UPGRADING SECURITY

Moving to Cloud-Based SCADA

A guide for water systems deciding whether to modernise — and how to buy it well

---

If you have already worked through the Security Self-Audit Workbook, you have a list of findings and you have fixed the ones that were free. Some of what is left cannot be fixed with a setting change. This guide is about those.

If you have not, you can still read this — but the audit is a better first step, and it is free.

THIS GUIDE NAMES NO PRODUCTS

Not even ours. Everything here is written as capabilities to require, so that any solution meeting them qualifies. We would rather raise the standard for every bidder than steer you toward one — including in the bid you may eventually put us in.

Prepared by:

1-844-365-8647

info@d6labs.com

Contents

Before You Start

Who this is for

Operators, superintendents, managers and board members at community water systems. Most utilities arrive at a document like this for one of five reasons.

- **Our equipment is obsolete.** The manufacturer stopped making it. Spares come from used inventory on eBay, and there are no more firmware updates — including for security problems that are already public and will never now be fixed.
- **We do not want to be the next Minnesota.** We need a system that is genuinely secure, and we have no IT department and no budget to keep one under contract.
- **Our system has never worked right.** People come out, bill us, and the same problems come back. Nobody has ever got to the bottom of it.
- **We have a site with constant radio trouble** that nobody has been able to fix.
- **We did a security audit and found more than we can handle.** We do not have the staff or the experience to close the findings, and the ones left do not have a cheap fix.

What this is

An explanation of what cloud-based SCADA actually is, how it addresses the problems above, and practical guidance on how to specify and buy it so that what you get is genuinely secure rather than merely modern.

What this is not

- **Not a sales brochure.** No product is named, including ours.
- **Not an engineering specification.** Part 8 tells you what to require and why, in plain language. Exact contract wording belongs with your engineer — there is a companion document written for them.

Why there are no product names in here

We build SCADA equipment, so it is fair to ask why a guide from us does not mention it.

The reason is practical rather than noble. No single company can supply every water system in this country, and it would not be good for you if one could. What we would rather have is a market where the minimum standard is high — where whoever wins your bid has to meet real security requirements, because you knew to ask for them.

So this guide describes capabilities, not brands. Some of the requirements in Part 8 are easier for some suppliers to meet today, and that will change as the rest of the industry catches up. That is the intended outcome, not a problem.

The practical effect for you: you can hand Part 8 to any bidder, including us, and judge the answers on the same terms.

How to use it

If you are…	Read
Trying to understand what cloud SCADA is	Parts 2, 3 and 4
Building the case for your board	Parts 6 and 10, plus Appendix C
Ready to go to market	Parts 7, 8 and 9, plus Appendices A and B
The one who has to live with it afterwards	Parts 4 and 9

The Short Version

If you read nothing else.

The problem

Most small water systems run SCADA on a Windows computer, talk to their sites over unencrypted radio, and reach the system from outside using remote-access software. Every publicly reported attack on a US water system in recent years went through one of those three things. You can harden all of them — and you should — but you cannot make an internet-reachable, password-protected, Windows-dependent system stop being those things. And if the manufacturer no longer supports your equipment, no amount of hardening will ever produce another firmware fix.

What cloud SCADA changes

Control still runs in the controller at the site, locally, always. What moves is everything between the controller and your eyes: the screens, the history, the alarms, the configuration. Done properly, that lets the system connect outward only — so there is nothing sitting on the internet waiting to be found — and lets credentials be certificates that can be revoked rather than passwords that can be shared.

What you get beyond security

- No SCADA computer to patch, licence or replace, and no upgrade project to defer
- Controllers that update themselves, and that can be swapped by an operator in minutes rather than by an engineer in days
- Screens, alarms and reports you can change yourself without a service call
- Access from a browser or a mobile app on any laptop, tablet or smartphone, anywhere in the world, around the clock
- Cellular instead of licensed radio: no path studies, towers, repeaters or FCC licences

What it costs

Two parts: the installed system, once, and a subscription, annually. A properly designed cloud system should come in anywhere from about thirty per cent below a conventional SCADA installation to roughly comparable with it. Where it separates decisively is afterwards — operator-performed device swaps, automatic program updates and self-service screens and alarms remove most of the recurring cost a conventional system generates. Part 6 has indicative figures as of August 2026 and a worksheet for your own numbers.

Where to start

Confirm coverage and get a properly scoped quote. A competent supplier can establish, before anyone visits, whether the cellular bands the equipment needs are available at each of your sites — that is their

job rather than yours, and Part 4 explains why your mobile phone is not a reliable test. Then use Part 8 to write real requirements into your bid, so whoever wins has to meet them.

Part 1 · Why Hardening Has a Ceiling

If you completed the self-audit workbook, you fixed the free things: removed accounts belonging to people who had left, took the SCADA computer off email, turned on two-step login, changed default passwords, locked the panels. Those repairs matter and they prevent most opportunistic attacks.

This part is about the findings that were left over — the ones with no afternoon-sized fix.

Everything you hardened has to stay hardened

A hardened conventional system is not a state you reach. It is a state you maintain, indefinitely, with the crew you have.

Consider how a well-secured system drifts back over three years. A contractor is given access for a project and nobody removes it afterwards. A port is forwarded so a vendor can troubleshoot on a Saturday, and it stays open. A radio is replaced and the new one ships with encryption off. Somebody retires and their login lives on because it is tied to a scheduled task nobody wants to break. A patch is deferred during a busy season and then forgotten.

None of that is negligence. It is what happens when security depends on a small number of people getting every decision right, every time, forever, on top of running a water system.

Equipment you cannot buy any more

A great many water systems are running controllers, radios and operator panels the manufacturer discontinued years ago. The consequences build up quietly:

- Spares come from used inventory on auction sites, with no warranty and no way to know how a unit was stored or why it was pulled from service
- There are no more firmware updates — including for security flaws that are publicly documented and will never now be fixed
- The programming software may not run on any operating system you are still allowed to use
- The number of people who know the platform shrinks every year, and their rates go up

There is no hardening answer to this one. A device that cannot receive a firmware update cannot be protected against a flaw in its own firmware, and no amount of network design around it changes that.

The Windows problem, and the part nobody has joined up

Three facts, which most operators know separately and few have put together:

- **Windows 10 support ended on 14 October 2025.** Machines still running it no longer receive security updates.
- **Extended Security Updates are available, at an escalating price.** Roughly $61 per device for the first year, doubling each year after, and the programme ends entirely in October 2028. It is a delay, not a solution.

- **Newer SCADA software increasingly requires Windows 11,** which requires TPM 2.0, Secure Boot and a recent processor. A large share of SCADA computers in service physically cannot be upgraded.

PUT THOSE TOGETHER

Upgrading your SCADA software means a new version. The new version needs Windows 11. Windows 11 needs hardware your current machine may not have. So "we should update the SCADA software" quietly becomes new computer, new operating system, new software version, and the labour to move everything across.

The upgrade nobody actually does

A major SCADA software upgrade on a conventional platform commonly runs $15,000–20,000 once labour is included, because the installation is genuinely difficult and time-consuming. What it buys is parity: the same functionality you already had, on a version that is still supported.

Most small systems do not do it. The money is not there, it is disruptive, and the system still works. That is not a criticism — it is the rational decision at the time, made repeatedly.

But deferring does not delete the cost. It converts a budgetable expense into an unbudgetable risk, and it moves the decision from you to circumstance: the upgrade eventually happens when a motherboard dies, or when a regulator asks, or after an incident.

The truck-roll economy

On a conventional system, most changes require somebody to physically attend. A new alarm, a screen that no longer matches the plant, an added tag, a failed controller. Typical service charges are a $500 minimum and $200 per hour and up.

This is structural rather than a fault of your integrator. The architecture requires a person with a laptop standing next to the equipment, so the cost of a small change has a floor set by travel time.

The compounding effect is worse than the invoices. Because changes cost money and take scheduling, they do not get made. Alarms nobody could tune get switched off. Screens drift out of date with the plant. Operators stop trusting what is displayed and go back to the clipboard — and a system nobody trusts provides no protection at all.

Where the ceiling actually is

YOU CAN MAKE IT MUCH SAFER. YOU CANNOT MAKE IT SOMETHING ELSE.

A conventional SCADA system can be hardened a long way. What hardening cannot do is change these facts about it:

- Something has to listen for inbound connections, or you cannot reach it from outside — and anything that listens can be found and attacked.
- Credentials are passwords, and passwords can be shared, written down, reused, guessed and inherited by whoever comes next.
- The system depends on a general-purpose computer with an operating system somebody must keep patched forever.
- Radio between sites is a broadcast; unless every link is encrypted and stays encrypted, it can be listened to and injected into from outside your fence.
- Equipment the manufacturer no longer supports will never be patched again, whatever you build around it.

If your remaining audit findings are mostly in that list, you have reached the ceiling. That is the point at which it is worth asking what a different architecture looks like — which is Part 2.

Part 2 · What Cloud SCADA Actually Is

Before anything else, one misconception has to go, because every argument later in this guide is unbelievable while it is still in place.

YOUR PLANT DOES NOT RUN ON THE INTERNET

In a properly designed cloud SCADA system the control logic runs in the controller, at the site, locally, all the time. The cloud handles supervision, history, screens, alarms and configuration. Your pump does not wait for a network round trip to decide whether to start.

The two architectures, side by side

	Conventional	Cloud
At the site	PLC running the control logic	Controller running the control logic — same job
Between sites	Licensed radio, usually unencrypted; sometimes cellular with a port forward	Encrypted cellular, connecting outward only
Where the screens live	SCADA software on a Windows computer in your office	A hosted platform, reached with a browser or a mobile app
How you see it remotely	Remote-access software, or a VPN into the office network	Browser or mobile app on any laptop, tablet or smartphone, anywhere in the world, around the clock, each person with their own login
What listens for connections	The office computer, the remote-access tool, any forwarded port	Nothing at the utility
Who patches the software	You do, or your integrator, when there is time	The provider, continuously
What fails if the office burns down	Everything, including your history	Nothing at the sites; you open a browser somewhere else

Getting to it from anywhere

Worth stating plainly, because it is one of the differences operators notice first. Access is through a browser or a mobile app, on any laptop, tablet or smartphone, from anywhere in the world, twenty-four hours a day. There is no software to install on one particular machine, no remote-access tool listening at the office, and no single computer that has to be switched on before anyone else can see the plant.

It also means more than one person can be looking at once, each with their own login — which matters at three in the morning when the operator on call and the superintendent both need to see the same screen.

What happens when communications drop

This is the question everyone asks, and it deserves a proper answer rather than reassurance. Here is the sequence in a correctly designed system.

When	What happens
Immediately	Control continues. The controller was already running the logic and does not need the link to keep doing it. Pumps run, valves operate, setpoints hold, interlocks and failsafes work exactly as before.
Within seconds	The controller notices the link is down and begins storing readings locally instead of sending them.
While it is down	You lose remote visibility and remote control — live values and setpoint changes from your phone.
On reconnection	Stored readings are forwarded, so your history has no hole in it. Remote visibility and control return.

Sites that depend on data from other sites

A pump station that normally starts and stops on tower level has a genuine dependency: if it cannot hear the tower, it needs to know what to do. A well-designed system does not simply stop.

The answer is an automatic failover control mode. A pump station that loses tower level falls back to controlling on its own discharge pressure — starting and stopping to hold pressure in the system rather than level in the tank. The tower may end up higher or lower than you would choose, and if the pressure setpoints are wrong it could overflow, but the distribution system stays pressurised and your customers keep water.

Ask any supplier this directly: "what does each site do if it loses communication with the site it depends on?" The answer should be a specific, configured behaviour for every dependency — not a shrug, and not an assurance that the link never fails.

Worth comparing with what happens today: on most radio systems a lost link between a tower and a pump station produces exactly the same problem, frequently with no failover at all.

What it replaces, and what it does not

It replaces	It does not replace
The SCADA computer	Your control logic
The SCADA software licence	Interlocks and failsafes
The server, UPS and cooling	Local operator interface where fitted
The backup regime	Physically maintaining equipment
The upgrade project	Your judgment and your licence
The version you are behind on	Your responsibility for the system

A plain-language decoder

Vendor vocabulary gets in the way, and most people are too polite to stop the meeting and ask. Here is what the words mean.

Term	What it actually means
PLC / RTU / controller	The box at the site that runs the control logic. Different words, broadly the same job; "controller" is the modern general term.
HMI	The screen an operator looks at — a panel at the site, or a page in a browser.
SCADA	The whole arrangement: controllers, communications and the supervisory screens.
Gateway	A device passing data between a site and somewhere else. In some designs a separate box; in others the controller does it.
Historian	The database storing your readings over time so you can look back.
Tag	One value monitored or controlled — a level, a pressure, a pump run command.
MQTT / MQTTS	A messaging method suited to sending data efficiently over cellular. The S means it runs inside an encrypted connection.
TLS	The encryption protecting data in transit — the padlock in a web browser.
Certificate	A cryptographic credential proving a device is what it claims to be. Not a password: it cannot be read over the phone or written on a panel door, and it can be revoked individually.
Outbound-only	The device makes connections out; nothing can connect in. No door on the outside of the building.
CAT M1	A 4G cellular technology built for equipment rather than phones — long reach and low power rather than high speed. See Part 4.
Edge	Processing done at the site rather than centrally.
Segmentation	Separating the control network from the office network.

What does not change

Your process. Your pumps, valves and instruments. Your treatment decisions. Your licence. Your responsibility for the water.

Part 3 · The Security Case, in Detail

This part explains the mechanisms rather than asserting the benefits, because the mechanisms are what you will need in order to judge a supplier's claims in Part 8.

What the attacks actually exploited

Briefly, because you already know the story. Across the incidents that have been made public — including the coordinated attacks on more than thirty Minnesota systems in July 2026, and earlier events in Florida, Pennsylvania and Kansas — attackers did not defeat sophisticated protections. They used:

- Controllers reachable from the public internet, found by scanning and then accessed with the manufacturer's own engineering software
- Passwords that were still the factory default, or shared across a crew, or belonged to somebody who had left
- Consumer remote-access software with a single shared login and no second factor
- Office and control networks with nothing between them

Every one of those is a property of the architecture as much as of the utility. The four principles below remove the property rather than patching the instance.

The four principles

Principle 1 · Outbound-only connectivity

The mechanism. The equipment at your sites opens connections outward to the platform and keeps them open. Nothing at your utility listens for incoming connections. There is no port to forward, no service waiting for a login attempt, and no address that answers when somebody probes it.

What it kills. Discovery by internet scanning, and every attack that begins with finding an exposed device. A scanner sweeping the internet finds nothing to talk to. This is the biggest structural difference, because it removes the first step of the most common attack rather than making that step harder.

What to watch for. "Outbound-only" is sometimes claimed for systems that do listen on a port with a firewall in front of it. That is not the same thing: a firewall is a rule that can be misconfigured, and something is still answering behind it. Ask whether anything at the site accepts inbound connections at all.

Principle 2 · Encrypted transport between sites

The mechanism. Data between your sites and the platform travels inside an encrypted connection, normally TLS — the same technology that protects online banking. Some designs add a second layer, encrypting the message itself as well as the channel carrying it, so intercepted traffic yields nothing usable even if the outer layer were broken.

What it kills. Interception and injection. This is the direct answer to unencrypted radio, where somebody within a few miles can listen to your telemetry and send commands your controllers accept as genuine — with no alarm, no log entry and no trace.

What to watch for. Ask whether encryption is on by default and cannot be switched off, or whether it is an option somebody has to enable. A great many radios in service today support encryption and shipped with it turned off.

Principle 3 · Certificate-based, revocable credentials

Worth understanding properly, because it is the least familiar and it changes the most about day-to-day operation.

The mechanism. Instead of a password, each device holds a certificate — a cryptographic credential issued when the device is deployed. The device proves its identity with that certificate and the platform proves its identity back. The credential is never typed, never spoken, and never travels in a form anybody could reuse.

A password	A certificate
Can be read aloud, texted, or written inside a panel door	Cannot be usefully copied down or communicated
Is often the same across many devices	Is unique to one device
Leaves with anyone who knew it	Stays with the device it was issued to
Changing it means visiting or reconfiguring everything that used it	Revoking one is an administrative action that takes effect immediately
You rarely know who else knows it	You can list exactly what is trusted, and withdraw any of it

What it kills. Default passwords, shared crew logins, credentials that leave with departing staff, and the whole category of attacks beginning with a stolen or guessed password. It also changes what a compromise feels like: instead of an emergency in which nobody is sure what else that password unlocked, it becomes a short administrative task with a known scope.

What to watch for. Ask how a certificate is revoked, how quickly revocation takes effect, and who at your utility can do it. A system where only the vendor can revoke, on a ticket, during business hours, is considerably weaker than one where you can.

Principle 4 · No local attack surface

The mechanism. The device offers nothing to connect to locally: no USB, no Bluetooth, no Wi-Fi, and no network services beyond the minimum required. Where a device must have an Ethernet port for legitimate reasons — a plant with several controllers that talk to each other — it should accept connections only from equipment presenting a valid certificate from your own system.

What it kills. Walk-up attacks. The scenario where somebody cuts a padlock at a remote site, opens the panel, plugs a laptop into the programming port and leaves with your control logic altered. If there is nothing to plug into, physical access gains an attacker very little.

Why this matters more than it sounds. Most utilities protect remote sites with a fence and a padlock. That is adequate against casual trespass and irrelevant against anyone who came deliberately. Removing the local attack surface is the only control that does not depend on the lock holding.

"If it is in the cloud, isn't it less secure?"

This is the most common objection we hear, and it deserves a direct answer rather than reassurance, because the instinct behind it is reasonable: your data is somewhere you cannot see.

Here is what is actually on the other end. A well-designed platform runs on major cloud infrastructure — in our own case Microsoft Azure — and the servers running your system are dedicated to that purpose. What that infrastructure provides is a level of protection no individual water utility could fund on its own:

	What that means in practice
Physical security	Data centre access is controlled by multiple layers — perimeter security, biometric access control, escort requirements, continuous camera coverage and 24-hour staffing. Nobody walks in with a padlock key.
Personnel screening	Staff with any physical or logical access are background-checked, and access is granted on a least-privilege, audited basis rather than as a general permission.
Continuous scanning	The infrastructure is subject to continuous vulnerability scanning, intrusion detection and independent security auditing against recognised standards — work that goes on every day, not once a year when somebody remembers.
Patching	Underlying infrastructure is patched continuously by people whose whole job that is.

Now compare that honestly with the alternative. The realistic comparison is not a vault. It is a Windows computer in an office, running an operating system that stopped receiving security updates in October 2025, with a password several people know, in a building secured by an ordinary door lock, patched when somebody has time.

The question is not whether the cloud is perfectly secure. It is which of those two environments you would rather have your control system living in — and which one has somebody watching it at three in the morning.

Replay your own findings against this

Take the findings from your self-audit and work down the left-hand column.

A finding from your audit	What changes under this architecture
Equipment reachable from the internet	Nothing listens, so there is nothing to find by scanning
Default passwords on controllers or radios	No device passwords to leave at their default
One shared login the whole crew uses	Per-person credentials; shared logins are not needed to make it work
Access still active for somebody who left	Revoke that person's credential; device credentials are unaffected
Remote-access software with no second factor	No inbound remote-access tool; access is through the platform, per person
Unencrypted radio between sites	Encrypted cellular; interception yields nothing usable
A panel anyone could open and plug into	Nothing to plug into; the lock stops being the only control
Equipment that no longer receives firmware updates	Supported equipment that updates itself, remotely, as fixes are released
Nobody has tested a backup in years	Configuration and programs are retained centrally and continuously

What still needs you

No architecture removes the need for people to do a few things well. These remain yours:

- Phishing awareness — if somebody gives away their own login, no architecture can prevent it
- Your business network, email and billing systems, which are outside all of this
- Physical protection of the equipment itself against damage and theft
- An emergency response plan, and the ability to run manually

Part 4 · What Changes Day to Day

Security is the reason it is safer. This part is the reason most operators end up wanting it, and for many utilities it carries more weight than the security case.

Controllers that update themselves

On a conventional system, improving the program at a site means an engineer, a drive and an invoice — per change, per site. Four sites and one program change is four visits and a week of scheduling.

Where devices hold certificates and connect outward to a platform they trust, new programs can be delivered without anyone attending. Fixes and improvements arrive the way updates arrive on a phone. For a utility running discontinued equipment today, this is the single biggest change: you go from never receiving another fix to receiving them automatically.

Note the dependency, because it matters in Part 8: this capability is only safe because of the security architecture. A device that will accept new code with nobody present must be certain who is sending it. Without certificate-based identity and an authenticated outbound channel, automatic updates would be a vulnerability rather than a feature.

Replacing a failed device

How it goes today

Lightning takes out a controller at your furthest site at six on a Friday evening. You call the integrator. You wait until Monday, or pay emergency rates. An engineer drives out with a laptop, loads the program — and hopes the copy they have is the current revision. They test it and invoice you. Labour alone is commonly $650–1,300 before the hardware, and the site runs blind all weekend.

How it goes on a well-designed cloud system

- An operator — not an engineer — takes a spare unit off the shelf and swaps it in.
- They assign the new unit's serial number to that site, from a phone, on the spot.
- The device authenticates, pulls down its own configuration and program, and resumes exactly the job the failed one was doing.
- Total time on site is usually a few minutes.

The hardware had to be replaced either way. What changed is everything around it: no truck roll, no engineer, no wondering whether the program is the right revision, and no weekend with the site dark.

One spare covers everything

On a conventional system a useful spare is a pre-programmed unit configured for one specific site. Covering four sites properly means four spares, each kept current as programs change. Almost nobody does this, which is why a failure usually means waiting for a part.

Where a device's job comes from its identity in the system rather than from what was loaded into it beforehand, a spare is just a spare. One unit on the shelf can become the pump station, the tower or the filter controller depending on the role you assign it.

A REQUIREMENT WORTH ASKING FOR BY NAME

"Can any spare unit take on the role of any failed unit, without being pre-configured for that site?" A short question, and the answer separates systems designed for field replacement from systems where it is an afterthought.

Doing your own work

On most conventional systems, changing a screen, adding an alarm or building a report means a phone call, a quote, a scheduled visit and a wait. The result is predictable: the changes do not get made.

Where configuration lives in a platform you reach with a browser or a mobile app, an operator can reasonably expect to:

- Add the pump that was installed last month to the screen
- Move a value to where it can actually be seen from where people stand
- Adjust an alarm threshold that has been crying wolf at three in the morning
- Change who gets called, in what order, and after how long
- Build the trend needed for Tuesday's board meeting

Ask which of those five you can do yourself, because the answers vary more than the marketing does. Self-service should cover screens, alarms, reports and trends. Rewriting control logic remains engineering work, and it should.

Communications

Most small systems reach their sites over licensed radio. The costs are easy to overlook because they are spread out and have always been there: FCC licensing and renewals, frequency coordination, path studies, line-of-sight constraints dictating where a site can go, tower and antenna leases, climbs, repeaters at the least accessible location you own, seasonal path problems nobody can pin down, lightning damage, and radios on their own obsolescence cycle.

Cellular removes that entire category. No licence, no path study, no tower, no repeater and no climb, and coverage is the carrier's problem to maintain rather than yours.

A quieter benefit: no tall antenna

Radio telemetry wants height, which is why so many sites have a mast or a tower-mounted antenna with a long feedline running down into the panel. In a state that leads the country in lightning, that arrangement is an invitation — and a great deal of the storm damage water utilities pay for every year arrives down a feedline into the equipment it is connected to.

Cellular antennas are typically small and mounted at or near ground level, because the link does not depend on line of sight to another of your sites. That removes both the elevated strike target and the long conductive path into the panel. It is not lightning immunity — nothing is — but it materially reduces one of the most common causes of storm damage at remote sites.

Use CAT M1

This is the one piece of technology worth naming as a requirement, because it settles several questions at once.

CAT M1 is a 4G cellular technology designed for equipment rather than for phones. It trades speed — which telemetry does not need — for reach and efficiency, and the practical effects matter:

- **It is 4G, not 3G.** Utilities were burned when 3G networks were retired and equipment was stranded. Specifying a 4G equipment technology removes that worry from this generation of hardware.
- **Dramatically better link budget.** CAT M1 reaches into places ordinary cellular does not — inside buildings, below grade, and at distance from a tower. This is the technical reason coverage assumptions based on phones are unreliable.
- **Lower cost.** Both the radio hardware and the data plans are cheaper than full-speed cellular, because the technology is built for small, regular messages rather than video.

Confirming coverage — the vendor's job, not yours

The usual objection is coverage, and it is usually based on the wrong evidence.

YOUR MOBILE PHONE IS NOT A COVERAGE TEST

Phones use different radio technology — typically CAT 4 or better — on a single carrier, and they are optimised for speed rather than reach. A site where your phone shows no bars can be a perfectly reliable CAT M1 site, and we regularly install at locations where customers were certain there was no service.

The right way to answer the question is not to send an operator out with a handset. Cellular coverage by band and by technology is published information, and a competent supplier can determine from that data — before anybody drives anywhere — whether one or more towers carrying the correct bands cover each of your sites.

Put that obligation on the supplier, in writing. The customer does not have the tools or the reference data to make the assessment, and should not be asked to. A supplier who tells you to go and check your phone signal is telling you something about how they scope work.

Equipment able to use more than one carrier and switch between them automatically improves matters further: the site is connected if any one of the carriers reaches it.

Data you can actually use

A conventional historian on a local disk keeps whatever fits. When it fills, the oldest data rolls off.

That is why the awkward problems stay unexplained. A filter that will not backwash properly, or a pump station that occasionally fails to start, is only mysterious because the data you would need to explain it —

what everything was doing at three in the morning on the night it happened — was overwritten weeks ago. Long-retention history does not prevent those faults. It makes them findable, which is the difference between a permanent nuisance and a fixed problem.

This matters especially for the utility whose system "has never worked right." Recurring faults that nobody has been able to diagnose across multiple service visits are very often a data problem rather than an equipment problem — nobody has ever been able to see what happened.

When your superintendent retires

Around a fifth of utility employees are eligible to retire within five years, against an average vacancy rate near nine per cent. For most small systems the risk is not abstract: one person knows why the screens are laid out the way they are, which alarms can be trusted, and what the integrator did in 2014.

Where configuration lives in a platform rather than in a project file on somebody's laptop, it is documented by virtue of being configured, and a new operator can learn to change it in an afternoon. Continuity becomes a property of the system rather than of one person — a considerably easier thing to take to a board than "it is more convenient."

Part 5 · The Objections, Answered

Four questions come up in nearly every conversation. They are fair, and they deserve specific answers rather than reassurance.

“What if the internet goes down?”

Answered mechanically in Part 2. In summary: control continues locally, readings buffer and forward on reconnection, sites that depend on other sites fail over to a configured backup control mode, and what you lose is remote visibility and remote control for the duration.

Weigh that against what a radio path failure costs you now, and against what happens today when the SCADA computer itself fails — which for most utilities is a longer and more damaging outage, and takes the history with it.

“I do not want my data on somebody else's computer.”

Covered in detail in Part 3. The short version: the infrastructure a serious platform runs on has physical security, personnel screening, continuous scanning and continuous patching that no individual utility could fund. The realistic alternative is an unpatched computer in an unlocked office.

What you should insist on regardless is contractual rather than technical: that you own the data, that you can export all of it at any time in a documented format, and that this is written down rather than assured verbally.

“Subscriptions forever, instead of buying it once.”

The premise is worth examining. Conventional SCADA is not bought once either. Over ten years it involves support contracts, version upgrades, at least one computer replacement, and a steady flow of service calls at $500 minimum and $200 an hour. That is a subscription too — it is simply irregular, unpredictable and impossible to budget.

The distinction that matters is not recurring versus one-off. It is predictable versus unpredictable — a line item you can plan for, against a capital surprise that arrives when a motherboard dies.

“What if the vendor goes under, or puts the price up?”

A fair question — and one worth asking of every supplier, not only a cloud one. Conventional SCADA carries exactly the same exposure in a different costume. Small companies fail. Large companies discontinue product lines. Either way you end up with equipment nobody supports, which is precisely the position a great many utilities are in today with discontinued controllers and software that no longer receives fixes.

So the protection is the same in both cases, and it is contractual rather than architectural:

- A written right to export all your data, in a documented format, at any time

- Clarity on what happens to equipment you own if the service ends — does it keep running locally, and can it be repurposed?
- Notice periods, and any cap on price increases
- Whether the platform is built on widely-used infrastructure and standard protocols, or on something entirely proprietary

"My integrator says cloud is not ready for water."

They may be right, and they know your system. It is also fair to note that truck rolls and upgrade projects are their revenue, and a system that does not need them for routine changes is not obviously in their commercial interest. That is a conflict of interest rather than a character flaw, and you are allowed to name it.

Make it specific. Ask which requirements in Part 8 they believe cannot be met, and why. A good integrator will give you a substantive answer, and it may well change your mind.

Part 6 · What It Costs

ABOUT THE NUMBERS IN THIS PART

Figures are indicative market ranges as of August 2026, included so you can sanity-check a quote and build a rough case before going to market. They are not an offer from anybody and they will drift. Get current quotes.

The headline

Two claims, both testable against your own quotes:

- **At installation,** a properly designed cloud system should come in anywhere from about thirty per cent below a conventional SCADA installation to roughly comparable with it.
- **In operation,** it should cost less — in every case we have seen, from a single pump station and tower up to a system with a treatment plant. The savings come from three specific things: operator-performed device replacement, program updates that need no site visit, and screens, alarms, reports and trends the utility can change itself.

If a quote you receive does not show that pattern, it is worth asking why — either the conventional comparison is incomplete, or what is being offered is not delivering the operational advantages.

Three kinds of cost

Comparing the two as single numbers does not survive contact with reality, because the columns are made of different things.

	What it is	Do you feel it?
Costs you PAY	Service calls, storm damage, callouts, support contracts	Yes — on every invoice
Costs you DEFER	The software upgrade you should do and keep not doing	No — until you cannot defer it any longer
Costs you are EXPOSED TO	An incident you cannot recover from quickly	Not yet

Reading a hardware price properly

Hardware figures are the most commonly misread numbers in this market, because suppliers are not selling the same thing as each other.

COMPARE LIKE WITH LIKE

A modern controller of the kind discussed in this guide combines three things in one device: the programmable controller, the operator interface, and the cellular gateway. Conventional designs buy those as separate boxes. Indicatively, a PLC around $600, an operator panel around $1,200 and a reputable cellular gateway around $500 — roughly $2,300 for three devices, three power supplies, three sets of wiring and three things to fail. A combined controller should come in below that.

As a guide, and hardware only, combined controllers run from around $2,000 for a single-purpose site such as a pump station or tower, to around $4,000 for a unit capable of running a plant with multiple devices. Indicative, as of August 2026.

Those are equipment prices and nothing else. They do not include installation, engineering time, panel fabrication, instrumentation, the balance of materials, commissioning or support. Any comparison built on hardware prices alone will mislead you, in either direction. The worked examples below use fully installed figures instead.

Subscription

Commonly $200–600 per connected device per year, with simpler sites at the lower end and plant-class devices at the upper end. A pump station or a tower is normally one device; a treatment plant is often three or four.

Worked example — a distribution system

Two pump stations and two towers, four connected devices. These are

fully installed and provisioned figures — equipment, panel, installation, engineering, configuration and commissioning — which is the only basis on which a comparison means anything.

	Cost
Pump station, fully installed and provisioned — $12,000 each	$24,000
Tower, fully installed and provisioned — $9,000 each	$18,000
INSTALLED TOTAL	$42,000
Subscription, four devices, per year	about $960
Subscription over five years	about $4,800
FIVE-YEAR TOTAL	about $46,800

Against a conventional installation

On the headline above, a conventional installation of the same four sites should land somewhere between comparable and about forty per cent higher — call it $42,000 to $60,000 installed. The

difference then compounds, because the conventional system continues to generate costs the cloud system does not:

Over the following five years	Conventional	Cloud
Service calls for screen, alarm and report changes	Recurring	None — self-service
Engineer visits to replace failed devices	Recurring	None — operator swap
Program updates across the fleet	Per site, per change	None — remote
Storm damage restoration labour	Recurring	Reduced — operator swap
Software upgrade	$15,000–20,000 eventually	None
Subscription	None	about $4,800

A useful sanity check for a system this size: a single storm-damaged controller restored conventionally can cost more in labour alone than a year's subscription for every site you own.

Worked example — a system with a plant

The same pattern holds and the gap widens, because a plant has more devices, more instrumentation and more failure modes — so it generates more of exactly the costs that disappear.

Plant installations vary too much for a representative figure to be useful, and anybody who quotes you one without seeing the site is guessing. Price the distribution sites using the figures above, and have the plant quoted specifically. What should hold in the quote is the same headline: comparable to about thirty per cent below at installation, and clearly lower to operate.

The costs with no invoice

Not in either table, and worth naming because they are real money leaving your utility in forms that never get totalled.

- **Operational waste.** Water lost through a filter that will not backwash properly, energy burned by pumps cycling wrongly, chemical overdosed, equipment worn out early, and operator hours spent chasing intermittent faults nobody can see.
- **Regulatory exposure** if one of those becomes an excursion.
- **The security exposure you are carrying.** This is the largest of the three and the easiest to leave out of a spreadsheet, because it is nothing at all until it is enormous. A serious incident means emergency response, rebuilding a control system from nothing, running manually for days or weeks, legal and notification costs, and a boil order with your utility's name attached to it. Municipal ransomware recoveries have run into the millions. You are not paying that today. The question is whether you are prepared to keep carrying it, and for how long.

Paying for it — Oklahoma

Most utilities assume there is no money. Frequently there is, and it does not have to be cybersecurity money — SCADA is water infrastructure, and these programmes fund water infrastructure.

Programme	What it is	Worth knowing
REAP grants (OWRB)	Rural Economic Action Plan grants for water and wastewater infrastructure in rural communities.	Point-based. Priority to municipalities under about 1,750 population and rural water districts under about 525 taps. Applications open 1 July and close the first working day of September each year.
RIG (Oklahoma Rural Water Association)	Rural Infrastructure Grant programme, renewed under the state budget.	Applications are evaluated in the order received, and previously submitted applications remain in the queue. Worth asking ORWA where yours stands.
DWSRF (OWRB / ODEQ)	Drinking Water State Revolving Fund — low-interest loans for public water supply projects.	Principal forgiveness is available for qualifying systems, and awards to small Oklahoma rural districts are made regularly. This is the largest single source in the state.
CWSRF (OWRB)	Clean Water State Revolving Fund, for the wastewater side.	Same mechanism as DWSRF. Relevant if your project includes lift stations or treatment.
State water fund	Additional state money is expected to reach the revolving programmes in 2027.	Timing and mechanism are not final. Worth asking OWRB what is coming before you assume your only option is a rate increase.

On cybersecurity-specific funding: at the time of writing we are not aware of an Oklahoma programme dedicated to it. Several other states have created one, and federal programmes may follow. It is worth asking your state drinking water programme each budget cycle rather than assuming the answer has not changed.

Whichever route you take, a completed self-audit with a costed plan attached makes a considerably stronger application than a general appeal for funding. It shows what you assessed, what you found, and precisely what the money buys.

Part 7 · Are You Ready?

Triggers that mean it is worth looking now

Any one of these is a reasonable reason to start the conversation. Two or more and the question is really about timing rather than whether.

- ☐ Your equipment is discontinued and spares come from used inventory
- ☐ Your SCADA computer runs an operating system that no longer receives security updates
- ☐ You have a quote for a software upgrade that you cannot see a way to fund
- ☐ A controller keeps failing, or you are running one you cannot buy a replacement for
- ☐ You have a site the radio never reached properly, or one with no telemetry at all
- ☐ You are about to build a new site
- ☐ The person who understands the SCADA system is within a few years of retiring
- ☐ Your last audit left findings you have no affordable way to close
- ☐ Your insurer, regulator or board has started asking questions you cannot answer well

Readiness check

Being ready is not the same as being convinced. Work through these before going to market — the gaps are usually fixable, and they are much cheaper to fix now.

Question	Yes	No	Notes
Do we have a written list of our equipment and sites?			
Do we know who currently has access to what?			
Have we asked a supplier to confirm cellular coverage at each site from published band data?			
Do we know what our current system costs us per year, in total?			
Is there a budget path — REAP, RIG, DWSRF, rate case, capital plan?			
Do we know who will own this internally after it is installed?			
Do we know our contract position with our current integrator?			
Has anyone briefed the board that this conversation is coming?			

The first two come directly from Parts 1 and 3 of the self-audit workbook. If you have not done those, do them before you ask anyone for a quote — you cannot get a sound proposal for a system nobody has inventoried.

Note the third: confirming coverage is the supplier's job, not yours. See Part 4.

Pick your first site

If your funding arrives in stages rather than all at once, the order matters. The sites that make the best starting point are the ones where the benefit is most obvious and the risk is lowest.

A good place to start	Why
A new site you are about to build	No migration at all — you are buying something either way, so the comparison is clean
A controller that keeps failing	You are already paying for it repeatedly; replacing it has its own justification
A lift station or tank with no telemetry today	Pure gain. Nothing to lose, and the benefit is obvious to everyone
The site the radio never reached	It solves a problem that has resisted every previous attempt, which is persuasive
A site far enough away to be annoying	The travel saving shows up immediately and is easy to quantify

Part 8 · How to Specify It

This is the part that determines whether you get a secure system or merely a modern-looking one. The requirements below are written as capabilities, so any supplier meeting them qualifies — which is the point.

HOW TO USE THIS PART

These are requirements to put into your bid documents and questions to ask in evaluation meetings. They are written in plain language for scoping conversations, not as contract clauses. When you are ready to issue a bid package, your engineer should translate them into specification language — there is a companion document written for engineers that does exactly that.

The ten requirements that matter

1. Nothing at the utility listens for inbound connections

What to require: All equipment at your sites initiates connections outward. No inbound connection to any device at any site is required for normal operation, remote access or support.

Why it matters: This removes the first step of the most common attack. Equipment that does not listen cannot be found by internet scanning and cannot be connected to by anyone who finds its address.

A good answer sounds like: *"No device at any site accepts inbound connections. Remote access is through the platform, and the site equipment initiates every session."*

An evasive answer sounds like: *"It sits behind a firewall, so nothing can get to it — which means something is listening and a rule is protecting it."*

2. All communication between sites is encrypted, by default

What to require: Encryption of data in transit is enabled by default, cannot be disabled by configuration error, and covers every link between site equipment and the platform.

Why it matters: Unencrypted radio can be read and injected into from outside your fence, with no alarm and no log entry. Encryption that is optional is encryption that will eventually be off somewhere.

A good answer sounds like: *"Encryption is always on. There is no configuration in which a site communicates in clear text."*

An evasive answer sounds like: *"Encryption is supported and can be enabled during commissioning."*

3. Credentials are certificates, issued per device, and revocable individually

What to require: Device authentication uses certificates rather than shared passwords. Each device has its own credential. Any individual credential can be revoked, and revocation takes effect immediately.

Why it matters: Passwords get shared, written down, reused across sites and inherited by whoever comes next. A certificate cannot usefully be copied down, and revoking one is an administrative action rather than a site visit.

A good answer sounds like: *"Each device has a unique certificate issued at deployment. You can revoke any of them yourself, from the platform, and it takes effect at once."*

An evasive answer sounds like: *"Every device has a strong password, and we change them on a schedule."*

4. No unnecessary local attack surface

What to require: Site equipment provides no USB, Bluetooth or Wi-Fi interface usable for configuration, and no network services beyond those strictly required. Where a wired port is required for inter-device communication, connections are accepted only from equipment presenting a valid certificate from your own system.

Why it matters: Most remote sites are protected by a fence and a padlock. This is the only control that still works after somebody gets inside the fence.

A good answer sounds like: *"There is nothing to plug into. The Ethernet port accepts only certificate-authenticated connections from within your own deployment."*

An evasive answer sounds like: *"The panel is locked and the site is fenced."*

5. A replacement device provisions itself

What to require: Any spare unit of the correct type can take on the role of any failed unit without being pre-programmed for that site. Replacement can be performed by utility staff, and the device retrieves its own configuration and program.

Why it matters: This is where most of the day-to-day saving comes from, and it removes the dependency on somebody holding a current copy of each site's program. One shelf spare covers every site rather than one per site.

A good answer sounds like: *"Swap the unit, assign it the role in the platform, and it pulls its configuration down. Your operator can do it; typically a few minutes on site."*

An evasive answer sounds like: *"We keep a configured spare for each of your sites and can be on site next business day."*

6. Operators can change screens, alarms and reports without a service call

What to require: Utility staff can create and modify operator screens, alarm thresholds and notification rules, escalation and on-call routing, and reports and trends — without vendor involvement or additional charge. Access is by browser or mobile app from any laptop, tablet or smartphone, with per-person logins.

Why it matters: Where changes cost money and require scheduling, they do not get made. Alarms that cannot be tuned get switched off, screens drift out of date, and people go back to the clipboard.

A good answer sounds like: *"All four are self-service, from a browser or the app, wherever you are. Here is the training, and here is a system where you can try it before you buy."*

An evasive answer sounds like: *"We can make those changes for you quickly under a support agreement."*

7. You own your data and can export it at any time

What to require: The utility owns all operational and historical data. Complete export in a documented, non-proprietary format is available at any time, at no additional charge, and this is stated in the contract.

Why it matters: This is your protection if the relationship ends, the price rises, or the supplier does not survive. A verbal assurance is not protection — and this exposure exists with conventional suppliers too.

A good answer sounds like: *"You own it, you can export all of it whenever you like, and here is the clause that says so."*

An evasive answer sounds like: *"Of course you own your data — offered without any corresponding contract language."*

8. The update model is stated in writing

What to require: The supplier states who is responsible for updating platform software, site firmware and device programs; how often; whether updates are included; and what the utility must do.

Why it matters: On conventional systems this ambiguity is exactly how utilities end up years behind and then facing a five-figure upgrade — or with equipment that stops receiving fixes altogether. Settle it before you buy.

A good answer sounds like: *"We patch the platform continuously and it is included. Device firmware updates deploy remotely with your approval, also included. Here is our published record."*

An evasive answer sounds like: *"Updates are handled as part of ongoing support, with the terms left undefined."*

9. CAT M1 cellular, multi-carrier, with coverage confirmed by the supplier

What to require: Site communications use CAT M1 — a 4G technology built for equipment — with the ability to use more than one carrier and switch automatically. The supplier confirms, from published coverage and band data, that each of your sites is served, and states the expected support horizon for the technology.

Why it matters: CAT M1 has a far better link budget than the cellular in a phone, so it reaches sites where a handset shows nothing; it is 4G, which settles the stranded-equipment worry left over from the 3G shutdown; and the hardware and data both cost less. Confirming coverage requires reference data the utility does not have, so it belongs with the supplier.

A good answer sounds like: *"CAT M1 on multiple carriers with automatic switching. We have checked the bands against every one of your site coordinates — here are the results."*

An evasive answer sounds like: *"It uses LTE, and you should check whether your phone gets a signal out there."*

10. Local control continues through an outage, including at dependent sites

What to require: Control logic executes at the site and continues through a communications outage of any duration. Data buffers locally and forwards on reconnection with no gap in history. Sites that normally depend on data from another site have a defined, configured failover control mode — for example a pump station reverting to discharge pressure control when it cannot see tower level.

Why it matters: This is the difference between a control system and a monitoring system. The dependent-site question is the one most often missed, and it is the one that determines whether your customers still have water during an outage.

A good answer sounds like: *"Control never depends on the link. Each dependent site has a configured failover mode — here is what each of yours would do."*

An evasive answer sounds like: *"Our uptime is excellent and outages are very rare."*

Red flags

- **"We can put a firewall in front of it."** Means something is listening. A firewall is a rule that can be misconfigured.
- **"The password is unique per site."** Still a password. Ask what happens when the person who knows it leaves.
- **"We would need to schedule a visit for that."** For a screen or alarm change, the operational model has not changed at all.
- **"You own your data," with no contract language.** Ask for the clause. If it does not exist, the statement is decoration.
- **"Go and see whether your phone has a signal there."** Coverage verification is the supplier's job, and a phone is the wrong instrument.
- **Any suggestion that scanning your live control network is routine.** Scanners have knocked working PLCs offline. A supplier casual about this should not be near your plant.
- **Reluctance to let you talk to an existing customer of similar size.** Ask for a reference at a utility like yours, and call them.

Comparing proposals

Score each proposal against the ten requirements. A blank scoring sheet is in Appendix B.

Scoring	Meaning
2 — Fully meets	Demonstrated, documented, and stated in writing
1 — Partially meets	Claimed but not demonstrated, or met with conditions
0 — Does not meet	Not offered, or the answer was evasive

Anything scoring 0 on requirement 1, 2, 3 or 10 is not a secure-by-design system, whatever else it offers. Those four are the architecture. The rest is how well it is built around it.

What to hand your engineer

For a formal bid your engineer needs specification language rather than plain-English requirements. Ask them to write these ten into the specification as mandatory capabilities — described by function, not by manufacturer — so any bidder meeting them can compete and any bidder who cannot is excluded on merit. There is a companion document written for engineers with suggested language for exactly this.

Part 9 · Making the Move

A good decision still needs a good transition. This part covers the mechanics most guides skip.

Running old and new at the same time

Unless you cut over everything in one weekend, there will be a period with both systems live. Plan it rather than tolerating it.

- Decide which system is authoritative for each site, and make sure every operator knows.
- Watch for alarms arriving twice — or worse, from only one system while people assume both.
- Agree how long the overlap lasts, with a date. Open-ended parallel running is how utilities end up paying for two systems for three years.
- Keep the old system's history until you have confirmed what you needed has been carried across.

What happens to your existing controllers

A fair question, and one suppliers often gloss over. There are three honest answers, and which applies depends on your equipment.

- **They are replaced.** Cleanest, and usually right where controllers are old, discontinued or already failing. It is also where the field-replacement and self-updating benefits actually come from — keeping an unsupported controller keeps its limitations with it.
- **They are kept and connected through a gateway.** Preserves working equipment and reduces up-front cost. Understand what you are keeping: an old controller behind a modern gateway is still an old controller, with its own credentials and its own unpatched firmware.
- **They stay exactly as they are, for now.** Reasonable at sites not yet in scope. Just do not lose track of them.

Ask any supplier which of these they propose for each site, and why.

What to carry across, and what to rebuild

Item	Recommendation
Historical data	Carry across what you genuinely use. Most utilities need far less than they expect — decide deliberately rather than paying to migrate everything.
Operator screens	Rebuild rather than replicate. Your current screens reflect what was possible in the old system, and often what somebody assumed years ago.
Alarm setpoints	Carry across, then review. Transition is the best opportunity you will get to retire the alarms everybody ignores.
Control logic	Rebuild with engineering involvement. Do not copy logic nobody understands.
Reports	Rebuild. Usually quick, and the results are usually better.

Documentation and drawings	Update as you go. If it does not happen during the project, it will not happen.

Realistic timelines

Stage	Typical
Scoping, coverage confirmation, quotes	2–6 weeks
Procurement and approval	Highly variable — often the longest step
Site installation and commissioning	1–3 days per site
Operator familiarisation	Days, not months, if self-service works as promised
Full changeover across a small system	Weeks to months, depending on funding

Keeping leverage

- Tie payment to demonstrated acceptance criteria, not to installation dates.
- Write the Part 8 requirements into the contract, not just the bid. A capability promised in a proposal and absent from the agreement is not a capability you bought.
- Define what "done" means for each site in advance, in writing.
- Get the data export right stated in the contract before you sign, not at renewal.
- Keep your old system available until acceptance is complete.

What to keep from the old system

Keep the SCADA computer, powered down and stored, until you are certain everything you needed has been carried across — a year is not excessive. Keep a copy of every controller program, the historian export and the documentation. Then, when you are genuinely finished, dispose of the computer properly: it holds credentials, network details and process information, and it should be wiped rather than sold or skipped.

Part 10 · Taking It to Your Board

Most operators do not control the budget. This part is written so you can hand it upward, or use it to prepare for the meeting where the decision is actually made.

What boards and councils actually ask

The question	What they are really asking	How to answer
Why now?	Has something changed, or is this a wish list?	Point at the specific trigger: discontinued equipment, unsupported software, a failing device, an audit finding, a deadline.
What happens if we do nothing?	Is this optional?	Usually: nothing, for a while — then a forced replacement at a worse moment, or an incident.
Haven't we already spent money on this?	Why didn't the last thing fix it?	The free fixes were done and they helped. These are the findings that had no free fix.
Why is it a subscription?	Are we signing up to something forever?	We already pay every year — in service calls and storm repairs. This makes it predictable and puts it in a budget.
Who else is doing this?	Is this proven, or are we the experiment?	Name comparable utilities. Offer to have the board hear from one.
Where does the money come from?	Is this a rate increase?	Not necessarily — REAP, RIG and DWSRF all fund water infrastructure. Part 6 has the list.

Framing that works

- **Lead with the money you already spend.** Boards engage with a comparison of two costs far better than with a description of a technology.
- **Predictable versus unpredictable.** An annual line item they can plan for, replacing capital surprises that arrive unannounced.
- **A deferred bill, not a new one.** The upgrade has been postponed for years. It has not gone away.
- **Continuity.** When the person who understands the system retires, does the knowledge leave with them?
- **Name the funding route.** A request that arrives with a grant or loan programme identified is a different conversation from one that does not.

What not to do

- **Do not lead with the attack story.** It reads as fear-selling to a finance committee, and it invites "so why haven't we been hit?" — a question you cannot win. Put the security case second, as one reason among several.

- **Do not present a single large number with no funding route attached.** Show what you are asking for and where the money would come from.

One-page summary for your board

Fill this in and take it with you. One page, their language, your numbers.

What we are asking for	
What it costs, installed	
What it costs annually, once running	
What we spend today on the same problems	
Where the money would come from	
What specifically triggered this request	
What happens if we defer it another year	
How we will know whether it worked	
When we will report back	

Appendix A · Requirements Checklist

Tear this out and take it to supplier meetings. Full explanations are in Part 8.

#	Requirement	Meets?	Evidence seen
1	Nothing at the utility listens for inbound connections		
2	All site communication encrypted by default, cannot be disabled		
3	Certificates per device, revocable individually by us		
4	No usable USB / Bluetooth / Wi-Fi; wired ports certificate-controlled		
5	Any spare provisions itself into any role; staff can replace		
6	We can change screens, alarms and reports ourselves, from anywhere		
7	We own the data; export in writing, any time, no charge		
8	Update responsibilities and costs stated in writing		
9	CAT M1, multi-carrier, coverage confirmed by supplier		
10	Local control through an outage, including dependent-site failover		

Also ask

- ☐ Can we speak to a utility of similar size that has been live for more than a year?
- ☐ Can we see it working over cellular at a real site, not on your network?
- ☐ What does it cost to add a site later?
- ☐ What happens to our equipment if we stop the subscription?
- ☐ Who do we call at 2 a.m., and what is the answer time?
- ☐ What is your process if we report a security vulnerability in your product?

Appendix B · Proposal Scoring Sheet

2 = fully meets and demonstrated · 1 = partial or claimed only · 0 = does not meet

Supplier: _____ Date: _____ Scored by: _____

#	Requirement	Weight	Score	Weighted	Notes
1	Nothing listens for inbound connections	×3			
2	Encrypted by default, cannot be disabled	×3			
3	Per-device certificates, we can revoke	×3			
4	No local attack surface	×2			
5	Self-provisioning replacement devices	×3			
6	Operator self-service, from anywhere	×3			
7	Data ownership and export in writing	×2			
8	Update model in writing	×2			
9	CAT M1, multi-carrier, coverage confirmed	×2			
10	Local control and dependent-site failover	×3			
	TOTAL				

Maximum weighted score is 52. Treat any 0 on requirements 1, 2, 3 or 10 as disqualifying regardless of the total — those four are the architecture.

Beyond the scoring

Consideration	Notes
Reference from a similar utility — did you call them?	
Total five-year cost, installed plus subscription	
What is excluded from the price	
Contract terms: export rights, notice, price protection	
Do we believe they will still be here in ten years?	

Appendix C · Cost Worksheet

Aligned with the self-audit workbook so figures carry across.

1 · What we spend today

Last twelve months	$
Service calls — screen, alarm, report and configuration changes	
Service calls — diagnosis that remote visibility would have avoided	
Device failures needing an engineer (labour only)	
Storm and lightning damage — labour	
Emergency and after-hours callouts	
Radio: licences, path studies, tower and repeater work	
SCADA support or maintenance contract	
ANNUAL TOTAL	
FIVE-YEAR TOTAL (× 5)	

2 · What we have deferred

Item	$	Due when?
SCADA software upgrade		
Replacement computer		
Radio or controller replacement		
TOTAL DEFERRED		

3 · What a new system would cost

Use fully installed and provisioned figures from your quotes, not hardware prices.

Item	Count	Unit	Total
Pump stations, installed			
Towers / tanks, installed			
Plant, installed (quoted separately)			
INSTALLED TOTAL — one time			

Subscription — annual			
Subscription — five years			
FIVE-YEAR TOTAL			

4 · The comparison

	Five years
What we spend today, if nothing changes	
Plus what we have deferred	
TOTAL, staying as we are	
New system, installed plus subscription	
DIFFERENCE	

Appendix D · Glossary

Term	Meaning
CAT M1	A 4G cellular technology built for equipment rather than phones. Long reach and low power rather than high speed, with better building and distance penetration than the cellular in a handset.
Certificate	A cryptographic credential proving a device or person is who they claim to be. Cannot be usefully written down or spoken, and can be revoked individually.
Cloud SCADA	An arrangement where control runs in controllers at your sites while screens, history, alarms and configuration are hosted centrally and reached with a browser or mobile app.
Edge	Processing done at the site rather than centrally.
Failover mode	A configured alternative control strategy a site uses when it loses the data it normally relies on.
Gateway	A device passing data between a site and somewhere else.
Historian	The database holding your readings over time.
HMI	Human-Machine Interface — the screen an operator uses.
Link budget	How much signal loss a radio link can tolerate and still work. A better link budget means reaching further, or into harder places.
MFA / two-step login	Requiring something beyond a password, usually a code from a phone.
MQTT / MQTTS	An efficient messaging method suited to cellular links. The S indicates it runs inside an encrypted connection.
Multi-carrier	Equipment able to use more than one mobile network and switch between them automatically.
OT	Operational Technology — the equipment running the physical process, as distinct from office IT.
Outbound-only	Equipment initiates connections out and accepts none in.
PLC	Programmable Logic Controller — the device running control logic at a site.
Provisioning	Giving a device its identity and configuration so it can do its job.
Revocation	Withdrawing a credential so it no longer works, immediately.
RTU	Remote Terminal Unit — similar to a PLC, historically used at remote sites.
Segmentation	Separating the control network from the office network.
TLS	The encryption protecting data in transit; the padlock in a browser.

Appendix E · Where to Get Help

Free

Who	What
EPA	Cybersecurity assessments and technical assistance for water systems, through the Water Sector Cybersecurity Evaluation Program.
CISA	Free external vulnerability scanning for water utilities, plus regional advisors who will visit your plant. report@cisa.gov · 1-844-SAY-CISA
WaterISAC	Water-sector threat alerts and cybersecurity fundamentals.
Oklahoma DEQ	Drinking water programme: compliance deadlines and requirements.
Oklahoma Rural Water Association	Circuit riders, technical assistance, and the RIG programme.

Funding

Who	What
Oklahoma Water Resources Board	REAP grants, DWSRF and CWSRF loans with principal forgiveness for qualifying systems. Start here.
Oklahoma Rural Water Association	Rural Infrastructure Grant (RIG).

Related documents

- **Water System Security Self-Audit workbook.** Find your own weaknesses and fix the free ones. Do this before buying anything.
- **Engineer edition of this guide.** Suggested specification language for the ten requirements in Part 8, written for the engineer preparing your bid documents.

About This Guide

Published by D6 Labs and free to use, copy and share within your utility and with other water systems.

We build SCADA equipment, and this guide names no products — including ours. The reasoning is in the front matter: we would rather compete against a high standard than steer you toward a brand. If the ten requirements in Part 8 become normal in water sector bid documents, every utility gets a better outcome regardless of who wins the work.

If you put those requirements in a bid and we lose it to a supplier who meets them, that is a good result.

Disclaimer

This guide is provided for educational purposes only, without warranty of any kind, express or implied. It is not professional engineering, security, legal or financial advice, and is not a substitute for professional assessment or a properly prepared specification.

Cost figures are indicative market ranges as of August 2026, are not an offer, and will change. Obtain current quotations. Funding programme details change frequently; confirm with the administering agency.

Completing anything in this guide does not satisfy any regulatory requirement. You remain responsible for your own systems, procurement decisions and regulatory obligations. D6 Labs accepts no liability for any loss, damage or disruption arising from its use.

Version

Version 1.0 — August 2026

First edition. Corrections and suggestions are genuinely welcome — particularly from anyone who has used Part 8 in a real procurement and found it wanting.

Questions / Contact information

Digital Six Laboratories Inc

425 N Meridian Ave

Oklahoma City, OK 73107

1-844-365-8647

info@d6labs.com