CLAIMS

What is claimed is:

1. A computing device comprising:

a network interface controller (NIC); and

processing circuitry having access to storage media encoded with instructions that, when executed, cause the processing circuitry to:

receive, by a main packet processor, a packet from the NIC;

send, by the main packet processor, based on a virtual extensible local area network (VXLAN) tunnel endpoint (VTEP) indicated by the packet, the packet to a tenant-specific packet processor associated with the VTEP;

send, by the tenant-specific packet processor, at least a portion of the packet to a workload; and

process, by the workload, the at least a portion of the packet.

2. The computing device of claim 1,

wherein the main packet processor processes the packet with a first network stack, and

wherein the tenant-specific packet processor processes the packet with a different, second network stack.

3. The computing device of any of claims 1-2,

wherein the packet comprises a first packet,

wherein the tenant-specific packet processor comprises a first tenant-specific packet processor,

wherein the VTEP comprises a first VTEP,

wherein the workload comprises a first workload associated with a first tenant, and

wherein the processing circuitry is configured to:

receive, by the main packet processor, a second packet from the NIC;

send, by the main packet processor, based on a second VTEP indicated by the second packet, the second packet to a second tenant-specific packet processor associated with the second VTEP;

send, by the second tenant-specific packet processor, at least a portion of the second packet to a second workload associated with a second tenant; and

process, by the second workload, the at least a portion of the second packet.

4. The computing device of claim 3,

wherein the first VTEP is associated with a first virtual network address, and

wherein the second VTEP is associated with a different, second virtual network address.

5. The computing device of claim 3,

wherein the main packet processor processes the packet with a first network stack,

wherein the first tenant-specific packet processor processes the packet with a different, second network stack, and

wherein the first tenant-specific packet processor processes the packet with a different, third network stack.

6. The computing device of any of claims 1-5,

wherein the VTEP is associated with a virtual network address, and

wherein the main packet processor sends the packet to the tenant-specific packet processor associated with the VTEP based on a destination network address of the packet matching the virtual network address associated with the VTEP.

7. The computing device of any of claims 1-6, wherein the workload implements a tenant control plane for a tenant associated with the VTEP.

8. The computing device of any of claims 1-7, further comprising:

a flood list for a tenant associated with the VTEP, the flood list storing an entry specifying an underlay network address for another computing device,

wherein the packet comprises a first packet;

wherein the processing circuitry is configured to:

receive, by the tenant-specific packet processor, an original packet from the workload;

generate a second packet comprising:

a tunnel header comprising the underlay network address for the another computing device,

an outer header comprising a virtual network address associated with the VTEP,

a VXLAN header, and

the original packet;

output the second packet via the NIC to the another computing device.

9. The computing device of claim 8, wherein the storage media is further encoded with instructions that, when executed, cause the processing circuitry to:

receive, by an agent, the entry; and

store the entry to the flood list.

10. The computing device of any of claims 1-9,

wherein the main packet processor and the tenant-specific packet processor are configured with a shared memory interface, and

wherein the main packet processor is configured to send, via the shared memory interface, the packet to a tenant-specific packet processor associated with the VTEP.

11. The computing device of any of claims 1-10,

wherein the main packet processor and the tenant-specific packet processor execute in separate namespaces.

12. A computing system comprising:

a first computing device configured with a first virtual extensible local area network (VXLAN) tunnel endpoint (VTEP) for a tenant, the first VTEP having an interface configured with a virtual network address; and

a second computing device configured with a second VTEP for the tenant, the second VTEP having an interface configured with the virtual network address.

13. The computing system of claim 12, wherein the first computing device comprises:

a main packet processor configured to obtain a packet,

wherein the main packet processor is further configured to send, based on the packet having a destination virtual network address that matches the virtual network address configured for the interface of the first VTEP, the packet to a tenant-specific packet processor associated with the first VTEP;

the tenant-specific packet processor, configured to send at least a portion of the packet to a workload; and

the workload, configured to process the at least a portion of the packet.

14. A method comprising:

receiving, by a main packet processor of a computing device, a packet from a network interface controller (NIC) of the computing device;

sending, by the main packet processor, based on a virtual extensible local area network (VXLAN) tunnel endpoint (VTEP) indicated by the packet, the packet to a tenant-specific packet processor of the computing device, the tenant-specific packet processor associated with the VTEP;

sending, by the tenant-specific packet processor, at least a portion of the packet to a workload of the computing device; and

processing, by the workload, the at least a portion of the packet.

15. The method of claim 14,

wherein the main packet processor processes the packet with a first network stack, and

wherein the tenant-specific packet processor processes the packet with a different, second network stack.

16. The method of any of claims 14-15,

wherein the packet comprises a first packet,

wherein the tenant-specific packet processor comprises a first tenant-specific packet processor,

wherein the VTEP comprises a first VTEP,

wherein the workload comprises a first workload associated with a first tenant,

the method further comprising:

receiving, by a main packet processor, a second packet from the NIC;

sending, by the main packet processor, based on a second VTEP indicated by the second packet, the second packet to a second tenant-specific packet processor associated with the second VTEP;

sending, by the second tenant-specific packet processor, at least a portion of the second packet to a second workload associated with a second tenant; and

processing, by the second workload, the at least a portion of the second packet.

17. The method of claim 16,

wherein the first VTEP is associated with a first virtual network address, and

wherein the second VTEP is associated with a different, second virtual network address.

18. The method of any of claims 14-17,

wherein the VTEP is associated with a virtual network address, and

sending, by the main packet processor, the packet to the tenant-specific packet processor associated with the VTEP based on a destination network address of the packet matching the virtual network address associated with the VTEP.

19. The method of any of claims 14-18, further comprising:

storing, by the computing device, a flood list for a tenant associated with the VTEP, the flood list storing an entry specifying an underlay network address for another computing device,

wherein the packet comprises a first packet;

the method further comprising:

receiving, by the tenant-specific packet processor, an original packet from the workload;

generating a second packet comprising:

a tunnel header comprising the underlay network address for the another computing device,

an outer header comprising a virtual network address associated with the VTEP,

a VXLAN header, and

the original packet; and

outputting the second packet via the NIC to the another computing device.

20. The method of any of claims 14-19, wherein the main packet processor and the tenant-specific packet processor execute in separate namespaces.

21. A computer-readable storage medium encoded with instructions for causing one or more programmable processors to become configured to carry out the method of any of claims 14–20 or to become configured as the computing device of any of claims 1–13.