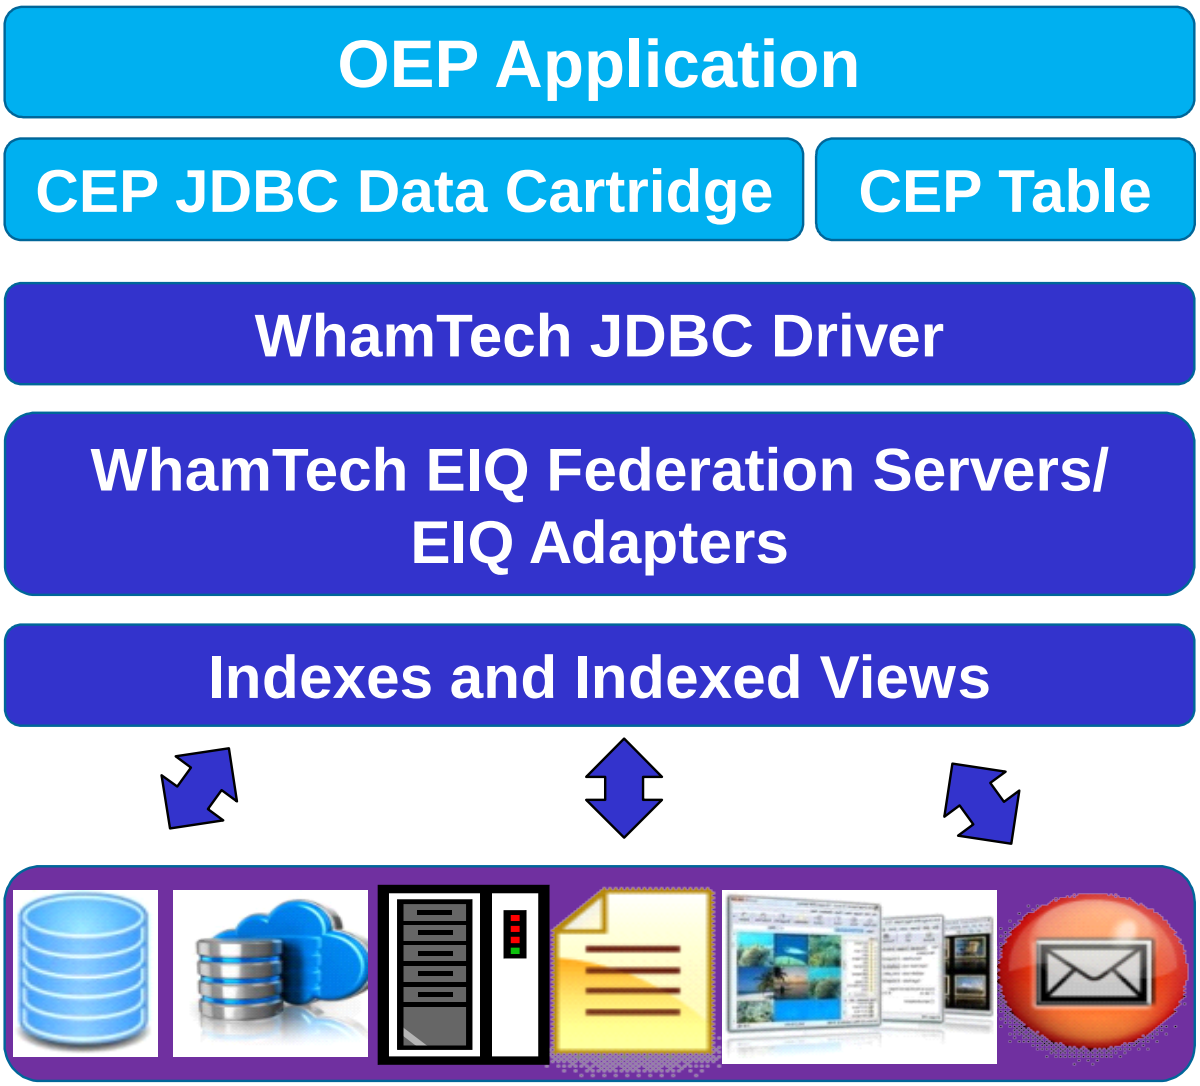

Oracle® Event Processing Using Indexed Virtualized and Federated Data Access from WhamTech™

Revision 1.8

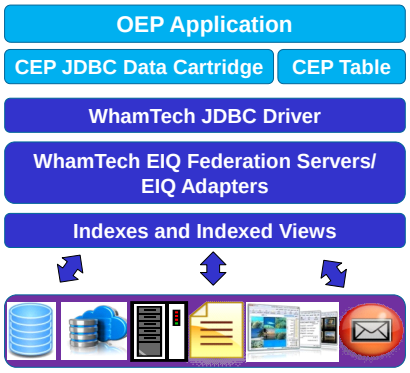

Contents

Executive Summary

Oracle Event Processing (OEP) partnered with WhamTech, Inc. ("WhamTech") to enable OEP application access to multiple disparate external data sources through index-based federated adapters and sub-middleware, called External Index and Query (EIQ) Adapters and EIQ Federation Servers, respectively. WhamTech enables OEP Query Language (CQL) calls to external data sources through Oracle CEP JDBC Data Cartridge and Oracle CEP Table functions that are executed in ANSI SQL on WhamTech federated adapter indexes. Queries yield result-set pointers that are then used to read results data from data sources.

Using WhamTech EIQ Adapters and EIQ Federation Servers, an OEP system can virtualize and federate data sources, and overcome almost any and all data and source system issues and limitations. OEP systems with WhamTech index-based federated adapters produce results that are clean, fast and impose almost no load on data sources. These adapters can be deployed in less time and cost than conventional approaches to data virtualization, federation and integration, and there are additional features, not commonly available with conventional approaches.

About Conventional Data Access and Integration Approaches

There are three conventional approaches to deal with data:

1. Copy it from the source to another structured repository, whether that is a cache-base, database/data warehouse, Big Data store, such as Hadoop, or cloud storage.
2. Leave it in the source and through conventional middleware and federated adapters, submit structured queries to, and obtain results from, the source.
3. Either of the two above approaches, but with associated unstructured text-oriented search with little to no adherence to the original structure of the source. Examples are search engines, content management and some Big Data stores, such as Hadoop.

Data copy raises issues of latency, ownership, responsibility, accountability, security, privacy, regulatory compliance, legal liability and additional costs. These are particularly important when there are schema changes and/or data is subject to cleansing, transformation and standardization through an extract, transform and load (ETL), or ELT, process. However, data copy approaches allow complete control over all aspects of data and query processing.

Leaving data in the source and using conventional middleware and federated adapters typically yields inefficient, load-imposing and low success queries with incomplete and/or inaccurate results. This is due to source system limitations and/or data format, type and quality issues. Typically, the process to establish adapters and optimize queries can be lengthy, complex and expensive, usually resulting in a series of queries in an attempt to

overcome these issues. However, the major advantage of federated adapter approaches is to allow data to remain in sources and, by definition, access the latest data available.

Unstructured text-oriented search does not allow, by definition, structured queries, such as SQL (and NoSQL), nor data cleansing, transformation and standardization, or semantic mapping to some form of data model. However, unstructured text-oriented search approaches allow all data to be findable and, of course, accommodate non-structured data.

About WhamTech Index-based Federated EIQ Adapters and Sub-middleware

WhamTech EIQ Adapters uniquely process queries against their own structured and unstructured indexes and then use results index pointers to read results data from sources. This approach combines the best of all three conventional data access and integration approaches previously discussed, and at the same time, overcomes the worst of each of the three approaches.

EIQ Adapters enable one of the main advantages of data copy through warehousing and associated ETL in that the data used to build and maintain indexes and for raw results, can be subject to a number of processes, including cleansing, transformation, standardization, aggregation, calculation, categorization, entity extraction, fuzzy matching and linking. Indexes, by default, do not store data, as results data comes from sources. There are also options for independent virtual and materialized views for pre-aggregation, pre-calculation, joins and master data. Indexes can be mapped to standard data models. And, at the same time, EIQ Adapters retain the main advantage of conventional middleware and federated adapters: Leaving data in the source.

EIQ Adapters enable another main advantage of data copy in that EIQ Adapters absorb almost 100% of the index and query load, and thus avoid overloading and slowing down source systems, such as used for transaction data. EIQ Adapters can be configured for near real-time alerts/notifications for event processing, decision support, business process management, dashboard updates and BI/analytics. And EIQ Adapter systems are linearly scalable through distributed parallel processing, segmented indexes, and independent and shared-nothing configurations.

EIQ Adapters can be accessed through standard drivers and accept ANSI SQL queries, with some variations of SQL (PL/SQL and other dialects) and other query languages through conversion (NoSQL, OQL and SPARQL).

EIQ Adapter structured indexes typically have the same schema as the data source and are columnar, i.e., a single index represents the values for one column in a data source table, with pointers that represent records/locations where these values are stored. EIQ Adapter unstructured indexes tend to be more global and can represent all values for many data

sources, e.g., text in multiple documents and types. EIQ Adapter structured queries and unstructured search can be combined in a single SQL statement.

Source data is at first discovered and profiled using raw (no processing) indexes on all or sample data. Data profiles are used to develop and test data cleansing, transformation and standardization ("data transforms"), and then applied to raw data as it is read from the source to build and maintain indexes. There are at least 12 different methods to update indexes, including in near real-time, depending on the data source and the business need. Figure 1 below illustrates a typical EIQ Adapter and sub-middleware components and process.

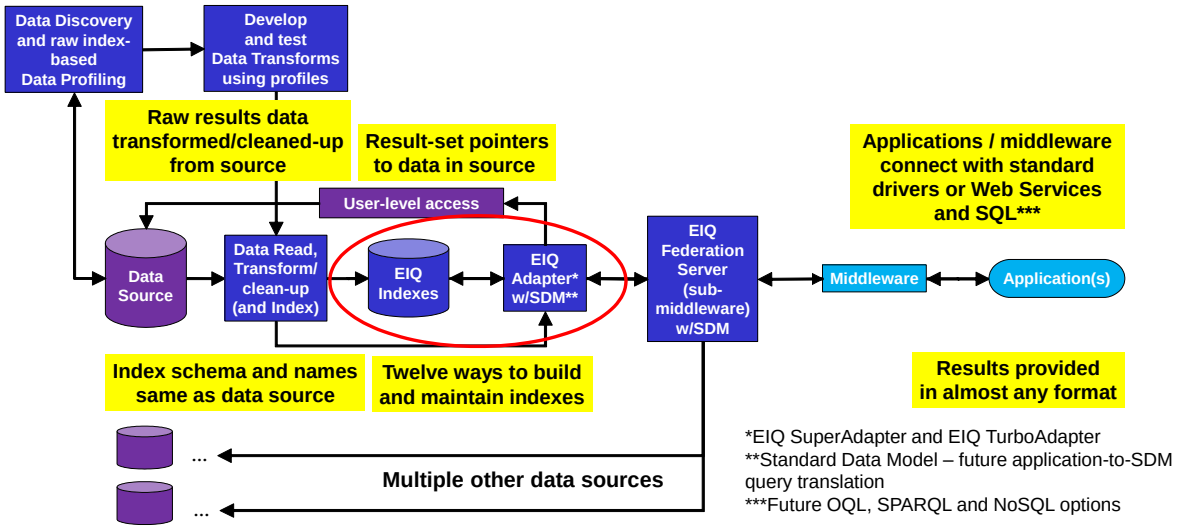

Figure 1: How EIQ Adapters and sub-middleware work.

Almost any and many data sources can be accessed through EIQ Adapters, covering the enterprise, Big Data, cloud and public, and includes mainframes, archived files, databases, file systems, Web pages, social media, office documents, files and Hadoop-based.

About Oracle Event Processing

Oracle Event Processing (OEP) provides a modular platform for building applications based on an event-driven architecture. At the heart of the OEP platform is the Continuous Query Language (CQL), which allows applications to filter, query and perform pattern matching operations on streams of data using a declarative, SQL-like language. Developers use CQL in conjunction with a lightweight Java programming model to write applications. Other platform modules include a feature-rich IDE, management console, clustering, distributed caching, event repository and monitoring, to name a few.

As event-driven architecture and complex event processing have become prominent features of the enterprise computing landscape, more and more enterprises have begun to build mission-critical applications using CEP technology. Today, mission-critical CEP applications can be found in many different industries. For example, CEP technology is being used in the

power industry to make utilities more efficient by allowing them to react instantaneously to changes in demand for electricity. CEP technology is being used in the credit card industry to detect potentially fraudulent transactions as they occur in real time. The list of mission-critical CEP applications continues to grow.

Use Cases for EIQ Adapters and Sub-middleware

There are many use cases where WhamTech EIQ Adapters and sub-middleware provide unique solutions that would otherwise be not possible or difficult to implement, as follows:

Query enabler – where data sources may…

- Not process SQL queries, e.g., archives or applications
- Need schema changes and associated indexes to process advanced SQL (or other QL) queries
- Need data transformation, entity extraction, advanced text search or other processing

Query enhancer – where data sources may…

- Need independent indexes, indexed views and queries to accelerate queries
- Need data cleansing and standardization to improve query success
- Be at capacity and cannot support additional external queries

Query federator – where data sources may…

- Need to be integrated with each other and existing systems without creating a data warehouse
- Not be copied
- Not have, but need real-time indexes and queries, e.g., streaming Fast Data and syslogs

Using OEP and EIQ Adapters and Sub-middleware in General

The above-mentioned use cases enable OEP and WhamTech EIQ Adapters to be highly complementary, as EIQ Adapters allow data to remain in sources and yet be fully accessible from OEP. There are two main external data source access options for OEP: Oracle CEP JDBC Data Cartridge and Oracle CEP Table. Access is enabled regardless of the data source location, driver/API, query language, query processing capabilities, data type, format or quality, or how

frequent updates occur. In fact, in simple cases, EIQ Adapters can support OEP and other applications with only configuration work.

CEP JDBC Data Cartridge allows for "ad hoc" queries from OEP. One CEP JDBC Data Cartridge accessing one EIQ Federation Server (sub-middleware) could access many EIQ Adapter-configured data sources and other EIQ Federation Servers.

CEP Table is populated when instantiated with a pre-determined query. Similarly to CEP JDBC Data Cartridge, one CEP Table accessing one EIQ Federation Server could access many EIQ Adapter-configured data sources and other EIQ Federation Servers.

Figure 2 below illustrates a simplified OEP/EIQ Adapter and sub-middleware system architecture.

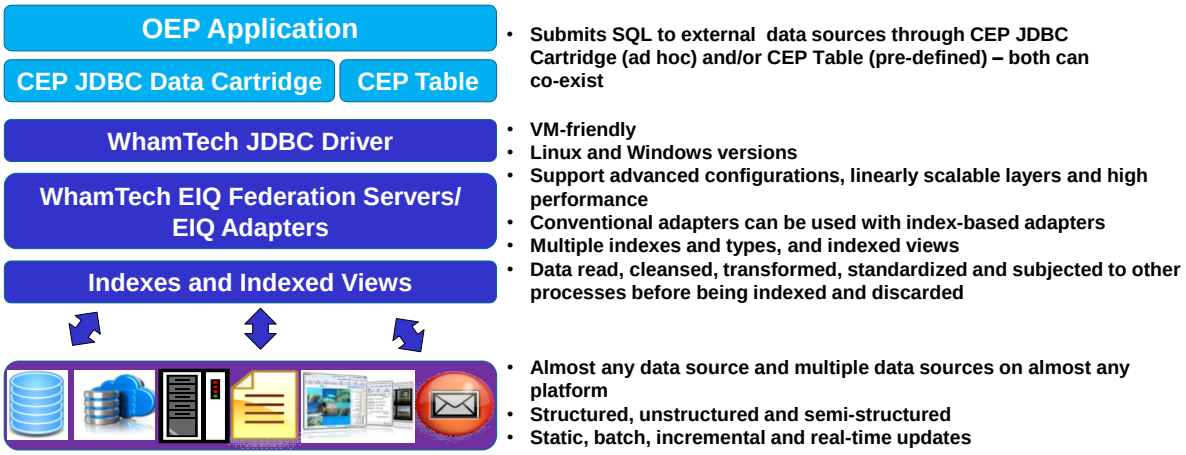

Figure 2: Simplified OEP/EIQ Adapter and sub-middleware system architecture.

<!– This is the configuration file used to register an EIQ Adapter and sub-middleware system as one external data source for use by an OEP application for both methods of Oracle CEP JDBC Data Cartridge and CEP Table -- >

```
<data-source>
  <name>myJdbcDataSource</name>
  <data-source-params>
     <global-transactions-protocol>None</global-transactions-protocol>
  </data-source-params>
  <connection-pool-params>
     <test-table-name>SQL SELECT 1 FROM Employee</test-table-name>
     <initial-capacity>5</initial-capacity>
     <max-capacity>10</max-capacity>
```

```
</connection-pool-params>
<driver-params>
    <url>jdbc:wham://127.0.0.1:1777/CEPORDERS3;SchemaInterface=native</url>
    <driver-name>com.wham.jdbc.Driver</driver-name>
    <properties>
        <element><name>user</name><value>admin</value></element>
        <element><name>password</name><value>admin</value></element>
    </properties>
    <use-xa-data-source-interface>false</use-xa-data-source-interface>
</driver-params>
</data-source>
```

Example Multiple Data Source, and EIQ Adapters and Sub-middleware System

Behind the two means that an OEP application can access and query external data sources, there can be a simple or more complex system consisting of multiple data sources, EIQ Adapters and sub-middleware

Figure 3 below illustrates how multiple types of adapters can be used to submit federated queries to, and obtain results from, multiple data sources, including index-based EIQ Adapters (EIQ SuperAdapters), non-indexed based EIQ Adapters (EIQ ConventionalAdapters), third-party adapters and sub-middleware, EIQ Federation Adapters. As can be seen, the capability for EIQ Federation Servers to access other EIQ Federation Servers allows WhamTech systems to scale. And because each EIQ Adapter and EIQ Federation Server is independently configurable, this avoids the requirement to have complex centralized command and control. In fact, independent configurations can lead to the situation Figure 3 illustrates, where the EIQ SuperAdapter for Hadoop would receive the same query from two separate EIQ Federation Servers; WhamTech has approaches to manage such situations. The multiple routes a query can take to a data source, for example, could provide benefits, such as redundant routing to a critical data source or for performance/load balancing, and designed into the configuration.

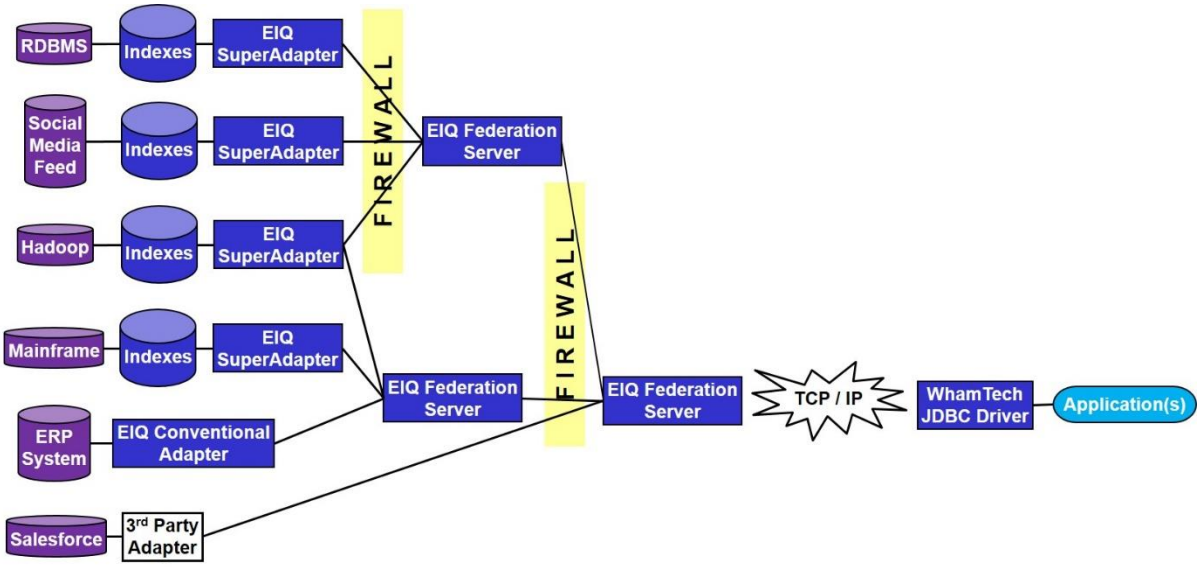

Figure 3: Example EIQ Adapters and sub-middleware system multiple data source configuration.

Other EIQ Adapter and Sub-middleware Configurations

Typically, at least one EIQ Adapter is required for each data source, but WhamTech allows for other configurations, including:

- Multiple EIQ Adapters to be used for one data source, each with the same indexes for load balancing, failover and backup
- Multiple EIQ Adapters to be used for one data source, each with different, segmented indexes for performance and scalability
- Distributed indexes across multiple storage locations for performance
- A bundled server for up to five data sources with up to five EIQ Adapters and one EIQ Federation Server
- In the almost unique case of Oracle, multiple database schemas can be treated as a single data source and can be accessed by one EIQ Adapter if they are all running on a single Oracle database server

Figure 4 below illustrates an example shared-nothing configuration that could be used as a basis for several of the configurations listed above.

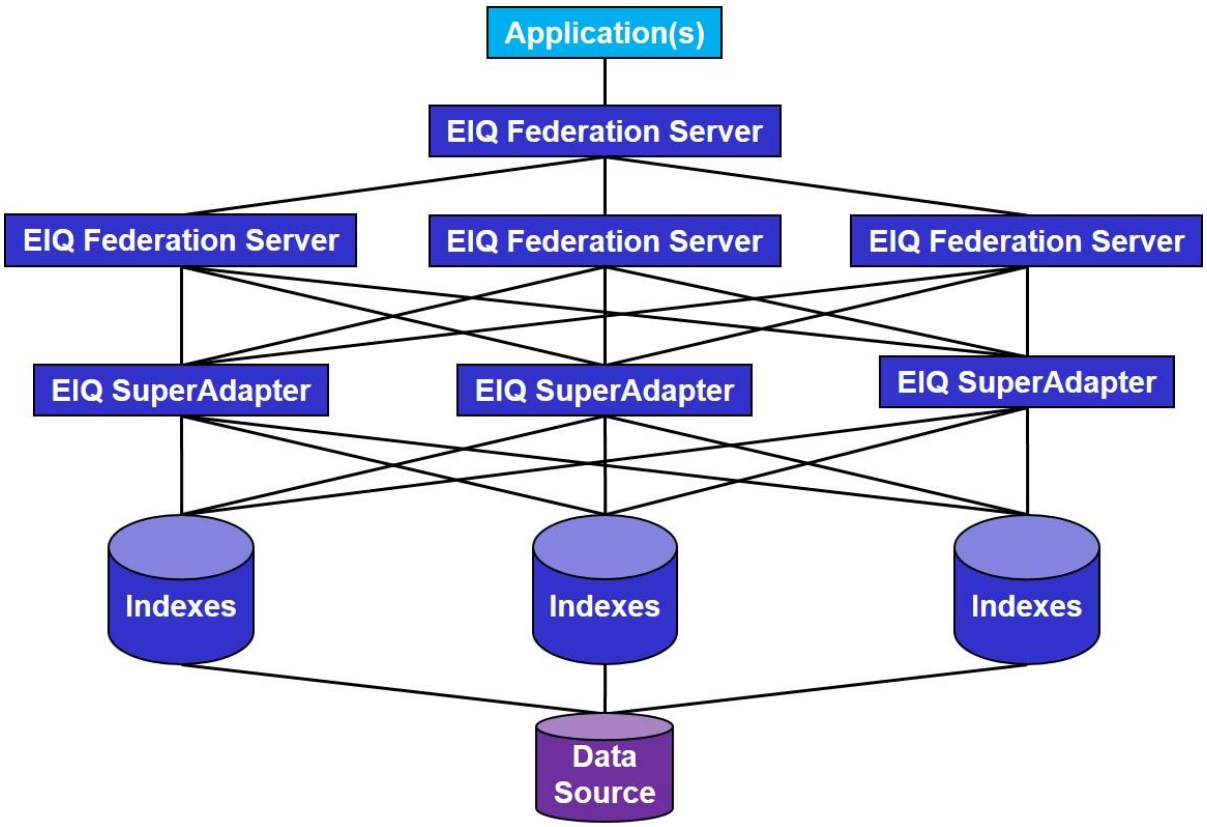

Figure 4: Example of a shared-nothing configuration showing multiple EIQ SuperAdapters, each with the same or segmented indexes that could be used for load balancing, failover and backup, or performance and scalability, respectively.

Using Oracle CEP JDBC Data Cartridge and EIQ Adapters and Sub-middleware

Oracle enables a seamless means to associate external data with the streaming data coming from Oracle CEP streams. The Oracle CEP JDBC Data Cartridge executes arbitrary (ad hoc) SQL queries against a database (or something that appears like a database) and uses the results in the CQL query. The process and parameters are describe as part of the Oracle Fusion Middleware documentation, available at: http://docs.oracle.com/cd/E17904_01/apirefs.1111/e12048/datacartjdbc.htm.

The use case in Figure 5 below illustrates an OEP application calling an EIQ Adapter system from a CEP process called “Proc” through an Oracle CEP JDBC Data Cartridge and submitting ad hoc SQL (in an Eclipse IDE).

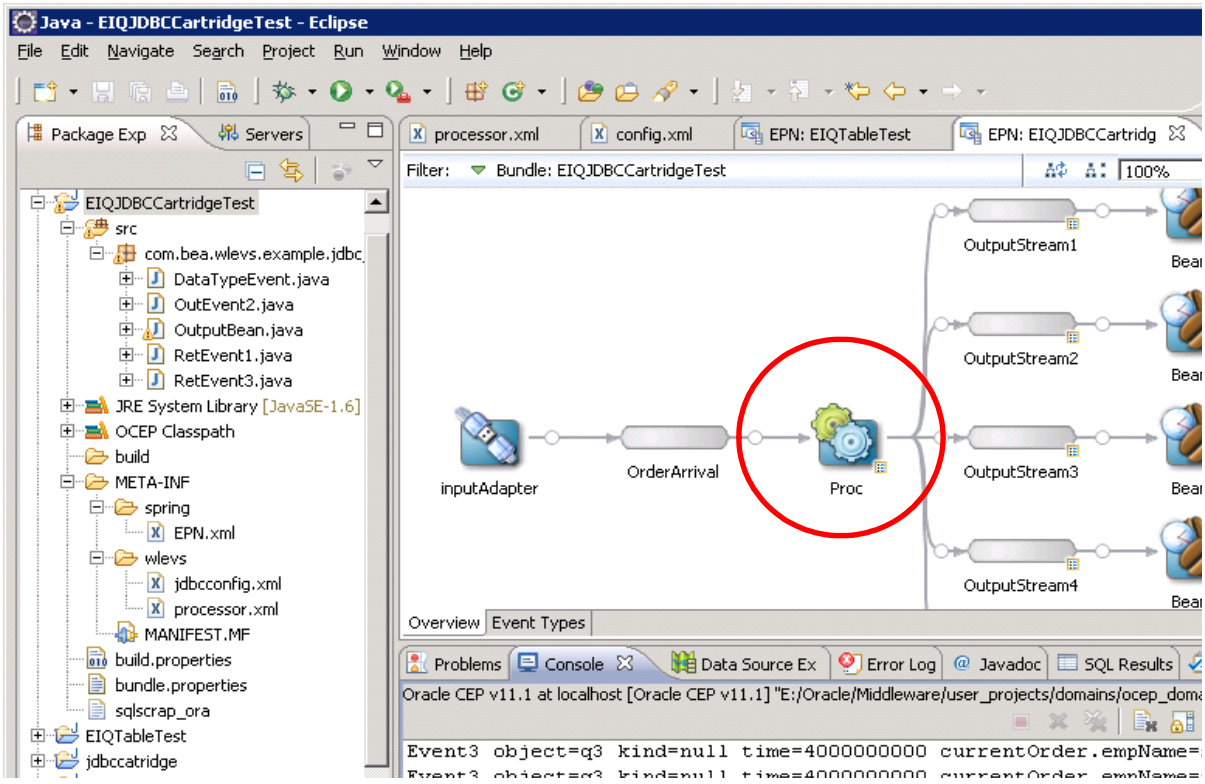

Figure 5: An OEP application calling an EIQ Adapter system from a CEP process called "Proc" through an Oracle CEP JDBC Data Cartridge.

<!– **This is the CEP process "Proc" using Oracle-defined standard CEP Query Language (CQL) to call the Oracle CEP JDBC Data Cartridge to submit an EIQ Adapter system SQL query** -- >

```
<processor>
  <name>Proc</name>
    <rules>
        <query id="q1"><![CDATA[
          RStream(
         select
          currentOrder.orderId,
          details.orderInfo.employeeName,
          details.orderInfo.concatEmpInfo() as empInfo
         from
          OrderArrival[now] as currentOrder,
```

```
        TABLE(getDetailsByOrderIdName@JdbcCartridgeOne(currentOrder.orderId,
        currentOrder.empName) as orderInfo) as details)
      ]]>
    </query>
```

<!– **This is the Oracle CEP JDBC Data Cartridge submitting an SQL query to the EIQ Adapter system** -- >

```
<jc:jdbc-ctx>
    <name>JdbcCartridgeOne</name>
    <data-source>myJdbcDataSource</data-source>
    <function name="getDetailsByOrderIdName">
        <param name="inpOrderId" type="int"/>
        <param name="inpName" type="char"/>
        <return-component-type>RetEvent1</return-component-type>
        <sql><![CDATA[
          SELECT
              Employee.empName as employeeName,
              Employee.empEmail as employeeEmail,
              OrderDetails.description as description
                  FROM
              PlacedOrders, OrderDetails , Employee
                  WHERE
            PlacedOrders.empId = Employee.empId AND
            PlacedOrders.orderId = OrderDetails.orderId AND
                      Employee.empName = :inpName AND
            PlacedOrders.orderId = :inpOrderId
        ]]> </sql>
    </function>
```

Using Oracle CEP Table and EIQ Adapters and Sub-middleware

The use case below illustrates an OEP application reading an EIQ Adapter system-prepopulated external CEP Table, called "DEPT", from a CEP process called "testOraProcessor" (in an Eclipse IDE). The CEP Table, "DEPT", is prepopulated by the EIQ Adapter system with a predefined SQL query when it is instantiated.

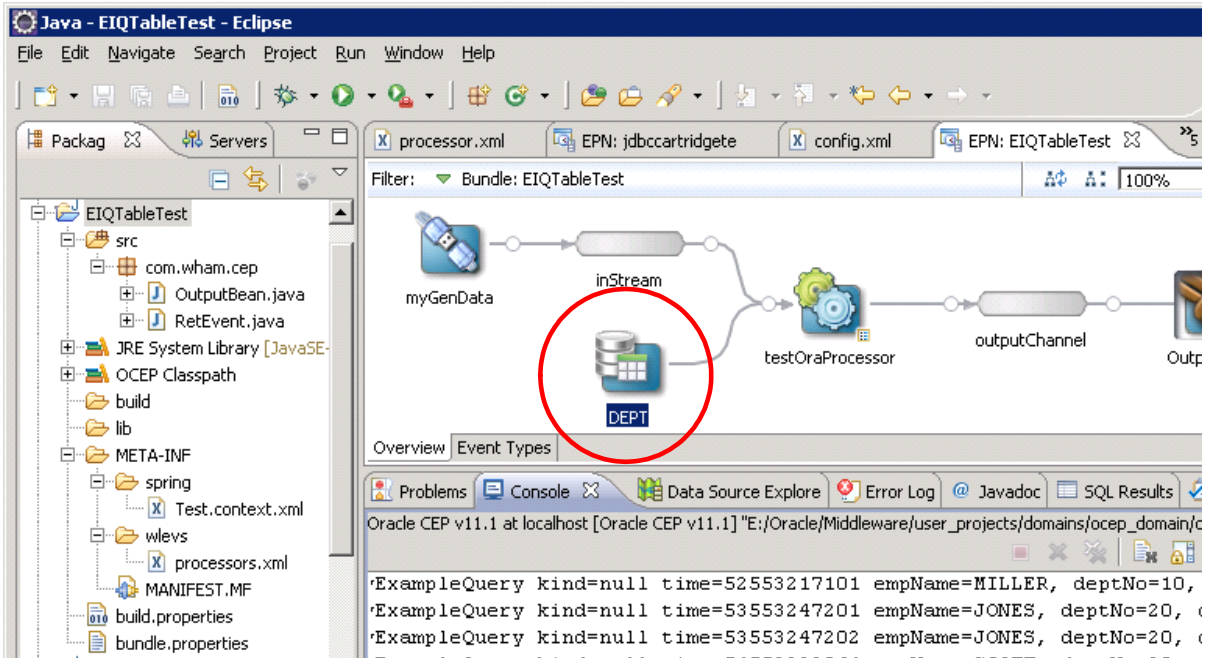

Figure 6: An OEP application reading an EIQ Adapter system-prepopulated external CEP Table, called "DEPT", from a CEP process called "testOraProcessor".

<!– This is the CQL to create the CEP Table, "DEPT", that in-turn created an SQL EIQ Table containing department information used as an event source -- >

```
<wlevs:table id="DEPT" event-type="deptInfo" data-source="ds_WhamTest" />
<wlevs:processor id="testOraProcessor">
   <wlevs:table-source ref="DEPT" />
</wlevs:processor>
```

<!– This the CQL used to JOIN the input stream with the CEP Table, DEPT, to obtain department details -- >

```
<processor>
   <name>testOraProcessor</name>
   <rules>
        <query id="myExampleQuery">
```

```
        <![CDATA[ select inStream.empName as empName , inStream.empDeptNo as
      deptNo , DEPT.DNAME as deptName, DEPT.LOC as deptLocation  from inStream
[now], DEPT where inStream.empDeptNo = DEPT.DEPTNO]]>
      </query>
  </rules>
</processor>
```

Top Ten Value Propositions for OEP Using EIQ Adapters and Sub-middleware

1. Plug-and-play in existing OEP architecture using two OEP external access options
2. Lower cost/TCO, shorter implementation time and better query metrics than alternatives
3. Avoid copying or moving data from sources
4. Support SQL (or other advanced query languages, including PL/SQL) on data sources that do not
5. Access the latest data available
6. Add near real-time indexes, indexed views and alerts/triggers (and functions and stored procedures in future)
7. Access, edge-process and integrate multiple data sources through a single standard data model/ontology
8. Bridge gaps among systems, e.g., Big Data, clouds, enterprise and government/public
9. Avoid high index or query loads on data sources
10. Can be used by other applications/solutions

Reasons Why WhamTech Is a Good Fit for Oracle in General

- Oracle has best-in-class event processing and other applications
- Oracle is defining standards for event processing and other solutions
- Oracle has strong support for third-party vendor-partners
- Oracle compatibility and experience in other areas, e.g., PL/SQL query emulation, Oracle Database Heterogeneous Services and Oracle Database Data Cartridges
- Oracle products well-established and complementary to WhamTech EIQ Products

Conclusions

Oracle and WhamTech established a unique combined capability to provide data warehouse quality OEP query access to data sources whose data remains in-place regardless of the data source location, driver/API, query language, query processing capabilities, data type, format or quality, or how frequent updates occur. WhamTech's unique approach to data virtualization and federation obviates the need for the conventional route of data warehousing with ETL or ELT and, of course, the commensurate time and cost of copying or moving data. Access is enabled through two methods: Oracle CEP JDBC Data Cartridge and CEP Table.

Best fit use cases are where data sources require queries (a) to be enabled where, e.g., SQL is not supported, (b) to be enhanced, e.g., cleansing and standardization, and/or (c) to be federated, e.g., sensitive data cannot be copied or moved. Flexible configurations and architectures accommodate the scale and complexity of data sources, and OEP application performance requirements. Two working examples test and illustrate Oracle CEP JDBC Data Cartridge and CEP Table used to access data sources through EIQ Adapters.

The main value propositions of OEP using EIQ Adapters and sub-middleware are that access to most data sources is through high performance independent index and query layers that cleanse, transform and standardize data for indexes and results, leave data in-place, and do not impose index or query loads. WhamTech has a history of working with Oracle database products and OEP is no exception. Oracle's leadership in defining standards and offering best-in-class tools and applications in event processing make Oracle the ideal partner for WhamTech.

About WhamTech

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 data virtualization, federation integration and interoperability technology software products. WhamTech develops these products to anticipate, meet and exceed the demands of customers seeking an alternative to the conventional approaches of data warehouses, federated data access with conventional adapters and enterprise search. Our goal is to provide a better and more seamless way to work with data, by changing the way data is accessed, analyzed, integrated, shared and interoperated. WhamTech adapter products are based on independent, cleansed indexes that execute both structured and unstructured queries for data warehouse capabilities, including results when data sources are unavailable and archive.

WhamTech customers are typically OEM sales to system integrators and platform companies, however, there have also been, and continue to be, a number of direct sales to end-customers.

Based on customer projects, WhamTech uses agile development methods with continuous integration to constantly improve EIQ Products.

WhamTech has entered into partnerships with system integrators and other companies to install, and provide support and training, for products.

Various patents and trademarks have been assigned to WhamTech.

Any questions on WhamTech products and capabilities, sales and investors can be obtained through whamtech.com.