Vac - Verifier of Administrative Role-Based Access Control Policies

Anna Lisa Ferrara[1], P. Madhusudan[2], Truc L. Nguyen[3], and Gennaro Parlato[3]

[1] University of Bristol, UK
[2] University of Illinois, USA
[3] University of Southampton, UK

Abstract. In this paper we present Vac, an automatic tool for verifying security properties of administrative Role-based Access Control (RBAC). RBAC has become an increasingly popular access control model, particularly suitable for large organizations, and it is implemented in several software. Automatic security analysis of administrative RBAC systems is recognized as an important problem, as an analysis tool can help designers check whether their policies meet expected security properties. Vac converts administrative RBAC policies to imperative programs that simulate the policies both precisely and abstractly and supports several automatic verification back-ends to analyze the resulting programs. In this paper, we describe the architecture of Vac and overview the analysis techniques that have been implemented in the tool. We also report on experiments with several benchmarks from the literature.

1 Introduction

Access control models allow to restrict access to shared resources by selectively assigning permissions to users. Role-based Access Control (RBAC) has become an increasingly popular access control model [5], it is standardized by NIST and is implemented in several software, such as Microsoft SQL Servers, Microsoft Active Directory, SELinux, and Oracle DBMS. RBAC reduces the complexity of user permissions administration by grouping users into roles and assigning permissions to each role. An *Administrative* RBAC User-Role Assignment (ARBAC-URA) policy defines a set of administrative roles and rules which specify how administrators *can assign* or *can revoke* roles to users [27].

Automatic security analysis of ARBAC systems is recognized as an important problem, as an analysis tool can help designers check whether their policies meet the expected security properties [25]. This is particularly desirable whenever policies need to be correct by design, for instance when accesses are not mediated by a monitor [6]. Most interesting security properties, such as privilege escalation and separation of duties, can be phrased as the *role-reachability problem* (i.e., is there a reachable configuration where some user can eventually be assigned to a target role?) [19, 7]. The role-reachability problem is known to be PSPACE-complete and hard to solve on real-world policies having hundreds of roles and rules and thousands of users [28].
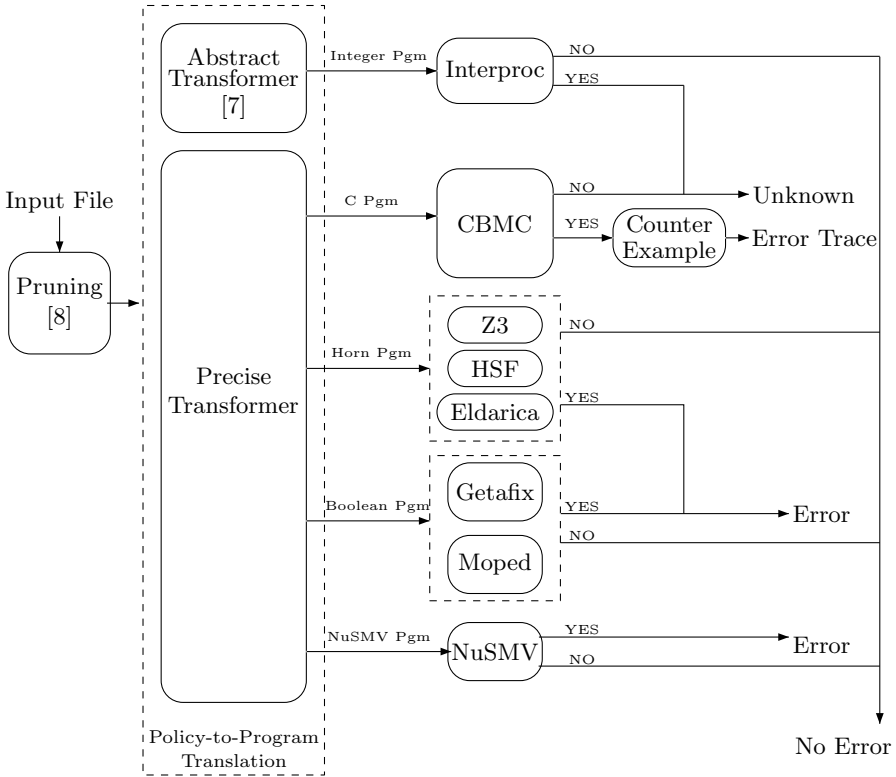
Fig. 1. Vac's Architecture

In this paper, we present Vac (*Verifier of Access Control*), an automatic and scalable tool for solving the role-reachability problem of ARBAC-URA policies. The main components of Vac are a pruning module that aims at simplifying the state space by reducing a policy to a smaller one that preserves the reachability of the target role, and a policy-to-program translation module that converts a policy to an imperative program that simulates the policy both precisely and abstractly. Vac supports a plethora of automatic verification back-ends for the analysis of the resulting programs and has a built-in counterexample generator.

In the rest of the paper, we describe the architecture of Vac and overview the analysis techniques implemented in the tool. Finally, we present experimental results showing the effectiveness of Vac on analyzing realistic and complex benchmarks from the literature.

2   Software Architecture and Verification Approaches

The high-level architecture of Vac is shown in Fig. 1. We first describe the input format of Vac and then its components.

Input Format. We refer to [8] for the syntax and the semantics of an ARBAC-URA system. To illustrate the input format of VAC[1] we use the toy example of an hospital policy shown in Fig. 2.

ROLES EMPLOYEE DOCTOR MANAGER PATIENT;
USERS ANNA LUKE STEVE LUCY;
UA ⟨ANNA, DOCTOR⟩ ⟨LUCY, MANAGER⟩ ⟨LUKE, DOCTOR⟩ ⟨STEVE, PATIENT⟩;
CR ⟨DOCTOR, PATIENT⟩ ⟨DOCTOR, MANAGER⟩;
CA ⟨DOCTOR, EMPLOYEE & -DOCTOR, MANAGER⟩ ⟨DOCTOR, TRUE, PATIENT⟩;
ADMIN ANNA LUKE;
SPEC DOCTOR;

Fig. 2. A VAC's input file

ROLES and **USERS** are keywords used to list roles and users, respectively. **UA** defines the initial user-role assignment, whereas **CR** and **CA** specify the administrative rules *Can-Revoke* and *Can-Assign*, respectively. A **CR** rule is a pair of roles. For instance, the **CR** rule ⟨ DOCTOR, PATIENT ⟩ says that any administrator with role DOCTOR can revoke the role PATIENT from any user. A **CA** rule also contains a precondition, that is, a Boolean formula written as a conjunction of literals over roles. For instance, the **CA** rule ⟨ DOC-TOR, EMPLOYEE & -DOCTOR, MANAGER ⟩ says that any administrator with role DOCTOR can assign any user u to the role MANAGER provided that u is member of EMPLOYEE and not a member of DOCTOR. The keyword **ADMIN** is used to list all users that are also administrators[2]. Finally, the keyword **SPEC** is used to specify the role-reachability query, namely the target role. VAC can also be used to check whether the target role is reachable by a specific user. For instance, the query **SPEC** LUCY DOCTOR allows to check whether LUCY can ever obtain role DOCTOR.

Pruning Module. This module takes as input a policy, which we refer as the original policy, and outputs a simplified one (in the same format as the input) that preserves the reachability of the target role. The module implements the pruning heuristic from [8] which is crucial for scalability. It eliminates roles, rules and users with the aim of reducing the state space to explore. This heuristic relies upon a fundamental theorem which states that the role-reachability problem can be solved by tracking only $k+1$ users, where k is the number of administrative roles [8]. Thus, the heuristic exploits sufficient conditions to eliminate administrative roles that are not relevant for the analysis. The effectiveness of the above method is amplified by a static pruning algorithm consisting of six pruning actions: the first three aim at discarding roles that are irrelevant to the reachability of the target role while the remaining ones identify administrative rules that can be combined or eliminated. Furthermore, whenever the target role is reachable within two steps in the intermediate policy, the pruning procedure terminates immediately returning the counterexample.

Policy-to-Program Translation. This module takes as input a policy and translates it into a program that simulates the evolution of the system. VAC provides the following two policy-to-program translations:

---

[1] VAC's input format is compatible with that of Mohawk [16].

[2] The list of administrators can be obtained from **CA** and **CR** rules. However, we include the keyword **ADMIN** to be consistent with Mohawk's input format [16].

Abstract Transformer. This module implements the policy-to-program translation proposed in [7]. A policy P is translated into an imperative non-deterministic while-program P' with an error location. P' uses only integer variables to capture the number of users in a subset of role combinations and abstractly simulates the evolution of the system in such a way that if the error location is not reachable then the role reachability problem has a negative answer on P. On the contrary, if the error location is reachable in P', this may correspond to a false positive as P' over-approximates the behaviour of P. P' is then analyzed by VAC using Interproc [18] with the box abstract domain.

Precise Transformer. This module translates a policy P into a Boolean program P' that precisely simulates the evolution of the system tracking at most $k+1$ users picked non-deterministically, where k is the number of administrative roles. The correctness of this approach relies on a fundamental theorem proven in [8]. The program uses $k+1$ blocks of n Boolean variables, where n is the number of roles in the policy. Each block tracks the role-membership of a selected user. The rest of the program consists of an infinite loop in which the administrative rules are non-deterministically simulated on a non-deterministically chosen user. The loop contains also an error location that is reachable whenever a tracked user reaches the target role. The role-reachability problem admits a positive answer on P if and only if the error location is reachable in P'. The reachability problem for Boolean programs is decidable and VAC supports several automated tools as back-ends for the analysis of P'. In particular, a complete analysis can be performed by using either (1) one of the following tools for Horn clauses: Z3 (μZ) [4, 13], HSF [11, 12], and Eldarica [15, 14], or (2) Moped [20, 29] and Getafix [23, 24] which are model checkers for Boolean Programs based on BDDs, or (3) NuSMV a model checker based on BDDs and SAT solvers [2, 3]. VAC uses the C bounded model checker CBMC [21, 22] for under-approximate analysis, particularly effective to find errors. If CBMC finds an error, it returns a counterexample showing how the error location can be reached in the program. Otherwise, VAC reports Unknown.

Counterexample Module. VAC implements an involved built-in counterexample generation module that takes as input the counterexample of the pruned policy returned by CBMC along with some information collected during the execution of the pruning, and outputs a counterexample (attack) of the original policy.

3   Implementation and Availability

Implementation. VAC is implemented in C and has dependencies with ANTLR (v3.2 for C), ROXML, and CCL libraries[3].

Availability. The source code, a set of benchmarks and static Linux binaries are available at: `http://users.ecs.soton.ac.uk/gp4/VAC`.

Usage. The shell command `./vac.sh <InputFile>` runs VAC with the default setting: (1) runs the abstract transformer and Interproc to prove correctness;

---

[3] ANTLR, ROXML and CCL are respectively available at `http://www.antlr.org/`, `http://www.libroxml.net`, and `https://code.google.com/p/ccl/`.

Table 1. VAC's results on realistic case studies

	name	\#roles	\#rules	\#admin	\#users	\#roles	\#rules	\#admin	\#users	Time	Answer	Time
		ARBAC Policy				Pruning					Reach	
1	Hospital1	13	37	5	1092	4	5	3	6	0.009s	No	0.029s
	Hospital2	13	37	5	1092	4	5	3	6	0.009s	No	0.023s
	Hospital3	13	37	5	1092	3	2	1	4	0.009s	Yes	0.103s
	Hospital4	13	37	5	1092	4	4	1	4	0.009s	Yes	0.110s
2	University1	32	449	9	943	6	7	3	13	0.009s	No	0.034s
	University2	32	449	9	943	6	8	3	13	0.004s	Yes	0.192s
	University3	32	449	9	943	4	5	1	6	0.006s	No	0.021s
	University4	32	449	9	943	12	37	4	31	0.004s	Yes	1.571s
3	Bank1	343	2225	1	2	3	2	1	2	0.007s	Yes	0.112s
	Bank2	683	4445	1	2	3	2	1	2	0.019s	Yes	0.139s
	Bank3	1023	6665	1	2	3	2	1	2	0.024s	Yes	0.167s
	Bank4	1363	8885	1	2	3	2	1	2	0.030s	Yes	0.168s
	Bank5	343	2225	1	2	3	2	1	2	0.044s	Yes	0.138s
	Bank6	683	4445	1	2	3	2	1	2	0.155s	Yes	0.247s
	Bank7	1023	6665	1	2	3	2	1	2	0.300s	Yes	0.435s
	Bank8	1363	8885	1	2	3	2	1	2	0.522s	Yes	0.663s
	Bank9	531	5126	1	2000	2	0	1	2	0.244s	No	0.253s
	Bank10	531	5126	1	2000	2	0	1	2	0.248s	No	0.254s
	Bank11	531	5126	1	2000	3	2	1	2	0.245s	Yes	0.396s
	Bank12	531	5126	1	2000	6	5	1	2	0.066s	Yes	0.223s

(2) if a proof cannot be provided, VAC runs the precise transformer and CBMC (with `unwind` set to 2) to find a counterexample; (3) if CBMC does not find an error, VAC runs μZ for complete analysis. VAC has options to print the translated programs and the simplified policies, and select the back-end for the analysis.

4   Experimental Results

We evaluate VAC, using the default setting, on several benchmarks from the literature. All experiments have been performed on a Linux 64-bit machine with Intel Core i7-3770 CPU and 16GB of RAM.

Table 1 shows the results on three sets of benchmarks based on realistic case studies. The first two case studies are carried out by Stoller at al. [30] and represent policies for a university and for an hospital, respectively. The third case study, conducted by Jayaraman et al. [16], models a bank with several branches[4]. While the first eight bank policies are from [16], we have built the last four from [16] by slightly modifying their policies to add more users and to make two of them correct. Table 1 reports the number of roles, rules, administrative roles and users of both the original policy and that after pruning. It also reports the tool's answer, the time taken by the pruning, and the overall analysis time.

Table 1 shows that the pruning module significantly reduces the size of these policies. Furthermore, VAC is extremely efficient in verifying these policies, regardless of whether the target role is reachable or not. More precisely, all benchmarks with a negative answer can be proved correct in less than a second. Similarly, on benchmarks with a reachable target the analysis takes less than 2 seconds including the time to generate the counterexample.

---

[4] The number of roles and rules depends on the number of branches considered. For instance, 343 roles corresponds to 10 branches and 1363 to 40 branches.

Table 2. Vac' s results on complex test suites

Size Policy		Vac								
		First Suite			Second Suite			Third Suite		
		Pruning		Verification	Pruning		Verification	Pruning		Verification
#roles	#rules	#roles	#rules	Time	#roles	#rules	Time	#roles	#rules	Time
4	10	3	1	0.080s	3	1	0.084s	2	1	0.085s
5	25	4	2	0.087s	4	2	0.096s	2	1	0.092s
20	100	3	1	0.099s	3	1	0.089s	3	2	0.087s
40	200	4	2	0.099s	4	2	0.096s	2	1	0.091s
200	1000	2	1	0.101s	2	1	0.088s	2	1	0.096s
500	2500	3	1	0.100s	3	1	0.104s	3	2	0.128s
4000	20000	2	1	0.239s	2	1	0.198s	4	3	0.252s
20000	80000	2	1	0.844s	2	1	0.579s	3	2	0.922s
30000	120000	2	1	1.288s	2	1	0.849s	2	1	1.285s
40000	200000	2	1	1.586s	2	1	1.100s	4	3	1.646s

Table 2 shows the results on three sets of complex test suites, synthetically generated by Jayaraman et al. [16], with the aim of capturing the complexity of real systems. Each suite consists of ten policies where the number of roles and rules ranges respectively from 4 to 40k and 10 to 200k. The role-reachability problem has a positive answer on all these benchmarks. Vac is very effective on these policies as well. The analysis takes less than 2 seconds on all policies and the pruning module reduces the policies to equivalent systems with a handful of roles and rules.

5   Conclusions

We have presented Vac an automatic and efficient tool for verifying security properties of administrative role-based access control policies. The main components of Vac are a pruning module which is essential for scalability, and a policy-to-program translation module that reduces the role-reachability problem to program verification problems. It supports several tools for the analysis, such as CBMC, Eldarica, Getafix, Interproc, Moped, NuSMV, HSF, and Z3 (μZ). Furthermore, it can provide counterexamples.

Related Work. Among the state-of-the-art tools for the analysis of ARBAC-URA systems, Vac is the only tool that simultaneously has the following features: (1) complete analysis (2) counterexample generation, and (3) scalable analysis on large policies. Mohawk [16] performs only under-approximate analysis, though it now considers thresholds for completeness [17]; RBAC-PAT [10] is unable to handle large policies. They also can only analyze policies with *separate administration* where administrators cannot change their role-membership; this is not realistic, but simplifies analysis as only a single user needs to be tracked.

ASASPXL is the latest tool developed by Ranise et al. for the analysis of ARBAC policies [26]. A previous version (ASASP [1]) was not able to scale on large policies. ASASPXL is mainly designed to handle large policies and does so by encoding the instances to MCMT [9] which is a model checker for infinite state systems based on SMT solvers and backward reachability. In contrast, Vac does not target any specific kind of instances, and handles large policies by carrying out an effective pruning that is independent of the verification technique used

for the analysis. VAC and ASASPXL can potentially handle the same kind of instances though they have different input formats.

All tools above do not generate counterexamples. Furthermore, VAC, on the policies of Section 4, has either the same performances or outperforms the tools mentioned above. VAC has also been used for the analysis of *temporal* RBAC [31].

References

[1] Alberti, F., Armando, A., Ranise, S.: ASASP: Automated Symbolic Analysis of Security Policies. In: Bjørner, N., Sofronie-Stokkermans, V. (eds.) CADE 2011. LNCS, vol. 6803, pp. 26–33. Springer, Heidelberg (2011)

[2] Cimatti, A., Clarke, E.M., Giunchiglia, E., Giunchiglia, F., Pistore, M., Roveri, M., Sebastiani, R., Tacchella, A.: NuSMV: A New Symbolic Model Checker, http://nusmv.fbk.eu

[3] Cimatti, A., Clarke, E.M., Giunchiglia, E., Giunchiglia, F., Pistore, M., Roveri, M., Sebastiani, R., Tacchella, A.: NuSMV 2: An OpenSource Tool for Symbolic Model Checking. In: Brinksma, E., Larsen, K.G. (eds.) CAV 2002. LNCS, vol. 2404, pp. 359–364. Springer, Heidelberg (2002)

[4] de Moura, L., Berdine, J., Bjorner, N.: Z3 High-performance Theorem Prover, http://z3.codeplex.com

[5] Ferraiolo, D., Kuhn, R.: Role-Based Access Control. In: 15th NIST-NCSC National Computer Security Conference, pp. 554–563. Springer (1992)

[6] Ferrara, A.L., Fuchsbauer, G., Warinschi, B.: Cryptographically Enforced RBAC. In: CSF, pp. 115–129. IEEE (2013)

[7] Ferrara, A.L., Madhusudan, P., Parlato, G.: Security Analysis of Role-Based Access Control through Program Verification. In: CSF, pp. 113–125 (2012)

[8] Ferrara, A.L., Madhusudan, P., Parlato, G.: Policy Analysis for Self-administrated Role-Based Access Control. In: Piterman, N., Smolka, S.A. (eds.) TACAS 2013 (ETAPS 2013). LNCS, vol. 7795, pp. 432–447. Springer, Heidelberg (2013)

[9] Ghilardi, S., Ranise, S.: MCMT: A Model Checker Modulo Theories. In: Giesl, J., Hähnle, R. (eds.) IJCAR 2010. LNCS, vol. 6173, pp. 22–29. Springer, Heidelberg (2010)

[10] Gofman, M.I., Luo, R., Solomon, A.C., Zhang, Y., Yang, P., Stoller, S.D.: RBAC-PAT: A Policy Analysis Tool for Role Based Access Control. In: Kowalewski, S., Philippou, A. (eds.) TACAS 2009. LNCS, vol. 5505, pp. 46–49. Springer, Heidelberg (2009)

[11] Grebenshchikov, S., Gupta, A., Lopes, N.P., Popeea, C., Rybalchenko, A.: HSF(C): A Software Verifier based on Horn Clauses, http://www7.in.tum.de/tools/hsf

[12] Grebenshchikov, S., Gupta, A., Lopes, N.P., Popeea, C., Rybalchenko, A.: HSF(C): A Software Verifier Based on Horn Clauses. In: Flanagan, C., König, B. (eds.) TACAS 2012. LNCS, vol. 7214, pp. 549–551. Springer, Heidelberg (2012)

[13] Hoder, K., Bjørner, N., de Moura, L.: μZ– An Efficient Engine for Fixed Points with Constraints. In: Gopalakrishnan, G., Qadeer, S. (eds.) CAV 2011. LNCS, vol. 6806, pp. 457–462. Springer, Heidelberg (2011)

[14] Hojjat, H., Konečný, F., Garnier, F., Iosif, R., Kuncak, V., Rümmer, P.: A Verification Toolkit for Numerical Transition Systems. In: Giannakopoulou, D., Méry, D. (eds.) FM 2012. LNCS, vol. 7436, pp. 247–251. Springer, Heidelberg (2012)

[15] Hojjat, H., Rümmer, P., Konecny, F.: A Predicate Abstraction Engine, http://lara.epfl.ch/w/eldarica

[16] Jayaraman, K., Ganesh, V., Tripunitara, M.V., Rinard, M.C., Chapin, S.J.: Automatic Error Finding in Access-Control Policies. In: CCS, pp. 163–174 (2011)

[17] Jayaraman, K., Tripunitara, M.V., Ganesh, V., Rinard, M.C., Chapin, S.J.: Mohawk: Abstraction-Refinement and Bound-Estimation for Verifying Access Control Policies. ACM Trans. Inf. Syst. Secur. 15(4), 18 (2013)

[18] Jeannet, B., Lalire, G., Argoud, M.: The Interproc Analyzer, http://pop-art.inrialpes.fr/interproc/interprocweb.cgi

[19] Jha, S., Li, N., Tripunitara, M., Wang, Q., Winsborough, W.: Towards Formal Verification of Role-Based Access Control Policies. IEEE Transactions on Dependable and Secure Computing 5(4), 242–255 (2008)

[20] Kiefer, S., Schwoon, S., Suwimonteerabuth, D.: A Model Checker for Pushdown Systems, http://www2.informatik.uni-stuttgart.de/fmi/szs/tools/moped

[21] Kroening, D., Clarke, E.: CBMC - Bounded Model Checking for ANSI-C, http://www.cprover.org/cbmc

[22] Kroening, D., Tautschnig, M.: CBMC – C Bounded Model Checker - (Competition Contribution). In: Ábrahám, E., Havelund, K. (eds.) TACAS 2014. LNCS, vol. 8413, pp. 389–391. Springer, Heidelberg (2014)

[23] La Torre, S., Madhusudan, P., Parlato, G.: Getafix: A Symbolic Model-checker for Recursive Programs, http://www.cs.uiuc.edu/~madhu/getafix

[24] La Torre, S., Madhusudan, P., Parlato, G.: Analyzing Recursive Programs using a Fixed-point Calculus. In: Hind, M., Diwan, A. (eds.) PLDI, pp. 211–222. ACM (2009)

[25] Li, N., Tripunitara, M.V.: Security Analysis in Role-Based Access Control. ACM Trans. Inf. Syst. Secur. 9(4), 391–420 (2006)

[26] Ranise, S., Truong, A., Armando, A.: Boosting Model Checking to Analyse Large ARBAC Policies. In: Jøsang, A., Samarati, P., Petrocchi, M. (eds.) STM 2012. LNCS, vol. 7783, pp. 273–288. Springer, Heidelberg (2013)

[27] Sandhu, R.S., Bhamidipati, V., Munawer, Q.: The ARBAC97 Model for Role-Based Administration of Roles. ACM Trans. Inf. Syst. Secur. 2(1), 105–135 (1999)

[28] Sasturkar, A., Yang, P., Stoller, S.D., Ramakrishnan, C.: Policy analysis for Administrative Role-Based Access Control. Theoretical Computer Science 412(44), 6208–6234 (2011)

[29] Schwoon, S.: Model-Checking Pushdown Systems. Ph.D. Thesis, Technische Universität München (June 2002)

[30] Stoller, S.D., Yang, P., Ramakrishnan, C.R., Gofman, M.I.: Efficient Policy Analysis for Administrative Role Based Access Control. In: CCS, pp. 445–455 (2007)

[31] Uzun, E., Atluri, V., Sural, S., Vaidya, J., Parlato, G., Ferrara, A.L., Madhusudan, P.: Analyzing temporal role based access control models. In: Atluri, V., Vaidya, J., Kern, A., Kantarcioglu, M. (eds.) SACMAT, pp. 177–186. ACM (2012)