

CLAUDE COMPETITOR BATTLECARD

Target: Anthropic Claude (Enterprise/API)

⚠ Landmine Questions (Expose Weaknesses)

? Data Sovereignty & Privacy

"How do you ensure our proprietary code and customer data never leave our VPC for training or inference?"

- Claude requires API calls to Anthropic servers (unless using AWS Bedrock, which still has limitations)
- Boudica offers 100% air-gapped, on-premise deployment
- Zero data leakage risk — no training on customer data

? Cost Predictability & Scaling

"What happens to our monthly bill if we scale from 10 to 100 developers?"

- Claude API is token-based — costs scale linearly with usage
- Boudica is per-user licensing — predictable, flat-rate costs
- No 'surprise' bills from unexpected usage spikes

✅ How We Win (Key Differentiators)

1. Full Stack Control

We provide the model, the inference engine, and the database. Claude is just an API. We offer complete control over the entire AI stack.

2. RAG Performance

Our hybrid search (keyword + semantic) outperforms Claude's standard RAG. We retrieve more relevant context, leading to better answers.

3. Cost Efficiency

For high-volume users, Boudica is 5-10x cheaper than Claude API. Our per-user model eliminates token-based cost anxiety.

4. Customization

We can fine-tune models on your specific domain data. Claude is a generalist; we are a specialist for your business.

📊 Quick Comparison

Feature	Claude	Boudica
Deployment	Cloud Only	On-Prem / VPC
Pricing Model	Per Token	Per User
Data Privacy	Shared Infrastructure	Isolated / Air-Gapped
RAG Quality	Standard	Hybrid (Keyword + Semantic)
Customization	Limited	Full Fine-Tuning