

Powered by Claude · GPT-4o · Gemini 1.5

Bridge Your Business Problem to the Right AI Solution

A multi-agent platform that interviews you, classifies your domain, researches the best tech stack, and generates a complete architecture + implementation roadmap — in minutes.

Start Discovery Session →

View Dashboard

BridgeAI Discovery Session

Welcome! I'm here to help you find the right data & AI solution. Can you describe the core business problem you're trying to solve?

We have sales data in multiple spreadsheets. We need real-time dashboards and want to predict which customers will churn.

Great context! That's a combination of Data Analytics + ML. How much data are we talking about — rows per day and total historical size?

...

Choose AI Model



Select provider and model for this session

- Google Gemini**
- Anthropic Claude
- OpenAI

Google API key

Stored in this browser and sent only when this provider is used.

Key required

Research API keys

Optional. Stored in this browser and sent only for web-enhanced research.

Taxonomy only

Brave Search

Exa

Firecrawl

Jina

Gemini 3.1 Pro Preview

Latest Gemini 3.1 Pro preview for multimodal reasoning, agents, and coding

Powerful

2M tokens

Gemini 3 Flash Preview DEFAULT

Frontier Gemini speed model with strong reasoning, search, and grounding

Balanced

2M tokens

Gemini 3.1 Flash-Lite Preview

Newest cost-efficient Gemini model for high-volume agentic tasks

Fast

1M tokens

Gemini 2.5 Flash

Current 2.5 Flash model for low-latency reasoning workloads

Fast

1M tokens

Gemini 2.5 Flash-Lite

Smallest 2.5 model for cost-efficient high-throughput workloads

Lite

1M tokens

Selected: Google Gemini — Gemini 3 Flash Preview

PIPELINE

Discovery

Classification

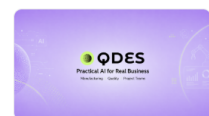
Research

Architecture

Roadmap

All complete!

View Architecture & Roadmap



Practical AI for manufacturing, quality, and project teams.



info@qdes.org

Denmark company - CVR 42351830

BridgeAI Discovery Agent

Online

View Results -->

We run 14 production lines across 3 factories. Every line generates sensor data — temperature, vibration, pressure, cycle counts — but it all sits in separate PLCs and SCADA systems that don't talk to each other. We have unplanned downtime roughly 8-10 times per month and we have no way to predict when a machine is about to fail. Our maintenance team still works off paper checklists. Success would be knowing 48 hours in advance that a specific machine needs attention, so we can schedule maintenance without stopping production.

That is a very clear and compelling picture of the challenge. Moving from paper checklists and 8-10 unplanned outages a month to a 48-hour predictive window would be a massive shift for your operations, likely saving significant revenue and stress. It sounds like the raw "Ingredients" (the sensor data) exist, but they are currently locked in silos.
To help me understand the urgency and the boundaries we're working within, what is your ideal "timeline" for getting a pilot or a first version of this predictive system up and running?

real time

I hear you—to catch those failures before they happen, the system needs to react as fast as the machines do.
Since we are talking about connecting factory floor hardware (PLCs and SCADA) to a centralized predictive system, are there specific "compliance or security requirements" we need to navigate? For example, does the data need to stay strictly on-premises, or is your organization comfortable with using a cloud platform like AWS or Azure for the analytics?

cloud ok

Discovery complete!
I now have a clear picture of your challenge. Let me summarise what I've learned:
Problem: 14 production lines across 3 factories generate sensor data that is siloed in separate PLCs and SCADA systems, preventing any centralized visibility or predictive capabilities.
Business Impact: Frequent unplanned downtime (8-10 times per month) and a reliance on manual, paper-based maintenance checklists, leading to operational inefficiency and lost production time.
Expected Output: 48-hour advance notification of machine failure, Centralized real-time monitoring of sensor data, Predictive maintenance alerts for specific machines
I'm now classifying your domain and researching the best technology stack for your needs. This will take a moment — I'll show you the architecture and roadmap shortly!

Discovery complete!
Your architecture and roadmap have been generated.

View Results -->

Manufacturing / IoT

Session #0d3664af •  Complete

Executive Summary

Architecture

Roadmap

Technology Stack



16

Total Weeks



83

Person-Days



5

Layers



7

Technologies



95%

Confidence

Problem Statement

14 production lines across 3 factories generate sensor data that is siloed in separate PLCs and SCADA systems, preventing any centralized visibility or predictive capabilities.

Business Impact

Frequent unplanned downtime (8-10 times per month) and a reliance on manual, paper-based maintenance checklists, leading to operational inefficiency and lost production time.

Domain Classification

PRIMARY DOMAIN

machine_learning

SUB-DOMAINS

anomaly_detection

time_series

streaming_ingestion

real_time_dashboards

The core objective is to transition to predictive maintenance (Machine Learning) using time-series sensor data. Achieving this requires breaking data silos (Data Engineering) and implementing real-time monitoring and alerting (Real-time Streaming/Analytics).

Recommended Approach

Feature Store + ML Training/Serving

This pattern was selected because the primary objective is predictive maintenance (anomaly detection) which requires complex time-series feature engineering (rolling windows, aggregations) and real-time inference. A Feature Store (Tecton) ensures consistency between training and serving, while the serving layer (BentoML) provides the low-latency predictions needed for 48-hour advance warnings.

Feature Store + ML Training/Serving

This pattern was selected because the primary objective is predictive maintenance (anomaly detection) which requires complex time-series feature engineering (rolling windows, aggregations) and real-time inference. A Feature Store (Tecton) ensures consistency between training and serving, while the serving layer (BentoML) provides the low-latency predictions needed for 48-hour advance warnings.

Layers

Data Ingestion & Validation Layer

Capture high-frequency sensor data from PLCs and SCADA systems and ensure data quality before processing.

Azure Event Hubs

Microsoft • streaming

Real-time ingestion of sensor data from 14 production lines across 3 factories.

Great Expectations

Great Expectations • quality

Automated data validation for sensor streams.

➡ In: PLC Sensors, SCADA Systems ➡ Out: Storage Layer, Feature Store (Online)

Storage & Historical Lake Layer

Provide a persistent, ACID-compliant repository for historical sensor data to support model training and retrospective analysis.

Azure Data Lake Storage Gen2

Microsoft • storage

Primary storage for raw and processed time-series data.

Apache Iceberg

Apache • storage

High-performance table format for the data lake to support real-time dashboards.

➡ In: Data Ingestion Layer ➡ Out: Feature Store (Offline), Real-time Dashboards

Feature Engineering Layer

Transform raw sensor streams into ML-ready features (e.g., 1-hour rolling average temperature) for both training and real-time inference.

Tecton

Tecton • feature_store

Managing and serving real-time features for anomaly detection.

➡ In: Azure Event Hubs (Online), Apache Iceberg (Offline) ➡ Out: ML Training Layer, Inference Layer

ML Development & Training Layer

Develop, train, and version anomaly detection models using a low-code approach suitable for the existing team skills.

16 weeks

Total Duration

83 days

Total Effort

4

Phases

8

Critical Tasks

Team Requirements**1.5** Data Engineer**1** ML Engineer**0.5** BI Developer**Timeline****Phase 1: Foundation & Connectivity**

Weeks 1-4 • Establish cloud landing zone and break data silos by connecting PLC/SCADA systems to Azure Event Hubs.

🔗 Data ingestion established for pilot production lines.

● Azure Infrastructure Provisioning 5d • Data Engineer • Infrastructure

● PLC/SCADA Ingestion Pipeline 10d • Data Engineer • Data Ingestion & Validation Layer

▸ Data Quality Guardrails 4d • Data Engineer • Data Ingestion & Validation Layer

Phase 2: Storage & Feature Engineering

Weeks 5-8 • Build the historical data lake and define the time-series features required for predictive maintenance.

🔗 Feature Store operational with historical sensor data.

● Iceberg Table Implementation 6d • Data Engineer • Storage & Historical Lake Layer

● Tecton Feature Store Setup 8d • ML Engineer • Feature Engineering Layer

▸ Historical Data Backfill 5d • Data Engineer • Storage & Historical Lake Layer

Phase 3: ML Development & Serving

Weeks 9-12 • Train anomaly detection models using low-code tools and deploy them for real-time inference.

🔗 Predictive model deployed and validated for lead-time accuracy.

● Low-Code Model Training 10d • ML Engineer • ML Development & Training Layer

● BentoML Inference Service 7d • ML Engineer • Inference & Alerting Layer

▸ 48-Hour Lead Time Validation 5d • ML Engineer • ML Development & Training Layer

Azure Event Hubs

Data Ingestion & Validation Layer

Microsoft • streaming

Real-time ingestion of sensor data from 14 production lines across 3 factories.

Score **0.98/10**

Alternatives

AWS Kinesis

Apache Kafka

🔥 consumption

Great Expectations

Data Ingestion & Validation Layer

Great Expectations • quality

Automated data validation for sensor streams.

Score **0.85/10**

Alternatives

Soda Core

Monte Carlo

🔥 open-source

Azure Data Lake Storage Gen2

Storage & Historical Lake Layer

Microsoft • storage

Primary storage for raw and processed time-series data.

Score **0.95/10**

Alternatives

AWS S3

Google Cloud Storage

🔥 consumption

Apache Iceberg

Storage & Historical Lake Layer

Apache • storage

High-performance table format for the data lake to support real-time dashboards.

Score **0.88/10**

Alternatives

Apache Hudi

🔥 open-source

Tecton

Feature Engineering Layer

Tecton • feature_store

Managing and serving real-time features for anomaly detection.

Score **0.92/10**

Alternatives

Feast

🔥 consumption

Azure ML

ML Development & Training Layer

Microsoft • training

Managed environment for developing and training time-series anomaly detection models.

Score **0.97/10**

Alternatives

AWS SageMaker

Google Vertex AI

🔥 consumption

BentoML

Inference & Alerting Layer

BentoML • serving

Packaging and deploying models for real-time inference.

Score **0.9/10**

Alternatives

FastAPI

Seldon Core

🔥 open-source

🚀 Quick Start

BridgeAI turns any business problem into a complete AI architecture and step-by-step implementation roadmap — no technical expertise needed.

01 Click New Session

Pick your AI model and provider. Google Gemini is recommended for getting started — it has a generous free tier.

02 Describe your problem

Explain your business problem in plain language. The more detail you give — data volumes, existing systems, team size, success criteria — the better the output.

03 Answer the AI's questions

BridgeAI conducts a structured discovery interview. Each question digs deeper to classify your domain and understand constraints.

04 Research phase runs automatically

Specialized agents search the web using Brave Search, Exa, Firecrawl, and Jina Reader to find the best tools for your specific use case.

05 Receive your blueprint

You get a full technology architecture with rationale, component decisions, and a phased implementation roadmap with effort and cost estimates.

🤖 AI Model API Keys

BridgeAI uses **your own API keys** — your usage, your billing, full control. Keys are stored only in your browser and sent directly to the AI provider. They are never stored on our servers.

To add keys: click the [model badge](#) in the sidebar → enter your key → click Activate.

🌟 Google Gemini

Starts with AIza...

Get key →

Free tier with generous limits. Best for getting started. Gemini 2.5 Flash is fast and cost-effective. Gemini 3 Flash Preview offers the strongest reasoning.

🔍 Research Tool API Keys

During the **Research phase**, BridgeAI agents autonomously search the web to find the best tools, vendors, and approaches for your specific problem. These keys are **optional** — without them the app still works using the AI's built-in knowledge, but live web research produces significantly better and more up-to-date architecture recommendations.

Add research keys in the same place as AI model keys: click the [model badge](#) in the sidebar.



Brave Search

Primary web search engine

2,000 queries/month free

Get key →

Starts with BSA...

Brave Search has an independent search index (not relying on Google or Bing), making it ideal for fresh, unbiased tech research. It leads AI agent benchmarks with a 14.89 Agent Score and the lowest average latency at 669ms. Used by the Tech Research Agent to find current tools, pricing pages, and technology comparisons.



Exa AI

Semantic / neural search

1,000 searches/month free

Get key →

Starts with exa...

Exa uses neural embeddings to understand the meaning of a query rather than matching keywords. This is critical for the discovery phase of deep research — finding conceptually similar tools and architectures even when they use different terminology. Particularly powerful for GitHub repository discovery and finding open-source alternatives.



Firecrawl

Deep web crawling & structured extraction

500 credits/month free

Get key →

Starts with fc...

Firecrawl crawls entire documentation sites, wikis, and product pages and converts them to clean Markdown or structured JSON for the AI to analyse. It powers the structured comparison feature — when BridgeAI needs to compare 5 tools across 10 attributes, Firecrawl extracts the data from each vendor's site in a single pass. Also used for building internal knowledge bases from docs.



Jina Reader

Instant URL-to-Markdown conversion

1M tokens/month free

Get key →

Starts with jina...

Jina Reader converts any URL to clean, AI-readable Markdown in a single API call with no setup. It's used for quick fact verification — when the AI needs to check a pricing page, a changelog, or a specific documentation section. Much faster than Firecrawl for single-page lookups and requires no credits for basic use.

Recommended setup: Add at least **Brave Search + Jina Reader** for good research coverage. Add **Exa** for deeper semantic discovery and **Firecrawl** if you want structured multi-tool comparisons extracted from vendor websites.

Predictive maintenance

Manufacturing / IoT

We run 14 production lines across 3 factories. Every line generates sensor data — temperature, vibration, pressure, cycle counts — but it all sits in separate PLCs and SCADA systems that don't talk to each other. We have unplanned downtime roughly 8-10 times per month and we have no way to predict when a machine is about to fail. Our maintenance team still works off paper checklists. Success would be knowing 48 hours in advance that a specific machine needs attention, so we can schedule maintenance without stopping production.

Try example

Churn prediction

Retail / Customer Analytics

We have 6 years of transaction data in our ERP, a loyalty programme in a separate CRM, and our e-commerce clickstream going into Google Analytics. Nobody has ever connected these three. Our marketing team is sending the same promotions to everyone and our churn rate has gone up 18% in the last year. We think we're losing our best customers but we genuinely don't know who they are or why they're leaving. Success is being able to identify at-risk customers before they churn and trigger personalised retention actions automatically.

Try example

Risk review automation

Financial Services / Compliance

We're a mid-size asset manager. Our compliance team manually reviews hundreds of client emails, trade notes, and internal communications every week looking for potential market abuse or mis-selling. It takes three full-time people and they still miss things. Regulators are asking us for audit trails we can't produce. Success means automated flagging of suspicious communications with an explainable risk score, a searchable audit log, and a significant reduction in manual review hours.

Try example

No-show risk scoring

Healthcare / Operations

We operate 9 outpatient clinics. Patient no-show rates are running at 23%, which is costing us roughly €2.1 million per year in lost appointment slots. We have 4 years of appointment history, patient demographics, referral source, GP data, and weather records. We've tried reminder SMS but it's not targeted — we send the same message to everyone. Success is a model that scores each upcoming appointment by no-show risk so our admin team can apply the right intervention — double-booking high-risk slots, sending extra reminders, or calling patients directly.

Try example

Semantic knowledge search

Professional Services / Knowledge Management

We're a 400-person consulting firm. When a new project starts, our consultants spend 2-3 days searching through old proposals, past project deliverables, methodology documents, and email threads trying to find relevant prior work. Everything is in SharePoint, email, and a mix of Word and PDF files, completely unsearchable in any meaningful way. Half the time they give up and start from scratch. Success is a system where a consultant types what they're working on and gets back the 5 most relevant past projects, the key findings, and the people who worked on them — in under 10 seconds.

Try example



Practical AI for manufacturing, quality, and project teams.



info@qdes.org

Denmark company - CVR 42351830