

N9LYA TECHNICAL GUIDES

Stop the Dataminers

Blocking Telemetry at Home

Jerome M. Kutche

N9LYA

Second Edition · September 2026

About This Book

Many smart TVs, routers, printers, thermostats, apps, and connected devices send operational and usage data to manufacturers or service providers. The industry calls it “telemetry.” Depending on the product and settings, that can include device identifiers, usage patterns, viewing information, diagnostics, and network activity; some recipients may also use or share data for advertising and analytics.

None of this requires you to have done anything wrong, or even anything unusual. Collection frequently begins during setup or first connection, under choices and privacy notices that vary widely in clarity. The practical issue is not that every transmission is malicious. It is that users often lack a clear inventory of what leaves the network, why it is sent, how long it is retained, and which other parties receive it.

This book walks through four levels of defense, from a five-minute browser fix to a full homelab-grade setup. Pick the level that fits your skills and time. **Even Level 1 alone makes a real difference.** None of these levels require you to throw out your devices or give up convenience — the goal is visibility and control, not a return to dial-up.

What Makes This One Different

There is no shortage of privacy advice online. Most of it has one of three problems.

It’s abstract. Long on why privacy matters, short on which setting to change.

It’s absolutist. Advice that assumes you’ll abandon every convenience, which nobody does, so the advice gets ignored entirely.

Or it’s a link farm — a list of tools with no explanation of what each actually blocks, what it misses, and how they layer.

This book is written from a working setup. **Every blocklist, custom rule, and configuration note here is pulled from a live network** — a segmented home lab with a virtualized DNS filter, a router doing IP-level blocking as a second layer, and IoT devices isolated on their own VLAN. The numbers are real numbers. The rules are rules that are actually loaded.

And it’s honest about the limits. DNS filtering has blind spots, and this book names them. Some devices bypass everything you do. Some blocking breaks things, and the troubleshooting chapter exists because it happened. A guide that only tells you what works isn’t a guide — it’s marketing.

Who This Is For

Anyone with a home network. Level 1 takes five minutes and requires no technical skill.

People with smart TVs, streaming boxes, and IoT devices — the category with the worst telemetry behavior and the least ability to run any protection on the device itself.

Homelabbers and self-hosters. Part V is written for you specifically.

Small business owners running a network with no IT department.

And anyone who has looked at a privacy policy, felt vaguely uneasy, and had no idea what to actually do about it. That’s the reader this book was written for.

How to Use It

Part I explains the problem — what telemetry actually is, who collects it, what it reveals, and what the law does and doesn’t do about it. Read it, or skip to the level you want and come back.

Parts II through V are the four levels, in order of increasing effort and coverage. **Each level stands on its own.** You can stop at any point and have made real progress.

Part VI is operating it — verification, troubleshooting, and long-term maintenance, which is where most setups quietly fail.

Each chapter ends with a **Block Check** — things to verify or try on your own network. Do them. The gap between “I set that up” and “it’s actually working” is where most privacy tooling lives, and the checks are what close it.

The appendices are the reference material: blocklist selection, device-specific notes, rule syntax, and troubleshooting trees. **Those are what you’ll come back to.**

A Note on Currency

Tool versions, blocklist maintainers, and specific domains change. Versions cited here reflect 2026 and will age.

The mechanisms don’t change. How DNS filtering works, why hardcoded IPs bypass it, what encrypted DNS does to your visibility, and why segmentation matters are all durable. **Parts I and the conceptual sections will outlast the specifics.** Verify current versions and list URLs against the projects’ own pages before deploying anything you’ll depend on.

About the Author

Jerome M. Kutche holds the amateur radio callsign N9LYA, first licensed in 1991. He spent over three decades in industrial electrical and controls work — Allen-Bradley and Siemens PLC systems, motor control, plant electrical maintenance and supervision — holding NFPA 70E and thirty-nine other professional certifications. From 1990 to 1996 he built RF test and measurement equipment for the telecommunications and cable television industries.

He serves as Indiana’s 44Net/AMPRNet coordinator and is trustee for two amateur radio clubs. His home network runs segmented VLANs behind a MikroTik router, with a virtualized DNS filter, reverse proxying, site-to-site VPN, and rather more monitoring than is strictly necessary.

The setup described in Part V is his own, and most of the troubleshooting advice comes from having broken it first.

Part I — Understanding the Problem

Chapter 1: What Telemetry Actually Is

Before blocking anything, it’s worth being precise about what’s being blocked — because the word “telemetry” is doing a lot of work to make this sound routine.

The Definition

Telemetry is data a device or application collects about its own operation and transmits to the manufacturer or a third party, automatically, without a specific active choice by the user each time it happens.

That last clause is the load-bearing one. **You may have consented once, in a policy you didn’t read, at setup.** You are not consenting each time, and in most cases you have no visibility into what’s being sent or when.


What Gets Collected

The categories, roughly in order of how commonly they appear:

Device and system information. Model, serial number, firmware version, hardware identifiers, storage capacity, screen resolution. **Individually mundane; collectively a fingerprint.**

Where Household Data Travels

Telemetry is a chain, not a single destination



You can interrupt the chain at more than one layer.

N9LYA TECHNICAL GUIDES

Figure 1: The household telemetry path

Usage patterns. What apps you open, how long you use them, what you click, when you use the device. On a TV: what you watch, for how long, and whether you skipped the ads.

Network information. Your IP address, your network name, and — this is the one that surprises people — **an inventory of other devices on your network.** Some devices scan the local network and report what they find.

Location. Explicitly via GPS on mobile devices, or inferred from IP address, WiFi network names in range, or your account address.

Performance and diagnostics. Crashes, errors, battery, temperature. **This is the category that genuinely helps manufacturers make products better**, and it's the category everything else gets bundled with.

Content identification. On smart TVs specifically — what's actually on screen, including from an HDMI-connected device. Chapter 3.

Account linkage. Your device data tied to an account, which is tied to your name, email, payment method, and everything else that account touches.

The Frequency Nobody Mentions

The thing that changes the picture: this isn't occasional.

A smart TV sitting powered on and idle, with nothing playing, generates outbound telemetry requests continuously — commonly every few minutes, sometimes more often. **A device you aren't using is still reporting.**

You can watch this happen. Once you have Level 3 running, a query log filtered to one device shows you exactly how often it's phoning home. **For most people this is the moment the abstract becomes concrete**, and it's why Chapter 18 spends time on reading the log rather than just installing the tool.

What Telemetry Isn't

Being fair about the distinction:

Not all data transmission is telemetry. A streaming app contacting a content server to fetch video is doing the job you asked it to do. A device checking for a firmware update is legitimate. **The goal here isn't blocking all outbound traffic** — it's blocking the surveillance layer riding alongside the functionality.

Not all telemetry is malicious. Crash reporting genuinely improves software. Some manufacturers collect narrowly and handle it responsibly.

The problem is that you can't tell the difference from the outside, the collection is bundled so you can't accept one and refuse the other, and the industry norm has drifted so far toward collecting everything that "reasonable" collection is now the exception rather than the default.

Why This Matters

It's worth being specific about what's actually at stake, because "privacy" as a word has been used so often it's started to feel abstract. Three concrete reasons this is worth an afternoon of your time:

The data compounds. A single data point — one show watched, one device connected — is nearly meaningless on its own. But telemetry isn't collected once; it's collected continuously, indefinitely, and correlated across every device in your house and often across other accounts you hold with the same company. **A pattern built from months of continuous collection is a genuinely detailed picture of your household's habits, presence, and schedule** — far more than any single data point would suggest.

You don't control where it goes after collection. Most telemetry doesn't stay with the manufacturer. It's routed to analytics platforms, advertising networks, and data brokers, who aggregate it

with data from other sources and resell it. **Once your data leaves your network, you have no practical way to know who has it or what they've combined it with.**

Blocking it costs you almost nothing. This is the part that makes the effort worth it: unlike a lot of privacy trade-offs, blocking telemetry rarely breaks the actual functionality of a device. Your smart TV still plays video. Your printer still prints. **What stops working is the surveillance layer riding alongside the feature you actually wanted** — which is exactly the layer that shouldn't have been bundled in without a real choice in the first place.

The Prime Example: Amazon eero

Amazon's eero WiFi system is one of the most aggressive data collectors you can plug into your home network. It continuously transmits device inventory and usage data back to Amazon's servers, and Amazon's own privacy policy openly states this data is shared with third parties.

"ET phoned home once. The eero never hangs up."

This isn't unique to eero — most consumer smart-home gear behaves the same way. **Eero is simply one of the most visible and well-documented examples, in part because it sits at the center of the network by design**, giving it visibility into every other device you own.

And that position is the point. A router occupies a privileged place: it can observe connected devices and routed flows, plus DNS queries that use its resolver or traverse it without encryption. VPNs, encrypted DNS, proxies, and relays can limit the detail it sees, but the placement still deserves careful vendor and policy scrutiny.

Block Check

1. Pick one device in your house. Find its privacy policy and read the data collection section.
2. Count how many categories from this chapter it mentions.
3. Look for the phrase "select partners" or "third parties." Note whether they're named.
4. Check whether the device has a telemetry opt-out, and how many menus deep it is.

Chapter 2: Where the Data Goes

Collection is only half the picture. What happens after matters more, and it's the part manufacturers describe least specifically.

The Chain

Your device → the manufacturer → analytics and advertising platforms → data brokers → buyers.

Each hop adds parties, and the further along the chain, the less any policy you agreed to constrains what happens.

The Players

Manufacturers collect first-party. Their policy governs, and it typically reserves broad rights to share.

Analytics platforms are the intermediaries most devices report to. **A single app frequently contains several analytics SDKs**, each reporting independently to a different company. The app developer may not fully know what each one collects.

Advertising networks receive data used to target ads, and — critically — **to build profiles that persist across devices and services.**

Data brokers are the least visible and most consequential. They aggregate data from many sources — device telemetry, purchase records, public records, loyalty programs, location data — and build combined profiles they sell.

The broker industry is where the compounding in Chapter 1 actually happens. No single source knows very much. The aggregate is a different thing entirely.

The Identifiers That Tie It Together

Data from separate sources only compounds if it can be linked. The linkage mechanisms:

Advertising IDs — a resettable identifier on phones and streaming devices, designed for exactly this purpose.

IP address — your household’s shared identifier. **Every device on your network presents the same public IP**, which is how a profile gets assembled across devices that share nothing else.

Account email — the strongest link, because it’s the same across services.

Device fingerprints — combinations of characteristics distinctive enough to identify a device without any explicit identifier. Chapter 8.

Hashed identifiers. Worth understanding: a phone number or email run through a hash function is frequently described as “anonymized.” **It isn’t.** The same input always produces the same output, so it works perfectly well as a linking key — it just isn’t human-readable. **“Hashed” means “still a unique identifier.”**

On “Anonymized” and “Aggregated”

Two words that appear in nearly every privacy policy and mean less than they sound.

“Anonymized” typically means direct identifiers were removed. **Re-identification from supposedly anonymous data is a well-established result** — a small number of data points about location, timing, or behavior is frequently enough to single out an individual within a large dataset. **Anonymization is a spectrum, and most of what’s labeled anonymous sits at the weak end.**

“Aggregated” means combined with other people’s data. Genuinely aggregated statistics are fine. **The word is also used loosely for data that’s merely stored alongside other data**, which is not the same thing.

The practical guidance: treat both words as describing intent rather than a technical guarantee, unless the policy specifies the actual method.

Why the Chain Is Hard to See

Three structural reasons you can’t audit this:

Policies describe categories, not parties. “We share with service providers and advertising partners” is compatible with three recipients or three hundred.

Downstream parties aren’t bound by upstream policies. A broker who buys data isn’t operating under the manufacturer’s privacy policy.

And there’s no expiry. Data collected in 2019 doesn’t disappear because you changed a setting in 2026. **The only reliable control is preventing collection**, which is what this book is about.

The Practical Conclusion

You cannot manage this after the fact. There’s no meaningful way to audit who has your data, no reliable way to have it deleted across the chain, and no practical recourse when it’s resold.

Which means the only effective intervention is at the point of transmission — your own network. That’s not a philosophical position; it’s the only place in the chain where you have actual control.

Block Check

1. On a phone or streaming device, find the advertising ID setting. Reset it, and note whether the option to disable it entirely exists.
2. Look up whether your TV manufacturer names its data partners specifically.
3. Count the number of distinct devices in your house sharing your public IP.
4. Check one privacy policy for the words “anonymized” or “aggregated” and see whether the method is specified.

Chapter 3: What the Data Actually Reveals

The abstract case for privacy is unpersuasive to most people. The concrete case is much stronger, and this chapter is the concrete case.

Presence and Schedule

Telemetry timestamps reveal when your house is occupied.

A TV that reports when it’s on, a thermostat reporting adjustments, a voice assistant reporting activations, a smart plug reporting switching — each is a timestamped signal of human presence.

Weeks of this produces your household’s schedule with considerable accuracy: when you wake, when you leave, when you return, when you sleep, when you’re away for several days.

Who wants that: advertisers targeting by daypart, insurers modeling occupancy, and — in the aggregate-and-resell market — anyone who buys the resulting profile.

Household Composition

Device inventory plus usage patterns implies who lives there.

A device used only in the afternoon on children’s content. A second phone appearing on the network every evening. A gaming console with distinct usage windows. **The inference isn’t certain and it doesn’t need to be** — probabilistic household composition is exactly what the advertising industry buys.

What You Watch, Precisely

This is where smart TVs are categorically worse than other devices, because of one specific technology.

ACR — Automatic Content Recognition — works by taking periodic snapshots or audio fingerprints of whatever is currently displayed on screen and matching them against a database.

The critical detail: it captures content from HDMI-connected devices, not just the TV’s own apps.

That means a TV with ACR enabled reports what you’re watching on your cable box, your game console, your Blu-ray player, and your laptop. **Content that never touched the TV’s software is still identified**, because the identification happens at the screen rather than in the app.

It’s usually enabled by default, described in the setup wizard in language that sounds like a feature (“personalized recommendations,” “content insights”), and **buried several menus deep to disable** — under names that don’t obviously relate to what it does.

This is exactly why blocking its outbound domains at the DNS level is more reliable than hunting for a toggle in the TV’s UI. The toggle may not exist, may not be honored, or may reset after a firmware update. **A blocked domain stays blocked.**

Inferred Attributes

The category that most people find genuinely uncomfortable.

Behavioral data supports inference of attributes you never disclosed: approximate age, household income bracket, health interests, political leaning, religious observance, relationship status, and — in one much-discussed and well-documented category — pregnancy, inferred from purchasing and browsing patterns before it was announced.

These inferences are frequently wrong. That's not reassuring. **A wrong inference in a profile that's sold and acted upon is a problem you can't see or correct.**

The Aggregation Example

A concrete walk-through, because the abstract case never lands.

Take one household over one month, using only data that is routinely collected and routinely sold:

From the TV: on at 6:40 a.m. most weekdays, off by 7:15. On again 6 p.m. to 11 p.m. Content identified by ACR — a mix of children's programming early evening, then adult drama after 8:30.

From the thermostat: setback ends 6:20 a.m., away setting 8:05 a.m. to 5:40 p.m. weekdays. No away setting on Tuesdays.

From smart plugs and lights: a bedroom light on at 9:15 p.m. and off at 9:40, every night.

From phones on the network: two adult devices, one arriving home at 5:40, one at 6:15. A third device present only on weekends.

From the router: a fourth device that appears every weekday at 3:20 p.m.

What that composite says, without a single piece of it being sensitive on its own: two working adults with different schedules, at least one school-age child arriving home mid-afternoon, a younger child with a 9:40 bedtime, one adult working from home on Tuesdays, a visitor or family member on weekends, and **an empty house from 8:05 to 3:20 on four weekdays a week.**

Nobody collected "when is this house empty." It fell out of the aggregate.

And that composite is a product. It's assembled, priced, and sold, and the household never saw a question they could have answered differently.

The Downstream Uses

Where profiles actually get used, beyond advertising:

Pricing. Differential pricing based on inferred willingness to pay, device type, or location.

Insurance and credit modeling, where permitted, using behavioral proxies.

Employment and tenant screening through background-check services that incorporate purchased data.

Security consequences. A detailed profile is raw material for targeted social engineering. **A phishing attempt referencing your actual devices, schedule, and household is dramatically more effective than a generic one.**

And breach exposure. Data collected is data that can leak. **Every broker holding your profile is a breach surface you didn't choose and can't audit.**

The Argument That Doesn't Hold

"I have nothing to hide" fails for three reasons worth stating plainly:

It assumes the data is about wrongdoing. It isn't. It's about schedule, habits, household, and health — none of which is wrongdoing and all of which is nobody's business.

It assumes you control the framing. You don't. An inference drawn about you, by a system you can't inspect, used in a decision you're not told about, is not something "having nothing to hide" protects you from.

And it assumes today's context holds. Data collected under one set of norms persists into a different one. **The relevant question isn't whether you object to who has it now.**

Block Check

1. Find the ACR setting on your TV. Note what it's called and how many menus deep it is.
2. Turn it off, then check whether the TV still queries its ACR domains. (Requires Level 3.)
3. List what an observer could infer about your household from device timestamps alone.
4. Check whether any device in your house has ever asked permission before collecting.

Chapter 4: The Legal Landscape

What the law does and doesn't do, and why the answer is "less than you'd hope" almost everywhere.

The General Shape

Two broad regulatory models exist.

Comprehensive frameworks — the European approach — establish a general rule that personal data may only be processed with a lawful basis, grant individuals rights over their data, and apply across sectors. **The GDPR is the reference implementation** and it has influenced legislation well beyond Europe.

Sectoral and state-level frameworks — the US approach — regulate specific sectors (health, financial, children's data) federally, with general consumer privacy handled at the state level. **The result is a patchwork:** a growing number of US states have comprehensive privacy laws, with meaningfully different requirements, thresholds, and enforcement.

The practical consequence for a household: your rights depend heavily on where you live, and in most of the US they're narrower than people assume.

What Rights Typically Exist

Where comprehensive law applies, common rights include:

Access — request what a company holds about you. **Deletion** — request removal, subject to exceptions. **Correction** — fix inaccurate data. **Portability** — receive your data in a usable format. **Opt-out of sale or sharing** — increasingly including a recognized universal opt-out signal. **Non-discrimination** — not being penalized for exercising these rights.

These are real and they're worth using. Chapter 31 covers how.

Why They Don't Solve This

Five structural limits, stated plainly:

Consent is theater. A dense policy presented at setup, with the only alternative being not using the device, is not meaningful choice. **Regulators increasingly agree; enforcement lags.**

Rights apply to identified data. Data a company classifies as anonymized or aggregated may fall outside the scope entirely. Chapter 2 — those words are doing more work than they should.

Enforcement is thin. Regulators are resourced for a small number of large cases, not for the long tail of device manufacturers.

Deletion doesn't propagate. A deletion request to a manufacturer doesn't reach the brokers who bought the data years ago.

And the burden is on you, per company, forever. Exercising rights across every company holding your data is not a realistic ongoing practice for an individual.

None of this means the law is useless. It means it's a complement to technical control, not a substitute.

The Legality of Blocking

A question people ask, so worth answering directly.

Configuring DNS, firewall rules, and content blocking on equipment you control is generally ordinary network administration. It can still affect warranties, support, employment systems, managed devices, service terms, or local legal obligations. This book is technical guidance, not legal advice; use extra care on networks or devices you do not solely own.

You are not intercepting anyone else's traffic, not modifying anyone else's systems, and not accessing anything you're not entitled to.

Two edges worth knowing:

Terms of service. A manufacturer's terms may purport to require telemetry. **Blocking it is a contract question, not a criminal one,** and the practical consequence is at most a feature not working.

Networks you don't control. Everything here applies to your own network. **Filtering an employer's network, a rental property's shared network, or anyone else's is a different question** and not one this book addresses.

The Direction of Travel

Regulation has been tightening, with more jurisdictions adopting comprehensive frameworks, growing attention to connected devices specifically, and increasing recognition of universal opt-out mechanisms.

And the collection has been growing faster. More devices, more sensors, more inference capability.

The honest read: technical control is the thing you can actually implement this weekend. Regulatory improvement is worth supporting and it operates on a different timescale than your network does.

Block Check

1. Determine which privacy law applies where you live, and what rights it grants.
2. Check whether your browser supports a universal opt-out signal, and enable it.
3. Submit one data access request to a manufacturer whose device you own. Note how long it takes.
4. Read your ISP's privacy policy. Note what it says about DNS query data.

Chapter 5: Threat Modeling for a Household

Deciding how far to go, and against what. A short chapter that saves a lot of wasted effort.

The Question

"How private should I be" is unanswerable. "What am I protecting against" is answerable.

For a typical household, the realistic adversaries are, in order:

Commercial data collection. Manufacturers, analytics platforms, advertising networks, brokers. **Continuous, automated, indiscriminate, and the thing this book addresses.**

Household Privacy Threat Model

Protect what matters instead of chasing perfect invisibility

ASSETS

- Location
- Habits
- Contacts

COLLECTORS

- Apps
- Platforms
- Devices

EXPOSURE

- Profiling
- Breach
- Misuse

CONTROLS

- Minimize
- Block
- Verify

N9LYA TECHNICAL GUIDES

Figure 2: A practical household privacy threat model

End of the free sample

Stop the Dataminers — Second Edition · by Jerome M. Kutche

That was the first 12 of 68 pages. The full book continues with the practical detail, worked examples and recovery steps that don't fit in a sample.

Get the full book (\$5.99, PDF + EPUB):

<https://booksrus.biz/product/stop-the-dataminers/>

Also on Gumroad:

<https://n9lya4.gumroad.com/l/oefaav>



All 17 Second Edition guides in one bundle for \$89.99:

booksrus.biz/product/n9lya-technical-guides-complete-bundle

This sample may be shared freely in unmodified form. © Jerome M. Kutche. Second Edition.