Assignment 5.1: Analyzing Cyber Attacks Using the MITRE ATT&CK Matrix

Joshua A. Wortz

University of San Diego

CSOL-580

Dr Hollandsworth

03/23/2026

Introduction

The 2023 cyber attack against MGM Resorts International represents a watershed moment in modern cyber warfare, illustrating the devastating effectiveness of identity-centric intrusion over traditional technical exploitation. While the cybersecurity industry has historically focused on hardening the network perimeter through firewalls and intrusion prevention systems, this incident proved that the most significant vulnerability in any enterprise remains the human element. Orchestrated by the threat group Scattered Spider (UNC3944) in collaboration with the ALPHV/BlackCat ransomware-as-a-service (RaaS) operation, the breach did more than just disrupt operations; it fundamentally challenged the industry's reliance on legacy authentication and support protocols. The resulting systemic operational paralysis across dozens of properties and a total financial impact exceeding $100 million underscores the reality that even the most well-resourced organizations are vulnerable to sophisticated social engineering.

To effectively analyze and mitigate such multifaceted threats, cybersecurity professionals must move beyond static indicators of compromise (IoCs) and toward a behavioral understanding of the adversary. The MITRE ATT&CK™ framework serves as the essential bridge in this transition. It provides a globally accessible, granular taxonomy of adversary tactics and techniques based on real-world observations, allowing analysts to move from a reactive "alert-based" posture to a proactive "intelligence-led" strategy. By mapping the MGM attack to this matrix, defenders can deconstruct the adversary's playbook, identifying not just *what* happened, but the underlying operational patterns that allowed the threat to persist and escalate. This structured analysis is critical for engineering resilient detections that remain effective even as attackers rotate their infrastructure or modify their malware.

Background

MGM Resorts International is a titan in the global hospitality and gaming world. They manage a massive, complex operation that includes everything from thousands of physical slot machines to millions of sensitive customer records and cloud-based identity systems. In early September 2023, this entire infrastructure was hit by a highly coordinated campaign from a group known as Scattered Spider. This group is different from your typical hackers; they are often younger, Western-based, and absolute experts at high-speed social engineering and "living off the land"—using a company's own legitimate tools to stay under the radar. Because they speak fluent, unaccented English and deeply understand how Western corporate IT help desks work, they are incredibly effective at tricking people over the phone.

The campaign was built on meticulous reconnaissance. During the first week of September, the threat actors utilized open-source intelligence (OSINT) to map the organizational hierarchy of MGM. By leveraging professional networking sites like LinkedIn, they identified the names, specific roles, and reporting structures of highly privileged IT system administrators (Bhosale, 2025). On September 8, 2023, the attackers executed the decisive initial breach. This was not achieved through an unpatched vulnerability or a sophisticated zero-day exploit, but through a ten-minute vishing call to the MGM IT help desk (Thompson, 2023). By impersonating a legitimate employee and leveraging the information gathered during reconnaissance, the attacker successfully manipulated the help desk operator into resetting the target's password and Multi-Factor Authentication (MFA) device.

Following this initial entry, the attack's velocity increased exponentially. By September 10, the attackers had moved from a single compromised account to achieving "Super Administrator" levels within Okta and "Global Administrator" status in Microsoft Azure (Serov, 2026). Recognizing the severity of the intrusion, MGM's security teams took the drastic containment

measure of shutting down internal networks and identity synchronization servers. This defensive action, while necessary to sever the attackers' access, triggered the immediate failure of digital room keys, slot machines, and reservation systems worldwide. In retaliation for these containment efforts, the threat actors deployed ALPHV/BlackCat ransomware on September 11, specifically targeting the virtualization layer to make recovery more difficult (Comps, n.d.).

Attack Analysis using MITRE ATT&CK Matrix

The power of the MITRE ATT&CK framework lies in its ability to categorize disparate events into a cohesive adversarial narrative. The MGM breach demonstrates a mastery of the identity-based attack lifecycle, where the adversary treats identity as the new perimeter.

Reconnaissance (TA0043)

The attack began with the Tactic of Reconnaissance, specifically the Technique of Gathering Victim Org Information: Identify Roles (T1591.004). In an era of public professional data, Scattered Spider demonstrated that an attacker does not need to probe a firewall to understand a target’s defenses. By scouring social media and corporate bios, they identified the "gatekeepers", those with the authority to reset credentials. This preparation was not merely about gathering names; it was about understanding the "personae" of the victims to ensure the subsequent social engineering attempt carried the weight of authority and urgency (Bhosale, 2025).

Initial Access (TA0001) and Defense Evasion (TA0005)

The pivot from reconnaissance to Initial Access was achieved through Phishing: Voice Phishing (T1566.004) and direct Impersonation (T1656). This stage highlights a critical gap in many defense strategies: the "human bypass." By convincing a help desk agent to reset a device,

the attacker essentially invited themselves in through the front door using legitimate credentials. This also served as a form of Defense Evasion, as the login appeared to be a standard, authorized user performing a password recovery, an event that rarely triggers a high-severity alert in a busy SOC (Thompson, 2023).

Persistence (TA0003) and Privilege Escalation (TA0004)

Once entry was secured, the adversary's priority shifted to preventing eviction. They utilized the Tactic of Persistence through Domain or Tenant Policy Modification: Trust Modification (T1484.002). By configuring a rogue "inbound federation" Identity Provider (IdP) within MGM's Okta tenant, the attackers created an invisible backdoor (Ropek, 2024). This technique is particularly dangerous because it allows the attacker to act as their own certificate authority; they could authenticate as any user within the environment by simply "asserting" that user's identity from their own rogue server. This bypassed the need for further password resets or MFA prompts, granting them a persistent, high-privilege beachhead that was decoupled from MGM's standard security controls.

Credential Access (TA0006)

To move from the cloud-based identity provider to the on-premise infrastructure, the attackers required local credentials. They employed a combination of OS Credential Dumping (T1003) and Network Sniffing (T1040). By lurking on Okta Agent servers, which serve as the bridge between the cloud and the local Active Directory, they were able to capture clear-text passwords during synchronization (The Silverfort Team, 2023). This gave them the "keys to the kingdom," allowing for seamless movement across the hybrid environment.

Lateral Movement (TA0008) and Exfiltration (TA0010)

With administrative credentials in hand, the group moved laterally using Remote Services (T1021) and Lateral Tool Transfer (T1570). They utilized legitimate remote monitoring and management (RMM) tools to evade detection, a classic LotL strategy. Before the final encryption phase, they moved 6 terabytes of sensitive data to external cloud storage using Exfiltration Over Web Service (T1567). This data exfiltration facilitated a "double extortion" model, ensuring that even if MGM could restore from backups, the threat of leaking 37 million customer records would maintain the adversary's leverage (Ramezanian, 2023).

Impact (TA0040)

The final Tactic was Impact, executed via Data Encrypted for Impact (T1486) and Inhibit System Recovery (T1490). The choice of target, the VMware ESXi hypervisors, was strategic. By encrypting the virtualization layer, attackers crashed thousands of guest virtual machines simultaneously, maximizing the operational blast radius (Comps, n.d.). By also deleting Volume Shadow Copy (VSS) snapshots and backups, they attempted to force MGM into a position where payment was the only viable path to recovery.

Defense and Mitigation

Analyzing the MGM incident through the MITRE ATT&CK framework reveals that the organization’s defenses were largely focused on technical indicators rather than behavioral patterns. While MGM likely had sophisticated firewalls and EDR solutions, these tools are often blind to a "Global Administrator" performing tasks that appear authorized.

Effectiveness of Existing Measures

MGM’s existing defenses were primarily reactive. The decision to shut down the network was an effective, albeit painful, containment strategy that likely prevented even further data loss

or the spread of encryption to isolated backup tiers. However, the initial reliance on "knowledge-based" verification at the help desk proved to be a fatal flaw. In the modern threat landscape, biographical data (SSNs, birthdays, and employee IDs) is frequently available on the dark web or via OSINT, making it an unreliable form of authentication.

Proposed Mitigations and Improvements

To prevent a recurrence of such an event, organizations must adopt a "Zero Trust" architecture that treats every identity as potentially compromised.

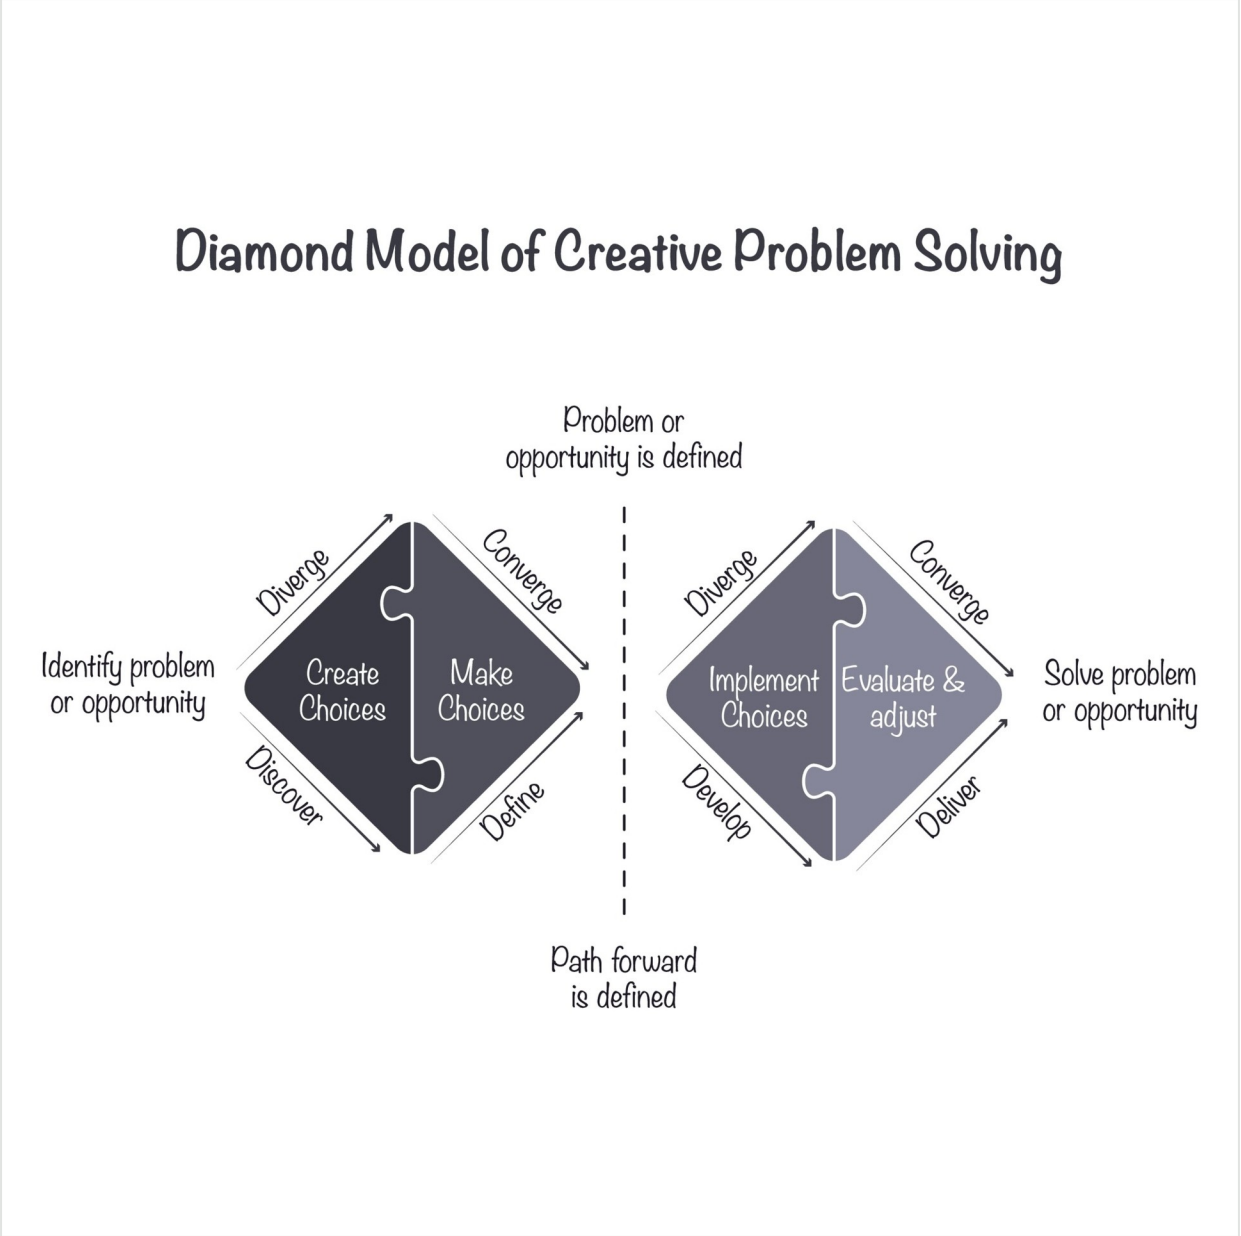
1. **Identity Governance and Help Desk Hardening:** To mitigate **Voice Phishing (T1566.004)**, help desks must implement cryptographically verified identity verification. This includes the use of "out-of-band" verification, where the support agent calls a pre-registered, trusted number on file, or requires a video-based identity check. Furthermore, any MFA reset for a privileged account should require a "dual-authorization" workflow, where a second administrator must approve the change before it takes effect (Bilbrey, 2025).
2. **Phishing-Resistant MFA (AAL 3):** The transition to FIDO2-compliant hardware security keys (such as YubiKeys) is no longer a luxury but a requirement for administrative accounts (1Kosmos, 2024). Unlike SMS or push-based MFA, which can be subverted via vishing or MFA fatigue attacks, hardware keys require physical possession and are bound to the specific origin of the website, making them immune to credential-sharing or proxy-based phishing.
3. **Monitoring Identity Trust and Federation:** Security teams must implement aggressive monitoring for **Trust Modification (T1484.002)**. Any change to "Inbound Federation"

settings or the addition of a new Identity Provider is a high-confidence indicator of a persistence attempt and should trigger an immediate incident response isolation.

4. **Just-in-Time (JIT) Privileged Access:** Organizations must eliminate "standing privileges." Using a JIT model, an administrator has zero elevated rights by default. When a task is required, they must request access, which is granted temporarily and automatically revoked upon completion. This significantly reduces the "window of opportunity" for an attacker who compromises an admin credential (Serov, 2026).
5. **Virtualization Layer Isolation:** To mitigate **Data Encrypted for Impact (T1486)**, hypervisors must be isolated from the general corporate domain. Administrative access to vCenter or ESXi hosts should require a separate, highly secured "jump box" that is entirely disconnected from the systems that standard employees use for email or web browsing (Comps, n.d.).

Conclusion

The 2023 MGM Resorts attack serves as a stark reminder that the modern cybersecurity battle is fought primarily in the realm of identity and trust. The adversary’s ability to pivot from a simple phone call to the total encryption of a global enterprise demonstrates an operational maturity that traditional signature-based defenses cannot counter. The impact was profound: a $100 million loss, the exposure of 37 million customer records, and a permanent shift in how the hospitality industry perceives risk.



Shutterstock

Integrating frameworks like MITRE ATT&CK into a comprehensive cybersecurity strategy is no longer optional; it is a foundational requirement for organizational resilience. While the Cyber Kill Chain provides a high-level timeline of an attack and the Diamond Model provides the relational context of the adversary, MITRE ATT&CK provides the granular technical playbook needed to engineer precise, behavior-based detections. By synthesizing these

frameworks, organizations can build a layered defense that recognizes the "human element" of risk while technically hardening the identity and virtualization layers. Ultimately, the lesson of MGM is that while technical exploits may open doors, it is the manipulation of identity and trust that brings the house down. Defensive strategies must evolve to meet this reality, prioritizing verifying every identity and monitoring every administrative action as if an adversary were already inside the network.

References

1Kosmos. (2024, October 10). *Unmasking the MGM Resorts cyber attack: Why identity-based authentication is the future*.

Abnormal AI. (2025, April 14). *Scattered Spider explained: Threat tactics and defense strategies for IT teams*.

Angelier, V. (2025, June 17). *Applying the Cyber Kill Chain to the SolarWinds attack*. https://doi.org/10.13140/RG.2.2.27607.94885

Bhosale, P. S. (2025, December 18). *Scattered Spider cyberattacks: Anatomy, impact, and defense strategies*. ProArch.

Bilbrey, T. (2025, September 10). *SCYTHE Labs threat intel: Scattered Spider (UNC3944) your help desk became a security risk*. SCYTHE.

Braithwaite, S. (2023, October 24). *ALPHV: Hackers reveal details of MGM cyber attack*. University of Hawai‘i–West O‘ahu.

Caltagirone, S., Pendergast, A., & Betz, C. (2013). *The Diamond Model of intrusion analysis*. Center for Cyber Intelligence Analysis and Threat Research.

Comps, J. (n.d.). *ATT&CK v17 and the increasingly targeted ESXi attack surface* [Presentation]. ATT&CKcon 6.0, MITRE ATT&CK.

Gomez, F. (2025, July 15). *MGM Resorts cyberattack: Data breach, class-action lawsuit & what businesses can learn*. Inszone Insurance.

Lockheed Martin. (2015). *Gaining the advantage: Applying Cyber Kill Chain® methodology to network defense* [White paper].

Naik, N., Jenkins, P., Grace, P., & Song, J. (2022). Comparing Attack Models for IT Systems: Lockheed Martin's Cyber Kill Chain, MITRE ATT&CK Framework and Diamond Model. *IEEE International Systems Engineering Symposium*.

Ramezanian, R. (2023, October 18). *MGM Resorts cyberattack–from cloud to casino floor*. Skyhigh Security.

Rapid7 Labs. (2025, July 3). *Scattered Spider: Rapid7 insights, observations, and recommendations*.

Ropek, A. (2024, March 5). *The Scattered Spider attack: Safeguarding your Okta infrastructure*. Panther.

Schrader, D. (2025, August 18). *An overview of the MGM cyber attack*. Netwrix.

Serov, V. (2026, March 22). *Breach lessons: The 2023 MGM breach*. The Defender's Initiative.

The Silverfort Team. (2023, September 21). *MGM breach takeaway: On-prem has become attackers' gateway to the cloud*. Silverfort.

Thompson, A. (2023, September 22). *The MGM Resorts attack: Initial analysis*. CyberArk.

Thompson, A. (2025, July 3). *Scattered Spider unmasked: How an identity-focused APT is redefining cyber threats*. CyberArk.