CLAIMS

1. A system comprising:

at least one processor; and

memory, operatively connected to the at least one processor and storing instructions that, when executed by the at least one processor, cause the system to perform a set of operations, the set of operations comprising:

receiving, from a client device, a request for computing functionality of a content distribution network (CDN);

generating a local performance metric associated with processing the request for computing functionality;

evaluating, based at least in part on a remote performance metric for a remote CDN node, the local performance metric to determine whether to process the request for computing functionality;

when it is determined to process the request for computing functionality:

generating a response to the request to the request for computing functionality; and

providing the response to the client device; and

when it is determined not to process the request for computing functionality:

generating, via an overlay network, a route to the remote CDN node; and

providing, to the remote CDN node, an indication to process, via the generated route, the request for computing functionality.

2. The system of claim 1, wherein:

the remote performance metric is a first remote performance metric;

the remote CDN node is a first remote CDN node; and

the first remote performance metric is stored in a latency table that further comprises a second remote performance metric for a second remote CDN node.

3. The system of claim 2, wherein evaluating the local performance metric further comprises determining that the first remote CDN node has less latency than the second remote CDN node based on comparing the first performance metric to the second performance metric.

4. The system of any preceding claim, wherein the remote performance metric for the remote CDN node is based on at least one of:

a network response latency for the remote CDN node;

a compute engine response latency for the remote CDN node; or

a software response latency for the remote CDN node.

5. The system of claim 4, wherein the network response latency for the remote CDN node is determined based at least in part on an Address Resolution Protocol (ARP) control message received from the remote CDN node via the overlay network.

6. The system of any preceding claim, wherein the set of operations further comprises:

receiving, from the remote CDN node via the generated route, a response to the request for computing functionality; and

providing the received response to the client device.

7. The system of any preceding claim, wherein the set of operations further comprises:

receiving another request from the client device associated with the computing functionality; and

providing, to the remote CDN node via the generated route, an indication of the another request.

8. A method for providing a compute resource in a content distribution network (CDN), the method comprising:

receiving, by a CDN node, a request for a compute resource, wherein the CDN node does not provide the compute resource;

determining, using a set of service rules, whether to reconfigure a compute engine of the CDN node to provide the compute resource; and

based on determining to reconfigure the compute engine:

identifying software associated with the requested compute resource;

executing the identified software using the compute engine; and

processing the request for the compute resource using the compute engine of the CDN node.

9. The method of claim 8, wherein the set of service rules comprises at least one of:

a service rule based on a time of day;

a service rule based on a day of week; or

a service rule based on an amount of requests for the compute resource within a predetermined period of time compared to a predetermined threshold.

10. The method of claim 8 or 9, wherein the software associated with the requested compute resource comprises at least one of:

a virtual machine;

a container; or

a software package.

11. The method of any one of claims 8 to 10, wherein identifying the software associated with the requested compute resource further comprises accessing the software from another CDN node.

12. The method of any one of claims 8 to 11, wherein the request for the compute resource comprises an indication that the request is for a compute computing functionality type.

13. The method of any one of claims 8 to 12, wherein the request for the compute resource further comprises an indication of the software associated with the requested compute resource.

14. A method for processing a request in a distributed content distribution network (CDN), the method comprising:

receiving, from a client device, a request for computing functionality of a content

distribution network (CDN);

generating a local performance metric associated with processing the request for computing functionality;

evaluating, based at least in part on a remote performance metric for a remote CDN node, the local performance metric to determine whether to process the request for computing functionality;

based on determining not to process the request for computing functionality:

generating, via an overlay network, a route to the remote CDN node; and

providing, to the remote CDN node, an indication to process, via the generated route, the request for computing functionality.

15. The method of claim 14, wherein:

the remote performance metric is a first remote performance metric;

the remote CDN node is a first remote CDN node; and

the first remote performance metric is stored in a latency table that further comprises a second remote performance metric for a second remote CDN node.

16. The method of claim 15, wherein evaluating the local performance metric further comprises determining that the first remote CDN node has less latency than the second remote CDN node based on comparing the first performance metric to the second performance metric.

17. The method of any one of claims 14 to 16, wherein the remote performance metric for the remote CDN node is based on at least one of:

a network response latency for the remote CDN node;

a compute engine response latency for the remote CDN node; or

a software response latency for the remote CDN node.

18. The method of claim 17, wherein the network response latency for the remote CDN node is determined based at least in part on an Address Resolution Protocol (ARP) control message received from the remote CDN node via the overlay network.

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

receiving, from the remote CDN node via the generated route, a response to the request for computing functionality; and

providing the received response to the client device.

20. The method of any one of claims 14 to 19, further comprising:

receiving another request from the client device associated with the computing functionality; and

providing, to the remote CDN node via the generated route, an indication of the another request.