Brief notes on my understanding of “LLDP Auto-attach” proposal

AAC: Auto-Attach Client (typically the IoT device which wants to attach)
AAS: Auto-Attach Server (typically the Backbone Edge Bridge providing core connectivity)

The envisaged operating mode is that:

1. AAC devices and AAS devices announce their presence to directly attached devices through LLDP, including an Element TLV which contains an element type (client or server), a 10-octet system ID and a management VID.
2. Optionally, AAC devices can receive the management VID info from the AAS and attach to that management network using DHCP. This would allow configuration before attaching to the main network.
3. The next part is not clear from the documentation, but I think the AAC and AAS use LLDP Service Assignment TLVs to advertise things to each other, and the content changes according to state machines operating in the background. Remember that LLDP is not a request/response protocol but just an announcement protocol, and so using it for things like this is a bit of a hack.
4. The AAC announces to the discovered AAS what mappings from VID to ISID it wants the AAS BEB to configure for it. The AAS announces (to anyone listening on the port) its VID to ISID mappings and their assignment status.
5. When the AAS receives "requests" from an attached AAC (to map certain VIDs to certain ISIDs; typically to map one particular VID to one particular ISID), it decides whether to accept them and changes its advertised mappings accordingly. If accepted, it makes the relevant mappings, "connecting" the attached client to the relevant core network.
6. The mappings time out if the LLDP frames stop coming.

The AACs are programmed by the administrator with a set of mappings between VLAN and ISID. This means that knowledge of the network context has to be programmed into each end-node prior to attaching to the relevant core network. This could be done by static configuration or through the management VLAN.