

# RSA<sup>®</sup>CONFERENCE2009

## Incident Handling in a Virtualized Data Center

### Brandon Gillespie


Virtualization and Storage Architect  
Contractor, Systems Implementers Inc

### Mike Neri

Deputy Director and CTO for the 75 Communications Group (Provisional)  
Ogden Air Logistics Center, Hill Air Force Base, UT

04/23/09 | Session ID: STAR-301

Session Classification:



# Disclaimer

---

This session does not represent endorsement by the US Air Force, Hill Air Force Base or any other component of the Department of Defense.

Approved for public release; distribution unlimited



# Origin: Hill AFB and Project Bonfire

---

- Modernization of the Datacenter
- Consolidation – 90 ESX hosts, 500 VMs
- Redesign Everything – Change the way we do business
- Change for people and process
- Virtualization has caused changes to security processes



# Origin: Hill AFB and Project Bonfire

---

## Igniting Change

*“Through Project BonFire, Hill Air Force Base updates data servers and storage systems to assure that 170 apps for jet and missile maintenance are available 24 x 7.*

*Design, reliability and screaming speed. When you think Air Force, these three things typically would be synonymous with jets. But they also drove the thinking of the Hill Air Force Base systems staff when they set out to upgrade the Air Force Materiel Command’s data center operations.*

*The command at the northern Utah base repairs and maintains F-16 and A-10 jet aircraft and intercontinental ballistic missiles. The mission: Keep these aircraft and missiles ever-ready for war. To do their jobs, military personnel rely on over 170 applications. Although the apps worked fine, in recent years, server sluggishness and downtime had become a problem, says Mike Jolley, chief of the Operational Policy Branch and program manager for the command’s computer center. “*

## Source: FedTech » Magazine » “Igniting Change”

[http://fedtechmagazine.com/article.asp?item\\_id=278](http://fedtechmagazine.com/article.asp?item_id=278)

## Articles about Project BonFire:

### Government Computing News » “Order out of chaos”

[http://www.gcn.com/print/26\\_16/44607-1.html](http://www.gcn.com/print/26_16/44607-1.html)

### Washington Technology » “Something to celebrate”

[http://www.washingtontechnology.com/print/22\\_14/31170-1.html](http://www.washingtontechnology.com/print/22_14/31170-1.html)

### Gartner » "Linux Case Study: The U.S. Air Force Goes From Big Unix 'Iron' to x86 Linux in 290 Days", By George J. Weiss

Gartner ID Number G00156338, Publication Date: 26 June 2008

### Business Wire » “Hill Air Force Base Named First Red Hat Innovator of the Year”

[http://findarticles.com/p/articles/mi\\_m0EIN/is\\_2007\\_June\\_12/ai\\_n27269703](http://findarticles.com/p/articles/mi_m0EIN/is_2007_June_12/ai_n27269703)



# Scope

---

- This presentation is specifically targeted to VMware ESX.
- The processes and procedures presented are specific to VMware ESX. However, the general concepts are applicable to any Virtualization solution.
- You can download this presentation from:

**[www.SystemsImplementers.com](http://www.SystemsImplementers.com)**



# VMware ESX Service Console

---

## Service Console (aka COS) is a VM with local hypervisor information

The Service Console adds a unique value that external API scripting cannot meet. The service console is a virtual Linux host running on each hypervisor. It provides **direct** internal access to the storage and files shared on the cluster, and also has multiple different forms of scripting functionality, including the bourne shell and perl shell.

## Critical to fill-in-the-gap functionality

The Service Console is critical to providing additional functionality to meet these security requirements, filling in the gaps where the software is incomplete and the APIs are not meeting your needs.

ESX 3.5 includes a Service Console. ESXi 3.5 does not.

## Service Console Future?

VMware intends to remove the Service Console in future releases of ESX (it will still be available in 4.0). The replacement is the *Vmware Management Appliance*, which in its 1.0 release does not have the direct vmfs access that the Service Console provides. Furthermore the VMA uses the standard API's, where COS has direct hooks in the VM Kernel. VMA will be running Red Hat 5.2, which has EAL4+ certification.

## VMware has been briefed on the situation

VMware has been briefed on these issues, and is considering alterations to their products to support these functionalities in the future, including API and Hypervisor changes.



# Virtualization and Security

---

- Hardening is a start on security, not the end.
- Incident Handling in a virtualized world changes the operations paradigm.
  - **Chain of Custody** – you have the benefit of being able to freeze original instances and forensically analyze copies, while protecting the Chain of Custody
- Four Categories of procedural risks:
  - **Tracking** – the methods for tracking users and processes
  - **Containment** – the methods for containing users and applications
  - **Seizure** – the methods for seizing equipment
  - **Sanitization** – the methods for verifiably deleting sensitive data from hardware



# Tracking

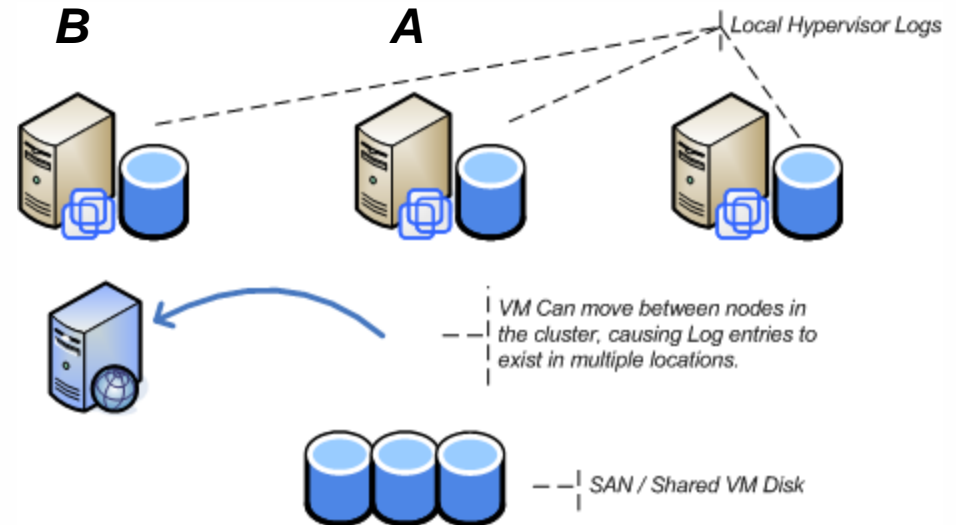
*The audit methods for tracking a user and their processing in a virtualized environment.*

## Goals:

- The ability for IA to locate a user based on IP and MAC, and which VMs this user may have logged into, and where these VMs were at the time of login (in the case of VMotion or VM migration).
- The ability to track a VMs movement in a VMware compute grid, as it is moved between physical locations.

## The Problem:

- Lag time exists in any environment between when an incident happens and when it is detected.
- Virtual Machines can move around from hardware to hardware, making it difficult to track all aspects of the VM lifecycle.



# Tracking: Log Centralization

---

## Multiple log locations:

- ESX & ESXi Hypervisor logs – *Syslog capable*
- VM Guest logs – *local disk only*
- VCenter Application Logs – *local disk only*
- VCenter Events and Task logs – *internal DB – third party or custom tool for centralization*
- Automation Logs – *varies by third party implementation*

## Policy for tracking should include:

- Virtual MAC address tracking in the network.
- Logfile centralization with syslog, where supported.
- Collection of non-centralized logfiles.



# Tracking: Applied Knowledge

Every ESX Hypervisor can be configured to use a central syslog facility for all logs—this only addresses the Hypervisor Logs. To implement this:

**Setup a syslog server.** If you are using a Unix host with rsyslog, this is as easy as:

- Enable TCP and UDP modules in rsyslog
- Add the lines at the end:  

```
$template DailyPerHostLogs, "/data/syslog/%HOSTNAME%/%%$YEAR%/%%$MONTH%/%%$DAY%.%programname%.log"  
*. * -?DailyPerHostLogs
```
- This will store logs for your system to /data/sylog. If you have SELinux enabled, you will also need to verify SELinux can store to this path.

**Hypervisor Logs:** Add a remote syslog line to /etc/syslog.conf on each service console pointing to your syslog server:

```
# Remote Log Server  
*.emerg;*.alert;*.crit;*.err;*.warning;*.notice{tab}@{remote-syslog-ip}
```

- *The {tab} character is important, and replace {remote-syslog-ip} with your syslog server ip address. Also, in some ESX updates, the syslog.conf file is replaced, removing your remote log server entry, so after each update, verify this entry is still there. Rsyslog and Syslog-NG both support TLS and can act as centralized logservers (the line above does not use TLS).*

**Guest VM Logs** must be archived centrally – as of ESX 3.5 Guest VM Logs are not handled by Syslog. A daily consolidation of Guest VM logs can be performed with scripting in the service consoles.

**vCenter Logs** must be archived centrally – as of ESX 3.5 logs are not run through a log utility for Virtual Center logs. Because of this, a daily process should be implemented to consolidate these logs.

**Third Party automation Logs** vary by vendor and implementation. Existing methods of tracking hosts by MAC address can still be used on the Virtual MAC address.



# Tracking: Applied Knowledge

---

## *Issues*

- ESX 3.5 does not support syslog with TLS.
- Guest VM Logs do not use syslog.
- vCenter Logs do not use syslog.
- There is no built-in way to export Event and Task logs from vCenter, other than third-party products or by using custom scripts with one of the toolkits.

If you have concern about any of these points, contact your VMware Sales or Support Representative.



# Containment

---

*Somebody has broken into your network. How do you contain the incident for forensic analysis without adding risk to other systems?*

## Goals:

- The ability to externally suppress and contain a compromised Virtual Machine without just shutting it down.
- Quarantining and containing users or compromised virtual machines, allowing for forensic analysis without further contamination.

## Without Virtualization:

- Conventionally this is handled with physical device bridging through a security device. With VM networking, this becomes more difficult.



# Containment: Implementation

---

## NULL Routing

Current methods of NULL Routing a MAC address are sufficient to suppress a host on the network. The VM MAC address can be determined through the arp tables. Do not disable Hypervisor ports.

## Custom Virtual Bridge

A Virtual bridging security device can be created. Today this can be implemented using Virtual Switches and a bridging security VM such as pfSense, in combination with recording the VM (API level). VMware also has vShield.

## Future

In the future, COTS solutions could take advantage of VMsafe and Virtual Private VLANs as these features are introduced in future versions of ESX.



# Containment: Applied Knowledge

---

## Implementation of a Virtual Bridge with pfSense

- Create a VM from the pfSense Virtual Appliance.
- Create two VLANs on all hypervisors, one for Production and one for Containment. The Containment VLAN should only be connected to the pfSense Virtual Appliance.
- Setup the pfSense Virtual Appliance as a bridge between the Containment VLAN and the Production VLAN, prebuilding rules to fit your needs.
- When an event occurs, the compromised VM network port can be moved into the quarantine VLAN, and the pfSense can be configured to bridge the session while also providing security control and views to the compromised VM.



# Seizure

---

A seizure event can occur through municipal requests (local law enforcement). Example:

- Virtual Desktop cluster of 1000 desktops.
- One user breaks the law. The sheriff arrives and wants to seize the users desktop.
- Possible downtime for 1000 users if the hardware is seized.

It is disastrous to your enterprise to have the entire cluster seized.

# Seizure: Process

---

ESX stores the hard drive as a file and can also store a snapshot of currently running memory as a file.

1. Suspend the VM. Wait for the VM to be fully suspended (its active memory is then stored on disk).
2. Create forensically safe non-repudiation signatures of the VM by signing each file using SHA-1 (FIPS standard hashing algorithm). Run from the ESX Service Console command-line:

```
cd /vmfs/volumes/*/ {vmname}
for f in *; do shasum $f > $f.sha1; done
```

3. Move the signed instance of the VM to a secured removable storage device which meets your network requirements, and can be presented to the authorities. Along with this signed instance, present a signed affidavit to maintain the chain of custody.
4. Archive the VM on the network until notified otherwise by the authorities.

# Seizure: Discussion Points

---

- A VMware ESX suspended VM stores not only the hard drive, but it can also store the running memory. This gives even greater advantage over a system which has been turned off, as the running memory is also available for analysis.
- The SHA1 signatures are a FIPS standard method of non-repudiation to demonstrate the data has not been tampered with since the VM was suspended in its execution environment. This can then be cloned and analyzed without violating the integrity or chain-of-custody.
- Ideally this is a built-in feature of vCenter.
- Forensic analysis of a seized VM must use the same version of ESX. Otherwise the VM image will need to be replicated and altered to an alternate VM format. However, while this altering does change the meta-data of the VM, it **does not** change the block data of the VM disk files. Since the original block-data is always available in an original copy preserving the chain of custody this process would only need to be used for run-time analysis of a system.
- An ideal feature of VMware would be to boot a signed VM as a linked clone, preserving the original state of the image and signatures.
- The SHA-1 signature is an accepted standard within the government to detect tampering, meeting FIPS-180-1 requirements. If additional levels are required, SHA-256 and SHA-512 can be used instead. SHA-256 and SHA-512 programs are not currently available in ESX 3.5, but can be added to the service console.



# Seizure: VM Separation

---

- Separate your VMs by type on completely discrete hardware.
- Conventional security recommendations suggest separating VMs by PROD, DEV, QA, etc.
- A step further and more important is to never mix Desktop VMs on the same clusters or Storage as other VMs. Ideally, clusters and storage are isolated by functional type, the most notable of which is Servers vs Desktops.
- VM Separation is also a preventative measure to help with Sanitization.

# Sanitization (Remanence Security)

---

- Every environment has sensitive data, such as credit card numbers, personal or other sensitive information.
- In the life cycle of a system, the data stores which hold this data need to be sanitized so the data cannot get into the wrong hands.
- Historically you could scrub a hard drive when a server was decommissioned. With virtual servers, this changes.
- If you have a VM with sensitive data and this VM is removed then the disk is reallocated to a new VM, can the new VM read the old VM's data? ESX zeroes unread blocks before read—what about other Hypervisors?
- Using the vmfstools feature to zero a filesystem is not an option—it does not follow any standard Remanence Security methodologies. Instead, use the tool shred.

# Sanitization: Key Data Targets

---

## VM Memory Swap file

Each VM has a swap file which can contain running memory of the VM (including unencrypted keys and other sensitive data). This is a concern and must be sanitized, but sanitizing it is difficult at best.

## VM Disk files

All VM files stored on VMFS can be reliably and discretely sanitized. Other storage types such as NFS cannot be reliably sanitized without destroying the entire volume.

## Backups

Follow your organization's current methods for sanitizing backups. Include VCB as well as host-level backups.



# Sanitization: Applied Knowledge

## VM Memory Swap file Proactive Steps:

- If VM has the possibility of containing sensitive information, its memory reservation can be set to reserve 100% of all memory in ESX. This keeps the VM from using a swapfile for memory on disk, where sensitive information would be written to disk (including unencrypted keys and other sensitive information). The size of the vswap file in VMware is inversely proportionate to how much memory is reserved (i.e. if 0 memory is reserved, the swapfile is equal to RAM, if 60% is reserved, it is 40% RAM. If 100% is reserved, it is a zero-sized file). The drawback is this makes it difficult for ESX to manage resources.
- In the case of an Incident, before powering off VM a memory scrubber should be run on the host to clean the virtual memory swap file (in ESX). The VM swapfile in ESX is stored as a file and may contain copies of physical RAM, including unencrypted keys and other sensitive information. Powering off VM deletes file, making it impossible to determine which blocks were affected and requiring a scrub of the entire LUN.

## VM Disk Files

- VM Disk files stored on VMFS. Disk files stored on NFS likely require the entire NFS mount point to be scrubbed, because NFS caches information. VMFS is verified by VMware to be a non-journaling file system which should work well with shredding disks.
- Disk cleanup:
  - The VM is “removed from inventory” (but not deleted).
  - Run a disk scrubber on each file (including snapshots, vmx, etc).
  - ESX service console includes the command shred. This command is not *certified* to meet DoD standards, but it can be run in a manner to meet at least AFSSI 8580 requirements, not including the 10% verify. Getting approval to use this process is up to your own organization:

```
shred -z -i 24 -v FILENAME
```
- In-Array Snapshots and other similar mechanisms may complicate this situation. Avoid using them if sanitization is a concern, for this reason.
- Desktop systems such as linked-clones create added complexity. We are still working on evaluating all the files involved in a linked-clone, and how to sanitize them.

## Backup cleanup:

- All backups of the VM (VCB and OS) should be identified and shredded, following existing SOP for cleaning backups.



# Sanitization: Discussion Points

---

- Ideally this is a built-in feature of the Hypervisor, which provides a mechanism to sanitize all files stored on disk for a VM:
  - Virtual disk files (vmdk), as well as Raw mapped LUNs
  - Delta files, before the delta file is deleted (perhaps mark a VM as sensitive, then always scrub delta files for a sensitive VM before they are deleted—this can happen out of band).
  - Memory and vmsn files; encrypted keys or sensitive information will be stored in memory as well, and if this memory is stored on disk, that disk needs to be sanitized.
  - Any other files relevant to the VM contents.
- In General ESX should also have an option to encrypt the memory files; keys can be stored in memory unencrypted, if this file is not protected it gives a vector for attack.
- A method needs to be researched and proven to sanitize the vswap file, even if it is run from ESX rather than the VM.
- These features have been requested to be provided by VMware. If this is something you would like to see, let your VMware representative know.



# Q&A

---

- **Brandon Gillespie**

Virtualization and Storage Architect  
Contractor, Systems Implementers Inc  
w. <http://www.systemsimplementers.com/>  
e. [brandon.gillespie@hill.af.mil](mailto:brandon.gillespie@hill.af.mil)

- **Mike Neri**

Deputy Director and CTO for the 75 Communications Group (Provisional)  
Ogden Air Logistics Center, Hill Air Force Base, UT



This presentation is available for download at  
[www.SystemsImplementers.com](http://www.SystemsImplementers.com)