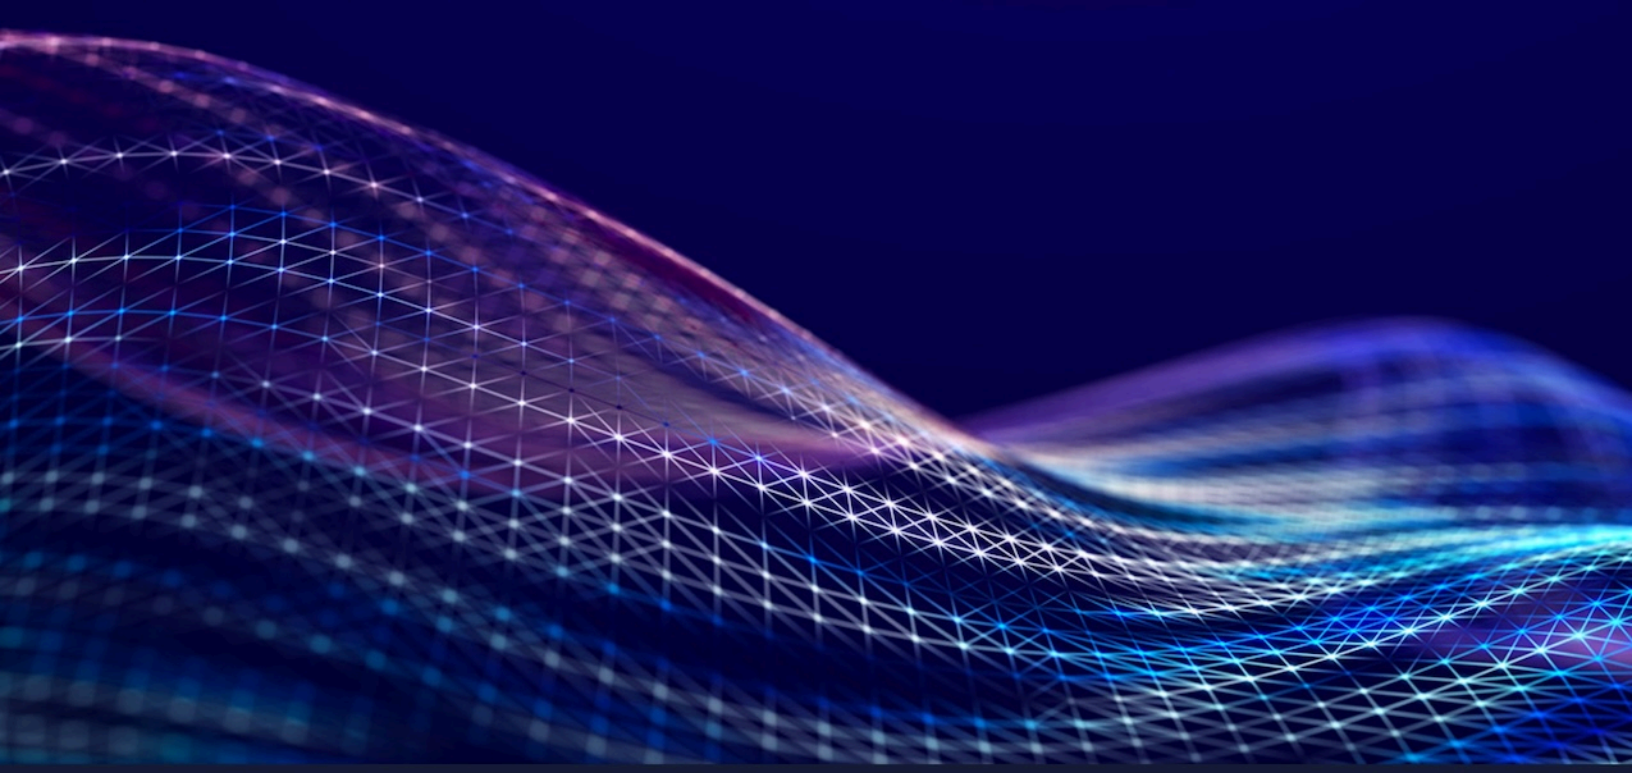

Tungsten TotalAgility Cloud
InfoSec Primer

TUNGSTEN AUTOMATION

Revision history

Date	Description
August 20, 2025	Adjusted image file
August 15, 2025	Updated to align with TotalAgility documentation set
October 13, 2025	Updated region failover content
October 14, 2025	Added section on AI model upgrades
November 7, 2025	Added US as Tungsten Document Library service location
November 17, 2025	Updated FAQ and business continuity content
December 2, 2025	Revised references to Azure content filtering and Azure Document Intelligence version
December 19, 2025	Added filtering details for Microsoft Azure Guardrails
February 17, 2026	Updated content for the Azure AI services section, and updated images
April 3, 2026	Updated Azure Document Intelligence section with reference to hosting regions

Table of Contents

Chapter 1

Executive summary

TotalAgility Cloud is a fully managed, hosted version of Tungsten TotalAgility offered as a Software-as-a-Service (SaaS) subscription model running on the Tungsten Cloud.

The Tungsten Cloud hosts all Tungsten Automation SaaS-based offerings and uses Microsoft Azure as the underlying hosting infrastructure. The Tungsten Cloud provides Tungsten Automation customers with a range of products and solutions to enable organizations to configure, deploy, and manage Intelligent Automation solutions to help solve use cases across the enterprise.

The Tungsten Cloud includes SaaS options for most of the leading Tungsten Automation products, including TotalAgility. With more than 8000 customers across these products, millions of documents and workflows are processed in the Tungsten Cloud every month.

Tungsten TotalAgility is a low-code platform with capabilities that enable organizations to improve efficiency, reduce costs, and increase operational scale based on the following primary pillars.

- **Document Intelligence**: Applies cognitive capture and artificial intelligence to structured, semi-structured, and unstructured content to automate and extract information and unlock data insights.
- **Connected Systems**: Brings together all the critical business systems, enterprise applications, legacy systems, mobile interactions, chatbots, and more across internal and external business processes.
- **Process Orchestration**: Orchestrates digital workflows in collaboration with users, systems, and data to gain real-time insights and greater governance across the digital workforce.

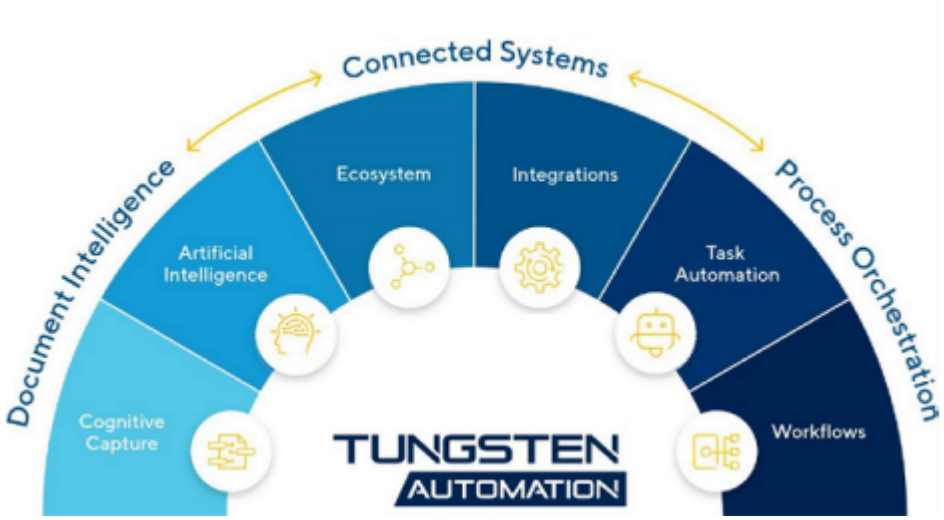

Chapter 2

TotalAgility Cloud technical overview

This chapter includes a technical overview of the TotalAgility Cloud, including information about:

Microsoft Azure hosted deployment

TotalAgility Cloud is a secure public Microsoft Azure-based deployment of Tungsten TotalAgility. TotalAgility Cloud customers use a shared infrastructure (computer, networking) but have their own dedicated tenant to ensure complete data segregation between deployments (storage, database).

Each tenant can have multiple environments with a separate URL for development, testing, and production. With an additional upgrade, clients can deploy on a "dedicated instance" with isolated infrastructure that is not shared with any other clients.

TotalAgility Cloud includes all the features available with the on-premises solution, along with cloud-specific components that leverage features such as elastic scalability. TotalAgility Cloud enables organizations to quickly deploy Intelligent Automation solutions in the cloud. TotalAgility Cloud can also be connected to an organization's on-premise or cloud-based systems of record. All of this is available in a low- or no-code, web-based environment. TotalAgility Cloud leverages the power of the Microsoft Azure security model to safely process and store sensitive customer information.

Tenant provisioning

TotalAgility Cloud is provisioned as a tenant in a multi-tenant, public cloud offering. By request and with an additional upgrade, TotalAgility Cloud can be deployed on a single-tenant Microsoft Azure dedicated Virtual Network (VNet). TotalAgility ensures logical separation between network zones, which includes the corporate network containing Development and QA environments, and the TotalAgility Cloud Production Network. The customer resides solely in the TotalAgility Cloud Production Network, which has no connection to the Tungsten Automation Corporate Network. Within the TotalAgility Cloud Production environment, Tungsten Automation deploys both a "Live" and "Dev" instance to host a customer's tenant.

Data in transit, at rest, and its processing flow

Connectivity to the TotalAgility Cloud uses TLS 1.2 over HTTPS port 443 and does not require any other ports to be open on the network. Once inbound traffic is received by the Azure Web Roles, it flows to the Azure Worker Roles for processing. Azure Worker Roles may pull from and temporarily store data in the dedicated Azure SQL Database, Azure Table Storage, and Azure Blob Storage devices. All data at rest is encrypted using AES-256 at the database level. Once Azure Worker Roles process data, customer data is deleted from Azure Table and Blob Storage in accordance with the timelines and policies defined by the customer within the TotalAgility application.

Information security and accessibility

Tungsten Cloud Services accesses the TotalAgility Cloud environment via the cloud portal, which is restricted based on IP address. Microsoft Azure Active Directory roles are used to manage security permissions within the TotalAgility Cloud environment itself. Virtual firewalls are used to apply rules for managing traffic within the TotalAgility Cloud environment, including inbound and outbound rules applied to the VNet.

The TotalAgility Cloud environment uses a variety of security tools, including a Web Application Firewall (WAF). Antivirus protection is installed in the production environment and acts as the next-generation anti-malware product, providing cloud security analytics and threat intelligence. Tungsten Automation performs vulnerability scans on the Microsoft Azure platform monthly, using a third-party scanning service. We are currently implementing solutions to perform authenticated internal vulnerability scans against the TotalAgility Cloud environment. Tungsten Automation also utilizes Trusec, a third-party service, to provide 24/7 monitoring and alerts on logs to detect potential cyber threats within the environment.

Tungsten Automation product development practices

Tungsten Automation performs application development using an agile development methodology that requires multiple levels of review and approval. This methodology also includes static and dynamic code scanning using Veracode, Burp Suite, and Qualys to ensure that code passes quality and functionality testing. Each customer is notified prior to an upcoming release, and updates are pushed to the production cloud environment.

Chapter 3

TotalAgility Cloud architecture

This chapter includes key details related to Microsoft Azure and third-party sub-processors.

High-level TotalAgility Cloud architecture on Microsoft Azure

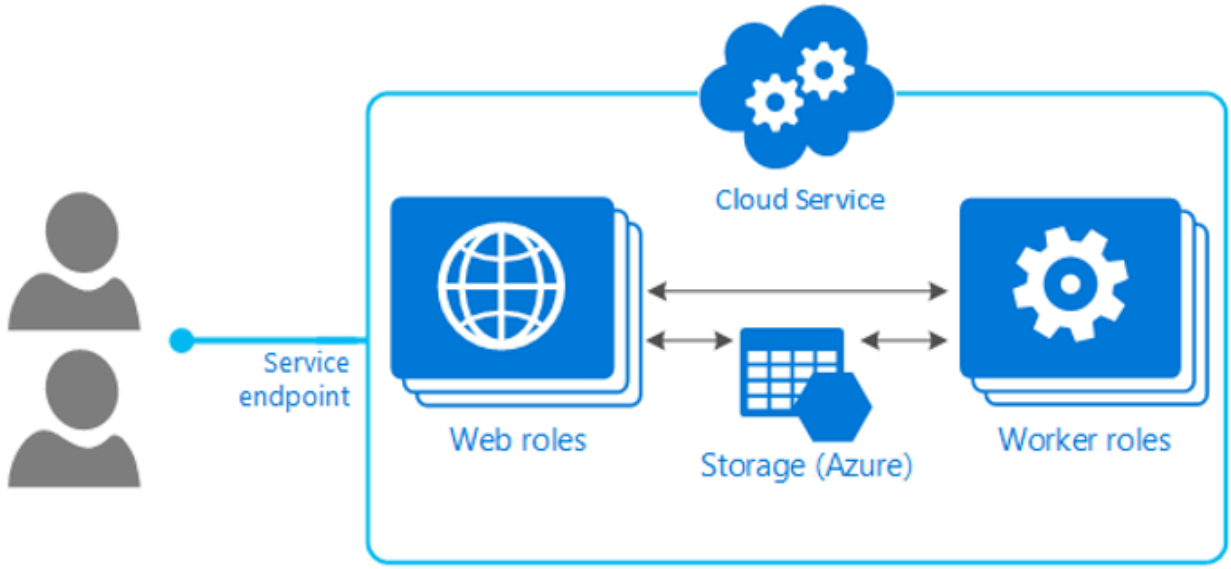

TotalAgility Cloud deployment on Azure with Availability Zones

TotalAgility Cloud is deployed in Azure Regions with Availability Zones to provide Zone Redundant Storage. The architecture diagram in this section shows a typical TotalAgility Cloud deployment in Azure.

In addition to zones, TotalAgility uses cloud services, web roles, and worker roles as its underlying architecture. Azure Document Intelligence, Azure OpenAI, and Azure AI Search are leveraged within the product to provide key functionality. Paired regions are used for Geo Redundant Storage. The Business Continuity section explains how this architecture provides resilience and high availability.

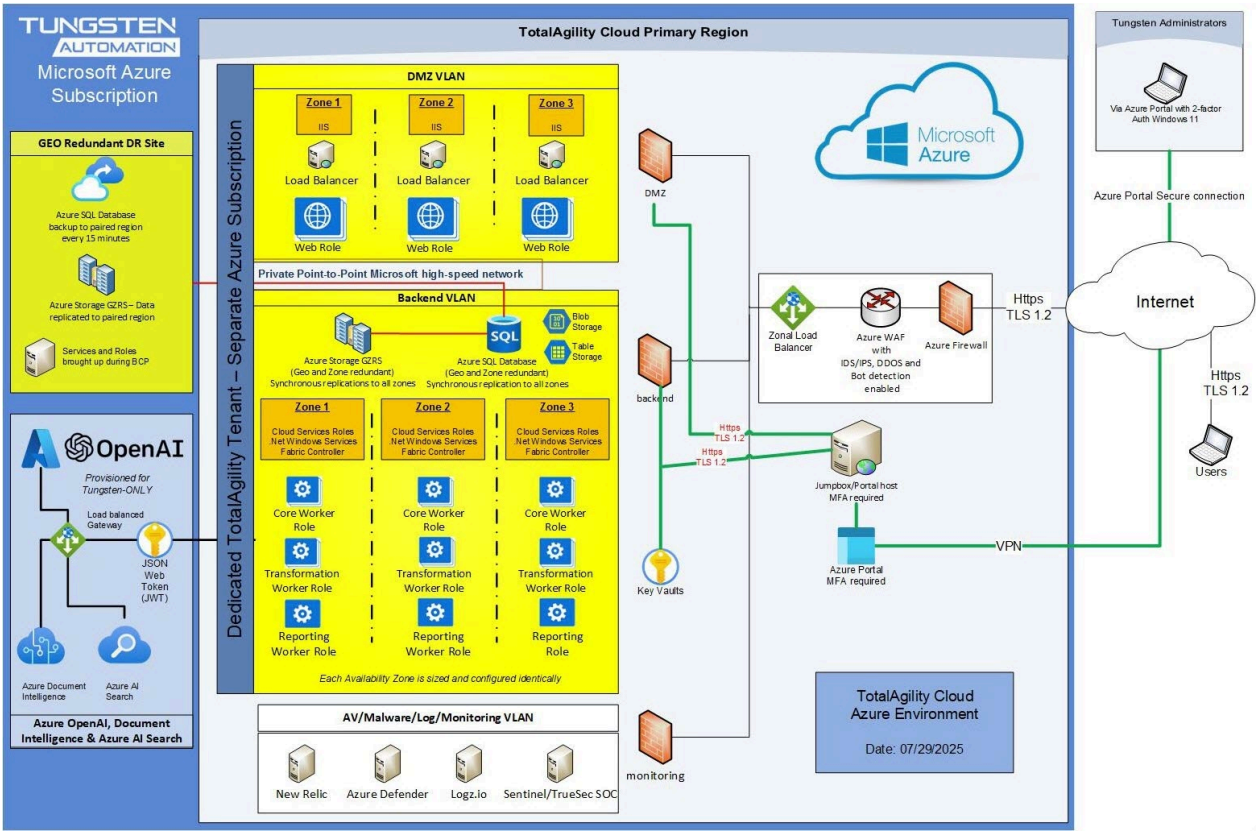

Microsoft Azure Cloud Services and TotalAgility Cloud

Three Cloud Services make up a deployment: live, dev, and the Tenant Management System. The cloud service is broken down into separate roles, including the web role and the worker roles.

Multiple worker roles make up the core components of TotalAgility. Each role contains a minimum of two instances in the live environment to support redundancy and scalability. The number of instances for each role is dynamically scaled to meet the processing demands of the environment.

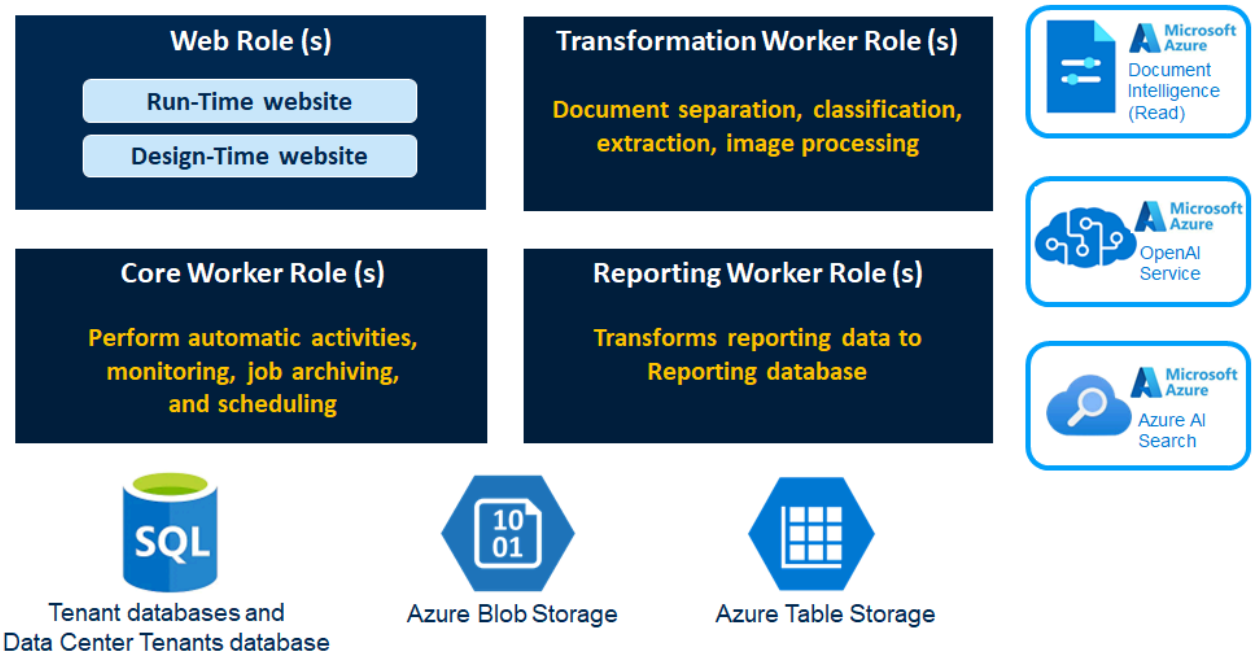

Azure SQL

Azure SQL is used to hold all databases required to operate TotalAgility. All databases are contained within an elastic pool, allowing ease of management, and scaling to accommodate additional load.

Azure table storage

Azure table storage is used to store transactional data. A single storage account is shared among tenants of an instance by default. If necessary, tenants can be separated into dedicated storage accounts.

Azure blob storage

Azure blob storage is used for documents and images that are processed through the TotalAgility solution.

Azure Document Intelligence

Azure Document Intelligence is an AI service is used for extracting insights and information from unstructured document data using advanced machine learning and natural language processing.

This automates document processing, extracts key information, and enables better decision-making.

Azure OpenAI

Azure OpenAI is used in your TotalAgility workflows to extract crucial information from documents, draw inferences, and more. This capability supports enhanced decision-making and more informed, efficient, and effective business operations. TotalAgility also supports other Generative AI providers such as OpenAI and custom LLMs.

Azure AI search

Azure AI Search is used to provide the Knowledge Discovery Service for Tungsten TotalAgility. This service allows users to search across multiple documents for answers to natural language queries within their solutions. Customers can create a knowledge base by uploading various documents that are stored in Azure Search. This knowledge base can then be queried using natural language.

Included environments

TotalAgility Cloud comes with three environments by default. Additional environments can be purchased.

Access to these environments is available only through the URLs listed in the table. For security and compliance purposes, no access is granted to the underlying file system.

Environment	Purpose	Associated URL
Non-production	Development	`https://<customer>-<env>.<totalagilityregion>.dev.tungstencloud.com/`
Non-production	User Acceptance Testing (UAT)	`https://<customer>-<env>.<totalagilityregion>.dev.tungstencloud.com/`
Live	Production processing	`https://<customer>.<totalagilityregion>.tungstencloud.com/`

Virtual Network

A single virtual network is used for all Cloud Service components. The virtual network includes subnets for production and development, with all development environments sharing the same subnet (as they also share the same Cloud Service). Deployments that span multiple regions (for high availability) use separate virtual networks for each region.

For an additional cost, Tungsten Cloud Services offers a "Dedicated Instance" of TotalAgility, with no shared infrastructure or networking.

Microsoft Azure regions

TotalAgility Cloud is available in the United States, hosted at:

- Primary: US West Region
- Backup: US East Region

TotalAgility Cloud is available in Canada, hosted at:

- Primary: Azure Canada Central (Canada - Toronto)
- Backup: Azure Canada East (Canada - Quebec)

TotalAgility Cloud is available in Europe (EU), hosted at:

- Primary: Azure Europe North (Ireland)
- Backup: Azure Europe West (Netherlands)

TotalAgility Cloud is available in the Asia Pacific, hosted at:

- Primary: Azure Australia East (New South Wales)
- Backup: Azure Australia Southeast (Victoria)

For TotalAgility Cloud deployments in other regions, contact the Tungsten Automation Account team or Product Management for the latest status and roadmap.

Third-party sub-processors

In addition to the software and services from Tungsten Automation subsidiaries and Microsoft Azure that provide the core functional capabilities for the TotalAgility Platform, some third-party sub-processors are used to provide specific functionality to extend the core platform, and to process information related to monitoring or support activities.

Tungsten Automation Cloud Services uses Salesforce to store support case data and track tickets related to managing the cloud environment. This data may include personal data about people involved in the case.

New Relic is used to track page load times and other web serving performance metrics (no client or personal data is collected, analyzed, or stored for this monitoring).

Depending on the product configuration, additional sub-processors may be used:

- Google Cloud Platform (if using Tungsten Clarity), specifically the Cloud Vision API.
  Tungsten Clarity can be configured to use US- or EU-based data processing.
- AuthenticID (if using Tungsten ID Verification), which is hosted on AWS.
  Tungsten ID Verification can be configured to use US- or EU-based data processing.

- Base64 (if using specific Document Libraries), see below for details of data processing in Document Libraries. The processor and data processing location options depend on the selected library/model.

For more information about sub-processors, refer to the data processing addendum provided with your contract.

Chapter 4

AI services in TotalAgility Cloud

Overview of AI services in TotalAgility

Customers have access to a range of AI services, algorithms, and capabilities within the TotalAgility Platform.

- Pre-trained Tungsten models (for OCR and Computer Vision) included in the software subscription, available for both self-hosted installation and Tungsten deployments.
- Tools so customers can easily train their own models using their own documents and data (for classification, separation, data locatio,n and data extraction from documents and scanned images).
- Pre-trained third-party models (for OCR, NLP, and Sentiment Analysis) included in the software subscription, available for both self-hosted installation and deployments hosted by Tungsten Automation.
- Access to third-party, cloud-hosted models (for OCR and Generative AI LLMs) included in the TotalAgility subscription.
- Optional third-party cloud-hosted services to extend the TotalAgility platform (pre-trained Document Libraries and OCR).
- Integration connectors for third-party or customer-hosted models (for Generative AI LLMs).

In all cases, the client has complete control over whether, what, and when data is passed to an AI model for processing. The client is the "data controller" with TotalAgility providing the data processing platform.

The following table provides a high-level summary of machine learning, computer vision, generative AI, and trainable algorithms available out of the box. Dependencies on Tungsten Automation-hosted or third-party Cloud Services are marked with an asterisk (*). All other models are included within the software (including those that are self-hosted in a customer environment).

Feature	Algorithms	Proprietary or third-party	Tungsten Training Data	Intended use
OCR	Various neural networks, including CNN	Tungsten Proprietary (OmniPage), optional third-party Azure Document Intelligence*, Google Vision*)	Proprietary, anonymized data	Read text from image documents.

Feature	Algorithms	Proprietary or third-party	Tungsten Training Data	Intended use
ICR	Various neural networks, including CNN	Third-party (ParaScript FormXtra, Google Vision*)		Read handwritten text from image documents.
Suggested Fields	Neural Network based on Yolo	Tungsten Proprietary	Proprietary, anonymized data	Provide the user with a set of potential fields to extract from a document.
Keyword-Value-Pair, Check Mark, and Table Locators (Computer Vision)	Neural Network based on Yolo	Tungsten Proprietary	Proprietary, anonymized data	Extract fields using keywords based on document layout. Locate and extract the check mark and table data.
All Copilots and GenAI Activities	LLM	Third-party (Azure OpenAI*)	How OpenAI foundation models are developed	Accelerate configuration of automation solutions.
Auto Extract Locator	LLM	Third-party (Azure OpenAI*)	How OpenAI foundation models are developed	Extract fields from documents using human language.
NLP	Various deep learning algorithms, such as CNN	Third-party (InMoment)		Extract named entities, sentiment, and summaries from documents.
Document Libraries (optional add-on)	KNN, Bayesian classifiers, LLM, Yolo-based NN	Tungsten Proprietary* or third-party (Microsoft*, Base64*)	Proprietary, anonymized data	Prebuilt document extraction models for common document types
Document Fraud Detection (optional add-on)		Third-party (Resistant.ai*)		Apply Document Forensics to PDF and other files to detect indicators of fraudulent behavior or editing.

All third-party suppliers must complete an extensive Supplier Due Diligence process and meet the requirements of our Secure Software Development Lifecycle methodologies and policies.

Tungsten AI and GDPR

Different types of models have different implications and considerations for GDPR and the general treatment of PII/PHI. For more information on these types, see Overview of AI services in TotalAgility.

- Pre-trained Tungsten models.
  - These proprietary models are pre-trained with Tungsten Automation data. Customer data is not used for these models, as these models are static and not further improved at runtime. Refer to Tungsten Automation: Our Commitment to Responsible AI and Compliance for more information on how Tungsten Automation manages data securely and privately.
- Tools that allow customers to train their own models using their own documents and data.
  - All document samples used with customer-trained models are provided by the client when the models are configured. Samples are then (optionally) updated using new samples as human operators correct and validate the algorithm's results. "Human in the loop" feedback is used to improve model training.
  - Tungsten Automation has no visibility or access to the models created in the TotalAgility platform by customers, or any access to the data sets used to train the models. The client is the data controller for these models and is responsible for ensuring appropriate legal permissions are obtained for all content included in the training sets, and for the ongoing maintenance of these training sets.
  - The platform provides the option to encrypt both the models and training datasets, so no data is stored in plaintext.
  - We recommend the following best practices:
    - Frequently retrain your models with current data/documents and remove old data from the model. This practice not only prevents model drift but also assures that most data comes from newer customers who are less likely to exercise their right to be forgotten.
    - Keep track of end customer IDs associated with each training sample, so you can find the documents and remove them from the applicable training set if any end customer executes the right to be forgotten.
- Pre-trained third-party models included with TotalAgility, and pre-trained third-party models accessed through a cloud service, including LLMs.
  - These models are pre-trained with data from the third party. Your data is not used to retrain them, as these models are static. Tungsten Automation assesses third-parties in detail when onboarding, assuring that data is managed privately and securely.
  - Often, model cards are available from the third-party vendor if more details are needed.

Tungsten AI and other regulations

In addition to GDPR, Tungsten AI is compliant with the following AI regulations.

- European Union: Artificial Intelligence Act
- California AB 2013: Artificial Intelligence Training Data Transparency
- California SB 942: California AI Transparency Act

For information about compliance with other AI regulations, refer to *Tungsten Automation: Our Commitment to Responsible AI and Compliance*.

European Union: Artificial Intelligence Act

The Artificial Intelligence Act categorizes AI systems or AI models into different risk levels to ensure that appropriate regulatory measures are applied. These categories include:

- **Unacceptable Risk:** AI systems that pose a clear threat to safety, rights, or ethical standards and are therefore prohibited.
- **High Risk:** AI systems that require strict compliance measures due to their significant impact on safety or fundamental rights. High-risk AI systems include those used in critical areas such as education, employment, law enforcement, and healthcare. These systems are subject to rigorous oversight and regulation.
- **Limited Risk:** AI systems with moderate impact are subject to transparency obligations and other minimal requirements to inform users about their operation.
- **Minimal Risk:** AI systems that pose little to no risk and are largely exempt from regulation, allowing for more flexibility in their deployment.
- **General Purpose AI (GPAI) models:** GPAI models designed for broad applications across various domains, often in the form of Large Language Models (LLMs). These GPAI models require the model provider to be responsible for implementing transparency and accountability, such as drawing up technical documentation or providing information for the downstream provider. The requirement also applies to those who fine-tune third-party GPAI models.

This framework helps ensure that AI technologies are developed and deployed responsibly, with appropriate safeguards based on their risk level.

TotalAgility falls under the **Limited Risk** category according to the Artificial Intelligence Act. TotalAgility AI models do not qualify for the Unacceptable Risk category, as they pose no threat to safety, rights, or ethical standards; nor do they function in areas covered by the High Risk category, such as education, employment, law enforcement, or healthcare. Additionally, Tungsten Automation does not provide the GPAI model provide or fine-tune third-party GPAI models.

Instead, TotalAgility operates within a specific, moderate-impact context that demonstrates a commitment to meeting transparency guidelines and minimal compliance regulations. This approach ensures responsible deployment without an extensive regulatory burden. AI system providers are required to make it transparent when a human interacts with AI, if not obvious. Tungsten Automation is committed to ensuring that the use of AI is obvious or explained in the product documentation.

In some cases, visuals in the user interface indicate that AI is at work. Also, providers are required to ensure that the outputs from AI systems are marked in a machine-readable format and detectable as artificially generated or manipulated, unless the AI systems are used for standard editing assistance or do not significantly change the input data or its meaning. Tungsten Automation is committed to ensuring that the outputs are marked accordingly.

California AB 2013: Artificial Intelligence Training Data Transparency

California AB 2013 focuses on transparency regarding the training data used to develop AI systems. Under this law, developers must provide a summary of the data sets used to train any AI system that generates synthetic content, such as text, images, video, or audio.

AI tools used in TotalAgility generate only text outputs, without creating or modifying images, audio, or video. Therefore, California AB 2013 requirements do not apply to TotalAgility.

California SB 942: California AI Transparency Act

California SB 942 targets generative AI systems that produce images, audio, and video, mandating additional transparency for these outputs. This law primarily affects providers with more than one million monthly users.

AI tools used in TotalAgility generate only text outputs, without creating or modifying images, audio, or video. As such, TotalAgility falls outside the current scope of California SB 942.

Generative AI use cases

TotalAgility provides a range of options for using Generative AI on the platform.

TotalAgility automation activities, agents, and chat controls in forms

Here, the client has complete control over the prompts and content passed to the model. The client is responsible for the results returned and associated with ongoing monitoring of the accuracy/appropriateness of responses from the model. This approach applies to tasks like zero-shot extraction, analysis, planning, and summarization. Review this chapter carefully for notes on security, deployment architecture, versions used, and other information.

Copilots for extraction and development

When using TotalAgility's built-in Copilot prompts, the images or files provided by the user/automation are combined with Tungsten Automation prompts and context.

Copilot for extraction: Uses patented technology to manage the process of breaking down content-rich documents into structured text/data, interacting with the LLM, and applying the user's prompt to return the desired information points.

Copilot for development: Enables users to describe the desired workflow, form, or data model they want to create in the TotalAgility Advanced Studio. Based on the description and/or image assets provided, the Copilot automatically generates a workflow/form layout. Users can add, edit, and adjust from this quick and easy starting point.

TotalAgility Enterprise: knowledge discovery

The **Knowledge discovery agent** activity available in TotalAgility Enterprise is a prebuilt agentic activity that combines all steps required to implement an advanced RAG pattern.

- Vector embedding the supplied search prompt/context/question.
- Searching the knowledge base (vector similarity search).
- Natural Language Understanding analysis of the prompt/context/question to identify the most appropriate results (result ranking), which improves the accuracy of the response.
- LLM-generated response/answer to the prompt/question.

AI model monitoring and management

Both Tungsten Automation models and customer-trained models support model monitoring.

Tungsten model monitoring

Tungsten Automation performs regular testing of all machine learning, computer vision, OCR, and generative AI models as part of our product development, regression testing, and quality assurance operations. Each release or update is tested and certified before being made available to customers to ensure any changes to the solution or models used do not impact model accuracy. This applies to all AI models modified in a release, as well as the capabilities that allow customers to train and configure their own models (for clarity, Tungsten Automation tests the tools provided to allow clients to create their own models, as opposed to testing the models that clients create).

Sample packages and testing data are used in benchmarking and testing to measure performance against our baseline benchmarks. A proprietary test set of example documents is used for benchmarking that includes a wide variety of content types and languages, including scanned images of different formats and qualities, as well as different types of documents with structured fixed layouts (tax forms, IDs, driver's licenses, medical cards), semi-structured layouts (invoices, bank statements, receipts, utility bills) and unstructured content (financial statements, annual reports, contracts, news articles). Different formats of documents are used for various aspects of AI functionality:

DocAI Studio (including OCR accuracy and Copilot for Extraction): BMP, GIF, JPEG, PDF, PNG, TIFF

Generative AI: BMP, GIF, JPEG, PDF, PNG, TIFF

AI Knowledge Base: BMP, DOCX, HTML, JPEG, PDF, PNG, PPTX, TIFF, XLSX

Copilot for Development: CSV, GIF, JPEG, PNG, TXT, WEBP, XML

Azure Document Intelligence OCR Provider: PNG (PNG is the format TotalAgility uses with Document Intelligence)

Tungsten Automation does not provide access to the datasets used for testing or individual test or benchmark results, although guidance on how clients can set up and run their own benchmarks and automated testing is provided in the *Tungsten TotalAgility Best Practices Guide*.

Specifically for Generative AI, Tungsten regularly tests new models (Large/Small Language Models, embedding models), including models from different vendors. When new models or model versions are available, Tungsten runs a series of internal benchmarks to ascertain potential benefits of using that model for the Copilot for Development, Copilot for Extraction and Knowledge discovery agent activities described above.

Methodologies used to test Generative AI models:

- **Human Evaluation**: Manual tests to ensure returned content is correct based on the prompt given and the content available in the documents used for testing.
- **Negative Testing:** Asking for information that is not available in the document being queried.
- **Prompt Injection Testing**: Attempting to misguide AI by using prompt injections and other adversarial techniques.
- **Automated Testing**: Unit and sanity testing of AI models is included in Build Verification Testing, including functionality associated with the use of third-party AI models such as OpenAI and Azure Document Intelligence.
- **Question and Answer benchmarks**: Automated benchmarks of Document and Knowledge Base RAG using validated correct question-answer pairs and source references from sample documents (test sets include public domain annual reports and Tungsten product documentation).

Besides urgent break/fix or security updates to models (which can be applied without notice), the consumption of a specific version of a Generative AI model, where a TotalAgility client is using a Tungsten-provided/hosted model, is linked to the version/release being used by the client. Any significant or noteworthy changes to the behavior of a Tungsten solution component relating to a change in the Generative AI (or other AI model) model used for that release/component will be detailed in the release notes accompanying that version update.

In addition to model testing, Tungsten separately runs performance, stability, and scaling tests against core software and all cloud infrastructure components. Again, these tests are run regularly as part of ongoing release management and QA processes.

Performance and scalability testing for AI models includes:

- **OCR and data extraction**: A test dataset of US and German invoices with a standardized hardware configuration to measure throughput/performance, memory usage, and precision/recall.
- **OCR and data extraction:** Performance, OCR, table extraction, and recognition testing to extract large files (100-500-page documents) in different deployment configurations.
- **OCR and data extraction**: Throughput/performance/stability/parallelization testing using 1000 and 100,000 batches of mixed documents (multi-million-page load tests).
- **Knowledge Base ingestion**: Performance testing using a large, 800+ page document with different chunking, enrichment, and processing configurations.
- **Knowledge Base ingestion**: Performance testing using a mix of PDF, JPG, and TIFF files.

 This information covers a subset of undertaken performance benchmarks that and is not an exhaustive list. It is provided as an illustration of some routine testing undertaken in a release cycle. Tungsten Automation does not publish individual benchmark or test results.

Customer-trained model monitoring and change management

For scenarios where clients use AI tools provided in TotalAgility for specific use cases, we highly recommend that the client independently unit-test, benchmark, and monitor outputs from the AI. This includes scenarios where clients have created custom document extraction models using our trainable, computer vision, and generative AI locators/extractors; agents/chatbots created in TotalAgility; knowledge base search activities; and where generative AI activities are used as steps in a wider process. Here, clients will have configured/trained models to meet their own requirements using their document samples, data, and prompts, and should ensure that any changes to the environment, configuration, code, or usage patterns do not negatively impact results. Similarly, clients should test their knowledge base configuration in the context of their requirements, since different document chunking configurations, seeds, temperature settings, prompts, and knowledge base content will produce different results from agentic search activities that embed LLM steps.

TotalAgility provides a suite of testing, benchmarks, and performance-monitoring tools for the ongoing monitoring and optimization of custom-trained models. Refer to Extraction benchmarks and Testing for more information on implementing and managing benchmarks within the platform, and how to creating test scripts, plans, and suites for automated testing. TotalAgility includes comprehensive tools for automated deployment, package management, conflict management within training data sets, and version control of assets within the platform.

TotalAgility also uses sampling and checking ratios as quality management measures to determine the extent of reviews required. Sampling can be applied to any process or agent, including those using customer-trained, generative AI, or third-party models for ongoing quality monitoring. Checking ratios can be dynamically adjusted depending on the skill level of the person performing a task or validation. Sampling and checking are used for oversight of "online learning" scenarios for customer-trained models, where the results of operator corrections or validations are fed back into the model and model training data sets. Refer to Set sampling and checking ratios.

In some scenarios, AI models/agents provide an option for automated testing or evaluation of complex, unstructured output and for generating synthetic test data. Detailed exploration of approaches and methodologies to using AI models in these ways is beyond the scope of this document.

Azure AI services

Tungsten TotalAgility customers have access to several Azure-hosted AI Services as part of their subscription, in addition to Tungsten and third-party models and algorithms provided as part of TotalAgility. Specifically, the following Azure services managed and provided by TotalAgility Cloud Services are available for use with TotalAgility:

- Azure Document Intelligence
- OpenAI Large Language Models and Embeddings Models

These services are hosted in the same region as the TotalAgility Cloud tenant (US, Canada, EU, or Australia), or for self-hosted TotalAgility clients using Tungsten-hosted LLM capacity and/or hosted Azure Document Intelligence as the OCR Recognition provider. Clients can configure their preferred region.

For both Azure Document Intelligence and OpenAI LLMs, the Tungsten-hosted services are accessed via a proxy service over an encrypted transport (TLS 1.2 or newer) using short-lived API key authentication based on the customer license details. Each TLS session uses a separate encryption key.

Clients always have the option to configure OpenAI or Document Intelligence deployments in their own Azure subscription as an alternative to using Tungsten-hosted services, including support for containerized on-premise deployments of Azure Document Intelligence.

Azure Document Intelligence

For **Azure Document Intelligence**, all input data and results are deleted within 24 hours and not used for any other purpose. For example, submitted data is not retained or used for model training. All data is encrypted in transit (TLS 1.2 or newer) and at rest using Azure Storage. See Azure AI search in the Tungsten Cloud for details on at-rest data encryption and key management in the TotalAgility Cloud. For additional information, refer to Data, privacy, and security for Document Intelligence - Azure AI services | Microsoft Learn.

The latest version of Azure Document Intelligence is used with TotalAgility by default. Users can select the prebuilt read or layout model, along with the hosting region: US, EU, or APAC.

OpenAI large language models

For **OpenAI Services**, customer interactions with the model are logically isolated and secured employing technical measures, including but not limited to transport encryption of TLS 1.2 or newer, compute security perimeter, tokenization of text, and exclusive access to allocated GPU memory. Prompts and completions are evaluated in real time for harmful content types, and content generation is filtered based on configured thresholds.

Prompts and responses are processed within the customer-specified geography (using the selected geography for the TotalAgility Cloud tenant: US, EU, or Australia), but may be processed between regions within the geography for operational purposes (including performance and capacity management).

The following diagram provides a high-level data flow from the TotalAgility environment to the LLM.

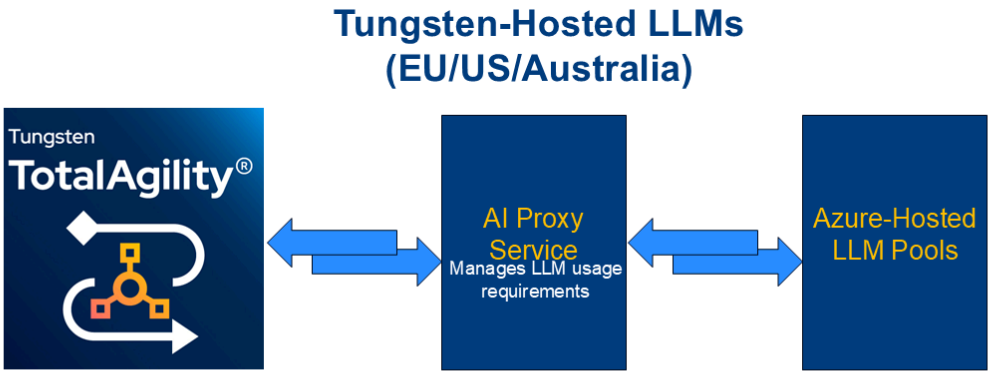

- Separate deployments (proxy + LLM pool) per geographic region
- Logical isolation of each call within LLM
- No data storage or retention (disabled Azure safety system)
- No data sharing with LLM provider for model training
- Redundant LLM regions within the selected geographic area
- TLS 1.2+ encryption for data in transit
- Unique encryption key and session generated per call
- Scalable infrastructure to meet client requirements

For details, refer to: https://learn.microsoft.com/en-us/legal/cognitive-services/openai/data-privacy

Tungsten Automation is approved for modified abuse monitoring (learn more at Azure OpenAI Service abuse monitoring), which means Microsoft does not store the prompts and completions associated with the TotalAgility Azure subscription. Therefore, the human review process described in the preceding link does not apply and is not performed. Microsoft Azure Guardrails for content filtering is activated and undesired content is blocked. "Block and Annotate" filters are configured as follows:

- For Violence, Hate, Sexual, and Self-Harm content at a low blocking level.
- For Prompt shields to prevent jailbreak or indirect attacks.

Clients can opt to replace the default configurations (hosted and managed by TotalAgility Cloud Services) with endpoints to access equivalent services in their own Azure subscriptions; or, in the case of OpenAI models, using a subscription hosted and provided by OpenAI, In this case, all data sent to these services would be processed in the customer's subscription, as opposed to the Tungsten-managed subscription (in these scenarios, the client is responsible for billing and costs for the services).

Similarly, clients can maintain multiple Large Language Model configuration. TotalAgility supports the following configuration options:

- **Tungsten Hosted LLM**
  - Included in your TotalAgility subscription license volume (Terms and Conditions apply)
  - Offers a choice of hosting regions: US, EU, and Australia.
  - Ensures no data retention/reuse to train models.
  - Utilizes OpenAI models hosted on Azure (updated to maintain version support and accuracy) GPT4o and GPT4o-mini for TotalAgility.
  - Scales to meet your volume requirements.

- **OpenAI Connectors**
  - Supports OpenAI GPT 3.5, 4, 4o and 4o-mini (refer to the TotalAgility documentation for supported versions).
  - Provides hosting options via OpenAI or Azure.
  - Supported for:
    - Copilot for Development and Extraction
    - Generative AI chat and Generative AI activities
- **Custom LLM**
  - Enables connection to any LLM, SLM, LCM, fine-tuned LM, or custom RAG/Chat implementation.
  - Supported for
    - Self-hosted/On-Premise LLM deployments.
    - Generative AI chat and Generative AI activities.
  - Connects Generative AI chat control to third-party agents

Tungsten RPA - Cloud AI services (optional add-on)

Users of Tungsten RPA Cloud can connect third-party AI services from leading cloud providers (Amazon Web Services, Microsoft Azure, and Google Cloud) for image detection, object identification, translation, and more. In these scenarios, the client provides the connectivity details (endpoints, API keys), and all data is processed in the client's subscription. For more information, refer to Cloud AI in the *Tungsten RPA Help.*

AI model upgrades

Tungsten Automation will periodically upgrade to newer Generative AI models from third-party AI service providers. Based on our standard testing procedures and lifecycle management practices, and to stay abreast of the latest technology, a secondary set of proxy services in each region will be used to test and verify new AI model versions prior to an upgrade.

Along with the default Production endpoint used with TotalAgility in your region of choice, every region also has an NVT (Next Version Test) endpoint. The NVT endpoint uses a newer version model, and if the customer elects to do so, it can be used in production. The primary purpose of the NVT endpoint is to test and verify a new model version before using it in production.

Upgrade timelines will be communicated to customers by Tungsten Automation Cloud Services.

AI services in TotalAgility Enterprise: Knowledge discovery

The information in this section only relates to users of TotalAgility Enterprise. In addition to the functionality provided in TotalAgility Standard (for IDP use cases using DocAI Studio) and TotalAgility Advanced, TotalAgility Enterprise provides new functionality for the storage and analysis of large, complex, unstructured documents.

The new functionality provides a complete suite of capabilities to prepare documents for use in advanced Retrieval Augmented Generation (RAG), supporting both AI Agents and Generative AI

chatbots for use cases, such as Research Automation, Risk Analysis, Question/Answer Data Mining, Helpdesk, and Response Automation. Large documents are chunked using semantic layout, or page, or fixed size configurations, including rich customization options, to create an advanced knowledge base for keyword, natural language, and/or semantic retrieval of relevant data to ground AI responses. This capability minimizes AI hallucinations, providing complete transparency and explainability into the sources used for a generative AI response.

Customers of TotalAgility Enterprise are provided with pre-configured AI and Storage options or can choose to provide the endpoints and keys to replace the Tungsten Cloud-hosted options with self-hosted options using the clients Azure/OpenAI subscriptions.

Third-party AI Services used in the provision of the TotalAgility Enterprise solution include:

- Azure Document Intelligence: See Azure Document Intelligence, which describes how this service is provisioned and secured for tenants in the Tungsten Cloud.
- OpenAI Chat Completions and Embedding models hosted in Azure: See the information above, which explains how this service is provisioned and secured in the Tungsten Cloud.
- Azure AI Search.

Azure AI search in the Tungsten Cloud

Azure AI Search is an enterprise-scale information storage and retrieval system storing data for use with both traditional search and vector search techniques.

The documents, case information, and data stored in the AI Search knowledge base are configured through low-code automation processes and agents in the TotalAgility Platform. The data that is stored and the duration it is retained are controlled by the client's implementation. Tungsten Automation has no visibility into the data that is stored in the knowledge base.

- Each TotalAgility tenant is configured with its own deployment of AI Search. There is no co-mingling of data with other clients in the AI Search knowledge base or shared API keys.
- All network traffic is encrypted using TLS 1.3.
- All data is encrypted at rest with AES 256 using Tungsten Cloud Services managed keys.
- Access to AI Search endpoints is secured using API keys that are managed and rotated by Tungsten Cloud Services in alignment with our security policies. API Keys are not visible in configuration screens or shared for external access. The only access to AI Search instances managed by Cloud Services is via the provided TotalAgility activities.
- The Tungsten Cloud Security Center monitors AI Search activities and resource logs. For details about the 24/7 security monitoring approaches and Security Information and Event Management (SIEM) technologies used, see Information security and accessibility.
- Data hosting and backup locations are the same as for the selected region (US, Canada, EU, or Australia).
- For additional information, refer to Security overview - Azure AI Search | Microsoft Learn.

Google AI services

Google Vision (Tungsten Clarity)

Clients may choose to use Google Vision for OCR and ICR, which is a separately licensed, optional add-on to the TotalAgility platform under the brand name Tungsten Clarity. Clients can specify their preferred hosting/data processing region when placing an order for Tungsten Clarity (choice of EU, US, or Australia). All data is encrypted in transport to this service (TLS1.2 or newer). Processed data is not retained for longer than the duration required to process the request and is not used for any other purposes (such as model training).

For additional information, refer to:

- Data Usage FAQ | Cloud Vision API | Google Cloud
- CMEK compliance in Vision API | Google Cloud

Chapter 5

Tungsten Document Library

Clients can add pre-built document models to their subscriptions. The Tungsten Document Library is an optional, separately licensed add-on to the TotalAgility Platform. The Tungsten Document Library is provided by Tungsten Automation, Azure Document Intelligence, and Base64. The specific document type selected determines the vendor model used and the location for data processing.

Location: The Tungsten Document Library service is located at the Tungsten Automation EU and US data centers. It can use third-party services located in the US for some document types.

Data retention: Tungsten Automation retains a redacted/anonymized copy of the documents processed by the Tungsten Document Library. The copy is used to improve and train models unless a customer selects the Opt-out option in the Tungsten Transformation Designer.

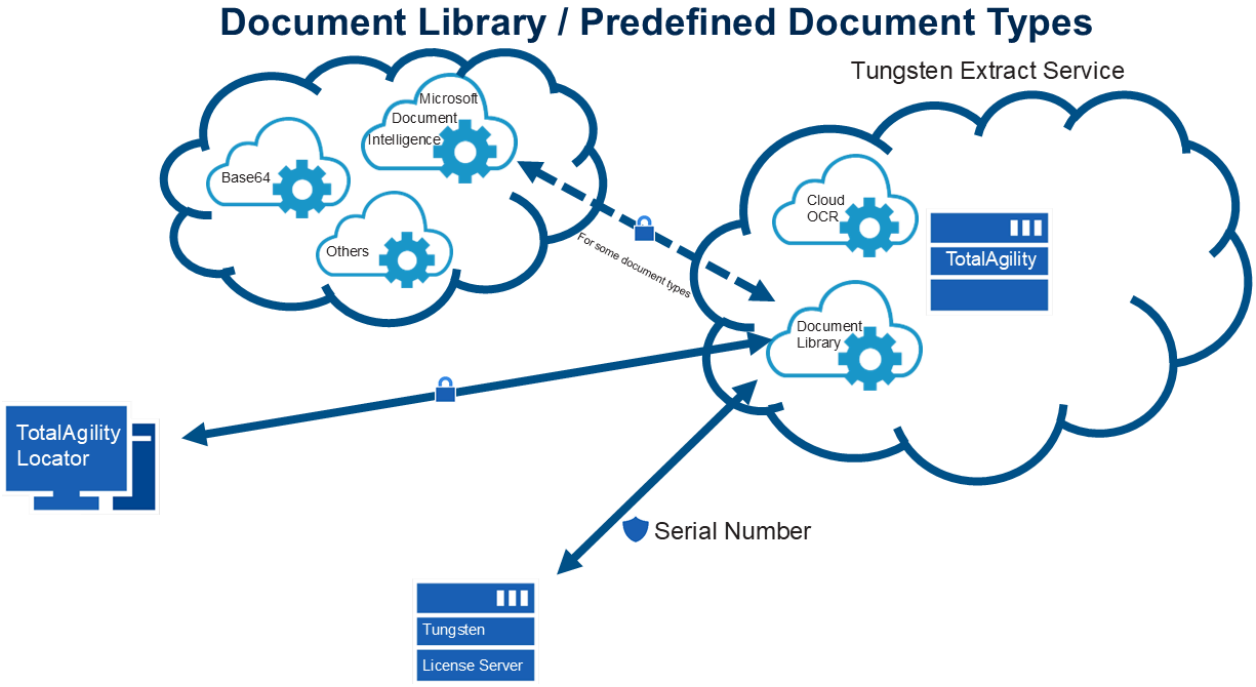

Chapter 6

Business continuity

This chapter explains how continuous availability and redundancy are ensured in the Tungsten Cloud environment.

Availability zones

Tungsten TotalAgility uses Regions with at least three Azure Availability Zones (AZs) to provide Business Continuity.

Network

AZs are designed to protect applications from network outages. Each zone has independent networking and is connected by high-speed, low-latency fiber links with redundant paths and multiple network providers to prevent a single point of failure. If one AZ experiences a network issue, others remain unaffected, and traffic is automatically rerouted using load balancers and health checks.

Storage

Storage services are designed to take full advantage of AZs to maintain data availability and integrity even during network or zone-level failures. For example, Azure offers Zone-Redundant Storage (ZRS), which replicates data synchronously across three separate AZs within the same region. This means that even if one zone suffers a network outage or physical disruption, the data remains accessible from the other two zones without downtime or data loss. ZRS ensures not just redundancy but also high read and write availability and is a synchronous action.

Processing

Cloud Services uses the modern deployment model that supports deployment across AZs, enhancing the fault tolerance and availability of hosted applications. When you deploy a Cloud Service using this model, you can specify AZs for your role instances, meaning that our web and worker roles are distributed across different physical zones. As a result, processing is not affected if a zone goes down. It could have a noticeable short slowdown of no more than 30 minutes, as not every zone has the same number of systems. In the case of a zone failover, autoscaling performs an immediate scale-up to resume continued function and maintain application availability. Azure automatically manages load balancing and failover between zones. Geo-Zone-Redundant Storage (GZRS) and ZRS are synchronous actions, as is the Azure SQL Zone Redundant Database (ZRD). This ensures that the data is in sync in each zone in case of failover. Each zone has at least one of each service running.

Zone redundant storage (ZRS) for Disaster recovery

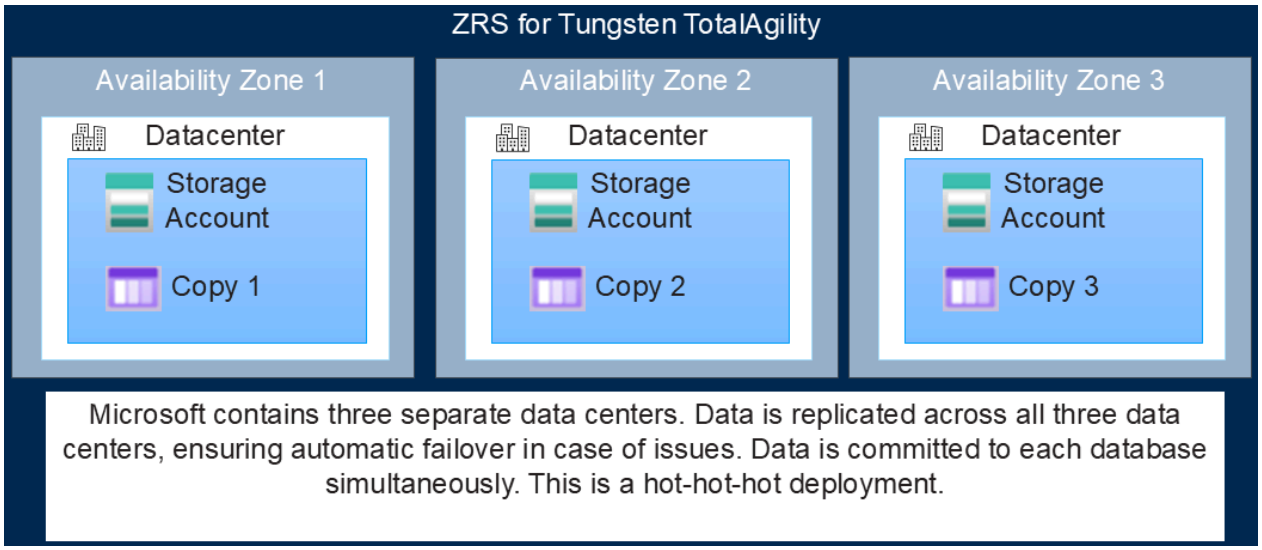

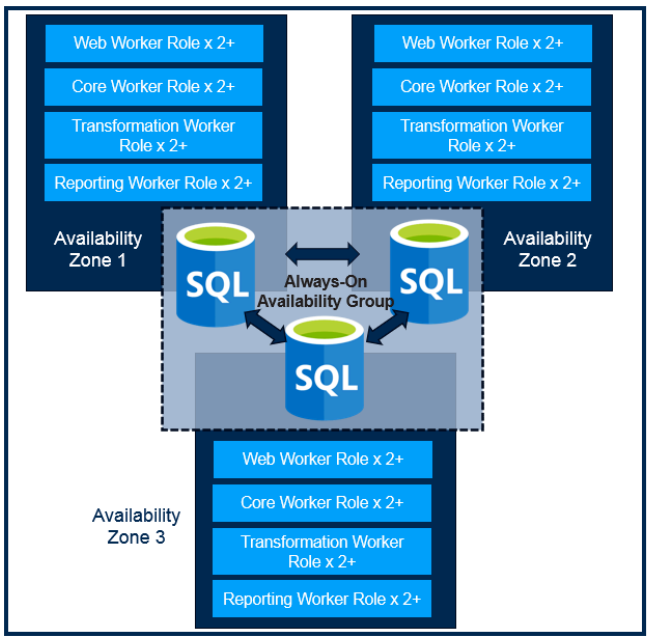

Tungsten TotalAgility Cloud uses Azure Regions with Availability Zones (minimum of three separate Zones). This is to ensure maximum resilience, high availability, and business continuity for Tungsten Cloud Customers.

- Each zone has an independent power source, networking, cooling, and storage.
- Each zone has redundant cloud service roles and a fault-tolerant SQL Server deployed.
- The TotalAgility databases are replicated across each zone as depicted in the diagram.
- The cloud services are distributed across the zones.
- If a zone goes down, Tungsten TotalAgility Cloud continues to operate seamlessly without any disruption in service for customers.

Region failover

Although failure of an entire region is highly unlikely, Geo Redundant Storage (GRS) and Paired Regions can be used to manage such an event.

Geo Redundant Storage (GRS)

GRS copies your data synchronously three times within one or more Azure availability zones in the primary region using Local Redundant Storage (LRS).

GRS then copies your data asynchronously to a single physical location in the secondary region. Within the secondary region, your data is copied synchronously three times using LRS.

If the primary region becomes unavailable, failover to the secondary region occurs. After the failover operation is completed, the secondary region becomes the primary region, restoring your ability to read and write data.

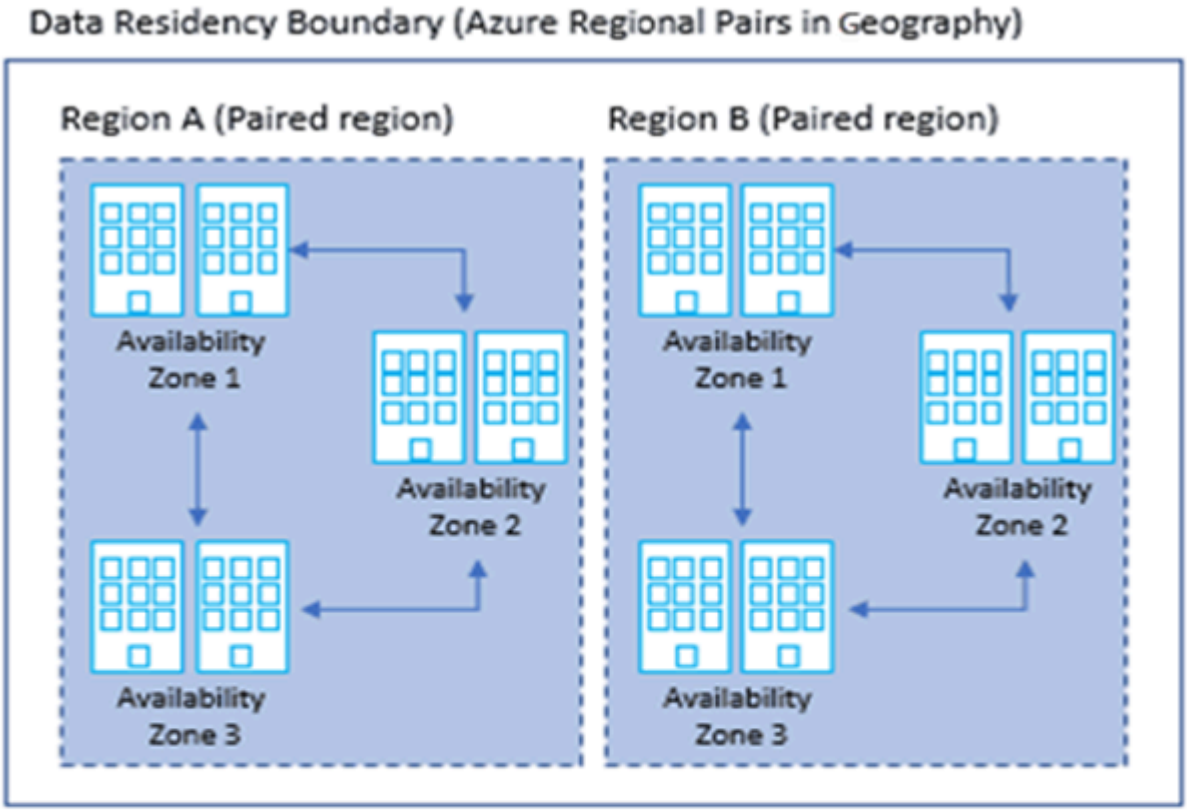

Primary	Secondary
West US	East US
North Europe	West Europe
Southeast Asia	East Asia

Examples of region pairs

Paired Regions

If a region failover occurs, the following automated steps are initiated on the paired US Central region using a manually launched Terraform script:

- Provision SQL Server and related SQL accounts
- Restore databases
- Deploy infrastructure, including Provision Virtual Network, Subnets, WAF, DNS, Load Balancer, and public IP addresses
- Deploy Tenant Management System Cloud Service
- Deploy Dev Cloud Service
- Deploy Live Cloud Service

TotalAgility continuity of business

The following diagram demonstrates region failover to ensure continuity of business. The continuity of business is only instantiated if ALL data centers within the primary region become unreachable.

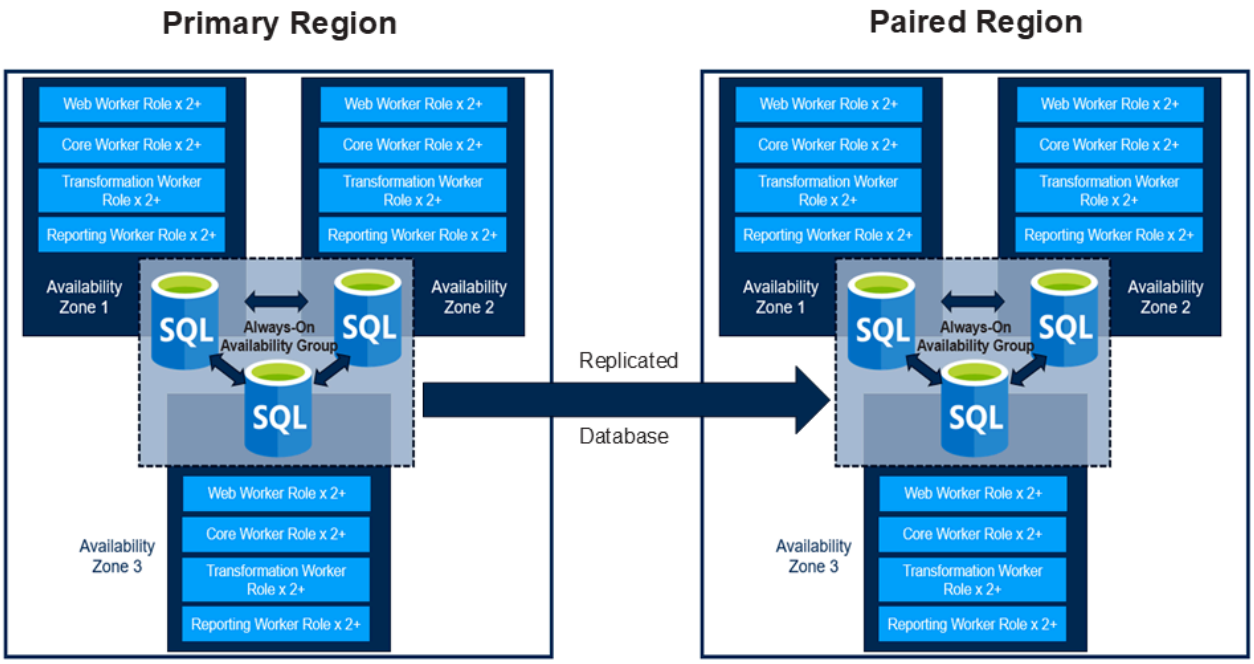

High availability and scaling

TotalAgility is a stateless application.

TotalAgility Cloud does not use VMs in Azure. Worker roles and Web roles are used. Behind the scenes, Azure does use VMs, and they are all built with full redundancy. Azure manages any failures seamlessly. For more information on Azure redundancy, refer to:

https://docs.microsoft.com/en-us/azure/virtual-machines/manage-availability

Azure auto scales the worker and web roles. TotalAgility Cloud always starts with a minimum of two (2) roles with no maximum number defined.

The Tungsten Cloud offers two deployment options for standard and enhanced high availability.

Chapter 7

Information security

This chapter describes how Tungsten Automation secures your data and manages access to it.

Data in transit

Connectivity to TotalAgility Cloud uses TLS 1.2 over HTTPS port 443, and no other ports are required to be open on the network.

TotalAgility Cloud integration with on-premise components

To enable a more comprehensive and configurable capability for communicating to and from TotalAgility Cloud, TotalAgility provides a separate component called an *Integration Server.* The Integration Server enables secure communication to/from a customer's data center. This allows a customer to import content on-premises and send it to TotalAgility Cloud, and/or then let TotalAgility Cloud to send data to an on-premises integration server for export to one or more systems of record.

The following diagram illustrates an overview of communication between TotalAgility Cloud and Integration Server.

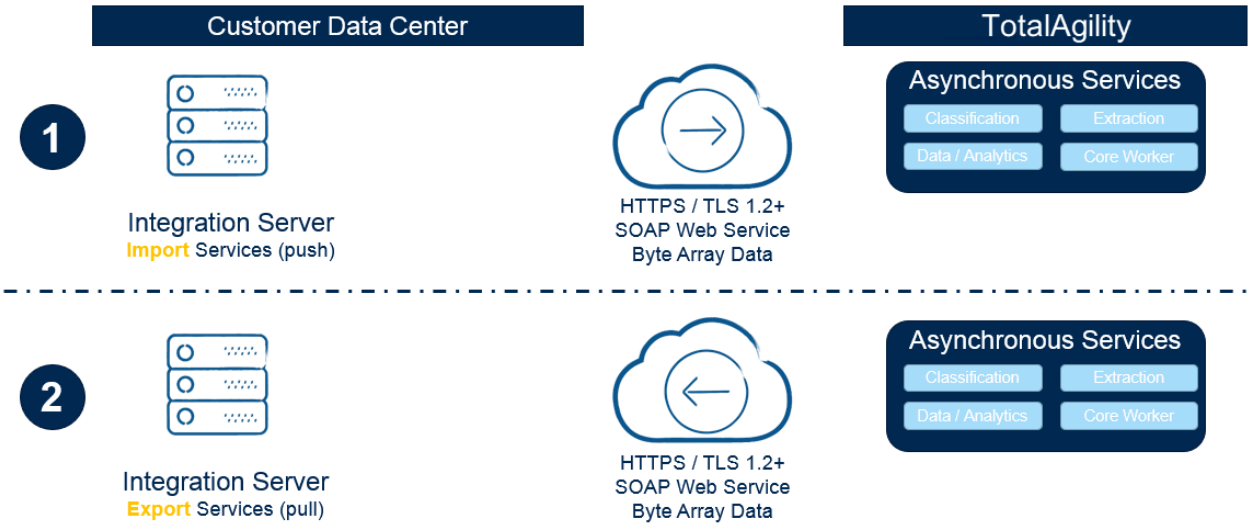

Integration Server notes:

- Integration Server provides a set of services configured during installation to point to TotalAgility Cloud.
- Integration Server does not store any data; all users and data are stored in TotalAgility Cloud. When logging in to an Integration Server, users are authenticated against TotalAgility Cloud.
- When logging into an Integration Server, users are authenticated against TotalAgility Cloud.
- Integration Server polls the TotalAgility Cloud database for work (Pull).
- Integration Server is primarily used for import/export and integration into on-premises systems of record.
- Multiple Integration Servers can point to the same TotalAgility Cloud tenant.
- Integration Servers can be scaled out for high availability and load-balancing.
- Integration Server runs automatic activities only.

The following diagram provides a detailed view of the service layers that depict Integration Server communication to/from the TotalAgility Cloud and the On-Premise system of record

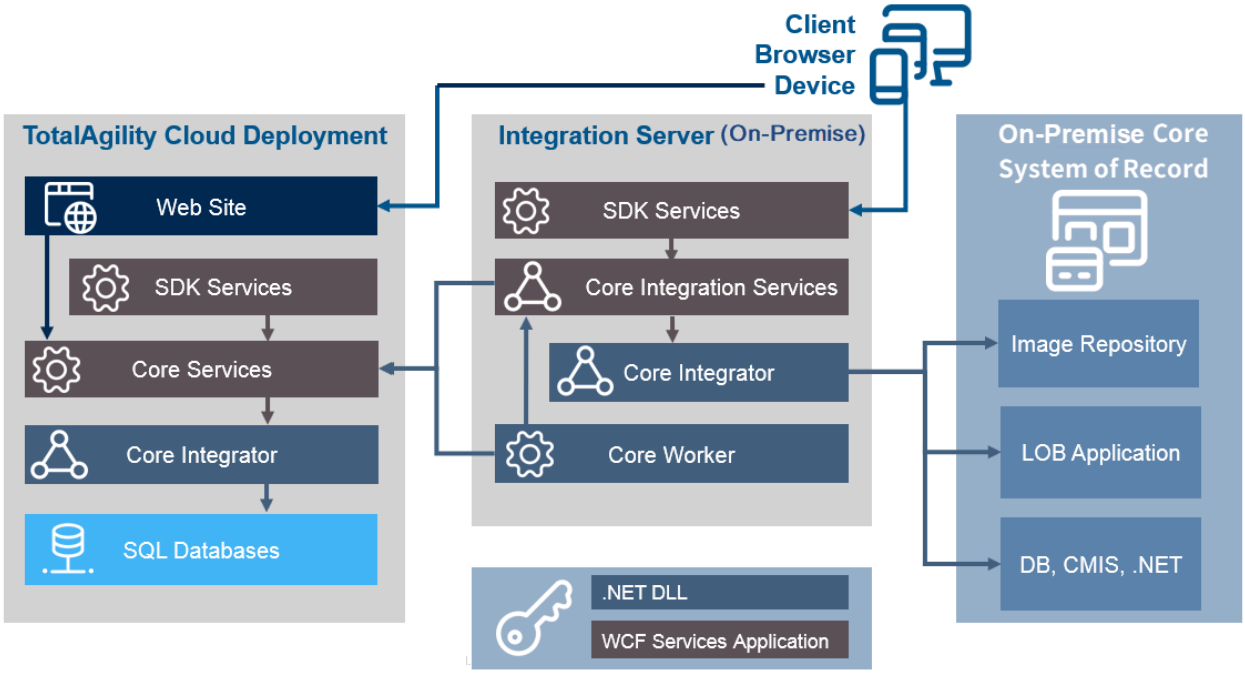

Data at rest

All data at rest is encrypted using AES-256 at the database level. Once Azure Worker Roles process data, customer data is deleted from Azure Table and Blob Storage in accordance with the timelines and policies defined by the customer within the TotalAgility application.

Authentication and Authorization

To perform any action within TotalAgility, you must be logged on as a registered user in TotalAgility. Once authenticated via a Logon action, a Session ID is generated and used in all subsequent interactions with TotalAgility. A user is authenticated in TotalAgility, typically through Single Sign-On (SSO) using Federated Security (such as using Active Directory FS).

Federated Security

The Federated Security integration with TotalAgility typically uses claims-based authentication that uses WS Federation or SAML. User claim mappings are then configured within TotalAgility. This determines how an existing user is found in TotalAgility when the logon is performed using the Authentication provider based on username or email address. The match-to selection must be mapped to a claim rather than entered manually. A set of rules is configured, indicating the Category and Groups to which a new user is added in TotalAgility after they have been successfully authenticated with the provider. This determines the privileges they have once they are authenticated. Access control lists can be configured to assign further privileges.

User privileges and access control lists

Several Access Control Lists (ACLs) determine which users can work with the design and implementation of transformation projects, cases, processes, business rules, dashboards, and resources. Typically, an administrator uses these ACLs to determine which users have access to specific areas of the platform.

Functional application - Level access

Functional access rights determine runtime privileges for end users.

Cases and processes - Types of access

- Functional access settings determine which users have permission to perform certain runtime application-level functions. (Creation, Suspend, Terminate, On Hold, Restart, Customize, Job Details).
- Maintenance access settings for cases and processes determine which users can change the design of a specific case or process type. This is used in conjunction with the ACLs above (Full Control, Read / Write, Read Only, None).
- Administrative Security is a pre-defined group within the platform called Administrators. Any member of this group can access the administrative functions for the categories to which they have acces (Operations, Archiving, Recovery).
- Hierarchy Security allows you to set up a hierarchy to reflect your organizational structure. Defining a user as a supervisor automatically gives them the rights (viewing work queues of

those people who report to them either directly or indirectly, or reassigning and delegating work).

- Work is assigned to people in many ways in TotalAgility. Typically, work is assigned to groups or roles of which certain individuals can be members and groups mirror a corporate LDAP group structure. Roles are business roles defined as part of a TotalAgility solution design. These roles can be fixed or floating. The key difference is that floating roles permit a more dynamic approach to work allocation. A key part of TotalAgility is this dynamic allocation of work. Typically, in the modern business world, work is not simply assigned to a team in advance but complex rules, including a person's skill set, security levels, experience, and suitability, are fed into a TotalAgility business rule to determine the most appropriate user for a specific task at runtime. Furthermore, TotalAgility employs exceptions so that if a required user is unexpectedly unavailable, it triggers alerts, and (if required) automatically reallocates work based on user defined business rules.

Key vault

Tungsten Automation uses Azure Key Vault to store encryption keys used by the customer environment. Each dedicated instance holds its own Key Vault, which includes separate keys and does not have any communication to the Multi-Tenant Key Vault system. All keys are rotated only annually and TotalAgility version upgrades. Standard Key Vault logging information is stored and transmitted to Log Rhythm. Tungsten Automation does not currently support "bring your own key."

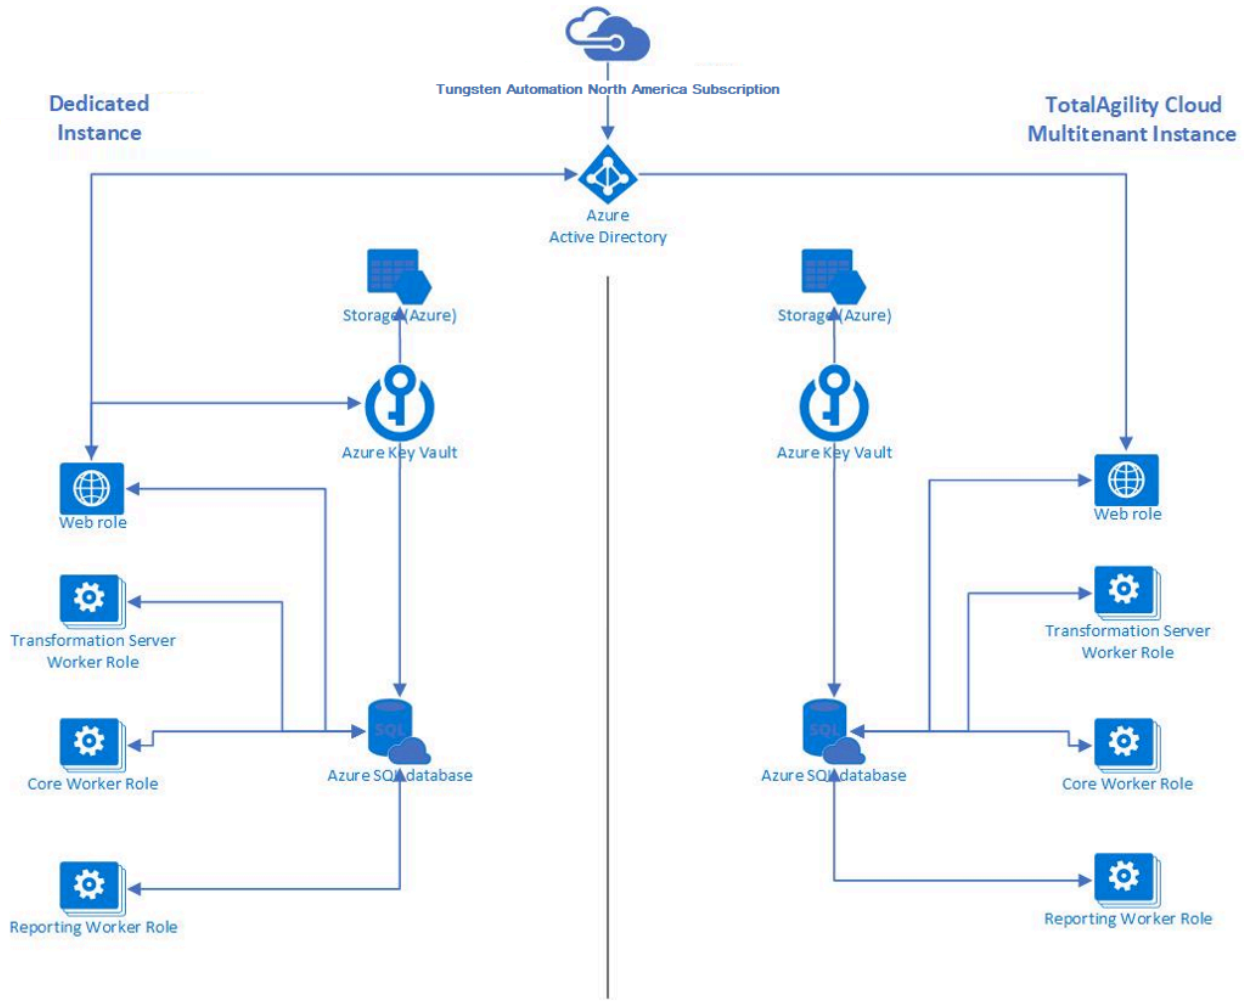

Access control

Tungsten Automation Cloud Services has a strict access control policy that restricts access to any cloud environment.

Access requests are submitted through a ticketing system, and justification is required for any access. The Director of Cloud Services approves all access requests. Only vetted and trained Cloud Services personnel are allowed into a TotalAgility Cloud environment. Personnel with access are those Systems Administrators and Cloud Architects assigned to the TotalAgility Cloud product inside Cloud Services. All Cloud Services personnel receive security training annually. All access to Cloud systems is tracked by Azure AD, and logs are also sent to Azure Sentinel SIEM and to Security Operations Center (SOC) for review.

There is no direct access to the Azure TotalAgility environment from the Tungsten Automation corporate network or by Tungsten Automation employees who are not members of the vetted

Cloud Services team. Users, whether on a Tungsten Automation network or remote, must log in with a username and password and use the Microsoft Authenticator application via mobile phone. Connection into the TotalAgility Azure environment requires a separate login from Tungsten Automation corporate systems, as it resides in a separate domain with no trust relationship to the Tungsten Automation corporate network.

Change management

Tungsten Automation Cloud Services uses a change management system that tracks all change requests. This includes any network, system, DNS, Certificates, Operating System (OS), or product changes. The change control process requires a peer review prior to management approval. All changes must include the reasoning for the change, along with the plan for the change, testing, and rollback processes. Management reviews every change ticket created and performs postmortems for all unsuccessful changes.

As stated earlier, only Cloud Services personnel have access to the TotalAgility Cloud environment. Customer Support does not have any access to the Azure environment and can only log into TotalAgility when the customer grants them access. Cloud Services created a tool that allows Customer Support to receive TotalAgility systems logs to assist in troubleshooting TotalAgility issues. These logs do not contain any customer-sensitive information.

The standard maintenance is performed one Saturday per month and begins at 22:00 Central Time in AMS, 22:00 GMT in EMEA, and 22:00 Central Time in AEST in APAC. The timing can be adjusted to begin at midnight if requested by the customer. Emergency patching has a window nightly from midnight to 4:00 Central Time in AMS, midnight to 4:00 GMT in EMEA and midnight to 4:00 AEST in APAC. This window is reserved only for emergency security patching. Tungsten Automation takes every measure to prevent any outages should an emergency patch be required. However, depending on the security patch, outages can occur.

Log management

All access and system logs are transmitted to the Azure Sentinel (SIEM) system for analysis and long-term storage. Logs are stored for at least 12 months. A security operations center (SOC) operates 24/7 to monitor and review all logs. Web logs are routed to the SOC's advanced security tool to determine if any threats have occurred. If a valid threat is suspected, the SOC notifies the Tungsten Automation network operations center via email and phone. Tungsten Automation network operations is staffed 24/7 and takes appropriate action in the event of a threat alert.

Network security

Tungsten Automation uses a variety of security measures to protect systems, including:

- Firewalls
- Azure WAF
- Security center

- Alerts for any configuration, web, or folder modification
- IDS
- Network-based anomaly detection
- Full monitoring 24/7 by our security operations center

Chapter 8

FAQs: TotalAgility

For answers to FAQs about Tungsten TotalAgility, see the following table.

Question	Answer
Can I control when software updates are performed?	Yes and No. Microsoft performs server patching automatically. Typically, major versions and patches for TotalAgility are applied to the TotalAgility Cloud six weeks after they are released for the core product. Clients are notified by email in advance of any updates, highlighting relevant release notes and deprecated functionality. At a minimum, a 36-hour notice is given for hot patches, and a 30-day notice is given for major version updates. Development environments are upgraded 30 days before test/ production environments, to allow clients time to test the updated version. Clients also have the option of purchasing a "dedicated instance" of TotalAgility Cloud, which runs on a dedicated (not shared) Azure infrastructure. For clients on dedicated instances, the upgrade timing can be agreed upon mutually. However, Tungsten Automation wishes to stay no more than two minor releases behind the current version. Urgent patches or hot fixes may be applied to TotalAgility Cloud outside of the standard notice period to address system stability, performance, or security issues.
Do you perform network and penetration testing?	Yes. Microsoft performs standard Azure tests. Tungsten Automation performs TotalAgility Cloud Penetration testing annually.
Do I have to use an Integration Server to communicate between on- premise systems of record?	No. However, the decision depends on the customer's security and compliance rules. Without an Integration Server, the customer would need to securely open a port on their firewall to make it publicly available to external connections. The customer could then apply a source IP restriction on this port to ensure that only the TotalAgility tenant can access it. Customer applications can connect to TotalAgility Cloud using published web services (REST and SOAP). This access is secured via TLS 1.2, SSO, secure sessions, and the network security measures explained in the Network Security section of this document.

Question	Answer
What access do Tungsten Automation personnel have to the TotalAgility Cloud environments?	The TotalAgility Cloud environment is hosted in Microsoft Azure, so all access to this system is remote. Access requires access approval from the Director, Cloud Services. Only TotalAgility Cloud Services personnel can access these systems. A username, a password, and a rotating multi-factor key are required to gain access. Access is provided for the initial security training and revoked if annual training is not completed.
Does TotalAgility Cloud have SOC 2 Compliance?	TotalAgility Cloud is SOC 2 Type 1 and 2 certified. Refer to Tungsten Automation Trust for more information about the security certifications undertaken by Tungsten Automation.
Does TotalAgility Cloud have ISO 27001 Compliance?	TotalAgility Cloud is ISO 27001 certified as of January 2023. Refer to Tungsten Automation Trust for more details about the security certifications undertaken by Tungsten Automation.
Does TotalAgility Cloud have PCI certification?	Out of the box, the TotalAgility Cloud is not PCI certified, but TotalAgility can be configured to support standards such as PCI and HIPAA. Contact Tungsten Automation Professional Services for guidance and more information.
Is there additional licensing required for the Integration Server?	No, it is included with TotalAgility Cloud.
What is the frequency that you back up the MS Azure databases?	Local database backups utilize Azure point-in-time restore. They are near real-time backups and are kept for up to 35 days. If the option for a secondary site is chosen, geographically redundant backups can be performed. The schedule for geographically redundant backups is customer-specific based on their requirements.
Are there options for a customer to engage a third-party to perform penetration tests of the production environment?	Yes, a third-party penetration test can be performed. Tungsten Automation would take any findings and work any severe, high, or medium findings into the release schedule. As a side note, Tungsten Automation also uses a third-party for penetration testing.

Question	Answer
Can TotalAgility solutions be packaged and automatically deployed?	Yes. The TotalAgility platform has a packages module that feeds into any DevOps plan. The packages module provides the capability to build packaged solutions and migrate them between environments. Deployment servers and schedules can be configured in TotalAgility to provide a continuous deployment capability based on the schedule. Packages can be created for specific environments, such as Development, User Acceptance Testing, or Production, as required. As part of the continuous deployment process, TotalAgility configured test plans can be run prior to deployment, and if test plans fail, the deployment is terminated. Furthermore, if there is an error during deployment, a rollback can be performed automatically if configured on the target server. This approach ensures that the target environment remains in a stable state. Continuous deployment can be configured to be executed on multiple target servers.
What is the outage schedule during planned maintenance?	The maintenance period is once per month and occurs on a Saturday starting at 22:00 Central Time in AMS, 22:00 GMT in EMEA, and 22:00 AEST in APAC. An emergency patch window exists nightly from midnight to 4:00 Central Time in AMS, midnight to 4:00 GMT in EMEA, and midnight to 4:00 AEST in APAC. Customers are notified at least 36 hours in advance of all maintenance activities.
What approach is taken for VNet IP address range allocation?	The selected IP address range must not overlap with any other Tungsten Automation IP ranges, and the range is communicated to the customer. If the customer has a conflict with the range, an alternate range can be selected.
SQL Azure Tenant and Data Center databases are shared by multiple TotalAgility deployments and environments. What measures are in place to control access to DB schema? How does encryption work on the schema level?	Access to the database and schema is restricted only to the application and the appropriate members of the Cloud Services team. The application is designed to use only the schema appropriate for its environment. Encryption occurs at the database level, not at the schema level.
Regarding SSL certificates for Integration Server connectivity with TotalAgility Cloud: Who generates and manages certificates that are imported to on-premise Integration Servers?	The Integration Server is only an outbound connector. It pushes or pulls information to/from the TotalAgility Cloud tenant and uses the certificate on the tenant site.
Is there annual/quarterly DR testing? What is the process to initiate it?	Annual testing is scheduled based on the calendar year. A full DR test is run approximately one month after the customer initiates the license, and then every calendar year afterward.

Question	Answer
What if I need a database to support my processes in TotalAgility?	On request, customers can have one additional database per environment. This database can be used by processes in TotalAgility, and it is backed up as part of the overall environment. The size of the database counts towards the overall storage usage and limits based on the client's contract.
What are the storage limits for TotalAgility Cloud?	Storage limits depend on the document and user volumes of each customer on TotalAgility Cloud. The storage assigned per tenant covers combined databases and storage usage and can be increased if required, subject to a commercial agreement.
Does the Tungsten Cloud offer a Real Time Transformation Service for TotalAgility?	Yes, Real Time Transformation (RTTS) can be selected as a configuration option. For example, the service can be used to support near real-time OCR, document extraction, or mobile ID verification workflows.
How long is data retained?	The Tungsten Automation Cloud Services team manages data retention after TotalAgility 8.0. By default, data is retained for 6 months, up to the contracted storage limit, and operational log data is retained for one year and 45 days. Clients can submit a ticket to Tungsten Automation if they need to adjust the data retention period. For more information about data retention at the end of the contractual period and other GDPR considerations, refer to the data processing addendum provided with your contract.

Question	Answer

Is there any option to test a new version of TotalAgility prior to the cloud environments being upgraded?	Upgrades to the TotalAgility Cloud environment are typically scheduled 30 days after the launch of a major version of TotalAgility.
	TotalAgility Cloud customers will have major and minor releases applied to their Development environments initially, then 60 days later applied to Test/Production instance/environments.
	• 30 days after a TotalAgility major/minor release, it is deployed by Cloud Services to:
	  • Customer Dev Tenants
	• 60 days later, the major/minor release is deployed by Cloud Services to:
	  • Customer Test Tenants
	  • Customer Production Tenants
	Tungsten Cloud clients should work with the Tungsten Automation Technical Support and/or Customer Success teams to report any issues encountered during testing.
	Clients are notified at least two weeks in advance of their environment being upgraded, with any scheduled maintenance downtime occurring outside of that region's working hours.
	If any bugs are identified during the testing period following an upgrade, they are reviewed by Tungsten Automation Technical Support and Tungsten Automation Cloud Services. If necessary, the upgrade of the test/production environments is paused while the identified issues are resolved.
	Customers using a "Dedicated Instance" have more flexibility when their instance is upgraded, but it is still expected to occur within a reasonable time frame. The upgrade cannot be delayed indefinitely.

Question	Answer
What measurements will Tungsten Automation accept regarding performance/responsiveness of the cloud system? What occurs if we experience a slow response while loading a TotalAgility form, or if OCR activities take a long time to execute?	TotalAgility Cloud uses auto-scaling technology to transparently scale resources to meet peaks in demand. The system performance and supporting infrastructure is continually monitored to proactively identify any performance or response time issues, so corrective action can be taken. If a client experiences persistently slow page load or process/activity execution times, they can open a support ticket for Tungsten Automation to investigate. Tungsten Automation does not explicitly specify the performance of individual TotalAgility components. Factors that affect Web page load times are beyond the control of Tungsten Automation. The exact configuration of processes/activities, network connection to the cloud, and size of payloads processed determine the response times experienced by end clients. Note that TotalAgility Cloud is by default optimized for throughput. The exact scheduling of execution tasks is influenced by many factors, including configured process priorities, thread pools, whether a process has been marked as synchronous or not, and more. Real-time document transformation options can be provided at additional cost. If real-time OCR is a requirement for the use case, the client should purchase sufficient volume of "Real Time Transformation Service" documents to meet this need. TotalAgility Cloud uses New Relic to track and monitor page load times for the TotalAgility Workspace portal.

Chapter 9

FAQs: Shared vs dedicated instance

For answers to FAQs about shared and dedicated instances of TotalAgility, see the following table.

Question	Public / Shared Instance Answer	Dedicated Instance Answer
How are the version upgrades managed between shared and dedicated instances?	Example: Development/User Acceptance Testing/Production upgrades are conducted 20-30 business days apart. Customers are notified 30 days in advance unless a security or emergency fix is required.	Example: Customer has more control on timing of upgrades. Important: Although the customer has control over timing, upgrades still must occur in a timely fashion to ensure system stability and security.
How is Physical Data split in terms of networks, servers, and subnets?	Each customer has its own database and storage systems, along with encryption keys.	The system is created in its own Azure subscription.
How is database performance affected in a shared environment vs. a dedicated instance?	As each customer has their own database, performance is monitored and can be scaled when required.	Same as public cloud.
What isolation mechanisms are in place to prevent data leakage between tenants in shared instances?	Each customer has their own databases and storage systems, and encryption keys.	The system is created in its own Azure subscription.
What are the backup and recovery options for databases in both shared and dedicated instances?	Databases backups occur every minute with 14 days of storage. The backups are stored in a separate Azure location for DR purposes.	Same as public cloud.
How does database scaling work in shared vs. dedicated environments?	As each customer has their own database, Tungsten CloudOps monitors the performance and scales it when needed.	Same as public cloud.
What are the network latency expectations for shared vs. dedicated instances?	There are no differences in network latency between public and dedicated.	There are no differences in network latency between public and dedicated.
How is network traffic managed and segmented between tenants?	Each tenant has its own URL, and this is how traffic is segmented.	The customer has its own URL.
Are there any bandwidth limitations or throttling policies for shared instances?	No.	No.

Question	Public / Shared Instance Answer	Dedicated Instance Answer
What network security measures are in place, such as firewalls and Virtual Private Networks (VPNs), for both shared and dedicated instances?	Firewalls, Web Application Firewall (WAF), Load Balancing, and Malware/anti-virus (AV) are all in place. VPNs are not allowed.	Same as public cloud.
Is storage shared among multiple tenants in a shared instance?	No, each tenant has its own storage with a separate encryption key.	Same as public cloud.
How is storage performance (I/O operations per second) managed in shared vs. dedicated instances?	Each tenant has its own storage with a separate encryption key.	Same as public cloud.
How does storage scalability differ between shared and dedicated instances? Describe any differences in data access speeds and latency.	There is no difference.	Same as public cloud.
How are computer resources (CPU, RAM) allocated and managed in shared versus dedicated instances?	Resource scaling is based on monitoring systems and the total number of users and documents being processed in the environment.	Same as public cloud.
What are the performance guarantees (such as dedicated CPU cores) for dedicated instances?	They are spelled out in the agreement, and there are no differences between public/shared and dedicated systems.	Same as public cloud.
Can I customize the computer resources (such as specific CPU models, RAM configurations) in dedicated instances?	No.	No.
How are security patches and updates managed in shared vs. dedicated instances?	Patching is an automated process that occurs monthly after each Microsoft Patch Tuesday. There are no outages with system patching.	Same as public cloud.
What monitoring and logging capabilities are available to track security events in both environments?	Extensive monitoring occurs for all events. They are sent to a SIEM and monitored by a Security Operations Center 24/7 in real time.	Same as public cloud.
How are resource allocation and prioritization managed in high-traffic scenarios in shared instances?	TotalAgility is designed for high-traffic processing in shared or dedicated instances. Also, the scaling and monitoring history allows the Cloud Services team to have systems right-sized and scaled at appropriate times.	Same as public cloud.

Question	Public / Shared Instance Answer	Dedicated Instance Answer
What options are available for automated scaling and load balancing in shared vs. dedicated instances?	They are managed by Cloud Services, and automated scaling is enabled by default.	Same as public cloud.
Are there any limitations on the types of processes or services that can be run on shared vs. dedicated instances?	No.	No.
Is Real time Transformation Service supported on shared vs. dedicated instances?	Yes.	Yes.
How easy is it to scale up or down based on future needs for both shared and dedicated instances? How long does it take to scale between the two options?	There is no difference in scaling, which is the responsibility of the Cloud Services team. Most services have auto-scaling enabled.	Same as public cloud.
What is the high availability and disaster recovery plan for each option?	See above for details and options.	Same as public cloud.
Is there a difference in support levels (priority support) between shared and dedicated instances?	No.	No.
What is the maximum capacity of the shared or dedicated instance?	There is no theoretical maximum capacity,	Same as public cloud.
Where are the support resources available for shared vs dedicated instances? Can support be local for dedicated instances?	Support and Cloud Services operate the same whether shared or dedicated.	Same as public cloud.

Question	Public / Shared Instance Answer	Dedicated Instance Answer
Why is the Hébergeurs de Données de Santé (HDS) certification required for entities such as cloud service providers that host the personal health data governed by French laws and collected for delivering preventive, diagnostic, and other health services?	When the EU region is selected to host the TotalAgility Cloud, data is processed/stored in North Europe, with West Europe for failover, backup, and resiliency. Microsoft's compliance with the HDS requirements for these Azure regions has been audited and certified by the BSI Group, an independent certifying body accredited by French authorities to conduct HDS audits. In combination with ISO 27001 compliance, these regions can be used for processing French health data. For more information, refer to: Health Data Hosting (HDS) France - Microsoft Compliance | Microsoft Learn	Same as public cloud, assuming a HDS certified region is selected.
What happens to TotalAgility source code in the event the company becomes insolvent?	Tungsten Automation has established a software escrow agreement (the "Escrow Agreement") with Innosafe, Inc. (the "Escrow Agent"). If Tungsten Automation becomes insolvent, the TotalAgility software source code and AI code can be obtained by customers who have signed up as the Escrow beneficiary.	Same as public cloud.

Appendix A

Responsibility matrix

Item	Tungsten Automation	Customer
TotalAgility deployment	✓	
Monitoring (logs, performance)	✓	
Backups	✓	
Security and logging	✓	
Patching	✓	
Upgrades	✓	
TotalAgility solution configuration		✓
TotalAgility user setup		✓