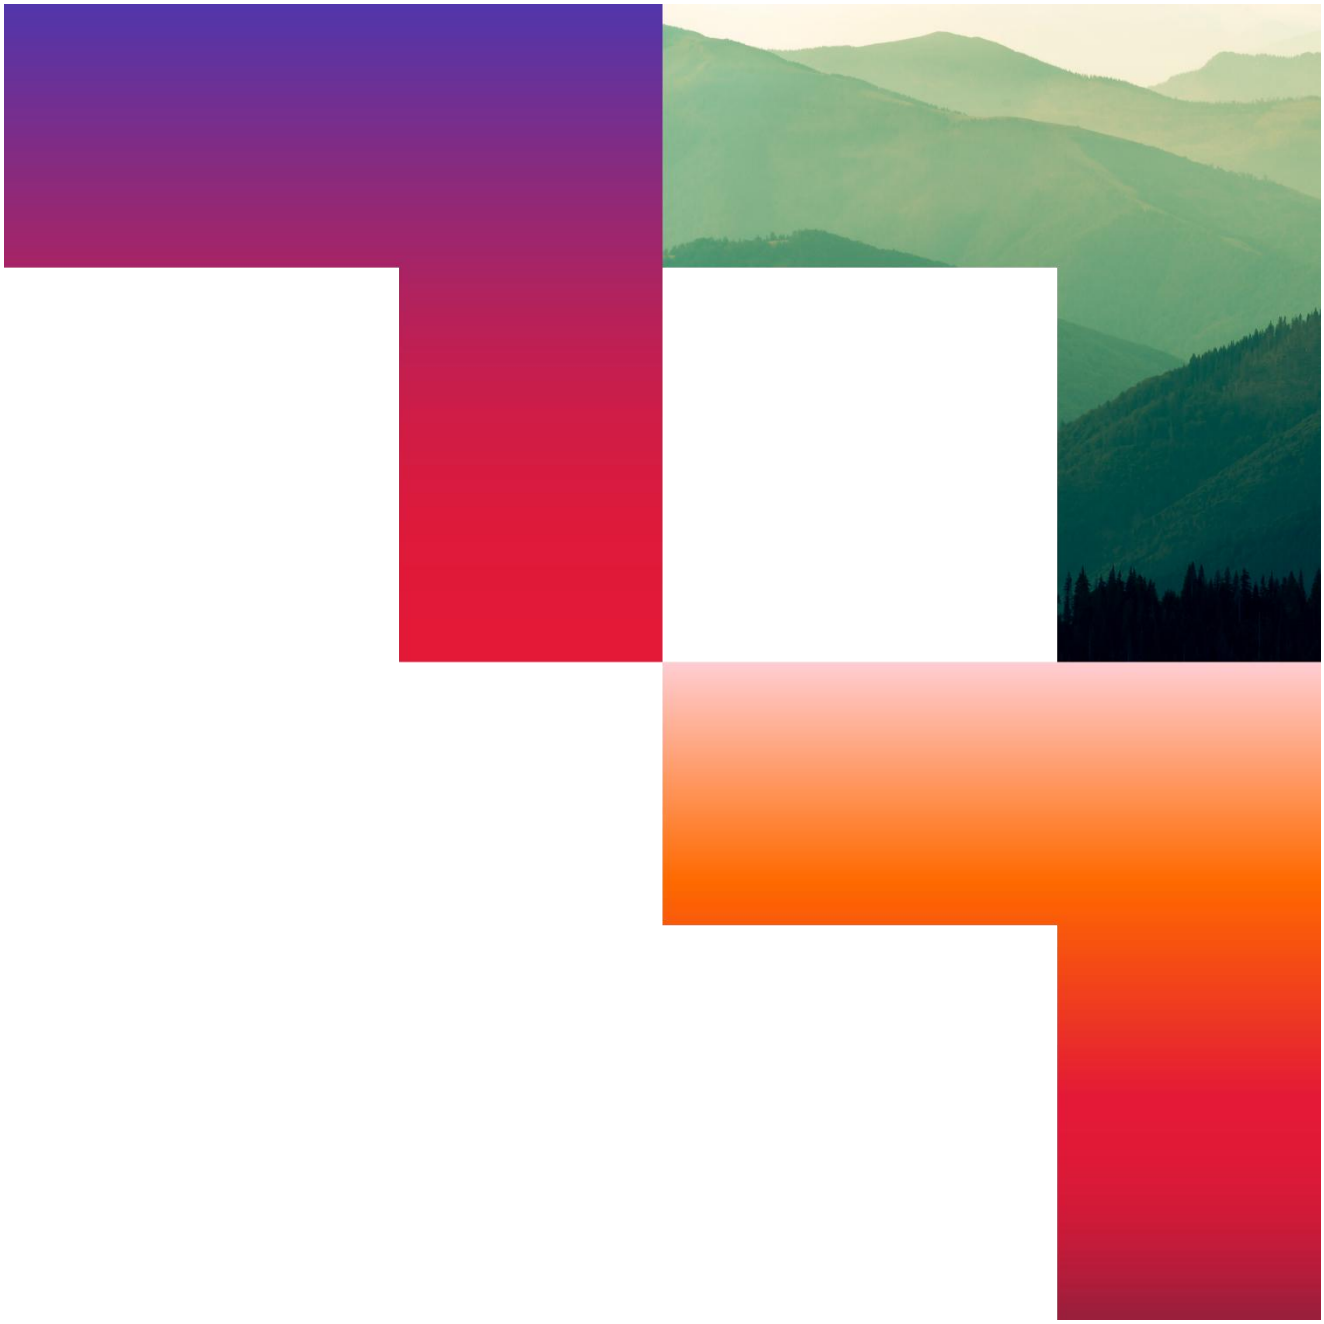

Machine Vision

Contracting body

30 Jan 2026

Issue 1.0

Table of contents

Service Definition overview

CGI Machine Vision is a proprietary visual generative AI platform designed to transform video and image data into real-time operational intelligence. The service enables automated interpretation of visual information to support monitoring, inspection, behavioural analysis, people dynamics and asset condition assessment across complex and dynamic environments.

Why CGI?

CGI brings together proprietary intellectual property, deep domain expertise, and enterprise delivery capability in the CGI Machine Vision service. Unlike point-solution providers, CGI combines advanced computer vision technology with an understanding of complex operational environments across public and regulated sectors.

CGI Machine Vision is built on CGI-owned Vision Transformer architecture and Large Vision Model technology, enabling adaptable, prompt-driven visual understanding without the need to rebuild models for each new use case. This approach allows the solution to evolve rapidly as operational needs change.

CGI’s experience across sectors such as government, transport, utilities, and infrastructure enables the solution to be aligned with real-world operational workflows rather than deployed as a standalone technical product. The agile, proof-of-concept–led delivery model reduces risk, demonstrates early value, and supports continuous refinement.

CGI’s scale, global presence, and experience delivering mission-critical services provide assurance that Machine Vision can be integrated, deployed, and evolved within enterprise and public-sector environments while maintaining flexibility and control.

1. Service Description

CGI Machine Vision is a proprietary visual generative AI platform designed to transform video and image data into real-time operational intelligence. The service enables automated interpretation of visual information to support monitoring, inspection, behavioural analysis, and asset condition assessment across complex and dynamic environments.

The solution is based on CGI-owned deep neural network and Vision Transformer architectures, forming a Large Vision Model analogous to large language models used in text-based AI, but optimised for images and video. The model is pre-trained on millions of images, providing broad visual understanding without the need to retrain models from scratch for each new use case.

Machine Vision supports natural-language prompting and contextual rules, allowing use cases to be adapted dynamically (for example, differentiating staff and visitors based on visual cues). This enables the platform to understand relationships between people, objects, and environments rather than performing isolated object detection.

The service is delivered using an agile, proof-of-concept–led framework, typically implemented in short delivery cycles to demonstrate early value, validate accuracy against real data, and allow rapid refinement or pivoting as operational needs evolve.

2. Onboarding and Offboarding

Onboarding typically begins with a discovery and qualification phase to define the client's business challenges, operational workflows, and target use cases. Clients may provide representative image or video data to support baseline evaluation of the Machine Vision foundational model.

Delivery commonly proceeds through a proof of concept or pilot phase, implemented using agile two-week sprints. This approach allows early validation of accuracy, responsiveness, and business value while enabling the solution to be refined or re-targeted as client needs evolve.

2.1. Onboarding

CGI Machine Vision is delivered through a structured project working with our clients as one team. We prioritise the ethical and safe use of AI which we apply through strict governance of its usage.

2.2. Offboarding

CGI Machine Vision can be off-boarded through a structured process to ensure business continuity and minimal disruption to client operations.

2.2.1. Access to data upon exit

Depending on whether the implementation is in CGI's secure cloud environment or the client's private cloud, all visual data resides with the client.

3. Service constraints

CGI Machine Vision requires access to suitable image or video sources and sufficient compute capability to support real-time inference. Deployments depend on the availability, positioning, and quality of camera infrastructure and on the presence of GPU-enabled edge devices or servers.

The effectiveness of the solution is influenced by environmental factors such as lighting conditions, weather, and camera angles. While the transformer-based architecture is designed to adapt to variation and occlusion, some highly nuanced or visually ambiguous use cases may require additional validation to confirm suitability.

Machine Vision is intended to complement existing operational workflows and is not a replacement for human decision-making where interpretation or judgement is required.

4. After sales support

After deployment, CGI supports continued optimisation and evolution of the Machine Vision solution in line with client needs. This includes refinement of analytics, adjustment of prompts or rules, and adaptation of the solution to new scenarios or environments as operational requirements change.

Support activities are aligned to the agile delivery framework used during implementation, recognising that use cases and priorities may evolve over time. The support model is flexible and tailored to the nature of the engagement rather than prescriptive.

5. Technical requirements

CGI Machine Vision requires access to image or video streams and sufficient GPU-enabled compute resources to perform inference. The solution is built using technologies including .NET Core, TensorFlow, CUDA, ONNX, and Python.

The platform supports deployment on low-power edge devices, GPU-accelerated edge servers, or hybrid architectures. Where deployed at the edge, models are pruned and optimised to operate efficiently on constrained hardware.

Connectivity requirements depend on the deployment model. In edge deployments, processing occurs locally, with only metadata and analytical outputs transmitted upstream in industry-standard formats for integration with enterprise systems, dashboards, or analytics platforms.

6. Hosting options and locations

CGI Machine Vision supports flexible deployment models, including client-hosted environments. The solution can be deployed directly on edge devices, on edge servers integrated with standard IP cameras, or within hybrid architectures that combine edge processing with centralised systems.

Edge-first deployment enables processing to occur close to the point of data capture, reducing latency, minimising bandwidth usage, and supporting environments with limited or intermittent connectivity. Hybrid and cloud-integrated options support centralised monitoring, reporting, and enterprise integration where required.

7. Security

CGI Machine Vision is designed with security and privacy considerations embedded into its architecture. Processing can occur entirely on-site at the edge, ensuring that raw video and image data does not need to leave the operational environment.

In edge deployments, only metadata, events, and analytical outputs are transmitted to downstream systems, reducing exposure of sensitive visual data and minimising bandwidth requirements. The platform supports integration with existing enterprise environments and security controls.

The solution focuses on behavioural and contextual analysis and does not analyse or store personal characteristics such as race or gender. This privacy-by-design approach supports use in regulated and operationally sensitive environments.

8. Service Pricing and Terms and Conditions

8.1. Pricing

Please refer to the G-Cloud 15 portal for the rates available.

8.2. Terms and Conditions

Please refer to the associated Terms and Conditions Document relevant for this Service including termination by the Buyer or Supplier.

9. Further information

For more information about this or any of our G-Cloud services, please contact us:

Phone: +44 0845 070 7765. Ask to speak to the G-Cloud team.

Email: uk.gen.ccsframeworks@cgi.com, including the following information:

1. The name of this service.
2. The name of your organisation.
3. Your name and contact details.
4. A brief description of your business situation.
5. Your preferred timescales for starting the work.

We will prepare a quotation for you. Once the quotation is agreed upon, we will issue you the necessary documentation (as required by the G-Cloud Framework) and ask you to provide us with a purchase order.

Once we have received your purchase order, the services will be configured to the requirements agreed with invoices issued to you for the services procured. On a monthly basis, we will also complete the mandated management information reports to Government Procurement.

PROPRIETARY AND CONFIDENTIAL

The information contained in this document is confidential to CGI and/or CGI group companies. This document shall not be reproduced in any form or by any mechanical or electronic means, including electronic archival systems, or submitted to a generative AI tool without the prior written approval of CGI. The receiving party may use this document and the information contained in it for the purpose of evaluating CGI's proposal only. Any personal data included in this document must not be replicated, stored or distributed beyond the immediate recipient or used for any purpose other than evaluating the document.

This proposal is subject to contract and shall not be binding unless and until execution by CGI and Contracting body of a final agreement to the proposal, containing the terms and conditions that will govern the relationship between the parties. Any final agreement is conditional (inter alia) upon due diligence and customary business investigations by CGI. The results of such due diligence and/or investigations may impact upon content of this proposal, including the business structure, business terms and financial arrangements.

If you have received this document by mistake, note that the reading, the reproduction or the distribution of this document is strictly forbidden. You are hereby requested to inform us by telephone at +44 0845 070 7765 and to return this document by special delivery marked for the attention of Simon Figg.

Except where indicated otherwise, all names, trademarks, logos and brands (registered or not) referred to in this document are the property of a company in the CGI group or its licensors.

The information in this proposal is submitted on 30 Jan 2026 - Issue 1.0 on behalf of CGI by the following authorised representative:

Simon Figg

Frameworks Manager

CGI IT UK Ltd

14th Floor, 20 Fenchurch Street, London, EC3M 3BY

Tel: +44 0845 070 7765

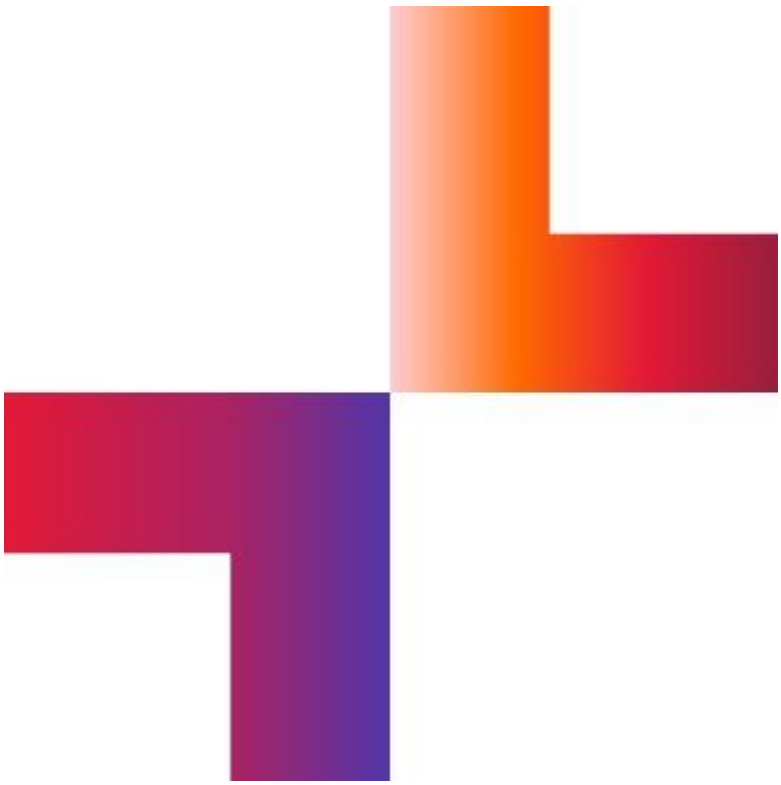

cgi.com/uk