Whitepaper

Database Security, Virtualization and Cloud Computing

The three key technology challenges in protecting sensitive data in modern IT architectures

Including:

- Limitations of existing database security approaches
- Security considerations when deploying virtualization
- How distributed monitoring best fits virtual and Cloud Computing environments

Provided by sentrigo™

Version 1.0, January 2010

www.sentrigo.com

Securing Information in Virtualization and Cloud Computing Architectures – Business as Usual?

Since virtualization is now in use across nearly all enterprises, and future plans to move some applications to Cloud Computing are also in the making, it is time to ask whether we need to update our IT security methodologies or continue to rely on the same tools we have relied on in the past.

Many would agree that these computing models require an entirely different approach. Often, it seems network security models lag behind the technology changes occurring in the systems and software environment they are intended to protect. For example, it took firewall vendors a long time to realize that applications are no longer easily protected by simply opening and blocking ports. As Enterprises began deploying VoIP and other complex protocols, they typically had to wait years before their firewall vendor allowed them to do so securely by analyzing SIP and other protocols.

This whitepaper suggests that with the recent changes in computing architecture from dedicated servers in the datacenter to virtualization and Cloud Computing, we need to rethink our IT security methodologies. And, while we focus in this paper on database security, many of the recommendations made here are applicable to securing most enterprise applications.

IT Security Trends – The Era of Appliance-based Solutions

In the last few years, many IT Security challenges centered around two major technological developments: very high performance networks and complex applications. In a very short time we moved from 10baseT networks with SMTP, FTP, HTTP and a few other relatively simple protocols to multi-gigabit networks with SIP, RPC, SOAP, various tunneling protocols and more. Some applications (in particular e-mail and web applications) were subject to many types of security threats, the result being that the same applications had to be checked multiple times, by multiple products.

To meet the challenges of high performance complex applications with relatively simple implementations, security companies introduced a range of network appliances – machines that could be positioned somewhere in the network (well, not quite, but we will leave the discussion of the art of positioning appliances for another paper) and inspect the traffic for either protocol violations, or malicious code, or viruses, or spam, or...

Even traditional software-oriented companies understood that often the easiest way to cope with complex, high performance networks was by deploying network-based appliances. Enterprises typically found themselves implementing multiple types of appliances – each aimed at mitigating a specific threat vector – and often multiple quantities of each type to handle requirements for scalability, performance, or network topology. This approach resulted in a security environment where customers are concerned with the sheer number of appliances they need to maintain and manage in order to inspect their traffic for all threats.

In this era of security appliances, solutions that met security threats utilizing host-based software models were by and large neglected. Network-based IDS and IPS won the battle against host-based solutions, and most enterprises do not add much to the OS-provided simple endpoint security (one noteworthy exception is 3rd party antivirus software that most enterprises continue to deploy on PC's even though they may have appliances at the perimeter providing protection). The concept of simply placing an appliance (or even a few appliances for that matter) in a rack and attaching it with a wire or two to a switch is very attractive, especially when resources available for security are limited.

And for the first generation of systems architectures, the assumptions surrounding security persuaded enterprises that network appliances were an acceptable solution. However for many applications today, especially those implemented in a distributed model, capturing application transactions on the network only identifies the majority of external threats. Driven by an increasing number of breaches by privileged insiders, and greater sophistication by external hackers, regulations are now requiring broader coverage of these threat vectors. Enterprises concerned with insider transactions (e.g. administrators working directly on servers, applications users abusing authorized access, etc.) have been leading the adoption of host-based solutions, either in conjunction with network appliances, or as their primary approach.

In addition to the increasing concerns about the insider threat, now that more and more applications are collapsing into virtual machines, (and in the case of Cloud Computing, outside the enterprise perimeter), it undermines the past assumptions that led to appliance-based security deployments. Even more challenging from a security perspective, these new databases may dynamically appear in new locations over the course of time, based on an organization's changing capacity requirements. These new architectures beg the questions of whether the network appliance approach will still be relevant when many transactions will not make it to the network, or whether a network monitoring approach is efficient when the application network moves from LAN to WAN.

Database Activity Monitoring and Attack Prevention – Case in Point

Perhaps owing to the fact that even networks apparently secured at the perimeter were still suffering breaches of sensitive data, in recent years enterprises started looking at providing aother layer of protection on their internal infrastructure. Whether to provide a safety net in the case of a perimeter breach, or to protect from malicious insiders, something more was needed. Previously, databases in most cases were not monitored or protected. For a number of reasons, including the prevalence of database breaches and stricter regulations regarding prevention and notification, customers are now investing a lot more time and effort in securing their databases.

It is not surprising that when IT security professionals were first faced with the challenge of protecting their databases they looked for the same solution they used for protecting their other assets: a network appliance. And indeed, several vendors quickly brought out appliances that inspected database network protocols and allowed auditing and protection of the database when accessed via the network. Enterprises were at first willing to forego the lack of visibility into database transactions occurring locally and from within the database servers. Later,

because of the large potential for damage that can be done on the local machine (e.g. via an SSH connection), it became clear to enterprises that to fully understand the threats to their databases, monitoring must also cover local and intra-db attacks.

At this point, appliance vendors were forced to add local agents to their solutions, making many of today's network-based solutions a hybrid network appliance and host-based solution. In most cases, these agents send local traffic back to the appliance for analysis – where each transaction that was originally done on the local host is measured against the appliance's policy. The hybrid approach is not ideal (for example, local access in breach of the security policy cannot be efficiently blocked since by the time it reaches the appliance it has already been processed by the database), but as long as most applications run on the network in plain sight of the appliance, some enterprises were willing to accept the risks.

These hybrid solutions lose many of the benefits of a pure network-based solution by introducing significantly more complex implementation requirements such as kernel-level installation of the agent, for example, requiring reboots to the DB server. And, as noted above, they still miss the sophisticated attacks generated from within the database itself – those attacks based on stored procedures, triggers, and views. However, more importantly, they also fail to address several key technical challenges when implemented in either a virtualized environment, or in the cloud. The next sections will introduce these challenges, and demonstrate how an architecture designed for distributed monitoring at the database memory level instead of network monitoring can best address them.

Challenge #1 – Visibility Into VM-to-VM Transactions

As organizations deploy applications and the databases that support them on virtual servers and in the cloud, a new complication frequently arises. In the past, an application was typically provisioned on one or more servers, and the databases housing the application were installed on separate networked servers. One of the benefits of virtualization (in a private datacenter or in the cloud) is the ability to share resources, resulting in environments where both the application and the databases are migrating to virtual machines, in many cases running on the same physical servers. In the diagram below, note that communication from the CRM application to the database storing customer data occurs entirely within the same physical server. In such a case, there is little to no network traffic as the transactions between the application and the database occur from VM-to-VM within the server. Network monitoring appliances will not see these transactions.

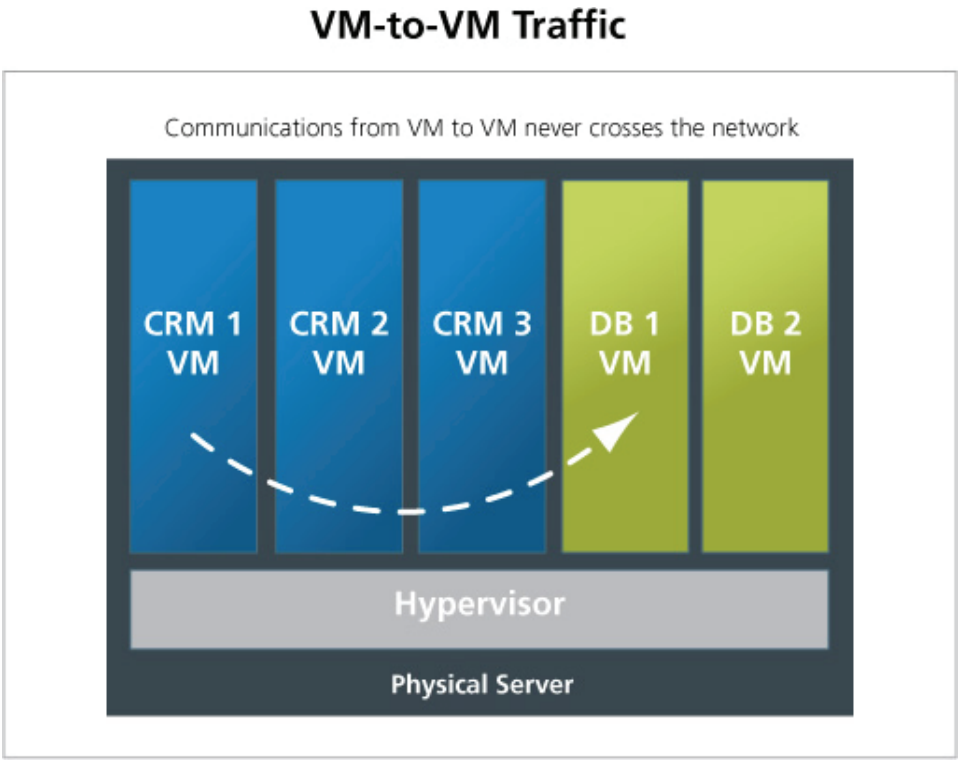

What can be done in such a case? Clearly the only solution is bringing the security inspection closer to its target. One solution is dubbed the "virtual appliance" where a virtual machine that runs the software formerly run by a dedicated appliance is installed on virtual servers and the servers are re-architected to send traffic through the virtual machine. This approach has two severe drawbacks: performance and architecture complications.

The performance problem of "virtual appliances" is as follows: dedicated appliances have the advantage of being able to cope with the enormous volumes of back office traffic (using dedicated NICs, dedicated hardware, and optimization of the software to fully utilize the dedicated hardware). When the software is running on a virtual machine, all the advantages are lost and the result more often than not is either a bottleneck that slows down databases when positioned inline (all database transactions pass through the virtual appliance), or missing a large percentage of the transactions when positioned outside the transaction path.

Challenge #2 – The Dynamic Systems Environment

The architecture complications arise from the fact that instead of creating a dynamic environment on virtual servers, enterprises need to plan ways in which all access to databases will pass through the virtual appliance. This is a complication of a problem that exists in standard (i.e. not virtual) networks as well. In most enterprises networks are not architected in a way where transactions to all databases can be monitored in a single location, so that enterprises face 3 difficult choices – re-architecting their network, using multiple security appliances, or protecting only some of the databases.

If virtual security appliances are far from an ideal solution for enterprise networks that run on virtual machines, they are even less relevant in cloud-based applications where networks are dynamic, hosts come and go, and adding virtual appliances to the mix is virtually impossible.

The only solution that works in all environments, including Cloud environment is a solution that comes up (and down) with every database – a solution that is based on sensors that run side by side with the database on every host machine that runs one database or more.

Challenge #3 – Performance Over Wide Area Networks

Especially in Cloud Computing deployments, but also in geographically distributed "private clouds based on virtualization, network bandwidth – and more importantly, network latency – will render off-host processing too inefficient. The basic concept of Cloud Computing prevents you from being able to co-locate a server close to your databases – you simply won't know where they are. Therefore, the time and resources spent sending every transaction to a necessarily remote server for analysis will bog down network performance, and prevent timely interruption of malicious activity.

To be efficient, the agent or sensor must be capable of implementing the necessary protection and alerting locally. This will ensure that the network doesn't become the gating factor for application or performance. For Cloud Computing (or in fact, for remote management of distributed datacenters), you'll also want to ensure that the system is designed to support wide area network (WAN) topologies. To limit exposure of sensitive data, it should encrypt all traffic between the management console and sensors, and for optimum performance compression techniques should be implemented so that policy updates and alerts are efficiently transmitted. Any architecture that requires off-host processing of database transactions to determine policy violations is unlikely to be responsive enough to prevent attacks.

Distributed Host-based Software Solutions – the Only Efficient Approach

The only way to secure databases on virtual machines or cloud environments, without sacrificing the huge benefits of these new architectures, is using software-based solutions that share the elasticity of virtual machines and Cloud Computing. The challenge is creating host-based solutions that do not suffer the same drawbacks that made old host-based solutions irrelevant. Namely, intrusive implementations, performance issues and the need to quickly adapt to new and changing environments (new OS versions, new application versions, etc.).

Next generation host-based solutions must not be based on kernel-level implementations and other risky / intrusive approaches, as this presents obstacles for deployment and management (requiring machine restarts, very sensitive to environment changes, etc.) Instead, the solutions have to be lightweight, user space software that can be easily added to the virtual machine where needed, and installed in parallel to the first database that is installed on a machine. This means that adding a layer of security does not require changes in architecture and does not rely on the virtualization technology in use. Moreover, as VM's running your database are provisioned (and de-provisioned) to balance capacity needs, no manual intervention is required at the management console.

Dynamic Provisioning in the Cloud

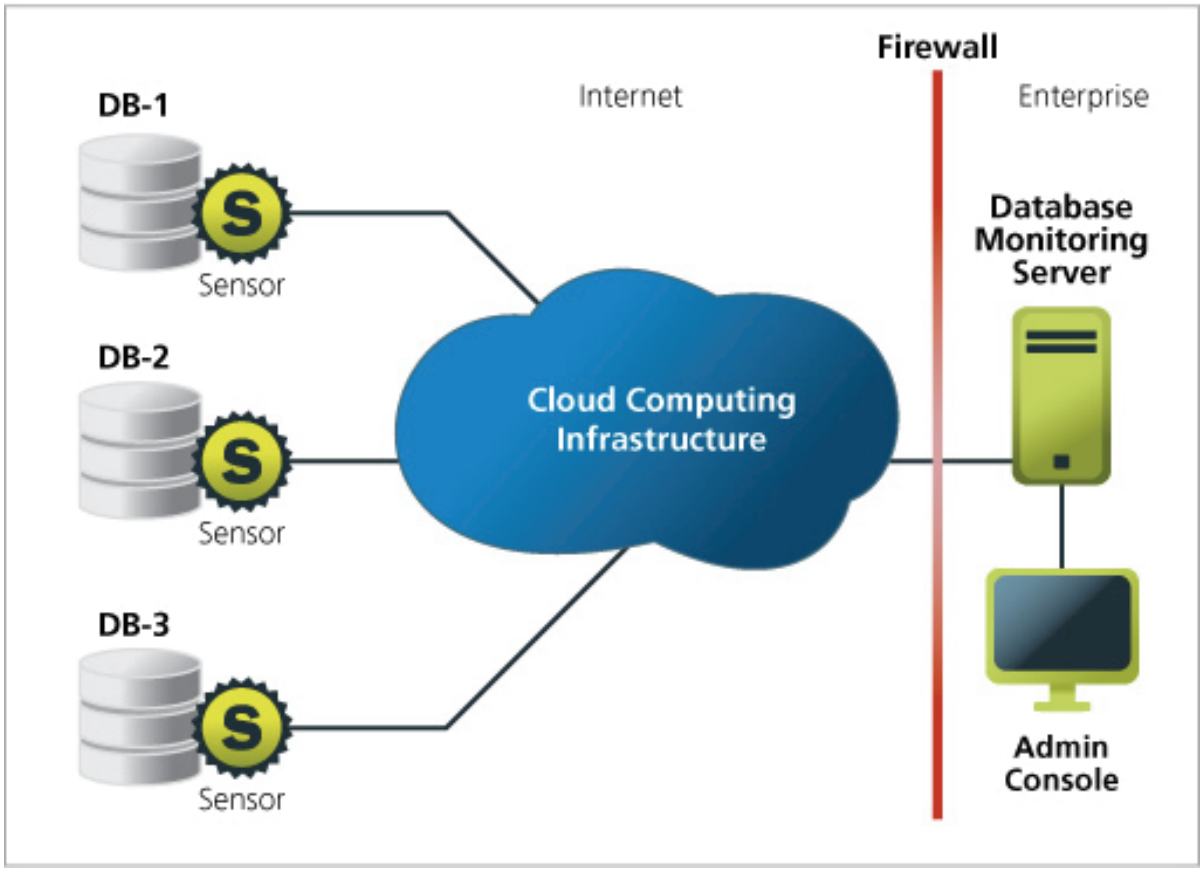

Sentrigo's solution was designed to meet these database security needs by utilizing a software-based sensor, a lightweight add-on that gets installed on the same virtual machine where the database or databases are installed. The sensor adds a process that monitors the database transactions as they occur in the memory. The memory-based sensors cover all the security requirements without the drawbacks – running with insignificant performance impact and no changes in the way the network is architected or how the databases function. In addition, by monitoring memory, the sensor protects against all attacks. Whether originating on the network, from a local privileged user, or even from inside the database itself via a sophisticated intra-DB attack, in the end, Sentrigo's local sensor will see the attempted exploit, and take action.

While the sensors are lightweight, they have all the logic required to make split second decisions on the legitimacy of databases transactions and can monitor, prevent attacks and audit database transactions in very much the same way they protect databases on regular servers. A sophisticated central management controls and receives information from all sensors. The only information that a new sensor needs is the location of the central management server. Whether the database is on a virtual machine, directly on a physical server, or somewhere in the cloud, as long as a sensor is installed on the same machine and the sensor has logical access to the central management, enterprises can enjoy a monitoring and attack prevention system that is not influenced by the underlying network or servers.

Sentrigo – Ready For Virtualization and Cloud Today

Many organizations have found themselves drawn towards virtualization and Cloud Computing architectures for their many benefits, only to find that the complexity of ensuring adequate data security was simply too great an obstacle. But, the movement towards these technologies is inevitable.

By deploying memory-based solutions for distributed database monitoring, enterprises will find that it is not only possible to protect sensitive information in these emerging computing models, but that the same architecture also provides both more effective and more efficient data security across their dedicated database servers as well.