

Water Quality Between the Plant and the Tap

Estimating residence time, water age and chlorine residual from the network data a utility already holds, and saying plainly that they are computed

PRODUCT

Grid Data Enhanced Analytics

AUDIENCE

Water quality managers, operations, regulatory compliance

Sampled at a few points, decided everywhere

A water utility knows a great deal about its water at the treatment plant. It measures disinfectant, turbidity and flow there continuously, and it samples at a set of sites across the distribution system on a schedule. Between those points lies most of the network: miles of main, storage tanks, dead ends and slow loops where water sits.

That is where quality is actually decided. Disinfectant residual fades as water travels and waits. Older water gives disinfection byproducts more time to form. A storage tank that doesn't cycle can hold water long after it should have moved on. A sample result tells the utility what was true at one tap on one day. It says little about the street next to it, and nothing about next week.

Quality managers know this well. The practical question is what they can reasonably know about the places they don't sample, using data they already have.

Why it's hard

Water age isn't measured anywhere. No sensor reports how long the water at a hydrant has been in the system. Age has to be computed from how water moves through the network, which means it depends on a network model and on flows.

Tank level can mislead. A reservoir whose level barely moves looks stagnant. But a steady level can simply mean inflow and outflow are balanced while plenty of water passes through. The reverse can also be true. Level is evidence about storage, not proof of how long water stays in it.

Computed numbers get mistaken for measured ones. Once an estimate appears on a screen next to real readings, it is easy for it to be read, quoted and filed as if it had been sampled. For anything that touches compliance, that confusion is a real risk.

Models have limits. A network model with a disconnected section, a missing pipe or a stretch with no flow cannot say anything reliable about the water there. A tool that fills those gaps with a number anyway is worse than one that says nothing.

Plant compliance is often reviewed after the fact. Turbidity at the plant is measured continuously, but it is frequently reviewed in a periodic summary, where a short excursion can be averaged away or found late.

Principles of a good approach

Start with what the utility already holds. The network model, the telemetry and the disinfectant readings a utility already takes are usually enough to estimate a lot. The first step should not be a new sensor program.

Estimate residence time, and keep level beside it. Tank level alone can't say how long water stays in a reservoir, so residence time should be estimated directly. Level movement is still worth showing next to residence time, so staff can see when the two tell different stories.

Label computed values as computed. Every estimate of water age or residual should carry a clear label that it comes from a model, not a sample. It should guide where to sample and what to watch. It should never stand in for a compliance result.

Withhold with a reason. When the model can't support an estimate for part of the network, the tool should say so and give the reason in plain words. A blank with a reason is more useful than a confident guess.

Judge the plant against the rule, reading by reading. Turbidity should be compared against the turbidity requirements of the treatment rule, and every exceedance should be listed with its time, not folded into an average.

What it looks like in GDEA

Grid Data Enhanced Analytics (GDEA) joins the utility's GIS network model, SCADA telemetry and meter data into one live picture of electric, gas and water networks, with a hydraulic solver for water. Its water quality view uses that picture to answer questions the sampling program can't answer alone.

Reservoirs. For each reservoir, GDEA estimates residence time and shows it beside how much the level turns over, along with the chlorine reading at storage. When the level sits still but water is clearly passing through, the view says so. A flat level stays visible as something to watch, without being mistaken for stagnation.

Water age and chlorine residual. Across the network, GDEA estimates how old the water is when it reaches each part of the system, using the utility's own network model. It also estimates how much chlorine residual is likely left when that water arrives. Both appear under a clear "computed, never measured" label. Where the model can't support them, GDEA withholds the estimate and states the reason.

Turbidity. At the plant, GDEA reads treated-water turbidity against the turbidity requirements of the treatment rule and lists every exceedance, so a short excursion is visible when it happens rather than at month end.

GDEA is monitoring-only. It reports what the data and the model show. Sampling plans, flushing, tank operation and regulatory reporting stay with the utility's own staff and systems.

AN ILLUSTRATIVE EXAMPLE (HYPOTHETICAL)

The following scenario is hypothetical. A utility's east storage tank has shown almost no level change for weeks. Operations suspects it is stagnant and is weighing a drain-and-fill. In GDEA, the tank's estimated residence time turns out to be short. Water is moving through steadily, and the flat level simply reflects balanced inflow and outflow. The level is still shown, and the tank goes on the watch list rather than the work list.

Meanwhile, the network view shows the oldest water in the system at the end of a long residential loop on the far side of town, with a low estimated residual. None of the utility's routine sampling sites sits on that loop. The

quality manager adds a temporary sampling point there and asks operations to review flushing for the area. The sample, not the estimate, is what goes in the record.

A short cul-de-sac shows no estimate at all, with a note that no modeled flow reaches it. That prompts the GIS team to check for a missing connection in the model.

At the plant, one brief turbidity excursion appears in the exceedance list with its time. The compliance team reviews it that day instead of discovering it in a monthly summary.

Questions to ask any vendor

- Does the product estimate water age and disinfectant residual from our existing network model and flows, or does it need new sensors first?
- Is every computed value clearly labeled as computed, everywhere it appears?
- When the model can't support an estimate, does the product withhold it and say why?
- Does reservoir analysis estimate how long water stays in storage, or does it rely on tank level alone?
- Is plant turbidity judged against the treatment rule reading by reading, with every exceedance listed?
- Can a quality manager trace each estimate back to the readings and the model behind it?
- Do sampling, flushing and reporting decisions stay with our staff and systems?

Closing

Water quality is sampled at a few points and decided everywhere in between. Estimates drawn from the network a utility already has, labeled honestly and withheld when they can't be supported, help put the next sample, the next flush and the next question in the right place.

Learn more about Grid Data Enhanced Analytics at perinimble.com/grid-data-enhanced-analytics/, or talk with our team at perinimble.com/contact/.

ABOUT PERINIMBLE

PeriNimble builds mission-critical grid data software for energy and utility companies. Founded in 2022 and headquartered in Bloomington, Minnesota, we support more than 15 million meters and over 1 billion interval reads per day in production.

Built for systems that can't afford failure. · perinimble.com · perinimble.com/contact