



(19) **United States**

(12) **Patent Application Publication**
Cannon et al.

(10) **Pub. No.: US 2019/0199828 A1**

(43) **Pub. Date: Jun. 27, 2019**

(54) **IMAGE DISTRIBUTION IN AN EDGE COMPUTING SYSTEM**

(52) **U.S. Cl.**

CPC *H04L 67/34* (2013.01); *G06F 9/54* (2013.01); *G06F 9/44505* (2013.01); *G06F 8/60* (2013.01)

(71) Applicant: **Virtuosys Limited**, Bath (GB)

(72) Inventors: **Mark Joseph Cannon**, Gloucestershire (GB); **Neil Edgar**, Wiltshire (GB)

(57)

ABSTRACT

(21) Appl. No.: **16/223,384**

A method (1000, 1100) in a mesh computing system (400) of pre-emptively positioning images at hosts, for retrieval from the hosts and for the creation of containers from the images, is described. The hosts are located in edge nodes of the mesh computing system (400). The mesh computing system (400) also comprises a central registry (490), located outside the mesh computing system (400), the central registry (490) holding the images. A swarm leader (424) located in the mesh computing system (400) identifies a need in a first host at a first edge node to run an image, the image being held in the central registry. The swarm leader (424) identifies a second host at a second edge node, and commands download of the image from the central registry (490) to the second host, via at least one wireless link of the mesh computing system (400). A mesh computing system (400) is provided.

(22) Filed: **Dec. 18, 2018**

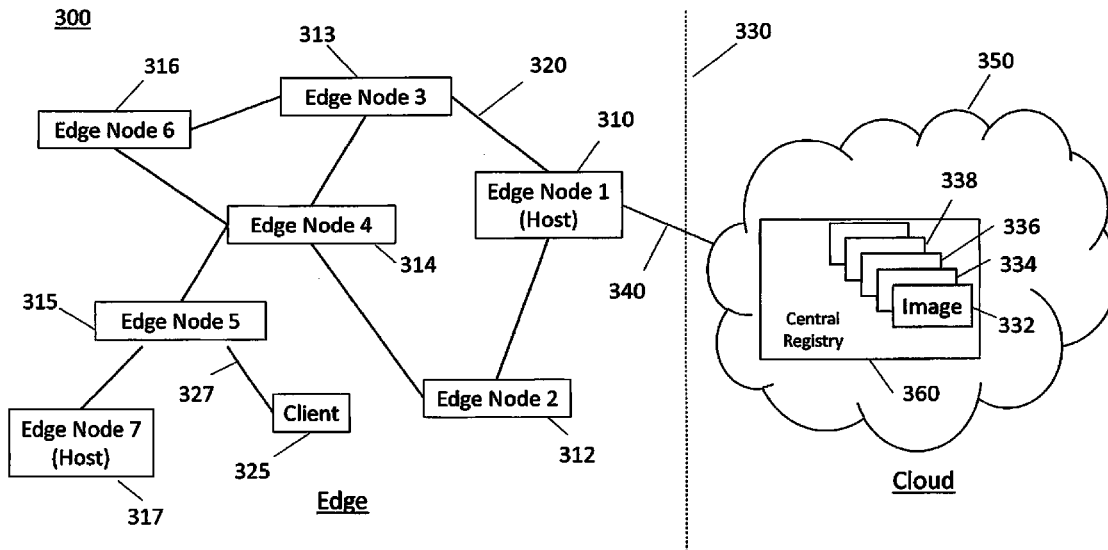
(30) **Foreign Application Priority Data**

Dec. 22, 2017 (GB) 1721776.1
May 25, 2018 (GB) 1808660.3

Publication Classification

(51) **Int. Cl.**

H04L 29/08 (2006.01)
G06F 8/60 (2006.01)
G06F 9/445 (2006.01)
G06F 9/54 (2006.01)






FIG. 1

100

130




FIG. 2

200




FIG. 3




FIG. 4




FIG. 6




FIG. 7

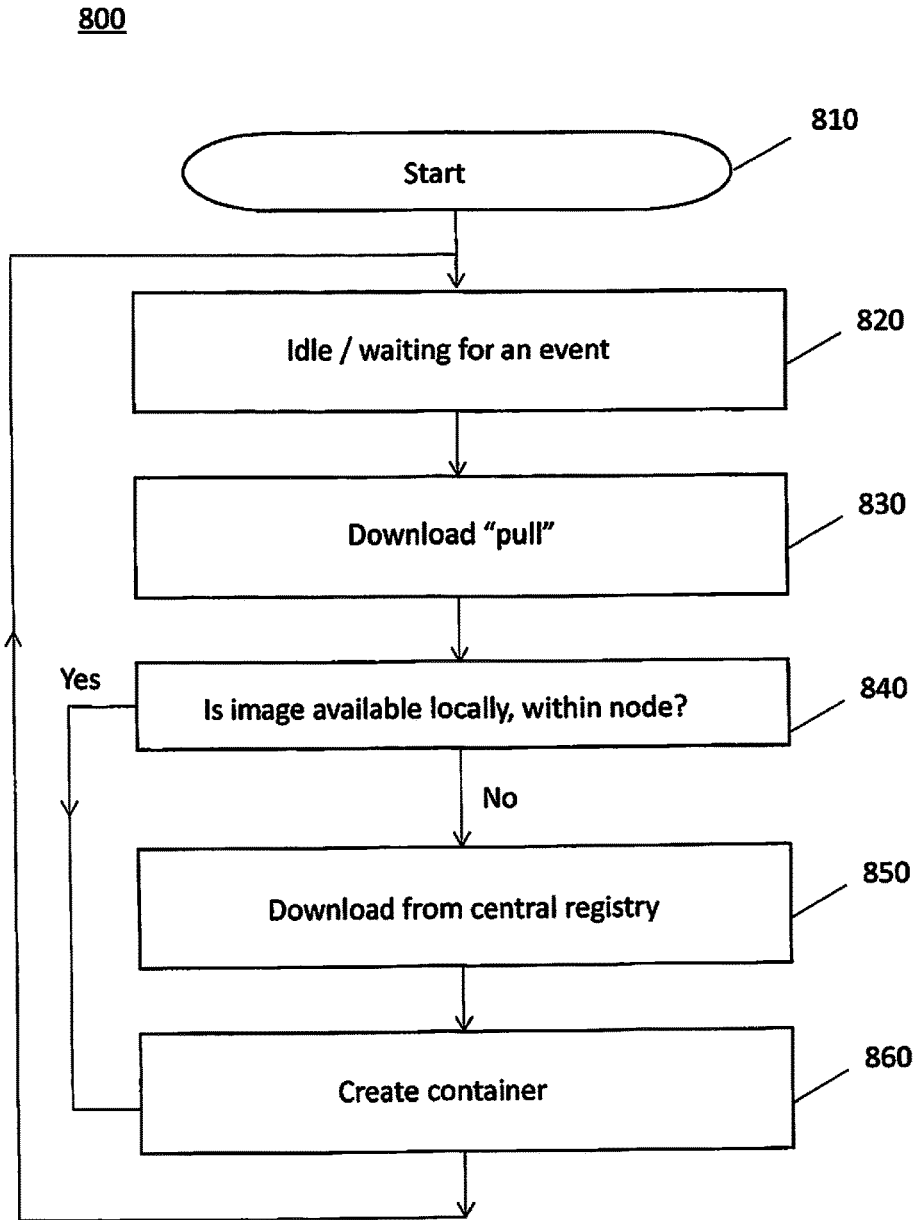


FIG. 8

900




FIG. 9

1000




FIG. 10

1100




FIG. 11

1200

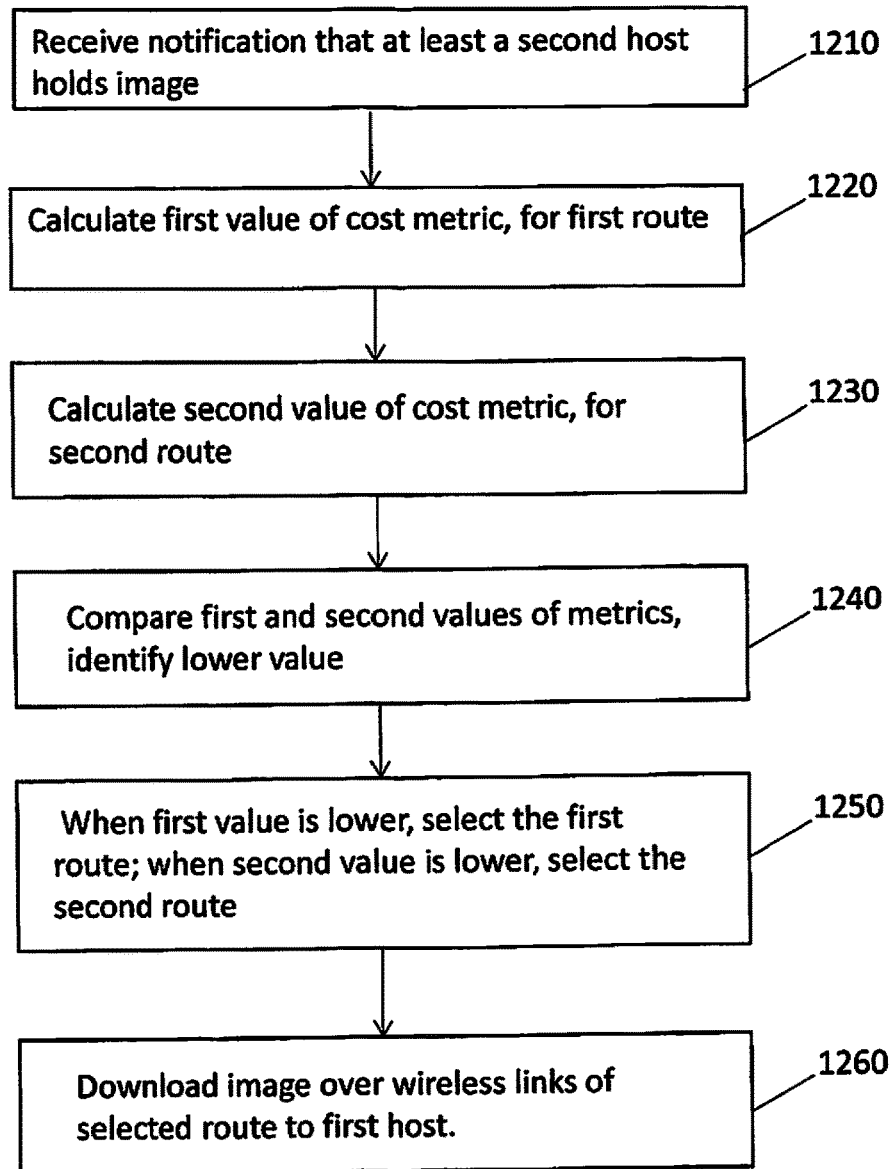


FIG. 12

1300

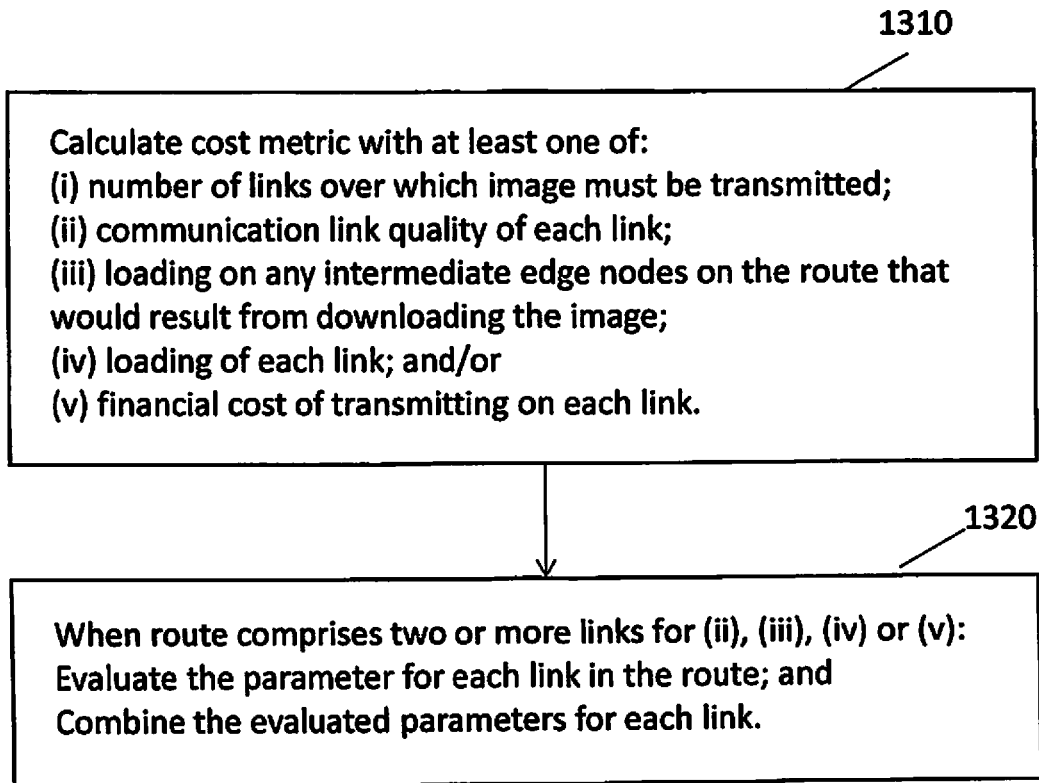


FIG. 13

1400

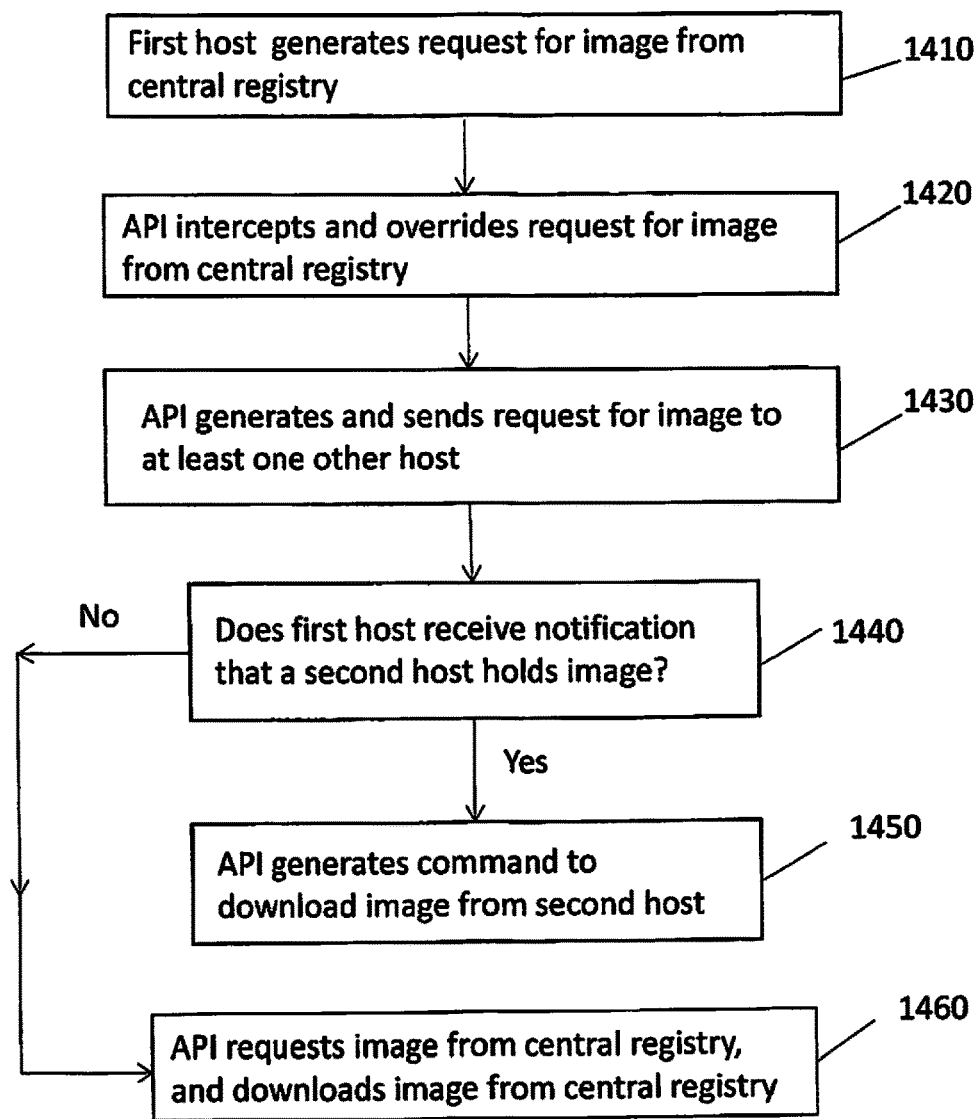


FIG. 14

1500




FIG. 15




FIG. 16

IMAGE DISTRIBUTION IN AN EDGE COMPUTING SYSTEM

RELATED APPLICATIONS

[0001] This application claims the benefit of Great Britain Application No. 1721776.1, filed on Dec. 22, 2017 and Great Britain Application No. 1808660.3, filed on May 25, 2018. The content of these applications are fully incorporated herein in their entirety.

TECHNICAL FIELD

[0002] The field of the invention is edge computing systems. In particular, the invention concerns image distribution in an edge computing system that implements virtualisation.

BACKGROUND

[0003] In some computing systems, Virtual Machines have been used to allow computing functions to be carried out at a desired location, which may be distant from a location where elements of a computing system were originally installed. Virtualisation at the level of the operating system in an edge computing system enables location-independent computing.

[0004] One approach to virtualisation is the use of ‘containers’. The aim of employing a container is to try to isolate an application and its dependencies into a self-contained unit that can be run anywhere. The container wraps up a piece of software in a complete file system. The file system then contains everything that it needs to run, such as code, runtime, system tools and system libraries. Containers have thus been used to provide operating-system-level virtualization. It is the very fact that the containerized application includes all the dependencies that enables it to run on any node into which it has been loaded. In turn, therefore, containers can be spread to new nodes relatively easily, and may be easier to update once deployed to multiple nodes.

[0005] One aspect of containerization is the use of ‘images’. An image is a package that can be executed to run of a piece of software. Thus, the image has the code for the software and information about configuration files. The image will also house libraries that are needed for the software, plus information on required variables. The advantage of images is that they can be moved around as a single entity. In terms of size and the requirements for transfer through a network, they are lightweight in comparison to more traditional Virtual Machines.

[0006] Open source projects have facilitated the deployment of various applications inside containers. A typical, but non-limiting, example of an open source approach is shown by ‘Docker™’, ‘Kubernetes™’ and ‘Coreos™ rkt’ provide aLTE™native approaches.

[0007] (i) Details of the Docker™ system are available on the website: <https://docs.Docker™.com/>

[0008] (ii) Details of the Kubernetes™ system are available on the website: <https://Kubernetes™.io/>

[0009] (iii) Details of the Coreos™ system are available on the website: <https://Coreos™.com/rkt>

[0010] For example, the Docker™ architecture has three main components. These components are the Host, Client and Registry. In addition, a ‘daemon’ is provided within the ‘Host’. In operation, the daemon selects an image, and creates a runnable instance of that image in a container.

Some images are locally available from within the Host itself. Any images that are not available in the Host will be downloaded from the Registry when needed, and then used to create a runnable instance of that image in a container.

[0011] FIG. 1 illustrates a possible arrangement of the Host, Client and Registry in a known Docker™ architecture. Computing system 100 in FIG. 1 comprises Docker™ Host 110, Docker™ Registry 120 and Docker™ Client 130.

[0012] Docker™ Host 110 comprises Docker™ daemon 112, multiple containers 114 and a Docker™ Repository 116. Within Docker™ Repository 116 are multiple Images 118.

[0013] Docker™ Registry 120 holds images 122. Although some images 118 are shown in Docker™ Repository 116, these may for example be images that have been used previously by Docker™ Host 110. In operation, Docker™ Host 110 may require further images at some point in the future. Such an image may be selected from images 122 in Docker™ Registry 120 and delivered from Docker™ Registry 120 to Docker™ daemon 112. The image may then be added to those stored in Docker™ Repository 116 of Docker™ Host 110 for use by Docker™ Host 110.

[0014] Docker™ Host 110, Docker™ Registry 120 and Docker™ Client 130 can all be on the same server. Alternatively, Docker™ Host 110, Docker™ Registry 120 and Docker™ Client 130 can be located on separate servers. A Docker™ client 130 can exist locally on the Docker™ Host 110 or in the ‘cloud’, i.e. at a remote location on the internet.

[0015] Clients are used in a wide variety of systems, including other systems than the Docker™ system 100 of FIG. 1. In such systems, in general, a client could be a command line interface running in any terminal locally. In all cases, clients communicate to a host through an API (Application Programming Interface).

[0016] FIG. 2 illustrates a LAN system 200 that uses containers. LAN system 200 illustrates a simple and typical implementation of containers. In the system of FIG. 2, there are multiple Hosts 210. Each of the Hosts 210 is located on a separate server. Registry 220 is also on another, separate server.

[0017] Hosts 210 and Registry 220 of the LAN system 200 of FIG. 2 are connected by an Ethernet LAN 205. A Docker™ Client 230 accesses a Host 210 through an Application Programming Interface (API) 235. Thus the Docker™ Client 230 may be local to the LAN system 200, as illustrated in FIG. 2. However, alternatively, the Docker™ Client 230 may be remote from the Host 210 that it is accessing. Thus the location of the Docker™ Client 230 does not play a significant role in the operation of the system of FIG. 2.

[0018] Returning to the definitions of the ‘image’ and ‘container’, their relationship in known systems can be understood with reference to FIGS. 1 and 2. A container 114 is a runnable instance of an image 118. The image 122 from Registry 220 in FIG. 2 is passed around a network such as LAN 205 in downloads. The image 118 is stored in the Docker™ repository 116. The Docker™ daemon 112 running on the Docker™ Host 110 takes the image 118 and creates a container 114 in the Docker™ Host 110. When the software runs on the Docker™ Host 110, it is the container 114 that is used.

[0019] In the Docker™ system 200 of FIG. 2, a group of computing devices that are running containers can be joined

together to form a cluster. That cluster may be termed a ‘swarm’. The execution of the commands to control the operation of the swarm may be carried out by a swarm leader. There are several types of nodes in a Docker™ swarm. Worker nodes do not participate in swarm management. There can be one or more manager nodes in the swarm. There is only one swarm leader, which is the primary manager node that makes all swarm management and orchestration decisions for the swarm. The ‘orchestrator’ is another term used in known systems for the swarm leader.

[0020] The Docker™ service allows a swarm of selected Docker™ nodes, such as hosts 210, to work together. Those nodes run a defined number of instances of a replica task, which is itself a Docker™ image. The instances of the replica task are identical, but they may carry out different processes on different Hosts. With this arrangement, the swarm leader is responsible for orchestrating or controlling the swarm. The swarm leader is in effect a controlling entity that decides which hosts will run the container. The swarm leader can run on any node.

[0021] Whilst the swarm takes care of managing the services, the end result is a “Docker™ pull”. The Docker™ pull causes the Host 210 to get an image 122 from the Registry 220, if that image 122 is needed and is not already in the Docker™ Repository 116 of that Registry 220.

[0022] In the LAN system 200 of FIG. 2, each Host 210 will operate by requesting an image separately when that image is not held by that Host 210. The request passes from the particular host 210 over LAN 205 to Registry 220. Each Host 210 may be equipped with its own repository 116, where the image is stored after download. The download of an image from Registry 220 will occur separately to each Host 210 that has requested that image. Potentially those downloads may occur in parallel.

[0023] An edge computing system may be implemented as a mesh network. A mesh network typically does not have a high bandwidth link such as LAN 205 of FIG. 2 available. Instead, some nodes of the mesh network will have direct wireless links between each other. The approaches that are used in LAN system 200 of FIG. 2 may not work optimally in such a mesh network. Hence a need exists to provide an edge computing system as a mesh network that can implement virtualization.

SUMMARY OF THE INVENTION

[0024] In accordance with a first aspect of the present invention, there is provided a method in accordance with appended claim 1. In accordance with a second aspect of the present invention, there is provided a method in accordance with appended claim 11. In accordance with a third aspect of the present invention, there is provided a mesh computing system in accordance with appended claim 18. In accordance with a fourth aspect of the present invention, there is provided a mesh computing system in accordance with appended claim 20. The dependent claims provide further steps and features of embodiments.

BRIEF DESCRIPTION OF THE DRAWINGS

[0025] Exemplary embodiments of the present invention will now be described, by way of example only, with reference to the accompanying drawings, in which:

[0026] FIG. 1 shows an example of a known Docker™ architecture, in schematic form.

[0027] FIG. 2 shows an example of a known LAN architecture using containers, in schematic form.

[0028] FIG. 3 illustrates the architecture of a known mesh computing system.

[0029] FIG. 4 illustrates the architecture of a known edge node and a host.

[0030] FIG. 5 illustrates the architecture of a mesh computing system in accordance with an embodiment.

[0031] FIG. 6 illustrates an embodiment of an architecture of an edge node, and a host of the embodiment of FIG. 5.

[0032] FIG. 7 illustrates an architecture of other edge nodes of FIG. 5.

[0033] FIG. 8 is a flow diagram of a method of operation of the known mesh computing system of FIG. 3.

[0034] FIG. 9 provides an illustration of a method of operation of a mesh computing system in accordance with an embodiment of the invention.

[0035] FIG. 10 provides an illustration of a method of operation a registry agent in accordance with an embodiment of the invention.

[0036] FIG. 11 provides an illustration of a method of operation of a registry agent in accordance with another embodiment of the invention.

[0037] FIG. 12 is a flow diagram of a method of downloading an image in accordance with an embodiment of the invention.

[0038] FIG. 13 provides an illustration of another method of operation in accordance with an embodiment of the invention.

[0039] FIG. 14 illustrates a method in accordance with an embodiment.

[0040] FIG. 15 illustrates an embodiment of a server of the edge node of the invention.

[0041] FIG. 16 illustrates a high level overview of a method of operation of an embodiment of the invention.

DETAILED DESCRIPTION

[0042] The parallel downloading of images 122 in the LAN system 200 of FIG. 2 may not provide a particular drain on resources when the system is implemented with a high capacity LAN 205. However, a mesh network may not be able to deliver comparable performance to that achievable with a LAN, when attempts are made to implement known virtualization systems on mesh computing systems. In particular, problems may arise when trying to move images through the mesh network.

[0043] The present invention addresses systems of hosts where there are constraints on the bandwidth of connections between the nodes in the system. When the hosts are not interconnected with a LAN, they may instead be connected by a variety of links that have significantly constrained bandwidths. In particular, a mesh network is an edge computing system that has wireless links between its nodes. The mesh network may therefore have links between nodes provided by interconnections that show significantly limited bandwidth. The bandwidth limitations may affect transmissions over the interconnections at all times, or just at particular times.

[0044] Communication devices in a mesh network communicate using mesh transport technology. That technology is typically either WiFi™ or a cellular link. Both technologies may be employed within one mesh network, for different interconnections between nodes.

[0045] Mesh routers may be provided to support routing within the mesh network. Mesh router devices provide routing functionality for data that originates from other devices, which may either be an end node or another mesh router. Such mesh routers can relay traffic from an end node or another mesh router towards an edge router. Mesh router devices also, typically, obtain an IPv6 prefix (address) from router advertisements sent from the edge router.

[0046] An attempt to implement a mesh computing system using containerisation may result in appreciable bandwidth demand over some of the links in the mesh. The bandwidth demand may also be variable, i.e. different at different times, which makes the demand unpredictable. Although the use of containers is lightweight in terms of resources compared to Virtual Machines (VM), nevertheless the size of such images can be large. The size of the images will lead to appreciable bandwidth demand in any links in the mesh network that have significant bandwidth limits.

[0047] FIG. 3 illustrates the architecture of a known mesh computing system 300, in which some links have significant bandwidth constraints. Mesh computing system 300 is illustrated as comprising seven edge nodes, each of which acts as a host. The first edge node has reference 310. Mesh computing system 300 also comprises a second edge node 312, a third edge node 313, a fourth edge node 314, a fifth edge node 315, a sixth edge node 316, and finally a seventh edge node 317.

[0048] First edge link 320 connects first edge node 310 and third edge node 313. The edge nodes are located in a portion of FIG. 3 that is generally labelled as the 'Edge', and which lies to the left of dividing line 330. Some or all of the edge nodes may correspond to Docker™ Host 110 in FIG. 1.

[0049] First edge node 310 is also connected via second link 340 to the internet, which is generally illustrated by global internet 350. Global internet 350 is located in a portion of FIG. 3 that is generally labelled as the 'Cloud', and which lies to the right of dividing line 330. Located within global internet 350 is central registry 360. Central registry 360 may correspond to Docker™ Registry 120 in FIG. 1. Second link 340 may take the form of a standard backhaul link. Central registry 360 is on a server in global internet 350. Central registry 360 holds images 332, 334, 336 and 338. Central registry 360 may in fact hold thousands of images.

[0050] Client 325 is shown directly linked to fifth edge node 315 by third link 327. In this example, therefore, client 325 is located on a separate server, i.e. not on fifth edge node 315. However, client 325 could be on the same server as fifth edge node 315. Alternatively, client 325 could be located on the same server as any of the other nodes. As a further alternative, client 325 could be located on a server in global internet 350.

[0051] In operation of an edge computing system such as that in FIG. 3, any of the seven edge nodes may require an image 332 that it does not already hold. Each edge node needs to download the required image 332 separately, i.e. independently of the activities of other nodes. The edge node that requires the image 332 will request the required image 332 from central registry 360.

[0052] A request by first edge node 310 for an image 332 will pass across second link 340. Second link 340 will then also return the sought image 332 from central registry 360 to first edge node 310. When any of the other six edge nodes

requests an image 332, the request will pass through one or more links such as first link 320, before reaching first edge node 310. Then the request will pass across second link 340 to global internet 350. Second link 340 will then return the sought image 332 from central registry 360 to first node 310, and the image 332 will pass to the requesting node through one or more links such as first link 320.

[0053] The size of any image 332 requested from Central registry 360 will vary, based on the container. For example, image sizes might be in the range of 100-400 MB for 'Ubuntu™' based containers. Ubuntu™ is an open source Linux™ language. The image size may, however, drop to about 50-60 MB for 'Alpine' containers. Alpine is another example of an open source Linux™ language.

[0054] Image sizes in the range 50-400 MB would not cause appreciable bandwidth demand in the known LAN system 200 of FIG. 2. Similarly, there may not be significant bandwidth demand over some high bandwidth links within mesh computing system 300, and hence no significant problems with transferring such images. However, when any links between nodes, such as first edge link 320 in FIG. 3, are constrained communication links, there will be bandwidth demand that may require a significant proportion of all the available bandwidth. There may well also be significant additional bandwidth demand over second link 340. For example, second link 340 may simultaneously need to download image 332 to several edge nodes, and image 334 to other nodes, and image 336 to other nodes, with each individual download requiring the image 332, 334 or 336 to pass in its entirety over second link 340.

[0055] Seventh edge node 317 can only receive an image 332 from central registry 360 if that image 332 passes over several links within the mesh. If any one of those links suffers demand for a significant proportion of its bandwidth, then the operation of seventh edge node 317 will be correspondingly affected.

[0056] FIG. 4 illustrates the architecture of a known edge node and a host. First edge node 310 and central registry 360 correspond to first edge node 310 and central registry 360 of FIG. 3.

[0057] A client 385 is shown connected to first edge node 310 via client link 395. Host 370 is located on first edge node 310. Within host 370, repository 380 acts to store received images. Link 390 provides for communication to and from central registry 360. Host 370 can request images from central registry 360 via link 390, and receive images directly into host 370 from central registry 360 via link 390. Client link 395 provides for communication between client 385 and host 370.

[0058] In known systems, host 370 will request an image from central registry 360 when that host 370 needs the image. That need arises when client 385 makes a first request, i.e. for that image. When the image arrives from central registry 360, host 370 places the image in its repository 380. Subsequently, client 385 may make a second request for an image. If host 370 finds that it can meet the second request by supplying one of the images that is already stored in repository 380, then the host 370 supplies the image immediately to client 385. However, if host 370 finds that it cannot meet the second request by supplying one of the images that is already stored in repository 380, then host 370 will request, from central registry 360, the image that is needed to meet the second request of client 385.

[0059] Considering the mesh computing system 300 of FIG. 3 in more detail, a further issue arises. When a particular image is to be downloaded from central registry 360 to seventh edge node 317, then that image will have to be ‘daisy chained’ through first edge node 310, second edge node 312, fourth edge node 314 and fifth edge node 315. The download of that image will be in a transmission that is not open to inspection by the intervening first edge node 310, second edge node 312, fourth edge node 314 and fifth edge node 315. Thus bandwidth will be consumed by the download of the image in each of the links between first edge node 310, second edge node 312, fourth edge node 314 and fifth edge node 315, without that consumption of bandwidth providing any advantage to those intervening nodes. The same image may at the same time, for example, have been requested separately by the fifth edge node 315, for its own use. In order to download the image from central registry 360 to the fifth edge node 315, the image will also have to be transmitted again from central registry 360 via first edge node 310, second edge node 312, and fourth edge node 314. The result in this situation is that the same image is transferred multiple times through the same links of the mesh computing system 300, possibly simultaneously. This results in high traffic on some links and through some nodes. The worst affected nodes are likely to be those closest to first node 310 and second link 340, the backhaul link. Second link 340 is likely to have the very highest load.

[0060] Further differences between a mesh computing system and a LAN system can be recognized. If a service is distributed to multiple nodes in the mesh computing system 300, all nodes will use the central registry 360 to obtain an image. When making a request for an image from central registry 360, any edge node in the mesh computing system 300 will then suffer bandwidth demand that may take a significant proportion of its available bandwidth. That demand may vary with time. The mesh will be bandwidth restricted to an extent that depends on the technology and the topology, and potentially on interference in the case of wireless links.

[0061] A comparison of the mesh computing system 300 of FIG. 3 and the LAN system 200 of FIG. 2 reveals the differences shown in summary in table 1 below:

TABLE 1

Differences between mesh computing and LAN systems	
Issue	Difference exhibited by mesh computing system 300 over LAN system 200
Architecture- Multiple links	Any image downloaded in mesh computing system 300 will potentially be passed over multiple ‘hops’ (links) on its way to its destination node, i.e. the host that requests it. The further the host is from the first node 310 and the second link 340, then the larger the number of links that the image must cross.
Technology- bandwidth per link	Each of the links in the mesh computing system 300 may be a constrained communication link, for example a cellular link or a Wi-Fi link.
Technology- interference	Cellular or Wi-fi™ links in mesh computing system 300 may suffer interference.
Process- simultaneous parallel downloads	When more than one node in the same part of mesh computing system 300 requests the same image at the same time, that image may be downloaded multiple times over the same link, or nearby links, simultaneously. This is not a significant issue within LAN 205 of FIG. 2.
Backhaul	The backhaul over second link 340 to the global internet 350 may also be a constrained communication link, for example a cellular link or a Wi-Fi link. This causes a significant bottleneck when an image is downloaded to each host from central registry 360 over this single link.

[0062] In mesh computing system 300 there are a wide variety of different usage scenarios. These different usage scenarios complicate the design of a mesh system that is to demonstrate acceptable bandwidth performance. In particular, in some usage scenarios, not all the edge nodes require a particular image. There are three different scenarios.

[0063] In the first scenario, the image does need to be downloaded to all hosts. This places large bandwidth demands across many or all of the links in mesh computing system 300, simultaneously. In the second scenario, the image needs to be downloaded only to specific hosts, i.e. one or more edge nodes. In the third scenario, the image needs to be downloaded to one or more ‘non-specific’ hosts. This third scenario may arise when a service is required to run, but it does not matter which particular host in mesh computing system 300 runs the service.

[0064] The present invention provides a mesh computing system in which hosts, i.e. edge nodes, can download an image from another host, e.g. a neighbouring host. With this approach, some of the disadvantages of the mesh computing system 300 can be overcome. A reduction in the need to download images from the central registry can thereby be achieved. Various embodiments of the invention, including various methods for achieving the invention, are outlined below.

[0065] FIG. 5 provides a schematic illustration of the architecture of a modified mesh computing system 400 of an embodiment of the invention. In the modified mesh computing system 400, some functions that were previously carried out by the central registry 360 of mesh computing system 300 have been distributed among the hosts in mesh computing system 400. In addition, in accordance with various embodiments of the invention, further functions may now be carried out by the various distributed registries of mesh computing system 400.

[0066] The various registry elements described in connection with FIG. 5 are not necessarily identical to central registry 360 of FIG.3. Each registry element that is co-located with a host may only contain some images, which the host is storing locally. FIGS. 6 and 7 provide illustrative embodiments of two different versions of the registry ele-

ments that are co-located with various ones of the hosts of FIG. 5. FIG. 6 shows a registry that has full functionality. The registry of FIG. 6 may have capabilities that are comparable to those of central registry 490. FIG. 7 shows a registry located at other nodes of mesh computing system 400, which has more limited functionality than the registry illustrated in FIG. 6. Thus an embodiment of FIG. 5 may have a mix of different designs of registry at different nodes, with some of the registries having the capability of the registry of FIG. 6 and other registries at other nodes having only the capability of the registry of FIG. 7.

[0067] After the description below of the structure of mesh computing system 400 in connection with FIGS. 5-7, various ways of operating mesh computing system 400 are then described with reference to FIGS. 9-14 and 16.

[0068] In mesh computing system 400, some links have significant bandwidth constraints. Mesh computing system 400 is illustrated as comprising six edge nodes, each of which is home to a host. Each edge node may, in a practical embodiment, be located on a server that is connected by wireless links to other nodes.

[0069] Mesh computing system 400 comprises first edge node 420. First registry element 422 is provided within first edge node 420. Swarm leader 424 is also provided within first edge node 420. In an LTE™natively arrangement that is not shown in FIG. 4, first registry element 422 could be outside first edge node 420, but connected to first edge node 420 by a link. That link would be a high bandwidth link. The swarm leader 424 of first registry element 422 carries out functions that are outlined in greater detail in the discussion of the method of operation of mesh computing system 400 below, and in connection with FIG. 6.

[0070] Second edge node 432 comprises second registry element 434. Third edge node 436 comprises third registry element 438. Fourth edge node 440 comprises fourth registry element 442. Fifth edge node 444 comprises fifth registry element 446. Sixth edge node 448 comprises sixth registry element 450. The edge nodes are located in a portion of FIG. 5 that is generally labelled as the 'Edge', and which lies to the left of dividing line 460. The registry elements shown in FIG. 5 may not be identical. In some embodiments, any individual registry may only hold a subset of the set of images that are held in central registry 490. Each of the first to sixth registry elements shown in FIG. 5 contains those images that are being stored by the host with which the registry element is associated.

[0071] In LTE™natively arrangements that are not shown in FIG. 5, any of second registry element 434, third registry element 438, fourth registry element 442, fifth registry element 446 or sixth registry element 450 could be outside of the edge nodes of which they currently form part, and could be connected to their edge node by a high bandwidth link.

[0072] First edge link L1 connects first edge node 420 and second edge node 432. Second edge link L2 connects first edge node 420 and third edge node 436. Third edge link L3 connects second edge node 432 and fourth edge node 442. Fourth edge link L4 connects third edge node 436 and fourth edge node 440. Fifth edge link L5 connects fourth edge node 440 and fifth edge node 444. Sixth edge link L6 connects fourth edge node 440 and sixth edge node 448.

[0073] First edge node 420 is also connected via seventh edge link L7 to global internet 480. Global internet 480 is

located in a portion of FIG. 5 that is generally labelled as the 'Cloud', and which lies to the right of dividing line 460.

[0074] Located within global internet 480 is central registry element 490. Central registry element 490 is on a server in global internet 480. Central registry element 490 may correspond to Docker™ Registry element 120 in FIG. 1. Seventh link L7 may take the form of a backhaul link.

[0075] Client 455 is shown directly linked to third edge node 436 by client link 457. Client 455 is located on a separate server. However, client 455 could be on the same server as third edge node 436. ALTE™natively, client 455 could be located on the same server as any of the other nodes, or could be located on a server in global internet 480.

[0076] The downloading of images from central registry 490 to individual registries in accordance with the method of the invention is orchestrated by swarm leader 424. Although swarm leader 424 is located at first edge node 420 in FIG. 5, it may instead be located at another node in mesh computing system 400. Swarm leader 424 is the originator of various commands, which tell various hosts which images they should download to their registries from central registry 490. The images do not need to pass through swarm leader 424 en route to the registry to which they should pass. Swarm leader 424 thus carries out 'command and control' functions.

[0077] FIG. 6 illustrates an architecture of an edge node and a host of the mesh computing system 400 of FIG. 5, in accordance with an embodiment. Client 455, central registry 490, first edge node 420 and fourth edge node 440 correspond to those elements in FIG. 5.

[0078] Third edge node 436 of FIG. 5 is represented in detail as third edge node 600 in FIG. 6. Host 602 is located on third edge node 600. In FIG. 6, the registry element 438 of FIG. 5 now takes the form of a Local Registry 610 located on third edge node 600. Local Registry 610 comprises an Application Programming Interface (API) 606. Local Registry 610 also comprises a repository 604, located on host 602. The dotted line in FIG. 6 that encompasses both API 606 and repository 604 shows them as part of Local Registry 610.

[0079] Client link 657 provides for communication between client 455 and host 602. Within host 602, repository 604 acts to store received images, i.e. images that have been received at host 602.

[0080] In the embodiment of FIG. 6, host 602 communicates with API 606 via internal link 608 within third edge node 600. API 606 communicates with: (i) First edge node 420, via external link L2; (ii) Central registry 490, via external links L2 and L7; and (iii) Fourth edge node 440, via external link L4. Local Registry 610 enables the host 602 to receive images into repository 604 via API 606, and to supply images from repository 604. The presence of local registry 610 enables functionality described below with reference to FIGS. 9-12.

[0081] Local Registry 610 may have equivalent functionality to central registry 490 of FIG. 4. In this case, repository 604 may hold a comprehensive database of images, which is as extensive as that held in central registry 490. However, repository 604 may hold a subset of a set of images that is held in central registry 490.

[0082] In one exemplary embodiment, all the registry elements in mesh computing system 400 may take the form of local registry 610 of FIG. 6. In a further exemplary embodiment, only some of the registry elements in mesh

computing system 400 take the form of Local Registry 610 of FIG. 6. In such an embodiment, other registry elements in mesh computing system 400 comprise a more basic design of registry agent, which is adapted to retrieve images from a repository located at the host. The more basic design of registry agent is illustrated in FIG. 7, and discussed below.

[0083] In a typical configuration of the invention, only a few nodes in mesh computing system 400 will be configured as shown for host 600 and local registry 610 of FIG. 6. A larger number of nodes will be configured with the more basic registry agent of FIG. 7.

[0084] In operation, mesh computing system 400 of FIG. 5 is able to pre-emptively position images from central repository 490 in local registry 610 of host 600. Swarm leader 424 makes the decisions about which images to pre-position in which local registries located in local registries of the form of local registry 610, and then sends out appropriate commands to the nodes on which those local registries are located. Those images are then available for supply from local registry 610 to other hosts of mesh computing system 400. The supply of the images from local registry 610 to other hosts of mesh computing system 400 is effected through API 606. However, swarm leader 424 does not pre-position images in the more basic local registries of the form illustrated in FIG. 7. The more basic registries do not have an API of the form of API 606 in FIG. 6, so are able to fetch an image from another node but not supply images to other nodes.

[0085] FIG. 7 illustrates the structure of the more basic registry agent. In the embodiment where only some of the registry elements in mesh computing system 400 take the form of Local Registry 610 of FIG. 6, the other registry elements in mesh computing system 400 comprise the more basic registry agent. In mesh computing system 400, fifth edge node 444, for example, may comprise the more basic registry agent illustrated in FIG. 7.

[0086] Edge node 700 of FIG. 7 comprises registry agent 738 and host 770. Repository 785 is provided on host 770. Link 784 connects registry agent 738 and host 770. Link L7 and central registry 490 correspond to link L7 and central registry 490 of FIG. 5. Fourth edge node 442 corresponds to fourth edge node 442 of FIG. 5.

[0087] In the illustration of FIG. 7, a client 755 is linked by link 757 to registry agent 738. Registry agent 738 is linked by link L5 to fourth edge node 442. L5 corresponds to link L5 on FIG. 5, and Fourth edge node 440 is shown linked by link L7 to central registry 490, although there are intermediate links between fourth edge node 440 and link L7 that have been omitted from FIG. 7 for ease of illustration.

[0088] Repository 785 is configured to hold images for the use of edge node 700. Registry agent 738 is configured to allow edge node 700 to request images directly from fourth edge node 440, central registry 490, or from other edge nodes of mesh computing system 400.

[0089] When host 770 receives a request from client 755 for a particular container, the registry agent 738 will first check whether the particular image is already in repository 785. When the particular image is not in repository 785, then registry agent 738 carries out a discovery process to decide from where host 770 can obtain the particular image. Host 770 may find that the particular image has been pre-positioned in the registry element 610 of first edge node 422 or another node with a registry corresponding to that illustrated in FIG. 6. In that case, registry agent 738 can request the

particular image from the relevant API, such as API 606. Each host, such as host 770, is able to discover which of the hosts with a registry corresponding to that illustrated in FIG. 6 already has the particular image.

[0090] This method can be accomplished by using a service discovery method to find a repository at any host that already has the particular image. The selected host that has and can supply the sought image will, preferably, be close to host 770 that is requesting the image. This can be viewed as a "repository as a service" function provided by the hosts that have an API such as API 606. As part of the service, a cost metric can be calculated. The cost metric indicates the resources that would be required to download a sought image to requesting host 770, from any particular neighbouring host.

[0091] Before further discussion of the method of operation of the mesh computing system of FIG. 5 and the edge nodes and hosts of FIGS. 6 and 7, FIG. 8 provides an explanation of how the known system of FIG. 3 operates.

[0092] FIG. 8 is a flow diagram of a method of operation of the mesh computing system 300 of FIG. 3. FIG. 8 will allow a clear comparison with the embodiments of the method of the invention that are described subsequently with reference to FIGS. 9-13.

[0093] FIG. 8 illustrates a flow diagram for a method of image download as would occur on mesh computing system 300 in FIG. 3. Method 800 is a method that would be used if containerization were implemented on mesh computing system 300 without using the invention. The references used in the discussion of method 800, below, are from FIG. 4.

[0094] Method 800 starts at step 810. At 820, the host 370 at edge node 310 is idling, i.e. waiting for an event. At step 830, client 385 has requested an image, which results in the generation of a command to 'pull' the image.

[0095] In step 840, host 370 ascertains whether the image is already held locally in the repository 380 of the host 370 itself. If the answer is 'No', then host 370 has to download the image from Central Registry 360 of FIG. 3. Then method 800 moves to step 860. If the answer is 'Yes' in response to the question in step 840, then the method bypasses step 850 and moves directly to step 860. At step 860, host 370 uses the image to create a container, in order to address the request from client 385. After step 860, method 800 returns to idle at step 820 until a new request is received from a client.

[0096] FIG. 9 is a flow diagram of a method of operation of the mesh computing system of FIG. 5, with reference to host 602 and local repository 610 of FIG. 6

[0097] Method 900 begins at step 910. At step 920, the host 602 receives a request from client 455 for an image. At step 930, the host 602 sends a request for the image to at least one other host of the mesh computing system 400.

[0098] When the host 602 receives 940 notification that at least a second host holds the image, method 900 moves to step 960. At step 960, the host 602 downloads the image from the second host to the host 602. However, when the host 602 does not receive 940 notification that at least a second host holds the image, method 900 moves to step 950. At step 950, the host 602 downloads the image from the central registry 490 to the host 602. After either step 950 or 960, host 470 creates 970 a container from the downloaded image.

[0099] With the method of FIG. 9 and the local registry 610 in FIG. 6, the first host 470 may be able to reduce

substantially the number of requests that it makes for images from central registry 490. Many requests from clients for images will be met by implementing step 960, i.e. downloading the image from another host in the mesh computing system 400. Downloading the image from another host in the mesh computing system 400 reduces traffic in the mesh computing system 400. When the image is downloaded from another host to host 602, there will be no request to central registry 490 and hence there is a reduction in traffic on seventh link L7, which is the backhaul link.

[0100] FIG. 10 provides an illustration of a method of operation 1000 of the swarm leader 424 of FIG. 5 in accordance with an embodiment of the invention.

[0101] At step 1010, the method 1000 starts. At step 1020, swarm leader 424 identifies a need in a first host at a first edge node of the mesh computing system to run an image, the image being held in the central registry 490. At step 1030, swarm leader 424 identifies a second host at a second edge node of the mesh computing system. The first and second hosts in method 1000 of FIG. 10 might be any of the hosts in mesh computing system 400 of FIG. 5. The second host would be constructed as third edge node 600 of FIG. 6. Swarm leader 424 may identify the second host as being a host that does not need to run the image. The need to run the image in the first host is a predicted need. Swarm leader 424 may identify the predicted need, when swarm leader 424 considers the first host as being likely to need to use the first image at some future timepoint.

[0102] At step 1040, swarm leader 424 commands download of the image from central registry 490 to the second host. The download occurs via at least one wireless link of the mesh computing system 400.

[0103] At step 1050, the first host receives a request from a client for the image. This request demonstrates that the prediction by swarm leader 424 of the need for the image was correct. At step 1060, the first host performs discovery to locate the image. The first host will discover that the first image is available on the second host.

[0104] At step 1070, the first host downloads the image from the second host. Such a download may involve download over few wireless links of the mesh network of the mesh computing system 400. The download step 1070 would therefore entail less bandwidth demand, in comparison with typical downloads performed in step 850 of method 800 in known mesh computing system 300.

[0105] FIG. 11 provides an illustration of another method of operation 1100 in accordance with another embodiment of the invention. Method 1100 operates with first and second hosts, each of which has an image storage repository and an API. The first and second hosts in method 1100 of FIG. 11 might be any of the hosts in mesh computing system 400 of FIG. 5, provided that they are constructed as third edge node 600 of FIG. 6.

[0106] At step 1110, method 1100 starts. At step 1120, the swarm leader identifies an image from the multiple images in the central registry 490 as being an image that is likely to be used by other hosts in the mesh computing system, than the first and second hosts.

[0107] At step 1130, under control of swarm leader 424, the first host downloads the image from central registry 490 to the first repository of the first host. At step 1140, also under control of swarm leader 424, the second host downloads the image from central registry 490 to the second repository of the second host.

[0108] At step 1150, a third host of the plurality of hosts receives a request from a client for the image. Step 1150 is shown dotted, as it is possible that step 1150 may occur much later than steps 1110-1140.

[0109] At step 1160, the third host performs discovery to locate the image. The image is now available from, each of the first and second hosts. At step 1170, the third host downloads the image from the first host or the second host. The decision about which of the first or the second hosts should provide the image may be based on a cost function, which is discussed further below. At step 1180, the third host creates a container from the image, and supplies the container to the requesting client.

[0110] FIG. 12 is a flow diagram of a method 1200 of downloading an image that may be comprised in step 960 of method 900, or step 1070 of method 1000.

[0111] At step 1210, the first host receives notification that at least a second host holds the image. At step 1220, the first host calculates a first value of a cost metric, for a first route over which the image could be downloaded to the first host. At step 1230, the first host calculates a second value of the cost metric, for a second route over which the image could be downloaded to the first host. Each route may comprise one or more of the links and corresponding edge nodes in FIG. 5.

[0112] At step 1240, the first host compares the first value of the cost metric and the second value of the cost metric, to identify the lower of the first value of the cost metric and the second value of the cost metric. At step 1250, when the first value of the cost metric is lower, the first host selects the first route. However, when the second value of the cost metric is lower, the first host selects the second route. At step 1260, the first host downloads the image over wireless links of the selected route, to the first host.

[0113] FIG. 13 is a flow diagram of a method 1300 of making the calculation in step 1220 or step 1230 of method 1200.

[0114] Method 1300 evaluates a situation where each of the first route and the second route comprises one or more wireless links. At step 1310, for each route, the cost metric calculation uses at least one of the following parameters: (i) a number of links over which the image must be transmitted; (ii) a communication link quality of each of the links; (iii) a loading on any intermediate edge nodes on the route that would result from downloading the image; (iv) a loading of each link; and/or (v) a financial cost of transmitting on each link.

[0115] At step 1320, when a route comprises two or more links, the cost metric is calculated in accordance with any of (ii)-(v) by evaluating the parameter for each link in the route; and combining the evaluated parameters for each link in the route to calculate a total cost metric for the route.

[0116] Method 1300 is illustrated for a situation where two routes, for example, start at the same second node. For the sake of completeness, the method 1300 can also be applied when the same desired image is available at both a second host and a third host. In this case, when at least a second host and a third host holds the image, method 1300 calculates: (i) the first value of the cost metric, for a first route over which the image could be downloaded to the first host from the second host; and (ii) the second value of the cost metric, for a second route over which the image could be downloaded to the first host from the third host. When the first value of the cost metric is lower, method 1300 selects the second host

to provide the image over the first route. When the second value of the cost metric is lower, method 1300 selects the third host to provide the image over the second route.

[0117] Method 1300 may, therefore provide selection of the route for downloading an image from a host, and/or may provide selection of the host from which to download the image.

[0118] FIG. 14 illustrates a method in accordance with an embodiment. Method 1400 should be read in conjunction with FIG. 6. At step 1410, in response to receiving the request to run the image, host 602 generates a request for the image from central registry 490.

[0119] At step 1420, API 606 intercepts and overrides the request for the image from central registry 490, before transmission of the request for the image. At step 1430, API 606 generates and sends the request for the image to at least one other host of the mesh computing system.

[0120] Two results are possible, as shown at step 1440. Host 602 may receive notification that a second host holds the image, in which case the method proceeds to step 1450. Alternatively, host 602 does not receive notification that any second host holds the image, in which case the method proceeds to step 1460.

[0121] At step 1450, when the first host receives notification that the second host holds the image, API 606 generates a command to download the image from the second host.

[0122] At step 1460, when the first host does not receive notification that the second host holds the image, the API sends a request for the image to the central registry 490, and downloads the image from the central registry 490. Step 1460 is the alternative to that shown in step 1450.

[0123] Steps 1420, 1430 and 1440 of method 1400 may be implemented by a plug-in software module of host 602. Such a plug-in software module would provide programming steps for a logic module of API 606 to perform the steps of: (i) intercepting and overriding the request for the image from host 602 to central registry 490; (ii) generating the request for the image, from the at least one other host; and (iii) generating the command to download the image from the second host.

[0124] In step 1440 of method 1400, and in other embodiments of the invention, the host 602 may receive the image from the second host using tar or ftp. Following step 1440 of method 1400, the host 602 may store the image in the repository 604 of host 602.

[0125] Subsequently to performing method 1400, host 602 may receive a request via API 706 from another host in the mesh network, for the image that is now stored in the repository 704 of host 702. API 706 can respond to a pull command from another host in the mesh network by subsequently accessing the image from repository 704 via an API pull command, and supplying the image to the requesting host.

[0126] FIG. 15 illustrates an embodiment of a server 1500 of the edge node of the mesh computing system of FIG. 4.

[0127] An antenna 1502, for receiving/transmitting cellular transmissions 1522, is provided. Transmission 1526 may, for example, be a WiFi™ or an LTE™ transmission to another edge node. The antenna 1502 is coupled to an antenna switch or duplexer 1504, which provides isolation between receive and transmit chains. One or more receiver chains, as known in the art, include receiver front-end circuitry 1506. Receiver front-end circuitry 1506 provides

reception, filtering and intermediate or base-band frequency conversion. Receiver front-end circuitry 1506 is coupled to a signal processor 1528, which may be realized by a digital signal processor (DSP).

[0128] A controller 1514 maintains overall operational control of the server 1500. Controller 1514 is coupled to receiver front-end circuitry 1506 and signal processor 1528. In some examples, the controller 1514 is also coupled to at least one memory, such as random access memory 1517 configured to store data content, and to a memory device 1516. Memory device 1516 selectively stores operating regimes, such as decoding/encoding functions, synchronization patterns, code sequences, and the like. A timer 1518 is operably coupled to the controller 1514 to control the timing of operations, such as transmission or reception of time-dependent signals, within server 1500.

[0129] As regards the transmit chain, the signal processor 1528 may process and formulate data for transmission to transmitter/modulation circuitry 1522 and thereafter a power amplifier 1524 and on to the antenna 1502 for transmitting the data. The transmitter/modulation circuitry 1522 and the power amplifier 1524 are operationally responsive to the controller 1514.

[0130] In accordance with example embodiments, signal processor 1528 of server 1500 has been configured to support WiFi™ operation. The various components within server 1500 can be realized in discrete or integrated component form, with an ultimate structure therefore being an application-specific or design selection.

[0131] In accordance with some example embodiments, the receiver 1506 of the server 1500 may be configured to receive signals via antenna 1502 and store data content when connected to other edge nodes, for example, or an Internet Protocol network. Server 1500 may also be configured to transmit retrieved data content to one or multiple mobile wireless communication units of mesh computing system 400, via antenna 1502, for example by WiFi. The at least one memory device 1516, operably coupled to the at least one receiver 1506, may be configured to store the received data content from other edge nodes and/or the mobile wireless communication units.

[0132] Thus FIGS. 5-7 and 15 provide details of a mesh network in accordance with the invention. As described above, the mesh network is configured to pre-emptively position images at hosts, for subsequent retrieval from the hosts and for the creation of containers from the images. The mesh network comprises hosts located in edge nodes of the mesh computing system. The edge computing system also comprises a central registry holding the images, the central registry located outside the mesh computing system. A swarm leader located in the mesh computing system is configured to identify a need in a first host at a first edge node of the mesh computing system to run an image, the image being held in the central registry. The swarm leader is also configured to identify a second host at a second edge node of the mesh computing system, and command download of the image from the central registry to the second host, via at least one wireless link of the mesh computing system

[0133] FIG. 16 illustrates a high level overview of a method of operation of an embodiment of the invention.

[0134] Method 1600 illustrates an embodiment where a swarm leader instigates supply of an image to a client, rather than a node supplying an image directly in response to a

request received at the node from a client. Method **1600** starts at step **1610**. At **1620**, a host at a node of the mesh computing system is in an idle or waiting mode. At step **1630**, the host receives a ‘pull’ command. The ‘pull’ command may be issued by a swarm leader within the mesh computing system. The swarm leader may control a plurality of nodes in the mesh computing system.

[**0135**] At step **1640**, a check is made whether the image is already available in the registry of the host. If the image is available, method **1600** moves to step **1690**. At step **1690**, a container is made from the image. The container is then supplied to a requesting host.

[**0136**] If, at step **1640**, the image is not available, then method **1600** moves to step **1650**. At step **1650**, the host send a service request. At step **1660**, a decision is made whether the image is available from at least one other host. If the image is available, the method moves to step **1680**. At step **1680**, cost metrics are calculated for the routes over which the image could be downloaded. Then method **1600** moves to step **1690**.

[**0137**] At step **1660**, if the image is not available from at least one other host, the method moves to step **1670**. At step **1670**, the host downloads the image from the central registry **490**, and moves to step **1680**.

[**0138**] As is clear from the preceding methods and apparatus, the invention provides options for pre-positioning and retrieving images from particular nodes within a mesh computing system. In comparison with known approaches, the overall traffic within the mesh computing system may be reduced. In particular, the traffic over links such as backhaul link **L7** to the central registry **490** will be lower when the invention is employed.

[**0139**] It will be further appreciated that, for clarity purposes, the described embodiments of the invention with reference to different functional units and processors may be modified or re-configured with any suitable distribution of functionality between different functional units or processors, without detracting from the invention. For example, functionality illustrated to be performed by separate processors or controllers may be performed by the same processor or controller. Hence, references to specific functional units are only to be seen as references to suitable means for providing the described functionality, rather than indicative of a strict logical or physical structure or organization.

[**0140**] Aspects of the invention may be implemented in any suitable form including hardware, software, firmware or any combination of these. The invention may optionally be implemented, at least partly, as computer software running on one or more data processors and/or digital signal processors. For example, the software may reside on non-transitory computer program product comprising executable program code to increase coverage in a wireless communication system.

[**0141**] Thus, the elements and components of an embodiment of the invention may be physically, functionally and logically implemented in any suitable way. Indeed, the functionality may be implemented in a single unit, in a plurality of units or as part of other functional units. Those skilled in the art will recognize that the functional blocks and/or logic elements herein described may be implemented in an integrated circuit for incorporation into one or more of the communication units.

[**0142**] Furthermore, it is intended that boundaries between logic blocks are merely illustrative and that aLTE™rnative

embodiments may merge logic blocks or circuit elements or impose an aLTE™rnative composition of functionality upon various logic blocks or circuit elements. It is further intended that the edge computing system and its elements depicted herein are merely exemplary, and that in fact many other edge computing systems and elements or architectures can be implemented that achieve the same functionality.

[**0143**] Although the present invention has been described in connection with some example embodiments, it is not intended to be limited to the specific form set forth herein. Rather, the scope of the present invention is limited only by the accompanying claims. Additionally, although a feature may appear to be described in connection with particular embodiments, one skilled in the art would recognize that various features of the described embodiments may be combined in accordance with the invention. In the claims, the term ‘comprising’ does not exclude the presence of other elements or steps.

[**0144**] Furthermore, although individually listed, a plurality of means, elements or method steps may be implemented by, for example, a single unit or processor. Additionally, although individual features may be included in different claims, these may possibly be advantageously combined, and the inclusion in different claims does not imply that a combination of features is not feasible and/or advantageous. Also, the inclusion of a feature in one category of claim does not imply a limitation to this category, but rather indicates that the feature is equally applicable to other claim categories, as appropriate.

[**0145**] Furthermore, the order of features in the claims does not imply any specific order in which the features must be performed and in particular the order of individual steps in a method claim does not imply that the steps must be performed in this order. Rather, the steps may be performed in any suitable order. In addition, singular references do not exclude a plurality. Thus, references to ‘a’, ‘an’, ‘first’, ‘second’, etc. do not preclude a plurality.

1. A method (**1000**), in a mesh computing system, for pre-emptively positioning images at hosts, for subsequent retrieval from the hosts and for a creation of containers from the images, the hosts located in edge nodes of the mesh computing system, the mesh computing system operably coupled to a central registry holding the images, the method (**1000**) comprising, at a swarm leader located in the mesh computing system:

identifying (**1020**) a need in a first host at a first edge node of the mesh computing system to run an image, the image being held in the central registry;

identifying (**1030**) a second host at a second edge node of the mesh computing system; and

commanding (**1040**) a download of the image from the central registry to the second host, via at least one wireless link of the mesh computing system.

2. The method (**1000**) according to claim 1, further comprising at the swarm leader: identifying the second host as not needing to run the image.

3. The method (**1000**) according to claim 1, further comprising:

commanding a download (**1070**) of the image to the first host, from the second host; or

commanding a download (**1170**) of the image from the first host via one or more further wireless links to at least one of: a third host and a fourth host located at an edge node of the mesh computing system,

whereby the image is cascaded from the central registry to the at least one of: the third host and the fourth host.

4. The method (1000) according to claim 1, wherein the second edge node of the mesh computing system is operably coupled via a backhaul link to the central registry, the backhaul link connecting the second host at the second edge node to the Internet.

5. The method (1000) according to claim 1, further comprising, at the first host:

using tape archive, tar, or a file transfer protocol, ftp, to download of the container image from the second host; and

storing the downloaded image in a repository, and subsequently accessing the image via an application programming interface, API, pull command.

6. The method (1000) according to claim 1, further comprising, at the second host:

storing the image in a repository of a registry; and

receiving a request for the image from the first host, and downloading (960; 1070) the image from the repository to a registry of the first host.

7. The method (1200) according to claim 1, further comprising at the first host:

sending (930) a request for the image to at least one other host of the mesh computing system;

receiving (1210) notification that at least the second host holds the image, and in response thereto calculating:

(i) a first value (1220) of a cost metric, for a first route over which the image could be downloaded to the first host; and

(ii) a second value (1230) of the cost metric, for a second route over which the image could be downloaded to the first host;

comparing (1240) the first value of the cost metric and the second value of the cost metric, to identify the lower of the first value of the cost metric and the second value of the cost metric; and

selecting the first route when (1250) the first value of the cost metric is lower;

selecting the second route when (1250) the second value of the cost metric is lower;

downloading (1260) the image over wireless links of the selected route.

8. The method (1300) according to claim 7, wherein each of the first route and the second route comprises one or more wireless links; and wherein for each route, the cost metric uses (1310) at least one of the following parameters:

(i) a number of links over which the image must be transmitted; and

(ii) a communication link quality of each of the links;

(iii) a loading on any intermediate edge nodes on the route that would result from downloading the image;

(iv) a loading of each link;

(v) a financial cost of transmitting on each link.

9. The method (1300) according to claim 8, wherein in response to identifying that a route comprises two or more links, the method further comprises calculating (1320) the cost metric in accordance with any of (ii)-(v) by:

evaluating the parameter for each link in the route; and

combining the evaluated parameters for each link in the route to calculate a total cost metric for the route.

10. The method (1300) according to claim 8, further comprising:

when at least a third host also holds the image, calculating a third value of the cost metric, for a third route over which the image could be downloaded to the first host from the third host; and

selecting the third host to provide the image over the third route, when the third value of the cost metric is lowest.

11. A method (1100) of pre-emptively positioning images at hosts in a mesh computing system, wherein the mesh computing system is operably coupled to a central registry that holds multiple images for use by the hosts of the mesh computing system, wherein the mesh computing system comprises:

a swarm leader;

a plurality of hosts located at a respective edge node in the mesh computing system, the plurality of hosts comprising:

a first host that has a first registry comprising a first image storage repository and a first application programming interface, API; and

a second host that has a second registry comprising a second image storage repository and a second API;

the method (1100) comprising:

at the swarm leader, identifying (1120) an image, from the multiple images, as being an image that is likely to be used by hosts in the mesh computing system other than the first host and second host;

at the first host, downloading (1130) the image from the central registry to the first image storage repository;

at the second host, downloading (1140) the image from the central registry to the second image storage repository; and

at a third host of the plurality of hosts, downloading (1170) the image from the registry of the first host or from the registry of the second host; and

at the third host, creating (1180) a container from the downloaded image.

12. The method (1100) according to claim 11, further comprising, at the swarm leader:

identifying the first host as being unlikely to need to use the first image; and/or

identifying the second host as being unlikely to need to use the first image; and

commanding a download of the image to the first host and the second host.

13. The method (1100) according to claim 12, further comprising the swarm leader identifying:

the first host and/or the second host as being located at a point in the mesh computing system proximal to a backhaul link between the mesh computing system and the central registry; and

subsequent to the swarm leader downloading the image from the central registry to the first image storage repository and to the second image storage repository, the method comprises, at the swarm leader:

commanding download of the image, from at least one of the first image storage repository and the second image storage repository, to a fourth image storage repository of a fourth host located at a fourth edge node in the mesh computing system; and

cascading the image to the fourth host located further from the central registry than the first host and the second host.

14. The method (1100) according to claim 11, further comprising at the swarm leader:

repeating the method (1100) of claim 11 for a second and a third image,

whereby the first image storage repository and the second image storage repository each hold a respective subset of a set of images held by the central registry, and each subset is available for future requests from other hosts in the mesh computing system that neighbour the first host and/or the second host.

15. The method (1200) according to claim 11, further comprising, at the third host:

sending (930) a request for the image to the first host of the mesh computing system;

receiving (1210) notification that at least the first host holds the image, and

calculating:

(i) a first value (1220) of a cost metric, for a first route over which the image could be downloaded to the third host; and

(ii) a second value (1230) of the cost metric, for a second route over which the image could be downloaded to the third host;

comparing (1240) the first value of the cost metric and the second value of the cost metric, to identify the lower of the first value of the cost metric and the second value of the cost metric; and

selecting the first route when (1250) the first value of the cost metric is lower;

selecting the second route when (1250) the second value of the cost metric is lower;

downloading (1260) the image over at least one wireless link of the selected route to the first host.

16. The method (1300) according to claim 15, wherein each of the first route and the second route comprises one or more wireless links and wherein for each route, the cost metric uses (1310) at least one of the following parameters:

(i) a number of links over which the image must be transmitted; and

(ii) a communication link quality of each of the links;

(iii) a loading on any intermediate edge nodes on the route that would result from downloading the image;

(iv) a loading of each link;

(v) a financial cost of transmitting on each link.

17. The method (1300) according to claim 16, wherein when a route comprises two or more links, the method further comprises:

calculating (1320) the cost metric in accordance with any of (ii)-(v) by:

evaluating the parameter for each link in the route; and

combining the evaluated parameters for each link in the route to calculate a total cost metric for the route.

18. The method (1300) according to claim 15, further comprising:

when the second host also holds the image, calculating a third value of the cost metric, for a third route over which the image could be downloaded to the third host from the second host;

selecting the second host to provide the image over the third route, when the third value of the cost metric is lowest.

19. A mesh computing system, configured to pre-emptively position images at hosts, for subsequent retrieval from the hosts and for a creation of containers from the images, the mesh computing system comprising hosts located in edge nodes of the mesh computing system, the mesh computing system operably coupled to a central registry that holds the images, the mesh computing system comprising a swarm leader configured to:

identify (1020) a need in a first host at a first edge node to run an image, the image being held in the central registry;

identify (1030) a second host at a second edge node of the mesh computing system; and

command (1040) a download of the image from the central registry to the second host, via at least one wireless link of the mesh computing system.

20. A mesh computing system, configured to pre-emptively position images at hosts located in edge nodes of the mesh computing system, for subsequent retrieval from the hosts and for a creation of containers from the images, the mesh computing system comprising a swarm leader and operably coupled to a central registry that holds multiple images for use at the hosts, wherein the mesh computing system comprises:

a first host, of the plurality of hosts, having a first registry that comprises a first image storage repository and a first application programming interface, API; and

a second host, of the plurality of hosts, having a second registry that comprises a second image storage repository and a second API;

a swarm leader, configured to:

identify (1120) an image, from the multiple images, as being an image that is likely to be used by hosts in the mesh computing system other than the first host and second host;

command the first host to download (1130) the image from the central registry to the first image storage repository;

command the second host to download (1140) the image from the central registry to the second image storage repository; and

a third host of the plurality of hosts, the third host configured to:

download (1170) the image from the registry of the first host or from the registry of the second host; and

create (1180) a container from the downloaded image.

* * * * *