(19) **United States**

(12) **Patent Application Publication** (10) Pub. No.: **US 2014/0177439 A1**

HOLNESS et al. (43) **Pub. Date:** **Jun. 26, 2014**

---

(54) **DYNAMIC PACKET TRAFFIC PERFORMANCE ADJUSTMENT SYSTEMS AND METHODS**

(71) Applicants: **Marc HOLNESS**, Nepean (CA); **David Wright MARTIN**, Stittsville (CA)

(72) Inventors: **Marc HOLNESS**, Nepean (CA); **David Wright MARTIN**, Stittsville (CA)

(57) **ABSTRACT**

A method, a network, and a network element use dynamic packet traffic performance adjustment techniques. In an exemplary embodiment, the dynamic resizing techniques utilize different packet connections providing connectivity to same sites between which bandwidth resizing is needed. Each of the packet connections has a separate and independent bandwidth profile that governs an amount of traffic that is dispatched over each packet connection. A network element sourcing traffic into the packet connections uses bridge functionality that dispatches client traffic onto all of the packet connections or an individual packet connection. This effectively means that the transport network bandwidth utilization is only consumed by a single packet connection, i.e., the packet connection-A (even through there are multiple configured). The network element sinking the traffic selects from a single active packet connection.

100

130

NETWORK

110a

110b

120   140   150

FIG. 1

FIG. 2

110a

330

BW Profile1 (Policer)

> CIR1

C

> EIR1

E

310

Client Ingress Traffic

320

Egress Port Queue

340

BW Profile2 (Policer)

> CIR2

C

> EIR2

E

FIG. 3

110a

330

BW Profile (Policer)

> CIR1

C

> EIR1

E

310

Client Ingress Traffic

320

Egress Port Queue

> CIR2

C

> EIR2

E

FIG. 4

500

PRECONFIGURE TWO PACKET
CONNECTIONS
502

OPERATE THE TWO PACKET
CONNECTIONS AT STEADY STATE
504

SOURCE NE PERFORMS A BRIDGE
FUNCTION DISPATCHING CLIENT
TRAFFIC TO BOTH PACKET
CONNECTIONS
506

SINK NE PERFORMS A SELECTOR
FUNCTION TO SELECT FROM THE
FIRST PACKET CONNECTION
508

CHANGE
PERFORMANCE
?
510

NO

YES

CHANGE THE SECOND PACKET
CONNECTION TO THE DESIRED BW
512

CHANGE SELECTOR FUNCTION TO
SELECT FROM THE SECOND PACKET
CONNECTION
514

CHANGE BW PROFILE OF THE FIRST
PACKET CONNECTION TO ZERO
516

FIG. 5

600

PRECONFIGURE TWO PACKET
CONNECTIONS
602

OPERATE THE TWO PACKET
CONNECTIONS AT STEADY STATE
604

SOURCE NE PERFORMS A BRIDGE
FUNCTION DISPATCHING CLIENT
TRAFFIC TO A SINGLE PACKET
CONNECTION
606

SINK NE PERFORMS A MERGE
FUNCTION TO EXTRACT FROM THE
PACKET CONNECTIONS
608

CHANGE
PERFORMANCE
?
610

NO

YES

CHANGE THE BRIDGE FUNCTION TO
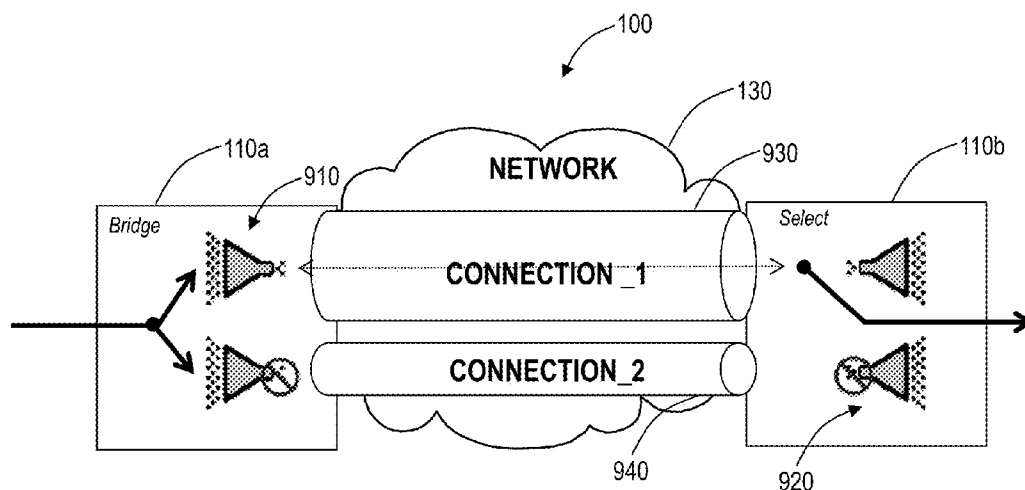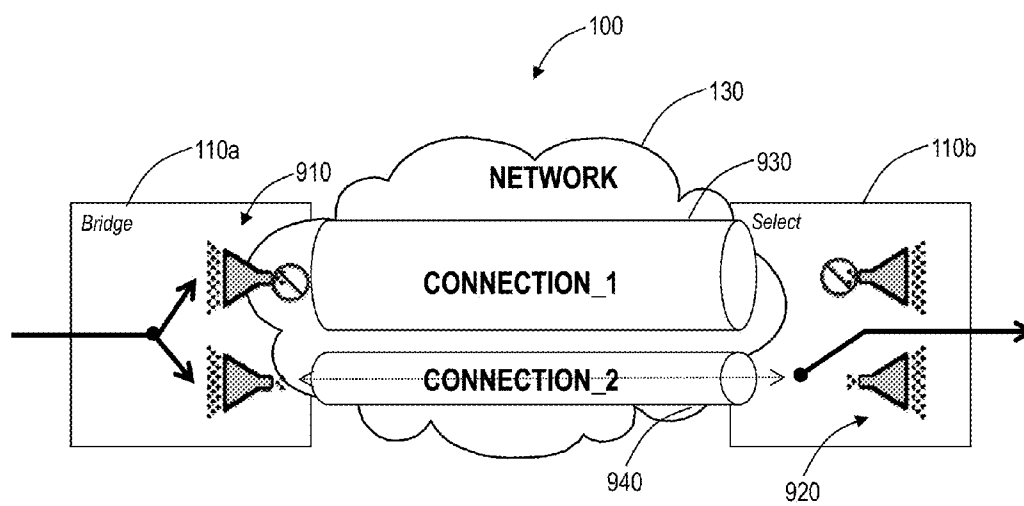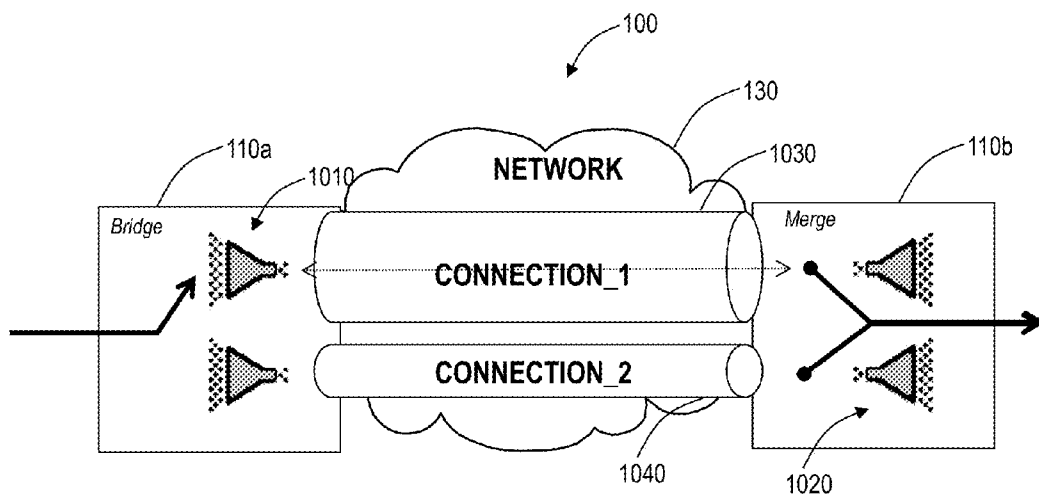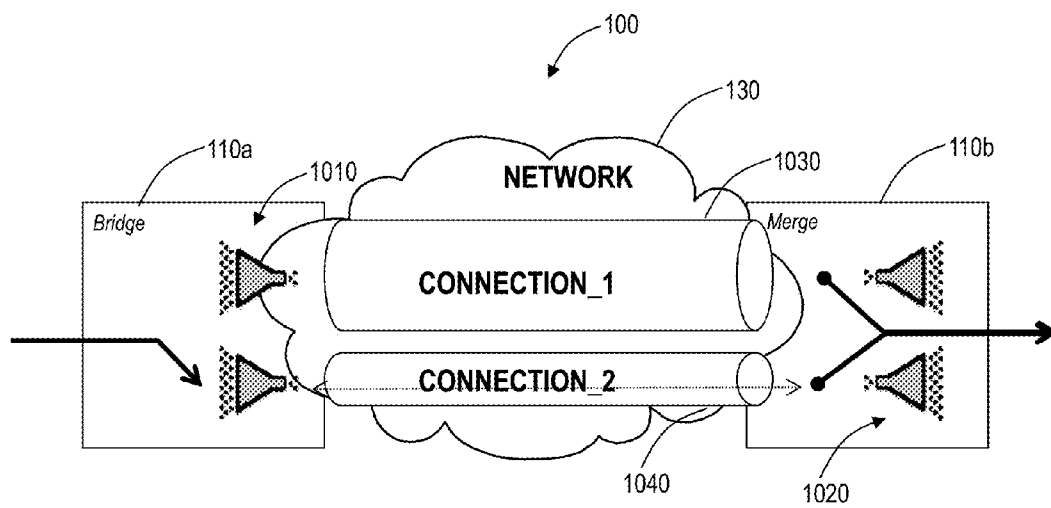THE SECOND PACKET CONNECTION
612

FIG. 6

FIG. 7

FIG. 8

FIG. 9

FIG. 10

FIG. 11

FIG. 12

DYNAMIC PACKET TRAFFIC PERFORMANCE ADJUSTMENT SYSTEMS AND METHODS

FIELD OF THE INVENTION

[0001] Generally, the field of art of the present disclosure pertains to networking systems and methods, and more particularly, to dynamic packet traffic performance adjustment systems and methods.

BACKGROUND OF THE INVENTION

[0002] Emerging network applications (e.g., Cloud services) require the need to dynamically adjust packet traffic performance (e.g., "resize" bandwidth) connections between network element end-points without impacting (or minimizing) the client traffic being transported over the connection. Given state of the art network element capabilities and internal components (e.g., hardware components such as Application Specific Integrated Circuits (ASICs)), there are currently two primary techniques to attempt to address the dynamic adjustment of packet traffic performance problem. First, network elements (NEs) can provide a capability allowing a user to change a bandwidth profile in-service. Current realizable implementations require the existing bandwidth profile (e.g., BW-Profile1), associated with a packet connection to be "re-sized", to be removed/deleted (thus impacting in-service traffic), and then a new First profile (e.g., BW-Profile2) to be added. This can be done, but the time taken to change the bandwidth profile affects traffic, given current implementation considerations. Second, NEs can include internal components that can support a sophisticated dynamic real-time bandwidth profile adjustment, while minimizing or nullifying the amount of in-service client traffic loss. In general, a typical switching device (especially those supported via current ASICs) is challenged to support such a sophisticated dynamic real-time bandwidth profile adjustment mechanism.

[0003] Thus, conventional techniques either result in a larger service traffic disruption, for NEs that require a bandwidth profile to be removed from a packet connection and then create a new BW profile and re-add it to the packet connection or require sophisticated dynamic bandwidth profile adjustment mechanisms, which are not typically supported in generally available ASICs. Consequently, a mechanism to adjust a packet connection bandwidth profile is required, without impacting or at least minimizing the in-service traffic impact (i.e., client traffic loss) is needed.

BRIEF SUMMARY OF THE INVENTION

[0004] In an exemplary embodiment, a method includes operating a first packet connection having a first bandwidth profile and a second packet connection having a second bandwidth profile, wherein the second bandwidth profile is set to zero; bridging ingress traffic at a source network element onto both the first packet connection and the second packet connection; selecting egress traffic at a sink network element from the first packet connection; responsive to a requirement for a bandwidth change in the first packet connection, modifying the second bandwidth profile accordingly at the first network element; selecting egress traffic at the sink network element from the second packet connection; and changing the first bandwidth profile to zero. The first packet connection and the second packet connection each can include a packet connection with the first bandwidth profile and the second bandwidth profile being separate and different. The first bandwidth profile can include a first Committed Information Rate and a first Excess Information Rate; and the second bandwidth profile can include a second Committed Information Rate and a second Excess Information Rate.

[0005] Prior to the requirement for the bandwidth change, the second Committed Information Rate and a second Excess Information Rate can be zero; and, subsequent to the modifying the second bandwidth profile, the second Committed Information Rate and a second Excess Information Rate can be set at values based on the requirement for the bandwidth change. The method can further include switching a selector at the sink network element at a time, t, subsequent to modifying the second bandwidth profile; wherein impact to client traffic from the source network element is limited to the time t.

[0006] The method can further include switching operating N packet connections each having an associated bandwidth profile, wherein the first packet connection includes one of the N packet connections; responsive to a requirement for a traffic performance change in the active packet connection, modifying the second bandwidth profile accordingly at the first network element; selecting egress traffic at the sink network element to the second packet connection for the any of the N packet connections; and changing a bandwidth profile to zero of the first packet connection.

[0007] The method can further include operating a first packet connection having a first bandwidth profile and a second packet connection having a second bandwidth profile; bridging ingress traffic at a source network element onto the first packet connection only; merging egress traffic at a sink network element from both the first and second packet connection; responsive to a requirement for a bandwidth change in the first packet connection, switching the bridge at the source network element to the second connection. The first packet connection and the second packet connection each can include a packet connection with the first bandwidth profile and the second bandwidth profile being separate and different. The first bandwidth profile can include a first Committed Information Rate and a first Excess Information Rate; and the second bandwidth profile can include a second Committed Information Rate and a second Excess Information Rate.

[0008] In another exemplary embodiment, a network includes a first packet network element; a second packet network element communicatively coupled to the first packet network element; a first packet connection between the first packet network element and the second packet network element having a first bandwidth profile being a first value; a second packet connection between the first packet network element and the second packet network element having a second bandwidth profile being set to zero; a bridge at the first packet network element configured to bridge ingress traffic onto both the first packet connection and the second packet connection; and a selector at the second packet network element configured to select traffic from one of the first packet connection and the second packet connection; wherein, responsive to a requirement to resize bandwidth of the first packet connection, the second bandwidth profile is set accordingly, the selector switches to the second packet connection, and the first bandwidth profile is set to zero.

[0009] The first packet connection and the second packet connection each can include a packet connection with the first bandwidth profile and the second bandwidth profile being

separate and different. The first bandwidth profile can include a first Committed Information Rate and a first Excess Information Rate; and wherein the second bandwidth profile can include a second Committed Information Rate and a second Excess Information Rate. Prior to the requirement to resize bandwidth, the second Committed Information Rate and a second Excess Information Rate can be zero; and, subsequent to the setting the second bandwidth profile, the second Committed Information Rate and a second Excess Information Rate can be set at values based on the requirement to resize bandwidth. Responsive to a requirement to resize bandwidth of the first packet connection, the selector can be configured to switch at a time, t, subsequent to the setting of the second bandwidth profile; wherein impact to client traffic from the source network element is limited to the time t.

[0010] In yet another exemplary embodiment, a network element includes at least one port; a packet switching fabric coupled to the at least one port; and a controller including instructions that, when executed, cause the controller, the at least one port, and the packet switching fabric to: operate a first packet connection having a first bandwidth profile and a second packet connection having a second bandwidth profile, wherein the second bandwidth profile is set to zero; bridge ingress traffic at the at least one port onto both the first packet connection and the second packet connection; responsive to a requirement for a bandwidth change in the first packet connection, modify the second bandwidth profile accordingly; and change the first bandwidth profile to zero subsequent to modifying the second bandwidth profile. The first packet connection and the second packet connection each can include a packet connection with the first bandwidth profile and the second bandwidth profile being separate and different. The first bandwidth profile can include a first Committed Information Rate and a first Excess Information Rate; wherein the second bandwidth profile includes a second Committed Information Rate and a second Excess Information Rate; wherein prior to the requirement for the bandwidth change, the second Committed Information Rate and a second Excess Information Rate are zero; and, subsequent to the modifying the second bandwidth profile, the second Committed Information Rate and a second Excess Information Rate are set at values based on the requirement for the bandwidth change.

[0011] The network element can include a source network element; wherein a second network element can include a sink network element for the first packet connection and the second packet connection; wherein the sink network element is configured to switch a selector from the first network element to the second network element at a time, t, subsequent to the modifying the second bandwidth profile; and wherein impact to client traffic from the source network element is limited to the time t. The instructions that, when executed, can further cause the controller, the at least one port, and the packet switching fabric to: operate N packet connections each having an associated bandwidth profile, wherein the first packet connection includes one of the N packet connections; responsive to a requirement for a bandwidth change in any of the N packet connections, modify the second bandwidth profile accordingly; and change a bandwidth profile to zero for the first packet connection.

BRIEF DESCRIPTION OF THE DRAWING(S)

[0012] Exemplary and non-limiting embodiments of the present disclosure are illustrated and described herein with reference to various drawings, in which like reference numbers denote like method steps and/or system components, respectively, and in which:

[0013] FIG. 1 is a network diagram of a network showing two interconnected packet switching network elements;

[0014] FIG. 2 is a block diagram of an exemplary implementation of a network element such as the packet switching network elements in the network of FIG. 1;

[0015] FIG. 3 is a block diagram of a first conventional technique for resizing bandwidth at a network element;

[0016] FIG. 4 is a block diagram of a second conventional technique for resizing bandwidth at a network element;

[0017] FIG. 5 is a flowchart of a dynamic packet traffic performance adjustment method;

[0018] FIG. 6 is a flowchart of an alternate dynamic packet traffic performance adjustment method;

[0019] FIG. 7 is a block diagram of a network element implementing the packet traffic performance adjustment method of FIG. 5;

[0020] FIG. 8 is a block diagram of a network element implementing the packet traffic performance adjustment method of FIG. 6;

[0021] FIGS. 9 and 10 are network diagrams of the network of FIG. 1 implementing the dynamic packet traffic performance adjustment method of FIG. 5; and

[0022] FIGS. 11 and 12 are network diagrams of the network of FIG. 1 implementing the dynamic packet traffic performance adjustment method of FIG. 6.

DETAILED DESCRIPTION OF THE INVENTION

[0023] In various exemplary embodiments, dynamic packet traffic performance adjustment systems and methods are described using dynamic resizing techniques. In an exemplary embodiment, the dynamic resizing techniques utilize different packet connections providing connectivity to same sites between which bandwidth resizing is needed. Each of the packet connections has a separate and independent bandwidth profile that governs an amount of traffic that is dispatched over each packet connection. A network element (NE) sourcing traffic into the packet connections uses bridge functionality that dispatches client traffic onto all of the packet connections. Only one of the packet connections will have a non-zero bandwidth profile (which governs the amount of traffic that is dispatched to the packet connection), i.e., a packet connection-A. This effectively means that the transport network bandwidth utilization is only consumed by a single packet connection, i.e., the packet connection-A (even through there are multiple configured). The network element sinking the traffic selects from the single packet connection. Management co-ordination is required so that the sink network element selects from the packet connection-A actively sending traffic.

[0024] When "re-sizing" of the bandwidth connection is required, the bandwidth profile of the "active" packet connection-A is moved from non-zero-A to zero (non-zero-A being the bandwidth amount of the packet connection-A), an alternate bandwidth profile packet connection-B is moved from zero to non-zero-B, and the "selector" function at the sink NE is moved from packet connection-A to packet connection-B. This changes the bandwidth profile from non-zero-A to non-zero-B in a minimally service-affecting manner. The actual disruption in client traffic being transported between sites is governed by the time it takes to change the selector function

at the sink network element, which is minimal, and can be easily realizable, by the most basic of packet switching network elements.

[0025] Beneficially, the dynamic resizing techniques support dynamic resizing of (a connection) bandwidth, while retaining the minimization of client service traffic loss/disruption, which can be readily realized by packet switching network elements. It is expected such functionality will be especially advantageous with respect to cloud services, data center interconnectivity, etc. which require rapid, dynamic resizing of bandwidth connections providing site connectivity, while minimizing the amount of disruption to the client traffic being transported. Packet technologies, Multiprotocol Label Switching (MPLS)-based, MPLS-Transport Profile (MPLS-TP)-based, Virtual Private Local Area Network Service (VPLS)-based, Internet Protocol (IP)-based, etc. are contemplated herewith.

[0026] Referring to FIG. 1, in an exemplary embodiment, a network diagram illustrates a network 100 showing two interconnected packet switching network elements 110a, 110b. The network elements 110a, 110b have a packet connection 120 there between over a network infrastructure 130 which can include optical equipment or the like forming physical links between the network elements 110a, 110b. The packet switching network elements 110a, 110b can include any type of network element such as, without limitation, a packet switch, an optical switch with packet switching capabilities, a service delivery switch, a service aggregation switch, a service module disposed within a larger system such as a wavelength division multiplexing (WDM) node or a microwave node, and the like.

[0027] The connection 120 can include a packet connection. For example, packet connections can be Ethernet Virtual Connections (EVCs) which are a service container described in the Metro Ethernet Forum's (MEF) technical specification 10.2, "Ethernet Services Attributes Phase 2," October 2009, the contents of which are incorporated by reference herein. Of course, the connection 120 can be other types of packet connections. In this exemplary embodiment of FIG. 1, the connection 120 can be at a first bandwidth 140 and it may be desired to resize to a second bandwidth 150. Note, the bandwidths 140, 150 can also be reversed, i.e., the dynamic resizing techniques described herein can support increases or decreases in bandwidth. Also, in an exemplary embodiment, the network elements 110a, 110b could be between data centers or the like in an application requiring the connection 120 to vary in size. Further, for this exemplary embodiment, the network element 110a is assumed to be a source for packet traffic and the network element 110b is assumed to be a sink for purposes of the following descriptions.

[0028] Referring to FIG. 2, in an exemplary embodiment, a block diagram illustrates an exemplary implementation of a network element 110 such as the packet switching network elements 110a, 110b in the network 100. In this exemplary embodiment, the network element 110 is a packet network switch for illustration purposes, but those of ordinary skill in the art will recognize the systems and methods described herein contemplate other types of network elements and other implementations providing dynamic traffic performance of packet connections. In this exemplary embodiment, the network element 110 includes a plurality of blades 202, 204 interconnected via an interface 206. The blades 202, 204 are also known as line cards, line modules, circuit packs, pluggable modules, etc. and refer generally to components

mounted within a chassis, shelf, etc. of a data switching device, i.e. the network element 110. Each of the blades 202, 204 may include numerous electronic devices and/or optical devices mounted on a circuit board along with various interconnects including interfaces to the chassis, shelf, etc.

[0029] Two exemplary blades are illustrated with line blades 202 and control blades 204. Of course, other types of blades are also contemplated. The line blades 202 generally include data ports 208 such as a plurality of packet ports. For example, the line blade 202 may include a plurality of physical ports disposed on an exterior of the blade 202 for receiving ingress/egress connections. Exemplary port types may include, without limitation, gigabit packet (GbE), 10GbE, 40GbE, 100GbE, Ethernet over SONET/SDH, Ethernet over Optical Transport Network, and the like. Additionally, the line blades 202 may include switching components to form a switching fabric via the interface 206 between all of the data ports 208 allowing data traffic to be switched between the data ports 208 on the various line blades 202. The switching fabric is a combination of hardware, software, firmware, etc. that moves data coming into the network element 110 out by the correct port 208 to the next network element. In general, the switching fabric may include switching units, or individual boxes, in a node; integrated circuits contained in the switching units; and programming that allows switching paths to be controlled. Note, the network element 110 is illustrated herein in a multi-blade configuration. Also, the network element 110 could be a single blade with line ports and the switching fabric contained therein. The single blade could also be part of another type of network element such as an optical switch, a WDM device, or the like.

[0030] The control blades 204 include a microprocessor 210, memory 212, software 214, and a network interface 216. Specifically, the microprocessor 210, the memory 212, and the software 214 may collectively control, configure, provision, monitor, etc. the network element 110. The network interface 216 may be utilized to communicate with a management system such as a Network Management System (NMS), Element Management System (EMS), and the like. Additionally, the control blades 204 may include a database 220 that tracks and maintains provisioning, configuration, operational data and the like. The database 220 may include a management information base (MIB) 222 which may include service related objects. Further, the control blades 204 may include a Simple Network Management Protocol (SNMP) Agent 224 configured to operate SNMPv2, SNMPv3, etc. or some other network management communication protocol. In this exemplary embodiment, the network element 110 includes two control blades 204 which may operate in a redundant or protected configuration such as 1:1, 1+1, etc. In general, the control blades 204 maintain dynamic system information including Layer two forwarding databases, protocol state machines, and the operational status of the ports 208 within the network element 110.

[0031] Referring to FIGS. 3-4, in conventional embodiments, the block diagrams illustrate a first conventional technique (FIG. 3) and a second conventional technique (FIG. 4) for resizing bandwidth at the network element 110a. The network element 110a has client ingress traffic 310 which is provided to an egress port queue 320. Note, other components are omitted for simplicity. Within the switching fabric of the network element 110a, there is a first bandwidth profile (BW profile 1) 330 for the connection 120. The profile 330 includes a first committed information rate (CIR1) and a first excess

information rate (EIR1) which provide the client ingress traffic **310** to the egress port queue **320** thereon.

[0032] In FIG. **3**, it is desired to change the profile **330** to a different, second bandwidth profile (BW profile **2**) **340**. The profile **340** includes a second committed information rate (CIR2) and a second excess information rate (EIR2) which provide the client ingress traffic **310** to the egress port queue **320** thereon. In this conventional embodiment, switching between the profiles **330**, **340** is possible, but the time required to change between the profiles **330**, **340** affects traffic. Specifically, in current realizable implementations, the profile **330** has to be removed (thus impacting in-service traffic) and then the profile **340** has to be added. Thus, conventional network element implementations require that client traffic is disrupted (i.e., lost) during the change from one bandwidth profile to the other associated with the packet connection being "re-sized".

[0033] In FIG. **4**, it is desired to change the profile **330** by changing the CIR1 and EIR1 rates to new, different CIR2 and EIR2 rates. In this exemplary embodiment, the network element **110***a* includes a sophisticated and dynamic real-time bandwidth adjustment mechanism that enables changing the profile **330** while minimizing an amount of in-service traffic loss. However, the network element **110***a* must include hardware, software, and/or firmware to implement this sophisticated and dynamic real-time bandwidth adjustment mechanism.

[0034] Referring to FIG. **5**, in an exemplary embodiment, a flowchart illustrates a dynamic bandwidth resizing method **500**. The dynamic bandwidth resizing method **500** contemplates operation in the network **100** between the network elements **110***a*, **110***b* as well as other networks and network elements. First, the dynamic bandwidth resizing method **500** includes preconfiguring two packet connections (step **502**). The first packet connection can have a first bandwidth profile (BW_Profile_1) and the second packet connection can have a second bandwidth profile (BW_Profile_2). The dynamic bandwidth resizing method **500** includes operating the two packet connections at steady state (step **504**). At this stage, the first bandwidth profile operates at <CIR, EIR> where CIR and EIR are values for the first bandwidth profile. The second bandwidth profile operates at <0, 0> where its CIR and EIR are each set to zero.

[0035] The dynamic bandwidth resizing method **500** includes a source network element performing a bridge function dispatching client traffic to both of the packet connections (step **506**). Conversely, the dynamic bandwidth resizing method **500** includes a sink network element performing a selector function to select from the first packet connection (step **508**). At this stage, the dynamic bandwidth resizing method **500** operates using the first packet connection at the first bandwidth profile. If there is a desired change in bandwidth (step **510**), then the dynamic bandwidth resizing method **500** includes changing the second packet connection to the desired bandwidth (step **512**). Specifically, the second bandwidth profile is changed to <CIR2, EIR2> where CIR2 and EIR2 are the new, different bandwidth amounts from the first bandwidth profile.

[0036] The dynamic bandwidth resizing method **500** includes changing the selector function at the second network element to select from the second packet connection (step **514**). Finally, the dynamic bandwidth resizing method **500** includes setting the bandwidth profile of the first packet connection to <0, 0> for the CIR and EIR. Note, the dynamic

bandwidth resizing method **500** includes a one for one correspondence between the two packet connections. Alternatively, there could be a 1:N relationship where the second packet connection is set to <0, 0> and waiting for any number, N, of working packet connections that want a corresponding bandwidth change. The trigger to switch to the second packet connection could be based on policy, on a schedule, on demand, etc.

[0037] Referring to FIG. **6**, in an exemplary embodiment, a flowchart illustrates a dynamic bandwidth resizing method **600**. The dynamic bandwidth resizing method **600** contemplates operation in the network **100** between the network elements **110***a*, **110***b* as well as other networks and network elements. First, the dynamic bandwidth resizing method **600** includes preconfiguring two packet connections (step **602**). The first packet connection can have a first bandwidth profile (BW_Profile_1) and the second packet connection can have a second bandwidth profile (BW_Profile_2). The dynamic bandwidth resizing method **600** includes operating the two packet connections at steady state (step **604**).

[0038] At this stage, the first bandwidth profile operates at <CIR, EIR> where CIR and EIR are values for the first bandwidth profile. The second bandwidth profile operates at <CIR, EIR> where its CIR and EIR are values for the second bandwidth profile.

[0039] The dynamic bandwidth resizing method **600** includes a source network element performing a bridge function dispatching client traffic to a single of the packet connections (step **606**). Conversely, the dynamic bandwidth resizing method **600** includes a sink network element performing a merge function to select from either of the packet connections (step **608**). At this stage, the dynamic bandwidth resizing method **600** operates using the first packet connection at the first bandwidth profile. If there is a desired change in bandwidth (step **610**), then the dynamic bandwidth resizing method **600** includes changing the bridge function to the second packet connection (step **612**).

[0040] The dynamic bandwidth resizing method **600** includes changing the bridge function at the first network element to dispatch to the packet connection with the desired traffic performance (e.g., bandwidth resizing). Note, the dynamic bandwidth resizing method **600** includes a one for one correspondence between the two packet connections. Alternatively, there could be a 1:N relationship where the second packet connection is selected by the source network element. The trigger to switch to the second packet connection could be based on policy, on a schedule, on demand, etc.

[0041] Referring to FIG. **7**, in an exemplary embodiment, a block diagram illustrates the network element **110***a* implementing the dynamic bandwidth resizing method **500**. Again, the network element **110***a* includes client ingress traffic **710** and an egress port queue **720**. The network element **110***a* is operating two packet connections **730**, **740** which forward traffic from the client ingress traffic **710** to the egress port queue **720**. The packet connection **730** has a bandwidth profile <CIR1, EIR1> and the packet connection **740** has a bandwidth profile <CIR2, EIR2>. At first, <CIR1, EIR1> equals some non-zero value for the packet connection **730** and <CIR2, EIR2>=<0, 0> for the second packet connection **740**. Upon a request to change the bandwidth profile of the first packet connection **730**, the second packet connection **740** is set to the desired value. In this manner, none of the client ingress traffic **710** is lost or dropped. At the second network

5

element **110***b*, the selector is switched after the bandwidth profile is set to the desired value for the packet connection **740**.

[0042] Referring to FIG. **8**, in an exemplary embodiment, a block diagram illustrates the network element **110***a* implementing the dynamic bandwidth resizing method **600**. Again, the network element **110***a* includes client ingress traffic **810** and an egress port queue **820**. The network element **110***a* is operating two packet connections **830**, **840** which forward traffic from the client ingress traffic **810** to the egress port queue **820**. The packet connection **830** has a bandwidth profile <CIR1, EIR1> and the packet connection **840** has a bandwidth profile <CIR2, EIR2>. At first, network element **110***a* will bridge client traffic to the first packet connection **830**. Upon a request to change the bandwidth profile of the packet connection **830**, the client traffic gets re-directed (i.e., bridged) to the second connection **840**. In this manner, none of the client ingress traffic **810** is lost or dropped. At the second network element **110***b*, the traffic is merged from both the first and second packet connections.

[0043] Referring to FIGS. **9** and **10**, in an exemplary embodiment, network diagrams illustrate the network **100** implementing the dynamic bandwidth resizing method **500**. In FIG. **9**, the two packet connections **930**, **940** are configured between the network elements **110***a*, **110***b*. At this stage, the packet connection **930** is operating at the first bandwidth profile with <CIR1, EIR1> equaling some non-zero value for the packet connection **930**. A bridging function **910** effectively sends all ingress client traffic to the packet connection **930** since the bandwidth profile of the packet connection **940** is <0, 0>. A selector function **920** at the network element **110***b* selects from the packet connection **930**. Thus, at steady state, network bandwidth resource utilization is only consumed by the packet connection **930**, since BW_Profile_2 is set to 0 (i.e., <CIR, EIR>=<0,022).

[0044] Note, as described herein, the bandwidth profiles are described for illustration purposes as <CIR, EIR>. However, a bandwidth profile, for a packet connection, can include <CIR, CBS, EIR, EBS, CF, CM> where CIR is Committed Information Rate, CBS is Committed Burst Size, EIR is Excess Information Rate, EBS is Excess Burst Size, CF is Coupling Flag, and CM is Color Mode. For this discussion only <CIR, EIR> are mentioned as they are the only parameters whose values would change. The expectation is that the assigned CoS and the per-CoS performance objectives (CPOs) <FLR, FD, MFD, FDR, IFDV, A, L, B> would not be changed. The bandwidth profile information for packet connections is described in detail in the MEF technical specification 10.2 which has been previously incorporated by reference herein.

[0045] In FIG. **10**, the dynamic bandwidth resizing method **500** has resized the packet connection **940** to the desired bandwidth profile for the change to the packet connection **930**. Again, the bridging function **910** effectively sends all ingress client traffic to the packet connection **940** since the bandwidth profile of the packet connection **930** is <0, 0> subsequent to the change in the packet connection **940**. The selector function **920** now selects from the packet connection **940**. When an application requires "dynamic traffic performance adjustment" of the packet connection **930** (i.e., dynamic adjustment of packet connection bandwidth requirements), then controlled actions occur: Change BW_Profile_2 of the packet connection **940** to desired <CIR, EIR>→Time=t1, Change selector at sink point to the packet

connection **940** Time=t2, and Change BW_Profile_1 of the packet connection **930** to 0 (i.e.,<CIR, EIR>=<0,0>) Time=t3. Thus, using the dynamic bandwidth resizing method **500**, Client traffic loss is effectively minimized to t2 and Network utilization overshoot is effectively t3. In general, time t2 is minimal, for current state of the art implementations.

[0046] Referring to FIGS. **11** and **12**, in an exemplary embodiment, network diagrams illustrate the network **100** implementing the dynamic bandwidth resizing method **600**. In FIG. **11**, the two packet connections **1030**, **1040** are configured between the network elements **110***a*, **110***b*. At this stage, the packet connection **1030** is operating at the first bandwidth profile with <CIR1, EIR1> and packet connection **1040** is operating at a second bandwidth profile with <CIR2, EIR2>. A bridging function **1010** sends all ingress client traffic to the packet connection **1030**. A merge function **1020** at the network element **110***b* selects from both the packet connections **1030** and **1040**. Thus, at steady state, network bandwidth resource utilization is only consumed by the packet connection **1030**, since the packet connection **1040** is operating at a second bandwidth profile with <CIR2, EIR2>=<0, 0>.

[0047] Note, as described herein, the bandwidth profiles are described for illustration purposes as <CIR, EIR>. However, a bandwidth profile, for a packet connection, can include <CIR, CBS, EIR, EBS, CF, CM> where CIR is Committed Information Rate, CBS is Committed Burst Size, EIR is Excess Information Rate, EBS is Excess Burst Size, CF is Coupling Flag, and CM is Color Mode. For this discussion only <CIR, EIR> are mentioned as they are the only parameters whose values would change. The expectation is that the assigned CoS and the per-CoS performance objectives (CPOs) <FLR, FD, MFD, FDR, IFDV, A, L, B> would not be changed. The bandwidth profile information for packet connections is described in detail in the MEF technical specification 10.2 which has been previously incorporated by reference herein.

[0048] In FIG. **12**, the dynamic traffic performance adjustment method **600** has resized the active packet connection **1030** to the desired packet connection **1040** performance. Again, the bridging function **1010** sends all ingress client traffic to the packet connection **1030**. The merger function **1020** will select/accept from both packet connections **1030** and **1040**. When an application requires "dynamic traffic performance adjustment" of the packet connection **1030** (i.e., dynamic adjustment of packet connection bandwidth requirements), then controlled actions occur: Network element **100***a* will have its bridge function changed to the packet connection **1040**. Thus, using the dynamic bandwidth resizing method **600**, Client traffic loss is effectively minimized to t2 and Network utilization overshoot is effectively t3. In general, time t2 is minimal, for current state of the art implementations.

[0049] It will be appreciated that some exemplary embodiments described herein may include one or more generic or specialized processors ("one or more processors") such as microprocessors, digital signal processors, customized processors, and field programmable gate arrays (FPGAs) and unique stored program instructions (including both software and firmware) that control the one or more processors to implement, in conjunction with certain non-processor circuits, some, most, or all of the functions of the methods and/or systems described herein. Alternatively, some or all functions

may be implemented by a state machine that has no stored program instructions, or in one or more application specific integrated circuits (ASICs), in which each function or some combinations of certain of the functions are implemented as custom logic. Of course, a combination of the aforementioned approaches may be used. Moreover, some exemplary embodiments may be implemented as a non-transitory computer-readable storage medium having computer readable code stored thereon for programming a computer, server, appliance, device, etc. each of which may include a processor to perform methods as described and claimed herein. Examples of such computer-readable storage mediums include, but are not limited to, a hard disk, an optical storage device, a magnetic storage device, a ROM (Read Only Memory), a PROM (Programmable Read Only Memory), an EPROM (Erasable Programmable Read Only Memory), an EEPROM (Electrically Erasable Programmable Read Only Memory), Flash memory, and the like. When stored in the non-transitory computer readable medium, software can include instructions executable by a processor that, in response to such execution, cause a processor or any other circuitry to perform a set of operations, steps, methods, processes, algorithms, etc.

[0050] Although the present disclosure has been illustrated and described herein with reference to preferred embodiments and specific examples thereof, it will be readily apparent to those of ordinary skill in the art that other embodiments and examples may perform similar functions and/or achieve like results. All such equivalent embodiments and examples are within the spirit and scope of the present disclosure and are intended to be covered by the following claims.

What is claimed is:

1. A method, comprising:
operating a first packet connection having a first bandwidth profile and a second packet connection having a second bandwidth profile, wherein the second bandwidth profile is set to consume zero bandwidth;
bridging ingress traffic at a source network element onto both the first packet connection and the second packet connection;
selecting egress traffic at a sink network element from the first packet connection;
responsive to a requirement for a bandwidth change in the first packet connection, modifying the second bandwidth profile accordingly at the first network element;
selecting egress traffic at the sink network element from the second packet connection; and
changing the first bandwidth profile to consume zero bandwidth.

2. The method of claim 1, wherein the first packet connection and the second packet connection each comprise a packet connection with the first bandwidth profile and the second bandwidth profile being separate and different.

3. The method of claim 2, wherein the first bandwidth profile comprises a first Committed Information Rate and a first Excess Information Rate;
wherein the second bandwidth profile comprises a second Committed Information Rate and a second Excess Information Rate.

4. The method of claim 3, wherein, prior to the requirement for the bandwidth change, the second Committed Information Rate and a second Excess Information Rate are zero; and
wherein, subsequent to the modifying the second bandwidth profile, the second Committed Information Rate

and a second Excess Information Rate are set at values based on the requirement for the bandwidth change.

5. The method of claim 1, further comprising:
switching a selector at the sink network element at a time, t, subsequent to the modifying the second bandwidth profile;
wherein impact to client traffic from the source network element is limited to the time t.

6. The method of claim 1, further comprising:
operating N packet connections each having an associated bandwidth profile, wherein the first packet connection comprises one of the N packet connections;
responsive to a requirement for a bandwidth change in any of the N packet connections, modifying the second bandwidth profile accordingly at the first network element;
selecting egress traffic at the sink network element from the second packet connection for the any of the N packet connections; and
changing a bandwidth profile to zero for the any of the N packet connections.

7. The method of claim 6, wherein the second packet connection is set to a zero bandwidth profile until one of the N packet connections requires a resizing of bandwidth.

8. A method, comprising:
operating a first packet connection having a first bandwidth profile and a second packet connection having a second bandwidth profile;
bridging ingress traffic at a source network element onto one of the packet connections;
selecting egress traffic at a sink network element from the set of packet connections; and
responsive to a requirement for a bandwidth change in the first packet connection, modifying the bridge function accordingly at the first network element.

9. The method of claim 8, wherein the first packet connection and the second packet connection each comprise a packet connection with the first bandwidth profile and the second bandwidth profile being separate and different.

10. The method of claim 9, wherein the first bandwidth profile comprises a first Committed Information Rate and a first Excess Information Rate;
wherein the second bandwidth profile comprises a second Committed Information Rate and a second Excess Information Rate.

11. The method of claim 8, further comprising:
switching the bridge at the source network element at a time, t, to the new active connection performance;
wherein impact to client traffic from the source network element is limited to the time t.

12. The method of claim 8, further comprising:
operating N packet connections each having an associated bandwidth profile, wherein the first packet connection comprises one of the N packet connections;
responsive to a requirement for a bandwidth change in any of the N packet connections, modifying the bridge selector to the second packet connection accordingly at the first network element;
selecting egress traffic at the sink network element from the set of N packet connections.

13. A network, comprising:
a first packet network element;
a second packet network element communicatively coupled to the first packet network element;

a first packet connection between the first packet network element and the second packet network element having a first bandwidth profile being a first value;

a second packet connection between the first packet network element and the second packet network element having a second bandwidth profile being set to zero;

a bridge at the first packet network element configured to bridge ingress traffic onto both the first packet connection and the second packet connection; and

a selector at the second packet network element configured to select traffic from one of the first packet connection and the second packet connection;

wherein, responsive to a requirement to resize bandwidth of the first packet connection, the second bandwidth profile is set accordingly, the selector switches to the second packet connection, and the first bandwidth profile is set to zero.

14. The network of claim **13**, wherein the first packet connection and the second packet connection each comprise a packet connection with the first bandwidth profile and the second bandwidth profile being separate and different.

15. The network of claim **14**, wherein the first bandwidth profile comprises a first Committed Information Rate and a first Excess Information Rate; and

wherein the second bandwidth profile comprises a second Committed Information Rate and a second Excess Information Rate.

16. The network of claim **15**, wherein, prior to the requirement to resize bandwidth, the second Committed Information Rate and a second Excess Information Rate are zero; and

Wherein, subsequent to the setting of the second bandwidth profile, the second Committed Information Rate and a second Excess Information Rate are set at values based on the requirement to resize bandwidth.

17. The network of claim **13**, wherein, responsive to a requirement to resize bandwidth of the first packet connection, the selector is configured to switch at a time, t, subsequent to the setting of the second bandwidth profile;

wherein impact to client traffic from the source network element is limited to the time t.

18. The network of claim **13**, further comprising:

N packet connections each having an associated bandwidth profile, wherein the first packet connection comprises one of the N packet connections;

wherein, responsive to a requirement to resize bandwidth in any of the N packet connections, the second bandwidth profile is set accordingly at the first network element;

wherein the second network element selects from the second packet connection for the any of the N packet connections; and

wherein the bandwidth profile is set to zero for the first packet connection.

19. A network, comprising:

a first packet network element;

a second packet network element communicatively coupled to the first packet network element;

a first packet connection between the first packet network element and the second packet network element having a first bandwidth profile being a first value;

a second packet connection between the first packet network element and the second packet network element having a second bandwidth profile being a second value;

a bridge at the first packet network element configured to bridge ingress traffic onto the first packet connection; and

a merge at the second packet network element configured to merge traffic from all of the packet connections;

wherein, responsive to a requirement to resize bandwidth of the first packet connection, the bridge function at the first packet network element is set to the second packet connection.

20. The network of claim **19**, wherein the first packet connection and the second packet connection each comprise a packet connection with the first bandwidth profile and the second bandwidth profile being separate and different.

21. The network of claim **20**, wherein the first bandwidth profile comprises a first Committed Information Rate and a first Excess Information Rate; and

wherein the second bandwidth profile comprises a second Committed Information Rate and a second Excess Information Rate.

22. The network of claim **21**, wherein, responsive to a requirement to resize bandwidth of the first packet connection, the bridge is configured to switch at a time, t;

wherein impact to client traffic from the source network element is limited to the time t.

23. The network of claim **19**, further comprising:

N packet connections each having an associated bandwidth profile, wherein the first packet connection comprises one of the N packet connections;

wherein, responsive to a requirement to resize bandwidth in any of the N packet connections, the bridge function at the first network element is set to the second packet connection;

wherein the second network element merges from the set of packet connections from all N packet connections.

24. A network element, comprising:

at least one port;

a packet switching fabric coupled to the at least one port; and

a controller comprising instructions that, when executed, cause the controller, the at least one port, and the packet switching fabric to:

operate a first packet connection having a first bandwidth profile and a second packet connection having a second bandwidth profile, wherein the second bandwidth profile is set to zero;

bridge ingress traffic at the at least one port onto both the first packet connection and the second packet connection;

responsive to a requirement for a bandwidth change in the first packet connection, modify the second bandwidth profile accordingly; and

change the first bandwidth profile to zero subsequent to modifying the second bandwidth profile.

25. The network element of claim **24**, wherein the first packet connection and the second packet connection each comprise a packet connection with the first bandwidth profile and the second bandwidth profile being separate and different.

26. The network element of claim **25**, wherein the first bandwidth profile comprises a first Committed Information Rate and a first Excess Information Rate;

wherein the second bandwidth profile comprises a second Committed Information Rate and a second Excess Information Rate;

wherein, prior to the requirement for the bandwidth change, the second Committed Information Rate and a second Excess Information Rate are zero; and

wherein, subsequent to the modifying the second bandwidth profile, the second Committed Information Rate and a second Excess Information Rate are set at values based on the requirement for the bandwidth change.

27. The network element of claim **24**, wherein the network element comprises a source network element;

wherein a second network element comprises a sink network element for the first packet connection and the second packet connection;

wherein the sink network element is configured to switch a selector from the first network element to the second network element at a time, t, subsequent to the modifying the second bandwidth profile; and

wherein impact to client traffic from the source network element is limited to the time t.

28. The network element of claim **24**, wherein the instructions that, when executed, further cause the controller, the at least one port, and the packet switching fabric to:

operate N packet connections each having an associated bandwidth profile, wherein the first packet connection comprises one of the N packet connections;

* * * * *