

OsmoMGW - Feature #4551

Dynamic/runtime allocation of new endpoints

05/14/2020 10:56 PM - ipse

Status:	New	Start date:	05/14/2020
Priority:	Normal	Due date:	
Assignee:		% Done:	0%
Category:			
Target version:			
Spec Reference:			

Description

Right now endpoints are pre-allocated on the osmo-mgw startup and their number can't be changed without restarting the process.

Reasoning: In a situation when osmo-mgw is hit with more load than expected, it can run out of available endpoints with no way to increase them. Your only solution right now is to change the configuration and restart BSC+MGW processes which leads to the service disruption in the middle of a peak hour.

Decreasing number of endpoints might be nice to have too but it might be more difficult to implement, given that MGW would need to wait for the endpoints to become idle to deallocate them. And it looks like decreasing is less important.

Increasing the number of endpoints can be either (1) manual from VTY or (2) automatic/dynamic.

In the dynamic case, we configure a maximum amount in the configuration file. Then osmo-mgw allocates new endpoints if it can't find an available one. But only up to the configured maximum amount. This way we can default to a significantly large default maximum value and don't bother about this until we hit physical server limitation.