Unique Prefix vs. Unique Mask for Minimizing SDN Flows with Transparent Edge Access

Josef.Hammer@aau.at, Hermann Hellwagner

Context Each regular (non-edge) destination IP would require a separate flow with *Transparent Access to Edge Services*

Challenge Flow memory in hardware switches is expensive and thus limited

Goal

Match many regular IPs with a single flow without matching a (virtual) edge service IP using a Patricia Trie

Unique Prefix / Mask

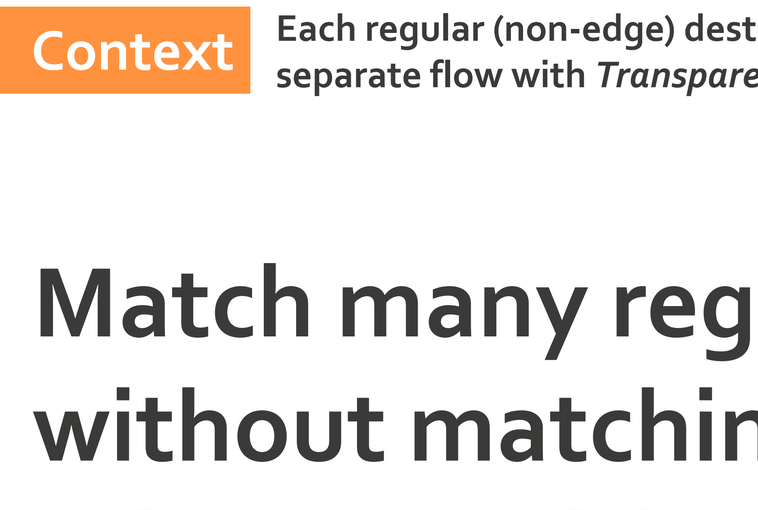

The *Unique Prefix* and the *Unique Mask* are used to generate a *match IP* for the switch flow. Both lead to a significant reduction in the number of flows by capturing many regular IP addresses with a single flow without also capturing a single virtual service address.

Algorithm

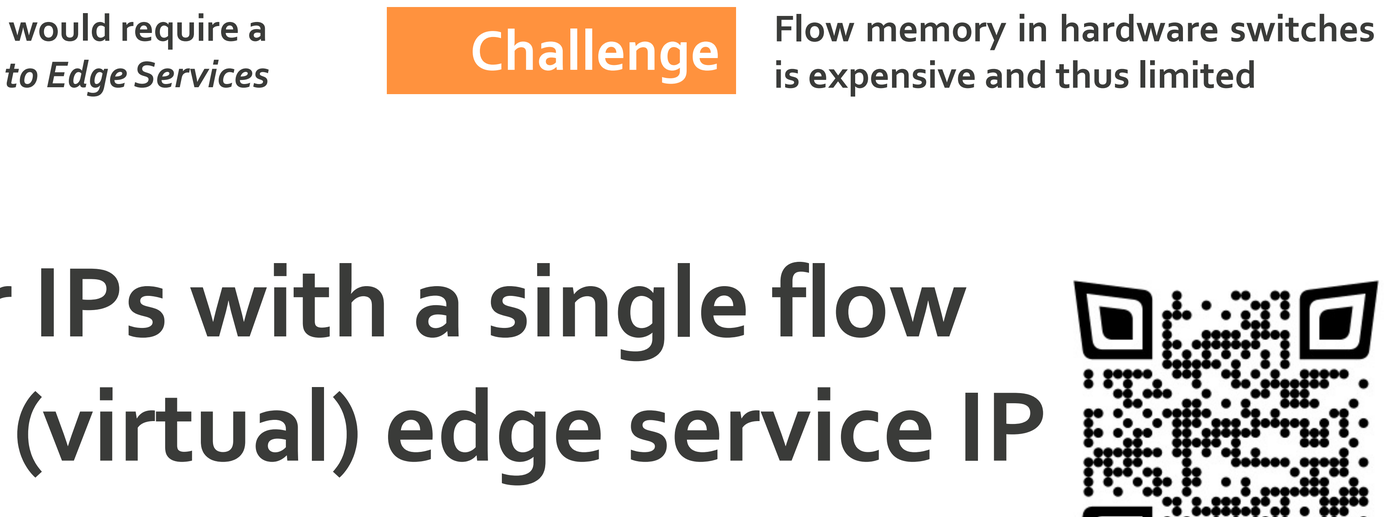

The *destination IP* is looked up in a Patricia Trie containing all registered virtual addresses. If the search ends up with a leaf node that does not contain the destination IP, we have a regular IP. This regular IP is then compared with the IP in the leaf node to calculate the *Unique Prefix* (shown above). By additionally using all parent prefixes to generate the IP mask, we get the *Unique Mask*. The generated *match IP* will not match any virtual address contained in the trie.

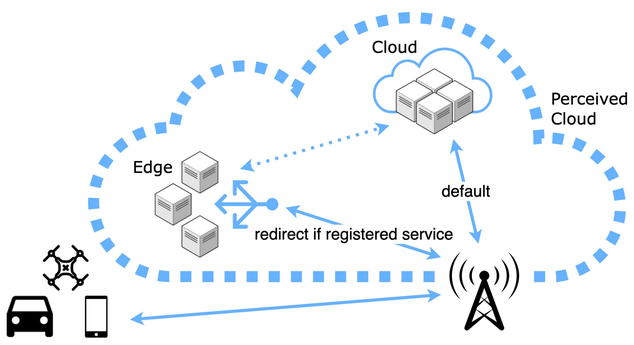

Transparent Access: All requests/responses look like cloud accesses to the client (UE) – the redirection to the edge is transparent

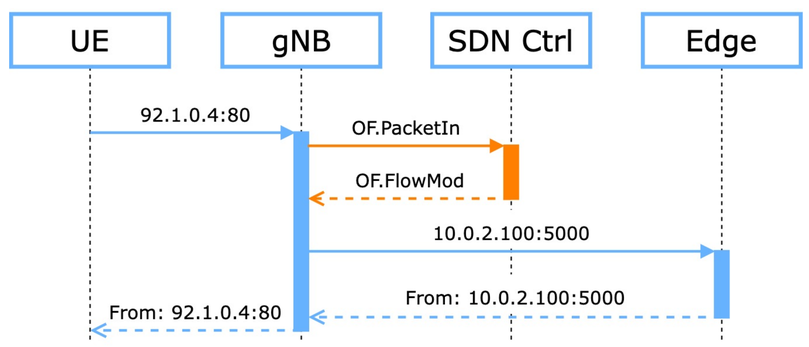

Routing with a registered service address (IP + port): Using OpenFlow (OF), the switch (gNB) transparently redirects the request to the edge server

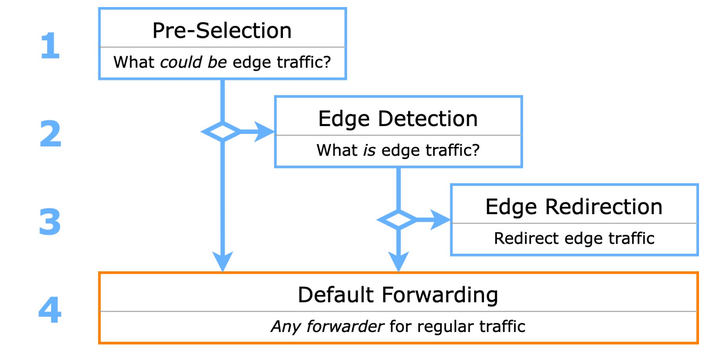

Four filter stages – our goal is to minimize the number of flows in the Edge Detection stage by matching many regular IPs with a single flow

Results

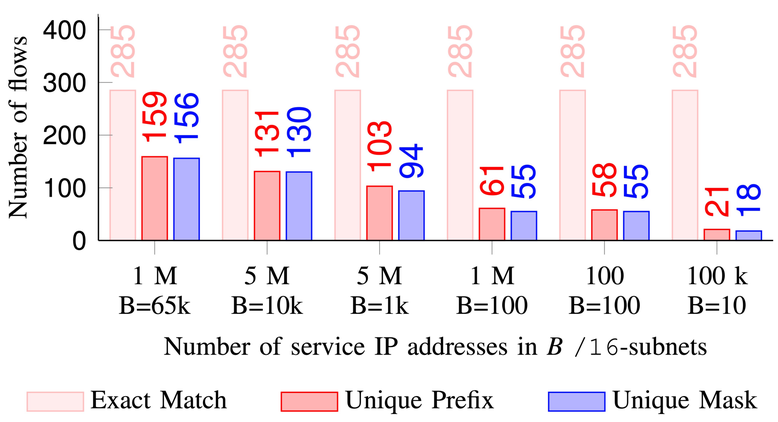

Number of concurrent flows (median) in the switch when replaying 269,616 requests to 1,595 public IP addresses from a five-minute real-world network traffic capture with an idle flow timeout of 5 seconds.

The random service IPs are distributed within a specific number of Class B (/16) subnets.

The random service IPs are distributed within a specific number of real-world subnets used by AWS/Azure/GCP.

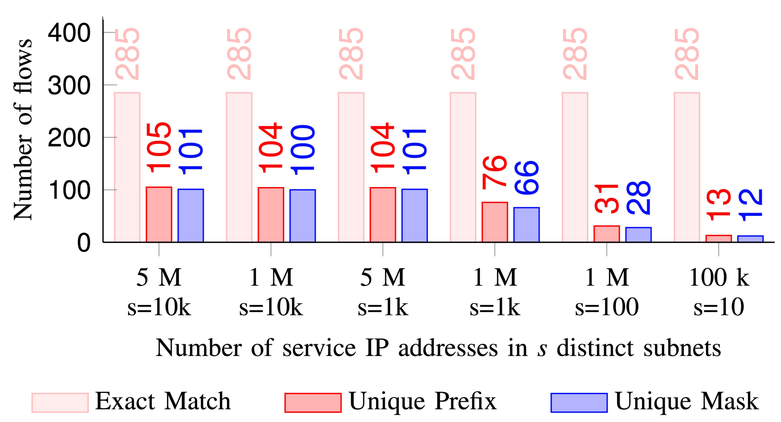

J. Hammer and H. Hellwagner, "Scalable Transparent Access to 5G Edge Services," 7th IEEE International Conference On Fog And Edge Computing (ICFEC), 2023