

(19) World Intellectual Property Organization
International Bureau



(43) International Publication Date
17 January 2008 (17.01.2008)

PCT

(10) International Publication Number
WO 2008/008883 A2

(51) International Patent Classification:
G06F 15/173 (2006.01)

(21) International Application Number:
PCT/US2007/073348

(22) International Filing Date: 12 July 2007 (12.07.2007)

(25) Filing Language: English

(26) Publication Language: English

(30) Priority Data:
11/487,130 14 July 2006 (14.07.2006) US

(71) Applicant (for all designated States except US): **DATA-SYNAPSE, INC.** [US/US]; 632 Broadway, 5th Floor, New York, NY 10012 (US).

(72) Inventors; and

(75) Inventors/Applicants (for US only): **BERNARDIN, James** [US/US]; 49 West 96th Street, Apt. 4a, New York, NY 10025 (US). **ISKOLD, Alexander** [US/US]; 24 Stoneham Drive, Livingston, NJ 07039 (US).

(74) Agents: **BEAN, Thomas, J.** et al.; Darby & Darby P.C., P.O. Box 770, Church Street Station, New York, NY 10008-0770 (US).

(81) Designated States (unless otherwise indicated, for every kind of national protection available): AE, AG, AL, AM, AT, AU, AZ, BA, BB, BG, BH, BR, BW, BY, BZ, CA, CH, CN, CO, CR, CU, CZ, DE, DK, DM, DO, DZ, EC, EE, EG, ES, FI, GB, GD, GE, GH, GM, GT, HN, HR, HU, ID, IL, IN, IS, JP, KE, KG, KM, KN, KP, KR, KZ, LA, LC, LK, LR, LS, LT, LU, LY, MA, MD, ME, MG, MK, MN, MW, MX, MY, MZ, NA, NG, NI, NO, NZ, OM, PG, PH, PL, PT, RO, RS, RU, SC, SD, SE, SG, SK, SL, SM, SV, SY, 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, LS, MW, MZ, NA, SD, SL, SZ, TZ, UG, ZM, ZW), Eurasian (AM, AZ, BY, KG, KZ, MD, RU, TJ, TM), European (AT, BE, BG, CH, CY, CZ, DE, DK, EE, ES, FI, FR, GB, GR, HU, IE, IS, IT, LT, LU, LV, MC, MT, NL, PL, PT, RO, SE, SI, SK, TR), OAPI (BF, BJ, CF, CG, CI, CM, GA, GN, GQ, GW, ML, MR, NE, SN, TD, TG).

Published:

— without international search report and to be republished upon receipt of that report

For two-letter codes and other abbreviations, refer to the "Guidance Notes on Codes and Abbreviations" appearing at the beginning of each regular issue of the PCT Gazette.

(54) Title: METHOD FOR ALLOCATING SHARED COMPUTING INFRASTRUCTURE FOR APPLICATION SERVER-BASED DEPLOYMENTS

(57) Abstract: A shared computing infrastructure includes a plurality of computing engines, applications servers, and computing domains. A broker component executes a method for dynamically allocating the computing engines among the computing domains. The allocation method begins with the step of determining an expected number of computing engines to be allocated to each of the computing domains as a function of a predetermined service policy and a predicted demand for the domain. While fewer than the expected number of computing engines has been allocated to each domain, the computing domains are sequentially selecting as a function of predetermined domain priorities. Unallocated computing engines are identified, and the unallocated computing engines are allocated to each selected computing domain according to predetermined selection rules for the domain. During an allocation improvement step, allocations among the computing domains are further adjusted to maximize a fitness statistic computed for the allocations.



WO 2008/008883 A2

METHOD FOR ALLOCATING SHARED COMPUTING INFRASTRUCTURE FOR APPLICATION SERVER-BASED DEPLOYMENTS

CROSS-REFERENCE TO RELATED APPLICATIONS

5 The present application claims priority to U.S. Patent Application Serial No. 11/487,130, filed July 14, 2006, the contents of which are incorporated by reference in its entirety.

FIELD OF THE INVENTION

 The present invention is directed to a method for dynamically allocating shared computing resources in a distributed computing environment. More specifically, the present
10 invention is directed to a method for dynamically allocating shared computing resources among a plurality of computing domains that host a plurality of software applications running in a plurality of run-time environments by means of a rules-based allocation policy.

BACKGROUND OF THE INVENTION

 The global marketplace is forcing companies to respond quickly to dynamic market
15 conditions while reducing costs. Businesses increasingly must have the ability to meet, or beat, competitors by introducing new and innovative products and services. These new offerings are often customer-facing and transaction-oriented, and introduce additional complexity and higher levels of volatility for the enterprise computing resources called upon to provision these products and services. Higher transaction volumes and demands for improved response times create an ever-
20 increasing need for computing resources.

 In a conventional enterprise computing environment, computing resources are usually manually assigned and provisioned to support various applications. This approach creates several problems. As the assigned resources are generally fixed at a point in time to meet a current demand level, the conventional enterprise computing environment is ill-equipped to adapt over time to meet
25 increasing demand levels for some applications and decreasing demand levels for others. In order to meet minimum service requirements, computing resources are often assigned and provisioned according to peak-level demands. As a result, during periods of less than peak-level demands, computing resources are underutilized.

With the advent of grid computing, conventional enterprise computing environments have been adapted to “virtualize” applications so that computing resources may be dynamically provisioned to applications in response to current demand levels. For example, the GRIDSERVER Virtual Enterprise Edition adaptive grid infrastructure software available from DataSynapse, New York, New York provides a computing operating environment that virtualizes application and data services, independent of specific system resources. Client applications submit service requests to the grid environment, and GRIDSERVER dynamically provisions services on specific system resources in the grid to meet the service requests. For example, requests from multiple client applications cause GRIDSERVER to create multiple service instances to handle the requests in parallel on different computing resource nodes in the computing resources grid. As a result, underutilization of resources can be substantially reduced, and service levels can be commensurately improved.

GRIDSERVER has been particularly effective at providing a virtualized computing environment that adapts to meet resource demands for computing-intensive processes. U.S. Patent Application No. 11/395,586, upon which the present application is based and which is incorporated by reference herein in its entirety, discloses a virtualized computing environment that effectively adapts to meet resource demands for high throughput, low latency transactional applications such as distributed web applications and other services-based application. The present application is directed to an inventive method by which computing resources in this virtualized computing environment may be dynamically and adaptively provisioned to optimally serve the web and other services-based applications.

SUMMARY OF THE INVENTION

The present invention is directed to a method for adaptively provisioning a shared computing infrastructure to support a plurality of software applications and a plurality of types of applications servers each providing a run-time environment for one or more of the software applications. The shared computing infrastructure includes computing engines assigned to execute instances of the plurality of software applications, clients accessing the computing engines to request and receive services from the software applications, and a broker that dynamically allocates and provisions computing engines to the clients for executing the software applications.

The broker includes an optimization module that periodically determines an optimal allocation of the computing engines to the software applications and applications servers. To

provision resources based on an optimal allocation, the broker device also includes a configuration manager for reconfiguring a computing engine by halting a current instance of a software application of a first type, and for loading and starting an instance of a software application of a second type.

- 5 According to the present invention, a method is performed by the optimization module of the broker for allocating the plurality of computing engines among the plurality of computing domains. The method begins with the step of determining an expected number of computing engines to be allocated to each of the plurality of computing domains, where the expected number is determined as a function of a predetermined service policy and a predicted demand for the computing domain.
- 10 Then, while one or more of the plurality of computing engines is unallocated and one or more of the plurality of computing domains has been allocated fewer than its expected number of computing resources, the method continues by sequentially selecting a next computing domain as a function of predetermined allocation priorities for the computing domains, selecting an unallocated computing engine according to engine selection rules for the selected computing domain, and allocating the
- 15 selected computing engine to the selected computing domain.

- After the computing engines have been initially assigned to the computing domains, the method enters an allocation improvement stage. First, a fitness value is calculated for the current allocation of computing engines to computing domains. A first computing domain is identified having fewer than its expected number of allocated computing engines, and a computing engine is
- 20 identified for re-allocation to the first computing domain. The identified computing engine may be an unallocated computing engine, or may be a computing engine allocated to a second computing domain to which the expected number of computing engines has been fully allocated. After re-allocation, the fitness value for the allocation of computing engines to computing domains is re-calculated. If the re-calculated fitness value exceeds the previously-calculated fitness value, the re-
- 25 allocation is maintained. Otherwise, the re-allocation is discarded, and the computing engine is returned to its previous allocation state. The re-allocation process repeats either for a predetermined number of iterations, or until no computing domain has fewer than its expected number of allocated computing engines.

The entire allocation method is repeatedly applied at predetermined intervals.

BRIEF DESCRIPTION OF THE DRAWINGS

The foregoing and other features of the present invention will be more readily apparent from the following detailed description and drawings of illustrative embodiments of the invention wherein like reference numbers refer to similar elements throughout the several views and in which:

5 Figure 1A provides a schematic diagram illustrating an architecture for the present invention;

Figure 1B provides a schematic diagram illustrating an alternate view of architecture for the present invention;

10 Figure 1C provides a schematic diagram illustrating a third view of the architecture for the present invention;

Figures 2, 3A and 3B illustrate exemplary web pages produced by an administrative tool used in conjunction with the present invention;

Figure 4A provides a schematic diagram illustrating domain types supported by the present invention;

15 Figure 4B illustrates a virtual gateway for balancing traffic for web applications and services;

Figure 4C illustrates an exemplary web page of a domain wizard of the administrative interface that identifies container types;

20 Figure 4D illustrates an exemplary web page of the domain wizard for creating or editing a web application domain

Figure 5A illustrates an exemplary policy editor page of the administrative tool of Figure 2 for setting minimum and maximum values for engine allocations to domains and groups;

Figure 5B provides an example of time-of-day dependent grid allocations as controlled by a policy of the present invention;

Figure 5C illustrates an exemplary policy wizard web page that may be used to set performance-based resource allocation constraints;

Figure 5D provides an example distribution of engines allocated to operating domains.

Figure 6A provides a schematic diagram illustrating elements of an engine component
5 of the present invention;

Figure 6B provides flowcharts illustrating the steps for executing an engine lifecycle;

Figure 7A provides a schematic diagram illustrating elements implementing a client component of the present invention

Figure 7B provides a schematic diagram illustrating the creation of threads on an engine
10 as managed by the client of Figure 7A;

Figure 8A provides a schematic diagram illustrating broker-initiated provisioning of engines dedicated to clients;

Figures 8B and 8C provide schematic diagrams illustrating broker-initiated provisioning of engines shared by clients;

Figure 8D provides a schematic diagram illustrating client-initiated provisioning of
15 engines;

Figure 9A provides a schematic diagram illustrating how performance statistics are compiled by the broker;

Figure 9B provides a schematic diagram illustrating how performance statistics are
20 collected by a client or engine;

Figures 9C and 9D provide a sample lists of the statistics compiled by the broker in Figure 9A;

Figure 9E illustrates an exemplary “dashboard” web page of the administrative tool;

Figure 9F illustrates an exemplary web page reporting a measured statistic for an engine;

Figure 10 presents a flow diagram illustrating an adaptive provisioning process according to the present invention;

5 Figure 11A presents a flow diagram illustrating additional steps in the provisioning step of the adaptive provisioning process of Figure 10.

Figure 11B presents a flow diagram illustrating additional steps in the step for determining expected allocations in the provisioning process of Figure 11A;

10 Figures 11C - 11E list pseudocode for a preferred method for implementing the step for determining expected allocations of Figure 11A;

Figure 11F presents a flow diagram illustrating additional steps in the step for making resource assignments in the provisioning process of Figure 11A;

Figure 11G illustrates a sequence of resource assignments made according to predetermined priorities for the domains;

15 Figure 11H lists pseudocode for a preferred method for implementing the step for making resource assignments of Figure 11A;

Figure 11I presents a flow diagram illustrating additional steps in the step for improving resource assignments in the provisioning process of Figure 11A; and

20 Figure 11J lists pseudocode for a preferred method for improving resource assignments in the provisioning process of Figure 11A.

DETAILED DESCRIPTION OF THE INVENTION

The present invention is directed to an application virtualization and provisioning (AVP) platform that creates a highly adaptive shared computing infrastructure. Within this infrastructure,
25 application servers and other service-oriented components are hosted and virtualized on shared

computing resources, and are adaptively provisioned and activated in response to demand. The shared computing resources may be geographically localized, or may be distributed over a wide geographic area and managed as a computing grid. See Mark Baker et al., *Grids and Grid technologies for wide-area distributed computing*,” Softw. Pract. Exper., John Wiley & Sons, Ltd.,
5 2002, which is hereby incorporated by reference.

Application virtualization enables the removal of static, host-specific configuration dependence from the application environment by using the AVP platform to automatically and adaptively allocate, configure, start, deploy, monitor and manage software applications and services. Explicit usage policies are defined that guide provisioning and dynamic allocation of the shared
10 computing resources. The configuration of software applications and services is in particular enabled by a broker component of the AVP platform that stores applications servers, configurations and software applications in a central repository for deployment to the shared computing resources as required according to the resource allocation, and manages the allocation of the shared computing resources based on the usage policies. This component also includes a facility for deploying new
15 application code to the shared computing resources while existing applications are running in the currently-allocated configurations.

The AVP platform is directed to managing “domains,” which may comprise an enterprise business application, utility service or data source that is allowed a certain percentage of available resources at a specified time, based on the usage policy. Domains consist of the artifacts
20 that make up the application, service or data source (e.g., web archive (WAR) files making up a web application), and may be classified among three domain types: service domains, data domains and web domains.

Service domains are used to virtualize Java programming objects, including plain old Java objects (POJOs), Spring Beans and Enterprise Java Beans (EJBs), by turning them into
25 services. These virtualized services may then be accessed by Java clients via dynamic proxies.

Data domains provide data services such as databases or scalable caching services. Data domains may preferably be implemented, for example, as JBOSS cache and TANGOSOL’s COHERENCE cache.

Web domains provide web applications and services such as web servers, messaging brokers, message-driven services and other services that typically require multiple running instances at any given point in time. Web domains include collections of application services that are accessible over the Internet via communications based on the hypertext transfer protocol (http).

5 Domains can be launched or hosted on one or more application servers (or “containers”). Containers are effectively “sandboxes” for hosting the domains. Each container type is capable of hosting one or more domain type, and a given domain can be launched by any container that supports its type. For example, a JBOSS container can host web application, web service and EJB service domains. Other container types may include but are not necessarily limited to APACHE TOMCAT containers, CAUCHO RESIN containers, IBM WEBLOGIC containers,
10 and other generic containers supported by the AVP platform.

A service-level policy, or consistent set of rules, is applied to dictate the operation and division of computing resources. The service-level policy may be defined for example by software application and/or by user group, and may be conditioned on certain performance requirements
15 including but not necessarily limited to response time, throughput and minimum/maximum allocation of computing resources (“percentage of grid”). In accordance with the defined service-level policy, the AVP platform operates to provision and activate services according to demand for improved performance and utilization of resources.

A system-level architecture for the AVP platform 100 is illustrated in Figure 1A. The
20 architecture includes four fundamental elements: clients 10, engines 20 associated with domains 40, and a broker 30. Clients 10a and 10b are software applications that access and utilize the domains 40. Engines 20 are processes that provision and run software applications in the domains 40. The broker 30 is a software application that carries out policy-driven resource allocation (e.g., allocation of engines 20 to domains 40 and clients 10a) and performance monitoring. Each of the clients 10a,
25 engines 20 and broker 30 may be implemented on conventional INTEL and/or SUN/SPARC hardware platforms running, for example, WINDOWS, WINDOWS SERVER, SOLARIS, RED HAT Linux or RED HAT Enterprise Linux operating systems

As illustrated in Figure 1A, clients 10a and engines 30 both interact with the broker 30. The broker 30 assigns engines 20 to domains 40, and provides information for example to JAVA
30 clients 10a that instructs the clients how to access to the engines 20. Thereafter, JAVA clients 10a

are able to submit service requests directly to the connected engines 20. Http clients 10b submit service requests via a router ("Vgateway 31"), which acts as a virtual gateway and load balancer for directing the service requests to engines 20 running web service or web application domains.

Figure 1B illustrates an alternate view of the architecture for the AVP platform 100.

5 AVP platform 100 dynamically assigns and provisions computing resources 21 among software applications 41 supported by application servers 41a by configuring domains 42, 43 and 44. AVP platform 100 optimizes the assignment of resources 21 among the applications 41 subject to constraints 60 which may include, for example, service-level policies associated with the domains 42, 43, 44, and/or with user groups seeking access to the domains, service level agreements
10 ("SLAs") associated with the domains 42, 43, 44 and or user groups, performance statistics periodically collected from engines, clients and other components of the AVP platform 100, and service demands predicted from the usage statistics.

Figure 1C illustrates a third view of the architecture for the AVP platform 100.

Computing resources are represented by grid nodes 25, which may each include one or more host
15 computers. Broker 30 allocates and configures one or more engines 20 to run on each of the grid nodes 25. Each engine 20 manages a container 26 that serves as an environment for running an application, service or data source, and preferably collects and reports performance statistics for the application, service or data source (for example, by Java Management Extension (JMX) proxy for Java 2 Platform, Enterprise Edition (J2EE) applications), and preferably binds with a container
20 software development kit (SDK) within an administrative interface (not shown) that may be used to configure the containers 26.

Broker 30 also configures a daemon 22 that runs on each host computer in each grid node 26 that monitors the host, manages the engines 22 that are running on the host, and deploys binary code provided by the broker 30 for running a container (or applications server) 26 and/or an
25 application, service or data source to be run by the container 26. In addition, broker 30 collects performance statistics provided by the engines 20 (and/or by clients 10a and Vgateway 31) for storage in a database 39, for reporting and/or as inputs to the allocation optimization. Broker 30 may also provide failover services for reallocating an application, service or data source from a failed host computer to an operating host computer.

AVP platform 100 of Figures 1A, 1B and 1C further includes an administrative interface (not shown) of the broker 30 that enables a platform administrator to define, register and deploy domains, to manage workloads and to configure the AVP platform environment. By way of example, Figure 2 illustrates a broker web page of the administrative interface that provides access to a variety of wizards available for creating data, web and service domains, and for establishing policy.

In addition, the administrative interface allows the platform administrator to monitor and manage various performance metrics, including but not necessarily limited to throughput, latency, resource usage, and exceptions. For example, Figure 3A illustrates a “dashboard” page of the administrative interface that provides a pie chart indicating a current allocation of resources among domains, and Figure 3B illustrates a “broker monitor” page of the administrative interface that graphs the allocation of resources among domains over time.

Domains

As illustrated in Figure 4A, the AVP platform 100 is directed to manage three types of domains: service domains 45, web domains 46 and data domains 47.

Web Domains

Web domains 45 provide web applications and services, for example, including web servers, messaging brokers, message-driven services and other services that typically require multiple running instances. Web domains represent any collection of application services that are accessible via http, and can effectively represent any object or process that can be started, stopped and interrogated for its current load. Types of web domains include web applications accessible via a browser, and web services made available via simple object access protocol (SOAP) over http.

Web domains are preferably targeted for J2EE application servers. For example, an existing J2EE application may be represented as a web service, with an application cluster size varying between 2 and 5 nodes, base on policy. The physical locations of the web domain instances are decided by the AVP platform 100 at runtime, and the resources are provisioned dynamically. The AVP platform 100 then instantiates each domain on one or more grid nodes. The policy that

dictates how many instances are created, and at what time they are created, are dictated by a service-level policy that is maintained by the broker 30.

As illustrated in Figures 1A and 4B, web clients 10b may preferably access web domains 40 via a virtual gateway router (Vgateway 31). VGateway 31 is preferably implemented as part of the broker 30, and functions essentially as a smart load balancer, routing web service requests and responses from external clients to resource virtualized engines. Unlike conventional static load balancers, Vgateway 31 is informed when the configuration of host computers 23 and/or domains 40 changes, and adjusts its load balancing scheme accordingly.

Service Domains

Service domains 46 include a collection of interfaces that can be virtualized across distributed computing resources ("grid resources"). By grouping these resources within a service domain, specific policies can be set to determine how many resources each service domain will be allowed to consume. JAVA-based service domains 45 may be defined, for example, using J2EE, plain old JAVA objects ("POJOs") or the Spring framework

Service domains 45 may be used, for example, to define any standard JAVA class or Enterprise Java Bean (EJB). No proprietary application programming interface (API) or class format is required to virtualize the associated Java service. For example, POJOs can be defined with application context, or the complete and necessary environment for making the associated Java object instance work correctly. For example, a JAVA class that represents a business service would be defined with access to database connections and messaging services in order to perform the required processing. Preferably, a simple Extensible Markup Language (XML) format is provided for declaring the proper context for each object.

Among supported service domain types, the POJO domain type is the simplest to construct. Any JAVA class can be included in a POJO service domain. In addition, a Spring service domain type may preferably be supported. See Rod Johnson, *Introduction to the Spring Framework*, May 2005, available at www.theserverside.com/articles/article.tss?l=SpringFramework, which is hereby incorporated by reference. The Spring framework simplifies J2EE development by using POJOs instead of EJBs, and allowing for the abstraction and encapsulation of implementation dependent components (for example, Hibernate and JDBC mapping tools). In addition, this

framework allows for dynamic proxy-based aspect oriented programming (AOP). AOP is a programming facility that allows developers the ability to inject logging, transaction, security and transaction capabilities across modules and components. The Spring framework also uses AOP to provided declarative transaction management for POJOs. For legacy components that are packaged
5 as EJBs, an EJB service domain allows for virtualized access to EJB functionality.

Data Domains

Data domains 47 of Figure 4A provide data services such as databases or scalable caching services. These services are essentially clientless, as no gateway or proxy to the services in provided via the broker. Instead, the AVP platform may provide a centralized lookup, for example,
10 such as a Java Naming and Directory Interface (JNDI) that allows clients to discover and connect to these domains.

Creating Domains

According to the principles of the present invention, domains are launched or hosted on one or more applications servers, or containers. Each container type is capable of hosting one or
15 more domain types, and a given domain can be launched by any container that supports its type. Figure 4C provides an exemplary listing of container types supported by a domain wizard of the administrative interface. For example, a JBOSS container can support web application, web service and EJB service domains. Other container types include but are not limited to APACHE TOMCAT containers, CAUCHO RESIN containers, IBM WEBLOGIC containers, and other generic
20 containers supported by the AVP platform. The administrative interface preferably includes a container software development kit (SDK) enabling the addition of additional container types as required.

Domains may be created for example by means of a domain wizard within the administrative interface. Figure 4D illustrates an exemplary web page of the domain wizard for
25 creating or editing a web application domain that deploys a specified web application. As illustrated in Figure 4D, the web application domain may be newly created or modified by selecting and specifying an appropriate container, and by specifying domain settings, associated archive files (for example, JAVA archive (JAR), Enterprise archive (EAR) or web archive (WAR) files), servlets and

Enterprise JAVABEANS (EJBs). In addition, tracked statistics for associated service-level policies may be specified. For web applications and web services, URL patterns to be used by the Vgateway 31 of Figure 1A may also be specified.

Policy and Resource Allocation

5 The broker 30 of Figure 1A is configured with a service-level policy that provides a consistent set of rules for the assignment and allocation of computing resources from the resource grid. This policy enables the broker 30 to select a resource allocation among the domains. In the absence of this policy, the broker may operate to assign an equal percentage of grid resources to each of the domains.

10 The service-level policy defines a hierarchical, constraint-based division of computing resources in the grid. For example, a first level of partitioning may be by domain, followed by a second level of partitioning by user group. Additional and/or alternate partitioning criteria may be arbitrarily selected (for example, partitioning by work project), all of which are fully contemplated within the scope of the present invention.

15 The service-level policy will generally define a minimum number and a maximum engines that should be allocated for each domain, either in terms of a number of engines or a percentage of available engines. A “minimum allocation percent” specifies a least amount of resource always held by an associated domain. If no clients are running in a domain, the rule may be excepted in order to make resources available to other grid clients (i.e., the “minimum allocation
20 percent” is set to zero, so that no resources are assigned unless no other clients are running in other domains). However, these resources are relinquished as soon as a non-running client starts up in the affected domain. If the minimum allocation percents for grid clients do not sum to 100 %, and all types of grid clients are active, the broker may continuously redistribute resources to optimize service-level agreements (SLAs) for the grid clients or domain.

25 A “maximum allocation percent” specifies a cap on the amount of resources to be given to an associated domain. This maximum is preferably enforced even if idle resources are available. As illustrated in Figure 5A, the administrative interface preferably provides an editor for editing the minimum and maximum engine allocations for domains. In addition, as shown in Figure 5A, the

editor provides a means for specifying a priority ranking for each domain that is used by the broker
30 for deciding a sequence in which unallocated resources are allocated among the domains.

Additional rules may further limit or change the engine allocations performed by the broker. For example, rules may be administered through the administrative interface of the broker
5 30 to limit which engines may be assigned to individual domains. The engines may be limited, for example, according to engine identifier, engine configuration, engine memory characteristics (e.g., amount of free memory), number of central processing units (CPUs) supported by the engine, and engine operating system (OS).

Allocation of engines may also be influenced by service level agreements (SLAs). SLA
10 targets may be set through the administrative interface to provide a means for changing engine allocation based on the performance of domain. For example, if a domain requires a particular level of network throughput, an associated SLA target can be attached to the domain and monitored. If throughput falls below the target level, allocation of engines to the domain can change according to associated rules. These additional rules may also include assigned priorities, which will essentially
15 dictate the order in which the rules are processed.

When an engine fails to meet a target threshold for a rule, the broker 30 may perform an additional statistical analysis to determine whether a reallocation is warranted. For example, the broker may calculate a normalized variance of the performance of all engines in an associated domain (for example, a normalized geometric variance of a performance statistic calculated as the
20 ratio of the geometric mean to the arithmetic mean). In the case of the normalized geometric variance, a value close to 1.0 indicates little variance in performance among the engines in the domain, providing strong evidence of the need for additional computing resources. Conversely, a value near 0.0 indicates a wide variance (and lack of balance) in performance among the engines. In this case, instead of immediately reassigning resources, the broker waits to see if balance is
25 achieved, and thereby if underutilized resources enable thresholds to be met without further allocation of resources.

The broker may in addition apply a policy schedule that indicates how policies are to be change over the course of a day. As illustrated for example in Figure 5B, the grid resources assigned to a domain 48 for an application varies with time. Domain 48a at 8:15 AM is allocated four
30 computing engines 20 from the grid. At 8:30 AM, the number of allocated engines in domain 48b is

reduced to three engines 20. At 8:45 AM, the number of engines allocated to domain 48c is increased to five engines 20.

Once the minimum and maximum number of engines is established, the broker 30 proceeds to provision resources to a minimum level. The broker 30 may choose to allocate more than the minimum number of engines to a domain if there is a statistical rule that can be used to understand the performance of the application. For example, if queue size can be measured as an indication of current load on the application, a rule can be established for adding an additional engine when the average queue size for all engines in the domain over a specified time interval exceeds a specified level. In addition, a rule for relinquishing an engine can be established based on the average queue size falling below a specified level over the specified time period. Figure 5C illustrates a policy wizard web page of the administrative interface that may be used for setting statistical rule-based constraints.

The broker 30 allocates engines (or application server instances) to domains based on a current active policy and current application performance statistics. The broker reevaluates and reapplies this policy periodically (for example, every 5 minutes), and then decides to take one of three courses of action: a) to make no change in the current allocation, b) to assign available engines to some domains, or c) to re-assign some engines from some domains to other domains. Figure 5D illustrates an allocation of engines across domains.

Engines

As illustrated in Figures 1C and 6A, each engine in the AVP platform 100 manages a container 30 to host and run a domain. Further, as illustrated for example in Figure 6A, an engine service instance 21 is managed by an engine daemon 22, both installed on a host computer 23.

Engines create service instances on demand, based on scheduling decisions made by the broker. A service is created with the first client request for an operation having the created service type. After creating and running the requested operation, the engine stores the newly-created service in a cache. A scheduler is made aware of the contents of the cache, such that it will route other requests for that service to the engine.

By default, engines operate as single-threaded processes (“engine instances”) performing only one service operation at a given time. As a result, more than one engine instance 21 is generally running at one time on the host computer 23. Processes running on multiple engine instances are started and managed by an agent that also runs on the host (engine daemon 22).

5 Engine daemon 22 is capable of starting and stopping engines based on a pre-defined engine policy. Engine policies may for example be based on one or more of CPU utilization of the host, user activity (in the case that the host is a user’s desktop) or time of day. In most cases, the engine daemon 22 starts and monitors engine instances 21, and restarts the engine instances 21 in response to failures or reconfigurations.

10 One engine daemon 22 runs per host. In addition to starting engine instances 21, the engine daemon 22 preferably controls the configuration of engine instances 21. For example, when changes to an engine configuration are made by a platform administrator (for example, to configure a new application server), the changes may be automatically propagated to an engine instance 21 via the engine daemon 22. Engine daemons 22 may log into the broker 30 for administration and
15 configuration.

Engine instances 21 are the processes that perform tasks for executing application software in the domain. On multi-CPU hosts, an engine daemon 22 will be able to run multiple engine instances 21. In addition, more than one engine instance 21 may be run on a single CPU.

Engines 20 report to the broker 30 when they are available to perform work. After
20 logging in and synchronizing resources, the engines accept work assignments, perform tasks for executing the applications software, and notify the broker 30 when results are ready. Because the engine daemon 22 controls the state of configuration for each engine instance 21, and engine configuration can be controlled centrally via the administrative interface of the broker, it is easy to control and configure engine instances across the computing resource grid.

25 Engines can be configured to run in a variety of modes, depending upon the type of host machines 23 on which they will be installed. Dedicated machines are configured to run continuously, and are best suited for computing resources devoted to full-time processing on the grid. A non-dedicated mode may be enabled for host machines that are only used on a part-time basis on the

grid, and otherwise used for other purposes (for example, user PCs sometimes made unavailable to the grid for user process use).

Engines configured in the non-dedicated mode determine when to run based on two different modes. In the user interface (UI) idle mode, a non-dedicated engine will start running after user inactivity on the host machine. Alternatively, in CPU idle mode, the engine will start to run when CPU utilization is sufficiently low. Engines are installed only once on a host machine. As engines are centrally managed by an engine daemon 22, they can be easily upgraded when later versions to the AVP platform 100 are available by using the administrative interface. In addition, to gain additional efficiencies, configuration profiles may be created by the administrative interface which may be used by multiple host machines to synchronize configurations.

Figure 6B provides a flow diagram illustrating steps in the lifecycle of an engine. At step 601, the engine daemon 22 determines that an engine instance should be running on the host based on a state of the host 23 and an engine mode of the host (for example, if the engine is non-dedicated, an engine instance may be created only if no other user processes are currently running on the host 23). At step 602, the engine instance 21 established a connection to the broker 30 to identify to the broker 30 that the instance 21 is ready for work. At step 603, the broker 30 provisions the engine instance 21 to a domain.

At step 604, a client, having received information relating to the engine instance 21 and its associated domain from the broker 30, connects to the engine instance to run a service. At step 605, when the service has completed, the engine instance establishes another connection to the broker 30 to indicate that it has completed the service and to request another assignment.

At step 607, if the engine instance 21 is interrupted or otherwise fails gracefully, it connects to the broker 30 to send a message indicating that it has logged out. Otherwise, if the engine instance 21 fails unexpectedly, an engine monitor of the broker will log the engine instance off. In either case, if available, the broker will provision another engine instance to the associated domain to replace the failed instance.

Clients

As illustrated in Figure 1A, requests for access to service domains 40 may be forwarded to the broker 30 by Java clients 10a. The clients 10a for example may make a request to invoke a method for processing in a service domain using simple application programming interface (API) access. In the case of web domains, a web client 10b (for example, a web browser or other http client) may access a web application or web service via Vgateway 31. In this case, the client simply opens a uniform resource locator (URL) that is directed to Vgateway 31, and configured to run the selected application, virtualized on a web domain.

As illustrated for example in Figure 7A, a client 10 synchronously invokes a method for processing in a service domain by sending an invocation message including service/method call information and a wait lock to a corresponding service domain 11. The service domain 11 adds the message to an invocation queue 12. A thread running on the engine 20 is then blocked by the service domain 11 using the wait lock. The process for asynchronous invocation is similar, except a result listener message is sent in the invocation message, indicating that a listening process will be created by the client and wait until the engine 20 indicates that the task has been completed.

Communications between the engine 20 and the client 10 are managed by an engine proxy 13. As illustrated in Figure 7B, the engine proxy 13 creates a new thread 14 for each thread 24 that has been started on the engine 20. With these threads, the proxy 13 continuously asks for a next invocation process. The threads 14 will block if the queue 12 is empty. If the proxy 13 fails to process an invocation, it notifies the queue manager 16, which places the unprocessed invocation back in the queue 12.

Each client 10 has a performance monitor (not shown) for monitoring call requests and keeping statistics on recent history. Request statistics monitored by the client 10 preferably include but are not necessarily limited to total response time, time in the queue 12, time in transport and time in user code (i.e., direct application or service processing time). The client monitor calculates average statistics for each of the measured statistics, as well as average throughput. The average statistics are periodically sent by the client 10 to the broker 30, as further described herein.

Broker

The broker 30 provides policy-driven resource allocation and monitoring for the AVP platform 100. Specific broker tasks include message routing and client authentication, engine

allocation, engine provisioning and resource management, performance monitoring, and application and application server code distribution. Engines and clients are able to log in to the broker 30 in support of these tasks.

Figure 8A schematically illustrates a broker-driven process by which engines are
5 allocated to domains. Client 10 periodically sends average statistics to broker 30, which are received by statistics manager 33 and placed in client and engine statistics queues 34. Engine allocator 32 scans client and engine statistics queues at regular intervals, applies policy-based optimization algorithm 37, and decides either to make no changes to the current allocation of engines to clients 10, to assign currently available engines from engine pool 35 to some of the clients 10, and/or to re-
10 assign some engines previously assigned to clients 10 to other clients. Clients 10 are provided access to engines in engine pool 35 with the delivery of associated engine proxies 36 from the broker 30 to the clients 10.

The broker 30 provisions engines according to the service policy based on the operational mode of the broker, allocation policies and client activity. Schemas include “broker-
15 initiated” provisioning and “client-initiated” provisioning. Broker-based provisioning is useful for controlling engine allocation across domains, and is required to enable engine sharing.

As illustrated in Figure 8B, broker-based provisioning begins with a service domain-specific request transmitted by a client 10 to the broker 30. In response, the broker provides the client with an engine proxy that is already assigned to a specific service domain. With broker-based
20 provisioning, a client may not directly ask the engine to change the assigned domain.

Two kinds of engine allocation are supported by broker-based provisioning. With exclusive allocation, as illustrated in Figure 8B, engines are assigned with the delivery of associated engine proxies 36 to clients 10 such that each engine 20 provisioned in a domain 40 is assigned to perform work for exactly one client 10. With shared allocation, as illustrated in Figure 8C, two or
25 more clients 10a, 10a' may respectively use shared engine proxies 36a, 36b to send requests to the same engine 20 in domain 40.

As illustrated in Figure 8D, under the client-initiated provisioning schema, clients 10 receive “blank slate” engine proxies 36c, and are able to provision them with service domains of their choice. A service domain-independent request is first transmitted by the client 10 to the broker

30. In response, the broker 30 provides the client with unassigned proxy 36c, allowing the client to activate a service domain of its choice via engine 20. Under this schema, no sharing of engines is possible.

The broker performs a number of additional functions on behalf of the AVP platform 100. For example, the broker configures and stores a variety of operational settings and parameters, including but not necessarily limited to user identification, passwords, client information, routing properties and engine configuration. Using this stored data, for example, associated tasks may be carried out by platform administrators via the administrative interface of the broker 30.

An internal database of the broker stores reporting data, including for example user, engine, client and broker information, and an external reporting database is used to log events and performance statistics. Associated database configurations may be managed by platform administrators via the administrative interface.

Domain resources are staged on the broker 30, for deployment to engines. For example, files may be uploaded to deploy service, web and data domains using the domain wizard component of the administrative interface as illustrated in Figure 2. Domains can be deployed at the time of uploading, or can be deployed or undeployed at a later time.

Monitoring and Statistics Tracking

The broker 30 carries out a performance monitoring function with the assistance of the clients 10 and engines 20. Figure 9A schematically illustrates how the this function is performed.

At regular intervals, the broker 30 asks at least one of each client 10 and engine 20 associated with service and web domains (preferably, at least each engine 20) to collect and forward averaged performance statistics. The information is collected and forwarded by at least one of a statistics collector 10c of the client 10 and a statistics collector 20c of the engine 20 to the statistics manager 33 of the broker 30. This process may for example be facilitated by means of a JMX proxy for clients and engines running J2EE applications.

Figure 9B further illustrates schematically how statistics are collected by the client 10 and engine 20. The statistics collectors 10c, 20c of the client 10 and engine 20 hold a collection of statistics providers 60. At regular intervals, the statistics collector 10c, 20c asks each provider 60 to format its latest average statistics into a common statistics record 61, and forwards the common
5 statistics records 61 to the statistics manager 33 of the broker 30 as illustrated in Figure 9A. The forwarded information includes an invoking group, domain, service and method “signature,” as well as the average values of collected statistics.

The statistics manager 33 places the forwarded information in client and engine statistics queues 34 of Figure 9A. Periodically (for example, hourly), statistics persister 38
10 consolidates the collected data by averaging the accumulated data for each client 10 and engine 20, calculating statistics for the entire computing resource grid, and storing the calculated statistics in grid usage statistics database 39. Additional averaging and archiving is preferably performed periodically on the database 39 to further consolidate the data. The accumulated data in the database 39 may be displayed on a reporting screen 50 via the AVP platform administrative interface.

15 A sample list of statistics tracked is provided in Figures 9C and 9D. Statistics used will vary according to application. For example, load on a JAVA application serve may be assessed by statistics such as queue depth, service throughput or thread count rather than user response time.

With frequent collection of statistics from each client 10 and engine 20, large amounts of statistical data accumulate. Accordingly, at frequent intervals, the broker operates to average the
20 data collected for each client and engine, to calculate statistics for the entire grid, and to save the resulting records in the broker databases. Archiving may be performed after successive intervals, using similar averaging methods.

The collected statistics may be viewed in a variety of ways via tracking tools in the administrative interface of the broker 30. As illustrated in Figure 9E, for example, the administrative
25 interface may include a dashboard for displaying a variety of statistical data in summary form. The dashboard may provide, for example, pie charts indicating domain and group allocations of engines, measured statistics for the clients and engines, and platform alerts generated according to a comparison of the measured statistics to service level agreements (SLAs) defined in the service-level policies. In addition, the dashboard may provide links for viewing domains, wizards for

configuring various components of the platform, and links to other frequently used pages of the administrative interface.

For example, as illustrated in Figure 9F, a web page of the administrative interface illustrates a thread count over time for a selected engine.

5 Adaptive Provisioning

Figure 10 summarizes the adaptive provisioning process according to the present invention. As described above, and as illustrated at step 1010 of Figure 10, a service-level policy is defined and stored in a database 39 accessible to the broker 30. The policy includes minimum and maximum resource levels to be allocated, for example, to a domain or user group, by time of day.

10 The policy may also define priorities to be applied in the event of resource contention, and service-level policies relating to measured statistics for the system.

At step 1020, the resources are provisioned according to the policy. Engines are assigned to domains by the broker 30, and configured by downloading and installing associated application server and application software in the engines. Once configured, engine instances are
15 also started in response to the receipt of service requests, and stopped upon task completion.

At step 1030, the broker 30 periodically collects averaged performance statistics from one or more of the clients and the engines, and compares the averaged statistics with SLAs 1035 defined in service-level policies. The statistics may provide, for example, measures of throughput, response time, CPU occupancy, memory occupancy and other attributes that may preferably be
20 defined as JMX attributes. The policy is then again re-applied at step 1010 and the resources are re-allocated at step 1020. In addition, at step 1040, alerts indicating violation of the SLAs may preferably be reported to administrators via a “dashboard” of the administrative interface of the broker 30.

The reallocation step 1020 is further described in an exemplary manner with reference
25 to Figures 11A - 11J. In Figure 11A, reallocation process 1100 begins at step 1110 with the creation of an empty resource allocation map. The resource allocation map is a data structure which is used, for example, by the broker 30 for recording assignments of resources (e.g., computing engines 20) to resource users (e.g., domains 40) during the reallocation process 1100. At step 1130,

the broker 30 determines a number of expected or desired resource for each domain 20 based on a predetermined policy for the domain, SLAs for the domain, and current performance statistics reported to the broker 30 by one or more of engines 20 and clients 10 assigned to the domain 40. The predetermined policy for the domain, for example, may specify a minimum and maximum
5 number of engines 20 to be assigned to the domain 40, a priority for the domain 20, and other associated rules. The other rules may constrain the types of engines that may be assigned to the domain 40 (for example, according to type of operating system), and further specify rules pertaining to the SLAs.

Once the expected number of resources for each domain is determined at step 1130 and
10 recorded in the allocation map, the reallocation process 1100 proceeds at step 1150 to prepare a set of resource assignments assigning specific computing engines 20 to domains 40. Assignments of computing engines 20 to domains 40 are made in an order reflective of the priorities specified for the domains 40, and are then recorded in the allocation map.

At step 1170, a current fitness value is calculated for each assignment of computing
15 engines 20 to a domain 40, and by means of one or more conventional constraint programming methods, the assignments made at step 1150 are adjusted, fitness values are recalculated for the adjusted assignments, and the adjusted assignments are maintained and recorded in the allocation map when the fitness values for the adjusted assignments improve over the fitness values for the assignments made at step 1150.

At step 1190, based on the assignment information currently in the allocation map, the
20 broker 30 proceeds to release and re-acquire computing engines 20 that are to be reassigned to among the domains 40 in order to provision the computing engines 20 to the domains 40. At step 1191, after a predetermined time period (for example, in the range of three to five seconds), the reallocation process 1100 resumes at step 1110 by preparing an empty allocation map to be used in
25 the subsequent reallocation steps.

Figure 11B further depicts the process at step 1130 for determining an expected
allocation of computing engines 20 among domains 40. At step 1132, an allocation entry is created in the allocation map for each of the domains 40, and a minimum number of engines for the domain (established according to the domain policy) is initially set to be the expected allocation for the
30 domain and recorded in the allocation map. At step 1134, the minimum number of engines for each

domain is compared to an expected number of engines that was calculated in a prior allocation cycle according to the process at step 1130. If a prior expected number of engines for a domain exceeds the minimum number of engines, the expected number of engines is set to be the prior expected number of engines, and re-recorded in the allocation map.

5 At step 1136, a current demand is computed for each domain, for example, as a function of SLAs defined in the policy for the domain and performance statistics collected for the domain (see "Policy and Resource Allocation," *supra*). If there is a positive demand (i.e., the statistics indicate that performance falls below the thresholds established by the SLAs), and the demand exhibits stability (i.e., as may be indicated according to a normalized variance in demand among
10 engines 20 currently assigned to the domain 40), the expected number of engines is increased. For example, where a stable, positive demand is indicated, the expected number of engines may be increased by a fixed number (i.e., by one engine). At step 1138, the expected number of engines is finally set to be the lesser of a) the expected number of engines (as may have been increased due to demand) and b) the maximum number of engines (which has been set according to the domain
15 policy). This final expected number of engines is recorded in the allocation map at step 1140. A pseudocode representation of the process at step 1132 is illustrated in Figure 11C, and a pseudocode representation of the process at steps 1134 - 1138 is illustrated in Figure 11D.

Figure 11E presents a pseudocode representation for a preferred method of generating the normalized variance indicative of a stability of demand for a domain 40. As indicated in Figure
20 11E, a normalized geometric variance (GV) is calculated as a function of a value of a performance statistic for each of n engines 20 assigned to a domain D. Specifically, the normalized geometric variance is determined as the ratio of the geometric mean for values of the performance statistic divided by the arithmetic mean of the performance statistic values. The computed normalized geometric variance is then compared to a predetermined threshold value for determining whether
25 the demand as viewed by all engines 20 in the domain D is sufficiently balanced and stable to warrant an increase in the expected number of resources. For example, values of GV that are close to 1.0 indicate a well-balanced behavior, and values of GV that are close to 0.0 indicate a lack of balance. Applicants have determined that a threshold of approximately 0.85 is particularly suitable for use in the present invention.

Figure 11F further depicts the process at step 1150 for making the initial assignment of computing engines 20 to domains 40. Each of the domains 40 is assigned a priority value according to the domain policies. At step 1151, a maximum priority value (for example, from 1 to 9) is determined among domains 40 seeking resource assignments. As a result of steps 1152 and 1153, the assignment process begins for a priority value i equal to N . At step 1156, each domain 40 having a priority value equal to or greater than i and having fewer than its expected number of resources assigned is identified.

At step 1157, for each identified domain 40, a selection rule is determined. For example, the selection rule may only allow selecting the engine 20 that is currently assigned to a particular allocation map entry for the domain 40 ("same entry selection"). Alternatively, the selection rule may allow the selection of an engine 20 that is already assigned to the domain ("same domain selection"), may allow the selection of an engine so long as it is not required by another domain ("rule selection"), or may allow the selection of any available engine ("always selection"). Based on the policy for the identified domain, a single selection rule may be identified for the domain, or more than one selection rule may be selected in a predefined sequence.

At step 1158, it is determined whether an unassigned engine is available for the identified domain 40 that meets the current selection rule. If so, the available engine 20 is assigned to the identified domain 40 at step 1159. If no eligible engine is available, the process proceeds to step 1160 to determine whether the selection round has been completed. If not completed, the process returns to step 1156 to identify a next domain 40 in the selection sequence that is looking for an available unassigned computing engine 20. Once the current selection round is completed, the process returns to step 1153 to decrement the priority i . At step 1154, if it is determined that i is zero, the process terminates at step 1161. Otherwise, the process begins a next selection round.

Figure 11G depicts a sequence of nine selection rounds for domains A - E, the domains A - E having priority values ranging from 9 to 1. Figure 11H presents a pseudocode representation for a preferred method for carrying out the sequence of selecting computing engines 20 to domains 40.

Figure 11I further depicts the process at step 1170 for improving the selected assignments of computing engines 20 to domains 40. At step 1171, a fitness value is computed for the allocation map. The fitness value may be determined using one of a variety of conventional

objective functions having the property of providing either an increased or a reduced value when the number of engines 20 assigned to an allocation table entry for a domain is improved in relation to the expected number of engines 20 for the allocation table entry. For example, the objective function may assign a value of 3 to an allocation table entry having the expected number of engines 20, a value of 2 to an entry that has more than the minimum number of engines 20 but less than the expected number of engines 20, and a value of 1 to an entry that has less than the minimum number of engines 20. The objection function may alternatively or in addition provide an indication of improvement when a match of preferred engine types assigned to an allocation table entry is improved.

At step 1172, a number of iterations N is selected for carrying out the improvement process of step 1170. Alternatively, the process may continue until a performance objective is reached (for example, reaching no more than M unsatisfied entries in the allocation map).

At step 1174, a domain having an unsatisfied allocation table entry is identified. This may be accomplished by selecting entries in the table in a pre-established order, by random selection, or by some other scheme (for example, as a function of fitness for the entry, time since the entry was last selected, and the like). Once the unsatisfied allocation table entry is identified, an engine 20 that is not currently assigned to the identified allocation table entry is identified and selected at step 1175, and is assigned to the allocation table entry at step 1176. The engine 20 assigned to the allocation table entry may preferably be selected from a set of unassigned engines 20, or alternatively, be selected to be reassigned from another allocation table entry, for example, having an allocation of engines 20 that currently satisfies its expected number of allocations. In the latter case, a conventional constraint programming-based method may preferably be employed to efficiently identify engines for re-allocation.

At step 1177, the fitness value for the allocation table is recomputed in view of the reassigned engine 20. At step 1178, it is determined whether or not the recomputed fitness value is greater than the current fitness value. If the recomputed value is greater, at step 1179, the reassignment of engine 20 is maintained, the fitness value is set equal to the recomputed fitness value, N is decremented by 1 and the process returns to step 1173 to determine whether N is greater than 0. If the recomputed fitness value is less than or equal to the current fitness value, at step 1180, the reassignment of engine 20 is discarded (i.e., the current assignment of engine 20 is maintained),

N is decremented by 1 and the process returns to step 1173 to determine whether N is greater than 0. At step 1173, if N is equal to zero, the process terminates at step 1181. Figure 11J presents a pseudocode representation for a preferred method for the process at step 1170.

Thus, while there have been shown, described, and pointed out fundamental novel
5 features of the invention as applied to a preferred embodiment thereof, it will be understood that various omissions, substitutions, and changes in the form and details of the devices illustrated, and in their operation, may be made by those skilled in the art without departing from the spirit and scope of the invention. For example, it is expressly intended that all combinations of those elements
10 and/or steps which perform substantially the same function, in substantially the same way, to achieve the same results are within the scope of the invention. Substitutions of elements from one described embodiment to another are also fully intended and contemplated. It is also to be understood that the drawings are not necessarily drawn to scale, but that they are merely conceptual in nature. It is the intention, therefore, to be limited only as indicated by the scope of the claims
appended hereto.

15 All references, publications, pending and issued patents are herein each incorporated by reference in their entirety.

CLAIMS

What is claimed as new and desired to be protected by Letters Patent of the United States is:

1. In a system for provisioning a shared computing infrastructure including a plurality of
5 computing engines, a plurality of applications servers, a plurality of computing domains and a
broker, the shared computing infrastructure supporting a plurality of software applications and the
plurality of applications servers including a plurality of applications server types, each application
server type capable of providing a run-time environment for executing at least one of the plurality of
software applications; each computing domain being hosted by an application server of a given
10 type, and the broker being operative to dynamically allocate the plurality of computing engines
among the plurality of computing domains based on allocation policies provided to the broker, a
method for specifying the allocation policy for one of the plurality of domains via an administrative
interface to the broker, the method comprising the steps of:

specifying a desired minimum number of computing engines to be allocated to the one
15 domain in any broker allocation;

specifying a desired maximum number of computing engines to be allocated to the one
domain in any broker allocation; and

specifying a priority of allocation for the one domain.

20 2. The method of claim 1, further comprising the step of:

specifying a service level agreement (SLA) target, wherein performance of the one domain
against the SLA target may be judged by the broker in determining whether to allocate one or more
additional computing engines to the one domain.

25 3. The method of claim 1, further comprising the step of:

specifying one or more of the plurality of computing engines which may be allocated to the
one domain.

4. The method of claim 3, wherein the one or more computing engines are each specified by
30 an engine identification number or engine configuration information, the engine configuration

information including one or more of a free memory size, the number of central processing units (CPUs) or the engine operating system.

5. The method of claim 1, further comprising the step of:

5 specifying a daily time interval for the broker to apply the policy to the allocation of computing engines to the one domain.

6. In a system for provisioning a shared computing infrastructure including a plurality of computing engines, a plurality of applications servers, a plurality of computing domains and a broker, the shared computing infrastructure supporting a plurality of software applications and the plurality of applications servers including a plurality of applications server types, each application server type capable of providing a run-time environment for executing at least one of the plurality of software applications; each computing domain being hosted by an application server of a given type, a method performed by the broker for allocating the plurality of computing engines among the plurality of computing domains based on an allocation policy, the method comprising the steps of:

15 determining an expected number of computing engines to be allocated to each of the plurality of computing domains, wherein the expected number for each computing domain is determined as a function of a predetermined service policy of the computing domain and a predicted demand for the computing domain; and

20 while one or more of the computing engines is unallocated and one or more of the plurality of computing domains has been allocated fewer than its expected number of computing resources, sequentially selecting a next computing domain as a function of predetermined priorities for the computing domains, selecting an unallocated computing engine, and allocating the selected computing engine to the selected computing domain.

7. The method of claim 6, further comprising the steps of:

calculating a fitness value representative of the current allocation of computing engines to computing domains;

25 identifying a first computing domain for which fewer than the expected number of computing engines has been allocated;

30 selecting one of an unallocated computing engine or a computing engine allocated to a second computing domain for which the expected number of computing engines has been allocated;

re-allocating the selected computing engine to the first computing domain;
re-calculating a fitness value representative of allocation of computing engines to computing domains after re-allocation; and

discarding the re-allocation of the selected computing engine to the first computing domain
5 if the fitness value for the current allocation exceeds the fitness value for the re-allocation.

8. The method according to claim 7, wherein the selecting, re-allocating, re-calculating and discarding steps are repeated for a predetermined number of iterations.

10 9. The method of claim 7, wherein the selection of an unallocated computing engine or a computing engine allocated to a second computing domain is made randomly.

10. The method of claim 7, wherein the selection of an unallocated computing engine or a computing engine allocated to a second computing domain is made by application of a constraint
15 programming search algorithm.

11. The method of claim 6, wherein the predetermined service policy for each computing domain specifies a minimum number of computing engines to be allocated to the computing domain, a maximum number of computing engines to be allocated to the computing domain, and an
20 allocation priority for allocating computing engines to the computing domain.

12. The method of claim 11, wherein the predetermined service policy further specifies a service level agreement (SLA) target for judging the performance of the computing domain.

25 13. The method of claim 12, wherein the step of determining an expected number of computing engines for each domain further includes the steps of:

setting the expected number of computing engines equal to the minimum number of computing engines for the computing domain;

initially adjusting the expected number of computing engines to equal to a current number of
30 computing engines allocated to the computing domain when the current number of allocated computing engines exceeds the minimum number of computing engines;

determining whether the computing domain has a current unmet demand;

further adjusting the expected number of computing engines when there is a current unmet demand; and

finally adjusting the expected number of computing engines to equal a lesser of the maximum number of computing engines to be allocated to the computing domain and the further
5 adjusted expected number of computing engines.

14. The method of claim 13, wherein the current unmet demand is determined as a function of a statistical performance measure for the domain and an associated service level agreement (SLA).
10

15. The method of claim 14, wherein the current unmet demand is validated as a function of a variability of the statistical performance measure among computing engines in the domain.

16. The method of claim 15, wherein the variability is determined as a normalized geometric
15 variance computed as a geometric mean value of the statistical performance measure among computing engines in the domain divided by an arithmetic mean of the statistical performance measure among computing engines in the domain..

17. The method of claim 16, wherein the current unmet demand is validated when the
20 statistical performance measure falls below a first threshold based on the SLA, and the normalized geometric variance exceeds a second threshold.

18. The method of claim 17, wherein the second threshold is 0.85.

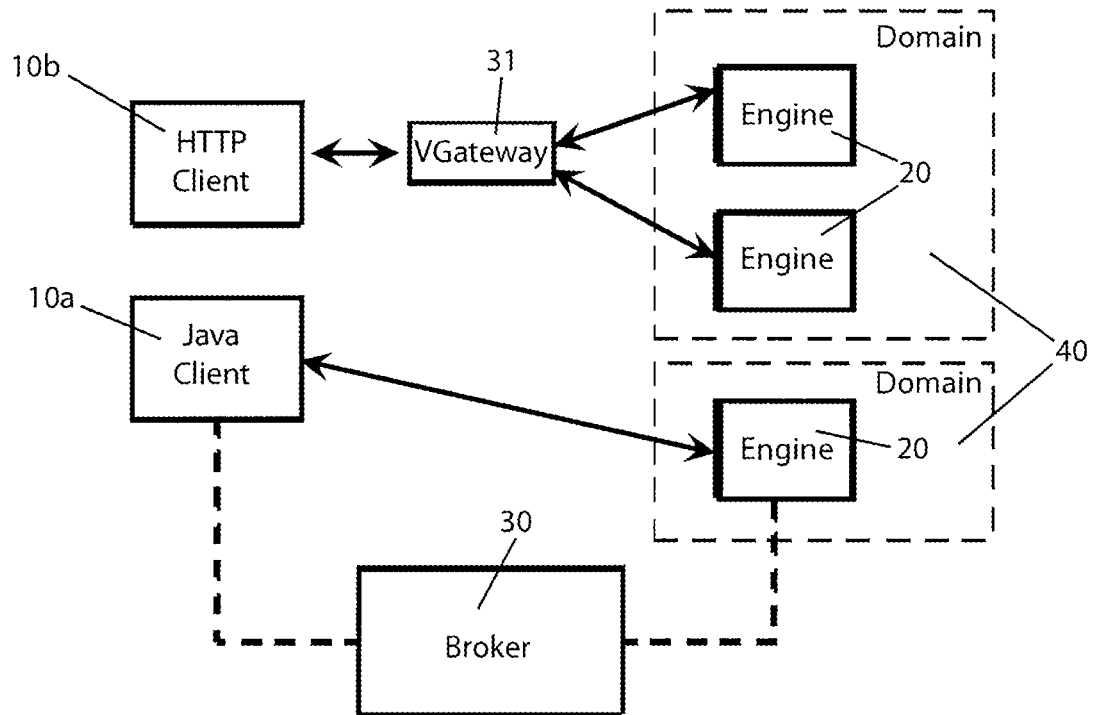
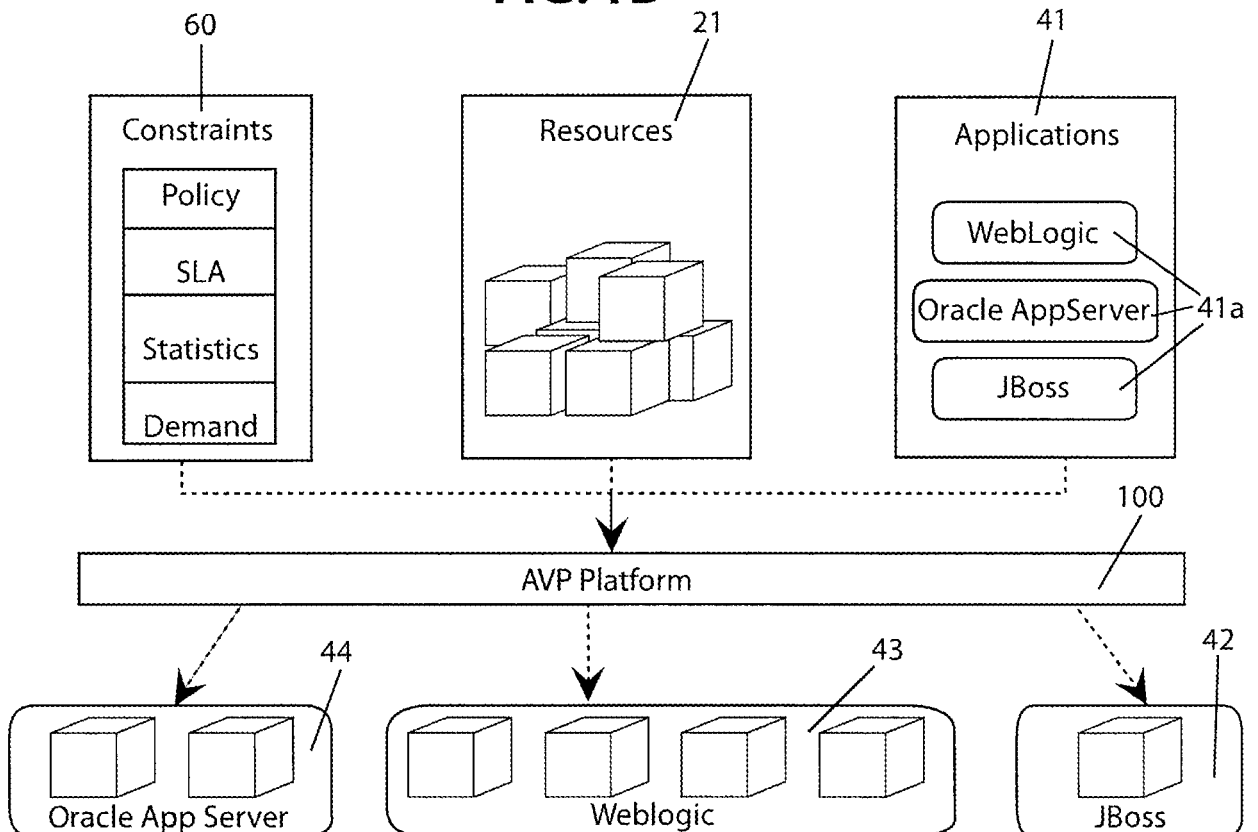
19. The method of claim 11, wherein the predetermined priorities provide an ordinal ranking
25 of computing domains, and the sequence of selected computing domains progresses from one or more computing domains having a highest priority ranking among the computing domains to one or more computing domains having a lowest priority ranking among the computing domains.

20. The method of claim 19, wherein the computing domain selection sequence comprises
30 one or more selection rounds, and wherein each selection round is associated with a priority ranking

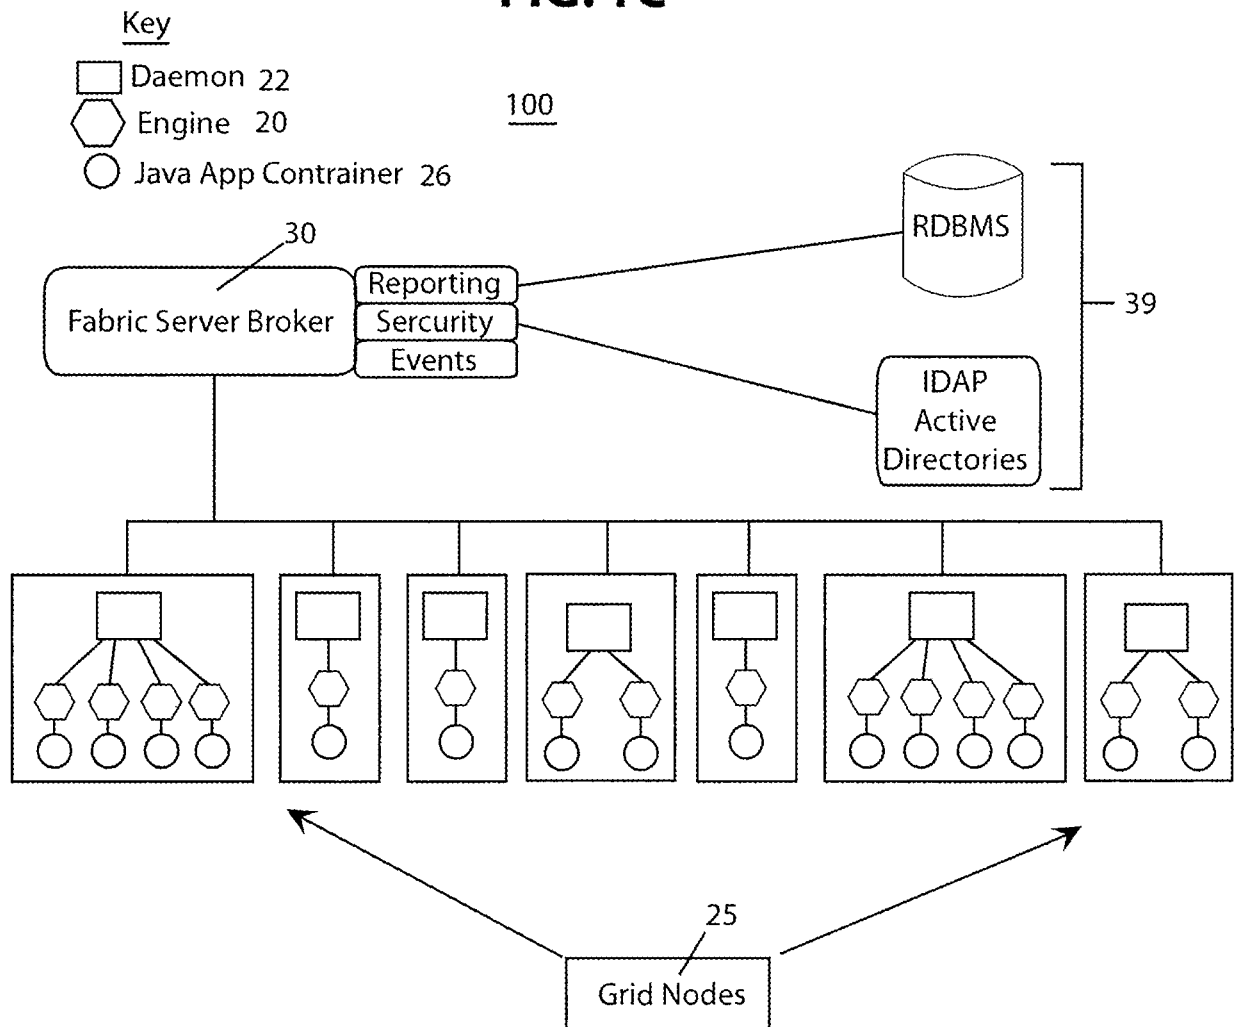
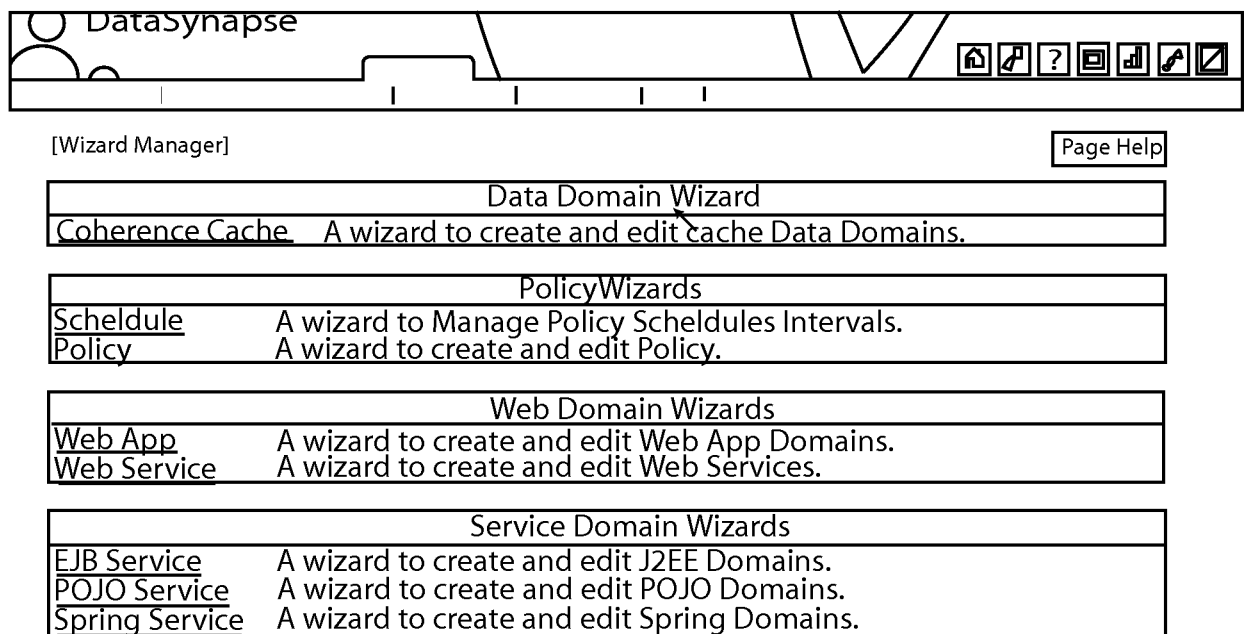
value, such that each of the plurality of computing domains having a priority ranking that is greater than or equal to the rank value is selected during the selection round.

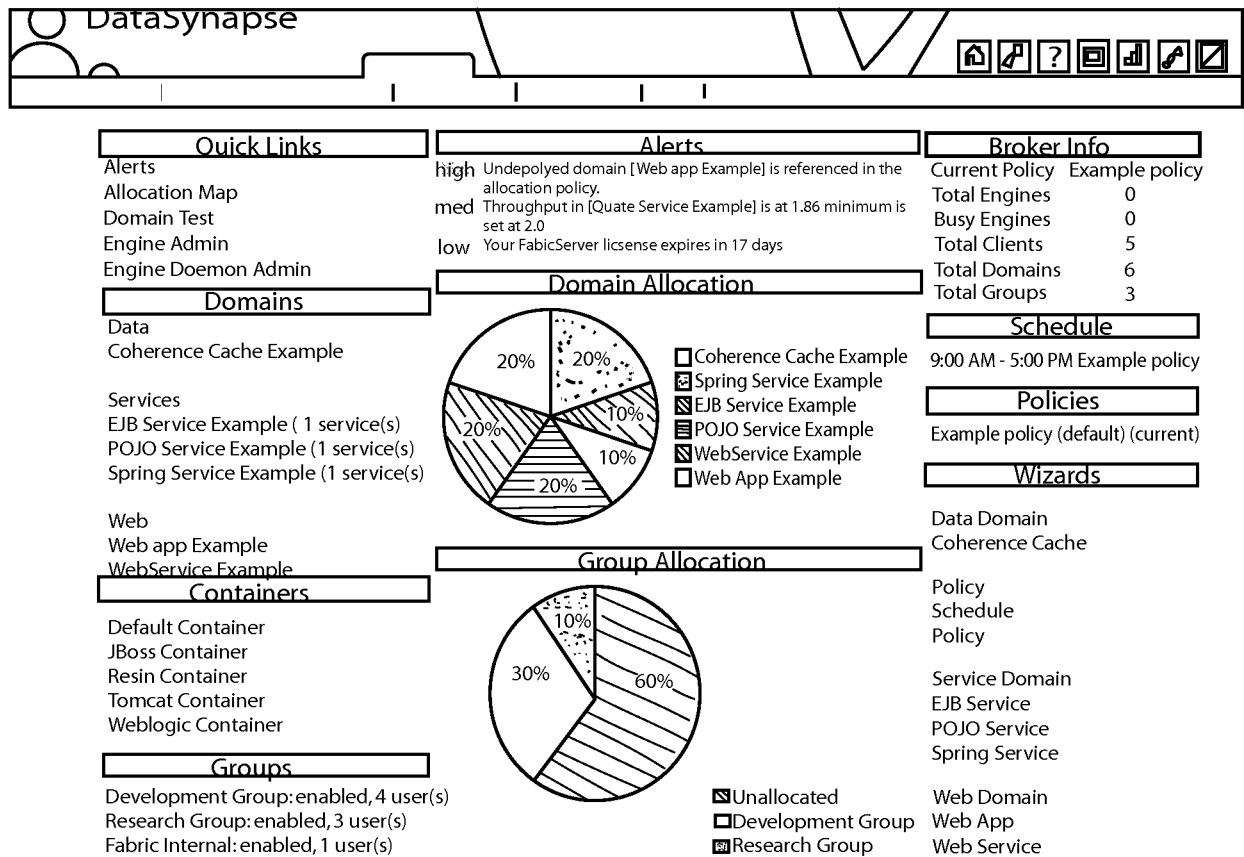

21. The method of claim 20, wherein the step of selecting an unallocated computing engine
- 5 for a selected computing domain in a current selection round is made according to a selection rule selected from the group consisting of:
- selecting a computing engine currently assigned to the selected computing domain;
 - selecting a computing engine previously assigned to the selected computing domain;
 - selecting a computing engine among unallocated computing engines that is not identified as
 - 10 being needed by another one of the plurality of computing domains; and
 - selecting a computing engine among unallocated computing engines.

1/24

FIG. 1A**FIG. 1B**

2/24

FIG. 1C**FIG. 2**

3/24
FIG. 3A**FIG. 3B**

4/24
FIG. 4A






FIG. 4B



5/24
FIG. 4C

[Containers]

Page Help

Name	Version	Description	Supported Dornain Types	Modified	Actions
Weblogic2 Container	8.1 SP5	Weblogic2 Container	Web App, Web Service	no	Actions ▾
Default Container	1.0	Basic container provided by DataSynapse	POJO Service, Spring Service	no	Actions ▾
Tomcat Container	4.0	Tomcat Container is the Fabric wrapper around Tomcat Application Server	Web App	no	Actions ▾
Weblogic Container	8.1 SP4	Weblogic Container	Webb App, Web Service	no	Actions ▾
Resin Container	n/a	Resin Container	Web App	no	Actions ▾
JBoss Container	4.0.2	JBoss Container is the Fabric wrapper around JBoss Application Server.	EJB Service, Web App, Web Service	no	Actions ▾

FIG. 4D

[Domain Wizard]


Web App
<p>I would like to:</p> <ul style="list-style-type: none"> . Add/edit Domain settings . Add/edit remove archive files . View servlet list . View EJB list . Add/edit tracked statistics . Add/edit advanced settings . Add/edit URL patterns . Add/edit container-specific config files . Add/edit container-specific runtime context variables
<div> <input type="button" value="Close"/> <input type="button" value="Previous"/> <input type="button" value="Next"/> <input type="button" value="Finish"/> </div>

6/24
FIG. 5A

[Policy Wizard]

Policy																							
Add or edit Domain allocations																							
<div style="display: flex; justify-content: space-around; margin-bottom: 5px;"> Add Edit Remove </div> <table border="1" style="width: 100%; border-collapse: collapse;"> <thead> <tr> <th style="text-align: left; padding: 5px;">Domain</th> <th style="text-align: center; padding: 5px;">Min</th> <th style="text-align: center; padding: 5px;">Max</th> <th style="text-align: center; padding: 5px;">Priority</th> </tr> </thead> <tbody> <tr> <td style="padding: 5px;">GenericWebService Domain</td> <td style="text-align: center; padding: 5px;">2</td> <td style="text-align: center; padding: 5px;">8</td> <td style="text-align: center; padding: 5px;">1</td> </tr> <tr> <td style="padding: 5px;">ClusteredRubisDomain</td> <td style="text-align: center; padding: 5px;">2</td> <td style="text-align: center; padding: 5px;">4</td> <td style="text-align: center; padding: 5px;">1</td> </tr> <tr> <td style="padding: 5px;">ClusteredGeneric Domain</td> <td style="text-align: center; padding: 5px;">1</td> <td style="text-align: center; padding: 5px;">3</td> <td style="text-align: center; padding: 5px;">1</td> </tr> <tr> <td style="padding: 5px;">ClusteredIODomain</td> <td style="text-align: center; padding: 5px;">2</td> <td style="text-align: center; padding: 5px;">3</td> <td style="text-align: center; padding: 5px;">1</td> </tr> </tbody> </table>				Domain	Min	Max	Priority	GenericWebService Domain	2	8	1	ClusteredRubisDomain	2	4	1	ClusteredGeneric Domain	1	3	1	ClusteredIODomain	2	3	1
Domain	Min	Max	Priority																				
GenericWebService Domain	2	8	1																				
ClusteredRubisDomain	2	4	1																				
ClusteredGeneric Domain	1	3	1																				
ClusteredIODomain	2	3	1																				
<div style="display: flex; justify-content: space-around; margin-top: 10px;"> Cancel Previous Next Finish </div>																							


FIG. 5B




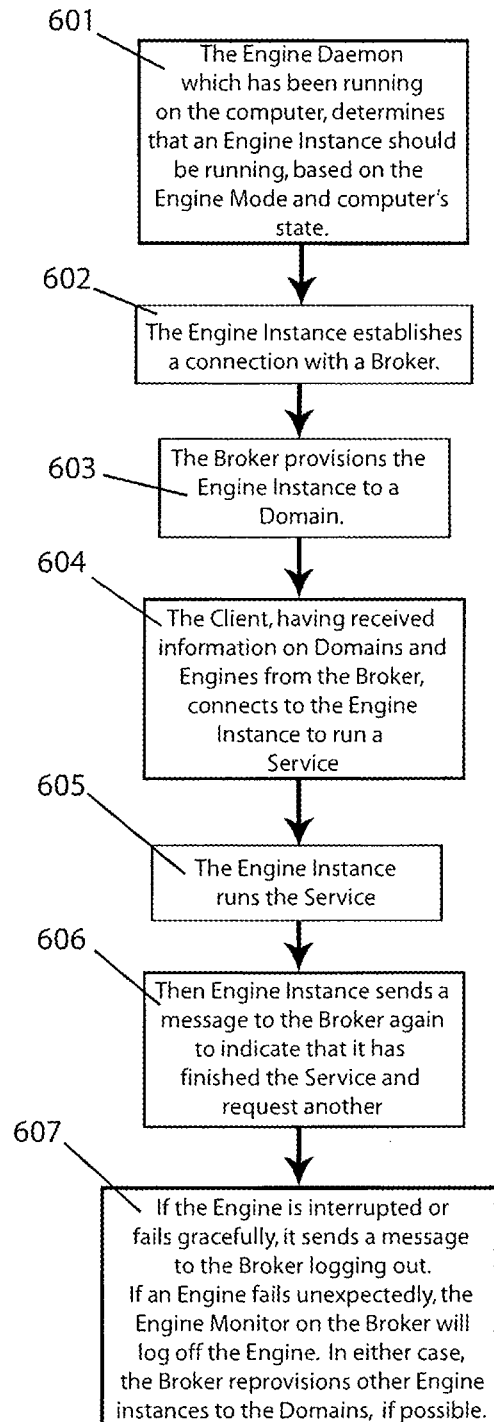
7/24

FIG. 5C

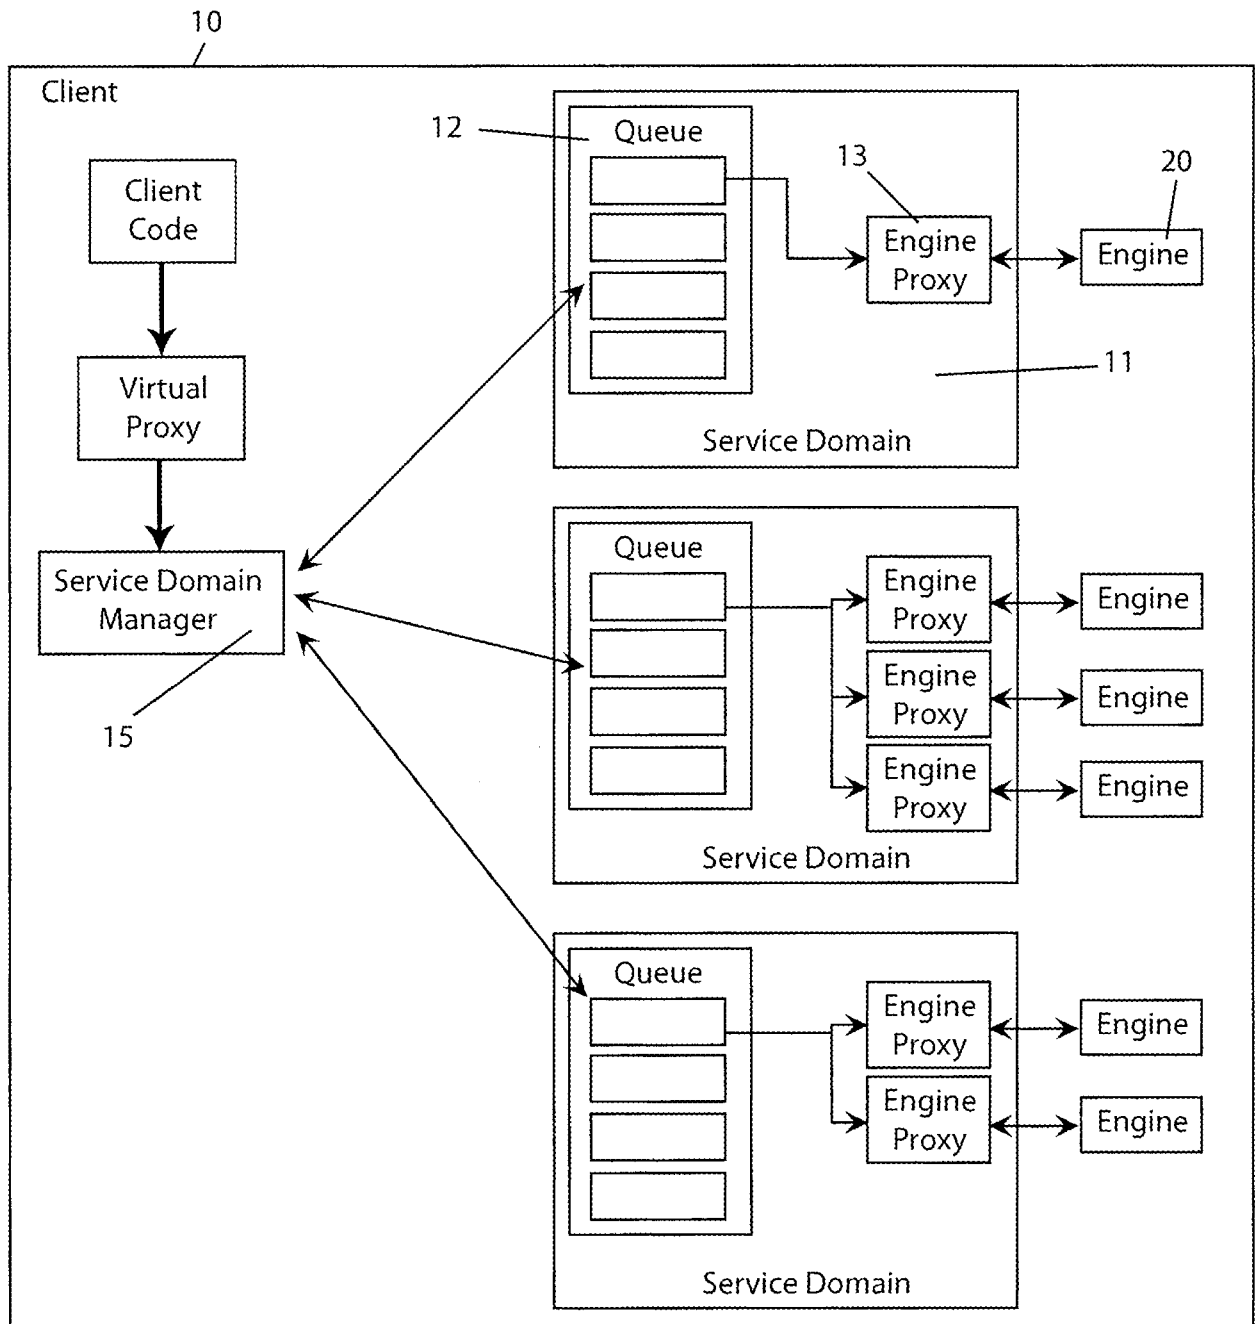
Enter rule definition	
Domain	<input type="text" value="ClusteredGenericDomain"/>
Rule	
Statistics	<input type="text" value="Weblogic Queue Size"/>
Sampling Window	<input type="text" value="120000"/>
Voilation Behavior	<input type="text" value="Add engine if > max/ Rmove engine if < min"/>
Min	<input type="text" value="40"/>
Max	<input type="text" value="10"/>
<input type="button" value="OK"/> <input type="button" value="Cancel"/>	

FIG. 5D



8/24

FIG. 6A**FIG. 6B**

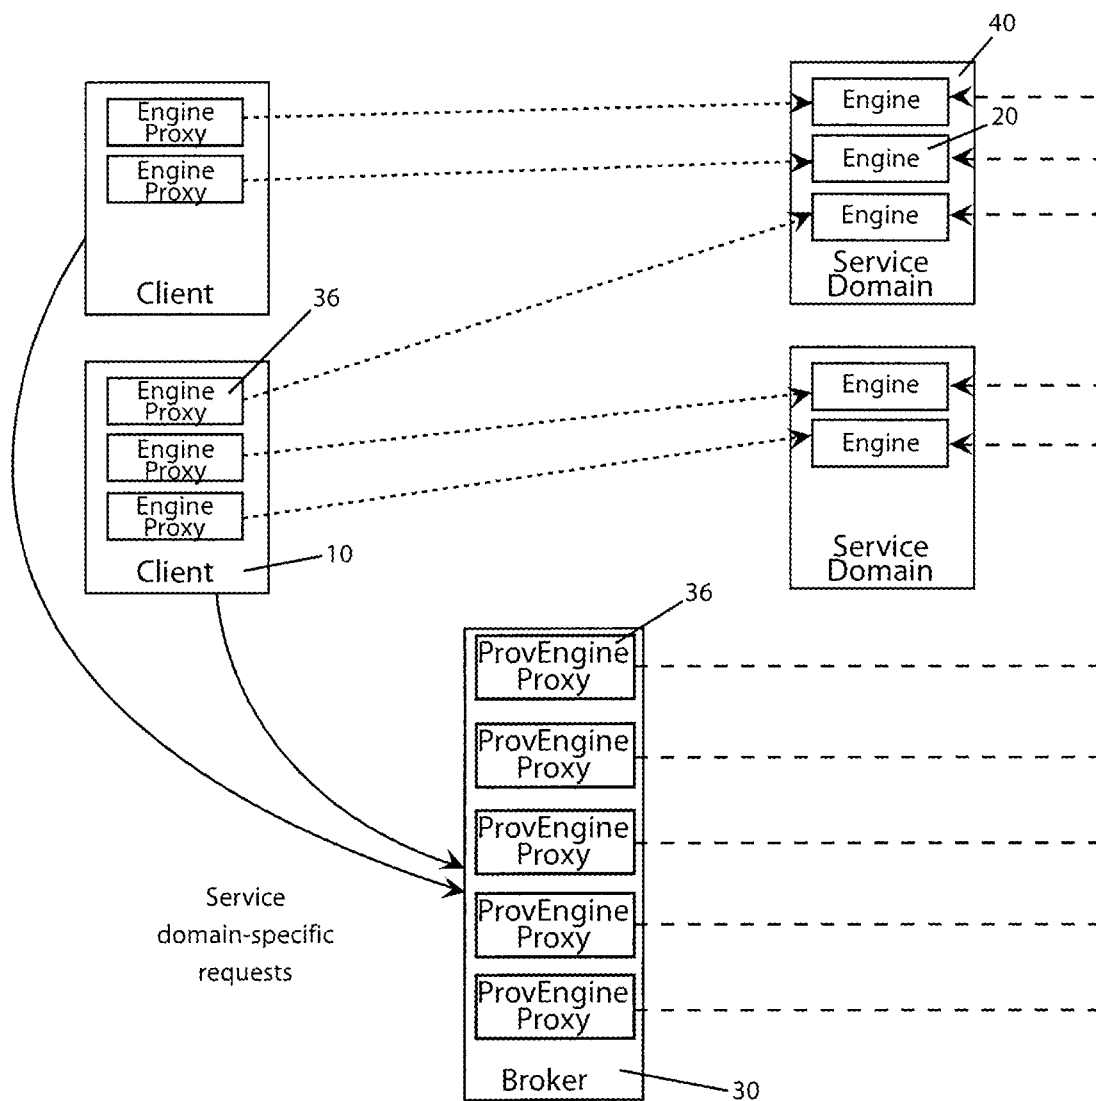
9/24

FIG. 7A

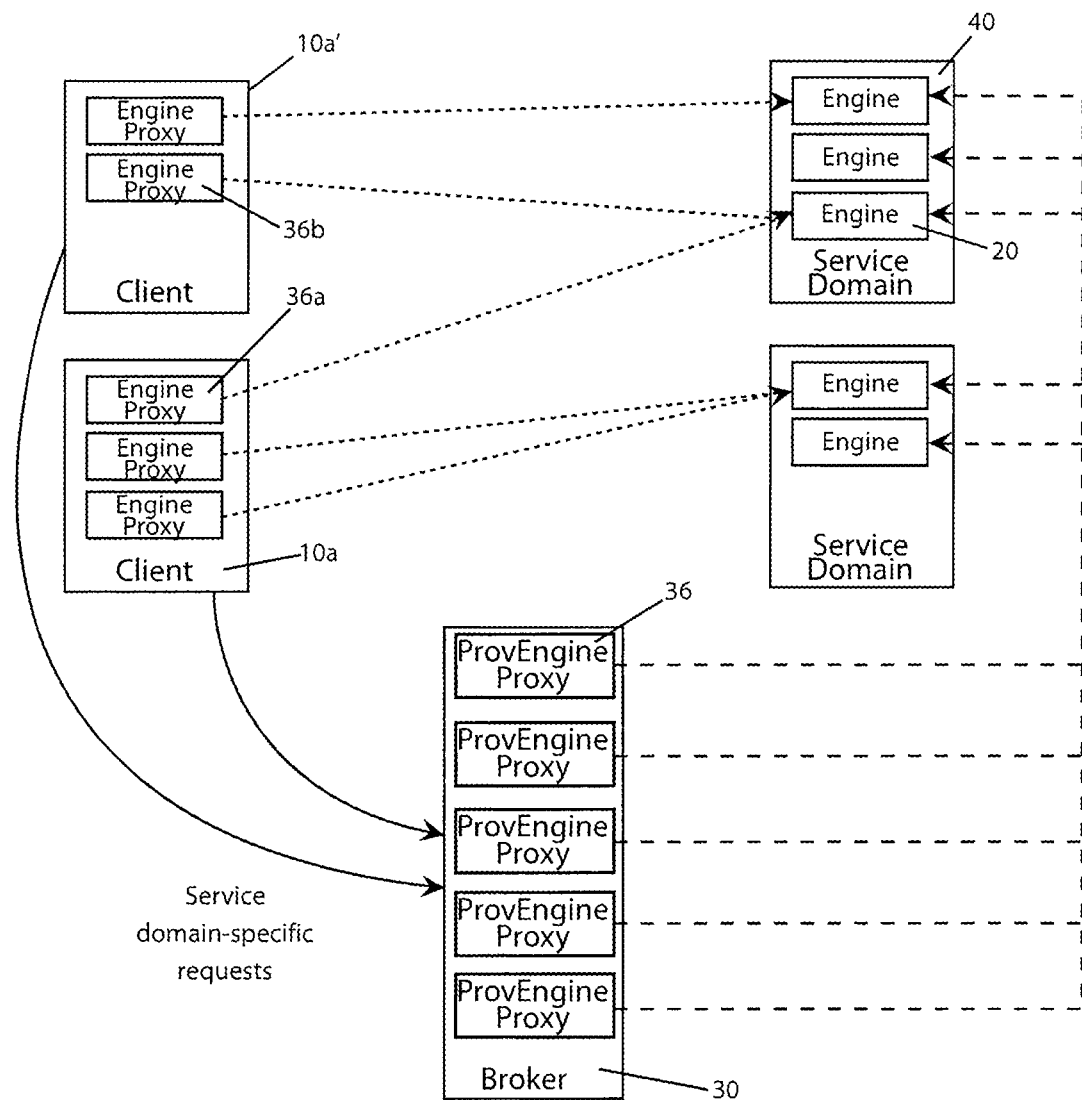
10/24

FIG. 7B**FIG. 8A**

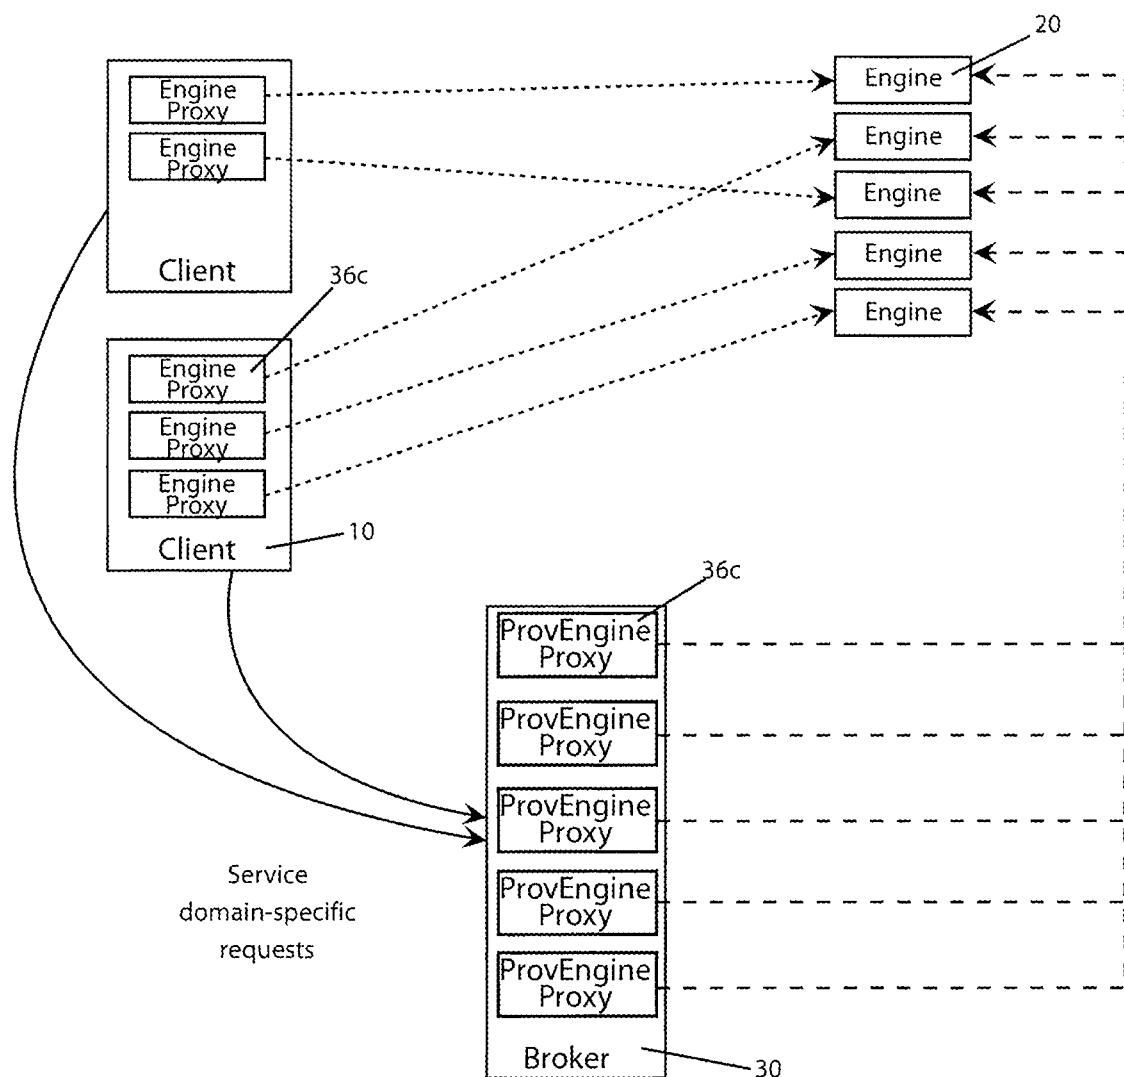
11/24

FIG. 8B

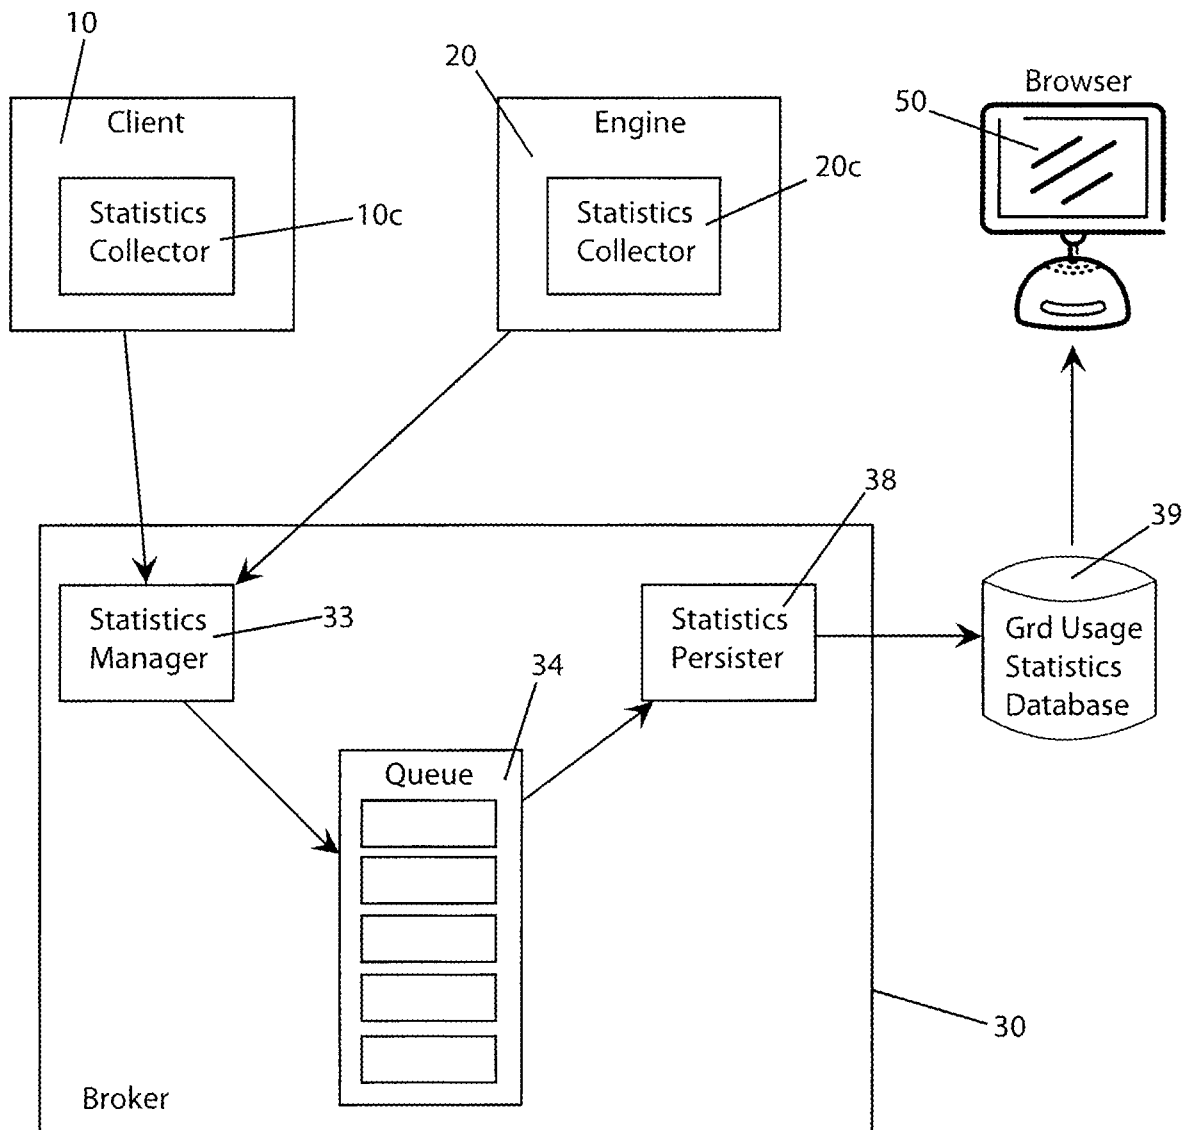
12/24

FIG. 8C

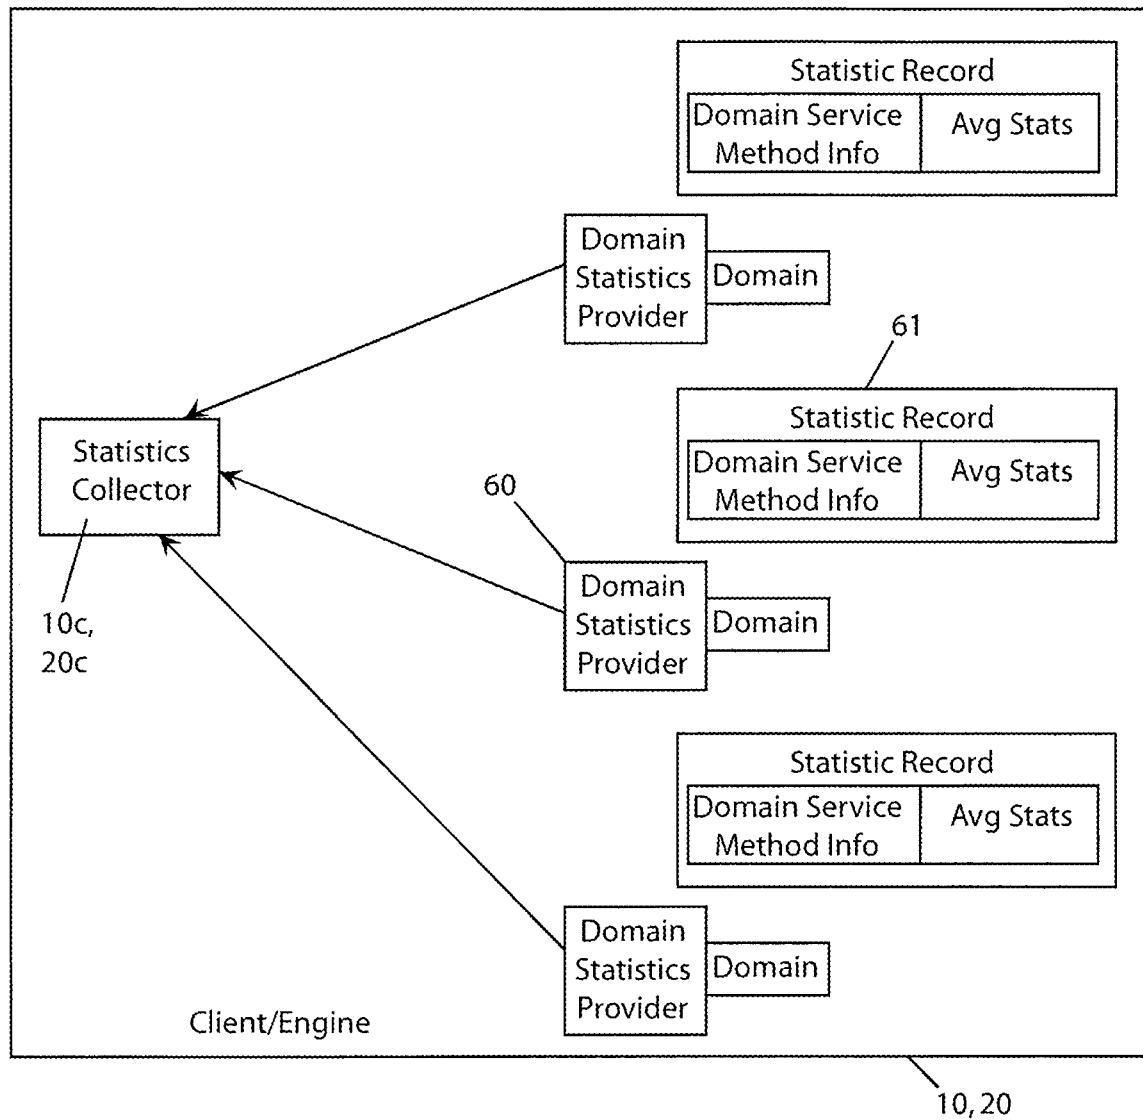
13/24

FIG. 8D

14/24

FIG. 9A

15/24

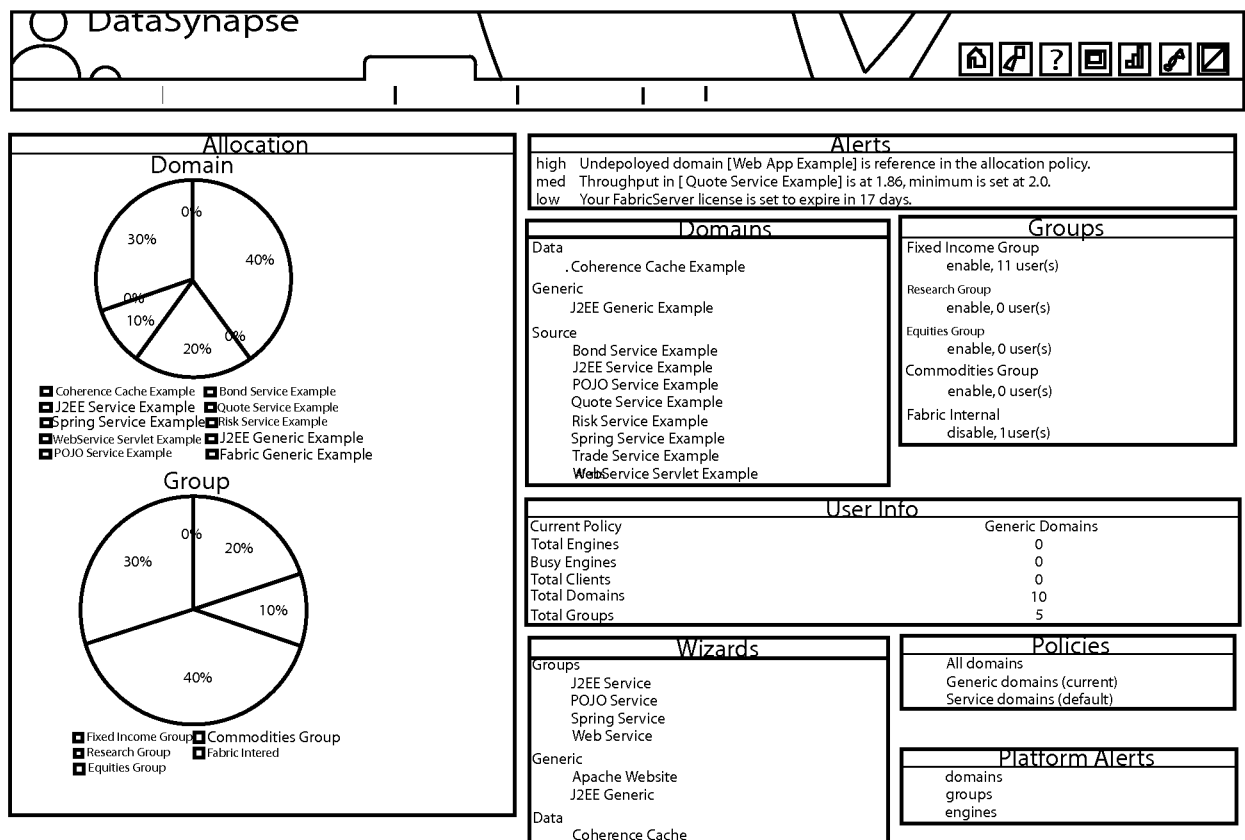
FIG. 9B**FIG. 9C**

Client	Engine	Client Grid	Engine Grid
Response time	Business method time	Response time	Business method time
Throughput	Throughput	Throughput	Throughput
Time on the queue	Throughput	Time on the queue	CPU Utilization
Queue size	Memory	Queue size	Number of engines
CPU Utilization	CPU Utilization	CPU Utilization	
Disk size	Disk size	Number of clients	

16/24

FIG. 9D

VGateway/VProxy	Domains
Response time	JMX-attributes
Throughput	Number of threads
Time on the queue	Free Memory
Queue size	CPU Utilization
	Free Disk space



FIG. 3E

17/24

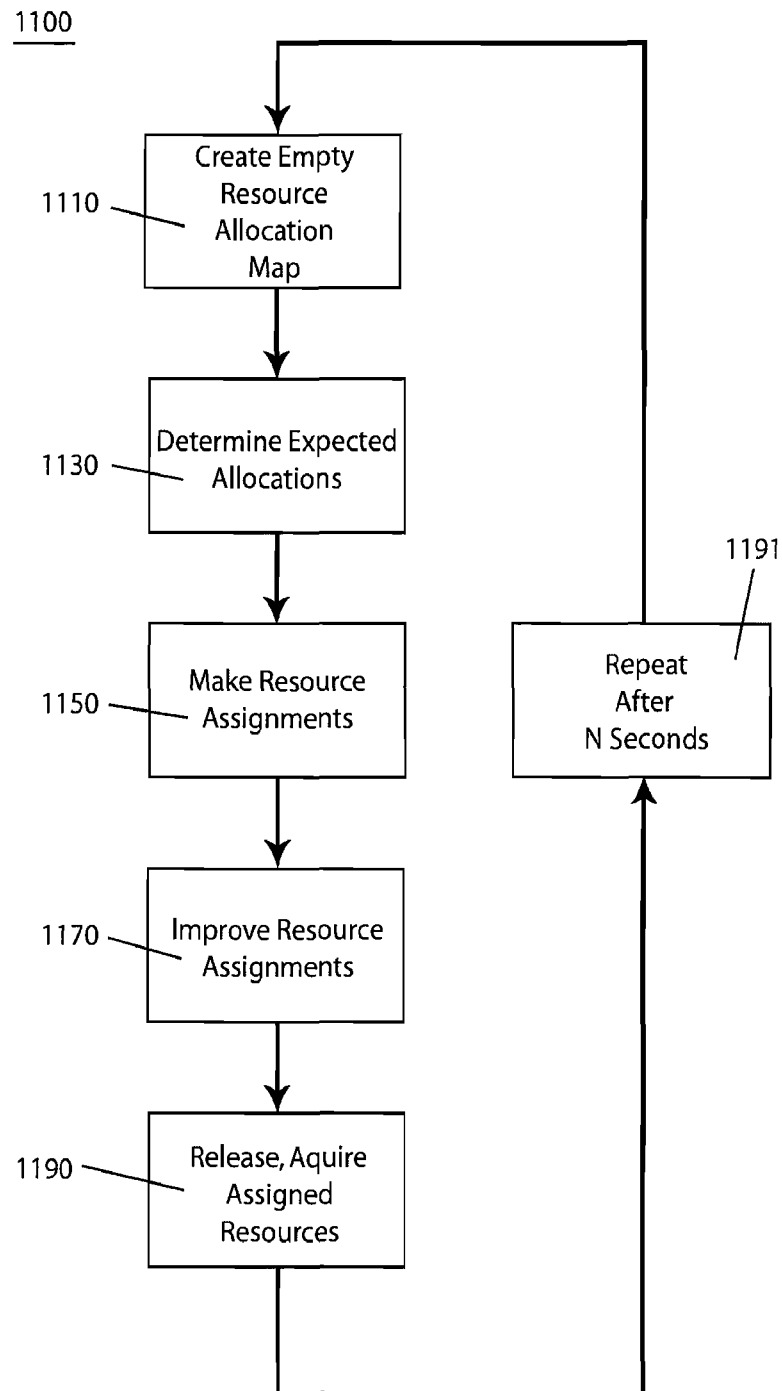
FIG. 9F

[Engine Charts]

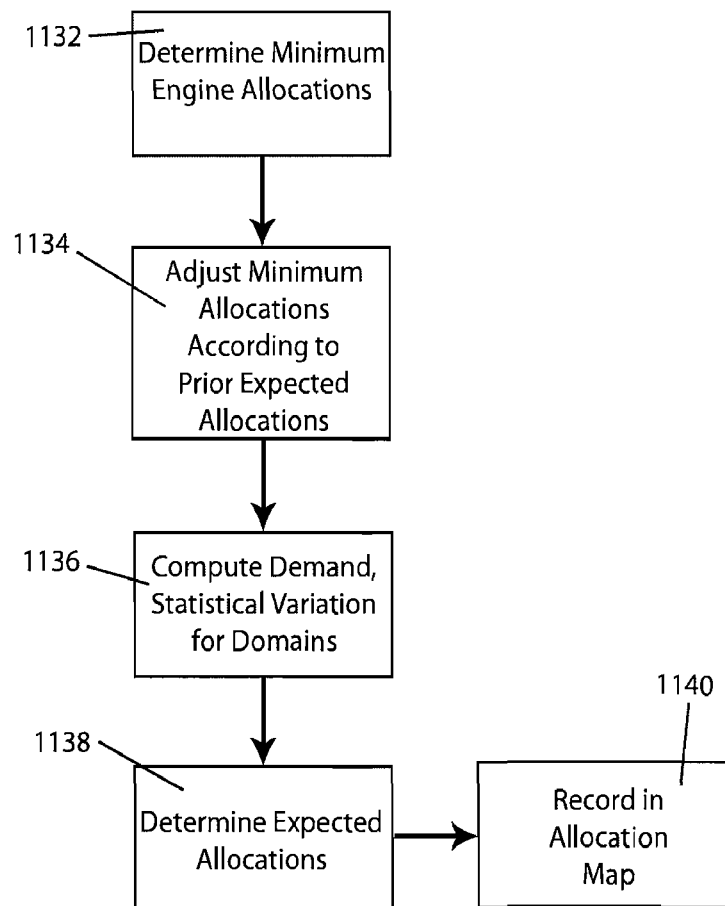
Choose Database		Choose Engine	
Database:	<input type="text" value="In Memory"/>	Engine:	<input type="text" value="PayHead-5640537194383352260"/>
Plot x-Axis		Plot y-Axis	
Statistic:	<input type="text" value="Time"/>	Statistic:	<input type="text" value="Thread Count"/>
<input type="button" value="Display"/>			

**FIG. 10**

18/24

FIG. 11A

19/24

FIG. 11B1130

20/24

FIG. 11C

```

for each domain {
    create allocation entry for this domain
    add interested clients to this entry
    set expected engines[ domain ] = min engines[ domain ]

```

FIG. 11D

```

for each allocation entry {
    if ( current expected engines[ domain ] > expected engines[ domain ] ) {
        set expected engines[ domain ] = expected engines[ domain ]
    }
    demand[ domain ] = compute demand ( set of statistics[ domain ] )
    adjust demand[ domain ] based on statistical variations between engines
    if ( demand[ domain ] > 0 ) {
        set expected engines[ domain ] = min( max engines[ domain ],
            expected engines[ domain ] + demand[ domain ] )
    }
}

```


FIG. 11E

Let E_1, \dots, E_n denote engines currently working on domain D
 Let V_{1i}, \dots, V_{ni} be the sequence of values for a statistic on the i th engine
 And let $A_i = \text{Avg}(E_i) = \sum V_{ji} / n$, with $1 < j < n$
 (i.e. average value of per engine)

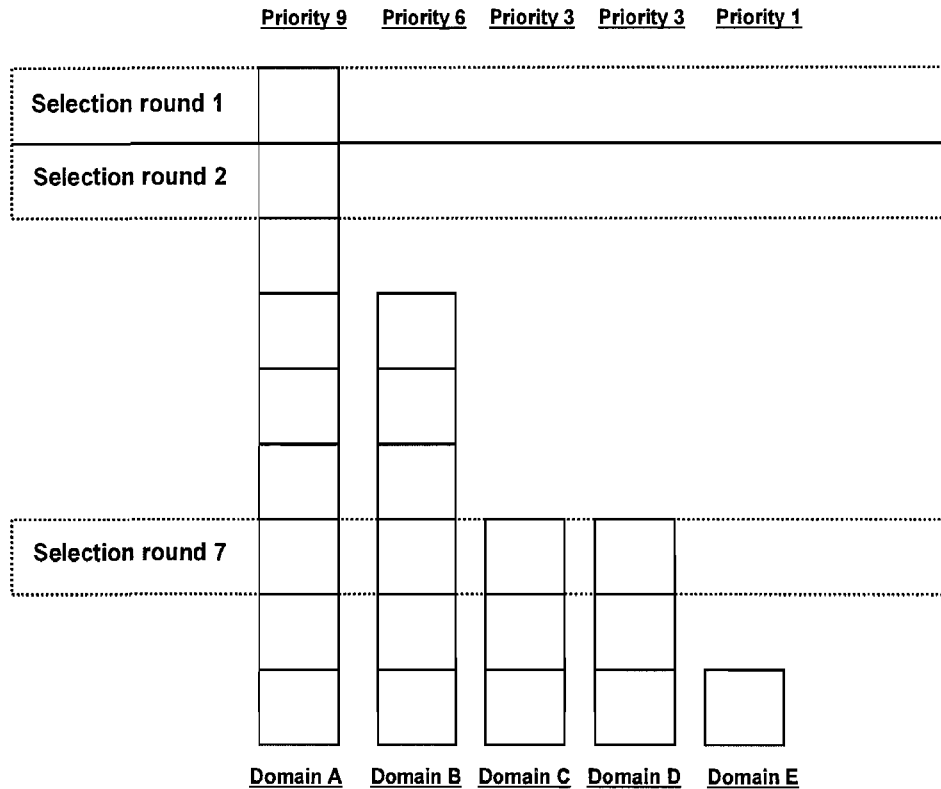
Then use normalized geometric variance as follows:

$GV = \text{power}(\prod A_i, 1/n) / (\sum A_i / n)$, with $1 < i < n$
 (i.e. geometric mean divided by arithmetic mean)

21/24

FIG. 11F

22/24

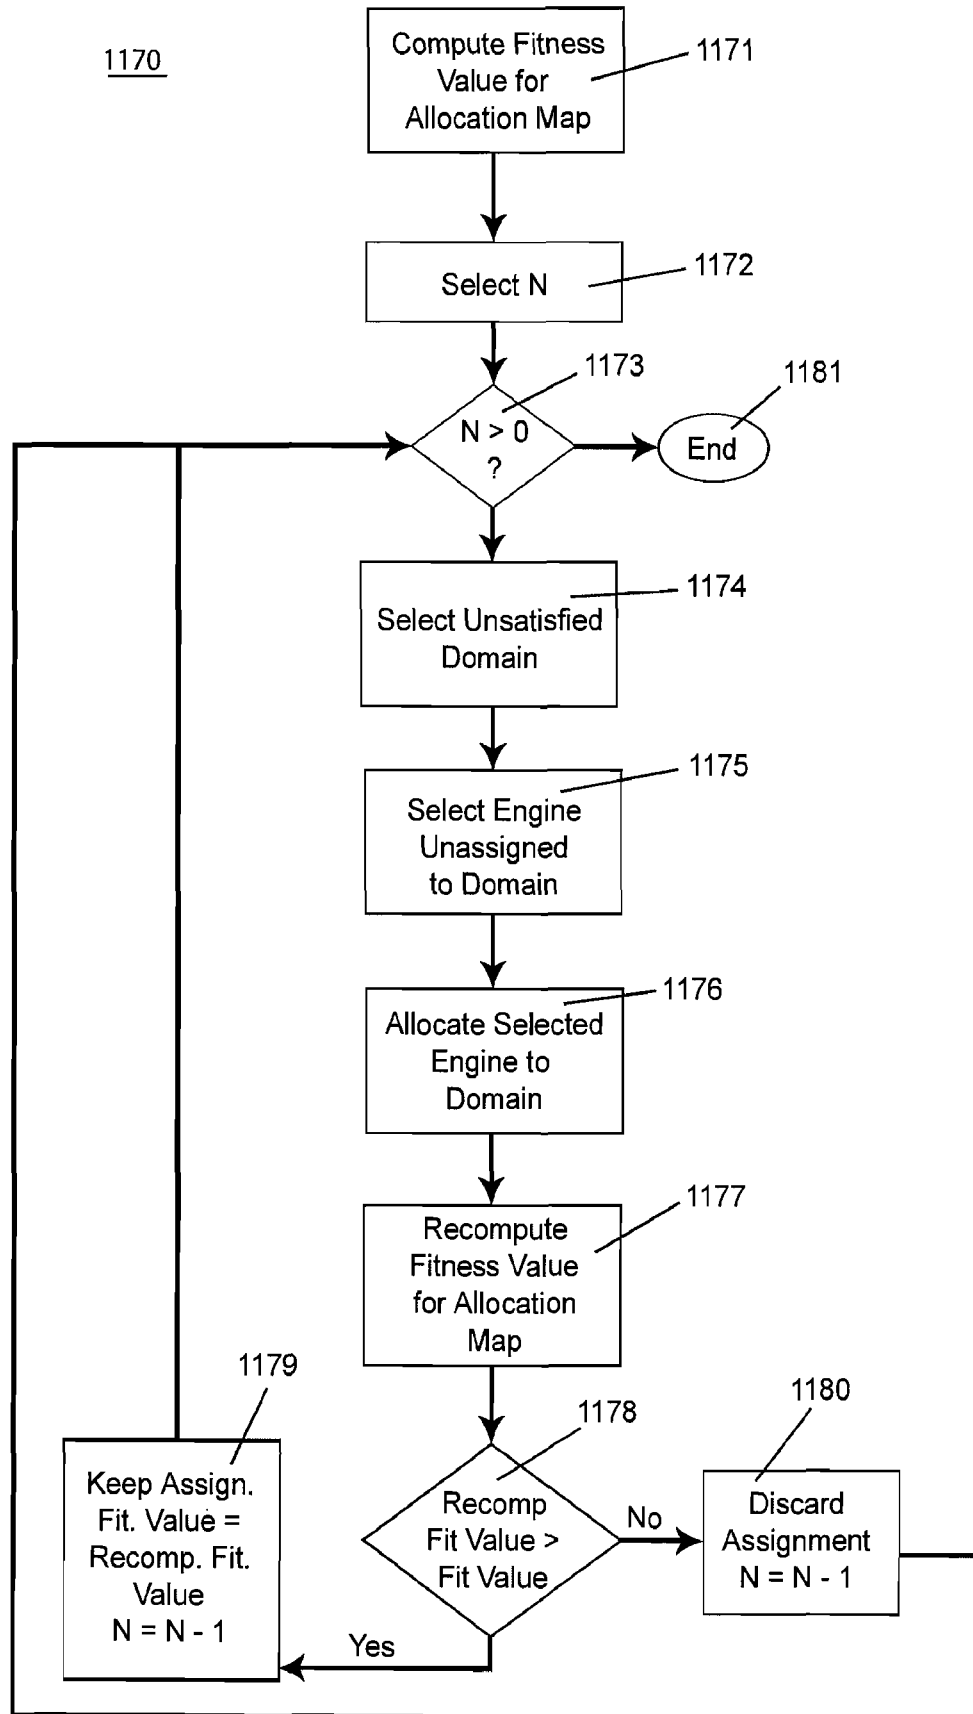
FIG. 11G**FIG. 11H**

```

partition/sort allocation entries into buckets by domain priority
set max priority = max( priority[ all domains ] )
while ( !done ) {
    for each priority from max priority down to 0 {
        for each allocation entry with domain priority >= current priority {
            filter the list of available engines based on current entry rules
            pick engine for current entry from filtered list using Selection Algorithms
        }
    }
    if ( there are no more engines || could not pick an engine ) {
        set done = true
    }
}

```

23/24

FIG. 11I

24/24

FIG. 11J

```
set current fitness = fitness ( Allocation Map )
repeat for N iterations or until there are unsatisfied entries {
    select unsatisfied Allocation Entry
    select a random engine from either {
        set of free engine or
        satisfied entry
    }
    assign the engine to unsatisfied entry
    if ( fitness( Allocation Map ) > current fitness ) {
        keep the assignment
        current fitness = fitness ( Allocation Map )
    } else {
        discard the assignment
    }
}
```