



(51) International Patent Classification:

G06F 11/34 (2006.01) G06F 11/30 (2006.01)

(21) International Application Number:

PCT/US2018/020815

(22) International Filing Date:

03 March 2018 (03.03.2018)

(25) Filing Language:

English

(26) Publication Language:

English

(30) Priority Data:

15/621,901 13 June 2017 (13.06.2017) US

(71) Applicant: **WESTERN DIGITAL TECHNOLOGIES, INC.** [US/US]; 5601 Great Oaks Parkway, San Jose, CA 95119 (US).

(72) Inventors: **DEMULDER, Timothy**; C/o Western Digital Technologies, Inc., 5601 Great Oaks Parkway, San Jose, CA 95119 (US). **DE ZAEYTIJD, Ruben**; C/o Western Digital Technologies, Inc., 5601 Great Oaks Parkway, San Jose, CA 95119 (US).

(74) Agent: **BOHN, Michel, E.** et al.; Patent Law Works LLP, 310 East 4500 South, Suite 400, Salt Lake City, UT 84107 (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, 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, SM, ST, SV, SY, TH, TJ, TM, TN, TR, TT, TZ, UA, UG, US, UZ, VC, VN, 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, 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))

(54) Title: RULE-BASED MONITORING ENGINE WITH TRACING CAPABILITIES FOR MULTI-THREADED LOGGING




Figure 1

(57) Abstract: A system and method embodied in a monitoring engine configured to create plug-ins to define functionality including rule-based configuration syntax adapted to intuitively and automatically launch the plug-ins as necessary in a software-based infrastructure or environment. The monitoring engine also includes a capability for tracing multi-threaded asynchronous logging onto a single file, by assigning a unique identification to each thread or rule initiated and facilitating each log according to a specific format and to formulate a combination of all the different thread identifications into a master unique identification that is easily traceable through an entire log file.

RULE-BASED MONITORING ENGINE WITH TRACING

CAPABILITIES FOR MULTI-THREADED LOGGING

5

10

BACKGROUND

15

[001] The present disclosure relates to systems and methods implemented in a monitoring engine configured with plug-ins to define functionality including rule-based configuration syntax adapted to intuitively and automatically deploy or launch the plug-ins as necessary in any software environment that requires a monitoring agent or daemon, a cloud infrastructure or the like. The present disclosure also relates to tracing multi-threaded asynchronous logging to a single file, by assigning a unique identification to each thread or process initiated and facilitating each logging according to a specific format to formulate a combination of all the different thread identifications into a master unique identification that is easily traceable through an entire log file.

20

[002]

25

[003] Typically, monitoring agents and engines are intelligent agent software that is used to observe and report on computer equipment or operations of complex computer networks. Monitoring engines may be used in any software-based environment or cloud infrastructure to monitor particular functionalities and applications. For example, organizations deploying cloud-computing services use monitoring tools to ensure that these services are operating smoothly. Existing monitoring agent designs are inflexible or do not permit those who have little programming knowledge to either quickly or easily come up with new monitoring functionality to monitor a particular operation within the cloud infrastructure because the actual functionality is tied to the monitoring engine code.

30

[004] Further, in a cloud infrastructure environment, multithreading is common practice in multitasking operating systems. Multithreading as a widespread programming and execution technique allows multiple threads to exist within the context of one process or operation that is being performed. These multiple threads share resources (hardware and

software), yet execute independently. This multithreading technology when applied to a single process enables parallel execution on a multiprocessor system. A single process may have many different functions executing concurrently, allowing the application to better use the available hardware (multiple cores/processors). By this, threads can communicate
5 between them, as they share the same memory. A significant problem that often arises in this situation is that every thread does not function in the same way when accessing shared objects and memory. Thus, although threading certainly helps with addressing complex, lengthy, independent problems, it introduces complexity in the operations, including use of multiple personal-identification numbers and internal-thread names. Multi-threaded
10 asynchronous logging to a single file can be difficult to untangle as the log messages are serialized in the actual file, but the source, if not always clear, for identifying time and order or logged events, becomes complicated and it becomes even more difficult when the child thread itself forks other processes.

[005] The present disclosure solves problems associated with ineffective monitoring
15 engines used in a cloud infrastructure, by providing a flexible platform capable of separating the actual engine from the functionality providers, i.e., the plug-ins, by making the plug-ins easily configurable for any particular functionality desired or for other new functionalities by defining functionality from a predetermined rule-based program. The present disclosure also solves the complexities of troubleshooting multi-threaded processes, by assigning, a unique
20 identification to each thread, process, or step of a rule and providing a logging interface in a specific format, so that the combination of all the identification numbers for the various threads or processes formulate a combined and unique identification, that is easily traceable through the entire log file.

SUMMARY

[006] According to one innovative aspect of the subject matter described in this disclosure, a monitoring engine created for use in a cloud infrastructure is configurable and executable with simple plug-ins (otherwise also referred to as plugins, add-ons, or extensions) that define functionality and use simple rule-based configuration syntax to launch or deploy the plug-ins
30 in an intuitive way. The writing instructions for this monitoring functionality are reduced to simply creating at least two plug-ins (e.g., a checker and an action plug-in) and using them in a rule. The present invention also solves the complexities of troubleshooting multi-threaded processes in a cloud infrastructure, by assigning a unique identification to each thread or process and providing a logging interface in a specific format so that the combination of all

the identification numbers for the various threads or processes (rules) formulate a combined, unique identification that is easily traceable through the entire log file and may be easily comprehended.

[007] According to another innovative aspect of the subject matter in this disclosure, the monitoring engine defines three plug-in categories, namely, checkers, validators, and actions and adds or installs a plug-in for each category. In some embodiments, the checker plug-in is tasked to return a certain value. In some embodiments, the validator plug-in is tasked to accept a value coming from a checker, compare it to something else, and return the status of the validation. In some embodiments, the action plug-in is tasked with performing a task.

The plug-ins serve in a rule that is configured to follow a simple syntax: 'checker' 'operator' 'validator': 'action0', 'action1'... This means when this rule is deployed or launched, the monitoring engine launches an instance of 'checker' and compares the returned value from the checker to 'validator', and if the result of the 'checker' 'operator' 'validator' is true, the engine launches the corresponding tasks (multiple actions defined by a single rule), e.g., 'action0', 'action1' and if the result of the first part is false, the engine does not launch the tasks. In some embodiments, the 'validator' field may be configured as a static value.

[008] In accordance with yet another innovative aspect of the present invention, multi-threaded asynchronous logging is tracked or traced by the rule-based system configured with the specific plug-in based rule system (implemented in the monitoring engine), by which each step in the rule is accorded a unique ID, and a combination of the three IDs is generated to formulate a unique ID that is specific to a particular rule run (otherwise referred to as a process, operation, or step). Other runs or steps using the same code may yield other IDs, making the multi-threaded process easy to follow. Although the threads all log to the same file in the non-linear fashion, the flow of events may be logically followed with ease by looking at the combined ID in the log file.

[009] In some embodiments, the monitoring and tracing system of the present invention may include a processor and memory storing instructions executable by the processor 1) to execute an application program interface to operate a monitoring engine configured to monitor a particular functionality within a software-based environment or cloud infrastructure based on a rule-based configuration, and configured 2) to define three separate categories of plug-ins to use the plug-ins in rules, the plug-ins including a checker plug-in to return a value output, a validator plug-in to receive the value output and compare the value output to a threshold standard and return a status of a validation operation performed by the validator

plug-in, and one or more action plug-ins configured to perform tasks after the validation operation.,

[0010] Other innovative aspects include corresponding systems, methods, apparatus, and computer program products.

5 [0011] It should be understood that language used in the present disclosure has been principally selected for readability and instructional purposes and not to limit the scope of the subject matter disclosed herein.

BRIEF DESCRIPTION OF THE DRAWINGS

10 [0012] The present disclosure is illustrated by way of example, and not by way of limitation, in the figures of the accompanying drawings in which like reference numerals are used to refer to similar elements.

[0013] Figure 1 is a high-level block diagram illustrating an application program interface (API) to a cloud infrastructure in accordance with the present invention including a rule-
15 based monitoring engine and functionality for tracing multi-threaded asynchronous logging to a single file.

[0014] Figure 2 is a block diagram illustrating an example plug-ins module within the monitoring engine of the Application Program Interface (API).

[0015] Figure 3 is a block diagram illustrating an example rule-queuing module within the
20 monitoring engine for tracing multi-threaded asynchronous logging.

[0016] Figure 4 is a block diagram illustrating an example Application Program Interface (API).

[0017] Figure 5 is a flow chart illustrating an example monitoring protocol for a rule run, process, or operation performed by the monitoring engine.

25 [0018] Figure 6A is a flow chart illustrating an example process for identifying a rule and determining its status.

[0019] Figure 6B is a flow chart illustrating an example method for generating individual IDs for each step in a rule and a combined ID for the rule to facilitate tracing of multi-thread operations relating to executing the rule.

30 [0020] Figure 7 is a flow chart illustrating an example validation protocol that receives a status report, receives validation and returns the validation.

[0021] Figure 8 is a flow chart illustrating an example protocol for performing an action based on a reported status and validation.

[0022] Figure 9A is a first part of the flow chart illustrating an example rule run or process for executing a rule from checking to executing the rule.

[0023] Figure 9B is a second part of the flow chart illustrating the example process rule run or process.

5

DETAILED DESCRIPTION

[0024] System and method architecture of a monitoring engine and functionalities for tracing of multi-threaded logging, e.g., implemented in an application program interface for interaction with a cloud infrastructure are described below. Although the systems and methods of the present disclosure are described in the context of particular system architectures that are illustrated, it should be understood that the systems and methods may be applied to other architectures and organizations of hardware with similar properties.

10

[0025] It should be understood that the present disclosure addresses the problem of existing monitoring agent designs that do not permit those with little programming knowledge to either quickly or easily come up with new monitoring functionality to monitor operations in a software-based environment or cloud architecture as the actual functionality is too tied in with the actual monitoring engine code base. The present disclosure also solves the complexities of multi-threaded processing, by assigning, a unique identification to each thread or process and providing a logging interface in a specific format, so that the combination of all the identification numbers for the various threads or processes formulate a big, unique identification, that is easily traceable through the entire log file.

15

20

[0026] For purposes of this disclosure, it should be understood that a plug-in refers to a software component that adds a specific feature or element to an existing computer program or structure. A plug-in may also be referred to as a plugin, an add-on, extension, or supplement. Plug-in applications typically are programs that may be easily installed and used as part an existing program. In addition, an application program interface (API) is a set of routines, protocols, and tools for building software applications for use with existing software/hardware structure or architecture (e.g. a cloud infrastructure). An application program interface (API) specifies how software components should interact with the existing architecture. A cloud infrastructure includes a network of remote servers hosted on the internet or a cloud product that is created to operate on a customer's premises to store, manage, and process data, to provide large-scale virtual resources. Cloud architecture typically provides cloud services that have several layers, namely, client, applications that are software-as-a-service (SaaS), platforms or platforms-as-a-service (PaaS), cloud infrastructure

25

30

or cloud infrastructure-as-a-service (IaaS), and servers. Each of these layers or categories in a cloud serves a different purpose and offers different products for businesses and individuals around the world.

[0027] Referring now to Figure 1, it should be recognized that Figure 1 is a block diagram illustrating an example system 100 with an application program interface (API) configured to interact with a cloud infrastructure. A monitoring-tracing engine 110 is illustrated that is rule-based and has a capability that traces multi-threaded logging into a file in the cloud infrastructure. The example system 100 may include a network interface (I/F) module 102 configured to connect with a cloud infrastructure via a network 105, a processor 104, a storage system 106 with a database 112, a memory 108, and an Application Program Interface (API) 114 with a monitoring-tracing engine 110 that has a rule module 152, a plug-ins module 154, and a rule-queuing module 156. The components of the system 100 are communicatively coupled to a bus or software communication mechanism 120 for communication with each other.

[0028] The example system 100 has an application program interface (API) 114 configured to access a cloud architecture (via the network 105) and to perform specific functions or operations (also referred to here as operations or runs) within the cloud architecture. The application program interface 114 has a monitoring-tracing engine 110, which is configured with plug-ins to define functionality including rule-based configuration syntax adapted to intuitively and automatically deploy or launch the plug-ins as necessary in a cloud infrastructure. The monitoring-tracing engine 110 is also configured to trace multi-threaded asynchronous logging to a single file within a cloud infrastructure, by assigning a unique identification to each thread or process initiated for a rule (e.g. validation) and facilitating each logging according to a specific format to formulate a combination of all the different thread identifications into a master or combined unique identification that is easily traceable through an entire log file. (e.g., to record all access and operation data). The monitoring-tracing engine 110 has a rule module 152, by which a particular rule is either identified or defined. The plug-ins module 154 is adapted to deploy a rule-based plug-in into the cloud infrastructure to perform monitoring tasks. The rule-queuing module 156 tracks the rule employed and the process or thread that relates to the rule (by asynchronous logging into a file). In some instances, each server within the cloud infrastructure may run its own monitoring-tracing engine 110 to operate local monitoring functions within each server. Each such monitoring-tracing engine 110 exposes an API of each server, allowing it to

remotely trigger, request, or otherwise interact with the server's local monitoring-tracing engine 110.

[0029] In some implementations, software communication mechanism 120 may be an object bus (e.g., Common Object Request Broker Architecture (CORBA)), direct socket

5 communication (e.g., Transmission Control Protocol (TCP)/Internet Protocol (IP) sockets) among software modules, remote procedure calls, User Datagram Protocol (UDP) broadcasts and receipts, HTTP connections, function or procedure calls, etc. Further, any or all of the

communication could be secure (Secure Shell (SSH), Hypertext Transfer Protocol Secure (HTTPS), etc.). The software communication mechanism 120 can be implemented on any

10 underlying hardware (for example, a network, the Internet, a bus, a combination thereof, etc.).

[0030] The network interface (I/F) module 102 is configured to connect system 100 to a network and/or other system, for example a network 105 coupled to a cloud infrastructure.

For example, network interface module 102 may enable communication through one or more of the Internet, cable networks, and other wired networks. The network interface module 102

15 links the processor 104 to a network that may in turn be coupled to other processing systems.

The network interface module 102 also provides other conventional connections to the network 120 for distribution and/or retrieval of files and/or media objects using standard network protocols such as TCP/IP, HTTP, HTTPS and Simple Mail Transfer Protocol

(SMTP) as will be understood. In some implementations, the network interface module 102
20 includes a transceiver for sending and receiving signals using Wi-Fi, Bluetooth®, or cellular communications for wireless communication.

[0031] The network interface module 102 may allow client and/or server devices to store data to and retrieve data from the storage system 106 (for example, across the network 120).

For example, the network interface module 102 may allow one or more components of the
25 system 100 to act as network accessed storage or as an array of storage devices.

[0032] The network 105 may be one of a conventional type, wired or wireless, and may have numerous different configurations including a star configuration, token ring

configuration, or other configurations. Furthermore, the network may include a local area

network (LAN), a wide area network (WAN) (e.g., the Internet), and/or other interconnected

30 data paths across which multiple devices may communicate. In some implementations, the network may be a peer-to-peer network. The network may also be coupled with or include portions of a telecommunications network for sending data using a variety of different

communication protocols. In some implementations, the network may include Bluetooth (or Bluetooth low energy) communication networks or a cellular communications network for

sending and receiving data including via short messaging service (SMS), multimedia messaging service (MMS), hypertext transfer protocol (HTTP), direct data connection, Wireless Application Protocol (WAP), email, etc.

[0033] The processor 104 may include an arithmetic logic unit, a microprocessor, a general-purpose controller or some other processor array configured to perform computations and provide electronic display signals to a display device in accordance with the present invention. In some implementations, the processor 104 is a hardware processor having one or more processing cores. The processor 104 is coupled to the bus 120 for communication with the other components of system 100. The processor 104 processes data signals and may include various computing architectures including a complex instruction set computer (CISC) architecture, a reduced instruction set computer (RISC) architecture, or an architecture implementing a combination of instruction sets. Although only a single processor is shown in the example of Figure 1, multiple processors and/or processing cores may be included. It should be understood that other processor configurations are also possible.

[0034] The storage system 106 is a data storage system that may store information on any type of attached array of writable storage media, such as magnetic disk or tape, optical disk (e.g., CD-ROM or DVD), flash memory, solid-state drive (SSD), electronic random access memory (RAM), micro-electro mechanical and/or any other similar media adapted to store information, including data and parity information. However, as illustratively described herein, the information is stored on disks of hard disk drives (HDDs). The storage device(s) of the storage system 106 may be included in one or more enclosures housed in the same location or physically distributed to disparate locations across a storage area network. The storage system comprises a database 112 and storage logic of the storage system 106 can provide computing functionalities, services, and/or resources to send, receive, read, write, and transform data from other entities of the system 100. In some embodiments, the storage logic can be a computing device configured to make a portion or all of the storage space available on the storage system 106. The storage logic is coupled to a switch for communication and cooperation with the storage system 106 and the database 112 of the system 100. In other embodiments, the storage logic transmits data between the storage system 106 and the memory 108 via the switch 118. It should be recognized that multiple storage logic units may be utilized, either in a distributed architecture or otherwise. For the purpose of this application, the system configuration and operations performed by the system are described in the context of a single storage logic and storage system 106.

[0035] The switch used in the storage system 106 may be a conventional type, wired, and may have numerous different configurations. Furthermore, the switch 118 may include an Ethernet, InfiniBand, Peripheral Component Interconnect Express (PCI-Express) switch, and/or other interconnected data path switches, across which multiple devices (e.g., storage devices 122) may communicate. Although the example of Figure 1 illustrates a single storage system 106, in practice one or a plurality of storage systems 106 can connect the entities.

[0036] The storage system 106 may include a non-transitory computer-usable (e.g., readable, writeable, etc.) medium, which may be any non-transitory apparatus or device that can contain, store, communicate, propagate or transport instructions, data, computer programs, software, code routines, etc., for processing by or in connection with a processor. In some embodiments, the storage system 106 communicates and cooperates with the storage logic by switch. The storage system 106 may be flash memory or may include a volatile memory such as a dynamic random access memory (DRAM) device, a static random access memory (SRAM) device, or some other memory devices.

[0037] The memory 108 stores instructions and/or data that may be executed by the processor 104. The memory 108 is coupled to the bus 120 for communication with the other components of the system 100. The instructions and/or data stored in the memory 108 may include code for performing any and/or all of the techniques described herein. The memory 108 may be, for example, volatile memory such as a dynamic random access memory (DRAM) device, a static random access memory (SRAM) device, or some other volatile memory device. In some implementations, the memory 108 also includes a non-volatile memory or similar permanent storage device and media, for example, a hard disk drive, a floppy disk drive, a compact disc read only memory (CD-ROM) device, a digital versatile disc read only memory (DVD-ROM) device, a digital versatile disc random access memory (DVD-RAM) device, a digital versatile disc rewritable (DVD-RW) device, a flash memory device, or some other non-volatile storage device. In some implementations, the memory 108 may store the data for the Application Program Interface (API) 114 and executable code to perform its functionalities.

[0038] Referring now to Figures 1 and 2, the monitoring-tracing engine 110 has a plug-ins module 154 for monitoring related scripts and a configuration file separation software for packaging separate files. The plug-ins module 154 has a single purpose and targets a single function, therefore, leaves a small footprint when it deploys a plug-in. The plug-ins module 154 defines plug-ins that are easy to write and support. The plug-ins module 154 has three

components – checker, validator, and action. The checker plug-in or module 202 checks a function within the cloud infrastructure and returns its status/result. An example routine performed by the checker module 202 is illustrated below.

Example: check: 'is A running?'

5 Returns: True

Check: 'how much disk space is in use on sda (in %)?'

Returns: 80

[0039] The plug-ins module 154 includes a validator plug-in or module 204 configured to receive a result generated by the checker module 202 and validate the result and return the
10 status of the validation.

Example:

receives '80' (from the disk checker)

validates it's bigger than '75' (arbitrary value)

returns True

15 [0040] The plug-ins module 154 executes a single line per rule, and with multiple rules uses the same checker function (by the checker module 202) and has a fixed format. An example is illustrated below:

Format:

<checker> <op> <validator>:<action 0>, <action 1>, ...

20 Example:

ch_disk.=va_disk:remove_home_oostj

ch_disk>=80:remove_home_oostj,raise_event

In a multi-threaded application, every plug-in runs in its own thread and there is no blocking of the processor 104 or of the IO or Network I/F module 102. There is no long-running plug-
25 in, and there is action plug-in collision detection. Furthermore, there is a single run of all plug-ins (persistent) and a multiple queuing system that can survive sudden read-only file systems.

[0041] In some scenarios, a REST-like API interface may control the monitoring-tracing engine 110 remotely and launch individual plug-ins and/or individual or multiple rules by the
30 plugins module 154, and retrieves results from the operations that are performed by the individual plug-ins.

[0042] Referring now to Figure 2, it should be recognized that the monitoring-tracing engine 110 does not provide any functionality by itself, but instead uses the plug-ins module 154, which is configured to define the three plug-in categories with three separate plug-ins. The

first, the checker module 202, the second, the validator module 204, and the third, the action module 206. The checker module or plug-in 202 has the task to return a certain value. The validator plug-in or module has the task to accept a value coming from the checker module 202, compare the value to something else and to return the status of this validation operation.

5 The third, action module 206 performs an action or task.

[0043] A rule is configured to follow a simple syntax: 'checker' 'operator' 'validator': 'action0', 'action1'... This simply means that when this rule is deployed, the monitoring-tracing engine 110 launch an instance of 'checker' and compares the returned value to 'validator', if the result of the 'checker' 'operator' 'validator' is true, 'action0', 'action1' will be
 10 launched, if the result of the first part is false; the monitoring-tracing engine 110 does not launch the actions. The 'validator' field may be a static value as well. An Example of a rule is indicated below:

disk_usage_in_percent>=90: send_event, rotate_logs

disk_usage_in_percent>=90: send_event, remove_logs

15 Defining the multiple rules, even by using the same checker module 202 is possible, further extending the flexibility of the system. The monitoring-tracing engine 110 (Figure 1) can either periodically run these rules to oversee operations or run them all once to act as a health-check or diagnostic run.

[0044] Combining these concepts, it becomes easy to define a new 'situation' to monitor and
 20 act on. The monitoring-tracing engine 110 has an unexpected favorable result as it serves as a diagnostic/health-check system, but the monitoring-tracing engine 110 may also serve to monitor action-based events, reducing development time.

[0045] Referring now to Figure 3, the present system 100 also addresses multi-threading and asynchronous logging functionalities within a cloud infrastructure. The
 25 monitoring-tracing engine 110 (Figure 1) represents a rule-based system that works with specific plug-ins (that define rules), as illustrated by the rule-queuing module 156, which comprises a thread-queuer/assign ID module 304 and a combine-ID module. The thread-queuer/assign ID module 304 identifies each step (or thread) in a particular rule and executes functionality for assigning a unique identification (ID) to each thread or step. The combine-
 30 ID module 306 generates a unique-combined identification that combines the three identifications (IDs) to create a unique identification for that rule operation. Other operations using the same code may yield other identifications (IDs), making the multi-threaded process (involving various steps or threads) easy to follow as illustrated below.

Example:

An example process identifies and illustrates three threads: 'A', 'B' and 'C', 'A' and 'b' are the same code but are separately identified as they are assigned a different identification (ID).

Each of the three threads spawn two other threads 'AY', 'AZ', 'BY', 'BZ', 'CY', 'CZ'.

For these threads, the main process log file may appear as below:

5 thread A: ID 'AAAAAA' assigned
thread B: ID 'BBBBBB' assigned
thread C: ID 'CCCCCC' assigned

[BBBBBB] spawning new thread with id 'BZBZBZ'

10 [AAAAAA] spawning new thread with id 'AYAYAY'

[CCCCCC] spawning new thread with id 'CYCYCY'

[AAAAAA] spawning new thread with id 'AZAZAZ'

[BBBBBB] spawning new thread with id 'BYBYBY'

[CCCCCC] spawning new thread with id 'CZCZCZ'

15

[BBBBBB-BYBYBY] doing something 0, assigning id 'BYCBYC'

[CCCCCC-CYCYCY] doing something 0, assigning id 'CYACYA'

[BBBBBB-BYBYBY] doing something 2, assigning id 'BYABYA'

[CCCCCC-CYCYCY] doing something 1, assigning id 'CYBCYB'

20 [AAAAAA-AYAYAY] doing something 0, assigning id 'AYAAYA'

[AAAAAA-AYAYAY] doing something 1, assigning id 'AYBAYB'

[BBBBBB-BYBYBY] doing something 1, assigning id 'BYBBYB'

[CCCCCC-CYCYCY] doing something 2, assigning id 'CYCCYC'

25 [AAAAAA-AYAYAY-AYBAYB] result of thread AAAAAA

spawning AYAYAY doing AYBAYB

[CCCCCC-CYCYCY-CYCCYC] result of thread CCCCCC

spawning CYCYCY doing CYCCYC

etc...

30

[0046] Even though the threads all log into the same file in a non-linear fashion (asynchronously), by this mechanism of identifying, it is logical to follow the flow of events with ease, by referring to the combined ID in the log file. By using this technique, it is relatively easy to follow specific log messages in a log file and to display them in a structured

way in a user interface or to use it to simplify troubleshooting. This application may be used in any multi-threaded application with ease.

[0047] Referring now to Figure 4, the Application Program Interface (API) 114 has a pause module 402, a fetch module 404, a launch module 406, a health-check module 408, and a listing module 410, each connected to the other by a bus 420. Each of these modules represents the hardware/software required for executing functions (by rules) in the cloud infrastructure. The pause module 402 may execute pause functions on the instructions executed with respect to a particular rule or launching of a plug-in. The fetch module 404 may fetch results from earlier rule runs. The launch 406 may execute a particular rule or deploy a plug-in. The health-check module 408 is configured to execute a diagnostics run. The listing module 410 may be configured to list the particular rules or plug-ins, each executable by the processor 104 (Figure 1).

[0048] Referring now to Figure 5, the operations of some embodiments of the plug-ins module 154 are described. In one example, as illustrated by process 500, configuring a rule follows a simple syntax: 'checker' 'operator' 'validator': 'action0', 'action 1'. When this rule is deployed, the monitoring-tracing engine 110 launches an instance of the 'checker' and compares the returned value to the 'validator', and if the result of 'checker' 'operator' validator' is true, 'action0', 'action 1' is launched, and if the result of the first part is false; the plug-ins module does not launch the action. The process 500, at step 502, runs a checker thread to return a value. The process 500 proceeds to the next decision block 504, at which point, a determination is made if the value is validated. If the answer is affirmative, the process 500 proceeds to the next step 506, at which point, an action is performed. From there, the process 500 proceeds to the next block 508, for queueing the action. If the determination at the decision block 504 is negative, the process 500 proceeds to block 508, for queueing the action. From block 508, the process 500 may return to block 502, to run the next checker thread (as illustrated by the broken lines). In some instances, the 'validator' field may be a static value as well.

An example of a rule is as follows:

disk_usage_in_percent >=90: send_event_rotate_logs (this translates to 'if disk usage is greater than or equal to 90 percent: send an event and rotate the logs')

An example of multiple rules using the same checker is as follows:

disk_usage_in_percent >=70: send_event_rotate_logs
disk_usage_in_percent >=90: send_event_remove_logs

Defining multiple rules, even using the same checker, is possible, and further extends the flexibility of the system.

[0049] The monitoring-tracing engine 110 is configured to either periodically run these rules to act as a monitoring agent or may run them all at once to act as a health check or
5 diagnostics module. By combining these concepts, it becomes very easy to define new 'situations' to monitor and act on. An unexpected yet favorable result of this monitoring-tracing engine is that it serves as a watchdog/health check system, yet may be used to implement action-based events, reducing development time. This system serves as a pluggable platform capable of providing the monitoring functionalities necessary.

10 Referring now to Figure 6A, the run-checker thread 502a that returns the value includes the step of identifying the rule 602a, determining the status of the rule 604a, and returning the status 606a. Referring now to Figure 6B, the run-checker 502b includes the step 602b, including one or more operations for identifying the rule and the steps in the rule. The process 502b proceeds to the next step 604b, including one or more operations for
15 assigning a unique identification (ID) to each step. The process 502b proceeds to the next step 606b, including one or more operations for combining all identifications (IDs) and generating a combined identification (ID) specific for that rule.

Referring now to Figure 7, the operations for validating the value are described in further detail. The operations include step 702, including one or more operations
20 for receiving the status. The process 504 proceeds to the step 704 of receiving validation, including one or more steps for receiving validation. From 704, the process 504 proceeds to the next step 706, including one or more operations, configured for returning the validation.

Referring now to Figure 8, the steps of performing actions by plug-ins, illustrated by block 506 are further described. The process 506, proceeds to step 802,
25 including one or more operations for receiving the status of a particular operation. From 802, the process 506 proceeds to step 804, including one or more operations for receiving validation (e.g., that the operation is performing without problem). From there, the process 506 may include one or more operations for performing an action (e.g., to correct a problem detected) based on the status and validation.

30 Referring now to Figure 9A, the process 800 of the rule-based system is described. The process begins by starting the rule run, at block 902. The process 800 includes the decision block 904, by which a determination is made to determine if the checker plug-in (executed by the checker module 202 in Figure 2) is already running and performing its intended purpose. If the answer is affirmative the process proceeds to the next block of

operations 906, by which the checker plug-in waits for the next checker iteration. If the answer output from decision block 904 is negative, the process 800 proceeds to the block 908, by which the checker plug-in is run and output is collected. The process 800 proceeds to a decision block 910 including one or more operations for determining if the checker run is successful. If the answer is affirmative, the process 800 proceeds to connector point A (connecting to continuing operations described in Figure 9B). If the answer at decision block 910 is negative, the process 800 proceeds to the next block 912, at which point, the plug-in run is placed in queue for the queue for the monitoring-tracing engine 110.

Referring now to Figure 9B, proceeding from connector point A, the process 800 proceeds to decision block 920, including one or more operations for determining if the validator is a plug-in. If the answer is affirmative, the process 800 proceeds to the next block 922, including one or more operations for launching the validator plug-in and collecting output data. From block 922, the process 800 proceeds to a next decision block 924, including one or more operations for determining if the validator plug-in run was successful. If the answer determined is negative, the process 800 proceeds to a connector C (connecting to block 912 in Figure 9A). If the answer at decision block 924 is affirmative, the process 800 proceeds to the next decision block 926, including one or more operations for comparing the checker output (by the checker plug-in) to the validator output (by the validator plug-in) to determine if the result from the comparison meets a desired outcome. If the answer generated at this decision block 926 is negative, the process proceeds to connector B (connecting to block 914 in Figure 9A). If the answer generated at decision block 926 is affirmative, it is determined that no action is required (e.g. to correct a problem detected by the monitoring engine 110).

[0050] In the depicted embodiments, the network 105 may connect to a cloud infrastructure (with cluster nodes) including a cloud computing system (not shown), which may include cloud storage including storage logic, a switch, and one or more storage devices. However, the present disclosure is not limited to this configuration, and a variety of different system environments and configurations can be employed and are within the scope of the present disclosure. Other embodiments may include additional or fewer components.

[0051] The systems and methods described above provide a highly reliable system for implementing cross device redundancy schemes. In the above description, for purposes of explanation, numerous specific details were set forth. It will be apparent, however, that the disclosed technologies can be practiced without any given subset of these specific details. In

other instances, structures and devices are shown in block diagram form. For example, the disclosed technologies are described in some implementations above with reference to user interfaces and particular hardware. Moreover, the technologies disclosed above primarily in the context of on line services; however, the disclosed technologies apply to other data sources and other data types (e.g., collections of other resources for example images, audio, web pages).

[0052] Reference in the specification to “one implementation” or “an implementation” means that a particular feature, structure, or characteristic described in connection with the implementation is included in at least one implementation of the disclosed technologies. The appearances of the phrase “in one implementation” in various places in the specification are not necessarily all referring to the same implementation.

[0053] Some portions of the detailed descriptions above were presented in terms of processes and symbolic representations of operations on data bits within a computer memory. A process can generally be considered a self-consistent sequence of steps leading to a result. The steps may involve physical manipulations of physical quantities. These quantities take the form of electrical or magnetic signals capable of being stored, transferred, combined, compared, and otherwise manipulated. These signals may be referred to as being in the form of bits, values, elements, symbols, characters, terms, numbers, or the like.

[0054] These and similar terms can be associated with the appropriate physical quantities and can be considered labels applied to these quantities. Unless specifically stated otherwise as apparent from the prior discussion, it is appreciated that throughout the description, discussions utilizing terms for example “processing” or “computing” or “calculating” or “determining” or “displaying” or the like, may refer to the action and processes of a computer system, or similar electronic computing device, that manipulates and transforms data represented as physical (electronic) quantities within the computer system’s registers and memories into other data similarly represented as physical quantities within the computer system memories or registers or other such information storage, transmission or display devices.

[0055] The disclosed technologies may also relate to an apparatus for performing the operations herein. This apparatus may be specially constructed for the required purposes, or it may include a general-purpose computer selectively activated or reconfigured by a computer program stored in the computer. Such a computer program may be stored in a computer readable storage medium, for example, but is not limited to, any type of disk including floppy disks, optical disks, CD-ROMs, and magnetic disks, read-only memories

(ROMs), random access memories (RAMs), Erasable Programmable Read-Only Memories (EPROMs), Electrically Erasable Programmable Read-Only Memories (EEPROMs), magnetic or optical cards, flash memories including Universal Serial Bus (USB) keys with non-volatile memory or any type of media suitable for storing electronic instructions, each
5 coupled to a computer system bus.

[0056] The disclosed technologies can take the form of an entirely hardware implementation, an entirely software implementation or an implementation containing both hardware and software elements. In some implementations, the technology is implemented in software, which includes but is not limited to firmware, resident software, microcode, state
10 machines, gate arrays, etc.

[0057] Furthermore, the disclosed technologies can take the form of a computer program product accessible from a non-transitory computer-usable or computer-readable medium providing program code for use by or in connection with a computer or any instruction execution system. For the purposes of this description, a computer-usable or computer-
15 readable medium can be any apparatus that can contain, store, communicate, propagate, or transport the program for use by or in connection with the instruction execution system, apparatus, or device.

[0058] A computing system or data processing system suitable for storing and/or executing program code will include at least one processor (e.g., a hardware processor) coupled directly
20 or indirectly to memory elements through a system bus. The memory elements can include local memory employed during actual execution of the program code, bulk storage, and cache memories which provide temporary storage of at least some program code in order to reduce the number of times code must be retrieved from bulk storage during execution.

[0059] When the disclosed technologies are embodied in the general context of computer-executable instructions, the embodiments may comprise program modules, executed by one
25 or more systems, computers, or other devices. Generally, program modules include routines, programs, objects, components, data structures, and the like, that perform particular tasks or implement particular abstract data types. The functionality of the program modules may be combined or distributed as desired in various embodiments.

[0060] Input/output or I/O devices (including but not limited to keyboards, displays, pointing devices, etc.) can be coupled to the system either directly or through intervening I/O
30 controllers.

[0061] Network adapters may also be coupled to the system to enable the data processing system to become coupled to other data processing systems or remote printers or storage

devices through intervening private or public networks. Modems, cable modems and Ethernet cards are just a few of the currently available types of network adapters.

[0062] Finally, the processes and displays presented herein may not be inherently related to any particular computer or other apparatus. Various general-purpose systems may be used with programs in accordance with the teachings herein to create special purpose apparatus, or it may prove convenient to construct a more specialized apparatus to perform the required method steps. The required structure for a variety of these systems will appear from the description below. In addition, the disclosed technologies were not described with reference to any particular programming language. It will be appreciated that a variety of programming languages may be used to implement the teachings of the technologies as described herein.

[0063] The foregoing description of the implementations of the present techniques and technologies has been presented for the purposes of illustration and description. It is not intended to be exhaustive or to limit the present techniques and technologies to the precise form disclosed. Many modifications and variations are possible in light of the above teaching. It is intended that the scope of the present techniques and technologies be limited not by this detailed description. The present techniques and technologies may be implemented in other specific forms without departing from the spirit or essential characteristics thereof. Likewise, the particular naming and division of the modules, routines, features, attributes, methodologies and other aspects are not mandatory or significant, and the mechanisms that implement the present techniques and technologies or its features may have different names, divisions and/or formats. Furthermore, the modules, routines, features, attributes, methodologies and other aspects of the present technology can be implemented as software, hardware, firmware or any combination of the three. Also, wherever a component, an example of which is a module, is implemented as software, the component can be implemented as a standalone program, as part of a larger program, as a plurality of separate programs, as a statically or dynamically linked library, as a kernel loadable module, as a device driver, and/or in every and any other way known now or in the future in computer programming. Additionally, the present techniques and technologies are in no way limited to implementation in any specific programming language, or for any specific operating system or environment. Accordingly, the disclosure of the present techniques and technologies is intended to be illustrative, but not limiting.

WHAT IS CLAIMED IS:

1. A system comprising:
a processor; and
memory storing instructions that, when executed, cause the processor to:
execute an application program interface to operate a monitoring engine
configured to monitor a particular functionality within a software-based infrastructure based on a rule-based configuration;
define, by the monitoring engine, three separate categories of plug-ins to
formulate the plug-ins in rules, the plug-ins including a checker plug-in
to return a value output after a checking operation, a validator plug-in
to receive the value output and compare the value output to a threshold
standard and return a status of a validation operation performed by the
validator plug-in, and an action plug-in configured to perform a task
after the validation operation indicates action is required; and
selectively launch the checker plug-in, the validator plug-in, and the action
plug-in into the software-based infrastructure.
.
2. The system of claim 1, wherein the monitoring engine comprises a plug-ins
module adapted to use at least the same checker plug-in once for defining multiple rules.
3. The system of claim 1, wherein the plug-ins module comprises a checker
module for adding the checker plug-in, a validator module for adding the validation plug-in,
and an action module for adding the action plug-in.
4. The system of claim 1, further comprising:
a first module adapted to identify each step in a particular rule or loggin thread and
assign an identification to each step.
5. The system of claim 4, further comprising:
a second module adapted to generate a unique-combined identification for all the
steps in the rule or logging identified.

6. The system of claim 1, wherein the status of the validation output determines if an action is required for a particular functionality identified in the software-based infrastructure.

7. The system of claim 6, wherein if the action is required and performed on a particular functionality identified in the software-based infrastructure, a further determination to evaluate if the action was successful is made.

8. A method comprising:
executing an application program interface to operate a monitoring engine configured to monitor a particular functionality within a software-based infrastructure based on a rule-based configuration;
defining three separate categories of plug-ins to formulate the plug-ins as rules, including a checker plug-in to return a value output from a particular situation, a validator plug-in to receive the value output and compare the value output to a threshold standard and return a status of a validation operation performed by the validator plug-in, and an action plug-in configured to perform a task after the validation operation; and
generating and launching the checker plug-in, the validator plug-in, and the action plug-in in sequence into the infrastructure.

9. The method of claim 8, wherein the monitoring engine comprises a plug-ins module adapted to launch the plurality of plug-ins in sequence upon determining a need for each plug-in.

10. The method of claim 8, wherein the plug-ins module comprises a checker module configured to formulate a checker plug-in, a validator module to formulate a validation plug-in, and an action module to formulate an action plug-in.

11. The method of claim 8, further comprising:
identifying each step in a particular rule or each login thread associated with each step and assigning an identification.

12. The method of claim 8, further comprising:
generating a unique-combined identification and assigning the unique-combined identification for all the steps in the rule or logging from multiple threads associated with the rule.
13. The method of claim 8, wherein the status of the validation output determines if an action is required to a particular functionality in the software-based infrastructure.
14. The method of claim 13, wherein if the action is required and performed, a further determination is made to evaluate if the action was successful.
15. A computer product comprising:
means for executing an application program interface to operate a monitoring engine configured to monitor a particular functionality within a software-based infrastructure based on a rule-based configuration;
means for defining three separate categories of plug-ins to use the plug-ins as rules, including a checker plug-in to return a value output, a validator plug-in to receive the value output and compare the value output to a threshold standard and return a status of a validation operation performed by the validator plug-in, and one or more action plug-in configured to perform a task after the validation operation; and
means for adding the checker plug-in, the validator plug-in, and the action plug-ins into the software-based infrastructure.
16. The computer product of claim 15, wherein the monitoring engine comprises a plug-ins module adapted to deploy the plurality of plug-ins in sequence.
17. The computer product of claim 15, further comprising:
means for defining tasks for plug-ins including a checker module configurable to define a first task for the checker plug-in, a validator module configurable to define a second task for the validator plug-in, and an action module configurable to define a third task for the action plug-in.
18. The computer product of claim 15, further comprising:

means for identifying each step in a particular rule or logging thread associated with the step and assigning an identification.

19. The computer product of claim 15, further comprising:

means for generating a unique-combined identification and assigning the unique-combined identification for all the steps in a rule.

20. The computer product of claim 15, further comprising:

means for determining, by the status of the validation output if an action is required in the software-based infrastructure.

21. The computer product of claim 15, further comprising:

means for further determining if the action was successful if the action is required and performed.




Figure 1




Figure 2




Figure 3




Figure 4




Figure 5

502a




Figure 6A

502b
↓




Figure 6B

504 →




Figure 7

506 →




Figure 8




Figure 9A




Figure 9B

INTERNATIONAL SEARCH REPORT

International application No

PCT/US2018/020815

A. CLASSIFICATION OF SUBJECT MATTER

INV. G06F11/34 G06F11/30
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

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, WPI Data

C. DOCUMENTS CONSIDERED TO BE RELEVANT

Category*	Citation of document, with indication, where appropriate, of the relevant passages	Relevant to claim No.
X	US 2007/168919 A1 (HENSELER DAVID A [US] ET AL) 19 July 2007 (2007-07-19) abstract; figures 15, 16, 21 paragraph [0151] - paragraph [0162] paragraph [0199] - paragraph [0223] the whole document	1-21
A	US 2016/036985 A1 (KOREN NIR [IL] ET AL) 4 February 2016 (2016-02-04) abstract; figures 1-6 paragraph [0001] - paragraph [0017] the whole document ----- -/-	1-21



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 August 2018

Date of mailing of the international search report

21/08/2018

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

Bozas, Ioannis

INTERNATIONAL SEARCH REPORT

International application No
PCT/US2018/020815

C(Continuation). DOCUMENTS CONSIDERED TO BE RELEVANT		
Category*	Citation of document, with indication, where appropriate, of the relevant passages	Relevant to claim No.
A	WO 2016/175851 A1 (APPDYNAMICS INC [US]) 3 November 2016 (2016-11-03) abstract; figures 1-4 paragraph [0002] - paragraph [0005] paragraph [0033] - paragraph [0039] the whole document -----	4,5,11, 12,18,19
A	US 2016/124783 A1 (PUVVADA SURAJ [US] ET AL) 5 May 2016 (2016-05-05) abstract; figures 1-4 paragraph [0002] - paragraph [0003] paragraph [0029] - paragraph [0034] the whole document -----	4,5,11, 12,18,19
A	US 2015/339210 A1 (KOPP MICHAEL [AT] ET AL) 26 November 2015 (2015-11-26) abstract; figure 4 paragraph [0083] - paragraph [0085] the whole document -----	4,5,11, 12,18,19

INTERNATIONAL SEARCH REPORT

Information on patent family members

International application No

PCT/US2018/020815

Patent document cited in search report	Publication date	Patent family member(s)	Publication date
US 2007168919 A1	19-07-2007	EP 1955151 A1	13-08-2008
		US 2007168919 A1	19-07-2007
		WO 2007064799 A1	07-06-2007

US 2016036985 A1	04-02-2016	NONE	

WO 2016175851 A1	03-11-2016	CA 2982272 A1	03-11-2016
		EP 3289455 A1	07-03-2018
		US 2016323362 A1	03-11-2016
		US 2018159724 A1	07-06-2018
		WO 2016175851 A1	03-11-2016

US 2016124783 A1	05-05-2016	US 2016124783 A1	05-05-2016
		US 2017147483 A1	25-05-2017

US 2015339210 A1	26-11-2015	NONE	
