Secure IP Infrastructures For Broadcasters

Essential Guide

EG

ESSENTIAL GUIDES

Introduction from HPE OEM Solutions

Rod Anliker.

Matt Quirk.

According to Gartner, by 2022, cybersecurity ratings will become as important as credit ratings when assessing the risk of existing and new business relationships[1]. If you're here, you're probably concerned about cybersafety and want to understand how you can secure your media business across all its touchpoints.

Organization-wide security starts with IP infrastructure security. As the protection of operating systems and networks become ever more of a priority against a backdrop of increasing data breaches, cyber-attackers are moving past these levels to target server firmware and other deep-level vulnerabilities. As a result, even devices with the most advanced firewalls and anti-virus software can be hacked – from within.

And that's why, to withstand current and future attacks, security needs to be integrated at the firmware and Basic Input/Output System (BIOS) levels. Equally important is an auditable product supply chain, which ensures every firmware update has a provenance and can be validated before it is run.

The Broadcast Bridge has created this Essential Guide to shed light on the latest firmware and low-level vulnerabilities in different devices – from hard disk drives to Solid State Drives (SSDs) – and the technological advances available to remedy them:

- **Part 1 – Secure Servers** introduces the concept of advanced server security, discusses the security challenges firmware presents and looks at possible solutions.
- **Part 2 – Lights Out Control** explores an out-of-band monitoring and control option that will significantly improve security for broadcast infrastructures.
- **Part 3 – Secure Virtualization** explains processor and server memory security, and discusses VM protection, hypervisor vulnerabilities and encrypted virtualization, and how it all fits in enhancing broadcast infrastructure security.

We hope this Essential Guide gives you valuable insights into broadcast system security and practicable next steps.

Sincerely,

Rod Anliker and Matt Quirk

Sponsored by Hewlett Packard Enterprise OEM Solutions

Hewlett Packard
Enterprise
OEM

[1] Gartner, Innovation Insight for Security Rating Services, September 2019.

Secure IP Infrastructures For Broadcasters

By Tony Orme, Technology Editor at The Broadcast Bridge

Part 1 - Secure Servers

In this Essential Guide, we investigate the underlying aspects of computer server design for high value security and 24-hour operation. In Part 1 we look at advanced server security, in Part 2 we understand how servers are controlled, and in Part 3 we gain a deeper understanding of virtualization and the benefits for secure operation.

As computer software and operating systems become more secure, cyber-criminals are looking for other methods of attacking them. These methods may not be immediately obvious but deep within the servers there are potentially hidden vulnerabilities. And special attention must be paid to their remedies.

In the abstract sense, the disk drive is string of sectors numbered sequentially. The servers operating system builds a reference table consisting of the files and their associated sector. If the user wants to read a file, then the operating system will determine from its table which sector the data is stored at and request the sector from the disk drive.

Intelligence

Although the sector size and file referencing may be a function of the operating system, the file mapping to the hard disk drive cylinder, head, and sector is not. The hardware configuration of the disk is most probably different across multiple vendors. Consequently, a method of converting the logical sector number to a physical cylinder, head, and sector is needed implying some intelligence.

Disk drives also have methods of self-diagnosis and fault detection. The SMART (Self-Monitoring Analysis and Reporting Technology) constantly monitors the health of the disk drive and determines if any parameters go out of specification. For example, the spin-up time of the spindle or the error rate of the read system. If these breach certain thresholds, they may inhibit the operation of the drive.

Sectors occasionally fail and the drive can determine the location of bad physical sectors. It marks them unusable to stop the operating system from writing to them. Again, this implies some form of intelligence in the disk drive.

To co-ordinate these tasks the disk drive contains three distinct components; the magnetic storage platters, the spindle and heads actuators and motors, and the disk controller, this is where the intelligence takes place and without the correct safeguards, is the potential target of the attack.

Unintended Vulnerabilities

To understand why the hard disk controller (HDC) can be the unintended source of vulnerability we need to understand more about the low-level operation of reading and writing data to the drive.

Simplistically, when the operating system reads data from the sector, the processor in the HDC will map the logical sector to the physical cylinder, head, and sector, read the data from the sector, store it in cache memory, and send it out on a SATA cable to the servers operating system. Although this is an effective method, it is also slow. For example, if the HDC is a 16bit processor running at 150MHz, the best data throughput we could hope for is 150*16MBits/sec = 2,400MBits/sec, or 2.4GBits/sec.

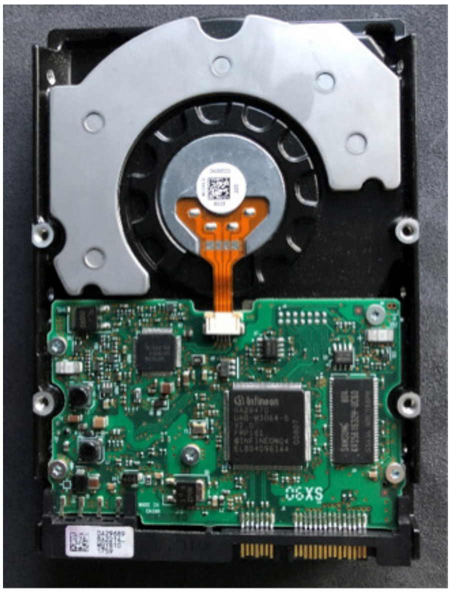

Fig 1 – The underside of this hard disk drive shows the controller PCB with processor, cache memory, and non-volatile memory.

However, the SATA (revision 2) bus can transfer data at 6GBits/sec, so the HDC is underperforming and acts as a bottleneck. To rectify this, a form of hardware acceleration is used called DMA (Dynamic Memory Access). DMA bypasses the HDC processor and copies the data directly from the sector to the cache memory on the HDC. When complete, the SATA DMA process transfers the data from the HDC cache back to the servers operating system. The same is true for writing to a sector but in the processes are reversed. The SATA DMA copies its data to the HDC cache and the HDC DMA copies the data from the cache to the cylinder, head, and sector on the platter. This improves data throughput significantly as there is no processor bottleneck to get in the way.

Consequently, there is a period when data exists in the HDC cache that can be accessed by an unsolicited hostile third party.

The HDC has its own firmware code stored in non-volatile memory on the HDC's circuit board, and it is possible for this firmware to be attacked. Furthermore, many hard disk manufacturers provide maintenance facilities to update the firmware over the SATA connection back to the server making the potential for compromise even higher.

If a hacker can infiltrate the firmware and run just a small snippet of their own code, then they can change the data in the HDC cache allowing them to effectively write to the hard disk drive. If this occurs, then the hacker has control of your server. From there on, they can wreak havoc in your broadcast facility. More importantly, this attack could go unnoticed for many days, weeks, or even months. It could be sat there just waiting for the cybercriminals to enable it.

Other Devices

This method of operation is not limited to just hard disk drives. Solid State Drives have intelligence built into them, along with a cache and ultimately access to the file system. Furthermore, how do you know if the disk drive, graphics card, network interface controller or even the power-supply doesn't already have some backdoor code hacked into it when you buy it?

One of the major advantages of moving to IP is that broadcasters can use COTS (Commercial Off the Shelf) equipment. Although this opens a whole new world of opportunities for broadcasters, the example detailed above demonstrates why we must be careful about our understanding of COTS. It certainly isn't an excuse to procure cheap x86 computers from your local store, put them in a rack, and expect them to perform with the reliability and security of a Tier-4 datacenter with 99.995% uptime.

As far as reliability and security is concerned, broadcasting is in the same arena as banking (especially high-frequency trading), telecommunications, and commerce websites. To understand the type of COTS servers we need to buy, we must look at the procurement decisions made by these industries and learn from them.

As well as having dual power supplies and redundant hardware, high-end COTS servers used in banking, telecoms and commercial websites use the concept of "Silicon Root of Trust". This provides two distinct levels of reliability and security; the procured hardware and firmware is known to be secure, and the server establishes secure operation at a hardware level.

To confirm your devices firmware hasn't been hacked before you even open the sealed delivery box, a reliable supply chain must be established. Enterprise OEM vendors generally provide this as part of maintenance and support packages.

Each device that is installed in the server has an audit trail of trusted suppliers who each validate their area of responsibility, whether it's hardware or firmware. For example, the key manufacture processes for each hard disk drive will be recorded and the vendor will be able to establish who loaded the firmware, where and when.

For broadcasters, this method of thinking is a major step away from the procurement, maintenance and service procedures of the past. Broadcast hardware vendors were trusted by implication as the industry is relatively small and everybody seems to know everybody else or has a friend who works at company XYZ. And much of the equipment has been traditionally hardware based with little opportunity for remote attack. It's only over the past few years computer IT equipment has been making serious inroads to television stations.

As we move to COTS, we must take a more proactive approach to security. Root of trust contracts are a critical component to achieving this and it's unlikely they will be found in local computer stores or at auction websites.

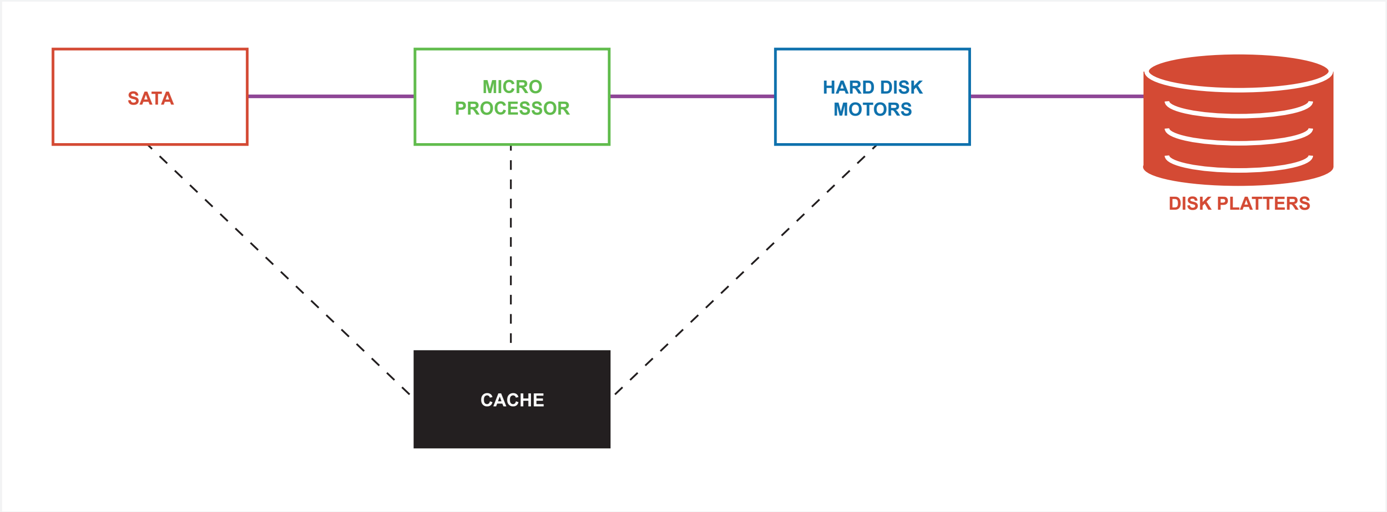

Fig 2 – The cache is used to hardware accelerate data transfer between the computer and hard disk drive, but it also has potential for security vulnerability if not correctly protected.

Another level of device validation takes place in the server itself. As part of the hardware design, vendors providing high-end COTS servers and infrastructure equipment have an extra level of security build into their hardware.

Firmware Certification

Long before the peripheral devices are accessed and the operating system is loaded, thousands of lines of embedded code are executed to interrogate each device to validate its firmware. Vendors work with device suppliers, through trusted partnerships, to provide certification keys for every version of firmware that is executed on the server to authenticate it. This low-level software tests each devices' firmware to confirm no hacks have taken place and no rogue firmware is running on them.

Only when the server has interrogated all the installed devices and authenticated every version of software does it load the boot sector on the primary disk and run the operating system. This takes security to a whole new level and again is something that will not be found in your local computer store. The server not only checks itself but is able to validate all the peripheral devices within it from disk drives and SSD's to network interface cards.

Security First

IP infrastructure security is much more than hiding everything behind a firewall and saying, "it's secure, we're safe". Broadcasters need to look at the whole infrastructure picture and understand vulnerabilities not just at the firewall and user software level, but deep into the hardware and firmware.

One method of mitigating against firmware issues as highlighted above is to work closely with an enterprise OEM vendor to establish a forensically auditable root of trust, and this should even be extended to your traditional broadcast equipment suppliers. If they provide a software and server package as a complete product, do you think you should ask them some searching questions about the root of trust that has been established with the server vendor and their suppliers? What is its provenance? Can they authenticate and validate every version of firmware running on the server no matter where the devices came from? Do they get updates of hacks and new firmware updates?

In Part 2 we continue our investigation into the hardware and look at how "Lights Out" management not only helps with maintenance but improves security too.

Part 2 - Lights Out Control

In Part 1 we looked at advanced server security and how the controller within a hard disk drive or SSD can be vulnerable to hacking even with the most advanced firewalls and anti-virus software. In Part 2, we delve deeper into the remedies and how Lights Out Control further ensures safe server operation.

As well as providing external control of a server, some out-of-band control systems also help significantly improve security. iLO (Integrated Lights Out) is one version of this and is a proprietary embedded design from Hewlett Packard that solves two challenges. Firstly, it allows servers to be controlled even when the operating system is not running, thus providing access to all the peripheral devices, and secondly, it provides unparalleled levels of hardware and server security.

In an old-school datacenter, IT engineers would be near the equipment to carry out certain hardware tasks. For example, if the server needed a power recycle or an operating system needed installing, they would require a CD-Rom or USB to be connected to it with the relevant media available. Many maintenance operations on the server needed a basic operating system to be running otherwise they wouldn't be able to communicate over the ethernet port.

As datacenters developed and their capacity increased, physically running up and down racks of equipment trying to find the correct server was both inefficient and dangerous. Power cycling the wrong server could have disastrous consequences and loading operating systems or Virtualization Code onto multiple servers could prove to be a challenge, especially if they were physically dispersed within a datacenter or multiple datacenters.

The iLO design is an embedded circuit board that sits inside the server but is its own independent system. It has its own ethernet port and IP address to facilitate an external connection, and more importantly, it operates even if the server is powered off.

It is possible to control the whole server remotely as if you were physically sat next to it. The IT engineer can log onto the iLO card and power the server on and off, connect virtual CD-Rom and USB memory devices for maintenance and for loading the operating system.

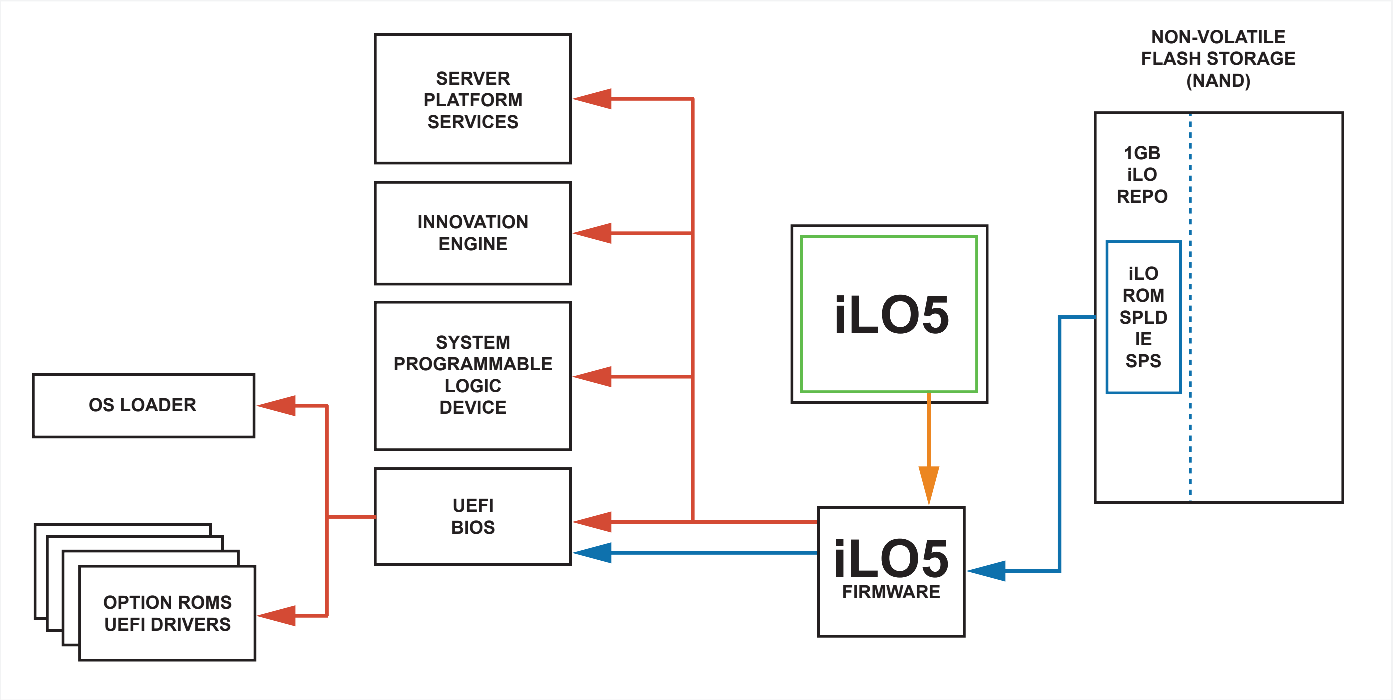

Fig 1 – iLO checks all peripheral devices including the UEFI and BIOS before loading the operating system. As well as providing better monitoring and control for maintenance, security is significantly improved.

Using management software, an IT engineer can load the operating system into multiple servers all over the world without having to leave their desk. The monitoring and diagnosis functions help flag any issues quickly so they can be rectified, and the monitoring and logging system provides detailed information about the server. This includes power supply voltage levels, CPU temperatures, fan speeds, memory capacity, whether the CPU is working or not, the CPU types fitted and much more.

But iLO has developed over the years to include deep security checking and validation of the server's own hardware and firmware.

Boot Validation Sequence

When the power supply to the server is first applied, the CPU motherboard is inhibited and the iLO firmware boots and enables its own ethernet/IP port to allow a web browser to open the password protected monitoring and administration pages. Immediately after this, it will check the UEFI (Unified Extensible Firmware Interface) and BIOS (Basic Input/Output System), and then test the firmware in each of the connected devices. Only when this has successfully completed does it load the operating system.

Back in 1981 when the first open architecture PC's hit the market, they did not have disk drives as standard. Instead, users would use compact cassette recorders to load software into the limited memory space.

Even today, servers and computers still harbor this legacy architecture in their design.

When a CPU is released from its reset state it jumps to a specific memory address where the first instruction of the program will run. As disk drives could not be assumed to be fitted, the CPU had to have some code to execute in place of the operating system and this was the primary role of BIOS. Included in the BIOS is POST (Power On Self Test) and this provided some basic checking of connected hardware.

The original designers of the PC actively encouraged third parties to build peripherals and write software, consequently, the POST could not assume any specific peripherals were fitted. POST tested for connected devices, so the system booted into a known state allowing the operating system to be loaded.

BIOS Legacy

The BIOS also included a method of providing common hardware access for some peripherals using the concept of software interrupts. In the x86 architecture, it is possible to create an interrupt by issuing a software instruction from anywhere in the user-space code. This had the advantage for early developers as they did not need to know where the device port map address of the CPU resides.

The interrupt vectors are loaded during boot by the BIOS so the relevant code is run to access devices such as the keyboard, mouse, screen and disk drive.

As PC architectures and operating systems have developed, software device drivers in the kernel of the operating systems have taken over the role of the BIOS. However, the BIOS is still needed to execute the first instruction of the CPU after reset, provide some basic peripheral testing and system testing, and then load the operating system from the boot device.

To keep up with advances in x86 architecture, and provide improved testing, control, and security, UEFI has been developed and is expected to take over from BIOS. Again, it contains the first executable instruction after the CPU reset but it addresses some of the limitations of the original 16bit 1Mbit x86 addressable space.

EEProm Access

The original BIOS code was stored in a PROM (Programmable Read Only Memory) and it was almost impossible to add code or change it. However, as computer hardware developed, BIOS and UEFI code was stored in EEPROM (Electrical Erasable Programable Read Only Memory). Although this provided some significant advantages for upgrades and bug fixes, it also potentially made these devices vulnerable to attack.

Fig 2 – UEFI is a significant improvement in low level hardware configuration for x86 servers.

During its secure boot, iLO checks the certification of the code in the BIOS or UEFI and validates it. In other words, it can ascertain whether the code has been tampered with. This process continues for all the devices with iLO validating each version of firmware.

Each version of firmware has its own unique key and iLO checks this against its own database. Through the silicon root of trust, vendors collaborate to guarantee each version of firmware is validated by the manufacturer of the device and that it has not been tampered with.

After all the tests have successfully completed, only then does iLO allow the processor to boot its first instruction from the BIOS/UEFI, and then the operating system or virtualized monitoring software. And even after the operating system is running, iLO continues to periodically check each devices firmware to test for malware or if the code has been tampered with.

Silicon Secure Root of Trust

At a first glance there might seem to be some duplication when UEFI is installed instead of the BIOS. That is, the UEFI provides firmware certificate validation as well as providing some external control and testing. However, the iLO system contains the concept of the silicon root of trust. Every component and firmware version can be traced back through a reliable audit trail. This is incredibly powerful in the fight against cybercrime.

Furthermore, iLO provides software interfaces with support for languages such as .NET, Java, and Powershell to facilitate remote control and automation. An array of centralized monitoring and logging tools use these interfaces to constantly monitor the state of each server in the infrastructure to provide speedy notifications if a hardware or security issue does occur. This monitoring information can be collated and tagged to form the basis of AI monitoring systems – a critical tool in the fight against cybercrime.

If it is suspected that a firmware version has been tampered with or has just become corrupted through some temporary hardware anomaly, a method of rolling back or reinstalling the firmware for each device is possible. A copy of trusted firmware versions is securely stored so that if a firmware breach is detected then a known good version of the firmware can be installed in the device thus mitigating any third-party attacks but also keeping the server up and running with the shortest possible downtime.

Cybercrime attacks have changed beyond all recognition over recent years and hardware manufacturers have not only raised their game to fight against these attacks but are increasingly proactive in doing so. Firmware validation and silicon root of trusts are critical in pursuit of this fight. In the same way commercial airplane manufacturers can trace the provenance of every single component on their aircraft, the same is becoming true of advanced high-end OEM server manufacturers. The need to be secure not only includes your server infrastructure but extends to the finished x86 server product a vendor may bring into your broadcast facility. They need to be secure too!

In Part 3 we continue our in-depth journey into infrastructure security and look at virtualization security.

Part 3 - Secure Virtualization

In Parts 1 and 2 we looked at advanced server security and out-of-band monitoring and control, especially with security validation of peripheral device firmware. In Part 3, we investigate virtualization further and its benefits for building secure broadcast infrastructures.

The cyber security landscape has changed beyond all recognition over the past twenty years. Back then, "script-kiddies" were responsible for hacking systems and large events would hit the news headlines once a year with a few thousand data records being compromised. Today, the cyber security landscape is very different.

Millions of records are now at risk daily. The adversary has changed completely with nation state and political activists attacking for ideological reasons. And cybercrime is now big business with some estimates showing a one trillion-dollar value. This demonstrates a lot of determination and resource on the part of the attacker.

In the past, firewalls were the first and most effective go-to level of defense. High-value and critical resources were enclosed within a secure firewall zone and we were relatively sure our systems were safe. But now, in a highly connected world, it's difficult to be certain firewalls will protect our systems completely.

Business Wide Responsibility

More now than ever, security must be considered the responsibility of the whole business and broadcast facility. It must be built into the core of everything we do. This is not just a challenge IT can solve but is a business wide initiative. It's about building secure applications, building into firmware and hardware the silicon root of trust, establishing trusted procurement paths, and understanding how we protect our data.

The primary benefit of virtualization is to make better and more efficient use of servers. For many user applications, the operating system spends a great deal of its time waiting for external events to occur. Disk reads, writes, keyboard and mouse actions, and network interface input and outputs are all relatively low frequency events compared to the speed of the processor. This is a highly inefficient use of expensive processor time.

But as we share the CPU, by implication, we must share all the server resources too including memory, disk drives, network interface cards, graphics cards etc. And the fundamental challenge we face is that the original x86 architecture was never designed to be operated natively in a virtualized mode.

Hypervisor Virtualization

The software running the virtualization service is called the hypervisor and runs on a host server. And the service that allocates machine resource is called the VMM (Virtualized Machine Manager), each virtual machine is referred to as a guest.

There are two types of virtualization topology; bare-metal and hosted. Servers running multiple machines in a data center tend to use bare-metal hypervisors and desktop computers running a hypervisor within their existing operating system are called hosted. For most data center broadcast infrastructures, the bare-metal approach (also called native) is the hypervisor of choice. It sits between the host hardware and guest operating systems.

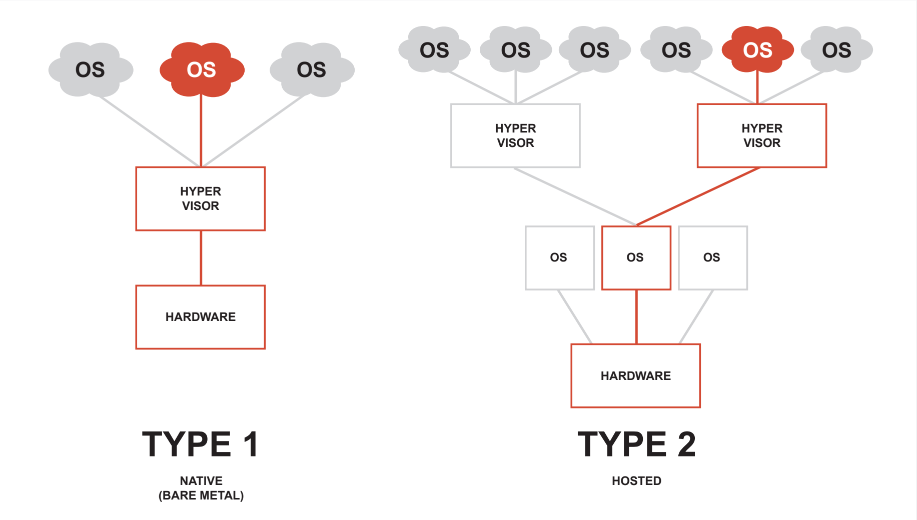

Fig 1 – Two types of virtualization are available. The hypervisor can either run on the bare-metal server, or on an existing OS.

When a user program wants to read a file from a SATA connected disk drive, it invokes a system-call in the operating system (OS). This in turn runs some low-level instructions to talk directly to the disk drive through the CPU's input/output bus. Without virtualization, that is just one OS running on a server, the OS has direct access to the hardware and will execute the necessary instructions. However, when multiple OS's are running on a virtualized server, they no longer have access directly to the hardware. Instead, the hypervisor detects they are trying to access the hardware and intervenes. But critically, the OS doesn't know the hypervisor has intervened, as far as each OS is concerned, they are running on their own server.

Before 2005, hypervisors used a method called binary translation to read through the executable file of the running program on the guest OS and determine which instructions were virtualization critical, and then replace them with the hypervisors code, to give the desired intervention. Although this method worked, there could be a significant overhead in providing binary translation.

Hardware Virtualization

From 2005, Intel and AMD both started to build hardware accelerated virtualized instructions. These would detect the code that was trying to access a peripheral device such as a network interface card and provide a hook back into the hypervisor so it could take over. This is the modern method of providing virtualization on servers but requires CPU's that have this feature built into them.

For both methods, from the point of view of the OS and the program it is executing, it has no idea a third party has effectively diverted control directly away from the peripheral device and is running a proxy instead, often referred to as I/O virtualization. And this is one reason why virtualization is incredibly secure. Combined with memory, network interface, and CPU virtualization, the hypervisor has a trusted level of control and monitoring. If a guest OS has had a security breach then not only can it be disabled, but a snapshot of its resources can be taken for later forensic analysis.

The hypervisor is constantly cycling between the guest OS's to make the most efficient use of the servers' resources. Although physical devices such as memory and hard disk drives are shared across multiple guest OS's, each OS runs in its own shell and cannot communicate or exchange data with the others through these shared resources.

Memory is divided into three categories; virtual, physical and machine. In this instance, the physical memory is an abstraction to give the guest OS the illusion that it is running on the machine directly. The hypervisor translates between the physical and machine memory so that it alone police's access for each guest OS. Every guest OS running on the virtualized server uses a zero-based address space. Therefore, each OS thinks it's running on the hardware memory from address zero. It's not, it's running whatever area of memory the hypervisor has allocated.

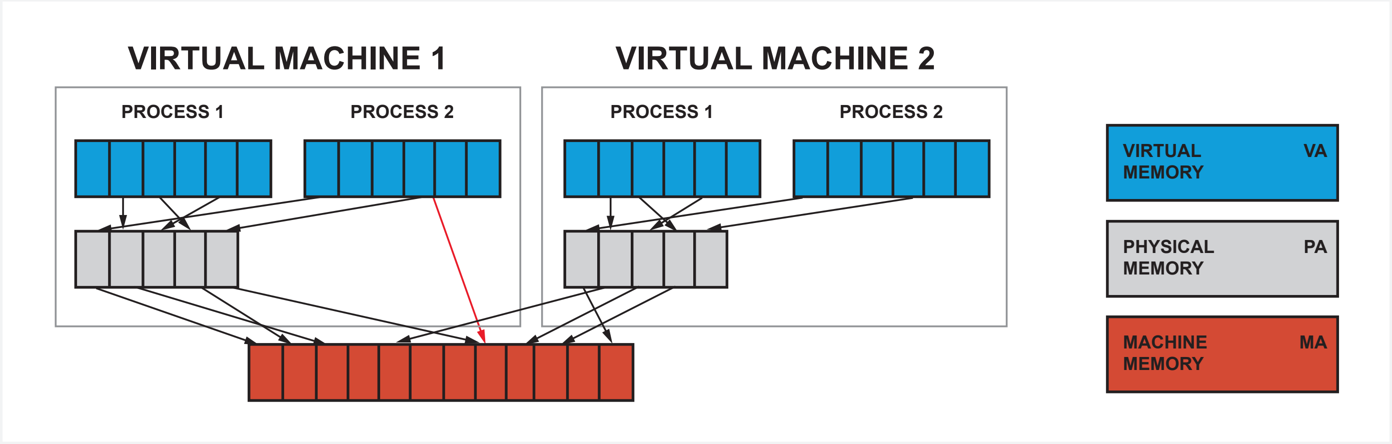

Fig 2 – Memory is abstracted away one layer to allow the hypervisor to check for attacks but at the same time it makes the OS think it's the only system running on the server.

If a guest OS does get compromised and starts to try and find other guests within the virtualized server by attempting memory buffer overflows or segmentation faults, the hypervisor will detect and stop it from accessing memory it doesn't have authorization to use. And importantly, it will log this event and send an alarm to the IT administrator to investigate.

Detecting and logging events is a core part of security as there are two aspects to it. First, there is the detection and second, there is the forensic analysis if a breach takes place. It's great news to be able to stop an attack, but it's equally important to know how it occurred and if there is still some malicious code lurking somewhere in the system. Logging helps enormously with this.

Virtual Drives

Although guest OS's think they have a disk drive to themselves, in reality, their disk drive is a file on the server's disc or connected storage device. If the server uses a disk that is 8TB and each OS has 1TB allocated to it, each 1TB allocation is in effect just a file giving the impression of a complete disk. When a guest OS accesses a disk drive, the hypervisor intervenes and first checks to confirm it has access to that virtual drive. If it does not then an exception is created back to the guest OS, an entry is made in the logging system and the IT administrator is notified.

As the virtual drive is a file, it can be transferred to another area if forensic analysis is required. And this is also the principal of snapshots and why guest OS's can be moved between physical servers. As well as transferring the virtual disk, the hypervisor takes a copy of the CPU state, input/output state, and other connected peripherals and system configurations.

Network interface cards can also be virtualized. A virtual network adapter is created, and multiple guest OS's can access it, each thinking they are the only OS that has access to the NIC. Multiple NIC's can be combined to aggregate the link to make each OS think it has a high-speed connection. Not only does this provide better redundancy, but it also gives much higher connection speeds. As the hypervisor has abstracted away the physical hardware from the guest OS's and in doing so has given another level of security.

Improved NIC Security

The virtualized NIC is a software abstraction of the underlying hardware so it can have extra levels of security added to it. For example, it can be enabled to provide VLAN layer 2 validation and firewall protection. IP access to the guest OS's can be limited to specific protocols such as HTTPS thus restricting an external hacker from trying to run TELNET or some other control software.

As well as providing more efficient use of the server's resource, the hypervisor provides an unprecedented level of security between guest OS's and to the connected world. It provides layers of monitoring, logging and detection, so that if a guest OS does become compromised, the IT administrators can be notified quickly without jeopardizing the rest of the infrastructure.

The Sponsors Perspective

Media Companies: Advance Your Security And Innovation Lifecycles

Hackers are always improving the level of sophistication and constantly finding new surface areas to attack – resulting in the surging volume and frequency of cyberattacks. Between 2016 and 2018 alone, ransomware attacks increased 15-fold, and going forward, global cybercrime is expected to cost $6 trillion by 2021[2].

The true cost of cybercrime extends beyond stolen data or revenue lost due to downtime. It can cost a business its reputation. More than 80% of consumers say a business with a good data security history is a key characteristic they look for when deciding who to buy from[3].

To protect data, there must be a security mindset present from the beginning – businesses cannot afford to let security be an afterthought.

Sponsored by Hewlett Packard Enterprise OEM Solutions

Hewlett Packard
Enterprise
OEM

[2] Forbes, "Hewlett Packard Enterprise Releases iLO Amplifier Pack With Server System Restore," February 2018.

[3] https://www.cbi.org.uk/media-centre/articles/almost-9-out-of-10-people-say-businesses-that-protect-theirdata-will-win-their-custom/

That's why HPE is committed to protecting the very core of their IT infrastructure – its servers. HPE's ProLiant Gen10 servers are the world's most secure industry-standard servers[4], and are built with the most advanced security capabilities available today. The foundation of HPE's latest server iteration is built on the fifth generation Integrated Lights Out (iLO 5) management chipset.

The iLO chipset, included in HPE servers for years, provides secure out-of-band management functionality regardless of the server hardware or OS status, and it is available whenever the server is connected to a power source, even if the server's main power switch is off. iLO offers strong authentication, highly configurable user privileges with robust authorization processes, and encryption of data, keystrokes, and security keys.

Engineering A Silicon Root Of Trust

The iLO 5 chipset breaks new ground by adding silicon-level security to the Gen10 servers, enabling an unparalleled level of hardware security. Dubbed the silicon root of trust, the chipset is built straight into the silicon hardware itself, making it impossible to alter and allowing firmware authentication to extend as far back in the supply chain as possible. It provides a secure start-up process and, most importantly, provides firmware runtime validation and secure recovery in the unlikely event of a security breach.

As we've learned in this Essential Guide, the firmware is becoming an increasingly attractive target as operating systems, applications, and hypervisors become more secure. Because the firmware always loads over a million lines of code before the operating system even boots, the firmware and BIOS must be protected.

The silicon root of trust that is burned into the silicon components on the motherboard makes it literally impossible to compromise, because the system cannot boot without this circuitry. This means that the safety of your broadcast systems is in the best hands possible.

Maintain A Pipeline Of Media Technology Innovation

As a broadcast technology provider, minimizing risk is only one area of concern. The arrival of 5G and challenges from over-the-top (OTT) media providers provide others. Against this backdrop, remaining relevant and competitive is becoming ever more pressing for traditional broadcasters and new entrants looking to differentiate themselves. Many are updating their products to meet new demands for always-on, high-definition content that can stream on multiple devices with minimal latency.

Media providers are also turning to HPE OEM Solutions to bring these innovative technologies to market faster. A partnership with HPE OEM Solutions is more than a traditional manufacturer/supplier relationship – it takes a long-term view and offers end-to-end advice, support and planning, to ensure you become and remain a media industry leader.

An HPE OEM Solutions partnership can help your broadcast technology business:

- Access global support: Whatever your project, query or concern, HPE subject matter experts and technical support teams are on hand to help in the first instance.
- Power Edge analytics: Create solutions that manage, manipulate and deliver content more intelligently with Edge solutions, which enable real-time analysis and actioning of operational data, and optimize capabilities and processes.
- Maximize uptime: Predict and prevent problems across your IT infrastructure stack with the world's first self-healing infrastructure technology – HPE InfoSight – to reduce both unplanned downtime and operating costs.
- Expand globally: Grow your business with confidence, assured that your products are being produced and delivered at the highest standards by HPE's worldwide supply depots and logistics networks.

Receive In-depth Training And Product Knowledge

Ongoing development goes hand in hand with continuous advancement. At HPE, we understand the value of education and training, which is the reason behind the newly launched HPE OEM Education Services.

Delivered via a digital platform, this training service provides our OEM partners with the targeted technical knowhow needed to install, configure and support HPE technology.

Sponsored by Hewlett Packard Enterprise OEM Solutions

Hewlett Packard Enterprise
OEM

[4] https://www.hpe.com/uk/en/servers/gen10-servers.html

Not only will this reduce our partners' time to productivity with HPE technology, but it will also enhance their operational efficiency, enable the delivery of support services their customers need and ultimately, help them better serve their customers in the field on an ongoing basis – at less cost and disruption to the business.

The value of training has been demonstrated by Konica Minolta, a Japanese technology company and HPE OEM Solutions partner. They approached HPE OEM Solutions requesting training that would help them successfully deploy the HPE servers embedded within their technology. After determining the salient knowledge gaps, HPE OEM Solutions selected six courses – including one that was tailored specifically for Konica Minolta. This enabled Konica Minolta to deploy servers faster and accelerate the roll-out of systems to customers.

At HPE OEM Solutions, we have the solutions, services, specialists and supply chain to help take your business further and stay ahead of the industry and competition. If you're ready to expand your business reach and accelerate go-to-market for your broadcast technology solutions, an HPE OEM Solutions partnership could be ideal for you.

Get in touch or find out more at www.hpe.com/solutions/oem-media.

Sponsored by Hewlett Packard Enterprise OEM Solutions

Hewlett Packard Enterprise

OEM

WP

WHITE PAPERS

EG

ESSENTIAL GUIDES

MEDIA

CS

CASE STUDIES

Find Out More

For more information and access to white papers, case studies and essential guides please visit:

thebroadcastbridge.com

9/2019

Sponsored by Hewlett Packard Enterprise OEM Solutions

Hewlett Packard
Enterprise

OEM