Virtualisation: There is no spoon

Michael Kemp
clappymonkey@gmail.com

Who am I?

- UK based security consultant, application vandal, and occasional researcher
- Make a living conducting security assessments and providing security consulting for a range of freelance clients throughout the world
- A very lazy guy that occasionally has some ideas above my station...
- I'm not here to sell anything – I just had a few ideas I wanted to run past you all

Disclaimer

- It should be noted that any ideas, views or opinions expressed in this presentation or supporting materials, are in to way indicative, reflective or representative of the views, opinions, or ideas held by my current or any previous employer. Additionally should you use any of the ideas in this talk and end up bricking anything, I hold no responsibility.

```
/end disclaimer
```

Disclaimer #2

- I am an 'app' guy – and networks (virtual or otherwise) aren't really my thing
- A lot of the topics discussed in this talk are the result of semi-drunk conversations and 'interpretation' (thanks to a lot of kit being beyond my budget)
- This is preliminary research with all that this entails – If you know better tell me, and I'll shut up (preferably after I've finished ranting)

`/end disclaimer` #2

Agenda

- The playing field
- Threats, Vulnerabilities and other Grooviness
- Memory of Future Past
- The Becoming?

Before we begin...

- Myself and a colleague noticed something weird a while ago (absolutely nothing to do with virtualisation btw)
- Standard logic goes that the maximum password / username length in MS-Win32 environments is 127 chars (password reset)
- On a client site it transpired that they had no maximum username length set in Domain Security Policy / AD
- Once authed, by performing a password reset, users could input any string length they wanted – which was passed local and to Active Directory

Before we begin...

- So, by passing a huge string it was *may* be possible to make the AD fall over
- MS said only 999 chars are allowed... I got in 8000+ and then got bored – MS didn't bother responding to that...
- Possibly a fault in MS – but definitely a faulty DSP / AD config ...
- Can't help but wonder how many others there are out there...

Before we begin...

The Playing Field

"Virtualization is both an opportunity and a threat..."

Patrick Lin
Senior Director - Product Management – VMWare.

"If we knew what it is we were doing, it would not be called research, would it?"

Albert Einstein
Genius

The Playing Field

- Many talks about virtualised technologies focus on one area
- This talk will be discussing more than one
- Virtualisation can be broken down to a number of distinct aspects:
  - Platform virtualisation (including app virtualisation, hardware enabled virtualisation etc.)
  - Resource virtualisation (including grid computing, VPNs, and clusters)

The Playing Field

- Platform virtualisation can be broken down yet further...
- However it is set up you will typically find some combination of virtual machines, guest OSs, hardware (well duh), Virtual Machine Monitor (VMM) and host OS all playing nicely (or not) together in a Virtual Machine Environment
- Platform virtualisation is used daily by researchers, legitimate companies and attackers alike, and serves a variety of purposes (nefarious or otherwise)

The Playing Field

- A lot of research has been dedicated to exploiting platform virtualisation (about more in a moment) but little thought is given to resource virtualisation techs

- The use of virtual servers and grid computing is growing (thanks HP...) and there are some real business drivers behind it (“Hey, why do we need all these boxes? – let's just put them in one! We'll save a fortune!... “)

- Small problem with that is that it may be rather more difficult to implement securely than conventionally thought

The Playing Field

- Before talking about the threats to virtual environments and platforms, it's worth considering what constitutes a 'secure' environment

- Simple really: – users shouldn't be able to break out of the 'cave', platforms and partitions should be separate and isolated and use separate resources (e.g. Memory) and it should all be auditable, scalable and flexible

- This is all very nice in theory – but in practice... well...

Threats, Vulnerabilities and other Grooviness

- Breaking stuff is fun, and a lot of people a lot cleverer than me have dedicated their time to breaking virtualized environments
- A *lot* of research has been directed at VMWare, as it's cheap to get hold of, and amusing to break should you be so inclined...
- Before going any further, it's worth having a quick recap of how things have been broken, by whom, ideas about breaking things further and what the vendor has had to say about all this...

Threats, Vulnerabilities and other Grooviness

- VMWare threats can be broken down into three distinct camps:
  - Establishing the cave
  - Breaking out of the cave
  - Smashing the cave up

Threats, Vulnerabilities and other Grooviness

- The cave analogy is used a lot in relation to Virtual environments – I'm going to steal it too...
- Before attacking anything, you first need to know where you are to attack it – hence, establishing the cave, or Virtual Machine Detection
- Lots of effort and research has gone into this...

Threats, Vulnerabilities and other Grooviness

- Attackers can find out if they are in the cave by looking for artefacts in memory e.g. Running processes, Registry keys, Files – and if they fancy getting around rootkit-lite techniques specific memory artefacts such as the Interrupt Table Descriptor (IDT) using CheckIDT or similar (please read more Phrack...)
- If feeling lazy, an attacker can also use the infamous Red Pill by Joanna Rutkowksa unleashed in 2004
- Beyond that, there's a host of other tools out there (VMDetect, Jerry, Scrappy, & Doo
- There's more than one way to skin a cat, and more than one way out of a cave...

Threats, Vulnerabilities and other Grooviness

- The cave can also be established remotely too...
- Research is currently afoot to detect virtual environments across the network using pattern identification
- According to some packet headers may have discernible patterns, and perhaps more promisingly there is evidence to suggest that there is a discernible lag in the timestamps of virtual environments (no great surprise)
- From personal experience I know that auditing on a virtualised resource environment is tricky as the timestamps are way off...
- There are no public release tools yet (that I know of...) – but undoubtedly they will be forthcoming at some point

Threats, Vulnerabilities and other Grooviness

- So, once an attacker knows they are in a cave (either by analysing stalactites or groping around in the dark) what's next?
- Finding a way out of it...
- In July last year, Ed Skoudis and the IntelGuardians guys demonstrated a range of tools geared towards breaking out of the cave and playing in the field beyond...
- Worth a look in more depth...

Threats, Vulnerabilities and other Grooviness

- To cut a long story short a number of vulnerabilities were found in VMWare Workstation 4/5 and may well exist in 6
- By exploiting the ComChannel between host and client, a number of highly entertaining escape routes from the cave were found...
- These included a suite of as yet unreleased tools, including: VMChat, VMCat, VMDrag-n-Hack, VMDrag-n-sploit and Vmftp
- I won't discuss each in detail, but VMChat is kind of interesting...

Threats, Vulnerabilities and other Grooviness

- VMChat allows for a simple chat application to be implemented between both Host and Client OS
- By injecting a DLL on the Host OS, an attacker can access the memory space of the Client OS and in doing so establish a shared memory buffer that can then be used as a communications channel
- This is largely the fault of the VMWare ComChannel itself and can be exploited whether or not VMWare Tools are installed
- By subverting x86 instructions, VMWare effectively allows shared communications between Host and Client and memory separation becomes a redundant notion....

Threats, Vulnerabilities and other Grooviness

- Before coming to this con, I was promised a standing ovation if I replicated VMChat....
- You can all remain seated!
- I did get in touch with Ed Skoudis and ask what DLL was injected and how, rather unsurprisingly I have heard nothing back
- In my desperation for applause though, I did some more research, and came across the same stuff as the guys from CoreLabs, namely the awesome VM Back from Ken Kato (http://chitchat.at.infoseek.co.jp/vmware/)

Threats, Vulnerabilities and other Grooviness

- Ken Kato has been exploring the Backdoor/IO 'functionality' of VMWare for years...
- Backdoor/IO is the same ComChannel used by VMWare for VMWare tools, and Ken has utilised it to do some interesting things...
- Be utilising Backdoor/IO functions, a number of command line tools have been generated, including a generic backdoor access program that allows users to copy and paste between Host and Client (amongst other things) and Vmftp, that allows for host and guest to exchange files through VM Shared Folders...

Threats, Vulnerabilities and other Grooviness

- Now, although I've not done this yet, it may well be possible to use VMftp to set up an interesting scenario
- It *may* be possible to set up a chat server on the Host, and a client on the Client(s)
- If that could be done (I don't know) then maybe another way of doing VMChat is at hand?
- Even if it's not, separation between Host and Client OS can be bypassed as demonstrated by CoreLabs Path Traversal in Shared Folders

Threats, Vulnerabilities and other Grooviness

- So, the cave of virtualisation can not only be discovered, but also escaped from
- As Tavis Ormandy found (and indeed the guys from ERNW can confirm), by utilsing crashme and iofuzz (which he will share by all accounts) virtualised platforms could be made to fall over with comparative ease
- Not only can a user find the cave, escape from it, but they can also smash it up too...

Threats, Vulnerabilities and other Grooviness

- As highlighted earlier in this talk a *lot* of research has been focused on platform virtualisation
- A lot of vulns have been found, and the seperation and isolation promised by virtualisation may or may not exist
- There has been little recent focused security research on resource virtualisation...

Memory of Future Past

"Looks like the ankle-biters have learned to read technical manuals"

Tsutomu Shimomura / John Markoff

Memory of Future Past

- It is a difficult proposition to discuss virtualisation without referencing computer memory design
- Typically platform virtualisation takes advantage of virtual memory models
- One of the most exiting aspects of recent years is NUMA (Non-Uniform Memory Access)
- For those of you unfamiliar with memory architecture in relation to virtualisation (I was in that position about two months ago – so I'm no guru!), here's a quick recap:
- BTW: In case you are wondering – I am going somewhere with this...

Memory of Future Past

- Processors run faster than the memory they are attached to
- To avoid CPUs having to stall and wait for memory access a number of projects have been undertaken to allow for high speed access to memory
- NUMA is just one of many (including Distributed Shared Virtual Memory, Multics, etc.) and attempts to solve this problem by providing seperate memory for each processor
- Limited real world implementations of NUMA thanks to plentiful programming difficulties

Memory of Future Past

- To get around the issues with NUMA, Cache Coherent Non Uniform Memory Access / Allocation was introduced
- As the name suggests ccNUMA seeks to maintain the integrity of data stored in local caches of a shared resource
- In essence ccNUMA seeks to maintain the integrity of data stored in the memory of multiprocessor systems
- Used in a lot of cluster computing models, and in virtualised resources and servers (e.g. HP Superdomes)

Memory of Future Past

- A typical analogy for understanding NUMA (and ccNUMA) is the process of baking a cake... I don't bake, so I have thought one up concerning drinking

- Imagine you want to throw a legendary party. To accomplish this you will need lots of lager (memory pages) to complete this task (process). You have a lot of lager in your own fridge (local memory) but your neighbour has fantastic whisky which you borrow (remote memory). You need to get drunk quickly, so storing alcohol locally is good. You can only fit so much alcohol locally though (physical nodal memory) so you may need to get your neighbour to hold onto some (if local memory full, allocate memory pages remotely)

- Basic premise is to complete a process quickly it may be necessary to allocate memory pages remotely as well as in local memory...

Memory of Future Past

- As promised – I will now discuss where I was going with all this...
- I first became aware of ccNUMA as a concept on the latter part of 2007
- A client at the time was considering replacing nice flat networks and attendant apps and DBs with a virtualised resource, namely a HP Superdome
- Before they could do this, I had to perform a risk assessment of the suggested architecture and hence ran across ccNUMA and the thorny topic of resource virtualisation...

Memory of Future Past

- Before going any further let's discuss how memory works in a HP Superdome (this is taken from one of their marketing documents which I am assuming is correct...)

There is no spoon

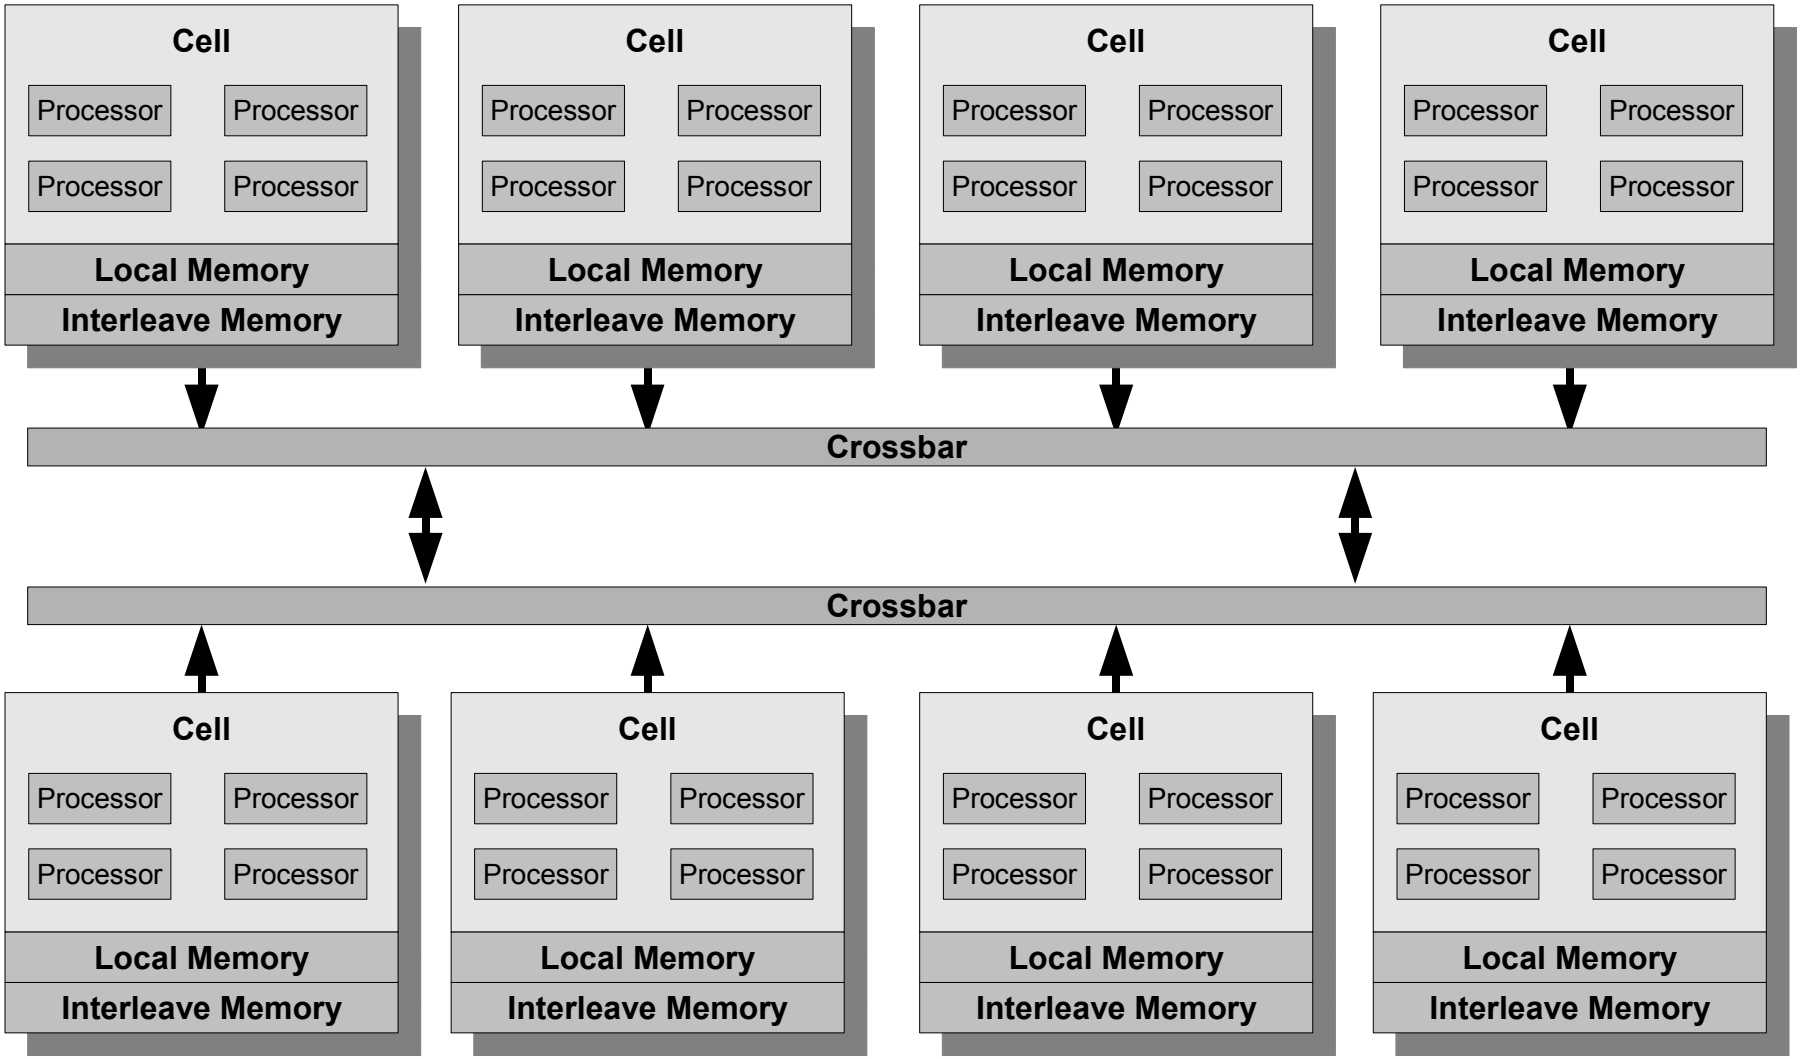

Memory of Future Past

- As discussed Superdomes can be broken down to individual cells, containing individual processors (64 in total I believe), local memory, interleave memory and relevant I/O connections
- Cells are connected via crossbars which form the interconnect
- If memory can't be written to on a local processor cell, it'll try it's neighbours with the proviso that memory transactions should not cross more than two crossbars
- At this point I could bore you senseless about Locality Domains (LDOMS) but at the basic level an LDOM is a related collection and processors and memory (apart from interleave memory – as that contains no processors), and with that in mind, I'll shut up about it now...

Memory of Future Past

- It is however, worth discussing the difference between Local Memory and Interleave Memory in a Superdome / ccNUMA model

- Local Memory is memory restricted to the storage of private objects and data structures. That said it can be accessed by any processor, but the further that processor is from the particular cell the longer access takes

- Interleaved Memory is used for the storage of shared objects and data subjects, and has uniform latency (in terms of time to access) regardless of the cell / processor accessing it. Interleaved Memory is accessed in a Round Robin fashion

Memory of Future Past

- Are you spotting the flaw in all this? I did, and that's why I'm here...
- From what I can tell, if an attacker can gain access and control over one processor they can access the memory space of all other processors not only in a cell but across the virtual resource as a whole
- Instead on injecting malicious code into the memory of one isolated machine, attackers can now inject into the memory of everything else too
- If an application falls over and becomes a processor hog (or an attacker creates this state – think malicious XML on MS-XML) then more than processor can get to come to the party...
- Basically if I can own one processor, I get to own everything!

Memory of Future Past

Memory of Future Past

- So what? This is all good in practice but in RL it will never happen...
- I know of at least three major banking institutions moving to virtualised resources, a couple of utility companies and at least two system of critical national infrastructure
- How are they coping? Well they are putting normal LANs into a virtual environment and using VRF (Virtual Routing and Forwarding) to take care of troublesome things such as network layer separation (or not in a number of cases)
- As to adequate protections, well getting real, tangible, legally admissible auditing and protection mechanisms implemented is a challenge, and one that is often overlooked in the rush towards saving cost and delivering projects

Memory of Future Past

- Because I don't want to get into a Christofer Hoff / Joanna Rutkowska RSA style brawl, it's worth pointing out that this is *just* an idea
- It's an idea based around my reading of HP Superdome 9000 too (I've not looked at HP Integrity Servers, IBM NUMA-Q, IBM P-Series, or Sun Enterprise 15000 either)
- I'm not saying it will work (after all I can't afford the 350k required to prove it) I am saying that it might – I can't emphasise this enough; as I have no desire to spread FUD like manure...
- I am also saying that this may well present a real challenge in terms of securing memory space in virtualised resources...
- PS: Please Hoff don't hurt me...

Memory of Future Past

- Call me cynical and jaded if you will, but I realised that a few of you might have a problem with all this...
- So, I spoke to HP
- At first they were all warm and fuzzy – and then I explained what I thought the issue was within Superdomes...
- To cut a long story short, this is what they had to say...

Memory of Future Past

"The configuration of resources (processors, memory, etc.) in a Superdome (as in all Integrity servers) is done via mechanisms which are protected from being re-programmed by malicious users. In general, ccNUMA machines are no more or less vulnerable than the same number of processors associated with a monolithic memory. There, too, if an attacker can get privileged access to a processor, they can write to the memory that all the processors share and corrupt their flow of execution."

R.F,

HP Software Security Response Team,

HP Cupertino

Memory of Future Past

- My HP contact also confirmed that 'protections' mostly existed at OS level
- You see now I hope why I discussed the AD misconfiguration issue earlier
- What happens if the OS config is out in this instance??
- I asked HP what specific protection mechanisms are built into their architecture, as I couldn't find any in the literature.... They still haven't told me....
- I've also spoken about this issue with specialist security architects, technical designers, security researchers and technical consultants and they agree with me (so if I'm wrong on this, I'm not alone at least....)

The Becoming?

- At this point a number of you are probably thinking that I have a grudge to bear against virtualisation (either at platform or resource level) – in short, I don't
- The implementation of secure virtualisation (if such a thing can ever truly exist) is a challenge at the best of times
- It's made especially problematic when research into platform virtualisation wins over resource virtualisation
- I understand *why* this is happening (filthy lucre) but I don't think it helps...

The Becoming?

- Reality is now intruding on virtualisation in a big way
- Many organisations are now moving towards a 'single box' (or single point of failure) solution
- Because all the separate components of an IT infrastructure are now being lumped together it's all getting jumbled
- Applications, Operating Systems, Network topologies and the Security of all these is being bundled together into a potentially unmanageable mass...

The Becoming?

- The problems don't stop there though...
- Sensible network layer separation is being supplanted by VRF, and vendors are falling over themselves to offer the magic bullet (Hypervisors, VMSafe etc.) – unfortunately one size does not fit all
- From a practical level conducting research on virtualised resources is a feat beyond most (thanks to the costs) and even running security assessment and protection apps on virtual resources presents its own heap of pain
- Virtualisation offers a lot to the bean counters (in terms of speed of delivery and adaptability) but to the security concious it's a potential nightmare waiting to happen

The Becoming?

- Organisations need to seriously start thinking about developing secure architectures for virtual resources
- Bean counters need to start thinking about risk as well as cost and delivery speed
- Researchers need to focus in on weird stuff like memory allocation and page mapping
- If that happens then maybe virtualisation can deliver what it promised...

The Becoming?

- Questions?
- Comments?
- Abuse?

Thanks for Listening

- Anything further, don't hesitate to email clappymonkey@gmail.com or visit me at www.clappymonkey.com