CLAIMS

1. A distributed bandwidth allocation system for managing bandwidth allocation to two or more entities that are subject to a total bandwidth allocation, the system comprising:

a distributed bandwidth limiter that:

receives first usage data describing a first bandwidth usage value for a first entity; and

receives second usage data describing a second bandwidth usage value for a second entity; and

based on the first usage data, second usage data, and a total bandwidth allocation, determines a first bandwidth allocation specifying bandwidth available for network traffic for the first entity and a second bandwidth allocation that specifies bandwidth available for network traffic for the second entity, wherein a sum of the first bandwidth allocation and the second bandwidth allocation does not exceed the total bandwidth allocation;

a first throttler that:

receives from the distributed bandwidth limiter the first bandwidth allocation;

throttles bandwidth available for network traffic for the first entity according to the first bandwidth allocation; and

reports to the distributed bandwidth limiter the first usage data; and

a second throttler that:

receives from the distributed bandwidth limiter the second bandwidth allocation;

throttles bandwidth available for network traffic for the second entity according to the second bandwidth allocation; and

reports to the distributed bandwidth limiter the second usage data,

wherein the distributed bandwidth limiter determines the first bandwidth allocation and the second bandwidth allocation by determining, based on prior first bandwidth usage values, a first entity bandwidth demand, and based on prior second bandwidth usage values, a second entity bandwidth demand,

wherein the distributed bandwidth limiter iteratively determines the first bandwidth allocation and the second bandwidth allocation and, for each iteration, provides

the first bandwidth allocation to the first throttler and the second bandwidth allocation to the second throttle, and

wherein for each iteration the distributed bandwidth limiter:

determines a first entity scale factor based on a current first bandwidth usage value and the first entity bandwidth demand, and a second entity scale factor based on a current second bandwidth usage value and the second entity bandwidth demand; and

determines the first bandwidth allocation and the second bandwidth allocation based on the respective first entity scale factor and second entity scale factor.

2. The system of claim 1, wherein:

for each iteration the distributed bandwidth limiter:

determines whether the current first bandwidth usage value meets a first entity demand threshold based on the first entity bandwidth demand; and

determines whether the current second bandwidth usage value meets a second entity demand threshold based on the second entity bandwidth demand;

for each iteration for which the current first bandwidth usage value meets the first entity demand threshold, the distributed bandwidth limiter increments the first entity scale factor, and for each iteration for which the current first bandwidth usage value does not meet the first entity demand threshold, the distributed bandwidth limiter resets the first entity scale factor to an initial first entity scale factor value; and

for each iteration for which the current second bandwidth usage value meets the second entity demand threshold, the distributed bandwidth limiter increments the second entity scale factor, and for each iteration for which the current second bandwidth usage value does not meet the second entity demand threshold, the distributed bandwidth limiter resets the second entity scale factor to an initial second entity scale factor value.

3. The system of claim 2, wherein:

the distributed bandwidth limiter determines the first entity bandwidth demand based on a product of the first entity scale factor and at least one of a plurality of prior first bandwidth usage values; and

the distributed bandwidth limiter determines the second entity bandwidth demand

based on a product of the second entity scale factor and at least one of a plurality of prior second bandwidth usage values.

4. The system of claim 1, wherein for iterations for which the distributed bandwidth limiter determines a sum of the first entity bandwidth demand and second entity bandwidth demand exceeds the total allocation bandwidth:

the distributed bandwidth limiter sets the first bandwidth allocation to a guaranteed first bandwidth allocation and sets the second bandwidth allocation to a guaranteed second bandwidth allocation when both the first entity bandwidth demand exceeds the guaranteed first bandwidth allocation and the second entity bandwidth demand exceeds the guaranteed second bandwidth allocation;

the distributed bandwidth limiter sets the second bandwidth allocation to the second entity bandwidth demand and then sets the first bandwidth allocation to a difference of the total allocation bandwidth and the second bandwidth allocation when the second entity bandwidth demand is less than the guaranteed second bandwidth allocation; and

the distributed bandwidth limiter sets the first bandwidth allocation to the first entity bandwidth demand and then sets the second bandwidth allocation to a difference of the total allocation bandwidth and the first bandwidth allocation when first entity bandwidth demand is less than the guaranteed first bandwidth allocation.

5. The system of claim 1, wherein for iterations for which the distributed bandwidth limiter determines a sum of the first entity bandwidth demand and second entity bandwidth demand is less than the total allocation bandwidth:

the distributed bandwidth limiter sets the first bandwidth allocation to a guaranteed first bandwidth allocation and sets the second bandwidth allocation to a guaranteed second bandwidth allocation when both the first entity bandwidth demand is less than the guaranteed first bandwidth allocation and the second entity bandwidth demand is less than the guaranteed second bandwidth allocation;

the distributed bandwidth limiter sets the first bandwidth allocation to the first entity bandwidth demand and then sets the second bandwidth allocation to a difference of the total allocation bandwidth and the first bandwidth allocation when the first entity bandwidth demand is greater than the guaranteed first bandwidth allocation; and

the distributed bandwidth limiter sets the second bandwidth allocation to the second entity bandwidth demand and then sets the first bandwidth allocation to a difference of the total allocation bandwidth and the second bandwidth allocation when the second entity bandwidth demand is greater than the guaranteed second bandwidth allocation.

6. The system of any one of the preceding claims, wherein:

the first entity is a storage system and the second entity is a virtual machine for which the storage system provides data storage;

the first usage data is storage system usage data and the first bandwidth usage value is a storage bandwidth usage value for network traffic for the storage system; and

the second usage data is virtual machine usage data for a virtual machine and the second bandwidth usage value is for network traffic for the virtual machine;

the first bandwidth allocation is a storage bandwidth allocation specifying bandwidth available for network traffic for the storage system;

the second bandwidth allocation is a virtual machine bandwidth allocation that specifies bandwidth available for network traffic for the virtual machine;

the first throttler is a storage system throttler that throttles bandwidth available for network traffic for the storage system according to the storage system bandwidth allocation; and

the second throttler is a virtual machine throttler that throttles bandwidth available for network traffic for the virtual machine according to the virtual machine bandwidth allocation.

7. The system of claim 6, wherein:

the distributed bandwidth limiter and the storage system throttler are instantiated in a hypervisor that manages the virtual machine and network traffic for the storage system storage for the virtual machine; and

the virtual machine throttler is implemented in a virtual switch through which network traffic for the virtual machine is routed to a network external to a host machine on which the virtual machine is instantiated.

8. A computer implemented method performed by a data processing apparatus, the method comprising:

receiving, by the data processing apparatus and from a first throttler for a first entity, first usage data describing a first bandwidth usage value for the first entity; and

receiving, by the data processing apparatus and from a second throttler for a second entity, second usage data describing a second bandwidth usage value for the second entity;

based the first usage data, second usage data, and a total bandwidth allocation, determining, by the data processing apparatus, a first bandwidth allocation specifying bandwidth available for network traffic for the first entity and a second bandwidth allocation that specifies bandwidth available for network traffic for the second entity, wherein a sum of the first bandwidth allocation and the second bandwidth allocation does not exceed the total bandwidth allocation;

providing, by the data processing apparatus, the first bandwidth allocation to the first throttler, wherein the first throttler throttles bandwidth available for network traffic for the first entity according to the first bandwidth allocation; and

providing, by the data processing apparatus, the second bandwidth allocation to the second throttler, wherein the second throttler throttles bandwidth available for network traffic for the second entity according to the second bandwidth allocation,

wherein determining the first bandwidth allocation and the second bandwidth allocation comprise determining, based on prior first bandwidth usage values, a first entity bandwidth demand, and based on prior second bandwidth usage values, a second entity bandwidth demand,

wherein the determining the first bandwidth allocation and the second bandwidth allocation comprises iteratively determining the first bandwidth allocation and the second bandwidth allocation and, for each iteration, providing the first bandwidth allocation to the first throttler and the second bandwidth allocation to the second throttle, and

wherein for each iteration, the method comprises:

determining a first entity scale factor based on a current first bandwidth usage value and the first entity bandwidth demand, and a second entity scale factor based on a current second bandwidth usage value and the second entity bandwidth demand; and

determining the first bandwidth allocation and the second bandwidth allocation based

on the respective first entity scale factor and second entity scale factor.

9. The computer-implemented method of claim 8, wherein:

for each iteration the method comprises:

determining whether the current first bandwidth usage value meets a first entity demand threshold based on the first entity bandwidth demand; and

determining whether the current second bandwidth usage value meets a second entity demand threshold based on the second entity bandwidth demand;

for each iteration for which the current first bandwidth usage value meets the first entity demand threshold, incrementing the first entity scale factor, and for each iteration for which the current first bandwidth usage value does not meet the first entity demand threshold, the resetting the first entity scale factor to an initial first entity scale factor value; and

for each iteration for which the current second bandwidth usage value meets the second entity demand threshold, the incrementing the second entity scale factor, and for each iteration for which the current second bandwidth usage value does not meet the second entity demand threshold, the resetting the second entity scale factor to an initial second entity scale factor value.

10. The computer-implemented method of claim 9, further comprising:

determining the first entity bandwidth demand based on a product of the first entity scale factor and at least one of a plurality of prior first bandwidth usage values; and

determining the second entity bandwidth demand based on a product of the second entity scale factor and at least one of a plurality of prior second bandwidth usage values.

11. The computer-implemented method of claim 8, wherein for iterations for which a sum of the first entity bandwidth demand and second entity bandwidth demand exceeds the total allocation bandwidth:

setting the first bandwidth allocation to a guaranteed first bandwidth allocation and setting the second bandwidth allocation to a guaranteed second bandwidth allocation when both the first entity bandwidth demand exceeds the guaranteed first bandwidth allocation and the second entity bandwidth demand exceeds the guaranteed second bandwidth allocation;

setting the second bandwidth allocation to the second entity bandwidth demand and then setting the first bandwidth allocation to a difference of the total allocation bandwidth and the second bandwidth allocation when the second entity bandwidth demand is less than the guaranteed second bandwidth allocation; and

setting the first bandwidth allocation to the first entity bandwidth demand and then setting the second bandwidth allocation to a difference of the total allocation bandwidth and the first bandwidth allocation when first entity bandwidth demand is less than the guaranteed first bandwidth allocation.

12. The computer-implemented method of claim 8, wherein for iterations for which a sum of the first entity bandwidth demand and second entity bandwidth demand is less than the total allocation bandwidth:

setting the first bandwidth allocation to a guaranteed first bandwidth allocation and setting the second bandwidth allocation to a guaranteed second bandwidth allocation when both the first entity bandwidth demand is less than the guaranteed first bandwidth allocation and the second entity bandwidth demand is less than the guaranteed second bandwidth allocation;

setting the first bandwidth allocation to the first entity bandwidth demand and then setting the second bandwidth allocation to a difference of the total allocation bandwidth and the first bandwidth allocation when the first entity bandwidth demand is greater than the guaranteed first bandwidth allocation; and

setting the second bandwidth allocation to the second entity bandwidth demand and then setting the first bandwidth allocation to a difference of the total allocation bandwidth and the second bandwidth allocation when the second entity bandwidth demand is greater than the guaranteed second bandwidth allocation.

13. The computer-implemented method of any one of claims 8 to 12, wherein:

the first entity is a storage system and the second entity is a virtual machine for which the storage system provides data storage;

the first usage data is storage system usage data and the first bandwidth usage value is a storage bandwidth usage value for network traffic for the storage system; and

the second usage data is virtual machine usage data for a virtual machine and the

second bandwidth usage value is for network traffic for the virtual machine;

the first bandwidth allocation is a storage bandwidth allocation specifying bandwidth available for network traffic for the storage system;

the second bandwidth allocation is a virtual machine bandwidth allocation that specifies bandwidth available for network traffic for the virtual machine;

the first throttler is a storage system throttler that throttles bandwidth available for network traffic for the storage system according to the storage system bandwidth allocation; and

the second throttler is a virtual machine throttler that throttles bandwidth available for network traffic for the virtual machine according to the virtual machine bandwidth allocation.

14. A non-transitory computer storage device storing instructions that upon execution by a data processing apparatus cause the data processing apparatus to perform operations comprising:

receiving from a first throttler for a first entity, first usage data describing a first bandwidth usage value for the first entity; and

receiving from a second throttler for a second entity, second usage data describing a second bandwidth usage value for the second entity;

based the first usage data, second usage data, and a total bandwidth allocation, determining a first bandwidth allocation specifying bandwidth available for network traffic for the first entity and a second bandwidth allocation that specifies bandwidth available for network traffic for the second entity, wherein a sum of the first bandwidth allocation and the second bandwidth allocation does not exceed the total bandwidth allocation;

providing the first bandwidth allocation to the first throttler, wherein the first throttler throttles bandwidth available for network traffic for the first entity according to the first bandwidth allocation; and

providing the second bandwidth allocation to the second throttler, wherein the second throttler throttles bandwidth available for network traffic for the second entity according to the second bandwidth allocation,

wherein the distributed bandwidth limiter determines the first bandwidth allocation and the second bandwidth allocation by determining, based on prior first bandwidth usage

values, a first entity bandwidth demand, and based on prior second bandwidth usage values, a second entity bandwidth demand,

wherein the distributed bandwidth limiter iteratively determines the first bandwidth allocation and the second bandwidth allocation and, for each iteration, provides the first bandwidth allocation to the first throttler and the second bandwidth allocation to the second throttle, and

wherein for each iteration the distributed bandwidth limiter:

determines a first entity scale factor based on a current first bandwidth usage value and the first entity bandwidth demand, and a second entity scale factor based on a current second bandwidth usage value and the second entity bandwidth demand; and

determines the first bandwidth allocation and the second bandwidth allocation based on the respective first entity scale factor and second entity scale factor.