OP2: A richer network abstraction

The current internet routing is relying IP packets being routed between BGP routers. An endpoint cannot have any direct impact on the path its packets take. One might want to intuitevely increase the size of routing tables to allow many alternate routes, that can be selected based on a new parameter inside the header of IP Packets, allowing it to bypass some countries . However, this has the drawback of being relativly costly and not scalable due to the potentially exponential grow of such a table. This can be solved by creating “regions”, providing an IP packet with a list of whitelisted/blacklisted region solves slightly the issue but still isn’t an optimal solution.

The idea would be to allow the endpoints to have to some degree the possibility to affect the route of packets they sent (and receive), while also allowing the network to restrict and controll slightly that freedom.

The implementation is inspired by SCION[1], a network protocol designed at ETH with such capability and better scalability. We can resume it to the following key points:

- IP Packets contain path segments
- Traditional routing is done from the source until the start of the path segment
- Routers simply forward the packet along the provided route when inside the provided path segment
- Traditional routing is done from the end of the path segment until the receipient.

This allows the endpoints to specify a path portion cutting out or avoiding networks as they desire, while avoiding the requirement for routers to hold gigantic routing tables .

More in details, the implementations would be as follows:

- From a data view, the main modification is adding a field to the IP packet holding path segments. Present as a list of list containing one more more path segments allowing the packet to be routed ignoring routing tables and instead following the provided path, avoiding the desired networks. Along with a marker indicating where the packet is in the provided paths (before, after, at a given node,...)
- From a control view, there are multiple modifications:
    - If the marker is outside of the path segments, we follow usual routing and forward packets to either the start of the next path segment or the receipient.
    - If the marker is inside one of the path segments, we forward the packet to the router provided by the packet and move the marker accordingly.
    - The receipient, to respect our routing desire, can invert the path segments when sending a packet back to follow a similar route back.

One drawback might be routers being down along the provided route. At that point, two solutions exists. Either we disreguard the provided path segment and try to route to the next node, potentially passing through a node we desired to avoid. Or we deem the packet unroutable, and dropping it.

However it has the advantage that endpoint control segments of the route and routers have in some cases an easier work to do, with the packet providint the next hop.

1 **SCION: A Secure Internet Architecture,** https://scion-architecture.net/pdf/SCION-book.pdf