Living-Off-the-Tenant: MDM/EMM Infrastructure as Enterprise Wiper Attack Surface

Systemic Risk in Cloud Device Management Platforms

Unofficial AI-assisted Research

Cloud Security Alliance AI Safety Initiative

2026-03-19

Executive Summary

Enterprise device management underwent a fundamental architectural transformation over the past decade. Organizations migrated from on-premises management infrastructure to cloud-hosted Mobile Device Management (MDM) and Enterprise Mobility Management (EMM) platforms that centralize control of every enrolled corporate and personal endpoint in a single administrative tenant. That transformation delivered real operational benefits — unified policy enforcement, over-the-air configuration, remote troubleshooting, and global enrollment visibility — but it created an attack surface that the security community is only beginning to fully characterize.

The defining incident of this threat category emerged on March 11, 2026, when the group publicly claiming responsibility as Handala — tracked by Microsoft as Storm-0842 and assessed by CrowdStrike as Banished Kitten, with assessed ties to Iran's Ministry of Intelligence and Security — executed a global destructive attack against Stryker Corporation by accessing the company's Microsoft Intune tenant and issuing mass remote wipe commands to what Handala claimed were more than 200,000 enrolled corporate and BYOD devices spanning 79 countries — a figure that has not been independently verified, with other reporting characterizing the scope as thousands of devices [1][2][22][23]. The attackers are assessed to have obtained valid administrative credentials — though the specific theft mechanism has not been publicly confirmed — and accessed a legitimate management interface, using it as intended at destructive scale. No novel exploit was used. No vulnerability in Microsoft Intune was required. The entire endpoint estate was disabled in a matter of minutes, patient-critical healthcare services were disrupted, and Stryker issued a public statement to customers acknowledging ongoing operational disruption [3][24].

The Stryker attack is the most visible data point in a pattern that has been developing for years. Security researchers at SpecterOps published the Maestro framework in October 2024, demonstrating how Microsoft Intune's Graph API can be abused for lateral movement and code execution across device fleets [4]. Google Cloud Threat Intelligence and IBM X-Force separately documented techniques for abusing Intune administrative permissions within Microsoft Entra ID environments [5][6]. Ukraine's critical infrastructure was struck in 2025 by PathWiper, a destructive payload delivered through a legitimate endpoint administration console rather than through conventional malware propagation [7]. VMware Workspace ONE was subject to active exploitation of a critical server-side template injection vulnerability in 2022 that enabled unauthenticated remote code execution against management

infrastructure serving enterprise device fleets [8]. The pattern is consistent: the device management plane is an asymmetric attack target, offering adversaries administrative leverage over the entire endpoint estate through a single point of compromise.

This whitepaper examines the structural dimensions of MDM/EMM infrastructure risk, catalogues the attack vectors across major platforms, reviews the documented incident record, and provides an architecture and governance framework for organizations seeking to treat device management platforms with the security rigor their administrative reach demands. The central argument is not that these platforms are poorly designed, but that organizational security programs have systematically under-invested in protecting management plane credentials, under-deployed monitoring on management plane activity, and failed to apply separation-of-duties controls commensurate with the destructive capabilities that management plane access enables.

Introduction: When the Control Plane Becomes the Weapon

Enterprise security programs have spent the better part of two decades building layered defenses designed to stop attackers who enter through identifiable vectors — phishing payloads, exploited vulnerabilities, compromised credentials — and then attempt to escalate, move laterally, and execute malicious code. The architecture of most enterprise security stacks reflects this threat model: endpoint detection and response (EDR) to catch malicious process execution, network detection tools to observe lateral movement traffic, security information and event management (SIEM) platforms correlating signals from across the environment. This model is well-calibrated against attackers who bring their own weapons.

It is poorly calibrated against attackers who use yours.

The emergence of cloud-hosted device management platforms created an administrative control layer that sits above conventional endpoint security controls. An administrator with valid credentials to a Microsoft Intune tenant, a Jamf Pro instance, or a VMware Workspace ONE environment can push software to thousands of endpoints simultaneously, modify security configurations across an entire fleet, wipe devices remotely, and deploy arbitrary scripts — all through a browser-based console whose activity generates no suspicious process on the endpoint, no malicious binary on disk, and no anomalous network traffic that conventional security tools are calibrated to detect. This capability was designed for IT administrators managing distributed workforces. It is equally useful to an adversary who has obtained those credentials.

The shift toward cloud-native unified endpoint management has intensified this dynamic. Organizations that previously distributed management authority across multiple on-premises servers — with geographic and architectural boundaries limiting the blast radius of any single compromise — now frequently operate a single cloud tenant that governs every enrolled device globally. A single set of administrative credentials, once compromised, provides the key to the entire castle. The question is no longer whether adversaries have identified this attack surface; the Stryker attack and the body of adversarial research documented in subsequent sections confirm they have. The question is whether enterprise security programs have adapted their controls accordingly.

This whitepaper is organized to provide security architects and enterprise risk leaders with the technical foundation and governance roadmap needed to make that adaptation. It begins with the historical context of wiper attacks and their evolution toward management plane abuse, proceeds through an analysis of documented attack vectors across the major MDM/EMM platforms, reviews the incident record, analyzes the detection and visibility gaps that allow these attacks to succeed, and concludes with a structured recommendations framework grounded in CSA's security control architecture.

The Evolution of Wiper Attacks Toward Management Plane Abuse

Destructive cyberattacks designed to render systems inoperable — commonly classified as wiper attacks for their tendency to overwrite or corrupt data at the storage layer — have been a fixture of nation-state offensive operations since the Shamoon campaign targeted Saudi Aramco in 2012. Understanding the evolutionary arc of this attack category is essential context for appreciating why MDM/EMM platforms have become the preferred delivery mechanism for the current generation of destructive operations.

The classical wiper model, exemplified by NotPetya in June 2017, relied on propagation through network protocols and the execution of destructive payloads on compromised hosts. NotPetya leveraged EternalBlue (CVE-2017-0144), a stolen NSA exploit targeting the SMBv1 protocol, to spread across networks at remarkable speed, ultimately causing an estimated $10 billion in global economic damage [9]. The technique was powerful but detectable: EDR tools observed the destructive process executing, incident responders could analyze the payload, and the attack signature could be developed and blocked — eventually. The technical debt of deploying novel malware included the exposure of that malware to reverse engineering and detection rule development.

HermeticWiper, deployed against Ukrainian banking and critical infrastructure in February 2022 at the opening of Russia's full-scale invasion, represented a refinement of this model. Rather than exploiting a network vulnerability for propagation, HermeticWiper leveraged a signed driver from EaseUS Partition Master software — a legitimate commercial product — to corrupt partition tables and render systems unbootable [10]. The use of a legitimately signed driver appeared designed to evade signature-based detection [10]. The driver was trusted; what it was instructed to do was not. This reflected a broader adversarial preference for legitimate tools and signed components that anticipated the increasing maturity of endpoint behavioral detection.

WhisperGate, deployed against Ukrainian government systems in January 2022, presented as ransomware but functioned as a wiper — masking its destructive intent behind familiar ransomware iconography while overwriting the master boot record and corrupting files beyond recovery [11]. Analysts assessed that the deceptive framing was likely intended to slow forensic response by creating an initial impression that recovery via ransom payment was possible [11].

The 2022 attacks on Albanian government infrastructure by Void Manticore (Handala) introduced the next evolutionary step. That campaign combined initial access via an unpatched SharePoint vulnerability (CVE-2019-0604) with 14 months of persistent network presence before the destructive phase, deploying a wiper designated ZEROCLEAR alongside the ROADSWEEP ransomware family and the CHIMNEYSWEEP backdoor [12]. The extended pre-positioning phase — which enabled the attacker to maximize reach and completeness of destruction before triggering defensive responses — demonstrated the shift toward patient, carefully planned destructive operations rather than opportunistic attacks.

PathWiper, documented in 2025 against Ukrainian critical infrastructure, introduced what may be the most significant tactical evolution in this lineage: deployment via an administrative console that the target organization used for legitimate endpoint management [7]. The adversary did not execute malicious code in the traditional sense; they issued administrative commands through a trusted management interface. The boundary between "attacker action" and "authorized administrative activity" was deliberately obscured, presenting defenders with an event that appeared, to conventional security controls, indistinguishable from routine IT operations.

The Stryker attack in March 2026 removed any remaining ambiguity about the direction of this evolution. In that incident, purpose-built wiper malware was unnecessary. The attacker needed only valid credentials to a legitimate administrative platform and the knowledge that "remote wipe" was an available command. The attack surface is not a vulnerability; it is an architectural feature weaponized through credential compromise.

MDM/EMM Architecture: The Single-Plane Attack Surface

The Centralization Model and Its Security Implications

Modern enterprise device management platforms are architected around the principle of centralization. Whether the platform is Microsoft Intune, Jamf Pro, VMware Workspace ONE (formerly AirWatch), IBM MaaS360, SOTI MobiControl, Kandji, Mosyle, or any of the other major MDM providers, the fundamental architecture is similar: a cloud-hosted management console provides administrators with a unified interface to enroll, configure, manage, and monitor every endpoint in the organization's fleet. Devices check in with the management server on a scheduled basis, receive policy updates, and comply with commands issued from the console — including, critically, commands to wipe local data and restore to factory state.

This architecture is well-suited to the operational realities of the modern distributed enterprise. IT teams can enforce security baselines on every enrolled device — requiring screen lock, enabling full-disk encryption, pushing certificate authorities, configuring VPN profiles — without touching individual machines. Remote wipe capability allows organizations to protect sensitive data on lost or stolen devices. Application management enables controlled deployment of software without local administrator rights. These are legitimate and valuable capabilities.

The security implication of this architecture is that the administrative console — and the credentials that provide access to it — represents a single point of leverage over the entire endpoint estate. An adversary who obtains administrative access to an organization's MDM tenant inherits the same capabilities as the organization's own IT administrators: the ability to push arbitrary scripts, deploy software, modify security configurations, and issue wipe commands to every enrolled device simultaneously. The scale of this capability grows proportionally with the size of the enrolled fleet. For enterprises with tens or hundreds of thousands of enrolled endpoints, MDM administrative access represents one of the highest-impact single points of failure in the endpoint security architecture, particularly for organizations with fully enrolled global fleets.

The BYOD Dimension

The Bring Your Own Device (BYOD) enrollment model, which has become widely adopted across enterprises seeking to reduce hardware costs and accommodate remote workforces, significantly expands the destructive potential of MDM compromise. When employees enroll personal devices —

smartphones, tablets, personal laptops — under corporate MDM management to gain access to corporate email, applications, and resources, those personal devices become subject to the same administrative commands as corporate-issued equipment.

The Stryker attack made the BYOD dimension concrete. Employees across multiple countries reported losing personal devices to the wipe operation — not because they were specifically targeted, but because their devices were enrolled in the same Intune tenant as corporate equipment, and the attacker's mass wipe command did not discriminate [2]. The destruction of personal data on employee devices adds a dimension of harm that extends well beyond operational disruption: employees lose irreplaceable personal content, and the organization faces legal exposure under data protection frameworks in jurisdictions where enrolled personal devices contain personal data subject to those frameworks.

Administrative Role Architecture and Blast Radius

The blast radius of an MDM tenant compromise is mediated by the specific administrative role held by the compromised account. In Microsoft Intune, the Intune Service Administrator role provides full management capabilities across all enrolled devices. The Global Administrator role in Microsoft Entra ID (formerly Azure Active Directory) provides this access and more. The DeviceManagementConfiguration.ReadWrite.All permission in Microsoft Graph API grants programmatic equivalents of console-level administrative capability.

Organizations frequently assign these roles persistently — meaning accounts hold the role continuously rather than only when performing a specific task. Persistent role assignment dramatically expands the window of exposure: if a credential is phished or stolen, the attacker has standing access to destructive capabilities for as long as the credential remains valid. The absence of just-in-time role elevation, multi-factor authentication resistant to phishing (such as FIDO2 hardware tokens), and privileged access workstations are organizational decisions that, in aggregate, determine whether an attacker who obtains an administrator credential can actually act on it.

Documented Attack Vectors

Credential Compromise of Administrative Accounts

The primary documented attack path for MDM-enabled destructive operations is the compromise of administrative credentials for the management console, whether through credential phishing, credential theft from compromised endpoints, or purchase from initial access brokers who specialize in selling pre-

compromised administrative accounts.

The Stryker attack followed this pattern, with investigators assessing initial access through credential compromise of Entra administrative accounts [1][2]. The specific credential theft mechanism has not been publicly confirmed, but the operational model — phishing targeting administrators with Intune or Entra administrative roles, followed by direct access to the management console — requires no platform vulnerability and is available to any threat actor with phishing capability. Phishing-resistant MFA based on hardware security keys or certificate-based authentication is the primary architectural control against this vector, and its absence from Stryker's administrative accounts is assessed as the most critical remediable gap — one that, if addressed, would have required the attacker to overcome significantly higher authentication barriers.

SpecterOps researchers documented an additional credential theft pathway relevant to Entra-managed environments: the extraction of Primary Refresh Token (PRT) cookies from compromised endpoints held by administrators [4]. PRT tokens, which provide long-lived authentication to Entra services without requiring password re-entry, can be extracted from a compromised admin workstation using tools such as ROADtools or through memory acquisition techniques. An attacker in possession of a PRT token can authenticate as the token holder to any Entra-connected service — including Intune — and can do so in a manner that bypasses Conditional Access policies that evaluate device compliance, because the token was issued to a compliant device even if it is now being used from an adversary-controlled system [4].

Lateral Movement via Management APIs

For adversaries who have obtained initial access to an enterprise environment at a non-administrative level, MDM platforms represent a high-value lateral movement target. An account that reaches Intune administrative privilege can transition from a typical post-exploitation footprint — code execution on a single compromised endpoint — to global administrative access over the entire enrolled device fleet.

Google Cloud Threat Intelligence documented the abuse of Intune's integration with Microsoft Entra ID for lateral movement in 2024, describing how the DeviceManagementConfiguration.ReadWrite.All permission, if granted to a compromised service principal or user account, enables device reconnaissance, configuration modification, and PowerShell script deployment to enrolled devices [5]. IBM X-Force published a complementary analysis documenting detection approaches for these movements, noting that the activity generates log entries in Microsoft Entra audit logs but does not surface in EDR telemetry, making it invisible to organizations relying exclusively on endpoint-based detection [6].

The SpecterOps Maestro tool, released as an open-source post-exploitation framework in October 2024, explicitly packages these techniques for practical use in authorized red team engagements [4]. Maestro provides a command-and-control interface for device reconnaissance, PowerShell script deployment, Windows application installation, and user enumeration through the Intune administrative API. Its release demonstrated that these techniques had progressed from theoretical to operationally accessible — any red team, and by extension any motivated adversary, could implement them.

Enrollment Token and Certificate Abuse

MDM platforms establish trust relationships with enrolled devices through enrollment tokens and certificates. These trust mechanisms, if improperly controlled, create additional attack pathways that do not require compromise of administrative credentials.

Organizations that permit open or insufficiently restricted device enrollment — without requiring specific device identity attestation — allow adversaries who have obtained an enrollment token to register rogue devices as trusted members of the enterprise device fleet. Rogue devices enrolled in this manner inherit the trust relationship established by enrollment: they receive enterprise certificates, can access applications and resources restricted to managed devices, and may be visible to administrators as legitimate corporate endpoints. In environments where MDM enrollment grants access to sensitive internal resources — such as email, VPN, or internal applications — rogue device enrollment can serve as an authentication bypass mechanism.

Certificate abuse presents a related risk. MDM platforms commonly deploy enterprise root certificates to enrolled devices to enable traffic inspection, internal resource access, and application identity verification. If an adversary compromises the certificate management infrastructure associated with the MDM platform, they may obtain the ability to impersonate internal services, perform man-in-the-middle interception of traffic from enrolled devices, or issue certificates on behalf of legitimate device identities. NCC Group research on iOS MDM enrollment documented the specific trust model established by user enrollment certificates and its security implications in multi-tenant environments [13].

Apple's Managed Device Attestation (MDA) framework, introduced to address weaknesses in the Simple Certificate Enrollment Protocol's inability to cryptographically verify device identity at enrollment, provides a more robust alternative: enrolled devices attest their identity using a hardware-bound private key, proving to the MDM server that the enrolling device is a genuine Apple product with a specific hardware identity [14]. Organizations managing Apple device fleets should evaluate MDA adoption as a control against rogue device enrollment attacks.

Platform Vulnerabilities

Beyond credential and enrollment abuse, the MDM platforms themselves present vulnerability surfaces that attackers have exploited to gain unauthorized access to management infrastructure.

VMware Workspace ONE Access was the subject of active widespread exploitation beginning in April 2022 following the public disclosure and proof-of-concept publication for CVE-2022-22954, a critical server-side template injection vulnerability with a CVSS score of 9.8 [8]. The vulnerability existed in the Apache FreeMarker template engine integrated into the Workspace ONE Access web interface, allowing unauthenticated attackers with network access to the management console to inject template directives that executed arbitrary shell commands as the VMware service account. Rapid7 observed active exploitation within hours of public proof-of-concept release, with attackers dropping cryptocurrency miners on vulnerable systems — but the same code path that allowed miner deployment could be used to modify MDM policies or issue destructive commands to enrolled devices [8]. CISA issued an advisory confirming active exploitation and urging immediate patching [15].

VMware Workspace ONE was also subject to CVE-2022-31656, an authentication bypass affecting Workspace ONE Access, Identity Manager, and vRealize Automation that allowed unauthenticated network-adjacent attackers to obtain administrative access to the management console without valid credentials [16]. Separately, CVE-2022-31685 documented a hard-coded valid authentication token for a super-administrator service account embedded in the product's database on every installation — allowing any attacker with network access to authenticate as a super-administrator without credentials. CVE-2022-31686 and CVE-2022-31687 documented additional authentication bypass and access control vulnerabilities in Workspace ONE Assist that further expanded this unauthenticated access surface [16].

IBM MaaS360 was found to contain hard-coded credentials in the Android MDM agent (CVE-2024-35118) in versions 8.60 and prior, where credential values embedded in the application binary could be recovered from physical device access and used to perform unauthorized MDM operations [17]. Jamf Pro has been subject to server-side request forgery vulnerabilities (CVE-2021-39303, CVE-2021-40809) that allowed authenticated users to leverage the Jamf infrastructure to reach internal network resources not directly accessible from the attacker's position [18], as well as a deserialization remote code execution vulnerability in versions prior to 10.15.1 that enabled denial of service and arbitrary code execution on Jamf Pro server infrastructure (a standalone CVE identifier for this vulnerability has not appeared in available public research).

The pattern across these platform vulnerabilities is consistent: authentication bypasses, hard-coded credentials, and code execution vulnerabilities in management plane software grant attackers the same standing as legitimate administrators, with all the destructive leverage that entails.

Cloud Infrastructure Management APIs

The MDM/EMM attack surface extends beyond dedicated device management platforms to encompass cloud provider-native infrastructure management services, which perform functionally equivalent roles for cloud-hosted servers and virtual machines.

AWS Systems Manager (SSM) provides a fleet management capability for EC2 instances and on-premises servers with the SSM agent installed, enabling administrators to execute commands via Run Command, establish interactive sessions via Session Manager, and manage configuration state via State Manager — all without requiring direct network access or SSH keys. Mitiga Security documented that the SSM agent can be abused by adversaries who obtain sufficient AWS IAM credentials: by re-registering the agent against an attacker-controlled AWS account, the adversary can issue commands to the victim's instances while the connection appears to originate from AWS infrastructure [19]. Additionally, EC2 instances assigned the AmazonSSMRoleForInstancesQuickSetup managed policy — a default in Quick Setup configurations — hold permissions that allow instances to execute commands against other instances in the same account, creating a lateral movement pathway through the management plane that does not traverse conventional network controls.

Azure Arc extends Azure management capabilities to on-premises, multi-cloud, and edge environments, connecting servers outside Azure to Azure Resource Manager and enabling policy enforcement, extension deployment, and command execution through the Azure management plane. As with AWS SSM, the security boundary protecting Arc-connected resources is the Azure IAM control plane rather than conventional network segmentation; an adversary who compromises sufficient Azure permissions can issue commands to any Arc-connected machine regardless of its network location [20].

Incident Record: MDM-Enabled Destructive Operations

The following table summarizes the key documented incidents in which device management infrastructure served as either the delivery mechanism or a significant attack vector in destructive operations.

Incident	Date	Actor	Mechanism	Scale	Source
Stryker Corporation	March 2026	Handala/Void Manticore	Microsoft Intune admin	Attacker-claimed	[1][2] [3][22]

Incident	Date	Actor	Mechanism	Scale	Source
		(assessed MOIS ties)	credential compromise; mass remote wipe	200,000+ devices, 79 countries (not independently verified; other reporting characterizes scope as thousands)	[23] [24]
PathWiper (Ukraine critical infrastructure)	2025	Assessed Russian state actors	Destructive payload deployed via administrative console	Undisclosed critical infrastructure organization	[7]
Albanian Government (ZEROCLEAR)	July–September 2022	Void Manticore / HomeLand Justice (MOIS)	Pre-positioned access; ZEROCLEAR wiper + ROADSWEEP ransomware; no MDM-specific delivery confirmed	Government ministries, parliament, telecoms, airline	[12][21]
Workspace ONE Exploitation	April 2022	Multiple; nation-state suspected in targeted cases	CVE-2022-22954 server-side template injection; unauthenticated RCE on management infrastructure	Widespread; confirmed exploitation at multiple organizations	[8][15]

Incident	Date	Actor	Mechanism	Scale	Source
HermeticWiper (Ukraine)	February 2022	Sandworm / Russian GRU	Signed driver abuse; MBR overwrite; pre-positioned via management infrastructure in some cases	Banks, energy, IT sector Ukraine	[10]

The Albanian government incident is notable for what it is not: a confirmed MDM-plane attack. MOIS actors in that campaign used SharePoint exploitation for initial access and deployed malware through conventional methods. It is included in this table because the same threat actor (Void Manticore) subsequently executed the Stryker attack using the MDM-plane technique, and the Albania campaign demonstrates the group's willingness to invest in extended pre-positioning before destructive operations — a methodology that transfers directly to cloud tenant compromise scenarios.

Detection and Visibility Gaps

The operational success of MDM-plane attacks is significantly enabled by the detection gap they exploit. Traditional enterprise security architectures are overwhelmingly oriented toward endpoint and network visibility. EDR tools observe process execution, file system activity, and network connections on managed endpoints. Network detection tools analyze traffic flows for known-malicious patterns. SIEM platforms correlate these endpoint and network signals to detect attack chains.

Management plane activity does not generate the signals these tools are calibrated to observe. When an Intune administrator issues a remote wipe command, no process is spawned on the victim device that EDR can flag. No unusual binary is written to disk. No lateral movement traffic is generated that network detection can catch. The action occurs in the cloud management plane — in Microsoft's infrastructure — and its effects appear on the endpoint as a legitimate device management operation. The distinction between an authorized IT administrator issuing that command and a nation-state actor doing so is largely invisible to endpoint-based EDR and network detection tools absent dedicated cloud audit log analysis.

Visibility into management plane activity requires cloud audit log collection and analysis. Microsoft Entra ID and Intune both generate detailed audit logs: Entra ID captures sign-in events, role assignments, and policy changes; Intune captures device management operations including the DeviceRemoteAction category, which records wipe and retire commands. These logs are available but are not automatically forwarded to external security monitoring; organizations must explicitly configure diagnostic settings to export logs to a Log Analytics workspace or SIEM, and then build detection rules calibrated to the specific signals that indicate abuse.

IBM X-Force identified the following key Intune audit log signals as indicators of potential management plane attack activity: mass remote wipe commands (high volume of DeviceAction_RemoteWipe events in a short time window); bulk device configuration policy modifications from an account that does not routinely perform such operations; PowerShell script deployments to large numbers of devices simultaneously; sign-in events to Intune from IP addresses or geographic locations inconsistent with the user's baseline; and privilege escalation within Entra (assignment of Intune Service Administrator or Global Administrator roles to accounts not previously holding them) [6].

The absence of cloud management plane monitoring — which appears common in organizations primarily invested in endpoint and network detection — is among the most consequential gaps this threat category exploits. Organizations that have invested heavily in EDR, SIEM, and network detection have built a perimeter that is irrelevant to attacks conducted entirely within the management plane.

Organizational Risk Factors

Several organizational patterns significantly elevate exposure to management plane wiper attacks. Security architects and enterprise risk leaders should assess their organizations against each of these dimensions.

Centralized single-tenant MDM architecture. Organizations that have consolidated all endpoint management into a single cloud MDM tenant, with no architectural boundaries between geographic regions or business units, present the maximum blast radius for a management plane compromise. An attacker with tenant-level administrative access can issue commands to every enrolled device globally, with no architectural control limiting the scope of destruction. Multi-tenant architectures, while operationally more complex, bound the blast radius of any single credential compromise.

Persistent privileged role assignment. The most dangerous organizational pattern is the persistent assignment of high-privilege roles — Global Administrator, Intune Service Administrator — to accounts that are accessed daily through conventional workstations and web browsers. These accounts are

exposed to the full range of credential theft attack vectors, including phishing, credential harvesting malware, and browser session hijacking, without the protections that privileged access workstations and just-in-time role elevation provide.

Software-based or absent MFA on administrative accounts. Phishing-resistant MFA based on hardware security keys (FIDO2) or certificate-based authentication is the primary defensive control against the credential phishing attacks that provide initial access in management plane attacks. Organizations relying on SMS-based MFA or TOTP authenticator apps for administrative accounts remain vulnerable to adversary-in-the-middle phishing frameworks (such as Evilginx2 or Modlishka) that intercept and relay authentication tokens in real time, bypassing these weaker MFA forms.

Unrestricted BYOD enrollment. The enrollment of personal employee devices under corporate MDM management without strong attestation controls or enrollment restrictions expands both the destructive scope of a wipe attack and the legal exposure under data protection frameworks. Organizations without enrollment restrictions — which limit enrollment to known corporate devices with specific hardware identities — cannot control which personal devices are subject to management plane commands.

Absence of separation of duties for destructive operations. The MDM remote wipe capability is a destructive action with no equivalent in any other enterprise administrative domain. Organizations apply separation of duties controls to financial transactions, production deployments, and privileged database operations. Organizations rarely apply equivalent controls to MDM remote wipe — requiring, for instance, approval from two named administrators or a formal change request before mass wipe operations can be issued. This absence of process control allows a single compromised account to execute irreversible destruction at scale.

Limited or absent cloud audit log monitoring. As described in the preceding section, management plane attacks are largely invisible to endpoint and network security tools. Organizations that do not collect and actively monitor Entra ID and Intune audit logs — specifically for anomalous bulk device management operations — have no detection capability for this attack class.

Strategic Recommendations

Architecture and Access Control

The most consequential changes organizations can make to reduce MDM/EMM attack surface are architectural and access control improvements that reduce the blast radius of a credential compromise before an attack occurs.

Organizations should implement Privileged Identity Management (PIM) in Microsoft Entra ID (or equivalent just-in-time access controls in other platforms) to eliminate persistent assignment of Intune Service Administrator and Global Administrator roles. Under a PIM model, administrators request role elevation when performing specific tasks, receive time-limited access, and are logged for those elevations — reducing the window during which a compromised account carries destructive capability from indefinite to the duration of the approved elevation period, typically one to eight hours. Privileged access workstations dedicated exclusively to MDM administration should be required for Intune console access and should be distinct from the general-purpose workstations used for daily work.

All MDM administrative accounts should be protected by phishing-resistant MFA. Hardware FIDO2 security keys or certificate-based authentication should be mandatory for any account holding Intune Service Administrator, Global Administrator, or equivalent roles. Conditional Access policies should be configured to block console access from non-compliant or non-corporate devices, and named location restrictions should be applied where operationally feasible to limit administrative console access to known network egress points.

Organizations should implement approval workflows for high-impact MDM operations — particularly remote wipe commands affecting more than a defined threshold of devices simultaneously. In Microsoft Intune, bulk wipe operations can be flagged through audit log monitoring and workflow automation; while Intune does not natively require multi-person approval for wipe commands, SIEM-based alerting can trigger a secondary approval gate before operations are confirmed. Organizations with the ability to implement custom workflows through Microsoft Graph API event triggers should build these controls explicitly.

Monitoring and Detection

Security operations teams should treat Entra ID and Intune audit logs as first-class monitoring data equivalent in priority to EDR telemetry. These logs must be exported in real time to an external SIEM or Log Analytics workspace that is isolated from the Entra tenant — a tenant administrator compromised by an attacker should not be able to delete or modify audit log exports. Detection rules should be implemented for: anomalous Entra sign-in patterns for administrative accounts; bulk device management operations (remote wipe, retire, or compliance policy changes) exceeding defined thresholds; Intune Service Administrator or Global Administrator role assignments to new accounts; PowerShell script deployment to large device populations; and Entra Conditional Access policy modifications.

Baseline behavioral profiles for administrative accounts performing MDM operations — typical activity patterns, geographic locations, time-of-day distributions, and operation types — provide the foundation for anomaly-based detection that can flag the first signs of management plane compromise. The

initiation of mass wipe commands at unusual hours, from atypical locations, or under accounts that have never previously issued bulk device operations, would present a detectable anomaly against a properly established baseline.

Organizations should also implement out-of-band alerting for MDM administrative sessions. If an administrator account authenticates to the Intune console from an IP address or geographic location inconsistent with their baseline, a real-time alert to the security operations center — not dependent on the Entra or Intune infrastructure that may be under attacker control — provides an opportunity to revoke the session before destructive actions are completed.

Enrollment Control and BYOD Architecture

Organizations should implement strong device identity attestation at enrollment to prevent rogue device registration. For Apple devices, Apple Managed Device Attestation (MDA) provides cryptographic proof that an enrolling device is a genuine Apple product with a specific hardware identity [14]. For Windows devices, Autopilot hardware hash enrollment combined with Windows Hello for Business certificate binding provides analogous attestation. Enrollment should be restricted to devices with confirmed hardware identities registered in the organization's asset inventory.

The BYOD enrollment model should be reevaluated in light of the Stryker incident. Organizations that enroll personal devices under the same MDM tenant as corporate equipment should implement a segregated enrollment profile that explicitly limits the device management commands applicable to BYOD enrollments — specifically, ensuring that mass wipe operations cannot be issued to BYOD-enrolled personal devices without individual per-device confirmation. App-based management containerization, which limits MDM management to a corporate work container rather than the entire device, provides an architectural alternative that preserves corporate data protection without extending organizational control — and organizational risk — to the personal portions of employee devices.

Cloud Infrastructure Management Controls

For AWS environments, organizations should review the permissions granted to SSM-related IAM roles on EC2 instances and restrict the default AmazonSSMRoleForInstancesQuickSetup managed policy where it grants cross-instance command execution capability beyond operational requirements. SSM Session Manager sessions should be logged to CloudTrail and CloudWatch Logs, with anomaly detection rules configured for bulk Run Command executions targeting large numbers of instances simultaneously. Least-privilege IAM policies should explicitly scope SSM permissions to the specific instance groups and operation types required by each administrative role.

For Azure Arc-connected environments, Azure Policy and Azure RBAC should be applied to Arc-connected machines with the same rigor applied to native Azure resources. The management plane access that Arc provides should be explicitly inventoried and included in privilege access reviews. Activity logs for Arc-connected machines should be monitored for unusual command execution patterns.

Business Continuity and Recovery

The Stryker attack demonstrated the inadequacy of business continuity planning that assumes a recovery path beginning from an operational IT infrastructure. When the MDM platform itself is the attack vector, the enrolled devices that would normally be used to coordinate recovery may themselves be unavailable. Organizations should maintain offline, air-gapped recovery documentation — device re-enrollment procedures, out-of-band communication channels, and manual operational procedures for critical workflows — that does not depend on the availability of managed endpoints or cloud management infrastructure.

Healthcare organizations and others operating life-safety-critical services should conduct tabletop exercises specifically addressing the scenario of simultaneous loss of all MDM-enrolled devices and associated cloud services. The Lifenet disruption experienced during the Stryker attack illustrates that the patient safety consequences of MDM-delivered destruction can extend well beyond IT operations into direct harm to individuals who depend on those services.

CSA Resource Alignment

The threat category described in this whitepaper intersects with multiple Cloud Security Alliance frameworks and publications, each providing specific control guidance applicable to different dimensions of the MDM/EMM attack surface.

CSA's **AI Controls Matrix (AICM)** establishes the governance and technical control framework most broadly applicable to cloud administrative environments, including device management infrastructure. AICM control domain IAM-02 (Identity Governance and Administration) addresses privileged access governance requirements that directly apply to MDM administrative roles, including role assignment reviews, separation of duties, and just-in-time provisioning. IAM-08 (Multi-Factor Authentication) specifies phishing-resistant MFA requirements for administrative accounts. LOG-04 through LOG-08 (Audit Logging and Monitoring) establish the cloud audit log collection and monitoring requirements

that the Stryker attack demonstrated to be absent in many enterprise environments. Organizations should conduct an AICM assessment scoped to cloud identity and endpoint management administrative controls as a structured mechanism for identifying their most critical gaps against this threat category.

The **MAESTRO** framework for agentic AI threat modeling [26] is relevant to this whitepaper's threat category in an emerging dimension. Note that "MAESTRO" as used here refers to the CSA AI Agent Threat Modeling Framework; the SpecterOps post-exploitation tool discussed elsewhere in this document is styled as "Maestro" (mixed case) and is an entirely separate product. Enterprise MDM environments are increasingly incorporating AI-assisted automation — agents that manage device enrollment flows, generate remediation scripts, enforce compliance policies, or summarize fleet security posture. As these AI agents are granted administrative permissions to MDM APIs, they inherit the same attack surface described throughout this document: a compromised or manipulated AI agent with Intune administrative privilege is functionally equivalent to a human attacker with the same credential. MAESTRO Layer 5 (Agent Control) guidance on constraining the operational blast radius of autonomous administrative agents — including bounds on what actions agents can take without human confirmation — should be incorporated into the design of any AI-assisted IT operations architecture that touches MDM administrative APIs.

CSA's **Zero Trust guidance** addresses the foundational architectural principles most directly applicable to the credential compromise vectors documented in this whitepaper. Zero Trust's premise — that no entity, whether human or machine, should be granted standing trust based on authentication alone, and that every access request should be evaluated continuously against the full context of the requesting entity's current state — maps precisely to the gap exploited in the Stryker attack. An architecture in which administrative access to destructive capabilities requires not just valid credentials but also verified device compliance, consistent behavioral context, and proximity to established baselines is architecturally resistant to the credential theft and phishing-based access patterns that have enabled management plane attacks.

CSA's **Cloud Controls Matrix (CCM)** [25] provides granular mapping of specific security controls to the gaps documented in this whitepaper. Beyond the IAM and LOG domains already cited, CCM control IVS-01 (Infrastructure and Virtualization Security) and IVS-04 address management plane security and network segmentation requirements. SEF-02 through SEF-07 (Security Incident Management) are directly applicable to organizational preparedness for the management plane attack scenario, including the requirement for offline incident response procedures that do not depend on managed endpoints or cloud services that may be unavailable during the attack. Organizations that have not conducted a CCM assessment with specific attention to cloud identity and endpoint management infrastructure should treat this as a near-term priority given the evolving threat environment. Per CSA's current framework guidance, AICM is a superset of CCM and should serve as the primary framework reference, with CCM providing supplementary control-level detail.

Finally, the **CSA STAR** program (Security Trust Assurance and Risk) provides a cloud service assurance framework relevant to organizations evaluating the security posture of their MDM/EMM vendors. Before selecting or continuing with an MDM platform, organizations should request STAR registry entries or equivalent security assessment documentation covering the vendor's access control architecture, vulnerability disclosure practices, audit log availability, and incident response history. The platform vulnerabilities documented in the VMware Workspace ONE, Jamf Pro, and IBM MaaS360 cases demonstrate that MDM vendor security maturity is directly relevant to organizational risk exposure.

Conclusion

The Stryker attack of March 2026 was not an isolated incident. It is better understood as the predictable result of an adversarial evolution that has proceeded from malware-based wiper deployment to signed-driver abuse to management console weaponization — each step representing adversaries adapting to the increasing maturity of conventional endpoint security and exploiting the organizational investments that enterprises themselves have made in centralizing administrative control. The MDM/EMM platform that a healthcare technology firm deployed to efficiently manage hundreds of thousands of endpoints across dozens of countries was weaponized, without modification, to destroy those endpoints in minutes.

The security community's response to this evolution requires a fundamental reorientation of where protective investment is directed. The endpoint has been the organizing principle of enterprise security architecture for two decades. The management plane — the administrative control layer above the endpoint — has been treated as a supporting IT function rather than a security-critical surface. The accumulated consequence of that misallocation is visible in the detection gap that allowed the Stryker attack to proceed to completion without triggering a meaningful defensive response.

Closing this gap requires coordinated action across architecture, access control, monitoring, and governance dimensions. Phishing-resistant MFA and just-in-time privilege elevation protect against credential compromise. Separation of duties and approval workflows for destructive operations reduce the blast radius when credentials are compromised. Cloud audit log collection and calibrated detection rules provide the visibility layer that makes management plane attacks detectable. BYOD enrollment controls and device attestation limit the population of devices subject to mass wipe commands. And governance frameworks — particularly CSA's AICM, Zero Trust guidance, and CCM — provide the structured control architecture within which these technical improvements can be planned, measured, and assured.

The management plane is the attack surface that adversaries have already found. The question for every enterprise security program is whether their defenses have found it too.

References

[1] Krebs on Security, "Iran-Backed Hackers Claim Wiper Attack on Medtech Firm Stryker," KrebsOnSecurity, March 2026. https://krebsonsecurity.com/2026/03/iran-backed-hackers-claim-wiper-attack-on-medtech-firm-stryker/

[2] TechCrunch, "Stryker says it's restoring systems after pro-Iran hackers wiped thousands of employee devices," TechCrunch, March 17, 2026. https://techcrunch.com/2026/03/17/stryker-says-its-restoring-systems-after-pro-iran-hackers-wiped-thousands-of-employee-devices/

[3] Stryker Corporation, "A Message to Our Customers," Stryker News, March 2026. https://www.stryker.com/us/en/about/news/2026/a-message-to-our-customers-03-2026.html

[4] SpecterOps, "Maestro: Abusing Intune for Lateral Movement Over C2," SpecterOps Blog, October 2024. https://specterops.io/blog/2024/10/31/maestro-abusing-intune-for-lateral-movement-over-c2-2/

[5] Google Cloud Threat Intelligence, "Abusing Intune Permissions for Lateral Movement and Privilege Escalation in Entra ID Environments," Google Cloud Security Blog, 2024. https://cloud.google.com/blog/topics/threat-intelligence/abusing-intune-permissions-entra-id-environments

[6] IBM X-Force, "Getting 'in tune' with an enterprise: Detecting Intune lateral movement," IBM Security Blog, 2024. https://www.ibm.com/think/x-force/detecting-intune-lateral-movement

[7] The Hacker News, "New PathWiper Data Wiper Malware Disrupts Ukrainian Critical Infrastructure in 2025 Attack," The Hacker News, 2025. https://thehackernews.com/2025/06/new-pathwiper-data-wiper-malware.html

[8] Rapid7, "Widespread Exploitation of VMware Workspace ONE Access CVE-2022-22954," Rapid7 Blog, April 2022. https://www.rapid7.com/blog/post/2022/04/29/widespread-exploitation-of-vmware-workspace-one-access-cve-2022-22954/

[9] Wired, "The Untold Story of NotPetya, the Most Devastating Cyberattack in History," Wired, August 2018. https://www.wired.com/story/notpetya-cyberattack-ukraine-russia-code-crashed-the-world/

[10] ESET Research, "HermeticWiper: New data-wiping malware hits Ukraine," ESET Blog, February 2022. https://www.welivesecurity.com/2022/02/24/hermeticwiper-new-data-wiping-malware-hits-ukraine/

[11] Microsoft Security Intelligence, "WhisperGate: Destructive malware targeting Ukraine," Microsoft Security Blog, January 2022. https://www.microsoft.com/en-us/security/blog/2022/01/15/destructive-malware-targeting-ukrainian-organizations/

[12] CISA and FBI, "Iranian State Actors Conduct Cyber Operations Against the Government of Albania," Joint Cybersecurity Advisory AA22-264A, September 2022. https://www.cisa.gov/news-events/cybersecurity-advisories/aa22-264a

[13] NCC Group, "iOS User Enrollment and Trusted Certificates," NCC Group Research, 2021. https://www.nccgroup.com/research/ios-user-enrollment-and-trusted-certificates/

[14] Smallstep, "Unlocking Zero Trust Security: Apple MDA for MDM Device Enrollment," Smallstep Blog. [Note: URL returns 308 redirect to homepage as of 2026-03-19; consult Apple's official Managed Device Attestation documentation in the Apple Platform Deployment Guide for authoritative guidance.] https://smallstep.com/blog/zero-trust-with-apple-mda/

[15] CISA, "Alert (AA22-138B): Threat Actors Chaining Unpatched VMware Vulnerabilities for Full System Control," CISA Advisory, May 2022. https://www.cisa.gov/news-events/cybersecurity-advisories/aa22-138b

[16] GreyNoise, "VMware Workspace ONE Vulnerabilities: CVE-2022-31656 and CVE-2022-31659," GreyNoise Blog, August 2022. https://www.greynoise.io/blog/vmware-workspace-one-vulnerabilities-cve-2022-31656-and-cve-2022-31659

[17] IBM, "Security Bulletin: IBM MaaS360 MDM Android Application CVE-2024-35118," IBM Support, 2024. https://www.ibm.com/support/pages/security-bulletin-ibm-maas360-has-identified-vulnerability-maas360-mdm-android-application-cve-2024-35118-0

[18] Assetnote, "Discovering Full-Read SSRF in Jamf (CVE-2021-39303, CVE-2021-40809)," Assetnote Research, 2021. https://www.assetnote.io/resources/research/discovering-full-read-ssrf-in-jamf-cve-2021-39303-cve-2021-40809

[19] Mitiga, "Mitiga Security Advisory: Abusing the SSM Agent as a Remote Access Trojan," Mitiga Security Blog, 2023. https://www.mitiga.io/blog/mitiga-security-advisory-abusing-the-ssm-agent-as-a-remote-access-trojan

[20] XYBYTES, "Abusing Azure Arc: From Service Principal Exposed to Reverse Shell," XYBYTES Research Blog. https://xybytes.com/azure/Abusing-Azure-Arc/

[21] Google Cloud Threat Intelligence (Mandiant), "ROADSWEEP Ransomware — Likely Iranian Threat Actor Conducts Politically Motivated Disruptive Activity Against Albanian Government Organizations," Google Cloud Security Blog, September 2022. https://cloud.google.com/blog/topics/threat-

intelligence/likely-iranian-threat-actor-conducts-politically-motivated-disruptive-activity-against/

[22] Palo Alto Networks Unit 42, "Insights: Increased Risk of Wiper Attacks," Unit 42 Blog, March 2026. https://unit42.paloaltonetworks.com/handala-hack-wiper-attacks/

[23] SecureWorld, "Iran-Linked Hacktivist Group Weaponizes Microsoft Intune in Destructive Wiper Attack on Stryker," SecureWorld News, March 2026. https://www.secureworld.io/industry-news/iran-linked-hacktivist-group-weaponizes-microsoft-intune-in-destructive-wiper-attack-on-stryker

[24] Layer 8 Security, "When Device Management Becomes the Attack Surface: Lessons from the Stryker Cyberattack," Layer 8 Security Blog, March 2026. https://layer8security.com/when-device-management-becomes-the-attack-surface-lessons-from-the-stryker-cyberattack/

[25] Cloud Security Alliance, "CCM: Cloud Controls Matrix v4.0," CSA Research, 2021. https://cloudsecurityalliance.org/research/cloud-controls-matrix/

[26] Cloud Security Alliance, "MAESTRO: AI Agent Threat Modeling Framework," CSA AI Safety Initiative, 2025. https://cloudsecurityalliance.org/ai-safety-initiative

