dri™
The Voice of the Defense Bar

Data and Security Dispatch

The newsletter of the
Cybersecurity and Data Privacy Committee

6/19/2018

Volume 4, Issue 1

Committee Leadership

Chair
James H. Kallianis, Jr.
James H. Kallianis, Jr.
Hinshaw & Culbertson LLP

Vice Chair
Alexander E. Potente
Clyde & Co US LLP
San Francisco, CA

Publications Chair/Newsletter Editors

Heyward Dodkin Bonyata
Nelson Mullins Riley & Scarborough
Columbia, SC

Wendy B. Degerman
Nelson Mullins Riley & Scarborough
Columbia, SC

Click here to view entire Leadership

In This Issue

Leadership Note

Back to Contents

From the Vice Chair

By Alexander E. Potente

Welcome DRI cyber liability readers! We have a great newsletter for you. We have three interesting articles on very current topics. The first involves AI and the legal profession and the very important question of whether we will have jobs in a decade: "Can Machines Think Like Lawyers?"   The second, "Three Ways that Counsel Can Assist Defense Contractors to Achieve Proactive Compliance with the Department of Defense's Newly Effective Cybersecurity Requirements," addresses new cybersecurity requirements for defense counsel and ways in which we as lawyers add value to that compliance process. The final article, "Seeing Enterprise IT in the Clouds," addresses outsourcing enterprise IT and ways it has become integral to our practices. Furthermore, as I hope all of you know, planning for the DRI Cybersecurity and Data Privacy Seminar is well underway. Please save the date for it in Chicago on September 6–7, 2018. We look forward to seeing all of you there!

Alexander E. (Alex) Potente, partner of Clyde & Co US LLP in San Francisco, is an experienced trial lawyer who represents insurers in complex commercial insurance litigation matters including disputes pertaining to general liability and professional liability policies, with an emphasis on bad faith litigation and coverage issues arising from claims involving class actions, product defects, public sector liability and environmental and other long-tail insurance coverage disputes.. Alex currently is the vice chair of DRI's Cybersecurity and Data Privacy Committee.

Can Machines Think Like Lawyers?

By Daniel S. Marvin

Can machines think? That question was first posed by famed British mathematician Alan Turing in 1950, who postulated that at some point in the future, computers would be able to be trained to act like humans. While computer scientists are still pondering Turing's question, there is a far more relevant question for us to consider: can machines be trained to think like lawyers? After all, when performing legal analysis, lawyers employ a host of human-based analytical tools such as intuition, reason and emotion. Can that be replicated by computers? And if so, can lawyers leverage such technology to their advantage? We can look to the field of computer science known as "Machine Learning" in an attempt to answer those questions.

Machine Learning is a subcategory of artificial intelligence where computers are able to "learn;" that is, they are able to use complex algorithms in order to form predictions based on data that is fed into their system. While the concept has been around since the time of Turing, it is only the past few years that have brought us computer processing powerful enough to analyze the enormous amounts of data that are being created and attempt to answer his question. Examples of Machine Learning are all around us in our everyday lives. For example, Amazon utilizes Machine Learning to predict what items you are likely to buy based upon your (and others) past purchases in order to suggest future purchases; Netflix uses Machine Learning to predict what movies to suggest to you based upon your viewing habits; Uber also uses Machine Learning to predict pick-up times and traffic patterns; and perhaps the best example, Spam filters use Machine Learning in order to predict what email messages are junk and should be kept from your inbox.

These are examples of "supervised" Machine Learning, meaning that the computers learn based on a user-provided data set. Over time, as more and more samples are entered into its database, the computer will get smarter. With spam filters, for instance, email providers feed in known spam emails, such as those containing the phrases "online pharmacy," "make money quick" or "looking for a date?" An algorithm will then not only flag similar emails as Spam, but also look for additional patterns amongst the sample set, such as if they originated from common IP addresses or locations, and then flag those messages as Spam as well. The system will then add those emails to its learning set, and over time, continue to get smarter and be able to all-but eliminate Spam from a user's inbox. Lawyers are already utilizing similar technology to code documents during discovery. As relevant documents are coded, Machine Learning not only filters out documents with similar coded data points, but also sees patterns amongst those documents that perhaps the reviewer missed. In just hours, Machine Learning can cull down millions of documents to a relevant subset that would take an army of lawyers months to discover.

A more interesting aspect of lawyers' use of Machine Learning is can be seen in case modeling. When a matter arises, clients rely on attorneys to analyze both facts and law, and arrive at suitable risk-reward based recommendations. In turn, attorneys rely on their experience, including familiarity with the law, juries, judges and opposing counsel in order to arrive at reasoned recommendations, such as whether to settle a case or take it to trial. The more experience an attorney has to call upon, the better and more systematic of a recommendation can be made. Attorneys must be able to synthesize all of the data they have gleaned through their experience in order to evaluate potential liability, likelihood of settlement, amount of a likely settlement and litigation costs. But even the best attorneys have a limited data set on which to rely. That is where Machine Learning comes into play. With the technology, attorneys are able to create data-driven predictive analysis using not only historical information from their firm, but also from published opinions. Once the data set is created, algorithms can comb all of information contained in it and arrive at case predictions based upon past occurrences.

Let's envision a law firm that has a medical malpractice defense practice. With Machine Learning, such a firm could enter data concerning their various matters into a centralized system. This data could include information such as the type of doctor who is a defendant; the nature of the alleged injury; the age, occupation and other statistics on the injured person, including pre-existing conditions; the initial demand; whether it settled and for how much; the stage of litigation at which it settled; if the matter went to trial and what the jury verdict was; information about the court and judge; and information on any unique legal issues. Using that data set, a Machine Learning system can evaluate the information and perform analytics in order to predict the outcome of future cases. As with a Spam filter or shopping patterns, the larger the data set the better the predictions. A system with thousands of cases worth of data can see hard-to-detect correlative relationships in-and-amongst the data. While attorneys inherently engage in a similar analysis when evaluating a case, a computer program can analyze immense volumes of data and connect certain dots that even the sharpest legal minds might miss.

The potential of such a system is assisting attorneys and improving client outcomes should not be underestimated. For example, algorithms with a large enough data set might notice that defendant doctors of a certain level of experience are 24 percent less likely to have a jury verdict rendered against them; or it may determine that a settlement agreed upon within the first three months of a litigation would likely be 15 percent lower than those in months three through six; or it might see that injured parties who were out of work due to their injury for six months or more were 22 percent less likely settle a case prior to trial than those who missed no work; or it might recognize that a certain judge or forum is more or less hostile to defendant doctors. These are just a few illustrations of the types of conclusions that could be drawn out using Machine Learning, and counsel can use this data to assist their clients in making informed decisions and determine at what point in a litigation is the optimal time to settle, or if settlement should not be considered at all.

At the end of the day, while computers can conceptually outdo attorneys with their sheer ability to compute huge swaths of data, they likely will never be able to call upon gut-instincts or read people the way humans (and lawyers) can. So while Alan Turing's vision may never be completely realized, lawyers need to begin to at least think about how Machine Learning can be incorporated into their practices in order to both improve results and enhance efficiency. As technology continues to improve and becomes cheaper, clients will to expect that these tools be implemented in order to give them every possible advantage, and lawyers should be ready to meet those expectations.

Daniel S. Marvin is a partner in the New York City office of Morrison Mahoney, where his practice focuses on data privacy, data security and cyber-insurance matters.

Three Ways that Counsel Can Assist Defense Contractors Achieve Proactive Compliance with the Department of Defense's Newly Effective Cybersecurity Requirements

By Ty Dedmon and Niya McCray, CIPP/US

Although the Department of Defense (DOD) has long required its contractors to provide "adequate security" to protect "Covered Defense Information," beginning on January 1 of this year, the Department specified that "adequate security" means compliance with all 109 of the security controls described in NIST 800-171. See Defense Federal Acquisition Regulation Supplement (DFARS) 252.204-7012. These requirements apply regardless of a contractor's size or amount of business with the government. Failure to meet these standards can result in legal actions against the contractor (for breach of contract or under the False Claims Act) as well as termination, suspension, and debarment from federal programs. While many of the required security controls are highly technical, this article will discuss a few basic ways that counsel to a defense contractor can provide substantial value in a client's efforts to meet the NIST 800-171 standards.

Breach Response Planning

In the wake of growing concerns over potential data breaches, the DOD has tightened requirements for its contractors and the ways that they implement protocols and respond to data incidents. The new DFARS clause introduces a 72-hour reporting deadline for cyber incidents, while also introducing additional handling procedures like the submission of malicious software in accordance with a contracting officer's direction and the preservation and protection of images of affected information systems.

The DFARS cyber rules are not based on the question of "if" an incident will occur, but rather "when" it will occur and how contractors can best prepare for the road to recovery. The best practice for government contractors is to update their plans to reflect the more specific DFARS requirements. The initial hours following a data breach are the most crucial. Contractors should already have an established set of protocols and plans that they can immediately enact upon discovery of a data incident. The first step should be to protect/privilege the data breach investigation to allow for a free flow of information between key players. From the onset, contractors should have their in-house or outside counsel assessing the facts and determining the potential risks and liabilities they may face.

Next, there should be an immediate establishment of exactly what types of data and how much data have been affected by the breach. Contractors should already have a cyber-forensic team and additional technology experts on retainer. As a general practice, it is best to negotiate those agreements before a breach occurs so that there is no artificial pressure or unfair leverage created by seeking help in a time of crisis.

The final considerations that contractors should keep in mind when updating and applying data incident responses is how best to communicate during the event of a breach. Although DFARS institutes a 72-hour reporting requirement, contractors must consider whether they will need to provide additional disclosure to customers, state attorney generals and/or legislators, employees, the press, and, in some instances, law enforcement. These are highly-complex determinations that can change based on a number of factors – legal counsel is essential to helping contactors formulate an appropriate plan for their organization. The types and contents of these communications should be prepared well in advance and they should also be ready to transmit within a reasonable time following the incident. Beyond the channels of communication, though, contractors should also contemplate business continuity plans that will allow them to maintain essential functions despite the disruption of certain platforms and applications. Nevertheless, attorneys should make sure that their contractor clients have assembled a team capable of making the best business and legal decisions as the incident unfolds and process of investigating, responding, and recovering begins.

Vendor Management

Another concern of the DFARS rules is access control. The DFARS clause requires that primary contractors "flow" the clause down to subcontractors at any level who are involved in the processing of covered defense information.

Covered defense information ("CDI") means unclassified controlled technical information or other information, as described in the Controlled Unclassified Information (CUI) Registry, that requires safeguarding or dissemination controls pursuant to and consistent with law, regulations, and government-wide policies. The information must be marked as CDI and/or processed in support of the contractor's performance of a government contract.

Essentially, the DFARS flow-down requirement compels both contractors and subcontractors, alike, to provide adequate security pursuant to NIST 800-171. In practice, primary contractors are tasked with vendor management, making sure that subcontractor work is performed on compliant systems. Some ways that a primary contractor can tackle this complex task include: engaging in direct communications with the subcontractor about the specific requirements of DFARS; conditioning subcontract work on the provision of evidence that subcontractors have engaged in a full NIST 800-171 security assessment and have developed, updated, and/or implemented security plans to remediate any shortcomings; or, providing assistance to subcontractors to ensure, firsthand, that they are in compliance with DFARS. Creative and informed legal counsel can be a contractor's best weapon in negotiating downstream contracts that reduce a contractor's compliance risk due to failures by a subcontractor.

While there are several different approaches to vendor/ subcontractor management, it is incumbent on the primary contractor to decide which method is the most feasible based on the extensiveness of the subcontractor's role to the contract. Contractors should also consider the types of CDI that respective subcontractors will be handling. Given that the prime contractor is ultimately liable for any violations of the DFARS rules, contractors should be wary of subcontractors who are lax in their cybersecurity or those who are completely unversed to the DFARS and NIST control requirements. It is virtually impossible to exclude subcontractors altogether; however, the addition of a non-DFARS compliant subcontractor could lead to unnecessary liability in the wake of a data incident.

Employee Awareness and Training

Lastly, attorneys can assist their clients' compliance efforts by addressing every contractor's highest cybersecurity risk: humans. While software can be updated and systems patched, employee carelessness can only be mitigated by repeated efforts to train the entire organization on sound security practices. This area of risk is so significant that NIST 800-171 devotes an entire family of controls (3.2) to "Awareness and Training" of system users. Counsel to defense contractors should, at minimum, become conversant in the most common types of attacks targeting employees, including phishing, malware, and social engineering. However, breaches commonly occur without instigation by a third-party – misplaced or lost laptops and phones are a risk area that must be addressed through employee training and hardware policies. In addition, counsel should work closely with a client's human resource department to ensure that disgruntled or departing employees cannot remove covered defense information from the company's systems.

While this article highlights select areas of cybersecurity compliance for defense contractors, the NIST 800-171 standards are far more comprehensive. In addition to these security controls, attorneys advising defense contractors should be mindful that the specific agreements between the DOD and its contractors may provide more specific compliance and certification obligations (including an obligation for contractors to self-certify their compliance or seek accommodations for areas where they are not yet compliant). As with many complex business problems, contractors can benefit from the perspective and protection of legal counsel and a privileged deliberation process for their compliance strategies. Both the reality of today's data-driven business environment and the DFARS regulatory requirements mean that defense contractors must be proactive in assessing and mitigating their cyber risk—parties who are purely reactive in addressing data issues are only preparing to fail in these critical obligations.

Ty Dedmon is a partner with the law firm of Bradley Arant Boult Cummings, LLP where he focuses on all areas of healthcare litigation and data management. Ty can be reached at tdedmon@bradley.com. Niya McCray, CIPP/US is an associate with the law firm of Bradley Arant Boult Cummings, LLP where she practices business litigation and cybersecurity/data privacy. Niya can be reached at nmccray@bradley.com

Seeing Enterprise I.T. in the Clouds

By Adam Cohen, CISSP CEH CCSP

Legal scholarship based upon an abstract concept of cloud computing (the “Cloud”) is abundant. However, lawyers should focus on the *facts* about the cloud services to which their clients are migrating. Amazon (Amazon Web Services or AWS) and Microsoft (Azure) are leading a movement rapidly replacing traditional I.T. environments. Lawyers need a solid understanding of these services and their impact on legal practice and process.

AWS and Azure Lead the New Cloud Order

Enterprise customers of AWS, Azure and competing platforms from other galactic tech giants, such as Google Cloud, enter a world of super-technology and buzzwords—*serverless*, *containerized*, *orchestrated* and *hyper-converged*, *software-defined* with “*infrastructure as code*.” Legal advisors new to this cloudscape may be startled by the strange jargon, unfamiliar contracts, seemingly infinite I.T. power and unique pricing structures. Their minds will spin with questions about security, e-discovery, regulatory compliance and more.

The clients flocking to AWS and Azure adopt a model where I.T. resources are treated like a utility, with payment metered to use. Resources include fundamentals of I.T., like storage or computer processing capability, delivered through services offering menus of optional capabilities and characteristics, each with corresponding cost, performance and information security implications. Providers like AWS and Azure are offering capabilities like artificial intelligence with voice recognition and response (*e.g.*, Alexa), mobile device farms to use for testing purposes (*e.g.*, for mobile app developers) and streaming analytics on real-time data (*e.g.*, video game developers can adjust online gaming environments as play unfolds).

These tech titans also curate a marketplace in compatible systems and services from independent providers (like the “app store” on your phone but enterprise cloud style). These offerings promise to make the cloud experience easier, faster, more resilient, secure, compliant, efficient, capable, versatile, scalable, transparent, elastic, accessible, less expensive and otherwise better in every way. This eco-system encapsulates the constant acceleration of change in technology, as each new development engenders multiple offspring.

The huge scale of the infrastructure the major providers like are building, aggressive competition and technological developments improving have benefited customers in the form of decreasing prices tied to precise units of time (seconds), volume and other such pricing structures. AWS dominates the market. However, AWS cannot rest on its laurels in the enterprise sector, with Azure wielding the inherent advantage of the Windows corporate client empire provides. In greater numbers every day, businesses run critical applications in AWS and Azure. Many businesses also store their most sensitive customer data with these providers.

A Preview of Modern Cloud I.T.

The basic storage services offered by AWS and Azure provide an amuse-bouche. Clients interact with them through application programming interfaces (APIs), which allow systems of different types, which use different languages, to communicate and interact (*e.g.*, a mobile shopping app and a payment system regarding a customer/user order). APIs are not unique to the Cloud; they are also used internally to enterprise I.T. environments to accomplish cross-system interactions. In the Cloud, the API is a point of information transfer between customer and provider, which makes API security particularly important.

The basic storage service in AWS is S3 (“simple storage service”), while Azure has “Blob.” S3 can be used for different data and a wide variety of uses. It can handle virtually unlimited data and numbers of concurrent users. However, unlike most systems lawyers are used to, S3 does not use files and folders. Instead, “objects” uniquely identified by “keys” are stored in “buckets.” While objects and buckets are analogous to files and folders, they are not the same. Such object storage is used in the Cloud because of performance advantages, including scalability and ease of replication. The details can be important for security, as well as critical legal process functions like search and retrieval.

An unlimited number of objects, each up to 5TB, can be stored in a bucket. According to AWS, the data portion of the object is “opaque” to AWS. The metadata includes system-generated, default metadata such as last modified data, but the user can also provide “custom” metadata when the object is stored in S3.

Blob also stores "object" data, but a blob is a "*file* of any type and size" (emphasis added). Blobs are stored in "containers" (not buckets) and come in three flavors for different uses: "block" blobs for files like documents and media (each block blob can contain a maximum of 50,000 blocks, each not more than 100MB, which means the maximum total size is slightly above 4.75TB); "append" blobs for logging (maximum of 50,000 blocks of not more than 4MB, total size just over 195GB) and "page" blobs for efficiency where frequent read/write operations are performed, for example, Azure Virtual Machines uses page blobs as OS disks (up to 8TB maximum).

These are merely a few tidbits of information about a basic resource, storage, but already there is a new vocabulary. Understanding the full range of *hundreds* of types and variations of services, even for *one* of these providers (especially AWS) is a major undertaking for anyone. Moreover, this environment is always changing—fast.

Virtualization, Aggregation, and Legal Risk

Certain legal risks arise directly from the superpowers granted by these services. The systemic efficiencies that enable them are rooted in the magic of virtualization. Virtualization allows different systems, such as operating systems (OS), to share the same hardware resources. A simple example of the traditional use of virtualization would be to run the Windows Operating System on your Apple computer. In the language of the Cloud, running virtual machines from different customers on the same physical hardware is "multi-tenancy."

More recently, cloud use of a technology called containerization has exploded. With containers, applications are packaged with all of the accoutrements that allow them to run on different hosts, except without their own OSs. Instead, multiple containers share resources of an underlying OS. This means that memory resource demand is substantially less than with virtual machines. There are security implications to this arrangement, however—access to the OS can potentially provide access to all containers sharing its resources. Virtual machines, in contrast, are protected by another layer of isolation, independent actors with their very own operating systems (and the accompanying memory resource demands).

You can keep other tenants of the cloud landlord off your computerized lawn, but only by retreating somewhat to the traditional model where expensive resources sit idly waiting for a surge. The major CSPs are masters at automated shifting of "workloads" dynamically within their massive, global infrastructure to assure maximum efficiency. But this becomes optimized at the scale of the tech titan providers only where many customers, separated virtually, are co-habiting physically. Multi-tenancy raises at least the concern of attacks that utilize access to shared resources to hop over virtual fences ("VM hopping" or "VM escape").

Virtualization and containerization are part of the movement towards "software-defined everything" in I.T. This is the "mind-body" relationship for enterprise I.T.; the mind is the software and the body is the hardware:

> The movement towards a software-defined infrastructure is about decoupling the bare metal that executes the point data transactions from the software layer that orchestrates them. The hope is that by separating the smarts from the brawn, the underlying hardware can become cheaper and interchangeable (avoiding vendor lock-in) while the overarching software becomes more capable and faster-evolving...Rather than individual elements (compute, storage, and networking), infrastructure will be treated as a set of resources required for specific workloads. In this worl d, the application, the end user, and hopefully the business are king.

See https://www.wired.com/insights/2013/05/are-you-ready-for-software-defined-everything/.

The buzzword for the "cheaper and interchangeable" underlying hardware is "hyper-converged." The use of virtual machines and containers, as well as the automated management of such resources, called "orchestration," is what allows the "decoupling" the author describes.

More recently, the "serverless" or "Function-as-a-Service (FaaS)" approach has burst onto the cloud scene. Introduced by AWS with its Lambda service and by Azure as Functions, it enables a customer to run code without managing server systems or applications. Of course, there is a server somewhere, but not in the usual way as a dedicated, persistent machine standing by to process requests. AWS and Azure can spin servers up or down in milli-seconds, which doesn't look persistent (thus, "serverless").

Getting this magic to work for real is more complicated. Conceptually, components of the tasks the server performs are split up and distributed between the client and third parties like AWS or Azure, triggered based on programming. There is no need to worry about whether the server instance can handle the number of users; scaling is handled automatically behind the curtain and the customers are charged only for the time the code is running, not for the server to run it. There potential use cases typically involve a compute operation triggered by a condition or event, with unpredictable spikes in traffic. Lawyers, are you ready for serverless e-Discovery?

Attack Surface Everywhere

In considering the dominance of the two largest cloud service providers, it is hard not to worry that the unprecedented degree of aggregation of customer systems and data under these providers presents a potential target of irresistible proportions. AWS and Azure are at least among the most sophisticated organizations ever when it comes to cyber-security, and they have and deploy awesome resources towards their defenses. But ultimately, they are to an important (if diminishing) degree controlled by human beings, at least some of which may not be perfectly immune from mistakes or bad intent.

Given the trust enterprises of all types are placing in these providers, it would be naive to believe that intelligence agencies are not cultivating provider personnel as "assets." Moreover, intelligent and successful people do things for reasons having nothing to do with state-sponsored espionage, including money problems, social engineering, etc. There is simply no way to ensure that unauthorized access never occurs in companies this size.

Even if the leading CSPs were controlled by benevolent, super-secure robots, they are portals to the great unwashed masses—their customers. Not all of them practice the best information security hygiene. These customers have their own customers, likely retain data about these customers, and digitally contact other third parties like vendors. The attack surface is endless.

Sharing Responsibility for Security

AWS and Azure proclaim a "shared security" or "shared responsibility" model, in which each participant has allocated roles and responsibilities for cyber-security. AWS has the mantra that the customer is responsible for security "in" the cloud, while the provider has responsibility for security "of" the cloud (artist's rendering below):

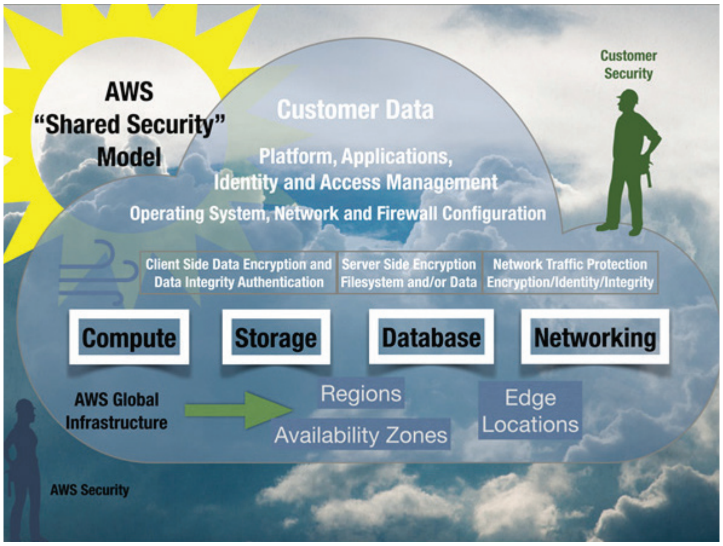
AWS
"Shared Security"
Model
Customer Security
Customer Data
Platform, Applications,
Identity and Access Management
Operating System, Network and Firewall Configuration
Client Side Data Encryption and Data Integrity Authentication
Server Side Encryption Filesystem and/or Data
Network Traffic Protection Encryption/Identity/Integrity
Compute
Storage
Database
Networking
AWS Global Infrastructure
Regions
Availability Zones
Edge Locations
AWS Security

Legally, the allocation is in a contract, practically, sharing is inherent in the integrated and interactive nature of the relationship. Nonetheless, misunderstandings on this topic are rampant.

The overly optimistic are under the delusion that, because these tech giants are the ultimate examples of technological sophistication and have the business incentives to be concerned about security, data in their services is more secure. (Period.) Even if true, using these services does not mean that the customer or user cannot impact information security. Moreover, each party's potential security impact varies depending on the particular service and how it is configured. The provider may drive the bus, but a passenger may still find some way to interfere with the driver's safe operation of the vehicle.

On the other side of the spectrum are the skeptics, believing that anything entrusted to a third-party is automatically less secure. For many lawyers, it may be difficult to overcome years of familiarity with the word cloud in connection with floating, fluffy, formless giant cotton-balls in the sky, to associate it with strong cyber-security. But in a world where there is no perfect security, security is relative. The notion that traditional corporate data centers are more secure or less vulnerable than AWS or Azure—as a blanket matter—is not supportable. Many criminal defendants have found out the hard way that representing themselves may not be wise and the same may apply to enterprise I.T.

Despite the growing number of publicly reported incidents where data in one of the big services was compromised, the ability to draw conclusions from these reports is limited. Like the universe of reported data breaches more generally, public information about these incidents is the tip of an unknowably large iceberg. Among those responding to a breach, "mum's the word" for obvious reasons and there are many business, legal and security rationales, real or imagined, for withholding disclosure.

Given the size and complexity of the major Cloud providers' infrastructures, it's easier not to think about the number of security incidents the providers evaluate; rather, stick with the understatement that the number is higher than what is publicly reported and enjoy blissful ignorance. Visibility into provider infrastructure is severely limited anyway—only the service provider has a complete picture. It turns out that substantial blindness is a serious side effect of getting the biggest and best in tech to provide and manage I.T. infrastructure.

Providers understand customers' need for monitoring and offer services that generate logs, such as CloudTrail and CloudWatch on AWS and Azure Monitor. But it may be up to the client to activate logging, as well as enable log file validation—important, because hackers alter log files to cover their tracks. Access to the logs themselves should be protected with measures like multi-factor authentication and attempts to access the logs should be logged. Logging should involve notifications to personnel who can take appropriate action. So, while the tools are offered, the customer has to know where and how to use them. Clients need to meticulously ensure that oversight of such matters penetrates to the level hands-on, implementation by mouse-click.

Regarding the increasing number of incidents in the news, we may charitably correlate it with accelerating cloud migration and the attendant learning curve. Reported breaches describe a similar, if not identical, story: a customer or their agent mistakenly designates a cloud repository as public rather than restricting access. Then, the period in which the company might realize the mistake and close the gates rapidly evaporates. Eventually, the vultures monitoring the public cloud for such exposed instances swoop in to gather a trove of data to sell. If only someone had closed the door.

To make matters worse for the victim's self-esteem, providers offer abundant guidance and tools to prevent catastrophes, although still a rising chorus demands they do more to protect customers from themselves. Security, even if presented as default setting, is not fully automated--a human being can usually click to render it powerless. Unauthorized access is widely viewed as the biggest threat to cloud security. In the nerd version of "Naughty By Nature", the problem is OPP—overly permissive permissions.

A thorough discussion of the information security assistance provided by the major services is for a later publication, as it is extensive. From network security, to resource management, secure virtual machine templates, encryption of data in transit and at rest, encryption key management, and much, much more, the protections are either in place or made available for use. This includes access to guidance and expertise from human beings. Finally, the major service providers have multiple certifications like ISO 27001, FedRAMP, DoD CSM, PCI DSS, etc.

The customer is responsible for making sure they configure their services securely. The services provide security tools, but if the customer either doesn't turn them on, or turns them off…and of course, customers are responsible for their own employees' internal compliance.

This is part of the bargain when you sign up for AWS or Azure. Regardless of the market power imbalance enjoyed by the tech giants, more legal and technical analysis of the risk allocation and its implications for clients would be constructive. Lawyers and clients need to understand the specific security measures available for the specific enterprise I.T. cloud services they use and how to ensure that they are implemented properly. Otherwise, the data there will be compromised. Plaintiffs and regulators will soon arrive to dance on the grave, citing the ample security measures offered by the provider for a click and maybe a few bucks.

Adam I. Cohen, CISSP, CEH, is a Managing Director at Berkeley Research Group LLC in New York. He is a Certified Ethical Hacker (CEH), Certified Information Systems Security Professional (CISSP), and former practicing attorney who for more than 20 years has advised clients on the intersection of technology and law, including cybersecurity, electronic discovery, and information governance. Adam has authored several books and his work has been cited in several landmark federal court opinions. He has served as a court-appointed neutral eDiscovery expert in the U.S. District Court, Southern District of New York and is an active member in professional organizations and educational institutions.