



(51) International Patent Classification:

G06F 21/10 (2013.01) H04N 21/241 (2011.01)

H04N 21/222 (2011.01) H04N 21/835 (2011.01)

H04N 21/24 (2011.01)

(21) International Application Number:

PCT/US2021/035690

(22) International Filing Date:

03 June 2021 (03.06.2021)

(25) Filing Language:

English

(26) Publication Language:

English

(30) Priority Data:

63/193,476 26 May 2021 (26.05.2021) US

(71) Applicant: ARRIS ENTERPRISES LLC [US/US]; 3871 Lakefield Drive, Suwanee, GA 30024 (US).

(72) Inventors: VU, Tracy; 300 Main St, Boxford, MA 01921 (US). NAGUIB, Elhamy; 3933 222nd Pl Se, Bothell, WA 98021 (US). RAJALINGARI, Karthik; 20664 Acadia Ct, Cupertino, CA 95014 (US). THAKORE, Priyanki; 48 Ar-

bor Way, Groton, MA 01450 (US). MCKEOWN, Kieran; 33 Norwood Park, Cobh, Cork, P24A267 (IE).

(74) Agent: RUSSELL, Kevin, L. et al.; Chernoff, Vilhauer, McClung & Stenzel, LLP, 111 Sw Columbia Street, Suite 725, Portland, OR 97201 (US).


(81) Designated States (unless otherwise indicated, for every kind of national protection available): AE, AG, AL, AM, AO, AT, AU, AZ, BA, BB, BG, BH, BN, BR, BW, BY, BZ, CA, CH, CL, CN, CO, CR, CU, CZ, DE, DJ, DK, DM, DO, DZ, EC, EE, EG, ES, FI, GB, GD, GE, GH, GM, GT, HN, HR, HU, ID, IL, IN, IR, IS, IT, JO, JP, KE, KG, KH, KN, KP, KR, KW, KZ, LA, LC, LK, LR, LS, LU, LY, MA, MD, ME, MG, MK, MN, MW, MX, MY, MZ, NA, NG, NI, NO, NZ, OM, PA, PE, PG, PH, PL, PT, QA, RO, RS, RU, RW, SA, SC, SD, SE, SG, SK, SL, ST, SV, SY, TH, TJ, TM, TN, TR, TT, TZ, UA, UG, US, UZ, VC, VN, WS, ZA, ZM, ZW.

(84) Designated States (unless otherwise indicated, for every kind of regional protection available): ARIPO (BW, GH, GM, KE, LR, LS, MW, MZ, NA, RW, SD, SL, ST, SZ, TZ, UG, ZM, ZW), Eurasian (AM, AZ, BY, KG, KZ, RU, TJ, TM), European (AL, AT, BE, BG, CH, CY, CZ, DE, DK,

(54) Title: LICENSE MANAGEMENT FOR VCORES

700

FIG. 8



(57) Abstract: A cable distribution system includes a head end connected to a plurality of customer devices through a transmission network that includes a remote fiber node, that converts received data to analog data suitable to be provided on a coaxial cable for the plurality of customer devices. The cable distribution system includes a license manager.

EE, ES, FI, FR, GB, GR, HR, HU, IE, IS, IT, LT, LU, LV,
MC, MK, MT, NL, NO, PL, PT, RO, RS, SE, SI, SK, SM,
TR), OAPI (BF, BJ, CF, CG, CI, CM, GA, GN, GQ, GW,
KM, ML, MR, NE, SN, TD, TG).

Published:

— *with international search report (Art. 21(3))*

LICENSE MANAGEMENT FOR vCORES

CROSS REFERENCE TO RELATED APPLICATIONS

[0001] The present application claims priority under 35 U.S.C. § 119(e) to U.S. Provisional Application No. 63/193,476 filed on May 26, 2021, which is incorporated by reference herein in its entirety.

BACKGROUND

[0002] The subject matter of this application relates to secure access to license management for vCores.

[0003] Cable Television (CATV) services provide content to large groups of customers (e.g., subscribers) from a central delivery unit, generally referred to as a "head end," which distributes channels of content to its customers from this central delivery unit through an access network comprising a hybrid fiber coax (HFC) cable plant, including associated components (nodes, amplifiers and taps). Modern Cable Television (CATV) service networks, however, not only provide media content such as television channels and music channels to a customer, but also provide a host of digital communication services such as Internet Service, Video-on-Demand, telephone service such as VoIP, home automation/security, and so forth. These digital communication services, in turn, require not only communication in a downstream direction from the head end, through the HFC, typically forming a branch network and to a customer, but also require communication in an upstream direction from a customer to the head end typically through the HFC network.

[0004] To this end, CATV head ends have historically included a separate Cable Modem Termination System (CMTS), used to provide high speed data services, such as cable Internet, Voice over Internet Protocol, etc. to cable customers and a video headend

system, used to provide video services, such as broadcast video and video on demand (VOD). Typically, a CMTS will include both Ethernet interfaces (or other more traditional high-speed data interfaces) as well as radio frequency (RF) interfaces so that traffic coming from the Internet can be routed (or bridged) through the Ethernet interface, through the CMTS, and then onto the RF interfaces that are connected to the cable company's hybrid fiber coax (HFC) system. Downstream traffic is delivered from the CMTS to a cable modem and/or set top box in a customer's home, while upstream traffic is delivered from a cable modem and/or set top box in a customer's home to the CMTS. The Video Headend System similarly provides video to either a set-top, TV with a video decryption card, or other device capable of demodulating and decrypting the incoming encrypted video services. Many modern CATV systems have combined the functionality of the CMTS with the video delivery system (e.g., EdgeQAM - quadrature amplitude modulation) in a single platform generally referred to an Integrated CMTS (e.g., Integrated Converged Cable Access Platform (CCAP)) - video services are prepared and provided to the I-CCAP which then QAM modulates the video onto the appropriate frequencies. Still other modern CATV systems generally referred to as distributed CMTS (e.g., distributed Converged Cable Access Platform) may include a Remote PHY (or R-PHY) which relocates the physical layer (PHY) of a traditional Integrated CCAP by pushing it to the network's fiber nodes (R-MAC PHY relocates both the MAC and the PHY to the network's nodes). Thus, while the core in the CCAP performs the higher layer processing, the R-PHY device in the remote node converts the downstream data sent from the core from digital-to-analog to be transmitted on radio frequency to the cable modems and/or set top boxes, and converts the upstream radio frequency data sent from the cable modems and/or set top boxes from analog-to-digital format to be transmitted optically to the core.

BRIEF DESCRIPTION OF THE DRAWINGS

[0005] For a better understanding of the invention, and to show how the same may be carried into effect, reference will now be made, by way of example, to the accompanying drawings, in which:

[0006] FIG. 1 illustrates an integrated Cable Modem Termination System.

[0007] FIG. 2 illustrates a distributed Cable Modem Termination System.

[0008] FIG. 3 illustrates a layered network processing stack.

[0009] FIG. 4 illustrates a server system with a resource allocation manager and a container orchestration system.

[0010] FIG. 5 illustrates a server system with containers and a container orchestration system.

[0011] FIG. 6 illustrates a server with vCores and a vCore management interface.

[0012] FIG. 7 illustrates a server with vCores, a vCore management interface, and a license manager.

[0013] FIG. 8 illustrates an embodiment of the license manager managing licenses for vCores.

[0014] FIG. 9 illustrates an embodiment of the license manager managing licenses for vCores.

[0015] FIG. 10 illustrates an embodiment of the license manager managing licenses for vCores.

DETAILED DESCRIPTION

[0016] Referring to FIG. 1, an integrated CMTS (e.g., Integrated Converged Cable Access Platform (CCAP)) 100 may include data 110 that is sent and received over the

Internet (or other network) typically in the form of packetized data. The integrated CMTS 100 may also receive downstream video 120, typically in the form of packetized data from an operator video aggregation system. By way of example, broadcast video is typically obtained from a satellite delivery system and pre-processed for delivery to the subscriber through the CCAP or video headend system. The integrated CMTS 100 receives and processes the received data 110 and downstream video 120. The CMTS 130 may transmit downstream data 140 and downstream video 150 to a customer's cable modem and/or set top box 160 through a RF distribution network, which may include other devices, such as amplifiers and splitters. The CMTS 130 may receive upstream data 170 from a customer's cable modem and/or set top box 160 through a network, which may include other devices, such as amplifiers and splitters. The CMTS 130 may include multiple devices to achieve its desired capabilities.

[0017] Referring to FIG. 2, as a result of increasing bandwidth demands, limited facility space for integrated CMTSs, and power consumption considerations, it is desirable to include a Distributed Cable Modem Termination System (D-CMTS) 200 (e.g., Distributed Converged Cable Access Platform (CCAP)). In general, the CMTS is focused on data services while the CCAP further includes broadcast video services. The D-CMTS 200 distributes a portion of the functionality of the I-CMTS 100 downstream to a remote location, such as a fiber node, using network packetized data. An exemplary D-CMTS 200 may include a remote PHY architecture, where a remote PHY (R-PHY) is preferably an optical node device that is located at the junction of the fiber and the coaxial. In general the R-PHY often includes the PHY layers of a portion of the system. The D-CMTS 200 may include a D-CMTS 230 (e.g., core) that includes data 210 that is sent and received over the Internet (or other network) typically in the form of packetized data. The D-CMTS 200 may also receive downstream video 220, typically in the form of packetized data from an operator video aggregation system. The D-CMTS 230 receives and processes the received data 210 and downstream video 220. A remote Fiber node 280 preferably include a remote PHY device 290. The remote PHY device 290 may

transmit downstream data 240 and downstream video 250 to a customer's cable modem and/or set top box 260 through a network, which may include other devices, such as amplifier and splitters. The remote PHY device 290 may receive upstream data 270 from a customer's cable modem and/or set top box 260 through a network, which may include other devices, such as amplifiers and splitters. The remote PHY device 290 may include multiple devices to achieve its desired capabilities. The remote PHY device 290 primarily includes PHY related circuitry, such as downstream QAM modulators, upstream QAM demodulators, together with pseudowire logic to connect to the D-CMTS 230 using network packetized data. The remote PHY device 290 and the D-CMTS 230 may include data and/or video interconnections, such as downstream data, downstream video, and upstream data 295. It is noted that, in some embodiments, video traffic may go directly to the remote physical device thereby bypassing the D-CMTS 230. In some cases, the remote PHY and/or remote MAC PHY functionality may be provided at the head end.

[0018] By way of example, the remote PHY device 290 may convert downstream DOCSIS (i.e., Data Over Cable Service Interface Specification) data (e.g., DOCSIS 1.0; 1.1; 2.0; 3.0; 3.1; and 4.0 each of which are incorporated herein by reference in their entirety), video data, out of band signals received from the D-CMTS 230 to analog for transmission over RF or analog optics. By way of example, the remote PHY device 290 may convert upstream DOCSIS, and out of band signals received from an analog medium, such as RF or linear optics, to digital for transmission to the D-CMTS 230. As it may be observed, depending on the particular configuration, the R-PHY may move all or a portion of the DOCSIS MAC and/or PHY layers down to the fiber node.

[0019] I-CMTS devices are typically custom built hardware devices that consist of a single chassis that include a series of slots, each of which receives a respective line card with a processor, memory, and other computing and networking functions supported thereon. Each of the line cards include the same hardware configuration, processing capabilities, and software. Each of the line cards performs the functions of the I-CMTS

device, including the MAC and PHY functionality. As the system increasingly scales to support additional customers, additional line cards are included with the system to expand the processing capability of the system. Unfortunately, it is problematic to dynamically scale the number of line cards in a real-time manner to meet the demands of a particular network.

[0020] The computational power of microprocessor based commercial off the shelf (COTS) server platforms are increasing while the expense of such systems is decreasing over time. With such systems, a computing system may be, if desired, virtualized and operated using one or more COTS server, generally referred to herein as a virtual machine. Using container technologies running on the COTS server and/or virtual machine, the COTS server may operate with only a single operating system. Each of the virtualized applications may then be isolated using software containers, such that the virtualized application may not see and are not aware of other virtualized applications operating on the same machine. Typically, each COTS server includes one or more Intel / AMD processors (or other processing devices) with associated memory and networking capabilities running an operating system software. Typically the COTS servers include a framework and an operating system where user applications are run on such framework and the operating system is abstracted away from the actual operating system. Each virtual machine may be instantiated and operated as one or more software applications running on a COTS server. A plurality of software containers may be instantiated and operated on the same COTS server and/or the same virtual machine. A plurality of COTS servers is typically included in one or more data centers, each of which are in communication with one another. A plurality of COTS server may be located in different geographic areas to provide geo-redundancy. In some embodiments, the container may include the same functionality as a virtual machine, or vice versa. In some embodiments, a grouping of containerized components, generally referred to as a pod, may be in the form of a virtual machine.

[0021] In some embodiments, the COTS servers may be “bare metal” servers that typically include an operating system thereon together with drivers and a portion of a container orchestration system. One or more containers are then added to the “bare metal” server while being managed by the container orchestration system. The container orchestration system described herein may likewise perform as, and be referred to as, a virtual machine orchestration system, as desired. In some embodiments, “bare metal” servers may be used with pods running on the operating system thereon together with drivers and a container orchestration system. In some embodiments, virtual machines may be omitted from the COTS servers.

[0022] Selected software processes that are included on a line card and/or a remote PHY device may be run on a “bare metal” server and/or virtual machine, including software containers, running on a COTS server, including both “active” and “back-up” software processes. The functionality provided by such a “bare metal” server and/or virtual machine may include higher level functions such as for example, packet processing that includes routing Internet packet provisioning, layer 2 virtual private networking which operates over pseudowires, and multiprotocol label switching routing. The functionality provided by such a “bare metal” server and/or virtual machine may include DOCSIS functions such as for example, DOCSIS MAC and encapsulation, channel provisioning, service flow management, quality of service and rate limiting, scheduling, and encryption. The functionality provided by such a “bare metal” server and/or virtual machine may include video processing such as for example, EQAM and MPEG processing.

[0023] Each of the COTS servers and/or the virtual machines and/or software containers may contain different hardware profiles and/or frameworks. For example, each of the COTS servers and/or “bare metal” servers and/or virtual machines and/or software containers may execute on different processor types, different number of processing cores per processor, different amounts of memory for each processor type, different amounts of memory per processing core, different cryptographic capabilities,

different amounts of available off-processor memory, different memory bandwidth (DDR) speeds, and varying types and capabilities of network interfaces, such as Ethernet cards. In this manner, different COTS servers and/or “bare metal” servers and/or virtual machines and/or software containers may have different processing capabilities that vary depending on the particular hardware. Each of the COTS servers and/or “bare metal” servers and/or the virtual machine and/or software containers may contain different software profiles. For example, each of the COTS servers and/or “bare metal” servers and/or virtual machines and/or software containers may include different software operating systems and/or other services running thereon, generally referred to herein as frameworks. In this manner, different COTS servers and/or “bare metal” servers and/or virtual machines and/or software containers may have different software processing capabilities that vary depending on the particular software profile.

[0024] Referring to FIG. 3, for data processing and for transferring data across a network, the architecture of the hardware and/or software may be configured in the form of a plurality of different planes, each of which performing a different set of functionality. In relevant part the layered architecture may include different planes such as a management plane 300, a control plane 310, a data plane 320, and switch fabric 330 to effectuate sending and receiving packets of data.

[0025] For example, the management plane 300 may be generally considered as the user interaction or otherwise the general software application being run. The management plane typically configures, monitors, and provides management, and configuration served to all layers of the network stack and other portions of the system.

[0026] For example, the control plane 310 is a component to a switching function that often includes system configuration, management, and exchange of routing table information and forwarding information. Typically, the exchange of routing table information is performed relatively infrequently. A route controller of the control plane 310 exchanges topology information with other switches and constructs a routing table

based upon a routing protocol. The control plane may also create a forwarding table for a forwarding engine. In general, the control plane may be thought of as the layer that makes decisions about where traffic is sent. Since the control functions are not performed on each arriving individual packet, they tend not to have a strict speed constraint.

[0027] For example, the data plane 320 parses packet headers for switching, manages quality of service, filtering, medium access control, encapsulations, and/or queuing. As a general matter, the data plane carries the data traffic, which may be substantial in the case of cable distribution networks. In general, the data plane may be thought of as the layer that primarily forwards traffic to the next hop along the path to the selected destination according to the control plane logic through the switch fabric. The data plane tends to have strict speed constraints since it is performing functions on each arriving individual packet.

[0028] For example, the switch fabric 330 provides a network topology to interconnect network nodes via one or more network switches.

[0029] As the system increasingly scales to support additional customers, additional COTS servers and/or “bare metal” servers and/or virtual machines and/or software containers are included with the system to expand the processing capability of the overall system. To provide processing redundancy, one or more additional COTS servers and/or “bare metal” servers and/or virtual machines and/or software containers may be included that are assigned as “back-up” which are exchanged for an “active” process upon detection of a failure event. The scaling of the data plane 320 on COTS servers and/or “bare metal” servers and/or virtual machines and/or software containers to service dynamically variable processing requirements should be performed in such a manner that ensures sufficiently fast processing of data packets and sufficient bandwidth for the transmission of the data packets to ensure they are not otherwise lost.

[0030] It is desirable to virtualize the data plane, and in particular a portion of the Remote PHY functionality on a COTS server and/or “bare metal” servers. In this manner, the MAC cores for the cable distribution system may run on COTS servers and/or “bare metal” servers. By way of reference herein, a virtualized Remote PHY MAC Core may be referred to herein as a vCore instance.

[0031] Referring to FIG. 4, it is desirable to incorporate platform as a service that uses operating system level virtualization to deliver software in packages, generally referred to as containers 410. Each of the containers are isolated from one another and bundle their own software, libraries, and configuration files. The containers may communicate with one another using defined channels. As a general matter, one or more applications and its dependencies may be packed in a virtual container that can run on a COTS server and/or “bare metal” server and/or a virtual machine. This containerization increases the flexibility and portability on where the application may run, such as an on-premises COTS server, a “bare metal” server, a public cloud COTS server, a private cloud COTS server, or otherwise. With each container being relatively lightweight, a single COTS server and/or “bare metal” server and/or a virtual machine operating on a COTS server and/or “bare metal” server may run several containers simultaneously. In addition, the COTS server and/or “bare metal” server and/or the virtual machine and/or the containers may be distributed within the cable distribution system.

[0032] A COTS server and/or “bare metal” server and/or a virtual machine may include a container orchestration system 420 for automating the application deployment, scaling, and management of the containers 410 across one or more COTS servers and/or “bare metal” servers and/or virtual machines. Preferably the computing device running the container orchestration system 420 is separate from the computing device providing the containers for the dataplane applications. It is to be understood that the virtual machine illustrated in FIG. 4 may be omitted, such as the COTS B. The application deployment, scaling, and management of the containers may include clusters across multiple hosts, such as multiple COTS servers. The deployment, maintaining, and

scaling, of the containers may be based upon characteristics of the underlying system capabilities, such as different processor types, different number of processing cores per processor, different amounts of memory for each processor type, different amounts of memory per processing core, different amounts of available off-processor memory, different memory bandwidth (DDR) speeds, different frameworks, and/or varying types and capabilities of network interfaces, such as Ethernet cards. Moreover, the container orchestration system 420 may allocate different amounts of the underlying system capabilities, such as particular processor types, a selected number of processors (e.g., 1 or more), a particular number of processing cores per selected processor, a selected amount of memory for each processor type, a selected amount of memory per processing core, a selected amount of available off-processor memory, a selected framework, and/or a selected amount and/or type of network interface(s), such as Ethernet cards. A corresponding agent for the container orchestration system 420 may be included on each COTS server (e.g., COTS A and/or COTS B).

[0033] The container orchestration system 420 may include a grouping of containerized components, generally referred to as a pod 430. A pod consists of one or more containers that are co-located on the same COTS server and/or “bare metal” server and/or the same virtual machine, which can share resources of the same COTS server and/or “bare metal” server and/or same virtual machine. Each pod 430 is preferably assigned a unique pod IP address within a cluster, which allows applications to use ports without the risk of conflicts. Within the pod 430, each of the containers may reference each other based upon a localhost or other addressing service, but a container within one pod preferably has no way of directly addressing another container within another pod, for that, it preferably uses the pod IP address or otherwise an addressing service.

[0034] A traditional D-CMTS RPHY Core may be implemented as a speciality built appliance including both software and hardware to achieve desired performance characteristics, such as ensuring the timing of the transfer of data packets. The specially built appliance is not amenable to automatic deployment nor automatic scaling due to the

fixed nature of its characteristics. In contrast to a specially built appliance, the vCore instance is preferably implemented in software operating on a COTS server and/or “bare metal” server on top of an operating system, such as Linux. The vCore instance is preferably implemented in a manner that readily facilitates automation techniques such as lifecycle management, flexible scaling, health monitoring, telemetry, etc. Unfortunately, running a vCore instance on a COTS server and/or “bare metal” server tends to result in several challenges, mostly related to the data plane components. One of the principal challenges involves ensuring that data is provided to the network in a timely and effective manner to achieve the real time characteristics of a cable data distribution environment. The cable data distribution environment includes real time constraints on the timing of data packet delivery, which is not present in typical web-based environments or database environments.

[0035] Each vCore instance is preferably implemented within a container, where the size (e.g., scale, memory, CPU, allocation, etc.) of each container translates into the amount of server hardware and software resources assigned to the particular vCore instance. The amount of server hardware and software resources assigned to each particular vCore instance is preferably a function of the number of groups of customers (e.g., service groups) and/or number of customers that the vCore instance can readily provide RPHY MAC Core services to. For example, a limited amount of server hardware and software resources may be assigned to a particular vCore instance that has a limited number of groups of customers and/or customers. For example, a substantial amount of server hardware and software resources may be assigned to a particular vCore instance that has a substantial number of groups of customers and/or customers. For example, selected server hardware resources are preferably allocated among the different vCore instances in a non-overlapping manner so that each vCore instance has a dedicated and predictable amount of server hardware resources. For example, selected software resources are preferably allocated among the different vCore instances in a non-

overlapping manner so that each vCore instance has a dedicated and predictable amount of software resources.

[0036] For example, the number of CPU cores preferably assigned to each vCore instance (C_c) may be a function of the total USSG (upstream service groups – groups of customer modems and/or set top boxes) (USsg) and the total DSSG (downstream service groups – groups of customer modems and/or set top boxes) (DSsg) connected through that vCore instance. This may be represented as vCore: $C_c = f_1(\text{USsg}, \text{DSsg})$. Other hardware and/or software characteristics may likewise be assigned, as desired.

[0037] For example, the network capacity assigned to each vCore instance (C_{bw}) may be a function of the of the total USSG (upstream service groups – groups of customer modems and/or set top boxes) (USsg) and the total DSSG (downstream service groups – groups of customer modems and/or set top boxes) (DSsg) connected to that vCore instance. This may be represented as $C_{bw} = f_2(\text{USsg}, \text{DSsg})$. Other hardware and/or software characteristics may likewise be assigned, as desired.

[0038] The scaling of the vCore instance may refer to the capability to automatically create and deploy a vCore instance within a container on a COTS server and/or “bare metal” server and/or virtual machine that is appropriately sized to serve a particular set of remote physical devices and/or service groups (e.g., sets of cable customers) and/or cable customers. The scaling of the vCore instance may also include, in some cases, the capability to automatically modify the hardware and/or software characteristics of an existing vCore instance within a container on a COTS server and/or “bare metal” server and/or virtual machine to be appropriately sized to serve a modified particular set of remote physical devices and/or service groups (e.g., sets of cable customers) and/or cable customers.

[0039] A resource allocation manager 470 may assign or reallocate a suitable amount of hardware and software of the COTS server and/or “bare metal” server resources to each particular vCore instance (e.g., CPU cores, and/or memory, and/or network

capacity). The amount of such COTS server and/or “bare metal” server hardware and software resources assigned to or reallocate to each vCore instance may be a function of its scale and also other features, such as various other resource allocations. A corresponding agent for the resource allocation manager 470 may be included on each COTS server (e.g., COTS A, COTS B).

[0040] The vCore instance includes data plane software for the transfer of data packets and other functions of the data plane.. The data plane software may include a set of data plane libraries and network interface controller (NIC) drivers that are used to manage the data packets for the data plane. Preferably, the data plane software operates in user space, as opposed to Kernel space like typical network processing software, thus it does not make use of the operating system kernel and container management network drivers and plugins. For example, the data plane software may include a queue manager, a buffer manager, a memory manager, and/or a packet framework for packet processing. The data plane software may use CPU cores that are isolated from the Kernel, meaning that the operating system scheduled processes are not running on these isolated CPU cores. The separation of the CPU cores between the data plane software and the operating system software ensures that tasks performed by the operating system software does not interfere with the data plane software processing the data packets in a timely manner. In addition, the separation of the CPU cores between the data plane software and the operating system software enables both to use the same physical central processing unit, albeit different cores, of the same physical central processing unit. In addition, other hardware and/or software capabilities may likewise be separated, such as for example, selected processors (e.g., 1 or more), particular number of processing cores per selected processor, selected amount of memory for each processor type, selected amount of memory per processing core, selected amount of available off-processor memory, selected framework, and/or selected amount and/or type of network interface(s).

[0041] It is also desirable for each vCore instance to have dedicated network bandwidth capability apart from other vCore instances and the operating system software.

To provide dedicated network bandwidth for a vCore instance, the physical network interface cards may be virtualized so that a plurality of different software applications can make use of the same network interface card, each with a guaranteed amount of bandwidth available. The network interface cards are preferably virtualized using a single root input / output virtualization technique (SR-IOV). The SR-IOV partitions the NIC physical functions (e.g., PFs) into one or more virtual functions (VFs). The capabilities of the PFs and VFs are generally different. In general, the PF supports queues, descriptions, offloads, hardware lock, hardware link control, etc. In general, the VF supports networking features based upon queues and descriptors.

[0042] The automated creation, deployment, and removal of vCore instances may be performed by the container orchestration system 420.

[0043] Referring to FIG. 5, the vCore instances 530 may operate on a COTS server and/or “bare metal” server 500 acting as a remote PHY MAC core for one or more remote physical devices connected over a converged interconnect network, normally located in the same hub. The vCore instances 530 may include data plane software 532. Each of the vCore instances 530 as generally referred to as a POD. In some cases, multiple vCores may be included in a POD. The COTS server 500 may communicate with the Internet 560, a set of networking switches 570, to remote physical devices 580, and the customers 590. The COTS server and/or “bare metal” server including the vCore instances operating thereon is typically a relatively high performance server that has one or more of the following characteristics:

[0044] Hardware:

[0045] At least one management NIC 510 is connected to, usually, a separate management network 512. The management NIC 510 is primarily used for orchestration and management of the server application, which may also manage the data traffic.

[0046] Preferably at least two (for redundancy) data plane NICs 514 (i.e., data plane physical network interfaces) together with SR-IOV and PTP (IEEE 1588) 522 are included for hardware timestamping capabilities of the data packets. The data plane NICs 514 are used to provide connectivity to the remote physical devices and the customer modems and/or set top boxes / consumer premises equipment behind such remote physical devices. The vCore instances 530 may each include a virtual function 534 network interface to each of the data plane NICs 514.

[0047] In addition, the hardware may include dedicated devices for DES encryption.

[0048] Software:

[0049] Preferably the operating system on the COTS server and/or “bare metal” server is a LINUX OS such as Ubuntu, Redhat, etc.

[0050] The COTS Server and/or “bare metal” server and/or virtual machine includes container software.

[0051] The COTS Server and/or “bare metal” server and/or virtual machine and/or other server includes at least a part of a container orchestration system.

[0052] The COTS Server and/or “bare metal” server and/or virtual machine and/or other server includes a resource allocation manager (RAM) 520 that manages, at least in part, the server allocation of software and/or hardware resources for vCore instances, including for example: CPU Cores, memory, VFs, MAC addresses, etc. The RAM 520 may also provide server configuration, including OS configuration, driver support, etc., diagnostics and health monitoring. The COTS Server and/or “bare metal” server and/or virtual machine and/or other server may include an orchestration app 540 that manages, at least in part, the management of the vCores (e.g., containers and/or pods).

[0053] The COTS Server and/or “bare metal” server and/or virtual machine and/or other server may run the PTP application 522 that synchronizes the system clock of the

COTS Server and/or “bare metal” server and/or virtual machine and/or vCore instances 520 based upon a grand master clock for the system as a whole. For increased accuracy, the PTP application 522 is preferably based upon hardware time stamping and a Precise Hardware Clock that is present on the NICs 514. As a general matter, the vCore preferably performs the MAC layer functionality.

[0054] Referring to FIG. 6, a simplified illustration shows each of the vCores 600 traditionally includes a command line interface 610 to receive instructions thereto for configuration, diagnostics, control, and management. The command line interface processes commands in the form of lines of text. A command line interface may be provided using a shell to provide interactive access to the corresponding operating system functions or services. A plurality of vCores, such as 32 vCores, may be installed on a COTS Server and/or “bare metal” server and/or virtual machine and/or other server 620 (generally referred to herein as a “server”). Each of the vCores 600 on the server 620 includes a respective command line interface 610. A vCore management interface 630 provides access to each of the respective vCores 700, such as opening up a respective terminal window, to facilitate access to the respective command line interface 610. By way of example, an operator may open a terminal window to access the command line interface on a first vCore and execute commands thereon. For example, bidirectional interactive text-oriented communications may be used for such command execution, such as Telnet. In order to provide such configuration, diagnostics, control, and management for each of the vCores, at least 32 such terminal windows are opened and 32 sets of commands are executed. In addition, to enable each of the vCores to provide selected services, one or more license keys are installed on the respective vCores. For a large system, there may be many servers each having a substantial number of vCores, therefore resulting in potentially thousands of vCores that periodically need to be accessed on an individual basis, and in particular to install one or more license keys on each of the vCores. This is a burdensome task to undertake using the management interface 630.

[0055] With such a substantial number of potential vCores 600, together with a substantial number of potential configurations of each of the vCores 600, the number of potentially different license combinations among a multitude of different vCores is burdensome to manually manage. By way of example, a license may be required for basic vCore functionality, an additional license may be required for one or more channels supported by the vCore, and/or an additional license may be required for one or more remote physical devices, and/or different license may be required for different upstream data bandwidth and for different downstream data bandwidth provided by the vCores and/or remote physical device(s), and/or a different number of users. Referring to FIG. 7, the vCore management interface 630 may include a license manager 700. The license manager 700 may likewise be located at any suitable location within the network that may communicate with one or more of the vCores of one or more of the servers.

[0056] Referring to FIG. 8, to simplify the license management, the license manager 700 may include a set of available credits 800 that are purchased or otherwise available for use by an operator that may be used for suitable licenses of vCores and/or features thereof. Based upon the available credits 800, the license manager may generate one or more license keys 810 for one or more vCores and/or features thereof. After generation of the license key(s) 810, the license key(s) 810 is installed on the corresponding vCore 820 to enable the vCores and/or features thereof. The number of available credits are reduced 830 based upon the license key(s) used. The license manager 700 tracks the licenses that are applied to vCores which are made available to the operator, as desired, for tracking purposes. When the operator removes 850 one or more features from a vCore and/or decommissions one or more vCores, the license manager disables, modifies, or otherwise removes the license key(s) 860. After disabling, modifying, or otherwise removing the license key(s) 860, the license manager supplements the set of available credits 870 that are available for use by the operator based upon the license key(s) disabled, modified, or otherwise removed. The license manager 700 enables the operator to manage the licenses applied across a plurality of vCores in a manner that is efficient,

together with license mobility as a result of reclaiming licenses that are no longer being used in the form of credits, that may be deployed in another manner. It is noted that, if credits are not available, then the corresponding services to be provided by the vCore(s) and/or feature(s) thereof are not enabled.

[0057] Referring to FIG. 9, to further simplify the license management, when a vCore is instantiated for service and/or a feature thereof 900, the system automatically checks the license manager to determine if there are available credits 910. If there are available credits, the license manager automatically generates a suitable license key(s) 920 that is then automatically applied to the instantiated vCore and/or feature thereof 930. In a similar manner, when a vCore is decommissioned from service and/or a feature thereof 950, the system automatically determines the licenses that were applied to the decommissioned vCore and/or feature thereof 960. The number of credits associated with the licenses of the decommissioned vCore(s) and/or feature(s) thereof are automatically determined 970, and are automatically supplemented to the set of available credits 980.

[0058] Referring to FIG. 10, one exemplary embodiment includes an operator 1000 that may send and/or receive a request to or from a user interface 1010 to add / deploy / modify / remove vCores, add / deploy / modify / and/or remove remote physical devices, and/or services thereof. The user interface 1010 may forward requests to and receive requests from a vCore manager 1020. The vCore manager 1020 may populate a list of required licenses and provide license related requests to a license manager 1030. The license manager 1030 may generate license keys, as needed. The user interface 1010 may query license data from the license manager 1030 for providing such data to the operator 1000. The vCore manager 1020 may configure and deploy the vCore devices 1040. The license manager 1030 may install and uninstall licenses on the vCore(s) 1040. Each of the vCores 1040 may have one or more associated remote physical devices 1050. The vCore manager 1020 and the license manager 1030 may use a database 1060 to store and retrieve data from.

[0059] The license manager also permits the operator to manually obtain license keys and manually deploy those license keys to corresponding vCores, if desired. The license manager also permits the operator to manually uninstall license keys on corresponding vCores, if desired. The license manager also permits the operator to supplement the available credits for the uninstalled license keys. However, over time, there tends to be unused license keys installed on various vCores and the management thereof can be problematic. To reduce the effects of manual management of the license keys, the system may include an auto-cure process that interrogates all of the vCores to determine which licenses are being used, update the database of the used licenses, and update the corresponding available credits. In this manner, the system may update itself to alleviate license management that may occur as a result of manual license management.

[0060] Moreover, each functional block or various features in each of the aforementioned embodiments may be implemented or executed by a circuitry, which is typically an integrated circuit or a plurality of integrated circuits. The circuitry designed to execute the functions described in the present specification may comprise a general-purpose processor, a digital signal processor (DSP), an application specific or general application integrated circuit (ASIC), a field programmable gate array (FPGA), or other programmable logic devices, discrete gates or transistor logic, or a discrete hardware component, or a combination thereof. The general-purpose processor may be a microprocessor, or alternatively, the processor may be a conventional processor, a controller, a microcontroller or a state machine. The general-purpose processor or each circuit described above may be configured by a digital circuit or may be configured by an analogue circuit. Further, when a technology of making into an integrated circuit superseding integrated circuits at the present time appears due to advancement of a semiconductor technology, the integrated circuit by this technology is also able to be used.

[0061] It will be appreciated that the invention is not restricted to the particular embodiment that has been described, and that variations may be made therein without

departing from the scope of the invention as defined in the appended claims, as interpreted in accordance with principles of prevailing law, including the doctrine of equivalents or any other principle that enlarges the enforceable scope of a claim beyond its literal scope. Unless the context indicates otherwise, a reference in a claim to the number of instances of an element, be it a reference to one instance or more than one instance, requires at least the stated number of instances of the element but is not intended to exclude from the scope of the claim a structure or method having more instances of that element than stated. The word "comprise" or a derivative thereof, when used in a claim, is used in a nonexclusive sense that is not intended to exclude the presence of other elements or steps in a claimed structure or method.

CLAIMS

1. A cable distribution system comprising:
 - (a) a head end connected to a plurality of customer devices through a transmission network that includes a remote fiber node, that converts received data to analog data suitable to be provided on a coaxial cable for said plurality of customer devices, where said head end includes at least one server each of which includes a respective processor;
 - (b) a vCore suitable to be instantiated on one of said at least one servers of said head end configured to provide services to a plurality of customers;
 - (c) a license manager that manages available credits for said vCore;
 - (d) said license manager provides a license key to said vCore to enable said vCore and/or services for said plurality of customers based upon having sufficient said available credits, and decreasing the number of said available credits based upon providing said license key to said vCore;
 - (e) said license manager disables a license key for said vCore to disable said vCore and/or services for said plurality of customers, and increasing the number of said available credits based upon disabling said license key for said vCore.
2. The cable distribution system of claim 1 wherein said license manager provides a corresponding license key to a plurality of different said vCores.
3. The cable distribution system of claim 2 wherein said license manager generates said corresponding license key.
4. The cable distribution system of claim 1 wherein a number of available credits used for said license key provided to said vCore varies based upon the number of channels provided by said vCore.

5. The cable distribution system of claim 1 wherein a number of available credits used for said license key provided to said vCore varies based upon the bandwidth of an upstream channel provided by said vCore.

6. The cable distribution system of claim 1 wherein a number of available credits used for said license key provided to said vCore varies based upon the bandwidth of a downstream channel provided by said vCore.

7. The cable distribution system of claim 1 wherein a number of available credits used for said license key provided to said vCore varies based upon a number of remote physical devices supported by said vCore.

8. The cable distribution system of claim 1 wherein a number of available credits used for said license key provided to said vCore varies based upon the number of said customers supported by said vCore.

9. The cable distribution system of claim 1 wherein said license manager automatically applies said license key to said vCore in response to instantiating said vCore and/or services thereof.

10. The cable distribution system of claim 9 wherein said license manager automatically generates said license key for said vCore in response to said instantiating said vCore and/or services thereof.

11. The cable distribution system of claim 10 wherein said license manager automatically modifies said available credits in response to generating said license key for said vCore.

12. The cable distribution system of claim 1 wherein said license manager automatically disables said license key of said vCore in response to modifying said vCore and/or services thereof.

13. The cable distribution system of claim 12 wherein said license manager automatically modifies said available credits in response to disabling said license key of said vCore.

14. The cable distribution system of claim 13 wherein said modifying includes at least one of disabling, modifying, and removing.




FIG. 1




FIG. 2

3/9




FIG. 3




FIG. 4

4/9




FIG. 5

5/9




FIG. 6

6/9




FIG. 7

7/9




FIG. 8

8/9




FIG. 9

9/9




FIG. 10

INTERNATIONAL SEARCH REPORT

International application No
PCT/US2021/035690

A. CLASSIFICATION OF SUBJECT MATTER

INV. **G06F21/10 H04N21/22 H04N21/24 H04N21/241 H04N21/835**
ADD.

According to International Patent Classification (IPC) or to both national classification and IPC

B. FIELDS SEARCHED

Minimum documentation searched (classification system followed by classification symbols)

G06F H04N

Documentation searched other than minimum documentation to the extent that such documents are included in the fields searched

Electronic data base consulted during the international search (name of data base and, where practicable, search terms used)

EPO-Internal

C. DOCUMENTS CONSIDERED TO BE RELEVANT

Category*	Citation of document, with indication, where appropriate, of the relevant passages	Relevant to claim No.
Y	US 2021/152386 A1 (CHARI SANTHANA [US] ET AL) 20 May 2021 (2021-05-20) figures 6,7 -----	1-14
Y	US 2018/321927 A1 (BORTHAKUR APOLAK [US] ET AL) 8 November 2018 (2018-11-08) paragraph [0070] -----	1-14
Y	US 2014/068792 A1 (KONDOU HIROSHI [JP] ET AL) 6 March 2014 (2014-03-06) paragraphs [0004], [0091], [0092] -----	1-14
A	US 2006/106727 A1 (YELLAI PRABHAKARA R [US] ET AL) 18 May 2006 (2006-05-18) paragraphs [0042], [0043]; figures 1-3 -----	1-14
A	US 2008/069206 A1 (YANG CHARLIE X [US]) 20 March 2008 (2008-03-20) figures 1,4 -----	1-14

☐ Further documents are listed in the continuation of Box C.

☒ See patent family annex.

* Special categories of cited documents :

"A" document defining the general state of the art which is not considered to be of particular relevance

"E" earlier application or patent but published on or after the international filing date

"L" document which may throw doubts on priority claim(s) or which is cited to establish the publication date of another citation or other special reason (as specified)

"O" document referring to an oral disclosure, use, exhibition or other means

"P" document published prior to the international filing date but later than the priority date claimed

"T" later document published after the international filing date or priority date and not in conflict with the application but cited to understand the principle or theory underlying the invention

"X" document of particular relevance; the claimed invention cannot be considered novel or cannot be considered to involve an inventive step when the document is taken alone

"Y" document of particular relevance; the claimed invention cannot be considered to involve an inventive step when the document is combined with one or more other such documents, such combination being obvious to a person skilled in the art

"&" document member of the same patent family

Date of the actual completion of the international search

1 February 2022

Date of mailing of the international search report

09/02/2022

Name and mailing address of the ISA/
European Patent Office, P.B. 5818 Patentlaan 2
NL - 2280 HV Rijswijk
Tel. (+31-70) 340-2040,
Fax: (+31-70) 340-3016

Authorized officer

Doyle, Walter

INTERNATIONAL SEARCH REPORT

Information on patent family members

International application No

PCT/US2021/035690

Patent document cited in search report	Publication date	Patent family member(s)	Publication date
US 2021152386 A1	20-05-2021	US 2021152386 A1	20-05-2021
		WO 2021101898 A1	27-05-2021

US 2018321927 A1	08-11-2018	CA 2990377 A1	05-11-2018
		EP 3399701 A1	07-11-2018
		US 2018321927 A1	08-11-2018
		US 2018321928 A1	08-11-2018
		US 2019163456 A1	30-05-2019
		US 2020401390 A1	24-12-2020

US 2014068792 A1	06-03-2014	EP 2704050 A1	05-03-2014
		JP 6074955 B2	08-02-2017
		JP 2014048907 A	17-03-2014
		US 2014068792 A1	06-03-2014

US 2006106727 A1	18-05-2006	US 2006106727 A1	18-05-2006
		US 2010241568 A1	23-09-2010
		US 2014207909 A1	24-07-2014

US 2008069206 A1	20-03-2008	US 2008069206 A1	20-03-2008
		US 2019007652 A1	03-01-2019
		WO 2008036232 A2	27-03-2008
