RIK FARROW

musings

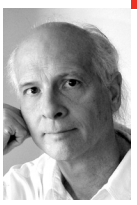

Rik is the Editor of *;login:*.

rik@usenix.org

I CAN SEE DARK, OMINOUS CLOUDS out my office window. It's been unusually dry here, although not nearly as dry as Robert Ferrell's home base, San Antonio. Perhaps the clouds I see will produce some much needed rain.

But it's not rain clouds, or the lack of them, that has sysadmins concerned these days. Instead, it's cloud computing that worries many. Cloud computing appears to be storming over the IT world, replacing local servers with ones somewhere "out there." If cloud computing takes over, many fear another wave of sysadmin job losses.

Appropriate use of cloud computing can save money as well as be more energy efficient. And, since it's the latest buzzword, every boss is wondering when his IT department will move "into the cloud," if only so that he can tell his golfing buddies about it.

I have my own worries about cloud computing, concerns over the security of data that will be stored and processed in the cloud. And I am not alone, either.

HotCloud

The HotCloud workshop summaries are included in this issue. I suggest reading these excellent summaries, in particular the report of the panel discussion in which Stefan Savage discusses some security concerns. Savage, like many others, pointed out that data stored off-site gets different US legal treatment from data stored on premises. A subpoena, something a judge must approve, may be required for access to some data stored off-site (Stored Communications Act [1]). Unless a cloud provider can guarantee that data will not be stored outside of the EU, and particularly not in the US, European Union users cannot use that cloud provider to store any confidential data.

Savage also pointed out that unless you are using Infrastructure as a Service (IaaS), you are relying on the cloud provider for privacy, storage availability, integrity, durability, and retention limits. Savage told of a cloud provider that lost client data, and the client had no recourse for the recovery of that data or for damages due to its loss.

Before I read the HotCloud summary, someone I know asked about the security of cloud computing and I came up with a different set of concerns. First, when you run your own servers, you con-

trol (or fail to control) the physical security of your servers. You have access to network infrastructure, file and backup storage, and servers themselves. Physical security is the base for computer security, and cloud computing turns this over to someone else.

You might be thinking that won't be a problem. After all, cloud server farms do have physical security, and will in many cases be able to arrange for better physical security than your organization could afford. But this brings about another dark idea. You do not get to hire the people running the cloud server farm, including those whose job it is to replace dead servers or drives in the hot, noisy racks.

You also lose the ability to monitor and log network traffic outside the hosted "server." Even if you don't routinely log traffic, you probably have had to do it when debugging a server on which performance suddenly dropped for no apparent reason. Running a tool like Argus [2] to collect connection logs is not only a good debugging tool, but also great for security audits. But in the cloud you have to hope your provider will do this for you. Right. A good cloud provider will keep logs, but sharing them with you will be difficult, because those logs reveal information about other hosted servers.

And your firewall will be included within the server, not outside it. You likely remember the mantra *security in depth*, but you no longer have an *outside* where you can put the firewall. An attacker who can elevate privileges can delete all logs (or perhaps just `rm -rf` everything), and you will no longer have a second source of logs outside the affected server. Unless you wisely have been saving logs locally, and not in the cloud, you will have lost your logs as well.

While few sites perform real forensics after an incident, your ability to look at drives after an incident will be gone too. Even if you want to find those deleted log files, one of the easiest things to do with disk forensics tools such as Sleuthkit [3], your deleted files really will be gone.

Even the little blinky lights on networking equipment that let you know that there are still packets reaching your server will be gone.

Virtual World

Servers in the cloud are hosted with other servers, sitting on top of VMs. The vendors of virtual machine monitor (VMM) solutions do their best to prevent exploits that can escape the boundaries of the virtual machine into the VMM, but it has happened. I just learned of an incident last week where a hosted server was properly secured, but another hosted server on the same hardware was exploited. The attacker then exploited the VMM and wiped all the other systems hosted under it, including the one properly secured. Of course, there are no disks handy where someone could perform forensics and prove this, but my acquaintance has been kicked out by his hosting provider for "attracting trouble."

Most server hardware today runs Intel or AMD processors, and these processors were not designed with virtualization in mind. These processors *do* have extensions to support virtualization, but these are for performance more than security. Real hardware support for virtual machines means that virtual machines are segregated using hardware beyond the memory management mechanisms (MM) used today. MM was designed to segregate processes, not virtual machines, but this is how it is being used today.

I published a column about virtualization security one year ago [4], and that column is still good reading today.

Virtualization certainly has its place. But if you care about the confidentiality of your data or are legally required to provide auditable, secure data store, you should not be moving into the cloud.

The Lineup

We start this issue with an article by Tom Limoncelli about software. Well, not quite, as Tom expounds on design decisions that fail to take into consideration installation, debugging, and maintenance as they affect the system administrators who manage this software.

Christoph Hellwig describes XFS, one of the file systems supported by Linux versions. During FAST '09, I overheard someone asking a Linux vendor why there needed to be more than one file system type, and I felt more than a little embarrassed. Hellwig explains how XFS is different from the default Linux file systems and when it should be used, and he provides some performance graphs to back up his assertions.

Kristaps Džonsons makes a strong case for creating better documentation. Džonsons says that mdoc comes closest to meeting his set of criteria for good documentation. He includes both syntactic regularity and semantic encapsulation, so that machines can interpret data and so that the documentation also works better for its human users.

Brandon Salmon and his co-authors have also written about file systems but take a very different perspective from Hellwig's. Salmon points out that users look at file systems very differently from the way software engineers and sysadmins do. iTunes, for example, groups music in its GUI very differently from how it lives in the hierarchically organized file systems where the data is actually stored. Perspective, their project, leverages semantic information for data management for home systems that includes intelligent file migration and backup.

Tasneem Brutch has written a survey of tools useful for compiling, profiling, and debugging programs destined for multicore systems. Parallel programming is hard, but there are a growing number of tools designed to make the task easier; Brutch does a great job of covering available tools, both open source and commercial.

Rudi Van Drunen continues his hardware series by discussing the trouble with static. Did you know that before you can even feel a static discharge the voltage has reached 3,000 volts? Rudi explains the sources of static, graphically shows the effects of static when frying microscopic circuitry, then covers countermeasures.

David Blank-Edelman completes his exposition of the Perl Web application framework, CGI::Application, begun in the August issue. He certainly makes things look easy.

Pete Galvin has written an extensive comparison of two new virtualization systems, VMware vSphere and Microsoft's Hyper-V (release 2). Pete, who has previously compared different Solaris-specific forms of virtualization [5], does a thorough job of comparing these two new offerings.

Dave Josephsen reveals a solution to authorized access using OpenVPN, OpenLDAP, and PF that he built in-house with a coworker. You can find the sources for the glue that makes this very cool system work in the online *;login:* at http://www.usenix.org/publications/login/2009-10/.

Robert Ferrell appears to be just as fond of cloud computing as I am, but has a very different manner of expressing his feelings.

Elizabeth Zwicky has reviewed the second edition of David Blank-Edelman's *Automating System Administration with Perl* in her usual style. She also describes a second book, on leadership, as "mostly painless." Then Dave Josephsen covers a book on the Android programming environment, followed by Brandon Ching on a book on OpenSolaris.

For summaries, we begin with the 2009 Annual Technical Conference, followed by the excellent summary of HotCloud written by Alva Couch and Kiran-Kumar Muniswamy-Reddy. Finally, we were fortunate to get some summaries from BSDCan, compiled by Royce Williams.

Those dark clouds I was watching never did produce any rain, unfortunately. And I suspect that everyone rushing into cloud computing will look back in one or two years and wonder why they were so eager to put most of their IT infrastructure, and precious data, into the cloud.

REFERENCES

[1] EFF on the Stored Communications Act: http://ilt.eff.org/index.php/Privacy:_Stored_Communications_Act.

[2] Argus network audit and analysis: http://www.qosient.com/argus/.

[3] Sleuthkit and Autopsy open source forensics tools: http://www.sleuthkit.org/.

[4] Rik Farrow, "Musings," *;login:*, October 2008, vol. 33, no. 5: http://www.usenix.org/publications/login/2008-10/openpdfs/musings.pdf.

[5] Peter Galvin, "Solaris Virtualizations," *;login:*, April 2009, vol. 34, no. 2: http://www.usenix.org/publications/login/2009-04/pdfs/galvin.pdf.