

Hawkrose

Procurement Masterplan

Service Definition





Copyright

The Copyright in this work is vested in Hawkrose Ltd and the document is issued in confidence for the express purpose for which it is supplied. It must not be reproduced, in whole or in part, or be used for any other purpose without prior written consent being obtained from Hawkrose Ltd, and then only on the condition that this notice be included in any such reproduction. No information as to the contents or subject matter of this document, or any part thereof, arising directly or indirectly there from shall be given orally or in writing or communicated in any manner whatsoever to any third party without the prior written consent of Hawkrose Ltd.

Copyright © Hawkrose 2026



Table of Contents



1	Procurement Masterplan.....	4
1.1	Service Description.....	4
1.1.1	Optimisation	5
1.1.2	Customisation	5
1.1.3	Large Language Model (LLM) element to tool.....	6
1.1.4	Onboarding.....	8
1.2	Service Support	8
1.2.1	Support Levels	8



1 Procurement Masterplan

1.1 Service Description

Hawthorn's Procurement Masterplan software provides Public Sector Procurement Project Managers with a digital masterplan which can predict, and integrate, the time, cost, performance, and risk parameters for an investment project. Underpinned by HMT Green Book rules – this system allows users to highlight to both industry and independent scrutiny the Critical Path for project delivery.






For more information see the Hawkrose website: www.hawkrose.com

1.1.1 Optimisation

The Optimiser uses genetic algorithms to search for the optimal solution to a problem – e.g. selecting the optimum critical path to deliver the project to time, cost and performance parameters. Having an A.I tool predict the likely future programme plan allows users to test and adjust their plan accordingly.



The best results can then be saved and assessed using the Evaluator.

Have we done this before?

- This Tool was developed for the Royal Navy and is a pan-Defence Lines of Development (DLOD) “Programme Masterplan” which links programme schedule with forecast cost.
- It contains a Genetic Algorithm Optimiser and is currently deployed on UK MOD Servers – and can be reconfigured for any procurement or support use case.
- Hawkrose has also delivered a schedule optimisation tool for a major UK Airport using machine learning to model delay curves based only on schedule inputs, providing a baseline model which can be used to understand the impact of proposed changes.
- Critically, Hawkrose use our core product set and then reconfigure our tools based on the specific domain, industry and customer requirements.

1.1.2 Customisation

The Procurement Masterplan is designed to be highly customisable and configurable. Customisable elements include:

- Assets
- Reference data
- Data structure



- Page flow
- Optimisation constraints
- Scenarios, scenario scoring
- Scenario linkages and nested scenarios

The platform is designed to be reconfigured and customised for each decision, transformation programme analysis or investment evaluation.

New data can be configured through the UI or uploaded via CSV.

1.1.3 Large Language Model (LLM) element to tool

This product – and many products in HM Government – needs to be supported by data inputs which are often stored in PDF and other CSV files. To address this challenge, there is a need for a system that can assist the analysts in the generation of these proposals using conversational AI. This system would allow analysts to input a set of requirements and quickly generate a preliminary proposal with estimated costs and timelines. By leveraging artificial intelligence, this solution aims to reduce the time burden on analysts by primarily reducing the research time required, improve the accuracy and consistency of estimates by leveraging domain specific past data, and reduce the proposal generation time by partial/complete automation.

Hawkrose can meet this requirement by using state of the art models and technologies available in the market. At its core it will be a **large language model (such as GPT or Llama)** that has been augmented with domain specific data through **RAG (Retrieval Augmented Generation)**, **fine tuning**, or a combination of the two. The exact approach will be determined at the proof of concept stage utilising clear and robust set of evaluation criteria. The development will process will utilise **LangChain**, an established framework for creating applications with LLMs, simplifying and streamlining the development process by exploiting availability of pre-built interfaces/ components.




Figure 1: Basic RAG Framework

Credit: <https://www.linkedin.com/pulse/how-build-your-own-rag-chatbot-using-langchain-streamlit-sri-laxmi-fktbc>




Figure 2: LangChain

Credit: <https://datasciencedojo.com/blog/understanding-langchain/>

LLM models comparison

While there are several models available on the market, there are few that always dominate the conversation. These are Open AI's ChatGPT/ GPT family, Meta's Llama, TII's Falcon, and Google's Mistral. At the proof of concept stage, the models will be compared against each other using the chosen evaluation criteria but also taking into account factors such as:

- Compute cost
- Hosting cost
- Data privacy
- Ease of integration etc.

An overview of the comparison based on research:

Criteria	OpenAI GPT (GPT-4)	LLaMA (Meta)	Falcon	Mistral
Performance	Best in class for text generation, comprehension, and reasoning.	Good, but slightly below GPT-4.	Strong text generation, slightly below GPT-4.	Efficient, high-quality text generation, competitive with LLaMA and Falcon.
Customization	Fine-tuning available (enterprise only), easy integration with Langchain/RAG.	Fully customizable via open-source fine-tuning.	Open-source, can be fine-tuned for specific tasks.	Fully open-source, supports fine-tuning.
Cost	Higher cost due to API usage.	Free to use, self-hostable.	Free to use, self-hostable.	Free to use, self-hostable.



Hosting	Cloud-based (no self-hosting).	Requires self-hosting, more control over data.	Requires self-hosting, more control over data.	Requires self-hosting, more control over data.
Data Privacy	Data passes through OpenAI servers.	Full control over data with self-hosting.	Full control over data with self-hosting.	Full control over data with self-hosting.

The LLM element of this product can be purchased alongside the Procurement Masterplan tool, or separately.

1.1.4 Onboarding

The Procurement Masterplan can be delivered as a hosted managed service or deployed into a customers' hosting environment. If deployed into a customer hosted environment a deployment guide is provided and training given to support system upgrades.

User guides are provided as standard, and training can be purchased as an additional service. This can be online or onsite, depending on customer requirements. Hawcrose also offer continuous user support packages where Hawcrose consultants can act as expert users.

1.2 Service Support

1.2.1 Support Levels

Support models are tailored to customers' requirements and based on agreed SLAs. Response times and support coverage is agreed and documented in a managed service agreement and tailored to the requirements of each customer and system deployment.

The Hawcrose standard approach is to offer three levels of support:

- **Bronze** – UK Office hours support with email-based fault notification
- **Silver** – Extended hours (7:00 – 22:00) 7 days per week with email-based fault notification and trouble ticketing system
- **Gold** – 24/7 support 7 days per week with email and phone-based fault notification and trouble ticketing system

Bespoke support packages can be agreed based on a Service Level Agreement with time to respond targets if required.

All customers are allocated a technical account manager to act as a point of escalation if required.