

onnai

Services As Software

Enterprise Solutions

onnai.ai

2025

echo "what you want" | onnai

The Problem

Current AI Workflows

- Browser-based, context-switching
- Copy/paste between tools
- Complex Python/ML stack required
- Configuration overhead
- Disconnected from existing workflows

Developer Time Lost

- 23 min avg. to regain focus after context switch
- Manual integration for every AI task
- Knowledge siloed in chat interfaces
- No composability with existing tools

English in. Unix power out.

- Pipe in, pipe out — standard Unix I/O
- No commands to memorize
- No configuration files
- No Python, no ML stack
- Just symlinks

```
curl -sSL onnai.ai/install | sh
```

How It Works

```
# Add specialized AI personas ("Mimics")
$ onnai --add Coder

# Use like any Unix tool
$ git diff | Coder "explain these changes"

# Chain with existing workflows
$ tail -f app.log | grep ERROR | onnai "diagnose"

# Clipboard workflow (macOS)
$ pbpaste | onnai "fix grammar" | pbcopy
```

Use Cases Across the Organization

Engineering

- Code review in git hooks
- Log analysis pipelines
- Documentation generation
- Context-aware development

Operations

- System diagnostics
- Incident analysis
- Automated reporting
- Process monitoring

Knowledge Work

- Data analysis
- Research summaries
- Content transformation
- Learning & training

onnnai CLI

- Resolves symlinks via argv[0]
- Walks directory tree for .context
- Unix socket to daemon (mode 0600)
- Auto-forks daemon on ECONNREFUSED

onnnaid Daemon

- Event-driven C, non-blocking
- kqueue (macOS) / epoll (Linux)
- Persistent HTTP/2 pool
- Double-fork daemonization

Zero latency. Secure credential isolation. No cold starts.

Bring Your Own Key (BYOK)

Your infrastructure. Your models. Your data.

- Set OPENAI_API_KEY, CEREBRAS_API_KEY, etc.
- Zero additional configuration
- onnai.ai never sees your API keys or traffic
- Support for local and cloud inference providers
- Private model endpoint integration

Enterprise Solutions

Custom Personas

- Proprietary knowledge
- Brand voice alignment
- Character design for users
- Persona marketplace

Deployment

- Full data control
- Audit capabilities
- On-premise Appliance
- Private model endpoints

Compliance

- SOC2
- HIPAA
- GDPR
- Custom requirements

Pricing

	BASIC	PRO	MAX
Monthly	\$9	\$45	\$119

30-Day Trial: \$10 for \$50 in credits

Enterprise pricing available upon request

API for Applications

```
POST https://onnai.ai/api/chat
Content-Type: application/json
Authorization: Bearer YOUR_API_KEY

{
  "Model": "Volcanara",
  "message": "What wisdom do you share?"
}
```

Perfect for: VR/AR, Video Games, Mobile Apps, Web Applications

Ready to Transform Your Workflow?

- Schedule a discovery session
- Custom pilot program
- Enterprise deployment planning

Contact: enterprise@onnai.ai

<https://onnai.ai>