



G-CLOUD 14

**TORTUS AI
SERVICE DEFINITION DOCUMENT**

**Ambient voice transcription and AI generation of notation
and documentation for clinicians**

TABLE OF CONTENTS

Introduction	3
Company Overview	3
Value Proposition	3
What the Service Provides	3
Social Value	5
Products in Detail	5
Associated Services	5
End User Support	5
Fault Reporting	5
Service Parameters	5
Email Support Option	5
Information Assurance	5
Data Back-up, Data Restoration and Disaster Recovery	6
Business Continuity Statement/Plan	6
Privacy By Design	6
Using the Service	6
Ordering and Invoicing	6
Pricing Overview	6
Trail Service	7
End User Ease of Use	7
Training	7
Technical Help	7
Implementation Plan	7
Service Management	7
Performance	7
Maintenance and Upgrades	10
Service Constraints	11
Technology Refresh	11
Service Levels	11
Prioritisation of Reports	11
Target Resolution Times	12
Major Incident	12
Outage and Maintenance Management	12
Monitoring	12



Change Management	12
Financial Recompense Model for Not Meeting Service Requirements	13
Provision of the Service	13
Buyer Responsibilities	13
Technical Requirements and Client-Side Requirements	13
Outcomes/Deliverables	13
After-Sales Account Management	14
Termination Process	14
Our Experience	14
Case Studies	14
Clients	17
Contact	17

INTRODUCTION

COMPANY OVERVIEW

At TORTUS AI, we are proud to introduce our AI-enhanced tool that improves interactions between clinicians and Electronic Health Records (EHRs). TORTUS uses AI to listen to the consultation, create instant documentation such as notes and letters, and clinical coding such as ICD-10 and SNOMED, saving clinicians time and improving data quality for systems.

Our agent, TORTUS, is a combination of AI technologies that mimics human behaviour. Today, TORTUS automates medical documentation, by combining highly accurate speech-to-text AI with Large to create instant notes and letters, which are customisable to match your voice and templates for notes. Beyond that we are adding database access, to allow TORTUS to summarise a patient history for you. Next year, TORTUS will also be able to use the system for you, such as placing orders. The system is capable of automatically generating codes, and we are working on this for a few UK systems right now (ICD-10, OPCS, SNOMED/QOF, HRG4 for ED).

Our commitment to scientific innovation is demonstrated through a large-scale NHS England backed study involving 5,000 patients across outpatients, ambulance services, emergency departments, and mental health settings. This study is set to showcase the robustness and versatility of TORTUS in various clinical environments. We are supported by the UK Department of Business and Trade, emphasising our alignment with national healthcare and business objectives.

TORTUS was co-founded by Dr Dom Pimenta, a physician with sixteen years clinical experience (~ten years NHS), pharmaceutical physician and AI healthcare academic, and a second time founder and Christopher Tan, a machine learning scientist. Both met on the prestigious Entrepreneur First programme (LD19).

The team today is composed of NHS clinicians, engineers, machine learning scientists and commercial operators, from Microsoft, Google and Babylon. TORTUS HQ is based in Holborn Town Hall in central London.

VALUE PROPOSITION

TORTUS is integrated with existing NHS systems such as Epic (and EMIS by Q1 2025), offering seamless functionality without disrupting established workflows. We are:

- DTAC assured
- DSPT compliant
- GDPR compliant
- CyberEssentials Plus Certified
- An MHRA Class 1 Medical Device
- CREST-approved pen-tested

We do not store or retain any patient data.

In addition to G Cloud 14, we are a Crown Commercial Supplier on DPSRM6200 as well, and we are positioned to provide value across the NHS with ease of procurement.

Evidence shows that TORTUS can reduce consultation time by up to 25% in both primary and secondary care, allowing clinicians more time for patient care. Through cost capture in clinical coding, improved clinician retention, and increased throughput and volume of patients seen, TORTUS offers immense value to any clinical setting. Currently we are trialling TORTUS across primary care, outpatient secondary care, A&E, the ambulance service, mental health and paediatrics as well as social care.

WHAT THE SERVICE PROVIDES

TORTUS includes essential tools for recording and summarizing clinical consultations. These tools consist of speech-to-text conversion, text summarization, and the generation of referral letters. It is designed for use by clinicians in clinic settings.

The workflow for any physician consultation includes taking a verbal history, performing a physical examination, formulating a diagnosis and treatment plan, and often writing correspondence to another physician. In over 90% of healthcare settings in the UK, this documentation is done digitally. Approximately 30-50% of a clinician's time is spent on documentation.

TORTUS provides "ambient intelligence" by listening to the verbal history-taking process. It can directly convert this information into text, thereby saving the physician time on documentation as it occurs simultaneously with the consultation. Additionally, TORTUS automates the drafting of notes and letters, tasks that can take a physician between 6 to 60 minutes, depending on the complexity of the patient and the care setting. We have been measuring the benefits in various forms throughout TORTUS development, but thorough testing has shown utilising TORTUS in the clinical workflow yields an approximate 20-30% increase in time efficiency and documented increases in documentation quality.

In our testing with physicians, we compared performance before and after utilizing the functionality for clinical consultations in multiple simulations. Results indicated that overall time was reduced by 55% when using the AI compared to when it was not used. We assessed the quality of the documentation produced using a Visual Analogue Scale, which serves as a direct measure of TORTUS functionality. The quality of documentation was consistent and improved by 53% above the baseline. Additionally, we measured cognitive load using the NASA TLX scoring system. This framework allowed us to evaluate the cognitive burden of AI-assisted consultations versus non-AI consultations, revealing a 55% reduction in cognitive burden when using AI.

Key features and benefits:

Key features include:

- Speech to text AI
- Instant medical note generation
- Templates customisation
- Instant referral letter generation
- Instant patient communication content generation
- Quick, accurate dictation
- Instantaneous suggested codes

Key benefits include:

- Reduced clinician's burnout
 - Reduced administrative burden for clinicians
 - Increased direct patient care
 - Decreased clinician cognitive load
 - Improved clinician experience
- Improved patient experience
- Increased cost capture
- Increased consistency and quality of documentation for audit
- Improved coding


- Improved clinician's productivity
- Reduced consultation's time

Case Study 1: Primary Care

Dr David Triska has used TORTUS for over a year, performing over 5,000 consultations. Compared to the same year long period post:

- Consultation time has reduced by >4mins
- Over run reduced by >60%
- Under-run increase by 300% (aka going home early)

And Dave reports huge improvements in quality of life. As a tool for improved patient experience and clinician retention, TORTUS has very high ROI in primary care.



Case Study 2: Outpatient Secondary Care


TORTUS performed a multi-site, year long clinical trial of time-in-motion in OPD departments in London. Data showed users:

- Preferred clinic with TORTUS to without 82% of the time
- Direct care (time with the patient) increase by up to 60%
- Total time spent with patients decreased by up to 9min per consultation
- OPD users spent 20% more time with the EHR and far less time after clinic writing notes
- Letter turnaround time in some cases was instant

As a tool to increase capacity in secondary care OPD and make the experience lighter for both clinicians and administrators, TORTUS is unparalleled.

Results

Top responders **saved ~7 mins** per consultation



Clinicians' Testimonials see the benefit of using TORTUS in their day-to-day

"Avg patient rating 4.75/5 in over 2.5k remote consultations last year. Since Jan I've had a consistent Patient rating of 5/5 every month. Feedback often mentions **listening**. The only thing I've changed is using and changing my consultation style to fit around TORTUS"

GP

"Really impressive software. Was bowled over by the accuracy when we interacted with this last week and more importantly how it recorded the consultation despite the relatively noisy conference hall."

CNIO, London

"I finally started using TORTUS with the S1 template you set up for me this week and I just wanted to let you know that it is a real godsend! I can't stress enough how it transformed my day – I was much more relaxed and really enjoyed speaking to patients more than I have done for a long while! I actually think it made me give better care as well as I was happy to explore more with my questioning and get a wider history without the stress of having to worry about how long it would take to document all the information. I am definitely not consulting without it anymore!"

GP Partner

SOCIAL VALUE

Our mission is to make medicine safer, smarter and kinder by giving every clinician a personal AI clinician assistant. We are building an AI agent for healthcare, a digital worker that can take over all computer tasks from a clinician such as documentation, requesting tests, access and summarising patient information. Our ultimate aim is to build an AI assistant for every clinician, making medicine safer, faster and more humane.

Our first company value is “patient first”. By supporting clinicians, reducing their admin workload, it means an overall safer, smarter healthcare system, and better outcome for patients.

We have a product user group to ensure users are engaged and that we’re gathering their feedback regularly.

We have thoroughly considered the accessibility of TORTUS and its impact on health inequalities to ensure our innovation does not exacerbate disparities. Recognising potential biases in speech-to-text AI models, especially regarding accents, age, and other characteristics, we have implemented strategies to address these issues.

Our speech-to-text models have been tested across various simulated accents in both controlled and real-world settings. They generally capture non-native English speakers' voices as accurately as native speakers, except for medication names. To address this, our Intended Use document recommends that physicians repeat medication names and verify them at the consultation's end.

PRODUCTS IN DETAIL

TORTUS is a web-app that uses ambient technology to listen to the consultation, and automatically create the relevant documentation (clinical notes, letters and suggested codes).

The TORTUS web-app is a cross-platform application that works on multiple OS and browsers.

The TORTUS API is the cloud backend, called via a secure HTTPS by the desktop application. The API is hosted on Google Cloud Run and secured by the Google API gateway and Auth0-powered authentication layer. It is hosted in the user’s region and manages connections to upstream services, also in the same region. Certain endpoints on the TORTUS API call additional services hosted on AWS, Azure and OpenAI, each provisioned in the same region as the end user.

ASSOCIATED SERVICES

END USER SUPPORT

The end user can contact the support team with any questions or queries, responsive from 9am-6pm, Monday–Friday excluding bank holidays.

FAULT REPORTING

Faults can be reported via contact to the support team with any questions or queries, responsive from 9am-6pm, Monday–Friday, with a 24-hour response time.

SERVICE PARAMETERS

EMAIL SUPPORT OPTION

The end user can contact the support team with any questions or queries, responsive from 9am-6pm, Monday–Friday excluding bank holidays.

INFORMATION ASSURANCE

We are certified to CSA CCM version 3.0 and ISO/IEC 27001. We have three in-house Clinical Safety Officers (inc the CEO) plus an external compliance partner who are embedded through the feature development lifecycle, to ensure clinical risk and safety are acknowledged and mitigated throughout. Furthermore, we have a clinical risk register where all changes affected are scored by our clinical team, and once deemed safe, only then can be embedded into the product. From an engineering POV, we have a CI/DC pipeline, reviews of code must be carried out by someone else who didn't initially write the code. Our Chief Technology Officer manages a team of senior engineers and junior engineers all under this umbrella.

DATA BACK-UP, DATA RESTORATION AND DISASTER RECOVERY

Our database is Google Cloud managed, and we have daily back-ups across our storage systems, and point in time recovery for our database.

Priority Level	Description	Response Time	Resolution Time
Critical	Service completely unavailable or major functionality impaired	30 minutes	4 hours
High	Significant impact on service with no workaround available	1 hour	8 hours
Medium	Moderate impact on service with a workaround available	4 hours	24 hours
Low	Minor issues or general inquiries with no service impact	8 hours	3 business days

BUSINESS CONTINUITY STATEMENT/PLAN

Our Business Continuity Plan is in accordance with the NHS DSPT requirements.

Plans are developed to respond to reasonably foreseeable business continuity events or incidents which may impact the organisation. As a minimum plans are developed for the following events:

- Loss of Premises
- Loss of computer system/essential data
- Loss of telephone system
- Incapacity of staff
- Staff Sickness
- Notifiable Data Security and Protection Event

In the event of data loss from a fault in our systems, we have the ability to instantly roll back our services to any point in time. For further data loss, we can restore from our daily back-ups. These back-ups are encrypted and only accessible for our database to restore from.

In the event of a larger systemic outage, such as GCP, we will be unable to provide the complete TORTUS functionality. The app will enter “offline” mode, where users can still record consultations, where they can resume full functionality when our services return to normal.

We will notify users of an outage via email, with regular updates as we progress with the incident.

PRIVACY BY DESIGN

TORTUS stores no patient data whatsoever. We keep some user meta data such as user profile information and usage metrics. Audio blobs that are sent for transcription are never stored in memory in the cloud, and only kept locally during the transcription process. No sensitive

information is ever stored by us or our providers. Local data is deleted after 24 hours automatically, or upon user log off, which can be varied on request by each organisation.

USING THE SERVICE

ORDERING AND INVOICING

An organization-wide contract is established between the organization and the Company, subject to any free or pilot trial phase. For primary care, invoices are issued either monthly or annually, depending on the patient population size.

For secondary care facilities or larger organizations, an annual license fee is required based on an agreed-upon number of hours for the entire organization. Our minimum contract under this framework is £200,000 per year for a typical hospital system. This fee covers the license and API calls for processing up to 100,000 hours of audio. Any extra hours will incur an overage fee, and unused hours can be rolled over to the next year's contract.

PRICING OVERVIEW

- See pricing document attached

TRIAL SERVICE

The user can sign up to TORTUS for a free trial of capped usage at 10h per month. To use TORTUS in a clinical setting, the user needs to sign a contract to confirm no misuse. For a larger organisation we allow organisations to have a 90-day no obligation break clause to allow for evaluation time if desired, with only the initial 90-day fee due.

END USER EASE OF USE

TORTUS is a web-app. Clinicians can access the tool on devices with in-built or connected microphones (desktop and mobile).

Without EHR integration

- At the start of a consultation, the clinician will press the 'Consultation' button, followed by the 'Start Recording' button, which is represented by a microphone symbol. The clinician can then proceed with the consultation without the need to take notes.
- During the consultation, the clinician has the option to select the microphone, allowing TORTUS to transcribe the conversation locally.
- Once the consultation is complete, the clinician should press the "Stop Listening" button. After that, pressing the "Finish & Create Documents" button will prompt TORTUS to process the audio and generate a transcription, summary, and letter, which can be accessed through the "Transcript", "Summary", and "Letter" buttons.
- All resulting outputs can be individually copied to an EHR system using the 'Copy to clipboard' button, represented by a double-page icon. Additionally, the summary can be copied to the EHR system section by section using the "Copy to clipboard" button located next to each segment.
- Finally, the 'Start a new consultation' button will clear the local screen, allowing the clinician to begin a new consultation.

With EHR Integration (Epic, EMIS Q1-25)

- In your EHR system, such as Epic, you can access a patient's outpatient record and find the "Launch TORTUS" option located on the left-hand side.

- After clicking the “Launch TORTUS” link, the TORTUS application will open, and a consultation will automatically begin, displaying the patient’s name and date of birth. The clinician should verify that they are seeing the correct patient.
- Once the consultation is complete, the clinician should press “Stop Listening,” and then select “Finish and Create Documents.” TORTUS will generate the most relevant template based on the clinician's selections.
- After reviewing the note and letter, the clinician should click the green button labeled “Approve & Write to Record.” This action will save the note to the patient’s record, indicated by a green tick.
- A confirmation banner will then appear, stating that the note has been successfully added to the record with the message “Saved to Record.”

TRAINING

Users are sent “how to” videos, a digital user guide and are offered demos and set up support - online or in person. The support also includes template generation personalised to the clinician’s specialty and care setting. Where relevant, we offer in person training, and encourage the training of in house “super users” to be champions of the tool.

TECHNICAL HELP

Users have access to the support email inbox, and we aim to reply within 24-hours. Often we triage problems via email, and if not resolved, support via a call and walking through the problem.

IMPLEMENTATION PLAN

Individual User

For individual users who wish to sign up, they can register directly on the website, and an account will be created for them. Once the account is verified by the user, TORTUS can be accessed online. The self-serve onboarding process is supported by user guides and implementation information, such as templates and microphone settings, which will be sent via email. For clinical use, users must obtain the necessary permission from their employer and sign a contract to confirm that they will not misuse the service.

Primary Care and Secondary Care

If you are part of an organization that has signed up with TORTUS, the contract will be signed by your administrator. The implementation plan for trusts and secondary care providers is as follows:

1. The TORTUS IT team will set up accounts under the organization’s name.
2. Clinicians will receive either on-site or virtual training before the accounts are issued. This training will be conducted by the TORTUS implementation team or by designated "super-user" trainers from headquarters.
3. There will be a familiarization period during which TORTUS will work alongside your regular practice to help build confidence in using the tool.
4. Our implementation manager will conduct weekly check-ins with the operations lead to ensure a smooth implementation process.

TORTUS employs an AI-driven toolset that includes a speech-to-text AI model, which is integrated with two large language model steps to create a medical note and a corresponding letter.

Quality Assurance is implemented at each stage of this process, with particular emphasis placed on the clinical entities discussed during the consultation, especially regarding clinical risk. The following aspects are evaluated for quality:

- The quality of the audio file's transcription to text.
- The accuracy of the subsequent medical note, focusing on hallucinations and omissions.
- The accuracy of the final letter, again assessing for significant hallucinations and omissions.

Our dataset consists of 321 transcribed consultations sourced from the Primock-57 and NatureOSCE datasets, covering a wide range of primary care specialties.

We annotate these transcripts for information relating to protected characteristics such as sex, age, disability status, and employment. We extract this information to evaluate potential biases in the large language model's generation of documentation, as well as its factuality and precision. The characteristics we label include patient age, sex, doctor age, sex, disability status, profession, and any patient addresses mentioned.

Speech-to-Text Evaluation

Our speech-to-text (S2T) system is evaluated through three key steps: we start by processing audio, then we refine the transcription, and finally, we check how accurate it is compared to an accurate reference transcript. One of the main ways we measure accuracy is by looking at the Word Error Rate (WER), which reflects how many mistakes are made in the transcription. We also consider how quickly the system produces results.

In our latest evaluation, in October 2024, we used a sample of 10 transcripts from the Primock dataset to compare different speech recognition models, including Google Cloud and AWS. Their WERs were 30.9% and 31.3%, meaning they had a higher percentage of transcription errors. In contrast, our model performed impressively, achieving a much lower mean WER of just 13.8%, alongside faster transcription times compared to AWS.

We also looked at various factors that could influence the accuracy of the transcriptions, such as the speaker's characteristics. Our analysis showed that while certain factors like gender and accent can have some impact on the results, they do not create significant differences on their own. Overall, we are pleased with how well our models performed in terms of accuracy and speed, making them reliable tools for use in clinical settings.

Medical Documentation Accuracy

The next step in the pipeline from the transcript is to produce a medical document, in this case a medical note.

We last evaluated this system in Oct 2024 - to see how well large language models (LLMs) summarize clinical notes using 321 transcribed consultations from various primary care specialties.

Key Data Points:

- ****Transcripts Analyzed:**** 321

Precision (Accuracy of generated information): 97% (with a margin of error of 11%)

Recall (Ability to capture all relevant information): 87% (with a margin of error of 13%)

We also assessed various characteristics such as patient and doctor sex, age, disability, and profession to check for any biases affecting LLM performance. However, none of these factors significantly impacted the accuracy of the summaries.

Deployment Success:

To date, none of our live user feedback has suggested there are any biases from the models. The system has been updated in late Q4 2024 and a repeat study is performed every three months at minimum. TORTUS now also deploy continuous accuracy monitoring of the outputs, and organisational reports can be produced on request.

Risk mitigation:

Despite having very high precision and recall rates, we have implemented risk mitigation strategies. It is essential for users to verify the accuracy of every note, as we place the responsibility on physicians, similar to how administrative staff would draft a letter for review.

We are developing automated systems to further minimize this error rate, which includes:

- Continuously improving existing prompts to reduce errors
- Adding additional AI steps to identify and eliminate hallucinations
- Collecting more ground-truth data validated by clinicians to enhance our capabilities in this area.

Conclusion

We have created a reliable evaluation framework for assessing the quality of the AI tools we currently use. Although our error rates are relatively low, we are committed to continually evaluating and improving the accuracy of all our tools as new vendors and technologies emerge, and as we receive user feedback. Whenever we introduce new tools or vendors, we will establish new evaluation metrics and continue our assessments.

MAINTENANCE AND UPGRADES

As TORTUS is a web-app, we will release on a regular cadence with relevant training and notice given to users. Minor changes will occur more frequently, larger modifications will come with a change management plan.

SERVICE CONSTRAINTS

We currently have two layers of technical redundancy built into every stage of our pipeline - if our external speech to text vendor fails, we have a back-up, and similarly if our external large language model vendor fails, we have a back-up. This architecture is reflected in our DPIA. The user may experience an additional delay in returning outputs (typically 1.5mins up to 5 mins) but should not lose any consultation data. The exception to this is network issues - we cannot control the local network availability and this is up to the user to check before using TORTUS that they have internet access.

TECHNOLOGY REFRESH

The company leadership and clinical team meets on a regular basis to review all planned development issues. New functionality is discussed and assessed with respect to the scope and intended use to ensure appropriate criteria and testing is applied. The risk analysis process is then completed following the risk management plan. All new functionality is subject to SWIFT analysis, with the output stored in the risk management file. If the safety profile is affected, relevant updates are made to the clinical safety documentation by the CSO and approved by the named personnel. As well as new features, software releases typically include many small, iterative changes to functionality.

SERVICE LEVELS

PRIORITISATION OF REPORTS

Clinical safety incidents must be reported to the CSO in accordance with the incident management procedure and documented within the incident log. The clinical safety team will make an assessment of the clinical safety risk and determine the appropriate response, which may include any or all of; notification to all or specific users, notification to all or specific deploying healthcare organisations, revert to previous release, unscheduled down-time of the application.

- TORTUS AI must establish, document and maintain a process to collect and review reported safety concerns and safety incidents.
- TORTUS AI must assess the impact of any such information on the ongoing validity of the Clinical Safety Case.
- Where any such evidence is assessed to undermine the safety case TORTUS AI must take appropriate corrective action in accordance with the RMP and document it in the CSR.
- TORTUS AI must ensure safety-related incidents are reported and resolved in a timely manner.
- A record of safety incidents, including their resolution, must be maintained by the manufacturer in a Safety Incident Management Log.

The purpose of the safety incident management process is to:

- Enable users to report safety incidents they have had, or they consider may have an impact on patient safety.
- Enable TORTUS AI personnel to report incidents they have had, or they consider may have an impact on patient safety.
- Provide a communication mechanism within TORTUS AI, and, where appropriate, the Customer Organisations that are using the product.
- Ensure appropriate and sufficient resources are allocated to manage and resolve reported incidents.
- Enable issuing of safety alerts or bulletins to Customer Organisations and relevant authorities to advise users of potential safety incidents and application of associated mitigations.

To achieve this, the safety incident management process needs to provide:

- A point of contact or mechanism where the incident is logged
- A mechanism by which a clinical risk assessment can be made. In practice, this will be the same criteria defined in the RMP. As part of this assessment, the existing safety case should be reviewed.
 - Does the reported incident constitute a new hazard?
 - Is the incident a realisation of a recorded hazard?
 - Have clinical risk controls failed?
- A mechanism by which the user community can be advised of the safety incident.
- A mechanism through which effective root cause analysis can be conducted. This mechanism will need to consider:
 - the provision of appropriate resource
 - collaboration with appropriate organisations.
- Corrective changes have been approved by the appropriate authority as identified in the Risk Management Plan.
- Appropriate key point indicators to ensure effective management.
- A mechanism by which the user community can be advised that the safety incident has been resolved.

TARGET RESOLUTION TIMES

We currently have two layers of technical redundancy built into every stage of our pipeline - if our external speech to text vendor fails, we have a back-up, and similarly if our external large language model vendor fails, we have a back-up. This architecture is reflected in our DPIA. The user may experience an additional delay in returning outputs (typically 1.5mins up to 5 mins) but should not lose any consultation. We aim to resolve all issues immediately, and rank issues based on priority via severity to the user.

MAJOR INCIDENT

A major incident is declared if TORTUS cannot be used by most users or there is a data breach.

Automatic alerting of incidents from our monitoring tool Datadog and Google Cloud Platform. If data has been breached, we assess the impact, provide a fix, and alert users via email if necessary. We then also work with our Pentest organisation 8Fold Governance to ensure we have correctly patched the breach.

We would isolate users who are affected and send comms to them, and let other users know of the risk while it exists. We would send further comms once resolved.

OUTAGE AND MAINTENANCE MANAGEMENT

MONITORING

We have active monitoring and alerting, once an issue is flagged, those on support will triage the problem and assign it to the appropriate team. We use DataDog for alerting.

CHANGE MANAGEMENT

Changes are managed via ticketed system, and are peer reviewed, and subject to automated and manual testing, before merging, additional testing is performed on a staging environment before promoting to production. Features are guarded by a feature flag and monitored closely. All clinical changes have to go through our clinical evaluation process. We do a phased roll out to users to ensure stability of changes.

FINANCIAL RECOMPENSE MODEL FOR NOT MEETING SERVICE REQUIREMENTS

How are customers able to apply for a refund or compensation for a loss of service? How is the refund or compensation calculated?

Individual customers applying for a refund can apply to support@tortus.ai with the exact dates and times of outages and compensation requirements, and username. These can then be verified on our own cloud system that outages fell below the contracted SLA, and a refund calculated on a daily basis (e.g. 1/30 days affected) and a partial refund given based on number of seats affected within that organisation.

In secondary care, outages need to be reported within 30 days and verified as outside the SLA requirements by emailing support@tortus.ai and receiving a formal acknowledgement.

Compensation is calculated as follows:

For outages between the hours of 8am-5pm:

- Up to 1 hour - no compensation due, recovery plan in place
- 1–4 hours - 50% pro rata compensation due
- 4+ hours - 100% pro rata compensation due

For outages between the hours of 5pm-8am:

- Up to 6 hours outage - no compensation due
- >6 hours outage - 50% pro rata compensation due

PROVISION OF THE SERVICE

BUYER RESPONSIBILITIES

TORTUS is designed for use by clinicians and healthcare professionals during clinical consultations in hospitals and primary care settings that involve verbal communication. Only clinicians who have been granted approved access can register to use this tool. Access can be managed internally or through an authorized administrator who has the ability to add, remove, or edit users.

Users must manually review all data generated by the tool, including summaries and letters, and make any necessary edits before incorporating them into an Electronic Health Record (EHR) or sending a referral letter.

TORTUS can be used as an intelligent dictation system for areas where ambient capture is not possible to, or in addition to for additional patient context. For example with non-verbal patients.

All associated hardware must be properly maintained, and alternative methods for transcribing consultations should be available. Users are also required to apply clinical judgment when determining the content of patient records and letters.

TECHNICAL REQUIREMENTS AND CLIENT-SIDE REQUIREMENTS

Clients require a system running one of the following browsers:

- Chrome (v121 or above)
- Edge (v122 or above)
- Firefox (v124 or above)
- Safari (v17.3 or above)

OUTCOMES/DELIVERABLES

- see the case studies above

AFTER-SALES ACCOUNT MANAGEMENT

All users will have access to an account manager for support, queries, but no formal participation from the account holder required. We will support onboarding, training and implementation.

TERMINATION PROCESS

The user will need to place a written request to end the contract, and once it comes to an end, their user account will have its access revoked. The user will be unable to log into TORTUS and use it.

OUR EXPERIENCE

CASE STUDIES

TORTUS AI published a recent study to evaluate its impact in conjunction with Great Ormond St Hospital. 8 GOSH clinicians carried out 48 simulated consultations with professional medical actors using dummy patient information, in a mock clinic environment. Clinicians completed consultations using current standard Electronic Health Record (EHR) GOSH workflows, which were compared to an equal number of matched consultations carried out using the TORTUS tool. All consultations were carried out in a Proof of Concept EHR environment.

Phase 2 evaluation showed that TORTUS could provide care that was:


- **Smarter**
 - Using TORTUS, above the use of standard EHR capability alone, clinic note quality improved from 43% to 100% of those described as good or very good.
 - Using TORTUS, above the use of standard EHR capability alone, clinic letter quality improved from 29% to 70% of those described as good or very good.
- **Kinder:** Through questionnaire responses, clinicians perceived that interaction with the patient and family was enhanced:
 - “Overall, I felt the patient/family had a good consultation with me today” (81% agreed with EHR alone vs 100% agreed with TORTUS use)
 - “Overall, I was able to give each patient and family my full attention during the consultation today” (66% agreed with EHR alone vs 100% agreed with TORTUS use)
- **Better**
 - Enhanced operational efficiency
 - Appointment duration (defined as the moment a patient entered room to when they left the clinic room):
 - 25.8% reduction in consultation length
 - On average, 3 min 13 seconds was saved from each 12 min consultation

Clinician survey results

- Overall, I felt the patient/family had a good consultation with me today **(81% agreed with EHR alone vs 100% agreed with TORTUS use)**
- Overall, I was able to give each patient and family my full attention during the consultation today **(66% agreed with EHR alone vs 100% agreed with TORTUS use)**
- 100% strongly or somewhat disagreed that computer based tasks were stressful during the consultation (28% with EHR alone)

Patient/carer survey results

- The percentage of respondents who strongly disagreed that ‘The clinician was distracted by the computer...’ **increased from 58% to 86%**
- The % of respondents who strongly or somewhat agreed that ‘The clinician spent enough time with me and my child’ **increased from 92% to 100%**
- The % of respondents who strongly agreed that ‘The clinician gave me and my child their full attention...’ **increased from 75% to 87%**



“What three words would you use to describe using TORTUS?”

Dr Dave Triska, a GP and GP partner from Surrey, explains the impact of incorporating the TORTUS system into his practice.

“Previously, clinic days were exhausting marathons filled with consultations, paperwork, and inevitable mental fatigue at the end of every week. The extensive note-taking required of each consultation was often a distraction from care, and in some cases caused oversight of important details. “

Seeking a solution to improve efficiency and patient engagement, Dave discovered TORTUS, an AI interface that works alongside the clinician on their computer.

This marked a significant shift in Dave’s practice, as it allowed him to fully concentrate on patients without the distraction of manual note-taking. TORTUS not only records conversations but also employs AI to summarise and capture the nuances of these interactions. Including both the clinician’s observations and the patient’s input, it helped to enhance follow-up care and ensure continuity.

One of the major benefits Dave has found is TORTUS’s ability to produce accurate, patient-centred notes that include comprehensive views of the discussions, plans, and differential diagnoses. These notes are structured in simple language, making them accessible and understandable for patients, thereby increasing transparency and trust in the healthcare process.”

He has felt a significant reduction in his day-to-day cognitive load. He allows time to check AI transcriptions for accuracy, but now regularly finishes a session 10 to 15 minutes early. “If you’re taking notes yourself, you are using two different bits of your brain: the memory bit is working really hard to structure and organise information, while your creative thinking and processing bit is trying to come up with novel solutions”. Tasks such as filling out referrals and completing paperwork, which were previously time-consuming, have become more manageable thanks to TORTUS’s precise documentation. This efficiency has led to:

- **Shorter consultation times** - compared to baseline at 15-min consultation, to 11-min consultation after using scribe. The time-save allows him to manage administrative duties without the usual rush, improving the overall productivity.

- **Shorter waiting room times** - compared to baseline at 8-min on average of waiting room times, using TORTUS brought a 50% reduction time in the waiting room to 4 minutes on average.
- **Reduction in overrun consultations** - there was a significant reduction in overrun consultation (ie >25% of the average baseline time) from 15% of the consultations to 9% of the consultations.

Consequently, Dave has noticed the positive effect that these changes have also had on his well-being. The reduction in post-clinic exhaustion has been replaced with a sense of energy and focus, which in turn has improved his quality of life.



CLIENTS



Witley and Milford Medical Partnership

Stanley Medical Centre

Golborne Medical – West Ten GPs
16 Golborne Road, Ladbrooke Grove, W10 5PE



CONTACT

support@tortus.ai

commercial@tortus.ai