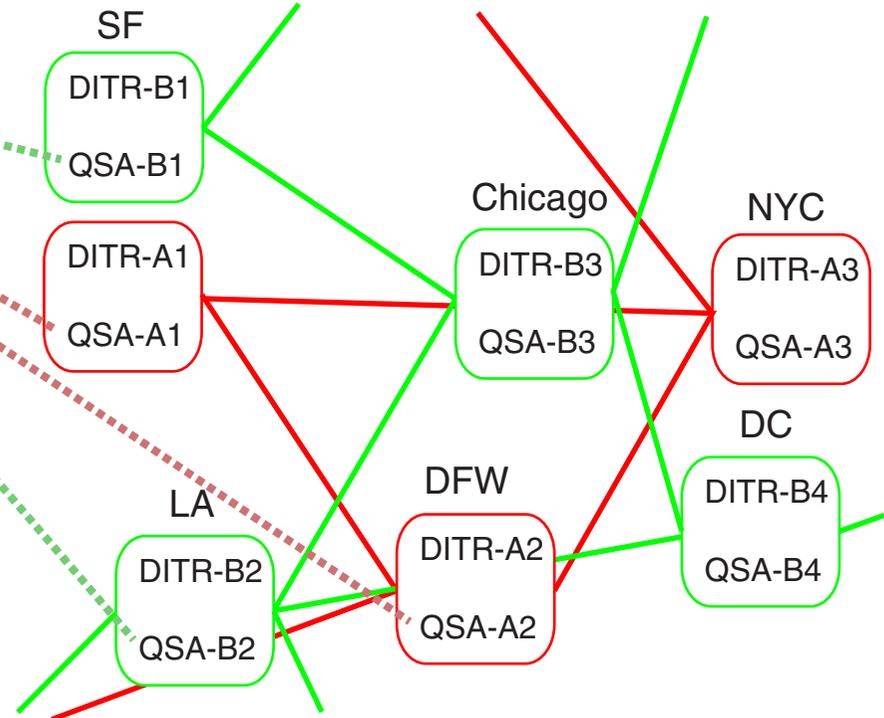
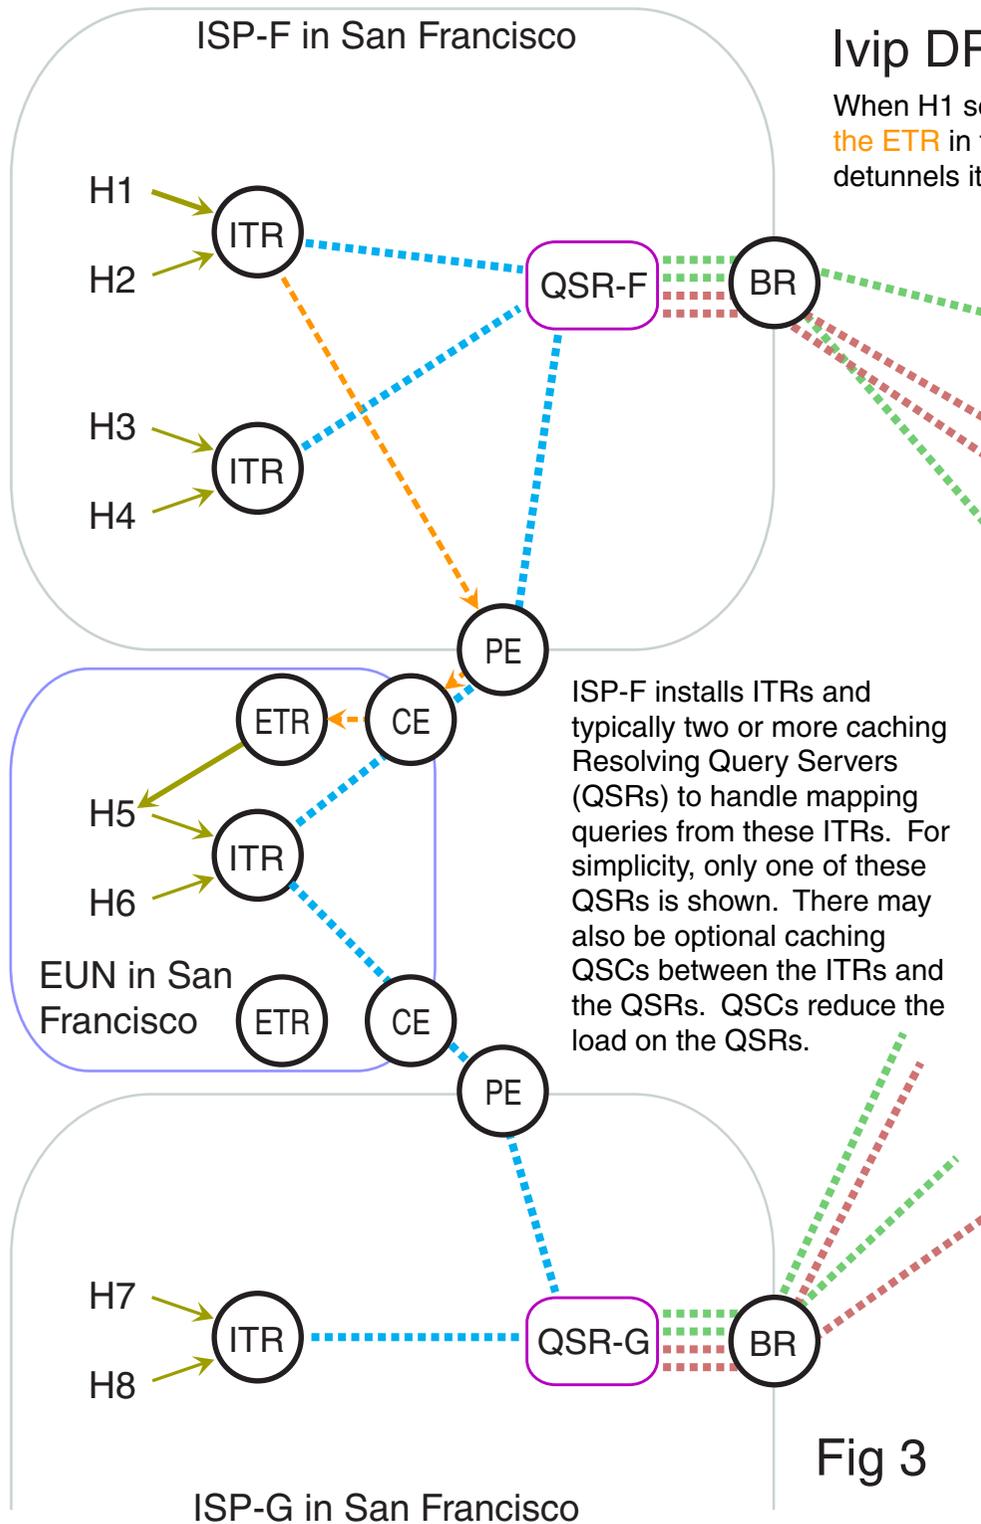


# Ivrip DRTM Stage 2: Add ITRs with caching QSRs

When H1 sends a packet to H5, it travels to an ITR in ISP-F's network, which tunnels it to the ETR in the EUN (End User Network) which uses the ISP-F PA address. The ETR detunnels it and forwards it within the destination network to the H5 destination host.



Assuming neither the ITR or QSR have cached mapping for a micronet which covers H5's address, the ITR queries the QSR (blue broken line), and the QSR sends a similar query (one of the green or brown broken lines) to a typically nearby QSA. Map Replies and any later Map Updates are secured by a nonce from the query. The QSR obtains mapping quickly (a few tens of milliseconds) and reliably from one of two or so typically nearby Authoritative (full-database, continually real-time updated) QSA servers, at DITR sites of the DITR network which handles the MAB which covers the EUN's SPI space. Each QSR has previously chosen two or so typically nearby QSAs for each MAB.

If the mapping changes during the caching time, the QSA sends a Map Update message securely to its querier - the QSR - and the QSR repeats this process to its one or more queriers (QSCs or ITRs). Any QSCs do the same. So all ITRs which are tunneling packets receive updated mapping within a fraction of a second of it being changed at the DITR site's QSA - which will typically be within a fraction of a second of the EUN, or some company it appoints, changing the mapping.

Fig 3