The Stalkerware Market Hides in Plain Sight. OSINT Can Map It — but Detection Is Only Half the Problem

Maria Cattini | 24/09/2026 | OSINT

Commercial surveillance software does not always arrive through an exploit, a clandestine intelligence service or a sophisticated intrusion campaign. Sometimes it is sold as a subscription, advertised as parental control or employee monitoring, supported by customer-service teams and installed on an ordinary smartphone.

That overlap between apparently legitimate monitoring software and covert surveillance is precisely what makes stalkerware difficult to investigate. The same technical capabilities can have very different meanings depending on consent, disclosure and use. An application that reports a device's location is not, by that fact alone, stalkerware. An application designed to monitor another person without their knowledge, hide its presence and transmit sensitive information to a remote user is something very different.

A recent open-source investigation by Intel Focus brought the issue back into view. Published on 16 September 2026, it examined how applications presented as monitoring or parental-control tools can enable technology-facilitated abuse. The investigation arrives against a wider body of evidence showing that stalkerware has neither disappeared nor become a purely historical cybersecurity problem.

For OSINT investigators, however, there is another question worth asking: how much of this ecosystem can be understood without compromising a device, accessing private accounts or interacting with victims?

Quite a lot.

The websites, marketing language, corporate structures, app policies, security tests, regulatory records and technical relationships surrounding commercial surveillance products leave public traces. Taken separately, none necessarily proves abuse. Combined carefully, they can reveal how a market presents itself, how products are distributed, what capabilities are documented, where safeguards fail and where the boundary between legitimate monitoring and covert surveillance begins to collapse.

The important part is not finding a suspicious app. It is building an evidence chain that explains what the available information actually establishes.

The first problem is definition

"Stalkerware" is sometimes used loosely to describe almost any invasive application. That creates an immediate methodological problem.

The Coalition Against Stalkerware uses a narrower definition: software made directly available to individuals that enables someone to monitor another user's device without that person's consent and without explicit, persistent notification, in circumstances that may facilitate surveillance, harassment, stalking or abuse.

That distinction matters because monitoring functions are not inherently evidence of malicious use.

Google's current developer policy illustrates the boundary. Monitoring applications intended for parental or enterprise use can be permitted under specific conditions, but they must disclose their monitoring function, display persistent notification while operating, identify themselves clearly and must not present themselves as tools for secret surveillance. Google also states that monitoring applications cannot be used to track a spouse, even with that person's knowledge and permission.

For an OSINT investigation, this gives us something more useful than a label: criteria.

Instead of asking whether an application "looks like spyware", an investigator can ask what the company claims the product is for, whether monitoring is disclosed to the person being monitored, whether the application is designed to remain visible, what information it collects and whether its documented behaviour corresponds with the platform rules under which legitimate monitoring software is supposed to operate.

That shift from labels to observable characteristics is fundamental.

Follow the public footprint, not the suspicion

A useful investigation can begin with the vendor's own public material.

Product pages, terms of service, privacy policies, installation documentation, archived versions of websites, company records, app-store listings and regulatory documents can establish different parts of the picture. Search engines and web archives may reveal how marketing language changes over time. Public technical research can establish capabilities that promotional material leaves vague.

But these sources should not be treated as interchangeable.

A marketing page tells us how a company describes a product. It does not independently prove what the software actually does.

A privacy policy may document categories of collected information. It does not prove that every installation collects every category.

A security researcher may demonstrate a capability in a tested version. That finding should not automatically be projected onto every version of the product.

An enforcement action establishes what a regulator alleged or determined in that particular case. It does not establish that every monitoring vendor behaves in the same way.

This separation is where OSINT becomes investigation rather than accumulation.

The US Federal Trade Commission provides a concrete example. In a case involving Support King and its SpyFone service, the FTC alleged that the company sold applications allowing purchasers to monitor devices secretly, required users to weaken device security protections, and collected information including photographs, text messages, browsing histories, location and physical movements. A 2021 order banned Support King and its CEO from offering surveillance applications and required deletion of data that the FTC said had been secretly collected.

In December 2025, the FTC rejected the CEO's petition to vacate or modify that order.

That is strong evidence about one company and one regulatory case. It is not evidence that every parental-control application is stalkerware.

Maintaining that boundary is essential.

The detection gap is measurable

One of the most useful pieces of evidence comes not from advertising or regulation but from controlled testing.

AV-Comparatives and the Electronic Frontier Foundation tested 13 Android mobile-security products against 17 commercial stalkerware samples in 2025. The test used a non-rooted Samsung Galaxy A36 running Android 15, and researchers installed the applications following the vendors' instructions before evaluating whether security products produced a clear warning.

The results varied substantially.

Detection rates ranged from 53% to 100% across the products tested. One detected all 17 samples, while Google Play Protect detected 53% in that particular test. Several products achieved rates above 90%, while others missed a significant portion of the test set.

These figures need context. They describe a defined 2025 test involving 17 selected applications and particular versions of security products. They are not universal detection rates for all stalkerware, nor do they tell us how the same products would perform against a different sample today.

What they demonstrate is more important: detection cannot be assumed.

The researchers also found evidence of shared infrastructure and rebranding. Some tested applications appeared to be variations of the same underlying product, reusing elements such as payment systems, backend infrastructure, administrative dashboards or APK files. Identical APK hashes were treated as one product to avoid counting duplicates.

For OSINT, that detail changes the investigative unit.

The visible brand may not be the real boundary of the ecosystem.

If multiple products reuse technical or commercial infrastructure, an investigation that searches only by brand name may fragment what is actually a connected network. Domains, certificates, company records, payment relationships, application identifiers, archived pages and publicly documented technical indicators can become more useful than the storefront name.

But infrastructure overlap is evidence of a relationship, not automatically evidence of common ownership. A shared service provider, reseller or white-label platform can produce similar traces. Attribution requires additional corroboration.

The market keeps changing

Kaspersky's research provides another view of scale, although its numbers must also be read according to their methodology.

For 2024–2025, the company reported more than 34,000 users affected by stalkerware in its telemetry and said it had identified 33 previously unseen stalkerware families. Affected users were detected in more than 160 countries. The figures come from aggregated threat statistics collected through the Kaspersky Security Network and therefore represent visibility within that ecosystem rather than a census of global stalkerware use.

That limitation is important.

Security-vendor telemetry can demonstrate that a phenomenon exists at significant scale within the observed population. It cannot tell us how many stalkerware victims exist worldwide.

The real number could be different for several reasons: not every device uses participating security software, products may differ in what they classify as stalkerware, some applications may remain undetected, and technology-facilitated surveillance increasingly includes mechanisms other than dedicated stalkerware.

An OSINT article that simply repeats “34,000 victims” would therefore turn a measurable dataset into a claim the dataset does not support.

The defensible statement is narrower: Kaspersky detected more than 34,000 affected users in its 2024–2025 telemetry.

That difference may look small. Methodologically, it is enormous.

An application is only one layer of the investigation

The temptation in a stalkerware investigation is to concentrate on the APK. For malware researchers that may be necessary, but an OSINT investigation can work differently.

Think of the product as the centre of a public ecosystem.

Around it may exist a company, domains, support documentation, reseller pages, social accounts, advertising material, corporate registrations, privacy policies, payment services, application-store histories, archived websites, technical analyses, court documents and regulatory actions.

Each answers a different question.

Corporate records may help establish who operates a service. Archived pages can show how a product was marketed at a particular time. Platform policies provide criteria against which public product behaviour can be compared. Security testing can demonstrate capabilities or detection failures. Regulatory records can establish documented allegations, findings and enforcement actions.

The investigation becomes stronger when these layers corroborate one another.

It also becomes safer.

There is rarely a legitimate reason for an OSINT researcher to install commercial surveillance software on another person’s device, attempt to access a customer dashboard or obtain victim data simply to prove that a surveillance service exists. Public-source investigation should exhaust public evidence before crossing into intrusive technical testing, and any laboratory testing requires an appropriate legal and ethical framework.

The objective is not to reproduce the surveillance.

It is to document the ecosystem around it.

What absence of evidence means

Stalkerware also demonstrates one of the most persistent mistakes in digital investigations: treating a failed detection as proof of absence.

The AV-Comparatives/EFF test makes that particularly clear. If security products can miss some applications in a controlled environment, a clean scan cannot logically establish that covert monitoring is impossible.

The reverse is also true.

Battery drain, unusual data consumption, unfamiliar applications or unexpected device behaviour can be indicators worth investigating, but they are not individually proof that stalkerware is installed.

The FTC warns that stalkerware is deliberately difficult for device owners to detect and advises people who suspect monitoring to consider their safety before taking action, because searching for help or removing software from a monitored device can alert an abusive partner.

This introduces an unusual constraint for investigators.

In many cybersecurity situations, the obvious response to malicious software is to identify it and remove it. In a technology-facilitated abuse case, immediate removal may have consequences outside the device.

Technical remediation and personal safety are not always the same problem.

AV-Comparatives and EFF reached a similar conclusion in their testing. They found that most security products recommended removal without explaining the potential risks of doing so. The researchers argued that automatic removal would be inappropriate because the affected person should be able to decide when action is safe.

That is not merely a user-interface issue. It is an evidence-handling issue.

An investigator may discover something whose removal destroys evidence, changes the behaviour of another person or creates a new risk for the individual being monitored.

Knowing when not to act is part of the investigation.

The OSINT workflow is therefore evidence-first

A responsible public-source investigation into a suspected stalkerware vendor can be structured around a series of questions rather than a collection of tools.

What exactly is the product claiming to do?

Who appears to operate it?

How has it been marketed over time?

What permissions or monitoring capabilities are independently documented?

Does the product disclose monitoring to the device user?

What do platform rules require?

Have security researchers tested it?

Do regulators, courts or consumer-protection authorities have records concerning the company?

Are apparently separate products connected by independently observable corporate or technical evidence?

And, crucially, which conclusions remain unsupported?

The last question prevents an investigation from becoming an accusation built from circumstantial traces.

A shared server does not necessarily mean shared ownership. Similar website language does not prove that two services have the same operator. A hidden application icon may be relevant evidence, but it does not by itself establish how a particular customer used the software. A company disclaimer requiring consent does not prove that consent actually existed.

OSINT can establish relationships and contradictions. It cannot manufacture the missing link between them.

The larger lesson is about visibility

Stalkerware is particularly revealing because much of the surrounding business can operate in public.

The software may be covert on a victim’s device while the company selling it maintains websites, documentation, subscriptions and customer-facing infrastructure. Regulation may restrict particular vendors without eliminating the underlying demand. Security products may detect many applications while still missing others. New families continue to appear.

This is why the investigative question should not be reduced to “Can antivirus detect stalkerware?”

The more useful question is how the surveillance ecosystem can be observed across different layers of public evidence.

OSINT is well suited to that task because it can connect fragments that were never designed to be read together: marketing claims with platform policies, corporate identities with domains, historical webpages with current products, security research with regulatory records.

But the method only works if those connections retain their evidentiary limits.

In an investigation involving covert surveillance, the most dangerous error is not failing to find one more data point. It is converting an indicator into proof, a technical relationship into attribution, or a suspicious capability into an allegation about how a specific person used it.

The public footprint can reveal a great deal.

The investigator’s job is to make equally clear what it does not reveal.
Commercial surveillance software does not always arrive through an exploit, a clandestine intelligence service or a sophisticated intrusion campaign. Sometimes it is sold as a subscription, advertised as parental control or employee monitoring, supported by customer-service teams and installed on an ordinary smartphone.

That overlap between apparently legitimate monitoring software and covert surveillance is precisely what makes stalkerware difficult to investigate. The same technical capabilities can have very different meanings depending on consent, disclosure and use. An application that reports a device’s location is not, by that fact alone, stalkerware. An application designed to monitor another person without their knowledge, hide its presence and transmit sensitive information to a remote user is something very different.

A recent open-source investigation by Intel Focus brought the issue back into view. Published on 16 September 2026, it examined how applications presented as monitoring or parental-control tools can enable technology-facilitated abuse. The investigation arrives against a wider body of evidence showing that stalkerware has neither disappeared nor become a purely historical cybersecurity problem.

For OSINT investigators, however, there is another question worth asking: how much of this ecosystem can be understood without compromising a device, accessing private accounts or interacting with victims?

Quite a lot.

The websites, marketing language, corporate structures, app policies, security tests, regulatory records and technical relationships surrounding commercial surveillance products leave public traces. Taken separately, none necessarily proves abuse. Combined carefully, they can reveal how a market presents itself, how products are distributed, what capabilities are documented, where safeguards fail and where the boundary between legitimate monitoring and covert surveillance begins to collapse.

The important part is not finding a suspicious app. It is building an evidence chain that explains what the available information actually establishes.

The first problem is definition

“Stalkerware” is sometimes used loosely to describe almost any invasive application. That creates an immediate methodological problem.

The Coalition Against Stalkerware uses a narrower definition: software made directly available to individuals that enables someone to monitor another user’s device without that person’s consent and without explicit, persistent notification, in circumstances that may facilitate surveillance, harassment, stalking or abuse.

That distinction matters because monitoring functions are not inherently evidence of malicious use.

Google’s current developer policy illustrates the boundary. Monitoring applications intended for parental or enterprise use can be permitted under specific conditions, but they must disclose their monitoring function, display persistent notification while operating, identify themselves clearly and must not present themselves as tools for secret surveillance. Google also states that monitoring applications cannot be used to track a spouse, even with that person’s knowledge and permission.

For an OSINT investigation, this gives us something more useful than a label: criteria.

Instead of asking whether an application “looks like spyware”, an investigator can ask what the company claims the product is for, whether monitoring is disclosed to the person being monitored, whether the application is designed to remain visible, what information it collects and whether its documented behaviour corresponds with the platform rules under which legitimate monitoring software is supposed to operate.

That shift from labels to observable characteristics is fundamental.

Follow the public footprint, not the suspicion

A useful investigation can begin with the vendor’s own public material.

Product pages, terms of service, privacy policies, installation documentation, archived versions of websites, company records, app-store listings and regulatory documents can establish different parts of the picture. Search engines and web archives may reveal how marketing language changes over time. Public technical research can establish capabilities that promotional material leaves vague.

But these sources should not be treated as interchangeable.

A marketing page tells us how a company describes a product. It does not independently prove what the software actually does.

A privacy policy may document categories of collected information. It does not prove that every installation collects every category.

A security researcher may demonstrate a capability in a tested version. That finding should not automatically be projected onto every version of the product.

An enforcement action establishes what a regulator alleged or determined in that particular case. It does not establish that every monitoring vendor behaves in the same way.

This separation is where OSINT becomes investigation rather than accumulation.

The US Federal Trade Commission provides a concrete example. In a case involving Support King and its SpyFone service, the FTC alleged that the company sold applications allowing purchasers to monitor devices secretly, required users to weaken device security protections, and collected information including photographs, text messages, browsing histories, location and physical movements. A 2021 order banned Support King and its CEO from offering surveillance applications and required deletion of data that the FTC said had been secretly collected.

In December 2025, the FTC rejected the CEO’s petition to vacate or modify that order.

That is strong evidence about one company and one regulatory case. It is not evidence that every parental-control application is stalkerware.

Maintaining that boundary is essential.

The detection gap is measurable

One of the most useful pieces of evidence comes not from advertising or regulation but from controlled testing.

AV-Comparatives and the Electronic Frontier Foundation tested 13 Android mobile-security products against 17 commercial stalkerware samples in 2025. The test used a non-rooted Samsung Galaxy A36 running Android 15, and researchers installed the applications following the vendors' instructions before evaluating whether security products produced a clear warning.

The results varied substantially.

Detection rates ranged from 53% to 100% across the products tested. One detected all 17 samples, while Google Play Protect detected 53% in that particular test. Several products achieved rates above 90%, while others missed a significant portion of the test set.

These figures need context. They describe a defined 2025 test involving 17 selected applications and particular versions of security products. They are not universal detection rates for all stalkerware, nor do they tell us how the same products would perform against a different sample today.

What they demonstrate is more important: detection cannot be assumed.

The researchers also found evidence of shared infrastructure and rebranding. Some tested applications appeared to be variations of the same underlying product, reusing elements such as payment systems, backend infrastructure, administrative dashboards or APK files. Identical APK hashes were treated as one product to avoid counting duplicates.

For OSINT, that detail changes the investigative unit.

The visible brand may not be the real boundary of the ecosystem.

If multiple products reuse technical or commercial infrastructure, an investigation that searches only by brand name may fragment what is actually a connected network. Domains, certificates, company records, payment relationships, application identifiers, archived pages and publicly documented technical indicators can become more useful than the storefront name.

But infrastructure overlap is evidence of a relationship, not automatically evidence of common ownership. A shared service provider, reseller or white-label platform can produce similar traces. Attribution requires additional corroboration.

The market keeps changing

Kaspersky's research provides another view of scale, although its numbers must also be read according to their methodology.

For 2024–2025, the company reported more than 34,000 users affected by stalkerware in its telemetry and said it had identified 33 previously unseen stalkerware families. Affected users were detected in more than 160 countries. The figures come from aggregated threat statistics collected through the Kaspersky Security Network and therefore represent visibility within that ecosystem rather than a census of global stalkerware use.

That limitation is important.

Security-vendor telemetry can demonstrate that a phenomenon exists at significant scale within the

observed population. It cannot tell us how many stalkerware victims exist worldwide.

The real number could be different for several reasons: not every device uses participating security software, products may differ in what they classify as stalkerware, some applications may remain undetected, and technology-facilitated surveillance increasingly includes mechanisms other than dedicated stalkerware.

An OSINT article that simply repeats “34,000 victims” would therefore turn a measurable dataset into a claim the dataset does not support.

The defensible statement is narrower: Kaspersky detected more than 34,000 affected users in its 2024–2025 telemetry.

That difference may look small. Methodologically, it is enormous.

An application is only one layer of the investigation

The temptation in a stalkerware investigation is to concentrate on the APK. For malware researchers that may be necessary, but an OSINT investigation can work differently.

Think of the product as the centre of a public ecosystem.

Around it may exist a company, domains, support documentation, reseller pages, social accounts, advertising material, corporate registrations, privacy policies, payment services, application-store histories, archived websites, technical analyses, court documents and regulatory actions.

Each answers a different question.

Corporate records may help establish who operates a service. Archived pages can show how a product was marketed at a particular time. Platform policies provide criteria against which public product behaviour can be compared. Security testing can demonstrate capabilities or detection failures. Regulatory records can establish documented allegations, findings and enforcement actions.

The investigation becomes stronger when these layers corroborate one another.

It also becomes safer.

There is rarely a legitimate reason for an OSINT researcher to install commercial surveillance software on another person’s device, attempt to access a customer dashboard or obtain victim data simply to prove that a surveillance service exists. Public-source investigation should exhaust public evidence before crossing into intrusive technical testing, and any laboratory testing requires an appropriate legal and ethical framework.

The objective is not to reproduce the surveillance.

It is to document the ecosystem around it.

What absence of evidence means

Stalkerware also demonstrates one of the most persistent mistakes in digital investigations: treating a failed detection as proof of absence.

The AV-Comparatives/EFF test makes that particularly clear. If security products can miss some applications in a controlled environment, a clean scan cannot logically establish that covert monitoring is impossible.

The reverse is also true.

Battery drain, unusual data consumption, unfamiliar applications or unexpected device behaviour

can be indicators worth investigating, but they are not individually proof that stalkerware is installed.

The FTC warns that stalkerware is deliberately difficult for device owners to detect and advises people who suspect monitoring to consider their safety before taking action, because searching for help or removing software from a monitored device can alert an abusive partner.

This introduces an unusual constraint for investigators.

In many cybersecurity situations, the obvious response to malicious software is to identify it and remove it. In a technology-facilitated abuse case, immediate removal may have consequences outside the device.

Technical remediation and personal safety are not always the same problem.

AV-Comparatives and EFF reached a similar conclusion in their testing. They found that most security products recommended removal without explaining the potential risks of doing so. The researchers argued that automatic removal would be inappropriate because the affected person should be able to decide when action is safe.

That is not merely a user-interface issue. It is an evidence-handling issue.

An investigator may discover something whose removal destroys evidence, changes the behaviour of another person or creates a new risk for the individual being monitored.

Knowing when not to act is part of the investigation.

The OSINT workflow is therefore evidence-first

A responsible public-source investigation into a suspected stalkerware vendor can be structured around a series of questions rather than a collection of tools.

What exactly is the product claiming to do?

Who appears to operate it?

How has it been marketed over time?

What permissions or monitoring capabilities are independently documented?

Does the product disclose monitoring to the device user?

What do platform rules require?

Have security researchers tested it?

Do regulators, courts or consumer-protection authorities have records concerning the company?

Are apparently separate products connected by independently observable corporate or technical evidence?

And, crucially, which conclusions remain unsupported?

The last question prevents an investigation from becoming an accusation built from circumstantial traces.

A shared server does not necessarily mean shared ownership. Similar website language does not prove that two services have the same operator. A hidden application icon may be relevant evidence, but it does not by itself establish how a particular customer used the software. A company disclaimer requiring consent does not prove that consent actually existed.

OSINT can establish relationships and contradictions. It cannot manufacture the missing link between them.

The larger lesson is about visibility

Stalkerware is particularly revealing because much of the surrounding business can operate in public.

The software may be covert on a victim’s device while the company selling it maintains websites, documentation, subscriptions and customer-facing infrastructure. Regulation may restrict particular vendors without eliminating the underlying demand. Security products may detect many applications while still missing others. New families continue to appear.

This is why the investigative question should not be reduced to “Can antivirus detect stalkerware?”

The more useful question is how the surveillance ecosystem can be observed across different layers of public evidence.

OSINT is well suited to that task because it can connect fragments that were never designed to be read together: marketing claims with platform policies, corporate identities with domains, historical webpages with current products, security research with regulatory records.

But the method only works if those connections retain their evidentiary limits.

In an investigation involving covert surveillance, the most dangerous error is not failing to find one more data point. It is converting an indicator into proof, a technical relationship into attribution, or a suspicious capability into an allegation about how a specific person used it.

The public footprint can reveal a great deal.

The investigator’s job is to make equally clear what it does not reveal.