

TinkerPopProposal

Apache Tinkerpop Proposal

<http://www.tinkerpop.com/docs/3.0.0-SNAPSHOT/images/tinkerpop-logo.png>

A. Abstract

TinkerPop is a graph computing framework written in Java. A **graph** is a data structure composed of vertices and edges and is useful for modeling complex domains with arbitrary relations (edges, links, lines) between entities (vertices, objects, dots). **TinkerPop** provides a core API that graph system vendors can implement. There are various types of graph systems including in-memory graph libraries, OLTP graph databases, and OLAP graph processors (see [Open Graph Computing](#) for more information). Once the core interfaces are implemented, the underlying graph system can be queried using the graph traversal language Gremlin and processed with **TinkerPop**-enabled algorithms. For many, **TinkerPop** is seen as the **JDBC** of the graph computing community.

B. Proposal

TinkerPop was formed in 2009 and is currently in the milestone series of 3.0.0. From the start, **TinkerPop** has provided its software open source and free to use for which ever reason (commercial or otherwise). Initially the license was BSD, but as of **TinkerPop3**, the license was changed to Apache2. The **TinkerPop** team is composed of developers, evangelists, and representatives from graph system vendors (see [TinkerPop Contributors](#) for more information). **TinkerPop** has done its best to remain vendor agnostic and works closely with all vendors to ensure that the constructs within **TinkerPop** are able to accommodate the requirements of the underlying graph system. To date, 12 "TinkerPop recognized" graph system vendors provide **TinkerPop** implementations. We believe that by joining The Apache Software Foundation, our vendors, users, and contributors will feel more comfortable in terms of legal protected, in terms of wider-adoption, and in terms of project stability.

C. Background

TinkerPop has had steady, active development since 2009 when it was founded. Over the years, the Gremlin query language within **TinkerPop** has been adopted by various JVM languages and as such, there exists Gremlin-Groovy, Gremlin-Scala, Gremlin-Clojure, Gremlin-JavaScript, and the like. In many ways, Gremlin is seen as a traversal style that can be readily adopted within the programming constructs of the developer's native language — both on and off the JVM. **TinkerPop** is not bound to the JVM in that developers wishing to interact with a **TinkerPop**-enabled graph system can leverage Gremlin Server which provides "over the wire" communication as well as the entry point for non-JVM language bindings. **TinkerPop** is being used in production graph-based applications around the world and is only getting better with age.

D. Rationale

The graph computing space has grown over the years to encompass numerous graph database and graph processing systems. **TinkerPop** was created as a unifying framework for interoperability, language standardization, and data model standardization. This framework makes it simple to "plug and play" the back-end graph implementation without affecting the developer's code. This is analogous to the way in which the JDBC allows users to swap relational databases while keeping the same programming interface. **TinkerPop** also brings together OLTP systems (graph databases) and OLAP systems (graph processors) by providing a single query language (Gremlin) for executing graph algorithms transparently over either type of system. The seamless support of single-machine systems and distributed systems provides an identical end-user experience within the boundaries of the space and time constraints of the underlying graph system.

E. Initial Goals

The goal of this proposal is to migrate **TinkerPop** to The Apache Software Foundation to

1. get more exposure to the project
2. legally protect the developers and contributors of the project
3. ensure that **TinkerPop** has a vendor-neutral, stable home for the future

F. Current Status

TinkerPop is planning to release **TinkerPop** 3.0.0.GA early 2015. We would like to release **TinkerPop** 3.0.0.GA at the Apache Software Foundation.

G. Meritocracy

TinkerPop already has an established community that values meritocracy. Anyone is welcome to contribute to **TinkerPop** and voice their opinions/thoughts on the direction of the project. If an individual provides beneficial pull requests, documentation updates, theoretical ideas, etc. then they are asked to join **TinkerPop** as committers. Moreover, **TinkerPop** is a business friendly community. It provides a forum where vendors collaborate, via individual representatives, by contributing ideas, code and content.

We believe the values of the community align well with the ASF, and we are willing to adapt our practices when necessary.

H. Community

The [TinkerPop mailing list](#) has 1500+ subscribers with a flow of approximately 10 emails a day. In addition to the initial committers there are about two dozen folks that have contributed code to the project, and 1/2 of those are from Graph vendors. Note that there are individuals that participate in the [TinkerPop](#) ecosystem by providing language bindings and system implementations. While not contributing to the core codebase per se, they support [TinkerPop](#) via their software and lecture outreach.

I. Core developers

The initial committer list is small, and we hope to grow that into a larger, vibrant community at the Apache Software Foundation.

J. Alignment

[TinkerPop](#) is licensed Apache2. The OLAP reference implementation that distributes with [TinkerPop](#) is based on [Hadoop](#) and [Giraph](#). There has been a push from vendors and commercial organizations interested in ensuring [TinkerPop](#) has the legal support of a software foundation. We believe The Apache Software Foundation to be a great home for [TinkerPop](#).

K. Known Risks & Abandoned projects

[TinkerPop](#) is currently at [TinkerPop 3.0.0.M6](#). There is a significant amount of code developed since [TinkerPop's](#) inception. We have done our best to secure CLAs from contributors. To date did not have any legal or political issues. [TinkerPop](#) has 12 official graph vendors that maintain or support a [TinkerPop](#) implementation for their system. If, lets say, 75% of those vendors decide against providing a [TinkerPop](#) interface, then this would be detrimental to [TinkerPop's](#) efforts. With that, it is very important that [TinkerPop](#) remain neutral and engaged with vendors.

L. Inexperience with Open Source

The primary developers of [TinkerPop](#) have been producing [TinkerPop](#) open source code and releases in full public view since 2009 via both our [GitHub](#) repository and our public mailing lists. We pride ourselves in professional software, documentation, and aesthetics. With its integration to numerous graph systems in the market, we see [TinkerPop](#) as a masterpiece in the open source space. We are working with both open source and commercial entities to ensure a solid, respected product.

M. Homogeneous Developers

The initial set of committers all hail from different employers.

N. Reliance on Salaried Developers

[TinkerPop](#) is open source and no individual is paid for their contributions except for our logo designer. All the logos and documentation graphics are released Apache2.

O. Relationships with Other Apache Projects

[TinkerPop](#) currently leverages Hadoop, Giraph, and in the near future, will integrate with [Spark](#). Internally, the project leverages Maven3, Apache Configuration, and Apache Commons.

P. An Excessive Fascination with the Apache Brand

[TinkerPop](#) is interested in The Apache Software Foundation for the legal support and protection it can offer our developers. Numerous large software organizations have urged [TinkerPop](#) to join a foundation so that they may more easily provide contributors from their organization and moreover, maintain confidence in building products atop [TinkerPop](#). With that said, [TinkerPop](#) does respect the Apache Software Foundation brand and believes that by joining The Apache Software Foundation, we will further boost the recognition and prestige of our product. Finally, note that a [TinkerPop book](#) is slated for 2015. Combining the Apache Software Foundation and a long awaited book should ensure a new uptake of contributors, users, vendors, and promoters.

Q. Documentation

[TinkerPop](#) takes pride in its documentation:

AsciiDoc	http://www.tinkerpop.com/docs/3.0.0-SNAPSHOT/
JavaDoc (core)	http://www.tinkerpop.com/javadocs/3.0.0-SNAPSHOT/core/
JavaDoc (full)	http://www.tinkerpop.com/javadocs/3.0.0-SNAPSHOT/full/

R. Initial Source

[TinkerPop](#) is currently hosted on [GitHub](#): <https://github.com/tinkerpop/>.

The following repositories would like to be migrated to ASF.

TinkerPop3	https://github.com/tinkerpop/tinkerpop3
Blueprints (TinkerPop2)	https://github.com/tinkerpop/blueprints
Pipes (TinkerPop2)	https://github.com/tinkerpop/pipes
Frames (TinkerPop2)	https://github.com/tinkerpop/frames
Gremlin (TinkerPop2)	https://github.com/tinkerpop/gremlin
Rexster (TinkerPop2)	https://github.com/tinkerpop/rexster

S. Source & Intellectual Property Submission Plan

TinkerPop has required CLAs from contributors in the past to ensure solid IP provenance. TinkerPop plans to submit a Software Grant for the content in the following repositories: <https://github.com/tinkerpop/tinkerpop3>

We plan to transfer to the ASF the TinkerPop trademark as well as the commissioned artwork for TinkerPop logos and the <http://tinkerpop.com> and <http://tinkerpop.org> domains.

T. External Dependencies

TinkerPop relies on a number of open source dependencies. We believe these to comply with the third-party licensing guidelines. If any dependencies are unfit, we will make the necessary changes.

U. Cryptography

None.

V. Required Resources

W. Mailing Lists

TinkerPop requests the following mailing list:

User Mailing List	users@tinkerpop.incubator.apache.org
Developer Mailing List	dev@tinkerpop.incubator.apache.org
Private Mailing List	private@tinkerpop.incubator.apache.org
Commits Mailing List	commits@tinkerpop.incubator.apache.org

X. Subversion Directory

TinkerPop will use subversion for it's website and requests:

- <https://svn.apache.org/repos/asf/incubator/tinkerpop>

Y. Git Repository

TinkerPop wishes to have the following created:

- <https://git-wip-us.apache.org/repos/asf/incubator-tinkerpop.git>

And to have all of these mirrored to Github with integration taken care of.

Z. Issue Tracking

Please help us setup a JIRA instance for both issue tracking and code review (TINKERPOP).

AA. Initial Committers and Affiliations

Marko A. Rodriguez (2009)	Aurelius
---------------------------	----------

Stephen Mallette (2011)	Nidomics
James Thornton (2013)	Electric Speed

BB. Sponsors

- Champion: David Nalley
- Mentors: Rich Bowen, Matt Franklin, Daniel Gruno, David Nalley, Hadrian Zbarcea
- Sponsoring Entity: We request the Incubator to sponsor this project.