

Richard Swinbank

richard@RichardSwinbank.net


Innovative, business-focused Data Platform Engineer with proven, in-depth skills and experience. Microsoft Data Platform MVP, author, client recommendations at <https://tinyurl.com/yc823f8s>.

Key skills

- Azure data platform (including IaC, Terraform) & Microsoft SQL Server
- Azure Data Factory (ADF), SSIS (including BIML, BimlScript)
- Azure Databricks, Azure Synapse Analytics, Microsoft Fabric
- SSAS tabular, AAS, Power BI, semantic modelling
- DevOps toolchain components (git/svn/TFS, Azure Pipelines, PowerShell)
- Expert T-SQL for queries, analysis, administration and process engineering
- Configuration/metadata-driven dynamic T-SQL
- C# (including SQL CLR assemblies, SSIS script tasks/components, BimlScript), Java
- Requirements-driven data modelling and warehouse/lakehouse development
- Effective communicator, able to convey complex technical issues to a range of audiences. Blog <https://richardswinbank.net>, X/Twitter <https://twitter.com/richardswinbank>.
- Data-driven, analytical problem solver
- Active in Microsoft Data Platform community: Organiser, Volunteer, Speaker (Data Relay, SQLBits)

Achievements

- Author of *Azure Data Factory by Example* (<https://www.amazon.com/dp/B096Z9JW86>)
- [Microsoft Data Platform MVP](#)
- Microsoft Certified Data Engineer Associate DP-203
- Built greenfield cloud-native, infrastructure-as-code analytics platform, enabling business to interrogate and understand performance of product features for B2B clients.
- Implemented value (ETL) and innovation (CI/CD) pipelines with high levels of automation and resilience using git (Azure DevOps Repos, Bitbucket, GitHub) and Azure DevOps Pipelines.
- Recruited high-performing, agile team of Data Engineers & BI Analysts, operating with high degrees of trust, collaboration & psychological safety.
- Built metadata-driven extract and staging framework, supporting rapid, repeatable implementation and maintenance of platform staging objects.
- Built ETL process controller for data platform in Microsoft Azure using Azure Durable Functions (C#)
- Designed multi-tenant data warehouse platform, allowing client to process multiple customers' data in a single instance, reducing administration overhead and customer cost.
- Enabled end-user master data management using Master Data Services (MDS) integrated into ETL.
- Built configurable master data management framework in T-SQL, allowing multiple source systems to be integrated into operational data store's global data dictionary (in the absence of MDS).
- Re-engineered existing data warehouse ETL in situ, replacing legacy corporate system sources with new ones while allowing valuable reporting assets to remain in use.
- Author of Sprockit (<http://sprockit.info>) – a free, open-source ETL process controller built in pure T-SQL with adaptive end-to-end ETL duration reduction and one-step resume after error.
- Contributor to and cited by Microsoft Learn Azure platform documentation



Richard Swinbank

richard@RichardSwinbank.net

Experience

May 2022 – November 2023 (permanent role)

Senior Data Platform Engineer – Avanade

Technical leadership for a range of projects for national and multinational clients

- Lift & shift migration of on-premises analytics platform into Azure, rebuilding, reusing and reorchestrating components for enhanced resilience and improved performance (Azure SQL, Azure Data Factory including SSIS-IR, Key Vault, Terraform).
- Reverse-engineered undocumented metadata-driven legacy platform, enabling reimplementations.
- Refactored existing data transformation processes, reducing execution times by 50-70% (T-SQL).
- Implemented cloud-native integration of heterogeneous mixed cloud/on-premises source systems (Azure Data Factory including managed VNet IRs).
- Operated highly automated data lakehouse using metadata-driven ingestion and transformation processes, including lifecycle management for personal data (Databricks, Azure Data Factory, Power BI).
- Empowered and led mixed (Avanade/client) engineering teams in agile, self-organised delivery of data analytics platforms.
- Transformed business requirements into deliverable work items, facilitated technical refinement.
- Mentored and advised junior engineers in processes, patterns, and technologies.

July 2020 – May 2022 (permanent role)

Senior Data Engineer, Head of Data – Boomin


Construction of greenfield cloud-native enterprise analytics platform in Microsoft Azure. Recruitment and leadership of data platform team.

- Built enterprise analytics platform comprising Azure Data Lake Gen2, Azure SQL DB, Azure Analysis Services & Power BI
- Automated platform construction and development using Terraform, Azure DevOps YAML pipelines and a range of deployment tools (SSDT, YAML templates, PowerShell scripts)
- Implemented metadata-driven process orchestration framework using ADF
- Built metadata-driven extract and staging framework for rapid staging process development
- Authored Azure Function Apps to encapsulate non-standard platform processes (e.g. Google OAuth token acquisition, Firebase user extracts)
- Automated and version-controlled platform configuration (including user access) using JSON metadata serialisation and custom-built Azure pipeline tasks
- Improved developer experience using PowerShell scripts and git hooks
- Recruited and led high-performing team of Data Engineers & BI Analysts, nurtured a culture of collaboration, trust & psychological safety
- Managed complex backlog of business requirements, writing user stories & guiding task refinement

September 2019 – February 2020 (one renewal)

Data Engineer – Oakbrook Finance

Implementation of next-generation enterprise data platform in Microsoft Azure.



Richard Swinbank

richard@RichardSwinbank.net

- Designed and built ETL process controller using Azure Durable Functions (C#), to manage execution of Azure Data Factory (ADF) pipelines and other processes
- Designed and built metadata transformation library, enabling ETL metadata to be both stored under version control as JSON, and imported to (and re-exported from) Azure SQL DB
- Built proof-of-concept Azure DevOps YAML pipelines to create disposable development environments in Azure using Terraform, SSDT deployments and metadata deserialisation
- Automated inference of normalised staging schemas from JSON source data in Azure blob storage
- Implemented extensive test suites for SQL objects (TSQL, tSQLt) and ETL process controller (C#, NUnit, Moq, FluentAssertions)
- Automated functional tests for ADF pipelines using NUnit and Azure DevOps YAML pipelines

January 2018 – September 2019 (four renewals)

Data Warehouse Developer – Dyson

Stabilisation and productionisation of SQL Server 2016 data warehouse platform.

- Instituted robust, well-documented deployment pipeline and change control processes
- Replaced centralised TFS-based version control with git feature branch and one-step deploy
- Refactored SSIS & TSQL ETL components, reducing nightly processing time by 60%
- Implemented SSRS dashboard and tools for routine ETL monitoring and management
- Extended Kimball data model to keep pace with fast-evolving reporting requirements
- Navigated end-user communication to elucidate data model designs from raw data requests
- Delegated management of integration controls to end users using Master Data Services
- Integrated data from heterogeneous source systems including Teamcenter, MS Project Server, Active Directory (LDAP queries), SharePoint (CSOM API), JIRA (REST API) & Exchange (EWS API)
- Built API extract interfaces using SSIS script source components and JSON shredding
- Streamlined DW integration by calling Tableau extracts directly from ETL via REST API

November 2017 – January 2018

DW/BI Consultant – Counter Intelligence Retail

Reimplementation and automation of SQL Server 2016 sales reporting platform to release analyst and developer time, permit integration with corporate master data and facilitate report suite redevelopment.


- Abstracted complex legacy reporting model into simplified star schema
- Prepared model for replacement of legacy dimensions with corporately-managed master data
- Automated data integration from a variety of source system Excel and flat file outputs using SSIS, saving 5+ analyst days per month
- Populated new model from historical data (T-SQL) and new sources (SSIS), demonstrating consistency with existing report outputs and ensuring continued availability of legacy BI assets

August 2016 – August 2017 (one renewal)

Data Warehouse Consultant – North Staffordshire Combined Healthcare NHS Trust

Integration of data from new Lorenzo patient system into existing SQL Server warehouse platform, enabling seamless reporting across source systems and retention of value in existing reporting assets.

- Built SSIS packages for flat file ETL
- Designed and built ETL metadata repository using SSIS to sync supplier metadata



Richard Swinbank

richard@RichardSwinbank.net

- Used BimlScript to generate 850 SSIS packages to standard ETL pattern from metadata
- Automated T-SQL build of staging objects from metadata, supporting schema upgrades in-situ
- Implemented metadata-driven ETL application in C#
- Built searchable end-user view of historical system data using SSRS and legacy Oracle back end
- Developed master patient record in operational data store, automatically combining records from different sources

March 2015 – July 2016 (three renewals)

ETL Developer – Staffordshire & Shropshire Health Informatics Service

Incorporation of new acute hospital tenant into multi-tenant SQL Server warehouse platform.

- Acquired and refined reporting requirements from known NHS data sets and by liaising with customer stakeholders from a variety of professional backgrounds (technical, clinical, non-clinical)
- Integrated data from numerous source systems in a variety of technologies (SQL Server, PostgreSQL, SQLBase, mixed-type CSV files) using T-SQL & SSIS
- Designed integration transformations from a variety of source systems and data models
- Implemented library of CLR functions and stored procedures to perform specialised tasks in C# (e.g. INSERT/EXEC nesting, regular expression matching, inline transformation to 1NF)

October 2014 – March 2015 (one renewal)

BI Developer – Birmingham & Solihull Mental Health NHS Foundation Trust

Implementation and extension of new SQL Server reporting data model.

May 2013 – September 2014 (six renewals)

Data Warehouse Developer – Staffordshire & Shropshire Health Informatics Service

Architecture and development of shared, multi-tenant SQL Server data warehouse platform, hosted by a shared service for the use of multiple customer organisations.

Education

PhD, Computer Science

Virtual forced splitting in multidimensional access methods (<http://etheses.bham.ac.uk/213/>)

University of Birmingham, 2008

MSc, Computer Science

Distinction

University of Birmingham, 2001

