



AUTOMOTIVE CRM TECHNICAL OVERVIEW



LOYALTY LOGISTIX
AUTOMOTIVE INTELLIGENCE

CONTENTS

Introduction.....	2
Architecture.....	2
Hosting infrastructure	3
Data model	4
Customer Relationships.....	4
Wholesale v retail view of customer data.....	5
Data provider integration	8
Deduplication	8
Customer deduplication	8
Customer deduplication rules	9
Vehicle deduplication.....	10
Multi-channel campaign management.....	10
Campaigns	10
Filters	10
Channels.....	10
Merge Fields.....	11
Schedules	11
Campaign Monitor	11
Email delivery service and email event tracking	11
Rewards management.....	12
Reward Schemes	12
Types of Rewards	12
Creating Rewards	13
Redeeming Rewards.....	13
ABI	14
Dealer CSV Files.....	14
LLFTP	16
Data Processing / ETL Staging Server	16
Blob Storage	16
Publication / Subscription Services	16
SQL Databases	17
Dashboards.....	17
Reports	17
Data Exports	17
Conclusion	18

INTRODUCTION

Loyalty Logistix is uniquely placed as a technology provider with its dedicated CRM offering for the automotive industry. Leveraging its many years of specialist expertise in the sector, Loyalty Logistix has created a system especially engineered for vehicle manufacturers, national sales companies, and retailers.

Key features of the suite include:

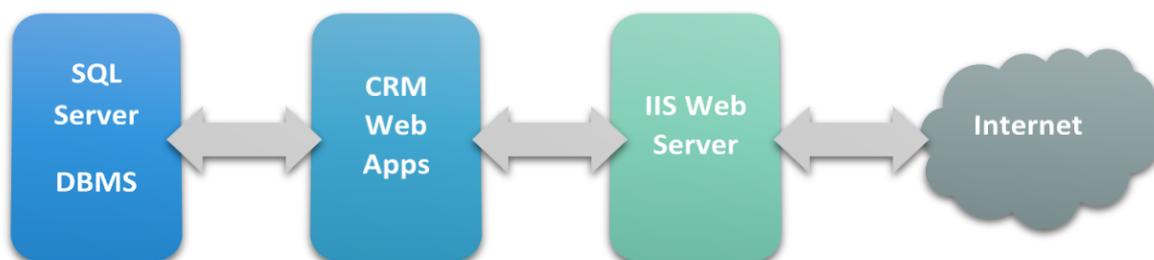
- Master Data Management,
- Lead Management,
- Campaign Management,
- Rewards Platform,
- and Customer Experience via customised vehicle-owner web sites.

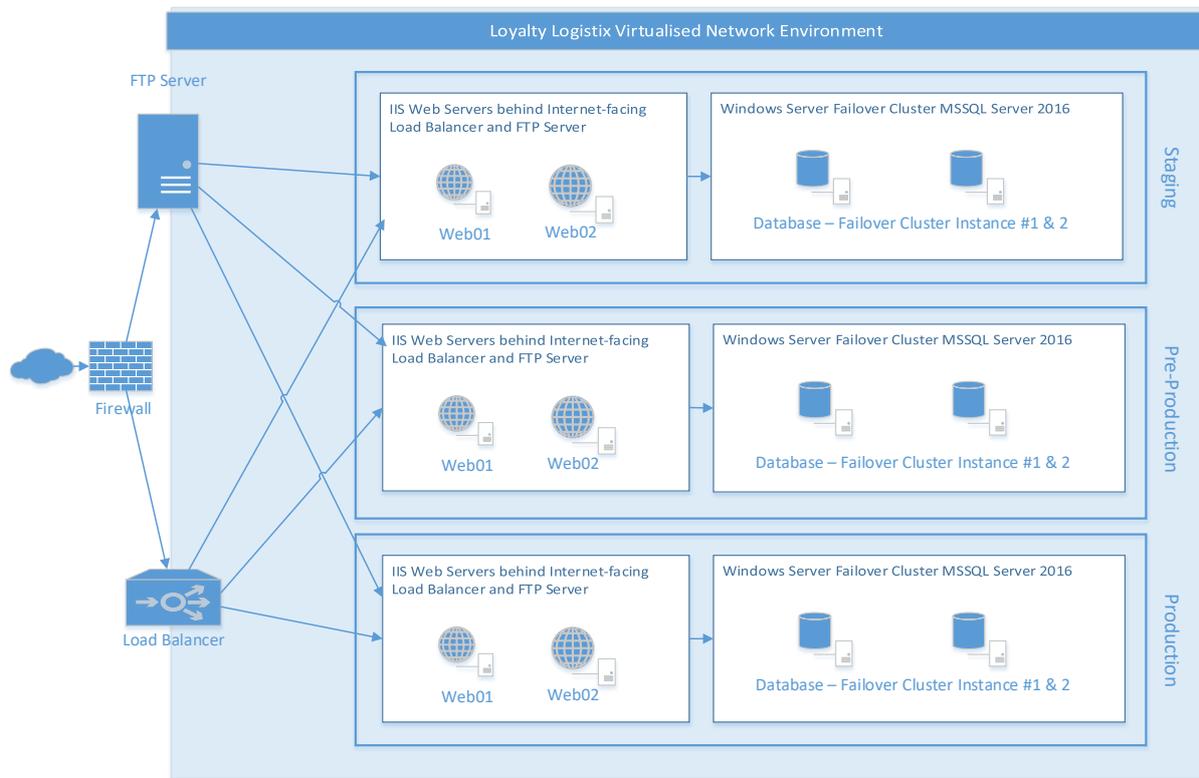
This document provides a high-level technical overview of the system, its interactions with 3rd party systems and its approach to data security and system resilience, and assured data quality.

ARCHITECTURE

The system is built on an ecosystem of Microsoft technologies:

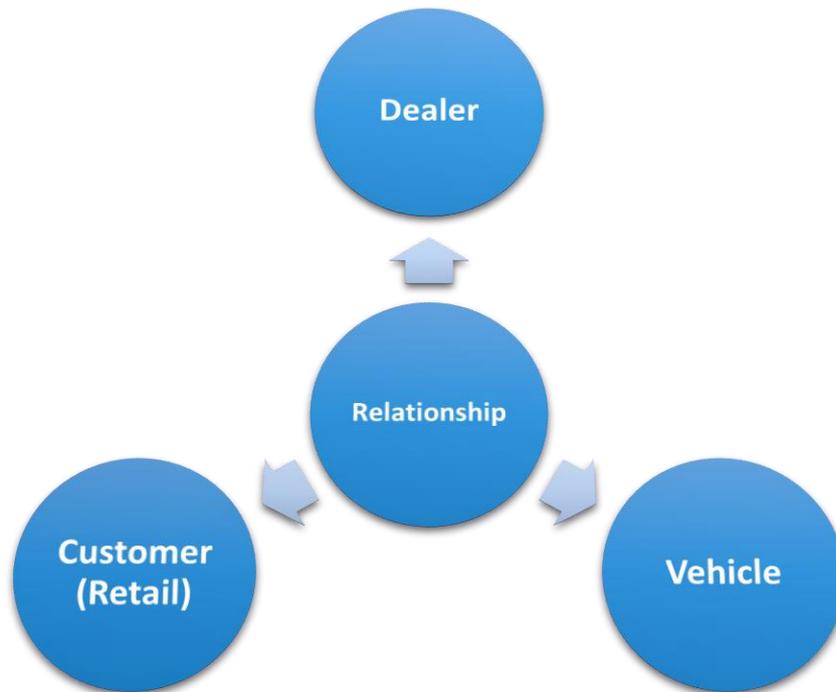
- SQL Server database management system,
- IIS web server,
- .NET framework,
- ASP.NET and MVC.





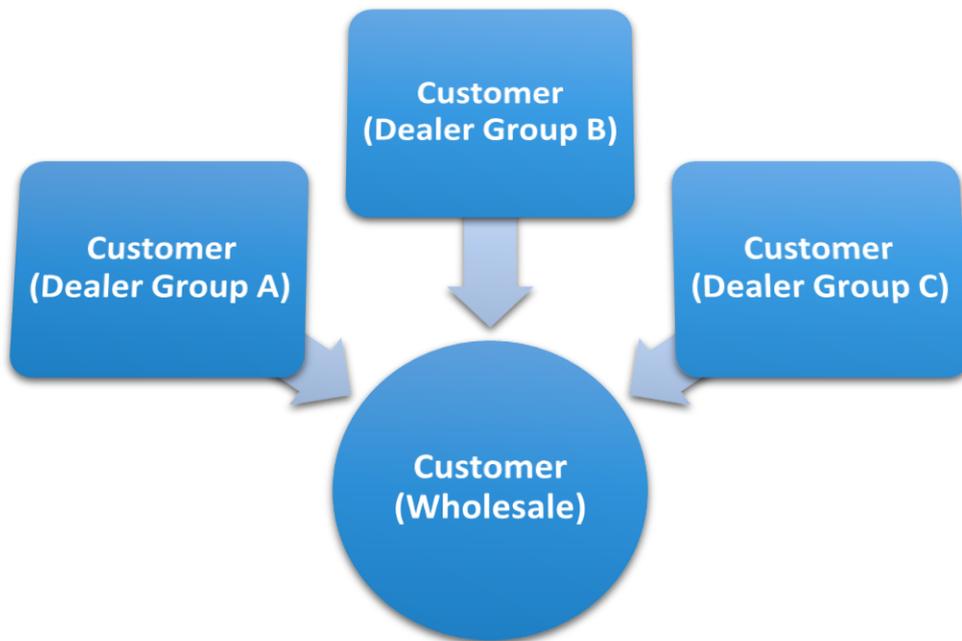
The above diagram describes a typical hosting solution. Each environment requires at least a separate web and SQL database server. However, depending on each organisation's appetite towards risk, server down time and the need for optimal performance, additional web and database servers can be added. Of course, additional servers can be added only to the production environment to reduce the overhead, but this would run the risk of issues occurring when updating the live environment, as like for like testing will not have been possible. It is advisable always to have pre-production and Staging environments that are identical.

Customer data will only be held on the Pre-Production and Production servers and access to these environments will be restricted to the maintenance and deployment personnel. Obfuscated data will be used on the staging environment.

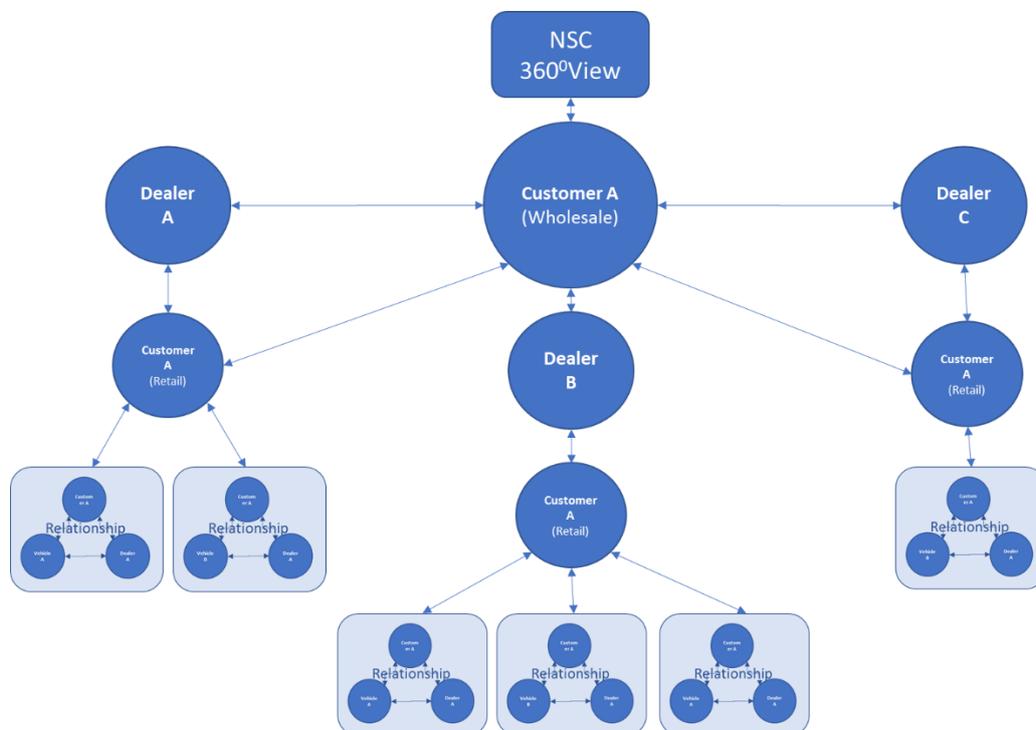


Central to the system is the concept of the Relationship between a Dealer, a Customer and a vehicle. This triangular relationship is always unique at any point in time in the system. Various types of Relationship are supported, such as Selling, Servicing, Prospect, Contact only, etc. A central point is that customers may know several dealers, buying a car in one dealer, servicing it at another, having a breakdown at another, etc., and of course customers change their cars over time, and the car may get a new owner, who also knows other dealers! The Relationship between all three parties is key, and the date when a relationship terminates, is the date at which the customer sells his car, whether it is a trade-in or a straight sale to a third party.

Wholesale v retail view of customer data

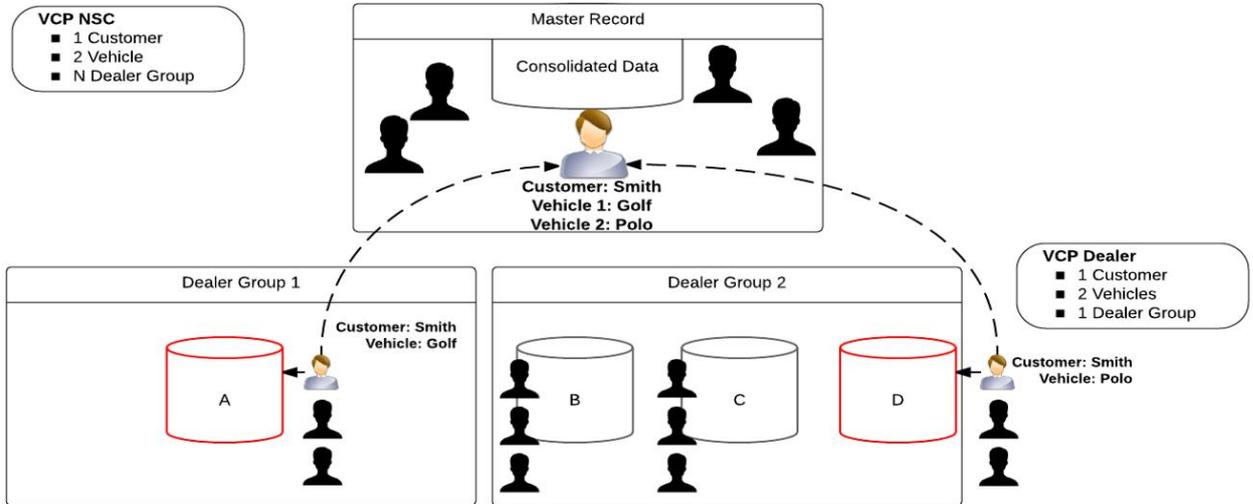


For cases where a customer has relationships with a number of different dealers within distinct Dealer Groups, data confidentiality is maintained by having separate customer records for each Dealer Group. These separate records are, however, aggregated into an overall record visible to wholesale-level users of the system. This is automatically kept up to date whenever changes are made to any of its associated retail customer records.

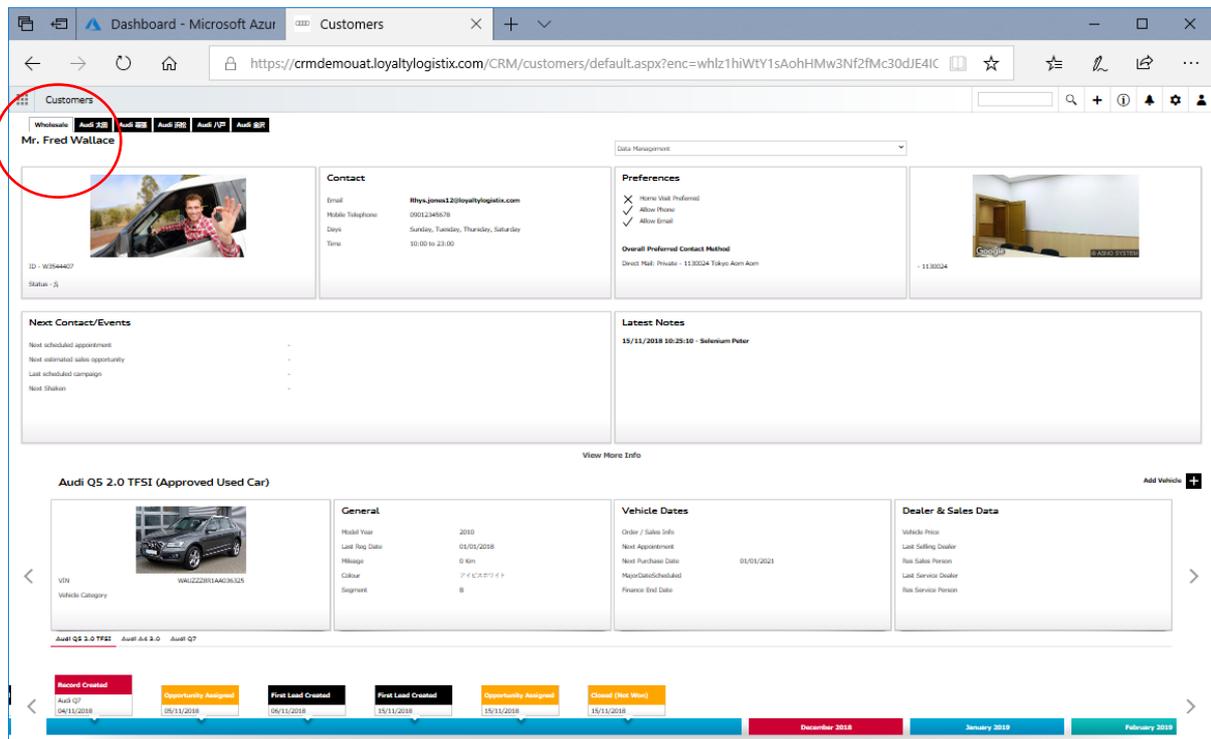


NSC VS RETAIL RECORD (EXAMPLES BELOW)

Customer Wholesale and Retail Record



WHOLESALE VIEW



RETAIL VIEW FROM 1ST DEALER

Dashboard - Microsoft Azure Customers

https://crmdemouat.loyaltylogistix.com/CRM/customers/default.aspx?enc=whlz1hiWY1sAohHMw3Nf2fMc30dJE4IC

Customers

Mr. Fred Wallace

Res Sales Person: Audi 2018, Audi 2017, Audi 2016, Audi 2015, Audi 2014, Audi 2013

Res Dealer: Audi 2018

Next Contact/Events

Next scheduled appointment	-
Next estimated sales opportunity	03/01/2018
Last scheduled campaign	-
Next Shaken	-

Latest Notes

Audi Q7

Vehicle Category: V7N

General

Model Year	03/01/2018
Last Reg Date	03/01/2018
Mileage	0 Km
Colour	
Segment	

Vehicle Dates

Order / Sales Info	
Last Appointment	
Next Purchase Date	01/01/2021
MajorDateScheduled	
Finance End Date	

Dealer & Sales Data

Vehicle Price	
Last Selling Dealer	
Res Sales Person	
Last Service Dealer	
Res Service Person	

Record Created

Audi Q5	WAUZZZ8R1AA036325	03/13/2018
Audi Q7		04/13/2018

Timeline: November 2018, December 2018, January 2019, February 2019, March 2019, April 2019

RETAIL VIEW FROM 2ND DEALER

Dashboard - Microsoft Azure Customers

https://crmdemouat.loyaltylogistix.com/CRM/customers/default.aspx?enc=whlz1hiWY1sAohHMw3Nf2fMc30dJE4IC

Customers

Mr. Fred Wallace

Res Sales Person: Audi 2018, Audi 2017, Audi 2016, Audi 2015, Audi 2014, Audi 2013

Res Dealer: Audi 2018

Next Contact/Events

Next scheduled appointment	-
Next estimated sales opportunity	-
Last scheduled campaign	-
Next Shaken	-

Latest Notes

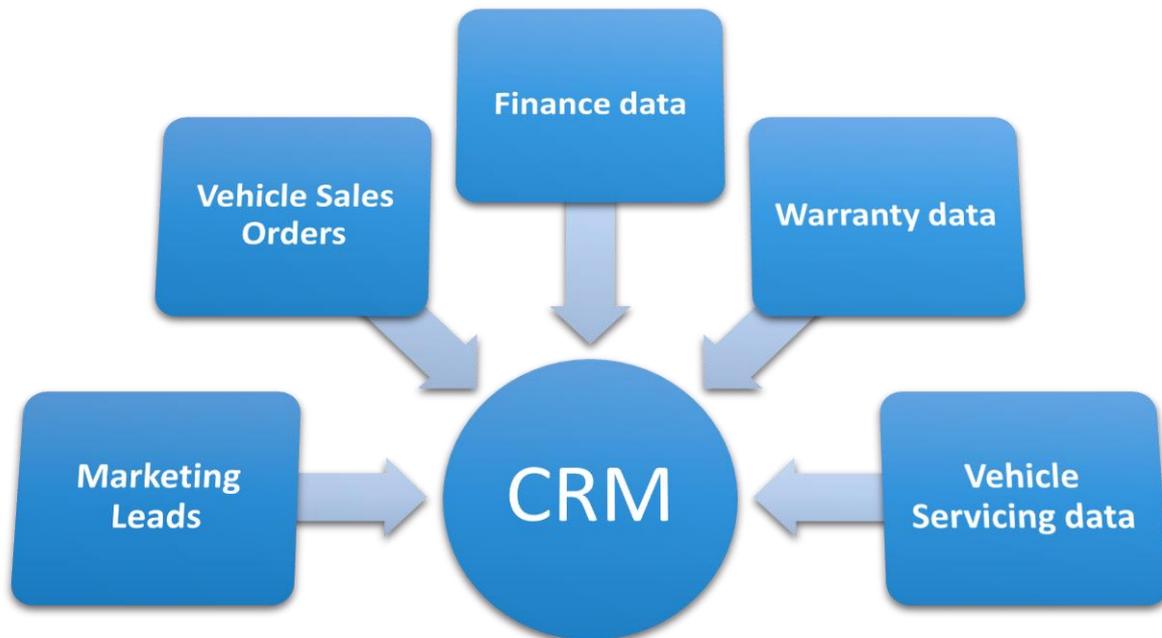
Audi Q7

Record Created

03/13/2018	
------------	--

Timeline: November 2018, December 2018, January 2019, February 2019, March 2019, April 2019

DATA PROVIDER INTEGRATION



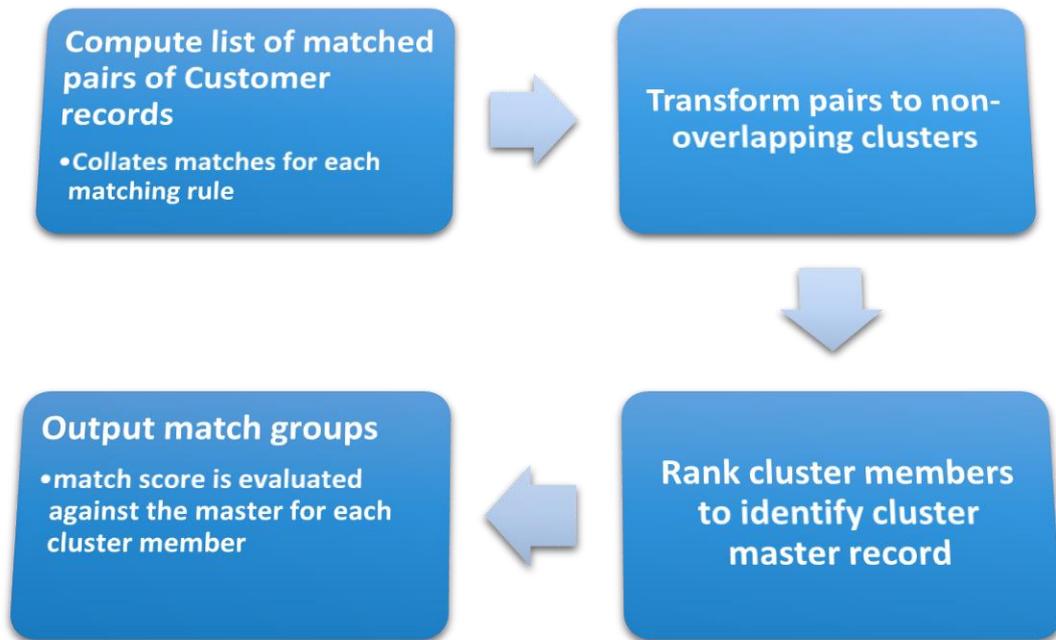
The CRM system integrates with a range of external data providers, acting as a central master data store. As it aggregates data from several sources with differing formats, crucial to its operation is its deduplication capability, which is the subject of the next section.

DEDUPLICATION

The system's inbuilt deduplication features provide facilities for automatically identifying and merging duplicate customer records, both at the retail and wholesale levels. Also available is the facility to detect and merge duplicate vehicles per customer.

CUSTOMER DEDUPLICATION

Customer deduplication is performed using a proprietary rule-based attribute-matching algorithm which computes numeric match scores. The main steps of the process are illustrated below.



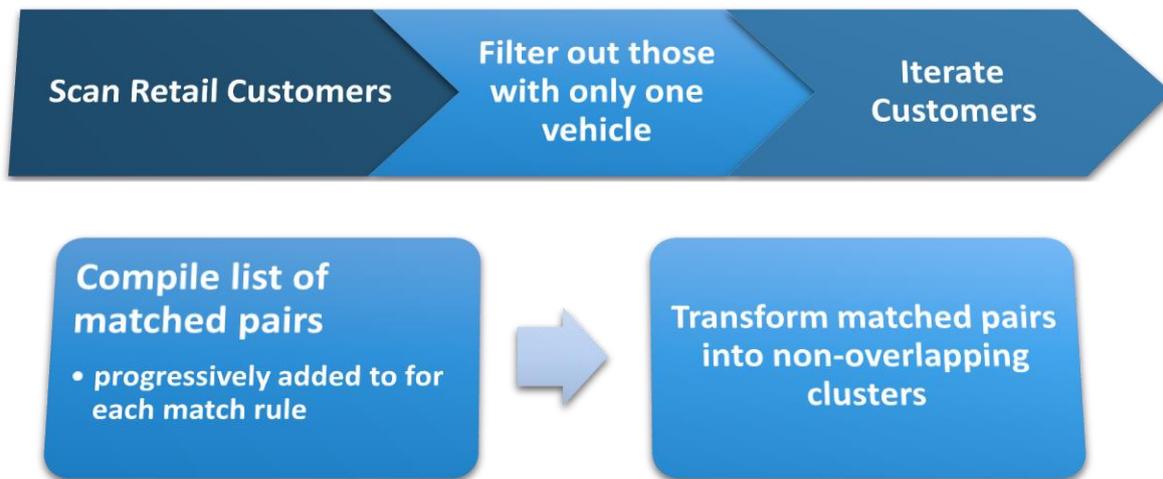
CUSTOMER DEDUPLICATION RULES

Rule ID	Mandatory field	required additional field	Match Lvl	Identifier level	data import/interface	manual entry/user	migration
	Name						
DupR1	*mandatory	Telephone number (either one) *only full match	High	additive identifier	add to duplicate list	add to duplicate list	add to duplicate list
DupR2		email Address (either one) *only full match	Very High	100% identifier	Update existing	add to duplicate list	Update existing
DupR3		Date of Birth *only full match	High	additive identifier	add to duplicate list	add to duplicate list	add to duplicate list
DupR4		Address (either one) *only full match	Very High	100% identifier	Update existing	add to duplicate list	Update existing
DupR4.1		ZIP	High	additive identifier	add to duplicate list	add to duplicate list	add to duplicate list
DupR4.2		City	High	additive identifier	add to duplicate list	add to duplicate list	add to duplicate list
DupR4.3		Street	High	additive identifier	add to duplicate list	add to duplicate list	add to duplicate list
DupR5		External ID *only full match	Very High	100% identifier	Update existing	add to duplicate list	Update existing
MigR1		Vehicle					
MigR1.1		VIN					
MigR1.2		Selling Dlr/Group					
MigR1.3		Registration Date					
			Very High	100% identifier			Update existing
MigR2		Vehicle					
MigR2.1		VIN					
MigR2.2		Selling Dlr/Group					
MigR2.3		Registration Date					
			High				add to duplicate list

maybe we can use this as failsafe:
 1. flag all cases that fail just because of either of these criteria not matching
 2. abort vehicle relationship creation for the duplicate and write all data to a log, so it can be checked.
 3. this way, we don't create accidental used car-status' and can check later.
 4. if not even these conditions are met, data quality is simply too bad and the system is not at fault

VEHICLE DEDUPLICATION

The duplicate detection process is depicted below



MULTI-CHANNEL CAMPAIGN MANAGEMENT

CAMPAIGNS

Users can create named campaigns with which to market customers using multiple channels. These can be set up either by a dealer user, targeting only that dealer's customers, or as an NSC (National Sales Company) user, targeting customers from all dealers, with an option to do this on a wholesale level, as the NSC, or on a retail level, on behalf of the customer's responsible dealer.

FILTERS

It is possible for a user to decide which customers should be included as part of a campaign by filtering on properties of the customer data, allowing a user to strategically target their marketing efforts. Most properties relating to the customer, their vehicles, leads and opportunities are available for filtering on. As an example, it is possible to create a campaign targeting customers with a specific existing vehicle or interest. The user is informed how many customers will be targeted using the filters and can preview which customers these are. It is possible to combine multiple filters to create a filter-set which can optionally be saved as a template for use in other campaigns.

CHANNELS

Multiple channels can be used for marketing, such as email, mobile email, direct mail, phone. A WYSIWYG editor has been provided to define the html used in emails. Labels can be created for sending direct mail and attack lists can be created for phone calls.

MERGE FIELDS

Customer data can be used to personalise the output from a channel using merge fields. In the editor, this appears as the name of the field to be used. This allows us to address the customer personally and display information relevant to them.

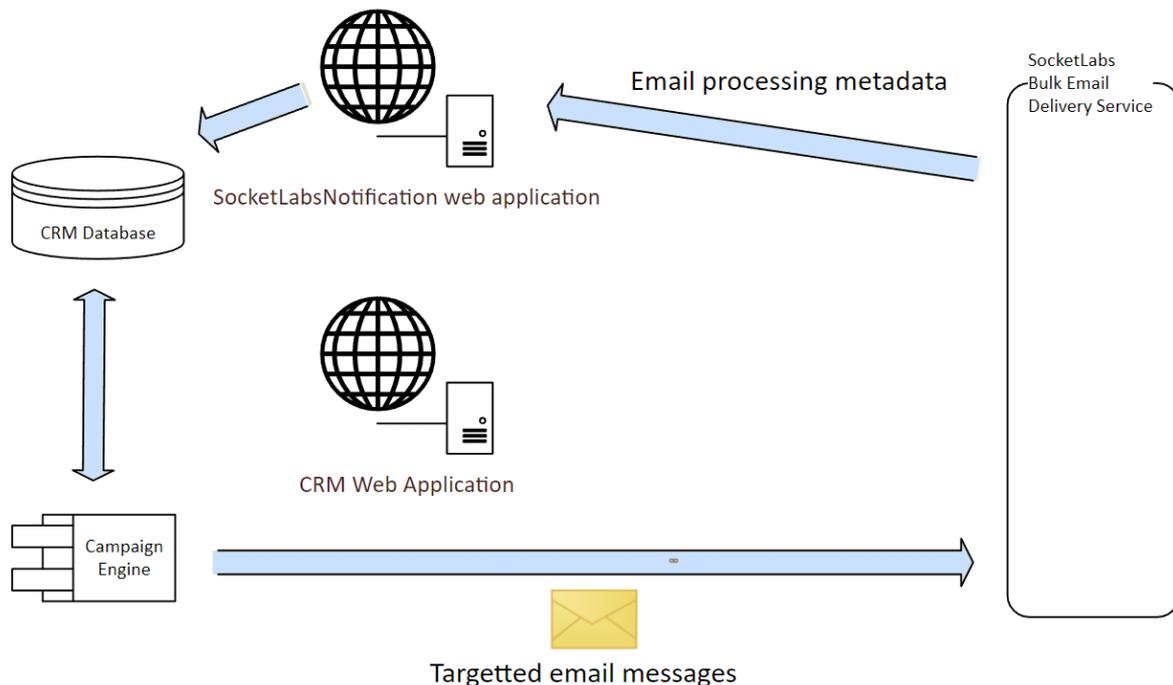
SCHEDULES

Campaigns can be set to run either manually or automatically. The automatic version includes options to schedule the campaign to run at a future date and time either once or on a recurring basis.

CAMPAIGN MONITOR

It is possible to monitor the campaigns being sent out via a monitor that reports on how far the campaign has progressed and if it has completed.

EMAIL DELIVERY SERVICE AND EMAIL EVENT TRACKING



In order to be able to report on activity (deliveries, failures, complaints and opens/clicks) related to email messages sent from CRM Campaign Management via SocketLabs' delivery service, associated email metadata is stored in a set of tables in the CRM database.

These tables are populated via SocketsLabs' Notification API. This API pushes individual notifications of email events to a handler application, which must be registered with SocketLabs when the API is enabled for a given SocketLabs server.

A web application, SocketLabsNotification, has been created for integration with CRM.

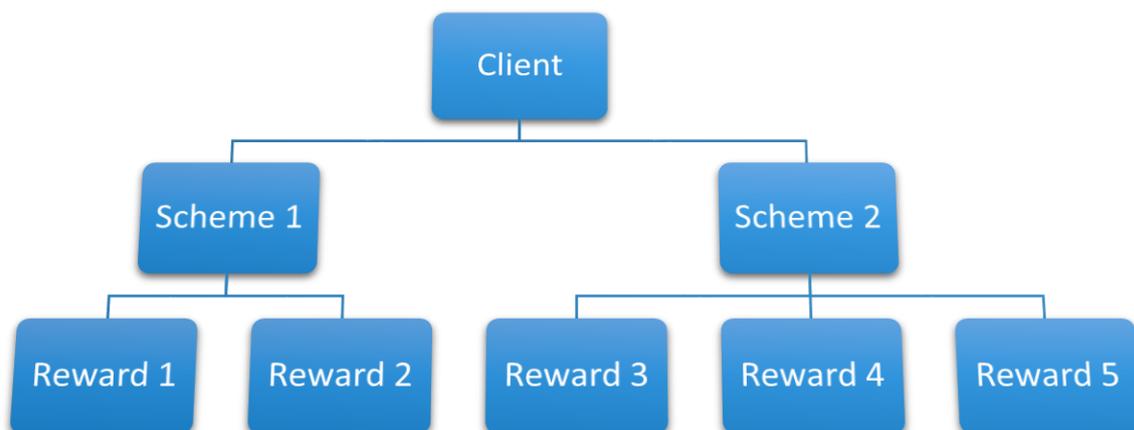
SocketLabsNotification has a single page, Receiver.aspx, which accepts the requests. The API passes data to the handler page via POSTed form variables. Receiver.aspx first of all authenticates the request and then parses the email event type and finally writes a record to the appropriate table in the database.

REWARDS MANAGEMENT

The system is configurable to enable customers to opt into an aftersales rewards scheme, setting up a personal account through which they can be targeted with marketing material informing them of special offers that are applicable to them. The system uses a sophisticated filtering system to allow highly targeted offers, which can be combined into a marketing campaign (see Section 6) to reduce repetition on the user's behalf.

REWARD SCHEMES

The client can create a single or multiple reward schemes that a customer can opt into. By choosing to have multiple schemes, such as Gold, Silver and Bronze programs, it is possible to configure attributes such as introductory bonuses, whether the scheme is available to new and/or existing customers, and how long the scheme is available for, amongst other things. Each reward is associated with a specific scheme by default, with customers only able to view and redeem rewards that are part of their scheme. There is also the ability to set whether the reward scheme is purely electronic, or whether the customer will have a physical membership card that they can use at participating dealer sites.



TYPES OF REWARDS

Each reward is completely customisable by the user. The system can use both a points and monetary setup, where the conversion ratio is configured for the scheme, and the two currencies are interchangeable. Therefore,

each reward can grant the customer a monetary discount, points added to their account, a physical item to collect from the dealership such as car accessories, or even a combination of all three if desired.

A customer can only redeem a reward if they have satisfied the conditions associated with it. These come in a variety of different forms, examples of which are given below:

- When the customer has accrued a certain number of points
- The reward is available for a limited time only
- One reward's redemption triggers the availability of another reward
- The customer owns a certain make or model of vehicle
- A promotion code is submitted by the customer that enables the reward

In addition to this, rewards can be configured to be automatically redeemed as soon as the customer is eligible for them, or to appear in a list of available rewards that are grouped into categories.

CREATING REWARDS

Users of the system are assigned different roles, and this creates a hierarchy when it comes to creating rewards. Network level users can create rewards that target customers nationally, dealer group admin users can have their rewards redeemed at all dealers that are part of their group, whereas dealer users can only create rewards that can be redeemed by customers at their dealership.



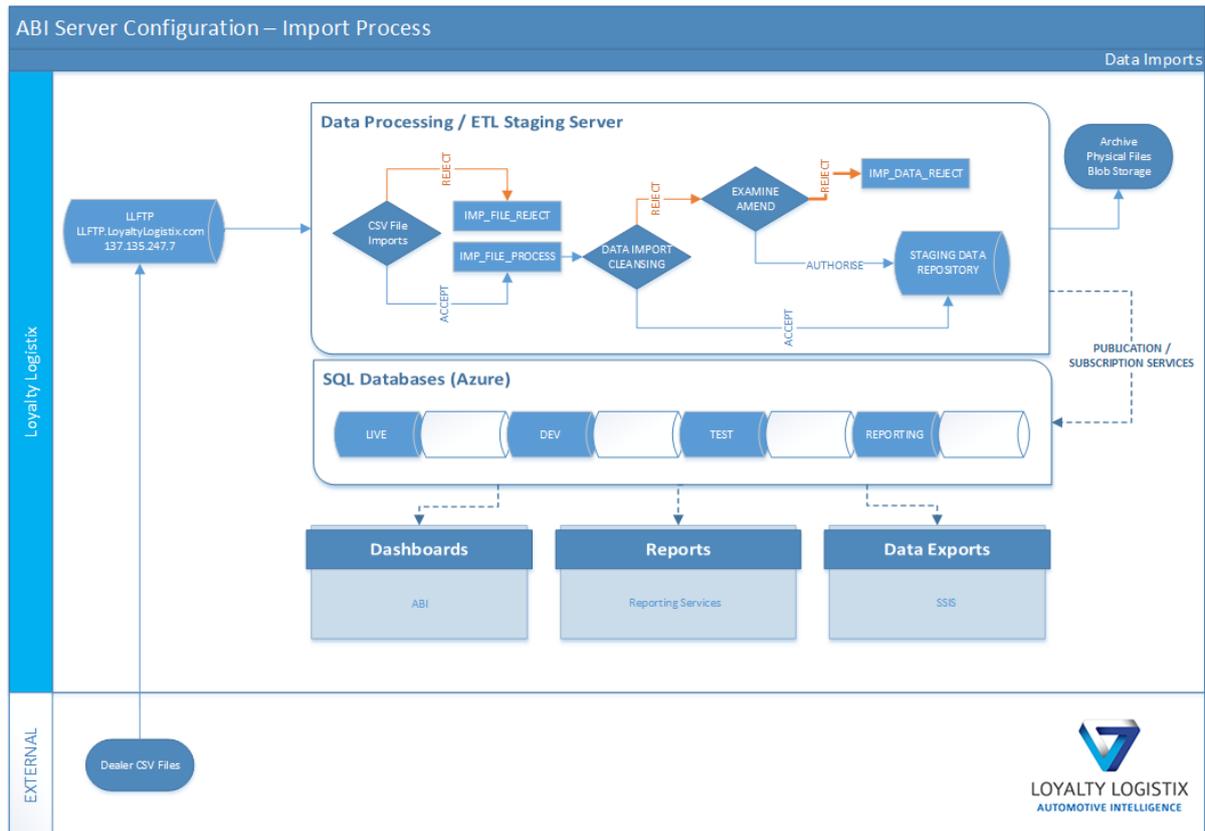
The system allows the reward owner to individually select participating dealers and groups, meaning a dealer group can still target specific dealers if desired. National campaigns can be set to target specific groups or individual dealers, ensuring ultra-flexibility.

REDEEMING REWARDS

Rewards can be redeemed when a customer visits a dealership, at which point the dealer can log into the system and perform the redemption on behalf of the customer. Alternatively, the system is able to connect to an online web portal, through which the user can log on themselves and redeem the reward to their account. For this, the customer can either have a specific login for the portal or use single sign-on.

ABI

DMS Data Journey

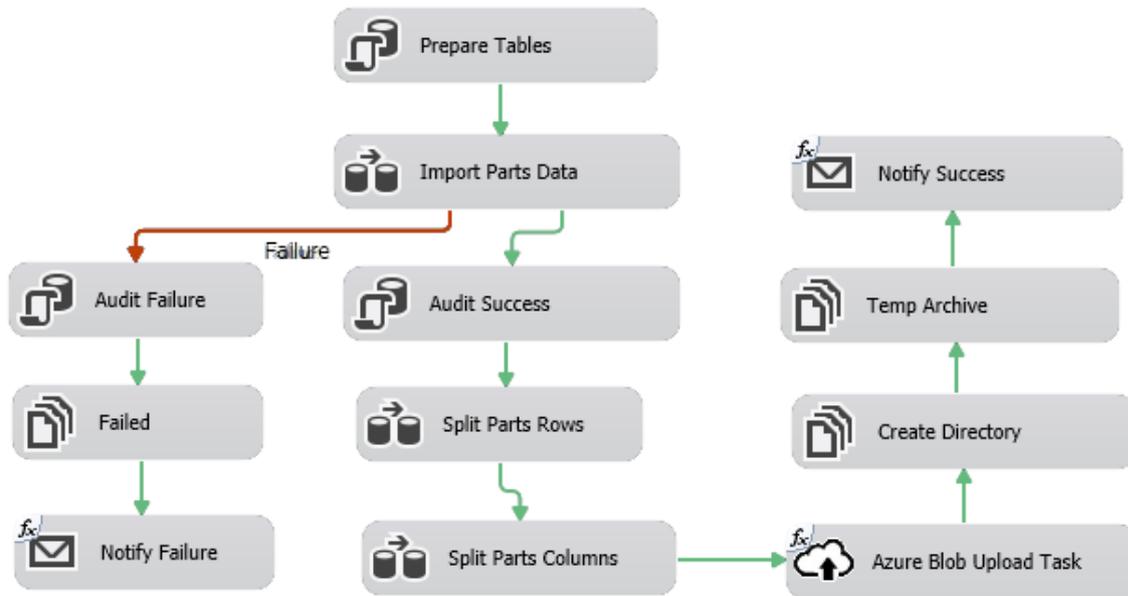


DEALER CSV FILES

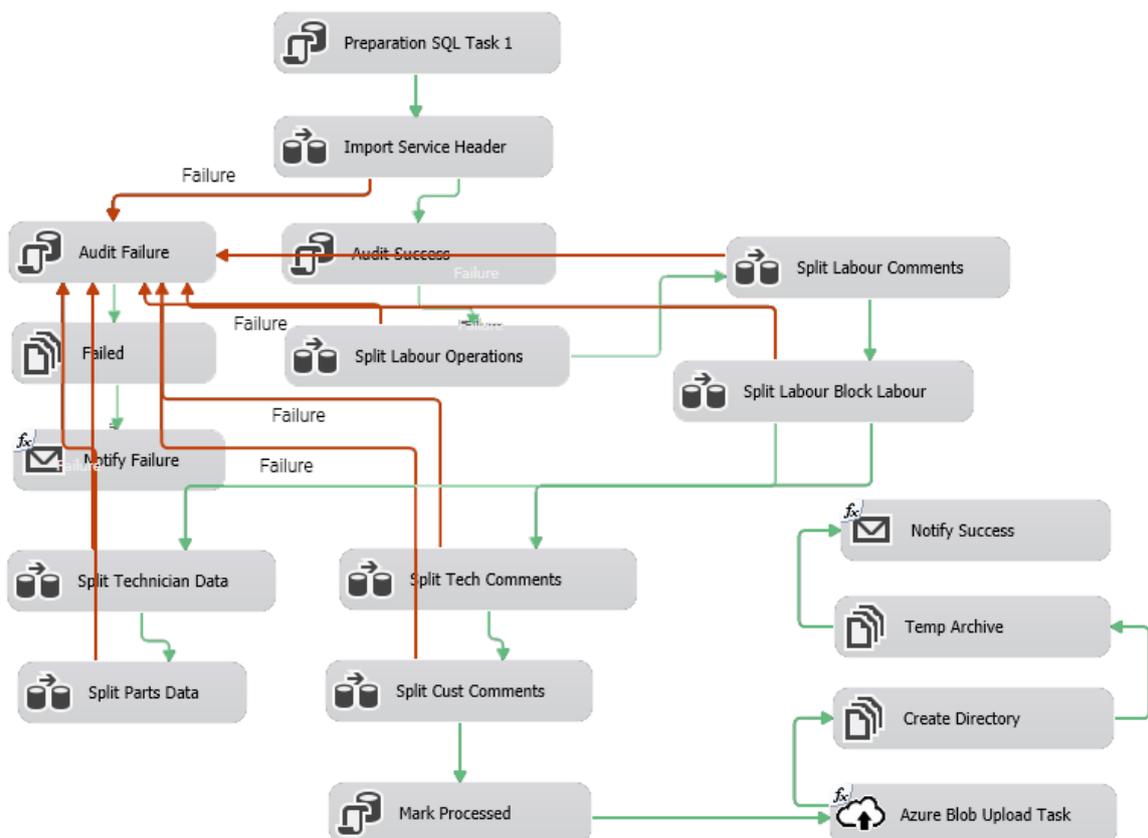
Dealer DMS Systems send in the required file types, usually consisting of the main 7, Service History, Labour Lines, Part Lines, Customers, Vehicles, New Stock, Used Stock. We are also now collecting Future Bookings and Workshop Transactions to provide additional functionality within the ABI dashboard.

Files are sent in via sftp using automated cron jobs to be received by LL prior to 06:30am for our main daily import routines.

Parts File



Vehicle Servicing File



LLFTP

LLFTP is our main Incoming ftp Server. All clients are setup with individual sftp accounts with allocated storage locations and can be configured independently. All accounts are by default configured as sftp using port 22. We also have the ability to setup public keys if required.

The sFTP server has a static external IP that can be used for firewall end to end protocols and we can restrict to specific IPs for specific accounts.

DATA PROCESSING / ETL STAGING SERVER

Our core ETL server (Extract, Transform, Load) is a VM machine with SQL 2016. We utilise SSIS to handle the Incoming csv files from the sftp server. All import routines for each file type and each DMS System are kept within a specific SISS project separate from any other. Once a routine is fully configured and tested we publish it to our central package repository which is what is used for the live routines. It is not possible to develop on these routines and publication is controlled. Every contract has its own database and its own security access. All databases are encrypted as standard using SQL 2016.

The process consists of pulling over the specific file type from the sftp server (and storing details of the incoming file within an audit table, we also add a date and timestamp to the end of every filename for reference.

The file goes through an initial pass to ensure the structure conforms to the specification (no of columns, line delimiters etc. We can setup separate import routines per DMS Provider so can tweak the routine if the system has certain restrictions re outputs compared to another (delimiters, CRLF etc). All the incoming files are then loaded into an initial staging table and the detail of the files and number of rows etc. are added to the audit tables (For KPI reporting).

We then take the data from the staging table and insert into the main data table using Merge Into routines, we use this to enable to handle duplicated records. Duplicates are hived off into a separate table purely for reference and archiving and the remaining legitimate records are inserted. Again, the audit table is updated with the number of records successfully inserted for auditing... We also stored the originating filename within each record of the system for traceability back to the source file in case of any queries relating to data.

BLOB STORAGE

Files that are rejected due to non-conformance are placed into a failed folder and a notification is generated to the systems team for further investigation and either fixing or referring to the DMS Provider. All successful files are then stored within a Cold storage account for archive. Every contract has its own storage area and all files are encrypted.

PUBLICATION / SUBSCRIPTION SERVICES

We have been using a mix of SQL stored procedures and SSIS packages as well as CDC (Change Data Capture) to migrate data to their respective destination databases (as explained below) but have now developed and currently implementing the new Publication and Subscription services within SQL 2016 for migration of the data from the core ETL databases to their respective databases. This gives us a lot more flexibility not only providing

scheduled data transferral but also instantaneous as and when required and the ability to spin up new servers for test and training purposes. We can also setup specific publications e.g. for Dashboard Only data or Reporting Specific datasets which are smaller and a lot faster.

SQL DATABASES

The Publication and Subscription services give us the ability now to utilise the new cloud sql database instances (rather than full vm's) opening a lot more functionality. Performance is monitored and scales automatically, all data is fully encrypted, instant data recovery etc. It also makes it far easier to have multiple database instances for the various activities required. A live and dev database environment with corresponding dashboards, a separate database specifically for reporting purposes and the ability to create ad hoc instances e.g. for training or specific testing purposes aside of the dev environment.

DASHBOARDS

Our dashboards are now being developed using PowerBI (replacing Silverlight) as our main tool with for more functionality, performance as well as improved aesthetics. Currently dashboards are hosted on VM's but we are currently developing them as cloud webapps again for performance and the added functionality offered for performance, availability etc. This also gives us more flexibility on provisioning ad hoc dashboards for testing and training and can be deployed fairly instantly.

REPORTS

Aside from the reports provided within the dashboard applications we also utilise SSRS Reporting Services for provision of automated reports and KPI for both our own internal use and monitoring as well as providing detailed information to our clients, e.g. daily KPI report detailing files and records imported and current system statuses....

DATA EXPORTS

Depending on the contract requirements we can also provide data extracts, either direct from the live or from the reporting server. We usually send out via our core sftp server direct to the clients sftp server with a IP route configured. We can also supply sftp accounts specific for clients to login to us and pull-down files as required usually for smaller scale or ad hoc exporting.

CONCLUSION

This document has attempted to provide an introductory overview of the Customer Platform. The reality is that LLL has collectively spent decades developing systems and processes for use in the Retail Motor Trade. The Customer Platform is the culmination of these decades of experience. It is believed to be unique in the sector, in that it offers an end to end solution, starting with new Leads for Sales or After-Sales through to after-sales marketing and Customer Rewards.

The document has described the fundamental principles of the operation of the system, as well as some of the detail of the technical infrastructure that customers enjoy, together with the nature of interactions with 3rd party systems vendors.

The document has attempted to avoid any sales or marketing related content, focusing on how the systems and processes interact to provide a solid and workable solution to a market that looks seductively simple until one gets into the depths of understanding of the real requirements of any National Sales Company or Importer operating in the current technical, economic and political climate.