Security of Cloud Computing

James Davenport [1]

University of Bath & BCS Council

26 March 2012

[1] Thanks to Ali Kaafarani for the briefing

How secure is cloud computing?

"How secure is $\not{c}\not{l}\not{o}\not{u}\not{d}$ computing?" — the two questions are essentially equivalent, and as unanswerable

It all depends on the application!

Furthermore, we must never forget that security is a "whole system" property

"A chain is as strong as its weakest link" is a common aphorism, but misleading in security

"A chain is no stronger than its weakest link" is more accurate

[1] combines industry-standard RSA with industry-standard TCP/IP, and produces a system which is very weak, and increasing the key length makes it weaker

[5] shows that VirtualBox (100%) and VMWare (30%) allow one to break **correct** TLS implementations due to bad interaction of virtualisation and random numbers

With cloud, do you know the whole system, and will you even be told when it changes? See Holt's talk

Never forget the CIA!

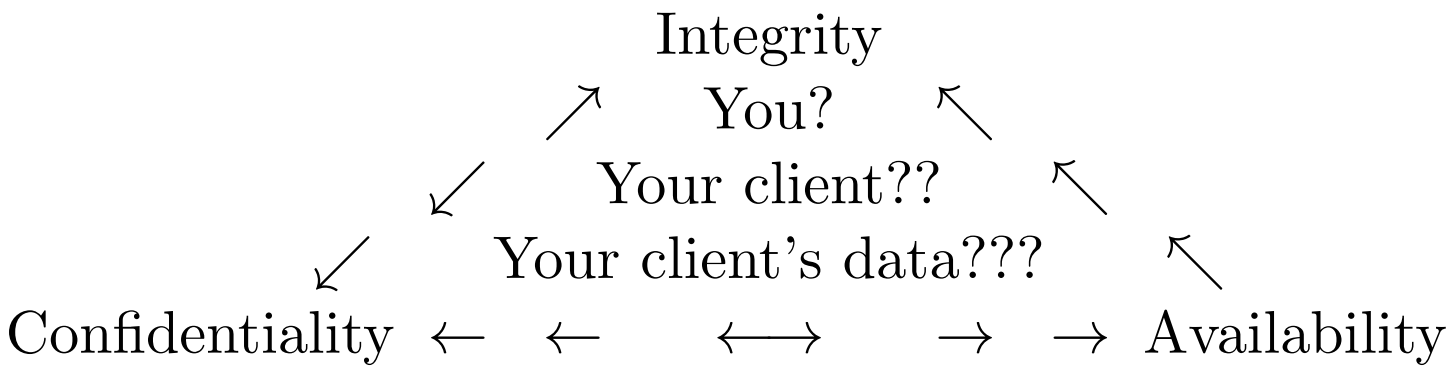

Any two are trivial:

- CI A locked box
- IA CD-ROM in the post to everyone.
- CA A hardware random number generator

But we want a combination.

Some views, and refutations

- No-one knows where my data are in the cloud

X [4] show that, *legally*, they can achieve 40% probability of co-residence with a named target user of EC2

$\surd$ This is, technically, no worse than 1970s bureaux, of course

- These clouds are built with the best technology, and staff

?$\surd$ Quite likely true about the staff

X But security is a whole system property, and his staff aren't responsible for your application

XX The technology of the components is no guarantee

A broad categorisation

But never forget that the whole system needs to be considered

AaaS	Application as a service
SaaS	Software as a service
STaaS	STorage as a service: DropBox etc.
PaaS	Platform as a service
IaaS	Infrastructure as a service

For this purpose, I (like Ben Halstead) don't see much difference in the last two

Let's consider a not-quite-hypothetical analogue:

GaaS	Gaming as a service

GaaS: Gaming as a service

- Many people like computer games as recreation
- But *homo sapiens* is pretty competitive — league tables etc
- And *homo sapiens* is (sometimes) pretty social — chat rooms etc.
- This is all good recreational stuff,
- So let's build a service round this, and charge for it

We've just invented the Sony Playstation Network
(and probably its leak of 11 months ago!)
Some peope would argue this wasn't "cloud computing", but

- Sony stored the data on systems customers could access (surely a cloud) (2)
- It was apparently hacked via a cloud (3)

AaaS: Application as a service

- Has been around since the payroll bureaux of the 1960s
- "No-one got fired for buying IBM"
- "No-one got fired for using Centrefile" [Bath 1966–2005]

But Centrefile had their own (or NatWest's) computers, under their control, in a secure location ...

? Can you make the same guarantees about today's AaaS vendors?

Ben (quite rightly) encouraged these start-ups to go for cloud hosting

SaaS: Software as a service

- BrightPearl: "High availability and fault tolerance — people might be staking their careers and businesses on your software, it needs to be available and reliable"
- JHD would argue "If people are staking their careers on it, it had better be secure as well"
- Where in the CIA triangle is this service?
- The answer is service-dependent!

STaaS: Storage as a service

- Not commonly referred to as such, but
- Many examples (DropBox, . . .)
- Technically a troubling paradigm: [3] shows that common theoretical techniques don't apply, and that strong hashing techniques can break in this setting
- The business model (client-side deduplication) is hard to reconcile with privacy [2]

S/STaaS: Stored Serviced Documents

Again, not a paradigm that's been investigated

- Most people's use of GoogleDocs etc.
- It's the storage and sharing that makes it worth while, not the massive features of the word processor :-)
- The "credential sharing/authentication" mechanism, i.e. via existing Google accounts, is *probably* much more susceptible to social attacks than technical ones

P/IaaS: Platform/Infrastructure as a service

- This tends to underpin the others
- "Trivial" differences, e.g. which virtualisation system, can have massive effects on system security [5]

But The user of AaaS/SaaS probably has little control over these

And Typical contracts make no provision for metadata

Challenges of Forensic Investigations in Cloud Computing

Cloud Service Providers artificially make the cloud opaque, for commercial confidentiality and other reasons
The traditional SAP (Securing, Analyzing, Presenting) Cycle, working off a bit-wise copy, is not possible any more
The CSP will have log data that is not accessible to you, particularly in an SaaS context
In IaaS, one should think that the user could install forensic tools, but, at least in EC2, if the machine shuts down, the [forensic] data is lost (this has happened in real life)
The Amazon SLA (and others) offer no help. Hence cloud forensics is basically impossible. If you don't trust the CSP, then leave the cloud!

Data and Metadata

- The "[forensic] data" is in fact metadata — data about our data
- An in-house service provides metadata automatically: we need to think about metadata in the cloud
- It's not until things go wrong that we worry about metadata
- This distinction needs to be clearly made: who owns the metadata?
- "date of last access", "userid of last accessor" ...
- While computation on encrypted data is a theoretical toy,
- computation with encrypted metadata may be more significant
- VaaS (Voting as a Service) is *all* about metadata: anyone can add up votes!

Conclusions

- Security is a "whole system" property

- and this is the underlying, **good**, reason for uneasiness about Cloud computing
- The technology does not exist to do good verification of a cloud application handling sensitive data
- Metadata are a very unclear area
- The contractual provisions are not in place for sensible forensics

Paper Do you scrap, or feel obliged to shred?

Cloud If I felt obliged to shred the paper, I'd be very uncomfortable putting the data on an external cloud

Links

1. http://people.bath.ac.uk/masjhd/Meetings/Cloud2011.pdf
2. http://www.theregister.co.uk/2011/04/26/sony_playstation_network_security_breach/
3. http://www.bloomberg.com/news/2011-05-13/sony-network-said-to-have-been-invaded-by-hackers-usin html

References

P.A. Crouch and J.H. Davenport.

Lattice Attacks on RSA-Encrypted IP and TCP.
In B. Honary, editor, *Proceedings 8th. IMA Conf. Cryptography and Coding*, pages 329–338, 2001.

D. Harnik, B. Pinkas, and A. Shulman-Peleg.

Side Channels in Cloud Services: Deduplication in Cloud Storage.
IEEE Security & Privacy 6 / DOI: 10.1109/MSP.2010.187, 8:40–47, 2010.

T. Ristenpart, H. Shacham, and T. Shrimpton.

Careful with Composition: Limitations of the Indifferentiability Framework.
In *Proceedings Eurocrypt 2011*, pages 487–506, 2011.

T. Ristenpart, E. Tromer, H. Shacham, and S. Savage.

Hey, you, get off of my cloud: exploring information leakage in third-party compute clouds.
In *Proceedings 16th ACM conference on Computer and communications security - CCS '09*, pages 199–212, 2009.

T. Ristenpart and S. Yilek.

When Good Randomness Goes Bad: Virtual Machine Reset Vulnerabilities and Hedging Deployed Cryptography.
In *Proceedings ISOC NDSS 2010*, 2010.