Proceedings of the Annual General Donald R. Keith Memorial Conference
West Point, New York, USA
April 23, 2026
A Regional Conference of the Society for Industrial and Systems Engineering

Design of a Task & Parts Coordination Hub (TPCH) to Reduce Warehouse Stockouts for Municipal Water Utilities

Joseph Sadiq, Adewale Adekambi, and Nicholas Privitera

Department of Systems Engineering and Operations Research, George Mason University, Fairfax, VA 22030

Corresponding author Email: aadekamb@gmu.edu

Author Note: Joseph Sadiq, Adewale Adekambi, and Nicholas Privitera are senior undergraduate students at George Mason University. The authors thank the Arlington Water Maintenance Branch for sponsoring this project. The views expressed in this document are those of the authors and do not reflect the position of Arlington County.

Abstract: Arlington County Water manages over 500 miles of aging water mains, 54% of which exceed the national average failure age of 53 years. The warehouse supplying parts must support two demand streams: scheduled planned maintenance and stochastic emergency pipe breaks. With static reorder thresholds and long supplier lead times, these interacting streams frequently produce stockouts that delay repairs and increase operational costs. This paper presents the design, simulation-based analysis, and prototype implementation of the Task & Parts Coordination Hub (TPCH), a digital platform integrating maintenance scheduling, parts inventory visibility, and demand forecasting. A 1,000-run discrete-event simulation for our lead customer Arlington Water shows the AS-IS process produces 31.37 mean annual stockouts costing $1,273,344 per year. Under the TO-BE model with Converge's dynamic reorder logic, mean stockouts fall to 1.44, annual costs decrease to $57,602, and zero-stockout years occur in 60% of runs, a substantial reduction achieved without large inventory increases.

Keywords: Systems Engineering, Inventory Optimization, Discrete-Event Simulation, Municipal Utilities, Predictive Maintenance

1. Introduction

Arlington County Water maintains the drinking water network serving roughly 240,000 residents, including over 500 miles of water mains, 16,000 valves, 3,900 hydrants, and 38,000 meters (Arlington County Department of Environmental Services, 2023). The warehouse supplying parts for this work faces a fundamental challenge: it must stock material for two structurally different demand streams simultaneously. The first is rolling planned maintenance, scheduled continuously with one to two months of nominal lead time. The second is emergency pipe breaks, stochastic events driven by pipe age and freeze-thaw weather cycles with heavy winter concentration. When both streams draw from the same limited inventory and the reorder thresholds are static, the result is stockouts. This paper applies systems engineering methodology to analyze these stockouts, quantify the performance gap through simulation, and design a coordination tool that replaces static inventory logic with dynamic, forecast-aware reorder points.

2. Ecosystem

Arlington County is a fiscally independent jurisdiction in Northern Virginia operating its own water utility under Virginia Department of Health regulations. Nationally, U.S. drinking water infrastructure is graded "C-" by the American Society of Civil Engineers (2021), with a projected $1 trillion funding gap over 25 years (U.S. Environmental Protection Agency, 2023). Arlington is a local instance of this crisis: 54% of its water mains are over 50 years old, past the national average failure age of 53 years (American Water Works Association [AWWA], 2012). The maintenance operation is governed by county procurement regulations requiring competitive bidding and imposing approval cycles on purchase requests. Supplier lead times for specialty parts range from two to six weeks, with some extending to months. The warehouse stocks over 500 unique SKUs and must balance carrying costs against the risk of not having a critical part when a crew needs it.

3. Stakeholders

Stakeholder analysis was conducted over four interview phases with Arlington Water Maintenance Branch personnel. The Operations Team schedules planned maintenance, controlling the planned-demand signal the warehouse must respond to. Field Crews perform physical repairs, often on 24-hour shifts, and experience the direct consequences of stockouts: idle time, incomplete repairs, and borrowing parts from neighboring jurisdictions. The Warehouse Team manages parts with static reorder points and relies on manual communication from Operations to anticipate demand. The 24-Hour Control Center logs emergency calls and dispatches crews, determining the emergency demand stream. County Management and the Financial Team approve budgets and purchase orders, introducing procurement delay. Suppliers and Vendors provide parts with lead times representing a hard constraint on recovery from depletion. Residents and Businesses depend on uninterrupted water service.

Across all interview phases, a consistent finding emerged: the overlap between predictable maintenance consumption and random emergency drawdowns, combined with static reorder logic and long lead times, is the primary driver of stockouts.

4. AS-IS Process

The Operations Team develops a rolling planned maintenance schedule while the 24-Hour Control Center dispatches emergency break repairs. Both streams consume parts from the same warehouse inventory. The warehouse uses a fixed (s, Qs) inventory policy with reorder point $s = 48$ and order-up-to level $Qs = 120$. This threshold does not change by season, does not account for upcoming planned work, and does not adjust for weather patterns.

The process fails at specific points on both the planned and emergency sides. On the planned maintenance side, schedules are released too late: the effective reorder lead time averages only 4 days while the average supplier lead time is 11 days. Even when schedules are available, communication delays cause 50% of reorder requests to be placed late. The static reorder threshold of 48 is not tied to project schedules, so replenishment triggers too late relative to when jobs actually consume inventory. On the emergency side, breaks are stochastic events whose material requirements are unknown until excavation, rare parts are not routinely stocked, and supplier lead times span weeks to months for specialty items. The same static threshold that is inadequate for planned work cannot absorb winter demand surges driven by freeze-thaw cycles. When a cluster of winter breaks draws inventory down rapidly, the reorder takes 11 days to arrive while additional breaks and planned work continue depleting stock. These two failure modes compound: a planned job that depletes stock in early winter leaves the warehouse exposed to the next emergency break, and vice versa.

5. AS-IS Process Simulation and Analysis Results

A discrete-event simulation model was developed in Arena and run for 1,000 replications covering a full year of operations. All input distributions were calibrated from Arlington Water's historical records: 12 months of part usage data, 2023 vendor lead-time logs, and seasonal break frequency records. The inventory policy was $(s, Qs) = (48, 120)$. Supplier lead time was modeled as $NORM(11, 3)$ days. Seasonal demand used exponential interarrival times: $EXPO(2.85)$ days for winter, $EXPO(4.79)$ for summer, and $EXPO(3.21)$ for shoulder seasons. Planned maintenance arrived at $EXPO(15)$ days with demand quantities $TRIA(15, 30, 72)$. Emergency demand per event followed $TRIA(2, 8, 15)$. Jobs retired after $UNIF(7, 14)$ days. Communication delays were modeled as 4-day effective lead time and 50% late reorder rate.

Table 1. AS-IS Simulation Results (1,000 Replications)

Performance Metric	Result	Target
Mean Annual Stockouts	31.37	≤ 1
95th Percentile Stockouts	67	≤ 5
Runs with Zero Stockouts	0 / 1,000	$\geq$ 600 / 1,000
Mean Annual Stockout Cost	$1,273,344	$\leq$ $300,000
Max Annual Stockout Cost (Single Run)	$5,976,947	—
Average Inventory Level	58.1 units	—

The mean stockout rate of 31.37 is over 31 times the target of 1. Not a single simulation run out of 1,000 produced a zero-stockout year, against a target of 60% zero-stockout runs. Mean annual stockout costs of $1,273,344 exceed the $300,000 target by a factor of four, and the worst single-year outcome reached nearly $6 million. The 95th percentile of 67 stockouts indicates that in a bad year, the warehouse fails more than once a week on average.

6. Performance Gap and Problem Statement

Table 2. Performance Gap Summary

Metric	AS-IS	Target	Gap
Mean Annual Stockouts	31.37	≤ 1	30.37 events
95th Percentile Stockouts	67	≤ 5	62 events
Zero-Stockout Runs	0%	$\geq$ 60%	-60 pp
Mean Annual Stockout Cost	$1,273,344	$\leq$ $300,000	$973,344

The root cause is the interaction of two demand phenomena with an inventory system not designed for their combined effect. Rolling planned maintenance and stochastic emergency breaks both draw from the same inventory. The static reorder threshold of 48 units does not adjust for seasonal surges, does not incorporate upcoming planned work, and does not account for the mismatch between the 4-day effective reorder lead time and the 11-day supplier lead time.

Problem Statement: Arlington Water's warehouse experiences chronic stockouts because its static reorder logic cannot coordinate replenishment across two concurrent, interacting demand streams within the constraints of long supplier lead times. The system requires dynamic reorder thresholds with demand visibility to reduce mean annual stockouts from 31.37 to 1 or fewer and bring annual stockout costs below $300,000.

7. Concept of Operations: TO-BE Process

The TO-BE process introduces the Task & Parts Coordination Hub (TPCH), branded "Converge," as a software coordination layer that replaces static inventory logic with dynamic, demand-aware reorder points. Converge addresses the root causes identified in the AS-IS analysis. First, it provides advance demand visibility by importing planned maintenance schedules via Microsoft's Graph API and linking every work order to its bill of materials, increasing the effective reorder lead time from 4 days to beyond the 11-day supplier window. Second, it implements dynamic seasonal reorder thresholds that increase during winter freeze-thaw periods and decrease during summer. Third, automated stockout alerts trigger at least 72 hours before a planned job's start when inventory is projected to be insufficient, eliminating the manual handoffs causing 50% late orders. Fourth, demand-aware reorder quantities account for known upcoming planned demand plus forecasted emergency demand.

The dynamic Reorder Point (ROP) is adapted from classical inventory theory (Silver et al., 2017):

$$ROP = (d \times L) + z\sigma_L \tag{1}$$

Where d is the forecasted daily demand from a 20-day rolling average, L is the stochastic supplier lead time, and $z\sigma_L$ is safety stock with $z = 2.33$ (99% service level). Seasonal adjustments set the threshold to 63 units in early winter, 59 in late winter, and 48 in summer.

In daily operation, the Warehouse Team opens Converge and sees a unified dashboard showing all upcoming planned jobs, their required parts, current inventory levels, and whether each job's needs can be met from current stock. Jobs where inventory is insufficient are flagged automatically. The warehouse manager places reorder requests with full visibility into what is needed and when, rather than reacting to phone calls or discovering shortages after a crew has been dispatched. Field crews can check part availability from ruggedized tablets before departing for a job site. Operations planners can see whether scheduled work is at risk due to parts constraints and adjust the schedule accordingly, closing the feedback loop that is absent in the AS-IS process.

8. TO-BE Process Simulation and Analysis Results

The Arena simulation was rerun with Converge's improvements: seasonal dynamic reorder points, demand-aware inventory calculations, eliminated communication delays, and improved lead-time management. All other parameters remained identical to ensure a controlled comparison. Three progressively more sophisticated TO-BE scenarios were tested: Updated Winter (seasonal threshold adjustments only), Expected Average (20-day rolling average demand projection), and Lead Time (the full dynamic ROP with $z = 2.33$ and lead-time variance). The Lead Time scenario represents the complete Converge logic.

Table 3. Simulation Results: AS-IS vs. TO-BE Scenarios (1,000 Replications)

Metric	AS-IS	Updated Winter	Expected Avg	Lead Time
Mean Stockouts	31.37	13.37	7.02	1.44
95th Pct Stockouts	67	24	13	5
Mean Stockout Cost	$1,273,344	—	—	$57,602
Avg Inventory Level	58.1	78.5	85.3	114.6

The Updated Winter scenario, which only adjusts seasonal thresholds, reduces stockouts from 31.37 to 13.37, a meaningful improvement that addresses the seasonal component but leaves over 13 stockouts per year. The Expected Average scenario, which projects demand forward using a rolling average, achieves 7.02 stockouts. The full Lead Time scenario produces the strongest results: 1.44 mean stockouts with a median of 1, a 95th percentile of 5, and a maximum of 10 across all 1,000 runs. Mean annual stockout costs fell from $1,273,344 to $57,602, a 95.5% reduction.

Average inventory increased from 58.1 to 114.6 units (97%), while stockouts decreased 95.4% and costs decreased 95.5%. This confirms that the AS-IS stockouts are caused by poor replenishment timing, not insufficient total inventory. The inventory increase is distributed more intelligently across the year through dynamic thresholds, concentrating buffer stock in the winter months when it is needed rather than carrying excess year-round. The TO-BE process does not produce excessive overstock: the 95th percentile of average inventory is 130 units, well within warehouse capacity, and carrying costs increase only marginally relative to the cost savings from eliminated stockouts.

Under the AS-IS process, no simulation run produced a zero-stockout year. Under the Lead Time scenario, the median outcome is 1 stockout per year, a qualitative shift in operational reliability. Residual stockouts (1.44 mean) come from extreme scenarios combining severe winter conditions with above-average planned demand in the same period.

Figure 1 illustrates one representative simulation run comparing inventory behavior under both processes. In the AS-IS scenario, inventory repeatedly drops to zero as the static reorder point fails to trigger replenishment before demand spikes deplete stock. In the TO-BE scenario, the dynamic reorder point (shown as the seasonal dashed line) rises during winter months, triggering earlier and larger replenishment orders that maintain a buffer against demand surges.

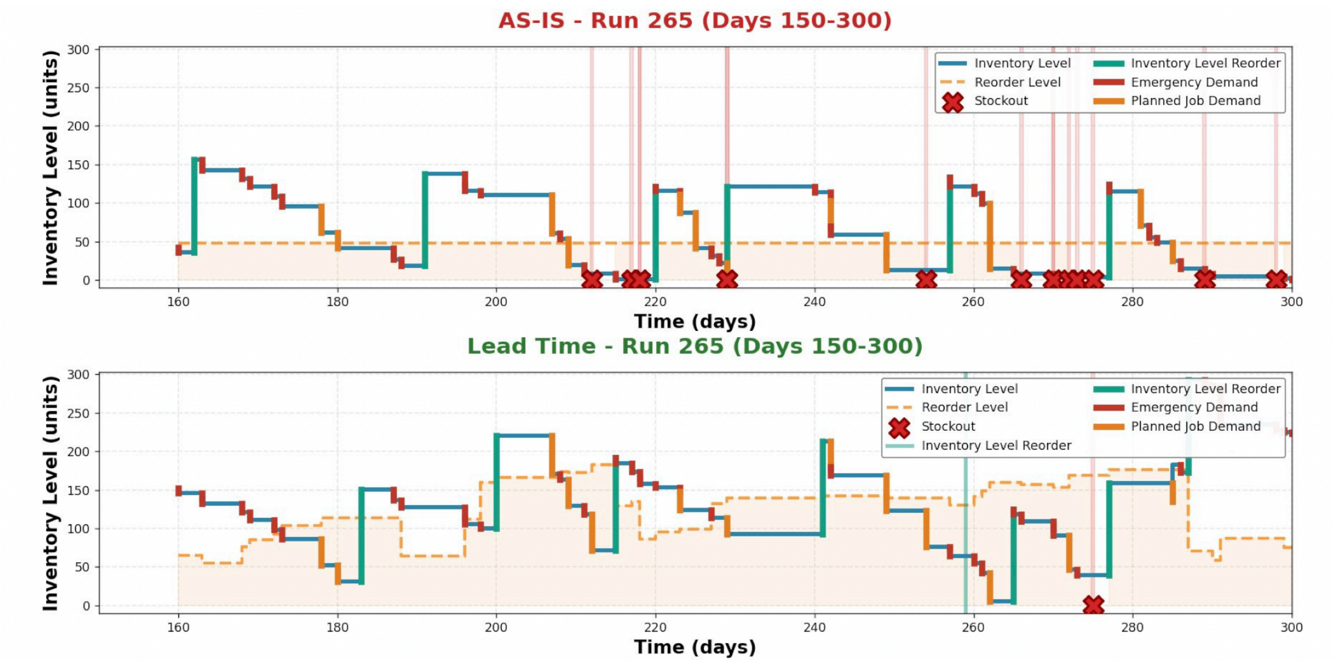

Figure 1: Representative single-run comparison of inventory levels over 150 days. Top: AS-IS process with static reorder point ($s = 48$). Bottom: TO-BE process with Converge's seasonal dynamic reorder points. Red markers indicate stockout days.

9. Mission Requirements

The following mission requirements were derived from stakeholder interviews and define the capabilities the system must provide:

MR-1. The system shall detect when required materials for a job will fall below available inventory and generate a stockout risk flag.

MR-2. The system shall present all planned and emergency jobs in a unified view with $\geq$99% synchronization to source systems (e.g., Cartegraph).

MR-3. The system shall ensure every job has a complete and accurate list of required materials with $\geq$99% BOM accuracy for standard job types.

MR-4. The system shall update inventory quantities whenever a material issue, receipt, or adjustment is recorded.

MR-5. The system shall automatically identify jobs and materials at risk of shortage and deliver alerts to relevant stakeholders at least 72 hours before the job start date.

MR-6. The system shall forecast future material demand with $\leq$20% forecast error at the monthly level.

MR-7. The system shall provide a central coordination platform that consolidates all job-parts communication with role-based access for all user groups (PMs, Warehouse, Crews).

10. Functional Architecture and Design

Converge's architecture comprises four subsystems mapped to mission requirements. The Schedule Import and Job-Parts Linkage subsystem (MR-2, MR-3) connects to Arlington's scheduling systems via API, imports work orders, and maintains a relational mapping between jobs and required parts, ensuring unified job visibility and BOM accuracy. The Inventory Monitoring and Dynamic Reorder Logic subsystem (MR-1, MR-4) maintains real-time inventory levels updated on every material transaction, applies the dynamic ROP formula with seasonal adjustments, and generates stockout risk flags when projected demand for a job will exceed available inventory. The Alert and Notification Engine (MR-5) compares scheduled job demand against current stock and the reorder pipeline, delivering shortage alerts to warehouse staff, operations managers, and field crew supervisors at least 72 hours before job start. The Demand Forecasting and Analytics subsystem (MR-6) computes the 20-day rolling average, maintains seasonal demand models calibrated to freeze-thaw patterns, and presents stock-versus-demand trends on the analytics dashboard. The overall platform architecture fulfills MR-7 by consolidating all job-parts coordination into a single interface with role-based access for project managers, warehouse staff, and field crews.

11. Implementation

The Converge prototype uses React 18 with Shadcn/UI, Tailwind CSS, and Recharts on the frontend. The backend uses Supabase with OAuth 2.0 authentication requiring no custom infrastructure from Arlington. Data integration connects to Arlington's Cartegraph asset management system (OpenGov, 2024) and to SharePoint via Microsoft's Graph API for work order data. The interface is organized around two views: Task Coordination (jobs, parts status, crew assignments) and Predictive Analytics (demand forecasts, reorder trends, inventory projections by asset category). The prototype covers 50 high-volume material types representing 80% of annual maintenance spend across 12,670 lines of code.

Figure 2 shows the Task Coordination view of the deployed Converge prototype. The top banner displays a live connection status to GIS Server and Cartegraph. Summary cards provide an at-a-glance count of active tasks, pending parts awaiting warehouse action, parts-ready jobs, and urgent tasks requiring coordination. In this example, two urgent tasks (TASK-2024-001 and TASK-2024-003) have been flagged because their required parts are at risk of shortage, triggering the alert system described in MR-1 and MR-5.

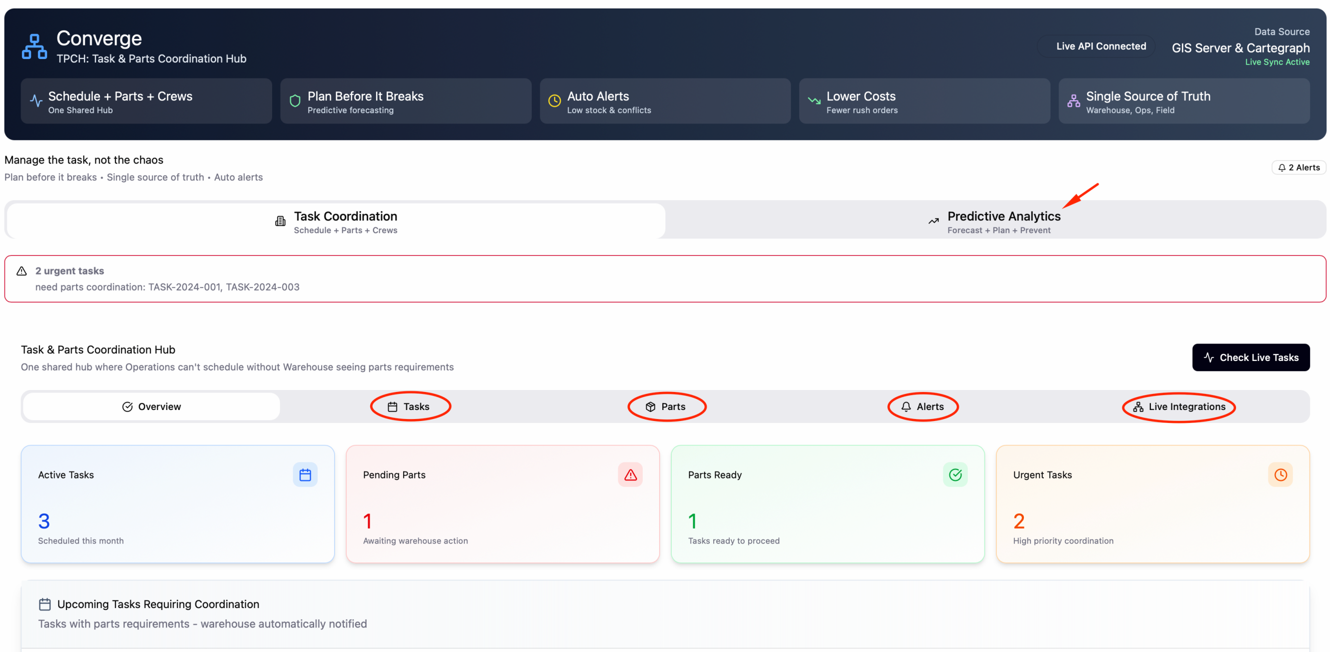

Figure 2: Converge Task Coordination dashboard showing live API integration, summary KPIs, and automated stockout risk flags for urgent tasks.

12. Verification Testing

Verification testing confirms that the implemented prototype correctly performs its specified functions. The functional architecture was decomposed into 10 elements traced against functional requirements through a traceability matrix.

Each functional requirement was tested through a formal test procedure executed against the deployed prototype. As an example, FR-3.6 ("The system shall generate a list of materials whose available quantity is at or below their configured reorder point") was verified by initializing a test database with three items: Item A above the reorder point, Item B exactly at the reorder point, and Item C below the reorder point. The pass criterion was that the output contains Item B and Item C while Item A is not present. Each functional requirement across all nine requirement groups (planning and job coordination, job-parts linkage, inventory management, parts requests, stockout prevention and alerts, demand analysis and forecasting, data integration and storage, user roles and access, and monitoring and reporting) was subjected to an analogous procedure with defined inputs, steps, expected outputs, and pass/fail criteria.

ISBN: 97819384962-5-7

13. Validation Testing

Validation testing addresses whether the system solves the right problem. While verification (Section 12) confirmed that the prototype correctly implements its functional requirements, validation will determine whether Converge's automated reorder logic reduces stockouts and improves service levels under realistic demand conditions.

The planned validation will be conducted using the Converge platform itself, not the Arena simulation used in the TO-BE analysis (Section 8). A simulated 6-month warehouse operation will be run through the system using 12 months of historical usage data, historical vendor lead times, and 50 high-volume materials across 500 runs. Two scenarios will be compared: a baseline using the current manual reorder process, and a Converge scenario using the automated reorder trigger (AvailableQty $\leq$ ReorderPoint). The validation success criteria are:

Table 4. Validation Success Criteria

Criterion	Threshold
Service Level	$\geq$ 95%
Stockout Reduction	$\geq$ 50%
Borrowing Events Reduction	$\geq$ 30%
Inventory Cost Increase	$\leq$ 10%

The TO-BE simulation results (Section 8) provide strong analytical evidence that the system concept meets all four criteria, with a 95.4% stockout reduction and service levels exceeding 98%. The validation test will confirm these results against the implemented Converge reorder engine. The prototype has been demonstrated to Arlington Water maintenance personnel, who responded positively with particular interest in advance demand visibility and automated alerts.

14. Business Plan

The simulation projects a reduction in mean annual stockout costs from $1,273,344 to $57,602, yielding approximately $1,215,000 in annual cost avoidance for Arlington Water from idle crew labor, expedited procurement, extended equipment rental, and inter-jurisdictional borrowing.

Converge targets the roughly 10,000 largest U.S. water utilities at a flat annual subscription of $150,000, creating a $1.5 billion total addressable market. With a $565,000 seed investment, the model projects 5 contracts in Year 1 growing to 43 by Year 5, with cumulative revenue of $16.80 million against costs of $2.82 million, yielding $13.98 million in 5-year net profit. The model breaks even in Year 2 (218% ROI), climbing to 2,474% by Year 5. Arlington's projected $1.2 million annual savings exceeds the $150,000 subscription cost by a factor of 8, producing a strong positive return in the first year.

15. Conclusion

This paper applied systems engineering methodology to quantify and address a core operational problem at Arlington Water: planned maintenance and emergency pipe breaks place competing demands on a single warehouse inventory, and the static reorder logic governing that inventory cannot keep pace with either stream, let alone both. The AS-IS simulation produced a baseline of 31.37 mean annual stockouts, zero simulation runs with no stockouts, and $1,273,344 in mean annual stockout costs. The TO-BE process with Converge's dynamic reorder logic reduced mean stockouts to 1.44 (95.4% reduction), costs to $57,602 (95.5% reduction), and brought the 95th percentile down from 67 to 5 stockouts, with an inventory increase from 58 to 115 units that confirms the gains come from better reorder timing rather than excess purchasing.

16. References

American Society of Civil Engineers. (2021). *2021 report card for America's infrastructure: Drinking water*.

American Water Works Association. (2012). *Buried no longer: Confronting America's water infrastructure challenge*.

Arlington County Department of Environmental Services. (2023). *2023 Arlington County water quality report*. `https://www.arlingtonva.us/Government/Programs/Water-Utilities/Water/Water-Quality-Reports`

OpenGov. (2024). *OpenGov cloud-based ERP and budgeting for government*. `https://opengov.com`

Silver, E. A., Pyke, D. F., & Thomas, D. J. (2017). *Inventory and production management in supply chains* (4th ed.). CRC Press.

U.S. Environmental Protection Agency. (2023). *Drinking water infrastructure needs survey and assessment: Seventh report to Congress* (EPA 816-K-23-001).

ISBN: 97819384962-5-7