WHITE PAPER • 2023

BAI Technical White Paper

An Overview of BenevolentAI's Security Programs and Best Practices

Benevolent[AI]

Table of Contents

Who we are

BenevolentAI Technology Ltd is a subsidiary company of BenevolentAI (AMS: BAI), a leading company that combines cutting-edge Artificial Intelligence (AI) and science to accelerate biopharma discovery and development and is listed on the Euronext Amsterdam stock exchange. We unite AI and cutting-edge science to discover and develop new medicines for complex diseases by creating purposeful technology in the service of science to augment human intelligence and empower scientific discovery.

The Benevolent Platform™ is a flexible and scalable AI-enabled drug discovery engine that enables scientists to formulate new hypotheses and rapidly discover high-quality drug targets based on a better understanding of diseases. We seek to improve patients' lives by applying technology designed to generate better data decision-making and in doing so lower drug development costs, decrease failure rates and increase the speed at which medicines are generated. For the purpose of applicable data protection laws, we are the data controller for the personal information you share with us during the registration/sign-up process and the data processor for any data you upload to the plugin for the purpose of carrying out the required analysis. Our representative in the European Economic Area (the "EEA") is BenevolentAI and our representative in the United Kingdom ("UK") is BenevolentAI Limited.

Overview

Welcome to our Technical Security White Paper - a comprehensive guide to the measures we undertake to safeguard your data and ensure the integrity of our services. As a provider of services for our customers, we are fully committed to maintaining high standards of cybersecurity and data privacy to identify and mitigate potential product security risks and drive security awareness across our company.

Our services are hosted entirely on Amazon Web Services (AWS), the world's most comprehensive and broadly adopted cloud platform. We leverage AWS's advanced infrastructure to offer a secure, reliable, and scalable environment for our customers. The platform's robust security features, coupled with our meticulous security practices, ensure that your data is protected at all times.

In this white paper, we will detail our security strategies, which encompass everything from software development and encryption to infrastructure security and architecture. As technology evolves, so does the nature and sophistication of potential threats. Risks associated with infrastructure, intellectual property (IP), and data privacy are in a constant state of flux. In response to this ever-changing threat landscape, we have implemented stringent controls across each area to provide a secure and compliant platform for our customers.

Whether it's the encryption of data, both at rest and in transit, the isolation provided by single-tenant architecture, or the rigorous testing in our software development lifecycle, every step is designed to protect your data and ensure the integrity of our platform.

The purpose of this document is to detail how BenevolentAI's security and privacy practices have been applied to our systems, what you should know about maintaining the application's security, and how we can partner with you to ensure security throughout your use of our products.

Products

Our product offering will be detailed on our website **www.benevolent.ai**

Architecture

Our platform's architecture is designed with a central focus on security and isolation, ensuring that each of our customers' data remains secure, isolated, and always available when they need it. At the heart of our architecture design is the adoption of single-tenancy, where we operate an independent AWS account for each of our customers. A high-level illustration of our architecture is detailed on the next page for reference.

This design strategy allows us to provide an additional layer of security and separation. Each customer's data and services are hosted in a dedicated environment, minimising the risk of data leakage between different clients. This unique approach to infrastructure ensures that the resources are not shared, enhancing the integrity and security of our customers' data.

Our platform runs on a set of Amazon EKS clusters, which provide the backbone for our containerised application deployment. Each customer's services are further isolated with unique endpoints, adding another layer of data security. This means each customer's interaction with our platform is limited to their own unique environment, providing a secure and personalised experience.

A robust authentication layer within our platform provides secure access controls, ensuring that only authorised users can access the appropriate resources.

The automated deployment of our architecture also allows us to scale easily, enabling us to swiftly accommodate our customer's growth or additional requirements.

Our architecture is purposefully designed to prioritise security, isolation, and flexibility. We have leveraged the best of cloud technologies to ensure that our customers can trust us with their data, and can depend on our platform for their crucial business operations.

Architecture

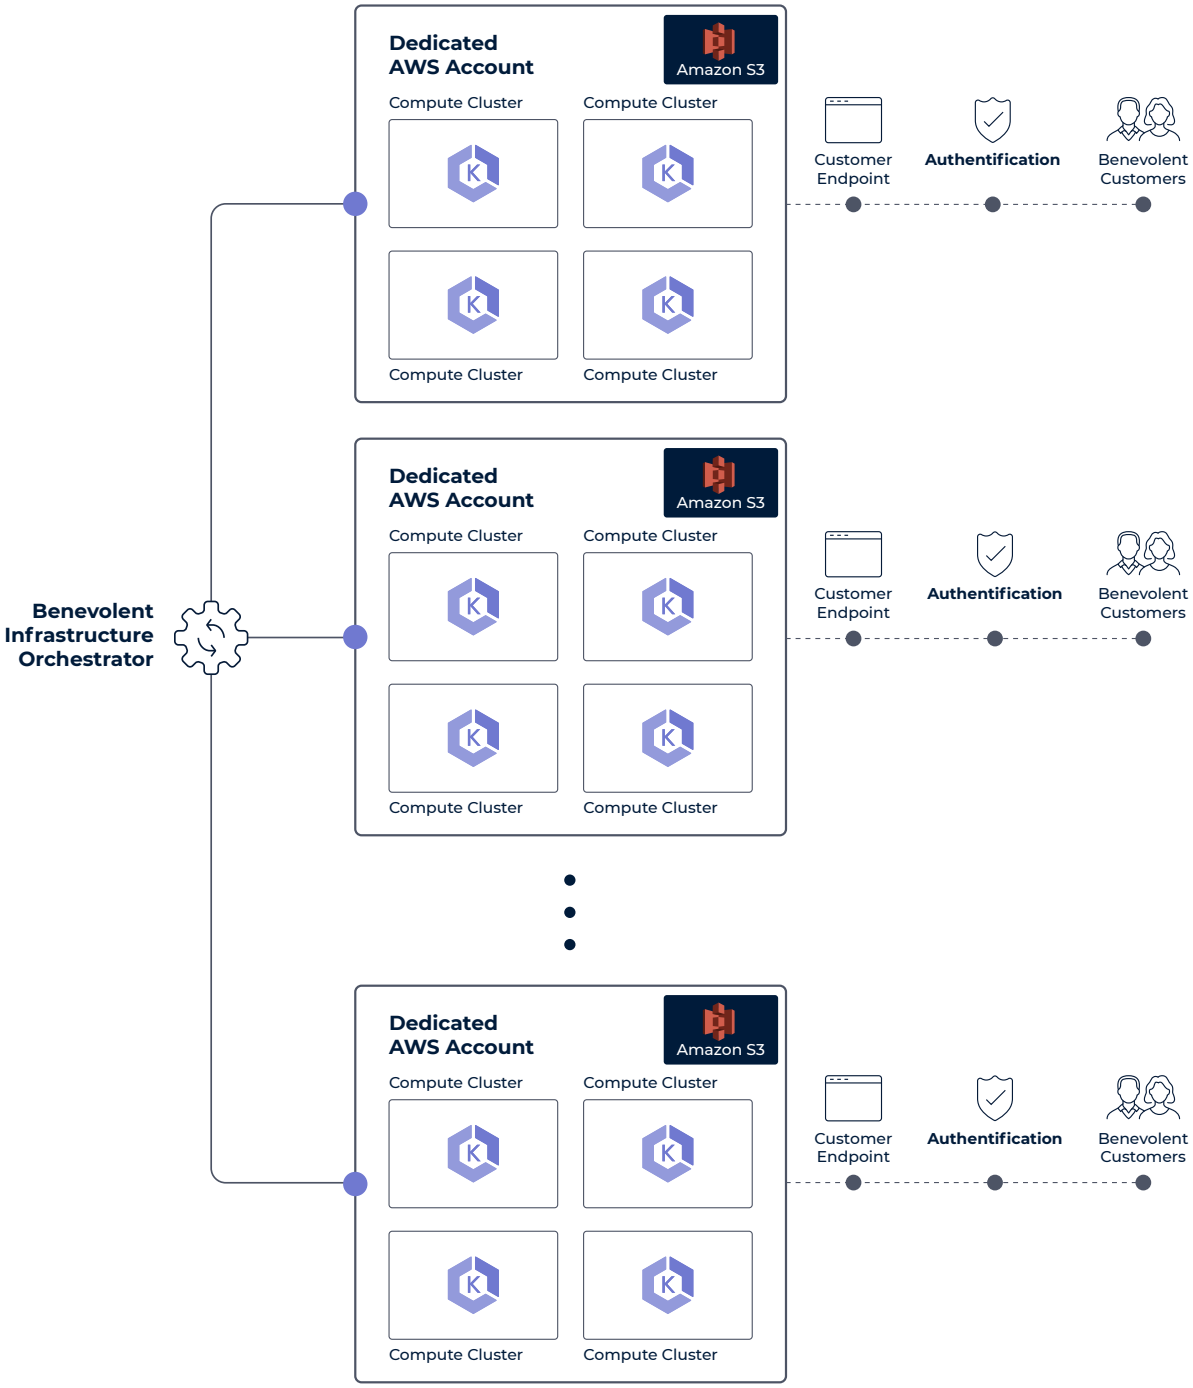

Data Flow and Transmission

Ensuring the security and confidentiality of our customer data is paramount to us. As such, we have designed a rigorous data flow and transmission architecture to safeguard the data that our customers entrust to us, hence ensuring privacy and security by design as an essential part of our overall architecture.

Central to our data handling strategy is the concept of account-level isolation. Each customer's data is stored, processed, and managed within a dedicated AWS account. This approach creates a clear boundary around each customer's data, ensuring no cross-contamination or unintended access.

Whether in motion or at rest, customer data is constrained at all times within the confines of a dedicated AWS account. It does not traverse any shared systems or networks, nor does it mingle with data from other customers. In essence, your data is stored, processed, and used exclusively within your dedicated account. This ensures the integrity and confidentiality of all data throughout its lifecycle, while still allowing you to leverage the power of our platform and services.

We understand that our customers entrust us with their most sensitive and valuable data, and we take that responsibility very seriously. Our data flow, transmission, and retrieval protocols are reflective of this understanding and are built to provide the highest level of security and peace of mind which is reflected in the contractual documentation we enter into with customers.

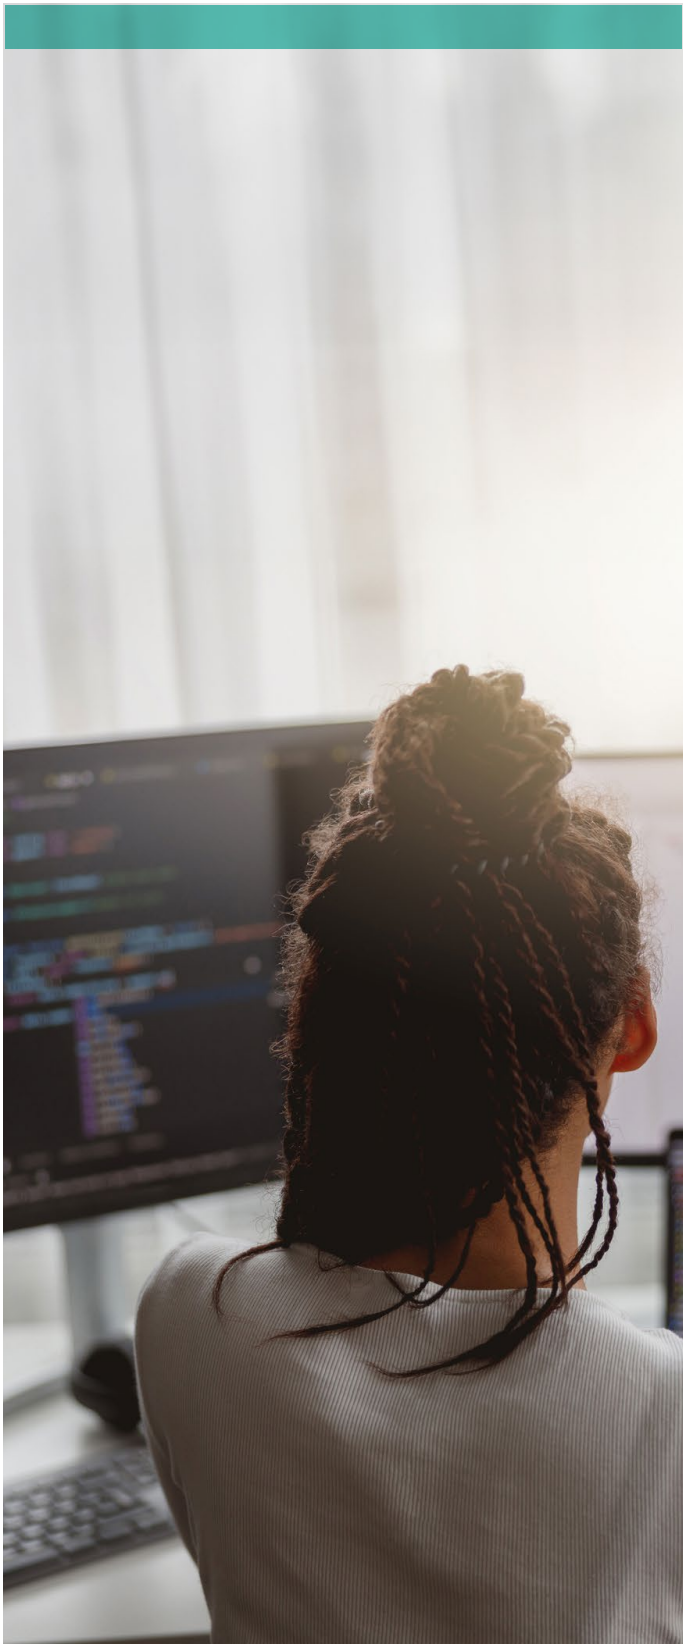

Data Storage and Encryption

At BenevolentAI, we take security seriously and have implemented a collaborative and comprehensive approach to security management, that involves a multi-functional team's strategic framework

Our Security team, Site Reliability Engineers (SRE), Software Engineers, Compliance, Privacy, and Legal teams all work hand in hand to ensure high levels of data protection. Each team brings its expertise to the table, from the development of secure software and the management of robust and resilient systems, to ensuring our practices are compliant with all necessary regulations.

This approach enables us to weave security into the fabric of our operations, making it an integral part of our organisation. By bridging the knowledge and efforts of these diverse teams, we ensure your data is not just encrypted and stored securely, but is also managed and handled in a manner that upholds our commitment to security, reliability, and compliance.

To this end, we have deployed a comprehensive encryption strategy that protects data at all times, both in transit and at rest.

All of our data is warehoused in the AWS London Region. Our storage services, including AWS's S3 for object storage and Elastic Block Store (EBS) for block storage, are fortified with automatic encryption. With the aid of AWS's standard encryption tools, every piece of data is encrypted before it is stored, ensuring its confidentiality and security.

Our encryption strategy does not stop at rest. When data is on the move - either across our internal systems or over the internet - it is shielded by industry-standard encryption protocols. This means that all communications and transfers are secured, maintaining the privacy and integrity of your data while it is in transit.

Infrastructure Security

At the forefront of our business operations is a commitment to security, a principle deeply embedded within our infrastructure. Leveraging best-in-class tools offered by AWS, we have implemented robust security measures that protect, monitor, and ensure the integrity of our customer's data.

We deploy tools such as AWS Config, GuardDuty, and Security Hub among others, which work to safeguard our platform. These advanced technologies continuously monitor our systems, track configuration changes, and detect any potential threats or unusual activity, allowing for immediate response to secure our infrastructure.

Adding another layer of protection, we adhere to the Zero Trust security model. This approach assumes no implicit trust, verifying each request as though it originates from an open network, regardless of where it's initiated. Internal access to our infrastructure is granted strictly via a secure VPN, adding an extra layer of data protection and access control.

Our policies are underpinned by the principle of 'least privilege', which means access rights are limited to the bare minimum that each user needs to perform their tasks. This helps reduce the attack surface and further safeguards our infrastructure.

Enforced Multi-Factor Authentication (MFA) is another critical aspect of our security strategy. This additional layer of access control ensures that user identities are thoroughly verified before they are allowed access to our systems.

These tools and principles work collectively to offer a holistic view of our security and compliance posture. Despite the complexity of the processes, our objective remains simple and clearcut to provide a secure, reliable, and compliant environment for our customers.

Software Development Lifecycle

At BenevolentAI, we understand that the foundation of secure operations lies in the way we develop our software. Our Secure Software Development Lifecycle (S-SDLC) follows industry best practices, ensuring the delivery of reliable, secure, and efficient software solutions to our customers by identifying and mitigating security risks.

Our software development process starts from the ground up, with a dedicated AWS account for each stage - development, staging, and production. Each of these environments are meticulously separated, ensuring an additional layer of data security and minimising the risk of unintended access or data breaches.

At the design stage, every new product and a major feature is reviewed by the security team and subjected to comprehensive risk assessments and threat modeling. Then, in our development environment, our experienced software engineers work on new features, enhancements, and bug fixes. They implement security best practices from the very beginning, adopting a “security by design” philosophy. Each code commit undergoes rigorous automatic and manual code reviews, including Static Application Security Testing (SAST), to ensure it meets our stringent quality and security standards.

Once the new software is ready, it moves to our staging environment where it is subjected to comprehensive testing. We use this opportunity to identify and rectify any potential vulnerabilities, bugs, or performance issues, prior to being deployed in the production environment.

Production deployment follows a controlled and well-documented release management process. This includes the use of continuous integration and continuous deployment (CI/CD) techniques, ensuring seamless and error-free transitions of software updates. All state-altering changes to production environments are performed through an automatic process with manual interventions being only allowed during emergencies and limited to authorised senior engineering personnel.

All these processes, checks, and balances that we have instituted in our Software Development Lifecycle (SDLC) are strategically designed to identify and mitigate any potential issues as early as possible. The goal is to ensure that by the time a software update or new feature reaches our customers, it has undergone exhaustive testing, review, and quality assurance measures. This early and proactive approach to risk management in our SDLC ensures that our software releases are safe, secure, and reliable.

Audit Logging

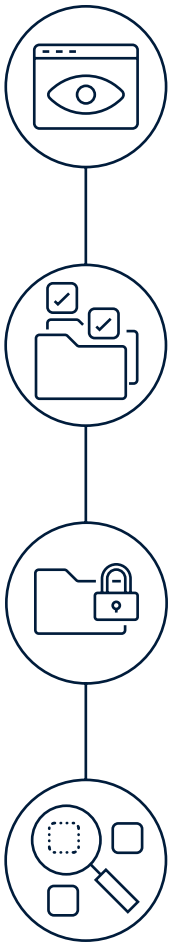

At the heart of our security strategy lies our commitment to complete transparency and accountability. We have built a comprehensive audit logging system that captures all activity within our platform. This extensive coverage equips us with the ability to scrutinise our operations meticulously, ensuring that all activities align with our high security standards.

Our strategy involves leveraging industry-standard tools, including AWS Flow Logs, CloudTrail, Kubernetes Audit Logs, and logs from our Zero Trust Access (ZTA) solution. Each of these tools plays a vital role in the audit logging landscape.

AWS Flow Logs and CloudTrail deliver extensive visibility into our network operations and user activities within AWS. Kubernetes Audit Logs augment this by providing us with detailed records of the sequential activities in our Kubernetes clusters. Lastly, with our ZTA layer, we not only manage access to our infrastructure but also capture each action and command executed by our personnel, adding an extra layer of security and transparency.

Once recorded, the data within these logs cannot be modified, deleted, or tampered with, even by administrators, ensuring a trustworthy and unalterable record of all activities. This immutability provides an added layer of security and maintains the integrity of our audit trails.

Finally, our audit logs are monitored and automatically analysed to spot patterns, detect anomalies and address potential security concerns promptly. By maintaining robust audit trails, we assure our customers of our commitment to safeguarding their data, upholding transparency, and ensuring the highest level of accountability.

Governance

BenevolentAI's Governance Framework (BGF) is the overarching structure that defines the policies, processes, training, and other controls we have implemented to ensure our product's responsible and ethical use of customer data. With the BGF, we have implemented adequate administrative, technical, and physical safeguards to help protect against security incidents and privacy breaches involving our product, provided these products are used in strict adherence to the Product's terms of use.

We have established processes and standards for the design, development, testing, and validation of our AL models, including processing, feature selections, and performance validation. We are also cognisant that systems and threats evolve very rapidly, and that no system can be protected against all vulnerabilities. Therefore we consider our customers and collaborators the most important partners in maintaining security and privacy safeguards.

These controls include:

- Privacy and Security by Design (PbD/SbD) Policy
- Model development and Validation
- Product and Supplier Risk Assessment
- Secure IP Management
- Vulnerability and Patch Management
- Secure Coding Practices and Analysis
- Transparency
- Vulnerability Scanning and Third-Party Testing
- Access Controls Appropriate to Customer Data
- Risk and Incident Response Management
- Continuous Evaluation and Improvement

Please keep in mind that the way the BGF controls are applied may differ depending on the size of your organisation and the regulatory environment you operate in.

BenevolentAI reserves the right to modify any information in this technical security white paper at any time, and the modified content will be added to the new version of this technical security white paper without separate notice but the most recent version will be posted to our website.

For any other information, please visit our website

Business Continuity

Our Business Continuity best practices involve ensuring uninterrupted operations, data integrity, incident management, and disaster recovery. Here's a summary of key practices:

- Risk Assessment and Business Impact Analysis
- Disaster Recovery Planning
- Continuous Monitoring and Alerts Management
- Incident Response
- Vendor/Third-parties Management
- Regular Testing and Exercising
- Employee Awareness and Training
- Documentation Management

These best practices enhance the resilience of our product, minimise downtime, and ensure the continuity of critical processes in the face of potential disruptions or disasters.

For any further technical security information regarding our Business Continuity visit Benevolent.ai

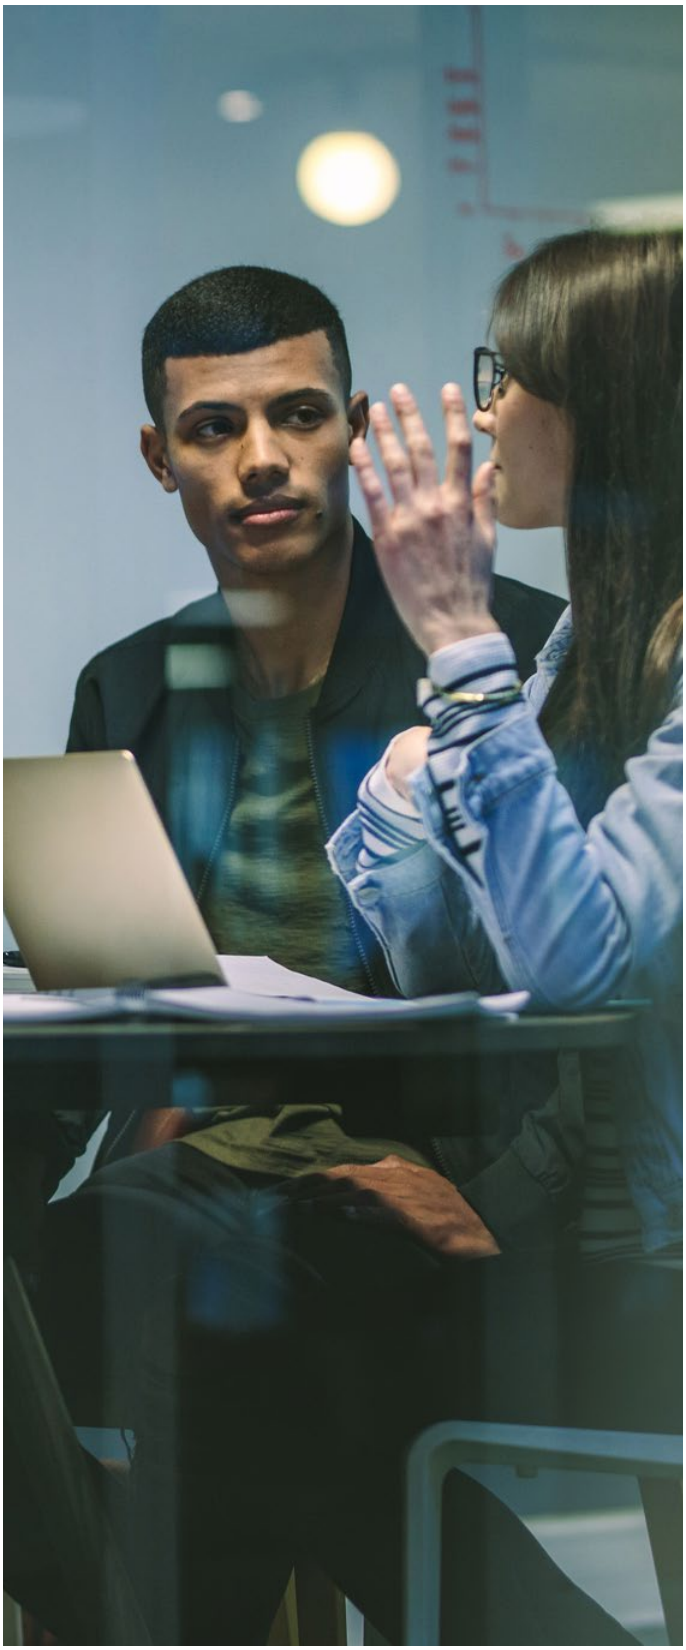

Service Handling Practices

We have set up specific strategies and protocols to ensure that our available products run smoothly and is effectively managed. These practices include:

- **User Support:** We have a robust system to support our users. Our dedicated support team can quickly address any questions, concerns, or feedback. We provide clear documentation and communication channels to help users navigate the application and overcome any challenges they may face.

- **Maintenance and Updates:** To guarantee optimal performance and security of the AI application, we have prioritised regular maintenance and updates. This includes bug fixes, feature enhancements, and system upgrades, following industry best practices for software development to maintain the application's reliability.

- **Scalability and Performance:** To adequately manage increased user demand over time, we have implemented strategies such as load testing, infrastructure optimisation, and monitoring system performance to ensure the application can handle growing user traffic and maintain responsiveness.

- **Data Handling and Privacy:** Proper data handling practices are in place through our robust Governance framework. This involves data protection measures, secure storage, and access controls.

- **Incident Detection and Response:** A comprehensive incident response plan has been implemented to aid proactive monitoring, threat detection, and effective incident management protocols to minimise downtime and protect user data.

- **Documentation and Training:** We maintain a well-documented process and training materials to support the service-handling practices of the application. This includes comprehensive guides, tutorials, and knowledge bases to assist users and internal teams in effectively and efficiently utilising the application.

End-of-Life & End-of-Support

Where a version of an application becomes retired, we implement plans to support all our customers in migrating to newer versions or alternative solutions. Our robust migration framework will facilitate a seamless transition and ensure customers can continue to receive support, updates and most importantly continue working on ongoing projects.

We will ensure timely notification to customers of any “End-of-life and end-of-support” timelines to allow time for any operational adjustments that might be required, hence minimising disruptions.

Benevolent AI

©BenevolentAI Limited or its affiliates. All rights reserved.

No part of this publication may be reproduced or transmitted in any form or for any purpose without the express permission of BenevolentAI Limited. The information contained in this Technical Security White Paper, including, but not limited to, specified products and services and third-party platforms and security vendors, may be changed without prior notice. This Technical Security White Paper is provided for informational purposes only, without representation or warranty of any kind and BenevolentAI or its affiliates shall not be liable for any errors or omissions contained herein. The information in this Technical Security White Paper is not a commitment, promise, or legal obligation to deliver any product or service. Any purchase of products and/or services will be subject to a separate written agreement between the customer and BenevolentAI. The only warranties in relation to any BenevolentAI products and services are those that are in express warranty statements accompanying such products and services, if any. Some products marketed by BenevolentAI may include proprietary software components of other software companies.

For more information, please visit **Benevolent.com**