

CLAIMS

What is claimed is:

1. An orchestration system operated by a data center provider for a data center, the orchestration system comprising:
 - processing circuitry coupled to a memory;
 - logic stored in the memory and configured for execution by the processing circuitry, wherein the logic is operative to:
 - compute an aggregate bandwidth for a plurality of flows associated with a tenant of the data center provider and processed by a virtual network function, assigned to the tenant, executing on a server of the data center; and
 - modify, based on the aggregate bandwidth, an allocation of compute resources of the server for executing the virtual network function.
2. The orchestration system of claim 1 further comprising logic configured to compute an average bandwidth for each flow based upon flow statistics.
3. The orchestration system of claim 1 further comprising logic configured to compare an allocation of bandwidth to the aggregate bandwidth to produce a comparison result and store the comparison result for trending.
4. The orchestration system of claim 1, wherein the logic is further operative to map packet flow virtual network interfaces (VNIs) to VNI space of a plurality of tenants of the data center, wherein the tenant of the data center is associated with at least one VNI of the VNI space, and identify flow statistics corresponding to the tenant based upon the at least one VNI.
5. The orchestration system of claim 1, wherein the logic is further operative to compute a bandwidth utilization rate based upon an average packet size of one or more of the plurality of flows.
6. The orchestration system of claim 1, wherein the logic is further operative to modify, based on an aggregate bandwidth trend, the allocation of compute resources.

7. The orchestration system of claim 1, wherein the logic is further operative to modify the allocation of compute resources based upon a resource allocation table.
8. The orchestration system of claim 1, wherein the logic is further operative to:
 - compute an average bandwidth by multiplying a number of packets per second by an average packet size; and
 - compute the aggregated bandwidth as a summation of each average bandwidth of a plurality of packet flows.
9. The orchestration system of claim 1, wherein the logic is further operative to:
 - compute an average bandwidth by dividing a total number of bytes by an amount of time in seconds; and
 - compute the aggregated bandwidth as a summation of each average bandwidth of a plurality of packet flows.
10. A method of an orchestration system operated by a data center provider for a data center, the method comprising:
 - computing, by the orchestration system, an aggregate bandwidth for a plurality of flows associated with a tenant of the data center provider and processed by a virtual network function, assigned to the tenant, executing on a server of the data center; and
 - by the orchestration system, modifying, based on the aggregate bandwidth, an allocation of compute resources of the server executing the virtual network function.
11. The method of claim 10, wherein computing the aggregated bandwidth further comprises computing an average bandwidth for each flow based upon flow statistics.
12. The method of claim 10, wherein modifying, based on the aggregate bandwidth, the allocation of compute resources further comprising comparing an allocation of bandwidth to the aggregate bandwidth to produce a comparison result and store the comparison result for trending.

13. The method of claim 10, wherein computing the aggregated bandwidth further comprises determining an average packet size based on the packet processing statistics.
14. The method of claim 10, wherein modifying, based on the aggregate bandwidth, the allocation of compute resources further comprising mapping packet flow virtual network interfaces (VNIs) to tenant VNI space of a plurality of tenants of the data center, and identify flow statistics with at least one VNI associated with each tenant.
15. The method of claim 10, wherein modifying, based on the aggregate bandwidth, the allocation of compute resources further comprises modifying, based on an aggregate bandwidth trend, the allocation of compute resources.
16. The method of claim 10, wherein modifying, based on the aggregate bandwidth, the allocation of compute resources further comprises modifying the allocation of compute resources based upon a resource allocation table.
17. The method of claim 10, wherein computing the aggregated bandwidth further comprises:
- computing an average bandwidth by multiplying a number of packets per second by an average packet size; and
 - computing the aggregated bandwidth as a summation of each average bandwidth of a plurality of packet flows.
18. The method of 10, wherein computing the aggregated bandwidth further comprises:
- computing an average bandwidth by dividing a total number of bytes by an amount of time in seconds; and
 - computing the aggregated bandwidth as a summation of each average bandwidth of a plurality of packet flows.
19. An interconnection system comprising:
- at least one interconnection facility and a programmable network platform,
 - the at least one interconnection facility including:

a cluster comprising one or more computing devices to host virtual network functions (VNFs) for tenants, wherein each tenant corresponds to one or more VNFs forming a virtual router through which packet flows are exchanged between a service and a customer,

the programmable network platform being configured to:

compute an aggregate bandwidth for a plurality of flows associated with a tenant of the tenants and processed by a VNF of the VNFs, assigned to the tenant, executing on a computing device of the cluster; and

modify, based on the aggregate bandwidth, an allocation of compute resources of the computing device executing the virtual network function.

20. The interconnection system of claim 19, wherein the programmable network platform further configured to:

mapping packet flow virtual network interfaces (VNIs) to VNI space of a plurality of tenants of the data center, wherein the tenant of the data center is associated with at least one VNI of the VNI space, and identify flow statistics corresponding to the tenant based upon the at least one VNI.