



Nicholas Westenberg

Senior Software and Artificial Intelligence Engineer

Burgemeester de Raadsingel 48, 3311 JG Dordrecht, The Netherlands

✉ nicholas.westenberg@techsphere.nl | 🌐 www.techsphere.nl | 📧 [@nickwestenberg](https://twitter.com/nickwestenberg) | 📺 [nicholaswestenberg](https://www.youtube.com/nicholaswestenberg)

Summary

As a Software Engineer and Artificial Intelligence Consultant, I bring extensive experience in designing scalable software solutions, cloud-native applications, and intelligent data-driven systems across modern technology platforms. I value professionalism, collaboration, innovation, and continuous learning, while adapting quickly to emerging technologies and evolving industry requirements.

I am currently expanding my expertise in Generative Artificial Intelligence, Large Language Models (LLMs), and Intelligent Automation. My recent work focuses on Retrieval-Augmented Generation (RAG), vector databases, AI agents, MLOps, and cloud-based machine learning pipelines. Professional certifications and hands-on experience with modern AI frameworks, cloud services, and data engineering technologies continue to strengthen my ability to deliver scalable and production-ready intelligent solutions.

Projects

NorthRiver Analytics

AI platform modernization on Azure – Data Engineer

2025

Goal: Consolidate fragmented analytics workloads and design/build a modern, scalable data platform on Microsoft Azure.

Description: NorthRiver Analytics initiated a broader data-maturity program to centralize operational data, modernize reporting pipelines, and support future Artificial Intelligence use cases. The selected architecture combined Azure Data Lake Storage, Azure Synapse Analytics, and Power BI as the primary data and reporting platform. The initial project phase focused on consolidating data from several business applications, engineering a reliable synchronization process from an on-premises PostgreSQL environment to Azure, implementing production-ready data pipelines with Databricks and Airflow, and defining an intermediate data strategy and technical learning roadmap for the internal data team.

Key Contributions:

- Assessed the organization's existing knowledge, tools, reporting requirements, and long-term platform objectives. Produced an intermediate architecture recommendation while simultaneously delivering the first phase of the data-platform implementation.
- Developed Python-based ingestion processes using JDBC and Parquet files to extract relational data from PostgreSQL and load it into Azure Data Lake Storage and Synapse Analytics. Resolved schema compatibility, data-type conversion, incremental loading, and pipeline-execution challenges across heterogeneous source systems.
- Guided the customer in adopting dimensional data modeling, automated testing, and transformation workflows using DBT.
- Designed reusable data marts and semantic models for executive dashboards and operational reporting in Microsoft Power BI.

Tools used: Terraform, Docker, Azure Synapse Analytics, Azure Data Lake Storage, Data Warehousing, Data Governance, Azure Key Vault, Azure DevOps Pipelines, Git CLI, Apache Airflow, DBT Core, Databricks, PySpark, SQL, and Python.

NovaTech Solutions

Cloud-native AI platform development and DevOps – Engineer

2024

Goal: Develop a scalable cloud-native platform to support Artificial Intelligence applications and modernize legacy deployment processes using automated DevOps practices.

Description: NovaTech Solutions launched a strategic initiative to redesign its software platform for intelligent business applications. Working within an international Agile team, the project focused on building a secure serverless architecture, automating software delivery pipelines, and supporting rapid feature deployment. Responsibilities included participating in Scrum ceremonies, infrastructure design decisions, CI/CD implementation, and deployment automation while collaborating closely with software developers and solution architects.

Key Contributions:

- Designed and implemented multi-stage CI/CD pipelines using GitHub Actions and Infrastructure as Code, enabling automated testing, validation, and deployment across multiple cloud environments.
- Established automated deployment notifications and monitoring integrations to provide rapid feedback to development teams, improving software quality and reducing deployment failures.
- Integrated automated unit testing, integration testing, security scanning, and code quality analysis into the software delivery pipeline.
- Investigated deployment bottlenecks, optimized build performance, and improved infrastructure provisioning through reusable Infrastructure as Code templates.
- Developed and deployed event-driven cloud services using containerized microservices, serverless functions, REST APIs, object storage, and managed databases to support modern AI-powered applications.

Tools used: GitHub Actions, Docker, Kubernetes, Terraform, Azure DevOps, Azure Functions, Azure Storage, Azure API Management, PostgreSQL, Redis, Python, TypeScript, Node.js, and Git.

MedCore Health Services

Secure cloud platform for emergency response – Cloud Architect/Engineer

2023

Goal: Design and implement a highly available cloud solution to provide healthcare professionals with secure access to critical patient information during emergency situations.

Description: Served as the Cloud Architect and Engineer for a mission-critical healthcare modernization project focused on business continuity and disaster recovery. The objective was to build a resilient cloud platform capable of securely synchronizing and delivering essential electronic medical records whenever the primary clinical information system became unavailable. Working closely with healthcare professionals, infrastructure engineers, and security specialists, I designed and implemented a secure event-driven architecture that satisfied stringent healthcare compliance and availability requirements.

Requirements in the project:

- Designed a highly available, event-driven cloud architecture capable of securely processing sensitive healthcare information while complying with international information security standards.
- Implemented secure automated synchronization and controlled distribution of medical documents to authorized healthcare professionals through encrypted cloud storage services.

Results / Challenges:

- Successfully deployed a secure serverless cloud platform that continuously synchronized patient information and ensured rapid access to critical medical records during service disruptions.
- Developed automated file-transfer workflows, monitoring processes, operational procedures, and technical documentation to support reliable business continuity operations across multiple departments.
- Resolved challenges related to event-driven processing, lifecycle management, storage optimization, security controls, and operational cost management while maintaining high system reliability.

Tools used: Amazon S3, AWS Lambda, Amazon EventBridge, AWS Step Functions, AWS DataSync, AWS Transfer Family, Amazon FSx, AWS KMS, CloudWatch Logs, IAM, Terraform, Python, and Git.

GGZ InGeest (Healthcare)

Build a new AWS data platform – Platform/Data Engineer

2023

Goal: Build the basis for a new Data Platform on AWS and empower the Data Engineering team to explore and determine what is efficient and possible when using AWS as a data platform.

Description: GGZ InGeest is a mental healthcare provider operating across 13 locations in the wider Amsterdam region. The organization requested leveraging AWS as their data platform, with Amazon Redshift as the central data warehouse. They sought end-to-end guidance throughout the process.

As a Consultant/Engineer, advised the customer about the possibilities of using AWS and positioned the Glue Data (Hive) Catalog as the central component of their data platform together with AWS Lake Formation. This included an assessment of the impact and usage of their AWS Landing Zone and guidance on the usage and estimated costs of Amazon Redshift (including Spectrum) and other related services.

Challenges:

- Proprietary database Progress (OpenEdge) used by MijnQuarant is not natively supported by AWS Glue Crawlers or Glue Job bookmarks.
- Change Data Capture increments were not directly accessible on tables (via SQL logging and/or timestamps). Solved this by computing differences using temporary SQL views.
- Speed of computing differences using SQL views was not ideal but enabled incremental replication and maintained dataset consistency.
- Cost considerations around Amazon Redshift.

Results / Facts:

- Engineered a solution for replicating the primary database (MijnQuarant) to AWS (220 GB healthcare data) and created a starting point for copying other data services via Glue Jobs/JDBC to S3.
- Set up a Glue Data Catalog (Crawler via S3) enabling faster loading and interoperability between AWS services.

Tools used: CloudFormation, S3, Parquet, DataFrames (Pandas), Glue (Data Catalog, Crawler, Studio), ECS Fargate, Data Hub, Lake Formation, Athena, Redshift (Spectrum). Language used: Python.

Vertex Digital Solutions

Modernize production database platform on Azure – Cloud Architect/Engineer

2025

Goal: Guide and support the migration of a legacy on-premises PostgreSQL database to a fully managed Azure database platform with improved scalability, security, and operational efficiency.

Description: Led the architecture, planning, and execution of a mission-critical database modernization project for a rapidly growing Software-as-a-Service (SaaS) company delivering enterprise workflow and digital business solutions. The platform serves customers across multiple countries, processing thousands of daily business transactions while supporting high-volume application workloads hosted on Kubernetes. Responsibilities included solution architecture, migration planning, performance validation, risk assessment, disaster recovery design, and production deployment coordination.

Key Contributions:

- Designed and documented a comprehensive migration runbook covering rollback procedures, failover strategies, disaster recovery scenarios, and business continuity requirements aligned with defined RPO/RTO objectives.
- Coordinated migration validation activities, production readiness assessments, database verification procedures, and deployment time windows with engineering and operations teams.
- Minimized production downtime by executing multiple rehearsal migrations within staging environments before the final production cutover.
- Resolved data migration challenges involving schema compatibility, replication strategies, performance tuning, and database optimization throughout the migration lifecycle.
- Provided architectural guidance on Azure Database for PostgreSQL, Azure SQL Database, migration services, high-availability configurations, monitoring, and infrastructure cost optimization.

Results:

- Successfully completed the production migration with minimal downtime and no measurable data loss.
- Reduced infrastructure and operational costs while improving database performance and long-term maintainability.
- Enhanced platform scalability, reliability, and disaster recovery capabilities through managed cloud database services.
- Strengthened operational confidence through detailed documentation, migration planning, and technical leadership throughout the project lifecycle.

Tools used: Azure Database for PostgreSQL, Azure SQL Database, Azure Database Migration Service, Azure Monitor, Kubernetes, Terraform, Docker, PostgreSQL, SQL, and Git.

Zwijssen (Publishing)

Build and migrate Web Stack to AWS – Engineer

2024

Goal: Create a highly available and scalable architecture on AWS to replace the on-premises environment, future-proof the application stack, enable on-demand pricing and improve operational efficiency.

Description: As an Engineer was responsible for designing, deploying, testing and guiding the customer in the migration to their new AWS platform. The customer's education stack (Backend/API based on ASP.NET served by Microsoft IIS web servers) needed to scale up to 40,000 concurrent connections/students across approximately 80 public domains during peak hours. The project was successfully launched during the summer and significantly improved the platform capabilities.

Results:

- Developed multiple CDK stacks to improve maintainability and continuous improvement of the target environment, separated into Acceptance, Production and Domains/SSL certificates.
- Implemented an AWS Image Builder mechanism to patch AMIs outside peak hours, reducing human intervention and operational risk while automating updates for ELB environments.
- Provided training for executing application updates and cloning AMI images to the new AWS platform.
- Guided the customer on integrating their stack into their operational workflow (WoW), including handling scaling challenges, CloudWatch metrics, alerts and improving their security posture.

Tools used: CDK (IaC), Elastic Load Balancers (ELB), RDS Microsoft SQL, EC2, Windows, AWS Networking, AWS Image Builder, CloudWatch, SSM Automation Documents, AWS Lambda, KMS Encryption, ACM Certificate Manager, Route53. Languages used: Python and Typescript.

Work Experience

Xebia / Oblivion

Cloud Architect / Cloud Engineer

2019 – 2025

During the 6 years at Xebia/Oblivion, supported numerous customers with Cloud Architecture, Engineering, CI/CD, (AWS) Landing Zones, Cloud Journey's, this includes Cloud and Network migrations. Am well known with the Well Architected Framework and executed WAR Reviews and Cloud Cost optimization.

- Create CI/CD pipelines to deploy/deliver secure, resilient, and scalable -versioned- infrastructure and code.
 - Architect, engineer and deliver Cloud Native (serverless) and Event-Driven and serverless solutions with effective monitoring.
 - Design, develop, and implement data workflows and strategies to integrate legacy (DB) systems with modern cloud storage.
- Contributed (plus creation of) to the Managed Service dept. – created project estimates, either helped customers in support, which includes CI/CD, infrastructure as code, networking, security- (iam, events/posture) processing and monitoring, mostly from the start in 2019 (when I joined).

In 2023 Xebia acquired Oblivion Cloud Control. In 2018, Amazon Web Services recognized Oblivion as the first local AWS Premier Consulting Partner in the Benelux and have delivered many successful projects for clients as (E.g.) ABN-AMRO, Aegon, Wehkamp, KNMI, MN-Services, GGZ-InGeest, Provincie Noord-Holland, PostNL, Pathe, Brill, Zwijsen and DSM.

Ticketscan – Alkmaar

Software/DevOps Engineer (Part-time, shareholder)

2014 – 2019

Ticketscan started in the early days of (online) event ticketing/entrance management software. I joined the tech-startup at a later stage, gained much Web 2.0, API-backend and DevOps experience. In its peak the transaction-based platform delivered almost 100.000 event -and festival tickets annually. Worked with many Atlassian developer products. Sold the customer base to competitor Eventix after marketsegment saturated and margins dropped.

- Worked on different tech (and business) projects which led to faster delivery of new platform functionality while enhancing (external) developer velocity.
- Introduced a new WoW using DevOps principles, Git and CI/CD pipelines with server provisioning via Ansible to DTAP environments. Big new code (and platform) versions were released to production using a blue-green deployment methodology (manually) with a Linode Node/LoadBalancer.
- Introduced versioning to all products, migrated codebases (C# Desktop-Windows mobile (IOT) and web API back/frontend) to Git repositories. Used Jira to track dev. progress / product roadmap.
- Created a cross-platform -Company IP isolated- local dev environment for ext. Developers using (e.g.) Ansible, Vagrant, VirtualBox and later Docker.
- Consolidated and migrated the database backend from Akamai/Linode to AWS.
- Co-created new responsive web shop for IOS/Android and Implemented Buckaroo payments provider.
- Implemented Google Tag Manager (GTM)/Ad-words in the multi-tenant webshop and portal backend.
- Advised customers in marketing automation in Google ad-words/analytics and Meta inc. campaigns.

Tools used: JavaScript, jQuery, Bootstrap, CodeIgniter Framework (PHP) - Incl API/Routing, C#, Bash, Ansible, Git, Atlassian Bitbucket (Git)/Bitbucket Pipelines, VirtualBox, Docker/Docker Hub, Linode Load Balancers / VMs, Debian, MySQL BinLog Replication, MySQL Aurora Serverless, S3, Kanban, Confluence/Jira.

Education

Delft University of Technology

Professional Certificate in Cloud Architecture

2024

University of Twente

Computer Science (B.Sc.)

2018

Technical Skills

Languages	Typescript, JavaScript, Python, Bash, ANSI SQL
Databases	Microsoft SQL, Oracle, MySQL, Redshift, DynamoDB, Big Query, Databricks
Developer Tools	MacOS, Linux, Windows, VIM, VS Code, DevContainers, Git, GitHub Actions, Jira
Technologies / Frameworks	Docker, Kubernetes (auto-mode), Data Lake House Architecture, DBT, Airflow
Infrastructure as Code	CloudFormation, CDK, Ansible, SST (Pulumi), Terraform
Methodologies / Concepts	Scrum, ITIL, CI/CD, Event-driven Architecture, Design / Enterprise Patterns

Certifications and Training

Microsoft Azure

- Microsoft Certified: Azure Solutions Architect Expert
- Microsoft Certified: Azure DevOps Engineer Expert
- Microsoft Certified: Azure Administrator Associate
- Microsoft Certified: Azure Developer Associate
- Microsoft Certified: Azure AI Engineer Associate
- Microsoft Certified: Azure Fundamentals

Databricks

- Databricks Data Engineer Professional
- Databricks Data Engineer Associate
- Databricks Spark Developer Associate

Others

- Astronomer Apache Airflow 3 Fundamentals
- HashiCorp Terraform Associate