AMENDED CLAIMS

What is claimed is:

1. A method comprising:

determining, by a controller of a network, (1) that a device is performing handover from a first edge data center to a second edge data center based on a signaling message that the device has attached or has requested to attach to the second edge data center instead of the first edge data center, or (2) that the device is to perform handover from the first edge data center to the second edge data center based on telemetry data about the device;

in response to determining that the device is performing or is to perform handover from the first edge data center to the second edge data center, instructing, by the controller, the first edge data center to replicate or move an application workload hosted on the first edge data center to the second edge data center,

computing, by the controller a path mapped to a traffic class to route traffic between the device and the second edge data center hosting the application workload, the traffic being routed along the path to one or more modules hosted on the second edge data center,

wherein the path is computed based on:

(1) element registration information for each module of one or more modules hosted on a set of interconnected edge data centers including the first edge data center and the second edge data center,

(2) one or more network devices that interconnect the modules in the network, and

(3) one or more routing metrics of the network; and

sending, by the controller, an indication of the path to the second edge data center to cause the one or more modules hosted on the second edge data center to use the path mapped to the traffic class to route traffic to the application workload hosted on the second edge data center.

2. The method of claim 1, wherein each of the modules hosted on the set of interconnected edge data centers provides at least one of a radio access network (RAN)

function, a 5th Generation of Mobile Networks (5G) core function, and an edge computing function.

3. The method of any of claims 1–2, further comprising:

receiving the signaling message from the second edge data center.

4. The method of any of claims 1–2, further comprising:

receiving the telemetry data about the device from the device or a 5th Generation of Mobile Networks (5G) core function, wherein the telemetry data comprises at least one of a location of the device and direction of travel of the device.

5. The method of claim 1, further comprising:

receiving, from the one or more modules of the second edge data center, network insight data,

wherein determining that the device is to perform handover further comprises determining, based on the telemetry data from the device and the network insight data from the modules of the second edge data center, that the device is to perform handover.

6. The method of any of claims 1–5, wherein the routing metrics of the network comprises at least one of latency, bandwidth, hop count, link utilization, reliability, throughput, load, packet loss, and path parameters of an optical network, wherein the path parameters of the optical network comprises at least one of span latencies, capacities, and protection states.

7. The method of any of claims 1–6, wherein the set of interconnected edge data centers is configured as a domain of one or more edge computing functions.

8. The method of any of claims 1–7, further comprising:

assigning, by the controller and based on the element registration information, a segment identifier for each of the modules hosted on the second edge data center, and for each of the one or more network devices that interconnect the modules in the network; and

generating, by the controller and from the segment identifiers, one or more segment routing label stacks mapped to the traffic class to route traffic for the application workload between the device and the second edge data center according to the path mapped to the traffic class.

9. The method of claim 8, wherein the one or more segment routing label stacks comprises an uplink segment routing label stack, wherein sending the indication of the path comprises sending the uplink segment routing label stack to the second edge data center to cause a target centralized unit of the second edge data center to encapsulate traffic with the uplink segment routing label stack to route traffic for the application workload from the device to the second edge data center.

10. The method of claim 8, wherein the one or more segment routing label stacks comprises a downlink segment routing label stack, wherein sending the path further comprises sending the downlink segment routing label stack to the second edge data center to cause the second edge data center to encapsulate traffic with the downlink segment routing label stack to route traffic for the application workload from the second edge data center to the device.

11. A controller for a network, the controller comprising:

one or more processors operably coupled to memory, wherein the one or more processors are configured to:

determine (1) that a device is performing handover from a first edge data center to a second edge data center based on a signaling message that the device has attached or has requested to attach to the second edge data center instead of the first edge data center, or (2) that the device is to perform handover from the first edge data center to the second edge data center based on telemetry data about the device;

instruct, in response to determining that the device is performing or is to perform handover from the first edge data center to the second edge data center, the first edge data center to replicate or move an application workload hosted on the first edge data center to the second edge data center,

compute a path mapped to a traffic class to route traffic between the device and the second edge data center hosting the application workload, the traffic being routed along the path to one or more modules hosted on the second edge data center,

wherein the path is computed based on:

(1) element registration information for each module of one or more modules hosted on a set of interconnected edge data centers including the first edge data center and the second edge data center,

(2) one or more network devices that interconnect the modules in the network, and

(3) one or more routing metrics of the network; and

send an indication of the path to the second edge data center to cause the one or more modules hosted on the second edge data center to use the path mapped to the traffic class to route traffic to the application workload hosted on the second edge data center.

12. The controller of claim 11, wherein each of the modules hosted on the set of interconnected edge data centers provide at least one of a radio access network (RAN) function, a 5th Generation of Mobile Networks (5G) core function, and an edge computing function;

13. The controller of any of claims 11–12, wherein the one or more processors are further configured to:

receive the signaling message from the second edge data center.

14. The controller of any of claims 11–12, wherein the one or more processors are further configured to:

receive telemetry data about the device from the device or a 5th Generation of Mobile Networks (5G) core function.

15. The controller of claim 11, wherein the one or more processors are further configured to:

receive, from the one or more modules of the second edge data center, network insight data,

wherein to determine that the device is to perform handover, the one or more processors are further configured to determine, based on the telemetry data from the device and the network insight data from the modules of the second edge data center, that the device is to perform handover.

16. The controller of claim 14, wherein the one or more processors are further configured to:

assign, based on the element registration information, a segment identifier for each of the modules hosted on the second edge data center, and for each of the one or more network devices that interconnect the modules in the network; and

generate, from the segment identifiers, an uplink segment routing label stack and a downlink segment routing label stack each mapped to the traffic class to route traffic for the application workload between the device and the second edge data center according to the path mapped to the traffic class;

send the uplink segment routing label stack to the second edge data center to cause a target centralized unit of the second edge data center to encapsulate traffic with the uplink segment routing label stack to route traffic for the application workload from the device to the second edge data center; and

send the downlink segment routing label stack to the second edge data center to cause the second edge data center to encapsulate traffic with the downlink segment routing label stack to route traffic for the application workload from the second edge data center to the device.

17. A method comprising:

receiving, by a controller of a network, a first request to route traffic between a first device and an edge computing function hosting an application workload via a first edge data center, and a second request to route traffic between a second device and the edge computing function via a second edge data center;

wherein the first edge data center and the second edge data center are each provided by different mobile network providers,

wherein a third edge data center hosting the application workload is provided by a common application provider to the different mobile network providers;

computing, by the controller, a first path mapped to a traffic class to route traffic between the first device and the edge computing function hosting the application workload via one or more modules hosted on the first edge data center and a second path mapped to the traffic class to route traffic between the second device and the edge computing function hosting the application workload via one or more modules hosted on the second edge data center,

wherein the first path and second path are computed based on:

(1) element registration information for modules hosted on a set of interconnected edge data centers including the first edge data center, the second edge data center, and the third edge data center,

(2) one or more network devices that interconnect the one or more modules hosted on the first edge data center or the one or more modules hosted on the second edge data center in the network, and

(3) one or more routing metrics of the network;

sending, by the controller, an indication of the first path to the first edge data center to cause the one or more modules hosted on the first edge data center to use the first path mapped to the traffic class to route traffic to the edge computing function hosting the application workload; and

sending, by the controller, an indication of the second path to the second edge data center to cause the one or more modules hosted on the second edge data center to use the second path mapped to the traffic class to route traffic to the edge computing function hosting the application workload.

18. The method of claim 17, wherein each of the modules hosted on the set of interconnected edge data centers provide at least one of a radio access network (RAN) function, a 5th Generation of Mobile Networks (5G) core function, and an edge computing function.

19. The method of any of claims 17–18, wherein the first request and the second request each include a common traffic class identifier.

20. The method of any of claims 17–19, further comprising:

in response to receiving the first request and the second request, sending, by the controller, an application workload activation request to the edge computing function hosting the application workload to activate the edge computing function for access by the first device via the first edge data center and the second device via the second edge data center.