CyberHive Trusted Cloud Product Datasheet

CyberHive's unique solution uses a combination of hardware-based cryptography and advanced whitelisting to protect servers from all unauthorised activity and malware in a way that conventional solutions cannot match. On average, data breaches go undetected for as much as 6 months before companies identify that they have been attacked. Trusted Cloud can cut this down to a matter of seconds.

The patented intrusion detection technology has been co-developed over several years with The University of Oxford. It offers several key benefits that no other security systems provide, including:

1. **Stops attacks that conventional security systems cannot detect.** Standard endpoint security solutions based on techniques such as blacklisting and heuristics often fail to detect bespoke attacks and advanced hacking. Trusted Cloud will detect any unauthorised software changes, no matter how small or well disguised.
2. **Our technology is virtually impervious to hacking.** Our intrusion detection software makes use of hardware encryption modules, built into most commercially-available servers, to prevent any unauthorised software from running on the servers.
3. **Protect against human error.** Our distributed verification technology prevents a simple security lapse on the part of any one person from compromising the security of your servers.
4. **Verify the integrity of your servers independently of the cloud provider.** The technology allows a company to independently verify that their servers are secure, not only from external attacks, but also from malicious intent from the cloud service provider, or even from their own employees.

What is Trusted Cloud?

Trusted Cloud protects your critical data by ensuring that no unauthorised code can run undetected on your critical server infrastructure.

The technology employed is a cutting edge approach called distributed hardware-backed whitelisting. Every few seconds, each server in your estate will generate an audit list of all programmes and configuration running on the server. This list is digitally signed using a cryptographically secure hardware module. This audit data is then sent to one or more verification servers where it will be cross-checked against a previously generated and signed whitelist. Even a single line of unauthorised code will be instantly detected and flagged, allowing your security team to stop an attack in its tracks before any damage can be inflicted or data lost.

Simplified System Diagram

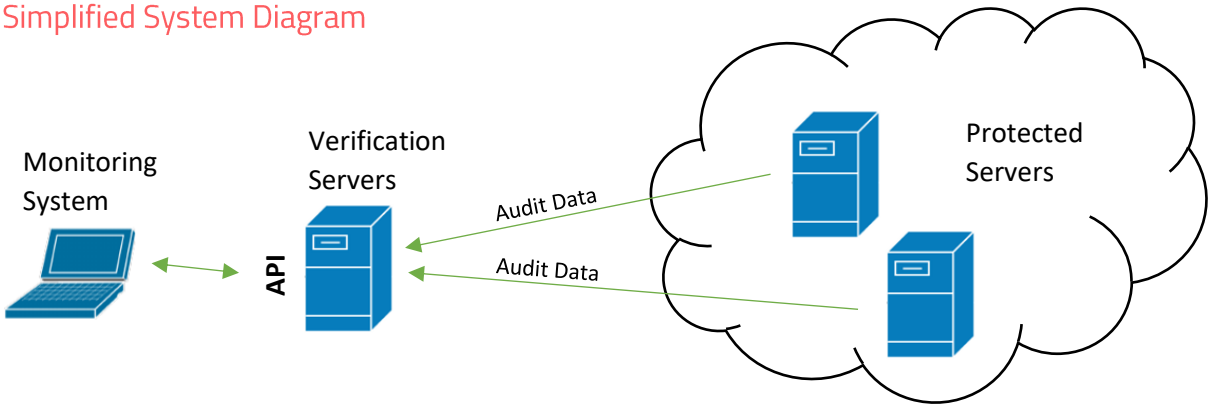

Trusted Cloud makes use of the Intel Trusted Platform Module (TPM) chip which is already installed on most server motherboards to record and digitally sign the audit data. The use of the hardware TPM chip to sign the files makes it virtually impossible to falsify the signature, since the key is only available in hardware in the TPM chip. The use of multiple Verification Servers to check against the whitelist prevents intrusions into any one server due to compromised security - any intruder would need to simultaneously attack multiple machines. Furthermore, one or more of the Attestation Servers can be placed in a third-party data centre and even administered independently, preventing unauthorised tampering from the service provider or even from company's own IT staff.

A Monitoring GUI is used to display the trust status and audit data to the security team. This connects to the Verification Servers using an API, so the customer can choose to integrate the monitoring into their existing Security Operations Centre monitoring solution instead of requiring a separate dashboard if appropriate.

System Components & Requirements

Protected Servers

Trusted Cloud can be installed to protect physical "bare metal" servers and Virtual Machines (VMs).

CyberHive can host VMs in our world-class robust cloud platform, hosted from Tier3 datacentres in central London. Alternatively Trusted Cloud can be installed on customer machines on premise or in a 3rd party 'private cloud'.

Customer Servers which need to be protected are required to install a small agent to generate, sign and transmit the audit data. No further configuration is required on these servers after initial set up which takes only a few minutes. Trusted Cloud introduces no significant processing or memory overhead to the server in most use cases.

Supported Linux Operating Systems:

- Ubuntu 16, 18
- Centos 7
- Debian 9
- RedHat

Additional Linux distributions can be supported on request. Support for Microsoft Windows is under development and will be available in 2020.

To provide the full benefit of hardware-backed verification for Virtual Machines (VMs), the Hypervisor used must support access to the TMP hardware.

Supported Hypervisors:

- KVM (OpenStack)
- VMWare (available from Q2 2019)

Trusted Cloud will also be available on Microsoft's Azure Platform from late 2019.

If an alternative hypervisor is used, or if the VMs to be protected are hosted on a 3rd party public cloud, Trusted Cloud can be installed using a 'Virtual TPM'. This offers the same benefits, however the level of robustness to attack through a compromised Hypervisor is slightly diminished.

Verification Servers

The verification service requires a small cluster of virtual machines to be provisioned to carry out the cross checking of audit data against the whitelist. These servers can be hosted by CyberHive and provided as 'Security as a Service' (SECaaS).

Alternatively, one or more verification clusters can be hosted by the customer as an on-premise solution. CyberHive would install and configure the verification cluster as part of the initial setup.

The verification service is configured through an API. A simple Web UI is provided which connects to this API to perform all required management functions including adding or removing servers from the infrastructure, creating and updating whitelists etc.

Whitelist creation can be carried out automatically and takes only a few seconds. The service has been designed to integrate with a DevOps approach, introducing minimal overheads to existing development and server management processes.

Monitoring System

Most standard monitoring and logging systems can be simply configured to connect to the Verification Cluster APIs, allowing the security status of customer servers to be displayed and monitored through existing customer systems. Additionally, CyberHive can provide a dedicated web-based dashboard to present the security status information to the customer security team if required.