

The Building That Runs the Cloud

Protecting the systems that keep data centers and their AI running.

Eldad Matityahu

Founder and CEO, NetOps IQ

Published September 19, 2026

ILLUSTRATIVE

1:42 a.m.

A cooling controller starts talking to an address it has never contacted before.

The security team can cut it off in seconds.

But will the server hall still be cool at 2:00?

Inside the building.

A Letter From Eldad

How to Read the Stories in This Book

The Four Ideas This Book Hangs On

Chapter 1: The Building Is Part of the Cloud

Chapter 2: Ten Hours Without a Trade

Chapter 3: When the Cloud Goes Dark

Chapter 4: The Controls Outside the Firewall

Chapter 5: Six Doors Into the Uptime Chain

Chapter 6: See Without Touching the Hall

Chapter 7: Three Questions Before You Touch Anything

Chapter 8: Contain the Controller, Keep the Hall Cool

Chapter 9: The Alarm Has to Work Too

Chapter 10: Ten Minutes With the Operations Director

Closing: The Building Keeps Running

Notes

Start With One Critical System

People talk about the cloud like it floats.

It doesn't. It lives in buildings. Big ones, full of chillers, pumps, batteries, generators, sensors, cameras and doors that lock.

Those buildings keep the servers alive. And most of the systems inside them now have a small computer. Software. A login. A network connection.

Most of those systems were bought by facilities. They're maintained by a vendor. They sit on a network IT runs. And they protect an uptime promise that operations owns.

Security often can't see any of it.

That creates a terrible choice at the worst possible moment. A cooling controller does something strange at night. You can leave it alone and hope it's harmless. Or you can cut it off and hope the server hall stays cool.

Neither is good enough.

This book is about a third option. See every device that keeps the building running. Know what each one can reach and what breaks if it stops. Take only the action you approved. Then check two things: the bad path is closed, and the building is still doing its job.

This is the fourth book in our series. The first was about hospitals. The second was about warehouses and supply chains. The third was about water. Different buildings, same question.

Can we contain the risk without causing the outage?

I wrote this for the people who answer that question at 2:00 a.m.: data center operators, facilities leaders, security teams and the executives responsible for keeping the service running. For colocation providers and AI cloud companies, the same question applies to the building systems their GPU workloads depend on.

Eldad Matityahu
Founder and CEO, NetOps IQ

How to read the stories.

Every story carries one of three labels, so you always know what kind of evidence you're reading.

Observed. Something our team saw during a NetOps IQ pilot, assessment or controlled demonstration. Names and identifying details are removed.

Independent evidence. Public reporting or research from a named outside source. You'll find the sources in the notes at the back, with publication dates where available.

Illustrative. A realistic scenario we built to explain how the problem works. It is not a real incident, and we never present it as one.

One more rule. When we describe a real outage or a real attack, we never claim Findavi would have stopped it. We use real events to show the stakes. We use our own work to show the method.

Start with the trunk.

Every story in this book is a branch off four ideas.

One. The cloud lives in a building.

Chillers, batteries, generators, sensors and locked doors keep the servers alive. When the building fails, the digital service fails with it.

Two. Building systems have computers too. Most have software, a login and a network connection. Facilities or a vendor runs them, so security often can't see them.

Three. In a data center, the defense can cause the outage. Cut off the wrong device and the server hall heats up. So every response has to be narrow, and decided before the night it's needed.

Four. Done means two checks pass. The bad path is closed. The building is still doing its job. One without the other leaves the case open.

Chapters 1 to 4 show why these ideas are true. Chapters 5 to 7 show where the third one bites. Chapters 8 to 10 show the method that delivers the fourth, and what it's worth.

The Building Is Part of the Cloud

The servers are easy to see. The systems that keep them alive often are not.

Walk into a data center and listen.

Fans. Pumps. Air moving through the room. A generator running its weekly test outside. Doors clicking open and shut. Rows of servers with blinking lights. Screens glowing in the operations center.

What you can't hear is the conversation underneath.

The cooling controller telling a chiller what to do. A sensor reporting from the hottest corner of a server hall. The uninterruptible power supply, or UPS, reporting to the power monitoring system. The badge reader deciding who gets through the secure door. The camera watching the loading dock. The vendor connection waiting quietly for someone to log in from far away.

Every time you ask an AI a question, a cooling system somewhere works a little harder.

A data center is a building that keeps servers alive. The cooling plant holds the temperature. The UPS and generators carry the load when the grid drops. Badge readers and cameras keep the right people in and the wrong people out. Sensors and alarms warn the team before a drift becomes an outage.

When all of that works, nobody thinks about it. That's the whole idea.

Less room for error

An AI rack runs hot. Really hot. AI chips, called GPUs, pack a lot of heat into a small space. Many new systems pipe liquid right to the hottest parts and use air for the rest.¹⁰

Uptime Institute's July 2026 survey found typical racks getting denser slowly, while more operators now run some racks at 30 kilowatts or more. The shift is real, but every hall is different.⁹

Liquid cooling also adds more small computers to watch. Pumps move the coolant. A cooling unit hands the heat from the rack's loop to the building's cooling.¹² Sensors report leaks, temperature, flow and pressure.¹¹ Some connect through a controller instead of having their own login.

How much time do you get if cooling stops? It depends on the load, the cooling design and the backup. Know your own number before the night you need it.

For an AI cloud company, stopped work can mean computing time it can no longer bill. For a colocation provider, the cost depends on what its contracts promise. Use your own numbers.

The scope stays clear. This book is about the building around the GPUs. Findavi does not protect the GPUs themselves or the fast network between them. Support for each liquid-cooling device, and each action, is confirmed at your site first.

Five teams, one building

Five teams. Five budgets. One building that has to keep running.

The security team usually doesn't own these devices. Facilities owns the cooling plant. Data center operations owns uptime. A specialist vendor maintains the generator, the UPS or the building management system. Physical security owns the cameras and badge readers. IT runs the network they all share.

When a device belongs to five teams, its security can belong to none of them. The gaps hide in ordinary places. An old controller that can't be patched without a shutdown. A vendor account that still works years after the install. One password shared across a whole fleet of sensors. A management screen reachable from a network it should never touch.

None of these announce themselves. They sit there doing their jobs, until one of them becomes the way in.

Just facilities?

It's tempting to call a cooling controller or a badge reader "just facilities."

That's like calling the engine "just part of the truck."

A cooling device doesn't need to run an application to take one down. If it fails, gets changed by the wrong person, or gets cut off from the system it reports to, the temperature in a server hall starts climbing. The same goes for the systems that watch power, control access, raise alarms and coordinate an emergency.

A problem in the building becomes a problem for everyone using the cloud.

The question at 1:42 a.m.

ILLUSTRATIVE

Picture a quiet night.

A cooling controller starts reaching for an address it has never contacted before. It should only talk to the building management system, the approved monitoring tools and, during scheduled maintenance, one known vendor service. This address is none of those.

The security team sees the alert. They could cut the controller off right now.

Before acting, they need an answer.

If we isolate this controller, will the server hall still be cool at 2:00?

Without knowing what depends on the controller, the team faces two bad options. Act fast and hope nothing important breaks. Or wait for someone to investigate while a possible attacker keeps moving.

Teams can do better. Know what the device can reach. Decide ahead of time which actions are safe. Contain only what needs containing. Then check two things. Is the unauthorized path closed? Is the building system still doing its job?

In a data center, a security fix that causes an outage is just a different kind of outage.

Try this

Could your team answer these five questions by tomorrow morning?

Can you list every connected device that affects cooling, power, alarms, environmental monitoring or physical access?

Do you know who owns each one from a security point of view?

Do you know which vendor connections are active, which are approved, and which should have been shut off years ago?

If a cooling controller acted strangely tonight, could you contain it without losing readings, alarms or temperature control?

Could you prove the service still worked afterward?

If any answer is "I'm not sure," start there.

Ten Hours Without a Trade

A cooling problem can become a global business outage.

INDEPENDENT EVIDENCE

Thanksgiving night, 2025. Aurora, Illinois.

A chiller plant at CyrusOne's CHI1 data center failed and took multiple cooling units down with it. That building is the main hub for CME Group, the world's largest derivatives exchange.²

By Friday morning, markets around the world had stopped moving.

CME's trading was down for more than ten hours.¹ Futures on the S&P 500, the Nasdaq 100 and U.S. Treasuries stopped updating. Currency trading went dark. Crude oil, gold and even palm oil in Malaysia felt it. Traders in Asia and Europe were hit first, because the halt landed in their working day.

Inside the building, temperatures reportedly climbed past 100°F. The site had redundancy. It even used free outside cooling when temperatures dropped below 30°F, and that night was cold. It still wasn't enough.⁴

CME had a disaster recovery site in the New York area. It chose not to switch, because early information pointed to a short outage. The outage was not short.

Afterward, CyrusOne said it added more backup cooling at the site.³

Be clear about what this was

This was a mechanical failure. Nobody has reported it as a cyberattack, and we don't suggest it was one.

So why open a security book with it?

Because it shows exactly what's on the line. One cooling plant in one suburban building stopped a slice of the global financial system for most of a day.

Ten hours. One cooling plant. No hacker.

Now imagine the same chiller plant run by a controller that someone else can reach.

The lesson

A cooling controller is a part of customer uptime. Treat it with the same seriousness you give the servers it protects.

And remember what slowed the response that night: uncertainty. Nobody knew how long the problem would last, so nobody made the big call. Knowing what depends on each system helps the team make that decision sooner.

When the Cloud Goes Dark

The attack took one region. Uncertainty took the controls for every region.

IDCF CLOUD, OCTOBER 7, 2026

EAST JAPAN REGION 1

tesla SERVERS STOPPED	henry SERVERS STOPPED
pascal SERVERS STOPPED	joule SERVERS STOPPED
radian NO BREAK-IN FOUND	newton NO BREAK-IN FOUND

OTHER REGIONS

East 2 NO BREAK-IN	East 3 NO BREAK-IN	West 1 NO BREAK-IN
------------------------------	------------------------------	------------------------------

Every region: customer console switched off.

Source: IDC Frontier, October 8, 2026.

INDEPENDENT EVIDENCE

October 7, 2026. Around 3:40 a.m. in Japan.

Ransomware hits IDCF Cloud. It is run by IDC Frontier, a SoftBank company, and businesses and local governments across Japan run on it.¹⁴

Four zones in its East Japan Region 1 go down. Virtual servers stop. They will not restart.¹³

The count: 495 companies and local governments.¹³

Ibaraki Prefecture's government and police websites go dark. The prefecture puts up a temporary site so it can keep sharing disaster information.¹⁵

A software company called Six Apart has 31 servers stuck inside. A business marketplace can't be reached. Inventory and price-update services stop.^{15, 16}

Then comes the hardest line in IDC Frontier's update. The customer data in those four zones will be hard to retrieve or restore. Customers can rebuild only from backups they kept themselves.¹³

Screenshots posted by customers show the attacker bragging that it took seven minutes to get in. That is the attacker's claim, not a verified fact.¹⁴

Two outages, not one

Look closely and there are two outages here.

The first is the attack. Ransomware locked up the servers in four zones. No containment decision caused that.

The second is caution. IDC Frontier cut Region 1 off from the network. Then it switched off the customer management console in every other region too, where no break-in had been found.¹³

Customers in those regions still had running servers. They just couldn't touch them. To start or stop their own machine, they had to ask IDC Frontier to do it for them.¹³

That is not a criticism. With an attacker inside and the way in still unknown, it was the responsible call. It is what any team does when it cannot yet prove what is safe.

But it shows the price of not knowing.

The attack hit four zones.

The uncertainty reached every customer.

Be clear about what this was

This was an attack on cloud software, not on a building. Nothing public ties it to a cooling controller, a power monitor, a door system or any device Findavi watches. IDC Frontier has not said how the attackers got in.¹³

We don't claim Findavi would have stopped it.

So why is it in this book?

Because it is idea three, playing out in public: in a shared facility, the defense can cause the outage. When you can't see the safe boundary, you draw it wide.

The customers who were ready

One detail matters more than the rest.

Six Apart kept seven daily backups of every server on a different cloud, Google Cloud.¹⁷ The copy it used was taken at 1:00 a.m., less than three hours before the attack. By 9:00 p.m. the next day, all 31 servers were running again on another provider.¹⁸

Other customers found out that their data lived in one place. And that place was gone.

The lesson

One lesson for each side of the contract.

If you run the facility: know your safe boundary before the bad night. Which devices, paths and controls can be isolated on their own? Which must stay on? Who can approve the change? Decide it on a calm Tuesday, not at 3:40 a.m.

If you rent the cloud: your provider's bad night is your bad night. Keep a backup somewhere else. Then prove you can rebuild from it.

The rest of this book is about the first lesson, for the building systems underneath. Small actions. Approved in advance. Checked twice.

The Controls Outside the Firewall

The easiest way in may be through the room around the server.

INDEPENDENT EVIDENCE

In 2026, researchers at TrendAI asked a simple question. How many industrial and building control devices can anyone on the internet reach, right next to America's data centers?

They looked within one kilometer of 1,063 U.S. data centers. After five rounds of filtering, they found about 6,300 devices they rated high-confidence. Building automation controllers. Heating and cooling systems. Power management screens. All answering on the public internet.⁵

The mix tells the story. BACnet devices, the language of heating, cooling, lighting, access and fire systems, made up 58 percent. Niagara building platforms made up another 23 percent. The list also included equipment built specifically for data center power and cooling.

The researchers also found 548 of the exposed building devices were running products past end of life. No more security fixes are coming for them, ever.⁶

One more surprise. Data centers permitted in 2021 or later showed more nearby exposure than older ones, about 13 percent versus 5 percent. The researchers didn't say why. But the trend points the wrong way.

What this research does not prove

We want to be careful here, because the researchers were.

It does not prove every device belonged to a data center. Location by internet address is rough, and some devices may sit in a neighbor's office building. It does not show that any data center was attacked.

What it does show is a pattern. A lot of controls for cooling, power and building systems are reachable from anywhere, in exactly the places where uptime matters most. And since the study left out smaller private facilities, the researchers called their count a floor.

Why it matters to you

An exposed port is a technical detail. An exposed path to the system that keeps your server hall cool is a business risk.

Ask both questions: "Is this device on the internet?" and "If someone reached it, what could they change, and what would stop working?"

You can only answer that if you can see the device, its connections, and what depends on it.

Six Doors Into the Uptime Chain

An attacker doesn't need your server rack if they can reach what keeps it running.

The front door has guards, badges and cameras. **The thermostat has none.**

The doors in this chapter are like that. Nobody guards them, because most people don't know they're doors.

Door one: The vendor who never left

The HVAC contractor needed remote access during the install. So did the UPS vendor and the generator service team. Years later, those connections are still live. Some still use the original password. Nobody remembers who approved them.

A maintenance path becomes an unmanaged way in.

Door two: The controller on the internet

A building controller gets a public address so a technician can check it from home. Now anyone can find it. Chapter 4 showed how common that is.

Direct reach to the systems that control the room.

Door three: The shared password

Two hundred sensors. One password. It came from the factory, or from the installer who set them all up in an afternoon.

One stolen credential opens the whole fleet.

Door four: The flat facilities network

Cooling, power monitoring, cameras and badge readers all share one network. Sometimes it touches the office network too. Every device can talk to every other device.

One weak camera becomes a stepping stone to the cooling plant.

Door five: The alarm nobody checked

Security blocks a suspicious connection. Good. But that connection also carried temperature readings to the monitoring screen. Now the room looks calm because the numbers stopped coming.

The team loses sight of the building while believing everything is fine.

Door six: The fix that caused the outage

Someone isolates a controller in a hurry. It turns out that controller managed half the cooling units in Hall B. The attacker is out. The hall is heating up.

The security response becomes the uptime incident.

Closing the doors

Look closely at doors five and six. The danger there comes from the defense itself. That's what makes data centers different from an office.

One rule runs through the rest of this book. A response is complete when two things are true. The unauthorized path is closed. And the physical service still works.

Blocking an address is step one, never the finish line.

See Without Touching the Hall

In a high-uptime building, looking too hard can be its own risk.

Most security tools find devices by poking them. They send probes, scan ports and ask each device what it is.

On a laptop, that's fine. On an old cooling controller, it can be a problem. Some building and industrial devices were never built to handle unexpected traffic. A scan meant to protect them can slow them down, confuse them, or knock them offline.

In a data center, nobody wants to find out which ones.

Listen first

Findavi takes a different approach. It connects through a passive monitoring port and listens. It never sends a single packet to the devices it's learning about.

From that traffic, it builds a record of what it can see. Which devices are on the network. What they are. Who they talk to. How often. Which communication methods they use. It can also connect to an RF switch, which lets it hear wireless links too.

Then it learns what normal looks like. A cooling controller should talk to its building management system, its system for recording readings and its approved management tools. A badge reader should talk to its access server. When either one reaches somewhere new, that stands out.

That's how you catch the 1:42 a.m. moment. A new destination. A vendor session outside its window. A device that suddenly reaches into a zone it never touched before. A brand new device nobody registered.

Your data stays home

Data center teams ask a fair question about any new box on their network. What does it send out?

Findavi's analysis runs on the device itself. The AI runs locally, with no outside AI service. The traffic it learns from is stored on the box, inside your network. Nothing goes to the cloud to be analyzed.

Be honest about the view

Passive listening shows the traffic that reaches the monitoring point.

It can't see a device that never sends traffic past the monitoring point. A cellular modem bolted to a generator may never touch your wired network at all.

So every NetOps IQ assessment says three things plainly. What we saw. What we couldn't see. And what needs a closer look before anyone acts.

Three Questions Before You Touch Anything

Not every device problem deserves the same response.

A suspicious camera in the parking lot and a suspicious cooling controller in Hall B are both alerts. They are very different decisions.

Before anyone acts, three questions decide what happens next.

Question one: What can it reach?

Can this device reach the building management system? The power monitoring system? The tenant network? The internet? Another facilities zone?

A device with few connections may give an attacker fewer places to go. But it can still be critical if it controls cooling, power or an alarm. Check what it does as well as what it can reach.

Question two: What breaks if it fails?

Could losing this device cut cooling capacity? Hide power readings? Silence an alarm? Lock staff out of a room? Delay a fire response? Break a tenant's service?

This is the question security teams usually can't answer alone. Facilities can. That's why the answer needs to be written down before the night it's needed.

Question three: Can we safely change it right now?

Is there a backup unit carrying the load? Is this a maintenance window? Does a facilities engineer or the vendor need to be on site? Has this exact action been approved ahead of time?

If the answers are unclear, the right move may be to watch closely and call the right person. If they're clear, act fast.

Why process is uptime

INDEPENDENT EVIDENCE

Uptime Institute's 2026 outage analysis found that power is still the leading cause of serious data center outages. UPS systems, transfer switches and generators lead the list.⁷

The same research found something just as important. When people cause outages, the most common reason is that someone didn't follow an established procedure.

That matters for security. A rushed, unplanned security action near power or cooling is exactly the kind of change that causes outages.

Write down the approved response while the team has time to review it. At 2:00 a.m., that gives the person on call a plan to follow.

Contain the Controller, Keep the Hall Cool

A secure controller that no longer runs the cooling is a failure.

This chapter shows how Findavi handles the 1:42 a.m. moment, step by step.

The night, replayed

ILLUSTRATIVE

At 1:42 a.m., Findavi sees the cooling controller reach for an address it has never contacted. It compares that with the controller's normal pattern. Building management system, yes. System for recording readings, yes. Approved vendor service, only during its window. This new address matches none of them.

Findavi checks the site's rules for this device. Facilities and security agreed on them weeks ago. For this controller, the rule says: restrict, don't disconnect, and ask a person first.

It recommends one narrow action. Keep the controller connected to its building management system, its monitoring tools and its management console. Block everything else.

The on-call engineer gets the request on her phone at 1:44. She reads the evidence and approves at 1:46.

Findavi pushes the change.

Then it runs the security check. Can the controller still reach the unknown address? No. The path is closed.

Then it runs the service check. Is the controller still sending temperature readings? Yes. Are the alarms still live? Yes. Is it still talking to the building management system? Yes.

At 2:00, the hall is right where it should be.

Findavi records all of it. What it saw, what it recommended, who approved, what changed, both checks, and anything left to investigate in the morning.

Contain the controller. Keep the data center.

What Findavi can actually do

The available actions depend on the equipment and the approved connections at your site.

For supported devices, Findavi can block a device, isolate a network port, or change a setting. We confirm the supported models, software versions and connections at your site before enabling an action. Each action must follow your rules and pass both a security check and a service check.

Those changes cover the connections and devices around your building controls, like passwords and network access. Findavi doesn't change how the cooling plant runs. Control and safety decisions stay with your facilities team.

You can stop, undo or switch off any agent at any time. Every action is recorded so you can audit it.

Three modes, your choice.

Every site, and every device type, gets the mode its owners are comfortable with.

Observe. Findavi watches, learns and reports. It changes nothing. Most sites start here.

Approve. Findavi recommends a specific action with its evidence. A named person says yes or no.

Autonomous. For narrow situations you've approved ahead of time, Findavi acts on its own, then runs both checks and reports.

A camera in the parking lot might run in Autonomous. A cooling controller might stay in Approve for a long time. Both are good answers. You decide.

From our own work

OBSERVED

In a controlled demonstration for a logistics customer, Findavi detected and stopped an attack on a heating and cooling system.

In a pilot at a live water site, our team ran a simulated attack on the plant's controller screens. Findavi blocked the attacking connection and isolated a port. Both the security check and the service check passed. The controls kept running for the operator.⁸

OBSERVED IN AN AUTHORIZED TEST, POWER GENERATION

In an authorized test at a power-generation facility, NetOps IQ recorded a blocked connection and an available camera feed. The notes do not confirm uninterrupted service or successful recording. This result covers the equipment and scope tested.

Those tests were in logistics, water and power generation, not an AI data center. They show the method: close the attack path, then check that the service still works. They do not establish support for a liquid-cooling pump, leak sensor or controller. We must confirm support for the equipment at your site.

The Alarm Has to Work Too

Uptime depends on what you can see as much as what's running.

A chiller can be humming along while the screen that watches it has gone blank.

That's a dangerous place to be. If temperature readings stop arriving, nobody knows the room is warming until a server throttles itself. If an alarm path breaks, the first warning may be a tenant on the phone.

So in a data center, the service check has to cover more than "is the device running?" It has to ask whether the team can still see it, hear from it and reach it the way they need to.

Two checks, every time

Every Findavi action ends with two checks. Each one has to pass on its own.

Two checks. Both must pass.

After this action	Security check	Service check
Restrict a cooling controller	The unknown address can't be reached	Temperature readings, alerts and the link to the building management system still work
Shut off an unapproved vendor session	The session is blocked and can't reconnect	The approved emergency access process still works
Isolate an access control device	The attack path is contained	Authorized staff and emergency exits keep the access they need
Push a new device setting	The old weak setting is gone	The device reconnects normally to its approved systems

When to keep the case open

If either check fails, the issue stays open. If either check can't be proven, the issue stays open too.

That sounds strict. It should be. "We think it's fine" is how door five happens.

A closed ticket should mean two things you can show anyone. The bad path is gone. The building is still doing its job.

Ten Minutes With the Operations Director

Skip the security budget pitch. Ask what one hour of uncertainty costs.

Operations directors want one thing from security: fewer bad nights.

So start the conversation with the night CME's markets stopped. More than ten hours. Trading halted across stocks, bonds, currencies and commodities, all from one cooling plant. Then ask the director one question. What would ten hours cost us?

Count the real cost

INDEPENDENT EVIDENCE

Uptime Institute's 2025 survey found that 57 percent of operators said their most recent major outage cost more than \$100,000. One in five said the last outage that hurt their business cost more than \$1 million.⁷

Your own number will be different. Work it out together. Start with the availability promises in your contracts and the service credits they trigger. Add the hours your customers would feel it. Add revenue or transactions at risk. Add staff time and emergency vendor call outs. Add the risk to hardware from a heat event. Then add the hardest one to price: trust, and what it does to renewals.

Last, add the cost of the broad shutdown. That's the outage you cause when you can't contain one device, so you take down a whole zone to be safe.

Three kinds of value

When you report results to the director or the board, keep three kinds of numbers apart. Mixing them is how security reports lose credibility.

Found. What the assessment measured. Devices nobody had listed. Exposed management screens. Unexpected connections. Devices with no clear owner. Every one of these is counted.

Saved. Time your team no longer spends hunting for device owners, chasing alerts, coordinating vendors and writing up incidents. Measured with your own labor rates.

Reduced risk. Downtime and incident scope you likely avoided. This one is an estimate, and it should always say so, with the assumptions shown.

Keep the report clear about what was measured and what was estimated. That gives the operations director something they can check.

The building keeps running.

ILLUSTRATIVE

Back to 1:42 a.m.

The controller reached for an address it had no reason to contact. The path was closed. The controller kept reporting to the building management system. The temperature alarms stayed live. The cooling plant kept running.

Nothing dramatic happened. No tenant called. No workload moved. No outage report got written.

That's what good security looks like in a data center. The risk is contained, and the building keeps doing its job.

Every AI answer depends on a building that stays cool. Keep the building safe, and the cloud stays up.

Eldad Matityahu

Founder and CEO, NetOps IQ

Sources behind the stories.

All sources checked October 10, 2026.

CHAPTER 2 / TEN HOURS WITHOUT A TRADE

1. CME trading halt after a cooling issue at CyrusOne data centers, with markets affected across currencies, Treasuries, commodities and stocks.

Reuters, November 28, 2025

<https://www.reuters.com/business/cme-trading-halted-due-cooling-issue-data-centers-2025-11-28/>

2. CyrusOne statement that its CHI1 facility experienced a chiller plant failure affecting multiple cooling units on November 27, 2025.

TechRepublic, December 1, 2025

<https://www.techrepublic.com/article/news-global-futures-cme-data-center-cooling-failure/>

3. Outage of more than ten hours, temperatures above 100°F, CME's decision not to switch to its New York area backup site, and CyrusOne bolstering backup cooling.

Bloomberg via Insurance Journal, December 1, 2025

<https://www.insurancejournal.com/news/midwest/2025/12/01/849319.htm>

4. Location of the data center complex in Aurora, Illinois, and its redundancy and free cooling design.

Fortune, November 30, 2025

<https://fortune.com/2025/11/30/cme-outage-financial-markets-data-center-cyrusone-aurora-illinois>

CHAPTER 4 / THE CONTROLS OUTSIDE THE FIREWALL

5. About 6,300 high-confidence internet-accessible ICS and BAS devices within one kilometer of 1,063 U.S. data centers, protocol mix, facility age findings, and study limits.

TrendAI Research, "An Invisible Attack Surface," 2026

<https://www.trendaisecurity.com/en-us/resources-insights/deep-research/an-invisible-attack-surface-thousands-of-industrial-control-systems-exposed-near-data-centers>

6. 548 exposed building automation devices running end-of-life products.

Cyber Press, 2026

<https://cyberpress.org/548-building-devices-unpatched/>

CHAPTERS 7 AND 10 / UPTIME AND COST

7. Power remains the leading cause of impactful outages; failure to follow procedures leads human error outages; outage cost findings from the 2025 survey.

Uptime Institute, Annual Outage Analysis 2026, May 13, 2026

<https://www.businesswire.com/news/home/20260513289344/en>

CHAPTER 8 / NETOPS IQ OBSERVATIONS

8. HVAC demonstration and water site pilot results come from NetOps IQ's own work, with customer details removed.

NetOps IQ

AI data centers: sources.

CHAPTER 1 / LESS ROOM FOR ERROR

9. Typical rack density and high-density peaks.

Uptime Institute Global Data Center Survey 2026, July 24, 2026

<https://intelligence.uptimeinstitute.com/resource/uptime-institute-global-data-center-survey-2026>

10. Hybrid air and liquid cooling in an AI rack system.

NVIDIA, Key Components of the DGX SuperPOD

<https://docs.nvidia.com/dgx-superpod/reference-architecture-scalable-infrastructure-gb200/latest/dgx-superpod-components.html>

11. Leak detection, flow, pressure and temperature measurements.

NVIDIA, Mission Control Software Administration Manual, section 12

<https://docs.nvidia.com/base-command-manager/manuals/11/nvidia-mission-control-manual.pdf>

12. How the rack and building cooling loops connect.

NVIDIA, Maximizing AI Factory Performance per Watt with NVIDIA DSX MaxLPS

<https://developer.nvidia.com/blog/maximizing-ai-factory-performance-per-watt-with-nvidia-dsx-maxlps/>

The cloud goes dark: sources.

CHAPTER 3 / WHEN THE CLOUD GOES DARK

13. IDC Frontier's third report: time of the attack, the four affected zones, 495 customers, data recovery only from customer backups, consoles stopped in other regions, intrusion path under investigation.

IDC Frontier, October 8, 2026 (Japanese)

<https://www.idcf.jp/news/topics/20261008001>

14. Ransomware confirmed in East Japan Region 1; IDC Frontier is a SoftBank Group subsidiary; the attacker's claim of a seven-minute breach.

BleepingComputer, October 2026

<https://www.bleepingcomputer.com/news/security/ransomware-attack-disrupts-japans-idcf-cloud-used-by-govt-clients/>

15. Ibaraki Prefecture and prefectural police websites, the NETSEA marketplace, and Six Apart's affected servers.

iThome, October 2026 (Traditional Chinese)

<https://www.ithome.com.tw/news/179548>

16. Six Apart's 31 affected servers and stopped inventory and price-update services.

Impress Watch, October 8, 2026 (Japanese)

<https://www.watch.impress.co.jp/docs/news/2146693.html>

17. Six Apart keeps seven daily backups of every server on Google Cloud, separate from IDCf Cloud.

Six Apart, second report, October 8, 2026 (Japanese)

<https://www.sixapart.jp/movabletype/news/2026/10/08-0900.html>

18. Six Apart restored from its 1:00 a.m. October 7 backup and moved all 31 servers to Sakura Cloud by 9:00 p.m. October 8.

Six Apart, third report, October 9, 2026 (Japanese)

<https://www.sixapart.jp/movabletype/news/2026/10/09-0900.html>

Start with one critical system.

Securing an entire data center starts with one zone you understand completely.

The **Findavi Uptime Blind-Spot Assessment** looks at one tightly scoped area. One server hall. One cooling zone. One building management segment. One generator and UPS monitoring environment. One vendor remote access path. One access control or camera system. You pick.

Findavi listens through a passive port in Observe mode. It changes nothing on your network.

At the end, you get a map of the connected devices and how they talk. You get an ownership map that shows what each device keeps running. You get the unexpected and exposed paths worth investigating, ranked by what matters most. And for every action we propose, you get the security check and the service check that would prove it worked.

Then you decide what moves to Approve, what might one day run in Autonomous, and what stays in Observe.



Request an assessment

Scan or click

e-Mail: info@netopsiq.com

The blind spot, secured.

NetOpsIQ™

netopsiq.com
