



SOLUTION BRIEF

OpenShift AI Readiness Review

Move from AI experimentation to enterprise-scale production—
faster, safer, and with confidence

01

Executive summary

Most organizations are stuck in AI proof-of-concept mode.

6-12 months

required for models to reach production

40-60%

of data scientists' time spent on infrastructure issues

3-6 month

delays introduced by security and compliance reviews

\$250K-\$500K+

in remediation costs due to poor platform decisions

Gruve's **OpenShift AI Readiness Review** enables organizations to confidently deploy Red Hat OpenShift AI by delivering a rapid, expert-led assessment and actionable roadmap in as little as 1-2 weeks. In addition to platform readiness, Gruve evaluates your preparedness to support scalable, governed Agentic AI workflows, including distributed inference, tool integration, and lifecycle governance.



OPENSIFT

02

The AI production gap

The problem

Organizations are investing heavily in AI — but few are realizing measurable production value. As organizations begin experimenting with agents when it comes to scaling out additional challenges emerge.

Common barriers include:

- Disconnected data science and platform teams
- Unclear GPU and storage requirements
- Inconsistent model lifecycle management
- Security and compliance uncertainty
- Overprovisioned or misaligned infrastructure
- Unstructured agent orchestration
- Uncontrolled tool access patterns
- Unclear ownership of agent lifecycle governance
- Limited visibility into inference performance and cost

The result?

Extended deployment cycles

Low executive confidence in AI ROI

Budget overruns

Delayed competitive advantage

Without a structured readiness framework—including assessment of emerging AI workflows—AI platform investments become high-risk initiatives.

03

Why this matters now

Competitive pressure is increasing

Organizations successfully deploying AI workflows are:

- Moving models and workflows to production in weeks, not months
- Improving operational efficiency
- Enhancing customer experiences
- Accelerating innovation cycles

The window of opportunity

The convergence of:

- Hybrid cloud adoption
- Kubernetes/OpenShift maturity
- Enterprise AI platform standardization
- Emerging AI regulatory frameworks
- Rapid adoption of LLM-based agent architectures

creates both opportunity and risk.

Organizations that prepare properly accelerate safely. Those that do not face costly rework, compliance gaps, stalled agent initiatives, and rising GPU costs. Early assessment of infrastructure readiness significantly reduces long-term platform risk.

04

The solution

Gruve's OpenShift AI Readiness Review

A structured, rapid engagement designed to ensure successful OpenShift AI deployment—including readiness for scalable AI workflows.

Duration: 1–2 weeks

Outcome: A prioritized, actionable roadmap for production AI and agent-enabled systems.

What we assess

1. Infrastructure readiness

- OpenShift/Kubernetes configuration
- GPU and accelerator capacity planning
- Storage architecture for ML workloads
- Networking and hybrid connectivity
- Scalability and performance requirements
- Distributed inference readiness

3. Data & integration architecture

- Data pipeline structure
- Feature engineering capabilities
- Data governance alignment
- Integration with enterprise data platforms
- Bottleneck and latency analysis
- RAG retrieval orchestration

5. Organizational readiness

- Data science and platform team alignment
- Skills assessment and training gaps
- Tool familiarity
- Change management readiness
- Cross-functional collaboration workflows

2. AIOps & ML lifecycle maturity

- Model versioning and registry capabilities
- CI/CD and automation maturity
- Deployment workflows
- Monitoring and observability
- Benchmarking against industry best practices

4. Security & compliance posture

- Identity and access management
- Secrets management
- Model security controls
- Alignment with HIPAA, GDPR, SOC 2, and industry regulations
- Secure tool invocation controls
- Access boundaries for agent workflows
- Audit traceability across multi-step reasoning processes



Deliverable: Your implementation roadmap

At the conclusion of the engagement, clients receive:

- Prioritized remediation plan
- Effort and timeline estimates
- Reference architecture aligned to their environment
- ROI projections and cost optimization analysis
- Structured evolution toward enterprise-grade AI workflows
- Phased implementation strategy

05

Measurable business impact



Accelerated time to production

- Reduce AI deployment timelines by 40–60%
- Move from months-long cycles to production in weeks
- Establish structured governance for agent-enabled systems



De-risked platform investment

- Avoid \$250K–\$500K+ in remediation costs by validating infrastructure and workflow readiness before implementation.



Optimized infrastructure spend

- Reduce infrastructure waste by 25–35% through right-sized GPU, inference, and storage planning aligned to both ML and Agentic AI workload patterns.



Enhanced security & compliance

- Prevent 3–6 month delays caused by failed security reviews by proactively designing governance for both models and agents.



Increased team productivity

- Enable data scientists and AI engineers to focus on innovation — **delivering 50% faster model and workflow iteration cycles with defined lifecycle controls.**

06

Service options

Foundation Assessment

5–7 days | 10–15 customer hours

Best for:

- Early-stage AI initiatives
- Rapid executive evaluation
- High-level roadmap development

Includes:

- High-level infrastructure review
- Current-state maturity benchmark
- Gap identification
- Executive-ready findings summary

Comprehensive Readiness Review

10–12 days | 20–30 customer hours

Best for:

- Enterprise-scale environments
- Regulated industries
- Multi-cloud or hybrid architectures
- Strategic AI investments

Includes:

- In-depth infrastructure analysis
- Detailed AIOps capability mapping
- Security and compliance assessment
- 2–3 collaborative workshops
- Detailed reference architecture
- ROI projections
- Skills gap analysis and training plan
- Change management strategy



Before & after: The transformation

Before

- Isolated notebook-based experimentation
- No standardized platform
- 3–6 month deployment cycles
- GPU overprovisioning
- Reactive security reviews
- Unstructured agent pilots
- Low executive confidence in AI ROI

After

- Production-ready OpenShift AI architecture
- Clear GPU and infrastructure strategy
- Accelerated deployment timelines
- Embedded compliance controls
- Phased AI and Agentic AI maturity roadmap
- Executive-aligned ROI and investment clarity

07

Take the next step

AI operationalization should not be a gamble.

Most assessments begin within two weeks of engagement.

In as little as **5-12 days**, your organization can gain:

- Infrastructure clarity
- Security confidence
- Financial justification
- A defined roadmap to production AI

Schedule your OpenShift AI Readiness Review

 **Website:** <https://www.gruve.ai/>

 **Email:** info@gruve.ai