

(51) International Patent Classification:
G06F 9/44 (2018.01)(21) International Application Number:
PCT/IB2021/056289(22) International Filing Date:
13 July 2021 (13.07.2021)

(25) Filing Language: English

(26) Publication Language: English

(30) Priority Data:
16/930,874 16 July 2020 (16.07.2020) US(71) Applicant: **INTERNATIONAL BUSINESS MACHINES CORPORATION** [US/US]; New Orchard Road, Armonk, New York 10504 (US).(71) Applicants (for MG only): **IBM (CHINA) INVESTMENT COMPANY LTD.** [CN/CN]; 25/F, Pangu Plaza, No. 27, Central North 4th Ring Road, Chaoyang District, Beijing 100101 (CN). **IBM DEUTSCHLAND GMBH** [DE/DE]; IBM-Allee 1, 71139 Ehningen (DE).(72) Inventors: **OKADA, Masakuni**; c/o IBM Japan, 9-21 Nihonbashi-Hakozakicho, Chuo-ku, Tokyo 103-8510 (JP). **UENOHARA, Hayato**; c/o IBM Japan, 9-21 Nihonbashi-Hakozakicho, Chuo-ku, Tokyo 103-8510 (JP).**MATSUNAGA, Takuya**; c/o IBM Japan, 9-21 Nihonbashi-Hakozakicho, Chuo-ku, Tokyo 103-8510 (JP).(74) Agent: **FERARA, Nina**; c/o IBM Deutschland GmbH, Patentwesen und Urheberrecht, IBM-Allee 1, 71139 Ehningen (DE).

(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, 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,

(54) Title: POOL MANAGEMENT FOR IN-VEHICLE DEVICE APPLICATION STARTUP

(57) Abstract: A method for pre-pooling virtual machine components prior to application start-up. The method includes reading a virtual machine pool manifest by a virtual machine pool manager (310). The method also includes launching a plurality of base virtual machines into a virtual machine pool and based on the initial number provided by the virtual machine pool manifest (320). The method further includes allocating initial resources to a portion of the base virtual machines based on a virtual machine resource definition in the virtual machine pool manifest (330). The method includes loading core program packages into the portion of the base virtual machines (340).




FIG. 3

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))*
- *in black and white; the international application as filed
contained color or greyscale and is available for download
from PATENTSCOPE*

POOL MANAGEMENT FOR IN-VEHICLE DEVICE APPLICATION STARTUP

BACKGROUND

[0001] The present disclosure relates to application startup speed in an in-vehicle device, and more specifically, to improving application startup speed by dividing the application launch process using efficient pool management techniques.

[0002] Information of Things (IoT) technology allows for function extensibility by providing the capability of adding additional applications into an in-vehicle device (e.g., a navigation system). Additionally, application development in vehicles is expanding to allow applications to control system functions (e.g., dashboard camera, air conditioner, vehicle information management) within a vehicle. These applications are developed as an integral part of the core functionality of an in-vehicle device allowing for system functions that can be installed as replaceable applications within the in-vehicle device.

SUMMARY

[0003] Various embodiments of the present disclosure are directed to a computer-implemented method of pre-pooling virtual machine components prior to application startup. The computer-implemented method includes reading a virtual machine pool manifest by a virtual machine pool manager. The virtual machine pool manifest includes an initial number of virtual machines to launch. The virtual machine pool manifest can also include how many virtual machines to allocate resources to, the number of resources to allocate based on a virtual machine resource definition, the number of virtual machines without a resource allocation, a default environment, and a parameter that indicates a maximum number of executable applications at a given time. The computer-implemented method also includes launching a plurality of base virtual machines into a virtual machine pool and based on the initial number provided by the virtual machine pool manifest. The base virtual machines lack initial application assignments. The computer-implemented method further includes allocating initial resources to a portion of the base virtual machines based on a virtual machine resource definition in the virtual machine pool manifest. The computer-implemented method includes loading core program packages into the portion of the base virtual machines.

[0004] Additional embodiments of the present disclosure include a computer program product of pre-pooling virtual machine components prior to the application startup, which can include a computer-readable storage medium having program instructions embodied therewith, the program instructions executable by a processor to cause the processor to perform a method. The method includes reading a virtual machine pool manifest by a virtual machine pool manager. The virtual machine pool manifest includes an initial number of virtual machines to launch. The virtual machine pool manifest can also include how many virtual machines to allocate resources to, the number of resources to allocate based on a virtual machine resource definition, the number of virtual machines without a resource allocation, a default environment, and a parameter that indicates a maximum number of executable applications at a given time. The method also includes launching a plurality of base virtual machines into a virtual machine pool and based on the initial number provided by the virtual machine pool manifest. The base virtual machines lack initial application assignments. The method further includes allocating initial resources to a portion of the base virtual machines based on a virtual machine resource definition in the virtual machine pool manifest. The method includes loading core program packages into the portion of the base virtual machines.

[0005] Further embodiments are directed to a virtual machine pool management system for pre-pooling virtual machine components prior to the application startup and configured to perform the methods described above. The present summary is not intended to illustrate each aspect of every implementation of, and/or every embodiment of the present disclosure.

BRIEF DESCRIPTION OF THE DRAWINGS

[0006] These and other features, aspects, and advantages of the embodiments of the disclosure will become better understood with regard to the following description, appended claims, and accompanying drawings where:

[0007] FIG. 1 is a block diagram illustrating an in-vehicle device in a vehicle, in accordance with embodiments of the present disclosure.

[0008] FIG. 2 is a block diagram illustrating a virtual machine pool management system, in accordance with embodiments of the present disclosure.

[0009] FIG. 3 is a flow diagram of pre-pooling virtual machine components prior to application startup, in accordance with embodiments of the present disclosure.

[0010] FIG. 4 is a high-level block diagram illustrating an example computer system that may be used in implementing one or more of the methods, tools, and modules, and any related functions, described herein, in accordance with embodiments of the present disclosure.

[0011] FIG. 5 depicts a cloud computing environment, in accordance with embodiments of the present disclosure.

[0012] FIG. 6 depicts abstraction model layers, in accordance with embodiments of the present disclosure.

[0013] While the present disclosure is amenable to various modifications and alternative forms, specifics thereof have been shown by way of example, in the drawings and will be described in detail. It should be understood, however, that the intention is not to limit the particular embodiments described. On the contrary, the intention is to cover all modifications, equivalents, and alternatives falling within the scope of the present disclosure. Like reference numerals are used to designate like parts in the accompanying drawings.

DETAILED DESCRIPTION

[0014] The present disclosure relates to application startup speed in an in-vehicle device, and more specifically, to improving application startup speed by dividing the application launch process using efficient pool management techniques. While the present disclosure is not necessarily limited to such applications, various aspects of the disclosure may be appreciated through a discussion of various examples using this context.

[0015] Information of Things (IoT) technology allows for function extensibility by providing the capability of adding additional applications into a vehicle's in-vehicle device (e.g., a navigation system). Additionally, application development in vehicles is expanding to allow applications to control system functions (e.g., dashboard camera, air conditioner, vehicle information management) within a vehicle. These applications are developed as an integral part of the core functionality of an in-vehicle device allowing for system functions that can be installed as replaceable applications within the in-vehicle device.

[0016] A typical startup procedure for an in-vehicle application includes an application platform or a core program to operate an application, being initialized first. The initialization process includes reading an application from a storage device, generating a process, starting

the core program of the application, and allocating resources required for the application. The application is started when control is shifted over to the application logic.

[0017] Limitations on application startup remain, however, as in-vehicle devices have limited resources (e.g., memory, storage, processor). Thus, in-vehicle applications can have a resource upper-limit restriction. Current techniques to accelerate application startup include technologies such as hibernation. In hibernation, the operating system can copy the running state of an application to memory before powering off. Accordingly, the operating system can restart the application by restoring the copy of the running state upon power-up. However, in vehicles, power interruption is triggered by an accessory off command, which may not provide the operating system an opportunity to create a copy of running states of all the applications that may be executing. Further, there is no distinction between different types of power-offs for in-vehicle devices, as opposed to traditional computing devices. For example, notebook computers can differentiate between power off commands. A power-off command from a press of the power button can elicit a different response from the notebook as opposed to a power-off command when the notebook is closed. Distinguishing between different types of power-offs can make it possible for the operating system on such devices to hibernate applications that are currently executing when a device is going to sleep as opposed to shutting down. Because in-vehicle devices cannot distinguish a sleep-type power-off event from a shutdown power-off event, for example, technologies such as hibernation are not applicable for in-vehicle devices.

[0018] Other techniques, such as hardware performance improvement and program run speed using ahead-of-time compilation (AOT), attempt to also accelerate application startup. However, these techniques may incur additional overhead costs and merely provide limited improvement to the overall startup speed of an in-vehicle application.

[0019] Embodiments of the present disclosure may overcome the above and other problems by using a virtual machine pool management system for pre-pooling virtual machine components prior to application startup. Additionally, resource allocation and application logic shifts are performed for an application. They are thereby allowing the application pool management system to accelerate the startup time of application in an in-vehicle device. The virtual machine pool management system includes a virtual machine pool manager configured to manage and launch virtual machines needed by an in-vehicle device. The virtual machine pool manager is also configured to allocate initial resources to the virtual machines based on a virtual machine pool manifest. The virtual machine pool

management system also includes an application package manager configured to install and load core program packages into the virtual machines allocated with resources. These virtual machines can stay within a virtual machine pool and be in a ready state until they are requested by an in-vehicle device to execute an application.

[0020] In some embodiments, the virtual machine pool management system can launch additional virtual machines that do not receive an allocation of resources. The virtual machine pool management system can place these virtual machines into a wait configuration until such a time as they are needed. In this way, when resources become available, the virtual machine pool manager can allocate resources to the virtual machines in the wait configuration and place these virtual machines into a ready configuration.

[0021] In some embodiments, the virtual machine pool manager allocates initial resources to the virtual machines based on a resource assignment pattern. The resource assignment pattern may indicate a typical pattern an in-vehicle device has during operation. For example, an in-vehicle device may execute five applications immediately upon startup and then another six applications after a predetermined time. The resource assignment pattern can mimic the need of the in-vehicle device so to anticipate its need.

[0022] Embodiments of the present disclosure include an application life cycle manager configured to manage the life cycle of the applications being executed in an in-vehicle device. Additionally, the application life cycle manager can determine which virtual machines to assign applications to. This can be based on the initial resources and core programming loaded onto the virtual machine. For example, if a virtual machine is allocated with a sufficient amount of resources to allow for the operation and execution of a requested application, then the application life cycle manager can select that virtual machine and assign it that specific application. Otherwise, the application life cycle manager can request that the virtual machine pool manager launch additional virtual machines or reallocate appropriate resources.

[0023] Referring now to FIG. 1, shown is a high-level block diagram of a vehicle 100 with an embedded in-vehicle device 105. Additionally, the vehicle 100 includes a sensor 160, and a camera 170. The in-vehicle device 105 includes a control unit 110, a storage unit 120, a communication unit 130, an input unit 140, and an output unit 150. The input unit 140 is communicatively connected to the sensor 160 and the camera 170.

[0024] The in-vehicle device 105 is a component of the vehicle 100 configured to receive and transmit information relating to the vehicle 100. Additionally, the in-vehicle device 105 can operate vehicle system functions such as driving mode, cruise control, cameras, sensors, headlights, and the like. The in-vehicle device 105 can also operate various vehicle controls. These vehicle controls include, for example, navigation controls, heating, ventilation, and air conditioning (HVAC) controls, and audio/video (A/V) entertainment controls.

[0025] The control unit 110 is a component of the in-vehicle device 105 configured to use one or a plurality of central processing units (CPU) or a multi-core CPU and includes read-only memory (ROM), a random access memory (RAM), an input/output interface, a timer, and the like. The control unit 110 is a determining unit that controls operations of respective constituent units by executing a program stored in an embedded ROM, and executes the virtual machine pool management system 200 described below.

[0026] The storage unit 120 is a component of the in-vehicle device configured to store information pertaining to the vehicle 100 and the in-vehicle device 105. In some embodiments, the storage unit 120 is a non-volatile memory (e.g., flash memory, electrically erasable programmable read-only memory (EEPROM)). The storage unit 120 can include one or more memories. Each of the memories in the storage unit 120 is, for example, a semiconductor memory, a magnetic memory, solid-state memory, or optical memory. Each memory functions, for example, as the main storage device, an auxiliary storage device, or a cache memory. The storage unit 120 can store any information used for the operation of the in-vehicle device 105. For example, the storage unit 120 can store applications executable by the in-vehicle device 105.

[0027] The communication unit 130 is a component of the in-vehicle device 105 configured to communicate data with a server. The communication device 130 has one or more communication modules. The communication module includes, for example, a module compatible with mobile communication standards such as the 4th generation (4G), the 5th generation (5G), Bluetooth, Wi-Fi, global system for mobile (GSM), code division multiple access (CDMA), long-term evolution (LTE), infrared (IR), and the like. In addition, the communication unit can have a communication device, such as a data communication module (DCM). The communication unit 130 connects to the in-vehicle device 105 to a network to communicate with a server. In some embodiments, the communication unit 130 includes a communication module that is a global positioning system (GPS) receiving

module, in which the in-vehicle device 105 receives GPS signals by the communication unit 130.

[0028] The input unit 140 is a component of the in-vehicle device 105 that is an interface that inputs a signal from the outside of the in-vehicle device 105. The input unit 140 can receive information a sensor 160 or a camera 170, and the control unit 110 can acquire the information from the sensor 160 and the camera 170.

[0029] The output unit 150 is a component of the in-vehicle device configured to output a signal indicating an operation functionality the vehicle 100 is capable of. In some embodiments, the output unit 150 is a communication unit 130 connected to a local area network (LAN) or can be a communication unit 130 in combination with the input unit 140.

[0030] It is noted that FIG. 1 is intended to depict the major representative components of an exemplary in-vehicle device 105. In some embodiments, however, individual components may have greater or less complexity than as represented in FIG. 1, components other than or in addition to those shown in FIG. 1 may be present, and the number, type, and configuration of such components may vary.

[0031] FIG. 2 is a block diagram illustrating a virtual machine pool management system 200 for pre-pooling virtual machine components prior to the application startup, in accordance with embodiments of the present disclosure. The virtual machine pool management system 200 includes a virtual machine pool manager 210, a virtual machine pool manifest 220, a virtual machine pool 230, application life cycle manager 240, applications 250-1, 250-2, 250-N (collectively “applications 250”), where N is a variable integer representing any number of possible applications 250, and an application package manager 260. The virtual machine pool includes base virtual machine 233-1, 233-2, 233-N (collectively “base virtual machines 233”), where N is a variable integer representing any number of possible base virtual machines 233, and resources 235-1, 235-2, 235-N (collectively “resources 235”), where N is a variable integer representing any number of possible resources 235. The application lifecycle manager 240 includes an application logic injection manager 244 and an application resource assigner 248. The applications 250 include resource information 251-1, 251-2, 251-N (collectively “resource information 251”), where N is a variable integer representing any number of possible resource information 251.

[0032] The virtual machine pool manager 210 is a component of the virtual machine pool management system 200 configured to manage and launch virtual machines. The virtual machine pool manager 210 is responsible for resource management of the virtual machines. It can maintain a list of which designated applications 250 run on which base virtual machines 233. For example, the list can have an entry that designates that navigation software runs on navigation specific virtual machines and another entry that designates that non-user interface (UI) applications run on non-UI specific virtual machines.

[0033] The virtual machine pool manager 210 can run a series of virtual machine management utilities. These utilities include, for example, functions that: create, destroy, put to sleep, and wake up virtual machines. The utilities can also maintain a list that matches applications to virtual machines.

[0034] The virtual machine pool manifest 220 is a component of the virtual machine pool management system 200 configured to provide initial virtual machine establishment requirements. The virtual machine pool manifest 220 can include an initial number of base virtual machines 233 to be placed in the virtual machine pool 230. For example, the initial number may indicate that seven base virtual machines 233 are required at startup. Additionally, the virtual machine pool manifest 220 can include the number of base virtual machines 233 that are to be allocated with resources 235.

[0035] In some embodiments, the virtual machine pool manifest 220 indicates an allocation of resources 235 based on a resource definition. A resource definition defines the type of resources needed basing on the type of application loaded to the virtual machine 233. For example, a resource definition may indicate the resources needed for a display buffer. A display buffer can indicate that an application requires the use of a UI. Applications that may require UIs include, for example, navigation, radio, system settings, speedometer, and the like. Another resource definition may indicate that a display buffer is not used, which may be indicative of fewer resources being needed for a particular application. These types of applications may be background related applications that an in-vehicle device 105 may use during the operation of a vehicle 100.

[0036] In some embodiments, the virtual machine pool manifest 220 includes a default environment resource assignment. The default environment resource assignment may include a default allocation of resources when a resource definition is not defined for a particular virtual machine. The default resource assignment can indicate the amount of memory, storage, and CPU usage to be allocated to a default base virtual machine 233.

[0037] In some embodiments, the virtual machine manifest 220 includes a parameter indicating a threshold number of applications an in-vehicle device 105 can have running at any given time. The threshold number can be based on the resources available to the in-vehicle device 105 and manageable by the virtual machine pool management system 200.

[0038] The virtual machine pool 230 is a component of the virtual machine pool management system 200 configured to store base virtual machines 233. The virtual machine pool 230 can be any storage capable of storing software. In some embodiments, the virtual machine pool 230 stores a single virtual machine that can be configured by the virtual machine pool manager 210. The virtual machine pool 230 can store any number of distinct base virtual machines 233 that can be configured by the virtual machine pool manager 210.

[0039] The base virtual machines 233 are components of the virtual machine pool management system 200 configured to mimic the performance of a device (e.g., a computer). The base virtual machines 233 can be retrieved from the virtual machine pool 230. The base virtual machines 233 include resources 235 required to operate the base virtual machines 233. For example, the resources include the amount of memory, storage, and CPU usage required by the base virtual machines 233. Additionally, the base virtual machines 233 can also include a core program package loaded onto the base virtual machines 233 once they are launched onto the virtual machine pool 230. The core program package can provide a base operating system required to support and execute the applications 250 as well as various application components included in the base virtual machines 233.

[0040] In some embodiments, the base virtual machines 233 include a client that corresponds with a server. The client can communicate with the server via the base virtual machines 233 to create, clone, or otherwise access a virtual machine instance according to the various configuration definitions included in the virtual machine pool manifest 220.

[0041] The application life cycle manager 240 is a component of the virtual machine pool management system 200 configured to manage the life cycle of the applications 250. The application life cycle manager 240 can request a base virtual machine 233 to load an application on. Additionally, the application life cycle manager 240 can determine which applications are loaded in the order in which the applications are loaded. In some embodiments, the application life cycle manager 240 requests additional base virtual machines 233 to be added to the virtual machine pool 230. For example, numerous applications 250 can be in operation, and to avoid a slowdown in operation, the application life cycle manager 240 can request base virtual machines 233 in wait mode be allocated

resources and loaded with core program packages. Additional base virtual machines 233 can also be launched within the virtual machine pool 230 and placed in a wait mode configuration. These base virtual machines 233 may not have assigned resources or program logic.

[0042] The application logic injection manager 244 is a component of the application life cycle manager 240 configured to load unique application logic onto the base virtual machines 233. The application logic injection manager 244 can identify the application 250 that is to be loaded onto a base virtual machine 233 and load application-specific information to allow for the execution of the application 250. For example, application classes can be loaded, and injection points can be attached to receptors in the base virtual machine 233.

[0043] The application resource assigner 248 is a component of the application life cycle manager 240 configured to allocate resources required by an application 250 onto a virtual machine 233. The application resource assigner 248 can compare the initial resources allocated to a base virtual machine 233 with the resource information 251 of an application 250 that is being loaded onto the virtual machine 233. If the resources 235 differ, then the application resource assigner 248 can reallocate resources to the virtual machine 233. For example, a base virtual machine 233 may include 50 megabytes of storage and an application 250 is loaded onto the base virtual machine 233 requires 65 megabytes of storage, then the application resource assigner 248 can reallocate at least 65 megabytes of storage to the base virtual machine 250 to allow the application to be loaded and executed.

[0044] In some embodiments, if the initial resources adequately provide enough resources based on the resource information 251 for an application 250, then the application resource assigner 248 does not reallocate any resources. In some embodiments, if the initial resources adequately provide enough resources based on the resource information 251 for an application 250, then the application resource assigner 248 evaluates the resources and deallocates any unnecessary resources. For example, if the initial resources include 200 kilobytes of memory and the application 250 only requires 100 kilobytes of memory, then the application resource assigner 248 can deallocate at most 100 kilobytes of memory from the virtual machine 233.

[0045] The applications 250 are software of the virtual machine pool management system 200 that are capable of being executed by an in-vehicle device 105. The applications include features such as telematics services, road information, traffic information, weather information, fleet management, health checks, emergency assistance, music/video,

insurance, news, and infotainment. Each application 250 additionally includes resource information 251 relating to the application 250. The resource information 251 can include the number of resources required to execute and operate a particular application 250. For example, an application 250 that provides weather information may require a certain amount of memory, storage, an CPU usage that the resource information 251 can store.

[0046] The application package manager 260 is a component of the virtual machine pool management system 200 configured to install and update applications 250 provided by the virtual machine pool management system 200. The application package manager 260 can automate the process of installing, upgrading, configuring, and removing applications 250 from an in-vehicle device 105. The application package manager 260 can deal with packages that are distributions of software and data in archive files. Packages can contain metadata, such as an application's 250 name, description of its purpose, version number, vendor, checksum, and a list of dependencies necessary for the application 250 to run properly (e.g., resource information 251). Upon installation, metadata can be stored in a local package database. The application package manager 260 can maintain a database of application 250 dependencies and version information to prevent software mismatches and missing prerequisites.

[0047] It is noted that FIG. 2 is intended to depict the major representative components of an exemplary virtual machine pool management system 200. In some embodiments, however, individual components may have greater or less complexity than as represented in FIG. 2, components other than or in addition to those shown in FIG. 2 may be present, and the number, type, and configuration of such components may vary.

[0048] FIG. 3 is a flow diagram illustrating a process 300 of pre-pooling virtual machine components prior to application startup, in accordance with embodiments of the present disclosure. The process 300 begins by the virtual machine pool manager 210 reading the virtual machine pool manifest 220 upon initiation of the virtual machine pool management system 200. This is illustrated at step 310. The virtual machine pool manager 210 can read the initial virtual machine 233 requirements. These requirements can include, for example, an initial number of base virtual machines 233 to be placed in the virtual machine pool 230, an allocation of resources 235 for each virtual machine 233 based on a resource definition, a default environment resource assignment, and a parameter indicating a maximum number of applications an in-vehicle device 105 can have running at any given time.

[0049] The virtual machine pool manager 210 launches base virtual machines 233 into the virtual machine pool 230, as indicated in the virtual machine pool manifest 220. This is illustrated at step 320. The virtual machine pool 230 can store any number of distinct base virtual machines 233 that can be configured by the virtual machine pool manager 210. Upon launch, the base virtual machines 233 lack a resource allocation and do not have core program packages loaded. In some embodiments, the base virtual machines 233 are launched according to a timeline indicated by the virtual machine pool manifest 220. For example, the virtual machine pool manifest 220 may indicate that twelve base virtual machines 233 are launched upon initiation, and then another five base virtual machines 233 are to be launched after a predetermined time. This series of predetermined launches of base virtual machines 233 can occur as many times as the virtual machine pool manifest 220 is configured to do.

[0050] The virtual machine pool manager 210 allocates initial resources 235 to a portion of the base virtual machines 233 launched at step 320. This is illustrated at step 330. The base virtual machines 233 that are allocated with resources 235 can be predetermined by the virtual machine pool manifest 220. The number of base virtual machines 233 that are also allocated resources can also be an optimized number based on the number of applications 250 that are required by a vehicle 100 upon startup. These applications include, for example, as telematics services, road information, traffic information, weather information, fleet management, health checks, emergency assistance, music/video, insurance, news, and infotainment. The remaining portion of base virtual machines 233 can be placed in a wait configuration until such a time as to when they are needed.

[0051] By having base virtual machines 233 placed in a wait configuration, the virtual machine pool management system 200 need not spend additional time launching a virtual machine 233 when an application 250 is required to be launched. Additionally, not all of the base virtual machines 233 that are allocated with resources 235 may not be assigned an application immediately, and instead can act as a buffer during operation of the in-vehicle device 105. As base virtual machines 233 are requested and used by applications 250, more virtual machines can be launched into the virtual machine pool 230 to maintain a buffer and speed up the startup process of an application 250.

[0052] The application package manager 260 loads core program packages onto the portion of base virtual machines 233 allocated with resources 235. This is illustrated at step 340. The core program packages can include a Java virtual machine environment without

application logic. Other core program packages include, for example, virtual machine base code, class libraries (e.g., Java), service application programming interface (API), and UI life cycle management software. The core program packages can include receptors that can represent shells for application logic to load into. Once a virtual machine 233 has been allocated resources 235 and loaded with core program packages, it is placed into a ready configuration and can be assigned and loaded with an application.

[0053] The application life cycle manager 240 assigns an application 250 a base virtual machine 233 that is in a ready configuration. This is illustrated at step 350. In some embodiments, the application life cycle manager 240 determines which applications are loaded in the order in which the applications are loaded. The order in which the application life cycle manager 240 selects and assigns an application 250 can be based on needing an urgency of an application 250 needed by an in-vehicle device 105. For example, an application 250 that provides system diagnostics upon startup of a vehicle 100 can be assigned a virtual machine 233 prior to an application 250 that provides infotainment to a passenger.

[0054] In some embodiments, the application life cycle manager 250 assigns an application to a base virtual machine 233 in a ready configuration based on the application type relating to the application 250 being launched. For example, an application 250 that does not require a UI may be launched first because the resources information 251 of the application 250 matches that of the initial resources 235 allocated to the base virtual machine 233. Thus, no additional resources 235 need to be reallocated to the virtual machine 233, and the application 250 can launch faster than an application 250 that may require additional resources.

[0055] The application logic injection manager 244 launches the application 250 in the base virtual machine 233. This is illustrated at step 360. Application logic includes, for example, application-specific classes, methods, and additional code. The application logic injection manager 244 can load application logic relating to the application 250 into the virtual machine 233. Additionally, the injection points can be attached to the receptors included in the virtual machine 233 that were loaded with the core program packages.

[0056] In some embodiments, the application resource assigner 248 analyzes the resource information 251 of the application 250 and initial resources 235 assigned to the virtual machine 233. Based on the analysis, the application resource assigner can reassign resources 235 to the virtual machine 233 to allow the application to launch successfully. For example,

if the initial resources 235 contain 45 kilobytes of available memory and the application 250 is launched requires 100 kilobytes of memory, then the application resource assigner 248 can reallocate additional memory to the virtual machine 233 so that the application 250 can properly launch.

[0057] In some embodiments, the application resource assigner 248 analyzes the resource information 251 of the application 250 and initial resources 235 assigned to the virtual machine 233. Based on the analysis, the application resource assigner can reassign unneeded resources 235 back to a resource pool. Once the application 250 has the application logic loaded and injection points attached, the application life cycle manager 240 executes the application 250. This is illustrated at step 370. The application life cycle manager 240 can monitor the application 250 during operation and execution and determine if and when to terminate the application 250.

[0058] Referring now to FIG. 4, shown is a high-level block diagram of an example computer system 400 (e.g., the virtual machine pool management system 100) that may be used in implementing one or more of the methods, tools, and modules, and any related functions, described herein (e.g., using one or more processor circuits or computer processors of the computer), in accordance with embodiments of the present disclosure. In some embodiments, the major components of the computer system 400 may comprise one or more processors 402, a memory 404, a terminal interface 412, an I/O (Input/Output) device interface 414, a storage interface 416, and a network interface 418, all of which may be communicatively coupled, directly or indirectly, for inter-component communication via a memory bus 403, an I/O bus 408, and an I/O bus interface 410.

[0059] The computer system 400 may contain one or more general-purpose programmable central processing units (CPUs) 402-1, 402-2, 402-3, and 402-N, herein generically referred to as the processor 402. In some embodiments, the computer system 400 may contain multiple processors typical of a relatively large system; however, in other embodiments, the computer system 400 may alternatively be a single CPU system. Each processor 401 may execute instructions stored in the memory 404 and may include one or more levels of on-board cache.

[0060] The memory 404 may include computer system readable media in the form of volatile memory, such as random-access memory (RAM) 422 or cache memory 424. Computer system 400 may further include other removable/non-removable, volatile/non-volatile computer system storage media. By way of example only, storage system 426 can

be provided for reading from and writing to a non-removable, non-volatile magnetic media, such as a “hard drive.” Although not shown, a magnetic disk drive for reading from and writing to a removable, non-volatile magnetic disk (e.g., a “floppy disk”), or an optical disk drive for reading from or writing to a removable, non-volatile optical disc such as a CD-ROM, DVD-ROM or other optical media can be provided. In addition, the memory 404 can include flash memory, e.g., a flash memory stick drive or a flash drive. Memory devices can be connected to memory bus 403 by one or more data media interfaces. The memory 404 may include at least one program product having a set (e.g., at least one) of program modules that are configured to carry out the functions of various embodiments.

[0061] Although the memory bus 403 is shown in FIG. 4 as a single bus structure providing a direct communication path among the processors 402, the memory 404, and the I/O bus interface 410, the memory bus 403 may, in some embodiments, include multiple different buses or communication paths, which may be arranged in any of various forms, such as point-to-point links in hierarchical, star or web configurations, multiple hierarchical buses, parallel and redundant paths, or any other appropriate type of configuration. Furthermore, while the I/O bus interface 410 and the I/O bus 408 are shown as single respective units, the computer system 400 may, in some embodiments, contain multiple I/O bus interface units, multiple I/O buses, or both. Further, while multiple I/O interface units are shown, which separate the I/O bus 408 from various communications paths running to the various I/O devices, in other embodiments some or all of the I/O devices may be connected directly to one or more system I/O buses.

[0062] In some embodiments, the computer system 400 may be a multi-user mainframe computer system, a single-user system, or a server computer or similar device that has little or no direct user interface but receives requests from other computer systems (clients). Further, in some embodiments, the computer system 400 may be implemented as a desktop computer, portable computer, laptop or notebook computer, tablet computer, pocket computer, telephone, smartphone, network switches or routers, or any other appropriate type of electronic device.

[0063] It is noted that FIG. 4 is intended to depict the major representative components of an exemplary computer system 400. In some embodiments, however, individual components may have greater or lesser complexity than as represented in FIG. 4, components other than or in addition to those shown in FIG. 4 may be present, and the number, type, and configuration of such components may vary.

[0064] One or more programs/utilities 428, each having at least one set of program modules 430 (e.g., the virtual machine pool management system 100) may be stored in memory 404. The programs/utilities 428 may include a hypervisor (also referred to as a virtual machine monitor), one or more operating systems, one or more application programs, other program modules, and program data. Each of the operating systems, one or more application programs, other program modules, and program data or some combination thereof, may include an implementation of a networking environment. Programs 428 and/or program modules 430 generally perform the functions or methodologies of various embodiments.

[0065] It is to be understood that although this disclosure includes a detailed description on cloud computing, implementation of the teachings recited herein is not limited to a cloud computing environment. Rather, embodiments of the present invention are capable of being implemented in conjunction with any other type of computing environment now known or later developed.

[0066] Cloud computing is a model of service delivery for enabling convenient, on-demand network access to a shared pool of configurable computing resources (e.g., networks, network bandwidth, servers, processing, memory, storage, applications, virtual machines, and services) that can be rapidly provisioned and released with minimal management effort or interaction with a provider of the service. This cloud model may include at least five characteristics, at least three service models, and at least four deployment models.

[0067] Characteristics are as follows:

[0068] On-demand self-service: a cloud consumer can unilaterally provision computing capabilities, such as server time and network storage, as needed automatically without requiring human interaction with the service's provider.

[0069] Broad network access: capabilities are available over a network and accessed through standard mechanisms that promote use by heterogeneous thin or thick client platforms (e.g., mobile phones, laptops, and PDAs).

[0070] Resource pooling: the provider's computing resources are pooled to serve multiple consumers using a multi-tenant model, with different physical and virtual resources dynamically assigned and reassigned according to demand. There is a sense of location

independence in that the consumer generally has no control or knowledge over the exact location of the provided resources but may be able to specify location at a higher level of abstraction (e.g., country, state, or datacenter).

[0071] Rapid elasticity: capabilities can be rapidly and elastically provisioned, in some cases automatically, to quickly scale out and rapidly released to quickly scale in. To the consumer, the capabilities available for provisioning often appear to be unlimited and can be purchased in any quantity at any time.

[0072] Measured service: cloud systems automatically control and optimize resource use by leveraging a metering capability at some level of abstraction appropriate to the type of service (e.g., storage, processing, bandwidth, and active user accounts). Resource usage can be monitored, controlled, and reported, providing transparency for both the provider and consumer of the utilized service.

[0073] Service Models are as follows:

[0074] Software as a Service (SaaS): the capability provided to the consumer is to use the provider's applications running on a cloud infrastructure. The applications are accessible from various client devices through a thin client interface such as a web browser (e.g., web-based e-mail). The consumer does not manage or control the underlying cloud infrastructure including network, servers, operating systems, storage, or even individual application capabilities, with the possible exception of limited user-specific application configuration settings.

[0075] Platform as a Service (PaaS): the capability provided to the consumer is to deploy onto the cloud infrastructure consumer-created or acquired applications created using programming languages and tools supported by the provider. The consumer does not manage or control the underlying cloud infrastructure including networks, servers, operating systems, or storage, but has control over the deployed applications and possibly application hosting environment configurations.

[0076] Infrastructure as a Service (IaaS): the capability provided to the consumer is to provision processing, storage, networks, and other fundamental computing resources where the consumer is able to deploy and run arbitrary software, which can include operating systems and applications. The consumer does not manage or control the underlying cloud

infrastructure but has control over operating systems, storage, deployed applications, and possibly limited control of select networking components (e.g., host firewalls).

[0077] Deployment Models are as follows:

[0078] Private cloud: the cloud infrastructure is operated solely for an organization. It may be managed by the organization or a third party and may exist on-premises or off-premises.

[0079] Community cloud: the cloud infrastructure is shared by several organizations and supports a specific community that has shared concerns (e.g., mission, security requirements, policy, and compliance considerations). It may be managed by the organizations or a third party and may exist on-premises or off-premises.

[0080] Public cloud: the cloud infrastructure is made available to the general public or a large industry group and is owned by an organization selling cloud services.

[0081] Hybrid cloud: the cloud infrastructure is a composition of two or more clouds (private, community, or public) that remain unique entities but are bound together by standardized or proprietary technology that enables data and application portability (e.g., cloud bursting for load-balancing between clouds).

[0082] A cloud computing environment is service-oriented with a focus on statelessness, low coupling, modularity, and semantic interoperability. At the heart of cloud computing is an infrastructure that includes a network of interconnected nodes.

[0083] Referring now to FIG. 5, illustrative cloud computing environment 500 is depicted. As shown, cloud computing environment 500 includes one or more cloud computing nodes 510 with which local computing devices used by cloud consumers, such as, for example, personal digital assistant (PDA) or cellular telephone 520-1, desktop computer 520-2, laptop computer 520-3, and/or automobile computer system 520-4 may communicate. Nodes 510 may communicate with one another. They may be grouped (not shown) physically or virtually, in one or more networks, such as Private, Community, Public, or Hybrid clouds as described hereinabove, or a combination thereof. This allows cloud computing environment 500 to offer infrastructure, platforms and/or software as services for which a cloud consumer does not need to maintain resources on a local computing device. It is understood that the types of computing devices 520-1 to 520-4 shown in FIG. 5 are intended to be illustrative only and that computing nodes 510 and cloud

computing environment 500 can communicate with any type of computerized device over any type of network and/or network addressable connection (e.g., using a web browser).

[0084] Referring now to FIG. 6, a set of functional abstraction layers 600 provided by cloud computing environment 500 (FIG. 5) is shown. It should be understood in advance that the components, layers, and functions shown in FIG. 6 are intended to be illustrative only and embodiments of the invention are not limited thereto. As depicted, the following layers and corresponding functions are provided:

[0085] Hardware and software layer 610 includes hardware and software components. Examples of hardware components include mainframes 611; RISC (Reduced Instruction Set Computer) architecture-based servers 612; servers 613; blade servers 614; storage devices 615; and networks and networking components 616. In some embodiments, software components include network application server software 617 and database software 618.

[0086] Virtualization layer 620 provides an abstraction layer from which the following examples of virtual entities may be provided: virtual servers 621; virtual storage 622; virtual networks 623, including virtual private networks; virtual applications and operating systems 624; and virtual clients 625.

[0087] In one example, management layer 630 may provide the functions described below. Resource provisioning 631 provides dynamic procurement of computing resources and other resources that are utilized to perform tasks within the cloud computing environment. Metering and Pricing 632 provide cost tracking as resources are utilized within the cloud computing environment, and billing or invoicing for consumption of these resources. In one example, these resources may include application software licenses. Security provides identity verification for cloud consumers and tasks, as well as protection for data and other resources. User portal 633 provides access to the cloud computing environment for consumers and system administrators. Service level management 634 provides cloud computing resource allocation and management such that required service levels are met. Service Level Agreement (SLA) planning and fulfillment 635 provide pre-arrangement for, and procurement of, cloud computing resources for which a future requirement is anticipated in accordance with an SLA.

[0088] Workloads layer 640 provides examples of functionality for which the cloud computing environment may be utilized. Examples of workloads and functions which may be provided from this layer include mapping and navigation 641; software development and

lifecycle management 642 (e.g., the virtual machine pool management system 100); virtual classroom education delivery 643; data analytics processing 644; transaction processing 645; and precision cohort analytics 646.

[0089] The present invention may be a system, a method, and/or a computer program product at any possible technical detail level of integration. The computer program product may include a computer-readable storage medium (or media) having computer readable program instructions thereon for causing a processor to carry out aspects of the present invention.

[0090] The computer-readable storage medium can be a tangible device that can retain and store instructions for use by an instruction execution device. The computer-readable storage medium may be, for example, but is not limited to, an electronic storage device, a magnetic storage device, an optical storage device, an electromagnetic storage device, a semiconductor storage device, or any suitable combination of the foregoing. A non-exhaustive list of more specific examples of the computer-readable storage medium includes the following: a portable computer diskette, a hard disk, a random access memory (RAM), a read-only memory (ROM), an erasable programmable read-only memory (EPROM or Flash memory), a static random access memory (SRAM), a portable compact disc read-only memory (CD-ROM), a digital versatile disk (DVD), a memory stick, a floppy disk, a mechanically encoded device such as punch-cards or raised structures in a groove having instructions recorded thereon, and any suitable combination of the foregoing. A computer-readable storage medium, as used herein, is not to be construed as being transitory signals per se, such as radio waves or other freely propagating electromagnetic waves, electromagnetic waves propagating through a waveguide or other transmission media (e.g., light pulses passing through a fiber-optic cable), or electrical signals transmitted through a wire.

[0091] Computer-readable program instructions described herein can be downloaded to respective computing/processing devices from a computer-readable storage medium or to an external computer or external storage device via a network, for example, the Internet, a local area network, a wide area network and/or a wireless network. The network may comprise copper transmission cables, optical transmission fibers, wireless transmission, routers, firewalls, switches, gateway computers and/or edge servers. A network adapter card or network interface in each computing/processing device receives computer readable program instructions from the network and forwards the computer readable program

instructions for storage in a computer readable storage medium within the respective computing/processing device.

[0092] Computer readable program instructions for carrying out operations of the present invention may be assembler instructions, instruction-set-architecture (ISA) instructions, machine instructions, machine dependent instructions, microcode, firmware instructions, state-setting data, configuration data for integrated circuitry, or either source code or object code written in any combination of one or more programming languages, including an object oriented programming language such as Smalltalk, C++, or the like, and procedural programming languages, such as the “C” programming language or similar programming languages. The computer readable program instructions may execute entirely on the user’s computer, partly on the user’s computer, as a stand-alone software package, partly on the user’s computer and partly on a remote computer or entirely on the remote computer or server. In the latter scenario, the remote computer may be connected to the user’s computer through any type of network, including a local area network (LAN) or a wide area network (WAN), or the connection may be made to an external computer (for example, through the Internet using an Internet Service Provider). In some embodiments, electronic circuitry including, for example, programmable logic circuitry, field-programmable gate arrays (FPGA), or programmable logic arrays (PLA) may execute the computer readable program instructions by utilizing state information of the computer readable program instructions to personalize the electronic circuitry, in order to perform aspects of the present invention.

[0093] Aspects of the present invention are described herein with reference to flowchart illustrations and/or block diagrams of methods, apparatus (systems), and computer program products according to embodiments of the invention. It will be understood that each block of the flowchart illustrations and/or block diagrams, and combinations of blocks in the flowchart illustrations and/or block diagrams, can be implemented by computer readable program instructions.

[0094] These computer readable program instructions may be provided to a processor of a computer, or other programmable data processing apparatus to produce a machine, such that the instructions, which execute via the processor of the computer or other programmable data processing apparatus, create means for implementing the functions/acts specified in the flowchart and/or block diagram block or blocks. These computer readable program instructions may also be stored in a computer readable storage medium that can direct a computer, a programmable data processing apparatus, and/or other devices to function in a

particular manner, such that the computer readable storage medium having instructions stored therein comprises an article of manufacture including instructions which implement aspects of the function/act specified in the flowchart and/or block diagram block or blocks.

[0095] The computer readable program instructions may also be loaded onto a computer, other programmable data processing apparatus, or other device to cause a series of operational steps to be performed on the computer, other programmable apparatus or other device to produce a computer implemented process, such that the instructions which execute on the computer, other programmable apparatus, or other device implement the functions/acts specified in the flowchart and/or block diagram block or blocks.

[0096] The flowchart and block diagrams in the Figures illustrate the architecture, functionality, and operation of possible implementations of systems, methods, and computer program products according to various embodiments of the present invention. In this regard, each block in the flowchart or block diagrams may represent a module, segment, or portion of instructions, which comprises one or more executable instructions for implementing the specified logical function(s). In some alternative implementations, the functions noted in the blocks may occur out of the order noted in the Figures. For example, two blocks shown in succession may, in fact, be accomplished as one step, executed concurrently, substantially concurrently, in a partially or wholly temporally overlapping manner, or the blocks may sometimes be executed in the reverse order, depending upon the functionality involved. It will also be noted that each block of the block diagrams and/or flowchart illustration, and combinations of blocks in the block diagrams and/or flowchart illustration, can be implemented by special purpose hardware-based systems that perform the specified functions or acts or carry out combinations of special purpose hardware and computer instructions.

[0097] The terminology used herein is for the purpose of describing particular embodiments only and is not intended to be limiting of the various embodiments. As used herein, the singular forms “a,” “an,” and “the” are intended to include the plural forms as well, unless the context clearly indicates otherwise. It will be further understood that the terms “includes” and/or “including,” when used in this specification, specify the presence of the stated features, integers, steps, operations, elements, and/or components, but do not preclude the presence or addition of one or more other features, integers, steps, operations, elements, components, and/or groups thereof. In the previous detailed description of example embodiments of the various embodiments, reference was made to the

accompanying drawings (where like numbers represent like elements), which form a part hereof, and in which is shown by way of illustration specific example embodiments in which the various embodiments may be practiced. These embodiments were described in sufficient detail to enable those skilled in the art to practice the embodiments, but other embodiments may be used and logical, mechanical, electrical, and other changes may be made without departing from the scope of the various embodiments. In the previous description, numerous specific details were set forth to provide a thorough understanding the various embodiments. But the various embodiments may be practiced without these specific details. In other instances, well-known circuits, structures, and techniques have not been shown in detail in order not to obscure embodiments.

[0098] When different reference numbers comprise a common number followed by differing letters (e.g., 100a, 100b, 100c) or punctuation followed by differing numbers (e.g., 100-1, 100-2, or 100.1, 100.2), use of the reference character only without the letter or following numbers (e.g., 100) may refer to the group of elements as a whole, any subset of the group, or an example specimen of the group.

[0099] Further, the phrase “at least one of,” when used with a list of items, means different combinations of one or more of the listed items can be used, and only one of each item in the list may be needed. In other words, “at least one of” means any combination of items and number of items may be used from the list, but not all of the items in the list are required. The item can be a particular object, a thing, or a category.

[00100] For example, without limitation, “at least one of item A, item B, or item C” may include item A, item A and item B, or item B. This example also may include item A, item B, and item C or item B and item C. Of course, any combinations of these items can be present. In some illustrative examples, “at least one of” can be, for example, without limitation, two of item A; one of item B; and ten of item C; four of item B and seven of item C; or other suitable combinations.

[00101] Different instances of the word “embodiment” as used within this specification do not necessarily refer to the same embodiment, but they may. Any data and data structures illustrated or described herein are examples only, and in other embodiments, different amounts of data, types of data, fields, numbers and types of fields, field names, numbers and types of rows, records, entries, or organizations of data may be used. In addition, any data may be combined with logic, so that a separate data structure may not be necessary. The previous detailed description is, therefore, not to be taken in a limiting sense.

[00102] The descriptions of the various embodiments of the present invention have been presented for purposes of illustration but are not intended to be exhaustive or limited to the embodiments disclosed. Many modifications and variations will be apparent to those of ordinary skill in the art without departing from the scope and spirit of the described embodiments. The terminology used herein was chosen to best explain the principles of the embodiments, the practical application or technical improvement over technologies found in the marketplace, or to enable others of ordinary skill in the art to understand the embodiments disclosed herein.

[00103] Although the present invention has been described in terms of specific embodiments, it is anticipated that alterations and modification thereof will become apparent to the skilled in the art. Therefore, it is intended that the following claims be interpreted as covering all such alterations and modifications as fall within the true spirit and scope of the invention.

CLAIMS

1. A computer-implemented method of pre-pooling virtual machine components prior to in-vehicle application startup, the computer-implemented method comprising:
 - reading a virtual machine pool manifest by a virtual machine pool manager, wherein the virtual machine pool manifest includes an initial number of virtual machines to launch;
 - launching a plurality of base virtual machines into a virtual machine pool and based on the initial number prior to launching an application;
 - allocating initial resources to a portion of the base virtual machines based on a virtual machine resource definition; and
 - loading core program packages into the portion of the base virtual machines.
2. The computer-implemented method according to the preceding claim, further comprising:
 - assigning the application to a first base virtual machine from the portion of the base virtual machines into an application life cycle manager based on an application type relating to the application;
 - launching the application in the first base virtual machine; and
 - executing the application.
3. The computer-implemented method according to the preceding claim, wherein assigning the application to the first base virtual machine is also based on the initial resources allocated to the first base virtual machine.
4. The computer-implemented method according to any of the preceding claims and with features of claim 2, wherein launching the application comprises:
 - loading an application logic relating to the application, wherein the application logic includes injection points;
 - attaching the injection points to receptors included in the first base virtual machine; and
 - reassigning resources to the first base virtual machine based on a resource information relating to the application.

5. The computer-implemented method according to any of the preceding claims, wherein the plurality of base virtual machines other than the portion of the base virtual machines are placed into a wait configuration.
6. The computer-implemented method according to any of the preceding claims, wherein allocating the initial resources is based on a resource assignment pattern.
7. The computer-implemented method according to any of the preceding claims, wherein the core program packages include a Java runtime environment.
8. The computer-implemented method according to any of the preceding claims, wherein the application is an in-vehicle application.
9. A computer program product of pre-pooling virtual machine components prior to application startup, the computer program product comprising:
 - one or more computer readable storage media, and program instructions stored on the one or more computer readable storage media, the program instructions comprising:
 - program instructions to read a virtual machine pool manifest by a virtual machine pool manager, wherein the virtual machine pool manifest includes an initial number of virtual machines to launch;
 - program instructions to launch a plurality of base virtual machines into a virtual machine pool and based on the initial number prior to launching an application;
 - program instructions to allocate initial resources to a portion of the base virtual machines based on a virtual machine resource definition; and
 - program instructions to load core program packages into the portion of the base virtual machines.
10. The computer program product according to the preceding claim, further comprising:

program instructions to assign the application to a first base virtual machine from the portion of the base virtual machines into an application life cycle manager based on an application type relating to the application;

program instructions to launch the application in the first base virtual machine; and

program instructions to execute the application.

11. The computer program product according to the preceding claim, wherein the program instructions to assign the application to the first base virtual machine is also based on the initial resources allocated to the first base virtual machine.
12. The computer program product according to any of the three preceding claims and with features of claim 10, wherein the program instructions to launch the application comprises:
 - program instructions to load an application logic relating to the application, wherein the application logic includes injection points;
 - program instructions to attach the injection points to receptors included in the base virtual machine; and
 - program instructions to reassign resources to the first base virtual machine based on a resource information relating to the application.
13. The computer program product according to any of the four preceding claims, wherein the base virtual machines other than the portion of the base virtual machines are placed into a wait configuration.
14. The computer program product according to any of the five preceding claims, wherein the program instructions to allocate the initial resources is based on a resource assignment pattern.
15. The computer program product according to any of the six preceding claims, wherein the core program packages include a Java runtime environment.

16. The computer program product according to any of the seven preceding claims, wherein the application is an in-vehicle application.
17. A system for pre-pooling virtual machine components prior to application startup, the system comprising:
 - a memory;
 - a processor;
 - a storage having stored thereon computer executable program code;
 - a virtual machine pool manager configured to launch a plurality of base virtual machines into a virtual machine pool, the virtual machine pool manager further configured to allocate initial resources to the base virtual machines and load core program packages onto the base virtual machines;
 - an application life cycle manager configured to assign applications to the base virtual machines; and
 - an application package manager configured to install and update the applications.
18. The system according to the preceding claim, wherein a portion of the base virtual machines are placed into a wait configuration.
19. The system according to any of the two preceding claims, wherein the initial resources are based on a resource assignment pattern.
20. The system according to any of the three preceding claims, wherein the applications allocated to the base virtual machines are also based on the initial resources allocated to the base virtual machines.




FIG. 1




FIG. 2




FIG. 3




FIG. 4




FIG. 5




FIG. 6

INTERNATIONAL SEARCH REPORT

International application No.

PCT/IB2021/056289

A. CLASSIFICATION OF SUBJECT MATTER

G06F 9/44(2018.01)i

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

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)

CNKI,CNPAT,WPI,EPODOC:virtual,machine,pool,startup,launch+,allocat+,load+, vehicle

C. DOCUMENTS CONSIDERED TO BE RELEVANT

Category*	Citation of document, with indication, where appropriate, of the relevant passages	Relevant to claim No.
A	US 2018239624 A1 (RED HAT INC.) 23 August 2018 (2018-08-23) the whole document	1-20
A	US 2015324216 A1 (NETAPP INC.) 12 November 2015 (2015-11-12) the whole document	1-20
A	US 9600316 B2 (MICROSOFT TECHNOLOGY LICENSING L.L.C.) 21 March 2017 (2017-03-21) the whole document	1-20
A	US 2016048409 A1 (BRIEL, Peter et al.) 18 February 2016 (2016-02-18) the whole document	1-20
A	US 9239814 B2 (ORACLE INT. CORP.) 19 January 2016 (2016-01-19) the whole document	1-20
A	CN 108292233 A (INTEL CORP.) 17 July 2018 (2018-07-17) the whole document	1-20

☐ 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

13 October 2021

Date of mailing of the international search report

27 October 2021

Name and mailing address of the ISA/CN

National Intellectual Property Administration, PRC
6, Xitucheng Rd., Jimen Bridge, Haidian District, Beijing
100088
China

Authorized officer

WANG,Bo

Facsimile No. (86-10)62019451

Telephone No. 010-53962643

INTERNATIONAL SEARCH REPORT
Information on patent family members

International application No.

PCT/IB2021/056289

Patent document cited in search report			Publication date (day/month/year)	Patent family member(s)		Publication date (day/month/year)
US	2018239624	A1	23 August 2018	US	10445126 B2	15 October 2019
US	2015324216	A1	12 November 2015	None		
US	9600316	B2	21 March 2017	US	2014109088 A1	17 April 2014
US	2016048409	A1	18 February 2016	None		
US	9239814	B2	19 January 2016	US	2014143389 A1	22 May 2014
				US	8639787 B2	28 January 2014
				US	2012266169 A1	18 October 2012
CN	108292233	A	17 July 2018	US	2017177377 A1	22 June 2017
				US	9898307 B2	20 February 2018
				WO	2017112126 A1	29 June 2017
				DE	112016005868 T5	11 October 2018