IPv4 multicast Setup in Campus Networks

Best Practice Document

Produced by UNINETT led working group
on campus networking
(UFS 130)

Authors: Trond Skjesol, Einar Lillebrygfjeld,
Stig Venås, Vidar Faltinsen

March 2013

Document No: GN3-NA3-T4-UFS130
Version / date: 22.03.2013
Original language: English
Original title: IPv4 multicast Setup in Campus Networks
Contact: campus@uninett.no

UNINETT bears responsibility for the content of this document. The work has been carried out by a UNINETT led working group on campus networking as part of a joint-venture project within the HE sector in Norway.

The research leading to these results has received funding from the European Community's Seventh Framework Programme (FP7/2007-2013) under grant agreement n° 238875, relating to the project 'Multi-Gigabit European Research and Education Network and Associated Services (GN3)'.

Table of Contents

Executive Summary

This document gives a recommendation for multicast setup at the campus level for higher education institutions in Norway. A general introduction to multicast is given. Both Any Source Multicast (ASM) and Source Specific Multicast (SSM) with typical deployment scenarios are covered. Layer 3 and Layer 2 challenges are discussed. Security issues are taken into consideration. Configuration examples for Cisco, Juniper and HP are provided. An overview of web-based and command line troubleshooting tools is included.

1 An Introduction to Multicast

Almost all communications on the Internet today is unicast. At the IP level, each sent packet is forwarded to the destination host identified by the destination IP address in the IP packet header. The IP routers have routing tables specifying where to forward packets based on this destination address.

For multicast this is different. Then the IP destination address refers to a group of IP hosts. For multicast the idea is that a packet that is sent to the multicast group address, should reach all hosts in the group. The principles are the same for both IPv4 and IPv6 multicast. This document focuses on IPv4. For more information on IPv6 multicast, see i.e. [CBPD119].

For IPv4 multicast the prefix `224.0.0.0/4` is reserved. This covers the IPv4 addresses from 224.0.0.0 to 239.255.255.255. Address assignments from within this range are specified in [RFC 5771]. Appendix A lists the most important multicast address scopes.

Multicast is well suited for one-to-many communication. Good examples are IP television broadcasting and distribution of software. Multicast may also be used for many-to-many communication, i.e. for videoconferencing. The advantage of using multicast is that it efficiently uses network capacity. With thousand listeners to a TV program using unicast, a thousand streams are sent over the network. With multicast, only one stream is sent.

For UNINETT, the Norwegian Research and Education Network, multicast traffic is expected to increase substantially as IP TV goes from pilot to production in 2013. There is also a huge potential for multicast-based live streaming of lectures from Norwegian universities and colleges.

1.1 Any Source Multicast and Source Specific Multicast

The basic IP multicast model is described in [RFC1112]. The idea is that any host can join a given multicast group G, and any host can send packets to the destination address of this group, and have it delivered to all group members. The sender itself does not need to be a member of the group. This classical model is referred to as "Any Source Multicast" (ASM).

Another model is "Source Specific Multicast" (SSM). With SSM a host joins the tuple (S,G) where S is a unicast IP address, and G is a multicast group. Here S specifies a specific source S, and by joining (S,G) the host specifies that it wants packets from source S and not any other source.

Another way of explaining this, is that when a source S sends a packet, it should reach all hosts that have either joined G or (S,G). The pair (S,G) is referred to as *a channel*. A channel can have many listeners, but only one sender. Note that there are specific address ranges for SSM use (see appendix A).

SSM is quite useful when there is only a single source, for instance a television broadcast. The source address can be announced together with the group address, and it solves possible issues with other people sending to the same group. A TV broadcaster would typically want to avoid others from also sending to the channel/group they are using. In general SSM is fine with a relatively small and fixed set of sources. Each receiver must in

some way be informed what the set of sources is, and possibly allow for this set to change at any time. This is done at the application level. With ASM the whole source discovery problem is solved at the network level, which simplifies applications with many sources, like e.g. video conferencing, quite a bit. For more information on SSM, see [RFC3569].

Even though SSM has obvious advantages compared to ASM, ASM is still by far the most deployed model. In UNINETT all IP TV services have so far been using ASM. This may change in the future, as IGMPv3 support in operating systems and applications keep improving, and as the IGMPv3 snooping capabilities in the layer 2 networks are more complete.

1.2 IGMP – the protocol between hosts and routers

In order for a host to send multicast, no protocol between host and router is needed. The host simply sends the multicast packets out on the network and provided the TTL (or hop limit) is greater than one, it will be forwarded by a multicast router.

In order to receive multicast, a host needs to join the multicast group it wishes to receive from. In the case of SSM, the host should also specify the source address. This is done using IGMP for IPv4 or MLD for IPv6. A multicast router will periodically send queries to the hosts on the subnet asking which groups they are listening to. A host should send a join message when it joins a new group and a leave message, if possible, when it leaves. For a host to specify both source and group as needed by SSM, IGMPv3 for IPv4 or MLDv2 for IPv6 is required. Both IGMPv3 and MLDv2 are backwards compatible with earlier versions. For more information on IGMPv3 and MLDv2, see [RFC4604].

Note that a multicast router does not care exactly which hosts are members of which groups, nor how many. All it wants to know is whether there is at least one listener for a given group or a given source/group pair behind a given router interface. If so, the multicast packets for the group in question will be forwarded.

1.3 Multicast routing

Routing multicast packets is relatively complex. There are several protocols with different pros and cons. We will describe how IPv4 multicast is deployed in UNINETT.

1.3.1 Intra-domain routing

The most common solution for intra-domain routing is Protocol Independent Multicast (PIM). PIM makes use of the unicast routing table, and is agnostic to which unicast routing protocols are used to populate the table. There are two variants of PIM, where UNINETT uses PIM Sparse-Mode (PIM-SM), see [RFC4601]. It is called "sparse mode" because it only forwards multicast packets where there are registered listeners. There is also a dense mode variant (PIM-DM) that regularly flood the network with multicast packets, and rely on "prune" messages from those that do not want to receive it.

In order to support ASM, PIM-SM takes care of source discovery by using a rendezvous point (RP). This is a router within the PIM domain, or to be precise, it is a unicast IP address defined as an extra /32 loopback interface at the rendezvous point router. The RP IP address must be reachable within the entire multicast domain.

Simply said, the RP keeps track of which multicast sources exist, and which multicast groups have listeners. Initially when a new source starts sending, a multicast router on the same link as the source will start sending "PIM register" messages to the RP (unless it is sending to an SSM group). They are sent as unicast and contain the multicast packet. The RP will then usually send a "PIM register-stop" message back. If it knows of any listeners it will before sending the stop message, send "PIM (S,G)-join" messages towards the source, and

make sure it receives the multicast packets natively. The routing table is used to determine where to send the (S,G)-joins, using S. The RP will send it to a neighbouring PIM router, who again will send a new join towards S. Finally it will reach a router on the same link as S, and the packets sent by S will be forwarded along the path built by the joins. Note that PIM "register" and "register-stop" messages are sent as unicast directly to the RP or the edge router respectively. Other PIM messages are sent as multicast only on the link, only reaching other PIM routers on that link.

When a PIM router learns that it has directly connected hosts listening to a group G or receives a (*,G)-join from another router, it will itself send a (*,G)-join towards the RP. Each router will look up the RP-address in the routing table to determine where to send the join. This builds a so-called shared tree, or RP tree (RPT), from the RP to the receivers. The multicast packets received by the RP will be forwarded down this tree, see the RPT part of figure 1.

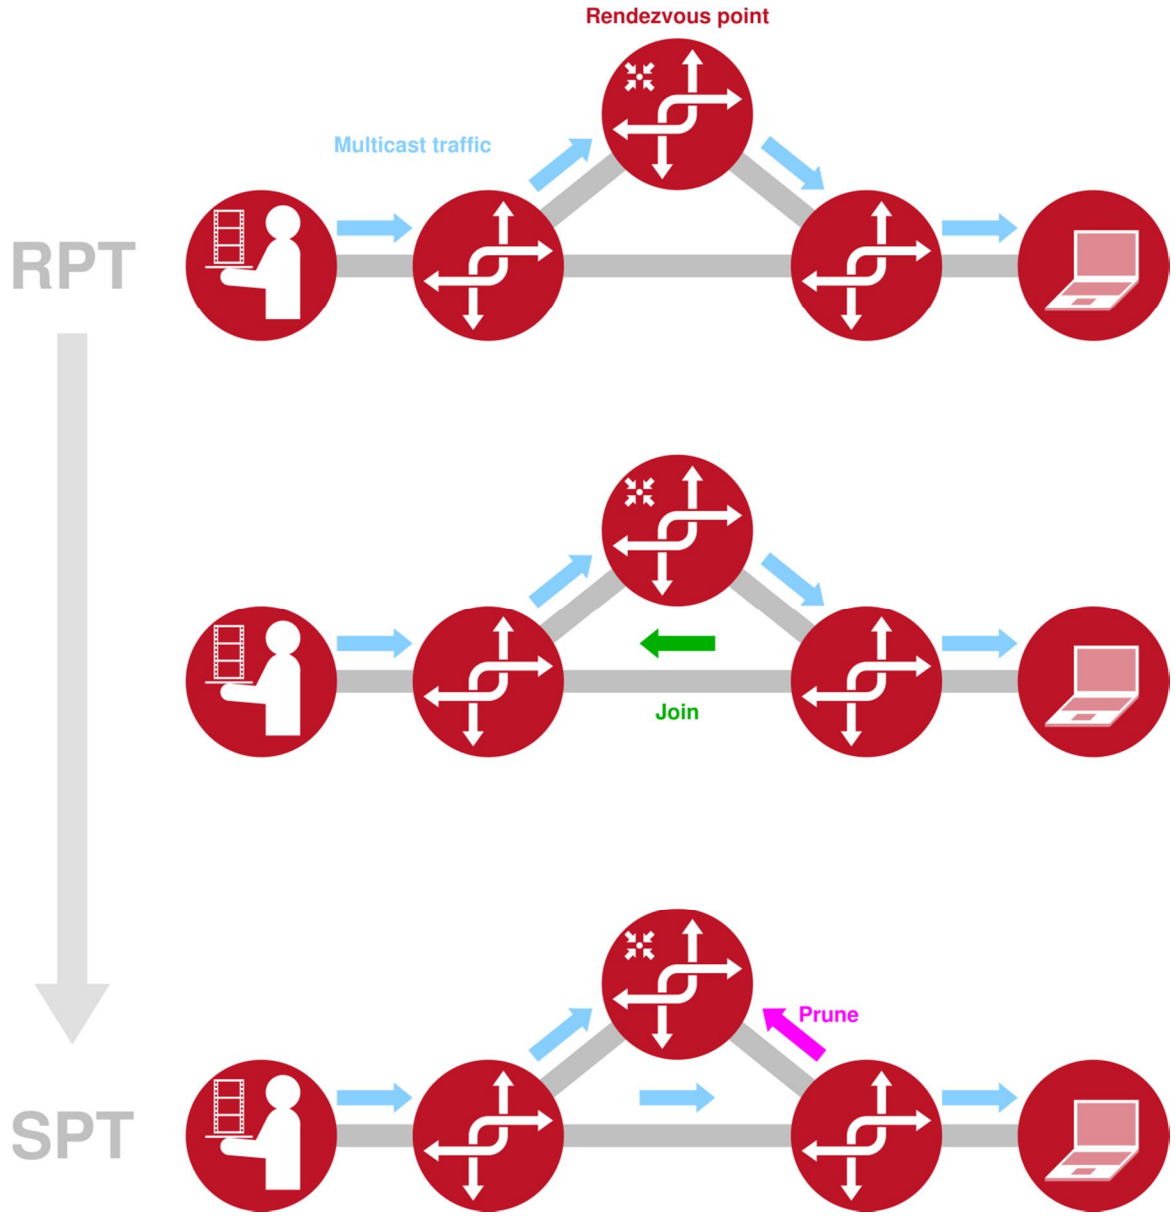

Figure 1: Initially the multicast stream follows the RPT, but may converge to a shortest-path tree (SPT)

When an edge router, one with connected listening hosts, starts receiving packets on the RPT, it may as an optimization build so-called shortest-path trees (SPT's) towards the sources (see figure 1), instead of receiving data through the RP. It does this by sending (S,G)-joins towards the source, similar to what the RP does, and when it starts receiving on the SPT, it can prune the source from the shared tree. Some routers do this when

they receive the first packet from a new source, others do it if the data rate from the source is above a certain threshold, and some never do. Note that it is possible to use SPT for some sources of a group, and use RPT for the other sources. Also, if there are hosts joining explicit sources using SSM, the router can build the SPT's at once without first receiving packets from the RP. Thus if all hosts used SSM, there would be no need for the RP.

There may be more than one RP within a PIM domain. All PIM routers in the domain need to know which RP to use for which multicast group. In order to have full connectivity throughout the domain it is essential that all PIM routers use the same group-to-RP-mapping. To keep this mapping consistent, the Bootstrap Router Protocol becomes useful.

1.3.2 Bootstrap Router Protocol

In the Higher Education (HE) sector in Norway, UNINETT have implemented a model with a sector-wide RP and local RPs at each campus that are used for local groups (i.e. for software distribution). This is illustrated in figure 2. The local RPs are set up using the same IP address on each campus. This is important in order to maintain a consistent group-to-RP mapping for sector-wide and local groups.

Historically the multicast design in UNINETT was based on a static configuration with access lists on each router defining which multicast groups should be kept local and which should be sector-wide. As new multicast groups constantly entered the scene, this solution became hard to manage. Today UNINETT has implemented the Bootstrap Router Protocol (BSR). BSR is standards-based dynamic protocol available with PIMv2. The protocol performs the same function as Cisco's proprietary Auto-RP, i.e. disseminates group-to-RP information, see [RFC5059]. With BSR UNINETT can centrally maintain the multicast group-to-RP mapping and through PIM distribute this to the customer networks on campus. This significantly simplifies the configuration at the campus level.

It should be noted that UNINETT's implementation of BSR is uncommon. Typically BSR is used within a single administrative domain with a BSR border setup towards the customer networks. This would in turn involve MBGP and MSDP peerings (see 1.3.3) with the customers. Based on the good working relations in the Norwegian HE sector, the UNINETT design aims for simplicity rather than formality. There are security considerations in the design, see section 2.7, but these are also manageable.

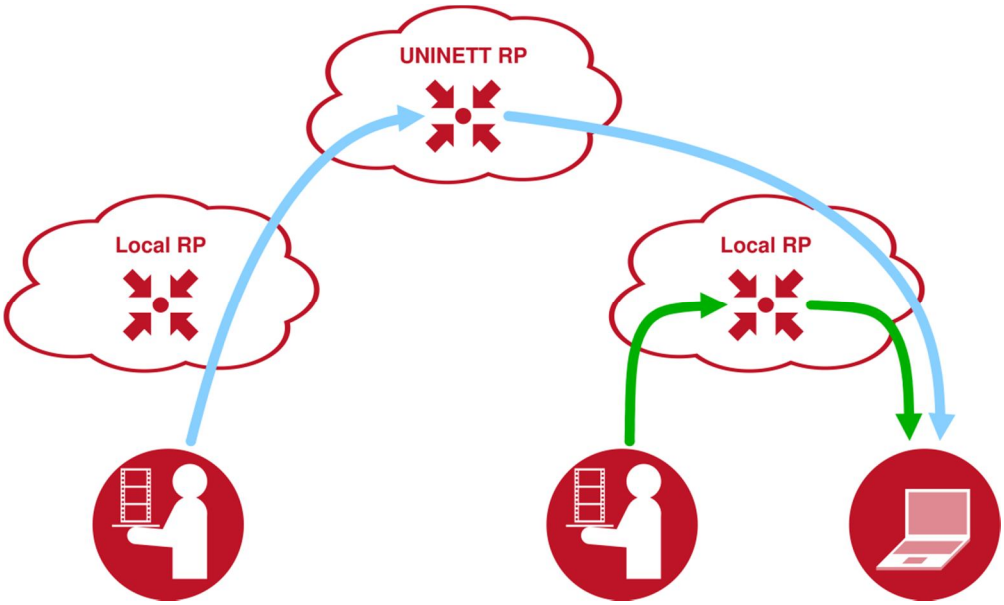

Figure 2: UNINETT's multicast setup with sector-wide and local rendezvous points

1.3.3 Inter-domain routing

Different multicast domains will have their own RP's. Hosts that join specific sources may still be able to receive multicast from sources in other domains. However if a host in one domain joins a group, and a host in another sends to the same group, the data will not reach the host. Data arrives at one RP, and the shared tree is built from another, waiting for data that never arrive. This is where MSDP comes in, see [RFC3618].

With MSDP peerings between pairs of RPs are set up. The peerings are TCP sessions. When an RP learns of a new source from PIM-SM, it will announce it to its MSDP peers. Also, a router receiving a source announcement from one peer, will forward it to its other peers. In this way, the source announcements can be flooded throughout a network of peers. When an RP receives a source announcement for a group with local interest, which means that someone has previously sent a (*,G)-join to this RP, it will send (S,G)-join towards the source S, building a SPT from the source in the other domain. Data received on the SPT can then be forwarded as usual.

1.4 Multicast at the Ethernet layer (layer 2)

Ethernet has built-in support for multicast. Similarly to IP, Ethernet has unicast MAC addresses for reaching a single host, and multicast MAC addresses for a group of hosts. There is a standardized way of mapping IPv4 and IPv6 multicast addresses to multicast MAC addresses. So when a host joins or sends to an IP multicast group, it will listen for or send Ethernet packets with the corresponding MAC destination address. However, the mapping is not one-to-one, which means that two hosts on a link might join two different IP multicast groups, but listen for Ethernet packets with the same MAC address and receive packets for the wrong group. Hence a host should still check the IP destination address.

Multicast on Ethernet seems pretty simple; but the use of layer 2 switches introduces new challenges. When using switches one would like to limit multicast traffic, so that it does not reach all hosts on the vlan.

As explained in section 1.2, IGMP is used between hosts and routers on the subnet (vlan), so that a router can know which multicast groups the hosts on the subnet want to be a member of. The switches deployed in the campus network should support IGMP snooping, and to support SSM, IGMPv3 snooping must be supported. With IGMP snooping enabled, only sources that has subscribed to a multicast group will receive the multicast stream.

Figure 3 below illustrated where IGMP snooping fits into the overall multicast picture.

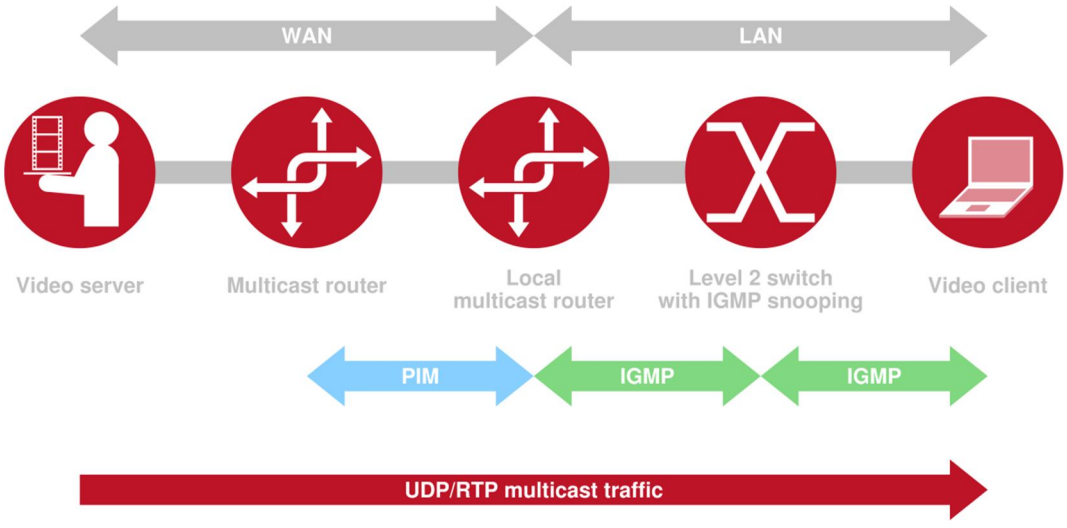

Figure 3: Protocols, equipment and mechanisms involved along the path of a multicast stream

2 Multicast configuration on campus

2.1 Layer 3 considerations

On your routers at campus you need to enable multicast routing and specify which interfaces should participate in the PIM sparse-mode setup that UNINETT provides. You should also configure support for IGMPv3 at the interface level. As long as your layer 2 switches (see section 2.2) have support for IGMP snooping, there is no problem enabling multicast on all subnets, wireless being a potential exception. Multicast over wireless is challenging, but possible. It requires good radio planning and a strictly enforced policy on your wireless controllers that limits multicast traffic to a manageable level.

Since UNINETT is running BSR (see section 1.3.2) and thus provides the multicast group-to-RP mapping there is no need for a BSR-setup or a static RP-mapping at the campus level. In fact this should *not* be set up – as this may or will introduce conflicts.

UNINETT provides a central rendezvous point. In addition you should have a local rendezvous point for local groups at your campus. For university colleges that cover multiple sites, UNINETT has set up a local rendezvous point that is able to support local multicast group (i.e. for software distribution) that spans multiple campuses. If your local multicast scope *does not* include multiple UNINETT-connected campuses, UNINETT requires that the local rendezvous point is announced at the campus level.

To announce the local RP locally, simply set up an extra loopback interface at a central campus router and on this interface announce the local rendezvous point host route (as a /32 prefix). Make sure the prefix is distributed with your local IGP.

2.2 Layer 2 considerations

At layer 2, make sure your switches support IGMP snooping for IPv4 (MLD snooping for IPv6). To support SSM snooping, the switches must support IGMPv3 (MLDv2 for IPv6). Since there is little deployment of SSM today, you should for the time being be fine if only IGMPv2 snooping is supported, but this is likely to change. We definitely recommend that procurement of new switches require IGMPv3 support.

Since IGMPv3 messages are different from the messages used in IGMPv2, using SSM with switches that only support IGMPv2 snooping will not work. If a switch does not recognize IGMPv3 messages, hosts will not correctly receive traffic if IGMPv3 is being used. Disabling IGMP snooping would solve the problem, but then

multicast traffic will be flooded to all hosts on the vlan. Alternatively you could implement a transition technique on the router (see next paragraph). The best solution is definitely to deploy switches with IGMPv3 support.

In SSM deployment cases you may also run into challenges if a non-IGMPv3 capable host joins the group. The host will then send an IGMPv2 join, and as [RFC4604] recommends, the default behaviour for the involved router is to transform the multicast group into ASM mode. Today all major operating systems support IGMPv3, so this is gradually becoming less of problem, still one should consider implementing an available transition solution. We recommend SSM mapping that is supported by both Cisco and Juniper, see section 2.6 for details. Alternative techniques available on Cisco are URL Rendezvous Directory (URD) and IGMP Version 3 lite (IGMP v3lite), see [v3lite]. HP has a similar solution called 'host join only IGMP snooping'.

2.3 Layer 3 configuration examples

2.3.1 Cisco

2.3.1.1 *Global configuration for Cisco IOS*

To enable IPv4 multicast on Cisco IOS router simply configure:

```
ip multicast-routing
```

'BSR listen' is enabled per default. No BSR configuration should be done.

2.3.1.2 *Global configuration for Cisco NX-OS*

```
feature pim
ip pim bsr listen forward   ! Required for receiving and forwarding BSR info
```

2.3.1.3 *Configuration at the interface level*

IPv4 multicast needs to be enabled per interface. The configuration is the same for IOS and NX-OS; configure PIM sparse-mode and IGMP version 3:

```
ip pim sparse-mode
ip igmp version 3
```

If you have applied access lists on the interface, make sure PIM, IGMP and traffic destined for multicast addresses are permitted:

```
permit pim any any
permit igmp any any
permit ip any 224.0.0.0 15.255.255.255
```

2.3.2 Juniper JUNOS

2.3.2.1 *Global configuration*

On a Juniper router multicast routing is enabled per default.

Please be aware that the initial multicast packets are tunnelled to the rendezvous point. On some Juniper platforms this requires dedicated hardware. On the EX platform tunnelling is done in the RE-CPU, while on the MX-platform the tunnelling can be enables on the TRIO-chip. On M- and T-series routers you need a dedicated tunnel-services PIC.

2.3.2.2 *Configuration at the interface level*

In the example below we enable PIM sparse-mode on all interfaces, except the management interface (fxp0.0):

```
protocols  {
    pim  {
        interface all  {
            mode sparse;
        }
        interface fxp0.0  {
            disable;
        }
    }
}
```

2.3.2.3 *BSR setup*

To prevent unwanted Cisco Auto-RP announcements to interfere with BSR, filter out the Auto-RP multicast groups.

```
protocols {
    pim {
        import pim-join-filter;
    }
}

policy-options {
    policy-statement pim-join-filter {
        term bad-groups {
            from {
                route-filter 224.0.1.40/32 exact;
                route-filter 224.0.1.39/32 exact;
            }
            then reject;
        }
        term last {
            then accept;
        }
    }
}
```

Then use BSR to map rendezvous point to local and global groups:

```
protocols {
  pim {
    rp {
      family inet {
        bootstrap {
          import bsr-import;
          export bsr-export;
        }
      }
    }
  }
}
policy-options {
  policy-statement bsr-import {
    term all {
      then accept;
    }
  }
  policy-statement bsr-export {
    term all {
      then accept;
    }
  }
}
```

2.3.3 HP

To support IPv4 multicast on an HP router, configure the following.

2.3.3.1 *Global configuration*

```
ip multicast-routing

router pim
   enable
exit
```

Note that ‘BSR listen’ is enabled per default.

2.3.3.2 *Configuration at the interface level*

```
 vlan x
    ip pim-sparse
    ip-addr any
    exit
 exit
```

2.4 Layer 2 configuration examples

2.4.1 Cisco IOS, NX-OS and CatOS

Recent Cisco software supports IGMPv3 snooping per default. No configuration is needed. Be aware that for some switches the IGMPv3 snooping only keeps track of which hosts are listening to which multicast groups. The switch does not keep track of the SSM sources. Newer Cisco switch models will not have this limitation.

2.4.2 Juniper JUNOS

Juniper has IGMPv2 snooping enabled per default per vlan. You should enable IGMPv3 snooping like this:

```
protocols {
    igmp-snooping {
        vlan all {
            version 3;
        }
    }
}
```

2.4.3 HP ProCurve

IGMP snooping is by default turned off and should be enabled. This is configured per vlan with the following command:

```
# vlan x ip igmp
```

HP supports IGMPv3, but does not keep track of the SSM sources.

2.5 Multicast in a redundant campus network

Since PIM uses the unicast routing protocol, there is no problem in enabling PIM in a redundant network design. When IP Multicast traffic is pulled through the network the paths are determined by the Designated Router (DR) that sends the PIM joins from the edges of the network. In VRRP/HSRP scenarios, where hosts have multiple routers on the local subnet, care should be taken. The DR should always be the active VRRP/HSRP router. On Cisco routers one can use the ‘ip pim dr-priority’ command to dictate which router is selected as DR. The default value is 1 and the router with highest dr-priority wins. There may however be scenarios where the active VRRP/HSRP router changes and multicast traffic malfunctions. Cisco has therefore implemented “HSRP aware PIM” on some router platforms (and has plans for “VRRP aware PIM”). Unfortunately, this is not yet available on the Catalyst 6500 that is deployed on many campus networks in Norway.

2.6 SSM mapping

As mentioned in section 2.2 there are challenges when IGMPv2 hosts joins an SSM group. This will by default transform the group into ASM mode. To avoid this, SSM mapping can be implemented. With SSM mapping a IGMPv2 join for group G will be transformed into an (S,G) join on the router. The implementation can either be static or DNS-based, where the latter is the most flexible solution. A simple static solution is given below; when an IGMPv2 host tries to join the group 232.1.2.1 the router will transform this into an SSM join for (192.168.2.3, 232.1.2.1).

The necessary Cisco IOS configuration is:

```
access-list 5 permit 232.1.2.1

ip igmp ssm-map enable
no ip igmp ssm-map query dns        ! enables static mapping instead of DNS-based
ip igmp ssm-map static 5 192.168.2.3
```

On Juniper the equivalent configuration is:

```
policy-options {
    policy-statement ssm-policy-example {
        term A {
            from {
                route-filter 232.1.2.1/32 exact;
            }
            then {
                ssm-source 192.168.2.3;
                accept;
            }
        }
    }
}
```

Apply the SSM-mapping on an interface:

```
protocols {
    igmp {
        interface fe-0/1/0.0 {
            ssm-map-policy ssm-policy-example;
            }
    }
}
```

2.7 Security considerations

2.7.1 ASM security measures

As already pointed out, the main weakness with ASM is that an illegitimate source may interfere on a multicast group, i.e. an ASM based TV broadcast. This can be prevented at layer 3 by blocking end users from sending data to the multicast groups one wishes to protect. This does however have an administrative overhead, and it will not give protection within a subnet on campus.

2.7.2 Securing PIM

We recommend that PIM is limited to the campus core network. Unwanted PIM neighbours can be filtered out at the interface level. For subnets with only one router this can be easily done on Cisco IOS with:

```
(config-if)# ip pim passive
```

On Juniper the equivalent configuration is:

```
protocols {
    pim {
        interface ge-0/0/0.0 {
            hello-interval 0;
        }
    }
}
```

If there are more than one router on the subnet, i.e. in cases where VRRP is implemented, the configuration is a bit more elaborate. For Cisco IOS configure:

```
(config-if)# ip pim neighbor-filter 5

(config)# access-list 5 permit host 192.168.1.3
(config)# access-list 5 deny   any
```

In the example only the IP address 192.168.1.3 is allowed as PIM neighbour.

Similar configuration on Juniper is:

```
policy-options {
    prefix-list nbrGroup {
        192.168.1.3/32;
    }
    policy-statement nbr-policy {
        term permitted-sources-groups {
            from {
                prefix-list nbrGroup;
            }
            then accept;
        }
        then reject;term deny-everything-else {
            then reject;
```

```
        }
    }
}
protocols {
    pim {
        rp {
        dr-register-policy dr-filter;
        }
    }
}
```

2.7.3 BSR security measures

The most secure way to implement BSR is to set BSR border towards the campus networks and then implement MBGP and MSDP. This has not been necessary in Norway as good trust relationships are established between UNINETT and its customers. When in addition PIM is blocked in the periphery of the campus network, as explained in section 2.7.2, the chance of introducing unwanted BSR announcements from illegitimate sources is reduced to a manageable level.

3 Troubleshooting IPv4 multicast

There are a variety of tools available for troubleshooting multicast. We divide them in two groups; end user tools and network administrator tools.

3.1 End user tools

UNINETT offers a simple tool to verify if your computer is connected to a multicast enabled subnet.
Simply go to http://forskningsnett.uninett.no/tv/mtester/ with a Java-enabled browser and do the test. A screenshot (in Norwegian) is included in figure 4 below.

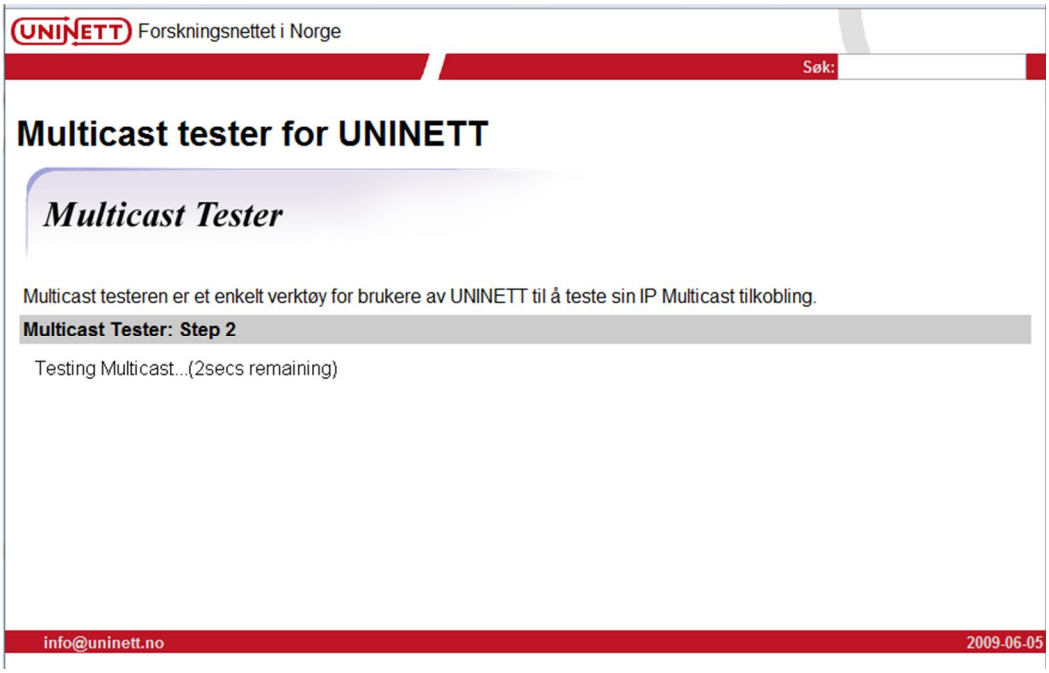

Figure 4: Web based multicast tester offered by UNINETT

A slightly more advances Java tool, called Mikkle's MulticastTest, can be downloaded from http://www.mikkle.dk/multicasttest/index.html . The tool lets you listen to and/or send data on specific multicast groups.

3.2 Network administrator web tools

The tools in this section may be useful for end users as well, although they do require a deeper insight and understanding of networking.

UNINETT has deployed over thirty special purpose monitoring servers (so-called measurement beacons) at the border between the UNINETT backbone network and various campus networks around the country. The measurement beacons are running a suite of monitoring applications. Relevant for multicast are the multicast beacon and the ssmping looking glass.

3.2.1 Multicast beacon

The multicast beacon is a great open source tool [mbeacon] to verify general multicast reachability, both for ASM and SSM, in UNINETT. Each measurement beacon sends regularly data to an ASM multicast group and an SSM multicast group. At the same time they listen to the same groups. Based on data they receive, a web based matrix gives an overview of multicast reachability. As an added feature the matrix also presents TTL counts, traffic loss, delay and jitter.

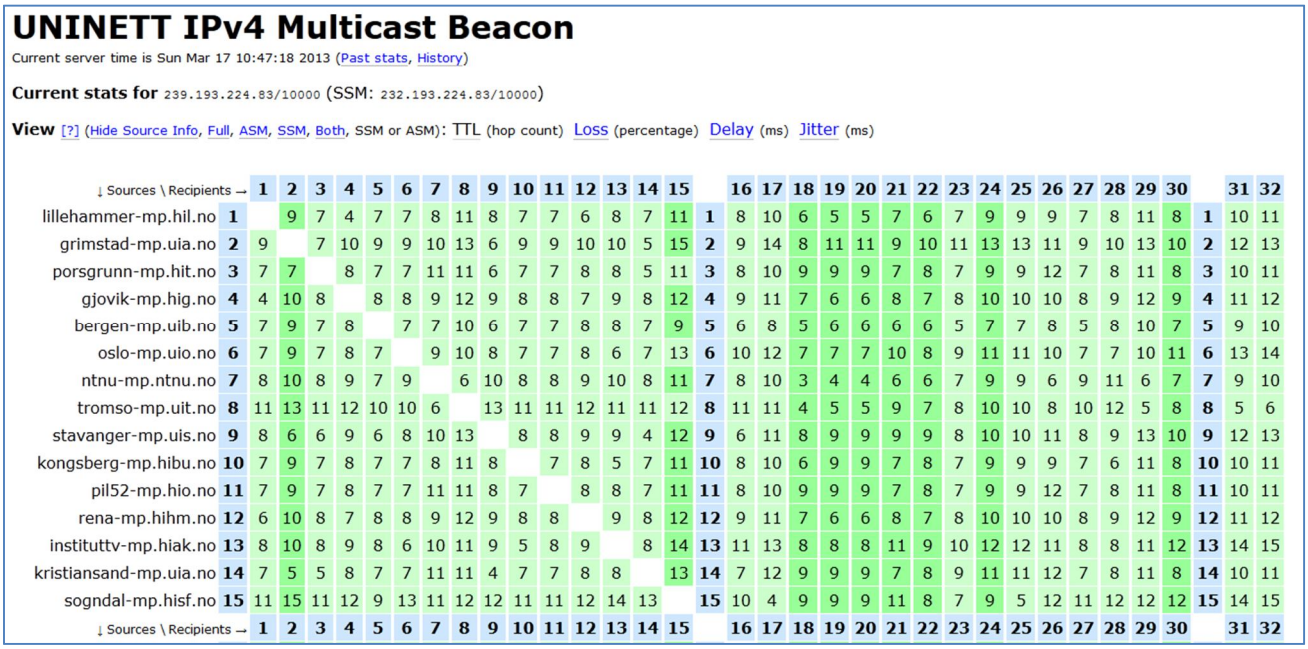

UNINETT IPv4 Multicast Beacon

Current server time is Sun Mar 17 10:47:18 2013 (Past stats, History)

Current stats for 239.193.224.83/10000 (SSM: 232.193.224.83/10000)

View [?] (Hide Source Info, Full, ASM, SSM, Both, SSM or ASM): TTL (hop count) Loss (percentage) Delay (ms) Jitter (ms)

↓ Sources \ Recipients →	1	2	3	4	5	6	7	8	9	10	11	12	13	14	15		16	17	18	19	20	21	22	23	24	25	26	27	28	29	30		31	32
lillehammer-mp.hil.no	1		9	7	4	7	7	8	11	8	7	7	6	8	7	11	1	8	10	6	5	5	7	6	7	9	9	9	7	8	11	8	1	10	11
grimstad-mp.uia.no	2	9		7	10	9	9	10	13	6	9	9	10	10	5	15	2	9	14	8	11	11	9	10	11	13	13	11	9	10	13	10	2	12	13
porsgrunn-mp.hit.no	3	7	7		8	7	7	11	11	6	7	7	8	8	5	11	3	8	10	9	9	9	7	8	7	9	9	12	7	8	11	8	3	10	11
gjovik-mp.hig.no	4	4	10	8		8	8	9	12	9	8	8	7	9	8	12	4	9	11	7	6	6	8	7	8	10	10	10	8	9	12	9	4	11	12
bergen-mp.uib.no	5	7	9	7	8		7	7	10	6	7	7	8	8	7	9	5	6	8	5	6	6	6	6	5	7	7	8	5	8	10	7	5	9	10
oslo-mp.uio.no	6	7	9	7	8	7		9	10	8	7	7	8	6	7	13	6	10	12	7	7	7	10	8	9	11	11	10	7	7	10	11	6	13	14
ntnu-mp.ntnu.no	7	8	10	8	9	7	9		6	10	8	8	9	10	8	11	7	8	10	3	4	4	6	6	7	9	9	6	9	11	6	7	7	9	10
tromso-mp.uit.no	8	11	13	11	12	10	10	6		13	11	11	12	11	11	12	8	11	11	4	5	5	9	7	8	10	10	8	10	12	5	8	8	5	6
stavanger-mp.uis.no	9	8	6	6	9	6	8	10	13		8	8	9	9	4	12	9	6	11	8	9	9	9	9	8	10	10	11	8	9	13	10	9	12	13
kongsberg-mp.hibu.no	10	7	9	7	8	7	7	8	11	8		7	8	5	7	11	10	8	10	6	9	9	7	8	7	9	9	9	7	6	11	8	10	10	11
pil52-mp.hio.no	11	7	9	7	8	7	7	11	11	8	7		8	8	7	11	11	8	10	9	9	9	7	8	7	9	9	12	7	8	11	8	11	10	11
rena-mp.hihm.no	12	6	10	8	7	8	8	9	12	9	8	8		9	8	12	12	9	11	7	6	6	8	7	8	10	10	10	8	9	12	9	12	11	12
instituttv-mp.hiak.no	13	8	10	8	9	8	6	10	11	9	5	8	9		8	14	13	11	13	8	8	8	11	9	10	12	12	11	8	8	11	12	13	14	15
kristiansand-mp.uia.no	14	7	5	5	8	7	7	11	11	4	7	7	8	8		13	14	7	12	9	9	9	7	8	9	11	11	12	7	8	11	8	14	10	11
sogndal-mp.hisf.no	15	11	15	11	12	9	13	11	12	12	11	11	12	14	13		15	10	4	9	9	9	11	8	7	9	5	12	11	12	12	12	15	14	15
↓ Sources \ Recipients →	1	2	3	4	5	6	7	8	9	10	11	12	13	14	15		16	17	18	19	20	21	22	23	24	25	26	27	28	29	30		31	32

Figure 5: UNINETT's multicast beacon

The multicast beacon matrixes are publicly available.
Visit one placed in Trondheim at http://mi6.uninett.no/matrix/

3.2.2 ssmping looking glass

The measurements beacons are also running an ssmping looking glass [ssmping]. Use the looking glass to verify either ASM or SSM reachability from the looking glass server to an IPv4 or IPv6 address of interest. A screenshot is provided below. Test yourself at http://mi6.uninett.no/ssmping/.

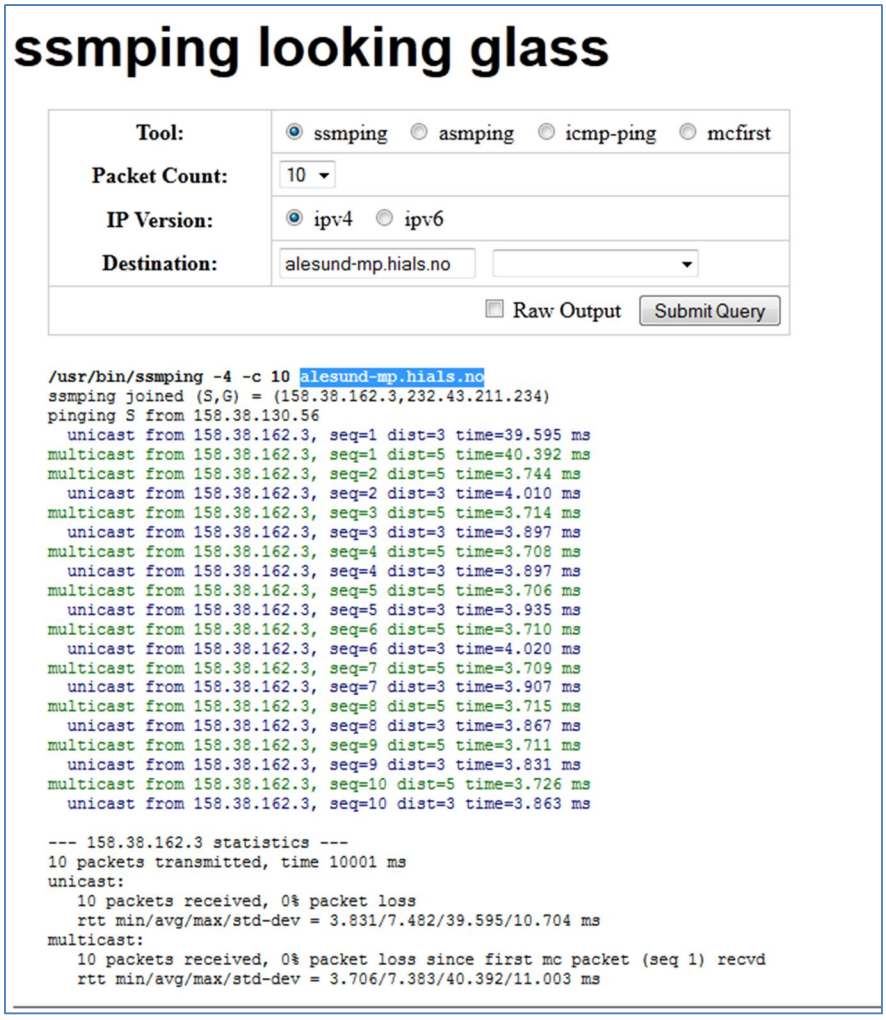

Figure 6: ssmping looking glass

3.3 CLI-based troubleshooting

3.3.1 Cisco IOS

Cisco provides a good multicast troubleshooting guide at
http://www.cisco.com/en/US/tech/tk828/technologies_tech_note09186a0080093f21.shtml

We provide an overview of the most useful commands.

To see your PIM neighbour routers, upstream and downstream:

```
show ip pim neighbor
```

To verify that you are receiving the RP-mapping as you should from UNINETT:

```
show ip pim rp mapping
```

In order to troubleshoot multicast routing an *active* sender and listener to a given multicast groups is needed. As an example we use the ASM-group 233.121.29.1 where active traffic is seen:

```
ios-gsw#show ip mroute 233.121.29.1
IP Multicast Routing Table
Flags: D - Dense, S - Sparse, B - Bidir Group, s - SSM Group, C - Connected,
     L - Local, P - Pruned, R - RP-bit set, F - Register flag,
     T - SPT-bit set, J - Join SPT, M - MSDP created entry, E - Extranet,
     X - Proxy Join Timer Running, A - Candidate for MSDP Advertisement,
     U - URD, I - Received Source Specific Host Report,
     Z - Multicast Tunnel, z - MDT-data group sender,
     Y - Joined MDT-data group, y - Sending to MDT-data group
     V - RD & Vector, v - Vector
Outgoing interface flags: H - Hardware switched, A - Assert winner
 Timers: Uptime/Expires
 Interface state: Interface, Next-Hop or VCD, State/Mode
(*, 233.121.29.1), 00:00:15/stopped, RP 128.39.0.85, flags: SJC
  Incoming interface: TenGigabitEthernet8/5, RPF nbr 128.39.47.49, Mbgp,
Partial-SC
  Outgoing interface list:
    Vlan20, Forward/Sparse, 00:00:15/00:02:49, H
(82.117.34.67, 233.121.29.1), 00:00:15/00:02:54, flags: JT
  Incoming interface: TenGigabitEthernet8/5, RPF nbr 128.39.47.49, Mbgp, RPF-MFD
  Outgoing interface list:
    Vlan20, Forward/Sparse, 00:00:15/00:02:49, H
```

The output verifies that the router has an operative RP tree towards the rendezvous point and that a multicast stream from the source is forwarded out on interface Vlan20 (the interface being in sparse-mode).

List active listeners (reporters) on a given subnet base on IGMP data. Also see which switch ports (interfaces) the reporters are connected to on layer 2:

```
uninett-gw#show ip igmp snooping explicit-tracking vlan 20

Source/Group                    Interface       Reporter        Filter_mode
-------------------------------------------------------------------------------
0.0.0.0/239.192.83.80           Vl20:Gi4/21     158.38.62.211   EXCLUDE
0.0.0.0/239.255.255.250         Vl20:Gi4/21     158.38.62.211   EXCLUDE
0.0.0.0/230.0.0.3               Vl20:Gi4/21     158.38.62.209   EXCLUDE
0.0.0.0/239.192.83.80           Vl20:Gi4/21     158.38.62.51    EXCLUDE
0.0.0.0/239.255.255.250         Vl20:Gi4/21     158.38.62.51    EXCLUDE
0.0.0.0/233.121.29.1            Vl20:Te8/4      158.38.62.66    EXCLUDE
```

3.3.2 Cisco NX-OS

To see your PIM neighbour routers, upstream and downstream:

```
show ip pim neighbor
```

To verify that you are receiving the RP-mapping as you should from UNINETT:

```
show ip pim rp mapping
```

Show active multicast traffic for the ASM-group 233.121.29.1:

```
ios_nx-gsw# sh ipow mroute 233.121.29.1
IP Multicast Routing Table for VRF "default"

(*, 233.121.29.1/32), uptime: 00:00:35, pim ip igmp
  Incoming interface: Vlan3, RPF nbr: 128.39.46.153
  Outgoing interface list: (count: 1)
    Vlan10, uptime: 00:00:02, igmp

(82.117.34.67/32, 233.121.29.1/32), uptime: 00:00:34, ip mrib pim
  Incoming interface: Vlan3, RPF nbr: 128.39.46.153
  Outgoing interface list: (count: 1)
    Vlan10, uptime: 00:00:02, mrib
```

List active listeners (reporters) on a given subnet base on IGMP data. Also see which switch ports (interfaces) the reporters are connected to on layer 2:

```
ios_nx-gsw# show ip igmp snooping explicit-tracking
IGMPv3 Snooping Explicit-tracking information
Vlan Source/Group
10   */239.193.224.83
     Eth1/2          158.37.153.34 6d05h     00:01:29  00:02:50
10   */239.255.255.250
     Eth1/2          158.37.153.26 00:15:29  00:01:37  00:02:42
```

3.3.3 Juniper JUNOS

Juniper provides a good guide for troubleshooting multicast at:
http://kb.juniper.net/InfoCenter/index?page=content&id=KB21149&cat=SSG520M&actp=LIST

Appendix A Multicast Assignment Practice

Address assignments for IPv4 multicast are detailed in [RFC 5771]. An overview of the most important scopes is given here.

Address Range	Size	Designation
224.0.0.0 - 224.0.0.255	(/24)	Local Network Control Block
224.0.0.2	(/32)	All Routers on this Subnet
224.0.1.0 - 224.0.1.255	(/24)	Internetwork Control Block
224.0.1.1	(/32)	NTP, Network Time Protocol
224.0.2.0 - 224.0.255.255	(65024)	AD-HOC Block I
224.1.0.0 - 224.1.255.255	(/16)	RESERVED
224.2.0.0 - 224.2.255.255	(/16)	SDP/SAP Block
224.3.0.0 - 224.4.255.255	(2 /16s)	AD-HOC Block II
224.5.0.0 - 224.255.255.255	(251 /16s)	RESERVED
225.0.0.0 - 231.255.255.255	(7 /8s)	RESERVED
232.0.0.0 - 232.255.255.255	(/8)	Source-Specific Multicast (SSM) Block
233.0.0.0 - 233.251.255.255	(16515072)	GLOP Block
233.252.0.0 - 233.255.255.255	(/14)	AD-HOC Block III
234.0.0.0 - 238.255.255.255	(5 /8s)	RESERVED
239.0.0.0 - 239.255.255.255	(/8)	Administratively Scoped Block
239.192.0.0 - 239.192.255.255	(/16)	UNINETT scope
239.193.0.0 - 239.194.255.255	(/16)	NORDUnet scope
239.194.0.0 - 239.194.255.255	(/16)	GEANT scope
239.255.0.0 – 239.255.255.255	(/16)	Organization local scope

To set up a local multicast group at campus, please use an address from the 239.255.0.0/16 scope.

References

[CBPD 119] Support for the operation of IPv6 multicast and anycast
http://www.terena.org/activities/campus-bp/pdf/gn3-na3-t4-cbpd119.pdf

[mbeacon] The open source tool Multicast Beacon
http://sourceforge.net/projects/multicastbeacon/

[RFC 1112] Host Extensions for IP Multicasting
http://tools.ietf.org/html/rfc1112

[RFC 3569] An Overview of Source-Specific Multicast (SSM)
http://tools.ietf.org/html/rfc3569

[RFC 3618] Multicast Source Discovery Protocol (MSDP)
http://tools.ietf.org/html/rfc3618

[RFC 4601] Protocol Independent Multicast - Sparse Mode (PIM-SM)
http://tools.ietf.org/html/rfc4601

[RFC 4604] Using Internet Group Management Protocol Version 3 (IGMPv3) and Multicast Listener Discovery Protocol Version 2 (MLDv2) for Source-Specific Multicast
http://tools.ietf.org/html/rfc4604

[RFC 5059] Bootstrap Router (BSR) Mechanism for Protocol Independent Multicast (PIM)
http://tools.ietf.org/html/rfc5059

[RFC 5771] IANA Guidelines for IPv4 Multicast Address Assignments
http://tools.ietf.org/html/rfc5771

[ssmping] The open source tool ssmping
http://www.venaas.no/multicast/ssmping/

[v3lite] Documentation on Cisco's proprietary IGMP v3lite and URD
http://www.cisco.com/en/US/docs/ios/12_2/ip/configuration/guide/1cfssm.html

Glossary

ASM	Any Source Multicast
BGP	Border Gateway Protocol
BPD	Best-Practice Document
BSR	Bootstrap Router
CLI	Command Line Interface
DR	Designated Router
HE	Higher Education
HSRP	Host Standby Router Protocol
IGMP	Internet Group Management Protocol
IGP	Interior Gateway Protocol
IPv4	Internet Protocol version 4
IPv6	Internet Protocol version 6
LAN	Local Area Network
MAC	Media Access Control
MBGP	Multiprotocol BGP (or Multicast BGP”)
MLD	Multicast Listener Discovery
MSDP	Multicast Source discovery Protocol
NREN	National Research and Education Network
PIC	Physical Interface Card
PIM	Protocol Independent Multicast
PIM-DM	PIM Dense Mode
PIM-SM	PIM Sparse Mode
RP	Rendezvous Point
RPT	Rendezvous Point Tree (or shared tree)
RTP	Real-time Transport Protocol
SAP	Session Announcement Protocol
SPT	Shortest Path Tree
SSM	Source Specific Multicast
TCP	Transmission Control Protocol
TTL	Time To Live
UDP	User Datagram Protocol
VRRP	Virtual Router Redundancy Protocol
WAN	Wide Area Network

Complete BPDs available at www.terena.org/campus-bp/

campus-bp-announcements@terena.org