

Formal Verification of a Certified Policy Language

Amir Eaman
Amy Felty

School of Electrical Engineering and Computer Science (SITE)
University of Ottawa
Ottawa, Canada

Amir



Amy



Access Control is a security service that guards protected resources against unauthorized access (Granting or rejecting access to resources is a very important aspect of computer systems).

Policy Language is a language to express authorization policies (in the form of set of rules)

Type Enforcement An access control mechanism which exploits security context of resources to regulate accesses (subject-object set of rules).

Certified Language By certified policy language, we mean a policy language with formal semantics and mathematical proofs of important properties.

TEpla a new certified TE access control policy language called *TEpla*.

Access control policy languages, especially those that are widely used in practice, often do not have formal semantics. Inconsistencies and contradictions in the design of a language can lead to possibly serious unintended errors, especially as policies grow large.

Motivation: Analysis of Access-Control SELinux Policies

The motivation comes from our previous work [Eaman, Sistany, Felty, MCETech, 2017]

- For having a verified security policy, it is crucial to formally reason about the policy language in which the policy is written.
- The SELinux policy language requires third-party analysis tools to help security administrators write policies and check various properties.
- The inherent complexity of the SELinux policy language as well as its lack of formal semantics have led to the development of many policy analysis tools to try to translate SELinux policies to other intermediate language.
- There is no proof for the correctness of policy analysis tools to make sure their results are reliable.
- Overall SELinux lacks clarity as an access control language. The clarity of an access control policy language can provide better decision making for incremental policy writing, ease of analysis, and ease of reasoning.

Using a Certified TE Policy Language to Express TE policies

- Our solution to the challenges mentioned earlier is a certified Type Enforcement policy Language
- A small and certifiably correct Type Enforcement policy language can be a good candidate for SELinux style access control.
- TEpla is a certified policy language.
- The Coq Proof Assistant has been used to develop proofs for theorems
- We analyzed the behavior of the language by defining different ordering relations on policies, queries, and decisions.
- The behavior of the language are presented by a set of formal properties including *order preservation*, *independent composition*, *non-decreasing*, and *determinism* (defined in [Tschantz, Krishnamurthi, SACMAT, 2006]). A deterministic language always produces the same decision for the same policies and queries.
- This insight into language behavior provides a formal way to analyze and reason about language specifications, i.e., policies written in the language.

Outline

- 1 Syntax in Coq
- 2 Ordering Relations on Decisions, Policies, Queries
- 3 Semantics in Coq
- 4 Conditions on Predicates
- 5 An Example Predicate to Express Separation of Duty (SoD)
- 6 Formal Language Properties of TEpla
- 7 Future Work and Conclusion

Syntax in Coq

The Coq Proof Assistant

- A proof assistant to develop machine-checked proofs
- Often used to verify the correctness of programs: Certified programs
- Properties and proofs are formalized in a special language
- Proofs are developed in a semi-interactive manner dependent on human guidance.
- A further strength of Coq is that executable programs can also be extracted from constructive proofs.



TEpla Syntax and its Encoding in Coq

```
Definition C := N.
Definition P := N.
Definition basicT := N.
(* examples *)
Definition File : C := 600.
Definition mail_t : basicT := 300.
Definition Read : P := 702.
Definition networkManager_ssh_t : basicT := 302.

Definition G : Set := list (basicT).
Inductive T : Type :=
| singleT : basicT → T
| groupT : G → T.

(* examples *)
Definition program_G : G := [mail_t;http_t].
Definition program_T : groupT program_G.
```


TEpla Syntax and its Encoding in Coq (contd.)

```
Inductive R : Set :=
  | Allow : T * T * C * P * B → R
  | Type_Transition : T * T * C → R.

(* example *)
Definition R_A : R :=
  Allow (groupT program_G, singleT mail_t, File, Read, true).

Inductive CSTE: Set :=
  | Constraint : C * P * T * T * list T *
    (list R → list T → C → P →
     T → T → T → T → B) → CSTE.

(* example *)
Definition CSTE_SoD : CSTE := Constraint(File, Read, groupT program_G,
  singleT networkManager_ssh_t, [], Prd_SoD).
```

TEpla Syntax and its Encoding in Coq (contd.)

```
Inductive TEPLCY: Set :=
  | TEPolicy : list R * list CSTE → TEPLCY.

Definition Q : Set :=
  T * T * C * P.
(* example *)
Definition sampleQ : Q := (singleT mail_t, singleT http_t, File, Write).

Inductive DCS: Set :=
  | Permitted | NotPermitted | UnKnown.
```

Ordering Relations on Decisions, Policies, Queries

Ordering Relations (poset)

(TEPLCY, \lesssim) We define the binary relation \lesssim on two policies p_1, p_2 , where $p_1 \lesssim p_2$ whenever p_2 has more information than p_1 . More formally:

$$\forall (p_1, p_2 \in \text{TEPLCY}), p_1 \lesssim p_2 \text{ iff } \text{length}(p_1) \leq \text{length}(p_2) \wedge p_1 \subseteq p_2.$$

(Q, $\ll =$) In TEpla, Two queries $Q_1 = (\text{SourceT_}Q_1, \text{DestT_}Q_1, C_1, P_1)$ and $Q_2 = (\text{SourceT_}Q_2, \text{DestT_}Q_2, C_2, P_2)$ are in relation $Q_1 \ll = Q_2$ if and only if $(\text{TSubset SourceT_}Q_2 \text{ SourceT_}Q_1)$ and $(\text{TSubset DestT_}Q_2 \text{ DestT_}Q_1)$ hold.

(DCS, $\langle ::$) We define the binary relation " $\langle ::$ " on this three element set as *NotPermitted* $\langle ::$ *Permitted* $\langle ::$ *UnKnown* to define the poset (DCS, $\langle ::$)

The *UnKnown* decision arises from conflicts in policies. The semantics of TEpla is a *homomorphism* on the posets we defined on TEPLCY, Q and DCS.

Semantics in Coq

Semantics in Coq - Evaluating Queries Against Policy Rules

Evaluating a query against single rule:

```
Definition R_EvalTE (R_policy:R) (q:Q) : DCS:=
  match R_policy with
  | Allow (alw_srcT, alw_dstT, alw_C, alw_P, alw_B) =>
    match q with
    |(qsrcT, qdsT, qC, qP) =>
      if ((TSubset qsrcT alw_srcT) && (TSubset qdsT alw_dstT) &&
          (Nat.eqb qC alw_C) && (Nat.eqb qP alw_P) && (alw_B))
      then Permitted else NotPermitted
    end
  | Type_Transition (trn_srcT, trn_dstT, trn_C) => ...
end.
```

Evaluation of a query against a list of rules:

```
Fixpoint listR_EvalTE
  (listR: list R) (qry: Q) : DCS:=
  match listR with
  | rule_h :: rule_body =>
      maximalDCS (R_EvalTE rule_h qry)
                  (listR_EvalTE rule_body qry)
  | [] => NotPermitted
  end.
```

Semantics in Coq (contd.)-Evaluating Queries against Constraints

After evaluating a query against the list of rules in the policy, we evaluate the query against constraints. This function evaluates a query against a constraint.

```
Definition CSTE_EvalTE
  (constraint_rule : CSTE) (Q_to_constr : Q) (listR : list R) : DCS :=
  match constraint_rule with
  | Constraint (cstrn_C, cstrn_P, cstrn_T_arg1, cstrn_T_arg2,
               cstrn_listT, cstrn_PRDT) =>
    match Q_to_constr with
    |(Q_srcT, Q_dstT, Q_C, Q_P) =>
      if (Nat.eqb Q_C cstrn_C && Nat.eqb Q_P cstrn_P) then
        match (cstrn_PRDT listR cstrn_listT cstrn_C cstrn_P
              Q_srcT Q_dstT cstrn_T_arg1 cstrn_T_arg2) with
        | true => Permitted
        | false => UnKnown
        end
      else NotPermitted
    end
  end.
```


Semantics in Coq (contd.) - Finding the decisions of policies

This function takes two decisions as arguments and applies the logic of combining the decisions of the two components of policies.

```
Definition maximalpolicy_DCS
  (compont_R compont_CSTE : DCS) : DCS :=
  if (Permitted <:: compont_R) then
    maximalDCS compont_R compont_CSTE
  else NotPermitted.
```

Semantics in Coq (contd.) - Evaluating a query against a policy

The functions for evaluating a query against a set of rules and a set of constraints and calling the helper function `maximalpolicy_DCS` to combine the results.

```
Fixpoint TEPLCY_EvalTE
  (policy: TEPLCY) (q: Q) : DCS:=
  match policy with
  | TEPolicy (CompOne_RList, CompTwo_CSTEList) =>
    maximalpolicy_DCS
      (listR_EvalTE CompOne_RList q)
      (listCSTE_EvalTE CompTwo_CSTEList q CompOne_RList) end.
```

Conditions on Predicates

Conditions on Predicates


Predicates defined by policy writers must satisfy three conditions and they must prove these conditions are met. The conditions are:

- 1 `Predicate_Q_condition`: Evaluating predicates given any two queries Q_1 and Q_2 such that $Q_1 \ll = Q_2$, if d_1 and d_2 are the decisions resulting from evaluating Q_1 and Q_2 , respectively, then $d_1 <:: d_2$
- 2 `Predicate_plc_cdn`: states that the same result is obtained from applying the predicate on the two lists of rules, whenever the two lists differ only in the order of the rules.
- 3 `Predicate_plc_cdn_Transition`: given a query Q and two rule lists $\text{list}\mathbb{R}$, and $\text{list}\mathbb{R} ++ \text{list}\mathbb{R}'$, if d_1 and d_2 are the decisions resulting from evaluating Q against the list of constraints, respectively, then $d_1 <:: d_2$

An Example Predicate to Express Separation of Duty (SoD)

Encoding SoD security goal

```
Fixpoint Prd_SoD (listR:list R) (ListT:list T) (sClass:C) (perm:P)
  (QSrcT:T) (QDesT:T) (PRDTsrcT:T) (PRDTDesT:T) : B:=
  if (TSubset QSrcT PRDTsrcT && TSubset QDesT PRDTDesT)
  then is_emptylistT ( IntersectionList
    (listRSearch_subjectTs listR PRDTsrcT)
    (listRSearch_subjectTs listR PRDTDesT))
  else true.
```



Verifying the predicate Prd_SoD

Verifying the predicate Prd_SoD satisfies the three conditions on predications:

```
Lemma qry_condition_SoD PRDT:  
  Predicate_Q_condition Prd_SoD.
```

```
Lemma plc_conditionS_SoD PRDT:  
  Predicate_plc_cdn Prd_SoD.
```

```
Lemma plc_conditionF_SoD PRDT:  
  Predicate_plc_cdn_Transition Prd_SoD.
```

Formal Language Properties of TEpla

Order Preservation of TEpla Queries

Order Preservation

Of particular importance is the preservation of order on decisions with respect to queries: if $q_1 \ll = q_2$, then the decisions d_1 and d_2 that result applying function `TEPLCY_EvalTE` on q_1 and q_2 , respectively, are in the relation $d_1 <:: d_2$.

Theorem `Order_Preservation_TEpla :`

```
∀ (listR:list R) (listCSTE:list CSTE) (q q' : Q),
  (q <<= q') ∧ const_imp_prd_List listCSTE →
  ((TEPLCY_EvalTE (TEPLCY (listR, listCSTE)) q) <::
   (TEPLCY_EvalTE (TEPLCY (listR, listCSTE)) q')) = true.
```

Non-Decreasing Property of TEpla Policies

Non-Decreasing

This theorem states that adding a policy `Single_pol`, to any list of policies `Pol_list` can change the decisions only according to the order relation $<::$ on decisions. The \oplus operator extracts the rule lists of all the policies in its argument list of policies and combines them into one list, and similarly for constraints, forming a single policy from these rules and constraints.

```
Theorem Non_Decreasing_TEpla :  
∀ (Pol_list: list TEPLCY) (Single_pol:TEPLCY) (q:Q) (d d': DCS),  
validCnstrtListPolicy Pol_list ∧ validConstrt Single_pol →  
  (TEPLCY_EvalTE (⊕ (Pol_list)) q) = d →  
  (TEPLCY_EvalTE (⊕ (Single_pol::Pol_list)) q) = d' →  
  (d <:: d') = true.
```

Independent Composition of TEpla Policies

Independent Composition

The independent composition theorem states that whenever a pair of lists satisfies this property, then the decision obtained by evaluating the combined policy on q is the maximum of the decisions resulting from evaluating each policy independently.

```
Theorem Independent_Composition :  
∀ (PLCY_DCS_pair : list (TEPLCY * DCS)) (q : Q) (dstar : DCS),  
  Foreach q (map fst PLCY_DCS_pair) (map snd PLCY_DCS_pair) ∧  
  (TEPLCY_EvalTE (⊕ (map fst PLCY_DCS_pair)) q) = dstar →  
  (maximum (map snd PLCY_DCS_pair) <:: dstar) = true.
```

Future Work and Conclusion

Conclusion and Future Work

- The infrastructure as well as formal properties of TEpla is encoded in the Coq proof assistant.
- TEpla is certified in terms of formal semantics and machine-checked proofs of a particular set of properties.
- TEpla provides the language constructs for allowing security administrators to encode different security goals in policies.
- TEpla is a certified solution for challenges and drawbacks that security administrators face in developing or analyzing security policies.
- There are some limitations in the language, such as the limited number of arguments for constraints
- The certification of TEpla done so far will help with the development of the next versions of TEpla.
- Program extraction is possible to OCaml, Scheme, Haskell

Thank you for your attention. Questions?

Links:

- TEpla Coq Code: [TEpla Coq code](#)
- Modeling and Analysis of Access-Control Policies (SELinux):
[Modeling and Analysis of access-control policies](#)

References:

- Eaman, A.: TEpla: A Certified Type Enforcement Access Control Policy Language., Ph.D. thesis, University of Ottawa (2019), [TEpla: A Certified TE Policy Language](#)
- Eaman, A., Sistany, B., Felty, A.: Review of existing analysis tools for SELinux security policies: Challenges and a proposed solution. In: 7th International Multi-disciplinary Conference on e-Technologies (MCETECH). pp. 116–135 (2017) [SELinux Policy Challenges](#)