Management Briefing:

Virtualisation: Don't let security lag behind technology maturity

Virtualisation is mainstream, with virtualised servers, storage, networking and software featuring in most enterprise IT departments. Applications and workloads are assigned to virtual machines (VMs) and moved from machine to machine as a matter of course, and as a result, deploying physical servers is now the exception rather than the rule.

While this trend promises significant productivity improvements, it raises the question of whether security practices are in danger of lagging behind the day-to-day reality of the IT department. Fortunately, there is a resolve amongst cloud service providers and technology suppliers to work together to ensure that virtualised environments are protected against the active and continual security threat to virtualised infrastructure.

However, end user organisations must also strengthen their virtualisation security strategies, says Forrester Research senior analyst Rick Holland. "The technology is mature and enterprise adoption is high, yet information security does not have a significant focus on virtual security. Given the converged nature of virtual environments, security incidents can result in significant damage; therefore, it is critical that security professionals redouble their efforts and make securing their virtual infrastructure a priority," he argues.

The scale of the threat

The virtualisation security threat is both broad and varied, with cyber criminals targeting users and operators of virtualised and cloud systems, including software-as-a-service (SaaS), platform-as-a-service (PaaS) and infrastructure-as-a service (IaaS). The threat can come from outside the service provider or organisation, or, worse still, from the inside.

Dr. Kevin Curran, senior member of the IEEE (Institute of Electrical and Electronics Engineers) and a reader in computer science at Ulster University, says that most resources to protect virtualised infrastructures are spent on protecting the perimeter network, but this is not an ideal strategy, because the approach offers reduced protection against internal threats.

"If one exploit succeeds in compromising a single virtual server then that compromised server could be used as a base to attack other servers on the same host." The key strategy is to protect the hypervisor – the software system that manages and orchestrates the VMs, says Curran.

But he adds, "It is increasingly difficult to protect virtualised infrastructures as the demand to create new virtual servers, to service dynamic workload spikes, can lead to simple administrative errors allowing backdoors to the attackers."

Rootkits are of particular concern (a rootkit being a maliciously-modified set of administrative tools for an operating system, granting 'root' access), as these can lead to the privileges on exploited systems to launch deep attacks. "It is also not helpful to administrators that modern rootkits are formed from sophisticated code where the signature is difficult to detect and thus defend against," notes Curran.

Another point of vulnerability within a virtualised infrastructure occurs where virtual machine images are being moved between physical servers. "When these virtual machines are residing on the network between secured perimeters, they are vulnerable to attack as hackers could plant malicious code so as to gain access to a destination data centre," says Curran.

Kate Craig-Wood, managing director and co-founder of technology services firm Memset, says the external cyberthreat to cloud and virtualised computing infrastructures is relentless. "We host 20,000 of Britain's busiest and largest websites. In just the last week our automated denial of service (DOS) protection system, affectionately known as the 'DOS-squasher', blocked just over 200 attacks aimed at our clients. None of them even knew there was a problem."

She adds, "On my personal server alone in the last week there were over 50 break-in attempts, all automatically deflected by a combination of good password choice, operating system lock down and firewalling. Multiplied up across our entire server estate that amounts to someone trying to compromise one of our customers' servers every few seconds."

In the case of Memset, its IT infrastructure is ISO27001 certified, as is the case with many enterprise-grade IT services firms and organisations. ISO27001 incorporates a systematic security risk and vulnerability assessment process, and implements information security controls and regular hardware and other checks.

"We recently passed a penetration test carried out by Incryption, who are certified by CREST and Tiger certified, who confirmed that our hypervisor layer is very secure," says Craig-Wood. Crest is a globally recognised competency certification organisations and Tiger is a commercial certification scheme for technical security specialists. Craig-Wood claims that the firm uses no security products at all, but designs its systems using open source software. "We achieve user segregation via VLAN, hypervisor layer and OpenStack. We passionately believe that open source is more secure," she says.

Memset is currently moving beyond ISO27001 to gain CESG (Communications-Electronics Security Group) Impact Level 2 accreditation, a government accreditation, which will allow it to provide IaaS/PaaS services under the Government's G-Cloud project, delivered through the Government's public cloud service which uses a Microsoft Azure cloud infrastructure.

Security accreditation and best practices are hot topics for the cloud and virtualisation industry, with vendors and users working together through industry groupings such as the Trusted Cloud Initiative (TCI).

The TCI recently released its Reference Architecture Model, and Reference Architecture Mapping, a methodology and set of tools to help organisations assess the security level of their internal cloud services, and that of their service providers. TCI's reference model is a vendor-neutral

architecture, used for the secure design and assessment of cloud infrastructure, and incorporating areas including governance, identity, authentication, authorisation and 'auditability'.

With initiatives such as these, cloud users and providers aim to boost the security of their individual infrastructures, and making trusted cloud-based networks both interoperable and secure.

Security technologies and techniques

As well as good security management and policies, there are a number of technologies that can be deployed to protect virtualised environments.

Organisations can use a combination of software and hardware approaches, for example using traditional network security appliances, or newer processor-based security and management features designed for virtualisation, such as that found in Intel's new Xeon E5 processors.

Dr. Curran says that when architecting virtualised infrastructures, it's still a good idea to concentrate on traditional security mechanisms such as malware and intrusion detection systems, alongside dedicated virtualisation firewalls on components such as the hypervisor. "Defending the global monitor against deactivation is crucial. Segmenting networks into isolated virtual clusters is to be recommended. Inspection of network traffic in real-time can lead to effective action, in the case of violated security policies," he says.

Curran adds that the hypervisor should be protected through access control, automatic updating, networking, and introspection on guest OSs (introspection is a technique for checking the security, physical location, patch status, etc, of individual VMs).

"Image management, especially with regards to migration, can be done through strong storage and network encryption so that sensitive data does not leak from the images. Finally, companies should not overlook setting file permissions, controlling users and groups, and setting up logging and time synchronisation, in addition to routine inspection for hardware failures and out of date systems in the physical infrastructure," says Curran.

Other security strategies include detecting new virtual servers automatically upon creation, which can prevent the arrival of 'rogue' instances; and enforcing a strict security patch management policy to protect the environment from new security vulnerabilities.

In addition, virtualisation infrastructure managers should monitor internal staff actions, says Curran. This can be achieved through simple logging of key actions taken by IT administrators across the entire infrastructure. This monitoring is also important for auditing purposes.

One firm that uses a multi-layer security strategy for its virtualised infrastructure is Ansarada. Financial and legal firms such as PwC, KPMG, Investec, and Linklaters, use Ansarada's ISO 27001-accredited virtual data rooms for multi-million pound deals.

As a result, Ansarada is required, by its customers and their industries, to maintain watertight security across its systems and for its end-users, says Harry Gill, head of European operations.

It uses a range of hardware and software security and procedures to help protect the virtualised

infrastructure. These include 256bit or 128bit encryption for documents; and enterprise-level encryption for data - both in transit and at rest.

The servers have multiple levels of redundancy, with no single point of failure, and run in several data centres. The two main reasons for this are to ensure security and availability, says Gill: "Clients want to know that their data won't get lost."

As for the deal room users themselves, they must be personally invited to the deal room by the lead participant, and each one has a secure login using a complex password, with the system using IP address tracking or only allowing users to connect using a specified IP address. Session times are logged for auditing purposes.

The company is considering deploying processor-level virtualisation security in the future, which will speedily encrypt data being used in its VMWare environment. "We try to use the best technologies in the marketplace, whether that's enterprise level encryption or solid state disks," says Gill.

"But the biggest issue with security is that there is a trade-off between speed and reliability and ease of use. You can have three or four factor authentication, with one-time secure passwords being sent to mobile phones, but it might make the system too slow for end users," he adds.

This trade-off is a key consideration for infrastructure managers. Effective hardware and software security technologies and techniques now exist to fully secure virtualised environments, and a multilayered approach is undoubtedly the best. But the big question remains: how much security is enough?