

Case Study: Cost Reduction and Performance Increase

From **\$87.5K** to **\$7K** Per Month Without
Compromising Performance

Table Of Content

- Executive Summary
- About SkillGigs
- The AWS Infrastructure Challenge
- The Cloud Logically Partnership
- Comprehensive Assessment & Discovery
- Optimization Strategy & Implementation
- Remarkable Results
- Business Impact and Strategic Value
- Partnership Excellence
- Sustained Success and Future Roadmap
- Client Testimonial
- Key Takeaways
- Ready to Transform Your Cloud Costs?

How We Saved **\$966K** Annually

How Cloud Logically Reduced SkillGigs' AWS Spend By 92%

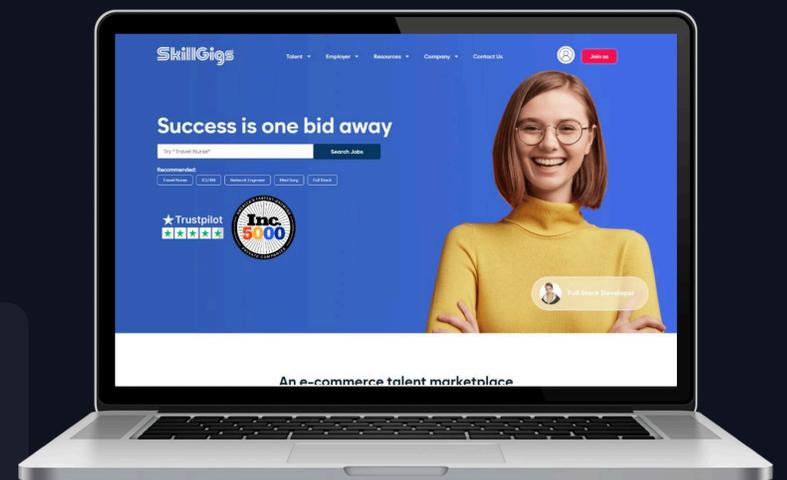
From **\$87.5K** to **\$7K** Per Month Without Compromising Performance

Executive Summary

- **The Challenge:** SkillGigs, a fast-growing talent marketplace serving 250,000+ users, faced escalating AWS costs that threatened profitability and investor confidence.
- **The Solution:** Cloud Logically's comprehensive cost optimization strategy delivered immediate and sustainable savings while improving platform performance.

Result

92% reduction in AWS spend (**\$87.5K → \$7K monthly**), 3-month ROI payback, and enhanced system performance that enabled strategic reinvestment in growth initiatives.



About SkillGigs

SkillGigs is a revolutionary online talent marketplace that leverages patented 3D resume AI technology to connect experienced professionals with employers. Having achieved recognition on the Inc. 5000 list, the platform serves over 250,000 active users and has become a critical bridge between talent and opportunity across multiple industries.

Key Business Metrics

- 250,000+ active users
- Inc. 5000 fastest-growing company
- Expanding from IT staffing to healthcare talent markets
- Revenue model based on marketplace commissions



Business-Critical AWS Dependencies

SkillGigs relied heavily on AWS for:

AI/ML Infrastructure: Their patented 3D resume technology required extensive computing power for resume analysis, skill extraction, and talent-job matching through AWS SageMaker and GPU instances.

Scalability: Quarter-million active users created fluctuating traffic patterns, especially during peak hiring seasons, requiring robust auto-scaling capabilities.

Data Management: Massive volumes of resumes, job postings, and user interactions needed secure, scalable storage using RDS, DynamoDB, and S3.

Marketplace Operations: Real-time bidding, notifications, and search functionality demanded high-availability infrastructure with minimal downtime tolerance.

The AWS Infrastructure Challenge

Rapid Growth, Escalating Costs

As SkillGigs scaled their innovative platform, their AWS infrastructure costs grew at an unsustainable rate:

- **Monthly AWS Spend:** \$87.5k+ and climbing
- **Growth Pattern:** Infrastructure costs increased 400% while user growth only tripled
- **Cost per User:** Rising monthly, threatening unit economics
- **Resource Utilization:** Massive over-provisioning with 70-80% idle capacity

The Tipping Point

By 2022, several critical issues converged:

- **Financial Risk:** Operating costs were decimating profitability and cash runway
- **Investor Concerns:** Rising infrastructure costs raised questions about scalability and unit economics
- **Engineering Drain:** 30% of development time spent on infrastructure firefighting rather than product innovation
- **Performance Issues:** Over-provisioned resources paradoxically created performance bottlenecks

The Cloud Logically Partnership

Discovery and Selection

SkillGigs discovered Cloud Logically through the AWS Partner Network, recommended by their AWS Technical Account Manager due to escalating costs and rapid user growth.

Why Cloud Logically

- Marketplace-specific expertise understanding complex workload patterns
- Comprehensive approach combining technical optimization with business strategy
- AWS Advanced Tier Consulting Partner certification
- Proven track record with marketplace and SaaS platforms

Target Outcomes



Comprehensive Assessment & Discovery

4-Week Deep Dive Analysis

Cloud Logically conducted an exhaustive evaluation across multiple dimensions:

Infrastructure Audit

- **Scope:** Complete AWS account review across EC2, RDS, S3, VPCs, load balancers, and CDN
- **Spend Analysis:** Detailed cost breakdown by service and business function
- **Usage Patterns:** Traffic analysis, seasonal hiring cycles, and geographic distribution

Technical Architecture Review

- **Scalability Assessment:** Auto-scaling configuration and database optimization opportunities
- **Performance vs. Cost Analysis:** Over-provisioning evaluation and right-sizing potential
- **Architecture Fit:** Monolithic vs. microservices impact on scaling and costs

Marketplace-Specific Considerations

- **Business Model Mapping:** Commission-based revenue correlation with infrastructure costs
- **AI/ML Workload Analysis:** Compute-intensive resume processing and matching algorithms
- **Seasonal Patterns:** Employment cycle impact on infrastructure requirements



Key Findings

Massive Over-Provisioning (70-80% idle capacity)

- **Web servers:** CPU utilization averaging 15-25%
- **Database instances:** Sized for peak load, idle 70% of time
- **Development environments:** Production-scale resources with <10% usage

Poor Resource Selection

- General-purpose instances for AI/ML workloads instead of GPU-optimized
- Older generation instances (t2/m4) vs. cost-effective newer generations
- Memory-intensive instances for CPU-bound matching algorithms

Network and Data Transfer Inefficiencies

- CDN misconfiguration serving non-cacheable assets
- Unnecessary cross-region data replication
- Inefficient geographic resource placement

Storage Waste

- S3 Standard storage for archived resume files
- Unoptimized backup retention policies
- Poor lifecycle management for development environments

Optimization Strategy & Implementation

Weeks 1-2: Assessment and Quick Wins

- Immediate cleanup of orphaned resources
- Comprehensive environment evaluation
- Comprehensive environment evaluation
- Immediate Impact: 15% cost reduction

Weeks 3-8: Core Optimizations

Compute Optimization:

- Right-sized EC2 instances (m5.2xlarge → m5.large for web servers)
- Implemented comprehensive Reserved Instance strategy
- Optimized auto-scaling policies for marketplace traffic patterns

Storage Transformation:

- Migrated to S3 Intelligent-Tiering (60% of resume files auto-moved)
- Implemented lifecycle policies (Glacier archiving after 1 year)
- EBS optimization (gp2 → gp3, 500GB → 200GB average sizing)

Database Optimization:

- Aurora Serverless for dev/staging environments
- Read replica implementation for production workloads
- Multi-AZ to Single-AZ transition for non-critical databases

Weeks 9-10: Automation and Monitoring

- Scheduled start/stop for development environments
- Automated resource lifecycle management
- Real-time cost monitoring and alerting dashboards
- Predictive scaling based on business patterns

Weeks 11-12: Knowledge Transfer and Documentation

- Comprehensive FinOps training for engineering teams
- Documentation of optimization playbooks
- Establishment of ongoing governance processes

Zero-Downtime Implementation

Savings by Category:

- **Right-Sizing & Auto-Scaling:** 40% of total savings
- **Storage Optimizations:** 25% of total savings
- **Reserved Instances/Savings Plans:** 30% of total savings
- **Automation & Governance:** 5% of total savings

Timeline Of Savings Realization



Performance Improvements

Enhanced System Performance:

- **Uptime Improvement:** 99.5% → 99.8%
- **API Latency Reduction:** 200ms → 150ms average response time
- **Peak Load Response:** 5 minutes → 30 seconds auto-scaling response
- **Backup Operations:** 40% faster backup and restore times

Operational Efficiency:

- **Engineering Time Saved:** 30 hours/week freed from infrastructure management
- **Fire-Fighting Reduction:** 60% decrease in emergency infrastructure issues
- **Development Velocity:** 25% increase in feature development speed

Business Impact

Financial Transformation

Immediate Financial Benefits:

- **Monthly Savings:** \$87.5K freed for strategic reinvestment
- **ROI Achievement:** 3-month payback on Cloud Logically engagement fees
- **12-Month ROI:** 420% return on optimization investment
- **Infrastructure Cost Ratio:** Reduced from 12% to 8% of revenue

Timeline Of Savings Realization

Week 2: 15% immediate savings from cleanup and initial right-sizing

Month 1: Additional 10% through Reserved Instance implementation

Month 3: Additional 8% from storage optimizations and automation

Month 6: Full 92% savings realized through comprehensive governance



Partnership Excellence

Collaborative Approach

Hands-On Collaboration:

- Daily Slack updates and transparent communication
- Weekly executive summaries with financial dashboards
- Bi-weekly technical deep dives with engineering teams
- Joint decision-making on all architectural changes

Trust Building Through Results

Credibility Established Early:

- 15% savings achieved within first two weeks
- Complete transparency in billing and methodology
- Zero surprises with clear change explanations
- Immediate ROI demonstration built long-term partnership foundation

Timeline Of Savings Realization

Extended Cash Runway: 8 additional months of operational runway, providing crucial breathing room for growth initiatives and fundraising activities.

Reinvestment Opportunities

From \$87.5K to \$7K Per Month savings enabled SkillGigs to:

Team Expansion: Hired 3 additional engineers focused on AI matching improvements

Marketing Investment: Increased marketing budget by 40% for user acquisition

Vertical Expansion: Funded expansion into healthcare talent markets

Security Enhancement: Invested in enterprise-grade security and compliance features

Investor And Market Position

Enhanced Investment Story:

- Demonstrated operational excellence and capital efficiency
- Improved unit economics and scalability metrics
- Extended runway increased valuation discussions
- Capital-efficient growth became key fundraising differentiator

Competitive Advantage:

- Lower infrastructure costs enabled more competitive pricing
- Engineering resource reallocation accelerated feature development
- Improved margins funded growth investments
- Market positioning strengthened vs. cloud-inefficient competitors

Knowledge Transfer:

- Four comprehensive FinOps workshops
- Hands-on training for monitoring and automation tools
- Detailed runbooks and optimization playbooks
- Ongoing advisory support for strategic decisions



Sustained Success And Future Roadmap

Ongoing Governance

Continuous Optimization:

- Automated monitoring dashboards tracking cost drift
- Quarterly optimization reviews and rightsizing assessments
- FinOps team established for ongoing cost management
- Executive dashboards monitoring cost-per-user trends

Future-Proofing Measures:

- Automated tagging and lifecycle policies
- Mandatory approval workflows for resource provisioning
- Developer training on cost-conscious architecture patterns
- Predictive cost modeling for growth planning

Continued Partnership

Cloud Logically maintains an ongoing advisory relationship providing:

- Monthly infrastructure health checks
- Quarterly strategic optimization planning
- On-call support for major architectural decisions
- Annual comprehensive optimization reviews

Future Projects

Current Advisory Scope:

- AI/ML infrastructure optimization for enhanced matching algorithms
- Quarterly strategic Multi-region deployment strategy for global expansion
- Healthcare compliance infrastructure for vertical expansion
- Advanced FinOps maturity with predictive cost modeling

Client Testimonial



Kashif Aftab
CEO -
SkillGigs

"Cloud Logically didn't just cut our costs – they transformed our entire approach to cloud infrastructure. The 92% savings were incredible, but what really impressed us was their deep understanding of marketplace dynamics and their commitment to building our internal capabilities. They're not just consultants; they're strategic partners who invested in our long-term success."

Key Takeaways

For Growing SaaS Companies

Continuous Optimization:

- **Cost Optimization ≠ Performance Compromise:** Strategic optimization can simultaneously reduce costs and improve performance when executed properly
- **Early Intervention is Critical:** Addressing cost inefficiencies before they become critical business risks enables more strategic options and better outcomes.
- **Partnership Over Consulting:** Deep collaboration and knowledge transfer create sustainable improvements beyond one-time optimizations.

For Marketplace Platforms

- **Workload-Specific Optimization:** Understanding marketplace traffic patterns, seasonal cycles, and user behavior enables targeted optimization strategies.
- **AI/ML Cost Management:** Specialized expertise in optimizing compute-intensive AI workloads can yield dramatic savings without compromising algorithm performance.
- **Scalability Planning:** Right-sizing for current needs while maintaining scalability for growth requires sophisticated architectural planning.

Ready To Transform Your Cloud Costs?



Cloud Logically's proven methodology has helped companies like SkillGigs achieve dramatic cost reductions while improving performance and enabling strategic growth.

Contact us today to discover how we can optimize your AWS infrastructure, reduce costs, and accelerate your business growth.

Cloud Logically is an AWS Advanced Tier Consulting Partner specializing in cost optimization, performance enhancement, and strategic cloud architecture for high-growth SaaS and marketplace platforms.

Give Your Business the Cloud Logically Edge

Get free Consultation:

www.cloudlogically.com

support@cloudlogically.com

+1 (713) 315 7076

