IPv6 Multicast

beyond link local

Agenda

- IP Multicast History
- Unicast - Broadcast - Anycast - Multicast
- Multicast in IPv6 Networks
- Case-Study DHCPv6
- Case-Study X11
- Case-Study NTP
- Case-Study DNS
- Conclusion

IP Multicast History

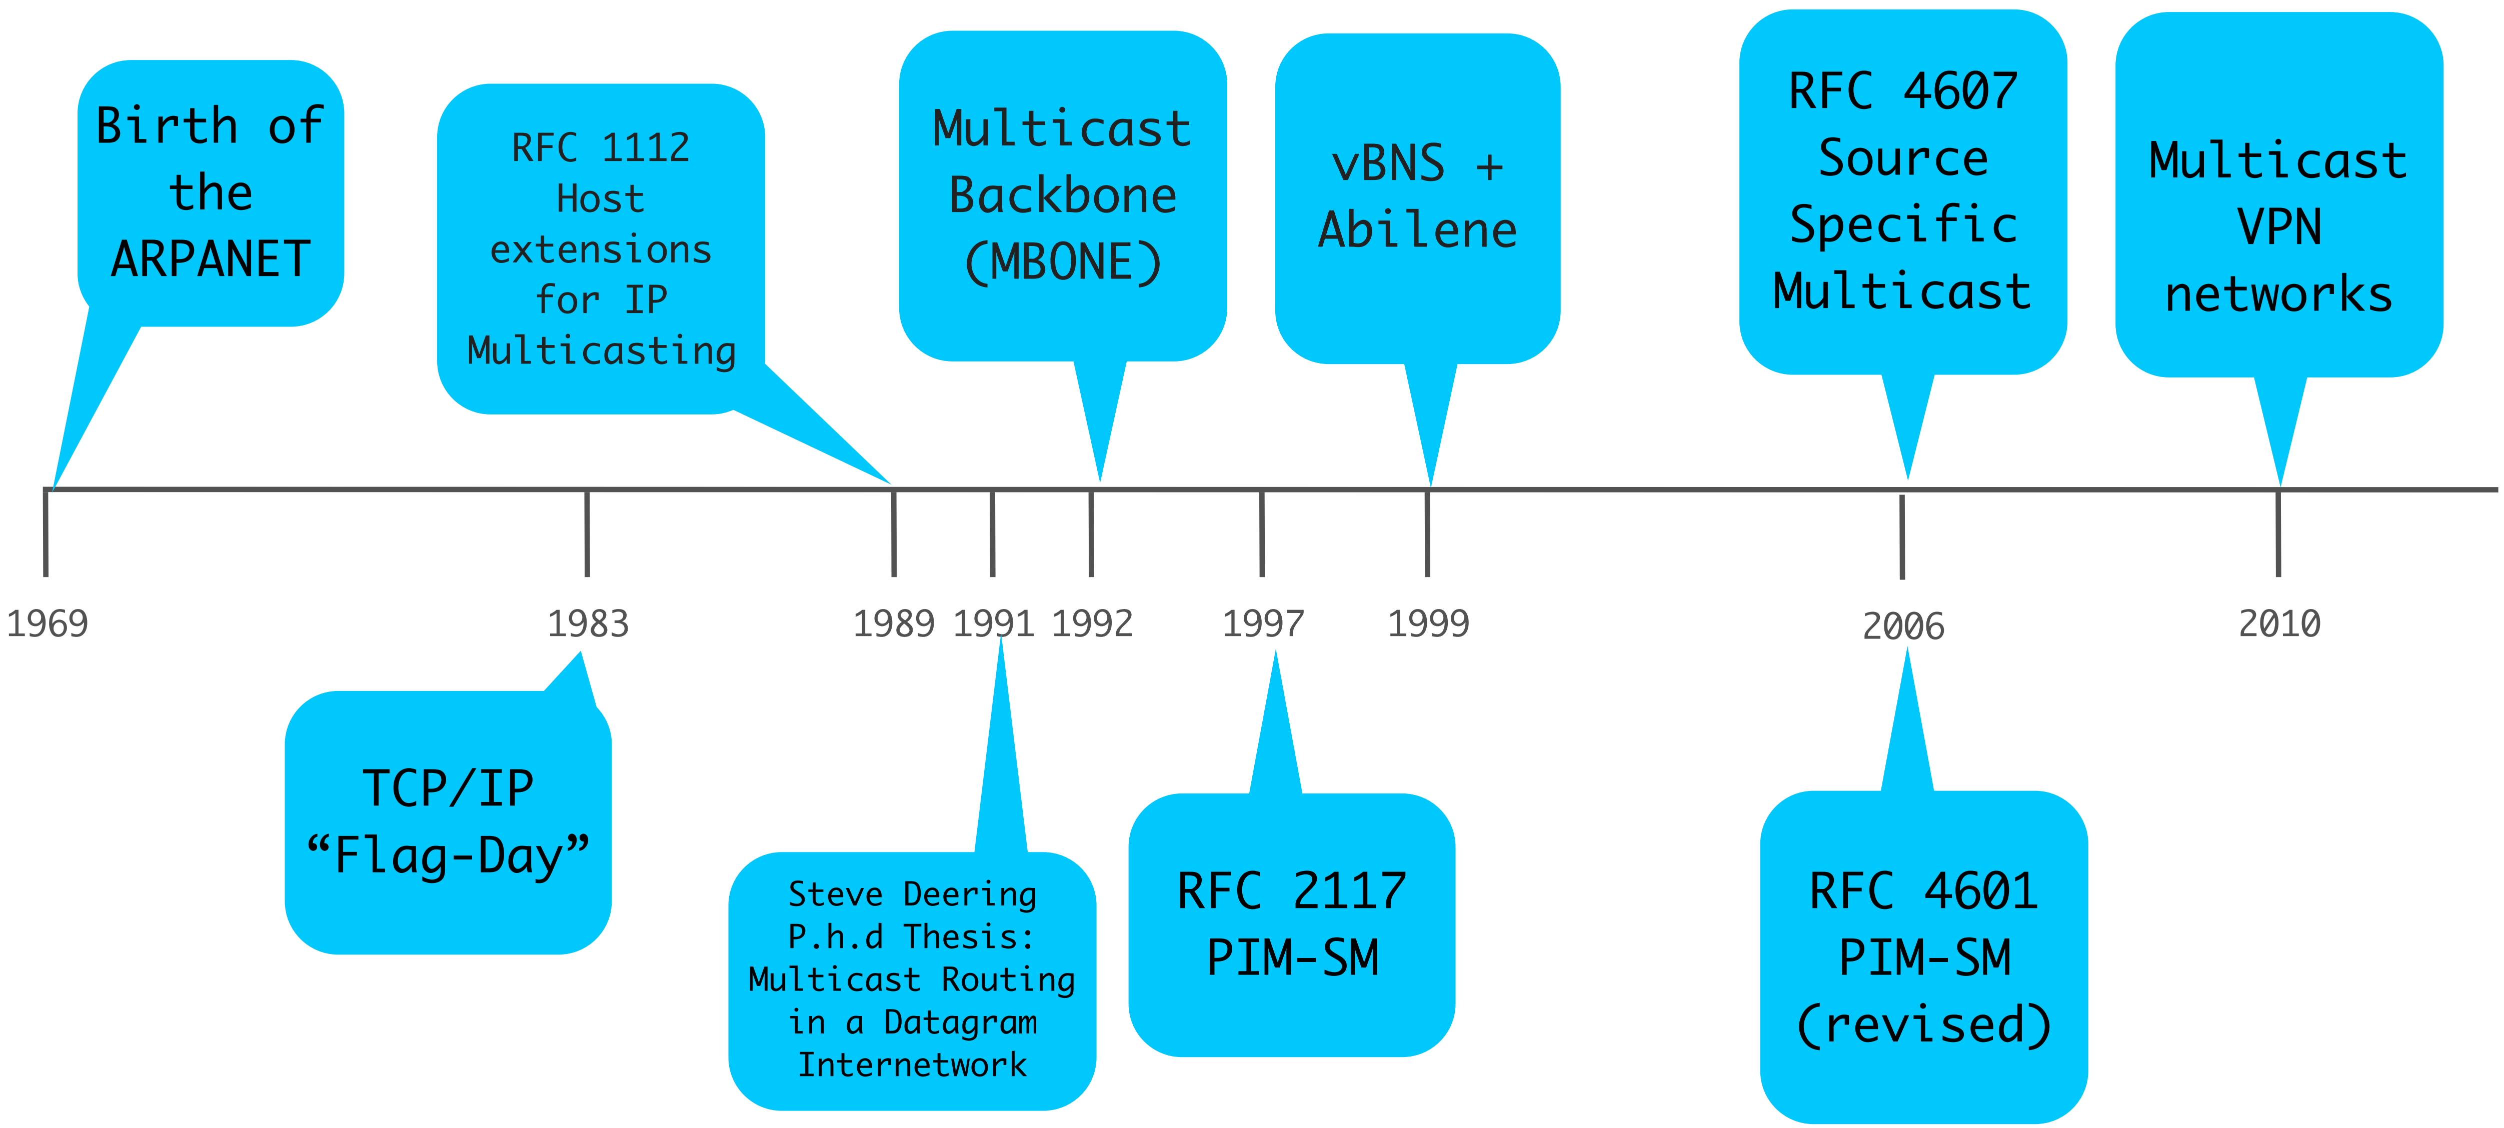

Definitions

- Unicast: Node to Node (1-1) communication
- Broadcast: Node to Network communication
- Anycast: Note to one Node from a group communication
- Multicast: Node to all Nodes from a group communication

Unicast

Broadcast

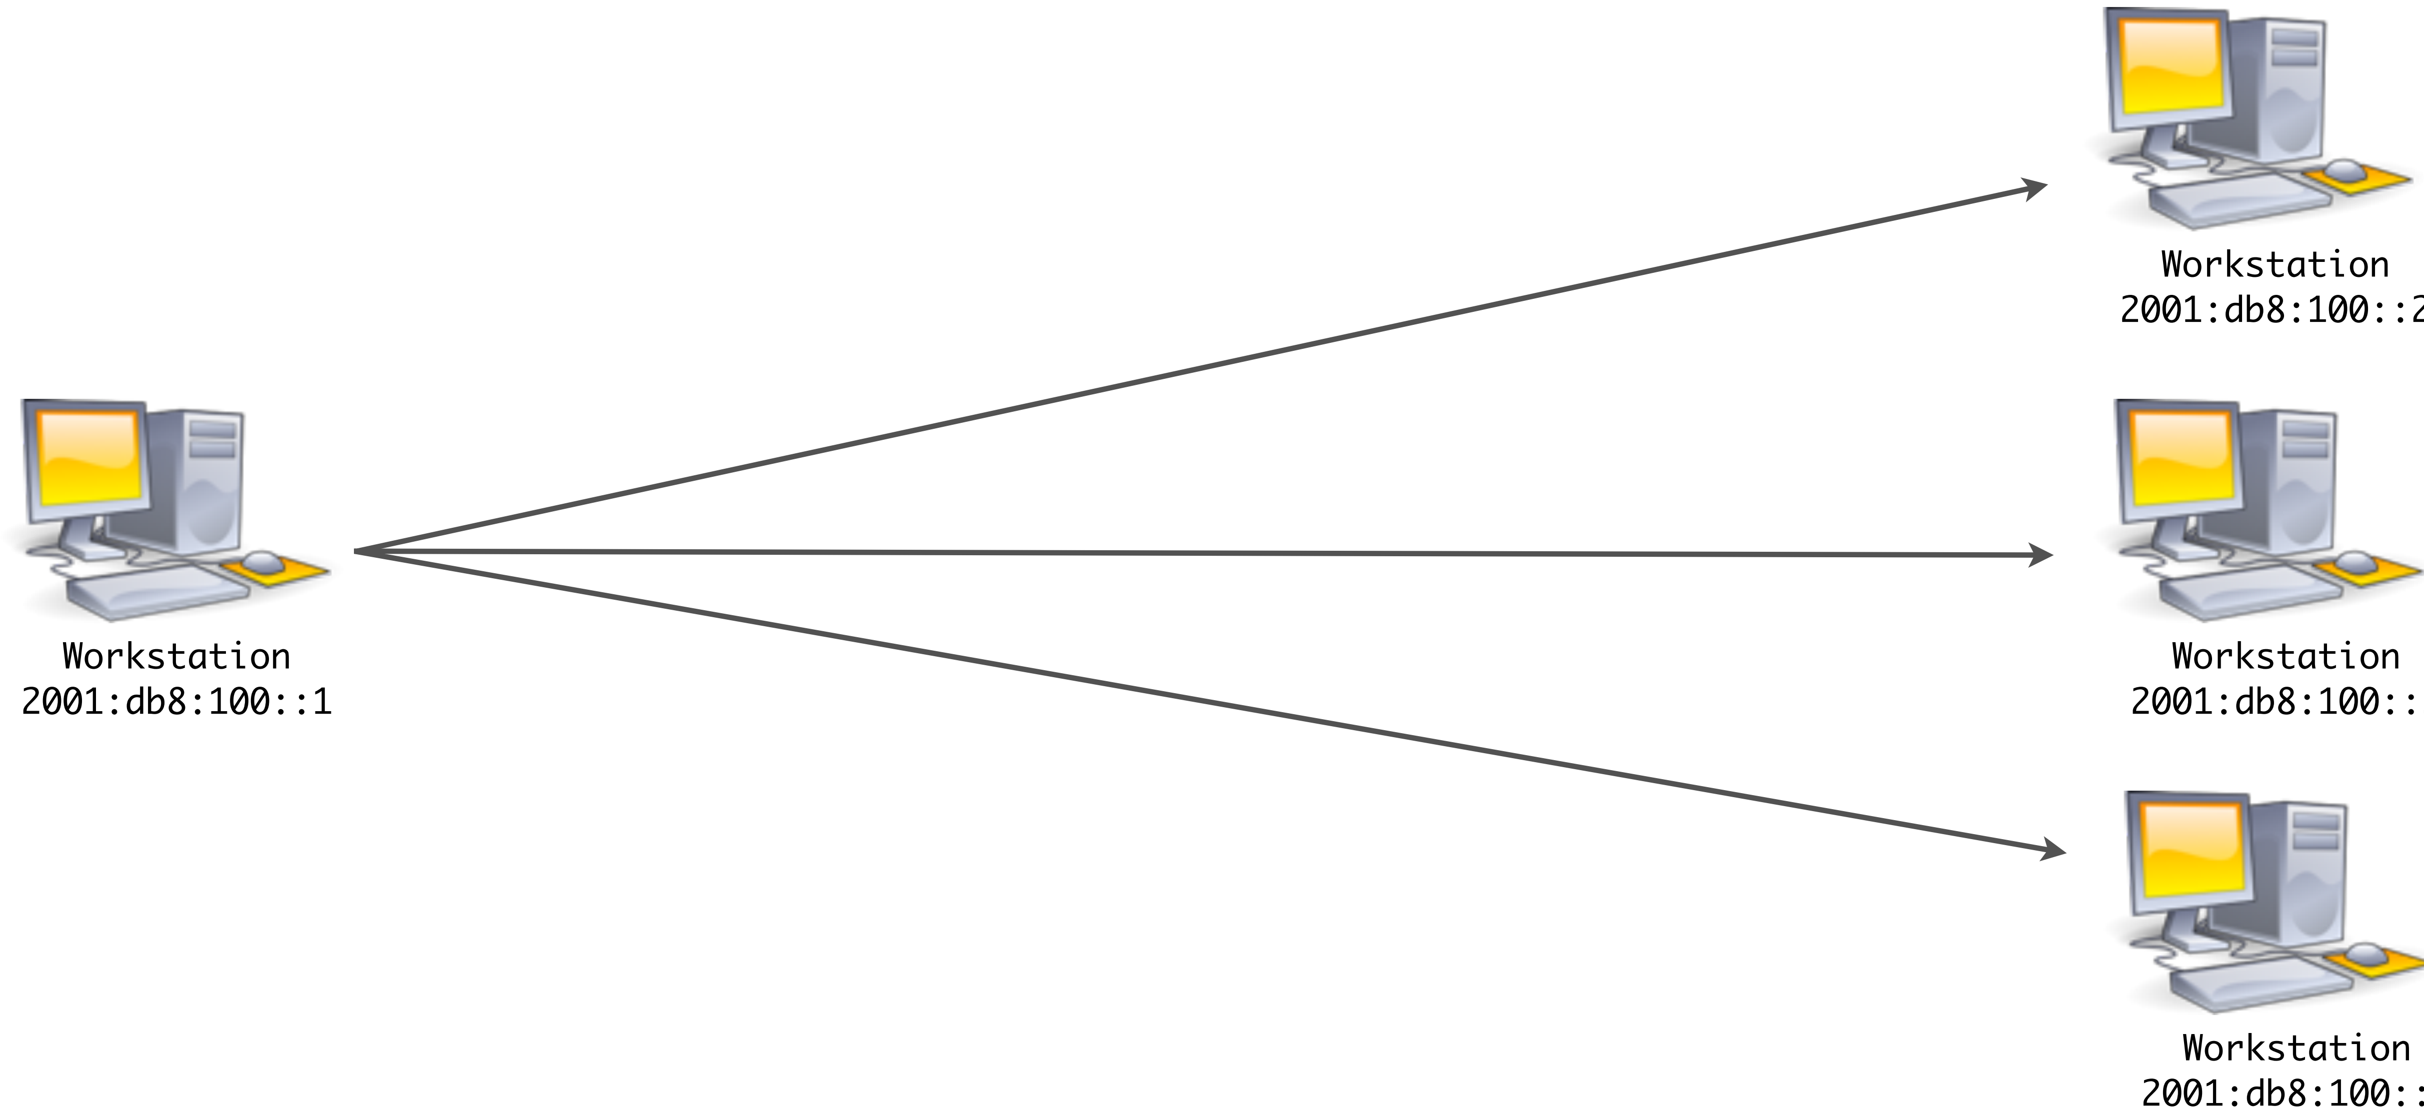
Workstation
2001:db8:100::1
Workstation
2001:db8:100::2
Workstation
2001:db8:100::3
Workstation
2001:db8:100::4

Anycast

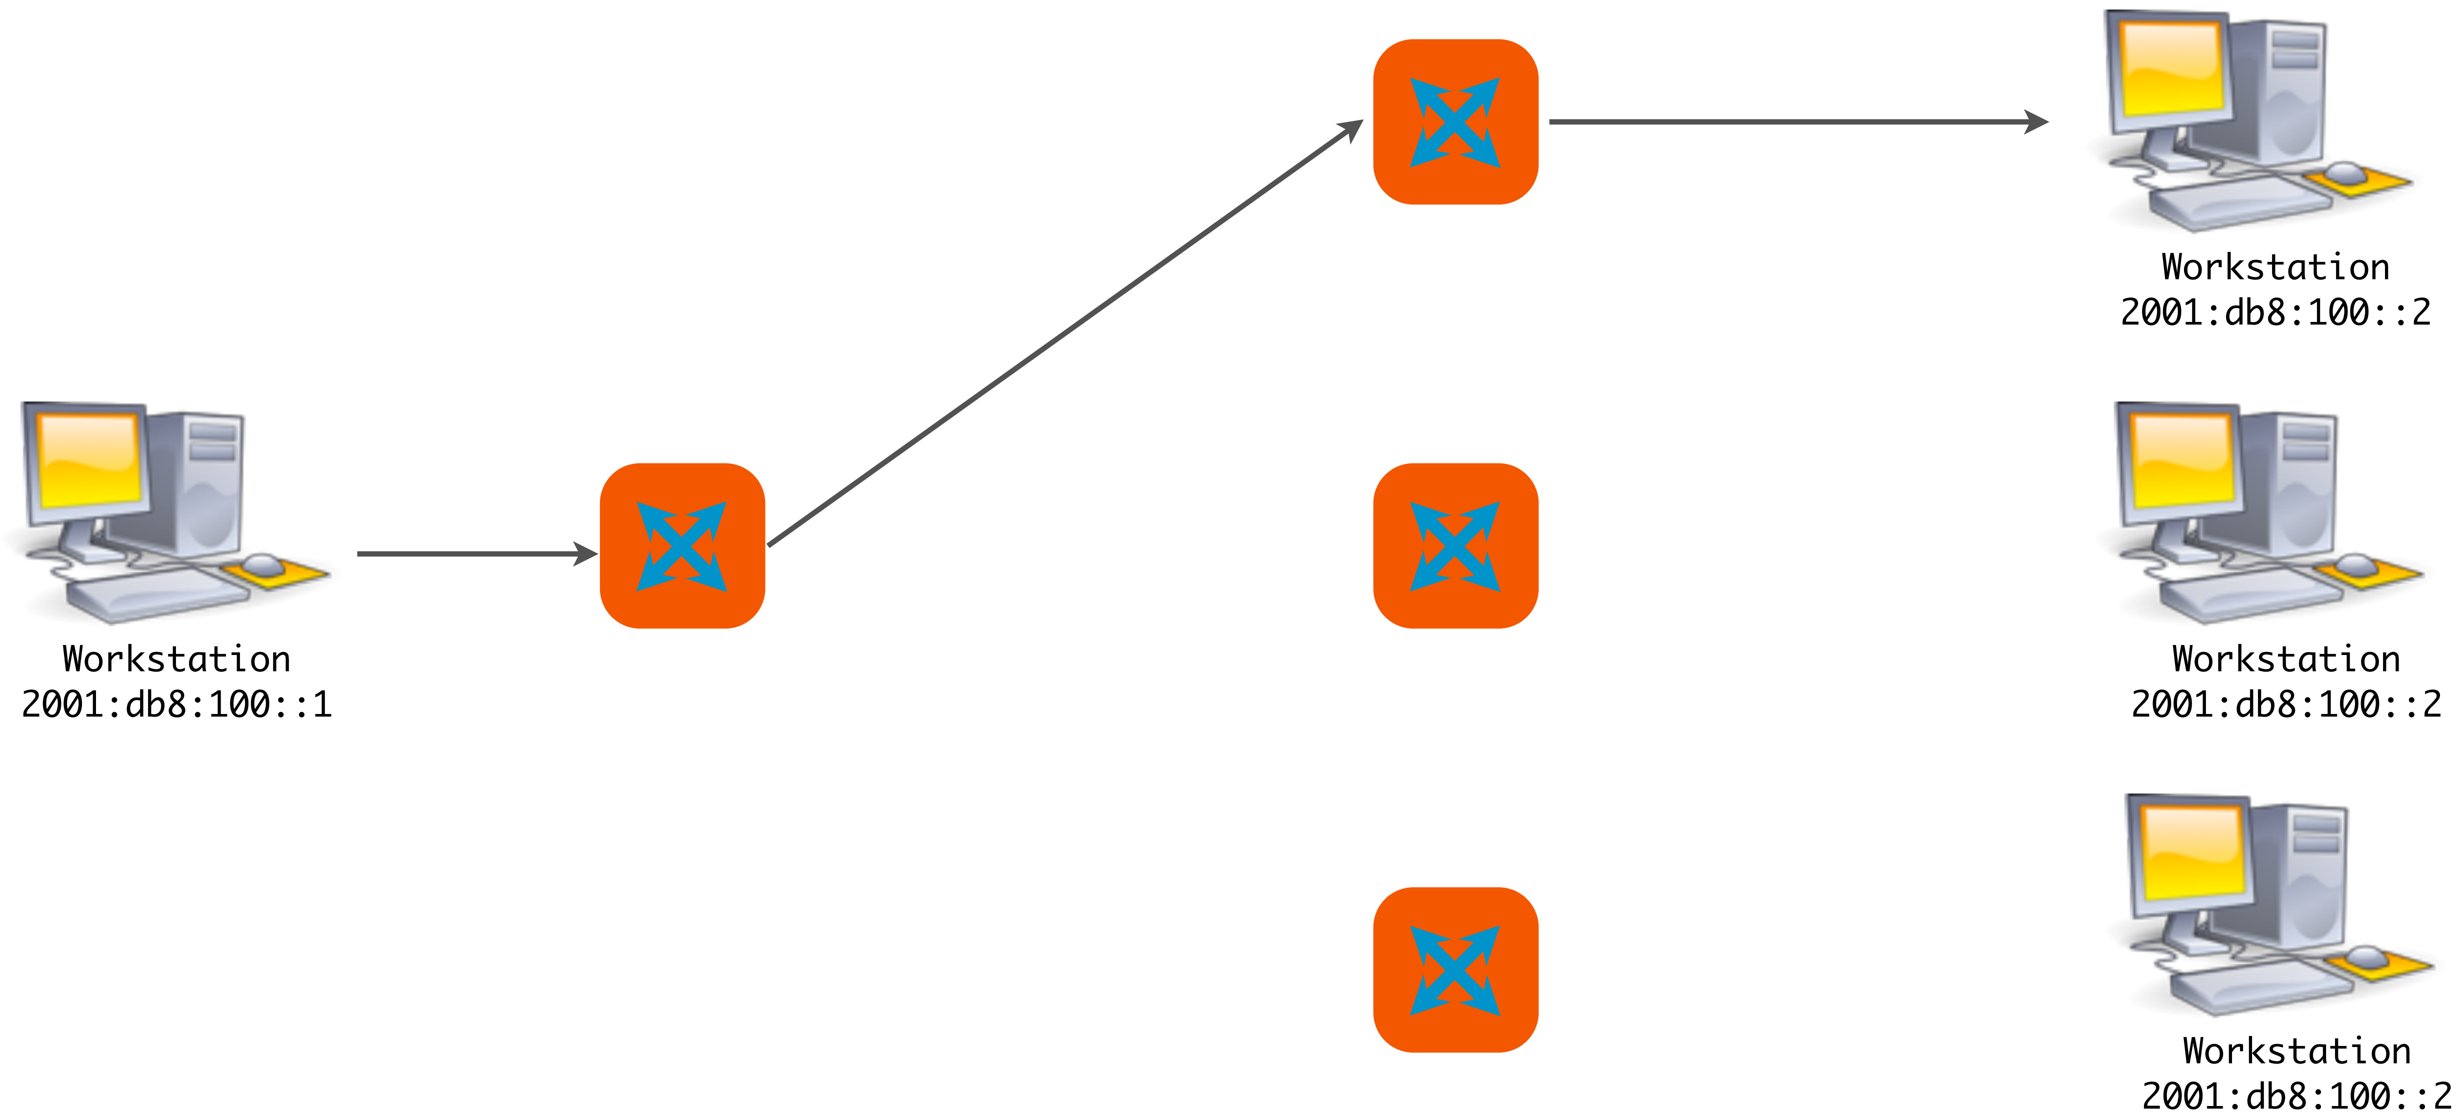

Workstation
2001:db8:100::1
Workstation
2001:db8:100::2
Workstation
2001:db8:100::2
Workstation
2001:db8:100::2

Multicast

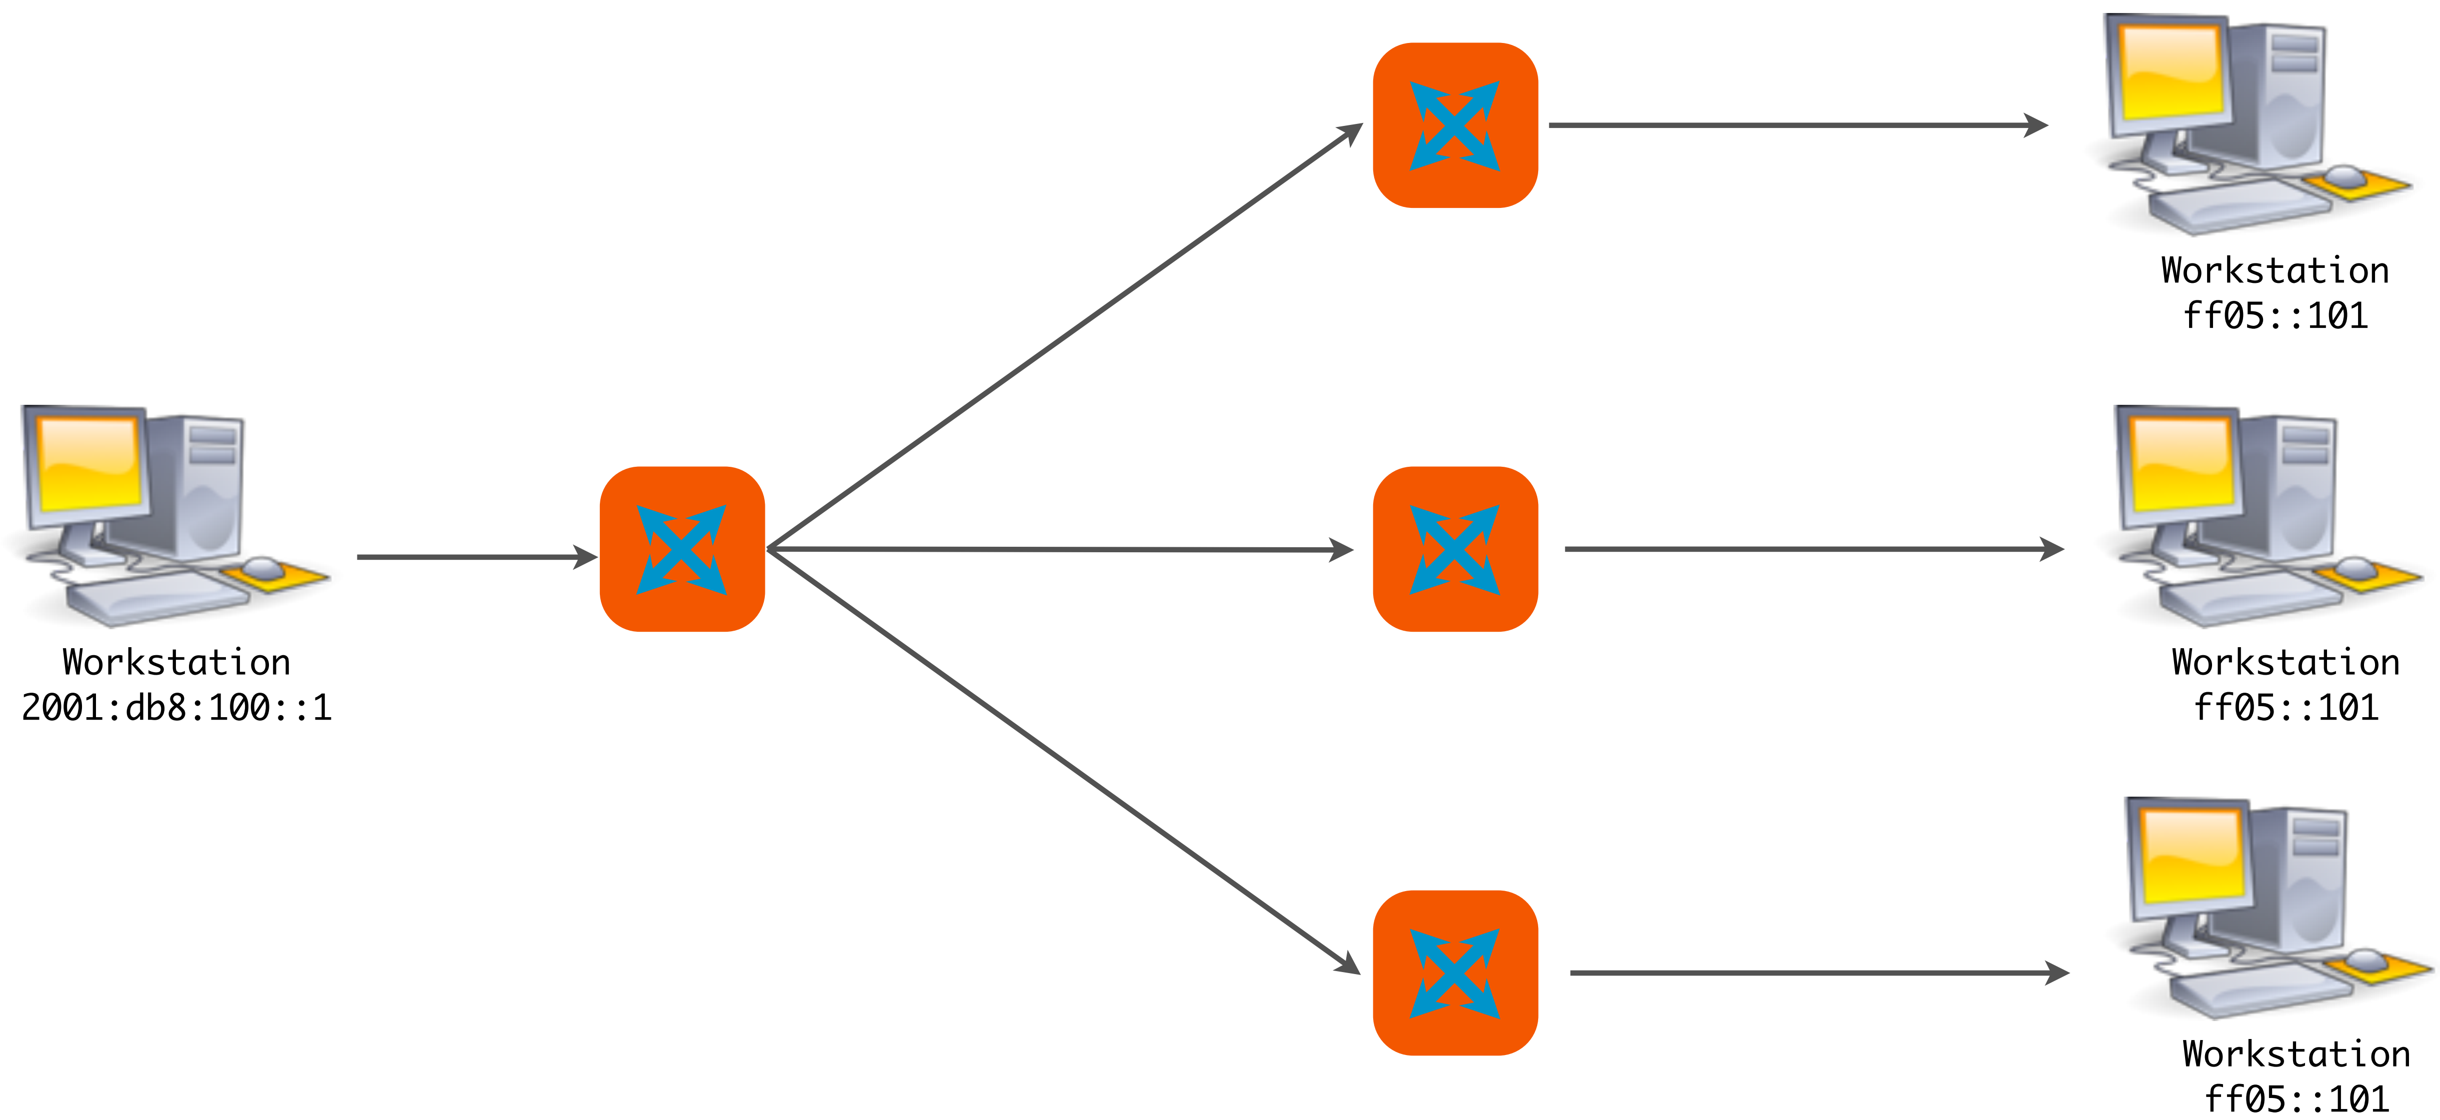

Workstation
ff05::101
Workstation
2001:db8:100::1
Workstation
ff05::101
Workstation
ff05::101

Multicast in IPv6 Networks

- Multicast is a mandatory part of IPv6
  - IPv6 is using Multicast on the protocol level (e.g. Neighborhood Discovery)
  - use of IPv6 Multicast beyond the link-local scope requires a multicast routing protocol
- official IANA Multicast IP Addresses:
  `http://www.iana.org/assignments/ipv6-multicast-addresses/ipv6-multicast-addresses.xml`

Multicast IPv6 Addresses

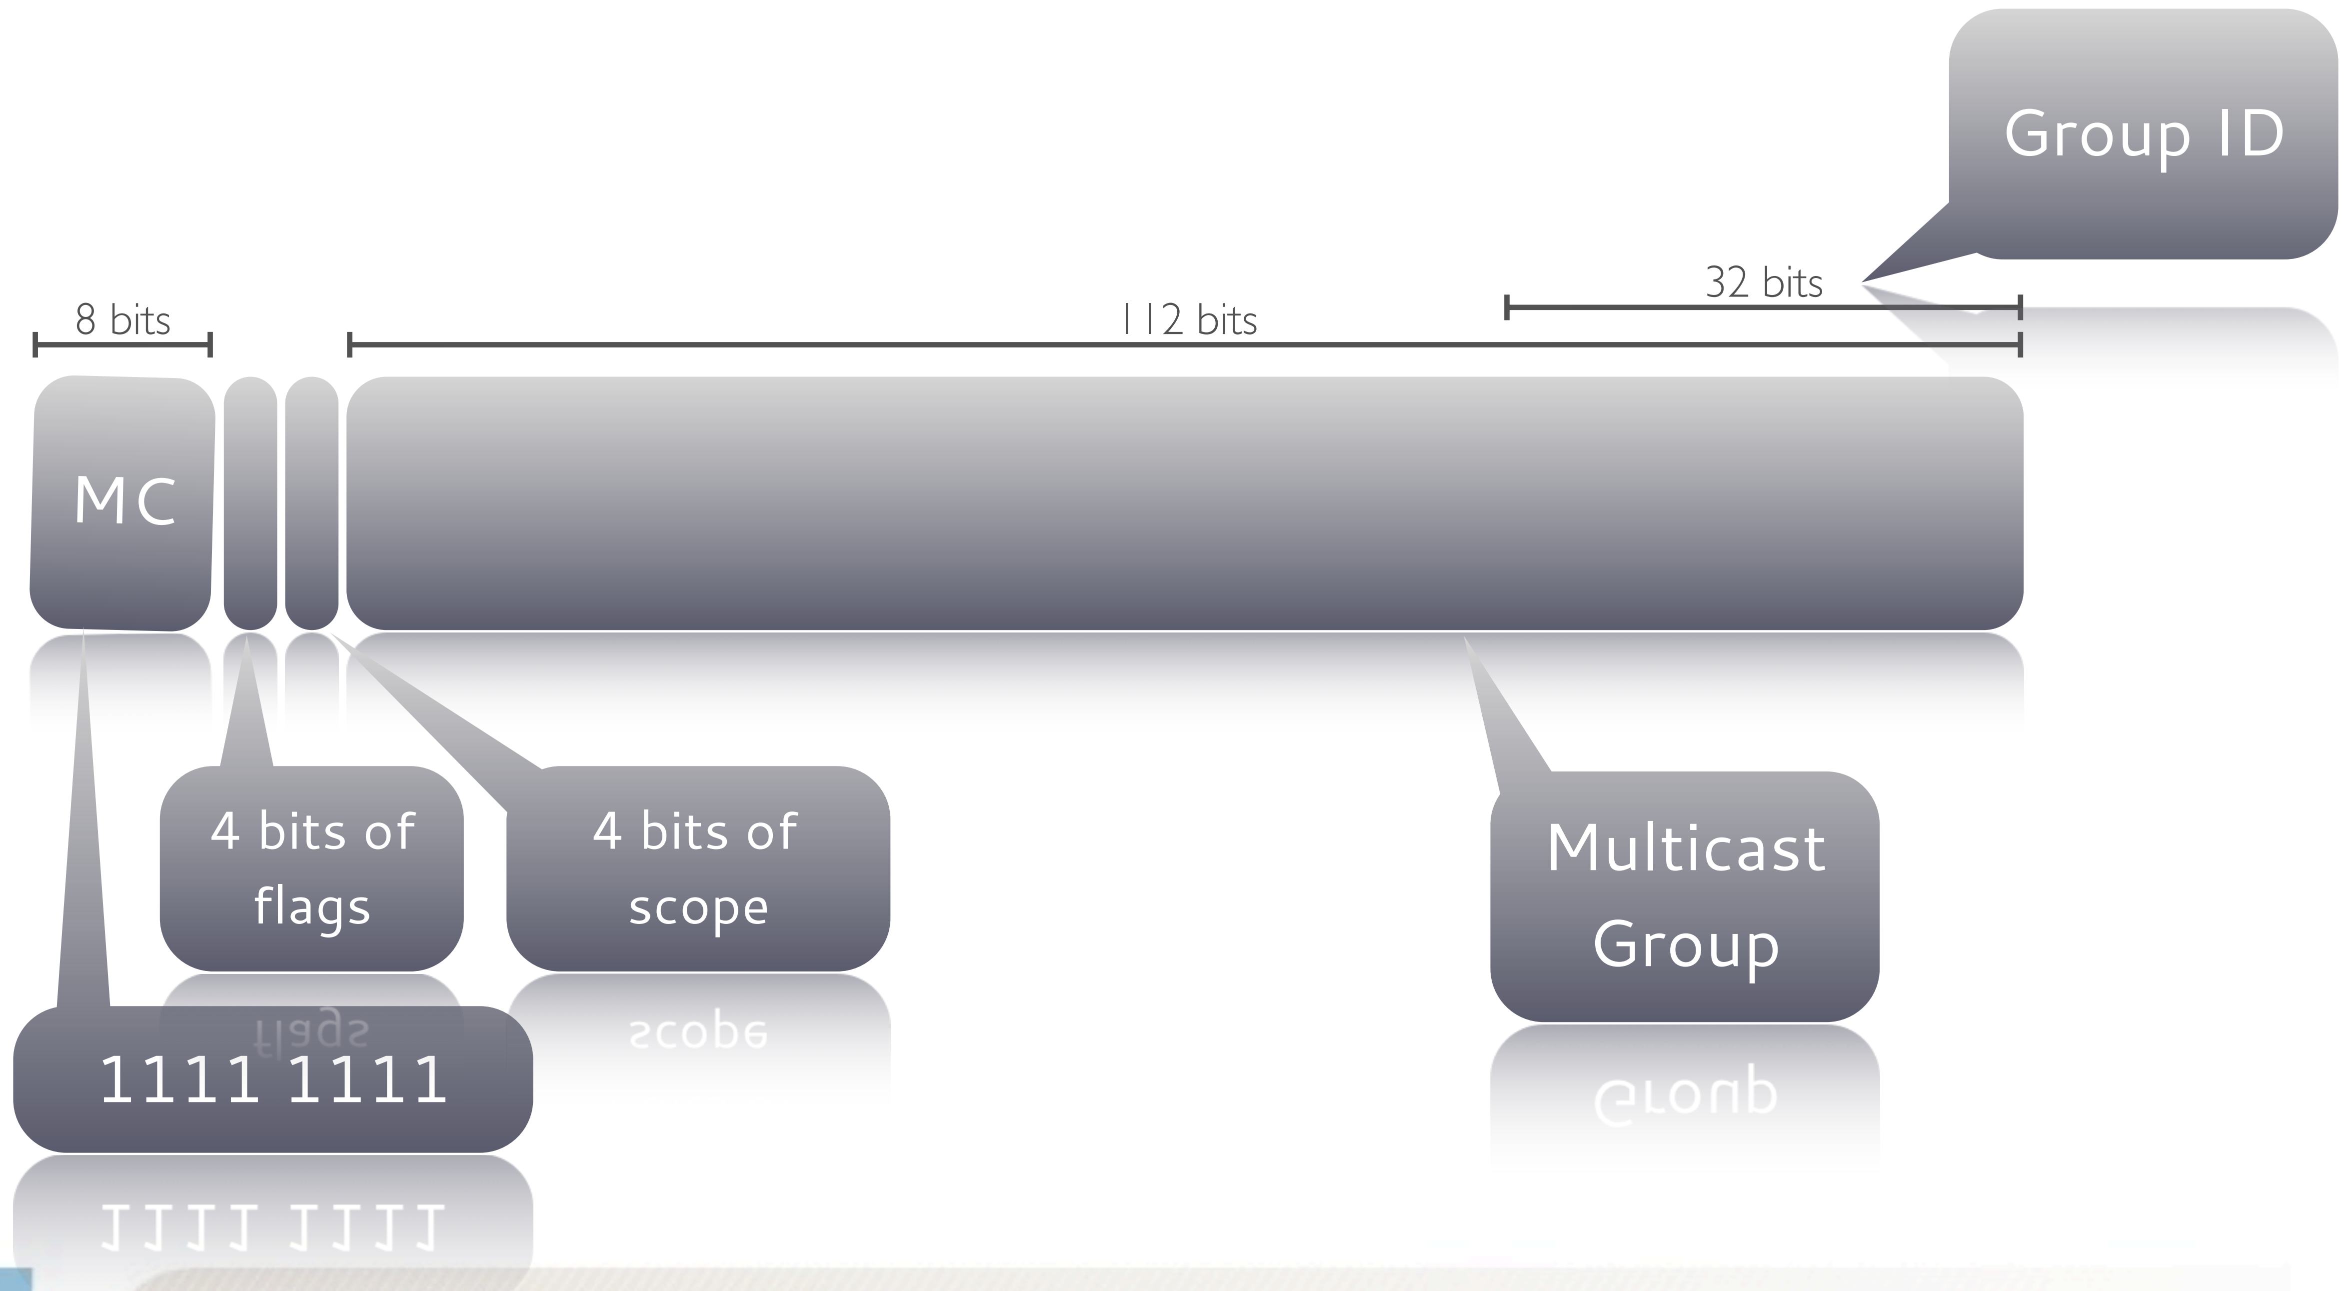

© Men & Mice http://menandmice,com

Multicast Addresses: Flags

Bit	Usage
0	Transient Flag: 0 = permanent address (IANA assigned, well-known) 1 = transient address
1	Prefix flag (RFC 3306/3307)
2	Rendezvous Point (RFC 3956)
3	undefined

Multicast Addresses: Scope

Value (Hex)	Scope
1	Interface/Node-local
2	Link-local
4	Admin-local
5	Site-local
8	Organization-local
E	Global

Node local Scope

Address	Use
ff01::1	all nodes address
ff01::2	all routers address
ff01::fb	mDNSv6

Link local scope

Address	Use
ff02::1	all nodes address
ff02::2	all routers address
ff02::fb	mDNSv6
ff02::F	UPnP Devices
ff02::1:1	Link Name
ff02::1:2	all DHCP agents
ff02::1:3	LLMNR
ff02::101	network time protocol (ntp)
ff02::1:ffXX:XXXX	Solicited-Node Address
ff02::2:ff00::/104	Node Information Queries

Site-Local/variable Scope

Address	Use
ff05::2	all routers address
ff05::fb	mDNSv6
ff05::1:3	all DHCPv6 Servers
ff0x::101	network time protocol (NTP)

Source Specific Multicast - SSM

- once a host joins a multicast group it will receive all traffic send from any machine to this group
  - this can be misused in open networks
  - SSM permits a host to select/deselect one or more multicast sources
  - SSM is based on “Protocol Independent Multicast - Sparse Mode” (PIM-SM) and requires MLDv2 (RFC3810)

Multicast with multiple sender

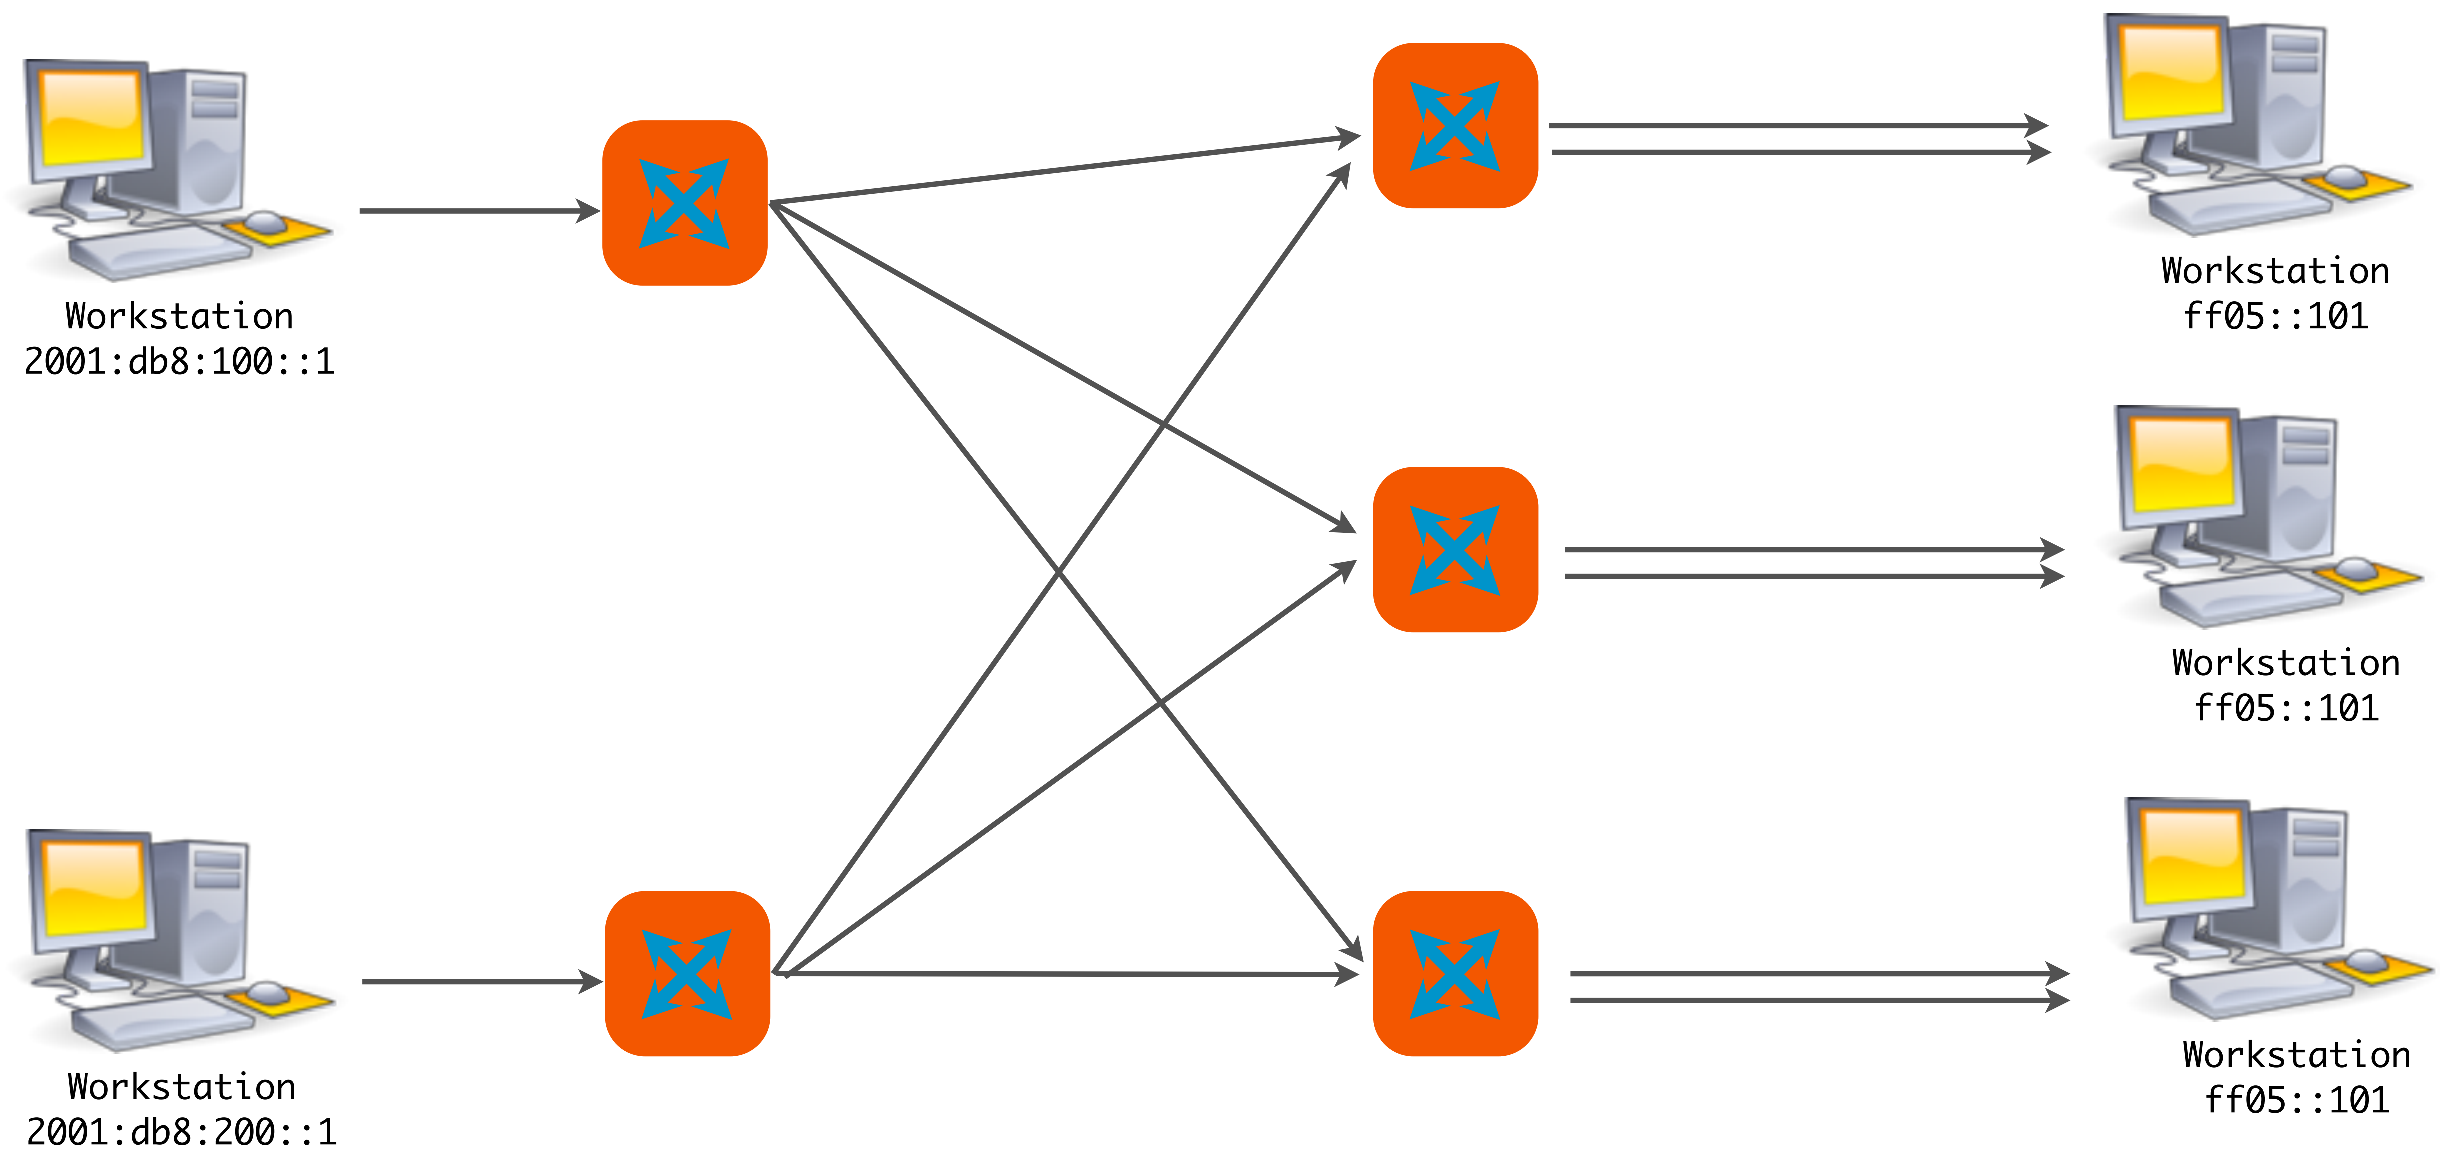

Source Specific Multicast -SSM

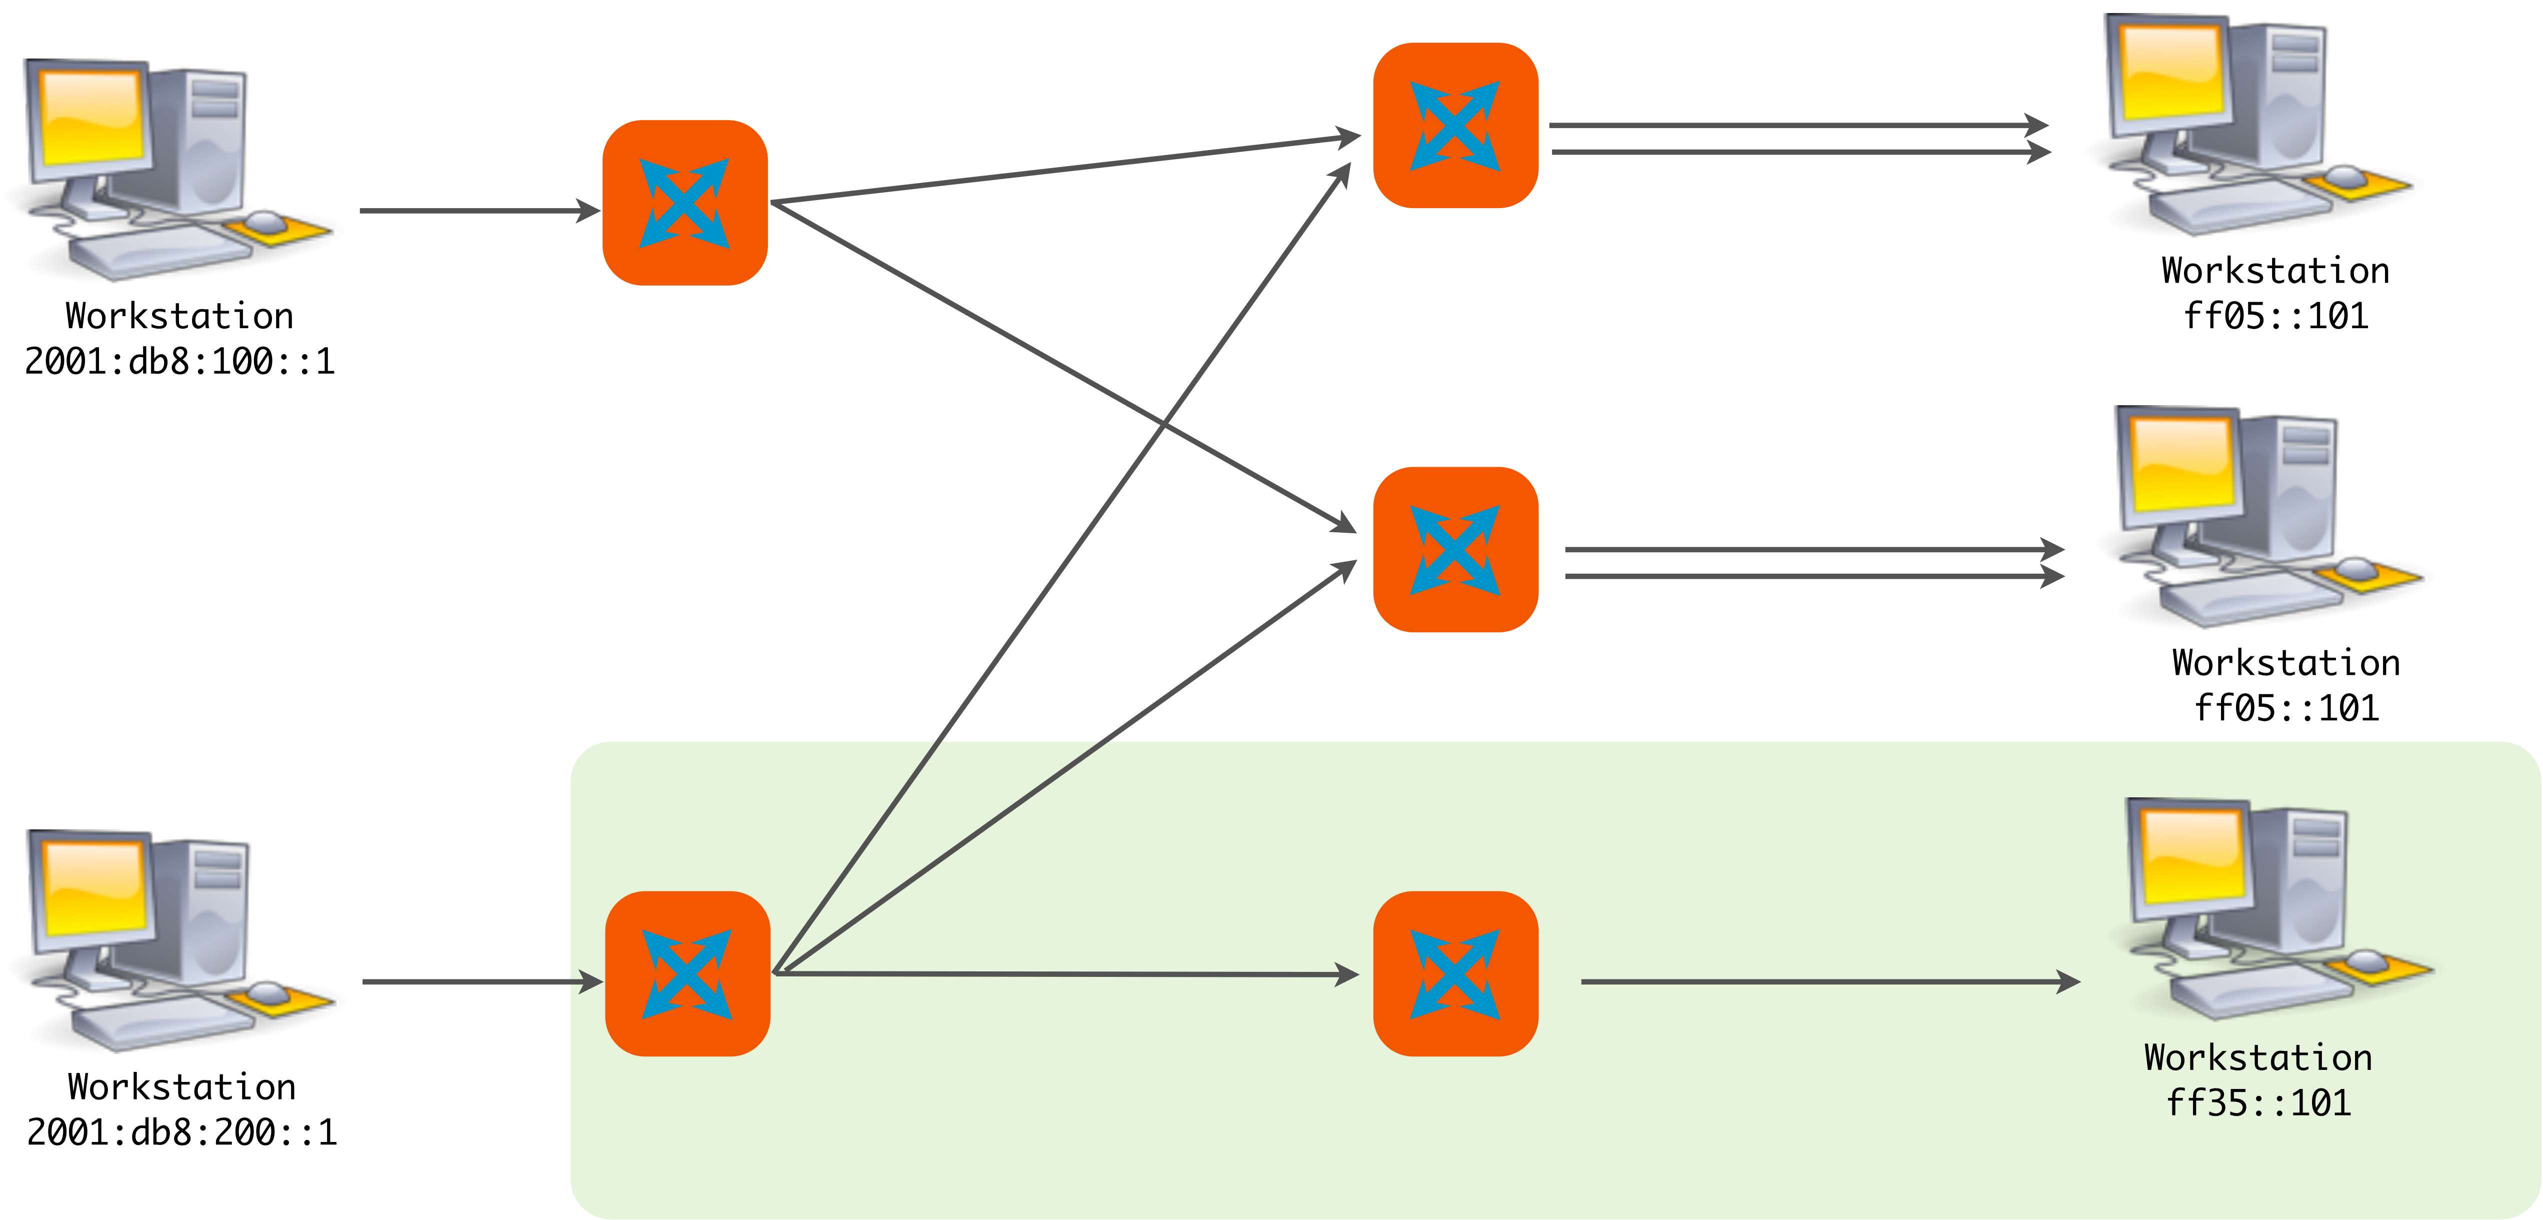

SSM IPv6 Addresses

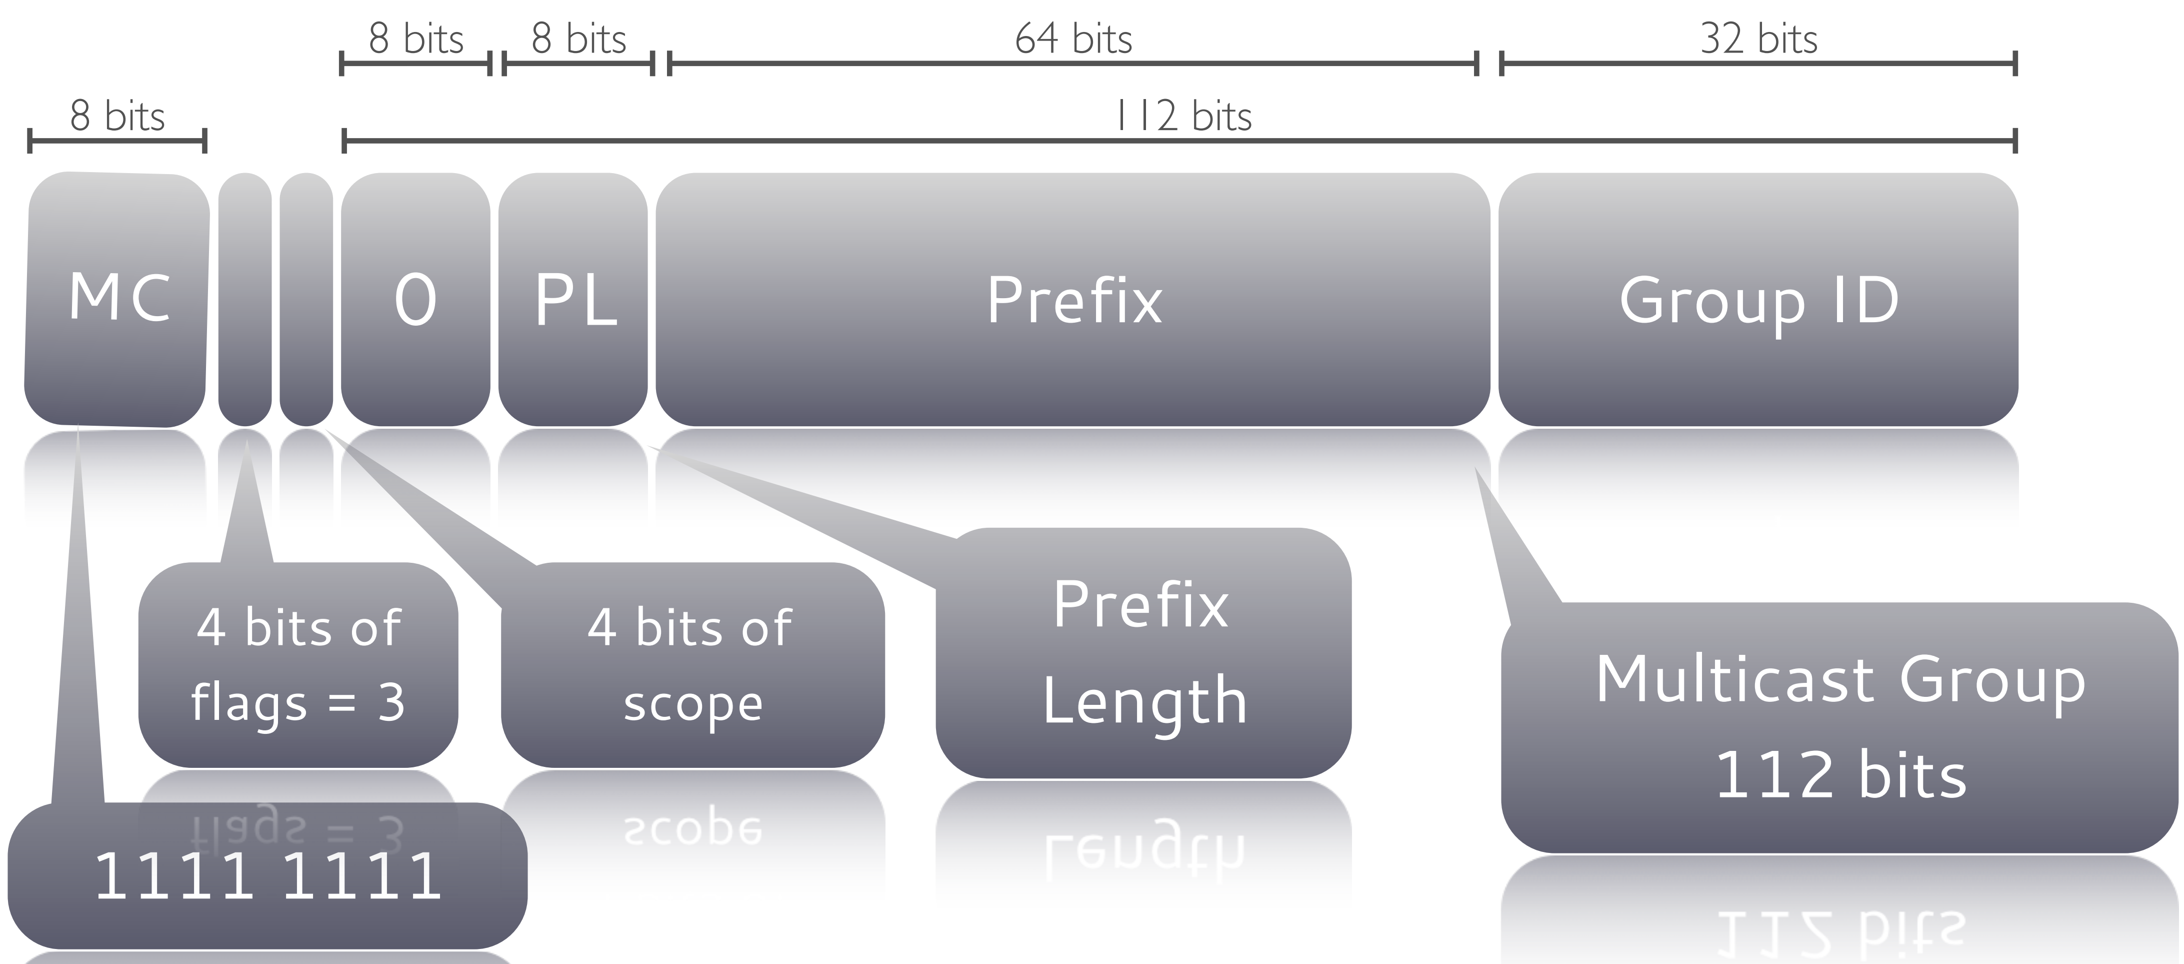

Example SSM IPv6 Address (with unicast prefix)

Prefix Length: 64

ff35:0040:2001:db8:100:0:534e:9ff0

1111 1111

SSM

scope: site

Reserved field: 0

Prefix

Multicast Group ID

Case Studies

- the case-studies look only at site/organization/admin local Multicast
  - not global Internet-Multicast
  - controlled network like enterprise or university campus networks

DHCPv6

- DHCPv6 is often deployed using relay-agents
  - relay-agents can use unicast or multicast
  - unicast configurations have scalability issues in case of a renumbering of DHCPv6 server or once new DHCPv6 servers are introduced
  - a multicast configuration scales without configuration-changes on the relay-agents

DHCPv6 using an unicast relay

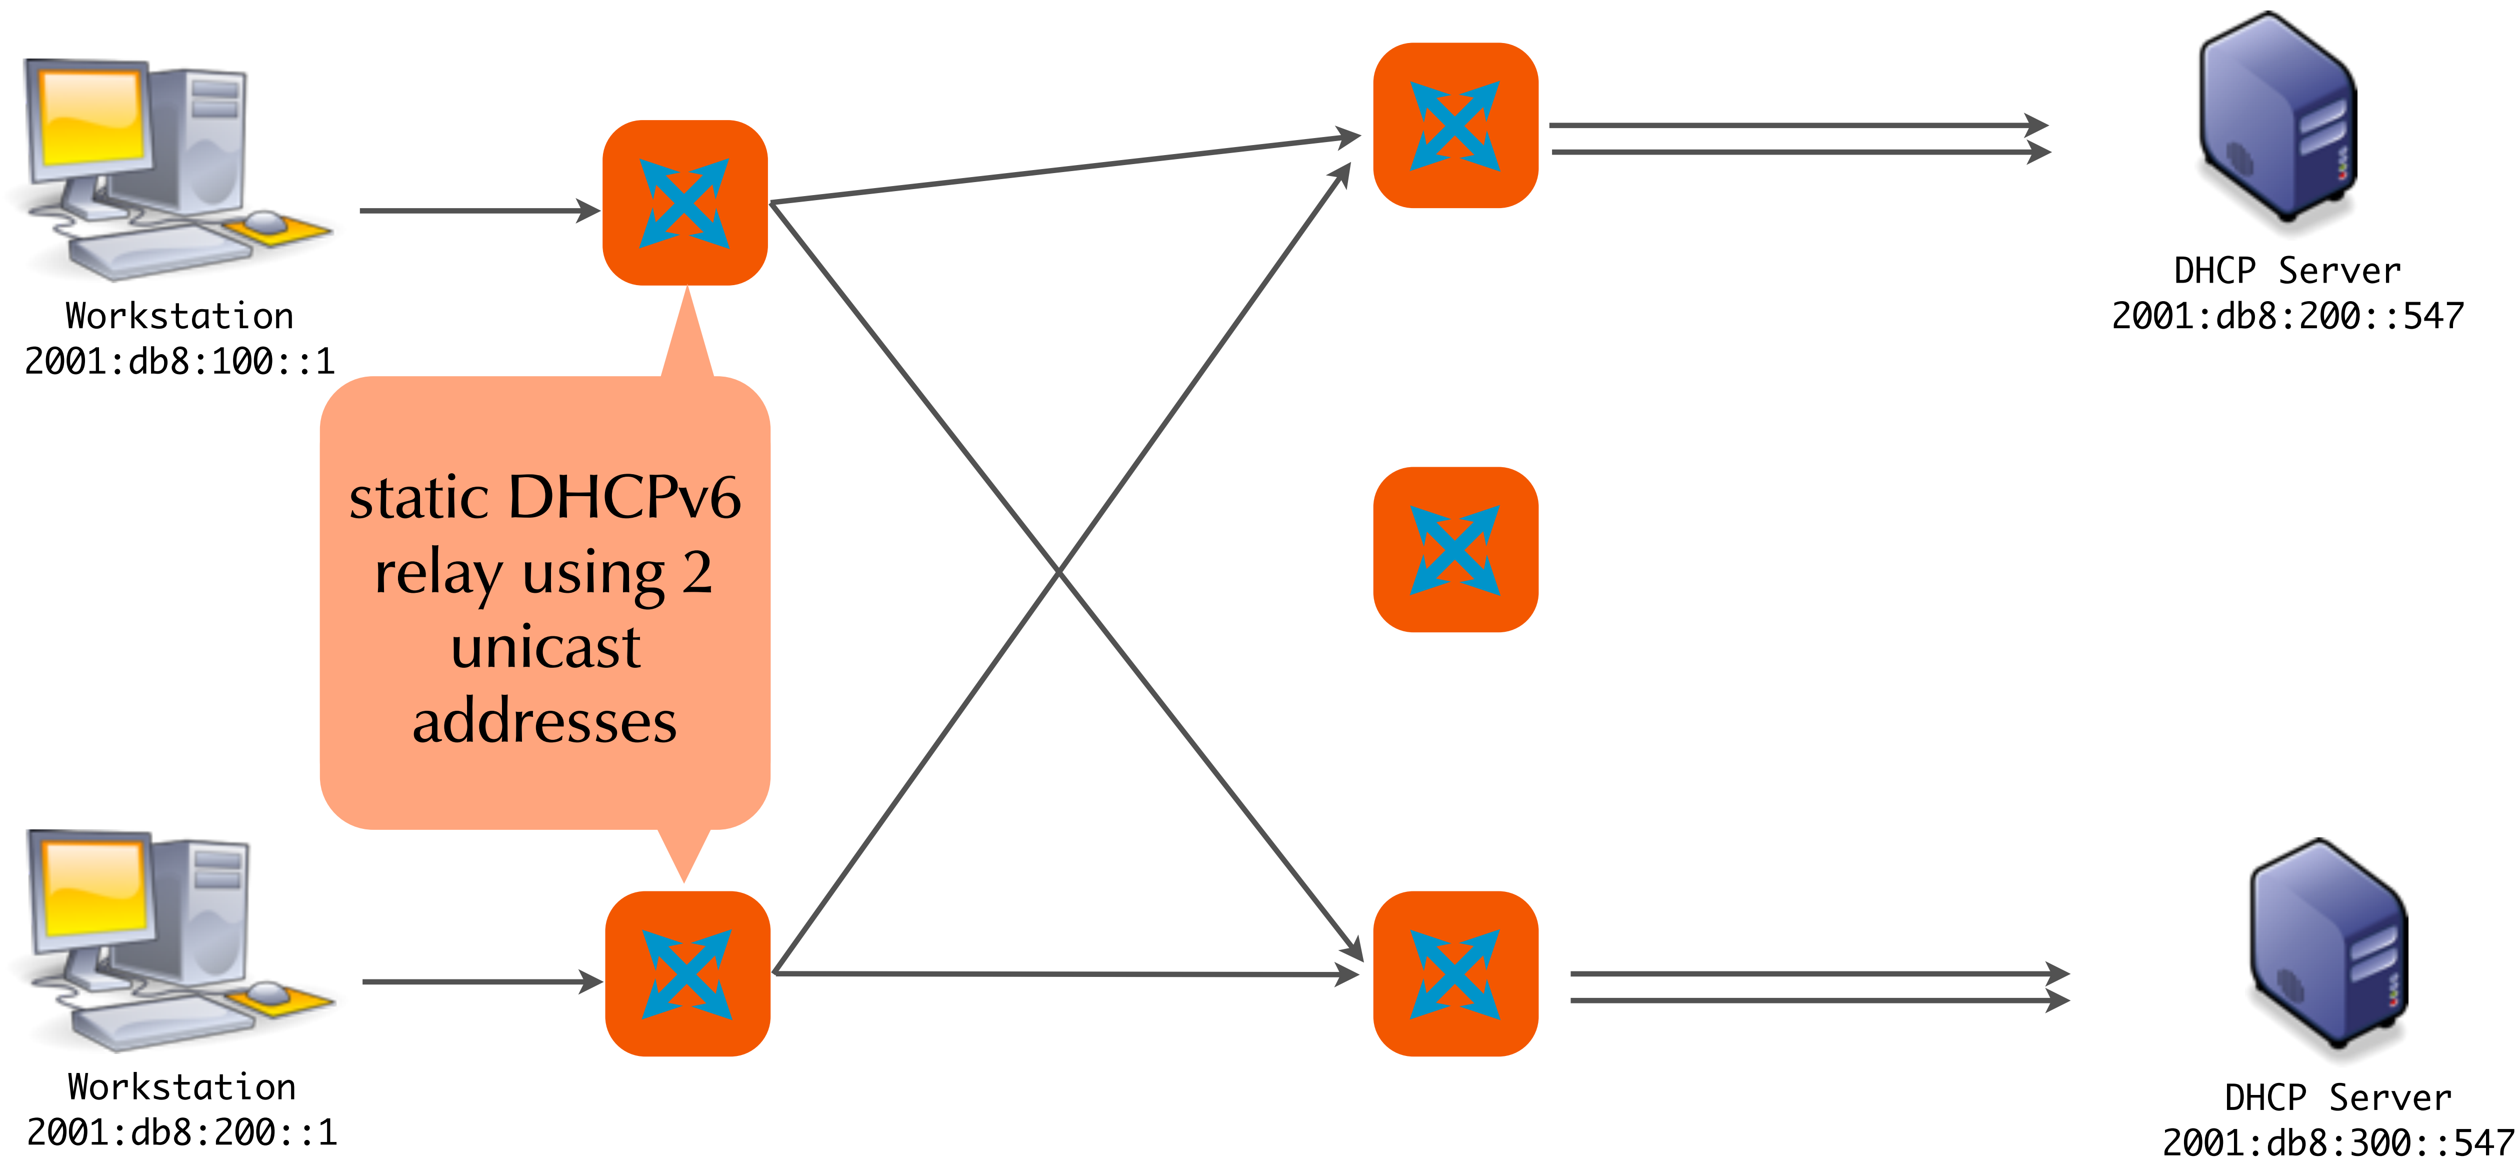

DHCPv6 using an unicast relay

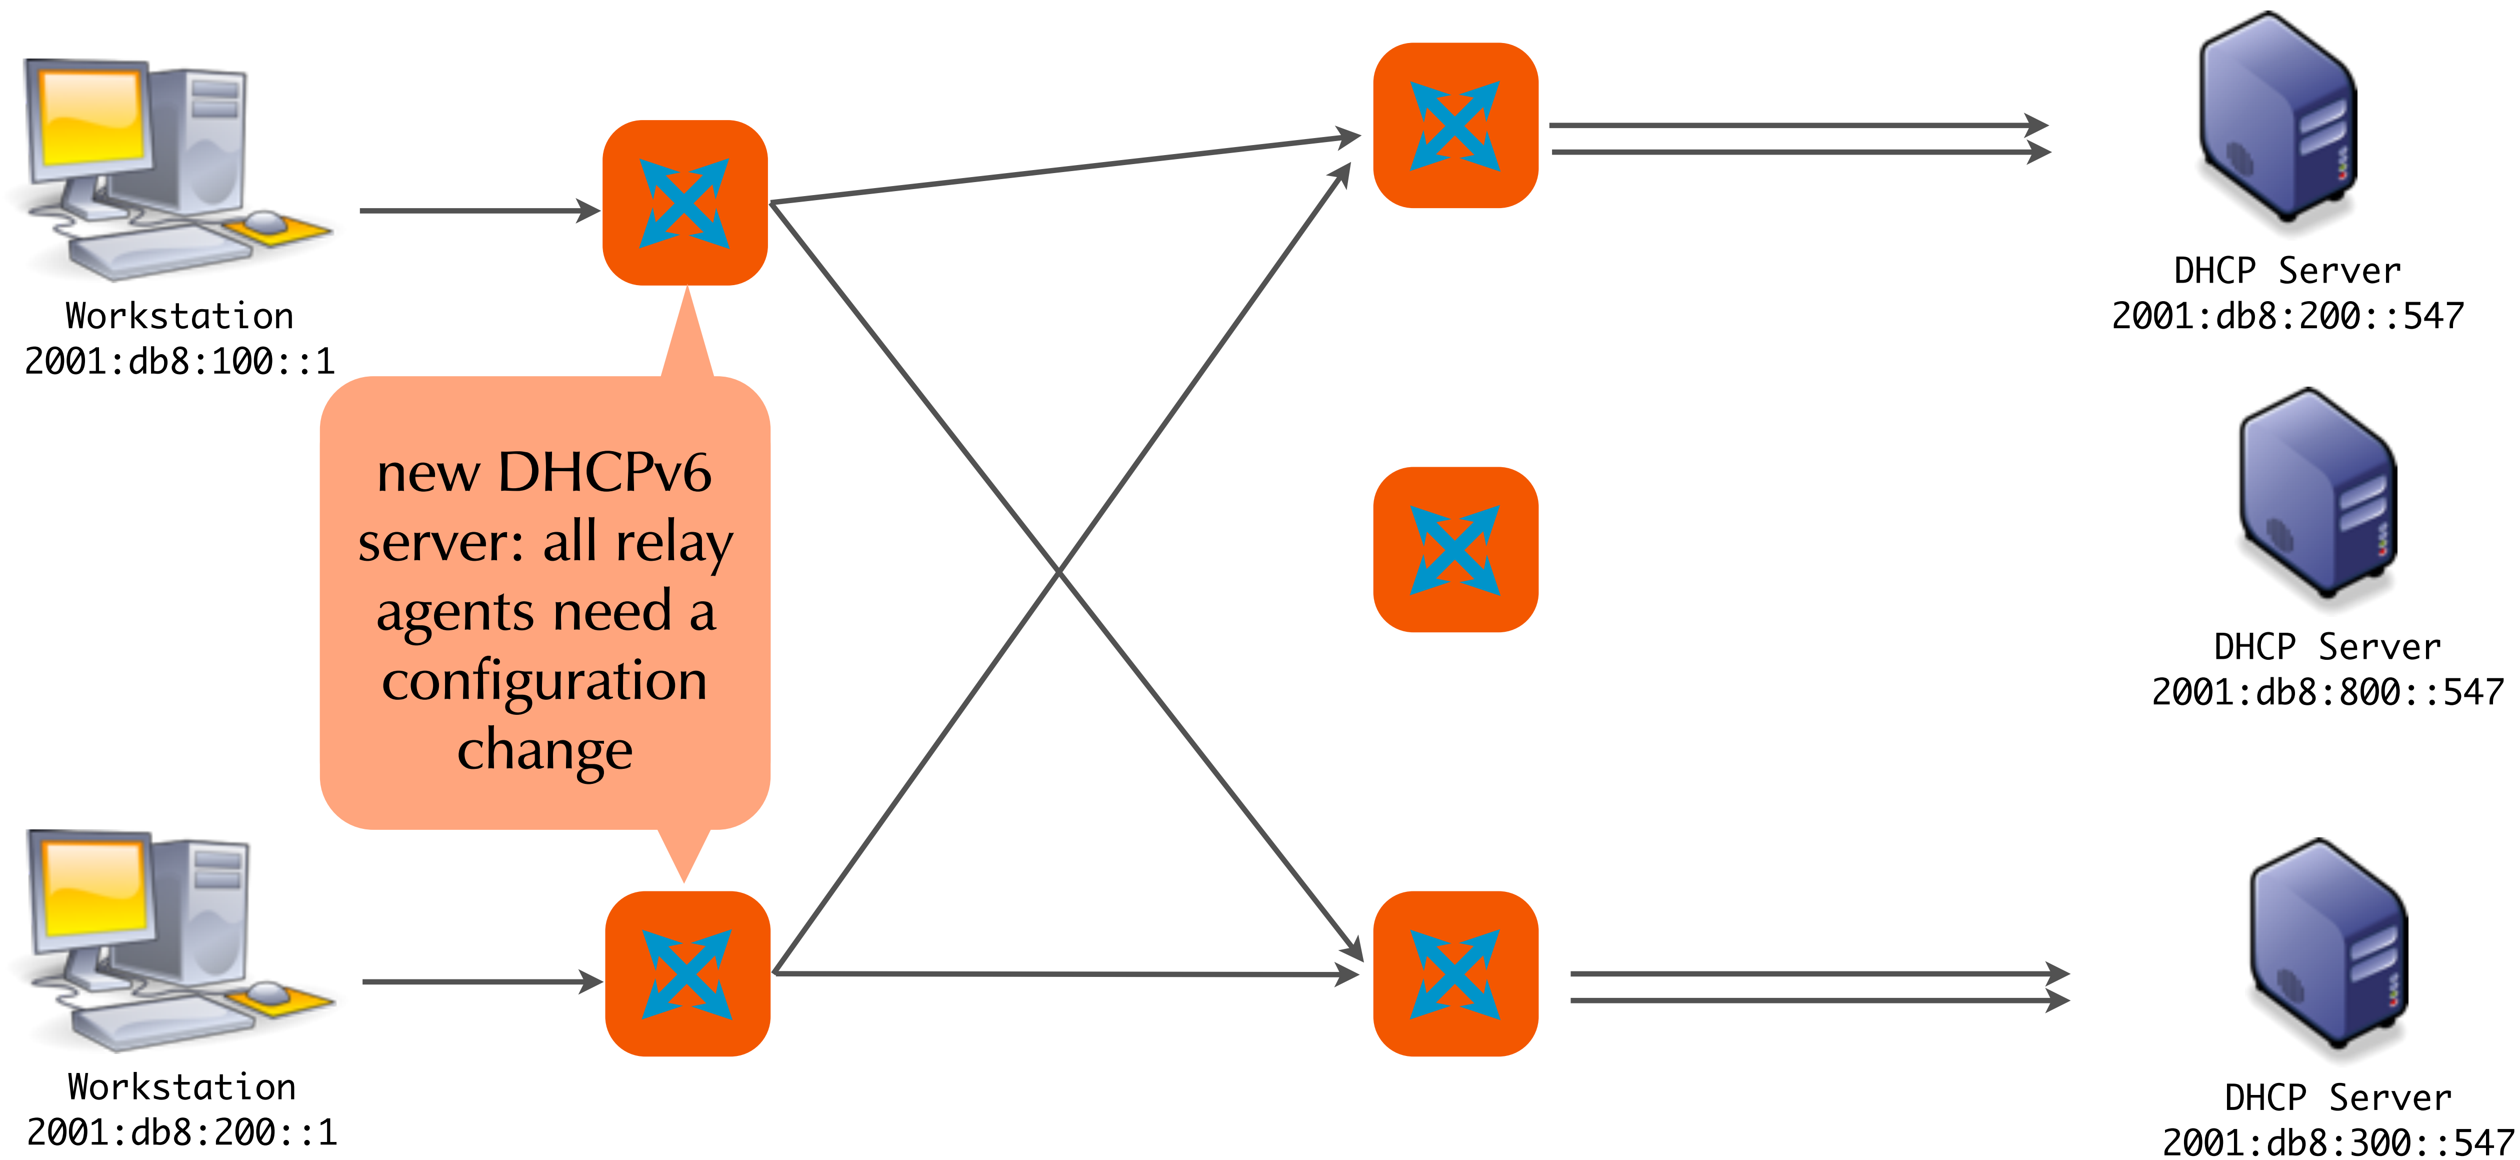

DHCPv6 using an unicast relay

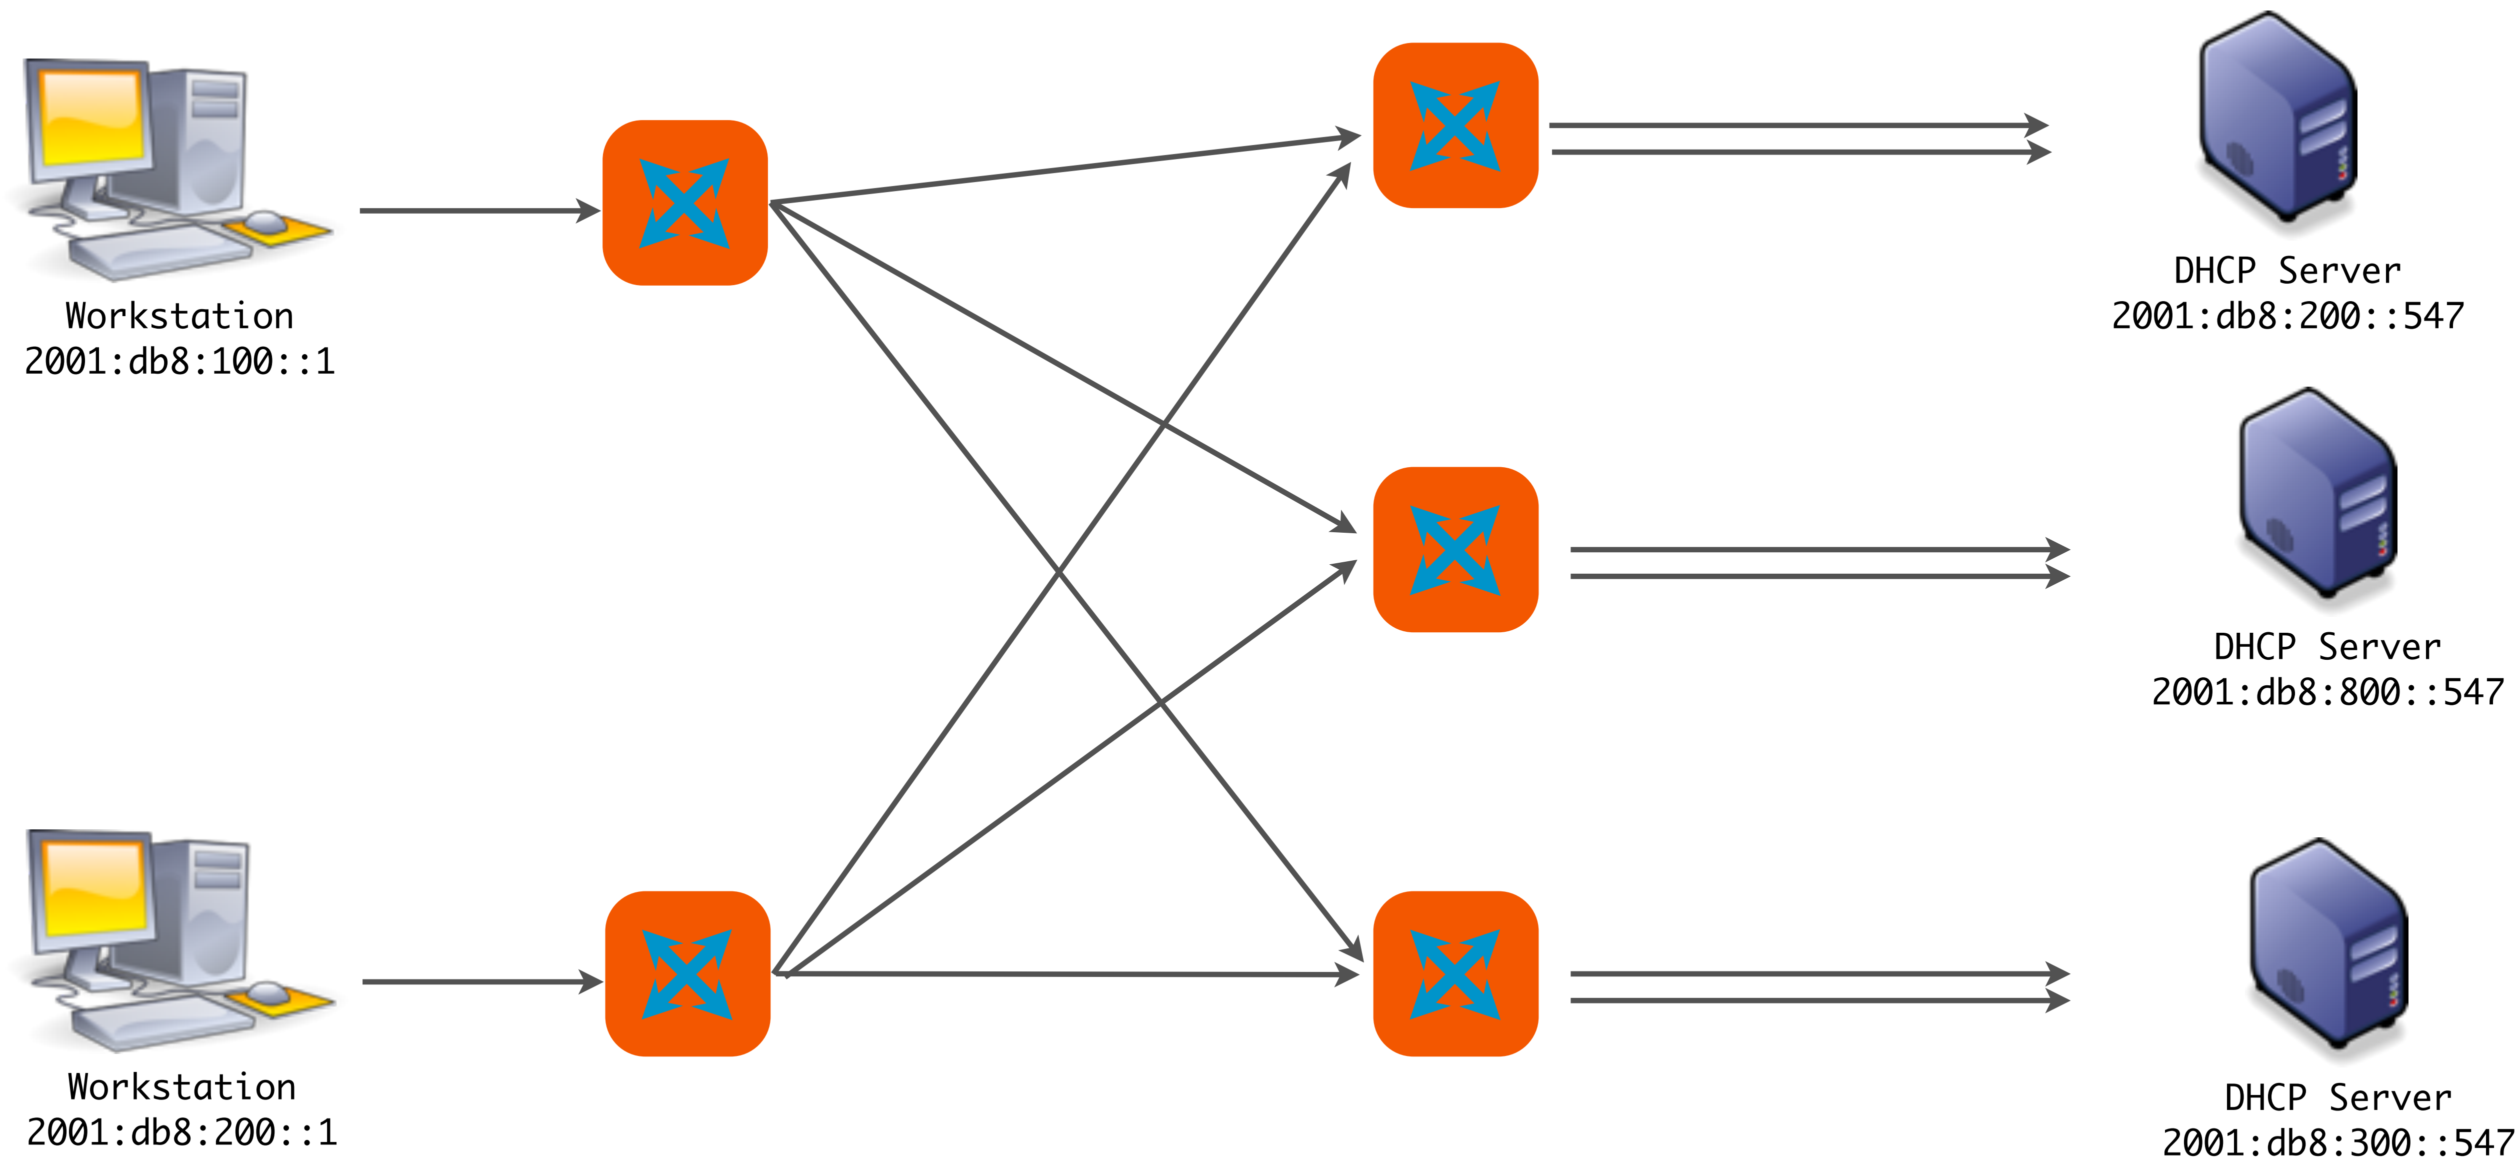

DHCPv6 using multicast relay

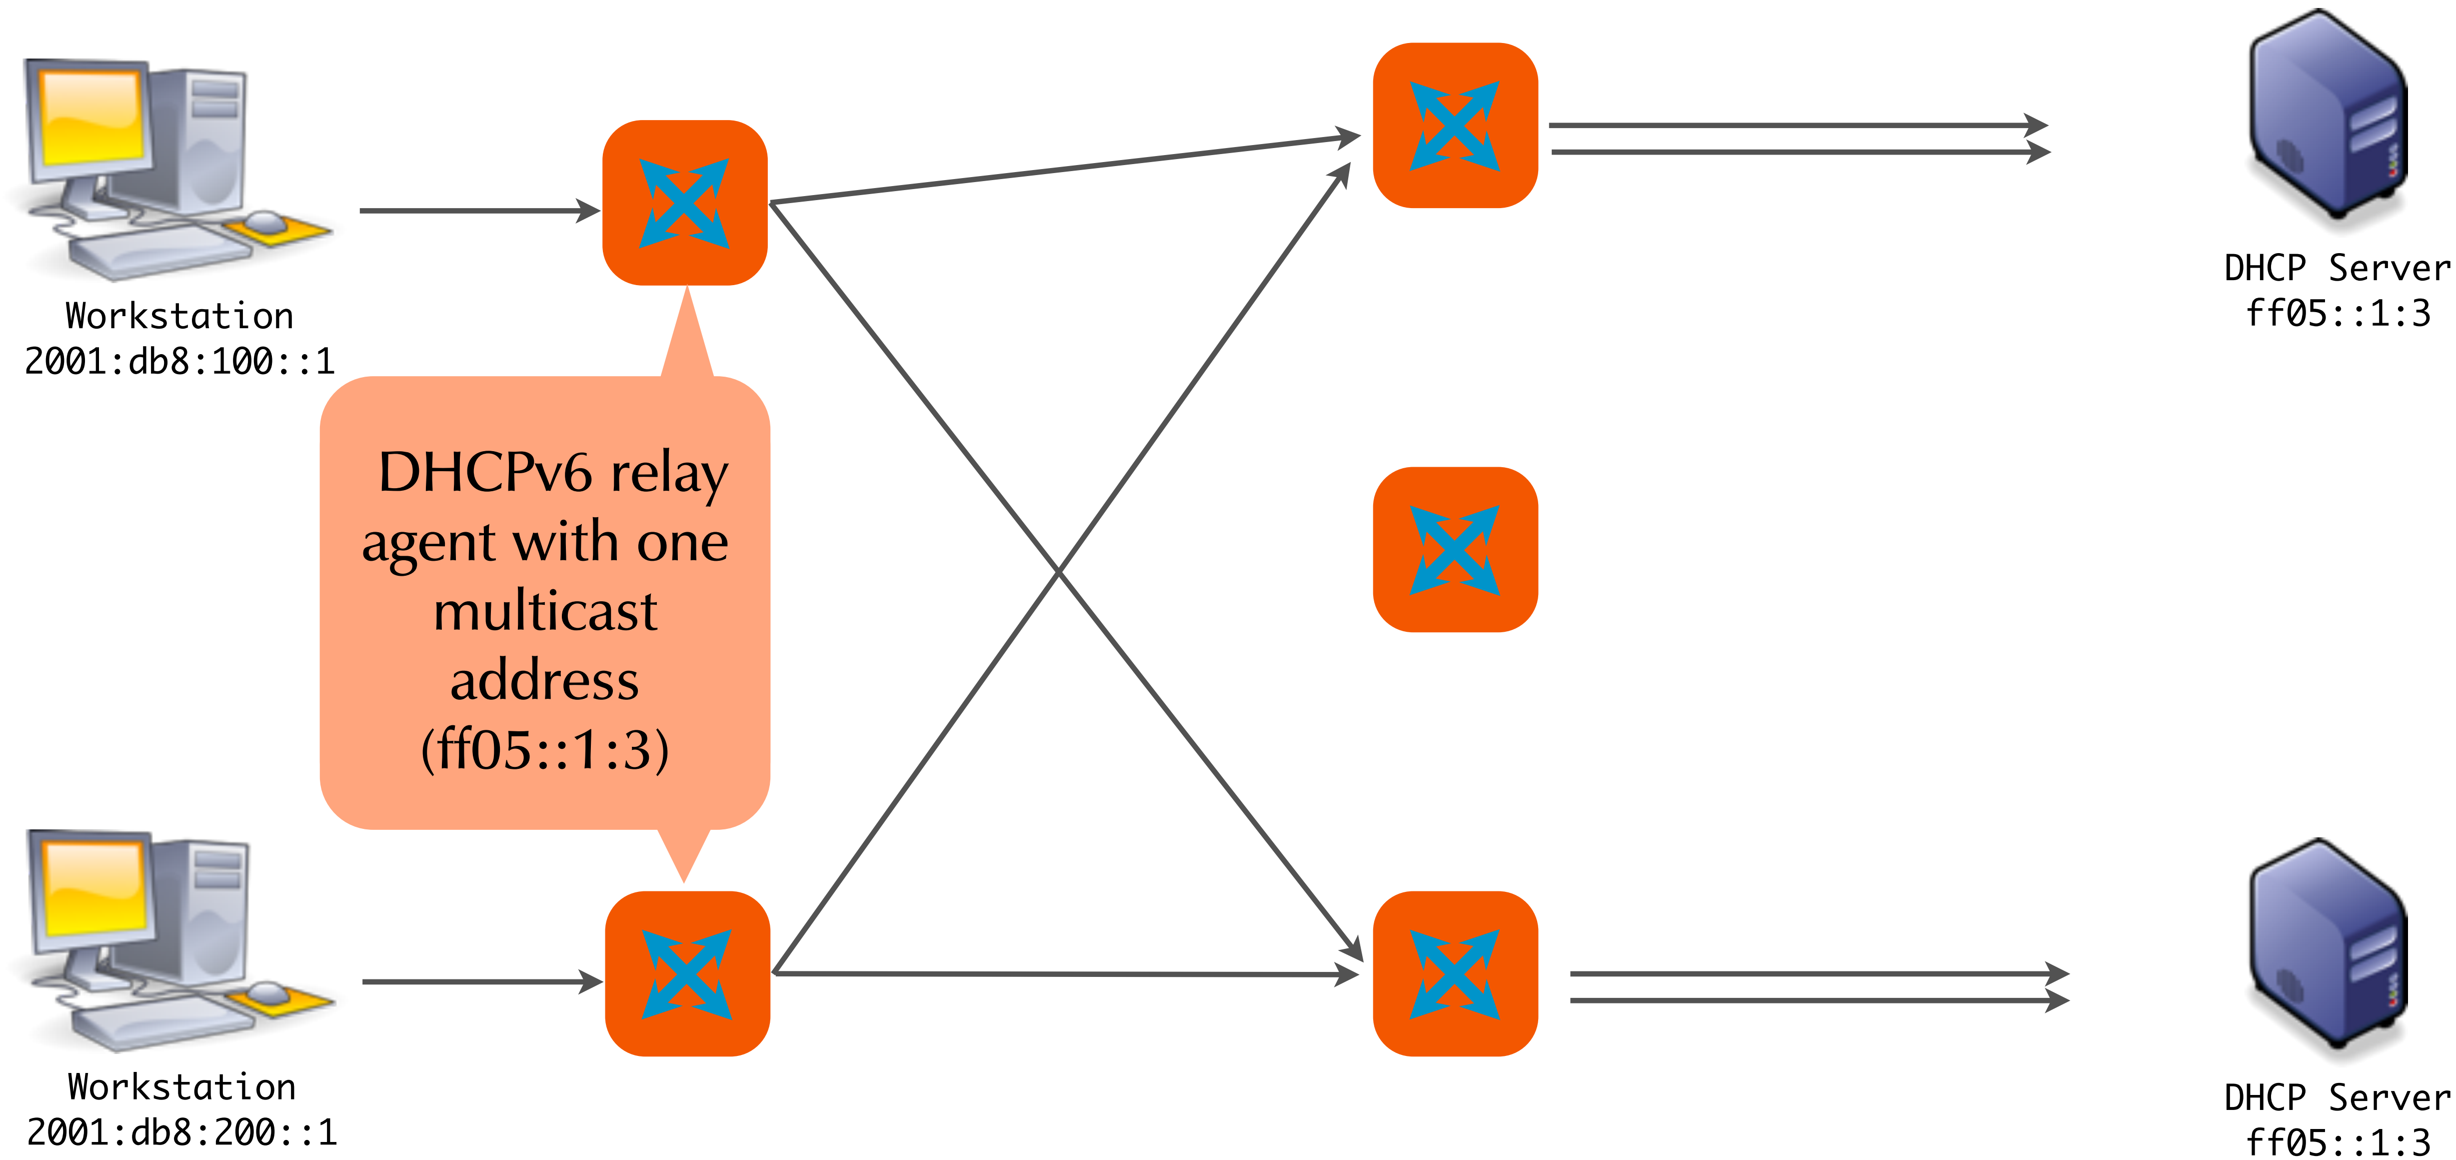

DHCPv6 using multicast relay

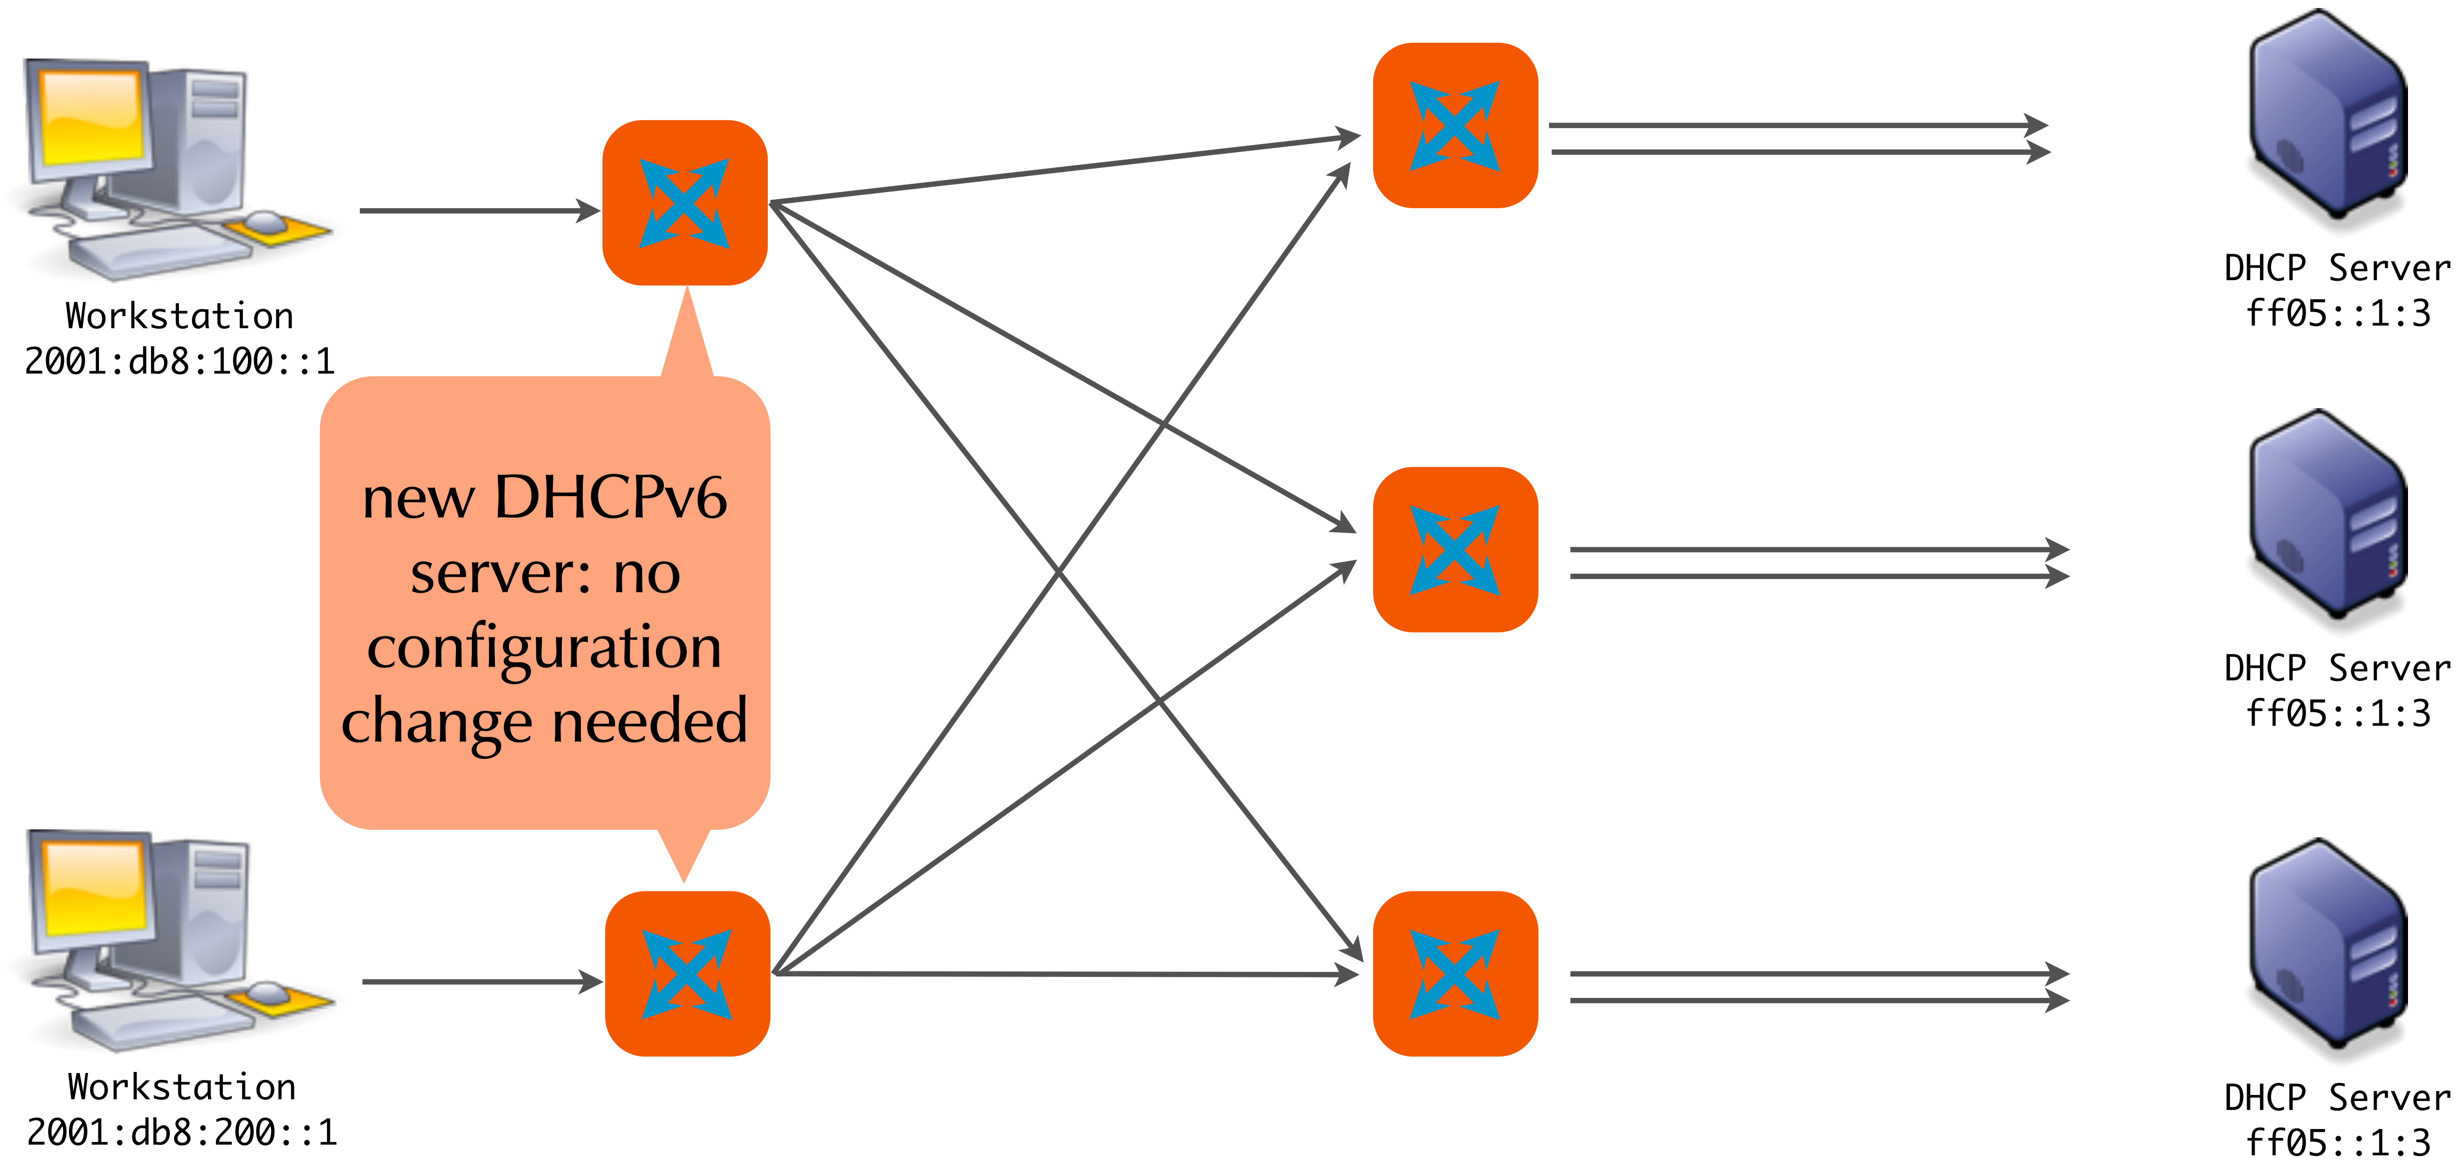

Case-study X11

- XDMCP - finds remote X11 applications (clients)
  - uses broadcast on IPv4
    - inefficient
    - only works in a local subnet
      - workaround: “chooser” machines that need configuration
  - XDMCP can use IPv6 Multicast

X11 and IPv6 Multicast

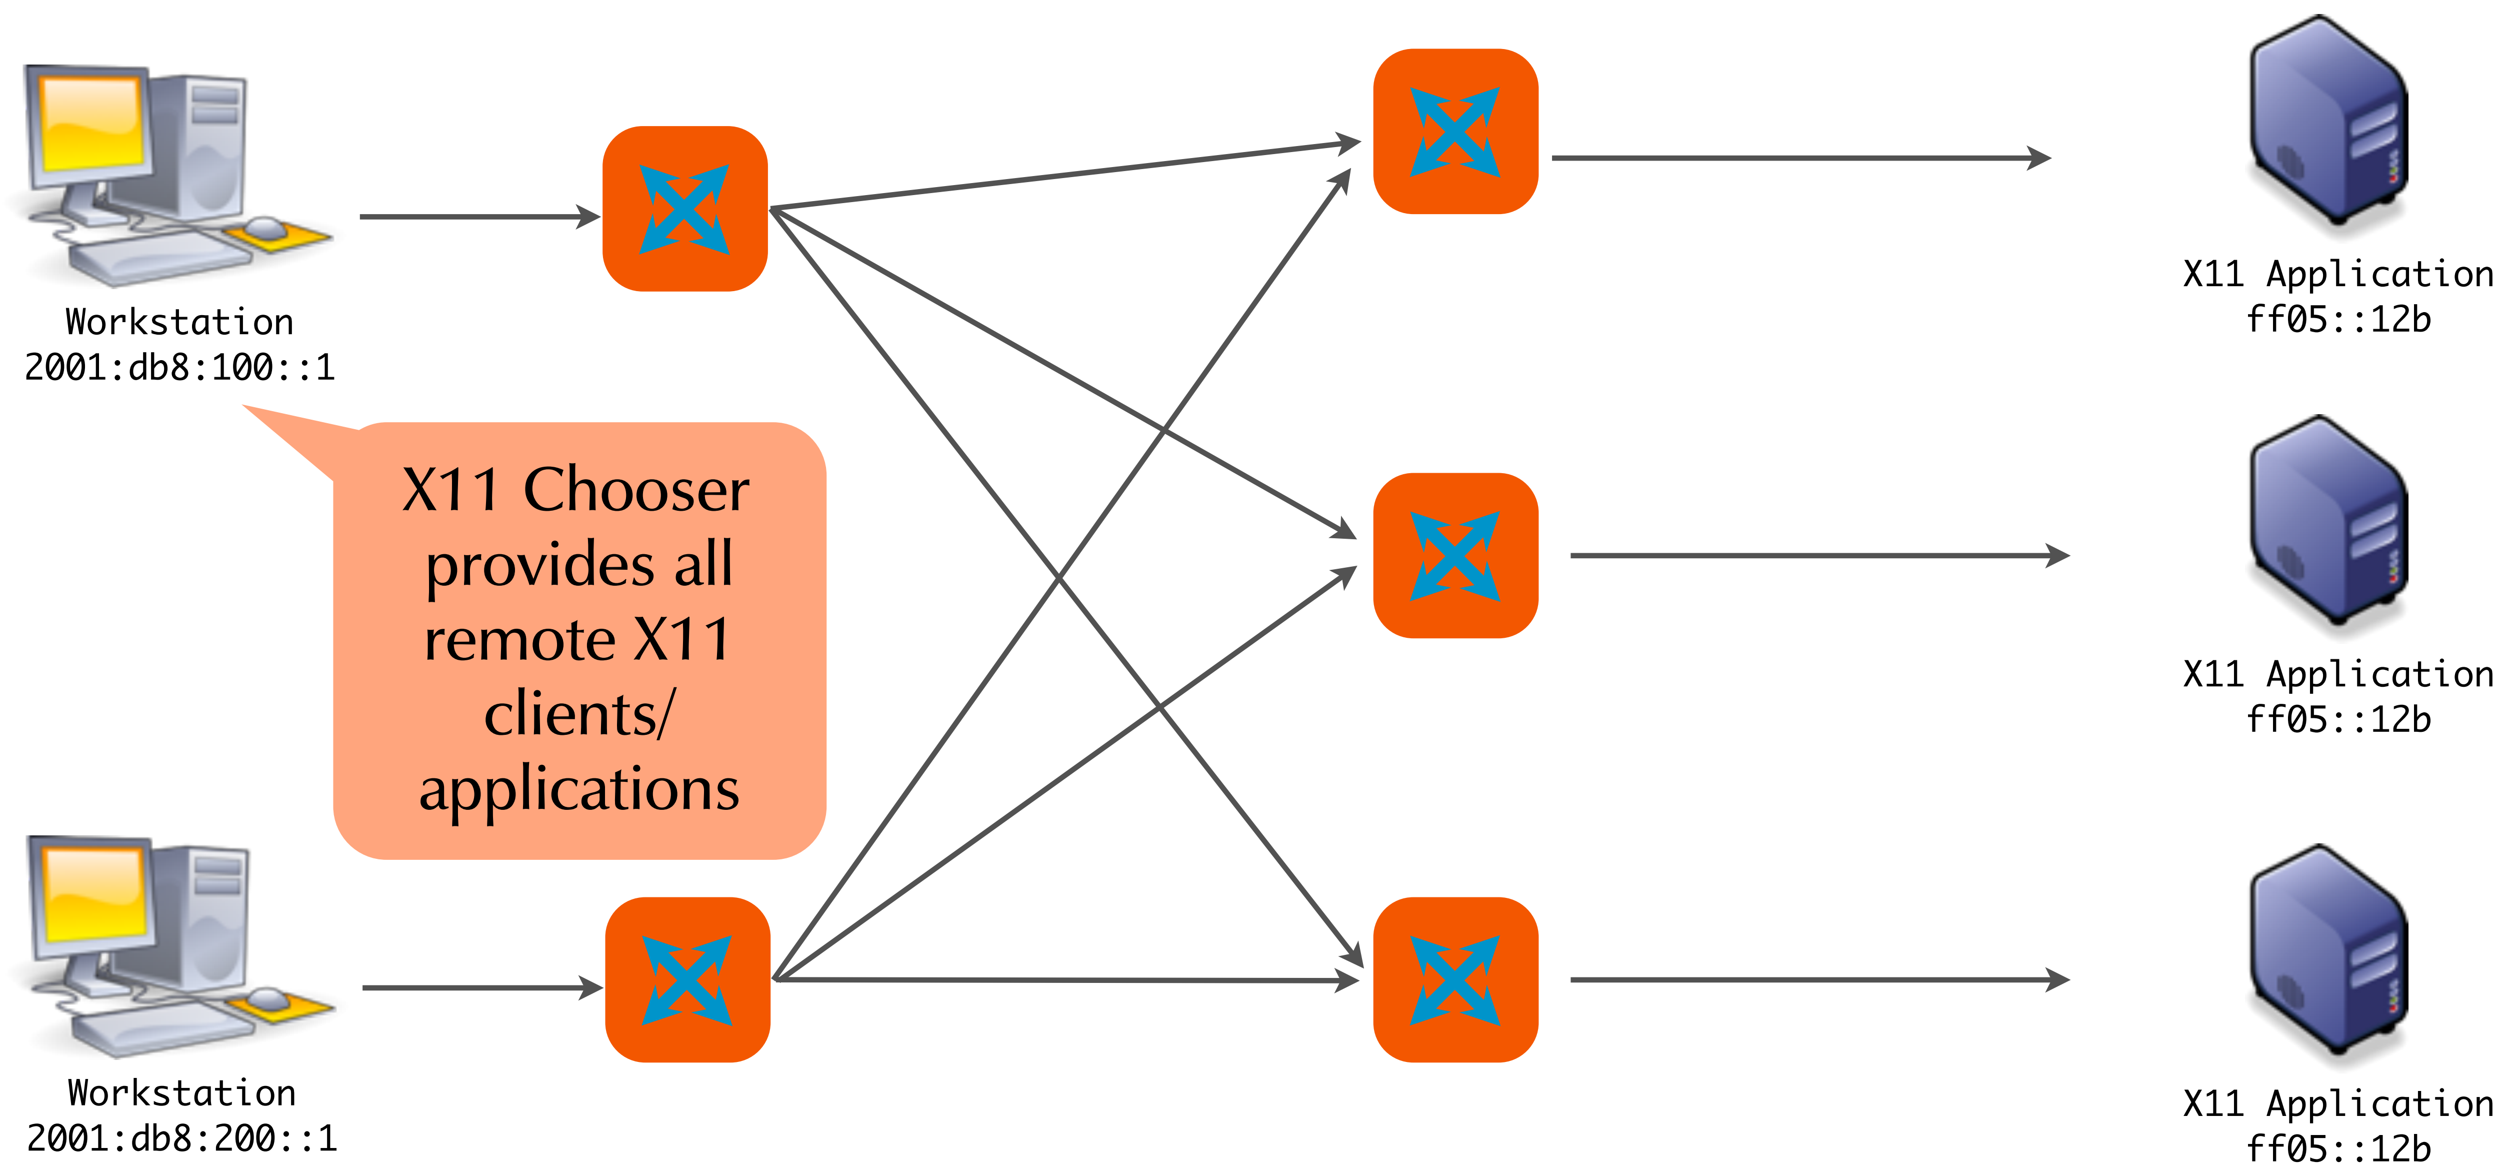

Case-study NTP

- NTP = Network Time Protocol
- NTP can use IPv6 multicast to distribute time information
  - new time sources can be added at any time without reconfiguration

NTP using Multicast

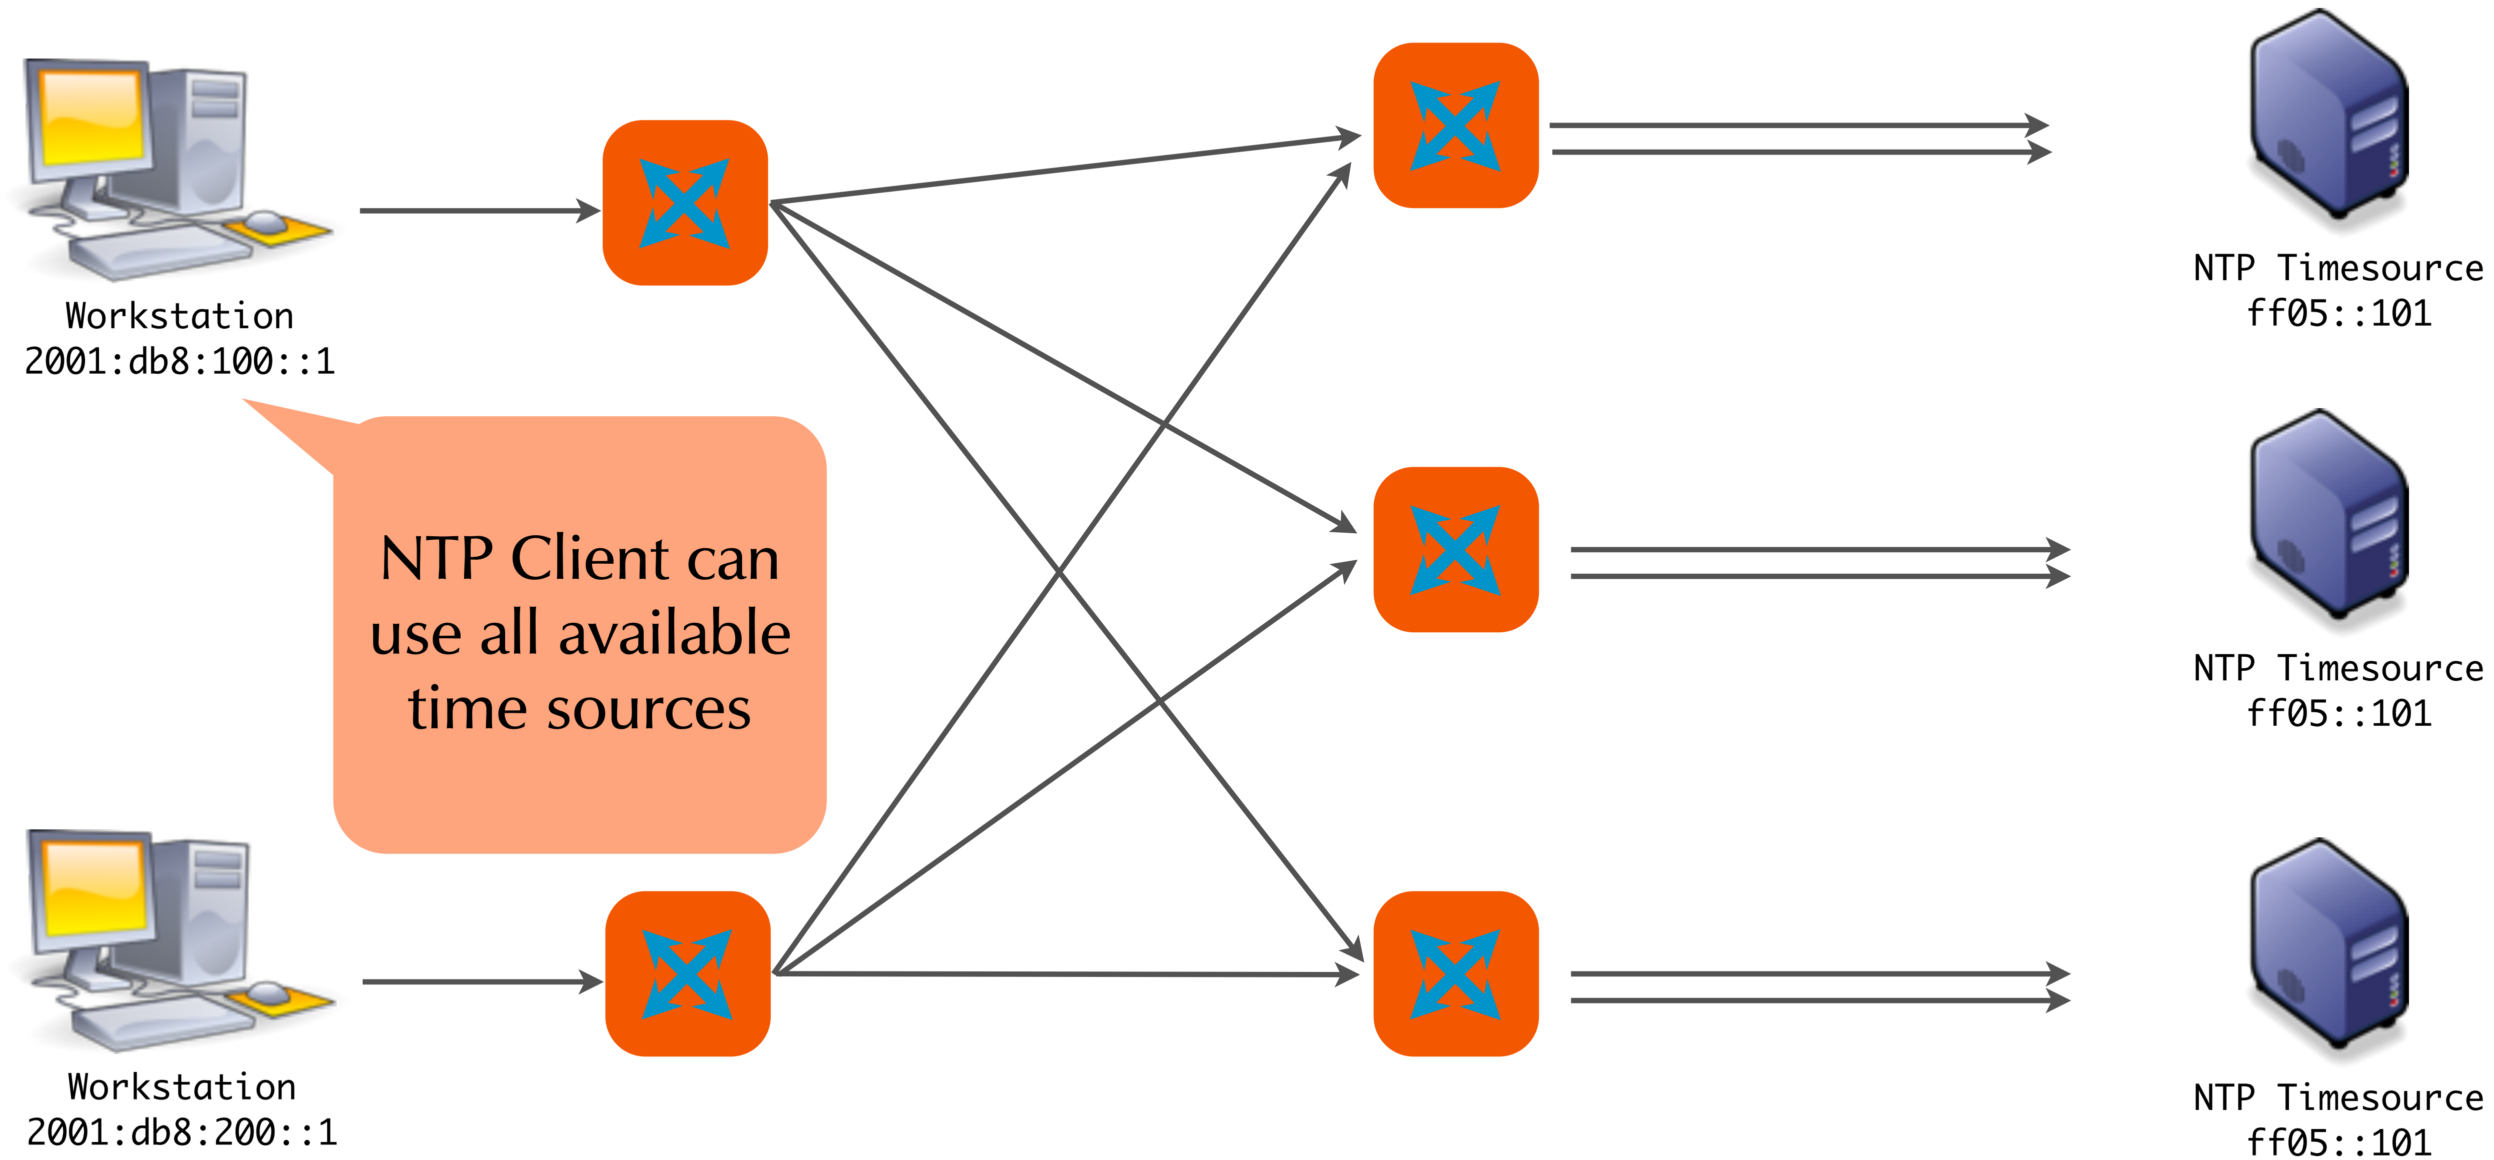

Case-study NTP Manycast

- NTP Manycast turns NTP-clients into NTP-Server (or sources)
  - the NTP-Protocol dynamically finds the best (most reliable/most accurate) time sources
  - every host is both multicast-client and -server
  - the NTP-Jargon calls this “Manycast”

NTP using Manycast

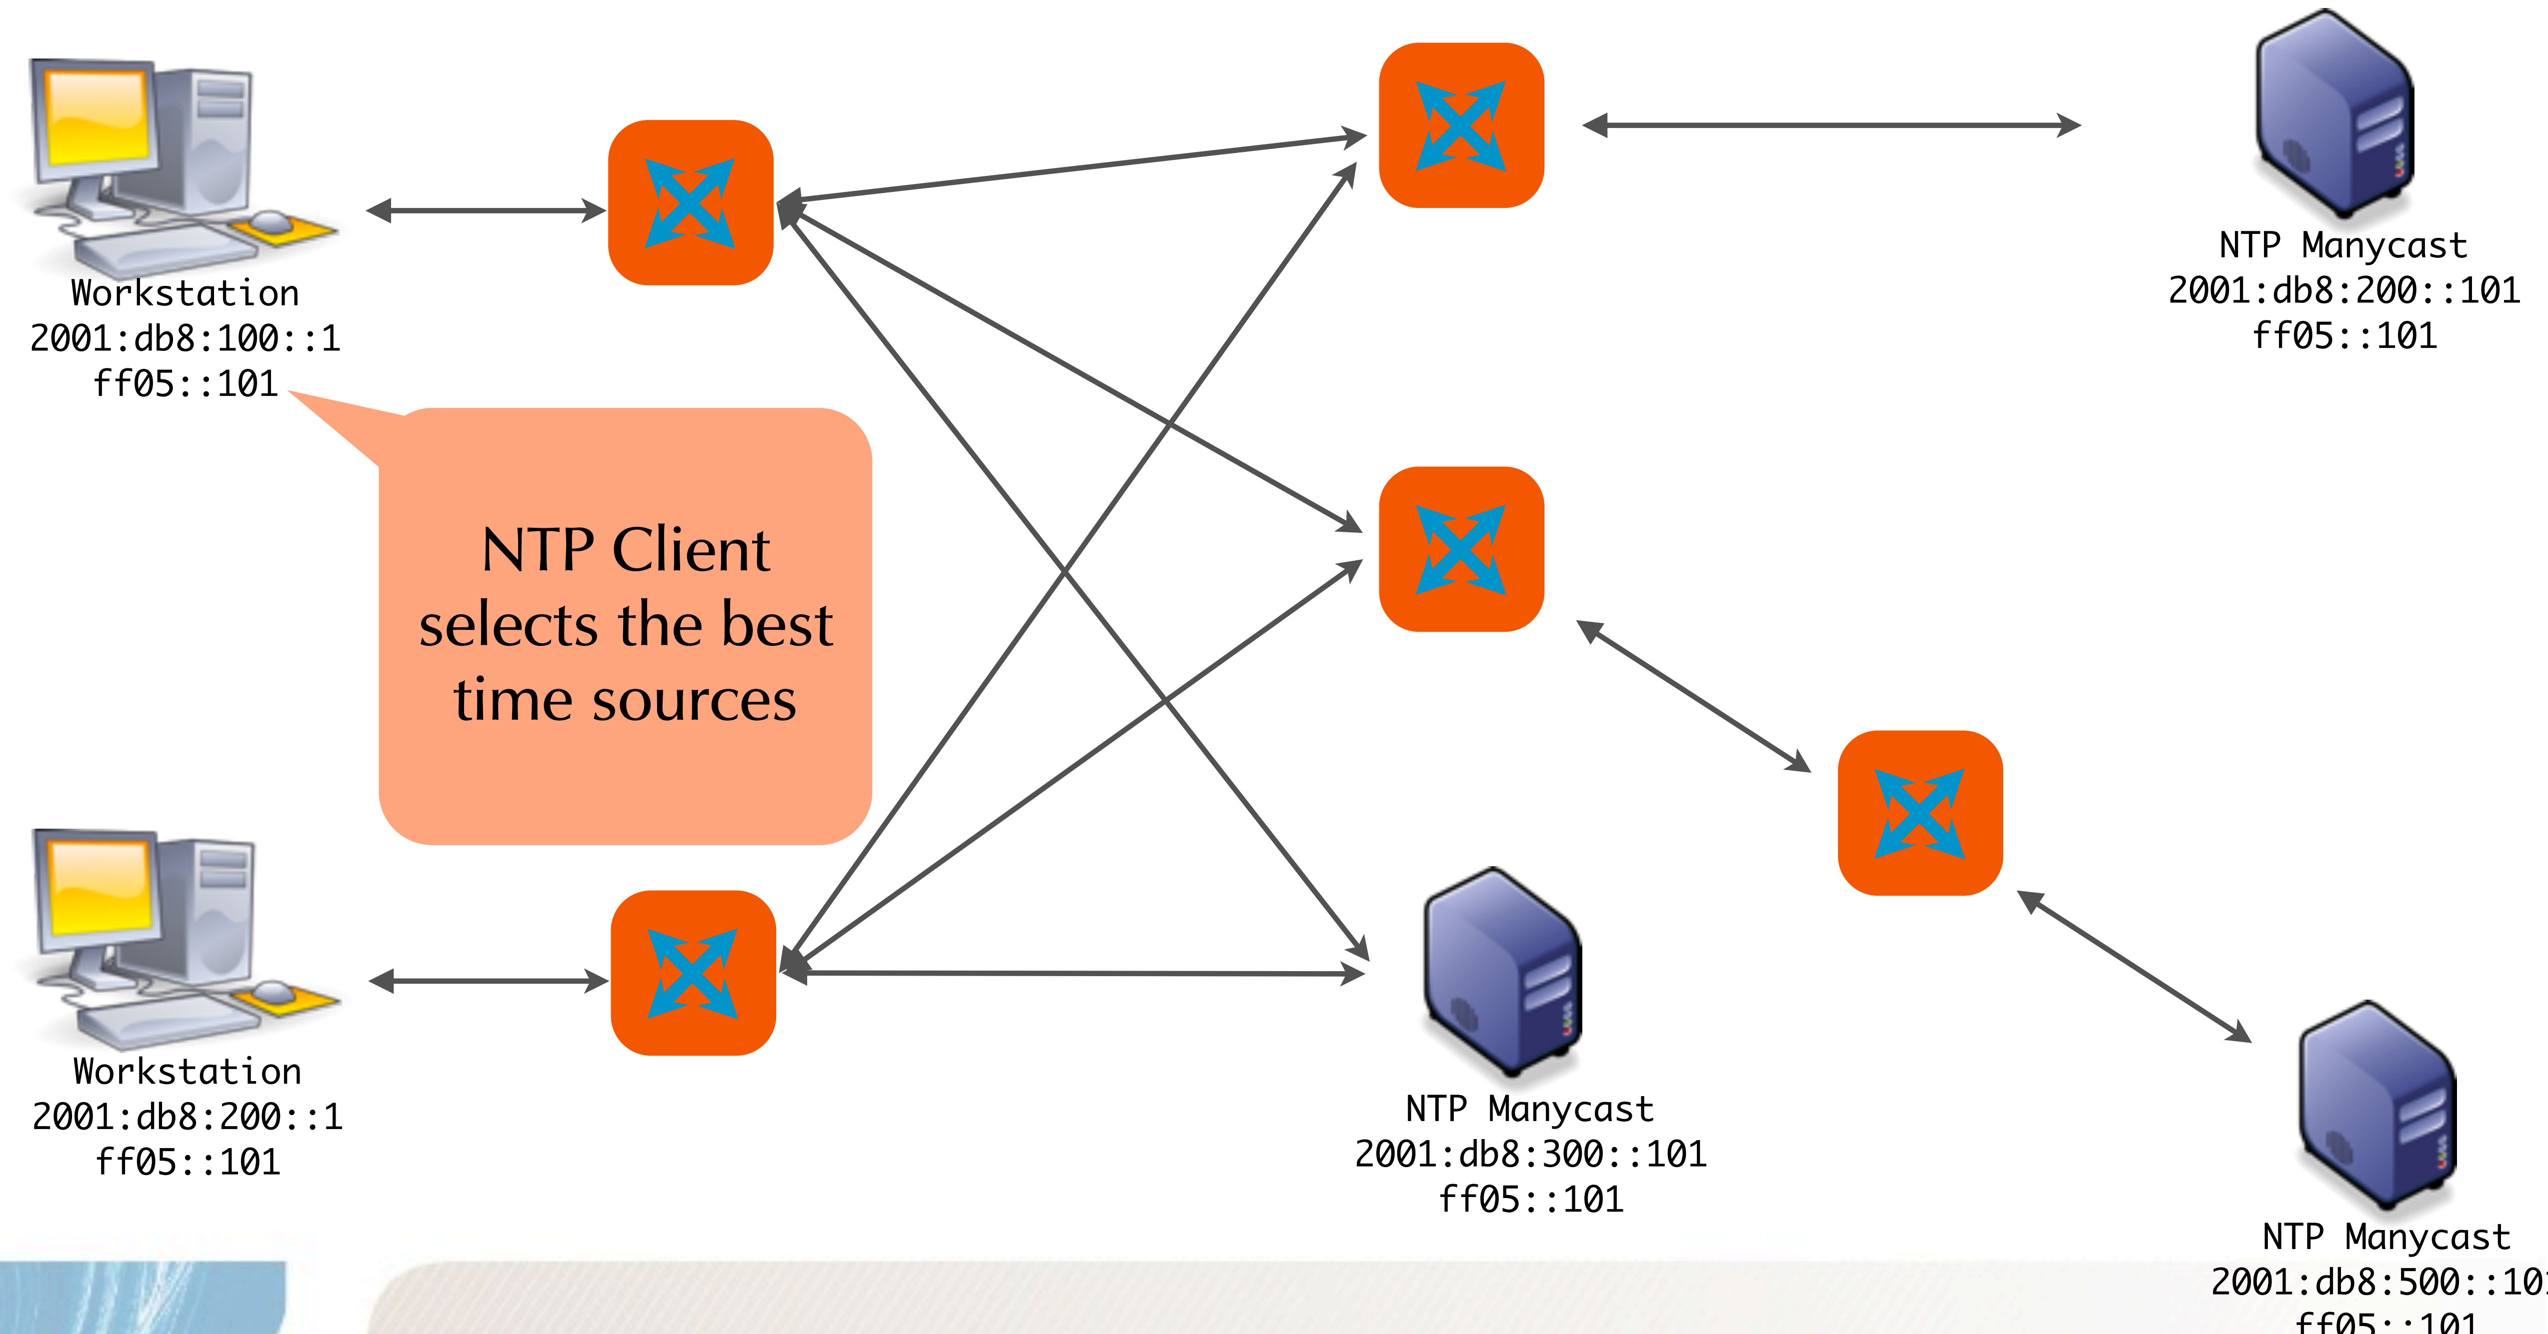

DNS using IPv6 Multicast

- this is experimental and not an Internet standard
  - more a “hack-that-works”
- don’t confuse with mDNS (Multicast DNS aka Apple Bonjour or Avahi)
- classic DNS communication between Stub-Resolver and DNS caching-server
- solves the problem of client DNS failover:
  - no extra Fail-Over or Load-Balancer necessary

DNS using Unicast

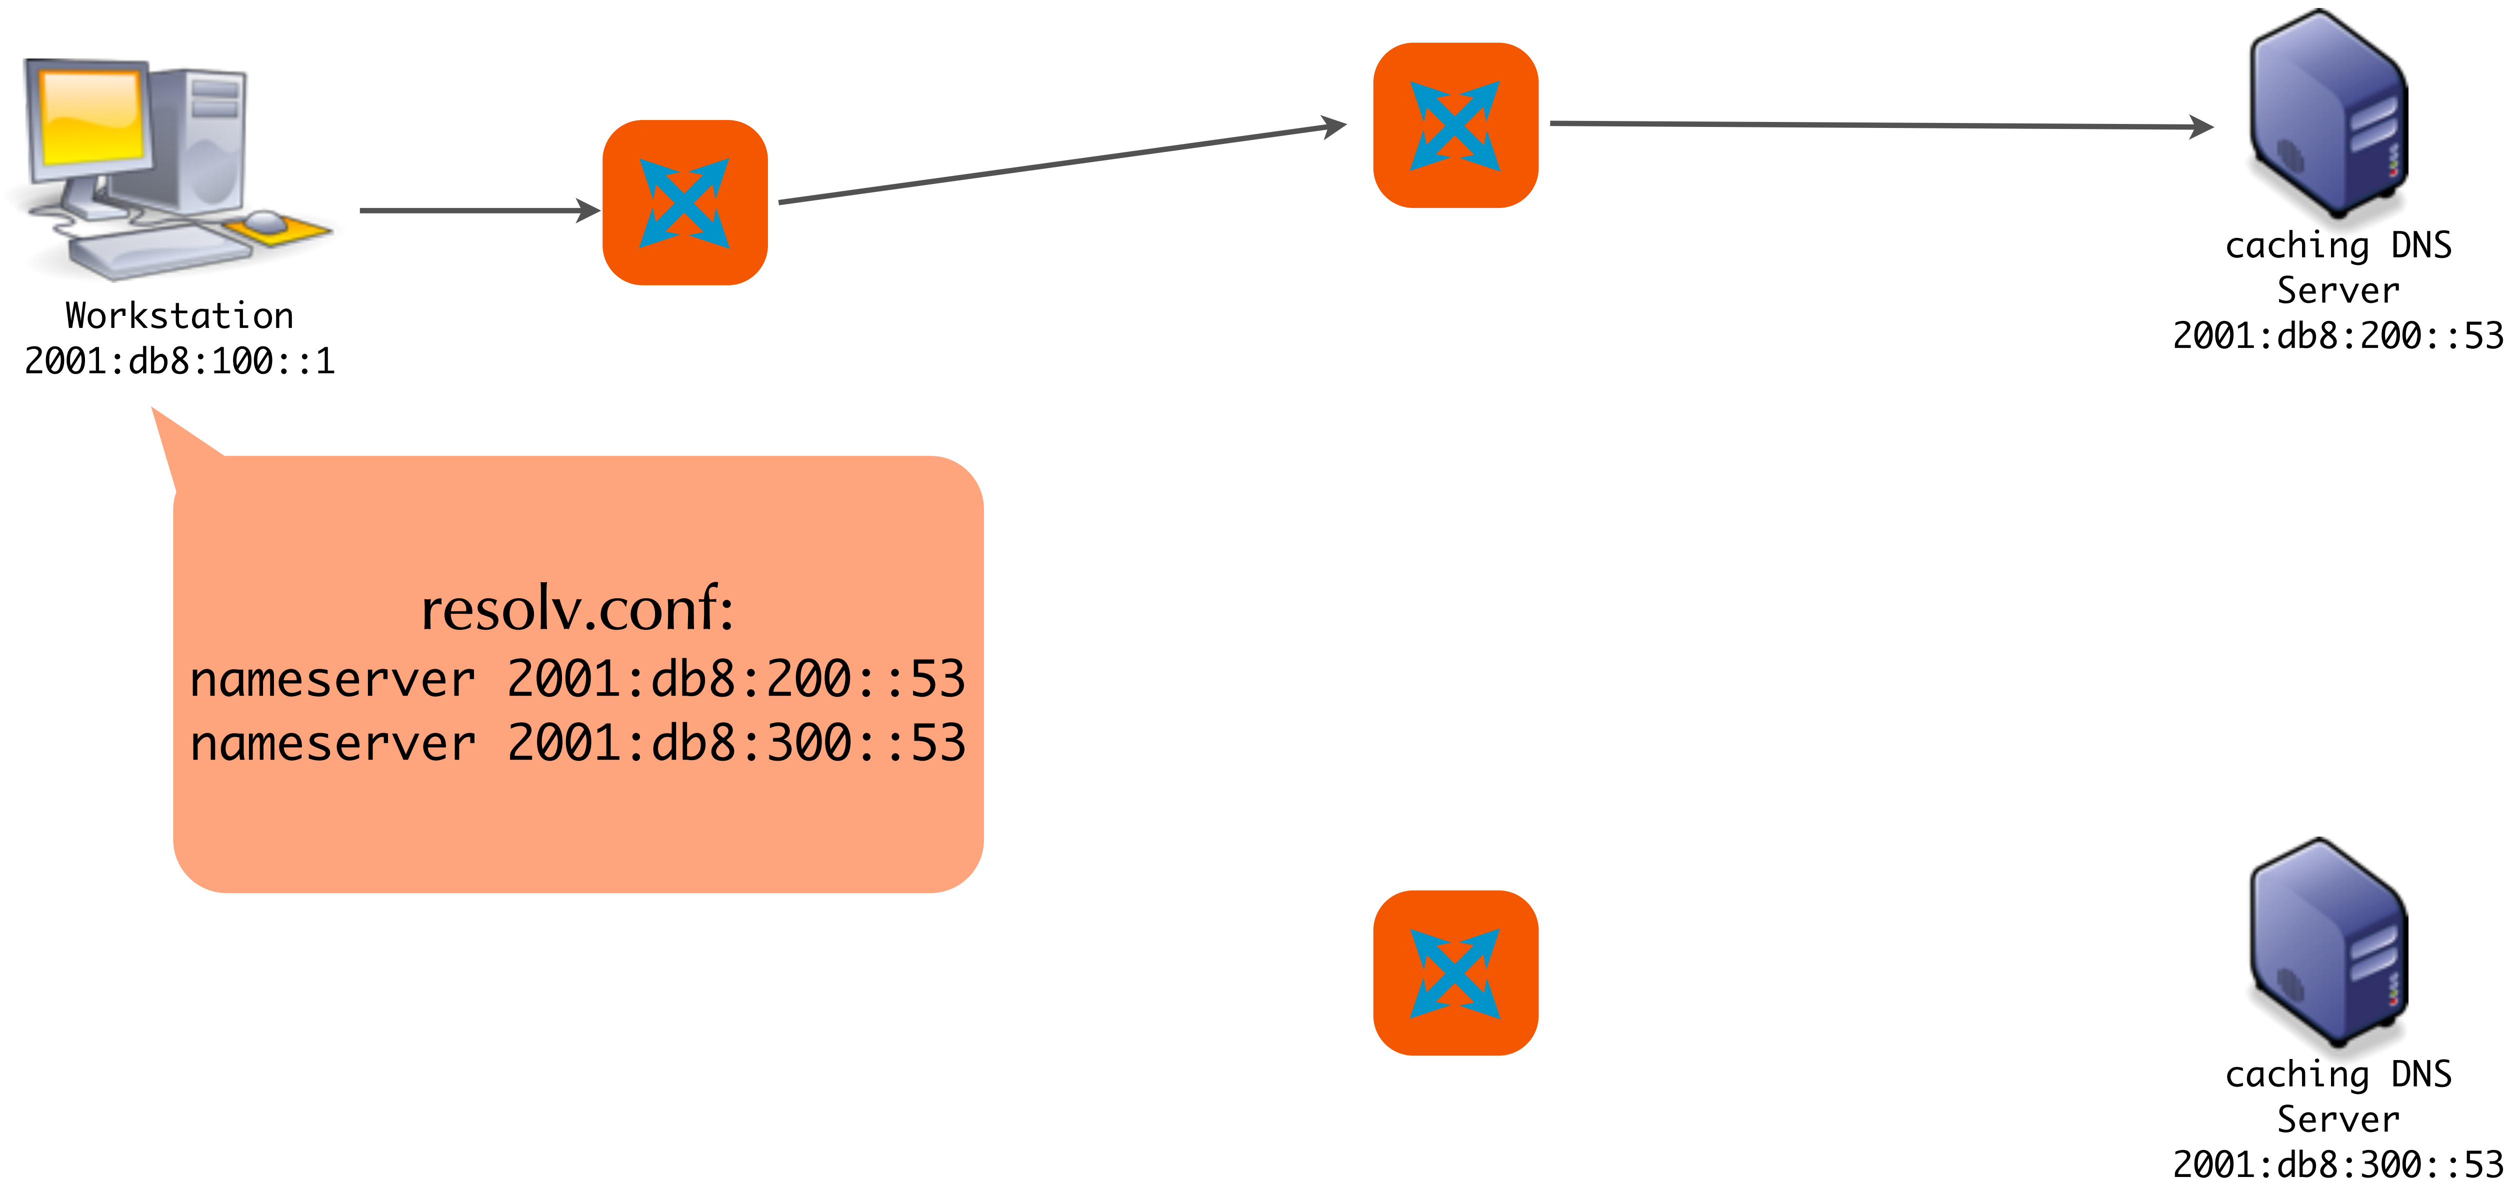

DNS using Unicast

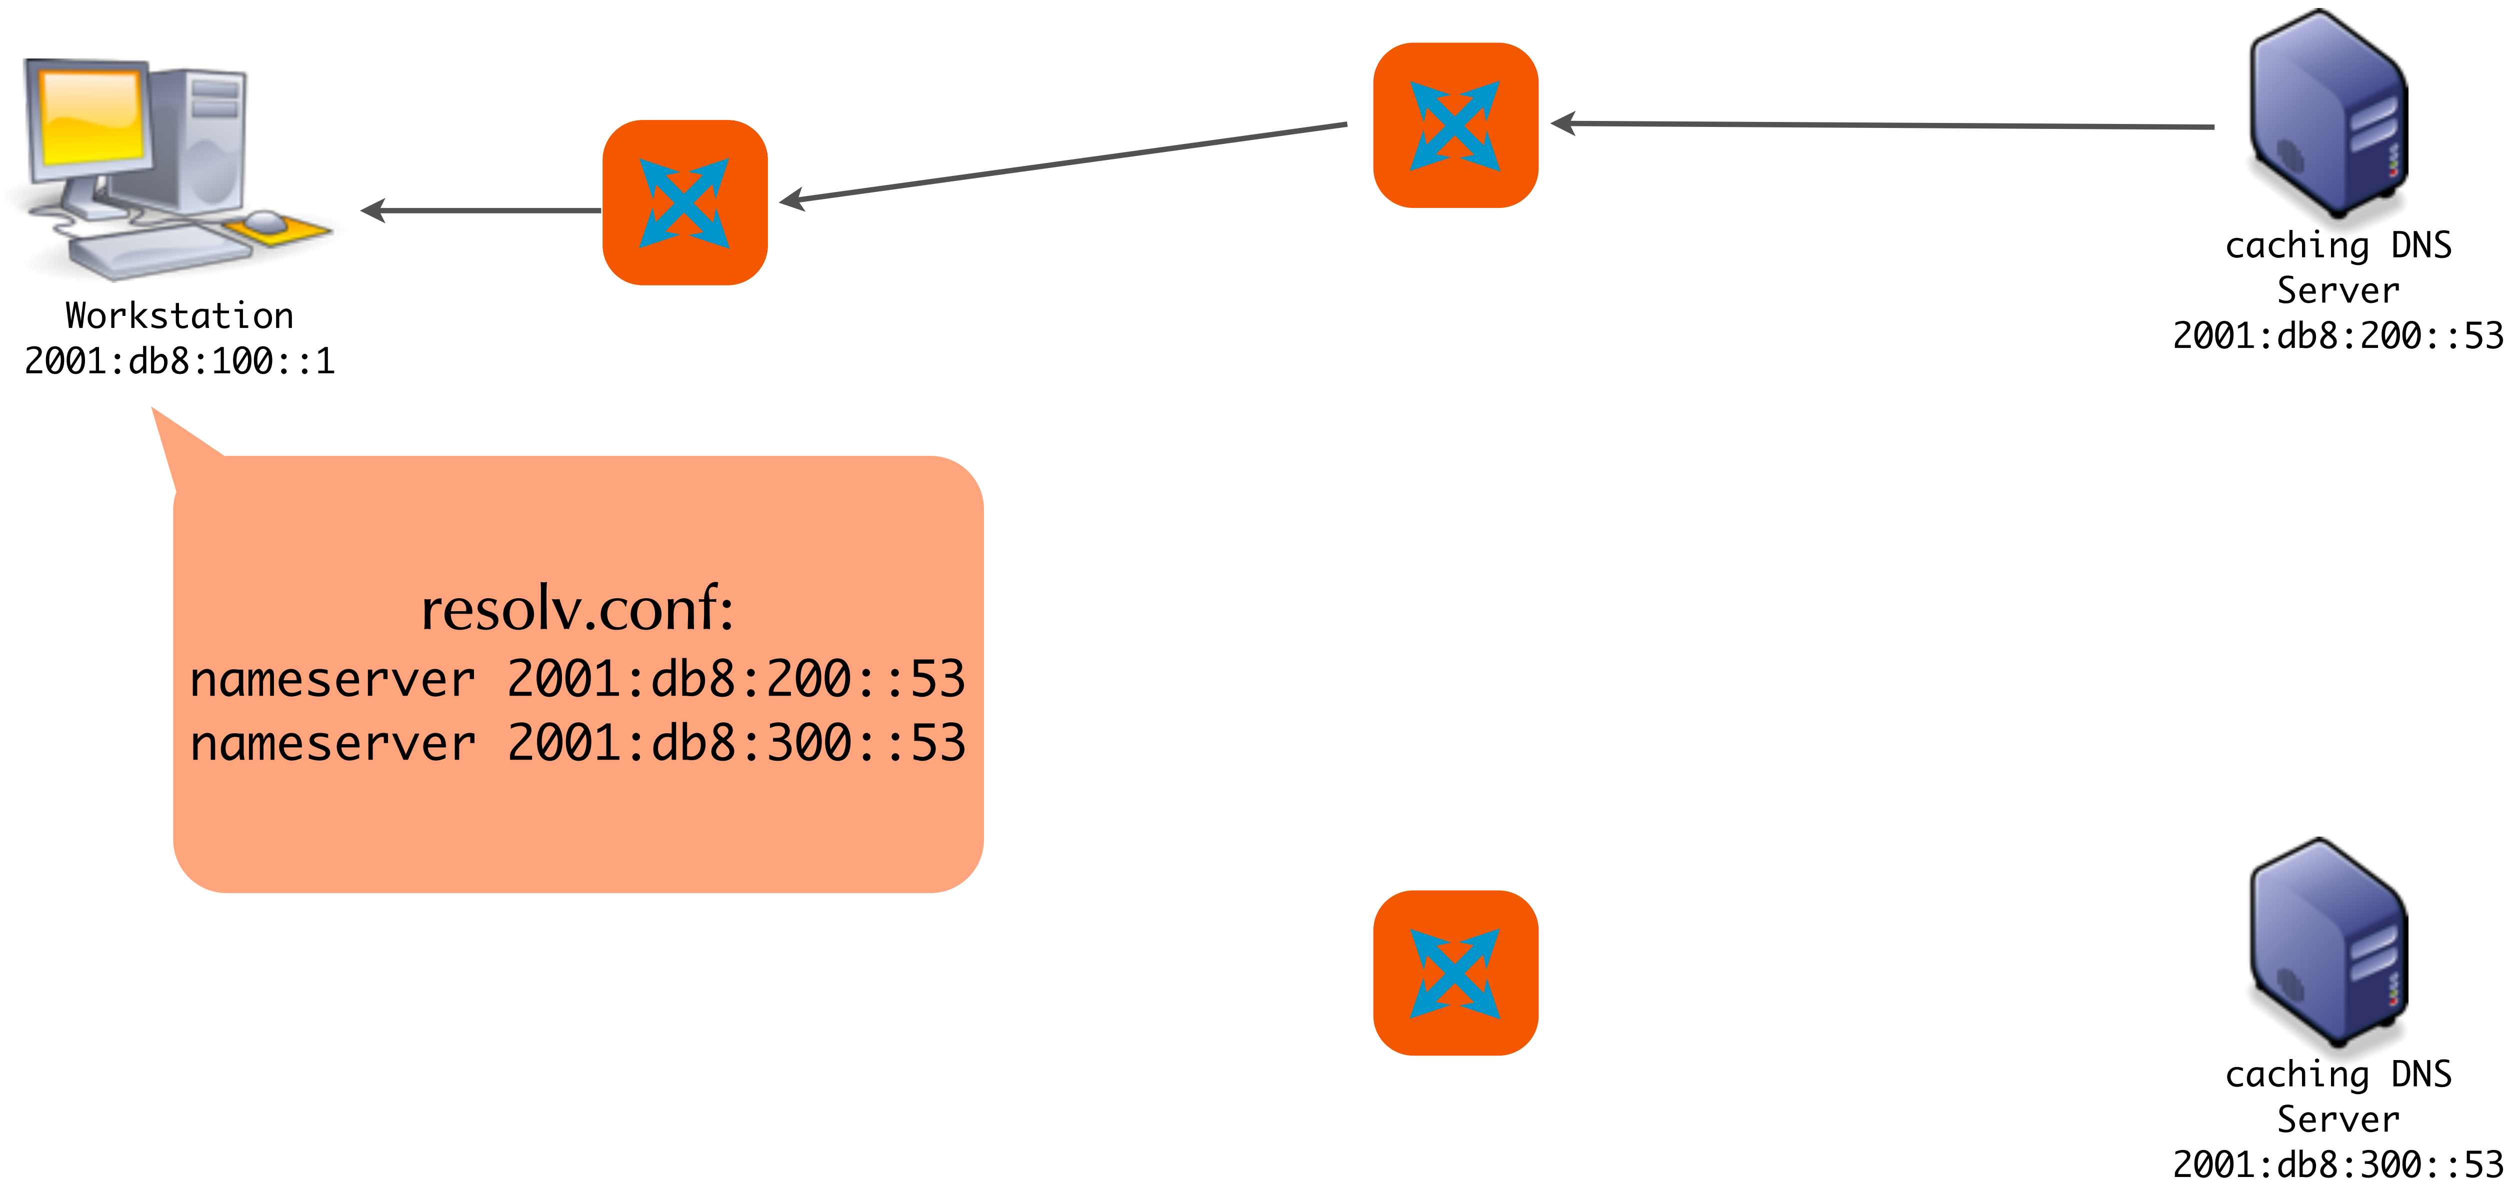

resolv.conf:
nameserver 2001:db8:200::53
nameserver 2001:db8:300::53

DNS using Unicast

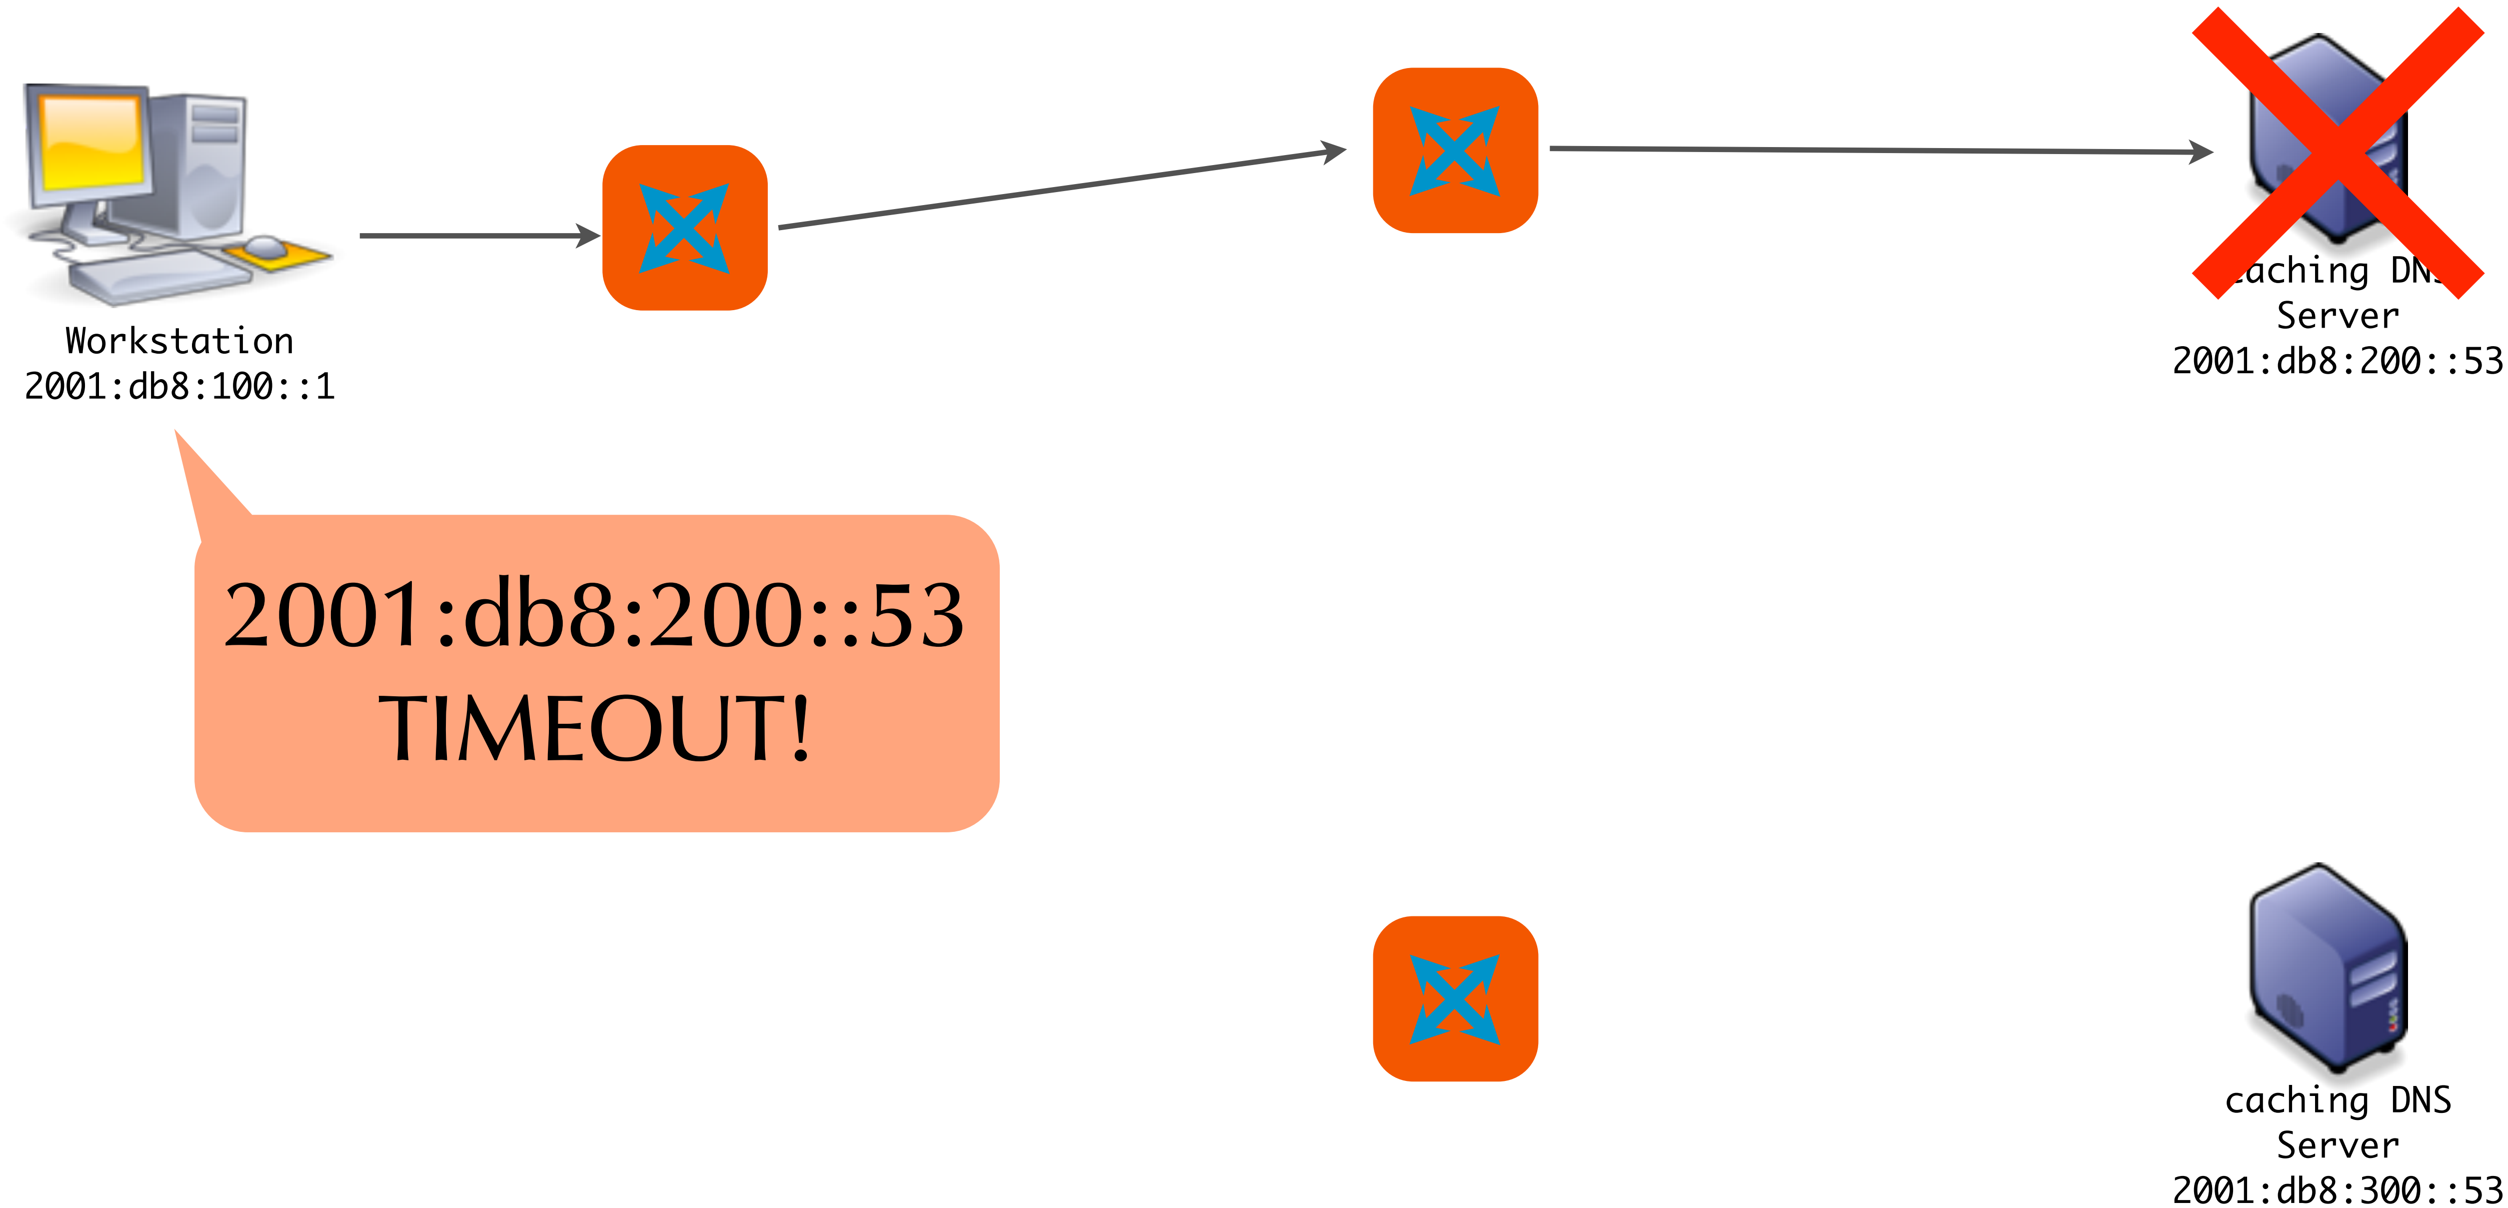

DNS using Unicast

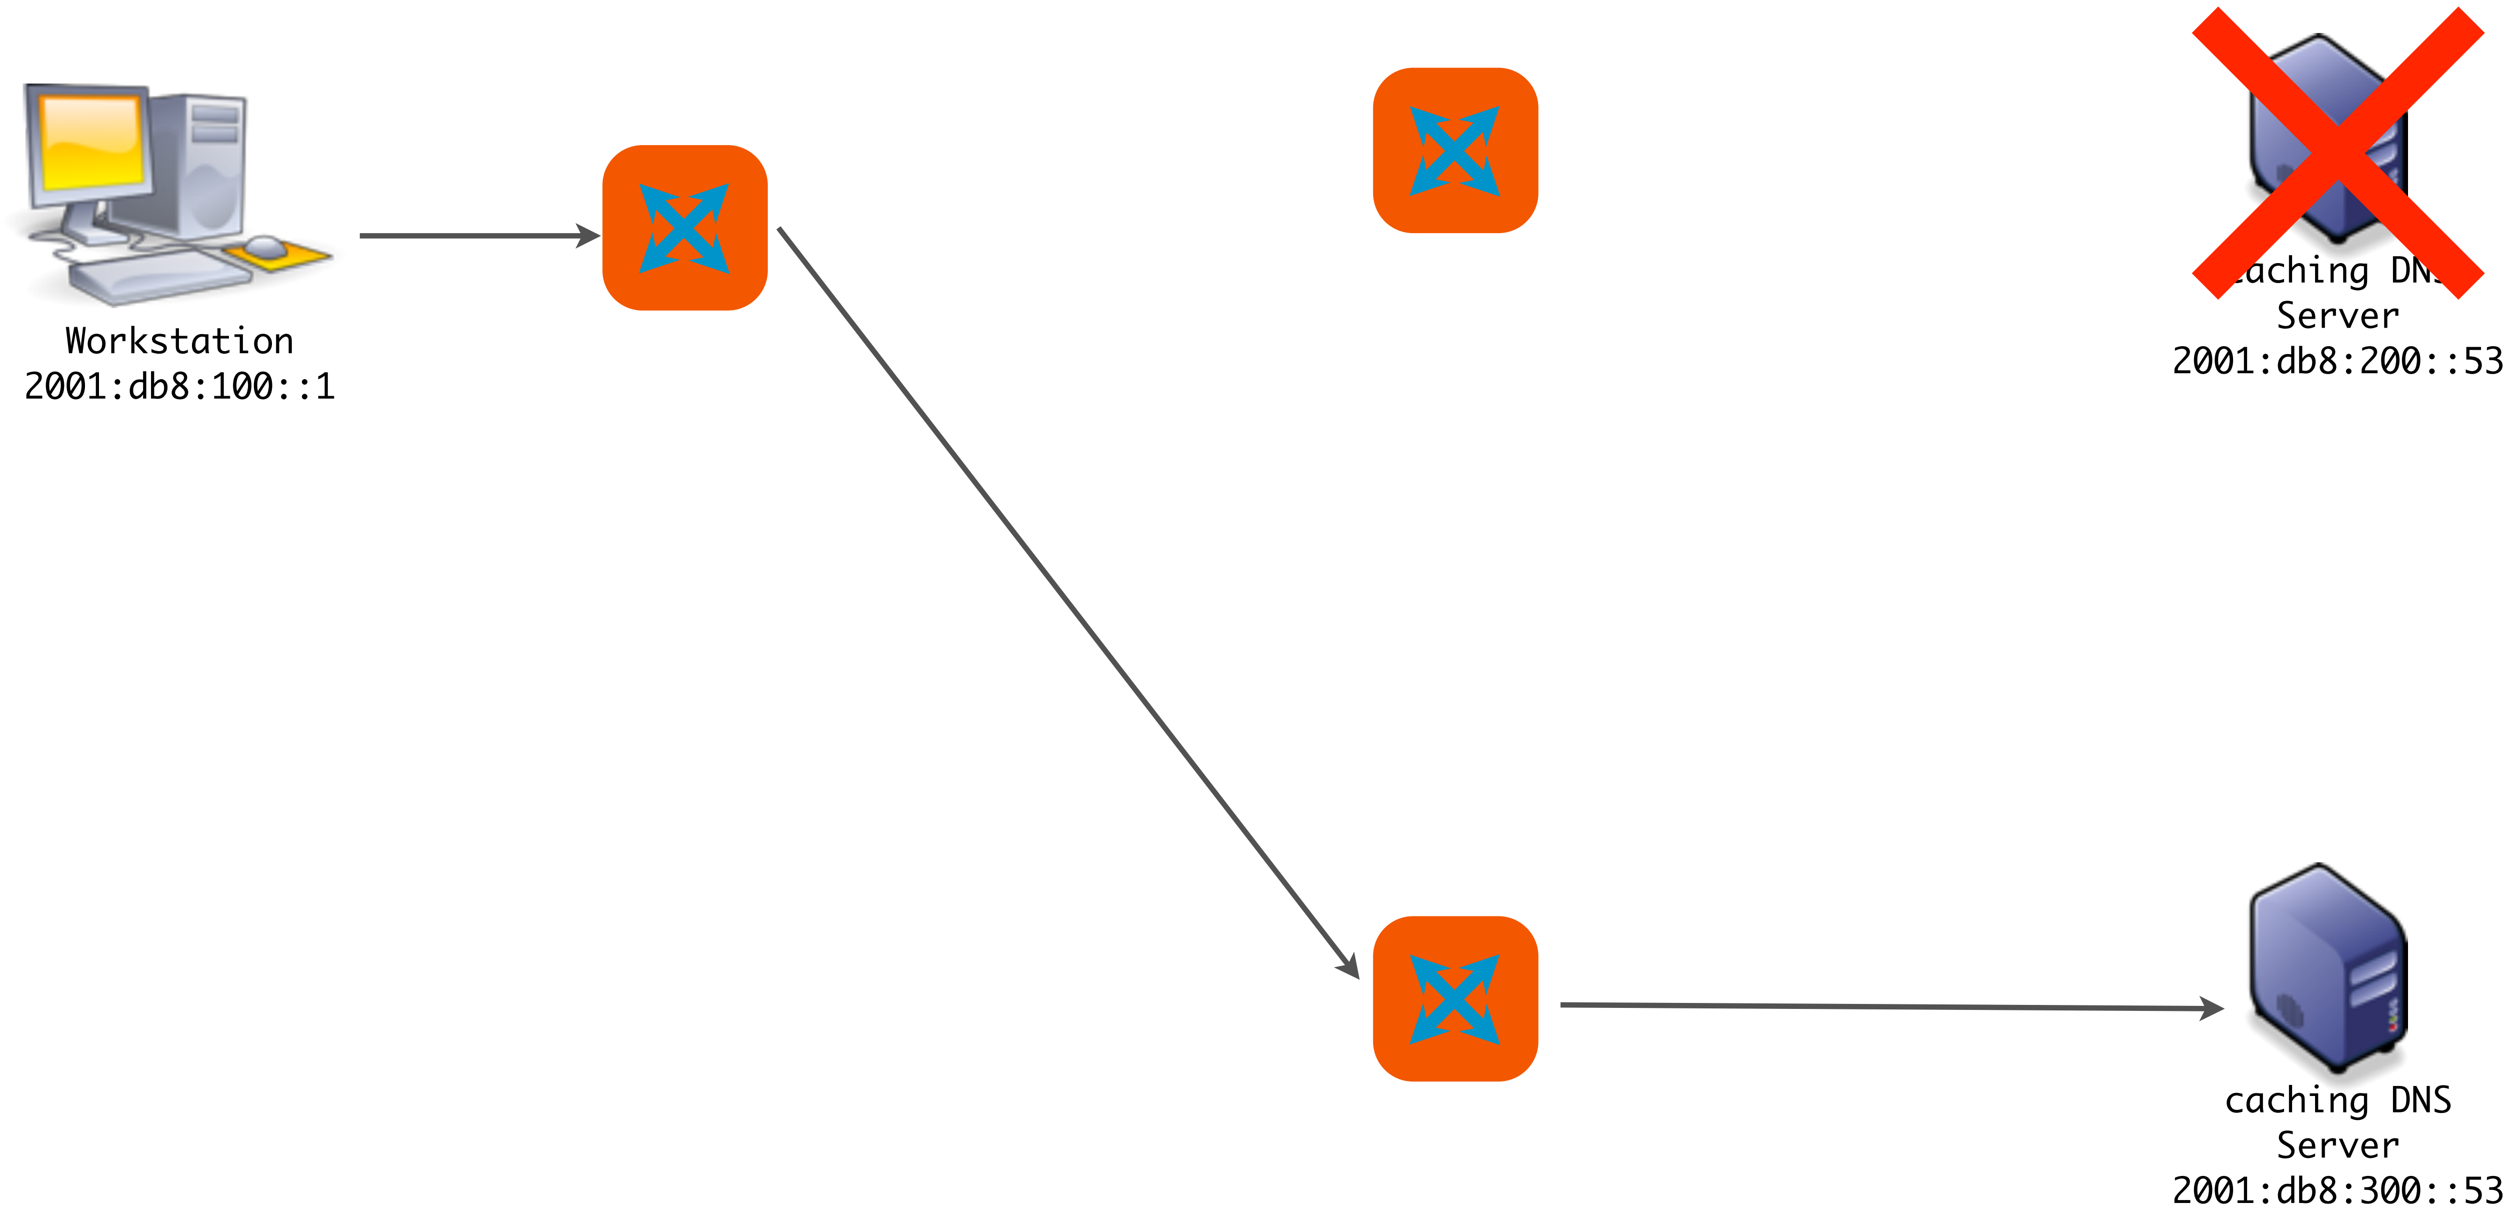

DNS using Unicast

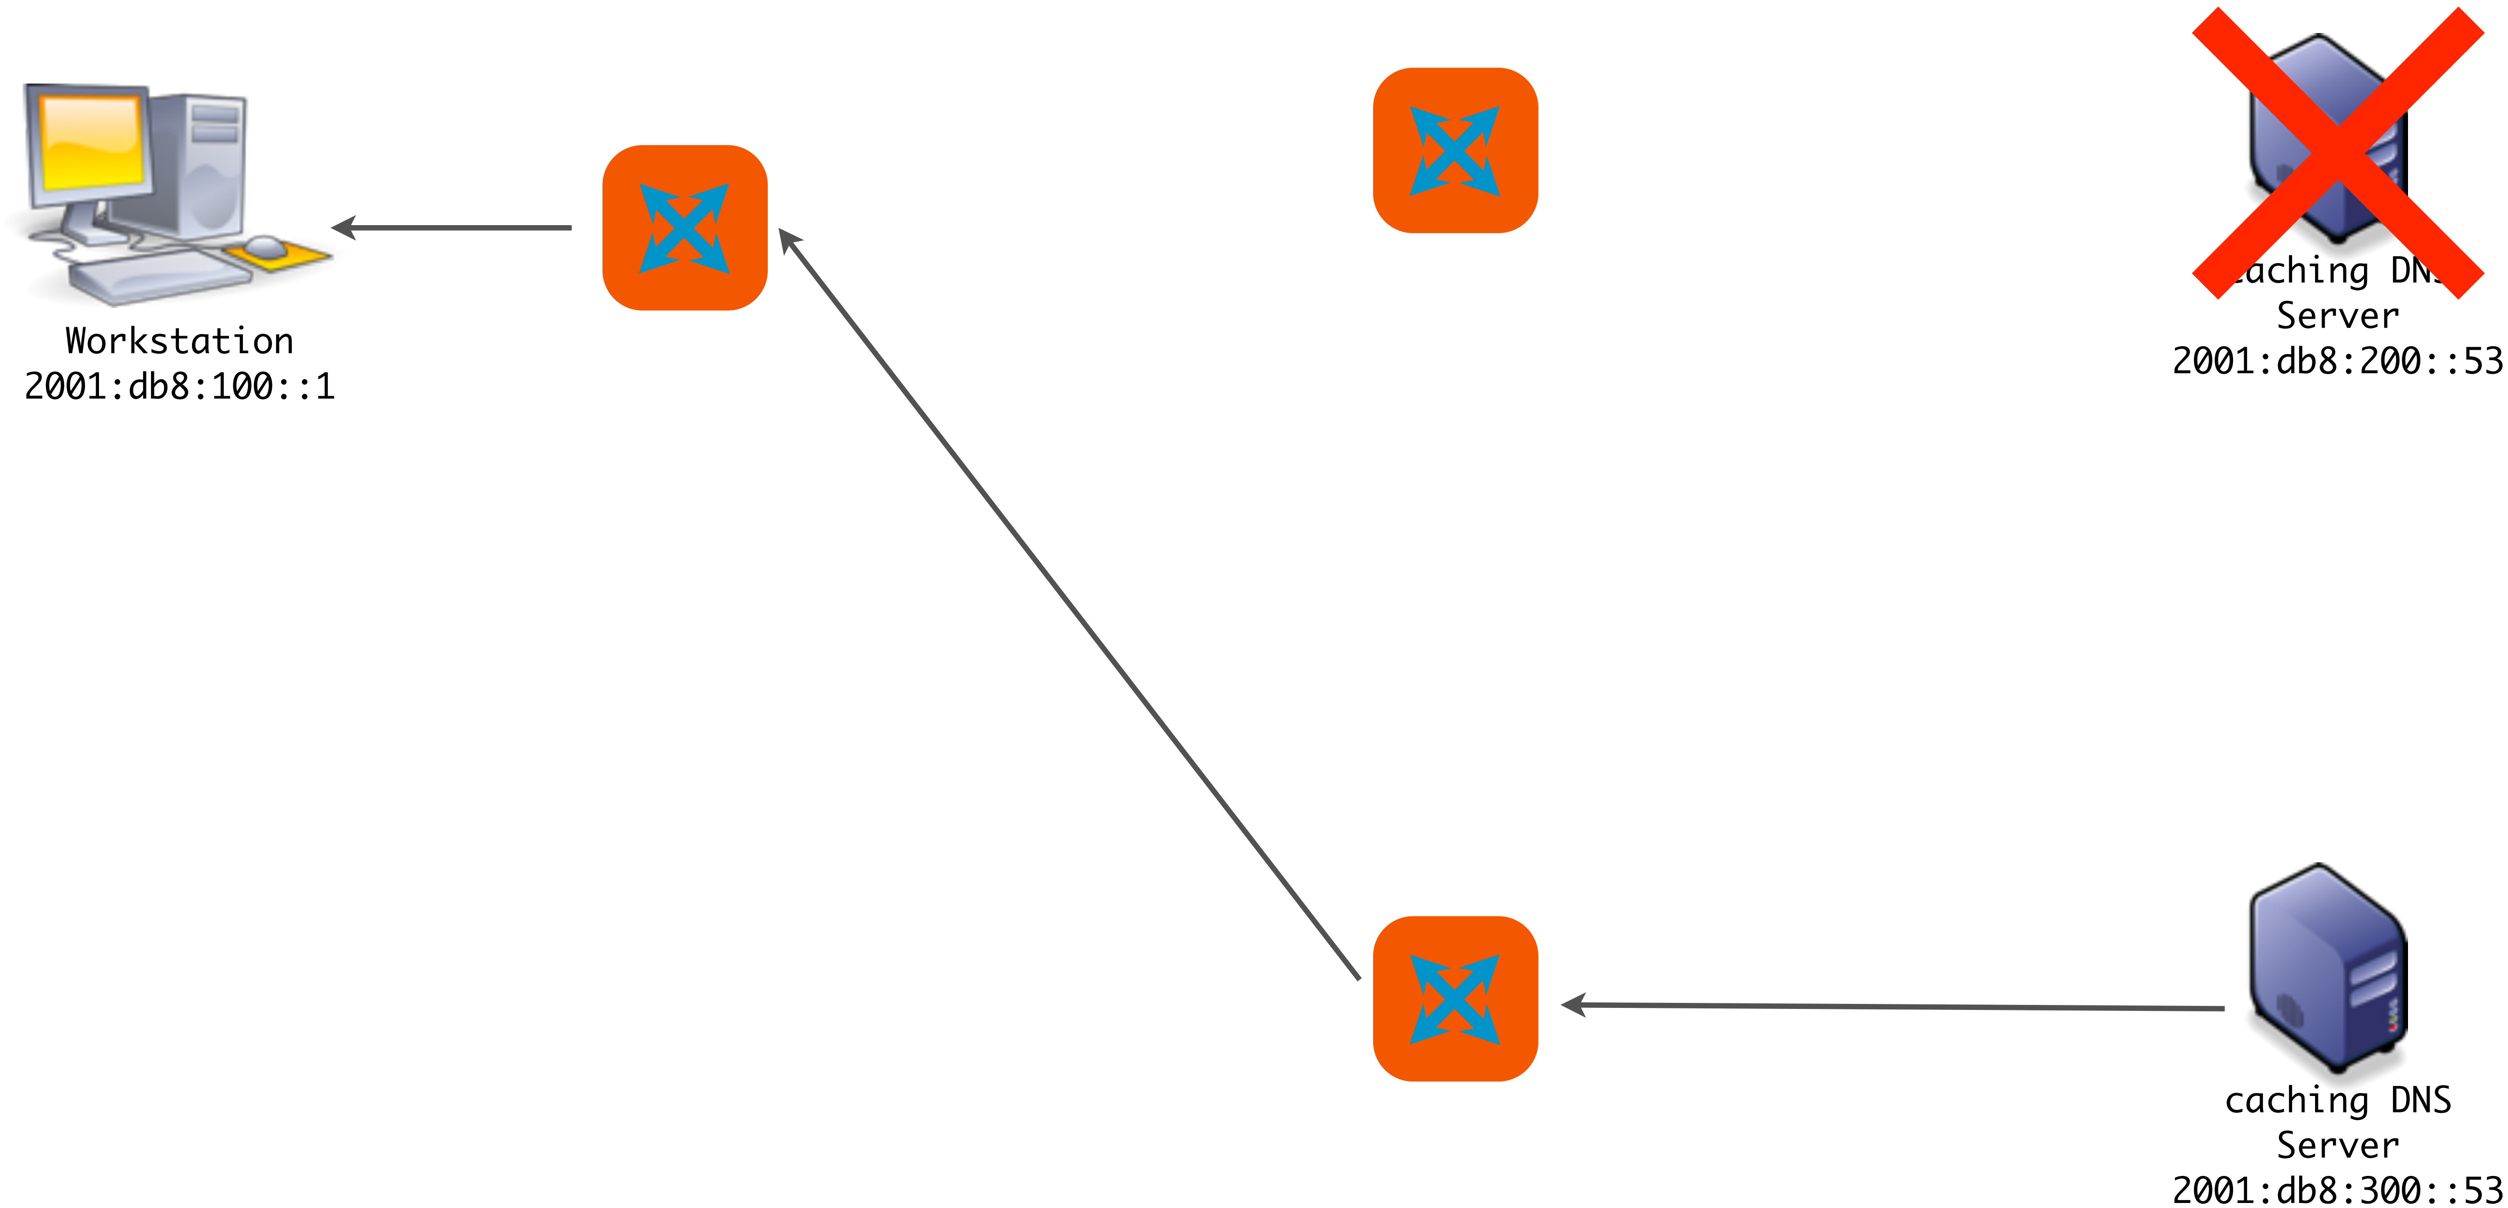

Classic load-balancer solution

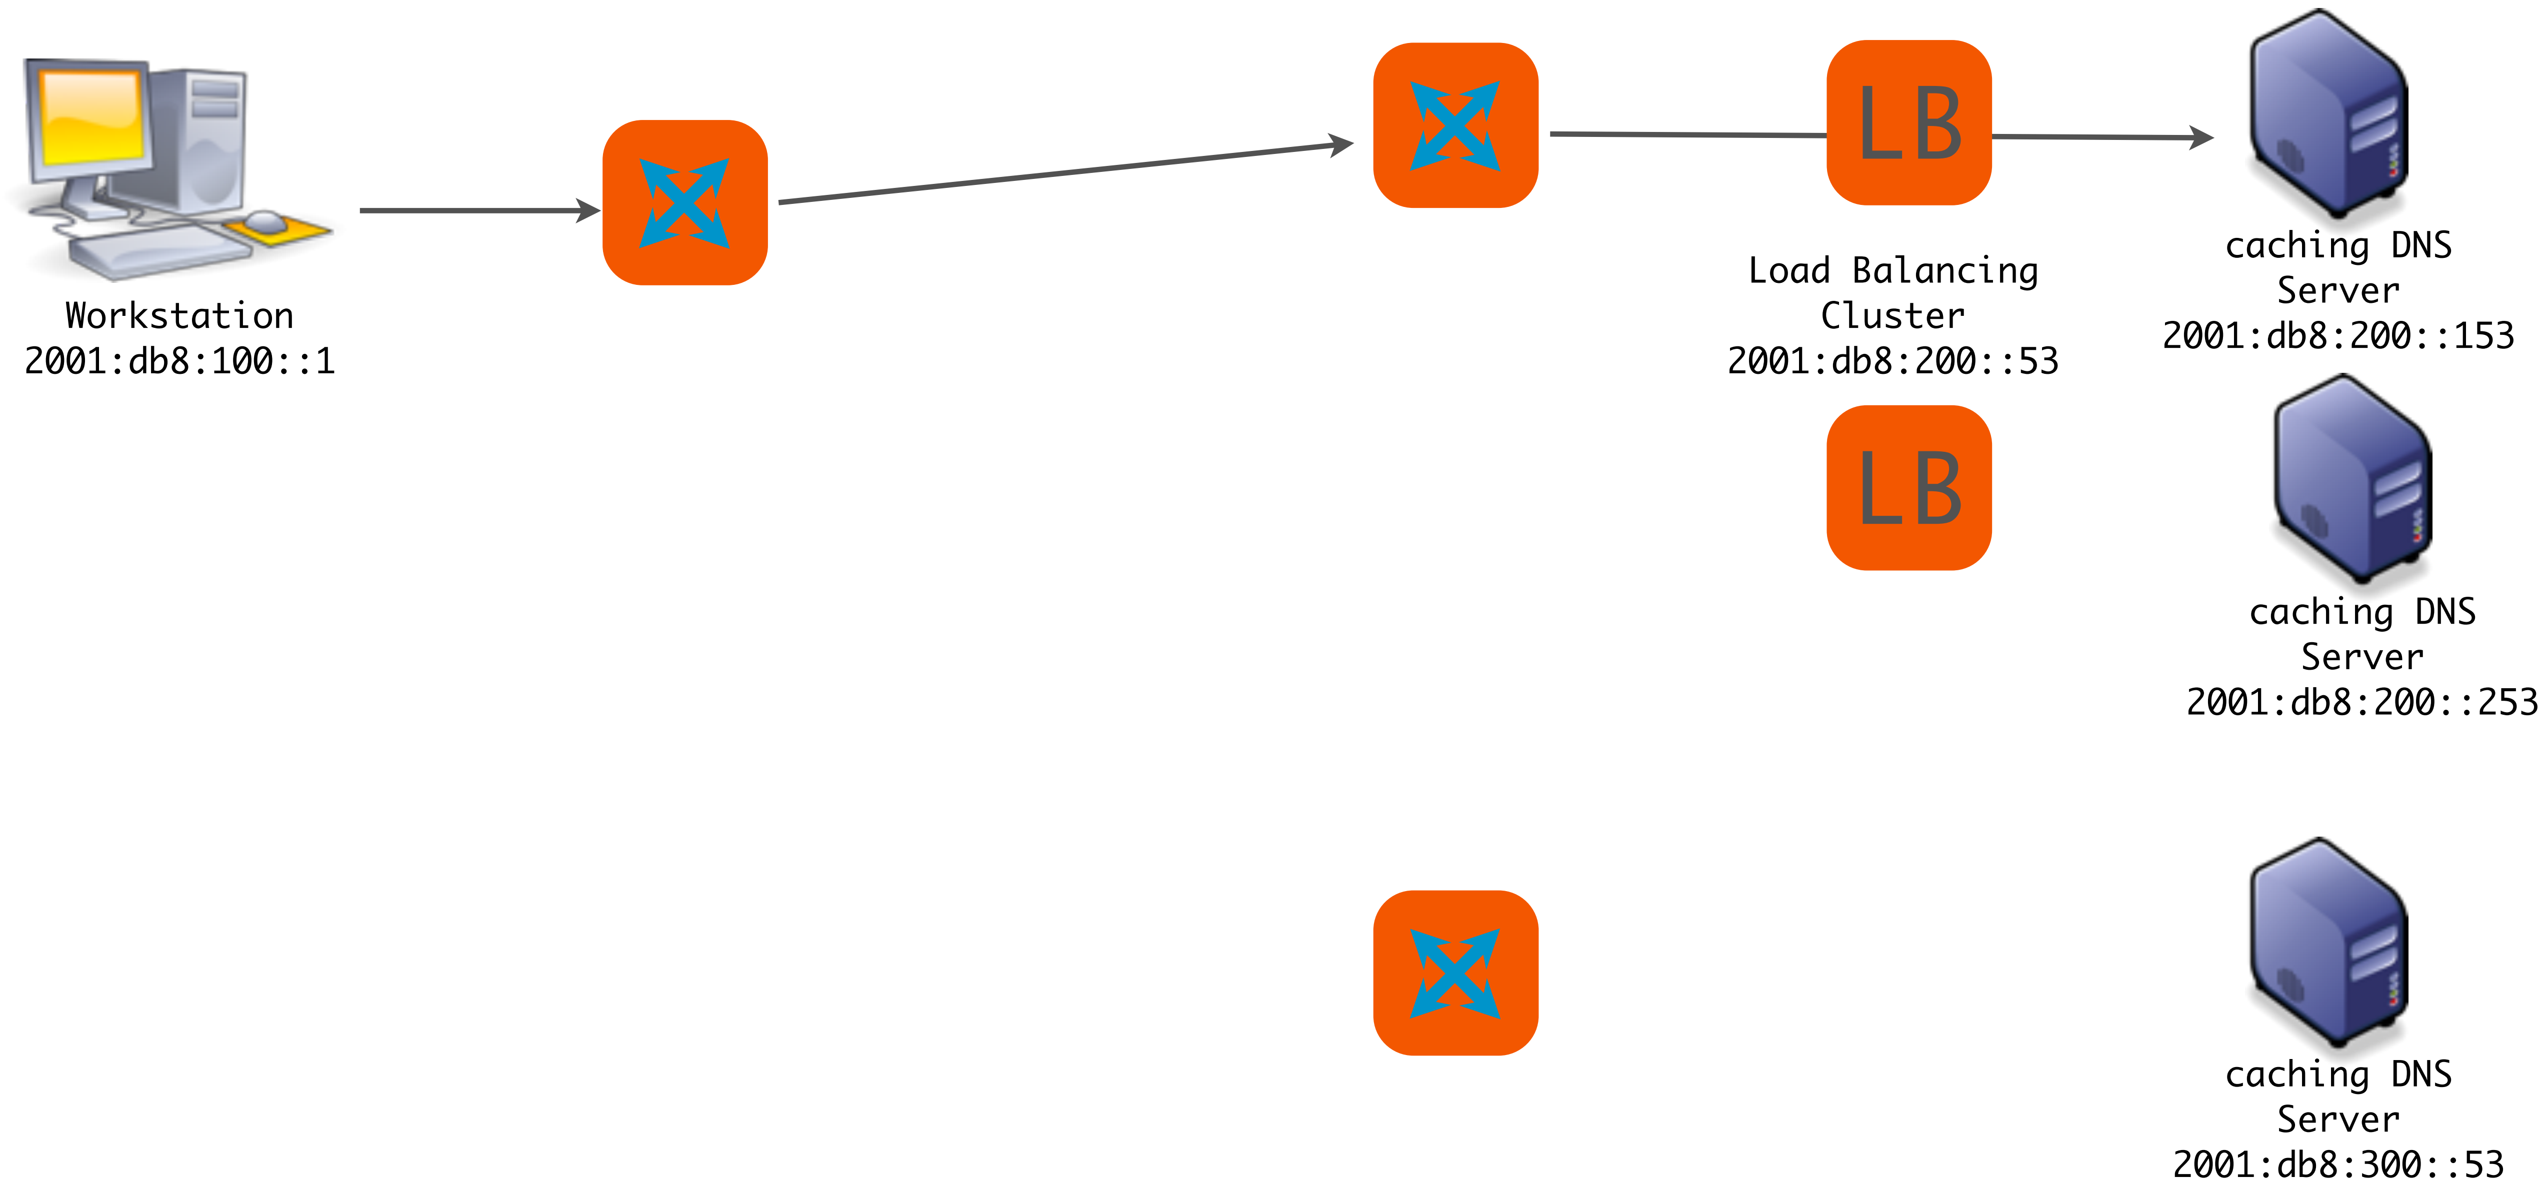

DNS and loadbalancing cluster

- DNS loadbalancing cluster:
  - additional costs
  - higher complexity

DNS using Multicast

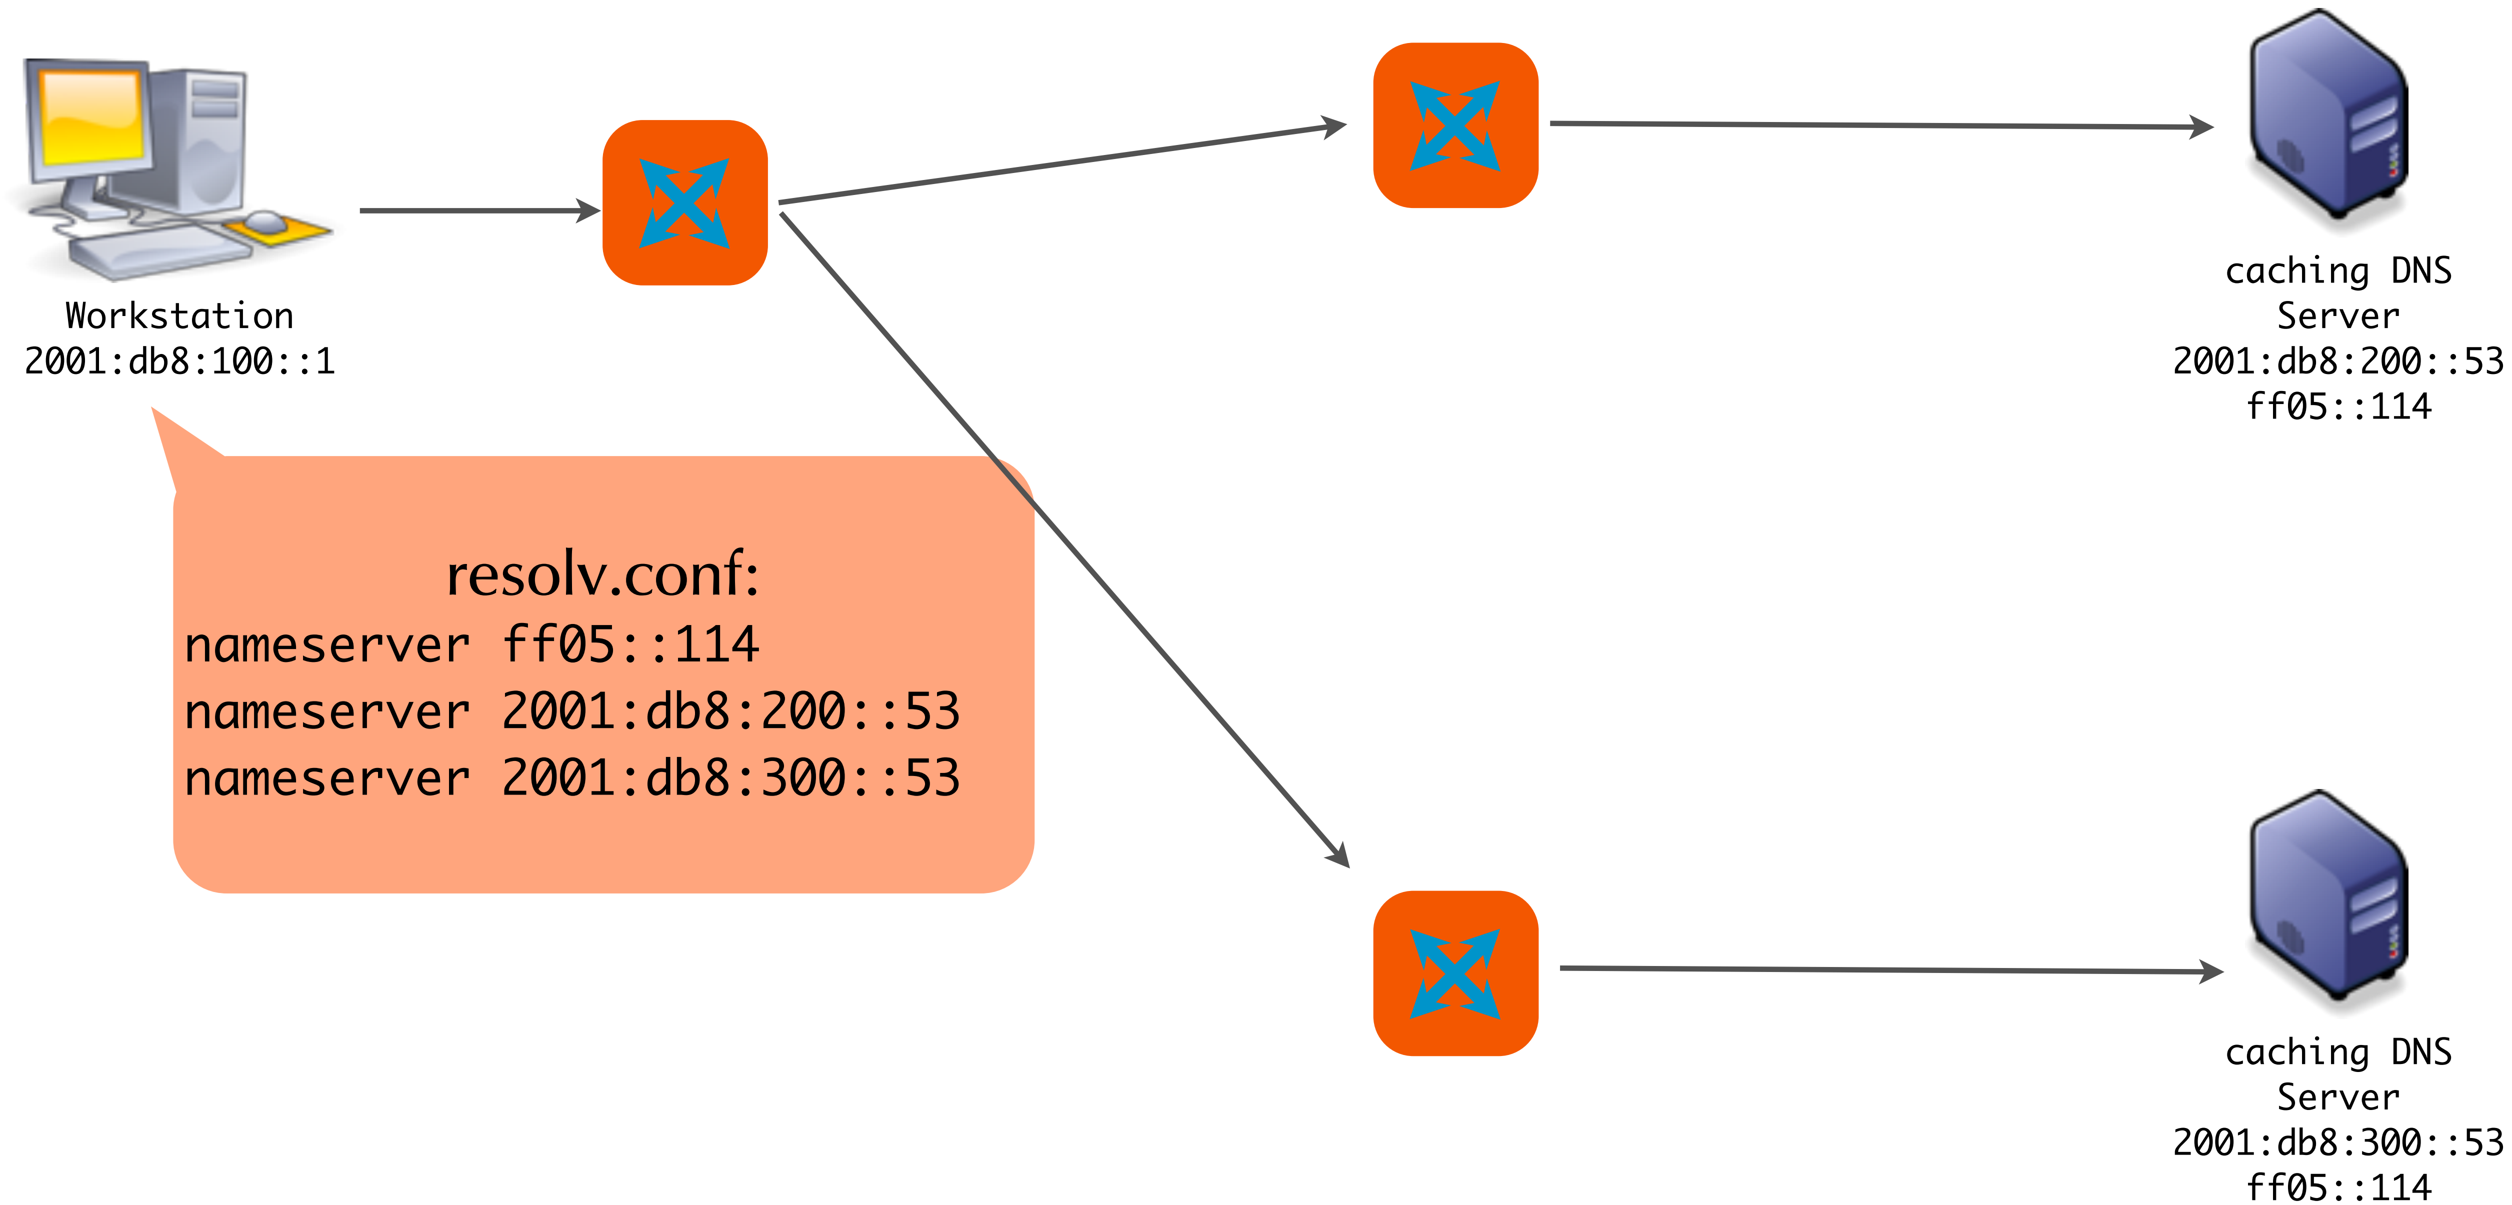

DNS using Multicast

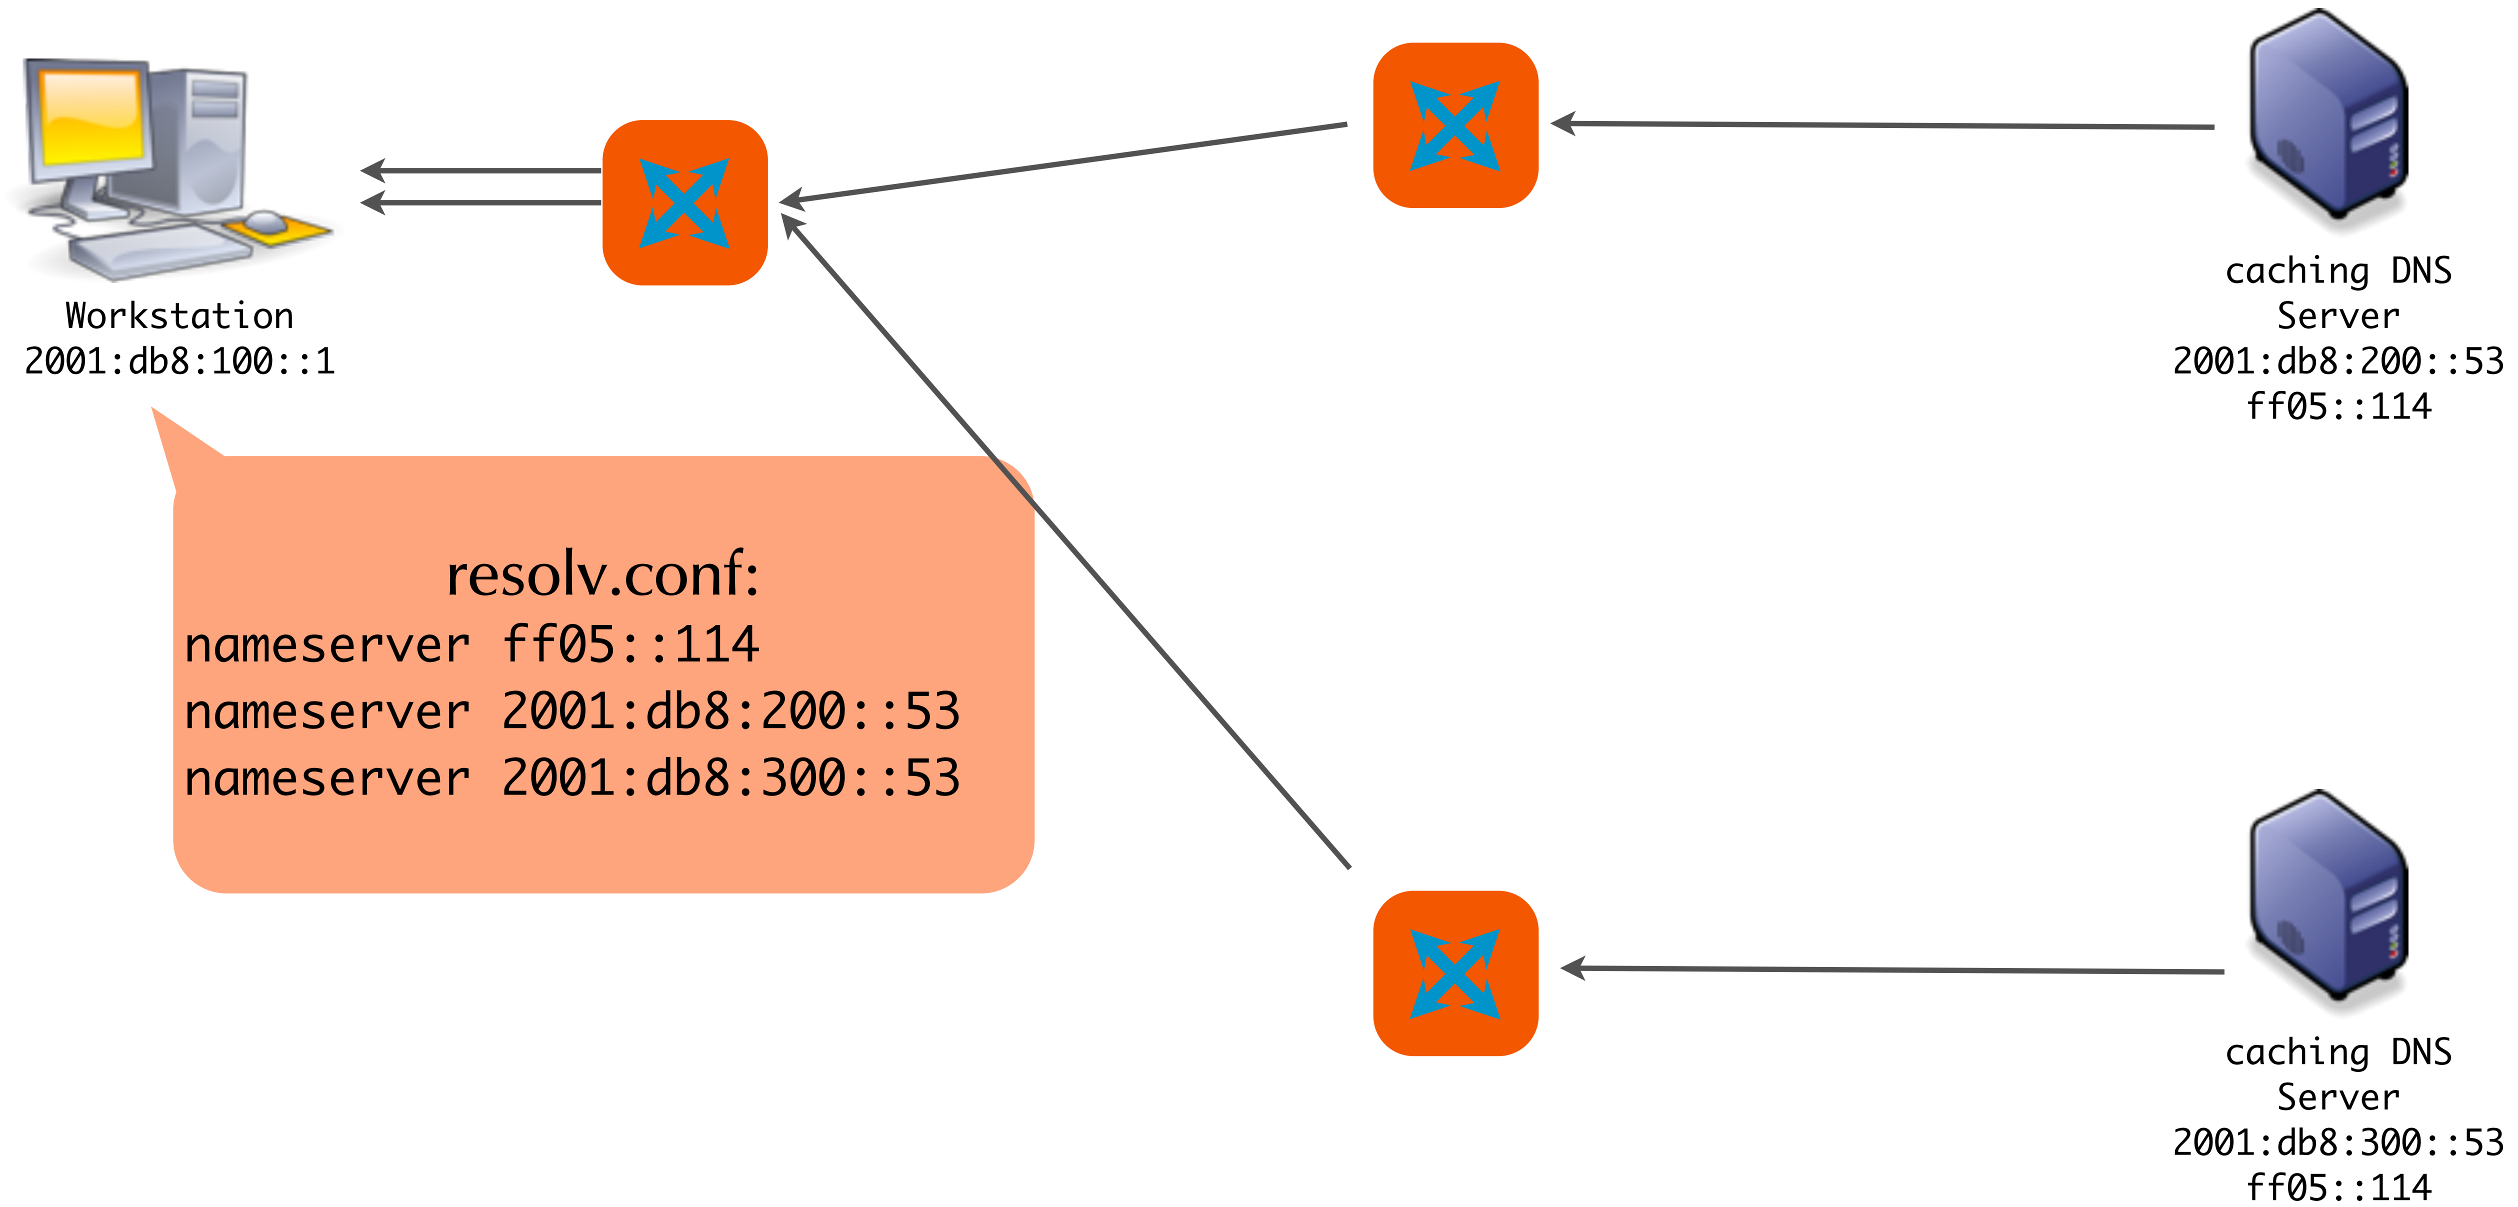

DNS using Multicast

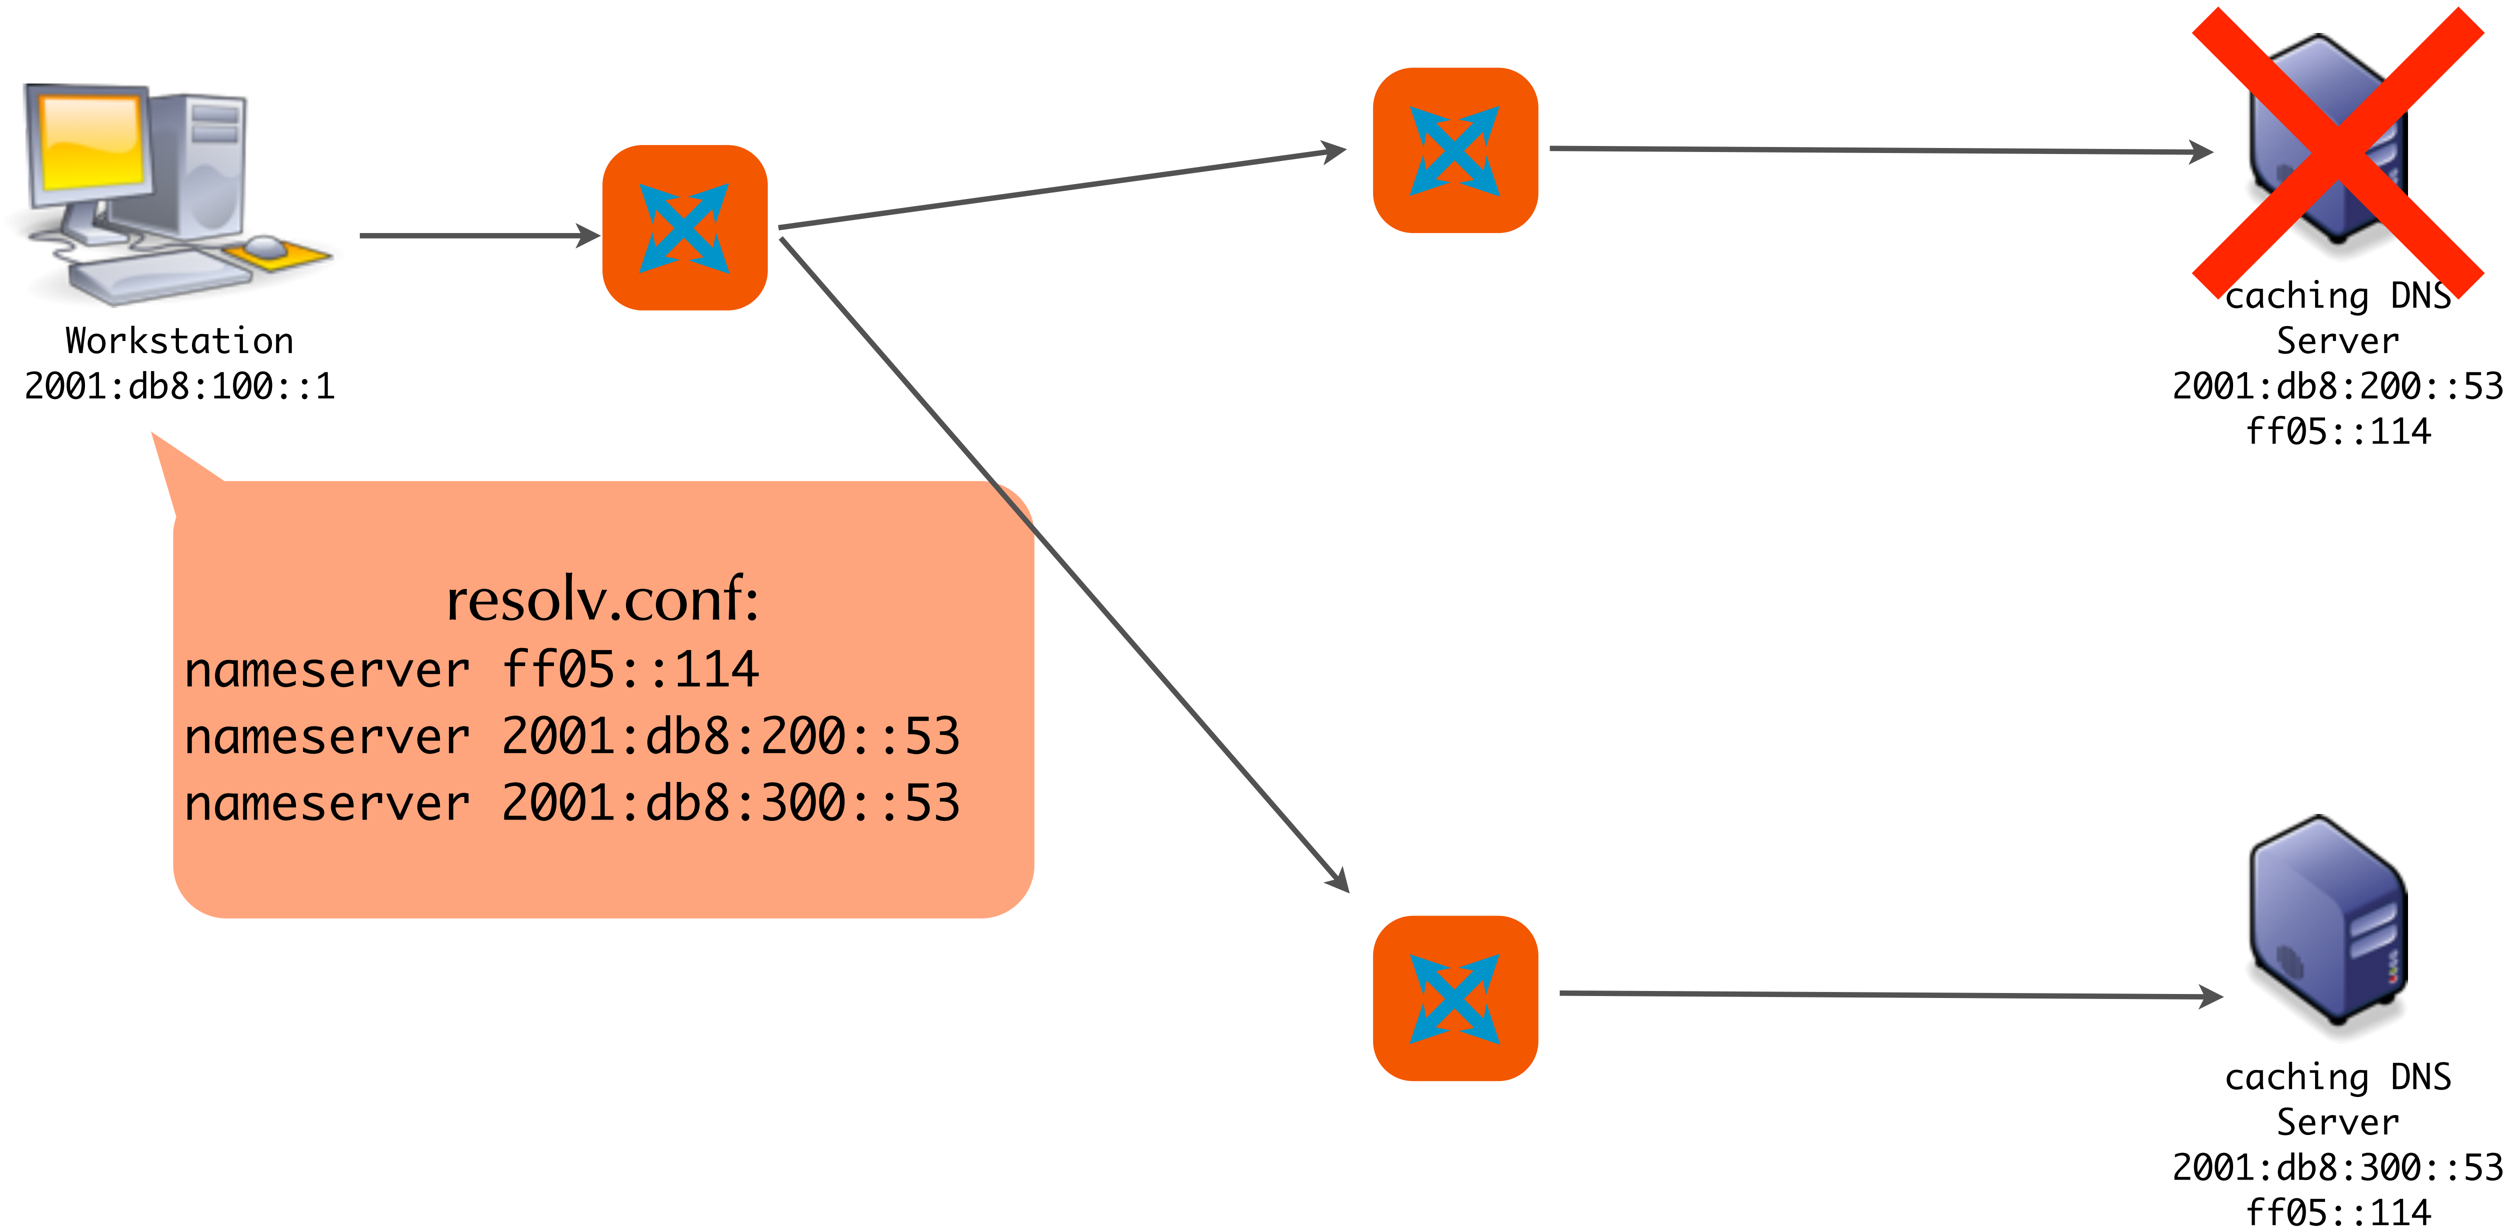

DNS using Multicast

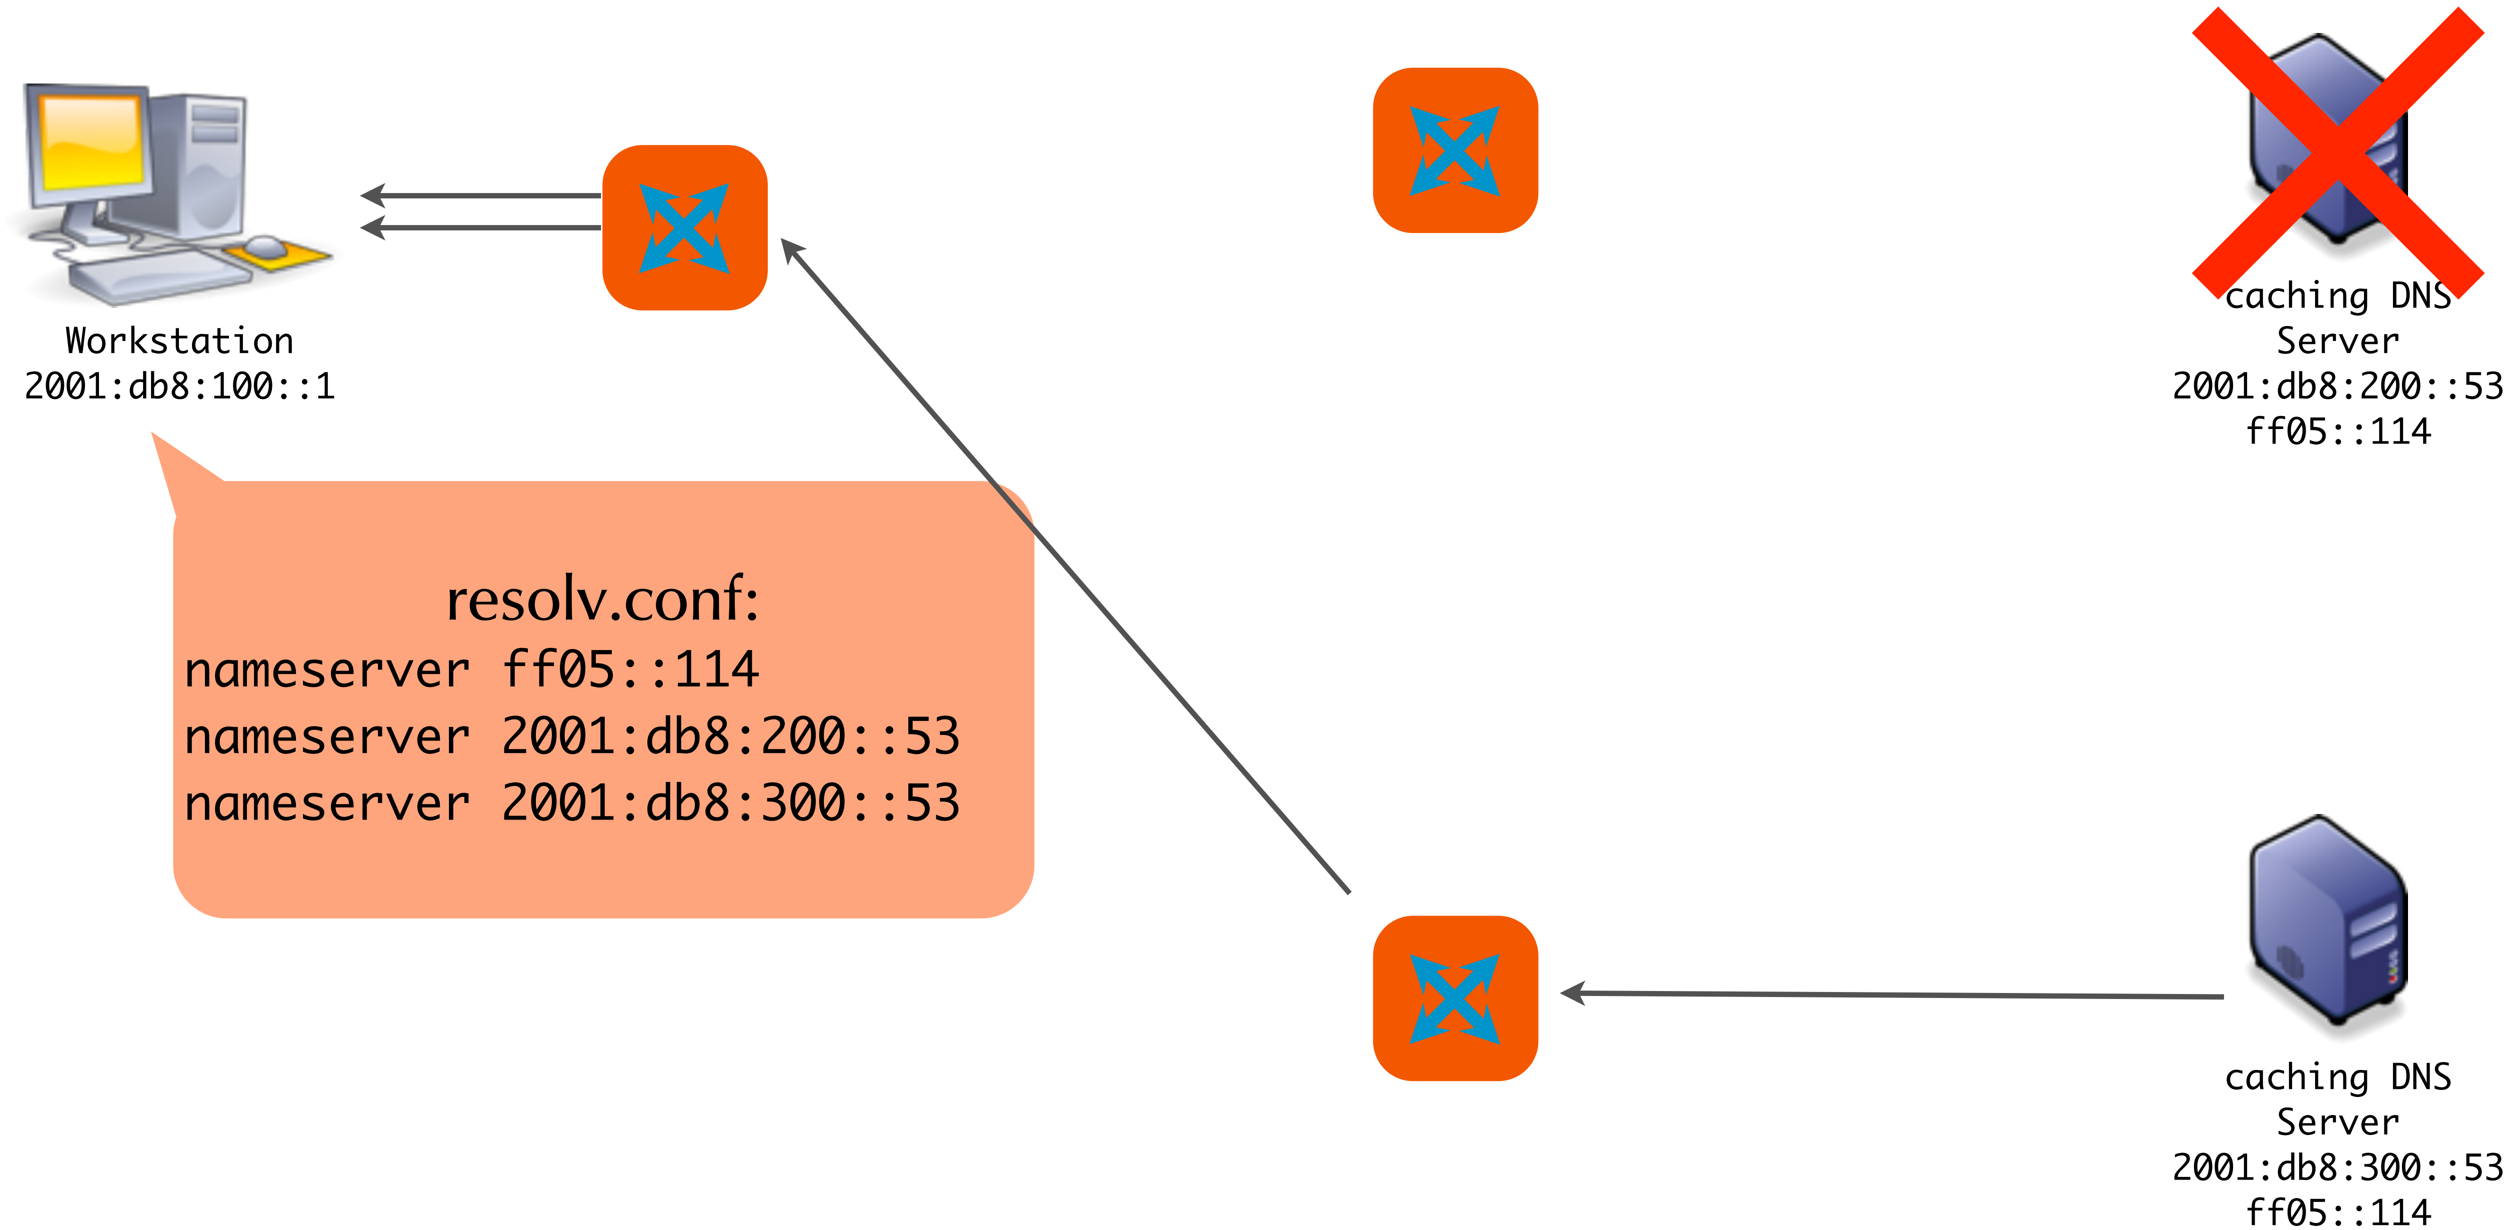

DNS over Multicast - Limits

- Fallback to TCP on large answers (EDNS0 / DNSSEC)
- Security (can be solved with SSM-Multicast)
- look for the “mcdnsProxy” Project
  `https://github.com/dnsworkshop/mcdnsProxy`

Conclusions

- Multicast can be used to simplify network configuration and designs in IPv6 networks
- knowledge on IPv6 Multicast-Routing is required
  - check on multicast support in your hardware (!)
  - but that is not much harder than IPv6 :)

Thank you!

E-Mail:
carsten@menandmice.com

MEN&MICE