

Copyright
by
Marcos Suguru Kajita
2009

The Report Committee for Marcos Suguru Kajita Certifies that this is the approved version of the following report:

**Google App Engine Case Study:
A Micro Blogging Site**

Committee:

Adnan Aziz, Supervisor

Sarfraz Khurshid

**Google App Engine Case Study:
A Micro Blogging Site**

by

**Marcos Suguru Kajita
BSCS**

Report

Presented to the Faculty of the Graduate School of
The University of Texas at Austin
in Partial Fulfillment
of the Requirements
for the Degree of

Master of Science in Engineering

**The University of Texas at Austin
December, 2009**

Google App Engine Case Study: Micro Blogging Site

Marcos Suguru Kajita, MSSE
The University of Texas at Austin, 2009

Supervisor: Adnan Aziz

Cloud computing refers to the combination of large scale hardware resources at datacenters integrated by system software that provides services, commonly known as Software as a Service (SaaS), over the Internet. As a result of more affordable datacenters, cloud computing has slowly made into the mainstream business arena and has the potential to revolutionize the IT industry. As more cloud computing solutions become available, it is expected that there will be a shift to what is called the Web Operating System. The Web Operating System along with the sense of infinite computing resources on the “cloud” has the potential to bring new challenges in software engineering. These challenges are the main motivation of this paper, which is divided into two parts. The first part gives a brief introduction and analysis of cloud computing. The second part focuses on Google’s cloud computing platform and evaluates the implementation of a micro blogging site using Google’s App Engine.

Table of Contents

List of Tables	vii
List of Figures	viii
Introduction.....	1
Cloud Computing.....	4
The Web Operating System	4
Impact on Business	5
Cloud Computing Challenges	7
Choosing the Right Cloud Computing Platform	8
Google App Engine.....	10
Java and Sandbox	10
GAE Plug-in for Eclipse	11
Google Web Toolkit	14
Datstore	15
Micro Blogging Site on Google App Engine.....	16
Motivations	16
Basic Approach	17
MyMicroBlog Architecture	17
MicroBlog	19
IBlogView	22
BlogEntry and ImageEntry Classes	24
BlogItem and ImageItem Classes	25
IBlogEntryService and IBlogServiceAsynch Interfaces	25
BlogEntryServiceImpl Class	26
Results	27
Development Experience	27
Performance	27
Lessons Learned	28

Future Work	28
Conclusion	30
Bibliography	31
Vita	33

List of Tables

Table 1:	List of Views in MyMicroBlog	22
----------	------------------------------------	----

List of Figures

Figure 1:	The Computing Platform Lifecycle	3
Figure 2:	Operating system market share	5
Figure 3:	The Cloud Computing Arena.....	6
Figure 4:	GAE menu buttons on Eclipse.....	12
Figure 5:	Eclipse Package Explorer	12
Figure 6:	GAE Web Application Wizard	13
Figure 7:	GAE Deployment Tool on Eclipse	13
Figure 8:	MVC Diagram	18
Figure 9:	MyMicroBlog Diagram	19
Figure 10:	Google Sign-in Screen	20
Figure 11:	MicroBlog Class Diagram	21
Figure 12:	Partial Class Diagram of View Classes in MyMicroBlog	23
Figure 13:	Layout of CreateBlogView using GWT Panels.....	24
Figure 14:	BlogItem Entity on Google App Engine Datastore	25
Figure 15:	BlogEntry RPC Implementation	26
Figure 16:	BlogEntry Servlet Configuration	27

Introduction

Cloud computing received its name from the representation of the Internet as a cloud on blackboard diagrams. Formally, cloud computing refers to the hardware, systems software, and datacenters that provide services over the Internet [1]. Cloud computing is not a new concept. In fact, it is in use in many places, but on a smaller scale. For example, researchers at university campuses perform complex computations on bio-chemical models and weather pattern predictions much faster on grid computers [2]. Users play the same online game from different geographical locations. Documents are created and shared by many users through Google Docs. These are examples of cloud computing concepts in use in smaller scale. Recently, cloud computing has made into the main stream business arena due to the commoditization of computers. Many companies have started to offer their own cloud computing solutions. Companies like Amazon, Google, Microsoft, IBM, and Sun have released innovative cloud computing solutions. These solutions vary from virtualized hardware to well designed API's. In a few months, users will have an increased number of options when more companies release their own cloud computing solutions. The business of providing cloud computing solutions is growing very fast because businesses taking advantage of such offerings are growing as well. As users move to the *cloud*, it is expected that there will a migration to the Web Operating System (Web OS). Cloud computing is setting the stage for a revolution in the IT industry. There is a possibility that desktop applications will be replaced by web applications and companies will use cloud computing solutions for their IT needs. At the moment, cloud computing is at its infancy. As any other platform development cycle (figure 1), there will be an adjustment phase followed by the standardization of cloud computing technologies. During the initial phase, it is important that users evaluate and

gain expertise on the *cloud* technologies because as in any other development cycle there will be a few technologies more successful than others. It is important that users become familiar with the different technologies in order to choose the prevailing technology when the industry stabilizes. Until the cloud computing industry stabilizes, the user needs to evaluate and choose the *cloud* technology that best fits his needs for the short term. For the long term, the user needs to compare and contrast different technologies. These experiments should not be a problem since the cloud technologies are cheap. The affordability of *cloud* solutions makes the technology even more attractive to users. For example, the Google App Engine (GAE) platform, which is evaluated in the second part of the paper, is free up to a point. With affordable solutions like GAE, there is no reason for users to stay away from the *cloud* technologies. In fact, there are plenty of reasons to stay up to date with the technology because the *cloud* enables creativity to drive the development of products and businesses. The sudden competition from fast-followers can cause a major impact on businesses in general. Therefore, careful evaluation of technologies like GAE helps keep one eye on the technology and the other on the competition.



Figure 1: The Computing Platform Lifecycle

Cloud Computing

THE WEB OPERATING SYSTEM

Hinchcliffe [4] describes the Web Operating System (Web OS) as the “overarching and emergent Internet Operating System.” The Web OS has become a reality with cloud computing. The Web OS is the environment where users run applications as if running them on the desktop. As the cloud computing facilities and services grow, there will be a migration from the desktop-centric to web applications. Most applications developed today target the desktop environment on specific Operating System (e.g., Windows, Mac, or Linux). On the Web OS, applications run from anywhere with a simple web browser. The Web OS takes today’s web services to a new level by providing data storage, services, API’s, and toolkits where everything is programmable and new products and businesses are created at will. The Web OS has opened the doors for new business opportunities. At the moment, only few companies like Amazon, Google, Microsoft, IBM, and Sun have released products with *cloud* capabilities, but there are many other companies lined up with similar solutions. The solutions from these companies range from hardware virtualization to flexible API’s. Each company is trying to differentiate from each other. As a result, there are plenty of options for users. The race for Web OS dominance has started and at the moment there is no clear winner [9]. Microsoft currently holds over 90% of the desktop operating system market. This is not true for the Web OS and there is plenty of the “pie” for any new player that is willing to capitalize on this new environment.




Figure 2: Desktop Operating System Market Share

IMPACT ON BUSINESS

Cloud computing and the promise of the Web OS is stirring the IT Industry. Businesses that traditionally had very little presence in IT have become major players on this new business endeavor. Amazon is a good example of such company. From its humble beginnings as an online store, Amazon has become one of the leading cloud computing providers. Many other companies are using their strengths to create successful cloud computing offerings. Some companies seem to have an advantage over the competitions since they built their company's DNA around the web. One example of such company is Google. Google has harnessed its experience and presence on the web to provide enterprise solutions on cloud computing. Traditionally, Google has not been involved in enterprise solutions but with its release of GAE it has plans to expand its

businesses beyond online advertisements [4]. Many cloud providers are offering new cloud solutions and the market does not seem to be limited to large companies only. The same commoditization of computers that has initially boosted the *cloud* business has also opened the doors for smaller companies who want to provide their unique solutions. The cloud has become a large centrifugal force that keeps attracting companies and businesses.

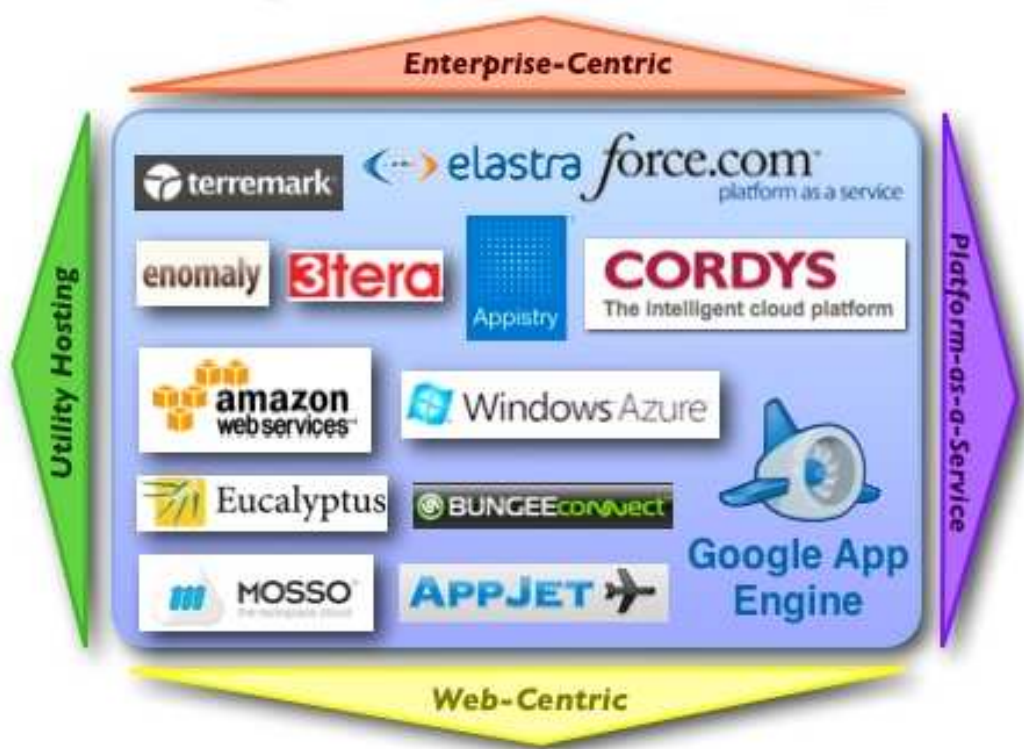


Figure 3: The cloud computing arena [10]

For users, the bottom line is that cloud computing is affordable. Amazon sells its service for only 10 cents per hour and its Scalable Storage Service (S3) for \$0.12 per gigabyte-month. The majority of today’s cloud computing clientele are medium and small businesses (i.e., startups). For these companies, affordable services like S3 provide

a golden opportunity to launch new products with significantly less capital upfront. The economics around cloud computing allow innovative companies to create products that were economically challenging or are significantly cheaper than the competition. Traditionally, many business ideas took time to realize because of the prohibitive cost of computing power, scalability issues, and primarily the high risk due to large capital. This has all changed and innovation is now the key.

CLOUD COMPUTING CHALLENGES

Most of the challenges will be on the cloud computing security. The idea of moving the entire company's private data to some remote data center controlled by a third party company may seem like exposing business secrets. It will take time to adjust to this new model and security will be one of the biggest challenges to cloud providers [10]. Imagine that a company has put all its data on Amazon's S3. Despite the fact that Amazon has built its reputation on reliability, there is still the risk of the service going offline for a few hours. If that happens, the user's data is no longer available and he is at the mercy of Amazon's ability to recover [10]. However, with heavier investments on systems software and hardware, Amazon's services should still be more resilient to system failures and secure than any in-house IT solution.

There is also the concern that "once you are in, you are really in" [5]. For many startups, the decision whether to adopt a cloud computing platform is simpler. However, for more established companies, the worry about lock-in is a real concern. Some companies are concerned that if they opt for a platform, they may lose their investment if the cloud provider goes under. There is always the risk of lock-ins if a technology is used before fully understanding it. For example, naïve implementation on Google's App Engine may result in a suite of applications that may be challenging to port to another

environment. If for some reason Google decides to increase their fees and the cost no longer makes economical sense for the client to stay with Google, the migration back to the client company's environment or another provider may turn out to be difficult and costly. The challenge is to understand the strengths and limitations of each cloud computing solution to make intelligent long term decisions.

CHOOSING THE RIGHT CLOUD COMPUTING PLATFORM

As mentioned above, there are several cloud computing solutions in the market today. The choices range from virtualized hardware to all-in-one API packages. The key to model computation, storage, communication, and the illusion of infinite resources is the level of abstraction these different solutions provide. Each solution seems to focus on different aspects of abstraction [1]. For example, Amazon's Elastic Compute Cloud (EC2) provides very low level abstraction to the hardware. EC2 virtualizes the hardware and offers a set of libraries that allow developers full control over the software stack. In contrast, Google's App Engine (GAE) provides a framework based on well defined interfaces and API's that allow quick development with quality. The development on GAE is done on either Java or Python. There is support for a few other languages but they are not as significant as for these two. For example, Google's data storage service supports Java Data Object (JDO), which completely abstracts the implementation of the database. At the moment, it is difficult to determine the right cloud computing platform. It is up to the user to determine whether he needs the level of hardware control provided by Amazon's EC2 or the easy migration of Java code to the App Engine environment. Hopefully, as the platform development cycle comes to a full circle, there will be enough standards that allow users to develop on one platform and store data on another. At that

point, the right choice will be the platforms that support the standards and allow users to migrate seamlessly from one provider to another.

Google App Engine

Google App Engine (GAE) is the platform that allows software developers to leverage Google's cloud computing infrastructure for web application development [3]. Similar computing resources (or services) used by Google Docs are available under the GAE. Google offers the same reliability and openness as its other flagship products like Google Maps and Google Earth. GAE is free up to a point, but there is plenty of room for the experiments that are needed to evaluate the technology. The following are some of GAE's key features.

- Dynamic web serving, with full support for common web technologies
- Persistent storage with queries, sorting, and transactions
- Automatic scaling and load balancing
- API's for authenticating users and sending email using Google Accounts
- A fully featured local development environment that simulates Google App Engine on your computer
- Task queues for performing work outside of the scope of a web request
- Scheduled tasks for triggering events at specified times and regular intervals [3]

JAVA AND SANDBOX

GAE released the support for Java only recently. Initially, it only supported Python, but since most of web enterprises have standardized to either .NET or Java, it was a matter of time to have Java support in GAE. The Java run time environment in GAE is feature rich and offers comprehensive support of Java's standard libraries and API's. All the development can be done in native Java and the application interacts with GAE via Java Servlet Interfaces. The Java environment provides Java 6 JVM, JavaMail,

JCache, and standard interfaces for GAE's datastore with Java Data Object (JDO) and Java Persistence API (JPA). Java standards make the development on GAE easy and familiar. There are some limitations due to sandboxing, so migrating existing code to GAE requires some minor changes. The JVM runs on a sandbox to provide security to applications. The sandbox assures that each application on GAE do not interfere with each other. It also guarantees performance and scalability.

GAE PLUG-IN FOR ECLIPSE

Eclipse is an open source community with the objective of building a standard development platform. The Eclipse Integrated Development Environment (IDE) provides Java as its primary programming language, but it supports plug-ins for many other languages such as C, C++, etc. Development on GAE could not be simpler with the GAE plug-in for Eclipse IDE. Eclipse manages the workspace and makes the development on GAE as simple as developing any other servlet-based web application. The Eclipse plug-in allows developers to create new GAE application from within Eclipse IDE. Provided that the user has a GAE account, the user can create, deploy, and run web applications directly from the IDE. The App Engine SDK provides a web server for testing applications locally in a simulated environment. The plug-in adds extra buttons to the Eclipse menu bar (figure 4) that allows the applications to run on the local web server in hosted mode.




Figure 4: GAE Menu Buttons on Eclipse

The management of client and server source files is also done from within the IDE (figure 5). Eclipse's development tools and debugging facilities are available for developing web applications. GAE application development on Eclipse is no different than writing a servlet or desktop application. Once the application is ready, the deployment is as easy as a button click.

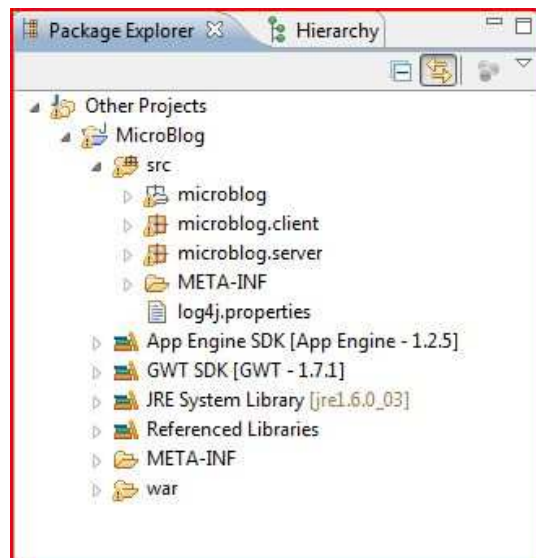


Figure 5: Eclipse Package Explorer

When creating a new web application, the wizard (figure 6) takes care of the creation of the necessary configuration files needed to run applications on GAE. The web.xml and the main HTML files are generated automatically by the GAE plug-in.

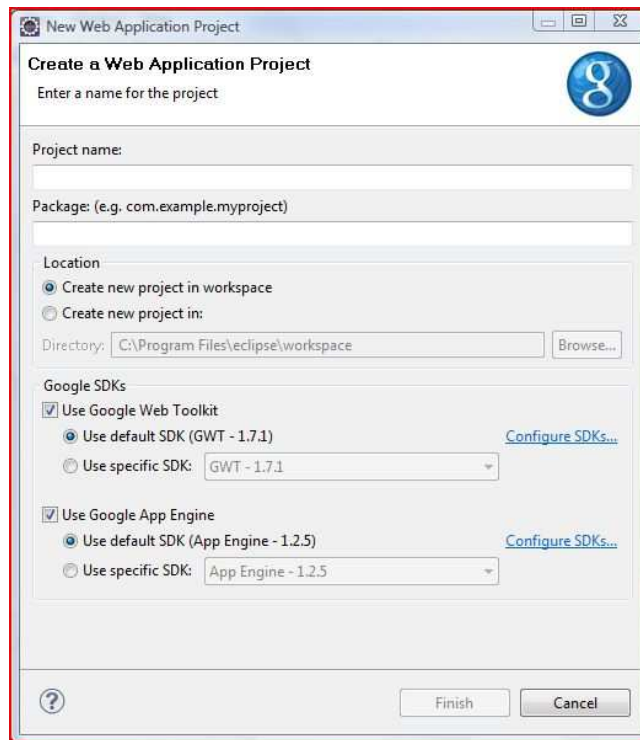


Figure 6: GAE Web Application Wizard

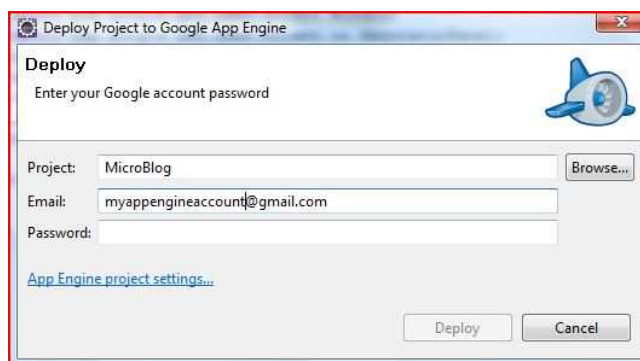


Figure 7: GAE Deployment Tool on Eclipse

GOOGLE WEB TOOLKIT

Developing web applications can be tedious and time consuming. It is fair to say that software developers spend most of their time trying to figure out little quirks in JavaScript or compatibility issues on certain browsers. Asynchronous JavaScript plus XML (AJAX) is usually a popular option for writing interactive and attractive web applications. AJAX was made popular in 2005 by Google and makes use of web standard such as JavaScript, XML, HTML, and CSS. However, writing scalable, reusable, and maintainable code in AJAX can be challenging. Google Web Toolkit (GWT) abstracts the complexity of JavaScript by providing a rich set of Java tools to develop attractive front end applications. Developers use GWT to design the layout of web applications and the display of Graphical User Interface (GUI) elements in pure Java. GWT's interfaces and libraries mimic those of existing User Interface (UI) frameworks such as Swing and Standard Widget Tool (SWT). The main difference is that GWT dynamically generates HTML instead of pixel based graphical elements. For example, GWT's Button results in HTML's '<button>' instead of an image or pixel based element. As a result, there are virtually no cross-browser incompatibility issues. Behind the scenes, GWT compiles the Java code into optimized JavaScript that runs on most of commercial browsers. Events, such as button clicks, are coded directly to the GUI element in Java so it follows object oriented programming. GWT code running on Eclipse in hosted mode enables programmers to use debugging tools allowing them to step into the code line by line. Writing interactive web applications with attractive GUI elements in GWT is very similar to writing desktop applications.

DATASTORE

GAE makes data persistency on Google's distributed data storage facilities very simple. GAE abstracts the administrative work required to maintain a distributed system and database such as Bigtable [6]. Bigtable is Google's solution for high availability distributed storage system. The system manages petabytes of structured data across multiple scalable commodity servers. As resources are needed more machines can be added to the clusters without interrupting the system. Depending on the type of web application, the user could be interacting with several servers at any given time. With GAE, the programmer does not worry about any of the intricacies needed to maintain and synchronize data in the datastore. The engine takes care of all the transactions, load distributions, and even backups. The access to GAE datastore is provided via standard API's, which result in a portable system. In other words, web applications built on these interfaces can be ported to and from Google environment effortlessly. This is an important characteristic of Google's datastore because the resulting system has no risk of lock-in with Google's BigTable. The engine supports the Java Data Object (JDO) and Java Persistence API (JPA). Both standards are from the DataNucleous Access Platform, which is an open source platform of many Java standards [7]. GAE also supports a SQL like query language called GQL. Because the datastore is not a traditional relational database, GQL provides a layer of abstraction with query notations similar to SQL. GQL simply provides familiarity to programmers accustomed with SQL statements on relational databases. Since GQL is a proprietary abstraction of Google's Bigtable, it is not recommended to design the system on it because of the risk of coupling and lock-in. The resulting system would not allow the user to port to a non-Google environment easily.

Micro Blogging Site on Google App Engine

MOTIVATIONS

The main motivation for implementing the micro blogging site is to evaluate GAE's services and technologies. A standard micro blogging site has the following features:

- User registration and secure login
- Client-server model
- Dynamic updates and data persistence
- Events and notifications
- Email communication

The goal is to implement the features listed above using the tools and API's provided by GAE development platform. This effort will provide an insight on the level of usability, scalability, portability, and flexibility in developing web applications on the GAE's platform.

The secondary motivation for developing the application is to design components for a micro blogging framework. The goal is to develop a set of modular components and building blocks for developing micro blogging sites. The use of micro blogs has made into corporate life where the tool helps create communities within the company environment [8]. The messaging and notification functionality of blogging tools help the exchange of information among colleagues at work. The exchange of messages does not need to be between people. Systems can be integrated to the blogging system so that it can notify users of events. For example, a blogging system could be used by a software engineering team to monitor the development efforts. When new code is checked into the Version Control System (VCS), a text message is sent to the software architect or

developer about certain events (e.g., build or compilation errors). Depending on the type of business the requirements for a micro blogging application can vary and commercial products may not support them out of the box. Hence, there is a need for a platform for quick development of customized features [8].

BASIC APPROACH

The MyMicroBlog micro blogging site is implemented entirely in Java on GAE using the Eclipse IDE. The Web Application Wizard, which is a tool provided by GAE plug-in for Eclipse, helped create all the preliminary client and server code, configuration files, and application settings (figure 6). Once the application was created, the deployment to Google's cloud environment was a two button click process using the Web Application Deployment tool (figure 7). Once the deployment completed, the site was running on <http://mymicroblog.appspot.com/>. The entire process of creating and deploying the web application did not take more than a few seconds. The client and server folder structure is created by the wizard and the environment is ready for development of the remaining pieces that makes the micro blogging site (figure 5).

MICROBLOG ARCHITECTURE

The basic architecture of MyMicroBlog is based on the Model View Controller (MVC) architectural pattern (figure 8). The pattern has the benefit of isolating the business logic from the presentation and data layers. The *model* encompasses the object or data structure used by the application. Any manipulation to the domain's data is dealt within the *model* and notification is sent to the views. The *view* presents the *model* (i.e., data) to the user. Typically, the *view* has only the implementation of the GUI elements

that visualizes the state of the *model*. The *controller* is the object that receives inputs (usually user inputs) and notifies the *model* or the *view*. The resulting application using the MVC pattern is modular and easy to test. Modular applications have the benefit of reusability, extensibility, and rapid development.




Figure 8: Model-View-Controller Diagram

The implementation of applications using MVC pattern usually uses the Observer design pattern. In the observer pattern, the object, called the subject, manages a list of “observers” and notifies them when an event occurs. In MVC, the controller is usually implemented using the observer pattern.

The MyMicroBlog does not use the observer pattern, but implements a controller class in the MVC pattern. The following diagram depicts the high level architecture of MyMicroBlog.

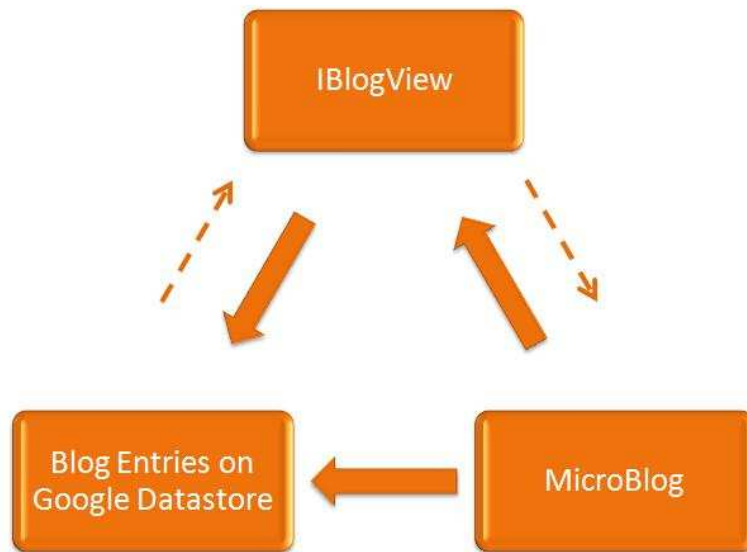


Figure 9: MyMicroBlog Diagram

MICROBLOG CLASS

The MicroBlog class is the main entry point for the application. It implements the EntryPoint interface from com.google.gwt.core.client package. The 'onModuleLoad' method from the EntryPoint interface is called when the application first starts. The user validation is performed in the onModuleLoad method. The user authentication is done via GAE's UserServieFactory, which is a service that authenticates users with Google accounts. If the user is validated successfully, the main MyMicroBlog page is loaded. If the user is not validated, the user is directed to the Google account sign in page (figure 10). The class is also the controller on the MVC pattern. The class handles the messages from the views and generates notifications. Figure 11 shows the MicroBlog class diagram.



Figure 10: Google Sign-in Screen

The main navigation bar for the application is also in the MicroBlog class. The navigation bar receives the user input and the MicroBlog class loads the appropriate views or re-routes the notifications to the model.

The current version of the MicroBlog class has unnecessary complexity. In order to improve the object, the controller logic can be replaced by an observer pattern. The navigation bar also adds complexity and makes the MicroBlog class less flexible. If the navigation bar logic is moved into its own class, it will allow for better customization and reuse. The result would be a lightweight MicroBlog class that serves as the entry point for the application and delegates the core functionality to the MicroBlog MVC architecture.

One important point to make about the class is the fact that it uses the EntryPoint interface. The use of EntryPoint can impose a problem for portability since the interface is part of the GWT package. However, after the two improvements mentioned above, the class should be very light weight and porting the application to a non-Google environment would require very little refactoring. In addition to that, GWT is open source and is available under the Apache 2.0 license.



Figure 11: MicroBlog Class Diagram

IBLOGVIEW

IBlogView implements the interface that allows the different views in the MyMicroBlog application to be displayed. The interface implements the following methods.

- void Initialize(MicroBlog)
- Widget GetWidget()
- void Reset()

The view that implements the IBlogView interface can be plugged into the MicroBlog main page seamlessly. Each of the view classes renders the GUI interface and handles the user inputs.

The following table shows the different views implemented in MyMicroBlog.

View Name	Description
AboutUsView	Displays the AboutUs page
CreateBlogView	Displays the create blog page with controls to allow new blog entries
DisplayBlogListView	Displays a page with a list of blogs
DisplayBlogView	Displays the blog entry
DisplayImageView	Displays the page for image loading and display
EditBlogView	Displays the page to edit blog entries
HelpView	Displays the Help
ReportAbuseView	Displays the report abuse page
ReportProblemsView	Displays the report problems page
SearchBloggerView	Displays the search page for bloggers
SearchByCategoryView	Displays the search page for categories
SendEmailView	Displays the email communication page
UserProfileView	Displays the user profile page

Table 1: List of Views in MyMicroBlog

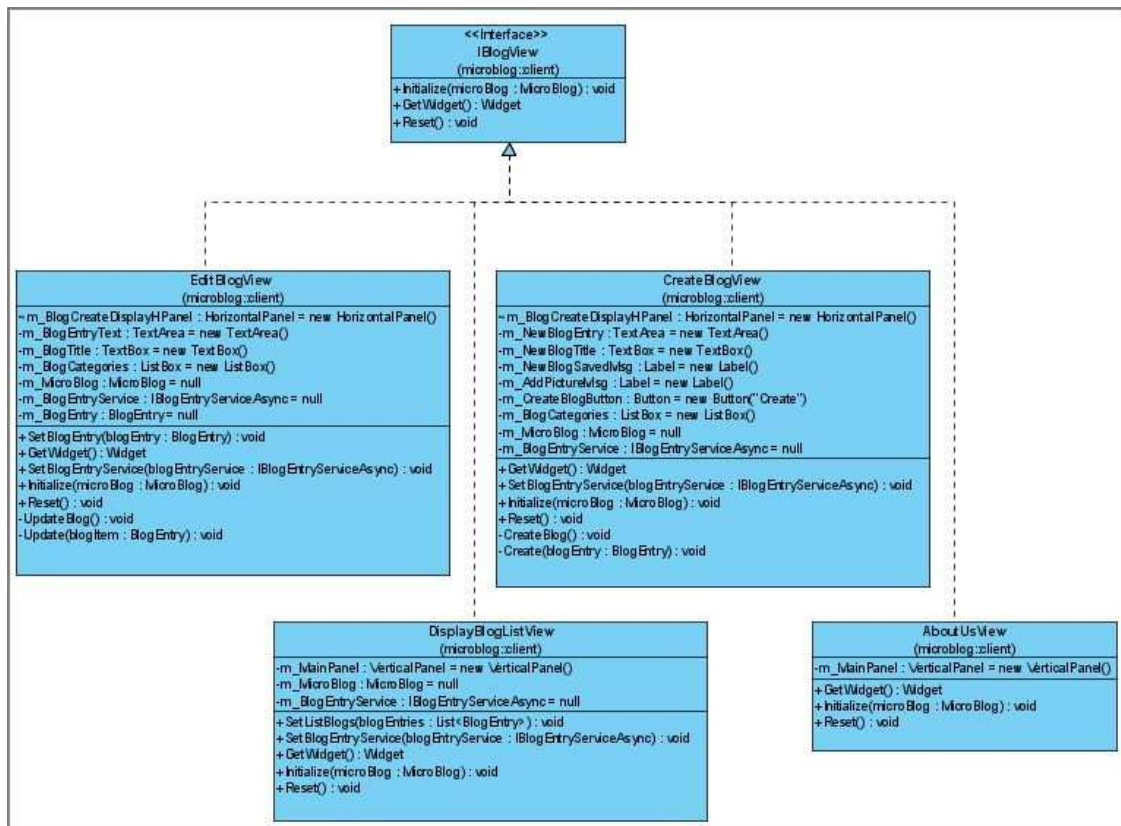


Figure 12: Partial Class Diagram of View Classes in MyMicroBlog

Each view class renders the GUI interface for a particular function in the system. For example, the `EditBlogView` class creates the controls needed to edit the blog entry and handles the user commands (i.e., Update or Cancel). The commands (events) are sent to the `MicroBlog` class, which in turn re-routes it to the `BlogEntry` object, which is an object in the *model* component on MVC.

The use of GWT panels and controls makes the development of attractive and interactive web pages extremely simple. All the micro blog view classes use GWT to create the GUI elements. The layout of the controls is done via the GWT panels. The concept of paneling allows developers to take a page layout and translate it into code. The

alignment of controls within a page is all done by specialized GWT panels. Having an upfront idea of the layout helps the selection of the panels. Figure 13 shows the different panels used in the layout of the CreateBlogView.

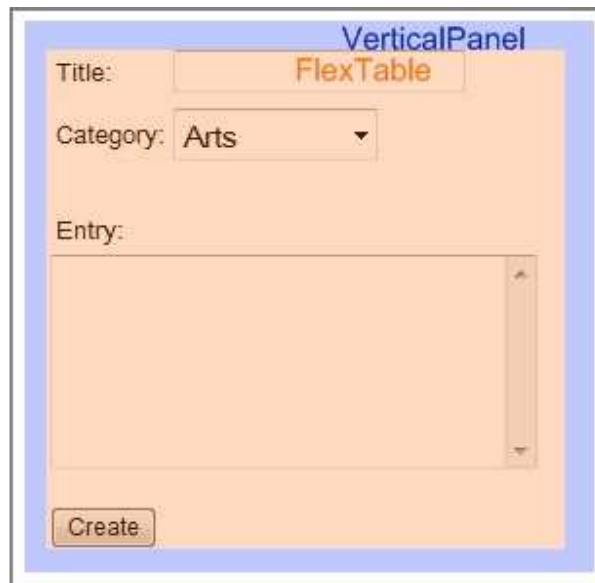


Figure 13: Layout of CreateBlogView using GWT Panels

Since the view classes make heavy use of GWT, the same observation regarding lock-in due to GWT applies. However, as long as the application conforms to the Apache 2.0 license, the code should be fairly portable.

BLOGENTRY AND IMAGEENTRY CLASSES

The BlogEntry and ImageEntry classes are the data structure that handles the blog and image entries. They both make the *model* component in the architecture. The BlogEntry class has the interfaces to create and edit a new blog entry. The ImageEntry class has the interfaces that allow the loading and storage of images in the application.

They both implement the `ISerializable` interface allowing them to be passed to the server by the Servlets.

BLOGITEM AND IMAGEITEM CLASSES

The `BlogItem` and `ImageItem` classes are the counterparts for the `BlogEntry` and `ImageEntry` on the server side. The `BlogEntry` and `ImageEntry` are stored in the datastore as instances of `BlogItem` and `ImageItem` respectively (figure 14). The main difference is that the item classes are simple data structures that allow the storage of blog and image information. Both `BlogItem` and `ImageItem` classes have only getter and setter functions. On the other hand, `BlogEntry` and `ImageEntry` can perform specialized functionality to handle events in the application.

BlogItem Entities							
ID/Name	m_CreateDate	m_LastUpdated	m_NickName	m_strCategory	m_strEntry	m_strImages	m_strTitle
id=7021	2009-10-27 03:53:07.696000	2009-10-29 02:28:50.588000	kajitams1	Auto	This is a test entry on mymicroblog.	7026	Test2
id=7025	2009-10-27 04:17:20.128000	2009-10-29 00:17:03.336000	kajitams2	Computers	For best visual results use IE 7 or 8. There are a few items on this page that do not render well on Chrome.	9005	MicroBlog1
id=9002	2009-10-27 03:24:23.559000	2009-10-29 02:29:09.923000	kajitams1	Business	Test1	9003,11002,12001	Test1

Figure 14: BlogItem Entity on Google App Engine Datastore

IBLOGENTRYSERVICE AND IBLOGSERVICEASYNCH INTERFACES

The `IBlogEntryService` and `IBlogServiceAsynch` interfaces makes up the code needed for the Remote Procedure Call (RPC). The interfaces are used to invoke the server and pass the `BlogEntry` object between the client and server. Although both interfaces are written in Java, GWT compiles then into JavaScript. Both the client and

server serialize and deserialize the data object to and from simple text form during the data exchange.

BLOGENTRYSERVICEIMPL CLASS

The `BlogEntryServiceImpl` class inherits from the `RemoteServiceServlet` and implements the `IBlogEntryService` (figure 15). `RemoteServiceServlet` is the servlet base class that manages the RPC calls and automatically performs deserialization on incoming data from the client and serialization on the outgoing data to the client. Ultimately, the `BlogEntryServiceImpl` is a Servlet that handles the requests from clients through the methods implemented by the `IBlogEntryService` interface.

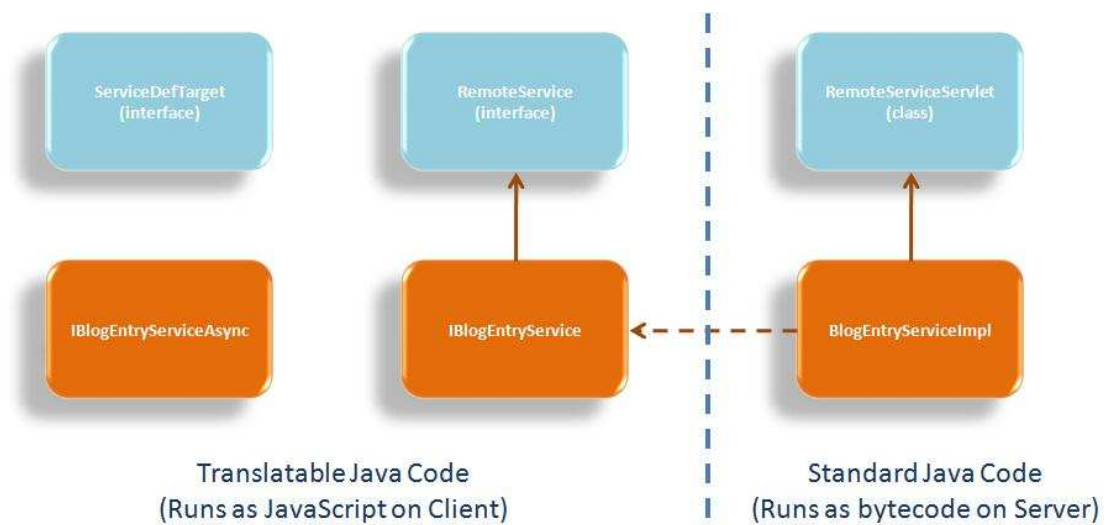


Figure 15: BlogEntry RPC Implementation

The configuration of the new Servlet is done on the `war.WEB-INF.web.xml` file (figure 16). After adding new Servlets the `web.xml` file needs to be updated and re-deployed.

```
<!-- Servlets -->
<servlet>
<servlet-name>MicroBlogServlet</servlet-name>
<servlet-class>microblog.server.BlogEntryServiceImpl</servlet-class>
</servlet>

<servlet-mapping>
<servlet-name>MicroBlogServlet</servlet-name>
<url-pattern>/microblog/microblogs</url-pattern>
</servlet-mapping>
```

Figure 16: BlogEntry Servlet Configuration

RESULTS

Development Experience

The overall development experience in GAE is extremely rewarding. The platform and GWT allow for quick development of interactive web applications. The GWT plug-in for Eclipse is a must have. There is very little to do as far as configuration and setup once you use the Create and Deploy Application Wizards. The developer can truly focus on the architecture and design of the web application.

Performance

The overall performance of GWT does not seem to affect the usability of the application. GWT compiles the GWT components in optimized JavaScript, so at run time there is no performance difference between JavaScript coded manually and auto generated by GWT. In fact, Google states that the JavaScript code generated by GWT performs better because of optimizations.

The only noticeable performance problem is during the initial creation of the PersistenceManagerFactory. The PersistenceManagerFactory used for JDO data transactions is very costly at first. In order to cope with the initial hit, a singleton class wraps the PersistenceManagerFactory instance. After the singleton class is created, the

PersistenceManagerFactory instance is kept in memory, and there is no significant performance hit thereafter. The performance hit is noticeable only if the singleton is unloaded due to low traffic. In that case, the user who visits the site for the first time will notice a 5~6 seconds of delay.

Lessons Learned

The Create Application and Deploy Application wizards are very useful. They automatically generate configuration files, folder structure, and source code. The configuration of the application is done once, so the auto configuration feature is very helpful. The auto generation of code is useful as well because in order to generate the generic greeting pages, the wizard also builds the package structure for the client and server. The only drawback of auto generation of code and configuration files is that the developer may end up with unnecessary code. Debugging through auto-generated code can be time consuming due to unfamiliarity and amount of unnecessary code. Errors in the configuration file introduced by new features added manually can be hard to detect, especially when programmer is a novice on Servlet code.

FUTURE WORK

There is room for improvement on the MyMicroBlog platform. The current implementation requires refactoring on the MicroBlog class in order to make the platform truly modular. For example, the complexity of the MicroBlog controller class does not allow for easy plug-in of new functionality in the navigation bar. The two refactoring points made in the MicroBlog Class section will result in a better design.

Security is also another aspect that needs attention in the platform. The calls to blog and image Servlets have no authentication at the moment. Any malicious web crawler could invoke the Servlets causing unnecessary loads on the server.

The user interface uses the default GWT styles. The default styles are regular Google styles and the page has the look and feel of any other Google application. Adding a mechanism to load a custom style sheet will allow the user to add personality to the web application. GWT controls have their own properties and need to be set on the main style sheet.

Conclusion

There is no doubt that cloud computing has opened new business opportunities and has started a new Internet revolution. For end users, the *cloud* promises unlimited computing power at very affordable prices. Creative companies can take advantage of the *cloud* services to deploy applications that before were too expensive to develop or maintain. There is the opportunity of getting rid of the IT department in the company all together resulting in savings and less overheads. For *cloud* providers, there is a new business opportunity where commoditized computer resources can be sold as pay-as-you-go model. The race is on for cloud providers and the market does not have a clear winner yet. All the cloud solutions available in the market today make the early stage of the platform development cycle. Until the cycle comes to a full circle, users should gain experience on the cloud technologies available. This is a very important step to create the sense of direction in the industry and the creation of standards. From the start of the cycle to the end when the cloud computing platform stabilizes, users will most likely be winners since the capital to experiment on technologies such as GAE is low. GAE has proven to be a well thought-out and designed platform. The use of standard API's in GAE allows the development of well designed and attractive web applications. The experience on GAE shows that users can focus on the business logic while developing flexible and portable web applications.

Bibliography

- [1] Armbrust, M., Fox, A., Griffith, R., Joseph, A. D., Katz, R., Konwinski, A., Lee, G., Patterson, D., Rabkin, A., Stoica, I., Zaharia, M., Above the Clouds: A Berkeley View of Cloud Computing, February 10th, 2009.
- [2] Delic, K. A., and Walker, M. A., Emergence of the Academic Computing Clouds, Hewlett-Packard Co., ACM Ubiquity, Volume 9, Issue 31, August 5th, 2008.
- [3] Google App Engine Main Site, <http://code.google.com/appengine/>
- [4] Hinchcliffe, D., How the Web OS Has Begun to Reshape IT and Business, September 6th, 2009, <http://blogs.zdnet.com/Hinchcliffe/?cat=78>.
- [5] Rogers, G., The Problem with Google App Engine, April 11th, 2008, <http://blogs.zdnet.com/Google/?p=1002>
- [6] Chang, F., Dean, J., Ghemawat, S., Hsieh, W. C., Wallach, D. A., Burrows, M., Chandra, T., Fikes, A., Gruber, R. E., Bigtable: A Distributed Storage System for Structured Data, OSDI'06: Seventh Symposium on Operating System Design and Implementation, Seattle, WA, November, 2006, <http://labs.google.com/papers/bigtable-osdi06.pdf>
- [7] DataNucleous Access Platform, <http://www.datanucleus.org/>
- [8] Hinchcliffe, D., Twitter in your Intranet: 17 Micro Blogging Tools for Business, June 1st, 2009, ZDNet, <http://blogs.zdnet.com/Hinchcliffe/?p=414>.
- [9] Hinchcliffe, D., Cloud Computing and the Return of the Platform Wars, March 26th, 2009, ZDNet, <http://blogs.zdnet.com/Hinchcliffe/?p=303#more-303>.

[10] Hinchcliffe, D., Cloud Computing: A New Era of IT Opportunity and Challenges, March 3rd, 2009, ZDNet, <http://blogs.zdnet.com/Hinchcliffe/?p=261#more-261>.

Vita

Marcos Suguru Kajita

kajitams@yahoo.com

Education

B.S. Computer Sciences, The University of Texas at Austin, May 2003

Work Experience

Software Engineer, Major Oil Service Company, 2003 – Present

Languages

Portuguese, Japanese, and English

Interests

Technology, sports, and movies

This report was typed by Marcos Suguru Kajita