A spectre haunts fintech

Frank Stajano warns of computer security risks lurking in the cloud and of other systemic threats that are rooted in the very design of the processor

Cloud computing is the practice of relocating data storage and computation from an organisation's own premises to general purpose servers, which an external operator runs somewhere on the internet. The main advantages include scalability, making the data accessible from anywhere and dumping on someone else the headaches associated with maintaining and upgrading the hardware.

But what about the security of such an arrangement? Traditionally, the three fundamental security properties are confidentiality, integrity and availability. Compared with in-house computing, cloud computing offers higher availability but arguably lower confidentiality and integrity, because the cloud provider is in a position to read and potentially even alter your data.

Webmail is an example of a cloud service that has security issues. The prevalent attitude of the public is one of trust: one thinks that the webmail providers are "big names" such as Microsoft, Yahoo and Google and they are bound to be well-behaved – they would stand to lose too much face if they did otherwise. There is some truth in this but a security expert will have a more cynical attitude. As Robert Morris, the late National Security Agency chief scientist and Unix security architect, pointed out: "A trusted party is one that can violate your security policy."

A crucial issue when talking of cloud security, therefore, is to have a clear idea of whether you are willing to trust the cloud provider; in other words, whether you have a good reason for paying a price, both in inconvenience and in cost of technical countermeasures (such as encryption), to guard against hypothetical attacks from the cloud provider. Note also that, if you are concerned about the confidentiality of your documents in cloud storage, you may assess the provider itself as trustworthy but not, for example, the government of the foreign country under whose jurisdiction its servers reside. Malign state actions could mean that your otherwise trustworthy provider may be compelled to break your confidentiality – and keep quiet about it.

Regardless of whether you trust the provider, and accept the risks that brings, you still have to consider third-party attackers: they will exist even if you run your own servers in-house. When it comes to external attackers, it is fair to say that an established cloud provider, with its dedicated team of security engineers and continuously monitored anomaly detection systems, is likely to have much stronger and more up-to-date defences than most SMEs. The greater investment in security for the large cloud provider is an economy of scale, amortised across all its clients.

But, even when security has been outsourced to a company for whom it is a core service, most systems have a glaring weak point: the password. In the all-too-frequent case that access to the cloud resources is controlled by a password, that password can be guessed, eavesdropped, brute-forced and so forth, or the reset mechanism can be gamed. Passwords remain a liability and a known weakness, and we ought to be looking for alternatives.

But another known weakness has emerged and is causing concern. The recently discovered vulnerabilities nicknamed Spectre and Meltdown affect almost every processor manufactured in the past two decades. Unlike most other known vulnerabilities, which reside in the application layer, the operating system layer, the protocol layer, the network layer,

the crypto layer and so forth, these are rooted in the very design of the processor. It is a systemic issue – a design flaw rather than an implementation flaw – as shown by the fact that it affects so many independently developed processor models and makes.

Modern computer systems are too complicated for anyone to comprehend at all levels, from semiconductor physics all the way up to internet applications. The way engineers handle that complexity is by breaking up the system into layers of abstraction. Even experts can only be experts at some of the layers. They rely on the abstraction for what is below and above. So processor designers aiming for greater performance have continued to offer the abstraction of a simple processor executing a linear sequence of instructions, while actually designing and building much more complicated hardware that, in its spare time, executes future instructions "just in case". This so-called "speculative execution" technique is responsible for dramatic increases in processing speed, but it has been exploited to read memory locations that should not have been accessible.

This was done by observing subtle side effects of the instructions that were only speculatively executed and then officially discarded. The fact that they had still been executed had left a faint trace, in the form of leftover data in the processor's cache, that affected the execution time of certain subsequent instructions. A full explanation would take several pages. The relevant point is that a malicious program exploiting these vulnerabilities can read areas of memory that it would not normally be allowed to read. This violates the operating system promise of "process isolation". The malicious program, if running on the same physical hardware as your program, could read your program's private data.

Cloud computing is advantageous because of economies of scale: a cloud provider consolidates the data centres of thousands of large customers and hundreds of thousands of smaller ones onto its large, distributed cluster of powerful machines. Each physical machine is dynamically partitioned into a large number of virtual machines, so as to exploit the hardware most efficiently.

Unfortunately, Spectre and Meltdown can "melt down" the barriers between different virtual machines running on the same processor. Unlike higher-level vulnerabilities, these cannot easily be fixed by a software patch because the flaws are not actually in the code but in the processor that executes the code.

The security patches that are available take the shape of workarounds that limit the use of speculative execution and caching (at some cost in performance) to avoid running into the problem. The underlying vulnerability still remains and will only disappear when the processors are replaced with redesigned ones that address the root cause. This will not be easy. Most of the performance gains that have been achieved by processor makers, once it was no longer physically possible to keep doubling the clock speed every couple of years, have been based on more and more aggressive parallelisation, while keeping up the abstraction that the code written by the programmer would execute sequentially.

> *"The malicious computer program can read areas of memory that it would not normally be allowed to access*

Although it is conceivable that speculative execution may be cleverly redesigned so as not to be vulnerable to Spectre and Meltdown, this will still require a radical rethink of processor architecture to achieve performance levels that customers have been used to.

Until all vulnerable processors are replaced by new, non-vulnerable ones, which have yet to be designed, a process that will take years, cloud providers will naturally be assuring people that their systems have been secured. For a fintech company with its crown jewels in the cloud (as in: unauthorised access to them would kill your business), it may be worth taking a belt-and-braces approach and insisting not only that the cloud provider install all the latest patches as a matter of course but also that it runs your processes on hardware dedicated solely to you, rather than on virtual machines co-hosted on hardware you share with other cloud customers.

This will cost you more. Whether in-cloud or not, you will also want to isolate your back-end from your web server. Even more importantly, you will want to re-read the fine print on who is liable should anything go wrong in the future because of a Spectre-like attack. ■

Frank Stajano *is professor of security and privacy in the department of Computer Science and Technology of the University of Cambridge, where he is the head of the Academic Centre of Excellence in Cyber Security Research. He is a founding director of two cyber security start-ups, Cambridge Cyber and Cambridge Authentication, and a founder of the Inter-ACE and Cambridge 2 Cambridge cyber security competitions*