

Open Source at Bloomberg in 2020

Engineering

Bloomberg

Hunter College – Department of Computer Science
March 26, 2020

Kevin P. Fleming
Head of Open Source Community Engagement

TechAtBloomberg.com

© 2020 Bloomberg Finance L.P. [Creative Commons Attribution 4.0 International License](https://creativecommons.org/licenses/by/4.0/).



What does Bloomberg do?



Bloomberg Professional Service

- **Bloomberg Professional Service**
- **Trading Systems**
- **Tradebook**
- **Bloomberg Enterprise**
- **News**
- **Media**
- **Bloomberg Law**
- **Bloomberg Government**

TechAtBloomberg.com

© 2020 Bloomberg Finance L.P. [Creative Commons Attribution 4.0 International License](https://creativecommons.org/licenses/by/4.0/).



Bloomberg

Engineering

Bloomberg by the numbers

- Founded in **1981**
- Over **325,000** subscribers
- Customers in **170** countries
- Nearly **20,000** employees in **167** locations
- **6,000+** engineers and growing
- Real-time data feeds from **hundreds** of exchanges
- News and pricing from **thousands** of contributors

TechAtBloomberg.com

© 2020 Bloomberg Finance L.P. [Creative Commons Attribution 4.0 International License](https://creativecommons.org/licenses/by/4.0/).

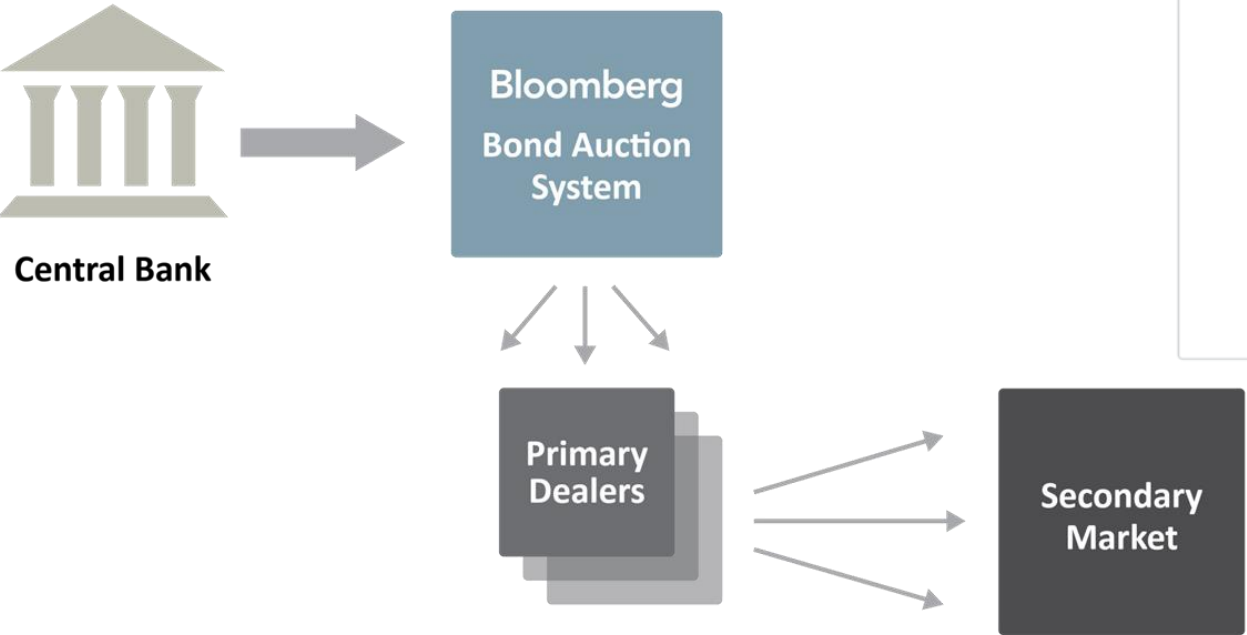


Bloomberg

Engineering



High Availability + Consistency + Low Latency + Large Scale



Philip O'Sullivan
@pdosullivan 

Bloomberg being down has finally made me understand how teenagers feel when Facebook crashes.



How do we use open source?

DevOps/Automation



GitHub



CHE

Cloud Computing



opens



kubernetes

Distributed Computing



Machine Learning



News Search



Quant Computing



heano



TechAtBloomberg.com

Bloomberg

Engineering



Let's talk about 'news search'

325K+

Terminal subscribers

News search

16M

queries per day

Stories available
for search in

~100ms

Average query
response time

<200ms

News volume

2M

stories per day

500

stories ingested
per second

650M

stories in index

News alerting

1.5M

subscriptions


500

stories matched
per second

Alerts delivered in

<100ms

And then there's market data...



TechAtBloomberg.com

© 2020 Bloomberg Finance L.P. [Creative Commons Attribution 4.0 International License](https://creativecommons.org/licenses/by/4.0/).



Bloomberg

Engineering



This leads to unique challenges for us

- For example, some infrastructure tools support clustering, replication, and load-sharing 'out of the box'
- However, the mechanisms used to detect node failures and re-issue queries might take as long as 1000ms (yes, one second)
- When the response must be provided to the user in less than 250ms, this is not acceptable
- Tradeoffs must be made to provide rapid failover, while ensuring data integrity



A small sample of Bloomberg's contributions

- JavaScript 'Private Fields' in TypeScript and Babel
- 'Learning to Rank' in Apache Solr
- Project Jupyter UX design
- 'PowerfulSeal' testing tool for Kubernetes
- 'solr-operator' – Apache Solr on Kubernetes

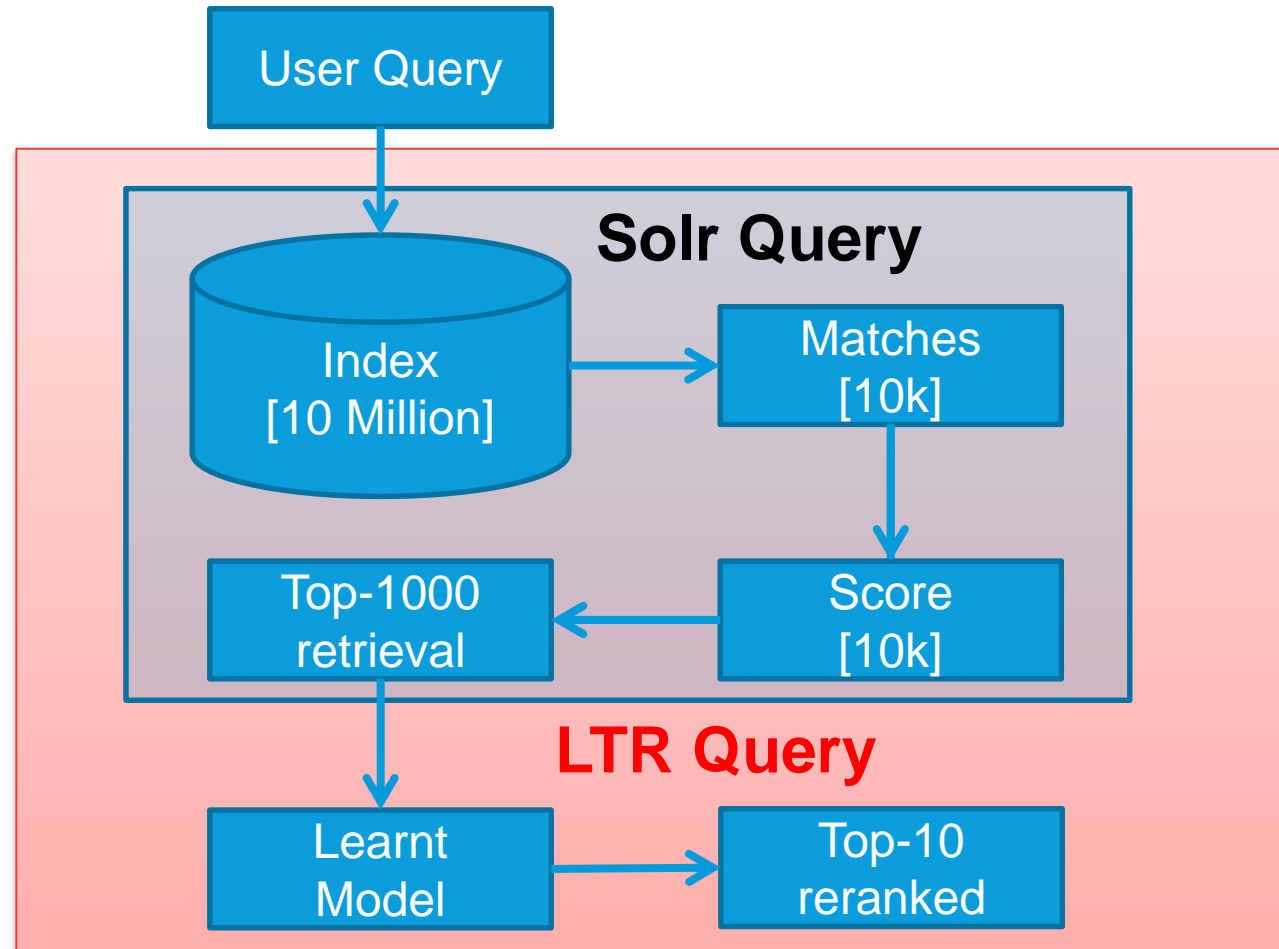
JavaScript 'Private Fields' in TypeScript

- New language feature which enables encapsulation of class data
- Currently in 'Stage 3' in ECMA TC39, awaiting implementation experience
- Bloomberg team has implemented support in TypeScript, a transpiler used throughout the JS ecosystem
- Bloomberg team has also collaborated on implementation in Babel, another transpiler

‘Learning to Rank’ in Apache Solr

- Plugin which allows execution of machine learning models inside Solr
- Executing models in Solr improves efficiency dramatically
- Models have access to entire universe of document features in the Solr index

'Learning to Rank' in Apache Solr



Project Jupyter UX Design

- Bloomberg is using Project Jupyter to power our BQNT<GO> product
- At JupyterCon 2017 and 2018, Bonnie John of Bloomberg's UX team held design sessions for JupyterLab
- Goal was to increase usability and effectiveness of the JupyterLab user interface

Project Jupyter UX Design



TechAtBloomberg.com

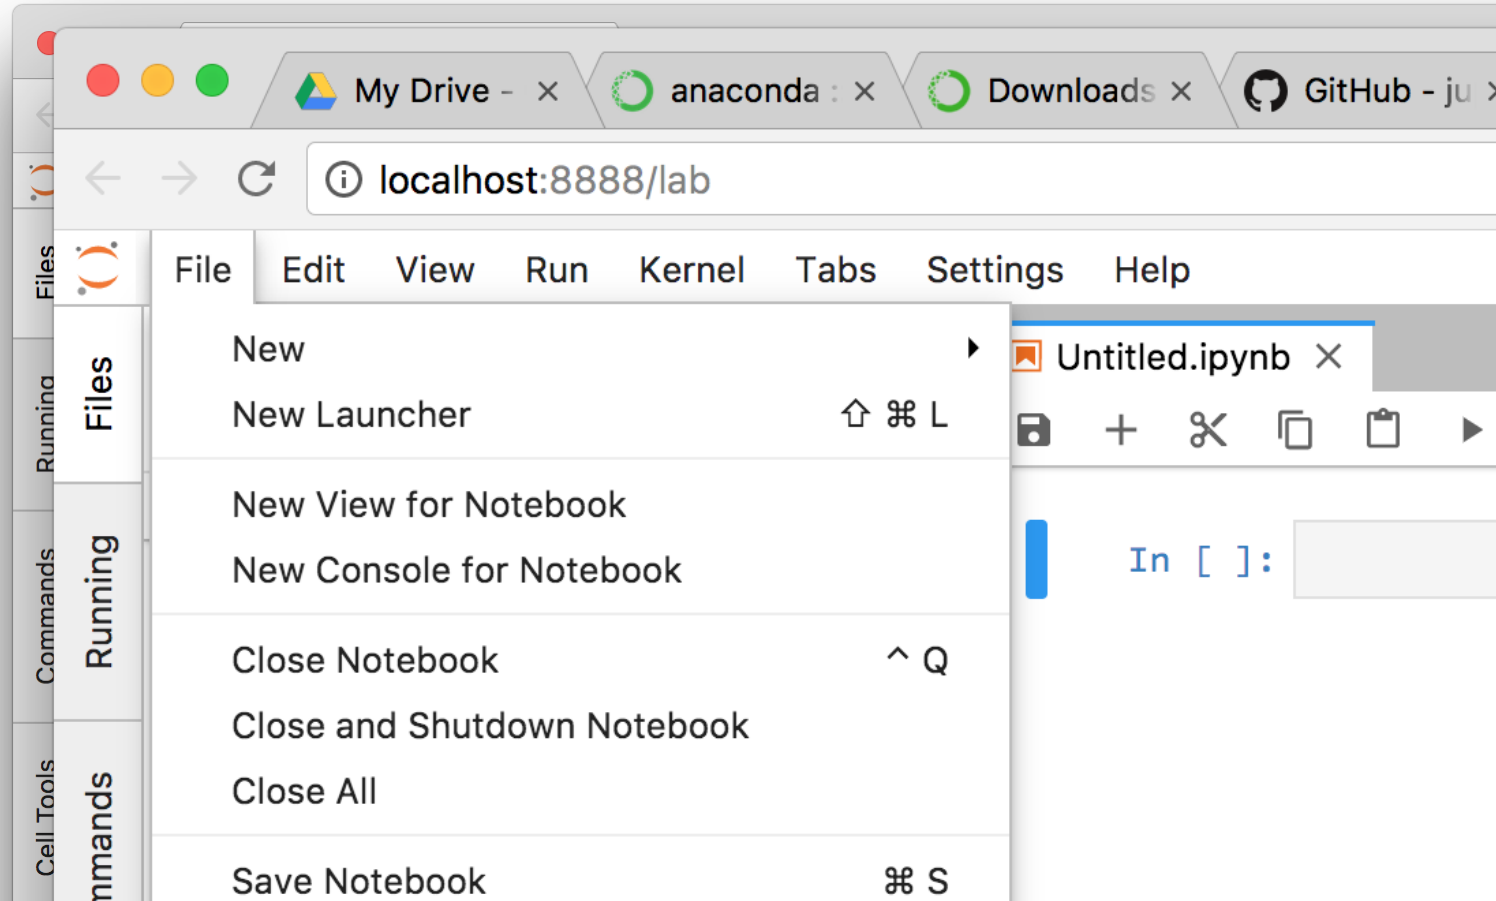
© 2020 Bloomberg Finance L.P. [Creative Commons Attribution 4.0 International License](https://creativecommons.org/licenses/by/4.0/).



Bloomberg

Engineering

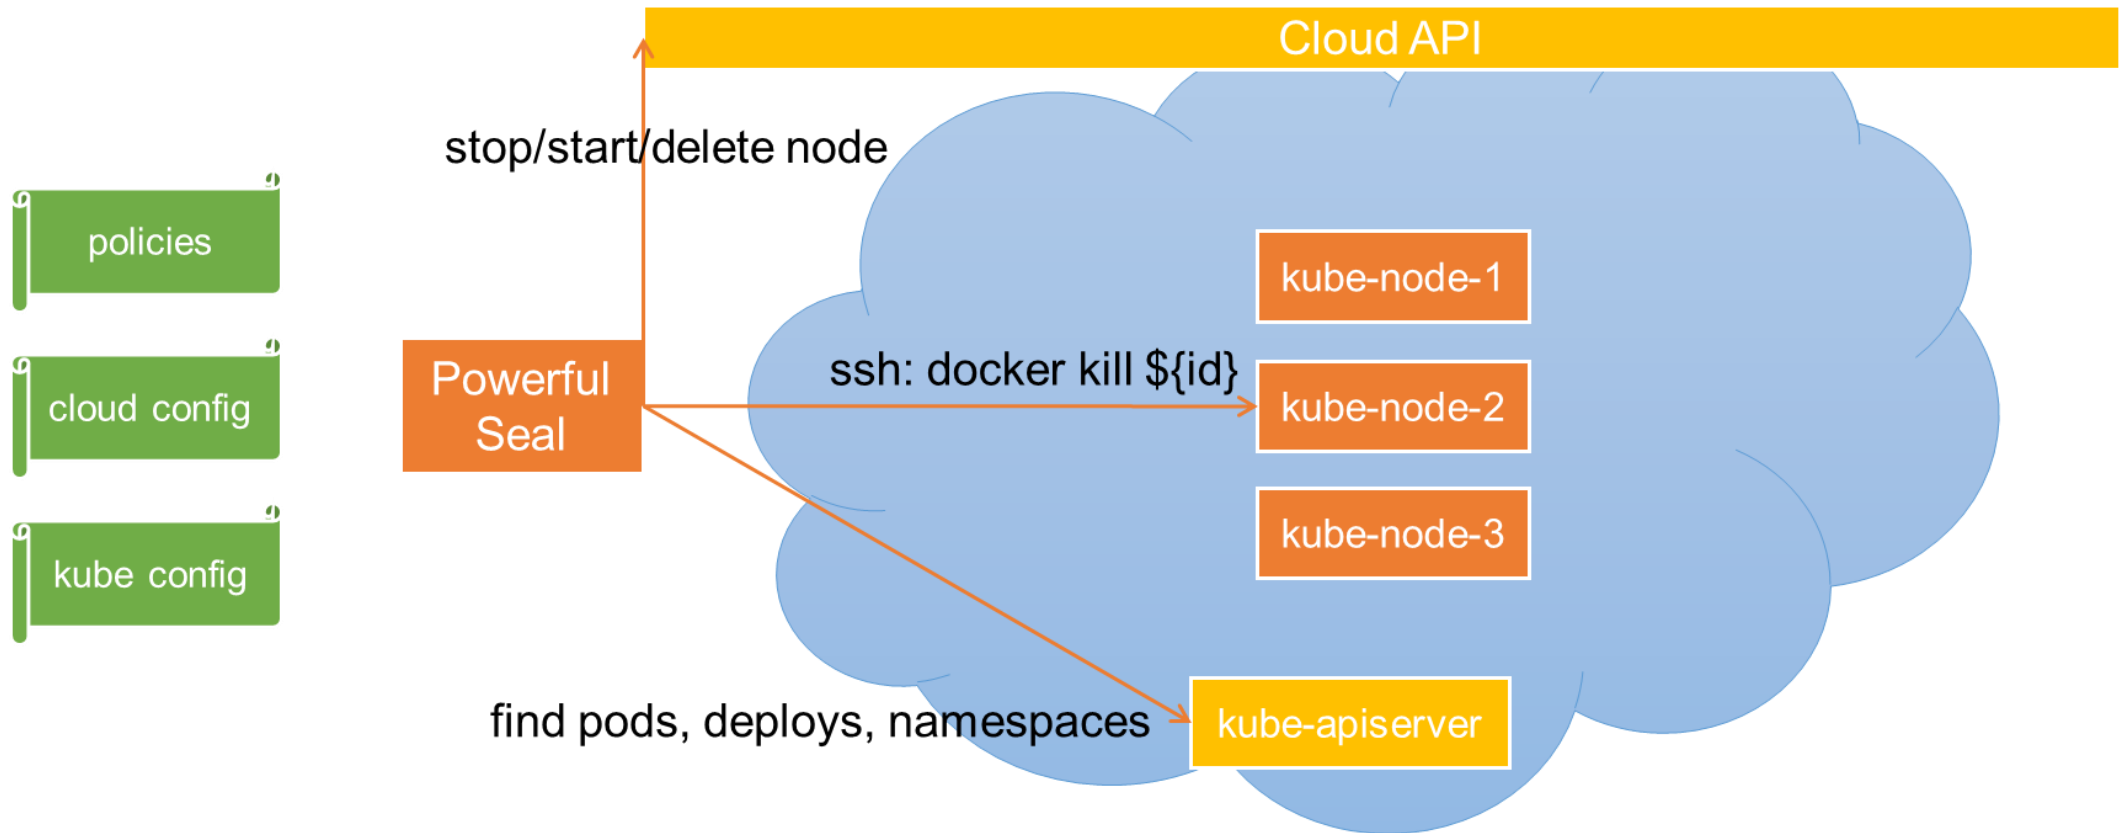
Project Jupyter UX Design



PowerfulSeal testing tool for Kubernetes

- Inspired by Netflix's 'Chaos Monkey'
- Tool to intentionally 'break' parts of a Kubernetes cluster
- Allows application developers to ensure that their application design can tolerate unplanned failures

PowerfulSeal testing tool for Kubernetes



Apache Solr on Kubernetes

- Apache Solr ‘operator’
- Allows a large-scale document index and search framework to utilize resources provided by Kubernetes clusters
- Separates data processing workflow management from compute resource management
- Leverages multi-tenancy support in Kubernetes to allow multiple Solr users to safely share compute resources

Support for open source organizations



TechAtBloomberg.com

© 2020 Bloomberg Finance L.P. [Creative Commons Attribution 4.0 International License](#).



Bloomberg

Engineering



Summary

- Our systems process billions of messages every day, and must meet our goals: Volume, Velocity, and Veracity.
- We employ open source infrastructure tools wherever we can.
- In many cases, our requirements cannot be met directly by the tools we choose, so we extend or enhance them.
- Whenever our extensions or enhancements would be useful to the user community, we contribute them to the project.

Internships, research, and full-time positions

- We're always looking for motivated, passionate people with engineering talent, for summer internships in San Francisco, New York, and London, and full-time positions in those cities and many more.
- If you're looking for a unique set of challenges in computing infrastructure or data science, we'd be happy to talk with you about our opportunities.

TechAtBloomberg.com

© 2020 Bloomberg Finance L.P. [Creative Commons Attribution 4.0 International License](https://creativecommons.org/licenses/by/4.0/).



Bloomberg

Engineering



Thanks!

If you're interested in learning more about our open source efforts, please reach out. I can be found in many ways:

- Email: kpflaming@bloomberg.net
- Twitter: [@realkpflaming](https://twitter.com/realkpflaming)
- GitHub: [kpflaming](https://github.com/kpflaming)
- LinkedIn: [linkedin.com/in/kpflaming](https://www.linkedin.com/in/kpflaming)

TechAtBloomberg.com

© 2020 Bloomberg Finance L.P. [Creative Commons Attribution 4.0 International License](https://creativecommons.org/licenses/by/4.0/).



Bloomberg

Engineering

