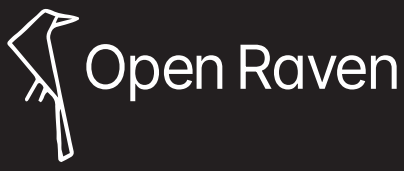

Open Source and Design at Open Raven

February 2020

Introduction: The state of Open Source

TABLE OF CONTENTS

There's no question that open source has changed the world. Open source projects are fueling software innovation, providing the essential foundation for projects ranging from self-driving cars, global banking and streaming TV to cloud hosting software, machine learning systems that detect cancer, spacecraft and more. It's telling that 69% of IT leaders say open source is very or extremely strategically important to their organization's overall enterprise infrastructure software plans (Red Hat).

Still, debate remains over the benefits of open source, the viability of this approach as a business model and the right way to cultivate related communities. Any startup embracing open source must have a clear understanding of these issues and a market approach that addresses them from the outset. History has taught us that open source strategies that are grafted on at later stages of a business have low likelihood of success.

This white paper will explore these considerations, examining them in historical context and how they underpin the design choices behind Open Raven.

Why Open Source?

For enterprises, there are a multitude of reasons for embracing open source technology. Open source software tends to be more flexible and agile, while also allowing the business to start small with lesser investments up front as they prove the concept – avoiding vendor l ock-in. Combined with the fact that there's shared maintenance responsibility, this type of software is often among the more cost-effective options, especially in the long term.

For software businesses themselves, the justification behind giving valuable intellectual property to the public for free is less clear. Among the philosophies and theories about why this model works, a few central ideas stand out:

CULTURE
TRANSPARENCY AND ACCOUNTABILITY
MORALITY
INDIRECT ECONOMICS

Culture

Software development is notorious for being dominated by opinionated, strong-willed, detail-orientated people. In recent years, there have been cases in the open source world where what can only be described as a toxic culture in some high-profile projects had become the norm (Business Insider, Vice). The open source model itself, however, has resulted in participants working together to define and create community guidelines that describe codes of conduct and ways to deal with specific circumstances (Microsoft, Google).

> It encourages a culture of participation, self-awareness and self-improvement.

Open source attracts people with common values and related skills, who join together working towards a shared goal to create the best software possible. In doing so, participants invest their time and effort and help and support others – typically without the guarantee of anything in return. Often they are motivated solely by the hope that their work will be recognized and acknowledged by their peers. The resulting culture is one of comradery, where individuals spend time working on the projects they believe in and have chosen to work on, with the people they have chosen to work with.

Transparancy and accountabliity

In many ways, the rising popularity of open source is also a reaction to what has happened with enterprise software. While many startups laud their large venture investment rounds across their homepages in the hopes of establishing financial confidence in prospective customers, it is nonetheless hard for a large established company to place a meaningful bet on any early stage company or technology given the amount of churn in the market.

Enterprises buying security products are far too often working to understand and integrate black box software without the ability to see what's happening under the hood. If the technology is failing them, they may not know it until it's too late. Furthermore, should the young company building the product fail, it can take its code down with it leaving the organization that invested in its software with an urgent hole to fill. Even in the event of an acquisition, the small company's product is many times greatly diminished due to the massive amount of distractions during and after the transition into the acquiring company.

Conversely, the open source model allows large enterprises to make bets on innovation in technology with far less risk.

> Open source code acts like an insurance policy; worst comes to worst, the organization has access and insight and can leverage their own engineers to extend solutions or pick up where their bought investments left off.

And from the outset, they benefit from improved transparency when using and integrating the software with their infrastructure.

Being the sole owner of the code base also means being in exclusive control of the features that will be built in the future and how much of it is led by user feedback (vs. partner demands, sales needs, etc.). In a commercial open source world, companies are forced to innovate and meet customer needs by the immediate, constant reality that if they don't, others will. While this can admittedly be uncomfortable, the obvious pressure to be engaged with user needs is an antibody versus complacency and ultimately, irrelevance.

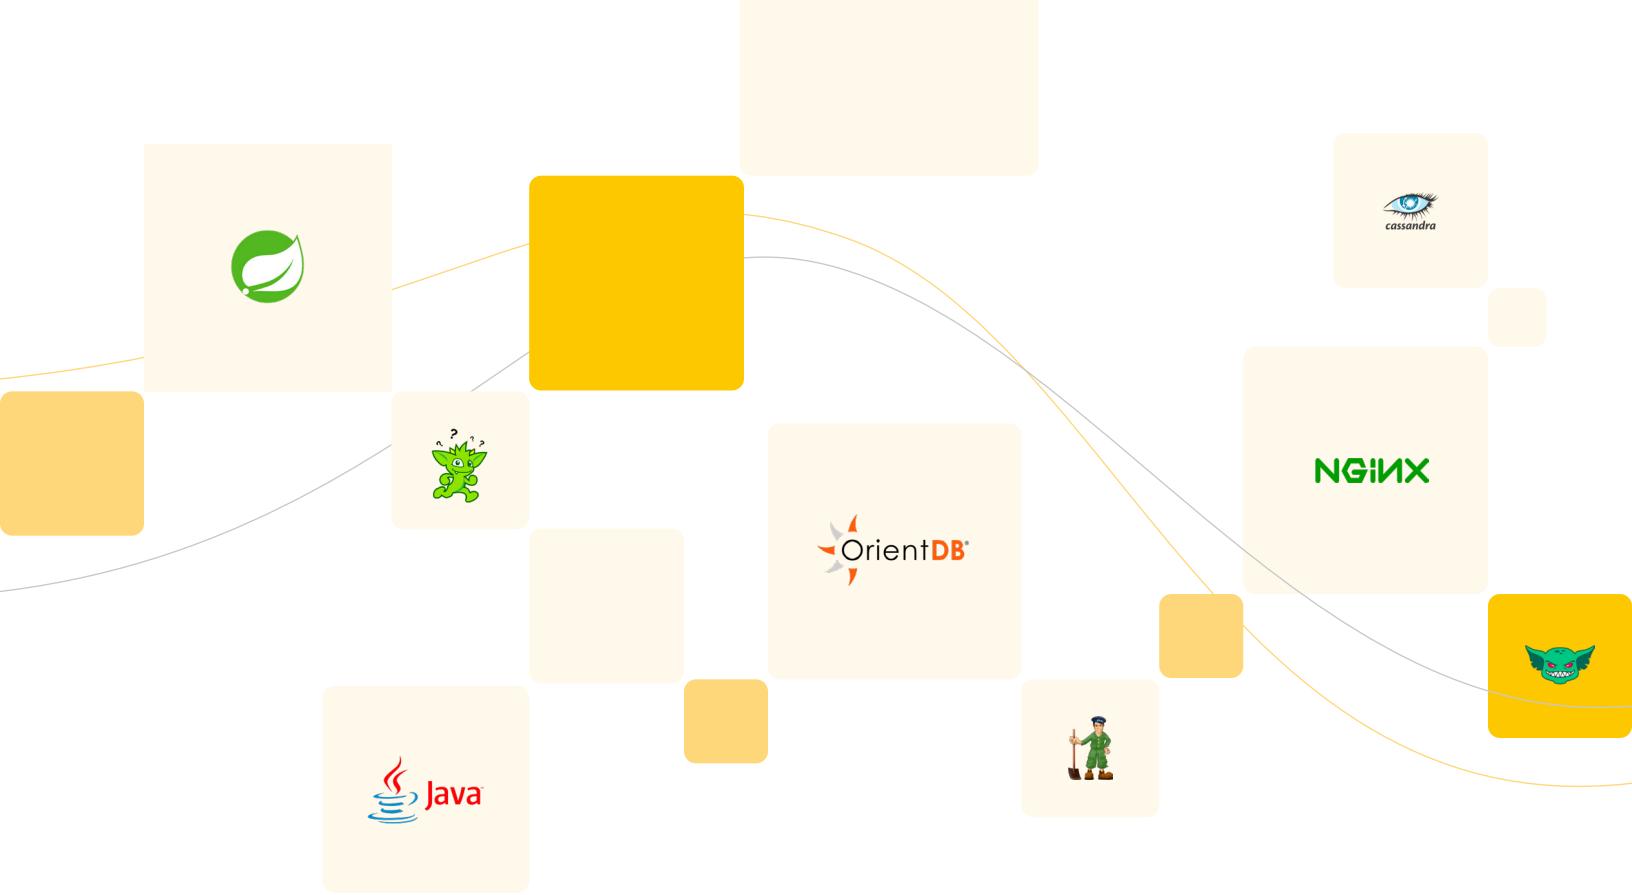

Why Open Source?

Morality

Open source works when everyone plays fair. It's about focusing on doing the right thing for the greater good.

> When building a platform using open source technologies, it's only fair to give back to those communities in return.

You'd be hard-pressed to find a modern security product that did not have at its foundations a substantial amount of open source code.

Open Raven could not be built if it wasn't for open source. We use Java, Spring, NGinx, OrientDB, Tinkerpop, Gremlin, Node.js, ReactJS, PostGres, Netty, ZMap, Zookeeper, Arcade Analytics, Cassandra and more. As well as open sourcing our core platform so that others can extend it, we are open sourcing many utilities we have built that others may find useful. We have and will continue to have a moral obligation to submit bug fixes and improvements to core technology we use. It's the right thing to do. This is the categorical imperative.

Indirect economics

The open source model has seen the rise of successful software companies like RedHat, MySQL, MuleSoft, Cloudera, Elastic, Gitlab, Kafka (Confluent) and Kong, Inc. Building successful open source software companies is now widely accepted as a viable commercial model. In fact, we now see regular examples traded on the public markets (Forbes, TechCrunch).

While there are some open source companies to reference in the security realm, we see no valid reason for there to be so few. Being able to audit the code and understand how protection measures are implemented should be a distinct advantage for security products. There is a protective force in transparency.

> Organizations shouldn't have to take the word of a salesperson; they should have the ability to look under the hood and understand what they're buying.

They should also know that what they use can withstand public scrutiny.Security vendors are often quick to claim their importance in defending an organization. We would argue anything of such importance deserves scrutiny. When the marketing claims reach far beyond the capabilities of a product, the mistakes can be costly, such as when a hot young antimalware product was tricked into thinking the prevalent WannaCry threat and other malware (Vice). While transparency can be uncomfortable, it makes scrutiny easier and when more people have a better understanding of how software works, ultimately everyone benefits.

When looking at successful open source companies, it's also evident that the commercial advantage has been one of indirect economics. Not one where third parties built software for free that was monetized by the company, but one where the indirect effects of being open source — such as recruiting, brand recognition and community goodwill — pay dividends. With open source, you "pay it forward" while "paying it back".

Open Source as a business model

Despite the compelling case for using open source as a business model, pitfalls abound. Many companies that have embraced open source fail to realize its potential and have been forced to change their original model to meet their commercial goals.

In most cases, this occurs as a result of the company selecting the wrong open source business model at the outset. Some companies have, for example, open sourced their entire code base and charged for a commercial license of the same code in the hopes that the companies relying on their code will pay. This rarely works as little is held back and the prospective customers discover that the free version is sufficient, leaving the company unable to monetize. Sentry comes to mind as an example.

Other companies have chosen restrictive copy-left licenses such as the GNU General Public License (GPL), which require customers to publically share all changes to the code no matter how sensitive or proprietary. As a result, established companies are unwilling to adopt their software due to the perception of unfriendly licensing terms that force disclosure of work they prefer kept private.

Other times, new technology paradigms like public cloud and SaaS negate or drastically reduce the originally intended protections offered by licenses like the GPL. Specifically, GPL's defensive provisions focus on distribution of the software itself, protection that is bypassed when GPL code is made available through network services such as SaaS. An example of such an issue and the resulting re-licensing effort is MongoDB.

With a range of open source business models available, no one model will satisfy everyone's philosophical beliefs. The important thing for a business is to be transparent about their open source model, specifically describing what it is and what it isn't.

Open Source at Open Raven

At Open Raven, we believe that the right model for a modern security company is what is commonly referred to as open core using a permissive Apache 2.0 license. In the open core model, we publish the source code for the foundations of our Community Edition under an Apache 2.0 license. This includes the discovery capabilities, data store fingerprinting and underlying graph data model. It does not include the user interface and back-end services we use to operationalize our platform for our users.

Community and Contributions

Open source project founders have learned time and again that if they don't continue to innovate on or maintain their core code base, then others will. Nonetheless, there is often well-placed concern with the open core model that others may build on top of the code base without giving back, or that they will only contribute features useful for them to build their own competitive advantage.

This situation has happened recently with Redis and CockroachLabs, forcing them to relicense and change to respective Commons Clause or Server Side Public Licenses. While these modified licenses come with additional issues for the vendor and the consumers, the organizations that use them must acknowledge that these issues may occur and deal with them if and when they do.

The Open Raven community

At Open Raven, we are transparent about what contributions we will take into core and what we will not. This is captured in a "Making Contributions" document that is posted in a readme file in each code repository on GitHub.

In general, the contributions we will take are:

- Bugs fixes
- Feature improvements or new features that we deem benefit all community users and or make the platform better
- Features or improvements that are on our community edition roadmap

We will not take contributions that compete with our commercial offerings (the planned Professional and Enterprise Editions) or contributions that we deem make the Community Edition code base harder to maintain. That doesn't mean others can't extend it as they see fit, but they will have to share the code and maintain it themselves.

Open Raven will be publishing a comprehensive Community Code of Conduct and reserves the right to refuse contributions from anyone we deem not following that code of conduct.

Creating change in cybersecurity with Open Source

Creating an open source security platform is an ambitious undertaking. As a business model, there are definite pitfalls on the road to profitability. In the security market, an open approach runs contrary to industry norms.

From a solutions standpoint, however, it's the most logical approach. An open approach leverages the intelligence of the entire community. Organizations can build what makes sense for them without building from the ground up and give back to the community in the process. Improvements can and will be made at a pace exceeding the capacity of any one business.

As a business ourselves, the open source approach keeps Open Raven honest and whole. If we aren't innovating, someone else will be. Therefore, as a business, Open Raven has consequences if we aren't living up to our promise.

However, what's truly exciting about this venture is the opportunity we have to make incremental but real gains on the issue of security. Harnessing the collective smarts and resources being poured into the security market in a collaborative way will enable us to move the needle rather than staying trapped in the buy-integrate-repeat cycle we've had to date.

As the Open Raven community grows and expands, it will be exciting to see the potential of our pooled resources realized.