



IBM

Open by Design

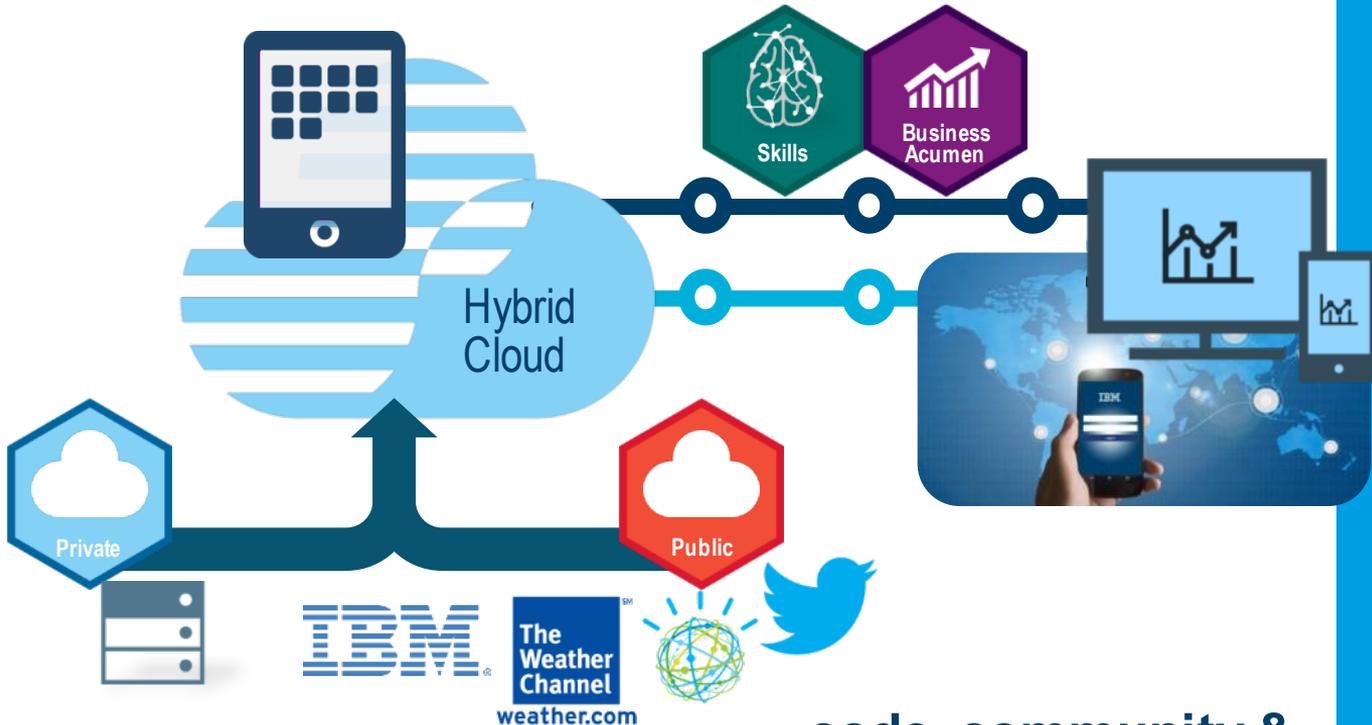
Jun 2016

Presented by:

Jin gi KONG 공진기
Bluemix Technical Evangelist
jgkong@kr.ibm.com

IBM **Cloud**

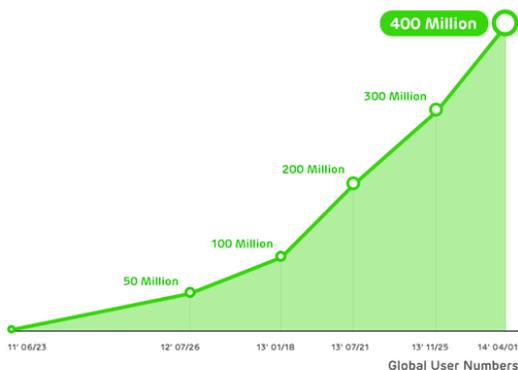
Technology is enabling a new era of innovation...



...code, community & culture are the catalyst

Open technology & open APIs are empowering the world with unprecedented access to data & function

Various Open Source Package for 500 million users



- Redis
- HBase
- MySQL
- Cassandra
- Arcus (memcached cluster)
- Hadoop
- Kafka
- Elastic Search
- Titan



<http://developers.linecorp.com/blog?p=1420>
<http://d2.naver.com/helloworld/809802>
<http://linecorp.com/en/pr/news/en/2014/714>

LINE Corp



Open source & commerce

AliSQL: Alibaba's fork of MySQL

- 40+ new bugs have been found & fixed
- 41 new features have been added
- 27 bottlenecks have been optimized

Widely used in Alibaba Group and Ant Financial (Alipay)

Singles' Day Global Shopping Festival 2015/11/11

- | | | |
|--|--|-----------------------------------|
| – 45M concurrent
online user
(max) | – 140,000 orders
per second (max) | – 14.3B\$ sales
for one day |
| | – 85,900 transaction
per second (max) | – 710M transaction
for one day |



Innovators can no longer afford to act as an island

Value & insight are derived by making connections between raw data, real ability & the power to act

By leading the IT renaissance, built on an Open Cloud Architecture, IBM offers Devs & Partners the opportunity to amplify their impact with the tools & techniques of this new era.

By 2020, 60% of
All companies will be
100% cloud based



Developers



Ecosystem



Partners

55% of CIOs
indicated they would
source all their critical
applications in the cloud
by 2020

Code • Community • Culture

Code: building the lynchpins of open innovation

Centers of open source gravity provide a wealth of innovation & access to valuable skills



- 2.5m lines of code by 3600+ contributors.
- IBM has been a top-3 contributor in the past 4 releases.
- IBM currently has 450+ developers & 20 core reviewers



- 3.2m lines of code by 2200+ contributors.
- IBM has the largest Cloud Foundry deployment
- IBM has 10 dedicated core contributors



- Leading dev of nightly builds for x86, Power & Mainframe systems
- #2 in Commits in 2015 behind Docker

Developers



Ecosystem



Partners

- 1.6b npm downloads in the past month
- IBM a founding member of the newly established foundation



- Pledged donation of IBM SystemML machine
- IBM pledged 3,500 devs researchers & to Spark projects across 12 labs
- Education for 1m + data scientists & engineers on Spark.



How do consumers make choices?

84% trust recommendations made through personal relationships





Level of adoption: Open Technology



Use

- Install
- Run
- Use
- Scale

Change

- Tune parameters
- Fix bugs
- Improve
- Optimize

Create

- Open existing project
- Create new one



What does bad open source look like?

IP assignment too limited or broad

Restrictive licensing

High committer turnover

Arbitrary hierarchies

Lack of usability

Single Source committer

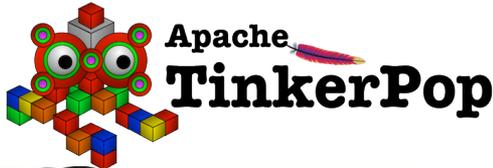
Large dumps of commercial code

How can you tell the difference? One or more of these characteristics may indicate problems



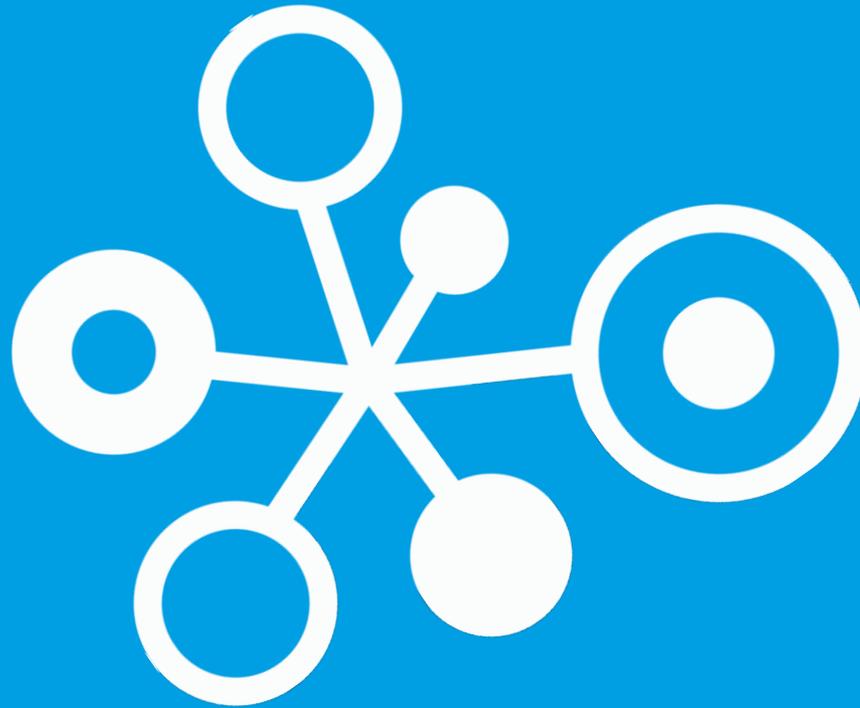
Open Source Projects - Databases

Led by foundation



Led by corporation





Pulling ahead with Open Technology centers of gravity



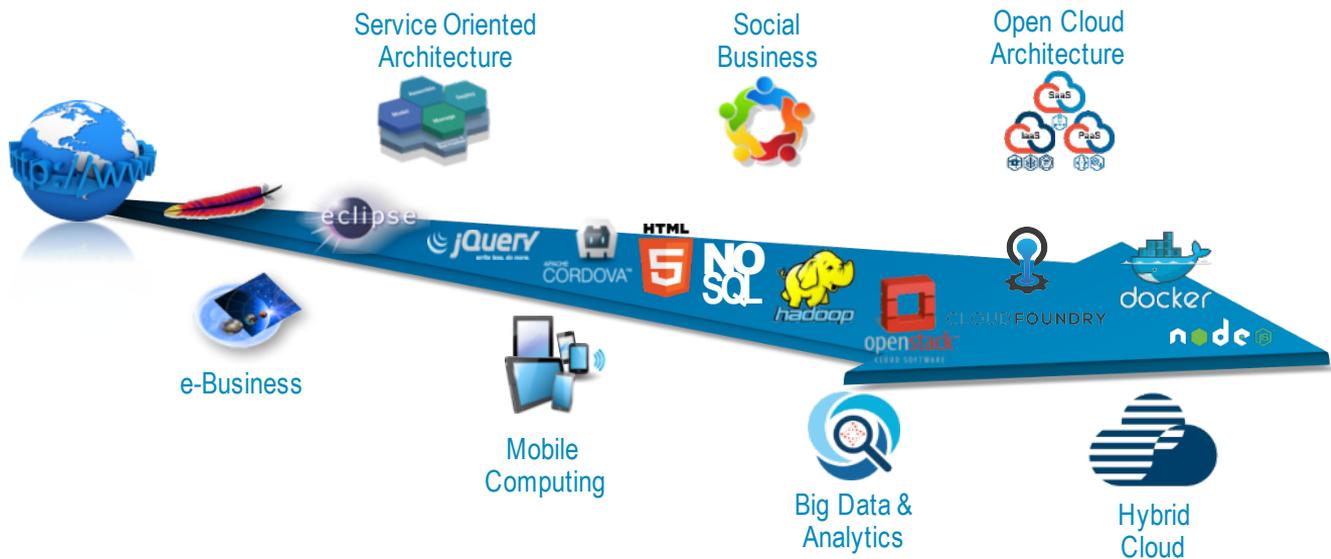


Open tech underpins the Internet & is directly responsible for its success.

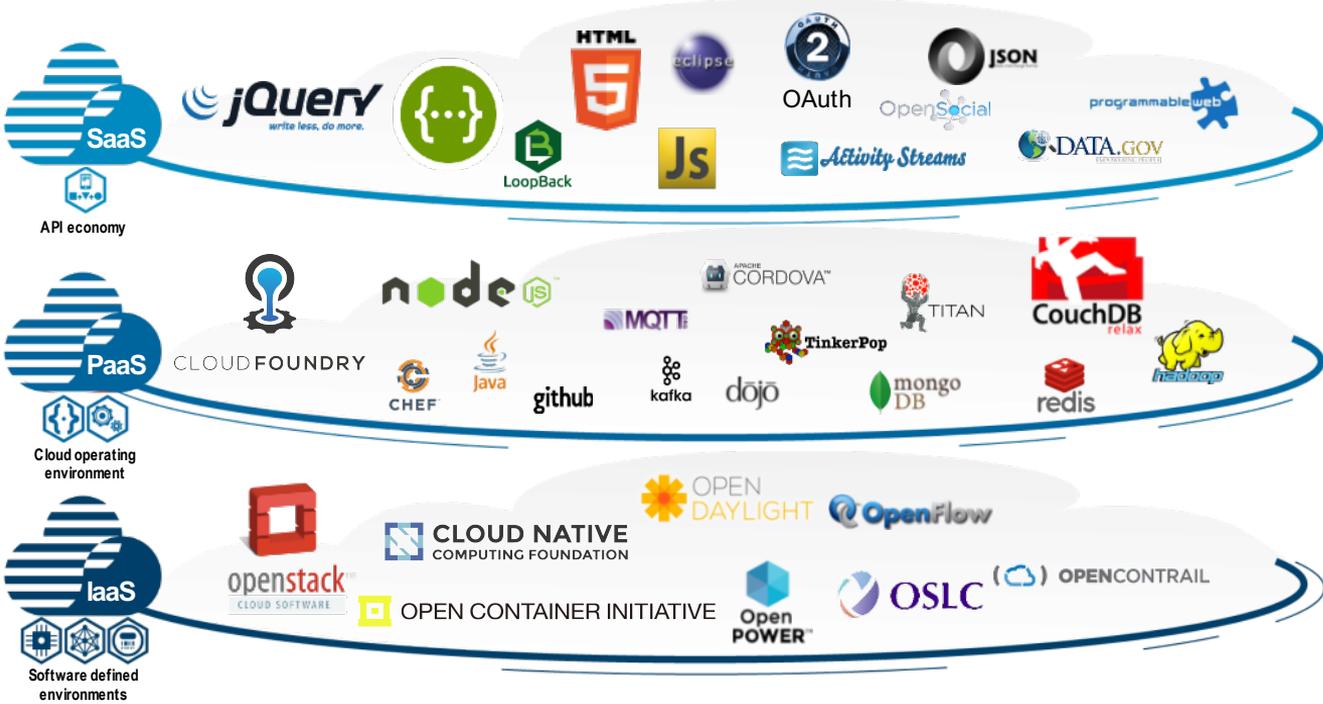
Cloud technology is doing the same thing – but faster.



Open technology has been delivering innovation



IBM's Cloud Architecture is **Open** *by design*™



An Open Cloud Architecture enables clouds, data & workloads to work together seamlessly



Bluemix Open Source Data services



Apache Spark

IBM



BigInsights for Apache

Hadoop

IBM



Cloudant NoSQL DB

IBM



dashDB

IBM



DataWorks

IBM



Elasticsearch by

Compose

IBM



Geospatial Analytics

IBM



IBM DB2 on Cloud

IBM



IBM Graph

IBM 베타



Insights for Twitter

IBM



Insights for Weather

IBM



MongoDB by Compose

IBM



PostgreSQL by Compose

IBM



Predictive Analytics

IBM



Redis by Compose

IBM



SQL Database

IBM



Streaming Analytics

IBM



ClearDB MySQL

Database

씨드파티



ElephantSQL

씨드파티



Namara.io Catalog

씨드파티



Redis Cloud

씨드파티

IBM has a long partnership with open technology

1999 - 2003

- IBM forms Linux Tech Center
- Leads Apache projects Xerces, Xalan, SOAP
- Starts ICU project
- Creates OSI-approved IBM Public License
- IBM OSDL founding member
- Founder of Eclipse.org & Eclipse Consortium
- Creates internal bazaar using OSS methodology
- Linux contributions to scalability (8-way+), reliability (stress testing, defect mgmt, doc)
- Leads Apache projects in Web Services
- Leads Eclipse projects GEF (editing), EMF (modeling), XSD/JML2 (XML Schema), Hyades (testing), Visual Editor, AspectJ, Equinox (OSGi bundles)

2004-2009

- Eclipse Foundation, Inc. becomes independent
- Pledged 500 patents to open source
- Starts Apache Derby database, supports Geronimo app server
- Contributions for Linux on Power, usability, security certifications
- Leads Apache projects Tuscany (SCA standard), OpenJPA, UIMA
- Contributes to Eclipse Higgins
- Partners with Zend PHP
- Accessibility code to Firefox
- IBM starts OpenAjax Alliance and joins Dojo Foundation
- IBM joins OpenOffice.org & creates ODF Toolkit Union
- IBM joins Open Health Tools, merging code from Eclipse OHF
- Contributes to Mozilla Bepin (web editor) & WebKit (browser engine)
- Lead Apache Aries (OSGi Enterprise)
- Contributes to Apache Shindig

2010-now

- Linux contributions to kvm, oVirt, & support Open Virtualization Alliance
- Supports Apache Hadoop (Big Data) – part of IBM BigInsights
- Eclipse: Orion (web-based tooling), Lyo (OSLC), Paho (M2M protocols)
- Cloud Standards Customer Council: Founding member
- Announces OpenJDK involvement
- Contributes to Apache Cordova (fka PhoneGap) (mobile app framework)
- Starts Dojo Maqetta (RIA tooling)
- Leads Apache OpenOffice, CouchDB (Cloudant)
- OpenStack: platinum sponsor of independent foundation; over 400 IBM contributors
- Increase OSS projects & visibility at JazzHub & GitHub
- Cloud Foundry: platinum sponsor of independent foundation; over 10 IBM contributors
- Open Data Platform: platinum sponsor
- Node.js: platinum sponsor of independent foundation
- Docker / Open Container Project: founding project sponsor
- Supports Apache Spark, Ambari & other data projects

More than 1000 IBM developers involved in OSS projects

IBM leads 80+ OSS projects

IBM leads & contributes to 200+ OSS projects



Cloud Standards Customer Council

Cloud Standards Customer Council

The Cloud Standards Customer Council is the key open, end-user advocacy group for cloud standards, interoperability & best practices.

IBM is a founding member of the CSCC & worked closely with OMG to establish a strong end-user presence. Today, IBM continues to be an active member of the Steering Committee & a regular contributor

- ### Top Initiatives
- CSCC adoption of application-specific reference architectures
 - Documentation of best practices around getting started with PaaS

500+ Members



Participating in the CSCC

40k Whitepaper



Downloads in 2015

4k Registrants



Attended 2014-15 Webinars

6 Papers Published in 2015



- Security for Cloud Computing V2 published March, 2015
- Practical Guide to Cloud Service Agreements V2 published April, 2015
- Practical Guide to PaaS targeted Sept., 2015
- Practical Guide to Hybrid Computing targeted 4Q15

- Cloud Reference Architectures
 - Web App Hosting published Jan., 2015
 - Mobile published May, 2015
 - Big Data Analytics published July, 2015
 - Additional architectures in 2H15 (Security, Hybrid, DevOps, Social)



OpenStack



Developed by the premier cloud community, OpenStack is an open source cloud operating system that is able to control large pools of compute, storage & networking resources throughout a datacenter

IBM is a founding member of the OpenStack foundation, and played an integral role in its establishment. Additionally, with over 400 developers contributing, IBM is a top contributor to the project

- Top Initiatives
- Improvements to security, federation & audit support
 - Swift & GPFS encryption
 - Extend hybrid storage functionality of Cinder
 - Trove usability with DB2 & CouchDB



OpenStack:
by the
numbers



In the
past
30 days



Since
the last
commit



Docker & The Open Container Initiative



Docker is an open platform to build, deploy & run applications using portable containers. June 22nd, The Open Container Initiative launched to ensure independent open governance of the community

IBM is a top contributor to the Docker open project, representing nearly 10 percent of all contributions. IBM is a founding member of the Open Container Initiative launched June 2015.

- Top Initiatives
- Grow the Open Container Project ecosystem
 - Harmonize container community around a single defacto standard

1,103 
Total contributors

150k 
Lines of code

Very High 
Project Activity

**Docker:
by the
numbers**

30 
New contributors

In the past 30 days

5 Days 
Since the last commit



Cloud Foundry



A robust development platform for the people, Cloud Foundry is a powerful & flexible PaaS cloud allowing coders to build & deploy apps on their own terms – using the languages & the tools they want

IBM is a founding member of the Cloud Foundry Foundation, and played an integral role in its establishment. IBM is a key contributor with 10 dedicated committers & code contributions second only to Pivotal

- Top Initiatives
- Growth of the CF foundation & the brand
 - Increase IBM contrib's & the number of dedicated contributors
 - Launch the IBM hosted CF Dojo in TRP
 - "Diego" into production



Cloud Foundry:
by the
numbers



In the
past
30 days



Node.js



Node.js is a JS runtime platform for easily building fast, scalable network apps. On June 16th, The Node.js Foundation to ensure open governance of the project & to reunite the forked codebase

IBM is a founding member of the Node.js Foundation, and played an integral role in its establishment. Currently, IBM is focused on the business & technical aspects of reunifying the project

Top Initiatives

- Reunify Node.js forks w/the original codebase & normalize community relations & process

726



Total contributors

2.0m



Lines of code

High



Project Activity

**Node.js:
by the
numbers**

5

New contributors



In the past 30 days

4

Days



Since the last commit



Apache CouchDB



CouchDB is an Apache project to build a robust, scalable, noSQL DB that embraces the web by using basic web technologies as its interface.

Cloudant is a leading contributor to the project & several of its developers sit on the Project Management Committee.

- Top Initiatives
- Horizontally scalable database clustering
 - Declarative query language
 - User & Admin Dashboards
 - Core storage engine enhancements

296



Total contributors

420k



Lines of code



Very High

Project Activity

CouchDB:
by the
numbers

1
New
contributors



In the
past
30 days

6
Days



Since
the last
commit



Apache Spark



Apache Spark is an open source cluster computing framework with in-memory primitives to significantly outperform Hadoop deployments

IBM has significant interests in Spark and has pledged 3.5k researchers to work related projects & education for 1m data scientists. Additionally, IBM plans to donate its SystemML machine learning to Spark.

- Top Initiatives
- Complete the donation of SystemML as OSS
 - Spark as a Service on Bluemix
 - Establish the Spark Technology Center

793



Total contributors

564k



Lines of code



Very High

Project Activity

Spark:
by the
numbers

53

New contributors



In the past 30 days

1 Days



Since the last commit



ODPi Foundation



The ODPi Foundation is a shared industry effort focused on promoting and advancing the state of Apache Hadoop and Big Data technologies for the enterprise

IBM is a founding member of the ODPi Foundation & is heavily involved in its establishment.

- Top Initiatives
- Create certifications for ODP compatible distributions
 - Create guidelines for ODP isvs / consumers
 - Introduce more big data projects into ODP

24
Member
companies



Up from 15 in
February

150
Bugbash
participants



Resolved
186 Hadoop
Issues

4 Hadoop
Platforms
available



- IBM Open Platform
- Infosys Information Platform
- Hortonworks Data Platform
- Pivotal HD

Based on a vision of a common ODP core

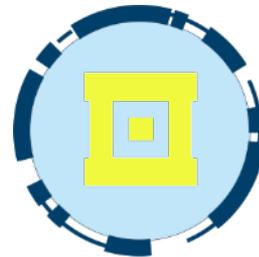


Emerging initiatives

Cloud native applications group is an effort to standardize container interfaces to connect runtimes with management & automation resources



The Open Container Project is a lightweight, open governance structure for that aims to open industry standards around container formats and runtime.



Swagger is an open, community driven framework to describe, produce & visualize RESTful APIs

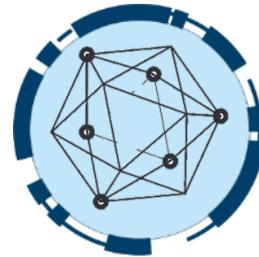


Next step of Open Technology

OpenWhisk, a distributed compute service for developing and executing application logic in response to events.



The Hyperledger Project is a collaborative effort created to advance blockchain technology for a cross-industry open standard.



IBM provides open sources at GitHub.com, GitHub.com and GitHub Enterprise is integrated with Bluemix for open development and adoption.



OpenWhisk

Welcome to the OpenWhisk project. This project contains the source code behind IBM Bluemix OpenWhisk

Have you heard?

Announcing IBM Bluemix OpenWhisk

Read More

Experimental now available!

Try out the new OpenWhisk service today on IBM Bluemix.

Try Now

High level architecture

OpenWhisk provides a distributed compute service to execute application logic in response to events. There are several key architectural concepts:

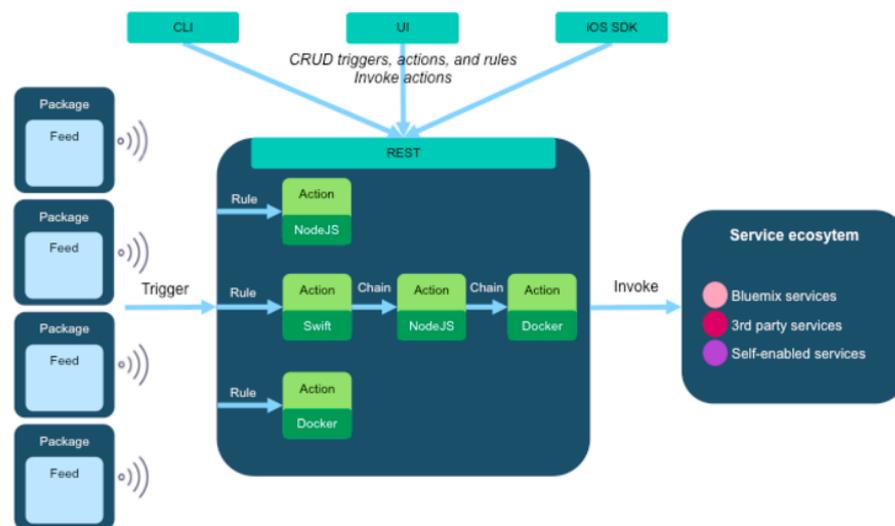
Triggers: A class of events emitted by event sources.

Actions: Encapsulate the actual code to be executed which support multiple language bindings including NodeJS, Swift and arbitrary binary programs encapsulated in Docker Containers. Actions invoke any part of an open ecosystem including existing Bluemix services for analytics, data, cognitive, or any other 3rd party service.

Rules: An association between a trigger and an action.

Packages: Describe external services in a uniform manner.

Combined these allow developers to compose solutions using modern abstraction and chaining which can be created, accessed, updated and deleted via the CLI. For convenience, special SDKs like our iOS SDK can be used for specific environments.





WHAT IS THE HYPERLEDGER PROJECT?

The Hyperledger Project is a collaborative effort created to advance blockchain technology by identifying and addressing important features for a cross-industry open standard for distributed ledgers that can transform the way business transactions are conducted globally.

WHY CREATE THE PROJECT?

Blockchain is a peer-to-peer distributed ledger technology for a new generation of transactional applications that establishes trust, accountability and transparency while streamlining business processes. Think of it as an operating system for interactions. It has the potential to vastly reduce the cost and complexity of getting things done.

<https://www.hyperledger.org/>



IBM Offerings Supporting Hyperledger

LINUX FOUNDATION
HYPERLEDGER
PROJECT

IBM Blockchain
ON IBM CLOUD

IBM Blockchain
SOLUTIONS

BLUEMIX SERVICE



Blockchain

Managed Service on IBM Cloud

Your private Blockchain network in 1-click

Learn with sample applications

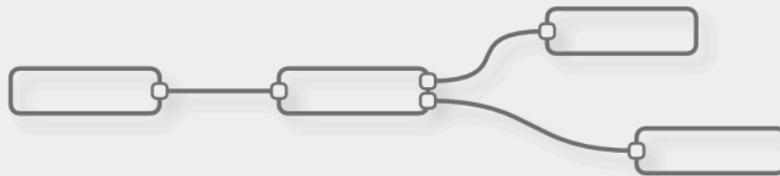
Develop your own Smart Contracts

<http://www.ibm.com/blockchain/>



Node-RED

A visual tool for wiring the Internet of Things



Node-RED is a tool for wiring together hardware devices, APIs and online services in new and interesting ways.

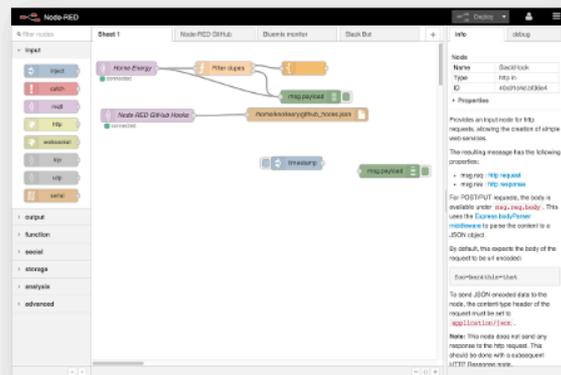
Getting Started

For Linux/OS X, if you already have Node.js installed, run:

```
$ sudo npm install -g node-red
$ node-red
```

Otherwise, head over to the [Getting Started guide](#).

Already got Node-RED installed and wonder where the download button has gone? Head over to the [Upgrading guide](#).



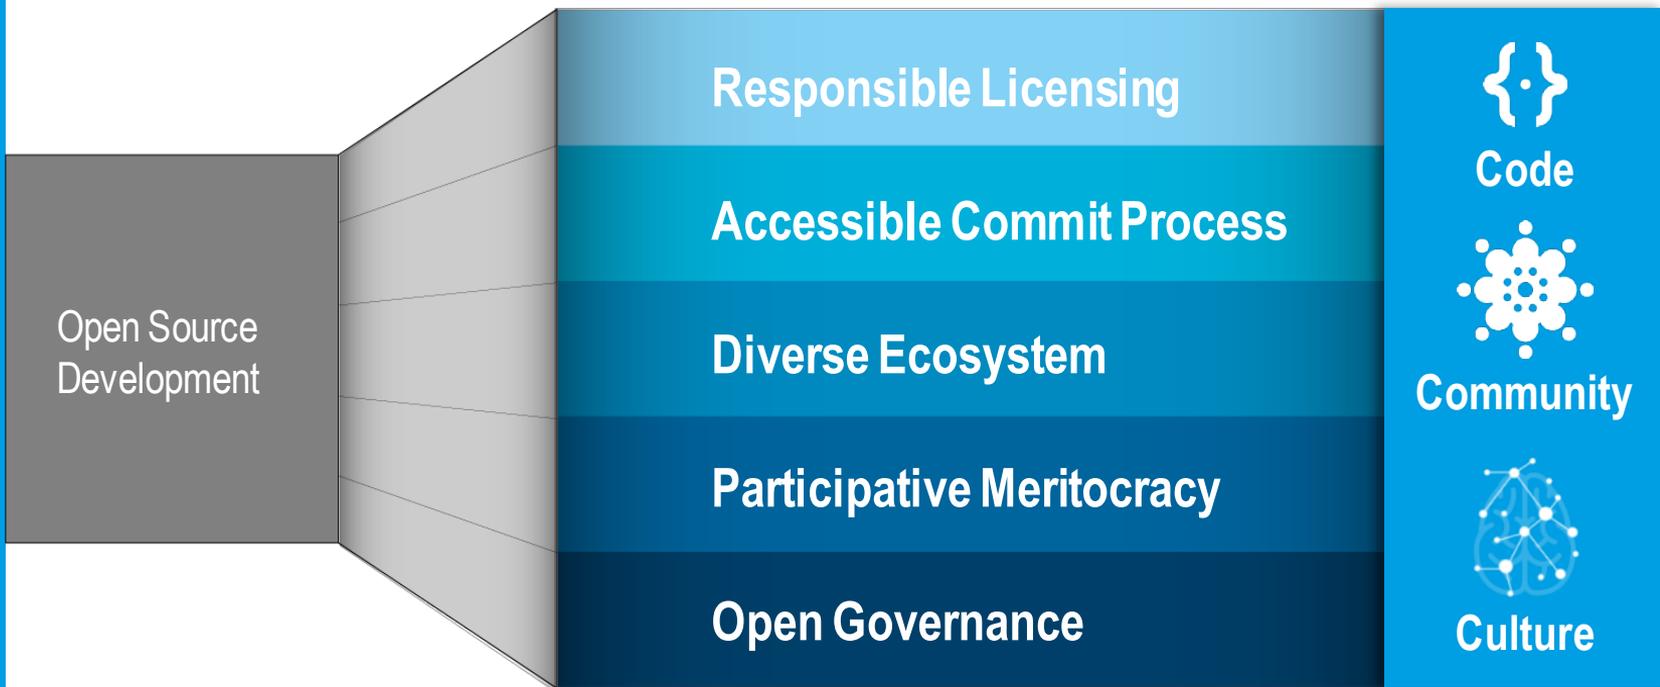
Latest version: [v0.13.4 \(npm\)](#)





Not all open tech is built the same

Unleash innovation with open & transparent communities





Open *by design*™

Open Tech in the Cloud

IBM's Cloud Offerings deliver faster development, interoperability with a vast ecosystem of compatible services, and fuel innovation, allowing customers of all sizes to engage in global markets in new and exciting ways.



CLOUDFOUNDRY



openstack™



OPEN CONTAINER INITIATIVE



CLOUD NATIVE COMPUTING FOUNDATION



APACHE



CouchDB
relax



HTML



jQuery
write less, do more



<https://developer.ibm.com/opentech/>

“IBM firmly believes that open source is the foundation of innovative application development in the cloud

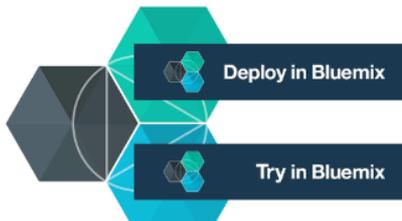


Open source from IBM

Let's work together - developerWorks Open is where today's cutting edge developers have a unique opportunity to collaborate and create tomorrow's big innovations. [READ MORE](#)

Many developerWorks Open projects are now integrated with Bluemix, IBM's cloud-based platform for building, managing, and running applications. Look for the Bluemix buttons in the "How to contribute" sections of projects like these:

<https://developer.ibm.com/open/>



IBM Open Source at GitHub

Enhance your GitHub experience with
Bluemix DevOps Services »

Deploy your application to
Bluemix cloud platform »

Engage with IBM Open Source
developerWorks Open »

SEARCH

RECENT

[java-sdk](#)

Java

:mortar_board: Java SDK to use the IBM Watson services. It supports Android.

[sample.batch.joblogevents](#)

An MDB application that demos the job log as an event functionality of the batch product. It uses the job log event messages to create a local job log

[streamsx.waterConservation.st...](#)

JavaScript

Starter kit for smart and connected sprinkler system using Apache Quarks, Streaming Analytics and Insights for Weather

[unity-sdk](#)

C#

:construction: Unity SDK to use the IBM Watson services.

REPOS(531)

[.NET-REST-Tool](#) (softlayer)
C#

A handy .NET-based REST troubleshooting tool with C# and vb.NET source available.

[a11y.js](#) (IBM-Watson)
JavaScript

A set of JS helpers for working with accessibility attributes

[acmeair](#) (acmeair)
Java

A Java implementation of the Acme Air Sample Application.

[acmeair-driver](#) (acmeair)
Java

A workload driver for the Acme Air Sample Application.

[acmeair-nodejs](#) (acmeair)
JavaScript

A Node.js implementation of the Acme Air Sample Application. With datastore support of MongoDB,

[acmeair-webcontent](#) (acmeair)

[acmeair.github.io](#) (acmeair)
CSS

[activitystreams](#) (OpenSocial)
Java

Activity Streams 2.0 Java Reference Implementation

<http://ibm.github.io/>





A New Way Forward

Type of Open Source

By community



By foundation



Linux



By corporation



Open

Open Source

Open Standards

Open Communities