

**The AI Talent  
You Can't Hire.**

**The Infrastructure  
You Can't Build.**

**The Time  
You Don't Have.**

Why Mission-Critical AI Needs More  
Than Software—It Needs Expert  
Deployment as a Service



**Neuron7.ai**

# AI That Solves Expert-Level Issues in Seconds

## The **AI Talent** You Can't Hire

- Expert AI engineering, deployed in your environment
- LLM orchestration & fine-tuning
- Custom model configurations optimized for your use cases, no PhD required
- Seamless enterprise integration
- Pre-built connectors for Salesforce, ServiceNow, Microsoft, SAP—deployed and maintained by experts
- Custom dashboards & configurations
- Tailored to your workflows, KPIs, and business logic—without draining your dev team
- Continuous optimization
- Ongoing model tuning and performance improvements—expertise that evolves with your needs

## The **Infrastructure** You Can't Build

- Enterprise-grade AI operations, without the capital investment
- Secure single-tenant deployment
- Dedicated infrastructure that meets your security and compliance requirements
- GPU, storage & compute—locked-in costs
- No surprise bills, capacity planning, or infrastructure headaches
- High-availability SLAs
- Mission-critical uptime guarantees backed by proactive monitoring
- Production-ready from day one
- Fully managed infrastructure that scales with your service operations

## The **Time** You Don't Have

- AI that solves expert-level issues in seconds
- Precision search at market-leading accuracy across regions, products, languages—from thousands of sources
- Turn-by-turn guidance (like GPS) for complex issues—no text-heavy back and forth
- Expert-level troubleshooting for anyone on your team—in seconds, not hours
- Instant analysis of service logs and telemetry data—no manual detective work
- Proactive optimization for parts, people, and processes—predict the next failure before it happens
- Real-time knowledge that learns from every interaction—getting smarter with every ticket

# Why AI-as-a-Service Wins in Complex Service Operations

When you service equipment that can’t afford downtime—industrial machinery, telecom systems, high-tech devices, medical equipment—you face a challenging choice:

### Build It Yourself

Cobble together OpenAI/Google/Anthropic APIs + hire AI specialists + buy GPUs + maintain infrastructure = 18+ month timeline, unpredictable costs, 67% failure rate, ongoing headcount burden

### Buy Generic Platform Agents

Salesforce Agentforce/ServiceNow/Microsoft = Great for CRM/ITSM ticket routing, but lack complex service domain expertise and require data migration to their cloud

### AI-as-a-Service

Purpose-built Service Resolution Intelligence + Full implementation + Managed infrastructure + Expert team = Production-ready in weeks, 90%+ accuracy, proven ROI

CHALLENGE	BUILDING IN-HOUSE	GENERIC PLATFORM AGENTS	NEURON7
Time to Production	12-18 months	4-8 weeks (limited scope)	6-12 weeks (full capability)
Accuracy for Complex Service	Must achieve through trial & error	60-80% (generic CRM/ITSM)	90%+ (service-specific)
Upfront Investment	\$500K-\$2M+	Platform fees + customization	Included in subscription
Team Required	AI specialists, prompt engineers, annotators, DevOps	Integration specialists + your IT	Included: service AI experts
GPU Costs	Fluctuating (+89%/year)	Included in platform (opaque)	Locked-in, predictable
Ongoing Maintenance	10-20% annually + \$5-20K/month	Platform updates (limited control)	Included + proactive optimization
Domain Expertise	Must build from scratch	Generic (CRM, sales, HR focus)	Purpose-built for complex service
Data Location	Your control	Vendor cloud (data migration)	Your secure single-tenant environment

# Everything Included. Nothing Left to Chance.

From Implementation to Ongoing Excellence. It's AI Deployment as a Complete Service

01

## Production-Ready AI Software

- Full Feedback Loop on what's working at scale
- CRM integration (SFDC/ServiceNow/Microsoft/SAP)
- Service-specific product roadmap
- Knowledge from data AND your people

02

## Cost Predictability

- GPU cost lock-in (no market volatility)
- No per-user or per-transaction licensing
- No reliance on internal IT to maintain
- All infrastructure costs included

03

## White Glove Implementation

- Data ingestion from cases, manuals, knowledge articles
- Configuration & testing
- CRM package install & support
- Training & change management

04

## Enterprise Grade Deployment

- Secure, scalable single-tenant deployment
- High-availability SLAs
- Proactive infrastructure monitoring
- SOC 2 Type II compliant

05

## 24x7 Support

- "Around the sun" support model
- Fast response & resolution times
- Software upgrades fully included
- Continuous monitoring

06

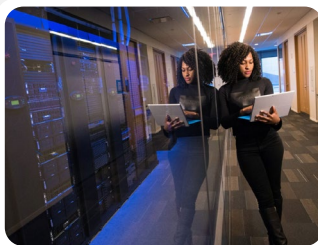
## Continuous Success

- Maintain 90%+ accuracy & minimize data drift
- LLM orchestration & fine-tuning per use case
- Custom dashboards & configurations
- Dedicated success team

## Key Industries



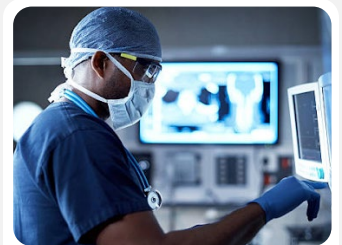
Industrial Equipment



High-Tech Device



Telecom Equipment



Medical Device

# Why Service Leaders Choose Neuron7

## Purpose-Built for Mission-Critical Service Environments Where...

- 3 hours of troubleshooting must **become 3 seconds**
- 86% of technicians “phone a friend” because **issues are too complex**
- Your service equipment spans multiple **regions and languages**
- Downtime costs **thousands per minute**
- Compliance and safety are **non-negotiable**

## Risk Mitigation

### Most (95%) enterprise AI pilots never make it to production

Tools designed specifically for service go live fast and deliver measurable results without long experimentation cycles

### No guesswork or trial-and-error deployments

Ready-to-use solutions integrate into existing workflows and start producing value immediately

### Proven outcomes for global service teams

Adopted across Fortune-1000 level service organizations with the backing of industry leaders and real world ROI

## Proven Results at Scale



**46% faster resolutions**

### Knowledge Capture

Reduced downtime and 14-point CSAT lift through AI-guided knowledge capture



**8,000+ technicians**

### Scale Best Practices

Consistent, high-quality service delivery across 60 regions



**24% lower parts cost**

### Reduce Parts Waste

Improve costs per work order with AI-driven parts & repair recommendations



**3 hours → 3 seconds**

### Faster Resolution

Complex diagnostics resolved instantly with explainable, real-time AI guidance

See how service leaders in your industry are **resolving complex issues faster with Neuron7**

