Claims

1. A method, comprising:

receiving customer information in a service request, wherein the customer information includes at least a desired computing service and customer location data for a customer site;

determining provider location data for a plurality of provider computing sites; determining a rough estimate of latency for each of the plurality of provider computing sites based on the customer location data and the provider location data;

providing a list of the provider computing sites based on the rough estimate of latency for each of the provider computing sites;

receiving a selection of one of the provider computing sites; and

providing the computing service to the customer site from the selected provider computing site.

2. The method of claim 1, further comprising:

receiving a latency requirement; and

determining, based on the latency requirement, a service radius;

wherein providing the list of the providing computing sites comprises including on the list only provider computing sites that are located within the service radius of the customer site.

3. The method of claim 1 or 2, further comprising:

determining whether a first provider computing site of the plurality of computing sites includes available resources necessary to provide the computing service;

when the first provider computing site does include available resources necessary to provide the computing service, including the first provider computing site on the list;

when the first provider computing site does not include available resources necessary to provide the computing service:

determining whether additional resources can be added to the first provider computing site to provide the computing service;

when the additional resources can be added to the first provider computing site to provide the computing service, including the first provider computing site on the list along with an indication of when the additional resources can be added;

when the additional resources cannot be added to the first provider computing site, excluding the first provider computing site from the list.

4. The method of any preceding claim, further comprising:

determining a first fine estimate of latency for a first provider computing site; and

determining a second fine estimate of latency for a second provider computing site of the plurality of computing sites;

providing a redetermined list of the plurality of provider computing sites based at least on the first fine estimate and the second fine estimate.

5. The method of claim 4, further comprising:

providing a client agent to a client device;

causing the client agent to perform a first latency test between the client device and the first provider computing site to determine the first fine estimate; and

causing the client agent to perform a second latency test between the client device and the second provider computing site to determine the second fine estimate.

6. The method of claim 5, further comprising:

causing a first latency test server to be instantiated at the first provider computing site;

causing a second latency test server to be instantiated at the second provider computing site;

wherein the first latency test is performed between the client device and the first latency test server and the second latency test is performed between the client device and the second latency test server.

7. The method of any one of claims 3 to 6, wherein the selected provider computing site is the first provider computing site and the first provider computing site does not include available resources necessary to provide the computing service, further comprising:

automatically causing the additional resources to be added to the first provider computing site.

8. The method of any preceding claim, wherein the customer location data comprises a street address of the customer site, further comprising:

translating the street address to a global positioning satellite (GPS) location; and

determining a distance between the GPS location and each of the provider computing sites;

wherein determining the rough estimate of latency is based on the distance and a speed of light.

9. The method of any one of claims 5 to 7, wherein at least the first latency test is a traceroute between the client device and the first provider computing site.

10. The method of claim 4, further comprising:

providing a client agent to a client device;

determining whether the client device is at the customer site;

when the client device is not at the customer site, sending a query to the client device to determine an internet protocol (IP) address of a second client device at the customer site;

receiving the IP address of the second client device;

performing a first latency test between the second client device and the first provider computing site to determine the first fine estimate;

performing a second latency test between the second client device and the second provider computing site to determine the second fine estimate; and

providing a redetermined list of the plurality of provider computing sites based at least on the first fine estimate and the second fine estimate.

11. A system, comprising:

at least one processor;

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 method, the method comprising:

receiving customer information in a service request, wherein the customer information includes at least a desired computing service and customer location data for a customer site;

determining provider location data for a plurality of provider computing sites;

determining a rough estimate of latency for each of the plurality of provider computing sites based on the customer location data and the provider location data;

providing a list of the provider computing sites based on the rough estimate of latency for each of the provider computing sites;

receiving a selection of one of the provider computing sites; and

providing the computing service to the customer site from the selected provider computing site.

12. The system of claim 11, wherein the method further comprises:

receiving a latency requirement; and

determining, based on the latency requirement, a service radius;

wherein providing the list of the providing computing sites comprises including on the list only provider computing sites that are located within the service radius of the customer site.

13. The system of claim 11 or 12, wherein the method further comprises:

determining whether a first provider computing site of the plurality of computing sites includes available resources necessary to provide the computing service;

when the first provider computing site does include available resources necessary to provide the computing service, including the first provider computing site on the list;

when the first provider computing site does not include available resources necessary to provide the computing service:

determining whether additional resources can be added to the first provider computing site to provide the computing service;

when the additional resources can be added to the first provider computing site to provide the computing service, including the first provider computing site on the list along with an indication of when the additional resources can be added;

when the additional resources cannot be added to the first provider computing site, excluding the first provider computing site from the list.

14. The system of any one of claims 11 to 13, wherein the method further comprises:

determining a first fine estimate of latency for a first provider computing site; and

determining a second fine estimate of latency for a second provider computing site of the plurality of computing sites;

providing a redetermined list of the plurality of provider computing sites based at least on the first fine estimate and the second fine estimate.

15. The system of claim 14, wherein the method further comprises:

providing a client agent to a client device;

causing the client agent to perform a first latency test between the client device and the first provider computing site to determine the first fine estimate; and

causing the client agent to perform a second latency test between the client device and the second provider computing site to determine the second fine estimate.

16. The system of claim 15, wherein the method further comprises:

causing a first latency test server to be instantiated at the first provider computing site;

causing a second latency test server to be instantiated at the second provider computing site;

wherein the first latency test is performed between the client device and the first latency test server and the second latency test is performed between the client device and the second latency test server.

17. The system of any one of claims 13 to 16, wherein the selected provider computing site is the first provider computing site and the first provider computing site does not include available resources necessary to provide the computing service, and wherein the method further comprises:

automatically causing the additional resources to be added to the first provider computing site.

18. The system of any one of claims 11 to 17, wherein the customer location data comprises a street address of the customer site, and wherein the method further comprises:

translating the street address to a global positioning satellite (GPS) location; and

determining a distance between the GPS location and each of the provider computing sites;

wherein determining the rough estimate of latency is based on the distance and a speed of light.

19. The system of claim 14, wherein the method further comprises:

providing a client agent to a client device;

determining whether the client device is at the customer site;

when the client device is not at the customer site, sending a query to the client device to determine an internet protocol (IP) address of a second client device at the customer site;

receiving the IP address of the second client device;

performing a first latency test between the second client device and the first provider computing site to determine the first fine estimate;

performing a second latency test between the second client device and the second provider computing site to determine the second fine estimate; and

providing a redetermined list of the plurality of provider computing sites based at least on the first fine estimate and the second fine estimate.

20. A method, comprising:

receiving customer information in a service request, wherein the customer information includes at least a latency requirement, a desired computing service, and customer location data for a customer site;

determining, based on the latency requirement, a service radius;

determining provider location data for a plurality of provider computing sites;

determining whether each of the plurality of provider computing sites is within the service radius of the customer site based on the customer location data and the provider location data;

providing a list of the provider computing sites that are within the service radius of the customer site;

receiving a selection of one of the provider computing sites from the list; and

providing the computing service to the customer site from the selected provider computing site.