

EXHIBIT

4

Claim Chart for U.S. Patent No. 8,014,394 | Claim 1 | Juniper

Claim Chart for U.S. Patent No. 8,014,394 | Claim 1 | Juniper

Claim Element	Juniper Networking Equipment (e.g., Juniper Networks MX240 Universal Routing Platform)
A method for communication, comprising:	To the extent the preamble is limiting, Juniper networking equipment, including, for example, the Juniper Networks MX240 Universal Routing Platform, perform a method for communication. Utilizing state-of-the-art software and hardware innovations, MX Series Universal Routing Platforms are helping network operators worldwide successfully transform their networks and services. Powered by the Juniper Networks Junos® operating system and the programmable Trio chipset, MX Series platforms support a broad set of automation tools and telemetry capabilities that enable a rich set of business- and consumer-oriented services with predictable low latency and wire-rate forwarding at scale, while providing the reliability needed to meet strict service-level agreements (SLAs). See e.g., MX Series Universal Routing Platforms Datasheet at 1 (available at https://www.juniper.net/content/dam/www/assets/datasheets/us/en/routers/mx-series-universal-routing-platforms-datasheet.pdf).
operating a network element, wherein the network element further comprises multiple interconnected processing units and multiple ports, wherein each processing unit is assigned to a different subset of the multiple ports of the network element, and wherein each processing	The Juniper Networks MX240 Universal Routing Platform performs the step of operating a network element, wherein the network element further comprises multiple interconnected processing units and multiple ports, wherein each processing unit is assigned to a different subset of the multiple ports of the network element, and wherein each processing unit contains a processor for processing packets communicated over the different subset of the multiple ports to which the processing unit is assigned.

Claim Chart for U.S. Patent No. 8,014,394 | Claim 1 | Juniper

Claim Element	Juniper Networking Equipment (e.g., Juniper Networks MX240 Universal Routing Platform)																				
unit contains a processor for processing packets communicated over the different subset of the multiple ports to which the processing unit is assigned;	Table 1. MX Series Universal Routing Platforms at a Glance <table><tr><th></th><th>MX960</th><th>MX480</th><th>MX240</th></tr><tr><td>Rack units</td><td>16</td><td>8</td><td>5</td></tr><tr><td>Systems per rack</td><td>3</td><td>6</td><td>9</td></tr><tr><td>Slots</td><td>11 MPCs</td><td>6 MPCs</td><td>2 MPCs</td></tr><tr><td>Per slot capacity</td><td>1.5 Tbps</td><td>1.5 Tbps</td><td>1.5 Tbps</td></tr></table> See e.g., MX Series Universal Routing Platforms Datasheet at 3 (available at https://www.juniper.net/content/dam/www/assets/datasheets/us/en/routers/mx-series-universal-routing-platforms-datasheet.pdf).		MX960	MX480	MX240	Rack units	16	8	5	Systems per rack	3	6	9	Slots	11 MPCs	6 MPCs	2 MPCs	Per slot capacity	1.5 Tbps	1.5 Tbps	1.5 Tbps
	MX960	MX480	MX240																		
Rack units	16	8	5																		
Systems per rack	3	6	9																		
Slots	11 MPCs	6 MPCs	2 MPCs																		
Per slot capacity	1.5 Tbps	1.5 Tbps	1.5 Tbps																		

Claim Chart for U.S. Patent No. 8,014,394 | Claim 1 | Juniper

Claim Element	Juniper Networking Equipment (<i>e.g.</i> , Juniper Networks MX240 Universal Routing Platform)
	MPC-10E Line Card The MPC-10E line card is a key contributor to the service provider transformation in the cloud era when deployed with MX960, MX480, and MX240 platforms in a Juniper Secure Automated Distributed Cloud environment. By providing the underlying network infrastructure with scale, agility, routing innovation, and pervasive security while incorporating universal (10/40/100/400GbE) ports, the MPC-10E protects existing investments with disaggregated software innovation and infinite programmability. Built-in automation enables rapid deployment without disrupting the existing MX960/MX480/MX240 footprint. The MPC-10E line card is powered by Juniper Trio 5 silicon, which enables the benefits highlighted in Table 2. <i>See e.g., id.</i> at 4.
storing in each processing unit a respective list, each list indicating all of one or more multicast packet streams that are permitted for forwarding by the network element, and	The Juniper Networks MX240 Universal Routing Platform performs the step of storing in each processing unit a respective list, each list indicating all of one or more multicast packet streams that are permitted for forwarding by the network element, and further indicating which of the one or more ports each multicast packet stream is permitted to be forwarded to.

Claim Chart for U.S. Patent No. 8,014,394 | Claim 1 | Juniper

Claim Element	Juniper Networking Equipment (<i>e.g.</i> , Juniper Networks MX240 Universal Routing Platform)
further indicating which of the one or more ports each multicast packet stream is permitted to be forwarded to;	IGMP version 3 (IGMPv3) supports inclusion and exclusion lists. Inclusion lists enable you to specify which sources can send to a multicast group. This type of multicast group is called a source-specific multicast (SSM) group, and its multicast address is 232/8. IGMPv3 provides support for source filtering. For example, a routing device can specify particular routing devices from which it accepts or rejects traffic. With IGMPv3, a multicast routing device can learn which sources are of interest to neighboring routing devices. <i>See e.g.</i>, Junos OS – Multicast Protocols User Guide at 28 (available at https://www.juniper.net/documentation/en_US/junos/information-products/pathway-pages/config-guide-multicast/config-guide-multicast.pdf). Distributed IGMP Overview Distributed IGMP works by moving IGMP processing from the Routing Engine to the Packet Forwarding Engine. When distributed IGMP is not enabled, IGMP processing is centralized on the routing protocol process (rpd) running on the Routing Engine. When you enable distributed IGMP, join and leave events are processed across Modular Port Concentrators (MPCs) on the Packet Forwarding Engine. Because join and leave processing is distributed across multiple MPCs instead of being processed through a centralized rpd on the Routing Engine, performance improves and join and leave latency decreases. When you enable distributed IGMP, each Packet Forwarding Engine processes reports and generates queries, maintains local group membership to the interface mapping table and updates the forwarding state based on this table, runs distributed IGMP independently, and implements the <code>group-policy</code> and <code>ssm-map-policy</code> IGMP interface options. <i>See e.g., id.</i> at 90.

Claim Chart for U.S. Patent No. 8,014,394 | Claim 1 | Juniper

Claim Element	Juniper Networking Equipment (<i>e.g.</i> , Juniper Networks MX240 Universal Routing Platform)
	How IGMP Snooping Works Devices usually learn unicast MAC addresses by checking the source address field of the frames they receive and then send any traffic for that unicast address only to the appropriate interfaces. However, a multicast MAC address can never be the source address for a packet. As a result, when a device receives traffic for a multicast destination address, it floods the traffic on the relevant VLAN, sending a significant amount of traffic for which there might not necessarily be interested receivers. IGMP snooping prevents this flooding. When you enable IGMP snooping, the device monitors IGMP packets between receivers and multicast routers and uses the content of the packets to build a multicast forwarding table—a database of multicast groups and the interfaces that are connected to members of the groups. When the device receives multicast packets, it uses the multicast forwarding table to selectively forward the traffic to only the interfaces that are connected to members of the appropriate multicast groups. <i>See e.g., id.</i> at 97. Support for IGMPv3 Multicast Sources In IGMPv3, a host can send a membership report that includes a list of source addresses. When the host sends a membership report in INCLUDE mode, the host is interested in group multicast traffic only from those sources in the source address list. If host sends a membership report in EXCLUDE mode, the host is interested in group multicast traffic from any source <i>except</i> the sources in the source address list. A host can also send an EXCLUDE report in which the source-list parameter is empty, which is known as an EXCLUDE NULL report. An EXCLUDE NULL report indicates that the host wants to join the multicast group and receive packets from all sources. Devices that support IGMPv3 process INCLUDE and EXCLUDE membership reports, and most devices forward source-specific multicast (SSM) traffic only from requested sources to subscribed receivers accordingly. However, you might see the device doesn't strictly forward multicast traffic on a per-source basis in some configurations such as: EX Series and QFX Series switches that do not use the Enhanced Layer 2 Software (ELS) configuration style <i>See e.g., id.</i> at 99.

Claim Chart for U.S. Patent No. 8,014,394 | Claim 1 | Juniper

Claim Element	Juniper Networking Equipment (<i>e.g.</i> , Juniper Networks MX240 Universal Routing Platform)
	IGMP Snooping and Forwarding Interfaces To determine how to forward multicast traffic, the device with IGMP snooping enabled maintains information about the following interfaces in its multicast forwarding table: • Multicast-router interfaces—These interfaces lead toward multicast routers or IGMP queriers. • Group-member interfaces—These interfaces lead toward hosts that are members of multicast groups. The device learns about these interfaces by monitoring IGMP traffic. If an interface receives IGMP queries or Protocol Independent Multicast (PIM) updates, the device adds the interface to its multicast forwarding table as a multicast-router interface. If an interface receives membership reports for a multicast group, the device adds the interface to its multicast forwarding table as a group-member interface. Learned interface table entries age out after a time period. For example, if a learned multicast-router interface does not receive IGMP queries or PIM hellos within a certain interval, the device removes the entry for that interface from its multicast forwarding table. NOTE: For the device to learn multicast-router interfaces and group-member interfaces, the network must include an IGMP querier. This is often in a multicast router, but if there is no multicast router on the local network, you can configure the device itself to be an IGMP querier. You can statically configure an interface to be a multicast-router interface or a group-member interface. The device adds a static interface to its multicast forwarding table without having to learn about the interface, and the entry in the table is not subject to aging. A device can have a mix of statically configured and dynamically learned interfaces. <i>See e.g., id.</i> at 100.
receiving from a client a request relating to reception of a given multicast packet stream;	The Juniper Networks MX240 Universal Routing Platform performs the step of receiving from a client a request relating to reception of a given multicast packet stream.

Claim Chart for U.S. Patent No. 8,014,394 | Claim 1 | Juniper

Claim Element	Juniper Networking Equipment (<i>e.g.</i> , Juniper Networks MX240 Universal Routing Platform)
	Distributed IGMP Overview Distributed IGMP works by moving IGMP processing from the Routing Engine to the Packet Forwarding Engine. When distributed IGMP is not enabled, IGMP processing is centralized on the routing protocol process (rpd) running on the Routing Engine. When you enable distributed IGMP, join and leave events are processed across Modular Port Concentrators (MPCs) on the Packet Forwarding Engine. Because join and leave processing is distributed across multiple MPCs instead of being processed through a centralized rpd on the Routing Engine, performance improves and join and leave latency decreases. When you enable distributed IGMP, each Packet Forwarding Engine processes reports and generates queries, maintains local group membership to the interface mapping table and updates the forwarding state based on this table, runs distributed IGMP independently, and implements the <code>group-policy</code> and <code>ssm-map-policy</code> IGMP interface options. <i>See e.g.</i>, Junos OS – Multicast Protocols User Guide at 90 (available at https://www.juniper.net/documentation/en_US/junos/information-products/pathway-pages/config-guide-multicast/config-guide-multicast.pdf).

Claim Chart for U.S. Patent No. 8,014,394 | Claim 1 | Juniper

Claim Element	Juniper Networking Equipment (<i>e.g.</i> , Juniper Networks MX240 Universal Routing Platform)
	How IGMP Snooping Works Devices usually learn unicast MAC addresses by checking the source address field of the frames they receive and then send any traffic for that unicast address only to the appropriate interfaces. However, a multicast MAC address can never be the source address for a packet. As a result, when a device receives traffic for a multicast destination address, it floods the traffic on the relevant VLAN, sending a significant amount of traffic for which there might not necessarily be interested receivers. IGMP snooping prevents this flooding. When you enable IGMP snooping, the device monitors IGMP packets between receivers and multicast routers and uses the content of the packets to build a multicast forwarding table—a database of multicast groups and the interfaces that are connected to members of the groups. When the device receives multicast packets, it uses the multicast forwarding table to selectively forward the traffic to only the interfaces that are connected to members of the appropriate multicast groups. <i>See e.g., id.</i> at 97.
distributing the request among the multiple processing units within the network element, and updating the respective list in each processing unit responsively to the request; and	The Juniper Networks MX240 Universal Routing Platform performs the step of distributing the request among the multiple processing units within the network element, and updating the respective list in each processing unit responsively to the request.

Claim Chart for U.S. Patent No. 8,014,394 | Claim 1 | Juniper

Claim Element	Juniper Networking Equipment (<i>e.g.</i> , Juniper Networks MX240 Universal Routing Platform)
	Distributed IGMP Overview Distributed IGMP works by moving IGMP processing from the Routing Engine to the Packet Forwarding Engine. When distributed IGMP is not enabled, IGMP processing is centralized on the routing protocol process (rpd) running on the Routing Engine. When you enable distributed IGMP, join and leave events are processed across Modular Port Concentrators (MPCs) on the Packet Forwarding Engine. Because join and leave processing is distributed across multiple MPCs instead of being processed through a centralized rpd on the Routing Engine, performance improves and join and leave latency decreases. When you enable distributed IGMP, each Packet Forwarding Engine processes reports and generates queries, maintains local group membership to the interface mapping table and updates the forwarding state based on this table, runs distributed IGMP independently, and implements the <code>group-policy</code> and <code>ssm-map-policy</code> IGMP interface options. <i>See e.g.</i>, Junos OS – Multicast Protocols User Guide at 90 (available at https://www.juniper.net/documentation/en_US/junos/information-products/pathway-pages/config-guide-multicast/config-guide-multicast.pdf).

Claim Chart for U.S. Patent No. 8,014,394 | Claim 1 | Juniper

Claim Element	Juniper Networking Equipment (<i>e.g.</i> , Juniper Networks MX240 Universal Routing Platform)
	How IGMP Snooping Works Devices usually learn unicast MAC addresses by checking the source address field of the frames they receive and then send any traffic for that unicast address only to the appropriate interfaces. However, a multicast MAC address can never be the source address for a packet. As a result, when a device receives traffic for a multicast destination address, it floods the traffic on the relevant VLAN, sending a significant amount of traffic for which there might not necessarily be interested receivers. IGMP snooping prevents this flooding. When you enable IGMP snooping, the device monitors IGMP packets between receivers and multicast routers and uses the content of the packets to build a multicast forwarding table—a database of multicast groups and the interfaces that are connected to members of the groups. When the device receives multicast packets, it uses the multicast forwarding table to selectively forward the traffic to only the interfaces that are connected to members of the appropriate multicast groups. <i>See e.g., id.</i> at 97. Support for IGMPv3 Multicast Sources In IGMPv3, a host can send a membership report that includes a list of source addresses. When the host sends a membership report in INCLUDE mode, the host is interested in group multicast traffic only from those sources in the source address list. If host sends a membership report in EXCLUDE mode, the host is interested in group multicast traffic from any source <i>except</i> the sources in the source address list. A host can also send an EXCLUDE report in which the source-list parameter is empty, which is known as an EXCLUDE NULL report. An EXCLUDE NULL report indicates that the host wants to join the multicast group and receive packets from all sources. Devices that support IGMPv3 process INCLUDE and EXCLUDE membership reports, and most devices forward source-specific multicast (SSM) traffic only from requested sources to subscribed receivers accordingly. However, you might see the device doesn't strictly forward multicast traffic on a per-source basis in some configurations such as: EX Series and QFX Series switches that do not use the Enhanced Layer 2 Software (ELS) configuration style <i>See e.g., id.</i> at 99.

Claim Chart for U.S. Patent No. 8,014,394 | Claim 1 | Juniper

Claim Element	Juniper Networking Equipment (<i>e.g.</i> , Juniper Networks MX240 Universal Routing Platform)
	IGMP Snooping and Forwarding Interfaces To determine how to forward multicast traffic, the device with IGMP snooping enabled maintains information about the following interfaces in its multicast forwarding table: • Multicast-router interfaces—These interfaces lead toward multicast routers or IGMP queriers. • Group-member interfaces—These interfaces lead toward hosts that are members of multicast groups. The device learns about these interfaces by monitoring IGMP traffic. If an interface receives IGMP queries or Protocol Independent Multicast (PIM) updates, the device adds the interface to its multicast forwarding table as a multicast-router interface. If an interface receives membership reports for a multicast group, the device adds the interface to its multicast forwarding table as a group-member interface. Learned interface table entries age out after a time period. For example, if a learned multicast-router interface does not receive IGMP queries or PIM hellos within a certain interval, the device removes the entry for that interface from its multicast forwarding table. NOTE: For the device to learn multicast-router interfaces and group-member interfaces, the network must include an IGMP querier. This is often in a multicast router, but if there is no multicast router on the local network, you can configure the device itself to be an IGMP querier. You can statically configure an interface to be a multicast-router interface or a group-member interface. The device adds a static interface to its multicast forwarding table without having to learn about the interface, and the entry in the table is not subject to aging. A device can have a mix of statically configured and dynamically learned interfaces. <i>See e.g., id.</i> at 100.
selectively forwarding multicast packets associated with the given	The Juniper Networks MX240 Universal Routing Platform performs the step of selectively forwarding multicast packets associated with the given multicast packet stream by the processing units in accordance with the respective updated lists.

Claim Chart for U.S. Patent No. 8,014,394 | Claim 1 | Juniper

Claim Element	Juniper Networking Equipment (<i>e.g.</i> , Juniper Networks MX240 Universal Routing Platform)
multicast packet stream by the processing units in accordance with the respective updated lists,	How IGMP Snooping Works Devices usually learn unicast MAC addresses by checking the source address field of the frames they receive and then send any traffic for that unicast address only to the appropriate interfaces. However, a multicast MAC address can never be the source address for a packet. As a result, when a device receives traffic for a multicast destination address, it floods the traffic on the relevant VLAN, sending a significant amount of traffic for which there might not necessarily be interested receivers. IGMP snooping prevents this flooding. When you enable IGMP snooping, the device monitors IGMP packets between receivers and multicast routers and uses the content of the packets to build a multicast forwarding table—a database of multicast groups and the interfaces that are connected to members of the groups. When the device receives multicast packets, it uses the multicast forwarding table to selectively forward the traffic to only the interfaces that are connected to members of the appropriate multicast groups. <i>See e.g.</i>, Junos OS – Multicast Protocols User Guide at 97 (available at https://www.juniper.net/documentation/en_US/junos/information-products/pathway-pages/config-guide-multicast/config-guide-multicast.pdf).

Claim Chart for U.S. Patent No. 8,014,394 | Claim 1 | Juniper

Claim Element	Juniper Networking Equipment (<i>e.g.</i> , Juniper Networks MX240 Universal Routing Platform)
	Support for IGMPv3 Multicast Sources In IGMPv3, a host can send a membership report that includes a list of source addresses. When the host sends a membership report in INCLUDE mode, the host is interested in group multicast traffic only from those sources in the source address list. If host sends a membership report in EXCLUDE mode, the host is interested in group multicast traffic from any source <i>except</i> the sources in the source address list. A host can also send an EXCLUDE report in which the source-list parameter is empty, which is known as an EXCLUDE NULL report. An EXCLUDE NULL report indicates that the host wants to join the multicast group and receive packets from all sources. Devices that support IGMPv3 process INCLUDE and EXCLUDE membership reports, and most devices forward source-specific multicast (SSM) traffic only from requested sources to subscribed receivers accordingly. However, you might see the device doesn't strictly forward multicast traffic on a per-source basis in some configurations such as: • EX Series and QFX Series switches that do not use the Enhanced Layer 2 Software (ELS) configuration style <i>See e.g., id.</i> at 99.

Claim Chart for U.S. Patent No. 8,014,394 | Claim 1 | Juniper

Claim Element	Juniper Networking Equipment (<i>e.g.</i> , Juniper Networks MX240 Universal Routing Platform)
	IGMP Snooping and Forwarding Interfaces To determine how to forward multicast traffic, the device with IGMP snooping enabled maintains information about the following interfaces in its multicast forwarding table: • Multicast-router interfaces—These interfaces lead toward multicast routers or IGMP queriers. • Group-member interfaces—These interfaces lead toward hosts that are members of multicast groups. The device learns about these interfaces by monitoring IGMP traffic. If an interface receives IGMP queries or Protocol Independent Multicast (PIM) updates, the device adds the interface to its multicast forwarding table as a multicast-router interface. If an interface receives membership reports for a multicast group, the device adds the interface to its multicast forwarding table as a group-member interface. Learned interface table entries age out after a time period. For example, if a learned multicast-router interface does not receive IGMP queries or PIM hellos within a certain interval, the device removes the entry for that interface from its multicast forwarding table. NOTE: For the device to learn multicast-router interfaces and group-member interfaces, the network must include an IGMP querier. This is often in a multicast router, but if there is no multicast router on the local network, you can configure the device itself to be an IGMP querier. You can statically configure an interface to be a multicast-router interface or a group-member interface. The device adds a static interface to its multicast forwarding table without having to learn about the interface, and the entry in the table is not subject to aging. A device can have a mix of statically configured and dynamically learned interfaces. <i>See e.g., id.</i> at 100.

Claim Chart for U.S. Patent No. 8,014,394 | Claim 1 | Juniper

Claim Element	Juniper Networking Equipment (<i>e.g.</i> , Juniper Networks MX240 Universal Routing Platform)
	General Forwarding Rules An interface in a VLAN with IGMP snooping enabled receives multicast traffic and forwards it according to the following rules. IGMP traffic: • Forward IGMP general queries received on a multicast-router interface to all other interfaces in the VLAN. • Forward IGMP group-specific queries received on a multicast-router interface to only those interfaces in the VLAN that are members of the group. • Forward IGMP reports received on a host interface to multicast-router interfaces in the same VLAN, but not to the other host interfaces in the VLAN. Multicast traffic that is not IGMP traffic: • Flood multicast packets with a destination address of 233.252.0.0/24 to all other interfaces on the VLAN. • Forward unregistered multicast packets (packets for a group that has no current members) to all multicast-router interfaces in the VLAN. • Forward registered multicast packets to those host interfaces in the VLAN that are members of the multicast group and to all multicast-router interfaces in the VLAN. <i>See e.g., id.</i> at 101.
wherein operating the network element comprises aggregating two or more of the ports to form a high-capacity port, and	The Juniper Networks MX240 Universal Routing Platform performs the step of wherein operating the network element comprises aggregating two or more of the ports to form a high-capacity port.

Claim Chart for U.S. Patent No. 8,014,394 | Claim 1 | Juniper

Claim Element	Juniper Networking Equipment (<i>e.g.</i> , Juniper Networks MX240 Universal Routing Platform)
	You can group or bundle multiple Ethernet interfaces together to form a single link layer interface known as the aggregated Ethernet interface (aex) or a link aggregation group (LAG). The IEEE 802.3ad standard defines link aggregation of Ethernet interfaces and provides a method by which you can group or bundle multiple Ethernet interfaces. Bundling multiple interfaces together enables you to increase the supported bandwidth. The device treats the aggregated Ethernet interface or LAG as a single link instead of a combination of multiple links. <i>See e.g.</i>, Junos OS – Ethernet Interfaces User Guide for Routing Devices at 64 (available at https://www.juniper.net/documentation/us/en/software/junos/interfaces-ethernet/interfaces-ethernet.pdf).
wherein aggregating the ports comprises setting an upper limit on a number of the multicast packet streams that are permitted for forwarding over the high-capacity port, wherein the request indicates that the given multicast packet stream is to be forwarded over the high-capacity port, and wherein updating the list comprises modifying the list to indicate that the given multicast packet stream is permitted for forwarding over the high-capacity port only when	The Juniper Networks MX240 Universal Routing Platform performs the step of wherein aggregating the ports comprises setting an upper limit on a number of the multicast packet streams that are permitted for forwarding over the high-capacity port, wherein the request indicates that the given multicast packet stream is to be forwarded over the high-capacity port, and wherein updating the list comprises modifying the list to indicate that the given multicast packet stream is permitted for forwarding over the high-capacity port only when the number of the multicast packet streams that are permitted for forwarding over the high-capacity port, including the given multicast packet stream, does not exceed the upper limit.

Claim Chart for U.S. Patent No. 8,014,394 | Claim 1 | Juniper

Claim Element	Juniper Networking Equipment (<i>e.g.</i> , Juniper Networks MX240 Universal Routing Platform)
the number of the multicast packet streams that are permitted for forwarding over the high-capacity port, including the given multicast packet stream, does not exceed the upper limit.	Understanding Aggregated Ethernet Load Balancing The link aggregation feature is used to bundle several physical aggregated Ethernet interfaces to form one logical interface. One or more links are aggregated to form a virtual link or link aggregation group (LAG). The MAC client treats this virtual link as if it were a single link. Link aggregation increases bandwidth, provides graceful degradation as failure occurs, and increases availability. In addition to these benefits, an aggregated Ethernet bundle is enhanced to provide load-balancing capabilities that ensure that the link utilization among the member links of the aggregated Ethernet bundle are fully and efficiently utilized. The load-balancing feature allows a device to divide incoming and outgoing traffic along multiple paths or interfaces in order to reduce congestion in the network. Load balancing improves the utilization of various network paths and provides more effective network bandwidth. <i>See e.g.</i>, Junos OS – Ethernet Interfaces User Guide for Routing Devices at 102 (available at https://www.juniper.net/documentation/us/en/software/junos/interfaces-ethernet/interfaces-ethernet.pdf). To enable per-flow load balancing based on hash values, include the <code>per-flow</code> statement at the at the [edit interfaces aeX unit logical-unit-number forwarding-options load-balance-stateful] hierarchy level. By default, Junos OS uses a hashing method based only on the destination address to elect a forwarding next hop when multiple equal-cost paths are available. All Packet Forwarding Engine slots are assigned the same hash value by default. To configure the load-balancing algorithm to dynamically rebalance the LAG using existing parameters, include the <code>rebalance interval</code> statement at the [edit interfaces aeX unit logical-unit-number forwarding-options load-balance-stateful] hierarchy level. This parameter periodically load balances traffic by providing a synchronized rebalance switchover across all the ingress Packet Forwarding Engines (PFEs) over a rebalance interval. You can specify the interval as a value in the range of 1 through 1000 flows per minute. To configure the load type, include the <code>load-type (low medium high)</code> statement at the [edit interfaces aeX unit logical-unit-number forwarding-options load-balance-stateful] hierarchy level. <i>See e.g.</i>, <i>id.</i> at 106.

Claim Chart for U.S. Patent No. 8,014,394 | Claim 1 | Juniper

Claim Element	Juniper Networking Equipment (<i>e.g.</i> , Juniper Networks MX240 Universal Routing Platform)						
	Table 34: Types of Port Speed Configuration <table> <tr> <th>Port Speed Configuration Type</th><th>Description</th></tr> <tr> <td>PIC or MIC Level</td><td>You can configure all the ports in one PIC or MIC to operate at the same speed. For example, you can configure all the ports of a PIC that support port speed of 100 Gbps to operate at 100 Gbps speed. If you do not specify the speed by using the <code>pic-mode</code> statement, then the port operates in the default speed.</td></tr> <tr> <td>Port Level</td><td> You can configure each port to operate at a different speed and thus enable each port. When you configure the port speed at the port level, you have the flexibility of operating the ports of the line card at different supported speeds. When you change the speed of a specific port in a given PIC using the <code>speed</code> statement, then only the speed of that port is modified. All other ports in the PIC remain unaffected. For example, you can configure four 10-Gigabit Ethernet interfaces on port 0, one 40-Gigabit Ethernet interface on port 1, and one 100-Gigabit Ethernet interface on port 2. </td></tr> </table> <i>See e.g., id.</i> at 282.	Port Speed Configuration Type	Description	PIC or MIC Level	You can configure all the ports in one PIC or MIC to operate at the same speed. For example, you can configure all the ports of a PIC that support port speed of 100 Gbps to operate at 100 Gbps speed. If you do not specify the speed by using the <code>pic-mode</code> statement, then the port operates in the default speed.	Port Level	You can configure each port to operate at a different speed and thus enable each port. When you configure the port speed at the port level, you have the flexibility of operating the ports of the line card at different supported speeds. When you change the speed of a specific port in a given PIC using the <code>speed</code> statement, then only the speed of that port is modified. All other ports in the PIC remain unaffected. For example, you can configure four 10-Gigabit Ethernet interfaces on port 0, one 40-Gigabit Ethernet interface on port 1, and one 100-Gigabit Ethernet interface on port 2.
Port Speed Configuration Type	Description						
PIC or MIC Level	You can configure all the ports in one PIC or MIC to operate at the same speed. For example, you can configure all the ports of a PIC that support port speed of 100 Gbps to operate at 100 Gbps speed. If you do not specify the speed by using the <code>pic-mode</code> statement, then the port operates in the default speed.						
Port Level	You can configure each port to operate at a different speed and thus enable each port. When you configure the port speed at the port level, you have the flexibility of operating the ports of the line card at different supported speeds. When you change the speed of a specific port in a given PIC using the <code>speed</code> statement, then only the speed of that port is modified. All other ports in the PIC remain unaffected. For example, you can configure four 10-Gigabit Ethernet interfaces on port 0, one 40-Gigabit Ethernet interface on port 1, and one 100-Gigabit Ethernet interface on port 2.						

Claim Chart for U.S. Patent No. 8,014,394 | Claim 1 | Juniper

Claim Element	Juniper Networking Equipment (e.g., Juniper Networks MX240 Universal Routing Platform)																	
	Table 65: Port Speed for MPC7E-MRATE <table><tr><th>PIC</th><th>Port</th><th>Port Speed Supported</th></tr><tr><td rowspan="3">PIC 0 (Fixed-port PIC)</td><td>0, 1, 3, 4</td><td>40-Gigabit Ethernet</td></tr><tr><td>2 and 5</td><td>100-Gigabit Ethernet</td></tr><tr><td></td><td>Default: 10-Gigabit Ethernet</td></tr><tr><td rowspan="3">PIC 1 (Fixed-port PIC)</td><td>0, 1, 3,4</td><td>40-Gigabit Ethernet</td></tr><tr><td>2 and 5</td><td>100-Gigabit Ethernet</td></tr><tr><td></td><td>Default: 10-Gigabit Ethernet</td></tr></table> See e.g., id. at 319. Configure the operating speed of all ports on the MPC7E-MRATE MPC, MIC-MRATE MIC, MX10003 MPC, and MX204 routers. (MX240, MX480, MX960, MX2010, and MX2020 routers with MPC7E-MRATE) To configure 100 Gbps, 10 Gbps, and 40 Gbps speed on all supported ports, specify 100G, 10G, or 40G, respectively, as the speed for the specified PIC. All the six ports of PIC 0 and PIC 1 of an MPC7E-MRATE MPC support 10-Gbps and 40-Gbps speeds. However, only ports 2 and 5 of PIC 0 and PIC 1 support 100-Gbps speed. Therefore, if you configure 100G as the operating speed of a PIC, the PIC is rebooted and the ports 0, 1, 3, and 4 are disabled. See e.g., id. at 1028.	PIC	Port	Port Speed Supported	PIC 0 (Fixed-port PIC)	0, 1, 3, 4	40-Gigabit Ethernet	2 and 5	100-Gigabit Ethernet		Default: 10-Gigabit Ethernet	PIC 1 (Fixed-port PIC)	0, 1, 3,4	40-Gigabit Ethernet	2 and 5	100-Gigabit Ethernet		Default: 10-Gigabit Ethernet
PIC	Port	Port Speed Supported																
PIC 0 (Fixed-port PIC)	0, 1, 3, 4	40-Gigabit Ethernet																
	2 and 5	100-Gigabit Ethernet																
		Default: 10-Gigabit Ethernet																
PIC 1 (Fixed-port PIC)	0, 1, 3,4	40-Gigabit Ethernet																
	2 and 5	100-Gigabit Ethernet																
		Default: 10-Gigabit Ethernet																

Claim Chart for U.S. Patent No. 8,014,394 | Claim 1 | Juniper

Claim Element	Juniper Networking Equipment (<i>e.g.</i> , Juniper Networks MX240 Universal Routing Platform)
	How IGMP Snooping Works Devices usually learn unicast MAC addresses by checking the source address field of the frames they receive and then send any traffic for that unicast address only to the appropriate interfaces. However, a multicast MAC address can never be the source address for a packet. As a result, when a device receives traffic for a multicast destination address, it floods the traffic on the relevant VLAN, sending a significant amount of traffic for which there might not necessarily be interested receivers. IGMP snooping prevents this flooding. When you enable IGMP snooping, the device monitors IGMP packets between receivers and multicast routers and uses the content of the packets to build a multicast forwarding table—a database of multicast groups and the interfaces that are connected to members of the groups. When the device receives multicast packets, it uses the multicast forwarding table to selectively forward the traffic to only the interfaces that are connected to members of the appropriate multicast groups. <i>See e.g.</i>, Junos OS – Multicast Protocols User Guide at 97 (available at https://www.juniper.net/documentation/en_US/junos/information-products/pathway-pages/config-guide-multicast/config-guide-multicast.pdf).

Claim Chart for U.S. Patent No. 8,014,394 | Claim 1 | Juniper

Claim Element	Juniper Networking Equipment (<i>e.g.</i> , Juniper Networks MX240 Universal Routing Platform)
	Support for IGMPv3 Multicast Sources In IGMPv3, a host can send a membership report that includes a list of source addresses. When the host sends a membership report in INCLUDE mode, the host is interested in group multicast traffic only from those sources in the source address list. If host sends a membership report in EXCLUDE mode, the host is interested in group multicast traffic from any source <i>except</i> the sources in the source address list. A host can also send an EXCLUDE report in which the source-list parameter is empty, which is known as an EXCLUDE NULL report. An EXCLUDE NULL report indicates that the host wants to join the multicast group and receive packets from all sources. Devices that support IGMPv3 process INCLUDE and EXCLUDE membership reports, and most devices forward source-specific multicast (SSM) traffic only from requested sources to subscribed receivers accordingly. However, you might see the device doesn't strictly forward multicast traffic on a per-source basis in some configurations such as: • EX Series and QFX Series switches that do not use the Enhanced Layer 2 Software (ELS) configuration style <i>See e.g., id.</i> at 99.

Claim Chart for U.S. Patent No. 8,014,394 | Claim 1 | Juniper

Claim Element	Juniper Networking Equipment (<i>e.g.</i> , Juniper Networks MX240 Universal Routing Platform)
	IGMP Snooping and Forwarding Interfaces To determine how to forward multicast traffic, the device with IGMP snooping enabled maintains information about the following interfaces in its multicast forwarding table: • Multicast-router interfaces—These interfaces lead toward multicast routers or IGMP queriers. • Group-member interfaces—These interfaces lead toward hosts that are members of multicast groups. The device learns about these interfaces by monitoring IGMP traffic. If an interface receives IGMP queries or Protocol Independent Multicast (PIM) updates, the device adds the interface to its multicast forwarding table as a multicast-router interface. If an interface receives membership reports for a multicast group, the device adds the interface to its multicast forwarding table as a group-member interface. Learned interface table entries age out after a time period. For example, if a learned multicast-router interface does not receive IGMP queries or PIM hellos within a certain interval, the device removes the entry for that interface from its multicast forwarding table. NOTE: For the device to learn multicast-router interfaces and group-member interfaces, the network must include an IGMP querier. This is often in a multicast router, but if there is no multicast router on the local network, you can configure the device itself to be an IGMP querier. You can statically configure an interface to be a multicast-router interface or a group-member interface. The device adds a static interface to its multicast forwarding table without having to learn about the interface, and the entry in the table is not subject to aging. A device can have a mix of statically configured and dynamically learned interfaces. <i>See e.g., id.</i> at 100.

Claim Chart for U.S. Patent No. 8,014,394 | Claim 1 | Juniper

Claim Element	Juniper Networking Equipment (<i>e.g.</i> , Juniper Networks MX240 Universal Routing Platform)
	General Forwarding Rules An interface in a VLAN with IGMP snooping enabled receives multicast traffic and forwards it according to the following rules. IGMP traffic: • Forward IGMP general queries received on a multicast-router interface to all other interfaces in the VLAN. • Forward IGMP group-specific queries received on a multicast-router interface to only those interfaces in the VLAN that are members of the group. • Forward IGMP reports received on a host interface to multicast-router interfaces in the same VLAN, but not to the other host interfaces in the VLAN. Multicast traffic that is not IGMP traffic: • Flood multicast packets with a destination address of 233.252.0.0/24 to all other interfaces on the VLAN. • Forward unregistered multicast packets (packets for a group that has no current members) to all multicast-router interfaces in the VLAN. • Forward registered multicast packets to those host interfaces in the VLAN that are members of the multicast group and to all multicast-router interfaces in the VLAN. <i>See e.g., id.</i> at 101.