Cyber AI & Darktrace Cloud

Securing New Computing Models, Applications, Users, and Devices

Contents

Executive Overview

The rapid adoption of cloud and SaaS services has transformed the digital business and fundamentally reshaped the challenge of defending the enterprise against advanced attacks.

Driven initially by the need to cut costs and increase efficiency, the transition to the cloud now serves as an essential conduit for digital transformation projects – from applying advanced analytics to big data sets, to supporting edge computing and devices that underlie everything from smart cities to connected cars. Yet from a security perspective, these new computing models have expanded the attack surface at an alarming rate, introducing new threat vectors across an increasingly dispersed corporate network.

This trend presents a special challenge for strained security teams, who must now cope with an environment where they have limited visibility and control, and where their familiar on-premise security tools are often not applicable. Additionally, the ease with which developers can spin up a cloud instance and bypass the IT or security team can expose the business to considerable risk, demanding a new DevSecOps approach which may be unfamiliar to teams who have grown up on the traditional on-premise network model.

More generally, the security challenges presented by the cloud are largely governed by a Shared Responsibility Model, which delineates the respective areas of the cloud that providers and customers are expected to manage and secure. While the customer's portion of the Shared Responsibility Model varies across IaaS and SaaS, the general thrust of the Model plainly illustrates that outsourcing certain IT processes to the cloud does not amount to outsourcing your security function altogether.

Most organizations recognize this reality but few, if any, are satisfied with the cloud-specific security solutions available on the market, nor can they immediately pivot their teams to a DevSecOps approach as an alternative. While many IaaS and SaaS providers offer native security controls to help customers secure their own portion of the Shared Responsibility Model, these controls are often limited in scope and tend to be useful for compliance, rather than proactive and real-time cyber defense. Even within this limited scope, native security controls can only be effective if they have been adequately deployed by the cloud customer.

To supplement native controls, third-party solutions like Cloud Access Security Brokers (CASBs) and Cloud Workload Protection Platforms (CWPPs) can also be useful for enforcing policies and providing cross-cloud visibility, but they struggle to detect subtle threats and anomalous behaviors not captured by pre-defined rules or policies. As threats develop and become more sophisticated, organizations require a fundamentally new approach to securing cloud environments against advanced attacks, before they have time to escalate into a crisis.

Beyond these native and third-party controls, what scenarios do you need to consider for cloud security? Some additional use cases include:

- Spinning up of an instance bypassing IT and security
- Production data being moved to less secure test systems
- Lateral traffic within the cloud
- Office 365 and Salesforce users moving data outside the cloud
- Edge devices that use cloud as a conduit
- Unusual, subtle changes in user behavior
- Attacks moving at machine-speed requiring immediate action
- Correlation and contextual analysis to detect anomalies
- Neutralizing malicious emails in Office 365, especially from trusted senders along the supply chain

> Darktrace Cloud represents a new frontier in AI-based cyber defense. Our team now has complete, real-time coverage across our SaaS applications, cloud containers, and city-wide distributed sensors.
>
> **City of Las Vegas**

Powered by machine learning and AI algorithms, Darktrace's cyber AI technology extends beyond the security that native controls and third-party tools provide. Using software and sensors, Darktrace Cloud covers all the use cases stated above by analyzing rich traffic and data flows across cloud and SaaS environments. Darktrace's AI learns the normal 'pattern of life' for every user, device, and container – without relying on prior assumptions or manual input.

This evolving understanding of 'normal' allows the platform to autonomously detect and respond to external attacks and insider threats in real time, while providing complete visibility across the digital business in a single pane of glass. This white paper examines the security gaps that Darktrace's platform fills using machine learning and AI. By learning the full extent of your organization's evolving 'pattern of life', Darktrace's AI is uniquely suited to detect and neutralize subtle deviations indicative of a threat in the cloud, in concert with your ever-changing digital estate.

Threat Vectors in the Cloud

The Shared Responsibility Model in the cloud outlines the respective security roles of Cloud Service Providers (CSPs) and customers across the main service models: Infrastructure-as-a-Service (IaaS) for systems and storage, and Software-as-a-Service (SaaS) for business applications.

With IaaS, the CSP is broadly responsible for securing the basic infrastructure components – including networks, servers, VMs, and containers – while the customer is expected to manage the guest operating system, any application software, and the configuration of native security controls. With SaaS, the CSP is responsible for the infrastructure and applications, while the customer must ensure that user and network activity is properly managed and secured. This leads to a specific set of threat vectors that the cloud customer must be able to defend against, but which most native security controls and third-party offerings are ill-equipped to detect at an early stage.

Gartner predicts that by 2022, at least 95% of cloud security failures will have occured in the customer's portion of the Shared Responsibility Model.This is a startling figure, but by unpacking the main threat vectors through which these failures might occur, we can better understand what the cloud customer can do to mitigate these risks effectively.

Insider Threat

Most of the industry recognizes that leading CSPs and SaaS vendors are highly resistant to security breaches, at least in their portion of the Shared Responsibility Model. Yet the unique challenges introduced by the cloud - from a lack of visibility and control, to the new unfamiliar mindset required by the agility and speed of digital business – have magnified the traditional risks that fall in the customer's area of responsibility.

In particular, insider threat represents a dangerous attack vector that has always posed a risk, but which has taken on a new dimension and agility via the cloud. These types of attacks originate from within the organization - through disgruntled, careless, or compromised employees, cloud consultants, and other business associates who abuse their access to internal systems.

Malicious insiders in particular have the advantage of familiarity with the systems they manipulate, and can take their time in preparing and perpetrating the attack. By leaking or manipulating data slowly over days and weeks, these actors are uniquely positioned to compromise entire cloud environments and evade rule-based security tools designed to monitor abnormal activity, to the extent that these have been deployed at all.

High-profile cases of insider threat – from Edward Snowden's leaks in 2013 to Tesla's experience of insider sabotage in 2018 – have sent shockwaves across the IT security industry and beyond. The fact that supposedly the most secure networks could be breached by those with sufficient motivation and technical know-how has sent a clear signal to security professionals that the threat may already be inside. With the advent of the cloud, security teams are now faced with a critical blind spot in a highly sensitive area of the business, where insiders can often operate without triggering suspicion.

Compromised Credentials

For similar reasons, the risk of an external attacker using legitimate credentials to gain access has also become a critical risk for organizations with little to no visibility in the cloud. By using the right set of credentials and evading traditional security controls, these threats have the potential to jeopardize an entire organization's critical assets, especially as employees continue to re-use passwords across personal and professional accounts.

More often than not, employee login credentials are retrieved through data breaches, exposures, or phishing campaigns and sold to the highest bidder on the dark web. Once purchased, credentials can be used to move laterally within the cloud to access critical systems and data. Attack missions vary from data exfiltration or manipulation, to corporate espionage.

With IaaS in particular, user credentials for system administrators are often seen as the keys to the cloud kingdom, giving hackers access to sensitive data in production and test environments, and even management of cloud infrastructure itself. Beyond stealing or altering critical data, cyber-criminals can use system administrator credentials to leverage the cloud's compute power for their own nefarious purposes, spinning up cloud instances to launch extensive crypto-mining operations or Distributed Denial of Service attacks. For cloud-only businesses in particular, this threat poses an existential risk.

Misconfigurations

Beyond direct cyber-attacks, one of the most common threat vectors in the cloud continues to be critical misconfigurations in IaaS environments. While human error can never be completely avoided, misconfigurations are often a natural consequence of the agility of deployment and rapid instantiation of test containers and data sets facilitated by the cloud, which often leads users to move quickly at the expense of security.

In today's digital business, developers have the ability to deploy a cloud instance in minutes, without having to consult security teams, IT, or QA. In the past, a slower workflow meant that these functions could afford to work in siloes, but the simplicity and speed of the cloud requires learning an agile and all-inclusive DevSecOps approach, which would ideally ensure that security considerations are brought to bear on a given cloud instance without slowing down the developer.

And yet, not all organizations are equipped to quickly adopt a radically new mindset, and this rocky transition has often led to critical misconfigurations that leave the business vulnerable to attack. In many cases, these misconfigurations occur against a backdrop of 'pet' and 'cattle' cloud deployments, where the 'pet' represents the well-secured and sanctioned production environments, and the 'cattle' represents disposable and rogue test environments which are often spun up without security concerns in mind. The resulting misconfigurations can range from forgetting to deploy native security controls, to configuring a test environment to be public-facing when it shouldn't be, or even forgetting the environment was meant to be disposed of at all. In this latter case, developers might even leave real data or credentials out in the open, where they can be picked up by routine scans from hackers looking to make a profit on the dark web.

> "Misconfiguration of cloud platforms is the number one threat to cloud security."
>
> **Crowd Research Partners**

Unsecured APIs

Unsecured APIs have become one of the most impactful misconfigurations in the cloud, being listed in the Top 10 OWASP Application Security Risks in 2017. An application's API is ultimately the interface to back-end data, so any vulnerability in error response handling would naturally be an attractive target for cyber-criminals with a range of motivations. As with other misconfigurations, developers working at the speed of digital business are often not working hand-in-hand with security, and can sometimes fail to harden APIs well enough to account for potential abuse – from spoofing applications for legitimate users, to coding the API to erroneously respond to attackers with highly sensitive data.

Edge Computing

In a wide range of areas, the enhanced computing power provided by the cloud has furnished a springboard for the development of new and innovative technologies.

Edge computing in particular stands as one of the cloud's most notable off-shoots, even as it represents a radical shift away from processing data in central nodes - bringing computing logic closer to the physical data sources to reduce latency and increase bandwidth.

Data analysis is thus increasingly being brought to the edge, where distributed sensors, IoT devices, and even applications can make rapid, real-time decisions, only having to send processed data back to the cloud when it's ready to be stored or further aggregated.

While the cloud may be taking a backseat in this use case, the explosion of edge computing and devices across manufacturing plants, smart cities, and oil rigs continues to expand the attack surface through which threats can ultimately find their way back to the central cloud.

Limitations of Native and Third-Party Security Tools

Against this backdrop of evolving threats, CSPs and third-party vendors have developed a range of security tools to help defend the customer's portion of the Shared Responsibility Model. While these solutions can provide some measure of protection, they are generally ill-equipped to defend against advanced threats in the cloud.

CSP-Native Security Controls

Apart from securing their own portion of the Responsibility Model, most cloud providers offer native solutions to help customers implement basic cyber hygiene in the cloud. These can span from firewalls, two-factor authentication and IAM tools, through to log monitoring and threat intelligence integrations.

While these native controls are a good start and can contribute to your organization's overall defense in depth strategy, they are often not sufficient in practice. As organizations continue to adopt cloud services from multiple providers, native controls cannot be relied on to provide comprehensive coverage, as they are often exclusively designed for the cloud environment of the specific provider.

Most businesses migrating their workloads to the cloud use multiple IaaS providers, while a study conducted in early 2018 found that the average number of cloud apps used in the enterprise has risen to nearly 2,000. Now that 'multi-cloud' deployments have become the norm, a stove-pipe approach to cloud security has quickly become outdated, and demand for provider-agnostic solutions that cut across the full range of cloud and SaaS environments is on the rise. It is also worth emphasizing that all native security controls must be adequately deployed by the customer, who may not be familiar with configuring these tools in the cloud.

> "With a new threat to face every day, traditional tools designed to spot known threats are no longer sufficient."
>
> **Inphi**

Yet even if deployed and configured properly in a single-cloud enterprise, most native security controls also tend to be more useful for regulatory compliance than cyber security. While the logs they collect can be given to auditors to demonstrate some level of visibility in the cloud, this view is often retrospective and unlikely to catch the most serious threats. As cyber-criminals continue to exploit blind spots in the cloud, log monitoring will hardly be enough to catch the silent, stealthy attackers that lurk beneath the surface.

Third-Party Cloud-Specific Tools

Third-party vendors have also begun to develop cloud-specific security solutions like Cloud Access Security Brokers (CASBs) and Cloud Workload Protection Platforms (CWPPs) to fill in the gaps left by native controls. CASBs are largely designed to secure SaaS applications, providing features for cross-cloud visibility, DLP, and compliance. CWPPs aim to cover hybrid IaaS environments, delivering application controls based on whitelists, as well as hybrid cloud visibility and network segmentation functionality within containers and workloads.

Broadly speaking, CASBs and CWPPs specialize in preventative controls, compliance, and visibility within the scope of their respective use cases.While these capabilities have their place, they will often fail to catch subtle or targeted attacks. Most vendors in this area have recognized this fact, and some have sought to implement rudimentary anomaly detection functionality as an add-on, by establishing a static baseline and pre-defining benign and malicious behavior. Yet this legacy approach cannot be expected to provide sufficient protection against advanced attacks, even as these markets continue to mature.

Darktrace Cloud

Darktrace's cyber AI technology brings a fundamentally unique approach to real-time cyber defense in the cloud. Built on a foundation of unsupervised machine learning and AI, Darktrace Cloud analyzes rich data flows within and across cloud workloads and SaaS applications, learning a normal 'pattern of life' for every user, device, and container. By correlating subtle deviations in behavior in real time, Darktrace Cloud can spot and stop the full range of cyber-threats in the cloud, from malicious insiders and external attacks, through to critical misconfigurations that can expose the business to high-impact compromise across the digital estate.

The power of Darktrace's technology lies in its self-learning approach, which does not rely on pre-defining 'benign' or 'malicious' behavior in advance. Instead, Darktrace Cloud models the normal behavior of users, containers, and devices in relation to their past, their peer group, and the wider organization, continuously revising its calculations in light of new evidence, and correlating weak indicators to establish an evolving measure of threat probability.

Darktrace's approach is critical in this new age of cloud-based cyber-threat, where insiders with privileged access and external actors with admin credentials can sweep through an entire cloud infrastructure without setting off alarms. The cloud provider cannot (and should not) be expected to secure the cloud against trusted connections, while third-party tools with anomaly detection capabilities can only do so in a blunt and flat-footed fashion. By relying on fixed learning periods and pre-defined notions of 'benign' and 'malicious', these tools can only detect the most obvious threats. In contrast, Darktrace's unsupervised machine learning and AI can go beyond what humans already know or can imagine, and detect subtle deviations that may point to a developing threat.

Instead of relying on pre-defined rules and policies, Darktrace Cloud embraces the uncertainty inherent in today's complex digital environment. All significant deviations are seen and correlated, resulting in the detection of genuine threats, without producing floods of false positives.

> "Darktrace Cloud works perfectly for AWS. It's lightweight, easy to use, and makes us feel much more comfortable spinning up our cloud infrastructure."
>
> **Innovating Capital**

Darktrace Cloud for IaaS and SaaS

Darktrace Cloud can easily integrate with your diverse digital estate, including IaaS environments like AWS and Azure, and SaaS applications like Salesforce, Box, G Suite, Dropbox, and Office 365.

box
salesforce
Microsoft Azure
aws
Office 365
Dropbox
rackspace
the open cloud company
Google Cloud Platform
G Suite
EGNYTE

Darktrace Antigena: Autonomous Response in the Cloud

Darktrace's AI not only detects but can also autonomously respond to in-progress cyber-threats in the cloud. Darktrace Antigena, the platform's autonomous response capability, uses artificial intelligence to take targeted, measured action in response to high-confidence cyber-threats – stopping their spread in real time, and giving the security team time to catch up.

The types of actions Darktrace Antigena can take vary depending on the specific cloud environment or SaaS application being used, as illustrated in the lists below which are not exhaustive nor definitive across all SaaS or cloud platforms.

To neutralize in-progress attacks in cloud environments like AWS and Azure, Darktrace Antigena can:

- Terminate a virtual machine or edit its properties
- Edit S3 bucket permissions in AWS
- Temporarily disable a user's programmatic access
- Reset user passwords to disable management access
- Edit user permissions
- Temporarily stop sharing a document

In SaaS applications like Office 365, Salesforce, G-Suite, and Box, Darktrace Antigena can:

- Kill a user's active sessions
- Temporarily disable users
- Restrict or delete file sharing settings from certain files and folders
- Restrict a user from accessing certain parts of the cloud environment
- Suspend members from teams and hence their access to certain shared files (in Dropbox, for example)

Darktrace Threat Visualizer: Complete Visibility

Most organizations migrating infrastructure and applications to the cloud struggle with migrating visibility and control along with it. Even security teams that properly configure and deploy native and third-party tools rarely have access to granular, real-time visibility to achieve continuous monitoring for interactive and contextualized threat investigations.

To provide this visibility across your digital infrastructure, Darktrace's graphical Threat Visualizer interface provides a single pane of glass from which anomalous activity in cloud workloads, SaaS applications and elsewhere can be visualized and investigated in real time. The Threat Visualizer is designed for users of all maturity levels, from forensic security experts, to business executives and less experienced members of the IT team.

A wealth of information can be variously queried and exposed using the interactive features within the Threat Visualizer, including a dynamic dashboard where users can filter incidents based on their level of severity, and an interactive Play-Back tool which lets users replay incidents and zero in on the real-time context around each event.

> "With Darktrace Cloud, we are shining a flashlight into the darkest corners of our network."
>
> **Addivant**

Technology Deployment Scenarios

Hybrid Cloud (IaaS)

For cloud, edge, and physical deployments, Darktrace's lightweight, host-based OS-Sensors are installed on each cloud endpoint and configured to send intelligent copies of network traffic to a local vSensor deployed in the same cloud environment. The receiving vSensor processes the data and feeds it back to Darktrace software in the enterprise, which correlates behavior across the organization's cloud and physical environments.

AWS and Azure customers additionally have the option of using Darktrace Connectors to monitor system administrator activity that may not be seen at the OS-Sensor level, such as logins, file changes or data transfers.

> "When we activated Darktrace Cloud, it was like flipping on a switch in a dark room."
>
> **TRJ Télécom**

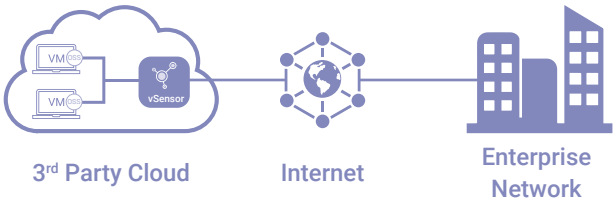

Cloud-Only (IaaS)

For organizations that run their infrastructure exclusively in the cloud, Darktrace can manage the deployment as a dedicated service, installing vSensors and OS-Sensors in the organization's cloud environment and feeding data back to a managed Darktrace cloud instance for analysis.

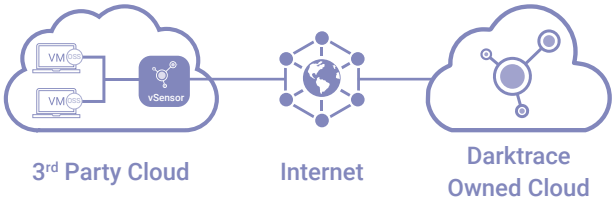

SaaS

For SaaS deployments, Darktrace SaaS Connectors are remotely installed in the enterprise network and interact directly with the SaaS vendor's security API, via HTTPS requests. This allows user interactions to be processed and monitored by Darktrace within minutes of creation, whether they originate inside the network or from remote locations.

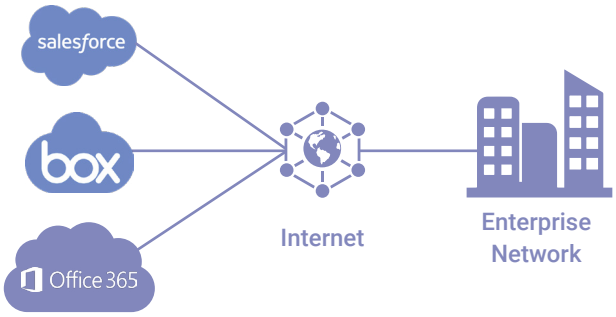

Real-World Threat Discoveries

Internal Data Theft From the Cloud

A retailer decided to restructure its IT department. In so doing, they had to let a number of employees go. One of the affected employees – an IT manager – downloaded contact details and credit card numbers from the customer database. Darktrace detected data transfers to a home server via that company's regular data transfer service. The employee was likely intending to sell the information for a profit.

The database was held on a third-party cloud service in order to enable flexible working and reduce hardware costs. The retailer's business model was based heavily on the usage of cloud synchronization, storage, and file transfer services. However, this IT manager demonstrated how cloud services can be exploited for insider data exfiltration.

The company's marketing department frequently used this cloud service, but it was highly unusual for an IT manager to send data externally through the cloud. Darktrace was able to make this distinction because it continually learns normal activity for every user and device, and compares behavior between devices to identify similarities.

Darktrace's technology detected this slight deviation from the normal 'pattern of life', enabling the platform to identify this threatening and subtle behavior even though the cloud service was regularly used for legitimate purposes.

Darktrace detected these anomalies in real time and provided the company with detailed information on the precise nature of the compromise. The employee's credentials were revoked, and the company quickly retrieved and secured the customer data.

External Attack on Cloud Perimeter

A financial services organization was hosting a number of critical servers on virtual appliances in the cloud, some of which were meant to be public-facing, and some of which were not.

When configuring their cloud deployment, they mistakenly left an important server exposed to the Internet when it was meant to be isolated behind the firewall. This could have happened for a variety of reasons, possibly because of a quick and chaotic migration, or because the security team simply was not as familiar with the native firewall provided by their CSP.

The exposed server was being continuously bombarded by attacks from malicious third-parties trying to gain access to that device, and through there, pivot into the cloud and potentially pivot back into the center of their physical network. Crucially, the customer didn't know about this because they didn't have visibility into what was going on in their cloud.

Yet after a swift installation, Darktrace quickly detected that the device was receiving an unusual amount of incoming connection attempts from a wide range of different external sources.

Darktrace identified the pattern of attack and alerted the customer to the ongoing risk, and they were able to shut down the hole in their security and secure their cloud perimeter, before they became victim to a more serious Denial of Service attack, or before this attack had actually managed to gain access and exfiltrate data from there.

By giving them that visibility, it was very easy for Darktrace to help them quickly understand what was happening in the cloud. This wasn't a complicated problem to fix, but without any visibility they would never have been able to detect it as it unfolded.

Conclusion

As organizations increasingly rely on cloud services and SaaS applications to streamline business practices, the familiar paradigm of the network perimeter has dissolved, leaving a porous and ever-changing digital estate in its wake.

While the benefits of cloud computing will ensure that migrations continue apace, the unique security challenges presented by the cloud will not only require a more agile mindset, but also self-learning technologies that can move at the speed of cloud deployments and spot subtle deviations indicative of a threat, while providing complete, real-time visibility across the digital business.

Darktrace's world leadership in the field of artificial intelligence for cyber security makes it the most effective and proven solution to detect unprecedented threats and anomalous cyber-incidents in the cloud. Whether faced with an insider threat, an attacker targeting sensitive data in test containers, or a significant misconfiguration that could be exploited in the future, Darktrace's cyber AI platform helps eliminate blind spots and protect your data, wherever it resides.

> "Darktrace detects and responds to threats that other tools miss."
>
> IDC

Learn more

darktrace.com | @darktrace | LinkedIn