

David Morgan-Gumm

D A T A P L A T F O R M M A N A G E R

C O N T A C T



+447949905787



davidmorgangumm@gmail.com



[david-morgan-gumm-450751133](https://www.sqlsquared.co.uk)



www.sqlsquared.co.uk

K E Y S K I L L S

- ✓ C-Suite / Senior stakeholder management
- ✓ Production of technical roadmaps and strategies
- ✓ Cloud data architecture and administration
- ✓ Database development and engineering
- ✓ DevOps implementation
- ✓ Positive and inspirational team leader
- ✓ Coaching and mentoring
- ✓ Public speaking
- ✓ AGILE project management

T E C H N I C A L P R O F I C I E N C I E S

- ❖ Azure Data Enterprise Architecture
- ❖ Relational & non-relational data storage (SQL/No-SQL)
- ❖ Azure Synapse
- ❖ Azure Data Factory
- ❖ Azure Data Lake storage
- ❖ Azure Cosmos DB
- ❖ SQL, KQL
- ❖ PowerBI, Excel
- ❖ Python, C#, JavaScript
- ❖ JSON, XML, Parquet, CSV

E D U C A T I O N

BSc (Hons) Computer Science - 2:1

University of Liverpool

Management & Leadership (Level 5)

CMI

In Progress

A B O U T M E

I have a proven track record of being a strong, pivotal program and team leader who spearheads large scale warehousing and integration projects, guiding teams through complex processes from conception to completion. With a focus on delivering results within tight deadlines, I am adept at navigating challenges while driving continuous improvement.

I am a seasoned Data Developer, Engineer, and Data Platform Manager with a wealth of experience in designing, implementing, and managing modern cloud architectures on Azure. My expertise encompasses a broad range of Azure services, which enable me to build scalable, secure, and cost-effective data and analytics solutions.

W O R K P L A C E E X P E R I E N C E

Data Platform Manager

Oliver James | Manchester, UK | 2021 - Present

- With a strong influence over the technical strategies for Data and Analytics I am a key resource in deciding the direction of tech focussed objectives in the business. My team deliver enterprise scale data architectures while steering the business towards quality data governance and literacy processes.
- My team and I have delivered cutting edge cloud-based distributed data solutions bringing the technical capabilities of the business forward into the modern data-first world.

Founder & Host

SQL_SQUARED | Manchester, UK | 2024 - Present

- I create quality productions (written and video) for the benefit of the global Data community via a personal website, blog, podcast and YouTube. They range from technical deep dives into complex data and AI solutions for experts, to light heart overviews for those new to the industry.
- I aim to grow a well-loved community under the sql_squared brand while evolving my own personal brand and expertise. Sql_squared should give me the platform

SQL Developer & Software Tester

Evolution Recruitment | Warrington, UK | 2018 - 2021

- My primary role was SQL Data Warehouse development for a web-based CRM being created in-house.
 - I was highly recognised for many innovative and ambitious ideas centred on producing easy to use software and developing an enhanced reporting stack.
 - I was a key resource in the business's initial migrations of their data platform to Azure to make use of a more modern stack.
- My secondary role involved gathering business requirements, translating them into technical requirements and testing the results of CRM development.

PROJECT EXPERIENCE OF NOTE

Azure Analytical Data Platform

Technologies: Azure SQL Managed Instance, Azure SQL Databases, Azure Data Factory, Azure Cosmos DB, Azure Logic Apps, Azure Function Apps, Azure Blob Storage, Azure Log Analytics

Languages & Formats: MS SQL, KQL, Python, C#, JSON

I was the lead data engineer on the development of a new cloud first SQL data warehouse to replace legacy on-premises systems. I scoped, architected and developed the warehouse to be highly functional and saleable for the business's analytical needs. It has a distributed design with a flexible ingestion layer, a core entity-based dimensionalised data store and forked data marts to serve our various reporting suites. Data is ingested from many systems in a variety of formats where it is required to be cleansed, matched and stored efficiently.

The warehouse is built as an Azure SQL Managed Instance and orchestrated by Azure Data Factory. Azure SQL Databases were tuned for use as PowerBI data marts while Cosmos DB stored data for our large self-serve web-based application. Azure Log Analytics, Logic Apps and custom-built maintenance pipelines were used to report on the health and stability of the warehouse and it's operations.

Azure Data Lake for Enterprise Master Data Management

Technologies: Azure Synapse, Azure Data Lake Storage, Azure Data Factory, Azure Logic Apps, Azure Function Apps, Azure Log Analytics

Languages: MS SQL, KQL, JSON, Parquet

We built a Data Lake within Azure Synapse Analytics using dedicated SQL pools to act as the front door to the analytical data platform. This was built as an enhancement to the platform to solve the business need for better master data management across all of the business's core functions and to improve the sync time of data to 3rd party systems. The data lake follows your standard medallion architecture with full traceability of data changes for audit purposes. Integrations of a non-analytical nature were migrated to this system with the ability to run off data as early in the pipeline as required.

Reporting on global Telephony Data

Technologies: Azure SQL Managed Instance, Azure Data Lake Storage, Azure Data Factory, Azure Function Apps, Azure Log Analytics

Languages & Formats: MS SQL, C#, JSON, REST, Event Streaming

Oliver James's telephony data from multiple providers around the world is the only data required on a more frequent schedule than all other reporting. We built a call data reporting suite allowing global management to monitor call statistics in near real-time and later implemented automatic linking of contact records in Salesforce, and AI summaries of call transcriptions. To achieve this we implemented real-time event based services to pull all the required call events from the various providers. These events were processed on read into our SQL warehouse structures to be visible in the reports almost immediately.

Warehousing and Reporting on Daily Global LinkedIn Job Market Data

Technologies: Azure Synapse, Azure Data Lake Storage, Azure Data Explorer, Kusto Clusters, Azure Data Factory

Languages: MS SQL, KQL, JSON

I worked closely with an external company that scraped and stored publicly available job information on LinkedIn. We were provided a huge dump of data every night that contained a global unrefined dataset of jobs that were active on the platform. My goal was process and store the data we were interested in as a business to provide and sell market analysis. To achieve the desired output in a cost efficient manner I ingested the data dumps in Azure Data Explorer kusto clusters to refine the data close to source. The result of this process was offloaded to our data lake to follow the standard warehousing implementations we had in place.

Salesforce Data integrations

Technologies: Azure Data Factory

Languages: MS SQL, SOQL

Most sales-based activity was produced within a custom Salesforce environment which required loading into the data platform for analytical reporting across the business. This data and the reports created on the back of it acted as the backbone for both core business defining decisions by upper management and general daily management of consultants' sales targets. I have become very accustomed to the back-end data structures of Salesforce and the functions required to make the most of the data the system produces.

Data was extracted from the system using Salesforce connectors in Azure Data Factoring with custom SOQL written to delta the data loads. The data was dimensionalised and linked within the Azure SQL Managed Instance and the Azure Data Lake depending on it's purpose.

Financial Data integrations with SAP ByDesign

Technologies: Azure Data Factory, SQL Server, Azure App Service

Languages & Formats: MS SQL, C#, XML, SOAP

I was a developer on a complex suite a data integrations required for the successful implementation of SAP ByDesign as the business's ERP platform. Much of the financial processing within the business was done in differing systems that needed to be consolidated. The warehousing and preparation of the data was done on SQL Server 2016 so we could easily interface with the SOAP endpoints of the SAP system. XML payloads were created in the required format to push data to the relevant objects. Mapping of data, the tracking of required change and the management of errors were major hurdles we needed to overcome which required innovative solutions to solve correctly.

Implementation of Azure DevOps and CI/CD for Data Warehouses

Technologies: Azure DevOps, DevOps Pipelines

Languages & Formats: git, YAML

When I joined Oliver James all data and software releases were reviewed and deployed manually outside of core working hours. As the team grew and the frequency of deployment became higher, I identified that proper DevOps both for our databases, APIs and web-based applications was vital to enable efficient feature growth. I introduced Azure DevOps and CI/CD to the business eventually growing a small DevOps function to manage all of the processes effectively. 100% of our software and data resources are now managed in this way allowing us to maintain our codebase easily, deploy and test change in safe UAT environments and release confirmed feature change to production with minimal downtime to the business and minimal manual input from developers.

Improving AGILE development practices

Technologies: JIRA, Confluence, Azure DevOps

When I first joined Oliver James the IT Development function had half baked AGILE and SCRUM processes in place that very easy got overridden by senior members of the business outside of the team. Over multiple years I made many major improvements to the team process to help us protect our workstack. This included:

- Solidifying our sprint process, schedules, planning and ceremonies. This allowed the team to work on a consistent cycle with progress and upcoming work always visible to the business.
- Implementing a development change board to control the tasks coming into the team and to report on the progress of roadmap deliverables at a high level.
- Matured the team's attitude towards change logging and release planning. Change that we were implementing became logged effectively, so we always had a trace of what was changing an why. The business became more aware of planned release schedules improving on the transparency of when they would get requested features.
- Created a centralised documentation repository in Confluence. This contained technical documentation of large scale projects, a knowledge base on our analytical metrics, and details of decisions made so leadership could always be help to account if there were future problems.