



Interested in learning
more about security?

SANS Institute InfoSec Reading Room

This paper is from the SANS Institute Reading Room site. Reposting is not permitted without express written permission.

Its 10PM...Do you know where your cloud is?

With advances in the virtualization market, there are a multitude of techniques to secure workloads that migrate from privately hosted clouds to publicly hosted clouds. With products such as VMware NSX, virtual firewalls, and VSANS becoming more capable on what seems like a daily basis, there is an increasing importance placed on both security and availability as equal considerations. Additionally within this space, securing workloads that use public cloudburst capacity during peak periods across hybrid clouds is another...

Copyright SANS Institute
Author Retains Full Rights

AD

DEEPARMOR®

Its 10PM...Do you know where your cloud is?

Securing virtualized workloads across federated (hybrid) clouds

GIAC SEC579 Gold Certification

Author: Robert J. Mavretich, bmav@rocketmail.com
Advisor: Richard Carbone

Accepted: July 13, 2014

Abstract

With advances in the virtualization market, there are a multitude of techniques to secure workloads that migrate from privately hosted clouds to publicly hosted clouds. With products such as VMware NSX, virtual firewalls, and VSANS becoming more capable on what seems like a daily basis, there is an increasing importance placed on both security and availability as equal considerations. Additionally within this space, securing workloads that use public cloudburst capacity during peak periods across hybrid clouds is another area to design in the appropriate security ahead of the curve. This paper will examine the leading edge of Software Defined Networking and examine the steps necessary to ensure that secure best practices from the physical world can translate to the virtual world in order to ensure seamless customer experiences, business continuity, as well as seamless support for those entrusted with maintaining the environment.

1. Introduction

From the time that Dr. Gordon Moore, the legendary founder of Intel postulated his theory that the number of transistors on an integrated circuit would double approximately every two years, the far off 21st century always seemed to hold the promise of flying cars and robotics making individual's lives easier. While the focus was mostly on hardware back then, it was a Harvard dropout in the late 1970's who staked his vision (and eventual fortune) on the idea that software was the future. While Dr. Moore might have seen his "Moore's Law" become less and less applicable in the present future of that past, Bill Gates' (the dropout) predictions are still relevant and remain readily available for public consumption and profit. Moore's Law focused on the advancement of the manufacturing of bare metal (hardware) technology, while the overlaying technology (software) has actually become the far-reaching enabler.

In the post-financial crisis business climate of 2008, concentration on cost cutting was intense. As expected, most companies concentrated on their greatest expense – labor. Technology practitioners, who were previously impervious to reductions in staff over the previous decade, were now suddenly finding themselves unemployed. Those who were left in Information Technology departments after the wave of layoffs were often overburdened with acute increases in workloads and no time increase in their day with which to perform those duties. While computers can run around the clock, administrators and engineers can only do without sleep for so long.

This exodus also revealed that there was a great need for highly skilled programmers, as software could be leveraged to quickly solve business challenges. With advances in hardware such as hot swappable drives, chassis advances from a racked server perspective, FireWire and USB connectors, as well as solid-state drives, the removal, repair, and transportation has become much easier and less dependent on vendor-specific training to use that hardware. In most cases, after the information was retrieved from failed hardware or physical drive space was upgraded, the old hardware was simply destroyed. If a company concentrates on the construction of a reliable and

scalable software platform there might be a greater return on investment in that product, rather than the hardware that needed (perceived) constant upgrading and co-location.

It also opened up an opportunity to offload the responsibility of owning and operating multiple physical data centers for redundancy purposes. This paradigm shift has allowed smaller organizations with fewer than ten employees (for example) to immediately compete with larger organizations to provide services to a wider range of clients. To be fair, there were certainly those small businesses who could not comprehend the value of the information superhighway building momentum before them, but there were also many who simply preferred not to participate. In the world of Information Technology, a conscientious objector who refuses to evolve with the platform-based trends of technology will quickly lose. These days, clouds are that technology and they can move fast.

Despite the promise of cloud-based technology, there have been growing pains in this space as well. Many companies saw the value of virtualization and developed their own cloud within their corporate borders to increase operational efficiency. Need a testing environment? Normal procurement time to process the order, receive it, determine rack space placement and power needs, installation of the operating system and software deployment, patching, etc., could take a month or more, at best, depending on the size of the organization. With virtualization the process could take ten minutes, assuming there was already a baseline virtual image with recent patches and software already applied and ready to be cloned. Companies realized that they could reduce the number of hardware assets required, and repurpose the ones they had already committed capital budget to in order to continue to receive benefits from them (rather than sell them for a loss or pay to dispose of). Private clouds are great cost savers that allow companies to stretch their scarce dollars. Moreover, they are also starting to look a lot like centralized mainframes that house all the valuable data in one place. However, one of the key differences in the architectural model is that the mainframe is much more silo-like in its implementation.

Private clouds aim to reduce the time to implementation while cutting across all business data silos, to increase awareness of that data and the ability to mine it in order to discover relationships that could produce an actionable business plan to drive revenue. It

has become a mandate of “find the needle in the haystack” as all corporate data can be useful - depending on the business perspective. Moore’s Law has made its way into the software arena with a slight twist – it now seems to describe the speed of both the new data and metadata created, and the proliferation of devices that can access that data.

While companies are making strides in accessing, provisioning, and controlling their data, the desire for faster provisioning to enable access to it has led us down the path of Software Defined Networking, almost eliminating the need for a one-to-one physical hardware position. While most companies have a relatively stable architecture that enables access into and out of the organization (firewalls, routers, WAN links) these hardware components are also subject to the same obsolescence problem. This is especially problematic, as these assets are being required to handle an increasing number of tasks at a speed they were not designed for, as well as their proprietary operating systems recently coming under an increasing number of attacks from malware.

The solution to this has been for some organizations to offload the heavy lifting to public cloud providers to aid in the defense and availability of their infrastructure, commonly known as Infrastructure as a Service (IaaS). With the rise of malware attacks from nation-states and other nefarious actors, and increasingly larger and more damaging DDoS attacks, companies can benefit from collaborating with a provider who can cut through infrastructure challenges quickly to direct unwanted traffic away from company “digital shores” and reduce the risk to customer and corporate data. Similarly, they need a partner who can assist in times of great bandwidth need, such as the holiday season sales push.

Not to be lost in the shuffle however, is the security of that data. “Cyber criminals are both taking advantage of the cloud’s benefits to launch attacks as well as targeting organizations’ cloud services. Exploits in private cloud and public cloud are being fueled with the ever increasing vulnerabilities in virtualized technologies which have nearly doubled between 2008 and 2010.” (Henry, P., VandenBrink, R., & Shackelford, D., 2013) Privacy concerns have been increasing as this technology has permeated every facet of an individual’s life as “big data” becomes bigger and wider, to the point that even non-technical users notice its presence in their daily lives.

Robert J. Mavretich, bmav@rocketmail.com

While companies are working to leverage their workloads and the information that it is based on across environments that stretch beyond their internal borders to those of a service provider like Amazon, for example, extreme care should prevail. “In a 2009 paper, four researchers from the University of California, San Diego, and Massachusetts Institute of Technology determined that there were significant risks to VMs running in public cloud environments. The risk stemmed from the multi-tenant nature of virtualization-based cloud environments. One major revelation the team exposed for the first time was the general lack of trust in cloud environments; in other words, how can you determine who your ‘neighbors’ are on a shared virtualization platform in a cloud provider environment? Criminals and other malicious attackers can provision VMs as easily as you do, and this can potentially have negative consequences.” (Shackleford, 2013) It would seem that technological advances have further virtualized social engineering capabilities. Attacks historically conducted via phone have moved to email, to embedded web links, to now entire platforms.

The only (but incredibly material) difference is that the data is now capable of moving effortlessly and instantly across the world to another hosting/data center for availability and support purposes, popularly called a “follow the sun” model. It is important to “bake” security into every business process to ensure that the data a company has been entrusted with will remain secure, no matter the abstracted platform on which that data resides, or where in the world that abstracted platform currently is.

Its 10PM...do you know where your cloud is?

2. Virtualization changes everything

When initially thinking about the changes in computing that have occurred since the late 1960's it is amazing to think how far we have come in the last ten years alone. Venerable corporations have been turned on their heads or turned into vapor and disappeared, due to the speed of technology that likely fueled those companies' initial ascension. Gone are the days when anti-virus programs would update once per month, then once per week, then once per day. It now seems strange if it does not update every hour. Patch Tuesday for Microsoft products needs to be a weekly event, not a monthly one with all the zero-day exploits created on a daily basis over the last few years. Data centers once clogged with rows and rows of network gear and server racks are now masquerading as barren wastelands with a few racks remaining in one small section of the entire floor. In addition, these remaining racks are capable of running, around the clock, a Fortune 50 operation effortlessly, and represent the promise of both the cloud and cloud computing quite nicely from a marketing perspective.

“There are a number of foundational characteristics that describe cloud computing, and most have come to rely on those defined by NIST, as a reasonable starting point.” (Henry, P., VandenBrink, R., & Shackelford, D., 2013)

These characteristics are defined as follows (unless other wise specified are from (Henry, P., VandenBrink, R., & Shackelford, D., 2013):

- “Broad Network Access: The cloud should be highly available from anywhere.
- Rapid Elasticity: Resources can be quickly provisioned to provide additional capacity, or scaled back when needed.
- Measured Service: Service Level Agreements concerning performance and function with respect to cost should be measurable, etc.
- On-Demand Self Service: Cloud users (customers) should be able to provision resources themselves, independent of a central administration function.
- Resource Pooling: “Resources should be pooled together for savings and efficiency.” (Mell & Grance, 2011)

Robert J. Mavretich, bmav@rocketmail.com

VMware has been the 800-pound gorilla of the virtualization space from the beginning of this technological advance, although Microsoft has been making head waves in recent years with HyperV. At a very basic level, VMware takes the entire operating system and boils it down to a representation of files, called the VMDK. This is cataloged and loaded every time a virtual machine is loaded, essentially making the platform virtual software. When not tied down to a single physical disk for processing, the speed and ability to complete intended tasks increases. Being able to share physical space on a SAN or NAS, or even one physical hard disk drive¹ has the capability to allow fast scalability, previously hindered by physical drive size constraints. SaaS has now taken over the desktop/laptop - and not just the vendor application software needed to run on it. These tie in with the “rapid elasticity” and “resource pooling” concepts already touched upon, now becoming a de facto standard that is seeing large-scale adoption.

Now extrapolate that concept across a data center and benefits are quickly gained from the “on the fly” provisioning of network services, as well as tying in the “Broad Network Access” component. “Network virtualization is a concept of combining the available resources in a network by splitting up the available bandwidth into channels, each of which is independent from the others, and each of which can be assigned to a particular server or device in real time.” (Asay, 2013) By releasing network equipment such as routers, switches, and firewalls from the shackles of bandwidth choking and access barriers, the possibilities increase as to what can be accomplished by opening the platform up to end-users so they can consume and somewhat direct the evolution of a company’s product. Most organizations are interested in revenue generating activities, and even though IT has been slowly losing its ill-gotten reputation as a “sunk cost” center as of late, the primary focus is still not on building a network, but rather a “nerve center” between a company and its customers. What used to be a burdensome cost is now getting attention from management because they can now do more with less, and need fewer capital dollars allocated for it.

As we move towards this future, most organizations have realized that they cannot go it alone. This is where “measured service” and “on-demand service” come into

¹ It would have to be fast and large enough to meet the minimum requirements for multiple VMDK’s.

play. Microsoft, Amazon, Google, and most of the other major companies that are at the leading edge of this phenomenon have offerings to allow end-users and corporations alike to participate in the virtualization benefits that the cloud can provide, especially *their* cloud offering. Now armed with the ability to outsource any function from a SQL database and its queries, credit card processing, email hosting and administration, federation of access management across companies, and everything in between, companies need to stop and ask themselves a key question: While it is available to them and they *can* do this...*should* they do this?

Of course, this question assumes that a company has a good idea of what its key assets are and that data classification is completed. “As companies embrace cloud computing, many are finding it advantageous to use external clouds to host non-critical IT services and data while keeping business-critical applications on the internal cloud infrastructure. However, this hybrid approach can create significant management challenges. The clouds must tightly integrate with one another, and legacy systems and data and workflows must be managed across the clouds and systems.” (Grimes, 2012) The annual Verizon Data Breach Report for 2014 indicates that for an alarming number of large and medium size organizations, a breach of critical assets occurred. Data breaches are a significant management challenge!

Amazon Web Services and their EC2 cloud offering, as well as Google Docs, have become highly entrenched in the business plans and capabilities of many organizations. One would hope that those organizations are not using a cloud provider with a public component for anything related to their confidential intellectual property without proper encryption of that data. An increasing number of cloud agreements will plainly state that customer or regulatory data shouldn't be housed in the public cloud and that they as service providers accept no responsibility for companies doing so. “In a survey of cloud services provider contracts, we note that nearly all cloud services providers give protection against third party suits for intellectual property infringement. This is common with software license agreements as well. Other indemnification, such as indemnification for violation of laws or for breach of contract, is far less common, and the existence of such a term should be viewed as a beneficial non-industry standard term for the customer.” (Kearns, 2014)

Robert J. Mavretich, bmav@rocketmail.com

Unfortunately, most companies do not engage in either obfuscation or encryption due to the time and money it would take to perform, as well as latency concerns. Despite high profile breaches including Hannaford Brothers grocery stores, TJ Maxx, and recently Target that have resulted in now tens of millions of customers having their personal information stolen, encryption is still not as pervasively implemented as it needs to be. “The added complexity breeds weaknesses that can be taken advantage of by a rogue employee or an external hacker...In many cases, there are too many different technologies pieced together that don't interface together...We're talking about a patchwork of solutions at many organizations that don't talk together.” (Westervelt, 2014) This will become more evident as companies shift their data towards a cloud-based model as cost-savings pressure is applied by senior management. On a somewhat positive note, latency concerns and pricing constraints are becoming less valid as an excuse for *not* encrypting, due to the advances in speed of delivery provided by cloud computing virtualization. Moreover, with vendors such as Akamai and other Content Delivery Network (CDN) vendors, their capabilities continue to grow quickly.

Even as the need for physical space decreases, the concerns surrounding security and compliance could not be larger. “...the opportunity that comes with ease of administration is the capability to secure and ensure compliance in a way that capitalizes on the fundamental concepts of Software Defined Networking. Ensuring security controls are multifunctional and adaptive and can react to change events in the network is an essential component of the converged data center. Software-defined security (SDS) meets these needs and protects the network from within the virtual infrastructure...” (Asay, 2013) It seems surprising that so many organizations are jumping into this technology without fully understanding the threat landscape or valuing/classifying their data appropriately to take advantage of the benefits of the software converged data center. Research from cloud security provider CloudLock (formerly Aprigo) from 2011 backs up this assertion. “What the company discovered...was the sheer number of businesses using the Internet cloud as their primary platform for computing—and how fast that market was growing. So CloudLock decided to drop its on premise, subscription-based software and sell only a cloud-based version.” (Huang, 2011)

The figure below indicates that from the early stages of cloud computing, the vulnerabilities have arrived at a faster pace year over year. The inconvenient details of the “IDS events” line may never be fully known but seem to indicate that this is a lagging indicator - a 500% increase in one year. This is not a particularly positive indication of both IDS capabilities and the forward-looking benefits of placing valuable information assets within a publicly based cloud infrastructure.

This increase may help to explain recent mergers in the industry such as Akamai (historically a content delivery provider) purchasing Prolexic, a leader in DDoS protection solutions. Customers may not know that they need DDoS solutions to protect their content from nefarious actor’s intent on doing harm, but these vendors are foreseeing the coming wave of exploits and looking to head them off in service to their customers, with an increase in their revenue as a bonus.

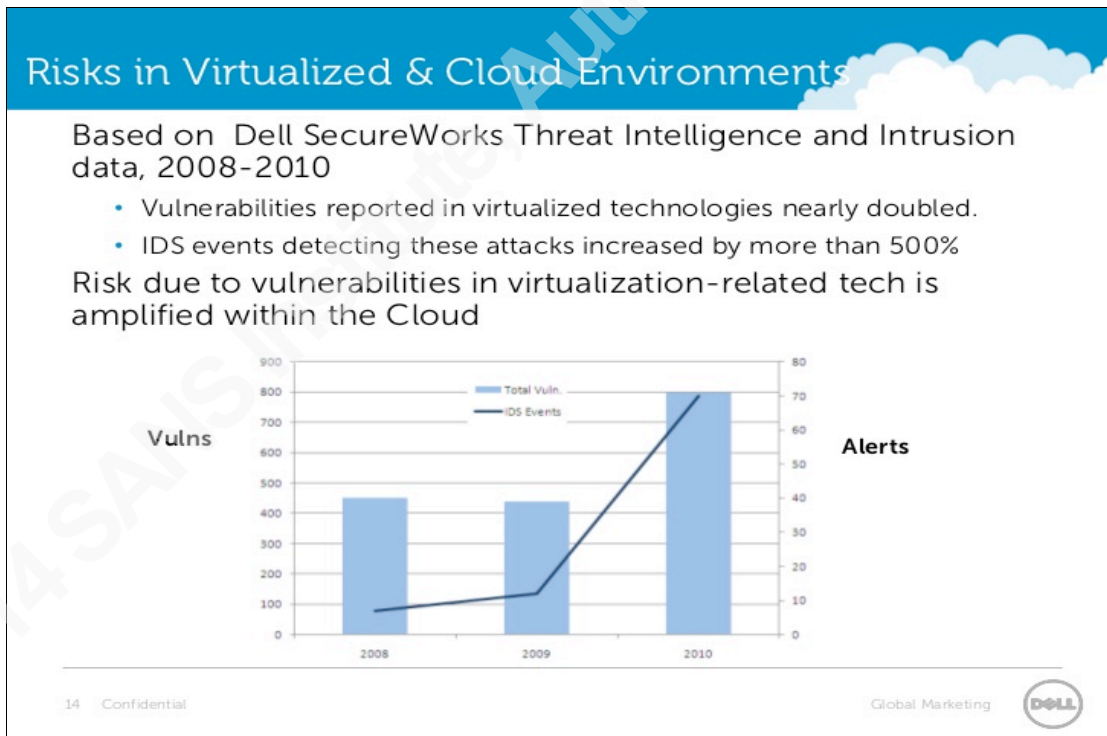


Figure 1: Risks in Virtualized & Cloud Environments (Photo Credit - Vance, 2011).

The way in which these vendors are protecting data is by leveraging the cloud. When a content delivery network serves up content to customers, they are effectively standing between the corporate network and the customer's browser. These vendors, such as Prolexic, "scrub" the data of malicious payloads. "When a DDoS attack is detected, our DDoS protection services are implemented within minutes. Upon activation of DDoS protection, a Prolexic customer routes in-bound traffic to the nearest Prolexic scrubbing center, where proprietary DDoS filtering techniques, advanced routing, and patent-pending anti-DoS hardware devices remove DDoS traffic close to the source of the botnet activity. Clean traffic is then routed back to the customer's network." (n.d., 2014) Once the traffic has had the threats removed, it continues on its transactional path towards various customers. Today, in a high-tech virtualized world, malicious data payloads are literally being sent to the car wash.

The ability for companies to perform these on-the-fly defensive maneuvers has become harder and harder as the threat landscape grows larger and larger everyday – Moore's Law at work again in the software arena. Companies are well served by concentrating on their own product (representing the growth engine of their company) and its abilities, rather than trying to become a CDN or DDoS provider themselves. This is no time to try to do things without a partner, even if staffed with a competent team. The "follow the sun" model of support has mandated that IT operations (and therefore the threats to it) run continuously. Unless a company possesses a large staff who have these competencies working around the clock, it is prudent to engage in outsourcing of important services such as defense and delivery.

Vendors themselves have been moving towards the near total virtualization of their own data centers and platforms, so that they may eventually take full advantage of Software Defined Networking for both their internal and external customers. Private clouds from vendors allow for the fast provisioning of environments to serve both development (internal employees) and production (client facing) situations, and they are quickly proving that the model can work very well. "Using Cisco Intelligent Automation for Cloud, a cloud management tool that provides a self-service Web portal, service catalog and orchestration, 'We built an easy to use Web portal' that lets users easily build sophisticated environments in short order, Cribari says. 'Let's say you want a virtual data

center with 100 VMs. You go in and select how much storage you want and in what increments you want to grow that storage, and then you select other things like, what kind of network do you need? Do you need a DMZ or an Internet facing network, or do you just want an internal production network?’ Users can then layer on Platform as a Service (PaaS) options, such as an Oracle database schema or an Apache server.” (Dix, 2014)

While the focus again is on the actual implementation of this time and cost saving concept, the issue of compliance can also be assisting with the technology. As easy as it is to provision VM’s, it is also easy to establish parameters that will enforce compliance without overwhelming staff or relying on their follow up. When provisioning in private clouds, the ability to “fill or kill” (to use a stock trading term) underutilized VM’s will help ensure that the sprawl that might have easily occurred is automatically culled on an ongoing basis. It also allows wasted space within a fabric network to be reclaimed to ensure efficiency. “Infrastructure capacity planning in this agile, virtual environment is done the same way as in traditional data centers, Cribari says. If the engineers see that a virtual environment is hitting a threshold -- 70% in production, 60% in non-production -- we get a team together to figure out what’s the next logical upgrade? How much do we need? How do we design it and provision it before we need it?” (Dix, 2014)

At this point, companies that have made the jump to move business data to the public cloud have realized that they need assistance in corralling all this data. This may be a little late as they have effectively thrown everything (and the proverbial kitchen sink) out to the cloud that they do not fully own or fully control. “To minimize risk exposure, many companies find that the best course of action is to start by moving lower-risk applications and data to the external cloud, so that any breach will not involve sensitive processes and information. Once confidence is established, the vendor can assume greater responsibility for other systems and databases.” (Grimes, 2012) A commercial during the Super Bowl of 2000 presented data proliferation challenges in the context of “herding cats” to convey the importance of getting a hold of meaningful data. Almost 15 years later, that commercial’s message still rings true. Vendors such as Solutionary (now an NTT company) are playing an increasingly important role in the monitoring required when there are multiple users and vendors in a corporate environment.

Robert J. Mavretich, bmav@rocketmail.com

Data Loss Prevention (DLP) software was a relatively new concept less than ten years ago, and is another example of marketing (and a solution in this case) that was ahead of its time. At the time, few companies actually took advantage of it because they saw it as serving one purpose, and not actually solving their problems but making them worse. The time to set up and “train” the software was often deemed too lengthy for most organizations, and too confusing and time consuming to their end-users. “The goal is to give the department’s direct access to information about how their records are being used as well as the ability to produce needed reports. The hard part of the process is bringing all the participants together under agreed procedures...” (Messmer, 2011)

For those companies who did take the time to purchase this type of software, the unintended benefits may have produced actionable metadata – such as where the valuable data actually resides within their environment. In order to normalize DLP solutions, a company was forced to categorize and classify valuable data to track its path around (and not out of) the organization. For most organizations, this (sometimes) multi-year exercise was too painful; but if taken on, completed, and incorporated as an ongoing business process, the benefits can be massive as virtualization tasked that knowledge across massive datasets. Knowing what and where data are across multiple platforms and databases will assist a company in ensuring that they can take the step towards a hybrid or public model for those processes that may contain less sensitive information and as a result can be outsourced to gain efficiencies. With the arrival of Software Defined Networking, there is no better time to corral these cats.

3. Software Defined Networking

When asked what is Software Defined Networking, one should consider that many ideas may define it, but the Open Networking foundation seems to have accurately captured it when it states: “Software-Defined Networking (SDN) is an emerging architecture that is dynamic, manageable, cost-effective, and adaptable, making it ideal for the high-bandwidth, dynamic nature of today's applications. This architecture decouples the network control and forwarding functions enabling the network control to become directly programmable and the underlying infrastructure to be abstracted for applications and network services” (n.d., 2013).

The SDN architecture is defined as (n.d., 2013):

- **“Directly programmable:** Network control is directly programmable because it is decoupled from forwarding functions.
- **Agile:** Abstracting control forwarding² lets administrators dynamically adjust network-wide traffic flow to meet changing needs.
- **Centrally managed:** Network intelligence is (logically) centralized in software-based SDN controllers that maintain a global view of the network, which appear to applications and policy engines as a single, logical switch.
- **Programmatically configured:** SDN lets network managers configure, manage, secure, and optimize network resources very quickly via dynamic, automated SDN programs, which they can write themselves because the programs do not depend on proprietary software.
- **Open standards-based and vendor-neutral:** When implemented through open standards, SDN simplifies network design and operation because instructions are provided by SDN controllers instead of multiple, vendor-specific devices and protocols.”

² Forwarding is another term for data packet movement




Figure 2: SDN - 3 Layers (Photo Credit: Open Networking Foundation, 2013).

The Open Networking Foundation’s “...OpenStack is a compilation of software components specifically designed for operating a cloud computing environment. For vendors, it allows them to leverage the work of a community of hundreds of developers who have worked to advance the code. For users, OpenStack allows the potential to have the same platform for your private cloud as multiple different public clouds, enabling a federated hybrid cloud.” (Butler, 2013) This is important because as the recent Heartbleed bug shows, using an open source product that lacks the financial backing of paying member organizations may leave corporate consumers of those open source tools holding the bag, incurring losses in customers and revenues. This is no secret as “...security experts and even the open-source movement’s biggest advocates acknowledge that Heartbleed revealed that some crucial open-source systems are underfunded and suffering from a lack of resources.” (Perloth, 2014)

As much as corporate products are frequently maligned for their deficiencies, they are backed by an organization that can quickly muster resources to fix identified problems and take some of the heat off those organizations that are consumers of those products. Therefore, the switch to an open source, vendor neutral format of business is a

new twist; the concept is predicated on no single organization having control. However, in a nod to the importance of the OpenStack concept to enable commerce and communication, many of the major organizations (VMWare, Microsoft) have taken care to ensure their products work with it. When the first major vulnerability is discovered, it will be interesting to see which organizations rise to the occasion to remedy it the quickest.

The central management of this technology needs to swiftly move towards a single pane of glass, and products like the old OpenView (by HP and since re-branded) may eventually come to be used as a verb, much like Google. The ability to move data seamlessly from one platform to another will be as easy as pointing and clicking, and not necessarily having to wonder if the end-user will get an error message stating: “the version of the software you are using appears to be an outdated version, please upgrade.” Of course, the ability to point and click can also create situations whereby the platform itself can be replicated across the globe. It must be verified that as the platform is moving, the ACLs, access management attributes, and protection of the underlying data is maintained and does not end up in a less secure state from where it started. Established RFC processes should enable this confirmation and maintain the dataset configuration “lift and shift” regardless of when, where, and how the lifting and shifting occurs.

The paradigm shift of how long our relationships endure with technology has changed irrevocably over the past fifteen years. Companies that push the latest technology through programs such as Verizon’s “New Every Two” are creating a critical mass of smart devices and unintentionally orphaned data left behind on them. “As we move closer to what some are calling an ‘Internet of Things,’ there will be more devices and systems to protect and, equally worrisome, more that can be used to launch attacks. And most imminently, there are the ongoing threats to businesses including the efforts to intercept communications, shut down networks, and harvest know-how from American companies...These efforts are becoming increasingly sophisticated, and they pose a growing challenge to both the competitiveness of American industry and the prosperity of our economy.” (Brennan, 2014)

Robert J. Mavretich, bmav@rocketmail.com

There is also privacy concerns, which most users largely ignore, then complain about vehemently when notified by a company that their personal information might have been breached. Security expert Bruce Schneier distills it perfectly for the casual user: “Take the Nest thermostat, which connects to the Internet. All your heating and cooling data are stored in the cloud, meaning on the company's servers. The company knows when you are home, when you are not. You might have said, ‘Well, that's a small company.’ But Google just bought Nest. Now Google has that data, along with its other data. [Nest's CEO has said that data is used only for Nest services and if this changes, users will be asked to ‘opt in.’]. You really don't have a choice. It's hard to live without Facebook or a cellphone. We're dealing with immediate gains vs. long-term, nebulous losses. Those are hard tradeoffs for people.” (Poppick, 2014)

However, in the case of Nest, they are using the aggregated data as told by their privacy policy. “We may share your aggregated and anonymous information in a variety of ways, including publishing trends about energy use and conservation, to help utilities provide demand-response services and to generally improve our system. We’ve taken steps to ensure that the information cannot be linked back to you and we require our partners to keep all information in its anonymous form.” (NEST Security Policy, 2014) This is the true promise of the cloud – enabling technology that is perceptive enough that it can positively affect and inform customers – without necessarily knowing that it is *a specific identified customer*.

In the pursuit of revenue sometimes, the only way to get there is by a rabid cost containment/reduction methodology. Unfortunately, if this means exposing customer data, as long as the Terms and Conditions spell it out in a less than clear manner (or require a law degree to interpret), some companies will not hesitate. This is a classic example of where public policy and business strategy can seem diametrically opposed. In addition, everything with a positive intent that starts to utilize cloud services, can have an equally negative potential outcome by using the cloud. Perhaps it is time for a Hippocratic Oath for business, with the foremost principle being “do no customer harm” rather than “profits first.”

A developing answer has been the concept of data lakes. “Enterprises from government agencies to large concerns and on down use big data inside public

multitenant cloud environments. All the risks of multi-tenancy apply in these scenarios including the vulnerabilities that come with the weaker security of another tenant, potential access by users of an adjacent tenant, PII/PHI exposure, and other regulatory non-compliance. Data lakes could protect big data from all the perils of the public cloud.” (Greer, 2014) The fact that there is admitted “peril” in the public cloud should again give companies pause on whether the data they are shipping off and its control is truly worth the risk. There are a multitude of systems still in use from another era, such as control systems around nuclear power plants and other critical national infrastructure. While most of these systems have been immune to the siren call of anything other than the closed system they use to function, it begs the question – how much longer can they hold out with the cloud hovering over them?

The cloud can enable protection as much as it can endanger it, but the *customer* has to take the lead in requiring it first. This is where the Master Services Agreement must spell out the specific expectations and requirements that surround the custody and use of the data (in-line or as an attached appendix). Without a specific instructive statement, most cloud providers will not provide a single tenant environment (where the hardware, storage and network are dedicated to a single customer), and co-locate corporate data on an instance with other companies’ data. This is important because there have been reports of side channel attacks against VMware that could endanger confidential data or platforms through the hypervisor itself. If the data is housed right next to another company’s data, how is an organization going to find that out until the breach occurs?

This makes data lakes seem more like Wall Street proprietary trading systems, which are hardly democratic and free market efficiency systems. “...between the public stock exchanges and the dark pools – private exchanges created by banks and brokers did not have to report in real time what trading activities took place within them...” (Lewis, 2014) As a result it is important to also maintain an open relationship with the cloud provider, so that at least once per year a company is allowed to “peek under the covers” to see what is being done with their data. Unfortunately, at present, attestations do not go far enough in what happens with each specific data set. While they are helpful and provide a level of comfort around the *general* security practices of a cloud provider, well-

meaning network administrators sometimes have accidents, and malicious activities are usually *specific* in nature, and *random and unplanned* as far as timing go.

This brings to light one of the greatest benefits of Software Defined Networking – the configuration of the environment can be made consistent. The entire infrastructure is now capable of moving without the need to configure an identical data center/host on the opposite side of the globe. This virtually eliminates the risk potential for mishaps such as the incorrect configuration of the destination hardware/software, because an organization is relying on its own staff of engineers. This also allows immediate feedback on whether or not things are working (and a quick look at disaster recovery capabilities). If not using the SaaS model but rather PaaS or IaaS, it removes the possibility of a rogue admin in the cloud from accessing the data because the organization controls the data and the routes it flows over if networking and security equipment has been virtualized as well. This is important to keep track of when making the changes.

As a result, auditing will play a greater role in cloud service delivery. “Auditing is a systematic process of objectively obtaining and evaluating evidence regarding assertions about economic actions and events to ascertain the degree of correspondence between those assertions and established criteria and communicating the results to interested users.” (Robertson & Davis, 1982) While this definition, taken from an accounting text could be challenged in this context, if the data are effectively made into a revenue stream and placed onto the cloud to parse, use, and move around, hasn’t its basic function been changed to one that is economic? Cloud providers (as of now) are also trying to be as *subjective* as possible regarding their capabilities in order to gain market share in an already crowded field.

In looking at most of the noteworthy breaches from the past few years (including the recent Target breach) a large number of them are aimed at large companies, because they continue to remain “crown jewels” as they have the most customers and revenue to collect. As Target showed, there is too much data cached locally while not being protected properly. This is why penetration testing is so important and gaining popularity. However, if data can be quickly moved by leveraging Software Defined Networking it makes it harder for the bad actors to hit that moving target. Regular failover of the datacenter could become not only a sound continuity of business/disaster recovery

Robert J. Mavretich, bmav@rocketmail.com

planning, but also a pro-active defense mechanism. A popular statement in security is that a black hat hacker needs to be successful *only once* to get in and damage the network, but the organization needs to get *lucky all the time* based on the number of attacks and attackers targeting the network. Where's your cloud? Only your organization knows. Black hats may be left scratching their heads as their reconnaissance showed them that a data center was here today, but will be gone tomorrow.

Penetration tests have taken on a greater role in the compliance and reporting arena over the past few years to provide that point in time objectivity, and may have to become even more prevalent as clouds expand their reach. Customer's should be encouraged and allowed to perform penetration testing access to their own data – before someone else does. This would largely be application level penetration testing, which would complement the provider's annual network penetration testing. Marrying the two together could give a greater assurance that data sensitive in nature to one company is being adequately protected in accordance with the contractual terms. It would also likely engender a large amount of goodwill and possibly new business opportunities for a truly “open” cloud provider.

As of late, a lot of penetration testing has been occurring *from* the cloud as well. The ability to leverage Software Defined Networking instead of clunky hardware boxes could increase the frequency of these tests, as there would (theoretically) be less of an issue setting up the test, configuring physical assets, reporting, etc. Because the abstraction layers are now on a level playing field, network controllers would know the best time to perform a test. During the scanning if there were any intrusive scanning activities affecting operational areas, the software defined switches and routers could either re-route the traffic or send a pause to the testing host. The test could effectively be able to learn (and inform people) how to run within the environmental constraints once set up with requirements and target hosts. A penetration test may eventually evolve into its own Software as a Service offering.

Furthering the cloud service model are companies that are interested in performing protection activities that allow end-users to determine what is actually important. DLP solutions historically required the purchase of an expensive hardware

device (though that has changed in the last few years from the original DLP providers) which may explain why it was so hard to do – funding required. However, vendors have gotten the message, and have noticed the trend of datasets making their way to the cloud and have started to provide solutions. “Aprigo's (renamed to CloudLock) cloud-based security service is a tool for Google Apps that lets administrators place access controls over their documents by specific user or groups, deciding who is allowed to share what with whom. The service allows for the selection of data owners or auditor rights and delegation controls. It can enforce policies around who can copy local files to Google Docs, and who has edit permissions. It offers a way to do data inventory, and keep track of which documents are shared outside the organization on a more public basis...” (Messmer, 2011)

Identity and access management will play a larger role in the coming years as federation models are leveraged to extend trust relationships across normally closed off corporate borders. By using the concepts behind Software Defined Networking such as Service Oriented Architecture, the provisioning, use, and maintenance of the required access (add, modify, delete) should become more prolific. This will continue as companies seek partnerships to share costs and increase revenue (popularly called a joint venture) by leveraging “write once, use multiple times” software code, instead of provisioning expensive and potentially limiting physical hardware.

One of the more damaging breaches of the past few years was the loss of an end-user laptop by a Veterans Administration hospital employee that contained the Personally Identifiable Information (PII) of thousands of hospital patients, including social security numbers. “The names, dates of birth and Social Security numbers of about 26.5 million active duty troops and veterans were on the laptop and external drive, which disappeared while in the custody of a Veterans Affairs data analyst in 2006.” (Frieden, 2009) This actually underscores another one of the positive benefits of using a cloud provider – no more physical security hardware losses for end-users. While this seems counterintuitive, end-users have proven that in general they treat company property with little to no regard, and do not follow basic practices of security even with their own devices that carry lots of their own personal information. In a deployment of Office360, (a public model),

Robert J. Mavretich, bmav@rocketmail.com

organizations may investigate writing data to a private cloud that offloads much of the overhead from a primary data center, while doing a better job of protecting the data in a non-decentralized way (as opposed to an administrator losing a laptop and exposing important information such as PII or company confidential data).

Unfortunately, the current reality is that Software Defined Networking will not result in a net increase in the security of data until companies and their consumers mature. Companies like Google and Facebook have been pushing the envelope in terms of privacy violations for years in order to enhance their revenue stream. This is why signing up for Facebook is “free and always will be” – if you are not paying for the product, you *are* the product. This is contrast to a corporate environment where there is some control over the terms and conditions of the products because they are subject to purchase agreements.

Because there is a clear “caveat emptor” statement concerning any non-private cloud offering, many firms have been investigating the possibility of purchasing breach insurance due to the massive breaches of the past decade. This is becoming very difficult to consistently and accurately place a dollar value (or even an appropriate range) on the potential risk. When including considerations such as brand reputation and other “goodwill” type balance sheet line items, insurance companies will likely start to challenge claims should they start to proliferate. However, unlike natural disasters, breaches occur daily and without any advanced warning. If the insurance industry can develop a model to account for this and provide a product that corporations would be interested in, this may be positive for the industry as a whole. Perhaps as part of a claim, the company in a breach situation would be required to report their breaches within a certain amount of time and at a certain level of detail, in order to claim insurance funds. This is why the contractual language portion of the contract can potentially make or break a cloud model relationship. Used properly, it may reduce the financial sting down to an acceptable level of risk. If you use a public cloud model, you cannot fully protect the data, and must act accordingly with this knowledge.

Robert J. Mavretich, bmav@rocketmail.com

4. Conclusion

When factoring in the complexity of Software Defined Networking, there will likely be a best effort by security firms looking to leverage the public cloud to reduce both operating and capital costs. It is suggested by this author that the following recommendations be followed, as well as any company-specific processes be accounted for and re-architected in this ongoing review when it involves anything other than public data and it's movement while under control of a third party (and not the data owner itself):

- Classify company data by using DLP software in-house;
- Perform site visits of as many potential Cloud vendor locations where possible, and require multiple levels of attestation/audits, such as PCI compliance, SSAE, SOC1/SOC2, external privacy audits, proof of monitoring and configuration such as Tripwire or Retina logs, etc.;
- Use FEDRAMP as a guide and minimum requirements bar for selecting cloud vendors;
- Be cognizant of the fact that nothing can fully secure the platform, and only move company data accordingly after appropriate classification;
- Monitor evolving cloud computing law by participating in organizations such as the Cloud Security Alliance (CSA).

Companies should make a thoughtful and consistent effort to fund security tools and processes that adhere to secure best practices, combined with encryption and multi-factor authentication models that stress security over convenience. It is only by leading in this way that organizations and their data remain as safe and secure as possible as it moves effortlessly across the globe through a brave new software defined world.

5. References

- Asay, R. (2013, 12 30). *The ripple effects of SDN: How it will change data center it*. Retrieved from <http://www.datacenterjournal.com/it/ripple-effects-sdn-change-data-center/>.
- Butler, B. (2013, 04). *OpenStack grows up: But is it grown up enough for enterprise it?* Retrieved from <http://www.networkworld.com/news/2013/041013-openstack-268604.html>.
- Cooney, M. (2014, 02 27). *CIA chief: Internet of things, infrastructure attacks are big security headaches*. Retrieved from <http://www.networkworld.com/community/blog/cia-chief-internet-things-infrastructure-attacks-are-big-security-headaches>.
- Dix, J. (2014, 04 05). *Inside cisco's private cloud*. Retrieved from http://www.networkworld.com/news/2014/030514-cisco-private-cloud-279416.html?source=NWWNLE_nlt_cloud_security_2014-03-06.
- Frederick, P. (2014, 02 20). *Red Hat CEO addresses big data's chicken-and-egg question*. Retrieved from <http://www.networkworld.com/community/blog/red-hat-ceo-addresses-big-datas-big-chicken-and-egg-question>.
- Greer, D. (2014, 03 16). *Can data lakes solve cloud security challenges?* Retrieved from http://www.networkworld.com/news/2014/031814-can-data-lakes-solve-cloud-279824.html?source=NWWNLE_nlt_cloud_security_2014-03-20.
- Grimes, D. (2012, 11 26). *Meeting the challenges of hybrid cloud computing infrastructures*. Retrieved from <http://www.networkworld.com/news/tech/2012/111512-hybrid-cloud-264322.html>.
- Henry, P., VandenBrink, R., & Shackelford, D. (2013). *Virtualization security architecture and design*. (V2013_1205 ed., Vol. 579.1). Baltimore: The SANS Institute.

- Henry, P., VandenBrink, R., & Shackleford, D. (2013). *Virtualization and private cloud infrastructure security*. (V2013_1205 ed., Vol. 579.2). Baltimore: The SANS Institute.
- Henry, P., VandenBrink, R., & Shackleford, D. (2013). *Virtualization offense and defense (part 1)*. (V2013_1205 ed., Vol. 579.3). Baltimore: The SANS Institute.
- Henry, P., VandenBrink, R., & Shackleford, D. (2013). *Virtualization offense and defense (part 2)*. (V2013_1205 ed., Vol. 579.4). Baltimore: The SANS Institute.
- Henry, P., VandenBrink, R., & Shackleford, D. (2013). *Virtualization and cloud integration: policy, operations, and compliance*. (V2013_1205 ed., Vol. 579.5). Baltimore: The SANS Institute.
- Henry, P., VandenBrink, R., & Shackleford, D. (2013). *Confidentiality, integrity, and availability with virtualization and cloud*. (V2013_1205 ed., Vol. 579.6). Baltimore: The SANS Institute.
- Huang, G. (2011, 04 06). From Aprigo to CloudLock: Data protection startup talks strategy shift, new name. Retrieved from <http://www.xconomy.com/boston/2011/04/06/from-aprigo-to-cloudlock-data-protection-startup-talks-strategy-shift-new-name/>.
- Kearns, W. (2014, 05 13). *cloud computing: Evolving contracting practices*. Retrieved from <http://www.dwt.com/Cloud-Computing-Evolving-Contracting-Practices-05-13-2014/>.
- Lewis, M. (2014, 04 06). The wolf hunters of Wall Street: How a band of outsiders discovered that the stock market was rigged - and set out to change it. The New York Times Magazine, DOI: www.nytimes.com.
- Mell, P., & Grance, T. (2011, 09). *The nist definition of cloud computing*. Retrieved from <http://csrc.nist.gov/publications/nistpubs/800-145/SP800-145.pdf>.
- Messmer, E. (2011, 02 13). How one municipality is securing, managing Google apps, docs. Retrieved from <http://www.networkworld.com/news/2011/022311-panama-google-apps-docs.html>.

- Metzler, J. (2012, 08 29). *What is software defined networking (sdn)?*. Retrieved from <http://www.networkworld.com/news/2012/082912-insider-sdn-262010.html>.
- Millard, C. (2013). *Cloud computing law*. (1st ed.). New York, NY: Oxford University Press.
- Owano, N. (2012, 11 09). *Vm researchers post rude awakening about virtualization security*. Retrieved from <http://phys.org/news/2012-11-vm-rude-awakening-virtualization.html>.
- Perloth, N. (2014, 04 24). *Companies back initiative to support openssl and other open-source projects*. Retrieved from http://bits.blogs.nytimes.com/2014/04/24/companies-back-initiative-to-support-openssl-and-other-open-source-projects/?_php=true&_type=blogs&_r=0.
- Poppick, S. (2014, 04 14). *Is my data safe?*. Retrieved from http://money.cnn.com/2014/04/01/technology/security/data-security.moneymag/index.html?iid=SF_M_River.
- Robertson, J., & Davis, F. (1982). *Auditing*. (3rd ed., p. 3). Plano, TX: Business Publications, Inc.
- Shackleford, D. (2013). *Virtualization security: protecting virtualized environments*. Indianapolis, IN: John Wiley and Sons.
- Strunk, William Jr. & White, E. B. . *The Elements of Style*. New York, NY: Longman
- Vance, A. (2011, 06 06). *Watch out for the dark cloud: Cloud computing and cyber threats*. Retrieved from <http://www.slideshare.net/DellServices/watch-out-for-the-dark-cloud-cloud-computing-and-cyber-threats0020>.
- Walters, S. (2014, 03 10). *Is the cloud the answer? what coke's recent breach teaches us*. Retrieved from <http://www.datacenterjournal.com/it/cloud-answer-cokes-breach-teaches/>.
- Warlick, David (2004). *Son of citation machine*. Retrieved February 17, 2009, from Son of citation machine Web site: <http://www.citationmachine.net>.

Westervelt, R. (2014, 01 27). *Coca-cola laptop breach a common failure of encryption, security basics*. Retrieved from

<http://www.crn.com/news/security/240165711/coca-cola-laptop-breach-a-common-failure-of-encryption-security-basics.htm>.

n.d. (2013) Software-Defined Networking: The New Norm for Networks". *White paper*. Open Networking Foundation. April 13, 2012. Retrieved August 22, 2013.

n.d. (2014)*Why prolexic: Scrubbing center*. (2014). Retrieved from

<http://www.prolexic.com/why-prolexic-best-dos-and-ddos-scrubbing-centers.html>.

6. Glossary

ACL:	Access Control List
Amazon EC2:	Amazon's Cloud Offering for Consumers
CDN:	Content Delivery Network
DDoS:	Distributed Denial of Service
DLP:	Data Loss Prevention
DMZ:	De-militarized Zone
FEDRAMP:	Streamlined certification process for providers who meet strict government data protection standards
HyperV:	Microsoft's version of the Hypervisor management interface for virtual machines
IaaS:	Infrastructure as a Service
IDS:	Intrusion Detection System
NAS:	Network Attached Storage
NIST:	National Institute of Standards and Technology
PaaS:	Platform as a Service
PHI:	Personal Health Information
PII:	Personally Identifiable Information
RFC:	Request for Change
SaaS:	Software as a Service
SAN:	Storage Area Network
SDN:	Software Defined Networking
SDS:	Software Defined Security
SOC1:	Service Organizations Control Report (replaces SAS 70)

SSAE: Statements on Standards for Attestation Engagements

SQL: Structured Query Language

VM's: Virtual Machines

VMDK: Virtual Machine Disk, coined by VMWare

VSANS: Virtual Storage Area Network

WAN Links: Wide-Area Network



Upcoming SANS Training

[Click Here for a full list of all Upcoming SANS Events by Location](#)

SANS Paris June 2018	Paris, FR	Jun 25, 2018 - Jun 30, 2018	Live Event
SANS Minneapolis 2018	Minneapolis, MNUS	Jun 25, 2018 - Jun 30, 2018	Live Event
SANS Vancouver 2018	Vancouver, BCCA	Jun 25, 2018 - Jun 30, 2018	Live Event
SANS London July 2018	London, GB	Jul 02, 2018 - Jul 07, 2018	Live Event
SANS Cyber Defence Singapore 2018	Singapore, SG	Jul 09, 2018 - Jul 14, 2018	Live Event
SANS Charlotte 2018	Charlotte, NCUS	Jul 09, 2018 - Jul 14, 2018	Live Event
SANSFIRE 2018	Washington, DCUS	Jul 14, 2018 - Jul 21, 2018	Live Event
SANS Cyber Defence Bangalore 2018	Bangalore, IN	Jul 16, 2018 - Jul 28, 2018	Live Event
SANS Pen Test Berlin 2018	Berlin, DE	Jul 23, 2018 - Jul 28, 2018	Live Event
SANS Riyadh July 2018	Riyadh, SA	Jul 28, 2018 - Aug 02, 2018	Live Event
Security Operations Summit & Training 2018	New Orleans, LAUS	Jul 30, 2018 - Aug 06, 2018	Live Event
SANS Pittsburgh 2018	Pittsburgh, PAUS	Jul 30, 2018 - Aug 04, 2018	Live Event
SANS San Antonio 2018	San Antonio, TXUS	Aug 06, 2018 - Aug 11, 2018	Live Event
SANS August Sydney 2018	Sydney, AU	Aug 06, 2018 - Aug 25, 2018	Live Event
SANS Boston Summer 2018	Boston, MAUS	Aug 06, 2018 - Aug 11, 2018	Live Event
Security Awareness Summit & Training 2018	Charleston, SCUS	Aug 06, 2018 - Aug 15, 2018	Live Event
SANS Hyderabad 2018	Hyderabad, IN	Aug 06, 2018 - Aug 11, 2018	Live Event
SANS New York City Summer 2018	New York City, NYUS	Aug 13, 2018 - Aug 18, 2018	Live Event
SANS Northern Virginia- Alexandria 2018	Alexandria, VAUS	Aug 13, 2018 - Aug 18, 2018	Live Event
SANS Krakow 2018	Krakow, PL	Aug 20, 2018 - Aug 25, 2018	Live Event
SANS Chicago 2018	Chicago, ILUS	Aug 20, 2018 - Aug 25, 2018	Live Event
Data Breach Summit & Training 2018	New York City, NYUS	Aug 20, 2018 - Aug 27, 2018	Live Event
SANS Prague 2018	Prague, CZ	Aug 20, 2018 - Aug 25, 2018	Live Event
SANS Virginia Beach 2018	Virginia Beach, VAUS	Aug 20, 2018 - Aug 31, 2018	Live Event
SANS San Francisco Summer 2018	San Francisco, CAUS	Aug 26, 2018 - Aug 31, 2018	Live Event
SANS Copenhagen August 2018	Copenhagen, DK	Aug 27, 2018 - Sep 01, 2018	Live Event
SANS SEC504 @ Bangalore 2018	Bangalore, IN	Aug 27, 2018 - Sep 01, 2018	Live Event
SANS Wellington 2018	Wellington, NZ	Sep 03, 2018 - Sep 08, 2018	Live Event
SANS Amsterdam September 2018	Amsterdam, NL	Sep 03, 2018 - Sep 08, 2018	Live Event
SANS Tokyo Autumn 2018	Tokyo, JP	Sep 03, 2018 - Sep 15, 2018	Live Event
SANS Tampa-Clearwater 2018	Tampa, FLUS	Sep 04, 2018 - Sep 09, 2018	Live Event
SANS MGT516 Beta One 2018	Arlington, VAUS	Sep 04, 2018 - Sep 08, 2018	Live Event
SANS Cyber Defence Canberra 2018	OnlineAU	Jun 25, 2018 - Jul 07, 2018	Live Event
SANS OnDemand	Books & MP3s OnlyUS	Anytime	Self Paced