

DCON (Distributed Conferencing) BOF @ IETF 82

Feasibility analysis


Lorenzo Miniero

Taipei, Taiwan, November 14, 2011

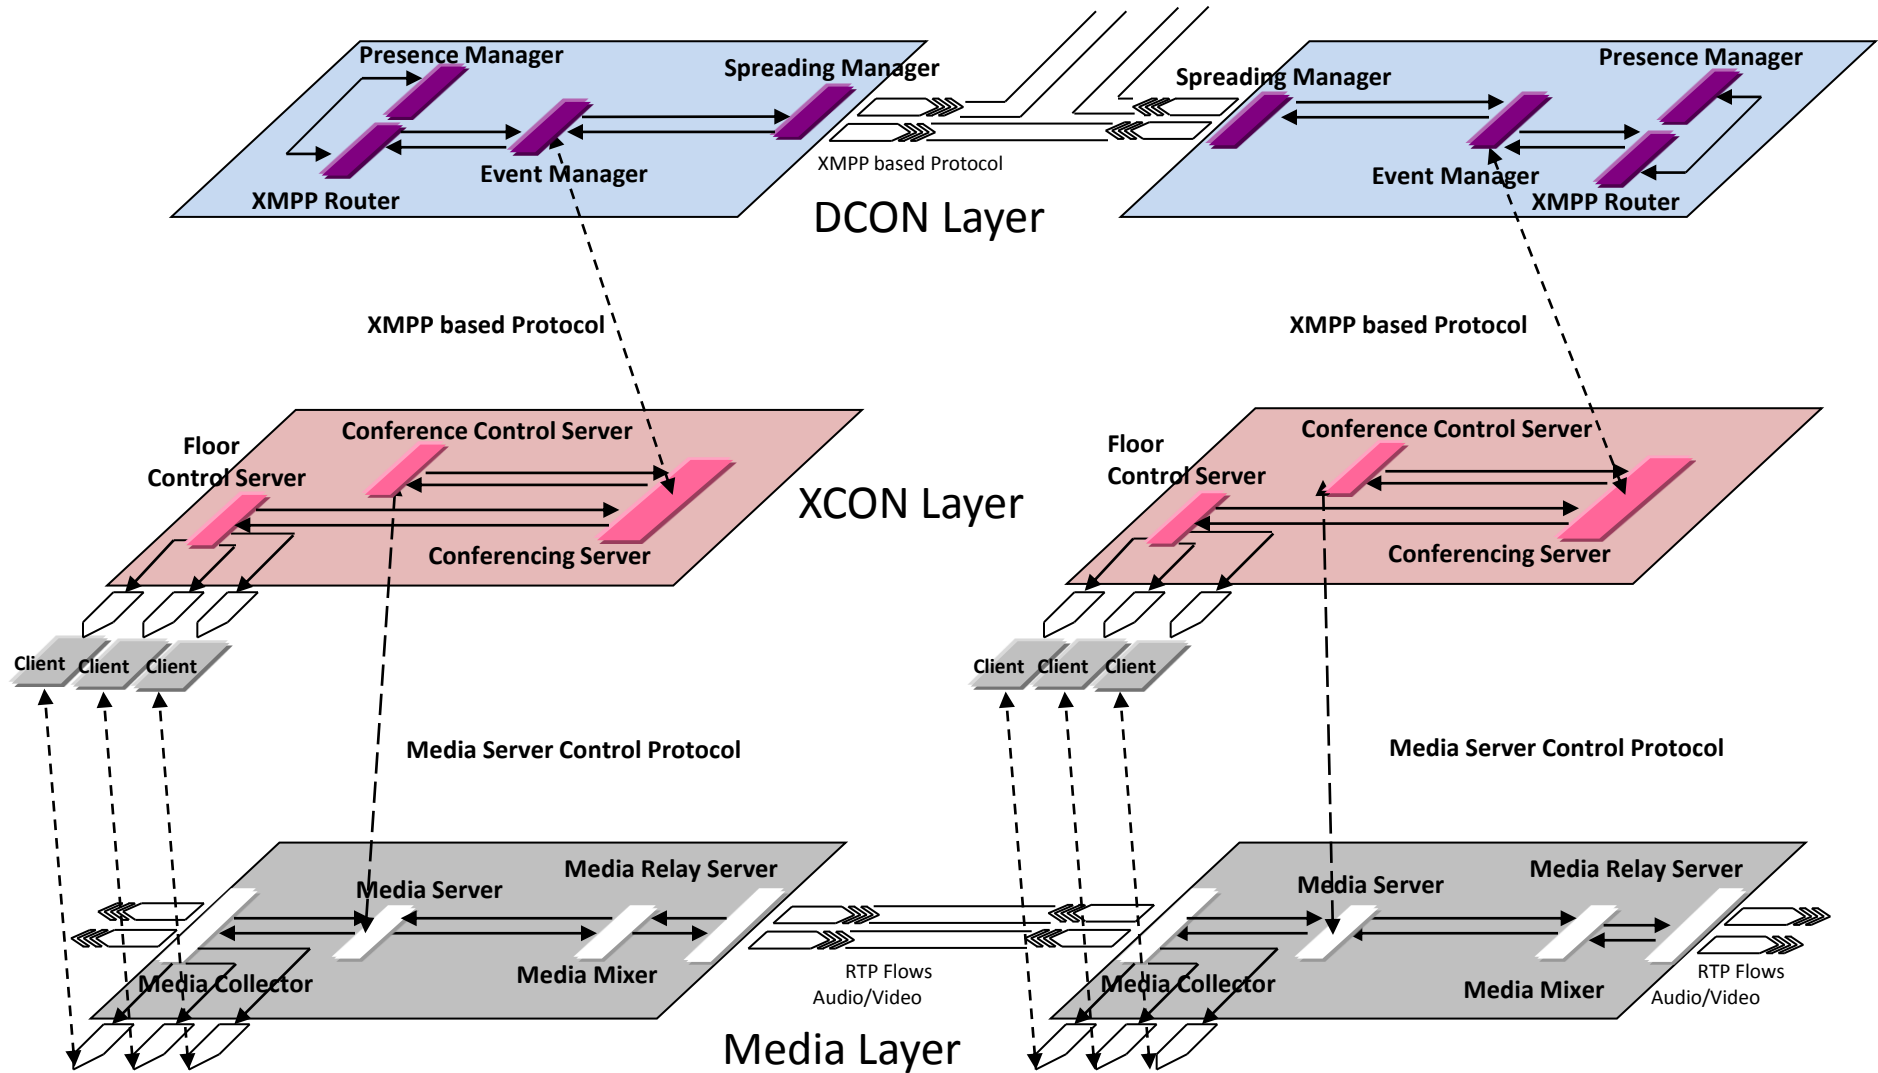
DCON feasibility analysis

- Proof-of-concept prototype realized by extending the Meetecho platform
- Focus on **scalability**
- Exploits XMPP Server-to-Server (S2S) channels for the overlay network
 - Spreading of conferences information and events
 - Dispatching of centralized protocols (e.g., BFCP)
- Leverages presence information for focus discovery
- BFCP-driven local mixing

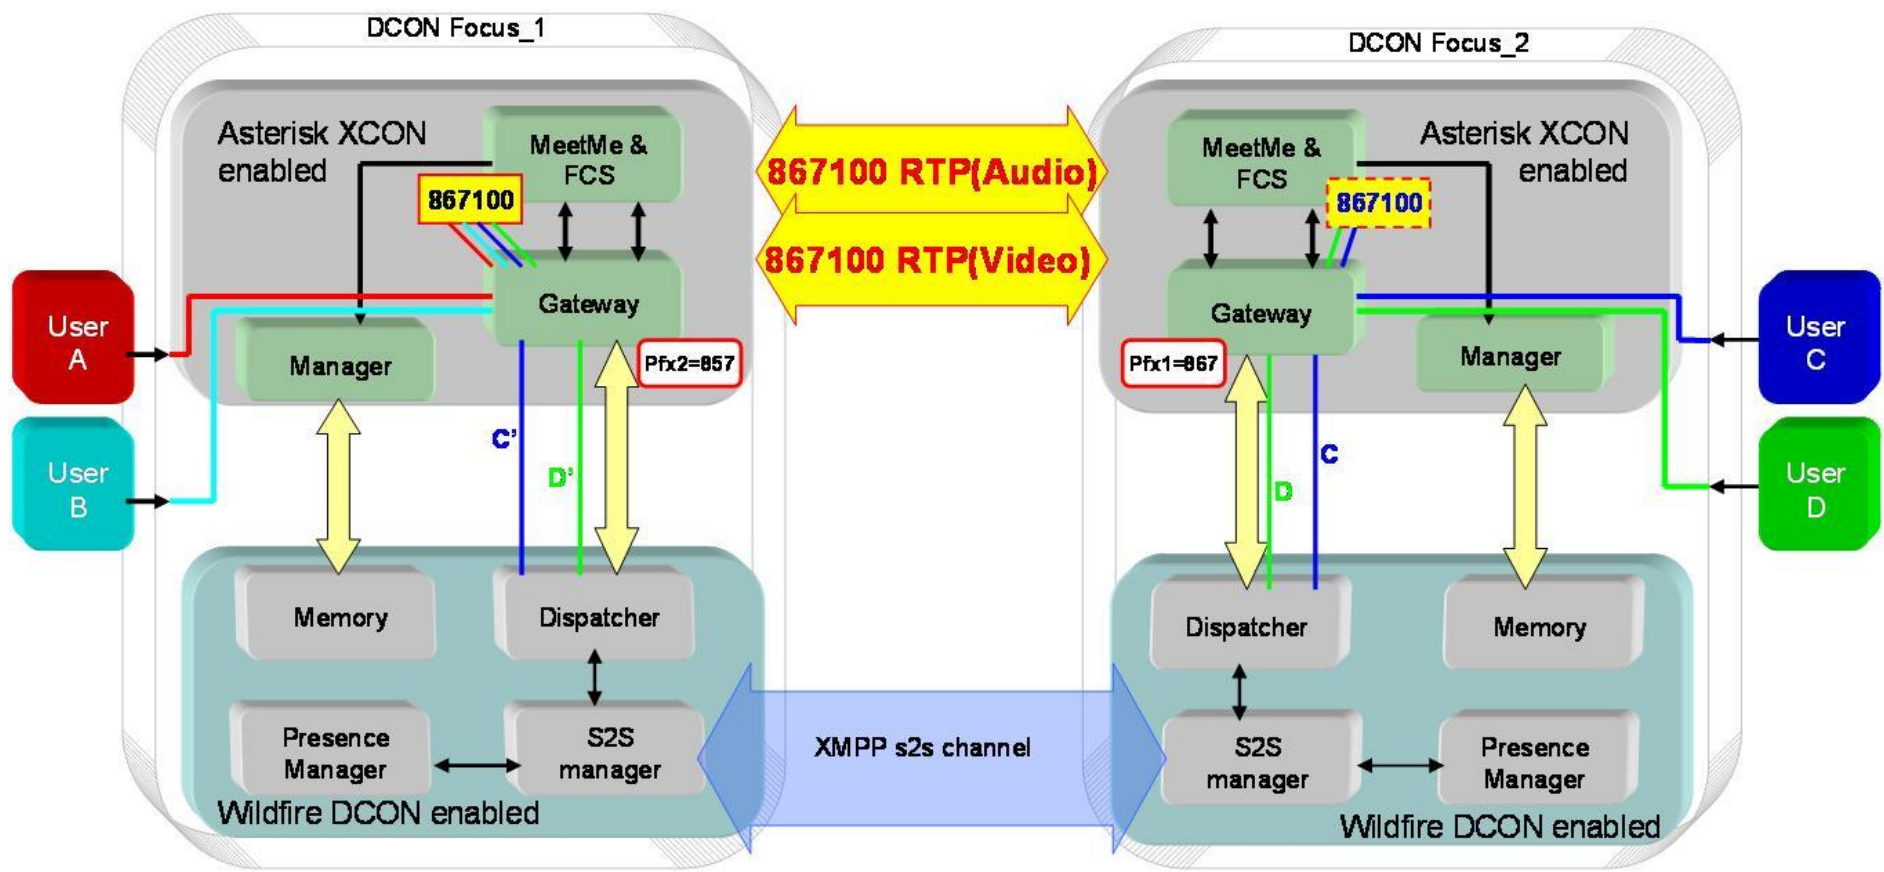
DCON layering structure



DCON layers interaction




Proof-of-concept implementation



Performance assessment: Centralized case

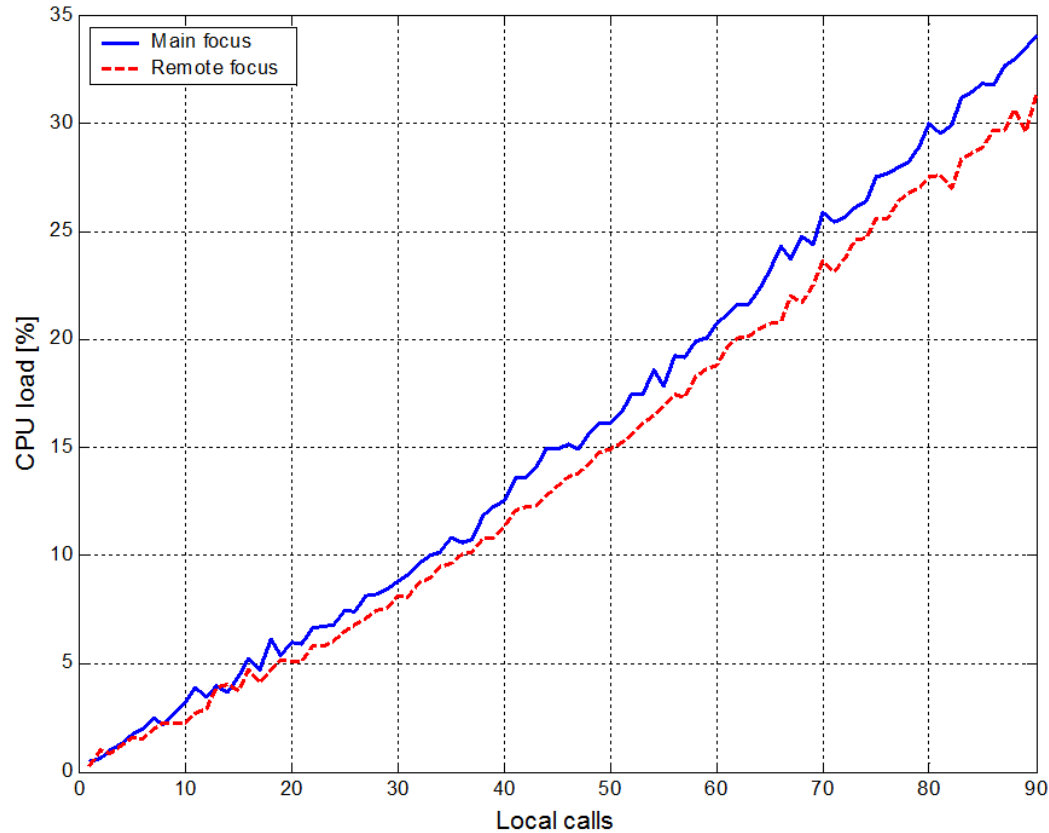
- Monitored parameter: CPU load of the focus/foci
- Each user requests and obtains the audio floor



- 180 as the peak value in the presence of BFCP functionality
- Might be quite restrictive
- A benchmark for the following tests


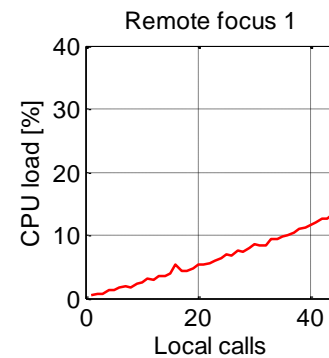
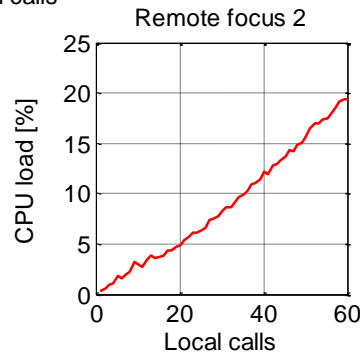

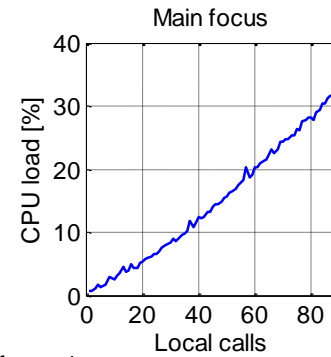
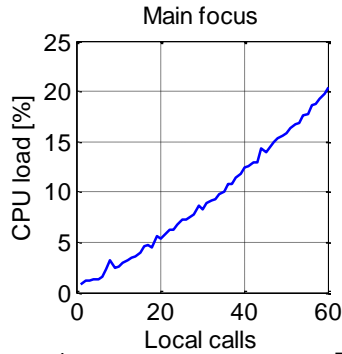
Focus	Calls	CPU load (%)
Main	180	≈100

Performance assessment: 2 islands case



Focus	Calls	CPU load (%)
Main	90	34.0
Remote	90	31.6

Performance assessment: 3 islands case



Focus	Calls	CPU load (%)
Main	60	21
Remote 1	60	20
Remote 2	60	20

Focus	Calls	CPU load (%)
Main	90	34.4
Remote 1	45	13
Remote 2	45	13

Scalability: figures in summary

Islands	Local users	Remote users	Main focus CPU load	Remote focus 1 CPU load	Remote focus 2 CPU load
1	180	-	≈100%	-	-
2	90	90	34%	31.6%	-
3	60	120	21%	20%	20%
3	90	90	34.4%	13%	13%

- Migration towards a distributed paradigm allows for a huge reduction in the load of the primary focus
- The sum of the CPU levels of all involved foci is less than the CPU level of the single focus in the centralized case
- Given a fixed number of local users, remote users distribution among multiple islands adds negligible overhead to the main focus

Considerations

- Distribution of components brings to a considerable improvement in terms of CPU load
- The study we presented just focused on **scalability**, but...
- ...what about other functionality?
 - **Load balancing:**
 - Fairly (and transparently) distribute users among a set of available conference servers
 - **Resiliency:**
 - Transparently migrate users to a new server should the one they are currently exploiting experience a fault
 - **Federation:**
 - Allow for heterogeneous servers (i.e. belonging to different vendors/organizations) to smoothly interoperate