



COMPARISON OF SMARTDATA FABRIC® WITH CISCO® DATA VIRTUALIZATION (FKA COMPOSITE SOFTWARE)

REVISION 1.9

Introduction

Cisco Data Virtualization (fka Composite Software) is probably the closest competitor WhamTech SmartData Fabric® has, as from their sales and marketing material, the Cisco solution seems similar to WhamTech, however, Cisco's approach is very different from WhamTech's. The reason for Cisco being considered closer than other data virtualization, federation and integration vendors to WhamTech is that Cisco assumes a distributed approach similar to WhamTech, versus a centralized approach of these other vendors.

Fundamental Difference in Approaches

Cisco is a combination of a middleware product and conventional federated adapters, similar to others offered by other vendors, e.g., Denodo®, IBM® InfoSphere®, Oracle® Fusion Middleware, Red Hat® JBoss® Data Virtualization and Tibco®. Queries are submitted by applications to the middleware, from the middleware to the conventional federated adapters and executed by data source systems. Query success and performance are limited by data quality and data source system capabilities, and query loads are absorbed completely by data source systems. Cisco can improve data quality on results and can add value by caching result sets and making them available to subsequent queries for that data.

WhamTech SmartData Fabric® EIQ Adapters™, in contrast to Cisco and all other data virtualization, federation and integration vendors, query WhamTech's own independent indexes and indexed views that are built and maintained in near real-time using cleansed, transformed and standardized data. This obviates the need for much of the work that conventional federated adapters and middleware perform to overcome both data and data source system deficiencies. And alleviates almost all loads on data source systems, significantly improves query success and performance, and allows far more capabilities.

Cisco in General

Cisco is primarily a federated data access system with middleware, conventional adapters and results cache:

- G1. Cisco relies on conventional federated adapters to submit queries to data source systems for execution.
- G2. Cisco builds materialized indexed views of past query results data for subsequent queries – elements of a distributed data warehouse/data mart.
- G3. Cisco establishes PK-FK relationships within data sources and across data sources based on schema-level keys, e.g., LAST_NAME in one data source = FAMILY_NAME in another.
- G4. Cisco performs ETL operations to applications, data warehouses and data marts that can include cleansing, transforming and standardizing data AFTER results data retrieved.
- G5. Cisco OEM their products to Informatica, among others.

Cisco Advantages Compared to WhamTech

- A1. Cisco is an established vendor with a large customer base.

WhamTech has worked in different type of projects, including specialist products for Web and enterprise search, and eDiscovery, and for large federal system integrators on DOD and

intelligence community projects. While WhamTech has had commercial customers in the past and is still working with large federal system integrators, in the last few years, WhamTech has primarily focused on commercial markets and is fast-gaining a reputation for being disruptive, particularly in healthcare, with some of the largest healthcare companies are customers.

- A2. Cisco probably has 80% solutions for standard applications such as SAP, Siebel, Oracle E-Business Suite, PeopleSoft and salesforce.com.

WhamTech is building its portfolio of data sources it can work with, including most common data warehouses/databases, mainframe data files, office files, Web pages, application data, email, social media, Big Data, streaming and IoT devices. WhamTech can work with almost any data source and almost any readable data, regardless of the data source or data within a data source being structured, unstructured or semi-structured.

- A3. Cisco does not need to build and maintain indexes for data sources, as it relies on data sources to execute queries through conventional federated adapters. Once results are obtained, they can be cached and indexed for subsequent queries in materialized indexed views.

WhamTech usually does have to build and maintain its own indexes and execute its own queries, but that is a positive capability for many reasons, as listed in the next section “Cisco Disadvantages Compared to WhamTech”.

- A4. Cisco does not rely on data source systems once cached result materialized indexed views are built, as long as the results data does not change in the data source, but that introduces the problem of cache management. See separate white paper “#WhamTech Competitor Comparison vs. Cisco, Denodo, Red Hat and Other Conventional Federated Data Access Vendor's Use of Cache”.

WhamTech generally returns to data sources to obtain results, but can store data locally and/or cache latest results data if needed. WhamTech prefers to leave data in sources, as this is the main benefit, and maintains indexes instead of the more difficult task of maintaining results cache. While locally stored data is simple to maintain, results cache is not, depending on how source data is changing. Results cache is particularly difficult to maintain in the Big Data world of large volume and high velocity of data, unless results consist of historic data that does not change. Again, see the above-referenced white paper for a more in-depth discussion of the reasons for and the problems of cache.

Cisco Disadvantages Compared to WhamTech

- D1. Cisco does not solve the initial data source access problem that plagues conventional federated data access – potentially poor data quality, and limited indexes and query processing. This directly impacts query success on original data sources, query performance and scalability.

WhamTech cleanses, transforms and standardizes data for indexes and results, has its own independent indexes and scales for high performance distributed parallel query processing on these indexes.

- D2. Cisco makes copies of potentially confidential/sensitive data in indexed cache/databases.

WhamTech leaves most data in sources until results need to be retrieved, subject to the appropriate role-based access control, and dynamic data masking, tokenization and encryption.



- D3. Cisco imposes significant query loads on data sources and query performance is completely dependent on data source systems.

WhamTech almost completely absorbs the query processing load. There is a small data source system load for sequential reads of final and sometimes interim results data, but this is minimal (< 5%) compared to full query processing. WhamTech has built out systems capable of simultaneously indexing and query processing 10s of billions of records and 10s of TBs per day.

- D4. Cisco has no obvious mechanism to update the results cache, although it is assumed that some mechanism is in-place.

WhamTech in a near real-time index update architecture, has the latest index data and results data available, with several options to account for any latency differences.

- D5. Cisco has no obvious capabilities for derived value indexes such as pre-aggregations, pre-calculations, pre-joins, virtual and material indexed views, fuzzy match, text search, information geometry (for structured and unstructured data categorization), entity extraction, and Link Indexes for link mapping, link analysis/graph database and link visualization.

WhamTech offers all of the above capabilities. Part of the benefits of pre-aggregations, pre-calculations, pre-joins and indexed views is that these can be maintained in near real-time and automatically used by WhamTech SmartData Fabric® EIQ Adapters when part of a query. This avoids relying on potentially stale results cache or the adapter requesting raw data from a source and performing on-the-fly aggregation, calculation or view materialization, which can considerably slow down reporting, BI and analytics.

WhamTech future development will allow the automatic creation of some of the derived value indexes and indexed views when their use would be beneficial to a query. Such indexes and indexed views could be deleted if not reused in a certain timeframe.

- D6. Cisco has so-called real-time monitoring, but that represents a constant polling of data sources, and therefore query load, to check if any relevant changes have occurred.

WhamTech directly monitors real-time updates to its own indexes, including derived value indexes mentioned in D5, and can resubmit relevant stored queries, automatically update operational/real-time monitoring dashboards, and/or send alerts/notifications to subscriber/subscribing applications. The combination of event processing, stored procedures at the SmartData Fabric® level and open source BPM software allows for extraordinary capabilities usually addressed by additional separate applications.

- D7. Cisco has no form of Link Indexes that allow record-level relationship data to be combined with content data.

WhamTech uses combinations of queries on content indexes and Link Indexes to enable high performance internal and external joins, degrees of separation queries, link mapping, link analysis and ontological data model representations, as well as enabling virtual solutions, such as CDI-MDM/MPI, CRM and other solutions.

- D8. Cisco seems to present potentially complex data models to applications, requiring significant schema mapping work.



WhamTech's approach is to employ relatively simple reference data models, such as NIEM, HL7, FHIR, ACORD and XBRL – same data from multiple data sources are mapped to the same reference data view. The reference data view can, in-turn, be accessed as though (i) data objects (Web/data services and ontological view), (ii) virtual relational view and (iii) Big Table/single flat file view. Various associated query languages can be used with conversion to SQL, e.g., OQL, SPARQL, PL-SQL and NoSQL.

D9. Cisco, in the worst-case scenario when a data source is unavailable, can only access cached results data.

WhamTech can invert any and all data from indexes, assuming that all needed data is indexed. WhamTech indexes and/or own or third-party storage can also be used for archive and/or time-slice data. This last approach is used to store highly compressed and aggregated IoT device data that is accessed using standard drivers, APIS and Web/data services, and SQL.

D10. Cisco is probably very expensive to scale using such an approach, plus, installing and configuring such a system to overcome data and data source system deficiencies are probably manually intensive and therefore very expensive as well. Typically, at least five times as much cost over an extended period is needed to achieve the basics associated with conventional data virtualization, federation and integration solutions, because of all the data and query issues that are attempted to overcome.

WhamTech inherent approach overcomes many of the time consuming and therefore expensive processes needed for Cisco to work.

Cisco and WhamTech Working Together?

This question has been asked a few times. Cisco could replace conventional federated adapters with WhamTech's EIQ adapters, but that might require an internal shift away from their own perceived competitive advantage. WhamTech has experienced this before with a large middleware company that realized that WhamTech obviates the need for much of what they leverage and market, leading to an internal not invented here (NIH) rejection, even with customers expressing dissatisfaction with their product and approach. One such large middleware company sold for a relatively small amount shortly after we were in discussions with them.

For more information, please contact: Gavin Robertson, CTO & SVP, gavin.robertson@whamtech.com, (972) 991-5700 x206



About WhamTech, Inc.

WhamTech, Inc. (WhamTech) is a privately-held US-owned Delaware Corporation established in October 2000 and based in Dallas, Texas. WhamTech's mission is to develop indexed adapter-based distributed virtual data management, master data management, analytics and security platform software products. WhamTech develops these products to anticipate, meet and exceed the demands of customers seeking an alternative to the conventional approaches of data warehousing, federated data access with conventional adapters, and enterprise search. WhamTech's goal is to provide an improved and more seamless way to work with data, by leaving it in sources and changing the way fundamental and advanced data management is addressed. Most WhamTech adapter products leverage independent, cleansed, transformed and standardized indexes that execute both structured and unstructured queries, and seamlessly and automatically integrate master data management to provide capabilities normally associated with multiple separate solutions, including providing results when data sources are unavailable and for archive.

Information on WhamTech solutions, sales and services, and partnership and investment opportunities can be obtained through whamtech.com.

Copyright © 2017, WhamTech, Inc. All rights reserved. This document is provided for information purposes only and the contents hereof are subject to change without notice. This document is not warranted to be error-free, nor subject to any other warranties or conditions, whether expressed orally or implied in law, including implied warranties and conditions of merchantability or fitness for a particular purpose. We specifically disclaim any liability with respect to this document and no contractual obligations are formed either directly or indirectly by this document. Product names may be registered trademarks of their respective owners.