Enterprise IPTV Validated Reference Design Guide

VITEC

Alcatel·Lucent Enterprise

Table of Contents

Introduction 3
About EZ TV 3
Overview 3
Design Guidelines 5
- Understanding the IPTV Solution 5
- Bandwidth Requirements 5
- Link Aggregation 6
- Virtual Chassis 7

Protocol-Independent Multicast (PIM) Variants 7
Multicast Groups 8
Multicast Flows 9
Deployment Guidelines 9
- General Guidelines 9
- UNP (Universal Network Profile) 10
- IP Multicast Switching 11

IP Multicast Routing 12
Conclusion 14

Introduction

Enterprise IPTV solutions bring broadcast-quality IPTV distribution and digital signage to the Enterprise. This Validated Reference Design Guide (VRDG) provides guidelines and best practices to help network engineers design and deploy ALE's wired and wireless networks for best IPTV performance. This document focuses on Vitec's EZ TV Enterprise IPTV platform and the guidelines and configuration snippets are taken from interoperability and performance testing conducted at ALE's Solution Lab.

About EZ TV

VITEC's Enterprise Grade IPTV & Digital Signage Platform combines broadcast-quality IPTV distribution with Digital Signage capabilities into an all-in-one integrated solution. Designed to seamlessly integrate with any IT environment and run on all types of networks: LAN, WAN and Wireless - EZ TV is the ideal solution for enterprise customers, sports stadiums and arenas, medical facilities and government agencies looking for a secure, scalable cost-effective way to distribute video and display dynamic signs using their IP network.

Overview

IPTV and digital signage are found in Enterprises of all sizes. In Figure 1 below, we have used a medium-large deployment, such as a multi-building site, as a reference, but the guidelines in this VRDG are equally valid in smaller or larger deployments.

Figure 1 – Test Bed

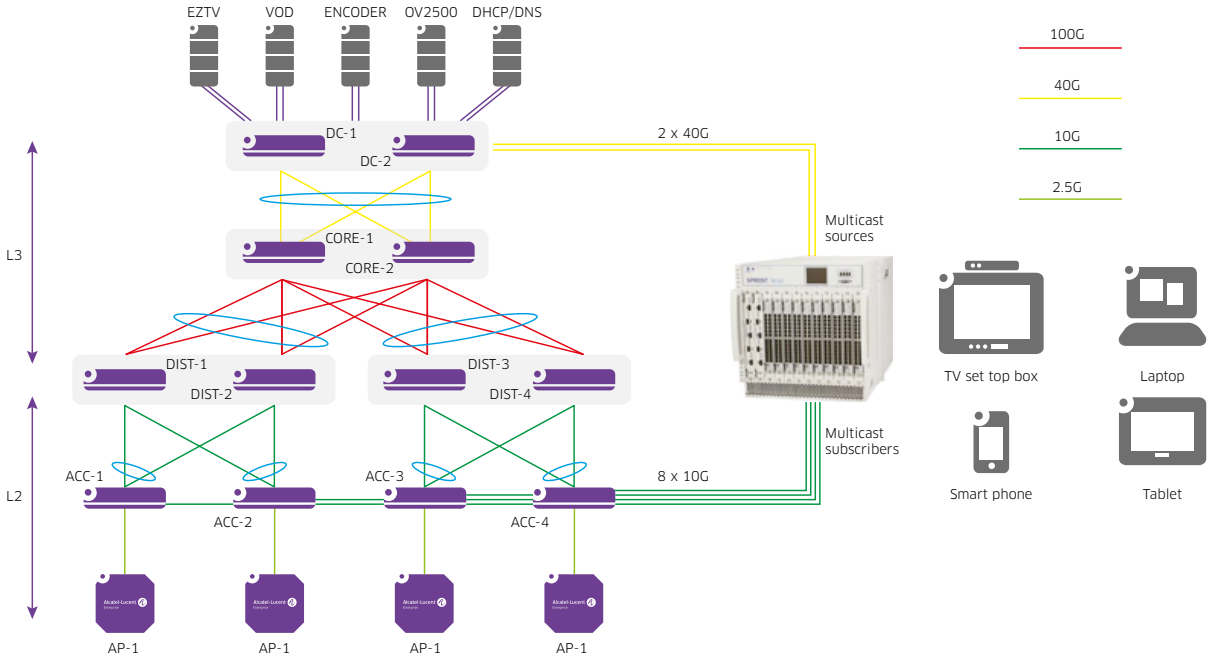

We can distinguish the following blocks:

Table 1 - Functional Blocks

	Description	Reference Family
Server Block	This block is comprised of two OmniSwitch 6900 switches deployed as a Virtual Chassis and provides connectivity to IPTV management, VOD, DHCP/DNS, OmniVista and other servers. All servers are virtualized and hypervisors are dual-attached through LACP aggregates. The Server Block collapses L3 and L2 functions in a single block. In a larger DC deployment however, L3 and L2 functions can be split between spine and leaf nodes. In a smaller deployment, the Core Block function can be collapsed on the same block as well. The Server Block connects to the Core Block through point-to-point OSPF interfaces. PIM is enabled on all L3 interfaces.	OmniSwitch 6860 OmniSwitch 6900
Core Block	This block is comprised of two OmniSwitch 6900 switches deployed as a Virtual Chassis. The Core Block interconnects building Distribution blocks with the Server Block. The Core Block is deployed as pure L3 and connects to Server and Distribution Blocks through point-to-point OSPF interfaces over LACP aggregates. The Core Block is also the PIM Rendez-vous Point because all traffic between the Multicast Sources and Subscribers needs to transit through this block. PIM is enabled on all L3 interfaces.	OmniSwitch 6900 OmniSwitch 9900
Distribution Blocks	One Distribution Block per building. Each Distribution Block is comprised of two OmniSwitch 6900 switches deployed as a Virtual Chassis. Distribution Blocks aggregate connections from Access Switches and provide default gateway routing for all wired and wireless access subnets. Distribution Blocks connect to the Core Block through point-to-point OSPF interfaces over LACP aggregates and to Access Switches through 802.1Q-tagged LACP aggregates. PIM is enabled on all L3 interfaces.	OmniSwitch 6900 OmniSwitch 6860(N) OmniSwitch 6865
Access Switches	OmniSwitch 6560 Access Switches are deployed as pure L2 devices and all wired and wireless VLANs are trunked to the Distribution Blocks where IP default gateway interfaces reside. Access Switches have IGMP snooping enabled to optimize distribution of multicast streams only to interested parties.	OmniSwitch 6860(N) OmniSwitch 6560 OmniSwitch 6450 OmniSwitch 6465
Access Points	OmniAccess Stellar Access Points connect to Access Switches through 802.1Q-tagged ports or aggregates. APs are multicast-aware and convert multicast streams into unicast streams as needed for improved performance.	OmniAccess Stellar APs – all models
Applications	Applications include Active Directory, DHCP, DNS and OmniVista 2500 as well as IPTV-specific applications: the video encoder, the EZ TV management application, the set-top-box management application and the Video on Demand server.	
Client Devices	Wired and wireless client devices are used for functional testing. This includes browser-based clients, mobile apps and hardware set-top-boxes.	EZ TV PC-client EZ TV Android client EZ TV IOS client EZ TV Endpoint
Test Appliance	The test appliance brings realism to the test by emulating thousands of multicast sources and subscribers, stressing not only the data plane, but crucially, the control plane. The test appliance evaluates various performance metrics such as group-join/prune latency, loss and flood/filtering behaviour.	Spirent TestCenter

Design Guidelines

In this section, we will provide design guidelines for all blocks. We will concentrate on multicast-related aspects. A functional unicast-enabled network is assumed as the starting point and we will not cover those aspects in this document (e.g. VLAN, LACP, OSPF…) except when related to multicast.

Note: Throughout this document we refer to multicast "flow" as a multicast stream identified by the combination of multicast group and source IP addresses.

Understanding the IPTV Solution

A good understanding of the IPTV solution and its deployed options is paramount. This includes the number of streams, the stream bitrate, MTU and whether IGMPv3 and SSM (Source-Specific Multicast) are used or not.

In the case of EZ TV:

- 30-500 streams in a typical deployment
- SD-quality stream bitrate ~ 4Mbps
- HD-quality stream bitrate (H.264 encoding)
  - ¬ Typical: ~6Mbps
  - ¬ Critical or fast-motion: ~8-10Mbps
  - ¬ Low-latency (e.g. sports): ~12Mbps
- Mobile-quality stream bitrate ~ 2-3 Mbps
- MTU: 1316 bytes
- IGMP: v2 in most deployments
- SSM: Supported on PC-clients only at the time of this writing
- One multicast group per channel

Bandwidth Requirements

We need to understand a few key IP multicast concepts:

- Multicast routing uses Reverse Path Forwarding (RPF). In a nutshell, what this means is that incoming multicast traffic received on a specific interface is not forwarded unless that interface is also on the path of unicast traffic back to the packet's source address.
- A single copy of any multicast packet is forwarded down any L3 interface and flows are only forwarded for as long as at least one downstream subscriber is interested in receiving the flow.
- A single copy of any multicast packet is forwarded down any untagged L2 interface and flows are only forwarded for as long as at least one downstream subscriber is interested in receiving the flow.
- Tagged L2 links may carry multiple replicas of the same multicast packet if downstream subscribers are on different VLANs because traffic needs to be replicated at the last L3 hop.

Let's analyse the example from Figure 1…

On all DC and L3 links, a simple rule of thumb is that we need enough bandwidth to cater for all streams from sources reachable (RPF) across the link. This is normally a good estimation assuming that all streams have at least one downstream subscriber. It is also a simple requirement to cater for in most Enterprise environments when using links of 10G capacity or higher, as we will see later. But if we were talking about multicast traffic across a WAN link however, a finer analysis would be warranted.

On access layer uplinks, we need to consider the following:

- The number of downstream clients.
- The number of multicast streams that each client subscribes to. A mosaic display may subscribe to as many as 32 individual streams.
- The total number of unique multicast streams subscribed to by all downstream subscribers.
- In case of 802.1q-tagged uplinks, the client allocation across VLANs. This is because per-VLAN replication is performed at the last L3 hop. Having clients spread across different VLANs can result in multiple replicas of the same frame being forwarded over the uplink (one for each VLAN with active subscribers).

Except in specific use cases, the number of unique multicast streams with active subscribers may be difficult to forecast (like predicting TV ratings) and a simple upper bound is used instead.

To put this in context, let's consider this scenario:

- One hundred TV channels simultaneously transmitting in HD and mobile quality at 8Mbps and 3Mbps, respectively.
- On DC or L3 links, this means no more than 100 x (8Mbps + 3Mbps) = 1.1Gbps of multicast traffic on these links.
- On each L2 access layer block, we will assume multicast clients are spread across 2 different VLANs (one for wired clients and another one for wireless clients). This is a reasonable assumption for a L2 access layer block of up to ~250 ports. In a larger access-layer block, an additional wired VLAN may be needed. More on VLAN allocation in Section 5.1.6.
- We will also assume that each wired client subscribes to a single HD-quality stream and each mobile client subscribes to a single mobile-quality stream.
- With the assumptions above, the upper bound for multicast traffic on each access-layer uplink is [8Mbps x 100 streams] x 1 wired_VLAN + [3Mbps x 100 streams] x 1 wireless_VLAN = 1.1Gbps. But in a larger access-layer block with 500 ports we will need an additional wired VLAN for a total of 1.9Gbps.

Note: this analysis is only considering the multicast streams. VOD as well as any other unicast traffic requirements need to be added on top.

Link Aggregation

Except in specific use cases, the general guideline is that Link Aggregation can be used for redundancy and design simplification but not to increase capacity. From a capacity point of view, few higher-speed links are better than many lower-speed ones (e.g. 2 x 40G is better than 8 x 10G).

This is because load balancing of multicast traffic is more difficult to predict and achieve. Load balancing across linkagg member ports depends on hashing of either source and destination address (in brief mode) or source and destination address and port (in extended mode). In many use cases, particularly with few multicast sources and/or multicast groups, there is not enough entropy for traffic to be evenly balanced across linkagg member ports.

Note that, load balancing of multicast, or more generally, non-unicast traffic, is not enabled by default. By default, all non-unicast traffic utilizes the primary port in the linkagg. Please refer to Section 5.1.1 for details on how to enable load balancing of multicast traffic in a linkagg.

Virtual Chassis

When a linkagg's member ports are spread across different virtual slots in a VC, we need to consider that some multicast traffic may need to be forwarded across the VFL and size the VFL capacity accordingly.

By default, load balancing of non-unicast traffic is disabled and all non-unicast traffic uses the primary port in the linkagg. If multicast traffic is received on slot X but the primary port is on slot Y, then this multicast traffic needs to be forwarded across the VFL.

When load balancing of non-unicast traffic is enabled, unlike what happens in the unicast case, no preference is given to local linkagg ports. Non-unicast traffic is load balanced across all member ports whether they are local or not. If multicast traffic is received on slot X and needs to be forwarded down a linkagg with member ports across both slots X and Y, then a portion of that multicast traffic will have to be forwarded across the VFL.

VFL link speeds can range from 10G to 100G depending on product family. Even a 10G VFL provides sufficient bandwidth in most situations but, higher speeds a recommended when supported. Note that direct-attached cables make high-capacity VFLs very cost-effective.

Protocol-Independent Multicast (PIM) Variants

In addition to a unicast routing protocol such as OSPF, PIM is required whenever multicast sources and subscribers are not on the same subnet, which is the case in all except but the smallest deployments (<250 devices).

PIM comes in different variants: Sparse Mode (SM), Dense Mode (DM), Source-Specific Multicast (SSM) and Bidirectional (BIDIR). The first three are valid choices for Enterprise IPTV and in this section, we provide guidelines to help make the right choice. Please refer to the OmniSwitch AOS 8 Advanced Routing Configuration Guide for a detailed description of PIM and its different variants.

PIM - Sparse Mode

PIM-SM is the right choice in most situations. PIM-SM only requires IGMPv2 support on client devices, which is broadly supported. In IGMPv2, subscribers join a group without specifying, or even knowing, the source's IP address. It is for this reason that PIM-SM builds trees rooted on a Rendez-vous Point (RP) node. Traffic between source and destination will initially flow through the RP but will automatically switch to the Shortest Path Tree (SPT) after (once the source IP address is known). PIM-SM does not forward traffic unless subscribers explicitly join the group. PIM-SM is a good fit when the subscriber density for each stream is relatively small (e.g. <10% of client population) but can also cater to environments in which the subscriber density is high.

PIM - Dense Mode

PIM-DM is not recommended. Unlike PIM-SM, PIM-DM forwards traffic over all links until prune messages from nodes without interested downstream subscribers are received. Due to this flood-prune behaviour, PIM-DM can temporarily saturate network links. Note that the flood-prune cycle is repeated periodically and therefore links may saturate periodically too. In addition, PIM-DM places a higher load on network node's control plane because state information is maintained even after pruning. PIM-DM can be a viable choice when the subscriber density for each stream is relatively high (e.g. > 10% of client population). That being said, PIM-SM can also cater to such environments without PIM-DM's downsides.

PIM - Source-Specific Multicast

PIM-SSM and IGMPv3 go hand-in-hand. In IGMPv3, subscribers specify not only the group they want to join but also the specific source(s) they want to receive traffic from. Subscribers learn channels as well as their associated multicast groups and source IP addresses from the IPTV management solution. PIM-SSM is similar to PIM-SM but a RP is not needed because the source IP addresses are known and therefore, the multicast distribution tree is rooted directly at the source's first-hop router. PIM-SSM forwards traffic over the SPT starting from the first packet and, for this reason, is more efficient than PIM-SM and does not require configuration of a RP. In addition, PIM-SSM enhances security compared to PIM-SM and PIM-DM. PIM-SSM is less susceptible to intentional or unintentional disruption of the stream and/or exploitation of vulnerabilities: spoofing the multicast group address is not enough to "hijack" the stream, the source IP address needs to be known and spoofed by the attacker as well. On top of that, since multicast traffic is not forwarded unless it passes the RPF (Reverse Path Forwarding) check, the rogue traffic would need to either originate from the same subnet as the legit traffic or, the attacker would need to tamper with the unicast routing table. In short, hacking or disrupting PIM-SSM streams is harder. For these reasons, PIM-SSM may be the best choice, but only if the application and clients support it.

PIM - Summary

To summarize: If the application supports SSM and the clients support IGMPv3, choose PIM-SSM. If not, choose PIM-SM. Note that some groups can be selectively mapped to use SSM whilst others use SM, both options can coexist.

Multicast Groups

An important but sometimes overlooked factor is the choice of multicast group addresses. Multicast group addresses are class D (224.0.0.0 – 239.255.255.255).

First, any reserved or well-known multicast group addresses should not be used. The list of reserved multicast addresses is maintained by IANA.

Second, mapping of multicast IP address to multicast MAC address is oversubscribed on a 1:32 ratio. What this means is that 32 different multicast group IP addresses map to the same MAC address and, consequently, multicast traffic can be incorrectly forwarded to uninterested or unintended recipients. Without getting down to the nitty-gritty, a simple way of avoiding this overlap is to fix the first octet and the first bit of the second octet.

For instance: the 225.[0-127].x.x block contains more than 8M non-overlapping multicast IP addresses. On the other hand, 225.0.1.1, 225.128.1.1, 226.0.1.1 and 29 more group addresses map to the exact same MAC address.

Multicast Flows

Network devices maintain state information for each multicast flow. The number of multicast flows supported varies with the switch model. This number is quoted in the OmniSwitch AOS Release 8 Specifications Guide, IPMS section. Note: (*,G) refers to PIM-SM or PIM-DM and (S,G) refers to PIM-SSM.

Deployment Guidelines

In this section, we will provide deployment guidelines along with configuration snippets for all blocks.

General Guidelines

STP Edge Ports (aka PortFast)

Ports connecting end-host devices such as set-top-boxes, servers and encoders should be configured as edge ports so that these ports will transition directly to a forwarding state and not trigger an unwanted topology change when a device is connected to the port. An edge port is a port at the edge of a LAN that does not receive STP BPDUs and has only one MAC address learned. Edge ports, however, will operationally revert to a point-to-point or a no-point-to-point connection type if a BPDU is received on the port. Configuring these ports as edge ports allows these devices to connect to the network immediately, thus avoiding timeouts. Note that "auto-edge" is the default port configuration on an OmniSwitch and therefore the recommendation is to not change it.

Load Balancing Non-Unicast Traffic Across Linkaggs

By default, non-unicast traffic is not load-balanced across link aggregate member ports and is forwarded on the primary port only.

Load balancing for non-unicast traffic is not necessarily recommended, this is not a general recommendation. Refer to Section 4.3 to understand load balancing of non-unicast traffic in a Linkagg in general and Section 4.4 to understand load balancing of non-unicast traffic in a Virtual Chassis.

But if load balancing of non-unicast traffic across a Linkagg is desired, then it is recommended that the extended hash control and tunnel-protocol options be enabled also.

Load Balancing Non-Unicast Traffic across Linkaggs

```
➔ hash-control extended
➔ hash-control load-balance non-ucast enable
➔ hash-control load-balance non-ucast tunnel-protocol enable
```

TTL

Make sure that the TTL setting in multicast sources is larger than the maximum number of IP hops between source and subscriber or otherwise traffic will be dropped.

MTU

To avoid fragmentation or dropping of large packets, make sure that the MTU used on the network is large enough compared to the MTU used by the source. The default MTU is 1500 bytes but can be increased up to 9198 bytes. The MTU is configured on a per-VLAN basis.

MTU

```
➔  vlan 100 mtu-ip 9198
```

Storm Control

Storm control is enabled by default on some platforms. Make sure thresholds are commensurate to the amount of expected multicast traffic. Refer to the OmniSwitch AOS Release 8 CLI Reference Guide for details.

Storm Control

```
➔ interfaces {slot chassis/slot| port chassis/slot/port[-port2]} flood-
  limit {bcast | mcast | uucast | all} rate {pps pps_num| mbps mbps_num
  | cap% cap_num | enable | disable | default} [low-threshold low_num]
```

Client VLAN Allocation

There are two different guidelines depending on whether the VLAN will be used for wired or wireless clients.

For wired clients, the guidelines are:

- Use different wired-client VLANs in each access block
- Minimize the number of wired-client VLANs per access block

These two guidelines above will minimize the load on uplinks. Remember that:

- When subscribers of the same multicast stream are spread across different VLANs, then the distribution layer switch needs to create a replica for each VLAN and this places a higher load on the uplinks.
- Role-based access does not necessarily mean different VLANs are needed for different roles. Two different roles can be mapped to the same VLAN but have different security or QoS policies thanks to UNP.

For wireless clients, the guidelines are:

- Use a single wireless-client VLAN (coupled with a large enough subnet) for all wireless clients and enable the same VLAN across all access blocks in each location. Avoid using VLAN pooling.
- Enable Broadcast Filter All / Broadcast Filter ARP on the SSID. This will block all broadcast traffic except for DHCP and ARP. The AP will proxy the ARP response when the target MAC address is known to the AP and only forward ARP when the MAC is not known to the AP.

These two guidelines above will minimize traffic on uplinks whilst ensuring wireless clients can perform L2-roams instead of L3-roams and keeping broadcast traffic to a minimum.

UNP (Universal Network Profile)

A UNP-enabled port or SSID is one that can dynamically adjust its parameters (VLAN, ACLs and QoS policies) to the devices connected, or associated, to it. With UNP, settings no longer need to be statically applied to the switch port or SSID and they automatically adjust as needed. UNPs eliminate the need for manual Moves Adds and Changes whilst at the same time increase security with role-based access. User- or Device-to-NP mappings are based on authentication (802.1x or MAC), classification rules (e.g. MAC OUI) or device “signature” profiles (combination of MAC OUI, DHCP options an HTTP user agent).

When configuring UNPs, make sure policies allow communication with all required systems and ports with sufficient bandwidth. Please refer to EZ TV documentation for details.

As an example, a UNP for an EZ TV endpoint should allow the following:

- DHCP
- ARP
- DNS
- NTP
- SSH and HTTPs (with Management Station)
- AD / LDAP
- HTTP (with EZ TV Portal)
- RTSP (with EZ TV VOD)

IP Multicast Switching

IP Multicast Switching (IPMS) configuration is required only on network devices and on VLANs to which subscribers connect. In our example from Figure 1, this means only the Access and Distribution switches need this configuration.

Enabling Multicast Switching

Multicast switching can be enabled on the system as a whole or on a per-VLAN basis.

Enabling Multicast Switching

```
➔  ip multicast admin-state enable
➔  ip multicast vlan 100 admin-state enable
```

Setting the IGMP Version

The IGMP version can be configured on the system as a whole or on a per-VLAN basis. IGMP version 2 is the default. The IGMP version should be changed to v3 only when all clients support it and when using PIM-SSM.

Setting the IGMP Version

```
➔  ip multicast version 3
➔  ip multicast vlan 100 version 3
```

Configuring an IGMP Querier

The IGMP querier's role is to maintain IGMP membership state up to date by periodically querying subscribers about their memberships. The IGMP querier function is implicitly enabled on PIM-enabled interfaces. In L2 deployments however (no L3 hop between sources and subscribers), no interface will be enabled for PIM and therefore this function needs to be explicitly enabled. Note that in our example from Figure 1 it is not needed because the default-gateway IP interfaces are PIM-enabled as we will see later in Section 5.4.2. The multicast querier can be enabled on the system as a whole or on a per-VLAN basis.

Configuring an IGMP Querier

```
➔  ip multicast querying enable
➔  ip multicast vlan 100 querying enable
```

IP Multicast Routing

IP Multicast Routing (IPMR) configuration is required on all devices that route multicast traffic. In our example from Figure 1, this means DC, Core and Distribution nodes require IPMR configuration.

Loading PIM

PIM needs to be loaded on all nodes that route multicast traffic.

Loading PIM

```
➔  ip load pim
```

Enabling PIM on IP Interfaces

PIM needs to be enabled on all IP interfaces that need to forward multicast traffic or that will be configured as RP. In our example from Figure 1, this means all IP interfaces on all nodes. If using PIM-SM, this also includes Loopback0 interfaces advertised as RP (refer to Section 5.4.6).

Tip: Make sure not to enable PIM on external interfaces such as an Internet uplink.

Enabling PIM on IP Interfaces

```
➔  ip pim interface "VLAN100"
➔  ip pim interface "VLAN101"
➔  ip pim interface "Loopback0"
```

Configuring the PIM Mode

For each PIM mode below, the configuration is required on all switches running PIM. In our example from Figure 1, this means DC, Core and Distribution nodes.

PIM-SM

PIM Sparse mode is enabled on a per-node basis.

Configuring PIM-SM Mode

```
➔  ip pim sparse admin-state enable
```

PIM-DM

PIM Dense mode on the other hand, needs to be enabled not only on the switch as a whole but also for the specific groups that need to be mapped to the PIM-DM mode.

Configuring PIM-DM Mode

```
➔  ip pim dense admin-state enable
➔  ip pim dense group 225.0.0.0/24
```

Tip: It is best practice to be as exact as possible when defining multicast group mappings and only advertise those groups which are needed. This minimizes exposure to rogue multicast traffic.

Tip: Don't forget to also map multicast groups required for digital signage, end-point commands, catering and score-board receivers. Please refer to EZ TV documentation.

PIM-SSM

Lastly, PIM-SSM requires enabling PIM-SM and mapping specific groups to the SSM mode.

Configuring PIM-SSM Mode

```
➔  ip pim sparse admin-state enable
➔  ip pim ssm group 234.0.0.0/24
```

Tip: It is best practice to be as exact as possible when defining multicast group mappings and only advertise those groups which are needed. This minimizes exposure to rogue multicast traffic.

Tip: Don't forget to also map multicast groups required for digital signage, end-point commands, catering and score-board receivers. Please refer to EZ TV documentation.

Static or Dynamic RP

RPs are only required when using PIM-SM and not for PIM-DM or PIM-SSM. RPs can be statically configured or dynamically learned. Specific multicast groups can be statically mapped to specific RPs. However, for nodes to be able to forward multicast traffic, RP configuration needs to be consistent across the network. For this reason, dynamic RP is preferred and we will not cover static RP here.

Choosing a RP

As explained in Section 4.5.1, multicast traffic will initially be forwarded through the RP until switchover to the SPT occurs. For this reason, it is best if the RP is already part of the SPT. When the location of sources and receivers is known, as in our example, picking an RP is easy. In our example from Figure 1, this is the Core block.

C-RP and C-BSR

Dynamic configuration of the RP relies on Candidate-RP (C-RP) and Bootstrap Router (BSR). A C-RP sends periodic advertisements to the BSR proposing itself as the RP for specific groups. The BSR keeps all other PIM-enabled routers up-to-date on multicast group to RP mappings by way of bootstrap messages. BSR nodes are elected. A Candidate-BSR (C-BSR) is a PIM-enabled router which is eligible for the BSR role. The C-BSR with highest priority (lowest number) and/or highest IP address is chosen as the BSR, which provides redundancy in case of BSR failure. When configuring a router as C-RP for a specific group, a priority value can be set. In this manner, other non-RP routers will choose the C-RP with the highest priority (lowest number) and/or highest IP address. This also provides for redundancy and load-balancing because different C-RPs can advertise their candidacies for the same group but with different priorities. Note that the BSR does not need to be in the way of traffic.

In our example from Figure 1, we will configure the Core block as both C-RP and C-BSR.

Configuring C-RP and C-BSR

```
➔  ip pim cbsr 10.1.1.1
➔  ip pim candidate-rp 10.1.1.1 225.0.0.0/24
```

In the snippet above, 10.1.1.1 is the Loopback0 IP address. It is assumed that this address is already created and advertised through the IGP routing protocol. Using Loopback addresses is not mandatory, but it is a best practice.

Note that, in the example above, RP and BSR redundancy is already taken care of by the VC feature. We could however configure two different nodes as C-BSR and two different nodes as C-RP. C-RP-1 and C-RP-2 can advertise their candidacy for the same groups but with different priorities providing resiliency and load balancing.

Tip: It is best practice to be as exact as possible when defining multicast group mappings and only advertise those groups which are needed. This minimizes exposure to rogue multicast traffic.

Tip: Don't forget to also map multicast groups required for digital signage, end-point commands, catering and score-board receivers. Please refer to EZ TV documentation.

Conclusion

ALE's OmniSwitch and OmniAccess Stellar networking product lines implement a broad range of multicast protocols and features. In this VRDG, we have provided guidelines and best practices to help network engineers design and deploy networks to support Enterprise IPTV applications such as Vitec's EZ TV with the best possible performance.

www.al-enterprise.com The Alcatel-Lucent name and logo are trademarks of Nokia used under license by ALE. To view other trademarks used by affiliated companies of ALE Holding, visit: www.al-enterprise.com/en/legal/trademarks-copyright. All other trademarks are the property of their respective owners. The information presented is subject to change without notice. Neither ALE Holding nor any of its affiliates assumes any responsibility for inaccuracies contained herein. © Copyright 2020 ALE International, ALE USA Inc. All rights reserved in all countries. DID 20052001 (May 2020)

Alcatel·Lucent Enterprise