

BOOK SIX / GOVERNMENT FACILITIES

The Doors That Guard the Building

Securing the cameras, badge readers and
building systems behind government facilities.

NetOpsIQ™

Eldad Matityahu

Founder and CEO, NetOps IQ

Published October 1, 2026



Five minutes.

Someone outside spent five minutes trying to break into a government network through its cameras.

Findavi caught the attempt. A person approved the recommended block. Findavi closed the connection and checked that the cameras were still recording.

The attacker never got past the cameras.

The cameras never stopped watching.

Confirmed by NetOps IQ's CTO. Published with customer approval. Agency and location withheld. Full account in Chapter 7.

Inside the building.

A Letter From Eldad	04	07 Five Minutes at the Camera	17
A Note on Evidence	05	08 The Camera Still Has to Record	20
01 The Building Has a Network	06	09 Ten Minutes With the Facility Director	23
02 A Computer With a Lens	08	Closing: The Cameras Kept Watching	25
03 The Door System	10	Sources	26
04 Six Doors Into the Facility	12	Start With One Building	29
05 Listen First	14		
06 Three Questions Before You Touch Anything	16		

To the people who keep public buildings open.

Every morning, someone unlocks a public building and expects it to work.

The doors open for the right badges. The cameras record the lobby. The gate lifts for the delivery truck. The heat comes on before the first visitor walks in. Nobody thinks about any of it, and that's how it should be.

But every one of those things runs on a small computer. The camera has software and a password. So does the badge reader. So do the gate, the intercom, the visitor kiosk and the box that runs the heating. Many of them share a network with everything else. Some were installed by a contractor years ago and never touched again.

Attackers have noticed government. In 2025, at least 44 U.S. states reported cyber incidents that hit state and local government systems. The devices that guard a building are computers too. They are often the ones nobody is watching.

This book is about those devices. It's for the people who answer for them: security, IT, facilities, physical security, procurement and the leaders who sign off on all of it.

It makes one promise. When something goes wrong at one camera or one door, you should be able to shut down that one problem and leave the rest of the building running.

Then you should be able to prove two things. The attacker's path is closed. And the door still opens for the right person.

Eldad Matityahu

Founder and CEO, NetOps IQ

What is real. What is illustrated.

The opening story is an observed NetOps IQ deployment account, confirmed by its CTO and approved by the customer for publication. The agency and location are withheld.

Independent evidence means public records support the event or vulnerability described.

Illustrative scenes explain a decision or workflow. They are constructed scenarios, not customer incidents. All cinematic artwork is conceptual and does not depict the customer site.

Product descriptions reflect NetOps IQ material. Device support, permissions and service checks must be agreed for each deployment.

This book follows one practical problem: the devices that guard a building are computers, and making them safer must not stop them doing their jobs. Responsibility needs a named owner. A closed attack path needs evidence. A working door, camera or alarm needs its own check.

The Building Has a Network

The security system runs on computers. The computers run on a network.

Take a walk through a government building at 5:30 in the morning. Nobody is there yet. The building is already busy.

A camera in the parking lot is sending video to a recorder in a closet. The vehicle gate is waiting for a code. At the loading dock, an intercom is ready to ring the security desk. Inside the lobby, a visitor kiosk sits ready to print the day's first temporary badge.



Connected doors may have readers linked through access-control panels to a server. The server keeps a log of who went where and when. Upstairs, a building controller is warming the offices and starting the air handlers. Somewhere, a vendor has a remote login so they can fix the system without driving in.

Count them up. The list gets long fast.

Twenty years ago, a secure building meant locks, guards and a few cameras on tape. Today it means computers you can walk past without noticing. Each one has software. Each one has a password. Each one has a connection to something else.

Now ask who owns them.

IT runs the network. Physical security runs the cameras and the doors. Facilities runs the heating and the air. A contractor installed half of it and still holds the keys. Procurement bought it all from different vendors over different years.

Everyone touches these devices. Nobody owns the risk of all of them together. That gap is where attackers like to work.

Ask your team: Could you list every connected camera, reader, gate, controller and vendor login in one building by tomorrow morning?

A Computer With a Lens

The camera that watches the door can become the door.



CONCEPTUAL ARTWORK

Look up at the camera over the public entrance. It has a lens. It also has software, a login, a password and a network cable. That makes it a computer. It just happens to point at a door.

Attackers know this. Three public records illustrate the risk.

INDEPENDENT EVIDENCE

In August 2024, CISA warned about an AVTECH security camera. A flaw let someone send commands to the camera over the network and make it run them. Attackers were already using it. CISA also reported that the maker had not answered its requests to help fix the problem.

In June 2025, CISA warned about pan-tilt-zoom cameras from PTZOptics and several other brands. Some shipped with administrator passwords that were easy to crack. Owners could not change those passwords, and could not switch off the remote login either. CISA listed government services and facilities among the places these cameras are used. It had no reports of attacks on these flaws at the time.

In September 2024, the FBI, the NSA and their partners described a China-linked botnet associated with actors known as Flax Typhoon. By June 2024 it had taken over more than 260,000 devices, nearly half of them in the United States. Its targets included IP cameras and video recorders, the same kind of equipment that guards a public building.

One caution. None of these records describes a break-in at a specific government building. What they show is simple. Cameras can be taken over from far away, and attackers do it at scale.

When a camera is turned, three things can go wrong. You lose the picture you were counting on. Someone else gains a quiet seat inside your network. And the camera can start talking to places it was never meant to reach.

Then comes the hard part. You can't unplug every camera just to be safe. That leaves the building blind.

Core idea: A camera is a computer with a lens.

Ask your team: Which of our cameras can reach more than their recorder and the system that manages them?

The Door System

The system that lets the right people in must keep the wrong people out.

A badge reader has two jobs, and they pull in opposite directions.

It has to stop the wrong person. And it has to let the right person through, every time, at the right moment. The employee at 7:00 a.m. The contractor at the loading dock. The firefighter at 2:00 a.m.



Behind every reader sits a bigger system. A controller or server checks the badge, depending on the system. A database holds the names and the rules. A log records who opened which door and when. That system is software, and software has flaws.

INDEPENDENT EVIDENCE

In July 2024, CISA published two warnings about C-CURE 9000, a widely used access-control platform from Johnson Controls. One flaw could let an attacker reach administrator access through weak installer passwords. The other could let an attacker get at the passwords the system itself uses. CISA updated the second warning twice in 2025. In both cases, CISA said it had no reports of anyone exploiting the flaws.

So this is a warning. Nobody broke in. It still matters, because of what comes next.

Say your team reads an alert like this on a Tuesday afternoon. The fix sounds simple: update the software, change permissions, tighten the network. Now picture doing it in a hurry.

A permission change locks out the account the readers depend on. A network rule cuts the reader off from its server. The badges stop working at the side entrance, and nobody finds out until the morning shift is standing outside. Or the doors still open, but the log stops recording who came in.

The flaw is closed. The building is worse off.

That's why patching is only half the job. The other half is proof that the door still works for the people who belong there.

Core idea: A secure door still has to open for the right person.

Ask your team: After our last change to the access-control system, what proved that the right badges still opened the right doors?

Six Doors Into the Facility.

Attackers look for the device someone forgot to own.

DOOR	WHAT IT LOOKS LIKE	HOW YOU CLOSE IT
1. The factory password	Factory or fleet-wide passwords remain in use.	Use unique credentials; verify service.
2. The open control screen	Management pages are reachable from unapproved networks.	Allow only approved access paths.
3. The flat network	Cameras, doors, controls and office systems can reach each other.	Separate systems; permit required traffic.
4. The vendor who never left	An installer’s old remote login remains active.	Name an owner; approve and expire access.
5. The radio nobody watches	RFID, wireless and cellular links sit outside wired visibility.	Define supported radio coverage and its limits.
6. The fix nobody checked	Containment silently breaks entry, recording or a building service.	Require a security check and a service check.

See. Prove. Control.

The first five doors are about how attackers get in. The sixth is about how defenders hurt themselves on the way out. It's the one most teams never measure. Both belong in the same plan.

See starts with the devices and connections in the agreed scope. Map what is observed, who owns it and what the assessment cannot see. An invisible segment is still an unanswered question.

Prove connects each finding to evidence. Establish what the device can reach, which service depends on it and who can approve a change. A severity score alone cannot tell you whether a public entrance will stop working.

Control means choosing a supported action, applying the approval policy and checking the result. Close the risky path, then separately check the service. Keep the decision and both checks in the record.

Which of these six doors would you be least comfortable swearing is shut in every building? Start the conversation there.



CHAPTER 05

Listen First

Looking too hard can break the thing you're looking at.

ILLUSTRATIVE

It's Monday morning. The lobby is filling up.

The security team wants to scan the access-control network to see what's on it. The facilities director asks one fair question.

"What happens if the door panel doesn't like being scanned?"

Nobody in the room has a good answer.

She's right to ask. Many of the devices that guard a building are old, built by small vendors and never designed to be poked. An aggressive scan can freeze a controller or knock a reader offline. The very act of checking for risk becomes the risk.

So Findavi starts by listening.

The Findavi appliance sits on site and connects to a port that copies the network traffic you approve. It doesn't send anything to your cameras or readers. It hears what they already say to each other, every day, and builds a picture from that.

From listening alone, it can tell a lot. Which devices are on the network. What they are and who makes them. Who each one talks to. Which ones still send passwords in the clear. When logins aren't encrypted, it can even spot one password shared across a whole fleet. Which ones are suddenly talking somewhere new.

The AI that does this runs on the Findavi appliance itself. No outside AI. The network data never leaves your site.

The Findavi appliance can hear the radio too. Connected to an RF switch, it watches RFID and wireless activity in range of that switch.

One honest limit. Findavi sees only the parts of the network it's connected to. A camera on a segment nobody mirrored is still invisible. Every assessment says plainly what it covered and what it didn't.

Core idea: Listen first. Don't poke the system that guards the building.

Ask your team: When did someone last check our list of cameras, readers and controllers against what's actually plugged in?

Three Questions Before You Touch Anything

The first look at a building usually turns up a long list. Old firmware. Shared passwords. A vendor login nobody remembers. Cameras talking to places they shouldn't.

The list feels urgent. It's also useless until you decide what comes first. Severity scores help, but they don't know your building. A flaw rated critical on a camera in an empty storage room can wait. A medium flaw on the reader at the main entrance can't.

Three questions sort the list.

1. What can it reach?

Can this device reach the access-control server? The video recorders? The building controls? A vendor's network? The internet? A device that can reach the systems that run the building moves to the top. A device boxed into its own corner moves down.

2. What breaks if it fails?

Would people get stuck outside? Would an emergency exit stop working the way it should? Would the cameras stop recording? Would an alarm stop reaching the people who respond to it? Would the heat or the air go down in a building full of staff? The bigger the answer, the more care the fix needs.

3. Can we change it safely today?

Is the fix one Findavi supports for this device? Has the owner said yes? Is there a quiet window to make the change? Is there a way back if it goes wrong? Do we know exactly how we'll check that the attack path is closed, and that the device still does its job?

If any answer is no, the fix waits. Keep watching the device and limit what it can reach in the meantime.

Core idea: Severity tells you something is wrong. What breaks if it fails tells you what comes first.

Ask your team: Which five devices in our most important building would cause the most trouble if they stopped working today?

What a broad shutdown costs.

When you can't tell which device is the problem, you're left with one move. Shut everything down.

INDEPENDENT EVIDENCE

On July 25, 2025, the city of Saint Paul, Minnesota, detected attackers inside its systems. Within days, the mayor declared a local emergency and the governor called in the National Guard's cyber protection team. The city shut down much of its network to contain the attack, and city services were disrupted. The Guard's cyber team stayed for 17 days. A ransomware gang called Interlock later took credit. The city refused to pay.

Officials have not said how the attackers got in. Nothing in the public record ties Saint Paul to a camera or a door. It belongs here for one reason: it shows what a broad shutdown costs a government that has to keep serving the public.

Five minutes.

That's how long someone outside spent trying to break into a government network through its cameras.

They wanted a way in. The cameras looked like one.

A Findavi appliance was already on that network, listening. It caught the attempt and sent the security team one recommendation: block the connection.

A person read it and said yes. Findavi blocked it.

Then it checked two things. Is the way in shut? Yes. Are the cameras still recording? Yes.

The attacker never got past the cameras. The cameras never stopped watching.



MINUTES

Confirmed by NetOps IQ's CTO. Published with customer approval. Agency and location withheld.

Seven steps. Two checks.

Every Findavi action follows the same seven steps.

1. Spot. See behavior that doesn't fit the device.
2. Check the rules. Is this action allowed for this device, at this time?
3. Ask. If a person must approve, send one clear recommendation and wait.
4. Act. Block the device, isolate a port or push a configuration change, for supported devices only.
5. Check security. Test the path again. Is it closed?
6. Check service. Is the device still doing its job?
7. Record. Keep the finding, the approval, the action and both results in one place.

You decide how much Findavi does on its own

Observe. Findavi finds problems and recommends fixes. It changes nothing.

Approve. Findavi proposes the fix. A person says yes before it runs. That's the mode the cameras above were in.

Autonomous. Findavi acts on its own, within limits your team set in advance.

You choose the mode for each device and each action. Your team can stop or switch off any Findavi agent at any time. With a person's approval, Findavi can restore a device's previous settings. Every action is recorded.

Core idea: Close the path. Keep the building open.

Ask your team: If one camera acted strangely tonight, could we cut off just that camera? Or would we have to take down the whole system to be safe?

The Camera Still Has to Record

Blocking the attacker is only half the job.



CONCEPTUAL ARTWORK

Too many security teams call it a win the moment the bad connection is blocked. In a government building, that's halftime.

A camera can be safe and no longer recording. A door can be locked down and no longer open for the staff. A vendor link can be gone and the heating system can stop reporting its faults. An intercom can be protected and unable to reach the security desk.

Each of those failures is silent. Nobody notices until someone needs the thing that broke.

The service still has to work.

The Findavi workflow calls for two checks after each supported action. Some service checks require the system owner on site.

Notice who does what. Findavi checks what it can see on the network: the path is closed, the video arrives, the badge events land. Some checks need a person on site. Only someone standing at the door can confirm the door opened. Findavi records those checks too, and keeps the issue open until they're done.

Core idea: If the path is closed but the service isn't proven, the job isn't done.

Ask your team: After our last security change to a camera, reader or building system, what proved nothing quietly broke?

Two checks before the job is done.

DEVICE	SECURITY CHECK	SERVICE CHECK
Camera	Risky path is blocked. Old credentials fail, where tested.	New video reaches the approved recorder.
Badge reader	Unapproved connection is cut.	Events reach access control. The owner tests a valid badge at the door.
Gate / intercom	Unexpected connection is closed.	Calls reach the desk; the gate opens on approval.
Building controller	Vendor access is restricted to approved use.	Facilities checks the equipment and monitoring.
Alarm / sensor	Risky connection is removed.	Alerts reach the response team.

Ten Minutes With the Facility Director

She thinks in open doors, working cameras and a building that opens on time.



CONCEPTUAL ARTWORK

The facility director's questions are simple. Will the building open at seven? Will the staff get in? Will the cameras be recording when something happens in the lobby? Will the alarms reach the people who respond? Can a contractor still fix the heating on a cold morning?

If you have ten minutes with her, talk in those terms. Then show results she can check, in three levels.

Show results she can check.

Found means evidenced discoveries: devices missing from inventory, exposed control screens, unowned vendor access and unexpected connections. Credential findings need an authorized way to verify them.

Saved means a measured reduction against your own baseline. Count time no longer spent hunting for devices, chasing vendors, documenting incidents or checking fixes by hand. Value that time at your own labor rates.

Estimated means potential disruption avoided. It is a judgment, not a measured saving. Show every assumption and let the facilities team change it.

Keep those three apart. Found is measured. Saved is measured against your own numbers. Estimated should look like an estimate.

Then make one small ask. Choose one building or one tightly scoped system where trouble would hurt most: an entrance, a camera fleet, a building-control segment, a vendor login or a supported radio environment.

Agree on the scope, the owner and what proof will look like before the assessment starts. Ask for a device map, evidenced exposures, a ranked plan, supported actions under your approval rules, and a security check and service check for every proposed change.

The cameras kept watching.

At the government facility in Chapter 7, a person approved the block. Findavi closed the connection and checked that recording continued.

The attacker never got past the cameras. The cameras never stopped watching.

That is the result to ask for in your building: a closed attack path, a working service and a record of both checks.

Close the path. Keep the building open.

Sources, by chapter.

All sources checked October 12, 2026.

Letter

1. House Committee on Homeland Security, "Threat Snapshot: Cyber Threats Remain Heightened Amid Lapse in Information Sharing Authorities, Government Shutdown," October 31, 2025. At least 44 U.S. states reported cyber incidents affecting state and local government systems in 2025.

homeland.house.gov/?p=7684

Chapter 2

2. CISA, ICS Advisory ICSA-24-214-07, "AVTECH IP Camera," August 1, 2024. CVE-2024-7029. Commands can be injected over the network; known public exploitation; the vendor did not respond to CISA.

cisa.gov

3. CISA, ICS Advisory ICSA-25-162-10, "PTZOptics and Other Pan-Tilt-Zoom Cameras," June 12, 2025. Hard-coded passwords users cannot change; SSH or telnet cannot be disabled; government services and facilities listed among sectors. No known public exploitation at publication.

cisa.gov

4. FBI, Cyber National Mission Force and NSA, Joint Cybersecurity Advisory, September 18, 2024. A botnet linked to Flax Typhoon held more than 260,000 devices as of June 2024, including IP cameras and video recorders.

ic3.gov

Sources, by chapter.

Chapter 3

5. CISA, ICS Advisory ICSA-24-191-04, "Johnson Controls Software House C-CURE 9000," July 2024. Weak installer credentials could allow administrative access. No known public exploitation.

[cisa.gov](https://www.cisa.gov)

6. CISA, ICS Advisory ICSA-24-191-05, "Johnson Controls Software House C-CURE 9000," July 2024; updated January 16 and July 17, 2025. Incorrect default permissions could expose credentials used for application access. No known public exploitation.

[cisa.gov](https://www.cisa.gov)

Chapter 7

7. City of Saint Paul, presentation to the Minnesota Legislative Commission on Cybersecurity, August 27, 2025. Attack detected July 25, 2025; Minnesota National Guard cyber team deployed for 17 days.

lcc.mn.gov

8. Infosecurity Magazine, August 2025. The mayor confirmed the Interlock ransomware group was responsible and that the city would not pay. Officials have not said how the attackers got in.

[infosecurity-magazine.com](https://www.infosecurity-magazine.com)

Sources, by chapter.

9. Observed: a NetOps IQ deployment at a government facility. Confirmed by NetOps IQ's CTO. Published with the customer's approval. Agency and location withheld.

Findavi capabilities (Chapters 5, 7 and 8)

10. NetOps IQ, Findavi platform. On-site appliance, AI on the appliance, supported actions, Observe, Approve and Autonomous modes, separate security and service checks. Device and action support is confirmed for each deployment.

netopsiq.com/product.html

Start with one building.


Securing every facility starts with one area you understand completely.

The Findavi Blind Spot Assessment looks at one tightly scoped area: one public entrance and its door system, one camera fleet, one building-control segment, one vendor login or one supported radio environment. You pick.

Findavi listens through an approved monitoring connection in Observe mode. Device changes require separate authorization.

You get a map of the devices and connections in scope, evidence of exposures, named owners and a plan ranked by operational impact. For every proposed change, agree on the supported action, approval, security check and service check.

Then you decide what moves to Approve, what might run within predefined Autonomous limits, and what stays in Observe.



Visit netopsiq.com

Scan or click

Request an assessment

netopsiq.com/assessment

info@netopsiq.com

The blind spot, secured.

NetOpsIQ™

netopsiq.com