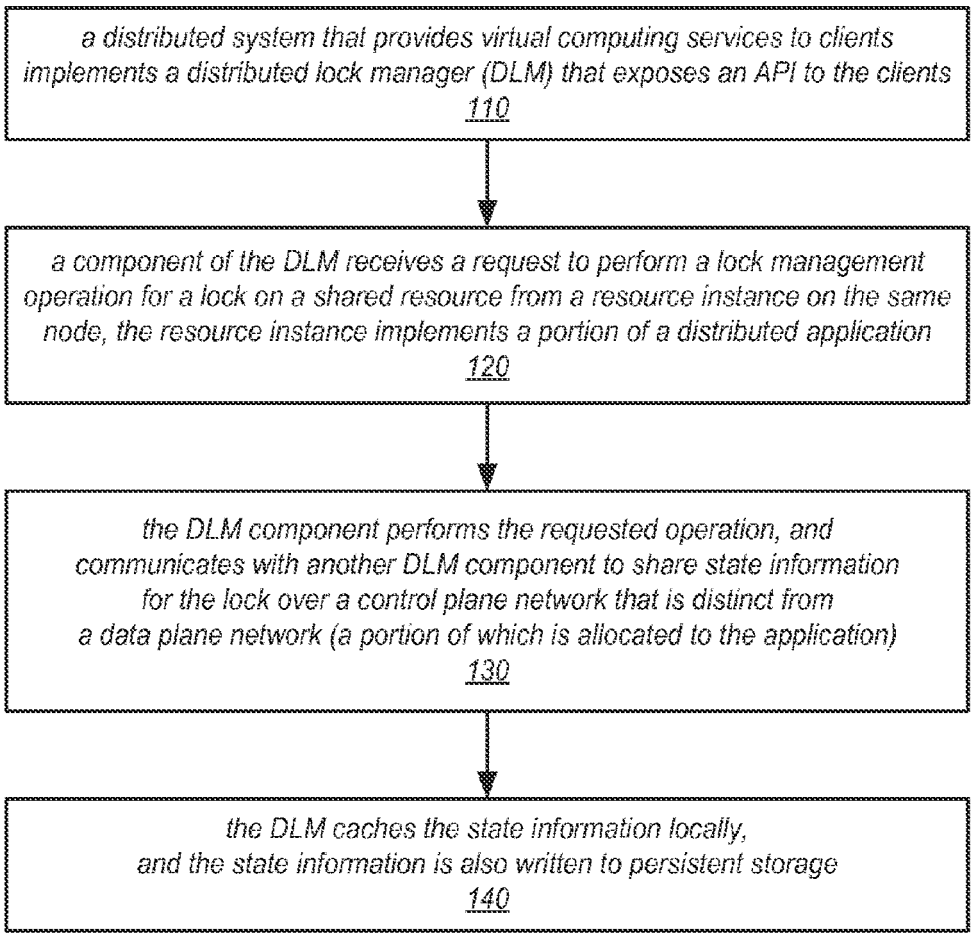
a distributed system that provides virtual computing services to clients implements a distributed lock manager (DLM) that exposes an API to the clients
110
a component of the DLM receives a request to perform a lock management operation for a lock on a shared resource from a resource instance on the same node, the resource instance implements a portion of a distributed application
120
the DLM component performs the requested operation, and communicates with another DLM component to share state information for the lock over a control plane network that is distinct from a data plane network (a portion of which is allocated to the application)
130
the DLM caches the state information locally, and the state information is also written to persistent storage
140

FIG. 1

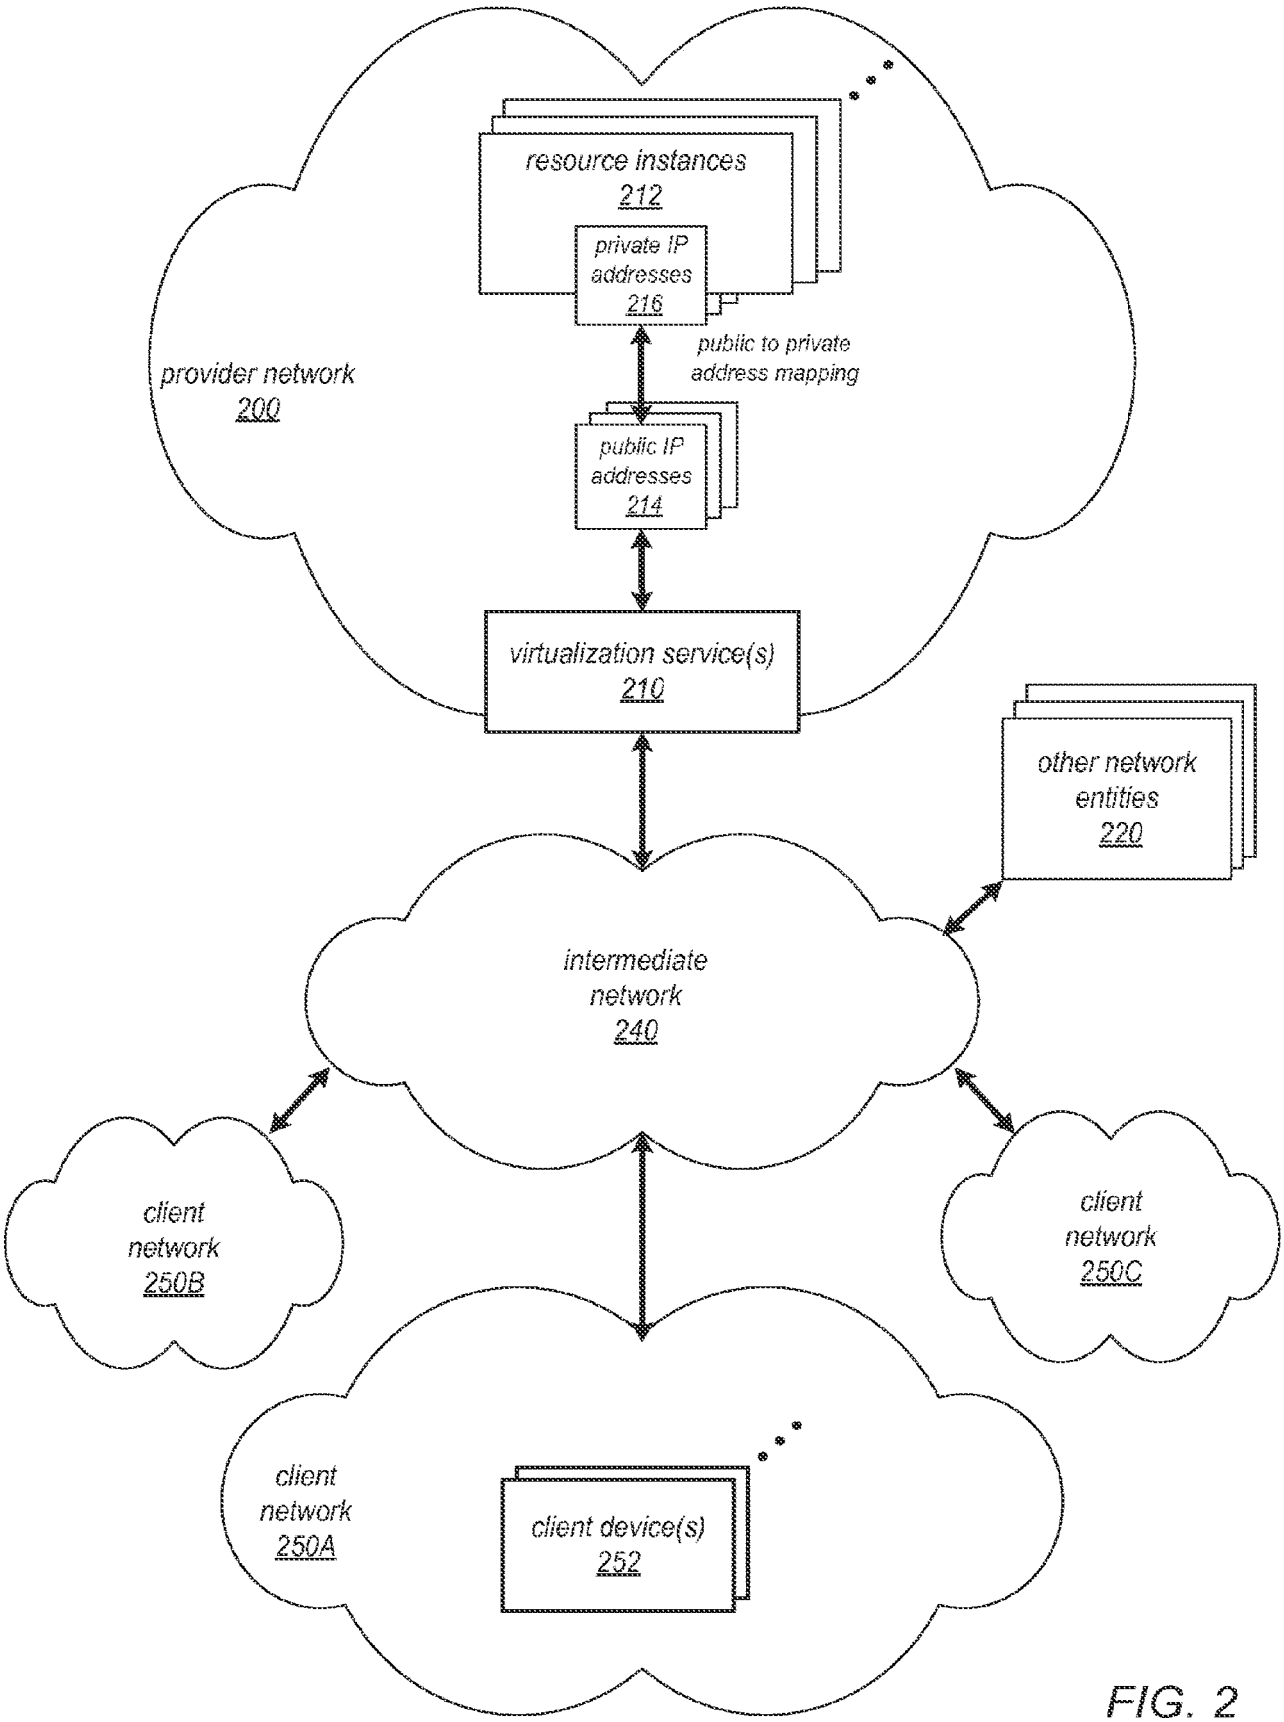
resource instances
212
private IP addresses
216
public to private address mapping
provider network
200
public IP addresses
214
virtualization service(s)
210
other network entities
220
intermediate network
240
client network
250B
client network
250C
client network
250A
client device(s)
252

FIG. 2

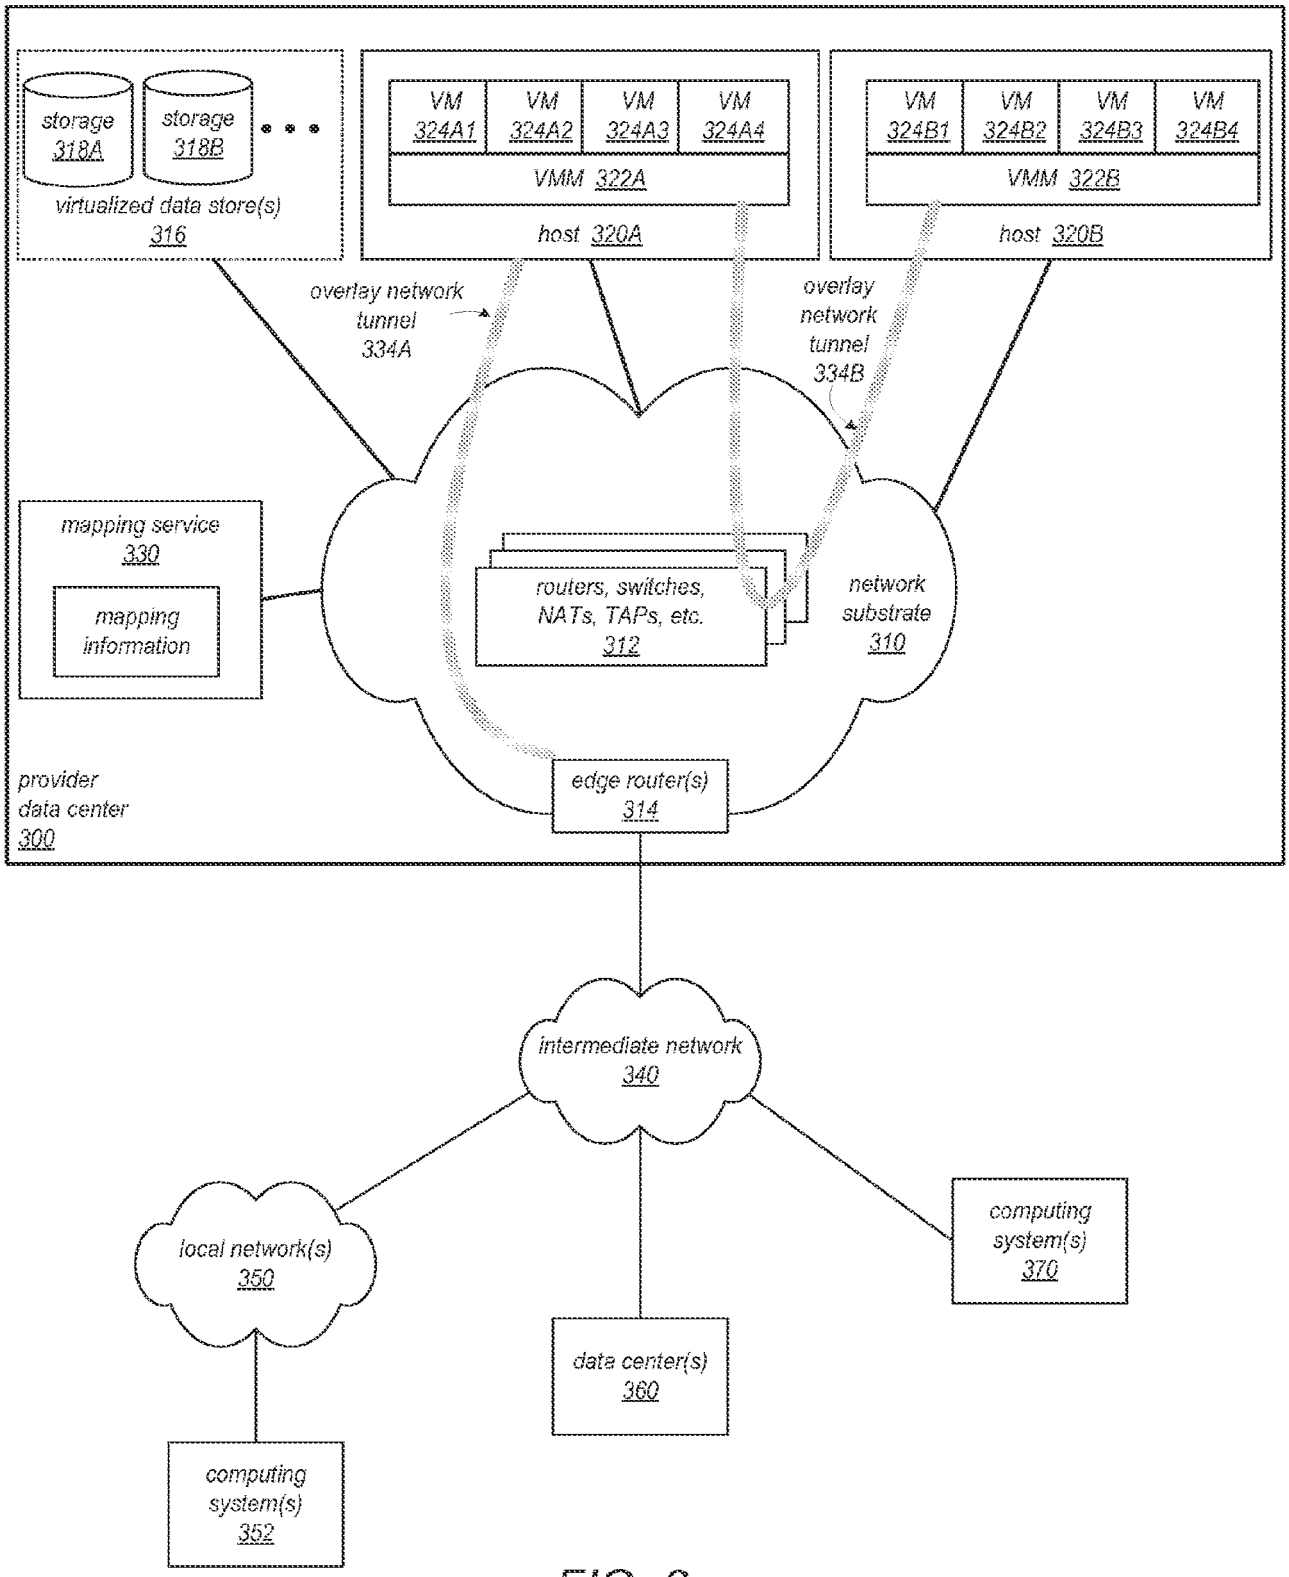
storage 318A
storage 318B
virtualized data store(s) 316
VM 324A1
VM 324A2
VM 324A3
VM 324A4
VMM 322A
host 320A
VM 324B1
VM 324B2
VM 324B3
VM 324B4
VMM 322B
host 320B
overlay network tunnel 334A
overlay network tunnel 334B
mapping service 330
mapping information
routers, switches, NATs, TAPs, etc. 312
network substrate 310
edge router(s) 314
provider data center 300
intermediate network 340
local network(s) 350
computing system(s) 370
data center(s) 360
computing system(s) 352

FIG. 3

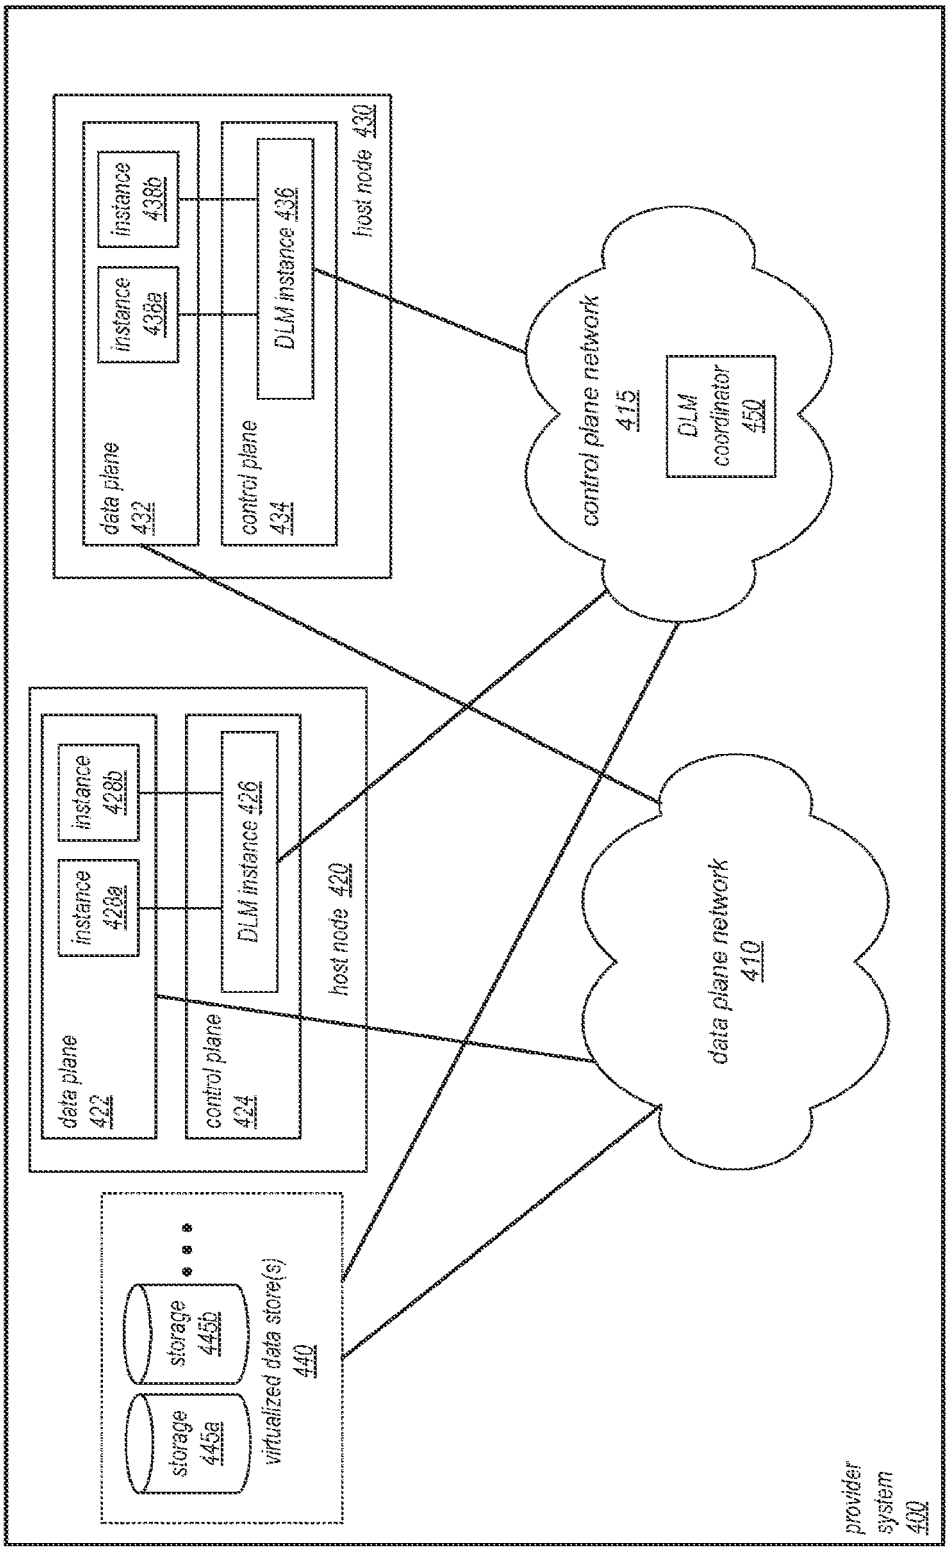
provider system 400
virtualized data store(s) 440
storage 445a
storage 445b
host node 420
data plane 422
instance 428a
instance 428b
control plane 424
DLM instance 426
host node 430
data plane 432
instance 438a
instance 438b
control plane 434
DLM instance 436
data plane network 410
control plane network 415
DLM coordinator 450

FIG. 4

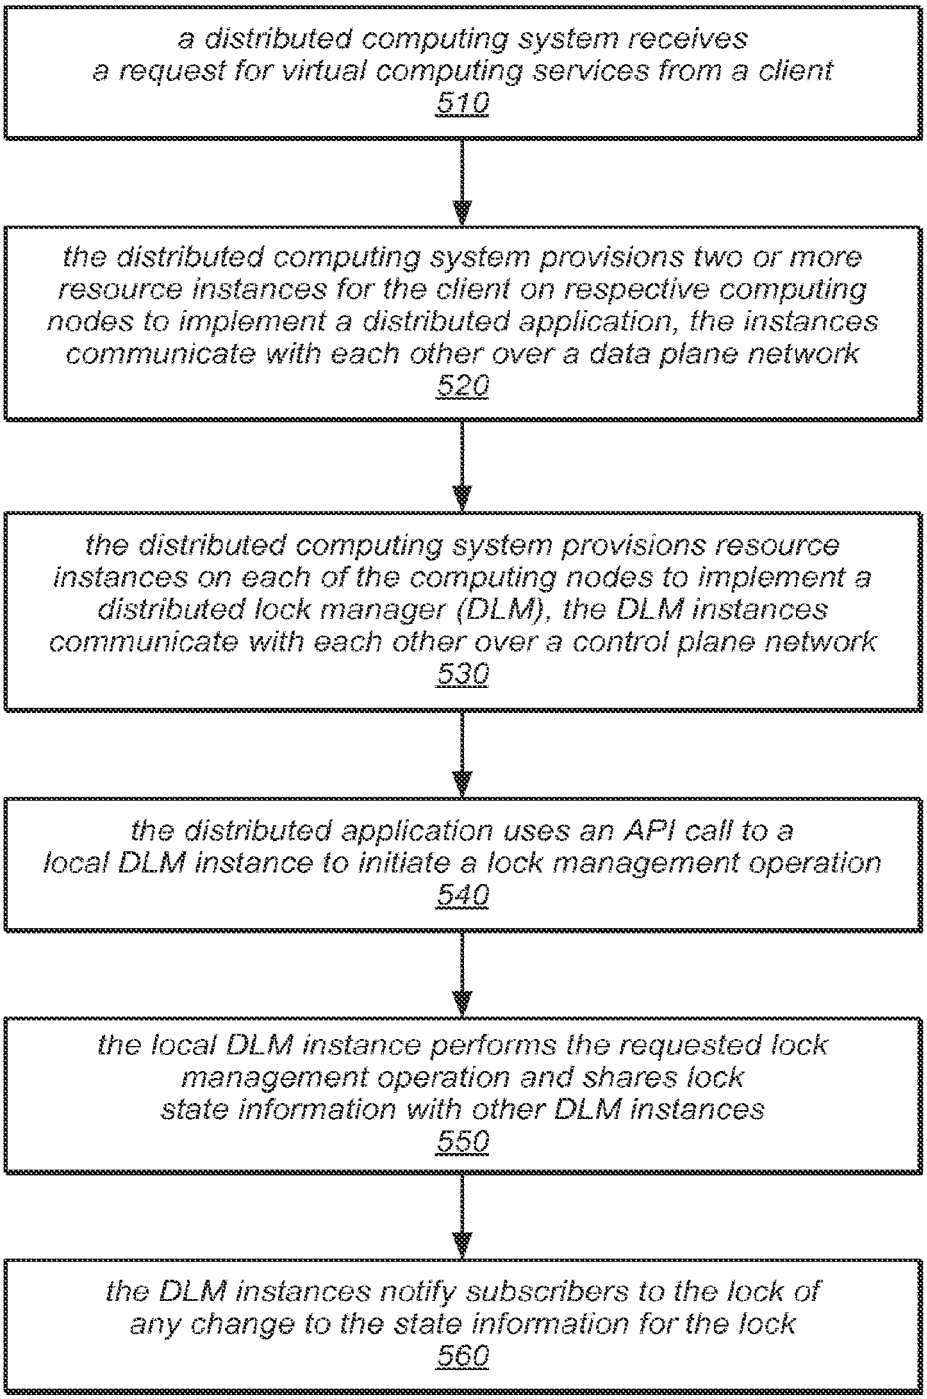
a distributed computing system receives a request for virtual computing services from a client
510
the distributed computing system provisions two or more resource instances for the client on respective computing nodes to implement a distributed application, the instances communicate with each other over a data plane network
520
the distributed computing system provisions resource instances on each of the computing nodes to implement a distributed lock manager (DLM), the DLM instances communicate with each other over a control plane network
530
the distributed application uses an API call to a local DLM instance to initiate a lock management operation
540
the local DLM instance performs the requested lock management operation and shares lock state information with other DLM instances
550
the DLM instances notify subscribers to the lock of any change to the state information for the lock
560

FIG. 5

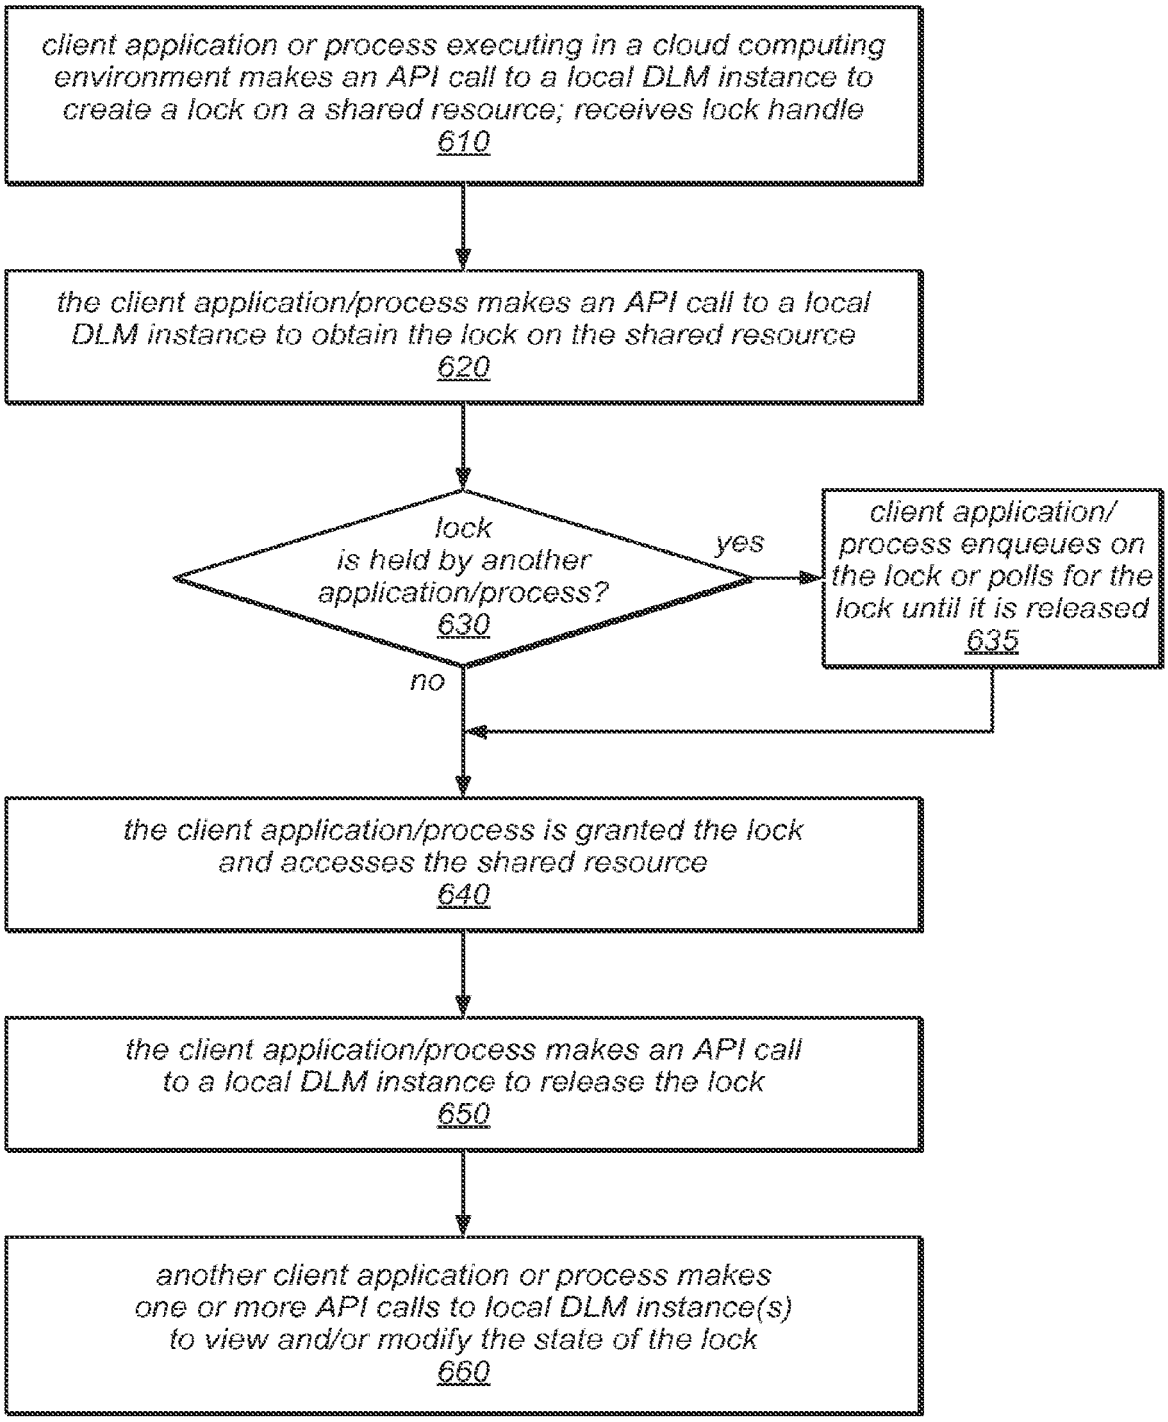
client application or process executing in a cloud computing environment makes an API call to a local DLM instance to create a lock on a shared resource; receives lock handle
610
the client application/process makes an API call to a local DLM instance to obtain the lock on the shared resource
620
lock is held by another application/process?
630
yes
client application/ process enqueues on the lock or polls for the lock until it is released
635
no
the client application/process is granted the lock and accesses the shared resource
640
the client application/process makes an API call to a local DLM instance to release the lock
650
another client application or process makes one or more API calls to local DLM instance(s) to view and/or modify the state of the lock
660

FIG. 6

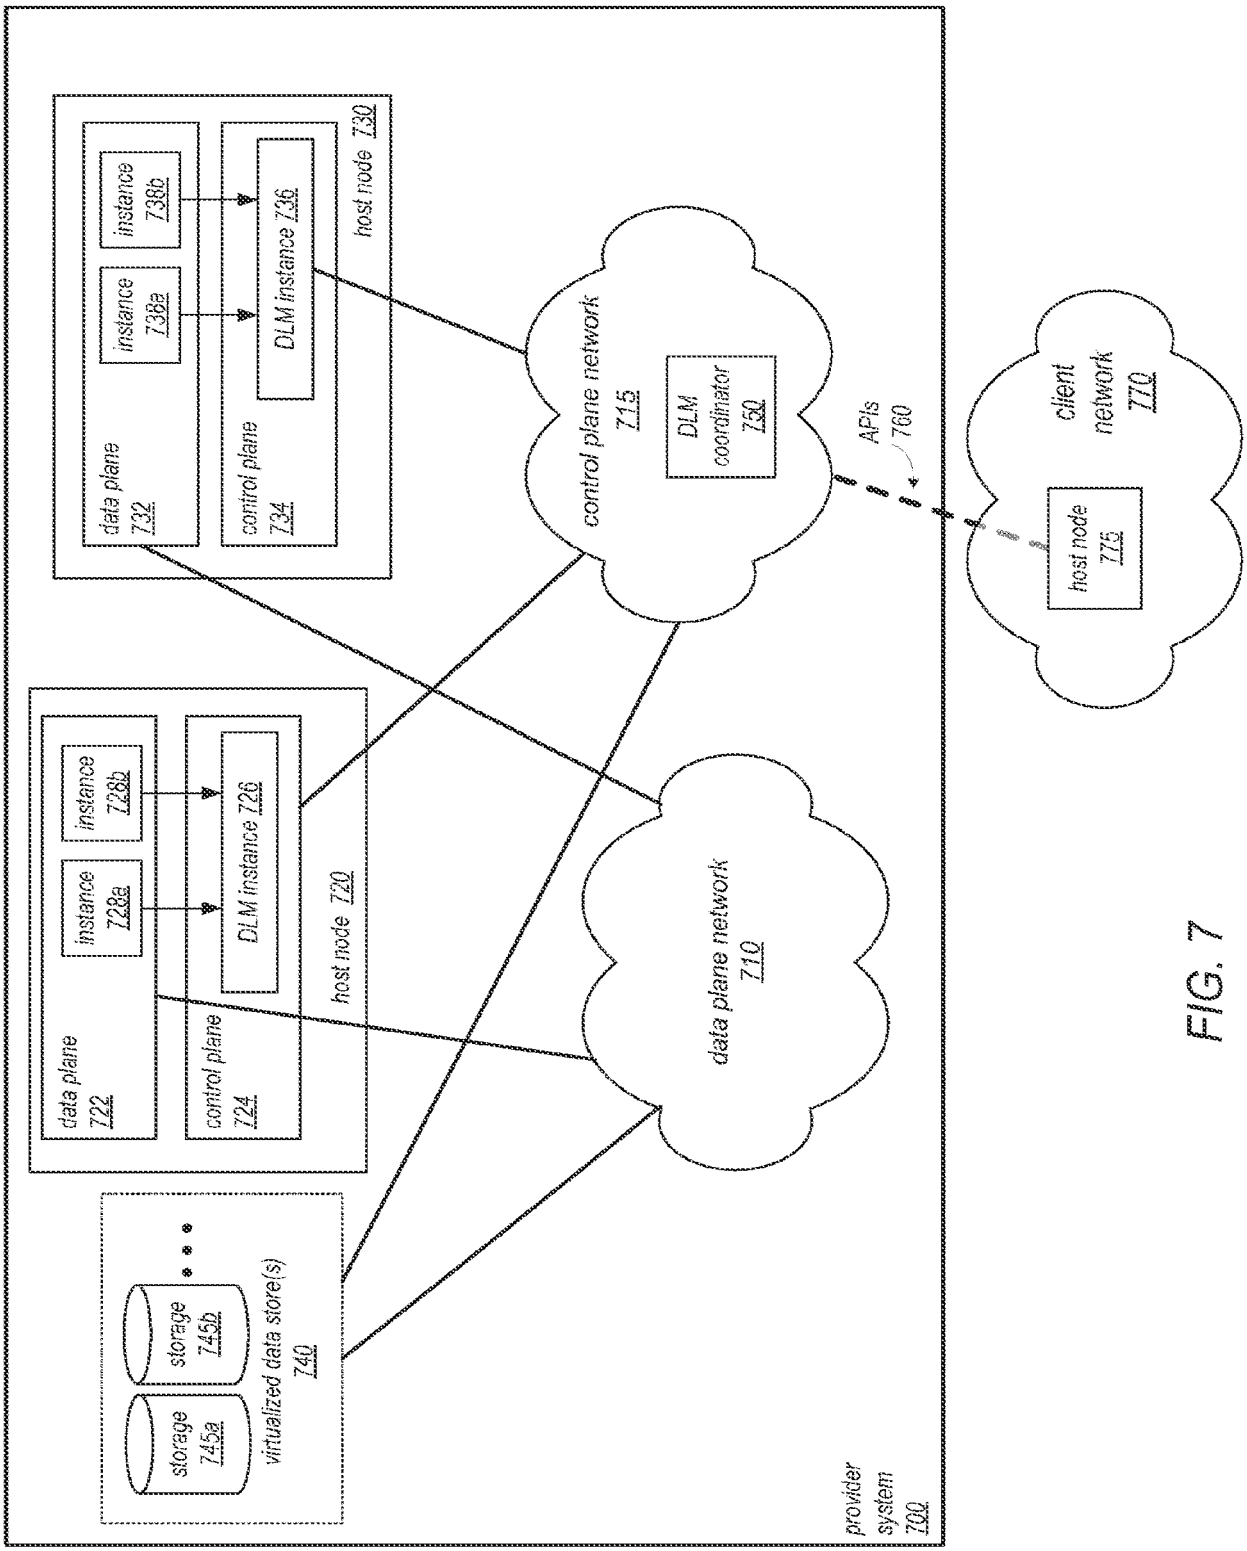
provider system 700
virtualized data store(s) 740
storage 745a
storage 745b
host node 720
data plane 722
instance 728a
instance 728b
control plane 724
DLM instance 726
data plane network 710
host node 730
data plane 732
instance 738a
instance 738b
control plane 734
DLM instance 736
control plane network 715
DLM coordinator 750
APIs 760
client network 770
host node 775

FIG. 7

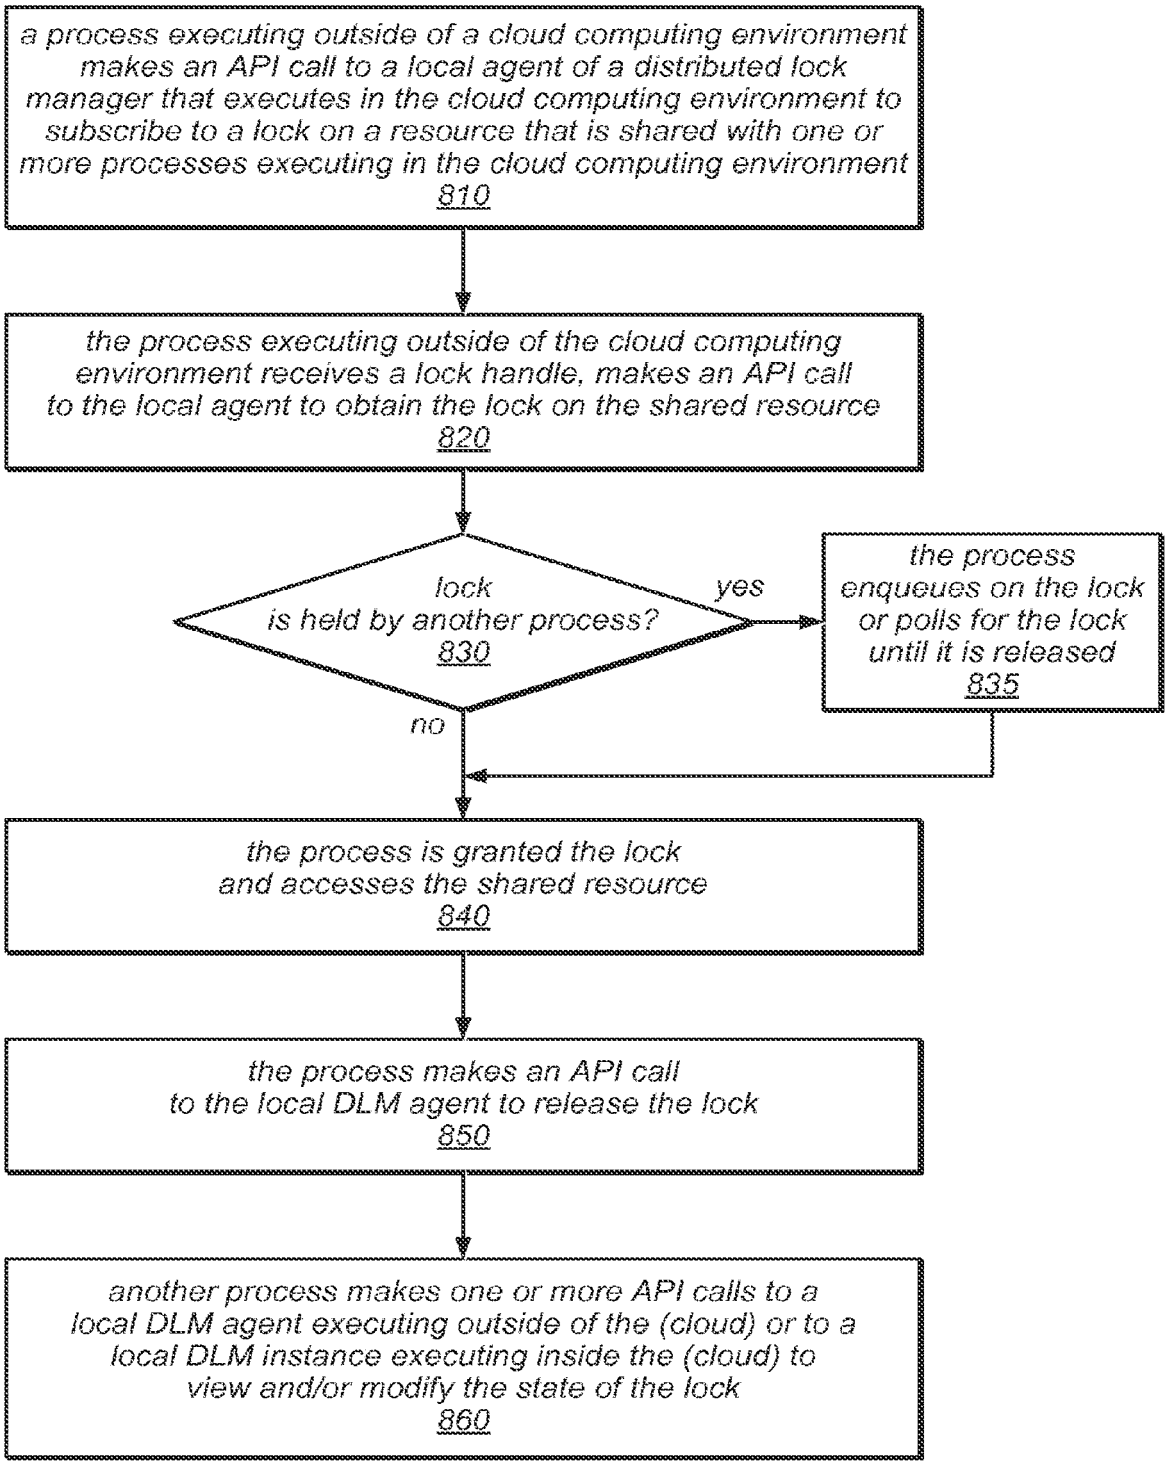
a process executing outside of a cloud computing environment makes an API call to a local agent of a distributed lock manager that executes in the cloud computing environment to subscribe to a lock on a resource that is shared with one or more processes executing in the cloud computing environment
810
the process executing outside of the cloud computing environment receives a lock handle, makes an API call to the local agent to obtain the lock on the shared resource
820
lock is held by another process?
830
yes
the process enqueues on the lock or polls for the lock until it is released
835
no
the process is granted the lock and accesses the shared resource
840
the process makes an API call to the local DLM agent to release the lock
850
another process makes one or more API calls to a local DLM agent executing outside of the (cloud) or to a local DLM instance executing inside the (cloud) to view and/or modify the state of the lock
860

FIG. 8

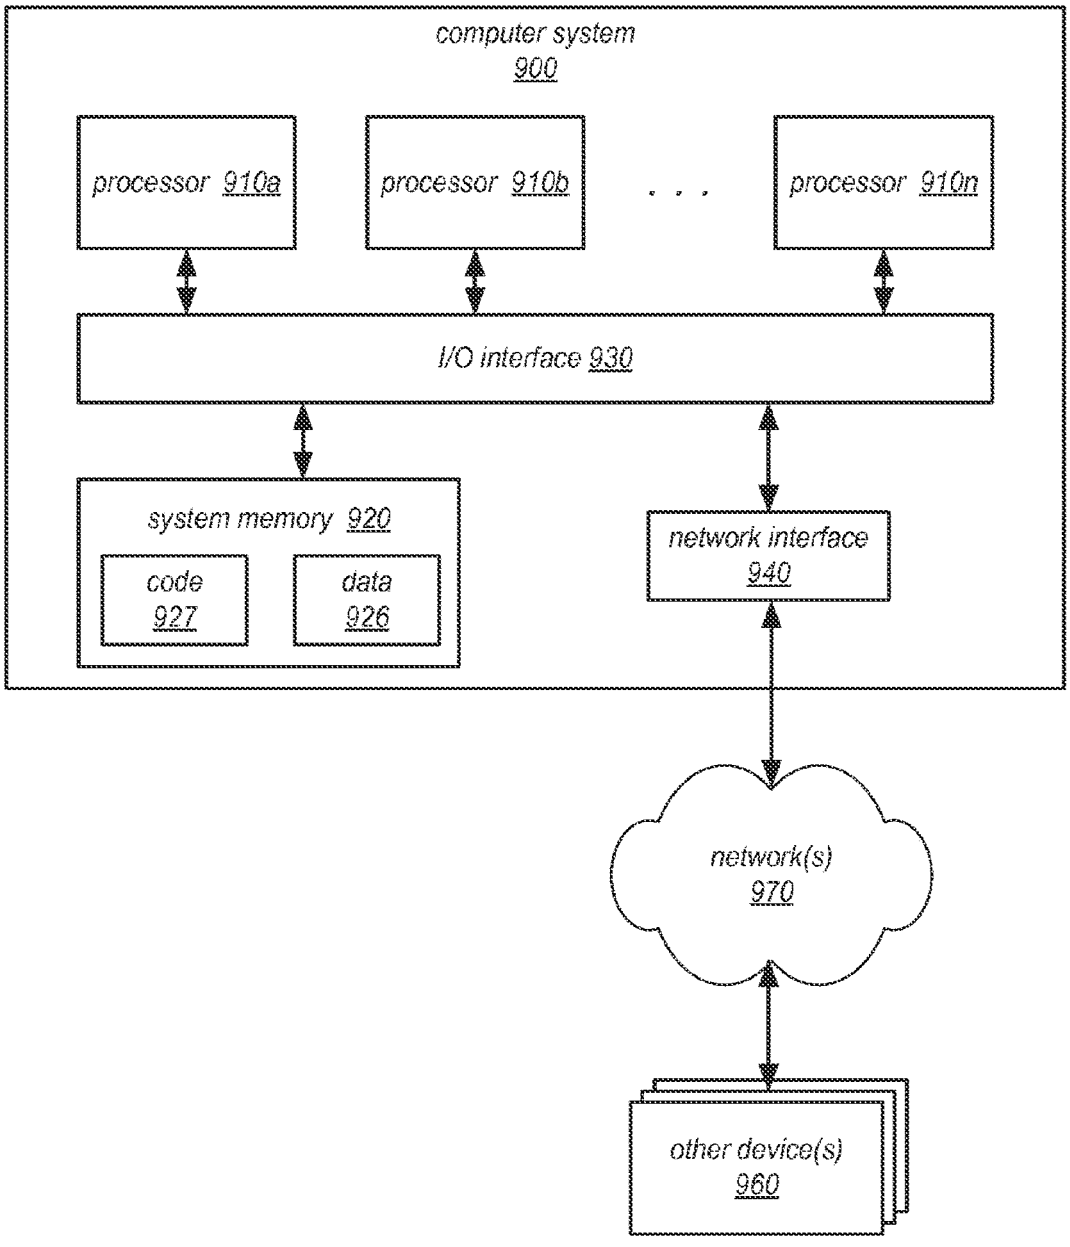
computer system
900
processor 910a
processor 910b
. . .
processor 910n
I/O interface 930
system memory 920
code
927
data
926
network interface
940
network(s)
970
other device(s)
960

FIG. 9