An evaluation checklist for an on-premises API gateway

Thirteen questions to ask any vendor, G8 included, and where the answer is visible in G8: on the site, or in the console during a proof of concept. Two questions no page can answer come first: how does it behave under your traffic, and how does the vendor answer when something breaks? Bring this list to the demo; we go through it row by row on a real install.

- [ ] **Where does the management plane run, and does it work without internet access?**
  Where to see it: Site: Deployment. Console: the Environment form in a closed network; the signed license file.

- [ ] **Who changed what, and was the change recorded before it was applied?**
  Where to see it: Console: the audit log; a change, a refused change, before and after values, the request id.

- [ ] **Can a role be limited to one environment, and does it come from our directory?**
  Where to see it: Console: a custom role's grid; Active Directory groups mapped per Environment; a Viewer's refused save.

- [ ] **How is configuration promoted from staging to production?**
  Where to see it: Console: Copy to environment; the preview of dependencies, dropped secrets and changed lines; the record in both Environments.

- [ ] **Which identity methods work on one route at the same time?**
  Where to see it: Console: the openid-connect, ldap-auth, saml-auth, jwt-auth and key-auth forms; multi-auth on one Route.

- [ ] **Can limits be shared across gateway nodes?**
  Where to see it: Console: limit-count with the Redis policy; the 429 in Traffic Explorer; the 4xx share in Analytics.

- [ ] **Where do request records and metrics live, and who can read a body?**
  Where to see it: Console: Traffic Explorer settings; your Elasticsearch or OpenSearch, retention days, masking; the body permission and its audit entry.

- [ ] **Do analytics match our own Prometheus?**
  Where to see it: Console: Analytics beside a Prometheus query for the same window. Site: Why G8, how we verify.

- [ ] **What does the gateway do for LLM traffic?**
  Where to see it: Console: ai-proxy-multi with two instances and a forced failure; ai-rate-limiting on tokens; ai-prompt-guard.

- [ ] **Can our existing APIs be served to AI agents without new code?**
  Where to see it: Console: openapi-to-mcp on a Route; the tool list an MCP client sees; key-auth and a limit on the same Route.

- [ ] **Which versions of our stores are supported, and for how long are records kept?**
  Where to see it: Site: Deployment. The Deployment and Security Brief, page 3: stores, retention.

- [ ] **What does the installation need from us, and how is it secured?**
  Where to see it: The Deployment and Security Brief: what you provide; non-root image, restricted profile, AES-256-GCM for stored keys, your CA.

- [ ] **What does it cost, and what is in each edition?**
  Where to see it: Site: How to buy. The Editions and Licensing one-pager; ask which edition the proof of concept runs.