



AI FAQ

→ What AI technology does Nodes & Links actually use?

We deploy two distinct AI technologies with fundamentally different purposes: Machine Learning (ML) for risk prediction and Generative AI (GenAI) for report automation. The ML component uses specialized algorithms originally developed for mission-critical applications like epidemic control, where both accuracy and interpretability are non-negotiable requirements. GenAI handles report generation, eliminating the need to start documentation from scratch. Our GenAI leverages our proprietary AI models and builds on top to ensure reporting reflects the unique nature of your project and the respective analytics.

→ How is your ML approach different from typical "black box" AI?

We use interpretable ML specifically designed for high-stakes decision-making environments. Unlike opaque neural networks, our approach provides transparent reasoning that project professionals can validate against their expertise. This isn't experimental technology. It's peer-reviewed, published in top-tier academic journals with MIT co-authors, and has been validated across hundreds of live projects.

→ What evidence exists for your accuracy claims?

Independent testing on live projects demonstrated 7x improvement in accuracy over traditional forecasting methods and 2x greater robustness against optimism bias. In milestone forecasting, our AI outperformed human predictions in 15 out of 17 cases. These aren't lab results. They're from actual construction projects deployed right now. You can read the case study [here](#).

→ How does the AI learn from my specific project?

For every project you create in Nodes & Links, we use that data to automatically create your own tailored prediction model. The AI requires approximately 10% of activities within the project to be completed before making reliable predictions. It classifies similar activities based on context (WBS, resources), time (float, timestamps), and connectivity (paths, relationships) and additional parameters, then applies weighted performance rates from completed activities to forecast future work. The precise components of this classification form the basis of our peer-reviewed innovation. The model subsequently retrains automatically with every schedule update.

In cases where there are not enough completed activities within a project's schedule, our AI prediction engine will fall back to a more general recommendation mode, which relies on a combination of our published findings to propose recommendations (you can find our published research [here](#)).

→ Does my data train models for other customers?

Absolutely not. Our multi-tenant architecture ensures complete logical separation between customer tenants. Project A's data trains models exclusively for Project A within your isolated workspace. No customer data is shared with other Nodes & Links customers, nor is it used to benefit anyone else's models.



→ **What happens to my data with your GenAI features?**

Your data never leaves Nodes & Links cloud infrastructure and is never shared with third-party providers. We run our own AI infrastructure under AWS's shared responsibility model specifically to maintain this boundary.

We may sometimes combine our AI with third-party foundational models (like OpenAI's GPT5, AWS' Nova and Anthropic's Claude 4.5) which we have fine-tuned and distilled with best practices, from industry leading boards incl. PMI and AACEI. All processing occurs within your encrypted, logically isolated tenant using only project schedule data, no PII and personal data is involved in our AI.

→ **How frequently do you update the AI models?**

The AI retrains with every project schedule you upload. More frequent updates yield more frequent and accurate predictions because the model continuously refines its understanding of your project's unique performance patterns. There's no manual intervention required. It happens automatically as part of your normal workflow.

→ **What technical controls protect my data?**

All data is encrypted at rest and in transit, with Web Application Firewall (WAF) protection on all endpoints. We maintain ISO 27001 certification covering all AI capabilities, evidencing both technical and organizational controls. Annual third-party penetration tests specifically assess AI guardrails and security measures, and we can provide those reports.

→ **How do you validate ongoing model performance?**

We use a backtesting methodology with crossfold validation: the AI makes predictions on historical data, then we compare those predictions against what actually happened. This constant feedback loop ensures models remain refined and tailored to each project's evolving reality. It's the same validation approach used in high-stake cases like drug testing and finance.

→ **Is your AI security cleared?**

Yes. Our ISO 27001 scope includes all AI capability, and our annual penetration tests explicitly assess AI-specific concerns, from code injections to data safeguards. Importantly, we process only project schedule data, no demographic or personnel information that could introduce discriminatory and related ethical outcomes.

→ **What control do I retain over AI recommendations?**

Complete control. You choose whether to engage, adopt, or discard all AI features and outcomes. AI recommendations are exactly that: recommendations you can override based on your professional judgment and site-specific knowledge. The AI augments your expertise; it doesn't replace it.



→ **What's your uptime guarantee?**

Our SLA commits to 99.5% availability measured over rolling 30-day periods, covering both AI and non-AI features equally. Current uptime is tracked live [here](#), typically reporting 100% uptime over the last 90 days.

→ **How is data residency handled?**

We support geographic-specific data storage in US, EU, UK, and Australia regions. This ensures compliance with regional data sovereignty requirements.

→ **What happens to my data if I terminate the service?**

Voluntarily closed accounts enter a 90-day retention period during which you can download your data. After 90 days, data is permanently removed. For involuntarily suspended accounts (payment issues), there's a 30-day grace period to restore access before the account closes and enters the standard 90-day deletion cycle.