Configuration parameters:

1) smac and saddr to use for sending ARP requests and masquerading tap packets (two of them, one per bridge iface)
2) IP address of the gateway and the netmask
3) MACr, IPr – virtual router for TAP interface

	Bridge	**Tap**
ARP Request	Create or update ARP flow Forward to another interface	Reply with MACr if taddr = IPr Otherwise drop
ARP Reply	Update ARP flow and forward to another interface Drop if no matching ARP flow found[1]	Drop
ICMP Echo[2]	Create IpTuple from the packet Find IpFlow with the same ingress iface/tuple or create new To create a new flow: - use another port as oif - create egTuple by cloning igTuple *- resolve ICMP id clashes by allocating new id in egTuple*[3] - create a new flow: {iif, igTuple, oif, egTuple} Take egTuple from found or newly created flow Transform the packet according to the egTuple Forward the packet to oif interface	If it is Echo Request to IPr, just reply with MACr. Otherwise: Create IcmpTuple from the packet Find IpFlow with the same ingress tuple/iface or create new To create a new flow: - resolve[4] daddr to get tmac and oif - choose smac and saddr depending on oif - create egTuple from igTuple: translate smac, saddr, dmac *- resolve ICMP id clashes by allocating new id in egTuple* - create a new flow: {iif, igTuple, oif, egTuple} Take egTuple from found or newly created flow Transform the packet according to the egTuple Forward the packet to oif interface
TCP/UDP	Identical to ICMP, but use sport/dport instead of ICMP id	
Other ICMP	*Extract the offending IP packet from ICMP packet, create a tuple from it* *Find IpFlow with the egress iface matching the iface ICMP packet came from and egress tuple matching offending one* *If no flow found, drop the packet.* *Otherwise, take igTuple from the flow and transform the ICMP packet and the offending IP packet* *Reassemble ICMP message and forward it to iif*	

1 We cannot build full ARP flow entry from single ARP reply packet because we are don't know where the original ARP request came from.

2 Currently we don't distinguish between Echo Requests and Replies.

3 Clashes may occur if two computers (one connected to bridge interface and another to tap one) decide to use the same ICMP id to ping the same target.

4 Resolution is done by searching for ARP flows having the same saddr or taddr. Next we check if daddr is not directly connected and use gateway as daddr. We send out ARP requests (with masqueraded smac and saddr) and drop the packet, because we don't know where to forward it to. Better approach is to enqueue the packet try to process it again when we receive ARP reply.