Load balancing for SIP media applications

Fostel 2007

Raphael Coeffic (raphael.coeffic@tekelec.com)

Problem statement

- Media servers consume usually **a lot** of ressources (SEMS is not an exception).
- It is then more common to have a **media server farm**.
- How are we supposed to implement the **load balancing** in such a farm?

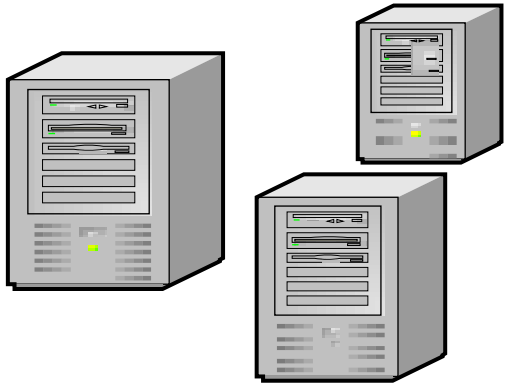

Different solutions

TEKELEC

- At the IP layer:
  - use traditional **IP load balancers** (UDP/TCP).
  - You will need at least some **basic support from the SIP side** to get that to work.
  - You could have troubles with NATs or firewalls.

- SIP-wise:
  - **leverage SIP routing** to spread the requests.
  - You will need a **very fast** SIP proxy / load balancer.
    - Guess what? **SER** is doing more than **5k transactions / second**.

SIP load balancing

TEKELEC

- Background
  - Media applications are mostly addressed by request URI (as recommended by relevant RFCs).

- How?
  - Round-robin style:
    - You may also want to implement some more complicated stuff, maybe based on load feedback.
    - Requests are just spread, and the load balancer does not need to stay in the signaling path.
  - Hashing based method:
    - Some value in the SIP message is hashed to determine the target.
    - The load balancer could even sit on the signaling path and stay stateless.
    - Statistics sometimes play tricks if the population is small.

Example: Voice2email

Round-robin load balancing

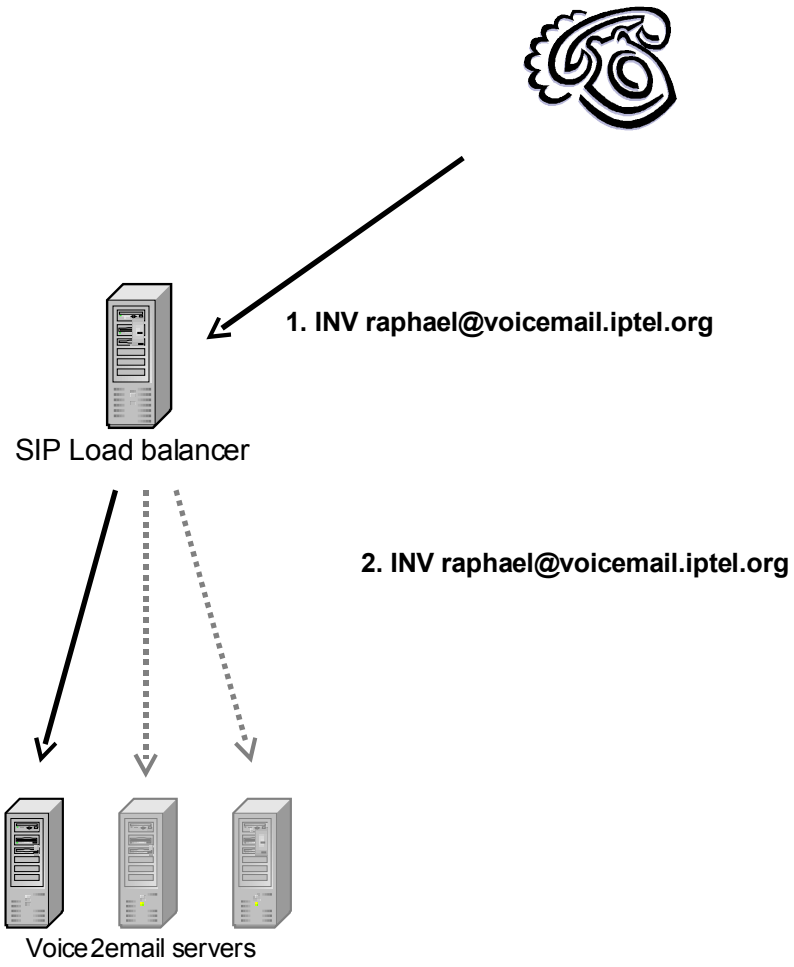

Example: PSTN conferencing

Hashing based load balancing

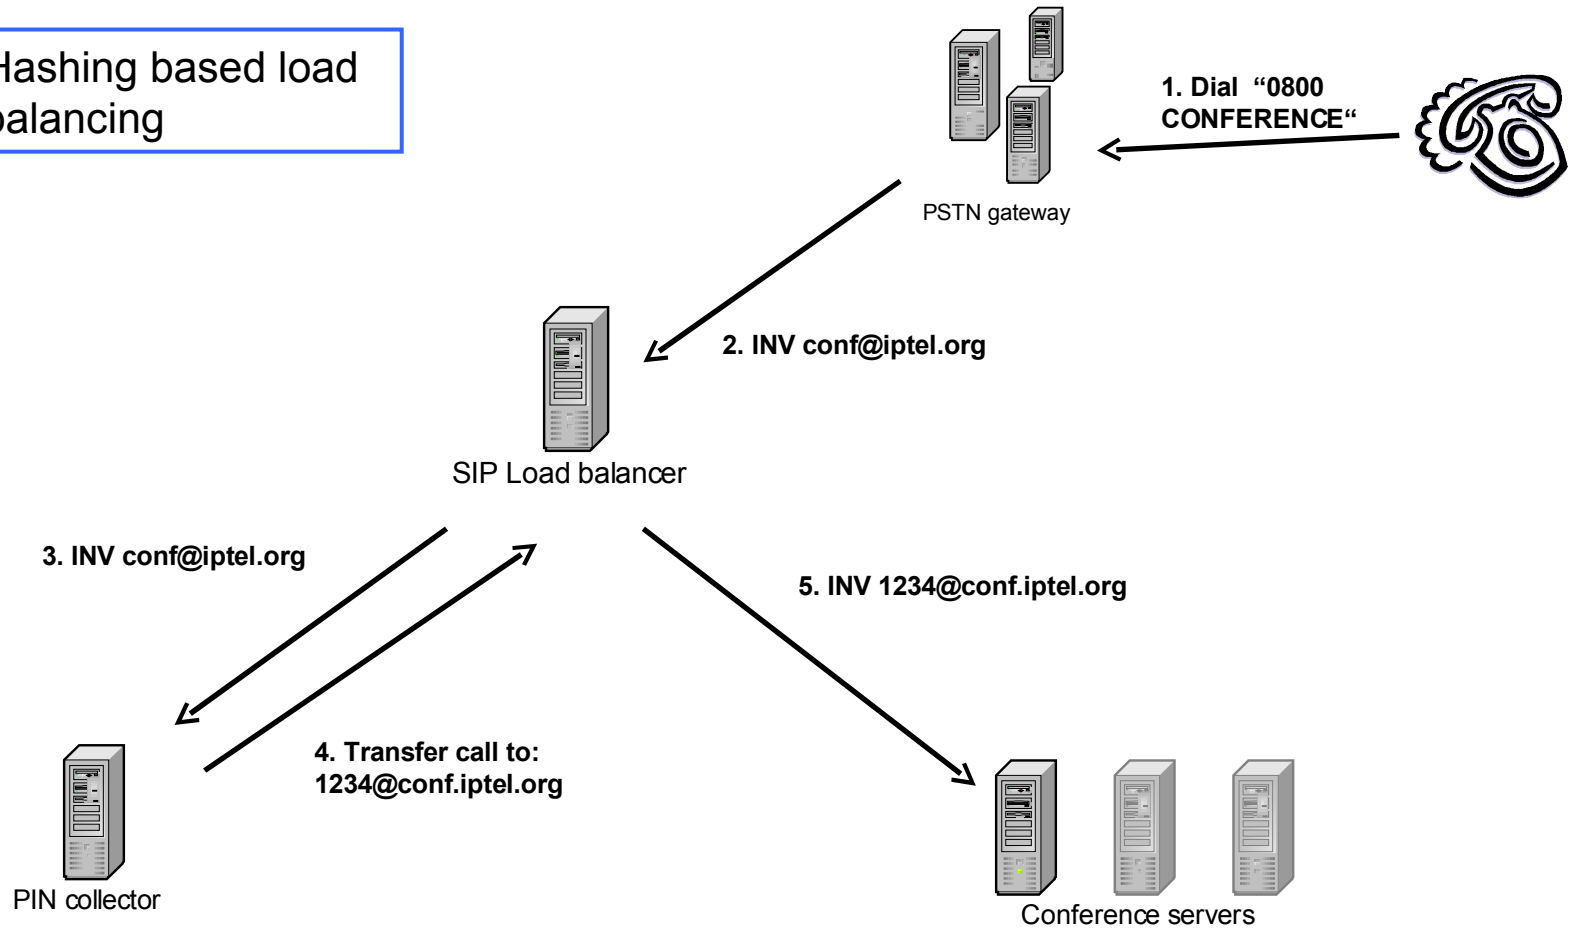

Thank you!
Questions?

Fostel 2007

Raphael Coeffic (raphael.coeffic@tekelec.com)