OSS in the Enterprise

A Primer on the Commercialization of Open Source Software

v1.6 – March 2020

Shea & Company

Enterprise Open Source: A Primer

- In the late 1980s, Saturday Night Live ran a commercial for the First Citywide Bank of Change, a commercial bank with one line of business: making change. ("If you come in with a dollar bill, we'll give you four quarters.") The skit ends with a great line: "Our customers always ask, 'how do you make money doing this?' The answer is simple: volume."

- Until recently, prevailing wisdom seemed to be the open source business model was a bug, not a feature; Red Hat was an anomaly. Open source adoption was driven by cost considerations and a signal of the commoditization of a market segment. More pointedly: *how can anybody make money from selling something that's free?* But spurred by the rise of cloud, rapid crowdsourced development and the economies of a freemium go-to-market, open source software now underpins the businesses of a broad array of very successful (and money-making) businesses.

- While the open source model is not new, the last two years have seen an incredible inflection in market acceptance. Before 2017, Red Hat's 1999 IPO and Oracle's 2008 acquisition of mySQL were the signature transactions; 2018 alone witnessed ten multi-billion-dollar liquidity events, led by Red Hat's $34 billion acquisition by IBM.

- Generally, open source projects emphasize features and functionality, rather than enterprise-grade robustness. Integration with other critical enterprise assets is typically the responsibility of IT staff. Communities cannot provide 24/7 critical care support and are not deploying patches for enterprise use. This gap between software functionality and enterprise requirements provides the path to monetization for open source vendors.

- Open source is not just a viable business model, but potentially the preferred path to business building for vendors of enterprise infrastructure software.

> *"We didn't open source MongoDB to get help from the community, to make the product better. We open sourced as a freemium strategy to drive adoption."*
>
> Dev Ittycheria (CEO, MongoDB)

> *"95% of enterprises use open source within mission-critical segments of their IT portfolios, particularly within infrastructure areas, as well dev tools and frameworks."*
>
> Gartner (2018 estimate)

Shea & Company

Open Source in the Enterprise

OSS Impacts Virtually All Aspects of ITOM and App Development …

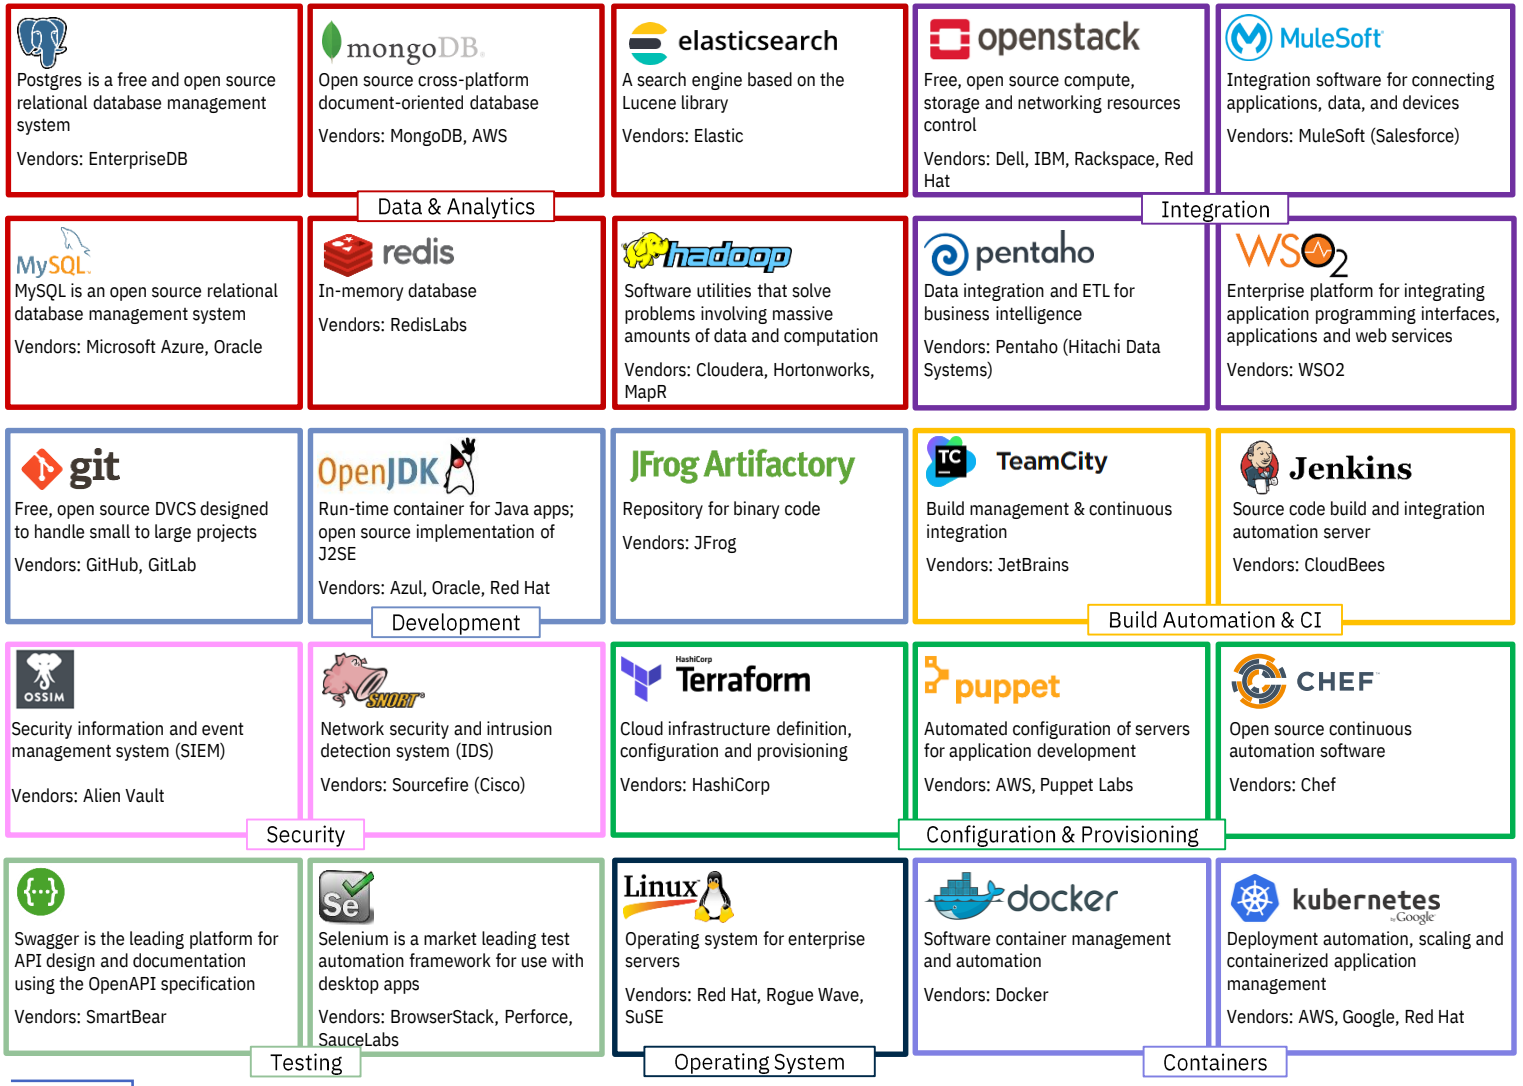

Shea & Company

For additional context, see Battery Ventures Open Source (BOSS) Index

Open Source in the Enterprise

… and Has Driven Many of the Recent Enterprise Software Success Stories

Pivotal	Acquired by VMware for **$2.7 billion** (January 2020); raised $555m via IPO (April 2018)
ACQUIA	Acquired by Vista Equity Partners for **$1.0 billion** (September 2019)
NGINX	Acquired by F5 for **$670 million** (March 2019)
HashiCorp	Raised $100m funding round led by IVP at **$1.9 billion** post-money (November 2018)
redhat	Acquired by IBM for **$34.0 billion** (October 2018)
elastic	Market cap of **$4.7 billion**, raised $252m via IPO (October 2018)
ALIEN VAULT	Acquired by AT&T for an estimated **$600 million** (July 2018) [a]
SUSE	Acquired by EQT for **$2.5 billion** (July 2018)
GitHub	Acquired by Microsoft for **$7.5 billion** (June 2018)
Magento	Acquired by Adobe for **$1.7 billion** (May 2018)
MuleSoft	Acquired by Salesforce for **$6.5 billion** (March 2018); raised $221m via IPO (March 2017)
cloudera	Market cap of **$2.5 billion**, raised $225m via IPO (April 2017)
mongoDB.	Market cap of **$9.0 billion**, raised $192m via IPO (October 2017)

Shea & Company

Notes:
Market capitalization data current as of March 2020
[a] 451 Research estimate

Favorable Market Dynamics Encourage Adoption by IT Organizations

Innovation is Under Pressure

30% Developers say they'd like to change their app stack but have no time or budget

54% Professional developers who select jobs based on the languages and frameworks they'd work with

Release Velocity Continues to Rise

67% Developers deploy code into production 1x per week

52% Applications are built in 3 months, but average lifespan is 20 years

Developers Favor Open Source

90% Developers believe Open Source software is of the same or better quality than proprietary software

65% Professional developers contribute to open source projects once a year or more

Critical Need for Support of OSS Packages

80% Issues with utilization of Open Source code are due to lack of knowledge or environment issues outside the open source package

51% Increase in reported open source vulnerabilities in 2017

Shea & Company

Source: Stack Overflow developer survey, 2019

Open Source Monetization Models

	Business Model	Representative Vendors
Distributions & Support Maintenance & Support for OSS	▪ No proprietary code ▪ Distributions of certified patches and releases, typically with enterprise-grade stability and security ▪ Fee for maintenance, support and installation, which can be either bundled or unbundled from the application itself ▪ Premium support serves as an “insurance policy”	
Open Core Proprietary code augmenting open source “core” product	▪ Mix of free “core” open source software, surrounded by a proprietary software layer ▪ Often associated with single-vendor (i.e. MongoDB, MuleSoft) rather than multi-vendor (Git, Linux) open source projects ▪ Open source “community” edition is free to use or modify, and offers unlimited free trials for developers to experiment before upgrading to a paid version ▪ Paid solutions provide features and capabilities required for enterprise-class production usage, or wrappers of adjacent (proprietary) functionality ▪ Often includes support for entire stack, including OSS core	

Shea & Company

Open Source Offers Several Advantages Over Proprietary Licensing

User Benefits

1 Practitioner-Friendly

- Technical users can download, try, modify and use open source components with little friction
- 90% of developers said open source code was of equal or better quality than proprietary offerings [a]
- Developers enjoy the ability to interact with others outside their enterprise, and 65% contributed to an open source project in the past year [a]

2 Development Flexibility

- Allows users to start small with community versions, then upgrade to enterprise-grade over time
- Allows the ability to deploy components of the code or build on top of it, to customize completely for the enterprise environment

3 Cost

- TCO (and of course upfront cost) is generally lower than proprietary solutions, for equivalent or superior capability
- Purchasers of proprietary software are in most cases limited to the support offerings from the vendor themselves – conversely, open source often enables enterprises to choose from many alternatives

4 Future & Future-Proof

- Future architectures are increasingly likely to be built open source
- Little risk for end-of-life for open source

Vendor Benefits

1 Built-in "Freemium" Adoption Path

- As enterprises become more "digital", IT practitioners are exerting growing influence on software and tooling decisions
- Once an OSS core is embedded, the upsell to an "enterprise edition" is dramatically accelerated
- This model has proven to be more efficient in some cases than traditional proprietary models (see Elastic vs. Splunk, next page)

2 Ability to Leverage User Community for Lead Generation

- Technical purchasers are hard to "sell to," but can be reached via user communities (e.g. those attached to OSS projects)
- Source of "warm leads" of a known group of potential purchasers, product evangelists and self-help resources, providing the "top of the funnel" for a freemium sales motion

3 Better Products

- Linus' Law: "*Given enough eyeballs, all bugs are shallow*" (named for Linus Torvalds, creator of Linux)
- The pace of feature/function expansion in many projects has been notably faster than commercial offerings – and enterprises have long cited the "technical superiority" of open source offerings [b]

4 Cloud-Friendly

- The cloud will increasingly be built on open source – with the notable exception of Microsoft Azure, the major cloud platform vendors are running open source stacks

Shea & Company

Notes:
[a] Stack Overflow developer survey, 2019
[b] PC World, Linux on the Rise for Business

Illustration: Business Model Benefits of Open Source Can Be Significant

	splunk> (FY 2013)	elastic (FY 2018)
Revenue	$ 302.6	$ 271.7
growth	*52%*	*70%*
COGS	35.8	78.0
Gross Profit	266.8	193.6
gross margin	*88%*	*71%*
Sales & Marketing	215.3	147.3
percent of revenue	*71%*	*54%*
Research & Development	75.9	101.2
percent of revenue	*25%*	*37%*
General & Administrative	53.9	46.5
percent of revenue	*18%*	*17%*
Total Operating Expenses	345.1	295.0
percent of revenue	*114%*	*109%*
EBITDA	$ (78.3)	$ (101.4)
EBITDA margin	*(26%)*	*(37%)*

Splunk (NASDAQ: SPLK) and Elastic (NYSE: ESTC) are competitive, market-leading vendors of software for search, monitoring and analysis of machine-generated data for large enterprises. Elastic commercializes the open source ELK stack (Elasticsearch, Logstash, Kibana) via a freemium model; Splunk is a direct competitor, via a proprietary offering.

- Elastic cited over 350 free million downloads of its products last year, then efficiently converted much of that trial interest into paid enterprise-class subscriptions and hosting sales
- Investing $65m less than Splunk at comparable scale, Elastic grew more than 70% year-on-year, in comparison to a still-impressive 52% growth rate at Splunk
- Elastic was then able to more-aggressively invest in R&D (contributing those innovations to the open source community)

Shea & Company

h/t: Battery Ventures partner Dharmesh Thacker, via Forbes

Common Open Source Licensing Schemes

- **Permissive Licenses (Apache, BSD, MIT)** are "free software" licenses virtually devoid of restrictions on the use and distribution of software; only requirement is reproduction of the license notice
  - Examples: Android operating system, Apache, Chef, Docker, Node, OpenSSL, Puppet
- **Weak Copyleft Licensing (Mozilla, LGPL)** is primarily used for software libraries/components, and allows developers to freely integrate open source into proprietary offerings, but *modifications to the specific open source component must be shared with the community*
  - OSS components are typically distributed by commercial software vendors separately from the proprietary application (often via a link) to ensure separation between proprietary and open source code bases
  - Examples: Adobe Flex, Firefox and LibreOffice, as well as many smaller development components providing discrete functionality
- **Strong Copyleft Licensing (AGPL, GPL)** requires any derivative application which includes, embeds or utilizes open source code be distributed under the same open source license terms (using the code means *the entire application must be shared with the community)*
  - Use of GPL code in proprietary applications is considered toxic, and was the genesis of BlackDuck's code-scanning business
  - Examples: Linux, MariaDB, MySQL, WordPress
- **Restrictive Licensing (SSPL)** is a newer, controversial variant of strong copyleft licensing designed to close the "loophole" which allows delivery the open source application as a commercial cloud service (e.g. AWS DocumentDB)
  - Examples: MongoDB, Redis

"Copyleft"
An arrangement whereby software or artistic work may be used, modified and distributed freely on condition that anything derived from it is bound by the same condition.

A leading provider of source code scanning and open source compliance, Black Duck was acquired by Synopsys for $565m in 2017

Shea & Company

Open Source in the Enterprise

Case Study: Alien Vault

Financial Highlights

- **Model:** Open Core
- **Status:** Acquired by AT&T for $600m [a]
- **FY18E Revenue:** $150m

Company Overview

- Founded in 2007 in Madrid, AlienVault is a developer of both commercial and open source solutions to manage cyber attacks, and is the sponsor of Open Threat Exchange, the world's largest open source IT security platform
- AlienVault offers Unified Security Management, a security monitoring product based on the open source Open Source Security Information Management (OSSIM) system, which combines and correlates event logs from IT devices, IT asset discovery/management and information from threat detection systems into a unified view of an enterprise's security environment
- The company's paid solutions add additional features including as cloud hosting, security monitoring features for AWS and Azure, continuous delivery of patches and support services
- Alien Vault also offers managed security monitoring services, managed detection and response (MDR), SIEM-as-a-Service (security information and event management) and compliance management
- Open Source projects include:
  - Open Threat Exchange – a global community of more than 100,000 threat researchers and security professionals who contribute threat indicators and collaborative research daily as well as automatic updates to security infrastructure from the collected data
  - Alien Vault OSSIM – AlienVault's open source SIEM project, which features event collection, normalization, and correlation via open source tools backed by the Open Threat Exchange
- AlienVault monetizes by offering a cloud-hosted, enterprise-grade subscription-based version of their open source security solution
- The company's competitors include Splunk, LogRhythm, Rapid7, Fortinet, Dell RSA, and Exabeam

Shea & Company

Note: [a] 451 Research estimate

Case Study: Azul Systems

Financial Highlights

- **Model:** Support + Open Core/Proprietary
- **Status:** Private, owned by Vitruvian Partners
- **Revenue:** undisclosed

Company Overview

- Sunnyvale, CA based Azul Systems is a leading independent provider of commercial support, secure distributions and software applications for Java runtimes, built off the OpenJDK project
- OpenJDK is a free and open source implementation of the Java Platform Standard Edition (Java SE) – a runtime environment providing a container for Java based applications allowing portability across desktop and server environments
- Oracle distributes two versions of the JDK: Oracle Java SE, a fully-supported commercial software offering (closed source) and Oracle OpenJDK, an open source Java build
- Through Oracle's decision to heavily monetize Java in 2016, a need opened in enterprises for an affordable, open-source JDK and JVM that Azul saw as an opportunity
- Azul offers two distinct solutions:
  - Zing – patented, proprietary, high performance JVM for specific applications based on OpenJDK, but with unique improvements to speed, performance, scalability and consistency
  - Zulu – an open-source Java SE alternative that is offered in three forms: (i) Enterprise – paid version with included commercial support, (ii) Embedded – paid OpenJDK built specifically for embedded devices (e.g. IoT, bundling, redistribution, etc.) and (iii) Community – free OpenJDK without support
- Within the market for JDK distributions and support, Azul commonly competes with Oracle, Red Hat (IBM), BellSoft and OpenLogic (Perforce); indirectly several cloud vendors offer builds of OpenJDK running on their infrastructure, including Amazon Corretto; Microsoft Azure runs Azul's Zulu Enterprise build

Case Study: Chef

Financial Highlights

- **Model:** Support
- **Status:** Private, has raised a total of $105m from venture investors including Amplify, Battery, DFJ, Ignition and Scale
- **FY19E Revenue:** estimated $50-$75m [a]

Company Overview

- Founded in 2008 in Seattle, Chef is a leader in continuous automation and helps over 40,000 organizations automate deployment, configuration and provisioning of IT infrastructure resources in connection with software application deployment
- Chef's Enterprise Automation Stack allows companies to define infrastructure, security policies and application dependencies as code, deliver the stack via an automated pipeline to any platform, and deploy, observe and manage the stack over its lifecycle. Solutions include:
  - Chef Infra - Infrastructure automation to provision, harden and maintain configuration state
  - Chef Habitat – Automated application dependency management
  - Chef Inspec – Security and compliance automation
  - Chef Automate - DevOps dashboard to help gain visibility and control over automation efforts
  - Common use cases include application delivery, cloud migration, patch management, delivering to Kubernetes and technology lifecycle management
- Additionally, the company has begun offering an Effortless Infrastructure Suite to codify infrastructure, security and compliance, and to audit and manage architectures
- Historically, Chef had leveraged an open core model with a number of proprietary products that complemented its open source tools, but in April 2019 the company began open-sourcing all of its software under the Apache 2 license
- While Chef's software is completely open source, the company charges an annual fee that covers hosting and support
- Traditional competitors in the IT automation space include Puppet and SaltStack, as well as Ansible (Red Hat), Hashicorp, Bamboo (Atlassian) and TeamCity (JetBrains); more recently increasing adoption of containers (Docker, Kubernetes) have impacted the share of wallet for IT automation broadly

Shea & Company

Note: [a] Shea & Company estimate

Case Study: GitHub

Financial Highlights

- **Model:** Open Core
- **Status:** Acquired by Microsoft for $7.5b
- **FY18 Revenue:** $375m

Company Overview

- Founded in 2005 as a community for development and collaboration based on the open source Git repository, GitHub today is one of the largest code hosts in the world, with over 100 million private, public and open source development projects and more than 30 million users
- The company was acquired by Microsoft for $7.5 billion in 2018
- At GitHub's core is the popular Git source code repository – an application that allows developers and teams to store and control access to application source code, and most importantly provides "version control" allowing multiple developers to check in / check out code from the same base, resolve conflicts and easily track changes to code over time
- GitHub hosts one of the most popular open source communities, built around the largest publicly-available repository of open source projects in the world – it is so ubiquitous, data around access to these projects within GitHub is used as an industry-standard measure of a given project's popularity
- The company monetizes Git via an open core model, offering significant proprietary technology and services which significantly extend functionality well beyond the open source core – from an initial foundation of cloud hosting, GitHub offers a number of solutions for individuals, teams and enterprises, including security, integrations to a broad array of developer solutions and collaboration
- The company utilizes a freemium model, offering basic solutions (beyond raw Git repositories) which are free for individual developers or small teams, and heavily leverages their community for lead-gen, first-line support and customer advocacy
- The company's competitors include other version control systems including other Git-based services (GitLab), other open source competition (Atlassian Bitbucket, Subversion) and proprietary solutions from Perforce and IBM

Case Study: MuleSoft

Financial Highlights

- **Model:** Open Core
- **Status:** Acquired by Salesforce for $6.5b
- **FY18E Revenue:** $410m [a]

Company Overview

- Founded in 2006, MuleSoft is a leading provider of solutions which integrate and orchestrate on-premise, hybrid and cloud applications based on APIs – the company's platform is arguably the most popular integration platform for enterprise integration
- The "mule" in the company's name originated from the "donkey work" of data integration
- MuleSoft's solutions core solutions include Mule ESB, which manages and secures the data flow between cloud and on-prem applications, and CloudHub, an iPaaS providing a global, fully-managed, multi-tenanted, secure and highly available platform for APIs and integrations
- Mule ESB is available in an open source "community edition" targeted at evaluation and pre-production use; deploying the open source edition in an enterprise production environment requires significant internal resources and expertise to support and manage the ESB infrastructure
- The company was initially launched as a middleware and messaging focused ESB, but soon evolved the business to provide the industry's leading integration platform as a service (iPaaS) with a particular focus on API-based integration; the company's solutions include:
  - Anypoint Design Center (developer solution for API design)
  - Anypoint Exchange (shared API library)
  - Anypoint Management Center (web UI to analyze, manage and monitor APIs and integrations)
  - Mule ESB (runtime engine for connecting enterprise applications)
- MuleSoft was long a darling of the investor community at all stages of its corporate lifecycle – prior to its IPO, MuleSoft raised $259 million in venture capital from a top-tier syndicate of VCs; the company went public in March 2017 and saw a 45% increase in the stock price on its first day of trading; less than a year later Salesforce announced the acquisition of MuleSoft for $6.5 billion
  - Salesforce led MuleSoft's $120m Series G financing in 2015
- The company's competitors include modern solutions from Apigee, Boomi (Dell), JitterBit, SnapLogic and Workato as well as legacy integration vendors Informatica, Oracle, SAP and Tibco

Shea & Company

Note: [a] estimate prior to acquisition by Salesforce

Case Study: Red Hat

Financial Highlights

- **Model:** Support
- **Status:** Acquired by IBM for $34b
- **FY19E Revenue:** $3.36b

Company Overview

- Red Hat was founded in 1993 as an commercialization of the open source-based Linux operating system, which was itself an alternative to Unix; Linux was first released in 1991 and is perhaps the most widely adopted open-source project in the world
- Red Hat began distributing commercial builds of Linux via CD, then migrated to servers and professional services – 100% of Red Hat products are available as FOSS (free open source)
- While Red Hat invests heavily in R&D, ultimately all modifications are contributed back the community and distributed as GPL (freely-available) open source; the company is the leading contributor to the Linux open source project and in particular is noted for quick releases of security patches The company was acquired by IBM in the largest software acquisition and third-largest technology acquisition in history in 2018 for 11x revenue
- Red Hat competes with other freely-available distributions of its products (e.g. the CentOS or Ubuntu builds of Linux) as well as large enterprise software vendors such as Microsoft, Oracle and VMware
- The company's flagship product, RHEL (Red Hat Enterprise Linux) launched in 2003, is an enterprise-class Linux distribution which includes security and enterprise features; RHEL is freely distributed in source format (as is required by the Linux license) – Red Hat's primary value add to its customers are i) first-line support for any technical issues; and ii) distribution of patches and software upgrades to its customers in binary format
- Red Hat is also a leader in open source cloud technology – the company is a sponsor of OpenStack (IaaS) and OpenShift (container-based development), and notably has made several recent cloud-oriented acquisitions, including Ansible (DevOps), CodeEnvy (container/cloud-native development) and InkTank (software-defined storage)

Case Study: SmartBear

Financial Highlights

SMARTBEAR

- **Model:** Open Core + Proprietary
- **Status:** Private, owned by Francisco Partners
- **FY19E Revenue:** undisclosed

Company Overview

- SmartBear was formed in 2010 following the three-way merger of AutomatedQA (test management software), Pragmatic Software (ALM software) and the original SmartBear (code review software)
- The company was acquired by Francisco Partners in 2017 and continues to have an aggressive organic and inorganic growth strategy targeting both open source and proprietary technologies, with a particular strength in QA and developer-led testing
- SmartBear's product suite includes leading solutions covering API design, quality, monitoring and testing; browser and mobile device test automation; test management; and UI functional and performance testing
- Across its portfolio, SmartBear utilizes a freemium go-to-market model, offering easy-to-try solutions which are free for individual developers or small teams, which the company then leverages via a land-and-expand motion for efficient department and enterprise-wide selling
- The company supports several open source projects including Swagger (OpenAPI), Cucumber and SoapUI, and offers solutions built on top of these projects / communities:
  - SwaggerHub: a commercial offering built on top of the OpenAPI specification and Swagger community, is the leading open-core API design and development platform
  - SoapUI Pro: the most advanced REST and SOAP API testing tool and, in combination with Swagger, positions SmartBear as the leader in open source-based API development and testing tools
  - Cucumber: an advanced behavior driven development (BDD) open source test framework targeting developers, on top of which are built enterprise-grade tools (Cucumber for Jira) integrated with SmartBear's proprietary BDD tool HipTest
- The company's competitors include vendors such as Apigee (Google), Parasoft, Postman, Sauce Labs and Tricentis as well as incumbents including CA, IBM and Micro Focus