

# PinotProposal

## Pinot Proposal

### Abstract

Pinot is a distributed columnar storage engine that can ingest data in real-time and serve analytical queries at low latency. There are two modes of data ingestion - batch and/or realtime. Batch mode allows users to generate pinot segments externally using systems such as Hadoop. These segments can be uploaded into Pinot via simple curl calls. Pinot can ingest data in near real-time from streaming sources such as Kafka. Data ingested into Pinot is stored in a columnar format. Pinot provides a SQL like interface (PQL) that supports filters, aggregations, and group by operations. It does not support joins by design, in order to guarantee predictable latency. It leverages other Apache projects such as Zookeeper, Kafka, and Helix, along with many libraries from the ASF.

### Proposal

Pinot was open sourced by [LinkedIn](#) and hosted on [GitHub](#). Majority of the development happens at [LinkedIn](#) with other contributions from Uber and Slack. We believe that being a part of Apache Software Foundation will improve the diversity and help form a strong community around the project.

[LinkedIn](#) submits this proposal to donate the code base to Apache Software Foundation. The code is already under Apache License 2.0. Code and the documentation are hosted on Github.

- Code: <http://github.com/linkedin/pinot>
- Documentation: <https://github.com/linkedin/pinot/wiki>

### Background

[LinkedIn](#), similar to other companies, has many applications that provide rich real-time insights to members and customers (internal and external). The workload characteristics for these applications vary a lot. Some internal applications simply need ad-hoc query capabilities with sub-second to multiple seconds latency. But external site facing applications require strong SLA even very high workloads. Prior to Pinot, [LinkedIn](#) had multiple solutions depending on the workload generated by the application and this was inefficient. Pinot was developed to be the one single platform that addresses all classes of applications. Today at [LinkedIn](#), Pinot powers more than 50 site facing products with workload ranging from few queries per second to 1000's of queries per second while maintaining the 99th percentile latency which can be as low as few milliseconds. All internal dashboards at [LinkedIn](#) are powered by Pinot.

### Rationale

We believe that requirement to develop rich real-time analytic applications is applicable to other organizations. Both Pinot and the interested communities would benefit from this work being openly available.

### Current Status

Pinot is currently open sourced under the Apache License Version 2.0 and available at [github.com/linkedin/pinot](http://github.com/linkedin/pinot). All the development is done using [GitHub](#) Pull Requests. We cut releases on a weekly basis and deploy it at [LinkedIn](#). mp-0.1.468 is the latest release tag that is deployed in production.

### Meritocracy

Following the Apache meritocracy model, we intend to build an open and diverse community around Pinot. We will encourage the community to contribute to discussion and codebase.

### Community

Pinot is currently used extensively at [LinkedIn](#) and Uber. Several companies have expressed interest in the project. We hope to extend the contributor base significantly by bringing Pinot into Apache.

### Core Developers

Pinot was started by engineers at [LinkedIn](#), and now has committers from Uber.

### Alignment

Apache is the most natural home for taking Pinot forward. Pinot leverages several existing Apache Projects such as Kafka, Helix, Zookeeper, and Avro. As Pinot gains adoption, we plan to add support for the ORC and Parquet formats, as well as adding integration with Yarn and Mesos.

## Known Risks

### Orphaned Products

The risk of the Pinot project being abandoned is minimal. The teams at [LinkedIn](#) and Uber are highly incentivized to continue development of Pinot as it is a critical part of their infrastructure.

### Inexperience with Open Source

Post open sourcing, Pinot was completely developed on [GitHub](#). All the current developers on Pinot are well aware of the open source development process. However, most of the developers are new to the Apache process. Kishore Gopalakrishna, one of the lead developers in Pinot, is VP and committer of the Apache Helix project.

### Homogenous Developers

The current core developers are all from [LinkedIn](#) and Uber. However, we hope to establish a developer community that includes contributors from several corporations and we are actively encouraging new contributors via the mailing lists and public presentations of Pinot.

### Reliance on Salaried Developers

It is expected that Pinot development will occur on both salaried time and on volunteer time, after hours. The majority of initial committers are paid by their employer to contribute to this project. However, they are all passionate about the project, and we are confident that the project will continue even if no salaried developers contribute to the project. We are committed to recruiting additional committers including non-salaried developers.

### Relationships with Other Apache Products

As mentioned earlier, Pinot uses several Apache Projects such as Kafka to ingest data in real-time, Zookeeper and Helix for cluster management. Pinot also uses Maven for build and release. We foresee adding support for the Parquet and ORC formats. Adding the ability to deploy on Yarn and Mesos clusters is another interesting project we might pursue.

### An Excessive Fascination with the Apache Brand

While we respect the reputation of the Apache brand and have no doubts that it will attract contributors and users, we believe ASF is the right home for Pinot to foster a great community that will lead to a better outcome in the long term.

## Documentation

- Code: <https://github.com/linkedin/pinot/>
- Documentation: <https://github.com/linkedin/pinot/wiki>
- User group: [https://groups.google.com/forum/#!forum/pinot\\_users](https://groups.google.com/forum/#!forum/pinot_users)

## Initial Source

The current Pinot codebase is hosted on Github and licensed under the Apache License V2. The source tree is self contained and relies on Maven as its build and dependency resolution mechanism.

## External Dependencies

All dependencies in Pinot have licenses that are compatible with Apache License V2, except for the org.json library, which will be removed prior to Apache incubation. The list below summarizes the external dependencies of Pinot grouped by license and ASF license category.

Dependencies from the ASF Category A

### Apache License 2.0

- com.101tec:zkclient:0.7
- com.alibaba:fastjson:1.1.24
- com.clearspring.analytics:stream:2.7.0
- com.fasterxml.jackson.core:jackson-annotations:2.8.0
- com.fasterxml.jackson.core:jackson-core:2.8.0
- com.fasterxml.jackson.core:jackson-databind:2.8.0
- com.google.code.findbugs:jsr305:3.0.0
- com.google.guava:guava:19
- com.ning:async-http-client:1.9.21
- com.yammer.metrics:metrics-core:2.2.0
- commons-beanutils:commons-beanutils:1.8.3
- commons-cli:commons-cli:1.2
- commons-codec:commons-codec:1.6
- commons-configuration:commons-configuration:1.6
- commons-fileupload:commons-fileupload:1.2.2

- commons-httpclient:commons-httpclient:3.1
- commons-io:commons-io:2.1
- commons-validator:commons-validator:1.4.0
- io.netty:netty-all:4.1.4.Final
- io.swagger:swagger-jaxrs:1.5.10
- io.swagger:swagger-jersey2-jaxrs:1.5.10
- it.unimi.dsi.fastutil:6.5.16
- joda-time:joda-time:2
- log4j:log4j:1.2.17
- me.lemire.integercompression:JavaFastPFOR:0.0.13
- nl.jqno.equalsverifier:equalsverifier:1.7.2
- org.apache.avro:avro:1.7.6
- org.apache.commons:commons-compress:1.9
- org.apache.commons:commons-lang3:3.5
- org.apache.commons:commons-math:2.1
- org.apache.hadoop:hadoop-client:2.7.0
- org.apache.hadoop:hadoop-common:2.7.0
- org.apache.helix:helix-core:0.6.8
- org.apache.httpcomponents:httpclient:4.1.3
- org.apache.httpcomponents:httpclient:4.2.5
- org.apache.httpcomponents:httpcore:4.2.5
- org.apache.httpcomponents:httpmime:4.2.5
- org.apache.kafka:kafka\_2.10:0.9.0.1
- org.apache.thrift:libthrift:0.9.1
- org.apache.zookeeper:zookeeper:3.4.9
- org.codehaus.jackson:jackson-core-asl:1.9.6
- org.codehaus.jackson:jackson-mapper-asl:1.9.6
- org.json:json:20080701
- org.roaringbitmap:RoaringBitmap:0.5.10
- org.testng:testng:6.0.1
- org.twitter4j:twitter4j-core:4.0.3
- org.webjars:swagger-ui:2.2.2
- org.xerial.larray:larray:0.2.1
- org.yaml:snakeyaml:1.16
- xml-apis:xml-apis:1.0.b2

## Dual license (Apache License 2.0 + LGPL 2.1), using under the Apache License

- org.codehaus.jackson:jackson-jaxrs:1.9.6
- org.codehaus.jackson:jackson-xc:1.9.6

## BSD

- com.jcabi:jcabi-log:0.17.1
- org.antlr:antlr4-annotations:4.3
- org.antlr:antlr4-runtime:4.3

## MIT

- com.github.nkzawa:socket.io-client:0.5.1
- org.mockito:mockito-core:2.10.0
- org.slf4j:slf4j-api:1.7.7
- org.slf4j:slf4j-log4j12:1.7.7

## Dependencies from the ASF Category B

Dual license (CDDL 1.1 + GPL 2 w/ CPE), using under the CDDL

- com.sun.jersey:jersey-client:1.19.2
- javax.servlet:javax.servlet-api:3.0.1
- org.glassfish.jersey.containers:jersey-container-grizzly2-http:2.23
- org.glassfish.jersey.core:jersey-common:2.23
- org.glassfish.jersey.core:jersey-server:2.23
- org.glassfish.jersey.media:jersey-media-json-jackson:2.24
- org.glassfish.jersey.media:jersey-media-multipart:2.23

## Dependencies from the ASF Category X

JSON License

- org.json:json:20080701 (to be removed before Apache incubation)

## Cryptography

None

## Required Resources

### Mailing lists

- pinot-private (with moderated subscriptions)
- pinot-user
- pinot-dev
- pinot-commits

### Git repository

- [git://git.apache.org/pinot](https://git.apache.org/pinot)
- <https://git-wip-us.apache.org/repos/asf/incubator-pinot.git>

### Issue Tracking

A JIRA Issue tracker (PINOT)

### Other Resources

The existing code already has unit and integration tests and we use travis to test the patch before committing it to master. We would like to have an instance of Jenkins to achieve similar functionality.

### Initial Committers

- Kishore Gopalakrishna
- Ravi Aringunram
- Jean-François Im
- Mayank Shrivastava
- Subbu Subramaniam
- Adwait Tumbde
- Xiaotian Jiang
- Jennifer Dai
- Seunghyun Lee
- Xiang Fu
- Dhaval Patel
- Neha Pawar
- Alex Pucher
- Jihao Zhang
- Akshay Rai
- Sunitha Beeram
- Jialiang Jack Li
- James Shao
- Marcel Siegrist

### Affiliations

- Kishore Gopalakrishna (LinkedIn)
- Ravi Aringunram (LinkedIn)
- Jean-François Im (LinkedIn)
- Mayank Shrivastava (LinkedIn)
- Subbu Subramaniam (LinkedIn)
- Adwait Tumbde (LinkedIn)
- Xiaotian Jiang (LinkedIn)
- Jennifer Dai (LinkedIn)
- Seunghyun Lee (LinkedIn)
- Xiang Fu (Uber)
- Dhaval Patel (Uber)
- James Shao (Uber)
- Chinmay Soman (Uber)
- Neha Pawar (LinkedIn)
- Alex Pucher (LinkedIn)
- Marcel Siegrist
- Akshay Rai (LinkedIn)
- Jihao Zhang (LinkedIn)
- Jialian Jack Li (LinkedIn)

## Sponsors

### Champion

- Olivier Lamy <olamy at apache dot org>

### Nominated Mentors

- Olivier Lamy <olamy at apache dot org>
- Kishore Gopalakrishna <kishoreg at apache dot org>
- Jim Jagielski <jim at apache dot org>
- Roman Shaposhnik <rsv at apache dot org>

### Sponsoring Entity

The Apache Incubator