ESnet
ENERGY SCIENCES NETWORK

Using Machine Learning for Intent-based Provisioning in High-Speed Science Network

Hocine Mahtout, Mariam Kiran,
Anu Mercian, Bashir Mohammad
Lawrence Berkeley National Lab
HPE

SNTA 2020

U.S. DEPARTMENT OF ENERGY
Office of Science

BERKELEY LAB

Problem Statement

Intent-based networking research

Tell me WHAT not HOW

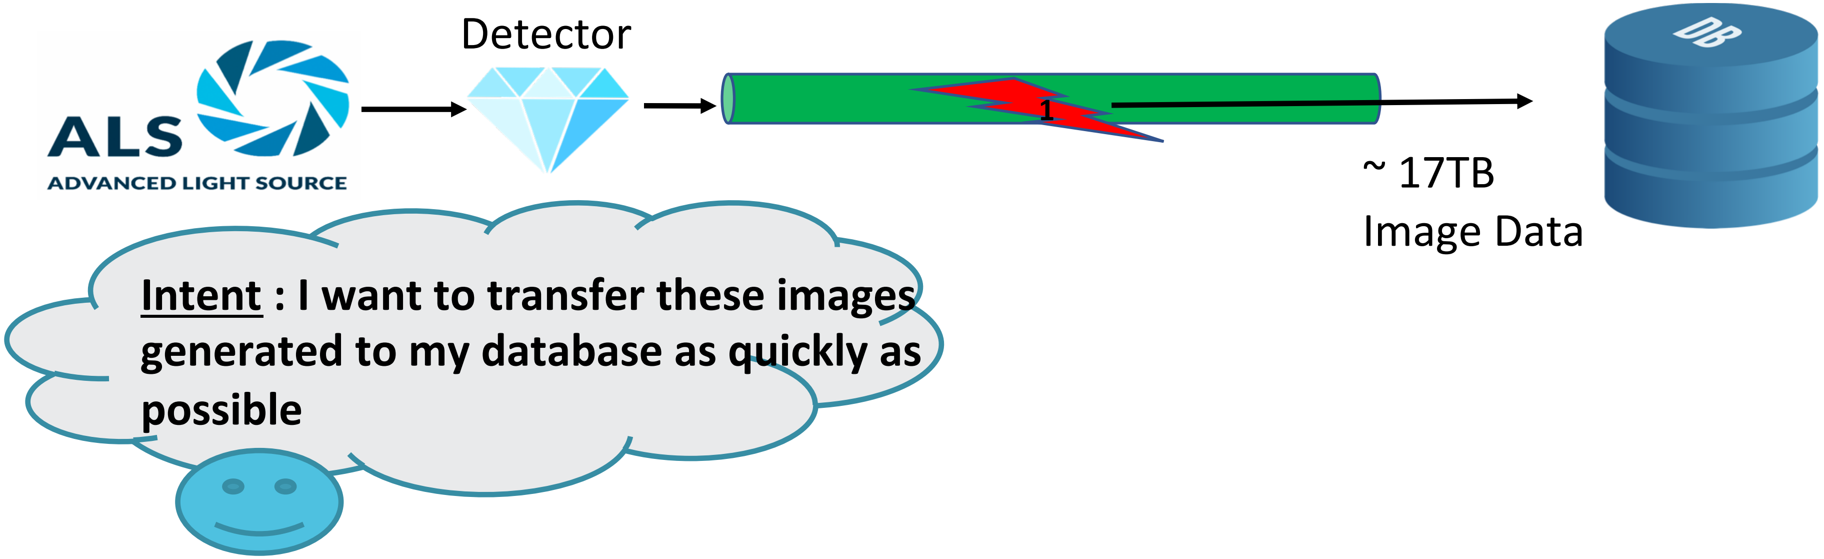

__Intent__ : I want to transfer these images generated to my database as quickly as possible

Overview

- Introduction and Motivation
  - Comparison of intent-based networking projects
- Machine Learning (Natural Language Processing = NLP)
- Evian Architecture
- Results & Conclusion

Introduction :
Focused on User intent

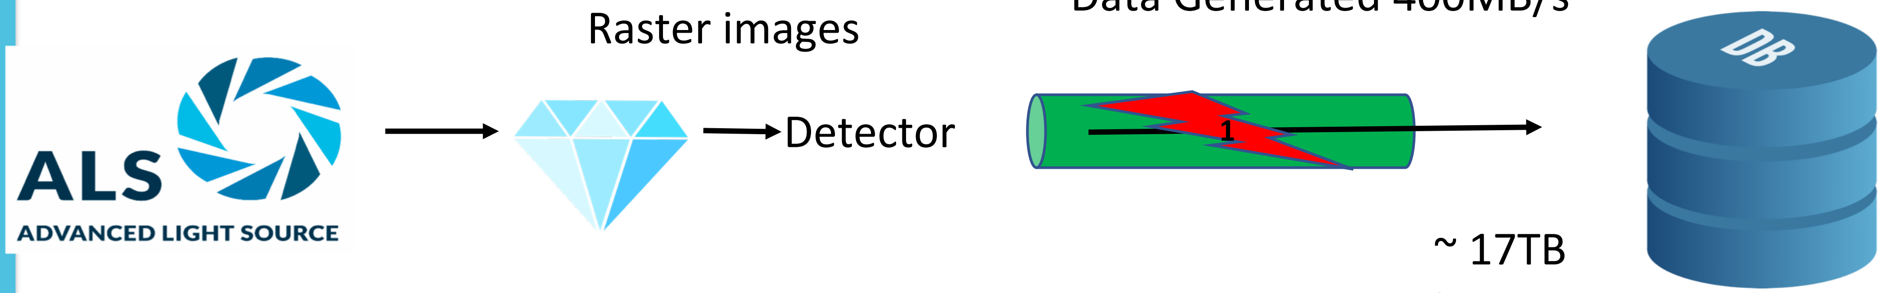

X-ray microscopic data
Raster images

Detector

Data Generated 400MB/s

~ 17TB
Image Data

<u>Intent</u> : I want to transfer these images generated to my database as quickly as possible

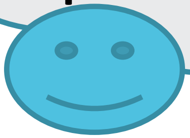

- Complex infrastructure
- Call engineers to set up these links
- Complex work around GUI upgrades
- And more….

Softwarization of Networks

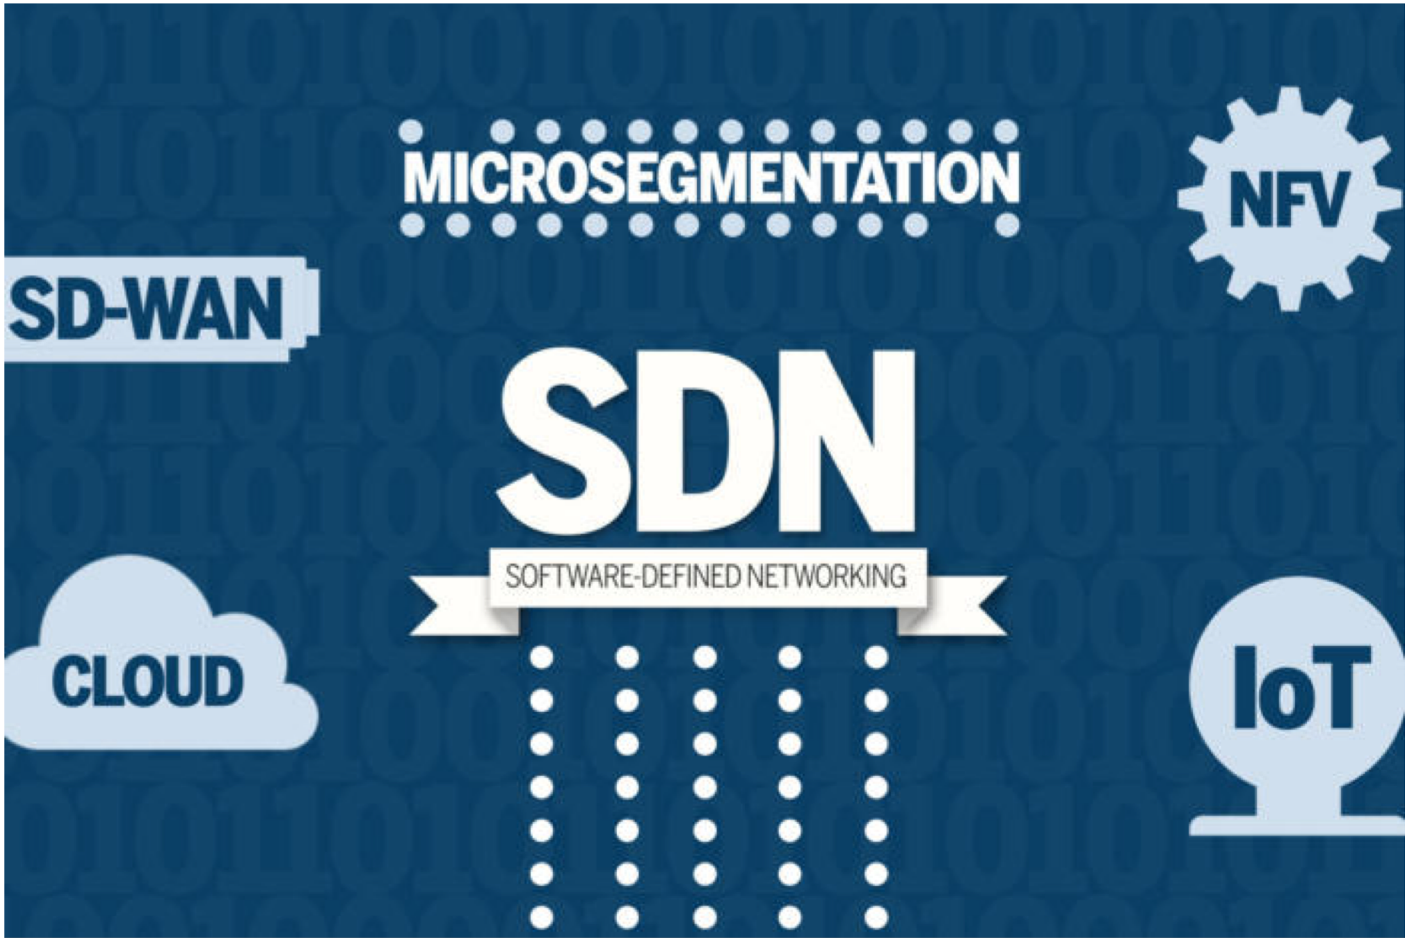

Some intent-based projects

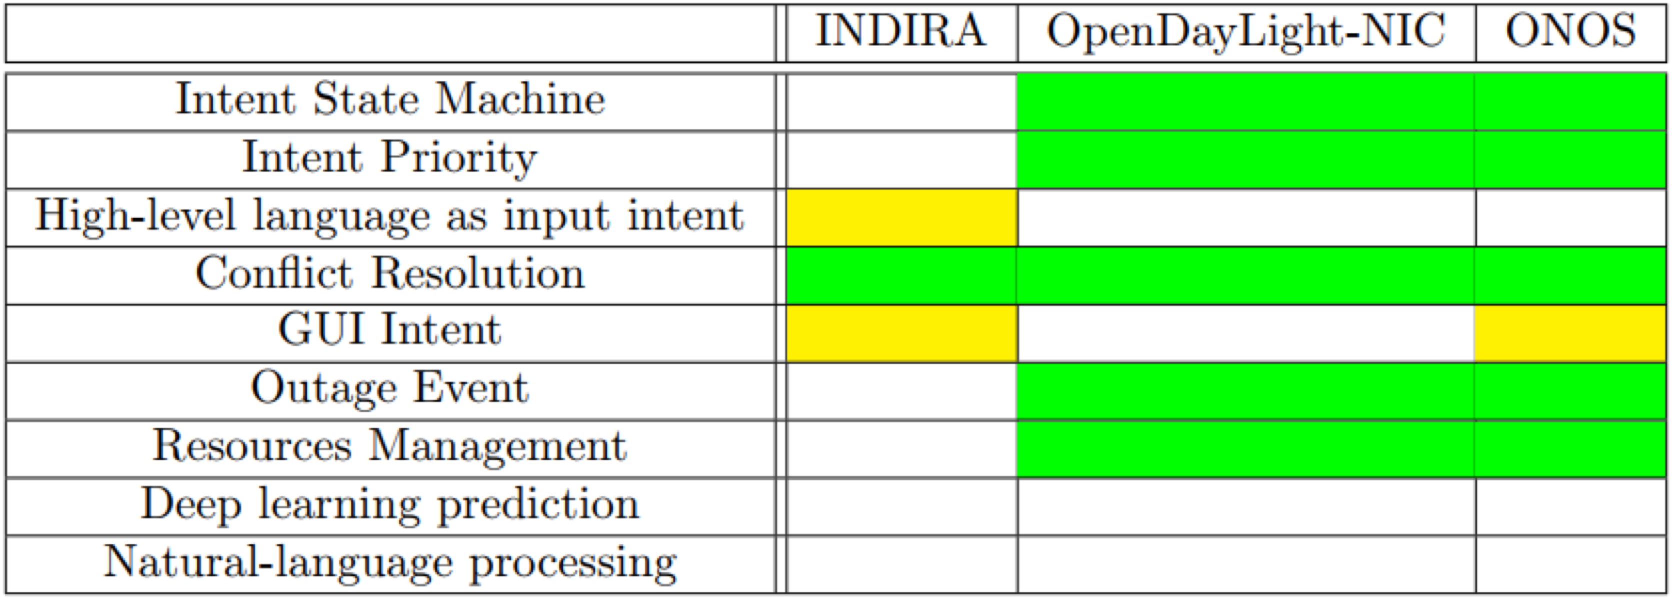

	INDIRA	OpenDayLight-NIC	ONOS
Intent State Machine	⬜	🟩	🟩
Intent Priority	⬜	🟩	🟩
High-level language as input intent	🟨	⬜	⬜
Conflict Resolution	🟩	🟩	🟩
GUI Intent	🟨	⬜	🟨
Outage Event	⬜	🟩	🟩
Resources Management	⬜	🟩	🟩
Deep learning prediction	⬜	⬜	⬜
Natural-language processing	⬜	⬜	⬜

🟩 Feature implemented

🟨 Feature almost implemented

⬜ Feature not implemented

Definition : Intent State Machine

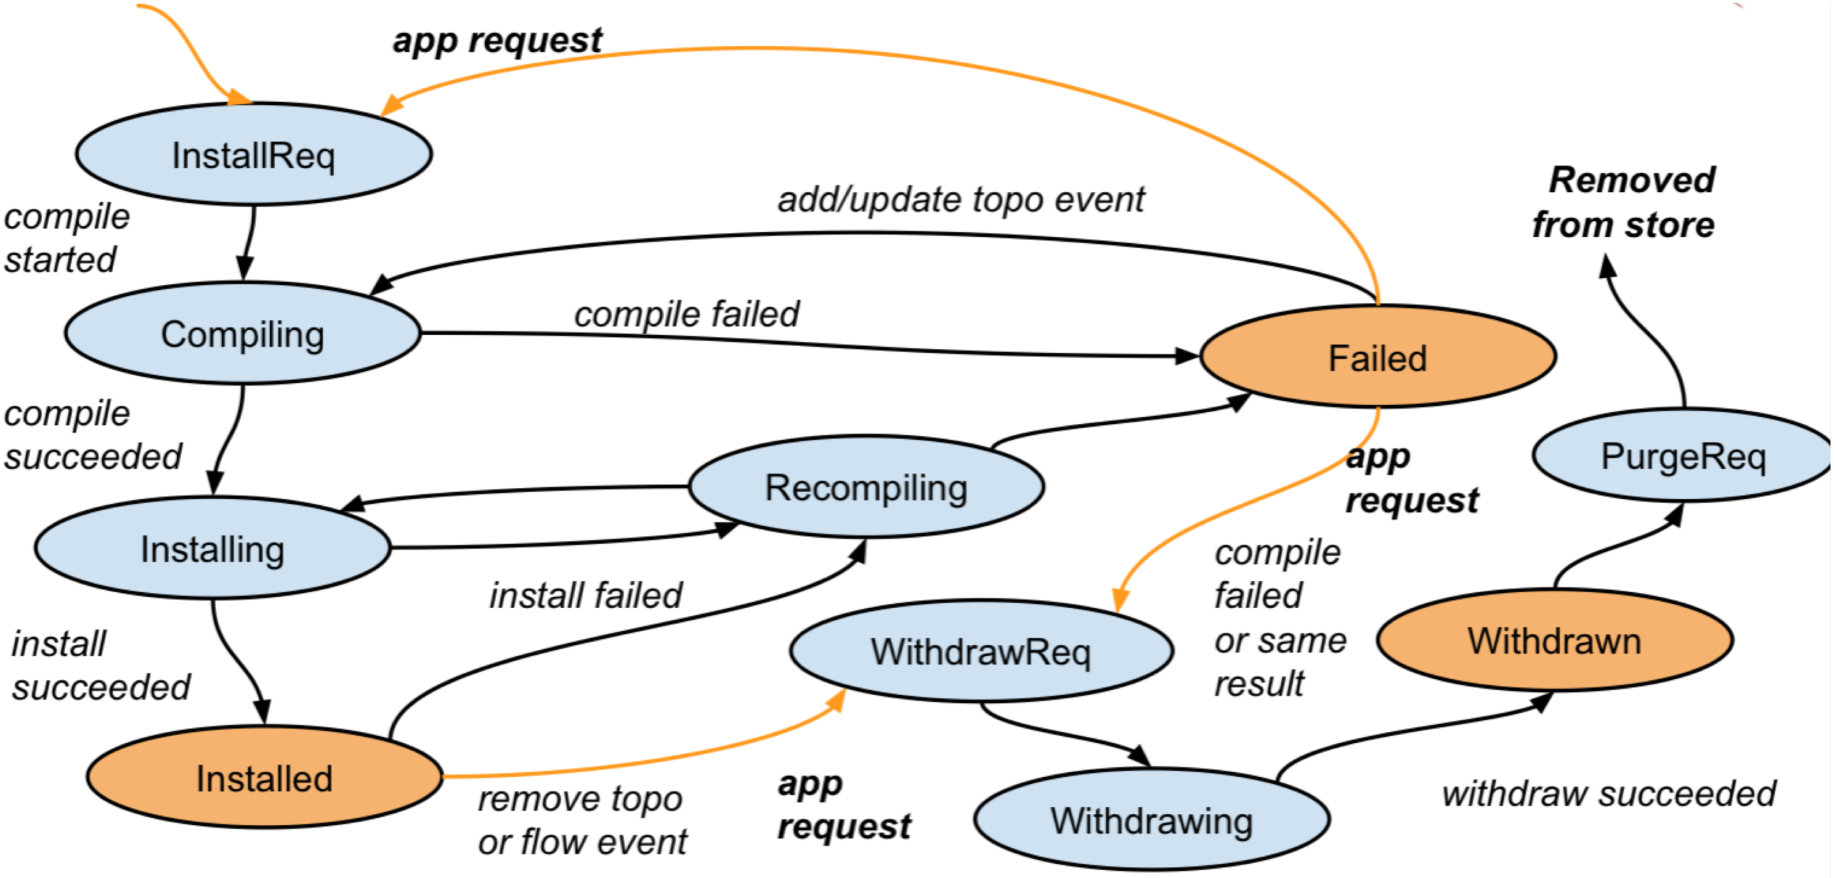

INDIRA Intent-tool: Lesson learned

- Presented at SC16

- Good:
  - Easy to talk manner
  - Automatic code rendering
  - Ontologies and Resource Description Graphs (RDF)

- Bad:
  - Very specific to certain tools: not extendable
  - No monitoring
  - No machine learning
  - Difficult to deploy

EVIAN’s new architecture

Handles issues from INDIRA

- Stripped it down to new code
- Machine learning for speech (NLP research)
- EVIAN bot can now ‘discuss’ options with users and negotiate
- EVIAN server is able to optimize configurations
- Multi tool capability
- Easy to deploy and also decouple different parts
- Bot communicates back to engineers before automating everything

Server-Client-Renderer Architecture

EVIAN Deployment

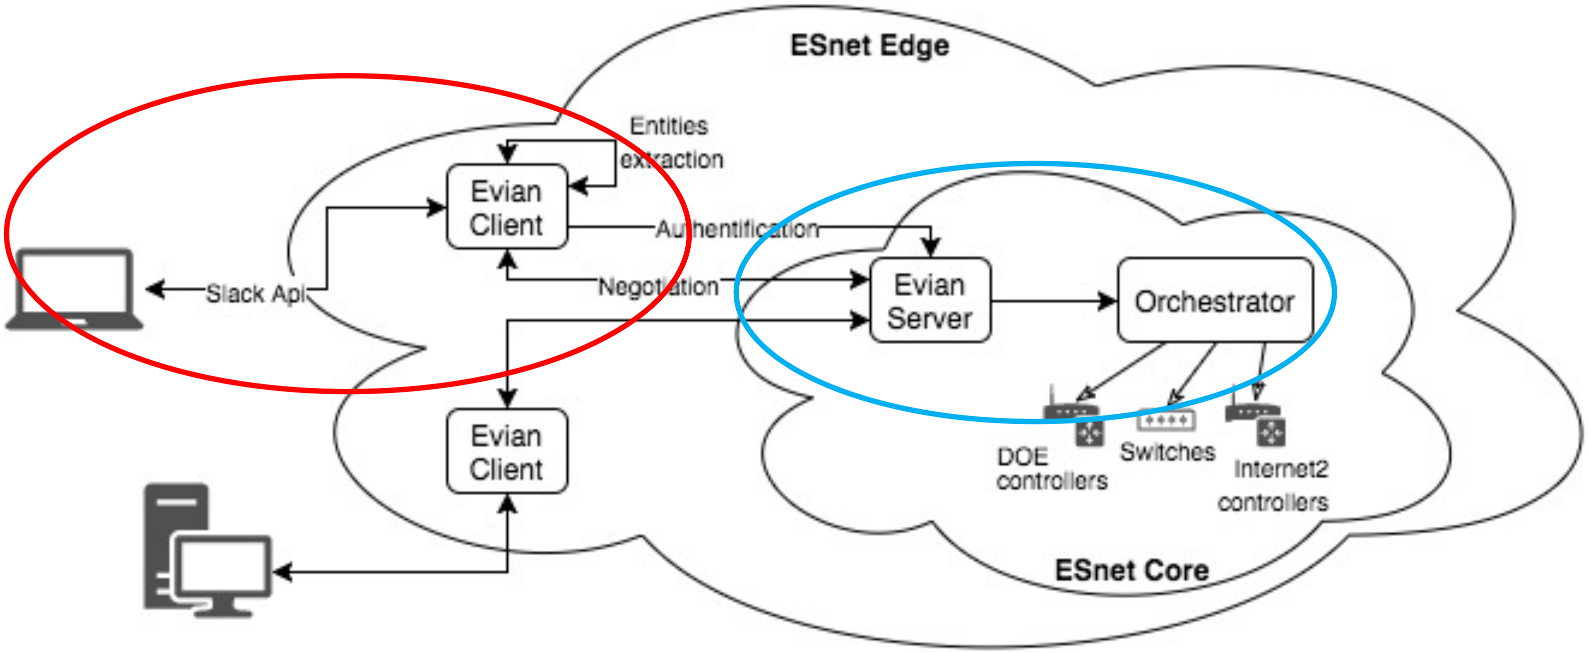

Mixing of many techniques

- EVIAN BOT :
  - RASA NLU Open source code
- EVIAN CLIENT:
  - SLACK API
- EVIAN SERVER:
  - Optimization for resource management

- RDF graphs to store data across all stages

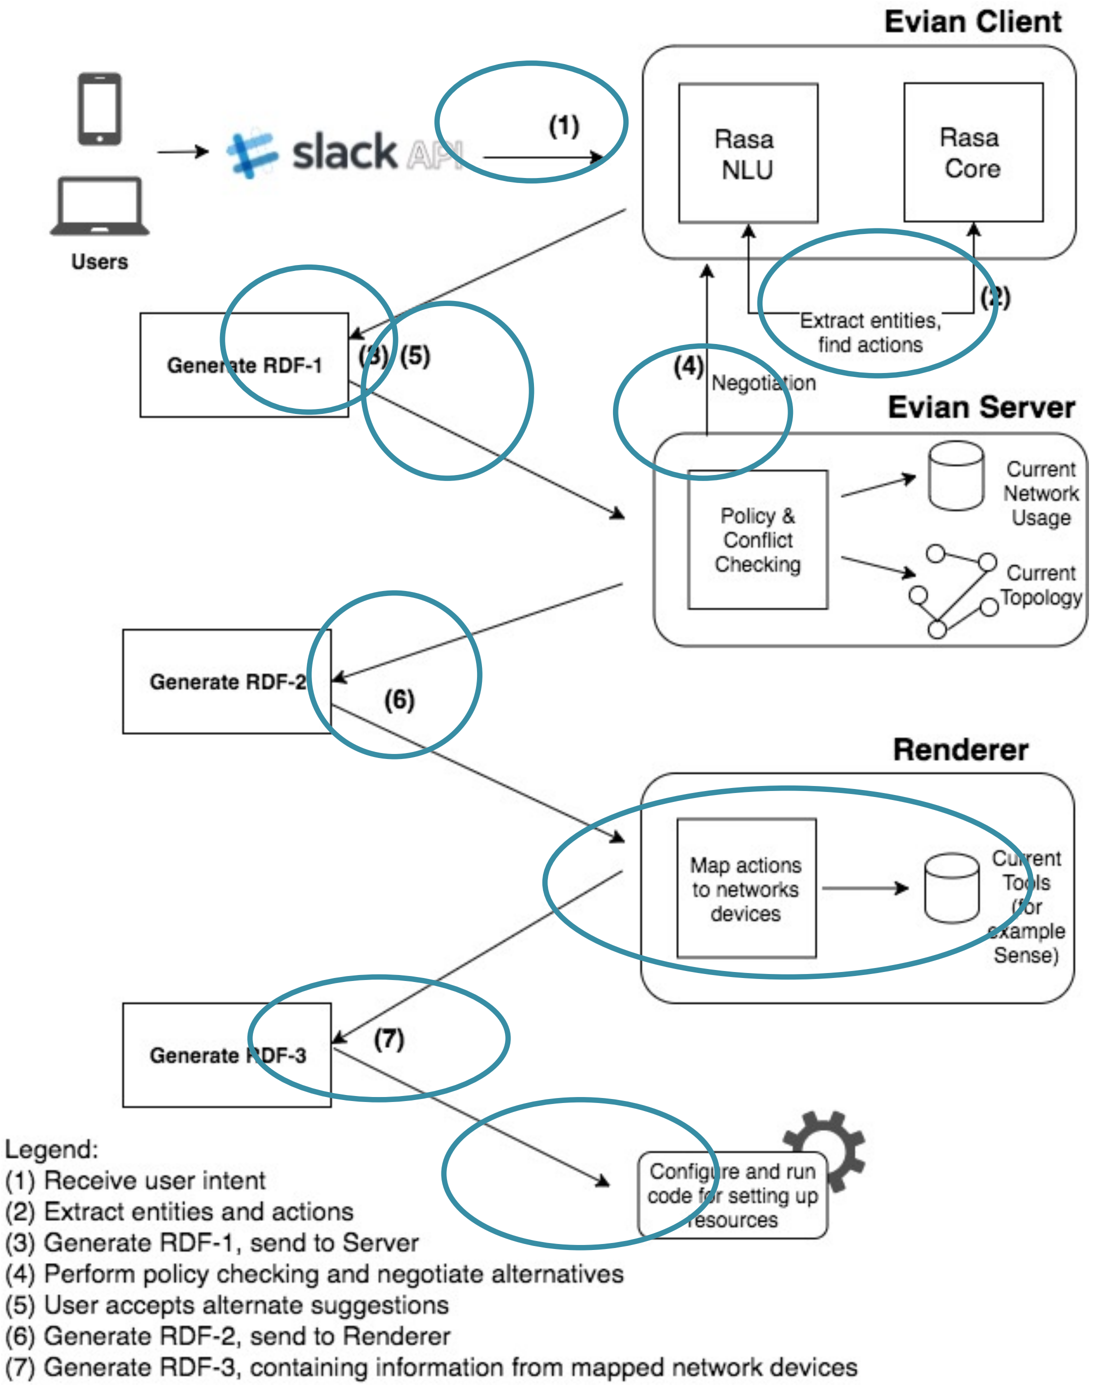

Legend:
(1) Receive user intent
(2) Extract entities and actions
(3) Generate RDF-1, send to Server
(4) Perform policy checking and negotiate alternatives
(5) User accepts alternate suggestions
(6) Generate RDF-2, send to Renderer
(7) Generate RDF-3, containing information from mapped network devices

- Overall Architecture
- Server-client-renderer

Add more functionality to either,

- Server with various optimization algorithms
- Client with more intelligence
- Renderer with many more tools: agile network

Designing the BOT:
Rasa - Open Source NLP System

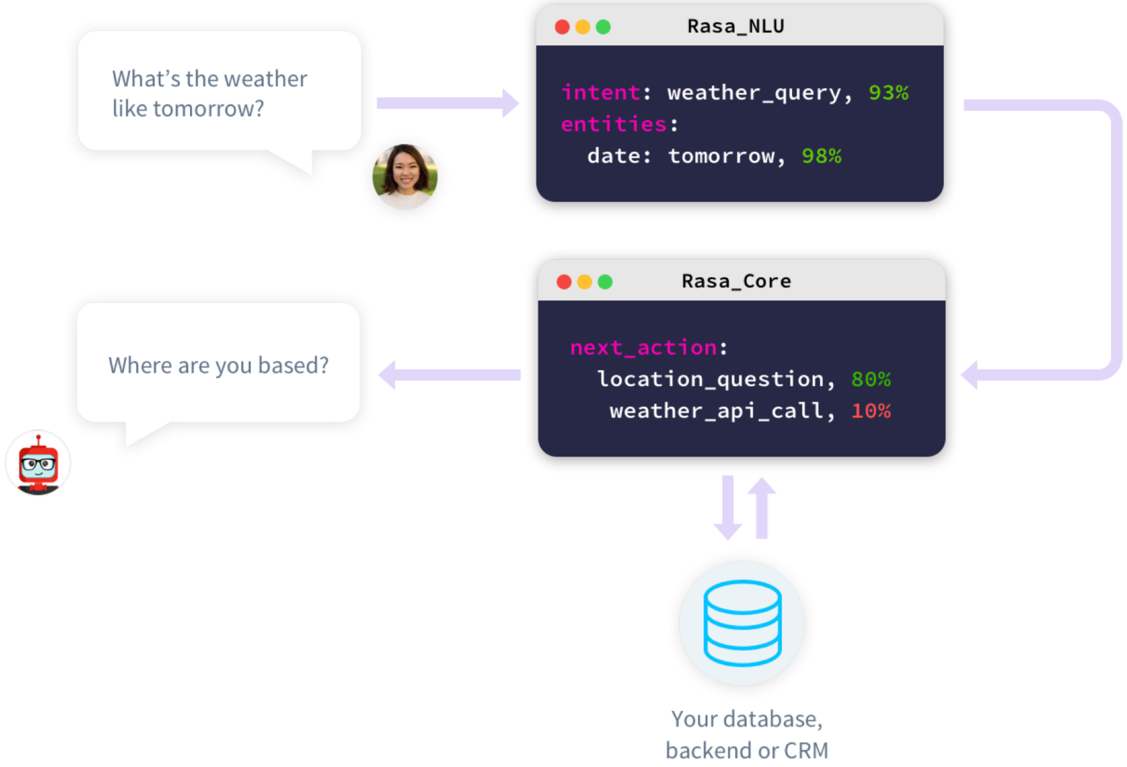

What's the weather like tomorrow?

Rasa_NLU

```
intent: weather_query, 93%
entities:
  date: tomorrow, 98%
```

Rasa_Core

```
next_action:
  location_question, 80%
   weather_api_call, 10%
```

Where are you based?

Your database,
backend or CRM

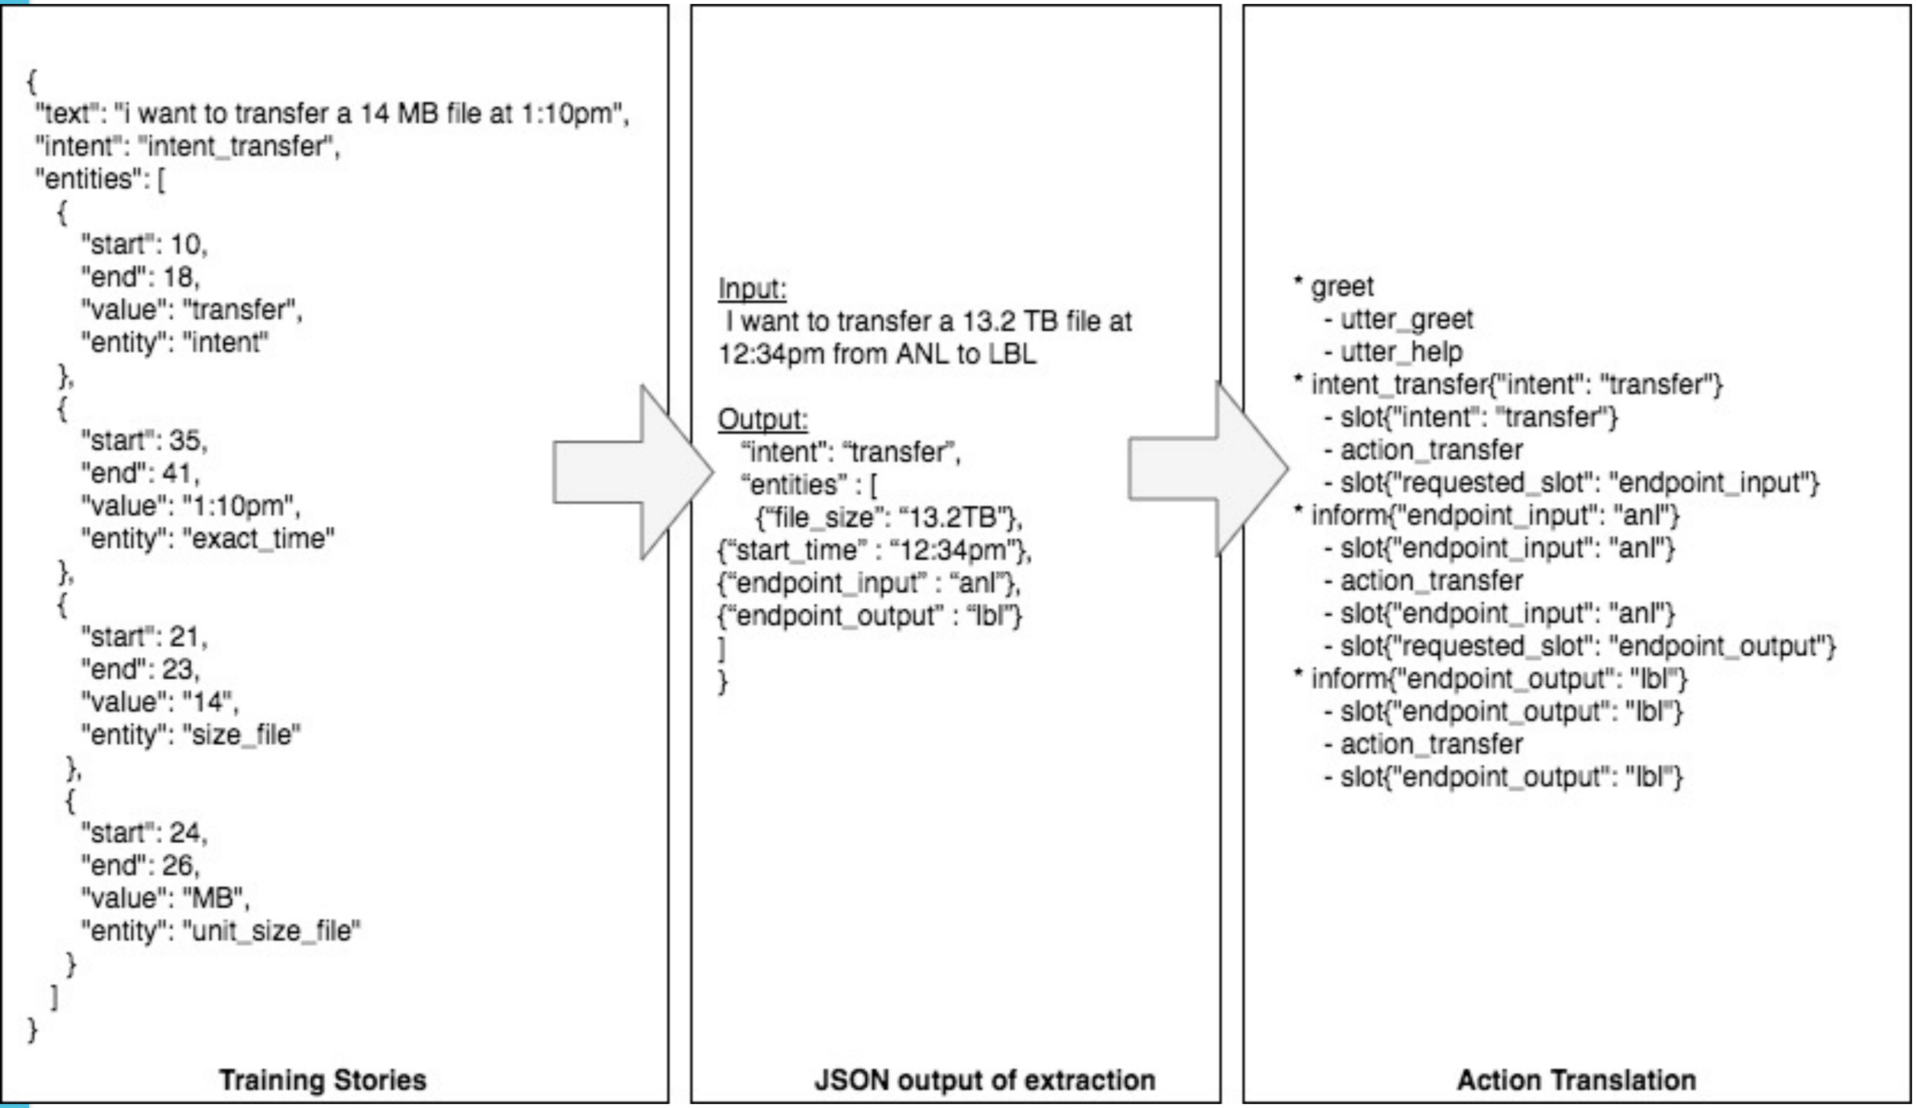

```
{
 "text": "i want to transfer a 14 MB file at 1:10pm",
 "intent": "intent_transfer",
 "entities": [
   {
     "start": 10,
     "end": 18,
     "value": "transfer",
     "entity": "intent"
   },
   {
     "start": 35,
     "end": 41,
     "value": "1:10pm",
     "entity": "exact_time"
   },
   {
     "start": 21,
     "end": 23,
     "value": "14",
     "entity": "size_file"
    },
    {
     "start": 24,
     "end": 26,
     "value": "MB",
     "entity": "unit_size_file"
    }
  ]
}
```

Training Stories

```
Input:
 I want to transfer a 13.2 TB file at
12:34pm from ANL to LBL

Output:
  “intent”: “transfer”,
  “entities” : [
    {“file_size”: “13.2TB”},
{“start_time” : “12:34pm”},
{“endpoint_input” : “anl”},
{“endpoint_output” : “lbl”}
]
}
```

JSON output of extraction

```
* greet
  - utter_greet
  - utter_help
* intent_transfer{"intent": "transfer"}
  - slot{"intent": "transfer"}
  - action_transfer
  - slot{"requested_slot": "endpoint_input"}
* inform{"endpoint_input": "anl"}
  - slot{"endpoint_input": "anl"}
  - action_transfer
  - slot{"endpoint_input": "anl"}
  - slot{"requested_slot": "endpoint_output"}
* inform{"endpoint_output": "lbl"}
  - slot{"endpoint_output": "lbl"}
  - action_transfer
  - slot{"endpoint_output": "lbl"}
```

Action Translation

RDF to store the intent

First RDF at Descriptive layer	First sentence with requirements
Second RDF at Prescriptive layer	Intent with user/network information
Third RDF at Rendering layer	Linking intent to rendering tools

Example

Bot > Hi there ! How can i help you ?
User> I would like to transfer a file please
Bot > From where ?
User> From where can i send this file ?
Bot> Here are your available endpoints :
Anl, lbl
User> From anl
Bot > To where ?
User> to lbl
Bot > What is the path in the input endpoint?
User> /dir/file
Bot > What is the path in the output endpoint?
User> ~/directory/file

Example (cont.)

Bot > What is the size of file ?

User> 13.4 TB

Bot > When do you want to start your transfer?

User> at 9pm

Bot > You asked for transfer, here are the information you have provided :

- endpoint_input : anl
- endpoint_output : lbl
- path_output : ./directory/file
- path_input : /dir/file
- size_file : 13.4
- unit_size_file : tb
- exact_time : 9pm

First RDF graph

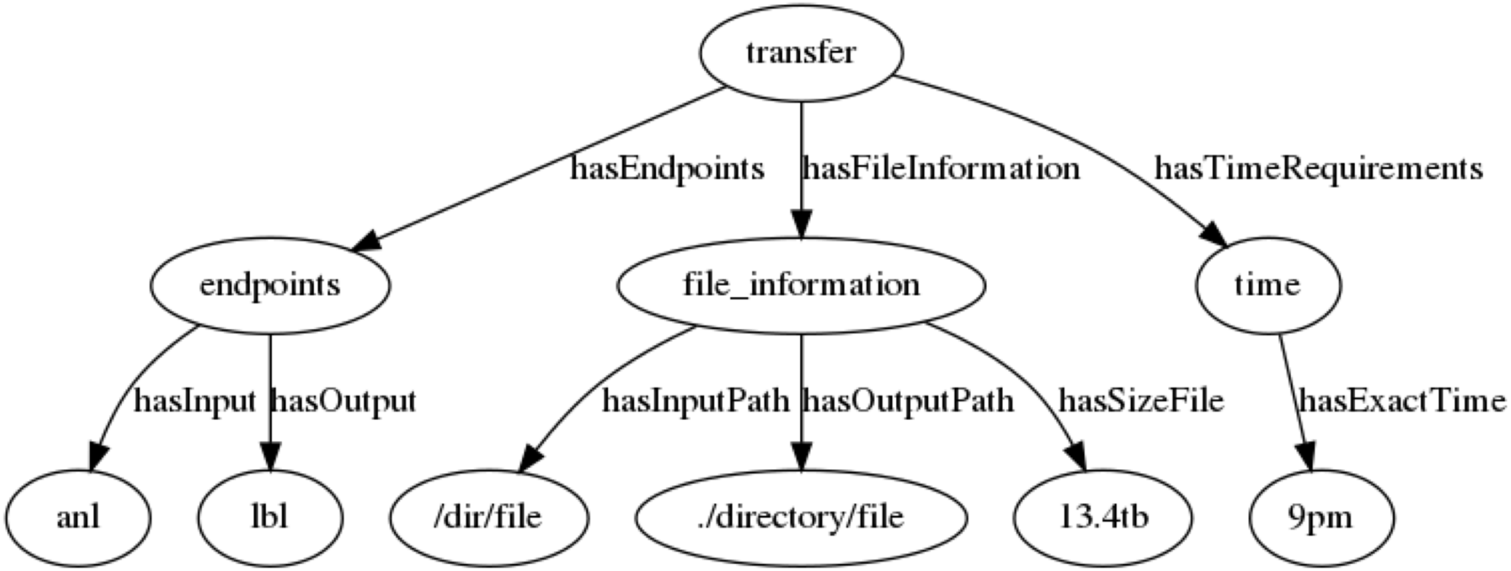

Second RDF graph

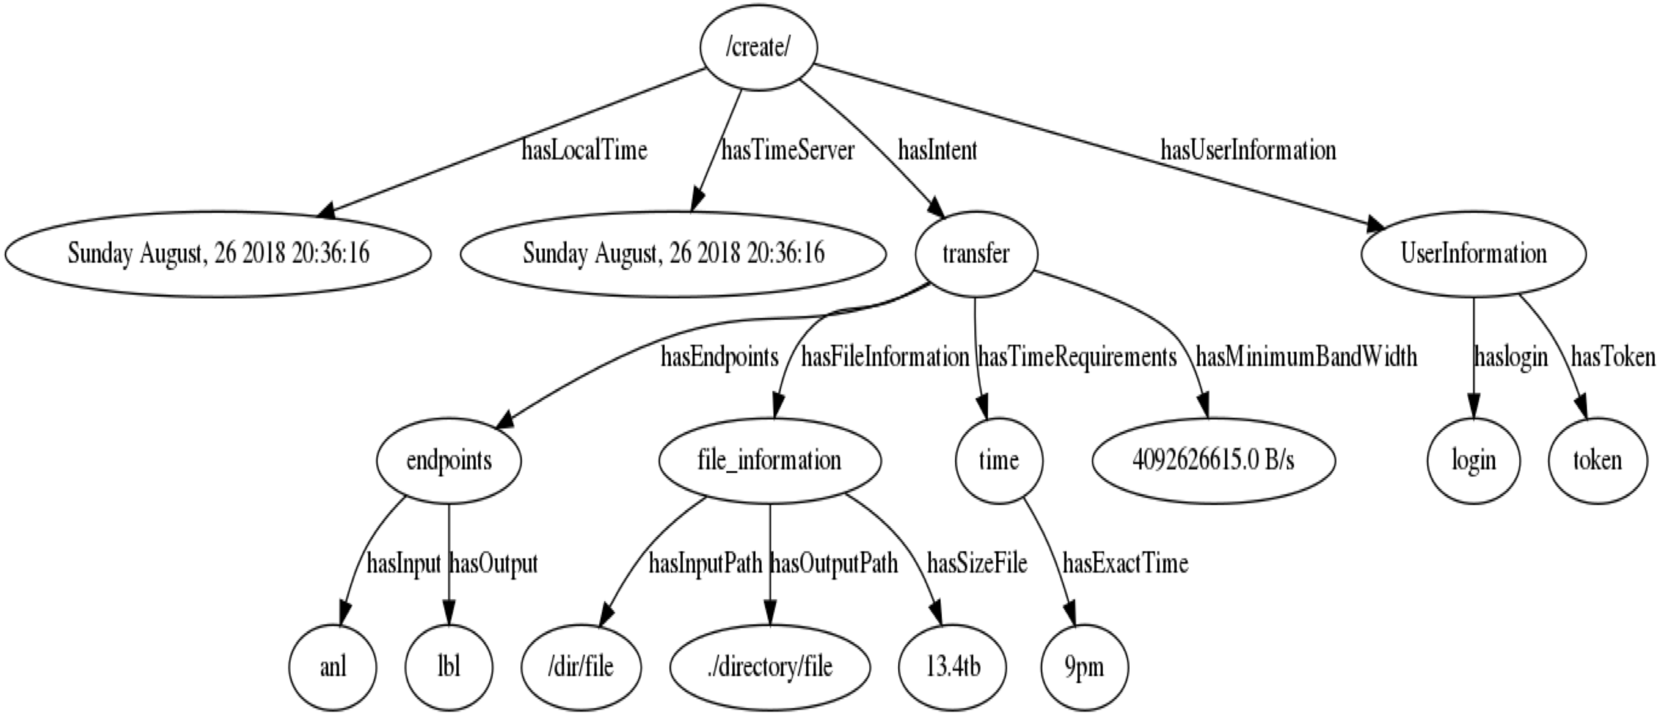

Example (cont.)

Bot> I will check if i can set this transfer for you

Your intent is not installable. Here are other time options :

1. 10pm
2. 11am

Or if you want to quit, enter quit

User> Send it at 10pm then

Bot> The network is set for your transfer

User> Bye !

Bot> Good Bye !

- SLACK API

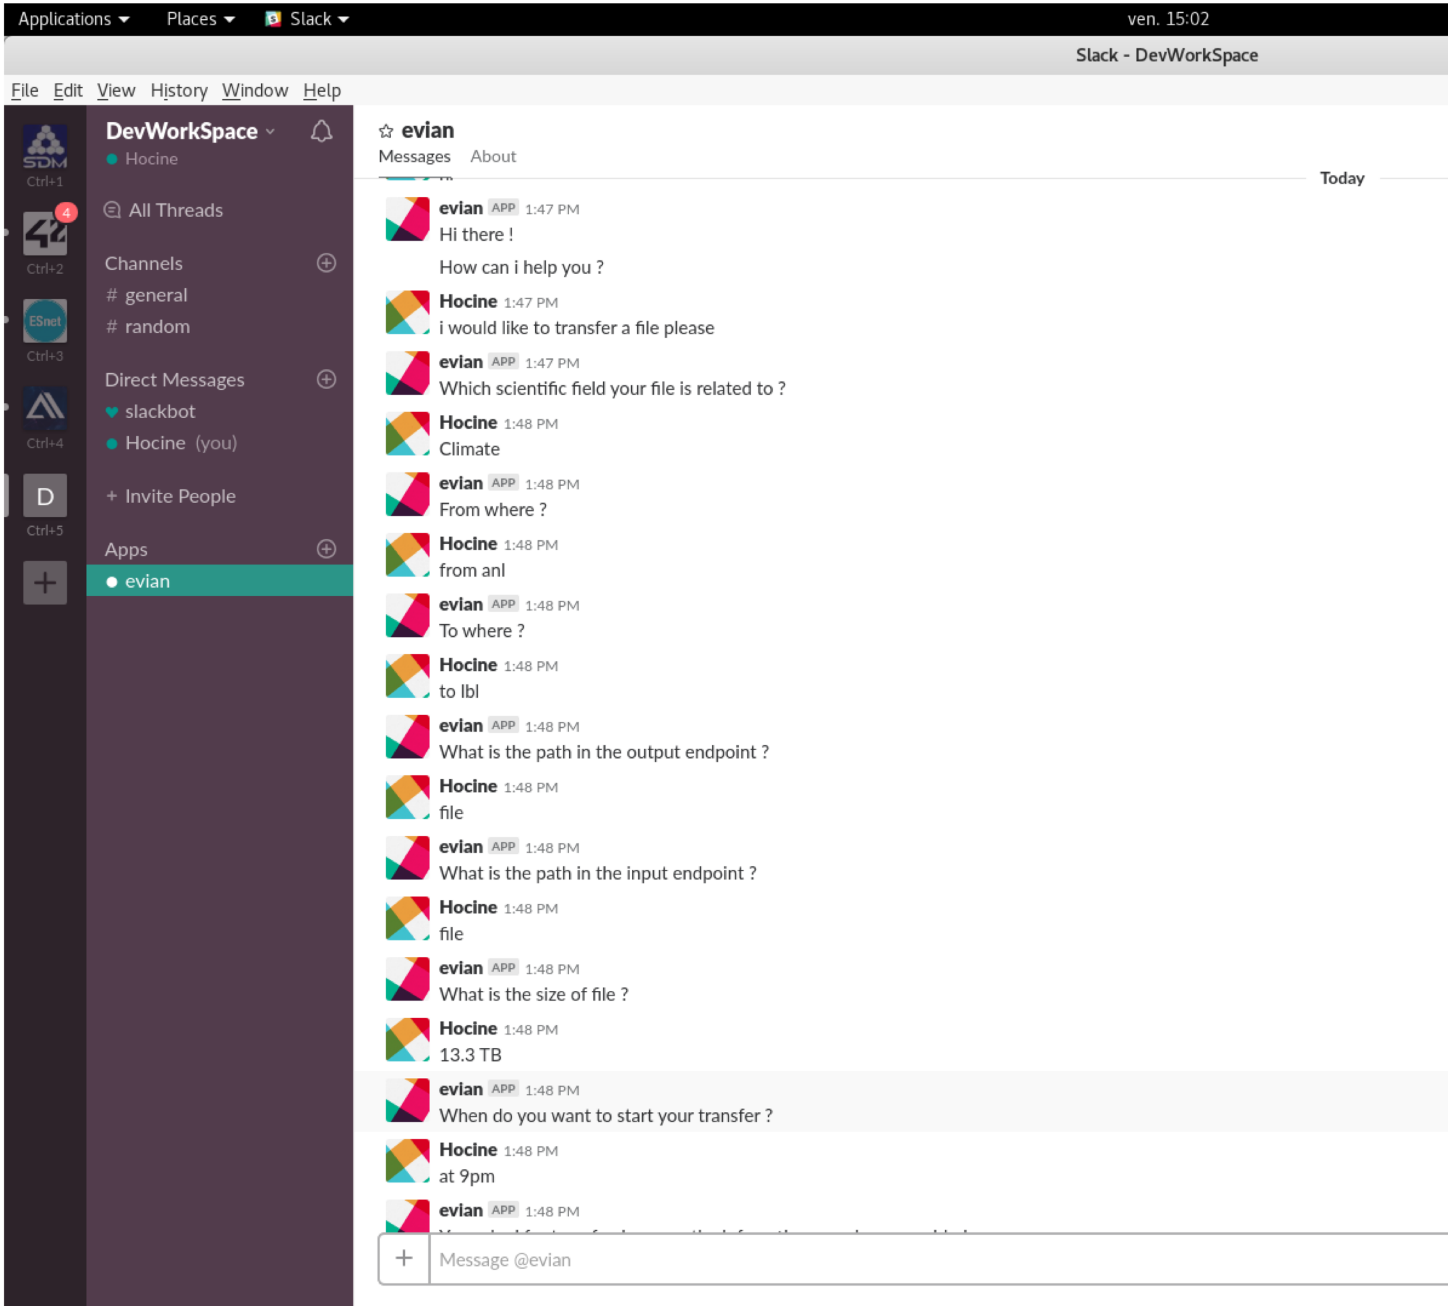

evian APP 1:47 PM
Hi there !
How can i help you ?

Hocine 1:47 PM
i would like to transfer a file please

evian APP 1:47 PM
Which scientific field your file is related to ?

Hocine 1:48 PM
Climate

evian APP 1:48 PM
From where ?

Hocine 1:48 PM
from anl

evian APP 1:48 PM
To where ?

Hocine 1:48 PM
to lbl

evian APP 1:48 PM
What is the path in the output endpoint ?

Hocine 1:48 PM
file

evian APP 1:48 PM
What is the path in the input endpoint ?

Hocine 1:48 PM
file

evian APP 1:48 PM
What is the size of file ?

Hocine 1:48 PM
13.3 TB

evian APP 1:48 PM
When do you want to start your transfer ?

Hocine 1:48 PM
at 9pm

evian APP 1:48 PM

Results

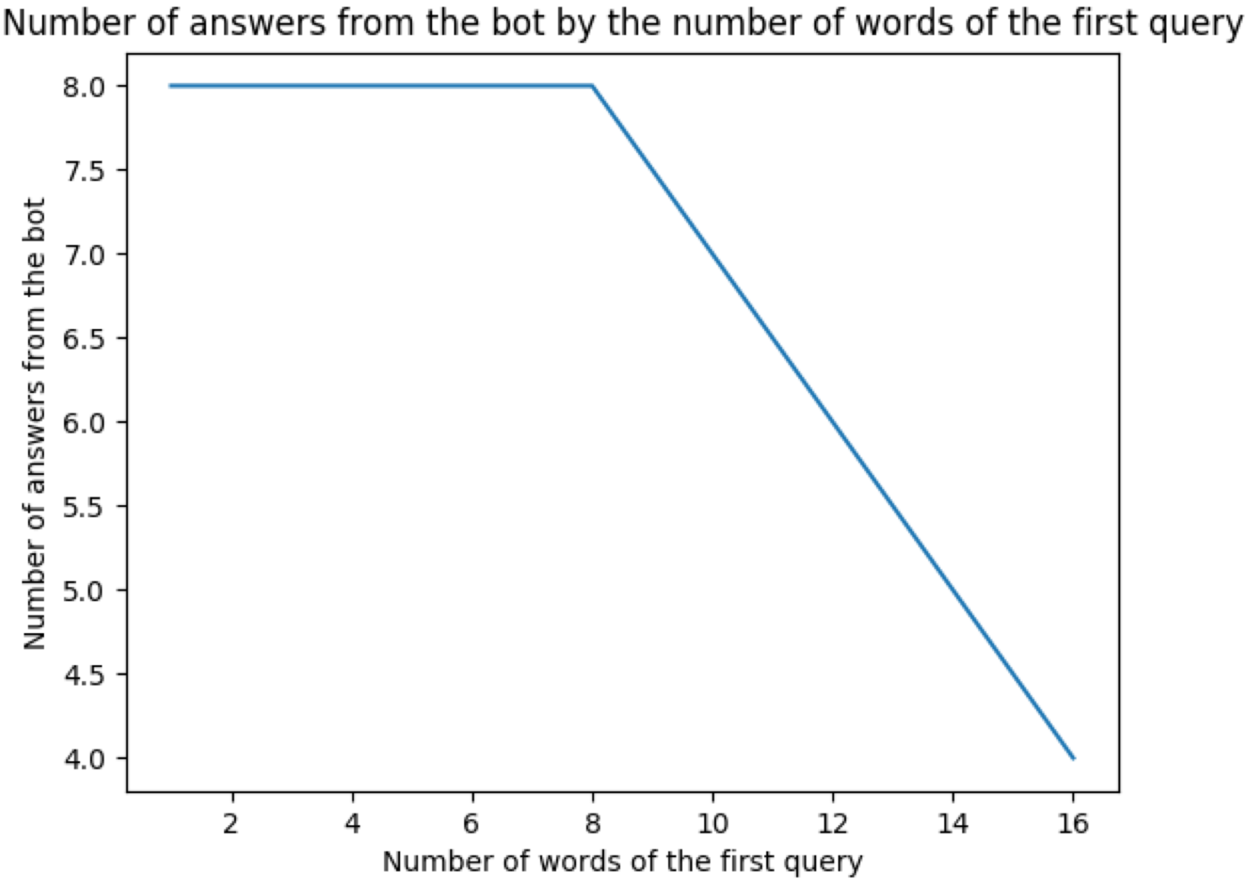
Number of answers from the bot by the number of words of the first query
Number of answers from the bot
8.0
7.5
7.0
6.5
6.0
5.5
5.0
4.5
4.0
2
4
6
8
10
12
14
16
Number of words of the first query

Results (cont.)

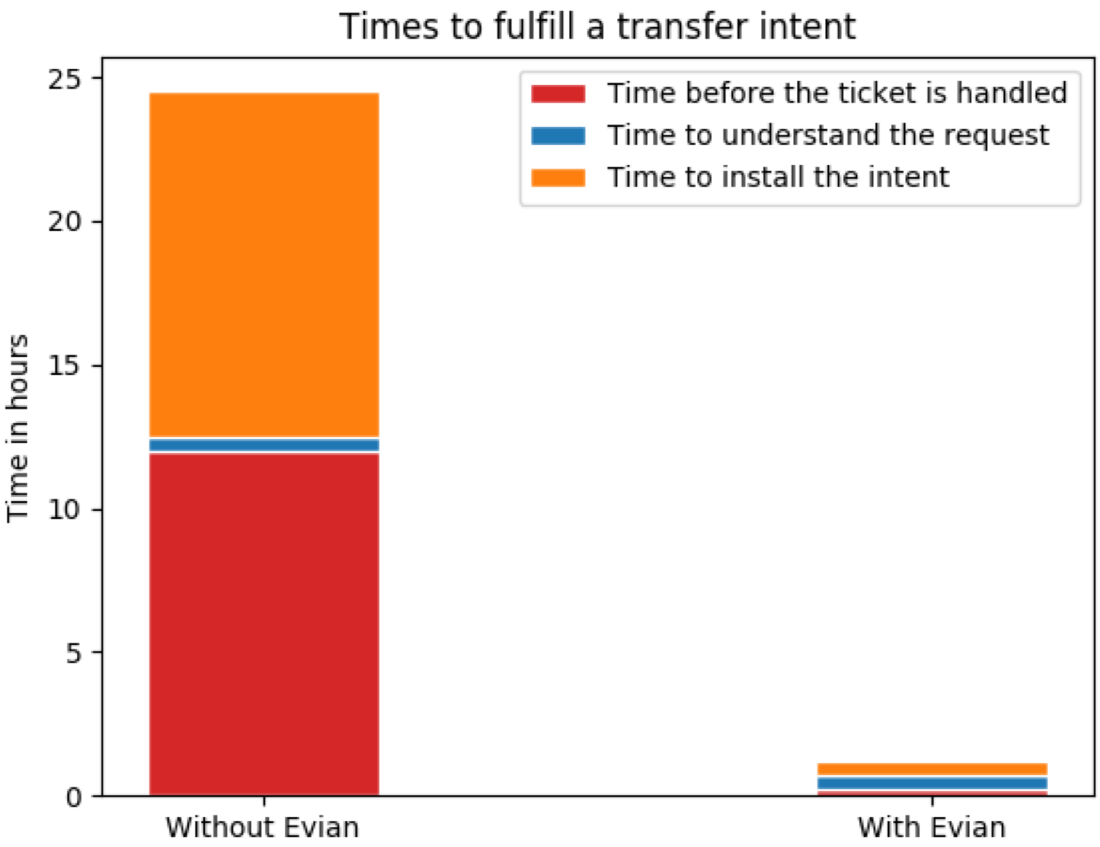

Ways forward

- EVIAN can connect to controllers through intent API
- Security permissions:
  - Slack client was an issue
  - Server has access to orchestrators might be an issue
  - Want to work with security team in ironing out these details
- Add conflict and policy checking
- Add machine learning predictions to bot responses
- Easy to change the slack API into a GUI on top
- Architecture allows more tools to be added and automation, might also write out Ansible code in future!

ESnet

Any questions ?

Project DAPHNE:

Developing Machine Learning Solutions for High-performance networks

mkiran@es.net