

the **Availability Digest**

www.availabilitydigest.com

Attunity Integration Suite

December 2010

Though not focused primarily on high availability in the uptime sense, the Attunity Integration Suite (www.attunity.com) provides extensive functionality necessary to ensure the availability of disparate data sources and application services to the enterprise as a whole.

In today's corporate environments, many of the critical applications and much of the corporate data is still contained in legacy systems. The Butler Group estimates that 80% of all corporate data is held in nonrelational databases. IDC estimates that 40% of all application development efforts are still focused on expanding existing legacy data.¹

Large and mid-sized firms alike have millions of lines of code executing on proprietary platforms. These systems process a large proportion of the transactions that are the lifeblood of the company. However, as today's competitive situation becomes more intensive and more complex, it is mandatory that these silo data sources and applications be able to interoperate with each other.

This is the goal of the Attunity Integration Suite.

The Need for Integration

The Attunity Integration Suite comprises three products:

- Attunity Connect
- Attunity Federate
- Attunity Stream

These services address the following four fundamental needs of data and application integration.

Information Access

Applications must be able to access information in remote databases, whether those databases are modern SQL databases or legacy nonrelational databases. Information access is provided by Attunity Connect.

¹ W. H. Inmon, Dan Meers, Maximizing the "E" in Legacy Extract, Transform, & Load (ETL), *White Paper*, December 2003.

Application Access

Applications must be able to access the services of other applications hosted on remote systems. Remote systems range from Windows to mainframe platforms. The applications that must be accessed may be written in any language from COBOL and RPG to C# and Java.

Application access is provided also by Attunity Connect.

Single View of Disparate Data Sources

A single view of multiple relational and nonrelational data sources must be possible. This function is provided by Attunity Federate.

Data Movement

Data must be moved efficiently from data sources to other platforms. Only data changes should be moved, not entire tables or files. This function is provided by Attunity Stream using Attunity's Change Data Capture (CDC) technology.

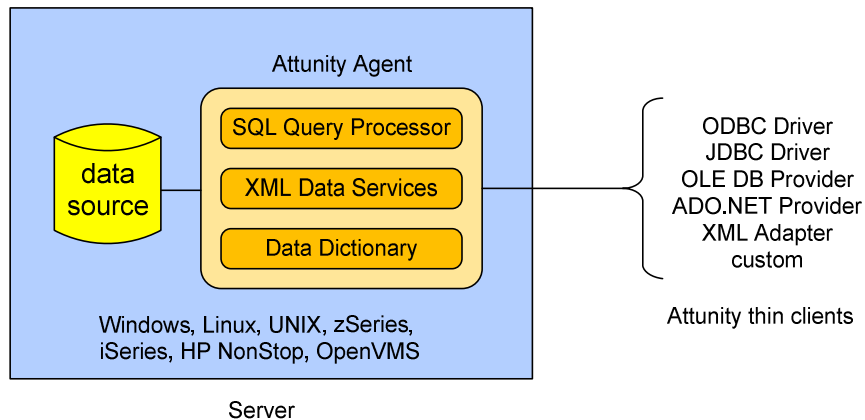
Attunity Connect

Attunity Connect provides two connection services:

- Access to remote data sources.
- Access to remote applications.

Data Connectivity

Attunity Connect provides direct read/write access to a wide range of SQL and nonrelational databases. Connectivity is provided by an Attunity Agent that resides on the database server and that interfaces with the data source. There is an agent for every supported data source.



Attunity Connect – Data Connectivity

Data sources supported by Attunity Connect include most common nonrelational and relational data sources in use today. Agents are provided for the following data sources:

Relational	Nonrelational
Oracle	IMS/DB
DB2	VSAM
SQL/MP	ISAM
SQL Server	Enscribe
Sybase	RMS
Ingres	Adabas
Informix	text files
Oracle Rdb	
CODASYL	

An Attunity Connect Agent imports the schema for its data source and stores it in its Data Dictionary. For nonrelational data sources, it creates a relational model and schema that can be used for SQL operations.

A SQL Query Processor optimizes SQL queries and ensures against “bad queries” that can be detrimental to performance. XML Data Services provide an XML interface for Web services and Service Oriented Architecture (SOA) access.

Via the Connect Agents, clients can read and write data via SQL operations, execute SQL queries, execute transactions, and participate in two-phase distributed transactions. Attunity Connect provides several thin client agents that enable universal connectivity by providing a variety of interfaces to the data sources:

- The ODBC Driver provides 32-bit and 64-bit connectivity using ANSI SQL-92.
- The JDBC Driver makes data sources accessible by any Java or J2EE application.
- The OLE DB Provider provides data access for Windows applications.
- The ADO.NET Provider makes data sources accessible by applications written in Visual Basic, ASP.NET, and C#.
- The XML Data Services Adapter defines XML operations on the data source in support of SOA and Web services.
- The SDK Driver allows developers to develop custom agents for proprietary data structures.

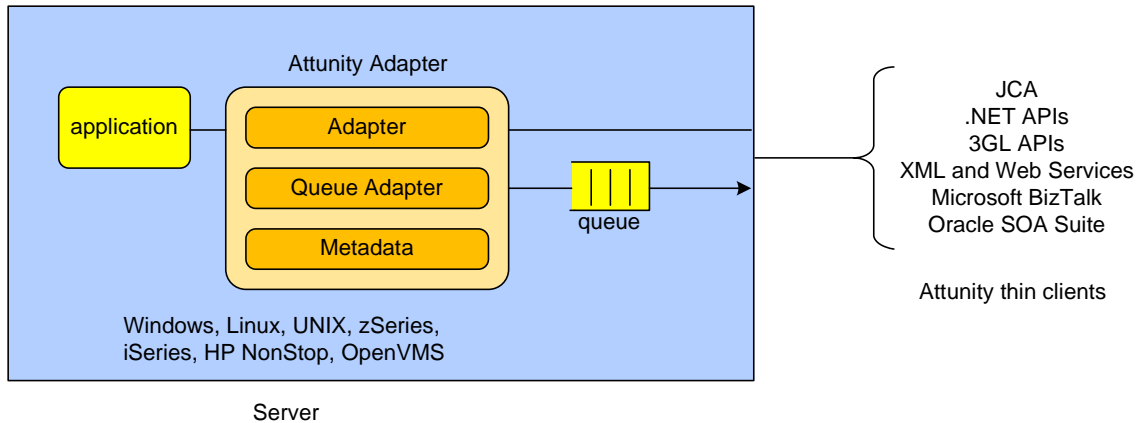
Application Connectivity

Attunity Connect application adapters provide integration with legacy applications running on platforms ranging from Windows to mainframes. The adapters expose existing applications as services that can be used by Java, .NET, and web applications.

Attunity Connect supports many application-integration platforms such as Oracle SOA Suite and Microsoft BizTalk. Local and global transaction interfaces are supported via the XA interface.

The Attunity adapters provide an XML interface to many application environments and queues. Metadata is imported from the application, and additional metadata can be defined for multiple interactions. The interfaces can be used by applications running in any computing environment.

Adapters are provided for CICS and CICS Queue, IMS/TM, COBOL, RPG, Tuxedo and Tuxedo Queue, and Pathway applications. For those environments in which queues are supported, queued events are converted to XML messages for delivery to enterprise applications.



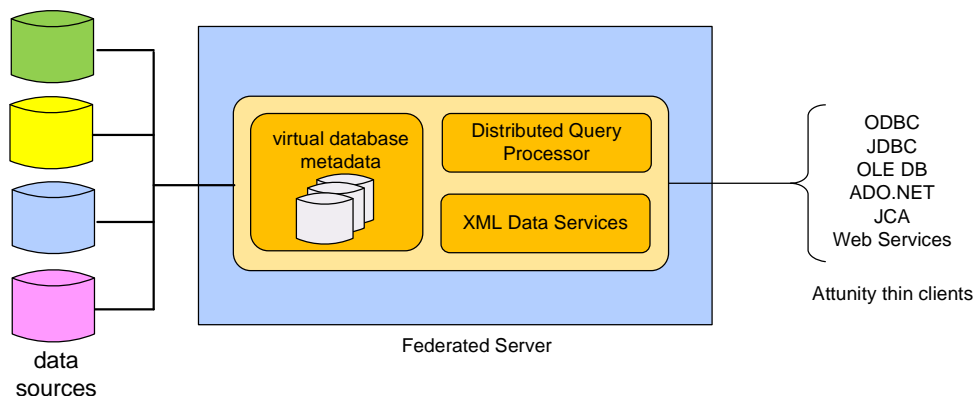
Attunity Connect – Application Connectivity

For each of these environments, the following interfaces are available:

- JCA
- .NET APIs
- 3GL APIs
- XML and Web Services (SOAP, WSDL)
- Microsoft BizTalk
- Oracle SOA Suite

Attunity Federate

Attunity Federate provides virtual data federation for enterprise information integration by allowing the definition of virtual views across data stored in disparate and geographically dispersed databases. Any set of databases supported by Attunity Connect can be consolidated into single views by Attunity Federate, where distributed queries can be executed as if working with a single consolidated database. This can be done across operational databases or for federating historical data from a data warehouse with real-time operational data from an OLTP system. With Federate, there is no need to move the data to a secondary data store so that it can be integrated.



Attunity Federate

Attunity Federate communicates with the Attunity Agents on each of the database servers to create a data network from which it can access data from any of the network's data sources. It

uses the metadata for each data source imported by Attunity Connect and allows the further definition of metadata for virtualized views across the data network.

Attunity's distributed query processor optimizes queries made on the virtualized federated database. It optimizes the execution of queries while delegating processing to the Attunity agents that reside on the database servers. Therefore, it is the agents on the database servers that do the heavy processing, thus distributing the query processing load across the data network.

The Metadata Repository contains the metadata definitions for the physical and virtual databases. It also contains the metadata for the XML data services.

The distributed query processor and the other Attunity Federate components are distributed across the database servers. There is no independent platform required for database federation.

The Attunity federated database can be accessed via the SQL and XML interfaces provided in the application platform from which the query originates. Data-processing applications can use ODBC, JDBC, ADO, or ADO.NET to access the virtual databases. Service-oriented applications can use XML with Web services (SOAP, WDSL), JCA, and .NET. 3GL APIs may also be used.

Attunity Stream

Attunity Stream moves database changes from one database to another. Typical uses are to build integrated data stores such as data warehouses and operational data stores, to feed such platforms as ETL (Extract, Transform, and Load), EAI (Enterprise Application Integration), and BI (Business Intelligence), or to keep different systems synchronized.² In addition, the target database can be used for reporting, for writing backups, or for batch processing, thus eliminating batch windows.

Stream is based on Attunity's Change Data Capture (CDC) technology.³

Change Data Capture

Change Data Capture captures changes made to enterprise data sources and makes these changes available for immediate or later processing by other facilities. It does so non-intrusively to the application updating the database – no changes need to be made to the application or the database.

Attunity provides CDC capabilities for many databases, both relational and nonrelational. For databases that maintain a log of changes, CDC follows the log, reading changes as they are entered into the change log. All SQL databases and some nonrelational databases maintain such logs. They include Oracle Redo Logs, DB2 Journals, NonStop SQL/MP and Enscribe Audit Trails, CICS logs for VSAM files, SQL Server Transaction Logs, and Adabas pLogs.

Special techniques are required for nonrelational databases that do not provide accessible logs. For instance, Attunity provides its own loggers for VSAM files that are updated by batch programs and for IMS/DB files. RMS updates are intercepted by a special logger that interacts with OpenVMS system services.

The CDC capture facility provides the change stream to some other facility that will process it. In the case of Attunity Stream, this facility is the CDC Router. Other users of CDC are ETL (Extract,

² Though some companies have opted to use Attunity Stream for data replication to a disaster-recovery site, Attunity points out that Stream is architected to support BI and similar applications. It is not ideal for DR because it is not optimized for transactional integrity.

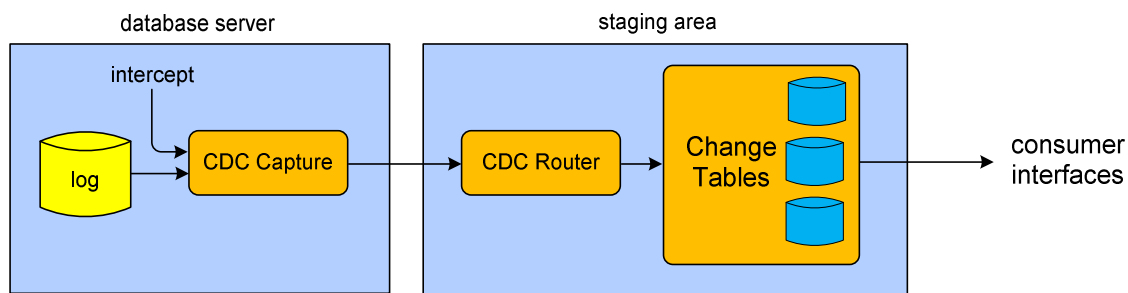
³ Efficient and Real Time Data Integration With Change Data Capture, *Attunity White Paper*, 2009.

Transform, and Load), EAI (Enterprise Application Integration), and BI (Business Intelligence) utilities.

CDC also can be used to capture and distribute events. With Event Data Capture (EDC), an entire transaction representing the event is packaged as an XML message including all relevant changes, even across several tables. Events are pushed via APIs such as JMS to event-driven applications.

Stream Processing

In Attunity Stream, changes are sent to the CDC Router in the Staging Area. The Staging Area provides a place to store changed data and to apply intelligent filters and services. It offloads the data servers and makes it easier to support multiple consumers of the streamed data changes. The Staging Area is typically resident on the same server as the tool being driven by Stream (ETL, EAI, BI, etc.) or on a dedicated server.



Attunity Streams

The CDC Router is the heavy lifter in Stream. It stores changes in Change Tables for use by different applications. Prior to storing data changes, it appends certain metadata information to each change, including the type of change (insert/updated/delete), the time of the change, and the application that made the change.

In addition, the CDC Router can be configured to apply certain processing rules to cleanse and transform the data. Changes can be filtered based on tables, the type of change, and the contents of certain fields.

From the Change Tables, applications can poll for changes (pull) using SQL (typically for ETL); or changes can be sent to the application (push) either immediately or according to a schedule via XML (typically for EAI). Pushing changes and events is done via a Publisher. For instance, the JMS Publisher publishes events to JMS queues, TIBCO queues, and MQ Server queues.

Attunity Stream supports several interfaces. For SQL, access may be accomplished via ODBC, JDBC, OLE DB, ADO, and ADO.NET clients. For XML, Stream supports .NET, JCA and 3GL clients.

In addition, Stream supports leading ETL and EAI tools, including:

- Microsoft SQL Server Integration Services (SSIS)
- SAP BusinessObjects Data Integrator/Services
- IBM InfoSphere DataStage
- Oracle Data Integrator
- Talend

- Cognos DecisionStream
- Informatica PowerCenter
- SAS Data Integration Server
- SyncSort DMX
- Open Text/Hummingbird Genio
- Embarcadero DT/Studio
- Microsoft BizTalk Server
- Oracle SOA Suite

Attunity Stream is fault-resilient. It maintains a pointer in the change stream to the last change that was successfully processed by each application. Stream positions are persistent. In the event of a failure, Stream can pick up and continue processing from the last change that was processed. In addition, change consumers can reposition in the change stream to facilitate recovery,

Attunity Studio

The Attunity Studio is a GUI-based facility for configuring and managing the components of the Attunity Integration Suite. Using an intuitive GUI interface, the Studio allows the developer to import existing metadata, define new metadata, configure agents and adapters, configure the CDC Router processing rules, deploy the Attunity components, and manage the production environment. It provides single-console access to the entire Attunity network throughout the enterprise.

Attunity

Attunity has been providing components of the Attunity Integration Suite for two decades. Its products are resold through the world's four largest software vendors – IBM, Oracle, Microsoft, and SAP/Business Objects. Through these vendors, other partners, distributors, and direct sales, Attunity products are installed at over 1,000 leading companies around the world.

Headquartered in Boston, Massachusetts, USA, Attunity provides global support via its offices in North America, Europe, the Middle East, and Asia.

Summary

The three products making up the Attunity Integration Suite – Connect, Federate, and Stream – provide the facilities necessary to integrate disparate applications and data stores across the enterprise in order to implement important enterprise-wide functions.

Attunity Connect and Attunity Federate make all enterprise data accessible to all enterprise applications either for individual data stores or as a unified virtual view of many data stores. Attunity Stream provides real-time or scheduled propagation of database changes for synchronizing databases, for feeding other utilities such as BI, ETL, and EAI facilities, and for providing timely updates to data warehouses and operational data stores.

