

The Oracle Cloud Breach Exposed a Harsh Truth: Public Cloud Infrastructure Is a Growing AI Attack Surface For Cybercriminals



Executive Summary

In early 2025, a threat actor known as Rose87168 claimed responsibility for a major breach of Oracle Cloud Infrastructure, allegedly compromising over 6 million records across 140,000 tenants. While Oracle has not officially confirmed the breach, researchers believe the attacker exploited either a zero-day vulnerability or a misconfiguration in the OAuth2 authentication layer, granting unauthorized access to highly sensitive information such as single sign on (SSO) credentials, LDAP passwords, OAuth keys, and tenant-specific data.



This incident is more than a one-off security event—it is a wake-up call for every organization relying on public cloud platforms to host their mission-critical or high-value data. It highlights a systemic risk in the shared architecture of public cloud platforms. Even when organizations follow best practices—strong IAM policies, least privilege access, and continuous monitoring—they remain exposed if the cloud provider makes a single architectural or operational mistake.

Zero-day vulnerabilities are immune to a cloud providers diligence. Worse still, cross-tenant attacks—where a breach in one tenant cascades into others—can occur when the underlying cloud control plane is compromised. While rare, the Oracle incident illustrates that AI-enhanced threat actors are increasingly capable of identifying and exploiting such vulnerabilities at scale.

This brief examines how and why in our AI world, the public cloud has become the primary attack vector for cybercriminals, how cross-tenant breaches can occur even with strong internal controls, and why mission-critical or high-value data should be isolated from public cloud environments to mitigate these emerging risks.

The AI Cybercrime Playbook: Why Public Cloud is the Ideal Target

The rise of artificial intelligence (AI) in cybercrime has dramatically increased the frequency, sophistication, and success of attacks on public cloud environments. These platforms—while offering scale and convenience—also present the highest-value, centralized targets for threat actors. AI tools now enable attackers to rapidly scan for vulnerabilities (like zero-day weaknesses), craft highly targeted phishing lures, automate intrusion attempts, and attack cloud operations to facilitate cross-tenant contamination. The result: public cloud infrastructure is more exposed than ever.



AI cybercriminals are increasingly zeroing in on public cloud platforms because they combine four attractive qualities:

Factor	Why it attracts Attackers
Massive Attack Surface	Cloud platforms are always-on and internet-facing. Misconfigured APIs, open ports, and exposed metadata services provide ample entry points to explore and exploit. AI tools can scan, analyze, and launch attacks on thousands of cloud environments simultaneously.
Concentration of High Value Data	Public clouds host billions of high value records across legal, finance, healthcare, government, and critical infrastructure. Customers are willing to pay the greatest and quickest for high value data.
Multi-tenant Infrastructure	Public clouds share the same hardware, virtualization platforms, and management consoles across multiple tenants. Shared infrastructure means attackers can exploit misconfigurations, API weaknesses, or zero-day hypervisor bugs to move laterally across accounts —especially in the presence of a control plane flaw.
High ROI for Attackers	Once inside a cloud provider’s infrastructure, attackers gain access to massive amounts of data across hundreds or thousands of victims simultaneously - maximizing payout potential for ransomware, data theft, or credential resale. One successful attack can sustain a cybercriminal operation for months.

How AI is Attacking the Public Cloud

AI has dramatically accelerated how cybercriminals discover, exploit vulnerabilities and avoid detection in the public cloud. The most dangerous vulnerabilities are zero-day vulnerabilities. This is a software or hardware flaw that is unknown to the vendor at the time it is exploited. Because there is no patch or defense available, it allows attackers to:

- Exploit systems before detection.
- Target control planes, APIs, or hypervisors to achieve privileged access.
- Cross-tenant access, data manipulation, encryption, or deletion in a cloud setting.

In the Oracle breach, it’s believed that a zero-day in OAuth2 allowed unauthenticated cross-service access—allegedly compromising millions of records across 140,000 tenants.

AI-Driven Techniques for Breaching the Cloud:

Whether deployed with **Proxmox VE**, **Scale Computing**, **VMware**, or **Hyper-V**, Nexsan Unity™ delivers the same secure, high-performance storage experience:

- **Automated Code Analysis:** Scanning source code for previously unknown flaws that have no patch. AI learns patterns that indicate weakness in logic, input handling, and memory management and applies this to scan new codebases at machine speed.
- **AI Generative Fuzzing:** “Fuzzing” is the process of sending malformed or unexpected inputs to software to cause a crash or unexpected behavior. This process can find bugs that may become zero-days. AI massively increases the efficiency and successful rate of finding exploitable bugs.
- **Automated Reconnaissance:** Two main targets 1. Scanning for open S3 buckets, misconfigured access keys, or weak IAM roles both at the cloud provider and tenants. 2. Scanning for known vulnerabilities that the cloud team or a tenant has not patched to fix.
- **Monitoring Patch Releases and Code Diffs:** By analyzing patches and “before/after” code changes, AI can reverse-engineer vulnerabilities—and weaponize exploits before patches are fully deployed (“N-day attacks”). This makes “N-day” attacks (on newly disclosed flaws) more dangerous—AI can weaponize them within hours.
- **Sophisticated Phishing:** AI crafts convincing login prompts and mimics executive communication to steal credentials.
- **Adaptive Evasion:** AI-driven malware changes signature and behavior in real time to avoid detection.

The Core Problem:

Cloud Assumes Trust in the Provider

Public cloud security rests on the shared responsibility model:

- You secure your workloads.
- The provider secures the infrastructure.

But when the provider’s infrastructure is the attack vector—as suspected in Oracle’s breach—the model collapses.

No matter how well an organization locks down IAM, encryption, and network security, it remains exposed if the cloud provider’s identity layer, control plane, or hypervisor is breached.

This isn’t theoretical. Oracle’s breach joins a growing list of major cloud-related incidents—highlighting that trust in a provider’s invisible layers can become a single point of catastrophic failure.



A Strategic Divide:

What Belongs in the Cloud—and What Doesn’t

The best way to reduce corporate risk toward cybercriminals is to decide what data to place in the public cloud and what data to hold in your on-premise data center. Please note that even data in your private data center is of course subject to a cyberattack, however your ability to control the attack vectors are 100% in your hands.

Public cloud platforms are well-suited for:

- Hosting websites
- Development environments
- Collaboration platforms
- Short-term analytics projects
- Non-sensitive or publicly available data

These workloads benefit from cloud scalability and cost-efficiency. But they don’t require bulletproof data integrity, immutability, or legal-grade audit trails.

High-Value Data, However, Requires More:

High-Value Data Attributes	Why Public Cloud Is Insufficient
Regulatory or Legal Requirements	Public cloud providers may not offer certified WORM, encryption, blockchain logging
Sensitive IP or PII	Shared infrastructure increases risk of unauthorized exposure
Long-Term Retention & Disposition	Public cloud lacks cryptographic deletion and air-gapped backup and may lack retention enforcement
Chain-of-Custody Requirements	Control-plane vulnerabilities break the integrity of digital evidence. No ability to tell if a file has been accessed but not changed.

Best Practices: Protecting High-Value Data Outside the Public Cloud

Organizations need to follow all best practices for protecting all their data and company from cybercriminals. These include end point detection and response, malware software, firewalls and keeping software patches updated. In addition to these, organizations managing critical, sensitive, or regulated data (High Value Data) should adopt the following strategy:

- Move High-Value Data to Air-Gapped, Immutable Storage**
 Use on-premise or hybrid solutions like Nexsan Assureon, which enforces file immutability, air-gapping, and blockchain-based audit logs.
- Enforce Retention with 3rd-Party Time Verification**
 Use a storage solution that integrates, accepts and enforces file level retention rules at the storage layer from your document management system. Prevent archive spoofing by using independent 3rd party time authorities for data timestamp validation and disposition controls.
- Use Zero Trust Architecture**
 Public clouds rarely implement **true network segmentation** at the storage layer. Assureon offers a zero-executable attack surface isolated from your corporate network.
- Employ Single-File Encryption**
 Instead of encrypting entire volumes, use unique keys per file for maximum privacy and to enable cryptographic deletion (critical for GDPR and data disposition mandates).
- Immutability, Versioning, and Shortcut Restoration**
 When file versioning is added to immutable storage, the ransomware encrypted file is not able to overwrite the original file. This it allows for a complete recovery from a ransomware attack by restoring the original version. In Assureon, this can be done by instant restoration of virtual shortcuts to the original files.

Assureon High Value Data

Purpose-built, high-security solution, designed to safeguard digital evidence in an AI-driven world.



Final Word: Oracle Is a Case Study—Not an Exception

The Oracle Cloud breach is a stark reminder that the weakest link in the cloud may not be you—it may be the provider. However, the breach occurred this demonstrates that even large, mature cloud providers are vulnerable to zero-day vulnerabilities and control-pane flaws that can result in cross-tenant data exposure – particularly dangerous in a multi-tenant environment where attackers can leap from one client’s environment to another. When not if this happens, their mistake becomes your crisis.

The cloud is still a powerful tool—but only for the right workloads. Its architecture, if not supplemented with immutability, air-gapping, and fine-grained access control, leaves organizations vulnerable to ransomware, data theft, and compliance failures. It’s time for enterprises to re-evaluate what data truly belongs in the cloud—and what data needs fortress-level protection. When it comes to protecting your organization’s most critical digital assets, the safest place is outside the blast radius.

If your organization is storing legal, regulated, evidentiary, or irreplaceable high-value data in public cloud platforms, eliminate the shared risk you may not be able to control—or contain by moving this High Value Data to an Assureon secure data storage appliance.

ABOUT NEXSAN

Nexsan® is a global leader in enabling customers to securely store, protect, and manage data. Established in 1999, Nexsan has earned a reputation for delivering the most highly reliable, secure, and cost-effective storage while always remaining agile to continuously deliver purpose-built storage and data management solutions that meet complex and ever-changing IT, business, and budgetary requirements. Nexsan’s patented technology is ideal for a variety of use cases including backup and recovery, content delivery and streaming, scientific lab data, virtualization, evidentiary data, digital video surveillance, regulatory compliance, and healthcare records. For further information, please visit www.nexsan.com.