



(19) **United States**

(12) **Patent Application Publication**

Joshi et al.

(10) **Pub. No.: US 2018/0007166 A1**

(43) **Pub. Date: Jan. 4, 2018**

(54) **PRIORITY-DRIVEN BOXCARRING OF ACTION REQUESTS FROM COMPONENT-DRIVEN CLOUD APPLICATIONS**

(52) **U.S. Cl.**
CPC *H04L 67/327* (2013.01); *H04W 4/001* (2013.01); *H04L 67/02* (2013.01); *H04L 67/42* (2013.01)

(71) Applicant: **salesforce.com, inc.**, San Francisco, CA (US)

(57) **ABSTRACT**

(72) Inventors: **Shweta Joshi**, Newark, CA (US); **Ashraya Raj Mathur**, Fremont, CA (US); **Ronnie Fong**, San Jose, CA (US); **Jianxun Jason Ding**, Cupertino, CA (US)

Improved perceived load time for browser and mobile application pages is achieved by adjusting boxcarring of action requests from coupled data consuming applications on the user device, using the priority level of regions and components in component-driven cloud applications. Priority labels differentiate among display regions rendered by the data consuming application and the priority labels further differentiate among components within respective display regions. The middleware application batches the action requests into batches based at least in part on the priority labels, into boxcars segregated by priority label according to a predetermined segregation schedule, and dispatches the boxcars of batched action requests to the server. Performance is also dynamically speeded up, by adjusting inter-boxcar intervals used to dispatch batches of action requests from the user device to a production server, based on the dynamically measured network communication latency between the user device and the server.

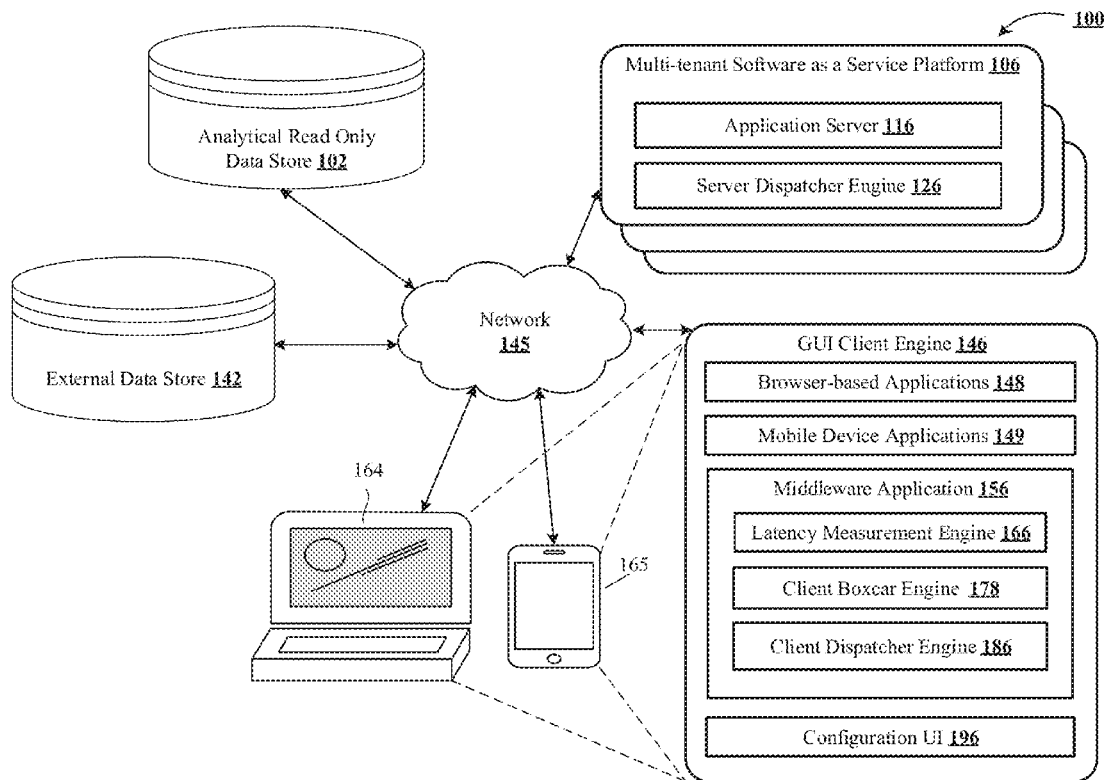
(73) Assignee: **salesforce.com, inc.**, San Francisco, CA (US)

(21) Appl. No.: **15/199,768**

(22) Filed: **Jun. 30, 2016**

Publication Classification

(51) **Int. Cl.**
H04L 29/08 (2006.01)
H04L 29/06 (2006.01)
H04W 4/00 (2009.01)



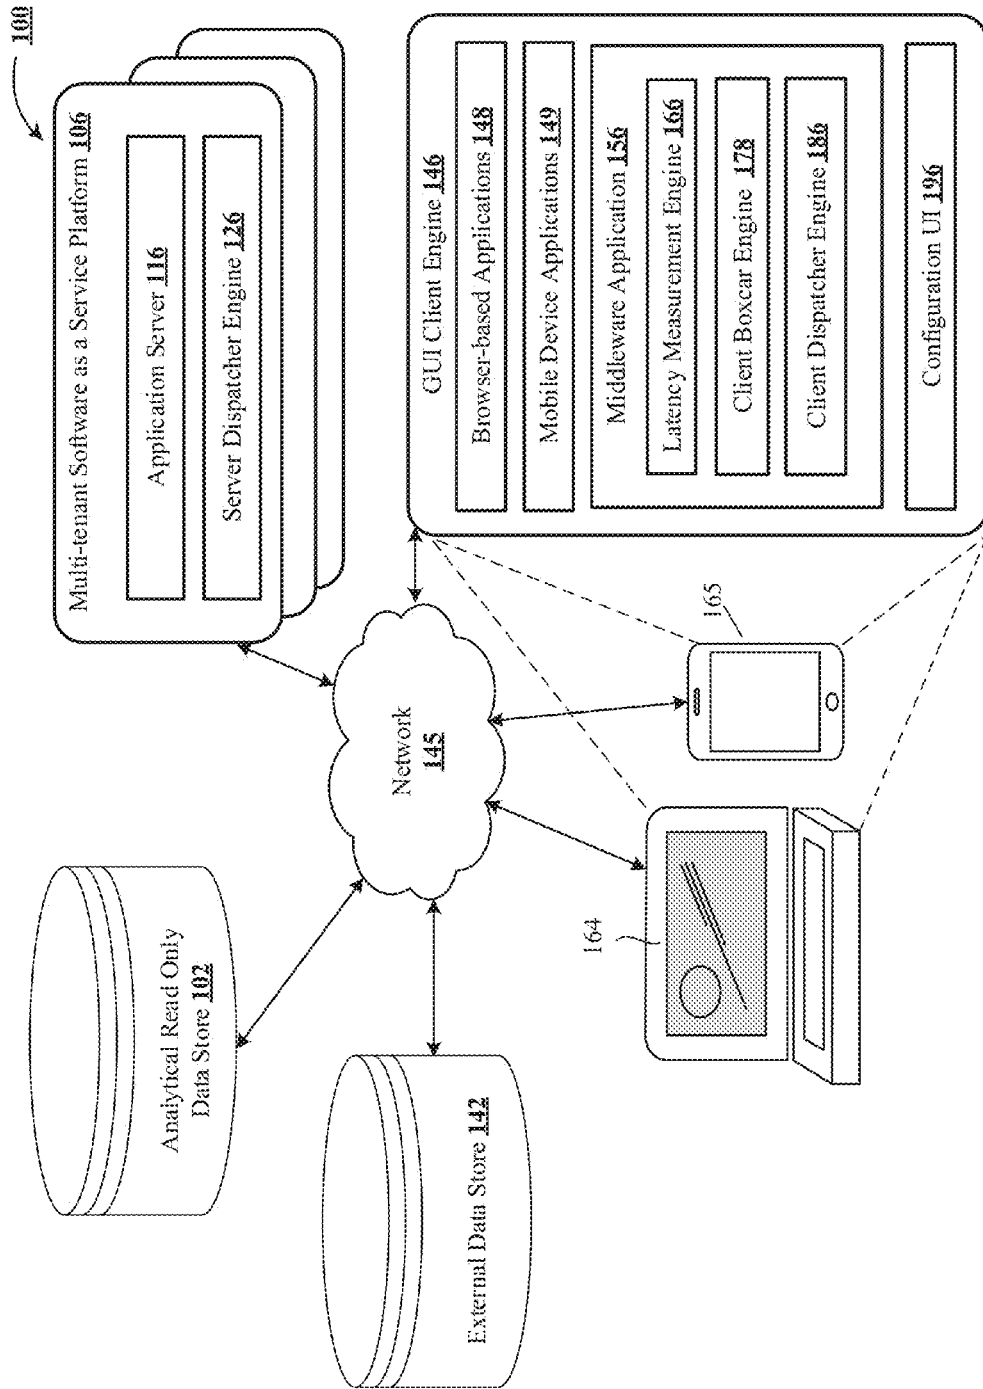


FIG. 1

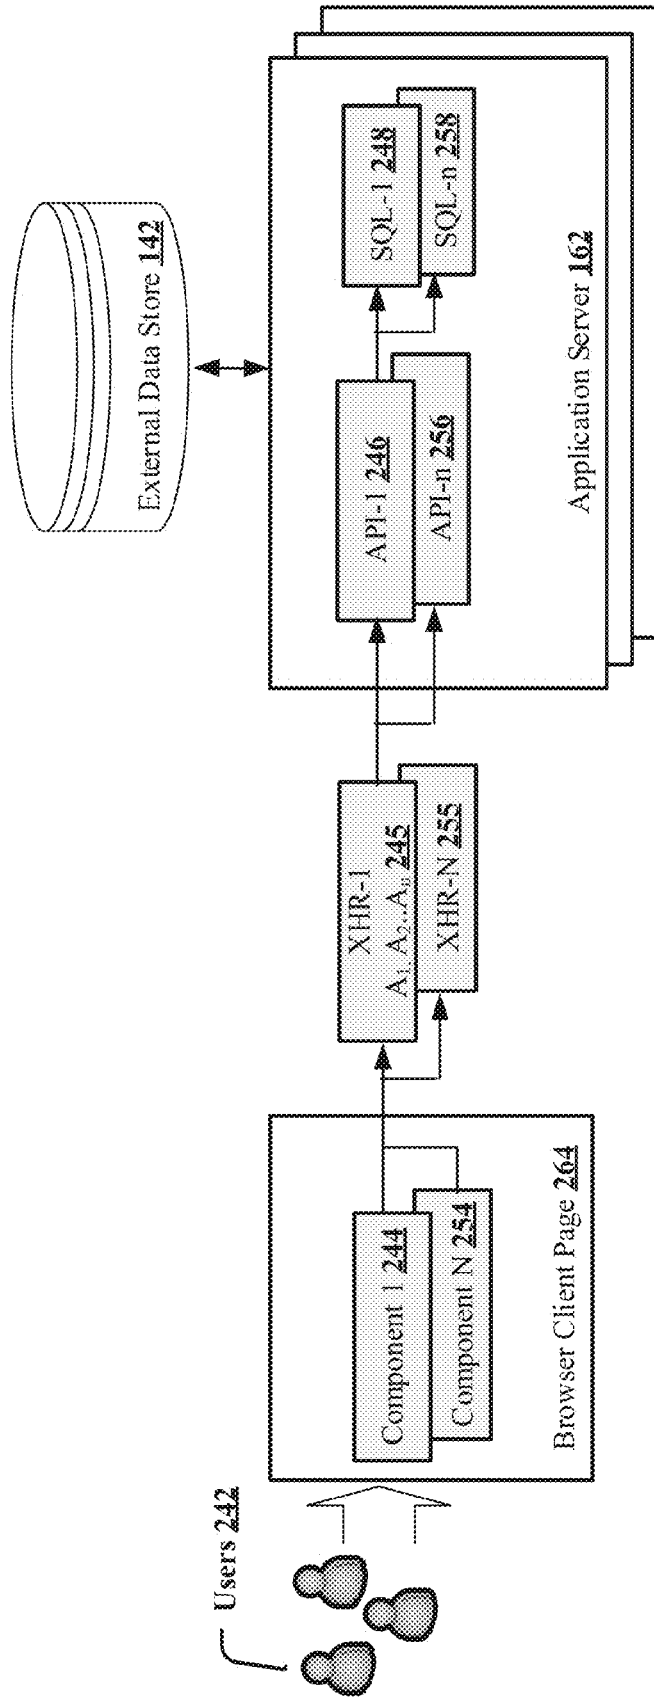


FIG. 2

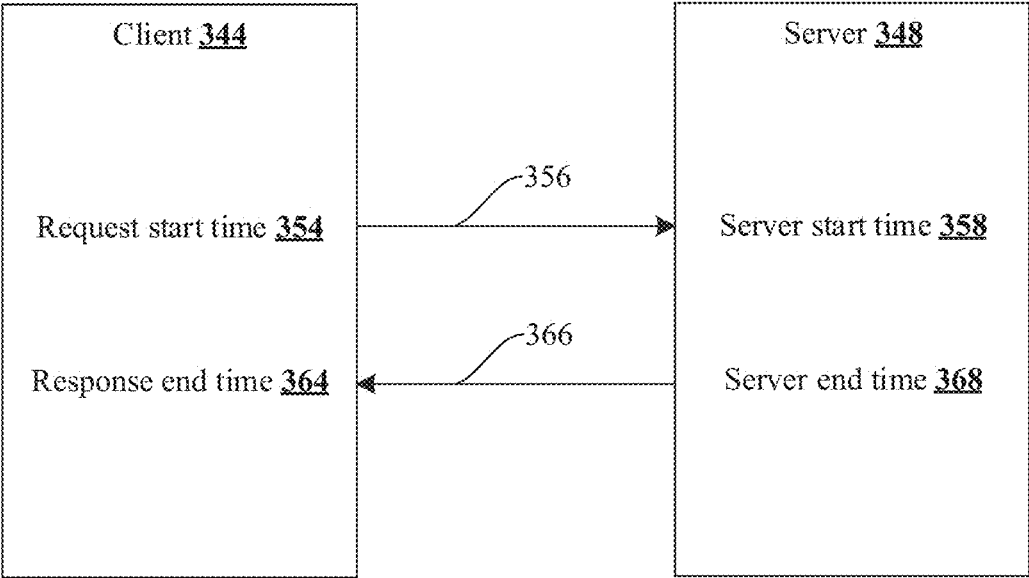


FIG. 3A

304 314 315

Top Browser name	Score	Performance	Connections per hostname	Max Connections	Script Sheet	Script Image	Script Frame	Async Scripts	CSS	CSS+ Inline Script	Cache Expires	Cache Redirects	Cache Resources Redirects	Link Prefetch	data URIs	Headers in trailer	# feeds
Chrome 32	12/16	yes	6	10	yes	yes	no	yes	yes	yes	yes	yes	yes	no	yes	no	345
Firefox 26	11/16	yes	6	17	yes	yes	no	yes	no	no	yes	yes	yes	yes	yes	no	305
IE 9	12/16	yes	6	35	yes	yes	no	no	yes	yes	yes	yes	yes	no	yes	no	1521
IE 10	12/16	yes	6	17	yes	yes	no	yes	yes	yes	yes	yes	yes	no	yes	no	344
IE 11	12/16	yes	12	17	yes	yes	no	yes	yes	yes	yes	yes	yes	no	yes	no	327
Safari 7.0.1	11/16	no	6	17	yes	yes	no	yes	yes	yes	yes	yes	yes	no	yes	no	28
Chrome 34	12/16	yes	6	10	yes	yes	no	yes	yes	yes	yes	yes	yes	no	yes	no	323
Firefox 27	11/16	yes	6	17	yes	yes	no	yes	no	no	yes	yes	yes	yes	yes	no	307
Android 2.3	8/16	no	3	10	yes	yes	no	no	yes	no	yes	no	yes	no	yes	no	333
Android 4	12/16	yes	6	17	yes	yes	no	yes	yes	yes	yes	yes	yes	yes	yes	no	322
BlackBerry 7	11/16	no	7	9	yes	yes	no	yes	yes	yes	yes	yes	yes	no	yes	no	7
Chrome Mobile 18	12/16	yes	6	16	yes	yes	no	yes	yes	yes	yes	yes	yes	no	yes	no	28
IE Mobile 9	11/16	yes	6	50	yes	yes	no	no	yes	yes	yes	no	yes	no	yes	no	3
IE Mobile 10																	0
iPhone 7																	0

FIG. 3B

400

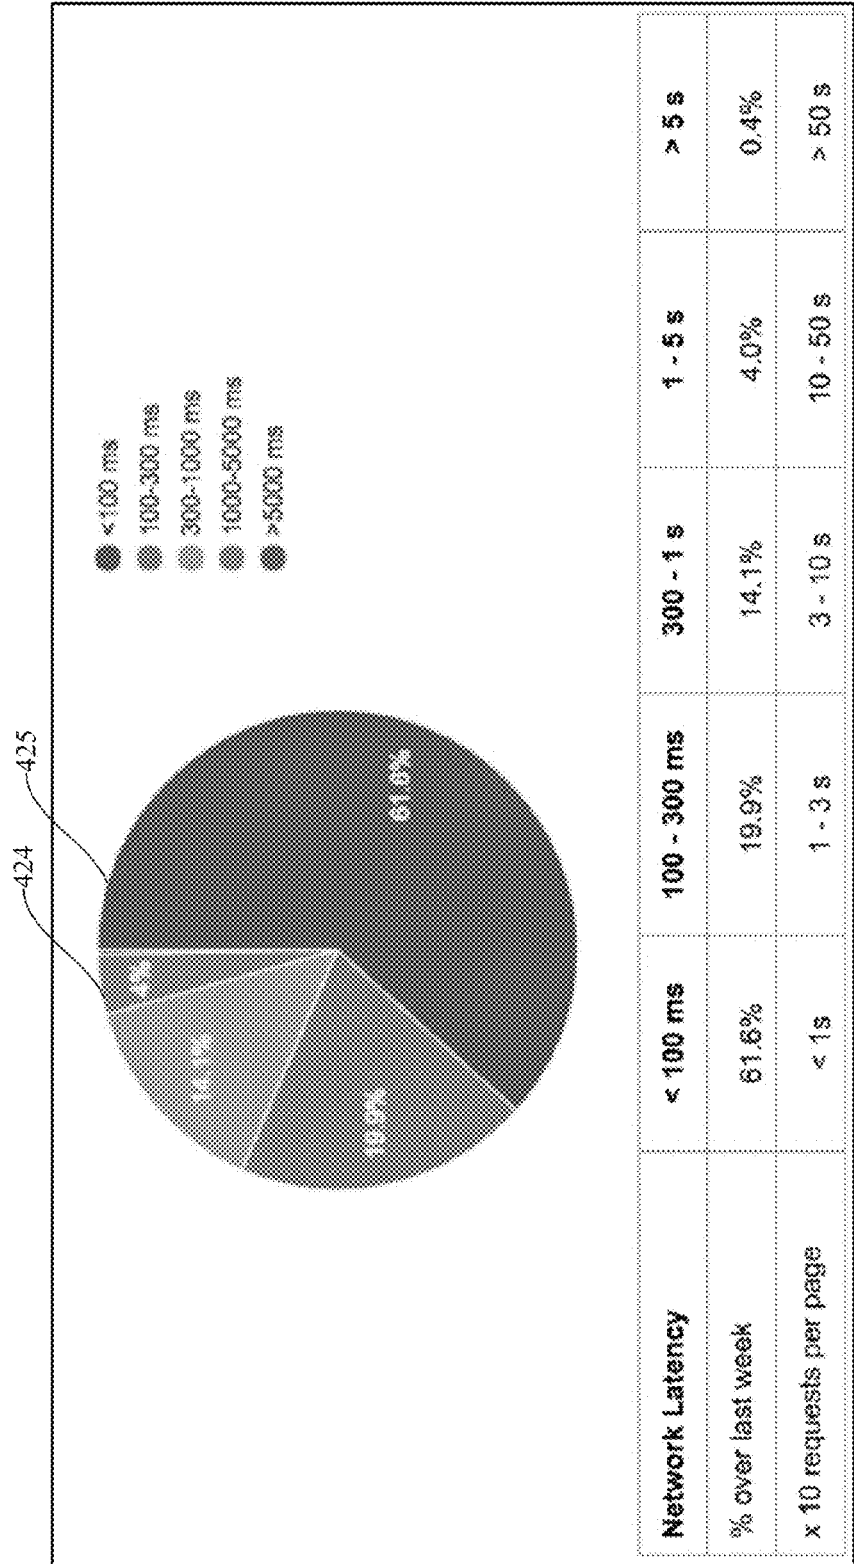


FIG. 4

500

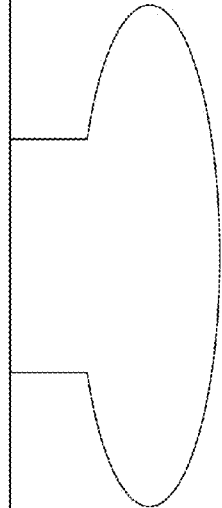
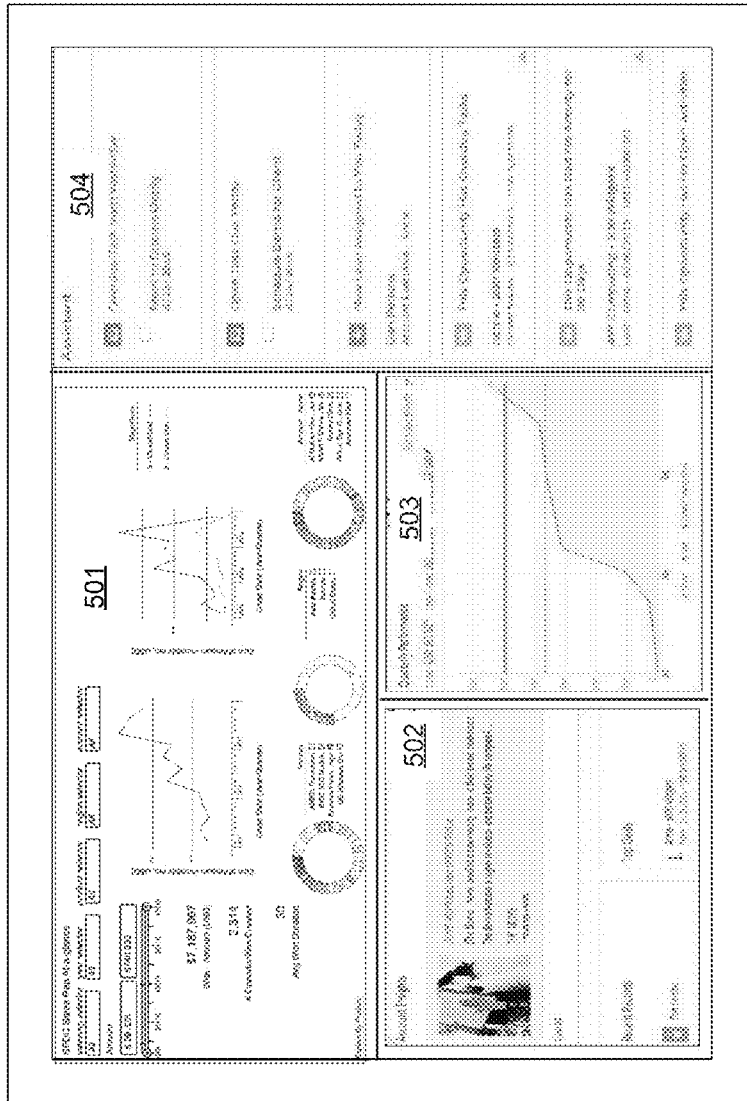


FIG. 5

```

600
<body>
  <div class="DESKTOP uiContainerManager flexipageEditorSurface">
    <div class="forceStyle flexipageEditorSurface">
      <div data-type="page" class="flexipageEditorPage">
        <div class="responsiveContents homeDesktopTemplate">
          <div class="colMain">
            632
            <div data-type="top">
              <div data-type="region"
                data-max-allowed="25"
                data-allow-drop="true"
                data-allowed="flexipage:
                availableForAllPageTypes,home:
                availableForDesktop"
                class="flexipageEditorRegion" priority="1">
                645
                <a class=
                  "flexipageEditorRegionPlaceholder" href="javascript:void(0);" >Add Component(s) Here</a>
              </div>
            </div>
            <div class="bottomWrapper">
              <div class="bottomLeft">
                652
                <div data-type="region"
                  data-max-allowed="25"
                  data-allow-drop="true"
                  data-allowed="flexipage:
                  availableForAllPageTypes,home:
                  availableForDesktop"
                  class="flexipageEditorRegion" priority="2">
                  <a class="flexipageEditorRegionPlaceholder" href="javascript:void(0);" >Add Component(s) Here</a>
                </div>
              </div>
            </div>
          </div>
        </div>
      </div>
    </div>
  </div>
  ...

```

FIG. 6

700

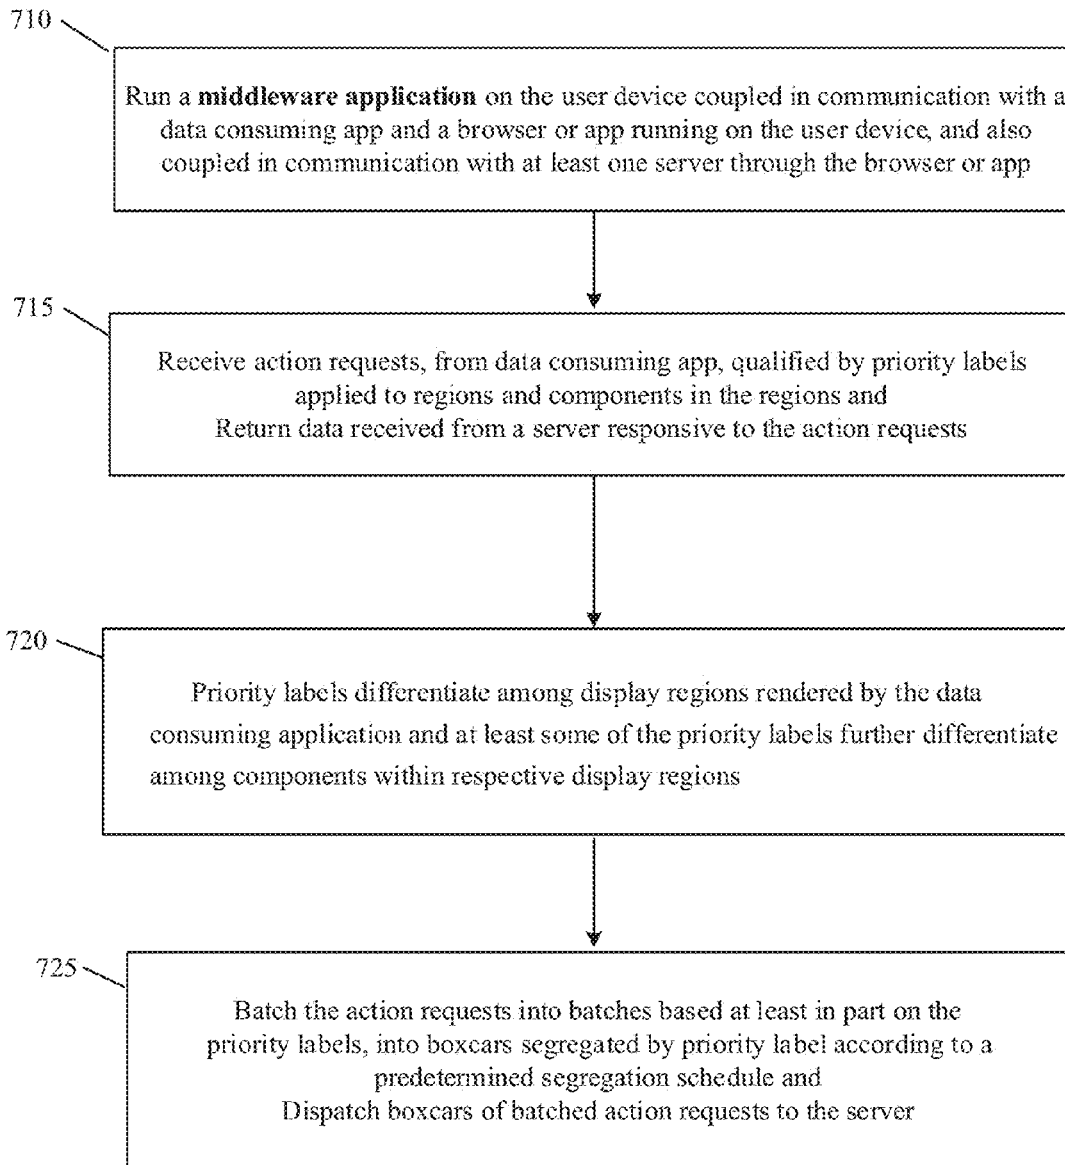


FIG. 7

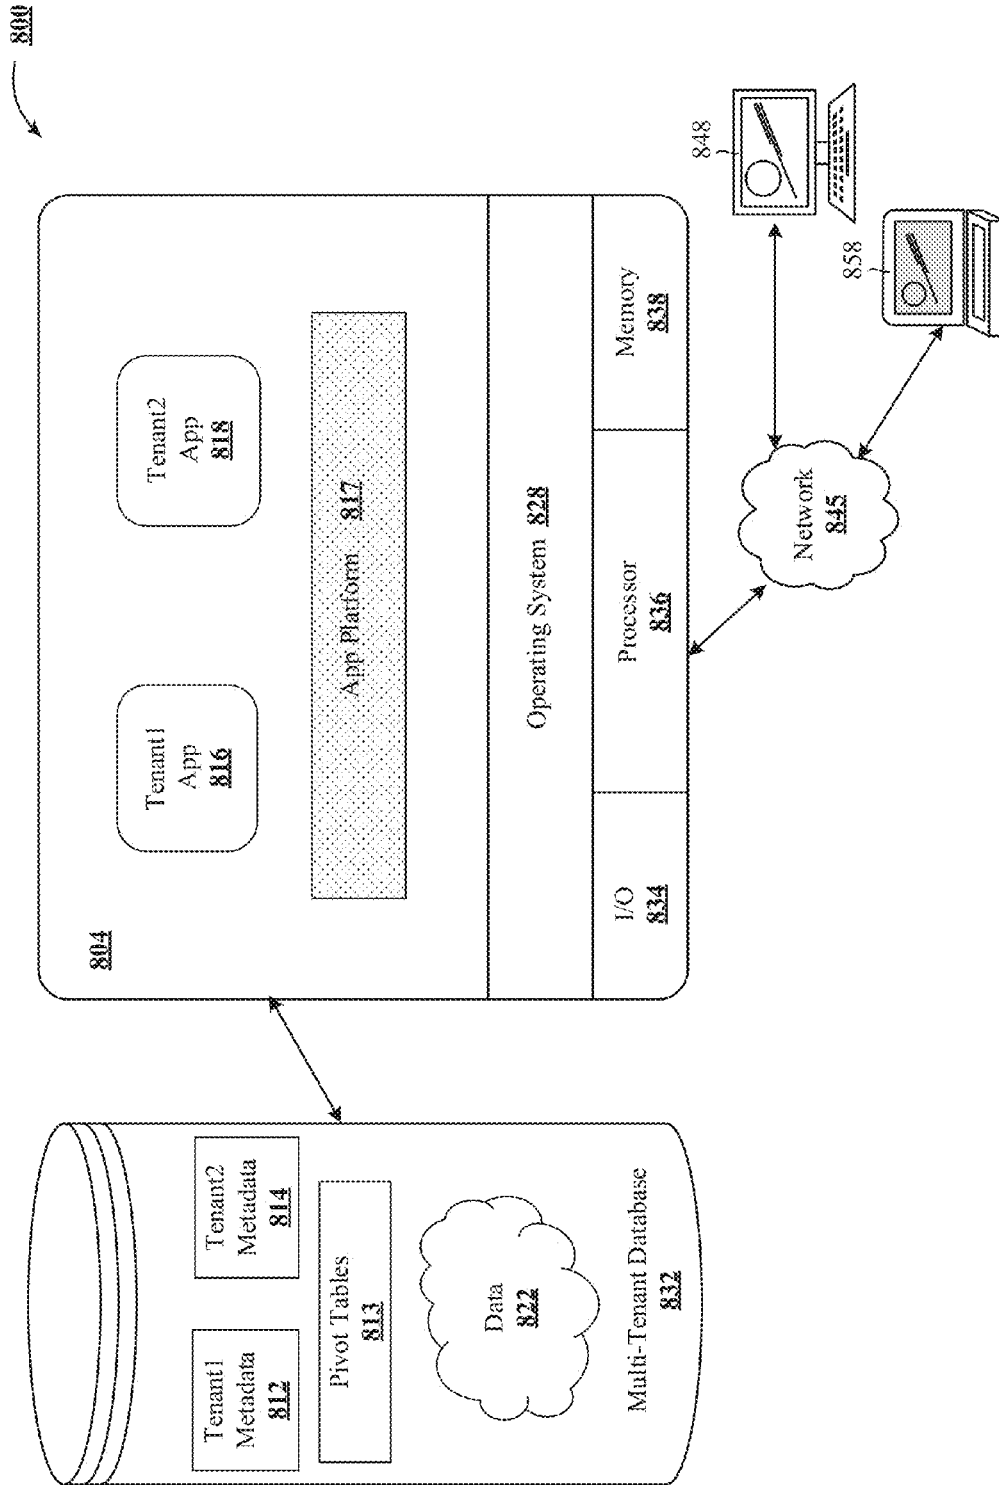


FIG. 8

**PRIORITY-DRIVEN BOXCARRING OF
ACTION REQUESTS FROM
COMPONENT-DRIVEN CLOUD
APPLICATIONS**

RELATED APPLICATION

[0001] This application is related to U.S. Patent Application entitled “DYNAMIC ADJUSTMENT OF BOXCARRING OF ACTION REQUESTS FROM COMPONENT-DRIVEN CLOUD APPLICATIONS” filed concurrently (Atty. Docket No. SALE 1163-1). The related application is hereby incorporated by reference for all purposes.

BACKGROUND

[0002] The subject matter discussed in the background section should not be assumed to be prior art merely as a result of its mention in the background section. Similarly, a problem mentioned in the background section or associated with the subject matter of the background section should not be assumed to have been previously recognized in the prior art. The subject matter in the background section merely represents different approaches, which in and of themselves may also correspond to implementations of the claimed inventions.

[0003] In today’s world, both developers and end users are dealing with large numbers of clients and huge data volumes, popularly referred to as “Big Data”. Web applications that serve and manage millions of Internet users, such as Facebook™ Instagram™, Twitter™, banking websites, or even online retail shops, such as Amazon.com™ or eBay™ are faced with the challenge of delivering applications as fast as possible so that the end users can be provided with a real-time experience.

[0004] Another major contributor the need to deliver fast applications is a concept and paradigm called “Internet of Things” (IoT). IoT is about a pervasive presence in the environment of a variety of things/objects that through wireless and wired connections are able to interact with each other and cooperate with other things/objects to create new applications and services. These applications and services are in areas likes smart cities (regions), smart car and mobility, smart home and assisted living, smart industries, public safety, energy and environmental protection, agriculture and tourism.

[0005] Transaction speed is a critical part of user experience—with more than three billion plus transactions daily for a single very large enterprise. A fast experience encourages users to return to the applications of an enterprise more often and to feel delighted about using them. Currently, there is a need for dynamic optimization of the speed of transactions between client and application servers, for users of component-driven multi-tenant cloud applications delivered as software as a service.

[0006] Therefore, an opportunity arises to speed up performance of a browser page or a client application with multiple components, including improving the network performance, and maintaining a performant user experience. The disclosed technology relates to dynamically adjusting boxcarring of action requests from component-driven cloud applications to speed up performance of a browser or application page with multiple components, in a cloud-based multi-tenant environment. The disclosed technology also

improves perceived page performance of browser and application pages with multiple regions and components.

SUMMARY

[0007] A simplified summary is provided herein to help enable a basic or general understanding of various aspects of exemplary, non-limiting implementations that follow in the more detailed description and the accompanying drawings. This summary is not intended, however, as an extensive or exhaustive overview. Instead, the sole purpose of this summary is to present some concepts related to some exemplary non-limiting implementations in a simplified form as a prelude to the more detailed description of the various implementations that follow.

[0008] Disclosed systems are usable for speeding up performance of a web page with multiple components rendered on a user device. A data consuming application, running on a user device, makes action requests to at least one server and consumes data received from the server responsive to the action requests. A middleware application, also running on the user device, measures network communication latency, adjusts inter-boxcar intervals used to dispatch batches of action requests; and dispatches boxcars of batched action requests. The measured network communication latency is calculated as dispatch-to-completed response time minus server processing time and the server processing time is received from the server for a boxcar of completed responses. The inter-boxcar intervals are adjusted, taking into account at least (a) a number of connections supported between the user device and the server and (b) the measured network communication latency. A browser or app running on the user device, in communication with the middleware application, supports the data consuming application; and the server runs a dispatcher and one or more application programs. The dispatcher receives the transmitted boxcar of action requests, forwards the action requests to the application programs, and returns the boxcar of completed responses; and calculates the server processing time as the difference between when the boxcar of action request is received and when the boxcar of completed responses is ready to be returned, and reports the calculated server processing time back to the middleware application.

[0009] A disclosed method for conserving connections between a browser or app running on a user device and one server includes running a middleware application on the user device, coupled in communication with a data consuming application and a browser or app running on the user device, and also coupled in communication with at least one server through the browser or app. The middleware application receives action requests from the data consuming application qualified by priority labels applied to regions and components in the regions and returns data received from the server responsive to the action requests. The priority labels differentiate among display regions rendered by the data consuming application and at least some of the priority labels further differentiate among components within respective display regions. The middleware application batches the action requests into batches based at least in part on the priority labels, into boxcars segregated by priority label according to a predetermined segregation schedule, and dispatches the boxcars of batched action requests to the server.

[0010] Other aspects and advantages of the technology disclosed can be seen on review of the drawings, the detailed description and the claims, which follow.

BRIEF DESCRIPTION OF THE DRAWINGS

[0011] The included drawings are for illustrative purposes and serve only to provide examples of possible structures and process operations for one or more implementations of this disclosure. These drawings in no way limit any changes in form and detail that may be made by one skilled in the art without departing from the spirit and scope of this disclosure. A more complete understanding of the subject matter may be derived by referring to the detailed description and claims when considered in conjunction with the following figures, wherein like reference numbers refer to similar elements throughout the figures.

[0012] FIG. 1 illustrates an example business information and analytics architecture environment capable of dynamically adjusting boxcarring of action requests from component-driven cloud applications.

[0013] FIG. 2 shows an example page delivery architecture block diagram for browsers and mobile device applications.

[0014] FIG. 3A shows a message diagram for measuring network communication latency between a client and a server.

[0015] FIG. 3B shows a dynamic browser scorecard for deducing network communication latency.

[0016] FIG. 4 shows a pie chart that displays network communication latency metric data that spans a week.

[0017] FIG. 5 shows an example use case of component loading prioritization for optimization of perceived page load time.

[0018] FIG. 6 shows an example code snippet for designing web pages for priority-driven boxcarring via a region-based app editor.

[0019] FIG. 7 shows an overview of the flow for adjusting boxcarring of action requests using the priority level of regions and components in component-driven cloud applications.

[0020] FIG. 8 is a block diagram of an example multi-tenant computer system capable of dynamically adjusting boxcarring of action requests from component-driven cloud applications.

DETAILED DESCRIPTION

Introduction

[0021] The following detailed description is made with reference to the figures. Sample implementations are described to illustrate the technology disclosed, not to limit its scope, which is defined by the claims. Those of ordinary skill in the art will recognize a variety of equivalent variations on the description that follows.

[0022] Application speed is a critical part of user experience. A fast experience will encourage users to return to applications more often and enjoy using the applications. For example, timely delivery of browser pages with “live” data analytics can illuminate puzzling business quandaries. The applications are available to users across smartphones, tablets, laptops and desktops, as well as other form factors.

[0023] Dynamic adjustment of boxcarring of action requests from component-driven cloud applications is

usable to speed up performance of a browser page or a client application with multiple components, to optimize the network performance, and maintain a performant user experience in a cloud-based environment. Dynamically adjusting boxcarring of action requests from component-driven cloud applications is usable to can speed up performance for a plurality of tenants who are managing their digital data on independent software instances, as well as for a group of users who share a common access with a specific set of privileges to a software instance of at least one application.

[0024] To increase perceived web page performance, the goal is to prioritize delivery of specified regions and components, in addition to adjusting the loading of component actions based on the user’s dynamic network communication latency.

[0025] The disclosed dynamic network optimization environment described next is usable to enhance both absolute and perceived page performance for applications, as a critical part of the user experience.

Environment

[0026] FIG. 1 illustrates one environment 100 for implementing dynamic adjustment of boxcarring of action requests from component-driven cloud applications. Environment 100 includes GUI client engine 146 which includes browser-based applications 148, mobile device applications 149 and middleware application 156. Middleware application 156 includes latency measurement engine 166, client boxcar engine 178 and client dispatcher engine 186. Latency measurement engine 166 dynamically determines the latency: the time required for a round trip for a boxcar packet, from the client dispatcher engine 186 of GUI client engine 146 to application server 116, for recent transactions. The network communication latency for each server request is the time between a client dispatching a request and the client receiving a response, excluding the time spent on the server. Client boxcar engine 178 arranges requests for action into a batch, and client dispatcher engine 186 sends each batch of action requests to the server as a single network request—in a boxcar—from the browser-based applications 148 and mobile device applications 149 to application server 166.

[0027] Perceived page performance includes the delivery of prioritized regions and components. GUI client engine 146 also includes client configuration UI 196 that provides a region-based app editor for adding priority rankings. The user can rank the components in a priority order of their choosing, for optimizing perceived page performance for customized pages. Using an app builder interface, a user is able to specify which components are assigned to which regions and also allocate a priority for each of these regions. Each component added within the specified region gets loaded based on the assigned priority for that region.

[0028] Also shown in the environment are analytical read-only data store 102, external data store 142 and multi-tenant software-as-a-service platform 106 which includes application server 116 and server dispatch engine 126. For a very large enterprise, as many as twenty to thirty multi-tenant software-as-a-service platforms 106 can handle user transactions across the globe simultaneously. Server dispatch engine 126 receives the transmitted boxcar of action requests, forwards the action requests to the application programs, and returns the boxcar of completed responses.

[0029] Analytical read-only data store 102 includes read-only datasets, with attributes of multiple users, usable for querying and viewing query results in real time, for large data sets being analyzed. Analytical read-only data store 102 includes datasets extracted from multi-tenant software-as-a-service platform 106 on a batch basis, in one example. The data extracted from large data repositories can be compiled into analytical read-only data store 102, and is usable to create “raw” datasets—read-only data structures for analytics—that can be augmented, transformed, flattened, etc. and published as customer-visible datasets for business entities.

[0030] External data store 142 can include data from sources that are not part of an enterprise’s content management system. Examples of external systems include but are not limited to SAP™, ORACLE E-BUSINESS™, PEOPLESOFT™, NETSUITE™ and WORKDAY™. This data can include customer purchase history, demographics, relationships, and preferences. In one example, data can be represented as comma separated values (CSV) that represent sales quotas for a competitor, provided in a spreadsheet format. Data can be received in other formats—including, but not limited to, other delimiter-separated formats, bit-mapped images, Ogg format containers for different types of multimedia, and proprietary file formats.

[0031] The environment shown in FIG. 1 includes application server 116, which runs all kinds of apps in the cloud—providing software as a service, such as customer relationship management, HR, accounting, and much more. Environment 100 also includes multi-tenant software-as-a-service platform 106, which includes application server 116 and can be of varying types including a workstation, server, computing cluster, blade server, server farm, or any other data processing system or computing device.

[0032] Additionally network 145 can be any network or combination of networks of devices that communicate with one another, and communicate among the data stores, servers, and engines described herein. For example, network 145 can be implemented using one or any combination of a LAN (local area network), WAN (wide area network), telephone network (Public Switched Telephone Network (PSTN), Session Initiation Protocol (SIP), 3G, 4G LTE), wireless network, point-to-point network, star network, token ring network, hub network, WiMAX, Wi-Fi, peer-to-peer connections like Bluetooth, Near Field Communication (NFC), Z-Wave, ZigBee, or other appropriate configuration of data networks, including the Internet. In other implementations, other networks can be used such as an intranet, an extranet, a virtual private network (VPN), a non-TCP/IP based network, any LAN or WAN or the like.

[0033] Analytical read-only data store 102, external data store 142, and multi-tenant software-as-a-service platform 106 can be implemented using a general-purpose distributed memory caching system. In some implementations, data structures can store information from one or more tenants into tables of a common database image to form an on-demand database service (ODDS), which can be implemented in many ways, such as a multi-tenant database system (MTDS). A database image can include one or more database objects. In other implementations, the databases can be relational database management systems (RDBMSs), object oriented database management systems (OODBMSs), distributed file systems (DFS), no-schema database, or any other data storing systems or computing devices. Analytical, read-only databases can implement response

times of under two seconds when searching over twenty million records and compiling aggregate statistics from selected records.

[0034] In some implementations, user computing device 164 can be a personal computer, a laptop computer, tablet computer, smartphone or other mobile computing device, personal digital assistant (PDA), digital image capture devices, and the like. In some implementations, user mobile device 165 can be a tablet computer, smartphone or other mobile computing device, personal digital assistant (PDA), digital image capture devices, and the like.

[0035] GUI client engine 166 can take one of a number of forms, running in a browser or as an application, including user interfaces, dashboard interfaces, engagement consoles, and other interfaces, such as mobile interfaces, tablet interfaces, summary interfaces, or wearable interfaces. In some implementations, it can be hosted on a web-based or cloud-based server in an on premise environment. In one implementation, GUI client engine 166 can be accessed from a browser running on a computing device. The browser can be CHROME™, INTERNET EXPLORER™, FIREFOX™ SAFARI™ OPERA™, ANDROID™, BLACKBERRY™ and the like. In other implementations, dispatcher engine 176 and batch engine 178 can run on a computer desktop application.

[0036] In other implementations, environment 100 for implementing dynamic adjustment of boxcarring of action requests from component-driven cloud applications may not have the same elements or components as those listed above and/or may have other/different elements or components instead of, or in addition to, those listed above, such as a web server and template database. The different elements or components can be combined into single software modules and multiple software modules can run on the same hardware. Communication between component-driven cloud applications and application servers is considered at multiple levels in the data flow for a system; one example is described next.

[0037] FIG. 2 shows an example page delivery architecture which includes users 242 of browser web pages and mobile app pages. Components are self-contained, reusable units of an app that represent a reusable section of the UI, and can range in granularity from a single line of text to an entire app. Users can interact with a data consuming application running on a user device via browser client page 264 includes component 1 244 to component N 254. Each component can request information from the server using “server actions” to make action requests, and each component can employ the use of multiple server actions to deliver the “component view” of the page to the user. The data consuming application running on the user device consumes the data received from the server responsive to the action requests. In one implementation, an enterprise employs a framework that includes a set of prebuilt components that can be assembled and configured to form new regions which are rendered to produce HTML DOM elements within the browser or for a mobile app UI.

[0038] Boxcarring is a method of grouping a set of action requests together and executing them at once. A page can employ multiple XHRs to deliver the “page view” to the user. Boxcars, XHR-1 245 to XHR-N 255, include sets of component action requests $A_1, A_2 \dots A_n$. In one use case, XHR data can be in the form of XML, JSON, HTML or plain text. The response data can be evaluated by client-side

scripting. For example, if a component was formatted as JSON by the web server, the component can convert the received JSON into a client-side data object for delivering the “component view” of the page to the user.

[0039] Middleware application 156, also running on the user device, measures network communication latency, adjusts inter-boxcar intervals used to dispatch batches of action requests, and dispatches boxcars of batched action requests. Multiple actions are boxed into an XMLHttpRequest (XHR). XHR is an API that provides client functionality for transferring data between a client and a server—for sending HTTP or HTTPS requests to a web server and loading the response data received from the server responsive to the action requests back to the client page.

[0040] Continuing with the description of FIG. 2, boxcars XHR-1 245 to XHR-N 255 transport the groups of action requests to application server 116 to API-1 246 to API-n 256, which use SQL 1 248 to SQL-n 258 to execute the component requests. Each action can call multiple different APIs on application server 116; and XHR-1 245 to XHR-N 255 can each include multiple actions, each calling multiple different APIs. A single XHR, with multiple actions batched inside, gets processed by a single application server. Multiple XHRs from the same client on the same page can get serviced by different app servers.

[0041] Server dispatcher engine 126 receives the transmitted boxcar of action requests, forwards the action requests to the application programs, and returns the boxcar of completed responses. Responses are grouped by server dispatcher engine 126 into the same boxcar as those in which they arrived, for return to the client page 264. Server dispatcher engine 126 calculates the server processing time as the difference between when the boxcar of action request is received and when the boxcar of completed responses is ready to be returned, and reports the calculated server processing time back to the middleware application 156.

[0042] Structured query language SQL 1 248 to SQL-n 258 can be used to access and manipulate data in external data store 142. In one use case, application servers use Salesforce Object Query Language (SOQL) to access and manipulate data. In other use cases, application servers use SQL to access and manipulate data in MySQL, SQL Server, Access, Oracle, Sybase, DB2, and other database systems.

[0043] Because most responses to component action requests are small—in the 2 kB to 5 kB range, based on analyzed experimental data—bandwidth typically has little effect on the speed of transactions for users. Next we describe the measurement of network latency for improving application performance.

[0044] The disclosed process of improving the application performance for a user includes dynamically measuring and characterizing the network latency for end users. Measured network latency is calculated as dispatch-to-completed response time minus server processing time, and the server processing time is received from the server for a boxcar of completed responses. In one application context, network latency for each server request is the time spent between the client making a request and the client receiving a response, excluding the time spent on the server.

[0045] FIG. 3A shows a message diagram 300 between client and server, with labeled time measurements usable for measuring network latency. In FIG. 3A, client 344 sends XHR request packet 356 at request start time 354. Upon

receipt of the packet, server 348 begins processing the contents, at server start time 358. When the server actions are complete, server 348 sends XHR response packet 366 to client 344. The XHR response packet 366 arrives at the client at response end time 364.

[0046] Total server time equals server end time minus server start time. Instrumenting the server 348 to capture the server start time 358 and server end time 368 makes it possible to measure the total server time. The data received by instrumenting the client app to record the time at which each XHR request packet 356 is sent and the time when each corresponding XHR response packet 366 is received is usable to calculate the measured server latency.

[0047] For an implementation of the disclosed technology, server metrics are appended to the server response and sent to the client 344, and the calculation of total network latency equals the response end time 364 minus the request start time 354 minus the total server time described earlier in this paragraph.

[0048] Determining the optimal path for component action requests for a user includes deducing the number of requests from components to boxcar together for the trip to the application server. Middleware application 156 adjusts the inter-boxcar intervals, taking into account the number of connections supported between the user device and the server and the measured network communication latency.

[0049] The server dispatcher engine 126 calculates the server processing time as the difference between when the boxcar of action requests is received and when the boxcar of completed responses is ready to be returned, and reports the calculated server processing time back to the middleware application 156.

[0050] In another implementation, to determine the optimal path, the total network latency can be calculated on the client and aggregated over a significant sample size of recorded latencies, for each action request, to characterize the user’s latency. In addition, this latency information can also be beamed to the server and logged. For some use cases, this information can be utilized on the server for future optimizations.

[0051] In some use cases, users 242 can configure their preferences for dynamic adjustment of boxcarring of action requests from component-driven cloud applications to be based on either absolute page performance or perceived page performance. Perceived page load time is the time required to fetch the components on the client page 264 that make the page partially usable for the user. This will be discussed further later.

[0052] Absolute page load time is the time required to fetch all the components on the page from the application server 116 and render them on the client page 264. In this configuration, the goal is to deliver the optimal absolute page performance based on the user’s network communication latency.

[0053] The disclosed technology includes speeding up performance of a page with multiple components, in a cloud-based environment, by delivering the most optimal absolute page performance. Based on the dynamic network latency of the user, the grading engine 158, in latency engine 146, categorizes the user’s transactions and assigns the user to one of several groupings. In one use case, the categories for grouping the users include low latency, average latency, and high latency—listed next, in milliseconds (ms).

[0054] Low latency <100 ms

[0055] Average latency >100 ms and <300 ms

[0056] High latency >300 ms

[0057] The number of categories and the groupings for the user's transactions can adapt over time, with the values determined empirically. The system adjusts latency categories according to feedback received as a browser or mobile device changes network connections or the network conditions change. Numbers for delay can also programmatically change or be tuned. That is, hard-coded numbers can change, with more buckets or different values learned with experience. In another use case, the number of latency categories for grouping the users may be different, based on the results of analyzing transaction data for over a period of time described by a number of transmissions, seconds, minutes, hours, days, weeks or months.

[0058] When the dynamic network latency falls in the low latency category of less than one hundred ms, action requests can be sent in parallel—utilizing the connections offered by the browser effectively. That is, each action request can be sent as soon as it is received. The number of nearly-concurrent connections used to service components of the web page can be rationed. Middleware application **156** determines the browser or app through which it is coupled in communication with the server and looks up the number of connections supported by the respective browser or app.

[0059] Most modern browsers support between six and thirteen concurrent connections per hostname (number of http/1.1 connections that can be opened for the browser), so a default number of nearly-concurrent connections can be initialized to a value of six and the number of available connections can take into consideration the network latency. In one implementation, Browserscope—a community-driven project for profiling web browsers—is usable for discerning the number of concurrent connections per hostname for a client server. Browserscope tracks browser functionality, and gathers latency and other test results from users “in the wild” <May 9, 2016: <http://www.browserscope.org/>>.

[0060] FIG. 3B shows an example dynamic Browserscope browser scorecard for deducing network latency. Users can choose to view results for only desktop browsers, only mobile device browsers, results for a particular family of browsers, or other optional groupings. Network **304** shows tests that address network performance issues, including connections per hostname **314**. The column labeled maximum connections **315** refers to the number of connections that the browser will open across all hostnames, which has a typical upper limit of 60.

[0061] In the low latency scenario, each component, from component **1 244** to component **N 254** of browser/client page **264**, makes its own server request directly without enqueueing it into a global action queue. This saves the queuing time and since each action is executed individually on the server, the server time is also optimized. This ensures that the server responses come, in the fastest way, back to the client—thereby delivering the most optimal network performance, based on the user's good latency.

[0062] Most users fall in the average latency category—with a measured latency of between 100 ms and 300 ms. For a composite component page, each component requests information from the server by enqueueing its actions into a global action queue. Each component is oblivious to other

components on the same page and can make their action requests at any time in their respective component lifecycle.

[0063] Client boxcar engine **178** collects all actions that are requested from the composite components around the same time and batches them into a single network request action boxcar XHR-N **255**. Client dispatcher engine **186** sends the batched network request to the application server **116**. Upon receiving the response boxcar from application server **116**, the client **264** decomposes the respective responses to each of the constituent components. Instead of making multiple XHR requests—XHR-1 **245** to XHR-N **255** to the server—one for each of the requested component actions, the disclosed technology dynamically adjusts box-carring of action requests from component-driven cloud applications—reducing the number of network requests to application server **116**. In the average case, the batching of action requests to be routed to the server improves network utilization.

[0064] The process for handling component requests includes measuring the network latency from initial XHR requests and establishing a network latency category, and dynamically calculating the queue wait time based on the established network latency category. The process includes continuously gathering actions into a collector within the queue wait time window, batching the actions, $A_1, A_2 \dots A_n$, **245** collected during the calculated queue wait time. The process also includes creating XHR requests **255** with these enqueued action requests $A_1, A_2 \dots A_n$, **245** and sending the boxcars to the application server **162**. The process includes adapting the queue wait time based on measured network latency feedback and action/XHRs in-flight—sent but not yet received, and in-queue. The code snippet listed next shows one implementation of the process. The following function for enqueueing actions gets activated from within an event loop for a component or region.

```
Aura.ActionService = function Aura.ActionService ( ) {
  this.NOOP = function ( ) { };
  var auraXHR = new Aura.Services.AuraClientService$.AuraXHR ( );
  var XHRLatency = new
  Aura.Services.AuraClientService$.XHRLatency ( );
  this.availableXHRs = [ auraXHR ];
  this.allXHRs = [ auraXHR ];
  this.actionStoreMap = { };
  this.collector = undefined;
  this.setQueueSize(this.Browser.MaxConnectionsPerHost);
  ...
  AuraClientService.prototype.enqueueAction = function(action) {
    $.A.assert(!$.A.util.isUndefinedOrNull(action), "EnqueueAction()
    cannot be called on an undefined or null action.");
    $.A.assert($.A.util.isAction(action), "Cannot call EnqueueAction()
    with a non Action parameter.");
    this.actionsQueued.push(action);
  };
};
```

[0065] The next code snippet shows the function for getting an available XHR.

```
AuraClientService.prototype.getAvailableXHR = function ( ) {
  var auraXHR = this.availableXHRs.pop ( );
  return auraXHR;
};
```

[0066] The next code snippet shows the function for releasing an XHR back into the pool.

```
AuraClientService.prototype.releaseXHR = function(auraXHR) {
  auraXHR.reset( );
  this.availableXHRs.push(auraXHR);
  if (this.inFlightXHRs( ) === 0) {
    this.processXHRIdleQueue( );
  }
};
```

[0067] The following code snippet shows the entry point for processing actions: creating a collector, and parceling out the action request handling. After this, action requests will be either getting values from storage or will be executed on the server, and the client action requests will be queued to be executed in the order set via `setTimeout`, giving server actions entry points to collect.

```
AuraClientService.prototype.process = function( ) {
  if (this.collector) {
    return;
  }
  this.collector = new
  Aura.Services.AuraClientService$AuraActionCollector( );
  this.continueProcessing( );
};
```

[0068] The process includes dividing actions into client actions and server actions. Server actions are further divided into stored and non-stored actions, and the following code snippet also protects against server actions collecting early. In the example, `actionsQueued` refers to the queue of actions that have yet to be processed.

```
AuraClientService.prototype.continueProcessing = function( ) {
  var i;
  var index = 0;
  var action;
  var actionList;
  this.collector.actionsToCollect += 1;
  actionList = this.actionsQueued;
  this.actionsQueued = [ ];
  for (i = 0; i < actionList.length; i++) {
    action = actionList[i];
    try {
      if (action.abortIfComponentInvalid(true)) {
        // action already aborted.
        // this will only occur if the component is no longer valid.
        continue;
      }
      if (action.getDef( ).isServerAction( )) {
        this.collector.actionsToCollect += 1;
        this.collector.collected[index] = undefined;
        this.collector.collecting[index] = action;
        this.collectServerAction(action, index);
        index += 1;
      }
    } catch (e) {
      var errorWrapper = new $A.auraError(null, e);
      errorWrapper.action = action;
      $A.logger.reportError(errorWrapper);
    }
  }
  this.collector.actionsToCollect -= 1;
};
```

[0069] Next in the process, the following code snippet shows collecting a single action into the list, finishing the collection process, and sending the XHR.

```
AuraClientService.prototype.collectServerAction =
function(action, index) {
  this.collector.collected[index] = action;
  this.collector.actionsToCollect -= 1;
  this.finishCollection( );
};
...
AuraClientService.prototype.finishCollection = function( ) {
  if (this.collector.actionsToCollect !== 0 ||
  this.collector.clientActions.length) {
    return;
  }
  if (this.collector.actionsCompleted) {
    this.fireDoneWaiting( );
  }
};
```

[0070] If a user has very high network latency greater than 300 ms, more conservative choices are implemented for network requests to improve the user's network performance. The time window for collecting component action requests can be increased to 300 ms to allow for more actions to be boxcarred into the same network request. This helps to mitigate the high latency cost of a user's slow network. In most scenarios, this reduces the number of network requests and thereby delivers a more optimal network performance within the user's bad latency constraints.

[0071] The process includes boxcarring into a single XHR composite actions enqueued around the same time, and for measured latency values, waiting for an additional 100 ms and then sending possibly additional actions via a single boxcar XHR, as shown in the following code snippet shows. If the measured latency increases, then the process includes waiting an additional 300 ms, and sending possibly additional actions as a single boxcar XHR.

```
var i = 0;
var length;
var collected = this.collector.collected;
this.collector.collected = [ ];
for (i = 0, length = collected.length; i < length; i++) {
  if (collected[i]) {
    this.actionsDeferred.push(collected[i]);
  }
}
if (this.actionsQueued.length) {
  this.continueProcessing( );
  return;
}
if (this.actionsDeferred.length) {
  if (XHR.Latency.latencyCategory === "average") {
    if (XHR.Latency.increasedWait) {
      XHR.Latency.increasedWait = false;
      this.sendActionXHRs(false);
    } else {
      setTimeout(function( ) {
        XHR.Latency.increasedWait = true;
        this.continueProcessing( );
        return;
      }, 100
    );
  }
} else if (XHR.Latency.latencyCategory === "high") {
  if (XHR.Latency.increasedWait) {
    XHR.Latency.increasedWait = false;
    this.sendActionXHRs(false);
  } else {
    setTimeout(function( ) {
      XHR.Latency.increasedWait = true;
      this.continueProcessing( );
      return;
    }, 300
  );
}
```

-continued

```

    );
  }
} else if(XHRLatency.latencyCategory === "low") {
  this.addActionXHRs(true);
}
}
}

```

[0072] The process includes sending requests with the maximum parallelism possible, in cases in which the measured latency is low, and ensuring that all actions get sent, marked as duplicates, or put into a deferred queue. Callback for an XHR for a set of actions includes correctly dealing with the case of interrupted communications, and aborts of requests. The next code snippet shows code for setting the XHR receive time, and a segment of the code used for catching errors, as needed.

```

AuraClientService.prototype.receive = function(auraXHR, timedOut) {
  auraXHR.XHREndTime = Date.now();
  var responseMessage;
  this.auraStack.push("AuraClientService$receive");
  try {
    responseMessage = this.decode(auraXHR.request, false, timedOut);
    if (responseMessage["status"] === "SUCCESS") {
      this.processResponses(auraXHR, responseMessage["message"]);
      this.processXHRLatency(auraXHR, responseMessage["perf"]);
    } else if (responseMessage["status"] === "INCOMPLETE") {
      this.processIncompletes(auraXHR);
    } else if (responseMessage["status"] === "ERROR") {
      this.processErrors(auraXHR, responseMessage["message"]);
    }
    this.fireDoneWaiting();
  } catch (e) {
    throw (e instanceof $A.auraError) ? e : new
    $A.auraError("AuraClientService.receive action callback failed", e);
  } finally {
    this.auraStack.pop();
    this.releaseXHR(auraXHR);
    this.process();
  }
  return responseMessage;
};

```

[0073] In this example case, if the latency is greater than 300 ms, the boxcarring is in the "high" category, if the latency is between 100 ms and 300 ms, the latency is in the average category, and if the latency is less than 100 ms, the latency category is "low". The next code snippet shows the receiving of the server start and end timestamps from the server performance instrumentation, and using the values to calculate network latency from the XHR to establish the network latency category.

```

AuraClientService.prototype.processXHRLatency = function(auraXHR) {
  var latency;
  var perf = responseMessage["perf"];
  if (perf) {
    auraXHR.ServerStartTime = perf.serverStartTime;
    auraXHR.ServerEndTime = perf.serverEndTime;
    var_totalServerTime = auraXHR.ServerEndTime -
    auraXHR.ServerStartTime;
    var_totalXHRTIME = auraXHR.XHREndTime -
    auraXHR.XHRStartTime;
    latency = _totalXHRTIME - _totalServerTime;
    XHRLatency.lastKnownLatency = latency;
    if (latency >= 300) {
      XHRLatency.latencyCategory = "high";
    }
  }
}

```

-continued

```

} else if (latency < 100) {
  XHRLatency.latencyCategory = "low";
} else {
  XHRLatency.latencyCategory = "average";
}
}
};

```

[0074] The process and code snippets described above for dynamic adjustment of boxcarring of action requests from component-driven cloud applications for a user show an implementation for dynamically measuring and characterizing the network latency for end users and for implementing boxcars, based on the measured network latency.

[0075] FIG. 4 shows a pie chart graph of network latency metrics **400** for a production page. This example set of metric measurement results covers a week of time, for ten requests per page. The graph **400** shows that 61.6% of requests **425** incurred a network latency of less than one hundred ms, so less than one second to transmit ten requests and receive ten component action requests. Only 0.4% of action requests **424** incurred a network latency of greater than five seconds, translating to more than fifty seconds—close to a minute—for ten requests.

Perceived Page Performance

[0076] For network utilization improvements based on perceived page performance, users **242** can configure their loading preferences to specify the order for the loading of regions for a page—that is, what components in what regions get loaded in what order. The goal is to deliver the most optimal perceived performance based on the user's network latency. The user perception can be based on configurable options.

[0077] Component loading based on a user's viewing choices can include an algorithm that identifies rules for choosing order, or can include an explicit loading order for components that specifies which components to prioritize in the component load order. For example, similar to the way one reads a book, the focus of a user's gaze moves from left to right, and then top to bottom. To give the user the perceived sense of responsiveness when viewing the book, the rules for prioritizing the component load order can specify starting at the top left and then moving across the displayed page, before moving toward the bottom right corner of the page.

[0078] FIG. 5 shows an example use case of component loading prioritization based on the typical direction of a user's gaze. The example page **500** includes four components: sales rep at a glance **501**, account insights **502**, quarterly performance **503** and assistant **504**, positioned as shown on the page. Rules for prioritizing the component load order can be expressed: Load sales rep at a glance **501**, then account insights **502**, then quarterly performance **503** and load assistant **504** last. In some use cases, to minimize perceived page load time latency, network latency and boxcarring of action requests can occur separately for the components within the region being loaded.

[0079] In another use case, component loading prioritization can be based on user-specified ranking of constituent components of a page—prioritizing one component over another. A client configuration UI can provide a list of all the components on the page to the user, in the page setup, and

the user can rank the components in a priority order of their choosing. This is useful in scenarios in which the user has a personal preference for a specific component to load as soon as possible, since that component is the most critical piece of information for them on the page. In one example, a user may want to view assistant **504** first, while the remaining regions of their page are loading.

[0080] When a new page is being designed or a page design is being optimized, the designer can use a region-based app editor for adding priority rankings that rank the components in a priority order of the user's choosing, for enhancing perceived page performance for customized pages. Using a region-based app builder interface, a user is allowed to specify which components are assigned to which regions and also allocate a priority for each of these regions. Each component added within the specified region will be loaded based on the assigned priority for that region.

[0081] FIG. 6 includes a snippet of code **600** for a region-based app editor, showing example priority labels for regions. In an iFrame, the body includes a series of classes that represent regions of a web page. In the example snippet of code **600**, the "top" region **632** is priority one **645** and the bottom left region **652** is priority two **665**. Similarly, additional classes can define regions such as bottom right with a priority of three and a region of column side with a priority of four (not shown).

[0082] Data consuming applications—browser-based applications **148** or mobile device applications **149** running on a user device **164**, **165**, make action requests qualified by the priority labels applied to regions and components in the regions and consume data received from at least one application server **116** responsive to the action requests. The priority labels differentiate among display regions rendered by the data consuming application and at least some of the priority labels further differentiate among components within respective display regions.

[0083] Middleware application **156**, also running on the user device, batches the action requests into batches based at least in part on the priority labels **645**, **665**, and dispatches boxcars of batched action requests, in which the priority labels differentiate among display regions rendered by the data consuming application and at least some of the priority labels further differentiate among components within respective display regions. The boxcars are segregated by priority label according to a predetermined segregation schedule; and the action requests that are qualified by higher priority labels are dispatched ahead of the action requests with lower priority labels. The server dispatcher engine **126** receives the transmitted boxcar of action requests, forwards the action requests to the application programs, collects responses from the application programs, and returns a responsive boxcar of completed responses. The server dispatcher engine **126** calculates the server processing time as the difference between when the boxcar of action request is received and when the boxcar of completed responses is ready to be returned, and reports the calculated server processing time back to the middleware application **156**.

[0084] Region load order configuration is very useful for enterprise customers who need to optimize customized pages for enterprise applications in which their clients can add any number of new and third party components of their liking to a page. Predicting the runtime performance of a completely re-architected page can be challenging, and this

client configuration UI feature empowers the designer to fine tune the page load time for a page, independently.

System Flow

[0085] FIG. 7 illustrates a flowchart of one implementation **700** of implementing a system for priority-driven boxcarring of action requests from component-driven cloud applications. Flowchart **700** can be implemented at least partially with a database system, e.g., by one or more processors configured to receive or retrieve information, process the information, store results, and transmit the results. Other implementations may perform the steps in different orders and/or with different, fewer or additional steps than the ones illustrated in FIG. 7. The actions described below can be subdivided into more steps or combined into fewer steps to carry out the method described using a different number or arrangement of steps.

[0086] At action **710**, run a middleware application on the user device, coupled in communication with a data consuming application and a browser or app running on the user device, and also coupled in communication with at least one server through the browser or app.

[0087] At action **715**, the middleware application receives action requests from the data consuming application qualified by priority labels applied to regions and components in the regions and returns data received from the server responsive to the action requests.

[0088] At action **720**, the priority labels differentiate among display regions rendered by the data consuming application and at least some of the priority labels further differentiate among components within respective display regions.

[0089] At action **725**, the middleware application batches the action requests into batches based at least in part on the priority labels, into boxcars segregated by priority label according to a predetermined segregation schedule, and dispatching the boxcars of batched action requests to the server.

[0090] The technology disclosed can be implemented in the context of any computer-implemented system including a database system, a multi-tenant environment, or the like. Moreover, this technology can be implemented using two or more separate and distinct computer-implemented systems that cooperate and communicate with one another. This technology can be implemented in numerous ways, including as a process, a method, an apparatus, a system, a device, a computer readable medium such as a computer readable storage medium that stores computer readable instructions or computer program code, or as a computer program product comprising a computer usable medium having a computer readable program code embodied therein.

Integration

[0091] FIG. 8 presents a block diagram of an exemplary multi-tenant system **800** suitable for implementing dynamic adjustment of boxcarring of action requests from component-driven cloud applications in environment **100** of FIG. 1. In general, the illustrated system **800** of FIG. 8 includes a server **804** that dynamically creates and supports virtual applications **816** and **818**, based upon data **822** from a common database **832** that is shared between multiple tenants, alternatively referred to herein as a "multi-tenant database". Data and services generated by the virtual appli-

cations **816** and **818**, including GUI clients, are provided via a network **845** to any number of client devices **848** or **858**, as desired.

[0092] As used herein, a “tenant” or an “organization” refers to a group of one or more users that shares access to common subset of the data within the multi-tenant database **832**. In this regard, each tenant includes one or more users associated with, assigned to, or otherwise belonging to that respective tenant. Stated another way, each respective user within the multi-tenant system **800** is associated with, assigned to, or otherwise belongs to a particular tenant of the plurality of tenants supported by the multi-tenant system **800**. Tenants may represent users, user departments, work or legal organizations, and/or any other entities that maintain data for particular sets of users within the multi-tenant system **800**. Although multiple tenants may share access to the server **804** and the database **832**, the particular data and services provided from the server **804** to each tenant can be securely isolated from those provided to other tenants. The multi-tenant architecture therefore allows different sets of users to share functionality and hardware resources without necessarily sharing any of the data **822** belonging to or otherwise associated with other tenants.

[0093] The multi-tenant database **832** is any sort of repository or other data storage system capable of storing and managing the data **822** associated with any number of tenants. The database **832** may be implemented using any type of conventional database server hardware. In various implementations, the database **832** shares processing hardware with the server **804**. In other implementations, the database **832** is implemented using separate physical and/or virtual database server hardware that communicates with the server **804** to perform the various functions described herein. The multi-tenant database **832** may alternatively be referred to herein as an on-demand database, in that the multi-tenant database **832** provides (or is available to provide) data at run-time to on-demand virtual applications **816** or **818** generated by the application platform **817**, with tenant1 metadata **812** and tenant2 metadata **814** securely isolated.

[0094] In practice, the data **822** may be organized and formatted in any manner to support the application platform **817**. In various implementations, conventional data relationships are established using any number of pivot tables **813** that establish indexing, uniqueness, relationships between entities, and/or other aspects of conventional database organization as desired.

[0095] The server **804** is implemented using one or more actual and/or virtual computing systems that collectively provide the dynamic application platform **817** for generating the virtual applications. For example, the server **804** may be implemented using a cluster of actual and/or virtual servers operating in conjunction with each other, typically in association with conventional network communications, cluster management, load balancing and other features as appropriate. The server **804** operates with any sort of conventional processing hardware such as a processor **836**, memory **838**, input/output features **834** and the like. The input/output **834** generally represent the interface(s) to networks (e.g., to the network **845**, or any other local area, wide area or other network), mass storage, display devices, data entry devices and/or the like. User interface input devices **834** can include a keyboard; pointing devices such as a mouse, trackball, touchpad, or graphics tablet; a scanner; a touch screen incorporated into the display; audio input devices such as

voice recognition systems and microphones; and other types of input devices. In general, use of the term “input device” is intended to include possible types of devices and ways to input information into computer system **817**.

[0096] User interface output devices can include a display subsystem, a printer, a fax machine, or non-visual displays such as audio output devices. The display subsystem can include a cathode ray tube (CRT), a flat-panel device such as a liquid crystal display (LCD), a projection device, or some other mechanism for creating a visible image. The display subsystem can also provide a non-visual display such as audio output devices. In general, use of the term “output device” is intended to include all possible types of devices and ways to output information from processor **836** to the user or to another machine or computer system.

[0097] The processor **836** may be implemented using any suitable processing system, such as one or more processors, controllers, microprocessors, microcontrollers, processing cores and/or other computing resources spread across any number of distributed or integrated systems, including any number of “cloud-based” or other virtual systems. The memory **838** represents any non-transitory short or long term storage or other computer-readable media capable of storing programming instructions for execution on the processor **836**, including any sort of random access memory (RAM), read only memory (ROM), flash memory, magnetic or optical mass storage, and/or the like. The computer-executable programming instructions, when read and executed by the server **804** and/or processor **836**, cause the server **804** and/or processor **836** to create, generate, or otherwise facilitate the application platform **817** and/or virtual applications **816** and **818**, and perform one or more additional tasks, operations, functions, and/or processes described herein. It should be noted that the memory **838** represents one suitable implementation of such computer-readable media, and alternatively or additionally, the server **804** could receive and cooperate with external computer-readable media that is realized as a portable or mobile component or application platform, e.g., a portable hard drive, a USB flash drive, an optical disc, or the like.

[0098] The application platform **817** is any sort of software application or other data processing engine that generates the virtual applications **816** and **818** that provide data and/or services to the client devices **848** and **858**. In a typical implementation, the application platform **817** gains access to processing resources, communications interfaces and other features of the processing hardware using any sort of conventional or proprietary operating system **828**. The virtual applications **816** and **818** are typically generated at run-time in response to input received from the client devices **848** and **858**.

[0099] With continued reference to FIG. **8**, the data and services provided by the server **804** can be retrieved using any sort of personal computer, mobile telephone, tablet or other network-enabled client device **848** or **858** on the network **845**. In an exemplary implementation, the client device **848** or **858** includes a display device, such as a monitor, screen, or another conventional electronic display capable of graphically presenting data and/or information retrieved from the multi-tenant database **832**.

[0100] In some implementations, network(s) **845** can be any one or any combination of Local Area Network (LAN), Wide Area Network (WAN), WiMAX, Wi-Fi, telephone network, wireless network, point-to-point network, star net-

work, token ring network, hub network, mesh network, peer-to-peer connections like Bluetooth, Near Field Communication (NFC), Z-Wave, ZigBee, or other appropriate configuration of data networks, including the Internet.

[0101] The foregoing description is merely illustrative in nature and is not intended to limit the implementations of the subject matter or the application and uses of such implementations. Furthermore, there is no intention to be bound by any expressed or implied theory presented in the technical field, background, or the detailed description. As used herein, the word “exemplary” means “serving as an example, instance, or illustration.” Any implementation described herein as exemplary is not necessarily to be construed as preferred or advantageous over other implementations, and the exemplary implementations described herein are not intended to limit the scope or applicability of the subject matter in any way.

[0102] The technology disclosed can be implemented in the context of any computer-implemented system including a database system, a multi-tenant environment, or a relational database implementation like an ORACLE™ compatible database implementation, an IBM DB2 Enterprise Server compatible relational database implementation, a MySQL or PostgreSQL compatible relational database implementation or a Microsoft SQL Server compatible relational database implementation or a NoSQL non-relational database implementation such as a Vampire™ compatible non-relational database implementation, an Apache Cassandra™ compatible non-relational database implementation, a BigTable compatible non-relational database implementation or an HBase or DynamoDB compatible non-relational database implementation.

[0103] Moreover, the technology disclosed can be implemented using two or more separate and distinct computer-implemented systems that cooperate and communicate with one another. The technology disclosed can be implemented in numerous ways, including as a process, a method, an apparatus, a system, a device, a computer readable medium such as a computer readable storage medium that stores computer readable instructions or computer program code, or as a computer program product comprising a computer usable medium having a computer readable program code embodied therein.

Particular Implementations

[0104] Some particular implementations and features are described in the following discussion.

[0105] In one implementation, a disclosed system includes a user device and a server, coupled in communication, including a data consuming application, running on a user device, that makes action requests qualified by priority labels applied to regions and components in the regions and that consumes data received from at least one server responsive to the action requests. A middleware application, also running on the user device, batches the action requests into batches based at least in part on the priority labels, and dispatches boxcars of batched action requests. The priority labels differentiate among display regions rendered by the data consuming application and at least some of the priority labels further differentiate among components within respective display regions. The boxcars are segregated by priority label according to a predetermined segregation schedule; and the action requests that are qualified by higher priority labels are dispatched ahead of the action requests

with lower priority labels. The disclosed system also includes a browser or app running on the user device, in communication with the middleware application, that supports the data consuming application; and the server, running a dispatcher and one or more application programs; and the dispatcher receives the boxcar of action requests, forwards the action requests to the application programs, collects responses from the application programs, and returns a responsive boxcar of completed responses.

[0106] For some implementations of the disclosed system, the priority labels are set using an attribute of a DIV html tag. The disclosed system further includes the middleware application, further implementing measurement of network communication latency, setting of inter-boxcar intervals used to dispatch the batches of action requests, wherein the measured network communication latency is calculated as dispatch-to-completed response time minus server processing time and the server processing time is received from the server for a boxcar of completed responses. The inter-boxcar intervals are adjusted taking into account at least (a) a number of connections supported between the user device and the server and (b) the measured network communication latency.

[0107] In one implementation of the disclosed system, the middleware application repeatedly measures network communication latency and readjusts the inter-boxcar intervals used to dispatch batches of action requests based on updated measurements, and adjusts the inter-boxcar intervals used to time when to dispatch batches of action requests to 250-350 ms when the measured network communication latency exceeds 300 ms. The middleware application also can identify the browser or app through which it is coupled in communication with the server and looks up the number of connections supported by the identified browser or app. For the disclosed system, the middleware application receives the boxcar of completed responses grouped to respond to a respective boxcar of action requests.

[0108] For one implementation, a user device runs a middleware application, coupled in communication with a data consuming application and a browser or app running on the user device, and also coupled in communication with at least one server through the browser or app; wherein the middleware application receives action requests from the data consuming application qualified by priority labels applied to regions and components in the regions and returns data received from the server responsive to the action requests. The priority labels differentiate among display regions rendered by the data consuming application and at least some of the priority labels further differentiate among components within respective display regions. The action requests that are qualified by higher priority labels are dispatched ahead of the action requests with lower priority labels. The middleware application batches the action requests into batches based at least in part on the priority labels, into boxcars segregated by priority label according to a predetermined segregation schedule, and dispatches the boxcars of batched action requests. The middleware application further measures network communication latency and adjusts inter-boxcar intervals used to dispatch the batches of action requests, and calculates the measured network communication latency as dispatch-to-completed response time minus server processing time, using the server processing time received from the server for a boxcar of completed responses. The inter-boxcar intervals are adjusted, taking

into account at least (a) a number of connections supported between the user device and the server and (b) the measured network communication latency.

[0109] For some implementations of the user device, the middleware application repeatedly measures network communication latency and readjusts the inter-boxcar intervals used to dispatch batches of action requests based on updated measurements. The middleware application adjusts the inter-boxcar intervals used to time when to dispatch batches of action requests to 250-350 ms when the measured network communication latency exceeds 300 ms, and the middleware application identifies the browser or app through which it is coupled in communication with the server and looks up the number of connections supported by the identified browser or app. The middleware application receives the boxcar of completed responses grouped to respond to a respective boxcar of action requests.

[0110] In another implementation, a disclosed server runs a dispatching application, coupled in communication with a multiplicity (100 or more) of middleware applications running on respective user devices and one or more application programs running on the server. The dispatcher receives boxcarred batches of action requests from the middleware applications in priority-based batches, forwards the action requests to the application programs, and returns responsive boxcars of completed responses; and calculates a server processing time as a difference between when a particular boxcar of action requests is received and when a respective boxcar of completed responses is ready to be returned and further reports the calculated server processing time back to the middleware application for the respective boxcar of completed responses. The dispatcher further identifies respective application programs that will respond to the action requests received in a particular boxcar from one of the middleware applications and dispatches the action requests from the particular boxcar to the responsive application programs. In some implementations of the disclosed server, the dispatcher postpones returning responses to action requests from a particular boxcar until it has received (or timed out on) responses to all of the action requests in the particular boxcar.

[0111] In yet another implementation, a disclosed method conserves connections between a browser or app running on a user device and at least one server, including running a middleware application on the user device, coupled in communication with a data consuming application and a browser or app running on the user device, and also coupled in communication with at least one server through the browser or app. The middleware application receives action requests from the data consuming application qualified by priority labels applied to regions and components in the regions and returns data received from the server responsive to the action requests. The priority labels differentiate among display regions rendered by the data consuming application and at least some of the priority labels further differentiate among components within respective display regions; and the middleware application batches the action requests into batches based at least in part on the priority labels, into boxcars segregated by priority label according to a predetermined segregation schedule, and dispatches the boxcars of batched action requests to the server. The middleware application measures network communication latency and adjusting inter-boxcar intervals used to dispatch batches of action requests, and calculates the measured network com-

munication latency as dispatch-to-completed response time minus server processing time, using the server processing time received from the server for a boxcar of completed responses. The middleware application also adjusts the inter-boxcar intervals taking into account at least (a) a number of connections supported between the user device and the server and (b) the measured network communication latency. The disclosed method further includes repeatedly measuring network communication latency and readjusting the inter-boxcar intervals used to dispatch batches of action requests based on updated measurements.

[0112] For some implementations, the disclosed method further includes adjusting the inter-boxcar intervals used to time when to dispatch batches of action requests to 250-350 ms when the measured network communication latency exceeds 300 ms. The middleware application identifies the browser or app through which it is coupled in communication with the server and looking up the number of connections supported by the identified browser or app. The middleware application receives the boxcar of completed responses grouped to respond to a respective boxcar of action requests.

[0113] In one implementation a disclosed method conserves connections between a browser or app running on a user device and at least one server, including a dispatcher running on the server and coupled in communication with a multiplicity of middleware applications running on respective user devices and one or more application programs running on the server. The dispatcher receives boxcarred batches of action requests in priority-based batches from the middleware applications, forwards the action requests to the application programs, and returning responsive boxcars of completed responses; and calculates a server processing time as a difference between when a particular boxcar of action requests is received and when a respective boxcar of completed responses is ready to be returned and further reporting the calculated server processing time back to the middleware application for the respective boxcar of completed responses. The dispatcher identifies respective application programs that will respond to the action requests received in a particular boxcar from one of the middleware applications and dispatching the action requests from the particular boxcar to the responsive application programs. In some implementations, the dispatcher postpones returning responses to action requests from a particular boxcar until it has received (or timed out on) responses to all of the action requests in the particular boxcar.

[0114] While the technology disclosed is disclosed by reference to the preferred embodiments and examples detailed above, it is to be understood that these examples are intended in an illustrative rather than in a limiting sense. It is contemplated that modifications and combinations will readily occur to those skilled in the art, which modifications and combinations will be within the spirit of the invention and the scope of the following claims.

What is claimed is:

1. A system including a user device and a server, coupled in communication, including:

a data consuming application, running on a user device, that makes action requests qualified by priority labels applied to regions and components in the regions and that consumes data received from at least one server responsive to the action requests;

- a middleware application, also running on the user device, that batches the action requests into batches based at least in part on the priority labels, and that dispatches boxcars of batched action requests;
- wherein the priority labels differentiate among display regions rendered by the data consuming application and at least some of the priority labels further differentiate among components within respective display regions;
- wherein the boxcars are segregated by priority label according to a predetermined segregation schedule; and
- wherein the action requests that are qualified by higher priority labels are dispatched ahead of the action requests with lower priority labels;
- a browser or app running on the user device, in communication with the middleware application, that supports the data consuming application; and
- the server, running a dispatcher and one or more application programs; and
- wherein the dispatcher receives the boxcar of action requests, forwards the action requests to the application programs, collects responses from the application programs, and returns a responsive boxcar of completed responses.
2. The system of claim 1, wherein the priority labels are set using an attribute of a DIV html tag.
3. The system of claim 1, further including:
- the middleware application, further implementing measurement of network communication latency, setting of inter-boxcar intervals used to dispatch the batches of action requests,
- wherein the measured network communication latency is calculated as dispatch-to-completed response time minus server processing time and the server processing time is received from the server for a boxcar of completed responses; and
- wherein the inter-boxcar intervals are adjusted taking into account at least (a) a number of connections supported between the user device and the server and (b) the measured network communication latency.
4. The system of claim 3, wherein the middleware application repeatedly measures network communication latency and readjusts the inter-boxcar intervals used to dispatch batches of action requests based on updated measurements.
5. The system of claim 3, wherein the middleware application adjusts the inter-boxcar intervals used to time when to dispatch batches of action requests to 250-350 ms when the measured network communication latency exceeds 300 ms.
6. The system of claim 3, wherein the middleware application identifies the browser or app through which it is coupled in communication with the server and looks up the number of connections supported by the identified browser or app.
7. The system of claim 3, wherein the middleware application receives the boxcar of completed responses grouped to respond to a respective boxcar of action requests.
8. A user device running a middleware application, including:
- a middleware application, running on the user device, coupled in communication with a data consuming application and a browser or app running on the user device, and also coupled in communication with at least one server through the browser or app;
- wherein the middleware application receives action requests from the data consuming application qualified by priority labels applied to regions and components in the regions and returns data received from the server responsive to the action requests;
- wherein the priority labels differentiate among display regions rendered by the data consuming application and at least some of the priority labels further differentiate among components within respective display regions;
- wherein the action requests that are qualified by higher priority labels are dispatched ahead of the action requests with lower priority labels; and
- wherein the middleware application batches the action requests into batches based at least in part on the priority labels, into boxcars segregated by priority label according to a predetermined segregation schedule, and dispatches the boxcars of batched action requests.
9. The user device of claim 8, wherein:
- the middleware application further measures network communication latency and adjusts inter-boxcar intervals used to dispatch the batches of action requests;
- wherein the middleware application calculates the measured network communication latency as dispatch-to-completed response time minus server processing time, using the server processing time received from the server for a boxcar of completed responses; and
- wherein the inter-boxcar intervals are adjusted, taking into account at least (a) a number of connections supported between the user device and the server and (b) the measured network communication latency.
10. The user device of claim 9, wherein the middleware application repeatedly measures network communication latency and readjusts the inter-boxcar intervals used to dispatch batches of action requests based on updated measurements.
11. The user device of claim 9, wherein the middleware application adjusts the inter-boxcar intervals used to time when to dispatch batches of action requests to 250-350 ms when the measured network communication latency exceeds 300 ms.
12. The user device of claim 9, wherein the middleware application identifies the browser or app through which it is coupled in communication with the server and looks up the number of connections supported by the identified browser or app.
13. The user device of claim 9, wherein the middleware application receives the boxcar of completed responses grouped to respond to a respective boxcar of action requests.
14. A server running a dispatching application, including:
- the dispatcher running on the server and coupled in communication with a multiplicity (100 or more) of middleware applications running on respective user devices and one or more application programs running on the server;
- wherein the dispatcher receives boxcarred batches of action requests from the middleware applications in priority-based batches, forwards the action requests to the application programs, and returns responsive boxcars of completed responses; and
- wherein the dispatcher further calculates a server processing time as a difference between when a particular boxcar of action requests is received and when a

respective boxcar of completed responses is ready to be returned and further reports the calculated server processing time back to the middleware application for the respective boxcar of completed responses.

15. The server of claim **14**, wherein the dispatcher further identifies respective application programs that will respond to the action requests received in a particular boxcar from one of the middleware applications and dispatches the action requests from the particular boxcar to the responsive application programs.

16. The server of claim **14**, wherein the dispatcher postpones returning responses to action requests from a particular boxcar until it has received (or timed out on) responses to all of the action requests in the particular boxcar.

17. A method that conserves connections between a browser or app running on a user device and at least one server, including:

running a middleware application on the user device, coupled in communication with a data consuming application and a browser or app running on the user device, and also coupled in communication with at least one server through the browser or app;

the middleware application receiving action requests from the data consuming application qualified by priority labels applied to regions and components in the regions and returns data received from the server responsive to the action requests;

wherein the priority labels differentiate among display regions rendered by the data consuming application and at least some of the priority labels further differentiate among components within respective display regions; and

the middleware application batching the action requests into batches based at least in part on the priority labels, into boxcars segregated by priority label according to a predetermined segregation schedule, and dispatching the boxcars of batched action requests to the server.

18. The method of claim **17**, further including:

the middleware application measuring network communication latency and adjusting inter-boxcar intervals used to dispatch batches of action requests;

the middleware application calculating the measured network communication latency as dispatch-to-completed response time minus server processing time, using the server processing time received from the server for a boxcar of completed responses; and

the middleware application adjusting the inter-boxcar intervals taking into account at least (a) a number of connections supported between the user device and the server and (b) the measured network communication latency.

19. The method of claim **18**, further including repeatedly measuring network communication latency and readjusting the inter-boxcar intervals used to dispatch batches of action requests based on updated measurements.

20. The method of claim **18**, further including adjusting the inter-boxcar intervals used to time when to dispatch batches of action requests to 250-350 ms when the measured network communication latency exceeds 300 ms.

21. The method of claim **18**, further including the middleware application identifying the browser or app through which it is coupled in communication with the server and looking up the number of connections supported by the identified browser or app.

22. The method of claim **18**, further including the middleware application receiving the boxcar of completed responses grouped to respond to a respective boxcar of action requests.

23. A method that conserves connections between a browser or app running on a user device and at least one server, including:

a dispatcher running on the server and coupled in communication with a multiplicity of middleware applications running on respective user devices and one or more application programs running on the server;

the dispatcher receiving boxcarred batches of action requests in priority-based batches from the middleware applications, forwarding the action requests to the application programs, and returning responsive boxcars of completed responses; and

the dispatcher further calculating a server processing time as a difference between when a particular boxcar of action requests is received and when a respective boxcar of completed responses is ready to be returned and further reporting the calculated server processing time back to the middleware application for the respective boxcar of completed responses.

24. The method of claim **23**, further including the dispatcher identifying respective application programs that will respond to the action requests received in a particular boxcar from one of the middleware applications and dispatching the action requests from the particular boxcar to the responsive application programs.

25. The method of claim **23**, further including the dispatcher postponing returning responses to action requests from a particular boxcar until it has received (or timed out on) responses to all of the action requests in the particular boxcar.

* * * * *