

**INTEGRATED SYSTEM FOR RULE EDITING, SIMULATION,
VERSION CONTROL, AND BUSINESS PROCESS MANAGEMENT**

CLAIM OF PRIORITY

This application claims priority to U.S. Patent Application Serial No. 62/515,845, filed on June 6, 2017 and the content of the application is incorporated here by reference in its entirety.

BACKGROUND

[0001] The description relates to an integrated system for editing, simulation, version control, and business process management of rules.

[0002] In some examples, a conventional ruleset source code control tool allows a team of software programmers to manage the source code of several rulesets as the rulesets are developed and updated. The source code control tool maintains a central repository of the rulesets. A software programmer may check out a ruleset, use a rule editor to edit the ruleset, and store the ruleset locally on the programmer's computer. When the ruleset is checked out, the ruleset is locked and cannot be edited by another programmer. After the programmer finishes editing the ruleset, the programmer may use a rule simulator to simulate execution of the ruleset by applying the ruleset to test data. When the programmer is satisfied with the changes to the ruleset, the programmer checks in the ruleset.

SUMMARY

[0003] In one aspect, in general, the description features a method for managing one or more rules for use with an application program in a data processing system. The method includes: providing an integrated user interface to enable users to access an integrated application development system that enables version control and business process management of rulesets, in which each ruleset includes one or more rules, each rule including one or more rule cases that compute values for the rule's output depending

on the rule's input, and each rule case includes one or more input relationships in which a rule case is triggered by a given input record if one or more data values of the given input record meet the one or more input relationships; storing multiple versions of a first ruleset in a repository; storing at least one log data object having information about the multiple versions of the first ruleset in the repository; and receiving, through a first user interface provided by the integrated user interface, instructions from a first user associated with a first user identifier. The method includes in response to the instructions from the first user, checking out a particular version of the first ruleset from the repository, transitioning the particular version of the first ruleset from a first state of a plurality of states to a second state of the plurality of states, and checking in the particular version of the first ruleset back into the repository. The method includes storing in a log data object associated with the particular version of the first ruleset at least one of: (i) information specifying that a user associated with the first user identifier checked out the particular version of the first ruleset from the repository and the user associated with the first user identifier approved the transition of the particular version of the first ruleset from the first state to the second state, or (ii) information specifying that the checking out of the particular version of the first ruleset from the repository is associated with a first timestamp, and the transition of the particular version of the first ruleset from the first state to the second state is associated with a second timestamp.

[0004] Aspects can include one or more of the following features. The first state can indicate that the first ruleset is a draft that is still being edited, and the second state can indicate that the first ruleset is ready for review by a reviewer.

[0005] The first state can indicate that the first ruleset is ready for review by a user having authority to approve the first ruleset for release, and the second state can indicate that the first ruleset is finalized and ready to be released.

[0006] Each ruleset can correspond to at least one of a software module or a data processing entity that transforms input data to output data.

[0007] The integrated application development system can include a source code control module and a business process management module, the source code control module can process the checking in and checking out of the particular version of the first

ruleset, and the business process management module can process the transitioning of the particular version of the first ruleset from the first state to the second state.

[0008] The source code control module can process the storing in the log data object of at least one of (i) information specifying that the first user checked out the particular version of the first ruleset from the repository or (ii) information specifying that the checking out of the particular version of the first ruleset from the repository is associated with the first timestamp. The business process management module can process the storing in the log data object of at least one of (i) information specifying that the first user approved the transition of the particular version of the first ruleset from the first state to the second state or (ii) information specifying that the transition of the particular version of the first ruleset from the first state to the second state is associated with the second timestamp.

[0009] The method can include after checking out the particular version of the first ruleset from the repository and before checking in the particular version of the first ruleset back into the repository, locking the particular version of the first ruleset stored in the repository to prevent other users from modifying the particular version of the first ruleset in the repository.

[0010] The method can include receiving, through a second user interface of the integrated user interface, instructions from a second user associated with a second user identifier. The method can include, in response to the instructions from the second user, checking out the particular version of the first ruleset from the repository, and performing at least one of (i) associating an application program with the particular version of the first ruleset, and executing the application program to process input records and produce outputs according to one or more rules of the particular version of the first ruleset, or (ii) retrieving stored outputs produced from a prior execution of an application program that processed input records to produce the outputs according to one or more rules of the particular version of the first ruleset. The method can include checking in the particular version of the first ruleset back into the repository.

[0011] The method can include storing in the log data object associated with the particular version of the first ruleset at least one of: (i) information specifying that a user associated with the second user identifier checked out the particular version of the first

ruleset from the repository and the user associated with the second user identifier executed the application program using the particular version of the first ruleset to produce outputs, (ii) information specifying that a user associated with the second user identifier checked out the particular version of the first ruleset from the repository and the user associated with the second user identifier retrieved stored outputs produced from a prior execution of the application program using the particular version of the first ruleset, (iii) information specifying that the checking out of the particular version of the first ruleset from the repository is associated with a third timestamp, and the execution of the application program using the particular version of the first ruleset to produce outputs is associated with a fourth timestamp, or (iv) information specifying that the checking out of the particular version of the first ruleset from the repository is associated with a third timestamp, and the retrieval of stored outputs produced from a prior execution of the application program using the particular version of the first ruleset is associated with a fourth timestamp.

[0012] The integrated application development system can include a source code control module and a simulation module. The source code control module can process the storing in the log data object of at least one of: (i) information specifying that a user associated with the second user identifier checked out the particular version of the first ruleset from the repository, or (ii) information specifying that the checking out of the particular version of the first ruleset from the repository is associated with the first timestamp. The simulation module can process the storing in the log data object of at least one of: (i) information specifying that a user associated with the second user identifier executed the application program using the particular version of the first ruleset to produce outputs, (ii) information specifying that a user associated with the second user identifier retrieved stored outputs produced from a prior execution of the application program using the particular version of the first ruleset, (iii) information specifying that the execution of the application program using the particular version of the first ruleset to produce outputs is associated with the second timestamp, or (iv) information specifying that the retrieval of stored outputs produced from a prior execution of the application program using the particular version of the first ruleset is associated with the second timestamp.

[0013] The integrated application development system can include a source code control module and a simulation module. The source code control module can process the checking in and checking out of the particular version of the first ruleset, and the simulation module can process the execution of the application program to process input records and produce outputs according to rules of the particular version of the first ruleset.

[0014] The first user interface can be provided at a computer terminal of the first user and the second user interface can be provided at a computer terminal of the second user.

[0015] The method can include, in response to the instructions from the first user, editing the particular version of the first ruleset; storing the edited particular version of the first ruleset in the repository, and storing in the log data object associated with the particular version of the first ruleset at least one of: (i) information specifying that the first user edited the particular version of the first ruleset, or (ii) information specifying that the modification of the particular version of the first ruleset is associated with a third timestamp.

[0016] The integrated application development system can include a source code control module, a business process management module, and a rule editor. The source code control module can process the checking in and checking out of the particular version of the first ruleset, the business process management module can process the transitioning of the particular version of the first ruleset from the first state to the second state, and the rule editor can process the editing of the particular version of the first ruleset.

[0017] In another aspect, in general, a computer-readable medium storing a computer program for managing one or more rules for use with an application program in a data processing system is provided. The computer program including instructions for causing a computer system to: provide an integrated user interface to enable users to access an integrated application development system that enables version control and business process management of rulesets, in which each ruleset includes one or more rules, each rule including one or more rule cases that compute values for the rule's output depending on the rule's input, and each rule case includes one or more input relationships

in which a rule case is triggered by a given input record if one or more data values of the given input record meet the one or more input relationships; store multiple versions of a first ruleset in a repository; store at least one log data object having information about the multiple versions of the first ruleset in the repository; and receive, through a first user interface of the integrated user interface, instructions from a first user associated with a first user identifier. The computer program includes instructions for causing the computer system to: in response to the instructions from the first user, check out a particular version of the first ruleset from the repository, transition the particular version of the first ruleset from a first state of a plurality of states to a second state of the plurality of states, and check in the particular version of the first ruleset back into the repository. The computer program includes instructions for causing the computer system to: store in a log data object associated with the particular version of the first ruleset at least one of: (i) information specifying that a user associated with the first user identifier checked out the particular version of the first ruleset from the repository and the user associated with the first user identifier approved the transition of the particular version of the first ruleset from the first state to the second state, or (ii) information specifying that the checking out of the particular version of the first ruleset from the repository is associated with a first timestamp, and the transition of the particular version of the first ruleset from the first state to the second state is associated with a second timestamp.

[0018] In another aspect, in general, a system for managing one or more rules for use with an application program is provided. The system includes: a data storage system storing one or more rules; at least one output device or port for providing a graphical user interface; and a computer system coupled to the data storage system and the at least one output device or port. The computer system includes circuitry to: provide an integrated user interface to enable users to access an integrated application development system that enables version control and business process management of rulesets, in which each ruleset includes one or more rules, each rule including one or more rule cases that compute values for the rule's output depending on the rule's input, and each rule case includes one or more input relationships in which a rule case is triggered by a given input record if one or more data values of the given input record meet the one or more input relationships; store multiple versions of a first ruleset in a repository; store at least one

log data object having information about the multiple versions of the first ruleset in the repository; and receive, through a first user interface of the integrated user interface, instructions from a first user associated with a first user identifier. The computer system includes circuitry to: in response to the instructions from the first user, check out a particular version of the first ruleset from the repository, transition the particular version of the first ruleset from a first state of a plurality of states to a second state of the plurality of states, and check in the particular version of the first ruleset back into the repository. The computer system includes circuitry to: store in a log data object associated with the particular version of the first ruleset at least one of: (i) information specifying that a user associated with the first user identifier checked out the particular version of the first ruleset from the repository and the user associated with the first user identifier approved the transition of the particular version of the first ruleset from the first state to the second state, or (ii) information specifying that the checking out of the particular version of the first ruleset from the repository is associated with a first timestamp, and the transition of the particular version of the first ruleset from the first state to the second state is associated with a second timestamp.

[0019] In another aspect, in general, a system for managing one or more rules for use with an application program is provided. The system includes: means for providing an integrated user interface to enable users to access an integrated application development system that enables version control and business process management of rulesets, in which each ruleset includes one or more rules, each rule including one or more rule cases that compute values for the rule's output depending on the rule's input, and each rule case includes one or more input relationships in which a rule case is triggered by a given input record if one or more data values of the given input record meet the one or more input relationships; means for storing multiple versions of a first ruleset in a repository; means for storing at least one log data object having information about the multiple versions of the first ruleset in the repository; means for receiving, through a first user interface of the integrated user interface, instructions from a first user associated with a first user identifier; means for, in response to the instructions from the first user, checking out a particular version of the first ruleset from the repository; means for, in response to the instructions from the first user, transitioning the particular version

of the first ruleset from a first state of a plurality of states to a second state of the plurality of states; means for, in response to the instructions from the first user, checking in the particular version of the first ruleset back into the repository; and means for storing in a log data object associated with the particular version of the first ruleset at least one of: (i) information specifying that a user associated with the first user identifier checked out the particular version of the first ruleset from the repository and the user associated with the first user identifier approved the transition of the particular version of the first ruleset from the first state to the second state, or (ii) information specifying that the checking out of the particular version of the first ruleset from the repository is associated with a first timestamp, and the transition of the particular version of the first ruleset from the first state to the second state is associated with a second timestamp.

[0020] In another aspect, in general, a method for managing one or more rules for use with an application program in a data processing system is provided. The method includes: providing an integrated user interface to enable users to access an integrated application development system that enables editing, simulation, business process management, and version control of rulesets, in which each ruleset includes one or more rules, each rule including one or more rule cases that compute values for the rule's output depending on the rule's input, and each rule case includes one or more input relationships in which a rule case is triggered by a given input record if one or more data values of the given input record meet the one or more input relationships; storing multiple versions of a first ruleset in a repository; storing at least one log data object having information about the multiple versions of the first ruleset in the repository; and receiving, through a first user interface of the integrated user interface, instructions from a first user associated with a first user identifier. The method includes in response to the instructions from the first user, checking out a particular version of the first ruleset from the repository, editing the particular version of the first ruleset, associating an application program with the particular version of the first ruleset, and executing the application program to process input records and produce outputs according to rules of the particular version of the first ruleset, transitioning the particular version of the first ruleset from a first state of a plurality of states to a second state of the plurality of states, and checking in the particular version of the first ruleset back into the repository.

[0021] Aspects can include one or more of the following features. The method can include storing in a log data object associated with the particular version of the first ruleset at least one of: (i) information specifying that a user associated with the first user identifier checked out the particular version of the first ruleset from the repository, edited the particular version of the first ruleset, executed the application program using the particular version of the first ruleset to produce outputs, and approved the transition of the particular version of the first ruleset from the first state to the second state, or (ii) information specifying that the checking out of the particular version of the first ruleset from the repository is associated with a first timestamp, the editing of the particular version of the first ruleset is associated with a second timestamp, the execution of the application program using the particular version of the first ruleset is associated with a third timestamp, and the transition of the particular ruleset from the first state to the second state is associated with a fourth timestamp.

[0022] The integrated application development system can include a rule editor, a simulation module, a business process management module, and a source code control module. The rule editor can process the storing in the log data object of information specifying that the user associated with the first user identifier edited the particular version of the first ruleset, the simulation module can process the storing in the log data object of information specifying that the user associated with the first user identifier executed the application program using the edited particular version of the first ruleset, the business process management module can process the storing in the log data object of information specifying that the user associated with the first user identifier approved the transition of the particular version of the first ruleset from the first state to the second state, and the source code control module can process the storing in the log data object of information specifying that the user associated with the first user identifier checked out the particular version of the first ruleset from the repository.

[0023] The integrated application development system can include a rule editor, a simulation module, a business process management module, and a source code control module. The source code control module can process the storing in the log data object of information specifying that the checking out of the particular version of the first ruleset from the repository is associated with the first timestamp, the rule editor can process the

storing in the log data object of information specifying that the editing of the particular version of the first ruleset is associated with the second timestamp, the simulation module can process the storing in the log data object of information specifying that the execution of the application program using the particular version of the first ruleset is associated with the third timestamp, and the business process management module can process the storing in the log data object of information specifying that the transition of the particular ruleset from the first state to the second state is associated with the fourth timestamp.

[0024] The integrated application development system can include a rule editor, a simulation module, a business process management module, and a source code control module. The rule editor can process the editing of the particular version of the first ruleset, the simulation module can process the execution of the application program to process the input records and produce the outputs according to one or more rules of the particular version of the first ruleset, the business process management module can process the transitioning of the particular version of the first ruleset from the first state to the second state, and the source code control module can process the checking in and checking out of the particular version of the first ruleset.

[0025] In another aspect, in general, a computer-readable medium storing a computer program for managing one or more rules for use with an application program in a data processing system is provided. The computer program includes instructions for causing a computer system to: provide an integrated user interface to enable users to access an integrated application development system that enables editing, simulation, business process management, and version control of rulesets, in which each ruleset includes one or more rules, each rule including one or more rule cases that compute values for the rule's output depending on the rule's input, and each rule case includes one or more input relationships in which a rule case is triggered by a given input record if one or more data values of the given input record meet the one or more input relationships; store multiple versions of a first ruleset in a repository; store at least one log data object having information about the multiple versions of the first ruleset in the repository; and receive, through a first user interface of the integrated user interface, instructions from a first user associated with a first user identifier. The computer program includes instructions for causing the computer system to: in response to the instructions from the

first user, check out a particular version of the first ruleset from the repository, edit the particular version of the first ruleset, associate an application program with the particular version of the first ruleset, execute the application program to process input records and produce outputs according to rules of the particular version of the first ruleset, transition the particular version of the first ruleset from a first state of a plurality of states to a second state of the plurality of states, and check in the particular version of the first ruleset back into the repository.

[0026] In another aspect, in general, a system for managing one or more rules for use with an application program is provided. The system includes: a data storage system storing one or more rules; at least one output device or port for providing a graphical user interface; and a computer system coupled to the data storage system and the at least one output device or port. The computer system includes circuitry to: provide an integrated user interface to enable users to access an integrated application development system that enables editing, simulation, business process management, and version control of rulesets, in which each ruleset includes one or more rules, each rule including one or more rule cases that compute values for the rule's output depending on the rule's input, and each rule case includes one or more input relationships in which a rule case is triggered by a given input record if one or more data values of the given input record meet the one or more input relationships; store multiple versions of a first ruleset in a repository; store at least one log data object having information about the multiple versions of the first ruleset in the repository; and receive, through a first user interface of the integrated user interface, instructions from a first user associated with a first user identifier. The computer system includes circuitry to: in response to the instructions from the first user, check out a particular version of the first ruleset from the repository, edit the particular version of the first ruleset, associate an application program with the particular version of the first ruleset, and executing the application program to process input records and produce outputs according to rules of the particular version of the first ruleset, transition the particular version of the first ruleset from a first state of a plurality of states to a second state of the plurality of states, and check in the particular version of the first ruleset back into the repository.

[0027] In another aspect, in general, a system for managing one or more rules for use with an application program is provided. The system includes: means for providing an integrated user interface to enable users to access an integrated application development system that enables editing, simulation, business process management, and version control of rulesets, in which each ruleset includes one or more rules, each rule including one or more rule cases that compute values for the rule's output depending on the rule's input, and each rule case includes one or more input relationships in which a rule case is triggered by a given input record if one or more data values of the given input record meet the one or more input relationships; a repository for storing multiple versions of a first ruleset in a repository and at least one log data object having information about the multiple versions of the first ruleset; a first user interface of the integrated user interface, for receiving instructions from a first user associated with a first user identifier; means for, in response to the instructions from the first user, checking out a particular version of the first ruleset from the repository; means for, in response to the instructions from the first user, editing the particular version of the first ruleset; means for, in response to the instructions from the first user, associating an application program with the particular version of the first ruleset; means for, in response to the instructions from the first user, executing the application program to process input records and produce outputs according to rules of the particular version of the first ruleset; means for, in response to the instructions from the first user, transitioning the particular version of the first ruleset from a first state of a plurality of states to a second state of the plurality of states; and means for, in response to the instructions from the first user, checking in the particular version of the first ruleset back into the repository.

[0028] In another aspect, in general, a method for managing one or more rule containers for use with an application program in a data processing system is provided. The method includes: providing an integrated user interface to enable users to access an integrated application development system that enables version control and business process management of rule containers, in which each rule container including two or more rulesets, each ruleset includes one or more rules, each rule including one or more rule cases that compute values for the rule's output depending on the rule's input, and each rule case includes one or more input relationships in which a rule case is triggered

by a given input record if one or more data values of the given input record meet the one or more input relationships; storing multiple versions of a first rule container in a repository; storing at least one log data object having information about the multiple versions of the first rule container in the repository; and receiving, through a first user interface of the integrated user interface, instructions from a first user associated with a first user identifier. The method includes in response to the instructions from the first user, checking out a particular version of the first rule container from the repository, transitioning the particular version of the first rule container from a first state of a plurality of states to a second state of the plurality of states, and checking in the particular version of the first rule container back into the repository. The method includes storing in a log data object associated with the particular version of the first rule container at least one of: (i) information specifying that a user associated with the first user identifier checked out the particular version of the first rule container from the repository and the user associated with the first user identifier approved the transition of the particular version of the first rule container from the first state to the second state, or (ii) information specifying that the checking out of the particular version of the first rule container from the repository is associated with a first timestamp, and the transition of the particular version of the first rule container from the first state to the second state is associated with a second timestamp.

[0029] Aspects can include one or more of the following features. The first state can indicate that the first rule container is a draft that is still being edited, and the second state can indicate that the first rule container is ready for review by a reviewer.

[0030] The first state can indicate that the first rule container is ready for review by a user having authority to approve the first rule container for release, and the second state can indicate that the first rule container is finalized and ready to be released.

[0031] Each ruleset can correspond to at least one of a software module or a data processing entity that transforms input data to output data.

[0032] The particular version of the first rule container can include a pointer to each of the rulesets in the particular version of the first rule container.

[0033] The method can include, in response to the instructions from the first user, transitioning a ruleset in a particular version of a second rule container from an “open”

state to a “submitted” state, in which the “open” state indicates that the ruleset is a draft that is still being edited, and the “submitted” state indicates that the ruleset is ready for review.

[0034] The method can include transitioning the states of the particular version of the second rule container independently of transitioning the states of the ruleset in the particular version of the second rule container.

[0035] The method can include determining that the particular version of the second rule container is in the “open” state and all of the rulesets in the particular version of the second rule container are in the “submitted” state, and in response to the determination that the particular version of the second rule container is in the “open” state and all of the rulesets in the particular version of the second rule container are in the “submitted” state, automatically transitioning the particular version of the second rule container from the “open” state to the “submitted” state.

[0036] The method can include, in response to instructions from a second user, transitioning one of the rulesets in the particular version of the second rule container from the “approved” state to the “open” state, and automatically transitioning the particular version of the second rule container from the “approved” state to the “open” state in response to transitioning one of the rulesets in the particular version of the second rule container from the “approved” state to the “open” state.

[0037] The integrated application development system can include a source code control module and a business process management module, the source code control module can process the checking in and checking out of the particular version of the first rule container, and the business process management module can process the transitioning of the particular version of the first rule container from the first state to the second state.

[0038] The source code control module can process the storing in the log data object of at least one of (i) information specifying that the first user checked out the particular version of the first rule container from the repository or (ii) information specifying that the checking out of the particular version of the first rule container from the repository is associated with the first timestamp, and the business process management module can process the storing in the log data object of at least one of (i)

information specifying that the first user approved the transition of the particular version of the first rule container from the first state to the second state or (ii) information specifying that the transition of the particular version of the first rule container from the first state to the second state is associated with the second timestamp.

[0039] The method can include receiving, through a second user interface of the integrated user interface, instructions from a second user associated with a second user identifier. The method can include, in response to the instructions from the second user, checking out the particular version of the first rule container from the repository, and performing at least one of (i) associating an application program with the particular version of the first rule container, and executing the application program to process input records and produce outputs according to rules of the rulesets of the particular version of the first rule container, or (ii) retrieving stored outputs produced from a prior execution of an application program that processed input records to produce the outputs according to rules of the rulesets of the particular version of the first rule container. The method can include checking in the particular version of the first rule container back into the repository.

[0040] The method can include storing in the log data object associated with the particular version of the first rule container at least one of: (i) information specifying that a user associated with the second user identifier checked out the particular version of the first rule container from the repository and the user associated with the second user identifier executed the application program using the particular version of the first rule container to produce outputs, (ii) information specifying that a user associated with the second user identifier checked out the particular version of the first rule container from the repository and the user associated with the second user identifier retrieved stored outputs produced from a prior execution of the application program using the particular version of the first rule container, (iii) information specifying that the checking out of the particular version of the first rule container from the repository is associated with a third timestamp, and the execution of the application program using the particular version of the first rule container to produce outputs is associated with a fourth timestamp, or (iv) information specifying that the checking out of the particular version of the first rule container from the repository is associated with a third timestamp, and the retrieval of

stored outputs produced from a prior execution of the application program using the particular version of the first rule container is associated with a fourth timestamp.

[0041] The integrated application development system can include a source code control module and a simulation module. The source code control module can process the storing in the log data object of at least one of: (i) information specifying that a user associated with the second user identifier checked out the particular version of the first rule container from the repository, or (ii) information specifying that the checking out of the particular version of the first rule container from the repository is associated with the first timestamp. The simulation module processes the storing in the log data object of at least one of: (i) information specifying that a user associated with the second user identifier executed the application program using the particular version of the first rule container to produce outputs, (ii) information specifying that a user associated with the second user identifier retrieved stored outputs produced from a prior execution of the application program using the particular version of the first rule container, (iii) information specifying that the execution of the application program using the particular version of the first rule container to produce outputs is associated with the second timestamp, or (iv) information specifying that the retrieval of stored outputs produced from a prior execution of the application program using the particular version of the first rule container is associated with the second timestamp.

[0042] The integrated application development system can include a source code control module and a simulation module, the source code control module can process the checking in and checking out of the particular version of the first rule container, and the simulation module can process the execution of the application program to process input records and produce outputs according to the rules of the particular version of the first rule container.

[0043] The first user interface can be provided at a computer terminal of the first user and the second user interface can be provided at a computer terminal of the second user.

[0044] The method can include, in response to the instructions from the first user, editing the particular version of the first rule container; and storing the edited particular version of the first rule container in the repository. The method can include storing in the

log data object associated with the particular version of the first rule container at least one of: (i) information specifying that the first user edited the particular version of the first rule container, or (ii) information specifying that the modification of the particular version of the first rule container is associated with a third timestamp.

[0045] The integrated application development system can include a source code control module, a business process management module, and a rule editor. The source code control module can process the checking in and checking out of the particular version of the first rule container, the business process management module can process the transitioning of the particular version of the first rule container from the first state to the second state, and the rule editor can process the editing of the particular version of the first rule container.

[0046] In another aspect, in general, a computer-readable medium storing a computer program for managing one or more rule containers for use with an application program in a data processing system is provided. The computer program includes instructions for causing a computer system to: provide an integrated user interface to enable users to access an integrated application development system that enables version control and business process management of rule containers, in which each rule container including two or more rulesets, each ruleset includes one or more rules, each rule including one or more rule cases that compute values for the rule's output depending on the rule's input, and each rule case includes one or more input relationships in which a rule case is triggered by a given input record if one or more data values of the given input record meet the one or more input relationships; store multiple versions of a first rule container in a repository; store at least one log data object having information about the multiple versions of the first rule container in the repository; and receive, through a first user interface of the integrated user interface, instructions from a first user associated with a first user identifier. The computer program includes instructions for causing the computer system to: in response to the instructions from the first user, check out a particular version of the first rule container from the repository, transition the particular version of the first rule container from a first state of a plurality of states to a second state of the plurality of states, and check in the particular version of the first rule container back into the repository. The computer program includes instructions for causing the

computer system to: store in a log data object associated with the particular version of the first rule container at least one of: (i) information specifying that a user associated with the first user identifier checked out the particular version of the first rule container from the repository and the user associated with the first user identifier approved the transition of the particular version of the first rule container from the first state to the second state, or (ii) information specifying that the checking out of the particular version of the first rule container from the repository is associated with a first timestamp, and the transition of the particular version of the first rule container from the first state to the second state is associated with a second timestamp.

[0047] In another aspect, in general, a system for managing one or more rule containers for use with an application program is provided. The system includes: a data storage system storing one or more rule containers; at least one output device or port for providing a graphical user interface; and a computer system coupled to the data storage system and the at least one output device or port. The computer system includes circuitry to: provide an integrated user interface to enable users to access an integrated application development system that enables version control and business process management of rule containers, in which each rule container including two or more rulesets, each ruleset includes one or more rules, each rule including one or more rule cases that compute values for the rule's output depending on the rule's input, and each rule case includes one or more input relationships in which a rule case is triggered by a given input record if one or more data values of the given input record meet the one or more input relationships; store multiple versions of a first rule container in a repository; store at least one log data object having information about the multiple versions of the first rule container in the repository; and receive, through a first user interface of the integrated user interface, instructions from a first user associated with a first user identifier. The computer system includes circuitry to: in response to the instructions from the first user, check out a particular version of the first rule container from the repository, transition the particular version of the first rule container from a first state of a plurality of states to a second state of the plurality of states, and check in the particular version of the first rule container back into the repository. The computer system includes circuitry to: store in a log data object associated with the particular version of the first rule container at least one of: (i)

information specifying that a user associated with the first user identifier checked out the particular version of the first rule container from the repository and the user associated with the first user identifier approved the transition of the particular version of the first rule container from the first state to the second state, or (ii) information specifying that the checking out of the particular version of the first rule container from the repository is associated with a first timestamp, and the transition of the particular version of the first rule container from the first state to the second state is associated with a second timestamp.

[0048] In another aspect, in general, a system for managing one or more rules for use with an application program is provided. The system includes: means for providing an integrated user interface to enable users to access an integrated application development system that enables version control and business process management of rule containers, in which each rule container including two or more rulesets, each ruleset includes one or more rules, each rule including one or more rule cases that compute values for the rule's output depending on the rule's input, and each rule case includes one or more input relationships in which a rule case is triggered by a given input record if one or more data values of the given input record meet the one or more input relationships; a repository for storing multiple versions of a first rule container and at least one log data object having information about the multiple versions of the first rule container; a first user interface of the integrated user interface, for receiving instructions from a first user associated with a first user identifier; means for, in response to the instructions from the first user, checking out a particular version of the first rule container from the repository; means for, in response to the instructions from the first user, transitioning the particular version of the first rule container from a first state of a plurality of states to a second state of the plurality of states; and means for, in response to the instructions from the first user, checking in the particular version of the first rule container back into the repository. The system includes: means for storing in a log data object associated with the particular version of the first rule container at least one of: (i) information specifying that a user associated with the first user identifier checked out the particular version of the first rule container from the repository and the user associated with the first user identifier approved the transition of the particular version of the first rule container from the first

state to the second state, or (ii) information specifying that the checking out of the particular version of the first rule container from the repository is associated with a first timestamp, and the transition of the particular version of the first rule container from the first state to the second state is associated with a second timestamp.

[0049] In another aspect, in general, a method for managing one or more rules for use with an application program in a data processing system is provided. The method includes: providing an integrated user interface to enable users to access an integrated application development system that enables editing, simulation, business process management, and version control of rule containers, in which each rule container includes two or more rulesets, each ruleset includes one or more rules, each rule including one or more rule cases that compute values for the rule's output depending on the rule's input, and each rule case includes one or more input relationships in which a rule case is triggered by a given input record if one or more data values of the given input record meet the one or more input relationships; storing multiple versions of a first rule container in a repository; storing at least one log data object having information about the multiple versions of the first rule container in the repository; and receiving, through a first user interface of the integrated user interface, instructions from a first user associated with a first user identifier. The method includes in response to the instructions from the first user, checking out a particular version of the first rule container from the repository, editing the particular version of the first rule container, associating an application program with the particular version of the first rule container, and executing the application program to process input records and produce outputs according to rules of the particular version of the first rule container, transitioning the particular version of the first rule container from a first state of a plurality of states to a second state of the plurality of states, and checking in the particular version of the first rule container back into the repository.

[0050] Aspects can include one or more of the following features. The method can include storing in a log data object associated with the particular version of the first rule container at least one of: (i) information specifying that a user associated with the first user identifier checked out the particular version of the first rule container from the repository, edited the particular version of the first rule container, executed the

application program using the particular version of the first rule container to produce outputs, and approved the transition of the particular version of the first rule container from the first state to the second state, or (ii) information specifying that the checking out of the particular version of the first rule container from the repository is associated with a first timestamp, the editing of the particular version of the first rule container is associated with a second timestamp, the execution of the application program using the particular version of the first rule container is associated with a third timestamp, and the transition of the particular rule container from the first state to the second state is associated with a fourth timestamp.

[0051] The integrated application development system can include a rule editor, a simulation module, a business process management module, and a source code control module. The rule editor can process the storing in the log data object of information specifying that the user associated with the first user identifier edited the particular version of the first rule container, the simulation module can process the storing in the log data object of information specifying that the user associated with the first user identifier executed the application program using the particular version of the first rule container, the business process management module can process the storing in the log data object of information specifying that the user associated with the first user identifier approved the transition of the particular version of the first rule container from the first state to the second state, and the source code control module can process the storing in the log data object of information specifying that the user associated with the first user identifier checked out the particular version of the first rule container from the repository.

[0052] The integrated application development system can include a rule editor, a simulation module, a business process management module, and a source code control module. The source code control module can process the storing in the log data object of information specifying that the checking out of the particular version of the first rule container from the repository is associated with the first timestamp, the rule editor can process the storing in the log data object of information specifying that the editing of the particular version of the first rule container is associated with the second timestamp, the simulation module can process the storing in the log data object of information specifying that the execution of the application program using the particular version of the first rule

container is associated with the third timestamp, and the business process management module can process the storing in the log data object of information specifying that the transition of the particular rule container from the first state to the second state is associated with the fourth timestamp.

[0053] The integrated application development system can include a rule editor, a simulation module, a business process management module, and a source code control module. The rule editor can process editing of the particular version of the first rule container, the simulation module can process the execution of the application program to process the input records and produce the outputs according to the rules of the particular version of the first rule container, the business process management module can process the transitioning of the particular version of the first rule container from the first state to the second state, and the source code control module can process the checking in and checking out of the particular version of the first rule container.

[0054] In another aspect, in general, a computer-readable medium storing a computer program for managing one or more rules for use with an application program in a data processing system is provided. The computer program includes instructions for causing a computer system to: provide an integrated user interface to enable users to access an integrated application development system that enables editing, simulation, business process management, and version control of rule containers, in which each rule container includes two or more rulesets, each ruleset includes one or more rules, each rule including one or more rule cases that compute values for the rule's output depending on the rule's input, and each rule case includes one or more input relationships in which a rule case is triggered by a given input record if one or more data values of the given input record meet the one or more input relationships; store multiple versions of a first rule container in a repository; store at least one log data object having information about the multiple versions of the first rule container in the repository; and receive, through a first user interface of the integrated user interface, instructions from a first user associated with a first user identifier. The computer program includes instructions for causing the computer system to: in response to the instructions from the first user, check out a particular version of the first rule container from the repository, edit the particular version of the first rule container, associate an application program with the particular version of

the first rule container, and executing the application program to process input records and produce outputs according to rules of the particular version of the first rule container, transition the particular version of the first rule container from a first state of a plurality of states to a second state of the plurality of states, and check in the particular version of the first rule container back into the repository.

[0055] In another aspect, in general, a system for managing one or more rules for use with an application program is provided. The system includes: a data storage system storing one or more rules; at least one output device or port for providing a graphical user interface; and a computer system coupled to the data storage system and the at least one output device or port. The computer system includes circuitry to: provide an integrated user interface to enable users to access an integrated application development system that enables editing, simulation, business process management, and version control of rule containers, in which each rule container includes two or more rulesets, each ruleset includes one or more rules, each rule including one or more rule cases that compute values for the rule's output depending on the rule's input, and each rule case includes one or more input relationships in which a rule case is triggered by a given input record if one or more data values of the given input record meet the one or more input relationships; store multiple versions of a first rule container in a repository; store at least one log data object having information about the multiple versions of the first rule container in the repository; and receive, through a first user interface of the integrated user interface, instructions from a first user associated with a first user identifier. The computer system includes circuitry to: in response to the instructions from the first user, check out a particular version of the first rule container from the repository, edit the particular version of the first rule container, associate an application program with the particular version of the first rule container, and executing the application program to process input records and produce outputs according to rules of the particular version of the first rule container, transition the particular version of the first rule container from a first state of a plurality of states to a second state of the plurality of states, and check in the particular version of the first rule container back into the repository.

[0056] In another aspect, in general, a system for managing one or more rules for use with an application program is provided. The system includes: means for providing

an integrated user interface to enable users to access an integrated application development system that enables editing, simulation, business process management, and version control of rule containers, in which each rule container includes two or more rulesets, each ruleset includes one or more rules, each rule including one or more rule cases that compute values for the rule's output depending on the rule's input, and each rule case includes one or more input relationships in which a rule case is triggered by a given input record if one or more data values of the given input record meet the one or more input relationships; a repository for storing multiple versions of a first rule container and at least one log data object having information about the multiple versions of the first rule container; a first user interface of the integrated user interface, for receiving instructions from a first user associated with a first user identifier; means for, in response to the instructions from the first user, checking out a particular version of the first rule container from the repository; means for, in response to the instructions from the first user, editing the particular version of the first rule container; means for, in response to the instructions from the first user, associating an application program with the particular version of the first rule container; means for, in response to the instructions from the first user, executing the application program to process input records and produce outputs according to rules of the particular version of the first rule container; means for, in response to the instructions from the first user, transitioning the particular version of the first rule container from a first state of a plurality of states to a second state of the plurality of states; and means for, in response to the instructions from the first user, checking in the particular version of the first rule container back into the repository.

[0057] In another aspect, in general, a method for managing one or more rules for use with an application program in a data processing system is provided. The method includes: providing an integrated user interface to enable users to access an integrated application development system that enables editing, business process management, and version control of rule containers, in which each rule container includes two or more rulesets, each ruleset includes one or more rules, each rule including one or more rule cases that compute values for the rule's output depending on the rule's input, and each rule case includes one or more input relationships in which a rule case is triggered by a given input record if one or more data values of the given input record meet the one or

more input relationships; storing multiple versions of a first rule in a repository; storing at least one log data object having information about the multiple versions of the first rule container in the repository; and receiving, through a first user interface of the integrated user interface, instructions from a first user associated with a first user identifier. The method includes in response to the instructions from the first user, checking out a particular version of the first rule from the repository, in which the particular version of the first rule is included in a first rule container, editing the particular version of the first rule, and assigning an “approved” state to the particular version of the first rule. The method includes, after the particular version of the first rule has been assigned the “approved” state, automatically determining that there is a version of the first rule container that is in an “open” state or a state equivalent to the “open” state, in which the “open” state indicates that the first rule container is in a draft mode that can still be edited. The method includes in response to the determination that there is a version of the first rule container that is in an “open” state or a state equivalent to the “open” state, adding the particular version of the first rule to the version of the first rule container that is in an “open” state or a state equivalent to the “open” state.

[0058] Aspects can include one or more of the following features. The method can include, in response to the instructions from the first user, after the particular version of the first rule has been assigned the “approved” state, automatically determining that there is no version of the first rule container that is in the “open” state or a state equivalent to the “open” state. The method can include, in response to the determination that there is no version of the first rule container that is in the “open” state or a state equivalent to the “open” state, automatically generate a new version of the first rule container, assign the new version of the first rule container to an “open” state, and add the particular version of the first rule to the new version of the first rule container.

[0059] In another aspect, in general, a computer-readable medium storing a computer program for managing one or more rule containers for use with an application program in a data processing system is provided. The computer program includes instructions for causing a computer system to: provide an integrated user interface to enable users to access an integrated application development system that enables editing, business process management, and version control of rule containers, in which each rule

container includes two or more rulesets, each ruleset includes one or more rules, each rule including one or more rule cases that compute values for the rule's output depending on the rule's input, and each rule case includes one or more input relationships in which a rule case is triggered by a given input record if one or more data values of the given input record meet the one or more input relationships; store multiple versions of a first rule in a repository; store at least one log data object having information about the multiple versions of the first rule container in the repository; and receive, through a first user interface of the integrated user interface, instructions from a first user associated with a first user identifier. The computer program includes instructions for causing the computer system to: in response to the instructions from the first user, check out a particular version of the first rule from the repository, in which the particular version of the first rule is included in a first rule container, edit the particular version of the first rule, and assign an "approved" state to the particular version of the first rule. The computer program includes instructions for causing the computer system to: after the particular version of the first rule has been assigned the "approved" state, automatically determine that there is a version of the first rule container that is in an "open" state or a state equivalent to the "open" state, in which the "open" state indicates that the first rule container is in a draft mode that can still be edited. The computer program includes instructions for causing the computer system to: in response to the determination that there is a version of the first rule container that is in an "open" state or a state equivalent to the "open" state, add the particular version of the first rule to the version of the first rule container that is in an "open" state or a state equivalent to the "open" state.

[0060] In another aspect, in general, a system for managing one or more rule containers for use with an application program is provided. The system includes: a data storage system storing one or more rule containers; at least one output device or port for providing a graphical user interface; and a computer system coupled to the data storage system and the at least one output device or port. The computer system includes circuitry to: provide an integrated user interface to enable users to access an integrated application development system that enables editing, business process management, and version control of rule containers, in which each rule container includes two or more rulesets, each ruleset includes one or more rules, each rule including one or more rule cases that

compute values for the rule's output depending on the rule's input, and each rule case includes one or more input relationships in which a rule case is triggered by a given input record if one or more data values of the given input record meet the one or more input relationships; store multiple versions of a first rule in a repository; store at least one log data object having information about the multiple versions of the first rule container in the repository; and receive, through a first user interface of the integrated user interface, instructions from a first user associated with a first user identifier. The computer system includes circuitry to: in response to the instructions from the first user, check out a particular version of the first rule from the repository, in which the particular version of the first rule is included in a first rule container, edit the particular version of the first rule, and assign an "approved" state to the particular version of the first rule. The computer system includes circuitry to: after the particular version of the first rule has been assigned the "approved" state, automatically determine that there is a version of the first rule container that is in an "open" state or a state equivalent to the "open" state, in which the "open" state indicates that the first rule container is in a draft mode that can still be edited. The computer system includes circuitry to: in response to the determination that there is a version of the first rule container that is in an "open" state or a state equivalent to the "open" state, add the particular version of the first rule to the version of the first rule container that is in an "open" state or a state equivalent to the "open" state.

[0061] In another aspect, in general, a system for managing one or more rules for use with an application program is provided. The system includes: means for providing an integrated user interface to enable users to access an integrated application development system that enables editing, business process management, and version control of rule containers, in which each rule container includes two or more rulesets, each ruleset includes one or more rules, each rule including one or more rule cases that compute values for the rule's output depending on the rule's input, and each rule case includes one or more input relationships in which a rule case is triggered by a given input record if one or more data values of the given input record meet the one or more input relationships; a repository for storing multiple versions of a first rule and at least one log data object having information about the multiple versions of the first rule container; a

first user interface of the integrated user interface, for receiving instructions from a first user associated with a first user identifier; means for, in response to the instructions from the first user, checking out a particular version of the first rule from the repository, in which the particular version of the first rule is included in a first rule container, means for, in response to the instructions from the first user, editing the particular version of the first rule, means for, in response to the instructions from the first user, assigning an “approved” state to the particular version of the first rule, means for, after the particular version of the first rule has been assigned the “approved” state, automatically determining that there is a version of the first rule container that is in an “open” state or a state equivalent to the “open” state, in which the “open” state indicates that the first rule container is in a draft mode that can still be edited, and means for, in response to the determination that there is a version of the first rule container that is in an “open” state or a state equivalent to the “open” state, adding the particular version of the first rule to the version of the first rule container that is in an “open” state or a state equivalent to the “open” state.

[0062] In another aspect, in general, a method for managing one or more rules for use with an application program in a data processing system is provided. The method includes: providing a first integrated user interface to enable the checking in and out of the rule specification, the editing of the ruleset, and the receiving of the state transition information; receiving a first instruction from a first user associated with a first user identifier through the first integrated user interface; in response to the first instruction from the first user, checking out a rule specification from a repository, in which the rule specification has a first version number and specifies a ruleset that includes a set of rules, each rule including a set of rule cases that compute values for the rule’s output depending on the rule’s input, each rule case including one or more input relationships in which a rule case is triggered by a given input record if one or more data values of the given input record meet the one or more input relationships; receiving editing instructions from the first user through the first integrated user interface; revising the rule specification based on the editing instructions from the first user; associating the ruleset with a first state of a plurality of states; receiving a second instruction from the first user through the first integrated user interface; in response to the second instruction from the first user,

checking in the rule specification to the repository, and assigning a second version number to the rule specification; receiving transition information associated with the ruleset from a second user associated with a second user identifier through a second integrated user interface; based on the transition information and permitted transitions associated with the ruleset and the first state, associating the ruleset with a second state of the plurality of states, in which the second state is different from the first state; and storing, in the repository, first data that associates changes in the second version of the rule specification with the first user identifier, and second data that associates the transition from the first state to the second state with the second user identifier.

[0063] Aspects can include one or more of the following features. The method can include storing, in the log data object, a first timestamp indicating when the second version of the rule specification was checked into the repository, and a second timestamp indicating when the current state was associated with the rule specification.

[0064] The method can include associating an application program with the ruleset; and executing the application program to process the input records and produce the outputs according to rules of the ruleset. At least one of the first or second integrated user interface also enables the executing of the application program to process records according to the rules of the ruleset.

[0065] The method can include storing multiple versions of the rule specification in the repository; and for each version of the rule specification, storing first data having information about a user identifier associated with a latest modification of the rule specification, and storing second data having information about a user identifier associated with an assignment of a current state or a transition from a previous state to the current state associated with the ruleset specified by the rule specification.

[0066] The method can include for each version of the rule specification, storing a first timestamp indicating when the rule specification was checked into the repository, and a second timestamp indicating when the current state was associated with the rule specification.

[0067] The method can include retrieving a particular version of the rule specification from the repository; and displaying the user identifier associated with the latest modification of the rule specification, and the user identifier associated with the

assignment of the current state or the transition from the previous state to the current state associated with the ruleset specified by the rule specification.

[0068] The method can include providing, in an integrated user interface, a first displayed input element to enable a user to select a particular version of the rule specification; and providing, in the integrated user interface, a second displayed input element to enable the user to associate an application program with the ruleset specified by the selected version of the rule specification, and execute the application program to process input records and produce outputs according to rules of the ruleset specified by the selected version of the rule specification.

[0069] The method can include storing multiple versions of the rule specification in the repository; and for each version of the rule specification, storing first data having information about a user identifier associated with a latest modification of the rule specification, and storing second data having information about a user identifier associated with an assignment of a current state or a transition from a previous state to the current state associated with the ruleset specified by the rule specification.

[0070] The method can include for each version of the rule specification, storing a first timestamp indicating when the rule specification was checked into the repository, and a second timestamp indicating when the current state was associated with the rule specification.

[0071] The method can include storing multiple versions of the rule specification in the repository; and for each version of the rule specification, enabling one or more users to assign an initial state to the ruleset or to transition the ruleset from one state to another state, and each time the rule specification is checked into the repository with a modified state compared to when the rule specification was checked out, storing data that associates the assignment of the initial state to the ruleset or the transition of the ruleset from one state to another state with the user identifier of the user who modified the state associated with the ruleset, and storing a timestamp indicating when the state was modified.

[0072] The method can include providing, in an integrated user interface, a first displayed input element to enable a user to select a particular version of the rule specification; providing, in the integrated user interface, a second displayed input element

to enable the user to select a particular time; providing, in the integrated user interface, a third displayed input element to enable the user to associate an application program with the ruleset specified by the selected version of the rule specification; executing the application program to process input records and produce outputs according to rules of the ruleset specified by the selected version of the rule specification; displaying the state of the ruleset at the selected time; and displaying the user identifier associated with the user who caused the ruleset to be associated with the state at the selected time.

[0073] Storing the first data and the second data in the repository can include storing or modifying one or more tags associated with the ruleset, and the one or more tags can include a pointer to a corresponding version of ruleset and information of the state of the ruleset.

[0074] In another aspect, in general, a computer-readable medium storing a computer program for managing one or more rules for use with an application program in a data processing system is provided. The computer program includes instructions for causing a computer system to: provide a first integrated user interface to enable the checking in and out of the rule specification, the editing of the ruleset, and the receiving of the state transition information; receive a first instruction from a first user associated with a first user identifier through the first integrated user interface; in response to the first instruction from the first user, check out a rule specification from a repository, in which the rule specification has a first version number and specifies a ruleset that includes a set of rules, each rule including a set of rule cases that compute values for the rule's output depending on the rule's input, each rule case including one or more input relationships in which a rule case is triggered by a given input record if one or more data values of the given input record meet the one or more input relationships; receive editing instructions from the first user through the first integrated user interface; revise the rule specification based on the editing instructions from the first user; associate the ruleset with a first state of a plurality of states; receive a second instruction from the first user through the first integrated user interface; in response to the second instruction from the first user, check in the rule specification to the repository, and assign a second version number to the rule specification; receive transition information associated with the ruleset from a second user associated with a second user identifier through a second integrated

user interface; based on the transition information and permitted transitions associated with the ruleset and the first state, associate the ruleset with a second state of the plurality of states, in which the second state is different from the first state; store, in the repository, first data that associates changes in the second version of the rule specification with the first user identifier, and second data that associates the transition from the first state to the second state with the second user identifier.

[0075] In another aspect, in general, a system for managing one or more rules for use with an application program is provided. The system includes: a data storage system storing one or more rule containers; at least one output device or port for providing a graphical user interface; and a computer system coupled to the data storage system and the at least one output device or port. The computer system includes circuitry to: provide a first integrated user interface to enable the checking in and out of the rule specification, the editing of the ruleset, and the receiving of the state transition information; receive a first instruction from a first user associated with a first user identifier through the first integrated user interface; in response to the first instruction from the first user, check out a rule specification from a repository, in which the rule specification has a first version number and specifies a ruleset that includes a set of rules, each rule including a set of rule cases that compute values for the rule's output depending on the rule's input, each rule case including one or more input relationships in which a rule case is triggered by a given input record if one or more data values of the given input record meet the one or more input relationships; receive editing instructions from the first user through the first integrated user interface; revise the rule specification based on the editing instructions from the first user; associate the ruleset with a first state of a plurality of states; receive a second instruction from the first user through the first integrated user interface; in response to the second instruction from the first user, check in the rule specification to the repository, and assign a second version number to the rule specification; receive transition information associated with the ruleset from a second user associated with a second user identifier through a second integrated user interface; based on the transition information and permitted transitions associated with the ruleset and the first state, associate the ruleset with a second state of the plurality of states, in which the second state is different from the first state; and store, in the repository, first data that associates

changes in the second version of the rule specification with the first user identifier, and second data that associates the transition from the first state to the second state with the second user identifier.

[0076] In another aspect, in general, a system for managing one or more rules for use with an application program is provided. The system includes: means for providing a first integrated user interface to enable the checking in and out of the rule specification, the editing of the ruleset, and the receiving of the state transition information; means for receiving a first instruction from a first user associated with a first user identifier through the first integrated user interface; means for, in response to the first instruction from the first user, checking out a rule specification from a repository, in which the rule specification has a first version number and specifies a ruleset that includes a set of rules, each rule including a set of rule cases that compute values for the rule's output depending on the rule's input, each rule case including one or more input relationships in which a rule case is triggered by a given input record if one or more data values of the given input record meet the one or more input relationships; means for receiving editing instructions from the first user through the first integrated user interface; means for revising the rule specification based on the editing instructions from the first user; means for associating the ruleset with a first state of a plurality of states; means for receiving a second instruction from the first user through the first integrated user interface; means for, in response to the second instruction from the first user, checking in the rule specification to the repository, and assigning a second version number to the rule specification; means for receiving transition information associated with the ruleset from a second user associated with a second user identifier through a second integrated user interface; means for, based on the transition information and permitted transitions associated with the ruleset and the first state, associating the ruleset with a second state of the plurality of states, in which the second state is different from the first state; and means for storing, in the repository, first data that associates changes in the second version of the rule specification with the first user identifier, and second data that associates the transition from the first state to the second state with the second user identifier.

[0077] In another aspect, in general, a method for managing one or more rule containers for use with an application program in a data processing system, each rule

container including two or more rulesets, the method including: providing an integrated application development system to enable checking in and out of a rule container and rule specifications included in the rule container, and transition of the rule container from one state to another state, wherein the integrated application development system includes an integrated user interface that includes user interfaces for receiving instructions from users and providing information to the users; and receiving, through a first interface of the integrated user interface, instructions from a first user associated with a first user identifier. The method includes, in response to the first instruction from the first user, checking out a rule container from a repository, the rule container including at least a first rule specification and a second rule specification, in which the rule container has a first version number, the first rule specifications has a second version number, the second rule specification has a third version number, each of the first and second rule specifications specifies a ruleset that includes a set of rules, each rule including a set of rule cases that compute values for the rule's output depending on the rule's input, each rule case including one or more input relationships in which a rule case is triggered by a given input record if one or more data values of the given input record meet the one or more input relationships; associating a first ruleset specified by the first rule specification with a first state of a plurality of states for rulesets; associating the rule container with a first state of a plurality of states for rule containers; and checking in the rule container to the repository and assigning a fourth version number to the rule container, and assigning a fifth version number to the first rule specification.

[0078] Aspects can include one or more of the following features. The method can include receiving, through the first interface of the integrated user interface, first transition information associated with the rule container and second transition information associated with the first ruleset, the first and second transition information being provided by the first user; based on the first transition information and permitted transitions associated with the rule container, associating the rule container with a second state of the plurality of states for rule containers, in which the second state for rule containers is different from the first state for rule containers; based on the second transition information and permitted transitions associated with the first ruleset, associating the first ruleset with a second state of the plurality of states for rulesets, in

which the second state for rulesets is different from the first state for rulesets; and storing, in a log data object, first log data specifying that the first user checked out the rule container, second log data specifying that the first user approved the rule container to transition from the first state to the second state for rule containers, and third log data specifying that the first user approved the first ruleset to transition from the first state to the second state for rulesets.

[0079] The first log data in the log data object can specify that the first user checked out the rule container having the first version number, the second log data can specify that the first user approved the rule container having the first version number to transition from the first state to the second state for rule containers, and the third log data can specify that the first user approved the first ruleset having the second version number to transition from the first state to the second state for rulesets.

[0080] The method can include receiving, through a second interface of the integrated user interface, a second instruction from a second user associated with a second user identifier; in response to the second instruction from the second user, checking out the rule container from the repository, receiving, through the second interface of the integrated user interface, third transition information associated with the rule container and fourth transition information associated with the first ruleset, the third and fourth transition information being provided by the second user; based on the third transition information and permitted transitions associated with the rule container, associating the rule container with a third state of the plurality of states for rule containers, in which the third state for rule containers is different from the second state for rule containers; based on the fourth transition information and permitted transitions associated with the first ruleset, associating the first ruleset with a third state of the plurality of states for rulesets, in which the third state for rulesets is different from the second state for rulesets; and storing, in the log data object, fourth log data specifying that the second user checked out the rule container, fifth log data specifying that the second user approved the rule container to transition from the second state to the third state for rule containers, and sixth log data specifying that the second user approved the first ruleset to transition from the second state to the third state for rulesets.

[0081] The fourth log data in the log data object can specify that the second user checked out the rule container having the fourth version number, the fifth log data can specify that the second user approved the rule container having the fourth version number to transition from the second state to the third state for rule containers, and the sixth log data can specify that the second user approved the first ruleset having the fifth version number to transition from the first state to the second state for rulesets.

[0082] The method can include receiving, through a second interface of the integrated user interface, a second instruction from a second user associated with a second user identifier; in response to the second instruction from the second user, checking out the rule container from the repository, and receiving, through the second interface of the integrated user interface, a third instruction from the second user. The method can include, in response to the third instruction from the second user, performing at least one of (i) associating an application program with the rule container, and executing the application program to process input records and produce outputs according to rules of the rulesets in the rule container, or (ii) retrieving previously stored outputs produced from a prior execution of an application program using the rules of the rulesets in the rule container to process input records.

[0083] The method can include storing, in a log data object, first log data specifying that the first user checked out the rule container at a first time point, second log data specifying that the second user checked out the rule container at a second time point, and third log data specifying that the second user performed at least one of (i) executing the application program to process input records and produce outputs according to rules of the rulesets in the rule container, or (ii) retrieving previously stored outputs produced from the prior execution of the application program using the rules of the rulesets in the rule container to process input records.

[0084] The method can include receiving, through a third interface of the integrated user interface, a fourth instruction from a third user associated with a third user identifier; in response to the fourth instruction from the third user, checking out the rule container from the repository; receiving, through the third interface of the integrated user interface, transition information from the third user; and in response to the transition

information from the third user, transitioning the rule container to a release state indicating that the rule container is ready to be released.

[0085] The method can include storing, in a log data object, first log data specifying that the first user checked out the rule container at a first time point, second log data specifying that the second user checked out the rule container at a second time point, third log data specifying that the second user performed at least one of (i) executing the application program to process input records and produce outputs according to rules of the rulesets in the rule container, or (ii) retrieving previously stored outputs that are the result of the prior execution of the application program using the rules of the rulesets in the rule container to process input records, fourth log data specifying that the third user checked out the rule container at a third time point, and fifth log data specifying that the third user approved transitioning the rule container to the release state.

[0086] In another aspect, in general, a computer-readable medium storing a computer program for managing one or more rule containers for use with an application program in a data processing system is provided. The computer program includes instructions for causing a computer system to: provide an integrated application development system to enable checking in and out of a rule container and rule specifications included in the rule container, and transition of the rule container from one state to another state, wherein the integrated application development system includes an integrated user interface that includes user interfaces for receiving instructions from users and providing information to the users; and receive, through a first interface of the integrated user interface, instructions from a first user associated with a first user identifier. The computer program includes instructions for causing the computer system to: in response to the first instruction from the first user, check out a rule container from a repository, the rule container including at least a first rule specification and a second rule specification, in which the rule container has a first version number, the first rule specifications has a second version number, the second rule specification has a third version number, each of the first and second rule specifications specifies a ruleset that includes a set of rules, each rule including a set of rule cases that compute values for the rule's output depending on the rule's input, each rule case including one or more input relationships in which a rule case is triggered by a given input record if one or more data

values of the given input record meet the one or more input relationships; associate a first ruleset specified by the first rule specification with a first state of a plurality of states for rulesets; associate the rule container with a first state of a plurality of states for rule containers; and check in the rule container to the repository and assigning a fourth version number to the rule container, and assigning a fifth version number to the first rule specification.

[0087] In another aspect, in general, a system for managing one or more rule containers for use with an application program is provided. The system includes: a data storage system storing one or more rule containers; at least one output device or port for providing a graphical user interface; and a computer system coupled to the data storage system and the at least one output device or port. The computer system includes circuitry to: provide an integrated application development system to enable checking in and out of a rule container and rule specifications included in the rule container, and transition of the rule container from one state to another state, wherein the integrated application development system includes an integrated user interface that includes user interfaces for receiving instructions from users and providing information to the users; and receive, through a first interface of the integrated user interface, instructions from a first user associated with a first user identifier. The computer system includes circuitry to: in response to the first instruction from the first user, check out a rule container from a repository, the rule container including at least a first rule specification and a second rule specification, in which the rule container has a first version number, the first rule specifications has a second version number, the second rule specification has a third version number, each of the first and second rule specifications specifies a ruleset that includes a set of rules, each rule including a set of rule cases that compute values for the rule's output depending on the rule's input, each rule case including one or more input relationships in which a rule case is triggered by a given input record if one or more data values of the given input record meet the one or more input relationships; associate a first ruleset specified by the first rule specification with a first state of a plurality of states for rulesets; associate the rule container with a first state of a plurality of states for rule containers; and check in the rule container to the repository and assigning a fourth

version number to the rule container, and assigning a fifth version number to the first rule specification.

[0088] In another aspect, in general, a system for managing one or more rules for use with an application program is provided. The system includes: means for providing an integrated application development system to enable checking in and out of a rule container and rule specifications included in the rule container, and transition of the rule container from one state to another state, wherein the integrated application development system includes an integrated user interface that includes user interfaces for receiving instructions from users and providing information to the users; means for receiving, through a first interface of the integrated user interface, instructions from a first user associated with a first user identifier; means for, in response to the first instruction from the first user, checking out a rule container from a repository, the rule container including at least a first rule specification and a second rule specification, in which the rule container has a first version number, the first rule specifications has a second version number, the second rule specification has a third version number, each of the first and second rule specifications specifies a ruleset that includes a set of rules, each rule including a set of rule cases that compute values for the rule's output depending on the rule's input, each rule case including one or more input relationships in which a rule case is triggered by a given input record if one or more data values of the given input record meet the one or more input relationships; means for, in response to the first instruction from the first user, associating a first ruleset specified by the first rule specification with a first state of a plurality of states for rulesets; means for, in response to the first instruction from the first user, associating the rule container with a first state of a plurality of states for rule containers; and means for, in response to the first instruction from the first user, checking in the rule container to the repository and assigning a fourth version number to the rule container, and assigning a fifth version number to the first rule specification.

[0089] In another aspect, in general, a method for managing one or more rules for use with an application program in a data processing system is provided. The method includes: providing an integrated user interface to enable version control and business process management of rulesets, in which each ruleset includes a set of rule cases that compute values for the rule's output depending on the rule's input, each rule case

including one or more input relationships in which a rule case is triggered by a given input record if one or more data values of the given input record meet the one or more input relationships; receiving, through a first interface of the integrated user interface, instructions from a first user associated with a first user identifier; in response to the instructions from the first user, checking out a first ruleset having a first version number from a repository, transitioning the first ruleset from a first state to a second state, and checking in the first ruleset to the repository and assigning a second version number to the first ruleset; and receiving, through a second interface of the integrated user interface, instructions from a second user associated with a second user identifier. The method includes, in response to the instructions from the second user, checking out the first ruleset having the second version number from the repository, transitioning the first ruleset from the second state to a third state, and checking in the first ruleset to the repository. The method includes storing, in a common log, information specifying that the first user checked out the first ruleset having the first version number, the first user approved the first ruleset to transition from the first state to the second state, the second user checked out the first ruleset having the second version number, and the second user approved the first ruleset to transition from the second state to the third state.

[0090] Aspects can include one or more of the following features. The integrated user interface can be configured to also enable simulation of rulesets, and the method can include: receiving, through a third interface of the integrated user interface, instructions from a third user associated with a third user identifier. The method can include, in response to the instructions from the third user, checking out the first ruleset having the second version number from the repository, and performing at least one of (i) associating an application program with the first ruleset, and executing the application program to process input records and produce outputs according to rules of the first ruleset, or (ii) retrieving previously stored outputs produced from a prior execution of an application program using rules of the first ruleset to process input records. The method can include storing, in the common log, information specifying that the third user checked out the first ruleset having the second version number, and the third user performed at least one of (i) executing the application program to process input records and produce outputs according to rules of the first ruleset, or (ii) retrieving previously stored outputs produced

from the prior execution of the application program using rules of the first ruleset to process input records.

[0091] The integrated user interface can be configured to also enable editing of rulesets, and the method can include: receiving, through the first interface of the integrated user interface, editing instructions from the first user. The method can include, in response to the editing instructions from the first user, editing the first ruleset. The method can include storing, in the common log, information specifying that the first user edited the first ruleset.

[0092] The first interface can be shown at a display of the first user, and the second interface can be shown at a display of the second user.

[0093] In another aspect, in general, a computer-readable medium storing a computer program for managing one or more rules for use with an application program in a data processing system is provided. The computer program includes instructions for causing a computer system to: provide an integrated user interface to enable version control and business process management of rulesets, in which each ruleset includes a set of rule cases that compute values for the rule's output depending on the rule's input, each rule case including one or more input relationships in which a rule case is triggered by a given input record if one or more data values of the given input record meet the one or more input relationships; and receive, through a first interface provided by the integrated user interface, instructions from a first user associated with a first user identifier; in response to the instructions from the first user, check out a first ruleset having a first version number from a repository, transitioning the first ruleset from a first state to a second state, and checking in the first ruleset to the repository and assigning a second version number to the first ruleset; receive, through a second interface provided by the integrated user interface, instructions from a second user associated with a second user identifier; in response to the instructions from the second user, check out the first ruleset having the second version number from the repository, transitioning the first ruleset from the second state to a third state, and checking in the first ruleset to the repository; and store, in a common log, information specifying that the first user checked out the first ruleset having the first version number, the first user approved the first ruleset to transition from the first state to the second state, the second user checked out the first

ruleset having the second version number, and the second user approved the first ruleset to transition from the second state to the third state.

[0094] In another aspect, in general, a system for managing one or more rules for use with an application program is provided. The system includes: a data storage system storing one or more rules; at least one output device or port for providing a graphical user interface; and a computer system coupled to the data storage system and the at least one output device or port. The computer system includes circuitry to: provide an integrated user interface to enable version control and business process management of rulesets, in which each ruleset includes a set of rule cases that compute values for the rule's output depending on the rule's input, each rule case including one or more input relationships in which a rule case is triggered by a given input record if one or more data values of the given input record meet the one or more input relationships; receive, through a first interface provided by the integrated user interface, instructions from a first user associated with a first user identifier; in response to the instructions from the first user, check out a first ruleset having a first version number from a repository, transitioning the first ruleset from a first state to a second state, and checking in the first ruleset to the repository and assigning a second version number to the first ruleset; receive, through a second interface provided by the integrated user interface, instructions from a second user associated with a second user identifier; in response to the instructions from the second user, check out the first ruleset having the second version number from the repository, transitioning the first ruleset from the second state to a third state, and checking in the first ruleset to the repository; and store, in a common log, information specifying that the first user checked out the first ruleset having the first version number, the first user approved the first ruleset to transition from the first state to the second state, the second user checked out the first ruleset having the second version number, and the second user approved the first ruleset to transition from the second state to the third state.

[0095] In another aspect, in general, a system for managing one or more rules for use with an application program is provided. The system includes: means for providing an integrated user interface to enable version control and business process management of rulesets, in which each ruleset includes a set of rule cases that compute values for the rule's output depending on the rule's input, each rule case including one or more input

relationships in which a rule case is triggered by a given input record if one or more data values of the given input record meet the one or more input relationships; means for receiving, through a first interface provided by the integrated user interface, instructions from a first user associated with a first user identifier; means for, in response to the instructions from the first user, checking out a first ruleset having a first version number from a repository, transitioning the first ruleset from a first state to a second state, and checking in the first ruleset to the repository and assigning a second version number to the first ruleset; means for receiving, through a second interface of the integrated user interface, instructions from a second user associated with a second user identifier; means for, in response to the instructions from the second user, checking out the first ruleset having the second version number from the repository, transitioning the first ruleset from the second state to a third state, and checking in the first ruleset to the repository; and means for storing, in a common log, information specifying that the first user checked out the first ruleset having the first version number, the first user approved the first ruleset to transition from the first state to the second state, the second user checked out the first ruleset having the second version number, and the second user approved the first ruleset to transition from the second state to the third state.

[0096] In another aspect, in general, a method for managing one or more rules for use with an application program in a data processing system is provided. The method includes: providing an integrated user interface to enable version control and business process management of rule containers, in which each rule container includes a plurality of rulesets, each ruleset includes a set of rule cases that compute values for the rule's output depending on the rule's input, and each rule case includes one or more input relationships in which a rule case is triggered by a given input record if one or more data values of the given input record meet the one or more input relationships; and receiving, through a first interface of the integrated user interface, instructions from a first user associated with a first user identifier. The method includes: in response to the instructions from the first user, checking out a first version of a first rule container from a repository, in which the first version of the first rule container is associated with a first version of a first ruleset and a first version of a second ruleset, the first version of the first rule container includes a pointer to the first version of the first ruleset and a pointer to the first

version of the second ruleset, and the first version of the second ruleset is associated with a first state for rulesets, generating a second version of the first rule container, generating a second version of the second ruleset that is a copy of the first version of the second ruleset, in which the second version of the second ruleset is associated with a first state for rulesets, transitioning the first version of the second ruleset from the first state for rulesets to a second state for rulesets, adding, to the second version of the first rule container, a pointer to the first version of the first ruleset and a pointer to the second version of the second ruleset, and checking in the first and second versions of the first rule container back into the repository.

[0097] Aspects can include one or more of the following features. The method can include storing, in a log data object, at least one of (i) information specifying that the first user checked out the first version of the first rule container, generated the second version of the first rule container, and approved transition of the first version of the second ruleset from the first state to the second state, or (ii) information specifying that the checking out of the first version of the first rule container is associated with a first timestamp, the generation of the second version of the first rule container is associated with a second timestamp, and the transition of the first version of the second ruleset from the first state to the second state is associated with a third timestamp.

[0098] The first version of the first rule container can be initially associated with a first state for rule containers, and the method can include: in response to the instructions from the first user, transitioning the first version of the first rule container from the first state for rule containers to a second state for rule containers.

[0099] The method can include storing in the log data object, at least one of (i) information specifying that the first user approved transition of the first version of the first rule container from the first state to the second state, or (ii) information specifying that the transition of the first version of the first rule container from the first state to the second state is associated with a first timestamp.

[0100] The method can include: receiving, through a second interface of the integrated user interface, instructions from a second user associated with a second user identifier. The method can include: in response to the instructions from the second user, checking out the first version of the first rule container from the repository, and

performing at least one of (i) associating an application program with the first version of the first rule container, and executing the application program to process input records and produce outputs according to rules of the rulesets in the first version of the first rule container, or (ii) retrieving previously stored outputs produced from a prior execution of an application program using rules of the rulesets in the first version of the first rule container to process input records.

[0101] The method can include storing in a log data object, at least one of (i) information specifying that the second user checked out the first version of the first rule container and executed the application program in combination with the first version of the first rule container, (ii) information specifying that the second user checked out the first version of the first rule container and retrieved previously stored outputs produced from a prior execution of the application program using rules of the rulesets in the first version of the first rule container to process input records, (iii) information specifying that the checking out of the first version of the first rule container is associated with a first timestamp, and execution of the application program in combination with the first version of the first rule container is associated with a second timestamp, or (iv) information specifying that the checking out of the first version of the first rule container is associated with a first timestamp, and retrieval of previously stored outputs produced from a prior execution of the application program using rules of the rulesets in the first version of the first rule container to process input records is associated with a second timestamp.

[0102] The first interface can be provided to a display of the first user and the second interface can be provided to a display of the second user.

[0103] The method can include: receiving, through a second interface of the integrated user interface, instructions from a second user associated with a second user identifier. The method can include: in response to the instructions from the second user, checking out the first version of the first rule container from the repository, editing the first version of the first ruleset in the first rule container, transitioning the first version of the first ruleset from a first state for rulesets to a second state for rulesets, and checking in the first version of the first rule container back into the repository.

[0104] The method can include storing in a log data object, at least one of: (i) information specifying that the second user checked out the first version of the first rule container, edited the first version of the first ruleset, and approved transition of the first version of the first ruleset from the first state to the second state, or (ii) information specifying that the checking out of the first version of the first rule container is associated with a first timestamp, the editing of the first version of the first ruleset is associated with a second timestamp, and the transition of the first version of the first ruleset from the first state to the second state is associated with a third timestamp.

[0105] In another aspect, in general, a computer-readable medium storing a computer program for managing one or more rule containers for use with an application program in a data processing system is provided. The computer program includes instructions for causing a computer system to: provide an integrated user interface to enable version control and business process management of rule containers, in which each rule container includes a plurality of rulesets, each ruleset includes a set of rule cases that compute values for the rule's output depending on the rule's input, and each rule case includes one or more input relationships in which a rule case is triggered by a given input record if one or more data values of the given input record meet the one or more input relationships; and receive, through a first interface of the integrated user interface, instructions from a first user associated with a first user identifier. The computer program includes instructions for causing the computer system to: in response to the instructions from the first user, check out a first version of a first rule container from a repository, in which the first version of the first rule container is associated with a first version of a first ruleset and a first version of a second ruleset, the first version of the first rule container includes a pointer to the first version of the first ruleset and a pointer to the first version of the second ruleset, and the first version of the second ruleset is associated with a first state for rulesets; generate a second version of the first rule container; generate a second version of the second ruleset that is a copy of the first version of the second ruleset, in which the second version of the second ruleset is associated with a first state for rulesets; transition the first version of the second ruleset from the first state for rulesets to a second state for rulesets; add, to the second version of the first rule container, a pointer to the first version of the first ruleset and a pointer to the

second version of the second ruleset; and check in the first and second versions of the first rule container back into the repository.

[0106] In another aspect, in general, a system for managing one or more rule containers for use with an application program is provided. The system includes: a data storage system storing one or more rule containers; at least one output device or port for providing a graphical user interface; and a computer system coupled to the data storage system and the at least one output device or port. The computer system includes circuitry to: provide an integrated user interface to enable version control and business process management of rule containers, in which each rule container includes a plurality of rulesets, each ruleset includes a set of rule cases that compute values for the rule's output depending on the rule's input, and each rule case includes one or more input relationships in which a rule case is triggered by a given input record if one or more data values of the given input record meet the one or more input relationships; and receive, through a first interface of the integrated user interface, instructions from a first user associated with a first user identifier. The computer system includes circuitry to: in response to the instructions from the first user, check out a first version of a first rule container from a repository, in which the first version of the first rule container is associated with a first version of a first ruleset and a first version of a second ruleset, the first version of the first rule container includes a pointer to the first version of the first ruleset and a pointer to the first version of the second ruleset, and the first version of the second ruleset is associated with a first state for rulesets; generate a second version of the first rule container; generate a second version of the second ruleset that is a copy of the first version of the second ruleset, in which the second version of the second ruleset is associated with a first state for rulesets; transition the first version of the second ruleset from the first state for rulesets to a second state for rulesets; add, to the second version of the first rule container, a pointer to the first version of the first ruleset and a pointer to the second version of the second ruleset; and check in the first and second versions of the first rule container back into the repository.

[0107] In another aspect, in general, a system for managing one or more rule containers for use with an application program is provided. The system includes: means for providing an integrated user interface to enable version control and business process

management of rule containers, in which each rule container includes a plurality of rulesets, each ruleset includes a set of rule cases that compute values for the rule's output depending on the rule's input, and each rule case includes one or more input relationships in which a rule case is triggered by a given input record if one or more data values of the given input record meet the one or more input relationships; means for receiving, through a first interface of the integrated user interface, instructions from a first user associated with a first user identifier; means for, in response to the instructions from the first user, checking out a first version of a first rule container from a repository, in which the first version of the first rule container is associated with a first version of a first ruleset and a first version of a second ruleset, the first version of the first rule container includes a pointer to the first version of the first ruleset and a pointer to the first version of the second ruleset, and the first version of the second ruleset is associated with a first state for rulesets; means for, in response to the instructions from the first user, generating a second version of the first rule container; means for, in response to the instructions from the first user, generating a second version of the second ruleset that is a copy of the first version of the second ruleset, in which the second version of the second ruleset is associated with a first state for rulesets; means for, in response to the instructions from the first user, transitioning the first version of the second ruleset from the first state for rulesets to a second state for rulesets; means for, in response to the instructions from the first user, adding, to the second version of the first rule container, a pointer to the first version of the first ruleset and a pointer to the second version of the second ruleset; and means for, in response to the instructions from the first user, checking in the first and second versions of the first rule container back into the repository.

[0108] In another aspect, in general, a method for managing one or more rules for use with an application program in a data processing system is provided. The method includes: providing an integrated user interface to enable users to interact with an integrated application development system to perform version control and business process management of rule containers, in which each rule container includes a plurality of rulesets, each ruleset includes a set of rule cases that compute values for the rule's output depending on the rule's input, and each rule case includes one or more input relationships in which a rule case is triggered by a given input record if one or more data

values of the given input record meet the one or more input relationships; storing multiple versions of a first rule container in a repository, each version of the first rule container including multiple rulesets; and receiving, through a first interface of an integrated user interface, instructions from a first user associated with a first user identifier. The method includes, in response to instructions from the first user, checking out a particular version of the first rule container from the repository; performing at least one of (i) transitioning the particular version of the first rule container from a first state to a second state, or (ii) transitioning a particular ruleset in the particular version of the first rule container from a first state to a second state; and checking in the particular version of the first rule container back into the repository.

[0109] Aspects can include one or more of the following features. The method can include, storing in a log data object at least one of (i) information specifying that the first user checked out the particular version of the first rule container and approved the transition of the particular version of the first rule container from the first state to the second state, (ii) information specifying that checking out of the particular version of the first rule container is associated with a first timestamp, and the transition of the particular version of the first rule container from the first state to the second state is associated with a second timestamp, (iii) information specifying that the first user checked out the particular version of the first rule container and approved the transition of the particular ruleset in the particular version of the first rule container from the first state to the second state, or (iv) information specifying that the checking out of the particular version of the first rule container is associated with a first timestamp, and the transition of the particular ruleset in the particular version of the first rule container from the first state to the second state is associated with a second timestamp.

[0110] For each rule container that includes multiple rulesets, the rule container can include pointers to the rulesets, and each pointer can point to a location in the repository where the corresponding ruleset is stored.

[0111] At least two different rule containers can include pointers that point to a common ruleset.

[0112] The method can include receiving, through the first interface, a second set of instructions from the first user. The method can include, in response to second set of

instructions from the first user, checking out a particular version of a second rule container from the repository; performing at least one of (i) associating an application program with the particular version of the second rule container, and executing the application program to process input records and produce outputs according to rules of the rulesets in the particular version of the second rule container, or (ii) retrieving previously stored outputs produced from a prior execution of an application program that was associated with the particular version of the second rule container; and checking in the particular version of the second rule container back into the repository.

[0113] The method can include storing in a log data object at least one of (i) information specifying that the first user checked out the particular version of the second rule container and executed the application program in combination with the particular version of the second rule container, (ii) information specifying that the first user checked out the particular version of the second rule container and retrieved previously stored outputs produced from a prior execution of the application program in combination with the particular version of the second rule container, (iii) information specifying that the checking out of the particular version of the second rule container is associated with a first timestamp, and execution of the application program in combination with the particular version of the second rule container is associated with a second timestamp, or (iv) information specifying that the checking out of the particular version of the second rule container is associated with a first timestamp, and retrieval of previously stored outputs produced from a prior execution of the application program in combination with the particular version of the second rule container is associated with a second timestamp.

[0114] The method can include: enabling, through the integrated user interface, one or more authors of a rule container to edit the rule container and approve transition of the rule container from a first state to a second state, in which the first state represents that the rule container is being edited, and the second state represents that the rule container is ready for review. The method can include enabling, through the integrated user interface, one or more reviewers to: review content of the rule container; perform at least one of (i) associating an application program with the rule container, and executing the application program to process input records and produce outputs according to rules of rulesets in the rule container, or (ii) retrieving previously stored outputs produced from

a prior execution of an application program in combination with the rule container; and perform at least one of (i) rejecting the rule container and cause the rule container to transition from the second state to the first state, or (ii) approve the rule container to transition from the second state to a third state, in which the third state represents that the rule container is ready for release.

[0115] The method can include: enabling, through the integrated user interface, one or more users to manage multiple releases of the first rule container, the multiple releases including a current release and an earlier release; retrieve the earlier release of the first rule container from the repository; and perform at least one of (i) associating an application program with the earlier release of the first rule container, and executing the application program to process input records and produce outputs according to rules of rulesets in the earlier release of the first rule container, or (ii) retrieving previously stored outputs produced from a prior execution of an application program in combination with the earlier release of the first rule container.

[0116] The method can include enabling, through the integrated user interface, the one or more users to review a log data object having information about what modifications were made to the earlier release of the first rule container, who made the modifications, when the modifications were made, who approved the earlier release of the first rule container to transition from the first state to the second state, and who approved the earlier release of the first rule container to transition from the second state to the third state.

[0117] In another aspect, in general, a computer-readable medium storing a computer program for managing one or more rule containers for use with an application program in a data processing system is provided. The computer program includes instructions for causing a computer system to: provide an integrated user interface to enable users to interact with an integrated application development system to perform version control and business process management of rule containers, in which each rule container includes a plurality of rulesets, each ruleset includes a set of rule cases that compute values for the rule's output depending on the rule's input, and each rule case includes one or more input relationships in which a rule case is triggered by a given input record if one or more data values of the given input record meet the one or more input

relationships; store multiple versions of a first rule container in a repository, each version of the first rule container including multiple rulesets; and receive, through a first interface of an integrated user interface, instructions from a first user associated with a first user identifier. The computer program includes instructions for causing the computer system to: in response to instructions from the first user, check out a particular version of the first rule container from the repository; perform at least one of (i) transitioning the particular version of the first rule container from a first state to a second state, or (ii) transitioning a particular ruleset in the particular version of the first rule container from a first state to a second state; and check in the particular version of the first rule container back into the repository.

[0118] In another aspect, in general, a system for managing one or more rule containers for use with an application program is provided. The system includes: a data storage system storing one or more rule containers; at least one output device or port for providing a graphical user interface; and a computer system coupled to the data storage system and the at least one output device or port. The computer system includes circuitry to: provide an integrated user interface to enable users to interact with an integrated application development system to perform version control and business process management of rule containers, in which each rule container includes a plurality of rulesets, each ruleset includes a set of rule cases that compute values for the rule's output depending on the rule's input, and each rule case includes one or more input relationships in which a rule case is triggered by a given input record if one or more data values of the given input record meet the one or more input relationships; store multiple versions of a first rule container in a repository, each version of the first rule container including multiple rulesets; and receive, through a first interface of an integrated user interface, instructions from a first user associated with a first user identifier. The computer system includes circuitry to: in response to instructions from the first user, check out a particular version of the first rule container from the repository; perform at least one of (i) transitioning the particular version of the first rule container from a first state to a second state, or (ii) transitioning a particular ruleset in the particular version of the first rule container from a first state to a second state; and check in the particular version of the first rule container back into the repository.

[0119] In another aspect, in general, a system for managing one or more rule containers for use with an application program is provided. The system includes: means for providing an integrated user interface to enable users to interact with an integrated application development system to perform version control and business process management of rule containers, in which each rule container includes a plurality of rulesets, each ruleset includes a set of rule cases that compute values for the rule's output depending on the rule's input, and each rule case includes one or more input relationships in which a rule case is triggered by a given input record if one or more data values of the given input record meet the one or more input relationships; means for storing multiple versions of a first rule container in a repository, each version of the first rule container including multiple rulesets; means for receiving, through a first interface of the integrated user interface, instructions from a first user associated with a first user identifier; means for, in response to instructions from the first user, checking out a particular version of the first rule container from the repository; means for, in response to instructions from the first user, performing at least one of (i) transitioning the particular version of the first rule container from a first state to a second state, or (ii) transitioning a particular ruleset in the particular version of the first rule container from a first state to a second state; and means for, in response to instructions from the first user, checking in the particular version of the first rule container back into the repository.

[0120] In another aspect, in general, a method for managing one or more rules for use with an application program in a data processing system is provided. The method includes: providing an integrated user interface to enable receiving transition information, associating the application program with a ruleset that includes a set of rules, and executing the application program; receiving a rule specification that specifies the ruleset, each rule of the ruleset including a set of rule cases that compute values for the rule's output depending on the rule's input, each rule case including one or more input relationships in which a rule case is triggered by a given input record if one or more data values of the given input record meet the one or more input relationships; associating the ruleset with a first stage of a plurality of stages; receiving, through the integrated user interface, transition information associated with the ruleset; based on the transition

information and permitted transitions associated with the ruleset and the first stage, associating the ruleset with a second stage of the plurality of stages, in which the second stage is different from the first stage; receiving, through the integrated user interface, a first instruction to associate the application program with the ruleset; in response to the first instruction, associating the application program with the ruleset; receiving, through the integrated user interface, a second instruction to execute the application; and in response to the second instruction, executing the application program to process the input records and produce the outputs according to rules of the ruleset.

[0121] Aspects can include one or more of the following features. Receiving a rule specification can include receiving a spreadsheet having a plurality of rows that correspond to the rule cases.

[0122] The spreadsheet can include trigger columns that correspond to available data values, and the rule cases include sets of triggering criteria that relate to the available data values.

[0123] Each rule case can be configured such that the rule case applies to a given record if one or more data values of that record, for each trigger column in which the rule case has a criterion or criteria, meets the triggering criterion or criteria.

[0124] The spreadsheet can include one or more output columns, each output column can correspond to a potential output variable, and a value of the output variable can be determined by a cell in the output column, in which the cell can be at the row corresponding to the rule case that applies to the given record.

[0125] The cell can include at least one of an output value that is assigned to the output variable or an expression that is evaluated to generate the output value.

[0126] The plurality of stages can include a plurality of development stages of a development lifecycle.

[0127] The plurality of development stages can include at least one of a development stage or a test stage.

[0128] The ruleset can include a set of criteria for performing at least one of converting data from one format to another, making determinations about data, or generating new data based on a set of input data.

[0129] The rule specification can include a rule specification table having a plurality of rows, each row defining a rule case and containing a relationship for each of a set of one or more potential inputs.

[0130] At least one of the relationships can include one or more of having a value equal to a threshold, having a value above a threshold, having a value below a threshold, having a value belonging to a set of values, having a value matching a pattern, having a relationship to a value of another input, having a relationship to a value of an output of another rule specification, or having a relationship to a value in a memory.

[0131] At least one of the rows can also contain an output including one or more of a combination of values of the inputs, a pre-determined value, or a value computed from one or more of the values of the inputs.

[0132] The method can include converting the plurality of rule cases in the rule specification to a plurality of logical expressions, compiling the plurality of logical expressions into computer-executable code, and incorporating the computer-executable code into the application program.

[0133] The method can include converting the plurality of rule cases to a transform, and associating the transform with a component in a graph-based computation having data processing components connected by linking elements representing data flows.

[0134] The permissible transitions can include at least one of (i) a ruleset in an earliest stage in a development lifecycle can be promoted to a later stage but cannot be further demoted, (ii) a ruleset that is in neither the earliest stage nor the last stage in the development lifecycle can be promoted to a later stage or demoted to an earlier stage in the development lifecycle, or (iii) a ruleset in a last stage in the development lifecycle can be demoted to an earlier stage but cannot be further promoted.

[0135] The method can include providing the one or more displayed input elements in an editor interface that is part of the integrated user interface, including at least one of (i) when the ruleset is in the earliest stage of the development lifecycle, a first input element for promoting the ruleset is shown and no input element for demoting the ruleset is shown in the editor interface, (ii) when the ruleset is in neither the earliest nor the last development stage, the first input element for promoting the ruleset and a second

input element for demoting the ruleset are shown in the editor interface, or (iii) when the ruleset is in the last development stage, the second input element for demoting the ruleset is shown and no input element for promoting the ruleset is shown in the editor interface.

[0136] Each version of the ruleset can transition through the plurality of development stages.

[0137] The method can include receiving a plurality of rule specifications that specify different versions of a ruleset. The method can include displaying identifiers of the different versions of the ruleset, and the development stages associated with the different versions of the ruleset.

[0138] Displaying identifiers of the different versions of the ruleset and the stages associated with the different versions of the ruleset can include displaying a plurality of major or minor versions of the ruleset and the stages associated with the major or minor versions of the ruleset.

[0139] The method can include storing a copy of the ruleset and naming the copy of the ruleset based on the updated stage associated with the ruleset.

[0140] The method can include receiving a tagging specification that specifies a plurality of possible states of tags associated with the ruleset, one or more input elements associated with each of the states, and a target state associated with each of the input elements.

[0141] In response to a user interaction with a displayed input element, the tag can be associated with the target state that is associated with the displayed input element.

[0142] The method can include using a tag having a pointer to the ruleset to keep track of the state associated with the ruleset, in which the tag transitions among the states based on the permissible transitions.

[0143] The tag can store information about a review history associated with the ruleset.

[0144] The plurality of stages can include a plurality of stages in an approval process of the ruleset.

[0145] In another aspect, in general, a method for managing one or more rules for use with an application program in a data processing system is provided. The method includes: receiving a group of rule specifications, each rule specification specifies a

ruleset that includes a set of rules, each rule including a set of rule cases that compute values for the rule's output depending on the rule's input, each rule case including one or more input relationships in which a rule case is triggered by a given input record if one or more data values of the given input record meet the one or more input relationships; receiving a transition specification that specifies a first plurality of states that can be associated with individual rulesets, a second plurality of states that can be associated with a group of rulesets, permissible state transitions associated with individual rulesets, and permissible state transitions associated with the group of rulesets; associating each ruleset with one of the first plurality of states; associating the group of rulesets with one of the second plurality of states; receiving first transition information associated with one of the rulesets; based on the first transition information, permitted transitions associated with the ruleset, and the state currently associated with the ruleset, associating the ruleset with another one of the first plurality of states; receiving second transition information associated with the group of rulesets; based on the second transition information, permitted transitions associated with the group, and the state currently associated with the group, associating the group of rule specifications with another one of the second plurality of states; and associating an application program with the group of rule specifications, in which the application program is configured to process input records and produce outputs according to the rules of the rulesets specified by the rule specifications.

[0146] Aspects can include one or more of the following features. The first plurality of states can include at least one of (i) a plurality of development lifecycle states, or (ii) a plurality of approval process states.

[0147] The second plurality of states can include at least one of (i) a plurality of development lifecycle states, or (ii) a plurality of approval process states.

[0148] The method can include associating each ruleset with a first tag that contains information about the state associated with the ruleset, and updating the first tag when the state associated with the ruleset changes.

[0149] The method can include associating the group with a second tag that contains information about the state associated with the group of rulesets, and updating the second tag when the state associated with the group of rulesets changes.

[0150] In another aspect, in general, a method for managing one or more rules for use with an application program in a data processing system is provided. The method includes: receiving a rule specification that specifies a ruleset that includes a set of rules, each rule including a set of rule cases that compute values for the rule's output depending on the rule's input, each rule case including one or more input relationships in which a rule case is triggered by a given input record if one or more data values of the given input record meet the one or more input relationships; receiving a transition specification that specifies permissible transitions among a plurality of states associated with the ruleset, and for each transition, one or more users who are authorized to implement the transition; associating the ruleset with a first state of the plurality of states; receiving transition information associated with the ruleset; verifying that a user who provided the transition information is authorized to provide the transition information; after verifying that the user who provided the transition information is authorized to provide the transition information, based on the transition information and permitted transitions associated with the ruleset and the first state, associating the ruleset with a second state of the plurality of states, in which the second state is different from the first state; and associating an application with the ruleset, in which the application program is configured to process input records and produce outputs according to rules of the ruleset.

[0151] Aspects can include one or more of the following features. The plurality of states can include at least one of (i) a plurality of development lifecycle states, or (ii) a plurality of approval process states.

[0152] The method can include associating each ruleset with a tag that contains information about the state associated with the ruleset, and updating the tag when the state associated with the ruleset changes.

[0153] In another aspect, in general, a computer-readable medium storing a computer program for managing one or more rules for use with an application program in a data processing system is provided. The computer program includes instructions for causing a computer system to: provide an integrated user interface to enable receiving transition information, associating the application program with a ruleset that includes a set of rules, and executing the application program; receive a rule specification that specifies the ruleset, each rule of the ruleset including a set of rule cases that compute

values for the rule's output depending on the rule's input, each rule case including one or more input relationships in which a rule case is triggered by a given input record if one or more data values of the given input record meet the one or more input relationships; associate the ruleset with a first stage of a plurality of stages; receive, through the integrated user interface, transition information associated with the ruleset; based on the transition information and permitted transitions associated with the ruleset and the first stage, associate the ruleset with a second stage of the plurality of stages, in which the second stage is different from the first stage; receive, through the integrated user interface, a first instruction to associate the application program with the ruleset; in response to the first instruction, associate the application program with the ruleset; receive, through the integrated user interface, a second instruction to execute the application; and in response to the second instruction, execute the application program to process the input records and produce the outputs according to rules of the ruleset.

[0154] In another aspect, in general, a system for managing one or more rules for use with an application program is provided. The system includes: a data storage system storing one or more rule specifications and one or more transition specifications that define permitted transitions among stages; at least one output device or port for providing a graphical user interface; and a computer system coupled to the data storage system and the at least one output device or port. The computer system includes circuitry to: provide an integrated user interface to enable receiving of transition information, associating the application program with a ruleset that includes a set of rules, and executing the application program; retrieve from the data storage system a rule specification that specifies the ruleset, each rule of the ruleset including a set of rule cases that compute values for the rule's output depending on the rule's input, each rule case including one or more input relationships in which a rule case is triggered by a given input record if one or more data values of the given input record meet the one or more input relationships; associate the ruleset with a first stage of a plurality of stages; receive, through the integrated user interface, transition information associated with the ruleset; based on the transition information and permitted transitions associated with the ruleset and the first stage, associate the ruleset with a second stage of the plurality of stages, in which the second stage is different from the first stage; and receive, through the integrated user

interface, a first instruction to associate the application program with the ruleset. The computer system includes circuitry to: in response to the first instruction, associate an application program with the ruleset; receive, through the integrated user interface, a second instruction to execute the application; and in response to the second instruction, execute the application program to process input records and produce outputs according to the rules of the ruleset.

[0155] In another aspect, in general, a system for managing one or more rules for use with an application program is provided. The system includes: means for providing an integrated user interface to enable receiving transition information, associating the application program with a ruleset that includes a set of rules, and executing the application program; means for receiving a rule specification that specifies the ruleset, each rule of the ruleset including a set of rule cases that compute values for the rule's output depending on the rule's input, each rule case including one or more input relationships in which a rule case is triggered by a given input record if one or more data values of the given input record meet the one or more input relationships; means for associating the ruleset with a first stage of a plurality of stages; means for receiving, through the integrated user interface, transition information associated with the ruleset; means for, based on the transition information and permitted transitions associated with the ruleset and the first stage, associating the ruleset with a second stage of the plurality of stages, in which the second stage is different from the first stage; means for receiving, through the integrated user interface, a first instruction to associate the application program with the ruleset; means for, in response to the first instruction, associating an application program with the ruleset; means for receiving, through the integrated user interface, a second instruction to execute the application; and means for, in response to the second instruction, executing the application program to process the input records and produce the outputs according to rules of the ruleset.

[0156] In another aspect, in general, a computer-readable medium storing a computer program for managing one or more rules for use with an application program is provided. The computer program includes instructions for causing a computer system to: receive a group of rule specifications, each rule specification specifies a ruleset that includes a set of rules, each rule including a set of rule cases that compute values for the

rule's output depending on the rule's input, each rule case including one or more input relationships in which a rule case is triggered by a given input record if one or more data values of the given input record meet the one or more input relationships; receive a transition specification that specifies a first plurality of states that can be associated with individual rulesets, a second plurality of states that can be associated with a group of rulesets, permissible state transitions associated with individual rulesets, and permissible state transitions associated with the group of rulesets; associate each ruleset with one of the first plurality of states; associate the group of rulesets with one of the second plurality of states; receive first transition information associated with one of the rulesets; based on the first transition information, permitted transitions associated with the ruleset, and the state currently associated with the ruleset, associate the ruleset with another one of the first plurality of states; receive second transition information associated with the group of rulesets; based on the second transition information, permitted transitions associated with the group, and the state currently associated with the group, associate the group of rule specifications with another one of the second plurality of states; and associate an application program with the group of rule specifications, in which the application program is configured to process input records and produce outputs according to the rules of the rulesets specified by the rule specifications.

[0157] In another aspect, in general, a system for managing one or more rules for use with an application program is provided. The system includes: a data storage system storing one or more rule specifications and one or more transition specifications; at least one output device or port for providing a graphical user interface; and a computer system coupled to the data storage system and the at least one output device or port. The computer system includes circuitry to: receive from the data storage system a group of rule specifications, each rule specification specifies a ruleset that includes a set of rules, each rule including a set of rule cases that compute values for the rule's output depending on the rule's input, each rule case including one or more input relationships in which a rule case is triggered by a given input record if one or more data values of the given input record meet the one or more input relationships; receive from the data storage system a transition specification that specifies a first plurality of states that can be associated with individual rulesets, a second plurality of states that can be associated with a group of

rulesets, permissible state transitions associated with individual rulesets, and permissible state transitions associated with the group of rulesets; associate each ruleset with one of the first plurality of states; associate the group of rulesets with one of the second plurality of states; receive first transition information associated with one of the rulesets; based on the first transition information, permitted transitions associated with the ruleset, and the state currently associated with the ruleset, associate the ruleset with another one of the first plurality of states; receive second transition information associated with the group of rulesets; based on the second transition information, permitted transitions associated with the group, and the state currently associated with the group, associate the group of rule specifications with another one of the second plurality of states; and associate an application program with the group of rule specifications, in which the application program is configured to process input records and produce outputs according to the rules of the rulesets specified by the rule specifications.

[0158] In another aspect, in general, a system for managing one or more rules for use with an application program is provided. The system includes: means for receiving a group of rule specifications, each rule specification specifies a ruleset that includes a set of rules, each rule including a set of rule cases that compute values for the rule's output depending on the rule's input, each rule case including one or more input relationships in which a rule case is triggered by a given input record if one or more data values of the given input record meet the one or more input relationships; means for receiving a transition specification that specifies a first plurality of states that can be associated with individual rulesets, a second plurality of states that can be associated with a group of rulesets, permissible state transitions associated with individual rulesets, and permissible state transitions associated with the group of rulesets; means for associating each ruleset with one of the first plurality of states; means for associating the group of rulesets with one of the second plurality of states; means for receiving first transition information associated with one of the rulesets; means for, based on the first transition information, permitted transitions associated with the ruleset, and the state currently associated with the ruleset, associating the ruleset with another one of the first plurality of states; means for receiving second transition information associated with the group of rulesets; means for, based on the second transition information, permitted transitions associated with the

group, and the state currently associated with the group, associating the group of rule specifications with another one of the second plurality of states; and means for associating an application program with the group of rule specifications, in which the application program is configured to process input records and produce outputs according to the rules of the rulesets specified by the rule specifications.

[0159] In another aspect, in general, a computer-readable medium storing a computer program for managing one or more rules for use with an application program is provided. The computer program includes instructions for causing a computer system to: receive a rule specification that specifies a ruleset that includes a set of rules, each rule including a set of rule cases that compute values for the rule's output depending on the rule's input, each rule case including one or more input relationships in which a rule case is triggered by a given input record if one or more data values of the given input record meet the one or more input relationships; receive a transition specification that specifies permissible transitions among a plurality of states associated with the ruleset, and for each transition, one or more users who are authorized to implement the transition; associate the ruleset with a first state of the plurality of states; receive transition information associated with the ruleset; verify that a user who provided the transition information is authorized to provide the transition information; after verifying that the user who provided the transition information is authorized to provide the transition information, based on the transition information and permitted transitions associated with the ruleset and the first state, associate the ruleset with a second state of the plurality of states, in which the second state is different from the first state; and associate an application with the ruleset, in which the application program is configured to process input records and produce outputs according to rules of the ruleset.

[0160] In another aspect, in general, a system for managing one or more rules for use with an application program in a data processing system is provided. The system includes: a data storage system storing one or more rule specifications and one or more transition specifications; at least one output device or port for providing a graphical user interface; and a computer system coupled to the data storage system and the at least one output device or port. The computer system includes circuitry to: receive, from the data storage system, a rule specification that specifies a ruleset that includes a set of rules,

each rule including a set of rule cases that compute values for the rule's output depending on the rule's input, each rule case including one or more input relationships in which a rule case is triggered by a given input record if one or more data values of the given input record meet the one or more input relationships; receive, from the data storage system, a transition specification that specifies permissible transitions among a plurality of states associated with the ruleset, and for each transition, one or more users who are authorized to implement the transition; associate the ruleset with a first state of the plurality of states; receive transition information associated with the ruleset; verify that a user who provided the transition information is authorized to provide the transition information; after verifying that the user who provided the transition information is authorized to provide the transition information, based on the transition information and permitted transitions associated with the ruleset and the first state, associate the ruleset with a second state of the plurality of states, in which the second state is different from the first state; and associate an application with the ruleset, in which the application program is configured to process input records and produce outputs according to rules of the ruleset.

[0161] In another aspect, in general, a system for managing one or more rules for use with an application program is provided. The system includes: means for receiving a rule specification that specifies a ruleset that includes a set of rules, each rule including a set of rule cases that compute values for the rule's output depending on the rule's input, each rule case including one or more input relationships in which a rule case is triggered by a given input record if one or more data values of the given input record meet the one or more input relationships; means for receiving a transition specification that specifies permissible transitions among a plurality of states associated with the ruleset, and for each transition, one or more users who are authorized to implement the transition; means for associating the ruleset with a first state of the plurality of states; means for receiving transition information associated with the ruleset; means for verifying that a user who provided the transition information is authorized to provide the transition information; means for, after verifying that the user who provided the transition information is authorized to provide the transition information, based on the transition information and permitted transitions associated with the ruleset and the first state, associating the ruleset with a second state of the plurality of states, in which the second state is different from

the first state; and means for associating an application with the ruleset, in which the application program is configured to process input records and produce outputs according to rules of the ruleset.

[0162] In another aspect, in general, a method for managing one or more rules for use with an application program in a data processing system, the method including: receiving a rule specification that specifies a ruleset that includes a set of rules, each rule including a set of rule cases that compute values for the rule's output depending on the rule's input, each rule case including one or more input relationships in which a rule case is triggered by a given input record if one or more data values of the given input record meet the one or more input relationships; associating the ruleset with a first stage of a plurality of stages; receiving transition information associated with the ruleset; based on the transition information and permitted transitions associated with the ruleset and the first stage, associating the ruleset with a second stage of the plurality of stages, in which the second stage is different from the first stage; and associating an application program with the ruleset in which the application program is configured to process input records and produce outputs according to rules of the ruleset. The method includes executing the application program to process the input records and produce the outputs according to rules of the ruleset, and providing an integrated user interface to enable the receiving of the transition information, the associating of the application program with the ruleset, and the executing of the application program.

[0163] Aspects can include one or more of the following features.

[0164] Receiving a rule specification can include receiving a spreadsheet having a plurality of rows that correspond to the rule cases.

[0165] The spreadsheet can include trigger columns that correspond to available data values, and the rule cases include sets of triggering criteria that relate to the available data values.

[0166] Each rule case can be configured such that the rule case applies to a given record if one or more data values of that record, for each trigger column in which the rule case has a criterion or criteria, meets the triggering criterion or criteria.

[0167] The spreadsheet can include one or more output columns, each output column can correspond to a potential output variable, and a value of the output variable

can be determined by a cell in the output column, in which the cell can be at the row corresponding to the rule case that applies to the given record.

[0168] The cell can include at least one of an output value that is assigned to the output variable or an expression that is evaluated to generate the output value.

[0169] The plurality of stages can include a plurality of development stages of a development lifecycle.

[0170] The plurality of development stages can include at least one of a develop stage or a test stage.

[0171] The ruleset can include a set of criteria for performing at least one of converting data from one format to another, making determinations about data, or generating new data based on a set of input data.

[0172] The rule specification can include a rule specification table having a plurality of rows, each row defining a rule case and containing a relationship for each of a set of one or more potential inputs.

[0173] At least one of the relationships can include one or more of having a value equal to a threshold, having a value above a threshold, having a value below a threshold, having a value belonging to a set of values, having a value matching a pattern, having a relationship to a value of another input, having a relationship to a value of an output of another rule specification, or having a relationship to a value in a memory.

[0174] At least one of the rows can also contain an output including one or more of a combination of values of the inputs, a pre-determined value, or a value computed from one or more of the values of the inputs.

[0175] The method can include converting the plurality of rule cases in the rule specification to a plurality of logical expressions, compiling the plurality of logical expressions into computer-executable code, and incorporating the computer-executable code into the application program.

[0176] The method can include converting the plurality of rule cases to a transform, and associating the transform with a component in a graph-based computation having data processing components connected by linking elements representing data flows.

[0177] The permissible transitions can include at least one of (i) a ruleset in an earliest stage in a development lifecycle can be promoted to a later stage but cannot be further demoted, (ii) a ruleset that is in neither the earliest stage nor the last stage in the development lifecycle can be promoted to a later stage or demoted to an earliest stage in the development lifecycle, or (iii) a ruleset in a last stage in the development lifecycle can be demoted to an earlier stage but cannot be further promoted.

[0178] The method can include providing the one or more displayed input elements in an editor interface, including at least one of (i) when the ruleset is in the earliest stage of the development lifecycle, a first input element for promoting the ruleset is shown and no input element for demoting the ruleset is shown in the editor interface, (ii) when the ruleset is in neither the earliest nor the last development stage, the first input element for promoting the ruleset and a second input element for demoting the ruleset are shown in the editor interface, or (iii) when the ruleset is in the last development stage, the second input element for demoting the ruleset is shown and no input element for promoting the ruleset is shown in the editor interface.

[0179] Each version of the ruleset can transition through the plurality of development stages.

[0180] The method can include receiving a plurality of rule specifications that specify different versions of a ruleset, and displaying identifiers of the different versions of the ruleset, and the development stages associated with the different versions of the ruleset.

[0181] Displaying identifiers of the different versions of the ruleset and the stages associated with the different versions of the ruleset can include displaying a plurality of major or minor versions of the ruleset and the stages associated with the major or minor versions of the ruleset.

[0182] The method can include storing a copy of the ruleset and naming the copy of the ruleset based on the updated stage associated with the ruleset.

[0183] The method can include receiving a tagging specification that specifies a plurality of possible states of tags associated with the ruleset, one or more input elements associated with each of the states, and a target state associated with each of the input elements.

[0184] In response to a user interaction with a displayed input element, the tag can be associated with the target state that is associated with the displayed input element.

[0185] The method can include using a tag having a pointer to the ruleset to keep track of the state associated with the ruleset, in which the tag transitions among the states based on the permissible transitions.

[0186] The tag can store information about a review history associated with the ruleset.

[0187] The plurality of stages can include a plurality of stages in an approval process of the ruleset.

[0188] In another aspect, in general, a method for managing one or more rules for use with an application program in a data processing system is provided. The method includes: receiving a group of rule specifications, each rule specification specifies a ruleset that includes a set of rules, each rule including a set of rule cases that compute values for the rule's output depending on the rule's input, each rule case including one or more input relationships in which a rule case is triggered by a given input record if one or more data values of the given input record meet the one or more input relationships; receiving a transition specification that specifies a first plurality of states that can be associated with individual rulesets, a second plurality of states that can be associated with a group of ruleset, permissible state transitions associated with individual rulesets, and permissible state transitions associated with the group of rulesets; associating each ruleset with one of the first plurality of states; associating the group of rulesets with one of the second plurality of states; receiving first transition information associated with one of the rulesets; based on the first transition information, permitted transitions associated with the ruleset, and the state currently associated with the ruleset, associating the ruleset with another one of the first plurality of states; receiving second transition information associated with the group of rulesets; based on the second transition information, permitted transitions associated with the group, and the state currently associated with the group, associating the group of rule specifications with another one of the second plurality of states; and associating an application program with the group of rule specifications, in which the application program is configured to process input records

and produce outputs according to the rules of the rulesets specified by the rule specifications.

[0189] Aspects can include one or more of the following features.

[0190] The first plurality of states can include at least one of (i) a plurality of development lifecycle states, or (ii) a plurality of approval process states.

[0191] The second plurality of states can include at least one of (i) a plurality of development lifecycle states, or (ii) a plurality of approval process states.

[0192] The method can include associating each ruleset with a first tag that contains information about the state associated with the ruleset, and updating the first tag when the state associated with the ruleset changes.

[0193] The method can include associating the group with a second tag that contains information about the state associated with the group of rulesets, and updating the second tag when the state associated with the group of rulesets changes.

[0194] In another aspect, in general, a method for managing one or more rules for use with an application program in a data processing system is provided. The method includes: receiving a rule specification that specifies a ruleset that includes a set of rules, each rule including a set of rule cases that compute values for the rule's output depending on the rule's input, each rule case including one or more input relationships in which a rule case is triggered by a given input record if one or more data values of the given input record meet the one or more input relationships; receiving a transition specification that specifies permissible transitions among a plurality of states associated with the ruleset, and for each transition, one or more users who are authorized to implement the transition; associating the ruleset with a first state of the plurality of states; receiving transition information associated with the ruleset; verifying that a user who provided the transition information is authorized to provide the transition information; after verifying that the user who provided the transition information is authorized to provide the transition information, based on the transition information and permitted transitions associated with the ruleset and the first state, associating the ruleset with a second state of the plurality of states, in which the second state is different from the first state; and associating an application with the ruleset, in which the application program is configured to process input records and produce outputs according to rules of the ruleset.

[0195] Aspects can include one or more of the following features.

[0196] The plurality of states can include at least one of (i) a plurality of development lifecycle states, or (ii) a plurality of approval process states.

[0197] The method can include associating each ruleset with a tag that contains information about the state associated with the ruleset, and updating the tag when the state associated with the ruleset changes.

[0198] In another aspect, in general, a computer-readable medium storing a computer program for managing one or more rules for use with an application program in a data processing system is provided. The computer program includes instructions for causing a computer system to: receive a rule specification that specifies a ruleset that includes a set of rules, each rule including a set of rule cases that compute values for the rule's output depending on the rule's input, each rule case including one or more input relationships in which a rule case is triggered by a given input record if one or more data values of the given input record meet the one or more input relationships; associate the ruleset with a first stage of a plurality of stages; receiving transition information associated with the ruleset; based on the transition information and permitted transitions associated with the ruleset and the first stage, associate the ruleset with a second stage of the plurality of stages, in which the second stage is different from the first stage; and associate an application program with the ruleset in which the application program is configured to process input records and produce outputs according to rules of the ruleset. The computer program includes instructions for causing the computer system to execute the application program to process the input records and produce the outputs according to rules of the ruleset, and provide an integrated user interface to enable the receiving of the transition information, the associating of the application program with the ruleset, and the executing of the application program.

[0199] In another aspect, in general, a system for managing one or more rules for use with an application program is provided. The system includes: a data storage system storing one or more rule specifications and one or more transition specifications; at least one output device or port for providing a graphical user interface; and a computer system coupled to the data storage system and the at least one output device or port. The computer system includes circuitry to retrieve from the data storage system a rule

specification that specifies a ruleset that includes a set of rules, each rule including a set of rule cases that compute values for the rule's output depending on the rule's input, each rule case including one or more input relationships in which a rule case is triggered by a given input record if one or more data values of the given input record meet the one or more input relationships; associate the ruleset with a first stage of a plurality of stages; retrieve from the data storage system transition information associated with the ruleset; based on the transition information and permitted transitions associated with the ruleset and the first stage, associate the ruleset with a second stage of the plurality of stages, in which the second stage is different from the first stage; and associate an application program with the ruleset in which the application program is configured to process input records and produce outputs according to rules of the ruleset. The computer system includes circuitry to execute the application program to process input records and produce outputs according to the rules of the ruleset, and provide an integrated user interface to enable the receiving of the transition information, the associating of the application program with the ruleset, and the executing of the application program.

[0200] In another aspect, in general, a system for managing one or more rules for use with an application program is provided. The system includes: means for receiving a rule specification that specifies a ruleset that includes a set of rules, each rule including a set of rule cases that compute values for the rule's output depending on the rule's input, each rule case including one or more input relationships in which a rule case is triggered by a given input record if one or more data values of the given input record meet the one or more input relationships; means for associating the ruleset with a first stage of a plurality of stages; means for receiving transition information associated with the ruleset; means for, based on the transition information and permitted transitions associated with the ruleset and the first stage, associating the ruleset with a second stage of the plurality of stages, in which the second stage is different from the first stage; and means for associating an application program with the ruleset in which the application program is configured to process input records and produce outputs according to rules of the ruleset. The system includes means for executing the application program to process the input records and produce the outputs according to rules of the ruleset, and means for providing an integrated user interface to enable the receiving of the transition information, the

associating of the application program with the ruleset, and the executing of the application program.

[0201] In another aspect, in general, a computer-readable medium storing a computer program for managing one or more rules for use with an application program is provided. The computer program includes instructions for causing a computer system to: receive a group of rule specifications, each rule specification specifies a ruleset that includes a set of rules, each rule including a set of rule cases that compute values for the rule's output depending on the rule's input, each rule case including one or more input relationships in which a rule case is triggered by a given input record if one or more data values of the given input record meet the one or more input relationships; receive a transition specification that specifies a first plurality of states that can be associated with individual rulesets, a second plurality of states that can be associated with a group of ruleset, permissible state transitions associated with individual rulesets, and permissible state transitions associated with the group of rulesets; associate each ruleset with one of the first plurality of states; associate the group of rulesets with one of the second plurality of states; receive first transition information associated with one of the rulesets; based on the first transition information, permitted transitions associated with the ruleset, and the state currently associated with the ruleset, associate the ruleset with another one of the first plurality of states; receive second transition information associated with the group of rulesets; based on the second transition information, permitted transitions associated with the group, and the state currently associated with the group, associate the group of rule specifications with another one of the second plurality of states; and associate an application program with the group of rule specifications, in which the application program is configured to process input records and produce outputs according to the rules of the rulesets specified by the rule specifications.

[0202] In another aspect, in general, a system for managing one or more rules for use with an application program is provided. The system includes: a data storage system storing one or more rule specifications and one or more transition specifications; at least one output device or port for providing a graphical user interface; and a computer system coupled to the data storage system and the at least one output device or port. The computer system includes circuitry to receive from the data storage system a group of

rule specifications, each rule specification specifies a ruleset that includes a set of rules, each rule including a set of rule cases that compute values for the rule's output depending on the rule's input, each rule case including one or more input relationships in which a rule case is triggered by a given input record if one or more data values of the given input record meet the one or more input relationships; receive from the data storage system a transition specification that specifies a first plurality of states that can be associated with individual rulesets, a second plurality of states that can be associated with a group of ruleset, permissible state transitions associated with individual rulesets, and permissible state transitions associated with the group of rulesets; associate each ruleset with one of the first plurality of states; associate the group of rulesets with one of the second plurality of states; receive first transition information associated with one of the rulesets; based on the first transition information, permitted transitions associated with the ruleset, and the state currently associated with the ruleset, associate the ruleset with another one of the first plurality of states; receive second transition information associated with the group of rulesets; based on the second transition information, permitted transitions associated with the group, and the state currently associated with the group, associate the group of rule specifications with another one of the second plurality of states; and associate an application program with the group of rule specifications, in which the application program is configured to process input records and produce outputs according to the rules of the rulesets specified by the rule specifications.

[0203] In another aspect, in general, a system for managing one or more rules for use with an application program is provided. The system includes: means for receiving a group of rule specifications, each rule specification specifies a ruleset that includes a set of rules, each rule including a set of rule cases that compute values for the rule's output depending on the rule's input, each rule case including one or more input relationships in which a rule case is triggered by a given input record if one or more data values of the given input record meet the one or more input relationships; means for receiving a transition specification that specifies a first plurality of states that can be associated with individual rulesets, a second plurality of states that can be associated with a group of ruleset, permissible state transitions associated with individual rulesets, and permissible state transitions associated with the group of rulesets; means for associating each ruleset

with one of the first plurality of states; means for associating the group of rulesets with one of the second plurality of states; means for receiving first transition information associated with one of the rulesets; means for, based on the first transition information, permitted transitions associated with the ruleset, and the state currently associated with the ruleset, associating the ruleset with another one of the first plurality of states; means for receiving second transition information associated with the group of rulesets; means for, based on the second transition information, permitted transitions associated with the group, and the state currently associated with the group, associating the group of rule specifications with another one of the second plurality of states; and means for associating an application program with the group of rule specifications, in which the application program is configured to process input records and produce outputs according to the rules of the rulesets specified by the rule specifications.

[0204] In another aspect, in general, a computer-readable medium storing a computer program for managing one or more rules for use with an application program is provided. The computer program includes instructions for causing a computer system to: receive a rule specification that specifies a ruleset that includes a set of rules, each rule including a set of rule cases that compute values for the rule's output depending on the rule's input, each rule case including one or more input relationships in which a rule case is triggered by a given input record if one or more data values of the given input record meet the one or more input relationships; receive a transition specification that specifies permissible transitions among a plurality of states associated with the ruleset, and for each transition, one or more users who are authorized to implement the transition; associate the ruleset with a first state of the plurality of states; receive transition information associated with the ruleset; verify that a user who provided the transition information is authorized to provide the transition information; after verifying that the user who provided the transition information is authorized to provide the transition information, based on the transition information and permitted transitions associated with the ruleset and the first state, associate the ruleset with a second state of the plurality of states, in which the second state is different from the first state; and associate an application with the ruleset, in which the application program is configured to process input records and produce outputs according to rules of the ruleset.

[0205] In another aspect, in general, a system for managing one or more rules for use with an application program in a data processing system is provided. The system includes: a data storage system storing one or more rule specifications and one or more transition specifications; at least one output device or port for providing a graphical user interface; and a computer system coupled to the data storage system and the at least one output device or port. The computer system includes circuitry to receive, from the data storage system, a rule specification that specifies a ruleset that includes a set of rules, each rule including a set of rule cases that compute values for the rule's output depending on the rule's input, each rule case including one or more input relationships in which a rule case is triggered by a given input record if one or more data values of the given input record meet the one or more input relationships; receive, from the data storage system, a transition specification that specifies permissible transitions among a plurality of states associated with the ruleset, and for each transition, one or more users who are authorized to implement the transition; associate the ruleset with a first state of the plurality of states; receive transition information associated with the ruleset; verify that a user who provided the transition information is authorized to provide the transition information; after verifying that the user who provided the transition information is authorized to provide the transition information, based on the transition information and permitted transitions associated with the ruleset and the first state, associate the ruleset with a second state of the plurality of states, in which the second state is different from the first state; and associate an application with the ruleset, in which the application program is configured to process input records and produce outputs according to rules of the ruleset.

[0206] In another aspect, in general, a system for managing one or more rules for use with an application program is provided. The system includes: means for receiving a rule specification that specifies a ruleset that includes a set of rules, each rule including a set of rule cases that compute values for the rule's output depending on the rule's input, each rule case including one or more input relationships in which a rule case is triggered by a given input record if one or more data values of the given input record meet the one or more input relationships; means for receiving a transition specification that specifies permissible transitions among a plurality of states associated with the ruleset, and for each transition, one or more users who are authorized to implement the transition; means

for associating the ruleset with a first state of the plurality of states; means for receiving transition information associated with the ruleset; means for verifying that a user who provided the transition information is authorized to provide the transition information; means for, after verifying that the user who provided the transition information is authorized to provide the transition information, based on the transition information and permitted transitions associated with the ruleset and the first state, associating the ruleset with a second state of the plurality of states, in which the second state is different from the first state; and means for associating an application with the ruleset, in which the application program is configured to process input records and produce outputs according to rules of the ruleset.

[0207] Aspects can have one or more of the following advantages. An integrated user interface can be provided to enable a user to edit rules, manage development lifecycles and versioning of the rules, and perform simulations of applications using the rules in a unified manner. Two or more of the rule editor, business process management module, source code control module, and simulator can share a common user interface, such that the user can perform functions associated with two or more of the modules without having to switch among two or more different user interfaces. Two or more of the rule editor, business process management module, source code control module, and simulator can share memory, so as to reduce memory requirements, as compared to a conventional system that includes a suite of two or more separate modules for rule editing, business process management, source code control, and simulation. By integrating two or more of the rule editor, business process management module, source code control module, and simulator, the overhead for the modules can be reduced, as compared to a conventional system that includes a suite of two or more separate modules for rule editing, business process management, source code control, and simulation, each module having its own overhead.

[0208] By integrating two or more of the rule editor, business process management module, source code control module, and simulator, we can avoid deadlock or contention, in which different modules process the same ruleset or rule container at the same time, resulting in conflicts. Two or more of the rule editor, business process management module, source code control module, and simulator can share a common log

data object, so that a user can review the common log data object to view information related to two or more of the modules. The integrated system ensures that the timestamps for events associated with the rule editor, business process management module, source code control module, and simulator are consistent. By comparison, a conventional system that includes a suite of two or more separate modules for rule editing, business process management, source code control, and simulation may have inconsistent timestamps because the different modules may use different clocks that are not entirely synchronized. For example, in the convention system, a first event associated with a first module may occur before a second event associated with a second module, but because the clock of the first module is slower than the clock of the second module, the first event may have a timestamp later than the timestamp of the second event. This problem can be avoided by using the integrated system of the present invention in which the rule editor, business process management module, source code control module, and the simulator generate timestamps that are based on a common reference clock.

[0209] In some implementations, the rule editor, business process management module, source code control module, and simulator share a common clock. When the rule editor, business process management module, source code control module, and the simulator are executed on the same computer server, the timestamps generated by the rule editor, business process management module, source code control module, and the simulator are all based on the same system clock of the computer server.

[0210] In some implementations, two or more of the rule editor, business process management module, source code control module, and simulator can execute on two or more computer servers that each has its own system clock. The clocks can be synchronized with one another so that when two events happen on two computer servers at the same time, the timestamps for the two events will be the same.

[0211] In some implementations, two or more of the rule editor, business process management module, source code control module, and simulator can execute on two or more computer servers that each has its own system clock, and the clocks are not synchronized. Each of the two or more computer servers compares its own system clock with a reference clock to determine the difference. Each computer server generates timestamps based on its own system clock, taking into account of the difference between

the system clock and the reference clock, so that the timestamp will be the same as if generated based on the reference clock. For example, the system clock of a first computer server can be used as the reference clock, and the other computer servers can compare their system clocks with the system clock of the first computer server, such that the timestamps generated by the computer servers are all based on the system clock of the first computer. This way, the timestamps generated by the rule editor, business process management module, source code control module, and the simulator are all based on the same reference clock.

[0212] Further, business tags can be used to manage rulesets as they go through various stages of a development lifecycle or review process. It is advantageous to automatically generate and link richly interconnected metadata in a data repository and support custom workflows that use the data repository as their backing datastore. Metadata applications can be built out of the data repository itself.

[0213] Other features and advantages of the description will become apparent from the following description, and from the claims.

DESCRIPTION OF DRAWINGS

[0214] FIG. 1 is a block diagram of an example of an application development system.

[0215] FIG. 2 is a diagram of an example of an integrated user interface.

[0216] FIG. 3 is a diagram of an example of an editor interface.

[0217] FIG. 4 is a block diagram of an example of an application development system.

[0218] FIG. 5 is a diagram of an example of a workspace.

[0219] FIGS. 6A to 6D are diagrams of examples of workspaces and editor interface.

[0220] FIG. 7 is a block diagram of an example of a non-integrated application development system.

[0221] FIG. 8 is a block diagram of an example of a system for executing graph-based computations.

[0222] FIG. 9 illustrates computation graphs.

[0223] FIG. 10 is a screenshot.

- [0224] FIG. 11A illustrates transforming input data into output data in a graph-based computation environment.
- [0225] FIG. 11B illustrates an example of a graph-based computation.
- [0226] FIG. 12 illustrates a screenshot of generating, displaying and linking data objects stored in a business rule environment.
- [0227] FIGS. 13 and 14 illustrate a series of wizard screens.
- [0228] FIGS. 15 and 16 illustrate a series of screens associated an object editor.
- [0229] FIG. 17 is a flowchart of example operations of generating, displaying and linking data objects stored in a data repository.
- [0230] FIG. 18 is an example of a tagging specification.
- [0231] FIGS. 19A and 19B show examples of user interfaces that enable the user of release tags to manage rulesets.
- [0232] FIG. 19C is a diagram showing an example of permissible transitions among states.
- [0233] FIG. 20 shows an example of a tagging specification.
- [0234] FIG. 21 is a diagram of an example of a ruleset approval process.
- [0235] FIG. 22 is an example of a user interface to enable the use of release tags to manage groups of rulesets.
- [0236] FIGS. 23A and 23B show screenshots of examples of activity logs associated with business tags.
- [0237] FIGS. 24 and 25 are examples of screenshots of user interfaces for managing rules.
- [0238] FIG. 26 a flow diagram of an exemplary process for the generation of business rules for a campaign.
- [0239] FIG. 27 is a flow diagram of an exemplary process for implementing the step for data preparation.
- [0240] FIG. 28 is a flow diagram of an exemplary process for authoring and validation of campaign products.
- [0241] FIG. 29 is a flow diagram of an exemplary process for implementing an overlap process.
- [0242] FIG. 30 is an example of spreadsheet-based rule entry.

[0243] FIG. 31 shows examples of test results.

[0244] FIG. 32 shows examples of a rule and corresponding test results.

DESCRIPTION

[0245] In this document, we describe a novel integrated application development system for editing rules, simulating execution of application programs that incorporate the rules, managing and controlling multiple versions of the rules and multiple versions of containers of various combinations of the rules, and managing business processes for authoring, reviewing, risk-checking, and approving the rules. The integrated application development system provides an integrated user interface that enables users to perform the functions mentioned above in a convenient and efficient manner. In some implementations, the integrated application development system provides intuitive user interfaces that enable business people to develop sophisticated application programs that incorporate complex rules for processing records of, e.g., millions or billions of customers without requiring the business people to have deep knowledge about software programming. The integrated application development system enables the business people to efficiently author, revise, test, and optimize the rules so that businesses can launch new marketing or business campaigns frequently and respond quickly to changes in customer preferences, market conditions, and/or business environments.

[0246] Referring to FIG. 1, in some implementations, an application development system 100 facilitates processes for developing business rules, e.g., rules that can be used by airline companies, credit card companies, and/or online market retailers to assist in determining what products or promotional offers to give to customers, and/or which customers are eligible for which rewards. For example, credit card companies can use business rules to assist in detecting potential fraudulent credit card transactions. Banks can use business rules to assist in detecting potential money laundering and fraudulent monetary transactions. Banks can also use business rules to assist in determining credit risks, loan eligibility, and interest rates for customers. Insurance companies can use business rules to assist in determining which customers are eligible for which programs, and what insurance rates to apply to customers.

[0247] In some implementations, the application development system 100 includes a document editor 102, a business process management tool 104, a source code control tool 106, and an application simulator 108. The system 100 includes a repository of data (e.g., enterprise metadata) 110 that can be accessed by the above tools 102, 104, 106, and 108. The system 100 includes one or more common log data objects 116 that store log information provided by the tools 102, 104, 106, and 108. One or more users 90 can access the document editor 102, business process management tool 104, source code control tool 106, application simulator 108, and the one or more common log data objects 116 using an integrated user interface 112. For example, a log data object can include a file, a message, a report, and/or a record. A log data object can include a field in a file, a message, a record, and/or a report. A log data object can be linked to or associated with a particular ruleset or rule container, such that the user 90 can easily access all the logged information associated with the ruleset or rule container.

[0248] In some implementations, a user 90 uses input devices (e.g., keyboard, mouse, microphone) of one or more computers to provide inputs to the integrated user interface 112, and the integrated user interface 112 provides outputs to the user 90 using output devices (e.g., one or more displays and speakers) of the one or more computers. The integrated user interface 112 can show information generated by modules executing on one or more computers. Other input/output devices, and other modes of interactions between the user 90 and the integrated user interface 112, can also be used and will be described in later sections of this document.

[0249] The integrated user interface 112 has user-side components and server-side components. For example, the user 90 uses one or more web browsers to access the tools 102, 104, 106, and 108. The web browser includes widgets that exchange data with the server-side components and render interactive graphical elements that are displayed on one or more web pages shown to the user 90. Multiple users can use web browsers to access the tools 102, 104, 106, and 108 concurrently. The integrated user interface 112 provides an interface for each of the users 90 to enable the user 90 to interact with the tools 102, 104, 106, and 108. For example, a first user 90a can use a web browser to access the document editor 102 to edit a first ruleset, and access the business process management tool 104 to submit the first ruleset for approval. A second user 90b can use a

web browser to access the document editor 102 to review the first ruleset, access the application simulator 108 to perform a simulation of execution of an application program that incorporates the first ruleset, and access the business process management tool 104 to approve the first ruleset for release. A third user 90c (not shown in the figure) can use a web browser to access the source code control tool 106 to add the first ruleset to a rule container. A fourth user 90d (not shown in the figure) can use a web browser to access the source code control tool 106 to make a new version of the ruleset and access the document editor 102 to modify the new version of the ruleset, and so forth. The integrated user interface 112 allows multiple users to collaborate on developing rulesets using an interface that has a consistent look-and-feel when accessing different tools (e.g., 102, 104, 106, and 108).

[0250] Referring to FIG. 2, in some implementations, the integrated user interface 112 includes an editor interface 2720, a business process management interface 2722, a source code control interface 2724, an application simulation interface 2726, and a common log interface 2728. The interfaces 2720, 2722, 2724, 2726, and 2728 can include interactive graphical elements shown on one or more web pages. The editor interface 2720 allows the user to access the document editor 102. For example, the editor interface 2720 can show the rule specification to be edited, and provide tools for the use to edit the rule specification.

[0251] In some implementations, when a first user edits a ruleset stored in the repository 114, the user accesses the source code control tool 106 to check out the ruleset. The user accesses the document editor 102 to edit the ruleset. When the ruleset is in the “open” or “development” state, the document editor 102 stores a copy of the ruleset in a “sandbox” that can only be access by the first user (or other authorized users). The sandbox allows the first user to experiment with various parameters of the ruleset without affecting a production environment or the source code in the repository 114. When the first user finished modifying the ruleset, the first user submits or publishes the ruleset, a copy of the ruleset is stored in the repository 110, and the state of the ruleset is changed to “submitted” or “published.” At this time, other users can view, edit, simulate, approve, and/or perform other processing on the modified ruleset.

[0252] The business process management interface 2722 allows the user to access the business process management tool 104. For example, the interface 2722 can show the current state of the rule being edited in the editor interface 2720, and possible target states that the rule can transition to. The business process management interface 2722 can also show the current and target states of other rules related to the rule being edited in the editor interface 2720. The business process management interface 2722 can show the state or states of the container or containers that include the selected version of the ruleset. The business process management interface allows the user to transition the selected version of the ruleset from one state to another state, such as promoting the selected version of the ruleset from “Development” state to “Testing” state, or demoting the selected version of the ruleset from “Testing” state to “Development” state.

[0253] The source code control interface 2724 allows the user to access the source code control tool 106. For example, the source code control interface 2724 can show a list of the various versions of the rule being edited in the editor interface 2720. The source code control interface 2724 enables the user to manage various versions of the rule containers and rulesets. The user can select a particular version of a particular ruleset using this interface, and the selected version of the ruleset can be edited in the editor interface 2720. The source code control interface 2724 can show information regarding which user checked out the particular version of the ruleset from the repository 114, and a timestamp showing when the particular version of the ruleset was checked out from the repository 114.

[0254] The application simulation interface 2726 allows the user to access the application simulator 108 to simulate execution of an application program that incorporates the selected version of the ruleset that is being edited in the editor interface 2720 to process input data and produce output data. The application simulation interface 2726 also allows the user to retrieve stored output data produced from a prior simulated execution of the application program using the selected version of the ruleset.

[0255] Simulating execution of an application program using test data allows developers of the ruleset to identify and correct potential problems that may occur when the application program incorporating the ruleset is executed using actual data in a live or production environment, such as applying the application program incorporating the

ruleset to production data (e.g., actual customer data) to produce results. In the following, depending on context, the phrase “results of execution of the application program” can mean results from simulation of execution of the application program using test data, or results from execution of the application program using production data.

[0256] The common log interface 2728 allows the user to access the common log data object 116. For example, the common log interface 2728 can show logged information having timestamps within a certain period of time.

[0257] Depending on which functions of the application development system 100 are being accessed by the user, one or more of the editor interface 2720, business process management interface 2722, source code control interface 2724, application simulation interface 2726, and common log interface 2728, or portions of them, can be shown on the display screen to the user. For example, the user can access the functions of two or more of the tools 102, 104, 106, and 108 in a concurrent manner. The integrated user interface 112 may include a combination of two or more of the interfaces 2720, 2722, 2724, 2726, and 2728, or portions of them, that are presented to the user concurrently.

[0258] FIG. 2 shows the editor interface 2720, business process management interface 2722, source code control interface 2724, application simulation interface 2726, and common log interface 2728 as distinct components. However, two or more of the interfaces 2720, 2722, 2724, 2726, and 2728 can be combined as one module. Portions of two or more of the interfaces 2720, 2722, 2724, 2726, and 2728 can be combined as one module. Each of the interfaces 2720, 2722, 2724, 2726, and 2728 can have two or more sub-modules.

[0259] For example, when two users 90a, 90b concurrently use the application development system 100, the integrated user interface 112 can include a first graphical user interface that is shown at the computer display of the first user 90a and a second graphical user interface that is shown at the computer display of the second user 90b. Each of the first and second graphical user interfaces can include one or more of the editor interface 2720, business process management interface 2722, source code control interface 2724, application simulation interface 2726, and common log interface 2728, or portions of them. Similarly, if there are additional users who concurrently access the

application development system 100, the integrated user interface 112 can include additional interfaces that are shown at the computer displays of the additional users.

[0260] Referring to FIG. 3, in some implementations, an editor interface 2500 can include components (e.g., interactive items such as buttons, menus, links, and input boxes) for accessing the business process management tool 104, the source code control tool 106, the application simulator 108, and the common log data object 116. For example, when the user 90 executes the document editor 102 to edit a particular version (e.g., version “xyz”) of a particular ruleset (e.g., ruleset “abc”), the application development system 100 provides the editor interface 2500 and shows the rule specification for the particular version of the particular ruleset. FIG. 3 shows examples of components of the editor interface 2500. The editor interface 2500 can have fewer or more components than those shown in the figure.

[0261] For example, the editor interface 2500 can show the identifier 2508 of the user who is viewing or editing the ruleset. The editor interface 2500 can include a general menu bar 2510 that has general menus, such as, e.g., a menu for file management, a menu for editing tools, and a help menu. The editor interface 2500 can have an editing tool bar 2512 that includes various editing tools, such as, e.g., a copy tool, a paste tool, and a delete tool. The editor interface 2500 can show the rules (e.g., 2504a, 2504b, 2504c) of the ruleset being edited. Each rule (e.g., 2504a) can include one or more rule cases (e.g., 2506a, 2506b, 2506c, 2506d) that compute values for the rule’s output 2508 depending on the rule’s input, and each rule case includes one or more input relationships or criterion (e.g., 2514a, 2514b, 2514c) in which a rule case is triggered by a given input record if one or more data values of the given input record meet the one or more input relationships.

[0262] The editor interface 2500 can show the current state 2526 of the ruleset, and a tool 2528 to allow the user to change the state of the ruleset. For example, the tool 2528 can be a menu that has a list of target states for which the ruleset can transition to. The tool 2528 can include one or more buttons representing the target states that the ruleset can transition to. The editor interface 2500 can have a section 2516 that shows the rule containers that include the ruleset. In this example, version x.x of container AAA, version y.y of container BBB, and version z.z of container CCC include version xyz of

ruleset abc. The editor interface 2500 can have a section 2518 that shows tools and information related to the business process management tool 104. For example, the section 2518 can include tools that can be used to edit a tagging specification that specifies the permissible transitions of the tags among the states. The tagging specification will be described in more detail in later sections of this document.

[0263] The editor interface 2500 can have a section 2520 that shows tools and information related to the source code control tool 106. For example, the section 2520 can include tools that can be used to modify the version of the rule, generate a new version of the rule, or delete an old version of the rule. The editor interface 2500 can include a section 2522 that shows tools and information related to the application simulator 108. For example, the section 2522 can show results of simulating execution of an application program that incorporates the rule to process input records. The editor interface 2500 can include a section 2524 that shows tools and information related to the common log data object 116. For example, the section 2522 can show log information related to the version xyz of the ruleset abc, in which the log information can be written by the tools 102, 104, 106, and/or 108.

[0264] In some implementations, the business process management interface 2722 can include interactive items (e.g., buttons, menus, links, input boxes) that allow the user 90 to access one or more of the other modules, such as the document editor 102, the source code control tool 106, the application simulator 108, and the common log data object 116. The source code control interface 2724 can include interactive items (e.g., buttons, menus, links, input boxes) that allow the user 90 to access one or more of the other modules, such as the document editor 102, the business process management tool 104, the application simulator 108, and the common log data object 116. The application simulation interface 2726 can include interactive items (e.g., buttons, menus, links, input boxes) that allow the user 90 to access one or more of the other modules, such as the document editor 102, the business process management tool 104, the source code control tool 106, and the common log data object 116. The common log interface 2728 can include interactive items (e.g., buttons, menus, links, input boxes) that allow the user 90 to access one or more of the other modules, such as the document

editor 102, the business process management tool 104, the source code control tool 106, and the application simulator 108.

[0265] In this document, an “integrated” user interface refers to a user interface that enables the user to access (i) two or more of the document editor 102, business process management tool 104, source code control tool 106, and application simulator 108, or (ii) functions associated with two or more of the above tools 102, 104, 106, 108, in which the two or more of the above tools 102, 104, 106, 108 share one or more common log data objects 116. The integrated user interface can refer to a graphical user interface that is shown on one or more displays that are controlled by one or more computers. The integrated user interface can show information generated by modules executing on one or more computers.

[0266] Referring back to FIG. 1, each of the document editor 102, business process management tool 104, source code control tool 106, and application simulator 108 can write information to the common log data object 116. For example, the source code control tool 106 can write in the common log data object 116 information specifying that a user associated with a particular user identifier checked out a particular version of a particular ruleset from the repository 110, and a timestamp specifying the timing that the ruleset was checked out. The document editor 102 can write in the common log data object 116 information about the user identifier associated with the user who edited the ruleset, and a timestamp specifying the timing for when the editing was performed.

[0267] The business process management tool 104 can write in the common log data object 116 information specifying that the user associated with the particular user identifier approved the transition of the ruleset from a first state to a second state, and a timestamp specifying the timing for the transition of states. The application simulator 108 can write in the common log data object 116 information specifying that the user associated with the particular user identifier executed an application program using the ruleset to produce outputs, and a timestamp specifying the timing of the execution of the application program incorporating the ruleset. The application simulator 108 can write in the common log data object 116 information specifying that the user associated with the particular user identifier retrieved stored outputs produced from a prior execution of the

application program using the particular version of the particular ruleset, and a timestamp specifying the timing of the retrieval of the stored outputs. The above are just examples of the kind of information that the tools 102, 104, 106, and 108 can write to the common log data object 116. The tools 102, 104, 106, and 108 can also write other types of information to the common log data object 116. Logging information regarding, e.g., who accessed which rules, who edited which rules, who approved which rules, and when the above events occurred can help a company perform audits at a later time, e.g., to comply with legal requirements or identify potential sources of errors.

[0268] Referring to FIG. 4, in some implementations, an application development system 2700 is executed in a distributed computing environment that includes, e.g., a group of computer servers or processors connected through a communication network, e.g., a wide area network (WAN), a local area network (LAN), or a network in a multi-processor system or on a chip. Different tools are executed on different computers or processors. In this example, the document editor 102 is executed on a first computer 2702, the business process management tool 104 is executed on a second computer 2704, the source code control tool 106 is executed on a third computer 2706, and the application simulator 108 is executed on a fourth computer 2708. One or more users 90 interact with the document editor 102, business process management tool 104, source code control tool 106, and the application simulator 108 through the integrated user interface 112. Regardless of whether the tools 102, 104, 106, and 108 are executed on the same or different server computers, the integrated user interface 112 provides a user-friendly interface that hides the complexities of the back-end processes.

[0269] Referring back to FIG. 1, in some implementations, the document editor 102 can be linked to, or integrated with, the business process management tool 104, which manages the states of the rules using, e.g., tags described below. When the document editor 102 modifies a rule, the modification can propagate to the business process management tool 104 and affect the outputs of the tool 104. In some examples, the outputs of the business process management tool 104 are updated automatically in response to changes made by the document editor 102. In some examples, the user has to manually request that the changes made by the document editor 102 be propagated to the business process management tool 104.

[0270] Similarly, the document editor 102 can be linked to, or integrated with, the source code control tool 106 so that, e.g., when a version of the rule is approved for release, that version of the rule is stored by the source code control tool 106, which maintains a master repository 114 of various versions of the rules. The document editor 102 can be linked to, or integrated with, the application simulator 108, which enables simulation of execution of application programs that incorporate the business rules. When the document editor 102 modifies a rule, the modification can propagate to the application simulator 108 and affect the outputs of the simulator 108. In some examples, modification of the rule triggers execution of the application simulator 108, and the outputs of the simulator 108 are updated automatically. In some examples, the user has to manually request execution of the application simulator 108 after the rule has been modified by the document editor 102 in order to obtain updated output from the simulator 108.

[0271] The business process management tool 104 can be linked to, or integrated with, the source code control tool 106, which enables the user to retrieve any version of the rule, determine the state of the rule, and modify the state of the rule using the integrated user interface 112. The business process management tool 104 can be linked to, or integrated with, the application simulator 108, which enables the user to use the integrated user interface 112 to modify the state of a rule after performing simulation of the execution of an application program incorporating the rule. The source code control tool 106 can be linked to, or integrated with, the application simulator 108, which enables the user to use the integrated user interface 112 to incorporate any version of a rule into an application program and perform a simulation of the execution of the application program incorporating the rule.

[0272] The integrated user interface 112 allows the users 90 to access the document editor 102, the business process management tool 104, the source code control tool 106, the application simulator 108, and the common log data object 116 using interfaces that have a consistent look-and-feel, e.g., having similar overall design, color scheme, menu structure, and font characteristics.

[0273] The integrated user interface 112 can simplify tasks for processing business rules. For example, referring to FIG. 5, the application development system 100

can implement a “workspace” 2540 for each user that shows the work items that need to be processed by the user. The workspace 2540 can show a first work queue that includes a first list of rules 2542 in “open” state, meaning that the rules are draft documents that are still being edited. The workspace 2540 can show a second work queue that includes a second list of rules 2544 in “submitted” state, which indicates that the rule is ready for review. The second list includes rules that are ready for review by the user. The workspace 2540 can show a third list of rules 2546 in the “submitted” state, in which the rules have been edited by the user and are ready for review by others. The workspace 2540 can show a fourth list of rules 2548 in “approved” state, in which the rules have been edited by the user and have been approved.

[0274] The following describes an example in which a first user 90a edits a rule named “reward_rule_1” and a second user 90b reviews the rule. Referring to FIG. 6A, for example, when the first user 90a logs in the workspace 2550 associated with the first user 90a, the user sees the rule “reward_rule_1” in the first list of rules 2542 that are in the “open” state. Referring to FIG. 6B, the first user 90a can double-click on the rule “reward_rule_1” and invoke the document editor 102, causing the editor interface 2720 to be presented to the first user 90a and allow the first user 90a to edit the selected rule “reward_rule_1.” The editor interface 2720 can show the current state “Open” 2552 and a “Submitted” button 2554 next to the rule being edited. After the first user 90a finishes editing the rule “reward_rule_1,” the first user 90a can click on the “Submitted” button 2554 to change state of the rule “reward_rule_1” from “open” to “submitted.”

[0275] Referring to FIG. 6C, after the first user 90a changes the state of the rule “reward_rule_1” from “open” to “submitted,” the rule “reward_rule_1” appears in the third list of rules 2546 in the workspace 2550 of the first user 90b, indicating that the rule is waiting to be reviewed by others. At the same time, the rule “reward_rule_1” appears in the second list of rules 2544 in the workspace 2560 of the second user 90b, who is responsible for reviewing the rules submitted by the first user 90a.

[0276] For example, the second user 90b logs in the workspace 2560 associated with the second user 90b and sees that the rule “reward_rule_1” appears in the second list of rules 2544 waiting to be reviewed. Referring to FIG. 6D, the second user 90b can double click on the rule “reward_rule_1” to invoke the document editor 102 in order to

review the rule. The editor interface 2720 can show next to the rule a “Simulate” button 2556, a test file selection tool 2558 to allow selection of a test input file, and an application program selection tool 2560 to allow selection of an application program for simulation. The second user 90b can select a test input file “test_input_1” and an application program “application_1” and click on the “Simulate” button, causing the system 100 to invoke the application simulator 108. The application simulator 108 converts the rule cases in the rule “reward_rule_1” to a plurality of logical expressions, compiles the plurality of logical expressions into computer-executable code, incorporates the computer-executable code into the application program “application_1,” and executes the application program “application_1” to use the “reward_rule_1” to process test input records in the “test_input_1” file and generate output data.

[0277] The editor interface 2720 can show an “Approve” button 2562 and a “Reject” button 2564 next to the rule “reward_rule_1.” If the second user 90b is satisfied that the output data from the simulation are correct, the second user 90b can click on the “Approve” button 2562 to approve the rule “reward_rule_1” and change the state of the rule from “submitted” to “approved.” If the second user 90b thinks the rule has problems, the second user 90b can click on the “Reject” button 2564 to reject the rule “reward_rule_1” and change the state of the rule from “submitted” to “open.”

[0278] Let’s assume that the second user 90b rejects the rule. The editor interface can show a text input box next to the rule “reward_rule_1” to allow the second user 90b to write a comment, which is stored in the common log data object 116. In this example, the second user 90b writes a comment “rewards too high” indicating that the reason for rejecting the rule “reward_rule_1” is because the rewards are too high. After the state of the rule “reward_rule_1” is changed to “open,” the rule “reward_rule_1” appears in the first work queue of the first user 90a.

[0279] The first user 90a sees that the rule “reward_rule_1” has been rejected and returned to the “open” state. The system 100 retrieves the comments related to the rule “reward_rule_1” from the common log data object 116 and displays the comments next to the rule “reward_rule_1.” The first user 90a sees the comment “rewards too high,” invokes the document editor 102 to revise the rule cases to reduce the amount of rewards, and invokes the application simulator 108 to simulate execution of the rule. The first user

90a selects the application program “application_1” and the test file “test_input_1” for simulation. The application simulator 108 converts the revised rule cases in the rule “reward_rule_1” to a plurality of logical expressions, compiles the plurality of logical expressions into computer-executable code, incorporates the computer-executable code into the application program “application_1,” and executes the application program “application_1” to use the revised “reward_rule_1” to process the test input records in the “test_input_1” file and generate output data. The application simulator 118 saves the output data. The first user 90a reviews the output data, confirms that the rewards have been correctly reduced, and changes the state of the rule “reward_rule_1” from “Open” to “Submitted” again.

[0280] The second user 90b sees that the first user 90a re-submitted the rule “reward_rule_1” and invokes the application simulator 108 to check whether the rewards have been correctly reduced. The second user 90b selects the application program “application_1” and the test file “test_input_1” for simulation. The application simulator 108 determines that a previous simulation was performed in which the rule “reward_rule_1” was incorporated into the application program “application_1” and the “test_input_1” file was used for the test input. The application simulator 108 loads the previously generated output data and presents the output data to the second user 90b.

[0281] The second user 90b reviews the previously generated output data and confirms that the rewards have been correctly reduced, determines that there are no other issues with the rule “reward_rule_1,” and clicks on the “Approved” button next to the rule. After the state of the rule “reward_rule_1” is changed from “submitted” to “approved,” the first user 90a sees in the workspace 2550 that the rule “reward_rule_1” appears in the fourth list of rules 2548 that have been approved.

[0282] In the above example of processing of the rule “reward_rule_1” by the first and second users, the document editor 102, the business process management tool 104, and the application simulator 108 were invoked with only a few clicks by the users. The integrated user interface 112 is configured such that when processing the rules, appropriate interactive elements (e.g., buttons, menus, links, input text boxes) are provided to enable the users to conveniently access the document editor 102, the business

process management tool 104, the source code control tool 106, the application simulator 108, and/or the common log data object 116.

[0283] As described in more detail below, the application development system 100 has several improvements over a non-integrated application development system, such as system 2600 shown in FIG. 7. For example, the non-integrated application development system 2600 may be built by bundling separate tools together as an application suite in which the separate tools do not share a common set of user interfaces, and do not share a common log data object. For example, the system 2600 can include a document editor 2610 that is accessed through an editor interface 2602, a business process management tool 2612 that is accessed through a business process management interface 2604, an application simulator 2614 that is accessed through an application simulation interface 2606, and a source code control tool 2616 that is accessed through a source code control interface 2608. Each tool 2610, 2612, 2614, and 2616 in the non-integrated application development system 2600 has its own user interface (2602, 2604, 2606, and 2608, respectively), and the look-and-feel of different interfaces of different tools in the system 2600 can be different.

[0284] For example, when the different tools (e.g., 2610, 2612, 2614, 2616) in the system 2600 store files to non-volatile storage devices, the different tools can store the files to different local file directories associated with the tools. Thus, for example, in system 2600, if the document editor 2610 is executed on a first computer “Comp_1,” the document editor 2610 can save working documents to a local file directory, such as “Comp_1\usr\doc_editor\working_file\”. In system 2600, if the business process management tool 2612 is executed on a second computer “Comp_2,” the business process management tool 2612 can save working documents to a local file directory, such as “Comp_2\usr\bpm\working_file\”. In the system 2600, if the source code control tool 2616 is executed on a third computer “Comp_3,” the source code control tool 2616 can save working documents to a local file directory, such as “Comp_3\usr\source_code_control\working_file\”. In the system 2600, if the application simulator 2618 is executed on a fourth computer “Comp_4,” the application simulator 2618 can save working documents to a local file directory, such as “Comp_4\usr\app_sim\working_file\”.

[0285] In system 2600, the document editor 2610 may not know where the working files of the business process management tool 2612, the source code control tool 2616, and the application simulator 2614 are stored. Similarly, in system 2600, the application simulator 2614 may not know where the working files of the document editor 2600, the business process management tool 2612, and the source code control tool 2616 are stored. Thus, in the system 2600, the user needs to spend time to navigate among different working file directories in order to access the needed files. In the system 2600, the user may have to spend effort setting up correct file paths in order for the different tools to know the locations of the working file directories of the other tools.

[0286] By comparison, the system 100 has an integrated user interface 112 that allows the user to conveniently access the working files of the various tools when processing the rules without the need to remember the details of the working file directories of the various tools 102, 104, 106, 108, and 116.

[0287] For example, the application development system 100 can have other advantages compared to the non-integrated system 2600. For example, the application development system 100 has reduced memory requirements (as compared to the system 2600) because the tools 102, 104, 106, and 108 access the same data stored in memory, reducing the need to duplicate data, which may be needed in the system 2600. As shown in FIG. 7, each of the document editor, business process management tool, application simulator, and source code control tool store copies of files and data locally. The system 100 has a common log data object 116 that can be accessed by the tools 102, 104, 106, and 108, such that information entered into the common log data object 116 by one of the tools 102, 104, 106, and 108 can be accessed by another one of the tools 102, 104, 106, and 108. By reviewing the common log data object 116, the user can determine, e.g., who accessed, edited, approved, and/or simulated a particular version of a particular rule container or ruleset, and when those events happened. This is more convenient than the system 2600 in which each tool writes to its own log file. In system 2600, the user has to access multiple log files in order to obtain information about who accessed, edited, approved, and/or simulated a particular version of a rule container or ruleset, and when those events happened.

[0288] The system 100 can have an advantage in which the integrated user interface 112 enables the user to access functions of two or more of the tools 102, 104, 106, 108, and 116, and the outputs of two or more of the tools 102, 104, 106, 108, and 116 can be provided to the user, e.g., shown side by side. Changes made to data processed by one of the tools 102, 104, 106, and 108 can automatically trigger the update of outputs generated by another one of the tools 102, 104, 106, and 108.

[0289] For example, in FIG. 2, the user can select a test input file and an application program using the application simulation interface 2726. The user can edit a rule using the editor interface 2720. In some implementations, as the user modifies the rule, the application development system 100 automatically invokes the application simulator 108 to convert the rule cases of the rule to a plurality of logical expressions, compile the plurality of logical expressions into computer-executable code, incorporate the computer-executable code into the selected application program, and executes the selected application program to use the rule to process the test input records and generate output data. The application simulation interface 2726 shows the output data to the user. This provides quick feedback to the user, allowing the user to try several parameter values for the rule cases to find optimal values. In some implementations, the application simulator 108 is configured to wait for an explicit instruction from the user before executing the application program.

[0290] In FIG. 1, the document editor 102, the business process management tool 104, the source code control tool 106, and the application simulator 108 are represented as separate modules, this is merely an example. In some examples, two or more of the tools 102, 104, 106, and 108 can be configured as one unit or module. In some examples, the tools 102, 104, 106, and 108 can all be implemented as one integrated module. In some examples, each of the tool 102, 104, 106, and 108 can be implemented by two or more modules. In some examples, several modules may be used in which each module implements a portion of the functions of the tool 102, 104, 106, and/or 108. Each or all or some portions of the tools 102, 104, 106, and 108 may be implemented in software, hardware, or a combination of software and hardware.

[0291] In the following, we first describe a system for executing graph-based computations that can incorporate business rules (FIG. 8), modules that can be used to

generate objects in user interfaces for processing business rules, and examples of usages of these modules (FIGS. 9 and 10), examples of business rules (FIG. 11A), examples of computation graphs that can incorporate business rules (FIG. 11B), examples of user interfaces for processing business rules (FIGS. 12-16), examples of processes for linking objects when processing business rules (FIG. 17), examples of tagging specifications having information about business rules (FIGS. 18 and 20), examples of using tags to manage business rules (FIGS. 19A, 19B, 19C), examples of lifecycles of developing business rules (FIG. 21), examples of release tags having information about business rules (FIG. 22), and examples of activity logs (FIGS. 23A and 23B).

[0292] We then describe the modules of the application development system 100 in greater detail, describe how the application development system 100 can be used to develop business rules, describe how the document editor 102 and the source code control tool 106 can provide an integrated user interface to enable users to easily access and edit different versions of rules and compare different versions of rules (FIGS. 24 and 25). We also describe examples of how the application development system 100 can be used in developing business rules that are used in marketing or business campaigns (FIGS. 26-29).

Graph Based Computations

[0293] FIG. 8 is a schematic block diagram showing the interrelationship of parts of a computing system 120 for developing, executing and managing graph-based computations. A graph-based computation is implemented using a “data flow graph” that is represented by a directed graph, with vertices in the graph representing components (e.g., data files, processes, etc.), and the directed links or “edges” in the graph representing flows of data between components. In some examples, a graphic development environment (GDE) 122 provides a user interface for specifying executable graphs and defining parameters for the graph components. The GDE 122 communicates with a repository 124 and a parallel operating environment 126. Also coupled to the repository 124 and the parallel operating environment 126 are a User Interface module 128 and an external data store 130.

[0294] The repository 124 preferably is a scalable object-oriented database system designed to support the development and execution of graph-based applications and the interchange of metadata between the graph-based applications and other systems (e.g., other operating systems). In this arrangement, the repository 124 is referred to as Enterprise Meta Environment (EME) and is capable of storing one or more types of metadata, including documentation, record formats (e.g., fields and data types of records in a table), transform functions, graphs, jobs, monitoring information, etc. In particular, the EME 124 may store metadata associated with graph-based computations. In this context, each computer program can be considered as being implemented using a computational graph, also called a dataflow graph, or simply a graph. A graph contains one or more components (or vertices) that may be in communication (e.g., by a flow or edge) to allow data flow. In some arrangements, the graphs can execute in the parallel processing environment 126.

[0295] Implemented as a datastore, the EME 124 can track changes in development of graphs, perform statistical and dependence analysis, manage metadata pertaining to the development of graphs, and provide other similar functionality. In general, the storage of graph related metadata allows for data impact analysis, thereby providing a user with a visual sense of how the data is changing in a graph, and the impacts those changes have upon one or more other graphs. Additionally, the EME 124 may permit multiple versions of a graph to be stored, e.g., to track code changes and adjustments during development. The EME 124 also stores data objects that represent actual data to be processed by the computing system 120 including data stored in the external data store 130.

[0296] The parallel operating environment 126 accepts a specification of a data flow graph generated in the GDE 122 and generates executable instructions (e.g., computer instructions) that correspond to the processing logic and resources defined by the graph. The parallel operating environment 126 then typically executes those instructions on a plurality of processors (which need not be homogeneous).

[0297] The User Interface module 128 provides a browser-based view of the contents of the EME 124. Using the User Interface module 128, a user may browse objects, generate new objects, alter existing objects, etc. For example, the User Interface

module 128 may generate browser screens for a user to search for and view objects and information about objects stored in the EME 124.

[0298] Metadata stored in the EME 124 may include data objects for graph-based applications including graph components and other functional objects for building computation graphs. One or more types of metadata may also be stored in the EME 124, for example, “technical” metadata (e.g., application-related business rules, record formats, and execution statistics) and “enterprise” or “business” metadata (e.g., user-defined documentation of job functions, roles, and responsibilities) and other similar types may be stored.

[0299] The information stored in the EME 124 in the form of data objects enables various kinds of analysis regarding applications and the data processed by those applications. For example, a user can obtain answers to questions about data lineage such as where did a given value come from, how was the output value computed, and which applications produce and depend on the data. A developer can understand the consequences of proposed modifications (e.g., If this component changes, what else will be affected? If this source format changes, which applications will be affected?). A user or developer can also obtain answers to questions involving both technical metadata and business metadata, such as which user or developer groups are responsible for producing and using a particular set of data, or what is the latest update to a specific application.

[0300] The EME 124 may also be capable of tracking the status of stored data. For example, multiple versions of one or more data objects may be tracked and stored, thereby making it possible to examine the state of objects over a time period (e.g., over a week, month, or year), and to compare it with the current state of objects. For example, the EME 124 may collect job-tracking, or execution information which enables trend analysis (e.g., How fast is our data growing?) and capacity planning (e.g., How long did that application take to run? How much data did it process, and at what rate? What resources did the application consume? When will we need to add another server?).

[0301] A user can view (and optionally, edit) information contained in and associated with the stored data objects through the User Interface module 108. In one implementation, a user can define one or more “Views,” which provide information in a

browser page about a relevant subset of the data objects. A “View” is defined using a syntax for selecting one or more data objects based on relationships among those objects.

[0302] As shown in FIG. 9, an EME browser interface 202 may be generated by the User interface module 128 in the EME 124. In order to display details about a specific object in the EME 124, the EME browser interface 202 may generate at least a portion of a browser page 206 (e.g., including a HyperText Markup Language (HTML) form) using information contained in the object (e.g., object A). For example, in the screenshot shown in FIG. 10, details about a Document object named “example1” are displayed. Returning to FIG. 9, a call is made as part of this browser page 206 generation to an element renderer 204 that uses information from the object (e.g., object A) and returns a page or portions of a page to be rendered. The page may include various browser form elements, including display elements that display information associated with object A and include input elements for receiving user interaction. In this particular arrangement, a button 208 is provided as an element that is capable of receiving user interactions, however, other input elements may be used to receive any of a variety of user interactions. For example, menus (e.g., drop down menus), checkboxes and other similar types of elements may be utilized. In general, the element renderer 204 refers to reusable logic implemented by one or more software programs or routines that are usually defined for a particular activity that may be performed on any objects of a particular type. As such, the element renderer 204 may determine what actions are performed on an object, and subsequently determine which input element (e.g., button 208) should be rendered on the browser page 206 (displayed in the browser interface 202). The relationship between user actions and corresponding elements (e.g., buttons) is generally provided by the logic that is specified within the element renderer 204. If an element renderer has been configured to support multiple actions or different options, the user may pass parameters to the renderer when specifying its inclusion in the “View.”

[0303] In one arrangement, an element renderer (referred to as “file_upload.py”) can either add a new document to the EME 124 or replace an existing document. Referring to the screenshot shown in FIG. 10, elements to the right of the text “Replace Document” have been generated by the element renderer 204. However, which action to

take may be specified when configuring the element renderer 204 as part of a “View.” In one example, a function call element renderer (e.g., file_upload.py) may by default add a new document to the EME 124. The function call element renderer (e.g., file_upload.py, in which “action=replace”) may also replace the current document. In another example, multiple parameters may be specified for such function calls (e.g., file_update.py, in which “action=add&rule_name=my_doc_oid”). To determine that particular elements are to be displayed (e.g., the button 208) by the element renderer, logic associated with the renderer may use provided parameters (and associated with part of a “View”). The element renderer 204 and its associated programs are also typically stored in the EME 104. Since the programs are not dedicated to a single object, such programs may apply to multiple objects presented in the EME 124.

[0304] In addition to the visible portions of an HTML form that are displayed in the browser (e.g., the file “browse” control button 302 and the “submit” button 304), hidden information within the HTML form may also be used during submission of a file (or another similar operation). In general, when the button 208 is pressed (via user interactions), regardless of scenario or associated functionality (e.g., the “submit” button 304 in FIG. 10), a script 210 is invoked. The script 210 may respond to various types of user requests (e.g., HTTP-specific services) and in turn provides server-side processing in the EME 124. Browser pages (e.g., browser page 206) for objects in the EME 124 may be drawn based on a particular object or set of objects. For example, browser pages that are displayed in response to a user selecting one or more objects stored in the EME 124 may be renderings of information or various characteristics associated with the selected objects. The scripts (e.g., script 210) may be capable of generating custom renderings of pages or elements of a page whose content and/or layout depend on the objects. In some arrangements, the element renderer 204 may render one or more elements of a browser page based on a particular object, and other portions of the browser page may include one or more renderings of elements based on a different object. Various types of input may be used by the script 210, for example, information associated with an object (e.g., object A), user provided information (e.g., entered through the browser page 206) and other types of information. During execution, the script 210 may generate one or more other objects within the EME 124. For example, a

second object B 212 may be produced within the EME 124, however, multiple objects may be produced. Along with producing one or more objects, the script 210 may also establish links (referred to as relations) between objects. For example, a link 214 may be established between the object B 212 and the object A. To establish such links, one or more techniques may be used, for example, commands may be called (e.g., using a command-line interface associated with EME 124). Along with being linked to the object (e.g., object A) associated with its production, the object B 212 can also be linked to other objects as determined by the script 210 associated with the rendered element or elements (e.g., button 304) for object A.

[0305] In this arrangement, EME 124 is an object oriented repository, and includes objects and associated relations. In some examples, a newly generated object may be linked to the object (or multiple objects) corresponding to the browser page being displayed. In other cases, linking to an object (e.g., an existing object) may be determined from information provided to the HTML form. In some examples, objects A and B may be linked in EME 124 (e.g., the link 214), and the link may be displayed as a link on the browser page 206 for navigation. Such techniques provide one or more advantages, for example, the knowledge of which objects to link in the EME 124 can be considered as being determined in an automated manner. Each object within the EME 124 may be stored in a “domain” which contains information about the relations that an object is associated with other objects in the EME 124. When generating one of the links (e.g., link 214) information may be added to the domain that indicates the domain to which the object is associated (e.g., points to). The script 210 may also update the target object’s domain, e.g., to record an inverse relation (e.g., that another object points to it). Links such as link 214 could be shown for some or all of the documents that are stored in the EME 124 that relate to a particular object (e.g., object A or B). In some examples, the script 210 may check permissions associated with a user’s input to determine whether generation of a new object is permitted, and if not the user can be notified.

[0306] In some arrangements, the user may also be directed to a new browser page or a refreshed version of the current page (e.g., page 206) to show the results of the action that has been performed (e.g., effects actions performed by the script 210). In

some examples, multiple users may make modifications to the EME 124 approximately simultaneously, thereby causing the refresh to display the changes. In another example, due to the execution of the script 210, a new object (or multiple new objects) may be generated, some attributes may have been set (e.g., via direct user editing), and the new object (or objects) may have been linked to some existing objects. Based upon the refresh, a detailed view of the newly generated object may be displayed. This view may present to the user information that has been supplied by the user when editing the attributes. For example, the view of an object can take advantage of the links between objects to navigate to other objects in the EME 124 and gather additional information. Navigation may include performing a “walk” action from a starting object to another object linked to the starting object (e.g., either “upstream” or “downstream” based on a direction of the link). Multiple walk actions may be performed over successive links to navigate to one or more related objects. As a result of this walking, additional information can be presented to the user after a refresh.

[0307] The user may also be directed to a new browser page or refreshed version of the current page 206 to be prompted for additional information to be stored in the EME 124. For example, the user may be prompted to supply information associated with the newly generated object (e.g., object B) or a different object or objects related by navigating over relational links. An advantage of the automatic linking of the newly generated object is that the context of the object can help refine and enhance subsequent interaction with the user. For example, a user may provide input to generate a data object representing an employee on a browser page associated with one of multiple departments of an organization (e.g., based on a data object for the department). Upon refreshing the page, a menu (e.g., a dropdown menu) of choices can be provided for selecting a data object e.g., representing the supervisor of the added employee. If the script 210 automatically links the employee data object to the department data object, the list of choices in the menu can be limited to those existing employees whose data objects are linked to the department data object as members of that department.

[0308] The element renderer 204 may be implemented using, for example, a high level computer programming language, such as a python script or Korn shell scripts that are stored in the EME 124. The result of its execution is a fragment of the browser

page 206 that is dependent on a View design and can be inserted into the rest of the content generated by the EME browser interface 202. The element renderer 204 may produce HTML for one element of the “View.” In some examples, View elements are configured sequentially using a command-line interface (not shown). The browser page fragment returned by the element renderer 204 may appear to be as simple as a single button 208 in some cases, but in others the page portion may contain other form elements for data entry. The example illustrated in FIG. 10 has the file “browse” input element 302.

[0309] As mentioned above, when the user clicks to submit the file, the script 210 is typically invoked. The script 210 may be implemented as a python script or in another similar form. In some examples, the actions performed by the script 210 may be performed through the following steps. The script 210 may first generate a new object (e.g., object B 212) within the EME 124 datastore using the command line interface to the EME 124. Referring to the example illustrated in FIG. 10, associated with file uploading, the new object may be a file that the user has uploaded from their desktop computers. In other cases, the new object may initially be entirely empty, but generated as a placeholder to which other metadata or attributes can be attached. Next, the script 210 may link the new object to existing objects within the EME 124 (e.g., object A). Finally, the script 210 may direct the user to a new page or refresh the original page for the user. If the user is directed to a new page, that page is typically a web-based editor for the newly linked and generated object.

Examples of Linking Objects

[0310] Example 1: User Feedback

[0311] In one implementation, a button (e.g., button 208) can be displayed on a browser page (e.g., the browser page 206) for any object in the EME 124. When a user clicks the button, a script (e.g., the script 210) may be configured to execute to generate a new note object. The script may also generate a link for the original object (e.g., object A) pointing to the new note object. The new note object that is generated is presumed to be about the original object. In some examples, the new note may not be stored as part of the original object. As such, to find and display this note at a later time,

a link or relation between the original and new note objects may be generated. Thereby, browser pages or reports may be designed that display all of the notes relating to the object, or some larger collection of objects. For example, this implementation may be able to provide the feedback that users have entered on objects in a project. Further, a user may be redirected to an editor to enter the text content for the new note object.

[0312] In some implementations, the feature described above can be used to generate notes for business rules, or packages of business rules. It can be used to generate notes about why a rule was approved or rejected. When the note is generated, it is automatically linked to the relevant rules or package of rules.

[0313] Example 2: File Upload

[0314] In this example, the element render 204 is used to display a form containing a file browse control (e.g., a user selectable file on a computer system for upload to a server) and a button (e.g., the button 208) on a browser page for an EME object. In some examples, the element renderer 204 generates a browser page element that is associated with a unique identifier, and allows multiple file upload forms to be displayed on the same browser page. After the user browses to select a file (e.g., an Excel spreadsheet) relating to an object that is part of a project residing in the EME 124 (and clicks the button), the script 210 is called to load the file into the EME 124 as a new document object. Operations associated with the script 210 may generate a link from the object corresponding to the original browser page to the new document object. Subsequently, the user may be redirected to a browser page for the original object that shows a list of documents to which the object is linked for navigation and reporting purposes.

[0315] Example 3: Horizontal Versioning of Rulesets in a Business Rule Environment (BRE)

[0316] In one implementation, business rules (e.g., in the form of a subset of metadata) may be stored in the EME 124. Various types and aspects of business rules may be implemented, for example, the rules can be similar to those described in U.S. patent 8,069,129, entitled “Editing and Compiling Business Rules,” U.S. patent

8,386,408, entitled “Editing and Compiling Business Rules,” U.S. patent 8,380,651, entitled “Editing and Compiling Business Rules,” U.S. patent 9,002,770, entitled “Editing and Compiling Business Rules,” and U.S. patent 9,589,232, entitled “Editing and Compiling Business Rules.” The entire contents of the above patents (including US 8,069,129, US 8,386,408, US 8,380,651, US 9,002,770, and US 9,589,232) are herein incorporated by reference. Each business rule can be stored in a separate EME object. A business rule can be expressed as a set of criteria for converting data from one format to another, making determinations about data, or generating new data based on a set of input data. For example, in FIG. 11A, a record 402 in a flight reservation system includes values for fields that indicate the name 404 of a passenger, how many miles 406 the passenger has flown this year, the class 408 of the passenger’s ticket, and the row 410 in which the passenger is seated. A business rule may indicate that such a passenger should be assigned to a particular boarding group (e.g., boarding group 1). While a business rule may be generally understandable to a user, e.g., “first class passengers are in group 1,” such information may need to be translated into one or more other forms to be recognized and processed by a computing device (e.g., a computer system).

[0317] In one particular arrangement, to implement business rules in a graph-based computation environment, a transform 412 may be generated that receives input records, such as record 402, from one or more data sources (e.g., input dataset 400). The term “transform” refers to a mechanism, such as a software module or a data processing entity, that transforms input data to output data, in which the transformation of data may include, e.g., filtering, sorting, cleansing, integration, and extraction of data. A transformation of data can be any type of data manipulation by a data processing entity. The transform may insert this information into an output record, e.g., a record 414, which indicates the passenger’s name 404 and which group the passenger has been assigned (e.g., a group 418). The record 414 may be correspondingly inserted into an output dataset 420. In this example, the datasets (e.g., dataset 420) are illustrated showing an exemplary record, but in general the datasets may include any number and types of records. In some arrangements, input and output datasets may also be referred to as data streams, for example, as the data (from the datasets) may flow into or out of a graph.

[0318] The transforms may be incorporated into graph-based computations having data processing components and that may be connected by linking elements representing data flows. For example, a computation graph 430 illustrated in FIG. 11B includes two data sets 432, 434 (for example, frequent flier data and flight reservation data) that provide input for one or more computations. Format components 436, 438 format the input data, for example, using similar or different types of adjustments to produce commonly formatted data sets that be combined (e.g., by a join component 440) to produce an output data set 442 (that may be used with additional processing operations). A transform may itself be a graph-based computation, such as that in the graph 430, or may be implemented within a component of a graph, such as the individual components 436, 438, and 440 of which the graph 430 is composed.

[0319] In some arrangements, one or more techniques may be used to simplify generation of transforms (e.g., for non-technical users). For example, a tool may be provided for users to input a set of business rules (referred to as a ruleset) in a familiar format that provides a computer system with the one or more transforms to be executed. For example, a ruleset may be considered a set of rules that produce a single transform. A rule may be composed of one or more rule cases that compute different values for the rule's output depending on the input. A rule may also include multiple rules, which may produce values for additional or alternative outputs. A rule set may also contain other rulesets, which may be referred to as "included" rulesets.

[0320] For example, a ruleset may receive as input records of transactions, e.g., credit card transactions, and generate as output alert messages, e.g., messages indicating fraudulent transactions or messages alerting users of over balance. A ruleset may involve one or more mathematical formulas in which one or more fields of the input records are parameters of the formulas.

[0321] For a project in the EME 124, a browser page may list some or all of the rulesets contained in the project. The browser page in this case may display details that are related to a particular object, e.g., a particular project. Projects in EME 124 usually contain various kinds of technical metadata. In this example, the browser page may list all the rulesets that are contained in the project. For each ruleset in the list, a number of

buttons or other types of elements may be displayed, for example, through use of the same element renderer 204.

[0322] Referring to FIG. 12, a series of rows are included in an exemplary list, and each row in the list includes four buttons (e.g., buttons 502, 504, 506 and 508). Each of the displayed buttons are provided by calling an element renderer (e.g., the element renderer 204) with different options. More specifically, buttons 502, 504, 506, 508 respectively represent different actions that can be performed based upon the ruleset: “Promote Ruleset”, “Demote Ruleset”, “New Minor Version”, and “New Major Version”. Promoting a ruleset broadly may be associated with transitioning a particular ruleset from a development phase to a testing phase, and finally to a production phase (e.g., where the ruleset may be executed by a user for data processing). These four actions may be used to promote or demote the rule during the development lifecycle and to possibly generate different versions of the rule (e.g., a new major version, a minor version, etc.). Such a capability allows for previous versions of a ruleset (which may be relatively old) to be accessed (and used) while allowing the ruleset to evolve over a period of time. For example, rather than changing an existing rule or ruleset, a new copy of the rule or ruleset may be generated, adjusted and assigned a new version number.

[0323] In some arrangements, an element renderer (e.g., element renderer 204) may be configured to determine whether a particular button should be displayed based upon one or more permissions associated with a ruleset. For example, a permission may be associated with modifying the ruleset and its current development lifecycle. Referring to the development lifecycles as an example, different phases in that lifecycle may be development (“Dev”), testing (“Test”) and production (“Prod”). So the logic of the element renderer (e.g., element renderer 204) may sequentially step through these three phases (in that particular order). Once the “Prod” phase is reached, the ruleset may not be further promoted. Similarly, in the “Dev” phase the user may not further demote the rule. The element renderer 204 may also pre-calculate confirmation text that can be displayed if the user clicks the button. That is, when the user clicks on one of the buttons, a popup graphic that includes text may be displayed and confirm the action. For example, clicking on the button 502 in the first row may popup a dialog box with text indicative of “Are you sure that you want to Promote Ruleset-1-Dev to Test?” By

providing such a display, the user may confirm the action so the probability of accidental changes (from accidental clicks) is reduced.

[0324] In addition, when a user clicks one of the buttons, the script 210 is called and makes a copy of the corresponding ruleset. Already populated with content, the ruleset is copied and can be used as a starting point for subsequent changes.

Alternatively, the copy may serve as a backup, such that changes may be made to the original. The name for the new copy of the ruleset may be determined based on the original name of the rule and the action. Different names may also be used so that the user can distinguish them later. For example, if the rule name originally ends in “Dev” and it is being promoted, the new rule name might end in “Test”. As mentioned, different activities are being performed on the rule at different stages. When in the “Dev” stage, it can be considered as being modified by a developer. When in the “Test” stage, the rule or ruleset is generally being checked for appropriate behavior and is typically remains unchanged. In the “Prod” stage, the ruleset is being used, for example, for its intended purpose and typically is not changed. Different stages and versions may be necessary because as soon as a ruleset is promoted from “Dev” to “Test”, a developer may start making new changes. Based on the new name, the script 210 may need to update some information that is stored within the ruleset. In some examples, rulesets contain information about the name of the code file (e.g., the transform 412 for a graph in FIG. 11A). In some examples, this code file is named after the ruleset. As such, if a ruleset is generated with a new name, the name for the file generated is also updated. Subsequently, the user may be redirected to the list of rulesets in the project which has been updated to contain the new rule (and be allowed to view the result of the performed action). For example, if any of these actions are performed on a ruleset, the user is taken to a browser page for the new or modified rule.

[0325] Example 4: A Wizard for Object Creation.

[0326] Referring to FIG. 13, additional operations may be implemented for producing one or more objects, for example, based upon user interaction. In one arrangement, information may be gathered from a user, over a series of steps, for producing one or more objects. For example, user interaction (e.g., a user clicking on a

graphical representation of a button) may initiate the execution of a wizard (e.g., series of dialogs) or another type of UI screen sequence that guide a user through a series of steps. While stepping through the sequence, information is collected e.g., from the user that is used for producing one or more objects. The collected information may also be used for other operations, for example, the collected information may represent properties associated with one or more objects and may be used for modifying the object or objects.

[0327] As shown in the figure, a portion of a screen 600, for illustrative purposes, includes fields (e.g., columns) in a dataset (e.g., a table). In this example, a data masking rule (e.g., obfuscation, de-identification, etc.) is defined and may operate upon data associated with the fields. To initiate the process, a graphical representation of a button 602 (labeled “Add Masking Rule”) is selected by a user (through interactions with a pointing device such as a mouse). Upon being selected, a wizard is executed that includes a series of screens. For example, a portion of a first wizard screen 604 may be presented first for collecting user provided information. In this particular example, the wizard screen 604 presents two radio buttons and prompts the user to select the type of masking rule to be produced. Once the selection is made, the user may select a button 606 (labeled “Next”) to transition to the next wizard screen. To terminate the wizard, another button 608 (labeled “Cancel”) may be selected by the user. Next, in this scenario, another wizard screen is presented (as provided by a portion of a screen 610). For this particular screen 610, two radio buttons are presented for the user for selecting whether to create a new rule or to add to an existing rule. Based upon the selections provided by the user, the wizard may determine which particular path (e.g., series of screens) is followed. As such, the wizard may include multiple branched paths (e.g., tree structures) for various user directed scenarios. Similar to the screen 604, the user may select to proceed forward with the wizard by selecting a button 612 (labeled “Next”), move back one screen by selecting a button 614 (labeled “Back”) or halt the wizard by selecting a button 616 (labeled “Cancel”). Upon making the selection, the wizard executes appropriate actions.

[0328] Referring to FIG. 14, another portion of a screen 700, which is provided by the wizard, prompts the user to select from a list of objects to define the masking rule. For the selection, in this example, a list of functions is provided in a menu (e.g., a drop

down menu). Each of the selected functions may be associated with one or more objects that may be stored (e.g., in the repository 124). In this example, one or more masking functions (e.g., Hash Census 1) are associated with the stored objects. Along with using the menu for mask selection, the user can use buttons (similar to the buttons 612, 614 and 616) to control the direction traveled through the wizard. In this example, a button 702 (labeled “Next”) is selected to move forward to a screen 704 that provides a summary of the user selections. For this example, the field named “preferred” (as illustrated in screen 600) is associated the function “Hash Census 1”. Upon reviewing this information and selecting a button 706 (labeled “Finish”), a new object (or multiple objects) is produced and one or more graphical representations may be produced, e.g., to navigate to the object (or objects). For example, briefly returning to the screen 600 of FIG. 14, a link to the newly produced object may replace the graphical button 602 and upon being selected, the link may direct the user to the new object. In one arrangement, a screen 708 may be presented to the user to view information and parameters associated with the produced object. While this wizard example demonstrated the production of an object (or multiple objects), other operations may be provided by similar wizards. For example, object adjustments (e.g., function assignments, etc.), object deletions and other similar operations may be executed.

[0329] Example 5: Edit Renderer

[0330] Referring to FIG. 15, upon one or more objects being created, other operations such as operations associated with exchanging data may be executed. In one arrangement, the functionality of an element rendered may be expanded to include editing capabilities. For example, a script, program or other type of instruction collection may be executed to produce an object that may be associated with another program (e.g., a dedicated editor). Introducing an editor that is associated with an object allows editing capabilities to be expanded (e.g., editing metadata with constraints, launching external edit controls, etc.). Along with enhancing editing capabilities, the editor associated with the object may increase editing flexibility.

[0331] As illustrated in the figure, a screen 800 includes a portion of an edit screen. A field 802 included in the screen (labeled “Key”) has a control that is an edit

renderer. In this example, a value (e.g., a numerical value, text, function, etc.) cannot be directly entered or edited. Rather, a graphical button 804 may be selected by a user to launch an editor that is associated with the object of the Key field. In this arrangement, a screen 806 presents an exemplary editor 808 that is associated with the Key field and allows user selected information to be associated with the field. In this particular arrangement, the editor 808 is implemented with a software development product, which produces rich internet applications (e.g., user interfaces) for multiple platforms, such as Flex (from Adobe Systems of San Jose, California). In this particular example, the editor 808 provides the user with a set of selectable records 810 and selection tools 812 for identifying data for the Key field 802. Once selected, a graphical button 814 may be selected to return (and provide the selected data) to the previous screen 800.

[0332] Referring to FIG. 16, a screen 900 represents an updated version of the screen 800 with the Key field populated with the information from the editor 808 (shown in screen 806). In particular, an updated Key field 902 includes the two selected fields from the editor 808. As such, two sets of data are associated with screen 900, however, in some example additional sets of data may be present. One of the data sets may be provided from objects associated with the screen 900, for example, data selected with a pull down menu 904 and entered into a dialog 906 rendered in the screen 900. The second set of data is provided by the editor 808 included in the screen 806 (and associated with the Key field 802). This data may be passed, for example, to another screen by using one or more techniques. For example, the data sets may be placed into a similar form. In one arrangement, the data may be bundled together in a predetermined form or format to be submitted in a query string (e.g., a HTTP query string). For example, the user may select a graphical button 908 (labeled “Save Changes”) that initiates the bundling of the data into a query string. Upon being bundled into the query string, the string may be provided to one or more screens, objects, or other locations for additional processing and presentation. Various types of editors may also be implemented, for example, editors associated with time and date, graphical editors (e.g., for viewing various aspects of three dimensional objects) and other types of editors may be used.

Operations of Element Renderer

[0333] Referring to FIG. 17, a flowchart 1000 represents some operations implemented by an element renderer (e.g., element renderer 204) and a script (e.g., the script 210) for managing data objects stored in a data repository (i.e., the EME 124) and displaying elements associated with data objects on a graphical user interface (i.e., the user interface module 108).

[0334] As mentioned, richly interconnected metadata objects in the EME 124 may be referenced by relational links. In addition to enforcing and maintaining the logic relationships among multiple data objects, the EME 124 also uses a web-browser-based graphical user interface (e.g., interface 202) to display details of a specific data object. In particular, the information within an object is automatically used to generate 1002 a form (e.g., a web form) that may invoke corresponding operations on one or more data objects. This is typically implemented by the element renderer, which is applicable to almost all data objects that reside in EME 124. Upon receiving 604 user input represented by various user interface interactions regarding the data object, the script 210 may translate user input data in the web form to generate 1006 a new data object within the data repository. Rather than merely copying the input data or representing a temporary variable obtain during intermediate computation, this new data object includes contents that may persist in the EME 124. Subsequently, linking 1008 the new data object with one or more related data objects in the EME 124 may be automatically implemented based on the location where the user operation is invoked. As a result of above mentioned objects generation and linking, the browser interface 202 manages contents for display on the graphical user interface and enables the user to see changes in the EME 124 through the interface.

Business Tagging

[0335] In some implementations, rulesets can be managed by using tags, which are data objects that contain information about the corresponding rulesets. The tags used to manage rulesets can be, e.g., extensible markup language (XML) tags, but can also be other types of data objects. The following describes the use of tags to manage business

rules, but the same principle can also be applied to using tags to manage other types of rules.

[0336] A rule specification or a ruleset may go through several stages of a development lifecycle or approval process. For example, referring to the development lifecycles as an example, different phases or stages in that lifecycle may be development, testing, and production. For example, referring to the approval process as an example, different phases or stages in that process may be open, submitted (or review), and approved (or released). Each stage of the development lifecycle or approval process are also referred to as a “state.”

[0337] As shown in FIG. 12, a user interface may provide buttons that represent different actions that can be performed based upon the ruleset, such as promoting or demoting a ruleset or generating a new major or minor version of the ruleset. Whether a particular button should be displayed can be determined based upon one or more permissions associated with the ruleset. In some implementations, the permissible transitions may be specified by a transition specification, referred to here as a “tagging specification.”

[0338] Referring to FIG. 18, as an example, a tagging specification 1150 specifies permissible transitions for states “Development,” “Testing,” and “Production” associated with a ruleset. The permissible transitions include a transition from “Development” to “Testing,” a transition from “Testing” to “Production,” a transition from “Testing” to “Development,” and a transition from “Production” to “Testing.” We use tags to record the state associated with the ruleset. For example, when we say that the ruleset transitions from the “Development” state to the “Testing” state, the ruleset itself does not change, rather, the tag associated with (or pointing to) the ruleset changes from the “Development” state to the “Testing” state. The tagging specification 1150 includes code 1152 that specifies, when the tag associated with the ruleset is in the “Development” state, a button with a label “Promote” is to be shown in a user interface. Upon user interaction with the “Promote” button, the tag associated with the ruleset is to transition to the “Testing” state. The tagging specification 1150 includes code 1154 that specifies, when the tag associated with the ruleset is in the “Testing” state, a first button with a label “Promote” and a second button with a label “Demote” are to be shown in the user

interface. Upon user interaction with the “Promote” button, the tag associated with the ruleset is to transition to the “Production” state, and upon user interaction with the “Demote” button, the tag associated with the ruleset is to transition to the “Development” state. The tagging specification 1150 includes code 1156 that specifies, when the ruleset is associated with the “Production” state, no further modification is permitted.

[0339] As the tag associated with the ruleset transitions among the various possible states, the tag associated with the ruleset contains information that keeps track of the state associated with the ruleset. Thus, when a ruleset is promoted or demoted to be associated with a different state, the content of the tag is updated to reflect the current state associated with the ruleset.

[0340] In some implementations, two or more types of tags can be used to manage rulesets. For example, one type of business tags can be “object tags” that are used for application configurations and rulesets generated by business users. Business users can generate object tags when they publish application configurations or rulesets. One type of business tags can be “technical tags” that are used for generic graphs, application templates, record format files, and packages that are generated by developers. A technical tag can be used to identify the versions of the generic graphs, application templates, and other objects to be used for a release. One type of business tags can be “release tags” that are used to identify the technical tags and the object tags that define the contents of a release of a ruleset or an application that uses one or more rulesets.

[0341] The following describes the use of release tags for managing various versions or releases of a ruleset. For example, referring to FIG. 19A, a user interface 1100 displays the names of three release tags that point to three versions or releases of a business ruleset: “CEP_release_V1.0” 1102, “CEP_release_V1.1” 1104, and “CEP_release_V1.2” 1106. Each release tag (which points to a particular version of the business ruleset) has an associated state 1108, such as “Approved,” “Open,” or “Submitted.” In this example, the release tags “CEP_release_V1.0,” “CEP_release_V1.1,” and “CEP_release_V1.2” are associated with the states “Approved,” “Open,” and “Open,” respectively. When a user selects one of the release tags, a button that allows the user to modify the state of the release tag is displayed in the user interface 1100. For example, when the user selects the release tag

“CEP_release_V1.2,” which has a state of “Open” 1110, a button “Submit” 1112 is shown.

[0342] In this example, each version or release of the ruleset is associated with a release tag that includes information on the state associated with the ruleset. The tagging specification defines the various states (or stages) of a review process. The tag object contains a link to a specific version of a ruleset and, optionally, the review history. For example, the review history may indicate that on Monday, a particular developer generated a tag for a ruleset; on Tuesday, a manager rejected the tag; and on Wednesday, the developer resubmitted the tag. As each ruleset transitions among the various states, the tag associated with the ruleset is updated to keep track of the state associated with the ruleset. Thus, when a ruleset is promoted or demoted to be associated with a different state, the content of the tag pointing to the ruleset is updated to reflect the current state and, optionally, review history.

[0343] Referring to FIG. 19B, when the user clicks on the button “Submit” 1112 (shown in FIG. 19A), the user interface 1110 updates the state of the tag “CEP_release_V1.2” to “Submitted” 1114, and displays buttons “Approve” 1116 and “Reject” 1118. If the user clicks on the button “Approve” 1116, the tag “CEP_release_V1.2” transitions to “Approved” state, whereas if the user clicks on the button “Reject” 1118, the tag transitions to “Open” state.

[0344] Referring to FIG. 19C, for the example shown in FIGS. 19A and 19B, the permissible transitions among the states include a transition from “Open” 1200 to “Submitted” 1202, a transition from “Submitted” 1202 to “Open” 1200, and a transition from “Submitted” 1202 to “Approved” 1204. Which button(s) (e.g., “Submit” 1112, “Approve” 1116, and/or “Reject” 1118) is/are displayed depends on the current state associated with the ruleset and the permissible transitions specified by the tagging specification.

[0345] For example, referring to FIG. 20, a tagging specification 1120 specifies the permissible transitions of the tags among the states, such as those shown in FIGS. 19A and 19B. The tagging specification 1120 includes code 1122 that specifies, when the tag is in the state “Open” 1124, a button with the label “Submit” 1126 is to be displayed in the user interface. Upon user interaction with the button “Submit” 1126, the

tag is to transition to a target state “Submitted” 1128. The tagging specification 1120 includes code 1130 that specifies, when the tag is in the state “Submitted” 1132, a first button with the label “Approve” 1134 and a second button with the label “Reject” 1136 are to be displayed in the user interface. Upon user interaction with the “Approve” button, the tag is to transition to a target state “Approved” 1138, and upon user interaction with the “Reject” button, the tag is to transition to a target state “Open” 1140. The tagging specification 1120 includes code 1142 that specifies, for the state “Approved” 1144, no button is displayed, indicating that no further transition is performed once the tag is in the “Approved” state, indicating that the ruleset associated with the tag has been finalized and approved.

[0346] In some implementations, the tagging specification optionally includes scripts that are invoked either before or after an action (e.g., an action to approve a tag so that the ruleset becomes associated with the approved state). For example, a “pre-state transition script” refers to a script that is invoked before the action, in which the script may include conditions that can cancel the action. For example, a pre-state transition script can cause an attempt to approve a tag to fail if there are validation errors.

[0347] In some implementations, a company that manages the development, testing, and release of rulesets may require that, as the rulesets transition through various stages in the development lifecycle or approval process, the rulesets need to be handled by particular authorized staff member(s) at each stage. In this example, each stage may correspond to certain role(s) of employee(s) in the company. The tagging specification that defines the states may also specify the staff member(s) who is/are authorized to enable the transitions. In the example shown in FIG. 18, the tagging specification 1150 includes code 1158a, 1158b, and 1158c that specifies which staff member(s) is/are authorized to enable various transitions among stages. For example, code 1158a specifies that the staff member(s) who is/are authorized to enable the transition from “Development” to “Testing” are listed in “rolename-list1”. Code 1158b specifies that the staff member(s) who is/are authorized to enable the transition from “Testing” to “Production” are listed in “rolename-list2”. Code 1158c specifies that the staff member(s) who is/are authorized to enable the transition from “Testing” to “Development” are listed in “rolename-list3”.

[0348] For example, referring to FIG. 21, a ruleset approval process 1160 includes a development phase 1162, a review phase 1164, and a released phase 1166. In the development phase 1162, the ruleset is associated with a first state 1168. In the review phase 1164, the ruleset can be associated with one of two states: a second state 1170 and a third state 1172. In the released phase 1166, the ruleset is associated with a last state 1174. For example, the first state can be an “Open” state, the second state can be a “Review-1” state, the third state can be a “Review-2” state, and the last state can be a “Released” state. A tag associated with the ruleset keeps track of the state associated with the ruleset. The tag transitions through various states, representing the ruleset going through the various stages in the approval process.

[0349] When the tag associated with the ruleset is in the first state 1168, developers 1176 can generate and modify the ruleset, and promote (or submit) the ruleset so that the tag transitions 1178 to the second state 1170. When the tag is in the second state 1170, the ruleset is reviewed by a manager 1180, and the developers 1176 cannot modify the ruleset or change the state of the tag. The tag associated with the ruleset may contain information about the review history for the ruleset. The manager 1180 can either promote (or approve) the ruleset so that the tag transitions 1182 to the third state 1172, or demote (or reject) the ruleset so that the tag transitions 1184 back to the first state 1168. When the tag transitions 1184 back to the first state 1168, the developers 1176 can further modify the ruleset and then promote the ruleset for further review.

[0350] When the tag associated with the ruleset is in the third state 1172, the ruleset is reviewed by a senior manager 1186, and the developers 1176 and the manager 1180 cannot modify the ruleset or change the state of the tag. The senior manager 1186 can either promote (or approve) the ruleset so that the tag transitions 1188 to the last state 1174, or demote (or reject) the ruleset so that the tag transitions 1190 back to the first state 1168. When the tag transitions 1190 back to the first state 1168, the developers 1176 can modify the ruleset and promote the ruleset for further review. When the tag is in the last state 1174, the state of the tag can no longer be modified. For example, when the senior manager 1186 promotes the ruleset and the tag transitions to the last state 1174, it may indicate that the ruleset is ready to be released to a customer and no further modification to the ruleset is allowed.

[0351] In some implementations, two or more business rulesets can be packaged together in a “container,” “collection,” or “release” and be managed as a group. The container can be associated with one of several states, and each of the business rulesets in the container can also be associated with a state that can be either related to, or independent of, the state of the container. The container can be associated with a tag that has information about the state associated with the container.

[0352] For example, a software product may undergo several releases or versions, in which each release of the software product may include a plurality of business rulesets. Different releases may be associated with different combinations of business rulesets. A “release tag” or “group tag” may be used to store information about a group of rulesets related to a release of the software product. Different releases of the software product may be associated with different release tags.

[0353] For example, referring to FIG. 22, a user interface 1190 shows a release tag “CEP_release_V1.1” 1192 associated with a group of rulesets, in which the release tag “CEP_release_V1.1” 1192 includes information about a group of object tags 1194, such as “Investigate_Customer_Activity_V1.1,” “Transaction_Classification_V1.1,” “Review_Phone_Activity_V1.1,” “Review_Posted_Monetary_Activity_V1.1,” “Review_Web_Activity_V1.0,” and “Review_Non_Monetary_Activity_V1.0.” Each of the object tags 1194 points to a corresponding ruleset. The release tag “CEP_release_V1.1” 1192 is associated with a state “Open” 1196. The release tag “CEP_release_V1.1” 1192 includes information about the current state associated with the group of rulesets. The object tags 1194 associated with the rulesets, such as “Investigate_Customer_Activity_V1.1,” “Transaction_Classification_V1.1,” “Review_Phone_activity_V1.1,” “Review_Posted_Monetary_Activity_V1.1,” “Review_Web_Activity_V1.0,” and “Review_Non_Monetary_Activity_V1.0” are associated with the states 1198 “Submitted,” “Submitted,” “Submitted,” “Submitted,” “Approved,” and “Approved,” respectively. Each of the object tags “Investigate_Customer_Activity_V1.1,” “Transaction_Classification_V1.1,” “Review_Phone_activity_V1.1,” “Review_Posted_Monetary_Activity_V1.1,” “Review_Web_Activity_V1.0,” and “Review_Non_Monetary_Activity_V1.0” keeps track of the current state associated with the corresponding ruleset.

[0354] In this example, the state associated with the group release tag “CEP_release_V1.1” 1192 is different from the states 1198 associated with the rulesets in the group. In some implementations, the states associated with the individual rulesets in a group are independent of the state associated with the group. Thus, when the release tag transitions among various states, the states associated with the individual rulesets does not change (unless the user specifically modifies the states of the object tags associated with the rulesets). For example, when the group release tag “CEP_release_V1.1” 1192 transitions from “Open” state to “Submitted” state, from “Submitted” state to “Approved” state, or from “Submitted” state to “Open” state, the states associated with the rulesets 1194 remain the same. Similarly, when one or more object tags associated with one or more of the rulesets 1194 change states, the state of the group release tag 1192 does not change (unless the user specifically modifies the state of the group release tag 1192).

[0355] In some implementations, the states associated with the rulesets in a group are influenced by the state associated with the group. For example, a tagging specification can specify that the group can be associated with one of N states (N being a positive integer greater than or equal to 2), and each ruleset in the group can be associated with one of the N states. The tagging specification can further specify that, when the group release tag is promoted from the $(i-1)$ -th state to the i -th state, the object tags (associated with the rulesets) that have states later than the i -th state are demoted to the i -th state, whereas the object tags that have states earlier than the i -th state remain unchanged.

[0356] For example, the tagging specification can specify that when the group release tag is promoted from the $(i-1)$ -th state to the i -th state, the object tags (associated with the rulesets) that have states earlier than the i -th state are also promoted to the i -th state, whereas the object tags that have states later than the i -th state remain unchanged.

[0357] For example, the tagging specification can specify that when the group release tag is demoted from the $(i+1)$ -th state to the i -th state, the object tags (associated with the rulesets) that have states later than the i -th state are also demoted to the i -th state, whereas the object tags that have states earlier than the i -th state remain unchanged.

[0358] For example, the tagging specification can specify that when the group release tag is demoted from the (i+1)-th state to the i-th state, the object tags (associated with the rulesets) that have states earlier than the i-th state are also promoted to the i-th state, whereas the object tags that have states later than the i-th state remain unchanged.

[0359] The states that are allowed to be associated with the group release tag 1192 can be different from the states that are allowed to be associated with the object tags 1194 (associated with the rulesets). For example, the tagging specification may specify that the group release tag 1192 can be associated with one of “Open,” “Review-1,” “Review-2,” or “Released” states, whereas each object tag 1194 can be associated with one of “Open,” “Submitted,” or “Approved” states.

[0360] The tags, which are data objects, can include various types of information associated with the development or approval of the ruleset. The tags can provide the function of the common log data objects 116 in FIG. 1. The various types of information can include, e.g., what actions (e.g., promote, demote, submit, approve, or reject) were performed on the ruleset, who performed the actions, and when the actions were performed. The tags can also include comments or notes by the people who developed, tested, reviewed, and/or approved the rulesets. This way, at each stage of the development lifecycle or approval process of the ruleset, the person working on the ruleset can have useful information about the work that has previously been done on the ruleset, enabling the person to quickly focus on issues that need to be addressed. When something goes wrong, the user can look at the review history and attempt to identify what went wrong, where the mistake occurred, and take corrective action so as to avoid that type of problem in the future.

[0361] FIGS. 23A and 23B show screenshots 1210 and 1220 of examples of activity logs associated with business tags in which the logs include information about the operations that were performed, the users who performed the operations, comments by the users, and the dates and times that the operations were performed.

Integrated Application Development System

[0362] The following describes the integrated application development system 100 in more detail. Referring back to FIG. 1, the application development system 100

includes the document editor 102, the business process management tool 104, the source code control tool 106, the application simulator 108, the repository of data 110, and the integrated user interface 112. The document editor 102 allows a user to view and edit business rules, and may provide functions including but not limited to those shown in FIGS. 12-16, 19A, 19B, 22, 23A, and 23B.

[0363] The user may use the document editor 102 to modify entries in a spreadsheet having rows that correspond to the rule cases. The document editor 102 can be used to edit the triggering criteria for the rules. The document editor 102 is linked to the business process management tool 104, which manages the states of the rules. For example, the document editor 102 in combination with the integrated user interface 112 enable users to promote or demote a rule during the development lifecycle of the rule, and to approve the rule for release. In response to the user promoting, demoting, or approving a rule, the business process management tool 104 keeps track of the states of the rule using tags described above. Based on the rule specification associated with a rule, the business process management tool 104 determines which transitions of states are allowed, and communicates with the document editor 102 so that the editor 102 can control the integrated user interface to display appropriate items (e.g., “Submit,” “Approve,” “Reject,” “Promote,” and/or “Demote” buttons) in the user interface to enable the user to change the state of the rule.

[0364] The user may develop several rules, and each rule may have several versions. For example, different versions of a rule may have different triggering criteria. FIG. 24 shows an exemplary user interface provided by the document editor 102 showing several rules that a user is developing. In this example, the user is working on rules such as “Cash_Back_Notifications,” “Review_Non_Monetary_Actions,” “Review_Phone_Activity,” “Review_Posted_Monetary_Activity,” “Review_Web_Activity,” and “Transaction_Classification” rules. The user selected “Review_Posted_Monetary_Activity” rule 1242, and the document editor 102 shows information 1244 about various versions of the rule 1242. In this example, versions 1.0 and 1.2 of the rule have been approved, versions 1.1 and 1.3 of the rule have been submitted, and version 1.4 of the rule is open. Version 1.0 of the rule is used in CEP (the

name of the workspace) release version 1.0, version 1.1 of the rule is used in CEP release version 1.1, and version 1.3 of the rule is used in CEP release version 1.2.

[0365] The document editor 102 is linked to the source code control tool 106 so that when a version of the rule is approved for release, that version of the rule is stored by the source code control tool 106, which maintains a master repository of various versions of the rules. Subsequently, when a user intends to edit a particular version of the rule, the user can request permission to change the rule, and the source code control tool 106 will allow the user to check out the rule and edit the rule, and prevent others from editing the rule until the rule is checked back in. The source code control tool 106 also keeps a record of the user who made changes to the rule, and what changes were made to the rule.

[0366] Referring to FIG. 25, in some implementations, the document editor 102 can generate a ruleset difference report 1250 that compares different versions of a rule and provide information on the differences between the versions. In this example, the user compares versions 61 and 128 of the “Review_Posted_Monetary_Activity” rule, and the document editor 102 provides a list of differences 1252 between the versions.

[0367] A feature of the business tagging approach is that the user may pick and choose among different versions of various rules to assemble a ruleset container or package for use with an application, and simulate the execution of the application using the assembled ruleset container. For example, a developer of a credit card company may use the application development system 100 to develop rules for a credit card processing application. The application may use a “credit card approval” rule for deciding whether to approve credit card applications, a “credit card rewards” rule for deciding what rewards to give to credit card customers, and a “fraud detection” rule to evaluate the possibilities that certain credit card transactions are fraudulent. For example, the developer may have developed versions 1.0, 1.1, 1.2, and 1.3 for the “credit card approval” rule, versions 1.0, 1.1, 1.2, 1.3, 1.4, 2.0, and 2.1 for the “credit card rewards” rule, and versions 1.0, 1.1, 2.0, 2.1, 3.0 and 3.1 for the “fraud detection” rule.

[0368] For example, the developer may decide to simulate the execution of the credit card processing application using version 1.1 of the “credit card approval” rule, version 2.1 of the “credit card rewards” rule, and version 3.0 of the “fraud detection” rule. The document editor 102 may provide an interface to allow the user to select

particular rules, and a particular version of each of the selected rules to be packaged together in a container, and associate the container of rules with the credit card processing application.

[0369] Referring back to FIG. 1, the document editor 102 is linked to the application simulator 108. Using the example above, after the user associates the container of rules with the credit card processing application, the document editor 102 provides an option for the user to invoke the application simulator 108 to simulate execution of the credit card processing application. In some implementations, the application is a graph-based computation having data processing components connected by linking elements representing data flows. The application simulator 108 converts the rules in the container into a transform, associates the transform with a component in the graph-based computation, executes the graph-based computation, and displays the output results of the execution.

[0370] The user may examine the outputs of the simulation to determine whether further modifications of the rules are needed. For example, if the user determines that the application failed to detect many fraudulent transactions in the test data set, the user may use the document editor 102 to change the state of and/or edit version 3.0 of the “fraud detection” rule currently associated with the application. Suppose version 3.0 of the “fraud detection” rule is currently in the “Testing” state. The user may demote version 3.0 of the rule to the “Development” state, indicating that the rule needs further revisions. The user may select version 2.1 of the “fraud detection” rule to be included in the container of rules associated with the credit card processing application, and run the simulation again. If the user is satisfied that version 1.1 of the “credit card approval” rule, version 2.1 of the “credit card rewards” rule, and version 2.1 of the “fraud detection” rule function properly together, the user may approve the container for release, for example, as release version 1.0.

[0371] After the user finishes editing version 3.0 of the “fraud detection” rule, the user may promote the state of the rule to “Testing,” include this version of the rule in the container, and run the simulation of the application. If the user determines that version 1.1 of the “credit card approval” rule, version 2.1 of the “credit card rewards” rule, and

version 3.0 of the “fraud detection” rule function properly together, the user may approve the container for release, for example, as release version 1.1.

[0372] An advantage of the application development system 100 is that it provides an integrated user interface that allows the user to edit rules, manage development lifecycles and versioning of the rules, and perform simulations of applications using the rules in a unified manner. Tags and other information associated with the rules are accessible to the various tools in the system 100. For example, when the user uses the document editor 102 to select particular rules and particular versions of the rules for a container, the information about the selected rules can be provided to the application simulator 108, which loads the selected versions of the rules and generate transforms based on those rules.

[0373] For example, the simulation output dataset generated by the application simulator 108 can be provided to the document editor 102, which can display the rules alongside the output dataset. The document editor 102 may provide an interface that allows the user to select one or more values in the output dataset, and the interface may highlight the portions of the rules that were triggered to generate the selected output value(s). If the output value(s) are different from intended, the user can review the highlighted portions of the rules to determine if they need further revisions. The user can modify the rules, run the simulation again, and see if the updated output values are the same as intended, and further revise the rules as necessary.

[0374] The application development system 100 allows a business person who may not have special expertise in computer programming to be able to edit rules, execute application programs using the rules, determine if the outputs are the same as intended, and further edit the rules as necessary. For example, a financial analyst working at a credit card company may have expertise in finance but may not know how to write computer code to implement rules for processing credit card applications, rules for determines credit card rewards, and rules for detecting fraudulent credit card transactions. Using the application development system 100, the financial analyst can revise values and formulas in a spreadsheet (which many financial analysts are familiar with), and the rules specified based on those values and formulas can be transformed into code that are inserted into an application program by the application simulator 108 for execution.

[0375] In some implementations, the application development system 100 facilitates processes for developing business rules used in marketing or business campaigns, such as the processes 1260, 1350, 1360, and 1370 in FIGS. 26-29 below. The processes 1260, 1350, 1360, and 1370 are examples only. Different companies may have different processes and different numbers of team members who work together to develop business rules. For example, one company may use a single maker and a single checker to write and review the business rules and evaluate the risks. For example, another company may have three or more checkers to provide additional levels of review.

[0376] In some implementations, the application development system 100 enables users to conveniently develop business rules for facilitating various applications. For example, a financial institution (e.g., a bank, credit card company, or investment firm) may have a campaign that offers new products or promotions to customers or potential customers. In a first product offering, the financial institution may offer a new home equity loan to qualified customers. In a second product offering, the financial institution may offer new credit cards having various rewards programs. The offers of these campaign products may appear in advertisements shown on ATM machines, advertisements sent by e-mail or text messages, advertisements presented on web sites, or advertisements sent by postal mail. For different customers, the details of the offers may be different.

[0377] For example, when a customer accesses an on-line account at the financial institution using a web browser, the customer may be greeted with a welcome web page showing the various services offered by the financial institution. The web page may also display one or more promotional offers tailored toward the customer. The web page may be generated by a computer server of the financial institution. When the customer logs into the on-line account, the computer server at the financial institution determines what promotions to offer to the customer based on, e.g., financial and other information associated with the customer.

[0378] For example, if the customer has a home mortgage, the computer server may determine that the customer may be interested in a home equity loan. When the computer server generates the web page to be shown to the customer, a promotional offer from the first campaign product may be inserted into the web page. For example, the

offer may provide a new home equity loan with a low interest rate. For example, if the customer does not have a credit card or has a high credit score, the computer server may determine that the customer may be interested in a new credit card. When the computer server generates the web page to be shown to the customer, a promotional offer from the second campaign product may be inserted into the web page, in which the offer may provide information about a new credit card program that offers generous cash-back rewards.

[0379] For a given campaign, different offers may be served to different customers. For example, in the first campaign product, a new home equity loan with a 3.5% interest rate may be offered to a first customer having a high credit score and a large combined savings and checking account, a new home equity loan with a 4% interest rate may be offered to a second customer having a high credit score and a small combined savings and checking account, and a new home equity loan with a 4.5% interest rate may be offered to a third customer having a medium credit score.

[0380] In the second campaign product, a first type of credit card with an annual fee and slightly higher rewards may be offered to a fourth customer having a high income level, a second type of credit card with no annual fee and slightly lower rewards may be offered to a fifth customer having a medium income level, and a third type of credit card having higher reward points on dining and travel may be offered to a sixth customer who already has a travel credit card with lower reward points. The first and second campaign products may not give any promotional offers to customers who have very low credit scores or are a few months behind in credit card payments.

[0381] The financial institution may have hundreds of thousands or millions of customers, so it is useful to partially automate the process of determining what promotions to offer to which customers. In some implementations, business rules similar to those described in U.S. patent 8,069,129, U.S. patent 8,386,408, U.S. patent 8,380,651, U.S. patent 9,002,770, and U.S. patent application 14/677,063 can be used to partially automate the process of evaluating the financial and other data of the customers and determining what promotional offers to provide to the customers. For example, a set of business rules may be authored, such that the rules can be applied to the records of the customers containing financial data to generate outputs including information about

which promotional offers are presented to which customers. The business rules may be incorporated into a graph-based application in the graphic development environment 102, and the parallel operating environment 106 may generate executable instructions (e.g., computer instructions) that correspond to the logic and resources defined by the graph that incorporates the business rules.

[0382] Each campaign may run for a number of days, weeks, or months. The financial institution may launch a new campaigns every couple of days, weeks, or months. Before each new campaign is launched, the staff members of the financial institution may modify the business rules according the new criteria set by the campaign. In some examples, to ensure the quality of the campaign, three groups of staff members may work on modifying and verifying the business rules. For example, the first group of staff members, which we will refer to as the “Makers,” are responsible for activities such as joining data, generating data sets, performing analyses, authoring business rules, running simulations, and profiling data.

[0383] For example, the second group of staff members, which we will refer to as the “Checkers,” are responsible for activities such as validating that the activities performed by the Makers are consistent with process requirements and corporate policies. For example, the Checkers verify that the campaign or risk objectives are clearly identified and documented, validate the business rules individually and in the aggregate, run the simulations and review the logs for the simulations to confirm the results are consistent with objectives and policies, complete checkpoints, and approves the business rules.

[0384] For example, the third group of staff members, which we will refer to as the “Managers,” are responsible for reviewing the checkers’ comments and recommendations, review audit trails for the generation of the business rules by the Makers and the validation of the business rules by the Checkers, assess the adequacy of controls and validations, and approve or reject progression to the next step in the overall process.

[0385] For any given campaign, there may be more than one Maker, more than one Checker, and more than one Manager involved in the authoring, checking, and reviewing of the business rules. It is useful to keep track of who made what changes to

the business rules at what time, and keep records (including, e.g., source codes and notes) of different versions of the business rules as the rules are modified over time. Keeping good records not only facilitates debugging and finding sources of errors, it may also help meet certain regulatory requirements.

[0386] The following describes an exemplary process for generating business rules for a campaign. Referring to FIG. 26, a flow diagram 1260 represents a high level overview of a process for the generation of business rules for a campaign. At step 1262, data necessary for defining the business rules are prepared. For example, a determination may be made as to which data sources are used and whether data sets need to be joined, merged, filtered, aggregated, or further processed. At step 1264, the campaign for each product is defined. Rules are written to define potential offers related to the business objectives of the financial institution. For example, the criteria for determining when and what promotional offers to give to which customers are defined for each product of the campaign. At step 1266, the business rules for the various products of a campaign are overlapped and combined into an integrated package. At step 1268, the business rules for the entire campaign that includes various products are validated and simulated. Step 1268 ensures that the business rules of different products do not conflict with one another. After the business rules for the whole campaign are simulated and validated, the campaign is ready to be launched at step 1270.

[0387] Referring to FIG. 27, in some implementations, a process 1350 is provided for implementing the step for data preparation 1262. The process 1350 may include several steps that are performed by, e.g., a Marketing Team Maker 1280, a Marketing Team Checker 1282, a Risk Team Maker 1284, and a Risk Team Checker 1286. In this example, at step 1288 the Marketing Team Maker 1280 builds a list of data necessary to run the campaign, and selects the marketing data fields 1290 that may be used by the business rules. The Risk Team Maker 1284 at step 1316 builds a list of data that is required to run risk models, and selects risk data fields 1318 that may be used business rules. At step 1292, the system selects data in the marketing data fields 1290 and the risk data fields 1318 for joining. At step 1294, the system performs filtering, aggregation, and other processing on the joined data. At step 1296, the system performs tests on individual units of data. If the data are not acceptable, the process loops back to step 1292 to select

the appropriate data to join. If the data are acceptable, in step 1298 a profile of test data is generated, and the data are submitted for approval. At step 1304, a dataset that includes the prepared data 1300 and the prepared data profile 1302 are submitted for approval at step 1304. At step 1306, a business tag may be associated with the dataset that has been submitted for approval by the Marketing Team Maker 1280. For example, the business tag may include information about the state of the dataset, who (e.g., Marketing Team Maker 1280) approved the dataset, and when the dataset was approved for submission.

[0388] At step 1308, the Marketing Team Checker 1282 reviews the dataset submitted by the Marketing Team Maker 1280 and determines whether the dataset contains necessary and sufficient campaign data. If the dataset does not contain necessary and sufficient campaign data, the Marketing Team Checker 1282 sends the data back to the Marketing Team Maker 1280, who performs step 1292 again to select data for joining. If the dataset contains necessary and sufficient campaign data, the Marketing Team Checker 1282 approves the dataset. At step 1310, a business tag may be associated with the dataset that has been submitted for approval by the Marketing Team Checker 1282. For example, the business tag may include information about the state of the dataset, who (e.g., Marketing Team Checker 1282) approved the dataset, and when the dataset was approved for review by the Risk Team Checker 1286.

[0389] At step 1314, the Risk Team Checker 1286 reviews the dataset to determine whether it contains necessary and sufficient risk data. If the dataset does not contain necessary and sufficient risk data, the Risk Team Checker 1286 sends the dataset back to the Marketing Team Maker 1280, who performs step 1292 again to select data for joining. If the dataset contains necessary and sufficient risk data, the Risk Team Checker 1286 approves the dataset for use by the campaign. At step 1312, a business tag may be associated with the dataset that has been approved by the Risk Team Checker 1286. For example, the business tag may include information about the state of the dataset, who (e.g., Risk Team Checker 1286) approved the dataset, and when the dataset was approved for use by the campaign.

[0390] Referring to FIG. 28, in some implementations, a process 1360 is provided implementing the step for authoring and validation for each campaign product 1264. The process 1360 may include several steps that are performed by, e.g., the Marketing Team

Maker 1280, the Marketing Team Checker 1282, the Risk Team Maker 1284, and the Risk Team Checker 1286. At step 1320, the Marketing Team Maker 1280 identifies the candidate population for the campaign. For example, the candidate population may be customers who are in a certain age range, have annual income in a certain range, and live within a certain geographical region. At step 1322, the Marketing Team Maker 1280 writes the inclusion rules on the business rule engine (BRE). For example, the inclusion rules may include rules determines which customers are eligible for the promotional offerings.

[0391] At step 1338, the Risk Team Maker 1284 defines thresholds (e.g., volume, entity, and maxima) associated with the product offering. For example, one threshold may specify that the total amount of rewards given to all eligible customers for any day or week may not exceed a certain amount. Another threshold may specify that the total number of customers eligible for a product may not exceed a certain value. At step 1340, the Risk Team Maker 1284 writes risk rules for exclusion based on a wide range of records. For example, a risk rule may specify that a customer having a credit score below a certain value, or having a credit card debt above a certain value, is not eligible for the promotional offer.

[0392] At step 1342, the Risk Team Checker 1286 evaluates the risk rules. If the risk rules are not appropriate, the risk rules are sent back to step 1340 so that the Risk Team Maker 1284 can rewrite the risk rules. If the risk rules are appropriate, the Risk Team Checker 1286 approves the risk rules.

[0393] At step 1324, the inclusion rules are merged with the risk rules to generate a complete campaign scope. At this step, the rules for determining which customers are eligible for a campaign product offering and which customers are high risk and not eligible for the offering are merged together into a ruleset. At step 1326, a simulation is performed to apply the full set of rules to test data to generate simulation outcomes 1330. The simulation outcomes 1330 include outputs from the business rules as well as data for validation (also referred to as artifacts) that may be generated based on a processing of the outputs from the business rules.

[0394] The validation data may include the following. A rule coverage report is generated when the simulation is run. The report documents which rules fired, and how

many times. This ensures that each rule has been run at least a minimum number of times to guarantee simulation coverage was sufficient. Rules that don't fire enough (or at all) indicate insufficient simulation coverage (bad validation data) or bad rule design. A data profile of the campaign data is generated. This provides information on statistical distribution of data for each product campaign data set. Questions such as "how many customers qualified?" or "How many customers didn't qualify" or "were there invalid records" can be answered and compared against expected results. Statistical calculations are performed to generate statistical data. This data provides additional information regarding the population that was qualified for a campaign, and population that was not. Statistics such as mean, median, rank, standard deviation, and outliers for items such as qualifying customer income, qualifying customer credit score, non-qualified income, credit score, can be compared against expected values. Random samples of actual data records are selected for manual review. The sampled data provide a "sanity check" that the values found in the output data do not offend marketing or risk criteria. Campaign cost and overall impact estimates may be generated. This may include a series of simulated cost/benefit models run against a set number of assumed take-rates (e.g., 1%, 3%, 5%, 10%, and 25% ...) for calculating values to include expected campaign value (revenue lift) versus immediate cost (e.g., promotions such as statement credits) versus write-offs (with different credit risk cut-off assumptions) and overall net present value. These are compared with input model assumptions.

[0395] At step 1328, a business tag may be associated with the simulation results. For example, the business tag may provide information on who performed the simulation, when the simulation was performed, and what dataset was used to perform the simulation.

[0396] The simulation outcomes 1330 are provided to the Marketing Team Checker 1282 for review at step 1332. At step 1334, if the Marketing Team Checker 1282 determines that the simulation outcomes are not appropriate and not consistent with expected outcomes, the process 1360 loops back to step 1320 so that the Marketing Team Maker 1280 can revise the inclusion rules. At step 1336, the Marketing Team Checker may write notes, such as why the simulation outcomes do not meet the campaign goals, and the notes are stored as part of a business tag associated with the ruleset. If the

simulation outcomes are appropriate and consistent with expected outcomes, the Marketing Team Checker 1282 approves the rules for review by the Risk Team Checker 1286. At step 1344, the Risk Team Checker 1286 reviews the ruleset that includes the inclusion rules authored by the Marketing Team Maker 1280 and the risk rules authored by the Risk Team Maker 1284. If the Risk Team Checker 1286 determines that the risk is too high, the process loops back to step 1320 so that the Marketing Team Maker 1280 can revise the inclusion rules. At step 1346, the Risk Team Checker may write notes, such as why the simulation outcomes indicate too high a risk, and the notes are stored as part of a business tag associated with the ruleset. If the Risk Team Checker 1286 determines that the risk is acceptable, the ruleset is approved, and the process moves to the next step.

[0397] Referring to FIG. 29, in some implementations, a process 1370 is provided to implement the overlap process 1266. The process 1370 may include several steps that are performed by the Marketing Team Maker 1280, the Marketing Team Checker 1282, the Risk Team Checker 1286, and a Cross-Sell Account Executive or Manager 1372. Previously, in the process 1360 shown in FIG. 28, the ruleset for each product is evaluated and approved. The goal of the process 1370 is to verify that the rules for individual products, when combined together, work appropriately and do not produce any conflict. At step 1372, the Marketing Team Maker 1280 writes overlap rules on the business rules engine. For example, the rules may specify the conditions for which two or more products may be offered at the same time. At step 1374, a simulation is performed by applying the full ruleset for all campaign products to test data and generate simulation outcomes 1376. The simulation outcomes 1376 may include outputs from the business rules and data for validation (also referred to as artifacts) that may be generated from a processing of the outputs from the business rules.

[0398] The validation data in the simulation outcomes 1376 (which are for the whole package of business rules across various campaign products) are similar to the validation data in the simulation outcomes 1330 (which are for the business rules of individual campaign products), with the difference being that following overlap the information is provided on an aggregate basis (including distribution of each product) as well as on a per-product basis (similar to what was generated on the per-product basis).

This allows a comparison of how the values shifted from each product or to each product before and after the overlap process is executed.

[0399] At step 1382, a business tag may be associated with the simulation results. For example, the business tag may provide information on who performed the simulation, when the simulation was performed, what datasets were used to perform the simulation, and other notes provided by the Marketing Team Maker 1280 regarding the simulation.

[0400] The simulation outcomes 1376 are provided to the Marketing Team Checker 1282 for review at step 1378. At step 1380, if the Marketing Team Checker 1282 determines that the simulation outcomes are not appropriate and not consistent with expected outcomes, the process 1370 loops back to step 1372 so that the Marketing Team Maker 1280 can revise the overlap rules. Perhaps certain products may not overlap certain other products during certain periods, and thus the overlap rules need to be adjusted. At step 1384, the Marketing Team Checker may write notes, such as why the simulation outcomes do not meet the campaign goals, and the notes are stored as part of a business tag associated with the overall ruleset. If the simulation outcomes are appropriate and consistent with expected outcomes, the Marketing Team Checker 1282 approves the rules for review by the Risk Team Checker 1286. At step 1386, the Risk Team Checker 1286 reviews the entire ruleset that includes the overlap rules authored by the Marketing Team Maker 1280. At step 1388, if the Risk Team Checker 1286 determines that the risk is too high, the process loops back to step 1372. If the simulation outcomes are appropriate and the risk is acceptable, the Risk Team Checker 1286 approves the rules for review by the Manager 1372.

[0401] At step 1390, the Manager 1372 reviews the whole ruleset that includes the overlap rules authored by the Marketing Team Maker 1280. If the Manager 1372 does not approve the overall package of rules, the process loops back to step 1372 so that the Marketing Team Maker 1280 can revise the overlap rules. At step 1394, the Manager 1372 may write notes, such as why the overall package does not meet the campaign goals or the risk is too high, and the notes are stored as part of a business tag associated with the package of rulesets. If the Manager 1372 approves the overall package of rulesets, the campaign goes live in step 1394.

[0402] The processes 1350, 1360, and 1370 show that authoring, checking, and approving business rules for a campaign is a complicated process that involves multiple staff members of the financial institution. The rules may pass back and forth among the Marketing Team Maker 1280, the Marketing Team Checker 1282, the Risk Team Maker 1284, the Risk Team Checker 1286, and the Manager 1372 several times. At each iteration, each staff member needs to know the exact status of each rule, such as whether the rule has been submitted by the Marketing Team Maker 1280 or the Risk Team Maker 1284, whether the rule has been approved by the Marketing Team Checker 1282 or the Risk Team Checker 1286, and whether the rule has been approved by the Manager 1372. During each iteration, notes regarding the rules may be written and stored in business tags.

[0403] Several versions of the business rules can be stored. For example, a first version of the business rule for a first campaign product may conflict with the rules for a second campaign product, so the Marketing Team Maker 1280 may decide to store the first version of the business rule and write a second version of the business rule for the first campaign product that does not conflict with the rules for the second campaign product. The first version of the business rule may have certain parameter values that can be useful for later campaigns. By the time the package of rulesets for the whole campaign is approved by the Manager 1372, each campaign product may be associated with several versions of business rules, in which the different versions of the rules may use slightly different parameter values that may be useful under different circumstances.

[0404] One of the improvements of the application development system 100 over conventional systems is that the users can access functions provided by the document editor 102, the business process management tool 104, the source code control tool 106, and the application simulator 108 through the integrated user interface 112 without the need to switch between different user interfaces, and without the need to manually copy information from one application to the other. For example, the system 2600 of FIG. 7 has a business process management tool 2612 and a source code control tool 2616 that are “bundled” together in a software suite, but the business process management tool 2612 and the source code control tool 2616 can have different user interfaces and do not share metadata. Thus, for example, using the system 2600, the user can save the

source code for several versions of a business rule using the source code control tool 2616, and use the business process management tool 2612 to manage the transition of the rules among different states. The source code control tool 2616 keeps track of which user made what changes to a particular piece of code over time. However, the information about states that are generated by the business process management tool 2612 is not automatically carried over to the source code control tool 2616. Thus, when the source code of the various versions of the business rules is saved, the information about the states generated by the business process management tool 2612 is not saved along with the source code by the source code control tool 2616. Thus, using the system 2600, it can be difficult for the user to retrieve an earlier version of a business rule and determine the state of the rule, such as whether it was “submitted” or “approved,” or to determine who authorized the transition of the rule from one state to another. Using the system 2600, the user may have to compare the logs of the source code control tool 2616 and the business process management tool 2612 to match the notes for a particular version of a particular rule in the business process management tool 2612 with the notes of the same particular version of the particular rule in the source code control tool 2616, and try to manually deduce the required information based on the notes.

[0405] One of the improvements of the application development system 100 over the non-integrated application development system 2600 is that the application development system 100 can provide functions that cannot be achieved, or would be very difficult to achieve, using the system 2600. For example, an application may evolve over time such that a series of releases of the application are issued over time. Each release of the application may use a package of particular versions of particular rulesets, in which each particular version of the ruleset has a particular state, and the transition to the particular state was approved by a particular user. Different releases of the application may use different packages that includes different versions of the rules. Suppose after several releases of the application have been issued, it is determined that an earlier release of the application did not function properly. In order to determine the cause of the problem, it may be useful to retrieve the particular versions of the particular rulesets used by the particular release of the application that caused the problem, review the logs of the particular versions of the particular rulesets, and determine what changes were made to

the rulesets prior to the occurrence of the problem, who made those changes, who approved those changes, and why those changes were made. It may be useful to review the logs of the overall package of rulesets, and determine what changes were made to the package prior to the occurrence of the problem, who made those changes, who approved those changes, and why those changes were made. It may also be useful to execute a first simulation of the earlier release of the application using the particular versions of the particular rulesets used by the earlier release of the application, execute a second simulation of the application using different versions of the particular rulesets, and/or execute a third simulation of the application by changing some parameters of the rulesets. Using the non-integrated application development system 2600, it may be difficult to perform all of the above functions because the source code control tool 2616 does not store the metadata generated by the business process management tool 2612 along with the source code of the rules that are needed to implement the functions mentioned above. Using the non-integrated application development system 2600, while users can manually copy some information from the business process management tool 2612 to the notes section of the source code control tool 2616, because many different users may write, edit, review, and approve the rules, and different people may have different ways of writing notes, the information manually carried over from the business process management tool 2612 to the source code control tool 2616 may not be consistent across different users, making the information difficult to review. By contrast, in the application development system 100, the metadata stored in the repository 110 have standard formats and can be access by all the tools 102, 104, 106, and 108, enabling new functions that cannot be achieved using the non-integrated application development system 2600.

[0406] One of the improvements of the application development system 100 over the non-integrated application development system 2600 is that outputs of applying business rules to input records can be automatically updated and displayed after the business rules are edited by the document editor 102. For example, a rule can be entered in a spreadsheet format, as shown in FIG. 30. Trigger columns 2202, 2204, 2206, 2208 in spreadsheet 2200 correspond to available data values, and rows 2210a-h correspond to rule cases, i.e., sets of criteria that relate the available data values. A rule case 2210n applies to a given record (e.g., 402 in FIG. 11A) if the data values of that record, for each

trigger column in which the rule case has criteria, meets the triggering criteria. If a rule case 2210 n applies, output is generated based on one or more output columns 2212. A rule case that has all of its input relationships satisfied may be referred to as “triggered.” Each output column 2212 corresponds to a potential output variable, and the value in the corresponding cell of the applicable row 2210 n determines the output, if any, for that variable. The cell could contain a value that is assigned to the variable or it could contain an expression that must be evaluated to generate the output value, as discussed below. There may be more than one output column, though only one is shown in FIG. 30.

[0407] FIG. 31 shows an example of an input test table 2414, an output test table 2416, and a results table 2410 that show the outputs for each of several test cases. In table 2410, the outputs for each test case is displayed in a separate row 2412 i with columns 2413a-d showing the test case and the outputs calculated for each case. When one row is selected (in this example, the second case or row 2412b is selected and shown by a dashed line), the input values 2414a-2414g from the corresponding test case are shown in an input table 2414. The outputs 2416a-2416c for that test case are shown in an output table 2416. The outputs 2416 i ($i = a, b, c$) shown in the table 2416 could be the expected outputs from the test case, the actual outputs computed, or both. Various interface cues can be used to inform the user about the status of the test. For example, icons 2419 indicate that the output in rows 2412b and 2412d for test cases 2 and 4 have changed since the last time the rule set was tested. If there is no “default” rule to catch combinations of input values that don’t match any other rule, test cases having such inputs produce a null output 2418. If an input record is rejected with a processing error, such as division by zero, an error may be indicated in the same way. If more than one rule case is triggered by a test case, the number of triggered rule cases may be indicated (2422). By interacting with one of the displayed outputs, for example, by clicking on it, the user can be taken to a screen that details how that output was determined.

[0408] In some implementations, because the application development system 100 integrates the document editor 102 and the application simulator 108, the system 100 can provide a user interface in which the ruleset is shown beside the test inputs and test outputs, as illustrated in FIG. 32. For example, when the user changes the trigger columns 2202, 2204, 2206, 2208 in spreadsheet 2200, the system 100 can

automatically update the test output values 2416i (i = a, b, c) shown in the output test table 2416. This allows the user to experiment with different rules and quickly receive feedback, allowing the user to optimize the rule in an efficient manner. In the example above, when the user makes a change to the ruleset in the spreadsheet 2200 using the document editor 102, the system 100 automatically invokes the application simulator 108 to convert the ruleset into a transform, associate the transform with a component in the graph-based computation, execute the graph-based computation, and use the output results of the execution to update the output test table 2416. The system 100 can give the user the option of turning on or off the simulator 108. The system 100 can also give the user to option of either automatically updating the output when a change is made to the rules, or updating the output only when the user makes such a request.

[0409] The integrated user interface 112 (FIG. 1) enables the user to access the functionalities of the document editor 102, business process management tool 104, the source code control tool 106, the application simulator 108, and the common log data object 116 can be, e.g., one or more graphical user interfaces that is/are shown on one or more displays that is/are controlled by one or more computer servers. The integrated user interface 112 can show information generated by modules executing on one or more computer servers.

[0410] In the example of FIG. 2, the integrated user interface 112 includes the editor interface 2720, the business process management interface 2722, the source code control interface 2724, the application simulation interface 2726, and the common log interface 2728 that are shown on a computer display. Two or more of the interfaces 2720, 2722, 2724, 2726, 2728 can be shown to the user at the same time. The interfaces 2720, 2722, 2724, 2726, 2728 can also be shown to the user one at a time.

[0411] The integrated user interface 112 is configured to allow the user to conveniently use one of the tools and access another one of the tools. For example, the source code control interface 2724 can show a list of rules, some of the rules each having multiple versions. The user can right click on a particular version of a particular rule, e.g., version 1.0 of “Rule ABC,” and a menu will appear to enable the user to, e.g., edit the selected rule, simulate execution of the selected rule, change the state of the selected rule, or view log information related to the selected rule.

[0412] If the user selects “Edit rule” item in the menu, the document editor 102 is invoked and the editor interface 2720 is shown on the display to enable the user to edit the selected rule, which in this example is version 1.0 of “Rule ABC.” When the editor interface 2720 is shown on the display, the editor interface 2720 includes user selectable elements (e.g., buttons, menus, links) that enables the user to access the business process management tool 104, the source code control tool 106, the application simulator 108, and the common log data object 116.

[0413] For example, while the editor interface 2720 is shown on the display, the user may decide to change the state of the rule by invoking the business process management tool 104, and the business process management interface 2722 is shown on the display. The business process management interface 2722 includes user selectable elements (e.g., buttons, menus, links) that enables the user to access the document editor 102, the source code control tool 106, the application simulator 108, and the common log data object 116.

[0414] For example, while the business process management interface 2722 is shown on the display, the user may decide to simulate execution of the rule by invoking the application simulator 108, and the application simulation interface 2726 is shown on the display. The application simulation interface 2726 includes user selectable elements (e.g., buttons, menus, links) that enables the user to access the document editor 102, the business process management tool 104, the source code control tool 106, and the common log data object 116.

[0415] For example, when the application simulation interface 2726 is shown on the display, the user may decide to view the log information related to the rule by accessing the common log data object 116, and the common log interface 2728 is shown on the display. The common log interface 2728 includes user selectable elements (e.g., buttons, menus, links) that enables the user to access the document editor 102, the business process management tool 104, the source code control tool 106, and the simulator 108.

[0416] In some implementations, any combination of the editor interface 2720, the business process management interface 2722, the source code control interface 2724,

the application simulation interface 2726, and the common log interface 2728 can be shown to the user simultaneously.

[0417] The object generating, linking, business tagging, rule editing, rule simulation, business rule process management, and rule source code management approaches described above can be implemented using software for execution on a computer. For instance, the software forms procedures in one or more computer programs that execute on one or more programmed or programmable computer systems (which may be of various architectures such as distributed, client/server, or grid) each including at least one processor, at least one data storage system (including volatile and non-volatile memory and/or storage elements), at least one input device or port, and at least one output device or port. The software may form one or more modules of a larger program, for example, that provides other services related to the design and configuration of computation graphs. The nodes and elements of the graph can be implemented as data structures stored in a computer readable medium or other organized data conforming to a data model stored in a data repository.

[0418] The software may be provided on a medium, such as a CD-ROM, readable by a general or special purpose programmable computer or delivered (encoded in a propagated signal) over a network to the computer where it is executed. All of the functions may be performed on a special purpose computer, or using special-purpose hardware, such as coprocessors. The software may be implemented in a distributed manner in which different parts of the computation specified by the software are performed by different computers. Each such computer program is preferably stored on or downloaded to a storage media or device (e.g., solid state memory or media, or magnetic or optical media) readable by a general or special purpose programmable computer, for configuring and operating the computer when the storage media or device is read by the computer system to perform the procedures described herein. The inventive system may also be considered to be implemented as a computer-readable storage medium, configured with a computer program, where the storage medium so configured causes a computer system to operate in a specific and predefined manner to perform the functions described herein.

[0419] A number of embodiments of the description have been described. Nevertheless, it will be understood that various modifications may be made without departing from the spirit and scope of the description. For example, some of the steps described above may be order independent, and thus can be performed in an order different from that described. It is to be understood that the foregoing description is intended to illustrate and not to limit the scope of the description, which is defined by the scope of the appended claims.

[0420] In some implementations, the integrated user interface 112 can provide one or more formats of information, e.g., text, graphics, images, audio, video, and/or haptic feedback to the users 90 using, e.g., display devices, audio output devices, and/or haptic devices. The information can be presented to the users 90 using augmented or virtual reality systems. The integrated user interface 112 can accept user inputs from one or more input devices, such as keyboards, mice, trackballs, joysticks, touchpads, touch screens, microphones, cameras, eye tracking devices, gesture detection devices, and/or motion sensors.

[0421] For example, the user 90 can edit a ruleset shown on a computer display screen. The user 90 can wear augmented reality eyeglasses, and output data derived from simulations of executing an application that uses the ruleset to process input records can be projected to appear above the computer display screen. Information related to source code control can be projected to appear to the left of the computer display screen, and information related to business process management can be projected to appear to the right of the computer display. When the user 90 edits the ruleset, the revised output data is projected above the computer display screen. When the user 90 checks out a ruleset by interacting with the elements related to source code control projected to the left of the computer display screen, the selected ruleset is shown on the computer display screen, the relevant state information is projected to the right of the computer display screen, and simulation outputs are projected above the computer display screen.

[0422] Other embodiments are within the scope of the following claims.

CLAIMS

1. A computer-implemented method for implementing one or more rules for use with an application that includes a process portion, a version control portion, and an editor portion in a data processing system, the method including:

storing multiple versions of a rule in a system storage, in which the rule includes one or more rule cases that specify one or more values for the rule's output depending on the rule's input, and each rule case includes one or more input relationships in which the rule case is triggered by a given input record if one or more data values of the given input record meet the one or more input relationships;

wherein each version of the rule is configured to be associated with a plurality of states including a first state and a second, different state of the plurality of states;

storing one or more records that are each shared by the process portion, the version control portion, and the editor portion;

receiving instructions through a user interface that integrates, through the shared one or more records, the process portion, the version control portion, and the editor portion, with the process portion being presented in a first user interface portion, the version control portion being presented in a second user interface portion, and the editor portion being presented in a third user interface portion;

wherein the third user interface portion displays one or more tables with rows and columns, wherein the one or more tables represent a particular version of the rule and one or more of the rows represent a particular rule case;

in response to the receiving of the instructions, executing the received instructions, the executing of the received instructions including:

selecting, by the version control portion of the application, the particular version of the rule from the multiple versions of the rule stored in the system storage;

modifying, through the editor portion of the application, the particular version of the rule;

updating, at least one of the one or more stored records, with data specifying that the particular version of the rule is modified; and

based on the updating, transitioning, by the process portion of the application, the particular version of the rule from the first state of the plurality of states to the second, different state of the plurality of states; and

storing, in one of the one or more records shared by the version control portion of the application, the process portion of application, and the editor portion, information resulting from the executing of the instructions received through the user interface that integrates the version control, editor and process portions.

2. The method of claim 1, wherein the information resulting from the executing of the instructions includes the rule with the selected particular version and with the second state of the plurality of states.

3. The method of claim 1, wherein the process portion of the application is distinct from the version control portion of the application, and wherein the process portion of the application and the version control portion of the application are not configured to transmit data to each other and to receive data from each other.

4. The method of claim 1, wherein the process portion of the application is executed on a first server computer and the version control portion of the application is executed on a second server computer.

5. The method of claim 1, wherein the received instructions include one or more state instructions received through the process portion of the application presented in the first user interface portion and one or more version instructions received through the version control portion of the application presented in the second interface portion, and wherein:

the selecting, by the version control portion of the application, of the particular version of the rule from the multiple versions of the rule stored in the system storage is performed in response to the receiving of the one or more version instructions specifying the particular version; and

the transitioning, by the process portion of the application, of the particular version of the rule from the first state of the plurality of states to the second, different state of the plurality of states is performed in response to the receiving of the one or more state instructions specifying the state transition.

6. The method of claim 1, wherein the plurality of states are a plurality of states of a development lifecycle or an approval process of the rule.

7. The method of claim 1, further including:
executing, through the integrated user interface, the particular version of the rule in the second, different state by:
receiving, from the version control portion of the application or the process portion of the application, the particular version of the rule;
generating a transform for transforming data based on the particular version of the rule in the second, different state;
providing a data processing component including the transform to a dataflow graph; and
executing the dataflow graph including using the data processing component to apply the transform to data inputted to the dataflow graph.

8. The method of claim 1, wherein:
storing includes storing multiple versions of a ruleset in a system storage, in which the ruleset includes one or more rules, each rule including one or more rule cases that specify one or more values for the rule's output depending on the rule's input, and each rule case includes one or more input relationships in which the rule case is triggered by a given input record if one or more data values of the given input record meet the one or more input relationships;
selecting includes selecting, by the version control portion of the application, a particular version of the ruleset from the multiple versions of the ruleset stored in the system storage; and
transitioning includes transitioning, by the process portion of the application, the particular version of the ruleset from the first state of the plurality of states to the second, different state of the plurality of states.

9. One or more machine-readable hardware storage devices for implementing one or more rules for use with an application that includes a process portion, a version control portion, and an editor portion in a data processing system, the one or more machine-readable hardware storage devices storing instructions that are executable by one or more processors to perform operations including:

- storing multiple versions of a rule in a system storage, in which the rule includes one or more rule cases that specify one or more values for the rule's output depending on the rule's input, and each rule case includes one or more input relationships in which the rule case is triggered by a given input record if one or more data values of the given input record meet the one or more input relationships;

- wherein each version of the rule is configured to be associated with a plurality of states including a first state and a second, different state of the plurality of states;

- storing one or more records that are each shared by the process portion, the version control portion, and the editor portion;

- receiving instructions through a user interface that integrates, through the shared one or more records, the process portion, the version control portion, and the editor portion, with the process portion being presented in a first user interface portion, the version control portion being presented in a second user interface portion, and the editor portion being presented in a third user interface portion;

- wherein the third user interface portion displays one or more tables with rows and columns, wherein the one or more tables represent a particular version of the rule and one or more of the rows represent a particular rule case;

- in response to the receiving of the instructions, executing the received instructions, the executing of the received instructions including:

- selecting, by the version control portion of the application, the particular version of the rule from the multiple versions of the rule stored in the system storage;

- modifying, through the editor portion of the application, the particular version of the rule;

- updating, at least one of the one or more stored records, with data specifying that the particular version of the rule is modified; and

based on the updating, transitioning, by the process portion of the application, the particular version of the rule from the first state of the plurality of states to the second, different state of the plurality of states; and

storing, in one of the one or more records shared by the version control portion of the application, the process portion of application, and the editor portion, information resulting from the executing of the instructions received through the user interface that integrates the version control, editor and process portions.

10. The method of claim 1 or the one or more machine-readable hardware storage devices of claim 9, wherein the one or more shared records are linked to the multiple versions of the rule.

11. The method of claim 1 or the one or more machine-readable hardware storage devices of claim 9, wherein the third user interface portion of the integrated user interface displays plural entries representing a given rule that is associated with a given state of the rule and a given version of the rule, and wherein the process portion of the application is for submitting data specifying a state transition of the given rule, and the version control portion of the application is for specifying a new version of the given rule.

12. The method or the one or more machine-readable hardware storage devices of claim 11, wherein the method or the operations further include:

modifying a given one of the plural entries through the editor portion being presented in the third user interface portion of the integrated user interface;

generating a new rule based on the modified given entry;

saving by the data processing system the new rule as a new version of the given rule; and

automatically updating, by the data processing system, a given record of the one or more shared records with data specifying the new version of the given rule generated and an initial state of the new version of the given rule.

13. The method or the one or more machine-readable hardware storage devices of claim 11, wherein the plural entries are the rows that are rendered in the third user interface portion of the integrated user interface, and with each of the rows including cells, of which one or more cells are modified through a further user interface portion of the integrated user interface, and the version control portion is for specifying the new version of the given rule.

14. The method or the one or more machine-readable hardware storage devices of claim 13, wherein the method or the operations further include:

automatically updating by the data processing system a given record of the one or more shared records, with data specifying the new version of the given rule generated and an initial state of the new version of the given rule, based on actions performed through at least two of the first user interface portion, the second user interface portion, and the third user interface portion of the integrated user interface.

15. The method of claim 1 or the one or more machine-readable hardware storage devices of claim 9, wherein the instructions are first instructions, and the method or the operations further include receiving, through the integrated user interface, second instructions and in response to the receiving of the second instructions, executing the received second instructions, the executing of the received second instructions including:

selecting a group of two or more rules from a plurality of rules stored in the system storage, each rule in the group being a particular version of a plurality of versions of the rules; associating the group of selected rules with a release tag; and

transitioning the release tag from a first state of a plurality of states to a second state of the plurality of states.

16. The method of claim 1 or the one or more machine-readable hardware storage devices of claim 9, wherein the information resulting from the executing of the instructions includes log information that represents the selecting performed by the version control portion of the application responsive to the received instructions and the transitioning performed by the process portion of the application responsive to the received instructions.

17. The method or the one or more machine-readable hardware storage devices of claim 12, further including receiving a user-initiated selection of simulator portion of the application being presented in a fourth user interface portion of the integrated user interface;

in response to receiving the selection of the simulation portion of the application:

causing, by the simulator portion of the application being accessed by the version control portion of the application, the new rule represented by the entries to execute against a given input data record to perform a simulation of the execution of the new version of the given rule; and

automatically updating, by the data processing system, a given record of the one or more shared records with information resulting from the simulation.

18. The method or the one or more machine-readable hardware storage devices of claim 17, wherein the fourth user interface portion provides visualizations indicative of the simulation currently being performed.

19. The method or the one or more machine-readable hardware storage devices of claim 18, wherein the visualizations includes: input visualizations that indicate input data of the given input data record on which the new rule is executing during the simulation; or visualizations indicating which rule case of the new rule has been triggered during the simulation; or output visualizations indicating output data that results from the execution of the new rule against the given input data record during the simulation.

20. The method or the one or more machine-readable hardware storage devices of claim 17, wherein the information resulting from the simulation includes: data specifying that instructions were received to perform the simulation on the new version of the given rule; or data that is indicative of the results of the simulation.

21. The method or the one or more machine-readable hardware storage devices of claim 18, wherein the visualizations provided by the fourth user interface portion are indicative of the progress of the execution of the new rule in real-time.