

ADAPTIVE SERVICE TIMEOUTS

[0001] The following application claims priority to, and the benefit of, U.S. Patent Application entitled "Adaptive Service Timeouts" filed on September 18, 2012 having application number 13/622,094, which is incorporated herein by reference in its entirety.

BACKGROUND

[0002] Timeouts can prevent excessive use of system resources in processing high-latency service calls. Specifically, when processes time out, they no longer consume system resources.

BRIEF DESCRIPTION OF THE DRAWINGS

[0003] Many aspects of the present disclosure can be better understood with reference to the following drawings. The components in the drawings are not necessarily to scale, with emphasis instead being placed upon clearly illustrating the principles of the disclosure. Moreover, in the drawings, like reference numerals designate corresponding parts throughout the several views.

[0004] FIG. 1 is a drawing of a networked environment according to various embodiments of the present disclosure.

[0005] FIG. 2 is a flowchart illustrating one example of functionality implemented as portions of the timeout management application executed in a computing environment in the networked environment of FIG. 1 according to various embodiments of the present disclosure.

[0006] FIG. 3 is a flowchart illustrating one example of service dependency functionality implemented as portions of the timeout management application executed in a computing environment in the networked environment of FIG. 1 according to various embodiments of the present disclosure.

[0007] FIG. 4 is a schematic block diagram that provides one example illustration of a computing environment employed in the networked environment of FIG. 1 according to various embodiments of the present disclosure.

DETAILED DESCRIPTION

[0008] Service-oriented architectures enforce timeouts to prevent excessive use of system resources when making calls to a high-latency service. This prevents a high-latency service from

monopolizing or abusing system resources to resolve its calls at the expense of other services competing for the same system resources. Timeouts are often preset values, defined with respect to the executed service and the service which called upon the executed service. This precludes adapting the timeout values to various system states. For example, in periods of low overall system use, timeouts for services could be increased. This would allow the services more time to access system resources to resolve their calls. This also reduces the number of service call failures without affecting other services. As another example, in periods of high system resource usage, timeouts for high-latency services could be decreased. This would result in more failures for that service, but would allow the other services to complete their service calls.

[0009] According to various embodiments, a timeout management application allows the timeouts of services to be redefined in order to adapt to the state of the environment in which they are implemented. Timeouts can be increased or decreased based on service latency and the amount of system resources in use. The timeout management application can also take into account service dependencies when modifying timeout values. For example, a parent service may call several child services, and both the parent and child services have their own timeouts. The child or parent service timeouts can be modified to ensure the timeouts meet predefined criteria with respect to each other. These dependencies can also affect which services are modified so as to adapt to the amount of system usage. The timeout management application may enforce a preference for modifying independent services, to minimize the number of affected services. In another example, the timeout management application may want to modify timeouts of services upon which many other services depend to ensure maximum service availability.

[0010] Additionally, the adaptive ability of a timeout management application can take into account various tiers or priorities associated with the services to ensure that high priority or tier services have maximized availability. In the following discussion, a general description of the system and its components is provided, followed by a discussion of the operation of the same.

[0011] With reference to FIG. 1, shown is a networked environment 100 according to various embodiments. The networked environment 100 includes a computing environment 101, and a client 104, which are in data communication with each other via a network 107. The

network 107 includes, for example, the Internet, intranets, extranets, wide area networks (WANs), local area networks (LANs), wired networks, wireless networks, or other suitable networks, *etc.*, or any combination of two or more such networks.

[0012] The computing environment 101 may comprise, for example, a server computer or any other system providing computing capability. Alternatively, the computing environment 101 may employ a plurality of computing devices that may be employed that are arranged, for example, in one or more server banks or computer banks or other arrangements. Such computing devices may be located in a single installation or may be distributed among many different geographical locations. For example, the computing environment 101 may include a plurality of computing devices that together may comprise a cloud computing resource, a grid computing resource, and/or any other distributed computing arrangement. In some cases, the computing environment 101 may correspond to an elastic computing resource where the allotted capacity of processing, network, storage, or other computing-related resources may vary over time.

[0013] Various applications and/or other functionality may be executed in the computing environment 101 according to various embodiments. Also, various data is stored in a data store 111 that is accessible to the computing environment 101. The data store 111 may be representative of a plurality of data stores 111 as can be appreciated. The data stored in the data store 111 for example, is associated with the operation of the various applications and/or functional entities described below.

[0014] The components executed on the computing environment 101, for example, include a timeout management application 114, a plurality of services 117, and other applications, services, processes, systems, engines, or functionality not discussed in detail herein. The timeout management application 114 is executed to modify the timeouts 121 corresponding to the respective pairs of the executed services 117 based on the state of the computing environment 101.

[0015] Services 117 comprise functionality that, in response to a call, perform some action or return some data. Services 117 may comprise functionality to serve network pages, read data from or store data in a data store, perform data transformations, communicate with other applications, manage network traffic, or some other functionality. Services 117 may be

embodied in a representational state transfer (RESTful) architecture, or some other architecture. A service 117 may require the execution of one or more child services 117 to complete a call to the service 117. Services 117 may be executed and called to facilitate the handling of requests 122 sent by a client 104 to generate a response 123, or for other purposes.

[0016] Each of the services 117 executed in the computing environment 101 is associated with at least one timeout 121. A called service 117 may have many timeouts 121, each corresponding to a different service 117 which calls the called service 117. The timeout 121 defines a maximum threshold of time to complete a call to the service 117. If the service 117 fails to complete a call within the timeout 121, the service 117 returns an error to the source of the call.

[0017] Each of the services 117 also corresponds to a latency indicative of the time it takes to complete a call to the corresponding one of the services 117. The latency may be affected by the amount of used system resources in the computing environment 101 relative to the total capacity of system resources. For example, the computing environment 101 may have a high amount of used system resources when handling many service 117 calls, or when the services 117 called are computationally expensive, resulting in an increased latency in some services 117.

[0018] The data stored in the data store 111 includes, for example, data embodying service tiers 124, a service priority list 127, a dependency model 131, and potentially other data. The service tiers 124 represent a grouping of the executed services 117 into one or more groupings. The groupings may be mutually exclusive or allow for a service 117 to belong to multiple groups. The groupings may be themselves ranked or ordered in some hierarchy. The groupings may also be based on some qualitative assessment of the included services 117, such as by priority.

[0019] The service priority list 127 embodies a ranking of the executed services 117 based on their priority of availability. The service priority list 127 may comprise the entirety of the executed services 117 or a subset thereof. Multiple service priority lists 127 may be stored for retrieval as a function of one or more predefined conditions.

[0020] The dependency model 131 embodies relationships and dependencies between the executed services 117. The structure of the dependency model 131 may embody a tree, some

other acyclical directed graph, or another data structure. Nodes in the dependency model 131 representing services 117 may be related to parent or child nodes as a function of a service 117 calling other services 117 during execution. The dependency model 131 may embody a completely connected graph, or may allow independent nodes.

[0021] The client 104 is representative of a plurality of client devices that may be coupled to the network 107. The client 104 may comprise, for example, a processor-based system such as a computer system. Such a computer system may be embodied in the form of a desktop computer, a laptop computer, personal digital assistants, cellular telephones, smartphones, set-top boxes, music players, web pads, tablet computer systems, game consoles, electronic book readers, or other devices with like capability.

[0022] The client 104 may be configured to execute various applications such as a client application 142 and/or other applications. The client application 142 may be executed in a client 104 for example, to access network content served up by the computing environment 101 and/or other servers. The client application 142 may, for example, correspond to a browser, a mobile application, *etc.* The client 104 may be configured to execute applications beyond the client application 142 such as, for example, browsers, mobile applications, email applications, social networking applications, and/or other applications.

[0023] The client 104 may communicate requests 122 to the computing environment 101 to obtain responses 123. These requests 122 may comprise requests for network pages to be served by the computing environment 101, requests for data, requests for some function or application to be executed on the computing environment 101, or another type of request. The request may be communicated by hypertext transfer protocol (HTTP), file transfer protocol (FTP), simple object access protocol (SOAP), simple mail transfer protocol (SMTP), by another protocol, or by some other approach. Communicating the requests 122 may be facilitated by the client application 142. For example, a request 122 for a network page may be facilitated by a browser client application 142. Other client applications 142 may also be executed to facilitate the communication of requests 122 to the computing environment 101.

[0024] Next, a general description of the operation of the various components of the networked environment 100 is provided. To begin, the computing environment 101 executes one or more services 117 in order to facilitate the handling of the requests 122 communicated by the

client 104 and the generation of the responses 123. The timeout management application 114 obtains latency information corresponding to each of the executed services 117. In some embodiments, latency data may be obtained from a monitoring process. For example, services 117 may report latency data to a log accessible to the monitoring process, which then aggregates the data for communication to the timeout management application 114. The services 117 may communicate latency data to the logs at a predefined interval, or in response to some event. The latency data may comprise the most recent latency at the time of communication to the logs, or latency data aggregated over a period of time. Latency data may also be communicated to the logs by another approach.

[0025] In another embodiment, latency data may be obtained by querying functionality of the service 117. For example, a service 117 may comprise functionality to maintain latency data and functionality to return that latency data in response to a specified call to the service 117. Latency data may also be obtained by determining the latency as a function of data from entities which called the service 117. Other techniques may also be used to obtain the latency data by the timeout management application 114.

[0026] The timeout management application 114 then determines the amount of system resources used in the computing environment 101 relative to the total capacity of the computing environment 101. The amount of system resources may be determined as a function of the obtained latency data, as well as other data, information, and logs. The timeout management application 114 may then adjust the timeout 121 of a selected service in response to an event associated with the amount of used system resources.

[0027] In some embodiments, the timeout management application 114 may adjust a timeout 121 of a service 117 when the amount of system resources used in the computing environment 101 falls below a predetermined threshold or below a minimum value of a predefined range. In these embodiments, such a condition would indicate that the computing environment 101 has processing capacity available to allow for a service 117 to have its timeout increased.

[0028] In other embodiments, the timeout management application 114 may adjust a timeout 121 of a service 117 when the amount of system resources used in the computing environment 101 exceeds a predetermined threshold or above a maximum value of a predefined

range. In these embodiments, such a condition would indicate that the computing environment 101 is at or is approaching capacity, at which the executed services 117 may experience latencies which exceed their timeouts 121, thereby resulting in increased user failure. Decreasing the timeout 121 of a selected service 117 would increase failure for the selected service 117, but allow other services 117 to capitalize on available system resources of the computing environment 101.

[0029] The timeout management application 114 may also adjust a timeout 121 of a service 117 as a function of system reliability. For example, the timeout management application 114 may determine an amount of system reliability based on the number of service 117 calls which successfully execute, and/or other data. A timeout 121 modification may be triggered if the system reliability falls below a predetermined reliability threshold. Other events may also trigger a timeout 121 modification by the timeout management application 114.

[0030] When an event occurs that triggers a timeout 121 modification by the timeout management application 114, the timeout management application 114 must select a service 117 whose timeout 121 will be modified. In some embodiments, this selection made as a function of service tiers 124 associated with the executed services 117. As a non-limiting example, services 117 may be organized into multiple tiers such as high priority, medium priority, and low priority service tiers 124 depending on the priority of the service 117 completing a call without error. If the timeout management application 114 is modifying the timeout 121 of a service 117 due to high available service capacity, the timeout management application 114 may select a service 117 belonging to the high priority service tier 124 to maximize the availability of the high priority service 117. If the timeout management application 114 is modifying the timeout 121 of a service 117 due to low available service capacity, the timeout management application 114 may select a service 117 belonging to the low or medium priority service tiers 124 to minimize impact on high priority services 117. Service tiers 124 may be used in other techniques to select a service 117 for timeout 121 modification, as well.

[0031] The timeout management application 114 may also select a service 117 for timeout 121 modification based at least in part on a service priority list 127. The service priority list 127 may contain a ranked or ordered list of the executed services 117, and the selection of a service 117 is based on their placement on the service priority list 127. The timeout management

application 114 may refer to a single service priority list 127 or multiple service priority lists 127. For example, in embodiments in which multiple events can trigger a modification of a service 117 timeout 121, the timeout management application 114 may consult one or more of a plurality of service priority lists 127 depending on the event. Service priority lists 127 may also be used in other techniques to select a service 117 for timeout 121 modification.

[0032] The timeout management application 114 may select a service 117 for timeout 121 modification based at least in part on a dependency model 131. For example, in embodiments in which the dependency model 131 comprises independent services 117, the timeout management application 114 may preferentially modify the timeout 121 of independent services 117 to minimize the impact on other executed services 117. In another embodiment in which a timeout 121 is to be increased, the timeout management application 114 may preferentially modify the timeout 121 of a service 117 which is depended on by multiple services 117 to maximize the availability of more services 117.

[0033] In embodiments in which the timeout management application 114 selects a service 117 for timeout 121 modification based at least in part on a dependency model 131, the timeout management application 114 may recursively select further services 117 for timeout 121 modification. For example, if a child service 117 has its timeout 121 increased, then the timeout management application 114 may then recursively increase the timeouts 121 of the parent services 117. As another example, if a parent service 117 has its timeout 121 decreased, the timeout management application 114 may then recursively decrease the timeouts 121 of the child services 117.

[0034] The timeout management application 114 may also select a service 117 for timeout 121 modification based on the latency or amount of used resources associated with the execution of the service 117. For example, a high latency service 117 may have its timeout increased to decrease the chance of a service 117 call failure. As another example, a high latency service 117 may have its timeout decreased to increase failure of its calls, but also free a greater amount of the service capacity for other services 117. Latency or the amount of used resources associated with the execution of the service 117 may be used in other approach, as well.

[0035] The previously discussed examples of techniques to select a service 117 for timeout 121 modification are non-limiting examples, and other techniques may also be used to select a service 117 for timeout 121 modification. Additionally, any of the previously discussed techniques as well as other techniques may be used alone or in combination with one another. For example, services 117 may be organized into service tiers 124, which are ranked within each service tier 124 according to a service priority list 127. Other combinations of techniques may also be used.

[0036] Once a service 117 has been selected for timeout 121 modification, the timeout management application 114 assigns the service 117 a new timeout 121. The new timeout 121 may be a function of the timeout 121 to be modified. For example, the new timeout 121 may comprise a percentage of the original timeout 121. The new timeout 121 may also comprise the original timeout 121 incremented or decremented by a predetermined interval. In embodiments employing service tiers 124, the new timeout 121 may be based on the service tier 124 to which the service 117 belongs. For example, a service 117 in a high priority service tier 124 may be assigned a new timeout 121 that is five seconds greater than the original timeout 121, while a service 117 of a medium service tier 124 may be assigned a new timeout 121 that is three seconds greater than the original timeout 121. Service tiers 124 may also be used in another approach to assign the new timeout 121.

[0037] In embodiments in which the services 117 are related using a dependency model 131, the new timeout 121 may be a function of the timeouts corresponding to parent or child services 117. As a non-limiting example, a new timeout 121 may comprise some amount greater than the summation of the timeouts of child services 117. As another example, a new timeout 121 may comprise some amount that is less than the timeout 121 of the associated parent services 117. Other functions may also be used to set parent and child timeouts with respect to each other.

[0038] Additionally, the new timeout 121 value may be a function of maximum or minimum acceptable timeouts 121. For example, a timeout 121 may not be set below some minimum threshold. If the selected service 117 timeout 121 is already at that minimum threshold, then the timeout management application 114 may select a different service 117. Minimum and maximum timeout 121 thresholds may comprise single thresholds applied to all services 117, or

may comprise thresholds that vary depending on the associated service 117. Other approaches may be used to determine a new timeout 121 based at least in part on the dependency model 131.

[0039] After the timeout management application 114 has assigned a new timeout 121 to the selected service 117, the timeout management application 114 continues to monitor the latency data of the executed service 117. Though the previously operations of the timeout management application 114 were discussed in the context of selecting a single service 117 for timeout 121 modification, the timeout management application 114 may also select multiple services 117 for timeout 121 modification in response to an event in some embodiments.

[0040] The timeout management application 114 may repeatedly select services 117 and modify their timeouts 121 as a function of the state of the computing environment 101, such as until the amount of used system resources is within a predefined range, or by another approach.

[0041] Referring next to FIG. 2, shown is a flowchart that provides one example of the operation of a portion of the timeout management application 114 (FIG. 1) according to various embodiments. It is understood that the flowchart of FIG. 2 provides merely an example of the many different types of functional arrangements that may be employed to implement the operation of the portion of the timeout management application 114 as described herein. As an alternative, the flowchart of FIG. 2 may be viewed as depicting an example of steps of a method implemented in the computing environment 101 (FIG. 1) according to one or more embodiments.

[0042] FIG. 2 shows an example embodiment of the timeout management application 114 reconfiguring the timeouts 121 (FIG. 1) of executed services 117 (FIG. 1) while the amount of used service capacity is outside of a predefined acceptable range. Services 117 are iteratively selected and their timeouts 121 modified until the used service capacity is within the predefined range.

[0043] Beginning with box 201, the timeout management application 114 obtains latency data associated with at least one service 117 executed in the computing environment 101. In some embodiments, the executed services 117 update log files with data associated with their latencies. The data may comprise the latency at the instance that the log file is updated, an aggregate function of latency over a predefined time period, or other data. In other embodiments, the timeout management application 114 may directly query functionality of the

executed services 117 which returns latency data. Other approaches may also be used to obtain the latency data.

[0044] Next, in box 204, the timeout management application 114 calculates an amount of used service capacity of the computing environment 101. The amount of used service capacity may correspond to the entirety of the computing environment 101 or a designated subcomponent or set of subcomponents. The amount of used service capacity may be calculated as a function of the obtained latency data, data obtained from another monitoring or profiling service, or other data as can be appreciated.

[0045] In box 207, the timeout management application 114 determines if the amount of used service capacity is outside of a predefined range. The predefined range may be a function of the total available service capacity or other data, or may be a preset independent value. If the amount of used service capacity is inside of the predefined range, the process ends, indicating that the computing environment 101 is operating within acceptable capacity and the timeouts 121 of the services 117 do not need to be adjusted.

[0046] If the amount of used service capacity is outside of the predefined range, the process moves to box 211 wherein the timeout management application 114 selects a service 117 for timeout 121 reconfiguration. In some embodiments, the services 117 are associated with at least one service tier 124 (FIG. 1). In such embodiments, the selection may be made as a function of the service tiers 124. For example, if the amount of used service capacity was below the minimum value of the predefined range, the timeout management application 114 may select a service 117 in a high priority service tier 124 to have their timeout 121 increased so as to allocate to it more service capacity. As another example, if the amount of used service capacity was above the maximum value of the predefined range, the timeout management application 114 may select a service 117 in a low priority service tier 124 to have their timeout 121 decreased so as to reduce the capacity used by lower priority services 117. Service tiers 124 may also be used in other techniques to select a service 117.

[0047] In embodiments in which the services 117 are ranked in at least one service priority list 127, the selection by the timeout management application 114 may be made as a function of the service priority list 127. For example, if the amount of used service capacity was above the maximum value of the predefined range, the timeout management application 114 may select a

service 117 ranked lower on the service priority list 127 for modification to minimize the effects on higher priority services 117. If the amount of used service capacity was below the minimum value of the predefined range, the timeout management application 114 may select a service 117 ranked higher on the service priority list 127 to maintain reliability of those high priority services 117. Service priority lists 127 may be used in selecting a service 117 for timeout 121 modification in another approach.

[0048] In embodiments in which the services 117 are related by a dependency model 131 (FIG. 1), the selection may be a function of the dependency model. For example, if the amount of used service capacity is below the minimum value of the predefined threshold, the timeout management application 114 may select a service 117 which is depended on by other services 117 to improve the reliability of the dependent services 117. As another example, if the amount of used service capacity is above the maximum value of the predefined threshold, the timeout management application 114 may select a service 117 which is independent of other services to have their timeout 121 reduced, minimizing the effect on other services 117. Dependency models 131 may be used to select services 117 for timeout 121 modification in other approaches, as well.

[0049] The selection of a service 117 may also be based on the latency data. For example, a high latency service 117 may be selected to have its timeout 121 reduced to prevent overuse of service capacity by the high latency service 117. The selection of a service 117 may also be based on maximum or minimum timeout 121 thresholds. For example, the timeout management application 114 would not select a service 117 to have its timeout 121 reduced when the timeout 121 already equals a minimum timeout 121 threshold. Other data, functions, and techniques may also be used to select the service 117.

[0050] After the service 117 has been selected, the timeout 121 of the selected service 117 is reconfigured in box 214. The new timeout 121 value may be a function of the amount of used service capacity. For example, if the amount of used service capacity is above the maximum value of the predefined threshold, the new timeout 121 value may be less than the original timeout 121 value. In embodiments in which the services 117 are related by a dependency model 131, the new timeout 121 value may be a function of the timeout 121 values of the related services 117. For example, if the selected service 117 is having their timeout 121 decreased, the

new timeout 121 value may not be set below the sum or some other function of the timeouts 121 of services 117 upon which the selected service 117 depends. As another example, reconfiguring the timeout 121 of a service 117 may also require the iterative reconfiguration of timeouts 121 for services upon which the selected service 117 depends. Other approaches may also be used by the timeout management application 114 to determine the new timeout 121 value.

[0051] Once the timeout 121 of the selected service 117 has been reconfigured, the process repeats the steps of selecting services 117 and reconfiguring their timeouts 121 until the amount of used service capacity is within the predefined range.

[0052] Referring next to FIG. 3, shown is a flowchart that provides one example of the iterative timeout 121 (FIG. 1) reconfiguration using a dependency model 131 (FIG. 1) of the timeout management application 114 (FIG. 1) of box 214 (FIG. 2) according to various embodiments. It is understood that the flowchart of FIG. 3 provides merely an example of the many different types of functional arrangements that may be employed to implement the operation of the portion of the timeout management application 114 as described herein. As an alternative, the flowchart of FIG. 3 may be viewed as depicting an example of steps of a method implemented in the computing environment 101 (FIG. 1) according to one or more embodiments.

[0053] FIG. 3 describes one embodiment of the timeout management application 114 implementing a dependency model 131. A selected service 117 (FIG. 1) whose timeout 121 has been modified may be related to several child services 117 upon which it depends. The timeouts 121 of a child service 117 should sum to be less than the timeout 121 of the parent service 117. The timeout management application 114 recursively traverses the dependency model 131 to modify the timeouts 121 of child services 117 to ensure that the sum of the child service 117 timeouts 121 is less than their parent service 117 timeout 121.

[0054] Beginning with box 301, after a parent service 117 has had their timeout 121 reconfigured to a lower value in box 211 (FIG. 2), the timeout management application 114 obtains the child services 117 for the selected parent service 117 using the dependency model 131. Obtaining child services 117 may be accomplished through a tree traversal algorithm, a graph search algorithm such as a breadth first search or depth first search, or by another approach.

[0055] In box 304, if the selected service 117 has no child services 117, the process ends. Otherwise, in box 307, the timeout management application 114 sums the timeouts 121 of the obtained child services 117. If the sum of the child service 117 timeouts 121 with respect to the selected service 117 is less than the timeout 121 of the selected service 117, then the process ends. Otherwise, if the sum of the timeouts 121 of the child services 117 exceeds the timeout 121 of the selected service 117, the timeouts 121 of at least one the child services 117 must be reconfigured.

[0056] In box 311, at least one of the child services 117 is selected for timeout 121 reconfiguration. In some embodiments, the entirety of the child services 117 may be selected. In other embodiments, a subset of the child services 117 may be selected. The selection may be a function of latency data associated with the child services 117 (FIG. 1), service tiers 124 (FIG. 1), a service priority list 127 (FIG. 1), the dependency model 131, minimum or maximum timeout thresholds or other data, by some technique described with respect to box 211 (FIG. 2) or by another approach.

[0057] In box 314, once the child services 117 have been selected, their corresponding timeouts 121 are reconfigured as described in box 214, or by another approach. After the child services 117 have been reconfigured, the process iterates for each of the reconfigured child services 117. The process will continue until all services 117 including and descending from the original selected service 117 have timeouts 121 greater than the sum of the timeouts 121 of their child services 117 or have no child services 117.

[0058] With reference to FIG. 4, shown is a schematic block diagram of the computing environment 101 according to an embodiment of the present disclosure. The computing environment 101 includes one or more computing devices 401. Each computing device 401 includes at least one processor circuit, for example, having a processor 402 and a memory 404, both of which are coupled to a local interface 407. To this end, each computing device 401 may comprise, for example, at least one server computer or like device. The local interface 407 may comprise, for example, a data bus with an accompanying address/control bus or other bus structure as can be appreciated.

[0059] Stored in the memory 404 are both data and several components that are executable by the processor 402. In particular, stored in the memory 404 and executable by the

processor 402 are a timeout management application 114 (FIG. 1), one or more services 117 (FIG. 1) having a corresponding timeout 121 (FIG. 1), and potentially other applications. Also stored in the memory 404 may be a data store 111 (FIG. 1) comprising service tiers 124 (FIG. 1), a service priority list 127 (FIG. 1) or a dependency model 131 (FIG. 1), and other data. In addition, an operating system may be stored in the memory 404 and executable by the processor 402.

[0060] It is understood that there may be other applications that are stored in the memory 404 and are executable by the processor 402 as can be appreciated. Where any component discussed herein is implemented in the form of software, any one of a number of programming languages may be employed such as, for example, C, C++, C#, Objective C, Java[®], JavaScript[®], Perl, PHP, Visual Basic[®], Python[®], Ruby, Flash[®], or other programming languages.

[0061] A number of software components are stored in the memory 404 and are executable by the processor 402. In this respect, the term "executable" means a program file that is in a form that can ultimately be run by the processor 402. Examples of executable programs may be, for example, a compiled program that can be translated into machine code in a format that can be loaded into a random access portion of the memory 404 and run by the processor 402, source code that may be expressed in proper format such as object code that is capable of being loaded into a random access portion of the memory 404 and executed by the processor 402, or source code that may be interpreted by another executable program to generate instructions in a random access portion of the memory 404 to be executed by the processor 402, *etc.* An executable program may be stored in any portion or component of the memory 404 including, for example, random access memory (RAM), read-only memory (ROM), hard drive, solid-state drive, USB flash drive, memory card, optical disc such as compact disc (CD) or digital versatile disc (DVD), floppy disk, magnetic tape, or other memory components.

[0062] The memory 404 is defined herein as including both volatile and nonvolatile memory and data storage components. Volatile components are those that do not retain data values upon loss of power. Nonvolatile components are those that retain data upon a loss of power. Thus, the memory 404 may comprise, for example, random access memory (RAM), read-only memory (ROM), hard disk drives, solid-state drives, USB flash drives, memory cards accessed via a memory card reader, floppy disks accessed via an associated floppy disk drive,

optical discs accessed via an optical disc drive, magnetic tapes accessed via an appropriate tape drive, and/or other memory components, or a combination of any two or more of these memory components. In addition, the RAM may comprise, for example, static random access memory (SRAM), dynamic random access memory (DRAM), or magnetic random access memory (MRAM) and other such devices. The ROM may comprise, for example, a programmable read-only memory (PROM), an erasable programmable read-only memory (EPROM), an electrically erasable programmable read-only memory (EEPROM), or other like memory device.

[0063] Also, the processor 402 may represent multiple processors 402 and/or multiple processor cores and the memory 404 may represent multiple memories 404 that operate in parallel processing circuits, respectively. In such a case, the local interface 407 may be an appropriate network that facilitates communication between any two of the multiple processors 402, between any processor 402 and any of the memories 404, or between any two of the memories 404, *etc.* The local interface 407 may comprise additional systems designed to coordinate this communication, including, for example, performing load balancing. The processor 402 may be of electrical or of some other available construction.

[0064] Although the timeout management application 114, and other various systems described herein may be embodied in software or code executed by general purpose hardware as discussed above, as an alternative the same may also be embodied in dedicated hardware or a combination of software/general purpose hardware and dedicated hardware. If embodied in dedicated hardware, each can be implemented as a circuit or state machine that employs any one of or a combination of a number of technologies. These technologies may include, but are not limited to, discrete logic circuits having logic gates for implementing various logic functions upon an application of one or more data signals, application specific integrated circuits (ASICs) having appropriate logic gates, field-programmable gate arrays (FPGAs), or other components, *etc.* Such technologies are generally well known by those skilled in the art and, consequently, are not described in detail herein.

[0065] The flowcharts of FIGS. 2 and 3 show the functionality and operation of an implementation of portions of the timeout management application 114. If embodied in software, each block may represent a module, segment, or portion of code that comprises program instructions to implement the specified logical function(s). The program instructions may be

embodied in the form of source code that comprises human-readable statements written in a programming language or machine code that comprises numerical instructions recognizable by a suitable execution system such as a processor 402 in a computer system or other system. The machine code may be converted from the source code, *etc.* If embodied in hardware, each block may represent a circuit or a number of interconnected circuits to implement the specified logical function(s).

[0066] Although the flowcharts of FIGS. 2 and 3 show a specific order of execution, it is understood that the order of execution may differ from that which is depicted. For example, the order of execution of two or more blocks may be scrambled relative to the order shown. Also, two or more blocks shown in succession in FIGS. 2 and 3 may be executed concurrently or with partial concurrence. Further, in some embodiments, one or more of the blocks shown in FIGS. 2 and 3 may be skipped or omitted. In addition, any number of counters, state variables, warning semaphores, or messages might be added to the logical flow described herein, for purposes of enhanced utility, accounting, performance measurement, or providing troubleshooting aids, *etc.* It is understood that all such variations are within the scope of the present disclosure.

[0067] Also, any logic or application described herein, including the timeout management application 114, that comprises software or code can be embodied in any non-transitory computer-readable medium for use by or in connection with an instruction execution system such as, for example, a processor 402 in a computer system or other system. In this sense, the logic may comprise, for example, statements including instructions and declarations that can be fetched from the computer-readable medium and executed by the instruction execution system. In the context of the present disclosure, a "computer-readable medium" can be any medium that can contain, store, or maintain the logic or application described herein for use by or in connection with the instruction execution system.

[0068] Various embodiments of the disclosure can be described in view of the following clauses:

[0069] 1. A non-transitory computer-readable medium embodying a program executable in at least one computing device, comprising:

[0070] code that obtains latency data from a plurality of services, each one of the services associated with a plurality of timeouts and associated with one of a plurality of tiers, each of the

timeouts corresponding to a pair of the services, and at least one of the services being a parent one of the service associated with a subset of child services;

[0071] code that calculates an aggregate latency with respect to a time interval based at least in part on the latency data obtained during the time interval;

[0072] code that calculates an amount of used service capacity based at least in part on the aggregate latency;

[0073] code that selects one of the services based at least in part on the tiers, a service dependency, the aggregate latency, or a priority list;

[0074] code that, in response to the amount of used service capacity exceeding a first predefined threshold, decreases the one of the timeouts corresponding to the selected one of the services; and

[0075] code that, in response to the amount of used service capacity falling below a second predefined threshold, increases the one of the timeouts corresponding to the selected one of the services.

[0076] 2. The non-transitory computer-readable medium of clause 1, wherein the code that selects one of the services is further based at least in part on a reliability threshold.

[0077] 3. The non-transitory computer-readable medium of clause 1, wherein the selected one of the services is one of the child services, and further comprising code that reconfigures the one of the timeouts corresponding to the select one of the services based at least in part on the one of the timeouts corresponding to the parent one of the services associated with the selected one of the services.

[0078] 4. A system, comprising:

[0079] at least one computing device; and

[0080] a timeout management application executable in the at least one computing device, the timeout management application comprising:

[0081] logic that monitors a plurality of latencies, each latency corresponding to one of a plurality of services; and

[0082] logic that sets a timeout associated with a pair of the services based at least in part on the latencies.

[0083] 5. The system of clause 4, wherein the timeout management further comprises logic that calculates an aggregate latency relative to a time interval based at least in part on the latencies, and the timeout is further based at least in part on the aggregate latency.

[0084] 6. The system of clause 4, wherein the timeout is further based at least in part on an amount of used service capacity.

[0085] 7. The system of clause 6, wherein the logic that sets the timeout increases the timeout in response to the amount of used service capacity relative to a total service capacity falling below a threshold.

[0086] 8. The system of clause 6, wherein the logic that sets the timeout decreases the timeout in response to the amount of used service capacity relative to a total service capacity exceeding a threshold.

[0087] 9. The system of clause 6, wherein the timeout management further comprises logic that selects one of the services based at least in part on a subset of the amount of used system resources used by the selected service.

[0088] 10. The system of clause 4, wherein the timeout management further comprises logic that selects one of the services based at least in part on a service priority list, and the timeout is associated with the selected one of the services.

[0089] 11. The system of clause 4, wherein the timeout management further comprises logic that selects one of the services based at least in part on a reliability threshold.

[0090] 12. The system of clause 4, wherein the logic that monitors the plurality of latencies further comprises logic that obtains log data from the services.

[0091] 13. A method, comprising the steps of:

[0092] obtaining, in a computing device, latency data from a plurality of services, each of the services having a plurality of timeouts, and each of the timeouts being associated with another one of the services;

[0093] selecting, in the computing device, one of the services based at least in part on the latency data; and

[0094] reconfiguring, in the computing device, the one of the timeouts corresponding to the selected one of the services.

[0095] 14. The method of clause 13, wherein at least one of the services is dependent on a subset of the services.

[0096] 15. The method of clause 14, wherein the selected one of the services is one of the subset of the services, and reconfiguring the one of the timeouts is based at least in part on the one of the timeouts corresponding to the one of the services which depends on the selected one of the services.

[0097] 16. The method of clause 14, wherein selecting one of the services is further based at least in part on a degree of dependency.

[0098] 17. The method of clause 13, wherein selecting one of the services is further based at least in part on a priority list.

[0099] 18. The method of clause 13, wherein the services are each associated with one of a plurality of tiers, and the selecting one of the services is further based on the tiers.

[0100] 19. The method of clause 13, further comprising determining, in the computing device, an amount of used service capacity based at least in part on the latency data, and wherein the steps of selecting one of the services and reconfiguring one of the timeouts are performed in response to the amount of used service capacity falling outside of a predefined range.

[0101] 20. The method of clause 19, further comprising repeating the steps of selecting one of the services and reconfiguring the one of the timeouts until the amount of used service capacity falls within the predefined range.

[0102] The computer-readable medium can comprise any one of many physical media such as, for example, magnetic, optical, or semiconductor media. More specific examples of a suitable computer-readable medium would include, but are not limited to, magnetic tapes, magnetic floppy diskettes, magnetic hard drives, memory cards, solid-state drives, USB flash drives, or optical discs. Also, the computer-readable medium may be a random access memory (RAM) including, for example, static random access memory (SRAM) and dynamic random access memory (DRAM), or magnetic random access memory (MRAM). In addition, the computer-readable medium may be a read-only memory (ROM), a programmable read-only

memory (PROM), an erasable programmable read-only memory (EPROM), an electrically erasable programmable read-only memory (EEPROM), or other type of memory device.

[0103] It should be emphasized that the above-described embodiments of the present disclosure are merely possible examples of implementations set forth for a clear understanding of the principles of the disclosure. Many variations and modifications may be made to the above-described embodiment(s) without departing substantially from the spirit and principles of the disclosure. All such modifications and variations are intended to be included herein within the scope of this disclosure and protected by the following claims.

CLAIMS

Therefore, the following is claimed:

1. A system, comprising:
at least one computing device; and
a timeout management application executable in the at least one computing device, the timeout management application comprising:
logic that monitors a plurality of latencies, each latency corresponding to one of a plurality of services; and
logic that sets a timeout associated with a pair of the services based at least in part on the latencies.
2. The system of claim 1, wherein the timeout management further comprises logic that calculates an aggregate latency relative to a time interval based at least in part on the latencies, and the timeout is further based at least in part on the aggregate latency.
3. The system of claim 1, wherein the timeout is further based at least in part on an amount of used service capacity.
4. The system of claim 3, wherein the logic that sets the timeout increases the timeout in response to the amount of used service capacity relative to a total service capacity falling below a threshold.
5. The system of claim 3, wherein the logic that sets the timeout decreases the timeout in response to the amount of used service capacity relative to a total service capacity exceeding a threshold.
6. The system of claim 3, wherein the timeout management further comprises logic that selects one of the services based at least in part on a subset of the amount of used system resources used by the selected service.

7. The system of claim 1, wherein the timeout management further comprises logic that selects one of the services based at least in part on a service priority list, and the timeout is associated with the selected one of the services.

8. The system of claim 1, wherein the timeout management further comprises logic that selects one of the services based at least in part on a reliability threshold.

9. The system of claim 1, wherein the logic that monitors the plurality of latencies further comprises logic that obtains log data from the services.