



INFORMS Journal on Applied Analytics

Publication details, including instructions for authors and subscription information:
<http://pubsonline.informs.org>

Dynamic Route Optimization and Automation of Industrial Routes at WM

Hemachandra Pillutla, Seongbae Kim, Chao Zhou, Heesu Hwang, Sergiy Savchenko,
Stuart Greene, Marcel Dalby

To cite this article:

Hemachandra Pillutla, Seongbae Kim, Chao Zhou, Heesu Hwang, Sergiy Savchenko, Stuart Greene, Marcel Dalby (2026) Dynamic Route Optimization and Automation of Industrial Routes at WM. INFORMS Journal on Applied Analytics 56(1):95-112. <https://doi.org/10.1287/inte.2025.0279>

Full terms and conditions of use: <https://pubsonline.informs.org/Publications/Librarians-Portal/PubsOnLine-Terms-and-Conditions>

This article may be used only for the purposes of research, teaching, and/or private study. Commercial use or systematic downloading (by robots or other automatic processes) is prohibited without explicit Publisher approval, unless otherwise noted. For more information, contact permissions@informs.org.

The Publisher does not warrant or guarantee the article's accuracy, completeness, merchantability, fitness for a particular purpose, or non-infringement. Descriptions of, or references to, products or publications, or inclusion of an advertisement in this article, neither constitutes nor implies a guarantee, endorsement, or support of claims made of that product, publication, or service.

Copyright © 2026, INFORMS

Please scroll down for article—it is on subsequent pages



With 12,500 members from nearly 90 countries, INFORMS is the largest international association of operations research (O.R.) and analytics professionals and students. INFORMS provides unique networking and learning opportunities for individual professionals, and organizations of all types and sizes, to better understand and use O.R. and analytics tools and methods to transform strategic visions and achieve better outcomes.

For more information on INFORMS, its publications, membership, or meetings visit <http://www.informs.org>



THE FRANZ EDELMAN AWARD
Achievement in Operations Research

Dynamic Route Optimization and Automation of Industrial Routes at WM

Hemachandra Pillutla,^{a,*} Seongbae Kim,^a Chao Zhou,^a Heesu Hwang,^a Sergiy Savchenko,^a Stuart Greene,^a Marcel Dalby^a

^aCollection Operations & Business Optimization, WM, Houston, Texas 77002

*Corresponding author

Contact: hpillut2@wm.com,  <https://orcid.org/0009-0009-1808-8745> (HP); skim2@wm.com (SK); czhou@wm.com,

 <https://orcid.org/0009-0002-9138-470X> (CZ); hhwang@wm.com (HH); ssavche1@wm.com (SS); sgreene@wm.com (SG); mdalby@wm.com (MD)

Accepted: October 10, 2025

<https://doi.org/10.1287/inte.2025.0279>

Copyright: © 2026 INFORMS

Abstract. WM, the leading provider of environmental and sustainability solutions, faced significant challenges in optimizing its industrial waste collection routes following the coronavirus disease 2019 pandemic. Traditionally, these routes were planned manually, a process that was time consuming, labor intensive, and suboptimal with multiple factors that needed to be considered. WM embarked on a journey to develop and implement a dynamic route optimization program tailored to the unique requirements of its industrial waste collection operations. The industrial waste collection sector presents a complex routing problem because of the dynamic nature of customer service demand and several factors, such as customer-specific service requirements, different container types and sizes, and the disposal of waste materials. The dynamic route optimization system leverages data analytics, forecasting future service demand for effective capacity planning, and advanced algorithms using metaheuristics to automate the generation of routes that improve efficiency while ensuring safety and fulfilling customer service commitments. By effectively combining advanced analytics, optimization, and technology-led automation for the industrial line of business, WM realized best-ever efficiency gains and operating margins, and it set the foundations for driving increased value in the future.

Keywords: Franz Edelman Award • route optimization • metaheuristics • roll-off routing • industrial waste collection routing

Introduction

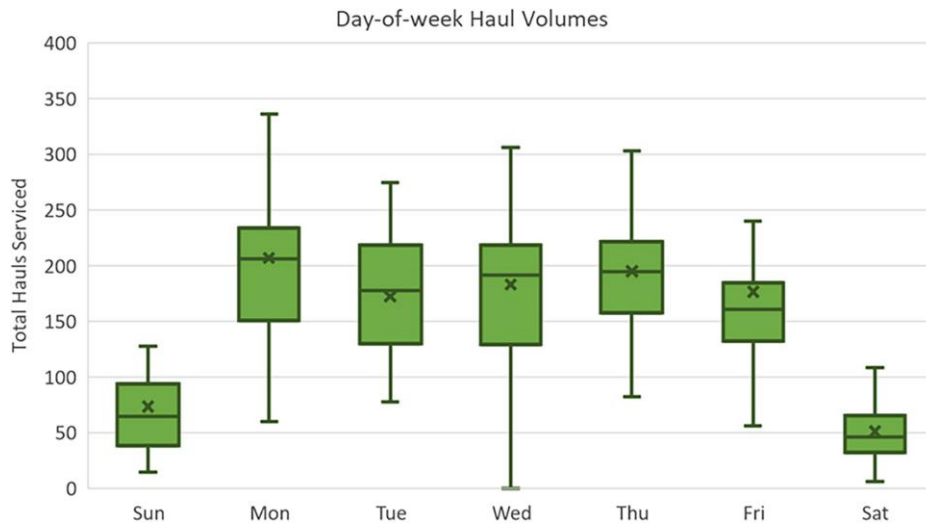
WM, previously known as Waste Management, is the leading provider of environmental solutions and the largest recycler in North America, with about 18,000 collection trucks in service. WM provides collection, recycling, and disposal services to millions of residential, commercial, industrial (also known as roll-off), and municipal customers throughout the United States and Canada. Residential collection addresses waste removal from private homes, with vehicles traversing neighborhoods collecting volumes of waste from multiple homes to gather and dispose/recycle the material collected. Commercial collection operations target businesses, such as restaurants and shopping centers, where specialized vehicles collect from designated containers. In the residential and commercial lines of business, customer services are recurring and typically scheduled on one or more designated service days. In contrast to residential and commercial

collections, industrial collection caters to waste producers, such as construction sites, and large commercial and mixed-use complexes. Day-to-day services for customers change, and the natures of customers and services are dynamic.

WM operations and routes are managed from 16 different market areas, encompassing more than 460 hauling sites (depots) across all lines of business, including over 3,800 industrial routes servicing more than 22,000 customer service tickets (also referred to as hauls) daily. The routes originate from the hauling site and can spread across 200 miles or more. The nature of the routes, including numbers of customers and types of service, varies substantially each day.

Haul volumes in the industrial line of business are influenced by the industrial economic environment. Day-of-the-week volume distributions differ greatly (Figure 1) both within and between sites. Volumes (Figure 2) may exhibit high volatility across the months

Figure 1. The Variability in Haul Volume by Day of the Week for a Site

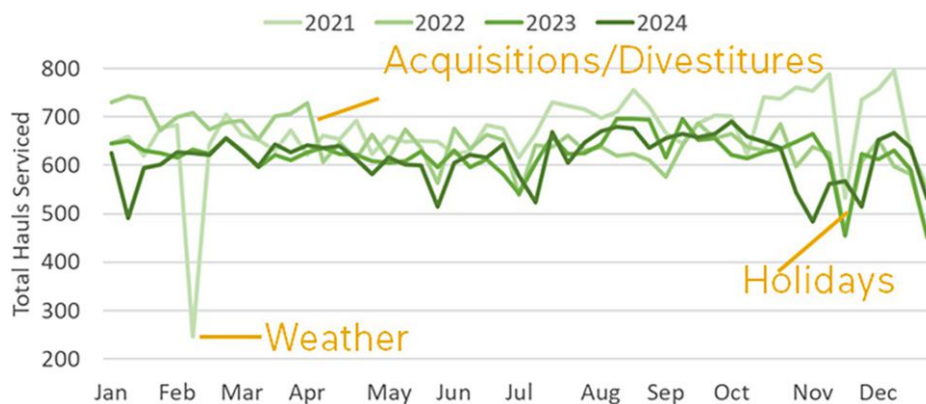


of the year because of external factors (e.g., weather, seasonal effects) and internal factors (e.g., staffing and vehicle constraints, major contract changes). Holiday schedules are not uniform across the enterprise, varying between the United States and Canada, and the days of the week on which several holidays, such as Independence Day and Christmas, are celebrated change each year. Sites take different approaches to holiday service depending on which day of the week the holiday falls. Some may schedule weekend work when they normally do not, and some may schedule extra routes before and/or after the holiday observance to cover for the lost time.

Hauling site operations range in size from a single route per day (i.e., fewer than 10 hauls) to over 50 routes per day (i.e., more than 500 hauls). Historically, industrial routes were planned and created manually

each day by over 150 route planners (i.e., routers) across our market areas for thousands of routes. The typical manual process would include estimation of daily haul volume demand expected and resources needed followed by the manual creation of routes. The anticipated volume of customer hauls was estimated based, at best, on moving averages. In order to create routes, the routers needed to consider many factors, including locations of customers, disposal facilities, over 20 potential service types per customer, more than 70 waste types, 100 or more container sizes and types, vehicle status and locations, driver- and vehicle-specific attributes, and customer service-level prioritization, while honoring several constraints, such as maximum route times, service time windows for customers, and facility operating hours. The process typically took four to five hours each day at each site

Figure 2. The Variability in Haul Volume by Month



Note. The graph depicts variations in haul volumes year over year across different months because of weather events, holidays falling on different days of week, and changes in volumes because of business acquisitions/divestitures.

per router. Human routers typically used visual tools, such as maps, to view the geographic locations of services and create routes. The quality of the routes produced by the routers was driven entirely by their experience and judgement as well as the quality of data available to them. All of these factors made it impossible for routers to optimally plan thousands of routes manually each day.

Given the dynamic nature of the daily volume of tickets in this line of business and same-day service requests, which much be estimated, proper forecasting and capacity planning are prerequisites for generating optimized routes. Improved demand forecasting, capacity planning, and technology-led route optimization and automation were needed to improve manual processes, build consistency in operations, better respond to changing conditions, and ensure that we meet our customer commitments.

Industrial Waste Route Optimization

Previous research (De Meulemeester et al. 1997, Bodin et al. 2000) has characterized the industrial waste collection challenge as the skip problem (i.e., a vehicle routing problem (VRP) where trucks deliver empty waste containers, collect full ones, and transport them to disposal facilities or depots), and the roll-on/roll-off collection and VRP are collectively referred to as the roll-on/roll-off vehicle routing problem (RRVRP). This paper similarly adopts the terms roll-off VRP, RRVRP, or roll-on/roll-off vehicle routing problem with time windows (RRVRPTW; incorporating time-window constraints).

In the roll-off problem, customers deposit waste or recyclable material into containers and request various services, such as collecting full containers, delivering empty containers, replacing full containers with empty ones, or changing container characteristics. Vehicles transport containers between customer locations and facility locations, including depots, landfills, and storage yards, to fulfill these demands.

The problem typically involves three types of facilities: hauling sites or depots, storage yards, and disposal facilities (i.e., landfills, transfer stations, or recycle facilities). Most vehicle routes originate and terminate at a

depot. Empty containers are stored at storage yards, whereas full containers are transported to disposal facilities.

Because of the considerable size and weight of the containers, most vehicles can carry only one container at a time except for an optional tandem-trailer configuration, which can handle two containers simultaneously. Tandem operations are feasible only for vehicles equipped to handle a trailer and when customer locations have sufficient space for two containers (see Figure 3). Tandem routes, although efficient, involve longer service times because of the need for the loading and unloading of two containers, and they require specially qualified drivers.

Storage yards maintain an inventory of empty containers of varying sizes and types, which are deployed based on customer requirements. The number of available containers for a given type/size at each yard may vary because there are multiple yards, and vehicles transport the containers between customer locations and facility locations. In addition, there are more than 70 types of waste, and the waste types must be matched to the landfills that are permitted to accept them. A customer's demand is assigned to a waste type; therefore, we must only use a permitted landfill for that customer's demand.

The logistics operations in the roll-off problem involve the coordination of vehicles to transport containers efficiently between customer locations, disposal locations, storage yards, and depots. These requirements include the delivery (DEL) of empty containers, exchange of full containers with new ones, transportation of full containers to disposal sites, and collection of empty containers from former customers among other tasks. Table 1 shows the major service types and provides a description that presents the movement of a vehicle for the specific service type. Other service types (total of 20) can be converted to one of the types listed in Table 1.

Service-type acronyms are listed in Table 1. Service types, such as empty and return (E/R) and switch out (S/O), comprise a significant portion (over 75%) of all services. For E/R services, a vehicle transports a full container at the customer location to a landfill that

Figure 3. (a) A Regular Vehicle for Industrial Waste Collection and (b) a Tandem Vehicle with Two Containers



Table 1. Examples of Service Types for Customer Demands

Type of service	Description
BTY (bring to yard)	Bring the container from a customer to the yard
BTS (bring to yard swap)	Swap the container at the customer location with a new one, and then, transport the existing container to the yard
E/R (empty and return)	Take a full container from the customer to a landfill, empty it, and return it
BTE (bring to yard empty)	Bring an empty container from the customer to the yard
FFN (full from yard no return)	Transport a full container from the yard to a landfill and empty it
FFY (full from yard)	FFN + deliver the empty container to the original customer
FDR (full from yard relocate)	FFY + bring an empty container back to a yard
DEL (delivery)	Deliver an empty container to the customer
DNR (do not return)	Transport a full container from the customer to a landfill and empty it
DNE (do not return empty)	Remove an empty container from the customer location
S/O (switch out)	Swap the full container at the customer location with a new one; then, transport the existing container to a landfill and empty it
REL (relocate)	Relocate the customer’s container to a different location within the same premises

accepts the same waste type, empties it, and returns the empty container to the customer. In contrast, the S/O services involve delivering a new container to the customer, leaving it in place, and when the customer has filled it, transporting the full container to a landfill for emptying. The empty container from the landfill can either be used for a different customer’s service or be returned to a container yard.

The operational requirements of each service-type necessitate variations in the vehicle’s status. For example, for S/O and DEL services, vehicles must arrive at the customer location with an empty container. In contrast, for do not return (DNR) and E/R services, vehicles are required to arrive at the customer location without a container.


Service types full from yard, full from yard no return, and full from yard relocate, which involve picking up and emptying a full container from the depot, are prioritized because they require a landfill visit that could not be completed the previous day, resulting in full containers being brought back to the depot.

The order in which customers are serviced significantly affects operational efficiency realized, where “efficiency” is WM’s primary productivity metric,

which is measured in the number of hauls completed per unit hour in the industrial line of business. Optimizing the sequence of customer services involves a process called “haul matching” at WM, through which empty containers are assigned to subsequent customers in a proper sequence based on the ending state of containers on a truck after servicing the current service and matching to the container requirement for the next customer service. The haul matching process helps reduce unnecessary trips to storage yards for containers while maximizing the use of limited container inventory. Figure 4 compares two routes serving the same three customers with DEL, E/R, and S/O services, thus demonstrating the efficiency gains of haul matching.

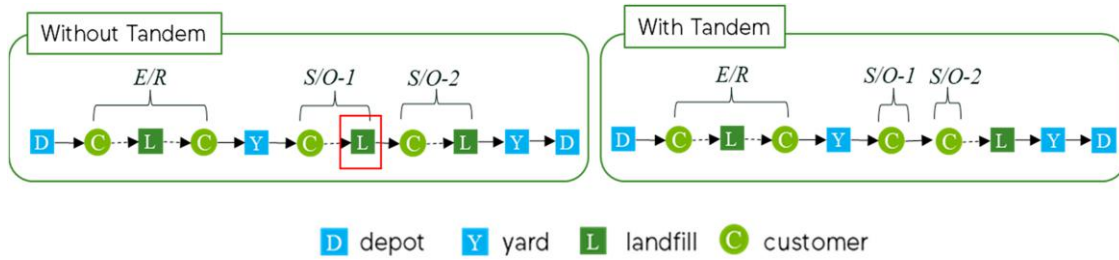
Each customer’s service requirements include designated waste material types; vehicles are routed to a landfill owned by either WM or a third party that accepts the same waste type to minimize the overall route time and improve efficiency. However, the model also maintains the flexibility to optimize based on total cost considerations in addition to time. In such cases, the choice of landfill is determined by cost factors, including disposal cost per ton and operating cost

Figure 4. Examples of Two Routes for Four Customers



Note. In this example, route 2, which incorporates haul matching between S/O and DEL services, eliminates unnecessary trips to the yard for a DEL customer, resulting in significantly improved efficiency compared with route 1.

Figure 5. Examples of Operations with Tandem Vehicles (Right Panel) and Without Tandem Vehicles (Left Panel) for Four Customers



Notes. In the tandem route, two landfill visits required for the two S/O services can be consolidated into a single visit. As a result, the tandem route is significantly more efficient compared with the route that does not use tandem vehicles.

of driver and truck per hour. Disposal costs may be considerably lower at WM-owned facilities compared with third-party facilities used, and overall costs may be lowered by maximizing disposal at our own facilities.

Additionally, the operations are subject to time-window constraints across vehicles, customers, landfills, yards, and depots. The working hours available for each vehicle and driver are variable, providing further volatility in scheduling and route planning.

Other Practical Considerations

In addition to the unique elements of the roll-off problem mentioned above, additional factors, including tandem routes that we describe in the Tandem Route section and carry-can operations that we describe in the Carry-can Operation section, must be considered in real-world operations.

Tandem Route

A portion of the vehicle fleet consists of tandem-capable vehicles, each equipped with a trailer for transporting an additional container. Tandem operations are subject to specific constraints, including the requirement for

qualified drivers trained to operate tandem trailers, permitted tonnage, and customer locations with sufficient space to accommodate a trailer with two containers. They also require extended loading and unloading service times compared with standard operations. Figure 5 illustrates nontandem and tandem routes for an E/R customer and two S/O customers.

Carry-Can Operation


Another practical use case is the carry-can operation. This involves temporarily parking an empty container at a designated customer location, thus enabling efficient use of containers in a remote area. As with tandem operations, not all customer locations are suitable for temporary container parking. Figure 6 illustrates the efficiency of carry-can operations in reducing unnecessary trips to storage yards.

Other Considerations

In addition to tandem and carry-can operational considerations, the following operational aspects are also incorporated in our model.

- *Service-type alteration.* E/R and S/O services may be interchanged when permitted by the customer and when doing so improves operational efficiency.

Figure 6. A Carry-Can Example with Three Customers



Notes. If three remote customer locations with different service types (e.g., E/R, S/O, and DEL) must be visited in that order, a vehicle with a carry-can significantly improves efficiency. Without the carry-can option, the vehicle needs to make a trip to the yard between E/R and S/O customers. With a carry-can operation and if an E/R customer is designated as a carry-can-available location, the vehicle departs with an empty container, can temporarily park the empty container at the E/R customer location, and can complete the E/R service. The vehicle then proceeds with the S/O and DEL services. This eliminates the unnecessary trip to the container yard, making the route more efficient.

- *Remote start/end location.* Vehicles may commence and/or conclude their operations at remote locations distinct from the depot locations. This flexibility allows drivers to start and end their shifts at these remote locations independent of their proximity to the depot, thereby enhancing operational feasibility.

- *Initial vehicle status consideration.* The model accounts for the initial status of each vehicle at the beginning of the day, including whether a container is attached and if so, whether it is empty or full.

- *Prioritization of premier and carryover customers.* Customers classified as carryover from the previous day or designated as premier customers are assigned high priority in meeting customer service commitments. These customers are scheduled for service in the morning whenever possible. In cases where the number of available vehicles is insufficient to handle all customer requests for the day, priority is given to premier and carryover customers to ensure that they are serviced and that their customer commitments are met.

The presence of practical considerations, such as tandem, carry-can, and service-type alternation requirements, combined with the inherent complexity of multiple service types distinguishes the roll-off

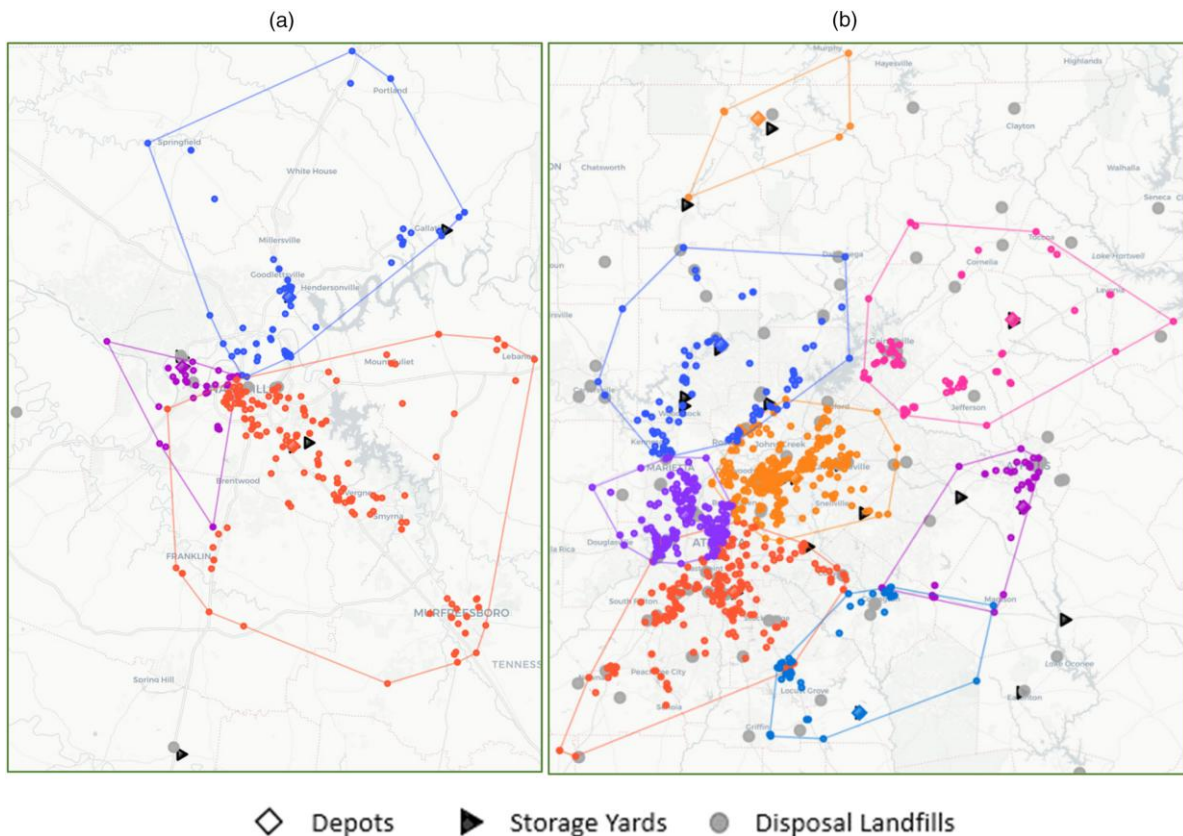
routing problem from conventional vehicle routing problems.

Cross-Site Optimization

Historically, most of our more than 460 hauling sites manage the routes/customers at a single-site level, with drivers and trucks serving as resources for routes from that site. WM operates in several large metropolitan areas (Figure 7) with multiple sites that are geographically close together, allowing drivers from one site to potentially service tickets from a different site.

Cross-site optimization offers the ability to treat one or many individual sites in neighboring geographies, cutting across current artificial boundaries as one cohort to pool work, drivers, and trucks (capacity pooling). This enables better capacity utilization and route planning across multiple sites while improving available driver and truck capacity imbalances and asset utilization across multiple sites. On the other hand, optimizing cross-site operations can significantly increase the problem size. For instance, some cross-site operations involve more than 1,000 customer service requests or hauls. Therefore, the

Figure 7. Examples of Geographies with Several Single Sites, Which Can Be Combined for Cross-Site Optimization



Notes. (a) The cross-site example features three sites with a total of 307 tickets. (b) The cross-site example features eight sites with 1,225 tickets.

solution approach must be able to handle large-scale problems efficiently while ensuring operational constraints are met.

Literature Review

As we state in Introduction, the waste collection industry encompasses three primary business lines: commercial, residential, and industrial (i.e., roll-off). In commercial waste collection, vehicles travel between commercial customers until the vehicle is fully loaded. Upon reaching capacity, the vehicle must visit a disposal facility to unload its contents and subsequently, return to the route. Typically, each vehicle undertakes two to three disposal visits per day. In residential waste collection, vehicles traverse residential streets to collect garbage. Because of these operational differences, commercial routes are modeled as node routing problems, whereas residential routes are treated as arc routing problems (Assad and Golden 1995, Kim et al. 2006, Bonomo et al. 2012, Ramamoorthy and Syrotiuk 2024).

Although both commercial and residential waste collection vehicle routing problems are treated as variants of the vehicle routing problem with time windows (VRPTW), the RRVRP presents unique challenges because of its complex constraints and specialized operational requirements.

Earlier studies primarily addressed simplified RRVRP variants, such as single-depot, single-facility scenarios. De Meulemeester et al. (1997) proposed a mathematical model addressing two customer service types—DEL and DNR. They proposed an exact enumerative algorithm and two simple heuristics. Bodin et al. (2000) defined the roll-on/roll-off vehicle routing problem, focusing on a single depot and a single landfill; their model included four customer service types—DEL, DNR, E/R, and S/O—and proposed a mathematical formulation, two lower bounds, and four heuristic algorithms. Baldacci et al. (2006) extended the model by accommodating multiple landfills and yards, adding a fifth service type—relocate (REL). They model the problem as a time-constrained VRP on a multigraph and develop an exact algorithm.

Later studies introduced more advanced approaches to address additional complexities. Raucq et al. (2019) employed a column generation approach for a tandem RRVRP with four service types and achieved notable improvements. Chiussi et al. (2022) formulated an RRVRP as a mixed-integer linear program and applied it to a real-world case in Italy. However, their model was limited to a maximum of 40 customers and did not consider carry-can operations or service-type alteration.

Heuristic and metaheuristic approaches have also been applied to RRVRPs. Wy and Kim (2013) explored

a scenario involving seven service types and service-type alterations (e.g., switching from E/R to S/O) using heuristic and metaheuristic algorithms. However, their methods did not address tandem or carry-can operations and considered fewer service types than required for practical deployment. Rabbani et al. (2016) applied simulated annealing metaheuristics to models with 10 service types. However, this approach did not account for time windows, tandem or carry-can operations, or service alterations between E/R and S/O.

To the best of our knowledge, no existing research or practical application documented in the literature fully addresses the additional complexities of RRVRPTW, such as tandem operations, carry-can requirements, large-scale instances, or service-type alterations. Given the large scale of our problem—serving more than 1,000 customer services per site across all sites—and run-time limitations because of business requirements, a heuristic-based approach offers the most practical solution. This strategy balances solution quality with computational efficiency, making it viable for enterprise-level deployment in a company of WM's size.

Analytical Methodology and Solution Approaches to Route Optimization

We used a two-step process of analytics-based demand forecasting and capacity planning followed by route optimization for the industrial line of business. The two-step approach provides the advantage of allowing local operations to meet demand expectations and customer commitments, and it accounts for anticipated anomalies, such as special event work (e.g., a rodeo in Houston or a local golf tournament), drivers calling in sick, or local weather challenges. In addition, route optimization improvements included analytically derived inputs to route optimization algorithms, such as origin-destination (OD) matrices for travel time and path estimates and the development of the optimization algorithms.

Demand Forecasting and Capacity Planning

The dynamic nature of customer demand makes the capacity planning process for the number of drivers and trucks needed every day to meet our customer commitments difficult for our planners. The number of routes is an input to the optimization step, and poor capacity planning relative to the hauls to be serviced could result in suboptimal routing solutions.

Routers are responsible for planning the number of routes required prior to the automated optimization runs each day. A route count recommendation based on the daily haul volume forecast is calculated using

two equations:

$$\text{Hours required} = \frac{\text{Hauls}}{\text{Planned Site Efficiency}}, \quad (1)$$

$$\text{Routes required} = \frac{\text{Hours required}}{\text{Average Route Length}}, \quad (2)$$

where “planned site efficiency” is the desired or target efficiency expected to be met by all of the drivers at the site and is measured as the total number of hauls completed per hour. The resulting recommended number of routes can be overridden by the operators, which has historically been done frequently because of a lack of trust in recommended haul forecasts, driven by a desire to avoid customer disruptions.

As we describe in Introduction, estimating daily roll-off haul volumes poses many challenges because of the variability across sites, such as volumes on specific days of the week. Prior to this program, WM used a trimmed moving-average method to estimate daily volumes. For each day of the week, volumes of the 10 prior weeks’ actual hauls were collected, the top and bottom two values were trimmed, and the simple average of the remainder was used as the forecast. Although it was easy to calculate and well understood by site routers, this approach did an inferior job of adjusting for seasonal variations in demand or accounting for special events.

To help improve the quality of the forecasts, we implemented autoregressive integrated moving-average (ARIMA) forecasting, which allowed us to react quickly to trends and improve our handling of seasonal effects. We built the ARIMA models separately for each day of the week and generated forecasts for a 14-day horizon based on actual haul data from the previous two years. Different types of forecasts were produced to account for daytime volumes and nighttime volumes as well as same-day requests for service from customers.

Generation of Travel Time Estimates and Route Paths

In WM logistics operations, safety is our top priority, and this principle must be reflected in route optimization and the resulting routes created. OD matrices with travel time and distance calculations between stops are critical inputs to route optimization algorithms; therefore, travel time estimates for routes should be calculated accurately for travel by our collection trucks on the appropriate roads with the correct speed assumptions. WM trucks are heavy duty compared with passenger vehicles and can carry up to eight-ton loads. They need roads with the ability to accommodate trucks’ dimensions, width, and weight as well as high clearance to avoid overhead bridge

strikes. The ability to design routes for safe roads, bridges, and driving speeds and to meet safety considerations is important to ensure that drivers can meet the planned route times while safely navigating the roads. We performed extensive analytical work based on multiple data sources, including millions of global positioning system data points from our trucks’ travel paths to calibrate speeds for our trucks, to ensure that travel paths and time estimates were geared to our industrial trucks. In the absence of commercial tools capable of meeting our specific business and scalability requirements, we developed proprietary routing algorithms and tools to generate OD matrices and route paths in a scalable way for our operations, incorporating the analytically derived speed data for our trucks and avoiding roads with bridge height and weight limitations not suitable for our trucks to ensure safe movement of our drivers on the roads.

Route Optimization Solution Methodology

A special instance of the RRVPTW, where all customer demands require only the REL service, reduces to the classical VRPTW, a problem widely recognized as NP hard. Consequently, the RRVPTW is also classified as NP hard, making the development of a theoretically efficient exact algorithm unlikely.

The RRVPTW poses significant challenges because of the specific service or operational requirements associated with customers, which extend beyond simple visits typically observed in standard vehicle routing problems. Modeling these complexities using mathematical programming is highly challenging. A common approach to address this challenge is to simplify the problem by transforming it into a traditional VRPTW. This transformation involves consolidating operational requirements for each customer based on that customer’s service type. For instance, an E/R service type (e.g., from the customer to the landfill and back) can be represented as a single node, with the entire service time (including travel) treated as one unit of service time. This approach reduces the problem’s complexity and facilitates its formulation, but it cannot capture all operational requirements. We provide the details of this transformation and the resulting simplified formulation in Appendix A.

Once the OD matrix has been prepared along with the estimated number of vehicles, the roll-off optimizer is employed to solve the roll-off routing problem. The objective is to minimize the total route time of all of the given vehicles. Table 2 outlines the additional input data required for the optimization process.

We employed a two-stage approach to solve the roll-off optimization problem as Figure 8 shows. Both stages use the same metaheuristic algorithm based on large neighborhood search (LNS) as we explain below.

Table 2. Input Data Required for Roll-off Routing Optimization

Data category	Data components
Vehicles	Route start/end times, route start/end locations, max/min route duration, special access ^a vehicle flag, tandem vehicle flag, beginning-of-the-day status
Disposal facilities	Locations, operating windows (first/second), ^b waste types, disposal service times
Customer tickets	Ticket time windows (first/second), ^b ticket mandate disposal locations, special access ^a ticket, preassigned ticket to vehicles, ticket service codes, ticket E/R and S/O swappable flag, ticket category (premium/normal), ticket type (tandem/normal), ticket waste type, ticket container type, ticket service times, special access tickets, carry-can available flag
Storage yards	Locations, hours of operation, inventory per container type/size, service times
Depot	Locations, pre-/postroute times
OD data	The OD matrix (time and distance) for each stop (defined by depot, ticket, disposal facility, and container yard)

^aSpecial access applies to customers and vehicles involved in servicing restricted locations (e.g., airports), requiring vehicles to have special permissions.

^bTime windows (first/second) indicate that two separate operating time windows are considered for all facilities and customers.

The first stage is to solve the problem of routing only tandem tickets with tandem-capable vehicles. After we create an initial solution, we solve the problem using the LNS-based metaheuristic algorithm. If unassigned tandem tickets or tandem vehicles remain—because of an excessive number of tandem tickets or tandem vehicles—these unassigned tandem tickets and unused tandem-capable vehicles are labeled as normal tickets/vehicles and considered in the second stage. The second stage starts with creating an initial solution using all of the unassigned tickets and all of the available vehicles. Those vehicles that already have tandem tickets assigned are not considered as available vehicles. Following the creation of this initial solution, we apply LNS-based improvements to refine the solution. A postprocessing phase follows the completion of the second stage. This phase ensures that the generated solutions meet all constraints, such as maximum vehicle (driver) hours and customer time windows.

Although local search heuristics improve a solution through incremental adjustments, the LNS-based algorithm that we employ in both stages applies significant perturbations to explore larger neighborhoods in a single iteration (Shaw 1998). LNS has demonstrated

strong performance for various vehicle routing problems (Ropke and Pisinger 2006, Pisinger and Ropke 2007). The overall procedure is summarized in Table 3.

The algorithm begins with a preprocessing phase, in which the necessary parameters are initialized (step 0 in Table 3). An initial solution is generated by randomly assigning tickets to vehicles while considering the vehicles’ status at the start of the day (step 1 in Table 3). In each iteration, parameters are updated (step 2 in Table 3), and the solution is perturbed by randomly removing a subset of tickets and reassigning them to different vehicles (step 3 in Table 3). To improve the perturbed solution, an intraroute improvement algorithm is run (step 4 in Table 3), employing heuristics, such as swapping, two-opt, and group ticket moves within a route. Further refinement is achieved through interroute improvement (step 5 in Table 3), which involves moving single or grouped tickets between routes and shuffling tickets across routes to identify better solutions. Service-type alterations are explored during both intraroute and interroute improvement steps. After each iteration, the best solution is updated (step 6 in Table 3), and the algorithm evaluates whether exit conditions are met (step 7 in Table 3).

Figure 8. We Solved the Roll-off Routing Problem with Tandem Tickets Using a Two-Stage Approach

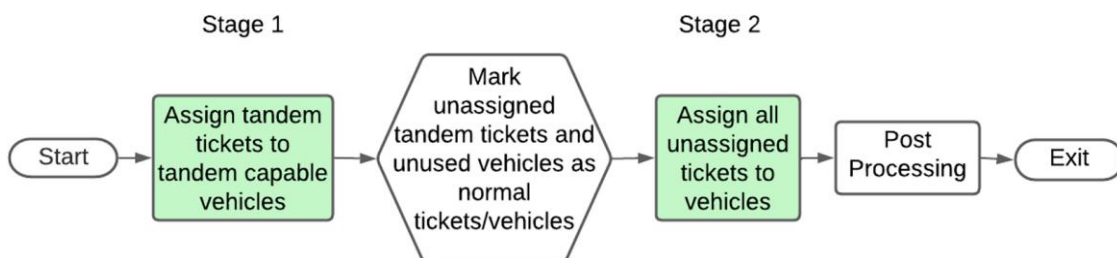


Table 3. LNS-Based Algorithm for Roll-off Routing

Steps	LNS-based algorithm for roll-off routing
0	Given parameter: I_{iter} = iteration count (set to 0) I_{none} = nonimprovement iteration counter (set to 0) I_{max_none} = maximum number of nonimprovement iterations C = number of intensification cycles, C_{count} = intensification cycle counter (set to 0) $F_{intensification}$ = intensification flag (set false), T_{max} = maximum run time
1	Create an initial solution by randomly assigning tickets to vehicles while satisfying beginning-of-the-day status of vehicles; set best solution = initial solution
2	Set $I_{iter} = I_{iter} + 1$ If $C_{count} \geq C$: $C_{count} = 0$, $F_{intensification} = !F_{intensification}$ Otherwise: $C_{count} = C_{count} + 1$
3	Randomly remove some tickets from some vehicles and then, randomly reassign tickets to vehicles
4	Perform intraroute improvements, and try E/R \leftrightarrow S/O service-type change
5	If $F_{intensification} = \text{true}$: perform interroute improvements, and try E/R \leftrightarrow S/O service-type change
6	If current solution < best solution: best solution = current solution, and set $I_{none} = 0$ Otherwise: $I_{none} = I_{none} + 1$
7	If run time < T_{max} and $I_{none} < I_{max_none}$: Go to step 2
8	Exit

Note. The algorithm's objective is to find the solution that minimizes the total route time.

Figure B.1 in Appendix B shows the flowchart of the LNS algorithm. To enhance the algorithm's performance, we implemented multithreading to leverage parallelism, thereby reducing optimization time, expanding the exploration of the solution space, and reducing the overall run time. Specifically, steps 3–5 in Table 3 are executed using multithreading techniques.

Incorporating Tandem and Carry-Can Operations

In both the route construction and the LNS phases, the decision to perform tandem or carry-can operations is made after determining the sequence of tickets in a route. For instance, consider a vehicle v that is assigned tickets t_1, t_2, t_3, t_4 in that order. If v is capable of tandem operations and all four tickets are tandem capable, five potential routes are evaluated:

1. tandem (t_1, t_2) followed by t_3 and t_4 ;
2. tandem (t_1, t_2) followed by tandem(t_3, t_4);
3. t_1 followed by tandem(t_2, t_3) and t_4 ;
4. t_1 followed by t_2 and tandem(t_3, t_4); and
5. t_1, t_2, t_3 , and t_4 without tandem operations.

The route with the minimum route time is selected as the final route for v , including tandem operations.

Similarly, carry-can operations are assessed after the sequence of tickets in a route is determined. For example, consider a route assigned to vehicle v with tickets $(\dots, t_{i-1}, t_i, t_{i+1}, \dots)$, where t_{i-1} and t_{i+1} are S/O tickets with identical container types and sizes and t_i is an E/R ticket. In this case, the route time is calculated for both possible scenarios: (1) performing a carry-can operation that includes t_{i-1}, t_i, t_{i+1} and (2)

not performing the carry-can operation. The option with the shortest route time is selected, and the corresponding status (i.e., whether the carry-can operation was performed) is recorded for vehicle v .

Evaluating the Performance of the Solution Approach

Benchmark test data sets specific to our problem are not readily available in the literature. To validate the performance of our algorithm, we applied it to actual WM data sets over extended periods and compared the calculated route time of the optimized solutions against those of manually sequenced routes. Table C.1 in Appendix C shows a detailed comparison of the manual versus optimized solutions.

Additionally, Table 4 provides examples of the time required to find solutions within 1% of the best solutions evaluated by route time and obtained from extensive runs. Although this is not proof of optimality, it indicates that our method's convergence behavior aligns with that of well-designed metaheuristic algorithms.

The performance test results provided insights into determining appropriate algorithm run-time settings aligned with business processes to produce routes in a reasonable time. In our practice, data sets are classified based on the number of vehicles, with each category assigned a specific maximum run time. These designated run times, which range from 15 to 60 minutes, serve as part of the exit criteria outlined in the LNS algorithm.

Table 4. Examples of Route Optimization Solution Performance on Various Problem Sizes

Problem size	No. of routes	No. of tickets	Run time (in minutes) to reach within 1% of best ERT ^a solution	ERT ^a (in minutes)
Small	12	60	<1	480
Medium	33	167	2	480
Large	49	323	31	720
Extra large ^b	155	1,151	41	900

Note. All results were generated on an AWS Graviton 3 server using eight threads.

^aERT solution refers to the best solution (in terms of route time) obtained from extensive computational runs, serving as a benchmark for solution quality.

^bExtra-large size is from a cross-site operations data set that merges nine nearby individual sites into a single cross-site run.

Implementation

The dynamic route optimization program took over three years to implement beginning in June 2021, with the completion of the enterprise rollout in quarter 3 of 2024. The program went through multiple phases of data gathering and centralization, design, development, proofs of concept, pilots, field training, business process improvements, deployment, and change management steps. The design and delivery relied on the support of WM executives and leaders, and they required inputs, participation, and support from functional experts and leaders in several functional areas within the company. Two previous attempts to implement route optimization technology in the industrial line of business failed to gain stakeholder acceptance and traction for multiple reasons, including technology gaps, a huge deficit of data (digitally and centrally) for optimization, consideration of various practical business needs, business processes, and lack of trust in the quality as well as the executability of the routes produced. We had to be more deliberate and thoughtful in our approach in the third attempt at incorporating the lessons learned from previous attempts while achieving the additional goal of automation. In this effort, our operational leaders and the routers were included in the design and development stages of the program. We focused significant efforts on the identification, digitalization, and centralization of data elements needed in addition to the building of data governance and quality controls. The most critical part of this effort was achieving buy-in from our frontline routers and drivers on route design, quality, and execution feasibility of the routes from our optimization algorithms. This required being transparent with field users, including providing the details of the algorithms. We made iterative improvements, utilizing insights and inputs from our routing and dispatch function field leaders with responsibility for planning and executing the routes safely.

In the initial stages of the project, the algorithms were tested by optimizing routes from 25 pilot sites

specifically selected to represent a diverse set of sites varying operational sizes and geographies across 14 of our 16 market areas; we analyzed these routes to compare the planned efficiency (measured as the number of hauls expected to be completed per unit hour) of manually created routes against routes generated from optimization algorithms (Appendix C). The study showed that dynamic route optimization-generated routes were 10.3% more efficient on an hours per haul basis and 11.6% better in miles per haul while meeting all business constraints compared with manually created routes. In contrast, manually created routes in this sample set violated time-window requirements over 11% of the time and maximum route time constraints by over 28% of the time. The planned efficiency improvements of optimized routes over manually created routes, albeit on a representative sample of routes, gave our stakeholders initial confidence in the value of route optimization. In practice, some planned route efficiencies—whether from manually created routes or optimized routes—will not be fully realized following actual driver execution through actual efficiency realization because of several factors. These include factors such as the driver’s skill-set level and experience and the driver’s ability to meet customer service requirements and route time expectations. Some other factors, such as driver staffing issues because of last-minute work absence notifications, difficult conditions on the road because of weather and traffic, longer wait times at facilities like landfills relative to planned times, and service-related issues at customer locations like blocked containers, also contribute to planned efficiencies not being fully realized and being reflected in actual efficiencies. Some of the controllable factors are managed typically through ongoing process improvements, including driver performance and coaching sessions to ensure full realization of planned efficiencies as much as possible after route execution. Gaps in actual efficiencies relative to planned efficiencies following actual execution tend to improve gradually as

drivers gain more familiarity with their new routes, a process that could take several months. As an essential service provider, our frontline team members prioritize service completion and customer satisfaction while improving efficiency.

We accomplished our automation objectives (Figure 9) through the design and development of scalable and highly reliable analytical software engines designed to automate the end-to-end process of route creation and delivery for execution.

The ARIMA forecasts generated with fine-tuned model parameters reduced the mean absolute percent error (MAPE) to 25% from the 30%–35% MAPE via the incumbent, trimmed-means approach. Further improvements to demand forecasting and capacity planning are in progress with two new models in development, including a time-series forecasting model using Prophet (Taylor and Letham 2018). Prophet provides several advantages over the current ARIMA model; these include explicit support for holidays, correction for major weather events, and piecewise trending to better handle seasonal special events and changes because of acquisitions/divestitures. Initial estimates show an MAPE of 15% from the Prophet models as compared with the incumbent ARIMA models with an MAPE of 25%. In addition, development is complete and testing is underway on a machine learning-based model, which will recommend the number of routes required to reduce subjectivity and increase consistency. Combined with the planned forecast model improvements, we expect these changes to further improve the route-count recommendations and the capacity planning process.

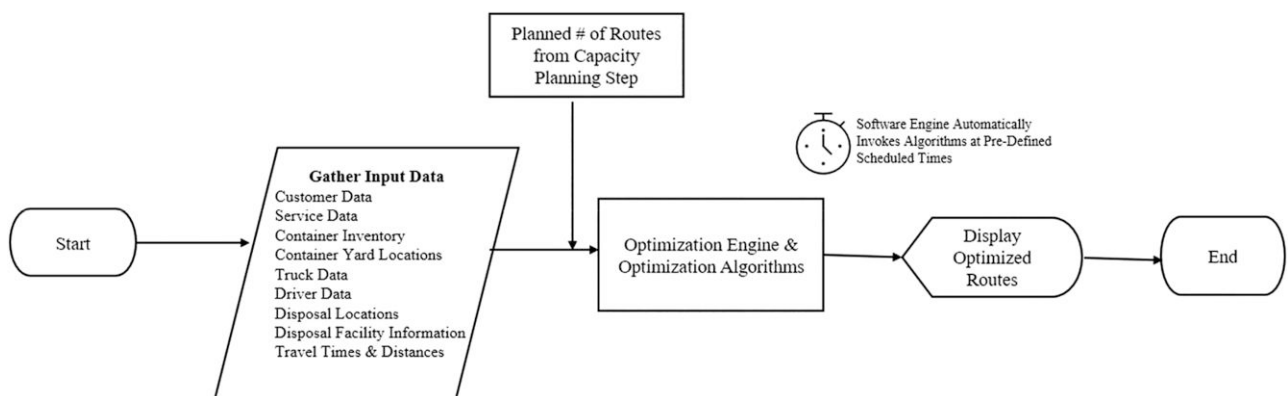
Benefits and Impact

During the period following the coronavirus disease 2019 pandemic, WM Chief Executive Officer (CEO) Jim Fish set a goal of turning our focus to manage for

the future using those challenging times as an opportunity to further differentiate ourselves through permanent enhancements in our business processes and service offerings to our customers. WM's overarching business strategy includes heavy investments in technology-led automation in multiple areas of our business to improve the overall cost structure and as a hedge against labor shortages in hard-to-fill labor positions in our industry. The industrial line of business, which generated close to 14% of WM's revenue (~\$3.1 billion) in 2024, constitutes an important part of our business.

The dynamic route optimization solution that we implemented is geared to design optimized routes for 100% of our roll-off routes, building efficient routes consistently and replacing manual processes. Our business improvements focus on safety, service, and savings. We measure these metrics year over year. In the area of safety, a 50% reduction in bridge overhead strike safety incidents was realized in our industrial line of business by using the routes designed for travel on roads appropriate for our trucks assisted by commercial truck navigation technology for our drivers. Industrial customer service reliability metrics in 2024 improved; occurrences of fulfilling service requirements at customer-requested service time windows of less than four hours improved by 10%. The efficiency-related savings consider labor hours as well as fuel and maintenance on our trucks. Using this first-of-its-kind technology suite and with near-complete enterprise deployment of our dynamic route optimization program in quarter 3 of 2024, WM's industrial line of business realized efficiency improvements of 1.24% year over year, which translated to an annual benefit of approximately \$11 million savings in 2024. In addition, savings of \$1 million were realized from reductions in some related job positions. A more significant benefit

Figure 9. Over 70 Data Elements Are Needed to Optimize Roll-off Routes



Notes. Optimization algorithms are invoked by automation engines. Recall that the planned numbers of routes are derived from the daily haul volume forecast for the site using given values of planned site efficiency and average route length.

was the expansion of margins driven by both price growth and operational expenses reduction. Our margins have improved 157 basis points year over year, which is the highest ever in the industrial line of business.

We expect to expand the benefits realized in the coming years by building on the initial success of the program. Following the completion of the single-site deployments, implementations are underway to roll out and scale cross-site optimization to drive better capacity pooling and efficiency gains. Sites at several urban areas, which amount to approximately 20% of the total number of industrial routes in large metropolitan areas, have been identified as primary opportunities for cross-site optimization. Benefits from cross-site optimization can include efficiency improvements as well as improved asset utilization through better resource pooling. We have piloted cross-site deployment at three sites, which generated efficiency improvements of 2%–3% in addition to improved asset utilization and reduced customer service carryovers. Field pilots are underway to implement total cost optimization considering overall operating costs as well as reductions in disposal costs achieved by moving these disposal services to internal WM facilities from more expensive third-party-owned facilities. Improvements from total cost-based optimization deployments over the next few years are expected to contribute from \$25 million to \$42 million in cost improvements. In addition, through sustained 2% improvements in industrial efficiency, we envision a reduction in new truck capital requirements from approximately \$26 million to \$30 million over the next several years based on the replacement life cycle of our trucks.

Other Benefits

The importance of good data required for strengthening analytics is now well understood across WM in multiple areas of the business. This program provided an impetus to create an enterprise data warehouse with different data marts for various functional areas as well as for collecting and centralizing data that were previously not available and digitized. These foundational capabilities have expanded our business intelligence insights, helped democratize data within our enterprise, and positioned us well for future data-driven insights. The evolution, development, and deployment of the program's capabilities have contributed to opportunities to improve our previous organizational design and roles. In addition to the creation of a business optimization function, organizational improvements were made, and new products, data, and optimization roles were created in both our corporate and field functions

to further improve and promote closer alignment and the success of future optimization initiatives.

WM typically supplies containers to our customers in the industrial line of business. One of the ancillary benefits from the broader program included the development of a container inventory system as a source of data to the optimization engines. Our optimization algorithms can maximize S/O service containers among customers when allowed rather than acquiring a container from the container yard and thus, reduce container yard trips and container inventory demand. As a result, we can reduce and more accurately estimate the demand for new containers to inform capital needs.

Automation was another key goal to reduce reliance on human judgement, introduce consistency, remove the need for local knowledge, and bring standardization into the process. The program has standardized the industrial routing process across the enterprise and significantly reduced the time involved in manually building routes from several hundred hours (four to five hours per site) to under 45 minutes per site through optimization and automation engines. The efforts of the routers have been shifted to maintaining, monitoring, and managing the data and spending time on other operational and customer needs. This has also enabled improved segregation of roles and responsibilities in some functions, enabling the foundations for future improvements in our dispatch operations.

In April 2024, Jim Fish, CEO of WM, remarked: "It was really 4, 5 years ago an admission by us, that our routing was not as efficient as it could be and if truly used technology could benefit. I think that's probably the case across the entire industry and so 4 or 5 years ago we decided that had to change and now you are seeing the fruits of that" (Seeking Alpha 2024a).

John Morris, Executive Vice President and Chief Operations Officer of WM, commented in October 2024: "The continued adoption of scheduling and planning tools, advanced mapping systems and dynamic routing is also driving efficiency and reducing operating costs. Our data-driven business decisions and technology investments are leading to greater operational efficiency and improved return on capital, which is reflected in the growth and margin performance of our Collections and Disposal operations" (Seeking Alpha 2024b).

Transportability

Although they are not as large as the industrial line of business, some smaller and emerging lines of business/areas within WM would be a good fit for some or all of the components. These include container delivery routes for new customers, which can be optimized using simplified versions of industrial

routing algorithms, and WM Energy Services dispatched operations, which involve trucking of liquid materials to and from oil and gas fracking sites and can benefit from the staging of full and empty containers. FOG2Fuel operations, where WM collects fats, oils, and greases from businesses, such as restaurants and hotels, and transports them to a plant for conversion to bio-fuels, is also a potential candidate for this technology. Our algorithm's ability to handle diverse service types and operational constraints makes it adaptable to different industries and problem settings. For example, in the oil and gas industry, water generated during production (i.e., production water) must be hauled away and recycled, reused, or disposed of in designated disposal wells. The operation is similar to roll-off routing with an S/O service code and single waste type. Our work can be directly applied to scheduling production water hauling for large oil and gas fields.

Conclusions

As part of this project, we embarked on a transformative journey to fundamentally change how we build routes in our most dynamic and complex industrial line of business, leveraging data, advanced analytics, and optimization techniques to transition from a manual approach to an automated approach. The success of this program at WM could serve as a blueprint for organizations planning similar initiatives. The key to our success was the commitment of executives and leaders in the business. Foundational investments in digitalization, centralized data, an understanding of business rules, transparency into how operations research-based algorithms function, and extensive change management efforts enabled the transition from disparate manual processes to a standardized enterprise-wide approach. Developing the roll-off routing algorithms required addressing the practical needs of the business, providing transparency, scaling them to the enterprise, overcoming initial skepticism, and meeting the expectations of our frontline team members. The process was supported by advanced data analytics, demand forecasting, and capacity planning. Automation of the analytical processes that the business depends on for day-to-day operations also required higher data quality and scalable technology solutions as we demonstrated in this project. The impact of the continued adoption of analytics and technology-led automation initiatives by WM's team members while providing safe and reliable service to our customers is reflected in the increased growth and performance of our operations.

Appendix A. Simplified Version of the Roll-off Problem

We formulate a simplified version of the roll-off problem using the framework proposed by Chiussi et al. (2022). This

formulation includes four service types: E/R, S/O, DEL, and DNR (Table 1). The simplified problem assumes the following conditions.

1. Each vehicle can transport only one container at a time (no tandem vehicles are allowed).
2. All routes have the same start time (e.g., 8:00 a.m.) and end time (e.g., 5:00 p.m.).
3. Each customer is associated with a predefined time window.
4. A single depot stores all spare containers.
5. The number of containers available at the depot is unlimited for all container types.
6. All disposal facilities and the depot operate continuously between the route start and end times, thereby precluding varying time windows for disposal facilities.

Let N represent the set of customers, M represent the set of disposal facilities, and d represent the depot. The subsets $E, S, D,$ and R denote the customers requiring service types E/R, S/O, DEL, and DNR, respectively, such that $N = E \cup S \cup D \cup R$. Additional parameters are defined as follows.

- t_{ij} Travel time between nodes i and j , where $i, j \in V = N \cup M \cup \{d\}$.
- s_i Service time at customer location i , where $i \in N$.
- s_d Service time (container loading or unloading time) at the depot d .
- l_i Service time at disposal landfill location i , where $i \in M$.

For all service types requiring a disposal trip—E/R, S/O, and DNR—all operations are consolidated into a single service time and a travel time. This involves computing new service times S_i for $i \in N$ and new travel times T_{ij} between two nodes i and j , where $i, j \in N \cup \{d\}$. The revised service times S_i and travel times T_{ij} are detailed in Tables A.1 and A.2, respectively. For service types requiring E/R, for example, the new service time accounts for all operations, including travel from the E/R customer to the disposal facility, the disposal service time, and travel back from the disposal facility to the E/R customer location (see case 2 in Table A.1). Conversely, for S/O and DNR service types, the disposal operation time is incorporated into T_{ij} rather than S_i . This is because following the disposal trip, the vehicle does not return to the customer, allowing the empty container to be utilized in the subsequent sequence (see cases 5 and 6 in Table A.2).

This transformation enables reducing the given problem to a standard VRPTW. The original problem defined on the network $G = (V, A)$, where $V = N \cup M \cup \{d\}$, is transformed into $G' = (V', A')$, where $V' = N \cup \{d\}$ and A' is a subset of A . It is important to note that all disposal facilities are excluded in the reduced VRPTW problem.

Table A.1. Examples of Adjusted Service Time S_i Incorporating Disposal Facility Visit for E/R Service

Case	i	S_i
1	$i \in S \cup D \cup R$	s_i
2	$i \in E$	$\min_{m \in M} \{t_{im} + l_m + t_{mj}\}$

Table A.2. Examples of Merged Travel Time T_{ij} Between Two Nodes Incorporating Disposal Facility Visits

Case	(i, j)	T_{ij}
1	$i = d, j \in E \cup R$	t_{dj}
2	$i = d, j \in S \cup D$	$t_{dj} + s_d$
3	$i \in E \cup D, j \in E \cup R$	t_{ij}
4	$i \in E \cup D, j \in S \cup D$	$t_{id} + t_{dj}$
5	$i \in S \cup R, j \in S \cup D, \alpha_i = \alpha_j$	$\min_{m \in M} \{t_{im} + l_m + t_{mj}\}$
6	$i \in S \cup R, j \in S \cup D, \alpha_i \neq \alpha_j$	$\min_{m \in M} \{t_{im} + l_m + t_{md} + s_d + t_{dj}\}$
7	$i \in S \cup R, j \in E \cup R$	$\min_{m \in M} \{t_{im} + l_m + t_{md} + s_d + t_{dj}\}$
8	$i \in S \cup R, j = d$	$\min_{m \in M} \{t_{im} + l_m + t_{md}\}$
9	$i \in E \cup D, j = d$	t_{id}

Each node in V' has a time window $[a_i, b_i]$ representing the earliest and latest possible starting service times, respectively. The depot's time window $[O, D]$ has the special meaning of the earliest possible departure from the depot and the latest possible arrival at the depot, and this time window will also serve as the vehicle's start and end times. The depot is assumed to have zero service time.

A nonnegative merged travel time, T_{ij} , is associated with each arc $(i, j) \in A'$. Let K represent the set of available vehicles. Let x_{ijk} , where $(i, j) \in A', k \in K$, equal one if arc (i, j) is used by vehicle k and zero otherwise. Additionally, variable w_{ik} , where $i \in V', k \in K$, represents the start service time at node i when it is served by vehicle k .

The key notation used in the simplified model is provided below followed by its mathematical formulation.

A.1. Sets

- V' Set of all locations, including the depot and customers
- A' Set of all possible edges (i, j) , representing travel between locations
- N Set of all customer locations
- K Set of available vehicles

A.2. Parameters

- T_{ij} Travel time from location i to location j
- S_i Service time required at location i
- a_i Earliest arrival time (time window) at location i
- b_i Latest arrival time (time window) at location i
- O Earliest possible start time (operating time window) for any vehicle
- D Latest possible finish time (operating time window) for any vehicle
- d Index representing the depot location

A.3. Decision Variables

- x_{ijk} Binary variable: one if vehicle k travels directly from location i to location j and zero otherwise
- w_{ik} Arrival time of vehicle k at location i

A.4. Simplified VRPTW Model

$$\min \sum_{k \in K} \sum_{(i, j) \in A'} T_{ij} x_{ijk}$$

subject to

$$\sum_{k \in K} \sum_{(i, j) \in A'} x_{ijk} = 1 \quad \forall i \in N, \quad (\text{A.1})$$

$$\sum_{j \in N} x_{djk} = 1 \quad \forall k \in K, \quad (\text{A.2})$$

$$\sum_{j \in V'} x_{jik} - \sum_{j \in V'} x_{ijk} = 0 \quad \forall k \in K, i \in N, \quad (\text{A.3})$$

$$w_{ik} + S_i + T_{ij} - w_{jk} \leq (1 - x_{ijk}) \text{Big}M \quad \forall k \in K, (i, j) \in A', \quad (\text{A.4})$$

$$a_i \sum_{j \in V'} x_{ijk} \leq w_{ik} \leq b_i \sum_{j \in V'} x_{ijk} \quad \forall k \in K, i \in N, \quad (\text{A.5})$$

$$O \leq w_{ik} \leq D \quad \forall k \in K, i \in V', \quad (\text{A.6})$$

$$x_{ijk} \in \{0, 1\} \quad \forall k \in K, (i, j) \in A'. \quad (\text{A.7})$$

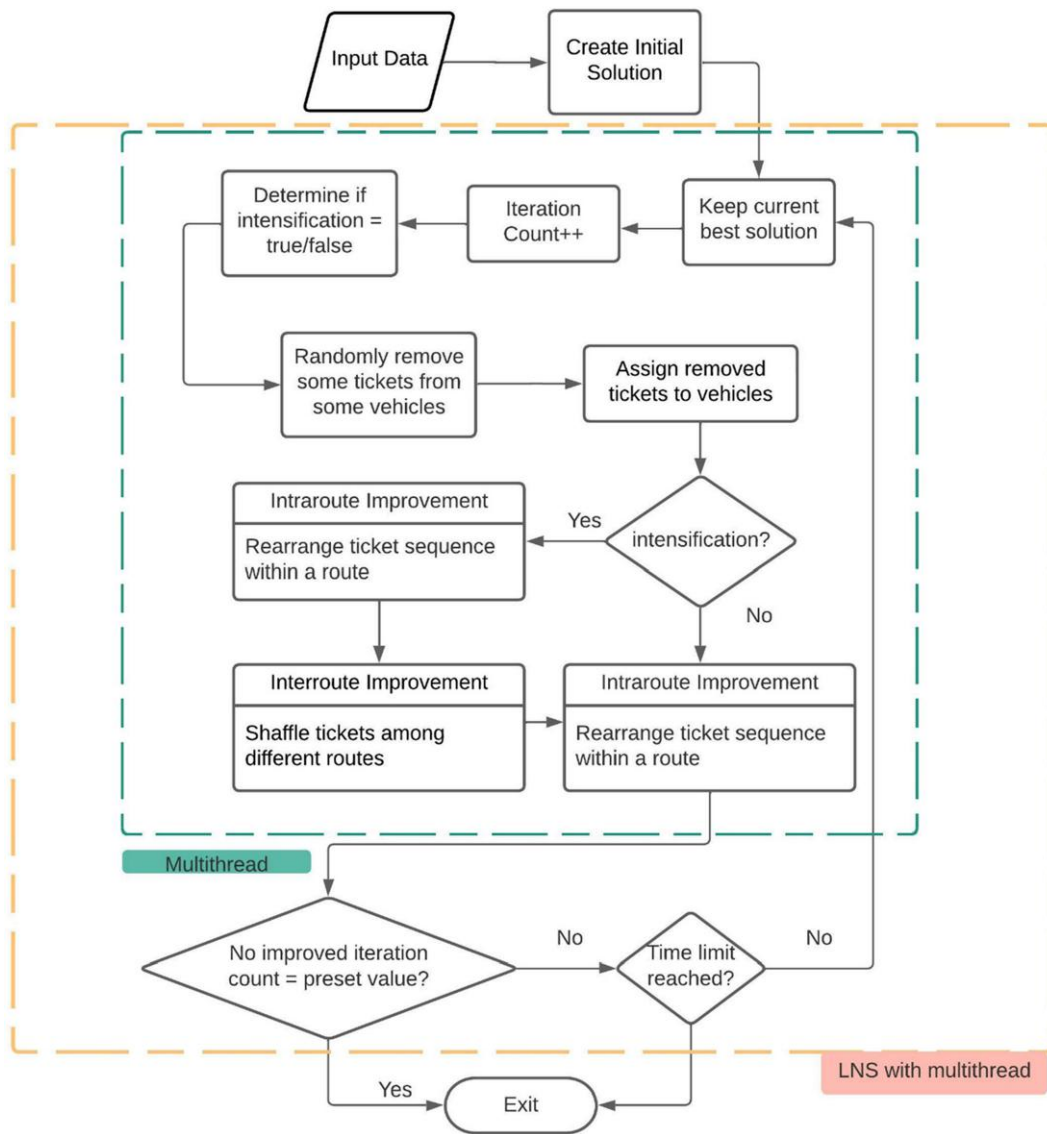
Constraints (A.1) enforce the requirement that each customer is serviced by exactly one vehicle. Constraints (A.2) ensure that each route begins at the depot. Constraints (A.3) are the classical flow conservation constraints; for any vehicle k and customer location i , if vehicle k enters location i , it must also leave, thus ensuring continuous and logical routes. Constraints (A.4), (A.5), and (A.6) ensure that the time constraints on both the route and vehicle are satisfied. These include the individual customer time windows (a_i (earliest) and b_i (latest) for customers) as well as the overall operating time window for vehicles defined by O and D . A tight Big M value, such as 129,600 (36 hours in seconds), can be approximated by the maximum possible value of the expression $w_{ik} + S_i + T_{ij}$, which is the latest possible arrival time at any node plus the service and travel times. Lastly, Constraints (A.7) impose binary conditions on decision variables.

The formulation utilizing the above transformation is inherently limited and cannot address the complexity of our problem. Specifically, it fails to accommodate tandem cases, service-type alterations, different waste types, and most critically, scenarios in which disposal facilities have distinct time windows.

Appendix B. LNS-Based Metaheuristic Algorithm for Solving Roll-off Routing Problems

The following flowchart in Figure B.1 illustrates the LNS-based metaheuristic algorithm for solving roll-off routing problems.

Figure B.1. (Color online) Flowchart of the LNS-Based Metaheuristic Algorithm for Solving Roll-off Routing Problems



Notes. Intensification refers to the phase executed a certain fixed number of times that is dedicated to rigorously exploring the neighborhood of high-quality solutions. This is achieved by applying a more thorough set of local search algorithms, including both intraroute and interrout improvements.

Appendix C. Comparison Between Manually Derived Routes and Optimized Routes for Route Times and Total Mileage

Table C.1 presents a comparison between manually derived routes and optimized routes for route times and total mileage.

Table C.1. Comparison of Manual vs. Optimize Solutions

Area site	Manual			Optimized			Improvement, %	
	Route time ^a	Eq. hauls ^b	Total miles	Route time ^a	Eq. hauls ^b	Total miles	Efficiency ^c	Total miles
EAST 001	879	492	16,300	762	492	13,636	15.3	16.3
EAST 002	332	193	7,000	290	193	5,942	14.4	15.1
EAST 003	601	308	13,234	530	308	11,939	13.4	9.8
EAST 004	31.4	21	308	30.7	21	294	2.3	4.3
EAST 005	3,496	2,707	44,025	3,261	2,707	39,059	7.2	11.3
EAST 006	819	657	10,926	730	657	9,620	12.2	12.0
EAST 007	318	129	6,569	273	129	5,398	16.5	17.8
EAST 008	196	98	3,801	156	98	2,773	25.6	27.0
EAST 009	1,491	934	33,851	1,316	934	30,281	13.3	10.5
EAST 010	369	165	9,658	333	165	8,471	10.8	12.3
EAST 011	123	75	2,152	109	75	1,842	12.5	14.4
EAST 012	159	74	3,420	147	74	3,034	8.5	11.3
WEST 001	713	380	13,758	598	380	11,155	19.3	18.9
WEST 002	845	725	13,611	777	725	12,150	8.7	10.7
WEST 003	749	601	12,200	687	601	10,419	9.1	14.6
WEST 004	385	286	5,855	347	286	5,089	11.0	13.1
WEST 005	1,096	590	19,162	1,015	590	17,518	8.0	8.6
WEST 006	603	329	11,306	577	329	10,852	4.6	4.0
WEST 007	298	213	5,488	267	213	4,943	11.5	9.9
WEST 008	472	412	5,623	457	412	5,358	3.3	4.7
WEST 009	295	214	3,706	265	214	3,353	11.4	9.5
WEST 010	443	412	5,366	402	412	4,682	10.2	12.7
WEST 011	828	445	17,246	765	445	15,689	8.2	9.0
WEST 012	638	334	15,335	571	334	13,699	11.8	10.7
WEST 013	208	113	4,972	197	113	4,675	5.7	6.0
Total	16,389	10,899	284,869	14,862	10,899	251,871	10.3	11.6

^aRoute time represents the total hours spent on all routes.

^bEquivalent hauls (Eq. hauls) refer to the number of customer stops weighted by their respective service requirements or complexity.

^cEfficiency is defined as total hauls per total route hours; this metric shows the number of hauls served per hour.

References

- Assad A, Golden B (1995) Arc routing methods and applications. Ball MO, Magnanti TL, Monma CL, Nemhauser GL, eds. *Handbooks in Operations Research and Management Science* (Elsevier, Amsterdam), 375–483.
- Baldacci R, Bodin L, Mingozzi A (2006) The multiple disposal facilities and multiple inventory locations rollon-rolloff vehicle routing problem. *Comput. Oper. Res.* 33(9):2667–2702.
- Bodin L, Mingozzi A, Baldacci R, Ball M (2000) The rollon-rolloff vehicle routing problem. *Transportation Sci.* 34(3):271–288.
- Bonomo F, Durán G, Larumbe F, Marengo J (2012) A method for optimizing waste collection using mathematical programming: A Buenos Aires case study. *Waste Management Res.* 30(3):311–324.
- Chiussi A, Felix G, Iori M, Santos AG (2022) Industrial waste collection optimization: A real-world case study in northern Italy. De Armas J, Ramalhinho H, Voß S, eds. *Proc. Comput. Logist. 13th Internat. Conf.* (Springer-Verlag, Berlin), 147–161.
- De Meulemeester L, Laporte G, Louveaux FV, Semet F (1997) Optimal sequencing of skip collections and deliveries. *J. Oper. Res. Soc.* 48(1):57–64.
- Kim B, Kim S, Sahoo S (2006) Waste collection vehicle routing problem with time windows. *Comput. Oper. Res.* 33(12):3624–3642.
- Pisinger D, Ropke S (2007) A general heuristic for vehicle routing problems. *Comput. Oper. Res.* 34(8):2403–2435.
- Rabbani M, Tabrizi H, Farrokhi-Asl H (2016) A hybrid metaheuristic algorithm for solving a roll-on roll-off waste collection vehicle routing problem considering waste separation and recycling center. *Internat. J. Appl. Oper. Res.* 6(2):19–31.
- Ramamoorthy M, Syrotiuk VR (2024) Learning heuristics for arc routing problems. *Intelligent Systems Appl.* 21(March):200300.
- Raucq J, Sörensen K, Cattrysse D (2019) Solving a real-life roll-on-roll-off waste collection problem with column generation. *J. Vehicle Routing Algorithms* 2(1):41–54.
- Ropke S, Pisinger D (2006) An adaptive large neighborhood search heuristic for the pickup and delivery problem with time windows. *Transportation Sci.* 40(4):455–472.
- Seeking Alpha (2024a) Waste Management, Inc. (WM) Q1 2024 earnings call transcript. Accessed August 9, 2025, <https://seekingalpha.com/article/4686256-waste-management-inc-wm-q1-2024-earnings-call-transcript>.
- Seeking Alpha (2024b) Waste Management, Inc. (WM) Q3 2024 earnings call transcript. Accessed August 9, 2025, <https://seekingalpha.com/article/4730577-waste-management-inc-wm-q3-2024-earnings-call-transcript>.
- Shaw P (1998) Using constraint programming and local search methods to solve vehicle routing problems. Accessed July 1, 2025, <https://neo.lcc.uma.es/radi-aeb/WebVRP/data/articles/CP4VRPshaw98.pdf>.

Taylor S, Letham B (2018) Forecasting at scale. *Amer. Statistician* 72(1):37–45.

Wy J, Kim B (2013) A hybrid metaheuristic approach for the rollon-roll-off vehicle routing problem. *Comput. Oper. Res.* 40(8):1947–1952.

Hemachandra Pillutla is the senior director of innovation and business optimization at WM leading the decision science function. His team is responsible for the development of operations research, machine learning, artificial intelligence, and advanced analytical capabilities in the areas of collection, fleet, disposal, recycling and sustainability enabling business strategy, optimization and automation outcomes. He has advanced degrees in engineering and management. Hema is co-inventor for ten patents.

Seongbae Kim is a senior strategist in business optimization at WM leading the operations research function. He is responsible for developing operations research algorithms focused on network, facility, and logistics optimization. Under his leadership, his team successfully developed core algorithms for dynamic route optimization projects over the past several years. He holds MS and PhD degrees in industrial engineering with a specialization in operations research from Texas A&M University. He is a 2004 Franz Edelman laureate and a coinventor for two patents.

Chao Zhou is a principal operations research scientist at WM, where he coleads the research, design, development, and support of algorithms embedded in automation engines for the dynamic route optimization project. With 20+ years of experience, he has developed advanced solutions for large-scale scheduling, vehicle routing, and network optimization challenges across various industries. He received his PhD in transportation and infrastructure systems engineering from Purdue University. He is the sole winner of the 2003 Milton Pikarsky Award from the Council of University Transportation Centers.

Heesu Hwang is a principal operations research scientist at WM. He is responsible for developing algorithms aimed at reducing the company's operational costs. He coleads the research, design, development, and delivery of algorithms embedded within automation engines, driving machine learning-based capacity planning as part of the dynamic route optimization project. He earned his PhD in industrial engineering from the University of Texas at Arlington.

Sergiy Savchenko is a principal operations research software architect at WM responsible for developing high-performance and scalable routing services as well as advanced geospatial analytics. He leads the design and delivery of routing algorithms for origin-destination matrix calculations and route path generation. These solutions support automation engines for safe and efficient routing in the dynamic route optimization project. He earned his master's degree in electromechanical engineering from Kharkov Aviation Institute. He is a coinventor for eight patents.

Stuart Greene is a data science manager at WM responsible for leading forecasting and advanced analytics initiatives in collection and fleet operations and network design. Over his 20 years at WM, he has supported the collection operations and digital teams as a leader and subject matter expert. He holds a bachelor's degree in naval architecture and marine engineering from Webb Institute and a master's degree in industrial engineering from the University of Houston.

Marcel Dalby is the vice president of collections operations and business optimization at WM overseeing the optimization efforts across company's collection, fleet, disposal, industrial automation, and recycling operations. His team is responsible for increasing efficiencies and improving overall customer service through the implementation of innovative technologies and improved business processes. He holds a Bachelor of Business Administration Management degree from Texas State University.