

## Assignment 4: Trusted Routing Zones: Building on the MANRS Core (9pts)

Due on Wednesday March 1, 2023 11:59pm

### 1 Introduction

Mutually Agreed Norms on Routing Security (MANRS) is an industry-led initiative to improve Internet routing security by encouraging participating networks to implement a set of best practices, called *actions*. These actions include registering one's IP prefixes in a trusted database, and using such information to prevent propagation of invalid routing information.

To take this idea further, imagine a connected region of the Internet composed of ISPs that commit to perform ROV on all announcements coming into that region. Then, within that region there will be no propagation of origin hijacks. Furthermore, any AS that directly attaches to that region (i.e., the AS is a customer of an ISP in that region, not necessarily in the region itself) cannot be the victim of an origin hijack inside that region. Similarly, if a new set of operational practices prevents the propagation of *path hijacks* in the region, then customers directly attached to the region will not be the victim of a path hijack inside that region. One might call this region a *zone of trust* because the protection arises at the perimeter of the zone. This protection requires that ASes in the zone be able to trust that the routers at the perimeter function correctly, which requires some degree of transparency and accountability.

The idea of a coherent perimeter around a zone is missing from today's interdomain routing system. Global deployment has always been a routing security protocol design assumption. Yet, recognition that ASes themselves can be the threat actors sheds doubt on any aspiration to make BGP *globally* secure. Creating a zone of trust through perimeter protection (a trust-but-verify regime) offers a more pragmatic approach for today's routing system. Most important is the incentive alignment. A zone of trust approach would be able to clearly articulate the benefit that the practices of that zone are bringing to their customers.

Could such a coherent topological region exist? Fortunately, it already does, in the context of the MANRS initiative. The current MANRS program is already centered around operational practices, and already manifests a significant zone of trust.

In this assignment, you will explore the network topology of MANRS members and calculate the MANRS Core. You will apply the data science knowledge from PA3 to complete the tasks in this assignment. You will be using the following datasets to complete this exercise.

- CAIDA AS Organizations Dataset
- CAIDA AS Relationships Dataset

- MANRS participants list

## 1.1 Due Date

**Wednesday, March 1st 11:59pm**

## 1.2 Submission Instructions

There will be 5 tasks in this assignment. You will complete task 1-4 in `python` files, and submit to the **PA4 Autograder** on Gradescope. Your scores will be immediately available upon submission and you have unlimited submission attempts until due date. The remaining one task will be in the form of a written homework, where you will answer a list of questions based on the readings from Week 7.

# 2 Identifying the MANRS Core

## 2.1 Background

In the FCC NOI response by Clark, Claffy, and Testart (see week 7 reading), there is a paragraph describing the MANRS Core:

*This core includes a number of Tier 1 providers, their customers, and so on. If the rule for being a part of the core is that the member must have at least one MANRS member as a provider, there are about 501 ASNs from 390 ISPs in the core.*

The goal of this section is to reproduce the methodology of the MANRS calculation and update the above numbers using more recent data. To clarify the methodology: a network is in the MANRS Core if it satisfies both of the following requirements:

- (a) It is a MANRS participant
- (b) It has at least one provider that is in the MANRS Core

## 2.2 Data Access

We will use the list of participants in the following two MANRS programs. You can see the description of each program and find the list of participants from the links below:

- MANRS Network Operator Program  
<https://www.manrs.org/netops/participants/>
- MANRS CDN and Cloud Provider Program  
<https://www.manrs.org/cdn-cloud-providers/participants/>

The files are also provided here for convenience:

MANRS ISPs: [https://cseweb.ucsd.edu/classes/wi23/cse291-e/pa4/manrs\\_isp\\_20230221.csv](https://cseweb.ucsd.edu/classes/wi23/cse291-e/pa4/manrs_isp_20230221.csv).

MANRS CDNs: [https://cseweb.ucsd.edu/classes/wi23/cse291-e/pa4/manrs\\_cdn\\_20230221.csv](https://cseweb.ucsd.edu/classes/wi23/cse291-e/pa4/manrs_cdn_20230221.csv)

You will use the following recent version of the AS Organizations and AS Relationship datasets:

AS Relationship: <https://cseweb.ucsd.edu/classes/wi23/cse291-e/pa4/20230201.as-rel.txt.bz2>

AS Organizations: <https://cseweb.ucsd.edu/classes/wi23/cse291-e/pa4/20230101.as-org2info.txt.gz>

## 2.3 Data Processing

### 2.3.1 MANRS Participants

The lists of MANRS participants are provided as `csv` files. You will need to identify the columns that contain ASNs in each file (MANRS Network Operators Program and MANRS CDN Program). Note that a MANRS organization may register multiple ASNs in the program.

**You will maintain a single list of the AS numbers of MANRS participants in both programs.**

### 2.3.2 AS Relationship – Tier 1 providers

Clark, Claffy, and Testart’s FCC response mentioned “a number of Tier 1 providers” in the MANRS Core. You can find a list of such Tier 1 providers in the AS Relationship dataset file. Look in the beginning of the file for the line that starts with `# input clique`. You will use the AS numbers in that line to start your MANRS Core calculation.

## 2.4 Task 1: Calculate the MANRS Core (1pt)

Your first task is to find all ASes in the MANRS Core. Your code should do the following:

- (1) Parse the participants lists in the MANRS ISP program and MANRS CDN program and keep a list of ASes participating in those programs. We refer to those ASes as MANRS ASes.
- (2) Parse the AS relationship file. Maintain a list of Tier 1 providers (§2.3.2) and a mapping of each AS to its customer ASes.
- (3) Create an empty set of ASes. We denote this set as the MANRS Core.
- (3) Find the list of ASes that are Tier 1 providers and MANRS ASes. Add them to the MANRS Core.
- (4) For each AS in the MANRS Core, find its customer ASes and add them to the MANRS Core if they are also MANRS ASes.
- (5) Repeat step (4) until the size of MANRS Core stops increasing.

- (6) Print the AS numbers, each separated by a newline, into `output.txt`. The ordering of AS numbers does not matter. The autograder expects the output file to be in the same directory as your `python` file. Example `output.txt`:

```
3
63
395326
40
396527
```

You will complete this task in `pa4q1.py`. Your code should be run by running the following command:

```
python3 pa4q1.py
```

## 2.5 Task 2: Map the MANRS Core ASes to Organizations (1pt)

Similar to PA2, parse the AS Organizations file and map the ASes to organization names. Print the organization names, each separated by a newline, into `output.txt`. The ordering does not matter. The autograder expects the output file to be in the same directory as your `python` file. Example `output.txt`:

```
Microsoft
Arelion
```

You will complete this task in `pa4q2.py`. Your code should run with the following command:

```
python3 pa4q2.py
```

# 3 Identifying Customers of the MANRS Core

## 3.1 Background

ASes that are direct customers of the MANRS Core may enjoy extra routing security benefits. In the FCC NOI response by Clark, Claffy, and Testart (see week 7 reading), there is a paragraph describe the MANRS Core direct customers:

*“Of the approximately 73K active ASes on the Internet today, about 23K are direct customers of a MANRS member that is part of the core. Of the 18K ASes that CAIDA identifies as being in the U.S., over 10K are direct customers of a MANRS member that is part of that core.”*

The goal of this section is to re-calculate the direct customers of the MANRS Core using more recent data. We will use a simplified calculation and you should expect the number of customer ASes to be smaller than the one above.

### 3.2 Task 3: Calculate the MANRS Core Customers (1pt)

In this task, your code is expected to do the following:

- (1) Parse the AS relationship file. Maintain a mapping of each AS to its customers.
- (2) Find all direct customers of the MANRS Core, i.e, of ASes in the MANRS core.
- (3) Print the AS numbers, each separated by a newline, into `output.txt`. The ordering of AS numbers does not matter. The autograder expects the output file to be in the same directory as your `python` file.

You will complete this task in `pa4q3.py`. Your code should be run by running the following command:

```
python3 pa4q3.py
```

### 3.3 Task 4: Identify the MANRS Core Customers based in the U.S. (1pt)

Similar to PA3, parse the AS Organizations file and map the ASes to organization names and its country. Find all ASes from Task 3 that are based in the U.S. Print the AS numbers, each separated by a newline, into `output.txt`. The ordering of AS numbers does not matter. The autograder expects the output file to be in the same directory as your `python` file.

You will complete this task in `pa4q4.py`. Your code should be run by running the following command:

```
python3 pa4q4.py
```

## 4 Week 7 Readings

Answer the 5 questions in the week 7 readings. You can find the week 7 readings here: [https://cseweb.ucsd.edu/classes/wi23/cse291-e/slides/week\\_7\\_readings.pdf](https://cseweb.ucsd.edu/classes/wi23/cse291-e/slides/week_7_readings.pdf)

### 4.1 Task 5: Answer week 7 reading questions (5pts)

Submit to PA4 Reading Summary on Gradescope.