

IBM Bluemix

The Cloud Platform for Creating and Delivering Applications

Raffaele Stifani



Cloud



International Technical Support Organization

**IBM Bluemix: The Cloud Platform for Creating and
Delivering Applications**

August 2015

Note: Before using this information and the product it supports, read the information in “Notices” on page vii.

First Edition (August 2015)

This edition applies to IBM Bluemix.

© Copyright International Business Machines Corporation 2015. All rights reserved.

Note to U.S. Government Users Restricted Rights -- Use, duplication or disclosure restricted by GSA ADP Schedule Contract with IBM Corp.

Contents

Notices	vii
Trademarks	viii
IBM Redbooks promotions	ix
Preface	xi
Author	xi
Now you can become a published author, too	xi
Comments welcome	xii
Stay connected to IBM Redbooks	xii
Chapter 1. Introduction to IBM Bluemix cloud app development platform	1
1.1 IBM Bluemix overview	2
1.1.1 Bluemix list of services	3
1.2 Bluemix architecture	4
1.2.1 Bluemix public	4
1.2.2 Bluemix Dedicated	6
1.2.3 How Bluemix works	7
1.2.4 Bluemix resiliency	10
1.2.5 Bluemix concepts	11
1.2.6 Integration with systems of record	14
1.3 Bluemix security	15
1.3.1 Security of the Bluemix platform	16
1.3.2 Functional security	16
1.3.3 Infrastructure security	17
1.3.4 Operational security	17
1.3.5 Physical security	18
1.3.6 Data security	18
1.3.7 Security of Bluemix applications	18
1.3.8 Bluemix security deployment	20
Chapter 2. Bluemix services	21
2.1 Boilerplates	22
2.1.1 Internet of Things Foundation Starter (IBM)	22
2.1.2 Java Cache Web Starter (IBM)	22
2.1.3 Java Cloudant Web Starter (IBM)	22
2.1.4 Java DB Web Starter (IBM)	22
2.1.5 Mobile Cloud (IBM)	22
2.1.6 MobileFirst Services Starter (IBM)	22
2.1.7 Node.js Cache Web Starter (IBM)	23
2.1.8 Node.js Cloudant DB Web Starter (IBM)	23
2.1.9 Personality Insights Java Web Starter (IBM)	23
2.1.10 Personality Insights Node.js Web Starter (IBM)	23
2.1.11 Node-RED Starter (Community)	23
2.1.12 Ruby Sinatra (Community)	23
2.1.13 Vaadin Rich Web Starter (Community)	23
2.1.14 WordPress on Bluemix (Community)	23
2.2 Runtimes	24
2.2.1 Liberty for Java	24

2.2.2	SDK for Node.js (IBM)	24
2.2.3	Go (Community)	24
2.2.4	PHP (Community)	24
2.2.5	Python (Community)	24
2.2.6	Ruby (Community)	24
2.2.7	Community buildpacks (Community)	24
2.3	IBM Watson	25
2.3.1	AlchemyAPI (IBM)	25
2.3.2	Concept Expansion (IBM)	25
2.3.3	Concept Insights (IBM)	25
2.3.4	Language Identification (IBM)	25
2.3.5	Machine Translation (IBM)	26
2.3.6	Personality Insights (IBM)	26
2.3.7	Question and Answer (IBM)	26
2.3.8	Relationship Extraction (IBM)	27
2.3.9	Speech to Text (IBM)	27
2.3.10	Text to Speech (IBM)	28
2.3.11	Tradeoff Analytics (IBM)	28
2.3.12	Visual Recognition (IBM)	28
2.3.13	Cognitive Commerce (third party)	29
2.3.14	Cognitive Graph (third party)	29
2.3.15	Cognitive Insights (third party)	30
2.4	Mobile	30
2.4.1	Advanced Mobile Access (IBM)	30
2.4.2	Cloudant NoSQL DB (IBM)	31
2.4.3	Mobile Application Security (IBM)	31
2.4.4	Mobile Data (IBM)	31
2.4.5	Mobile Quality Assurance (IBM)	33
2.4.6	Presence Insights (IBM)	33
2.4.7	Push (IBM)	35
2.4.8	Push iOS 8 (IBM)	35
2.4.9	Twilio (third party)	36
2.5	DevOps	36
2.5.1	Auto-Scaling (IBM)	36
2.5.2	Delivery Pipeline (IBM)	36
2.5.3	Monitoring and Analytics (IBM)	37
2.5.4	Track & Plan (IBM)	39
2.5.5	BlazeMeter (third party)	39
2.5.6	Load Impact (third party)	39
2.5.7	New Relic (third party)	39
2.6	Web and Application services	39
2.6.1	Business Rules (IBM)	40
2.6.2	Data Cache (IBM)	40
2.6.3	MQ Light (IBM)	41
2.6.4	Session Cache (IBM)	41
2.6.5	Workflow (IBM)	41
2.6.6	Workload Scheduler (IBM)	43
2.6.7	Box (third party)	43
2.6.8	CloudAMQP (third party)	43
2.6.9	DreamFace (third party)	43
2.6.10	Geocoding (third party)	43
2.6.11	Memcached Cloud (third party)	44
2.6.12	Namara.io Catalog (third party)	44

2.6.13	Reappt (third party)	44
2.6.14	Reverse Geocoding (third party)	44
2.6.15	SendGrid (third party)	44
2.6.16	Simplicite aPaaS for Bluemix (third party)	44
2.6.17	Statica (third party)	44
2.6.18	Travel Boundary Service (third party)	44
2.6.19	Ustream (third party)	45
2.6.20	Validate Address (third party)	45
2.7	Integration	45
2.7.1	API Management (IBM)	45
2.7.2	Cloud Integration (IBM)	46
2.7.3	Secure Gateway (IBM)	46
2.8	Data Management	48
2.8.1	Cloudant NoSQL DB (IBM)	48
2.8.2	DataWorks (IBM)	49
2.8.3	Object Storage (IBM)	50
2.8.4	Object Storage (v2) (IBM)	50
2.8.5	SQL Database (IBM)	52
2.8.6	ClearDB MySQL Database (third party)	52
2.8.7	ElephantSQL (third party)	52
2.8.8	MongoLab (third party)	52
2.8.9	Redis Cloud (third party)	53
2.9	Big Data	53
2.9.1	BigInsights for Apache Hadoop (IBM)	53
2.9.2	dashDB (IBM)	56
2.9.3	Geospatial Analytics (IBM)	60
2.9.4	IBM Analytics for Apache Hadoop (IBM)	60
2.9.5	Insights for Twitter (IBM)	63
2.9.6	Time Series Database (IBM)	64
2.10	Security	64
2.10.1	Application Security Manager (IBM)	64
2.10.2	AppScan Dynamic Analyzer (IBM)	66
2.10.3	AppScan Mobile Analyzer (IBM)	68
2.10.4	Mobile Analyzer for iOS (IBM)	71
2.10.5	Single Sign On (IBM)	74
2.10.6	Static Analyzer (IBM)	77
2.10.7	aPersona Adaptive Security Manager (third party)	80
2.11	Business Analytics	80
2.11.1	Embeddable Reporting (IBM)	80
2.11.2	Predictive Modeling (IBM)	83
2.11.3	Cupenya Insights (third party)	83
2.12	Internet of Things	83
2.12.1	Internet of Things (IBM)	83
2.12.2	flowthings.io (third party)	85
	Related publications	87
	IBM Redbooks	87
	Online resources	87
	Help from IBM	88

Notices

This information was developed for products and services offered in the U.S.A.

IBM may not offer the products, services, or features discussed in this document in other countries. Consult your local IBM representative for information on the products and services currently available in your area. Any reference to an IBM product, program, or service is not intended to state or imply that only that IBM product, program, or service may be used. Any functionally equivalent product, program, or service that does not infringe any IBM intellectual property right may be used instead. However, it is the user's responsibility to evaluate and verify the operation of any non-IBM product, program, or service.

IBM may have patents or pending patent applications covering subject matter described in this document. The furnishing of this document does not grant you any license to these patents. You can send license inquiries, in writing, to:

IBM Director of Licensing, IBM Corporation, North Castle Drive, Armonk, NY 10504-1785 U.S.A.

The following paragraph does not apply to the United Kingdom or any other country where such provisions are inconsistent with local law: INTERNATIONAL BUSINESS MACHINES CORPORATION PROVIDES THIS PUBLICATION "AS IS" WITHOUT WARRANTY OF ANY KIND, EITHER EXPRESS OR IMPLIED, INCLUDING, BUT NOT LIMITED TO, THE IMPLIED WARRANTIES OF NON-INFRINGEMENT, MERCHANTABILITY OR FITNESS FOR A PARTICULAR PURPOSE. Some states do not allow disclaimer of express or implied warranties in certain transactions, therefore, this statement may not apply to you.

This information could include technical inaccuracies or typographical errors. Changes are periodically made to the information herein; these changes will be incorporated in new editions of the publication. IBM may make improvements and/or changes in the product(s) and/or the program(s) described in this publication at any time without notice.

Any references in this information to non-IBM websites are provided for convenience only and do not in any manner serve as an endorsement of those websites. The materials at those websites are not part of the materials for this IBM product and use of those websites is at your own risk.

IBM may use or distribute any of the information you supply in any way it believes appropriate without incurring any obligation to you.

Any performance data contained herein was determined in a controlled environment. Therefore, the results obtained in other operating environments may vary significantly. Some measurements may have been made on development-level systems and there is no guarantee that these measurements will be the same on generally available systems. Furthermore, some measurements may have been estimated through extrapolation. Actual results may vary. Users of this document should verify the applicable data for their specific environment.

Information concerning non-IBM products was obtained from the suppliers of those products, their published announcements or other publicly available sources. IBM has not tested those products and cannot confirm the accuracy of performance, compatibility or any other claims related to non-IBM products. Questions on the capabilities of non-IBM products should be addressed to the suppliers of those products.

This information contains examples of data and reports used in daily business operations. To illustrate them as completely as possible, the examples include the names of individuals, companies, brands, and products. All of these names are fictitious and any similarity to the names and addresses used by an actual business enterprise is entirely coincidental.


COPYRIGHT LICENSE:

This information contains sample application programs in source language, which illustrate programming techniques on various operating platforms. You may copy, modify, and distribute these sample programs in any form without payment to IBM, for the purposes of developing, using, marketing or distributing application programs conforming to the application programming interface for the operating platform for which the sample programs are written. These examples have not been thoroughly tested under all conditions. IBM, therefore, cannot guarantee or imply reliability, serviceability, or function of these programs.

Trademarks

IBM, the IBM logo, and [ibm.com](http://www.ibm.com) are trademarks or registered trademarks of International Business Machines Corporation in the United States, other countries, or both. These and other IBM trademarked terms are marked on their first occurrence in this information with the appropriate symbol (® or ™), indicating US registered or common law trademarks owned by IBM at the time this information was published. Such trademarks may also be registered or common law trademarks in other countries. A current list of IBM trademarks is available on the Web at <http://www.ibm.com/legal/copytrade.shtml>

The following terms are trademarks of the International Business Machines Corporation in the United States, other countries, or both:

AppScan®	IBM®	QRadar®
BigInsights®	IBM MobileFirst™	Redbooks®
Bluemix™	IBM UrbanCode™	Redpaper™
Cloudant®	IBM Watson™	Redbooks (logo)  ®
Cognitive Commerce™	Informix®	SPSS®
dashDB™	InfoSphere®	WebSphere®
DataPower®	Insight™	z/OS®
DB2®	Insights™	
Guardium®	Optim™	

The following terms are trademarks of other companies:

AlchemyAPI, and AlchemyAI logo are trademarks or registered trademarks of AlchemyAI, Inc., an IBM Company.

Netezza, and N logo are trademarks or registered trademarks of IBM International Group B.V., an IBM Company.

SoftLayer, and SoftLayer device are trademarks or registered trademarks of SoftLayer, Inc., an IBM Company.

Linux is a trademark of Linus Torvalds in the United States, other countries, or both.

Java, and all Java-based trademarks and logos are trademarks or registered trademarks of Oracle and/or its affiliates.

Other company, product, or service names may be trademarks or service marks of others.

Find and read thousands of IBM Redbooks publications

- ▶ Search, bookmark, save and organize favorites
- ▶ Get up-to-the-minute Redbooks news and announcements
- ▶ Link to the latest Redbooks blogs and videos

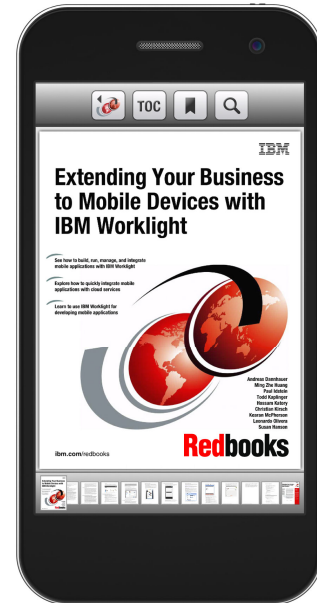
Get the latest version of the Redbooks Mobile App



Download
Now

iOS

Android



Promote your business in an IBM Redbooks publication

Place a Sponsorship Promotion in an IBM® Redbooks® publication, featuring your business or solution with a link to your web site.

Qualified IBM Business Partners may place a full page promotion in the most popular Redbooks publications. Imagine the power of being seen by users who download millions of Redbooks publications each year!



ibm.com/Redbooks

About Redbooks → Business Partner Programs

THIS PAGE INTENTIONALLY LEFT BLANK

Preface

This IBM® Redpaper™ publication gives readers a broad understanding of IBM Bluemix™ cloud application development platform capabilities. Providing a platform as a service (PaaS) environment as one of its run times, along with containers and virtual machines, Bluemix uses the Cloud Foundry project as one of its open source technologies to accelerate new application development and DevOps methods. It provides optimized and flexible workloads, enables continuous availability, and simplifies delivery and manageability of an application by providing prebuilt services and hosting capabilities.

The paper reviews the Bluemix architecture, explains how it works, describes key concepts and components, and provides an overview of Bluemix security. It also covers the various Bluemix service categories and the services within each category. This information will help anyone who is interested in exploring the potential and capabilities of Bluemix and its services.

Author

Raffaele Stifani is an IBM Executive Architect in Italy. He works as a Software Client Architect and provides support to enterprises, designing solutions in response to the business needs of the digital era. His current focus is on the innovative areas of cloud computing, big data and analytics, mobile computing, social business, security systems, business process management, and enterprise architecture. He is recognized inside and outside of IBM as an expert in complex systems integration, architectural design, and technology implementation. Raffaele has a wide IT experience across several industries and business organizations of all sizes. He has a proven ability to lead the implementation of innovative solutions in addition to establishing teams and groups to realize the strategic visions of the enterprise. He conducts seminars at universities, is the author of publications, and a speaker at conferences. He has a degree in engineering and holds the certification of Distinguished Architect by The Open Group. He is member of IBM Technical Expert Council in Italy and a board member of AICA (the Italian Computer Society). His 29-years career with IBM has included systems engineering, project management, defining architectures, and teaching positions.

Thanks to the following people for their contributions to this project:

LindaMay Patterson, Todd Kelsey, Vasfi Gucer, and Judith Broadhurst
IBM USA

Now you can become a published author, too

Here's an opportunity to spotlight your skills, grow your career, and become a published author, all at the same time. Join an ITSO residency project and help write a book in your area of expertise, while honing your experience using leading-edge technologies. Your efforts will help to increase product acceptance and customer satisfaction, as you expand your network of technical contacts and relationships. Residencies run from two to six weeks in length, and you can participate either in person or as a remote resident working from your home base.

Find out more about the residency program, browse the residency index, and apply online:

ibm.com/redbooks/residencies.html

Comments welcome

Your comments are important to us.

We want our papers to be as helpful as possible. Send us your comments about this paper or other IBM Redbooks® publications in one of the following ways:

- Use the online **Contact us** review Redbooks form:

ibm.com/redbooks

- Send your comments by email:

redbooks@us.ibm.com

- Mail your comments:

IBM Corporation, International Technical Support Organization
Dept. HYTD Mail Station P099
2455 South Road
Poughkeepsie, NY 12601-5400

Stay connected to IBM Redbooks

- Find us on Facebook:

<http://www.facebook.com/IBMRedbooks>

- Follow us on Twitter:

<http://twitter.com/ibmredbooks>

- Look for us on LinkedIn:

<http://www.linkedin.com/groups?home=&gid=2130806>

- Explore new Redbooks publications, residencies, and workshops with the IBM Redbooks weekly newsletter:

<https://www.redbooks.ibm.com/Redbooks.nsf/subscribe?OpenForm>

- Stay current on recent Redbooks publications with RSS feeds:

<http://www.redbooks.ibm.com/rss.html>



Introduction to IBM Bluemix cloud app development platform

This chapter provides an overview of IBM Bluemix cloud app development platform.

This chapter covers the following topics:

- ▶ IBM Bluemix overview
- ▶ Bluemix architecture
- ▶ Bluemix security

1.1 IBM Bluemix overview

Bluemix is the IBM open cloud platform that offers mobile and web developers access to IBM software for integration, security, transactions, and other key functions, as well as software from business partners. Built on Cloud Foundry open source technology, Bluemix makes application development easier with Platform as a Service (PaaS). Bluemix also provides prebuilt Mobile Backend as a Service (MBaaS) capabilities. The goal is to simplify the delivery of an application by providing services that are ready for immediate use and hosting capabilities to enable internal scale development.

Bluemix also offers cloud deployments that fit your needs. Whether you are a small business that plans to scale or a large enterprise that requires additional isolation, you can develop in a cloud without borders, where you can connect your dedicated services to the public Bluemix services available from IBM and third-party providers. All service instances are managed by IBM. You will get one bill for only what you choose to use. With the broad set of services and runtimes in Bluemix, the developer gains control and flexibility, and has access to various data options, from predictive analytics to big data.

Bluemix provides the following features:

- ▶ A range of services that enable you to build and extend web and mobile apps fast
- ▶ Processing power for you to deliver app changes continuously
- ▶ Fit-for-purpose programming models and services
- ▶ Manageability of services and applications
- ▶ Optimized and elastic workloads
- ▶ Continuous availability

Bluemix abstracts and hides most of the complexities that are associated with hosting and managing cloud-based applications. As an application developer, you can focus on developing your application without having to manage the infrastructure that is required to host it. For both mobile and web apps, you can use the prebuilt services that are provided by Bluemix. You can upload your web app to Bluemix and indicate how many instances that you want running. After your apps are deployed, you can easily scale them up or down when the use or load of the apps change.

You can use Bluemix to quickly develop apps in the most popular programming languages. You can develop mobile apps in iOS, Android, and HTML with JavaScript. For web apps, you can use languages such as Ruby, PHP, Java, Go, and Python.

Bluemix also provides middleware services for your applications to use. Bluemix acts on the application's behalf when it provisions new service instances, and then binds those services to the application. Your app can perform its real job, leaving the management of the services to the infrastructure.

In general, you do not have to worry about the operating system and infrastructure layers when running applications on Bluemix. Layers, such as root file systems and middleware components, are abstracted so that you can focus on your application code.

The services are grouped into the categories shown in Table 1-1, which the sections that follow describe.

Table 1-1 Categories of Bluemix services

<ul style="list-style-type: none">▶ Boilerplates▶ Runtimes▶ Watson▶ Mobile	<ul style="list-style-type: none">▶ DevOps▶ Web and Application Services▶ Integration▶ Data Management	<ul style="list-style-type: none">▶ Big Data▶ Security▶ Business Analytics▶ Internet of Things
---	---	---

1.1.1 Bluemix list of services

Table 1-2 lists Bluemix services by service name, grouped by category.

Table 1-2 Bluemix list of services

<p>Boilerplates</p> <ul style="list-style-type: none"> - Internet of Things Foundation Starter (IBM) - Java Cache Web Starter (IBM) - Java Cloudant Web Starter (IBM) - Java DB Web Starter (IBM) - Node.js Cloudant DB Web Starter (IBM) - IBM Personality Insights™ Java Web Starter (IBM) - Personality Insights Node.js Web Starter (IBM) - Node-RED Starter (Community) - Ruby Sinatra (Community) - Vaadin Rich Web Starter (Community) - WordPress on Bluemix (Community) <p>Runtimes</p> <ul style="list-style-type: none"> - Liberty for Java (IBM) - SDK for Node.js (IBM) - Go (Community) - PHP (Community) - Python (Community) - Ruby (Community) <p>Watson</p> <ul style="list-style-type: none"> - AlchemyAPI® (IBM) - Concept Expansion (IBM) - Concept Insights (IBM) - Language Identification (IBM) - Machine Translation (IBM) - Personality Insights (IBM) - Question and Answer (IBM) - Relationship Extraction (IBM) - Speech To Text (IBM) - Text to Speech (IBM) - Tradeoff Analytics (IBM) - Visual Recognition (IBM) - IBM Cognitive Commerce™ (third party) - Cognitive Graph (third party) - Cognitive Insights (third party) <p>Mobile</p> <ul style="list-style-type: none"> - Advanced Mobile Access (IBM) - Cloudant NoSQL DB (IBM) - Mobile Application Security (IBM) - Mobile Data (IBM) - Mobile Quality Assurance (IBM) - Presence Insights (IBM) - Push (IBM) - Push iOS 8 (IBM) - Twilio (third party) <p>DevOps</p> <ul style="list-style-type: none"> - Auto-Scaling (IBM) - Delivery Pipeline (IBM) - Monitoring and Analytics (IBM) - Track & Plan (IBM) - BlazeMeter (third party) - Load Impact (third party) - New Relic (third party) 	<p>Web and Application Services</p> <ul style="list-style-type: none"> - Business Rules (IBM) - Data Cache (IBM) - IBM MQ Light (IBM) - Session Cache (IBM) - Workflow (IBM) - Workload Scheduler (IBM) - Box (third party) - CloudAMQP (third party) - DreamFace (third party) - Geocoding (third party) - Memcached Cloud (third party) - Namara.io catalog (third party) - Reappt (third party) - Reverse Geocoding (third party) - SendGrid (third party) - Simplicite aPaaS for Bluemix (third party) - Statica (third party) - Travel Boundary Service (third party) - Ustream (third party) - Validate Address (third party) <p>Integration</p> <ul style="list-style-type: none"> - API Management (IBM) - Cloud Integration (IBM) - Secure Gateway (IBM) - Insights for Twitter (IBM) - IBM Informix® Time Series database (IBM) <p>Data Management</p> <ul style="list-style-type: none"> - Cloudant NoSQL DB (IBM) - DataWorks (IBM) - Object Storage (IBM) - Object Storage (v2) (IBM) - SQL database (IBM) - ClearDB MySQL database (third party) - ElephantSQL (third party) - MongoLab (third party) - Redis Cloud (third party) <p>Big Data</p> <ul style="list-style-type: none"> - IBM BigInsights® for Hadoop (IBM) - IBM dashDB™ (IBM) - Geospatial Analytics (IBM) - IBM Analytics for Hadoop (IBM) <p>Security</p> <ul style="list-style-type: none"> - Application Security Manager (IBM) - IBM AppScan® Dynamic Analyzer (IBM) - AppScan Mobile Analyzer (IBM) - Mobile Analyzer for iOS (IBM) - Single Sign On (IBM) - Static Analyzer (IBM) - aPersona Adaptive Security Manager (third party) <p>Business Analytics</p> <ul style="list-style-type: none"> - Embeddable Reporting (IBM) - Predictive Modeling (IBM) - Cupenya Insights (third party) <p>Internet of Things</p> <ul style="list-style-type: none"> - Internet of Things (IBM) - flowthings.io (third party)
---	--

1.2 Bluemix architecture

With Bluemix, you can access the public Bluemix platform, set up a dedicated Bluemix platform or use both.

1.2.1 Bluemix public

Bluemix is an environment to develop apps and use services that provide ready-to-use functions. Bluemix also provides an environment to host application artifacts that run on an application server, such as the IBM WebSphere® Application Server Liberty profile.

By using IBM SoftLayer® global cloud infrastructure, Bluemix deploys virtual containers that host each deployed app. In this environment, the app can use prebuilt services (including third-party services) to make app assembly easy.

As a developer, you can interact with the Bluemix infrastructure by using a browser-based interface. You can also use a Cloud Foundry command-line interface, called *cf*, to deploy web apps.

Clients can be mobile apps, apps that run externally, apps that are built on Bluemix, or developers who are using browsers. All interact with Bluemix-hosted apps. Clients use REST or HTTP APIs to route requests through Bluemix to one of the app instances or to composite services.

Figure 1-1 shows the high-level Bluemix architecture.

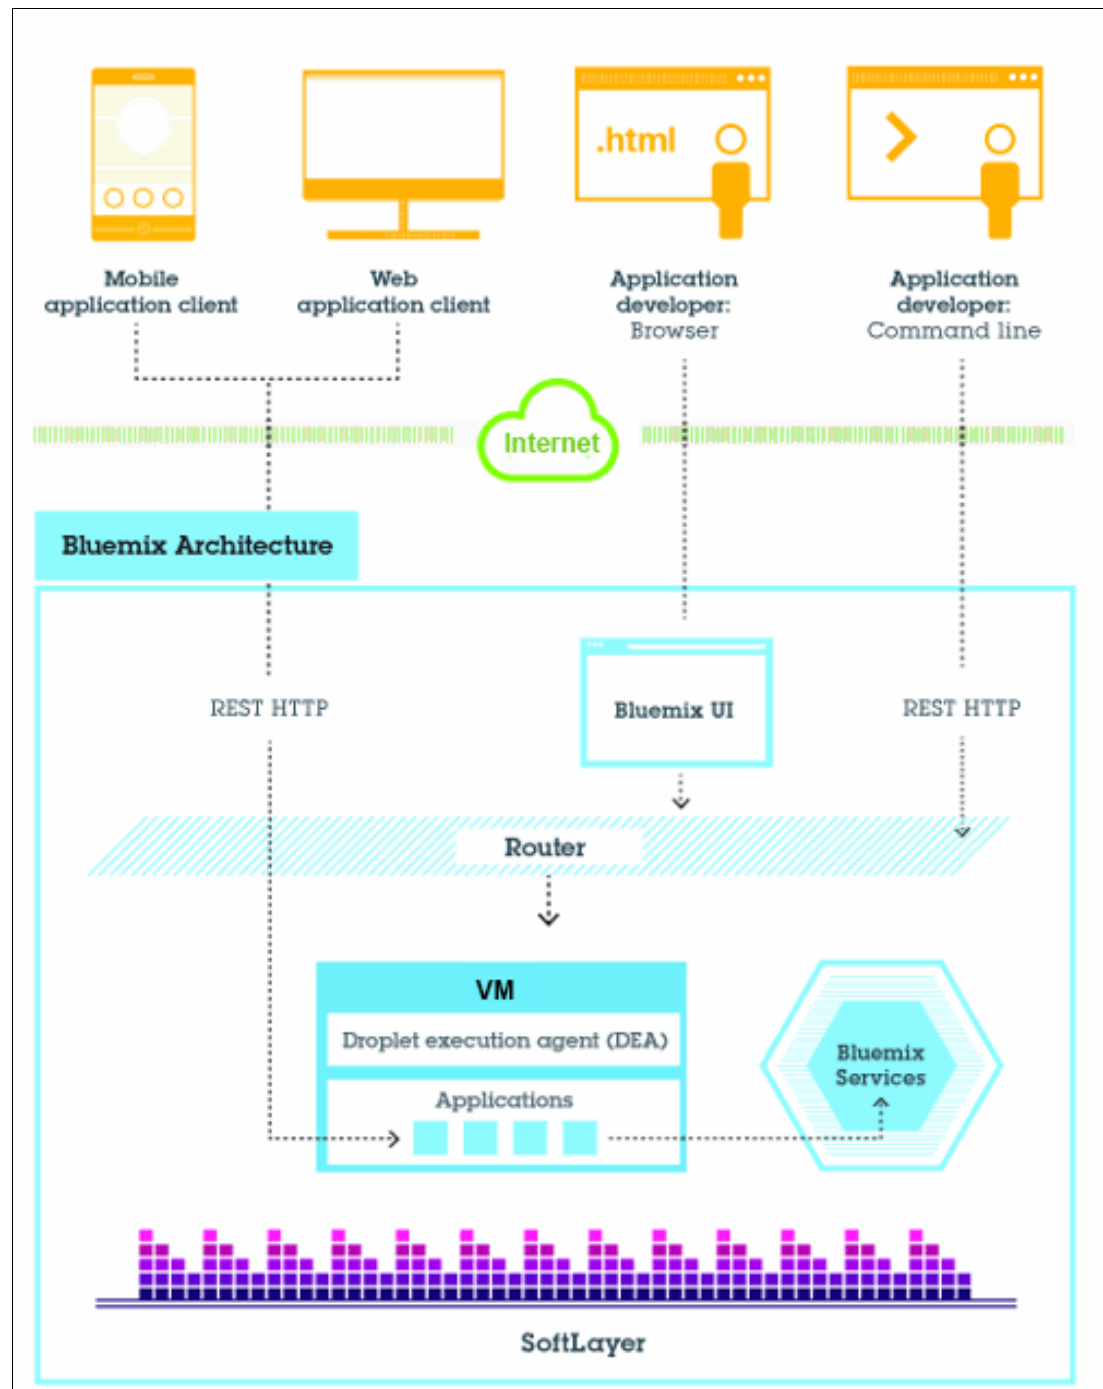


Figure 1-1 Bluemix architecture

You can deploy your apps to different Bluemix regions (Figure 1-2) for latency or security considerations. You can choose to deploy either to one region or across multiple regions.

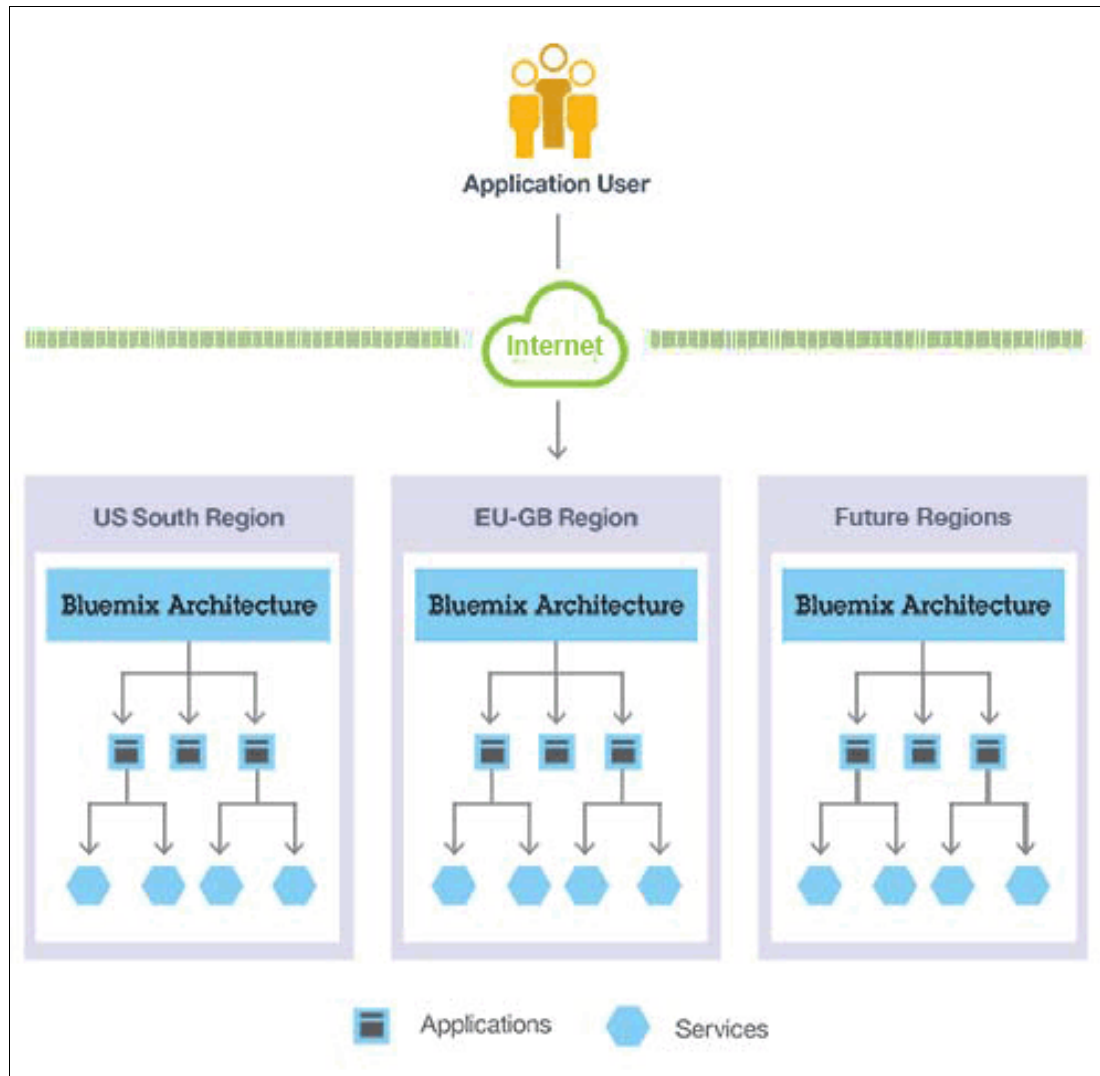


Figure 1-2 Multi-region application deployment

1.2.2 Bluemix Dedicated

Bluemix Dedicated is your own exclusive SoftLayer environment that's securely connected to both the public Bluemix and your own network. Bluemix Dedicated sits on your network through a virtual private network (VPN) or a direct network connection. Your single-tenant hardware can be set up in any SoftLayer data center around the world. IBM specialists manage the dedicated platform and dedicated services, so you can focus on building custom apps. They also perform all maintenance to dedicated instances by using a maintenance window that you select.

Several services and all run times are available in your dedicated environment. You can connect to all public services, too.

All dedicated deployments of Bluemix include the following features and benefits at no additional cost: VPN, private virtual local area network (VLAN), firewall, connectivity with your LDAP directory, ability to use existing on-premises databases and apps, 24x7 onsite security, dedicated hardware, and standard support.

Figure 1-3 shows an overview of the Bluemix Dedicated environment.

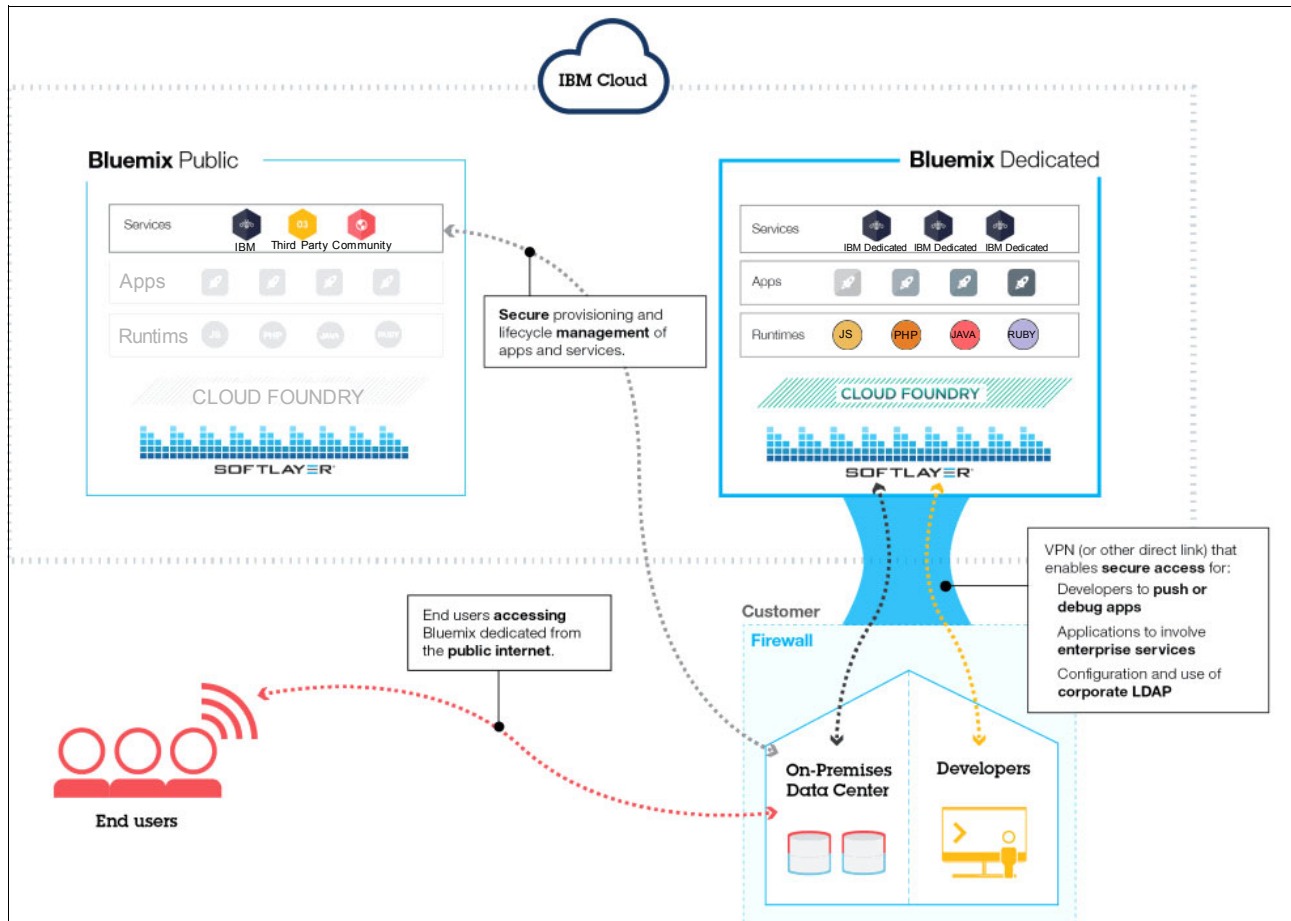


Figure 1-3 Bluemix Dedicated

1.2.3 How Bluemix works

When you deploy an app to Bluemix, you must configure Bluemix with enough of the essential information to support the app:

- ▶ For a mobile app, Bluemix contains an artifact that represents the mobile app's backend, such as the services that are used by the mobile app to communicate with a server.
- ▶ For a web app, you must ensure that information about the proper run time and framework is communicated to Bluemix, so that it can set up the proper execution environment to run the app.

Each execution environment, including both mobile and web, is isolated from the execution environment of other apps. The execution environments are isolated even though these apps are on the same physical machine.

Figure 1-4 shows the basic flow of how Bluemix manages the deployment of apps.

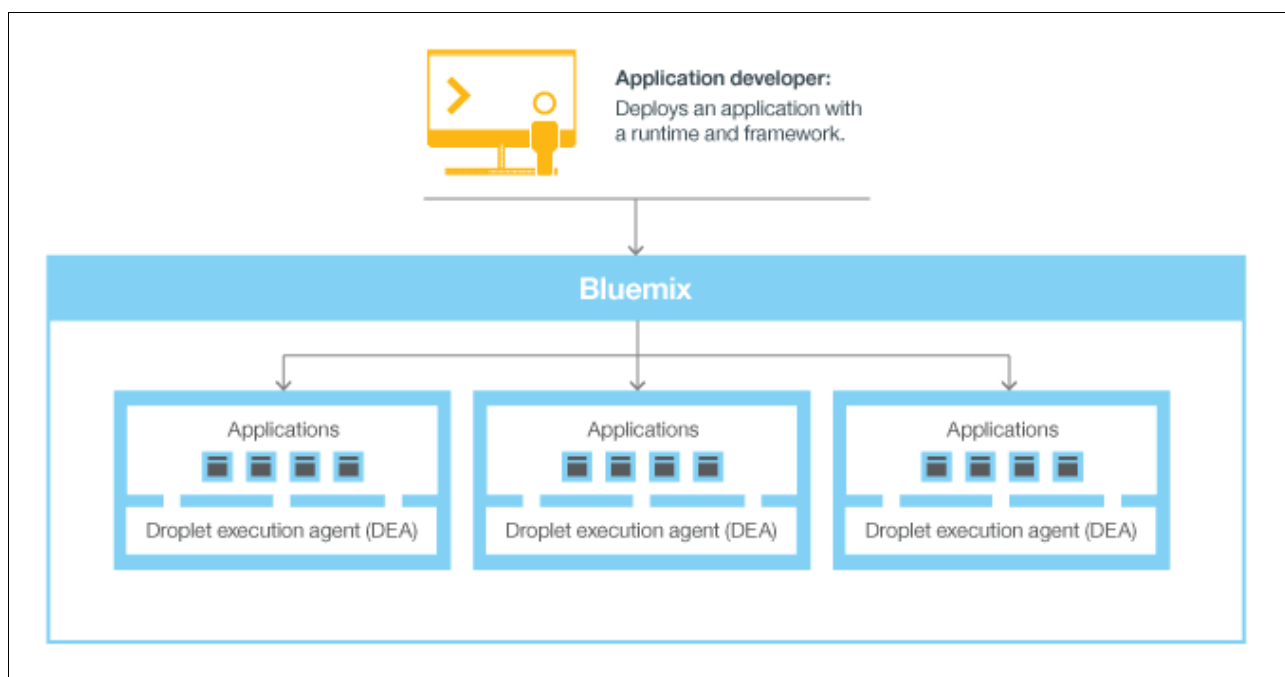


Figure 1-4 Deploying an application

When you create an app and deploy it to Bluemix, the Bluemix environment determines an appropriate virtual machine (VM) to which the app (or artifacts that the app represents) is sent. For a mobile app, a mobile backend projection is created on Bluemix. Any code for the mobile app that is running in the cloud eventually runs in the Bluemix environment. For a web app, the code running in the cloud is the app that the developer deploys to Bluemix. The determination of which VM to use is based on several factors, including the load already on the machine and run times or frameworks supported by that VM. After a VM is chosen, an application manager on each VM installs the correct framework and run time for the app. Then, the app can be deployed into that framework. When the deployment is complete, the application artifacts are started.

Figure 1-5 shows the structure of a VM, also known as *droplet execution agent* (DEA), that has multiple apps deployed to it.

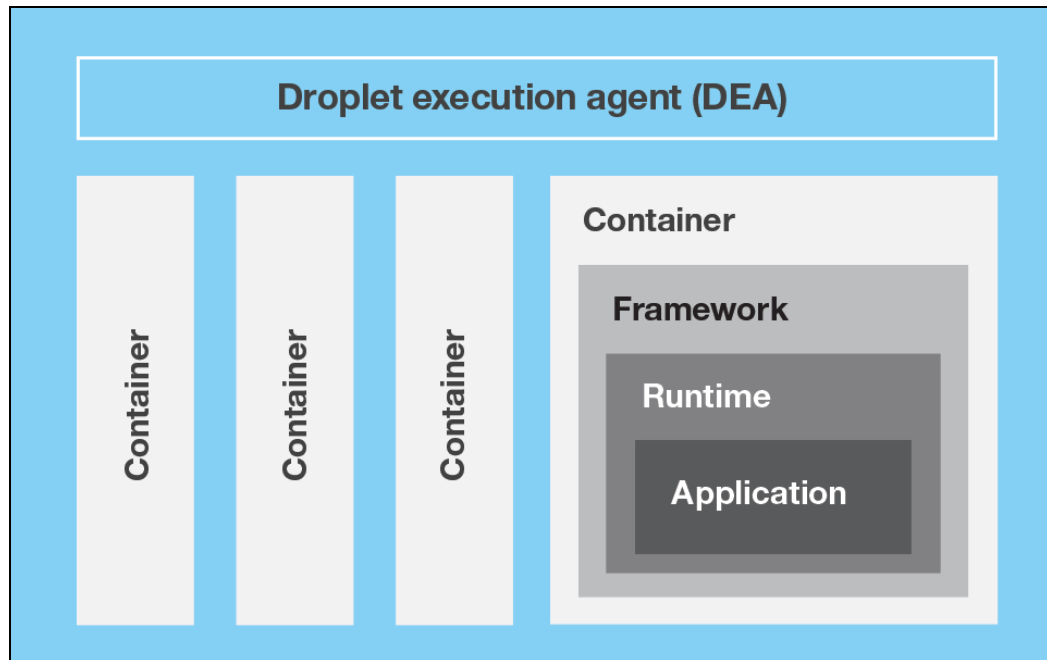


Figure 1-5 Design of a VM

In each VM, an application manager communicates with the rest of the Bluemix infrastructure, and manages the apps that are deployed to this VM. Each VM has containers to separate and protect apps. In each container, Bluemix installs the framework and run time that are required for each app.

When the app is deployed, if it has a web interface (as for a Java web app) or other REST-based services (such as mobile services displayed publicly to the mobile app), users of the app can communicate with it by using normal HTTP requests, as the diagram in Figure 1-6 on page 10 illustrates.

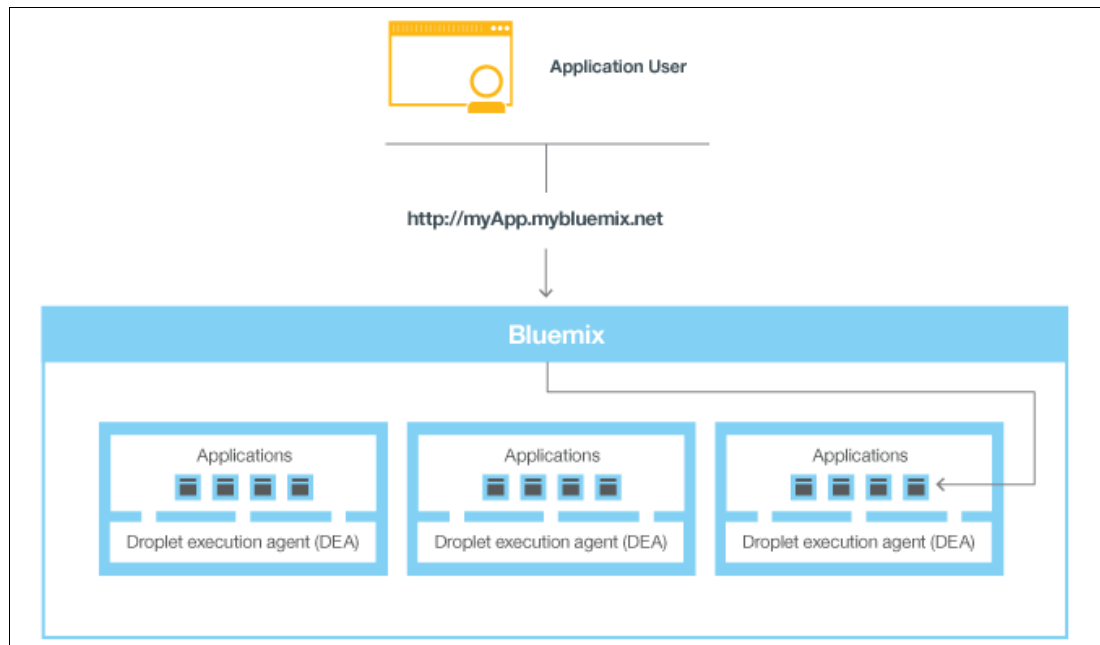


Figure 1-6 Invoking a Bluemix application

Each app can have one or more URLs associated with it, but all of them must point to the Bluemix endpoint. When a request comes in, Bluemix examines the request, determines which app it is intended for, and then selects one of the instances of the app to receive the request.

1.2.4 Bluemix resiliency

Bluemix is designed to host scalable, resilient apps and application artifacts that can scale to meet your needs and remain highly available and quick to recover from problems. Bluemix separates those components that track the state of interactions (stateful) from those that do not (stateless). This separation gives Bluemix the flexibility to move apps as necessary to maintain scalability and resiliency.

You can have one or more instances of your app running. When you have multiple instances of one app, the app is uploaded only once. However, Bluemix deploys the number of instances of the app requested and distributes the instances across as many VMs as possible.

You must save all persistent data in a stateful data store that is outside of your app, such as on one of the data store services that are provided by Bluemix. Because anything cached in memory or on disk might not be available even after a restart, you can use the memory space or file system of a single Bluemix instance as a brief, single-transaction cache. With a single instance setup, the request to your app might be interrupted because of the stateless nature of Bluemix. Therefore, it is best to use at least three instances for each app to ensure availability of your app.

All Bluemix infrastructure, Cloud Foundry components, and IBM-specific management components are highly available. Multiple instances of the infrastructure are used to balance the load.

1.2.5 Bluemix concepts

Bluemix consists of applications, services, buildpacks, and other components. You can deploy applications to different Bluemix regions by using one IBM ID.

Regions

A Bluemix *region* is a defined geographical territory that you can deploy your applications to. You can create applications and service instances in different regions with the same Bluemix infrastructure for application management and get the same usage details view for billing. You can select the region that is nearest to your customers and deploy your apps to this region to get low application latency. You can also select the region where you want to keep the application data to address security issues. When you build apps in multiple regions, if one region goes down, the apps that are in the other regions continue to run. Your resource allowance is the same for each region that you use.

If you are using the Bluemix user interface, you can switch to a different region to work with the spaces in that region. If you are using the Cloud Foundry command-line interface (CLI), you must connect to the Bluemix region that you want to work with by using the `cf api` command and specifying the API endpoint of the region. For example, enter the following command to connect to Bluemix Europe United Kingdom region:

```
cf api https://api.eu-gb.bluemix.net
```

If you are using the Eclipse tools, you must connect to the Bluemix region that you want to work with by creating a Bluemix server and specifying the API endpoint of the region. A unique prefix is assigned to each region. Bluemix provides the following regions and region prefixes (see Table 1-3).

Table 1-3 Bluemix region list

Region name	Region prefix	CF API endpoint	UI console
US South region	us-south	api.ng.bluemix.net	console.ng.bluemix.net
Europe United Kingdom region	eu-gb	api.eu-gb.bluemix.net	console.eu-gb.bluemix.net

Infrastructure

Bluemix offers three ways for you to run your code: Cloud Foundry, IBM Containers, and virtual machines. You can pick the correct infrastructure for deploying your applications.

Cloud Foundry

Apps running in the Cloud Foundry infrastructure work with existing Cloud Foundry apps, and can bind to any of the services that are available in the Bluemix catalog. With this infrastructure, you develop and manage your application code, and Bluemix takes care of the management and maintenance of the infrastructure that powers those apps.

IBM Containers

With the IBM Containers service, you can run your web app anywhere that supports container deployment. A *container* is an object that holds everything that is needed for an app to run. This infrastructure includes a private registry for your trusted images so that you can upload, store, and retrieve them. Then, you can make those images available in Bluemix and manage your containers on the platform. The IBM Containers infrastructure scales both horizontally and vertically.

You can use all of the images that are available in the public Docker Hub and use the docker API and command-line interface to manage your containers on Bluemix. IBM also provides some public images in the Containers Registry, which you can use and extend.

IBM Containers are used to run Docker containers in a hosted cloud environment. Docker adds an engine that deploys an app to the virtual environment that you use for running your containers. Docker also provides an environment that you can use to run your code. When you are ready, it provides the means by which you can transfer the code from your development environment, to your test environment, and then to your production environment.

Virtual machines (in beta)

The Bluemix virtual machines infrastructure gives you the ability to create and manage virtual machine groups on the IBM public cloud. You can also create and manage VM groups on your private IBM clouds that you have chosen to make available to Bluemix users. The infrastructure includes a guided experience to connect to your on-premises infrastructure. Support for monitoring and logging is integrated with Bluemix. You can deploy and manage your virtual machines by using either the Bluemix user interface or the cloud's OpenStack APIs.

Virtual machines on Bluemix support provisioning of virtual machine groups with auto-scaling. Through this support, the number of instances can be automatically increased or decreased, based on CPU load or the failure of an instance. Load balancing is also supported, which enables the assignment of virtual IP (floating IP) addresses as needed.

Applications

In Bluemix, an application, or *app*, represents the artifact that a developer is building. The application lifecycle in Bluemix and Cloud Foundry are identical, regardless of how you push the app to the Bluemix.

Mobile apps

Mobile apps run outside of the Bluemix environment and use services that the mobile apps are exposed to. These services typically act in concert and represent the backend projection of that app. Bluemix can also host app code that the developer would rather run on a backend server in a container-based environment.

Web apps

Web apps consist of all code that is required to be run or referenced at run time. Web apps are uploaded to Bluemix to host the app.

For languages, such as Java, where the source code is compiled into runtime binary files, only the binary files are required to be uploaded.

Services

A *service* is a cloud extension that is hosted by Bluemix. The service provides functions that are ready for use by the app's running code. The predefined services that are provided by Bluemix include database, messaging, push notifications for mobile apps, and flexible caching for web apps.

You can also create your own services in Bluemix, and they can vary in complexity. Your services can be simple utilities, such as the functions that you might see in a runtime library. Or they can be complex business logic that you might see in a business process modeling service or a database.

Bluemix simplifies the use of services by provisioning new instances of the service and binding those service instances to your app. The management of the service is handled automatically by Bluemix.

Boilerplates

In Bluemix, a *boilerplate* contains an app and its associated runtime environment and predefined services for a particular domain. You can use a boilerplate to quickly get up and running. For example, you can select the Mobile Cloud boilerplate to host mobile and web apps and accelerate development time of server-side scripts by using the Mobile App template and SDK.

Runtimes

A *run time* is the set of resources that is used to run an app. Bluemix provides runtime environments as containers for different types of apps. The runtime environments are integrated into Bluemix as buildpacks and are automatically configured for use.

Buildpacks

A *buildpack* is a collection of scripts that prepare your code for execution on the target cloud platform as a service (PaaS). A buildpack gathers the runtime and framework dependencies of an app. Then, it packages them with the app into a *droplet* that can be deployed to the cloud.

If you do not specify a buildpack when you deploy your app to Bluemix, one of three kinds of built-in buildpacks is used by default:

- Built-in IBM buildpacks

Table 1-4 shows the list of built-in buildpacks that are created by IBM.

Table 1-4 Bluemix built-in buildpacks

Name	Git repository URL
Liberty for Java	Built-in
Node.js	Built-in

- Built-in community buildpacks

In Bluemix, you can also use built-in buildpacks that are provided by the Cloud Foundry community. To list built-in community buildpacks, run the **cf buildpacks** command.

- External buildpacks

If you cannot find the run time or framework that you want in the built-in buildpacks that are provided by Bluemix, you can use an external, existing buildpack for your app. External buildpacks are provided by the Cloud Foundry community for you to use as your own buildpacks. You specify the buildpack when you deploy your app by using the **cf push** command.

Starters

A *starter* is a template that includes predefined services and app code that is configured with a particular buildpack. A starter might be app code that is written in a specific programming language or a combination of app code and a set of services. There are two types of starters: *boilerplates* and *runtimes*.

1.2.6 Integration with systems of record

Bluemix can help developers by connecting two broad categories of systems in a cloud environment: Systems of record and systems of engagement:

- ▶ *Systems of record* include applications and databases that store business records and automate standardized processes.
- ▶ *Systems of engagement* are capabilities that expand the usefulness of systems of record and make them more engaging to users.

By integrating a system of record with the application that you create in Bluemix, you can perform the following actions:

- ▶ Enable secure communication between the application and the backend database by downloading and installing a secure connector on premise.
- ▶ Invoke a database in a secure way.
- ▶ Create APIs from integration flows with databases and backend systems, such as customer relationship management system.
- ▶ Show only the schemas and tables that you want to be exposed to the application.
- ▶ As a Bluemix organization manager, publish an API as a private service that is visible only to your organization members..

To integrate a system of record with the applications that you create in Bluemix, use the Cloud Integration service. By using this service, you can create a Cloud Integration API and publish the API as a private service for your organization:

- ▶ Cloud Integration API

A Cloud Integration API provides secured access to the systems of record that reside behind a firewall through web APIs. When you create the Cloud Integration API, you choose the resource that you want to access through the web API, specify the operations that are permitted, and include software development kits (SDKs) and samples to access the API.

- ▶ Private service

A private service consists of a Cloud Integration API, SDKs, and entitlement policies. The private service might also contain documentation or other items from the service provider. Only the organization manager can publish a Cloud Integration API as a private service. To see the private services that are available to you, select the **Private** check box in the Bluemix catalog. You can select and bind a private service to an application without connecting to the Cloud Integration service. You bind private services to your application in the same way as you do for other Bluemix services.

Scenario: Create a mobile app and connect to system of record

Assume that you're an application developer who wants to evaluate the scenario of creating a rich mobile application to connect to your system of record. Bluemix provides a platform where you can integrate your mobile app, cloud services, and enterprise systems of record to provide an application that interacts with your on-premises data.

For example, you can build a mobile app to interact with your customer relationship management system that resides on-premises behind a firewall. You can invoke the system of record in a secure way and use the mobile services in Bluemix so that you can build a rich mobile application.

First, your integration developer creates the mobile backend application in Bluemix. He uses the Mobile Cloud boilerplate that uses the Node.js run time that he is most familiar with.

Then, by using the Cloud Integration service in the Bluemix user interface, he exposes an API through a secure connector. Your integration developer downloads the secure connector and installs it on your premises to enable secure communication between his API and the database. After he creates the database endpoint, he can look at all the schemas and extract the tables that he wants to expose as APIs to the application.

Your integration developer adds the Push service to deliver mobile notifications to interested consumers. He also adds a business partner service to tweet when a new customer record is created with a Twitter API.

Next, as the application developer, you can log in to Bluemix, download the Android development toolkit, and develop code that invokes the APIs that your integration developer created. You can develop a mobile application that enables users to enter their information on their mobile devices. The mobile app then creates customer records in the customer management system. When the records are created, the application pushes a notification to a mobile devices and initiates a tweet about each new record.

1.3 Bluemix security

Designed with secure engineering practices, the Bluemix platform has layered security controls across network and infrastructure. Bluemix also provides a suite of security services that can be used by application developers to secure their mobile and web apps. These elements combine to make IBM Bluemix a platform with clear choices for secure application development.

Bluemix ensures security readiness by adhering to security policies that are driven by preferred practices in IBM for systems, networking, and secure engineering. These policies include practices such as source code scanning, dynamic scanning, threat modeling, and penetration testing. Bluemix follows the IBM Product Security Incident Response Team (PSIRT) process for security incident management.

Bluemix uses the IBM SoftLayer infrastructure-as-a-service (IaaS) cloud service and takes full advantage of its security architecture. SoftLayer IaaS provides multiple, overlapping tiers of protection for your applications and data. Bluemix adds security capabilities at the PaaS layer in different categories: platform, data, and application.

1.3.1 Security of the Bluemix platform

Bluemix provides functional, infrastructure, operational, and physical security (through IBM SoftLayer) for the core platform.

The Bluemix environment on SoftLayer is compliant with the most restrictive IBM information technology (IT) security standards, including the following, which meet or exceed the industry standards:

- ▶ Network, data encryption and access control
- ▶ Application access control lists (ACLs), permissions, and penetration testing
- ▶ Identification, authentication, and authorization
- ▶ Information and data protection
- ▶ Service integrity and availability
- ▶ Vulnerability and fix management
- ▶ Denial of service and systematic attacks detection
- ▶ Security incident response

1.3.2 Functional security

Bluemix provides various functional security capabilities, including user authentication, access authorization, auditing of critical operations, and data protection:

- ▶ Authentication

Application developers are authenticated to Bluemix by using the IBM web identity.

For Bluemix Dedicated, authentication through LDAP is supported by default. On request, authentication through IBM web identity can be set up instead.

- ▶ Authorization

Bluemix uses Cloud Foundry mechanisms to ensure that each application developer has access only to the applications and service instances that they created. Authorization to Bluemix services is based on OAuth. Access to all Bluemix Platform internal endpoints are restricted to external users.

- ▶ Auditing

Audit logs are created for all successful and unsuccessful authentication attempts of application developers. Audit logs are also created for privileged access to Linux systems that host the containers where Bluemix applications run.

- ▶ Data protection

All Bluemix traffic goes through IBM WebSphere DataPower® SOA Appliances, which provide reverse proxy, SSL termination, and load balancing functions. The following HTTP methods are allowed: DELETE, GET, HEAD, OPTIONS, POST, PUT, TRACE.

If HTTP is inactive, it times out at two minutes.

- ▶ Secure development practices

Periodic security vulnerability scans are performed on various Bluemix components by using IBM Security AppScan Dynamic Analyzer and static analyzer offerings. Threat modeling and penetration testing are performed to detect and address any potential vulnerabilities. In addition, application developers can use the AppScan Dynamic Analyzer service to secure their web apps that are deployed on Bluemix.

1.3.3 Infrastructure security

Bluemix builds upon Cloud Foundry to provide a robust foundation for running your applications. Within the architecture, several components are provided for security and isolation:

- ▶ Environment segregation
Development and production environments are segregated from each other to improve application stability and security.
- ▶ Firewalls
Firewalls are in place to restrict access to the Bluemix network.
- ▶ Intrusion protection
Bluemix enables intrusion protection to discover threats so that they can be addressed. Intrusion protection policies are enabled on firewalls.
- ▶ Secure application container management
Each Bluemix application is isolated and runs in its own container that has specific resource limits for processor, memory, and disk.
- ▶ Operating system security hardening
IBM administrators regularly perform network and operating system hardening by using tools such as IBM Endpoint Manager.

In addition, change management and backup and recovery procedures are implemented to ensure integrity and availability.

1.3.4 Operational security

Bluemix provides a robust operational security environment with the following controls:

- ▶ Vulnerability scan
Bluemix uses the Tenable Network Security vulnerability scanning tool, Nessus, to detect any issues with network and host configurations so that the issues can be resolved.
- ▶ Automated fix management
Bluemix administrators ensure that fixes for operating systems are applied at appropriate frequencies. Automated fixes are enabled by using IBM Endpoint Manager.
- ▶ Audit log consolidation and analysis
Bluemix uses the IBM QRadar® Security Intelligence Platform tools to consolidate Linux logs to monitor privileged access on Linux systems. Bluemix also uses IBM QRadar security information and event management (SIEM) to monitor successful and unsuccessful login attempts of application developers.
- ▶ User access management
Within Bluemix, Segregation of Duties guidelines are followed to assign granular access privileges to users and to ensure that users have only the access that is required to perform their jobs according to the principle of least privilege.

Within a Bluemix dedicated environment, assigned administrators can manage roles and permissions for Bluemix user in their organization by using the Admin console.

1.3.5 Physical security

Bluemix relies on the *network-within-a-network* topology of SoftLayer for physical network security. This architecture ensures that network-within-a-network systems are fully accessible only to authorized personnel.

In SoftLayer network-within-a-network architecture, the public network layer handles public traffic to hosted web sites or online resources. The private network layer allows for true out-of-band management through a distinct stand-alone third carrier over Secure Sockets Layer (SSL), Point-to-Point Tunneling Protocol (PPTP), or Internet Protocol Security (IPSec) VPN gateways. The network layer from data center to data center provides free-of-charge and secure connectivity between servers that are housed in separate SoftLayer facilities.

Every SoftLayer data center is fully secured with controls that meet Statement on Standards for Attestation Engagements No. 16 (SSAE 16) and industry-recognized requirements, without exceptions.

1.3.6 Data security

With Bluemix, securing your data against unauthorized access is a joint effort between Bluemix and you.

Data that is associated with a running application can be in one of three states: data-in-transit, data-at-rest, and data-in-use.

Data-in-transit	Data that is being transferred between nodes on a network.
Data-at-rest	Data that is stored.
Data-in-use	Data not currently stored that is being acted upon at an endpoint.

Each type of data needs to be considered when you plan for data security.

The Bluemix platform secures data-in-transit by securing user access to the application by using SSL, through the network until the data reaches IBM DataPower Gateway at the boundary of the Bluemix internal network. DataPower Gateway acts as a reverse proxy and provides SSL termination.

Note: Security for both data-in-use and data-at-rest is your responsibility as you develop your application. You can take advantage of several data-related services available in the Bluemix Catalog to help with these concerns.

1.3.7 Security of Bluemix applications

As an application developer, you must enable the security configurations, including application data protection, for your applications that run on Bluemix.

You can use security capabilities that are provided by several Bluemix services to secure your applications. All Bluemix services that are produced by IBM follow IBM secure engineering development practices.

IBM Single Sign On for Bluemix is a policy-based authentication service that provides an easy to embed single sign-on capability for Node.js or Liberty for Java applications. To enable an application developer to embed single sign-on capability into an application, the administrator creates service instances and adds identity sources.

The Single Sign On service supports several identity sources where users' credentials are stored:

SAML Enterprise	A user registry with an exchange of Security Assertion Markup Language (SAML) tokens that completes the authentication
Cloud Directory	A user registry that is hosted on the IBM Cloud platform (see http://www.ibm.com/cloud-computing/ for more information)
Social identity sources	The user registries that are maintained by Google, Facebook, and LinkedIn

IBM AppScan Mobile Analyzer

This service provides a security analysis of Android mobile applications. To use this service, you must upload a compiled Android app as an Android application package (APK) file. When the security analysis scan is done, you can download a report.

AppScan Dynamic Analyzer

This service provides a security analysis of web applications with a dynamic analysis tool. The tool works on the deployed web app, not on the app source code, and it can scan any Bluemix web app regardless of its language or technology. You can scan only applications of the organizations that you belong to. To create a scan, you must configure the web app URL and the login credentials, if any. When the scan is finished, you can download a report.

Mobile Analyzer for iOS (beta)

The Mobile Analyzer for iOS service provides AppScan dynamic security analysis for iOS mobile applications. It helps you identify security issues in your iOS mobile apps.

Static Analyzer (beta)

The Static Analyzer service enables static application security testing on the cloud. It helps you find source code vulnerabilities early in the software development lifecycle, so that they can be fixed before deployment.

Static Analyzer enables you to scan Java and Java web content by using a command-line interface (CLI) on your local disk. In addition, you can run a small installer that adds Static Analyzer plug-ins to Eclipse or Maven. You can use the client utility to scan and gather information about your files in an archive file that you then submit to the cloud for scan results.

IBM UrbanCode plug-in for application security testing

The IBM Application Security Testing for Bluemix plug-in enables you to run security scans on your web or Android apps that are hosted on Bluemix. This plug-in is developed and supported by the IBM UrbanCode™ Deploy Community on the IBM Bluemix DevOps Services platform.

SQL Database

The SQL Database service adds a fully provisioned relational database to your app. This service uses IBM Security Directory Server for authentication and IBM InfoSphere® Guardium® Data Activity Monitor to protect the database that is accessed by applications. The connection between applications and the database is protected by the SSL certificate that DigiCert signs.

In certain plans with this service, you can use the SQL database console in Bluemix to get reports that contain the following information:

- ▶ Sensitive data that might exist in the database that is accessed by applications
- ▶ The application users who accessed the database within a specified period
- ▶ The application users who are accessing sensitive data that is in the database

To mask data by using SQL, applications can call the masking user-defined functions (UDFs) that are deployed together with the database. For example, you can mask the data that you want to use elsewhere for testing. The UDFs implement the data masking algorithms from IBM InfoSphere Optim™ data management solutions.

The premium plan for this service also includes data encryption.

dashDB

The dashDB service uses Security Directory Server for user authentication and IBM InfoSphere Guardium Data Activity Monitor to protect the database that is accessed by applications. The connection between applications and the database is protected by SSL certificates. This service uses the IBM DB2® database native encryption capability to automatically encrypt your deployed database and database backups. Master key rotation is automatic and happens every 90 days.

Cloud Integration

The Cloud Integration service enables you to integrate cloud and on-premises data. You can add a service to interact with backend databases, such as DB2, Oracle, and SAP. Next, you can move data or create REST APIs for Bluemix applications to access and use. The service enables secure communication with on-premises secure connectors and displays backend systems of record as REST APIs to be used by applications.

Secure Gateway

The Secure Gateway service enables you to securely connect Bluemix apps to remote locations, either on-premises or in the cloud. It provides secure connectivity and establishes a tunnel between your Bluemix organization and the remote location that you want to connect to. You can configure and create a secure gateway by using the Bluemix user interface or an API package.

1.3.8 Bluemix security deployment

Bluemix security deployment architecture includes different information flows for app users and developers to ensure secure access.

This is the information flow for Bluemix app users:

1. Through a firewall, with intrusion prevention and network security in place
2. Through the DataPower Gateway with reverse proxy and SSL termination proxy
3. Through the network router
4. Reaches the application run time in the droplet execution agent (DEA)



Bluemix services

This chapter describes the various services that are available with Bluemix. The services are presented by category and identified by service name.

This chapter includes details about the following Bluemix service categories:

- ▶ Boilerplates
- ▶ Runtimes
- ▶ IBM Watson
- ▶ Mobile
- ▶ DevOps
- ▶ Web and Application services
- ▶ Integration
- ▶ Data Management
- ▶ Big Data
- ▶ Security
- ▶ Business Analytics
- ▶ Internet of Things

2.1 Boilerplates

In IBM Bluemix, a *boilerplate* is a container for an application and its associated runtime environment and predefined services for a particular domain. You can use a boilerplate to quickly get up and running. This section identifies the following boilerplates and provides a brief description of each of them:

- ▶ Internet of Things Foundation Starter (IBM)
- ▶ Java Cache Web Starter (IBM)
- ▶ Java Cloudant Web Starter (IBM)
- ▶ Java DB Web Starter (IBM)
- ▶ Mobile Cloud (IBM)
- ▶ MobileFirst Services Starter (IBM)
- ▶ Node.js Cache Web Starter (IBM)
- ▶ Node.js Cloudant DB Web Starter (IBM)
- ▶ Personality Insights Java Web Starter (IBM)
- ▶ Personality Insights Node.js Web Starter (IBM)
- ▶ Node-RED Starter (Community)
- ▶ Ruby Sinatra (Community)
- ▶ Vaadin Rich Web Starter (Community))
- ▶ WordPress on Bluemix (Community)

2.1.1 Internet of Things Foundation Starter (IBM)

Get started with an Internet of Things Foundation application by using Node-RED in Bluemix. Try the sample flow with a simulator and customize it for your own devices.

2.1.2 Java Cache Web Starter (IBM)

This sample application demonstrates how to use the IBM DataCache service Java Native API with the Java Web runtime on the IBM Cloud platform.

2.1.3 Java Cloudant Web Starter (IBM)

This application demonstrates how to use the Cloudant NoSQL database service with the IBM WebSphere Application Server Liberty profile for Java runtime on IBM Cloud.

2.1.4 Java DB Web Starter (IBM)

This sample application demonstrates how to use Java JPA to connect to a SQL database.

2.1.5 Mobile Cloud (IBM)

Quickly get started with your next mobile app by using mobile services for shared data, push, and server-side scripting. This service includes software development kits (SDKs) for Android, iOS, and JavaScript.

2.1.6 MobileFirst Services Starter (IBM)

Start building your next iOS app with the IBM MobileFirst™ for iOS platform.

2.1.7 Node.js Cache Web Starter (IBM)

This application demonstrates how to use the IBM DataCache service with the Node.js runtime environment on the IBM Cloud platform.

2.1.8 Node.js Cloudant DB Web Starter (IBM)

This application demonstrates how to connect and work with a Cloudant DB service from a node web application, with a Node.js runtime environment on IBM Cloud.

2.1.9 Personality Insights Java Web Starter (IBM)

This application is a simple Oracle Java app that uses the IBM Watson™ Personality Insights service to analyze text to derive personality traits.

2.1.10 Personality Insights Node.js Web Starter (IBM)

This application is a simple Node.js app that uses the IBM Watson Personality Insights service to analyze text to derive personality traits.

2.1.11 Node-RED Starter (Community)

This application demonstrates how to run the Node-RED open source project within IBM Bluemix.

2.1.12 Ruby Sinatra (Community)

This boilerplate is a simple Sinatra app that will get you up and running quickly.

2.1.13 Vaadin Rich Web Starter (Community)

This application demonstrates how to use the Vaadin UI Framework to build rich HTML5 applications with the Liberty for Java runtime environment and an IBM DB2 database on IBM Cloud.

2.1.14 WordPress on Bluemix (Community)

A WordPress boilerplate that bundles clearDB and Object Storage to deploy WordPress on Bluemix in minutes.

2.2 Runtimes

Runtime services provide the environment to run an app in the language of your choice. The following services and languages are included in the Runtimes category:

- ▶ Liberty for Java
- ▶ SDK for Node.js (IBM)
- ▶ Go (Community)
- ▶ PHP (Community)
- ▶ Python (Community)
- ▶ Ruby (Community)
- ▶ Community buildpacks (Community)

2.2.1 Liberty for Java

Develop, deploy, and scale Java web apps. The Liberty profile is a highly composable, ultra-fast, ultra-light profile of IBM WebSphere Application Server designed for the cloud.

2.2.2 SDK for Node.js (IBM)

Develop, deploy, and scale server-side JavaScript apps. The IBM SDK for Node.js provides enhanced performance, security, and serviceability.

2.2.3 Go (Community)

This is a simple Go application that gets you up and running quickly.

2.2.4 PHP (Community)

This is a simple PHP application that gets you up and running quickly.

2.2.5 Python (Community)

This is a simple Python application that gets you up and running quickly.

2.2.6 Ruby (Community)

The focus of this one is on quickly creating web applications in Ruby with minimal effort.

2.2.7 Community buildpacks (Community)

This provides some custom, Cloud Foundry-compatible buildpacks.

2.3 IBM Watson

The IBM Watson category provides services to build cognitive apps that help enhance, scale, and accelerate human expertise. The following services are available in the Watson category:

- ▶ AlchemyAPI (IBM)
- ▶ Concept Expansion (IBM)
- ▶ Concept Insights (IBM)
- ▶ Language Identification (IBM)
- ▶ Machine Translation (IBM)
- ▶ Personality Insights (IBM)
- ▶ Question and Answer (IBM)
- ▶ Relationship Extraction (IBM)
- ▶ Speech to Text (IBM)
- ▶ Text to Speech (IBM)
- ▶ Tradeoff Analytics (IBM)
- ▶ Visual Recognition (IBM)
- ▶ Cognitive Commerce (third party)
- ▶ Cognitive Graph (third party)
- ▶ Cognitive Insights (third party)

2.3.1 AlchemyAPI (IBM)

Build smart apps that use natural language processing and computer vision to deeply understand the world's conversations, documents, and photos.

2.3.2 Concept Expansion (IBM)

The Concept Expansion service analyzes text and interprets its meaning based on use in other similar contexts. For example, it could interpret “The Big Apple” as meaning “New York City.” It can be used to create a dictionary of related words and concepts so that euphemisms, colloquialisms, or otherwise unclear phrases can be better understood and analyzed.

2.3.3 Concept Insights (IBM)

The Concept Insights service links documents that you provide with a pre-existing graph of concepts based on Wikipedia (for example “The New York Times” and “machine learning”). Two types of links are identified by explicit links when a document directly mentions a concept and implicit links that connect your documents to relevant concepts that are not directly mentioned in them. Users of this service can also search for documents that are relevant to a concept or collection of concepts by exploring the explicit and implicit links.

2.3.4 Language Identification (IBM)

The Language Identification service detects the language in which text is written. This service helps inform next steps, such as translation, voice to text, or direct analysis. The service can be used in tandem with the Machine Translation service. Today, the language identification service identifies plain text as having been written in one of the following languages:

Afrikaans, Albanian, Arabic, Azerbaijani, Bashkir, Belarusian, Bulgarian, Bengali, Bosnian, Chinese, Traditional Chinese, Czech, Chuvash, Danish, Dutch, German, Greek, English, Esperanto, Spanish, Estonian, Basque, Farsi or Persian, Finnish, French, Gujarati, Hebrew, Hindi, Haitian, Hungarian, Armenian, Indonesian, Icelandic, Italian, Japanese, Georgian,

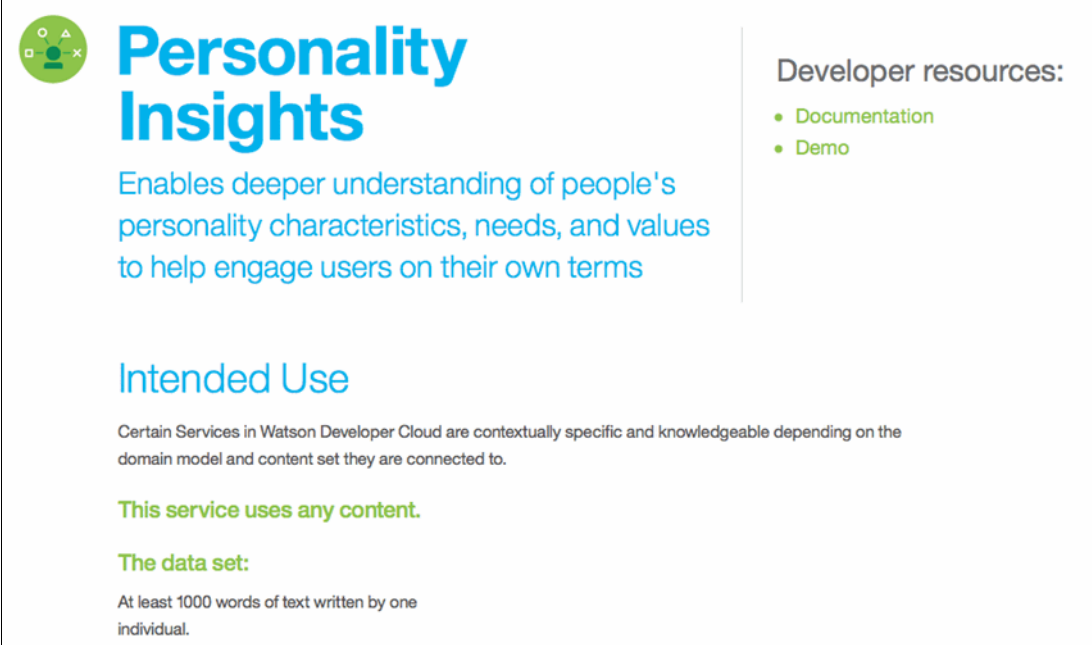
Kazakh, Central Khmer, Korean, Kurdish, Kirghiz, Lithuanian, Latvian, Malayalam, Mongolian, Norwegian Bokmal, Norwegian Nynorsk, Panjabi, Polish, Pushto, Portuguese, Romanian, Russian, Slovakian, Somali, Swedish, Tamil, Telugu, Turkish, Ukrainian, Urdu, and Vietnamese.

2.3.5 Machine Translation (IBM)

The Machine Translation service converts text input in one language into a destination language for the user. Translation is available for these languages: English, Brazilian Portuguese, Spanish, French, and Arabic.

2.3.6 Personality Insights (IBM)

The Personality Insights service (Figure 2-1) derives insights from transactional and social media data to identify psychological traits that influence purchase decisions, intent, and behavioral traits. The service is used to improve conversion rates.



Personality Insights

Enables deeper understanding of people's personality characteristics, needs, and values to help engage users on their own terms

Developer resources:

- [Documentation](#)
- [Demo](#)

Intended Use

Certain Services in Watson Developer Cloud are contextually specific and knowledgeable depending on the domain model and content set they are connected to.

This service uses any content.

The data set:

At least 1000 words of text written by one individual.

Figure 2-1 Personality Insights

2.3.7 Question and Answer (IBM)


This IBM Watson service interprets and answers user questions directly based on primary data sources (including, brochures, web pages, manuals, and records) that have been selected and gathered into a body of data or *corpus*. The service returns candidate responses with associated confidence levels and links to supporting evidence. The current data corpora on Bluemix focus on the travel and healthcare industries.

2.3.8 Relationship Extraction (IBM)

The Relationship Extraction service parses sentences into their various components and detects relationships between the components. It can process terms (that it has never analyzed before (such as people's names in a news feed) through contextual analysis. Sentence components include parts of speech (such as noun, verb, adjective, and conjunction) and functions (such as subjects, objects, and predicates). The service maps the relationships between the components so that users or analytics engines can more easily understand the meaning of individual sentences and documents.

2.3.9 Speech to Text (IBM)

The Speech to Text service (Figure 2-2) converts the human speech into written words. This easy-to-use service uses machine intelligence to combine information about grammar and language structure with knowledge of the composition of the audio signal to generate a more accurate transcription. The transcription is continuously sent back to the client and retroactively updated as more speech is heard. Recognition models can be trained for different languages and for specific domains.

 **Speech to Text**

Transcribes speech to text with low latency and high accuracy

Developer resources:

- [Documentation](#)
- [Demo](#)

Intended Use

Certain Services in Watson Developer Cloud are contextually specific and knowledgeable depending on the domain model and content set they are connected to.

[This service uses any content.](#)

The data set:

- Intelligible English speech

You can input:

streamed or recorded audio

And the service will output:

text transcriptions of speech

Figure 2-2 Speech to Text service

2.3.10 Text to Speech (IBM)

The Text to Speech service understands text and natural language to generate synthesized audio output, complete with appropriate cadence and intonation. It is available in three voices: Two US English voices, including the voice used by Watson in the 2011 Jeopardy match, and one Spanish voice.

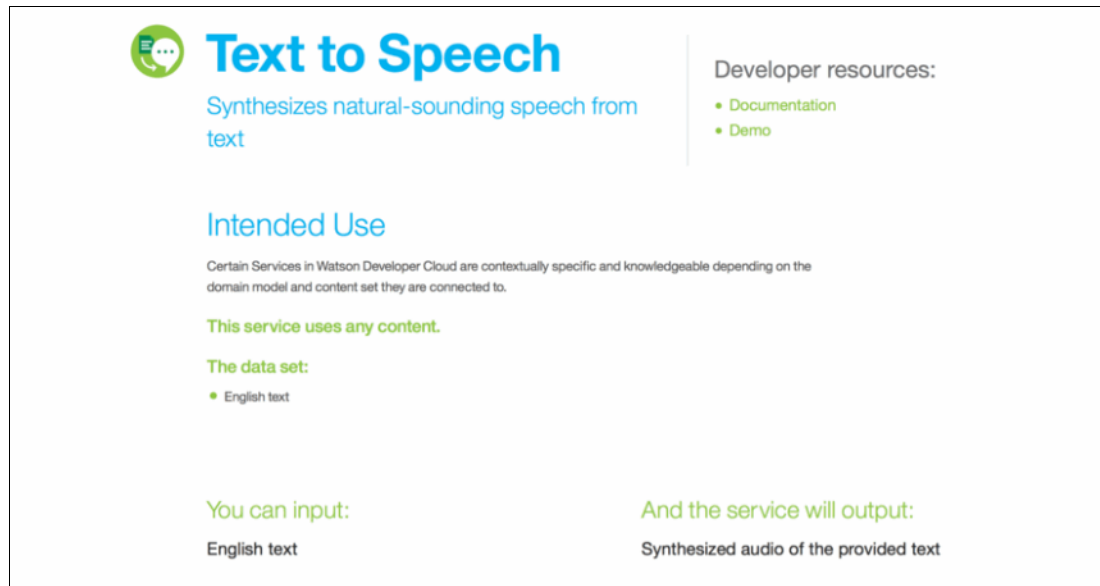


Figure 2-3 Text to Speech service

2.3.11 Tradeoff Analytics (IBM)

Tradeoff Analytics helps people make better choices while taking into account multiple, often conflicting, goals that matter when making that choice. The service can be used to help make complex decisions, such as what mortgage offer to take, and for helping with more everyday decisions, such as which laptop to purchase. Tradeoff Analytics uses Pareto filtering techniques to identify the optimal alternatives across multiple criteria. It then uses various analytical and visual approaches to help the decision maker explore the tradeoffs within the optimal set of alternatives. This ensures that the chosen option adheres to the goals and criteria that matter to the decision maker.

2.3.12 Visual Recognition (IBM)

The Visual Recognition service enables you to analyze the visual appearance of images or video frames to understand what is happening in a scene. Using machine learning technology, semantic classifiers recognize many visual entities, such as settings, objects, and events. The service applies these prelearned models to imagery that you uploaded to the service and returns a score for each image for each model, indicating the likelihood of that visual element being present in the image.

2.3.13 Cognitive Commerce (third party)

Cognitive Commerce, provided by CognitiveScale, provides insights and advice that can guide personalized and contextual commerce opportunities for consumers and businesses and provide more relevant and actionable recommendations. Cognitive Commerce, a service provided by CognitiveScale, provides the next level of fulfillment and helps drive commercial transactions.

Cognitive Commerce includes the following key capabilities:

- ▶ Intelligent, self-learning systems
 - Data mining, analysis, pattern recognition, and natural language processing
- ▶ Industry-specific content
 - Navigation through multi-structured data, including text, images and video to generate actionable insights from *dark data* (“Dark data” is data that is not machine readable, not easily accessible, or not used.)
- ▶ Location awareness
 - Surface context-relevant offers, deals and recommendations.
- ▶ Open standards foundation
 - Complete data security, sovereignty, and transparency while maintaining maximum flexibility and control over how and where data is stored and accessed

2.3.14 Cognitive Graph (third party)

A service provided by CognitiveScale, Cognitive Graphs are an encapsulation of knowledge that is sourced from third party, internal, and private data sources, using domain-specific models and placed in a graph representation that can be queried. Sourcing Agents pull, enrich, transform, and map multi-structured and dark data by using machine learning techniques. The Cognitive Graph can be projected in multiple ways to be applied to different problem sets.

A Cognitive Graph includes three main capabilities:

- ▶ Automatically evolves and learns
 - The graph is rapidly updated as new data streams in, models evolve, and new concepts are learned to create “living” applications.
- ▶ Universal understanding
 - The graph fuses domain-specific, concept-relationship ontologies with a representation of natural language, which provides a stable and unified basis for storing knowledge.
- ▶ Traceability
 - Knowledge represented in the graph can be traced all the way back to the original data sources, through all levels of enrichment, for authenticity and visibility into the data lineage behind the applications.

2.3.15 Cognitive Insights (third party)

A service provided by CognitiveScale, Cognitive Insights are contextually relevant and personalized observations or predictions that are presented with a recommendation, with the purpose of invoking user action. An insight augments the knowledge, perception, and awareness of the user, with the goals of improving their efficiency, decisions, and ability to react quickly to emerging conditions. When used in conjunction with a basis of knowledge and real-time data provided by the cognitive graph, this service can be used to add value to end-user applications.

Cognitive Insights includes these main attributes:

- ▶ Relevant
Insights are prioritized by comparing the impact and confidence in the prediction within the context of a domain model that is applied to real-time data.
- ▶ Personalized
An ever-growing user profile is used to improve the relevance of the insights that surface.
- ▶ Learning
Insights can get smarter when the user provides feedback to validate the insight.

2.4 Mobile

The Mobile category provides services to quickly get started with your next app. It includes these services, which are described in the subsections that follow:

- ▶ Advanced Mobile Access (IBM)
- ▶ Cloudant NoSQL DB (IBM)
- ▶ Mobile Application Security (IBM)
- ▶ Mobile Data (IBM)
- ▶ Mobile Quality Assurance (IBM)
- ▶ Presence Insights (IBM)
- ▶ Push (IBM)
- ▶ Push iOS 8 (IBM)
- ▶ Twilio (third party)

2.4.1 Advanced Mobile Access (IBM)

You can fine-tune Apple iOS 8 apps with operational analytics results from real-time performance and data use. You can also ensure that communications with your backend systems are secure, and let your app users log in with custom accounts or existing social accounts

2.4.2 Cloudant NoSQL DB (IBM)

Cloudant NoSQL DB provides access to a fully managed NoSQL JavaScript Object Notation (JSON) data layer that's always on. This service is compatible with CouchDB and accessible through a simple-to-use HTTP interface for mobile and web application models. This service includes the following main capabilities:

- ▶ **Ease of use**

Work with self-describing JSON documents through a RESTful API that makes every document in your Cloudant database accessible as JSON via a URL. Documents can be retrieved, stored, or deleted individually or in bulk and can also have files attached. IBM takes care of the provisioning, management, and scalability of the data store, which frees your time to focus on your application.

- ▶ **Powerful search, sync, and more**

With extremely powerful indexing, real time MapReduce and Apache Lucene-based full-text search, Cloudant NoSQL DB makes it easy to add advanced data analytics and powerful data access. Data access can also extend to Cloudant Sync to enable data access from mobile devices and client apps to run connected or offline.

2.4.3 Mobile Application Security (IBM)

Manage application access and implement a basic application security framework for the Mobile Cloud services with IBM Mobile Application Security (MAS) service for Bluemix. These important capabilities are part of this service:

- ▶ **Control access from the mobile app**

Identify users and devices that can access your Mobile Cloud services.

- ▶ **Block unwanted access**

Make sure that unauthorized, compromised, or lost devices can no longer reach your Mobile Cloud services and data.

2.4.4 Mobile Data (IBM)

Mobile Data (powered by Cloudant) provides simple-to-use SDKs that give you access to a scalable, fully managed database via familiar object-oriented APIs. Mobile Data includes these main capabilities:

- ▶ **Ease of use**

Work with data in the way that mobile developers are accustomed to, with iOS, Android, and JavaScript SDKs that provide powerful object-oriented APIs rather than raw database access. Provisioning, management, and scalability of the data store are taken care of by IBM, freeing your time to focus on your application.

- ▶ **Streamlined backend access**

No knowledge of databases, schemas, or SQL is required. Work with your data as objects that you can create, update, delete, and query, and let the SDKs handle cloud-based data storage and retrieval for you.

Figure 2-4 shows the Manage Data view.

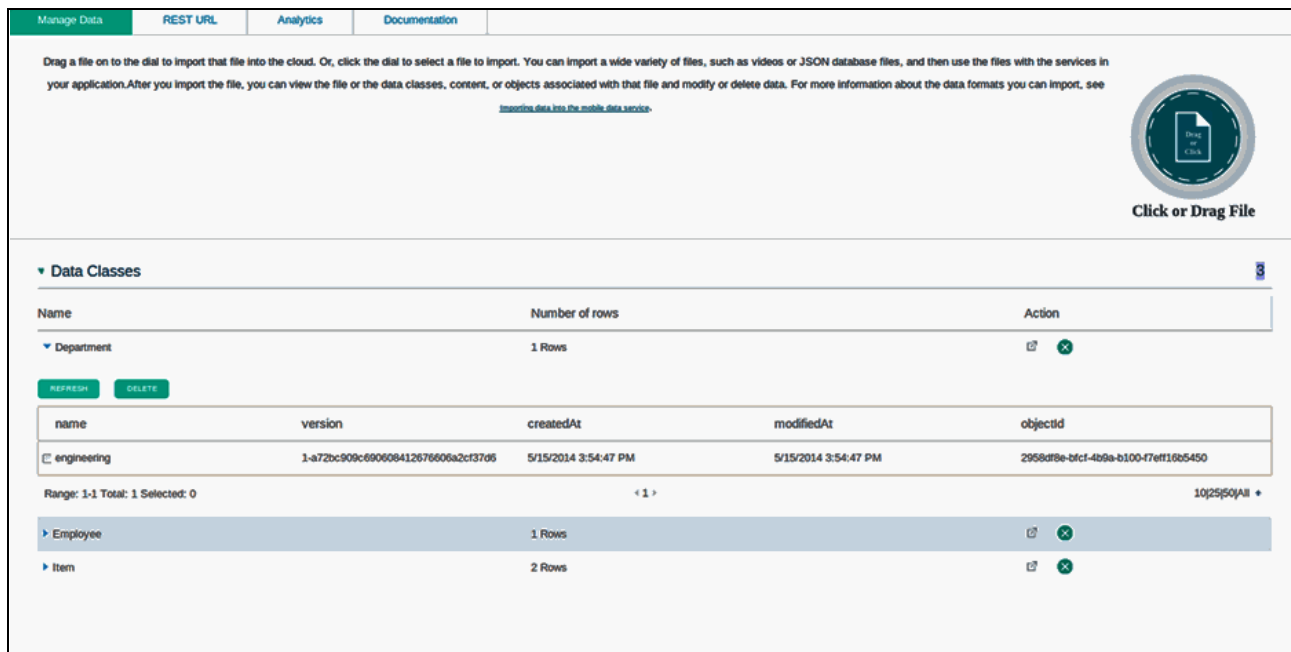


Figure 2-4 Manage Data view

Figure 2-5 shows the Manage Data Analytics view.

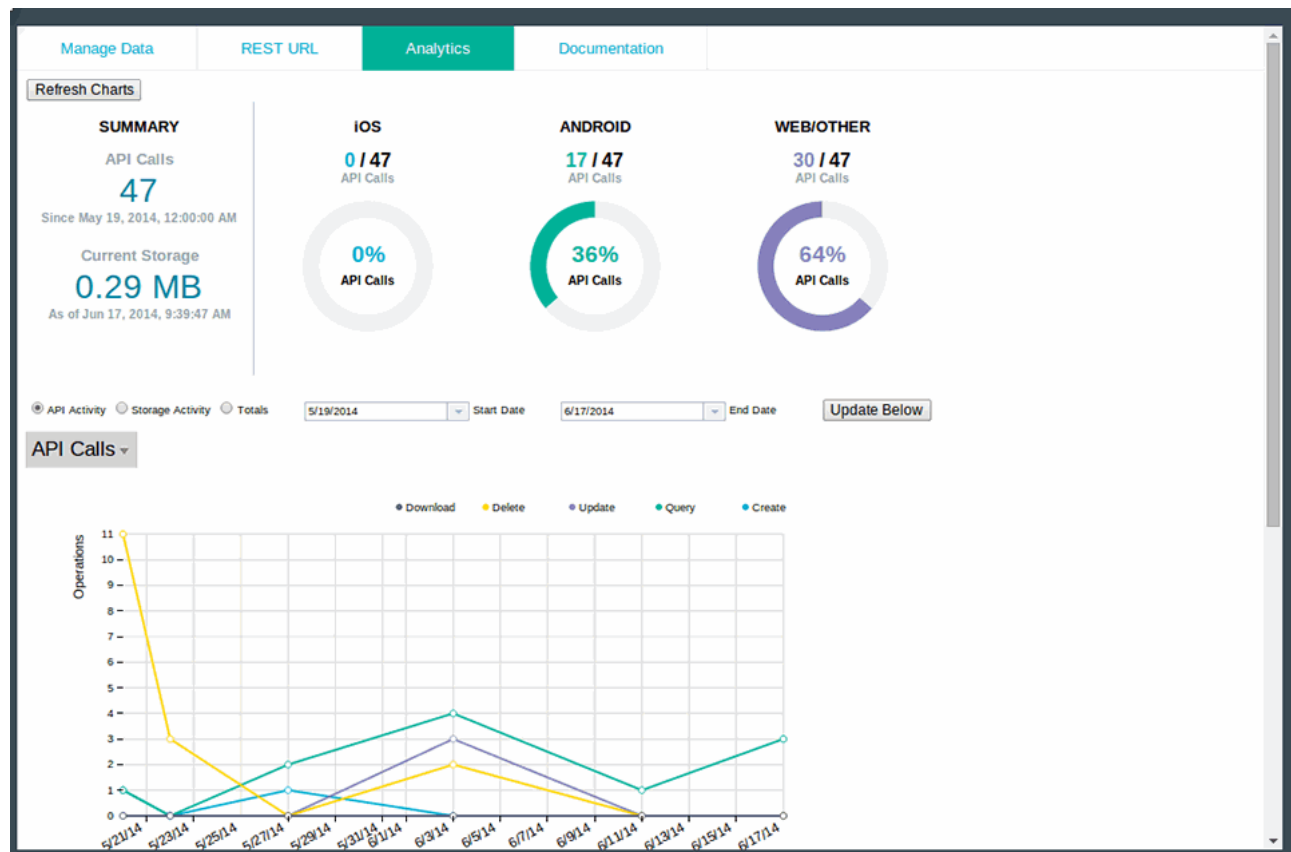


Figure 2-5 Manage Data Analytics

2.4.5 Mobile Quality Assurance (IBM)

IBM Mobile Quality Assurance enables mobile app testing, user validation, and streamlined quality feedback, plus these capabilities:

- ▶ Over-the-air build distribution
Get the latest build in the hands of the testers as soon as it is available.
- ▶ Automated crash reporting
Get aggregated crash logs from pre-production and production environments automatically.
- ▶ In-app bug reporting
Submit defect reports in seconds while using the application from your mobile device.
- ▶ User feedback and sentiment analysis
Get feedback about your app straight from your customers, and mine app ratings and reviews to extract actionable insights.

2.4.6 Presence Insights (IBM)

IBM Presence Insights helps you understand mobile activity in and around physical locations, such as malls, stadiums, airports, and stores. It extends customer service and inventory promotions through your customers' mobile devices, which offer you significant advantages:

- ▶ Deepen customer engagement
Presence Insights simplifies the process of gathering and analyzing real-time analytics. Use this data to better understand your customers' patterns and trends. These findings provide numerous opportunities to implement intelligent engagement strategies that enable a richer customer experience.
- ▶ Optimize onsite staff operations
By gaining a clear picture of activity within and around a location, such a specific department in a retail establishment, you can more strategically use your staff members and customer support opportunities to increase revenue.

Figure 2-6 illustrates setting up a new site to be monitored by Presence Insights.



Figure 2-6 Setting up a new site to monitor

Figure 2-7 shows where visitors are spending their time throughout the day.

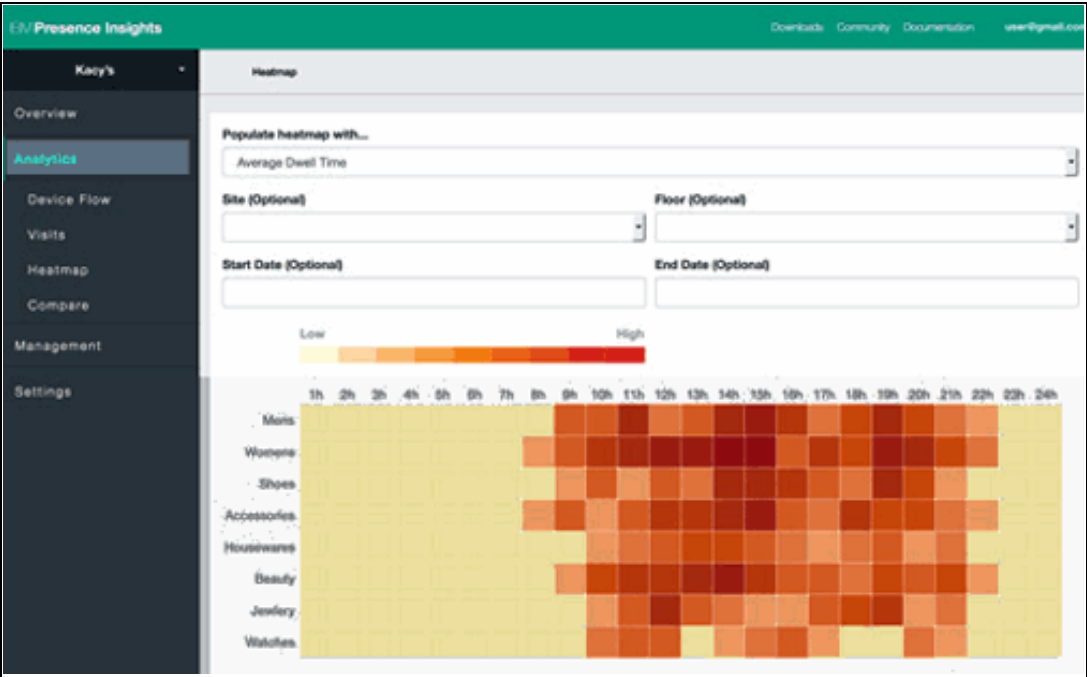


Figure 2-7 Discover where visitors are spending their time, by hours

Figure 2-8 shows how many visitors are visiting each area.

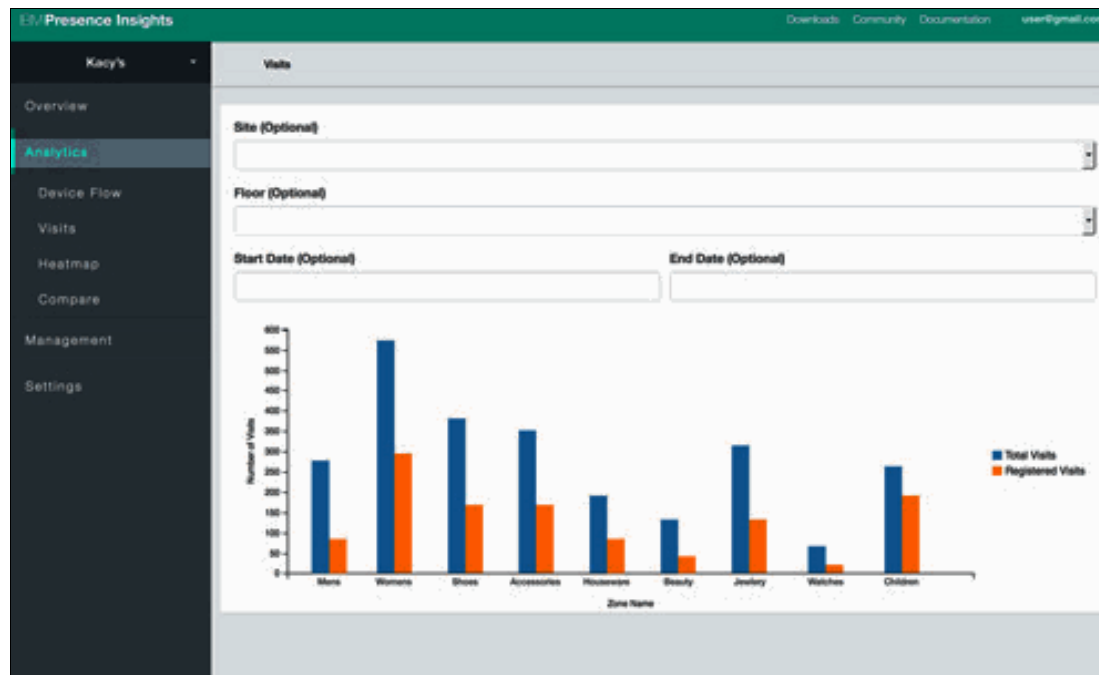


Figure 2-8 Discover how many visitors are visiting each area

2.4.7 Push (IBM)

Send relevant content to the right people at the right place and time with IBM Push for Bluemix. This service includes these main capabilities:

- ▶ Engage with users on their mobile devices
Push information to all application users or to a specific set of users and devices. You can even let users subscribe to specific tags or topics for notification. After users are engaged, you can analyze the number of devices that are registered to receive notifications, the number of notifications sent, and the platforms of devices that are receiving notifications.
- ▶ Use a single platform for notifying heterogeneous devices
Use simple and uniform REST APIs to configure, subscribe, send, and monitor push notifications to both iOS and Android devices.

2.4.8 Push iOS 8 (IBM)

The Push iOS 8 service helps to send push notifications to iOS devices. The devices can be targeted by using tags or device identifiers. Push iOS 8 includes these main capabilities:

- ▶ Notify iOS 8 devices
Use simple and uniform REST APIs to configure, subscribe, send, and monitor push notifications to iOS 8 devices.

- ▶ Send notifications by various criteria and analyze effects

Push information to all application users or to a specific set of users and devices. You can send interactive notifications and silent notifications. You can even let users subscribe to specific tags or topics for notification. After users are engaged, you can analyze the number of devices that are registered to receive notifications and the number of notifications sent.

2.4.9 Twilio (third party)

The Twilio service is used to build apps that communicate. It can integrate voice, messaging, and VoIP into your web and mobile apps.

2.5 DevOps

The DevOps category provides the following services to support automatic deployment:

- ▶ Auto-Scaling (IBM)
- ▶ Delivery Pipeline (IBM)
- ▶ Monitoring and Analytics (IBM)
- ▶ Track & Plan (IBM)
- ▶ BlazeMeter (third party)
- ▶ Load Impact (third party)
- ▶ New Relic (third party)

2.5.1 Auto-Scaling (IBM)

With the IBM Auto-Scaling for Bluemix service, you can automatically increase or decrease the compute capacity of your application. The number of application instances is adjusted dynamically, based on the Auto-Scaling policy you define. These are the main capabilities:

- ▶ Dynamic scaling
Automatically add or remove resources to match the current workload.
- ▶ Custom scaling policy
Define policy for metrics of interest.
- ▶ Metric statistics
Visualize the current and historical values of performance metrics.
- ▶ Scaling history
Query the scaling activities based on status, time and type.

2.5.2 Delivery Pipeline (IBM)

Use IBM DevOps Services to automate builds and deployments, test execution, configure build scripts, and automate execution of unit tests. Automatically build and deploy your application to Bluemix.

2.5.3 Monitoring and Analytics (IBM)

This service helps you gain the visibility and control you need over your applications. You can determine the response time that your users see, understand the performance and availability of the application components, and use analytics to keep your application operating and performing well. These are the main advantages:

- ▶ Effortless visibility

Get the visibility you need without spending any time or effort to learn or deploy yet another new tool.

- ▶ Innovate faster

Spend your valuable time innovating and delivering value to your users, not chasing bugs and performance issues.

- ▶ Diagnose problems 90% faster

Quickly identify the root cause of application problems with line-of-code diagnostics.

- ▶ Resolve problems with embedded analytics

Intelligently search through logs and metric data to find the answers to your application problems.

- ▶ Keep your app running

When your app is up, IBM can help you keep it up. Reduce your maintenance costs while improving availability.

Figure 2-9 shows the Application Availability view.



Figure 2-9 Application Availability view

Figure 2-10 shows the Log Analysis view.



Figure 2-10 Log Analysis view

Figure 2-11 shows the Performance Monitoring view.

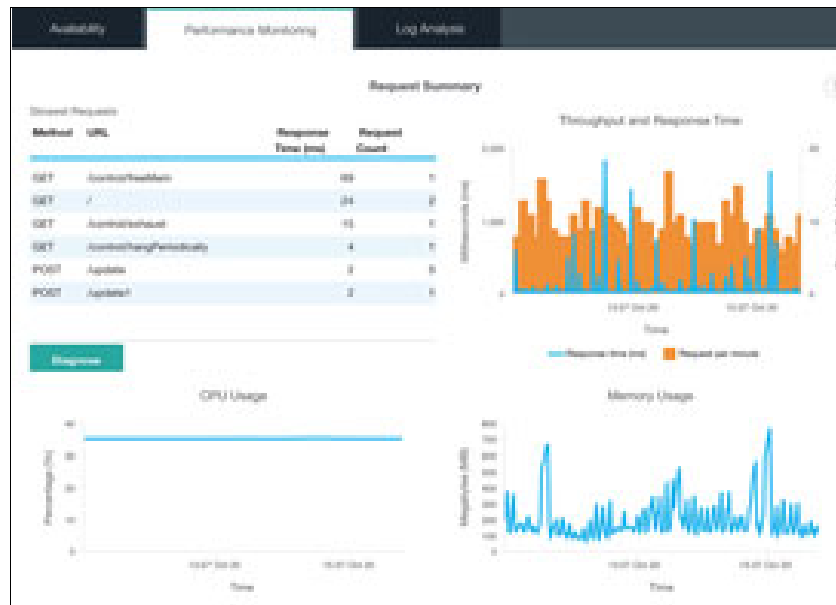


Figure 2-11 Performance Monitoring

Figure 2-12 shows the Stack Trace view.

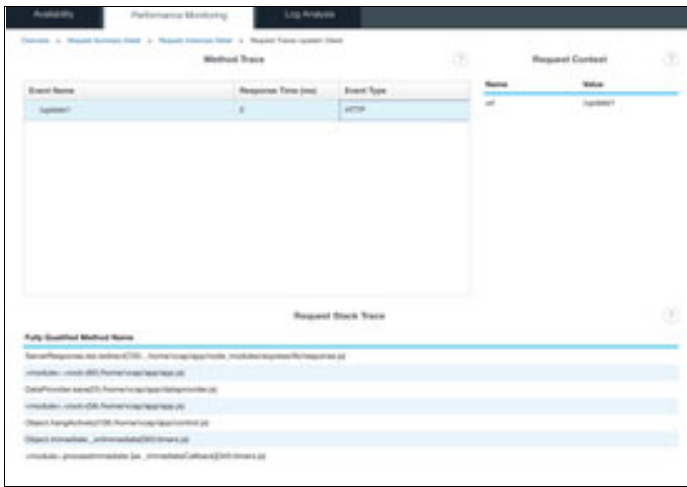


Figure 2-12 Stack Trace

2.5.4 Track & Plan (IBM)

Use IBM DevOps Services to create stories, tasks, and defects to describe and track project work, and use agile planning tools for the product backlog, releases, and sprints.

2.5.5 BlazeMeter (third party)

BlazeMeter is a self-service, web, API mobile load-testing platform (PaaS) that offers developers an enterprise grade, out-of-the-box load testing solution that is 100% compatible with Apache JMeter. Start testing in under five minutes, and scale tests to more than one million concurrent users.

2.5.6 Load Impact (third party)

Using this service, you can do unlimited testing on demand from multiple geographic locations. Create sophisticated tests by using the simple GUI or connect directly to the IBM platform via our API.

2.5.7 New Relic (third party)

New Relic is the all-in-one web app performance tool that lets you see performance from the end user experience, through servers, and down to the line of code.

2.6 Web and Application services

The web and application services help you deliver new web applications and mobile apps. This category includes the following services:

- ▶ Business Rules (IBM)
- ▶ Data Cache (IBM)
- ▶ MQ Light (IBM)
- ▶ Session Cache (IBM)

- ▶ Workflow (IBM)
- ▶ Workload Scheduler (IBM)
- ▶ Box (third party)
- ▶ CloudAMQP (third party)
- ▶ DreamFace (third party)
- ▶ Geocoding (third party)
- ▶ Memcached Cloud (third party)
- ▶ Namara.io Catalog (third party)
- ▶ Reappt (third party)
- ▶ Reverse Geocoding (third party)
- ▶ SendGrid (third party)
- ▶ Simplicite aPaaS for Bluemix (third party)
- ▶ Statica (third party)
- ▶ Travel Boundary Service (third party)
- ▶ Ustream (third party)
- ▶ Validate Address (third party)

2.6.1 Business Rules (IBM)

Enables developers to spend less time recoding and testing when the business policy changes. The Business Rules service minimizes your code changes by keeping business logic separate from application logic.

2.6.2 Data Cache (IBM)

Improve the performance and user experience of web applications by retrieving information from fast, managed, in-memory caches, rather than relying entirely on slower disk-based databases. These are the main capabilities:

- ▶ Distributed cache
- ▶ Improves the response time of a cloud-centric application's presentation layer by caching response to frequent user requests and helps build responsive applications
- ▶ Create, read, update, and delete operations

Remotely uses the caching capabilities of the data grid and lets you perform create, retrieve, update and delete operations
- ▶ NoSQL style in-memory grid

Provides a NoSQL style distributed key and value storage service in memory
- ▶ Linear and scalability

Provides linear scalability, predictable performance and fault tolerance of the web application's data requirements

2.6.3 MQ Light (IBM)

IBM MQ Light for Bluemix is a cloud-based messaging service that is built on the Advanced Message Queuing Protocol (AMQP). It provides flexible and easy-to-use messaging for Bluemix apps.

Develop responsive, scalable applications with a fully managed messaging provider in the cloud. Quickly integrate with application frameworks through easy-to-use APIs. These are the main capabilities:

- ▶ Easy to use

Connect applications simply and efficiently so they can offload work, share data, or push events with a simple API for Java and JavaScript and zero administration.

- ▶ Robust and scalable

Rely on MQ Light's data integrity and asynchronous delivery to ensure that your distributed applications are loosely coupled, robust, and scalable.

2.6.4 Session Cache (IBM)

Improve application resiliency by storing session state information across many HTTP requests. Enable persistent HTTP sessions for your application and seamless session recovery in event of an application failure. These are the main capabilities:

- ▶ Distributed session cache

Special purpose elastic in memory cache for storing HTTP session data and provides low latency data access and transactional semantics.

- ▶ Ease of use

No code change to applications using the J2EE standard HTTP session cache.

- ▶ Session failover

Data for each HTTP session survives on a server outage for an application.

- ▶ Session replication

Replicates the session data to avoid a single point of failure.

2.6.5 Workflow (IBM)

IBM Workflow for Bluemix makes it easy to create workflows that orchestrate and coordinate the REST-based services that you use in your apps. The JavaScript based Workflow language lets you define interactions between any services. By offloading all of the service interactions to the Workflow service, your application becomes easier to understand, maintain, and evolve. Your workflows are run and managed in a robust and scalable way, regardless of whether your workflow and services run for milliseconds or days.

Figure 2-13 shows authoring a workflow in IBM DevOps, using a flexible scripting language.

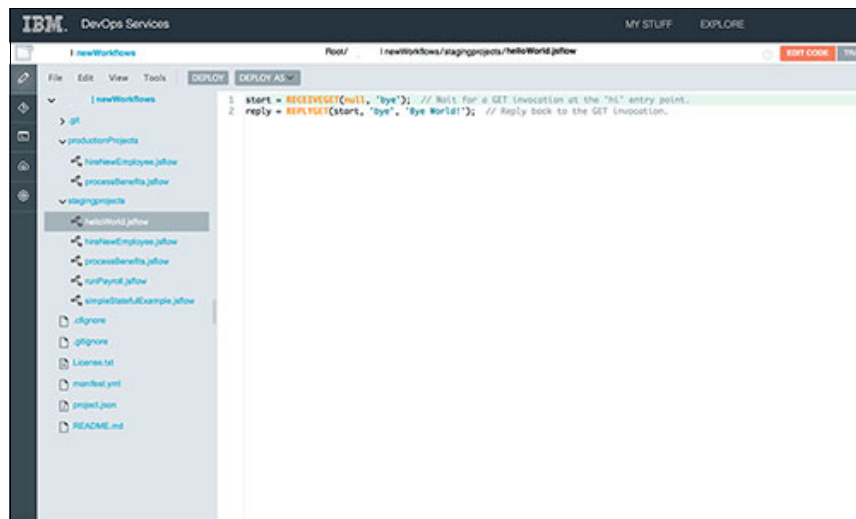


Figure 2-13 Author workflows in IBM DevOps, using a flexible scripting language

Figure 2-14 shows administering multiple workflows in the IBM Workflow Administration Console.

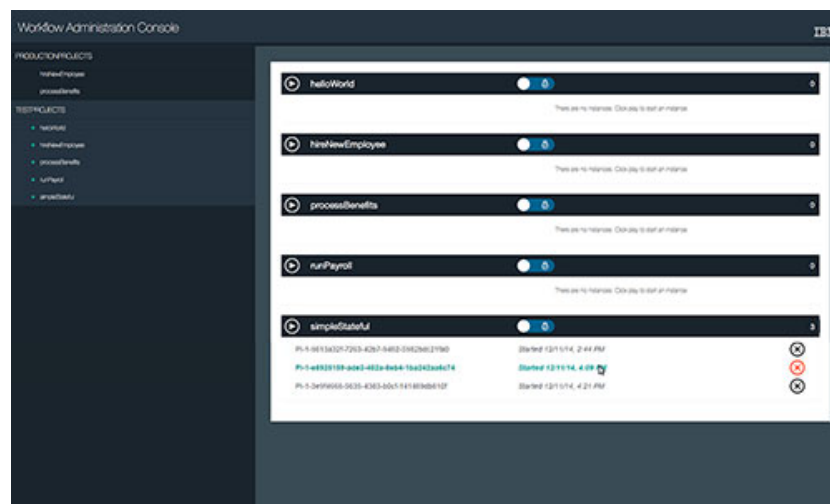


Figure 2-14 Administer multiple workflows in the IBM Workflow Administration Console

Figure 2-15 shows viewing specific workflow instance diagrams to find failed tasks.

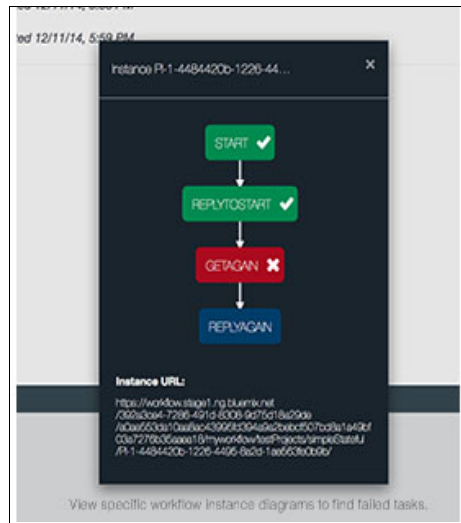


Figure 2-15 View specific workflow instance diagrams to find failed tasks

2.6.6 Workload Scheduler (IBM)

Exploit the IBM APIs to trigger the execution of simple or complex sequences of actions (such as REST call and database query). Leave repetitive activities to be automated by IBM and keep control of the execution through APIs or summarized views.

2.6.7 Box (third party)

Whether you are building a line of business app, content management software, or need to display content beautifully on web and mobile, the Box API can help power the content and data layer of your application.

2.6.8 CloudAMQP (third party)

Decouple, distribute, and scale your applications with the speed of CloudAMQP.

2.6.9 DreamFace (third party)

DreamFace is a cloud-based development platform to deliver enterprise-scale web and mobile applications with ease and speed. This service produces uncompromising user experiences and delivers the right data to the right user on the right device by secure and code-less use of any public or private API.

2.6.10 Geocoding (third party)

Bluemix address geocoding solutions help you to find exact geographical locations by inputting an address to get the location as a response.

2.6.11 Memcached Cloud (third party)

Memcached Cloud is a fully managed cloud service for hosting and running your Memcached dataset in Bluemix in a highly available and scalable manner, with predictable and stable top performance. It offers true high availability with data persistence, in-memory replication, and instant automatic failover.

2.6.12 Namara.io Catalog (third party)

The Namara.io platform aggregates available open data released by all levels of government and presents it to users in a single portal. It organizes and catalogs this public data, providing users with API access to high-value information.

2.6.13 Reappt (third party)

Reappt is the latest cloud-based data distribution solution by Push Technology. Reappt gives developers the tools to quickly and easily build reactive apps that can respond to event-driven activity at scale.

2.6.14 Reverse Geocoding (third party)

Reverse Geocoding enables you to look up an address for a GPS location. Precision Reverse Geocoding (by Pitney Bowes Software) empowers marketers to connect with consumers via location-based mobile marketing.

2.6.15 SendGrid (third party)

The SendGrid cloud-based email infrastructure relieves businesses of the cost and complexity of maintaining email systems.

2.6.16 Simplicity aPaaS for Bluemix (third party)

The Simplicity aPaaS for Bluemix is a dynamically model-driven platform for building agile custom enterprise applications in the cloud. *aPaaS* means *application platform as a service*.

2.6.17 Statica (third party)

Access resources behind firewalls and IP white-listed services by routing your traffic via a static IP. Your requests always originate from the same IP address, regardless of where your app is running or how many instances you have operating.

2.6.18 Travel Boundary Service (third party)

Generate isochrone bands around a location based on driving time or distance. This is essential for site selection analysis, logistics planning, and supply chain management, among other purposes.

2.6.19 Ustream (third party)

To stream live video, Ustream offers cloud storage, multi-screen compatibility, and customized video embedding.

2.6.20 Validate Address (third party)

Validate Address enables you to standardize and validate addresses, ensuring that your address data adheres to quality standards established by the postal authority. This increases the likelihood of your mail being delivered.

2.7 Integration

The Integration category provides the following services to extend existing investments and infrastructure:

- ▶ API Management (IBM)
- ▶ Cloud Integration (IBM)
- ▶ Secure Gateway (IBM)

2.7.1 API Management (IBM)

The IBM API Management for Bluemix service enables developers and organizations to manage and enforce policies related to the use of their business services. You can use an existing API or design a new API. Next, apply security controls, set rate limits, test APIs in place, and then publish these “managed APIs” on Bluemix. You can publish these managed APIs either to your organization or to selected developer organizations or to app developers outside of Bluemix. Share your APIs by using an available self-service portal that can be white-labeled and that provides built-in support for blogs, discussion forums, comments, ratings, FAQs, and the APIs that you choose to publish. This service includes API versioning, lifecycle management, and API use analytics.

The service provides these main capabilities:

- ▶ Discover existing APIs

Easily find and display REST and SOAP services by discovering them from custom registries, such as those on IBM z/OS® systems. Then, manage, track, and limit the discovered services.

- ▶ Design new APIs

Create a composite of multiple existing APIs to get a new, more powerful API. Import APIs from Swagger or Web Service Definition Language (WSDL) definitions, or build one from scratch.

2.7.2 Cloud Integration (IBM)

IBM Cloud Integration for Bluemix enables users to integrate cloud services with enterprise systems of record. Bluemix Cloud Integration displays the backend systems of record as ReST APIs to be used by applications. Cloud Integration provides these main capabilities:

- Fast and simple

This service enables users to quickly integrate cloud services with enterprise systems of record

- Connectivity

This service displays the backend systems of record as ReST APIs to be used by applications.

2.7.3 Secure Gateway (IBM)

The Secure Gateway Service brings hybrid Integration capability to your Bluemix environment. It provides secure connectivity from Bluemix to other applications and data sources running on-premises or in other clouds. A remote client is provided to enable secure connectivity. This service provides these main capabilities:

- Fast and simple

Set up gateways to other environments and monitor your traffic.

- Secure gateway

Create and manage local endpoint mappings to remote destinations.

Figure 2-16 shows the original splash screen when first launching the service.

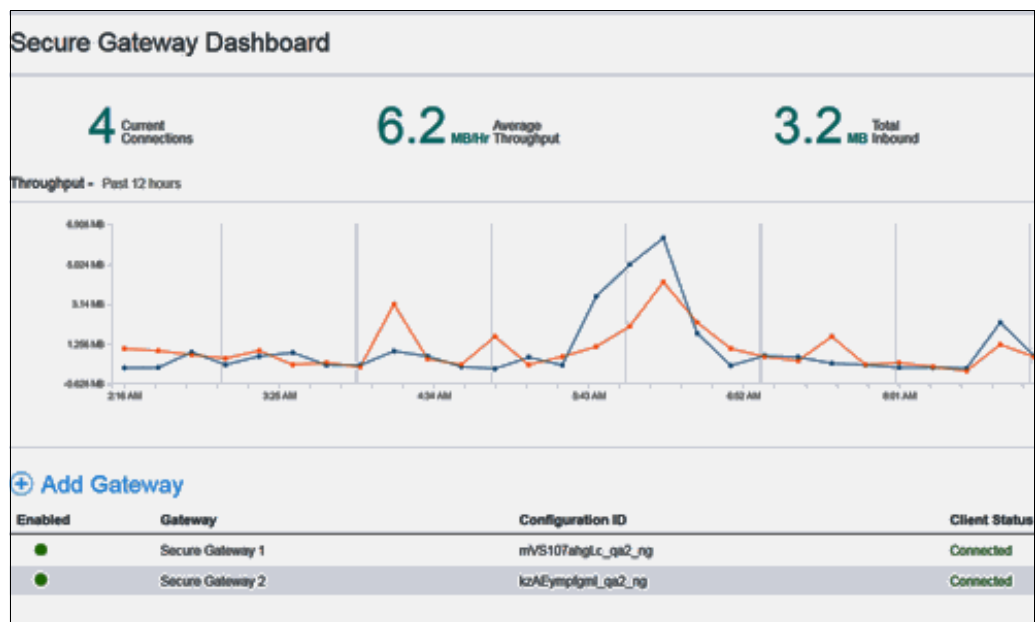


Figure 2-16 The original splash screen when first launching the service

Figure 2-17 shows how you can manage your secure gateways.

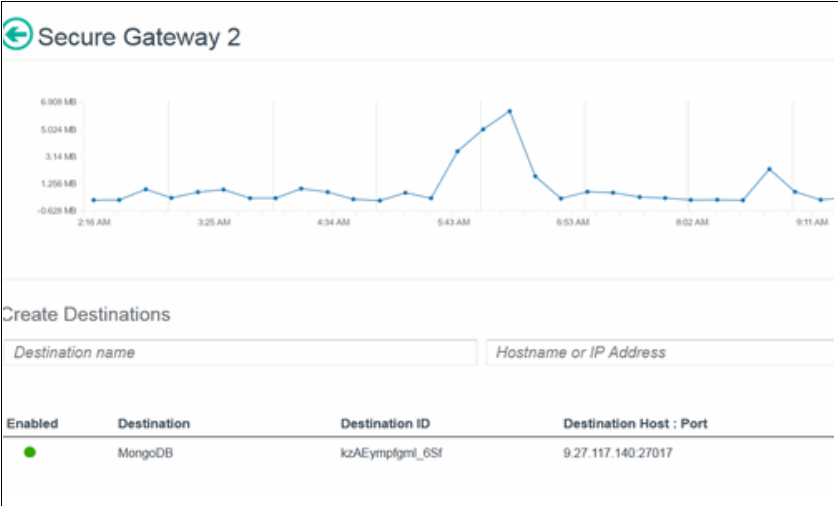


Figure 2-17 Manage your Secure Gateways

Figure 2-18 show how you can manage the destinations of a Secure Gateway.

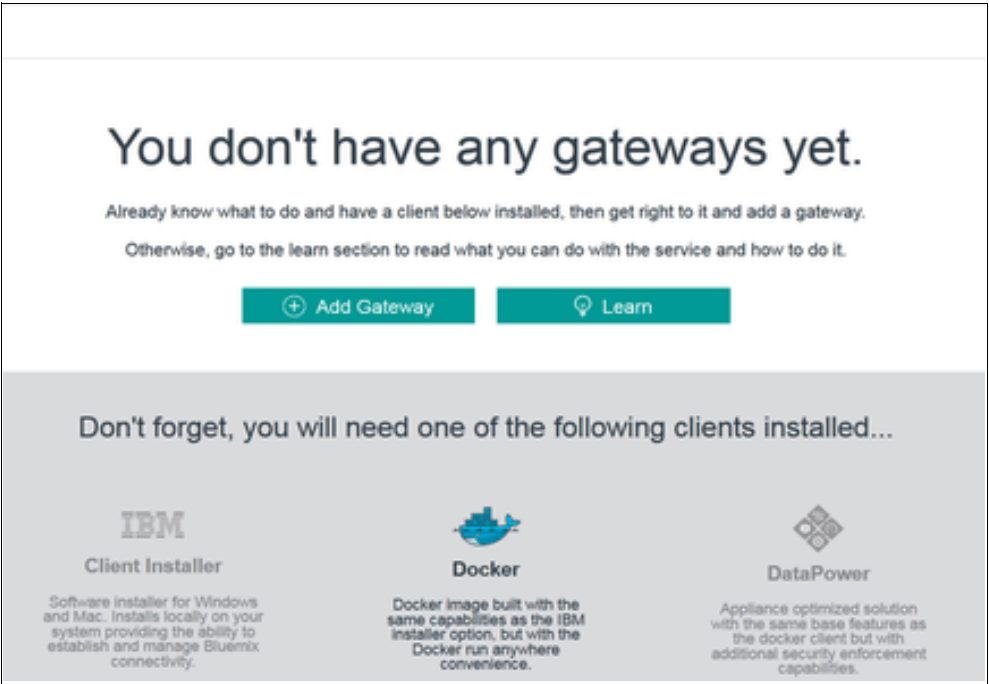


Figure 2-18 Manage the destinations of a Secure Gateway

2.8 Data Management

Data Management category provides services for agile data management and refinement. This service provides these main capabilities:

- ▶ Cloudant NoSQL DB (IBM)
- ▶ DataWorks (IBM)
- ▶ Object Storage (IBM)
- ▶ Object Storage (v2) (IBM)
- ▶ SQL Database (IBM)
- ▶ ClearDB MySQL Database (third party)
- ▶ ElephantSQL (third party)
- ▶ MongoLab (third party)
- ▶ Redis Cloud (third party)

2.8.1 Cloudant NoSQL DB (IBM)

Cloudant NoSQL DB provides access to a fully managed NoSQL JSON data layer that is always on. This service is compatible with CouchDB and accessible through a simple-to-use HTTP interface for mobile and web application models. This service provides these main capabilities:

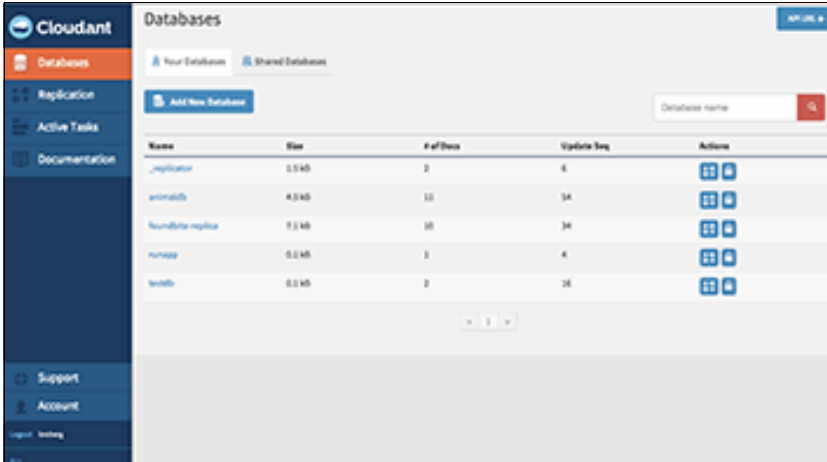
- ▶ Ease of use

Work with self-describing JSON documents through a RESTful API that makes every document in your Cloudant database accessible as JSON via a URL. Documents can be retrieved, stored, or deleted individually or in bulk and can also have files attached. IBM takes care of the provisioning, management, and scalability of the data store, freeing your time to focus on your application.

- ▶ Powerful search, sync and more

With extremely powerful indexing, real-time MapReduce and Apache Lucene-based full-text search, Cloudant NoSQL DB makes it easy to add advanced data analytics and powerful data access. Data access can also extend to Cloudant Sync, which enables data access from mobile devices and client apps to run connected or offline.

Figure 2-19 shows the Cloudant NoSQL DB database list.



The screenshot shows the Cloudant web interface. On the left is a navigation sidebar with links for Databases, Replication, Active Tasks, Documentation, Support, and Account. The main area is titled 'Databases' and has tabs for 'Your Databases' and 'Shared Databases'. Below the tabs is a search bar and a table listing databases. The table has columns for Name, Size, # of Docs, Update Seq, and Actions. Five databases are listed: _replicator, animals, foundata-replica, humapp, and testdb. Each database has a size, a number of documents, an update sequence, and icons for actions like refresh and delete.

Name	Size	# of Docs	Update Seq	Actions
_replicator	0.0 KB	0	0	[Refresh] [Delete]
animals	4.0 KB	10	34	[Refresh] [Delete]
foundata-replica	0.0 KB	0	34	[Refresh] [Delete]
humapp	0.0 KB	0	4	[Refresh] [Delete]
testdb	0.0 KB	0	34	[Refresh] [Delete]

Figure 2-19 Database list

Figure 2-20 shows Database Permissions screen.

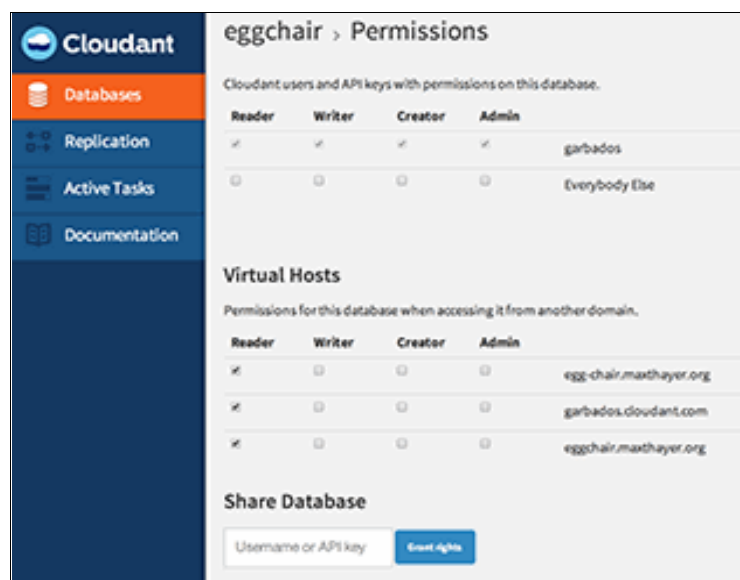


Figure 2-20 Database Permissions

Figure 2-20 shows the screen to create a search index.

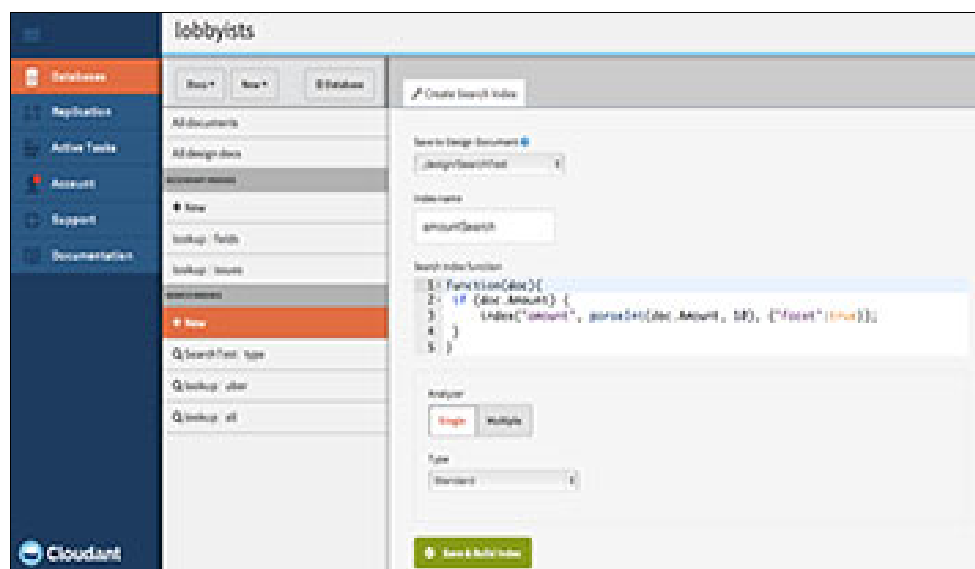


Figure 2-21 Create search index

2.8.2 DataWorks (IBM)

The IBM DataWorks data refinery transforms raw data into relevant information. It includes IBM DataWorks Forge, an app primarily for knowledge workers, as well as APIs for application developers. IBM DataWorks uses a highly performant and scalable engine to discover, profile, enrich, mask, and deliver data to applications.

DataWorks provides these main capabilities:

- **Forge (beta)**

A data-rich app that empowers knowledge workers (including business analysts, data scientists, and nontechnical users) to find data, visually represent it, and prepare it for use. By automatically profiling, classifying, and scoring data, Forge guides you through the process of enriching and improving the quality of data using actions, such as removing duplicates, filtering, and joining. After you prepare and enrich your data, Forge makes it easy for you to deliver data to applications and systems.

- **APIs**

Flexible, REST-based APIs enable developers to quickly access data and ensure it is fit for purpose. Using the IBM DataWorks APIs, you can quickly create higher quality applications that load data between data sources (such as SQL Database, Object Storage, dashDB, IBM Analytics for Hadoop, DB2, and Oracle), mask data while loading, securely load on-premises data to cloud environments, cleanse US postal addresses, and classify and profile data.

2.8.3 Object Storage (IBM)

The Object Storage service in Bluemix is based on SoftLayer Object Storage, which is based on OpenStack Swift. It has built-in support for provisioning independent object stores, and it creates an individual subaccount per object store.

2.8.4 Object Storage (v2) (IBM)

Object Storage provides access to a fully provisioned Swift Object storage account. Object Storage is ideal for cost effective, scale-out storage. Swift provides a fully distributed, API-accessible platform that can be integrated into your applications or used for back up. Topologies can be public or private:

- **Public**

Public object storage is available to all Bluemix organizations. A dedicated Swift account is provided to each organization. The account can be used with Bluemix applications, virtual machines, containers, or your on-premises locations.

- **Private**

Private object storage is available only to your Bluemix organization. Using private object storage helps protect your existing infrastructure investments. Private object storage can be used with your Bluemix applications, virtual machines, and containers.

Figure 2-22 shows the screen to choose a cloud.

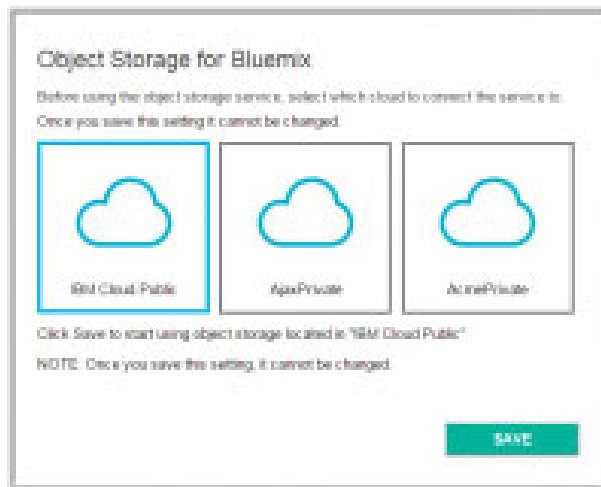


Figure 2-22 Choose a cloud

Figure 2-23 shows viewing account use.

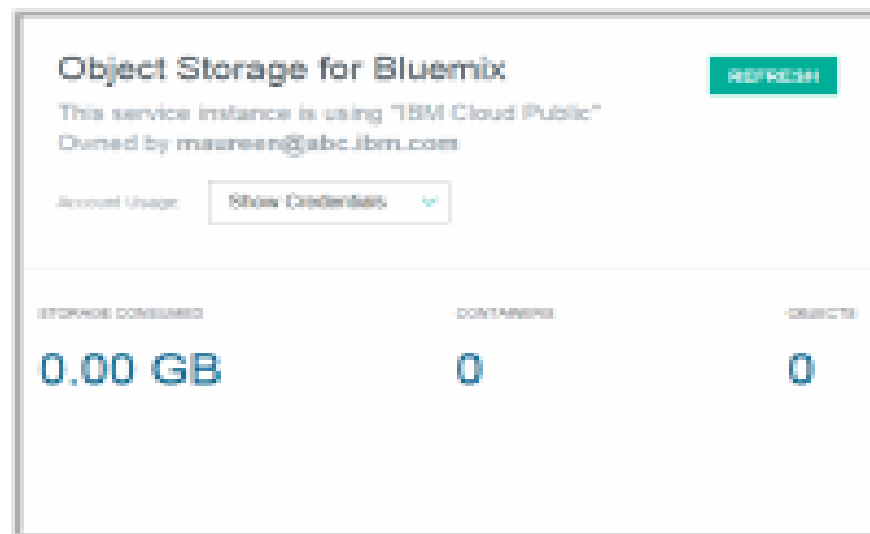


Figure 2-23 View account use

Figure 2-24 shows viewing account credentials.

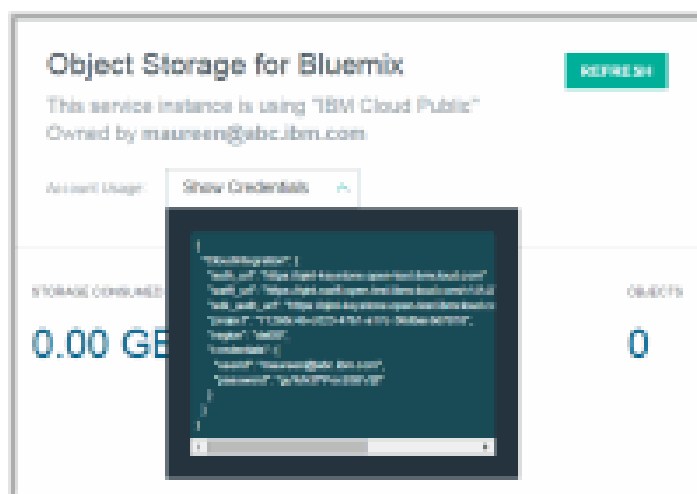


Figure 2-24 View account credentials

2.8.5 SQL Database (IBM)

SQL Database adds an on demand relational database to your application. Powered by DB2, it provides a managed database service to handle web and transactional workloads. SQL This service provides these main capabilities:

- ▶ Fully managed
Allows you to focus on your application rather than worrying about time consuming administration of your database server.
- ▶ Secure and private
SSL encryption will be standard on all plans. The Small Plan also includes sensitive data reporting, which is powered by IBM InfoSphere Guardium Data Activity Monitor.

2.8.6 ClearDB MySQL Database (third party)

ClearDB is a reliable, fault tolerant, geo-distributed database-as-a-service offering for your MySQL powered applications.

2.8.7 ElephantSQL (third party)

The ElephantSQL service will get a deployed open-source database with automatically configured database parameters hosted in the cloud.

2.8.8 MongoLab (third party)

MongoLab is a fully managed cloud database service featuring highly available MongoDB databases, automated backups, web-based tools, 24x7 monitoring, and expert support.

2.8.9 Redis Cloud (third party)

Redis is a fully managed cloud service for hosting and running your Redis dataset in Bluemix in a highly available and scalable manner, with predictable and stable performance. It offers true high availability with its in-memory dataset replication and instant auto-failover mechanism.

2.9 Big Data

The Big Data category provides services for big data management and refinement. It encompasses the following services, which the subsections that follow describe:

- ▶ BigInsights for Apache Hadoop (IBM)
- ▶ dashDB (IBM)
- ▶ Geospatial Analytics (IBM)
- ▶ IBM Analytics for Apache Hadoop (IBM)
- ▶ Insights for Twitter (IBM)
- ▶ Time Series Database (IBM)

2.9.1 BigInsights for Apache Hadoop (IBM)

Develop analytics applications by using open source Apache Hadoop and IBM BigInsights APIs without having to manage the platform. The service is managed and scaled for you by the big data experts at IBM. This service provides these advantages:

- ▶ Power of open source innovation
This service provides 100% open source Apache Hadoop through the IBM Open Platform. It is updated regularly to provide the best versions of the Apache Hadoop ecosystem components, including, among others, Ambari, YARN, Spark, Knox, HBase, Hive, and an encrypted Hadoop Distributed File System (HDFS).
- ▶ Focus on delivering insights
Key components of the platform, including the infrastructure, are proactively monitored by an IBM cloud operations team 24x7. Therefore, critical security fixes, high severity patches, and defect resolutions are promptly applied on the cluster. You can take advantage of high-value Hadoop analytics features, such as Big SQL, BigSheets, Text Analytics, Big R, and Machine Learning, to gain insight faster.

Figure 2-25 shows defining a BigInsights for Hadoop cluster.

Define Cluster MANAGE CLUSTERS

* Name

Description

Estimate Your Cluster Requirements

Hardware sizing per node [Compare hardware](#)

☒ Small
Recommended for basic data extraction, transformation, file processing, search

☐ Medium
Recommended for data warehouse augmentation

☐ Large
Recommended for advanced analytics

Choose Software Package to be installed [Compare Packages](#)

☐ IOP
Open Data Platform distribution

☒ BigInsights Data Analyst
Distribution recommended for data analyst

☐ BigInsights Data Scientist
Distribution recommended for data scientists

Number of compute nodes -OR- Data storage on compute nodes (TB)

Data center

Figure 2-25 Define BigInsights for Hadoop Cluster

Figure 2-26 shows the BigInsights for Hadoop Cluster Topology screen.



Figure 2-26 BigInsights for Hadoop Cluster Topology

Figure 2-27 shows the screen to manage BigInsights for Hadoop Clusters.

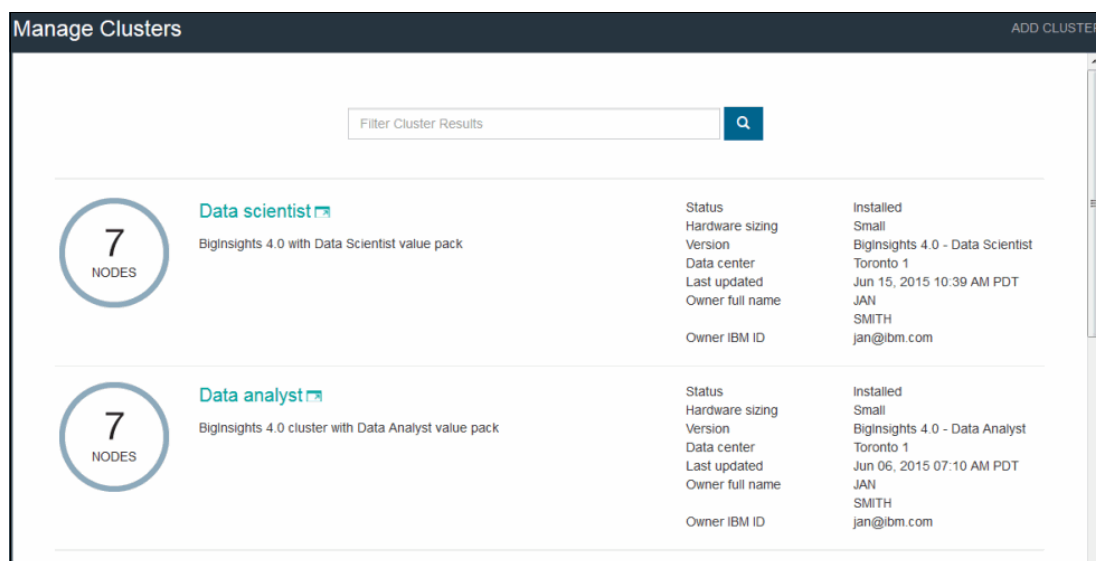


Figure 2-27 Manage BigInsights for Hadoop Clusters

Figure 2-28 shows the BigInsights for Hadoop Cluster details.

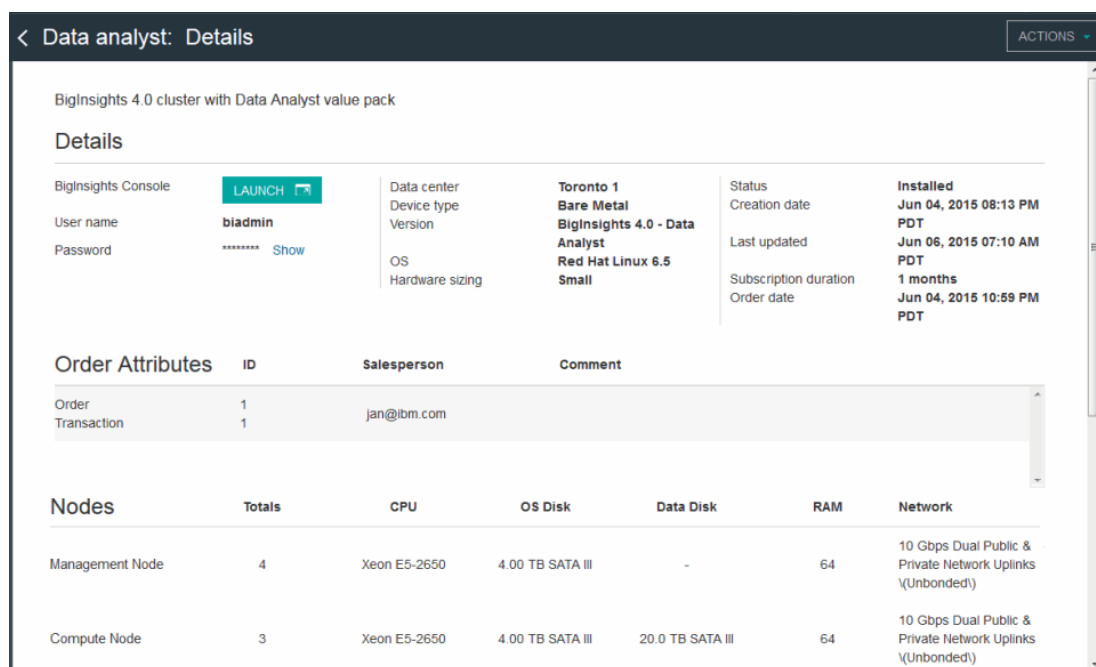


Figure 2-28 BigInsights for Hadoop cluster details

Figure 2-29 shows how you can visualize your data in BigInsights for Hadoop.

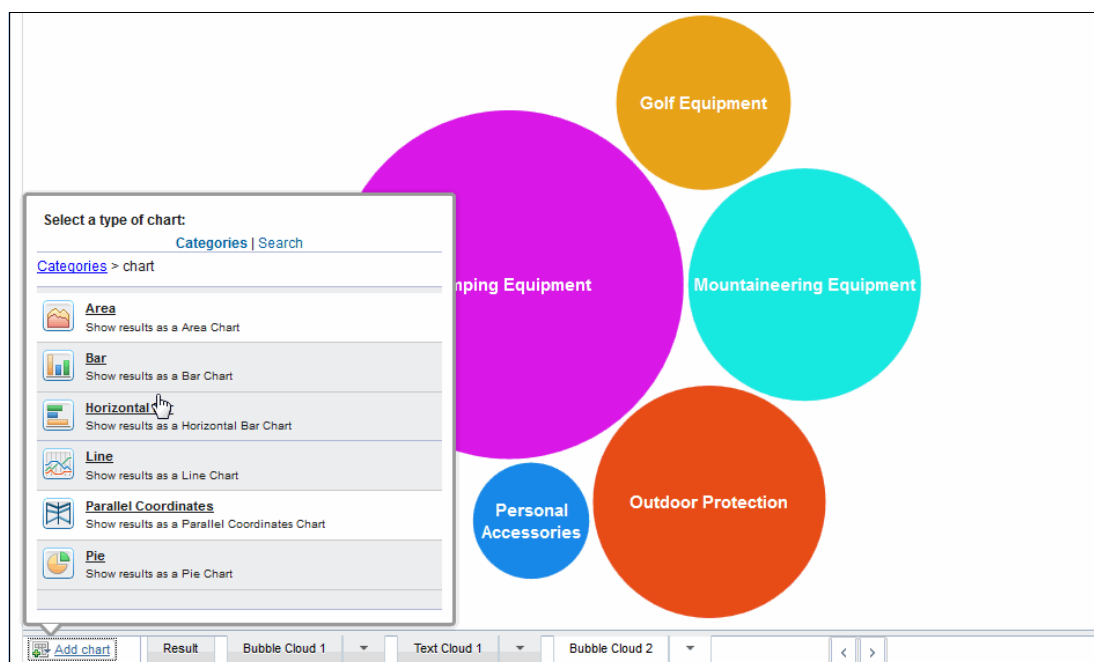


Figure 2-29 Visualize your data in BigInsights for Hadoop

2.9.2 dashDB (IBM)

dashDB is a data warehousing and analytics solution. Use dashDB to store relational data, including special types, such as geospatial data. Then, analyze that data with SQL or advanced built-in analytics, such as predictive analytics and data mining, analytics with R, and geospatial analytics. You can use the in-memory database technology to use both columnar and row-based tables. This service provides these main capabilities:

- **Powered by IBM BLU Acceleration and IBM Netezza® In-Database Analytics**
BLU Acceleration is fast and simple. It uses dynamic in-memory columnar technology and innovations, such as actionable compression, to rapidly scan and return relevant data. In-database analytic algorithms integrated from IBM Netezza appliance technology brings simplicity and performance to advanced analytics.
- **Connectivity**
dashDB is built to connect easily to all of your services and applications. You can start analyzing your data right away with familiar tools.

Figure 2-30 on page 57 shows the Welcome screen for first time users.

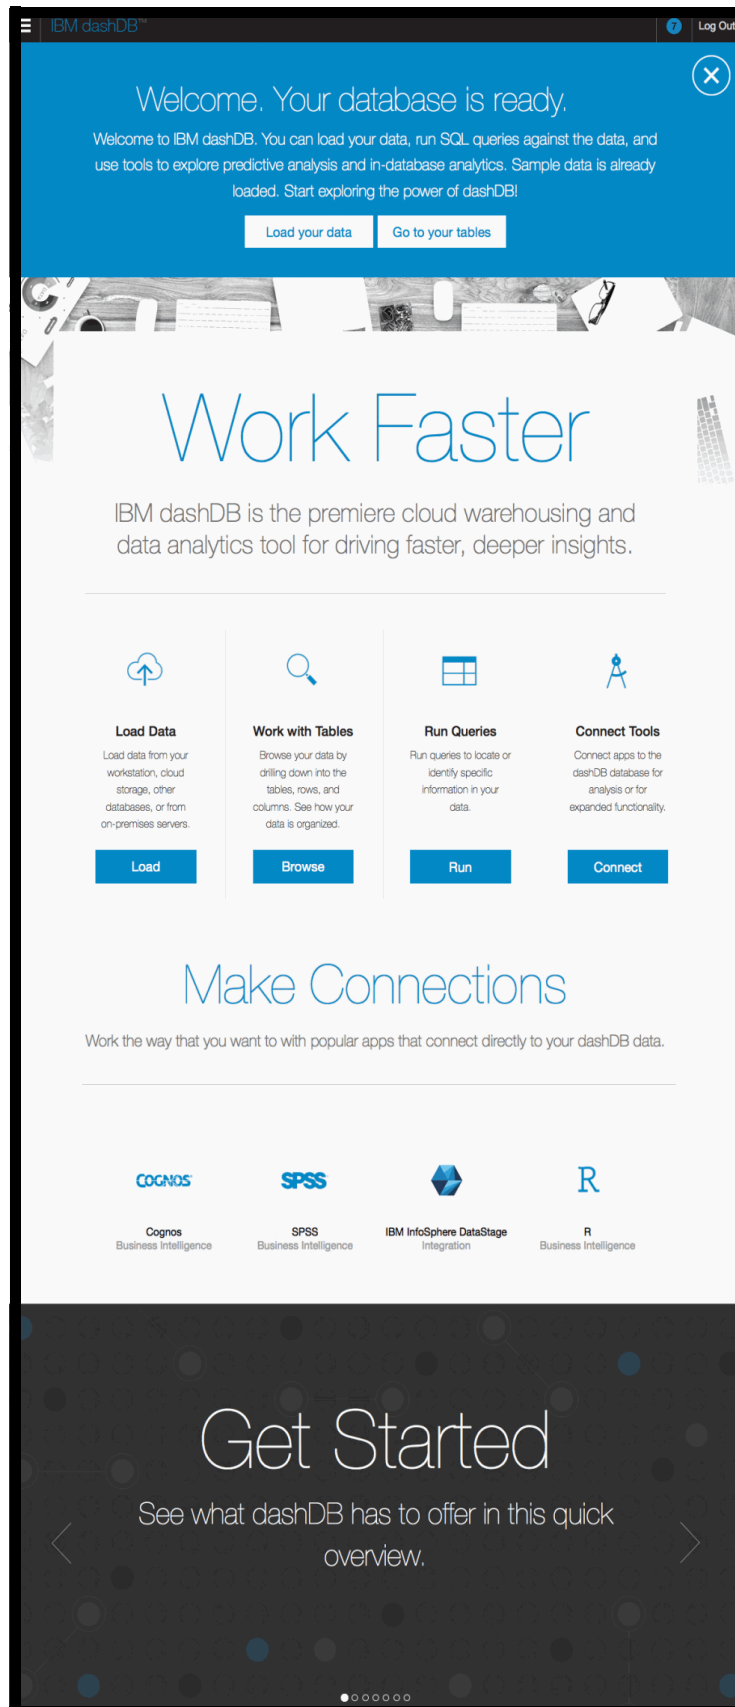


Figure 2-30 Welcome screen for first time users

Figure 2-31 shows how to rapidly you can get your local data loaded with Quick Load.

The screenshot displays the 'Load data into a table' interface in IBM dashDB. On the left, a sidebar contains a 'Quick load' tab (highlighted in blue) and three other options: 'Schedule a load', 'View or modify scheduled loads', and 'View the history and status of scheduled loads'. The main content area is titled 'Do a one-time load from an Excel file or from a delimited text file such as a comma-separated value (CSV) file'. It features a progress bar with four steps: '1. Upload a file', '2. Choose the target', '3. Select a table', and '4. Load complete'. Below the progress bar, it states 'Supported filetypes: Excel files, CSV'. The 'File Name' field shows 'No File Selected' with a 'Browse files' button. A note indicates 'Maximum file size 20 MB'. The next section, 'Specify the codepage, separator, and date or time formats of the source file', includes a 'Row one contains the column names' toggle set to 'Yes'. The 'Code page' is set to '1208' with a 'Learn More' link. For 'Separator character', 'comma' is selected. The 'Does the file have columns that contain dates or times?' toggle is set to 'No'. A 'Cancel' button is at the bottom.

Figure 2-31 Rapidly get your local data loaded with Quick Load

Figure 2-32 shows the Run R scripts that you can execute straight from the console.



Figure 2-32 Run R scripts straight from the console

Figure 2-33 shows how you can view all of your data in one place.

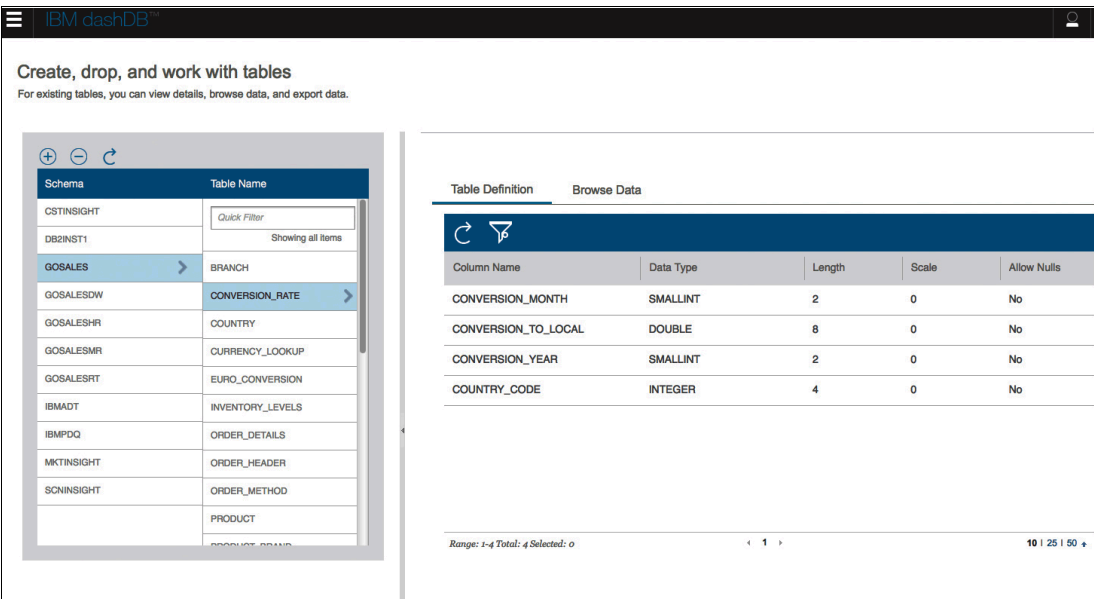


Figure 2-33 View all of your data in one place

Figure 2-34 shows how to easily connect your business intelligence or analytics-focused application.

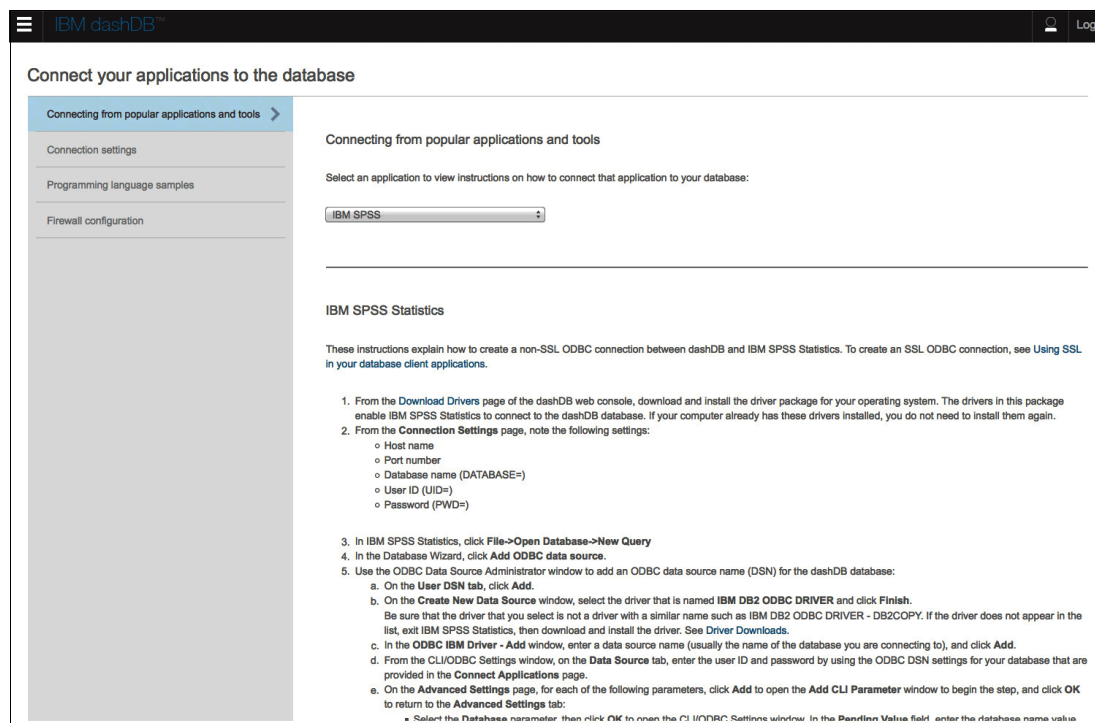


Figure 2-34 Easily connect your business intelligence or analytics-focused application

2.9.3 Geospatial Analytics (IBM)

Expand the boundaries of your application. Use real-time geospatial analytics to track when devices enter or leave defined regions. These are the main capabilities:

- Monitor device locations in real-time

Connect to data sources that support the MQ Telemetry Transport (MQTT) protocol and monitor devices as they move into geographic regions of interest.
- Control region monitoring using the geospatial API

Define geographic regions and control monitoring of regions using the geospatial application programming interface.

2.9.4 IBM Analytics for Apache Hadoop (IBM)

IBM Analytics for Apache Hadoop combines the power of 100% open source Apache Hadoop technology and Hadoop analytics capabilities from IBM BigInsights for Apache Hadoop to deliver insights across your organization. It provides a two-node Hadoop cluster. IBM Analytics for Apache Hadoop is an extremely quick and simple way to try out BigInsights features without having to worry about installing, configuring, or administering a Hadoop cluster. For production use or enterprise-level POCs, use the IBM BigInsights for Apache Hadoop Bluemix service and provision one or more scalable clusters on bare metal hardware.

This service provides these main capabilities:

- ▶ Immediately build big data applications
- ▶ This service provides an easy way to access data on Hadoop clusters, build applications, and analyze structured or unstructured data. Visualize your findings in charts and graphs; you can bring your data into Hadoop for analysis by using Big SQL, BigSheets, Text Analytics, Big R, or Machine Learning.
- ▶ Built on open source technology

This service is powered by IBM Open Platform, which is based on Apache Hadoop. It combines the power of 100% open source Hadoop APIs and other Hadoop ecosystem projects.

Figure 2-35 shows the Ambari Console screen.

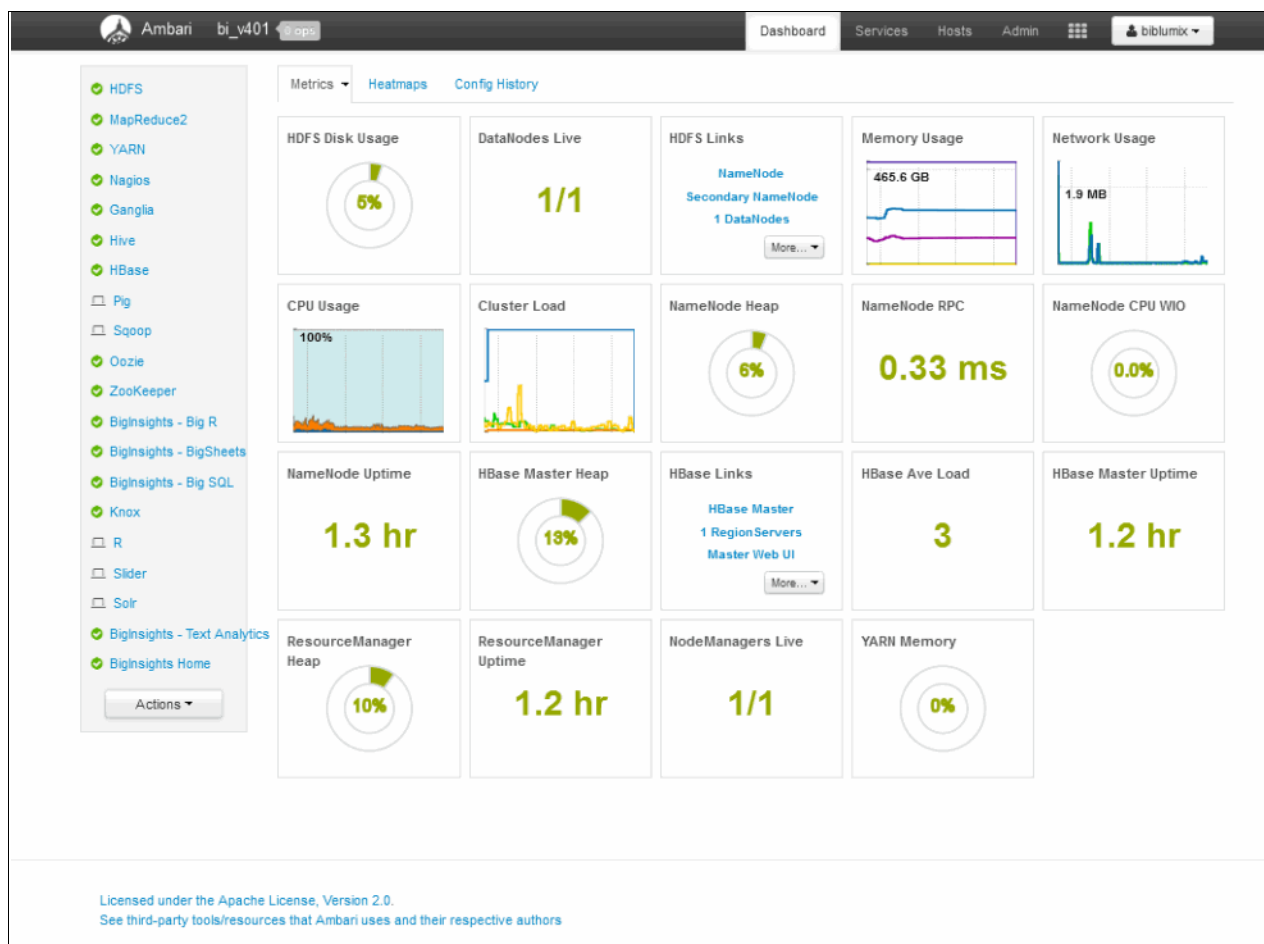


Figure 2-35 Ambari Console screen

Figure 2-36 shows the BigSheets screen.



Figure 2-36 BigSheets screen

Figure 2-37 shows the Big SQL screen.

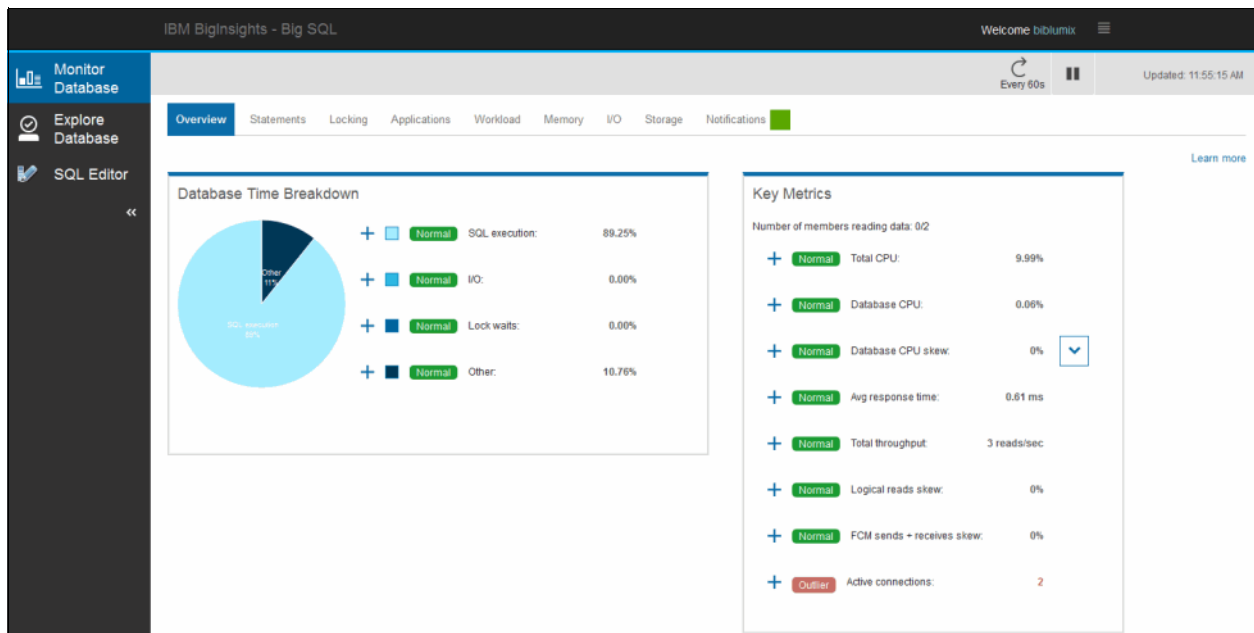


Figure 2-37 Big SQL screen

Figure 2-38 shows the Text Analytics screen.

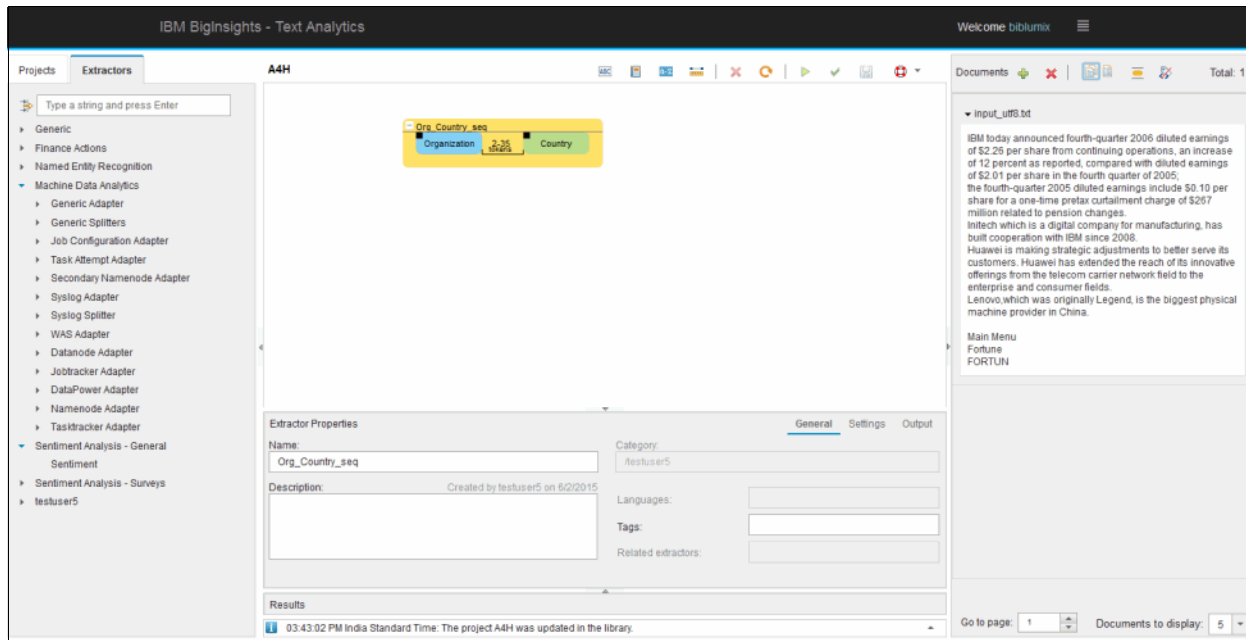


Figure 2-38 Text Analytics screen

2.9.5 Insights for Twitter (IBM)

Insights for Twitter provides sentiment and other enrichments for multiple languages, based on deep natural language processing algorithms from IBM Social Media Analytics. Real-time processing of Twitter data streams is fully supported and is configurable through a rich set of query parameters and keywords. Insights for Twitter includes RESTful APIs that allow you to customize your searches and returns Tweets and enrichments in JSON format. Insights for Twitter main capabilities are:

- Twitter data

Search Twitter content from the Twitter Decahose. The content store is frequently refreshed and indexed, making searches dynamic and fast.

- Enrichments

Get advanced enrichments based on real-time analysis of the Twitter Decahose stream, such as author with Gender and Permanent Location (defined by country, state, and city) and Sentiment (for example, positive, negative, ambivalent, or neutral for Tweets in English, German, French, and Spanish).

- Compliance checking

To validate IBM Insights for Twitter search results, the service provides a REST API method that confirms whether a particular Tweet is still accessible on Twitter.

2.9.6 Time Series Database (IBM)

Time Series Database (powered by IBM Informix) is purpose-built for fast and efficient storage and analysis of time series data. This services main capabilities are:

- ▶ Scalable performance

Consolidates and organizes massive amounts of time-stamped data for consistently fast analysis.

- ▶ Easy

Supports multiple APIs so applications and reports can seamlessly access time-stamped data alone or in combination with other data types (SQL, spatial, JSON). No knowledge of time series is required.

2.10 Security

The Security category provides services to build security into application design. This category includes the following services:

- ▶ Application Security Manager (IBM)
- ▶ AppScan Dynamic Analyzer (IBM)
- ▶ Mobile Analyzer for iOS (IBM)
- ▶ Single Sign On (IBM)
- ▶ Static Analyzer (IBM)
- ▶ aPersona Adaptive Security Manager (third party)

2.10.1 Application Security Manager (IBM)

To manage the challenge of addressing app security at the enterprise level, security teams must take a risk-based approach. This approach means that the team must prioritize assets, focus on identifying areas of highest risk, and then mitigate the risk. Large organizations might have thousands of apps that serve various purposes. The responsibility to assess and address application security typically belongs to a small security team.

Application Security Manager provides a set of capabilities that enable organizations to take a strategic, risk-based approach to the application security problem. Application security, just like any other area of security, is about understanding, managing, and mitigating risk to critical assets. Unless the approach taken to address application security is based on managing risk, an organization's application security initiative will probably never become effective.

Application Security Manager provides a risk-based approach through four simple steps:

1. Create an app. (Figure 2-39 on page 65 shows how to create an app.)
2. Import a CSV file of issues discovered by third-party tools (Figure 2-40).
3. Triage the issues. All issues are classified as *open* by default. You can see the issue classification by viewing the issue status (Figure 2-41).
4. Evaluate business risks based on the security risk rating that Application Security Manager calculates for each application (Figure 2-42).

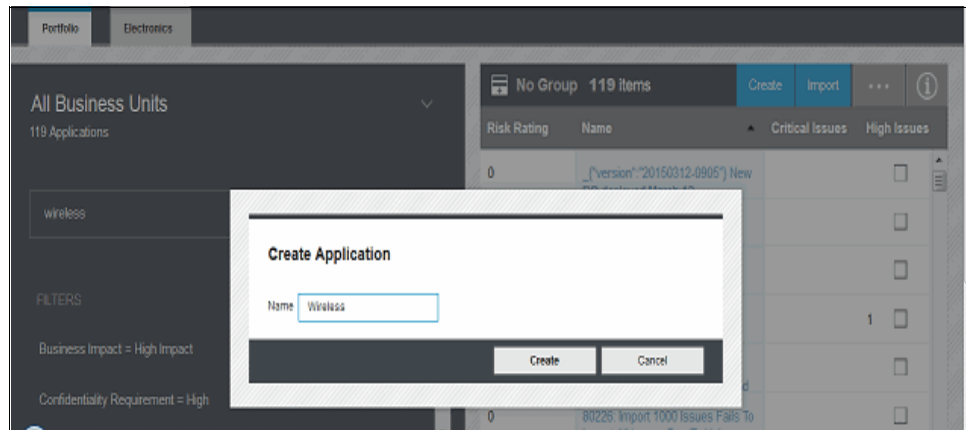


Figure 2-39 Create an app

Figure 2-40 shows importing issues into your app.

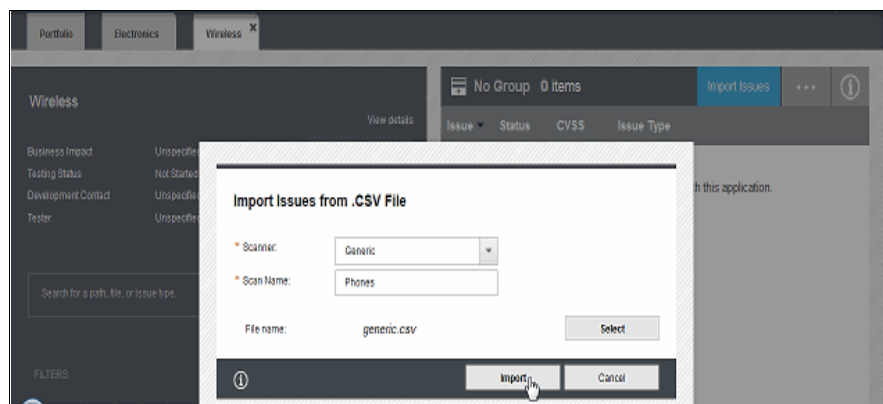


Figure 2-40 Import issues into your app

Figure 2-41 shows categorizing issues.

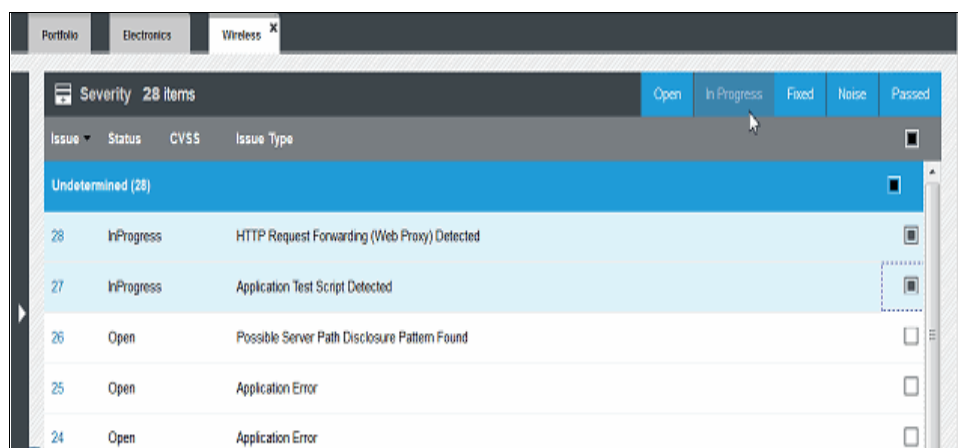


Figure 2-41 Categorize issues

Figure 2-42 shows evaluating an app's security risk.

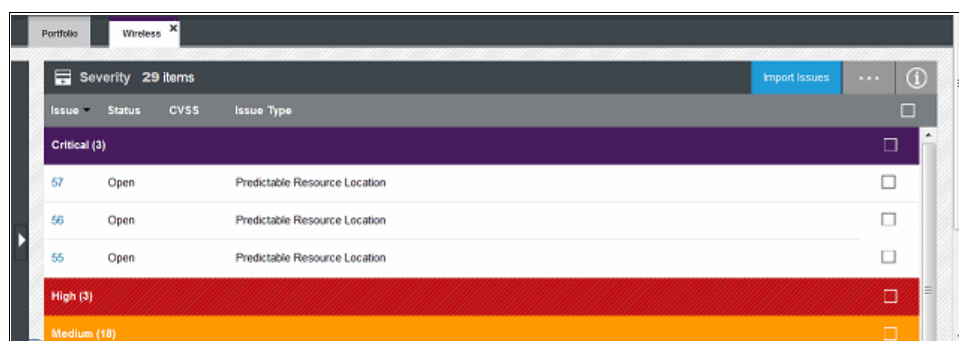


Figure 2-42 Evaluate an app's security risk

2.10.2 AppScan Dynamic Analyzer (IBM)

IBM AppScan Dynamic Analyzer brings the power of AppScan's proven Dynamic Application Security Testing to the cloud to help you identify security issues in web applications and keep them secure. Based on significant application security experience, *IBM Security AppScan* services are capable of identifying dozens of vulnerabilities with a high degree of accuracy.

To use AppScan Dynamic Analyzer, add the service and the starting URL for the scan. Simply select your web application (see Figure 2-43 on page 67), and, if needed, enter login credentials (Figure 2-44 on page 67).

When the scan is complete, the results are summarized in a graph and, if you selected that option, you receive an email notification. Use the **Download Report** link to download the full security report to your computer.

The security report (Figure 2-45 on page 68) shows the application's vulnerability to a range of security issues, indicating severity, number of issues found, associated security risks, suggested remediation tasks, and other details. The dashboard displays your scan results, subscription status, and re-scan options. Each scanned application is listed in its own row and shows the results of the most recent scan.

Following your first scan, if you fix issues, you can scan the application again multiple times.

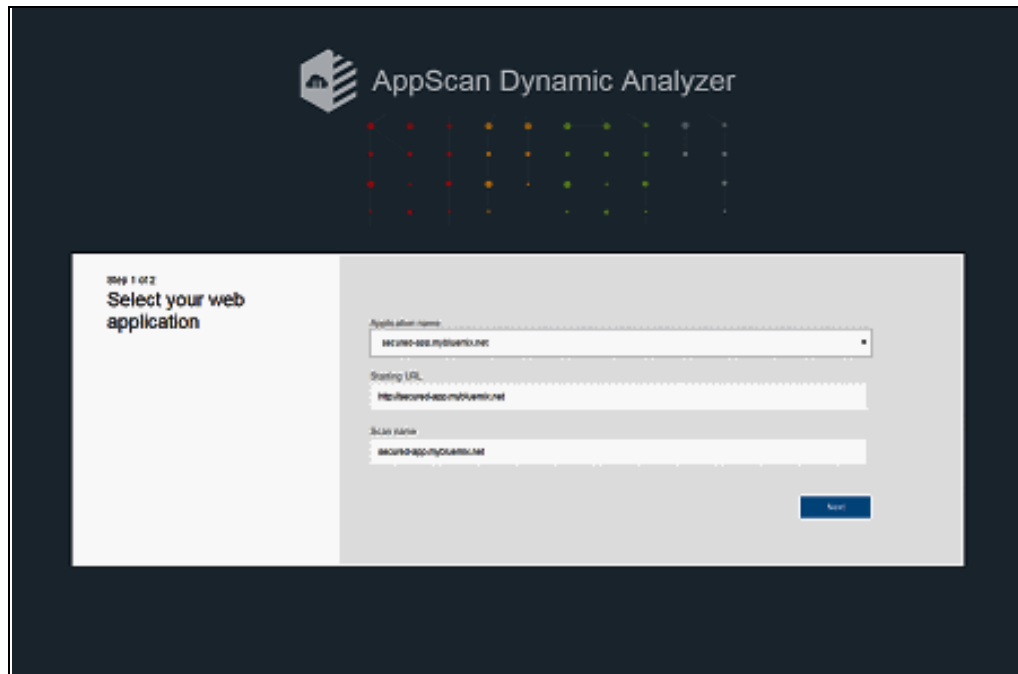


Figure 2-43 Select your web application

Figure 2-44 shows entering login credentials.

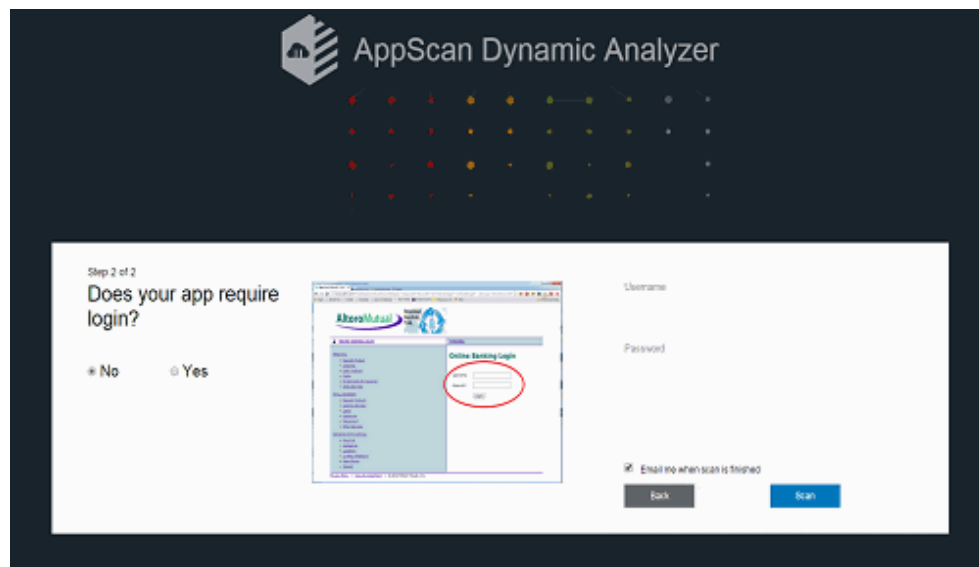


Figure 2-44 Enter login credentials

Figure 2-45 shows a security report.



Figure 2-45 Security report

2.10.3 AppScan Mobile Analyzer (IBM)

IBM Security AppScan Mobile Analyzer brings the proven power of AppScan to Android mobile applications to identify security issues in your Android applications and help you keep your applications secure. Based on IBM experts' significant application security experience, AppSan Mobile Analyzer service can identify dozens of vulnerabilities with a high degree of accuracy.

To use AppScan Mobile Analyzer, simply add the service and upload your Android application package (APK) file (Figure 2-46 on page 69) and name your scan (Figure 2-47 on page 69). The Scan list (Figure 2-48 on page 70) shows all applications that have been scanned, so far. After the scan completes, download a security report (Figure 2-49 on page 70, Executive Summary) to help you to remediate detected vulnerabilities.

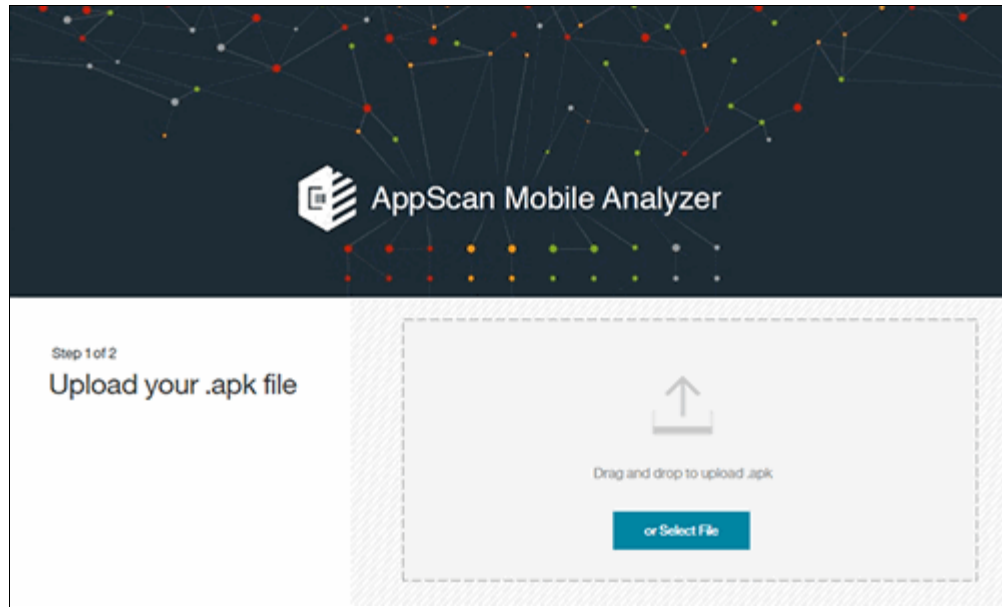


Figure 2-46 Select your APK file

Figure 2-47 shows naming your scan.

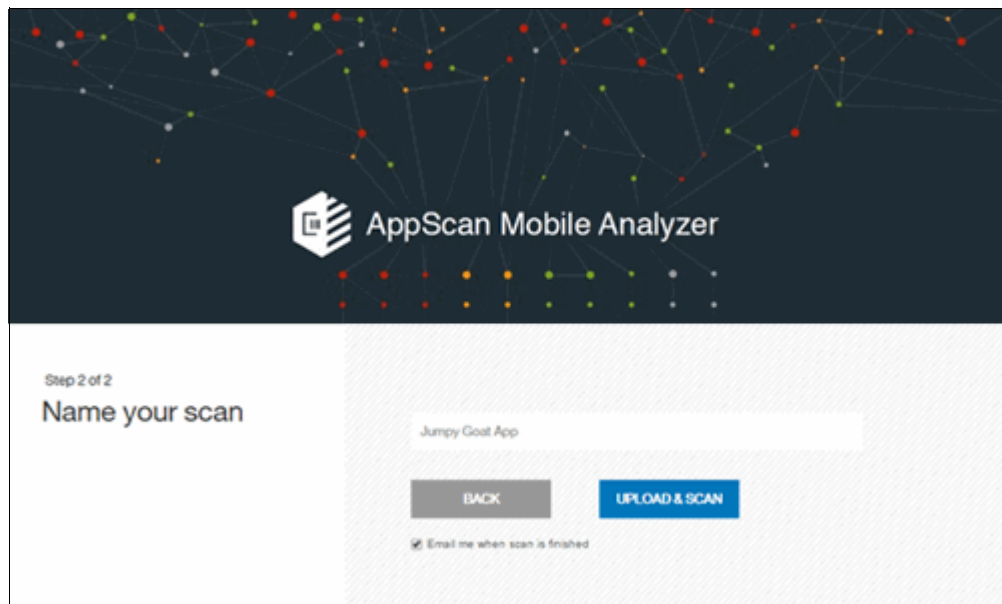


Figure 2-47 Name your scan

Figure 2-48 shows the Scan list.

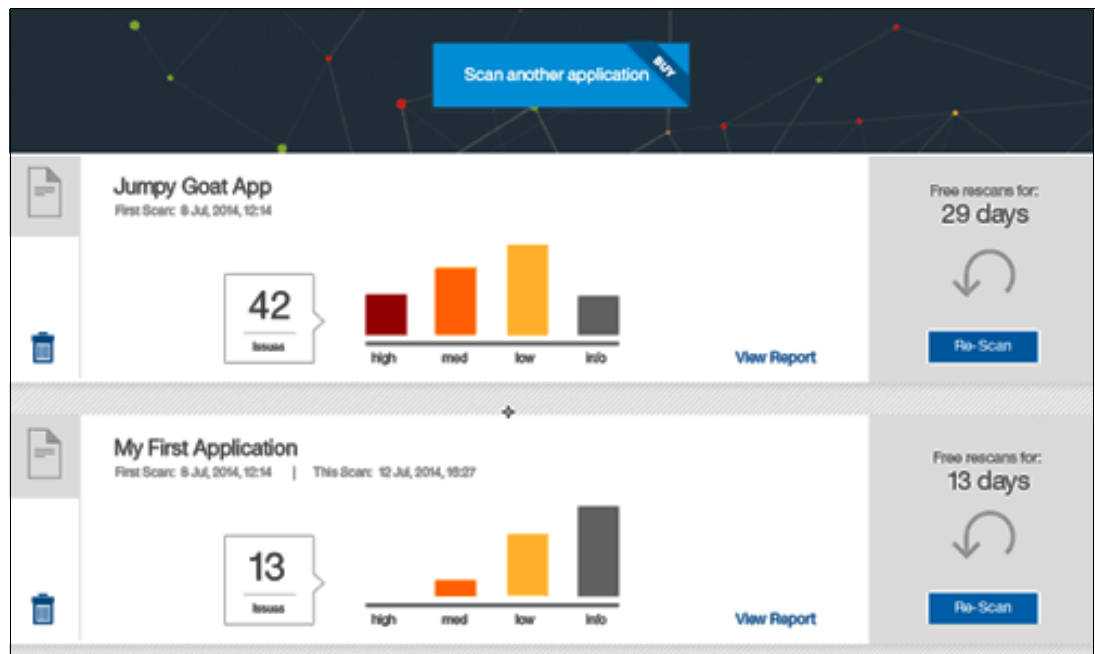


Figure 2-48 Scan list

Figure 2-49 shows the Executive Summary.

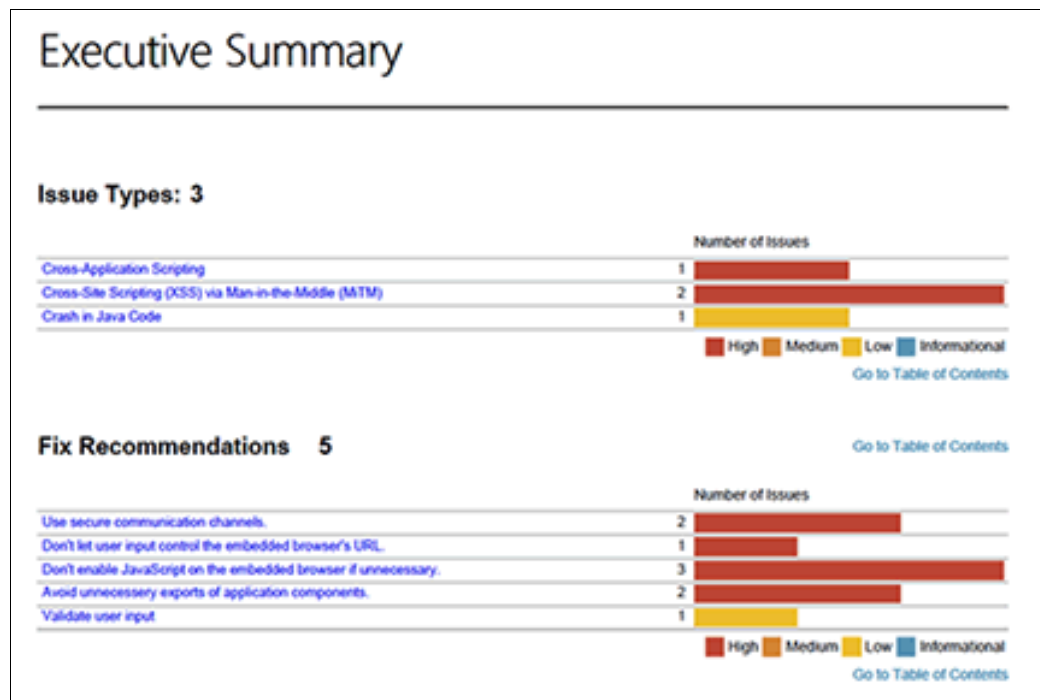


Figure 2-49 Executive Summary of the security report

Figure 2-50 shows the AppScan Mobile Analyzer.

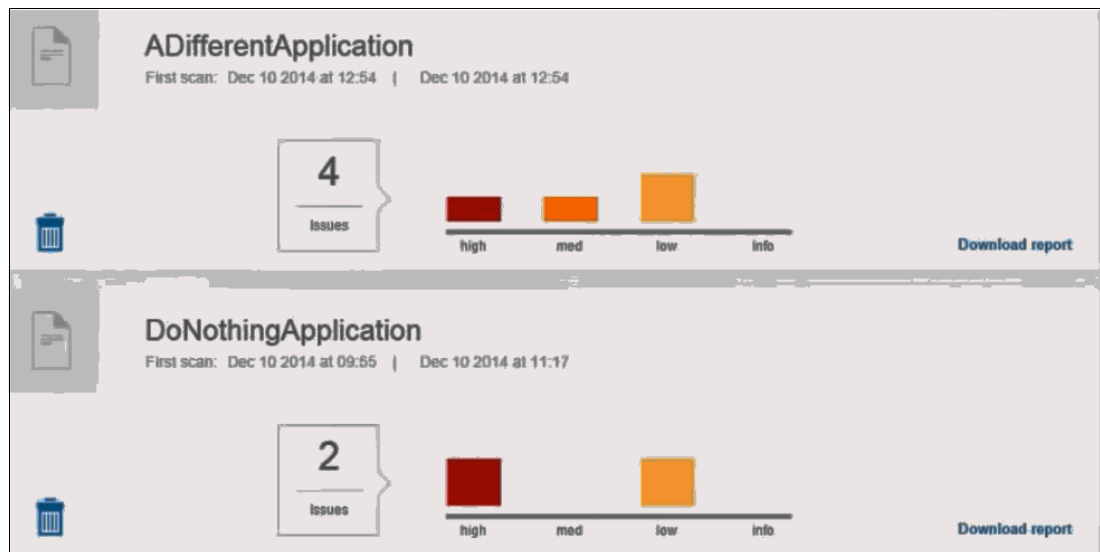


Figure 2-50 IBM Security AppScan Mobile Analyzer

2.10.4 Mobile Analyzer for iOS (IBM)

Mobile Analyzer for iOS brings the proven power of AppScan to iOS mobile applications. It identifies security issues in your iOS applications to help you keep them secure. The service scans Xcode projects or workspaces.

Based on IBM experts' significant application security experience, the Mobile Analyzer for iOS service can identify dozens of vulnerabilities with a high degree of accuracy.

There are just three simple steps to work with AppScan Mobile Analyzer for iOS:

1. Add the service.
2. Use the supplied IPAX generator to generate an IPAX file from your Xcode project.
3. Upload the generated IPAX file to the service.

After the scan completes, download a security report to help you remediate detected vulnerabilities. The dashboard displays your scan results, subscription status, and re-scan options. When you use the Scan Again option, the user credentials from the original scan are automatically used, and the new scan overwrites the original scan.

Figure 2-51 shows the Download IPAX generator.

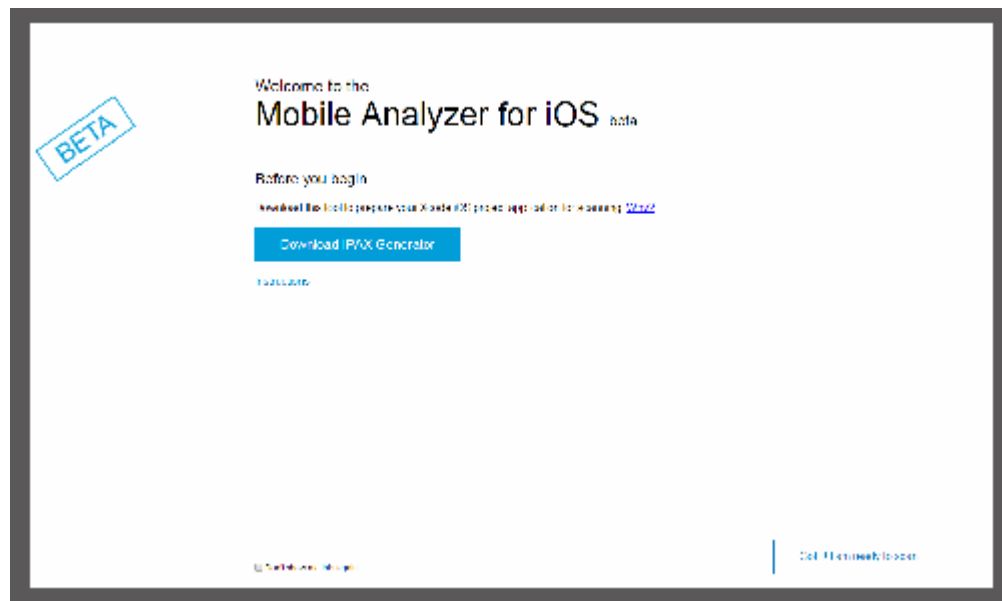


Figure 2-51 Download IPAX generator

Figure 2-52 shows selecting your IPAX file.

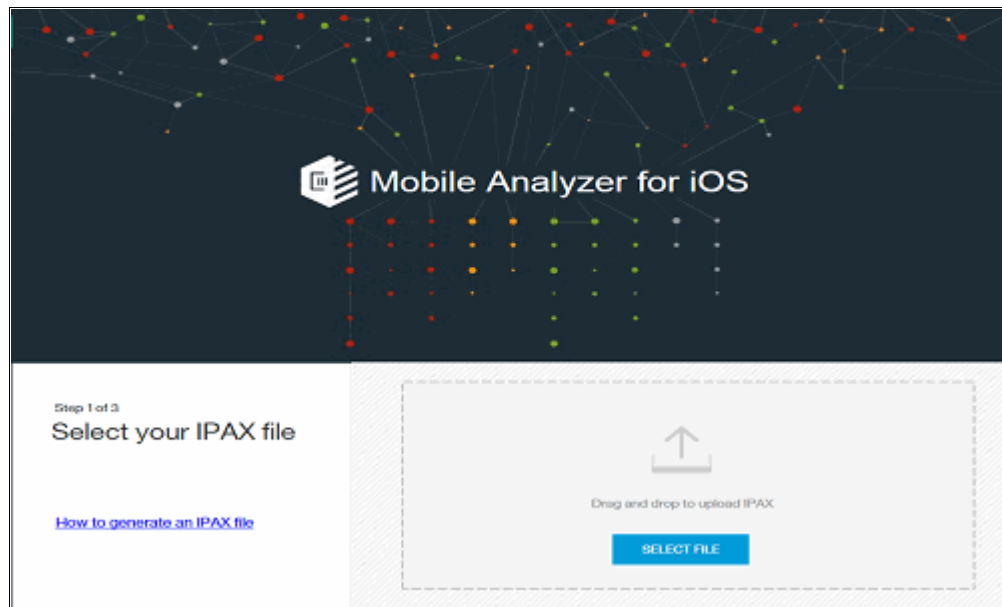


Figure 2-52 Select your IPAX file

Figure 2-53 shows the Scan List view.



Figure 2-53 Scan List

Figure 2-54 shows the security report.

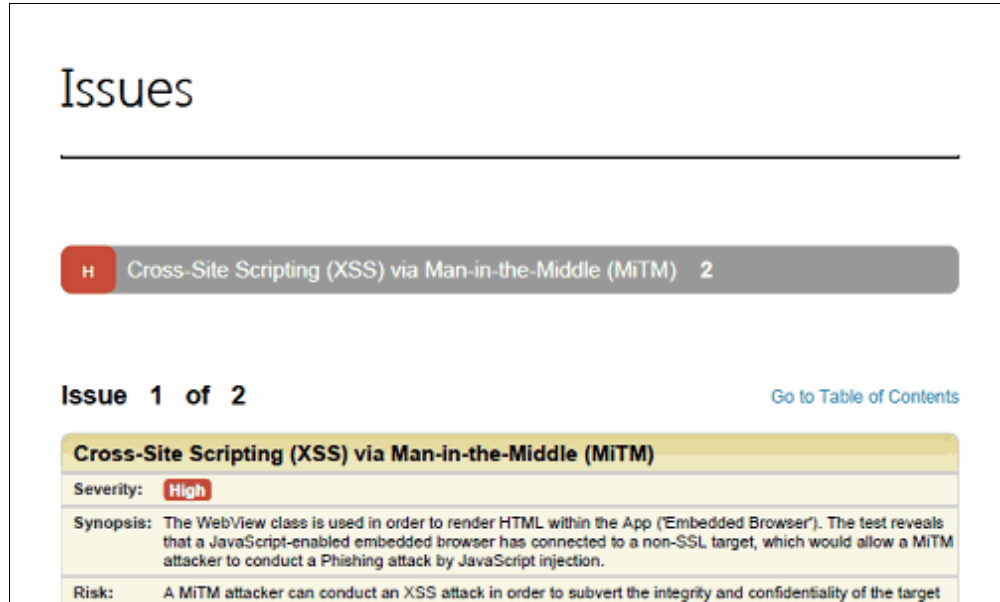


Figure 2-54 Security report

2.10.5 Single Sign On (IBM)

Single Sign On enables you to implement user authentication for your web and mobile apps quickly, using simple policy-based configurations. These are among the Single Sign On main capabilities:

- ▶ Secure apps with confidence, not a lot of coding.
Add user authentication to your apps with policy-based configuration options and an easy to use SDK. Are you writing Java apps? Take advantage of our zero-coding approach.
- ▶ You choose the identity sources and we do the rest.

Whether you are using an existing enterprise directory with Security Assertion Markup Language (SAML), popular social identity sources, such as Facebook, LinkedIn, and Google, or you want to create your own cloud directory, it's easy to set up. Apps do not require information about all of the sources that users might authenticate from.

Figure 2-55 shows selecting the identity sources to authenticate your users.

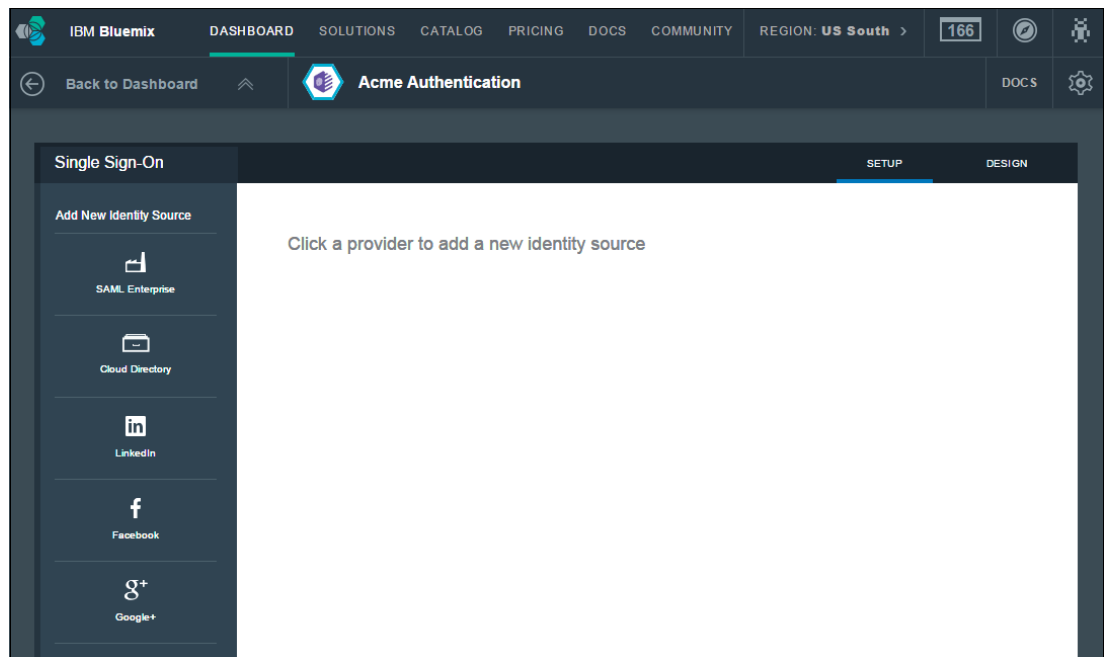


Figure 2-55 Select the identity sources to authenticate your users

Figure 2-56 shows integrating with an existing on-premises directory using SAML.

IBM Bluemix DASHBOARD SOLUTIONS CATALOG PRICING DOCS COMMUNITY REGION: US South > 166

Back to Dashboard Acme Authentication DOCS

Acme Enterprise Directory

SAML Enterprise

Save

Configure Identity Source

Step 1

From Identity Provider

The SAML Enterprise Identity Source allows an application to consume user identity information from a SAML Identity Provider.

Enter the URL to initiate a SAML Single Sign-On flow from the Identity Provider.

Upload the SAML 2.0 Identity Provider metadata file to enable the SAML Enterprise Identity Source to communicate with the Identity Provider.

Upload File

Step 2

To Identity Provider

The SAML Enterprise Identity Source acts as a SAML Service Provider to obtain identity information.

Download the SAML 2.0 Service Provider metadata file. Give the file to your identity provider to complete the configuration.

Download File

Download the Identity Bridge from the [IBM Fix Central](#) web site. Login to the site with your IBM ID.

[Click here to view additional information on SAML.](#)

Figure 2-56 Integrate with an existing on-premises directory using SAML

Figure 2-57 shows creating and managing users with your own in-cloud directory.

IBM Bluemix DASHBOARD SOLUTIONS CATALOG PRICING DOCS COMMUNITY REGION: US South > 166

Back to Dashboard Acme Authentication DOCS

Acme Cloud App Users

Cloud Directory

Save

Search text

Username	Email	Added
----------	-------	-------

Add User

Username
Acmeuser

Password

Verify password

Given name

Surname

Email

Figure 2-57 Create and manage users with your own in-cloud directory

Figure 2-58 shows configuring your service to authenticate users with social identity providers.

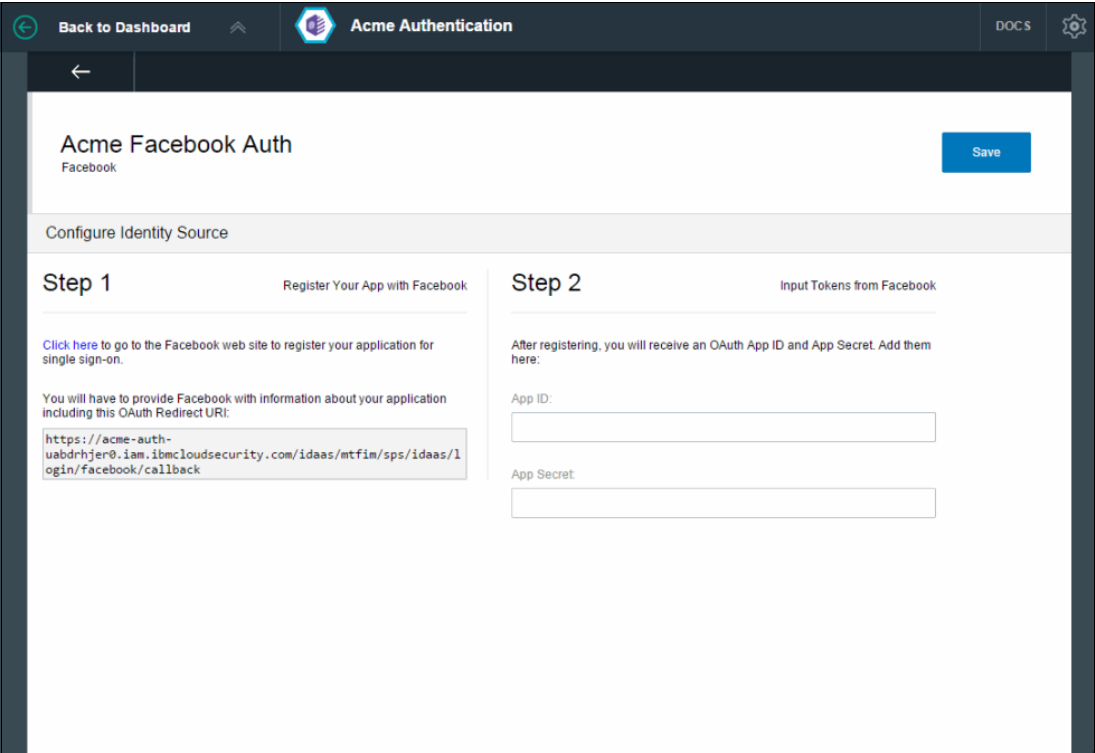


Figure 2-58 Configure your service to authenticate users with social identity providers

Figure 2-59 shows that after setting up the service, you bind it to your applications.

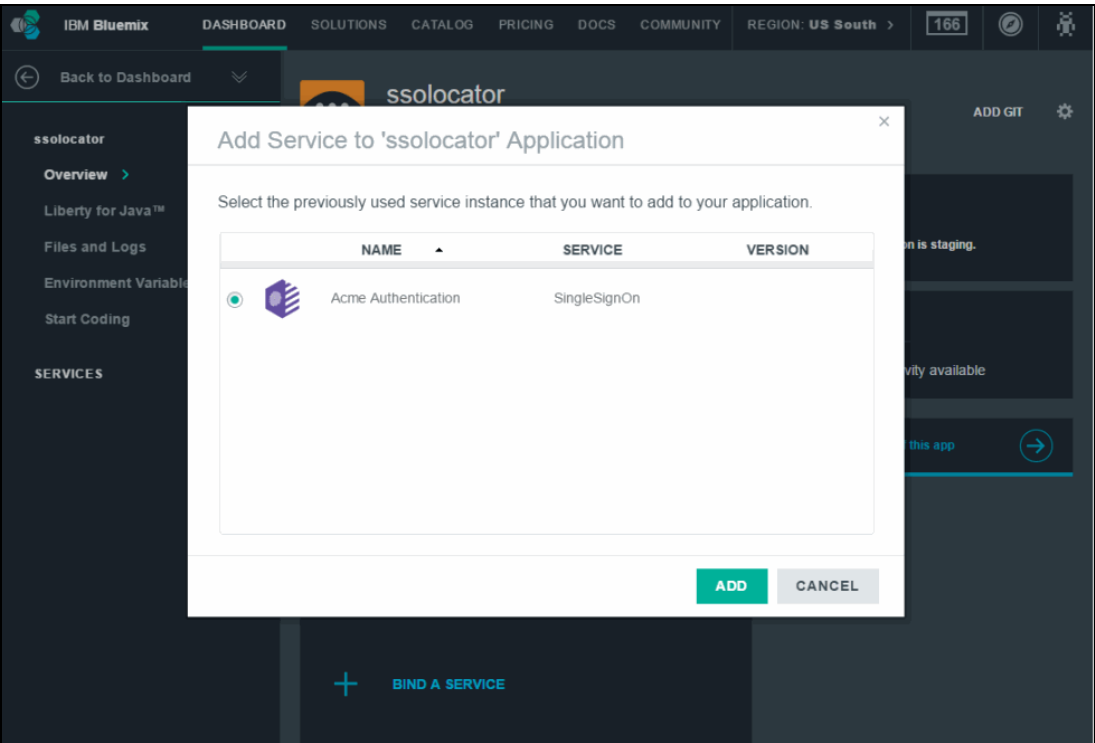


Figure 2-59 After setting up the service, bind it to your application

Figure 2-60 shows completing the setup by integrating the service into your application.

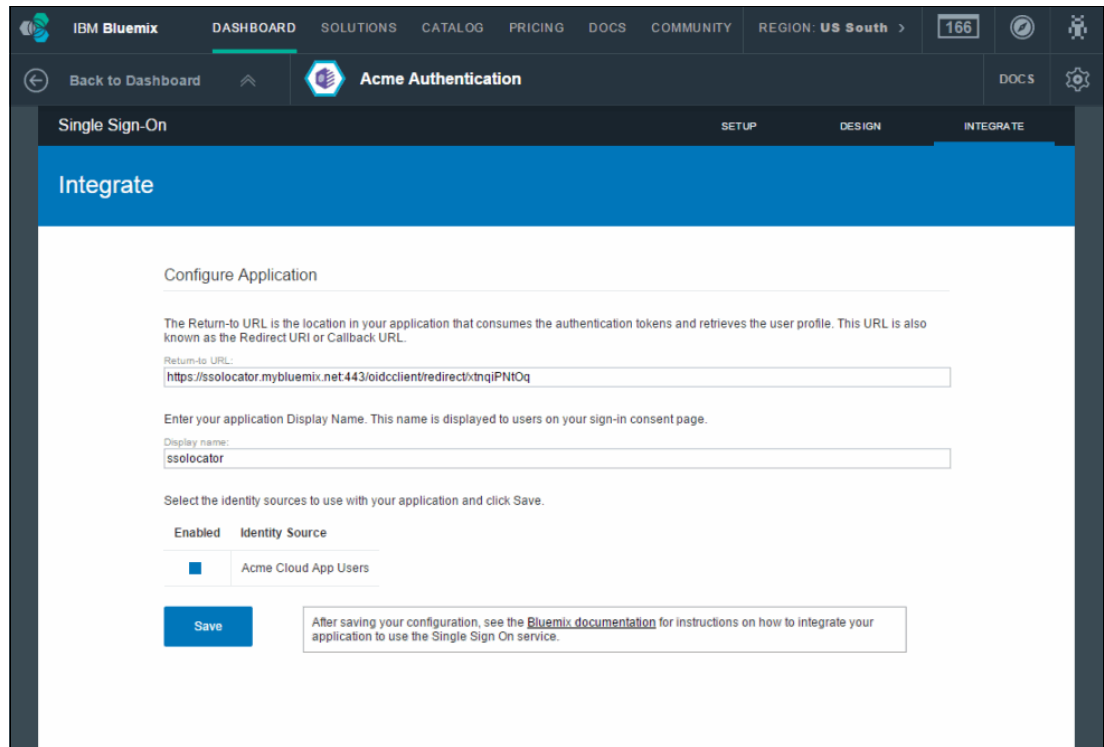


Figure 2-60 Complete the setup by integrating the service into your application

2.10.6 Static Analyzer (IBM)

The Static Analyzer service brings the power of static application security testing to the cloud. This service helps you find source code vulnerabilities early in the software development lifecycle so that these can be fixed before deployment.

Combining the IBM static analysis engine with new Intelligent Findings Analytics, Static Analyzer can identify key vulnerabilities with a high degree of accuracy.

Static Analyzer enables you to scan Java and Java web content by using a small client command-line interface (CLI) that you simply download and extract to your local disk. If you prefer, you can run a small installer that adds Static Analyzer plug-ins to Eclipse or Maven. After you have the Client Utility on your computer or applied to Eclipse or Maven, you can use it to scan and gather information about your files in the form of an archive file (Intermediate Representation, or IRX file). You then submit the IRX file to the cloud to receive scan results. (The IRX file is encrypted to keep your data safe.) You can connect to the cloud, upload IRX files, and check the status of analysis jobs in Bluemix. If you prefer, you can use the CLI for these tasks.

When a scan is complete, you receive email notification (if the scan is submitted via the CLI, you will automatically receive email; if the scan is submitted via the Bluemix user interface, you will receive email if you select the email notification option). The email includes a link to Bluemix so that you can log in to download your scan results.

Then, use the report to remediate detected vulnerabilities.

Figure 2-61 shows getting started with Static Analyzer.

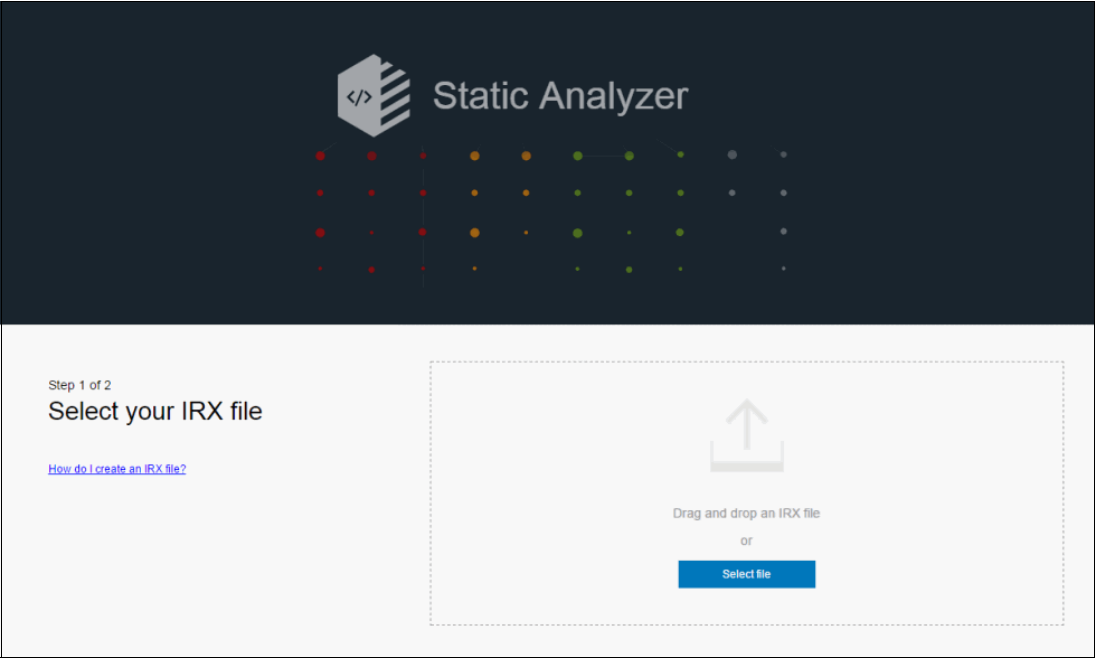


Figure 2-61 Getting started

Figure 2-62 shows the Scan list.

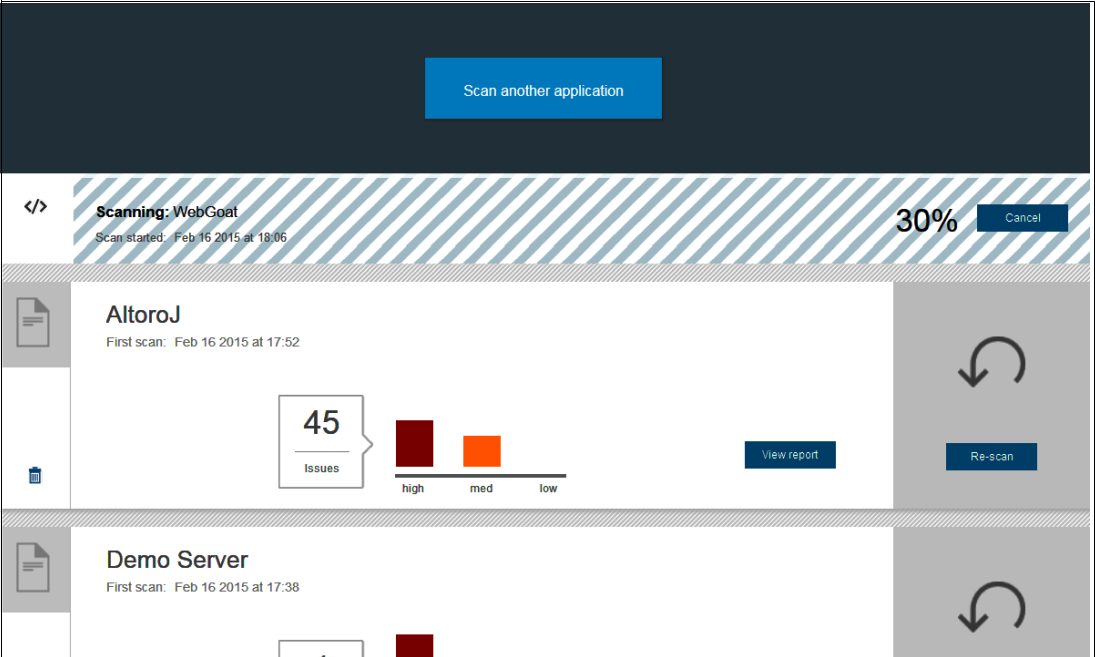


Figure 2-62 Scan list

Figure 2-63 shows a Static Analyzer report.



Figure 2-63 Static Analyzer report

2.10.7 aPersona Adaptive Security Manager (third party)

The Adaptive Security Manager (ASM) from aPersona is a multi-tenant frictionless adaptive multi-factor authentication platform and service that protects web based transactions (including logins) from fraud and account takeover by unauthorized users.

2.11 Business Analytics

Business Analytics category provides services to make easy powerful analytics made easy. This service provides these main capabilities:

- ▶ Embeddable Reporting (IBM)
- ▶ Predictive Modeling (IBM)
- ▶ Cupenya Insights (third party)

2.11.1 Embeddable Reporting (IBM)

This helps you include dashboards and reports in your web or mobile application by using a simple cloud editor, which removes the need to understand complex query syntax and expressions that are required for rich visualizations. Then, embed reports and dashboards, using a wide variety of languages, such as Node.js or Java, by using a RESTful API. This service offers these main advantages:

- ▶ Simple authoring environment

Developers, designers, and product managers can easily collaborate on data-based assets, such as reports to be included in the application. They can do so without the need to manage data sources, visualization engines, and query syntax across source code. By using our cloud authoring environment, the service helps you create assets ranging from a simple query for displaying a metric to joining data from multiple sources for a sophisticated dashboard.

- ▶ RESTful API

Embedding a report or dashboard is as easy as an HTTP GET via the RESTful API. Access all of the content that you have created the way that you prefer, with support for wide variety of languages, frameworks and tools. The service minimizes code changes that are required to embed reports, which enables you to modify the data in a report through parameters.

Figure 2-64 shows creation and management of reports for applications from a single view.

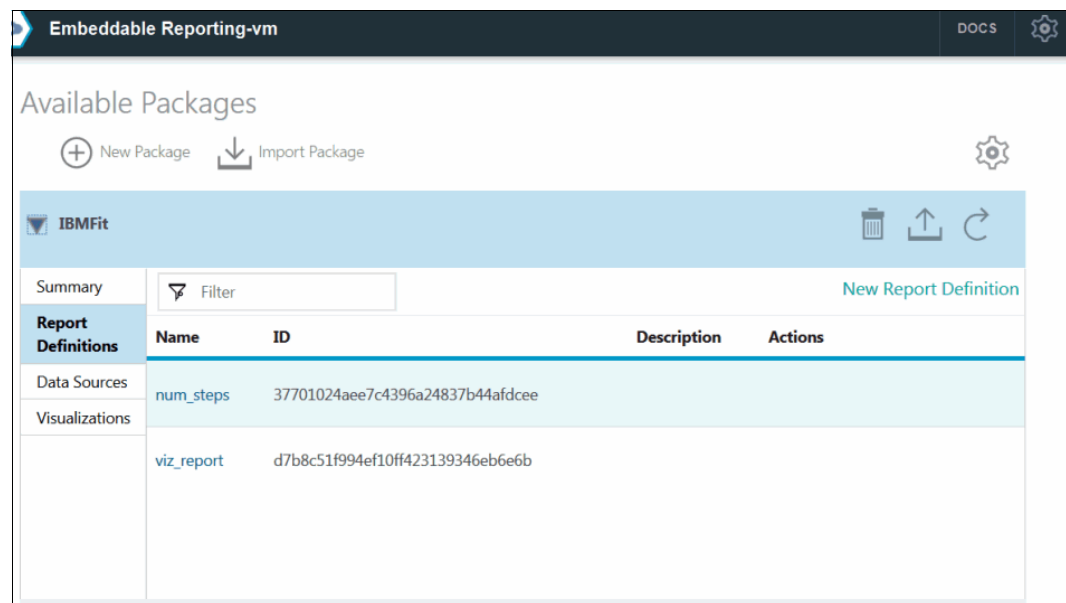


Figure 2-64 Create and manage reports for applications from a single view

Figure 2-65 shows how to easily create new report packages.

The screenshot shows a modal dialog titled 'Provide A Name And An Optional Description For Your New Package.' with a close button (X) in the top right corner. The dialog is divided into two sections: 'Package Information' and 'Connection Information For The Default Datasource'. The 'Package Information' section has two input fields: 'Name:' and 'Description:'. The 'Connection Information For The Default Datasource' section has a text block stating 'DashDB and SQL database are the 2 supported datasource types. They require the same connection information.' followed by three input fields: 'JDBC URI:', 'User Name:', and 'User Password:'. At the bottom right, there are two buttons: 'CANCEL' and 'CREATE'.

Figure 2-65 Easily create new report packages

Figure 2-66 shows how you can combine different data elements and visualizations into a single report.

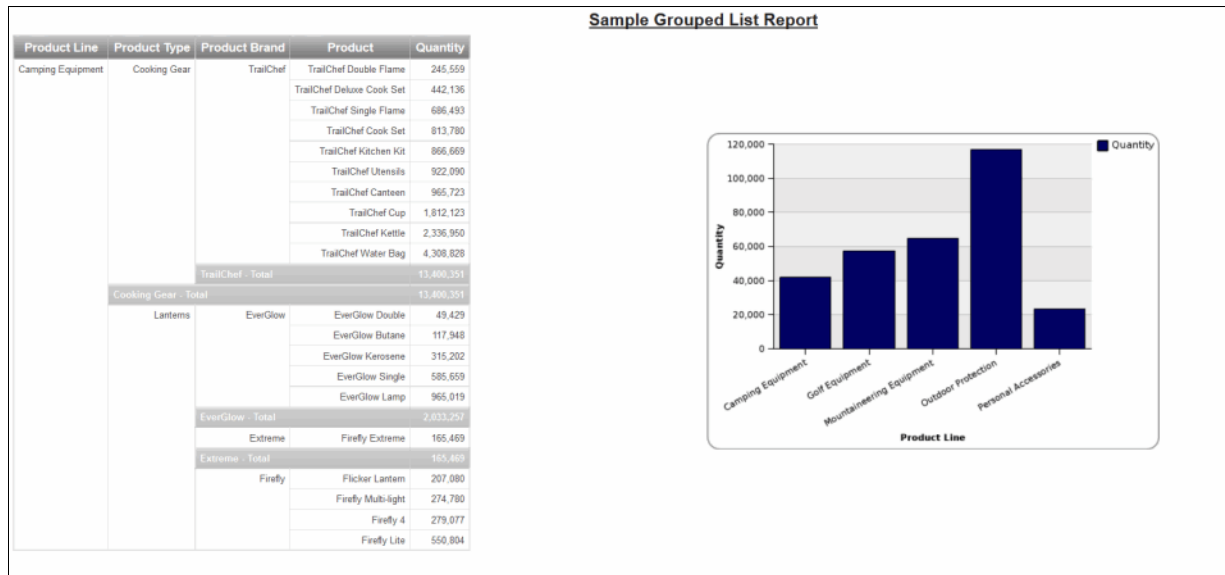


Figure 2-66 Combine different data elements and visualizations into a single report

Figure 2-67 shows samples that are available in JazzHub for quick starting point.

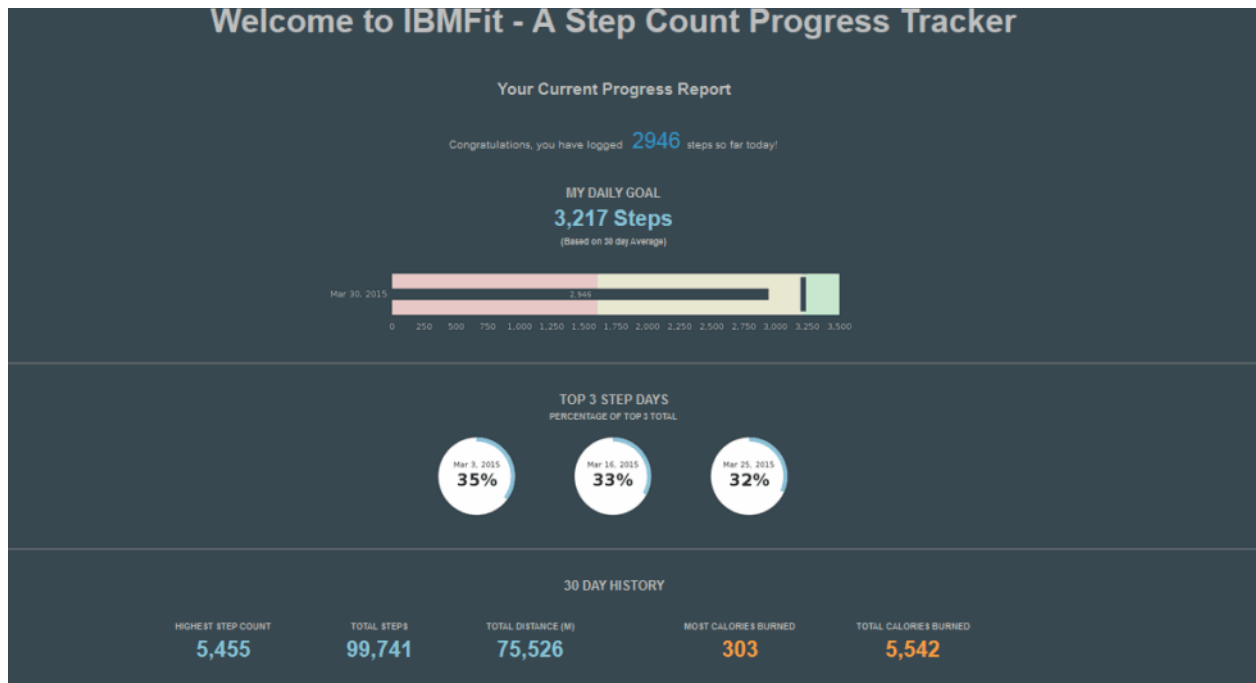


Figure 2-67 Samples available in JazzHub for quick starting point

2.11.2 Predictive Modeling (IBM)

The IBM Predictive Modeling service permits the integration of analytics with your applications. Bind your Bluemix apps to a Predictive Modeling service instance and generate the predictive analytics your app needs to deliver higher value to your users. You manage your models in the administration dashboard, and you may update or refresh your models by using the dashboard without stopping or redeploying your apps.

The focus of the Predictive Modeling service is deployment, and the IBM SPSS® Modeler Client software is required for authoring and working with models and streams.

2.11.3 Cupenya Insights (third party)

Cupenya Insight provides business analytics for the proactive enterprise. You can use it to connect and monitor business activities across your IBM Bluemix applications, define key performance indicator (KPI) targets, and receive an end-to-end performance overview of your business process.

2.12 Internet of Things

Internet of Things category provides services to create the new generation of applications that are capturing data from Internet. These are the two main services currently in the Internet of Things category:

- ▶ Internet of Things (IBM)
- ▶ flowthings.io (third party)

2.12.1 Internet of Things (IBM)

The IBM Internet of Things service lets your apps communicate with and retrieve data collected by your connected devices, sensors, and gateways. Our *recipes* make it easy to get devices connected to our Internet of Things cloud. Your apps can then use our real-time and REST APIs to communicate with your devices and use the data that you have set them up to collect. This service helps you reach two important goals:

- ▶ Connect your devices securely to the cloud

Before your apps can get to work, you need to get your devices connected. We have a set of verified instructions, or “recipes,” for connecting devices, sensors, and gateways from a variety of partners and individuals.

- ▶ Build an app that talks to your devices

Communication between your devices and the cloud is through the open, lightweight MQTT protocol. For example, you might have a sensor that collects and sends humidity readings every minute. Our REST and real-time APIs enable you to quickly pull that device’s data into your apps for further analysis.

Figure 2-68 shows how the capabilities Internet of Things work.

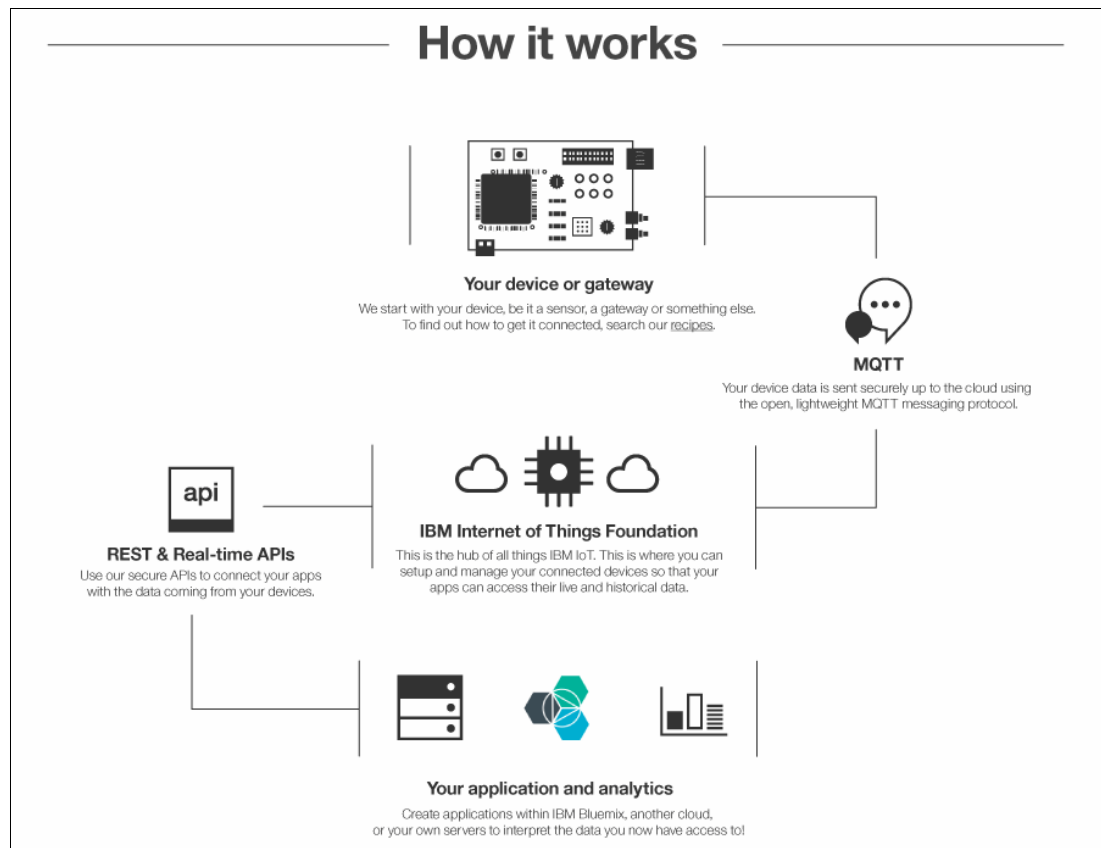


Figure 2-68 How Internet of Things works

Figure 2-69 shows getting started in the Internet of Things Foundation.

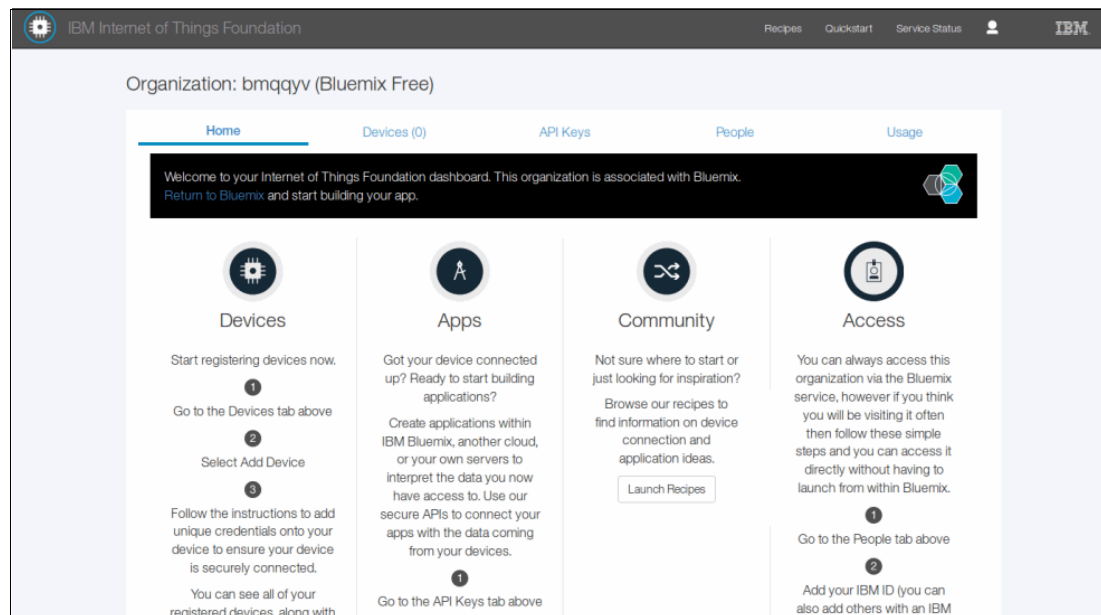


Figure 2-69 Getting started in the Internet of Things Foundation

Figure 2-70 shows device registration.

Figure 2-70 Device registration

Figure 2-71 shows the Devices dashboard.

Device ID	Device Type	Last Event	Message Rate	Date Added	Added By
1367892372389632	Intel Galileo	Just now	Every 5 seconds	Thu Oct 02 2014	User @email.com
hur006	Intel Galileo	Unavailable	-	Tue Aug 05 2014	User @email.com
hur002	mbed NXP LPC1768	2 hours ago	-	Tue Aug 05 2014	User @email.com
hur004	mbed NXP LPC1768	6 hours ago	-	Tue Aug 05 2014	User @email.com
hur001	Raspberry Pi	Unavailable	-	Tue Aug 05 2014	User @email.com
112233445577	CC2541 SensorTag	Unavailable	-	Tue Sep 09 2014	User @email.com
hur005	CC2541 SensorTag	Unavailable	-	Wed Aug 06 2014	User @email.com
hur007	System Utilization	Just now	Every 1 second	Tue Aug 05 2014	User @email.com

Figure 2-71 Devices Dashboard

2.12.2 flowthings.io (third party)

The flowthings.io service empowers any developer or organization to use the growing instrumentation of the physical world (that is, the Internet of Things) to discover information that can be turned into solutions. These findings might surprise you with their intelligence, contextual awareness, and effectiveness in operations and user experiences.

Related publications

The publications listed in this section are considered particularly suitable for a more detailed discussion of the topics covered in this paper.

IBM Redbooks

The following IBM Redbooks publications provide additional information about the topic in this document. Some publications referenced in this list might be available in softcopy only.

- ▶ *Accelerate Development of New Enterprise Solutions for the Cloud with Codename BlueMix*, REDP-5011
- ▶ *IBM Bluemix Architecture Series: Web Application Hosting on Java Liberty*, REDP-5184
- ▶ *IBM Bluemix Architecture Series: Web Application Hosting on IBM Containers*, REDP-5181
- ▶ *Secure Cloud-to-Mainframe Connectivity with IBM Bluemix*, REDP-5243
- ▶ *Microservices: From Theory to Practice. Creating Applications in Bluemix Using the Microservices Approach*, SG24-8275
- ▶ *Creating Applications in Bluemix using the Microservices Approach*, TIPS1309
- ▶ *Getting Started with IBM Bluemix: Web Application Hosting Scenario on Java Liberty*, TIPS1280

You can search for, view, download or order these documents and other Redbooks, Redpapers, Web Docs, draft and additional materials, at the following website:

ibm.com/redbooks

Online resources

These websites are also relevant as further information sources:

- ▶ An introduction to the application lifecycle on IBM Bluemix
<http://www.ibm.com/developerworks/cloud/library/cl-intro-codename-bluemix-video/>
- ▶ IBM Bluemix main page
<http://www.ibm.com/cloud-computing/bluemix/>
- ▶ IBM Bluemix solutions
<http://www.ibm.com/cloud-computing/bluemix/solutions/data-analytics/>
- ▶ IBM Bluemix product page
<http://www.ibm.com/software/bluemix/>
- ▶ IBM Bluemix documentation
<https://www.ng.bluemix.net/docs/>
- ▶ IBM Bluemix catalog
<https://console.ng.bluemix.net/catalog/>

- ▶ IBM Bluemix pricing
<https://console.ng.bluemix.net/pricing/>
- ▶ IBM Bluemix Developers Community
<https://developer.ibm.com/bluemix/>
- ▶ IBM Bluemix page on IBM developerWorks
<https://www.ibm.com/developerworks/cloud/bluemix/>
- ▶ IBM Bluemix DevOps Services
<http://bit.ly/1Jk2MUZ>

Help from IBM

IBM Support and downloads

ibm.com/support

IBM Global Services

ibm.com/services



REDP-5242-00

ISBN 0738454427

Printed in U.S.A.

Get connected

