Speed up your AI
FIXSTARS

Fixstars Corporation

Mid-term Management Vision

October, 2026 – September, 2029

From Service Company to

AI-Native Service + AI Inference + AI Agents Company

Leveraging AI agents and automated optimization AI to scale flow revenue independent of engineer headcount, while driving recurring revenue from inference and agent platforms.

November, 2026

Speed up your AI
FIXSTARS

01

About Fixstars

Where We Stand Today:

Revenue expanded from ¥7.0 billion to ¥10.6 billion over the past three years, successfully building a solid foundation for sustainable growth.
Challenges: transforming revenue structure.

Our Software Skills to Maximize Hardware Performance

Speed up your AI — We are experts in AI and software acceleration

Software Acceleration Skills

- Solve clients' business challenges by accelerating computing performance.
- Achieve equivalent processing power on hardware with significantly lower power consumption.

Hardware Exploitation Expertise

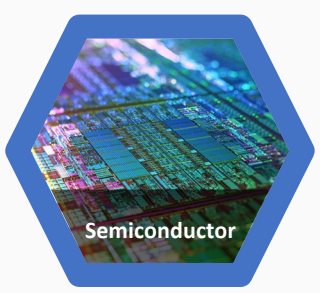
- Maximize the performance of emerging hardware through software optimization.
- Build and refine SDKs and libraries to enhance hardware competitiveness.

Target Hardware: Our Track Record

- GPU
- CPU
- FPGA
- AI inference accelerator
- Edge SoC
- Quantum computer / Ising machine

etc.

Our Main Industries











Official Corporate Purpose

The advent of the AI inference era amplifies the significance of our corporate purpose.

Through software optimization



Making infeasible computations possible



Lowering energy consumption for computing workloads

Surging AI adoption continues to drive up inference compute workloads and power consumption—making efficient execution technologies even more critical over the next three years.

The greater the demand for inference, the more our efficient execution technology translates directly into client cost savings and revenue opportunities for Fixstars.

Review of the Previous Vision

From “Accelerate software” to “Operate accelerated software”—the foundation for growth is established, with the revenue model transformation ahead.

Achievement

Business Growth

Net sales scaled from ¥7.0B (FY09/23) to ¥10.6B (FY09/26), while maintaining high profitability.

Launch SaaS Business

Expanded enterprise and academic adoption of Fixstars Amplify, our quantum and mathematical optimization cloud.

Rollout & Adoption of AI Infrastructure Products

Launched Fixstars Vega, an AI Agent-based development environment. Building our software optimization AI, “APEX” by using Vega internally.

Challenges

Revenue remains predominantly flow-based

Over 90% of revenue remains contract-based, with recurring product/cloud streams still in the investment-for-growth phase.

Inbound-driven sales structure

Sales remain predominantly driven by inbound client inquiries, leaving proactive market-creation capabilities underdeveloped.

Shortage of executive management talent

The recruitment and development of management talent have not kept pace with our rapid business growth.

The new Medium-Term Vision focuses on addressing these three remaining challenges.

Speed up your AI
FIXSTARS

02

Market Trends

Our Business Environment

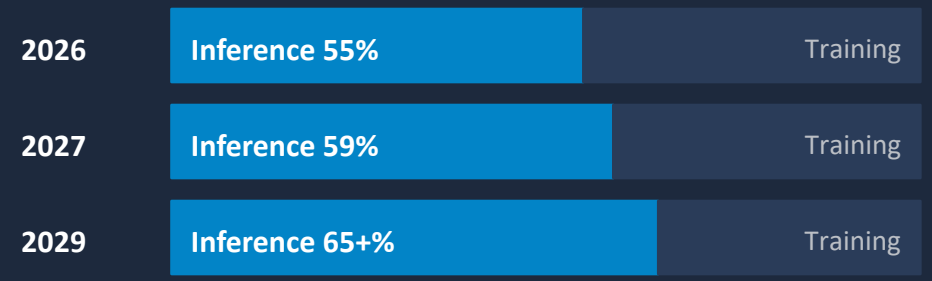
The Era of AI Inference
As AI models and hardware diversify, *the value of optimization rises*.

AI: From Development to Deployment

Computing demand continues to transition from model creation, “training” to model deployment and execution, “inference”.

In 2026, spending on AI inference overtakes training.

Breakdown of Global Cloud Spending for AI (AI-Optimized IaaS)



*All figures are projections.

In Japan,

AI inference server spending projected to overtake training by 2027

Source: Gartner (October, 2025 / August, 2026), IDC Japan (March, 2026)

Surging Inference Workloads Driven by AI Agents

From human-prompted AI to autonomous AI: Exponentially increasing inference required per single task.

Cost and Power Constraints on AI Adoption

Inference costs and power availability dictate the scale of enterprise AI adoption.

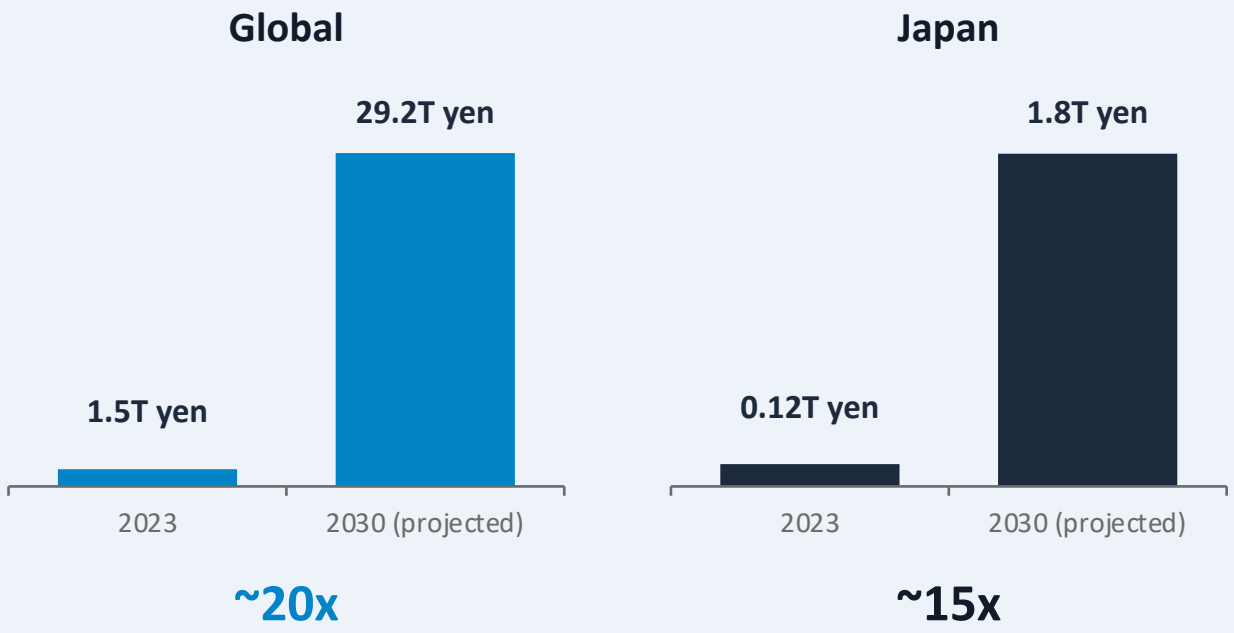
Same hardware, performance dictated by software

On identical GPUs, throughput and latency vary drastically depending on implementation.

Rise of Open Models: Expanding Inference and Creating Massive Market Potential

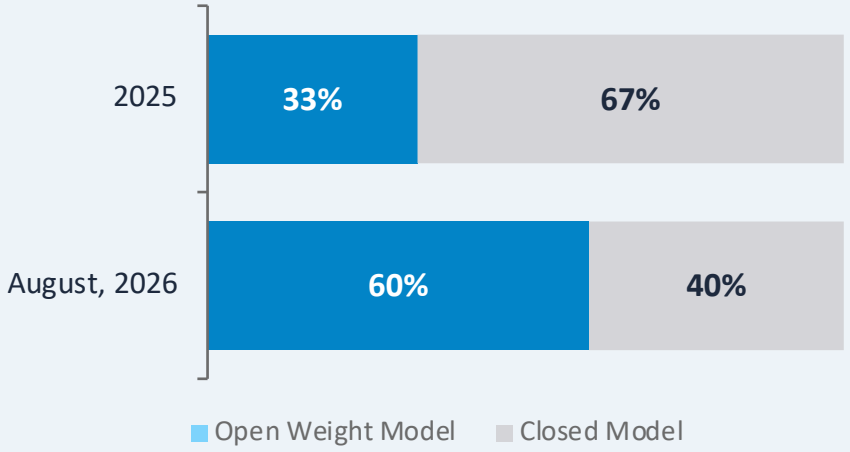
Global GenAI demand is projected to reach ~¥29T and Japan ~¥1.8T by 2030, with open models driving surging inference workloads.

① Surging Market Demand for Generative AI



② The Rise of Open Models in AI Inference

Breakdown of Tokens Consumed in Inference
(Based on OpenRouter Usage Data)



Open models already account for the majority in the U.S.

A Massive Market Dominated by Open Models.
Fixstars Aims to Provide the World's Leading Low-Cost Inference Infrastructure.

Source: ① JEITA (December, 2023. 1USD = 138.6 JPY) ② OpenRouter・a16z (December, 2025) , The New Stack (September, 2026 / Request originated from the US)

Three Structural Shifts Expand Our Growth Opportunities

The optimal configuration is constantly changing

Evolution and Diversification of Hardware

Rapid GPU generation cycles and the rise of dedicated inference chips and edge SoCs ensure that the optimal hardware continuously evolves.

The Rise of Open Models

A flood of high-performance open models makes optimal selection increasingly complex.

Surging Demand for In-House AI Operations

More enterprises demand internal control over their proprietary data and custom AI agents, driving surging demand for dedicated, private inference environments.

Which Model, on Which Hardware, and How to Execute

Lowering AI Costs and Power Consumption Through Our Optimization

Speed up your AI
FIXSTARS

03

Mid-term Management Vision

October, 2026 – September, 2029

Transitioning Custom Engineering to Headcount-independent Scalable Growth
Stacking High-Margin Recurring Revenues from Inference and Autonomous Agents

Speed up your AI
FIXSTARS

OUR VISION 2029

Towards

AI-Native Service + AI Inference + AI Agents Company: Leveraging Optimal Silicon of the Era to Provide Cost-Effective, Low-Power AI Inference and Leading Software Engineering via AI Agents



Low Cost



Low Latency



Low Power



AI Agents

Optimal Hardware for Each Era. Best-in-Class Models for Each Moment. AI Agents Driving Development on Top.

Core Strategic Principles

Scaling recurring revenues in AI inference and AI agents through an AI-native foundation

AI Agents (Agentic Development)

Fixstars Vega AI Agent Environment for Software Development

AI Agents Refined in Our Own Engineering Trenches: Unique Value Beyond Standard Inference APIs

Recurring Revenue based on AI agent solutions

AI Inference

Fixstars Inference Cloud (shared) to Private Inference Environment

Shared Inference Cloud › **Dedicated Cloud** › **Private On-premise**

Recurring revenue based on AI inference environments

AI-Native Service (Foundation)

AI-Native Engineering Service

APEX
Proprietary Internal AI for Automated Porting and Optimization

An optimization service that progressively leverages our internal AI platform, APEX, to deliver equivalent outcomes in shorter timeframes.

High-margin project-based revenue

Strategic Countermeasures to Challenges from the Previous Medium-Term Vision

Overcoming the three remaining challenges: a strategic transformation under this vision

Revenue remains predominantly flow-based	Establishing recurring revenue streams in inference and autonomous AI agents while elevating transactional revenue through AI.
Inbound-driven sales structure	Proactively unlocking latent demand through dedicated teams while proving AI inference superiority.
Shortage of executive management talent	Expanding management talent via tech-leader development, business hiring, and rapid M&A integration.

Our Competitive Moat

Four core competencies forged through proven technical track records and client trust

Hardware-Aware Optimization Technology

Over 20 years of acceleration and optimization for GPUs, FPGAs, and ASICs: rapidly unlocking maximum performance from Next-Gen hardware.

Automated Porting and Optimization AI: APEX

Applying real-world porting expertise to efficiently scale hardware compatibility for our automated AI engine and systematically funneling frontline insights into APEX's continuous evolution.

Proprietary AI Agent: Fixstars Vega

Daily in-house use and rapid iteration in client projects: enhancing product competitiveness through continuous client delivery.

Long-Standing Partnerships with Security-Critical Enterprises

Identifying needs for dedicated AI environments through decades of engagements with security-critical clients in semiconductors and automotive industries.

Synergizing Four Core Strengths to Build an Enduring Competitive Moat

AI-Native Engineering Service — Engineer x APEX Upgrading Engineering Workflows

Next-generation engineering services: fusing elite engineering talent with in-house AI "APEX" to accelerate development

Four Strategic Domains for AI Application

Software Acceleration and AI Inference Optimization

Our Core Competency: Accelerating the Discovery and Exploration of Optimization Techniques via Proprietary AI

Embedded, Automotive, and Edge AI Implementation

Ultra-Low-Power, Highly Efficient AI Inference on the Edge

Cybersecurity

End-to-End Vulnerability Assessment and Remediation

Mathematical Optimization

Leveraging AI to Streamline the Full Lifecycle: From Formulation to Execution and Operations

Transitioning from Hourly Billing to Value-Based Pricing Tied to Performance Gains and Accelerated Delivery

Engineer x APEX

Engineer's Roles

- Defining scope
- Architecting
- Client engagement

APEX's Roles

- Porting
- Finding optimization methods
- Validation and testing

Channeling Project Insights Back into APEX to Continuously Advance Our Optimization Technologies

Project Execution › **Feedback into APEX** › **Leveraging for future projects**

Transforming individual know-how into permanent, scalable organizational technical assets.

APEX — AI for Automated Porting and Optimization Across Diverse Hardware

A proprietary AI for porting and optimization: retained in-house as our core competitive moat

Inputs

AI Model

LLM etc.

Software

Existing programs

APEX

Internal-Use Only

Proprietary AI absorbing cross-silicon differences.

① Automated Porting

Deep architectural analysis enabling tailored model and software transformation for any silicon

② Automated Optimization

Automating chip-specific parameter exploration through repeated empirical on-device measurements.

Measure › Explore › Apply › Train

CLOUD

- GPU
- AI Inference Chip
- Accelerator

→ Strengthening competitiveness of Fixstars Inference Cloud.

EDGE

- Automotive SoC
- Embedded NPU
- FPGA

→ Streamlining embedded, automotive, and edge AI services.

Chip diversification multiplies porting and tuning workload.
AI automation and knowledge retention build structural moats.

Fixstars Inference Cloud — Pay-As-You-Go AI Inference Cloud

Scheduled for FY2027 release: world-class optimized open-model inference cloud—zero hardware procurement, zero model building

Users

- **Software Developers**
- **Enterprise Application and Services**
- **AI Agents**

Fixstars Inference Cloud

State-of-the-Art Open-Weight Model

- **Chat**
- **Code**
- **Image Recognition**
- **Inference**

APEX Automatic optimization of Model x Hardware

- **GPU**
- **AI Inference Chip**
- **Next-Gen Chip**

Values

Accelerated Inference

Low-Latency Inference Across Generative Text, Coding, and Computer Vision Workloads

Pay-As-You-Go

Zero-Upfront Investment with Per-Token Pay-As-You-Go Pricing

Low Cost

Sharing APEX Efficiency Gains Through Disruptive, Value-Aligned Pricing

Low Latency

High-Frequency Resilience: Ultra-Low Latency and High-Throughput Inference

Low Power

Maximizing Inference Throughput per Watt

AI Inference — Fixstars Inference Cloud to Your Private Inference Environment

From proven public cloud efficiency to high-impact dedicated AI environments

Public Service (Hands-On User Experience)

Enterprise (Own your Intelligence: Securing Proprietary Data and Dedicated AI Environments)

Shared **Fixstars Inference Cloud**

Target Customers

Developers/Enterprise Customers (Global)

- Instantly Accessible Shared Environments
- Experiencing Superior Cost-Performance and Throughput with Live Production Traffic
- Zero Initial CapEx: Effortless Testing and Validation of Leading AI Models

Dedicated **Private Cloud Environment**

Target Customers

Enterprises pursuing full-scale AI implementation

- Dedicated, Single-Tenant Cloud Inference Environments
- Workload Isolation in Single-Tenant Dedicated Environments
- Predictable Fixed-Fee Pricing to Lock In Operating Costs

Dedicated **On-premise Environment**

Target Customers

Enterprises enclosing AI as proprietary assets

- Fully Private On-Premises Inference Environments Built Within Client Facilities
- Fully Self-Contained In-House Data and AI Operations
- Sustained Long-Term Operation of Proprietary AI Assets

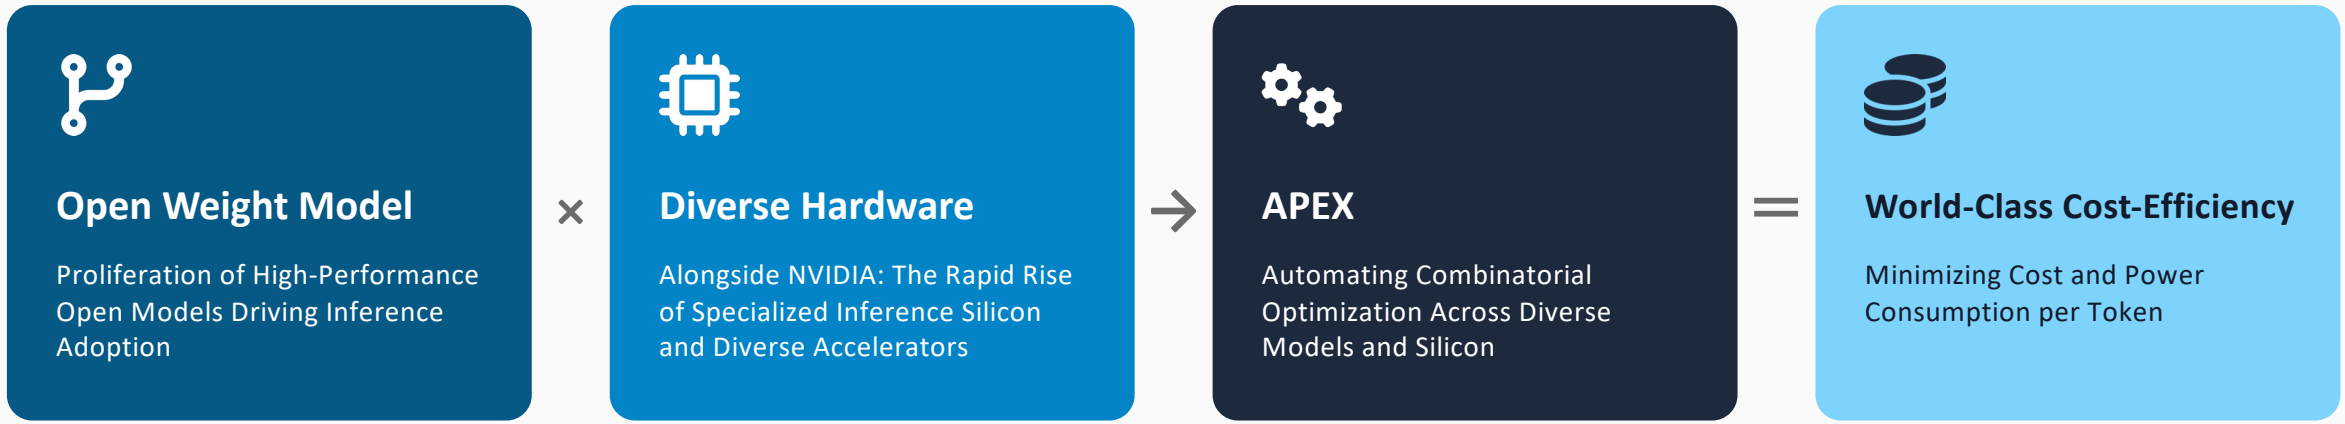
Lowering initial AI deployment barriers

Seamless migration with identical models and operations

Long-term recurring revenue transition

Evolving into the World's Most Cost-Efficient AI Inference Provider

Optimizing combinations across proliferating models and hardware: the foundation of inference cost competitiveness



Key Drivers Enabling World-Class Low Costs

1 Hardware Selection Agility

Silicon-agnostic strategy utilizing the best price-performance hardware of each era.

2 Maximizing Hardware Performance

Maximizing Hardware Throughput via Decades of Low-Level Engineering Mastery

3 Automating the End-to-End Optimization Process

A Widening Competitive Moat: Turning APEX Automated Optimization into Unrivaled Cost Reductions and Rapid Support for Evolving Models

Human Limits: The Unsustainable Bottleneck of Manual Model-Hardware Tuning.
A Widening Competitive Edge Driven by Proprietary Auto-Optimization.

Agentic Development — AI Agent Development Platform: "Fixstars Vega"

High-margin value creation: bundling in-house trained AI agents with low-cost inference as a strategic gateway

Values of AI Agents

Daily In-House Use and Field-Testing Across Client Projects

Turning Daily Engineering Practice into an AI Knowledge Asset

Delivery

Commercial Launch. Phased Rollout Across Dedicated and Shared Infrastructure

Why We Integrate AI Agents atop Our Inference Infrastructure

Inference Infrastructure Only

- Exposure to Commoditized Token Price Wars
- Minimal Switching Friction

Inference Infrastructure + AI Agent

- Chosen for Measurable Productivity and Business Outcomes
- Continuous Accumulation of Client-Specific Know-How within the Agent
- Decoupling Intelligence from Base Models: Clients Retain Proprietary Workflows and Domain Knowledge Across Changing Model Lifecycles

Provided AI Agents

Development Agents

End-to-End Engineering Execution: From Massive Codebase Comprehension to Testing, Infused with Embedded and Low-Level Mastery

Cybersecurity Agents

End-to-End Vulnerability Lifecycle: Seamless Execution from Diagnosis to Remediation and Regression Testing

Mathematical Optimization Agent

Accelerating Formulation from Real-World Challenges to Mathematical Models: Lowering the Adoption Barrier for Fixstars Amplify

Strategic Synergy with Contract Software Engineering Services

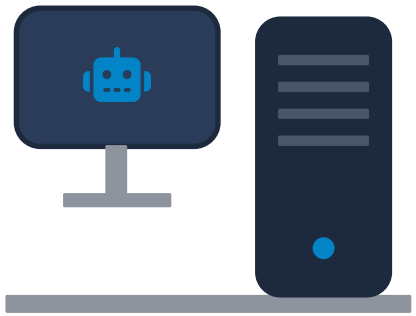
Optimal Workforce Balance: Autonomous Agents for Routine Tasks, Elite Engineers for High-Complexity Challenges

Fixstars Vega — A Dedicated AI Engineering Team by Your Side

Frictionless scalability: from local workstations to GPU clusters, from small-scale pilots to enterprise infrastructure

Fixstars Vega Universal Agent Architecture: Seamless Consistency Across Cloud, On-Premises, and Hybrid

- Development Agent
- Cybersecurity Agent
- Mathematical Optimization Agent



Workstation Type

By Your Side

- Deployable in standard office environments
- Adoption by small teams and workgroups
- Enabling bottom-up, agile pilots

Small-Scale



GPU Cluster Type

Large-Scale GPUs

- Centralized management across on-premise datacenters and dedicated cloud infrastructure
- Enterprise-wide concurrent access and resource sharing
- Parallel multi-agent processing

Enterprise-Wide

Own your Intelligence **Full-Stack Governance: Retaining Absolute Control from Compute to Proprietary Data Assets**

AI-Native Service + AI Inference + AI Agents: Transforming Our Revenue Structure

Transitioning from headcount-proportional growth to high-productivity services and recurring inference & agent revenues

To Date

Revenue =
# of Engineers
x Unit Price

Headcount-Driven Linear Revenue Growth

From Now **Revenue =**

AI-Native Service

High-Value Engineering Services Powered by Phased APEX Integration

Transactional | Optimization Service Fees

+

AI Inference

Providing Shared (Fixstars Inference Cloud) and Dedicated AI Inference Infrastructure

Recurring | Inference Fees

+

AI Agents

Providing "Fixstars Vega" AI Agent Environments for Enterprise Development

Recurring | AI Agent Fees

Structural Business Pivot: From Labor-Intensive Consulting to Recurring Software & Inference Foundations

Three Important KPIs

Track AI-native service + AI inference + AI agents progress by following three KPIs

① Revenue per Engineer

Engineering Services & Custom Development

What It Stands For

Value creation via AI-native services

Definition

Revenue per engineer driven by APEX and autonomous AI agents

② Annual Revenue of AI Business

Fixstars Inference Cloud・Fixstars Vega・Fixstars Amplify

What It Stands For

Progress in shifting to a recurring business model

Definition

Total annualized recurring revenue from inference, agents, and quantum.

③ Tokens for AI Inference Service

AI Inference Volume Managed and Operated by Fixstars

What It Stands For

Proven operational scale and compute processing track record

Definition

Aggregate tokens processed on our inference platform across all workloads.

Tracking Structural Revenue Transformation via These Three KPIs.

Four Strategic Initiatives to Accelerate Growth

Creating market demand through strategic alliances and proactive expansion

Branding & Marketing

Demonstrating technical superiority through objective track records on external inference platforms, while proactively cultivating enterprise demand with dedicated account teams

Strategic Partner Alliances

Strengthening strategic partnerships with chipmakers and cloud platforms, accelerating early silicon access and deep joint software-hardware optimization

M&A and Capital Alliances

Driving M&A and alliances with engineering organizations to boost productivity via APEX and Fixstars Vega, elevating value added and profitability to scale our business

Global Market Expansion

Expanding into North America's massive AI inference market backed by our proven domestic track record, unlocking enterprise demand through collaboration with local cloud and semiconductor leaders

Capital Allocation Policy **Phasing inference infrastructure investments in line with demand trends, while prioritizing abundant cash reserves for strategic M&A and human capital**

Culture-led Scaling and Cultivating Technology Management Talent

Anchoring on our culture to appoint key business and executive leaders from within and outside

Cultivating Leaders in Technology Management

Provide top engineers with opportunities to develop and apply management skills, establishing a clear pathway from engineering roles to executive leadership.

Strengthening Strategic Human Resource Functions

Enhancing talent acquisition, training, and evaluation frameworks to attract specialized leaders driving business development, proactive sales, and M&A in addition to engineers

Company Culture

Freedom and Discipline

Think independently and take action

Fostering an Inclusive Environment for a Diverse Workforce

Promoting a bilingual work environment to establish an onboarding infrastructure where exceptional talent can thrive regardless of nationality

Driving post-merger organizational integration

Accelerating early integration by deploying our AI environment to target companies and enabling rapid onboarding to scale the organization

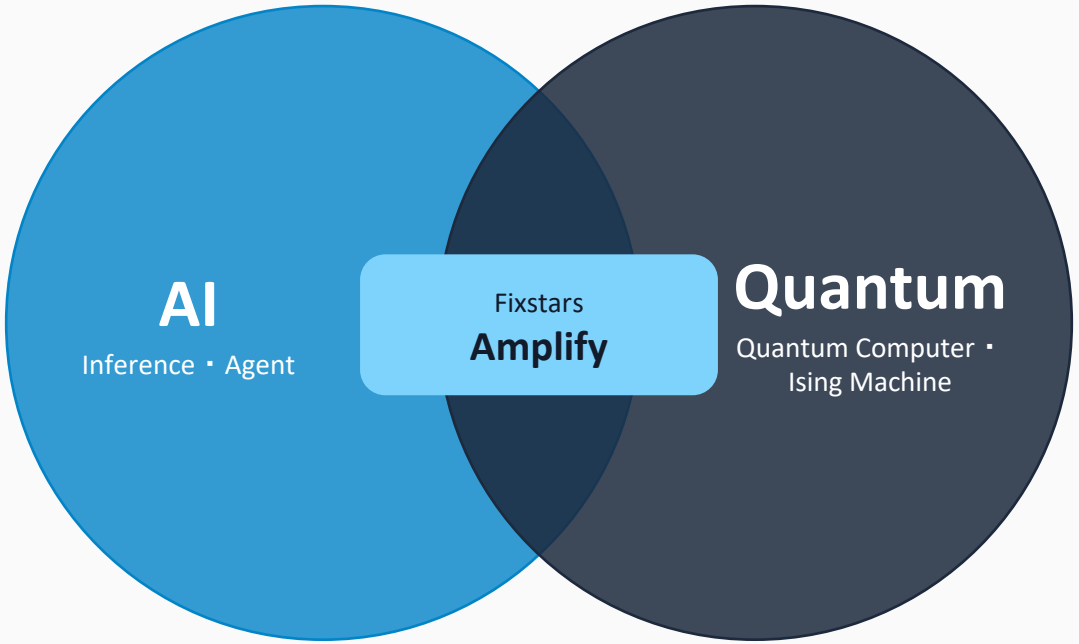
Three-year Roadmap

Building recurring revenues via inference infrastructure and AI agents over three years, powered by AI-native services

	FY2027 Building Foundations & Validating Value 2026/10–2027/09	**FY2028 Expanding Offerings and Accelerating Deployment** 2027/10–2028/09	**FY2029 Establishing a Core Pillar of Recurring Revenue** 2028/10–2029/09
Operation / Organization	• Leveraging AI across all projects and feeding insights back to APEX • Establishing dedicated client teams and launching our M&A execution team	• Scaling value creation per engineer into full gear • Driving M&A and post-merger integration of engineering organizations	• Driving full-scale expansion into the North American market backed by our domestic track record • Securing a leading position at the intersection of quantum computing and AI
AI Inference	• Launching Fixstars Inference Cloud (shared) and proving objective technological superiority • Build a track record of deploying dedicated AI inference environments	• Accelerating upselling from shared to dedicated environments • Expanding non-NVIDIA chip optimization	• Establishing a stable recurring revenue pillar through shared pay-as-you-go and dedicated subscription fees
Agentic Development	• Demonstrating productivity gains through full deployment across internal development environments • Launching bundled offerings of Fixstars Vega and dedicated inference environments	• Broadening our domain coverage to include cybersecurity and mathematical optimization • Launching Fixstars Vega on our shared cloud (Fixstars Inference Cloud) to expand our user base	• Expanding application scope across diverse industry sectors and use cases • Accelerating stable recurring revenue from agent platform fees in earnest

The Convergence of Quantum and AI: Becoming the Defining Company of the Era

Rapidly evolving AI and quantum moving toward practical application: driving next-generation computing platforms through deep-level hardware mastery



From AI Inference to Quantum Computing: A Unified Platform Integrating Optimal Hardware for Practical Deployment

An Established Business Foundation: Fixstars Amplify

- Commercializing a quantum and mathematical optimization cloud connected to leading quantum computers, Ising machines, and GPU solvers.
- Building proven enterprise track records and a solid customer base across industry and academia.

Automating Mathematical Modeling and Formulation Through AI

- Overcoming the biggest barrier to adoption: AI agents assist in formulating complex business problems
- Enabling seamless ingestion of operational optimization challenges into Fixstars Amplify

Hardware-Agnostic Scalability

- A common computing platform independent of specific hardware
- Ensuring seamless migration of existing assets to next-generation hardware as commercial quantum computing matures

VISION 2030

2030: AI Agents Operating Across Every Corner of Society, Driving an Explosive Surge in Inference Compute

Running AI Inference on the Era's Best Hardware and Driving Societal Software Development Alongside AI Agents

Enterprise Software Development, Security, and Optimization Client Front Lines

Fixstars Vega Agentic Development — Developer AI Agent Environments Powered on Inference Infrastructure

Fixstars Inference Cloud AI Inference — Shared ／ Dedicated Own your Intelligence

Cutting-edge GPU | Inference Accelerator | Edge SoC | Quantum Computer

The Era’s Best Hardware | Performance Underpinned by APEX and Optimization Service Expertise

*The future plans, goals, and other forward-looking statements described in this document are based on assumptions as of the date of this document and are subject to change due to future changes in the business environment and other factors.