APPENDIX A

Open Source Continuum

The Open Source community has gone full circle over the past few decades, from first being considered as hobbyist software that was buggy and not to be considered as professionally ready for prime time to making inroads into a myriad of businesses, governments, and organizations. With solution offerings that many times surpass the quality and capabilities of the most costly and proprietary based systems, Open Source solutions have proved themselves capable. From professional grade operating systems that are taking over in the world's data centers to business office oriented tools like spreadsheets and document editors, the Open Source world of solutions has matured and become a mainstay; not bad for once-labeled products of amateur hobby.

Large proprietary software and solution providers that had once dismissed Open Source solutions as inferior amateur products have been forced to come to terms. In the raw actuality the term Open Source is a more modern term for software sharing and collaborative development principles that have been around since the advent of the term software. It has become common knowledge that in the early days of computing, virtually all software was created by academics and researchers in the corporate world who were all working collaboratively and sharing the results of their endeavors openly and freely. It is literally amazing that the digital and computing world that modern times have experienced came from such humble beginnings.

W. Wood, *Migrating to MariaDB*, https://doi.org/10.1007/978-1-4842-3997-1

Open Source in the Data Center

If there was a defining moment in time that truly indicated the beginnings of the modern advent of what we now call Open Source, one could target the year 1991 as just such a moment. This was when Linus Torvalds released the first version of his Open Source Operating System, effectively going on to be named Linux, which was based in part on System V Unix and written in the C programming language. There are many other software solutions and operating systems of note that came before Linux; however, when looking back over time and comparing to what is driving the current data center and having a monumental impact on modern computing today, the Linux OS stands out.

Originally released just as an OS kernel, it has aged gracefully and migrated into the OS of choice in computing industry data centers around the world. Initially it took the majority of web-based server roles in large part to other Open Source driven works like Perl, PHP, and MySQL. Now the Linux OS has become the go-to for web hosting and web development for which terms like LAMP Stack were coined, which referred to the combination of Linux Apache MySQL and PHP/Perl. The population of Linux servers has literally exploded over the years, supplanting such mainstays as HP-UX, Windows Server, and Solaris just to name a few, to the point that in some statistics it's holding 60% to 70% of the rack space in web hosting environments. With the web hosting market as just the beginning, the Linux OS with its many distributions (distros) has not stopped there and has moved into enterprise application hosting as well.

There are several Linux distros to choose from; however, we are going to maintain our focus on the one pulled from the story about Vernon and his team at Financial Widgets Plus (FWP), RedHat. The RedHat organization is one of the largest success stories coming from the Open Source arena, both commercial and community based, via their Enterprise Linux and Fedora versions, respectively. Their team chose RedHat Enterprise Linux (RHEL) as their OS of choice when replacing their HP-UX environments, experiencing improvements in both performance and ease of maintenance when it came to hosting their Oracle RAC environment. Oracle even has released their own distribution, effectively named Oracle Linux, which has a striking resemblance to RedHat's RHEL with a kernel that is reportedly enhanced specifically for running the Oracle DBMS and similar products. This has resulted in wild speculation in media circles and authorships of a bit of hard feelings between the two organizations, but this will be left to the individual reader to interpret in regard to any significance therein.

RHEL is a rock solid and very well supported commercial version of the RedHat distribution, and along with many other distribution providers they have in recent years begun to supplant the historical inclusion of the MySQL ODBMS with their newer releases with MariaDB. This aversion to Oracle products can be interpreted in many ways as a storyline similar to that of the fictional account of the FWP team in regard to the fictional business practices alluded to by the Oracle Corporation in recent years. This is left to the readership to make their own decisions, do their own research, and form their own opinions on such matters. As to why businesses make the choices they do, in many cases, just as in Financial Widget Plus's case, the primary driver is that it made good business sense to migrate away from their high-cost proprietary legacy database solution.

Anyone would be hard pressed to claim that Oracle has not developed a less than desirable reputation over the years, which comes with growth and sometimes can be a by-product of success as well as competition driven. The fact is Oracle and their perceived business practices had a positive effect on the Open Source community with their purchase of Sun Microsystems, and by inclusion MySQL, in 2010. This was the catalyst for the creation of the MySQL fork by Michael Widenius that has become a perfectly viable replacement for Oracle's own Enterprise Edition DBMS, with the inclusion and support of encryption for data at rest released in version 10.1.

With many improvements since release 10.1, MariaDB has grown their solution into a fully dependable, professionally driven and secure database solution that is an entirely Open Source commercial product. There is a magnanimous amount of excitement growing in the computing industry about MariaDB, and it's future appears to be limitless with the strides in market share and improvements the folks there are making on a daily bases. Where RedHat has taken over the data centers in recent years, MariaDB now stands poised to take over the data.

Entrepreneurial Limits of Big Name Proprietary Systems

The costs as calculated and used by Vernon in the fictional example for FWP used pricing metrics that were derived from publicly available pricing lists published by the Oracle Corporation. These numbers beg the question as to how would any small business, enterprising startup, or in the case of an existing small company like FWP who is trying to expand, afford to do so. These costs are amazing in comparison and could easily break a business just from the initial upfront costs. This is where the commercially available versions of software from the Open Source community are not only viable solutions, but solutions that will assuredly garner a substantially growing market share in the future.

The entrepreneurial spirit and competition are easily squashed by high-priced and closed proprietary solutions that can run into the millions of dollars in cost just for the first five years of a new business. This is money that would be well spent in generating more revenue to help a business succeed in the first few years of their incarnation, and this is something that closed system proprietors like the Oracle Corporation do not seem to understand the concepts of and where companies like RedHat and MariaDB stand to make their mark. They see the bottom line relying on the sole benefit of the short-term hustle by charging as much as possible without benefit of looking at the longer term arrangement. What is suggested here is counterintuitive to the age old management and business practices that have been employed historically; however, the logic is sound.

Revenue generation and earnings impact can be expanded greatly over time with an approach centered on creating and developing a business relationship using a much more affordable pricing model to grow your earnings over time. This can be easily modeled using the fictional FWP and the information garnered from their experience. According to the Small Business Administration (SBA) one out of five businesses fail in the first year and approximately 50% fail within the first five years, so comparing Oracle's pricing model with that of MariaDB we have the following calculated costs for five years:

- Oracle: $1,562,280
- MariaDB: $187,500
- Difference: $1,374,780

This is a huge discrepancy of $1,374,780 comparing a standalone three-node setup in a clustered database environment. Imagine what any business could do with that extra money, much less an enterprising business bent on success. That would be money that could be reinvested back into the business for marketing, resources, staffing, and various other methods used to obtain more customers and augment existing solutions. This amount of savings being put back into the business in a beneficial manner could very well mean the difference between success and failure. In the end, would it be enough to improve the SBA statistics is not easily answered; however, it is certainly food for thought and a topic to be explored.

There is always a flip side to any kind of extruded benefit from this logic and that is that if 50% of all business are going to fail, then it would be better to reap the highest potential earnings possible for a large entity such as the Oracle Corporation. In this scenario, only one entity benefits from this strategy and can add another island to their holdings, whereas with the logic being presented the benefits are spread out exponentially along with the potential for a higher business success rate not only with the business, but also with their vendors, including their ODBMS vendor. With more revenue to invest in driving solutions to market resulting in more customers, and possibly providing more services to existing customers through increased product offerings, this will create a ratio to include having to increase and maximize their database footprint to account for this growth.

This is where the Open Source Continuum exists in all of its disjoint glory from the closed source proprietary systems. It allows for maximizing success due to lowered costs, while maximizing its own future at the same time in doing so, if leveraged appropriately with good business practices. Exploring this is as easy as envisioning a new start-up company, calling it OSC in honor of the Open Source Continuum, and analyzing the potential in a monetary timeline starting at the end of its first year.

Exploring the state of OSC, they effectually have broken even after inking a five-year deal with MariaDB for their database service on the single three-node production cluster leaving them over a million dollars to reinvest back into the business the second year. OSC has also contracted to provide their service to 55 customers in their first year, and they have found that they can easily run an optimum tuned database count of 50 customers on their three-node cluster with no performance impact. The state of the business is as such:

- Funds being reinvested the second year: $1,374,780
- Current number of customers: 60

- Customers in various stages of contract negotiations or deployment: 12
- Additional potential customers: 6

Let's imagine that OSC decides to spend part of their operating funds on marketing, attending conferences, and exploring additional business relationships in order to expand their business and need to add additional clusters to meet these needs. With a large influx in business comes the need to add additional database resources, thus expanding the footprint and increasing licensing requirements for MariaDB. This is a successful business model where many more benefit.

At the end of year two, OSC has expanded their customer base significantly and is becoming quite the success story. Their current state of business is:

- Current number of customers: 110
- Customers in various stages of on boarding: 24
- Potential customers: 16

OSC has had to expand their MariaDB licensing to account for this and now has three production clusters running. This has tripled their new ODBMS footprint and they are now spending almost $600,000 with their new vendor for licensing and support. In the meantime, the MariaDB organization has experienced corresponding growth as well, which means hiring more staff to handle the increase and an increase in revenue.

It is easy to see that Open Source embraces the entrepreneurial spirit and provides solutions that would otherwise be beyond reach. From the developer working in her/his free time on an idea to add to the code base, to the small start-up, and reaching all the way to the largest organizations the solutions that were once thought of as amateur are coming into their own. From RedHat taking over the data centers to MariaDB taking on the niche that a company like Oracle has held onto for so many years, it is an exciting time to be working in the technology sector.

Where Is Open Source Not Viable

All of this begs the question, is there anywhere that Open Source technologies cannot be considered a viable solution, and that answer is becoming harder to get to every day. The advent of taking Open Source solutions into the business world via commercially branded and supported versions, that by the way are still open, has changed the viewpoint to where proprietary closed systems are not seen as the only answer anymore. Community driven projects can be leveraged for virtually anything and their potential is limitless.

There will almost always be proprietary systems and code, as that is also part of the entrepreneurial spirit as well. In many cases intellectual rights to ingenious ideas, designs, and concepts do need to be protected and closely guarded, as they may very well relate to a business or entity's survival and earnings potential. The optimum solution is one that employs Open Source to drive these types of ideas and designs. Stealing an often heard buzzword, we could call these hybrid Open Source solutions.

An arguable point might be that highly secure systems for military and government use might not be a good area for Open Source solutions; however, that actually flies in the face of the entire concept behind it. It has been the author's opinion for many years that if someone can build something, they can certainly break it down, and since many of the Open Source solutions come from the global community there is certainly concern with terrorism, security breaches, and data theft. This is true with virtually any software that is pirated, and by using reverse engineering methodologies, even with closed proprietary systems. However, with Open Source solutions the code is much more available with potentially less time involved in finding a security flaw. On the other side of that same coin is that these solutions are Open Source and can be modified, ported, and expanded from their original base to add in additional security, validations, and capabilities.

In the early days of what we now refer to as Open Source, the risks were potentially very high with many solutions being buggy, lacking thorough documentation, and in essence much broader to an organization. Commercially marketed and supported solutions have mitigated many of these issues with dedicated resources from the vendors in regard to developers and engineers augmented by services and offerings such as 24/7 technical support, training, certifications, and even remote administration available to augment one's internal staff.

Benefits of Open Source

The benefits of Open Source as presented are boundless. Even though it has been around a very long time and was the basis of almost all initial software development, it has made a resurgence and has come full circle into its own with what are becoming standard deployments and go-to solutions. The growth in open source offerings will continue, with their commercially offered counterparts being the springboard into mainstream adoption.

The entrepreneurial spirit of open source solutions opens up the same spirits in small companies and start-ups that can leverage these capable but lower cost solutions to get into the marketplace with their products and thrive. This has been proven time and time again over the past two decades as open solutions have taken over, as vast amounts of the daily processing that occurs over the World Wide Web is being performed by solutions like Apache, MySQL, PostGres, Perl, PHP, and Linux. Fortunes have been made with online stores that initially had low cost and short implementation times to get up and running due to leveraging open source solutions to drive these virtual store fronts.

Anyone with a desire to learn technology can do so with open source technologies; whether they want to be a programmer, database administrator, web developer, systems administrator, or anything in between, the opportunity is there. It is free, it is open, and all one has to do is download it. The avenues are there to learn the software and even become a contributor for anyone who has an interest. The same cannot be said for proprietary closed systems; not that they don't have their place, but it isn't as easy to get into the nuts and bolts and improve them.

Index

W. Wood, *Migrating to MariaDB*, https://doi.org/10.1007/978-1-4842-3997-1