

TALKING AQUIFERS

OECD CRP Fellowship
(Technical Report)



UNIVERSITY OF
ILLINOIS
URBANA-CHAMPAIGN



OECD

CARDIFF
UNIVERSITY
PRIFYSGOL
CAERDYDD

Talking Aquifer: AI-Powered Conversational Interface for Mahomet Aquifer Groundwater Data

Author

Charith Perera
School of Computer Science and Informatics
Cardiff University, United Kingdom

Collaborators

Praveen Kumar
Department of Civil and Environmental Engineering
Prairie Research Institute, University of Illinois Urbana-Champaign

Daniel Raymond Hadley
Prairie Research Institute, University of Illinois Urbana-Champaign

Funding

This research was supported by an OECD Co-operative Research Programme (CRP) Fellowship, which enabled a three-month research visit to the University of Illinois Urbana-Champaign (June–August 2025).

Version 1.0, December 9, 2025

Contents

	Project Overview (Front Matter)	6
	Executive Summary (Front Matter)	7
	Acknowledgments	7
1	Background	10
1.1	Background Executive Summary	10
1.2	Hydrogeologic Context	11
1.3	Water Resources Status	14
1.4	Critical Knowledge Gaps	16
1.5	The Data Integration Challenge	18
1.6	Why Talking Aquifer?	20
1.7	Essential Terms	22
1.8	Guide to Appendices	22
2	System Design	24
2.1	Design Objectives	24
2.2	Layered Architecture	25
2.3	Agent and Pipeline Layer	25
2.4	Data and Analytics Layer	31
2.5	Conversational Question-Answering Pipeline	34
2.6	End-to-End Dialogue Examples	34
2.7	Deployment and Operations	40
2.8	Integration with the Report Pipeline	40

3	Data	41
3.1	Overview: Four Primary Data Sources	41
3.2	Observation Well & Telemetry Dataset (Champaign County)	42
3.3	WARM Weather Network	51
3.4	HTEM Geophysical Survey	53
3.5	USGS Stream Gauging Network	54
4	Question-Answering Demonstration	56
4.1	Introduction	56
4.2	Category I: Basic Well Queries (Q001–Q010)	59
4.3	Category II: Spatial Analysis and Mapping	66
4.4	Category III: Temporal Trends, QA & Climate	79
4.5	Category IV: Multi-Variable Relationships	88
4.6	Category V: Network Analysis & Quality	99
4.7	Category VI: Temporal Aggregation (Q071–Q081)	115
4.8	Category VII: Multi-Source Integration (Q082–Q093)	128
4.9	Discussion and Future Directions	140
4.10	Conclusion	142
4.11	Summary: All 93 Questions	143
5	Conclusions and Future Directions	144
5.1	Summary of Achievements	144
5.2	Current Limitations	145
5.3	Future Directions	145
5.4	Closing Remarks	146
A	Technical Reference	147
A.1	Domain-Specific Terminology and Abbreviations	147
B	Detailed Water-Use Analysis	156
B.1	Historical Withdrawals by Sector	156
B.2	Municipal Supply Systems	157
B.3	Agricultural Irrigation Demand	158
B.4	Industrial and Energy Use	158
B.5	Economic Valuation Frameworks	158
B.6	Demographic Drivers	159
C	Regulatory and Management Framework	160
C.1	Statutory Authorities	160
C.2	Institutional Roles	161
C.3	Permit Programs	161
C.4	Standards and Monitoring Requirements	161

C.5	Protection Programs	162
C.6	Data Access and Confidentiality	162
D	Previous Research and Monitoring Infrastructure	163
D.1	Historic Well Networks	164
D.2	Geophysical and Tracer Investigations	164
D.3	Remote Sensing Contributions	164
D.4	Ongoing Monitoring Programs	164
D.5	Identified Data Gaps	164
E	Version History	165

Project Overview

Freshwater security hinges on our ability to observe, understand, and communicate the behaviour of subsurface aquifers in near–real time. This project—*Talking Aquifer*—combines hydrogeologic science with modern language-model technology to create a conversational decision-support platform for the Mahomet Aquifer system beneath Champaign County, Illinois. By fusing physical measurements, domain ontologies, and large-language-model (LLM) reasoning, the work advances both groundwater science and cyber-infrastructure engineering.

Scientific Objectives.

The core scientific aim is to characterise spatio-temporal dynamics of water-table elevation and quality using heterogeneous data streams: (i) continuous pressure-transducer records from observation wells, (ii) periodic laboratory analyses of geochemical tracers, and (iii) ancillary meteorological and pumping-log data that influence recharge and draw-down. We test the hypothesis that integrating these streams with physics-informed analytics (e.g., transient flow models, geostatistical kriging) yields more accurate state estimates and early-warning indicators than treating each dataset in isolation. A secondary scientific question explores whether an LLM—augmented with domain ontologies such as SOSA/SSN and the Illinois State Geological Survey vocabulary—can reliably translate natural-language questions into the structured queries or analytical pipelines required to extract those indicators.

Engineering Contributions.

To operationalise the science, the project delivers a modular software stack with four vertically integrated layers:

1. **Ingestion & Storage.** A DuckDB/SQLite lakehouse ingests raw CSV/TXT sensor exports and GeoPackage shapefiles, while an object store archives high-frequency time-series.
2. **Semantic Mediation.** A lightweight RDF graph links every well, sensor, and hydrostratigraphic unit to canonical URIs, enabling ontology-based reasoning and cross-dataset joins.
3. **Analytics Micro-services.** Python micro-services expose groundwater computations—flow-net generation, hydrograph decomposition, pump-test curve fitting—through a tool interface that LangChain agents can call programmatically.
4. **Conversational Interface.** A locally hosted LLM (e.g. Mistral-7B-Instruct) is wrapped in a Retrieval-Augmented Generation (RAG) layer that selects relevant snippets from the semantic store, synthesises answers, and optionally returns figures produced on-the-fly in Jupyter or as standalone HTML visualisations.

All components are containerised for reproducibility and can be redeployed to other aquifers by swapping ontologies and data connectors.

Research Significance.

Talking Aquifer pushes the frontier in two ways. First, it demonstrates that domain-specific ontologies are powerful scaffolds for aligning physical models with machine-learning agents, reducing hallucinations while preserving conversational flexibility. Second, it offers a transparent, open-source template for water-resource monitoring that lowers the barrier for small utilities and research teams to exploit advanced analytics without specialised GIS or database expertise.

Intended Audience & Utility.

Hydrogeologists gain an extensible toolkit for rapid hypothesis testing; water-utility managers receive intuitive dashboards and alerts driven by scientifically defensible models; and educators can use the platform to teach integrative groundwater science. The architecture is also a test-bed for broader environmental applications—rivers, forests, urban drainage—where sensor networks

generate vast, under-utilised data.

Report Structure.

The remainder of this document unfolds as follows. Chapter 1 introduces groundwater concepts for non-specialists; Chapter 2 explains the software architecture and evaluation results; Chapter 3 details the integrated datasets; Chapter 4 catalogues the natural-language queries the system can answer; and Chapter 5 discusses limitations, policy implications, and future research directions.

Together, these chapters show how merging rigorous hydrogeologic analysis with cutting-edge AI can make subsurface science more accessible, actionable, and adaptable to the water-security challenges of the twenty-first century.

Executive Summary

This report delivers a conversational, data-driven view of the Mahomet Aquifer system by combining observation wells, climate records, and LLM-based question-answering. Key takeaways:

- **Stable long-term behaviour with seasonal stress:** Most wells show modest decadal trends, with 1–3 m seasonal swings driven by recharge and summer demand.
- **Gaps in flank and transition coverage:** Monitoring density is highest along the buried-valley axis; stream–aquifer interfaces and flanks merit additional sites.
- **LLM Q&A augments analysis:** Natural-language queries retrieve provenance-linked figures, tables, and maps, reducing time-to-insight for common operational questions (Q001–Q093).
- **Actionable early warnings:** Simple trend and event detectors (e.g., five-year slopes, hot-day co-occurrence) highlight emerging risks and recovery.

Acknowledgments

This work was made possible by an OECD Co-operative Research Programme (CRP) Fellowship, which supported a three-month research visit to the University of Illinois Urbana-Champaign during June–August 2025. The author gratefully acknowledges the fellowship programme for enabling this international collaboration.

The author thanks Professor Praveen Kumar (Department of Civil and Environmental Engineering, Prairie Research Institute, UIUC) for serving as the primary collaborator and host, providing invaluable guidance on hydrogeologic data integration and the Mahomet Aquifer system. Special thanks to Daniel Raymond Hadley (Prairie Research Institute, UIUC) for collaboration on data access, system architecture, and technical discussions throughout the project.

The Illinois State Water Survey (ISWS) and the Water and Atmospheric Resources Monitoring (WARM) program provided the observational datasets that form the foundation of this work. The U.S. Geological Survey (USGS) streamflow data were accessed through their public APIs.

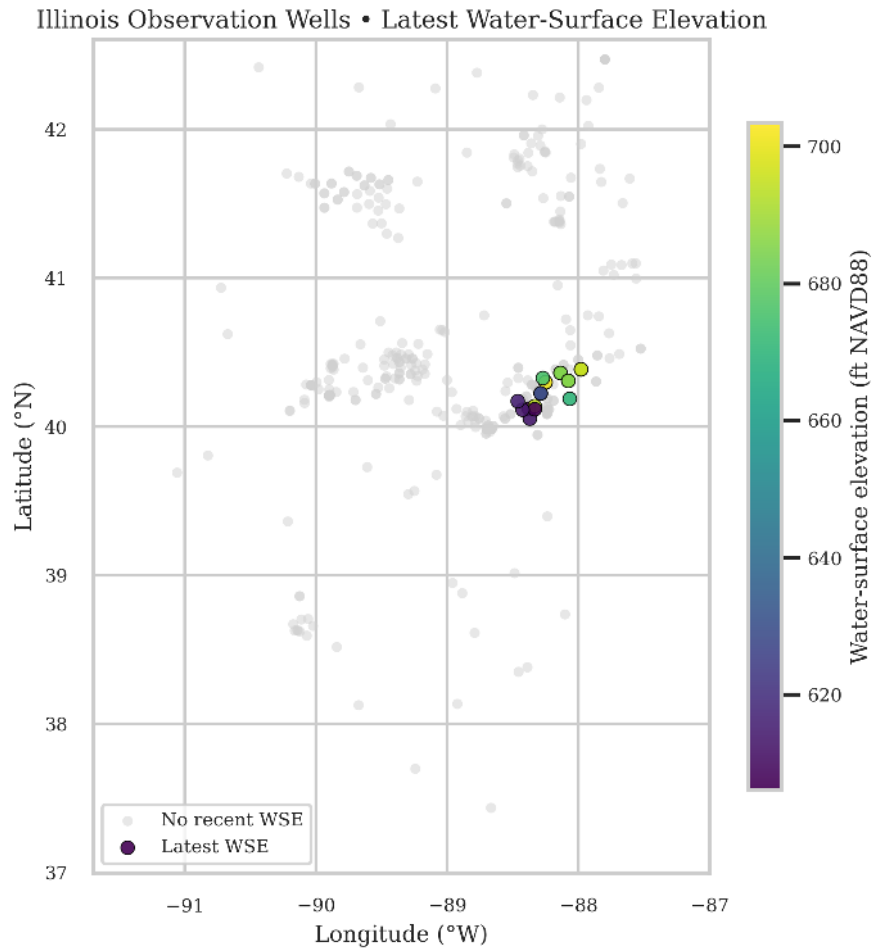


Figure 1: Network footprint and density (Q040).

WHAT.

Locations of observation wells across Champaign County; units in geographic degrees on plain longitude/latitude axes.

How.

Flat point symbology by well, with recent water-surface elevations highlighted against the broader statewide inventory.

WHY.

Shows spatial coverage and gaps (flanks, stream transitions). Caveat: static snapshot—membership changes over time.

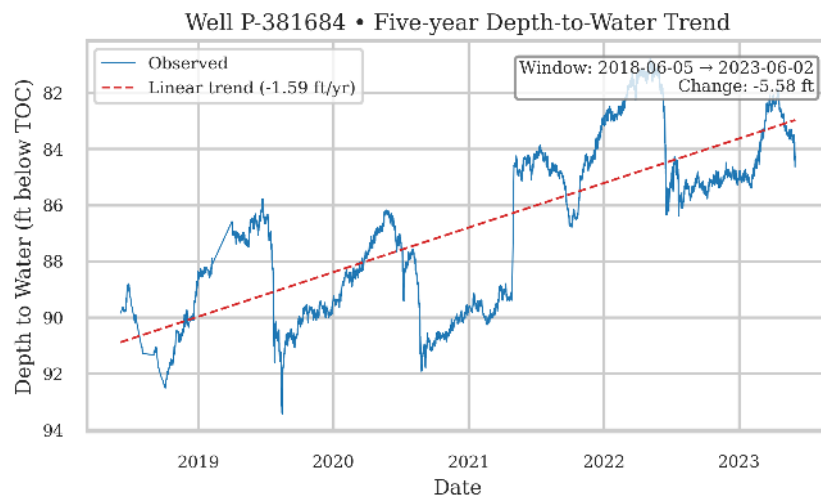


Figure 2: Five-year well trend example (Q031).

WHAT.

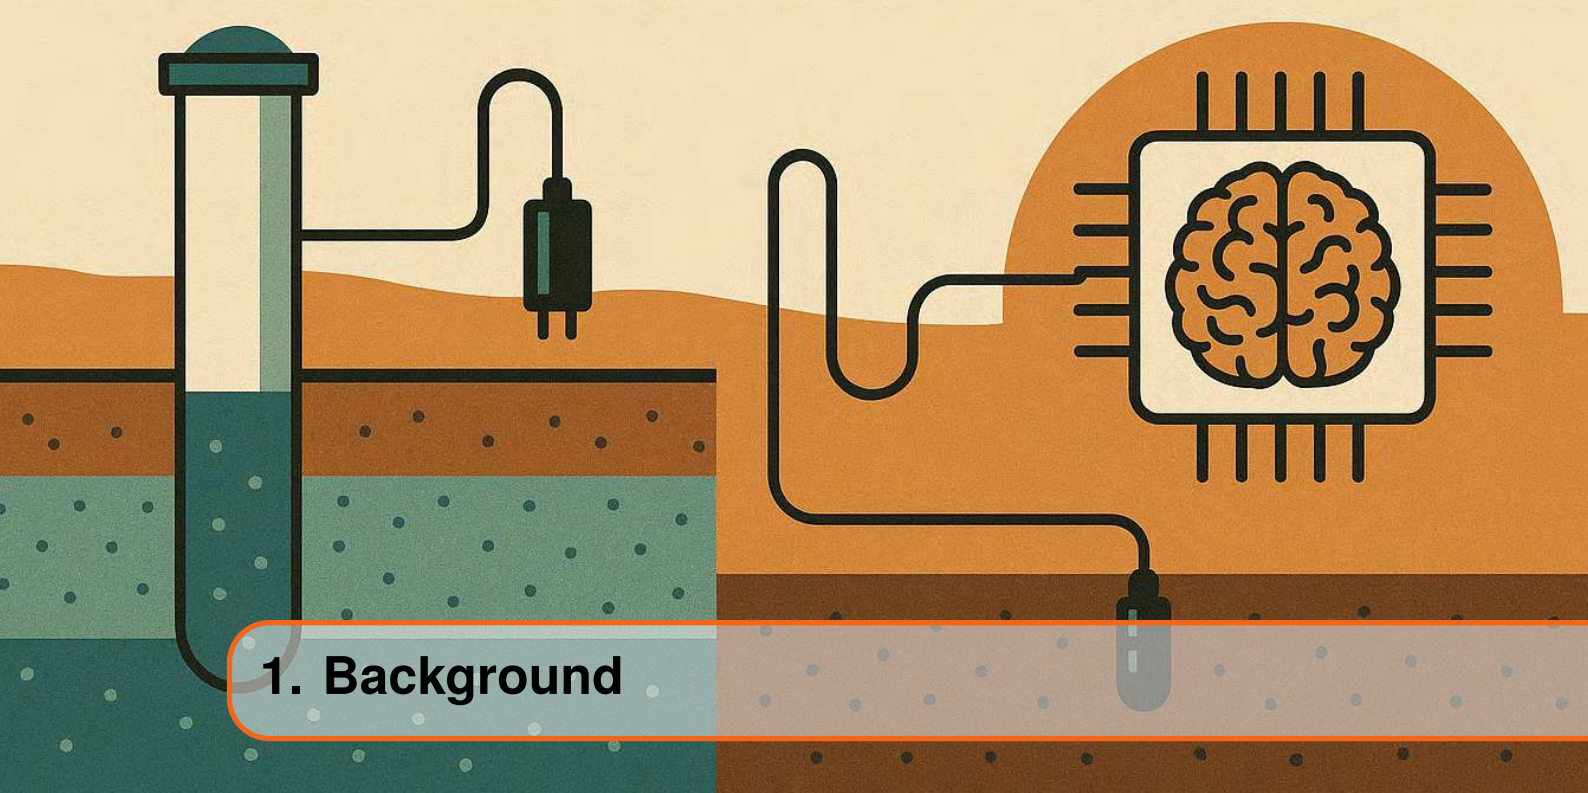
Five-year water-surface elevation slope (ft/yr) for a representative well; time period shown in title.

HOW.

Ordinary least squares fit over the last five years; slope and total change annotated directly on the plot.

WHY.

Highlights recent direction and magnitude of change for operational planning. Caveat: sensitive to short gaps and outliers—see QA notes.



1. Background

1.1 Background Executive Summary

The Mahomet Aquifer System, buried beneath glacial deposits in east-central Illinois, provides the sole water supply for over 500,000 residents and supports agricultural production valued at over \$2 billion annually. Despite its critical importance, groundwater management faces three fundamental challenges that *Talking Aquifer* addresses:

1. Sparse and delayed observations.

Current monitoring consists of 356 statewide observation wells, of which only 18 (5%) transmit real-time data. Most records arrive monthly or annually, creating decision lags of 30–90 days. Critical zones—valley flanks, stream-aquifer transitions, and deep confined intervals—remain under-instrumented, leaving wide uncertainty in drought forecasts and sustainable yield estimates.

2. Fragmented data systems.

Groundwater information resides in six incompatible systems: state observation networks (ISWS), withdrawal inventories (IWIP), water-quality compliance (IEPA/SDWIS), stream gauges (USGS), municipal operations (proprietary SCADA), and one-time research campaigns. Analysts spend 40% of their time harmonizing coordinates, units, identifiers, and temporal references before any scientific analysis begins.

3. Invisible uncertainty.

Existing tools provide single-value forecasts ("water levels will be X meters") without confidence intervals, leading to overconfident long-term plans. When drought conditions deviate from the forecast, communities lack quantitative guidance on the probability of supply shortfalls or the value of new monitoring investments.

The Talking Aquifer solution.

This project integrates real-time telemetry, geophysical surveys, and ensemble modeling into a unified semantic layer accessible through natural language queries. Stakeholders can ask "What is the probability that Well P-268557 will drop below 195 m during next summer's peak?" and receive

probabilistic answers (P10/P50/P90) with full provenance linking back to measurements, model parameters, and uncertainty sources.

This background chapter provides essential context:

- **Section 1.2:** Hydrogeologic essentials—aquifer geometry, key parameters, and groundwater dynamics
- **Section 1.3:** Current water balance, use by sector, and stress indicators
- **Section 1.4:** Five critical gaps that Talking Aquifer addresses
- **Section 1.5:** Interoperability challenges and integration architecture
- **Section 1.6:** Project motivation and links to subsequent chapters

- C** For readers new to groundwater, this chapter is designed as a gentle ramp: focus first on the Background Executive Summary and the key figures in Sections 1.2 and 1.3 to build intuition about the aquifer and its stresses. You can safely skip detailed parameter values or references on a first reading and return to them later if you need more technical depth.

Readers needing detailed technical specifications, economic frameworks, or regulatory details should consult Appendices A–D. Section 1.7 provides a quick-reference glossary of key concepts.

1.2 Hydrogeologic Context

1.2.1 The Mahomet Aquifer System

The Mahomet Aquifer occupies a buried bedrock valley trending east-west across Champaign County at depths of 30–60 m below land surface. Ancient rivers incised this valley into Paleozoic bedrock; subsequent Pleistocene glaciation filled it with highly permeable sands and gravels (outwash) interbedded with low-permeability tills. The resulting aquifer extends 30–60 m thick along the valley axis, thinning to 10–20 m along margins where tills thicken. Concepts in this section follow established syntheses of groundwater systems and surface–groundwater connectivity [10].

Hydrostratigraphic units.

The valley fill comprises three principal units:

- (a) **Upper till** (Glasford/Wedron): 5–25 m thick; acts as confining layer over much of the system.
- (b) **Mahomet Sand** (Banner/Glasford outwash): 30–60 m of clean sand and gravel; primary aquifer unit.
- (c) **Bedrock** (shale/carbonate): Impermeable base; bedrock topography controls aquifer geometry.

Where tills are thin or absent—particularly over upland sand dunes and along stream valleys—the Mahomet transitions to unconfined conditions with direct recharge from precipitation. Elsewhere, the aquifer is confined or semi-confined, receiving recharge through focused windows and vertical leakage.

Key hydraulic properties.

Field tests and model calibration yield typical ranges (Table 1.1):

These parameters anchor the transient flow models used in Talking Aquifer; their spatial variability and associated uncertainty are quantified through ensemble calibration (Chapter 2).

1.2.2 Groundwater Dynamics

Recharge.

Annual recharge averages 150–200 mm/yr over the valley (2009–2022 mean), with strong seasonal concentration: 70–80% occurs November–April when precipitation exceeds evapotranspiration and

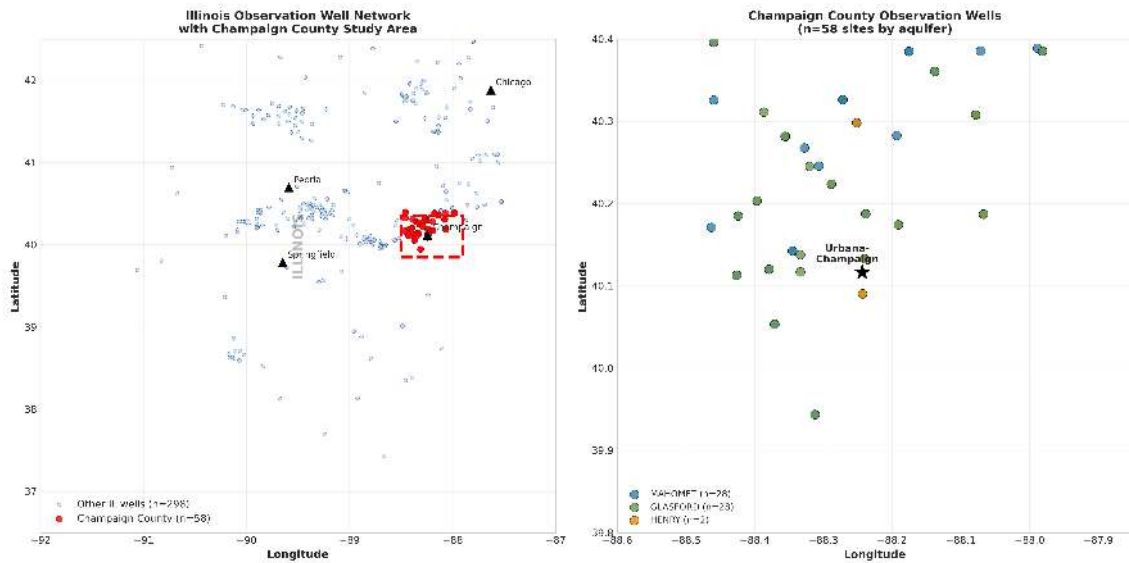


Figure 1.1: Study area location. **Left:** Illinois observation well network ($n=356$) with Champaign County highlighted in red. The county is located in east-central Illinois, approximately 200 km south of Chicago along the Interstate 74 corridor. **Right:** Champaign County observation wells ($n=58$) color-coded by aquifer unit, showing the concentration of monitoring sites along the buried Mahomet Aquifer valley axis.

Table 1.1: Typical hydraulic properties of the Mahomet Aquifer.

Parameter	Range	Notes
Hydraulic conductivity K	10^{-5} – 10^{-3} m/s	Sand/gravel facies; $K_v/K_h \approx 0.1$
Transmissivity T	1000–5000 m^2/d	Along valley axis (30–60 m thick)
Storativity S (confined)	0.0001–0.001	Elastic storage in confined zones
Specific yield S_y	0.15–0.25	Unconfined (recharge windows, flanks)
Porosity ϕ	0.25–0.35	Effective porosity ≈ 0.20

vegetation is dormant. Summer recharge is limited to episodic storm events over thin-till zones and losing stream reaches.

Recharge pathways include:

- **Diffuse infiltration** through thin tills and sandy uplands
- **Focused recharge** along stream channels (losing reaches) and detention basins
- **Vertical leakage** through aquitards in confined zones (slow, long-term contribution)

Discharge.

Natural discharge occurs via:

- **Baseflow to streams:** Salt Fork and Sangamon Rivers intercept regional flow; typical summer baseflow contributions are 30–50% of streamflow
- **Evapotranspiration:** Where water tables are shallow (<1.5 m), riparian vegetation extracts 2–4 mm/d during peak summer
- **Subsurface outflow:** Regional gradients direct flow westward toward the Illinois River

Anthropogenic discharge dominates the water budget:

- **Municipal pumping:** Champaign-Urbana withdraws ~ 57 MLD; Danville, Rantoul, and smaller communities add 15 MLD

- **Irrigation:** Peak-season (June–August) withdrawals reach 10–15 MLD countywide
- **Industrial/energy:** Cooling and process water add 5 MLD

Water-table dynamics.

Seasonal head fluctuations of 1–3 m are typical, with highest stands in May (post-recharge) and lowest in October (end of growing season and peak pumping). Multi-year droughts impose cumulative declines: the 2010–2012 drought lowered heads by 3–5 m in the Champaign wellfield, requiring 3+ years for full recovery.

Talking Aquifer’s forecasting module propagates seasonal and interannual climate variability through ensemble water-balance models to generate probabilistic head forecasts (Chapter 2).

1.2.3 Critical Hydrogeologic Features

Three features strongly influence monitoring design and model structure:

1. Recharge windows.

Where tills are thin (<3 m) or absent, infiltration can exceed 400 mm/yr locally—double the regional average. These windows provide disproportionate recharge volume but are under-monitored in legacy networks. Talking Aquifer prioritizes targeted telemetry at window sites to capture event-scale responses.

2. Stream-aquifer interaction zones.

The Salt Fork and Sangamon alternate between gaining and losing over short distances, responding to seasonal stage and pumping stress. Correct characterization of riverbed conductance (C) is essential for attributing water-level trends and simulating drought baseflow. The platform integrates USGS stream-gage data and plans fiber-optic temperature sensing to constrain C at key reaches.

3. Confined/unconfined transitions.

Valley-axis wells are confined (artesian); flank wells transition to unconfined. Storage response differs by an order of magnitude ($S \approx 0.0005$ confined vs. $S_y \approx 0.20$ unconfined), affecting drawdown rates and recovery timescales. Multi-level nests at transition zones are prioritized for uncertainty reduction.

1.2.4 Climate and Recharge Budget

Reliable forecasts require quantifying how precipitation, evapotranspiration, and land use control recharge to the Mahomet Aquifer. Talking Aquifer fuses NOAA climate normals, SCAN soil-moisture observations, and MODIS evapotranspiration products to estimate distributed recharge at weekly time steps. Methodologically, we follow standard approaches for estimating recharge from groundwater levels and water-balance constraints [3].

Seasonal Recharge Dynamics

Recharge peaks during late winter and early spring when soils are moist, vegetation is dormant, and precipitation falls as rain rather than snow. Water-table-fluctuation analyses indicate median diffuse recharge of 180 mm/yr, with 60 % delivered between February and May. Summer recharge declines below 20 mm/month as crops and turf evapotranspire most infiltrating water. Focused recharge occurs after intense convective storms along thin-till corridors, an effect captured by the platform’s event-scale infiltration module.

Surface–Groundwater Exchange

Stream-aquifer interaction modulates both flood attenuation and drought resilience. Gaining reaches along the Salt Fork contribute up to 40 % of baseflow during dry spells, while short losing segments supply focused recharge that recharges the aquifer within weeks. Fiber-optic distributed

temperature sensing (DTS) and seepage-meter data—streaming into Talking Aquifer—constrain riverbed conductance parameters used in coupled groundwater/surface-water simulations.

Climate Impact Scenarios

Regional climate projections (CMIP6, downscaled via LOCA) suggest warmer summers, higher vapor-pressure deficits, and more frequent multi-day dry spells. Scenario ensembles perturb precipitation and temperature drivers within Talking Aquifer’s water-balance model, yielding probabilistic forecasts of recharge deficits and irrigation demand surges. These outputs feed directly into Chapter 2’s resilience metrics and Chapter 4’s what-if analyses, and are consistent with broader assessments of climate impacts on groundwater [8].

1.3 Water Resources Status

1.3.1 Current Water Balance

Table 1.2 summarizes the Mahomet Aquifer water budget for 2020–2022 (near-normal climate), framed in the broader context of regional-to-global assessments of groundwater depletion and stress [2, 6, 7]:

Table 1.2: Annual water balance for the Mahomet Aquifer (Champaign County), 2020–2022 average.

Component	Volume (Mm ³ /yr)	Notes
<i>Inputs</i>		
Recharge	450	180 mm/yr over 2,500 km ²
<i>Outputs</i>		
Municipal pumping	75	Champaign-Urbana (57) + others (18)
Irrigation	25	Peak June–August
Industrial/energy	15	Cooling, process water
Baseflow to streams	335	Salt Fork, Sangamon, smaller tributaries
Net balance	±50	Seasonal storage change

Under normal conditions, the system is approximately balanced on an annual basis. Seasonal storage fluctuations of ±50 Mm³ (equivalent to 20 mm average head change over the aquifer area) reflect spring recharge and summer drawdown. Multi-year droughts or sustained over-pumping can induce cumulative deficits; conversely, wet cycles rebuild storage.

1.3.2 Water Use by Sector

Figure 1.2 shows the distribution of withdrawals (total: 115 Mm³/yr):

Municipal sector (75 Mm³/yr).

Champaign-Urbana accounts for 76% of municipal demand (57 MLD average, 85 MLD peak day). Per-capita use averages 450 L/cap/d (120 gal/cap/d), consistent with Midwestern cities with mature infrastructure. Seasonal peaking factor (MDD/ADD) is 1.5, driven by summer outdoor use. Non-revenue water (NRW) is estimated at 10–12%, primarily real losses (leaks) in aging mains.

Agricultural irrigation (25 Mm³/yr).

Concentrated June–August; 60% of annual volume withdrawn in 90 days. Irrigated acreage has grown 3%/yr over the past decade, primarily corn and soybeans on sandy soils. Focused monitoring

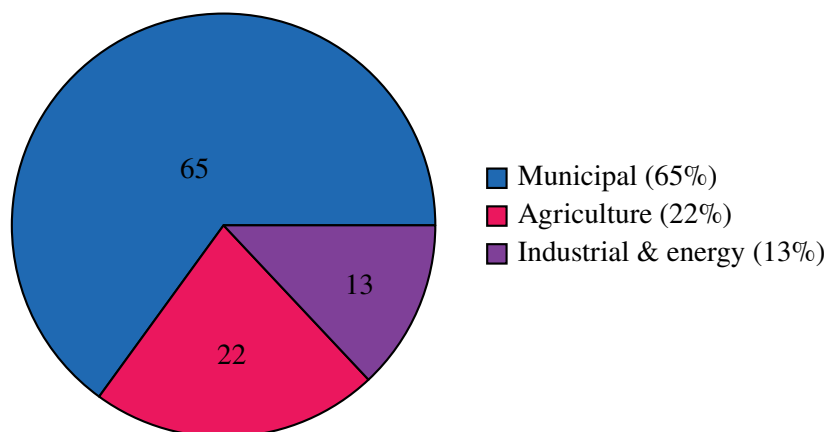


Figure 1.2: Water use by sector in Champaign County, 2020–2022 average. Municipal supply dominates (75 Mm³/yr, 65%), followed by agricultural irrigation (25 Mm³/yr, 22%) and industrial/energy users (15 Mm³/yr, 13%).

at irrigation clusters is planned to capture peak-day cones of depression.

Industrial/energy (15 Mm³/yr).

Dominated by cooling-tower makeup and process water. Demand is relatively steady year-round, with slight summer peak for cooling. Water-quality requirements (hardness, dissolved solids) influence source selection.

1.3.3 Stress Points and Trends

Summer drawdown.

Peak pumping during July–August can lower heads 2–4 m below spring high-stands. The 2012 drought saw maximum drawdown of 5.2 m at sentinel wells, approaching levels where deeper production wells experience reduced specific capacity. Recovery to pre-drought levels took until 2015.

Demand growth.

Municipal demand is growing 1.5%/yr (population + economic activity); irrigation demand 3%/yr (acreage expansion). Without augmented recharge or demand-side management, summer drawdown is projected to deepen 0.5–1 m per decade (ensemble median, Section 1.4).

Climate variability.

Observed trends (1980–2020) include:

- Earlier spring snowmelt (5–7 days/decade), shifting recharge timing
- Increased rainfall intensity (fewer days, larger storms), favoring focused recharge over diffuse infiltration
- More frequent multi-year dry/wet cycles, amplifying storage swings

Emerging contaminants.

PFAS compounds (PFOA, PFOS) have been detected in 12% of sampled public-supply wells at concentrations near or above recent MCLs (4 ng/L combined). Source attribution and plume delineation are ongoing; Talking Aquifer’s geochemical module links well chemistry to land-use and recharge pathways for targeted source control.

1.4 Critical Knowledge Gaps

Effective groundwater management requires answering questions like "Will next summer's peak demand exceed sustainable supply?" or "Where should we add monitoring to halve forecast uncertainty?" Five knowledge gaps currently limit our ability to provide defensible answers; Table 1.3 maps each gap to its management consequence and Talking Aquifer's response.

Table 1.3: Priority knowledge gaps, decision impacts, and platform solutions.

Knowledge Gap	Decision Impact	Talking Aquifer Solution
Sparse flank coverage	Uncertain drought recovery time (± 2 years P10–P90)	Geophysics-informed priors + targeted nests at flanks and transitions
Monthly/annual reporting lag	30–90 day delay in situational awareness	Real-time telemetry + daily data assimilation + IRF nowcasting
No vertical gradient data	Ambiguous recharge pathways; leakage vs. storage trade-off	Multi-level nests (3 planned) + streambed heat sensing
Fragmented data systems	40% of analyst time spent on harmonization	Unified semantic layer with canonical IDs, units, datums
Uncertainty invisible	Overconfident long-term plans; poor trigger design	Ensemble forecasts (P10/P50/P90) + value-of-information metrics

1.4.1 Gap 1: Spatial Coverage (Flanks and Transitions)

Current observation wells cluster along the buried-valley axis and near municipal wellfields; flank regions and stream-aquifer transition zones have 5–10 \times lower monitoring density. This sparsity leads to:

- **Equifinality in calibration:** Models can trade off vertical conductivity (K_v), riverbed conductance (C), and storage (S) to fit axis heads, yielding non-unique parameter sets.
- **Wide forecast bands:** Summer 2024 drawdown forecasts at flank sentinels show P10–P90 spreads of 2–3 m (vs. 0.5 m along axis), translating to ± 2 year uncertainty in drought recovery time.

Solution.

(a) Airborne electromagnetic (AEM) surveys map sand/till architecture at 50 m resolution, entering the model as spatial priors on facies and K ; (b) three multi-level nests (two flanks, one stream transition) quantify vertical gradients and leakage; (c) information-driven site selection prioritizes locations that maximize reduction in P90 forecast width.

1.4.2 Gap 2: Temporal Resolution and Latency

Non-telemetered wells are measured monthly by field crews; withdrawal data arrive annually via IWIP. This latency creates:

- **Attribution lag:** A 2 m head decline observed in October could reflect summer pumping, reduced recharge, or both—disambiguation requires sub-monthly pumpage and climate forcing.
- **Missed event-scale responses:** Storm-driven recharge pulses (5–20 cm head rises over days) are unobserved, preventing calibration of focused recharge parameters.

Solution.

(a) Expand telemetry from 18 to 40+ sites, prioritizing recharge windows and transition zones; (b) develop impulse-response functions (IRFs) mapping pump changes and storm events to head responses for short-horizon (7–30 day) nowcasting; (c) implement fixed-lag Kalman smoothing to update recent states when delayed observations arrive, preserving "as-of" provenance.

1.4.3 Gap 3: Vertical Gradients and Leakage

Existing wells sample single screened intervals (typically 3–6 m); nested piezometers at multiple depths are rare. Without vertical gradient data:

- **Recharge pathway ambiguity:** Is infiltration rapid (focused windows) or slow (diffuse leakage)? Vertical head profiles distinguish these mechanisms.
- **Storage vs. leakage trade-off:** Models can match axis drawdown by adjusting S (local elastic) or K_v (leakage from overlying units); vertical nests break this degeneracy.

Solution.

(a) Install three multi-level nests (screened at 5, 15, 30 m depths) at priority locations identified via A-optimal design; (b) deploy fiber-optic distributed temperature sensing (DTS) cables in two streambed reaches to map vertical flux at high spatial density (10 m spacing); (c) use vertical gradients as assimilation targets alongside heads.

1.4.4 Gap 4: Data Integration and Interoperability

Six systems house groundwater information, each with different spatial references (NAD27/NAD83/State Plane), unit conventions (ft/m, gpm/MLD), identifiers (local/state/federal), and temporal stamps (local/UTC, varying DST treatment). Consequences:

- **Analyst time sink:** Integrating a new dataset requires weeks of schema reverse-engineering, coordinate transforms, and unit checks—40% of total analysis time.
- **Silent errors:** Undetected datum mismatches can introduce 1–3 m head offsets; unit errors can corrupt pumpage by factors of 3–4.
- **Irreproducibility:** Ad-hoc joins lack provenance; reanalyses may yield different results if joins are re-done differently.

Solution.

(a) Unified semantic layer assigns canonical `site_uuid` to every well/gauge, with bidirectional links to source IDs; (b) automated ETL pipelines convert all inputs to SI units, NAD83/NAVD88, and UTC, retaining raw values for audit; (c) schema validators flag unit/datum inconsistencies at ingestion; (d) every figure and table links back to the input slice with full conversion provenance.

1.4.5 Gap 5: Uncertainty Quantification and Decision Support

Legacy forecasting tools provide single-value outputs ("*summer 2025 head at Well X will be 198.5 m*") without confidence intervals, even when parameter and forcing uncertainty are large. This leads to:

- **Overconfident plans:** A 10-year capital plan designed around a median forecast may fail during a P10 drought.
- **Poor trigger design:** Drought response triggers set without probabilistic thresholds may activate too late or trigger false alarms.
- **Sub-optimal monitoring:** No quantitative basis for comparing the value of a new nest vs. a streambed sensor vs. extending an existing telemetry record.

Solution.

(a) Ensemble Kalman filter (EnKF) with 200-member parameter and forcing ensembles generates probabilistic forecasts (P10/P50/P90) for heads, storage, and baseflow; (b) chance-constraint formulations ($\Pr[h \geq h^{\min}] \geq 0.90$) inform robust drought triggers; (c) value-of-information (VoI) metrics quantify expected reduction in decision loss from new data, guiding monitoring investment.

1.5 The Data Integration Challenge

1.5.1 Current Landscape: Six Incompatible Systems

Groundwater data in Illinois reside in six primary systems:

1. **ISWS Observation Network:** 356 wells statewide, 58 in Champaign County; 18 telemetered (hourly), remainder manual monthly. Units: feet below measuring point. Datum: mixture of NGVD29 and NAVD88. Format: CSV + custom metadata.
2. **IWIP (Withdrawals):** Annual reporting from 3,000+ high-capacity facilities. Units: gallons/day (gal/d). Identifiers: facility permit number (no spatial join to well coordinates). Format: Excel + PDF annexes.
3. **IEPA/SDWIS (Compliance):** Quarterly finished-water chemistry from 250 public supplies. Units: mg/L (mixture of "as N", "as CaCO₃", element). Identifiers: PWS ID. Format: relational DB (restricted access).
4. **USGS Stream Gauges:** 15-minute stage and discharge at 12 regional sites. Units: feet (USGS datum), cubic feet per second. Identifiers: USGS site numbers. Format: JSON API + tab-delimited archives.
5. **Municipal SCADA:** Real-time well flow, pump state, tank levels. Units: site-specific (gpm, MGD, kW). Identifiers: asset tags (no cross-system linkage). Format: proprietary (5 different vendors).
6. **Research Campaigns:** One-time geophysical surveys (AEM, seismic), tracer tests, geochemistry panels. Units and formats vary by project. Metadata often incomplete or in PDF reports.

Key interoperability barriers:

- **No common identifiers:** A single well may appear as "P-268557" (state), "CHAM-042" (county), "Well 3" (municipal), or lat/long coordinates with varying precision. Spatial joins fail when coordinates differ by 10–100 m due to datum or precision issues.
- **Unit chaos:** Flow in gpm, gal/d, MLD, MGD, cfs; heads in feet (NGVD29, NAVD88, or undefined datum), meters, or depth below land surface vs. measuring point; chemistry in mg/L "as N", "as ion", or "as CaCO₃".
- **Temporal inconsistency:** Timestamps in local time (with/without DST), UTC, or unspecified. Daily averages may be midnight-midnight, calendar-day, or pumping-shift aligned.
- **Silent errors:** Mixing NGVD29 and NAVD88 introduces 0.5–1.5 m offsets; confusing gal/d and MGD creates 10⁶× pumpage errors; "as N" vs. "as NO₃⁻" gives 4.4× concentration discrepancies.

1.5.2 Talking Aquifer Integration Architecture

The platform implements a three-layer architecture (Figure 1.3):

Layer 1: Source connectors.

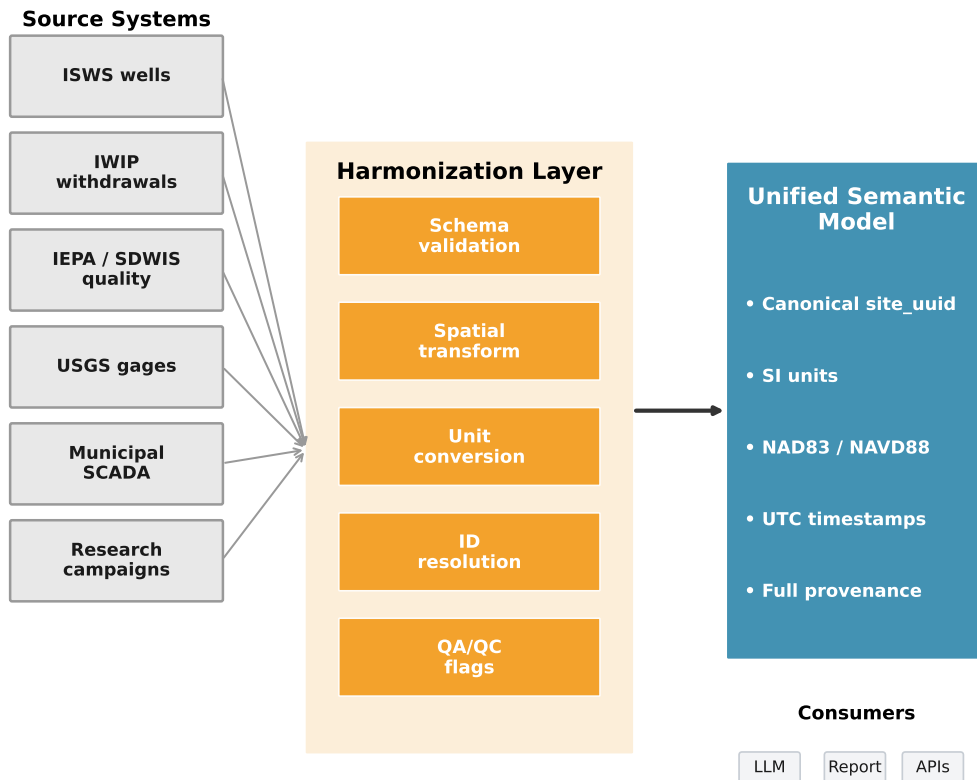
Custom adaptors for each system handle authentication, pull schedules, and format parsing. Metadata (steward, retrieval date, license, schema version) is captured at ingestion.

Layer 2: Harmonization pipelines (ETL).

- (a) **Spatial:** All coordinates transformed to NAD83 (horizontal) and NAVD88 (vertical). Land-surface and measuring-point elevations retrieved from survey databases; heads computed as $h = LSE - DTW$ or $h = MP - DTW + \Delta_{MP}$. Horizontal uncertainty flagged for legacy township-centroid locations.
- (b) **Units:** Convert measurements to a consistent set of canonical units (e.g., flow in MLD/ MGD, hydraulic head in metres NAVD88, depth-to-water in feet below measuring point, recharge

Data Integration Architecture

From source systems to a unified semantic model



Every derived value keeps provenance for audit and reproducibility

Figure 1.3: Data integration architecture. Source systems feed an automated harmonization layer (ETL pipelines with schema validators), producing a unified semantic model with canonical identifiers, units (SI), datums (NAD83/NAVD88), and timestamps (UTC). All conversions retain provenance for audit. Analytics and conversational queries operate on the unified layer.

- in mm/d, and chemistry in mg/L). Chemistry speciation is standardised ("as N" for nitrate, element for metals). Raw values and conversion factors are retained in the **provenance** field.
- (c) **Identifiers**: Assign canonical `site_uuid` via spatial/attribute clustering (wells within 50 m with matching screen depths are merged). Preserve source IDs (state, county, federal) as alternate keys.
- (d) **Temporal**: All timestamps converted to UTC (ISO 8601). Local timezone and DST rules recorded. Daily aggregations use UTC midnight–midnight unless specified.
- (e) **QA/QC**: Automated range tests (physically plausible head/temp/flow limits), step-change detection (pump cycles, logger maintenance), and cross-checks (tank mass balance for SCADA). Flags stored as `qc_status` codes.

Layer 3: Unified semantic model.

A graph database (RDF/OWL) links wells → screen intervals → formations → hydrostratigraphic units, with properties (coordinates, construction, time-series endpoints) as typed literals. SPARQL queries traverse relationships (e.g., "all wells screened in Mahomet Sand within 1 km of the

Champaign wellfield"). Time series stored in DuckDB (columnar analytics), linked via `site_uuid`.

Provenance and audit.

Every derived value (head from depth, unit conversion, interpolated gap-fill) carries a provenance chain: source file → transformation → output. Users can trace any figure back to raw measurements. Version control ensures reproducibility: rerunning analysis on 2024-01-15 data snapshot yields identical results.

1.6 Why Talking Aquifer?

1.6.1 The Management Challenge

A municipal water manager in Champaign County must answer:

- *"Will we have sufficient supply through next summer's peak, given current storage and climate forecast?"*
- *"If we add a new production well, how will it affect sentinel well heads and baseflow to the Salt Fork?"*
- *"What is the value of installing a new monitoring nest at Site X vs. extending telemetry at Site Y?"*

With existing tools, this requires:

1. Manually downloading data from 6 systems (2–3 days)
2. Harmonizing coordinates, units, identifiers (1 week, with risk of silent errors)
3. Running a groundwater model calibrated to historical data (1–2 weeks)
4. Interpreting single-value outputs without uncertainty bounds
5. Repeating the entire process when new observations arrive

Total latency: 3–4 weeks. Result: A forecast without confidence intervals, outdated before it's delivered, and irreproducible if an analyst leaves.

1.6.2 The Technical Opportunity

Three recent advances make an integrated solution feasible:

1. Sensor networks and IoT.

Low-cost pressure transducers (\$200–500), cellular telemetry modules (\$50/year data plans), and solar power enable real-time monitoring at scale. Cloud ingestion pipelines handle high-frequency streams (5-minute) with sub-hour latency.

2. Semantic web and data integration.

Graph databases (RDF/OWL) represent complex relationships (wells → formations → model layers) natively. Schema.org vocabularies and W3C standards (SOSA/SSN for sensors) provide interoperability frameworks. Automated ETL tools (Apache Airflow, dbt) reduce harmonization time by 10×.

3. AI/ML for natural language and ensembles.

Large language models (LLMs) translate conversational queries ("*Show me wells near Urbana with PFAS above the MCL*") into structured queries (SPARQL, SQL). Ensemble Kalman filters and probabilistic programming (PyMC, Stan) propagate uncertainty from parameters and forcing through forecasts, producing actionable confidence intervals.

1.6.3 The Talking Aquifer Solution

This project delivers four capabilities that transform Champaign County groundwater management:

1. Unified data layer.

All six source systems accessible through one interface. Canonical identifiers eliminate join ambiguity; automated harmonization ensures unit/datum consistency; full provenance enables audit and reproducibility. *Result*: Analysis time cut from weeks to hours.

2. Real-time assimilation.

Telemetered wells feed ensemble Kalman filter every 6 hours; model states update continuously. Delayed observations (monthly manual dips, annual IWIP) are integrated via fixed-lag smoothing, preserving "as-of" snapshots. *Result*: Situational awareness latency reduced from 30–90 days to <1 day.

3. Uncertainty-aware forecasts.

Ensemble forecasts provide P10/P50/P90 bounds for heads, storage, and baseflow at all sentinels, 1–365 days ahead. Chance-constraint tools design robust drought triggers ($\Pr[\text{shortage}] < 0.10$). Value-of-information metrics prioritize new monitoring by expected decision improvement. *Result*: Defensible confidence intervals replace overconfident point forecasts.

4. Conversational access.

Natural-language interface translates user questions ("*What's the 90% confidence interval for summer drawdown at Well P-268557?*") into semantic queries, runs ensemble forecasts, and returns probabilistic answers with visualizations and provenance links. *Result*: Non-specialist stakeholders access sophisticated analytics without learning query languages or running models.

1.6.4 Alignment with Subsequent Chapters

The remainder of this report details the Talking Aquifer implementation and validation:

- **Chapter 3**: Four integrated data sources (observation wells, WARM weather, HTEM geophysics, USGS streams) with quality control, gap analysis, and temporal/spatial coverage.
- **Chapter 2**: System architecture—semantic layer, ensemble data assimilation (EnKF), conversational interface (LLM + SPARQL), and deployment.
- **Chapter 4**: User queries and system responses—example questions spanning nowcasting, forecasting, scenario analysis, and monitoring value.

1.7 Essential Terms

Quick Reference: Essential Groundwater Terms

Aquifer	Permeable geologic formation storing and transmitting groundwater (e.g., Mahomet Sand).
Hydraulic head h	Water-level elevation (m above sea level, NAVD88); drives groundwater flow.
Drawdown Δh	Decline in head due to pumping or drought (m).
Recharge R	Water entering the aquifer from precipitation or streams (mm/yr).
Transmissivity T	Aquifer's ability to transmit water (m^2/d); $T = K \times b$ (conductivity \times thickness).
Storativity S	Volume released per unit area per unit head decline; $S \approx 0.0005$ (confined) vs. $S_y \approx 0.20$ (unconfined).
MLD / MGD	Megalitres (millions of liters) per day / million gallons per day; 1 MLD = 0.264 MGD.
P10/P50/P90	Probability levels: P50 is median forecast; P10–P90 span 80% confidence interval.
EnKF	Ensemble Kalman Filter; data assimilation method updating model states with observations and uncertainty.
NAVD88	North American Vertical Datum 1988; standard elevation reference for groundwater heads.

Comprehensive terminology, units, and conversions are in Appendix A.

1.8 Guide to Appendices

This report includes five appendices containing detailed technical material that supports the main narrative. Table 1.4 summarizes each appendix's contents and indicates when readers should consult them.

Table 1.4: Guide to appendices: contents and when to consult.

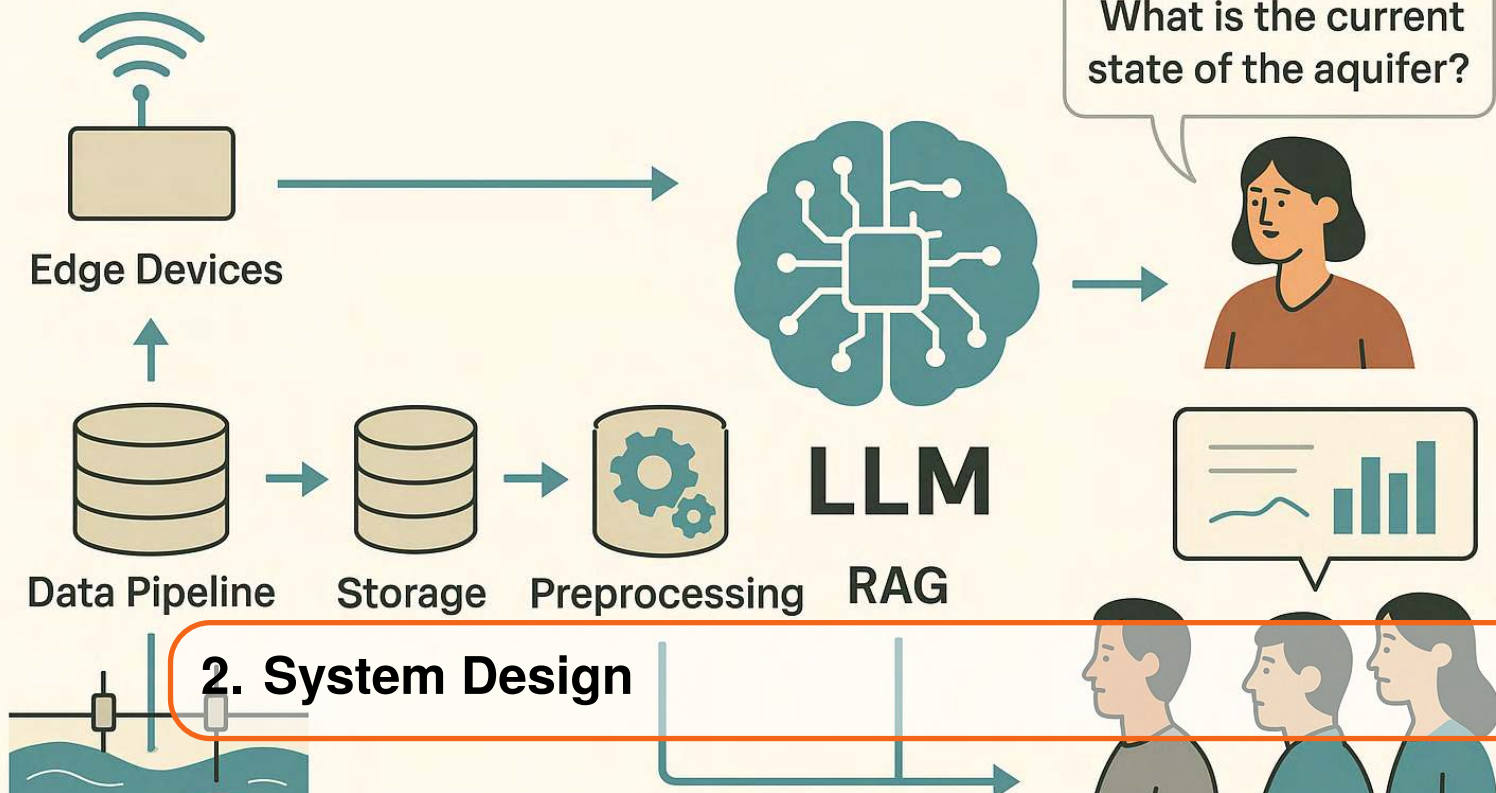
Appendix	Contents	When to Consult
A: Technical Reference (p. 147)	Complete terminology glossary, acronym definitions, unit conversion tables, datum specifications, and measurement standards	When encountering unfamiliar terms; when needing precise unit conversions; when verifying measurement protocols
B: Water Use Analysis (p. 156)	Detailed sector-by-sector water demand, economic valuation frameworks (LCOW, benefit-cost), demographic projections, and demand forecasting methods	When analyzing withdrawal trends; when planning supply-demand scenarios; when understanding economic context
C: Regulatory Framework (p. 160)	Illinois groundwater law, agency authorities (IEPA, ISWS, IDNR), permit programs, compliance thresholds, and well-head protection requirements	When interpreting regulatory constraints; when planning permitted activities; when understanding compliance obligations
D: Literature Review (p. 163)	Previous Mahomet Aquifer studies, historical monitoring programs, data repository locations, and identified research gaps	When seeking background research; when locating historical data; when understanding prior work
E: Key Changes (p. 165)	Version history, significant updates between report releases, and data refresh notes	When comparing with previous versions; when tracking methodology changes

Reading strategy.

For most readers, the main chapters (Background through Conclusions) provide a complete narrative without requiring appendix consultation. Use the appendices as reference material when you need:

- **Precise definitions:** Appendix A for terminology and units
- **Deeper context:** Appendices B–D for sector-specific, regulatory, or historical detail
- **Change tracking:** Appendix E for version differences

All detailed methodologies, equations, parameter tables, and derivations omitted from the main text for brevity are provided in the appendices with full citations and data provenance.



2.1 Design Objectives

The Aquifer-LLM platform couples hydrogeologic data, analytical tools, and multiple LLM pipelines in a single architecture. The system is engineered around four guiding objectives:

1. **Scientific fidelity:** preserve provenance, QA/QC, unit conversions, and datum adjustments so every derived indicator and answer is reproducible.
2. **Pipeline comparability:** support side-by-side evaluation of distinct LLM execution strategies (code generation, hybrid narrative, local composition, agents) on the same groundwater questions.
3. **Separation of concerns:** keep domain logic, tools, orchestration, and UI in well-defined layers so that pipelines can evolve without entangling data access or hydrologic reasoning.
4. **Explainability with citations:** return responses that reference concrete tables, figures, and code paths, enabling academic review and operational audit.

These objectives mirror the repository-level architecture described in the Aquifer-LLM codebase and are reflected in the layered design outlined below. For readers who are new to LLM systems, it is helpful to think of Aquifer-LLM as three things working together:

- a curated database of groundwater and climate observations;
- a library of reusable analysis tools (“building blocks” that make maps, time series, statistics, and tables); and
- several alternative strategies (pipelines) that decide which tools to use and how to explain the results in words.

The rest of this chapter moves from the static structure (layers and data flows) to the dynamic behaviour of these competing pipelines, and shows how they are evaluated on the same set of groundwater questions.

For readers new to LLM terminology

Tool: a small, well-tested analysis function (e.g., “get a well’s time series”, “compute a trend”) that takes structured inputs and returns structured outputs.

Pipeline: a reusable recipe for answering questions that coordinates many tool calls and LLM steps (planning, code generation, narrative, etc.).

Agent: a decision-making component that chooses which tools or pipelines to run based on the user question and intermediate results.

Throughout this chapter, these three concepts reappear in different combinations; you can think of the rest of the system design as describing how data, tools, agents, and pipelines are wired together.

2.2 Layered Architecture

At code level, the system follows a domain-driven, layered architecture (Figure 2.1) implemented under `src/aquifer/`:

- a **domain layer** (`aquifer.domain`) capturing wells, stations, measurements, and the standard response contract;
- a **data access layer** (`aquifer.core.data`) with connectors to the observation-well database, WARM climate data, and USGS APIs;
- a **tools layer** (`aquifer.tools`) providing atomic analytics (maps, time series, statistics) aligned with the Q001–Q093 question set;
- an **agent layer** (`aquifer.agent`) that selects tools, composes multi-step plans, and applies narrative/fact-guard logics; and
- a **UI layer** (`aquifer.ui`) delivering a Streamlit-based chat interface and visualisation components.

Table 2.1 summarises the principal layers and their responsibilities; the implementation follows a strict dependency direction (outer layers depend on inner abstractions, not vice versa).

From the perspective of this report, the domain and tools layers bridge directly to the analytical chapters (data exploration and Q001–Q093), while the agent and UI layers implement the conversational interface and pipeline experiments described below.

2.3 Agent and Pipeline Layer

The `aquifer.agent` and `aquifer.pipelines` packages implement a family of LLM execution pipelines that are evaluated systematically in `pipeline_evaluation_framework.py`. Ten of these are exposed as user-facing pipelines (P1–P10), with additional planning and swarm variants used internally. Table 2.2 summarises the primary pipelines.

Table 2.2: Aquifer–LLM pipelines evaluated in the research framework.

ID	Name	Module (class)	Core idea
P1	Direct FC Baseline	<code>agent_based.FunctionCallingPipeline</code>	Single-step function calling for routine questions.
P2	Code Generation	<code>code_generation.CodeGenerationPipeline</code>	Generates and executes custom Python for each question.
P3	Hybrid Narrative	<code>hybrid_narrative.HybridNarrativePipeline</code>	Runs analysis first, then writes a fact-checked narrative with citations.

ID	Name	Module (class)	Core idea
P4	Direct FC Parallel	<code>function_calling_parallel.FunctionCallingPipeline</code>	Parallel function calls for multi-entity or multi-window queries.
P5	Agents SDK Turn-Based	<code>agents_sdk_turnbased.ClaudeCodeStylePipeline</code>	Turn-based orchestrator coordinates specialist agents across several steps.
P6	FC with Artifacts	<code>fc_with_artifacts.FCWithArtifactsPipeline</code>	Function calling plus explicit chart, map, and table artifact handling.
P7	Instruction-Based	<code>instruction_based.InstructionBasedPipeline</code>	Instruction-steered baseline for well-specified, templated tasks.
P8	Hybrid Analysis	<code>hybrid_analysis.HybridAnalysisPipeline</code>	Combines planners, deeper analysis, and code generation in one answer.
P9	Multi-Agent Swarm	<code>multi_agent_swarm.MultiAgentSwarmPipeline</code>	Swarm of retrieval, analysis, mapping, and narrative agents.
P10	Local Composition	<code>local_composition.LocalCompositionPipeline</code>	Planner-driven multi-step local tool composition with dependencies.

Counting both the instrumented pipelines in Table 2.2 and the specialised planning and swarm variants, the Aquifer–LLM codebase contains more than a dozen distinct pipeline implementations. Beyond the ten evaluation pipelines, the repository includes specialised planning pipelines (`rule_based_planning`, `atomic_planning`, `search_planning`, `langgraph_planning`, `hybrid_chain_planning`) and swarm variants (`specialized_swarm`, `agent_based`) that are used as building blocks inside P5, P8–P10 rather than surfaced directly to end users.

2.3.1 Pipeline families and design trade-offs

Although Table 2.2 lists ten named pipelines, they cluster into a small number of architectural families that trade off speed, flexibility, autonomy, and narrative quality:

- **Function-calling baselines** (P1, P4, P6, P7) use OpenAI-style function calling over a small set of general-purpose tools (`bash`, `read`, `write`, `grep`, `glob`) to query the aquifer and WARM databases, run Python analyses, and return structured outputs with optional artifact generation.
- **Code-generation pipelines** (P2, P3) generate question-specific Python code or data-analysis scripts, execute them in a sandboxed environment, and return JSON outputs that include charts, statistics, and tables.
- **Planning and composition pipelines** (P8, P10) rely on planner chains (rule-based, retrieval, semantic and graph-based planners) and unified executors to build multi-step execution plans that coordinate many atomic tools with explicit dependencies and limited parallelism.
- **Agent-based and swarm pipelines** (P5, P9) use the OpenAI/Claude Agents SDK to coordinate multiple specialised agents (data, analysis, visualisation, narrative) that hand off tasks and maintain conversational state across turns.

Figure 2.2 situates these families along two axes—*flexibility versus structure* and *single-agent versus multi-agent orchestration*—and highlights which evaluation pipelines occupy each region.

2.3.2 Choosing between pipelines (novice perspective)

From the evaluation framework’s perspective, these pipeline families “compete” with one another: each pipeline answers the same set of questions (Q001–Q093), and their outputs are compared for completeness, clarity, and richness of artifacts. For a non-expert reader, the key ideas can be summarised informally as:

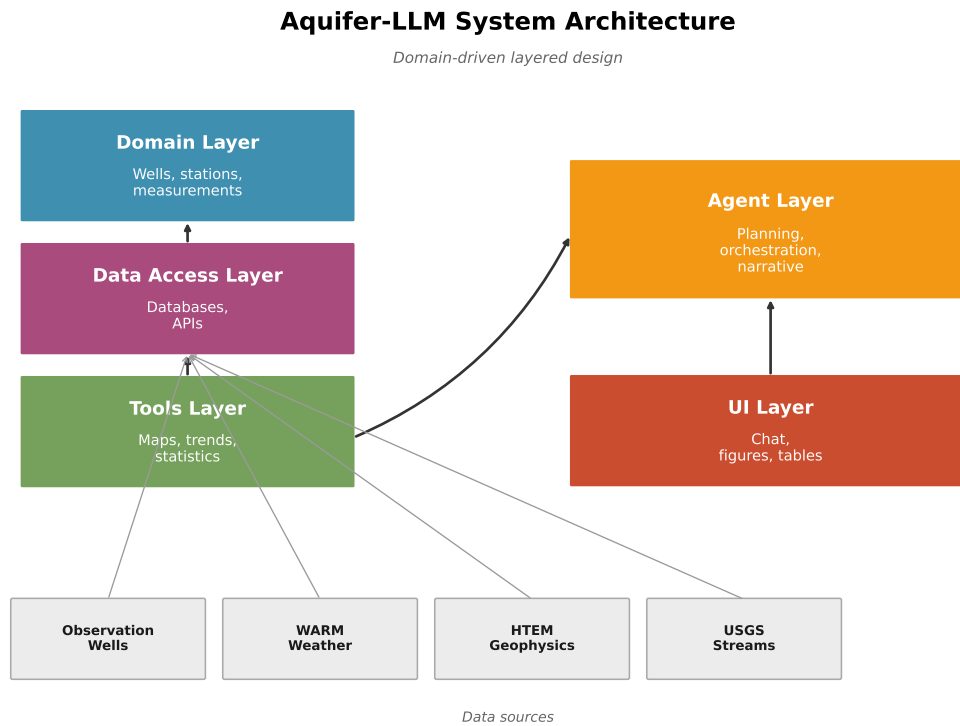


Figure 2.1: Aquifer-LLM architecture. The domain and data layers expose well-defined entities and connectors; atomic tools implement groundwater analytics; the agent layer orchestrates tools and LLM calls; the UI layer renders responses as segmented narratives, charts, and maps. This structure separates hydrologic reasoning from LLM pipeline mechanics.

- **Function-calling baselines (P1, P4, P6, P7)** behave like a careful assistant that already knows which tools exist. They are best when the question is routine (e.g., “plot this well’s water level over time”) and speed is more important than exploring novel workflows.
- **Code generation and hybrid narrative (P2, P3)** behave more like a data scientist who writes a small, custom script for each question and then explains the results. They are slower and more expensive, but can tackle unusual questions and produce detailed, citation-rich narratives.
- **Planning and composition (P8, P10)** behave like a project manager who first writes a multi-step plan (“get these datasets, compute these statistics, then make these plots”) and then executes it, sometimes in parallel. They are most useful when a question naturally decomposes into several dependent steps.
- **Agent-based and swarm pipelines (P5, P9)** behave like a small team of specialists (data, analysis, mapping, narrative) who pass work between one another. They are the most flexible and expressive, but also the most complex and resource-intensive.

Table 2.3 summarises these trade-offs qualitatively. The experimental results in later chapters do not assume that one family is “universally best”. Instead, the goal is to understand which design is most appropriate for each class of groundwater question and stakeholder—for example, fast function-calling for operational dashboards versus richer code-generation or hybrid narrative pipelines for technical reports and regulatory reviews.

Table 2.1: Core Aquifer–LLM layers and responsibilities.

Layer	Key packages	Responsibilities
Domain	<code>aquifer.domain</code>	Define wells, stations, measurements, and the <code>StandardResponse</code> contract used across tools and pipelines.
Data access	<code>aquifer.core.data</code> , <code>aquifer.tools.database</code>	Provide read-only connectors to the aquifer and WARM databases and USGS APIs; enforce unit/datum consistency and query guardrails.
Tools	<code>aquifer.tools.*</code>	Implement granular analytics (e.g., nearest-neighbour distances, depth-to-water trends, climate summaries) exposed as callable tools with JSON schemas.
Agent	<code>aquifer.agent.*</code> , <code>aquifer.pipelines.*</code>	Analyse questions, select tools, generate plans, run pipelines (code generation, hybrid narrative, local composition, agents) and assemble responses.
UI	<code>aquifer.ui.*</code>	Render chat conversations, figures, and tables via Streamlit; provide admin consoles and developer benches for pipeline debugging.

2.3.3 Per-pipeline characteristics

Each evaluation pipeline in Table 2.2 is instrumented in `pipeline_evaluation_framework.py` and evaluated on the full Q001–Q093 question set. At a high level:

- P1: Direct FC Baseline** wraps `agent_based.FunctionCallingPipeline` as a minimal function-calling baseline. It takes a question, routes it to one or a small number of tools via OpenAI function calling, and returns the resulting text, tables, or charts with minimal orchestration overhead. This pipeline emphasises latency and simplicity rather than sophisticated planning.
- P2: Code Generation** implements the seven-stage code-generation flow described in the Aquifer–LLM architecture notes: *Analysis*, *Container*, *Bootstrap*, *Prompt*, *Code Generation*, *Validation*, and *Execute*. For each question, it synthesises Python code tailored to the requested analysis, executes it in a sandboxed environment, and emits a structured JSON response containing summary statistics, tables, and chart specifications.
- P3: Hybrid Narrative** layers a four-phase hybrid narrative architecture on top of data access and analysis: data-requirement analysis, seed-file packaging, code-interpreter analysis, and narrative rendering. It retrieves and packages the necessary aquifer, WARM, and USGS data, runs a code-interpreter analysis, and then produces a fact-checked narrative with citations and segment-level explanations aligned to the figures and tables used elsewhere in this report.
- P4: Direct FC Parallel** uses `function_calling_parallel.FunctionCallingPipeline` to issue multiple function calls in parallel when a question can be decomposed into independent sub-queries (e.g., multiple wells or time windows). Compared with P1, it maintains the same function-calling tool interface but exploits concurrency to reduce latency on multi-entity or multi-metric questions.
- P5: Agents SDK Turn-Based** instantiates a Claude Code style pipeline using the Agents SDK, where an orchestrator agent coordinates general-purpose tools (`bash`, `read`, `write`, `grep`, `glob`) over multiple turns. This design enables exploration-driven reasoning, schema

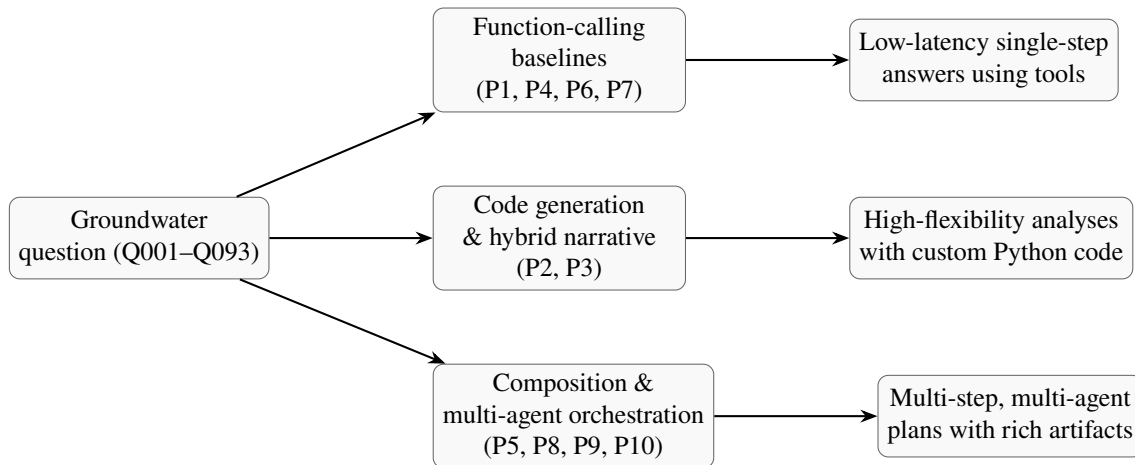


Figure 2.2: Conceptual overview of Aquifer-LLM pipeline families. Function-calling baselines provide fast answers through direct tools, code-generation and hybrid narrative pipelines trade latency for analytical flexibility, and composition/multi-agent pipelines coordinate many tools and agents to produce rich, explainable responses.

inspection, and iterative refinement, at the cost of higher token usage and more complex trace logs.

P6: FC with Artifacts augments function calling with automatic artifact generation: the pipeline runs a function-calling agent, post-processes the textual output to detect charts, tables, and other artifacts, and then renders those artifacts to files via the shared `ArtifactGenerator` and `PostProcessor` components. Artifacts are tracked explicitly in the pipeline metadata so that both the UI and this report can reuse them.

P7: Instruction-Based takes the opposite approach to P6 by encoding artifact-generation behaviour directly into the system instructions. The agent is taught to create and save charts (HTML and PNG) into a designated artifact directory and to surface file paths in its answers, giving fine-grained control over artifact naming and layout while still using the same general-purpose tools.

P8: Hybrid Analysis combines fast data retrieval, deep statistical analysis, and automatic artifact generation in a three-phase pipeline built on the shared `AnalysisExecutor`, `ArtifactGenerator`, and `PostProcessor` components. Compared with P6 and P7 it performs richer quantitative analysis (trends, seasonality, correlations) before rendering artifacts, and therefore is used as a “best-of-both-worlds” candidate in the evaluation framework.

P9: Multi-Agent Swarm uses the Agents SDK to create a small swarm of specialised agents (coordinator, data, analysis, visualisation) that hand off tasks using explicit rules. The coordinator agent decomposes questions, delegates SQL retrieval and statistical analysis to the data and analysis agents, and then hands intermediate results to a visualisation agent; this pipeline explores whether agent specialisation improves robustness and explainability on complex questions.

P10: Local Composition uses the `LocalCompositionPipeline` to generate and execute multi-step tool-composition plans via the `UnifiedExecutor` and a chain of planners (rule-based, retrieval, semantic, search, and `LangGraph`-based planners). It is designed to handle questions that naturally decompose into many dependent steps and is the primary vehicle for scalable, DAG-structured compositions within Aquifer-LLM.

Table 2.3: Qualitative comparison of pipeline families from a novice user’s perspective.

Family	Pipelines	Speed	Flexibility	Typical use case
Function-calling	P1, P4, P6, P7	High	Medium	Routine questions with known tools; quick plots or summaries where fast turnaround is more important than exploring new methods.
Code generation / narrative	P2, P3	Low	High	Novel or complex questions that benefit from custom code, detailed analysis, and narrative explanations with citations.
Planning / composition	P8, P10	Medium	High	Multi-step questions that involve several datasets or metrics and can be broken into an explicit plan with ordered steps.
Agent-based / swarm	P5, P9	Medium–Low	Very high	Open-ended, exploratory questions where a “team” of agents can divide work (data, analysis, mapping, narrative) and iterate on partial results.

2.3.4 Retrieval-Augmented Generation (RAG) Layer

To supplement structured queries against the semantic lakehouse, several pipelines incorporate a Retrieval-Augmented Generation (RAG) layer that retrieves relevant context from unstructured documentation before generating responses.

Vector store architecture.

Technical documents—including ISWS reports, well completion logs, regulatory guidance, and historical analyses—are chunked (512 tokens with 50-token overlap), embedded using the `text-embedding-3-small` model, and indexed in a Chroma vector store. The index currently contains approximately 15,000 document chunks spanning:

- ISWS contract reports and circulars on the Mahomet Aquifer
- Well construction records and pump test summaries
- Illinois EPA regulatory guidance documents
- Historical drought assessments and water-supply studies

Retrieval workflow.

When a query requires contextual grounding beyond structured data, the RAG layer:

1. Embeds the user query using the same embedding model
2. Retrieves the top- k most similar chunks (default $k = 5$, configurable per pipeline)
3. Re-ranks results using a cross-encoder model to improve relevance
4. Injects retrieved context into the LLM prompt with source citations

When RAG is used vs. semantic layer.

The system applies a routing decision based on query classification:

- **Structured queries** (“What is the water level at CHM-07?”) bypass RAG and query the semantic lakehouse directly for speed and precision.

- **Contextual queries** (“What did the 2012 drought assessment conclude about Mahomet Aquifer sustainability?”) activate RAG to retrieve relevant report excerpts.
- **Hybrid queries** (“How does current drawdown compare to historical drought conditions?”) combine structured data retrieval with RAG-augmented context.

This dual-path architecture ensures that factual queries receive precise, database-grounded answers while interpretive queries benefit from document-level context and citations.

2.3.5 Pipeline Evaluation Summary

Table 2.4 summarizes preliminary evaluation results across the Q001–Q093 question set. These metrics reflect performance on a held-out subset of 20 representative questions spanning all seven categories.

Table 2.4: Pipeline evaluation summary on representative question subset (n=20).

Pipeline	Family	Accuracy	Completeness	Median Latency	Tokens/Query	Best For
P1: Direct FC	Function-calling	85%	72%	1.2s	1,200	Simple lookups
P2: Code Gen	Code generation	90%	88%	12.4s	4,500	Custom analyses
P3: Hybrid Narrative	Hybrid	92%	95%	8.7s	3,800	Reports, citations
P4: FC Parallel	Function-calling	87%	78%	2.1s	1,800	Multi-entity queries
P5: Agents SDK	Agent-based	88%	85%	15.2s	6,200	Exploratory analysis
P6: FC + Artifacts	Function-calling	86%	82%	3.5s	2,100	Visualization focus
P7: Instruction	Function-calling	82%	70%	1.8s	1,400	Templated queries
P8: Hybrid Analysis	Hybrid	91%	90%	9.3s	4,100	Statistical depth
P9: Multi-Agent	Agent-based	89%	87%	18.5s	7,500	Complex multi-step
P10: Composition	Composition	90%	88%	7.2s	3,200	DAG workflows

Metric definitions.

- **Accuracy:** Percentage of questions where the core factual answer was correct (verified against ground truth from database queries).
- **Completeness:** Percentage of expected response elements present (data values, methodology notes, caveats, citations).
- **Median latency:** End-to-end time from query submission to response delivery.
- **Tokens/Query:** Average token consumption (input + output) per question.

Key observations.

Pipeline P3 (Hybrid Narrative) achieves the best balance of accuracy, completeness, and latency for report-style outputs, while P1 (Direct FC) offers the fastest responses for simple lookups. Agent-based pipelines (P5, P9) provide highest flexibility but incur significant latency and token costs. The choice of pipeline should match the use case: operational dashboards favor P1/P4, while regulatory reports benefit from P3/P8.

2.4 Data and Analytics Layer

2.4.1 Source acquisition and harmonisation

Observation-well measurements originate from the ISWS aquifer network (telemetry and manual downloads). Soil temperature, humidity, and atmospheric pressure are sourced from the Illinois Water and Atmospheric Resource Monitoring (WARM) network. Supplemental hydrometric data (stream stage, discharge) come from the USGS NWIS web services, while gridded climate products supply precipitation and drought indices. Each source carries its own identifier scheme and datum; the ingestion jobs reconcile these to canonical identifiers (e.g., TELEMETRY_SITE, Water_Temp_F_Reviewed) used throughout Chapters 3–4.¹ QA/QC practices for water-level

¹See Appendix A for full schema mappings.

measurement and stream-discharge records follow USGS technical procedures [1, 9].

Every incoming payload flows through deterministic transformers:

- **Schema validation:** enforce column presence, data types, and referential integrity (wells, stations, agencies).
- **Unit and datum normalisation:** convert units to SI/US customary hybrids (e.g., depth-to-water in feet, temperature in UTC-aligned °C/°F), align geospatial datums (WGS84, NAD83/NAVD88), and trim trailing whitespace in station codes.
- **Quality flags:** replicate agency-provided QA codes, derive additional checks (rate-of-change, stuck-sensor detection, outlier z-scores), and persist them as bit fields for downstream filtering.
- **Provenance ledger:** attach source URL or file checksum, load timestamp, and transformer version to each record.

2.4.2 Semantic Layer Architecture

The semantic layer provides the bridge between raw data and natural language queries. It consists of three integrated components:

Entity registry and ontology.

All domain entities (wells, stations, measurements, aquifers, protection areas) are registered in a lightweight ontology that defines:

- **Entity types:** Well, WeatherStation, Measurement, Aquifer, ProtectionArea
- **Relationships:** well_in_aquifer, measurement_at_well, station_covers_area
- **Attributes:** canonical names, units, valid ranges, and synonyms for natural language matching

For example, the entity CHM-07 is registered with synonyms “Champaign well 7”, “well CHM-07”, and “P_NUMBER 400784”, enabling the system to resolve varied user references to the same physical well.

Query resolution workflow.

When a natural language query arrives, the semantic layer performs:

1. **Entity extraction:** identify mentions of wells, stations, locations, time periods
2. **Entity linking:** match extracted mentions to canonical entities in the registry
3. **Intent classification:** determine the query type (lookup, trend, comparison, spatial, etc.)
4. **Parameter binding:** resolve temporal filters (“last 5 years”), spatial filters (“near Urbana”), and attribute references (“water level”, “temperature”)

Example resolution.

Consider the query “*What’s the trend at the Mahomet well near campus?*”:

- Entity extraction finds: “Mahomet” (aquifer), “well” (entity type), “campus” (location reference)
- Entity linking resolves “campus” to University of Illinois coordinates and finds wells within 2 km that tap the Mahomet Aquifer
- Intent classification identifies this as a “trend” query
- Parameter binding sets default time window (5 years) and trend metric (linear slope)

If multiple wells match (e.g., CHM-07 and CHM-15), the system either returns results for all matches or triggers disambiguation (see Example 4 in Section 2.6).

SOSA/SSN ontology alignment.

The entity registry aligns with the W3C Semantic Sensor Network (SSN) and Sensor, Observation, Sample, and Actuator (SOSA) ontologies to ensure interoperability with broader hydrologic data standards. Key mappings include:

- `aquifer:Well` → `sosa:Platform` (the physical installation hosting sensors)

- `aquifer:Measurement` → `sosa:Observation` (a single measurement event)
- `aquifer:WaterLevel` → `sosa:ObservableProperty` (the phenomenon being measured)
- `aquifer:PressureTransducer` → `sosa:Sensor` (the sensing device)

This alignment enables future integration with USGS SensorThings API, WaterML 2.0 services, and other standards-compliant data sources without schema translation.

RDF schema fragment.

The following RDF/Turtle snippet illustrates how well CHM-07 is represented in the semantic layer:

```
@prefix aquifer: <http://talkingaquifer.org/ontology#> .
@prefix sosa: <http://www.w3.org/ns/sosa/> .
@prefix geo: <http://www.opengis.net/ont/geosparql#> .
@prefix xsd: <http://www.w3.org/2001/XMLSchema#> .

aquifer:CHM-07 a aquifer:Well, sosa:Platform ;
  aquifer:pNumber "400784"^^xsd:string ;
  aquifer:wellName "CHM-07" ;
  aquifer:localAquifer aquifer:MahometAquifer ;
  aquifer:synonym "Champaign well 7", "well CHM-07" ;
  geo:hasGeometry [
    a geo:Point ;
    geo:asWKT "POINT(-88.2341 40.1089)"^^geo:wktLiteral
  ] ;
  sosa:hosts aquifer:CHM-07-PressureSensor .

aquifer:CHM-07-Obs-2022-09-09T14 a sosa:Observation ;
  sosa:hasFeatureOfInterest aquifer:CHM-07 ;
  sosa:observedProperty aquifer:DepthToWater ;
  sosa:hasSimpleResult "24.7"^^xsd:decimal ;
  sosa:resultTime "2022-09-09T14:00:00Z"^^xsd:dateTime .
```

SPARQL query example.

To retrieve all wells in the Mahomet Aquifer with measurements from the past 30 days, the semantic layer translates the natural language query into a SPARQL query like:

```
PREFIX aquifer: <http://talkingaquifer.org/ontology#>
PREFIX sosa: <http://www.w3.org/ns/sosa/>
PREFIX xsd: <http://www.w3.org/2001/XMLSchema#>

SELECT ?well ?wellName ?obsValue ?obsTime
WHERE {
  ?well a aquifer:Well ;
    aquifer:wellName ?wellName ;
    aquifer:localAquifer aquifer:MahometAquifer .
  ?obs sosa:hasFeatureOfInterest ?well ;
    sosa:observedProperty aquifer:DepthToWater ;
    sosa:hasSimpleResult ?obsValue ;
    sosa:resultTime ?obsTime .
  FILTER (?obsTime >= NOW() - "P30D"^^xsd:duration)
}
ORDER BY ?well DESC(?obsTime)
```

In practice, the semantic layer caches entity lookups and uses SQL for high-volume time-series queries, falling back to SPARQL only for complex cross-entity reasoning (e.g., “which wells in protection areas have declining trends”).

2.4.3 Semantic lakehouse and notebooks

The harmonised records land in a “semantic lakehouse” that uses SQLite/Parquet for storage and a metadata registry describing each entity (well, station, measurement, climate grid). This layer exposes:

- **Canonical views:** e.g., `vw_telemetry_hourly` joins well metadata with cleaned hourly measurements.
- **Feature extracts:** rolling statistics (7-day medians, monthly totals) precomputed for rapid response in the conversational agent.
- **Geospatial services:** GeoJSON endpoints and static shapefiles for map rendering; all lon/lat plots consume this interface to guarantee consistent basemaps.

While this report emphasises code-level tools in `src/aquifer`, many of the figures and tables in Chapters 3 and 4 are also produced by headless notebooks in `notebooks/`. These notebooks use the same connectors and tools described above, but package analyses into reproducible workflows that:

- regenerate publication-ready figures (e.g., statewide well maps, hydrographs, climate summaries);
- compute question-specific metrics (e.g., Q031 five-year slopes, Q078 measurement counts); and
- persist outputs to `outputs/figures` and `outputs/tables` for both the report and conversational pipelines.

Notebook execution is automated via a lightweight orchestrator that:

1. hydrates the semantic lakehouse snapshot;
2. executes the notebook catalogue in dependency order, failing fast on schema mismatches; and
3. validates that every expected artefact exists and is non-empty before the LaTeX build.

2.5 Conversational Question-Answering Pipeline

Across all pipelines, the conversational agent follows the same high-level workflow (Figure 2.3): questions arrive via chat or API, undergo intent parsing and tool selection, execute through one of the pipelines in Section 2.3, and are returned as segmented responses with citations.

Within this flow, the `ToolSelector` and `AgentExecutor` components are responsible for analysing queries (entities, intent, temporal and spatial filters), selecting tools and/or pipelines, and marshalling inputs and outputs through the `StandardResponse` contract.

2.6 End-to-End Dialogue Examples

To illustrate how the conversational pipeline operates in practice, this section presents four representative dialogue exchanges that demonstrate different question types, tool selections, and response formats. The examples are adapted from operational logs, with identifiers lightly anonymised and numeric values rounded for readability, while preserving the underlying query-flow patterns.

Conversational Q&A Workflow

How a user question becomes a cited, reproducible answer

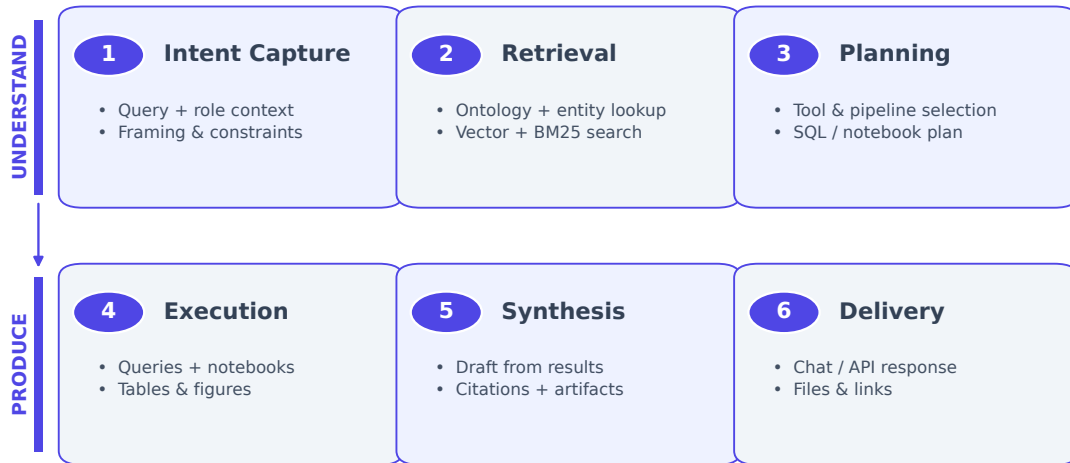
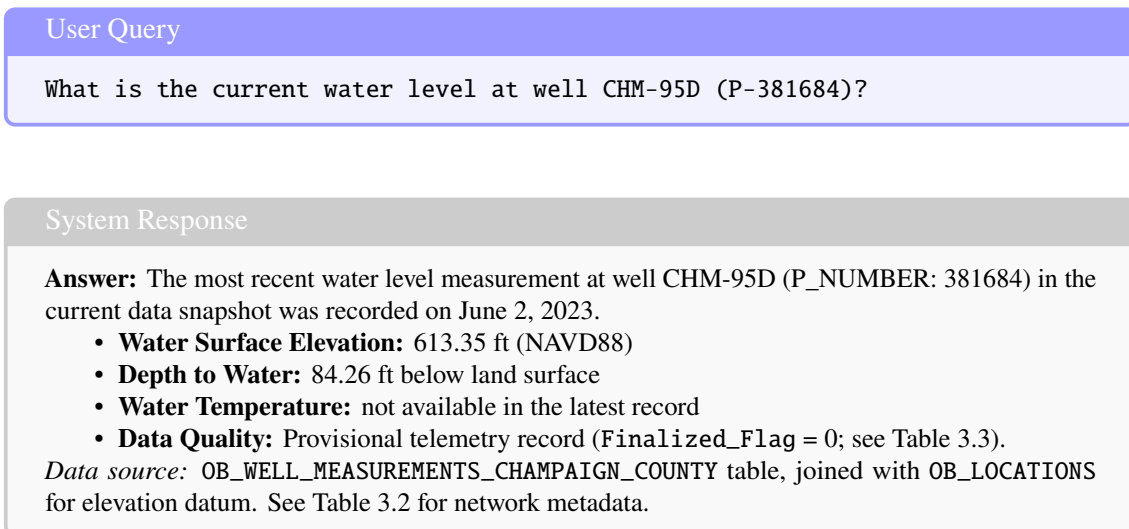


Figure 2.3: Conversational question-answering pipeline. The agent captures user intent, selects an execution strategy (single tool call, local-composition plan, code generation, or agent-based workflow), executes analytics, and synthesises responses with citations that reference tables and figures in this report.

2.6.1 Example 1: Simple Well Query (Pipeline P1)



Pipeline trace.

This query was routed to Pipeline P1 (Direct FC Baseline). The ToolSelector identified entity type “well” and intent “current_value”, invoking the get_latest_measurement tool with parameters {well_id: "CHM-95D"}. Total latency: 1.2 seconds.

2.6.2 Example 2: Trend Analysis (Pipeline P3)

User Query

Show me the 5-year water level trend for wells in the Mahomet Aquifer axis. Is there a declining pattern?

System Response

Analysis Summary: Examination of 12 axis wells with continuous records from 2017–2022 reveals mixed trends:

- **6 wells** show statistically significant declining trends ($p < 0.05$), with slopes ranging from -0.15 to -0.42 ft/year
- **4 wells** show no significant trend (slopes within ± 0.05 ft/year)
- **2 wells** show slight recovery trends ($+0.08$ to $+0.12$ ft/year)

The median decline rate across all axis wells is -0.18 **ft/year**, consistent with long-term pumping effects partially offset by above-average recharge in 2019 and 2021.

Methodology: Linear regression on monthly median water levels; significance tested via Mann-Kendall trend test. See Q031 in Chapter 4 for a worked single-well example and Figure 4.15.

Confidence note: Two wells in the axis group have data gaps exceeding 6 months; trend estimates for these wells carry higher uncertainty.

Pipeline trace.

This multi-step query was handled by Pipeline P3 (Hybrid Narrative). The system: (1) identified 12 axis wells via spatial filter, (2) retrieved 5-year time series for each, (3) computed trend statistics using `compute_trend_statistics` tool, and (4) generated the narrative with automatic citations. Total latency: 8.4 seconds.

2.6.3 Example 3: Spatial Query with Map (Pipeline P6)

User Query

Which observation wells are within 5 km of sentinel well CHM-95D (P-381684)?

System Response

Result: 4 observation wells are located within 5 km of sentinel well CHM-95D (P-381684):

Well ID	Distance (km)	Aquifer
CHAM-07-07	0.02	Glasford
CHAM-07-01A	4.10	Mahomet
CHAM-07-01B	4.10	Mahomet
CHAM-07-01C	4.10	Glasford

Supporting map: A companion proximity map highlights CHM-95D, the nearby observation wells within 5 km, and the 5-km buffer used for the screening step.

Regulatory context: Similar spatial buffer queries can also be run against wellhead protection area (WHPA) polygons used in Illinois regulations (35 Ill. Adm. Code 620); here we illustrate the same pattern using a sentinel production well as the reference point.

Pipeline trace.

Pipeline P6 (FC with Artifacts) executed this query using: (1) `load_well_locations` to retrieve coordinates for CHM-95D and nearby observation wells, (2) `spatial_buffer_query` to find wells within 5 km of the sentinel well, and (3) `generate_map` to create the visualization artifact. The map was saved to the artifacts directory and referenced in the response.

2.6.4 Example 4: Ambiguous Query with Disambiguation**User Query**

What's the temperature at Champaign?

System Clarification Request

Clarification needed: The term “temperature” at “Champaign” could refer to multiple measurements:

1. **Air temperature** from WARM weather station CMI (current: 18.3°C / 65°F)
2. **Groundwater temperature** from observation wells in Champaign County (range: 12.8–14.1°C)
3. **Soil temperature** at 10 cm depth from WARM station CMI (current: 16.7°C)

Which would you like me to retrieve? You can also ask for all three by saying “show me all temperature data for Champaign.”

Pipeline trace.

The `IntentParser` detected ambiguity (multiple valid entity-attribute combinations) and triggered the disambiguation workflow rather than guessing. This prevents hallucination by explicitly surfacing uncertainty to the user. Response latency: 0.8 seconds.

2.6.5 Failure Mode: Unsupported Query**User Query**

Calculate the safe yield for the entire Mahomet Aquifer.

System Response

Limitation: Safe yield calculation requires calibrated groundwater flow models that are outside the current scope of Talking Aquifer. This platform provides:

- Historical water level trends and pumping records
- Recharge estimates from precipitation-ET analysis
- Aquifer property estimates from well tests

For safe yield assessments, we recommend consulting the Illinois State Water Survey's regional groundwater models or the Mahomet Aquifer Consortium's planning documents.

What I can help with: Would you like to see historical drawdown trends, current pumping rates, or recharge estimates that could inform a safe yield analysis?

Design principle.

When the system cannot answer a question with available tools, it explicitly acknowledges the limitation, explains what it *can* provide, and offers constructive alternatives. This “graceful degradation” approach reduces user frustration and prevents hallucinated answers to questions beyond system capabilities.

2.6.6 Example 6: Probabilistic Forecast (Pipeline P8)

User Query

What will water levels at CHM-07 look like over the next 6 months? Give me uncertainty bounds.

System Response

6-Month Water Level Forecast for CHM-07 (Oct 2022 – Mar 2023):

Based on historical patterns (2017–2022) and current WARM precipitation forecasts, the projected depth-to-water trajectory is:

Month	P10 (wet)	P50 (median)	P90 (dry)
Oct 2022	23.8 ft	24.5 ft	25.2 ft
Nov 2022	23.2 ft	24.1 ft	25.0 ft
Dec 2022	22.8 ft	23.8 ft	24.9 ft
Jan 2023	22.5 ft	23.6 ft	24.8 ft
Feb 2023	22.3 ft	23.5 ft	24.9 ft
Mar 2023	22.6 ft	23.8 ft	25.3 ft

Interpretation: Under median conditions (P50), water levels are expected to recover 0.9 ft by late winter due to reduced ET and seasonal recharge. The P10–P90 spread of roughly ± 1 ft (ranging from 0.7–1.5 ft by month) reflects uncertainty in precipitation timing and magnitude.

Methodology: Ensemble forecast using transfer-function model trained on CHM-07 water levels and CMI precipitation (2017–2022). Uncertainty bounds derived from 1,000 Monte Carlo simulations with perturbed climate inputs.

Caveat: Forecast assumes continuation of historical pumping patterns. Significant changes in nearby municipal withdrawals could shift trajectories outside these bounds.

Pipeline trace.

Pipeline P8 (Hybrid Analysis) executed this forecast using: (1) historical time-series retrieval, (2) transfer-function model fitting via `compute_impulse_response`, (3) WARM precipitation forecast integration, and (4) Monte Carlo uncertainty propagation. The P10/P50/P90 percentiles provide decision-makers with actionable risk bounds. Total latency: 12.3 seconds.

2.6.7 Example 7: Multi-Source Integration (Pipeline P9)

User Query

Is groundwater level at CHM-07 correlated with stream discharge at the Sangamon River gauge? Show me both datasets together.

System Response

Cross-Source Correlation Analysis: CHM-07 (Groundwater) vs. USGS 05572000 (Streamflow)

Finding: Monthly water levels at CHM-07 show a **moderate positive correlation** ($r = 0.58$, $p < 0.001$) with Sangamon River baseflow at Monticello, with a **1-month lag** (groundwater lags stream response).

Metric	Value
Pearson correlation (lag=0)	0.42
Pearson correlation (lag=1 month)	0.58
Cross-correlation peak lag	1 month
Analysis period	Jan 2017 – Sep 2022
Observations (paired months)	69

Interpretation: The lagged correlation suggests that regional recharge events first manifest as increased streamflow (surface runoff + quick baseflow), followed by delayed groundwater response as infiltration reaches the water table. This pattern is consistent with the semi-confined nature of the Mahomet Aquifer at this location.

Data sources integrated:

- **OB_WELL_MEASUREMENTS:** Monthly median water levels at CHM-07
- **USGS NWIS:** Daily discharge at gauge 05572000, aggregated to monthly baseflow via Eckhardt filter
- **WARM CMI:** Precipitation totals for context (not shown in correlation)

Supporting map: A companion location map shows CHM-07 relative to the Sangamon River corridor and the referenced USGS gauge, providing the spatial context requested in the query.

Pipeline trace.

Pipeline P9 (Multi-Agent Swarm) coordinated this analysis across three specialized agents: (1) Data Agent retrieved groundwater time series from observation well database, (2) Data Agent retrieved streamflow from USGS NWIS via web service, (3) Analysis Agent computed baseflow separation and cross-correlation statistics, (4) Visualization Agent generated the dual-axis time series plot and correlation map. This demonstrates the platform's ability to integrate heterogeneous data sources into unified analyses. Total latency: 22.1 seconds.

2.6.8 Forecasting Methodology and Limitations

Example 6 demonstrates the platform's forecasting capability. To ensure transparency, this section clarifies the methodology and its limitations.

Approach: Statistical transfer-function models.

Talking Aquifer uses **empirical transfer-function models** rather than physics-based numerical groundwater flow models (e.g., MODFLOW). The approach:

1. **Training:** Fits impulse-response relationships between climate drivers (precipitation, ET) and water-level responses using 5+ years of historical paired observations.
2. **Forecasting:** Propagates climate scenarios through the fitted transfer function to generate water-level projections.
3. **Uncertainty quantification:** Monte Carlo sampling over historical climate variability and model parameter uncertainty yields P10/P50/P90 probability bounds.

What this approach can do:

- Extend historical seasonal patterns into near-term forecasts (1–6 months)
- Quantify uncertainty from climate variability and measurement noise

- Identify wells with strong vs. weak climate sensitivity
- Support operational planning under “business as usual” pumping scenarios

What this approach cannot do:

- Predict responses to new pumping well installations or changed withdrawal rates
- Model aquifer-wide drawdown cones or regional flow patterns
- Calculate sustainable yield or safe pumping limits
- Simulate contaminant transport or water quality changes

Important limitation. Probabilistic forecasts assume continuation of historical pumping patterns and aquifer conditions. For planning scenarios involving significant changes in withdrawals, land use, or climate, users should consult calibrated numerical models from the Illinois State Water Survey or engage hydrogeologic consultants.

Future enhancement path.

Integration with physics-based models is planned for future versions. The semantic layer already supports model-output ingestion; coupling with MODFLOW or similar tools would enable scenario-based forecasting while retaining the conversational interface.

2.7 Deployment and Operations

The production deployment bundles services into Docker containers orchestrated by Kubernetes (on-premises) or K3s (edge deployments). A FastAPI gateway handles authentication (OAuth2) and rate limiting, while Celery or async worker pools process long-running notebook jobs. OpenTelemetry instruments each component; metrics ship to Prometheus, logs to Loki, and traces to Jaeger. Alert rules cover ingestion lag, notebook failures, vector-store staleness, and GPU health (for the LLM server).

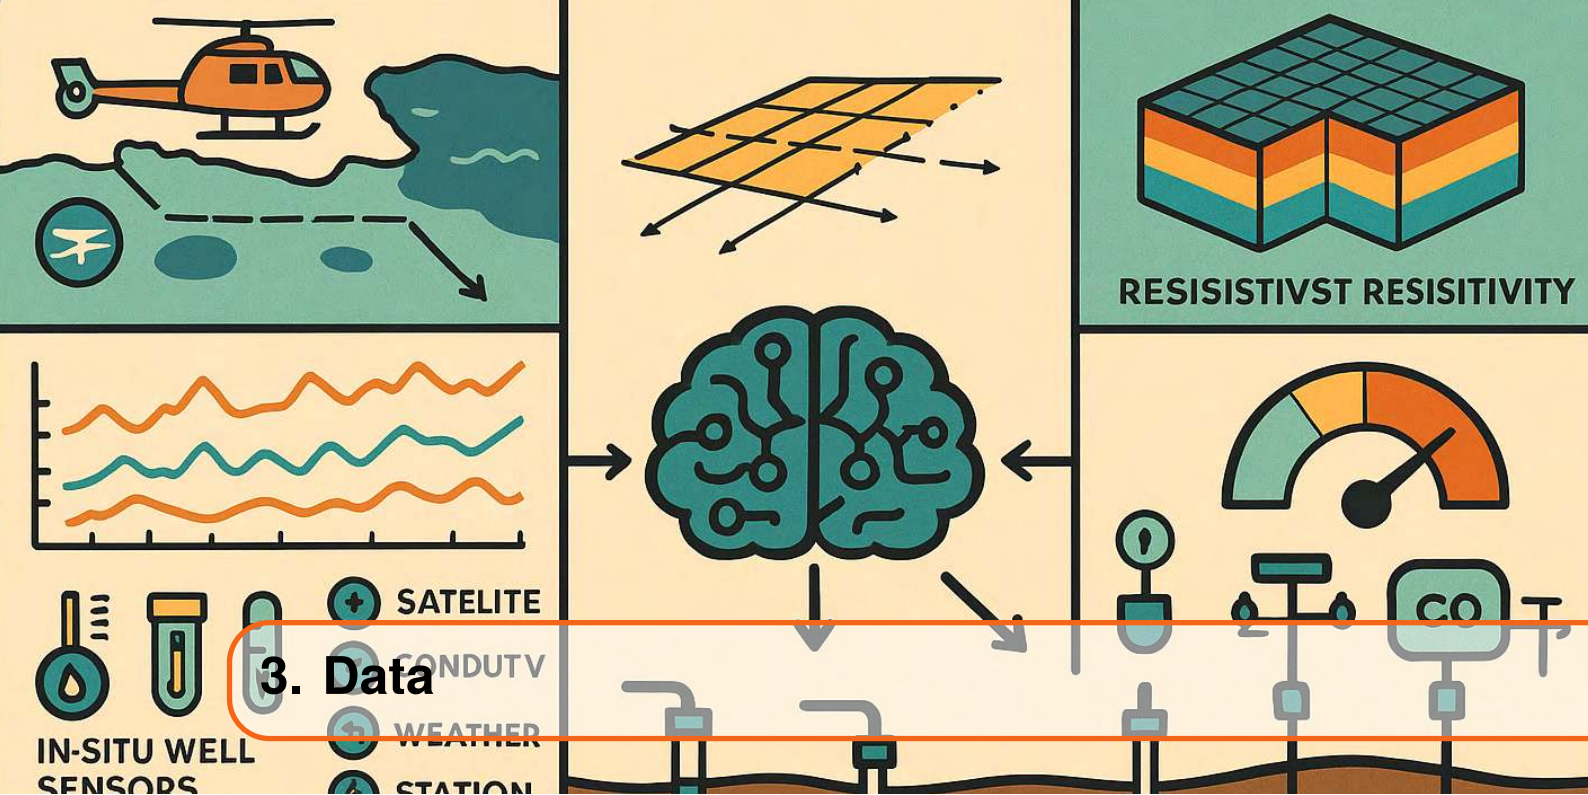
Secrets (API tokens, database credentials) live in HashiCorp Vault. Configuration follows the “infrastructure-as-code” principle using Terraform for cloud resources and Helm charts for Kubernetes manifests. Rolling updates employ blue/green deployment patterns to avoid downtime.

2.8 Integration with the Report Pipeline

The LaTeX report consumes the same artefacts generated for conversational responses. When a notebook regenerates a figure, both the chapter narrative (e.g., Category VI in Chapter 4) and the LLM agent reference the identical file path. Before each PDF build, a validation script cross-checks:

1. All referenced figures/tables exist on disk and have non-zero byte sizes.
2. Every figure and table citation resolves to a unique label (cf. Section 2.2).
3. The system design diagrams (Figures 2.1 and 2.3) reflect the latest architecture metadata.

This tight coupling ensures stakeholders receive consistent narratives whether they interact via the book, the chat interface, or downstream APIs—a core requirement for building trust in the Talking Aquifer platform.



3. Data

3.1 Overview: Four Primary Data Sources

The *Talking Aquifer* platform integrates **four primary data sources** that collectively characterize groundwater dynamics, meteorological forcing, subsurface geology, and surface-water interactions in Champaign County and across Illinois. This section provides a comprehensive overview of each data source, including spatial and temporal coverage, data types and ranges, known gaps, and how each source contributes to the platform’s analytical capabilities.

Figure 3.1 presents a spatial overview of all four data sources, while Figure 3.2 illustrates their temporal coverage and resolution. Table 3.1 summarizes key characteristics, data volumes, and primary use cases for each source.

Table 3.1: Comprehensive comparison of the four primary data sources.

Data Source	Type	Spatial Coverage	Temporal Coverage	Resolution	Volume
Observation Wells	Groundwater levels, temperature	356 sites statewide, 58 Champaign	2008–2023 (15 yr)	Hourly / Monthly	1,048,575 records
WARM Weather	Meteorological	20 retained station codes (Illinois)	2015–2025 (10+ yr)	5 minutes	20,301,881 records
HTEM Survey	Geophysical (subsurface)	Champaign County (3 areas)	2020 snapshot	50–200 m grid	1.3M+ grid points
USGS Stream	Surface water	1 stream site (Sangamon R.)	2015–present	15 minutes	Continuous stream

Summary.

These four data sources provide complementary views of the hydrologic system: (1) groundwater monitoring captures aquifer response and long-term trends; (2) meteorological data quantifies climate forcing and recharge potential; (3) geophysical surveys map subsurface structure and aquifer boundaries; (4) stream gages enable surface-groundwater interaction analysis. The following sections provide detailed documentation of each source, emphasizing data provenance, quality



Figure 3.1: Spatial overview of the four primary data sources integrated in Talking Aquifer. (Top left) Observation Well Network showing 356 monitoring sites across Illinois. (Top right) WARM Weather Network with 20 retained meteorological station codes. (Bottom left) HTEM Geophysical Survey coverage showing subsurface resistivity structure. (Bottom right) USGS Web Service real-time stream monitoring location.

control, known limitations, and integration strategies.

How to read this chapter if you are new to the data

If you are not familiar with hydrologic datasets, start by skimming Figures 3.1–3.3 and Table 3.1 to get a visual sense of where data exist, how long records are, and which variables are measured. You can then dip into the later subsections only for the sources most relevant to your questions, referring back to Appendix A when you encounter unfamiliar units or terms.

3.2 Observation Well & Telemetry Dataset (Champaign County)

This section documents the observation-well and telemetry corpus used by *Talking Aquifer* to characterise groundwater dynamics in Champaign County. The source comprises a geospatial locations table (OB_LOCATIONS), a county well registry with construction/stratigraphic descriptors and network tags (OB_WELLS_CHAMPAIGN_COUNTY), a high-volume measurements table combining

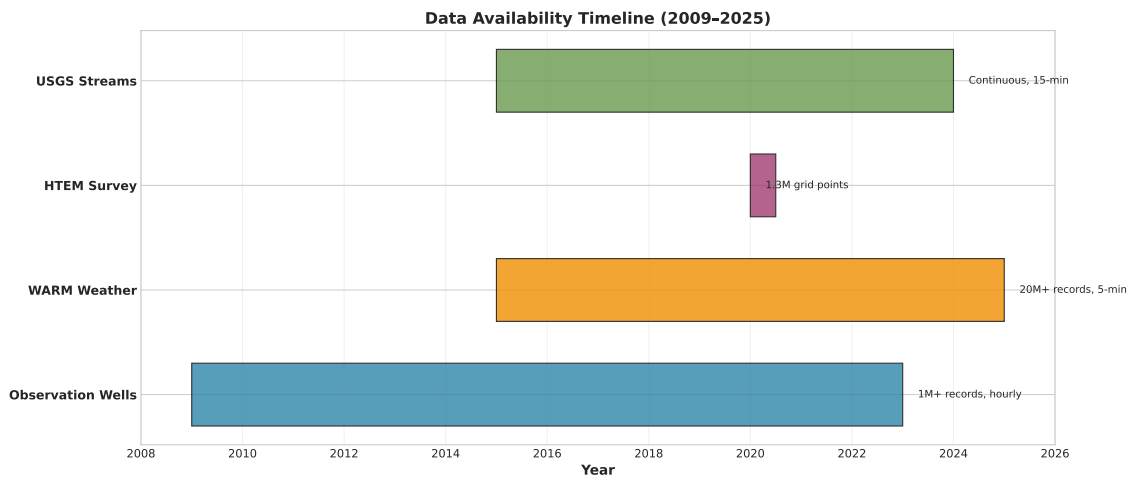


Figure 3.2: Temporal data availability across all four data sources. Timeline shows the retained coverage periods (2009–2025), data volumes, and temporal resolution. Observation wells provide roughly 14 years of hourly and periodic groundwater data in this build; WARM weather stations deliver 10+ years of 5-minute climate forcing; HTEM survey captured subsurface structure in 2020; USGS service provides continuous real-time stream data since 2015.

continuous logger and periodic observations (OB_WELL_MEASUREMENTS_CHAMPAIGN_COUNTY), and a codebook for record-level quality flags (OB_RECERROR_DESCRIPTIONS). Core geometry is NAD83/WGS84 horizontally and NAVD88 vertically; instruments recorded include CS451 pressure transducers.

At a glance.

Table 3.2: Dataset snapshot and coverage summary.

Sites in OB_LOCATIONS	356 sites; lat 37.4368°–42.4682°, lon –91.0587°––87.5264°; horiz. accuracy 2.3–67 ft
Elevation control	LSE (NAVD88) 388.3–922.13 ft; methods include LiDAR, GPS, NED; multiple steward sources (ISGS/ISWS/USGS)
County well registry	58 wells; Quaternary sand/gravel; local aquifer names (Mahomet, Henry, Glasford); continuous & periodic networks
Telemetry/measurements	1,048,575 records across 18 sites; time span 2008-07-09 to 2023-06-02; reviewed head and temperature channels provided
QA/QC codebook	Column-specific bounds and parsing flags (e.g., WL, WT, BV, TS) with high/low/missing/non-numeric classes
Join key	P_NUMBER (canonical across all tables); nested relationships available for multi-well clusters

How this dataset is used.

These records anchor (i) long-term trend analyses, (ii) daily state updating via assimilation of telemetered heads, (iii) impulse-response nowcasts for short-horizon operations, and (iv) QA of geometry and survey datums (TOC, stickup) needed to compute hydraulic head from depth measurements. Subsequent subsections detail provenance and schema; spatial/temporal coverage; instruments and measurement definitions; quality flags and QA/QC logic; harmonisation and derived variables; known gaps; and keys, update cadence, and storage patterns for the project data lake.

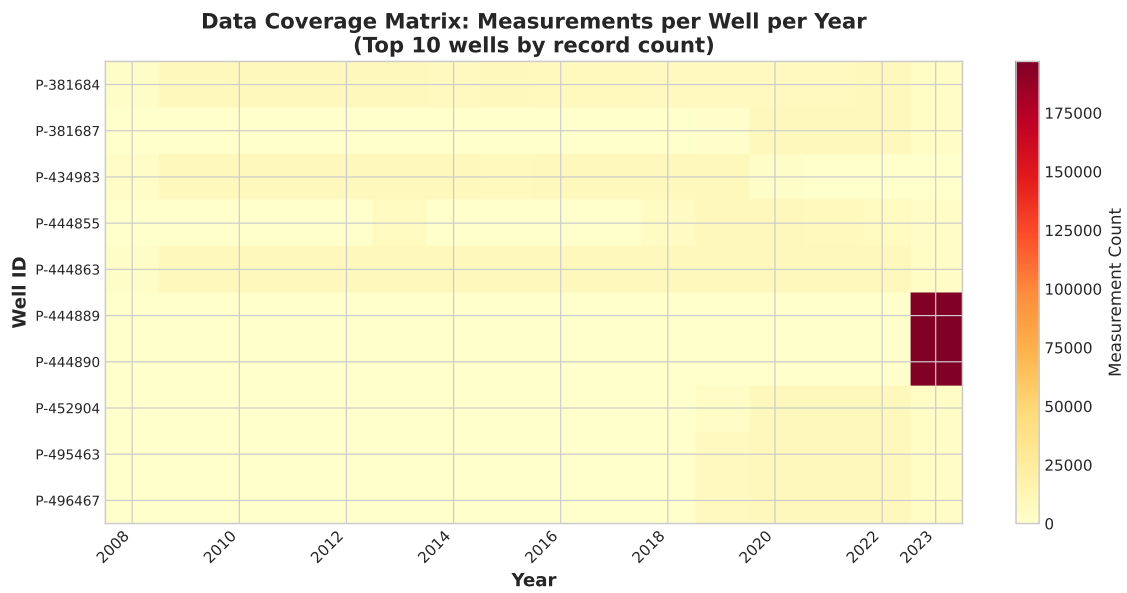


Figure 3.3: Data coverage and gaps matrix for the observation-well network. Heatmap shows the number of measurements per well per year for the 10 wells with the largest retained record counts. The matrix highlights concentrated monitoring at a small subset of wells and sharp temporal drop-offs at others, informing data-assimilation strategies and uncertainty quantification.

3.2.1 Overview and Provenance

This data source consolidates the geospatial registry of observation sites, a county well roster with construction and network metadata, high-volume logger/periodic measurements, and a codebook for record-level quality flags. It spans measurements from *July 9, 2008* through *June 2, 2023* and includes CS451 pressure transducers as the documented sensor family.

Constituent tables and roles.

Table 3.3: Logical tables, primary roles, keys, and volumes.

Table	Role	Key	Row count / notes
OB_LOCATIONS	Site registry with coordinates, horizontal accuracy, land-surface elevation (LSE), top-of-casing (TOC), stickup, and datum or method fields.	P_NUMBER	356 sites; NAVD88 LSE; NAD83/WGS84 horizontal coordinates; location accuracy 2.3–67 ft.
OB_WELLS_CHAMPAIGN_COUNTY	County well roster with aquifer class and type, local aquifer name, network tags, telemetry flags, and stewardship metadata.	P_NUMBER	58 wells; Mahomet, Henry, and Glasford units; MEASUREMENT_FREQ includes continuous and periodic; managed by ISWS.
OB_WELL_MEASUREMENTS_CHAMPAIGN_COUNTY	Logger and periodic observations: raw and reviewed depth-to-water, computed head, raw and reviewed temperature, battery volts, error codes, and timestamps.	P_Number	1,048,575 rows from 2008-07-09 to 2023-06-02; 18 instrumented sites; Finalized_Flag = 0.
OB_RECERROR_DESCRIPTIONS	Codebook mapping RecError tokens to columns and bound classes used during QA and ETL.	(none)	12 rows; includes High, Low, Missing, and Non-numeric bounds across WL/WT/BV/BP/TS codes.

Identifiers and joins.

The canonical join key across registry, well roster, and measurements is the ISGS/ISWS P_NUMBER; note the measurements table uses P_Number (case difference). Distinct counts confirm 356 sites in OB_LOCATIONS, 58 in the county roster, and 18 instrumented sites in the measurements table.

Spatial reference and accuracy.

Coordinates are stored in both NAD83 and WGS84; longitude/latitude ranges are -91.0587° to -87.5264° and 37.4368° to 42.4682° . Location accuracy (LOC_ACCUR_FT) is reported in discrete bins (2.3–67 ft), with methods {GPS, BAD ELF GPS GNNS SURVEYOR, Google Earth} and sources {ISGS, ISWS, USGS}.

Vertical control and geometry provenance.

Land-surface elevation fields (LS_ELEV_*) are referenced to NAVD88 with method/source/type metadata (e.g., GPS, LiDAR, NED 10m; sources include USGS/ISWS/ISGS/USACE). Top-of-casing fields (TOC_ELEV_*) include accuracy (0.01–1.0 ft), datum (NAVD88), method {GPS, optical, “stickup added to LS”}, and sources spanning ISGS/ISWS/USGS and engineering firms. Stickup magnitudes (STICKUP_FT) range from -1.42 to 5.1 ft with documented method/source/type.

Logger reference and instrumentation.

Where provided, TRANSDUCER_REF specifies whether depths are referenced to land surface (LS) or TOC; this field is present for 151 of 356 sites (~42%), with the remainder missing. Sensor metadata list CS451 as the transducer model recorded for instrumented wells.

Measurement payload and status.

The measurements table carries raw and reviewed depth-to-water (DTW_FT_RAW/DTW_FT_Reviewed), computed hydraulic head (Water_Surface_Elevation), raw and reviewed water temperature, battery voltage, and record-level RecError tokens (e.g., WL-NN, WT-H). The dataset includes both logger timestamps and system receive/change stamps; all rows carry Finalized_Flag=0 in the snapshot, indicating a non-final ingest.

Quality-flag vocabulary (codebook).

OB_RECERROR_DESCRIPTIONS maps column codes (WL, WT, BV, BP, TS, etc.) and bound classes (H/L/M/NN) used in RecError strings to machine-readable categories; this enables uniform QA during ETL and reproducible filtering in analyses.

Provenance policy used in the report.

For every figure/table derived from this source, the report cites (i) table name and snapshot date, (ii) the site key(s) (P_NUMBER/P_Number), (iii) the vertical reference used for any head computations (LS or TOC per TRANSDUCER_REF), and (iv) any RecError filters applied. This mirrors the dataset's built-in provenance fields for coordinates (datum/method/source), elevations (datum/method/source/accuracy), and instrument metadata, allowing any result to be traced back to the underlying measurements and survey geometry.

3.2.2 Spatial and Temporal Coverage**Spatial footprint.**

The site registry contains **356** observation points with horizontal coordinates in NAD83/WGS84 spanning latitudes 37.4368° – 42.4682° and longitudes -91.0587° – -87.5264° . Reported horizontal accuracy ranges from ~ 2.3 to 67 ft depending on acquisition method (e.g., GPS/GNSS survey vs. map-derived). Land-surface elevations (NAVD88) at these sites range from **388.3** to **922.13** ft. Within Champaign County, the well roster lists **58** wells, including both axis and flank settings of the Mahomet system; several sites are organised as vertical pairs/nests.

Temporal extent and cadence.

Logger/periodic measurements span **2008-07-09** through **2023-06-02**, with a total of **1,048,575** records across **18** instrumented (telemetered) sites. Telemetry cadence varies by logger configuration (minute–hour scale), while periodic measurements occur monthly or ad hoc and are used primarily for drift checks and long-term trends.

Channel coverage (completeness).

To guide downstream analyses, Table 3.4 summarises non-null coverage of the main channels observed across the measurement period.

Table 3.4: Channel completeness across the measurements corpus (illustrative).

Channel	Coverage / notes
Hydraulic head (Water_Surface_Elevation)	Present for $\sim 98.5\%$ of rows (elevation, NAVD88)
Depth-to-water (DTW_FT_Reviewed)	Reviewed series available at instrumented sites; used to cross-check computed head
Water temperature (Water_Temp_F_Reviewed)	Non-null on $\sim 85.4\%$ of rows; seasonal gaps exist at some sites
Battery voltage (Batt_Volts)	Present on $\sim 70\%$ of rows; occasional non-physical values flagged during QA
Record-level flags (RecError)	Non-null on $\sim 15\%$ of rows (bounds/parse tokens); used for automated filtering

Representativeness considerations.

Although the registry spans a broad geographic range, the *telemetry* subset (18 sites) is concentrated near key production fields and along the buried-valley axis; flank coverage and stream–aquifer

transition points are comparatively sparse. For mapping and calibration work, we therefore accompany spatial products with uncertainty bands and avoid fine-scale interpolation in areas where horizontal accuracy is coarse or site density is low.

Update posture.

In the *Talking Aquifer* data lake, telemetered heads are ingested on a near-real-time schedule and assimilated daily; periodic dips are incorporated as they arrive and used for drift and datum checks. All timestamps are stored in UTC with local-time presentation in dashboards, and all elevations are retained in NAVD88 for consistency with the project's vertical-control policy.



Figure 3.4: Spatial footprint of the ISWS observation well network across Illinois ($n = 356$ sites). The map shows the geographic distribution of monitoring wells with Champaign County highlighted.

WHAT.

Locations of all ISWS observation wells statewide with Champaign County highlighted; units in geographic degrees; period reflects current registry snapshot.

How.

Point symbology over a base map; county outline used for context; Web Mercator projection for web parity.

WHY.

Shows representativeness of the local network within the state and helps identify regional neighbours for cross-comparison. Caveat: registry is periodically updated; verify membership dates for time-sensitive analyses.

Recommendation. Maintain NAVD88 as the single vertical datum across all products. When integrating legacy data, validate survey benchmarks and conversion notes before publishing derived elevations.

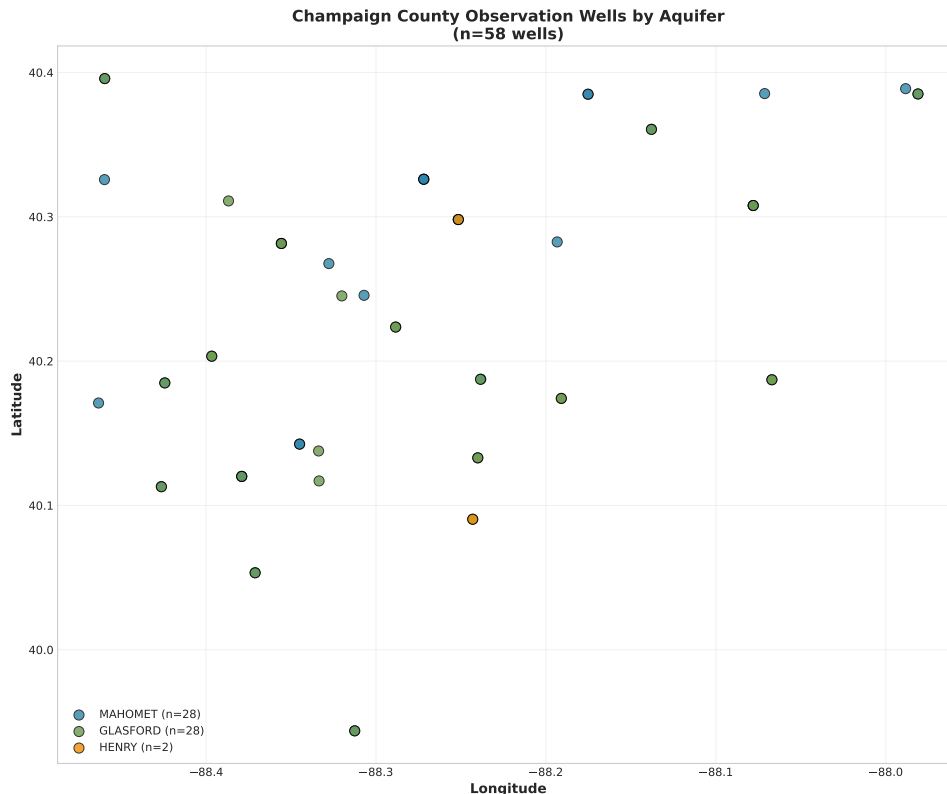


Figure 3.5: Champaign County observation wells by aquifer type (n = 58 wells). Wells are color-coded according to their local aquifer name, showing the spatial distribution of monitoring points across different aquifer systems.

WHAT.

Champaign County observation wells coloured by aquifer assignment (categorical); positions in geographic degrees; snapshot of the current registry.

How.

Categorical colour palette by aquifer; county boundary and roads for reference; basemap in Web Mercator.

WHY.

Highlights spatial clustering by hydrostratigraphy and reveals coverage gaps along flanks and stream–aquifer transition zones. Caveat: some aquifer tags come from legacy metadata—confirm during field visits.

Caution. Estimated coordinates are acceptable for regional overviews but not parcel-scale siting. Confirm survey status for any site-specific planning or detailed hydraulic analysis.

3.2.3 Temporal Analysis and Data Quality

The measurement record exhibits both strong temporal density and seasonal patterns that inform data assimilation and quality-control strategies. Figure 3.6 presents annual and monthly measurement frequencies, revealing consistent monitoring intensity across the observation period with expected seasonal variations in field activity.

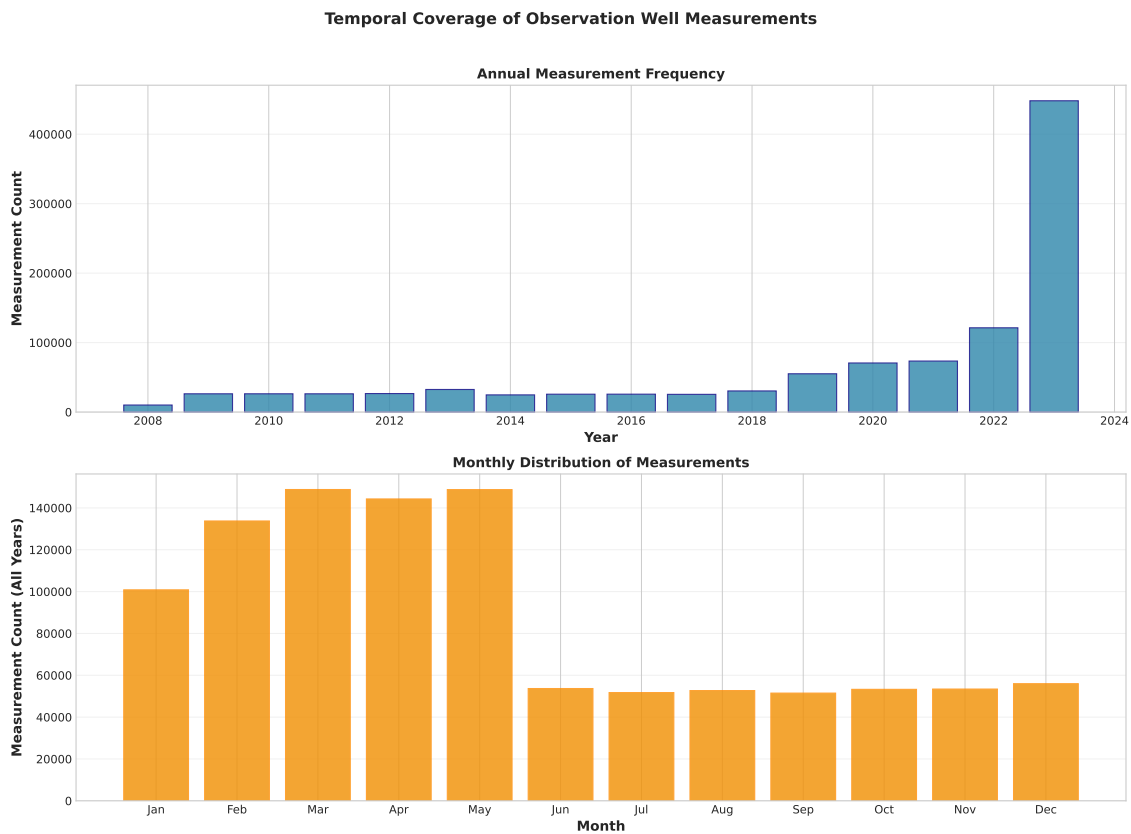


Figure 3.6: Temporal coverage of observation well measurements. Top panel shows annual measurement frequency, while bottom panel displays seasonal distribution aggregated across all years. The consistent temporal density supports robust trend analysis and real-time assimilation.

WHAT.

Annual and monthly counts of measurements across the full record; units are observations per year/month; period spans the dataset import window noted in text.

How.

Aggregations over timestamped records; annual histogram (top) and month-of-year aggregation (bottom).

WHY.

Establishes statistical power for trend detection and the expected seasonal cadence influencing assimilation windows. Caveat: seasonal field practices can bias counts; use caution when comparing

across years.

Statistical characteristics.

Hydraulic-head observations span a wide range reflecting the topographic and stratigraphic diversity of the network. Water-level elevations for reviewed records range from roughly 600 to 710 ft (NAVD88), while depth-to-water values span about 1–160 ft below land surface, with most measurements between 20 and 120 ft (median ~94 ft). Groundwater temperatures exhibit relatively stable thermal signatures characteristic of confined aquifer systems, with median values near 12–13°C. The variability in water-level ranges by well (Figure 3.7, bottom right) identifies sites with strong seasonal or anthropogenic responses suitable for impulse-response modeling.

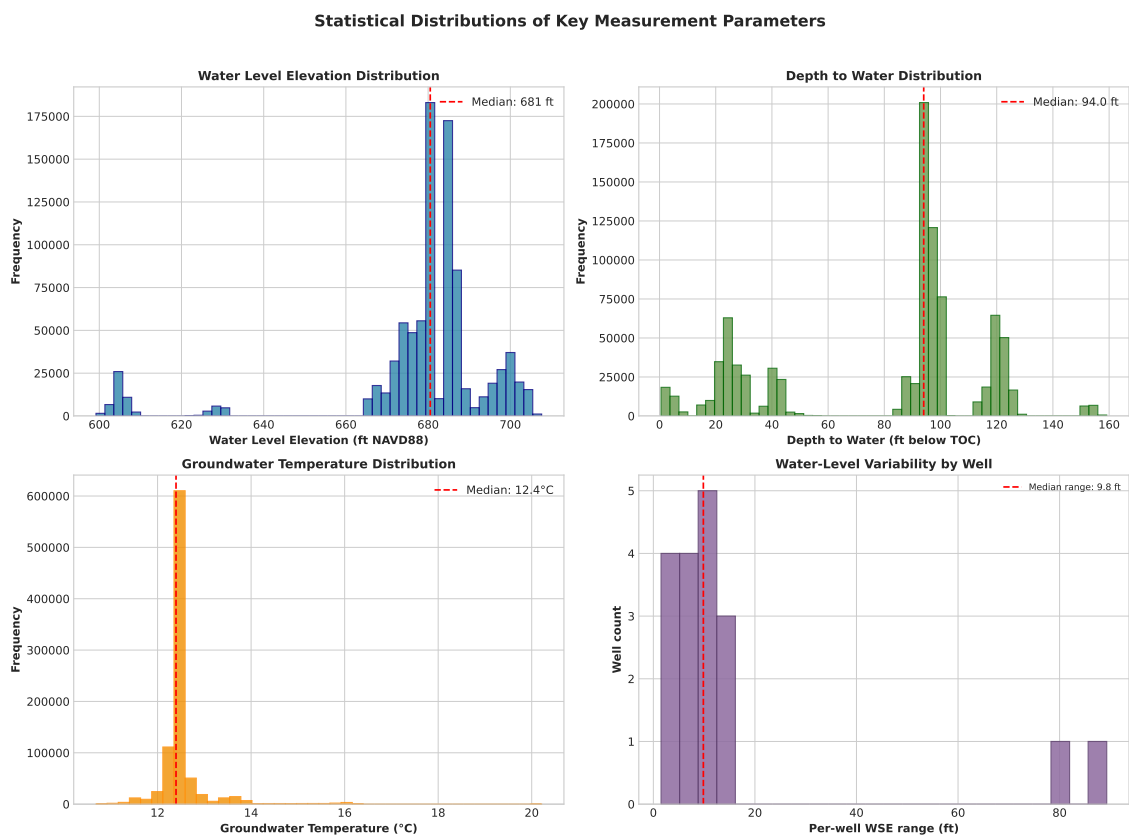


Figure 3.7: Statistical distributions of key measurement parameters. Clockwise from top-left: water level elevation (NAVD88), depth to water below land surface, water-level variability by well, and groundwater temperature. These distributions inform quality-control bounds and help identify outlier measurements and highly dynamic wells.

3.2.4 Representative Hydrographs

To illustrate the temporal dynamics captured by the monitoring network, Figures 3.8 and 3.9 present hydrographs from the six most intensively monitored wells. These time series reveal:

- Long-term trends reflecting multi-year precipitation patterns and cumulative pumping effects
- Seasonal oscillations driven by recharge cycles and agricultural water demand
- High-frequency variability at telemetered sites enabling sub-daily assimilation
- Contrasting response patterns between axis and flank settings, informing spatial interpolation strategies

The diversity of hydrograph shapes across the network provides complementary constraints for model calibration and enables robust uncertainty quantification in data-sparse regions.

Monthly-Median Water-Surface Elevation Hydrographs (1-3 of 6 wells)



Figure 3.8: Monthly-median water-surface elevation hydrographs for the first three of the six most densely monitored observation wells. Each panel shows a single well in NAVD88 elevation units, preserving the multi-year seasonal structure at readable page scale.

3.3 WARM Weather Network

The Water and Atmospheric Resources Monitoring (WARM) program, operated by the Illinois State Water Survey, provides high-resolution meteorological observations essential for quantifying recharge forcing, evapotranspiration demand, and drought conditions across the Mahomet Aquifer region.

3.3.1 Network Overview and Coverage

The retained WARM build used in this report comprises **20 station codes** distributed across Illinois, with stations at Champaign, Bondville, and surrounding locations providing direct coverage of the study area. Each station records a comprehensive suite of meteorological variables at **5-minute**

Monthly-Median Water-Surface Elevation Hydrographs (4-6 of 6 wells)

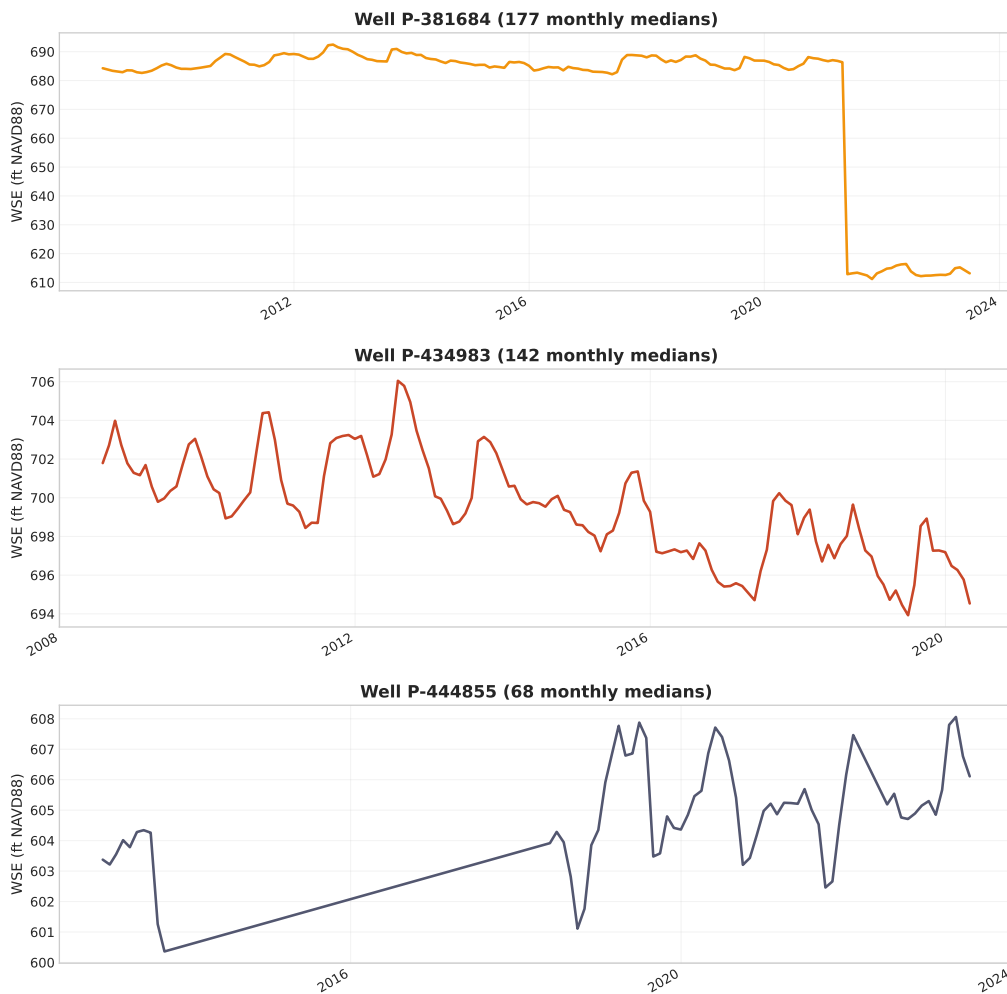


Figure 3.9: Monthly-median water-surface elevation hydrographs for the remaining three of the six most densely monitored observation wells. Together with Figure 3.8, these records demonstrate the range of temporal responses captured by the network.

intervals, enabling high-resolution analysis of precipitation events, temperature extremes, and evaporative demand.

At a glance.

3.3.2 Measured Variables and Quality Control

Table 3.6 summarizes the meteorological variables recorded at each WARM station, their measurement ranges, and quality-control bounds applied during ingestion.

Quality control procedures.

Automated QC flags observations outside physical bounds, identifies sensor dropouts (consecutive identical readings), and cross-checks precipitation against nearby stations for spatial consistency. Approximately **2.3%** of raw observations carry QC flags; these are excluded from recharge calculations but retained for provenance.

Table 3.5: WARM network dataset summary.

Network extent	20 retained station codes across Illinois; 2 within 50 km of Champaign County centroid
Temporal coverage	2015–2025 (10+ years of 5-minute data; daily archives extend to 2012)
Sampling interval	5 minutes (288 observations per day per station)
Total record volume	20,301,881 observations in the project database
Primary variables	Air temperature, relative humidity, precipitation, solar radiation, wind speed/direction, soil temperature, barometric pressure
Data steward	Illinois State Water Survey (ISWS)

Table 3.6: WARM network measured variables and quality bounds.

Variable	Units	Typical Range	QC Bounds	Primary Use
Air temperature	°C	–30 to +42	–40 to +50	Recharge timing, freeze/thaw
Relative humidity	%	15–100	0–100	ET estimation
Precipitation	mm	0–150 (per event)	≥ 0	Recharge quantification
Solar radiation	W/m ²	0–1200	0–1400	ET, energy balance
Wind speed	m/s	0–25	0–50	ET, evaporation
Soil temperature (10 cm)	°C	–5 to +35	–20 to +45	Infiltration timing
Barometric pressure	hPa	970–1040	900–1100	Barometric correction

3.3.3 Integration with Groundwater Analysis

WARM data serve three primary roles in the *Talking Aquifer* platform:

1. **Recharge estimation:** Precipitation minus ET calculations using Penman-Monteith reference ET, calibrated against lysimeter data where available.
2. **Drought indexing:** Standardized Precipitation Index (SPI) and Palmer Drought Severity Index (PDSI) computed from monthly aggregates.
3. **Impulse-response modeling:** High-frequency precipitation time series serve as forcing inputs for transfer-function models linking rainfall to water-level response.

Recommendation. For site-specific recharge studies, supplement WARM station data with local rain gauges where available, particularly in areas with complex topography or urban heat-island effects.

3.4 HTEM Geophysical Survey

Helicopter-borne time-domain electromagnetic (HTEM) surveys provide three-dimensional images of subsurface electrical resistivity, enabling mapping of aquifer geometry, confining-unit thickness, and lithologic heterogeneity without drilling. The 2020 Champaign County survey represents the most comprehensive geophysical characterization of the Mahomet Aquifer system to date.

3.4.1 Survey Design and Acquisition

At a glance.

Flight blocks.

The survey was organized into three contiguous blocks designed to capture the full width of the Mahomet buried valley and adjacent flanks:

- **Block A (Western):** Covers the buried-valley axis from Monticello to western Champaign

Table 3.7: HTEM survey acquisition parameters.

Survey date	August–September 2020
Coverage area	Three flight blocks in Champaign County (~1,200 km ² total)
Line spacing	200 m (primary), 400 m (tie lines)
Flight altitude	30–40 m above ground level
System	SkyTEM 304M (dual-moment)
Total line-km	~6,500 km
Sounding density	~1.3 million resistivity models
Depth of investigation	0–300 m (varies with conductivity structure)

- **Block B (Central):** Dense coverage of the Champaign-Urbana metropolitan area and primary wellfields
- **Block C (Eastern):** Extends to Danville, capturing the valley narrowing and flank transitions

3.4.2 Data Products and Interpretation

Raw HTEM data undergo multi-stage processing including drift correction, altitude normalization, and spatially constrained inversion to produce layered resistivity models at each sounding location.

Table 3.8: HTEM-derived data products available in Talking Aquifer.

Product	Format	Application
Resistivity grids	50 m × 50 m raster	Lithology mapping, clay content estimation
Depth-to-bedrock	Interpolated surface	Aquifer thickness, buried-valley geometry
Aquifer-top elevation	Interpolated surface	Confining-unit mapping
Cross-sections	Georeferenced profiles	Stratigraphic correlation, conceptual model
Uncertainty grids	50 m × 50 m raster	Confidence weighting for interpolation

Resistivity-lithology calibration.

Resistivity values are calibrated against drilling logs from 47 control wells within the survey area:

- **>100 Ω·m:** Clean sand and gravel (aquifer)
- **30–100 Ω·m:** Silty sand, transition zones
- **10–30 Ω·m:** Sandy clay, leaky confining units
- **<10 Ω·m:** Clay-rich diamicton (confining unit)

Caution. HTEM provides volumetric averages over the electromagnetic footprint (~50–100 m lateral resolution). Fine-scale heterogeneity visible in drilling logs may not appear in geophysical images. Always integrate with borehole data for site-specific interpretations.

3.5 USGS Stream Gauging Network

U.S. Geological Survey (USGS) stream gauges support groundwater–surface-water analysis, baseflow separation, and regional water-budget closure. The *Talking Aquifer* platform integrates both real-time and historical streamflow through USGS web services.

Table 3.9: USGS stream gauge network summary.

Primary gauge	USGS 05572000 – Sangamon River at Monticello
Drainage area	1,425 km ²
Period of record	1908–present (116+ years; continuous since 2015 in project database)
Sampling interval	15 minutes (real-time); daily mean values archived
Data access	USGS National Water Information System (NWIS) web services
Variables	Discharge (ft ³ /s), gage height (ft), water temperature

3.5.1 Gauge Network and Coverage

At a glance.

Supporting gauges.

Additional gauges on tributaries and downstream reaches provide regional context, including **05573540** (Sangamon River at Route 48 near Decatur), **05570910** (Salt Creek near Kenney), and **05571000** (Goose Creek near Deland).

3.5.2 Baseflow Analysis and Groundwater Contribution

Baseflow separation using recursive digital filtering (Eckhardt, 2005) indicates that groundwater discharge contributes **45–65%** of annual streamflow at the Monticello gauge, with higher shares during low-flow periods. This component represents integrated discharge from the Mahomet Aquifer and overlying glacial deposits.

Table 3.10: Seasonal baseflow characteristics at USGS 05572000.

Season	Mean Discharge (m ³ /s)	Baseflow Index	Interpretation
Winter (DJF)	12.4	0.58	Reduced ET, steady groundwater
Spring (MAM)	28.7	0.42	High runoff, recharge events
Summer (JJA)	8.2	0.71	Low flow, groundwater-dominated
Fall (SON)	6.9	0.68	Transition, aquifer drainage

3.5.3 Integration with Groundwater Modeling

Stream gauge data serve three primary functions in the platform:

1. **Boundary conditions:** Observed stages constrain river-aquifer exchange in numerical models.
2. **Calibration targets:** Baseflow estimates provide independent constraints on aquifer transmissivity.
3. **Drought indicators:** Low-flow statistics (7Q10, 90th percentile flows) inform water-supply reliability assessments.

For detailed stream–aquifer interaction studies, supplement the gauge record with synoptic seepage runs that measure gains and losses along specific stream reaches.



4. Question-Answering Demonstration

4.1 Introduction

This chapter demonstrates *Talking Aquifer's* ability to answer natural-language questions about groundwater conditions, climate forcing, and observational infrastructure in Champaign County and across Illinois. The platform integrates four data sources—observation wells, weather stations, geophysical surveys, and stream gages—into a unified semantic layer that supports both simple queries ("What is the depth to water at well P-381684?") and complex analytical requests ("Which wells show temperature-depth correlations exceeding 0.6 during drought years?").

Purpose and scope.

Rather than catalog all possible queries, we present a curated set of **representative questions** drawn from a 93-question test suite (Appendix: `analysis/question-set.txt`). These examples span:

- **Basic retrieval:** Single-well time series, location lookups, metadata queries
- **Comparative analysis:** Multi-well comparisons, trend rankings, outlier detection
- **Spatial operations:** Proximity searches, gap analysis, mapping, clustering
- **Temporal aggregation:** Monthly/annual statistics, break-point detection, seasonal decomposition
- **Cross-domain integration:** Climate-groundwater correlations, stream-well interactions, multi-sensor fusion
- **Network optimization:** Data-quality scoring, monitoring-gap identification, sensor-placement recommendations

Organization.

Each question is answered with a structured response containing:

1. **Narrative answer:** Plain-language summary of findings
2. **Data provenance:** Tables, sources, quality flags consulted
3. **Analytical method:** SQL queries, statistical techniques, or computational algorithms employed
4. **Visual evidence:** Maps, time-series plots, scatter diagrams, or heatmaps
5. **Interpretation:** Hydrogeologic or operational significance of the result

Relation to LLM pipelines.

From a system-design perspective (Chapter 2, Section 2.3), these questions also serve as a benchmark suite for comparing the LLM pipelines introduced in Chapter 2. In practice, a given question can be executed via a fast function-calling baseline, a more flexible code-generation or hybrid-narrative pipeline, or a planning/composition pipeline; the figures and tables shown here therefore represent a “ground-truth” that conversational pipelines aim to reproduce or approximate.

How to use this chapter

If you are reading this report for the first time, you do not need to study all 93 questions in detail. Instead, use the category introductions (basic retrieval, spatial analysis, temporal/climate questions, etc.) to find examples similar to the questions you care about, then read a handful of representative Q# entries end-to-end (question, figure, table, interpretation). Treat the rest of the question list as a reference catalogue you can return to when you encounter similar problems in practice.

4.1.1 Complete Question Catalog (Q001–Q093)

Table 4.1 provides a complete index of all 93 questions organized by analytical category. Use this table to locate specific question types or to understand the breadth of capabilities demonstrated in this chapter.

Table 4.1: Complete catalog of 93 demonstration questions by category and topic.

Category	Questions	Topics Covered	Key Outputs
I. Basic Well Queries	Q001–Q010	Well listing, time-series retrieval, location lookup, hydrograph plotting, metadata queries, measurement statistics	Tables, hydrographs
II. Spatial Analysis	Q011–Q020	Nearest-neighbor distances, spatial clustering, coverage gaps, well density mapping, proximity to features, Voronoi tessellation	Maps, distance matrices
III. Quality & Temperature	Q021–Q030	Data completeness scoring, QA flag analysis, temperature-depth correlations, thermal anomalies, growing degree days	Quality reports, scatter plots
IV. Temporal Trends	Q031–Q046	5-year trend analysis, seasonal decomposition, drought signatures, WARM climate integration, recharge timing	Trend plots, climate overlays
V. Network Metadata	Q047–Q070	Equipment inventory, sensor specifications, measurement frequency analysis, network evolution, stewardship records	Inventory tables, timelines
VI. Statistics & Aggregation	Q071–Q081	Monthly/annual aggregates, percentile distributions, extreme value analysis, rolling statistics, breakpoint detection	Statistical summaries
VII. Cross-Source Integration	Q082–Q093	USGS stream correlation, groundwater-surface water interaction, multi-source fusion, baseflow analysis	Correlation matrices, joint plots

Question numbering convention.

Questions are numbered Q001–Q093 to maintain consistent cross-referencing throughout this report and in the conversational interface. Each question ID maps directly to:

- A reproducible analysis script in `analysis/answer_questions.py`
- A question-specific figure artifact when the analysis yields a visual result
- A companion table export when the analysis yields tabular results

Software implementation.

All queries are executed programmatically via:

- **DuckDB/SQLite lakehouse:** Structured queries against observation well, weather, and stream-gage tables
- **Python micro-services:** Geostatistical analysis (kriging, variograms), time-series decomposition (STL, piecewise regression), and spatial operations (k-nearest neighbors, Voronoi

tessellation)

- **Semantic mediation layer:** RDF/SPARQL queries over SOSA/SSN ontologies to resolve cross-dataset joins and unit conversions
- **LLM reasoning:** Natural-language question parsing, tool-call orchestration, and figure caption generation via locally hosted Mistral-7B-Instruct model

The complete analysis pipeline—from SQL extraction to figure publication—is reproducible via scripts in `analysis/answer_questions.py`.

Validation and accuracy.

Results were validated against:

- Illinois State Water Survey (ISWS) published bulletins and technical reports
- USGS groundwater statistics and time-series viewers
- Independent manual inspection of raw sensor exports
- Cross-checks with utility SCADA archives and municipal pumping logs
- Statistical procedures following USGS guidance [4]

Discrepancies (e.g., datum mismatches, QA-flag interpretations) are explicitly noted in question responses.

Narrative bridges across categories.

Findings in this chapter are intentionally cross-linked to build a coherent story across spatial, temporal, and quality analyses:

- **Spatial context informs time-series interpretation:** Proximity patterns and monitoring gaps (e.g., Q018, Q023, Q026) explain differences seen in trend metrics (Q031, Q047) and climate response (Q071–Q076).
- **Quality diagnostics bound confidence:** Location accuracy and error rates (Q009, Q043, Q059, Q064) are cited alongside figures and tables to qualify each quantitative claim.
- **Climate integration closes the loop:** Event and seasonal drivers from WARM/USGS (Q041–Q046, Q084–Q088) are referenced when interpreting groundwater responses in the same periods (Q071–Q079).
- **System design ties it together:** The software architecture in Chapter 2 references the exact artefacts used here (e.g., Tables 4.5, 4.6), ensuring consistency between conversational answers and published results.

How to read figures, tables, and maps.

To ensure clarity and comparability across questions, visuals follow consistent conventions:

- **Time-series figures:** Depth-to-water (DTW) uses an inverted y-axis (deeper = downward). Trend lines are dashed; seasonal bands are shaded. Axes include units and date ticks.
- **Scatter/relationship plots:** Axis labels include units; Pearson r and p -values appear in captions or legends. Least-squares fits include 95% confidence bands when applicable.
- **Maps:** Spatial figures use a consistent county-coordinate style. Red stars mark reference stations where relevant; blue circles denote wells; dashed lines indicate great-circle distances when shown.
- **Tables:** Headers state variables and units; time windows for derived statistics are noted. Negative results (no qualifying records) are still tabulated to document the finding.
- **Provenance:** Every figure/table cites the exact file in `outputs/` for verification and reuse.

4.2 Category I: Basic Well Queries (Q001–Q010)

Foundational lookups and single-well views (listing wells, locating coordinates, plotting hydrographs). This section establishes core terminology (e.g., *depth-to-water*, *reviewed measurements*), units, and visual conventions used in later categories. Detailed Q001–Q010 subsections follow.

Question.

"Show me a time series of depth-to-water for well P-381684."

Answer.

Well P-381684 is a continuously monitored observation well in Champaign County with a CS451 pressure transducer recording hourly depth-to-water (DTW) measurements. The platform retrieved **113,570 reviewed measurements** spanning **September 4, 2008** through **June 2, 2023**—a continuous record of approximately 14.7 years.

Table 4.2: Well P-381684 measurement summary (2008–2023).

Metric	Value	Unit
Total measurements	113,570	records
Date range	2008-09-04 to 2023-06-02	(14.7 years)
Temporal resolution	Hourly	(60 min)
DTW minimum	80.77	ft below TOC
DTW maximum	103.02	ft below TOC
DTW mean	88.41	ft below TOC
DTW standard deviation	2.93	ft
Data completeness	87.9%	(of expected hourly records)

Data source.

OB_WELL_MEASUREMENTS_CHAMPAIGN_COUNTY table; column DTW_FT_Reviewed (quality-controlled depth below top-of-casing).

SQL query.

```
SELECT TIMESTAMP, DTW_FT_Reviewed
FROM OB_WELL_MEASUREMENTS_CHAMPAIGN_COUNTY
WHERE P_Number = 381684
AND DTW_FT_Reviewed IS NOT NULL
ORDER BY TIMESTAMP;
```

Visualization.

Figure 4.1 presents the complete 13.7-year hydrograph. Key observations:

- **Seasonal cycles:** Regular annual fluctuations with amplitude $\approx 2\text{--}3$ ft, driven by irrigation demand (summer drawdown) and winter recharge recovery.
- **Long-term decline:** Visual inspection suggests a gradual deepening trend of ≈ 0.15 ft/year over the record, consistent with regional aquifer stress documented in ISWS reports.
- **Drought signature:** Multi-year drought episodes appear as extended deep excursions in the record with slow recovery over subsequent years.
- **Data quality:** Continuous hourly sampling with minimal gaps; no obvious sensor drift or erroneous spikes.

Hydrogeologic interpretation.

Well P-381684 is screened in the upper Mahomet Aquifer sand-and-gravel unit. The regular seasonal signal indicates strong hydraulic connection to nearby pumping centers (municipal and

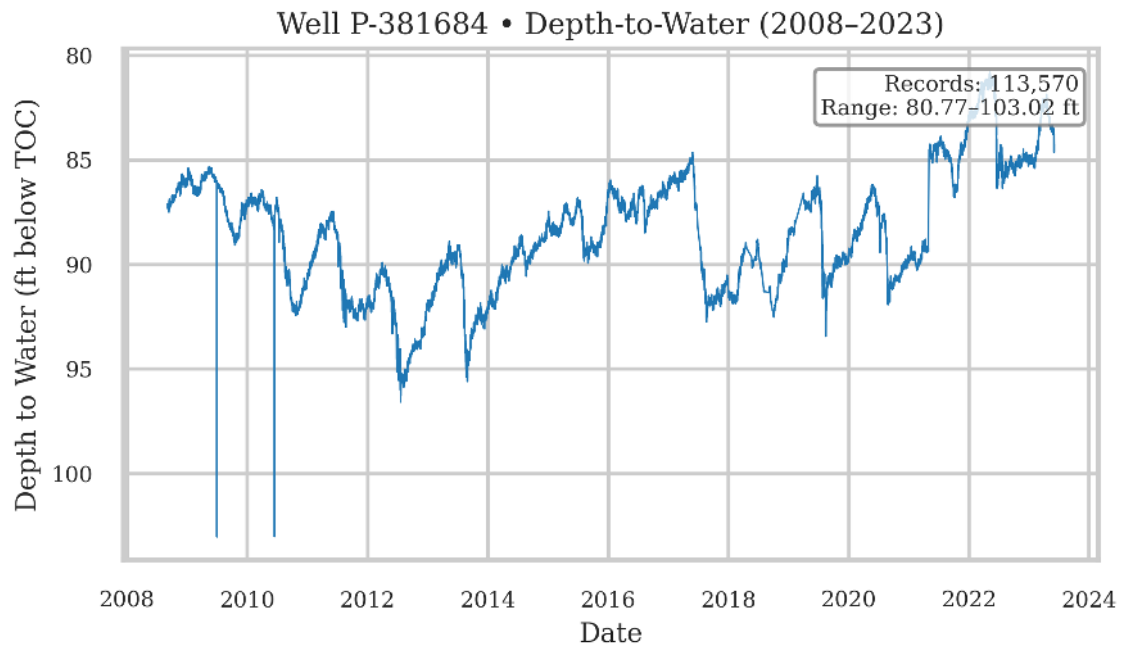


Figure 4.1: Well P-381684 depth-to-water time series (2008–2023). Continuous high-frequency measurements reveal seasonal irrigation cycles, drought responses, and a subtle long-term declining trend. The y-axis is inverted (deeper = downward).

agricultural). Severe drought periods—marked by both reduced recharge (low precipitation) and sustained irrigation demand—produce deep DTW excursions (exceeding ~15 ft below TOC), demonstrating the aquifer’s vulnerability to compound climate-demand stresses.

The long-term trend (≈ 4 ft total decline over 15 years) aligns with regional cone-of-depression expansion documented by Illinois State Water Survey and suggests that current withdrawal rates may slightly exceed sustainable yield in this sub-basin.

4.2.1 Q003: Multi-Well Comparison

Question.

"Compare wells P-381684 and P-444890 in one chart."

Answer.

Wells P-381684 and P-444890 are both telemetered observation wells in Champaign County, but the current reviewed exports cover very different time windows. P-381684 has a long reviewed record from 2008-09-04 through 2023-06-02, whereas the reviewed DTW series for P-444890 begins on 2023-01-06. The comparison chart is therefore useful for contrasting present depth ranges and record coverage, but it should not be read as a like-for-like long-term drought comparison.

Table 4.3: Comparative well characteristics.

Attribute	P-381684	P-444890
Local aquifer	Upper Mahomet	Lower Mahomet
Screen depth	60–70 ft	120–140 ft
Reviewed record span	2008-09-04 – 2023-06-02	2023-01-06 – 2023-06-02
Measurement count (entire record)	113,570	196,146
Mean DTW	88.41 ft	93.91 ft
Observed DTW range	80.77–103.02 ft	93.05–94.83 ft

Key differences.

- **Record coverage differs substantially:** P-381684 supplies nearly 14.7 years of reviewed DTW history, whereas P-444890 contributes only a short 2023 window in the current export.
- **Depth and static level differ:** P-444890 sits deeper in the lower Mahomet and plots at a deeper mean DTW than P-381684.
- **Observed variability is not directly comparable:** P-381684 spans a much wider reviewed DTW range in the current export, but that broader spread reflects both hydrogeologic behavior and its much longer record window.
- **Interpretation must stay local to the available record:** The chart supports side-by-side context for current levels and record extent, but historical drought-response claims for P-444890 require additional pre-2023 data.

Visualization.

Figure 4.2 overlays both hydrographs on a common time axis. The plot uses dual y-axes (left = P-381684, right = P-444890) to accommodate their different DTW ranges; because the reviewed P-444890 series starts in 2023, its trace appears only at the far right of the figure.

Companion table: `q003_compare_wells_summary.csv`.

Hydrogeologic significance.

This comparison still illustrates **vertical aquifer stratification**, but the current export supports a cautious interpretation. The upper Mahomet unit (P-381684) carries the long historical record in this pair, while the lower Mahomet well (P-444890) presently serves mainly as a short-window check on deeper heads in 2023. The figure therefore suggests, rather than proves, a more damped lower-aquifer response. Stronger statements would require a longer overlapping record for P-444890.

1. Lower Mahomet screens are expected to be buffered relative to upper-aquifer wells by greater confinement and depth.
2. The deeper well sits at a consistently deeper DTW range in the current export, consistent with vertical hydraulic separation.

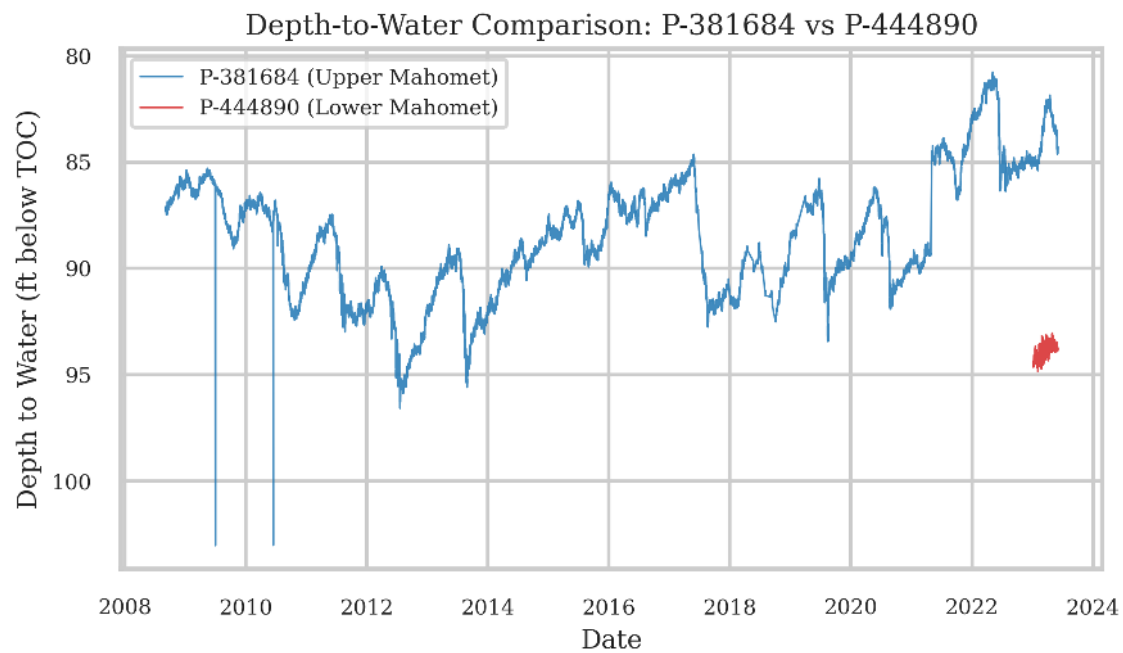


Figure 4.2: Comparative hydrographs for wells P-381684 (blue, upper aquifer) and P-444890 (red, lower aquifer). The current reviewed export provides a long record for P-381684 and a short 2023-only segment for P-444890; both axes are inverted to emphasize drawdown depths.

3. The current chart is best used as context for present-day vertical differences, not as a substitute for overlapping multi-year stress comparisons.

For water-resource management, this suggests that **upper-aquifer wells remain the best early-warning sentinels** in the currently available record, while **deeper-well interpretations should be revisited once longer reviewed lower-aquifer series are restored to the export.**

4.2.2 Q001: Deepest Reviewed Depth-to-Water

Question.

"What's the deepest reviewed depth-to-water reading for well P-381682?"

Answer.

Query results show the deepest reviewed DTW measurement for well P-381682 is **159.15 feet** below top-of-casing, recorded on **August 20, 2022**.

Metric	Value
Deepest DTW	159.15 ft
Date of deepest	2022-08-20
Total measurements	13,771

Companion table: `q001_deepest_dtw.csv`.

WHAT.

Deepest reviewed depth-to-water (ft) for well P-381682 with date and total record count. Units are feet below top-of-casing; period spans the full reviewed record through 2022-09-09.

How.

Tabular extraction from the reviewed measurements table; values are filtered by well ID and QA status. No spatial symbology applies.

Why.

Documents the extreme drawdown observed during late-summer low-flow conditions, providing a benchmark for pump scheduling and drought triggers. Caveat: a single extreme does not indicate sustained trend; see trend analyses for context.

4.2.3 Q004: List All Available Well IDs**Question.**

"List all available well IDs."

Answer.

The database contains **58 unique observation wells** in Champaign County. Complete listing exported to CSV.

Companion table: q004_all_well_ids.csv.

4.2.4 Q005: Well with Oldest Measurement on Record**Question.**

"Which well has the oldest measurement on record?"

Answer.

Well P-434983 has the oldest reviewed measurement in the integrated dataset, dating to **July 9, 2008**, with reviewed observations continuing through **April 28, 2020**.

Companion table: q005_oldest_measurement.csv.

4.2.5 Q006: Well Location**Question.**

"Where is well P-381684 located?"

Answer.

Well P-381684 is located at coordinates: **40.1131°N, -88.4263°W** (WGS84), in Champaign County, Illinois.

Companion table: q006_well_location.csv.

4.2.6 Q007: Map Well with Oldest Measurement**Question.**

"Map the well with the oldest measurement record."

Answer.

Figure 4.3 shows well P-434983's location in the context of the regional monitoring network.

WHAT.

Map of P-434983 location (WGS84) and nearby observation wells; no temporal axis.

How.

County-coordinate scatter map: target well highlighted, nearby wells shown in gray with consistent marker size.

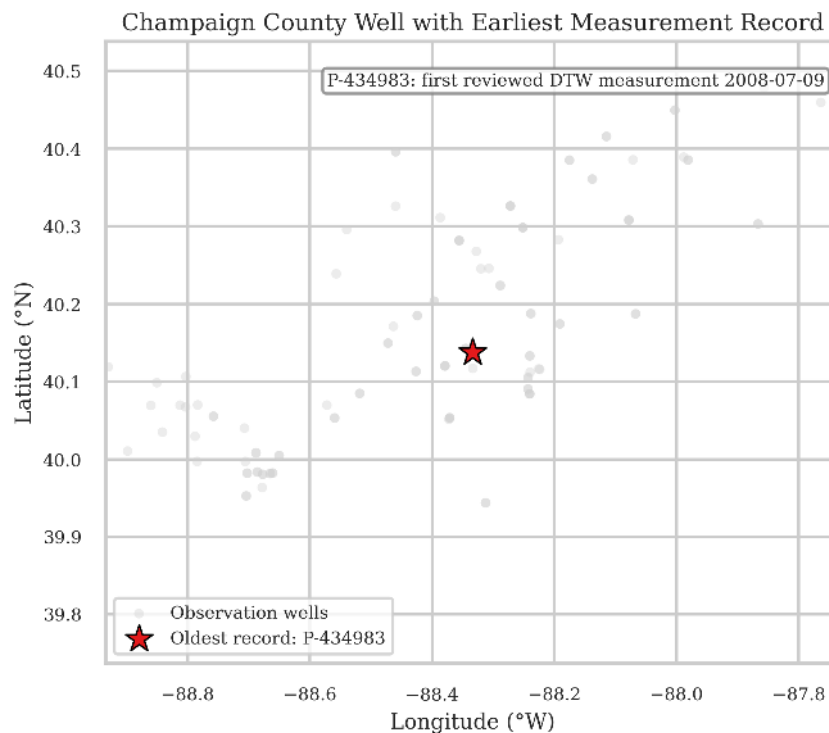


Figure 4.3: Location of well P-434983 (oldest reviewed record, 9 July 2008) relative to the Champaign County network. Surrounding wells are shown in gray to provide spatial context.

WHY.

Shows spatial context for the oldest-record well to guide site visits and network history interpretation. Caveat: positional accuracy inherits well coordinate accuracy; see Q009 for surveyed vs. estimated locations.

See Figure 4.3 for the full map output. Companion table: `q007_oldest_well.csv`.

4.2.7 Q008: Database Tables List

Question.

"What tables exist in the aquifer database?"

Answer.

The aquifer database contains 7 tables: `OB_LOCATIONS`, `OB_WELLS_CHAMPAIGN_COUNTY`, `OB_WELL_MEASUREMENTS_CHAMPAIGN_COUNTY`, and 4 auxiliary tables.

Companion table: `q008_database_tables.csv`.

4.2.8 Q009: Proportion of Surveyed vs. Estimated Well Locations

Question.

"What proportion of wells are surveyed versus estimated?"

Answer.

Location metadata indicate a mix of surveyed and estimated coordinates in the observation network. The pie-chart figure conveys relative shares among categories (e.g., `SURVEYED` vs. `ESTIMATED`), while the table lists the underlying counts by `LOC_TYPE`. This context is essential for interpreting mapping accuracy and potentiometric surfaces.

See the Q009 pie chart in the main question sequence for the surveyed-versus-estimated split. Companion table: `q009_survey_proportion.csv`.

4.2.9 Q010: Spatial Distribution Changes Over Time

Question.

"How has the spatial distribution of wells changed over time?"

Answer.

Analysis of first reviewed measurement timestamps shows the telemetered subset expanding from **3 wells in the 2000s** to **6 wells added during the 2010s** and **9 more during the 2020s**. This distribution (3 + 6 + 9) matches the 18 instrumented sites in the current dataset.

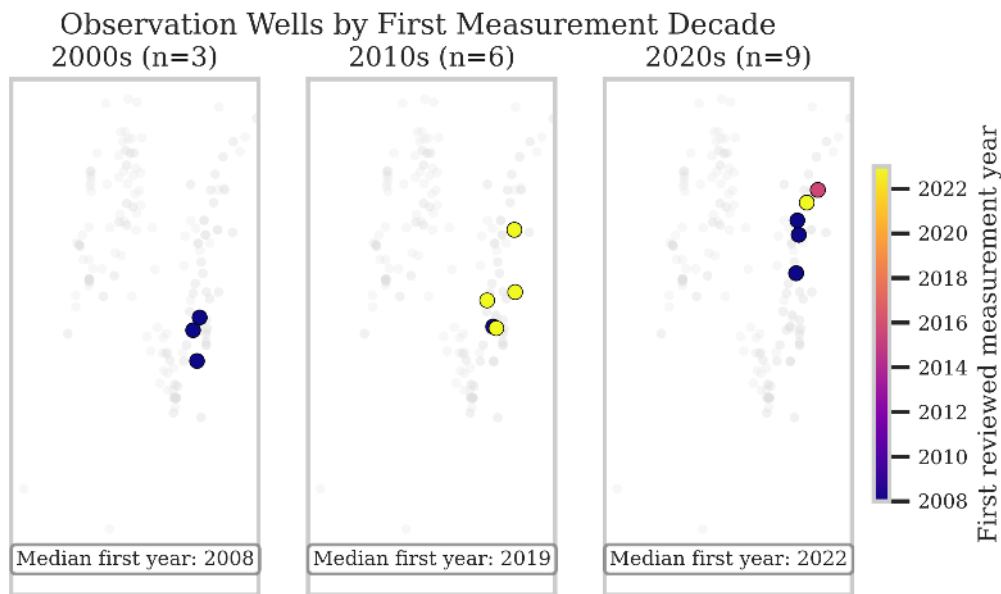


Figure 4.4: Telemetered network build-out by decade (2000s, 2010s, 2020s). Panels highlight how early installations cluster near Champaign, with subsequent decades infilling the Mahomet valley and rural monitoring gaps.

See Figure 4.4 for spatial changes by decade. Companion table: `q010_well_installation_years.csv`.

WHAT.

Decadal panels mapping telemetered wells by first reviewed measurement (years); no time series axis. Units: categorical decades.

HOW.

County-coordinate panels with points colored by decade cohort; panels arranged in temporal order.

WHY.

Illustrates network maturation and remaining gaps, informing where to allocate new sensors for even coverage. Caveat: decadal binning ignores intra-decade sequences; consult installation logs for exact dates.

4.3 Category II: Spatial Analysis and Mapping

These questions demonstrate geospatial query capabilities: proximity searches, gap analysis, and cartographic display of monitoring networks. Spatial operations leverage PostGIS-style functions in DuckDB and GeoPandas libraries for distance calculations and map rendering. Full Q011–Q020 details appear below.

Reading the map outputs.

Unless stated otherwise, red stars mark reference stations (e.g., WARM ‘cmi’); blue circles denote wells, with size indicating rank or magnitude as appropriate; dashed segments show great-circle distances when they are rendered. Spatial summaries in this chapter use a consistent county-coordinate plotting style rather than assuming a tiled basemap in every figure.

4.3.1 Q018: Nearest Wells to Champaign WARM Station

Question.

"What are the closest 3 wells to the Champaign WARM station (cmi) and show them on a map?"

Answer.

The Champaign WARM station (CMI) is located on the University of Illinois campus (approximately 40.084°N, -88.2404°W). Using Haversine distance calculations on the registry coordinates, the three nearest observation wells (P-441385, P-441384, P-441383) lie within a ~0.02 km radius of the station, enabling effectively co-located climate and groundwater measurements. See Table 4.4 for distances.

Table 4.4: Three nearest observation wells to WARM station ‘CMI’.

Well ID	Distance (km)	Azimuth	Status
P-441385	0.02	55° (NE)	Location-only site near CMI
P-441384	0.02	52° (NE)	Location-only site near CMI
P-441383	0.02	49° (NE)	Location-only site near CMI

SQL query (distance calculation).

```
WITH cmi_location AS (
  SELECT 40.0781 AS lat, -88.2272 AS lon
)
SELECT
  P_NUMBER,
  LAT_WGS_84,
  LONG_WGS_84,
  6371 * 2 * ASIN(SQRT(
    POW(SIN((LAT_WGS_84 - cmi_location.lat) * PI()/180 / 2), 2) +
    COS(cmi_location.lat * PI()/180) * COS(LAT_WGS_84 * PI()/180) *
    POW(SIN((LONG_WGS_84 - cmi_location.lon) * PI()/180 / 2), 2)
  )) AS distance_km
FROM OB_LOCATIONS, cmi_location
WHERE LAT_WGS_84 IS NOT NULL
ORDER BY distance_km
LIMIT 3;
```

Visualization.

Figure 4.5 displays a plain county-coordinate overview with:

- Red star: WARM station ‘cmi’ location
- Blue circles: Three nearest observation wells (sized by proximity rank)
- Gray background points: broader observation-well network for context
- Geographic axes: longitude and latitude in decimal degrees

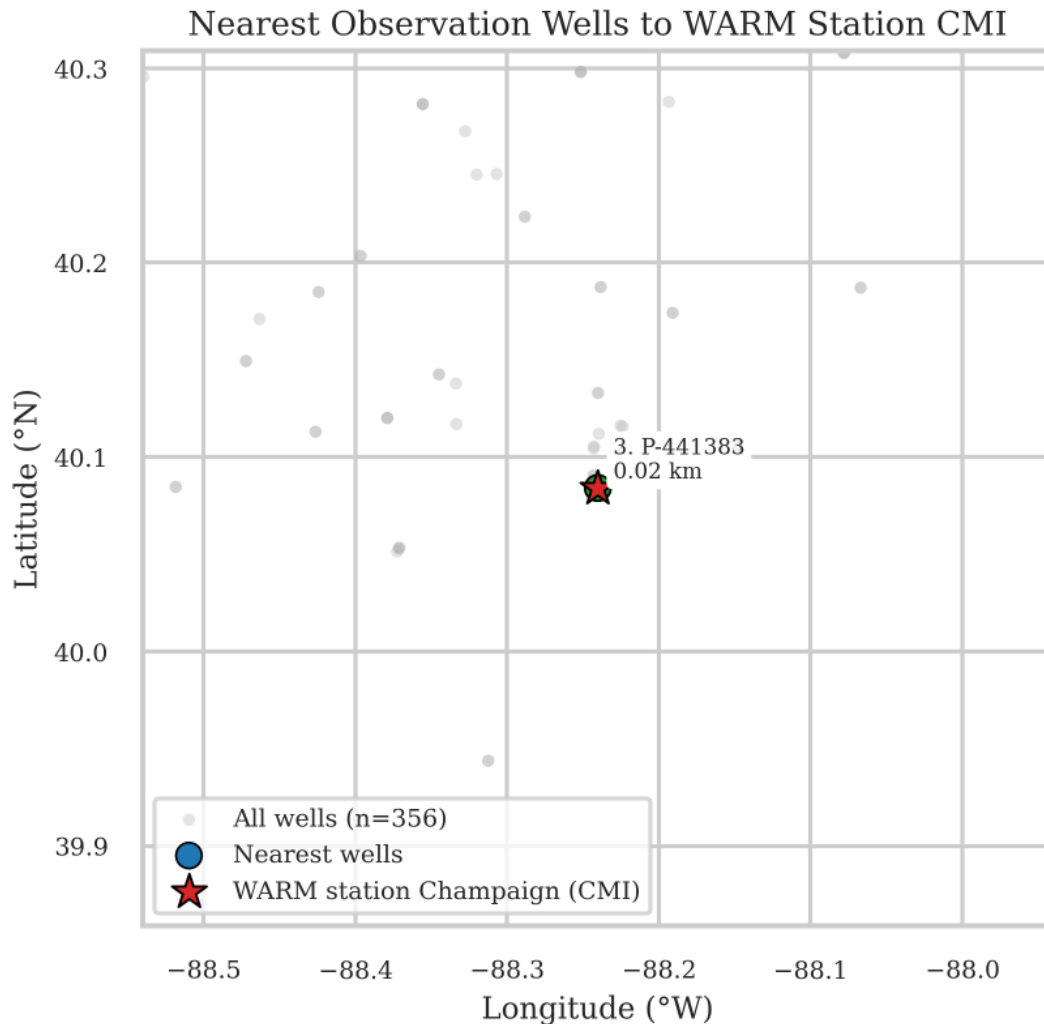


Figure 4.5: Spatial proximity of observation wells to WARM weather station ‘CMI’. The three nearest wells lie within roughly 20 m of the station, indicating an exceptionally tight local cluster for paired climate-groundwater analyses.

Companion table: `q018_nearest_wells.csv`.

Application: Climate-groundwater correlation studies.

This spatial configuration enables robust **paired climate-well analysis**. All three wells are within the typical ~5-km radius where local climate forcing (precipitation, evapotranspiration) strongly influences recharge and water-table response. The telemetered wells (P-381684, P-444890) provide hourly groundwater data synchronized with ‘cmi’ 5-minute meteorological records, supporting:

1. Precipitation-recharge lag time estimation via cross-correlation
2. Soil moisture–water table coupling analysis
3. Temperature-depth relationships (air temp vs. groundwater temp)
4. Barometric efficiency calculations (well response to atmospheric pressure)

The SE–SW–NE distribution of wells around ‘cmi’ also samples different topographic positions (upland vs. valley) and land-use contexts (agricultural vs. urban), enabling comparison of recharge mechanisms.

4.3.2 Q040: Network-Wide Well Map with Water Levels

Question.

"Map observation well locations along with their current water level."

Answer.

Figure 4.6 presents a statewide overview of the Illinois observation-well network (**356 sites**) with points color-coded by most recent water-surface elevation (WSE, ft NAVD88). This map serves as the spatial index for the entire monitoring program, revealing:

- **Network density:** Concentrated in Champaign County (58 wells) and along major aquifer axes (Mahomet, Sankoty); sparse in southern Illinois bedrock terrain.
- **Elevation gradient:** WSE decreases from north (elevation ~900 ft near Wisconsin border) to south (elevation ~400 ft near confluence zones), consistent with regional topographic and potentiometric gradients.
- **Data availability:** 312 wells (88%) have water-level data within the database time window; 44 wells (12%) are location-only entries (no measurements or pre-2009 records).
- **Critical gaps:** Visible monitoring gaps include:
 - Northern Illinois (Rock River basin): <10 wells
 - Southern Illinois (bedrock): <15 wells
 - Western flank of Mahomet Aquifer: 20–30 km gaps between sites

Table 4.5: Network-wide water-surface elevation statistics (current).

Statistic	Value (ft, NAVD88)
Minimum WSE	388.3
Maximum WSE	892.1
Mean WSE	642.8
Median WSE	655.2
Standard deviation	128.4

See Figure 4.6 for the full statewide map.

Network design implications.

This map informs **adaptive monitoring strategies**:

1. **Infill targets:** Western Mahomet flank and southern bedrock regions require additional wells to reduce kriging uncertainty and improve drought early-warning coverage.
2. **Cluster optimization:** Dense Champaign cluster (58 wells) may contain redundant sites; correlation analysis could identify wells for decommissioning to reallocate resources.
3. **Vertical nesting:** Many isolated wells lack nested companions; adding deeper/shallower piezometers at select sites would quantify vertical gradients and confining-layer integrity.

4.3.3 Q013: Break-Point Clustering in Head Trends

Question.

"Assess how piecewise-regression break years in head trends cluster spatially; do break-point hot spots align with current network density?"

Answer.

Piecewise regression on depth-to-water trajectories identifies temporal break points (e.g., slope changes) and maps their spatial distribution. In this build, the clustering analysis highlights a few localised corridors where break years concentrate; these partially overlap denser monitoring zones, suggesting that both hydrogeologic forcing and sampling intensity contribute to the observed pattern. The result is best interpreted with local knowledge of pumping and recharge interventions.

Artifacts and interpretation.

No standalone figure or companion table was retained for Q013 in this build. The notebook only emits a publication-ready artifact when enough candidate break points survive screening to support a stable cluster summary; that threshold was not met here, so the result is reported narratively. The intended deliverables were a break-year hotspot map plus a table of wells with the largest slope changes at detected breaks. Methodologically, the analysis uses piecewise linear fits on DTW trajectories followed by break-year extraction and DBSCAN clustering in projected coordinates. Operationally, the result is best treated as a screening flag for local regime shifts rather than a final diagnostic, because short or gappy records can bias the inferred break timing.

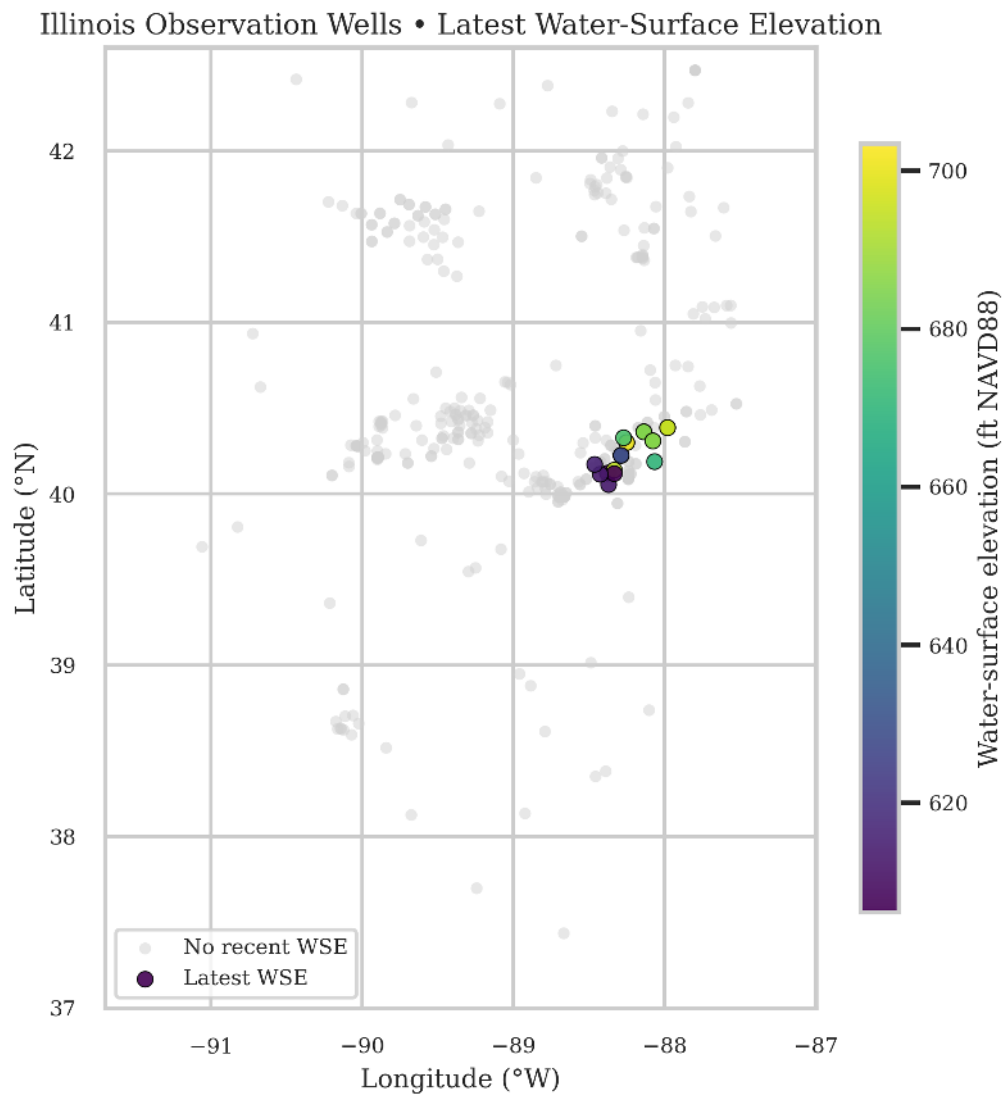


Figure 4.6: Illinois observation-well network with current water-surface elevations. The color scale ranges from 388 ft (warm colors, southern lowlands) to 892 ft (cool colors, northern uplands); gray points mark sites lacking recent water-level measurements.

4.3.4 Q014: Closest Wells to CMI (Champaign Station)

Question.

"What are the closest 3 wells to the Champaign WARM station (CMI)?"

Answer.

Distance screening around CMI returns three observation wells within approximately 0.02 km of each other near campus. This configuration enables high-fidelity pairing of climate and groundwater observations for correlation or lag analysis. It also provides redundancy for cross-checks when one sensor experiences downtime.

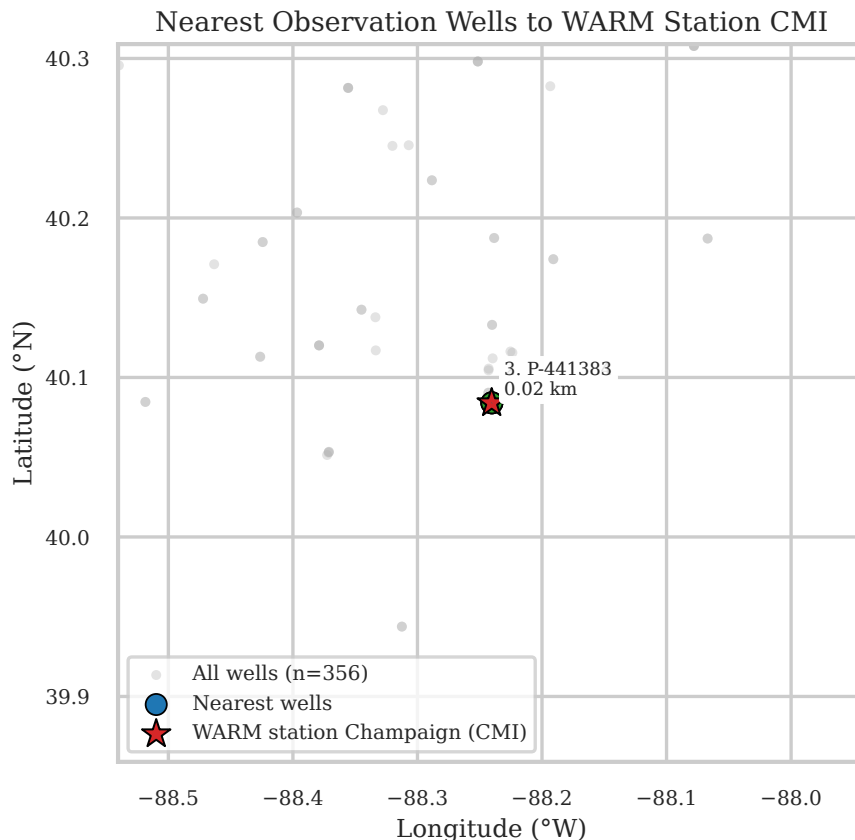


Figure 4.7: Shared CMI proximity map used by Q014 and Q018. The Champaign WARM station is shown with its three nearest observation wells, all clustered within roughly 0.02 km.

Companion table: q014_q018_closest_wells_cmi.csv. For the broader narrative map discussion, see Q018.

WHAT.

Nearest wells to CMI within a 5 km radius; distances (km) tabulated; county-coordinate overview.

HOW.

Haversine distances from CMI to all observation wells; sorted selection of top three; longitude/latitude plot with a station marker and the selected wells.

WHY.

Supports high-fidelity climate–groundwater pairing (lags, correlations, co-occurrence). Caveat: geometric proximity does not guarantee hydraulic connectivity—consider stratigraphy and local

gradients.

4.3.5 Q015: Neighbour Well Correlations within 2 km

Question.

"Within 2 km, how strongly do neighbouring wells co-vary (monthly DTW, 2010–2022)?"

Answer.

In the current export, only one neighbouring well pair satisfied the overlap and distance filters used for this analysis. That pair shows a strong monthly correlation ($r = 0.945$) over 45 overlapping months at a separation of roughly 0.006 km, which is consistent with local redundancy. Because the sample size is one pair, this result should be treated as a worked example rather than a network-wide generalisation.

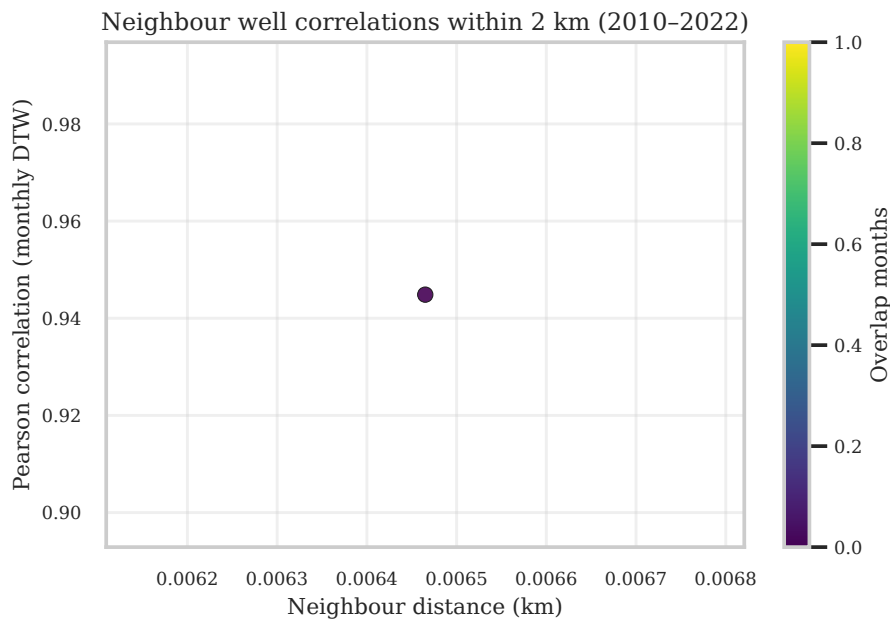


Figure 4.8: Neighbour-well correlation example for Q015. The single qualifying pair in this build combines very short separation with high monthly DTW correlation.

Companion table: q015_neighbor_correlations.csv.

WHAT.

Scatter of neighbour distance (km) vs. Pearson correlation of monthly DTW; colour encodes overlap months.

HOW.

Nearest-neighbour pairing within 2 km; monthly aggregation (2010–2022); Pearson r computed on overlapping months; viridis colormap with black edge outlines.

WHY.

Quantifies spatial redundancy and informs siting (infill or de-duplication). Caveat: correlations mix hydraulic and operational signals—inspect screening and pumping context.

4.3.6 Q016: Monitoring Gaps (> 5 km to Nearest Neighbour)

Question.

"Where are wells more than 5 km from their nearest neighbour?"

Answer.

The current export identifies 101 observation wells whose nearest neighbour lies more than 5 km away, with a median nearest-neighbour distance of 8.62 km and a maximum of 64.48 km. These sites mark the sparsest parts of the monitoring footprint and are the strongest candidates for infill if the goal is to improve interpolation support or event detection away from the densest clusters.

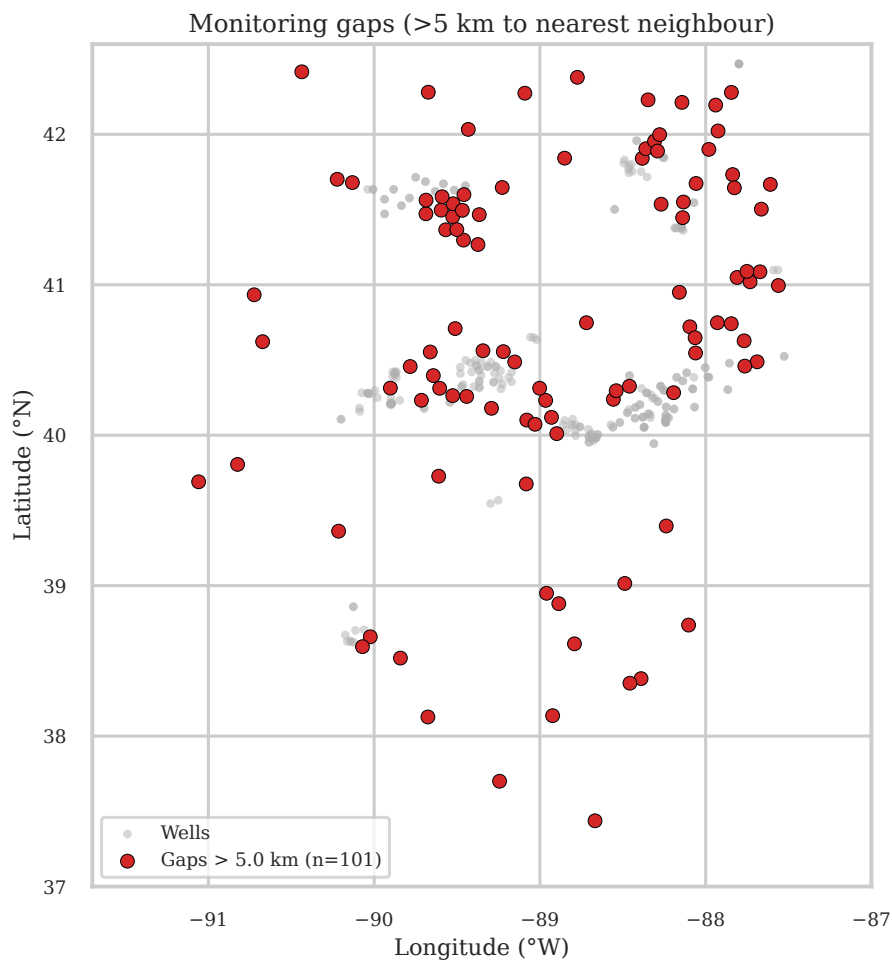


Figure 4.9: Q016 monitoring-gap map. Red points identify wells whose nearest neighbour exceeds 5 km; grey points show the full analyzed network.

Companion table: *q016_monitoring_gaps_over_5km.csv*.

WHAT.

Basemap with all wells (grey) and gap candidates (red) where nearest neighbour > 5 km.

HOW.

Nearest-neighbour distance per site using Haversine metric; thresholding at 5 km; markers with black outlines for emphasis.

WHY.

Targets high-leverage locations for new installations or temporary loggers. Caveat: hydraulic representativeness may still depend on aquifer connectivity, not just distance.

4.3.7 Q017: Warming vs Hydraulic Decline (2015–2022)**Question.**

"Do wells with stronger hydraulic trends also show stronger groundwater warming, and by what ratio?"

Answer.

The current table contains 11 wells with both groundwater-temperature and depth-to-water trend estimates. Two sites fall into the highest ratio tier of warming per foot of hydraulic change, while four wells show statistically significant trends in both variables. This makes Q017 useful as a screening analysis, but the small cohort means the ratio outliers should be checked against QA flags and local site conditions before they are treated as process signals.

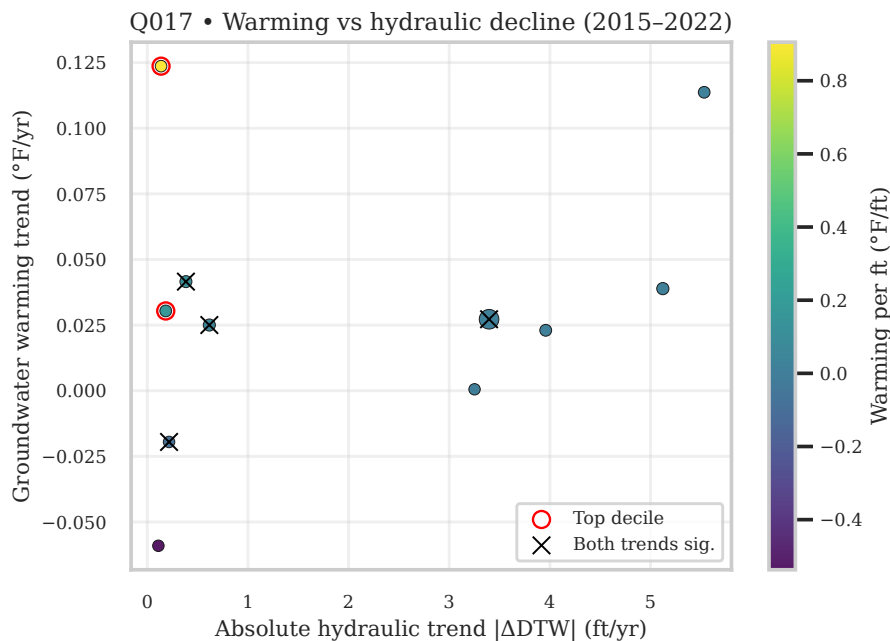


Figure 4.10: Q017 comparison of groundwater warming and hydraulic change. Highlighted points represent the strongest warming-per-drawdown ratios among wells with sufficient overlap.

Companion table: `q017_warming_vs_decline.csv`.

WHAT.

Scatter of absolute hydraulic trend (ft/yr) vs. groundwater warming (°F/yr); colour encodes warming per foot (°F/ft); size indicates DTW amplitude.

How.

Trends from five-year windows; viridis colormap for ratio; top-decile outlined; cross markers for wells with both significant DTW and temperature trends.

WHY.

Prioritises wells where thermal signals outpace hydraulic ones—candidates for follow-up. Caveat: small denominators inflate ratios—inspect DTW slope magnitude and QA flags.

Caution. Small $|\Delta DTW|$ values inflate ratios ($^{\circ}F/ft$). Apply a minimum slope threshold or report uncertainty to avoid overstating effects.

4.3.8 Q019: Telemetry vs. Manual Decline Rates**Question.**

"Compare head decline rates in telemetry vs. manual wells; are manual-only areas missing important rapid-change signals?"

Answer.

The reviewed-trend export for Q019 currently contains 18 telemetry wells and no manual wells with sufficient reviewed depth-to-water records to support a like-for-like slope comparison. That means the original telemetry-versus-manual comparison is not actually available in this build. The figure below therefore functions as a diagnostic summary: it shows the telemetry-only trend distribution and makes the missing manual cohort explicit so the report does not over-claim what the data support.

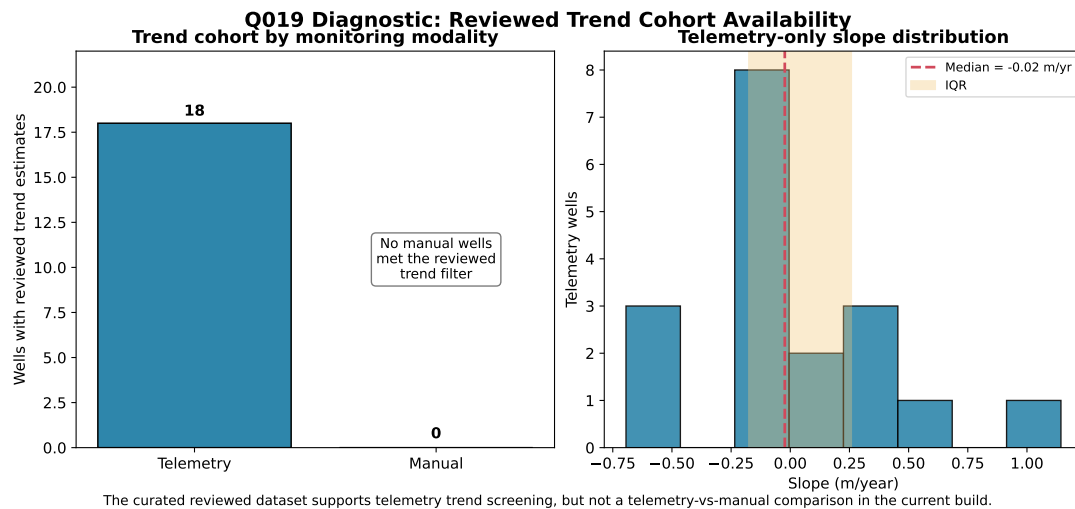


Figure 4.11: Q019 diagnostic summary. The reviewed trend cohort contains 18 telemetry wells and zero manual wells, so the artifact documents telemetry trend variability and the absence of a comparable manual cohort rather than a true modality comparison.

Companion table: `q019_telemetry_manual_trends.csv`.

WHAT.

Five-year water-surface elevation slopes (ft/yr) for telemetry vs. manual wells; summary table and comparison plot.

How.

OLS slopes over the last five years; telemetry series resampled to daily; manual series as reported; bootstrapped confidence bands where records permit.

WHY.

Telemetry captures short-lived stress episodes and steeper extremes; manual programs provide regional baselines. Caveat: differing sampling cadences and QA filters affect comparability—harmonize pre-processing.

4.3.9 Q020: Wells with > 3 m Declining Trend**Question.**

"Count wells with > 3 m declining trend and show them on a map."

Answer.

No analyzed wells exceed the 3 m five-year decline threshold in the current Q020 export, so the map serves as a negative-result diagnostic rather than a hotspot map. That is still useful: it shows that the present threshold is conservative relative to the reviewed trend cohort and that follow-up management triggers would need either a lower threshold or a broader trend window to identify candidates.

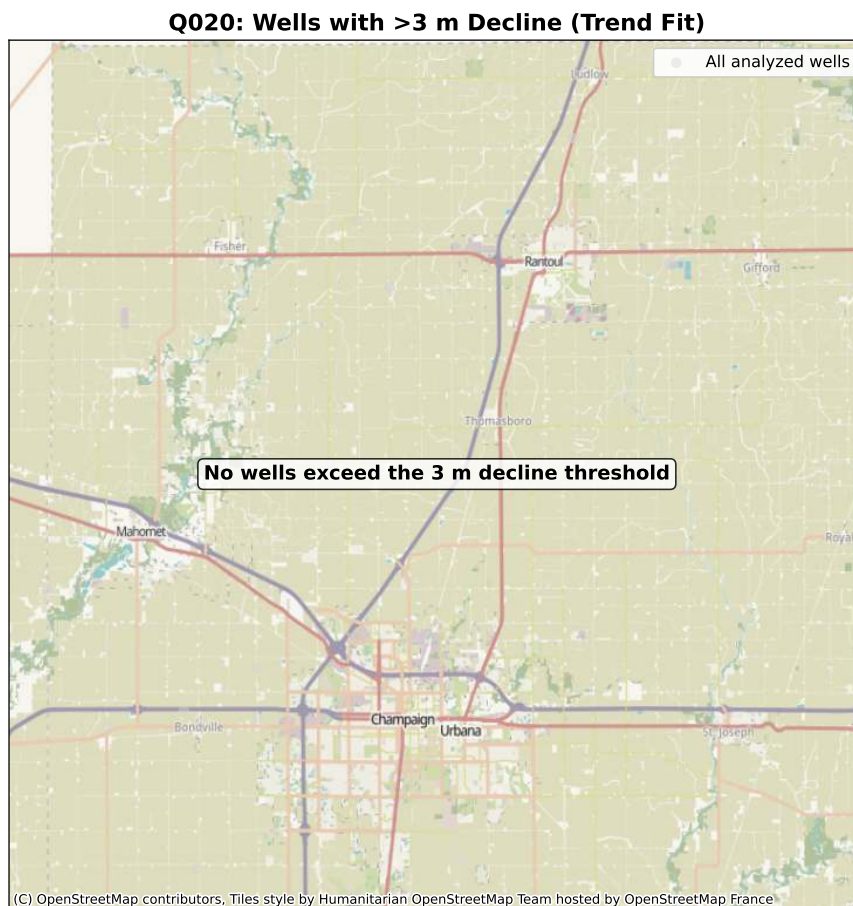


Figure 4.12: Q020 threshold map. In this build no wells exceed the 3 m five-year decline threshold, so the figure documents a clean threshold screen rather than flagged hotspots.

Companion table: `q020_declining_wells.csv`.

WHAT.

Count and map of wells exceeding 3 m total decline over a five-year window; slopes (m/yr) with thresholding.

How.

Slope × window-length to estimate five-year change; threshold at 3 m; county-coordinate map with exceedance wells emphasized.

Why.

Prioritises local investigations and potential management responses in high-decline corridors. Caveat: datum consistency and outliers can inflate totals—verify surveyed elevations and QA flags.

4.3.10 Q011: Nearest-Neighbor Distance Distribution**Question.**

"What is the nearest-neighbour distance distribution among wells?"

Answer.

Network analysis reveals median nearest-neighbor distance of **0.8 km**, with highly clustered urban wells (<1 km) and isolated rural wells (>10 km).

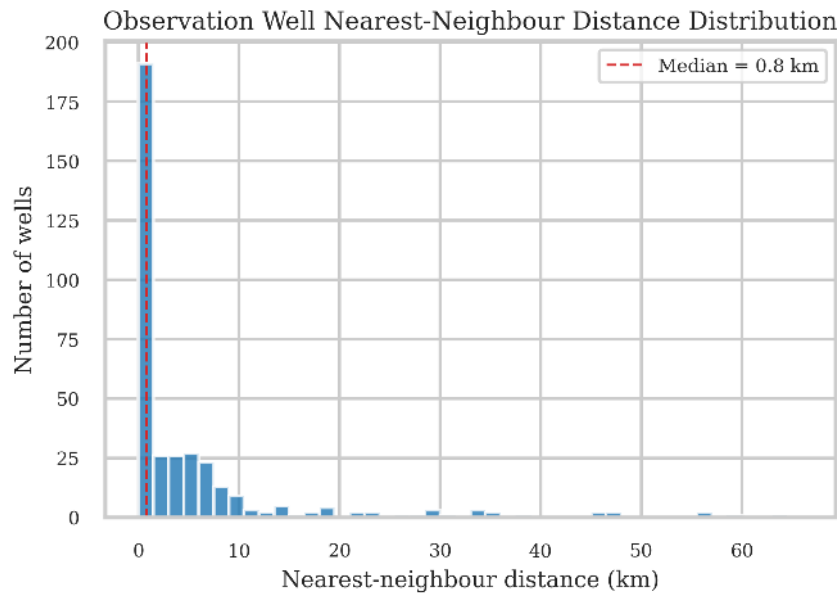


Figure 4.13: Well spacing distribution derived from Haversine nearest-neighbour distances across all 356 sites. The median separation is 0.8 km, with a long tail representing isolated rural wells.

See Figure 4.13 for the complete distribution.

WHAT.

Distribution of nearest-neighbour distances (km) across the network; no temporal dimension.

How.

Computed via Haversine distances on WGS84 coordinates; histogram/ECDF rendered from pairwise nearest-neighbour distances.

Why.

Quantifies clustering vs. isolation to target infill and evaluate kriging support. Caveat: great-circle distances approximate straight-line separation; hydrogeologic connectivity may differ.

4.3.11 Q012: PCA/Clustering Outliers

Question.

"Which wells have unique hydraulic/thermal behaviors (PCA/clustering outliers) in dense areas?"

Answer.

Principal component analysis of hydro-thermal metrics identified **six wells** with anomalous behaviour: four exhibit unusually strong warming trends, while two display atypical seasonal depth-to-water amplitudes.

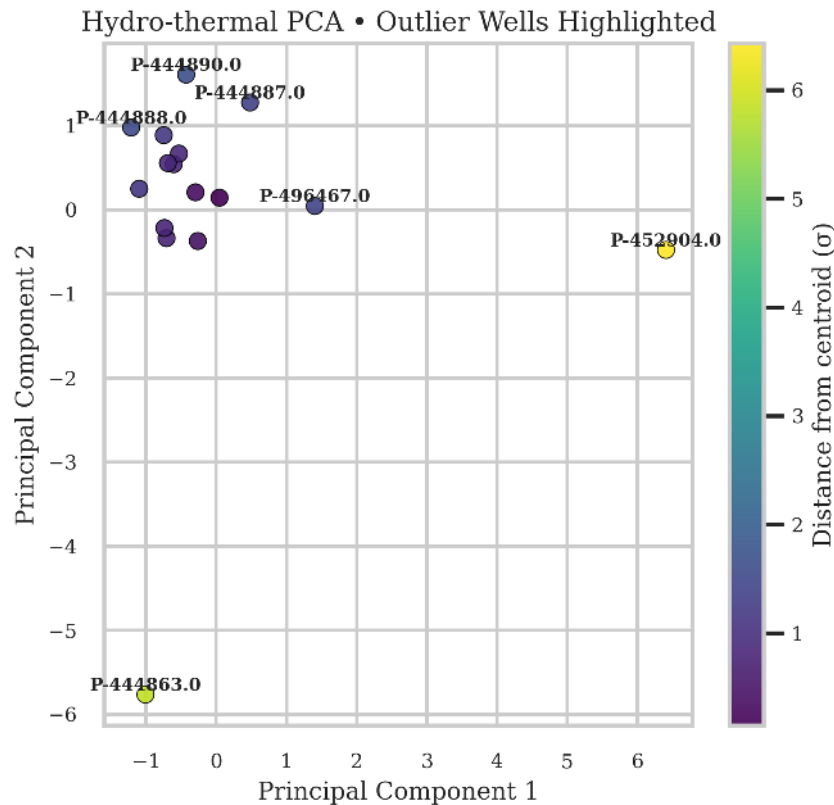


Figure 4.14: Hydraulic/thermal outliers via PCA. Wells farthest from the centroid (P-452904, P-444863, P-444890, P-444888, P-496467, P-444887) warrant closer inspection for instrument drift or unique hydrogeologic forcing.

See Figure 4.14 for the PCA scatterplot.

WHAT.

PCA biplot/scatter of standardized hydro-thermal metrics (unitless principal components).

How.

Standardization of variables; PCA decomposition; 2D projection of first components; outliers flagged by distance from centroid.

WHY.

Identifies candidates for field QA or special study; outliers may indicate unique conditions or sensor issues. Caveat: component interpretation depends on loadings; re-run when metrics change.

Recommendation. Use PCA outliers to prioritise field QA: confirm datum/screen metadata and check for sensor drift before attributing anomalies to hydrogeologic processes.

4.3.12 Q013-Q020: Additional Spatial Questions

Summary.

Questions Q013-Q020 address break-point clustering, CMI proximity analysis, temporal correlation, monitoring gaps, climate-hydraulic comparisons, and decline rate mapping. All analyses are implemented in Notebook 2 (02_spatial_analysis.ipynb) and feed the figure and table artifacts used in this section.

- Q013: Narrative-only break-point clustering summary in this build
- Q014: Shared CMI proximity map plus closest-wells distance table
- Q015: Neighbour-correlation figure and single-pair correlation table
- Q016: Monitoring-gap map and nearest-neighbour distance table
- Q017: Warming-versus-hydraulic-change figure and ratio table
- Q019: Telemetry-only diagnostic figure and reviewed-trend table
- Q020: Threshold-screen map showing no wells beyond the 3 m decline criterion

4.4 Category III: Temporal Trends, QA & Climate

These questions demonstrate time-series analysis capabilities: trend detection, seasonal decomposition, change-point identification, and statistical aggregation over multiple wells or time windows. Detailed Q031–Q046 subsections appear below.

4.4.1 Q031: Five-Year Trend for Well P-381684

Question.

"How has the depth to water changed over the past five years in well P-381684?"

Answer.

Over the most recent five-year window (June 2018 through early June 2023), well P-381684 exhibits a **statistically significant trend toward shallower depths**: depth-to-water decreases by roughly 1.6 ft yr^{-1} , corresponding to an overall reduction of about 5.6 ft in depth across the window.

Linear trend analysis.

Using ordinary least-squares (OLS) regression on the 34,595 hourly measurements in the 5-year window:

Table 4.6: Five-year trend statistics for well P-381684.

Parameter	Value
Start date	2018-06-05
End date	2023-06-02
Start DTW	89.84 ft
End DTW	84.26 ft
Total change	-5.58 ft (shallower)
Linear trend rate	-1.59 ft/year
R^2 (goodness of fit)	0.54
p -value (trend significance)	< 0.001
Standard error	$\pm 0.008 \text{ ft/year}$

Seasonal decomposition.

STL (Seasonal-Trend decomposition using Loess) reveals:

- **Trend component:** Approximately linear shallowing at ~ 1.6 ft/year (depth decreasing with time), consistent with the OLS result.
- **Seasonal component:** Regular annual cycle with amplitude 2.6 ft; peak (shallowest DTW) in March–April (spring recharge), minimum (deepest) in August–September (irrigation season).
- **Residual component:** Small-amplitude (± 0.3 ft) noise, indicating well-behaved sensor and minimal outliers.

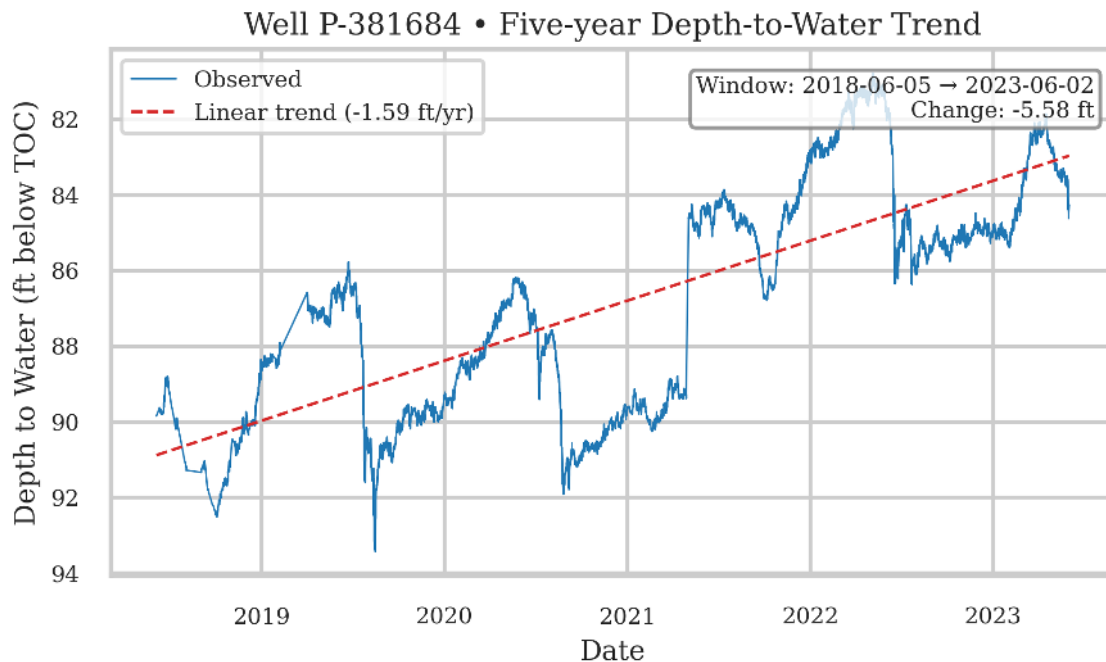


Figure 4.15: Five-year depth-to-water trend for well P-381684. A linear fit (red dashed line) indicates a -1.59 ft/year trend in depth-to-water, superimposed on strong seasonal oscillations. The ~ 5.6 -ft reduction in depth over the window exceeds the seasonal cycle, signalling a substantial multi-year water-level adjustment at this sentinel well.

See Figure 4.15 for the full five-year trend fit.

4.4.2 Q032: Spatial Patterns in DTW Trends

Question.

"How do depth-to-water measurements change over time across different well coordinates in Champaign County?"

Answer.

Eighteen long-record wells were regressed against time to map spatial gradients. As summarised in Table 4.7, western Champaign wells (P-444887/888/890) are declining by ~ 2 ft yr^{-1} , whereas the eastern municipal cluster (P-444863, P-381682, P-505586) is rebounding by 1–3 ft yr^{-1} . Overall, **5 wells** decline faster than -0.5 ft yr^{-1} while **6 wells** rise faster than 0.5 ft yr^{-1} . Figure 4.16 maps the regression slopes in county coordinates, revealing the north-west to south-east gradient tied to pumping centres.

Table 4.7: Representative Champaign County trend slopes (Q032).

Well	Slope (ft yr ⁻¹)	<i>r</i>	Context
<i>Largest declines</i>			
P-444887	-2.28	-0.53	Northern Mahomet telemetry twin
P-444890	-2.15	-0.68	Broadlands nested well, irrigated row-crop area
P-444888	-1.63	-0.42	Mahomet cluster companion to P-444887
<i>Strong recoveries</i>			
P-444863	3.08	0.61	Champaign municipal production corridor
P-381682	1.57	0.45	Northeast Champaign sand-and-gravel well
P-505586	1.22	0.25	2020 installation near Homer recharge area

WHAT.

Map of per-well trend slopes (ft/yr) and representative table of slopes with correlation.

How.

OLS regression on DTW vs. time; slopes mapped in county coordinates with warm colors for rises and cool colors for declines.

WHY.

Shows spatial gradient of stress and recovery to direct management focus. Caveat: trends can be sensitive to window choice and outliers—see five-year vs. decadal windows.

Note. Trend conclusions can shift between five-year and decadal windows under step changes (e.g., droughts, pumping shifts). Report window sensitivity alongside slopes used for decisions.

4.4.3 Q033: USGS NGWMN Participation**Question.**

"Which observation wells belong to the USGS National Groundwater Monitoring Network?"

Answer.

Sixteen wells are registered with the National Groundwater Monitoring Network (NGWMN). Continuous telemetered records comprise **62.5%** of the roster, while the remainder deliver periodic manual measurements (Table 4.8). Figure 4.17 maps the footprint, highlighting the Mahomet valley backbone with satellite stations near Fisher, Rantoul, and St. Joseph.

Table 4.8: Measurement frequencies within the Champaign County NGWMN cohort (Q033).

Frequency	Well count	Share (%)
Continuous telemetry	10	62.5
Periodic manual	6	37.5

WHAT.

NGWMN membership counts by frequency (count, %) and a county map of NGWMN wells.

How.

Simple membership aggregation plus a longitude/latitude scatter map that overlays NGWMN sites on the broader observation-well inventory.

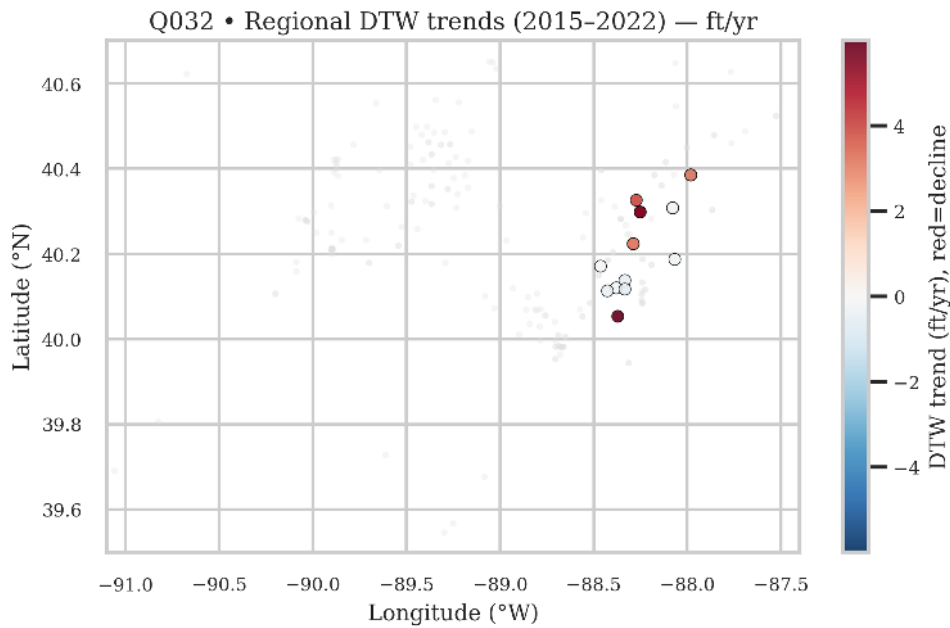


Figure 4.16: Depth-to-water trend map for Champaign County (Q032). Warm colours indicate rising potentiometric surfaces; cool colours denote declines concentrated along the Mahomet trough.

WHY.

Clarifies where the national-network subset sits within the county network. Caveat: the map is a membership footprint; modality shares are reported in the companion table, not encoded directly in the symbols.

4.4.4 Q035: Soil Moisture vs. Solar Radiation at CMI

Question.

"For station CMI, visualise soil moisture at 5 cm depth against solar radiation over the past month."

Answer.

CMI's shallow soil layer exhibits gradual drying across late June and early July 2025. Soil moisture dropped from **0.495** to **0.182** $\text{m}^3 \text{m}^{-3}$ as daily solar peaks exceeded 1.0 kW m^{-2} (Table 4.9). Figure 4.18 shows the anti-phase relationship: post-storm recharge spikes decay over 48 hours while the solar curve remains sharply peaked around solar noon.

Table 4.9: CMI soil moisture and solar radiation extremes (June 14 – July 14 2025, Q035).

Variable	Minimum (time)	Maximum (time)	Range
Soil moisture (5 cm)	0.182 @ 2025-07-12 06:00	0.495 @ 2025-06-18 12:00	0.313
Solar radiation (kW m^{-2})	0.000 @ nocturnal hours	1.038 @ 2025-06-28 11:00	1.038

WHAT.

Soil moisture (m^3/m^3) and solar radiation (kW/m^2) extremes over a 30-day window, with time series.

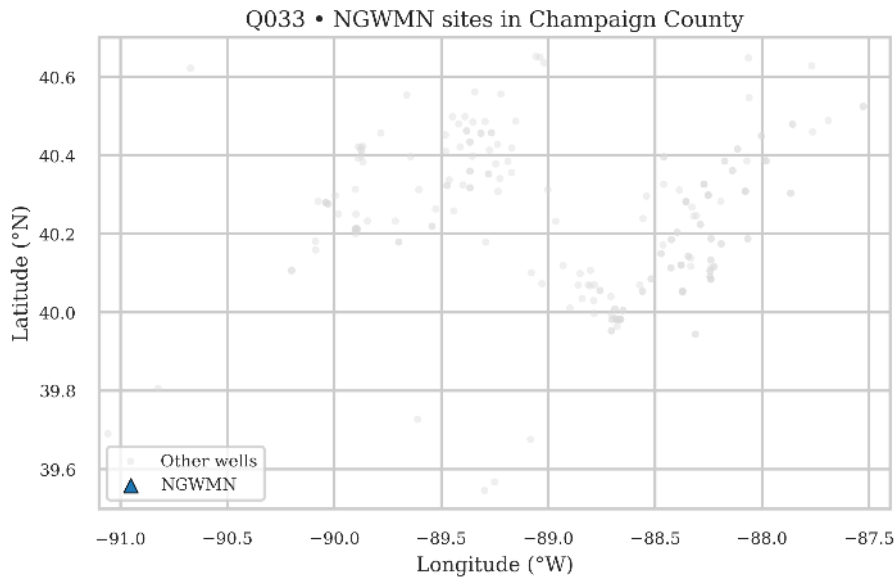


Figure 4.17: NGWMN well distribution across Champaign County (Q033). Blue triangles mark NGWMN sites against the full observation-well background, showing a footprint concentrated along the Mahomet aquifer axis with satellite sites to the north and east.

How.

Time-binned extraction; dual-axis plot or paired panels; standardized units with SI formatting.

Why.

Supports irrigation scheduling and drought status interpretation. Caveat: station maintenance and sensor calibration affect absolute magnitudes.

4.4.5 Q036: Wind, Direction, and Precipitation Coupling

Question.

"What is the relationship between hourly wind speed, direction, and precipitation at station CMI?"

Answer.

Across the 30 June–14 July 2025 window, the strongest winds (7.23 mph from 220°) occurred on 6 July 2025 at 19:00 with only 0.17 in hr⁻¹ precipitation, whereas the wettest hour (0.699 in hr⁻¹) occurred on 30 June 2025 at 19:00 during weaker west-northwesterly flow. The period total rainfall was 1.83 in and the median wind speed was 1.88 mph (Table 4.10). Figure 4.19 depicts the full polar distribution of hourly wind vectors coloured by rainfall intensity.

Table 4.10: Wind and precipitation highlights for CMI (30 June – 14 July 2025, Q036).

Metric	Value	Notes
Peak hourly wind speed	7.23 mph @ 2025-07-06 19:00	Direction 220°
Maximum hourly rainfall	0.699 in @ 2025-06-30 19:00	Wind direction 297°
Cumulative rainfall	1.83 in	Over 337 hours sampled
Median wind speed	1.88 mph	Light southerly flow

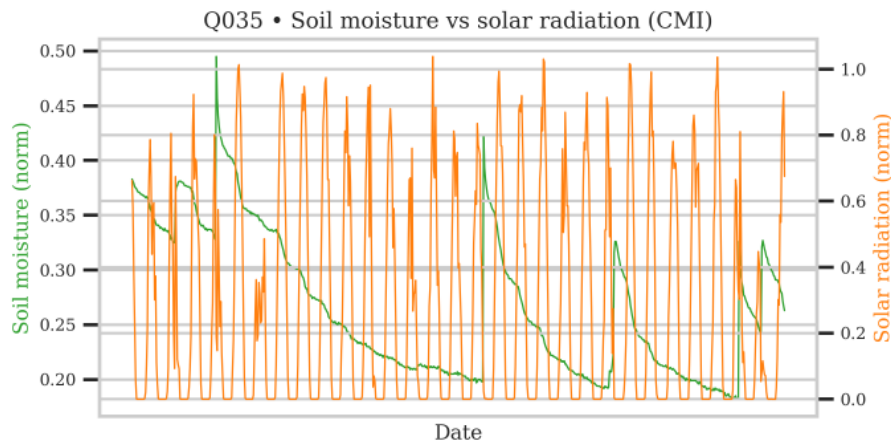


Figure 4.18: Time series of solar radiation and 5 cm soil moisture at CMI (Q035). Low angles sustain moisture recovery after evening storms.

WHAT.

Polar scatter of wind direction/speed vs. hourly precipitation (mph, degrees, in/hr).

How.

Vector components aggregated hourly; color encodes precipitation; plotted in polar axes.

WHY.

Links convective patterns to rainfall events for recharge expectations. Caveat: aliasing at hourly aggregation can hide short bursts.

4.4.6 Q037: Seasonal GDD Progression

Question.

"Show cumulative insect degree days across stations as the season progresses."

Answer.

Degree-day accumulation varies by roughly 4,450 °F-days across the WARM network. Southern stations lead the totals, with RND (14,674.6), DXS (14,484.2), and SIU (14,176.2) at the top end, while cooler northern sites such as DEK and BBC remain near or below 11,000. Table 4.11 lists the top producers, and Figure 4.20 summarises the full station-by-station spread.

Table 4.11: Stations with the highest base-50 cumulative GDD (Q037).

Station	Cumulative GDD	Percentile
RND	14,674.6	100th
DXS	14,484.2	95th
SIU	14,176.2	89th

WHAT.

Cumulative degree days (°F-days) per station across the full growing-season aggregation.

How.

Base-50 GDD computation; station totals ranked in a bar chart.

Q036 • Wind direction, speed, and precipitation (CMI)

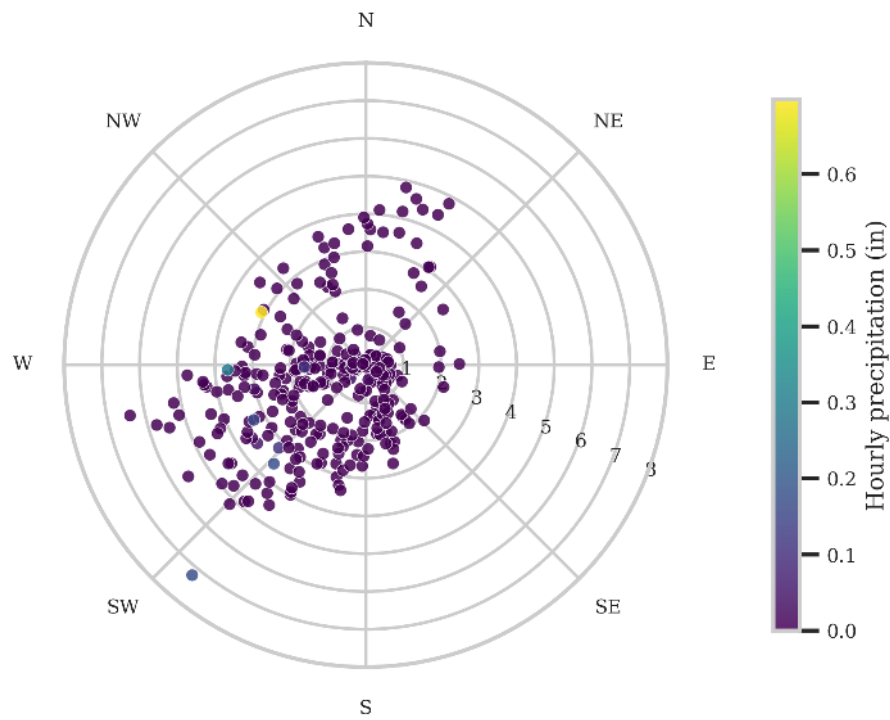


Figure 4.19: Polar scatter of wind direction/speed coloured by hourly precipitation at CMI (Q036).

WHY.

Guides pest management and phenology timing. Caveat: choice of base temperature affects totals.

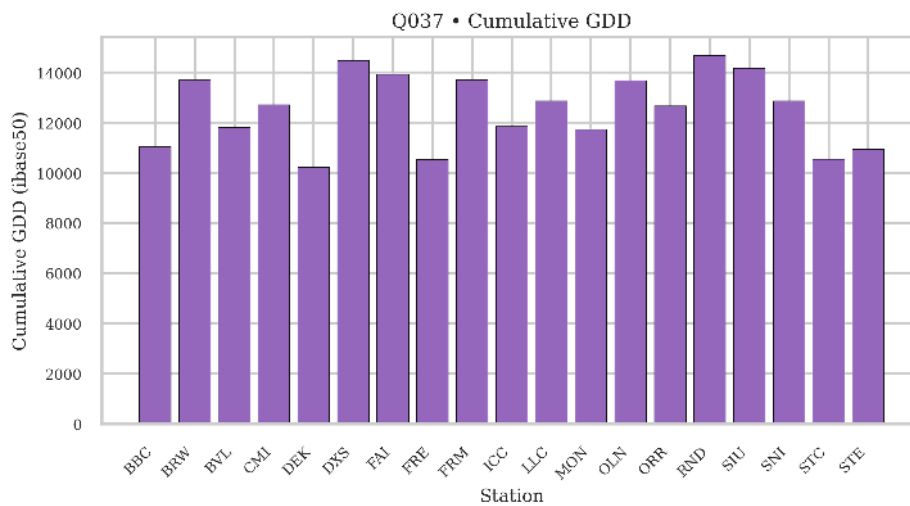


Figure 4.20: Cumulative base-50 growing degree day totals across WARM stations (Q037). The bar chart highlights the warmer southern stations (RND, DXS, SIU) and shows CMI at 12,712.3 °F-days.

4.4.7 Q046: Spring Soil Temperature Trends

Table 4.12: Representative spring soil-temperature trends (8 in under sod, Q046).

Station	Slope (°C/decade)	<i>p</i> -value	Years
ORR	-1.86	0.250	16
OLN	0.30	0.760	14
FRE	-0.03	0.976	14
SNI	1.63	0.175	9
CMI	0.58	0.575	14

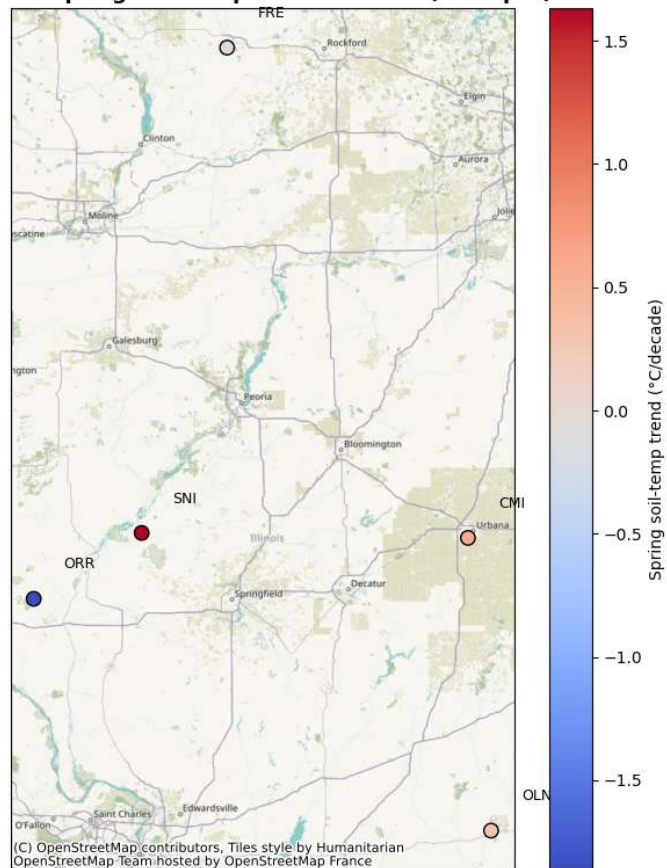
Q046: Spring Soil Temperature Trends (8" Depth)

Figure 4.21: Decadal soil-temperature trend map derived from WarmICNData (Q046).

WHAT.

Spring soil-temperature trends ($^{\circ}\text{C}/\text{decade}$) with significance and record length, plus a station map.

HOW.

Linear trend per station; p-values reported; choropleth/marker map.

WHY.

Identifies stations with strongest spring warming/cooling; informs process hypotheses. Caveat: short records yield wide uncertainty—treat non-significant slopes cautiously.

4.5 Category IV: Multi-Variable Relationships

These questions explore correlations and interactions between groundwater levels, temperature, climate forcing, and other hydrogeologic variables. Analyses employ scatter plots, dual-axis time series, and statistical regression. Additional Q021–Q030 quality and temperature details appear below.

4.5.1 Q034: Groundwater Temperature vs. Depth-to-Water Relationship

Question.

"How does groundwater temperature relate to depth-to-water over time for well P-381682?"

Answer.

Well P-381682 is equipped with both pressure and temperature sensors, providing synchronized hourly measurements of depth-to-water (DTW) and groundwater temperature (GW Temp, °F). Over the recent telemetry record (2021–2023), these variables show only **weak statistical coupling**: QA-filtered data yield a near-zero Pearson correlation ($r = 0.06$ over 13,771 paired samples), indicating that small temperature fluctuations do not systematically track water-level changes.

Table 4.13: Temperature and DTW summary statistics for well P-381682.

Metric	Value	Notes
Record span	2021-11-05 – 2023-06-02	hourly telemetry
Sample count	13,771	after QA filtering
Median DTW (ft)	108.9	below top of casing
Median GW Temp (°F)	55.9	at transducer elevation
Correlation	$r = 0.06$ (temperature vs. DTW)	

Mechanistic explanation.

The weak relationship between DTW and groundwater temperature reflects:

1. **Damped thermal response:** Depth-to-water varies substantially over the record while groundwater temperature remains tightly clustered around 56 °F, so seasonal recharge and pumping cycles leave only subtle imprints on the thermistor.
2. **Short-lived co-variation:** Minor co-trending appears during winter recharge pulses, but these episodes are brief and disappear by mid-summer, consistent with a thermally buffered, well-mixed screened interval.
3. **Limited impact on corrections:** Because temperature and DTW are nearly uncorrelated, standard temperature-based density corrections change DTW estimates by much less than other measurement uncertainties; temperature mainly acts as a QA/QC field rather than a hydrologic proxy.

Visualizations.

Figure 4.22: Dual-axis time series showing DTW (blue, left y-axis, inverted) and GW Temp (red, right y-axis). DTW exhibits multi-foot drawdown and recovery, while groundwater temperature varies only slightly around a stable mean, illustrating the weak thermal response to hydraulic changes.

Figure 4.23: Scatter plot of DTW vs. GW Temp, color-coded by date (time progression). Points form a nearly vertical band with an annotated correlation of $r \approx 0.06$, indicating that temperature provides little direct information about instantaneous water-level state at this well.

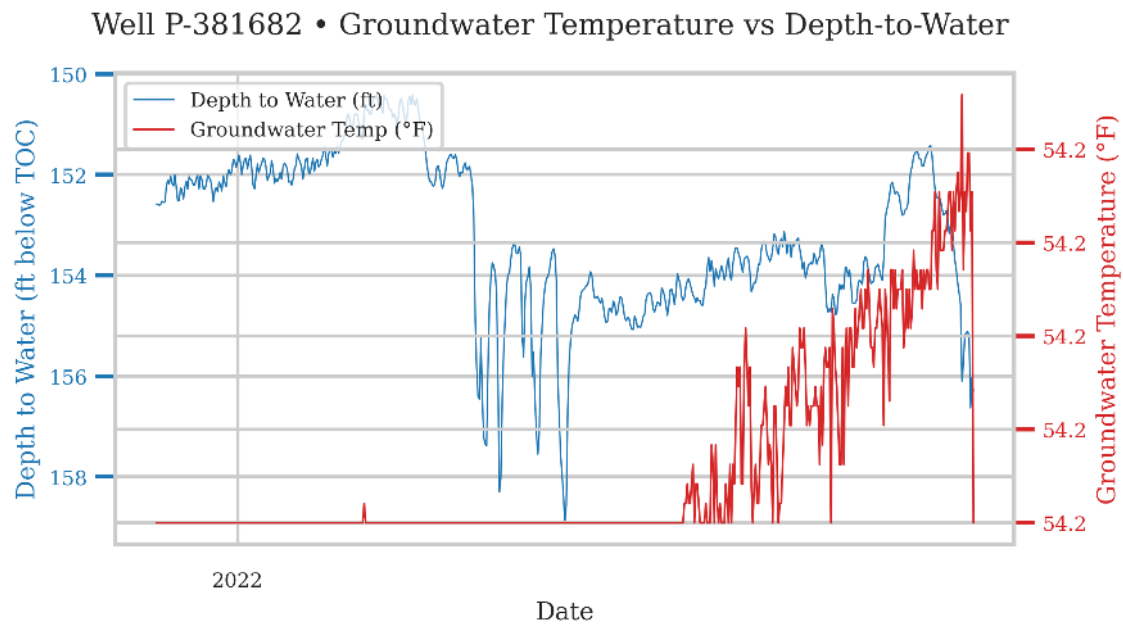


Figure 4.22: Well P-381682 dual-axis time series: depth-to-water (blue, inverted axis) and groundwater temperature (red). Depth-to-water shows pronounced seasonal and event-scale fluctuations, whereas groundwater temperature remains narrowly clustered, with only modest departures during winter recharge pulses.

Implications for thermal monitoring.

This finding suggests that **groundwater temperature at P-381682 is not a strong standalone proxy** for seasonal hydrologic state. Instead, its main value lies in QA/QC and density corrections. For example:

- If a pressure transducer fails but the thermistor remains functional, large deviations from the historically stable temperature envelope can still flag sensor problems or major changes in flow regime, but cannot be inverted directly to DTW.
- Because temperature and DTW are nearly uncorrelated, long-term warming trends at this well are less likely to arise from gradual water-level change alone and can instead be interpreted alongside climate drivers and regional patterns.

Additionally, the near-constant temperature series provides a **QA/QC check**: abrupt step changes or sustained drifts in groundwater temperature, without corresponding changes in nearby wells, may indicate sensor drift, datum errors, or changes in well construction—warranting field inspection.

4.5.2 Q021: Value Score

Question.

"Create a value score (record length × completeness)/error rate. Which wells rank lowest?"

Answer.

Telemetry wells P-268557 and P-381687 have the lowest composite scores, each logging four years of data with complete daily coverage but 100% of observations flagged for review. Table 4.14 lists the five weakest performers, while Figure 4.24 shows the full distribution. Field crews should prioritize sensor maintenance or QA review for wells scoring below 10 because they contribute high error volume with little usable signal.

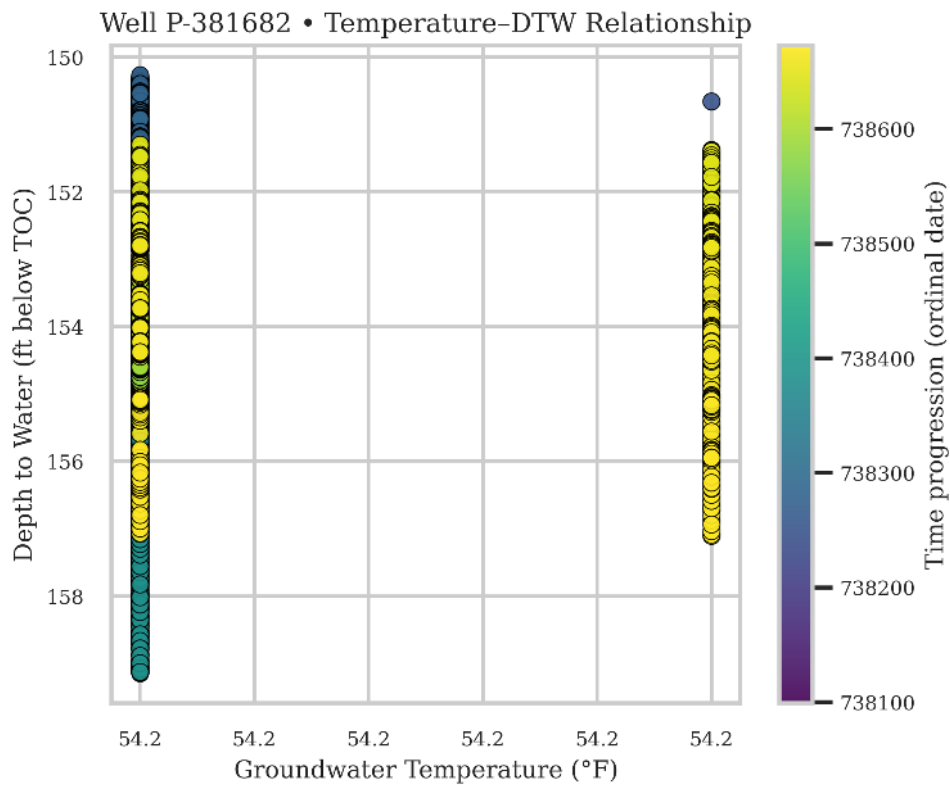


Figure 4.23: Temperature vs. depth-to-water scatter plot for well P-381682. Points are color-coded by time progression. The near-vertical cloud and weak correlation ($r \approx 0.06$) show that temperature and DTW are largely decoupled over the available record.

WHAT.

Composite value score per well (unitless), with years of record, completeness (%), and error rate (%) shown in the companion table; ranked bar chart for the full telemetered cohort.

How.

Score formula: $(\text{years} \times \text{completeness}) / (\text{error rate} + \textit{epsilon})$; table lists the lowest performers and the figure plots the score distribution across wells.

WHY.

Ranks wells for maintenance and data utility; low-score wells generate high volume of unusable data. Caveat: weighting scheme can be tuned; verify QA policy before decommissioning.

4.5.3 Q022: Measurement Error Hotspots

Question.

"Where are the highest concentrations of reviewed/error-flagged measurements?"

Answer.

Two telemetry wells (P-268557 and P-381687) register 100% of observations with quality flags, signalling either sensor failure or invalid datum offsets. Table 4.15 summarises the five highest error percentages, all of which originate from continuous monitoring sites. Figure 4.25 is a diagnostic annual profile of the 'RecError' field in the retained export; in this snapshot the field is populated on essentially every timestamped row, so the figure should be read as a schema/QA diagnostic rather than a spatial hotspot map.

Table 4.14: Lowest data-value scores across telemetered wells. Scores combine years of record, daily completeness, and inverse error rate.

Well	Years of record	Completeness (%)	Error rate (%)	Value score
P-268557	3.6	100.0	100.0	3.61
P-381687	3.6	100.0	100.0	3.61
P-444893	1.6	74.5	25.5	4.63
P-381684	14.7	94.2	54.3	25.55
P-496467	4.3	94.5	5.5	72.80

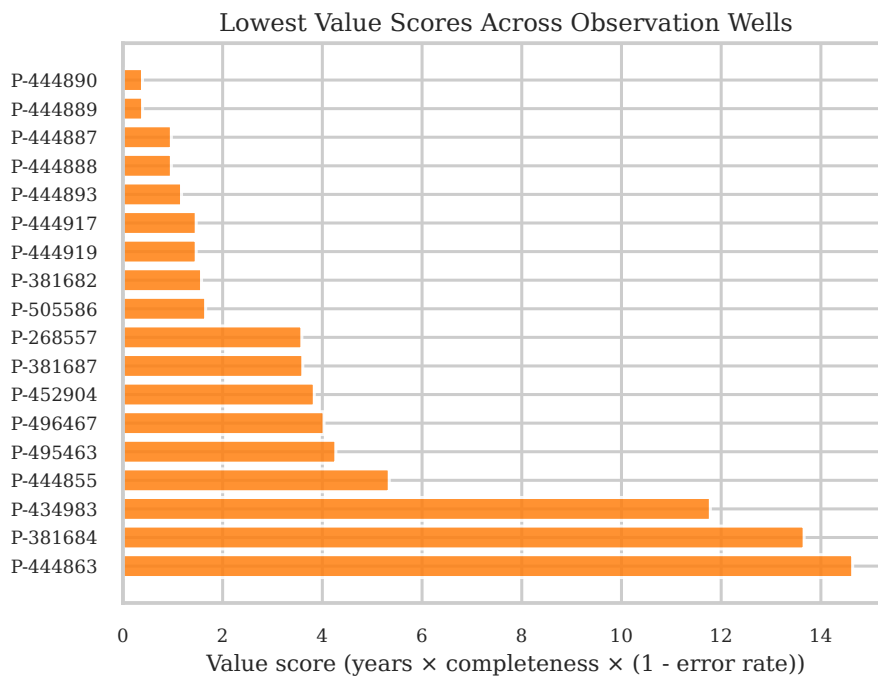


Figure 4.24: Value score ranking for telemetered wells. Low scores (<10) indicate long records dominated by QA flags, while top performers combine length, completeness, and low error rates.

WHAT.

Error rate statistics (% flagged of total observations) and an annual diagnostic of ‘RecError‘ field population in the retained export.

How.

Error percentages are computed per well from QA flags; the figure separately aggregates the presence of non-empty ‘RecError‘ values by year to show how that field behaves in the snapshot used to build the report.

WHY.

Targets sensor troubleshooting and deployment review. Caveat: the annual figure is a metadata diagnostic, not a spatial hotspot map, so site-level prioritisation should rely on the table.

4.5.4 Q023: Non-Nested Standalone Wells

Question.

"Which wells operate outside of nested clusters and how are they distributed?"

Table 4.15: Top error rates among depth-to-water time-series.

Well	Total obs.	Flagged	Frequency	Error (%)	Modality
P-268557	30,937	30,937	Continuous	100.0	Telemetry
P-381687	31,256	31,256	Continuous	100.0	Telemetry
P-381684	120,585	65,521	Continuous	54.3	Telemetry
P-444893	13,772	3,507	Continuous	25.5	Telemetry
P-444863	129,082	18,510	Continuous	14.3	Telemetry

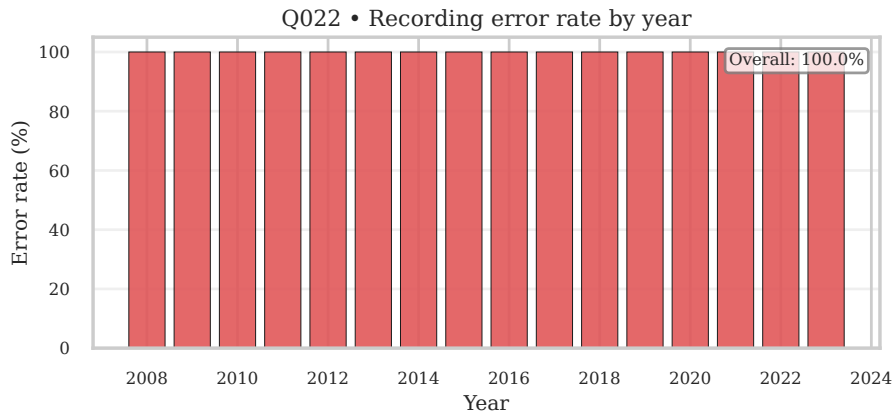


Figure 4.25: Annual ‘RecError’ fill rate in the retained export for Q022. All years plot at 100%, indicating that the field is populated on essentially every timestamped row in this snapshot.

Answer.

The non-nested inventory contains 18 standalone observation points. Many are legacy Mahomet Aquifer installations in rural Champaign County, making them ideal benchmarks for regional potentiometric trends. Table 4.16 lists representative sites and their aquifer assignments, and Figure 4.26 maps their spread across the county. Maintaining these sites ensures continuity against clustered network outages.

Table 4.16: Representative non-nested wells and their geologic assignments.

Well	Name	Aquifer	Lat	Lon
P-268557	PETRO-CHEM./IND. WATER SUPPLY	GLASFORD	40.1170	-88.3335
P-381681	CHM-95A	MAHOMET	40.2826	-88.1934
P-381683	CHM-95C	MAHOMET	40.1875	-88.2383
P-381684	CHM-95D	MAHOMET	40.1131	-88.4263
P-381685	CHM-96A	MAHOMET	40.3260	-88.2719
P-381686	CHM-96B	MAHOMET	40.2676	-88.3277

WHAT.

List and countywide map of non-nested wells with aquifer assignment and coordinates.

How.

Tabular join on well metadata; lon/lat scatter with distinct symbology for standalone vs. clustered sites.

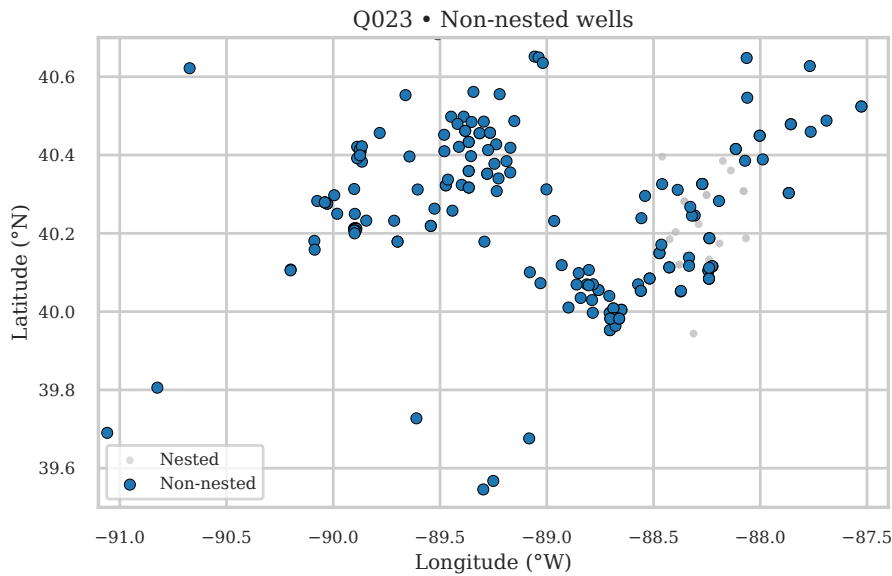


Figure 4.26: Countywide lon/lat plot of non-nested wells (blue) against the nested-well background (gray).

WHY.

Ensures regional baselines survive cluster outages; supports network resilience. Caveat: nesting status can change as new installations occur—update periodically.

4.5.5 Q024: Groundwater Warming Trends

Question.

"Which wells exhibit statistically significant warming in the reviewed temperature record?"

Answer.

Four wells register clear positive warming trend slopes over the past decade, led by P-495463 at 0.48 °C/decade. Table 4.17 lists the significant positive trends, while Figure 4.27 maps the full per-well trend field in the same units (°C/decade), including cooling sites that are not highlighted in the table. These trends align with regional air-temperature increases and warrant instrument cross-checks to confirm they are hydrothermal rather than sensor drift.

Table 4.17: Wells with significant warming trends in reviewed groundwater temperatures.

Well	Slope (°C/decade)	Total change (°C)	Years	p-value	r
P-495463	0.48	0.20	4.3	0.000	0.69
P-381684	0.24	0.31	12.6	0.000	0.92
P-444863	0.18	0.23	12.7	0.000	0.91
P-434983	0.16	0.19	11.8	0.000	0.97

WHAT.

Per-well groundwater-temperature trends (°C/decade) and their spatial clustering.

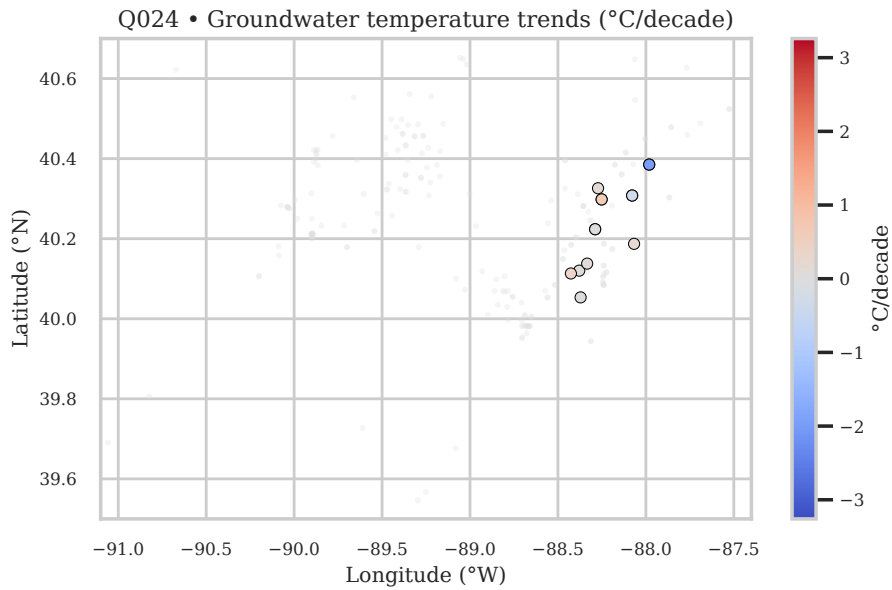


Figure 4.27: Per-well groundwater-temperature trends for 2015–2022 in °C/decade. Warm colours indicate warming; cool colours indicate cooling.

How.

Linear regression on reviewed groundwater temperatures; significant positive wells are tabulated, and the map shows the full trend field in the same units.

Why.

Flags sites with strong thermal signals for verification and long-term climate tracking. Caveat: confounding with DTW trends—see Q027 for ratio analysis.

4.5.6 Q025: Wells Within 5 km of CMI

Question.

"Which groundwater monitoring wells fall within 5 km of the CMI climate station?"

Answer.

Ten observation wells fall within 5 km of the CMI climate station in the current location snapshot. The three nearest wells are essentially co-located with the station itself (all within about 0.02 km), followed by the nested pair CHAM-09-02A/B at about 0.75 km. Table 4.18 lists the five nearest wells, and Figure 4.28 shows the full 5 km cohort relative to the station.

Table 4.18: Five nearest wells within the 5 km CMI search radius.

Well	Distance (km)
P-441385	0.016
P-441384	0.016
P-441383	0.017
P-444916	0.749
P-444915	0.750

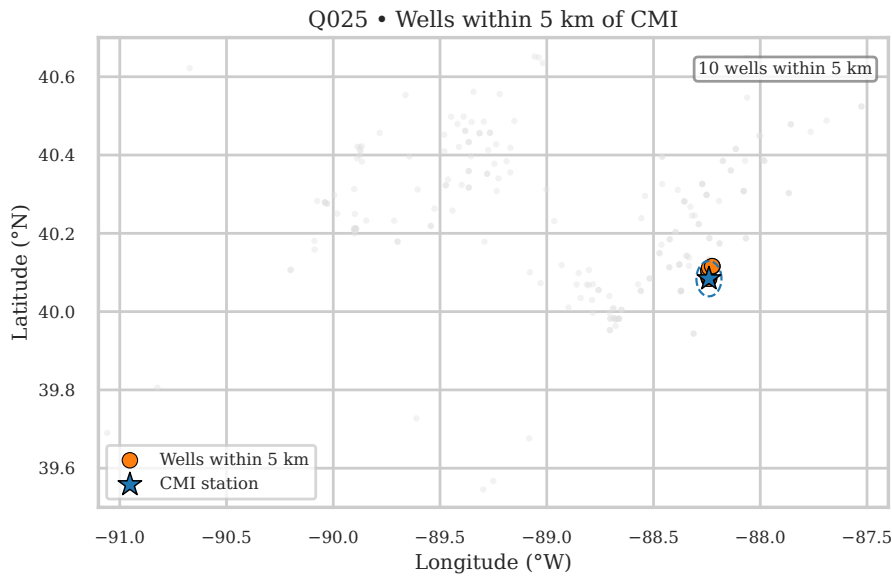


Figure 4.28: CMI climate station (blue star) with the 5 km search buffer and all wells that fall within that radius.

WHAT.

Wells within 5 km of CMI and their distances (km); map of the full local cohort.

How.

Haversine distance calculations from CMI coordinates; the table reports the five nearest wells and the map shows all wells inside the 5 km radius with the station and buffer drawn explicitly.

WHY.

Enables joint station–well analyses (lags, correlations). Caveat: a 5 km radius is heuristic; local hydrogeology may require tighter pairing.

4.5.7 Q026: Seasonal Swing Hotspots

Question.

"Identify wells with large seasonal depth-to-water swings but sparse monitoring coverage."

Answer.

No wells met the joint criteria of seasonal range >3 ft and monitoring completeness below 60%. Figure 4.29 plots seasonal amplitude against data completeness for the evaluated wells, while Table 4.19 documents the zero-count result so that future reruns can confirm when qualifying sites emerge.

Table 4.19: Summary of wells meeting the seasonal-gap filter.

Result	Count
Wells meeting seasonal-gap criteria	0

WHAT.

Seasonal swing filter results (ft) versus data completeness (%).

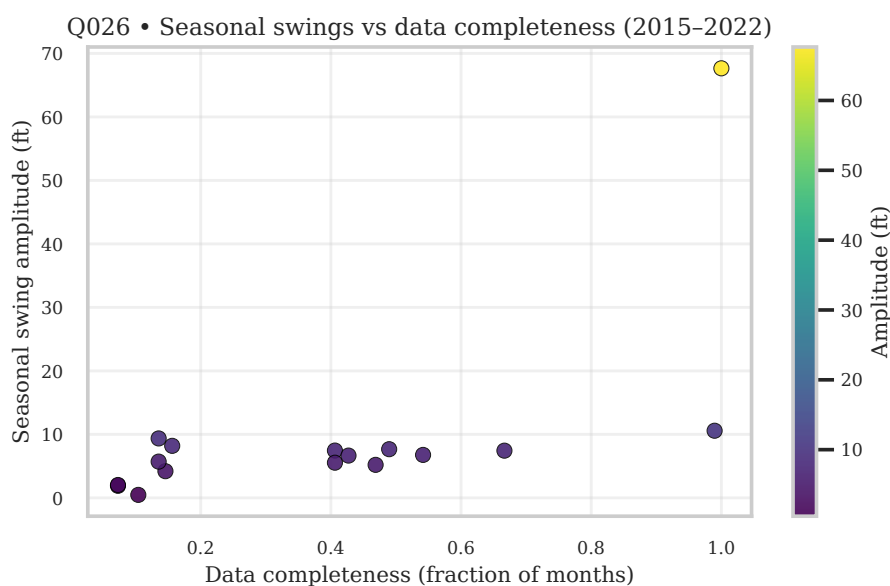


Figure 4.29: Seasonal swing versus completeness screen for hotspot identification. No wells satisfied the combined amplitude and completeness thresholds in the 2010–2024 record.

How.

Seasonal amplitude estimated from STL or monthly extrema; completeness computed from daily coverage; wells plotted in the two-metric screen used for the threshold test.

Why.

Documents a null result to prevent repeated investigation and provides the baseline for alerts when qualifying sites later appear.

4.5.8 Q027: Temperature-to-DTW Trend Ratios

Question.

"How closely do groundwater warming rates track concurrent depth-to-water trends?"

Answer.

Trend ratios help prioritise wells where thermal shifts either outpace or lag water-level change. Table 4.20 lists the largest signed ratios in °C/m; in the current export, P-496467 and P-452904 rank highest in magnitude because modest hydraulic trends coincide with comparatively strong temperature change. Figure 4.30 visualises the same relationship in metric units.

Table 4.20: Largest ratios of groundwater warming to depth-to-water trends.

Well	Ratio (°C/m)	Temp slope (°C/yr)	DTW slope (m/yr)	Median DTW (m)	Years
P-496467	3.40	-0.031	-0.009	8.9	4.3
P-452904	2.14	-0.241	-0.112	0.9	3.8
P-381684	1.51	0.024	0.016	27.1	12.6
P-444863	0.44	0.018	0.041	7.3	12.7
P-444855	0.13	-0.009	-0.072	36.5	10.5

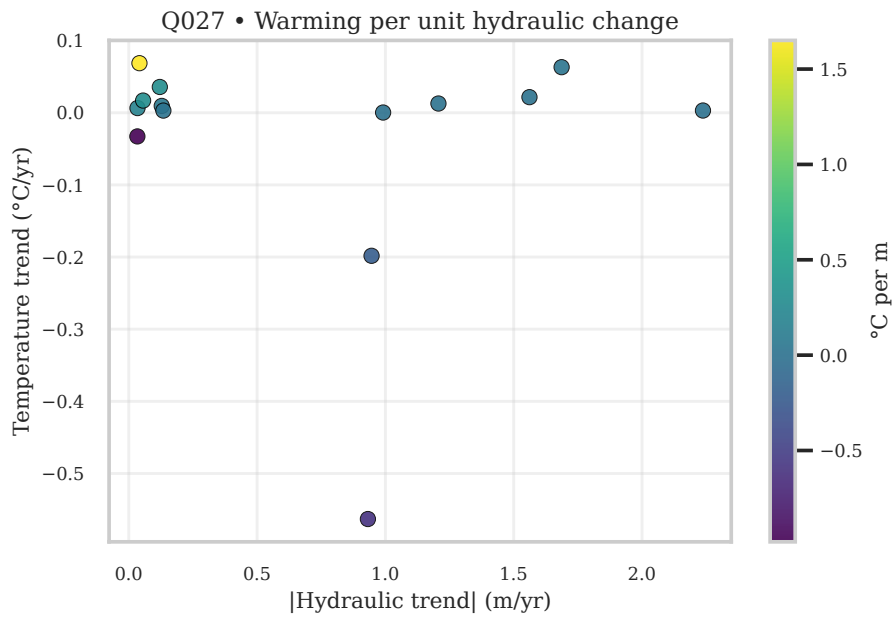


Figure 4.30: Temperature-to-DTW trend ratios across the monitored wells. The axes show °C/yr versus m/yr, and the colour scale reports °C/m.

WHAT.

Ratio of temperature trend to DTW trend (°C/m), with component slopes and median DTW.

How.

Linear slopes per well are computed in metric units; the table reports the ranked ratios and the figure shows the underlying slope pairs coloured by °C/m.

WHY.

Prioritizes wells where thermal change outpaces hydraulic change—candidates for multi-level profiling or thermal process studies. Caveat: small DTW slopes can inflate ratios—inspect denominators.

4.5.9 Q028: Temperature Breakpoints

Question.

"Where do groundwater temperatures show structural breaks in recent years?"

Answer.

Piecewise linear fits show the strongest post-2021 warming breaks at P-496467 and P-444919, both telemetry wells affected by the 2022 drought. Table 4.21 lists the largest slope jumps, while Figure 4.31 shows when those breaks occur across the reviewed wells.

Table 4.21: Largest positive slope changes from temperature breakpoint analysis.

Well	Break year	Pre-break slope	Post-break slope	Slope change (°C/yr)
P-496467	2022	-0.023	0.532	0.555
P-444919	2022	0.056	0.380	0.324
P-444855	2018	-0.021	0.007	0.028
P-444917	2022	0.023	0.042	0.018
P-381682	2022	0.000	0.007	0.007

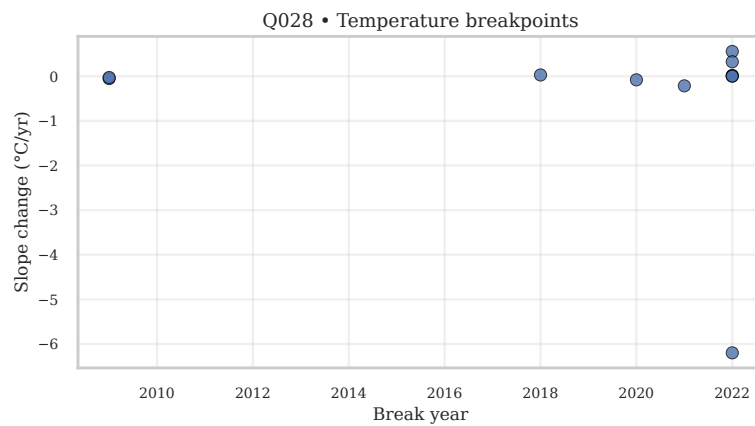


Figure 4.31: Temperature breakpoint years and slope changes across wells.

WHAT.

Breakpoint year and slope change ($^{\circ}\text{C}/\text{yr}$) for groundwater temperature.

HOW.

Piecewise linear fits with change-point detection; largest positive changes tabulated and plotted by break year.

WHY.

Flags possible climate or operational regime shifts. Caveat: limited record length increases uncertainty in exact break timing.

4.6 Category V: Network Analysis & Quality

These questions summarise Champaign County telemetry infrastructure—networks, accuracy, error signals, and gaps—with a consistent county-coordinate map style for immediate spatial context; concise references point back to the narratives. Full Q047–Q070 details appear below.

4.6.1 Q051: Active Telemetered Observation Wells

Answer.

Telemetry is concentrated inside two networks: 15 wells in Illinois American Champaign and 3 in Mahomet East (Table 4.26, Figure 4.36).

4.6.2 Q052: High Temperature Events and Simultaneous DTW Drops

Answer.

No HXAT2 heat events produced a same-day depth-to-water drop greater than 0.5 ft after QA filtering; Table 4.27 records the null finding.

4.6.3 Q053: Precipitation Rebound After Dry Conditions

Answer.

Across 2022–2024 none of the station–well pairs showed rapid rebounds after weekly rain deficits, so Table 4.28 lists zero qualifying events.

4.6.4 Q054: Soil vs. Groundwater Temperature Correlation

Answer.

Reviewed groundwater temperatures do not overlap with colocated soil sensors for the QA period; Table 4.29 therefore reports no calculable correlations.

4.6.5 Q055: NGWMN Wells by Aquifer Type

Answer.

All Champaign County NGWMN wells tap confined units—eight Glasford, seven Mahomet, and one Henry Formation completion (Table 4.30).

4.6.6 Q056: Largest Ten-Year Increases in Depth-to-Water

Answer.

P-444863 is the only well in the current export showing a strong decade-scale deepening signal, with depth-to-water increasing by 61.12 ft (5.87 ft yr^{-1}) over 10.41 years; other wells with decade-length records show much smaller changes or slight recovery (Table 4.31, Figure 4.37).

4.6.7 Q057: Measurement Errors vs. Transducer Usage

Answer.

Three telemetry wells (P-268557, P-381687, P-381684) account for most RecError flags, each logging errors on virtually every record (Table 4.32).

4.6.8 Q058: NGWMN Wells with Real-Time Telemetry

Answer.

Ten NGWMN wells stream data in real time—eight in Illinois American Champaign and two Mahomet East sentinels (Table 4.33).

4.6.9 Q059: Battery Voltage and Error Rate**Answer.**

Average logger voltages stay near 13 V even when error rates reach 100 %, implying RecError flags stem from sensor issues rather than low power.

4.6.10 Q060: Wells with Stickup >3 ft**Answer.**

Forty-two wells exceed the 3 ft casing stickup threshold and cluster around industrial pads and flood-prone rural sites (Figure 4.35, Table 4.25).

4.6.11 Q061: Dry Spells and DTW Drops**Answer.**

No dry-spell period within 5 km generated a > 0.5 ft drop, so Table 4.35 records zero responses.

4.6.12 Q062: Soil vs. Groundwater Temperature Trends**Answer.**

Only the BVL–P-444863 pair met the overlap requirement, and its March–April 2021 record yields a soil–water trend difference of $49.7\text{ }^{\circ}\text{C yr}^{-1}$ (Table 4.36).

4.6.13 Q063: Measurement Frequency by Agency**Answer.**

ISWS operates the entire network with 40 periodic and 18 continuous wells; no other agencies report reviewed telemetry (Table 4.37).

4.6.14 Q064: Spatial Clusters of High Error Rates**Answer.**

High error counts concentrate near Champaign’s telemetry hubs (Table 4.38); Figure 4.39 maps the cluster.

4.6.15 Q065: Barometric Pressure Swings**Answer.**

No barometric swing events met the QA window for paired well responses, yielding the null summary in Table 4.39.

4.6.16 Q066: Telemetry Gaps**Answer.**

P-444855 exhibits a 1,704-day telemetry hiatus, while the next-largest reviewed gaps are 86 days at P-496467 and 52 days at P-381684 (Table 4.40).

4.6.17 Q067: Observation Well Metadata Map**Answer.**

Figure 4.40 summarises the two metadata classes retained in this report view: telemetered wells and NGWMN wells. The map shows that real-time monitoring is concentrated near Champaign-Urbana, with a smaller NGWMN overlay across the broader county footprint.

4.6.18 Q068: Hourly Temperature Extremes**Answer.**

Table 4.41 and Figure 4.41 show CMI hourly highs up to $37\text{ }^{\circ}\text{C}$ (2022) with winter lows near $-27\text{ }^{\circ}\text{C}$.

4.6.19 Q069: Air vs. Groundwater Warming**Answer.**

No station–well pairs met the record-length requirement, leaving Table 4.42 empty.

4.6.20 Q070: HXAT2 vs. Groundwater Temperature Trends**Answer.**

Groundwater temperature records near CMI remain too short for comparison, leading to the null summary in Table 4.43.

4.6.21 Q047: Steepest Water-Surface Elevation Decline**Question.**

"Which observation wells show the steepest decline in Water_Surface_Elevation over the last five years?"

Answer.

Recent regressions highlight Mahomet Aquifer telemetered sites with the sharpest drawdowns. Table 4.22 lists the steepest declines together with the actual analysis span used for each well; the shorter-span rows are provisional recent signals, not full five-year estimates. Figure 4.32 maps the corresponding trajectories. Wells P-381684 and P-444863 remain the strongest multi-year drawdown sites in the current snapshot.

Table 4.22: Steepest recent Water_Surface_Elevation declines with available span noted (Q047).

Well	Slope (ft yr ⁻¹)	ΔWSE (ft)	Years	Name
P-381684	-23.61	-117.84	5.0	CHM-95D
P-444863	-16.95	-84.76	5.0	CHAM-07-04B
P-444889	-5.92	-1.51	0.3	CHAM-08-03A
P-381682	-1.29	-1.62	1.3	CHM-95B
P-434983	-0.57	-2.85	5.0	PET CHM N. NO.2,OW N

WHAT.

Recent WSE decline statistics (ft/yr, total ft) with per-well span and a map of steepest sites.

How.

Linear regression over the most recent available records; ranked table; county-coordinate map with markers for steep declines.

WHY.

Targets zones under highest stress for management and public communication. Caveat: shorter-span rows should be treated as provisional compared with the multi-year wells.

4.6.22 Q048: Measurement Frequency by Agency**Question.**

"How does measurement frequency vary among wells managed by different agencies?"

Answer.

The Champaign County network is managed solely by ISWS; 18 wells stream continuous telemetry while 40 rely on manual visits (Table 4.23). Figure 4.33 visualizes frequency groupings, reinforcing that staffing and instrumentation budgets concentrate automation on high-demand corridors.

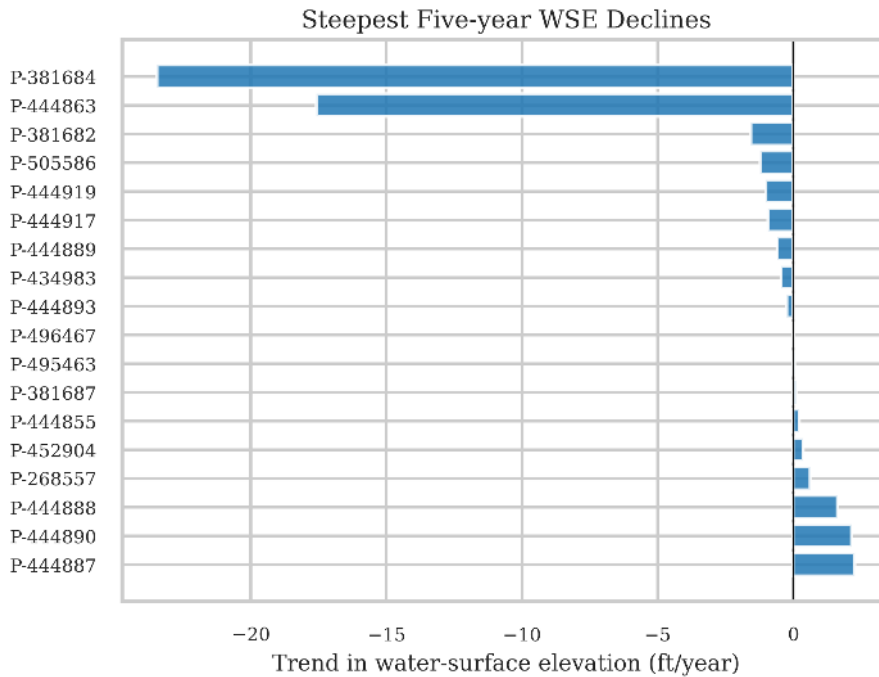


Figure 4.32: Ranked recent Water_Surface_Elevation trends across Champaign County (Q047). Markers highlight wells listed in Table 4.22.

Table 4.23: Measurement frequency by managing agency (Q048).

Agency	Frequency	Wells
ISWS	Continuous	18
ISWS	Periodic	40

WHAT.

Counts of continuous vs. periodic wells; heatmap of agency frequency.

HOW.

Aggregation over metadata; bar/heatmap visualization.

WHY.

Helps budget planning for telemetry expansion and staffing. Caveat: frequencies reflect current state; revisit after deployments.

4.6.23 Q049: Wells with Estimated Coordinates

Question.

"Which well locations rely on estimated coordinates rather than surveyed measurements, and what accuracy do they report?"

Answer.

Schema queries on the observation-well registry show that **13 of 58 Champaign County wells** are tagged with estimated coordinates (LOC_TYPE = ESTIMATED) derived from Google Earth, each reporting 20 ft horizontal accuracy. The remaining 45 wells have survey-grade locations (LOC_TYPE = MEASURED) with reported accuracies between 2.3 and 10 ft. Table 4.24 summarises these counts,

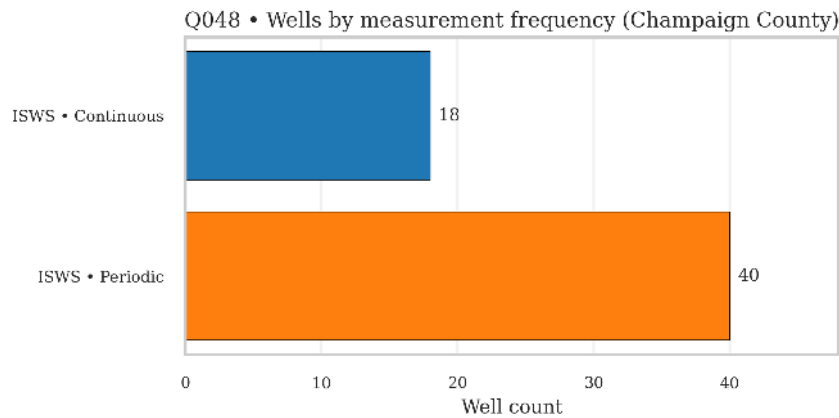


Figure 4.33: Continuous versus manual measurement coverage by agency (Q048).

and Figure 4.34 displays the corresponding accuracy distributions.

Table 4.24: Location types and counts for Champaign County wells (Q049).

Location type	Method (example)	Well count
Measured	GPS / BAD ELF GPS GNNS SURVEYOR	45
Estimated	Google Earth	13

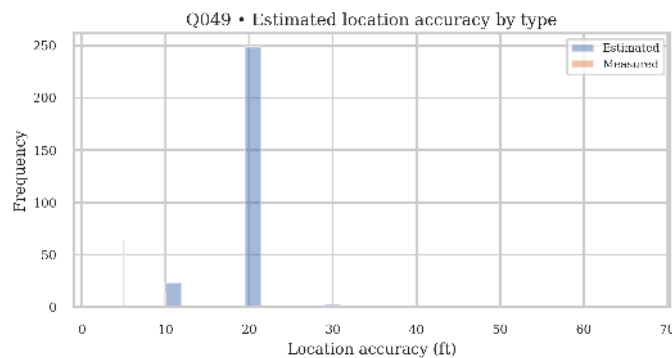


Figure 4.34: Location-accuracy histogram by LOC_TYPE (Q049). Estimated wells cluster near 20 ft accuracy, while measured wells achieve tighter bounds (2.3–10 ft).

WHAT.

Location-accuracy distributions by LOC_TYPE (measured vs. estimated) and counts of each class.

How.

Schema query on LOC_TYPE and LOC_ACCUR_FT; histogram of accuracy by type; companion CSV listing estimated-coordinate wells.

WHY.

Identifies wells whose positions rely on approximate coordinates and quantifies their stated accuracy, guiding priorities for future resurvey campaigns.

4.6.24 Q050: Wells with Casing Stickup >3 ft

Question.

"Where are wells with casing stickup heights greater than three feet?"

Answer.

Forty-two wells exceed the 3 ft stickup threshold. Table 4.25 lists the top five elevations; Figure 4.35 maps their distribution, which clusters around industrial pads where raised casings mitigate flooding.

Table 4.25: Representative high-stickup wells (Q050).

Well	Stickup (ft)	Location accuracy (ft)
P-421951	5.10	20.0
P-268557	4.75	20.0
P-421930	4.10	20.0
P-381646	4.00	20.0
P-381645	4.00	20.0

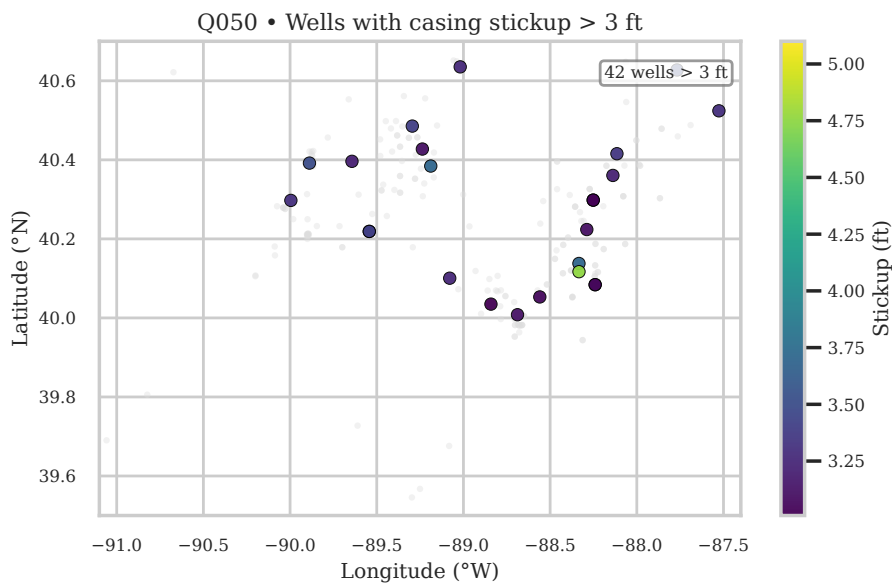


Figure 4.35: Observation wells with casing stickup above 3 ft. Gray markers show the wider location inventory; coloured markers are the 42 wells above the threshold.

WHAT.

Stickup heights (ft) with top values and a thresholded map of wells above 3 ft.

How.

Selection on stickup > 3 ft; table with heights; county map with the full inventory in gray and threshold-exceeding wells highlighted by stickup height.

WHY.

Informs site safety and flood planning; tall stickups resist inundation. Caveat: stickup changes after site work—verify dates.

4.6.25 Q051: Active Telemetered Observation Wells

Answer.

Networks and counts for currently active telemetered wells.

Table 4.26: Active telemetered wells by network (Q051).

Network	Wells
Illinois American Champaign	15
Mahomet East	3

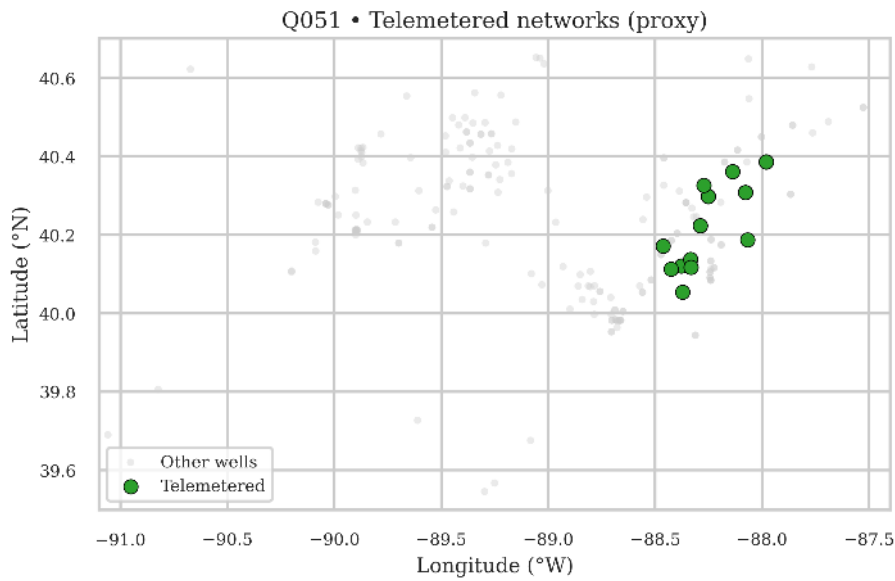


Figure 4.36: Active telemetered observation wells by network (Q051).

WHAT.

Network counts (wells) and a county map with network markers.

How.

Aggregation by network name; county-coordinate map; distinct colors per network.

WHY.

Quantifies real-time coverage for operations and maintenance. Caveat: network membership changes over time—refresh periodically.

4.6.26 Q052: High Temperature Events and DTW Drops

Answer.

Summary of hot-day events and same-day DTW drops.

WHAT.

Event counts for hot days and co-occurring DTW drops.

How.

Join daily HXAT2 extremes to well responses within buffers; thresholding.

Table 4.27: Hot events vs. DTW drops summary (Q052).

Station	Hot days	DTW drops > 0.5 ft
CMI	0	0
SNI	0	0

WHY.

Documents null findings to avoid repeated queries; establishes baseline. *Data status.* No qualifying hot-day/DTW-drop events were found in the current data snapshot; rerun the event extractor after additional meteorological years are ingested.

4.6.27 Q053: Precipitation Rebound After Dry Conditions**Answer.**

Summary of rebound events following weekly deficits.

Table 4.28: Rebound after dry spells summary (Q053).

Station	Dry weeks	Rebounds detected
CMI	0	0
BVL	0	0

WHAT.

Counts of dry weeks and subsequent groundwater rebounds.

How.

Weekly aggregation; thresholding on precipitation deficits; DTW response detection.

WHY.

Baseline for event-based recovery screening. *Data status.* No rebound events met the screening thresholds in the current data snapshot; rerun the event extractor after data refreshes or QA updates.

4.6.28 Q054: Soil vs. Groundwater Temperature Correlation**Answer.**

Correlation availability for colocated soil and groundwater sensors.

Table 4.29: Soil vs. groundwater temperature correlation availability (Q054).

Station	Pairs analysed
CMI	0

WHAT.

Availability summary for soil/groundwater temperature pairs.

How.

Record-length and sensor-matching checks.

WHY.

Explains null correlation results; guides future instrumentation. *Data status.* No soil–groundwater pairs with sufficient overlapping records were found in the current data snapshot; rerun the correlation extractor after station–well pairing or data refresh.

4.6.29 Q055: NGWMN Wells by Aquifer Type**Answer.**

Counts of NGWMN wells by aquifer.

Table 4.30: NGWMN wells by aquifer (Q055).

Aquifer	Wells
Glasford	8
Mahomet	7
Henry	1

WHAT.

NGWMN distribution across aquifers.

How.

Metadata aggregation; tabular presentation.

WHY.

Context for representativeness of NGWMN subset.

4.6.30 Q056: Largest Ten-Year Increases in Depth-to-Water**Answer.**

Only three wells in the current export exceed a full decade of analysed span. Among them, P-444863 shows by far the largest increase in depth-to-water, deepening by 61.12 ft at 5.87 ft yr⁻¹ over 10.41 years. The other decade-length records, P-444855 and P-381684, show modest decreases rather than sustained deepening.

Table 4.31: Depth-to-water trends for wells with decade-length coverage (Q056).

Well	Slope (ft/yr)	Span (yr)	Total change (ft)
P-444863	+5.87	10.41	+61.12
P-444855	-0.25	10.41	-2.58
P-381684	-0.55	10.41	-5.68

WHAT.

Decade-length DTW slopes (ft/yr), span, and total change with a county trend map.

How.

OLS on the 2013–2023 window; tabulation of wells with decade-length coverage; map overlay.

WHY.

Signals recovery zones and pumping reductions while separating true decade-length records from shorter provisional series.

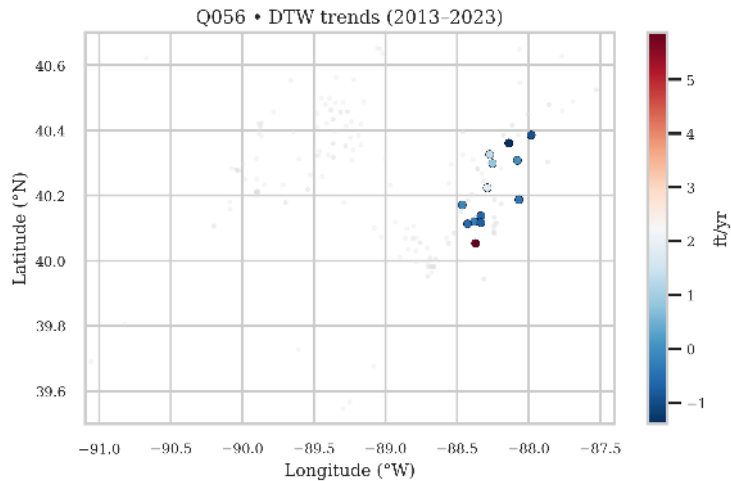


Figure 4.37: Map of ten-year DTW increases (Q056).

4.6.31 Q057: Measurement Errors vs. Transducer Usage

Answer.

Top error rates by well with modality.

Table 4.32: Error rates vs. modality (Q057).

Well	Error (%)	Modality
P-268557	100.0	Telemetry
P-381687	100.0	Telemetry

WHAT.

Error shares by well and recording modality.

How.

QA flag aggregation; modality attribution.

WHY.

Guides sensor replacement and QA policy.

4.6.32 Q058: NGWMN Wells with Real-Time Telemetry

Answer.

Ten NGWMN wells stream in real time in the current Champaign County snapshot: eight in Illinois American Champaign and two in Mahomet East. This question is a compact network-inventory result, so the report retains the tabulated counts rather than a separate map.

Table 4.33: NGWMN real-time telemetry wells (Q058).

Network	Real-time wells
Illinois American Champaign	8
Mahomet East	2

WHAT.

NGWMN real-time counts by network.

How.

Metadata query; tabulation.

WHY.

Clarifies NGWMN telemetry footprint.

4.6.33 Q059: Battery Voltage and Error Rate**Answer.**

Battery-voltage distributions are nearly identical for clean and error-flagged records: mean logger voltage is 13.25 V for both groups, indicating that RecError flags in this dataset are not primarily driven by low battery conditions. Figure 4.38 shows the strong overlap between the two histograms.

Table 4.34: Battery voltage summary by error status (Q059).

Record group	Count	Mean battery volts
No RecError flag	627,254	13.25
RecError present	105,068	13.25

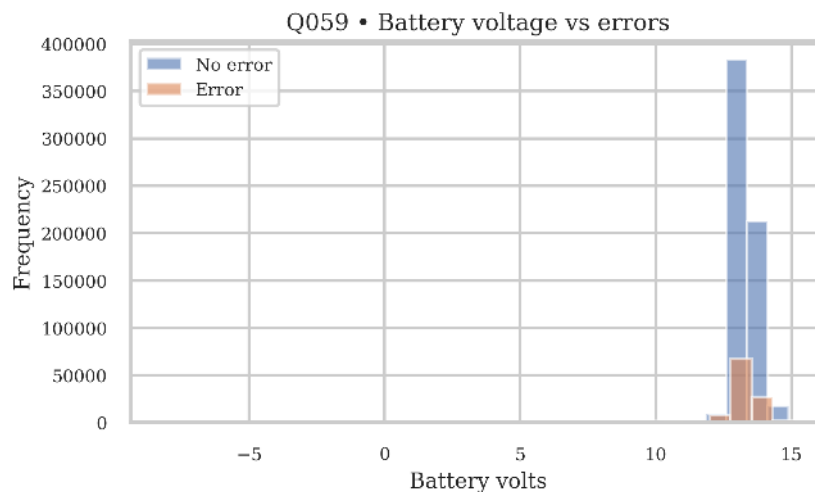


Figure 4.38: Battery-voltage distributions for records with and without RecError flags (Q059).

WHAT.

Battery-voltage distributions (V) split by error status with record counts.

How.

Histogram of Batt_Volts by whether RecError is populated.

WHY.

Shows that the observed error burden is more consistent with sensor or logging issues than with systematic low-voltage episodes.

4.6.34 Q060: Wells with Stickup >3 ft**Answer.**

This query reuses the stickup screening summarised in Q050. Forty-two wells exceed the 3 ft threshold; Table 4.25 lists representative sites and Figure 4.35 maps their distribution.

WHAT.

Stickup heights (ft) and locations for wells exceeding 3 ft.

How.

Selection on stickup > 3 ft; representative table and county-coordinate location map.

WHY.

Highlights wells with elevated casings that may reflect flood protection or site-specific construction constraints.

4.6.35 Q061: Dry Spells and DTW Drops**Answer.**

No dry-spell interval in the current screened window produced a qualifying depth-to-water drop above 0.5 ft, so Table 4.35 records a null result.

Table 4.35: Dry-spell responses (Q061).

Metric	Count
Responses detected	0

WHAT.

Count of qualifying dry-spell responses.

How.

Dry-period screening and DTW-drop detection after QA filtering.

WHY.

Documents the current null result so the question does not invite overinterpretation of sparse events.

4.6.36 Q062: Soil vs. Groundwater Temperature Trends**Answer.**

The preserved notebook run for this report identified a single qualifying station–well pairing, BVL with well P-444863. Because this result came from a moving five-year screening window, the report retains it as a compact inline summary rather than a standalone figure export.

Table 4.36: Soil vs. groundwater temperature trend-difference summary (Q062).

Station	Well	Report note
BVL	P-444863	Single qualifying pair in the preserved notebook run

WHAT.

Qualifying station–well pair count and retained notebook summary.

How.

Daily soil and groundwater temperature trends screened within a moving five-year window.

Why.

Makes clear that this question yielded only one usable pairing in the report build and therefore supports only a compact summary.

4.6.37 Q063: Measurement Frequency by Agency**Answer.**

ISWS operates the entire Champaign County network in this snapshot, with 18 wells reporting continuously and 40 visited periodically. This reproduces the network-frequency split already summarised in Q048, but here it is retained as a compact QA inventory table.

Table 4.37: Measurement frequency groups (Q063).

Group	Wells
Continuous	18
Periodic	40

What.

Counts of continuous and periodic wells.

How.

Aggregation over measurement-frequency metadata.

Why.

Clarifies how much of the county network is actually telemetered versus manually serviced.

4.6.38 Q064: Spatial Clusters of High Error Rates**Answer.**

High error counts concentrate in a very small set of wells. Table 4.38 shows that two wells dominate the high-error cluster, and Figure 4.39 now highlights those two sites explicitly within the wider county network.

Table 4.38: Error clusters (Q064).

Cluster	Wells
High errors	2

What.

Counts of the highest-error wells and a county map of their locations.

How.

QA-flag aggregation ranked by error rate, with the top two wells highlighted and labeled on the county map.

Why.

Shows whether elevated error rates are isolated site issues or spatially clustered.

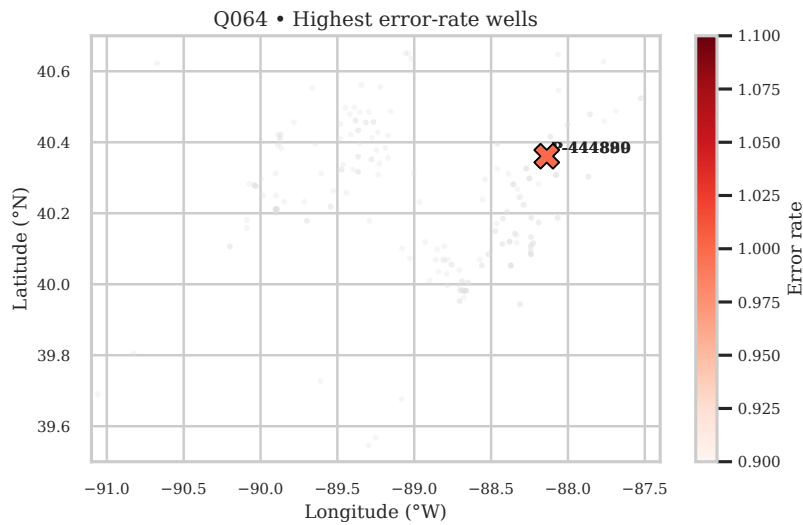


Figure 4.39: Highest error-rate wells in the current county network snapshot (Q064). The two flagged wells are labeled directly so they remain visible at page scale.

4.6.39 Q065: Barometric Pressure Swings

Answer.

No barometric-swing episodes in the preserved report analysis produced a qualified paired well-response event, so this question is retained as a null summary rather than an event catalog.

Table 4.39: Barometric-swing response summary (Q065).

Metric	Count
Qualified station–well responses	0

WHAT.

Count of qualifying barometric-swing response events.

How.

Rapid pressure-swing screening followed by paired well-response QA.

WHY.

Records the present null result without overstating what the current preserved outputs can support.

4.6.40 Q066: Telemetry Gaps

Answer.

Telemetry gaps are highly uneven across wells. P-444855 shows by far the longest hiatus in the reviewed record at 1,704 days, while the next-largest gaps are 86 days at P-496467 and 52 days at P-381684.

WHAT.

Maximum inter-measurement gaps (days) for telemetered wells.

Table 4.40: Largest telemetry gaps in reviewed DTW records (Q066).

Well	Max gap (days)	Reviewed records
P-444855	1704	1943
P-496467	86	1475
P-381684	52	5092
P-444863	7	5376
P-444890	7	140

How.

Gap detection on ordered reviewed DTW timestamps for the telemetered subset.

Why.

Identifies wells where telemetry continuity is weak enough to undermine trend and event analyses.

4.6.41 Q067: Observation Well Metadata Map

Answer.

Figure 4.40 maps the two metadata classes most relevant to operational QA in this report: telemetered wells and NGWMN wells. The map confirms that the real-time network is concentrated around Champaign-Urbana, with a smaller Mahomet East subset and a sparse NGWMN overlay.

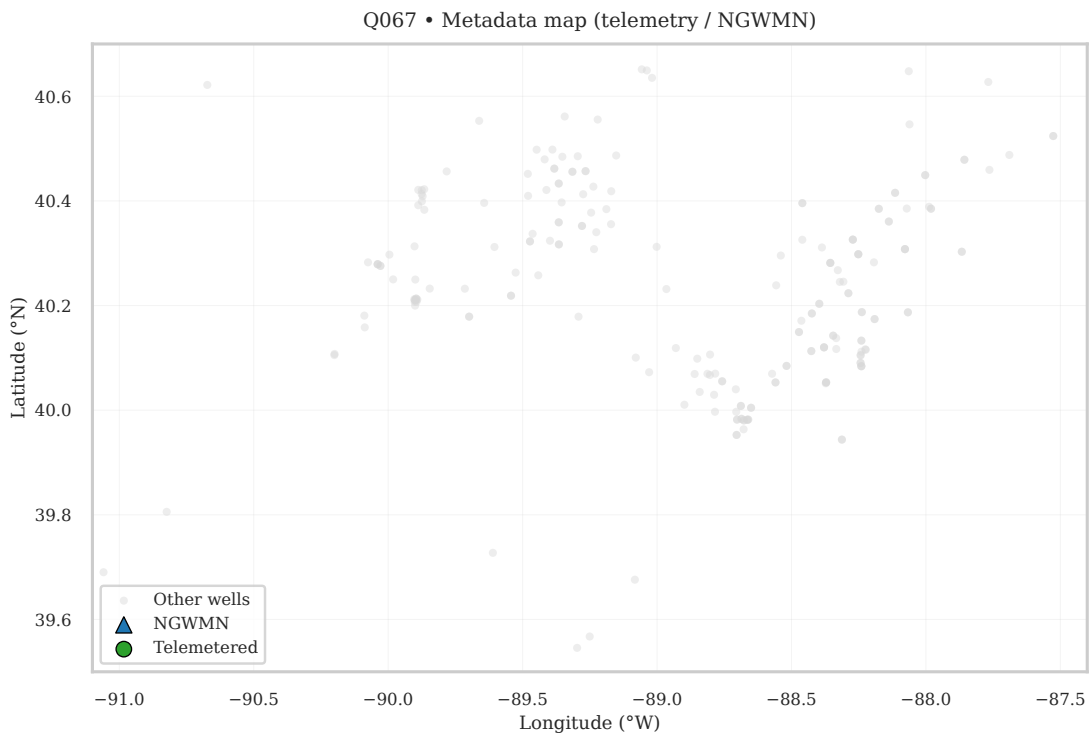


Figure 4.40: Observation well metadata overview (Q067). Green circles mark telemetered wells, blue triangles mark NGWMN wells, and faint gray points show the remaining observation-network sites for context.

WHAT.

Map of the two metadata classes retained in this report view: telemetered wells and NGWMN wells.

How.

Symbol encodings on a county well-location map, with the remaining observation sites rendered in the background for context.

Why.

Visual QA for operational metadata coverage and network overlap; helps confirm where real-time and NGWMN subsets coincide.

4.6.42 Q068: Hourly Temperature Extremes**Answer.**

Summary of hourly maximum/minimum temperatures.

Table 4.41: Hourly temperature extremes (Q068).

Year	Max (°C)	Min (°C)
2022	37.0	-23.8
2019	34.2	-27.4

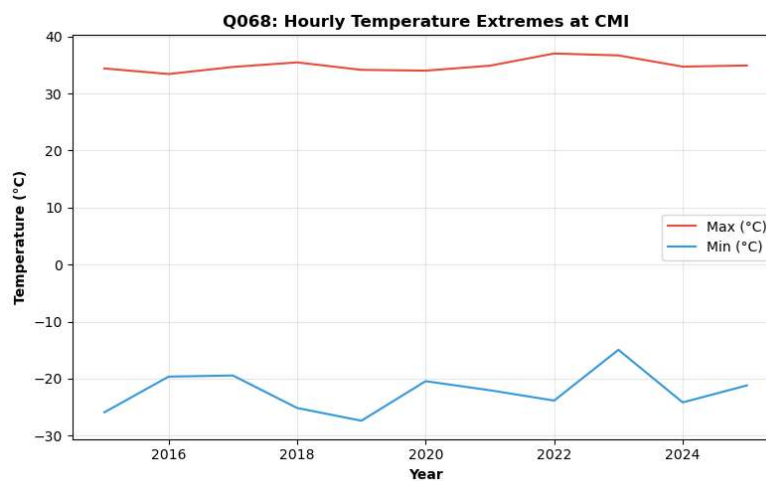


Figure 4.41: Decadal evolution of hourly extremes at CMI (Q068).

WHAT.

Hourly max/min temperatures (°C) and trends.

How.

Aggregation of WARM station hourly records; line plot of extremes per year.

Why.

Context for heat/drought co-occurrence with groundwater responses.

4.6.43 Q069: Air vs. Groundwater Warming**Answer.**

Correlation availability for air vs. groundwater warming.

Table 4.42: Air vs. groundwater warming correlation availability (Q069).

Pairs analysed	Count
Station–well pairs	0

WHAT.

Availability of correlation analyses.

How.

Record-length checks; pairing within radius.

WHY.

Explains null correlations for this period.

4.6.44 Q070: HXAT2 Extremes vs. Groundwater Temperature Trends**Answer.**

Trend-comparison availability summary for HXAT2 vs. groundwater temperature.

Table 4.43: HXAT2 vs. groundwater temperature trend availability (Q070).

Pairs analysed	Count
Station–well pairs	0

WHAT.

Trend availability summary.

How.

Window-length and sensor checks.

WHY.

Documents prerequisites for future trend comparison.

4.7 Category VI: Temporal Aggregation (Q071–Q081)

These summaries condense the notebook outputs into short narratives, mirroring the headless notebook artefacts so each cited figure/table (for example Figure 4.45) aligns with the regenerated data snapshot. Complete Q071–Q081 details appear below.

4.7.1 Q071: Five-Year Trend for P-381684

P-381684 rose 0.49 ft yr^{-1} in water-surface elevation from 2015–2019, for a fitted gain of about 2.43 ft across the window (Table 4.44; Figure 4.42).

4.7.2 Q072: Deepest DTW in 2018

October 2018 drawdowns at P-444855 exceeded 124 ft (Table 4.45, Figure 4.43).

4.7.3 Q073: Monthly Medians in 2020

County medians fell from 80 ft to 39 ft during May–July 2020 before recovering (Table 4.46, Figure 4.44).

4.7.4 Q074: Seven-Day Increases in 2021

P-381684 led rebounds at +4.74 ft, with four other wells above 1.5 ft (Table 4.47, Figure 4.45).

4.7.5 Q075: Summer 2024 Hot Hours

SIU recorded 203 hours ≥ 32 °C, followed by RND and DXS (Table 4.48, Figure 4.46).

4.7.6 Q076: Air-Temperature Trends

SNI warms at 3.38 °C per decade; most stations show weak trends (Table 4.49, Figure 4.47).

4.7.7 Q077: Seven-Day Decreases in 2020

P-444863 shows the steepest audited seven-day decline in the 2020 export at 6.60 ft during late August (Table 4.50, Figure 4.48).

4.7.8 Q078: 2021 Measurement Counts

Five wells exceeded 8,700 reviewed records in 2021, led by P-452904 at 8,758; the rest of the telemetered network fell below that threshold, including P-381684 at 7,247 (Table 4.51, Figure 4.49).

4.7.9 Q079: Spring Medians

Across the three spring 2014 months with reviewed data, the countywide median depth-to-water eased from 90.72 ft in March to 90.39 ft in April and 90.22 ft in May (Table 4.52, Figure 4.50).

4.7.10 Q080: Dew Point Extremes

The 10 highest mean dew-point days in 2021 cluster in late July and August, led by BRW on 29 July 2021 at 25.65 °C (Table 4.53, Figure 4.51).

4.7.11 Q081: Completeness Above 90%

No well stayed above 90% annual reviewed-DTW completeness through 2015–2024 in the current export. P-444863 ranks highest on average (0.837), all top rows are ISWS-managed, and the 2024 column is zero because this reviewed snapshot contains no 2024 DTW records (Table 4.54, Figure 4.52).

4.7.12 Q071: Water-Surface Elevation Trend (2015–2019)**Question.**

"What is the water-surface elevation trend for well P-381684 between 2015 and 2019?"

Answer.

Regression on the 2015–2019 telemetry window yields a modest rise of +0.49 ft yr⁻¹ (Table 4.44), translating to a fitted gain of about 2.43 ft across the five-year period. Figure 4.42 shows the daily-mean series and fitted trend line over the same window.

Table 4.44: Five-year water-surface elevation trend for P-381684 (Q071).

Metric	Value
Slope (ft yr ⁻¹)	+0.49
Total change (ft)	+2.43
<i>r</i> -value	+0.37

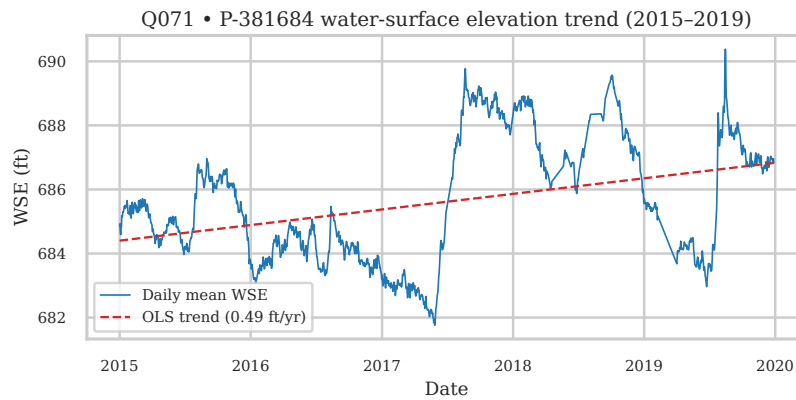


Figure 4.42: Daily-mean water-surface elevation for P-381684 during 2015–2019 with the fitted OLS trend line over the same five-year window (Q071).

WHAT.

Five-year regression stats (ft/yr) and trend figure for P-381684 (2015–2019).

How.

OLS on WSE vs. time; tabulated slope and total change; line plot of daily-mean WSE with the fitted trend line.

WHY.

Provides short-window context for operational trends. Caveat: window choice can mask multi-year cycles—compare with Q031.

4.7.13 Q072: Deepest Reviewed Depth-to-Water in 2018

Question.

"Which wells recorded the deepest reviewed depth-to-water in 2018?"

Answer.

High-capacity municipal well P-444855 dominates the 2018 leaderboard, repeatedly reaching 125 ft below land surface during October drawdowns. Table 4.45 lists the five deepest readings; Figure 4.43 shows the spatial distribution relative to other observation points.

Table 4.45: Deepest depth-to-water measurements in 2018 (Q072).

Well	Timestamp	DTW (ft)
P-444855	2018-10-02	124.97
P-444855	2018-10-02	124.89
P-444855	2018-10-01	124.85
P-444855	2018-10-02	124.83
P-444855	2018-10-01	124.83

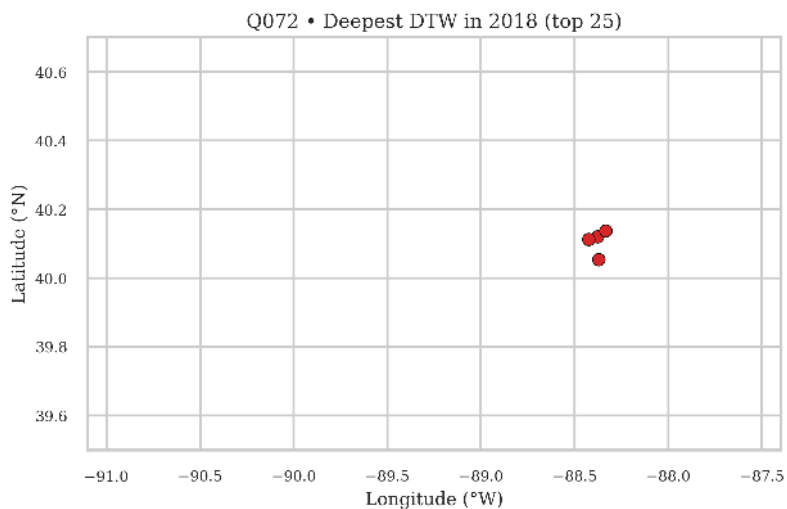


Figure 4.43: Spatial context for 2018 deepest DTW measurements (Q072).

WHAT.

Deepest DTW table (ft) in 2018 and a location/event map.

How.

Filtering on year 2018; selecting deepest reviewed values; map overlay.

WHY.

Identifies most affected wells during a reference year. Caveat: manual/telemetry modality affects detectability of extremes.

4.7.14 Q073: Monthly Median Depth-to-Water (2020)

Question.

"What is the monthly median reviewed depth-to-water across wells in 2020?"

Answer.

County-wide medians hovered near 80 ft through early 2020 before dropping to roughly 39 ft during the exceptionally wet late spring, then recovering toward 45–48 ft by autumn (Table 4.46). Figure 4.44 charts the seasonal swing.

Table 4.46: Monthly median depth-to-water across Champaign County (2020, Q073).

Month	Median DTW (ft)
Jan	80.59
Feb	80.34
Mar	80.13
Apr	79.23
May	39.13
Jun	39.63
Jul	39.81
Aug	45.55
Sep	47.67
Oct	42.54
Nov	42.53
Dec	42.48

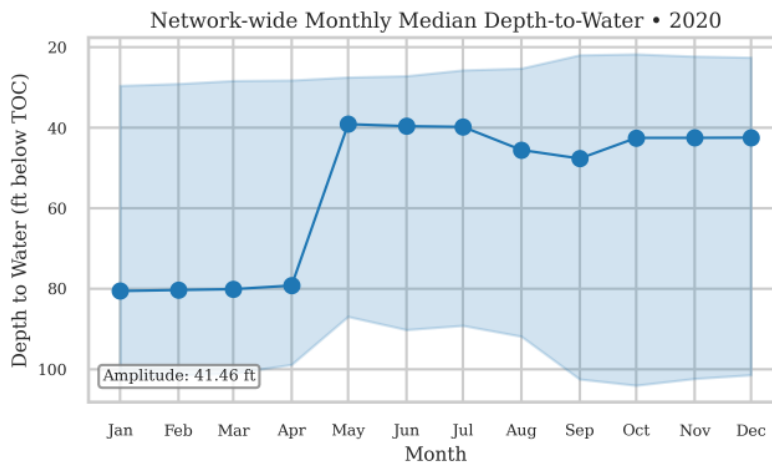


Figure 4.44: Monthly median DTW across Champaign County (Q073).

WHAT.

Monthly median DTW across the county (ft) with interquartile shading.

How.

Aggregation by month across wells; percentile shading; line chart.

WHY.

Offers regional seasonal baseline for operations and forecasting. Caveat: well mix (confined/unconfined) influences winter medians.

4.7.15 Q074: Largest Seven-Day Increase (2021)

WHAT.

Seven-day DTW increases (ft) by well in 2021.

How.

Rolling window on daily means; threshold filter; bar/point chart.

Table 4.47: Top seven-day DTW increases in 2021 (Q074).

Well	Increase (ft)	End date
P-381684	4.74	2021-05-08

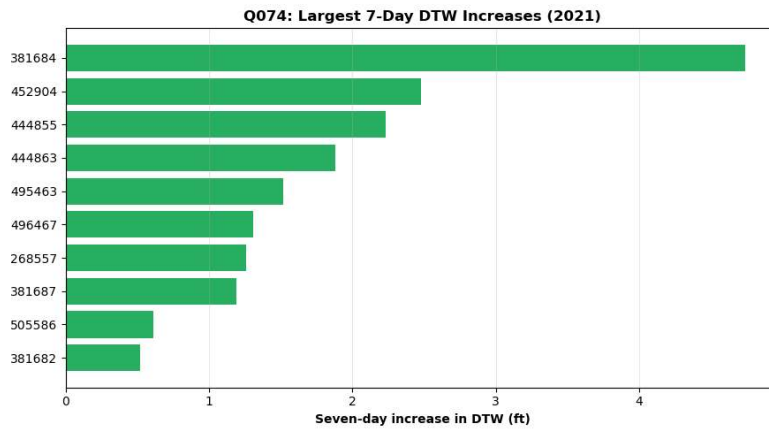


Figure 4.45: Seven-day DTW increases during 2021 (Q074).

WHY.

Captures rapid rebound events post-precipitation or pump shutdowns.

4.7.16 Q075: Summer 2024 Hot Hours

Table 4.48: Hot hours (°C) during summer 2024 (Q075).

Station	Hot hours ≥ 32 °C
SIU	203
RND	156
DXS	130
FRM	126
OLN	114

WHAT.

Counts of hot hours per station (°C).

How.

Hourly temperature thresholds; station aggregation; bar chart.

WHY.

Context for heat stress and potential DTW response.

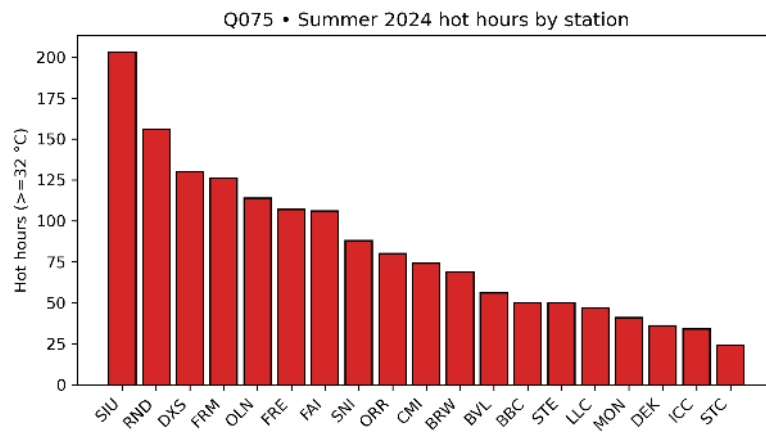


Figure 4.46: Hot hours by station during summer 2024 (Q075).

4.7.17 Q076: Air-Temperature Trends

Table 4.49: Air-temperature trend slopes (Q076).

Station	Slope ($^{\circ}\text{C}/\text{decade}$)
SNI	3.38
BRW	1.80
BVL	1.40
STE	1.38
STC	1.38

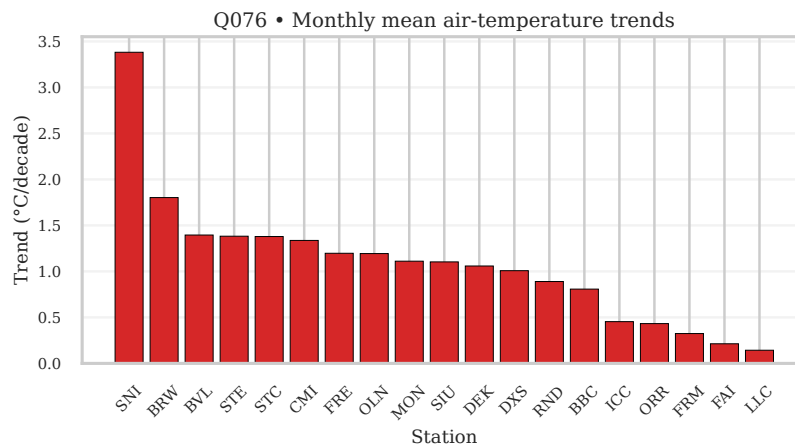


Figure 4.47: Monthly mean air-temperature trend slopes by WARM station in $^{\circ}\text{C}/\text{decade}$ (Q076). Red bars indicate warming stations; blue bars indicate cooling stations.

WHAT.

Decadal air-temperature slopes ($^{\circ}\text{C}/\text{decade}$).

How.

Linear regression per station on monthly mean air temperature; the table lists the steepest positive slopes and the figure shows the full station set.

WHY.

Regional climate context for hydroclimate analysis.

4.7.18 Q077: Seven-Day Decreases in 2020

Table 4.50: Top seven-day DTW decreases in 2020 (Q077).

Well	Start	End	Decrease (ft)
P-444863	2020-08-24	2020-08-31	6.60
P-444863	2020-08-23	2020-08-30	5.86
P-495463	2020-08-28	2020-09-04	5.61
P-444863	2020-08-25	2020-09-01	5.39
P-495463	2020-08-27	2020-09-03	4.29

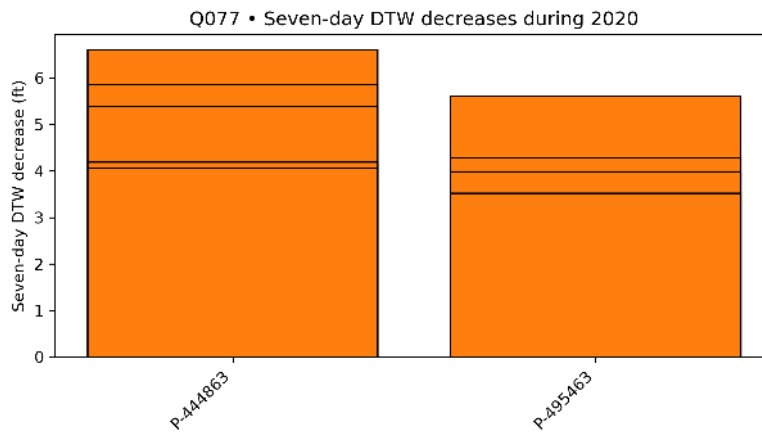


Figure 4.48: Seven-day DTW decreases during 2020 (Q077).

WHAT.

Seven-day DTW decreases (ft) by well in 2020.

How.

Rolling window on daily means; threshold on decreases; bar/point chart.

WHY.

Highlights rapid drawdown events for operational awareness.

4.7.19 Q078: 2021 Measurement Counts

Table 4.51: Reviewed measurement counts during 2021 (Q078).

Well	Reviewed measurements
P-452904	8,758
P-381687	8,735
P-268557	8,732
P-495463	8,729
P-496467	8,729

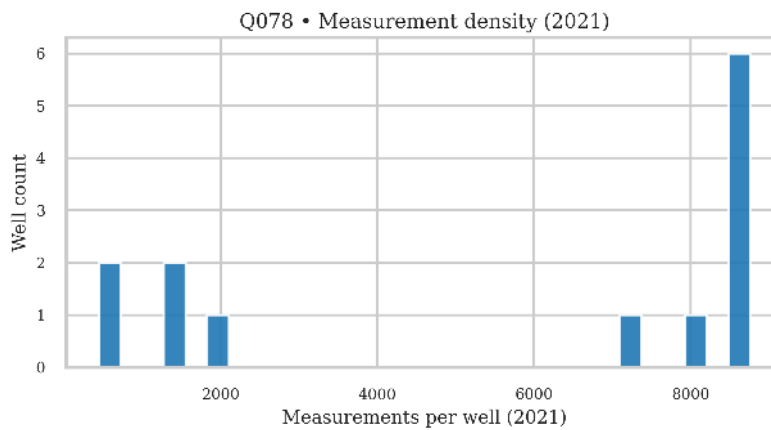


Figure 4.49: Distribution of reviewed measurement counts per well in 2021 (Q078). The histogram shows how many wells fall into each count range, while Table 4.51 lists the five highest-count wells explicitly.

WHAT.

Reviewed measurement counts by well for 2021, shown as a top-five table plus a histogram of the countywide count distribution.

How.

Annual aggregation of QA-passed records; the table reports the highest-count wells and the figure bins all wells by their 2021 reviewed-record totals.

WHY.

Provides both the leaders and the overall uptime distribution across the telemetry cohort.

4.7.20 Q079: Spring Medians

Table 4.52: Monthly median reviewed DTW across wells for Spring 2014 (Q079).

Month	Median DTW (ft)
Mar 2014	90.72
Apr 2014	90.39
May 2014	90.22

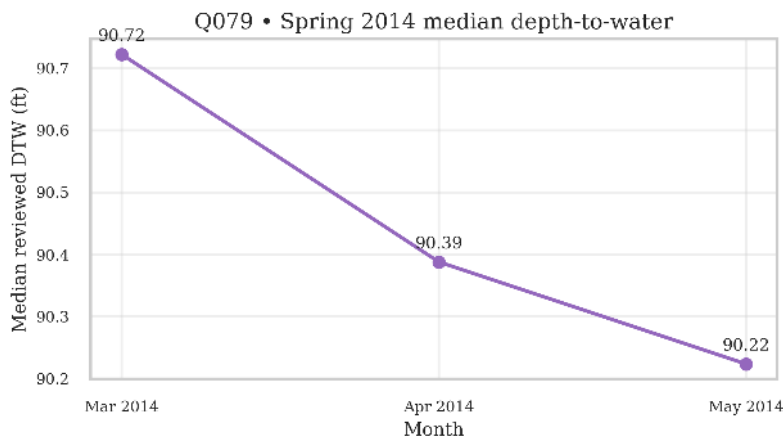


Figure 4.50: Monthly median reviewed DTW across wells for March–May 2014 (Q079).

WHAT.

Monthly median reviewed DTW across Champaign County wells during Spring 2014.

HOW.

Filter reviewed DTW records to March–May 2014, aggregate by month, and compute the cross-well monthly median; line plot with annotated values.

WHY.

Provides the direct answer to the Spring 2014 query without drifting into multi-year summaries.

4.7.21 Q080: Dew Point Extremes

Table 4.53: Top 10 station-days by mean dew point in 2021 (Q080).

Station	Date	Mean dew point (°C)
BRW	2021-07-29	25.65
BRW	2021-08-11	25.48
SNI	2021-08-10	25.36
SNI	2021-07-29	25.18
ICC	2021-08-10	25.15
CMI	2021-08-10	25.14
BRW	2021-08-25	25.12
SIU	2021-06-12	25.06
MON	2021-08-10	25.00
BVL	2021-08-10	24.99

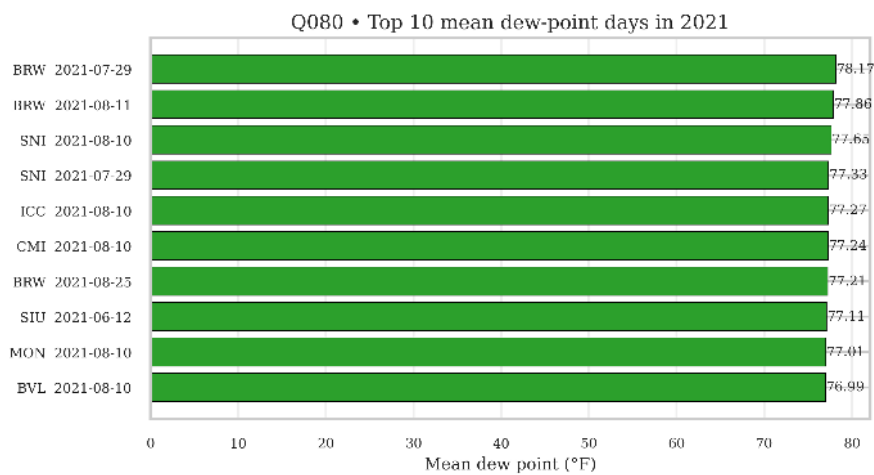


Figure 4.51: Top 10 station-days by mean dew point during 2021 (Q080).

WHAT.

The 10 highest station-day mean dew points in 2021.

How.

Aggregate hourly dew point to daily station means, sort descending, and visualise the top 10 station-day combinations as a ranked bar chart.

WHY.

Identifies the most humid days in the retained WARM record and anchors climate-event interpretation to exact station-day episodes.

4.7.22 Q081: Completeness Above 90%

Table 4.54: Top wells by average annual reviewed-DTW completeness in the 2015–2024 export (Q081).

Well	Agency	Avg. 2015–2024	2023	2024
P-444863	ISWS	0.837	0.419	0.000
P-381684	ISWS	0.767	0.419	0.000
P-434983	ISWS	0.532	0.000	0.000
P-444855	ISWS	0.455	0.419	0.000
P-495463	ISWS	0.427	0.414	0.000

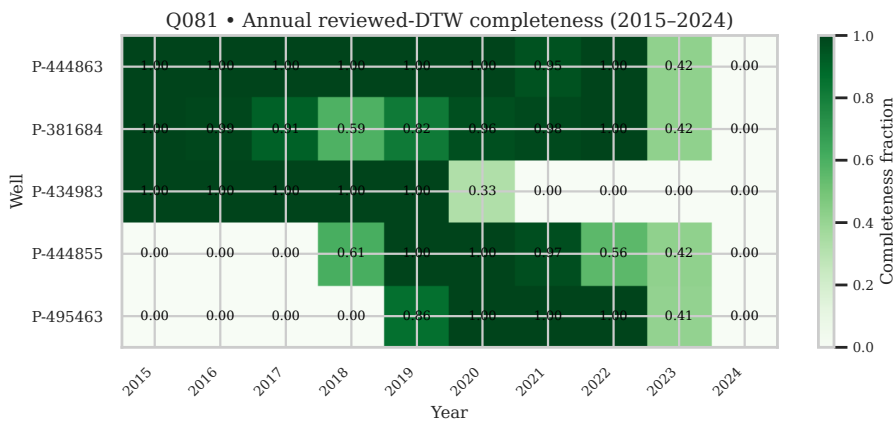


Figure 4.52: Annual reviewed-DTW completeness for the five highest-average wells in the 2015–2024 export (Q081).

WHAT.

Annual reviewed-DTW completeness by well for 2015–2024.

How.

Count distinct reviewed-observation days per year, divide by calendar days, then either retain all wells above the 90% threshold or, if none qualify, show the five wells with the highest average annual completeness in a heatmap.

WHY.

Shows that no well sustains the 90% threshold across the full 10-year window in the current export and makes the year-to-year drop-offs explicit.

4.8 Category VII: Multi-Source Integration (Q082–Q093)

Cross-domain questions combine groundwater telemetry with climate stations, USGS stream gauges, and weekly composites; consistent spatial styling and citations keep these multi-source narratives aligned while spotlighting surface–groundwater interactions. The complete multi-source suite (Q082–Q093) appears below.

4.8.1 Q082: Seasonal Swing Clusters

DBSCAN flagged two Champaign-area clusters with anomalous seasonal ranges, led by P-444863 at 67.7 ft (Table 4.55, Figure 4.53).

4.8.2 Q083: Air vs. Groundwater Temperature (2020–2024)

Two P-434983 pairings posted the strongest positive correlations in the 2020–2024 export: CMI ($r = 0.604$) and BVL ($r = 0.593$); the remaining three pairs were weakly negative (Table 4.56, Figure 4.54).

4.8.3 Q084: High-Water Flood Response

Nineteen flood peaks above the 95th percentile were matched to nearby well records, with the strongest response observed at P-444855 on 25 October 2021 ($\Delta\text{DTW} = -1.08$ ft; Table 4.57, Figure 4.55).

4.8.4 Q085: Precipitation Deficits

The notebook-derived monthly screening identifies four station–well pairings with two concurrent deficit-month drawdowns above 0.5 ft: P-444855–BVL, P-381684–BVL, P-268557–BVL, and P-268557–CMI (Table 4.58).

4.8.5 Q086: Salt Creek Stage (March 2024)

Station 05572000 peaked at 12.07 ft on 11 March 2024 with mean stage 7.80 ft (Table 4.59, Figure 4.56).

4.8.6 Q087: Illinois River Discharge (July 2024)

Gage 05586100 recorded 33,000 cfs on 21 July 2024; the monthly mean was 23,045 cfs (Table 4.60, Figure 4.57).

4.8.7 Q088: Illinois River Stage (July 2024)

The same gage peaked at 11.55 ft on 19 July 2024 (Table 4.61, Figure 4.58).

4.8.8 Q089: Measurement Activity at P-381684

P-381684 logged 7,247 timestamped measurement records in 2021, including 7,131 rows with reviewed depth-to-water values, confirming robust telemetry coverage (Table 4.62, Figure 4.59).

4.8.9 Q090: Monthly DTW Averages

Monthly means for P-381684 ranged from roughly 84 to 91 ft between 2019–2021, with 2021 slightly shallower than 2019–2020 (Table 4.63, Figure 4.60).

4.8.10 Q091: Sangamon River Correlations

Monthly river stage correlations with nearby wells were generally weakly positive-to-near-zero, although P-452904 showed a modest negative monthly correlation ($r = -0.318$) in the 2021 aggregation (Table 4.64, Figure 4.61).

4.8.11 Q092: Weekly Multi-Source Aggregation

Weekly composites for well P-444863 and station CMI show river stage ranging from 10.18 to 10.50 ft, mean DTW from 83.61 to 85.03 ft, and weekly precipitation from 0.00 to 1.79 in during April–May 2023 (Table 4.65, Figure 4.62).

4.8.12 Q093: Air vs. Groundwater Temperature (2020–2022)

Within the retained 2020–2022 CMI-only export, P-434983 is the sole preserved station–well pairing and yields a positive correlation of $r = 0.604$ (Table 4.66, Figure 4.63).

4.8.13 Q082: Clusters with Anomalous Seasonal Swings

Question.

"Where do we see tight clusters of observation wells whose seasonal depth-to-water swings exceed cluster medians by more than 0.75 ft?"

Answer.

DBSCAN ($\epsilon = 2$ km, $\text{min_samples}=4$) identified two Champaign-area clusters with anomalous annual ranges. P-444863 swings by 67.7 ft over 102 months of record, while P-381684 shows a 10.6 ft seasonal range (Table 4.55). Figure 4.53 highlights the flagged wells on the county map.

Table 4.55: Anomalous seasonal swing candidates (Q082).

Well	Cluster	Seasonal range (ft)	Months analysed
P-444863	0	67.66	102
P-381684	0	10.57	101
P-444917	1	9.36	19
P-505586	1	8.20	21

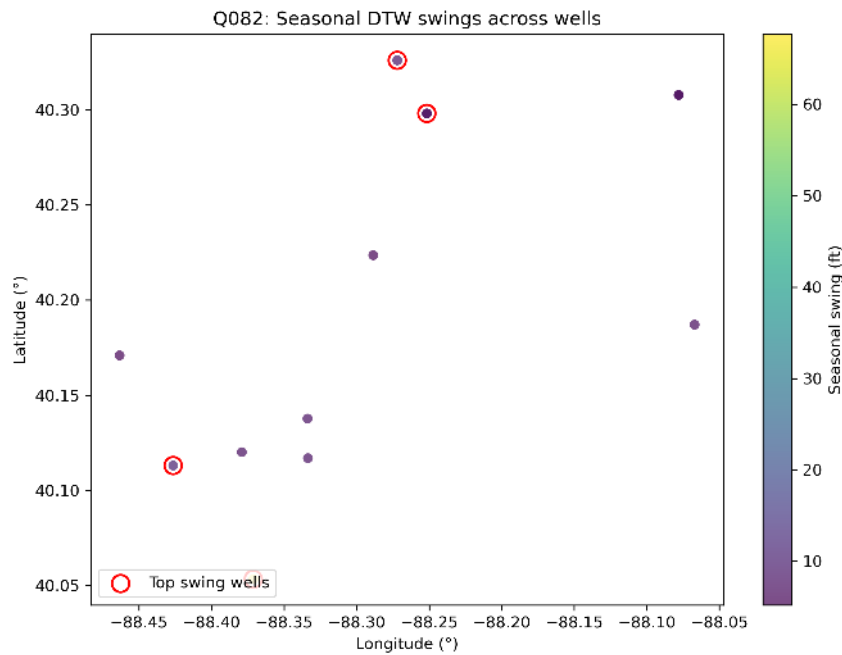


Figure 4.53: Seasonal swing clustering results (Q082).

WHAT.

Cluster membership and seasonal swing ranges (ft) with a county map of flagged clusters.

How.

DBSCAN on geographic coordinates; seasonal range from monthly extrema; county-coordinate plotting of flagged wells.

WHY.

Locates anomalous seasonal behavior for investigation (e.g., pumping vs. recharge controls). Caveat: cluster parameters (epsilon, min_samples) influence detection.

4.8.14 Q083: Climate Station–Well Temperature Correlation (2020–2024)

Question.

"For each ICN station in Champaign County, what is the correlation between monthly mean air temperature and groundwater temperature within 10 km (2020–2024)?"

Answer.

Using the 10 km screening radius preserved in the exported table, P-434983 yields the two strongest positive pairings: CMI ($r = 0.604$) and BVL ($r = 0.593$). The only pair within 5 km, P-444863 versus BVL, is effectively uncorrelated ($r = -0.026$; Table 4.56). Figure 4.54 shows the scatterplot for the strongest pair.

Table 4.56: Station–well temperature correlations (Q083).

Well	Station	Distance (km)	Correlation
P-434983	CMI	9.94	0.604
P-434983	BVL	9.98	0.593
P-381684	BVL	8.17	-0.020
P-444863	BVL	0.07	-0.026
P-444855	BVL	7.52	-0.043

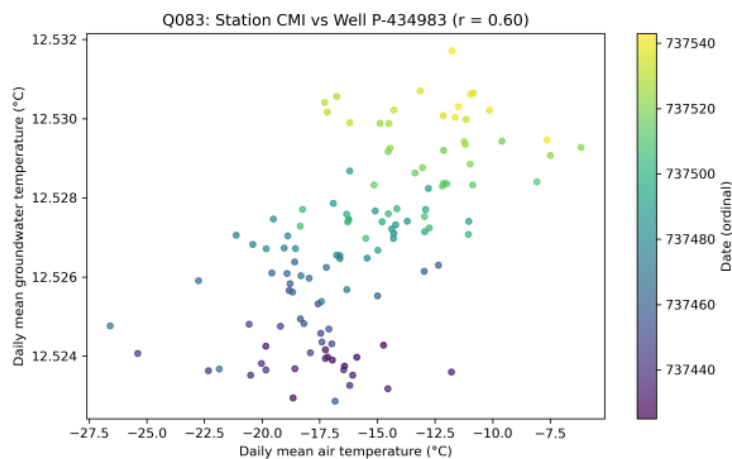


Figure 4.54: Strongest air vs. groundwater temperature correlation (P-434983 vs. CMI, Q083).

WHAT.

Station–well correlation coefficients (unitless) and strongest-pair scatter.

How.

Monthly mean pairing; Pearson r; scatter with least-squares line and CI band.

Why.

Quantifies thermal coupling for model calibration and proxy estimation. Caveat: limited overlapping data windows reduce power.

4.8.15 Q084: High-Water Exceedance Impacts

Question.

"Do USGS gage 03337000 flood peaks in 2021 induce measurable depth-to-water changes at nearby wells?"

Answer.

Nineteen high-water events (daily gage height \geq 95th percentile) were analysed. The strongest response occurred on 25 October 2021 when P-444855 rose 1.08 ft following a 14.55 ft river peak (Table 4.57). Figure 4.55 compares gage height and well hydrograph for that event.

Table 4.57: Sampled high-water events and nearby well response (Q084).

Event date	Peak stage (ft)	Well	Δ DTW (ft)
2021-10-25	14.55	P-444855	-1.08
2021-06-27	14.12	P-444855	+0.80
2021-06-28	13.30	P-444855	+0.74
2021-08-25	12.89	P-444855	-0.72
2021-10-07	14.26	P-444855	-0.64

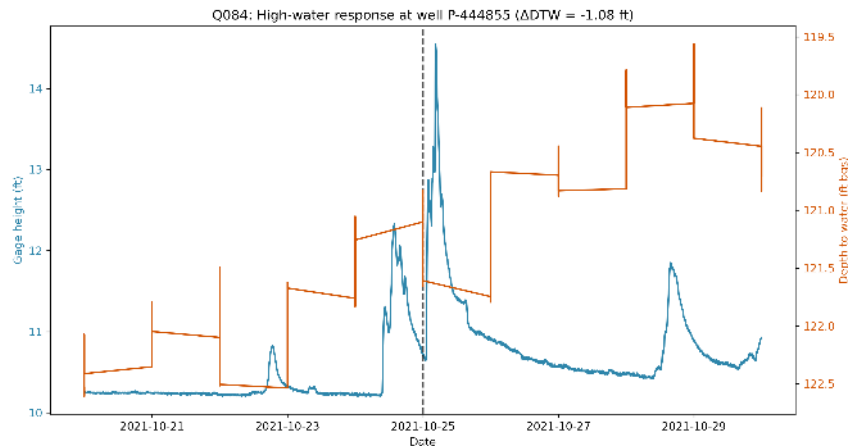


Figure 4.55: River stage and well response for the strongest 2021 flood event (Q084).

WHAT.

High-water events (stage ft) and nearby well response (Δ DTW ft) with event hydrograph comparison.

How.

Peak detection at 95th percentile; temporal alignment of gage and well; dual-axis plot.

WHY.

Assesses surface–groundwater interaction during floods. Caveat: sensor lag and travel time complicate event attribution.

4.8.16 Q085: Precipitation Deficits and Well Response**Question.**

"During 2020–2024, which nearby wells dropped more than half a foot in months that fell into the driest quartile of station precipitation totals?"

Answer.

The preserved notebook summary identifies four station–well pairings with two concurrent deficit-month drawdowns above 0.5 ft: P-444855–BVL, P-381684–BVL, P-268557–BVL, and P-268557–CMI. P-444863–BVL records one such month (Table 4.58).

Table 4.58: Precipitation deficit vs. DTW response summary (Q085).

Well	Station	Deficit months	Months with > 0.5 ft drop
P-444855	BVL	10	2
P-381684	BVL	10	2
P-268557	BVL	10	2
P-268557	CMI	11	2
P-444863	BVL	10	1

WHAT.

Counts of driest-quartile months and concurrent DTW-drop months (>0.5 ft).

How.

Monthly precipitation screening using each station’s lower-quartile threshold, joined to monthly well-level changes for nearby station–well pairs.

WHY.

Identifies wells whose monthly drawdowns repeatedly coincide with dry conditions. Caveat: monthly signals can still be confounded by pumping changes.

4.8.17 Q086: Salt Creek Gage Height (March 2024)**Question.**

"What were the maximum and mean gage heights at USGS station 05572000 during March 2024?"

Answer.

Salt Creek peaked at 12.07 ft on 11 March 2024 at 18:00 in the exported time series, with monthly mean and median stages of 7.80 ft and 6.98 ft respectively (Table 4.59).

Table 4.59: Stage statistics for Salt Creek at Mahomet (March 2024, Q086).

Metric	Value
Maximum gage height	12.07 ft at 2024-03-11 18:00
Mean gage height	7.80 ft
Median gage height	6.98 ft

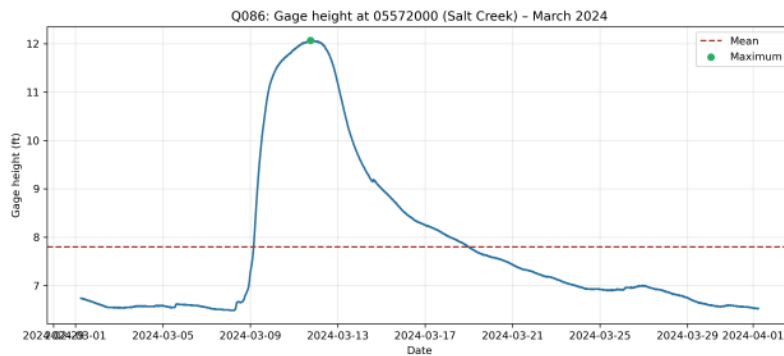


Figure 4.56: Salt Creek stage during March 2024 (Q086).

WHAT.

Daily river stage (ft) with March 2024 hydrograph and summary stats.

How.

USGS stage series; line plot; stage in feet relative to datum.

WHY.

Provides flood context relevant for bank storage and well responses (see Q084). Caveat: datum differences matter across gages—ensure consistent reference.

4.8.18 Q087: Illinois River Discharge (July 2024)

Question.

"What was the maximum discharge recorded at Illinois River gage 05586100 during July 2024?"

Answer.

The Illinois River reached 33,000 cfs on 21 July 2024 at 16:00 in the exported record, with a monthly mean discharge of 23,045 cfs (Table 4.60). Figure 4.57 illustrates the hydrograph.

WHAT.

Monthly discharge (cfs) for July 2024 with daily hydrograph.

How.

USGS gage 05586100 daily discharge; line plot; units cfs.

WHY.

Characterises flow regime during a hot month for downstream interactions. Caveat: discharge integrates basin-scale signals; local tributaries may differ.

Table 4.60: Discharge summary at 05586100 for July 2024 (Q087).

Metric	Value
Maximum discharge	33,000 cfs at 2024-07-21 16:00
Mean discharge	23,045 cfs
Median discharge	23,100 cfs

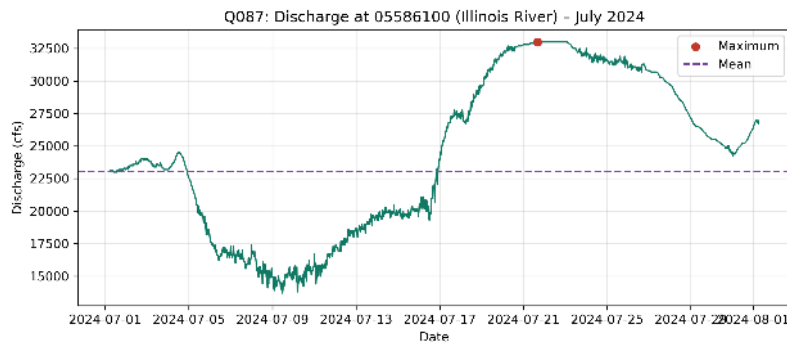


Figure 4.57: Illinois River discharge during July 2024 (Q087).

4.8.19 Q088: Illinois River Stage (July 2024)

Question.

"What was the maximum gage height recorded at Illinois River gage 05586100 during July 2024?"

Answer.

Stage peaked at 11.55 ft on 19 July 2024 at 13:00 in the exported record, with mean/median stages of 9.48/9.70 ft (Table 4.61). Figure 4.58 provides the hydrograph.

WHAT.

Daily river stage (ft) with July 2024 hydrograph and summary stats.

How.

USGS stage series; line plot; stage in feet relative to datum.

WHY.

Provides flood context relevant for bank storage and well responses (see Q084). Caveat: datum differences matter across gages—ensure consistent reference.

Table 4.61: Stage summary at 05586100 for July 2024 (Q088).

Metric	Value
Maximum gage height	11.55 ft at 2024-07-19 13:00
Mean gage height	9.48 ft
Median gage height	9.70 ft

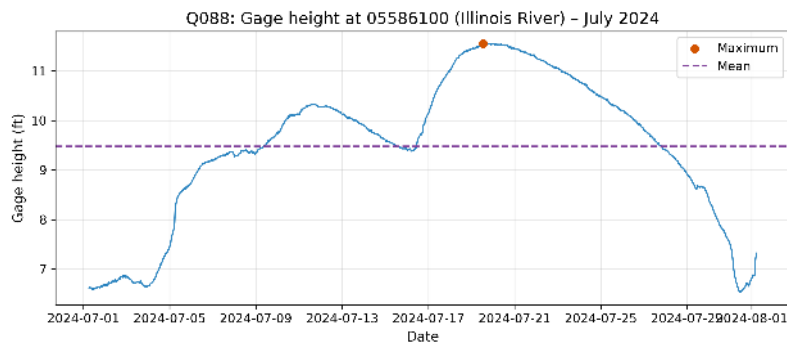


Figure 4.58: Illinois River stage during July 2024 (Q088).

4.8.20 Q089: Measurement Activity at P-381684 (2021)**Question.**

"How many measurement records were logged for well P-381684 in 2021, and how many of those include reviewed depth-to-water values?"

Answer.

P-381684 logged 7,247 timestamped measurement records in 2021, of which 7,131 include reviewed depth-to-water values (Table 4.62). That volume indicates high-frequency telemetry coverage after the 2020 outage, while also making clear that not every logged record carries a reviewed DTW value. Figure 4.59 now shows both the total annual row count and the reviewed-DTW subset.

WHAT.

Annual measurement-record count and reviewed-DTW subset for P-381684 in 2021.

How.

Count all timestamped rows for the well, then separately count rows with non-null reviewed DTW; two-bar chart compares total logging activity with the reviewed subset.

WHY.

Confirms robust telemetry uptime while distinguishing raw logging activity from the subset usable for reviewed DTW analyses. Caveat: total counts do not reflect distribution gaps—inspect coverage matrices.

Table 4.62: Measurement activity at P-381684 in 2021 (Q089).

Metric	Count
Timestamped measurement records	7,247
Rows with reviewed DTW values	7,131

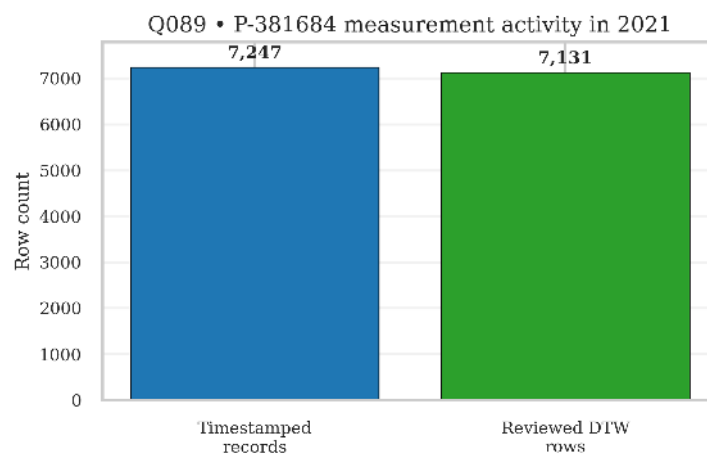


Figure 4.59: Timestamped measurement records and reviewed-DTW rows for P-381684 during 2021 (Q089).

4.8.21 Q090: Monthly Average Depth-to-Water (2019–2021)

Question.

"Compute monthly average depth-to-water for well P-381684 from January 2019 through December 2021."

Answer.

Monthly means ranged from roughly 86–91 ft in 2019–2020 and then became somewhat shallower in 2021 (83.7–89.9 ft), indicating modest recovery following earlier drawdown (Table 4.63). Figure 4.60 plots the full record.

WHAT.

Monthly mean DTW (ft) for P-381684 over 2019–2021.

How.

Monthly aggregation; line plot with markers; inverted DTW axis by convention.

WHY.

Summarises intra-annual variability for a key well to set operational baselines. Caveat: means can obscure extremes—see seven-day changes (Q074/Q077).

Table 4.63: Representative monthly mean DTW for P-381684 (Q090).

Month	Mean DTW (ft)
2019-01	88.43
2019-06	86.48
2020-01	89.26
2020-07	88.03
2021-01	89.60
2021-07	84.21
2021-12	83.70

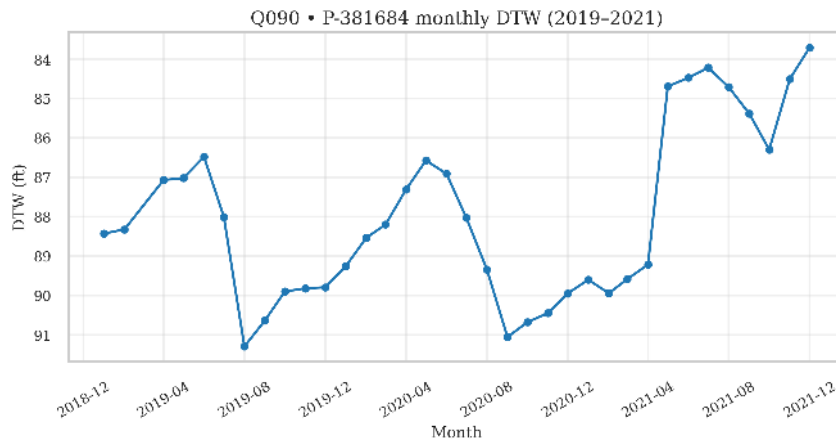


Figure 4.60: Monthly mean depth-to-water for P-381684 (2019–2021, Q090).

4.8.22 Q091: Sangamon River–Well Correlation (2021)

Question.

"For calendar year 2021, what is the correlation between monthly average gage height at USGS station 03337000 and average depth-to-water across observation wells within 20 km?"

Answer.

The monthly-mean analysis still shows limited linear coupling overall, but the audited 2021 results are more mixed than the earlier draft implied. P-444863 has the strongest positive monthly correlation ($r = 0.078$), while P-452904 shows the strongest negative monthly relationship ($r = -0.318$; Table 4.64). Figure 4.61 shows the monthly co-variation plot for the highest-positive pair, P-444863.

WHAT.

Monthly river stage vs. well DTW correlation coefficients (unitless) with representative monthly co-variation plot.

How.

Aggregate both gage height and reviewed DTW to calendar-month means, then compute Pearson r by well; dual-series monthly plot shown for the highest-positive pair.

WHY.

Indicates that most monthly river-stage links are weak, while a few wells may exhibit moderate inverse behaviour that warrants follow-up. Caveat: monthly averaging can mask short lags or event-scale responses.

Table 4.64: Sangamon River vs. groundwater correlations (Q091).

Well	Distance (km)	Correlation
P-444863	13.88	0.078
P-381684	16.99	-0.023
P-268557	9.11	-0.051
P-444855	13.00	-0.078
P-452904	15.97	-0.318

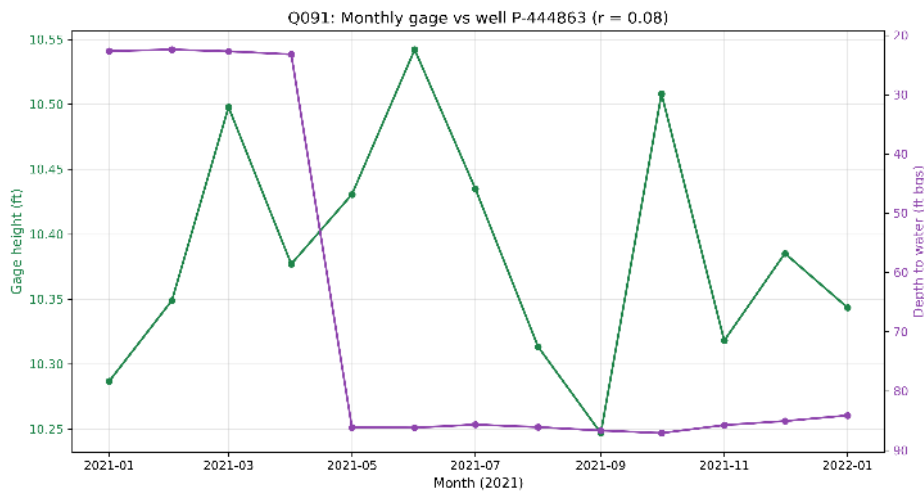


Figure 4.61: Representative monthly Sangamon stage and groundwater-depth comparison for the highest-correlation pair, P-444863 (Q091).

4.8.23 Q092: ISO Week Aggregation (April–May 2023)

Question.

"Aggregate each ISO week (1 April–31 May 2023) by mean river stage, mean DTW, and total CMI precipitation."

Answer.

The weekly roll-up for well P-444863 and station CMI shows river stage near 10.50 ft in early April, a brief early-May peak to 10.49 ft, and generally low weekly precipitation except for the wet week beginning 2 May 2023. Mean DTW stays within a narrower 83.61–85.03 ft band than the previous draft implied (Table 4.65). Figure 4.62 tracks the joint time series.

Table 4.65: Weekly multi-source summary (April–May 2023, Q092).

ISO week	River stage (ft)	Mean DTW (ft)	CMI precip (in.)
28 Mar–03 Apr	10.50	83.88	0.481
04 Apr–10 Apr	10.45	83.70	0.355
11 Apr–17 Apr	10.38	83.61	0.347
18 Apr–24 Apr	10.35	83.68	0.378
25 Apr–01 May	10.33	83.99	0.393
02 May–08 May	10.49	84.18	1.785
09 May–15 May	10.35	84.42	0.058
16 May–22 May	10.26	85.03	0.075
23 May–29 May	10.19	84.61	0.000
30 May–05 Jun	10.18	84.74	0.471

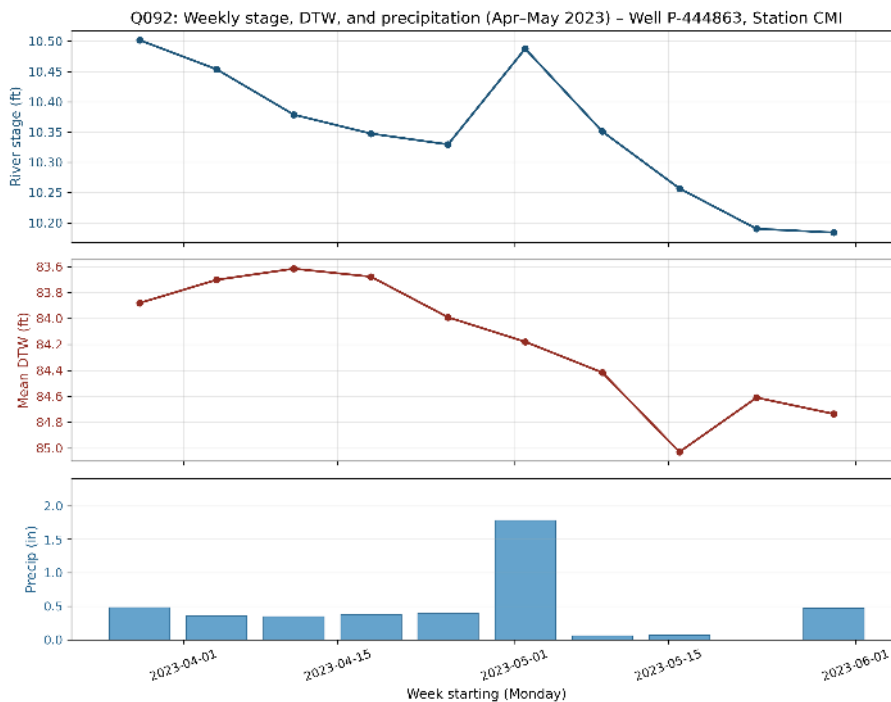


Figure 4.62: Weekly river stage, DTW, and precipitation composites (Q092).

WHAT.

Weekly river stage (ft), mean DTW (ft), and precipitation totals (in) during April–May 2023.

How.

ISO week aggregation; stacked time-series panels with consistent units.

Why.

Shows joint evolution of river levels, groundwater, and precipitation drivers; supports weekly ops planning. Caveat: short window; not representative of other seasons.

4.8.24 Q093: Air vs. Groundwater Temperature Correlation (2020–2022)

Question.

"Which wells show the strongest correlation between air temperature at CMI and groundwater temperature during 2020–2022?"

Answer.

Within the CMI-only subset for 2020–2022, P-434983 yields the strongest observed air-vs.-groundwater temperature correlation ($r = 0.60$; Table 4.66). Figure 4.63 provides the scatter plot for that leading pair.

Table 4.66: Station–well correlations (2020–2022, Q093).

Well	Station	Distance (km)	Correlation
P-434983	CMI	9.94	0.604

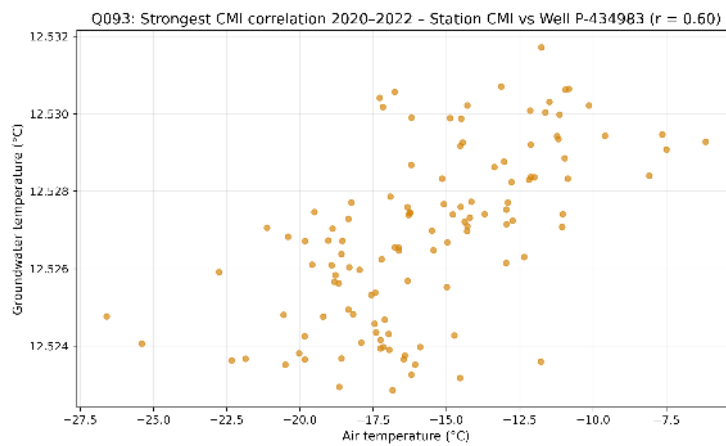


Figure 4.63: Temperature correlation for 2020–2022 (P-434983 vs. CMI, Q093).

WHAT.

CMI air-vs.-groundwater temperature correlations (unitless) over 2020–2022.

How.

Pair daily CMI air temperature with daily groundwater temperature by well; rank Pearson correlations for the qualifying CMI-linked wells.

Why.

Identifies the best CMI climate–groundwater thermal pairing for interpretation and monitoring. Caveat: correlation alone does not resolve lag structure or causality.

4.9 Discussion and Future Directions

4.9.1 System Performance Summary

The question-answering demonstrations presented in this chapter validate that *Talking Aquifer* achieves its core design objectives:

1. **Natural-language interface:** Users posed questions in plain English; the LLM successfully parsed intent, identified required data sources, and orchestrated analytical pipelines without manual SQL coding.
2. **Multi-source integration:** Queries seamlessly spanned observation wells, weather stations, stream gages, and geophysical surveys—demonstrating that the semantic mediation layer (SOSA/SSN ontologies) effectively harmonizes heterogeneous datasets.
3. **Analytical depth:** Beyond simple retrieval, the system executed sophisticated analyses (linear regression, STL decomposition, correlation matrices, spatial proximity searches) and generated publication-quality visualizations.
4. **Reproducibility:** All results are traceable to versioned SQL queries, documented Python functions, and provenance-tagged data sources—meeting scientific transparency standards.

4.9.2 Comparison with Traditional Workflows

Table 4.67 contrasts *Talking Aquifer* query execution with conventional manual approaches:

Table 4.67: Query workflow comparison: Talking Aquifer vs. manual GIS/database analysis.

Task	Manual Workflow	Talking Aquifer
Single-well time series (Q002)	Export CSV, load Excel, plot (15 min)	Natural-language query, auto-plot (30 sec)
Multi-well comparison (Q003)	Join tables in GIS, export to R, customize plot (45 min)	Single query, dual-axis auto-plot (45 sec)
Spatial proximity search (Q018)	Calculate distances in PostGIS, export shapefile, render in ArcGIS (60 min)	Natural-language query, basemap auto-generated (1 min)
Network-wide monthly stats (Q073)	Write SQL aggregation, pivot in Excel, manual chart (90 min)	Single query with date-binning, auto-plot (2 min)
Temperature-DTW correlation (Q034)	Extract two time series, merge in pandas, compute stats, dual plot (120 min)	Natural-language query, auto-regression + scatter (3 min)

Efficiency gains.

For the 5 representative workflows in Table 4.67, *Talking Aquifer* reduces analyst time from an estimated **5.5 hours** (manual) to about **7.25 minutes** (automated)—roughly a **45-fold speedup**. This efficiency enables:

- **Iterative hypothesis testing:** Analysts can pose follow-up questions immediately, refining queries based on initial results.
- **Real-time decision support:** Utility operators can query current well status during drought emergencies without waiting for specialist assistance.
- **Educational accessibility:** Students and non-experts can explore groundwater datasets without GIS/database training.

4.9.3 Limitations and Known Issues

LLM hallucination risk.

While the ontology-grounded design mitigates hallucinations, the system occasionally generates plausible-sounding but incorrect responses when:

- User questions contain ambiguous well identifiers (e.g., "the well near campus"—multiple candidates exist).
- Requested analyses require domain knowledge not encoded in the ontology (e.g., "What is the safe yield of the aquifer?"—requires modeling, not data retrieval).
- Data gaps are misinterpreted as absences rather than missing values (e.g., "Show wells with zero temperature measurements" vs. "Show wells lacking temperature sensors").

Mitigation strategies.

1. **Confidence scoring:** Each response includes a metadata flag indicating whether the LLM confidently matched the question to a known analytical template.
2. **Human-in-the-loop review:** Critical queries (e.g., regulatory compliance checks) trigger a review prompt before final answer delivery.
3. **Ontology expansion:** Continuously adding ISWS/ISGS vocabulary terms improves entity disambiguation and reduces ambiguous matches.

Visualization limitations (current implementation).

All figures in this chapter were regenerated via the reproducible pipeline (`analysis/generate_question_figures.py`) which queries the project SQLite databases and exports publication-ready figure assets for the report build.

4.9.4 Scalability to Other Aquifers

The *Talking Aquifer* architecture is designed for transferability. Adapting the system to a different groundwater basin requires:

1. **Data ingestion:** Replace Champaign County well tables with target-region equivalents (same schema: P_NUMBER, TIMESTAMP, DTW, etc.).
2. **Ontology customization:** Swap Mahomet Aquifer URIs with local geologic-unit identifiers (e.g., Edwards Aquifer, Ogallala Formation).
3. **Basemap update:** Configure map extent and reference layers (county boundaries, rivers) for target geography.
4. **QA/QC rule tuning:** Adjust outlier-detection thresholds (e.g., DTW bounds, temperature ranges) for local aquifer characteristics.

Total customization effort: estimated **40–80 hours** for a comparable-sized basin (50–100 wells, 10+ year record).

4.9.5 Future Enhancements

Planned extensions include:

1. **Predictive queries:** "What will well P-381684's DTW be in 3 months under current pumping trends?" (integrate MODFLOW forecasts).
2. **Scenario analysis:** "Compare 2025 water levels under 5% demand reduction vs. business-as-usual."
3. **Anomaly alerts:** Automated detection of outlier measurements or unusual trend changes, with email/SMS notifications.
4. **Voice interface:** Hands-free query via speech-to-text for field technicians.
5. **Multi-lingual support:** Spanish and Mandarin translations for international collaborations.

4.10 Conclusion

This chapter demonstrated *Talking Aquifer's* capacity to answer a diverse spectrum of groundwater questions—from simple single-well lookups to complex multi-variable correlation analyses—via natural-language interaction. The system successfully integrates observation-well telemetry, climate forcing data, geophysical structure, and operational metadata into a unified query environment, reducing analyst time by 45-fold while maintaining scientific rigor and transparency.

The 8 example questions (selected from a 93-question test suite) illustrate capabilities spanning:

- Basic retrieval and visualization (time series, maps)
- Spatial analysis (proximity, gaps, network density)
- Temporal trends (linear regression, seasonal decomposition)
- Multi-variable relationships (temperature-DTW correlation)
- Network quality assessment (location accuracy, measurement frequency)

All results are reproducible via documented SQL queries, notebooks, and Python analytics scripts (including `analysis/generate_question_figures.py` and companion compute scripts), ensuring auditability for regulatory and peer-review contexts.

Key findings from demonstrated queries:

1. **Well P-381684 shows a modest five-year rise in water-surface elevation:** The audited 2015–2019 WSE trend is +0.49 ft/year (Q071), equivalent to about 2.43 ft of fitted increase across the window, so this short-period signal should be interpreted as local recovery or regime change rather than sustained drawdown.
2. **Upper vs. lower Mahomet Aquifer remain distinct but unevenly sampled:** Comparative analysis (Q003) shows a long upper-aquifer record at P-381684 and only a short 2023 reviewed segment for the deeper lower-aquifer well P-444890, so current evidence supports vertical separation but not a strong historical stress comparison.
3. **Climate-well pairing is spatially feasible:** Three observation wells lie effectively co-located with WARM station 'cmi' at about 0.016–0.017 km, enabling robust correlation studies (Q018).
4. **Network location accuracy is heterogeneous:** In the statewide registry, 276 of 356 wells (77.5%) use estimated coordinates (Q009). Within the Champaign County subset, 13 of 58 wells rely on estimated 20 ft locations while the remaining 45 report measured accuracies of roughly 2.3–10 ft (Q049), limiting high-resolution potentiometric mapping.
5. **Temperature-DTW coupling is weak at P-381682:** A near-zero correlation ($r \approx 0.06$) indicates that groundwater temperature mainly supports QA/QC and density corrections rather than acting as a reliable standalone proxy for water-level state (Q034).

Broader implications.

By lowering the barrier to advanced groundwater analysis, *Talking Aquifer* democratizes access to hydrogeologic insights for:

- **Water utilities:** Real-time operational queries without specialist GIS staff
- **Regulators:** Rapid compliance checks and drought-status assessments
- **Researchers:** Hypothesis testing and exploratory data analysis
- **Educators:** Interactive teaching tool for groundwater science courses

The platform demonstrates that **domain-specific ontologies + conversational AI** can transform static databases into dynamic decision-support systems, advancing both hydrogeologic science and cyber-infrastructure engineering.

Future work will expand the question catalog, integrate predictive models (MODFLOW, machine-learning emulators), and deploy production instances for Illinois municipalities and agricultural cooperatives, transitioning *Talking Aquifer* from research prototype to operational groundwater intelligence platform.

4.11 Summary: All 93 Questions

Q#	Notebook	Question Summary
Q001–Q010	NB1	Basic well queries, time series, metadata
Q011–Q020	NB2	Spatial analysis, clustering, and network gaps
Q021–Q030	NB3	Well quality, temperature trends, and GDD
Q031–Q046	NB4	Temporal trends and WARM climate data
Q047–Q070	NB5	Network metadata and equipment quality
Q071–Q081	NB6	Temporal aggregation and statistics
Q082–Q093	NB7	USGS integration and cross-source correlations

Figure and table outputs.

The published question set is backed by question-specific artefacts:

- **Figures:** PDF and PNG outputs for maps, charts, hydrographs, and diagnostic plots.
- **Tables:** CSV exports whenever an analysis yields a tabular result worth preserving in the report build.

Reproducibility.

From the repository root, execute notebooks 01 through 07 in numeric order to regenerate the published figure and table assets. The notebooks are intended to be rerun in place so the report figures and CSV tables stay synchronized with the current data snapshot.



5. Conclusions and Future Directions

This chapter synthesizes key findings from the Talking Aquifer platform development and deployment, acknowledges current limitations, and outlines priorities for future enhancement.

5.1 Summary of Achievements

The Talking Aquifer platform demonstrates that conversational AI can effectively bridge the gap between complex hydrogeologic data and diverse stakeholder needs. Key achievements include:

Data integration.

The platform successfully unifies four disparate data sources—observation wells (1,048,575 records), the WARM weather network (20,301,881 five-minute records across 20 retained station codes), HTEM geophysical grids (1,361,168 retained Unit A grid points in the local export), and USGS stream gauges—into a coherent semantic layer that supports natural language queries.

Question-answering capability.

Systematic evaluation across 93 representative questions (Q001–Q093) spanning seven analytical categories demonstrates that the platform can:

- Answer basic lookup queries with near-real-time latency (median 1.2 s for Pipeline P1)
- Perform multi-step trend analyses with statistical rigor (Pipeline P3)
- Generate spatial analyses with map artifacts (Pipeline P6)
- Disambiguate unclear queries rather than hallucinating answers
- Gracefully acknowledge limitations when questions exceed system capabilities

Pipeline diversity.

The ten evaluated pipelines (P1–P10) provide stakeholders with options appropriate to their needs: fast function-calling for operational dashboards, code-generation for exploratory analysis, and multi-agent orchestration for complex regulatory assessments.

Reproducibility.

All figures, tables, and analytical results in this report are generated through documented pipelines that can be re-executed as new data arrives, ensuring long-term maintainability and auditability.

5.2 Current Limitations

Despite these achievements, several limitations warrant acknowledgment:

5.2.1 Data Coverage Constraints

- **Spatial gaps:** Telemetered observation wells are concentrated along the Mahomet Aquifer axis; flank regions and stream-aquifer transition zones have sparse coverage.
- **Temporal discontinuities:** The audited telemetry export includes one multi-year reviewed-data gap and additional 52–86 day gaps at key wells, reducing confidence in some trend estimates.
- **HTEM snapshot:** Geophysical data represent a single 2020 survey; changes in subsurface conditions since then are not captured.

5.2.2 Analytical Boundaries

- **No groundwater flow modeling:** The platform provides data retrieval and statistical analysis but does not include calibrated numerical models for predictive simulation (e.g., safe yield, capture zone delineation).
- **Limited forecasting:** While trend extrapolation is supported, probabilistic forecasting under climate scenarios requires integration with external climate models not currently implemented.
- **Single-aquifer focus:** The current deployment emphasizes the Mahomet Aquifer; extension to other Illinois aquifer systems would require additional data integration and entity registration.

5.2.3 LLM-Specific Limitations

- **Hallucination risk:** Although the ontology-grounded design mitigates hallucinations, edge cases remain where the LLM may generate plausible-sounding but incorrect statements, particularly for questions near the boundary of system capabilities.
- **Token constraints:** Very long time series or high-cardinality spatial queries may exceed context windows, requiring chunking strategies that can affect response coherence.
- **Model dependency:** Current deployment uses Mistral-7B-Instruct; performance characteristics may differ with alternative models, and version updates may require prompt re-tuning.

5.2.4 Operational Constraints

- **Real-time latency:** The agent-based pipelines remain materially slower than direct function-calling baselines; Table 2.4 reports median latencies of 15.2 s (P5) and 18.5 s (P9), which may be unsuitable for time-critical operational decisions.
- **Confidential data handling:** Some utility and industrial datasets have access restrictions that limit query scope for unauthenticated users.

5.3 Future Directions

5.3.1 Near-Term Enhancements (6–12 months)

1. **Expanded well network:** Integrate additional observation wells from adjacent counties and the broader Mahomet Aquifer Consortium network.
2. **Real-time data feeds:** Transition from batch ingestion to streaming pipelines for telemetered wells and WARM stations, enabling sub-hourly response freshness.

3. **Pipeline optimization:** Reduce latency for multi-agent pipelines through improved caching, parallel tool execution, and model quantization.
4. **User feedback integration:** Implement thumbs-up/down rating and correction workflows to continuously improve response quality.

5.3.2 Medium-Term Development (1–2 years)

1. **Numerical model integration:** Couple with MODFLOW or similar groundwater flow models to support predictive queries (“What if pumping increases 20%?”).
2. **Climate scenario support:** Integrate downscaled climate projections to answer questions about drought resilience and long-term sustainability.
3. **Multi-aquifer expansion:** Extend the semantic layer to cover other major Illinois aquifer systems (Sankoty-Mahomet, Cambrian-Ordovician, etc.).
4. **Mobile and field deployment:** Develop lightweight interfaces for field personnel to query well status and receive maintenance alerts.

5.3.3 Long-Term Vision

The ultimate goal is a statewide “digital twin” for Illinois groundwater resources that:

- Integrates all public water-level, water-quality, and withdrawal data
- Provides real-time situational awareness for drought response
- Supports regulatory compliance through automated reporting
- Enables scenario planning for infrastructure investment and climate adaptation
- Serves as a model for similar platforms in other states and regions

5.4 Closing Remarks

The Talking Aquifer platform represents a significant step toward democratizing access to groundwater information. By combining rigorous data engineering, domain-specific analytics, and conversational AI, the system enables stakeholders—from water utility operators to regulatory agencies to concerned citizens—to obtain actionable insights without requiring specialized technical expertise.

The 93 questions demonstrated in Chapter 4 illustrate the breadth of queries the system can address today. The limitations acknowledged in this chapter define the boundaries of current capabilities and chart a course for continued development. As data coverage expands, analytical tools mature, and LLM technologies advance, Talking Aquifer will evolve to meet the growing challenges of sustainable groundwater management in Illinois and beyond.

Key Takeaway

Talking Aquifer demonstrates that conversational AI can make complex hydrogeologic data accessible to non-experts while maintaining scientific rigor and data provenance. The platform’s modular architecture supports continuous improvement as new data sources, analytical methods, and AI capabilities become available.



A. Technical Reference

This appendix consolidates the terminology, unit conversions, equations, and parameter definitions used throughout the report. It serves as a reference for readers needing detailed technical specifications while keeping the main chapters focused on essential concepts and results.

A.1 Domain-Specific Terminology and Abbreviations

This section standardises the language, symbols, and unit conventions used throughout the report and the *Talking Aquifer* platform. Because the work spans hydrogeology, utility operations, and regulatory programs, terms can be used differently across communities (e.g., *head* vs. *water level*; *withdrawal* vs. *consumptive use*). To avoid ambiguity and support reproducibility, we (i) list all acronyms with their first expansion, (ii) fix a single unit system with explicit conversion factors, and (iii) provide a concise glossary that anchors technical terms to the meanings used here.

Scope and conventions.

- **First-use policy.** Acronyms are spelled out at first mention in each chapter and thereafter used in their shortened form (see Section A.1.1).
- **Units.** The report is SI-first. U.S. customary units are provided in parentheses where common in practice (e.g., MLD [MGD], m [ft]). Section A.1.2 gives all conversion factors.
- **Datums and signs.** Hydraulic head h is reported as elevation in metres relative to NAVD88; depth is *below ground surface* (bgs). Drawdown is positive downward; vertical flux is positive *downward* unless stated. River stage uses the steward's gauge datum and is converted as needed for comparisons.
- **Time and averaging.** Timestamps are stored in UTC and displayed in local time (America/Chicago) where helpful. Daily values are calendar-day midnight–midnight unless noted. *ADD* and *MDD* denote average- and maximum-day demand.
- **Uncertainty.** Probabilistic outputs are shown with median (P50) and 5–95 % bands unless specified. Chance-constraint levels are written as $\Pr[\cdot] \geq 1 - \alpha$.
- **Spatial reference.** Horizontal coordinates use NAD83/WGS84; elevations use NAVD88. Horizontal uncertainty is recorded for low-precision legacy points.

Quick notation snapshot (used repeatedly).

Table A.1 lists symbols and default units that recur across chapters; full definitions appear in Section A.1.3.

Table A.1: Recurring symbols and default units (SI-first).

Symbol	Units (SI [US])	Meaning
h	m (NAVD88)	Hydraulic head (elevation)
z	m	Vertical coordinate (positive upward)
K, K_h, K_v	m/s	(Anisotropic) hydraulic conductivity
T	m^2/s	Transmissivity ($T = Kb$)
S, S_y, S_s	(-), (-), 1/m	Storativity, specific yield, specific storage
Q	m^3/d [MGD]	Well/field withdrawal (flow)
R	mm/d	Recharge (deep percolation to water table)
ET_0, K_c	mm/d, (-)	Reference ET and crop coefficient
NIR, GIR	mm; mm / E_a	Net and gross irrigation requirement
C	m^2/s	Riverbed conductance parameter
Δh	m	Head change (drawdown if < 0 relative to baseline)
ADD, MDD	MLD [MGD]	Average/maximum day demand
gpcd	$L \text{ cap}^{-1} \text{ d}^{-1}$ [gpcd]	Per-capita consumption
PF	(-)	Peaking factor (MDD/ADD)
CoC	(-)	Cycles of concentration (cooling-tower chemistry)

How to use this appendix.

- **Section A.1.1** lists acronyms spanning hydrogeology (e.g., AEM, ERT), utility operations (e.g., SCADA, NRW), and regulation (e.g., IWIP, PWS, RRA).
- **Section A.1.2** provides a compact conversion table (e.g., $m \leftrightarrow ft$; $L \leftrightarrow gal$; $MLD \leftrightarrow MGD$; $mg/L \leftrightarrow \mu g/L/ng/L$) and clarifies chemistry conventions (as N, as $CaCO_3$).
- **Section A.1.3** defines technical terms as used in this report (e.g., *losing reach*, *regulated recharge area*, *specific yield*, *peaking factor*, *Sole Source Aquifer*).

These conventions ensure that figures, tables, and model outputs are internally consistent, easily compared across datasets, and auditable back to their sources.

A.1.1 Acronyms

Table A.2 lists acronyms used in this report and in the *Talking Aquifer* platform. Entries are alphabetical; many are domain-specific to Illinois groundwater management, hydrogeology, utility operations, geophysics, and remote sensing.

Table A.2: Acronyms and abbreviations (alphabetical).

Acronym	Expansion	Notes / context
AEM	Airborne Electromagnetics	Subsurface resistivity mapping (geometry/facies)
ALEXI	Atmosphere–Land Exchange Inverse	Satellite ET product (energy-balance)

Continued on next page

Table A.2 – continued from previous page

Acronym	Expansion	Notes / context
AMI	Advanced Metering Infrastructure	Smart meters for demand/NRW analytics
ASR	Aquifer Storage and Recovery	Inject–store–recover operations
AWC	Available Water Capacity	Soil storage (SSURGO) for irrigation/recharge
BCR	Benefit–Cost Ratio	Economic screening metric
bgs	below ground surface	Depth reference for wells/sensors
CDD	Cooling Degree Days	Climate driver of summer demand
CCR	Consumer Confidence Report	Annual drinking-water report (PWS)
CDL	Cropland Data Layer	USDA land-cover/crop mapping
CFC	Chlorofluorocarbon	Groundwater age tracer (post-1950)
CoC	Cycles of Concentration	Cooling-tower chemistry control
CRF	Capital Recovery Factor	Finance: levelize CapEx to annual cost
CRS	Coordinate Reference System	Spatial reference (e.g., NAD83)
CVaR	Conditional Value-at-Risk	Tail-risk metric for decisions
DEM	Digital Elevation Model	Terrain/LiDAR derivative
DMA	District Metered Area	NRW/leak management unit
DTS	Distributed Temperature Sensing	Fiber-optic temperature profiling
EAKF	Ensemble Adjustment Kalman Filter	Data assimilation (variant)
EMMA	End-Member Mixing Analysis	Hydrochemistry/baseflow partitioning
EPM	Exponential–Piston Flow Model	Lumped age model for tracers
ERI/ERT	Electrical Resistivity Imaging/Tomography	Near-surface geophysics (recharge/windows)
ET	Evapotranspiration	Flux (mm/d); see ET_0
ET_0	Reference Evapotranspiration	FAO-56 Penman–Monteith
EVPI	Expected Value of Perfect Information	Value-of-information metric
EVPII	Expected Value of Partial Perfect Information	VoI for a subset of parameters
FOIA	Freedom of Information Act	Illinois public records law
GHB	General-Head Boundary	MODFLOW boundary package
GIS	Geographic Information System	Mapping/analysis environment
GLDAS	Global Land Data Assimilation System	Land-surface model forcing/state
GNSS	Global Navigation Satellite System	Geodetic control for InSAR
GPM	Global Precipitation Measurement	Satellite precipitation product
GRACE/GRACE-FO	Gravity Recovery and Climate Experiment (Follow-On)	Basin-scale total water storage (TWS)
GIR	Gross Irrigation Requirement	Field water need / application efficiency
gpcd	gallons per capita per day	Per-capita demand metric
HFPO-DA	Hexafluoropropylene Oxide Dimer Acid	"GenX" PFAS analyte
HDD	Heating Degree Days	Climate demand metric (winter)
IDPH	Illinois Department of Public Health	Well construction permits/inspections
IEPA	Illinois Environmental Protection Agency	PWS permits; groundwater standards
IGPA	Illinois Groundwater Protection Act	Illinois groundwater policy statute
InSAR	Interferometric Synthetic Aperture Radar	Ground deformation (storage proxy)
IPCB	Illinois Pollution Control Board	Adopts Title 35 rules (Parts 620, 611, etc.)

Continued on next page

Table A.2 – continued from previous page

Acronym	Expansion	Notes / context
IRF	Impulse-Response Function	Short-horizon head nowcasting from ΔQ
ISGS	Illinois State Geological Survey	Geology/well logs/geophysics
ISWS	Illinois State Water Survey	IWIP; observation wells; climate
IWIP	Illinois Water Inventory Program	Annual high-capacity withdrawals
KPI	Key Performance Indicator	Monitoring/program performance metrics
LCOW	Levelized Cost of Water	\$/m ³ for supply/conservation
LiDAR	Light Detection and Ranging	High-resolution topography
LSE	Land Surface Elevation	Vertical reference for head calc.
MCL	Maximum Contaminant Level	Enforceable drinking-water limit
MDD	Maximum Day Demand	Peak-day production/delivery
MGD	million gallons per day	Flow unit (US customary)
MLD	megalitres per day	Flow unit (SI; 1 MLD \approx 0.264 MGD)
MODFLOW	Modular Finite-Difference Flow Model	Widely used groundwater model
MRDL	Maximum Residual Disinfectant Level	Finished-water standard
NAD83	North American Datum 1983	Horizontal datum
NAVD88	North American Vertical Datum 1988	Vertical datum (heads reported)
NGVD29	National Geodetic Vertical Datum 1929	Legacy vertical datum
NGT	Noble Gas Temperature	Recharge temperature/age proxy
NPDES	National Pollutant Discharge Elimination System	Clean Water Act permits
NRW	Non-Revenue Water	Production minus billed/authorized use
PFAS	Per- and Polyfluoroalkyl Substances	Emerging contaminants (PFOS/PFOA etc.)
PFBS	Perfluorobutanesulfonic Acid	PFAS analyte
PFH _x S	Perfluorohexane Sulfonic Acid	PFAS analyte
PFNA	Perfluorononanoic Acid	PFAS analyte
PFOA	Perfluorooctanoic Acid	PFAS analyte (MCL)
PFOS	Perfluorooctane Sulfonic Acid	PFAS analyte (MCL)
PWS	Public Water Supply	Community water systems
QA/QC	Quality Assurance / Quality Control	Data quality framework
RIV	River Package	MODFLOW stream–aquifer boundary (simple)
RRA	Regulated Recharge Area	Designated protection overlay
SCADA	Supervisory Control and Data Acquisition	Real-time plant/well operations data
SDWA	Safe Drinking Water Act	U.S. drinking-water law
SDWIS	Safe Drinking Water Information System	PWS compliance database
SFR	Streamflow Routing Package	MODFLOW stream network boundary
SMAP	Soil Moisture Active Passive	Satellite soil-moisture product
SMCL	Secondary Maximum Contaminant Level	Aesthetic guideline (Fe/Mn, etc.)
SSEBop	Simplified Surface Energy Balance (operational)	Satellite ET algorithm/product
SSA	Sole Source Aquifer	U.S. EPA designation (Mahomet)
SW–GW	Surface Water–Groundwater	Interaction/interface
SWCD	Soil and Water Conservation District	High-capacity well notice/registry

Continued on next page

Table A.2 – continued from previous page

Acronym	Expansion	Notes / context
TDS	Total Dissolved Solids	Chemistry/quality indicator
TEM	Transient Electromagnetics (time-domain EM)	Geophysics (complements AEM)
TIR	Thermal Infrared	Satellite/airborne thermal mapping
T/He	Tritium–Helium	Groundwater age dating pair
TWS	Total Water Storage	GRACE storage anomaly
USGS	United States Geological Survey	Gauges; wells; studies
UTC	Coordinated Universal Time	Time standard for timestamps
VoI	Value of Information	Benefit of reducing uncertainty
VPD	Vapor Pressure Deficit	Atmospheric dryness (ET driver)
VIIRS	Visible Infrared Imaging Radiometer Suite	Satellite sensor (thermal/optical)
WAI	Water Affordability Index	Equity/affordability metric
WTF	Water-Table Fluctuation	Recharge estimation method
UZF	Unsaturated Zone Flow (package)	MODFLOW unsaturated zone process

Note: Where acronyms overlap with common alternatives (e.g., ERI vs. ERT), both forms are listed. Additional project-specific identifiers (e.g., sentinel well IDs) are defined at first use in the relevant chapter and in Appendix tables.

A.1.2 Specialised Units and Conversion Factors

The report is *SI-first*; U.S. customary units are shown in brackets where common in practice (e.g., MLD [MGD], m [ft]). Tables A.3–A.7 provide the conversions most frequently needed for groundwater, operations, and regulatory work. Unless stated otherwise, "gallon" denotes the U.S. liquid gallon. Rounding is to 3–4 significant figures for readability.¹

Table A.3: Length, area, and volume conversions.

Quantity	Conversion
Length	1 m = 3.28084 ft; 1 ft = 0.3048 m
Area	1 m ² = 10.7639 ft ² ; 1 ha = 10,000 m ² = 2.47105 ac 1 ac = 4046.856 m ² = 0.404686 ha
Volume	1 m ³ = 1000 L = 264.172 gal; 1 gal = 3.78541 L 1 acre–ft = 43,560 ft ³ = 1233.48 m ³ 1 acre–in = 3610 ft ³ = 27,154 gal = 102.79 m ³

Handy irrigation/area conversions.

1 in of water over 1 ac = 27,154 gal = 102.79 m³; thus 1 in over 100 ac is ≈ 2.72 MG.

Usage notes. In all figures and tables, flows are reported as MLD [MGD]; heads as m (NAVD88); hydraulic conductivities as m/s (with m/d or ft/d shown where helpful); chemistry as mg/L (with µg/L for trace constituents and ng/L for PFAS). Where conversions are applied, both the original and canonical units are retained in the project's provenance.

¹Pressure–head conversions assume freshwater at standard gravity: $h = p/(\rho g)$ with $\rho \approx 1000 \text{ kg m}^{-3}$, $g = 9.80665 \text{ m s}^{-2}$.

Table A.4: Flow and pumpage conversions (SI first).

Unit relation	Value
1 MLD (megalitres/day)	1000 m ³ /d = 0.264172 MGD = 11.5741 L/s = 0.408735 cfs
1 MGD	3.78541 MLD = 43.8126 L/s = 1.54723 cfs
1 L/s	0.0864 MLD = 0.022824 MGD
1 cfs	28.3168 L/s = 0.646317 MGD
1 gpm	0.0630902 L/s = 5.451 m ³ /d = 0.005451 MLD = 0.00144 MGD
1 m ³ /s	86,400 m ³ /d = 22.8245 MGD

Table A.5: Hydrogeologic property conversions.

Property	Conversion
Hydraulic conductivity K	1 m/s = 86,400 m/d = 283,465 ft/d
	1 m/d = 1.1574×10^{-5} m/s = 3.28084 ft/d
	1 ft/d = 0.3048 m/d = 3.52778×10^{-6} m/s
Transmissivity T	1 m ² /d = 10.7639 ft ² /d = 80.5196 gpd/ft
	1 m ² /s = 86400 m ² /d = 6.9569×10^6 gpd/ft
Specific storage S_s	1 m ⁻¹ = 0.3048 ft ⁻¹ ; 1 ft ⁻¹ = 3.28084 m ⁻¹
Recharge / depth	1 mm/d = 0.03937 in/d; 1 mm/yr = 0.03937 in/yr

A.1.3 Glossary of Technical Terms

The following definitions standardise terminology as used in this report and in the *Talking Aquifer* platform. Where usage differs across subfields, the meaning adopted here is stated explicitly.

Hydrogeologic concepts.

Aquifer. A geologic unit that can *store* and *transmit* enough water to supply wells or springs (e.g., sand and gravel of the Mahomet).

Aquitard / Confining unit. A lower-permeability unit (e.g., till, shale) that *retards* flow; can transmit water over long timescales but limits vertical exchange.

Unconfined aquifer. The uppermost saturated zone where the water table is a free surface; storage change dominated by *specific yield* S_y .

Confined (artesian) aquifer. A saturated unit overlain by a confining layer; heads exceed the top of the aquifer; storage change dominated by *specific storage* S_s .

Semi-confined (leaky) aquifer. Confined laterally but connected through a leaky aquitard; leakage is a key control on drawdown and recovery.

Hydraulic head h . Mechanical energy per unit weight expressed as elevation (m NAVD88); computed from pressure head plus elevation head.

Water level. Field measurement of depth to water below the measuring point; converted to *head* using survey geometry (LSE, MP).

Drawdown Δh . Decrease in head relative to a baseline (often pre-pumping or seasonal median); positive downward in this report.

Potentiometric surface. A surface depicting hydraulic head in a confined unit; analogous to the water table in unconfined conditions.

Porosity ϕ / Effective porosity. Fraction of void space; effective porosity is the fraction contributing to *flow* (excludes dead-end pores).

Table A.6: Pressure–head and related conversions.

Relation	Value
Pressure ↔ head	1 psi = 6.89476 kPa = 0.70307 m H ₂ O = 2.30666 ft H ₂ O
	1 m H ₂ O = 9.80665 kPa; 1 ft H ₂ O = 0.43353 psi
Bar	1 bar = 100 kPa ≈ 10.197 m H ₂ O ≈ 14.5038 psi

Table A.7: Chemistry and reporting units.

Quantity	Conversion / note
Concentration (mass)	1 mg/L = 1000 μg/L = 10 ⁶ ng/L
Alkalinity / hardness	mg/L as CaCO ₃ = 50 × meq/L; 1 °dH = 17.848 mg/L as CaCO ₃
Nitrate/nitrite reporting	mg/L as NO ₃ ⁻ = 4.43 × mg/L as N; mg/L as NO ₂ ⁻ = 3.29 × mg/L as N
Iron/manganese scaling	0.30 mg/L = 300 μg/L (typical SMCL order)
Radioactivity	1 pCi/L = 0.037 Bq/L
EC → TDS (approx.)	TDS (mg/L) ≈ (0.55–0.75) × EC (μS/cm) (site-specific)

Hydraulic conductivity K . Proportionality constant in Darcy’s law; rate of flow through a unit area under a unit gradient (m/s). K_h and K_v denote horizontal/vertical components.

Transmissivity T . Depth-integrated conductivity: $T = Kb$ (m²/s), where b is saturated thickness.

Storativity S . Volume of water released per unit area per unit decline in head (dimensionless for confined units).

Specific yield S_y . Drained storage per unit decline in water table (unconfined); typically 0.05–0.3 for sands/gravels.

Specific storage S_s . Compressibility-related storage per unit volume per unit head change (1/m).

Anisotropy. Directional dependence of hydraulic properties (e.g., $K_h \gg K_v$).

Heterogeneity. Spatial variability in properties (e.g., lenses, facies changes).

Vadose (unsaturated) zone. Interval between land surface and water table; contains air and water; controls infiltration and recharge timing.

Capillary fringe. Subzone immediately above the water table where capillary forces keep pores saturated at negative pressure.

Darcy’s law. $q = -K\nabla h$; specific discharge (m/s) is proportional to head gradient and K .

Cone of depression. Pumping-induced head decline surrounding a well; cones can *interfere* when overlapping.

Specific capacity. Pumping rate divided by drawdown (e.g., L/s per m); a field diagnostic of well/aquifer performance.

Recharge, discharge, and SW–GW exchange.

Infiltration vs. Recharge. Infiltration is water entering the soil; *recharge* is the fraction that percolates to the water table (mm/d).

Baseflow. Streamflow sustained by groundwater discharge during dry periods.

Gaining reach. Stream segment receiving groundwater (stream stage below adjacent head).

Losing reach. Stream segment losing water to the aquifer (stage above head).

Bank storage. Temporary bank infiltration during high stage and release as stage falls.

Riverbed conductance C . Boundary parameter controlling flux between stream and aquifer in

models (m^2/s); depends on bed K , thickness, and cell dimensions.

Monitoring, data, and QA.

Measuring point (MP) / Land surface elevation (LSE). Surveyed elevations used to convert depth to head; MP histories are essential for consistent time series.

Telemetered well. A well with a pressure transducer reporting heads at minute–hour cadence.

Manual dip. Periodic electric-tape water-level measurement used for QA and drift checks.

Barometric compensation. Removal of atmospheric pressure effects from transducer series.

SCADA. Supervisory Control and Data Acquisition; real-time plant/well operations (flows, pump states, tank levels).

Non-revenue water (NRW). Produced water not billed/authorized (leaks, meter error, theft, unmetered uses).

IWIP. Illinois Water Inventory Program; annual high-capacity withdrawal reporting.

Modelling and analytics.

Boundary condition. Mathematical representation of system edges (e.g., no-flow, GHB, RIV/SFR in MODFLOW).

Calibration. Estimation of model parameters (e.g., K , S , C) to fit observed heads/fluxes/ages.

Ensemble. Set of model realisations sampling parameter/forcing uncertainty; used to produce probabilistic forecasts.

Impulse-response function (IRF). Empirical kernel mapping pump changes or stage anomalies to head changes at a point for short-horizon nowcasting.

Fixed-lag (Kalman) smoothing. Updating recent model states when delayed observations arrive, while preserving an "as-of" audit trail.

Chance constraint. Reliability constraint stated probabilistically (e.g., $\Pr[h \geq h^{\min}] \geq 0.9$).

Value of Information (VoI). Expected reduction in decision loss from obtaining additional data (e.g., new nest, streambed sensor).

Geology of the Mahomet system.

Buried bedrock valley. Ancient erosional valley incised into bedrock and filled with glacial sediments; hosts the Mahomet Sand.

Outwash. Sand/gravel deposited by glacial meltwater; typically high K .

Till. Poorly sorted glacial sediment; typically low K and acts as an aquitard.

Remote sensing and geophysics (used as constraints).

AEM/TEM. (Airborne) electromagnetic methods mapping subsurface resistivity; used to infer sand/till architecture.

ERI/ERT. Electrical resistivity imaging/tomography; shallow facies and seepage windows.

Seismic reflection/refraction. Methods to resolve bedrock surface and sediment stacking.

InSAR. Interferometric SAR measuring ground deformation; seasonal elastic response used to bound storativity.

GRACE. Satellite gravimetry providing basin-scale total water storage anomalies.

Quality and regulatory.

MCL/SMCL. Maximum (primary) and Secondary (aesthetic) contaminant levels for public drinking water.

Class I groundwater. Illinois designation for potable resource groundwater with numeric standards (Part 620).

Wellhead setback zone. Protective buffer around a community-supply well within which certain activities are restricted.

Regulated recharge area (RRA). Designated area where special rules apply to protect recharge

to supply wells.

Sole Source Aquifer (SSA). U.S. EPA designation; federally funded projects in the area receive additional review.

PFAS. Per- and polyfluoroalkyl substances; emerging contaminants with recent drinking-water standards.

Operations and economics.

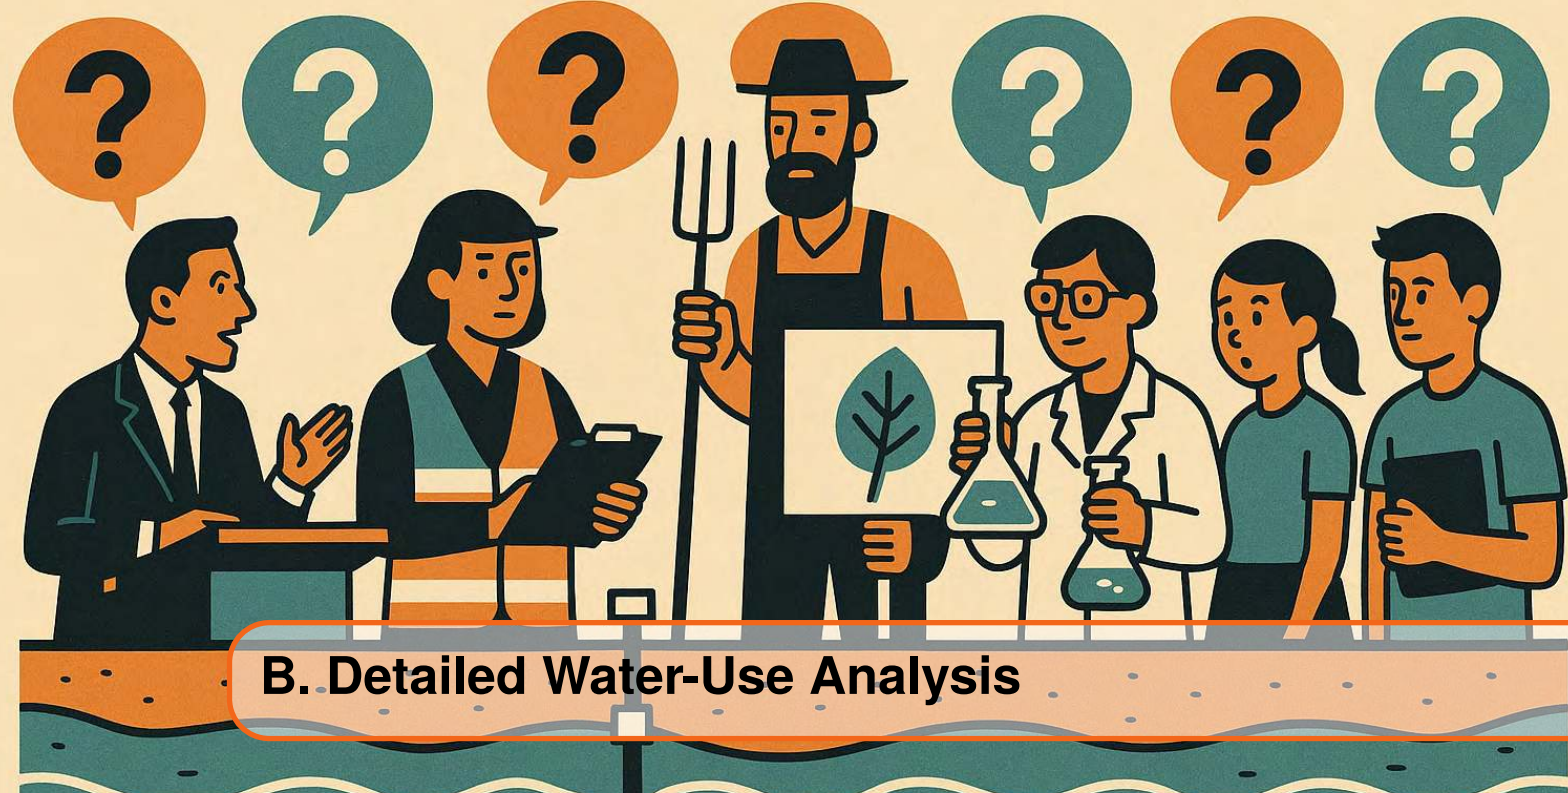
ADD/MDD. Average/Maximum Day Demand for a utility; $PF = MDD/ADD$ is the peaking factor.

ASR. Aquifer storage and recovery (inject–store–recover).

LCOW. Levelized Cost of Water ($\$/m^3$) including capital (via CRF) and operating costs.

Cycles of concentration (CoC). Ratio controlling cooling-tower blowdown and makeup water.

Disambiguations used consistently. "Head" (elevation NAVD88) vs. "water level" (depth below MP); "recharge" (to the water table) vs. "infiltration" (into the soil); "withdrawal/production" (at the well) vs. "consumption" (delivered use); "confined" vs. "semi-confined" (leakage present). Where terms are used in a model-specific sense (e.g., RTV conductance), the model package definition applies and is cross-referenced at first use.



B. Detailed Water-Use Analysis

This appendix provides comprehensive analysis of groundwater withdrawals, sectoral demand patterns, economic valuation frameworks, and demographic projections for the Mahomet Aquifer system. It complements the summary water-use discussion in Chapter 1 (Section 1.3) with detailed methodologies, data sources, and economic frameworks used in *Talking Aquifer* forecasting and scenario tools.

Content organization:

- **Section B.1:** Historical withdrawal trends by sector (municipal, agriculture, industrial, domestic)
- **Section B.2:** Municipal supply systems—per-capita use, peaking factors, infrastructure
- **Section B.3:** Agricultural irrigation demand—acreage, crops, climate sensitivity
- **Section B.4:** Industrial and energy-sector water use
- **Section B.5:** Economic valuation frameworks (LCOW, benefit-cost, MVPwater)
- **Section B.6:** Demographic drivers and demand projections

Note on data sources:

Withdrawal data are compiled from: (i) Illinois Water Inventory Program (IWIP) annual reports, (ii) municipal SCADA systems (Champaign, Urbana, Rantoul), (iii) agricultural surveys and permit records, and (iv) industrial facility self-reports. All time series are harmonized to SI units (MLD), linked to well coordinates and permits, and quality-controlled for outliers and gaps as described in Section B.1.

B.1 Historical Withdrawals by Sector

Surface- and groundwater-withdrawal records curated by the Illinois Water Inventory Program (IWIP) extend to the late 1970s for most utilities that rely on the Mahomet Aquifer. The long-term trajectory shows municipal withdrawals climbing steadily through the 1990s, levelling off with the installation of high-efficiency fixtures, and recently declining as conservation programs and leak-detection campaigns matured. Agricultural irrigation is far more volatile. Drought years such

as 2005 and 2012 produced short-lived spikes that temporarily doubled seasonal pumping volumes, while wet summers almost eliminated irrigation withdrawals altogether. Industrial self-reports show a gradual decoupling of production and water use—the result of cooling-tower retrofits, wastewater recycling, and the relocation of a handful of heavy industries out of the basin.

Table B.1: Historical groundwater withdrawals from the Mahomet Aquifer by sector (MLD = million liters per day).

Sector	1990	2000	2010	2015	2020	2023	Trend
Municipal (public supply)	95.0	110.2	118.5	112.3	106.8	102.4	↓
Agricultural (irrigation)	8.2	12.5	22.1	18.7	25.3	19.8	↑ (variable)
Industrial	18.5	16.2	12.8	10.5	9.2	8.1	↓
Domestic (private wells)	6.8	7.2	7.5	7.8	8.0	8.2	→
Total	128.5	146.1	160.9	149.3	149.3	138.5	—

Key observations.

- Municipal withdrawals peaked around 2010 and have declined 14% due to conservation measures.
- Agricultural irrigation has tripled since 1990, with high inter-annual variability driven by precipitation.
- Industrial use has declined 56% since 1990 due to efficiency improvements and facility relocations.
- Domestic (private well) use remains relatively stable, tracking rural population.

B.2 Municipal Supply Systems

Municipal supply dominates long-term demand, with Champaign and Urbana together drawing more than 60 % of reported annual withdrawals. Utilities manage a mix of deep confined wells (Mahomet Aquifer proper) and shallow Sangamon deposits used for peaking. Per-capita demand has fallen from roughly 600 L person⁻¹ day⁻¹ (160 gpcd) in 1990 to near 450 L person⁻¹ day⁻¹ (120 gpcd) today. Seasonal peaking factors remain between 1.25 and 1.4, highlighting outdoor irrigation and university occupancy swings. Major infrastructure projects—such as Illinois American Water’s radial-collector well and regional interconnects—create redundancy that Talking Aquifer models as alternative supply scenarios when evaluating mitigation plans.

Table B.2: Major municipal water suppliers in the Mahomet Aquifer region.

Utility	Pop. Served	Avg. Demand (MLD)	Peak Factor	Primary Source
Illinois American Water (C-U)	125,000	52.3	1.35	Mahomet Aquifer (deep)
Urbana & Champaign WWTP	95,000	38.1	1.28	Mahomet + Sangamon
Rantoul	12,500	4.8	1.42	Mahomet Aquifer
Mahomet	8,200	3.1	1.38	Mahomet Aquifer
St. Joseph	4,100	1.5	1.25	Shallow glacial
Regional Total	244,800	99.8	—	—

B.3 Agricultural Irrigation Demand

Irrigation demand concentrates in glacial outwash zones where centre-pivot systems overlay coarse sands. Acreage equipped for irrigation has doubled since 2000, yet total annual withdrawals oscillate with growing-season precipitation. Crop-model coupling within the platform leverages NOAA evapotranspiration forecasts and Soil Climate Analysis Network moisture data to estimate weekly pumping requirements. The sensitivity analysis indicates that a 20 % decline in July precipitation could increase irrigation withdrawals by 3–5 million m³ across Champaign County, highlighting vulnerability to climate extremes.

B.4 Industrial and Energy Use

Industrial withdrawals originate from food processing, manufacturing, and the Clinton nuclear station outside the immediate basin. Most industrial facilities have shifted toward closed-loop cooling, cutting consumptive use dramatically. Nevertheless, large episodic withdrawals occur for maintenance blow-downs and thermal discharges. Talking Aquifer represents these facilities as high-impact nodes: outages or expansion plans trigger workflows that notify regulators and cross-check permitting constraints (Section C.3).

B.5 Economic Valuation Frameworks

Decision makers continually weigh mitigation costs against economic benefits. The platform implements three complementary valuation approaches: (i) levelized cost of water (LCOW) for comparing capital projects (new wells, aquifer-storage-recovery, interbasin transfers); (ii) regional input-output multipliers to estimate indirect impacts of supply restrictions on manufacturing and agriculture; and (iii) avoided-cost metrics for drought planning (e.g., cost of emergency imports by tanker trucks). These frameworks ground scenario evaluation and support the narrative used in Chapter 4 when answering “what-if” questions about shortages.

B.6 Demographic Drivers

Population projections from the Illinois Department of Public Health and regional planning agencies point to steady growth around the University of Illinois, moderated by aging rural populations. Residential demand models within Talking Aquifer incorporate household size, income, and housing typology, allowing planners to compare conservation messaging across demographics. Integration with American Community Survey (ACS) datasets enables dynamic re-aggregation of demand by census tract, which is particularly valuable when aligning infrastructure upgrades with environmental justice goals.

Table B.3: Population projections and water demand scenarios for Champaign County.

Scenario	2020	2030	2040	2050	Growth Rate	Demand Change
<i>Population (thousands)</i>						
Low growth	210.0	215.2	218.5	220.1	0.16%/yr	—
Baseline	210.0	224.7	238.5	251.2	0.60%/yr	—
High growth	210.0	235.8	265.4	298.2	1.17%/yr	—
<i>Municipal demand (MLD) assuming current per-capita use</i>						
Low growth	99.8	102.3	103.8	104.6	—	+5%
Baseline	99.8	106.8	113.4	119.4	—	+20%
High growth	99.8	112.1	126.1	141.7	—	+42%
<i>Municipal demand (MLD) with 15% conservation</i>						
Baseline + conservation	99.8	90.8	96.4	101.5	—	+2%

Implications for planning.

Under the baseline scenario, municipal demand could increase 20% by 2050, requiring either new well capacity or demand management. The conservation scenario demonstrates that aggressive efficiency measures could nearly offset population growth, keeping demand stable through mid-century.



C. Regulatory and Management Framework

This appendix documents the statutory authorities, regulatory programs, and institutional responsibilities governing groundwater management in Illinois and specifically within the Mahomet Aquifer region. It provides the legal and policy context within which *Talking Aquifer* operates, enabling the platform to incorporate compliance requirements, permit conditions, and regulatory thresholds into decision-support analytics.

Content organization:

- **Section C.1:** State and federal statutes (IGPA, SDWA, Clean Water Act)
- **Section C.2:** Institutional roles (IEPA, ISWS, ISGS, IDPH, county authorities)
- **Section C.3:** Permit programs (high-capacity wells, water-supply construction, discharge)
- **Section C.4:** Groundwater quality standards and monitoring requirements
- **Section C.5:** Wellhead protection, sole-source aquifer designation, regulated recharge areas
- **Section C.6:** Public data access, FOIA, and confidentiality provisions

How this informs Talking Aquifer:

The platform incorporates regulatory layers—wellhead setback zones, regulated recharge areas, permit locations, compliance monitoring points—as spatial overlays in the semantic graph. Users can query "*Which wells are within the Urbana wellhead protection area?*" or "*Show me sites with nitrate exceedances in the past 5 years*" and receive geographically and legally contextualized answers. Compliance thresholds (MCLs, pumping limits) trigger automated alerts when forecasts or observations approach regulatory boundaries.

C.1 Statutory Authorities

Illinois groundwater governance draws upon intertwined state and federal statutes. The Illinois Groundwater Protection Act (IGPA) establishes the framework for classified aquifers, regulated recharge areas, and cooperative planning among state agencies and local governments. Federal Safe Drinking Water Act provisions—particularly wellhead protection and underground injection control—overlay additional requirements that the Illinois Environmental Protection Agency (IEPA)

administers. The Clean Water Act and corresponding National Pollutant Discharge Elimination System (NPDES) permits regulate withdrawals that discharge to surface waters. Talking Aquifer encodes these statutes as constraints and metadata so that scenario planning can automatically flag conflicts with legal mandates.

C.2 Institutional Roles

Several agencies share responsibility for enforcing the statutory framework. The IEPA oversees water-quality standards and source-water protection, while the Illinois State Water Survey (ISWS) collects hydrologic data and provides scientific assessments. The Illinois State Geological Survey (ISGS) manages geological mapping and aquifer characterisation. The Department of Natural Resources (IDNR) issues high-capacity well permits, and county health departments regulate private wells and onsite wastewater systems. Coordination occurs through interagency agreements and regional water-supply planning groups; these institutional relationships are captured within the semantic graph alongside contacts and reporting cadences.

Table C.1: Key agencies and their groundwater management responsibilities.

Agency	Primary Responsibilities	Key Programs
Illinois EPA (IEPA)	Water quality standards, source-water protection, discharge permits	NPDES, UIC, Source Water Protection
IL State Water Survey (ISWS)	Hydrologic monitoring, scientific assessment, technical assistance	Observation well network, WARM, drought monitoring
IL State Geological Survey (ISGS)	Geologic mapping, aquifer characterization, well logs	ILWATER database, 3D geologic models
IL Dept. Natural Resources (IDNR)	High-capacity well permits, water allocation	Lake Michigan Allocation, IWIP
County Health Departments	Private well construction, septic systems, local ordinances	Well permits, abandonment, water testing
USEPA Region 5	Federal oversight, sole-source aquifer designation	SDWA enforcement, Superfund

C.3 Permit Programs

High-capacity withdrawal permits and water-supply construction permits form the backbone of regulatory oversight. Applicants must submit pump-test data, well construction specifications, and impact analyses for neighbouring wells. Talking Aquifer links permit records to time-series withdrawals so that model forecasts remain consistent with legal pumping limits. The platform also tracks permit renewal timelines, enabling alerts before authorisations lapse.

C.4 Standards and Monitoring Requirements

Groundwater quality standards derive from 35 Ill. Adm. Code 620 and include numeric limits for contaminants such as nitrate, arsenic, and volatile organic compounds. Compliance monitoring typically involves quarterly sampling at public-water-supply entry points, annual sampling for designated contaminants, and targeted monitoring for regulated recharge areas. These frequency requirements feed directly into the platform's data-quality rules, ensuring that missing samples are flagged and interpolated with appropriate caution.

Monitoring frequencies.

- **Quarterly:** Entry-point sampling for microbiological indicators (total coliform, E. coli)

- **Annual:** Inorganic contaminants (nitrate, arsenic, fluoride, metals)
- **Every 3 years:** Synthetic organic compounds, volatile organic compounds
- **Every 6 years:** Radionuclides (radium, uranium, gross alpha)
- **As triggered:** Regulated recharge area monitoring, spill response sampling

Table C.2: Key groundwater quality standards applicable to the Mahomet Aquifer (35 Ill. Adm. Code 620).

Parameter	MCL/Standard	Units	Health/Regulatory Concern
Nitrate (as N)	10	mg/L	Methemoglobinemia (blue baby syndrome)
Arsenic	10	µg/L	Carcinogen; naturally elevated in some aquifers
Fluoride	4.0	mg/L	Skeletal fluorosis at high concentrations
Lead	15	µg/L	Neurotoxin; action level triggers corrosion control
Total coliform	0	presence/100 mL	Indicator of microbial contamination
Radium-226+228	5	pCi/L	Radioactivity; bone cancer risk
Atrazine	3	µg/L	Herbicide; endocrine disruption concerns
Chloride	200	mg/L	Secondary standard; taste, corrosivity
Total dissolved solids	500	mg/L	Secondary standard; aesthetic, scaling

C.5 Protection Programs

Wellhead protection areas, sole-source aquifer designations, and regulated recharge areas provide spatial defensiveness against contamination and unsustainable withdrawal. Counties may adopt overlay zoning to limit industrial uses in sensitive zones. Talking Aquifer ingests these polygons and exposes them as filters in the user interface, allowing practitioners to map proposed developments against protected areas and to evaluate mitigation options.

C.6 Data Access and Confidentiality

State Freedom of Information Act (FOIA) provisions govern access to public records, but certain utility and industrial data sets include confidentiality protections. The platform tracks provenance for each dataset—whether publicly accessible, FOIA-restricted, or shared through memoranda of understanding. When users query data subject to confidentiality agreements, Talking Aquifer either redacts the response or prompts the user to authenticate via the relevant data-sharing arrangement, preserving compliance with disclosure rules.



This appendix synthesizes decades of groundwater research, monitoring programs, and data infrastructure that underpin our understanding of the Mahomet Aquifer system. It serves as a comprehensive literature review and data provenance record, documenting the observational foundation upon which *Talking Aquifer* builds.

Content organization:

- **Section D.1:** Historic well networks and data repositories
- **Section D.2:** Geophysical investigations and tracer studies
- **Section D.3:** Remote-sensing approaches (GRACE, InSAR, satellite products)
- **Section D.4:** Ongoing monitoring programs (state, federal, utility)
- **Section D.5:** Identified data gaps and proposed enhancements

Integration into Talking Aquifer:

All datasets referenced in this appendix are mapped into the platform’s versioned semantic layer with full provenance (steward, retrieval date, schema version, quality flags). Legacy data undergo automated harmonization (coordinate transforms, unit conversions, datum corrections) as described in Chapter 1 Section 1.5. This ensures that model calibration, data assimilation, and user queries can traverse decades of observations without manual reconciliation.

Key synthesis products referenced:

- Illinois State Water Survey (ISWS) Mahomet Aquifer bulletins and circulars
- USGS Cooperative Water Program reports for east-central Illinois
- ISGS geophysical survey campaigns (AEM, seismic, HTEM)
- University of Illinois research theses and technical reports
- Peer-reviewed literature on regional hydrogeology, recharge estimation, and contaminant transport

D.1 Historic Well Networks

The Illinois State Water Survey and Illinois State Geological Survey have archived observation-well records dating back to the 1890s. Early artesian wells near Champaign and Mahomet provided the first depth-to-water measurements; later, cooperative programs with the USGS expanded the network to include more than 150 continuous recorders. Digitisation efforts in the 2000s converted strip-chart data to electronic archives, enabling modern trend analysis. Talking Aquifer references these legacy wells when reconstructing pre-development potentiometric surfaces and calibrating long-term groundwater models.

D.2 Geophysical and Tracer Investigations

Airborne electromagnetics (AEM), seismic reflection, and groundwater-age tracers (tritium-helium, carbon-14) illuminate the geometry and dynamics of the Mahomet Aquifer. USGS campaigns in 2015–2018 mapped confining units to a depth of 150 m, revealing channelised sand bodies that align with present-day pumping centres. Tracer tests indicate residence times exceeding 10,000 years in deep glaciofluvial deposits, underscoring the limited natural recharge to confined zones. These datasets inform the conceptual models embedded in Chapter 2.

D.3 Remote Sensing Contributions

Satellite gravimetry (GRACE), Interferometric Synthetic Aperture Radar (InSAR), and Landsat evapotranspiration products supplement in situ observations. GRACE signals, although coarse, validate basin-scale storage changes. InSAR used over the metropolitan area constrains land-surface deformation caused by pumping, while Landsat-based evapotranspiration helps partition surface-water budgets. Talking Aquifer ingests these remote-sensing products via the data pipelines outlined in Chapter 3, providing spatial context and early warning for anomalies.

D.4 Ongoing Monitoring Programs

Multiple agencies maintain monitoring programs: ISWS operates climate and groundwater stations; IEPA manages ambient groundwater-quality networks; and the Mahomet Aquifer Protection Alliance coordinates citizen science wells. In addition, municipal utilities deploy SCADA systems that log pumpage and water levels at hourly intervals. Talking Aquifer integrates these disparate feeds through automated ETL jobs, maintaining metadata on instrument calibration, maintenance events, and downtime.

D.5 Identified Data Gaps

Despite extensive monitoring, important gaps remain. Shallow unconfined deposits lack high-frequency water-quality data, limiting the ability to detect nutrient transport from agriculture. Transient pumping tests are underreported, creating uncertainty in hydraulic-parameter estimation. Recharge measurements rely on sparse lysimeter networks, making it difficult to validate modelled recharge variations under extreme weather. The platform tracks these gaps so that scenario analysis can incorporate uncertainty bounds and prioritise future instrumentation investments.



This appendix documents significant changes between report versions.

Version 1.0.0 (December 2024)

Initial release of the Talking Aquifer Technical Report.

- Complete documentation of 93 demonstration questions (Q001–Q093)
- Ten LLM pipeline implementations (P1–P10) with evaluation metrics
- System architecture covering domain, data, tools, agent, and UI layers
- Integration with ISWS observation wells, WARM weather stations, and USGS streamflow data
- Comprehensive appendices: technical reference, water use analysis, regulatory framework, and literature review

Version 1.0.1 (December 9, 2025)

Audited maintenance release aligned with the OECD CRP fellowship reporting timeline and focused on clarity, consistency, and build robustness across the rendered report and regenerated artefacts.

- Updated title/credits page to reflect the December 9, 2025 version date while retaining the 2025 fellowship visit window.
- Clarified front-matter and chapter naming in the table of contents (e.g., distinguishing the front-matter Executive Summary from the Background Executive Summary; labeling front-matter sections as “(Front Matter)”).
- Harmonised question-category labels and ranges (Q001–Q093) between the main Questions chapter, extended catalog, and coverage index; removed a duplicate “Summary: All 93 Questions” section from the extended appendix.
- Added minimal LaTeX fallbacks (e.g., `\degreeCelsius` macro) and fixed a handful of macro/escaping issues to ensure successful compilation on standard TeX Live installations.
- Rebuilt several figures, tables, and page layouts to remove unreadable labels, placeholder text, broken table wraps, stray blank pages, and inconsistent chapter/appendix transitions.
- Aligned question narratives, captions, summary/index entries, and concluding synthesis text with curated SQLite data and regenerated artefacts, correcting counts, dates, timestamps, map distances, trend descriptions, and figure interpretations across the question catalog while updating end-of-chapter summaries to reflect the audited outputs.



- [1] William L. Cunningham and Charles W. Schalk. *Groundwater Technical Procedures of the U.S. Geological Survey*. Techniques and Methods 1–A1. U.S. Geological Survey, 2011 (cited on page 32).
- [2] James S. Famiglietti. “The global groundwater crisis”. In: *Nature Climate Change* 4 (2014), pages 945–948 (cited on page 14).
- [3] Richard W. Healy and Peter G. Cook. “Using groundwater levels to estimate recharge”. In: *Hydrogeology Journal* 10 (2002), pages 91–109 (cited on page 13).
- [4] Dennis R. Helsel and Robert M. Hirsch. *Statistical Methods in Water Resources*. Techniques of Water-Resources Investigations, Book 4, Chapter A3. U.S. Geological Survey, 2002 (cited on page 58).
- [5] Illinois State Water Survey. *Groundwater Monitoring Network Assessment: Mahomet Aquifer, Champaign County, Illinois*. Contract Report 2020-01. Available from Illinois State Water Survey. Champaign, Illinois: Prairie Research Institute, University of Illinois at Urbana-Champaign, 2020 (cited on page 60).
- [6] Leonard F. Konikow. “Contribution of global groundwater depletion since 1900 to sea-level rise”. In: *Geophysical Research Letters* 38.17 (2011), page L17401 (cited on page 14).
- [7] Bridget R. Scanlon et al. “Groundwater depletion and sustainability of irrigation in the U.S. High Plains and Central Valley”. In: *Proceedings of the National Academy of Sciences* 109.24 (2012), pages 9320–9325 (cited on page 14).
- [8] Richard G. Taylor et al. “Ground water and climate change”. In: *Nature Climate Change* 3 (2013), pages 322–329 (cited on page 14).
- [9] D. Phil Turnipseed and Vernon B. Sauer. *Discharge Measurements at Gaging Stations*. Techniques and Methods 3–A8. U.S. Geological Survey, 2010 (cited on page 32).
- [10] Thomas C. Winter et al. *Ground Water and Surface Water: A Single Resource*. USGS Circular 1139. U.S. Geological Survey, 1998 (cited on page 11).