

⌘ AgentK Chat Platform — Delivery Guide v1.0

Product: AgentK Personal License (\$29 one-time)

Version: 1.0 Beta

Delivery Date: June 15, 2026

Support: support@tikkunolam.today

⚠ IMPORTANT: FEATURE STATUS

Feature	Status	Availability
Chat Interface	✓ WORKING	Available now
Cognitive LLM Routing	✓ WORKING	Available now
Model Registry	✓ WORKING	5 models tracked
Memory Persistence	✓ WORKING	Sled-based storage
Security Layer	✓ WORKING	AES-256-GCM encryption
Autonomous Agent Loop	☐ COMING Q3 2026	Roadmap
Mission Execution	☐ COMING Q3 2026	Roadmap
Auto-Revenue Generation	☐ COMING Q3 2026	Roadmap

What This Means: You can use AgentK as a sovereign chat interface today. The autonomous “agent swarm” features are under development and will ship in Q3 2026. You will receive free updates when they launch.



WHAT YOU RECEIVE

1. Binary Download

- Pre-compiled Linux binary (agentk executable)
- Dashboard web interface (built-in)
- No installation required — run directly

2. Access Credentials

- Dashboard URL (default: `http://localhost:5001`)
- Admin API key (for programmatic access)

3. Documentation

- This guide
 - API reference
 - Model configuration instructions
-

QUICK START (3 Minutes)

Step 1: Download the Binary

After purchase, download from: - Gumroad Downloads page (your library) - Or request GitHub repo access for source build

Step 2: Run AgentK

```
# Make executable
chmod +x agentk

# Start the service
./agentk --dashboard-port 5001

# Expected output:
# 🤖 AgentK v1.0 Beta starting...
# 📡 Dashboard: http://localhost:5001
# ✅ Ready
```

Step 3: Open the Dashboard

Navigate to: `http://localhost:5001`

You should see: - Chat interface - Model selection dropdown - Conversation history sidebar

Step 4: Send Your First Message

1. Select a model from the dropdown
 2. Type a message in the chat box
 3. Press Enter or click "Send"
 4. Receive AI response
-

CONFIGURATION

Environment Variables

```
# Optional customization
AGENTK_PORT=5001           # Dashboard port
AGENTK_HOST=0.0.0.0
# Bind address (0.0.0.0 = all interfaces)
AGENTK_MODEL_PATH=/path/to/models # Custom model
directory                  # Custom model
AGENTK_DB_PATH=./agentk.db    # Memory database
location
```

Model Configuration

AgentK supports multiple LLM backends:

Model	Provider	Use Case
default	Local Ollama	General chat
gemma4	Cloud	Complex reasoning
qwen3.5	Cloud	Code assistance
phi4-mini	Local	Lightweight tasks

Configure in ~/.agentk/config.toml:

```
[models]
default = "ollama/gemma4:e4b"
code = "ollama/qwen3.5:4b"
lightweight = "ollama/phi4-mini:latest"
```

API REFERENCE

Health Check

```
curl http://localhost:5001/health
# Expected: {"status":"ok"}
```

Send Chat Message

```
curl -X POST http://localhost:5001/api/chat \
-H "Content-Type: application/json" \
-d '{
  "message": "What is SoulTech?",
  "model": "default",
  "conversation_id": "my_conversation_123"
}'

# Expected: {
#   "response": "<AI response text>",
#   "conversation_id": "my_conversation_123",
#   "model_used": "default"
# }
```

List Available Models

```
curl http://localhost:5001/api/models
# Expected:
[{"name":"default","provider":"ollama"},...]
```

Get Conversation History

```
curl "http://localhost:5001/api/memory?
conversation_id=my_conversation_123"
# Expected: {"messages":[...]}
```

MEMORY & PERSISTENCE

AgentK stores conversations locally using Sled (embedded Rust database).

Location: ~/.agentk/agentk.db

Backup: Copy this file to preserve conversation history.

Privacy: All data stays on your machine. No cloud sync unless you configure it.

SECURITY

Feature	Implementation
Data Encryption	AES-256-GCM for stored conversations
API Authentication	Optional API key header
Network Binding	Configurable (localhost default)
Audit Logging	Request logs in ~/.agentk/logs/

TROUBLESHOOTING

Dashboard Won't Load

```
# Check if service is running
ps aux | grep agentk

# Check port binding
netstat -tlnp | grep 5001

# Restart the service
pkill agentk && ./agentk --dashboard-port 5001
```

Chat Returns Errors

1. Verify model is available: `curl http://localhost:5001/api/models`
2. Check model provider connectivity (Ollama running? Cloud API key set?)
3. Review logs: `tail -f ~/.agentk/logs/agentk.log`

High Memory Usage

AgentK loads models into memory. To reduce footprint:

```
# Use lightweight model for simple queries
curl -X POST http://localhost:5001/api/chat \
```

```
-H "Content-Type: application/json" \  
-d '{"message": "Hi", "model": "phi4-mini"}'
```

ROADMAP (Q3 2026)

Feature	Description	ETA
Autonomous Agent Loop	Agents execute missions without manual triggering	July 2026
Mission Engine	Define goals, agents create + execute plans	August 2026
Swarm Coordination	Multiple agents collaborate on complex tasks	September 2026
Revenue Integration	Real payment processing (not simulation)	Q3 2026

Free Updates: All roadmap features included in your \$29 license at no extra cost.

SUPPORT

Issue	Contact	Response Time
Installation help	support@tikkunolam.today	24 hours
Bug report	bugs@tikkunolam.today	48 hours
Feature request	features@tikkunolam.today	1 week
Roadmap questions	roadmap@tikkunolam.today	1 week

LICENSE

Personal License v1.0

-  Use for personal projects

- Modify for personal use
- Receive free updates
- Commercial resale
- SaaS hosting for others
- Remove copyright notices

Upgrade Path: Commercial licenses available at tikkunolam.today/store

Full license: <https://tikkunolam.today/agentk/license>

UPDATE HISTORY

Version	Date	Changes
1.0 Beta	2026-06-15	Initial release (chat functional, autonomous Q3)

⌘ **All Is One. One Is All.**

Delivered by SoulTech — Sovereign AI Infrastructure