CLAIMS

1. An orchestration system operated by a data center provider for a data center, the orchestration system comprising:

processing circuitry coupled to a memory; and

software instructions stored in the memory and configured for execution by the processing circuitry to:

map packet flow virtual network identifiers (VNIs) to a VNI space of a plurality of tenants of the data center, wherein a tenant of the data center is associated with a first VNI of the VNI space and a second VNI of the VNI space;

identify flow statistics corresponding to the tenant based upon the first VNI and the second VNI in the mapping;

compute, based on the flow statistics indicating a first packet flow having the first VNI associated with the tenant and a second packet flow having the second VNI associated with the tenant, an aggregate bandwidth for a plurality of packet flows associated with the tenant and processed by a virtual network function (VNF) that is assigned to the tenant and executes on physical compute resources of a server; and

modify, based on the aggregate bandwidth, an allocation of the physical compute resources of the server to re-assign a portion of the physical compute resources of the server between the VNF assigned to the tenant and one or more VNFs associated with a different tenant.

2. The orchestration system of claim 1, wherein the software instructions are configured for execution by the processing circuitry to compute an average bandwidth for each of the first packet flow and the second packet flow based upon the flow statistics.

3. The orchestration system of claim 1, wherein the software instructions are configured for execution by the processing circuitry 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 software instructions are configured for execution by the processing circuitry to compute a bandwidth utilization rate based upon an average packet size of one or more of the plurality of packet flows.

5. The orchestration system of claim 1, wherein the software instructions are configured for execution by the processing circuitry to modify, based on an aggregate bandwidth trend, the allocation of the physical compute resources.

6. The orchestration system of claim 1, wherein the software instructions are configured for execution by the processing circuitry to modify the allocation of the physical compute resources based upon a resource allocation table.

7. The orchestration system of claim 1, wherein the software instructions are configured for execution by the processing circuitry to:

compute an average bandwidth by multiplying a number of packets per second by an average packet size; and

compute the aggregate bandwidth as a summation of each average bandwidth of the plurality of packet flows.

8. The orchestration system of claim 1, wherein the software instructions are configured for execution by the processing circuitry to:

compute an average bandwidth by dividing a total number of bytes by an amount of time in seconds; and

compute the aggregate bandwidth as a summation of each average bandwidth of a plurality of packet flows.

9. A method of an orchestration system operated by a data center provider for a data center, the method comprising:

mapping packet flow virtual network identifiers (VNIs) to a VNI space of a plurality of tenants of the data center, wherein a tenant of the data center is associated with a first VNI of the VNI space and a second VNI of the VNI space;

identifying flow statistics corresponding to the tenant based upon the first VNI and the second VNI in the mapping;

computing, based on the flow statistics indicating a first packet flow having the first VNI associated with the tenant and a second packet flow having the second VNI associated with the tenant, an aggregate bandwidth for a plurality of packet flows associated with the tenant and processed by a virtual network function (VNF) that is assigned to the tenant and executes on physical compute resources of a server of the data center; and

modifying, based on the aggregate bandwidth, an allocation of the physical compute resources of the server to re-assign a portion of the physical compute resources of the server between the VNF assigned to the tenant and one or more VNFs associated with a different tenant.

10. The method of claim 9, wherein computing the aggregate bandwidth further comprises computing an average bandwidth for each of the first packet flow and the second packet flow based upon the flow statistics.

11. The method of claim 9, wherein modifying the allocation of physical 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.

12. The method of claim 9, wherein computing the aggregate bandwidth further comprises determining an average packet size based on the flow statistics.

13. The method of claim 9, wherein modifying the allocation of physical compute resources further comprises modifying, based on an aggregate bandwidth trend, the allocation of the physical compute resources.

14. The method of claim 9, wherein modifying, based on the aggregate bandwidth, the allocation of physical compute resources further comprises modifying the allocation of the physical compute resources based upon a resource allocation table.

15. The method of claim 9, wherein computing the aggregate bandwidth further comprises:

computing an average bandwidth by multiplying a number of packets per second by an average packet size; and

computing the aggregate bandwidth as a summation of each average bandwidth of the plurality of packet flows.

16. The method of 9, wherein computing the aggregate bandwidth further comprises:

computing an average bandwidth by dividing a total number of bytes by an amount of time in seconds; and

computing the aggregate bandwidth as a summation of each average bandwidth of the plurality of packet flows.

17. 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,

the programmable network platform being configured to:

map packet flow virtual network identifiers (VNIs) to a VNI space of a plurality of tenants of the at least one interconnection facility, wherein a tenant of the at least one interconnection facility is associated with a first VNI of the VNI space and a second VNI of the VNI space;

identify flow statistics corresponding to the tenant based upon the first VNI and the second VNI in the mapping;

compute, based on the flow statistics indicating a first packet flow having the first VNI associated with the tenant and a second packet flow having the second VNI associated with the tenant, an aggregate bandwidth for a plurality of packet flows associated with the tenant and processed by a virtual network function (VNF) that is assigned to the tenant and executes on physical compute resources of a server; and

modify, based on the aggregate bandwidth, an allocation of the physical compute resources of the server to re-assign a portion of the physical compute resources of the server between the VNF assigned to the tenant and one or more VNFs associated with a different tenant.