Solutions *Atlantic*

Regulatory Reporting System
Technical White Paper

Solutions Atlantic, Inc.

Overview

This white paper describes the capabilities and architecture of Solutions Atlantic's Regulatory Reporting System (RRS). RRS is a modern software system that enables financial institutions to automatically monitor and identify regulatory disclosure requirements. Once identified, the system provides compliance users with the necessary workflow tools to research, modify, format and ultimately submit regulatory documentation.

RRS has been designed from the ground up to increase accuracy and reduce work. As such, from the user's perspective, RRS is chiefly a workflow system that precisely models the business process associated with regulatory disclosure and filing submission. In addition, users have the capability to browse and report on the extensive aggregated holdings database maintained within RRS.

The system architecture consists of a core platform that can be used to identify and manage any regulatory obligations, as well as add-on modules to support specific jurisdictions.

Features and Benefits

The Regulatory Reporting Systems provides financial institutions with the following features and benefits:

Features:

- Aggregates positions across multiple accounts, trading books, funds, and/or asset managers
- Merges multiple source feeds into single consolidated feed with full exception handling and status reports
- Real-time, daily, monthly, quarterly, and annual position monitoring
- Easily maintainable user-defined business rules/securities law operations
- Produces workflow tasks to research, modify and approve identified filing requirements
- Automatically generates required filing documents
- Sophisticated security master incorporating all asset classes including convertibles and derivative contracts
- Legal entity model supports complex financial institution structures including all accounts under management
- Leverages existing Microsoft technologies, including VBA, IIS, COM+ and SQL

Benefits:

- Automates the regulatory disclosure process, ensuring timely and accurate disclosure
- Eliminates unnecessary disclosure
- Frees up valuable compliance resources by reducing processing time, manual errors, and labor costs
- Codifies senior management's response to regulatory disclosure requirements
- Enables complex data consolidation, roll-up aggregation and reporting across separate business units
- Improves global data quality through comprehensive implementation

- Complete security master asset class model supports accurate convertible security aggregation and testing
- Adaptable to internal and/or securities law changes
- Provides historical audit trail of filings

System Capabilities

RRS includes an automated process for identifying required disclosures and a workflow process for reviewing, approving and submitting filings. These two processes are tightly integrated. When the automated process determines the need to file a report, a task is entered into the workflow system.

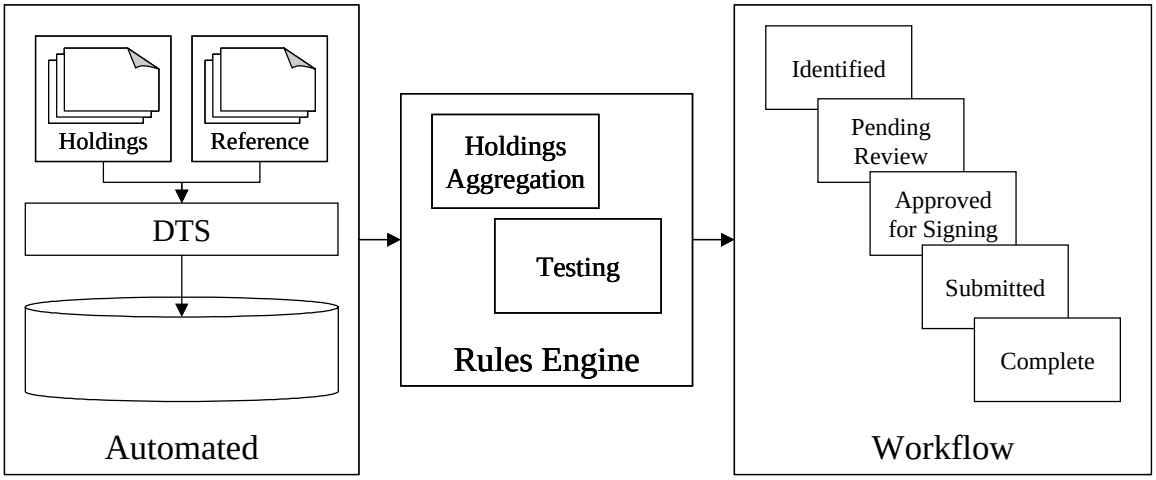

Figure 1 – System Processes

Automated Process

The automated process consists of three phases: data load, aggregation and testing. The following sections describe each phase:

Data Load

RRS gathers the data it needs from other sources within the institution. There are three types of data used by the system: holdings data, reference data and compliance data. Holdings data describes the positions in accounts held by various entities across the institution. This data is very volatile and needs to be refreshed frequently. Reference data describes securities, their issuers and the accounts holding the securities. This data is relatively stable; it is modified or enhanced much less often than holdings data. Compliance data is specific to the compliance function, for example, the address to which regulatory reports must be sent, and is not generally available from other sources within the enterprise. The compliance department usually maintains this data.

RRS can import data from a variety of sources. Typically, most data is delivered via database tables or flat files. Once loaded, the data is transformed into the required format and placed in the working tables of the RRS database.

RRS also has the capability to merge different sources of the same type. For example, for complex institutions with a variety of accounting systems, RRS can merge different holdings feeds into one consolidated feed for processing. Several exception and status reports are available to indicate the status of the merging process. Also, different feeds can be used for different aggregation and testing rules in different jurisdictions. This feature is supported by naming specific feeds and then associating the name with a rule. This is especially useful when holdings must be audited or a reporting period must close accommodating all corrections before the data can be used for evaluating exposure.

An additional feature is the ability to override data arriving during the feeds process. Users can override a specific element of data, and the override will be used for aggregation and testing from that point forward. When the override is introduced, the original value available on the feed is also stored with the override. Then, if the original value on the feed ever changes, the override is determined to be suspect, and the users have the opportunity to verify the new value. They can then use the new value, or override it yet again.

When new data appears on a feed, the system can assist users in evaluating the data for determining regulatory exposure. For example, if a new security listed in Toronto is loaded, the system will mark the issuer of the security as a "Reporting Issuer" in the Canadian provincial Ontario jurisdiction. (This is irrespective of whether or not the security previously was loaded and listed on other exchanges, or, for that matter, whether or not this issuer has other classes of listed or unlisted stock.) When a data is marked in this way, users have the ability to verify the system's decisions as part of the daily workflow process.

If a feed on a particular day was subsequently determined to be invalid for any reason, the system can be directed to reload the data from a new feed and re-aggregate positions. In the case when tasks were generated because of a previously run automated process, the system conservatively does not automatically regenerate the tasks. Instead, there is a standard process that administrative personnel follow to inform users of the potentially invalid tasks. At this point, users can either cancel the tasks, or alternatively, administrators can "back out" the tasks and re-run the process.

Aggregation

For the aggregation phase, the holdings data is aggregated and combined with reference data in order to build the aggregation results required for testing. These results are stored back into the database before being passed onto the testing phase. During aggregation, the system applies any special processing that may be required for multiply listed securities and convertible securities of all asset classes, including convertible bonds, options, rights, warrants and other derivatives, including one off contracts such as over-the-counter credit swaps. Aggregating holdings across convertibles and other related securities enables the system to correctly analyze exposure within an account and at any level up and across corporate entities. Aggregation is extremely fast, enabling the engine to be run either in a nightly batch or on-demand as positions are updated. Solutions Atlantic regularly performs stress testing on the aggregation engine with over 250,000 positions of all asset classes.

Testing

The testing phase applies regulatory rules to each aggregated position. Tests can be run for any arbitrary regulatory rule. For example, positions can be analyzed at the aggregate security, class of stock, issuer or even corporate entity levels.

Positive results cause the system to generate a task that is stored in a workflow queue. Intermediate results generated as a result of executing the rules are stored in the database. Finally, the system automatically generates all required documents. The supporting data and generated documents are available to the user when reviewing the task.

Workflow Process

The workflow process is implemented through a sophisticated graphical user interface consisting of the following: the RRS Home Page, task review forms, data maintenance screens and reports.

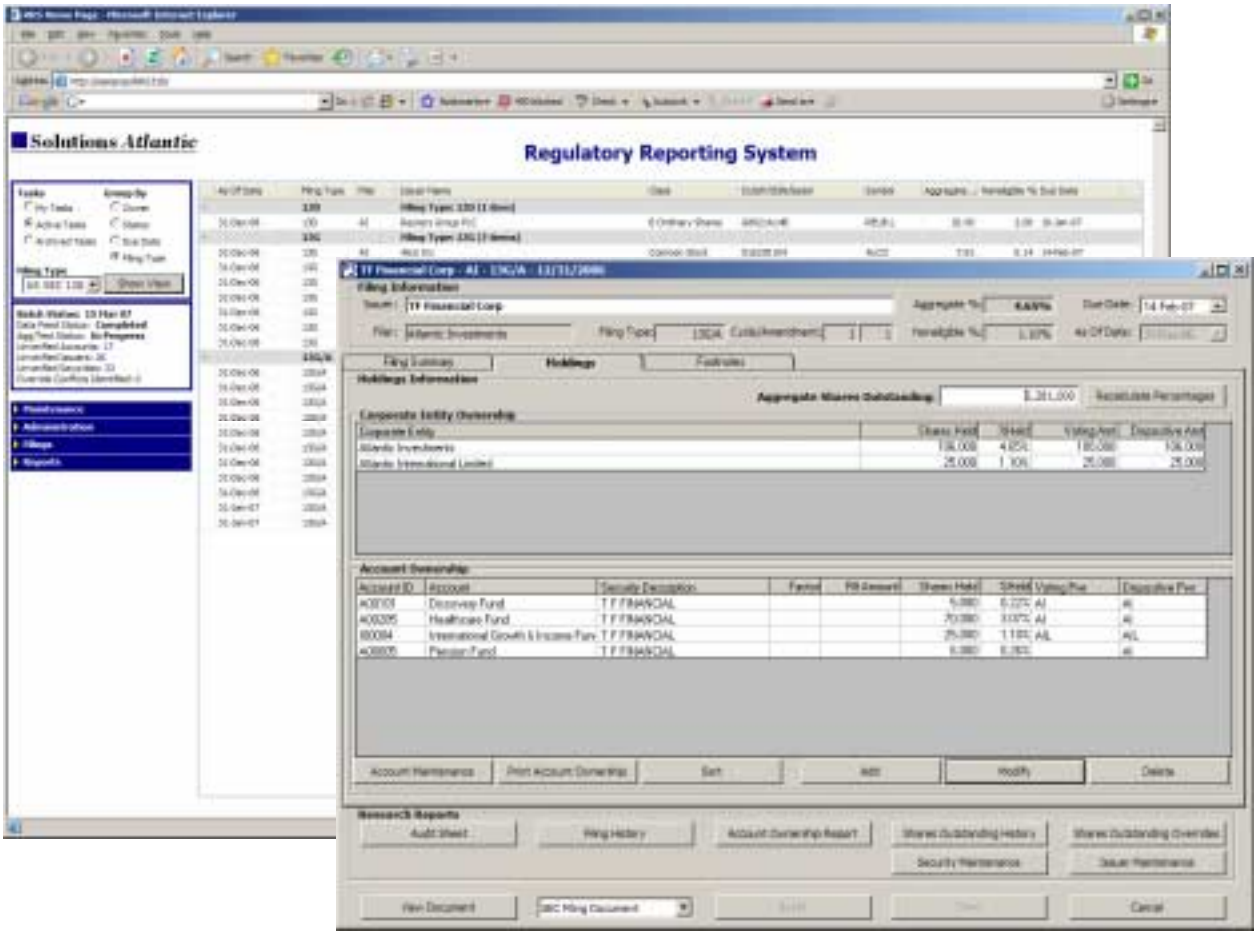

Figure 2 – Home Page and Task Form

The following sections describe each of these areas:

RRS Home Page

The RRS Home Page is the top-level user screen for the system. It provides facilities for task management and system navigation.

The Home Page lists all outstanding tasks in a clear, flexible tabular format, showing the user the current task status at a glance. Users can sort and filter the table according to their specific needs, organizing it by attributes such as filing type, status or assignee. Clicking on an entry in the table displays the corresponding task review form.

A menu on the Home Page gives the user rapid access to the RRS data maintenance screens and reports.

Task Review Forms

For each task, the system provides a form that allows the user to examine the results of the testing process and the documents associated with the task. If necessary, the user can update the holdings data, reference data or compliance data associated with the task. The system then automatically updates the aggregation data, the testing results and the generated documents. The form also allows the user to shepherd the task through the workflow process by updating its status or assigning it to another user. Comprehensive auditing captures all user actions and records them in the database.

In order to support the review process, the task review form provides links to relevant data maintenance screens and reports. These screens and reports are context sensitive. When invoked from the task form, they display only information that is relevant to the task. For example, with one click the user can display a report that shows the previous filing history just for the security in the task.

Data Maintenance Screens

RRS provides a numbers of screens that enable the users to augment or, when necessary, override the data that the system receives from the external data feed. Chief among these are the Security Maintenance screen for managing security master file data and the Account Maintenance screen for managing account master file data. The system also provides screens for managing system configuration data, such as user profiles and the holiday calendar.

Reports

The system provides a variety of reports to assist the user with the task review process. These reports can be run within the context of an individual task, as noted above, or can be used as general query tools with input parameters specified by the user. The subjects of the reports include historical data, such as filing history, investment information, such as beneficial ownership, or operational items such as the audit log.

Architecture

This section describes the RRS system architecture and the technologies used to implement it.

RRS is built using a modern n-Tier architecture on top of Windows DNA, a Microsoft recommended standard architecture for distributed applications. The system consists of a number of discrete software modules, each responsible for a specific function. These modules are then organized into tiers, which include the User Interface, the Database, and the Application Server. The Application Server is further broken down into the RRS Components, the Rules Engine, and the Documentation Generation module.

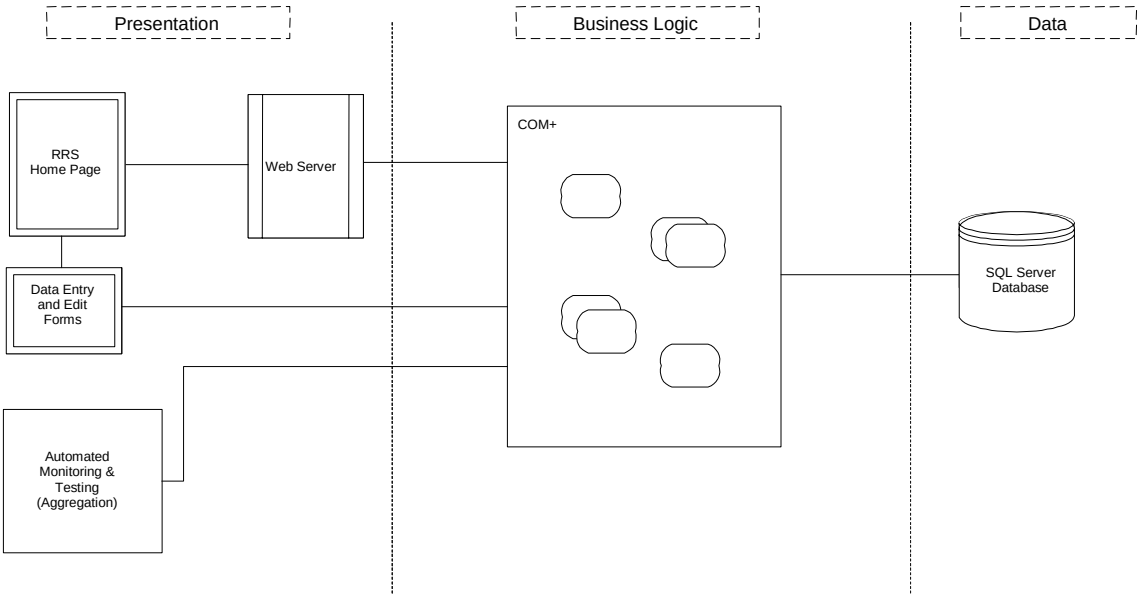

Figure 3 – 3-tier Architecture

The system uses Microsoft COM+ to enable the modules to locate and communicate with each other. This structure makes the system easily distributable across a variety of hardware configurations, including both back-end servers and front-end clients. For example, in a simple installation, the system can be configured so that all modules run on a single computer. Alternatively, for installations with more complex requirements, RRS can also be configured to be distributed across a database server, Web server, application server and several client desktops.

Solutions Atlantic's standard 3-tier development methodology enforces clean separation between presentation, logic and database software code. These logical tiers correspond to the User Interface, Application Server and Database in RRS. Because of this clean architecture, RRS is easily migrated to the latest stable Microsoft enterprise architectures. Solutions Atlantic plans to track and adopt Microsoft's .NET and related technologies according to our standard release schedule.

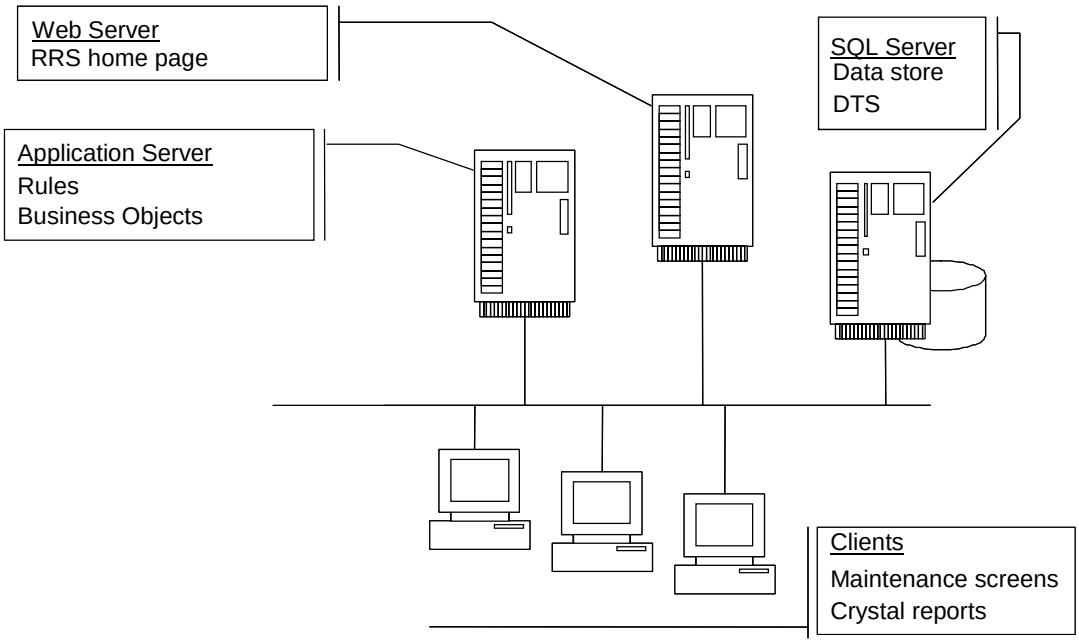

Figure 4 – Network Architecture

The 3-tier architecture also enables RRS to provide a sophisticated user authorization mechanism. Users are authenticated at multiple levels. In order to provide a seamless experience, RRS supports single sign on. Users do not have to log in twice as RRS obtains the user's credentials from the Windows operating system. Once a user is authenticated, they are then assigned application privileges based on their membership in a role, one of Manager, User or Viewer. Roles can be set up either through Windows ActiveDirectory (or other compatible directory service) or via the database internal to RRS.

In order to implement the automated process in customers' standard IT infrastructure, RRS can integrate with any batch scheduling system, including AutoSys, Tivoli, etc. Standard entry points are defined for the process, and simple scripts can be written to invoke the system.

The following sections describe the key software modules in the system.

Database

The database module is the internal working data store of RRS and is implemented using Microsoft SQL Server. It consists of two sub-modules, the schema and the loading engine.

Schema

The database schema is a relational model for storing holdings, accounts, corporate entities and security reference information. It also contains operational data, such as audit logs, user security, and submission history.

There are also intermediate results tables that are generated and used during the automated process and workflow process. These tables are optimized for use during task creation and submission. To support research and review, all results tables contain links back to the base data.

In addition, as RRS is a financial system, significant records are never deleted. Rather they are marked as "inactive", which supports the historical auditing necessary for compliance.

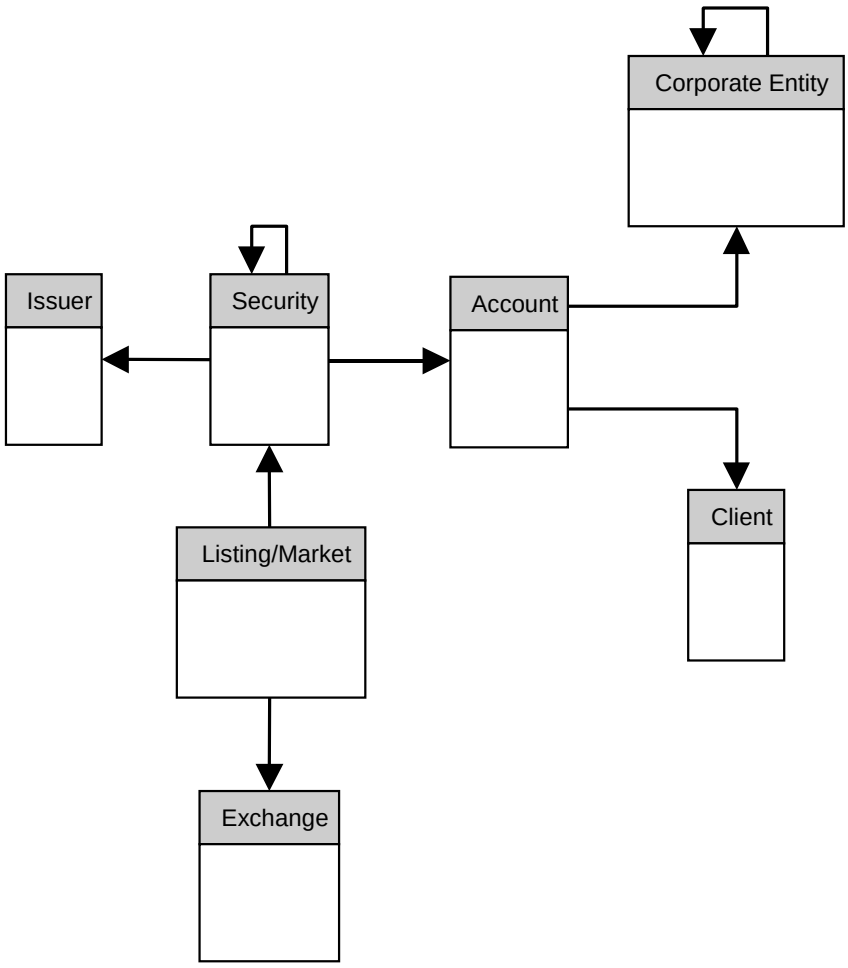

Figure 5 – Simplified Data Model

The high quality and sophistication of the RRS securities master is chiefly responsible for RRS's ability to report across different related asset classes. In order to model convertibility, optionality, exercise rights, etc., RRS abstracts these features into a generic model of relationships between securities. Different features of the relationship are captured, including factors, timing of the optionality, physical verses cash settlement and whether or not the "converted to" shares are issued at time of conversion. In this way, RRS can process any possible chain of conversion, including securities that convert into multiple securities, securities that convert to securities that convert to securities, and so on. In addition, RRS can handle over-the-counter derivatives, baskets, indices, and any other combination of instruments that are related to other securities. When exposure is evaluated, RRS takes into account any relevant features of the conversion, such

as whether or not a European-style option is exercisable within a timeframe specified by the regulations.

The RRS security master also contains a full security cross-reference capability, supporting the use of multiple identifiers at the issuer, the security and the market/listing levels. This enables the data load module to precisely identify each holding regardless of identifier used, resulting in accurate aggregation. For example, on different feeds, securities can be identified by CUSIP, SEDOL or an internal security ID, and the aggregation will be correct.

The hierarchy of corporate legal entities within an institution is also represented within the schema. In addition, contractual relationships between corporate entities and the accounts and trading books that hold positions are also modeled. This enables RRS to correctly aggregate positions and evaluate exposure across accounts, books and portfolios and up through corporate parents and holding companies. Entities can be internal or external to the financial institution, which enables money managers to communicate regulatory obligations to their clients. In addition, multiple relationships are modeled based on the nature of "roll up" from entity to entity. For example, investment discretion, voting authority and corporate ownership are all captured separately.

Support for the rules engine in the database consists of a flexible schema for describing a complete hierarchy of international regulations. Regulations are grouped by jurisdiction, including country, body of law, and testing rules. For example, for 13G filings, USA is the country, SEC is the body of law, and Rule 13(d) is the set of testing rules. The testing rules are then supported by a particular aggregation type.

The configuration for rules is stored in the database, and is read by the rules engine. The rules engine is responsible for evaluating the rules described in the database.

Adding regulations for a new country consists of adding the appropriate hierarchy to the database and describing the rules to the rules engine.

Loading Engine

To facilitate global aggregation, data from variety of sources must be gathered, integrated and loaded into the database. RRS uses Data Transformation Services (DTS) in Microsoft SQL Server to accomplish this task. DTS supports a great many types of databases and file formats, including SQL Server, Oracle, Sybase, Excel, Access, text files, etc. The loading process is encapsulated by storing the necessary logic in DTS packages. These packages are used to load data from multiple heterogeneous sources either interactively or automatically on a regularly scheduled basis.

In order to merge data that arrives from multiple sources, the working tables are populated in two steps. When the data feeds arrive, the first step loads the data into “staging” tables. The second step merges the data from the staging tables into consolidated tables for processing.

At any point in the process, exception and status reports can be generated to indicate results of the loading to both users and systems staff. Exception reports are generally specific to customer feed configurations and are created as part of the system implementation.

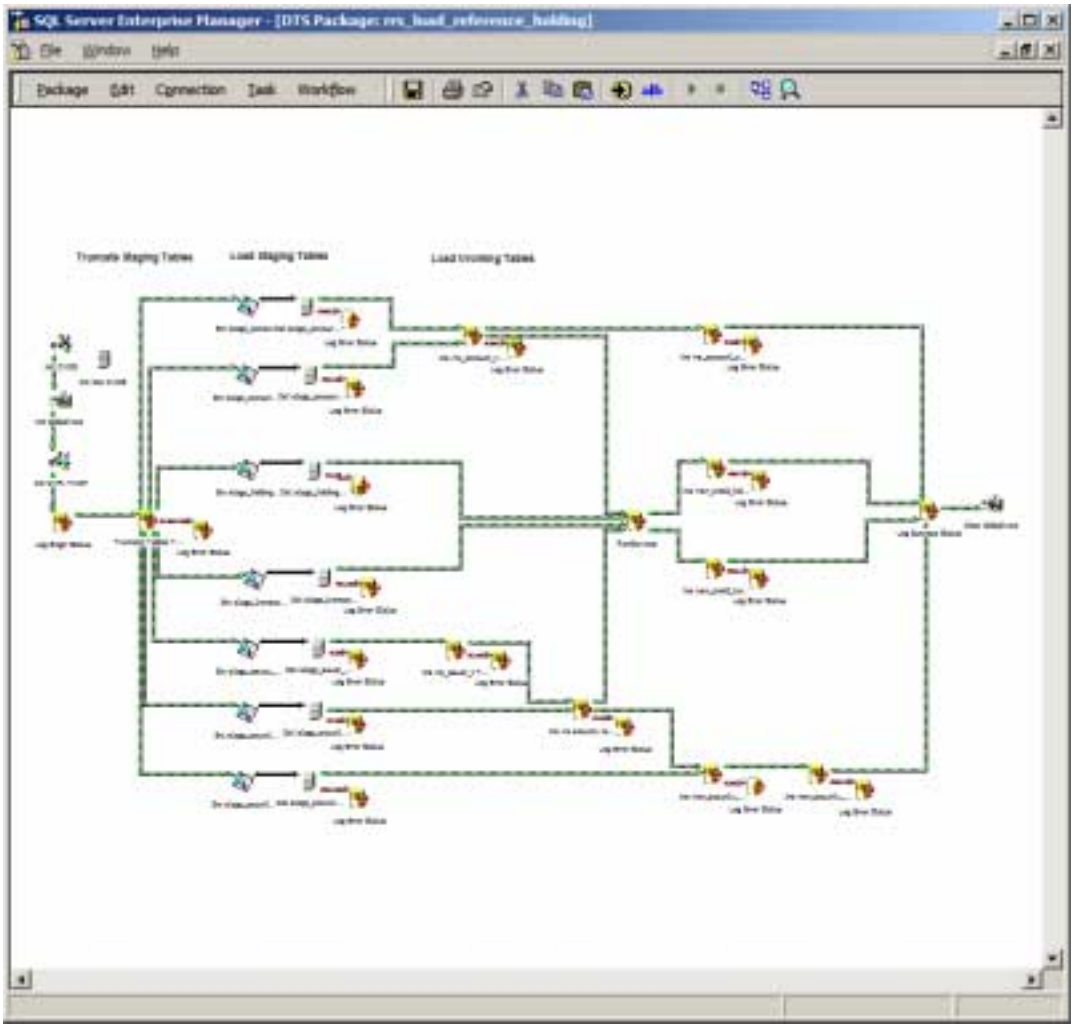

Figure 6 – Sample Feed

Rules Engine

The RRS rules engine is the mechanism for modeling business rules for regulatory reporting in a modular, flexible, extensible and easily maintainable fashion.

The rules engine is a high-performance multi-threaded driver for evaluating and processing data. It is used for aggregating holdings data, and also for testing the aggregated holdings against disclosure requirements. The rules engine is deployed in the Application Server tier.

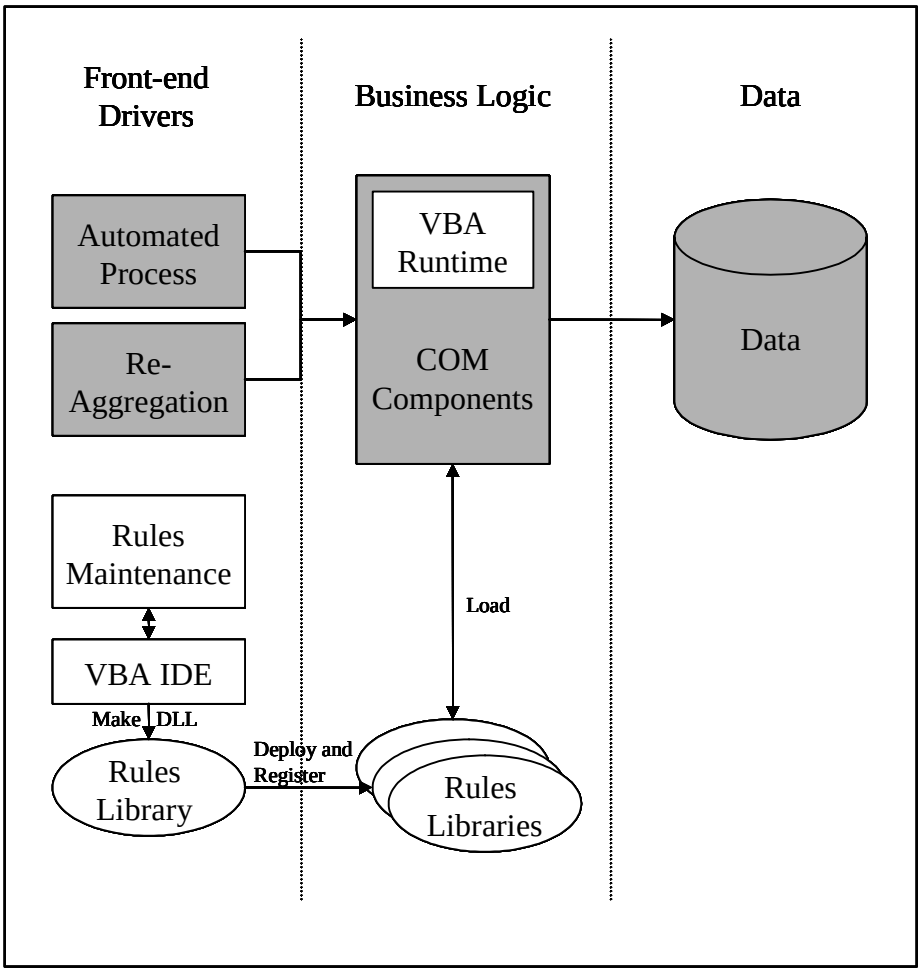

Figure 7 – Rules Engine Architecture

The business logic for aggregating holdings data and for testing the aggregated holdings against regulatory rules is expressed as a set of rules. By separating the processing engine from the rules definition, the system allows new rules to be added and existing rules to be modified without requiring a new release of the software.

The data model and internal organization of the regulatory rules are explicitly designed to support the flexibility and extensibility that are required in order to meet international filing requirements. The choice of business rules that govern a particular aggregation or testing process is entirely data-driven. The rules are organized in a country/jurisdiction/rule hierarchy.

Solutions Atlantic chose Microsoft Visual Basic for Applications (VBA) as the rules language for RRS for three reasons:

- VBA has the logical and computational operators required to express the logic found in regulatory rules.
- The VBA object model enabled Solutions Atlantic to construct objects that allow finance professionals to express rules using their natural vocabulary.

- VBA's use in Excel macros has made it a *de facto* standard in the financial industry.

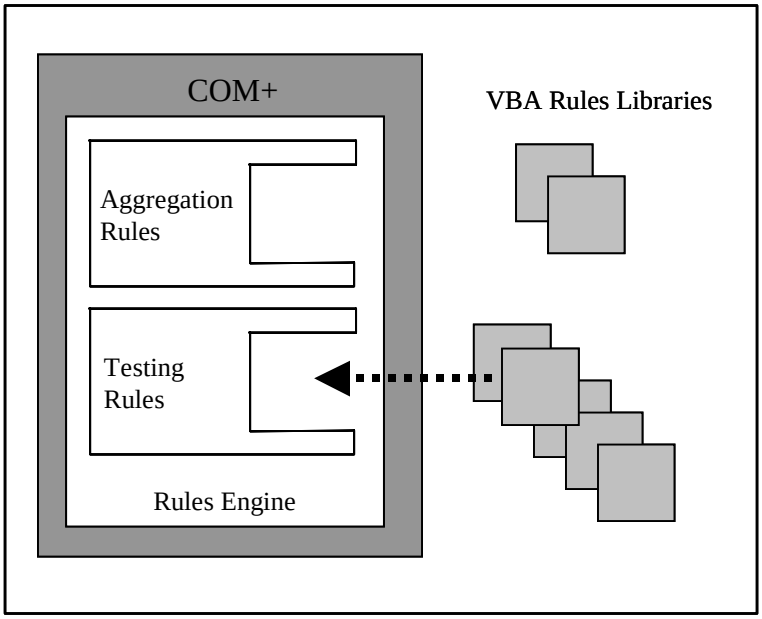

Figure 8 – Pluggable Rules Framework

When a testing rule indicates that a disclosure may be required, the Rules Engine creates a task in the workflow module. This task is supported in two additional ways. First, all of the supporting data for the task is captured and written to the database. Second, any required documents to be filed with the regulatory agency and any other associated documents, such as notices to the issuer of the security, are automatically generated and saved in electronic form.

As noted above, rules are easily maintained by legal and compliance analysts. When a rule is in need of modification, a standard process is followed to move the rule from development to quality assurance to staging and ultimately to production. A rigorous testing plan ensures that rules are correct interpretations of the regulatory language. Legal and compliance analysts with adequate credentials can review production rules at any time.

User Interface

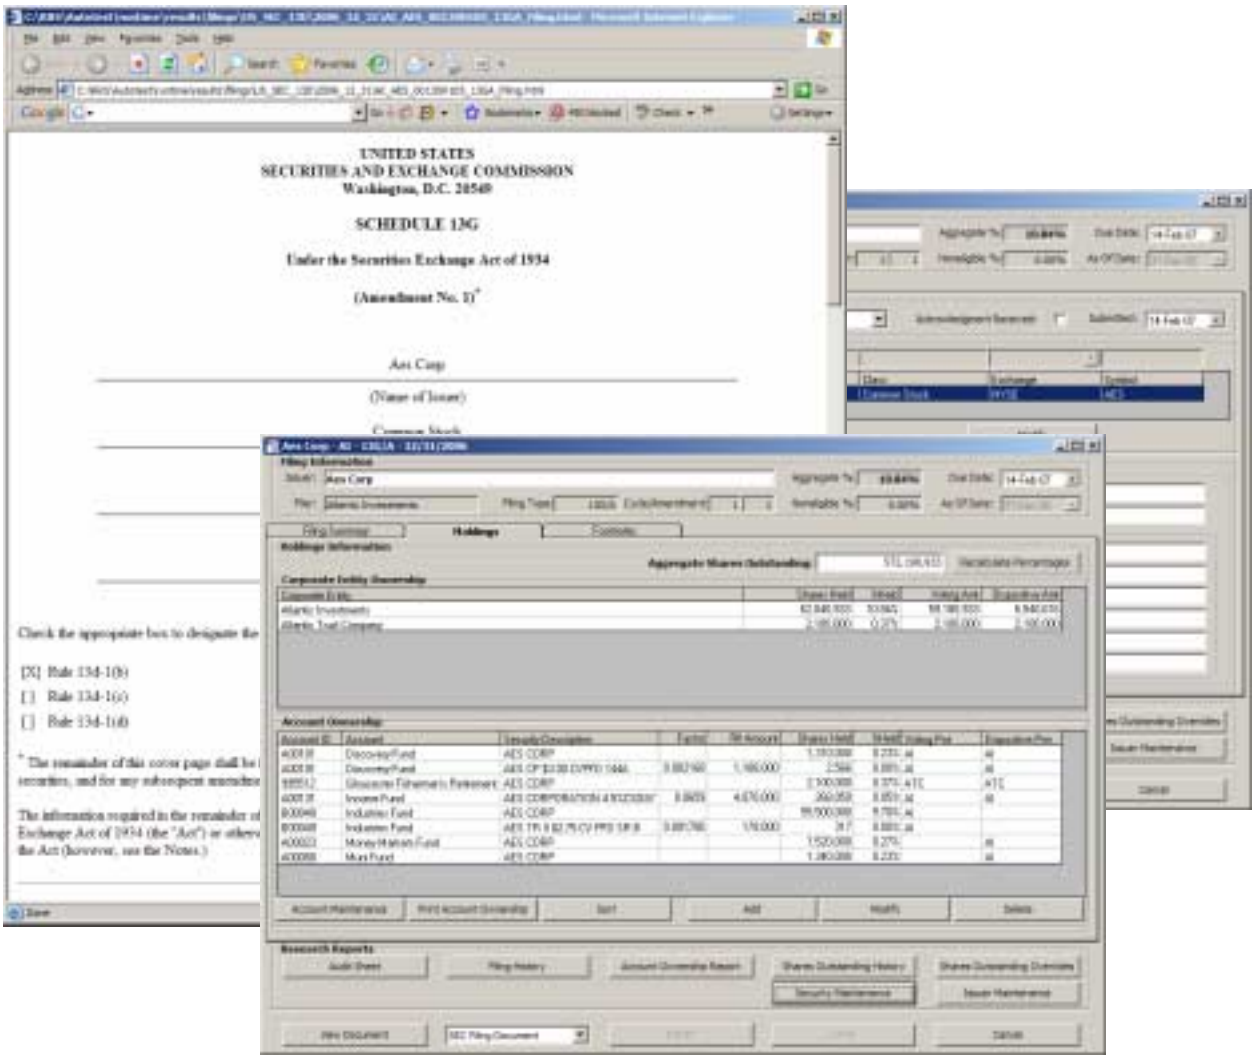

Figure 9 – User Screens

Home Page

The RRS Home Page is a web-server based home page that contains task items, summary information, and a menu to request data entry / edit functions. The Home Page is implemented using IIS Active Server Pages. Also, ActiveX components implementing the data entry/maintenance functions are called from the RRS home page.

User Screens

The database contains task items used to provide the basic workflow management data, such as task owner, task status and due date. The filing task forms and data maintenance screens are implemented using the standard Microsoft toolkits and ActiveX components, providing a full complement of familiar user interface objects.

Reports

The system uses Crystal Reports to implement the various reporting screens. RRS ships with a large number of canned reports as well as the capability for ad hoc queries created by the users, and delivered through a variety of mechanisms. Reports can be accessed from several locations in the system, including the Home Page and User Screens. Reports are also available within any web browser through the standard Crystal Reports mechanisms.

The following canned reports are supplied

- Mailing label printing
- Jurisdiction Securities Exposures
- Shares Outstanding Modifications/Research
- Account Holdings
- Consolidated Holdings
- Corporate Entity Holdings
- Issuer Holdings
- Security Holdings
- Quarterly Changes
- 13F Official List
- 13F De Minimis
- Overrides in Effect
- Account Beneficial Ownership

Document Generation

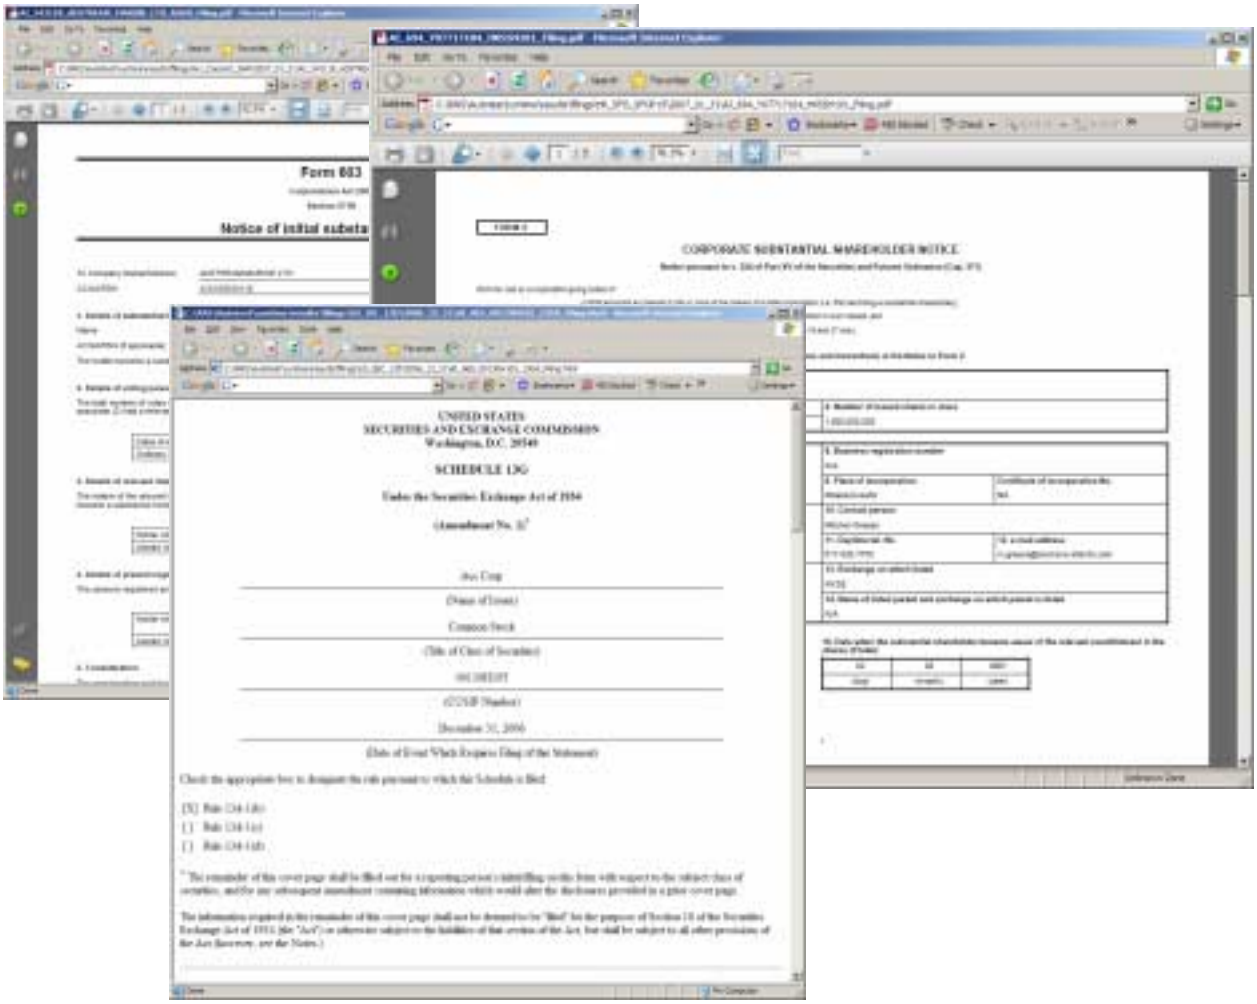

Figure 10 – Sample Documents

The documentation generation module of RRS can generate any required documentation to support the regulatory disclosure process. Documentation can include regulatory filing forms, letters to issuers, cover letters, etc. Documentation is created via templates, which can be customized according to the unique needs of every financial institution. By making use of standards-based XML document template technology, the system can populate the fields in the document, and even incorporate custom language provided by the institution.

Templates are represented as standard XSL documents, which are transformed into the necessary output in combination with an XML representation of the filing. The same XML data is used for all documentation, making the process of customizing and creating new templates easy and fast. The documentation generation module can generate ASCII text, HTML, PDF, RTF and many other file formats via XSL-FO. RRS is further configurable to support any off-the-shelf XSL-FO Processor to generate the finished output, including Apache FOP, Adobe, etc.

The system can also format and wrap the resulting documentation according to the specific needs of each regulatory body's electronic filing system. For example, all SEC filing submissions are correctly formatted in XFDL, eliminating the need for cumbersome, time-consuming and error prone use of the EDGAR filing software. Generated documents can be immediately submitted to EDGAR.

To support the workflow for document generation, document signatures are stored in the database. Information is captured about the power of attorney for each legal entity and associated filing types. If an individual with signature authority is ever unavailable, compliance analysts have the ability to temporarily supply an alternative signer with power of attorney.

Standard Jurisdictions

RRS comes complete with the US module, which includes SEC Rule 13(d) for beneficial ownership 13G and 13D filings and 13(f) for institutional investment manager disclosure Form 13F filings. Also available with the US module are reports supporting NASD Rule 2711 and NYSE Rule 472 research analyst 1% disclosure requirements. FDIC and OTS filings are also supported using the same aggregation rules. US filings are correctly formatted for direct submission to the EDGAR filing system.

Other jurisdiction modules currently available include Australia (Corporation Act, FATA (acquisitions and takeovers act)), Canada (Provincial Securities Acts and NI 62-103), Hong Kong, Japan, Korea, Singapore and the UK (POTAM Rule 8 and FSA DTR5). New jurisdiction modules will be added based on client needs. All international modules include monitoring and testing rules as well as filing document generation and submission formatting.

Additional jurisdiction modules are under development including Germany, Italy and France. New jurisdiction modules will be added based on client needs. All international modules include monitoring and testing rules as well as filing document generation and submission formatting.

Summary

Solutions Atlantic's Regulatory Reporting System enables financial institutions to meet their regulatory reporting and disclosure requirements. By providing automated monitoring and testing of holdings, RRS ensures that all filings are identified on time. The system's workflow capabilities assist compliance specialists to guarantee that submitted filings are correct.

The flexible rules engine gives the RRS platform two unique and powerful capabilities. Not only does it enable financial institutions to track changes and updates to existing rules, but also it is an extensible framework for adding rules, meeting the operational needs of global financial institutions.