



G-CLOUD 14

TORTUS AI O.S.L.E.R.

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 O.S.L.E.R., our AI-enhanced tool that improves interactions between clinicians and Electronic Health Records (EHRs). O.S.L.E.R. uses AI to listen to the consultation, create instant documentation notes and letters, coding and diagnoses, saving clinicians time and improving data quality for systems.

Our agent, O.S.L.E.R., is a combination of AI technologies that mimics human behaviour. In v1 we are taking over documentation, by combining highly accurate speech to text AI with Large language models to create instant notes and letters. In the roadmap we are adding customisation to match your voice and templates for notes. Beyond that we are adding database access, to allow O.S.L.E.R. to summarise a patient history for you. Once you've finished, next year O.S.L.E.R. will also start to be able to use the system for you, placing orders and filing diagnosis. The system is capable of automatically generating codes, and we are working on this for a few UK systems right now.

Our commitment to 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 O.S.L.E.R. in various clinical environments. Furthermore, our initiatives are supported by the Department of Business and Trade, emphasising our alignment with national healthcare and business objectives.

TORTUS was formed 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. Christopher Tan is a machine learning scientist and former research partner at DeepMind and Google Brain, focussed academically on agents in the AI space for the past several years. Both met on the prestigious Entrepreneur First programme (LD19). Team: Scientists, ex Google and Spotify principal engineer, ex Babylon project manager, post grad ML scientists (UCL/Imperial).

VALUE PROPOSITION

O.S.L.E.R. is on track to be fully integrated with existing NHS systems, offering seamless functionality without disrupting established workflows. We are DTAC assured and GDPR-compliant, ensuring that our solutions meet the highest standards of data security and privacy. As a Crown Commercial Supplier, we are positioned to provide value across the NHS with ease of procurement.

Evidence shows that O.S.L.E.R. can reduce administrative tasks by up to 25%, allowing clinicians more time for patient care. This not only improves efficiency but also increases patient throughput—potentially adding over £4 million in additional revenue annually per average hospital. Furthermore, the quality of clinical documentation is doubled, enhancing patient safety and care continuity.

WHAT THE SERVICE PROVIDES

O.S.L.E.R. encapsulates some of the core tools required for the task of recording and summarising a clinical consultation. These tools include Speech to Text, text summarisation and referral letter generation. It is expected to be used in clinic by clinicians.

The physician workflow for any consultation involves taking a verbal history, performing a physical examination and formulating a diagnosis and treatment plan, as well as writing onward correspondence to another physician in most cases. This documentation is performed digitally in >90% of healthcare settings in the UK, as well as the US. Approximately 30-50% of clinician time is taken up with this documentation. O.S.L.E.R. provides 'ambient intelligence' by listening to the verbal history taking, O.S.L.E.R. can then record this directly into text, saving the physician documentation time as the task is done in parallel to the consultation itself. It then automates the drafting of the note and the letter, tasks which take a physician between 2–10 minutes depending on the patient complexity.

We have been measuring the benefits in various forms throughout O.S.L.E.R. development, but thorough testing has shown utilising O.S.L.E.R. in the clinical workflow yields an approximate 20-30% increase in time efficiency and documented increases in documentation quality.

Throughout testing with physicians, performance was compared before and after using the functionality to perform a clinical consultation in multiple simulations. These showed overall time reduced by 55% using the AI vs without. Measuring the quality of documentation produced was performed using a Visual Analogue Scale, a direction measure of O.S.L.E.R. functionality. Documentation quality was consistent and improved above baseline by 53%. Cognitive load was measured by the NASA TLX scoring system. Using this framework, we assessed AI use vs. non-AI use and found AI-assisted consultation cognitive burden was reduced by 55%.

Anecdotal feedback was also positive following the product launch in Canada in 2023, with one physician noting it saved enough time to enable a lunch break for the first time in her career.

Early studies against similar products show clear benefits in terms of time saved, cognitive burden and patient-doctor quality of interaction.

Key features and benefits: *describe the features and benefits of each service component and any customisation possible.*

Key features include:

- Speech to text AI
- Instant medical note generation
- Customisation of templates
- Instant referral letter generation
- Instant patient communication content generation
- Quick, accurate dictation

Key benefits include:

- Reduced administrative burden
- Increased clinician time
- Decreased clinician cognitive load
- Improved clinician experience
- Double the quality of documentation
- Increased cost capture
- Increased consistency of documentation for audit
- Improved coding
- Better patient experience

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 also have an evaluation platform which is leveraged by over 100+ clinical volunteers. Whenever we are considering a change to our product, the changes go through the evaluation process to ensure hallucinations and omissions don’t occur - these are checked by the clinician volunteers, which also means we are

PRODUCTS IN DETAIL

The O.S.L.E.R. desktop application is a cross-platform application built with the Tauri framework. This consists of a Rust backend and secure, locally-hosted TypeScript frontend. The backend and frontend communicate with each other over a secure IPC layer.

The O.S.L.E.R. 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 O.S.L.E.R. API call additional services hosted on AWS, Azure and OVH cloud, 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 9-6pm, Monday–Friday.

FAULT REPORTING

Faults can be reported via contact to the support team with any questions or queries, responsive from 9-6pm, 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 9-6pm, Monday–Friday

INFORMATION ASSURANCE

We are certified to CSA CCM version 3.0 and ISO/IEC 27001. We have a CSO, clinical risk manager, 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 this 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. We have a Head of Engineering, who manages a tech lead, and senior engineers and junior

engineers all under this umbrella. The product manager works with the compliance manager, CSO, clinical risk manager, compliance manager.

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.

BUSINESS CONTINUITY STATEMENT/PLAN

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 O.S.L.E.R. 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

O.S.L.E.R. stores key functional data such as user profile information and consultation history. Audio blobs that are sent for transcription are only stored in memory during the transcription process, with metadata (audio length) saved in a database. No sensitive information is ever stored by us or our providers.

USING THE SERVICE

ORDERING AND INVOICING

An organisation wide contract is put in place between the organisation and the Company - subject to any free or pilot trial phase. For primary care - invoices are raised monthly or annually and dependent on number of full-time seats.

For secondary care or larger organisation - an annual licence is paid up to an agreed number of hours for the organisation as a whole. For example, £500,000 a year for a typical hospital system up to 250,000 hours of audio, with an overage fee and/or rollover of additional hours to the next year contract amount.

PRICING OVERVIEW

O.S.L.E.R. can be paid for:

- Monthly with unlimited consultations for **£200 a month.**
- Agreed annually with unlimited consultations for only **£100 a month, with the following added benefits:**
 - Free recommended mics (RRP £249.99)
 - 200m range
 - Noise cancelling and highest quality
 - Wearable
 - Early access to new features such as:
 - Coding
 - Actions (labs, ERS)
 - Document triage
 - Bespoke onboarding
- For secondary care, bespoke pricing dependent on activity will be formulated with a fixed cost annual licence with capped usage.

TRAIL SERVICE

The user can sign up to O.S.L.E.R. and download the app for a free trial of capped usage. To use O.S.L.E.R. in a clinical setting, the user needs to sign a contract to confirm no misuse.

END USER EASE OF USE

O.S.L.E.R. is a desktop app, and has to be downloaded to be used. Clinicians can access the tool on devices with in-built or connected microphones.

- At the start of a consultation, the clinician will press 'Start a new consultation', then press the 'Start Recording' button, a microphone symbol, before commencing with the consultation without the need to take notes.
- The clinician can now choose the microphone, and O.S.L.E.R. will transcribe the consultation locally.
- Once the consultation has concluded, the clinician can press the 'Stop Recording' button.
- The clinician can edit the transcription if their recording microphone did not pick up certain parts of the consultation.
- O.S.L.E.R. will then process the audio, and generate a transcription, summary and letter, accessed by the 'Open Transcription', 'Open Summary', and 'Open Letter' buttons.
- All the resulting outputs can be individually copied across to an EHR system using the 'Copy to clipboard' button (marked by a double page icon).
- The 'Start a new consultation' button will clear the local screen, allowing the clinician to start again.

TRAINING

Users are sent how to videos, and are offered demos and set up support - online or in person.

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

The user gets a link to download O.S.L.E.R., an account will be created, the user will need to verify their account.

If using O.S.L.E.R. in a clinical setting, the user will need to sign a contract to confirm no misuse. If the user is part of an organisation that is signing up, their administrator will sign the contract, so each individual user does not have to. The users can then download O.S.L.E.R. via a link, which they can access via their username and password. After downloading, they can log into O.S.L.E.R. and start using it immediately.

SERVICE MANAGEMENT

PERFORMANCE

O.S.L.E.R. utilizes an AI toolset that consists of a speech-to-text AI model and combines the output of this with two large language model steps to produce a medical note and letter respectively.

Quality Assurance takes place at each step in this chain, specifically, with specific weight given to the clinical entities discussed in the consultation in terms of clinical risk.

1. The quality of the audio file to text transcript produced
2. The accuracy of the subsequent medical note, particularly around hallucinations and omissions
3. The accuracy of the subsequent letter produced, again looking for significant hallucinations and omissions.

Speech-to-text evaluation

We curated a set of ground-truth pairs of transcripts and clinical audio (consultation between a single patient and single doctor) from published literature sources. We then internally verified the ground-truth transcript for clinical accuracy with qualified clinicians to produce a dataset of verified transcripts.

Not only that, but we then passed the same audio files to a series of speech-to-text models (AWS Bulk, AWS Real-Time, and Gladia) to evaluate the accuracy of the ground-truth human transcript versus the AI models.

Word error rate:

On 5th Dec 2023 we reevaluated all available bulk API S2T models-initially using a crude word error rate derived from the Levenshtein distances from the ground-truth (this is a metric where the measure is similarity between the characters in the ground-truth text and the transcript) - with the above results. Initial results showed comparable performance between Whisper and Gladia API and improved performance with AWS against previous benchmarks. Latency was comparable between Gladia, AWS and Whisper but much longer (and therefore not practical) with GCP at the time of testing.

Clinical entity matching

To go beyond standard word error differences which aren't clinically significant, we created a metric derived around clinical entities, by using Amazon Comprehend Medical.

Amazon Comprehend Medical is an API NLP service that extracts clinical entities from an unstructured text, including demographics, personal information, medications, conditions, symptoms etc.

We ran each human transcript through this service to generate a list of ground-truth clinical entities. We then ran each AI transcript through the same service to generate a list of AI clinical entities. We then identified which transcripts were omitted or added in the AI transcript versus the ground-truth.

Entities were then classified into True Positive, False Positives, True Negatives and False Negatives in the following table:

	Present in AI Transcript	Absent in AI Transcript
Present in Ground-Truth	True Positive	False Negative
Absent in Ground-Truth	False Positive	True Negative

And from this Precision (True Positive/True Positive+False Positive) and Recall (True Positive/True Positive+False Negative) were calculated and indexed into a single F1 score, which was then averaged across all speech-to-text models.

Generally, all models performed well at clinical entity capture versus ground truth. The average performance of AWS Real-Time was 0.88 (F1), while Bulk Transcription performed better (AWS Bulk - (5/12/23–0.905) and Gladia in this test suite was the most performant (5/12/23–0.93).

Decision:

Based on this, we can conclude that speech-to-text accuracy is high, and we decided to use Gladia as our primary model with AWS as a secondary redundancy.

Medical Note Accuracy

The next stage in the process is to convert the transcript into a medical note. This involves using a large-language model which is guided in it's task by a 'prompt' to generate a medical note that includes the relevant information.

At TORTUS we use a standardized SOAP note prompt to give a template and style guidance to the large language model (GPT-4 32K) to summarize the information in the transcript into the relevant headings.

In this setting, we think about errors using the same confusion matrix:

	Present in Medical Note	Absent in Medical Note
Present in Transcript	True Positive	False Negative (Omission)
Absent in Transcript	False Positive (Hallucination)	True Negative

Ground-truth transcripts were then converted into SOAP notes using our standard prompt and evaluated by hand by a clinician for hallucinations. This means reviewing each clinical entity in the SOAP note and finding the corresponding source in the transcript where that entity could be reasonably attributed.

Hallucinations are categorized clinically into minor and major:

- Minor hallucinations - do not alter clinical management in the context of diagnosis, investigation, treatment or prevention of disease.

- Major hallucinations - would alter clinical management in the context of diagnosis, investigation, treatment or prevention of disease if not corrected by the physician.

19 examples were evaluated by a human clinician for evidence of hallucinations. Across all examples, 5 notes contained at least one minor hallucination, and 2 notes contained at least one major hallucination. With an average of 50 clinical entities per note, this gives an average minor hallucination rate of 0.5% and an average major hallucination rate of 0.21%.

5th Dec 2023: The SOAP note was then updated to reduce these hallucinations and then new rate was [TBD]

Risk mitigation:

Despite the very low numbers of hallucinations per note, risk mitigation strategies are in place.

- User must always check every note for accuracy - we pass responsibility to the physician in the same way any secretarial or admin support staff would draft a letter to be reviewed
- We are building automated systems to iterate and reduce this error rate further, including:
 - Modifying the existing prompts to reduce errors
 - Adding additional AI steps to identify and remove hallucinations
 - Collecting more ground-truth data with clinician validation to further our abilities here

Letter accuracy

The same methodology was then applied to letter production.

With the previous architecture of passing the transcript directly to the letter minor hallucinations occurred in 3/3 initial tests. This was resolved by passing the SOAP note to the transcript as the source instead. The final results showed a minor hallucination in 2/23 examples and a major hallucination rate of 0/23 example. With an average of 30 clinical entities in each letter, this gives a final average hallucination rate as a percentage of total clinical entities to be 0.17% and 0% respectively.

Conclusion

We have developed a repeatable and consistent evaluation framework for reviewing the quality of the AI toolsets we currently use. Although error rates are relatively low already, we will continue to evaluate and improve the accuracy levels of all of our tools as new vendors and technologies become available, and new user feedback is garnered. As and when we add new tools or vendors, we will establish new evaluation metrics and continue to evaluate.

MAINTENANCE AND UPGRADES

We can push upgrades to users via the app, and the user can accept these and install them at a convenient time for them. We will notify users of new features and upgrades so they can plan to upgrade.

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. 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 O.S.L.E.R that they have internet access.

TECHNOLOGY REFRESH

The RMT 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 user.

MAJOR INCIDENT

A major incident is declared if O.S.L.E.R. 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

As a start-up, we don't have a change management process - any change to resolve a bug or issue will be resolved immediately, the Head of Engineering will review all code submitted, the process itself is less than 1 hour.

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 be 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

O.S.L.E.R. is to be used by clinicians and health professionals during outpatient hospital clinical consultations and primary care settings where a verbal consultation takes place. Only clinicians with approved access can be signed up to access the tool. Either the access is managed in-house, or via an authorised administrator with access to add/remove/edit users.

The user should check the required outputs before verifying any transcription when copying into an EHR or sending a referral letter.

O.S.L.E.R. is not to be used in inpatient settings, is not appropriate for consultations where non-verbal communication is the mainstay (e.g., sign language), and is not intended for autonomous use.

Associated hardware must be adequately maintained, with other methods of practice available to transcribe the consultation.

Users must perform a manual check of all data and edit as required (transcription, summary and generated letters).

Users must also apply clinical judgement to determine the content of patient records and letters.

TECHNICAL REQUIREMENTS AND CLIENT-SIDE REQUIREMENTS

Clients require a system running either Windows 10 OS or later, or macOS 11 or later.

OUTCOMES/DELIVERABLES

Outcomes - less burnt out clinicians, less time spent on admin, taking notes and writing lessons, increased clinician welfare and increased patient care.

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.

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 not be asked to delete the app, but they will be unable to log into O.S.L.E.R. and use it.

OUR EXPERIENCE

CASE STUDIES

Great Ormond Street Hospital (GOSH) is working with TORTUS AI to develop AI tools to support clinical care. 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 O.S.L.E.R. tool. All consultations were carried out in a Proof of Concept EHR environment.

Phase 2 evaluation showed that O.S.L.E.R. could provide care that was:

- **Smarter**
Using O.S.L.E.R, 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 O.S.L.E.R., 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 O.S.L.E.R. 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 O.S.L.E.R. 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
- **Safer**
Automatic ICD-10 coding of Medical Note
This AI capability facilitates the automatic generation of a ‘coding section’ at the end of each clinical note. This supports operational grip/enhances the depth of the electronic medical record (GOSH believe the Problem List to be the most important structured field, as it allows clinicians to create a clinical narrative, anchors decision support and supports ongoing care delivery).

Clarity section
This AI capability is able to automatically support the generation of a Clarity section, which provides patients and their family with a lay explanation of the medical terms


used in their letter, making patients active, engaged partners in their care management.

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 O.S.L.E.R. 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 O.S.L.E.R. use**)
- 100% strongly or somewhat disagreed that computer based tasks were stressful during the consultation (28% with EHR alone)

Patient/carers 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 OSLER 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 OSLER, 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. OSLER 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 OSLER’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. Tasks such as filling out referrals and completing paperwork, which were previously time-consuming, have become more manageable thanks to OSLER's precise documentation. This efficiency has led to shorter consultation times, allowing him to manage administrative duties without the usual rush, ultimately improving his overall productivity.

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



CONTACT

support@tortus.ai