LA Generation Service - user manual

version 0.2.0 - 13 December 2013

PoSecCo

http://www.posecco.eu

Contents

1 Introduction

This document provides an overview of the *LA Generation Service*, which is the tool used to refine IT level policies into logical associations (LA) in the PoSecCo workflow.

This service is composed by a set of different modules with their own user interface. All of these modules and their options are documented in the following sections.

In Sec. 2 a short description of policies and logical associations is provided. They are respectively the data processed and produced by the LA Generation Service.

For an introduction to the LA Generation process see Sec. 3.

Finally, Sec. 4 is devoted to explaining how the tool works, its functionalities, its internal modules and the graphical user interface.

2 Data models

Policies

In the PoSecCo vision, after that a security expert has extracted a set of policies from the business requirements, the next step is to perform the policy refinement that is used to create a less abstract representation of the requirements which is nearer the machine world and more distant from the human mind. In PoSecCo there exist two types of policies:

Authentication policies The authentication policies are used to instruct the system that a user or a group of users can access a particular service, that is these policies are used to configure the authentication systems such as the ACLs.

Authorization policies The authorization policies can be used, among other things, to allow or deny communications between network endpoints, possibly by specifying some protection constraints. These policies are actually implemented by access control mechanisms, as well as filtering (e.g., firewalls) or data protection (e.g., SSL, WS-Security) security controls.

As shown in Fig. 1, where the general PoSecCo refinement process is presented, the LA Generation Service is only involved in the translation of policies that require the configuration of infrastructure security controls, such as filtering devices. For this reason, the LA Generation Service has to process only authorization policies, since only this type of policies can be implemented by infrastructure security controls. Hereafter, the term policy will implicitly mean authorization policy.

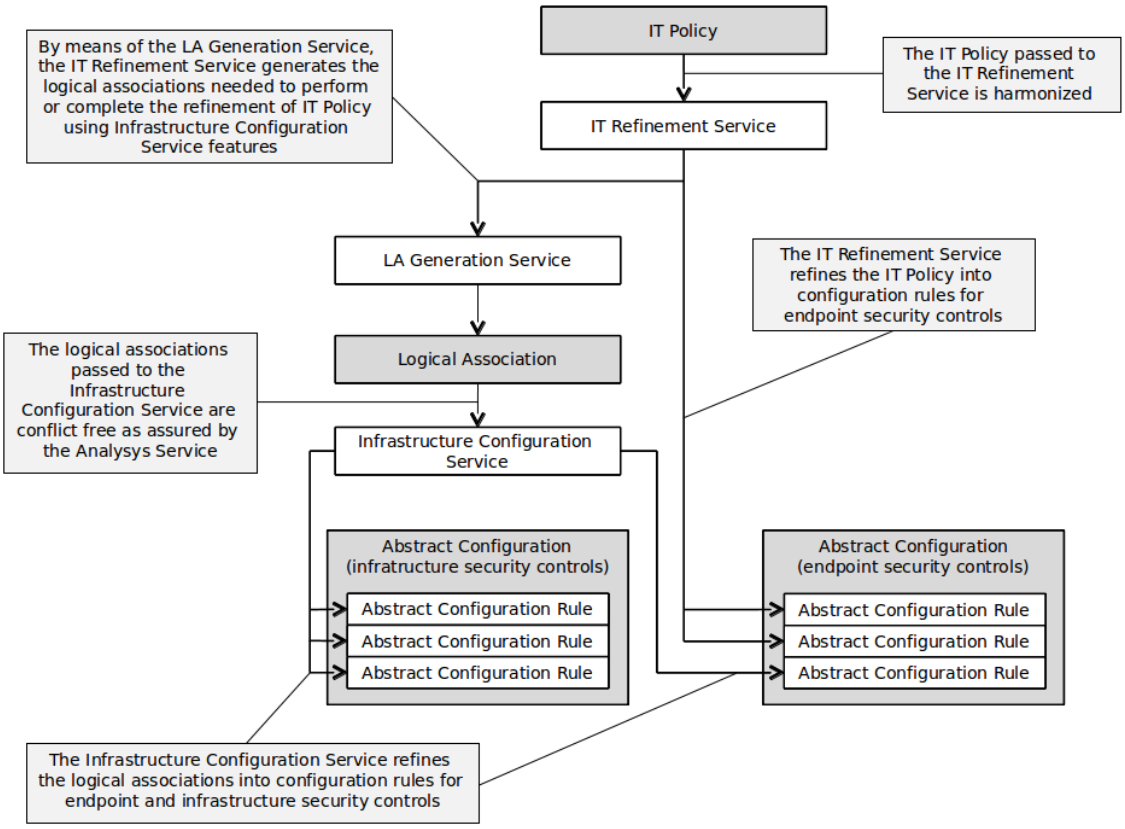

Figure 1: The workflow of the PoSecCo refinement process.

In the PoSecCo representation, an authorization policy can be defined in two different ways:

- by explicitly writing an *IT policy*. These policies are used to tell that a user or a group of users can access a service or a particular resource;

- by defining a *link* between two endpoints. These are implicit policies that are used to tell that the communication between two services or resources must be authorized.

The LA Generation Service supports both types of authorization policies (IT policies and links).

Logical associations

A *logical association*, or *LA* for short, is an intermediate format between the policies and the configurations. It is a lower-level directive than a policy, but it does not contain all the technical information of a configuration.

In general a logical association contains the following data:

- The communication endpoints, that are a (single or multiple) source and a (single or multiple) destination. Note that the LA endpoints are different from the policy endpoints. The policy endpoints are usually IT services, users and so on, while the LA endpoints are low-level elements which usually contain information such as IP addresses, ports and URIs.
- A *privilege*, that dictates the privilege between the two endpoints. For our case the privilege is always LAReach, that is used to state that the two endpoints can communicate.
- A set of security properties, used to select individual protections (e.g., confidentiality, data authenticity, key exchange or peer authentication) or aggregation of individual properties by means of templates (e.g., *HighSecurityDataProtection*). Additionally, other attributes are used to define constraints on the ISO/OSI layer where the protection has to be applied or the specific technology (e.g., SSL/TLS, WS-Security, . . .).

Example

In this section an example of policy and generated logical association is provided.

The policy is the following:

Service provider's application administrators can access the administration interface of Tomcat servers.

The presented policy, written in natural language to simplify the reading, can be represented as an XMI instance of the PoSecCo security meta-model. As you can see, the policy is quite abstract and it contains no network-aware detail or technological constraint. For this reason, to convert the policy into configuration rules for the security controls present in the landscape, a refinement process is required. By executing the LA Generation Service, the presented policy could be refined into the following logical association, also presented in natural language to simplify the reading:

Computer C1 securely reaches the Tomcat web administration interface running on server S1 at address 172.17.8.131:8080 using the SSL/TLS protocol.

As you can see, the generated LA contains network and technology-aware details. Other topology-related details will be added by the refinement process executed by the Infrastructure Configuration Service (see user manual [1]) that transforms logical associations into configuration rules for the security controls present in the landscape.

3 Refinement process

The *policy refinement* process consists of translating the policies into one or more logical associations. The translation is not as trivial as it seems, since several LA properties are not directly specified in the policies but must be inferred by analyzing the landscape and the policies themselves.

In order to perform the required deductions and translations, the LA Generation Service intensively uses an ontological and reasoning system based on the OWL standard and several reasoners.

The LA Generation workflow is composed by the following phases:

1. **Connect to the ontology**. The first step retrieves the landscape description and the IT level policies from the PoSecCo repository. Based on these data, it generates an internal ontology-based representation which also contains all classes and associations defined in the Logical Associations meta-model. This workflow step is internally split in two different phases. The first one, named the *TBox Phase*, is responsible for the creation of the TBox part of the ontology, i.e., the ontology portion that only contains classes and property definitions of all the different PoSecCo meta-models. The second one, named the *ABox Phase*, fills the ABox portion of the ontology, i.e., it converts instances read from the PoSecCo repository into ontology individuals. The information needed to contact the PoSecCo repository, such as the IP address and port, can be manually configured by the user.

2. **Select the items to refine**. Once the PoSecCo ontology is correctly formed, the LA Generation Service asks the user the IT policies and links to refine through a UI. By knowing in advance the items to be refined, it is possible to speed-up the next steps by selectively applying the algorithms only to the minimum number of objects in the ontology.

3. **Perform the low-level mapping**. Starting from the IT policies and landscape information described in the ontology, this workflow step manages the definition and maintenance of connections between policy subjects or objects and LA endpoints. This kind of connection between the two different abstraction layers is essential for the effective execution of the refinement process because it permits to attach more detailed infrastructure data to IT level policies. For instance, during this phase, all the users described in the IT policies are linked to the computers they work with. In this way, authorization policies used to grant users some access privilege can later be translated into filtering configuration rules having as source or target IP address those of the computers users work with.

4. **Enrich the ontology**. This workflow phase handles the execution of enrichment modules (EMs), whose purpose is to better classify information described in the ontology and to infer previously unspecified details. This kind of reasoning can help the tool to better understand policy requirements and, as a consequence, to generate more precise logical associations. For example, an enrichment module could be used to distinguish public and private services in the landscape i.e., services accessible from every user on the web in contrast to services that can only be used by a company's employees. Enrichment modules could also be used to add technological constraints to policies. For instance, an EM could attach to policies some requirements related to the usage of WS-Security to protect connections traversing untrusted third parties' networks.

5. **Extract the LAs**. This is the final workflow step and it is responsible for the actual generation of the logical associations. During this phase, the user can specify several LA properties such as the aggregation level of the generated LAs and the usage of specific security technologies. Once the generation process is completed, the logical associations are written back to the PoSecCo repository.

4 Use of the tool

Graphical User Interface

The LA Generation Service is used to refine IT level policies into logical associations. Fig. 2 represents the LA Generation main application. The left part contains the view that can be used to activate the execution of the LA Generation process, while the right part is composed by the *Refinable Items Explorer* tab, used to display the refinable items, i.e., IT policies and links, and the *LA Explorer* tab, used to display the generated logical associations.

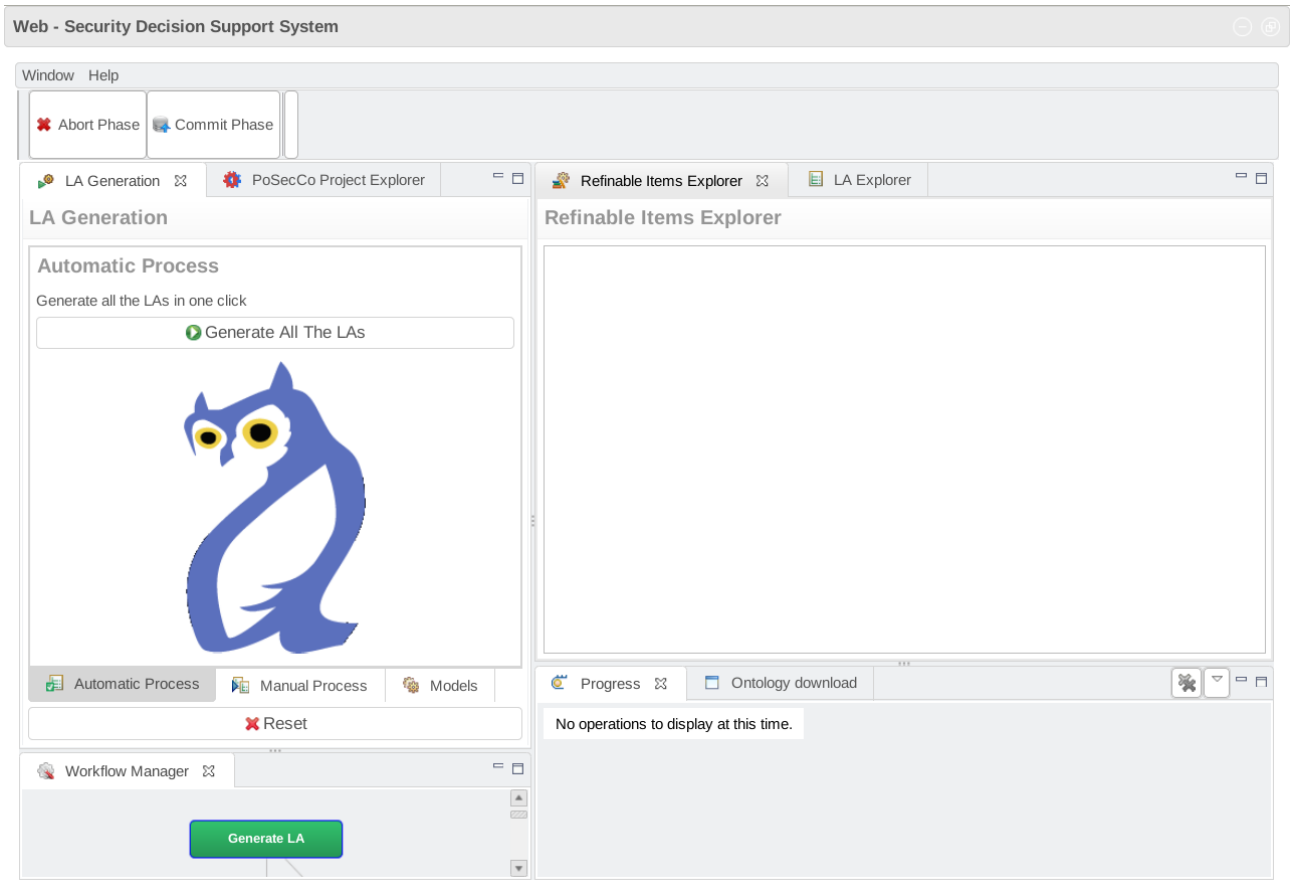

Figure 2: LA Generation Service main GUI

As you can see in Fig. 3, the LA Generation process can be automatic (you can press the *Generate All The LAs* button in the *Automatic Process* tab shown in Fig. 3a) or it can be manually performed step by step (the steps are accessible in the *Manual Process* view presented in Fig. 3b).

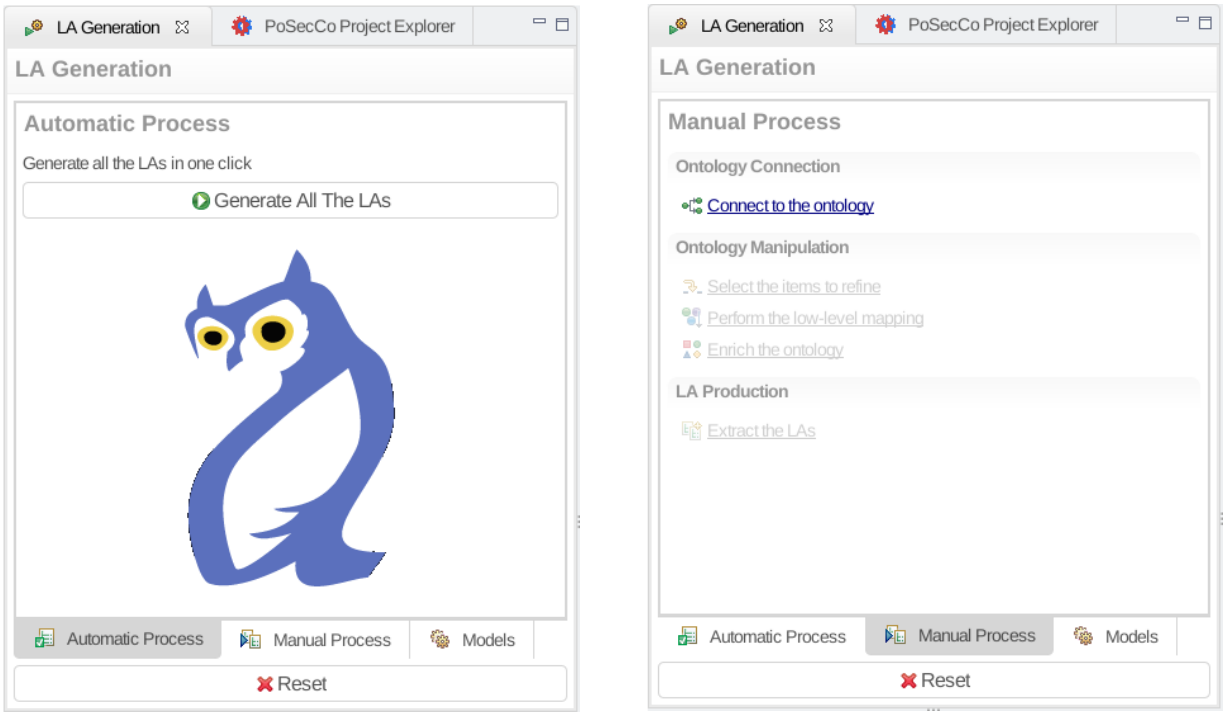

(a) Automatic execution

(b) Manual execution

Figure 3: View to execute the LA Generation workflow

The tool also offers the *Reset* button, that can be used to clear the LA Generation process. In this way you will be able to discard all execution data and start the process from scratch.

In addition, the tab shown in Fig. 4 allows the user to control the internal model of the LA Generation Service.

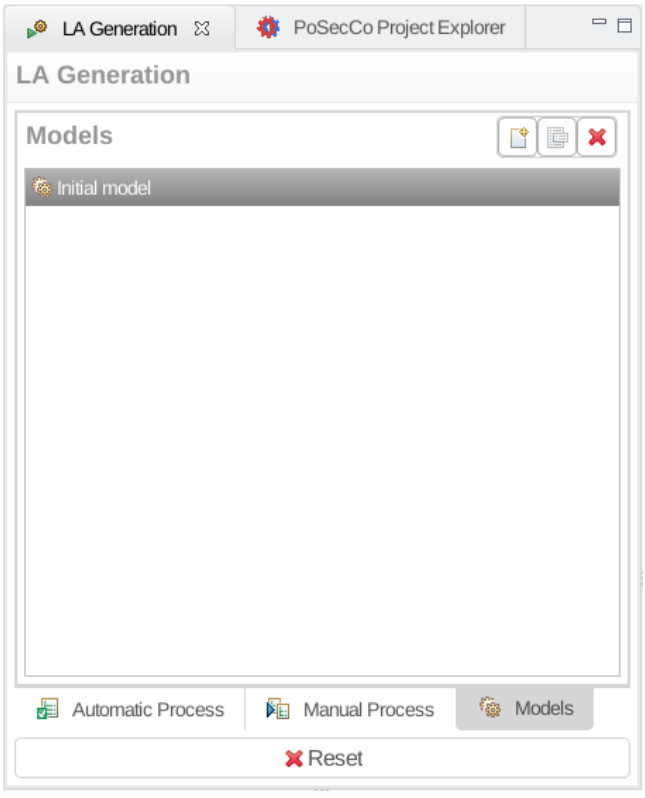

Figure 4: LA Generation Service model management

The LA Generation Service uses a multi-MVC approach, that means that the user can internally use a set of

models, shown in this tab. The buttons in the top-right corner allow the user to create a new empty model, to change the name of the selected one or to deleted an existing model.

LA Generation Service preferences

The tool offers the user the possibility to define some preference values for the LA Generation process. To access the LA Generation Service preference page, you can click on *Window* → *Preferences* and expand the *LA Generation* category as shown in Fig. 5.

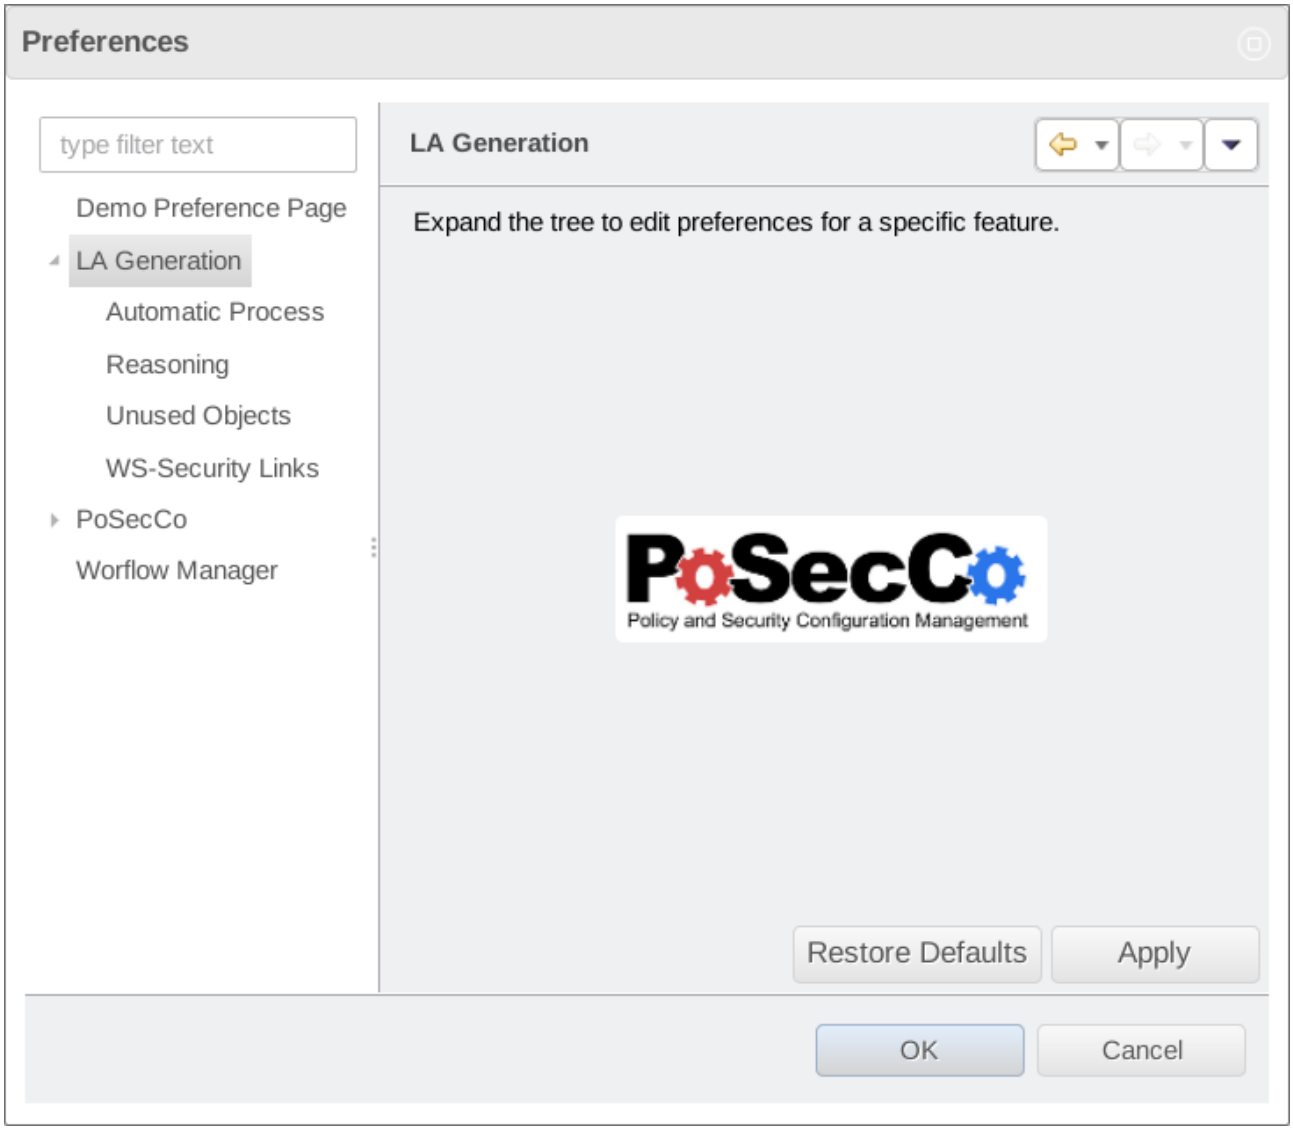

Figure 5: LA Generation Service preferences

Preferences are divided in different categories. The first one, *Automatic Process*, is shown in Fig. 6 and it lets you define preference values for the previously presented automatic execution mode. More precisely, you can specify preference values for the following workflow phases:

- Refinable items selection phase: you can specify how the tool should behave regarding the selection of which policies and links must be refined. The two options are the following:
  - *Select everything*: this value can be used to specify that the tool must perform the refinement process for all available policies and links;
  - *Ask about the refinable items*: by using this option, the tool will ask you to choose which items should be refined among the existing ones.

- Low-level mapping phase: you can define how the tool should behave if some missing mapping association is identified. The possible values are the following:
  - *Ignore all the missing relationships*: this option can be used to specify that you want to ignore all missing mapping associations;
  - *Ask about all the missing relationships*: you can use this option if you want the tool to ask you about all identified missing mapping associations.
- Enrichment phase: you also have the chance to control the behavior of the enrichment process. In particular, you can choose one of the following approaches:
  - *Never run any EM*: during the enrichment phase, none of the available EMs will be executed;
  - *Ask the EMs to run*: the tool will ask you to choose which enrichment modules you want to execute. This is the same behavior as in the manual execution mode;
  - *Execute all the EMs and validate the results*: the tool will execute all available enrichment modules and it will let you view and validate their results;
  - *Execute all the EMs without the validation*: the tool will execute all available enrichment modules without asking you to validate their results.
- LA extraction phase: for the final workflow phase, you can specify one of the following values:
  - *Ask about the LA options*: by choosing this value, you will be able to define further details for the generated logical associations, such as aggregation level and security technology constraints;
  - *Extract all the LAs without asking anything*: this option can be used to let the tool choose the following default values for the LA extraction phase:
    - where possible, LAs with aggregated endpoints will be generated. Aggregated endpoints are a particular class of LA endpoints that can be used to group together multiple (source and/or target) endpoints in a single logical association. This can be useful to reduce the number of configuration rules and, as a consequence, to achieve better system performances. For instance, if a policy applies to all computers of a network zone, it is possible to generate a single LA that will be later refined into just one configuration rule on the firewalls present in the landscape;
    - all generated logical associations will contain information integrity and confidentiality requirements;
    - no specific technology constraint will be added to the generated LAs, unless required by EMs' inferences.

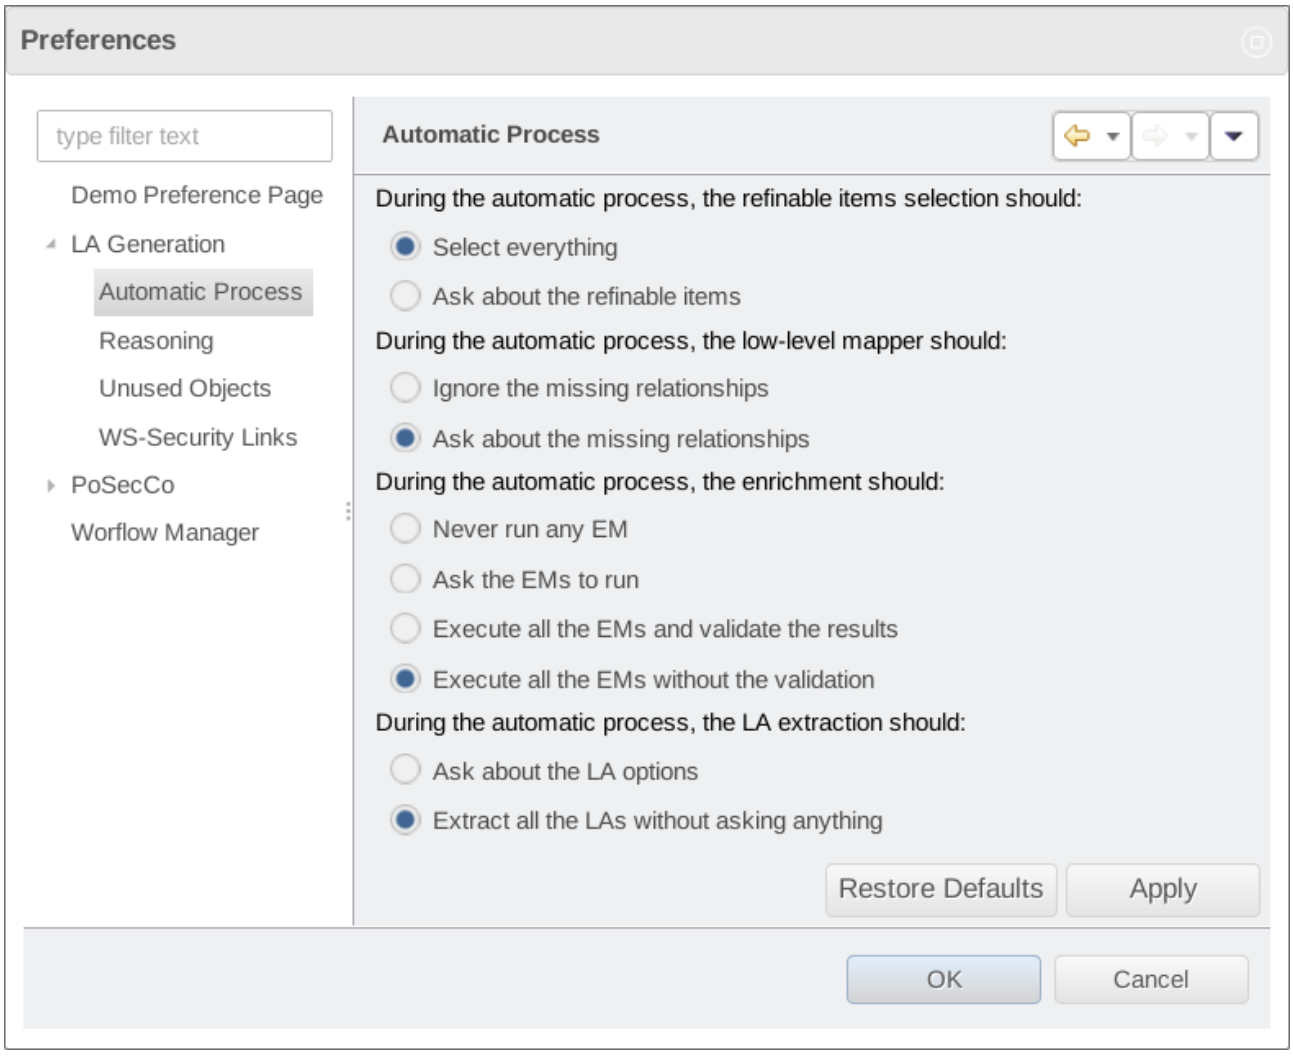

Figure 6: Automatic execution mode preferences

The second category, *Reasoning*, is shown in Fig. 7. In this category you can decide what reasoner you want the tool to use in the different phases of the LA Generation process. In particular, at the current stage of development, you can choose either *Pellet*[1] or *Hermit*[2]. Please note that you can choose different reasoners for different phases.

In addition to selecting the reasoner, you can also choose if you want the tool to perform an ontology consistency check after each of the LA Generation phases. This can be done by clicking on the checkboxes located to the left of each phase (*Perform a consistency check after...*).

[1] Pellet reasoner, http://clarkparsia.com/pellet/
[2] Hermit reasoner, http://www.hermit-reasoner.com/

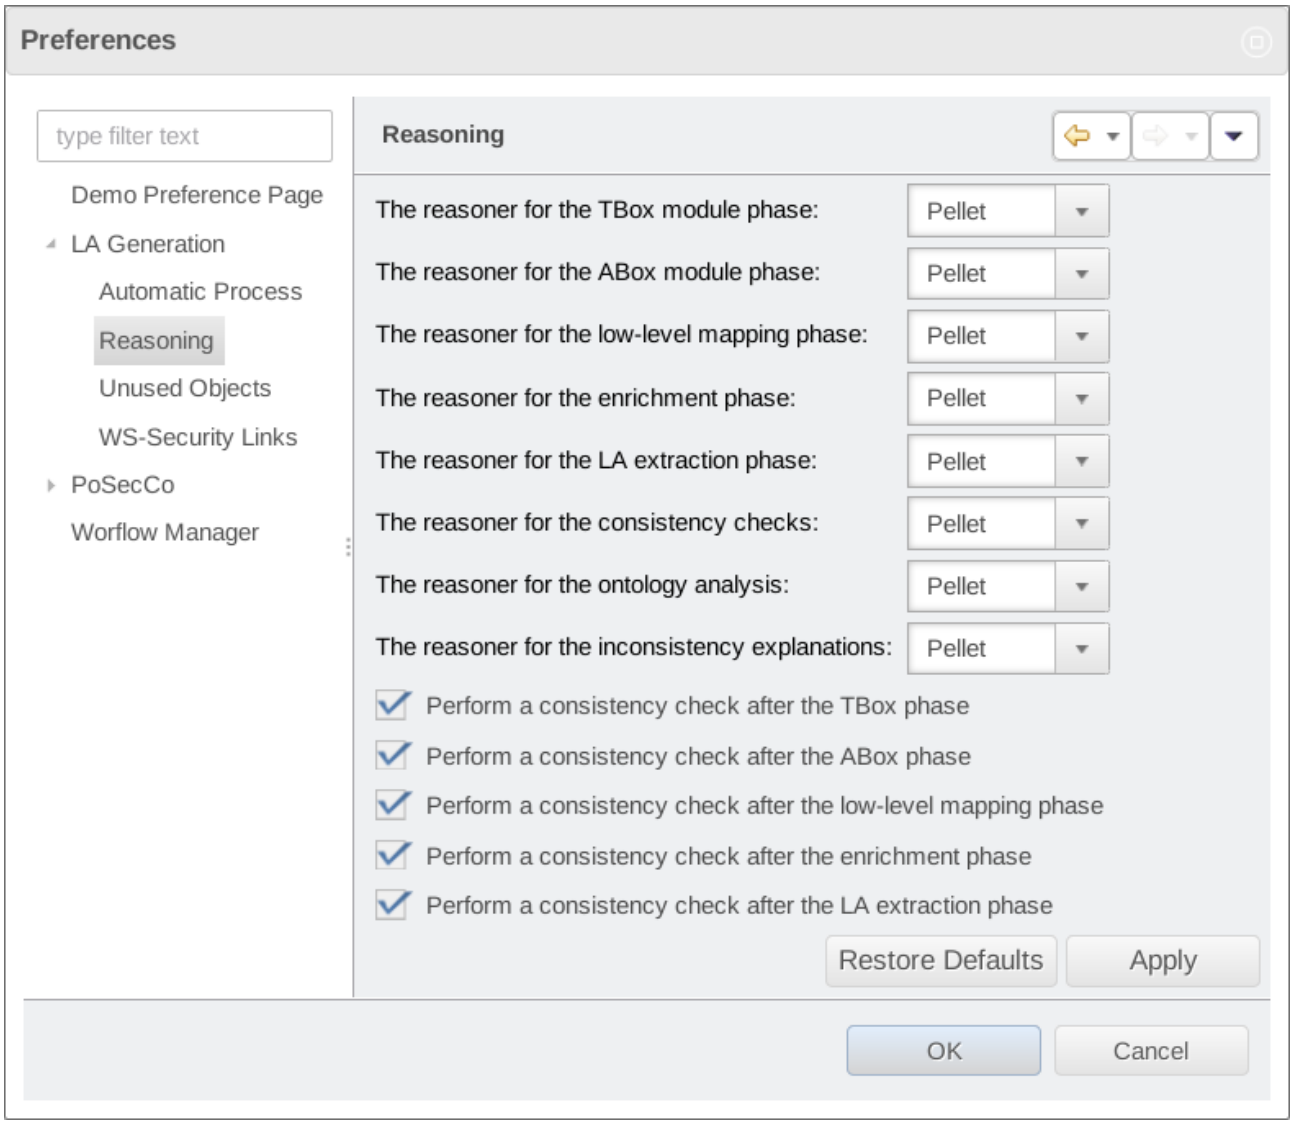

Figure 7: Reasoning preferences

The third category, *Unused Objects*, is presented in Fig. 8 and it corresponds to the custom preference page for the *Unused Objects Classifier* enrichment module. This EM can be used to distinguish IT level objects that are not involved in policies and links from those that are instead used. To this purpose, the preference page offers the ability to define how the EM should behave by specifying if it should identify:

- *The unused objects only*: in this way, the module will only identify the objects that are not involved in policies and links;
- *The used objects only*: by specifying this value, the module will detect only those objects that are actually involved in polices or links;
- *The used and unused objects*: this value can be chosen to make the tool identify both categories of objects.

Furthermore, it is also possible to choose the type of IT objects that should be analyzed (data models, IT interface models, IT resource models and IT service models).

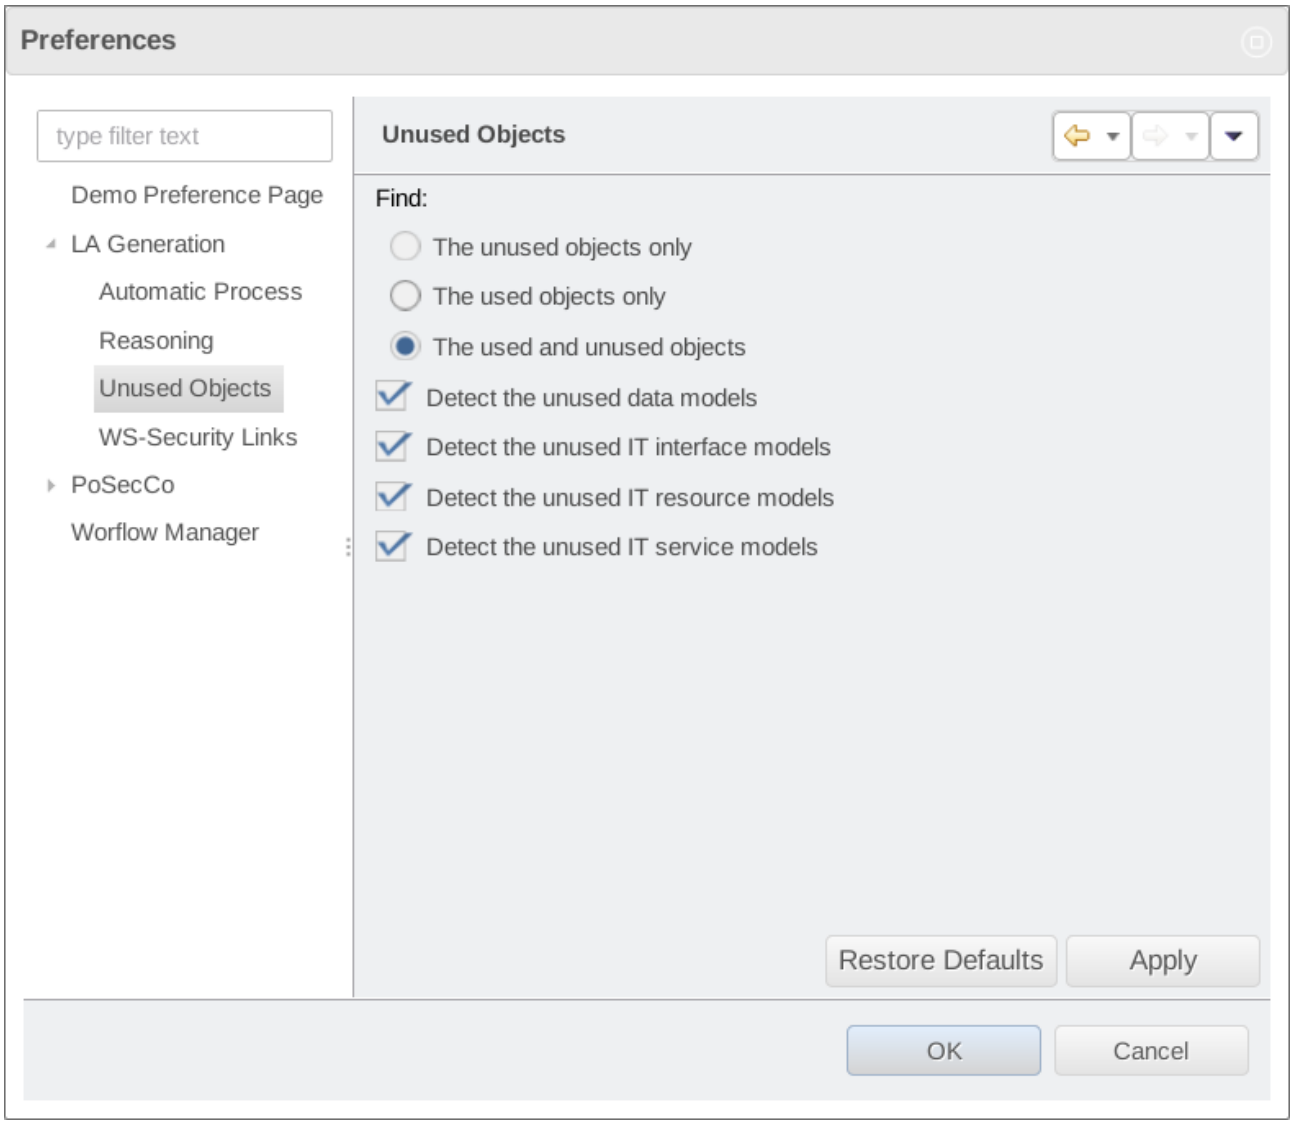

Figure 8: Unused Objects preferences

The final category, *WS-Security Links*, is presented in Fig. 9 and it corresponds to the custom preference page for the *WS-Security Link Classifier* enrichment module.

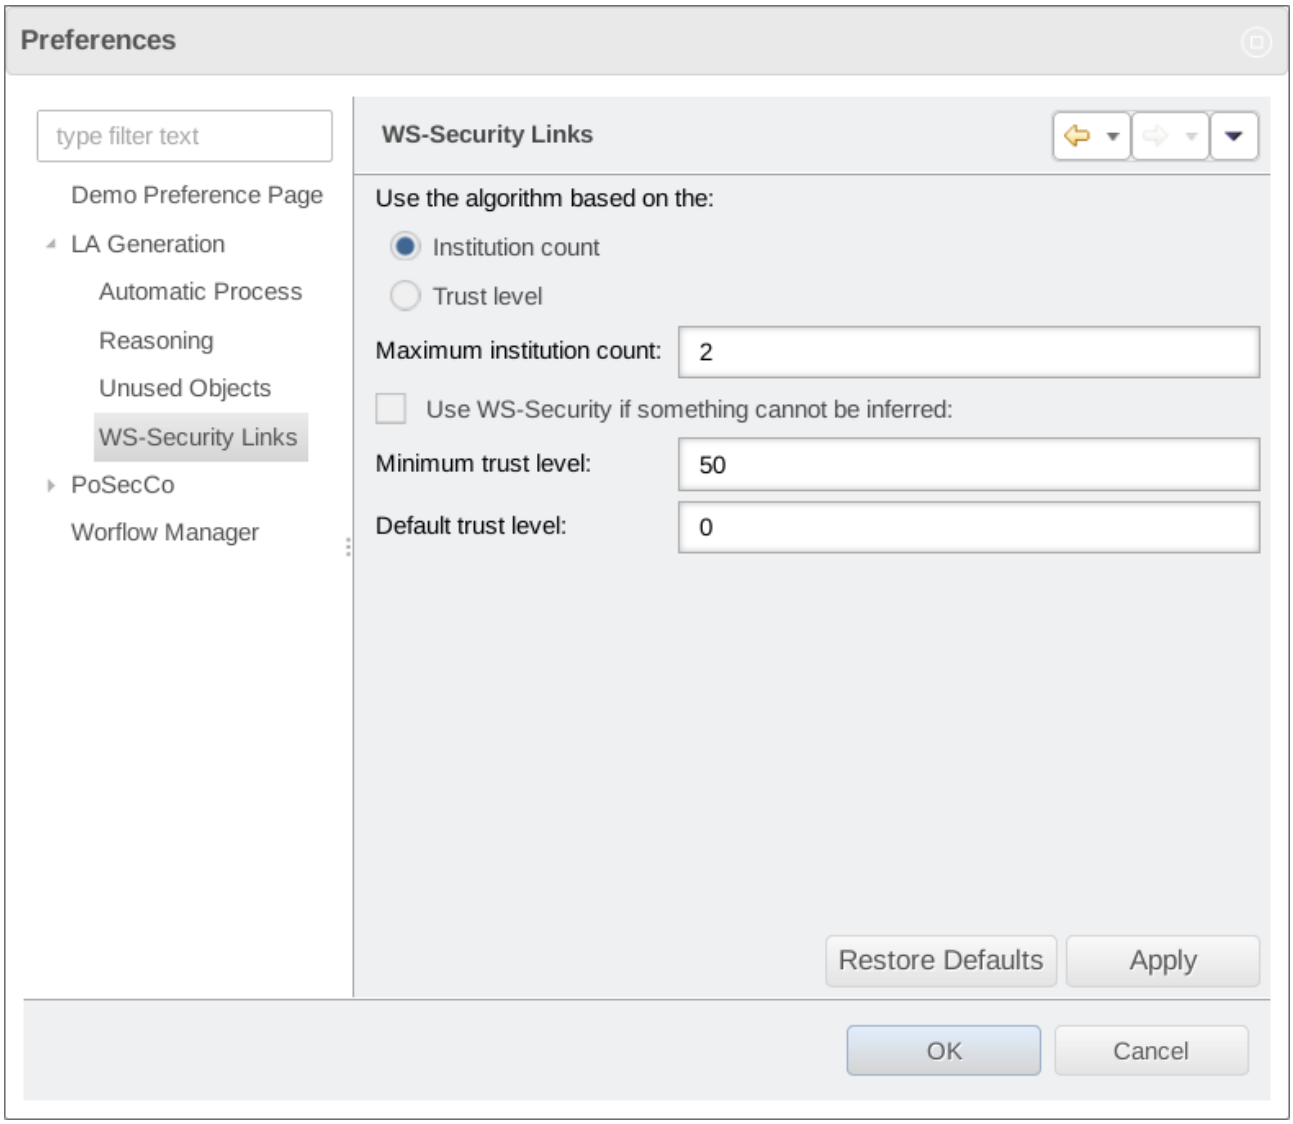

Figure 9: WS-Security Link Classifier preferences

The preference page shows different options but the most important is the first one, which allows the user to select the algorithm to identify the links that need WS-Security. The two algorithms are based on the:

- *Institution count.* If you select this algorithm the module will suggest the usage of the WS-Security protocol to protect links that involve a number of institutions bigger than a specified maximum value. The threshold can be specified by defining the *Maximum institution count* preference value.
- *Trust level.* If you select this algorithm the module will enable the WS-Security protocol by calculating a trust level[3] for a link. If a link's trust level is lower than a specified threshold, then WS-Security will be suggested. The option *Minimum trust level* allows you to specify the aforementioned threshold, while the *Default trust level* allows you to set a default level that will be used in case this information is not specified for a link or if it cannot be automatically inferred by the tool.

The *Use WS-Security if something cannot be inferred* value allows you to automatically enable this protocol if some properties about a link (its institution count or trust level) cannot be inferred due to some missing data in the ontology.

[3]The trust level is an integer where the lowest values are the less trusted ones.

Workflow execution

To introduce all views and information related to the LA Generation process, we consider the previously presented manual execution mode. Please note that the automatic mode is similar to the manual one, but some of the views may not be shown depending on the preferences you specified.

Step 1. By clicking on *Connect to the ontology* the tool loads the IT level policies and the landscape from the PoSecCo repository to the internal ontology. Note that the other menu items are disabled. During this phase, the tool also checks the existence of NAT configurations in the landscape description. If no NAT configuration is found, the user has the ability to upload a separate ontology containing the required configurations. This can be achieved by means of the *NAT Configurations Upload* wizard shown in Fig. 10. By clicking on the *Browse* button, you can select an ontology file in your file system. To perform the upload, you can click on the *Next* button. You will then be presented with the details of the uploaded NAT configurations (Fig. 11). Finally, to confirm the uploaded configurations, you can click on the *Finish* button and the configurations will be inserted in the internal ontology. Instead, if you don't need to upload any NAT configuration, you can simply ignore this step by clicking on the *Cancel* button.

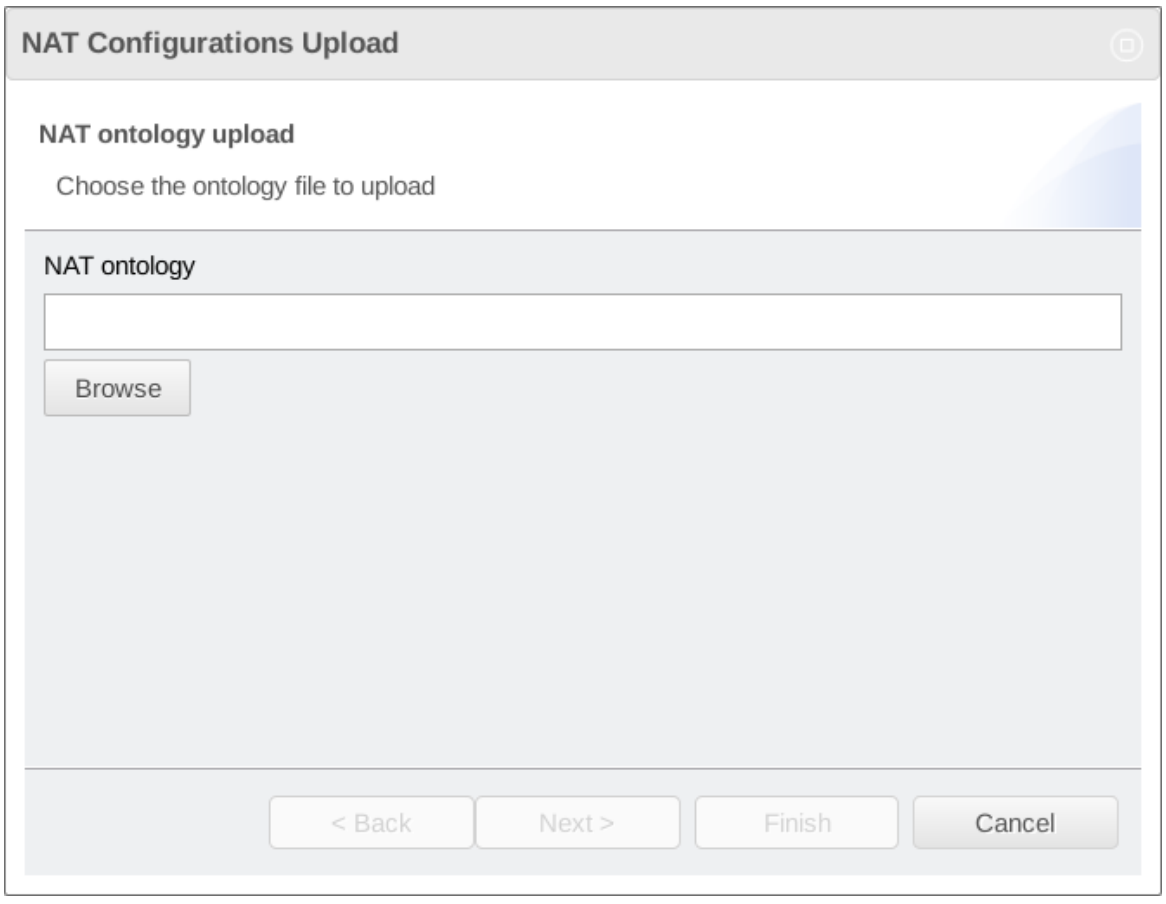

Figure 10: NAT upload

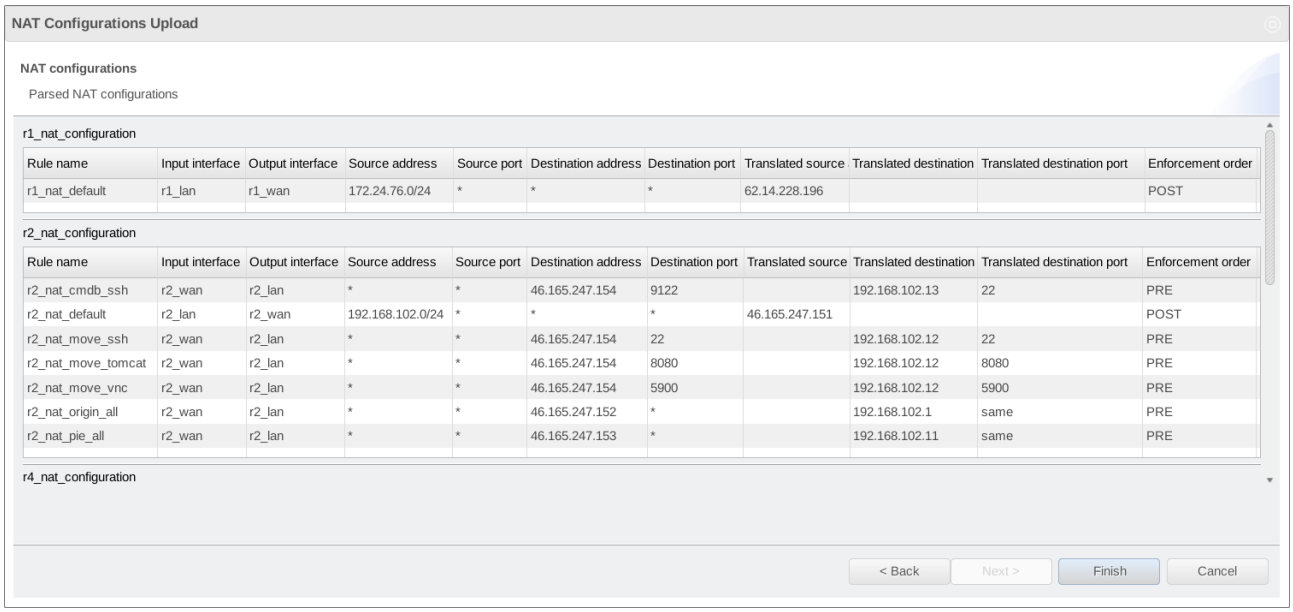

Figure 11: NAT view

Once the input operations described above are completed, the resulting set of refinable items, containing both policies and links, is displayed in the *Refinable Items Explorer* view (Fig. 12). For every IT policy, the tool shows the system authorization rules that compose it and, for each of them, the following information is presented:

- the sign, used to distinguish positive authorizations from negative ones;
- the subject who is granted or denied the authorization privilege;
- the target resource;
- the action specified by the authorization privilege.

In case of links, instead, the presented information is the following:

- the link endpoints (source and target);
- the type of data transported by the link.

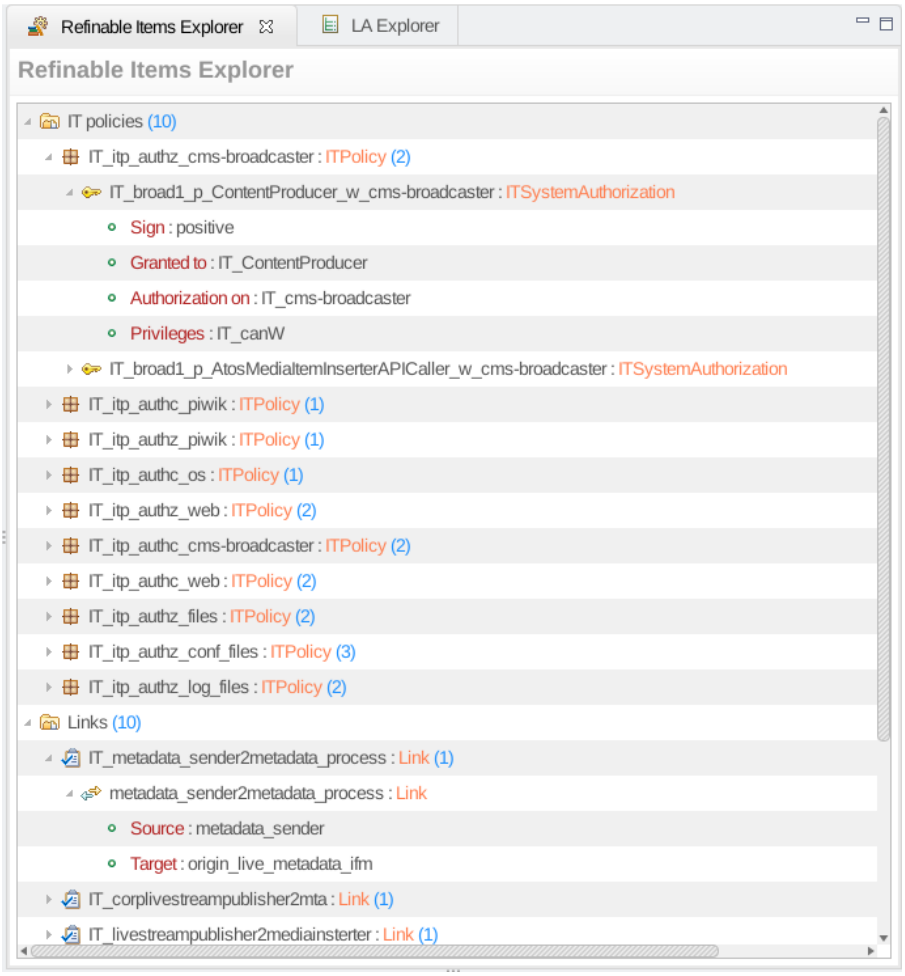

Figure 12: Refinable items details

Step 2. After that the PoSecCo ontology is created, you can click on *Select the items to refine* in order to choose the policies and links that will be refined. The Fig. 13 shows the wizard that will appear. The refinable items, divided in IT policies and links, are displayed. For each of them, it is possible to specify if the refinement process must be performed by ticking the corresponding checkbox. The view also offers the buttons *Select all*, to select all refinable items, and *Deselect all*, to clear the selection.

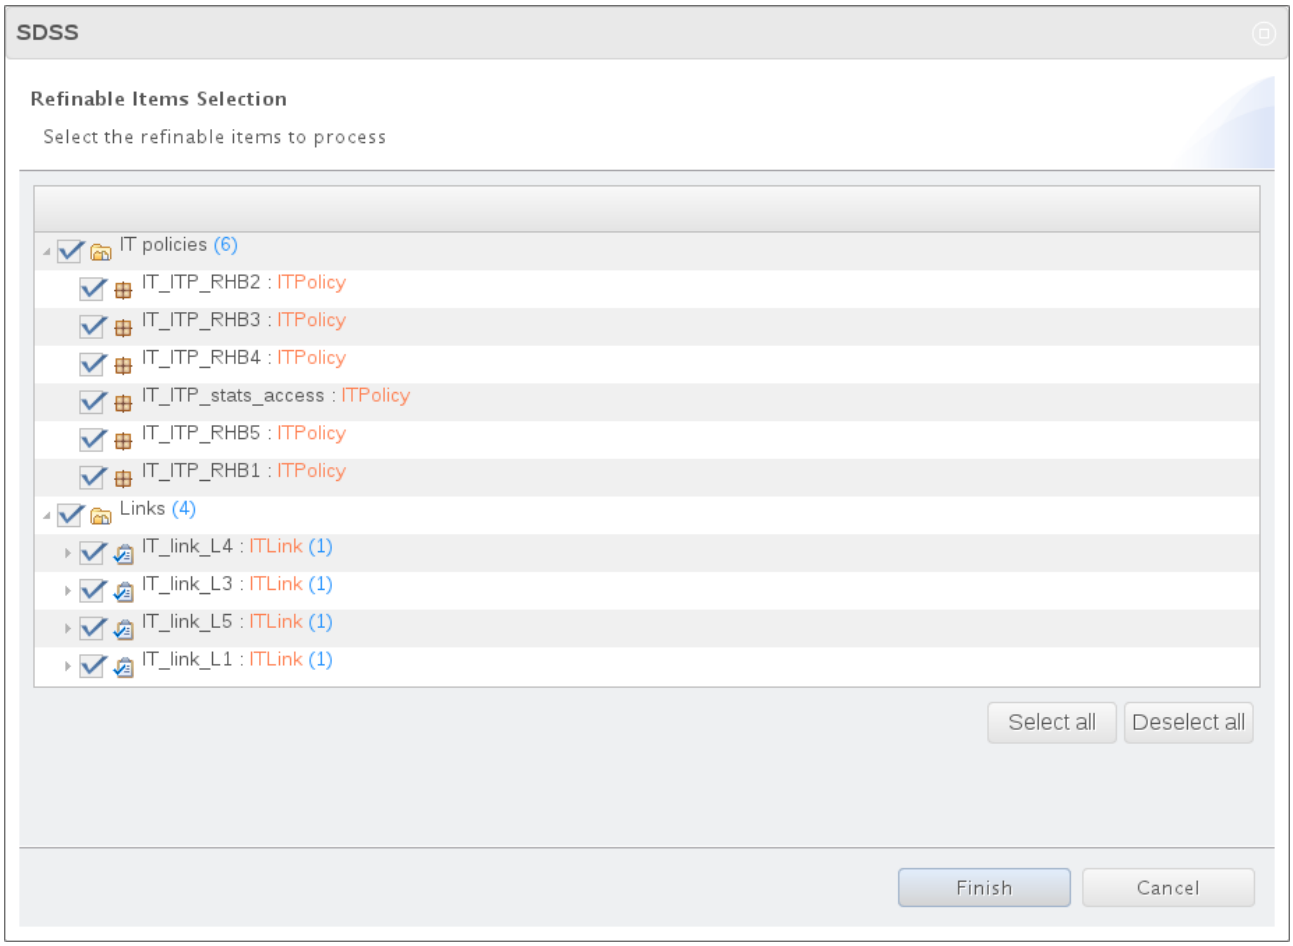

Figure 13: Refinable items selections

Step 3. Once the previous step is completed, you can click on *Perform the low-level mapping* to start the low-level mapping phase. During this phase the tool might detect some missing associations that are required for the completion of the mapping process. In this scenario, the user is asked to specify the required information by means of the *Missing object property assertions* view presented in Fig. 14. This view contains one or more missing associations, each of them identified by the following data:

- *Source individual*: the identifier of the individual for which a previously undefined association is required;
- *Source class*: the class of the individual;
- *Property*: the name of the missing association;
- *Target individual*: the possibly empty set of association values specified by the user;

When you select one of the missing associations, the tool computes a set of possible values and shows them in the *Candidate targets* portion of the view. You can choose one or more of the proposed values. Please note that you can also choose not to specify any value for a given missing association. In this scenario, the tool will remember this decision so that the same missing association will not be asked again during the same session. **However, given the importance of the low-level mapping process for the LA Generation process, please keep in mind that this kind of decision may affect the final results and you should follow this approach only in case you do not know which value must be specified**.

To confirm the specified association values, you can click on the *Finish* button. The *Cancel* button can instead be used to ignore all listed missing associations.

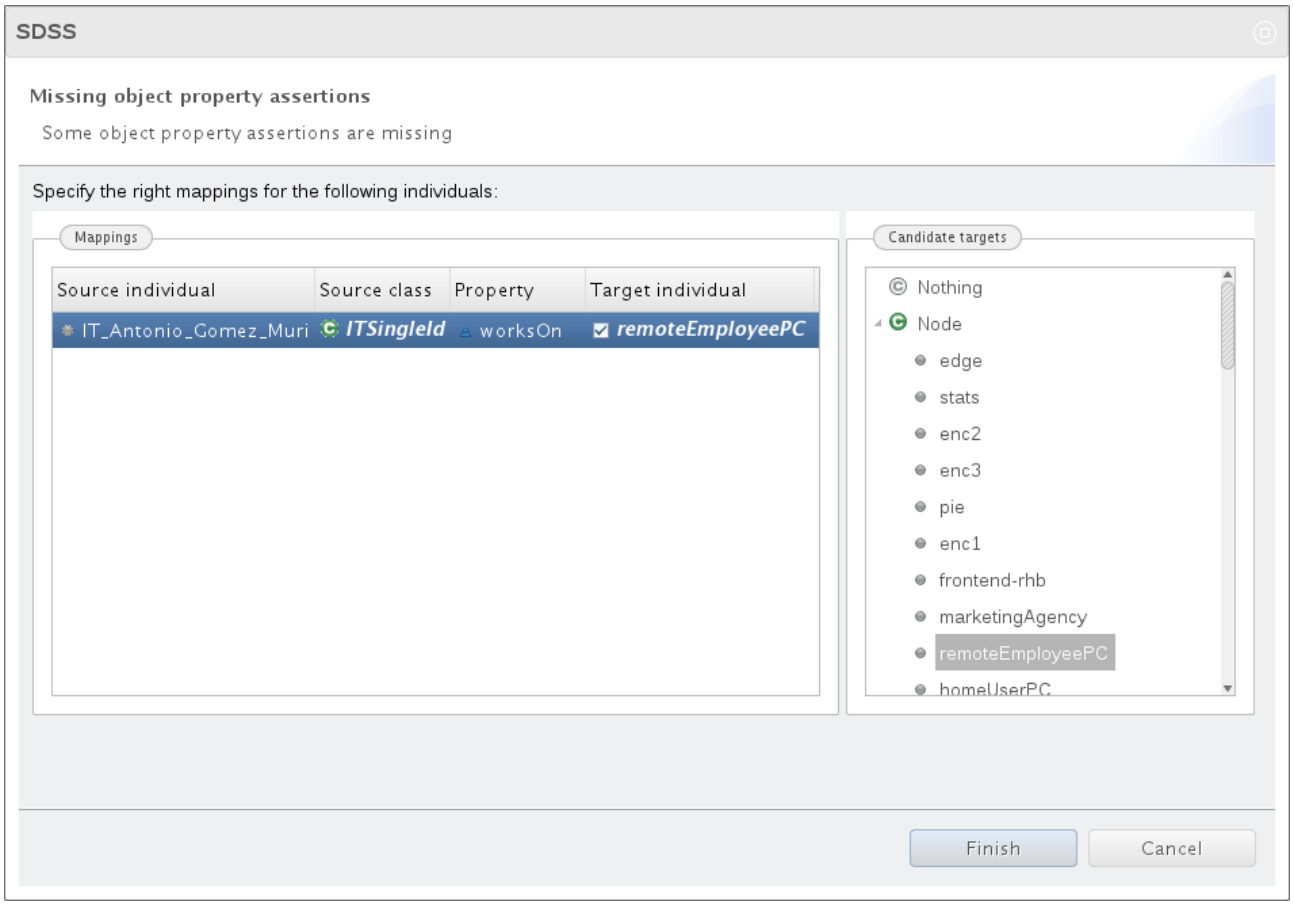

Figure 14: View to ask the user missing mapping associations

Step 4. After completing the low-level mapping phase, you can click on *Enrich the ontology* to execute enrichment modules. If any EM is available, the tool will show the *Enrichment Module Selection* view depicted in Fig. 15. This view contains a list of enrichment modules, each of them identified by its name (e.g., *WS-Security Link Classifier*). By clicking the checkbox at the left of the EM's name, you can specify whether you want to execute that specific enrichment module or not. Furthermore, in the *Validation* column, you can choose if you want to validate the EM's execution results or not, i.e., if you want to view the inferences suggested by the enrichment module and to either confirm or discard them. You can also use buttons *Select all* and *Deselect all* to execute all or none of the available EMs and buttons *Validate all* and *Validate none* to specify your validation preferences for all chosen enrichment modules. After choosing which EMs you want to execute, you can click on the *Finish* button to start the actual enrichment process.

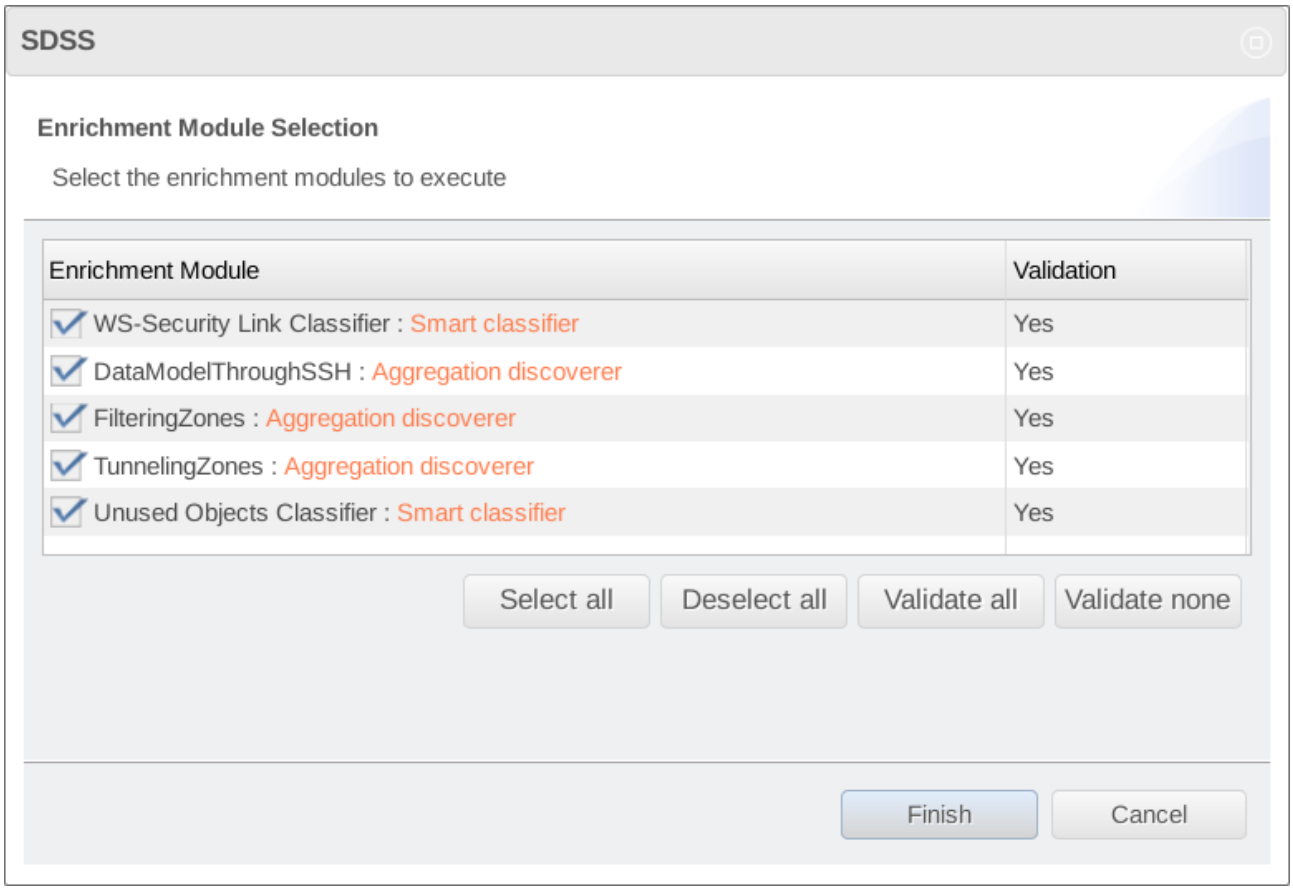

Figure 15: EMs selection

If you choose to validate an EM's results, you will be presented with a particular enrichment validation view that varies depending on the enrichment module type.

For the smart classifiers, the *Realization Validation* view depicted in Fig. 16 will appear. The view contains a list of enrichment inferences, each of them identified by the following information:

- *Individual*: identifier of the individual processed by the enrichment module;
- *Smart Class*: name of the new class used by the enrichment module to add details to individuals;
- *Explanation*: a string used to explain the enrichment result.

Please note that only selected individuals, i.e., those with a ticked checkbox, were suggested for reclassification by the enrichment module. You can modify an individual's enrichment result by either ticking its checkbox, to add it to the new class used by the EM, or clearing it, to keep the original classification. You can also use buttons *Select all*, to reclassify all listed individuals, or *Deselect all* to use the original classification.

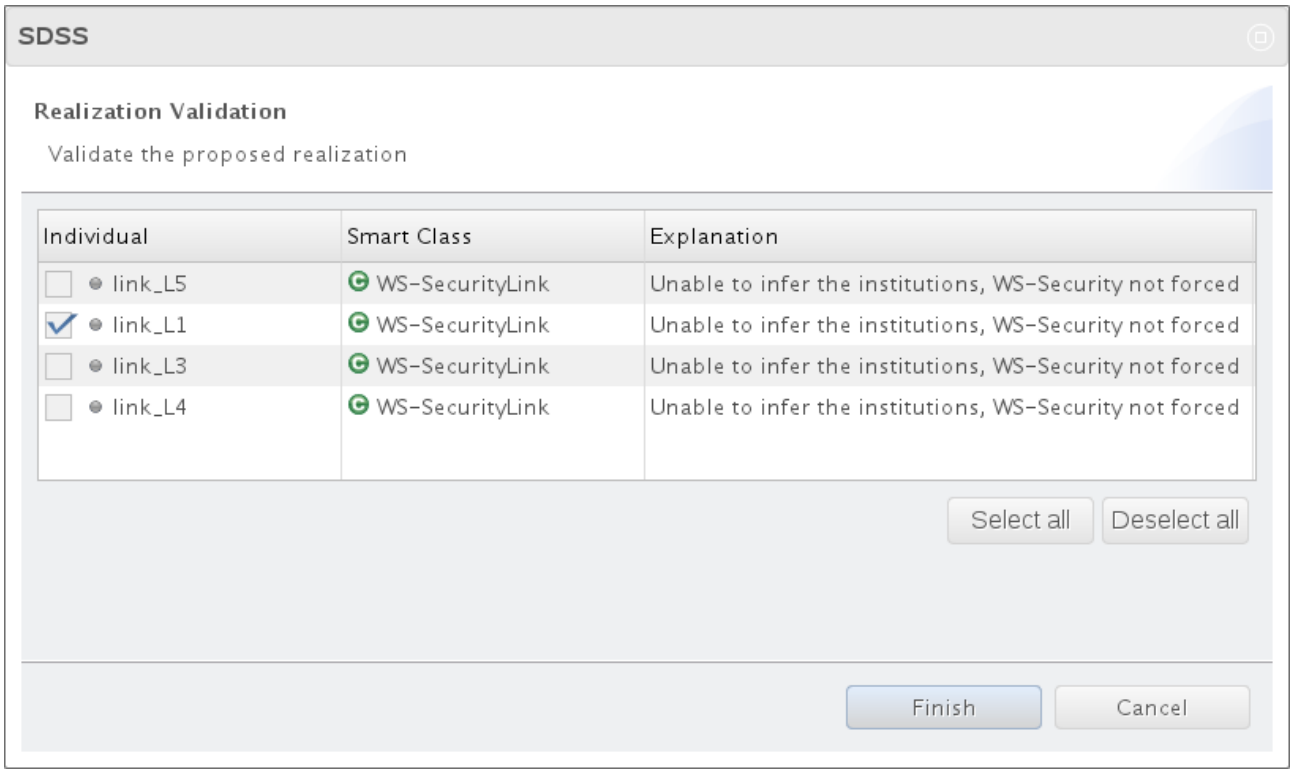

Figure 16: Smart classifiers results validation

For the aggregation discoverers, instead, the *Aggregation Validation* view depicted in Fig. 17 will appear. The view contains a list of enrichment inferences, each of them identified by the following information:

- *Aggregation*: the list of inferred aggregations and their class;
- *Aggregation Content*: the contents of the selected aggregation in the aggregation list (i.e., its internal individuals and properties).

You can enable or disable any aggregation inference, i.e., internal individuals, properties, and whole aggregations, by modifying the selection of the different checkboxes presented in the view. Please note that aggregation inferences can be related to each other. For this reason, it is possible that disabling one aggregation inference causes other ones to get disabled too.

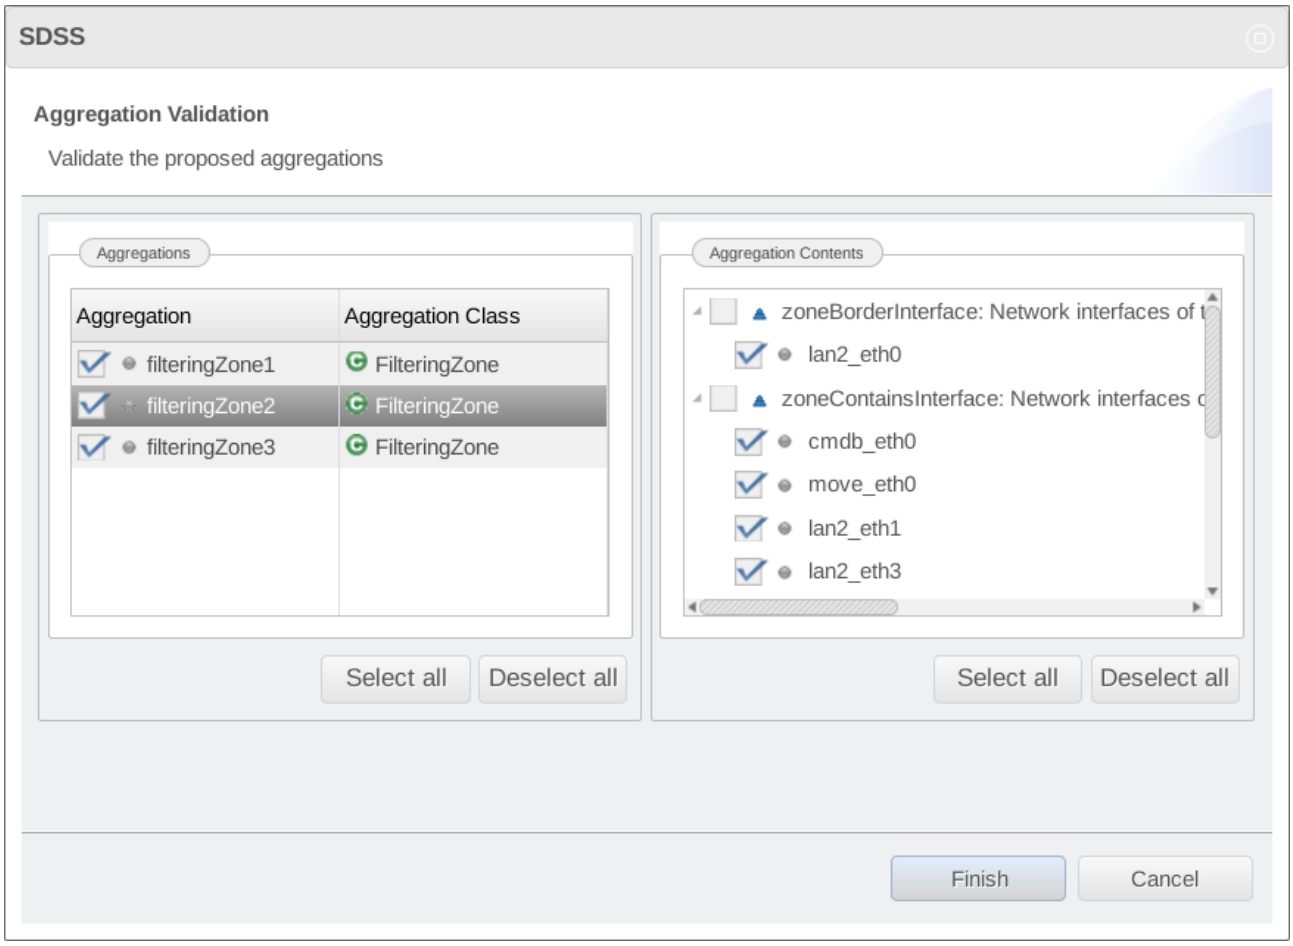

Figure 17: Aggregation discoverers results validation

Step 5. Once the enrichment process is completed, you can click on *Extract the LAs* to execute the final LA Generation step. The tool will show you the *Logical associations extraction* wizard, presented in Fig. 18, that will allow you to customize the produced LAs. In this view, you have the possibility to specify some further constraints for the LA Generation process. In particular, for every authorization rule and link, you can define values for the following properties:

- *Channels*: this option can be *Single*, to produce 1-to-1 LAs i.e., logical association with single subjects and objects, or *Aggregated* to produce LAs with aggregated endpoints;
- *Integrity*: this option can be turned *On* to produce logical associations with information integrity requirements;
- *Confidentiality*: this option can be turned *On* to produce logical associations with information confidentiality requirements;
- *Technology*: this field opens the wizard shown in Fig. 19 that can be used to force a particular set of technologies over a LA. If no technology is selected, then the best one will be chosen by the Infrastructure Configuration Service.

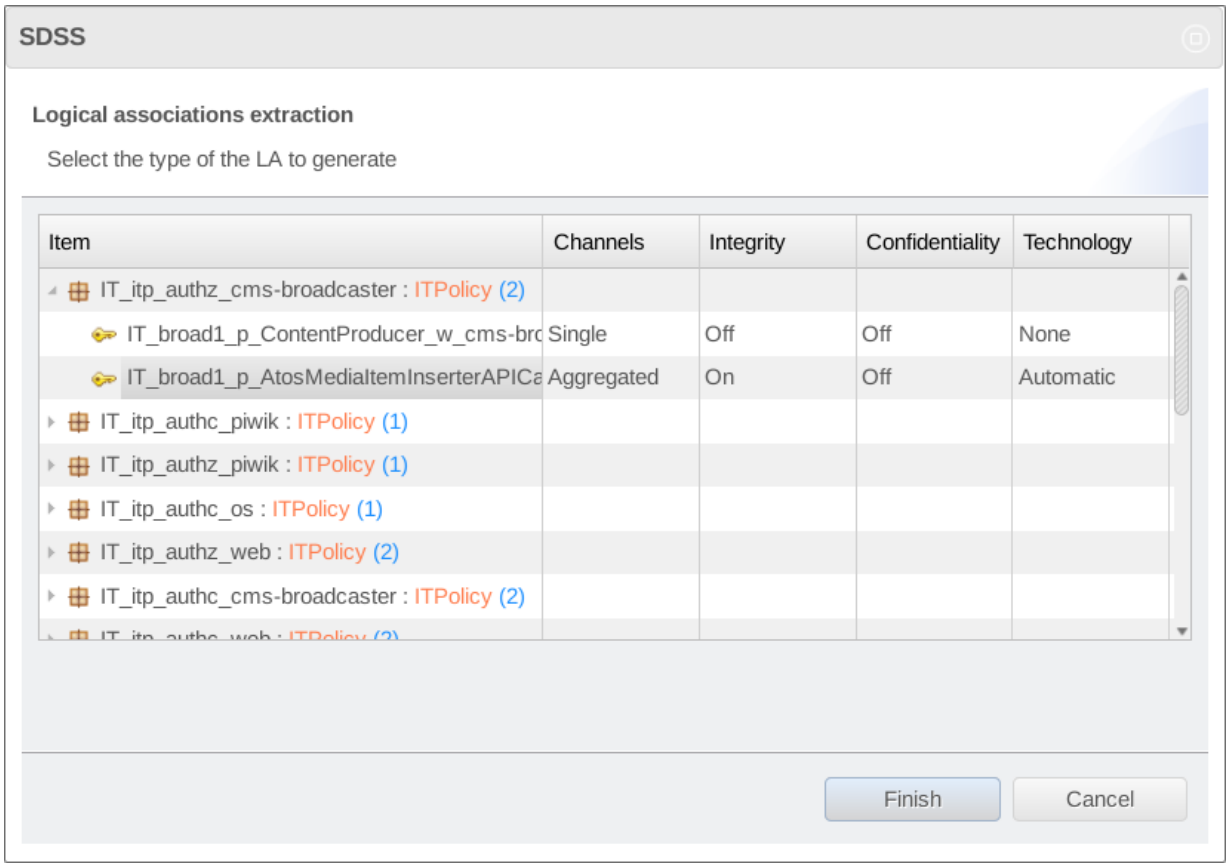

Figure 18: LA Type selection

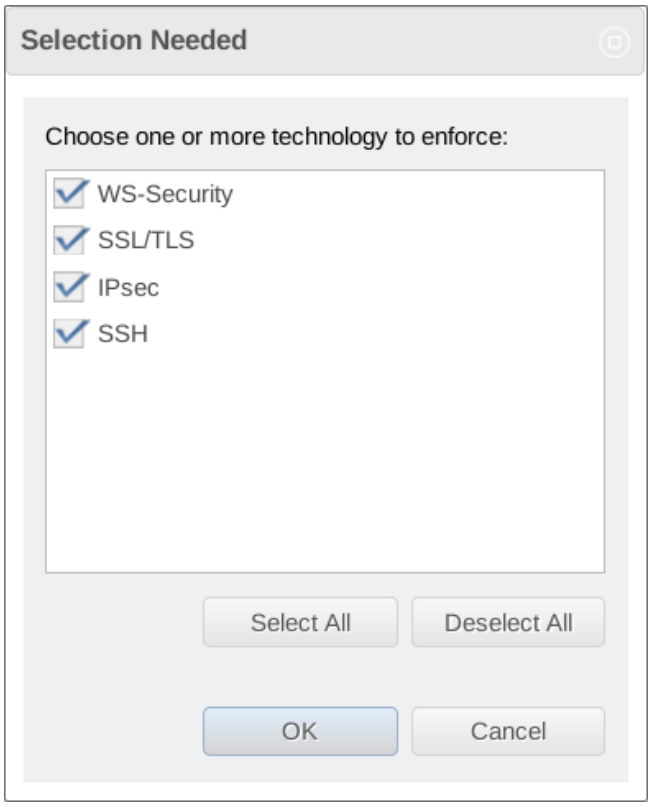

Figure 19: LA Technology selection

At the current stage of development, you can choose one of the following values:

- *None*: to specify no technology constraint;
- *WS-Security*: to use the WS-Security protocol to protect communications (this value can only be chosen for links);
- *SSL/TLS*: to use the SSL/TLS protocol;
- *IPsec*: to use the IPsec protocol;
- *SSH*: to use the SSH protocol.

Please note that possible values may be limited by previously executed enrichment modules or by the LA type you choose. For instance, if you don't specify integrity or confidentiality requirements, the only possible value will be *None*.

To proceed with the LA Generation process, you can click on the *Finish* button. The tool will process the chosen refinable items and, finally, it will show the generated logical associations in the *LA Explorer* view (Fig. 20).

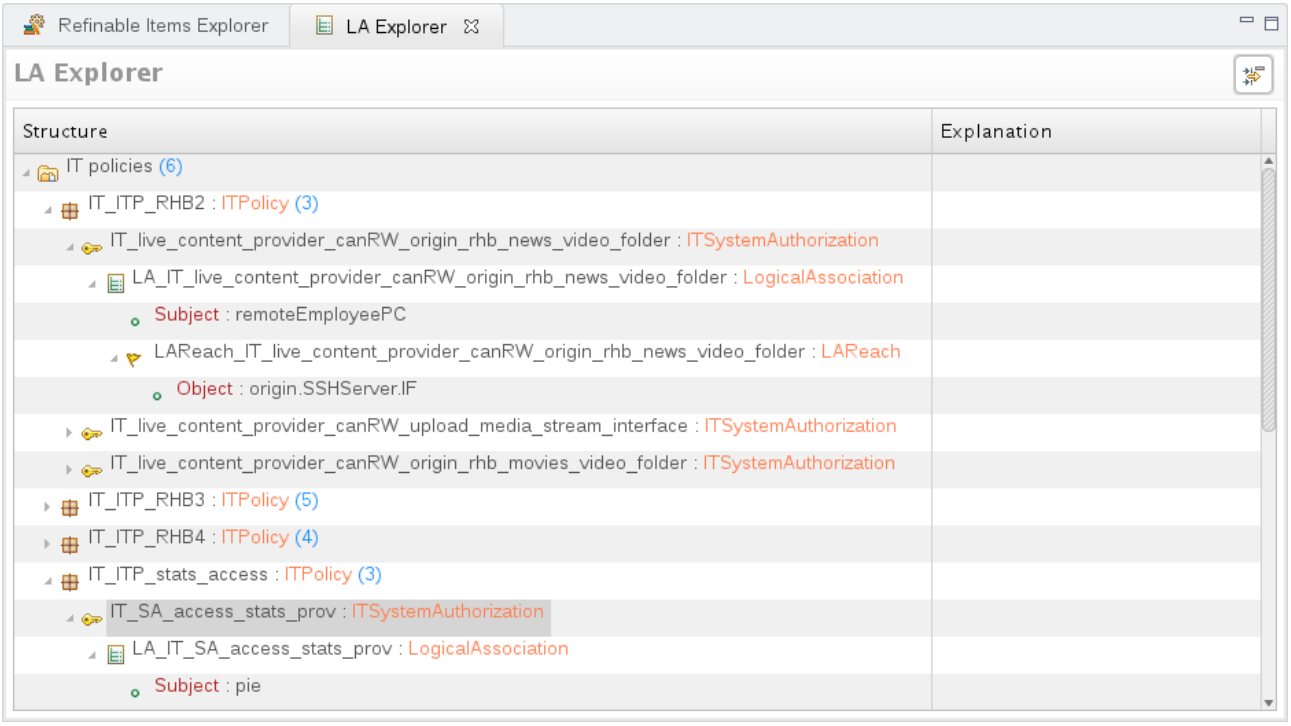

Figure 20: LA Explorer view

Logical associations are grouped according to the authorization rule or link from which they were generated. Furthermore, for every LA, the following information is displayed:

- the subject LA endpoint or, if you choose one of the aggregated LA types, the set of subjects;
- the privilege object or, also in this case, the set of objects for aggregated LAs;
- the additional protection property, if you specified a protected LA type;
- the selected security technology, if any.

If no logical association could be generated for a specific refinable item, the tool will display it with a red mark and an explanation message. This could happen, for instance, if a refinable item contains only authentication rules that are not handled by the LA Generation Service. If you want to view only the correctly refined items, you can click on the *Show only the refined items* button in the upper-right corner of the view. In this way, items that could not be refined will be hidden.

Other views

In addition to the previously presented views, the tool also offers some utility views that can be helpful for the user. The first one is the *Progress* view depicted in Fig. 21. It contains information regarding the execution state of the LA Generation process. More precisely, it shows the currently running phase, if any, and information about the execution time of other phases.

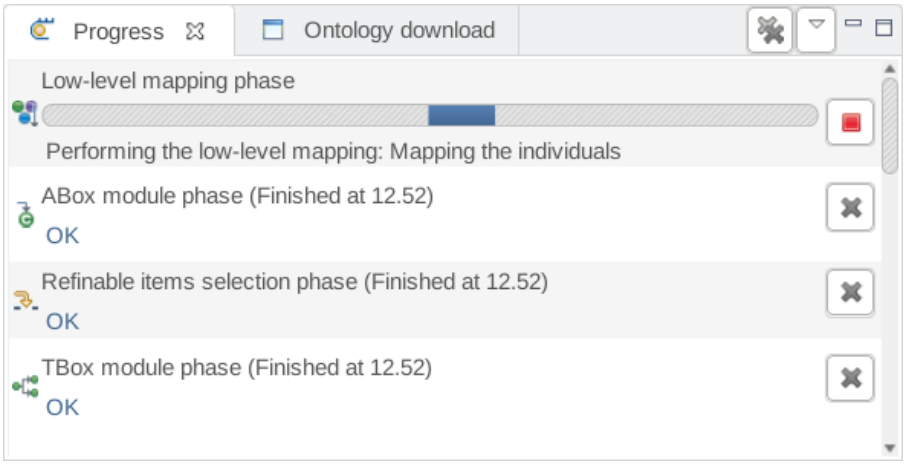

Figure 21: Progress view

The tool also offers the possibility to download the ontologies created during the LA Generation process. They contain the intermediate results produced and saved during the different workflow steps. These ontologies can be downloaded by means of the *Ontology download* view presented in Fig. 22.

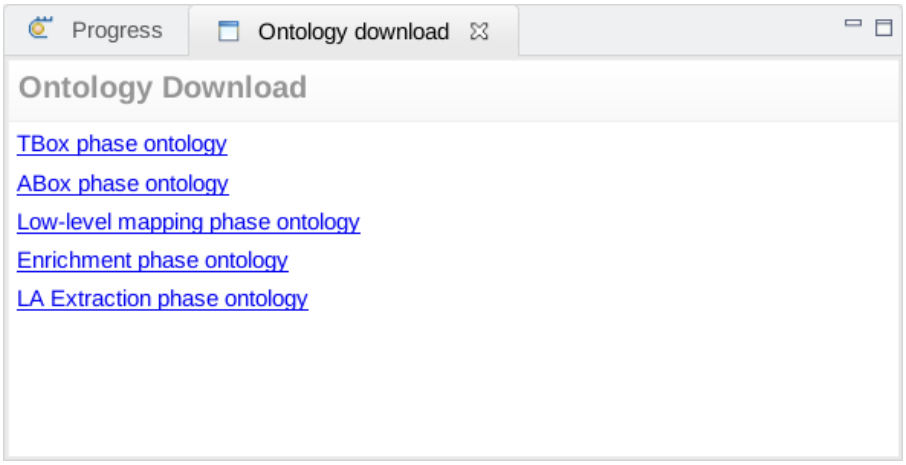

Figure 22: Ontology download view

This view contains a list of links that is dynamically populated during the LA Generation process. In general, at the end of the process, the view will contain the following links:

- *TBox phase ontology*: together with the *ABox phase* ontology, this ontology is created during the *Connect to the ontology* phase. It contains only the schema portion, i.e., classes and property definitions, of the PoSecCo data models;
- *ABox phase ontology*: like the previous one, this ontology is created during the *Connect to the ontology* phase. It contains the landscape and IT level policies gathered from the PoSecCo repository;
- *Low-level mapping phase ontology*: this ontology is created during the low-level mapping process. As such, it contains all mapping associations added during this phase;

- *Enrichment phase ontology*: this ontology is created during the enrichment phase and it contains all information inferred as a result of the EMs execution;
- *LA Extraction phase ontology*: this is the ontology corresponding to the final LA Generation phase and, as such, it contains all the generated logical associations.

To download one of the ontologies, simply click on the corresponding link and you will be presented with your browser's file download window. Downloaded ontologies are in RDF/XML format.

References

[1] The PoSecCo project. Infrastructure Configuration Service User Manual. http://security.polito.it/posecco/sdss/um-InfrastructureConfigurationService.pdf.