



EIGHTWIRE

TECHNICAL DOCUMENTATION

AUDIENCE

This document outlines the technical aspects of Eightwire and how it stores and moves data. It is intended for a technical audience who will have experience with databases, file systems and networking.

EXECUTIVE SUMMARY


Enterprises have always needed to share sensitive information across business and technology boundaries. Software products have met this need with aggregating data repositories into one platform or complex data engineering tools to move and transform data over the internet. These solutions don't work for many enterprises, so sensitive data sharing is often delivered with ad hoc tools and processes that fail to meet security, auditing and governance legislation.

Eightwire is a data sharing product that automates the movement and management of sensitive business data. It removes the majority of the work for customers to share data across boundaries while maintaining end-to-end security. With Eightwire, enterprises have a complete system-to-system solution to feed data platforms from external sources and internal silos that would be difficult to access.

EIGHTWIRE USES

Eightwire is ideal for:

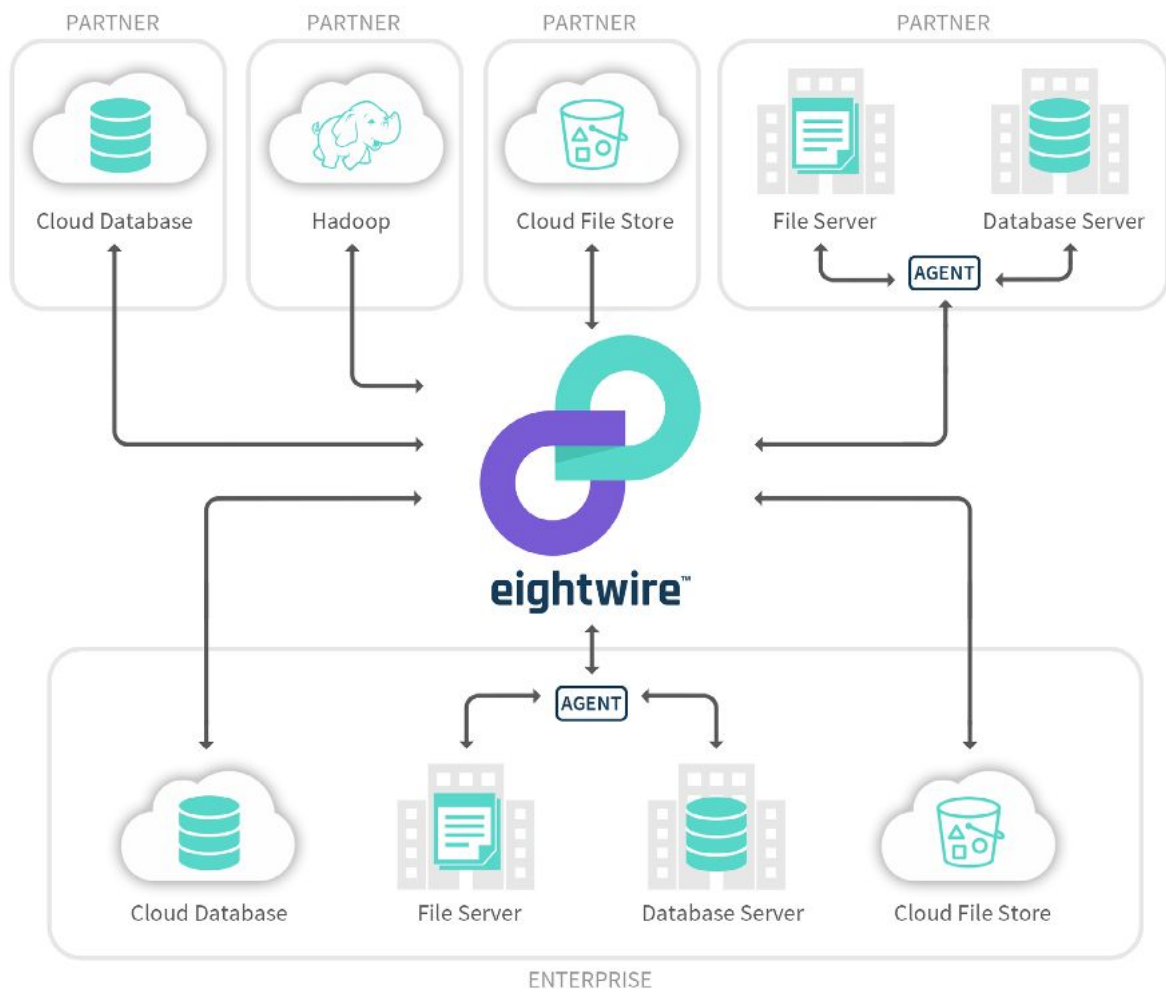
- Sharing sensitive information across separate divisions, partners, and customers.
- Populating cloud-based data warehouse platforms with sensitive data from on-premises systems.
- Aggregating data for data warehouse systems from across external partners.
- Deploying data sharing features for enterprise SaaS products that rely on sensitive data from legacy systems.



ARCHITECTURE

There are three key components of Eightwire: Agents, the Data Processor, and the Management Architecture. Agents are small, secure applications that customers download and install on Internet-connected computers that access data stores. Cloud-based data stores such as Amazon Redshift or Microsoft Azure do not require Agents and connect directly to the Data Processor. The Data Processors are controlled by Eightwire's management system and are the only part of Eightwire's infrastructure that touches customer data. Data Processors can be inside Eightwire's tenancy or on in the customer's data centre. The Management Architecture is the central nervous system of Eightwire that manages the Agents and Data Processors while providing an interface for customers to manage data sharing.

ARCHITECTURE



USING EIGHTWIRE

Eightwire automates the technical and business effort needed for system-to-system data sharing across business boundaries. All that is required is a simple 3 step process that can be carried out by developers, business analysts, or data analysts.

1. Connect to source and destination data stores

Data Stores can be cloud-based or on-premises data storage systems. This includes folders of delimited, fixed-width or Excel files, NoSQL databases or relational databases. For accessing databases on your network a small Agent is downloaded, usually with no changes to your existing firewall or infrastructure. The Agent runs as a Microsoft Windows service and relies on ODBC, OleDb, and vendor native drivers to connect to the database. The Agent can also connect to folders of spreadsheets, delimited and fixed width text files and raw content such as PDFs, images, and documents.

2. Create processes to move your data

Once the Data Stores are connected, Eightwire's algorithms carry out the mapping and transformation tasks between them. The user needs to:

Select the source and destination data stores,

Select the tables or objects to move or synchronise, and

Check source to destination mapping.

3. Execute processes to move your data

Finally, processes can be run individually or as a group. Defining a schedule allows you to control when your data is moved.

BEHIND THE SCENES

Processes are executed as batches. For each batch, incoming data is profiled to verify column names, data types, value ranges and formats. If any variation is found, Eightwire will attempt to fix the problem and continue. Likewise, any structural changes to the destination are accommodated and mapped automatically. The amount of variation tolerated is user configurable. If the data changes will result in records being rejected, Eightwire will continue to process all records and notify the user of the rejected records and the reason why.

Source data is streamed into the Eightwire Data Processor from the source database, file or application. If an Agent is used, the Agent will stream the data, if not, then Eightwire will work directly with the source systems via API calls.


The incoming data is profiled to determine structure, format and data types.

Samples from the destination system undergo the same profiling. The source and destination profiles are compared and dynamically mapped according to user-selected thresholds.

Data that will be moved to the destination is assembled into database inserts, API calls or files. These are then executed against or copied to the destination system, either by the Agent or via API calls from the Data Processor.

EXAMPLE EIGHTWIRE SOLUTION CONFIGURATIONS

GROUND TO GROUND



Using Eightwire for on-premises data sharing is carried out through installing an Agent on source and destination infrastructure and using the Eightwire web application to build and configure the processes. The Agent allows Eightwire to communicate through firewalls and securely manage data feeds from any web browser. With the Agents managing the data movement, the source and destination systems can be on any corporate networks anywhere in the world.

GROUND TO CLOUD



The solution design for ground to cloud integration requires an Agent on the internal system only. Using the Eightwire application, the on-premise data store is synced with a cloud destination data store. This can all be managed via the web application. The use of this application is the same for cloud data stores as it is for internal databases.

ENTERPRISE EIGHTWIRE



With an Enterprise Eightwire solution, the customer has its own version of Eightwire Processor installed on a server within their network. This can be on a physical server or a VM within their private cloud. The customer has the same user experience in a single tenant environment with the confidence that their data isn't leaving their network and they now have dedicated infrastructure for data management.

All communication between agent and server is encrypted and obscure — aside from the data itself, it does not contain any reference to accounts or users and can't be linked to individual customers by using information contained in the metadata related to the data transfer.

AGENT TO EIGHTWIRE SECURITY

When an Agent, acting as a source, asks for data from data stores on your network (whether it is data from a file or database) it compresses, encrypts, and sends the data to secure Eightwire servers for processing.

Eightwire has deployed a security model with perfect forward secrecy using two separate chains of trust. Every batch is encrypted with one randomly generated certificate for the HTTPS transport to the Data Processor using the TLS 1.2 cryptographic protocol. At the same time, the payload is encrypted with a separate, randomly generated AES 256 cipher. Generating two ciphers for each batch adds processing overhead, but allows Eightwire to:

1. Prevent an adversary from breaking both certs from being able to decrypt any other data batches and,
2. Lock Eightwire out of customer data so it cannot view or decrypt customer data that has passed through the Data Processor.

An Agent does not “listen” to the Internet for a program to unlock its security. The Agent only calls out to Eightwire servers. User and Agent security is handled and monitored through the Eightwire server and is designed not to be “spoofed” or intercepted by external or internal programs. All communication between Agent and server is encrypted and obscure—aside from the data itself, it does not contain any reference to accounts or users and can't be linked to individual customers by using information contained in the metadata related to the data transfer.

INTERNAL EIGHTWIRE SECURITY

When Eightwire processes data, it is deleted as soon as the process has completed. All data is encrypted at rest using the AES cipher and is only decrypted in the memory space during processing before being re-encrypted in the final format on the server. The data processing servers are not backed up so it is not possible for customer data to leak into backups and the memory space is overwritten between data processing batches. Finally, all communications between the Eightwire servers are encrypted and customer data is never shared between Management Architecture and Data Processing architectures.

COMPLEMENTARY TO MAJOR ETL PROVIDERS

Eightwire is designed to work with Enterprise Extract, Transfer, Load (ETL) products that customers may already be using in their organisation. The benefits of existing ETL platforms can still be realised without requiring any retooling or code migration. Eightwire significantly enhances enterprise ETL by surrounding the core enterprise applications with performant, cost-effective remote data access for sensitive or legacy data. In provisioning information from remote systems, Eightwire provides significant return on investment by replacing hand coding for the initial design, agile sprints, maintenance, and documentation efforts

