Harborline draft AI principles (internal, not a standard)

Working draft, 2027-05. Not board-approved.

1. We will know **purpose, data, and human role** before expanding an AI use.
2. We will not use AI to make **automated decisions that affect members** without a named owner and review. (The draft does not define "decision.")
3. Customer content is not used to train vendor models unless Privacy records an exception.
4. Staff remain accountable for messages sent in Harborline's name.
5. We reassess when **purpose, data, model, automation, or oversight** changes.
6. Marketing claims are not controls.

Use these **after** you have a use-case map and a governance plan, when you map principles and NIST AI RMF functions. They do not replace understanding Member Assist.