

Detecting Multi-Step IAM Attacks in AWS Environments via Model Checking

Ilia Shevrin¹, Oded Margalit²

¹Citi

²Ben-Gurion University



Background

Cloud adoption is on the rise, more data is stored in the cloud

Security posture of cloud applications is a growing concern

AWS introduced the shared responsibility model:

AWS is responsible for **infrastructure and hardware**

The customer is responsible for **application, data and IAM**

But IAM is notoriously hard to master due to its complexity

McGraw Hill's S3 buckets exposed 100,000 students' grades and personal info

Educator gets an F for security

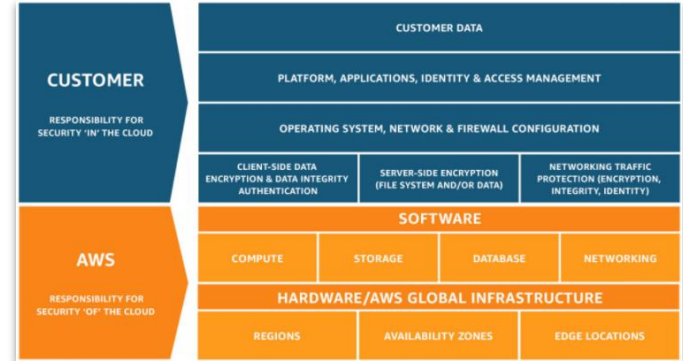


Image from

<https://docs.aws.amazon.com/whitepapers/latest/security-overview-of-amazon-codeguru-reviewer/the-shared-responsibility-model.html>

Cloud Misconfig Exposes 3TB of Sensitive Airport Data in Amazon S3 Bucket: 'Lives at Stake'

The unsecured server exposed more than 1.5 million files, including airport worker ID photos and other PII, highlighting the ongoing cloud-security challenges worldwide.

IAM Misconfiguration Example



Alice

```
"Effect ": "allow",  
"Action ": "*Role*",  
"Resource ": "*"
```

```
"Effect ": "deny",  
"Action ": "*Role*",  
"Resource ": "Alice"
```



bobs-bucket

```
"Effect ": "deny",  
"Action ": "*",  
"Resource ": "bobs-bucket/*",  
"Principal ": "Alice"
```



1. CreateRole Carol
1. PutRolePolicy Carol

```
{  
  "effect": "allow",  
  "action": "*",  
  "resource": "*" }  
}
```
1. AssumeRole Carol
1. GetObject bobs-bucket/*

Existing AWS IAM Security Tools

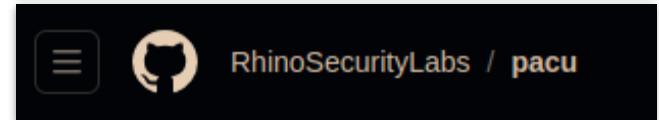
[Rhino Security Labs](#) identified 20+ AWS IAM privilege escalation techniques and released [Pacu](#) - an open source tool that scans policies and detects potential usage of these techniques

AWS developed [Zelkova](#) that mathematically verifies properties in IAM policies - for example checking if a bucket is public

Intro: AWS Privilege Escalation Vulnerabilities

lot of [penetration testing for AWS architecture](#), and invest heavily in related AWS new IAM Privilege Escalation methods - 21 in total - which allow an attacker to count to full administrative privileges.

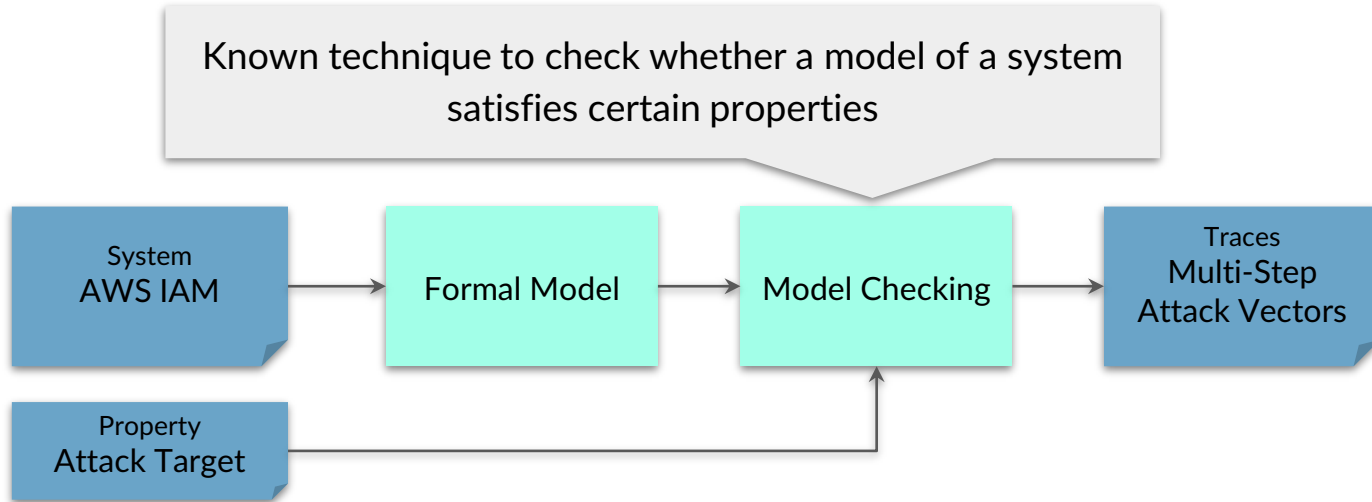
For privilege escalation routes, we've created a scanning tool ([available on Github](#)) to identify if you have an account with IAM read access for all users, the script can be run against a bucket in your account-wide.



Semantic-based Automated Reasoning for AWS Access Policies using SMT

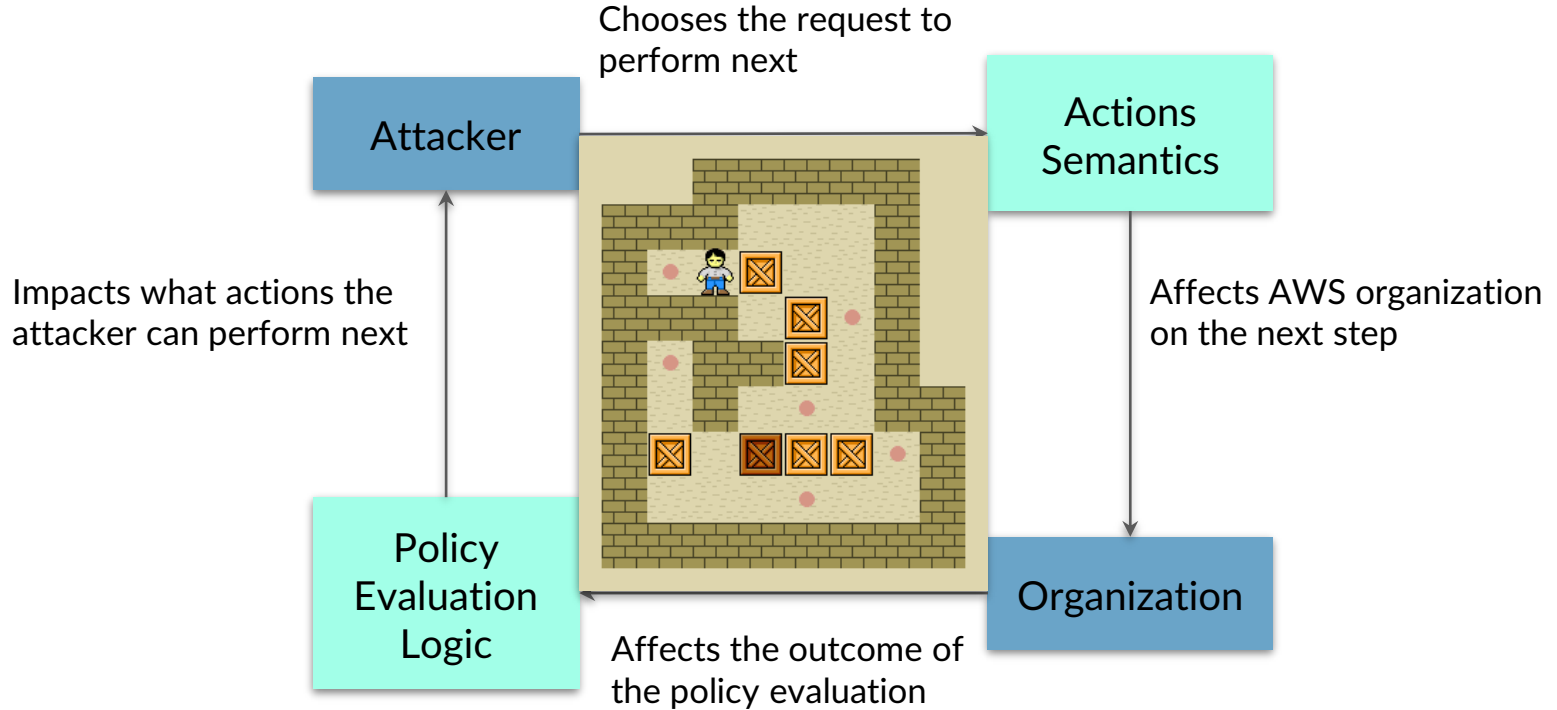
John Backes, Pauline Bolignano, Byron Cook, Catherine Dodge, Andrew Gacek, Kasper Luckow, Neha Rungta, Oksana Tkachuk, Carsten Varming
Amazon Web Services

A Model Checking Approach



Was already suggested in the context of network vulnerabilities by Ritchey & Ammann in 2000

AWS IAM Model



Actions Semantics Encoding

Recognize only *means* actions - affect IAM directly or indirectly (in total around 60 AWS actions)



Recognize a
privilege escalation
prone action



Understand AWS
documentation and translate
into an action semantics formula

12. Creating/Updating an inline policy for a role
PutRolePolicy permission can escalate privileges by creating or updating versions of that policy to the attacker.
The method might look like this:

```
aws iam put-role-inline-policy --role-name role_name --policy-name policy_name
```


A user can temporarily assume with [sts:AssumeRole](#).
To specify an arbitrary policy document with this method, the attacker can specify a resource, ultimately escalating to full administrator privileges.

PutRolePolicy
PDF
Adds or updates an inline policy document.
When you embed an inline policy in a role, the policy is created at the same time as the role. For more information about roles, see [IAM User Guide](#).
A role can also have a managed policy attached. For information about the maximum number of managed policies, see [IAM User Guide](#).

```
action = PutRolePolicy implies  
forall account in accounts:  
  forall role in account.IAMRoles:  
    ((resourceAccount = account.id and  
      resourceName = role.name) implies  
      role.MaximumPermissions' = true) and  
    ((resourceAccount != account.id or  
      resourceName != role.name) implies  
      role.MaximumPermissions' =  
      role.MaximumPermissions)
```

Rhino Security Labs website at
<https://rhinosecuritylabs.com/aws/aws-privilege-escalation-methods-mitigation/>

AWS documentation at
https://docs.aws.amazon.com/IAM/latest/APIReference/API_PutRolePolicy.html

Policy Evaluation Logic Encoding

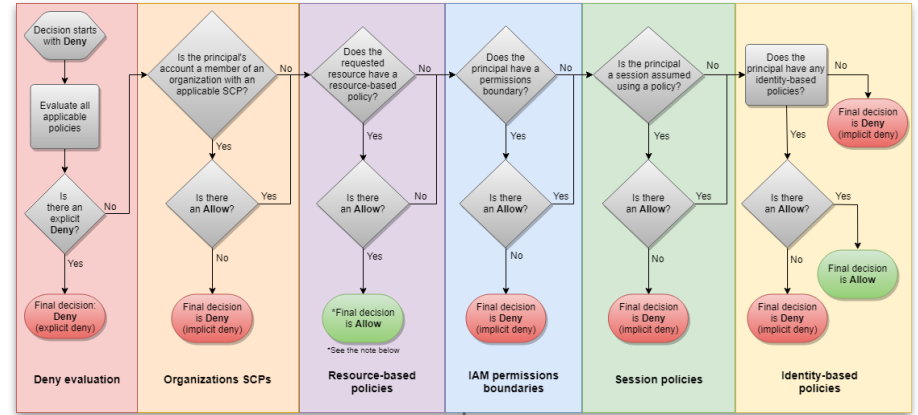
```
"Effect ": "allow",  
"Action ": "*Role*",  
"Resource ": "*"

"Effect ": "deny",  
"Action ": "*Role*"  
"Resource ": "Alice"
```

Encode formulas
similarly to Zelkova

```
allowStatements =  
  action in "*Role*" and  
  resourceType = Role

denyStatements =  
  action in "*Role*" and  
  resourceName = "Alice" and  
  resourceType = Role
```



```
ServiceControlPoliciesAllow and  
(ResourceBasedPoliciesAllow or  
(IdentityBasedPoliciesAllow and  
PermissionsBoundariesAllow)) and
```

```
not ServiceControlPoliciesDeny and  
not ResourceBasedPoliciesDeny and  
not PermissionsBoundariesDeny and  
not IdentityBasedPoliciesDeny
```


Implementation and Evaluation

We implemented the model checking process using Java + Z3 SAT Solver API. We used a bounded model checking algorithm (BMC) + an exhaustive version


We use a large pre-production AWS organization with ~100 accounts, with an average of ~200 IAM resources in each account

We ask how does the model checking process manage to detect existing misconfigurations in real world AWS environments


Gathered 141 different requests from security engineers, testing who can get access to data resources such as S3 buckets or SQS queues

Evaluation Results

Attack vectors by length



Actions performed in all the attack vectors



- ❑ Security engineers were satisfied with the results and fixed a lot of sneaky over-permissive policies
- ❑ Additional performance evaluation showed that the approach detects IAM attacks of up to 5 steps, in accounts with hundreds of resources, in under a minute

Thank you!

Questions?

iliashevrin@mail.tau.ac.il