SewerAI

Inside the Sewer Infrastructure Ecosystem:

Getting Utilities, Contractors, and Engineers on the Same Record

"By the time it makes it back to the utility, the video often doesn't match what's in the database. There's been all these corrections made." — Tim McGarry, Director of Sales at SewerAI

By the time a CCTV inspection makes it from a contractor's truck back to the utility that paid for it, the data has been touched, manipulated, uploaded, and downloaded three to six times. Field crew to office. Office to consultant. Consultant review. Consultant back to contractor for rejected work. Contractor back to consultant. Consultant back to utility. Sometimes still on a physical drive.

"By the time it makes it back to the utility, the video often doesn't match what's in the database," says Tim McGarry, director of sales at SewerAI. "There's been all these corrections made."

At every step in the chain, the next stakeholder knows the data they're receiving has some unknown degree of error baked in — so they re-review it. The contractor's data person QCs the field crew's work. The consultant QCs the contractor's QC. The utility QCs the consultant's recommendation. The same decisions get mulled over three or four times before anyone acts on them.

The result is a system where every party is doing more work than the last one and trusting it less. Engineers stamp recommendations they had to rebuild from scratch. Contractors deliver inspections that get rejected after the crew has moved on. Utilities approve capital plans they can't fully audit.

This whitepaper is for utility ops leaders who recognize that pattern and want a way out. The good news is that the way out doesn't require replacing anyone's expertise. It requires getting the experts to their decisions faster—by cutting the redundant review cycles that currently eat most of the project. The rest of this paper covers where the handoffs break down, what changes when the experts can trust the data they're handed, and what each party gets back when the system actually works.

Three parties, one asset, four handoffs

The utility owns the pipes. They contract two things out: the expertise to make decisions, and the work to execute them. That's the engineering consultant on one side and the contractor on the other. Three parties working on the same asset, on behalf of the utility.

"Decision-making is slow for that reason. It's expensive. And things get delayed or missed or incorrectly assigned because of all the back and forth," says McGarry.

The breakdowns happen at four specific handoffs. Each is a point where a deliverable changes hands and something measurable goes wrong.

- **Field to contractor's office.** Data doesn't make it off the truck cleanly. Files get corrupted. Notes the field crew took get translated into a coding system the office uses, and detail gets lost in the translation.
- **Contractor's office to engineering consultant.** Delivery delays. Format mismatches. As McGarry puts it: "Believe it or not, there's still a lot of physical data delivery on drives that's going on."
- **Consultant review.** First pass review leads to rejections. Submittals go back to the contractor for rework while new work is still being delivered. The contractor is now QCing yesterday's data while collecting today's.
- **Consultant to utility.** The utility rejects the budget or the method. The consultant has to rework recommendations. The three-year plan becomes a seven-year plan. Or an emergency repair drains the proactive budget and the scope shrinks.

By the end of that chain, the data has been reviewed by everyone, trusted by no one, and acted on by a utility that can't fully reconstruct how it got to the version they're looking at.

Why every party redoes the same review

Ask any engineer or contractor why they re-review data that's already been QC'd, and the answer is some version of: I have to put my name on this.

"Decision-making is slow for that reason. It's expensive. And things get delayed or missed or incorrectly assigned because of all the back and forth." — McGarry

Their reputation, their license, their contract, all depend on the accuracy of what they sign off on. So they take the data back to the start and rebuild trust in it from scratch.

"The same decisions are being mulled over multiple times before it finally reaches the end of that chain. Because confidence is low at each step. And every time it goes to a new stakeholder, they understand there's a varying degree of accuracy and a lot of subjectivity in the data, so they'll take it all the way from the beginning again," says McGarry.

The pattern is rational at the individual level and ruinous at the system level. There's no single unifying step where the data passes through one accepted process. So every step tries to be that unifying step, and the project pays for the same work three or four times.

Get the experts to the five-yard line

The people in this ecosystem are good at their jobs. Field crews can read a defect on a camera screen with experience that takes years to build. Engineers can look at a stretch of pipe data and know which assets need attention in the next budget cycle and which can wait. Utility ops leaders can look at a capital plan and know whether it'll survive a board meeting.

None of those skills are what's slowing programs down. What's slowing programs down is the work the experts are doing instead of using those skills.

"Technology shouldn't replace the decisions made by engineers or folks in the construction industry. It should help them make decisions faster and give them all the context they need — rather than bogging them down in aggregating, synthesizing, QCing the data. If we can get them to the five-yard line of the decision faster and just let them carry it in using their actual expertise, that's a win for everybody," says McGarry.

The five-yard-line frame is the right one. The expertise is what carries the ball into the end zone — the engineer's judgment

about risk, the contractor's read on field conditions, the utility's call on what makes the capital plan defensible. A shared system has a different job: get to the five-yard line faster, with data the next person in the chain can actually trust.

That trust comes from one straightforward shift: When the data carries a clear record of who did what and when, the next stakeholder doesn't have to redo it.

The contractor's coding doesn't get re-coded by the consultant from scratch. The consultant's analysis doesn't get re-litigated by the utility. The contractor keeps their CCTV equipment. The engineer keeps their analytical framework. The utility keeps its GIS. What changes is that the outputs of all three feed one shared record—and the time everyone used to spend reconciling versions goes back to the work that actually moves the project.

McGarry says it plainly: there'll be greater trust in the quality of the data. That's the whole game. When trust is there, decisions get made on the first pass. When trust is missing, every decision gets made three or four times—and the project pays for it.

What the ecosystem looks like when it works

The good version of this ecosystem is a story about three groups of professionals who each get to spend more time doing what they're best at—not about the technology that makes that possible.

Field crews and contractors
Crews running the camera bring something the office can't replicate: real-time judgment in the pipe. They see what the camera sees, but they also see what the camera misses—the smell, the flow rate, the way the structure looks above the manhole.

When the data they collect carries forward cleanly, with their notes intact and their codes attributable to them, the next person in the chain inherits all of that context. The field crew's work counts as the field crew's work. Contractors get to be paid for the inspection they delivered, not for the version that survived three rounds of office QC and consultant review.

Engineering consultants
Engineers bring the analytical framework that turns inspection data into a defensible recommendation. That's the work their license is on the line for, and it's the work they're hired to do. When they don't have to spend the first half of every project re-cleaning the data they were given—questioning every code, double-checking every defect classification—they get to do more of the work they're actually good at. More risk modeling. More capital planning. More of the strategic conversations utilities hire them for in the first place.

Utilities
The utility brings the asset knowledge and the accountability. They know the system better than anyone, because they've been running it for decades. When they have a record they can trust—one where every data point can be traced back to its source and every decision has an author—they get to be directive in the way that produces good outcomes. They set standards upfront. They review progress in real time instead of inheriting a finished report at the end. They go to the board with a capital plan that has the audit trail to back it up.

"A good system looks like when there's a question about an asset, everyone goes to the same source for the answer. Rather than the utility opening one tool, the consultant opening their proprietary tool, the contractor pulling out a PDF report—that's where decisions get muddled, stalled, or reworked," says McGarry.

None of these professionals need to be replaced or managed harder. They need a system that lets their expertise carry forward

"The same decisions are being mulled over multiple times before it finally reaches the end of that chain. Because confidence is low at each step. And every time it goes to a new stakeholder, they understand there's a varying degree of accuracy and a lot of subjectivity in the data, so they'll take it all the way from the beginning again." — McGarry

instead of getting redone at every step.

The cost of staying reactive

Most utilities operate reactively, not by choice but by phone call. A sinkhole opens. A homeowner reports a backed-up lateral. Crews go out and answer the calls. Each one is necessary. Each one is also the least efficient way to do inspections—scattershot across the service area, no ability to plan routes or batch work, and an emergency premium baked into every repair.

"Each time they're going out there on a reactive call, it's the least efficient way to do inspections. Versus having the time to go inspect an entire area, get a baseline, and make decisions on which assets are going to fail in the next five years—which costs on average $1 versus the $5 spent on the emergency repair," says McGarry.

Five-to-one is the gap between proactive and reactive spending. And it compounds: every reactive repair shrinks the proactive budget that would have prevented the next reactive repair. Breaking out of the loop requires defensible data the utility can take to a board to fund proactive work—which is what one shared record makes possible.

What this looks like in practice

›› HOUSTON: AI-assisted assessment at the largest scale in the country

The city of Houston operates the nation's largest sewer collection utility—6,000 miles of gravity conveyance and 129,000 manholes. In 2021, the city entered a federal EPA Consent Decree that requires systematic inspection and assessment of the entire collection system, with near-real-time reporting to federal authorities.

That program runs through six CCTV and cleaning contractors and three engineering firms. The exact ecosystem this paper has been describing, at maximum scale.

Since selecting SewerAI in 2021, Houston has increased inspected miles per year by 53%—a critical metric for hitting Consent Decree milestones. An independent third-party audit of 1,500 surveys validated 97% accuracy in defect coding, which is roughly eight times the average accuracy of contractor-submitted data. The same shared platform that makes that accuracy possible has cut the failure rate for contractor submittals by 55%, eliminating the rejection-and-rework cycle that used to consume project time across the contractor network.

The financial picture follows from the workflow change: roughly $1M per year in staff augmentation savings, another $100K in year-one savings from accurate contractor pay-items through QAI (Quality Assurance for Inspections). The condition data generated through the platform now directly supports $6 billion in planned infrastructure investments over 15 years.

"The data is readily available both in the app and via an API, making it easier for operations to make critical decisions required to maintain the collection system. SewerAI reduced the internal labor hours needed to code and review, and the data is more precise," says Greg Eyerly, Senior Assistant Director, City of Houston.

›› VILLAGE OF SCHAUMBURG: a utility and its engineer of record, looking at the same data

The village of Schaumburg, Illinois ran an accuracy study with RJN Group, the village's engineer of record. The two organizations took approximately 75,000 linear feet of existing inspection footage

"Technology shouldn't replace the decisions made by engineers or folks in the construction industry. It should help them make decisions faster and give them all the context they need — rather than bogging them down in aggregating, synthesizing, QCing the data. If we can get them to the five-yard line of the decision faster and just let them carry it in using their actual expertise, that's a win for everybody." — McGarry

"The data is readily available both in the app and via an API, making it easier for operations to make critical decisions required to maintain the collection system. SewerAI reduced the internal labor hours needed to code and review, and the data is more precise." — Greg Eyerly, Senior Assistant Director, City of Houston

and compared the Village's self-performed manual PACP coding against the same footage processed through SewerAI's AutoCode.

The results changed what both parties thought they knew about the collection system.

Manual coding identified 3.24 Grade 4 and 5 defects per mile. AutoCode identified 22.89—more than 600% more of the most severe defects, the ones that demand the most urgent operational and capital response. For structural defects specifically, manual coding found 6.48 per mile; AutoCode found 99.93—ore than 1,400% more.

The defects were always there. The capacity to find them at scale wasn't. With the manual baseline, the village's capital plan and emergency exposure were sized to a picture of the system that turned out to be substantially incomplete. With the new baseline—produced from a shared record both the utility and its engineer of record could see and validate together—the Village and RJN Group can size the program to what's actually in the ground.

This is the small-utility version of the same story Houston tells at scale: the experts in the ecosystem—the village's operators, RJN's engineers—get to spend their judgment on real conditions instead of on reconciling whose version of the data to trust.

Closing

When a sewer system makes the news, the utility shows up at the press conference. The asset owner is the public face of the program, and that accountability shapes how the work gets run. Some utilities outsource heavily and step back from the day-to-day. Others stay directive by setting standards upfront for what data gets collected, where it lives, and how decisions get made. Both end with the utility owning the outcome. Only one ends with the utility having shaped the outcome they wanted.

The utilities that handle this well share three habits. They get directive early on. They write data rights into the contract before kickoff. And they pick a system that earns adoption by integrating with the tools contractors and engineers already use. Budget size has less to do with the result than most people assume. What matters is whether every party can defend their work—and whether the utility can defend the program.

SewerAI's platform—Pioneer for data and project management, AutoCode for AI-coded inspections, and Risk & Rehab for capital planning—is built on the integration-first principles described in this paper. If you're a utility ops leader trying to move from passive to directive on data, we'd like to show you what that looks like in practice.

Book a demo at sewerai.com.