Dual-Stack Connections in 3GPP Networks

Jari Arkko and Fredrik Garneij, Ericsson

Goals for This Presentation

1. Introduction to 3GPP network access architecture
2. Evolution of the IPv6 access mechanisms in 3GPP
3. Reasons for adopting the dual stack connections

3GPP Access Network Architecture

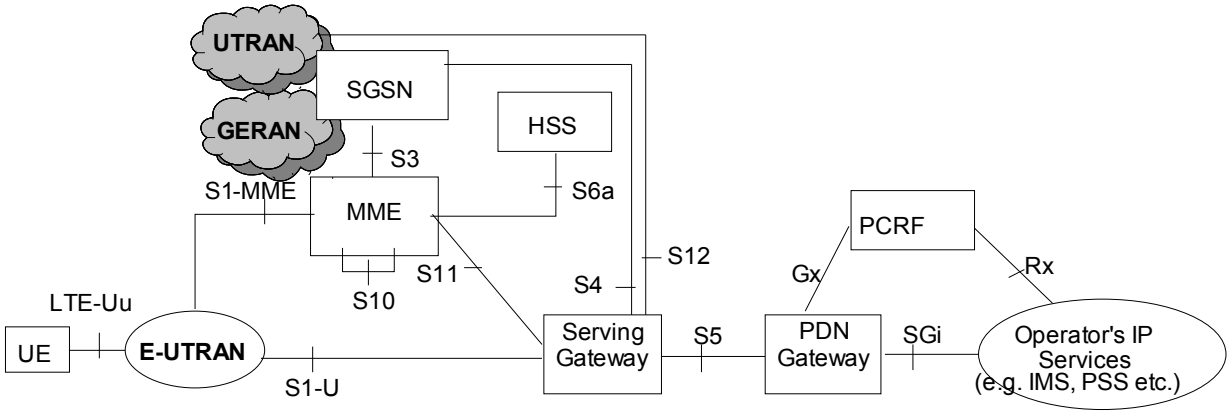

- End user equipment (UE) asks for a *connection*
- A *PDN connection* holds addresses, QoS info, APN, ...
- PDN connections are unaffected by movements; the *PDN Gateway* acts as an anchoring node and 1st hop router
- End user equipment sends native packets over radio; network tunnels packets to the PDN Gateway

Observations

- IPv6 deployment becomes mostly just a matter of endpoint changes – as long as the UE and PDN Gateway support and use IPv6, the underlying network stays unaffected
  - Upgrading the routers on the path is not an issue
- Much work still remains: making sure the operator's gateways are connected to the IPv6 Internet, that UEs support IPv6, ...

Evolution of the Connection Model

1. IPv6 support was originally defined in 1997 as part of general work for the network architecture – separate PDP context type was used (PDP Type IPv6)
2. In subsequent years there were potential use cases and interest; network nodes started supporting the IPv6 PDP context type in 2003
3. Despite some trials there was little usage. This changed from 2008 onwards with the IPv4 situation
4. With more attention, 3GPP worked hard on evolving and finetuning its architecture, determining recommended deployment models, looked if it needed new transition tools, etc.
5. A new PDN connection for PDN type IPv4v6, sometimes referred to as Dual-Stack PDP context/bearer was introduced in 2008 as a part of EPC in 3GPP Release 8. Legacy GPRS architecture adopted this also in 3GPP Release 9.

Motivation for the Dual Stack Connections

The 1990s model for connections was to have a separate connection for IPv4 and IPv6. This is still supported.

Why was it important to add a dual-stack connection? Why is the dual-stack connection now the recommended approach?

1. Dual stack migration strategy
2. Resources
3. Signaling
4. Load balancing side-effects
5. Operational problems
6. Charging

Resources and Signaling

- It is likely that most parameters in the connection are the same across IP versions – two connections may lead to resource waste
- A more serious issue is duplicating the need for signaling for connection setup, movements, QoS, filter updates
  - This is particularly bad over radio
  - And in most cases, we are doing this for no benefit

Operational Side-Effects

- If there are two separate connections, they may end up in different devices in a load-shared gateway farm

  This prevents implementations that share storage

- If something does not work, debugging is harder: two connections, possibly two different gateways, two resets needed to restore the situation, etc.

Charging

- Evolution in payment models and charging technology has given us diverse tools to deal with charging
  - Pre-paid accounts, credit limits and other innovations
- Charging operations for different connections but still part of the same service may be difficult
- Charging operations from different load balanced gateways would be even harder

Conclusions

- 3GPP learned some lessons the hard way
- Dual stack connections would have been much simpler

  The historical baggage from these lessons is heavy – the system could be a bit simpler if we had understood everything at the beginning
- Hopefully you can apply our experiences in your work!
  - Note that not all issues may apply to BBF networks
  - But some probably do
  - Fixed mobile-convergence?

ERICSSON