paloalto
TECHDOCS

Advanced Routing Engine Migration Reference

March 2023

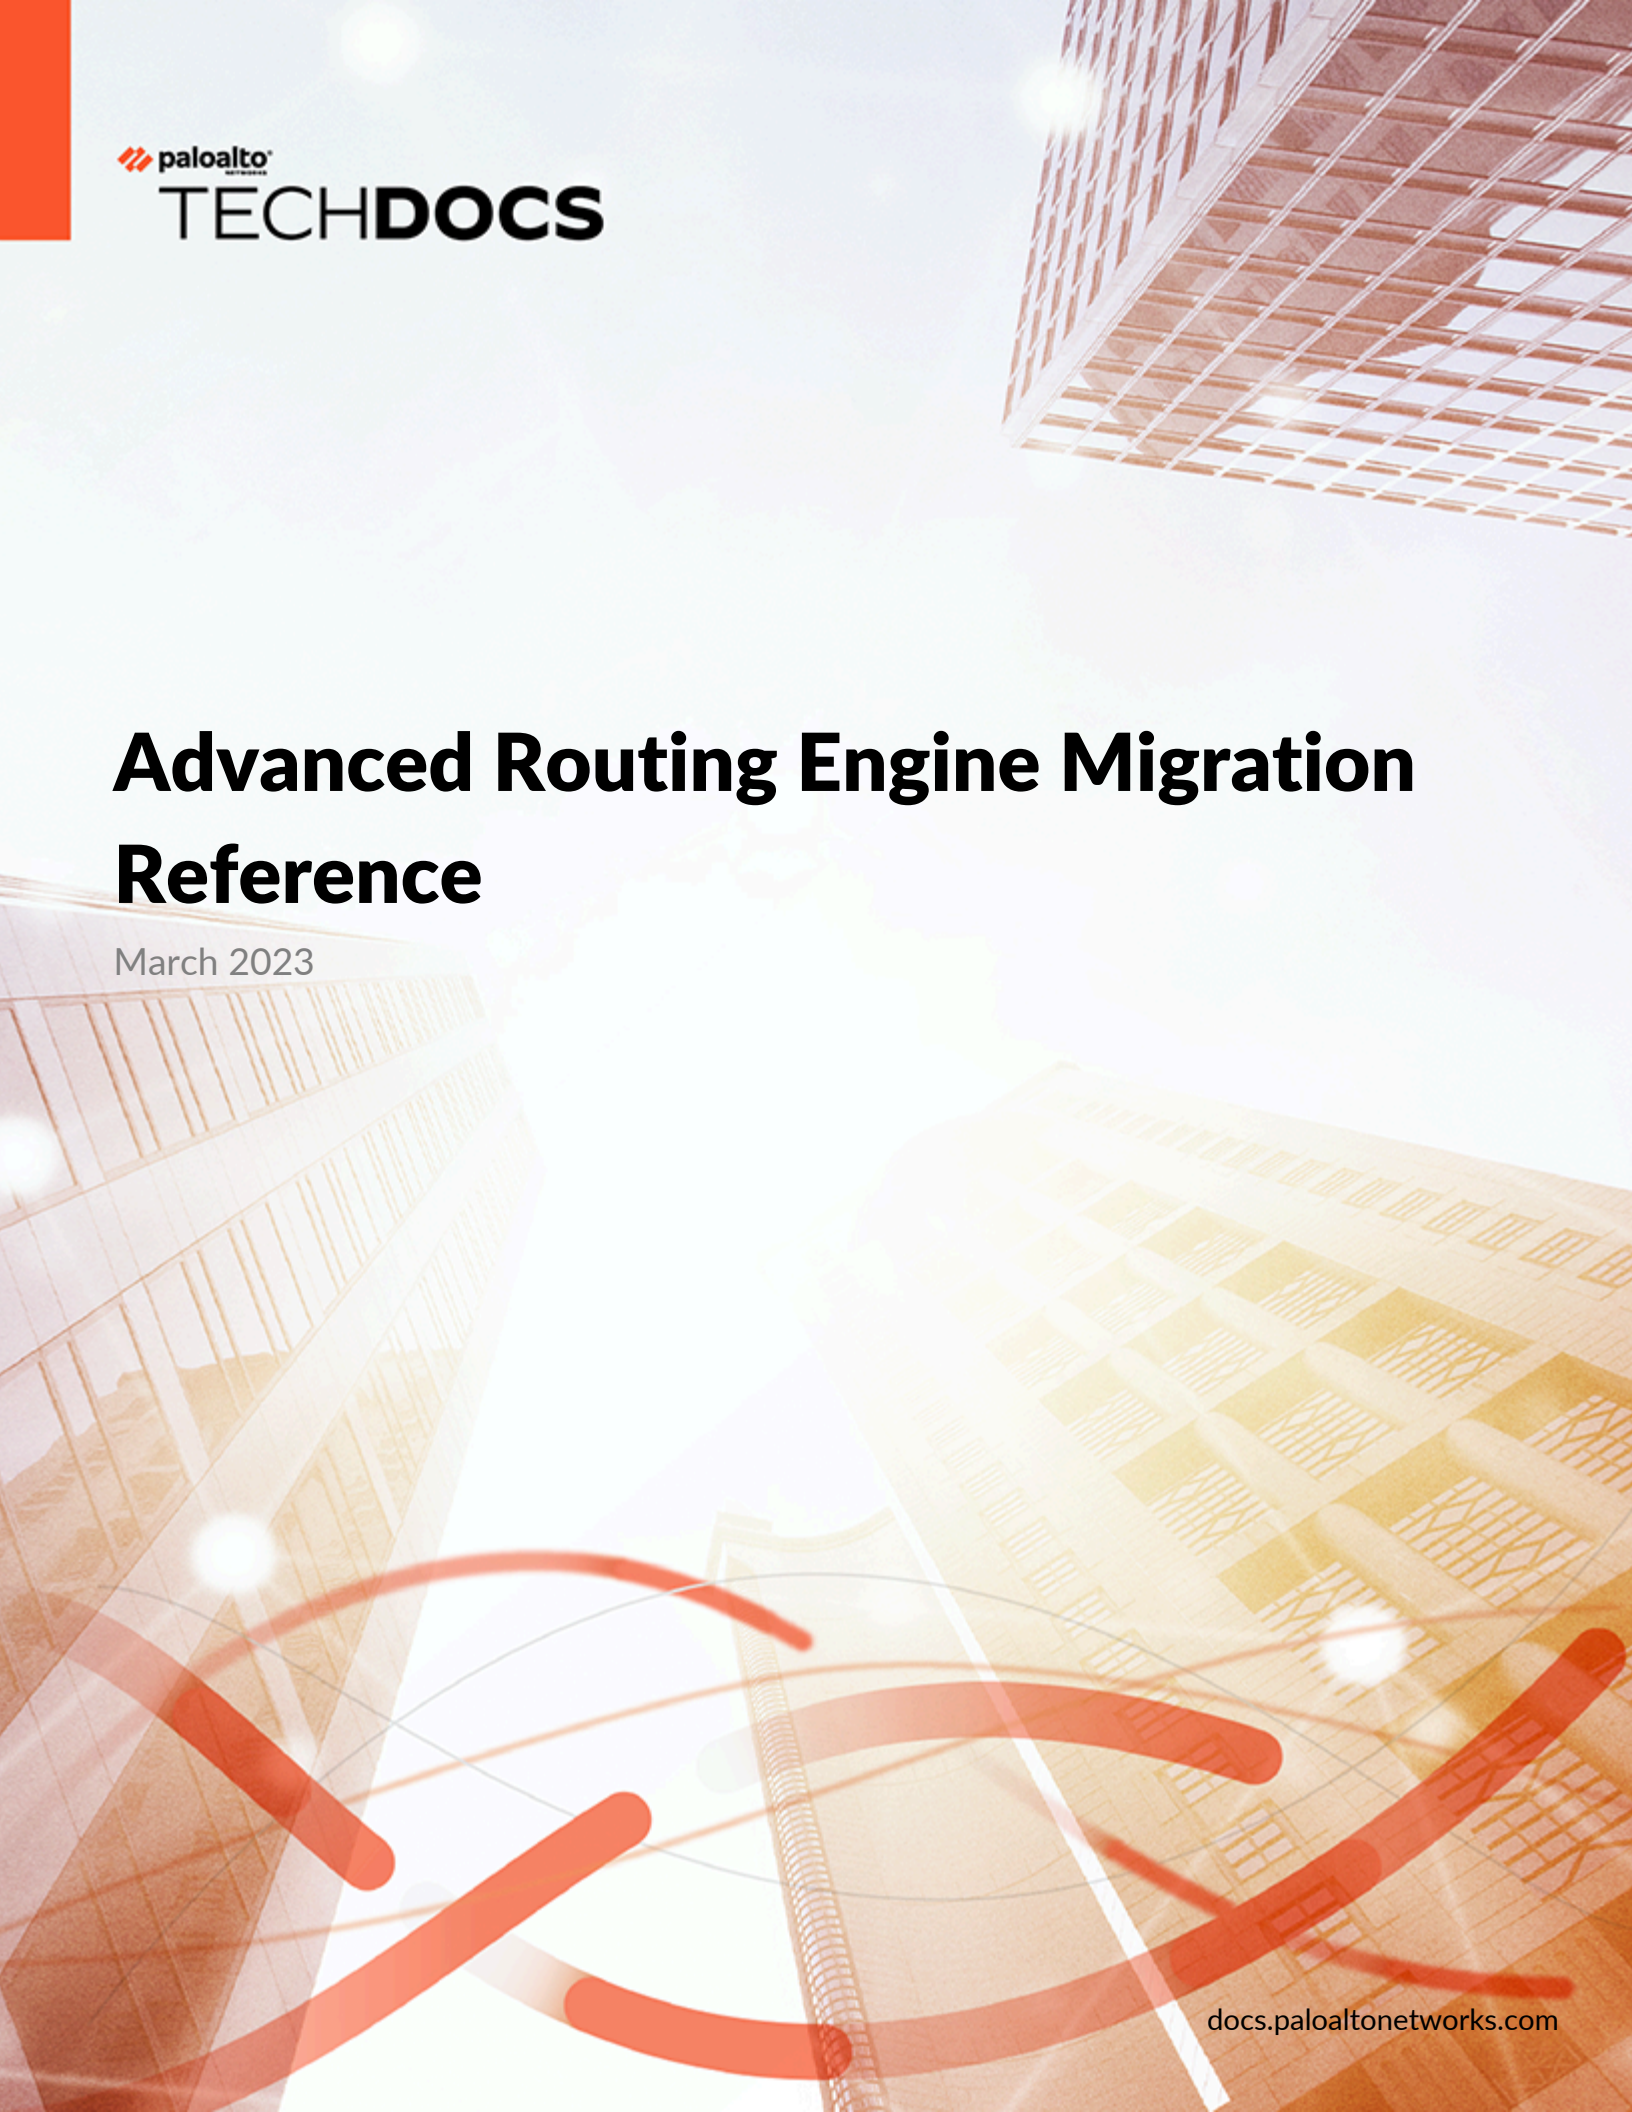

docs.paloaltonetworks.com

Contact Information

Corporate Headquarters:
Palo Alto Networks
3000 Tannery Way
Santa Clara, CA 95054
www.paloaltonetworks.com/company/contact-support

About the Documentation

- For the most recent version of this guide or for access to related documentation, visit the Technical Documentation portal docs.paloaltonetworks.com.
- To search for a specific topic, go to our search page docs.paloaltonetworks.com/search.html.
- Have feedback or questions for us? Leave a comment on any page in the portal, or write to us at documentation@paloaltonetworks.com.

Copyright

Last Revised

May 17, 2023

Table of Contents

Get Started with Routing Engine Migration

The Palo Alto Networks Next-Generation Firewall (NGFW) supports either the legacy routing engine or the advanced routing engine for routing operations; by default, the NGFW comes with the legacy routing engine. The PAN-OS® 10.2 release introduced an advanced routing engine that uses an industry-standard configuration methodology to reduce your learning curve.

The advanced routing engine contains new filtering concepts and relies heavily on routing profiles. As a result, the advanced routing engine configuration is significantly different that the legacy routing engine.

To simplify the advanced routing configuration, we help you migrate from the legacy routing engine to the advanced routing engine with a simple process:

- Invoke the migration script from your existing firewall directly; or
- Invoke the migration script from the Panorama™ management server (if manages the firewall).

Learn about the differences between these two routing engine configurations before you start the migration process.

- Plan Your Routing Engine Migration
- Learn the Differences Between Legacy and Advanced Routing Engine

Plan Your Routing Engine Migration

Beginning with PAN-OS® 10.2, we introduced an advanced routing engine for the Palo Alto Networks family of Next-Generation Firewalls that has more features and is more flexible than the legacy routing engine. You can enable the advanced routing engine and migrate an existing firewall configuration that uses the legacy routing engine to the advanced routing engine. However, there are significant differences between the legacy and advanced routing engines that may require you to monitor and provide information for your environment during the migration process.

To migrate Panorama managed firewalls, both Panorama and local firewall configuration must be migrated as follows:

- Enable **Advanced Routing** on Panorama to migrate and push the Panorama configuration to all devices.
- Disable **Advanced Routing** in each Panorama managed firewall and don't commit the changes.
- Enable **Advanced Routing** for each firewall managed by Panorama to migrate the local firewall configuration.
- Commit the changes on each firewall and restart the firewalls for the changes to take effect.

After you enable **Advanced Routing** on a firewall or through your Panorama management server, a built-in migration script will migrate your existing legacy routing configuration to the advanced routing engine configuration. When the script finishes, the **Migration Configuration** displays color codes that indicate the migration status.

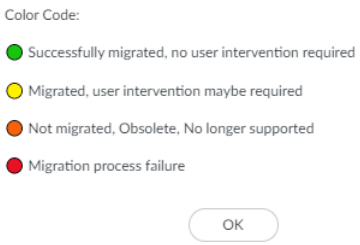

If a migration exception occurs, it's highlighted in yellow or orange depending on the action required. Additionally, any exceptions detected during the migration process result in an incomplete migration and you will need to resolve the issue before you can attempt the migration process again.

STATUS COLOR CODE	STATUS DESCRIPTION	ACTION
Green	No issues were encountered when migrating the existing protocol features into the advanced routing engine configuration.	No action required.

STATUS COLOR CODE	STATUS DESCRIPTION	ACTION
Yellow	The advanced routing engine supports one or more features that exist in the legacy routing engine configuration but uses different parameters.	No action is required because the alternate configuration guideline is available.
Orange	One or more features or settings are no longer supported.	Identify if any configuration correction is required and update the configuration after migration manually.
Red	A migration script failure occurred.	Generate the technical support file, log in to Palo Alto Networks Customer Support Portal, and report your issues to get help with your product.

This migration reference includes information about any exceptions and proposes solutions that help make the migration process go as smoothly as possible.

Review the specific protocols listed in this migration reference before you begin the migration process. If you find that there are exceptions, either fix them beforehand or identify the appropriate solution before the script pauses for your input. After the migration is complete, ensure that the key elements of your design are preserved by reviewing the final configuration.

To ensure an accurate migration process on the firewalls, including high availability (HA) active/active and active/passive configurations take snapshots of both the Routing Information Base (RIB) and Forwarding Information Base (FIB) before you begin the migration process and then compare the results after the migration.

Learn the Differences Between Legacy and Advanced Routing Engine

The advanced routing engine aligns with widespread configuration methodologies in the networking industry. Some of the configuration shortcuts available in the legacy routing engine have been deprecated, and corresponding well-known configuration guidelines are made available in the advance routing engine.

In the advanced routing engine, we replaced virtual routers with logical routers.

Virtual and logical routers are conceptually equivalent but a few configuration parameters may require action because the implementation differs.

LEGACY ROUTING ENGINE	ADVANCED ROUTING ENGINE	MIGRATION EXCEPTION
Supports virtual router.	Supports logical router.	Unlike the legacy routing engine, the advanced routing engine does not create a default (logical) router.
Routing profiles are unique to the virtual router and are not shared.	Routing profiles are used to configure route filters, route redistribution, routing protocol profiles, and the advanced routing engine relies extensively on these routing profiles. These routing profiles are shared among the logical routers within the same virtual system as well as across virtual systems. This makes it possible to reuse these profiles among logical routers and protocols in the logical routers. A BGP peer in a peer group can inherit routing profiles from that BGP peer group. In addition, BGP peers can have their own routing profiles, which are not inherited from their BGP peer group.	Routing profiles may create issues when translating profiles from the legacy routing engine to the advanced routing engine. In the legacy routing engine, profiles are unique to the virtual router and are not shared. You may encounter issues with profiles that use the same name in more than one virtual router.

LEGACY ROUTING ENGINE	ADVANCED ROUTING ENGINE	MIGRATION EXCEPTION
Enable you to create filters and policy rules without applying them to an object. For example, the legacy routing engine enables you to create a BGP import or export policy rule without applying it to any peers or peer groups.	Enable you to create filters and policy rules without applying them to an object, similar to the legacy routing engine.	To optimize the migration process, these orphan filters and profiles are not converted during migration.
Supports route tagging. The legacy routing engine uses a 32-bit dotted decimal value to represent the tag (such as 10.1.7.1).	Supports route tagging more completely and effectively than the legacy routing engine. The advanced routing engine uses 32-bit decimal notation to represent the tag (such as 1234567).	No exceptions; tags are migrated from dotted decimal format to decimal notation successfully.

Routing Protocol Migration Exceptions

The legacy and advanced routing engines have different capabilities and they each support different configuration options. Although most legacy routing engine and advanced routing engine configuration parameters are similar, they have some important differences that include different configuration parameters, parameter ranges, default values, and even features.

Review the different configuration parameters and feature support differences between the legacy and advanced routing engines for each of the following routing protocols.

- MP-BGP
- OSPF
- OSPFv3
- PIM
- IGMP

MP-BGP

The advanced routing engine provides the same functionality as the legacy routing engine but with enhanced capabilities. For example, PAN-OS 11.0 enables you to advertise IPv4 Network Layer Reachability Information (NLRI) with an IPv6 next hop address. As a result, you can deploy Palo Alto Networks Next-Generation Firewalls in a dual stack network using fewer peers.

There are several multiprotocol BGP (MP-BGP) configuration differences between the legacy and advanced routing engines.

- Dampening Profile
- Route Reflector Client Parameters
- Route Map
- Multicast
- AS Path Limit Attribute

Dampening Profile

Migration Exception: The virtual router dampening profiles on the legacy routing engine are incompatible with advanced routing engine dampening profiles. The dampening profile will migrate with only default values and will not be linked to any peer groups.

The following table compares the dampening profile parameters of the two routing engines:

CONFIGURED IN (LEGACY ROUTING ENGINE)	LEGACY ROUTING ENGINE	MIGRATED TO (ADVANCED ROUTING ENGINE)	ADVANCED ROUTING ENGINE
Network > Virtual Router > BGP > Advanced	Dampening parameters: • **Cutoff**—Specifies a route withdrawal threshold above which a route advertisement is suppressed (range is 0.0 to 1,000.0; default is 1.25). • **Reuse**—Specifies a route withdrawal threshold below which a suppressed route is reused (range is 0.0 to 1,000.0; default is 5). • **Max Hold Time (sec)**—Specifies the	**Network > Routing > Routing Profiles > BGP > BGP Dampening Profiles**	Dampening parameters: • **Suppress Limit**—Specifies the cumulative value of the penalties for flapping, at which point all the routes coming from a peer are dampened (range is 1 to 20,000; default is 2,000). • **Reuse Limit**—Controls when a route can be reused based on the procedure described for **Half Life** (range is

CONFIGURED IN (LEGACY ROUTING ENGINE)	LEGACY ROUTING ENGINE	MIGRATED TO (ADVANCED ROUTING ENGINE)	ADVANCED ROUTING ENGINE
	maximum length of time, in seconds, that a route can be suppressed, regardless of how unstable it is (range is 0 to 3,600; default is 900). • **Decay Half Life Reachable (sec)**—Specifies the length of time, in seconds, after which a route's stability metric is halved if the firewall considers the route to be reachable (range is 0 to 3,600; default is 300). • **Decay Half Life Unreachable (sec)**—Specifies the length of time, in seconds, after which a route's stability metric is halved if the firewall considers the route is unreachable (range is 0 to 3,600; default is 300).		1 to 20,000; default is 750). • **Half Life (min)**—Specifies the number of minutes for the half life time to control the stability metric (penalty) applied to a flapping route (range is 1 to 45; default is 15). The stability metric starts at 1,000. After a penalized route stabilizes, the half life timer counts down until it expires, at which point the next stability metric applied to the router is only half of the previous value (500). Successive cuts continue until the stability metric is less than half of the **Reuse Limit**, and then the stability metric is removed from the router. • **Maximum Suppress Time (min)**—Specifies the maximum number of minutes a route can be suppressed, regardless of how unstable it has been (range is 1 to 255; default is 60).

Route Reflector Client Parameters

To avoid routing table loops, interior BGP (iBGP) does not advertise iBGP-learned routes to other routers in the same session to avoid routing table loops. As a result, iBGP requires a complete mesh of all peers, which quickly becomes unscalable in large networks. Using route reflectors eliminates the need for full-mesh connectivity between iBGP peers.

Route reflectors broadcast routes announced by peers that are configured as clients to all other clients.

Migration Exception: The advanced routing engine supports only the route reflector client mode; no other modes are supported. The advanced routing engine receives routes from the route reflector in client mode and can send routes only to a route reflector when client mode is enabled.

CONFIGURED IN (LEGACY ROUTING ENGINE)	LEGACY ROUTING ENGINE	MIGRATED TO (ADVANCED ROUTING ENGINE)	ADVANCED ROUTING ENGINE
Network > **Virtual Router** > **BGP** > **Peer Group** > **Peer** > **Advanced**	Supported **Reflector Client** types: • **Non-Client**—Specifies that the firewall reflects all routes from non-client to all clients. • **Client**—Specifies that routes advertised by this client type are reflected to all non-client and client peers. • **Meshed-Client**—Specifies that routes advertised from a meshed client are reflected to all neighbors except for other meshed-client iBGP peers.	**Network** > **Routing** > **Routing Profiles** > **BGP** > **BGP Address Family Profiles**	Supported **Reflector Client** types: • **Route Reflector Client**—Specifies that routes advertised by this client type are reflected to all non-client and client peers.

Route Map

Palo Alto Networks recommends BGP route maps for filtering prefixes within BGP and both from and to another interior gateway protocol (IGP). However, BGP route maps do not support configuring extended communities in route maps set action.

CONFIGURED IN (LEGACY ROUTING ENGINE)	LEGACY ROUTING ENGINE	MIGRATED TO (ADVANCED ROUTING ENGINE)	ADVANCED ROUTING ENGINE
Network > **Virtual Router** > **BGP** > **Redist Rules**	• **Set Community**—Supports standard community. Enter a 32-bit value in decimal or hexadecimal, or enter a value in AS:VAL format, where Autonomous System (AS) and VAL are each in the range 0 to 65,525. You can enter a maximum of 10 entries. • **Set Extended Community**—Set Extended Community is a 64-bit value in hexadecimal or in TYPE:AS:VAL or TYPE:IP:VAL format. TYPE is 16 bits; AS or IP is 16 bits; VAL is 32 bits. You can enter a maximum of five entries.	**Network** > **Routing** > **Routing Profiles** > **Filters** > **Filters Route Map BGP**	• **Regular Community**—In the **Set** tab, for **Regular Community**, enter either AS:VAL pairs or well-known community names. • **Large Community**—Enables additional functionality and convenience over traditional community. The 32-bit Global Administrator field (GLOBAL) enables seamless use in networks using 4-byte Autonomous System Number (ASN). In the **Set** tab, **Large Community** has three components instead of two and each are 32-bits values. Define **Large Community** values in GLOBAL:LOCAL1:LOCAL2 format. Where, GLOBAL is a 32-bit Global Administrator field (commonly used as the AS number of the operator); LOCAL1 is a 32-bit Local Data Part 1 subfield (referred to as a function); LOCAL2

CONFIGURED IN (LEGACY ROUTING ENGINE)	LEGACY ROUTING ENGINE	MIGRATED TO (ADVANCED ROUTING ENGINE)	ADVANCED ROUTING ENGINE
			is a 32-bit Local Data Part 2 field (referred to as the parameter subfield). For example, 65551:1:10 represents AS as 65551, function as 1, and parameter as 10.

Multicast

Both the legacy and the advanced routing engines support the multicast subsequent address family identifier (SAFI) for IPv4 addresses.

Migration Exception: The advanced routing engine doesn't redistribute multicast source prefixes into MP-BGP (IPv4 address family) and multicast subsequent family; hence, it can't be used for reverse path forwarding (RPF) check.

LEGACY ROUTING ENGINE	ADVANCED ROUTING ENGINE
• Enables you to configure multicast as SAFI for specific IPv4 peers. • Enables you to redistribute a static IPv4 route used for multicast RPF verification into BGP.	• Enables you to select the multicast SAFI in the BGP Address Family Identifier (AFI) profile. After you select an multicast SAFI, the profile is applied to all peers assigned to this profile. • Does not support redistribution of static IPv4 routes used for multicast RPF verification into BGP.

AS Path Limit Attribute

The AS Path Limit is an optional path transitive attribute. It improves routing subsystem scalability by providing a maximum range of Autonomous System (AS) numbers where a prefix will propagate. If used improperly, this attribute can cause routing loops caused by inconsistent routing tables. As a result, the IETF didn't standardize this attribute.

Migration Exception: The advanced routing engine does not support AS path limit attribute; it will ignore the attribute and advertise the prefix without AS path limit attribute.

CONFIGURED IN (LEGACY ROUTING ENGINE)	LEGACY ROUTING ENGINE	ADVANCED ROUTING ENGINE
Network > **Virtual Router** > **BGP** > **Import or Export** > **Action**	Supports the AS path limit attribute. If configured, the AS path limit attribute will be exchanged with peers and is applied to prefixes.	Does not support AS path limit attribute. The advanced routing engine ignores the AS path limit attribute and advertise the prefix without AS path limit attribute. If you need this attribute, replace it with a route map that matches a specific AS path length (using a regular expression) and configure the community to not advertise.

OSPF

There are parameter setting differences between legacy and advanced routing engines when configuring OSPF settings.

Route Redistribution

OSPF handles route redistribution to another routing protocol differently in each engine. When redistributing from OSPF, filters are based on a variety of criteria.

Migration Exception: The advanced routing engine does not allow redistribution of routes based on Link State Advertisement (LSA) type (external-1, external-2, inter-area, or intra-area) or origin area. You can, however use prefixes and tags associated with LSA updates.

CONFIGURED IN (LEGACY ROUTING ENGINE)	LEGACY ROUTING ENGINE	MIGRATED TO (ADVANCED ROUTING ENGINE)	ADVANCED ROUTING ENGINE
Network > **Virtual Router** > **OSPF** > **Export Rules**	Supports the redistribution of routes that match the following criteria: • **Interface** • **Address** • **Next Hop** • **Path Type** • **Area** • **Tag**	**Network** > **Routing** > **Routing Profiles** > **OSPF** > **OSPF Redistribution Profile**	Supports **Redistribute Route-Map** that **Match** the following criteria: • **Interface** • **Address** (using **Prefix List** or **Access List**) • **Next Hop** (using **Prefix List** or **Access List**) • **Metric** • **Tag**

OSPFv3

There are parameter setting differences between legacy and advanced routing engines when configuring OSPFv3 settings.

Route Redistribution

OSPFv3 handles route redistribution to another routing protocol differently in the legacy and advanced routing engines. When redistributing from OSPFv3, filters are based on a variety of criteria.

Migration Exception: The advanced routing engine does not allow redistribution of routes based on Link State Advertisement (LSA) type (external-1, external-2, inter-area, or intra-area) or origin area. You can, however, use prefixes and tags associated with LSA updates.

CONFIGURED IN (LEGACY ROUTING ENGINE)	LEGACY ROUTING ENGINE	MIGRATED TO (ADVANCED ROUTING ENGINE)	ADVANCED ROUTING ENGINE
Network > **Virtual Router** > **OSPFv3** > **Export Rules**	Supports redistribution of routes that match the following criteria: • **Interface** • **Address** • **Next Hop** • **Path Type** • **Area** • **Tag**	**Network** > **Routing** > **Routing Profiles** > **OSPFv3** > **OSPFv3 Redistribution Profile**	Supports **Redistribute Route-Map** that **Match** the following criteria: • **Interface** • **Address** (using **Prefix List** or **Access List**) • **Next Hop** (using **Prefix List** or **Access List**) • **Metric** • **Tag**

Authentication Profile

Migration Exception: In the advanced routing engine, the authentication profile variables are retained globally for reuse which can cause an issue with the Security Policy Index (SPI). The advanced routing engine does not allow multiple interface authentication profiles within a single virtual system (vsys) to share the same SPI value.

CONFIGURED IN (LEGACY ROUTING ENGINE)	LEGACY ROUTING ENGINE	MIGRATED TO (ADVANCED ROUTING ENGINE)	ADVANCED ROUTING ENGINE
Network > **Virtual Router** > **OSPFv3** > **Auth Profiles**	Maintains the authentication profile variables within the context of virtual routers.	**Network** > **Routing** > **Routing Profiles** > **OSPFv3** > **OSPFv3 Interface Auth Profile**	Retains the authentication profile variables globally for reuse.

PIM

There are parameter setting differences between legacy and advanced routing engines when configuring Protocol Independent Multicast (PIM) settings.

Migration Exception:

- During migration, each entry in the access list is replaced with a sequence number.
- While a legacy routing engine can apply group permissions on specific interfaces, the advanced routing engine applies group permissions only globally. As a result, any interface-specific PIM group permissions on the legacy routing engine are not migrated to the advanced routing engine.

LEGACY ROUTING ENGINE	ADVANCED ROUTING ENGINE
A legacy routing engine refers to group permissions by name. This applies to both source-specific and non-source specific (any source) entries.	The advanced routing engine uses extended access lists as filters. Each entry in these access lists uses a sequence number instead of a name. As a result, the entry names are lost and are replaced by a sequence number in the access list during the migration.
A legacy routing engine can apply group permissions to specific interfaces.	The advanced routing engine can apply group permissions only globally—not to specific interfaces. Therefore, any interface-specific PIM group permissions configured on the legacy routing engine are not migrated.

PIM Timers

Migration Exception: The configurable values of interface timers vary between the legacy and advanced routing engines.

CONFIGURED IN (LEGACY ROUTING ENGINE)	LEGACY ROUTING ENGINE	MIGRATED TO (ADVANCED ROUTING ENGINE)	ADVANCED ROUTING ENGINE
Network > **Virtual Router** > **Multicast** > **Interfaces** > **PIM**	Specifies only a global set of timers for each function. These timers are applied to all interfaces participating in the PIM routing domain. PIM parameters for interface group:	**Network** > **Routing** > **Routing Profiles** > **Multicast** > **Multicast IPv4 PIM Interface Timer Profiles**	The PIM interface timer is part of the multicast profile that you can configure per interface. • **Assert Interval**—Specifies the interval (in seconds) between PIM Assert messages to elect a PIM Forwarder (range is 1 to

CONFIGURED IN (LEGACY ROUTING ENGINE)	LEGACY ROUTING ENGINE	MIGRATED TO (ADVANCED ROUTING ENGINE)	ADVANCED ROUTING ENGINE
	• **Assert Interval**—Specifies the interval (in seconds) between PIM Assert messages to elect a PIM Forwarder (range is 0 to 65,534; default is 177). • **Hello Interval**—Specifies the interval (in seconds) between PIM Hello messages (range is 0 to 18,000; default is 30). • **Join Prune Interval**—Specifies the number of seconds between the two PIM-Join messages and between two PIM-Prune messages (range is 1 to 18,000; default is 60).		65,534; default is 177). • **Hello Interval**—Specifies the interval (in seconds) between PIM Hello messages (range is 1 to 180; default is 30). A legacy routing range configured with a value greater than 180 is changed during migration to a value supported by the advanced routing engine. • **Join Prune Interval**—Specifies the number of seconds between PIM Join messages and between PIM Prune messages (range is 60 to 600; default is 60). A legacy routing range configured with a value greater than 600 is changed during migration to a value supported by the advanced routing engine.

PIM - Multicast Interfaces

Migration Exception:

- The legacy routing engine **BSR Border** settings are migrated to the **Send BSM** settings in the advanced routing engine.
- The legacy routing engine **Permitted PIM Neighbors** are migrated to the **Neighbor Filter** in the advanced routing engine.
- The **DR Priority** configuration range differs between the legacy and advanced routing engines.

CONFIGURED IN (LEGACY ROUTING ENGINE)	LEGACY ROUTING ENGINE	MIGRATED TO (ADVANCED ROUTING ENGINE)	ADVANCED ROUTING ENGINE
Network > **Virtual Router** > **Multicast** > **Interfaces** > **PIM**	PIM configuration parameters for multicast interfaces: • **BSR Border**—Enables use of the interface as the bootstrap border. • **Permitted PIM Neighbors**—Specifies the list of neighbors that can communicate using PIM. • **DR Priority**—Specifies the designated router priority for an interface group (range is 0 to 4,294,967,295; default is 1).	**Network** > **Routing** > **Logical Routers** > **Multicast** > **PIM - Interfaces**	PIM configuration parameters for multicast interfaces: • **Send BSM**—Enables propagation of Bootstrap Messages (enabled by default). • **Neighbor Filter**—Specifies the access list that determines the prefixes of devices that are or are not allowed to become PIM neighbors of the logical router (default is none—no access list). • **DR Priority**—Specifies the designated router priority for an interface group (range is 1 to 4,294,967,295; default is 1).

IGMP

There are parameter setting differences between legacy and advanced routing engines when configuring Internet Group Management Protocol (IGMP) settings.

Migration Exception:

- During migration, each entry in the access list is replaced with a sequence number.
- The advanced routing engine does not apply group permissions directly to specific interfaces.

LEGACY ROUTING ENGINE	ADVANCED ROUTING ENGINE
A legacy routing engine refers to group permissions by name. This applies to IGMPv1, IGMPv2, and IGMPv3 entries.	The advanced routing engine uses extended access lists as filters. Each entry in these access lists uses a sequence number instead of a name. As a result, the entry names are lost and are replaced with a sequence number in the access list during the migration.
A legacy routing engine can apply group permissions directly to specific interfaces.	The advanced routing engine applies group permissions per interface but does not apply group permissions directly to the interfaces. Instead, the advanced routing engine creates an access list (**Logical Routers** > **Multicast** > **IGMP** > **Dynamic** > **Group Filter**) and then applies the access list to specific interfaces.

IGMP Version Support

Migration Exception: The legacy routing engine supports IGMPv1, IGMPv2, and IGMPv3 but the advanced routing engine supports only IGMPv2 and IGMPv3.

LEGACY ROUTING ENGINE	ADVANCED ROUTING ENGINE
Supports IGMPv1 (RFC 1112), IGMPv2 (RFC 2236), and IGMPv3 (RFC 3376). IGMPv1 supports only two messages: membership query and membership reply. There isn't a separate message supported to announce when a host unsubscribes from a multicast group. Because there is not an explicit leave message, multicast streams can flood a specific segment even when there are no subscribers.	The advanced routing engine supports only IGMPv2 and IGMPv3—support for IGMPv1 is deprecated. However, IGMPv2 is backward compatible and can process IGMPv1 packets.

Differences between IGMPv1 and IGMPv2

FEATURE	IGMPv1	IGMPv2
Default Query Interval (in seconds)	60	125
Max Response Time (in seconds)	10 seconds (fixed)	Range is 0 to 25 seconds (configurable)
Leave Messages	No	Yes
IP Address for Leave Messages	—	224.0.0.2
Support for Group-Specific Query	No	Yes
Querier Election Mechanism	None (depends on PIM)	Router with the lowest IP address on the subnet.

IGMP Timers

The IGMP timer configuration range varies between legacy and advanced routing engines.

Migration Exception: The legacy routing engine IGMP timer value is converted during migration to a value within the supported advanced routing engine configuration range.

- If the IGMP timer configured in the legacy routing engine is a decimal value, then—during migration—it is rounded to the nearest integer value within the valid advanced routing range. This exception impacts the **Max Query Response Time** and **Last Member Query Interval**.
- If the IGMP timer configured in the legacy routing engine is an integer value or is migrated to an integer value that exceeds the advanced routing engine configuration range, then this value is changed to a value supported by the advanced routing engine.

CONFIGURED IN (LEGACY ROUTING ENGINE)	LEGACY ROUTING ENGINE	MIGRATED TO (ADVANCED ROUTING ENGINE)	ADVANCED ROUTING ENGINE
Network > **Virtual Router** > **Multicast** > **Interfaces** > **IGMP**	Supports deciseconds value for few parameters. The legacy routing engine also supports extended ranges for some IGMP parameters. • **Max Query Response Time**—Specifies the	**Network** > **Routing** > **Routing Profiles** > **Multicast** > **Multicast IPv4 IGMP Interface**	Supports the following IGMP timer configurations: • **Max Query Response Time**—Specifies the maximum response time in seconds within which a host must respond to an IGMP

CONFIGURED IN (LEGACY ROUTING ENGINE)	LEGACY ROUTING ENGINE	MIGRATED TO (ADVANCED ROUTING ENGINE)	ADVANCED ROUTING ENGINE
	maximum response time in seconds within which a host must respond to an IGMP query message (range is 0 to 3,174.4; default is 10). • **Query Interval**—Specifies the periodic interval in seconds that defines how often the querier sends IGMP host-query messages from an interface (range is 1 to 31,744; default is 125). • **Last Member Query Interval**—Specifies the value in seconds that IGMP uses when a router hears an IGMPv2 Leave report. If the interface is not configured for Immediate Leave, the firewall sends a set of group-specific queries at a specified interval. If the querier does not receive a response to the queries, it removes the group state and discontinues multicast transmissions (range is 0.1 to 3,174.4; default is 1).	**Query Profiles**	query message (range is 1 to 25; default is 10). • **Query Interval**—Specifies the periodic interval in seconds that defines how often the querier sends IGMP host-query messages from an interface (range is 1 to 1,800; default is 125). • **Last Member Query Interval**—Specifies the value in seconds that IGMP uses when a router hears an IGMPv2 Leave report. If the interface is not configured for Immediate Leave option, the firewall sends a set of group-specific queries at a specified interval. If the querier does not receive a response to the queries, it removes the group state and discontinues multicast transmissions (range is 1 to 25; default is 1 second).

IGMP Group Permission

Migration Exception: A legacy routing engine refers to access list entries by name. The advanced routing engine, instead, identifies access list entries by a sequence number. In the advanced routing engine, the access lists use the source to represent the multicast source and the destination to represent the group.

CONFIGURED IN (LEGACY ROUTING ENGINE)	LEGACY ROUTING ENGINE	ADVANCED ROUTING ENGINE
Network > **Virtual Router** > **Multicast** > **Interfaces** > **Group Permissions**	Supports the ability to restrict client access to specific groups or specific sources associated with a specific group. These group permissions lists identify each entry by a name and applies to both source-specific and non-source-specific (any source) entries.	Supports extended access lists that restrict access to specific groups or sources associated with a specific group. The advanced routing engine can apply separate access lists to each IGMP-enabled interface. These access lists are defined in the **Filters** routing profile section of the configuration.