ERNW
Living Security.

Virtualization Insecurity

Enno Rey, erey@ernw.de

Notice

- **Everything you are about to see, hear, read and experience is for educational purposes only. No warranties or guarantees implied or otherwise are in effect. Use of these tools, techniques and technologies are at your own risk.**

Agenda

- **A View into History & A Working Hypothesis**
- **Some Technical Terms and Concepts**
- **Objectives of “Secure Virtualization“**
- **Threats & Vulnerabilities**
- **Classification of Mitigating Controls**
- **Digging into *VMware ESX***

- **Appendix A: Hardening Guidelines**
- **Appendix B: Sample Policy**

Some introductory remarks

- **This presentation is based on an internal research project we started recently (in fact it‘s the first time this stuff is discussed in public)**
- **Thus parts of the presentation are... somewhat preliminary**
- **If interested in progress, pls follow event announcements on www.ernw.de.**

Goals of the research project

- **Step 1: Definition of a list of criteria for a “secure“ virtualization solution/environment/platform, ideally with some metrics (to make trustworthiness measurable)**
- **Step 2: Evaluation of existing platforms against these criteria**
- **Step 3: Compilation of platform specific *Hardening Guides* to close gap between *Step 1* criteria and *Step 2* results.**

A View into History – Remember *Argus Pitbull*?

ERNW Living Security.

- **Addon implementing *multilevel security* for common *NIX**
- **Once widely used, in particular in banks**
- **Software had/has ITSEC B1, CC EAL4+ certifications**
- **In 2001 NetBSD kernel bug affected Solaris as well**
- **“High secure“ system set up by vendor in the course of a hacking contest could be compromised publicly (within 24h)**
- **Lesson learned: ...**

Working Hypothesis

- **The security of virtualized environments depends highly on the "quality of the underlying common infrastructure".**
- **History shows that technologies that were introduced for some reason whatsoever are susceptible to "failure" when used for security purposes.**

- **Solutions that do not have security built in on an architectural/ design level should not be used in environments with high security requirements.**
- **Thorough risk analysis *must* be performed.**

A look into corporate reality – Question is: What do you want to achieve?

“Same security as on/with physical platforms“

- **What “security“ are you talking about? ;-) [just kidding]**
- **Next one: hmm... why don‘t you stay with the *physicals* then? [ok, just kidding again]**
- **Still: c‘mon... business won‘t accept this stance.**

This is what business is going to tell you...

ERNW Living Security.

- **“It must be easy to move this stuff to another data center, it‘s virtualized.“**

- **“What network problems? The twenty-something servers of that brand new PPS are on one physical platform, aren‘t they?“**

- **“Which costs (for security controls) are you talking about? The sales rep told me this thing was secure-by-design?“**

- **“Which policy restrictions? It‘s new, so what policy?“**

Virtual vs. Physical Security

- **Mendel Rosenblum in 09/2007:**
  "We're trying to [...] bring ourselves up to the level of security where physical machines are“. (see [8])

- **“Publicly, VMware insists that its core technology is yet to suffer from a serious breach.“ [8]**

- **Quoted article is from Sep 17.**
  On Sep 18 *VMware* issued a laarge advisory...

The Theoretical Tour

- **Asset**
- **Requirements / Goals**
- **Threats & Vulnerabilities**
- **Risks**
- **Mitigating Controls**

Quick Definition of Asset

"*Virtualization* is the creation of substitutes for real resources, that is substitutes that have the same functions and external interfaces as their counterparts, but that differ in attributes, such as size, performance, and cost. These substitutes are called *virtual resources*, and their users are typically unaware of the substitution." [2]

Virtualization, Hypervisor/VMM

- **Type I: running directly on the hardware.**
- **Type II: run on a host operating system.**
- **Hybrid**

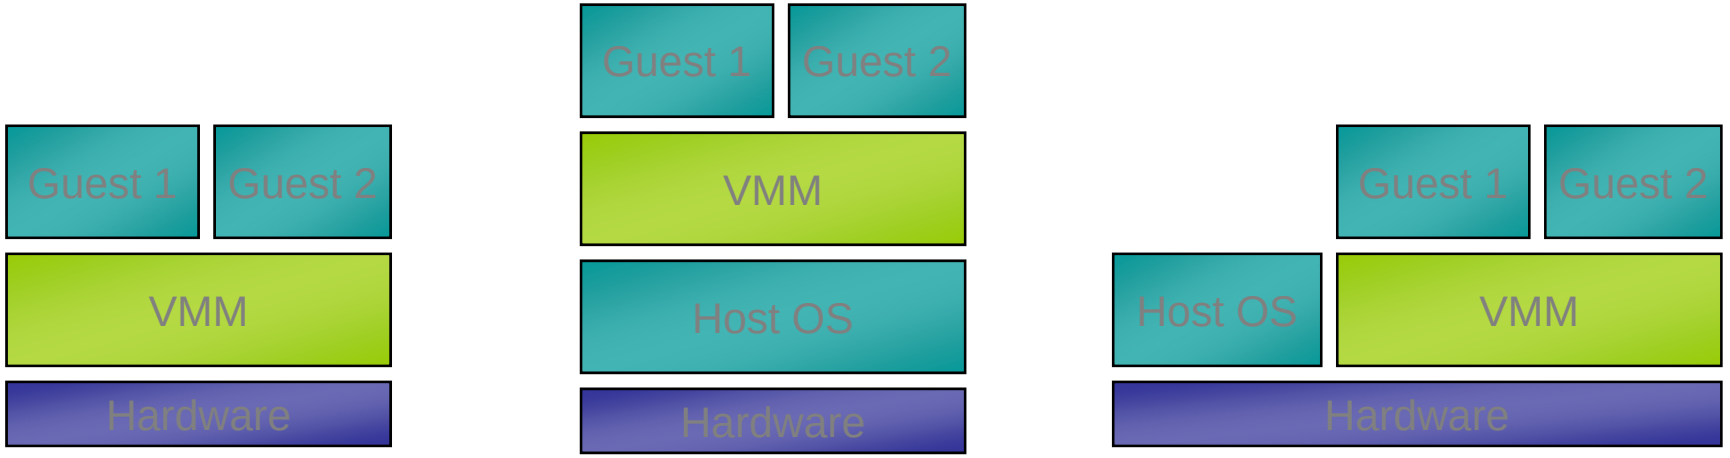

Virtualization, Overview (from [2])

ERNW Living Security.

Table 1. IBM systems and supported hypervisors

Hypervisor	Hypervisor type	IBM system that runs hypervisor
Microsoft® Virtual Server	Type 2	System x™ and BladeCenter®
POWER5™	Type 1	System i™ and System p™
Processor Resource /System Manager	Type 1	System z™
VMware ESX Server	Type 1	System x and BladeCenter
VMware GSX Server	Type 2	System x and BladeCenter
VMware Server	Type 2	System x and BladeCenter

Hypervisor	Hypervisor type	IBM system that runs hypervisor
Xen	Type 1	System x and BladeCenter
z/VM®	Type 1	System z

Security Objectives (1)

ERNW Living Security.

Sailer et.al. [Sailer2005] define the following security goals:

- **(SG1) strong isolation guarantees between multiple partitions**
- **(SG2) controlled sharing (communication and co-operation) among partitions**
- **(SG3) platform and partition integrity guarantees**
- **(SG4) platform and partition content attestation**
- **(SG5) resource accounting and control**
- **(SG6) secure services (e.g. auditing)**

Security Objectives (2)

We add:

- **( ) Flexibility and scalability**
- **( ) Compliance (this is a security presentation in 2007... you certainly had a fair expectation this buzzword was going to be dropped at least once, didn‘t you? ;-)**
- **But wait: what does “compliance“ mean here?**

Criteria for “Secure Virtualization“

ERNW Living Security.

- **Hypervisor storable in firmware or with TPM attestation**
- **Tamper-proof hypervisor**
- **Hypervisor must be able to use processor capabilities**

- **Mechanisms to assign resources (in particular CPU/memory) to guest-OS/domains must be present/configurable**
- **Mechanisms to segment I/O devices (network connectivity, storage etc.) must be present/configurable**

- **No sharing of memory pages whatsoever**
- **It must be possible to disable Copy+Paste operations between host-OS and guest-OS or between guest-OSs**
- **More to come as research project progresses…**

Threats ....

- **Attacks against mgmt_infrastructure**
- **Attacks against host-OS**
- **Attacks against guest-OS**
- **Human/configuration errors (it‘s new, it‘s business-driven)**
- **Violation of legal requirements or internal arch guidelines [compliance ;-)]**

Threats...

- **Attacks against mgmt_infrastructure**
  - from host-OS
  - via mgmt connections (SSH, Web-GUI)
  - sniffing
  - buffer overflows et.al.

- **Attacks against host-OS**
  - from guest-OS (besides simple DoS)
  - “classic stuff“

- **Attacks against guest-OSs**
  - from host-OS
  - from guest-OS
  - Direct Memory Access (DMA) attacks

Guest to Guest Attack

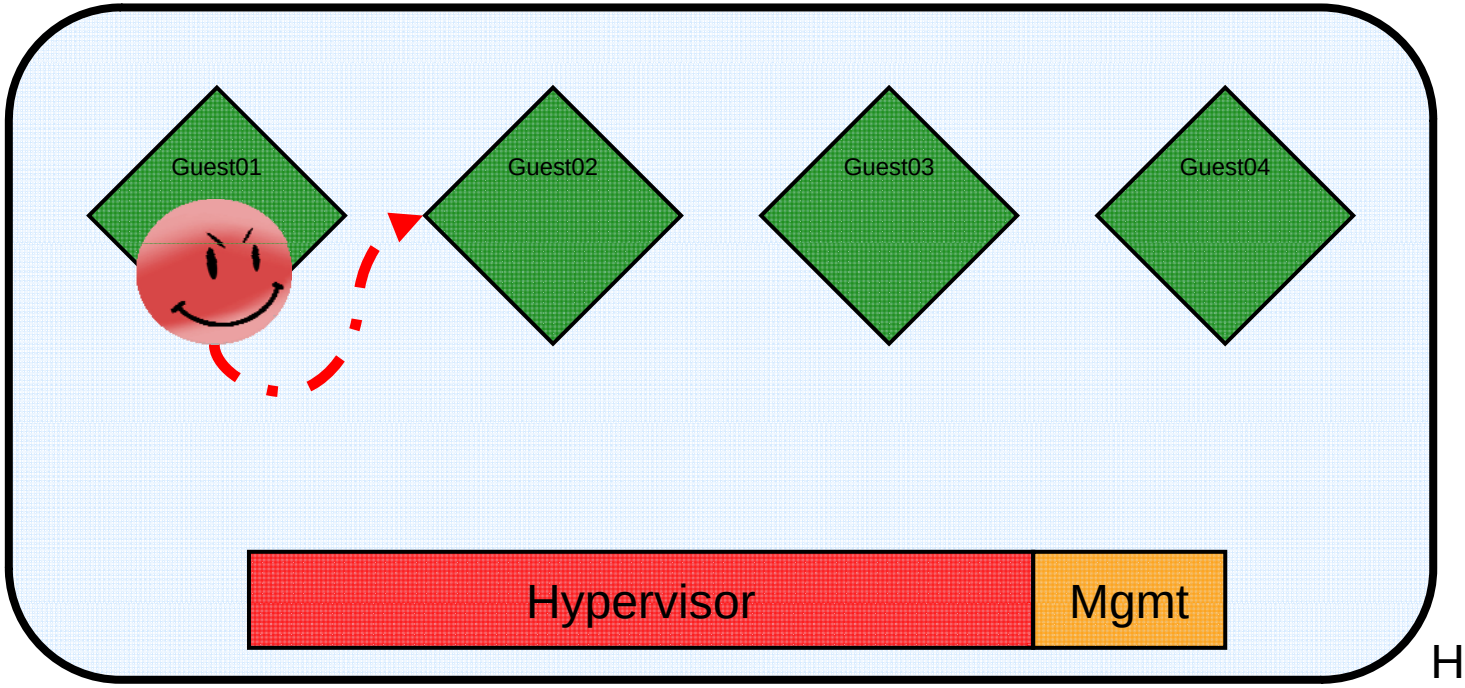

Guest to Hypervisor

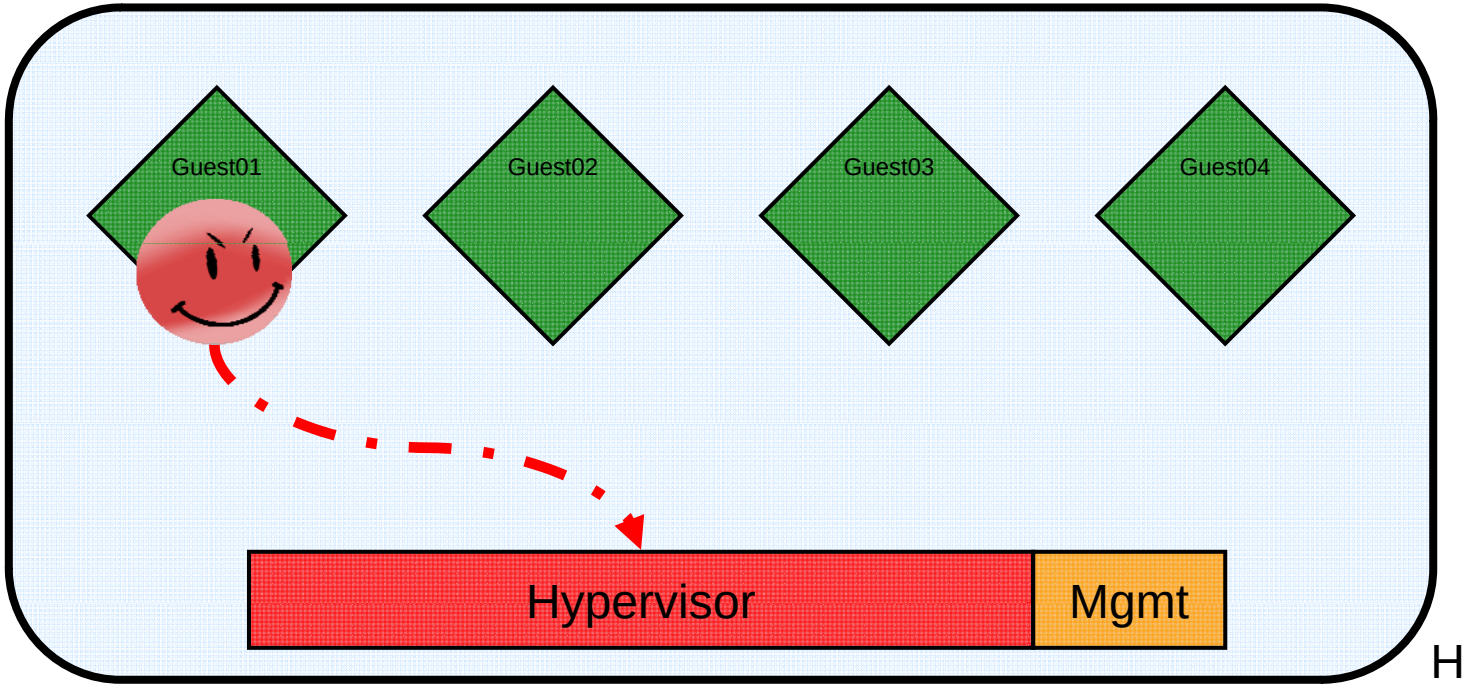

Guest to Mgmt

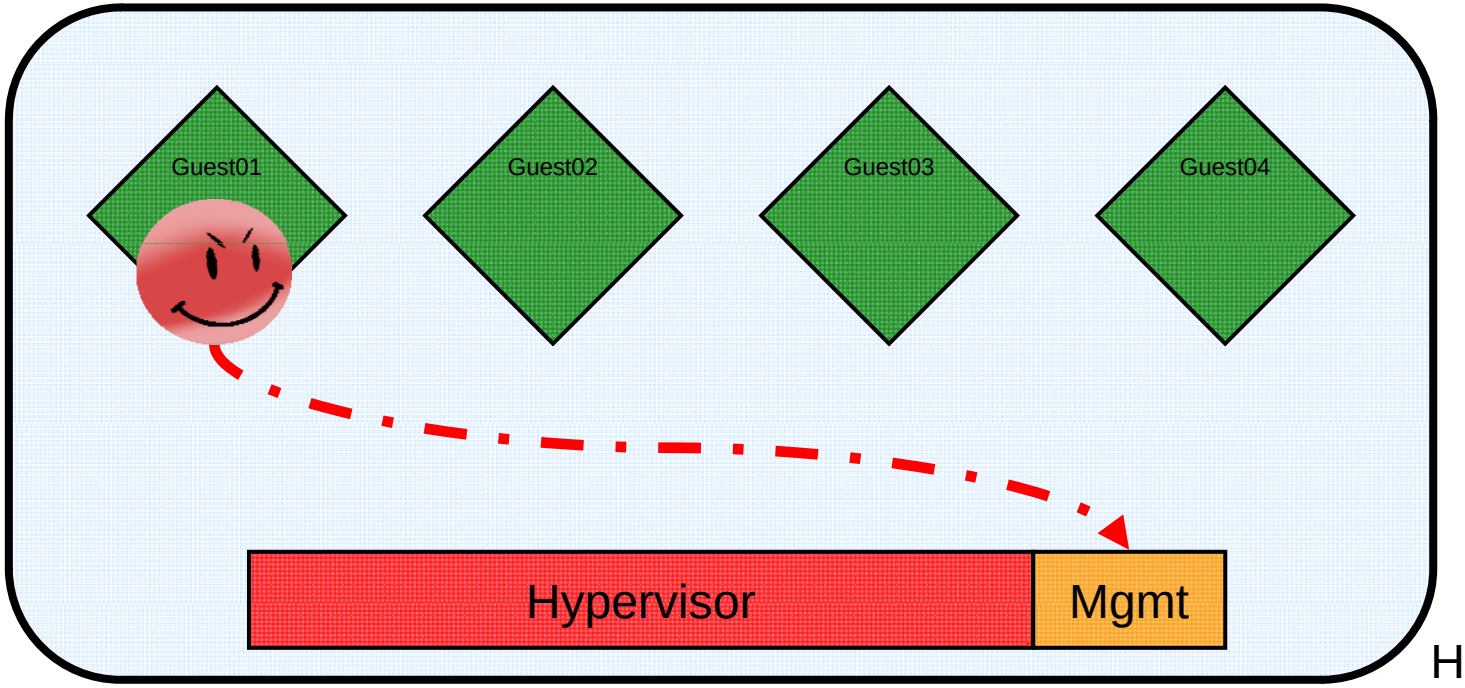

Attacking the host from the guest

VMware vulnerability in NAT networking Dec 21 2005 07:47AM
vmware-security-alert vmware com

-----BEGIN PGP SIGNED MESSAGE-----
Hash: SHA1

VULNERABILITY SUMMARY
A vulnerability has been discovered in vmnat.exe on Windows hosts and vmnet-natd on Linux systems.
The vulnerability makes it possible for a malicious guest using a NAT networking configuration to execute unwanted code on the host machine.

AFFECTED SYSTEMS:
VMware Workstation, VMware GSX Server, VMware ACE, and VMware Player.

RESOLUTION:
VMware believes that the vulnerability is very serious, and recommends that affected users update their products to the new releases or change the configuration of the virtual machine so it does not use NAT networking.

The new releases are now available for download at www.vmware.com/download

If you choose not to update your product but want to ensure that the NAT service is not available, you can disable it completely on VMware Workstation or VMware GSX Server by following the instructions in the Knowledge Base article (Answer ID 2002) at
http://www.vmware.com/support/kb.

VMware thanks Tim Shelton of ACS Security Assessment Engineering, Affiliated Computer Services, Inc., for reporting this vulnerability.

Attacking the host from the guest

- **From that mentioned Sep 2007 advisory:**

- **“This release fixes a security vulnerability that could allow a guest operating system user with administrative privileges to cause memory corruption in a host process, and thus potentially execute arbitrary code on the host. (CVE-2007-4496)“**

- **Hmm... do you mind me calling this a “serious breach“ ??**

& (some) Vulnerabilities

ERNW Living Security.

- **Strong security not enforceable due to processor or VMM architecture**
- **Hypervisor/VMM not built with security-in-mind**
- **Insecure default configurations (?)**

- **Compartmentalization capabilities not available in $VMM**

- **Partitions with different security requirements running on "same sec-level"**

Mitigating Controls Classification

ERNW Living Security.

- **Security from within guest-OS/domain**
- **Security inside the hypervisor**
- **Virtual Appliances**
- **Security as a plug-in to the hypervisor**

Choice of Processor Platform

- **Robin/Irvine in [Robin2000]:**

  “We conclude that current VMMs for the Intel architecture should not be used to enforce critical security policies. Furthermore, it would be unwise to try to implement a high assurance virtual machine monitor as a Type II VMM hosted on a generic commercial operating system.“

- **Note: in the interim there‘s stuff like Intel VT which should change things.**

Security from within guest-OS/domain

ERNW Living Security.

- **“If I‘m a virtual box I do not allow certain operations/ instructions, in order to protect my host.“**
- **Nothing existent so far, but could be done theoretically**
- **Trust problems arise quickly though (guest-OS usually untrusted from host-OS perspective)**
- **Compromised guest-OS may subvert protection stuff**

Security inside the hypervisor

ERNW Living Security.

- **Interesting approach: build some security functionality directly into the hypervisor.**
- **Hypervisor could even perform common “shared“ security services to avoid dozens of guest-OS doing the same stuff (local pf, AV etc.)**
- **Most probably we will see this in the near future (think of VMwares acquisition of *Determina*)**

Virtual Appliances

ERNW Living Security.

- **“Security Appliances“ are virtualized and can thus be “inserted“ into virtual environments**
- **Think of any physical box being part of a network**
- **Should be doable with most sw-based security solutions**
- **Some vendors (cleverly) market their stuff as “special virtual security products“**

- **See *VMware Virtual Appliance Marketplace* [http://www.vmware.com/vmtn/appliances/]**
- **Most prominent example: *Reflex VSA***

Virtual Appliances

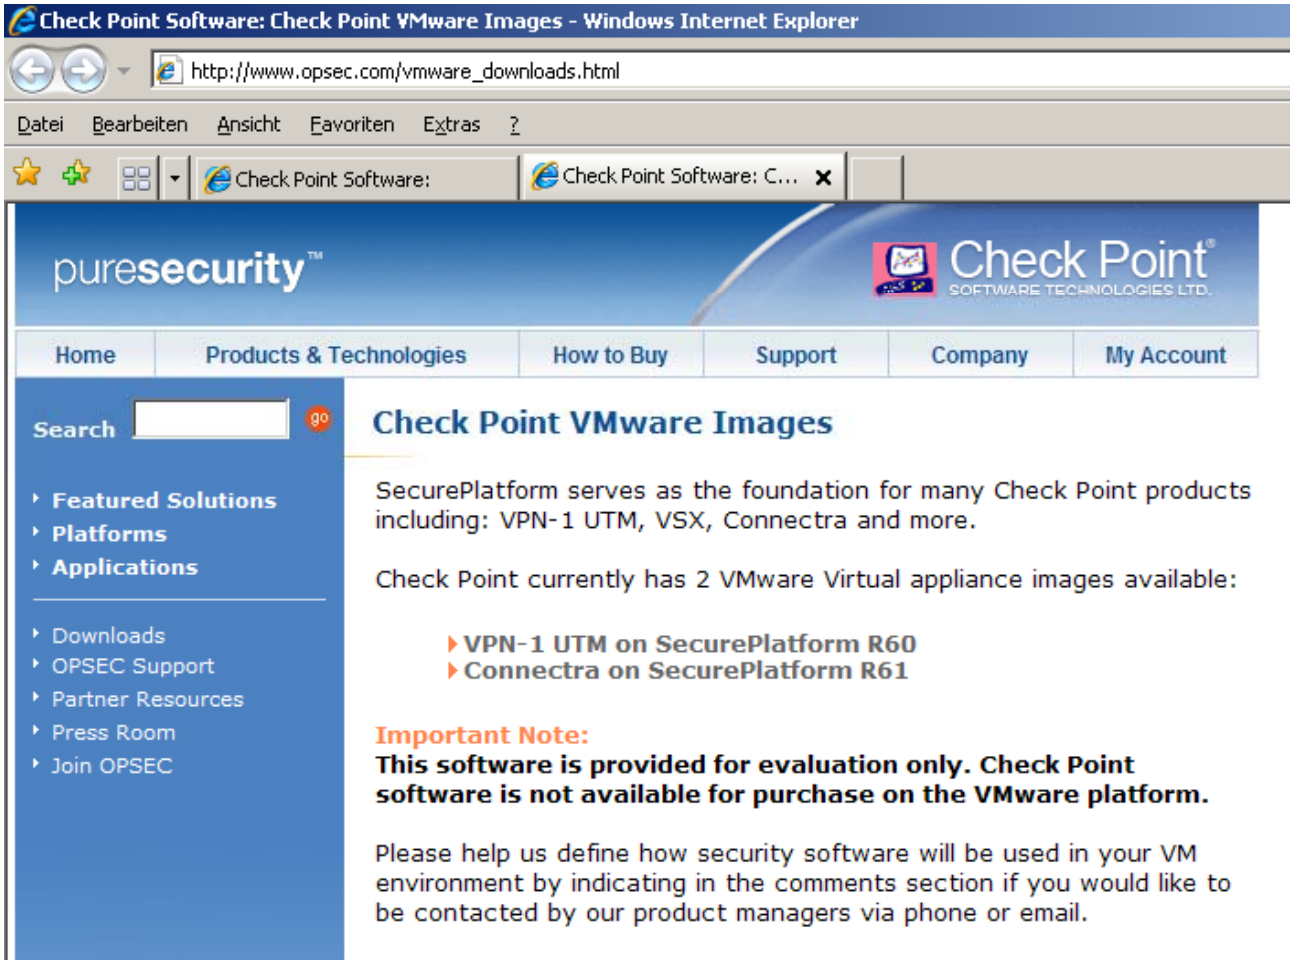

Security as a plug-in to the hypervisor

ERNW Living Security.

- ***Virtual Shield* approach**
- ***Secure Hypervisor* approach**

Virtual Shield approach

- **Extra layer added between guest-OS/domain and hypervisor**
- **This layer performs security functions as packet filtering (if needed), monitoring open ports, basic IDS/IPS etc.**
- **Most prominent example: *Blue Lane* VirtualShield. *Catbird's V-Agent* seems to work this way as well.**

sHype

- ***secure Hypervisor***
- **Research project at *IBM Research***
- **2600 lines of code**
- **Adds *Mandatory Access Control* (MAC) functionality to XEN hypervisor**
- **Designed to achieve medium assurance (CC EAL4)**
- **For details see [Sailer2005] or [4]**

Mitigating Controls
“Toolbox“

- **tbd, interim see hardening stuff**

- **CIS *Virtual Machine Security Guidelines* might be a good start. [10]**

The Practical Approach

- **Let‘s have a closer look at one widely-known product...**

- **Or, in short: break it! ;-)**

VMware Security in 2007 so far...

- April: Presentation “An Empirical Study into the Security Exposure to Hosts of Hostile Virtualized Environments“ by Tavis Ormandy at CanSecWest 2007
- July: Demonstration of some (probably) fancy tools at SANSFIRE 2007
- August: VMware acquires HIPS provider *Determina* [7]
- September: (Large) VMware Security Advis. VMSA-2007-0006
- September: VMwares announces API sharing program *Vsafe*

Tavis Ormandy Preso & Paper

- **Very interesting stuff**
- **Examined paths for attacks guest -> host, mainly by hardware fuzzing**
  - Well-known *crashme* tool
  - Self-developed *iofuzz* tool
- **Subject to analysis**
  - QEMU
  - Bochs
  - VMware workstation & server (_not_ ESX)
  - Others
- **Was able to fully compromise host in most cases**

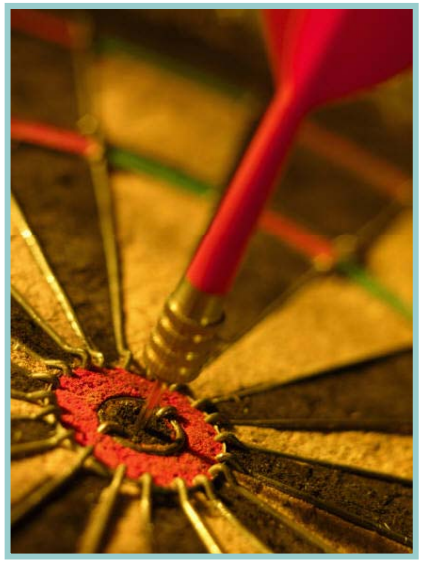

Tool Release at SANSFIRE 2007

ERNW Living Security.

- **Ed Skoudis:**

- **"The tools we presented have names based on their functionality: VMchat, VMcat, VM Drag-N-Hack (which undermines drag-and-drop, altering a file going from guest to host), VM Drag-N-Sploit (which alters a dragged file into something that shovels a shell into the guest), and, finally VMftp. That last one (VMftp) exploited the directory traversal flaw to provide FTP-style file access to the host from the guest, representing a true escape. I do not think that VMftp is an overhyped name."**

- **See http://www.cutawaysecurity.com/blog/archives/170**
- **Preso or tools not publicly available. No answer from guy.**

What we did

- **Had a “audit look“ at VMware ESX**

- **Followed the most promising attack paths ;-)**

  Guest -> Host

  Attack mgmt infrastructure

Audit Look (2.54): SSH

- **IP address for connections not specified**
- ***Root login* enabled**
- ***X11Forwarding* enabled**

- **By default all users in MUI can login via SSH [=> local exploits may be of interest, too]**

Audit Look: Users and Groups

- **Some users from /etc/passwd (albeit most without login shell)**
  - lp
  - news
  - uucp
  - Games
  - sync
  - gopher
- **/etc/group includes**
  - lp
  - news
  - uucp
  - man
  - games
  - gopher
  - floppy

Apache SSL config (2.54)

```
[root@esx /]# grep SSLCipherSuite
/usr/lib/vmware-mui/apache/conf/httpd.conf

SSLCipherSuite
ALL:!ADH:!EXP56:RC4+RSA:+HIGH:+MEDIUM:+LOW:+SS
Lv2:+EXP:+eNULL
```

- They improved this in v3

My sales rep told me *ESX server* was *Common Criteria* certified?

- **Uh, yes... indeed, it is...**
  EAL 2

- **To give you an idea... this means (from CC 3.1 part 3):**
  “EAL2 is therefore applicable in those circumstances where developers or users require a low to moderate level of independently assured security in the absence of ready availability of the complete development record. Such a situation may arise when securing legacy systems, or where access to the developer may be limited.“

  “The analysis is supported by independent testing of the TSF, evidence of developer testing based on the functional specification, selective independent confirmation of the developer test results, and a vulnerability analysis [...] demonstrating resistance to penetration attackers with a **basic attack potential**.“

- **Personal email I received from a guy leading one of the most important CC labs worldwide:**

  “At the moment, though, I do not think VMware can go any higher than EAL2.”

Attacks Guest -> Host

- **We took Tavis Ormandy‘s work as a starting point**
- **On ESX no interesting results with *crashme***
- ***iofuzz* not publicly available**
  - Why not ask the author for it? ;-)
  - We (in fact our student Muhammed) did...
  - And Tavis sent us an ISO-image. Kudos for that, guy! Great move.
- **Played around with it ... ;-)**

iofuzz

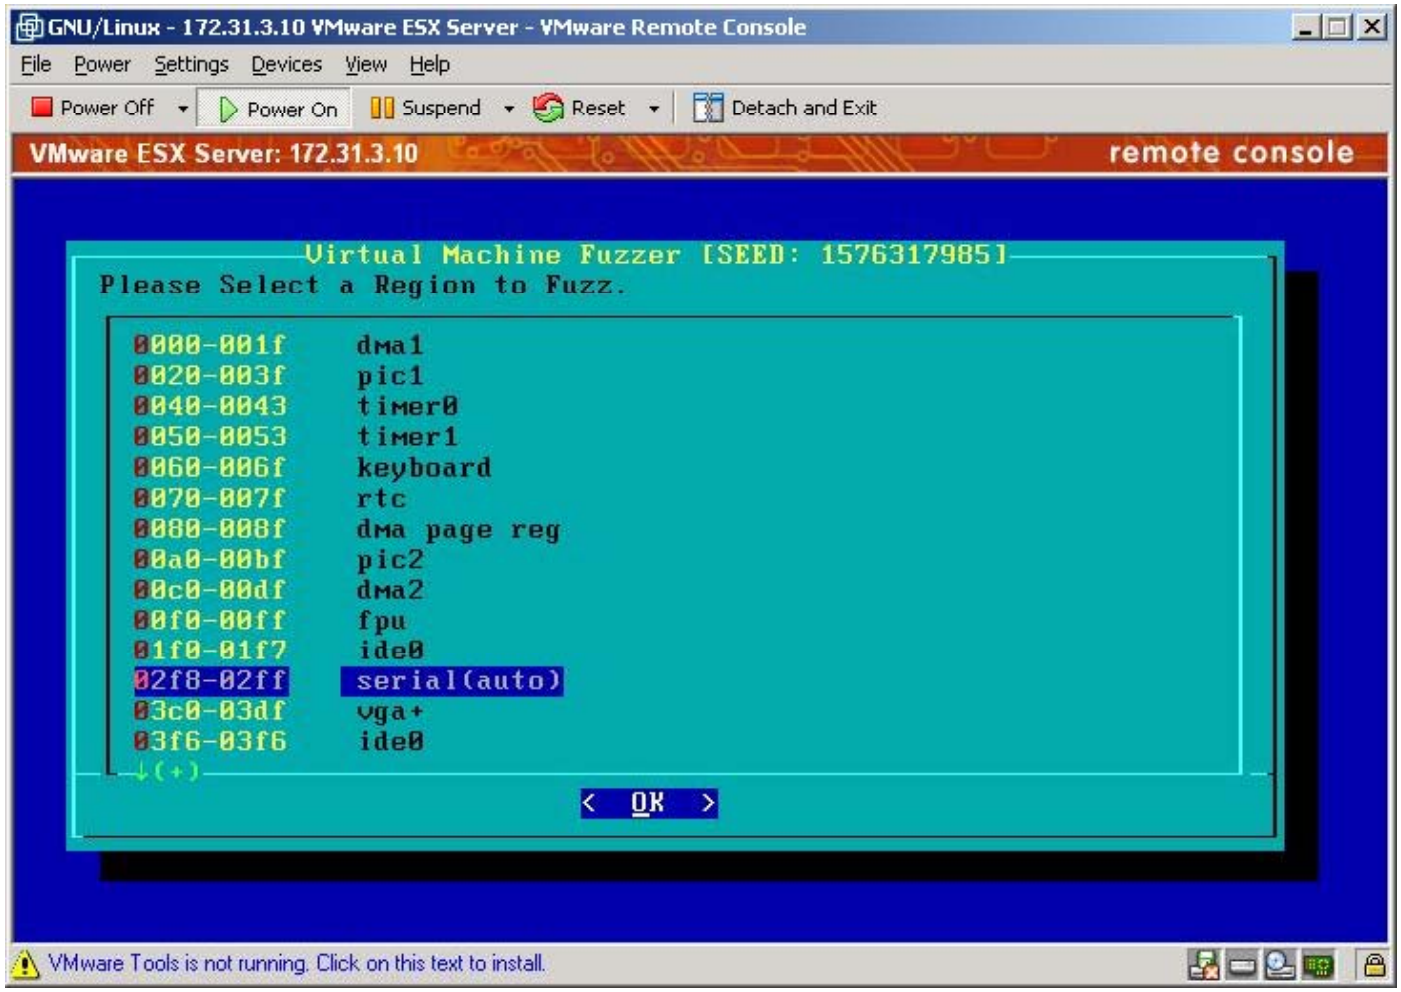
GNU/Linux - 172.31.3.10 VMware ESX Server - VMware Remote Console
File Power Settings Devices View Help
Power Off
Power On
Suspend
Reset
Detach and Exit
VMware ESX Server: 172.31.3.10
remote console
Virtual Machine Fuzzer [SEED: 1576317985]
Please Select a Region to Fuzz.
0000-001f dma1
0020-003f pic1
0040-0043 timer0
0050-0053 timer1
0060-006f keyboard
0070-007f rtc
0080-008f dma page reg
00a0-00bf pic2
00c0-00df dma2
00f0-00ff fpu
01f0-01f7 ide0
02f8-02ff serial(auto)
03c0-03df vga+
03f6-03f6 ide0
↓(+)
< OK >
VMware Tools is not running. Click on this text to install.

... and the impact

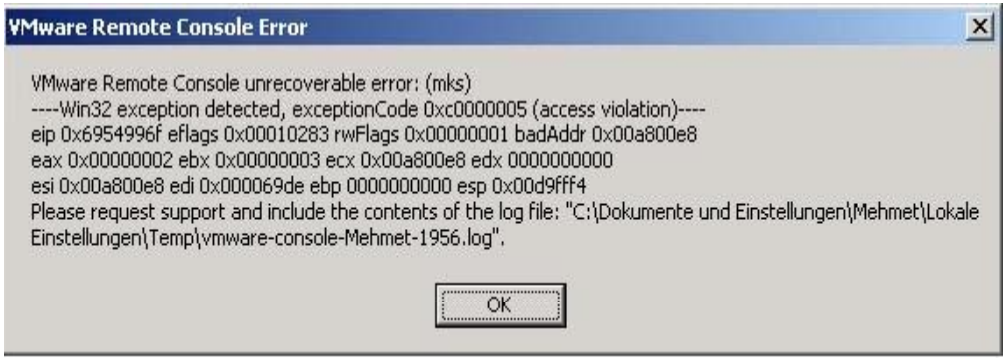

This certainly does not look good ;-))


```
Jun 29 11:05:15: vcpu-0| Backtrace[6] 0xbf7ffa94 eip 0x8084e7a
Jun 29 11:05:15: vcpu-0| Backtrace[7] 0xbf7ffab4 eip 0x807e848
Jun 29 11:05:15: vcpu-0| Backtrace[8] 0xbf7ffb24 eip 0x80e3d08
Jun 29 11:05:15: vcpu-0| Backtrace[9] 0xbf7ffbf4 eip 0x40047fb7
Jun 29 11:05:15: vcpu-0| Backtrace[10] 00000000 eip 0x4015acba
Jun 29 11:05:15: vcpu-0| Msg_Post: error
Jun 29 11:05:15: vcpu-0| [msg.log.vmxpanic] VMware ESX server unrecoverable
   error: (vcpu-0)
Jun 29 11:05:15: vcpu-0| BUG F(553):566 bugNr=431
Jun 29 11:05:15: vcpu-0| Please request support and include the contents of
   the
log file: "/root/Vmware/fuzz/Vmware.log".  We will respond on the basis of
   your support entitlement.
Jun 29 11:05:15: vcpu-0| ----------------------------------------
Jun 29 11:05:26: vcpu-0| VTHREAD thread 4 start exiting
Jun 29 11:05:26: vcpu-0| VTHREAD counting thread 0
Jun 29 11:05:26: vcpu-0| VTHREAD counting thread 1
Jun 29 11:05:26: vcpu-0| VTHREAD thread 4 exiting, 2 left
Jun 29 11:05:26: vmx| VTHREAD watched thread 4 "vcpu-0" died
Jun 29 11:05:26: vmx| VTHREAD thread 0 start exiting
```

Attacks on mgmt

- **Used available and home-grown attack tools on standard mgmt interface stuff (mainly HTTP / HTTPS based)**
- **First round on 2.54**
- **Promising results... so be decided to buy 3.01 [remember the ERNW research approach ;-)]**
- **Problems seem still to exist**
- **Currently researching on it (since two weeks another student writes his master thesis on this...)**

VMWare ESX - Management Structure

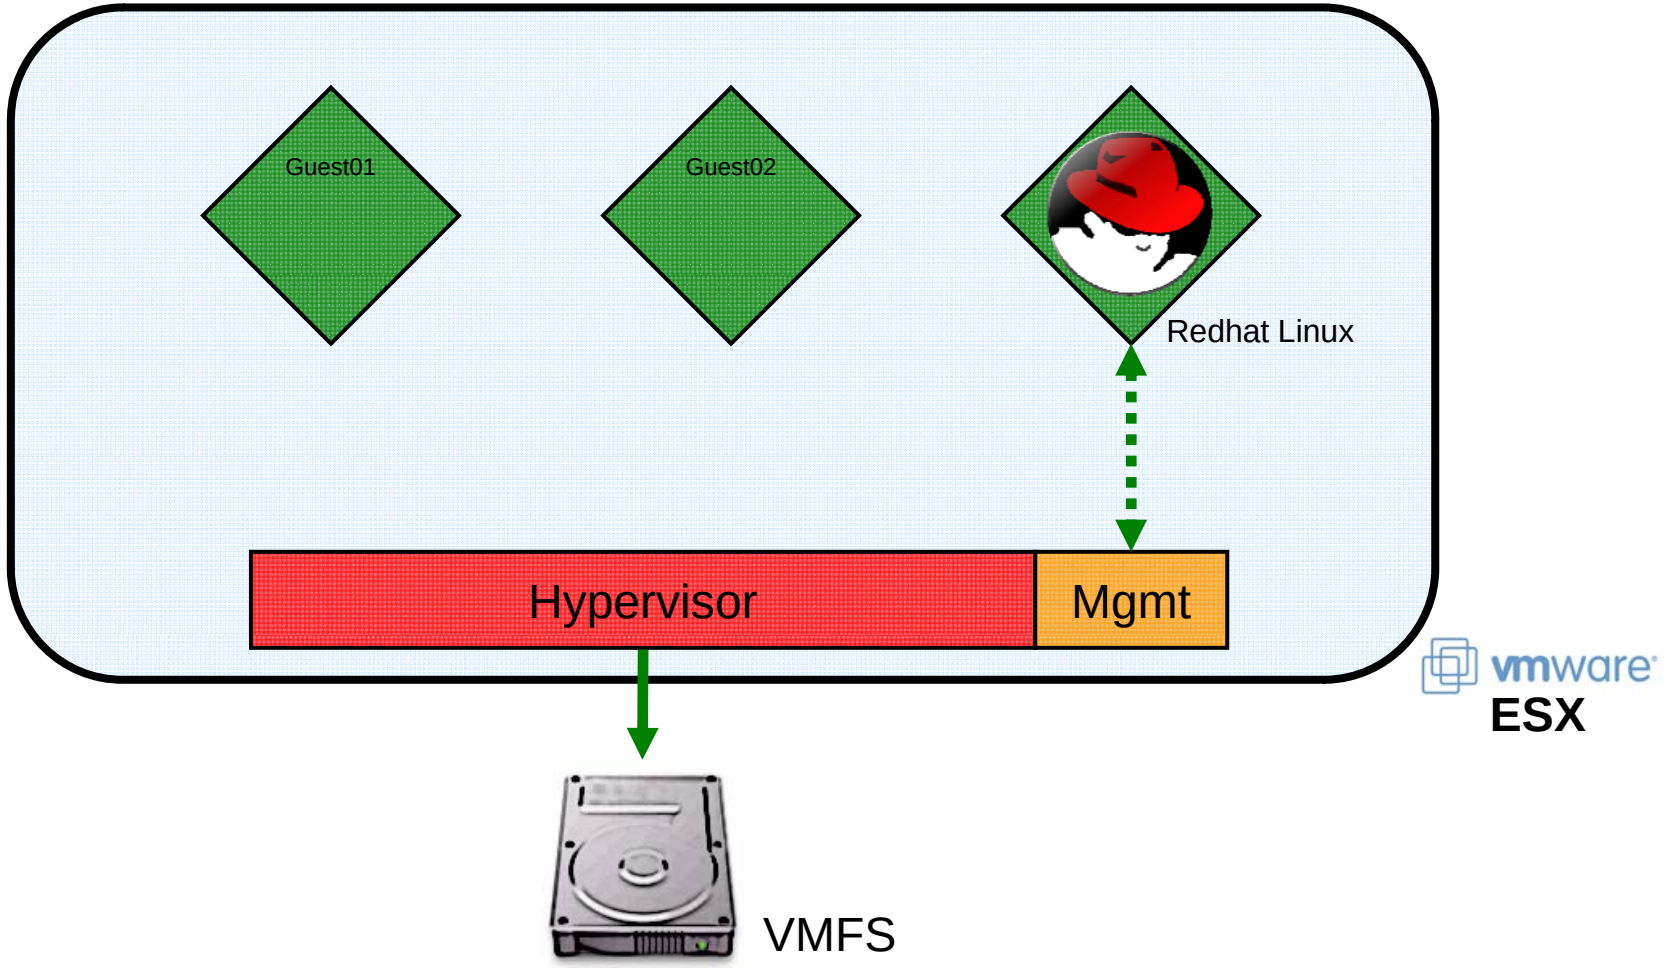

What‘s this? ;-))

```
[Mon Apr  2 18:49:26 2007] [error] Segmentation fault in
MUI VMDB
```

- **We are able to reproduce this!**
- ***VMware Console* takes no more connections then**
- **Process(es) still running**
- **Most probably exploitable**
- **Btw: it was *not* difficult to find!**
- **Currently researching it, pls don‘t expect PoC/exploit code before approx. dec07 (due to ERNW disclosure policy)**

On 3.02 it looks like this

```
[2007-10-09 03:08:03.057 'SoapAdapter' 51899312 verbose] Failed to
   write reply to connection; Broken pipe
[2007-10-09 03:08:06.179 'Vmomi' 124365744 info] Activation
   [N5Vmomi10ActivationE:0xb81b5f0] : Invoke done [waitForUpdates]
   on [vmodl.query.PropertyCollector:ha-property-collector]
[2007-10-09 03:08:06.180 'Vmomi' 124365744 info] Throw
   vmodl.fault.RequestCanceled
[2007-10-09 03:08:06.180 'Vmomi' 124365744 info] Result:
(vmodl.fault.RequestCanceled) {
   msg = ""
}
[2007-10-09 03:08:06.197 'SoapAdapter' 124365744 verbose] Failed to
   write reply to connection; Broken pipe
[2007-10-09 03:09:01.539 'Vmomi' 81202096 info] Activation
   [N5Vmomi10ActivationE:0xb48d468] : Invoke done [waitForUpdates]
   on [vmodl.query.PropertyCollector:ha-property-collector]
[2007-10-09 03:09:01.540 'Vmomi' 81202096 info] Throw
   vmodl.fault.RequestCanceled
```

Summary

ERNW Living Security.

- **Virtualization offers interesting features but will introduce new threats that are not yet fully understood.**
- **Thorough risk analysis needed in environments with high security requirements.**
- **Security currently not built-in; strive for security may add complexity and costs.**
- **When planning to implement virtualized environments you must understand architecture to select business reasonable mitigating controls.**

Questions?

ERNW Living Security.

Thanks for your attention!

ERNW
Living Security.

References

[2] IBM Whitepaper
http://publib.boulder.ibm.com/infocenter/eserver/v1r2/topic/eicay/eicay.pdf

[3] Tavis Ormandy, http://taviso.decsystem.org/virtsec.pdf

[4] Ralf Spenneberg on sHype at DFN CERT Workshop:
http://www.dfn-cert.de/events/ws/2007/dfncert-ws2007-f8.pdf

[6] http://www.apricot.net/apricot2006/slides/tutorial/monday/MPLS_Tutorial.zip

[7] http://searchsecurity.techtarget.com/originalContent/0,289142,sid14_gci1268544,00.html

[8] http://www.news.com/VMware-shares-secrets-in-security-drive/2100-1012_3-6208354.html

[Sailer2005]: IBM Research Report RC23511: *R. Sailer, E. Valdez, T. Jaeger, R. Perez, L. van Doorn, J. L. Griffin, S. Berger:* sHype: Secure Hypervisor Approach to Trusted Virtualized Systems.

[Robin2000]: Robin, John Scott & Irvine, Cynthia E.: Analysis of the Intel Pentium‘s Ability to Support a Secure Virtual Machine Monitor
http://www.usenix.org/events/sec2000/full_papers/robin/robin.pdf

References

- [9] http://www.news.com/VMware-shares-secrets-in-security-drive/2100-1012_3-6208354.htm
- [10] Center for Internet Security: Virtual Machine Security Guidelines, http://www.cisecurity.org/tools2/vm/CIS_VM_Benchmark_v1.0.pdf

Appendix A: Hardening Guidelines

- **Some unordered measures (more to come):**

- **Use TPM attestation**
- **Use processors with isolation capabilities**

- **Deactivate hyperthreading for guest-OSs.**

- **Disable Copy+Paste operations between host-OS and guest-OS or between guest-OSs**

- **Configure strict control of resources (in particular memory)**

Hardening 2

- **Segmentation of mgmt and host-OS network traffic**
- **Restriction of mgmt connections on IP_source basis**
- **SNMP ?? [VMware ESX does not support SNMPv3 out-of-the-box)**

- **Use esxcfg-firewall appropriately**

- **No root login, use sudo**
- **Implement password policy**

- **Tighten SSLCiphers if necessary**

Hardening, more stuff

- **Patching process, in particular for "offline-images"**
- **Filesystem integrity**
- **Enable NTP (KB article 1339)**
- **Extensive Logging**

- **Strong segregation of duties as for administration of host/administration of guests/auditing**

- **Reject MAC impersonation on virtual switch security profile [=> protection against ARP spoofing]**

- **Use LUN masking**

Appendix B: Sample Policy

- **On the following slides some building blocks can be found. A much more extensive template is available upon request (erey@ernw.de).**

Appendix B: Sample Policy

- **As a company, ERNW constantly challenges its employees to innovate. It is one of our core values. Technology often provides new ways to innovate and drive savings. One technology that shows great promise and is being driven by this spirit of innovation is virtualized computing. Cost savings, dynamic resource allocation and server consolidation are propelling this technology forward in many organizations. As this trend continues, we are sure to discover new and novel ways to use the technology.**
- **This paper addresses Virtual Computing Environment security requirements and standards to assist responsible personnel in engineering commercially reasonable solutions that meet the needs of the business regardless of the vendor solution. It draws upon current ERNW Security Policies, ERNW Computers and Network Security and the ERNW Network Security Architecture. Specific guidance is given with regard to required and prudent security controls.**

Appendix B: Sample Policy

- **As ERNW Business Units explore new and innovative ways to leverage Virtualization technology, we fully expect to evolve our understanding of risk and risk mitigation. This understanding will be reflected in security control affinities and minimum security requirements. It is our intent to document specific vendor implementations in Appendix A. Thereby, providing real world examples of compliant implementations specific to each vendor. Until that time, certain implementations will be prohibited pending a formal vetting to define residual risk and identify compensating controls.**
- **Implementations of VCE's that bridge/span security domains (ERNW Network Security Architecture)**
- **Implementations of VCE's that are managed by multiple operations groups**
- **Implementations of VCE's that provide access High loss value data belonging to multiple custodians.**
- **Implementations of VCE's that provide authentication services such as LDAP and AD (except for testing purposes must not contain real/live credentials).**