mxdusa.org
@mxdinnovates
info@uilabs.org

1415 N. Cherry Avenue
Chicago, IL 60642
(312) 281-6900

Manufacturing x Digital

MxD FINAL REPORT PROJECT 17-02-01

SUPPLY CHAIN RISK ALERT	
PRINCIPLE INVESTIGATOR / EMAIL ADDRESS	Jeremy Archbold, jarchbold@dow.com
PROJECT TEAM LEAD	Dow, Inc.
PROJECT DESIGNATION	17-02-01
MxD CONTRACT NUMBER	042018003
PROJECT PARTICIPANTS	Indiana University Purdue University-Indianapolis, Rochester Institute of Technology, ITAMCO, Microsoft
MxD FUNDING VALUE	$499,641
PROJECT TEAM COST SHARE	$628,415
AWARD DATE	July 9, 2018
COMPLETION DATE	December 22, 2019

SPONSORSHIP DISCLAIMER STATEMENT: This project was completed under the Cooperative Agreement W31P4Q-14-2-0001, between U.S. Army - Army Contracting Command - Redstone and UI LABS on behalf of the Digital Manufacturing and Design Innovation Institute. Any opinions, findings, and conclusions or recommendations expressed in this material are those of the author(s) and do not necessarily reflect the views of the Department of the Army.

TABLE OF CONTENTS

mxdusa.org
@mxdinnovates
info@uilabs.org

1415 N. Cherry Avenue
Chicago, IL 60642
(312) 281-6900

EXECUTIVE SUMMARY

The purpose of this report is to document MxD Project 17-02-01 and share the deliverables, technical approach, results, discussion, analysis, and conclusions. This project worked to improve and further automate the Supply Chain Event Management (SCEM) process by creating a framework to digitize, integrate, and automate the information pipeline and action workflow as well as offer recommendations based on prior mitigation actions. The primary use case is integration within a supply chain including an event management team, logistics teams, business decision makers, and customer service representatives, where event mitigation actions can be determined using support data, tracked for cycle-time and where the output can be utilized to make faster informed decisions in similar situations in the future.

Work on this project was split into modules that perform independent actions and when linked, deliver the desired results. The modules are:

- *Predictive Transit*: This module provides the shipment planner with an estimated transit time for future shipments based on source, destination, planned shipment date, product type, weather, and event data.
- *Risk Assessment*: This module provides a graphical user interface to allow a user to document current or future events and automatically compute a risk score for individual shipments in the affected area.
- *Mitigation Planning*: This module automatically sends a communication (text/SMS or email) to subscribed decision makers when a relevant shipment has been promoted from the Risk Assessment module. The communication gives them an overview of the shipment and the event impacting it. The user can enter mitigation information that is then recorded, storing the knowledge. After enough mitigation decisions are collected, a machine learning method can be trained to then automate recommended actions.
- *SIMBA Chain Communications*: This module integrates relevant data from different modules using a blockchain ledger and sends automated notifications to targeted individuals when risk thresholds are met.
- *Performance Analytics*: This module is a data enablement SQL table with the Azure framework with direct connection to a dashboard that calculates and displays Key Performance Indicators (KPIs) which are derived from the various modules. It provides aggregated data and an easy access point for monitoring and decision making.

The above modules were tested using Dow historical shipment data in the same form that data would be expected when being deployed for full use in daily operation. In addition to the above modules, this project delivered: technical demonstrations, a report on existing technology, documentation and user guides for the modules, a commercialization plan report, MxD technical deliverable acceptance checklists, a half day seminar, and this report.

This project delivered the expected functionality but was limited by a lack of a live production environment that would have allowed for a more complete demonstration of the mitigation planning module recommender system. The further exploration of a future pilot will be part of Dow’s next steps activities.

MxD members can access all the software on the MxD membership portal according to the membership agreement post-project. SIMBA Chain offers commercial Smart Contract as a Service (SCaaS) software for blockchain integration. The mitigation planning module and performance analytics modules are deployed as examples only and may be deployed using any preferred dashboard or web tools. The predictive Transit module requires least two years of historical shipment data .

PROJECT DELIVERABLES

The following list includes all deliverables created through this project. These deliverables will be referenced throughout this report and should be accessible on the MxD membership portal in accordance with the rights defined by the Membership Agreement. Specific deliverable types include, but are not limited to, the following items.

Table 1: Project Deliverables

#	DELIVERABLE NAME	DESCRIPTION	FORMAT OF DELIVERY
1	Report on Existing Technology_11.13.19	High-level overview of existing technologies that were evaluated for project integration	.docx
2	MxD Project Final Presentation_FINAL.pptx	Complete final presentation slide deck.	.pptx
3	Mitigation Planning Module	Source code for Web UI, APIs source code for pulls/pushes from SQL, example SQL structure	Folder with readme, licensing.txt, documentation and zip folders with project files
4	Predictive Transit Module	Admin UI, User UI, ML algorithms source code, Executables, example input and output tables, APIs source code for an SQL backend database.	Folder with readme, licensing.txt, documentation and zip folders with project files
5	Risk Assessment Module	Jupyter notebook w/ source code and UI, executables, example input and output tables structure, APIs source code for pulls/pushes from SQL, example SQL structure	Folder with installation and user guide .docx files
6	ITAMCO-SIMBA Communications Module	Contains installation instructions for an instance SIMBA Chain's Smart Contract as a Service (SCaaS) Azure Environment.	Word document
7	Performance Analytics Module	Sample dashboard and list of demonstrated and recommended KPIs. Azure SQL data structure (reference Mitigation module architecture)	.pbix and .xlsx
8	Implementation Guide - SIMBA Chain Integration	Contains developer guide for creating SIMBA Chain's Smart Contract as a Service (SCaaS) clients.	Word document, Python and .NET software examples, LICENCE.txt and README.md
9	Overall System Architecture	This shows the overall architecture of the technology deliverables	.pdf
10	Presentations	Collection of presentation and demonstrations from throughout the project duration, including the recording of the final presentation.	.pdf, .pptx, and .mp4 files

PROJECT REVIEW

This section includes use cases and a problem statement for the project, the scope and objectives, and planned benefits.

Use Cases & Problem Statement

The project aims to address the impact of external events on outbound logistics, an impact which is not currently addressed by Supply Chain Event Management (SCEM) systems. This gap includes:

- Sensing and providing advance warning of events (such as major weather or social political events) that would impact schedules or transit,
- Objectively assessing the risk associated with an event and shipment pairing,
- Proactively mitigating risk impact in an expedited and more data informed manner, and
- Recording what mitigation actions are successful for future knowledge and modelling.

The primary use case can be described from the point of view of a customer service representative (CSR) or shipment owner: As a CSR, I want to proactively address the risk associated with my shipment due to an unexpected event so that I can mitigate the cost associated with a delayed shipment and minimize the negative customer experience impact.

Currently, delayed shipments impacted by events are often reacted to after the fact or as the impact is happening. This leads to scenarios when a CSR is notified of a delay just as the customer is just finding out as well. In the case of events with longer lead times, such as a hurricane, there is some proactive work that can be in a “war room” style of group mitigation activity. What this kind of reaction typically fails to do is track and record what actions were taken and how effective they were, so some individuals come away with firsthand experience but there are no records or stored knowledge to use the next time a similar event occurs. Another use case from the point of view of a CSR: As a CSR, I want to have a recommendation for how to mitigate the risk associated with a shipment so that I can keep the shipment on schedule using proven best practices.

Scope & Objectives

The core objectives are to deliver a more systematic monitoring capability and better assessment of risk exposure, as well as a means by which mitigation decisions can be made more accurately and rapidly, so as to create a more positive customer experience and employee experience. The scope selected needs to ensure that this can be demonstrated effectively.

Dow is a very large global manufacturing company that does business in about 160 countries. Modes of Transportation (MOTs) include Air, Deep Sea, Inland Waterways, Pipeline, Postal Service, Rail, Road and Sea. Across these primary MOTs, there are more than a couple dozen various shipment types that help move hundreds of thousands of shipments. Setting realistic expectations for what could be tackled in this project, while covering sufficient complexity, impacted the scope of the work selected and performed on this project. The framework is modular and extensible so that manufacturers can integrate supply chain solutions suited to their business needs at desired scale of deployment, regardless of MOTs and shipment volume.

Necessarily we limited our project scope for the proof-of-concept shipment selection. We did, however, make a selection that would be perhaps the most impactful and likely most relatable to the majority of domestic manufacturing partners, as our intended scope was to help cover the most common and frequent shipping types. Shipment data was limited to North American (US/CAN) road and rail outbound shipments being tracked by origin and destination points. While it would have been ideal, we were not able to refine scope to only “real-time shipment visibility” data from electronic logging devices (ELDs), as this may have offered enhanced modeling data points. Two years of historical data wase used for feature selection, training and testing of the machine learning modules for transit time prediction, and the forward-looking planning window was limited to 5 days. Further, the solution was limited to internal communications only, but has full potential to incorporate supplier and customer communications through the framework.

Beyond shipment scope, there were additional scoping challenges during the project duration regarding the technology and models considered. Notably the decision to utilize native Azure connectivity to move and aggregate the module flow and key communications, although the block chain solution could also perform these tasks if working in a non-integrated environment, and indeed was developed to augment the native Azure capability, as demonstrated. Also, a scope decision was made regarding the mitigation module and the recommender logic. The decision came about as we lacked historical decision data in a format that was able to be organized to train the model. Therefore, we opted to scope out the recommender until such a point in time when the newly designed mitigation module could capture and store new mitigations decisions for this training. While this was unfortunate, there was not enough time originally scoped for the project for this additional manual data gathering and cleansing activity.

Overall the scope was still quite acceptable for the testing and validation of the event and risk modules and presented a good amount of variability so we could comfortably be assured that the solution presented would deliver value. This final project scope represented approximately a quarter of the total global shipments for Dow in 2019.

Planned Benefits

The planned benefits of the solution are increased customer experience and on-time delivery, and decreased response time, recovery time, and freight and manpower costs associated with disruptions. These benefits would be expected as soon as the solution is implemented and would improve over time as it becomes more integrated into work processes.

Benefits can be achieved by pursuing an agile and efficient supply chain infrastructure that is well monitored and maintained to meet, or exceed, customer expectations/needs during supply chain disruption events.

The benefit of such a system would allow the shift towards proactivity regarding event impact prediction, risk prioritization, mitigation planning, timely and consistent communications, with a robust system performance monitoring and value measurement capability.

Recently, Dow has conducted a comprehensive assessment of the expected value from such a solution and found that it can deliver reductions in disruption incurred freight costs and reductions in disruption related manpower costs with a total value of nearly $5MM over three years in the Dow's Outbound Logistics space for North America. Moreover, Dow's research into response time latency revealed an important connection between days in transit and cost – primarily through freight cost surges – but also in Manpower costs to manage a disruption. A faster response means fewer delays, fewer suboptimal shipments, and superior customer service during a crisis. The proposed solution will reduce freight costs based on a substantial improvement in response to disruptive events and better prioritization of shipments based on improved categorization of risks. The scalable, automated approach will also reduce the significant manpower costs associated with managing disruptions. Large scale events incur a greater manpower and shipping costs. The goal of the proposed SCEM solution is to ensure that significant disruptive events are handled as efficiently as possible.

TECHNICAL APPROACH

This section outlines the functionalities of each module (i.e., Predictive Transit, Risk Assessment, Mitigation Planning, and SIMBA Chain, Performance Analytics modules).

There is also a strong emphasis on the integration of the modules, as while they can each function independently, much of the end-to-end value comes from an integrated solution across the connected modules.

The way our project team approached this was by putting a strong focus on our Use Cases, and assuming the personas of our key users. By doing so we could better understand how what we were building would be utilized in a production environment. As an example, a CSR would not want to see various types of communications for all events when triggered, so creating a system that could consistently deliver only the required details in a timely and organized manner was key. The CSR also would like to have control over the event alerts they see, so having a flexible subscription capability was key technical objective.

Likewise, a business decision maker who has been alerted of an event that is of meaningful risk level would always be interested in reviewing decision support data, and thus would appreciate having that readily available. This happens to be a huge bottleneck as reporting access/skill can be restricted and is also a time-consuming activity. The technical approach was to pre-pull the multiple data sources periodically and aggregate within our system so it would be ready and available when needed, and then could be integrated into the mitigation module for direct review by the business decision maker, enabling much better and faster data driven decisions.

The data we included encompasses all the possible sources that are utilized today. These are resources that the business ‘wants’ to consider and utilize, but due to access and timing issues, they often would not spend the time to acquire and join/merge the support data with the impacted shipments. So in Dow’s case, the data was available, but was not utilized due to accessibility, but our technical approach for an integrated framework and data aggregation solved that problem.

Predictive Transit

The predictive transit module works on a route by route basis, training a model using historical shipment data and validating against more current data. A route is defined according to a source/destination zip codes. The route is geofenced and historical weather and events on that route are extracted. A model is trained using three groups of historical data: supplier shipments data, weather data, and social event data related to traffic. Once a model is trained for a given route, it can be used to estimate the transit time for a future shipment on the same route. The estimated transit time from the model and the baseline transit time established by the shipment planner are compared and presented to the Risk Assessment module for review and analysis. In the future the estimated transit time derived from Predictive Transit module can be used to establish realistic event-adjusted reference transit time during shipment planning.

Twitter Data	Count of tweets for each keywords: *road, event, accident, traffic*.
Weather Data	Maximum and minimum temperature, maximum and minimum rain, Maximum and minimum snow
Supplier Data	Shipment date
	Shipment type (e.g., full truck load, partial truck load, tank truck)
	Delivery priority
	Delivery item
	Carrier identification
	Dangerous good indicator (Yes/No)
	Loading point for the shipment

Table 1 Model input features.

In order to construct a model for a given route, the supplier data is complemented with two external datasets: social media and weather. The former is extracted from Twitter© and the latter from NOAA©. These two sources were selected because they provide publicly available historical data. The combination of the three datasets constitute the input features of the proposed model which are listed in Table 1.

In order to generate the data needed for each model, several steps are performed. The first step converts the source/destination zip codes into a sequence of geocodes and geofences along the route. Each geofence has a radius and a centre represented by a longitude and a latitude. This conversion is performed only once for each route by using geocoding and geofencing applications such as Geocode© and OSRM©, respectively. Tweets are then retrieved by using a date, a geofence from the route and a set of query keywords. The keywords are *road*, *event*, *accident*, and *traffic*. These query keywords were selected based on an examination of several tweets. For instance, contexts such as “road congestion” and “road construction” were aggregated under the keyword “road” because of the co-occurrence of the underlying terms. Once the above extraction is completed, duplicate tweets are deleted, and the total number of tweets that mention each keyword are summed. The frequency of occurrence of each keyword is used as an input feature to the model (Table 1). We opted for this approach in order to limit the computational complexity of the model given that a supplier may have in excess of a thousand routes where each route corresponds to a model.

In order to extract weather data, weather stations along each route are identified according to their proximity to the centers of the geofences. Daily minimum and maximum values are then extracted for temperature, rain and snow (Table 1). These values are aggregated along the route to create overall minimum and maximum daily values.

Finally, the supplier data include the shipment date. The remaining features in this group are categorical and include the shipment type (e.g., full truck load or partial truck load), the delivery priority which represents the priority level given to the shipment, the delivery item which represents the product being delivered, the carrier assigned to the shipment, a dangerous good indicator that identifies whether or not the product being delivered falls under the category of hazardous materials, and the loading point assigned to the shipment.

Using the above data, a machine learning model is developed for each route. This model consists of an ensemble of classifiers that are trained using a set of shipments and tested using a different set of shipments. Each classifier uses support vector machine (SVM) with radial basis function kernel. For the purpose of this study, the training shipments span a two year period and the validation was performed on shipments spanning a period of 6 months in 2019. This selection of training and testing shipments guarantees the separation between training and testing datasets. Moreover, it aligns with the aim of this application which is to use historical shipment data in order to estimate the transit time for future shipments.

The number of classifiers for each route is determined based on the range of the transit time. The transit time is measured in number of days where a transit time equal to 0 corresponds to the same day delivery, a transit time of 1 day corresponds to the next day delivery, etc. For each route, a valid range (i.e., minimum, maximum) for the transit time is defined from the historical supplier shipment data. For each value in the transit time range a classifier is developed. For example, if a given route has a transit time range from 0 days to 2 days, an ensemble of three different SVM classifiers are developed. The goal of the first classifier is to identify the shipments that have a transit time of 0 days. Similarly, the goal of the second and third classifiers is to identify the shipments that have a transit time of 1 and 2 days, respectively. The result of the classifiers are then combined to estimate the transit time for any shipment using the one-versus-all approach.

Risk Assessment

The Risk Assessment module considers a *subjective assessment of likelihood* (SAL) (manual decision from the SCEM team), a shipment priority score (based on business unit weights and customer variables), and transit information from the previous module to create a risk score for each shipment impacted by an event (source or destination is within the event geofence). If this risk score is above a set determined level, the shipment/event is promoted forward for review in the Mitigation Planning module.

The *shipment priority score* (SPS) is computed using critical customer variables weighted by the specific business group. Each business unit can customize the weights to best reflect their specific internal processes. The weights are internally normalized to permit user friendly units and different scales for different customer variables, while maintaining desired relative importance of the customer variables. The SPS is then combined with the SAL and nonlinearly mapped to produce the risk score, as illustrated in the figure below.

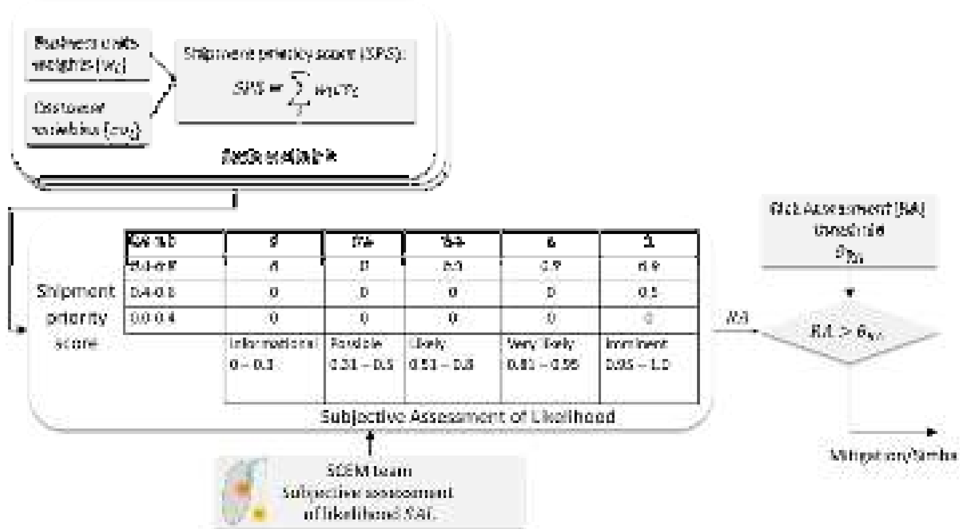

Figure .

This process is integrated with the manufacturer's database which provides data on shipments, customer variables, and business unit groups. Additionally, there is a connection to the PTM. The delayed shipment identified by the PTM are displayed per SCEM analyst's request, by overlaying markers on the map of previously identified shipments on a day.   When an event is submitted by the event manager, the system records the event data into the manufacturer's database which contains several tables, described in the stand alone document on the RA module (RA Module Final Documentation.docx). In total, the module employs eleven internal tables for storing the data and model parameters. The RA module notifies SIMBA chain, which in turn creates immutable records of events containing at-risk shipments that are above a prescribed risk threshold. The resulting records from an event creation using the RA module are currently used by the manufacturer to trigger emails about the at-risk shipments to business group managers for mitigation options.

Mitigation Planning

Making timely and informed mitigation decision around your impacted shipments is a critical opportunity for nearly all manufactures. Of course, the initial challenge is coming to know about what events are occurring (event module), and of those many events which are actually hitting your shipments in an impactful way (risk module), and only then can we know that action is required to perhaps alter the course the impacted shipment via mitigation planning.

The Mitigation Planning module will proactively notify decision makers when they have impacted shipments and collect their mitigation decision feedback, including final mitigation actions, primary decision factor that influenced the decision (used to modify future risk assessment weighting), cost to serve data (disruption incurred freight and manpower costs), and a timestamp to measure time to recover metrics. These are collected through a web form that is sent to subscribed users for a business, customer, location, route or other criterion. The webform currently provides relevant supporting customer data by querying the reporting database and will have the ability to provide a recommended action using a machine learning model trained with prior mitigation decision data.

With the project's objective towards the digitalization of the decision-making process, the proposed integrated framework will enable these actions to occur with a sequenced workflow that can operate with fewer human touches. This effectively accomplishes many of our digitalization goals, speeding up the process step handoffs, improving data capture of the process steps for time tracking, and enabling automated communications for rapid notifications.

The mitigation module is triggered when the Risk module is used and initiates action by sending an email and/or SMS to the business manager (reference MitigationImage#1).

MitigationImage#1

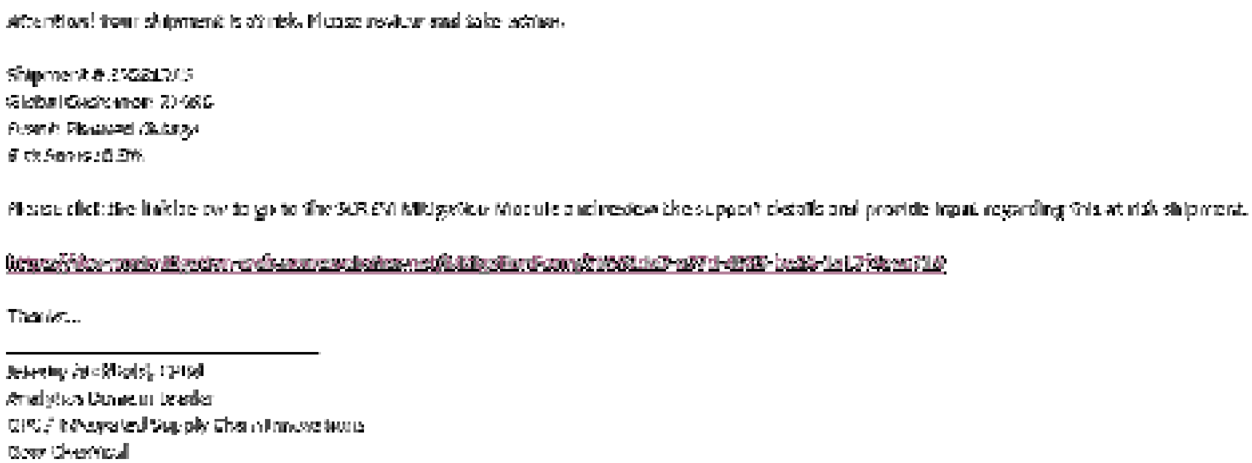

Attention! Your shipment is at risk. Please review and take action.

Shipment #: [illegible]
Global Customer: [illegible]
Event: [illegible]
Risk Score: [illegible]

Please click the link below to go to the SCREM Mitigation Module and review the support details and provide input regarding this at risk shipment.

[illegible]

Thanks...

[illegible]
Analytics Domain Leader
[illegible] Integrated Supply Chain Innovations
Dow Chemical

An entry is also written to the blockchain at this time for recordkeeping. When this email/SMS is received, it will always come in a consistent format for easy consumption. It will notify the business manager of the impacted shipment, the customer code, the event and the risk score. The email/SMS will also contain a unique hyperlink to the Mitigation module webform that enables the business manager to access a view that provides them with support decision data and further details on the customer and shipment impacted (reference MitigationImage#2).

MitigationImage#2

mxdusa.org
@mxdinnovates
info@uilabs.org

1415 N. Cherry Avenue
Chicago, IL 60642
(312) 281-6900

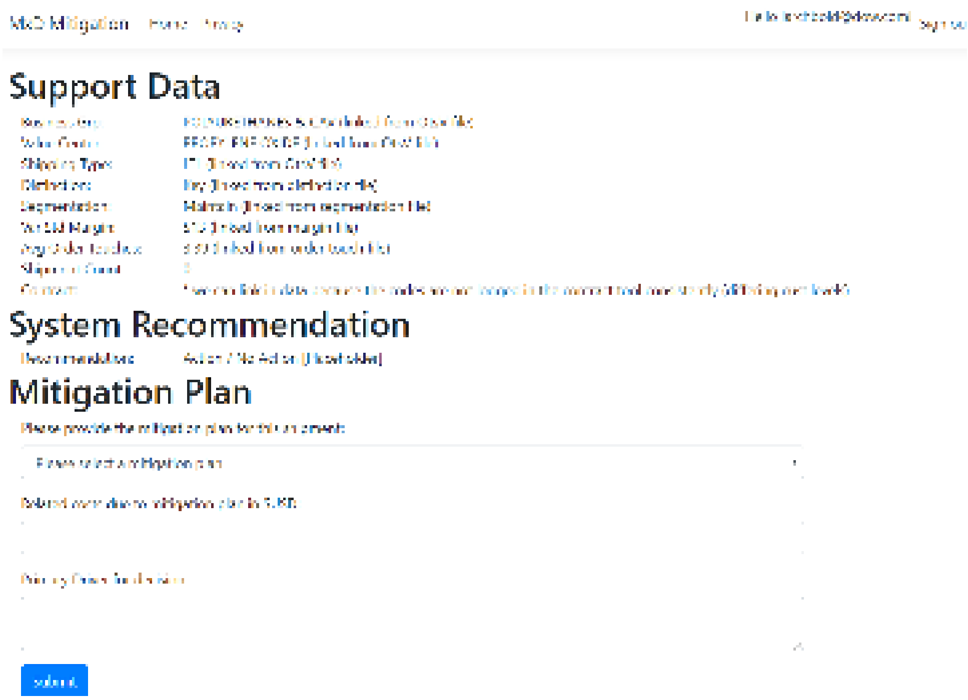

On this very simple one-page/no-scroll webform the business manager will log a few mitigation details that are critical for informing subsequent actions to be taken. Three questions are required for mitigation to be submitted: What is the mitigation plan? (reference MitigationImage#3) What is the cost associated with the mitigation action? What is the primary driver for the mitigation action? Once these are entered, the business manager may hit submit and their actions are completed for the impacted shipment/s.

MitigationImage#3

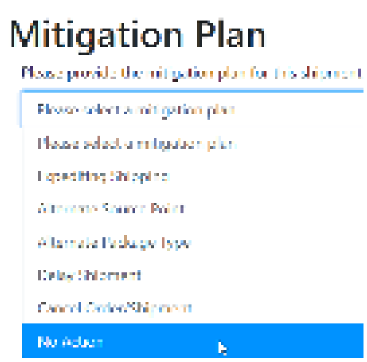

Options for the mitigation actions and support data are likely to be business specific, so the examples selected by Dow for our proof-of-concept may not apply to the options another manufacturer may opt to include. Likewise, the support data elements presented on the mitigation webform may not be the choices selected. This should attempt to replicate the current ideal state of optimal decision-making available at your business.

A key objective that we attempted to achieve but fell somewhat short of was the implementation of a fully active recommender system. This is noted on the Mitigation webform under System

Recommendation (placeholder). The challenge we encountered is that previous mitigation decisions were not recorded in such a manner that we could utilize the data for training. There was no consistency in the way Dow stored the decisions, nor were they categorized similarly. Dow also did not record the primary decision driver or the direct cost. We therefore decided that we could not implement this system, but rather *enable* the system moving forward, and as more data became available, we could turn the recommender on. This would occur in a few phases, starting with a recommendation for action or no-action. This would initially be based on risk score threshold (above the level to trigger the mitigation planning). Based on training from the business manager where they select some action or no action, the model would improve accuracy based on the available features (events, business, customer, etc.). The next phase would be where the recommender would propose a specific mitigation action plan (i.e. expedited shipping, alternate source point). This would further refine the prediction capabilities of the recommender model. The final proposed phase would be a design where the recommended mitigation action is preselected and taken, and only an action of the business manager would alter that recommendation. This would allow for the fastest decision making and could eventually bypass much of the human portion of the decision-making process once trained sufficiently.

Another important aspect of the design is that the steps have been purposefully inserted to create measurement points for cycle times. The system can measure when the mitigation module was triggered by the risk module. The system can measure when the email/SMS was sent, when the embedded link was clicked, when the mitigation plan was submitted. All these timestamps get logged as part of natural workflow and fill in a gap we have today where it is challenging to understand our current decision-making cycle time. This is not possible when making decisions via direct email or via phone calls, or whilst in war rooms… unless someone is dedicated to recording such measures.

A final point, the design intent was to keep everything very simple. Timing during these impactful events is critical, and many shipments are often impacted, so making the process as easy as possible was an objective. There was more decision support data we could have presented but we limited it to a half-dozen of the most critical factors. We also could have asked a few more mitigation decision questions, but then the process might take more than the expected 15-20 seconds and require page scrolling. These are small barriers, but they do impact adoption.

SIMBA Chain

The SIMBA Chain tracks the communication between each module and logs it in the blockchain. This allows otherwise non-integrated solutions to be able to both talk to each other and gives the framework a single repository of non-refutable record of the data flow. Rapid delivery of secure alerts are delivered using Blockchain secure communication and Blockchain API endpoints.

In developing the communications solution the team started with a service where members can subscribe to certain features stored on the blockchain, such as shipment delivery zipcode or customer number. This subscription service is setup through SIMBA Chain's user interface. The future vision is for the solution to then migrate to use online learning techniques to tailor a customized filter for each member based on historical data and user profile. However this was not yet implemented. Aligned with the differential workflows, alerts will also be tailored around the categorization of the risk with automated alerts for low risk events. For prioritized, high risk events, where predictive modeling has been enacted, the system allows for customized communications in support of collaborative working arrangements required to deal with predicted impacts. Moreover, the blockchain communications solution incorporates security features utilizing Blockchain technology in order to facilitate information sharing with authorized parties while adhering to compliance governing the sharing of sensitive regulated data and/or restricting access to competitive intelligence resulting from system identification of potential adverse Supply Chain events and their corresponding impact on broader business operations. Authorization is done through Microsoft Azure.

Performance Analytics

The Performance Analytics module is comprised of a data aggregation table in the primary database, where all the relevant data artifacts that support the KPIs are stored, and a visualization dashboard capability. Together, this technical solution enables near-real-time visibility into the system performance for the defined KPIs and measures.

The table that supplies the data pipeline to the dashboard is housed in the Azure environment, within an SQL database. This is the location that aggregates the data that supports the KPIs, and all other modules (PTM, Risk & Mitigation) and connects to the SIMBA blockchain communication module as well.

The dashboard we chose was built using Microsoft PowerBI, but just as well could have been built with Tableau, Excel or an existing internal reporting system which can accept database connections. The presentation was less about the actual visualization, and more about the connected capability.

Select categories of Key Performance Indicators and measures were selected for tracking and testing. They include: cycle-time, system integration, transit time estimate, event extraction, risk assessment, automatic alerts and general performance management.

Response Cycle-time: In this category, focus was put on creating timestamps around key activities and handoffs so that cycle times could be generated. For the proof-of-concept we set to focus on a couple key cycle-time (total cycle time & email response time), but the implementation should call out the specific cycle-times of interest for your implementation.

System Integration: In this category, the modules were reviewed for connectivity. As part of the technical design, we required that there be minimal human intervention between the modules, although necessarily within the modules some human input would be required. This was accomplished by having the integration build within the Azure environment and by having SIMBA blockchain also log the receipt of the completion of each module's handoff.

Estimated Transit Time: In this category, we reviewed the estimated transit times for given specific routes, and if the system was able to arrive at a modeled proposal, with respect to inputs such as location, weather, traffic and events.

Events Extracted Automatically: In this category, we reviewed the expected percentage of the events that could automatically be collected for the given routes modelled.

Risk Assessed Systematically: In this category, we looked at the impacted shipment and events combinations to ensure all could have a risk score assessed automatically.

Targeted Automatic Alerts: In this category, we were concerned with alerts (emails/SMS) getting to the proper users. As part of our testing, we setup subscribed users in various

mxdusa.org
@mxdinnovates
info@uilabs.org

1415 N. Cherry Avenue
Chicago, IL 60642
(312) 281-6900

subscriptions levels to test this capability, in both the Azure framework (using email and Twilio/SMS) and from SIMBA blockchain (email).

Performance Management System: And finally, in this category we assessed the capability of the reporting and measures, and KPIs for the entire integrated framework.

All these categories, along with bare measures/counts of events, cases, shipments are part of the total Performance Analytics module.

See diagram below for visualization of the capabilities and their relation to the specific value drivers they support.

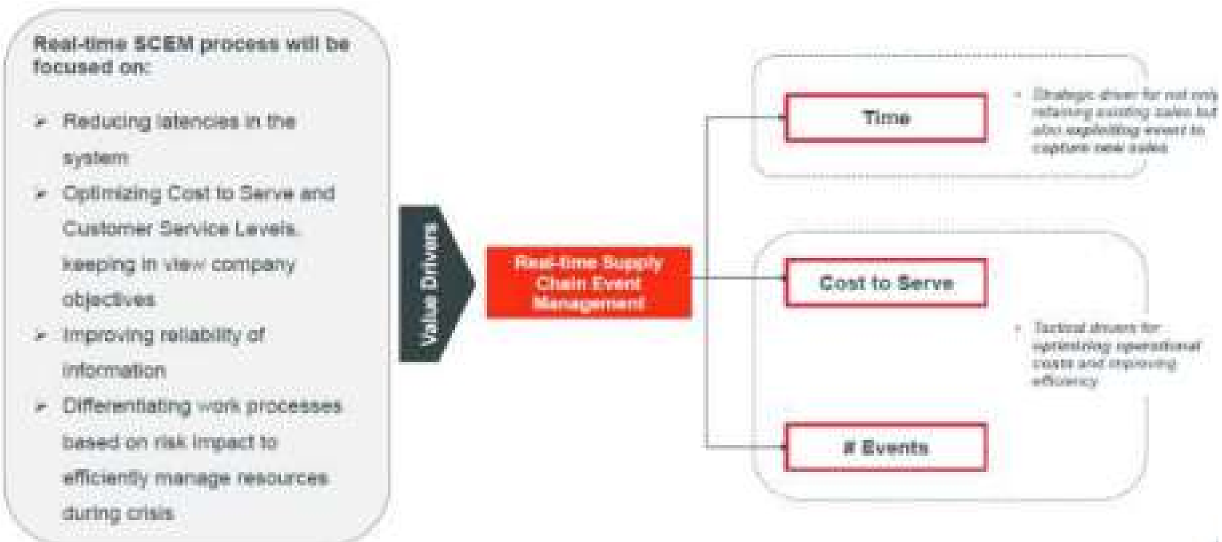

RESULTS

This section includes a list of technology deliverables, a system overview, the system requirements, the system architecture, the features and attributes of the technology deliverables, the target users and modes of operation, and software development and design documentation.

The table below will summarize these top-line deliverables, but we will cover each in more depth, including impact & output of each deliverable.

Table 2: Technology Deliverables

#	DELIVERABLE NAME	DESCRIPTION	FORMAT OF DELIVERY
1	Overall System Architecture	This shows the overall architecture of the software components that were developed and integrated through the framework for the proof-of-concept.	.pdf
2	Mitigation Planning Module	Source code for Web UI, APIs source code for pulls/pushes from SQL, example SQL structure	Folder with readme, licensing.txt,

			documentation and zip folders with project files
3	Predictive Transit Module	Admin UI, User UI, ML algorithms source code, Executables, example input and output tables, APIs source code for an SQL backend database.	Folder with readme, licensing.txt, documentation and zip folders with project files
4	Risk Assessment Module	Jupyter notebook w/ source code and UI, executables, example input and output tables structure, APIs source code for pulls/pushes from SQL, example SQL structure	Documentation
5	SIMBA Communications Modules	Smart Contracts for blockchain deployment, Source code for subscribe, APIs source code between PTM, RA, and MP modules	Folder with readme, licensing.txt, documentation and folders with project code
6	Performance Analytics Module	Sample dashboard and list of demonstrated and recommended KPIs	.pbix and .xlsx

System Overview

There were multiple gaps in the existing process for managing Supply Chain Events, so many areas needed covered. We could see upfront that this may not be the case for other manufacturers, who many have more mature systems or reliable partners handling some of these gap-areas for them. This drove us to develop a modular designed framework that would enable manufacturers to utilize individual modules as needed to close on their specific gaps. And furthermore, for those modules, we would leave the options open for customization to fit the business use cases as per the specific customer and markets they did business in.

The project developed a vendor-agnostic integration framework for connecting supply chain solutions in order to present supply chain managers with the information they need to assess and mitigate risks systematically and efficiently. The framework is modular and extensible so that manufacturers can integrate supply chain solutions suited to their business needs at the desired scale of deployment. The Blockchain technology enables the framework to keep an immutable record of all supply chain management events and serves as a centralized authority.

This project implemented this integration framework through a proof-of-concept for supply chain management of Dow's outbound logistics. Through a standards-based framework, the project applied machine learning models for predicting shipment transit time, developed software for systematically assessing risks with a human-in-the-loop, and developed software that allows supply chain managers to view all the data needed to make mitigation decisions and collect data about the mitigation action.

The predictive transit model uses machine learning algorithms that leverage data from weather, and social media to predict shipment transit time in order to give advance warning for planning. A risk assessment software module consists of backend algorithms for systematic risk assessment configurable to business metrics and a user interface for mapping those risks across shipments. A mitigation planning software module delivers automated email notifications to supply chain event managers to alert them of potential risks. This module directs users to a web-based interface to view all the data for high-risk shipments and capture data about mitigation actions taken. Once enough data is collected, the actual mitigation data will be used to train machine learning models that automatically recommend mitigation actions to the supply chain event manager.

Through the framework, the analytics results and support data are integrated and delivered automatically in targeted, tailored notifications to end users. All software components and data sources are connected leveraging an immutable blockchain ledger deployed with Smart Contract technology that easily scales to connect more solutions. Business managers can quickly visualize supply chain management performance analytics through a dashboard that aggregates information from all modules.

The system being delivered is a series of modules that, when interconnected, represent a complete framework for Supply Chain Event Management, designed to assign a risk score to impacted shipments that may be adversely affected by an event, notifies relevant subscribed parties, record mitigation action taken as knowledge, and tracks and visualizes KPIs. The framework has the ability to provide a recommended mitigation action based on past actions, but this functionality was not demonstrated, as the capability was not brought to a production environment since historical training data was not available. The functionalities of each of the modules is described in each respective document as summarized below.

The PTM module is documented in three manuals. The User Guide - PTMAUser.docx user manual describes the use of the module by shipment planners to estimate the transit time for future shipments on a given route.

User Guide - PTMAUser

The detail documentation of the Risk Assessment module is provided in User Guide - Risk Assessment

The document identifies three different users, viz. an SCEM analyst, a system administrator, a business unit manager, and a software development and provides specific details tailored to the three types of users in addition to the common overview of the module. The principle of operation and the underlying data structure that consists of eleven SQL tables, for storing both data and parameters, are described in detail. Finally, the document explains the software implementation, organized in three classes: the first for communication, the second for graphical user interface and risk computation, and the third for managing geofences.

A guide for implementing the SIMBA Chain blockchain communications module with sample code can be found here: User Guide - SIMBA Chain Communications Module

The Mitigation module is designed to accept a shipping risk event and then to notify the pertinent party of the event via email and provide information so they can make an accurate decision on a mitigation of the risk. The primary use case is risk of a shipment being delayed. The shipment risk is determined from sophisticated machine learning models that consider the shipping route, method and current events in the area to determine a risk score. If the risk score is above a certain threshold a shipment risk event is created and dispatched to the mitigation module. The mitigation module records the event and sends an email to notify the appropriate party.

User Guide - Mitigation Planning Module

System Requirements

In order to implement the framework being described in this report, there is fair amount of system requirements that should be recognized and arranged. While we will share many specific examples of systems and applications we utilized, these are by no way the only means by which such a framework could be implemented. For example, where as we opted to develop within a cloud environment, the very same solution could be stood up as an on-premise solution, which offers some pros and cons itself, but perhaps most notably could address network and security concerns. Likewise, instead of partnering with Microsoft on their Azure cloud environment, partnership with Amazon Web Services (AWS) or Google Cloud Platform (GCP) could have delivered a similar outcome. This continues for the other elements of the modules, such as our blockchain partnership with SIMBA, or for the development platform such as Jupyter notebooks vs Azure notebooks. Please use our framework as a guide, and not as a prescriptive design you must follow exactly.

For systems that would be ideal to stand up a similar framework to that which we enabled in our proof-of-concept, please review the listing below:

- SIMBA Chain Smart Contract as a Service (SCaaS) Deployed to an Azure Subscription
- Azure environment services:
    - Managed Postgres
    - Key Vault
    - Service Bus
    - Blob Store
    - Active Directory
    - Azure Managed Blockchain (Quorum)
    - SQL Database
- HTTP/REST Client side components to communicate with SCaaS

The files linked below explain the necessary requirements for usage of each module:

- Predictive Transit Module - The installation manual describes the configuration and deployment steps for the PTM module. The main requirements of the PTM module are an SQL database, Python and .net. Installation and Implementation Guide - PTM
- The installation manual for the risk assessment module is provided in "Installation and Implementation Guide - Risk Assessment". The document identifies non-standard libraries, not included in a typical Anaconda installation, and provides detailed installation instructions.
- User Guide - Mitigation Planning Module - This installation manual covers the setup requirements for the environment, the Visual Studio Workloads requirements, list of repositories, and details instructions to enable the Mitigation Module (database, code and webpage) - Installation and Implementation Guide - Mitigation Planning Module
- Communication Module - This installation manual covers the setup of SIMBA blockchain Implementation Guide – SIMBA Chain Integration.docx

System Architecture

The project developed a system architecture that meets the requirements of the project use case in the most efficient way while leveraging a modular design. This is presented below in the form of two charts which represent that very same architecture, only in varying degrees of detail.

While these diagrams do a good job of showing the systematic exchanges between the technical deliverables, the humans involved are shown in the diagrams as the personas noted: Supply Chain Event Management User, Business Unit User, Admin User, Planning User, Customer Service Rep User, and Business & Functional Partner Users. These are the common and universal players in a typical supply chain event management system for nearly any manufacturer.

For the design to function appropriately, engagement from these users would be required, and thus change management and training should scope in these users specifically, as well as any support teams that impact these user groups.

On the front end of the architecture we have setup engagement required from the SCEM user, BU user, Admin user and the Planning user. These are required before the system can work, as their engagement in the process is to setup many of the triggers and values that run the system. This includes defining route setups, subscribing users to communications, assessing risk thresholds, and determining the mitigation decision support data needed. The admin would also be making the necessary system changes behind the user interfaces designed (including ETL and Azure Data Factory work to migrate data into the environment from your ERP or reporting system).

In the middle we have continued engagement from these personas when it comes time to make decision support (business users and customer service reps) and from the SCEM user who has the option to manually create and trigger events that were not captured with the PTM module, such as internal events such as a weigh scale down which might delay all shipments from a site.

On the backend of the architecture a few personas are called out, but in reality, any partner who is interested in the system performance could gain access to the secured visualization capability with the Performance Analytics module. In the demo case, access is controlled via a MS PowerBI dashboard, where there is the option to generate multiple customer views as per users' requirements. This was explored, as the interests of the SCEM users and the business users may be somewhat different, and the ability to connect to a common data source via our developed pipeline that enables this customization was considered a valuable capability.

The architecture of the PTM module is a three-tier architecture. The back-end is a SQL database that implements several tables needed for the training and execution of the predictive transit module. This database has two main data tables: one for historical data and one for new shipments. The former is used to develop the predictive route for each route (source/destination). The latter, includes a list of future shipments that are processed by the module. In addition to these two tables, the backend maintains a set of administrative tables for categorical value conversion, static route information, and for standard metrics related to the predictive transit models. The middle tier consists of the machine learning engine that derives predictive models for each routes and the data extraction APIs. The data extraction APIs are source dependent. In fact that are three APIs one for supplier data, one for weather data and one for Twitter data. The first tier is a combination of two user interfaces. The first UI is an administrative UI that allows the application owner to develop or update the predictive models for each route. The second UI allows a general user with no administrative privileges to enter data related to a future shipment and estimate the transit time. All new shipments are maintained in a table within the PTM module. This table is accessed by the Risk Assessment module for further processing of shipments that are at risk for delay. In addition, once the transit time for a future shipment is derived by the PTM module, this estimate is compared to the transit time defined by the shipment planner. If the estimated transit time is higher than the anticipated transit time, the shipment information is posted to the Blockchain using an HTTP POST that includes the shipment information as well as the values of all the input features used to derive the estimate (Table 1.)

The RA module uses SQL to interface with the PTM module and the Mitigation module: it reads the data from the “Shipment” table provided by the manufacturer to identify shipments and reads the “newshipents” table from the PTM module to find the identified delays. The shipments identified by the PTM module are indicated on the map, at SCEM analyst’s request. The RA module stores the calculated risk results in the Event_shipments table and subscribes to the mitigation module’s rest API to trigger the module if the event contains a shipment with risk above the prescribed threshold.   After the event is created and risk scores computed for all shipments within the event, data is posted to SIMBA for the shipments whose risk is above a prescribed threshold (initially set at 0.5). Data is posted to SIMBA using the two fields *Event_ID* and *Update_No*, along with a (.csv) file containing individual shipments information (shown in Table 10). The shipments information includes the fields *Shipment_No*, *Business_Group_Code*, *Mode*, *Update_No*, and *Risk_Score.*

SIMBA Chain provides an immutable, non-refutable ledger of the outputs of the various modules in the system.

The Risk Assessment Module, Predictive Transit Model and Mitigation Planning Module (Recommender and Connectors) communicate with the SIMBA component via the exposed HTTP API, in order to push their results onto the blockchain. SIMBA receives the HTTP POST requests and populates the blockchain.

SIMBA also supports subscriptions and notifications for data pushed to the blockchain. This is used to send email notifications to interested parties based on the results and particular attributes (e.g., shipment number) of the messages from the Mitigation Planning Module.

The file linked below represents the visuals that describe the total system architecture used in this solution. Both the detailed and simplified version can also be found in the MxD Project Final Presentation_Final.pptx file.

System Architecture Diagram

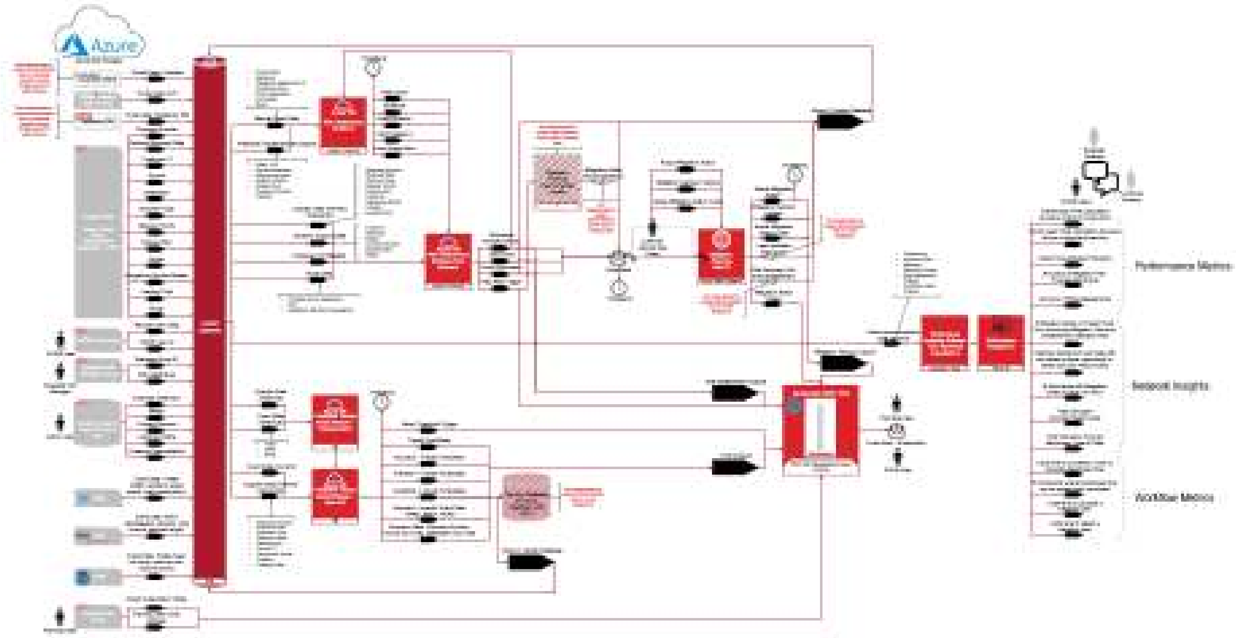

Detailed System Architecture

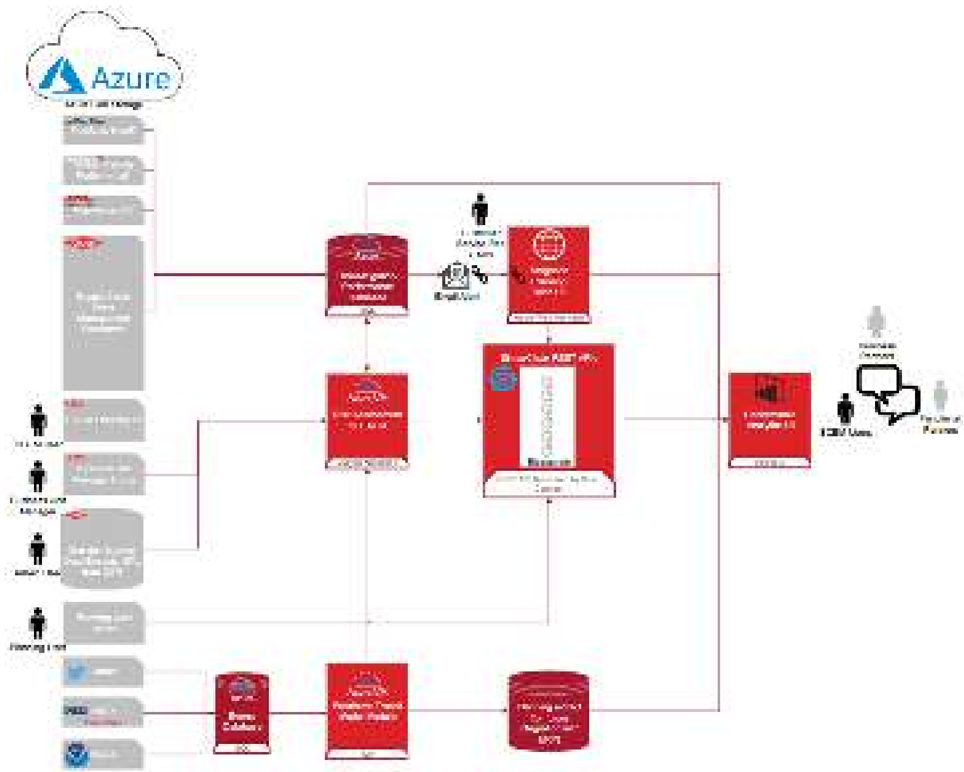

High-Level System Architecture

Features & Attributes

The main components of the PTM middle tier are as follows:

- Route configuration: this component converts a source/destination zip codes into a route of geofences. This configuration is performed once for each new route.
- Tweet extractor: For each geofence in the route, the tweet extractor returns the tweets that contain set of query keywords (i.e., *road*, *event*, *accident*, and *traffic)* on a given date. The radius of each geofence is five miles.
- Historical weather extractor: For each shipment route, a list of weather stations are identified along the route. Historical weather data associated with shipment date is then extracted from the NOAA year summary files. The parsed data is aggregated into maximum and minimum temperature, rain, and snow.
- Forecast weather extractor: The forecast weather extractor is similar to the historical weather extractor except it extracts the data from a different source (weather.gov) using the forecasting API.
- Transit time predictor: The transit time predictor use the SVM application to build multiple classifiers. This component is invoked by the PTM Admin UI to develop predictive model for each route. These modules are then used by the PTM user UI to estimate the transit time for a future shipment.

The three key components of the RA module are organized into three classes:

- Class 1 – dbToolsMxD is used to connect and communicate with the manufacturer's database. The class contains an initialization method for connection, and a simple query method which runs a query and returns all results as an array. Secure queries using the execute method directly from the database connection's cursor object were employed.
- Class 2 – Event_Creation_Tool is the main class of the module. It contains all the methods for interacting with the GUI, retrieving and displaying information requested by the analyst, setting off the risk score computation, saving the data to the database, and submitting events to the database, Simba Chain, and mitigation processes. To facilitate the future development and use of the code, each method has a "Notes" section in the doc-string which describe the step-by-step actions taken by each method. Additionally, these steps are listed before each section of code in the methods.
- Class 3 – Geofence serves to generate an elliptical geofence in the form required for ipyleaflet to display. The ellipse is parametrized with its size (two semi-axes), tilt and location. These are set by the user when interacting directly with the geofence on the map. The parameters are stored in the database, providing the capability to restore and modify the geofence, as well as to capture of the entire history of committed changes.

At a high level, SIMBA Chain provides an immutable, non-refutable blockchain ledger of the outputs of the various modules. This provides the mechanism for review and audit of all decisions made by the system. This is crucial to providing high levels of trust in, as well as adjustment of the automated decision process. Additionally, SIMBA Chain provides subscription and notification of blockchain events. These have been used for notifying interested parties of decisions.

SIMBA Chain Smart Contract as a Service (SCaaS) is used by the project. This is a SaaS offering that is deployed to Microsoft's Azure environment. It provides a tailored interface to the blockchain, defined by the smart contract used by the project modules. SCaaS provides two main services related to a smart contract:

1. An HTTP/REST API that models the methods and arguments in the smart contract. A simple POST to the API will result in a transaction on the Azure Blockchain. Supported MIME types are multipart/form-data and application/json. Details of the HTTP API are provided in the Implementation Guide - SIMBA Chain Integration deliverable.

2. The Explorer Interface to query for transactions on the block chain and set up subscriptions to particular smart contract methods and attributes. Details of the Explorer User Interface are provided in the SIMBA Chain Communications Modules deliverable.

SCaaS leverages Azure Active Directory (AD) for authentication and authorization. This provides a single point of access control that can integrate easily with other existing deployed services, as well as providing a client with full control over access to SCaaS services. Additionally, SCaaS integrates with Azure Key Vault which means that a client does not have to worry about creating or maintaining blockchain addresses or associated private keys. Instead, SCaaS creates a new blockchain address and private signing key for each user that accesses the HTTP API and reuses those details for each request they make. The secret (blockchain address and private key) are mapped to the user's AD Object ID in Key Vault. This also means that clients can easily track who has created which transactions on the blockchain by matching up the user's object ID in AD to the key to the secret in Key Vault.

SCaaS uses Azure's Service Bus Queue in order to create a robust mechanism for pushing data to the blockchain. Messages are pulled from the queue before pushing to the blockchain.

For off-chain data, SCaaS leverages Azure's blob store. This allows clients to store files associated with transactions outside of the blockchain itself and referenced by a hash value on the blockchain, providing a pointer to the data and ensuring data integrity via the hash value. When files are received via the HTTP API, they are zipped up and stored in the blob store. Along with the files, the zip file contains a manifest giving the hash of each individual file, the hash algorithm used (currently SHA256), the file's mime type and size. The filename of the zip itself is the hash of the zip file. The manifest also describes the algorithm for this hashing (currently also SHA256), allowing you to validate the zip and its contents for integrity.

SCaaS uses an Azure managed instance of Postgres to store transaction data, allowing fast query on the contents of the blockchain. The database is populated by a chain watcher component which checks for blocks being created on the blockchain and captures transactions on the smart contract that have been verified by the blockchain.

Finally, SCaaS leverages Azure's blockchain service as the underlying blockchain implementation. This is a Quorum blockchain which does not require tokens for executing transactions.

The documentation below outlines the features that each module enables.

<u>User Guide- Mitigation Planning Module </u>- MxD mitigation module is designed to accept a shipping risk event and then to notify the pertinent party of the event via email and provide information so they can make an accurate decision on a mitigation of the risk. The primary use case is risk of a shipment being delayed. The shipment risk is determined from sophisticated machine learning models that consider the shipping route, method and current events in the area to determine a risk score. If the risk score is above a certain threshold a shipment risk

event is created and dispatched to the mitigation module. The mitigation module records the event and sends an email to notify the appropriate party.

PADashboard.pbix - The Performance Analytics dashboard was developed in Microsoft PowerBI. This could have developed in Tableau or other dashboard tool, as there are no special capability requirements here with PowerBI specifically. The core deliverable as part of the technology is that a SQL database was developed with the Azure framework which has an aggregated Performance Analytics table which has all the live-connected links to enable the connection out to PowerBI (or your dashboard tool of choice).

User Guide - PTMAdminUser - The administrative user manual describes the procedures needed for training and validating the predictive models for each route.

User Guide - Risk Assessment Module – The user documentation for risk assessment tools including in-depth description of each feature.

User Guide - SIMBA Chain Communications Modules.docx – The document provides detailed documentation on the SIMBA Chain module configuration.

Target Users & Modes of Operation

There are five key target user or user groups for this solution. These were called out earlier in the architecture design section, but we did not get into the reason why these were target users, and perhaps their interest areas relative to such a solution as being presented here.

These five users are:

Business unit manager determines business metrics for unit/division which influences risk thresholds and feature weightings. Their interest in this solution would be that their goals are typically aligned to business sales performance, and they are motivated to enable to smoothest flow of order to shipment to invoice to cash. Any barrier to this conversion of sale to cash would be a priority to the business user. A risk encountered during events is the loss of orders, or the loss of customer confidence due to lack of performance under event duress. This system seeks to enable the business to remove such barriers.

Supply chain event manager has expert knowledge of event detection, risk, and mitigation strategies. Event managers collect event data through a variety of sources to determine risk and provide details to customer service representatives. Their interest is to most efficiently manage supply chain events. This includes ensuring that nearly 100% of events are identified and that the time to resolve is minimized. The system helps them here by automating event detection via the PTM module. It also simplifies the manual event creation process through the user interface within the Risk module. The time to resolve is minimized by having a system that supports targeted and automated communications. The Performance Analytics module ensure that specific performance can be reviewed in near-real-time so adjustments can be made as needed.

Customer service representatives receive data about supply chain risks and determine the appropriate mitigation plan, then act on the plan and record the resulting cost and decision features. Their interest is in providing the customer the very best service, which in this scenario would be primarily information sharing. Today the CSR has limited insight to delays that are impacting their specific shipments, and it is not unheard of that a CSR is notified of an impacted shipment via the customer themselves. The system proposed allows CSRs to subscribe so that their shipments are flagged so they can come to proactively know of impactful events prior to the customer becoming aware. This enables them to take charge of the information flow with the customer, which creates additional options when it comes to mitigation planning.

Supply chain planners use predicted transit times prior to a shipment date to inform logistics planning decisions. Their interest is a mix between the business unit manager and the CSR, as they are trying to optimize to meet customer orders, while minimizing inventory and related business costs (warehouse costs, logistic movements, extra handling, etc). This is done in service to both the customer (meeting their demand requirements) and the business (optimizing the sales & profits). The system when used properly would tend to find an equilibrium point that can be more easily measured using the Performance Analytics module, versus the path that if often followed today which is to do anything possible regardless of costs or efficiency for all customers, across all profitability levels.

System administrator maintains regular updates to the backend such as blob storage updates, predictive model training data, risk score weightings, and notification subscription authorization. Their interest is in operating a reliable, efficient and low-cost system that enables the business to make profit. They are a support team in the system being proposed, and engagement would not be necessary on a regular basis. They would be engaged to make code updates to the modules where adjustments are desired, to update ETL processes where new support data is needed, or to manage the system performance (increasing processing and disk speed in cloud to enable faster system performance). There is an equilibrium point for cost versus performance, and the admin would help find that balance.

Software Development/Design Documentation

The pre-requisite and deployment steps for the PTM module are included in the Installation and Implementation Guide - PTM manual. The software delivery also includes sample synthetic data that can be used to populate the database and perform initial tests and verifications. The steps outlined in User Guide - PTMAdminUser allows an application owner within the organization to configure PTM for specific shipment routes used by the manufacturer. The manual also provides procedures for developing and maintaining the predictive models. The model uses Support Vector Machine as the machine learning technique.

Details of configuring a SCaaS instance are described in the deliverable SIMBA Chain Communications Modules.

Details of the APIs and how to create SCaaS client software are described in the deliverable Implementation Guide - SIMBA Chain Integration. The software aspect of this deliverable provides example code for Python and .NET clients. The code shows how to POST transactions to the service and how to GET all transactions for a given smart contract method.

POST examples include sending data using the `application/json` miime type and sending data plus upload files using `multpart/form-data` mime type.

The interface uses simple HTTP POST and GET and is therefore easy to integrate into an existing software stack. The Azure specific code is around authentication using OAuth 2 implicit flow.

There are two approaches to defining identities for OAuth:

1. The `single-identity` directory contains examples for sending and getting data from SCaaS when you are happy for all connecting client to share the same blockchain identity (aka address). This is the case when all clients are trusted and write to their own contract methods so you don't have to distinguish between identities writing to the same method.

2. The `multiple-identity` directory contains examples for sending and getting data from SCaaS when you want to define separate identities for connecting applications.

DISCUSSION & ANALYSIS

This section includes discussion on the industry impact, key performance indicators, accessing the technology, workforce development, and lessons learned.

Industry Impact

Every industry encounters events that can impact their supply chain. Depending upon the breadth of the supply chain, this impact can be greatly magnified. For a large global company like Dow, the opportunity was enormous to explore how such a system would enable improved performance over the existing non-integrated framework event management system.

As it is delivered, this proposed framework solution would offer benefits across almost any large supply chain in North America, primarily large manufacturers. The solution provides the ability to integrate each segment of the event management process and improve the effectiveness of that process over time. This is achieved by capturing mitigation decisions and training a model to provide suggestions to mitigation opportunities in the future. This methodology could also be leveraged outside of industries with a large volume of outbound shipments by shifting the perspective of data collection from mitigation activities to other trackable decisions being made by informed employees that pull heavily from experience and background knowledge.

In a digitally-minded market, the ability to quickly react to risks is imperative to remaining competitive. There are no current solutions on the market that encompass the capabilities to predict, assess, and mitigate risks through one platform without sharing sensitive supply chain data with the vendor. This solution developed a flexible system that allows manufacturers to improve existing supply chain management processes while choosing solutions that meet their needs, without compromising data security or level of system plug-and-play.

This human-in-the-loop solution will impact supply chain management business processes by taking a skilled workforce away from tedious, repetitive tasks and instead empowering them to make more informed decisions using their skills. End user feedback on the developed system indicates significant improvement in employee experience once adoption and training is complete.

Additionally, this project is one of few successful real-world deployments of blockchain in supply chain. With a growing demand for improved communication within manufacturers' supply chains, blockchain technology has high potential for encouraging openness while still protecting data fidelity. Recording transactions on an immutable ledger has many potential benefits including freight dispute settlement, where there is an immutable record of transit times. This project serves as a tangible example of implementing this cutting-edge technology and will pave the path for future deployments across the industry.

Key Performance Indicators & Metrics

The table below summarizes the KPIs for this project, and how the goal from the project outset compares to the results at the project conclusion.

Table 3: KPI's and Metrics

METRIC	BASELINE	GOAL	RESULTS
Events Extracted Automatically — Automated data collection of open data sources % Collected	Baseline Process: Significant manual extraction or searching **None (0%)**	# Data sources collected automatically, applying APIs, for active monitoring **Most (>80%)**	The team built a human-in-the-loop event detection system that accepts both automated and manual event detection. **>80%**
Risks Assessed Systematically — Working systemic risk categorization in place % Systemic (initial categorization)	Baseline Process: Subjective risk categorization only **None (0%)**	Systemic risk categorization module operational for identified disturbances **Most (>80%)**	All risks assessed using combination of Subjective Assessment of Likelihood and objective Risk Score **All (100%)**
Targeted Automatic Alerts — Targeted, tailored and secure alerts system that can also provide cueing for high impact events	Baseline Process: Communications system not targeted or secure **None (0%)**	# Targeted, tailored and secure automated alerts system **Most (> 90%)**	Email & text alerts delivered to subscribers through blockchain or through Azure **All (100%)**
Performance Management System in Place — Performance Monitoring of Key SCEM Operational and Financial Metrics	Baseline Process: Limited Management Information on SCEM **None (0%)**	# Operational / Financial / Other Metrics, Events information viewable **All (>85% Viewable, remainder derivable)**	Integrated performance analytics dashboard draws from central database, customizable to metrics or platform of interest **All (100%)**
Response Cycle Time — Duration from detection of an event to notification of risk	Baseline Process: Response time for 'war room' manual event	Response time from event detection to risk assessment notification and mitigation actions	Instant notifications **Minutes to view supporting data, assess mitigation recommendation, and enter mitigation decision data all in web UI**

and mitigation option to a supply chain manager	detection, risk assessment and mitigation actions **>2-3 days**	**Event Detection to Alert: Scale of Minutes** **Alert to Mitigation Action: Scale of Hours**	*Mitigation plan recommender not tested
System Integration — Software modules are connected and require no manual intervention by user to “start” or to transfer data between modules	Baseline Process: Significant manual transfer of data between systems **None (0%)**	Level of integration between supply chain components **Fully Integrated (100% modules integrated)**	Integration via Azure and blockchain **Fully Integrated (100% modules integrated)**
Estimated Transit Time Transit times are estimated for a given route based on specific weather and events for a given shipment date.	Baseline Process: Transit time for each route is fixed and do not vary according, date, weather conditions, or traffic events.	Accurate transit times for specific shipments (Low mean average error)	**Module was tested on seven routes and in excess of 2000 shipments. The mean Average error (MAE) was less than 1 day for all routes.**

Referring to the diagram above, the KPIs noted were previously discussed categorically, and will be again summarized below, but there is still need to share project results and discuss the reasoning for selection and challenges encountered.

The seven categories of KPIs selected by the project team that were intended to cover all the desired performance measures for the performance of the system. They include: cycle-time, system integration, transit time estimate, event extraction, risk assessment, automatic alerts and general performance management.

The team did encounter first challenges in selecting these specific measures in that many required counting (as a baseline) that which were previously unable to be accounted for. This becomes a challenge when discussing percent improvement over baseline but will tend to make more sense when measured against itself while under future operational usage. Note those with a baseline of none/0%.

Summary and *breakdown discussion*:

Response Cycle-time: In this category, focus was put on creating timestamps around key activities and handoffs so that cycle times could be generated. For the proof-of-concept we set to focus on a couple key cycle-time (total cycle time & email response time), but the implementation should call out the specific cycle-times of interest for your implementation.

With the aggregation of the modules time/date triggers into a single database table, we enabled this capability. The team initially noted a dozen different cycle-times of interest for the various users, but limited calculation and presentation due to timing. The reasoning was that the capability was demonstrated and that further examples would be effectively duplicative for demonstration purposes. The Dow usage of the system was a proof of concept, using real data, but not within our live production system. That pilot usage would come later (see adoption section). No concerns over reported numbers.

System Integration: In this category, the modules were reviewed for connectivity. As part of the technical design, we required that there be minimal human intervention between the modules, although necessarily within the modules some human input would be required. This was accomplished by having the integration build within the Azure environment and by having SIMBA blockchain also log the receipt of the completion of each module's handoff.

Despite developing these modules individually, the team was able to accomplish this objective fully. In fact, we had two paths to completion here with both the aggregation and containment of the solution within a single connected cloud environment (Azure for our example) and with the usage of BlockChain to track the movement between modules. The team was able to design this upfront by setting up the single database for modules to deposit their outputs and with the BlockChain design to allow for very simple calls from each module. No concerns over reported numbers.

Estimated Transit Time: In this category, we reviewed the estimated transit times for given specific routes, and is the system was able to arrive at a modeled proposal, with respect to inputs such as location, weather, traffic and events.

For that which was in-scope, the solution met the mark on this measure. However due to the runtime effort to extract historical performance to compare against, the scope was limited to a much smaller sample of higher volume shipping routes. The team would have liked to run even more, maybe even up to the entire scope of the project (NAA road and rail) but this was not easily possible with the scale designed. Of course, with more time this could be easily accomplished with a scaled implementation of the solution. This initial historical load, after all, is a onetime task for routes, and moving forward repeat daily usage would be additive. No concerns over reported numbers.

Events Extracted Automatically: In this category, we reviewed the expected percentage of the events that could automatically be collected for the given routes modelled.

The PTM module was trained for seven routes using shipments data from 2017 and 2018. It was then tested on shipments from 2019 and the performance of the model was measured using the mean absolute error (MAE) defined as:

$$MAE = \frac{1}{n}\sum_{1}^{n}|predicted\ transit\ time - actual\ transit\ time| \qquad (1)$$

where n is the number of shipments in the test dataset.

The seven representative routes varied in distance, source, destination, heading, and transit time as shown in Table 2. Routes A and B start from the same source. As expected, long routes have a high average transit time whereas shorter routes have a low average transit time. Moreover, the transit times for certain routes can vary considerably. For example, one of the routes has an average transit time of 2.5 days with a standard deviation of 7.3 days.

The MAE value across all shipments in the test dataset was less than the standard deviation of the transit times for all routes. These tests also revealed that the PTM may be more useful for longer routes than shorter routes. For example, the model predicts the transit time for the above example route with an error of 1 day. However, for a short route, PTM estimated the transit time with an error of 0.59 days where the average transit time for the route is 0.7 days and the standard deviation is 0.7. Unless there are major disruptions, PTM will not improve on baseline prior transit time distribution probabilities for short routes. However, for long routes, it can help the shipment planner obtain more accurate estimates.

In today's environment, there is no automation of event extraction. It is primarily a manual task. Dow has independently explored options to procure this event data, but that did not fall into this project scope, but perhaps could be an option for those that wish to augment with additional data sources which prove to be more challenging to automate yourself. For the project scope the teamtouched the few top examples likely to be desired: social media, traffic and weather. For those selected in scope, the solution effectively extracted the necessary event data from Twitter, INRIX and NOAA. See PTM module for more details. It is recognized that certain internal events would still not be able to be automatically extracted, so the target was less than 100%. No concerns over reported numbers.

Risk Assessed Systematically: In this category, we looked at the impacted shipment and events combinations to ensure all could have a risk score assessed automatically.

With the scope for this KPI being that which feeds into the systems either automatically via the PTM module, or via the UI interface for the SCEM User, the team could ensure that 100% of these items would see risk assessment scoring. The concern here may be that unknown events exist that are not captured, and thus are not scored, so this KPI really does focus only on that which exists within the system. No concerns over reported numbers.

Targeted Automatic Alerts: In this category, we were concerned with alerts (emails/SMS) getting to the proper users. As part of our testing, we setup subscribed users in various subscriptions levels to test this capability, in both the Azure framework (using email and Twilio/SMS) and from SIMBA blockchain (email).

This KPI is somewhat driven with assumptions of proper setup and subscription of connected users (in business or customer service). These are part of the design setup and can be modified with assistance of the Admin User who makes modifications in the Mitigation Module (notification of mitigation needed), or where the users directly make updates via BlockChain interface (mitigation decisions alerts). The assumption then is with proper setup and maintenance, the capability is there, and it would be upon human change management to ensure the proper contacts are maintained. The team did explore more robust options where the system would read Dow's Organizational Charts to maintain business contacts and read the shipment/order details to subscribe CSRs actively, but were unable to scope these extensions into the project, but certainly would be possible. No concerns over reported numbers.

Performance Management System: And finally, in this category we assessed the capability of the reporting and measures, and KPIs for the entire integrated framework.

This KPIs is a little self-led, as this list of items ensures the completeness of such an objective. Therefore, it is pertinent to consider the options that were had with this item. The objective was to have an integrated, centralized and customizable system. What the system passes through the framework is somewhat variable, but no limitations were specifically designed to constrain the capability. In the goal the team supposed that some additional measures may be desired, *and perhaps the team would not pursue building the calculation for that new metric, but as the team developed in a visualization platform like MS PowerBI, the end user would have the capability to create their own measures in that platform. If it makes sense, that measure could then be built in the primary dashboard if it makes sense to leverage the measure into the core package. No concerns over reported numbers.*

Accessing the Technology

A report has been prepared that details what technology was in place prior to work beginning on this project, linked below. Further in each modules documentation there is more information on the expertise necessary to implement the technology.

MxD members can access all the software and APIs developed by this project on the MxD membership portal according to the membership agreement post-project. SIMBA Chain offers commercial Smart Contract as a Service (SCaaS) software for blockchain integration. If you are interested in licensing software, please contact MxD who will direct you to the appropriate team member. The mitigation planning module and performance analytics modules are deployed as examples only and may be deployed using any preferred dashboard or web tools. And to use the predictive analytics modules, you must have at least two years of historical shipment data available.

17-02-01 Report on Existing Technology.docx

Presentations

This section contains links to certain presentations created throughout the project duration, and a link to a folder with more than what is listed.

Predicting Distribution Transit Times for Outbound Logistics (M.S. Thesis, work in progress)

On-time product delivery is important since delays can cause major supply-chain disruptions. Because this disruption is primarily felt downstream, mitigating delay risks is typically addressed from an inbound perspective with either higher inventory levels or multi-supplier sourcing. This thesis describes a model that approaches distribution delays from an outbound perspective. At the core of the proposed approach is a supervised classifier which is based on support vector machines (SMV). The classifier is specific to a source-destination route between supplier and customer locations and is trained by using historical shipment, weather, and social media data. The proposed approach is validated using seven case study routes in the United States. These routes range in distance from 100 km to over 1,500 Km. Predictive models are developed for the seven routes. These models are able to estimate transit time for future shipments with a mean absolute error that is less than the standard deviation of the transit time. Moreover, an analysis of the models shows that the input features that contribute the most towards estimating the shipment transit time vary from one route to another. The main limitations of the proposed approach are the amount of available historical shipment data and the quality of the social media data.

This is the recording of our final project presentation at MxD. It covers the completed deliverables as well as a system demonstration. *Duration 1hr 39min.*

17-02-01 Final Presentation_2019.12.11.mp4

This is the presentation for the Workshop held in Indianapolis @ IUPUI on Nov 15th 2019. Guest speakers from INDOT, INRIX and Rolls Royce participated on panel discussions, in additional to participation from the project team. This workshop was not recorded.

11.15.19 Workshop.pptx

This is the presentation offered to the MxD membership as a technical review. It is a great summary of the project overview, problem statement, use cases, goals/objectives, KPIs and more. The MxD Project Member Technical Reviews were recorded and are available on the MxD portal.

MxD Project Member Technical Review_17-02-01_8.15.19.pptx

Lessons Learned

The team had many lessons learned, both positive in what went well, and learnings associated with what went poorly. They are summarized below.

What went well:

- Modular approach is key to encouraging adoption
- MxD project coordination and management
- Active technical support from manufacturer and cloud service provider (Microsoft)
- Azure is a great place to collaborate
- Highly engaged customer at Dow
- Availability of clean structured data with well-defined data catalog
- Key to meet with current Dow supply chain team (end users) to understand business processes and gaps early in the project
- Development approach of phase 1 independent modular development but then insisting on integration in phase 2
- Leaflet/IPyLeaflet is powerful free mapping software
- Defining use case/business case for event extraction (planning) distinction from risk assessment (SCEM), and business case for blockchain

What went poorly and associated learnings:

- Narrowing scope - caution implementers to fully understand available data and quality of data in advance
- Time complexity for external data extraction was underestimated
- Planning licenses for external data sources in advance
- Data sharing can take a long time because of security procedures (plan ahead)
- Should meet with current users and end users earlier (first month). This should be a required step in the project process.

- Azure notebooks didn't work out- built RA module using Jupyter Hub (in Azure) with built in authentication and that was great

Dow was very pleased with the lessons learned, as most were very positive. The agreement to develop the solution as modules versus one singular application was a key benefit for leveragability and general utility. It also helped speed up development, and ultimately no connectivity was sacrificed as a result. The project management aspect was also very well received from the team, as this was traditionally an extra role managed with the team resources, stretching the project principal thin. Dow also worked very hard on enabling quality data, cleansed and reliably available, to the project team. This was an extra effort to make available to the external partners, and the team encountered a few security barriers which were ultimately resolved, but it was a big plus to the project being able to share data efficiently.

There were some less than ideal events that came about during the project, and these merit call-outs as well. We underestimated the complexity to work with Dow's data security group, and the amount of time it takes to get approvals… very slow. The team also underestimated the time it would take to enable a custom data feed for the project, which took much longer than planned. Both issues were ultimately resolved, but the original estimates were dreadfully short. It would be ideal to get more upfront engagement from both the end-users (within the first month) and from the technology partners (to help with architecture design). These both occurred much later than they should have.

CONCLUSIONS & FUTURE WORK

This section includes the project conclusions, next steps and a transition planning aid. A transition plan document is included within the Final Presentation deliverable.

Next Steps

When considering next steps and how a MxD member would use these project deliverables, a good start would be to consider the available data sources. In order to create the predictive transit models, you would require a decent set of historical data with appropriate features (24+ months). You would also need to assemble a team of systems integration experts and developers that can help connect your digital resources, as well as functional experts who will own the functional process related to supply chain event management. You should also engage early with your data security team and determine the location you plan to house and access the code, as you can do this in the cloud or on premise. Additional preparation steps could include rationalizing the implementation scope (highest volume routes, highest risk regions, etc.) so that you can achieve ROI faster.

The developed framework is modular, so there are the options to bring in one, a couple or all of these modules, to address the specific needs/gaps which might exist within your existing system. The concept is that each module can run self-contained but would also work fully integrated as shown in the demonstrated framework. The approach assumes that certain manufacturers may only have a use case for parts of the framework or may already have a partial solution in place and need pieces of the framework to close on gaps within their existing processes.

Dow, perhaps like many other large manufacturers, is in a continuous improvement cycle where they are always entertaining and exploring beneficial platforms as part of our supply chain strategies. This includes visualization platforms which are in the process of being implemented (FourKites – Road, GVP/Transcore – Rail). Beyond these couple. there are additional platforms that are being considered which could impact the way Dow would internalize and scale the proof-of-concept demonstrated here. What if Dow had a relationship with DHL and pursued their Resilience 360 capability, and could import events that way? What if Dow had a relationship with Everbridge and used their event platform to push communications and track events? In which case a supported platform may already be in the implementation plan, it may make sense to utilize that platform and gain help in the way of scale and system support. The caveat to this is whether the manufacturer will be comfortable 'loading' supply chain 'crown jewel' data into a partner's cloud system. The concerns over security and leaks of manufacturer's most prized data would give pause to many leaders before they jump onto a commercial off the shelf solution and may be another reason one might want to build an internally hosted solution.

On the immediate horizon, next steps for Dow include additional demonstrations of the capability developed to further socialize and share the benefits. This includes sessions with Dow's internal business stakeholders beyond the project team members, and with the extended logistics team (mode and site logistics resources). Dowis holding off on an immediate scale up to pilot as Dow integrates this development with ongoing implementations, as the resources needed to do so are already aligned to ongoing efforts. Realistic expectations would be for some action to occur within six months, but technically there is no such barrier that Dow could not begin immediately, if strategy and resources were aligned to begin immediately.

Adoption

This capability represents a great example of a digital adoption opportunity. In such an example, we find ourselves using digital tools to their full extent. For many, this seems like a simple task, where the system offers a tested capability, the value case, and a framework to go get that value. However, there are challenges that cannot be overlooked.

These adoption challenges will encounter the traditional change management obstacles of People, Process and Tools/Technology. For this framework to deliver, and adoption strategy at each company would need to have a handle on these challenges.

There are also the concerns, when pursuing digital solutions, to lose some buy-in from some stakeholders who will perceive the change as just an excuse to use new tech for the purpose of using new tech, mostly ignoring the actual benefits of the technology. Blockchain and cloud have fallen into this category, but it simply means that additional efforts must be made in the change management plan to provide evidence to justify the need, which hopefully should be apparent from the demonstration provided.

Another adoption challenge that should be addressed is competing efforts. There are limited resources at a company, and their engagement on any given number of projects should be considered before setting targets for implementation. Best intentions are always in place, but if there are too many efforts requiring support from any key users, one may see adoption failure issues.

A final adoption challenge to note is the perceived value as converted to dollars. If the expected value to be realized with the framework is not valued the same, there can be confusion over if the investment is worth the effort. This can occur less in places where hard values are available (freight costs, man hour reduction) and more where we deal in soft values (customer experience, employee experience, reliability).

Some key points to consider when reviewing the framework developed are noted below. These would be a good spot to start the internal conversation with appropriate stakeholders about considering this framework.

- Getting Started
    - Identify existing internal manufacturer SCEM historical data sources (structured and unstructured) which will determine ability to deploy different modules
    - Assemble team consisting of system integrators, IT/OT, data security, system admins and supply chain event managers
    - Select data storage mechanism (cloud, on-premise)
    - Create rollout and training plan for supply chain managers who will pilot the system
- The system integration framework can be transitioned to any manufacturer and can support integration with new modules, data sources, and different vendor solutions
- All software modules require customization and integration to meet the needs of company-specific business processes
- The Predictive Transit module is self-contained with configurable integration tokens that needs minor adaptation to each manufacturer's specific data repository
- The Risk Assessment module is self-contained with configurable risk weightings
- The Mitigation and Performance Analytics modules are examples that can easily be deployed on preferred platforms
- SIMBA Chain is a commercially available SCaaS platform
- Modules currently use a SQL database

Transition Plan

Dow will deploy the SCEM Framework proof-of-concept at their Digital Fulfillment Center to gain stakeholder commitment to run a pilot for Dow business units. The blockchain integration will continue to be de-risked for commercialization by ITAMCO. And conversations will continue with supply chain visibility solution providers about the ability to connect their solutions using this framework.

The table below provides a catalog of all the project deliverables and their respective transition routes. Deliverables can transition through deployment at an industry member’s site, as an educational reference or through a commercialization effort. Each of these transition routes are detailed below.

Table 4: Deliverable Deployment Summary

#	DELIVERABLE FILE NAME	TECHNOLOGY INTEGRATION	EDUCATION	COMMERCIALIZE

1	Mitigation Planning Module	X	X	
2	Predictive Transit Module	X	X	
3	Risk Assessment Module	X	X	
4	ITAMCO-SIMBA CloseoutModule	X	X	
5	Performance Analytics Module	X	XX	
6	Overall System Architecture	X	X	
7	Report on Existing Technology_11.13.19		X	
8	MxD Project Final Presentation_FINAL.pptx		X	
9	Post Project BIP and IP Claims_12.11.19			
10	Presentations		X	

Transition Activities:

- **Mitigation Planning Module** – Dow intends to investigate this module further in it current state. Dow also intends to investigate alternative process scenarios where recommender logic may be of benefit for rapid decision making.
- **Predictive Transit Module** – Dow intends to look at available data sources and relationships with data partners, and then revisit the model developed.
- **Risk Assessment Module** – Dow intends to consider the existing developed module as an option for the SCEM team.
- **ITAMCO-SIMBA Closeout** – Dow is exploring the utility of blockchain in this specific use case, and the full benefit will depend upon the modules utilized. Dow is also exploring further use case where blockchain may be leveraged.
- **Performance Analytics Module** – Dow is carrying forward the KPIs derived from this project and will maintain them as the requirement for any solution which will be implemented.
- **Overall System Architecture** – Dow has selected Azure as the development platform of choice, so the majority of the system architecture will remain as designed.
- **Report on Existing Technology_11.13.19** – Dow will explore with their Enterprise Architects the available technologies that we partner with to determine which are strategic for the solution presented to be sustainable.
- **MxD Project Final Presentation_FINAL.pptx** – Dow will leverage the presentation to gain business support to carry the implementation beyond this proof-of-concept, into a scaled implementation.
- **Post Project BIP and IP Claims_12.11.19** – Dow will continue to partner with the project team and MxD to share as much as possible for the benefit of the MxD membership.
- **Presentations** – Education references.

mxdusa.org
@mxdinnovates
info@uilabs.org

1415 N. Cherry Avenue
Chicago, IL 60642
(312) 281-6900

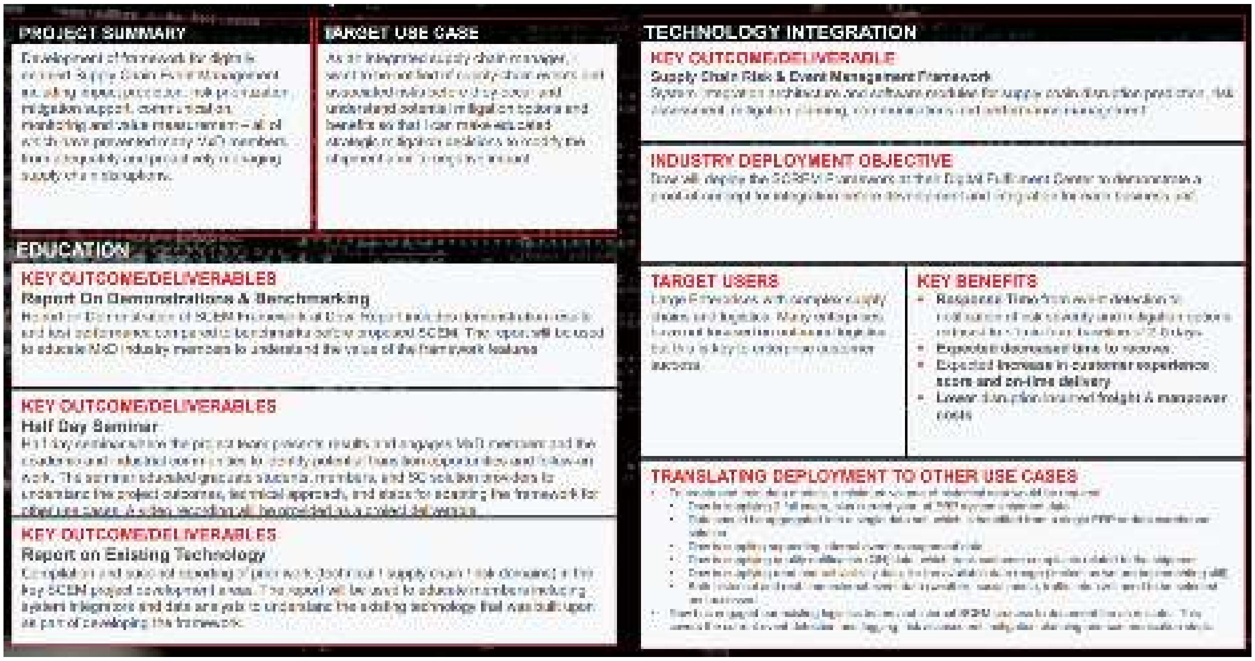

APPENDICES

Appendix A: Definitions

Decision Support Data

Business Group (from Order to Shipment Visibility report)

- Code representing a grouping of businesses at the company level

Value Center (from Order to Shipment Visibility report)

- Code representing a sub-grouping within a given business

Shipping Type (from Order to Shipment Visibility report)

- Code representing the type of shipment as Road or Rail

Distinction (from distinction report)

- Proposed customer identified, that helps differentiate by value importance

Segmentation (from segmentation report)

- Proposed customer identified, that helps differentiate by market or service needs

Variable Standard Margin (from Margin report)

- Average margin of customer at global level

Average Order Touches (from Order Touch report)

- Average times a customer makes changes to their orders, adding rework/challenges

Shipment Count (from Order to Shipment Visibility report)

- Sum of all shipments, count by individual deliveries

Contract (from contract tool report)

- Indicator if we are bound under contract to provide pounds or service within a period

Appendix B: Demos and Presentations

This section includes the key demonstration videos and presentations shared during the project. There were also internal Dow presentations offered to key business and functional partners, but as this included actual customer data, they cannot be shared. Some outcome from these sessions did include the generation of additional context points for potential Dow implementation, which was covered under the transition section.

MxD Project Final Presentation_FINAL.pptx

17-02-01 Final Presentation_2019.12.11_video

MxD Project Member Technical Review_17-02-01_4.5.19_video

MxD Project Member Technical Review_17-02-01_4.5.19_slides

MxD Project Member Technical Review_17-02-01_8.15.19_video

MxD Project Member Technical Review_17-02-01_8.15.19_slides

11.15.19 Half-Day Workshop

Appendix C: Validation & Testing

The framework was developed using the agile methodology. Unit testing for the modules was performed and demonstrated during the first three sprints of the project. Initially, the modules used synthetic data. After the third sprint, testing was based on production data. The first three sprints also enabled the definition of module integration requirements. Functional and integration testing and demonstration was the focus of the second set of three sprints. Final, integration and user acceptance testing was demonstrated during the last sprint of the project. Live demo using future shipments have also presented during project demonstration.

The PTM module was tested and validated using seven sample routes. The shipments for these routes spanned a period of two and half year starting January 2017. Unit testing was performed on each component including the machine learning components, the data extraction components, and the user interfaces. This testing verifies that data is extracted and inserted into the database correctly. This process was applied to nearly 10,000 shipments.

Integration testing within the PTM module was performed following the expected process flow. Operational shipment data was uploaded into the framework for the seven routes, weather and

mxdusa.org
@mxdinnovates
info@uilabs.org

1415 N. Cherry Avenue
Chicago, IL 60642
(312) 281-6900

Twitter data for all the shipments in the routes was extracted and populated in the database using the data extractors, finally, predictive transit models were developed for each route. Integration testing was then extended to the user interfaces in order to validate that user inputs are captured correctly and expected output are displayed as expected. This testing was performed at the individual shipment level as well as in batch mode.

Functional testing on the PTM was conducted for the predictive models. The predictive transit models for the seven routes were tested using a total of 2000 future shipments. The predicted transit times for these shipments were compared to the actual transit times and the mean average difference between the predicted and the actual transit times for the test shipments was less than 1 day.

Appendix D: User Resources

SIMBA Chain Communications Module Documentation Folder

Mitigation Planning Module Documentation Folder

Predictive Transit Module Documentation Folder

Risk Assessment Module Documentation Folder

System Architecture