

RSA Security Brief

March 2010



Infrastructure Security: Getting to the Bottom of Compliance in the Cloud

Authors

Sam Curry
CTO, Marketing
RSA, the Security Division of EMC

Jon Darbyshire
President & CEO, Archer Technologies

Douglas W. Fisher
Vice President & General Manager
Systems Software Division, Intel Corp.

Bret Hartman
Chief Technology Officer
RSA, the Security Division of EMC

Dr. Stephen Herrod
Chief Technology Officer and Senior Vice President of R&D,
VMware, Inc.

Vishal Kumar
Senior Product Manager, Emerging Markets and
Technologies, VMware, Inc.

Dr. Fernando Martins
Director of Strategic Technologies
System Software Division, Intel Corp.

Steve Orrin
Director of Security Solutions
Software Services Division, Intel Corp.

Dana Elizabeth Wolf
Senior Manager, New Business Development
RSA, the Security Division of EMC



The Security Division of EMC

RSA Security Briefs provide security leaders and other executives with essential guidance on today's most pressing information security risks and opportunities. Each Brief is created by a select response team of security and technology experts who mobilize across companies to share specialized knowledge on a critical emerging topic. Offering both big-picture insight and practical technology advice, RSA Security Briefs are vital reading for today's forward-thinking security practitioners.

Contents

Executive Summary	1
Cloud Compliance: Changing Expectations for Visibility & Control in Cloud Services	2
Cloud Infrastructure: the Next Frontier in Cloud Security & Compliance	4
Hardware Root of Trust: the New Sheriff for the Cloud Frontier	5
Cloud Gazing: What's on the Horizon in Cloud Security & Compliance	7
Practitioner Guidance for Improving Cloud Infrastructure Security	10
Harden all hypervisors	10
Set clear policies for co-residency and be equipped to enforce them	10
Evaluate whether cloud vendors can deliver on their promises	12
Assess cloud providers' methods for attesting to infrastructure security	12
Request automated dashboard services for monitoring & compliance	12
Summary & Conclusions	13
About the Authors	14
Jon Darbyshire	14
Sam Curry	14
Douglas W. Fisher	14
Bret Hartman	14
Dr. Stephen Herrod	14
Vishal Kumar	15
Fernando Martins	15
Steve Orrin	15
Dana Elizabeth Wolf	15
Solutions for a Secure Cloud Infrastructure	16
Governance, Risk & Compliance Solutions	16
Cloud Security Solutions	16
Hardware Root of Trust Solutions	17
Virtualization Layer Solutions	17
Cloud Consulting Services	18

Executive Summary

In cloud environments, one of the most pervasive and fundamental challenges for organizations in demonstrating policy compliance is proving that the physical and virtual infrastructure of the cloud can be trusted – particularly when those infrastructure components are owned and managed by external service providers.

For many business functions commonly run in the cloud – hosting websites and wikis, for example – it's often sufficient to have a cloud provider vouch for the security of the underlying infrastructure. For business-critical processes and sensitive data, however, third-party attestations usually aren't enough. In such cases, it's absolutely essential for organizations to be able to verify for themselves that the underlying cloud infrastructure is secure.

The next frontier in cloud security and compliance will be to create transparency at the bottom-most layers of the cloud by developing the standards, tools and linkages to monitor and prove that the cloud's physical and virtual machines are actually performing as they should. Verifying what's happening at the foundational levels of the cloud is important for the simple reason that if organizations can't trust the safety of their computing infrastructure, the security of all the data, software and services running on top of that infrastructure falls into doubt. There's currently no easy way for organizations to monitor actual conditions and operating states within the hardware, hypervisors and virtual machines comprising their clouds. At those depths, we go dark.

Cloud providers and the IT community are already preparing to address this problem. Groups of technology companies have banded together to develop a new, interoperable and highly secure computing infrastructure for the cloud based on a "hardware root of trust," which provides tamper-proof measurements of every physical and virtual component in the entire computing stack, including the hypervisor. Members of the IT community are exploring ways to use these measurements to improve visibility, control and compliance in the cloud.

They're collaborating on a conceptual IT framework to integrate the secure measurements provided by a hardware root of trust into adjoining hypervisors and virtualization management software. The resulting infrastructure stack would be tied into data analysis tools and a governance, risk & compliance (GRC) console, which would contextualize conditions in the cloud's hardware and virtualization layers to present a reliable assessment of an organization's overall security and compliance posture. This type of integrated hardware-software framework would make the lowest levels of the cloud's infrastructure as inspectable, analyzable and reportable for compliance as the cloud's top-most application services layer.

With this unprecedented level of visibility, we believe clouds can develop the infrastructure-level policy controls and the end-to-end security attestations to handle even the most demanding security requirements for applications and data. Ultimately, this will enable organizations to take advantage of the cloud's benefits in supporting a much broader range of business processes.

For many business functions commonly run in the cloud – hosting websites and wikis, for example – it's often sufficient to have a cloud provider vouch for the security of the underlying infrastructure. For business-critical processes and sensitive data, however, third-party attestations usually aren't enough.

Cloud Compliance: Changing Expectations for Visibility & Control in Cloud Services

Cloud computing delivers convenient, on-demand access to shared pools of data, applications and hardware. The cloud computing paradigm—made possible by sophisticated automation, provisioning and virtualization technologies—differs dramatically from today's IT model because it decouples data and software from the servers and storage systems running them and allows IT resources to be dynamically allocated and delivered as a service, either in component parts (where users subscribe to specific applications or simply lease computing power) or as an integrated whole.

While the cloud provides organizations with a more efficient, flexible, convenient and cost-effective alternative to owning and operating their own servers, storage, networks and software, it also erases many of the traditional, physical boundaries that help define and protect an organization's data assets. Physical servers are replaced by virtual ones. Perimeters are established not by firewalls alone but also by highly mobile virtual machines. Mitigating risk becomes more complex, as the cloud introduces ever expanding, transient chains of custody for sensitive data and applications.

For this reason, the vast majority of data and applications handled by clouds today isn't business-critical and has lower security requirements. Most organizations are already leasing computing capacity from an outside data center to host websites or corporate e-mail. Some have outsourced business functions such as sales force management to providers in the cloud. If the data from these applications

Differences Between Private and Public Clouds

"Private cloud" describes an IT infrastructure in which a shared pool of computing resources – servers, networks, storage, applications and software services – can be rapidly provisioned, dynamically allocated and operated for the benefit of a single organization. The organization needn't physically own or operate the IT assets that form its private cloud. Some assets can be outsourced to cloud providers – for instance, computing capacity may be leased from an outside data center. Nevertheless, the organization still effectively "owns" its private cloud by controlling and setting policies governing how virtual IT assets are operated, with cloud vendors guaranteeing specific levels of service and conformance to agreed-upon standards for information security and compliance.

If all the IT assets of a private cloud are physically owned and operated by the organization itself, the cloud is sometimes called an "internal cloud." In terms of monitoring and proving compliance with information security policies, organizations presumably have complete visibility, transparency and control over their internal clouds, because they own and maintain the entire cloud infrastructure, from servers to services.

"Public clouds" refer to shared cloud services that are made available to a broad base of users. Although many organizations use public clouds for private business benefit, they don't control how those cloud services are operated, accessed or secured. Popular examples of public clouds include Amazon's Elastic Compute Cloud (EC2), Google Apps and Salesforce.com.

For definitions of the security terminology included in this paper, please consult the [glossary available on RSA's website](#).

were compromised or the business processes became unavailable for a short period of time, the organization might be highly inconvenienced, but the consequences would probably not be disastrous.

Higher-value business data and processes, however, have been slower to move into the cloud. These business-critical functions – for example, the cash management system for a bank or patient records management within a hospital – are usually run instead on in-house IT systems to ensure maximum control over the confidentiality, integrity and availability of those processes and data. Although some organizations are using “internal clouds” (see sidebar for definition) for high-value information and business processes, they’re still reluctant to outsource the underlying IT systems, because of concerns about their ability to enforce security strategies and to use familiar security controls in proving compliance.

The cloud introduces four basic security and compliance challenges to organizations:

- **First, the cloud typically increases an organization’s reliance on cloud providers’ logs, reports and attestations in proving compliance.** When companies outsource parts of their IT infrastructure to cloud providers, they effectively give up some control over their information infrastructure and processes, even while they are required to bear greater responsibility for data confidentiality and compliance. While enterprises still get to define how information is handled, who gets access to that information and under what conditions in their private clouds (see sidebar for definition), they must largely take cloud providers at their word (or at their SLA) that security policies and conditions are indeed being met. The organization’s ability to monitor actual activities and verify security conditions within the cloud is usually very limited, and there are no standard, commercial tools to validate conformance to policies and service level agreements (SLAs).
- **Second, organizations running private clouds need to factor their cloud providers’ practices into their overall security and compliance assessments.** For these assessments to be thorough and reliable, organizations need to learn as much as possible about their cloud providers’ security policies, procedures, systems and controls, some of which may be different from or incompatible with their own. The general idea behind these detailed assessments is, although an organization may have very little visibility into the cloud provider’s operations and to actual states within the cloud, the organization can take assurance from verifying the cloud provider’s business practices.
- **The cloud introduces new risks resulting from co-residency, which is when different users within a cloud share the same physical equipment to run their virtual machines.** Creating secure partitions between co-resident VMs has proven challenging for many cloud providers. Challenges range from the unintentional – such as when a VM’s activities consume so much processing power and memory that it starves co-resident VMs of resources – to the deliberately malicious, such as when malware is injected into the virtualization layer, enabling hostile parties to monitor and control all the VMs residing on a system.
- **Finally, cloud services are typically virtualized, which adds a hypervisor layer to the traditional IT services stack.** Any new layer in the services stack introduces new opportunities for improving security and compliance, as well as new planes of exposure to risks. Organizations must evaluate the new monitoring opportunities and the evolving risks presented by the hypervisor layer and learn to account for them in policy-setting and compliance reporting. Later in this paper, we present several ideas for improving virtual-layer security for compliance.

For most organizations using cloud services today, the new compliance procedures introduced by the cloud are manageable. Mostly, that's because the data and applications running in the vast majority of clouds aren't governed by strict information security policies. The tougher challenges in cloud compliance won't arrive until higher-value business functions and data begin migrating to private clouds, creating stronger requirements for cloud security in order to prove compliance.

Most IT industry observers believe that it's inevitable for higher-value information and business processes to migrate to the cloud. The rate of migration should accelerate as the cloud's already-significant advantages in efficiency, cost and flexibility continue to increase, as compared to keeping those same IT capabilities in-house. As more and more sensitive data and business-critical processes move to cloud environments, there will be many wide-ranging implications for organizational leaders responsible for compliance and for information security.

Cloud Infrastructure: the Next Frontier in Cloud Security & Compliance

While an organization's compliance and security policies won't change when IT processes are shifted to the cloud, the way an organization enforces those policies and proves compliance will change dramatically.

For most compliance officers and information security professionals, the virtualized environment of the cloud is a "black box": it's hard to see into the cloud provider's services infrastructure to confirm that conditions are as they should be. Today, organizations rely primarily on cloud providers' reports, as well as certifications from outside auditors, to monitor performance against SLAs and to prove compliance in their private clouds. There can be considerable variability among cloud providers' and auditors' methodologies, and it's questionable if they're providing adequate evidence of the conditions that matter. In most cases, attestations prove compliance with a general, standardized process when what many organizations are looking for is proof that cloud providers' practices comply with their own internal policies and compliance guidelines. Sometimes, if organizations are running sensitive data or an important business process on private clouds, they may send in their own IT staff to directly examine their cloud provider's facilities, equipment, logs and incident reports. Conducting such inspections, however, can be very onerous. Making sense of the collected information can be even more so, as it's prohibitively difficult to isolate activity and performance data relating to an organization's private cloud from unrelated data about other tenants' cloud activities. The challenge is further compounded by the fact that true compliance is not a periodic or one-time evaluation. Instead, compliance monitoring should be continuous, so that errant conditions or problems can be recognized and addressed immediately.

Some organizations have deployed specialized software tools and management consoles that automate governance, risk and compliance (GRC) monitoring and enforcement across both physical and virtualized IT environments. These sophisticated GRC consoles aggregate all the disparate data and work flows affecting an organization's GRC posture and provide a summarized, contextual view of those conditions through a central dashboard. Such tools can coordinate and automate an organization's compliance processes, helping to identify and mitigate risks, reduce errors from manual processing, streamline compliance reporting and reduce the cost of audits.

For most compliance officers and information security professionals, the virtualized environment of the cloud is a "black box": it's hard to see into the cloud provider's services infrastructure to confirm that conditions are as they should be.

Although automated GRC tools can be very useful in providing a unified view of compliance processes within an organization's in-house and virtualized IT environments, such tools are still limited in their ability to report on conditions and activities within the deeper technology layers of the cloud. Regardless if it's a GRC console pulling summarized data feeds directly from a cloud provider or if it's an organization's IT staff manually sorting through their cloud provider's logs and reports, the information they're analyzing is generated at the application level or, at best, from operating systems (OS). There's currently no easy way for organizations to accurately monitor security conditions and operating states beneath the OS level of their private clouds. At sub-OS depths, we go dark. Even cloud providers are hard-pressed to provide VM- and hardware-level attestations for their private cloud clients, as it can be very laborious and time-consuming to isolate performance data or event logs generated only by the virtual machines or logical partitions that service an organization's private cloud.

This raises an obvious question: can an organization realistically control risks and attest to security in its private cloud infrastructure when there's almost no visibility into that infrastructure or what's actually happening within it? The simple answer is no.

The next frontier in cloud compliance will be attesting to the security of the cloud's virtual and physical machines: infrastructure-as-a-service (IaaS) components such as virtual machines, hypervisors and computing systems. Verifying the security of IaaS components is important for the simple reason that if organizations can't attest to the safety of their computing infrastructure, the security of all the data, software and services running on top of that infrastructure falls into doubt. Organizations must be able to verify the security posture of the entire technology stack, from the foundational hardware all the way through top-level application services.

Hardware Root of Trust: the New Sheriff for the Cloud Frontier

We believe organizations using cloud services will, in the very near future, push cloud providers to better secure the hardware layer and provide greater transparency into system activities within and below the hypervisor. This means cloud providers should soon be able to 1) give organizations greater visibility into the security states of the hardware platforms running the IaaS for their private clouds, 2) produce automated, standardized reports on the configuration of the physical and virtual infrastructure running customers' virtual machines, 3) provide measured evidence that their services infrastructure complies with security policies and with regulated data standards, 4) provide more granular views regarding their performance against SLAs, particularly as they relate to the allocation of cloud resources and the enforcement of co-residency restrictions, and 5) allow more customized, flexible provisioning of secure computing resources with the ability to monitor usage and bill accordingly.

While security of IaaS-level components have been relatively safe in the past, data centers are increasingly finding their servers under attack – not just by the more common viruses and Trojans, but by more sophisticated, coordinated security threats. New types of malicious software, or “malware,” have emerged that specifically target the foundations of the cloud: virtual machines and physical servers. Two of the better-publicized examples are the Blue Pill¹, which was first presented at the Black Hat Conference in 2006, and SubVirt², a lab project developed by University of Michigan researchers with funding from Microsoft. Blue Pill and SubVirt² are both examples of a class of malware called “virtual rootkits,” which shim themselves between operating systems and the system hardware. Virtual rootkits effectively masquerade as a virtual machine manager or hypervisor, shielding

¹ For more about the BluePill rootkit, please visit researcher [Joanna Rutkowska's website](#), InvisibleThings.org.

² For more about the SubVirt rootkit, see “[SubVirt: Implementing Malware with Virtual Machines](#)” by Samuel King & Peter Chen, University of Michigan; Yi-Min Wang, Chad Verbowski, Helen Wang & Jacob Lorch, Microsoft Research.

themselves from antivirus scans and other forms of detection, practically all of which are software-based and rely on the hijacked OS. Once in place, virtual rootkits can intercept any function of the operating system – such as someone entering a password – while performing almost any illicit activity imaginable.

It's important to note that such malware threats are very difficult to pull off, and there have been no publicized infiltrations by virtual rootkits to any major cloud (or, at least, none that anyone has admitted to yet). Nevertheless, cloud-based malware threats are evolving. As companies expand their use of clouds and as the exploitable value of information and business transactions handled within clouds continues to grow, it's only reasonable to expect that clouds will become stronger magnets for malware attacks. And because sub-OS attacks have such potent abilities and are difficult to detect, it's very likely that the cloud's infrastructure layer will become a highly desirable attack vector.

Cloud providers and the IT community are already preparing to address this problem. Entire ecosystems of technology companies are banding together to develop a new, interoperable, trusted computing infrastructure that leaves no room for "shimming" by malware such as virtual rootkits. The foundation of this new trusted computing infrastructure is the hardware root of trust, which establishes a bottoms-up security posture based on hardware components embedded with inalterable security technology.

In servers equipped with these secure computing chips – Intel's Trusted Execution Technology being one example – the embedded security technology examines and measures all processing components to attest to a trusted profile every time the server is turned on or reset. This "measurement," or expected security configuration, is stored not in software but in secure cryptoprocessors derived from the Trusted Computing Group's industry-standard, interoperable specifications. Collectively, the secure processing platform – defined in silicon and virtually tamper proof – creates a "hardware root of trust" that's used to authenticate each and every step of the boot sequence, from initializing the BIOS and launching the hypervisor all the way to the loading of the OS and application software. At each step in the process, the hardware root of trust measures the configurations of the hardware and software attempting to join the secure system. It then compares these measurements of actual configurations against the approved configurations defined for the secure system, which are stored within the hardware platform. If the measurements match up, then the hardware root of trust cryptographically "signs" each component, preserves a snapshot of the measurement and adds the hardware or software component to the secure chain of trust.

Because the hardware root of trust authenticates each and every part of the secure system, it eliminates the possibility that malware such as virtual rootkits can infiltrate the OS and penetrate the virtualization layer: the foreign software would instantly be rejected for not meeting the system's secure, recognized configurations. As clouds become more lucrative targets for attack and virtual-layer malware evolves and spreads, we anticipate the use of hardware roots of trust to validate each element in the secure computing stack will eventually become widespread, even commonplace.

In addition to inoculating the cloud against virtual-layer attacks, systems built on a secure chain of trust have many other valuable uses. The IT community has already begun evaluating a broad range of applications and services that could be built using the unique, measured capabilities of these highly secure systems.

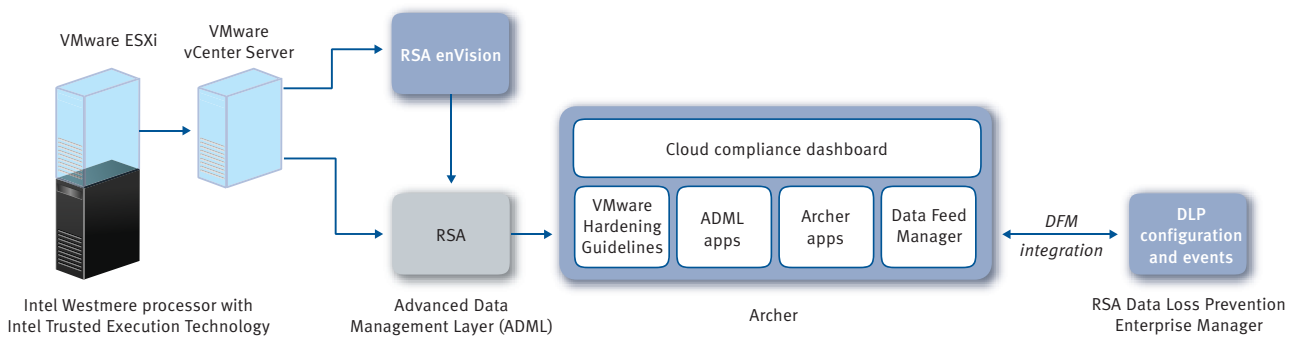
One of the most promising avenues of collaboration integrates the secure capabilities provided by a hardware root of trust into adjoining hypervisors and virtualization management software. The resulting integrated solution feeds highly precise data on physical hardware and virtualization layer conditions into specialized data analysis tools, which examine the data for events and conditions affecting security and compliance. The information is then evaluated in the context of larger business requirements by GRC software, which presents a unified, policy-based assessment of the organization's security and compliance posture through a central dashboard. This conceptual framework for a measurably secure

cloud computing environment – comprising a hardware root of trust, trusted virtualization environment, security information and event management tools and GRC management software – provides truly unprecedented visibility into actual conditions within the bottom-most layers of the cloud.

Archer Technologies, Intel, RSA and VMware have developed a proof-of-concept for a measured chain of trust, which demonstrates how to improve infrastructure-layer transparency and simplify security enforcement and compliance reporting for both internal clouds and private clouds.

One of the most promising avenues of collaboration integrates the secure capabilities provided by a hardware root of trust into adjoining hypervisors and virtualization management software.

Proof of Concept for Measuring and Monitoring Cloud Infrastructure Security



Cloud Gazing: What’s on the Horizon in Cloud Security & Compliance

Centralized controls and the specialized expertise of cloud services providers will enable security technologies for the computing infrastructure to be deployed far faster and more efficiently in cloud environments than if those same technologies were deployed in traditional enterprise IT environments. In fact, we believe by the end of this year, cloud providers will be able to introduce the first IaaS clouds built on measured trust environments. These new, highly secure clouds will give organizations more flexible, affordable and efficient alternatives for shifting high-value business processes and data into private clouds.

Although not every organization will need the high security afforded by a trusted computing environment, every organization using cloud services could benefit from the vastly improved control and transparency that a measured chain of trust enables. Simply being able to verify conditions in the cloud services stack down through the hypervisor is a huge step forward in providing visibility into actual states and activities within the cloud and in better regulating how cloud resources are managed. Internal and private clouds built on a measured chain of trust will 1) strengthen an organization’s ability to enforce differentiated policies in private clouds, 2) enhance monitoring for compliance at all layers within the cloud, 3) streamline the auditing process, and 4) allow for more flexible usage and billing for secure computing resources.

Here's a closer look at some of the foreseeable advantages of building cloud services on a trusted chain of computing resources:

- **Improve co-residency security by ensuring the launch of only trusted code.**
Protecting against untrusted software isn't just about malware; it also applies to more benign conditions, such as the improper migration or deployment of virtual machines. To illustrate, if load balancing software or a cloud administrator attempts to move virtual machines from an unsecured computing platform to a secure, trusted one, the secured platform would “red flag” the incoming VMs for not measuring up to the system's accepted configurations.
- **Prevent the unsafe transit of secure virtual machines.**
In the same way that unsecured VMs would not be allowed to move to secured platforms, secured VMs would not be allowed to move to unsecured ones. If, for instance, an administrator attempted to transfer a secured VM onto a new server, the virtualization management console would first perform a policy check on the outgoing VM and then measure the security configurations of the new server against accepted standards. If the new server couldn't meet the secure standards required to host the VM, the virtualization management console would block VM's move and log the attempt.
- **Maximize operational efficiency by creating trusted pools of systems.**
Once platform trustworthiness can be measured, cloud providers can put such measurements to use in building trusted pools of systems, all with identical security profiles. Hypervisors can then make more efficient use of secure clouds, moving VMs within zones of identically secured systems for load balancing and other administrative purposes – all while protecting data in conformance with regulated standards and policies.
- **Build secure clouds customized to comply with the most rigorous requirements.**
The secure cloud's ability to map high-trust zones of systems will enable organizations and cloud providers to customize their clouds to comply specifically with PCI DSS, HIPAA or other highly controlled information standards. Then, trusted pools of cloud-based resources – all compliant with the same set of information standards – could be dynamically allocated to optimize workloads. Such a scenario would extend the cloud's efficiency and scalability benefits to even the most strictly controlled business processes and heavily regulated industries.

Furthermore, cloud services could be fine-tuned to provide different levels of data security. For instance, two clouds could be proven HIPAA-compliant, with one cloud tuned to provide lower-level security at a lower cost for data such as patients' insurance information. The other HIPAA-compliant cloud, handling sensitive health information such as patient medical histories, could be tuned for maximum security. By tailoring cloud service levels, security and pricing to the value of information handled within each cloud, organizations provisioning private clouds can buy only what they need, making the cost benefits and business case for moving into the cloud even more compelling.

- **Create more granular co-residency controls through system mapping.**
Cloud services built on trusted measurements could also enable more granular controls to prevent co-residency conditions that are simply undesirable. For instance, in a cloud provider's environment, policies could be set within the virtualization management console to ensure that the VMs of two competing pharmaceutical companies couldn't simultaneously run on the same server. Instead, each company's VMs would be mapped to their own trusted pools of systems. In a similar vein, within an organization's internal cloud, policies could be set to preclude potentially risky co-residence among VMs from different business units with different privilege levels.

Taking this one step further, using IaaS built on a measured chain of trust could also help address another emerging threat vector in the cloud: “exfiltration” attacks. These are attack strategies in which hostile parties use IP addresses to identify the servers hosting a targeted organization's VMs,

then wage denial-of-service attacks to force the organization to provision additional VMs to handle the workload. As the organization brings additional VMs online, the attackers simultaneously provision VMs from the organization's cloud provider, greatly increasing their chances of collocating at least some of their VMs with the targeted organization's. Once they're running on the same machine and operating under the same hypervisor, the attacker's hostile VMs are able to monitor the activities of the target organization's VMs. Conceivably, hostile VMs could also gain access to shared system memory and corrupt it to "poison" or exploit other VMs on the hypervisor. Although exfiltration attacks are resource-intensive and fairly uncommon, they demonstrate the potential co-residency risks when organizations can't observe what's happening within the IaaS layers of the cloud. By offering trusted cloud services mapped to different security levels or zones, organizations can pick the service most suitable for their needs and reduce the risk of exfiltration attacks by either specifying security policies that make it difficult for untrusted parties to manipulate collocation or by specifying dedicated hardware to eliminate co-residency risks altogether.

– **Assess performance against SLAs with greater accuracy.**

Organizations could use the detailed measurements and attestations now available from the hardware and virtualization layers to assess how their cloud providers are performing against SLAs. Because the system is reporting on actual configurations and states, such reports can be useful, credible tools for ensuring cloud providers are living up to their commitments.

– **Provide unprecedented proof of compliance, even in the deepest levels of the cloud.**

Because trusted zones of computing resources can provide verifiable measurements of actual conditions at all levels of the infrastructure, including the hardware and hypervisor layers, they are, in effect, a powerful new tool for proving compliance. The hardware root of trust attests to the secure state of the hypervisor, verifying that virtual machines are running on the right hardware systems and comparing measurements of actual activity in the system with what's expected and acceptable. The secure hypervisor, in turn, is able to allocate computing resources within trusted zones, as well as monitor and log activities and events in the cloud. Collectively, the trusted hardware and virtualized platform can tell compliance officers and auditors just about anything they need to know – with verified metrics.

That said, analyzing and reporting the data is just as important as collecting it. IT companies are developing analytics software and data reduction tools to sort through the new mountains of data that will emerge from trusted cloud platforms. Ultimately, the goal would be to generate a concise summary of the most important events, which would then feed into a GRC dashboard that could frame cloud compliance within the organization's overall compliance posture.

While the advent of trusted computing environments built on a measured chain of trust is not a cure-all for cloud security and compliance, it does mark a very important milestone in the evolution of cloud services. The hardware and virtualization layers, formerly a "black box" within the cloud, now

Analyzing and reporting the data is just as important as collecting it. Ultimately, the goal would be to generate a concise summary of the most important events, which would then feed into a GRC dashboard that could frame cloud compliance within the organization's overall compliance posture.

become as inspectable, analyzable and reportable for compliance as the cloud's top-most application services layer. With this unprecedented level of visibility, clouds can develop the infrastructure-level policy controls and the end-to-end security attestations to handle even the most demanding security requirements for applications and data. Ultimately, this means that organizations will be able to take advantage of the cloud's benefits in supporting a much broader range of business processes.

Practitioner Guidance for Improving Cloud Infrastructure Security

Until IaaS offerings built on hardware roots of trust can be commercially developed, it will remain challenging for organizations to ensure security within the equipment, networks and virtual machines of their clouds, because there will continue to be very limited visibility into these levels, making actual conditions hard to verify. Nevertheless, just because you can't look for risks doesn't mean you can't be armed. There are several things that organizations can do right now to improve security and compliance within the hardware and virtual layers. Simultaneously, organizations can begin evaluating if and how they could benefit from deploying IaaS built on a measured chain of trust.

Harden all hypervisors

Despite some of the emerging security threats targeting the virtualization layer, hypervisors remain far less susceptible to malware and other problems afflicting general-purpose operating systems. Hypervisors, however, are not impervious to attack. One of the simplest, most effective steps that organizations can take to defend against unintended activities or malicious damage in the virtualization layer of their internal clouds is to "harden" their hypervisors. Organizations should also require the service providers running their private clouds to demonstrate they've conformed to recommended specifications for hardening their hypervisors.

Every hypervisor has its own recommendations for how best to secure the virtual stack. The most widely deployed hypervisors in the world, created by VMware, also offer a detailed set of hardening guidelines to mitigate security risks. In all, there are more than 100 specifications for hardening the VMware virtual infrastructure. They range from the simple (avoid accessing administrator functions through the ESX service console unless absolutely necessary) to the sophisticated (instructions for configuring the system to rotate or delete log files to keep from overwhelming the ESX Server Host). [The newest set of hardening guidelines](#) can be found on VMware's website.

Set clear policies for co-residency and be equipped to enforce them

Organizations will have to lay out acceptable conditions for co-residency for their private clouds. This will require that organizations first understand how their cloud providers enforce barriers between different customers and create separation of access. For example, how secure are the logical partitions created within a virtual system? What precautions do they take in setting up virtual LANs, storage areas and firewalls? How are virtual machines isolated between multiple zones?

Also, organizations may want to examine who their cloud providers' customers are, so they can identify potential business conflicts and set parameters for those with whom they shouldn't share hardware or virtual resources. Additionally, organizations should work with cloud vendors to ensure transferability of security controls. In other words, if and when data or virtual resources are moved to another server or to a backup data center, the security policies, including co-residency conditions, established for the original server or primary data center should automatically be implemented in the new locations. Finally,

specific performance and security metrics for co-residency should be written into managed service agreements and enforced with financial consequences if those agreed-upon performance conditions are not upheld.

Within internal clouds, organizations may want to evaluate co-residency rules for each business group. For example, the finance server may be required to run on dedicated hardware, while it may be acceptable to dynamically allocate computing resources between operations and sales.

Tips to Assess the Safety of Your Cloud Providers' Infrastructure

Organizations outsourcing portions of their IT infrastructure must be able to trust the companies providing them with cloud-based services. Trust cannot be granted on the cloud provider's reputation alone; it should be validated through thorough assessments to determine if the cloud provider needs to take additional steps to comply with the organization's information security policies and compliance requirements. Furthermore, performance conditions and standards must be written into SLAs and managed services agreements.

Following are some basic questions that organizations deploying private clouds should ask prospective cloud infrastructure providers.

- What practices do you employ to ensure safe co-residency? Have these practices been verified by a third-party auditor? If not, would you be willing to submit to an audit by either our staff or an independent auditor?
- Do you follow the hardening guidelines provided by your hypervisor vendor(s)?
- How do you ensure that my VMs and data are running only on authorized machines or data centers?
- What are the capabilities of the hardware supplied in your IaaS offering?
- Can I choose and control where my data is stored?
- Can I choose my own data encryption keys?
- Can I choose the key for protocols handling my data transfers?
- What monitoring tools do you use to manage your data centers?
- What are your processes for hiring privileged administrators, and how do you control their access?
- What certifications have you secured for your infrastructure?
- May we see a sample of your logs to gain a better understanding of what types of data can and will be reported?
- What specific audit rights, as well as liability controls and protections, do you typically offer in your managed services agreements?
- Are your data centers open for physical inspection? May we visit them if desired to assess physical environmental security?
- What provisions have you made for data mobility, in the event I choose to move from your infrastructure in the future?

The European Network and Information Security Agency provides a thorough checklist of security issues to consider when evaluating cloud services. [ENISA's Cloud Computing Security Risk Assessment](#) may be downloaded from ENISA's website.

Evaluate whether cloud vendors can deliver on their promises

Because information security is only as strong as its weakest link, it's essential for organizations to evaluate the quality of their cloud vendors. Having a high-profile "brand name" vendor and an explicit SLA is not enough: organizations must aggressively verify whether cloud vendors can deliver upon and validate their security claims. What do other customers have to say? What's their upgrade cycle or technology roadmap for investment? Some organizations have taken to conducting very thorough on-ramping audits for private cloud services providers. While such audits can be inconvenient for both parties, they are a necessary prerequisite for moving higher-risk data and business functions onto cloud providers' systems.

Assess cloud providers' methods for attesting to infrastructure security

Cloud providers have many different methods for attesting to infrastructure security. Some allow on-site inspections and penetration/vulnerability testing by their customers while others rely mostly on third-party security certifications. Regardless of how cloud providers may choose to prove the security of their infrastructure services — in other words, that hardware profiles meet requirements, that hypervisors are hardened, that there's safe, partitioned access to cloud resources and appropriate administrator access — organizations should remember that infrastructure security is not a one-time assessment. Especially given the dynamic nature of virtual machines, cloud providers must have an ongoing security posture that organizations understand and in which they can be confident.

The best way to verify infrastructure security is for organizations to demand maximum transparency into their cloud providers' operations. Log files and reports on administrator activities can be burdensome without context and some interpretation. For example, in addition to furnishing rosters of cloud administrators who have access to an organization's systems, can the service provider generate a timely report of what administrators tried to do? In addition to logging access events, are they analyzing the data for anomalies? For high-security clouds, perhaps those handling PCI or HIPAA data, can the service provider make regular attestations of IaaS security in a more automated fashion, perhaps through a dashboard application? Are they planning to deliver services built on a measured and trusted cloud infrastructure?

The best way to verify infrastructure security is for organizations to demand maximum transparency into their cloud providers' operations.

Request automated dashboard services for monitoring & compliance

Audit logging is critical to managing the security of any IT environment and a specific requirement of many government regulations and standards. Organizations deploying private clouds should coordinate with their various cloud providers to ensure the data needed to prove regulatory compliance is fed back into the organization. Cloud vendors' logs can be imported into the organization's security information and event management (SIEM) solution. This allows virtual events from the private cloud to be monitored and analyzed in the organization's central security operations console, alongside the organization's in-house IT infrastructure.

Additionally, the highly granular monitoring capabilities of hypervisors can provide a valuable view into what's occurring within the cloud. The challenge, again, is having the data analyzed, reduced and reported in a meaningful way. Cloud providers can deploy various monitoring tools to get a view of events in the cloud. Sometimes, these monitoring services can be summarized within a management console and extended to client-side environments.

Summary & Conclusions

The cloud's infrastructure will be the next frontier for improving information security. Verifying security at the bottom-most infrastructure layers of the cloud is essential to building secure conditions higher up the technology and services stack. After all, how can an organization claim to control risks and provide evidence of cloud security when there's almost no visibility into the underlying cloud infrastructure or what's actually happening within that infrastructure? IT companies and cloud service providers will soon be able to measure and monitor secure conditions within the cloud's physical and virtual machines.

Today, verifying secure conditions and operating states within the cloud's virtualization and hardware layers is prohibitively difficult. Usually, if evidence from these layers must be furnished for compliance, the auditing process has been extremely labor-intensive, ad hoc and time-consuming – in other words, highly variable and unscalable.

In order for organizations to move their higher-value business processes and regulated data into the cloud, they'll need to prove, on a continuous and reliable basis, that the infrastructure supporting their data and processes is secure. This cannot be achieved by using the impractical, one-off auditing processes of the past; instead, the IT industry and cloud providers will have to develop real-time monitoring solutions that provide configurable reports of actual conditions within the cloud. Getting detailed views of what's happening within the cloud's physical and virtual platforms will deliver several important benefits.

First, organizations using cloud services will finally gain the visibility and controls needed to prove compliance for even their most highly regulated data. This paves the way for organizations to migrate their mission-critical business processes and information to the cloud. It also allows them to finally take full advantage of the cloud's significant efficiency, cost and scalability benefits. Additionally, having a measured, trusted cloud computing environment would conceivably provide organizations' private clouds with these service enhancements:

- **Faster, more accurate and efficient auditing and compliance process** resulting from having measured evidence that their cloud services infrastructure complies with security policies and with regulated data standards
- **More granular views of cloud providers' performance against SLAs**, particularly as they relate to the allocation of cloud resources and the enforcement of co-residency restrictions
- **More customized, flexible provisioning of trusted computing services** with the ability to dynamically allocate resources, monitor usage and estimate spending on tiered levels of secure services
- **Finer-grained policy controls** over virtual machine environments and co-residency by being able to map virtual computing resources to trusted zones of secure systems

In turn, cloud providers would be able to create new, valuable service offerings by providing differentiated levels of security for various cloud environments. These variable security levels could be measured – with proof of service furnished – using automated monitoring and reporting tools that would penetrate to even the lowest layers of the cloud's infrastructure. Furthermore, having a secure computing environment that generates measured attestations of trust could ease cloud providers' pain points in meeting certain customer service requirements:

- **Customer compliance audits:** Streamlines the process for helping their customers prove policy compliance within the managed portions their private clouds
- **On-boarding audits:** Provides an automated, repeatable and more scalable way to demonstrate their conformance to prospective customers' selection criteria and to respond to detailed on-boarding audits

- **SLA performance:** Generates measured proof of performance against guaranteed service levels and other conditions covered by SLAs

Cloud providers and the IT community are gearing up to provide trusted computing platforms that build on the security benefits of a hardware root of trust. RSA, the Security Division of EMC, along with its newly acquired company, Archer Technologies, has collaborated with Intel and VMware to develop a vision for a measured, trusted computing environment that makes the cloud's bottom-most infrastructure as inspectable, analyzable and reportable for compliance as the cloud's top-most application services layer. We expect this conceptual framework to help improve visibility, control and compliance in the cloud and move the IT industry toward a more secure foundation for future cloud services.

About the Authors

Jon Darbyshire

President & CEO, Archer Technologies

Jon Darbyshire founded Archer Technologies in 2000 with a vision to create enterprise-wide IT risk and compliance management solutions that would replace traditional manual processes and disparate point solutions. Mr. Darbyshire's vision has evolved into Archer's award-winning enterprise governance, risk and compliance solutions built on the Archer SmartSuite™ Framework, which allows business users to create and tailor applications to meet their unique needs.

Sam Curry

CTO, Marketing, RSA, the Security Division of EMC

Sam Curry is the Chief Technology Officer for the Go-to-Market arm of RSA, the Security Division of EMC. Mr. Curry has more than 18 years of experience in security product management, marketing, product development, quality assurance, support, sales and marketing. Mr. Curry has also been a cryptographer, researcher and writer. Prior to his current role, he was Vice President of Product Management for two years, where he lead and set the strategic direction for all aspects of product management for RSA's solutions.

Douglas W. Fisher

Vice President, Software and Services Group & General Manager, Systems Software Division, Intel Corp.

Doug Fisher leads Intel's worldwide organization responsible for a broad range of development, architecture analysis and optimization efforts, including pre-boot firmware, operating systems, virtualization, middleware software, graphics, SoftSDV and client/server projections. In addition, he is the general manager for Intel's corporate initiative for virtualization. Prior to joining Intel, Mr. Fisher worked for 10 years at Hewlett-Packard.

Bret Hartman

Chief Technology Officer, RSA, the Security Division of EMC

Bret Hartman is responsible for defining the corporate security technology strategy for EMC, as implemented by the RSA division. Mr. Hartman has over 25 years of experience building information security solutions for major enterprises. His expertise includes Service Oriented Architecture (SOA) and Web Services security, policy development and management, and security modeling and analysis.

Dr. Stephen Herrod

Chief Technology Officer and Senior VP of R&D, VMware, Inc.

Steve Herrod is responsible for VMware's new technologies and collaborations with customers, partners and standards groups. Dr. Herrod joined VMware in 2001 and has led the VMware ESX group through numerous successful releases. Prior to joining VMware, he was Senior Director of Software at Transmeta Corporation co-leading development of their "Code Morphing" technology.

Vishal Kumar

Senior Product Manager, Emerging Markets and Technologies, VMware, Inc.

Vishal Kumar is a well-known authority in the areas of systems integration and security, having worked on various platform technologies and products at Microsoft and VMware. He is currently helping architect a secure next-generation cloud platform designed for public and private cloud workloads. Prior to VMware, Mr. Kumar was a senior manager at Microsoft, where he was responsible for security for several Microsoft product lines and helped create the Secure Development Lifecycle.

Fernando Martins

Director of Strategic Technologies and Corporate Virtualization Strategy,
System Software Division, Intel Corp.

Fernando Martins leads a team of strategists responsible for providing guidance to Intel's top executives on selected strategic technologies. He is also responsible for all aspects of Intel's corporate virtualization strategy, including defining roadmaps of silicon features and industry ecosystem relationships. Dr. Martins received Intel's highest honor – the Intel Achievement Award – twice for his contributions to Intel's virtualization strategy. Prior to Intel, Dr. Martins held appointments with Rockwell, Allen-Bradley, IBM Research and two successful startups. His work has been featured in more than 40 international publications, and he holds 23 patents.

Steve Orrin

Director of Security Solutions, Software Services Division, Intel Corp.

Steve Orrin is responsible for security platforms architecture, security strategy and product direction for SSG's SPI group at Intel. Prior to Intel, he held executive management positions at software companies specializing in security for web, SOA/XML and enterprise applications. Mr. Orrin was named one of InfoWorld's Top 25 CTOs of 2004. He was recently named a fellow at the Center for Advanced Defense Studies and is a recognized expert and frequent lecturer on enterprise security.

Dana Elizabeth Wolf

Senior Manager, New Business Development, RSA, the Security Division of EMC

Dana Wolf is responsible for creating and developing new security technologies and business opportunities for RSA from the Office of the CTO. She also manages CTO operations and RSA's advanced development engineering team. Ms. Wolf joined RSA in 2004 as a principal software architect and served two years as an Entrepreneur in Residence at RSA for her graduate school work on payment card security.

Solutions for a Secure Cloud Infrastructure

Cloud services are evolving rapidly, driven by an escalating threat environment and a raft of cutting-edge technologies and products. Organizations should continually assess new cloud services offerings and technology solutions to keep their internal and private clouds secure.

The products and services described below align with the best practices described in this RSA Security Brief. The solutions overview is not intended to provide a comprehensive list of applicable solutions from EMC, Intel, RSA or VMware. Rather, it's intended to serve as a starting point for compliance officers and security technology practitioners wanting to learn about some of the options available to them.

Governance, Risk & Compliance Solutions

- Archer SmartSuite Framework is a flexible software platform that is designed to help organizations implement a consistent, efficient and sustainable program for enterprise governance, risk and compliance. The framework incorporates nine preconfigured solutions that are engineered to allow organizations to manage the lifecycle of corporate policies, assess and respond to business risks, and measure and report compliance with controls and regulations. The Archer SmartSuite Framework is designed to provide a consolidated, business-level view of risk and compliance across the enterprise and reduces costs associated with GRC. Because it's easily customizable to an organization's unique business requirements, the framework has become a go-to GRC platform for many global companies.

Cloud Security Solutions

- The RSA® Advanced Data Processing Layer (ADPL) is designed to analyze and reduce high-volume data streams from a variety of sources, including the RSA enVision platform, VMWare vCenter Server, and Microsoft® Active Directory® service, to automate an organization's compliance assessments. RSA ADPL connects to Archer Technologies' innovative Archer SmartSuite Framework, providing summarized, contextualized reports of the most important data impacting an organization's compliance.
- The RSA® Data Loss Prevention Suite is engineered to provide policy-based approach to securing data in data centers, networks and end points, enabling customers to classify their sensitive data, locate and track data across the enterprise, enforce controls, and report and audit

activities to ensure policy compliance. The RSA DLP Suite is engineered to help reduce total cost of ownership with high scalability, automated data protection services and the most extensive data policy and classification library available in the industry. It features three components:

RSA® DLP Datacenter

RSA DLP Datacenter helps companies locate and track sensitive data no matter where it resides in the data center — on file systems, databases, email systems and large SAN/NAS environments.

RSA® DLP Network

RSA DLP Network is designed to monitor and control sensitive data leaving your network.

RSA® DLP Endpoint

RSA DLP Endpoint helps you discover, monitor and control sensitive information on endpoints such as laptops and desktops.

The RSA DLP suite works with the RSA enVision log management and analysis solution to simplify security operations by streamlining incident handling and workflows.

- The RSA enVision platform provides collection, alerting and analysis of log data that enables organizations to simplify compliance and quickly respond to high-risk security events. The RSA enVision *3-in-1* platform offers an effective security and information event management (SIEM) and log management solution, capable of collecting and analyzing large amounts of data in real-time, from any event source and in computing environments of any size. The RSA enVision platform is easily scalable, eliminating the need for filtering and to deploy agents. More than 1,600 customers, including major global enterprises and government agencies, have selected the RSA enVision *3-in-1* solution to simplify compliance, enhance security and optimize IT and network operations.

Hardware Root of Trust Solutions

- Intel® Trusted Execution Technology (Intel TXT) is a highly versatile set of hardware extensions to Intel processors and chipsets. In combination with other Intel solutions, it provides a higher, measured level of trust and control over computer systems. Intel TXT enables hardware-rooted security that helps defend systems against software-based attacks and protects sensitive information without compromising the usability of the computing platform.

- Intel® Virtualization Technology (Intel VT) provides “hardware-assisted virtualization” that boosts virtualization software performance and improves application response times. Intel VT reduces demands placed on virtualization software to help organizations consolidate more applications and heavier workloads per server, making the most of computing and software investments.

Virtualization Layer Solutions

- VMware® ESX™ and ESXi are the most widely deployed hypervisors in the world. Both allow enterprises to use their own security certificates when securing remote sessions. The user name, password and network packets sent to ESX Server over a network connection when using the VMware Remote Console or the VMware Management Interface are encrypted in ESX Server by default when medium- or high-security settings are activated for the server.
- VMware vCenter® Server gives IT administrators unprecedented visibility and centralized control of every level of the VMware vSphere virtual infrastructure. It provides granular privilege management that limits who can deploy virtual machines to specific clouds and storage devices. Combined with well-defined operational processes and work flows, these capabilities can provide maximum mobility for virtual machines while managing risk.
- VMware vCenter Lifecycle Manager enables IT administrators to track ownership of virtual machines and to keep records of when virtual machines are created, deployed and decommissioned. It gives IT administrators more control over virtual machine deployments and optimizes resource utilization for greater ROI.
- VMware vShield Zones enables enterprises to run applications efficiently within a pool of shared computing resources, while preserving network segmentation of users and data. It allows administrators to bridge, firewall or isolate virtual machines between multiple zones, as defined by organizational and trust boundaries. It also allows for convenient, centralized management by providing highly granular views of the entire virtual machine and virtual network deployment, easing configuration of zone-based policies and reducing the risk of errors.

- VMware vSphere™ is the market-leading virtualization solution that allows organizations to turn their infrastructure into an efficient and flexible internal cloud, potentially decreasing capital and operating costs by up to 60 percent while streamlining operations, increasing control over IT resources and improving flexibility. VMware vSphere employs federation and standards to bridge internal and external cloud infrastructures, helping organizations of all sizes achieve the full benefits of cloud computing. Furthermore, VMware vSphere 4 is Common Criteria certified at EAL4+, providing an objective measure of security that’s well understood by security professionals and compliance auditors.

Cloud Consulting Services

- EMC® Consulting provides strategic guidance and technology expertise to help organizations exploit information to its maximum potential. With worldwide expertise across organizations’ business, applications and infrastructure, as well as deep industry understanding, EMC Consulting guides and delivers revolutionary thinking to help clients realize their ambitions in an information economy. EMC Consulting drives execution for its clients, including more than half of the Global Fortune 500 companies, to transform information into actionable strategies and tangible business results. In particular, EMC Consulting provides consulting and technology deployment and optimization services to help organizations address their business and cloud infrastructure challenges. EMC Consulting’s infrastructure consultants draw on their deep expertise in EMC’s broad and deep solutions portfolio, industry-leading infrastructure technologies and IT operational best practices to help organizations accelerate their journey to the private cloud and develop and deliver IT services that support business growth and change. They help organizations reduce infrastructure costs and complexity and use those savings to develop solutions to improve data governance, optimize risk management and implement process and resource management automation.



RSA Security Inc.
RSA Security Ireland Limited
www.rsa.com

The Security Division of EMC

©2010 EMC Corporation. All rights reserved.

RSA, enVision and EMC are either registered trademarks or trademarks of EMC Corporation in the United States and/or other countries. Archer and SmartSuite are trademarks of Archer Technologies, Inc. in the United States and/or other countries. VMware, vCenter and vSphere are registered trademarks or trademarks of VMware, Inc. in the United States and/or other jurisdictions. Intel is a trademark of Intel Corporation in the U.S. and other countries. All other products and/or services mentioned are trademarks of their respective companies.

CCOM BRF 0310