

Kofax TotalAgility

Architecture Guide

Version: 7.4.1

Date: 2017-10-26



Legal Notice

© 2017 Kofax. All rights reserved.

Kofax is a trademark of Kofax, Inc., registered in the U.S. and/or other countries. All other trademarks are the property of their respective owners. No part of this publication may be reproduced, stored, or transmitted in any form without the prior written permission of Kofax.

Table of Contents

Legal Notice.....	2
Preface.....	5
Related Documentation.....	5
Training.....	6
Get help for Kofax products.....	6
Overview.....	8
Web/App tier.....	9
Server side Windows services.....	9
Database.....	9
Client side components.....	10
Real Time Transformation Service.....	11
Core Worker Service.....	13
Background System (Worker) tasks.....	13
Execution through threads in thread pools.....	14
Types of System (Worker) tasks.....	14
In-memory / cached System Worker tasks.....	14
Configuration.....	15
Performance Counter Usage.....	16
Transformation Server.....	17
Task Retriever.....	17
Task Dispatcher.....	17
Activity Processing Logic.....	18
Transformation Process Pool.....	18
Image Processing Process Pool.....	18
Ports and protocols.....	19
Custom forms configuration.....	20
Scaling and resiliency.....	21
TotalAgility Web and Application servers.....	21
Deployment options.....	21
Scaling.....	23
Clustering.....	23
Performance and sizing best practices.....	23
Transformation Server.....	24
Deployment options.....	24

- Scaling..... 24
- Performance and sizing best practices..... 24
- License Server..... 24
 - Deployment options..... 24
- Import Connector..... 25
 - Deployment options..... 25
 - Scaling..... 25
 - Performance and scaling best practices..... 25
- Reporting service..... 25
 - Deployment options..... 25
 - Scaling..... 26
 - Performance and scaling best practices..... 26
- Business continuity..... 27

Preface

This guide gives you an overview of the salient points of architecture in Kofax TotalAgility, covering various deployments for an on-premise, on-premise multi-tenancy or Azure environment.

Related Documentation

The product documentation set for Kofax TotalAgility 7.4.0 is available at the following location.

https://docshield.kofax.com/Portal/Products/en_US/740-uc0n6j0c5s/KTA.htm

In addition to this guide, the documentation set includes the following items:

- *Kofax TotalAgility Prerequisites Guide*: Provides system requirements for installing TotalAgility, instructions for running the prerequisite utility, and a software checklist for various installation types.
- *Kofax TotalAgility Installation Guide*: Describes how to install and configure TotalAgility.
- *Kofax TotalAgility Integration Server Installation Guide*: Describes how to install Kofax Integration Server and integrate it with other products.
- *Kofax TotalAgility On-Premise Multi-Tenancy Installation Guide*: Describes how to install and configure On- Premise Multi-Tenant system.
- *Kofax TotalAgility Configuration Utility Guide*: Explains how to use the Configuration Utility to update settings across various configuration files for different types of installation and deployment.
- *Kofax TotalAgility Administrator's Guide*: Provides information to the administrator on configuring and maintaining a TotalAgility installation.
- *Kofax TotalAgility Best Practices Guide*: Describes the best practices you must follow when using TotalAgility to improve performance, cost, maintenance, availability and security.
- *Kofax TotalAgility Features Guide*: Provides an overview of the TotalAgility features.
- *Kofax TotalAgility Migration Guide*: Provides information on TotalAgility upgrades from different versions and post upgrade configuration.
- *Kofax TotalAgility Help*: Provides details about using TotalAgility to design business jobs and cases, assign resources, create forms, integrate with external applications, and more. Access the help from the TotalAgility application by clicking the Help button.
- *Kofax TotalAgility Workspace Help*: Describes how to use the Workspace to manage activities, jobs, and resources. Access the help from the TotalAgility Workspace by clicking the Help button.
- *Kofax TotalAgility On-Premise Multi-Tenant System Help*: Describes how to create and manage tenants using the TotalAgility On-Premise Multi-Tenant system.
- *Kofax TotalAgility Web Capture Control Help* : Provides details on using a Web Capture control in creating multi-page documents, creating a new document in a new folder, deleting pages that have been incorrectly scanned, and more; also, describes the buttons available in a Web Capture control toolbar.

- *Kofax Analytics for TotalAgility Product Features Guide*: Provides an overview of the dashboards that help you track data through the workflow, analyze the effectiveness of the processes and resources, and address business problems.
- *Kofax TotalAgility Tables*: Describes the Kofax TotalAgility tables and fields used by Kofax Analytics for TotalAgility.
- *Migration From Kofax Products Guide*: Provides information about migrating TotalAgility files and Kofax Transformation Modules projects to TotalAgility.

Training

Kofax offers both classroom and computer-based training that will help you make the most of your Kofax TotalAgility solution. Visit the Kofax website at www.kofax.com for complete details about the available training options and schedules.

Get help for Kofax products

Kofax regularly updates the Kofax Support site with the latest information about Kofax products.

To access some resources, you must have a valid Support Agreement with an authorized Kofax Reseller/ Partner or with Kofax directly.

Use the tools that Kofax provides for researching and identifying issues. For example, use the Kofax Support site to search for answers about messages, keywords, and product issues. To access the Kofax Support page, go to www.kofax.com.

The Kofax Support page provides:

- Product information and release news
Click a product family, select a product, and select a version number.
- Downloadable product documentation
Click a product family, select a product, and click **Documentation**.
- Access to product knowledge bases
Click **Knowledge Base**.
- Access to the Kofax Customer Portal (for eligible customers)
Click **Account Management** and log in.

To optimize your use of the portal, go to the Kofax Customer Portal login page and click the link to open the *Guide to the Kofax Support Portal*. This guide describes how to access the support site, what to do before contacting the support team, how to open a new case or view an open case, and what information to collect before opening a case.

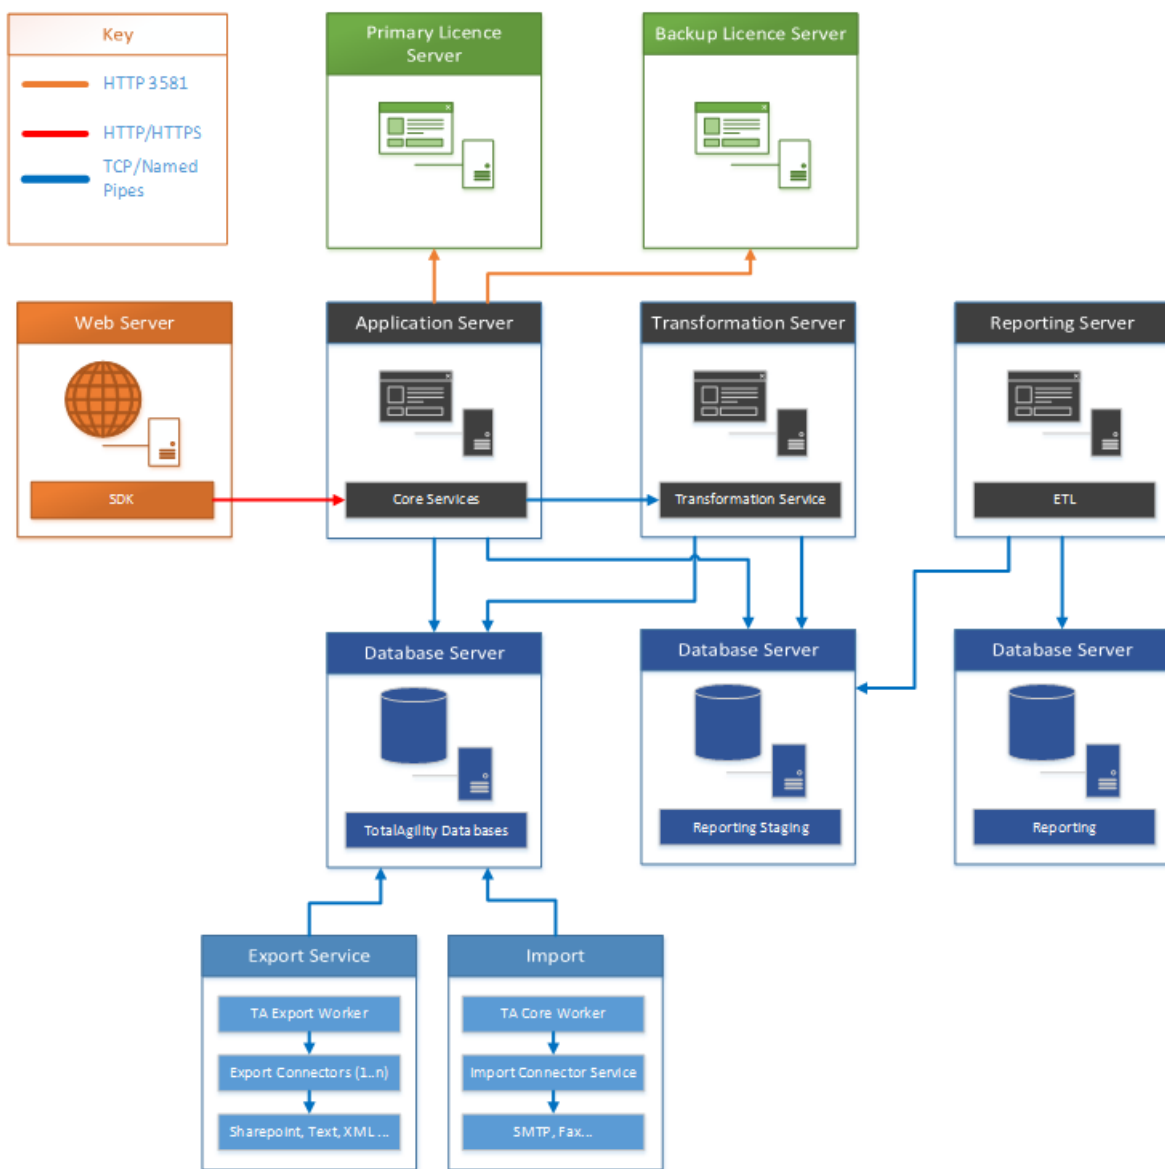
- Access to support tools
Click **Tools** and select the tool to use.
- Information about the support commitment for Kofax products
Click **Support Details** and select **Kofax Support Commitment**.

Use these tools to find answers to questions that you have, to learn about new functionality, and to research possible solutions to current issues.

Overview

The TotalAgility installation includes the components in the Logical Architecture diagram. Many components communicate with each other as shown in the Logical Architecture diagram. For more information, see the diagrams depicting communication in [Scaling and resiliency](#).

Logical Architecture



Web/App tier

The web tier contains the SDK Services and provides a fully documented API set for users to call. There is no business logic contained within the SDK Services. All API calls to the SDK services are forwarded to the Core Services. The Core Services connect to the database and also contain business logic. When a combined Web/App is deployed, the Core services reside on the Web server, and the calls from the SDK Services to the Core Services occur in memory. When a split Web/App is deployed, the Core Services reside on a separate Application server (the app tier).

Server side Windows services

The server side Windows services include:

- **CoreWorker Service:** Responsible for executing the background system tasks, such as Perform Auto Activities, Evaluate Jobs and more. See [Core Worker Service](#) for more details.
- **Transformation Service:** Polls the database for automatic capture activities to perform Image Processing, Separation, Classification, Extraction, and PDF Generation. As these operations are CPU intensive, this service must run on a dedicated server to perform its work when not running in Real Time Transformation Service mode. If running in Real Time Transformation Service mode, where it provides customers with enhanced features aimed at mobile/capture APIs, it works with in-memory documents to improve performance. See [Real Time Transformation Service](#) for more information.
- **Reporting Service:** Extracts the capture data from the Reporting Staging database, transforms it and then loads it into the Reporting Main database. This service should be run on a dedicated server as it performs CPU-intensive translation of data.
- **License Service:** Decrements user and volume license counts.
- **Streaming Service:** Used to import and export packages in the TotalAgility Designer. You must use a Windows service as IIS only supports a maximum of 2 GB for a stream. This service runs on the web server. In a split web/app, it runs on both the web tier and the app tier.
- **(Import) Message Connector:** Enables documents to be imported into TotalAgility using the configured import connectors.
- **Export Service:** Exports documents to the customer's choice of store. Once exported, the documents can be removed from the main TotalAgility database.

Database

SQL Server is the supported database.

TotalAgility on-premise supports SQL Server authentication or Windows authentication.

TotalAgility on-premise multi-tenancy supports SQL Server authentication.

Client side components

Kofax Web Capture Service is responsible for image processing, bar code and patch code detection, file import from disk, and image normalization.

Real Time Transformation Service

TotalAgility can be installed in a specialized configuration mode known as the Real Time Transformation Service. In this mode, TotalAgility provides synchronous classification and extraction capabilities that reduce response time and improve performance by storing and processing documents data in memory. You can persist this data into the document repository after classification and extraction.

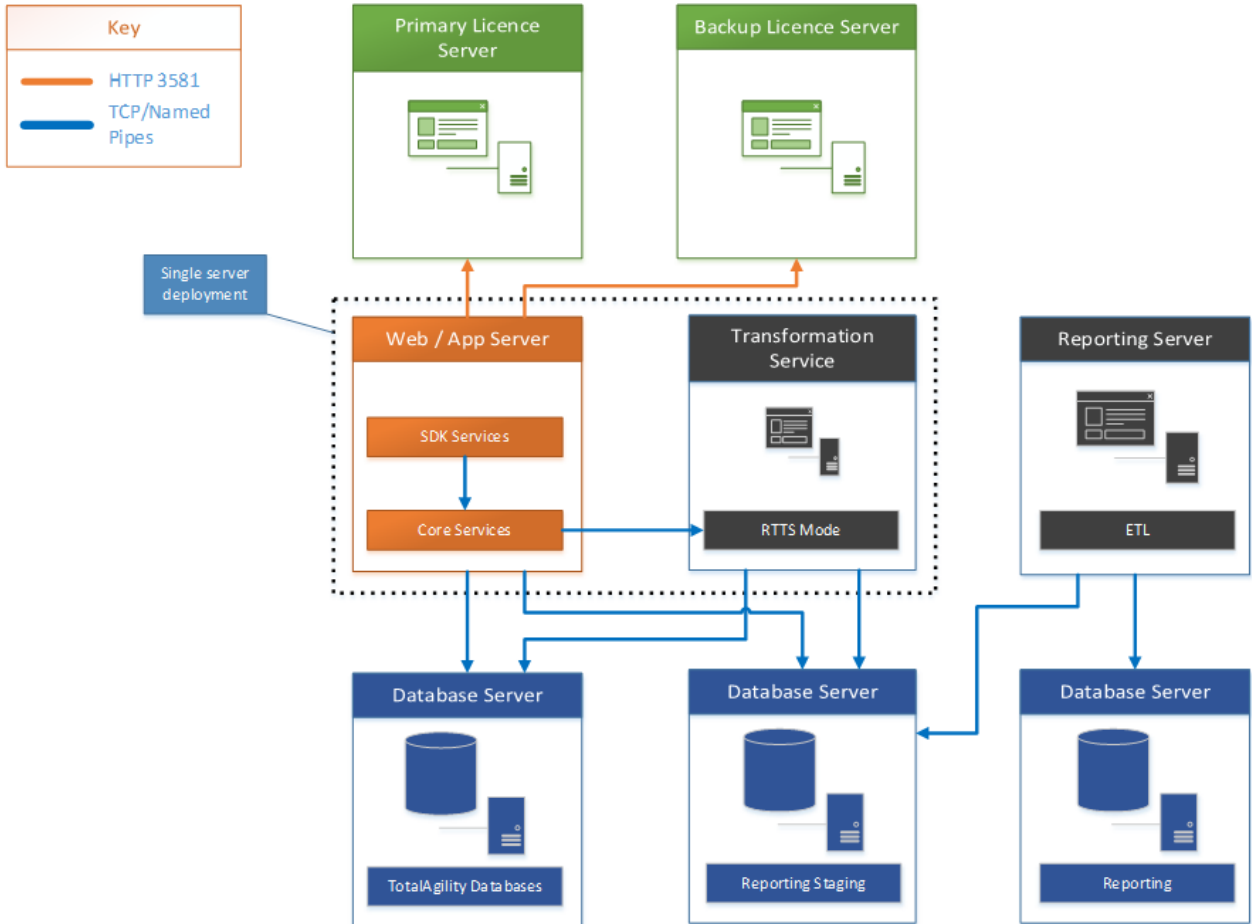
TotalAgility supports real time transformation and optimizes performance by:

- Supporting a configuration mode in Transformation Server that does not poll for pending capture activities and is used only for synchronous processing.
- Allowing Transformation Server to preload specific Transformation projects.
- Bypassing folder and document structure validation rules.

To use real time transformation, you must select the Real Time Transformation Service mode during the TotalAgility installation.

The following diagram depicts the Real Time Transformation Service architecture.

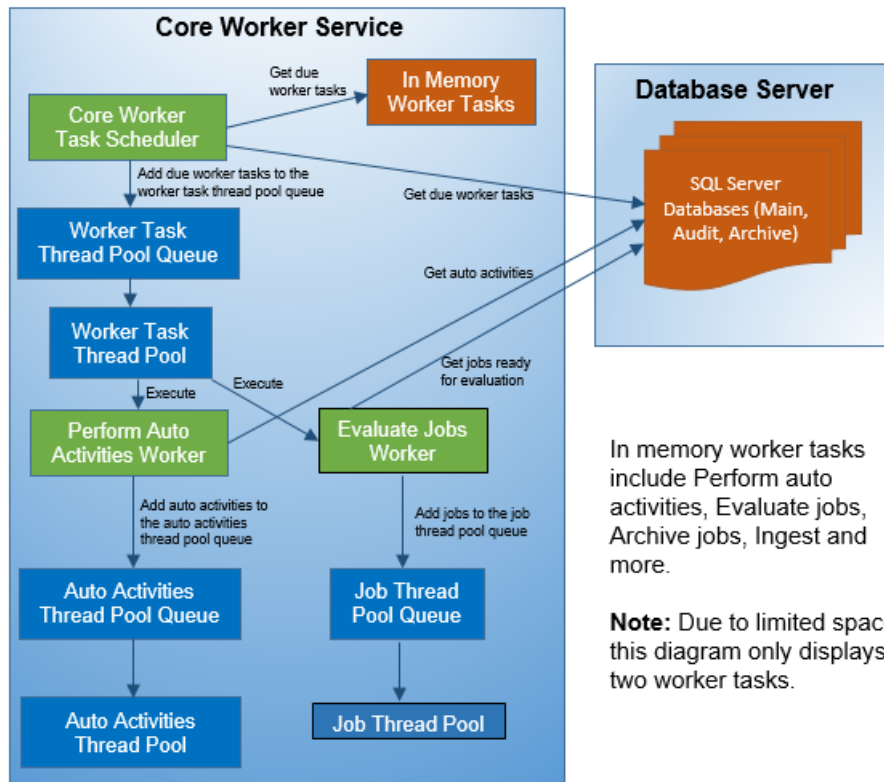
Real Time Transformation Service architecture



Core Worker Service

This section describes the inner workings of the Core Worker Service as shown in the diagram.

Core Worker Service



Background System (Worker) tasks

The Core Worker looks for non-capture System (Worker) tasks due for execution at preconfigured intervals. The frequency of their execution is based upon the interval, configured within the System Settings in the TotalAgility Designer.

In addition to a specific System task, some System tasks are also responsible for carrying out a non-system task. For example, the Evaluate Job System task is only responsible for evaluating jobs that are ready for evaluation, whereas the Archive Jobs System task not only archives completed jobs but also removes items that are marked for deletion by the Retention Policy System task.

Execution through threads in thread pools

For each System task due, a corresponding entry is added to the process Worker Task Thread Pool Queue. When a thread in the process Worker Task thread pool becomes available, the item at the head of the queue is removed and allocated to the available thread. This thread then performs the (due) System Worker task. Once the System Worker task is completed, the thread becomes available again.

For example, the thread might perform the Auto Activities Worker task. This worker task looks for automatic activities to perform. For each automatic activity found, an entry is added to the Auto Activities Thread Pool Queue. When a thread in the Auto Activities Thread Pool becomes available, the item at the head of the queue is removed and allocated to the available thread. This thread will perform the automatic activity.

Once the automatic activity is completed, the thread becomes available again. Essentially these automatic activities are executed asynchronously through threads in the Auto Activities Thread Pool.

Non-System tasks are executed asynchronously through thread pools, in a similar fashion to the automatic activities described above. In this example, both the jobs to be archived and the retention policy items to be deleted are serviced by the Job thread pool.

Types of System (Worker) tasks

This section describes the different System (Worker) task types.

In-memory / cached System Worker tasks

The Core Worker polls an in-memory collection containing the cached System worker tasks to identify which ones are due. When the Core Worker is scaled out horizontally across multiple servers, the in-memory Worker tasks can run concurrently. For example, the Auto Activities Worker task can run on both servers simultaneously.

The in-memory system tasks include:

- Perform Auto Activities
- Evaluate Jobs
- Archive Jobs
- Ingest
- Create Exceptions
- Create Jobs
- Process Activities
- Process State Actions

Non-Cached System Worker tasks

The Core Worker polls the database for the non-cached System Worker tasks to identify which ones are due. These types of System Worker tasks cannot execute concurrently across multiple servers. Example: Monitoring and Job scheduling.

Configuration

You can optimize the system for specific workloads by configuring settings, such as the number of threads in a thread pool or the queue size for a thread pool. We recommend that you keep the default configuration settings; however, you can change them based on the customer's requirements. Changing a single configuration setting may not have the desired effect because many of the configuration settings are coupled.

For example, the following settings relate to the processing of automatic activities:

- **Thread Pool size for Automatic Activities:** The number of threads in a thread pool, that is, the maximum number of concurrent automatic activities.

For TotalAgility On-Premise, you can set the thread pool size for automatic activities in the **Automatic Activity** setting available at **System Settings > Settings > Thread Pool** in the TotalAgility Designer.

For TotalAgility On-Premise Multi-Tenant and TotalAgility in Azure environment, it is not possible to configure thread pools in the TotalAgility Designer, as they are not tenant specific.

For these environments, you can set the thread pool size for automatic activities in the **AutoActivitiesThreadPoolSize** setting available in **Agility.Server.Core.WorkerService.exe.config**.

- **CoreWorkerMaxAutoActivityThreadPoolSize:** The maximum number of automatic activities queued to be serviced by the thread pool.
- **System task interval for Perform Auto Activities:** The frequency at which the Core Worker polls the database for automatic activities.
- **MaximumNumberOfActivitiesToLock:** The number of automatic activities to look for when executing the Auto Activities Worker, and subsequently pass to the Auto Activities Thread Pool to perform.

Other settings include:

- **CoreWorkerIntervallnSeconds:** Dictates how frequently the Core Worker polls for due in-memory System tasks.
- **CoreWorkerDBIntervallnSeconds:** Dictates how frequently the Core Worker polls the database for System tasks that are due.
- **CoreWorkerWorkQueueSize:** Specifies how many System tasks are requested each time Core Worker polls the database.

Example

Consider increasing the **MaximumNumberOfActivitiesToLock** parameter from 16 (default) to 50. It may result in higher throughput of activities processed; however, if the Core Worker is not processing automatic activities quickly, the thread pool queue may simply reach its limit (default: 100). In this case, the number of activities entering the thread pool queue increases in size, but the number of activities

being serviced by the thread pool remains the same. To remediate the issue, consider increasing the thread pool size and **CoreWorkerMaxAutoActivityThreadPoolSize** accordingly.

Performance Counter Usage

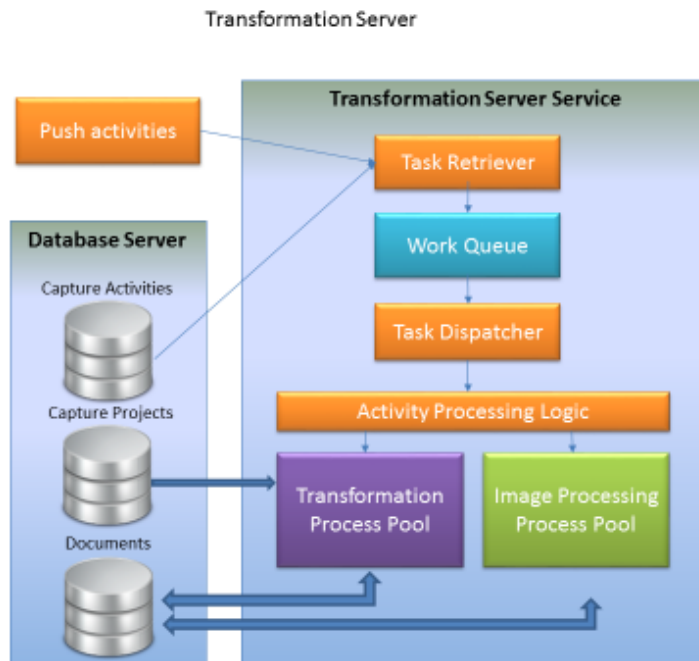
The Performance Counters are used to record application states and activities. Many are already included as part of the Windows operating system, and additional counters may be created programmatically. Counters are written and placed in an application where specific measurements need to be made, such as rate of throughput, count of items and others. Once the counters are installed on a target system, an application can write values to them. These values can then be read in real-time using the Windows Performance Monitor, a Microsoft Management Console (MMC) snap-in. TotalAgility has a number of Performance Counters; the placement of these within the application are used to measure Key Performance Indicators (KPIs) of a number of critical services.

For more information on TotalAgility Performance Counters, see the following guides:

- Kofax TotalAgility Administrator's Guide
- Kofax TotalAgility Installation Guide
- Kofax TotalAgility OnPremise MultiTenancy Installation Guide

Transformation Server

This section describes the inner workings of the Transformation Server as shown in the diagram.



Task Retriever

Fetches the capture work either from the database or directly across WCF calls from the TotalAgility Core Worker Service (known as Push activities). Work coming over the WCF channel is of high priority and is placed at the head of the work queue to be processed when resources are available. Processing resources can be reserved for these high priority activities to make sure there is no delay in processing them. The Task Retriever polls the database for more work when it anticipates that resources may be available for more processing.

Task Dispatcher

Maximizes the use of processing power on the server. The task dispatcher tries to consume all the server resources to maximize throughput unless it is configured to limit the number of documents that may be processed at the same time.

Activity Processing Logic

Handles the high level logic involved in Image Processing, Separation, Classification, Extraction, Validation, and PDF Generation.

Transformation Process Pool

Includes processes that perform Separation, Classification, Extraction, and PDF Generation using the same core transformation modules as the Kofax Transformation product.

Image Processing Process Pool

Includes processes that handle image processing tasks which use eVRS, AtalaSoft, and other libraries to perform image conversion, image processing, scaling, bar code detection and separation, patch sheet detection and separation, and others.

Ports and protocols

Refer to the following table for ports and protocols used in the internal components. Also refer to the Logical Architecture diagram in [Overview](#).

From	To	Protocol	Port
User	Web Tier	HTTP / HTTPS	80 /443
Web Tier	App Tier	HTTP / HTTPS	80 / 443
Core Services in the App Tier / Core Worker Service / Transformation Server / Reporting Service / Analytics	Database	Named Pipes / TCP	1433
Core Services in the App Tier / Core Worker Service / TS / Reporting Service	License Server	HTTP	3581
License Proxy	License Server	HTTP / HTTPS	80 /443
Core Services in the App Tier / Core Worker Service	Transformation Server	TCP	9001
App Tier	TRIM / SP / Dynamics CRM / Dynamic AX	HTTP / HTTPS	80 /443
App Tier	MS Exchange	TCP	25

Custom forms configuration

You can seamlessly extend your forms-based solution using custom code where necessary using the following two options.

1. Use the Custom Pages and Images features in the Form Designer section of TotalAgility to upload custom HTML/JS and Image files for use in any TotalAgility forms site.

This is the recommended approach as these pages and images are retained on upgrading TotalAgility.

2. Copy all custom pages to the following location:

```
<TotalAgility Installation>\TotalAgility\Forms\<sitename>
```

If you use a TotalAgility site, create a subfolder in this folder with the same name as the site and place all custom pages in this subfolder.

Note With this approach, ensure that you back up the files before upgrading TotalAgility and then copy over again after upgrade.

We recommend that you create custom forms only when the functionality required is not achievable within the Form Designer, or when the form logic becomes too complex to maintain.

Scaling and resiliency

This section describes the deployment options and technologies available in TotalAgility to scale the installation to meet high performance workloads and become resilient to interruptions and outages.

TotalAgility is a suite of stateless WCF.NET services. The architecture lets you scale your environment using standard load balancing technologies. These technologies include application clustering with MSC (Windows server 2008 server families) or WSFC (Windows 2012 server family), Application Routing Request (ARR), Network Load Balancing (NLB) and other load balancing hardware and software. TotalAgility allows each engine request to be serviced by any node in the farm that the load balancing software and hardware select. The state is maintained in a relational database, which you can scale further.

The web-based, mobile and MFP clients communicate with TotalAgility through TotalAgility Web server using HTTP/HTTPS protocols.

We recommend that you install each server component onto a separate box with the possible exception of the Kofax TotalAgility Web server and TotalAgility Application server that may be installed together.

TotalAgility Web server, TotalAgility Application server and Transformation Server can be scaled horizontally. You can improve throughput of your installation by using several instances of these services.

TotalAgility Web server is an IIS-based web application that can easily scale using standard load balancing technologies. These technologies include WLBS, NLB, Application server and other load balancing hardware and software. If you plan to install TotalAgility Web and Application servers separately, you also need to load balance traffic from the Kofax TotalAgility Web server to the TotalAgility Application server.

TotalAgility Web and Application servers

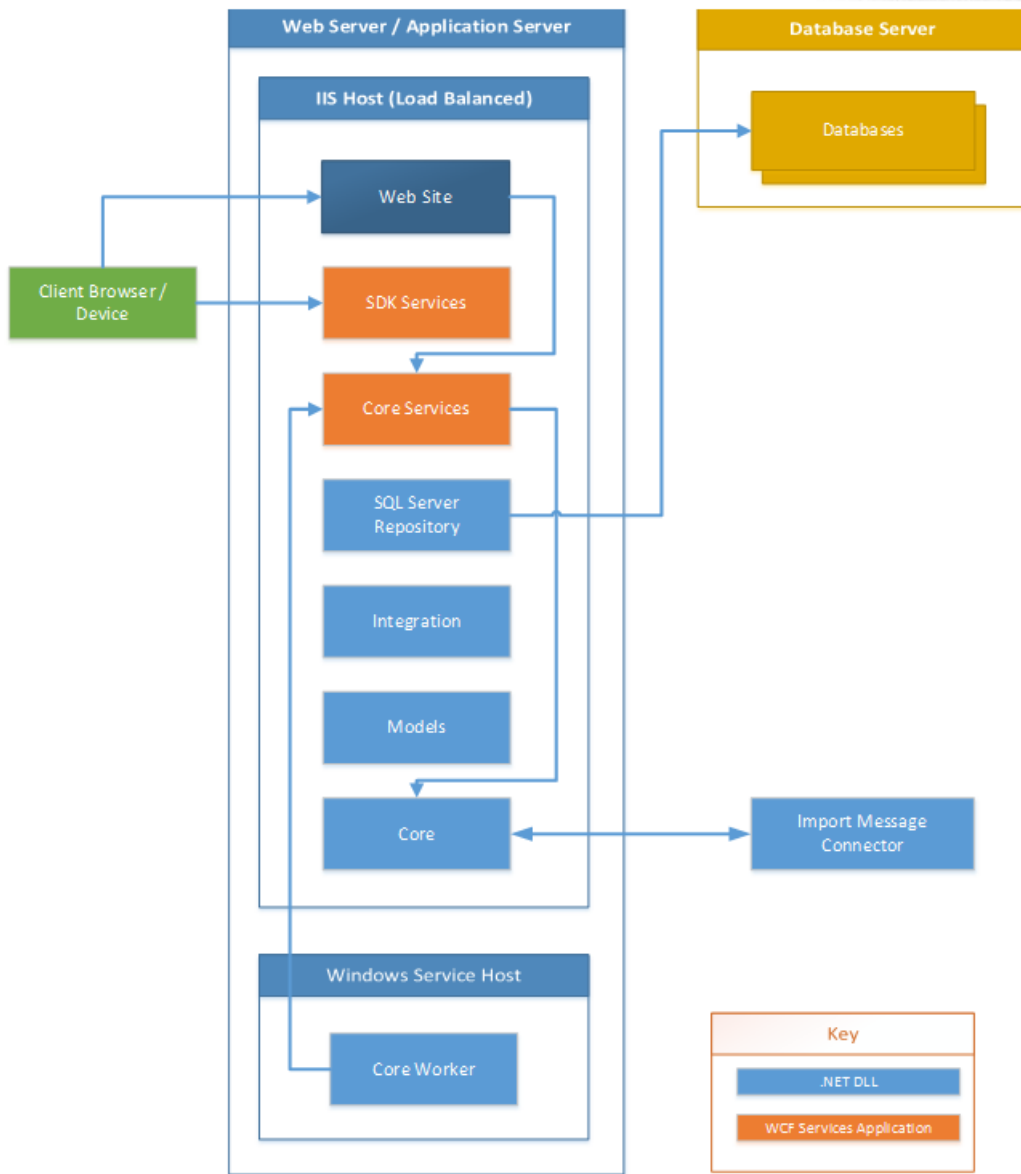
Deployment options

TotalAgility supports the following installation modes:

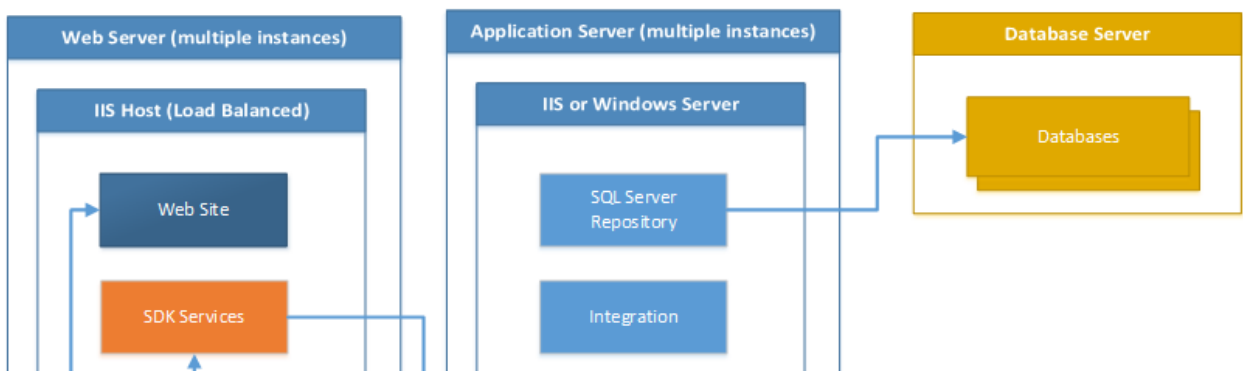
- **Combined Web and Application Server:** Installs the Web and Application servers on a single computer.
- **Web Server only:** TotalAgility Designer web pages, TotalAgility Workspace forms and SDK services.
- **Application Server only:** Core services

The TotalAgility Web Server and the TotalAgility Application server may be installed on the same computer or separately. See the diagrams below for combined Web App and separate Web and Application servers.

Combined Web App Server



Separate Web App Server



Install the TotalAgility Web server and the TotalAgility Application server on separate boxes if you have a need to exclude the possibility of accessing your enterprise internal network services (for example, databases and mail servers) from a web server, or if you want to be able to scale the TotalAgility Web server and the TotalAgility Application Server separately. Although separating the TotalAgility Web and the TotalAgility Application servers makes the installation more secure against Web attacks, having them installed on the same computer may result in better overall system performance.

Scaling

TotalAgility Web server is an IIS-based web application that can easily scale using standard load balancing technologies such as WLBS, NLB, Application server and other load balancing hardware and software. If you plan to install TotalAgility Web and Application servers separately, you also need to load balance traffic from the TotalAgility Web server to the TotalAgility Application server.

You can scale the TotalAgility Web and TotalAgility Application servers vertically (installing them on computers with more powerful hardware configurations, such as more RAM and more CPU cores) or horizontally (installing additional instances of TotalAgility Web and/or TotalAgility Application server service instances).

The TotalAgility Core Worker service on the Application server periodically polls the TotalAgility database for pending system tasks. Some examples of these system tasks are Perform Automatic Activities, Monitoring and Job Archiving.

When the system does a high volume of background processing, we recommend that you scale out the TotalAgility Application servers. Each Application server will have a separate Core Worker service instance. Each instance performs a share of the pending system tasks, thereby improving overall system throughput.

Clustering

Web servers can be hosted within a single or multiple NLB clusters, where all nodes are active at any one time, awaiting requests from the clients. NLB will route requests to the next available node (same IP address) within the respective cluster.

Use MSC (Windows server 2003 or 2008 server families) or WSFC (Windows 2012 server family) for the application server clustering where only one node is active at any one time. High availability is obtained through the creation of multiple MSC\WSFC clusters (each with its own IP address). The application cluster is scalable in that extra nodes can be added (for MSC up to eight nodes and for WSFC up to sixteen nodes).

Performance and sizing best practices

Many factors may influence the number of TotalAgility Web/Application servers that you should install:

- Hardware configuration (number of CPUs, RAM, IO and network throughput).
- Expected number of running jobs and complexity of TotalAgility processes.
- Expected load with background processing (unattended activities, system tasks, and more) that the TotalAgility Application Server should be running.
- Number of concurrent capture and non-capture user sessions (including mobile and MFP sessions).

We recommend monitoring the average CPU load for your TotalAgility Web and TotalAgility Application services. If the loads are higher than 70% for extended periods, you may need to deploy additional services. Otherwise, there may be performance degradation on peak loads.

Transformation Server

Deployment options

For better performance, the Transformation Server service should always be installed on a separate computer from other Kofax TotalAgility services.

Scaling

To improve transformation throughput you can install multiple instances of the Transformation Server service, each on a separate computer.

Performance and sizing best practices

Many factors may have an influence on how many Transformation Server instances need to be installed:

- The hardware configuration of the computer where the Transformation Server is being installed (number of CPUs, RAM, IO, and network throughput).
- Number of extraction groups and document types you have defined.
- Number of documents you plan to process hourly.
- Complexity of your document types in terms of the number of pages, number of fields, recognition rules, and so forth.
- Whether you want to use synchronous Transformation Server activity or fuzzy search queries.
- Whether you plan to use the Online Learning feature.
- Whether you configure a process to use Trainable Document Separation.
- Whether the classification is configured to use layout or content classification (the latter requiring OCR, which consumes time).

If you plan to configure several Transformation Server pools, the throughput of every pool should be calculated to determine the number of servers in it.

License Server

Deployment options

TotalAgility uses a license server to manage all licensing requests. Your Kofax serial number and product code are locked to the computer.

Note If installing a separate License Server on Windows Server 2012 R2 x64, you must have Microsoft .NET Framework 3.5.

TotalAgility supports one primary and one backup license server, and each license server needs its own activation code and license configuration.

Note The volume usage in the backup license servers can be very low (there only needs to be enough volume to last until the primary license server is restored).

If the primary license server fails, the system automatically switches to the backup license server.

The system can use the backup license server until it runs out of volume. When the primary license server is restored, TotalAgility automatically resumes using it.

Import Connector

Deployment options

The Import Message Connector is installed as part of the Application server. We recommend that you configure dedicated servers with the Import Message Connector when processing content in the production environments.

Scaling

As ingestion workload grows additional Application servers can be installed and their corresponding Import Message Connectors configured to accommodate the increase in workload.

Performance and scaling best practices

We recommend that you monitor the average CPU load for your Application server running the Import Message Connector. When CPU utilization is sustains 70% for extended periods, you may need to deploy additional Application servers.

Reporting service

The Reporting service is automatically installed with other Application server services. The Reporting service should be installed and run on its own dedicated server.

Deployment options

TotalAgility services require several Microsoft SQL Server databases:

- TotalAgility databases (main, simulation, archive, and so forth.)
- TotalAgility Document Repository database
- Kofax Reporting Staging database

- Kofax Reporting Main database

This list does not include databases installed by Kofax Analytics for TotalAgility.

TotalAgility does not support databases where the default collation of the server instance is case sensitive.

When the default collation is case sensitive, we recommend creating a new instance, ensuring the default collation of this instance is case insensitive. (for example, SQL_Latin1_General_CP1_CI_AS).

Scaling

TotalAgility supports database scaling technologies available to SQL Server such as Always-On, Active-Active, or Active-Passive deployment of database instances, and scaling via increased hardware resources.

Performance and scaling best practices


You may require high storage volume for the Reporting and Document Repository. You can create these databases on separate SQL server installations, from each other, and from TotalAgility databases, to lower the load on the SQL server hosting the TotalAgility databases.

Business continuity

In TotalAgility, business continuity means implementing a Disaster Recovery (DR) system as part of, or in lieu of, a fault tolerance or High Availability strategy.

Usually, many DR options are temporary, and many DR infrastructures do not have built-in redundancy, powerful hardware or fast network connections. In such scenarios, to resume normal production level, we recommend that you develop or augment an existing DR plan after deployment of a new TotalAgility system, specifically if you have a vertically and/or horizontally scaled-out system. See the sample WAN topology diagram.

WAN topology



Note DR is concerned with total system failure and its subsequent reinstatement; it is not concerned with failover to a node.

Typically, most customers view DR as an overly expensive, relatively redundant system until they really need it. You can consider one of the following approaches:

- Hot DR, which is fully automated.
- Warm DR with partial automation but still requiring a degree of manual intervention.
- Manual DR where every step requires human intervention.

You must perform the DR testing on a regular basis and must factor it into your DR plan.

In TotalAgility, the following factors influence your DR plan:

- On-premise or hosted, or a combination of both
- Physical hardware or virtualization, or a hybrid of both
- Deployment topology
- Licensing costs
- Regulatory and/or business compliance
- Database deployment

Some businesses outsource disaster recovery to a managed service provider or cloud hosting company, as they cannot make the investment in location diversity for disaster preparedness, or do not have experienced IT staff (or no IT staff). Virtualization is often the de facto choice as shown in the preceding diagram.

As TotalAgility is a stateless based application, the database is a key component. For an on-premise SQL Server, the following options are currently available.

- **AlwaysOn Failover Cluster Instances:** This option leverages the Windows Server Failover Clustering (WSFC) functionality to provide local high availability through redundancy at the server-instance level.
- **AlwaysOn Availability Groups:** This option is an enterprise-level, high-availability and disaster recovery solution, introduced in SQL Server 2012. Use this feature to maximize availability for one or more user databases. To use this option, the SQL Server instances must reside on Windows Server Failover Clustering (WSFC) nodes.
- **Database Mirroring:** Creates and maintains redundant copies of a database to ensure continuous data availability and minimize or avoid downtime. The future version of Microsoft SQL Server will not support database mirroring.