

Argentina.com : A Case Study

Carlos Horowicz

ch@argentina.com

\$FreeBSD: doc/en_US.ISO8859-1/articles/casestudy-argentina.com/article.sgml,v
1.2 2005/09/04 22:03:59 joel Exp \$

FreeBSD is a registered trademark of the FreeBSD Foundation.

CVSup is a registered trademark of John D. Polstra.

Intel, Celeron, EtherExpress, i386, i486, Itanium, Pentium, and Xeon are trademarks or registered trademarks of Intel Corporation or its subsidiaries in the United States and other countries.

XFree86 is a trademark of The XFree86 Project, Inc.

Many of the designations used by manufacturers and sellers to distinguish their products are claimed as trademarks. Where those designations appear in this document, and the FreeBSD Project was aware of the trademark claim, the designations have been followed by the “TM” or the “®” symbol.

Table of Contents

1 Overview	??
2 The Challenge.....	??
3 The FreeBSD solution.....	??
4 Results	??

1 Overview

Argentina.Com is an Argentine ISP with a small infrastructure of fewer than 15 employees and whose primary source of income originates in the free dialup business. It began operation in the year 2000 with barely one server for mail and chat.

It has since grown to a market presence in the Argentine free dialup market of 4.5 billion minutes annually. Its most popular product provides nearly half a million users with free e-mail with webmail, POP3 and SMTP access, and 300M disk space. Towards the end of 2002 there were around 50,000 mail users. After two and a half years of re-engineering and consistent technical improvements this ISP has grown by a factor of 3 in terms of billing, and by a factor of 10 with regard to the mail user base.

Our competitors in the Argentine market of free dialup include Fullzero which is owned by the Clarin Media Group, Alternativa Gratis, and Tutozia which is funded by IFX and promoted by Hotmail. Some of these large corporate competitors started their free dialup business with multi-million dollar investments and aggressive television and Internet ad campaigns. Argentina.Com does not rely on advertising like these other larger corporations. It has climbed to the fourth position and to an 8% market share during the last two years thanks to superior quality of service.

In Argentina and Latin America in general people who do not have computers at home go to so called “Locutorios” (Internet Centers), where for a few pesos they can use a computer connected to the Internet and usually read and write emails through popular webmails like Hotmail, Yahoo or Argentina.Com.

Due to limited financial resources, Argentina.Com made the decision to invest in a new email system instead of publicity in the media. This strategic decision opens the door to a future business in the corporate and paid email arena.

2 The Challenge

The main challenge for Argentina.Com is to achieve a dialup uptime of at least 99.95%, or less than 5 hours yearly downtime. Due to the high rotation and volatility in this business, things have to work correctly so the user does not switch -voluntarily or not- the dialup provider or the number he calls to connect. The dialup business involves a support structure to deal with the Telcos about telephony problems and quality of service, plus a technical structure where latency and packet-loss should be minimized due to the UDP nature of Radius and DNS, and where recursive DNS should always be available.

This also implies having a high uptime in the POP3 and SMTP services, and in the webmail. For POP3 and SMTP we estimated the need for an uptime equal to the one for dialup, whereas for the webmail we could live with 99.5% which means around two days of yearly downtime.

We decided to migrate the email to a proprietary, opensource architecture which should be horizontally scalable, and whose antivirus and antispam infrastructure should support more than just one type of mailstore or back-end.

The rough competition in the free email market, mostly due to the recent improvements introduced by Hotmail, Yahoo and Gmail, made it necessary to design the new system with at least 300M user disk space, but at a cost lower than 3 US dollars per GB with some degree of redundancy. Bear in mind that rackmountable hardware is hard to find in Argentina, and is between 30 and 40% more expensive than in the US. Our total budget for equipment acquisition in two years was 75,000 USD, which is only a fraction of our direct competitors’ investments.

With regard to the antispam service, it became necessary to develop a product that could compete with the systems offered by the big ones. Given the hostile conditions imposed by the existence of spam (dictionary attacks, spams with high degree of obfuscation and refinement, phishing, trojans, mail-bombs, etc.) it becomes very difficult to achieve an excellent uptime while repelling attacks. One must also be careful that the user does not lose mails because of false positives in the classification strategy, that he does not become flooded with spam or spam notifications, and dangerous mails do not make it through to his mailbox. In addition, the technical infrastructure for spam classification should not introduce noticeable delays in the delivery of mails. Finally, the mail system has to be protected from spammers who might misuse it to send spam.

The opensource paradigm tends to require hiring large teams of system administrators, operators and programmers who apply patches, correct bugs and integrate platforms. The opposed paradigm is also costly because of expensive software licences, the need for increasingly expensive hardware and a large support staff. So the challenge was to find the right mixture for scarce human and monetary resources, high stability and predictability, and quick and reliable deployment. In Buenos Aires, well-trained Computer Science professionals are hard to find, most of them live and work abroad, while the remaining have stable jobs either at the government or big companies.

3 The FreeBSD solution

3.1 Introduction

At the beginning of 2003 we had a CriticalPath mail system running on Solaris x86 plus a Redhat box for SMTP, Radius and DNS. The DNS and Radius services were constantly down and we were struggling with huge mail queues. There was an attempt to install CriticalPath for Linux into Redhat on an Intel box with a Megaraid card, but the disk latency was enormous and the mail application never really worked.

The first step depicted towards the "FreeBSD solution" consisted in migrating this hardware and commercial software to FreeBSD 4.8 with Linux emulation.

3.2 The choice of FreeBSD

The FreeBSD operating system is well-known for its great stability, plus its pragmatism and common sense to put applications on-line thanks to its excellent Ports System (<http://www.FreeBSD.org/ports>). We consider its release engineering process (<http://www.FreeBSD.org/releng>) to be easily understandable, while the users' community at the official mailing lists keeps a polite and civilized style when it comes to asking for support or reading other people's problems and solutions.

Another important feature is quick deployment. Fortunately, we could state our OS install policy around FreeBSD's great out-of-the-box capability. In a small company you sometimes need to run to a Datacenter and quickly setup a server for some service. In the last two years, Argentina.Com acquired around forty servers, most of them Pentium IV but also several double-Xeons and a few double-Opterons to be co-located in the Datacenters where we have dialup and hosting operation contracts. All of them run FreeBSD, ranging from 4.8 (there are a couple with two years uptime and zero trouble) til currently 6.0-BETA2.

The general policy for the operating system is to try to bring all servers periodically to the stable code branch by using RELENG_4, RELENG_5 and now RELENG_6. This regularity lets us be more prepared regarding possible exploits at the operating system or base software level, especially in web servers.

3.3 Basic re-engineering

The first re-engineering step was to put in place two FreeBSD 4.8 boxes whose unique task was to be authoritative DNS for all our domains. The chosen software was Bind9. Those boxes were co-located in different datacenters, taking care that there was good latency between them to avoid zone transfer problems, and making it possible to deal with TTLs between 60 and 600 seconds to have quicker response in case of trouble.

Second step was to deploy two more boxes of the same class, again in different Datacenters, to only deal with Radius and recursive DNS. The Network Access Servers at the Telcos were configured to send Radius Authorization and Accounting to those servers, and to assign these recursive DNSs to dialup users.

The third "golden rule" never to put SMTP incoming and outgoing in the same servers. We deployed separate FreeBSD boxes with postfix for incoming and outgoing mail.

3.4 Email migration

The email migration required careful planning due to the fact that we were going to migrate both mail front and back-ends. We first built a perimetral antispam and antivirus system in FreeBSD 4.x and 5.x based on postfix,

amavisd-new, clamav and SpamAssassin. These systems were to deliver mails to both the old and the new system until the new back-end was in place. In the meantime, we added small FreeBSD NFS boxes to increase CriticalPath's mailspool, without any problem.

At the frontline of incoming mail, we put in place several MXs of the Argentina.com domain to filter dictionary attacks (attempts to forward mail to nonexistent users) as well as a black-list derived from SURBL that resulted in almost no false positives. The mails are then multiplexed to a cluster of double-Xeons and double-Opterons where we run amavisd-new with MySQL based white and black-listing. We discarded the use of Bayes and Autowhitelisting at the global level because of great quantities of false positives and false negatives. We instead defined a few spam levels going from the least to the most tolerant, each one with cutoff or discard levels. Every email with a score below the one associated with the selected spam tolerance goes to the user's Inbox. Emails between this level and the cutoff level go to a user's folder named Spam, and those above the cutoff level get discarded because it is a very obvious spam. For the sake of simplicity, we transparently associated the use of the Address Book with the antispam system, so that every personal contact gets automatically whitelisted.

With the introduction of Spamassassin 3.x, the DNS traffic to query global blacklists grew considerably, so we signed agreements with SpamCop, Spamhaus and SURBL to install public mirrors of their databases in our FreeBSD equipment. Thanks to these mirrors that cost us between 1 and 2Mbps in traffic, we were able to dramatically cut down Spamassassin latency.

At the 3rd level there is the delivery to the mailedrops. As soon as we started building a new Cyrus-Imap back-end with MySQL authentication, we needed to multiplex incoming mail to users in both old and new mailedrop formats. Finally, we managed to migrate hundreds of thousands of mailspools to the new Cyrus architecture using a great tool named imapsync, which is directly installable from ports. We also put perdition, a POP3 and IMAP proxy, in the middle to assure a transparent migration and distribution of mailboxes across several servers. Briefly, all information of where a user's mailedrop is located resides in MySQL, and is being used by all software pieces in the chain.

With regard to the hardware for disk space, we currently use seven Cyrus-Imap loaded FreeBSD boxes with diverse hardware. The biggest are Pentium IV with 4G of RAM and 3ware cards in chassis with 12 hotswappable bays, organized in 3 RAID-5 units of 1 Terabyte each. The 3ware software sends you an email whenever the RAID is degraded -mostly because of a failing disk- and lets you rebuild the RAID with everything up and running. We use smartmontools in the cases where we have less redundancy, to have immediate alerts of disks with temperature problems or failing selftests.

As webmail software, we chose a commercial product named Atmail, which is available with perl sources and utilizes mod_perl. Under FreeBSD it is extremely easy to deal with perl modules, you do not even need to use the CPAN shell, you just have to choose the right port and run "make install". After several months of integration work, we integrated the Client-only version of Atmail that talks IMAP with our back-ends. We had to modify some parts of the code to adapt the product to our massive free environment, and to our antispam and antivirus perimeter, in addition to our specific customizations and translations.

3.5 Web migration

With the adoption of FreeBSD, there was almost no additional effort necessary to setup a working Apache, PHP and MySQL environment in minutes. Even the upgrades from PHP4 to PHP5 were painless. The ports system was again extremely useful in these cases, and permitted us to do things like compress text and html contents in Apache with just a few lines of documentation. In addition, we have experienced excellent performance and rock-solid stability and uptime.

4 Results

We managed to deploy a FreeBSD based email architecture that is horizontally scalable, using 3 Terabyte Intel based storage servers at a current cost of 3 dollars per Gigabyte with redundancy.

The great stability achieved enabled Argentina.Com to explore other fields like hosting for resellers and housing with presence in three Argentine Datacenters.

We offer now also corporate dialup for roaming users in Argentina and Peru thanks to our presence and contracts with most Telcos. Among our indirect customers, there are major American companies like Ford, Exxon and Reuters. We now run the free dialup business in Brazil, Chile, Colombia and Panama as well.