



Composite Data Virtualization

Composite Data Virtualization Usage Patterns

Composite Software

January 2011

TABLE OF CONTENTS

- INTRODUCTION 3**
- DATA FEDERATION 4**
 - PROBLEM – DATA CONSOLIDATION IS NOT THE ONLY SOLUTION 4
 - SOLUTION – FEDERATE DATA FASTER, FOR LESS WITH COMPOSITE 4
 - LEARN MORE 4
- DATA WAREHOUSE EXTENSION 5**
 - PROBLEM – BUSINESS CHANGE OFTEN OUTPACES ENTERPRISE DATA WAREHOUSE EVOLUTION. 5
 - SOLUTION – USE COMPOSITE DATA VIRTUALIZATION MAXIMIZE VALUE OF ENTERPRISE DATA WAREHOUSES..... 5
 - LEARN MORE 6
- DATA VIRTUALIZATION LAYER 7**
 - PROBLEM – BUSINESS SUCCESS REQUIRES SHARING DATA ACROSS THE ENTERPRISE 7
 - SOLUTION – COMPOSITE DATA VIRTUALIZATION LETS ENTERPRISES SHARE DATA MORE EASILY. 7
 - LEARN MORE 8
- CLOUD DATA INTEGRATION 9**
 - PROBLEM – CLOUD COMPUTING ADDS NEW DATA SILOS 9
 - SOLUTION – COMPOSITE DATA VIRTUALIZATION INTEGRATES CLOUD DATA..... 9
 - LEARN MORE 10
- CONCLUSION..... 11**

INTRODUCTION

Large enterprises and government agencies face similar challenges:

- Constant business change necessitates rapid IT response
- Growing data volumes and complexity increase risk and reduce agility
- Financial constraints necessitate cost-effective IT solutions

Following virtualization's established path of storage, servers, applications, and now data itself—data virtualization is the proven way to overcome these challenges and fulfill critical information needs, significantly faster with far fewer resources.

Data virtualization is used to integrate data from multiple, disparate sources - anywhere across the extended enterprise - in a unified, logically virtualized manner for consumption by nearly any front-end business solution, including portals, reports, applications, search, and more.

Data virtualization is typically deployed at two levels, often as a complement to other integration approaches such as consolidation or replication. At the project level, it can virtually integrate the data required in support of a specific application or use case. On an enterprise level, it can be implemented as common services or as a loosely-coupled data abstraction layer to share data across multiple solutions and use.

Composite Software is the only company that focuses solely on data virtualization. As our customers across multiple industries adopt our offerings, four common usage patterns have proven valuable including:

- Data Federation
- Data Warehouse Extension
- Data Virtualization Layer
- Cloud Data Integration

This white paper introduces these usage patterns, the business and IT challenges behind them, and the solution that Composite data virtualization delivers to overcome these challenges.

DATA FEDERATION

Problem – Data Consolidation Is Not the Only Solution

Everyday IT is challenged to provide their business colleagues with new information typically sourced from existing data silos. When choosing between physical consolidation and virtual approaches, architects have to consider a number of business, data source and data consumer considerations. Increasingly, data federation is proving the right approach, especially when:

- Business needs a solution fast
- Business and IT do not have a sizable budget to spend on infrastructure and staffing
- IT wants minimize the risk involved in deploying a new solution, queries and ETL scripts.

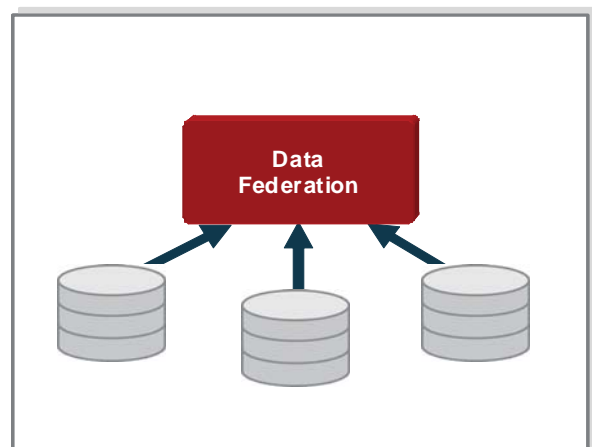
Technical challenges include:

- How do you ensure that queries perform across the network, database, and memory?
- How do you join XML and SQL structures at both design and at runtime?
- How do you apply source and consumer level security access and authentication rules?
- How do you retrieve data from diverse sources and complex application schemas including Relational, Files, Mainframe, Web Services, SAP, Oracle E-business, Siebel, Salesforce.com, and more?
- What happens when a source is not available 7x24?

Solution – Federate Data Faster, for Less with Composite

Composite helped pioneer data federation starting in 2002 and today provides the most proven data federation offering in the market. Further, Composite provides architects and developers with the most flexible pallet of federation options including:

- Federated Views
- Data Services
- Data Mashups
- Caches
- Virtual Data Marts
- Virtual Operational Data Stores



Learn More

For more insight on these data federation usage options including more specific data federation challenges, Composite data virtualization solutions, and actual customer examples, please go to <http://www.compositesw.com/index.php/solutions/data-federation/>.

DATA WAREHOUSE EXTENSION

Problem – Business Change Often Outpaces Enterprise Data Warehouse Evolution

Supporting critical, yet ever changing information requirements in an environment of ever increasing data volumes and complexity is a challenge well understood by large enterprises and government agencies today.

This inexorable pressure has and will continue to drive the demand for enterprise data warehouse (EDW) centric solutions, to support the array of business intelligence applications that rely on data sourced from the EDW.

However, business change often outpaces enterprise data warehouse evolution. And while useful for physically consolidating and transforming a large portion of enterprise data, significant volumes of enterprise data continues to reside outside the confines of the EDW. Further, enterprise data warehouses themselves require support throughout their lifecycle, driving demand for solutions that prototype, migrate, extend, federate and leverage EDW assets.

Technical challenges include:

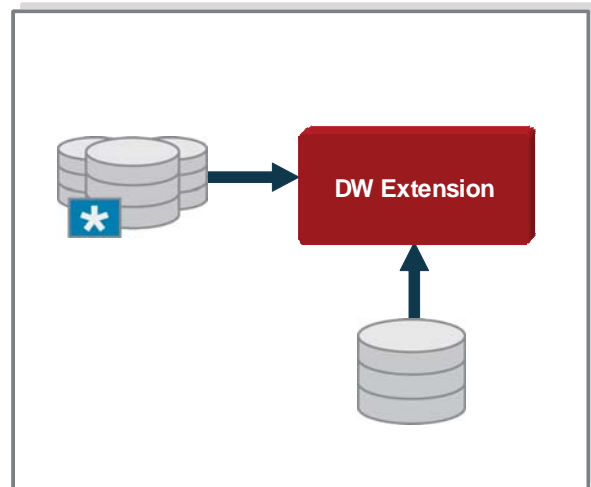
- Will your ETL developers be productive using complementary development tools?
- How do you ensure compatibility with the BI tools that consume the data?
- How do you efficiently and effectively join multi-dimensional, relational, and hierarchical structures?
- How do you retrieve data from diverse sources and complex application schemas beyond the EDW including Relational, Files, Mainframe, Web Services, SAP, Oracle E-business, Siebel, Salesforce.com, and more?
- How do you ensure reliable, scalable operations?

Solution – Use Composite Data Virtualization Maximize Value of Enterprise Data Warehouses

The Composite Data Virtualization Platform complements your EDW by providing a range of flexible data integration techniques that let you preserve and extend existing enterprise data warehouse investments. By extending the warehouse schema to include additional data, data virtualization can deliver greater business value from existing warehouse and non-warehouse data assets.

In practice, Composite customers use the Composite Data Virtualization Platform to augment their EDWs in the eight ways listed below. This flexibility allows users to solve immediate problems quickly while avoiding the costs, risks, and time required to modify the warehouse structures and supporting ETL themselves. These options include:

- Data Warehouse Augmentation
- MDM Hub Extension – 360 View
- Data Warehouse Federation
- Hub & Virtual Spoke
- Enterprise Architecture
- Complementing ETL
- Data Warehouse Prototyping
- Data Warehouse Migration



Learn More

For more insight on these data warehouse usage options including more specific data warehousing and business intelligence challenges, Composite data virtualization solutions, and actual customer examples, please go to <http://www.compositesw.com/index.php/solutions/data-warehouse-extension/>.

In addition, you can also read our white paper, Eight Ways Composite Data Virtualization Adds Value to the Enterprise Data Warehousing available at <http://www.compositesw.com/index.php/resources/white-papers-reports/> .

DATA VIRTUALIZATION LAYER

Problem – Business Success Requires Sharing Data Across the Enterprise

Enterprise data sharing is becoming increasingly popular in enterprises and government agencies where large numbers of information consumers use a range of analysis and reporting tools to access and analyze large amounts of diverse data, both from disparate sources and multiple geographical locations.

However, adopting a data virtualization layer at enterprise scale presents significant challenges. Existing legacy architectures, different project priorities, and previous SOA investments are but a few of the factors

Technical challenges include:

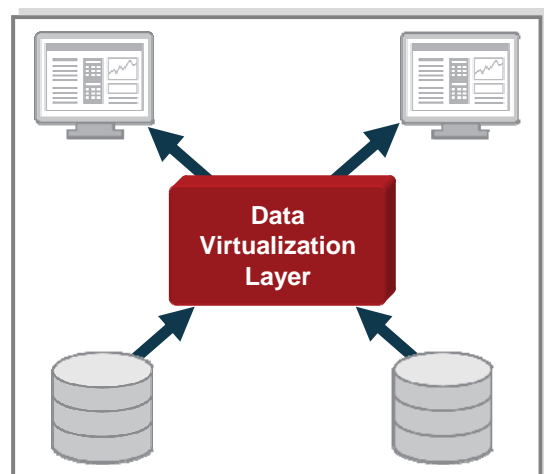
- How much data will you virtualize? And how much will you persist?
- Will standards allow your applications developers be productive using new tools?
- How do you support all the standards used by the diverse consuming applications including REST, SOAP, JMS, JSON, JDBC, ADO.Net and more?
- How granular should your data services be as you balance diversity versus reuse?
- How do you conform your data services to industry or internal standards?

Solution – Composite Data Virtualization Lets Enterprises Share Data More Easily

Composite combines SOA principles including, decoupling, reuse, and agility with key information governance principles such as abstraction, shared semantic models and data standards enabling organization to build and deploy enterprise-wide data virtualization layers in a simpler, faster, more consistent, and scalable manner.

The Composite Data Virtualization Platform supports several data virtualization layer deployment types so you can flexibly provide the information required to meet broader enterprise information requirements in a more consistent way including:

- Shared Data Services
- Data Abstraction Layer
- Standards-compliant Data Services Layer
- Distributed Data Virtualization Layer
- Forrester Information-as-a-Service
- Gartner Data Services Layer



Learn More

For more insight on these data virtualization layer usage options including more specific service oriented architecture and abstraction challenges, Composite data virtualization solutions, and actual customer examples, please go to <http://www.compositesw.com/index.php/solutions/data-virtualization-layer/> .

CLLOUD DATA INTEGRATION

Problem – Cloud Computing Adds New Data Silos

With the advent of Software-as-a-Service (SaaS), Platform-as-a-Service (PaaS), and Infrastructure-as-a-Service (IaaS), cloud computing offers a range of functionality and computing resources at an attractive, pay-as-you-go price.

However, if you're an enterprise or government agency seeking to leverage these capabilities, you've discovered introducing cloud computing increases your data integration complexity. Equal to a new data silo, each new cloud source must be integrated with your existing on-premise information sources. Further, cloud data integration requires new integration methods not typically supported by traditional approaches such as direct database queries and ETL scripts.

For example, when you integrate data from a SaaS provider such as salesforce.com with an internal customer data warehouse, you will need deep knowledge of salesforce's APIs, the ability to query data through a firewall across the Web, and on-demand, rather than batch mode, operation.

Technical challenges include:

- How do you optimize internet, internal network, database, memory so that queries perform across cloud and on-premise sources?
- Which access and authentication security rules apply? Is encryption required?
- How do you get data out of complex on on-premise and cloud application schemas such as SAP, Oracle E-business, Siebel, and Salesforce.com?
- What happens when on-premise or cloud source is not available?
- How do you ensure reliable operations in mixed on-premise and cloud environment?

Solution – Composite Data Virtualization Integrates Cloud Data

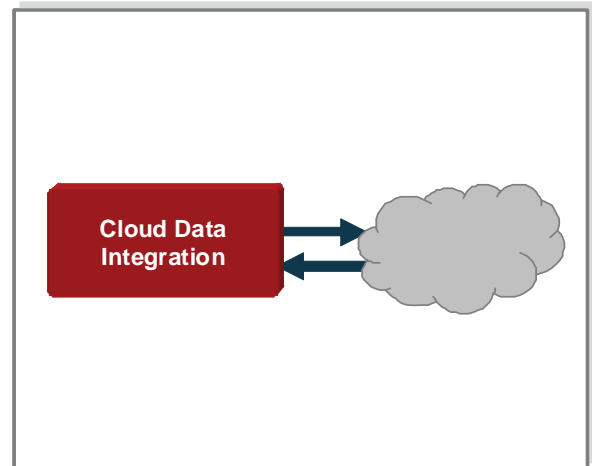
The Composite Data Virtualization Platform allows enterprises like yours to flexibly integrate on-premise and cloud data as required to meet a range of use cases. Among its key capabilities that are especially well suited to cloud data integration are:

- Composite supports both database developers and Java-centric developers with a relational development environment, Composite Studio, for those who prefer relational modeling-oriented development tools when building federated views and data services, and an Eclipse IDE, Composite Designer, for XML/SOA developers.
- Data integration services authored and run on the Composite Information Server are perfectly suited to operate across the Internet.
- Because Composite data virtualization accesses, federates, abstracts and delivers queried data on demand, no additional cloud-based data storage is needed.

- Further, Composite supports multiple security models to ensure secure access and delivery of appropriate data to authenticated and approved internal and external users.
- With pre-built Composite Application Data Services available for popular SaaS applications such as salesforce.com as well as IaaS cloud-hosted applications such as SAP and Oracle E-Business, Composite simplifies complex data models thereby unlocking data for use in the cloud.

Learn More

For more insight on these cloud data integration challenges, Composite data virtualization solutions, and actual customer examples, please go to <http://www.compositesw.com/index.php/solutions/cloud-computing/>.



CONCLUSION

Composite data virtualization, as described in this paper, is a proven data integration approach used by large enterprises and government agencies today to overcome business and IT challenges and fulfill critical information needs, significantly faster with far fewer resources.

In this paper, four commonly deployed data virtualization usage patterns were identified. Each pattern includes the challenge, the Composite data virtualization solution, diagrams, and guidance on where you can learn more.

If your enterprise is facing similar challenges, consider Composite Software, the gold standard in data virtualization.

ABOUT COMPOSITE SOFTWARE

Composite Software, Inc. ® is the data virtualization gold standard at ten of the top 20 banks, six of the top ten pharmaceutical companies, four of the top five energy firms, major media and technology organizations; and multiple government agencies.

These are among the hundreds of global organizations with disparate, complex information environments that count on the Composite to increase their data agility, cut costs and reduce risk.

Backed by nearly a decade of pioneering R&D, Composite is the data virtualization performance leader, scaling from project to enterprise for data federation, data warehouse extension, data virtualization layer and cloud data integration.

Founded in 2002, Composite Software is a privately held, venture-funded corporation based in Silicon Valley. For more information, please visit www.compositesw.com.