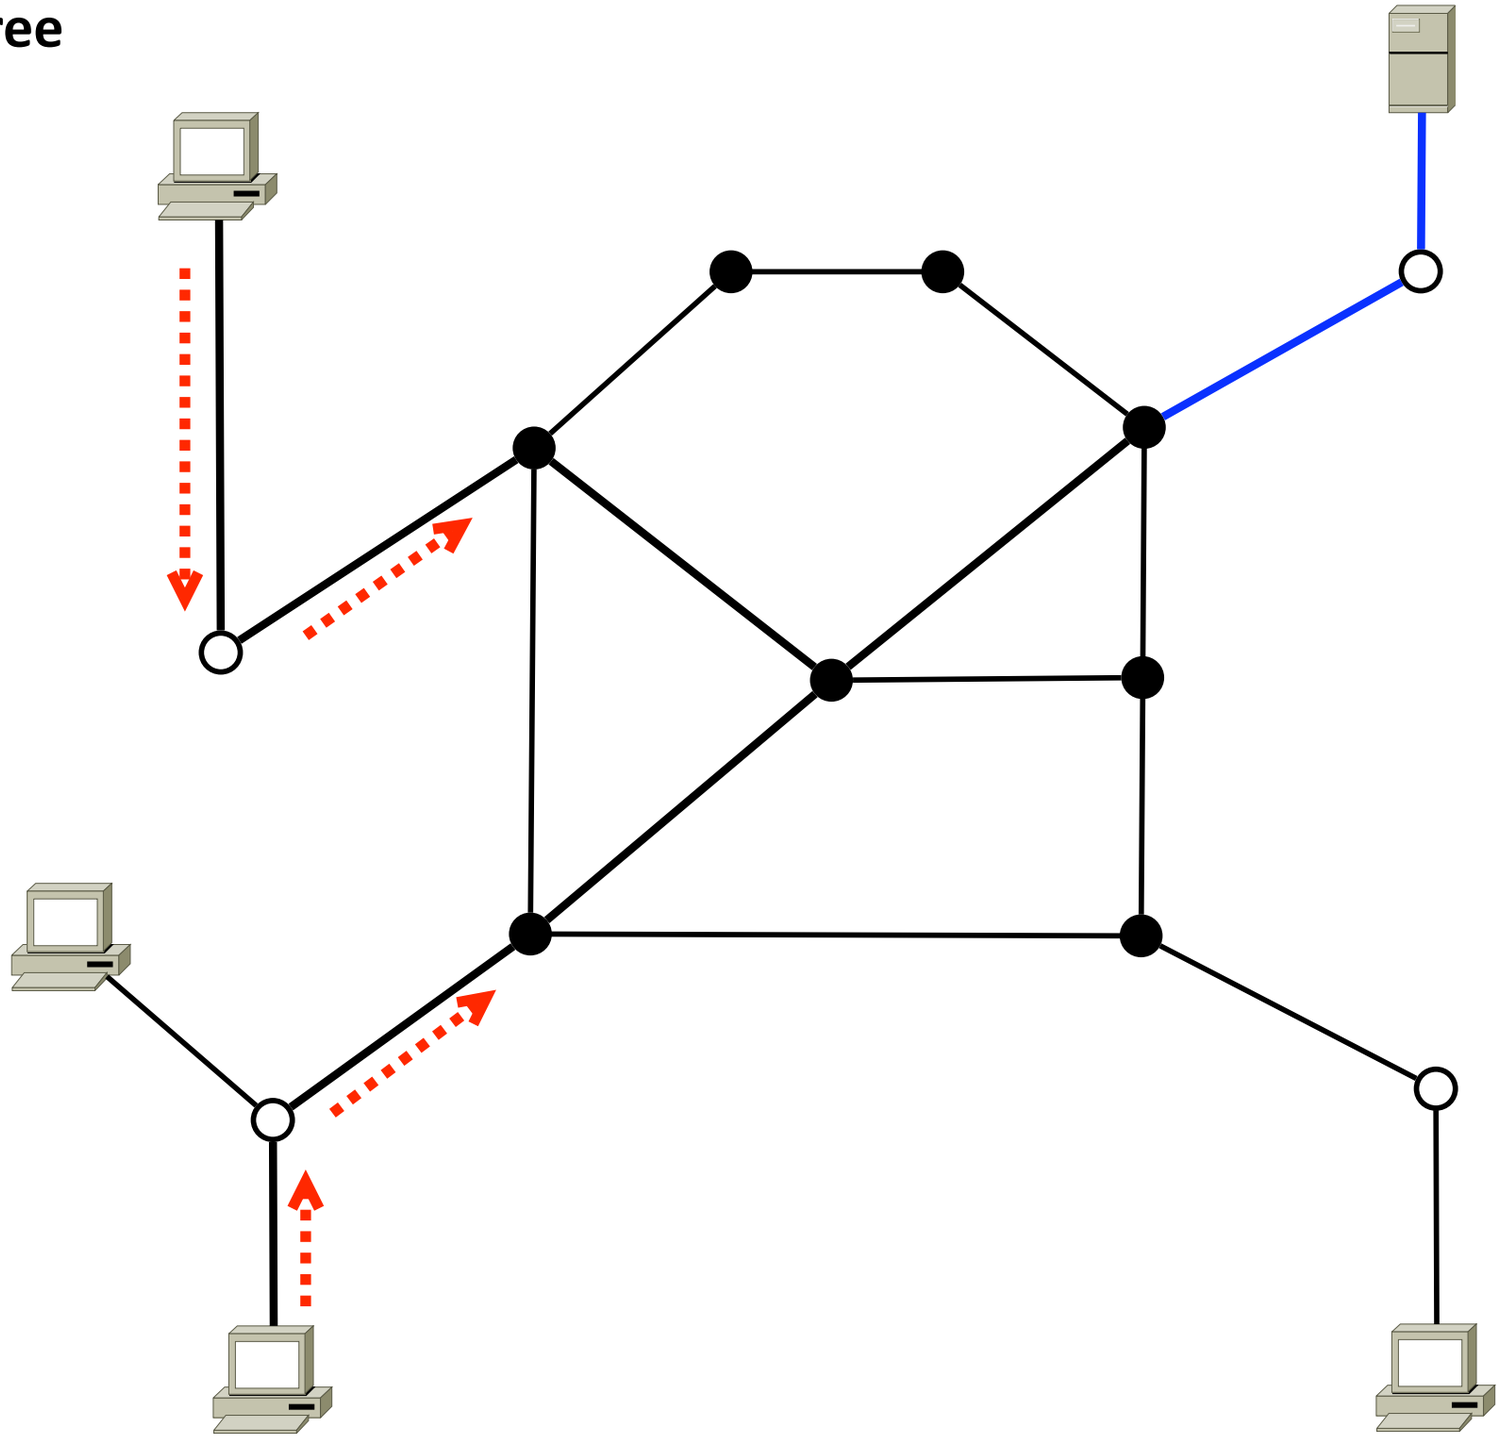

Multicast tree

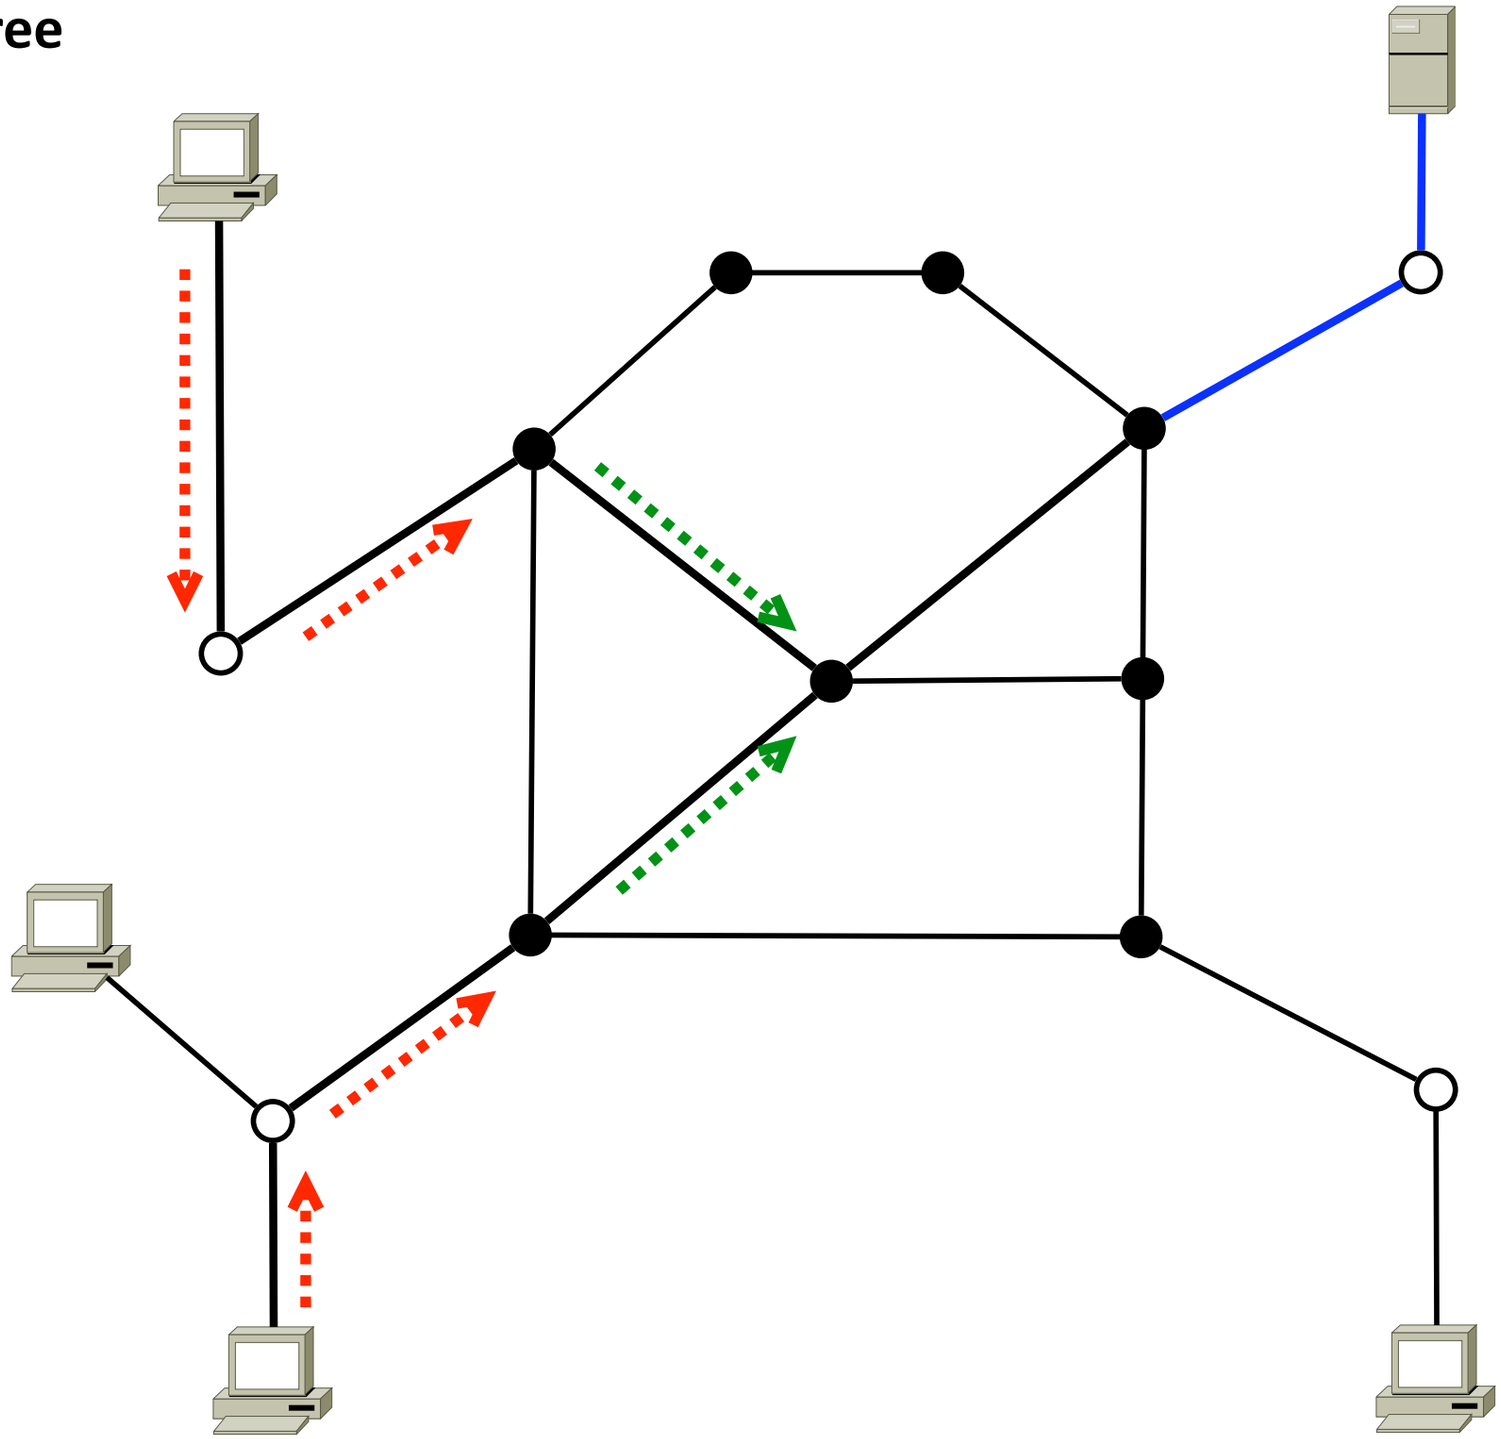

Multicast tree

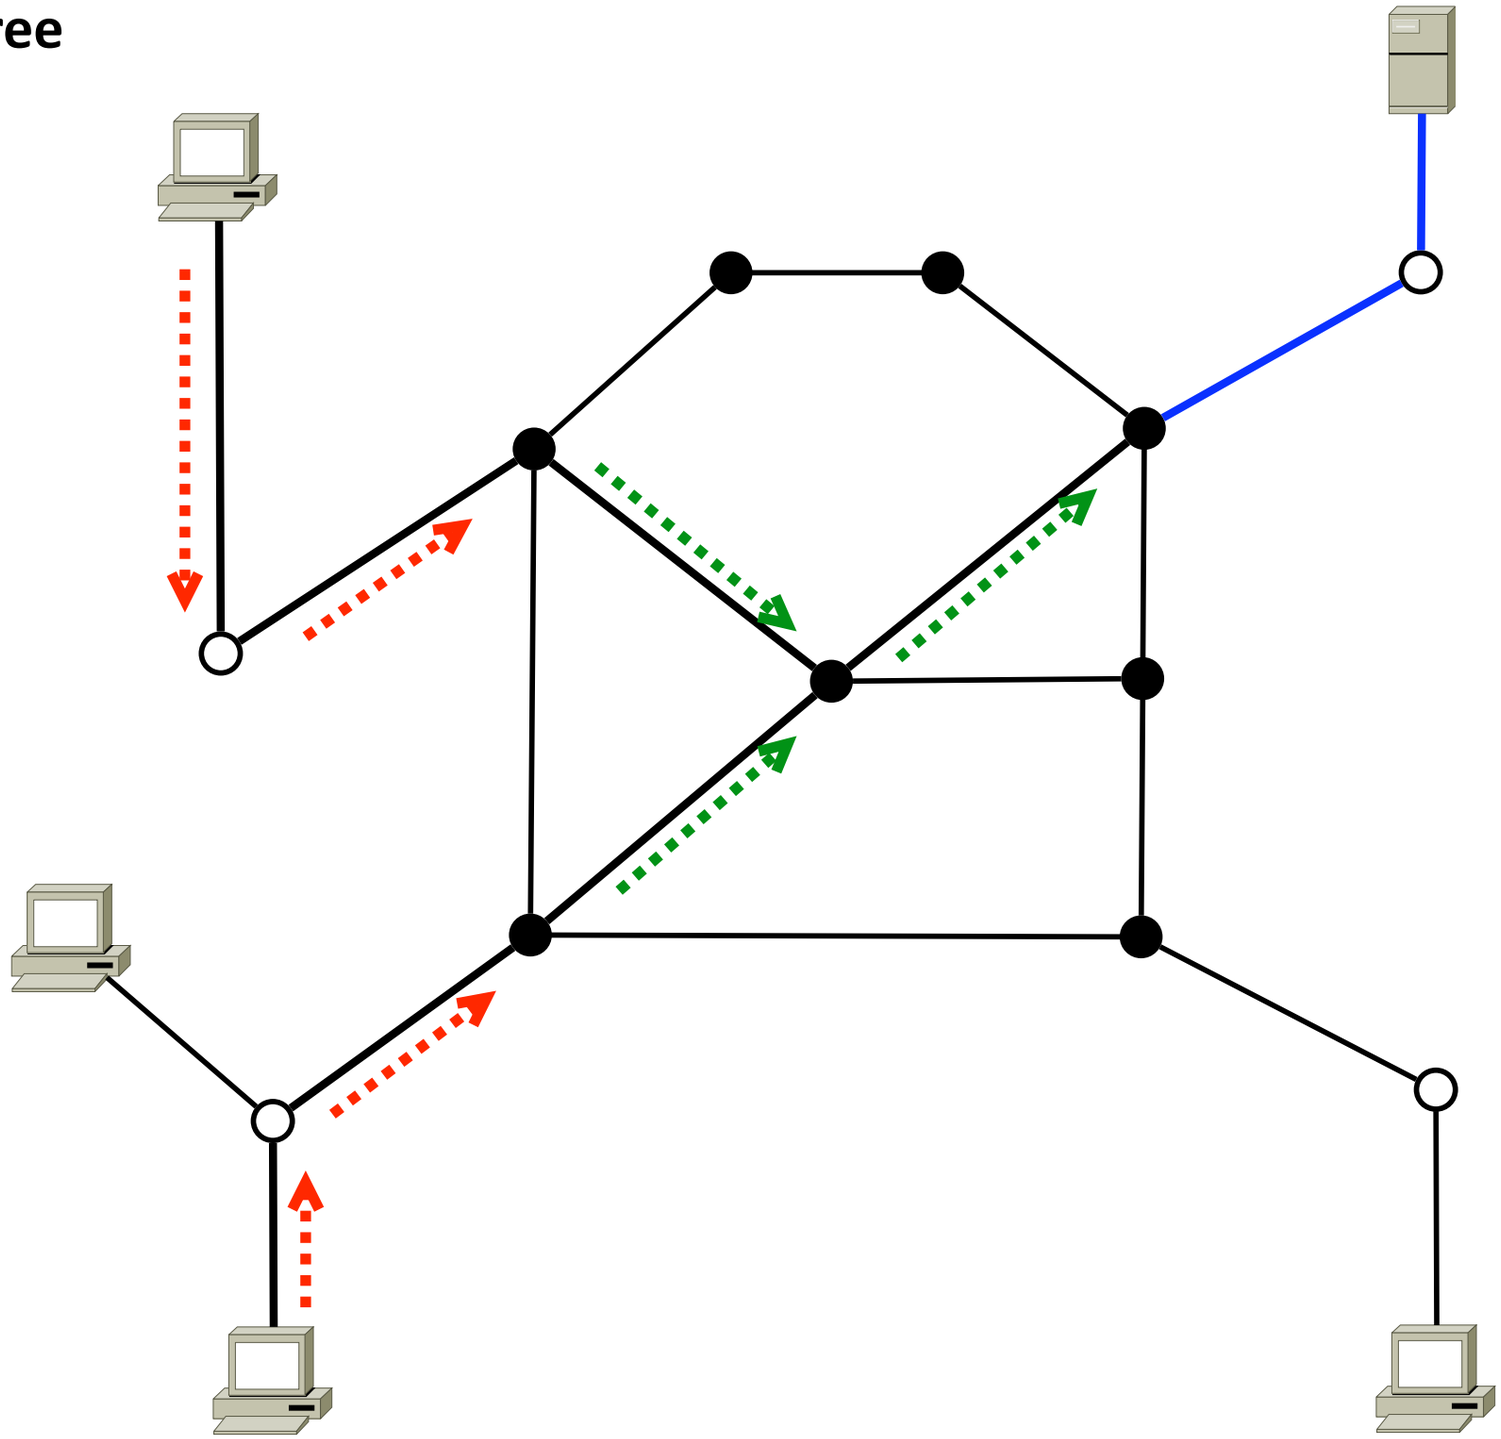

Multicast tree

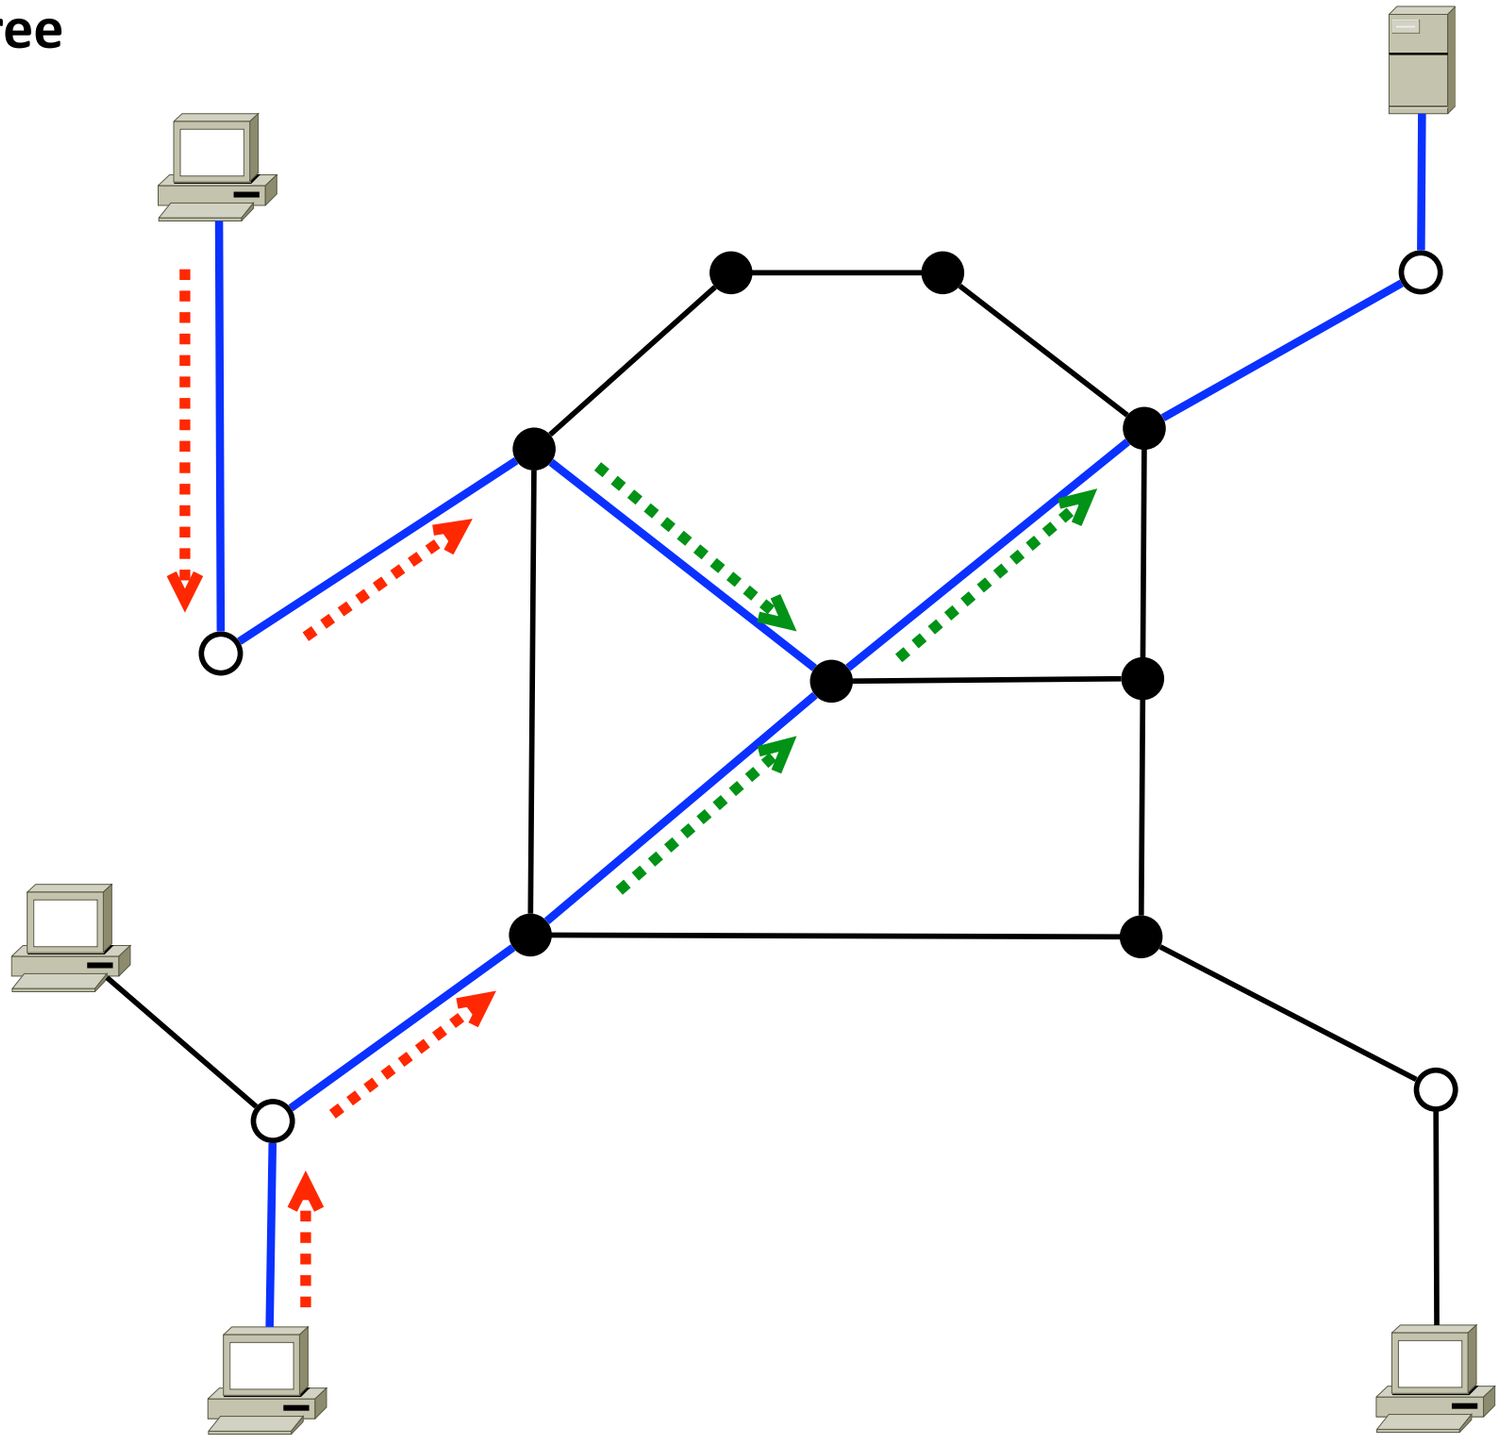

PIM Snooping

- Control of multicast traffic
- Opt-in for multicast
- No congestion by unwanted multicast traffic

```
router mpls
  vpls isp 601
    multicast passive
    multicast pimsm-snooping
wr mem
```

PIM Proxy

- Addition to PIM Snooping
- Blocking and aggregation of PIM messages

Scalability

Where in a VPLS environment are multicast packets being replicated? Which effect does this have on scalability?

Does PIM-snooping use the VPLS path information to direct the multicast streams only to the source PE?

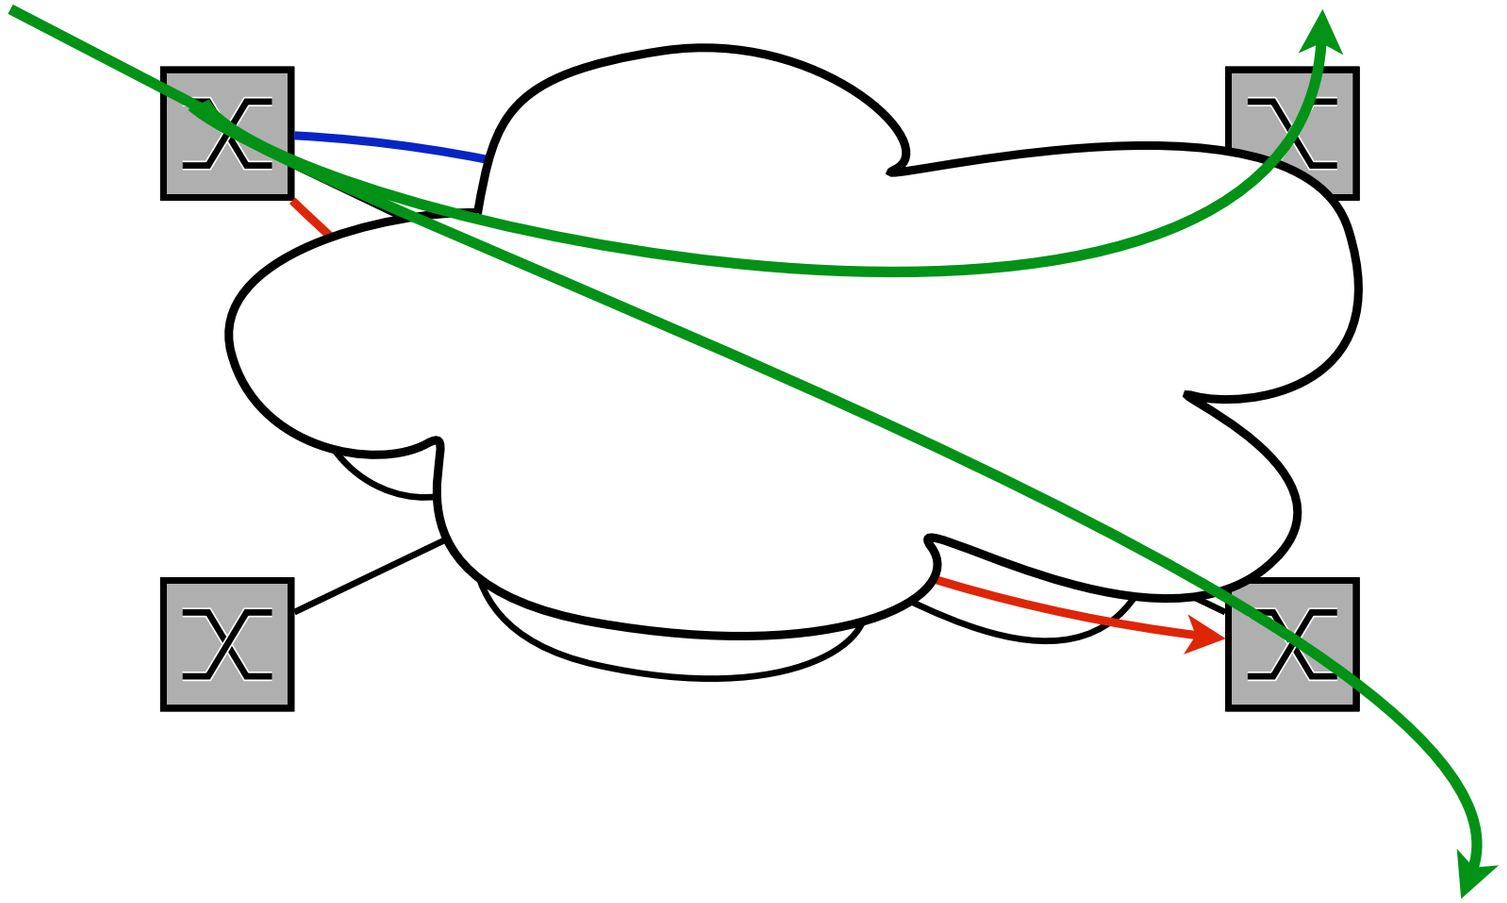

Resource usage

Which problems can be expected when enabling PIM-snooping on the AMS-IX in terms of routing, load, performance and availability?

- Effect on routing protocols
- Memory usage
- CPU usage

Implementation

Which metric does PIM-snooping on the PE's use for switching the multicast stream? Does this cause problems?

$$IPv4 : MAC \leftarrow 01005E000000_{16} \vee (group \wedge 7FFFFF_{16})$$
$$IPv6 : MAC \leftarrow 333300000000_{16} \vee (group \wedge FFFFFFFF_{16})$$

What is the difference between ASM and SSM in respect to the AMS-IX?

Abuse

What will happen when a PE receives an excessive amount of join messages?

How do the switches react on unexpected PIM messages?

Demo

Crashing an internet exchange with a single packet

attilla@laborantix: ~ — ssh — 83×5

```
laborantix:/home/yuri#
```

attilla@laborantix: ~ — ssh — 36×11

Terminal — telnet — 45×11

```
bambix#
bambix#
bambix#
bambix#
bambix#
bambix#
bambix#
bambix#
bambix#
bambix#
bambix#
```

Terminal — telnet — 66×20

```
transformix#
transformix#
transformix#
transformix#
transformix#
transformix#
transformix#
transformix#
transformix#
transformix#
transformix#
transformix#
transformix#
transformix#
transformix#
transformix#
transformix#
transformix#
transformix#
transformix#
```

Terminal — telnet — 66×20

```
infinix#
infinix#
infinix#
infinix#
infinix#
infinix#
infinix#
infinix#
infinix#
infinix#
infinix#
infinix#
infinix#
infinix#
infinix#
infinix#
infinix#
infinix#
infinix#
infinix#
```

Terminal — telnet — 66×20

```
larix#
larix#
larix#
larix#
larix#
larix#
larix#
larix#
larix#
larix#
larix#
larix#
larix#
larix#
larix#
larix#
larix#
larix#
larix#
larix#
```

Conclusion

- Scalability
- Resource usage
- Stability

Recommendations

- Do not yet enable PIM-snooping
- Do not use a PIM-proxy

Questions?