jamf

A Practical Guide to Apple Endpoint Telemetry

What you see is what you get

Visibility starts with telemetry

Have you ever walked into a room that's completely dark? Without light, you aren't able to clearly make out shapes, like where the wall ends and the furniture's edge begins. Endpoint security and by extension compliance are much like that. Without telemetry data, you can't protect what you can't see.

While this may sound obvious, the inability to capture current endpoint health data from managed Apple fleets represents one of the biggest challenges to modern organizations today. Sure, IT can easily glean how many macOS, iOS and iPadOS devices are enrolled, and security teams can generate a report detailing how many devices upgraded to the latest major version of macOS during a migration project six months ago.

But can your IT and security team validate the state of compliance across the Mac and mobile device fleet at this very moment? **Put another way, how confident are your teams in the data they're working with to verify:**

- Which processes are running?
- Is system integrity impacted?
- Who made those changes?
- Are endpoints compliant?

Notice the subtle but important distinction between the words *identify* and *verify*. Because, between those two, lies the difference between enrolled devices: those that have records in your Mobile Device Management (MDM) solution that only update during scheduled check-ins (identify). And compliant endpoints: active monitoring gathers and analyzes rich telemetry data to provide real-time insight into endpoint health statuses (verify).

The former is based on the most recent check-in, which, depending on schedule and resources, is stale and potentially introduces unknown risk to organizations. The latter relies on current health data to validate endpoint compliance in real-time, providing timestamped proof of device compliance to auditors.

For organizations of every size, the key to closing the gap between identifying and verifying is telemetry. **In this guide, we not only highlight the critical role telemetry plays in modern management and security strategies, but we also cover:**

- What it is and isn't
- Its role in driving fleet insight
- How visibility:
  - Unlocks comprehensive IT operations
  - Enables desired business outcomes
  - Delivers value through greater alignment
  - Strengthens organizational security postures

To set expectations, while a dedicated security operations center, a specialized team of analysts or a background in threat intelligence certainly help, they aren't required. The only thing necessary is to understand what telemetry provides, so teams build proactive workflows based on real-time, data-driven decision-making, not relying on reactive processes based on stale, out-of-date information.

What telemetry is

Looking at the etymology of telemetry, we find it's a compound of two words derived from Greek: tele, meaning "far off" and metron, meaning "measure." Roughly, this translates to "remote measurement."

When used in the context of endpoint security, telemetry is the stream of health and activity measurements a device reports about itself: what ran, what changed, who changed it, when and what it connected to.

As it pertains to macOS endpoints, telemetry commonly include the following:

- **Access tracking** records who logged in, when and from where, which resources they requested and which permissions were granted or denied, such as requests for privilege escalation.
- **Process execution** follows app launches, background tasks, related child processes and their interactions with kernel and user space.
- **System events** log configuration drift, update successes and failures, system file modifications and changes to hardware or software state.
- **Network activity** captures connection details, like encryption strength and network types, IP information, ports, protocols and the apps transmitting data.
- **Diagnostic logs** collect system operation and performance data, along with crash reports, code states and memory dumps.

When looking at telemetry, it's common to dismiss individual artifacts as being forgettable or easy to miss. However, much like a single puzzle piece doesn't do much on its own, the full picture is realized when all the pieces are combined. Together, they provide IT and security teams a detailed analysis that verifies endpoint health at that specific point in time.

What teams do with those measurements, whether that's threat hunting, patch management or validating compliance, is a separate step.

What telemetry is not

As with most security controls in a defense-in-depth strategy, admins must understand where the gaps exist so defenses can be layered to close them as effectively as possible. Likewise, you should be just as familiar with telemetry's limitations as its benefits to use it effectively.

Telemetry is not:

- Surveillance of a user's personal activity or browsing habits.
- Replacement for tooling that carries out management tasks or security remediation workflows.
- Blanket attestation of compliance throughout an endpoint's lifecycle.

Telemetry data is simply information that's been gathered about a device (or your entire device fleet) sitting in logs. Without further action, telemetry is of little use in and of itself. Telemetry delivers value when organizations use it: first by reviewing it or streaming it to their security information and event management (SIEM) solution for analysis and applying that context as data-driven decision-making to management, security and compliance workflows.

This distinction matters.

Collecting telemetry data without applying it to MDM and endpoint security solutions is like the single block of aluminum Apple uses to create each MacBook. The aluminum, by itself, doesn't help users accomplish work-related tasks. However, by applying the aluminum to their manufacturing processes, Apple is able to produce a laptop that enables users to fulfill their role responsibilities or creative endeavors.

Telemetry, in this analogy, is the raw material (aluminum). Only by applying it are organizations able to build their compliance program with it, and it begins with holistic visibility and the patterns that emerge from combining the individual data points.

How telemetry works: from query to insight

Now that we've explained what telemetry is, let's shift perspective slightly to go over how telemetry data sitting on endpoints transforms decision-making relating to device and organizational security postures.

To best understand how telemetry flows, let's take a closer look at each of the four stages of the pipeline:

1 Collection

The macOS platform generates a stream of system-level events that capture everything occurring on a device in real-time. This health data is gathered at the kernel level and from hardware sensors, recording performance metrics monitored by Apple's Endpoint Security API.

2 Aggregation

Once the raw data is collected, the next step is to parse this information through the Unified Logging System (ULS). During this phase, data captured is formatted to meet structure and storage requirements on each device.

3 Processing

After the data is logged, native macOS tools, like XProtect and MRT, alongside third-party security agents tap into the log stream in real time. They read and analyze the data stream, prioritizing high-signal events for the final stage.

4 Reporting

At this stage, rich telemetry data that has been structured and processed is transmitted via first-party standard protocols, such as Apple's diagnostic servers for reviewing crash reports and usage analytics, as well as third-party data analysis tools.

After telemetry data has flowed through the pipeline, it's common for IT and security teams to review device logs as they assess endpoint health to identify potential issues. Due to the number of logs generated by a single device, then multiplied by the number of devices across the fleet, it is highly recommended for organizations to feed this never-ending stream of data to their SIEM solution where it is collated based on severity levels or organizational compliance needs. In short, instead of sorting through thousands of entries looking for the proverbial needle in the haystack, the technology performs the heavy lifting by sorting the events of greatest importance to be reviewed first.

Depending on the size and resources available to your organization, this marks the point where cross-functional teams rely on the details of the telemetry data to resolve issues flagged.

For example, organizations with dedicated IT and security teams increasingly come together as one department, relying on their strengths to enforce a compliance policy on endpoints found to be missing a critical patch. Security verifies the finding as a true positive, while IT ensures the non-compliant devices are grouped accordingly within the MDM. Working in tandem marks the inflection point, where these teams ensure that automatic remediation is performed by the endpoint security solution and compliance is validated by collecting updated telemetry data.

Just to clarify, grasping how telemetry provides insight to teams doesn't require knowing every protocol, data type or API call involved in the underlying process by heart. It's about knowing how the data travels along the pipeline and what it tells IT and security teams about the current state of their devices, highlighting the most compliant path forward.

Visibility: the outcome telemetry enables

Telemetry and visibility differ in one important way: telemetry is the input and visibility is the output. Often, these terms are used interchangeably because they are similar but it's important to not conflate them because they are different enough to warrant distinction.

An easy way to identify the difference is: cause and effect. Telemetry illuminates the "cause" of an issue, like a misconfigured device. Visibility identifies what the misconfiguration is exactly and which devices are affected by it, also known as the "effect."

From an operational perspective, visibility means teams have a current and accurate assessment of what's happening across their fleet. It also prescribes an actionable path toward mitigating the issue(s) surfaced. It answers questions like:

- Is baseline behavior being observed?
- What devices are healthy?
- Which are at risk?
- Is there an active threat?

Telemetry is the collected data that makes visibility possible. It is made up of the information that is gathered from all managed endpoints across your fleet, providing details relating to device health statuses in real-time, such as:

- Missing system or security patches
- Unsanctioned apps installed
- Configurations that have changed from baseline

While telemetry enables visibility, it does not do so automatically. Organizations can constantly monitor their endpoints, gathering copious amounts of telemetry yet still lack visibility into fleet health and compliance status.

Telemetry = input

Visibility = output

How is it possible to collect telemetry but lack visibility?

Telemetry ≠ visibility.

The former enables and feeds the latter but only when the data is reviewed, contextualized and turned into a decision or action. Without data review and decision-making, telemetry is like a book that is not read — it just sits in storage, full of incredible knowledge, never benefiting the user.

Instead of thinking of visibility as an on/off switch, it's more accurate to view it as a spectrum. At one end, visibility exists while at the opposite end, it doesn't. With that in mind, consider your organization and where it falls on the visibility scale.

While you do that, let's look at one example company set across two distinct ends of the visibility spectrum:

No visibility: An organization requires that devices be enrolled in MDM and uses a default check-in time of every 12 hours to update device records, this scenario provides limited insight into device health and no real-time visibility. Because of the check-in delay, any data you collect will be at least 12 hours old, at best. Should devices miss a scheduled check-in, the collection will reoccur at the next scheduled window, making that information even more stale.

While MDM is capable of a lot, it wasn't designed to operate as a replacement for an endpoint security solution. This means risk indicators, like the presence of malware or monitoring of suspicious behaviors, are not being collected. Furthermore, containment workflows cannot be executed automatically, impacting security postures due to the lack of visibility.

Visibility enabled: An organization requires that devices be enrolled in MDM and uses a default check-in time of every 12 hours to update device records and integrated endpoint security that actively monitors endpoints. Together they provide comprehensive insight into device health and real-time visibility. Despite the MDM's scheduled check-in delay, active monitoring means the security agent is constantly transmitting telemetry data through the pipeline, enabling accurate endpoint health in real-time for each device across the fleet.

With endpoint security integrated alongside an MDM solution, risky behaviors, apps and code are always being monitored — and when detected, they are prevented automatically. Additionally, workflows like quarantining compromised endpoints and mitigating risk may also be performed, alongside providing teams with current health data to enforce compliance through policy-based management or validate compliance to auditors with timestamped evidence.

In short, visibility is the result of using telemetry as a genuine security advantage. It's not about adding more headcount or buying more tools, but rather about turning the telemetry your macOS, iOS/iPadOS devices are already generating into insight. Using that as the basis for scalable decisions and actions, based on the confidence of knowing the health status of your growing fleet — not what it could be, or was a day or two ago the last time it was checked.

Detection, investigation and response are powered by visibility.

Visibility isn't just valuable because of the insight it provides technical teams relating to compliance and its impact on security postures. Its value extends to everyday IT and Security operations, namely:

Detection

IT uses baselines to define normal operating procedures and performance levels, based on best practices and tailored to the organization. Telemetry, and by extension, visibility provide teams insight into device health, like if configurations align with baselines or have drifted away from them. Visibility also makes abnormal behavior stand out, including anything that has drifted from baseline. This is especially beneficial as fleets grow, by surfacing anomalies sooner rather than later. Or put another way, real-time visibility proactively detects issues that may otherwise lie dormant and go unnoticed until it becomes an incident.

Investigation

Arguably, this is where visibility plays its most impactful role. The moment an anomaly is detected, security teams switch into investigative mode to validate the issue and then answer a whole slew of questions, like:

- What happened?
- When did it occur?
- Why did it happen?
- How deep is the impact?

To put it bluntly, without telemetry, responses to these and other critical questions will be based purely on guesswork, conjecture and hearsay from user statements. This isn't to say those accounts are wrong, only that they place a heavier burden on technical teams to corroborate them. Whereas telemetry is evidence — it's the clear proof that explains what happened across the timeline of events.

Response

The best phrase to describe visibility's influence over this operation is *"where understanding becomes action."* After the issue has been detected and Security has reviewed the evidence, a telemetry-driven incident response workflow provides teams the insight needed to perform tasks in a repeatable order to maximize remediation while limiting exposure:

- Isolate compromised endpoints to contain impact.
- Prioritize threats by highest severity/ criticality levels.
- Mitigate the risk(s) to eradicate lingering threats.
- Remediate endpoints, restoring to baseline performance.

> > >

Responding: without and with visitiliby

Earlier in this e-book, we highlighted the contrast between an organization having visibility and no visibility. The key difference being actively applying telemetry data gathered in the former to gathering it but not putting it to use (or relying on stale data) for the latter. While examples like this show up clearly throughout this guide, nowhere is it more clearly illustrated than while technical teams are working under pressure during incident response.

The following points show the stark contrast between two teams (A & B), tasked with responding to the same incident: a vulnerable, unsanctioned application installed on a company-issued device leads to a ransomware attack.

Team A: Without Visibility

- **Works with partial information:** Can see apps installed but cannot verify vulnerability level.
- **Reacts as symptoms surface:** Device performance is sluggish at first; later, data becomes inaccessible.
- **Non-verifiable data sources:** Information presented from the affected user's perspective and cannot be attested.
- **Resolution requires additional work:** Absent corroborating indicators, more troubleshooting steps to remediate root cause.
- **Does not scale with fleet:** As device fleets grow, workloads increase exponentially as do risk factors and exposure windows.

Team B: With Visibility

- **Comprehensive data fills the gaps:** Sees the same incident from the inside, with cause, sequence and effect laid out clearly.
- **From detection to resolution:** Remediates the incident with a full understanding of the timeline — not guesswork.
- **Current, accurate and validated:** Information attests to endpoint's health status to teams and auditors alike.
- **Review data with context:** Information gathered includes indicators, collected from kernel and sensors paint a full picture.
- **Grow fleets — not workloads:** With scalability in mind, processes and workflows adjust without increasing administrative overhead.

Working smarter — not harder

Telemetry is a driver for automation.

When integrating mobile device management, endpoint security and identity and access management, telemetry-driven workflows can augment device management and security strategies in novel ways that leverage technology to perform the heavy lifting in a number of scenarios, freeing admins to focus their skills on delivering enhanced experiences and better aligning operations with business outcomes.

Consider a policy-based automation where endpoint security flags a vulnerable app and triggers remediation via your MDM solution. This simple workflow maintains the security posture and mitigates risk without disrupting the user, involving IT support or becoming an incident.

Modern organizations have seen how AI/ML technologies can become a force multiplier for teams of all sizes. AI-powered processes can hunt for sophisticated threats across a fleet at incredible speeds. It correlates telemetry data gathered from endpoints, analyzes and cross-references it in real-time against multiple open and closed-source threat intelligence databases to detect and alert teams of unknown threats. And with scalable automation options, keeping a human in the loop for incidents that require it is as easy to bake into your workflow as policy-based workflows that handle less severe issues on their own.

By being able to customize solutions to fit your organization's unique needs, incidents are resolved faster and with better information, which minimizes risk and manual effort while optimizing efficiency and maximizing efficacy.

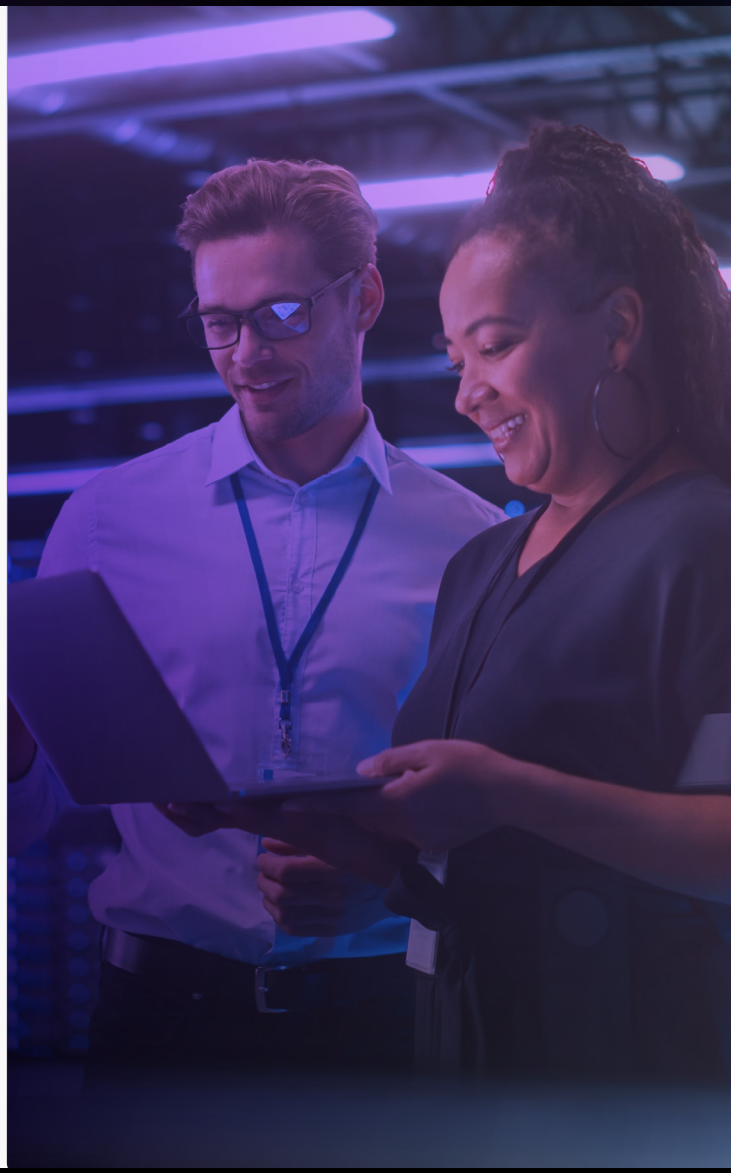

Telemetry in context: compliance, operations and organizational value

In addition to the threat hunting and incident response benefits discussed previously, two other areas deserve equal weight in how organizations think about telemetry.

Compliance

Industries like healthcare, finance and education face increasing scrutiny from regulations globally. Federal and regional mandates require organizations to demonstrate that data stays secure and that the devices touching it stay configured to compliance requirements.

One point bears repeating: claims of compliance must be verified. The default that most auditors defer to is rooted in this belief: if you cannot verify a device was compliant, then it wasn't. Simple as that.

Telemetry is the secret sauce that enables validation. As mentioned earlier on, telemetry in and of itself is nothing more than a collection of data. Without review, contextualization, decisions or actions, it does not enable visibility. The same holds true that without the proper security controls and policies in place first, telemetry does not guarantee compliance.

With context applied and controls in place, telemetry becomes the driver of those aspects that prove (or disprove) that they're working properly at the time of assessment. Outside of formalized audits, evidentiary data (telemetry) matters in other compliance matters: operations, which we'll discuss in the next section and cybersecurity insurance. In that specific case, underwriters often request organizations to demonstrate proof of their security practices before coverage begins or is renewed, or when an organization files a claim.

Operations

The same telemetry data also benefits IT operations, as the detection, investigation and response examples showed earlier. This helps IT and security teams carry out the tasks related to their roles, but also it gives IT leaders and executives a real-time view of organizational health and compliance without requiring a separate process, additional tasks or generating individualized reports.

Telemetry is multifaceted, and by basing reporting on real-time telemetry, it performs double duty: serving as operationally useful for technical teams day to day while also supporting business use cases through credible, evidence-backed documentation of compliance statuses and risk assessments to drive future growth conversations basing decision-making on accurate, current data.

Simply put: telemetry's underlying value doesn't change based on team size or organizational structure. Telemetry does the same job whether it's reviewed by a single IT Manager or a team of Security Analysts. It deserves a broader home. Rather than something "only security teams use," telemetry serves different purposes depending on who's asking the question.

Assessing your telemetry and visibility posture

Understanding telemetry and visibility conceptually is useful. How it helps organizations know where they stand is more useful. A self-assessment of your own telemetry and visibility standing is an excellent place for organizations to start, without requiring purchases of new tools, specialized training or changing the configuration of your existing stack.

Begin by answering the following:

1. Does your organization keep endpoint logging data on-device or stream it to a central location (SIEM solution) for analysis?
2. Could your team confidently determine the current security posture of any device in your fleet?
3. Are they able to reconstruct an incident timeline using evidence, or must they solely rely on the affected user's description of events?
4. Can your organization produce timestamped compliance verification for an audit, if requested today?
5. When compliance violations are detected, do they remediate automatically or require manual follow-up?

The answers to the questions above tell a lot about the security posture of your organization, and more to the point, **your team's ability to proactively mitigate risk versus reactively respond** to an incident that has already occurred. Below are a few common gaps to be mindful of:

- Rich telemetry stored on endpoints or shared with your existing stack introduces gaps in visibility that teams might miss.
- Forwarding some events (Windows) while siloing others (Apple) in cross-platform environments still exposes the organization to risk.
- Collecting telemetry is only part of a solution, it must be reviewed regularly and acted upon for it to unlock its true value.

Conversely, mature security stacks share the following consistent traits:

- Telemetry flows continuously and is built into their workflows, not just reviewed when someone remembers to check.
- Data collected is analyzed against multiple sources, correlated and prioritized based on severity. Manual sorting results in critical details possibly being missed.
- Teams with successful compliance programs aren't the largest or have bigger budgets. They're the ones that make use of the telemetry being generated before making decisions or acting.
- Visibility built on top of your existing data is almost always a faster path forward than visibility being built from new tooling that isn't integrated.

Conclusion

It bears repeating: telemetry, without review, is just data collected in a file.

The constant stream of timestamped facts about what a device is doing, its configuration state or vulnerability level only becomes useful when organizations use this information to act on the findings before a minor issue silently becomes a costly violation of compliance.

The ability for raw telemetry data to enable data-driven decision-making, which leads to timely action based on real-time health statuses — that is visibility. The accuracy and currency of the data add a layer of confidence to the data that follows it along: from the pipeline to the decision makers at its destination.

The gap between managed and secured is closable. It doesn't always require new purchases, extensive team-level expertise or significant changes to your existing stack. The devices in your fleet are designed to log what they're doing, and they're already telling you everything that's going on.

The real question is whether your organization is set up to listen and act on what it hears.

Ready to see it in action?

Experience Jamf today