

AI Studio & AI Guides

Genesys PureCloud AI Studio & AI Guides Documentation

Study Notes

Topic	Description
AI Studio	Centralized workbench for building, managing, and deploying AI-powered experiences
AI Guides	No-code feature that converts business instructions into Virtual Agents using natural language
Purpose	Enable business users to create intelligent virtual agents without coding expertise
Licensing	Requires Genesys Cloud AI Experience tokens (metered pricing)
Innovation Level	Level 4 agentic AI - semi-autonomous with defined boundaries and guardrails

Navigation

Admin → AI Studio OR Admin → Contact Center → Automation → AI Studio

AI Studio Overview

Genesys Cloud AI Studio provides a centralized workbench to build, manage and deploy AI experiences like Virtual Agents and Agent Copilots. It serves as the command center for organizations to create next-generation AI-powered customer experiences with governance, control, and scalability.

Key Capabilities

- Natural language editor with no-code and low-code tools for users without technical know-how
- Unified environment for designing AI agents across all channels
- Built-in governance and compliance controls
- AI Guides that allow business teams to create Virtual Agents that respond intelligently to customer context and adapt their behavior dynamically within conversations
- Customizable summaries for interactions
- Integration with existing Architect Virtual Agent flows
- Version control and deployment management
- Performance analytics and monitoring

Strategic Positioning

AI Studio represents a leap from simple execution to intelligent problem-solving - Level 4 where semi-autonomous systems are configured around specific objectives using reasoning, planning and memory to figure out how best to accomplish goals while still operating within clearly defined boundaries to meet compliance and policy requirements

Edition & Module Requirements

Requirement	Details
Minimum Edition	Genesys Cloud CX 1, CX 2, CX 3, or CX 4 license
Licensing Model	AI Experience tokens (metered, usage-based pricing)
Permissions	Role-based AI Studio permissions required
Additional	Access to Virtual Agent module for full functionality

Study Notes - AI Studio

Components

Component	Purpose	User Type
AI Guides	Create Virtual Agents from natural language prompts	Business users, CX teams
Guide Editor	Build and refine guide instructions and logic	All user levels
Virtual Agent Integration	Deploy guides to Virtual Agent flows	Technical users
Customizable Summaries	Shape interaction summaries to business needs	Admins, supervisors
Governance Controls	Ensure compliance and policy adherence	Admins, compliance
Performance Dashboard	Monitor guide usage and effectiveness	Managers, analysts
API Access	Programmatic guide and agent management	Developers

AI Guides Overview

Genesys Cloud AI Guides enables business users to create AI-powered Virtual Agents through natural language instructions without coding, facilitating conversational flows that mirror customer journeys, adapt dynamically to customer context, and combine structured logic with AI capabilities

How AI Guides Work

1. Business user describes goal in plain language or uploads process documentation
2. AI Guides use large language models (LLMs)-based natural language processing to interpret user prompts and documents, generating complete agentic flows including intents, slots and dialog logic
3. System generates draft virtual agent with conversation flow
4. User reviews, edits, and customizes the generated guide
5. AI executable instructions produced by AI Guides are fully editable within Genesys Cloud, allowing users to update messages, logic and backend integrations before publishing
6. Guide is published and connected to Virtual Agent
7. Virtual Agent uses guide to handle customer interactions

Key Features

- Natural Language, No Code Required - easily build or refine virtual agents using plain language or existing documentation with no coding skills needed
 - Build Once, Deploy Anywhere - design experiences once and deploy them across Genesys Cloud Virtual Agent and Copilots, and more to maintain consistency and reduce duplication of effort
 - Enterprise-Grade Collaboration - seamlessly connect front, middle and back-office systems to execute tasks, automate workflows and deliver measurable business outcomes
 - Guardrails Built In - implement configurable, testable safety controls to help enable increased accuracy, appropriate tone and policy compliance to support responsible adoption of agentic AI
 - Knowledge integration - guides can use either knowledge Workbench v2 knowledge bases or knowledge fabric configurations to answer customer questions at any point in a conversation
-

Implementation Guide

Step 1: Assessment & Planning

1. Identify suitable use cases for AI Guides
2. Document business processes and customer journeys
3. Gather process documentation or playbooks
4. Define success metrics for each guide
5. Plan escalation scenarios
6. Assess team readiness for AI automation
7. Plan change management approach

Step 2: Licensing & Setup

1. Ensure organization has Genesys Cloud CX license (CX 1, CX 2, CX 3, or CX 4)
2. Purchase Genesys Cloud AI Experience tokens
3. Add necessary AI Studio permissions
4. Set up role-based access controls
5. Configure Virtual Agent if not already enabled
6. Test integration with backend systems
7. Establish governance policies

Step 3: Guide Creation

1. Navigate to AI Studio
2. Create a guide using an AI prompt, convert a process document into a guide, or create from scratch by starting with a blank guide
3. Describe goal in natural language or upload documentation
4. Review AI-generated guide structure
5. Edit guide instructions and customize flow
6. Add variables and data integrations
7. Configure escalation paths

Step 4: Testing & Refinement

1. Preview guide behavior in test environment
2. Author preview before publish - preview real knowledge responses during Guide configuration to confirm accuracy and behavior
3. Test with sample customer scenarios
4. Verify escalation triggers work correctly
5. Validate data integrations
6. Test across all supported channels
7. Gather feedback from SMEs

Step 5: Publishing & Deployment

1. Publish guide to Virtual Agent
2. Connect the guide to Virtual Agent flows in Architect
3. Assign to production queues
4. Monitor initial interactions closely
5. Validate customer experience
6. Adjust guide parameters based on feedback
7. Scale to additional queues as needed

Step 6: Monitoring & Optimization

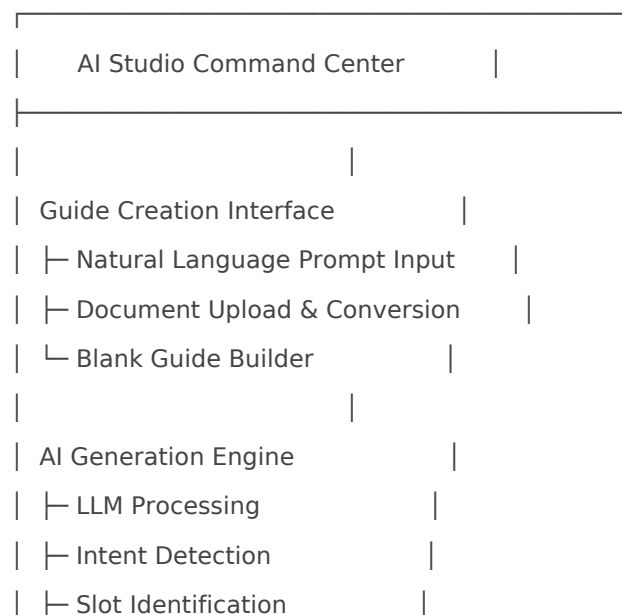
1. Monitor guide usage metrics daily
2. Track customer satisfaction and resolution rates
3. Review escalation patterns
4. Analyze customer feedback
5. Refine guide instructions based on data
6. A/B test different guide variations
7. Update regularly based on learnings

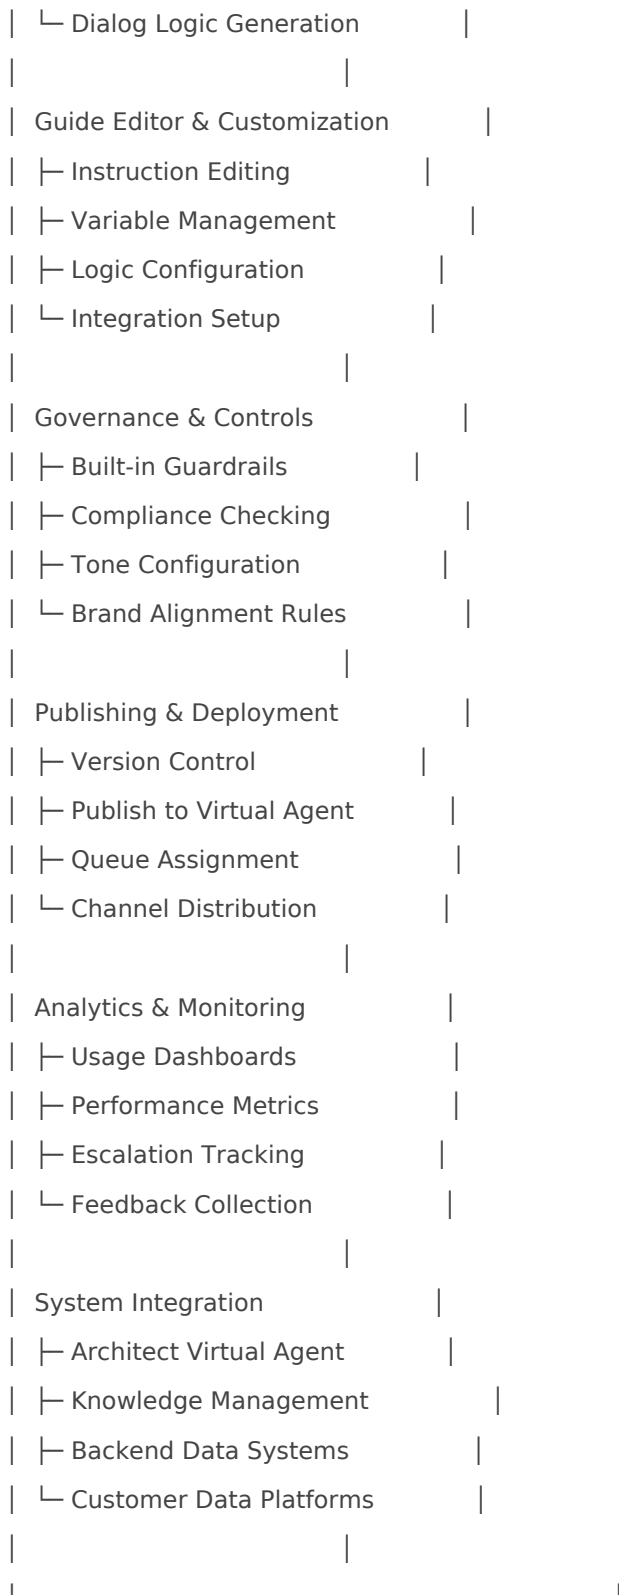
How to Implement

Phase	Description	Timeline
Planning	Identify use cases, document processes, assess readiness	Week 1-2
Setup	Activate licenses, configure permissions, enable integrations	Week 2-3
Creation	Build guides from prompts or documents, test	Week 3-5
Testing	Validate behavior, test scenarios, refine	Week 5-6
Pilot	Deploy to select queues, monitor closely	Week 6-7
Rollout	Expand to production, scale across organization	Week 7-8
Optimization	Monitor, analyze, improve continuously	Ongoing

AI Studio & AI Guides Architecture

AI Studio Centralized Workbench





↓

Virtual Agent Deployment

↓

Customer Interactions

AI Guides Use Cases & Examples

Use Case 1: Order Status & Tracking

Business Process Documentation:

1. Collect order number from customer
2. Look up order in system
3. Provide tracking information
4. Offer additional help options
5. Close or escalate as needed

AI Guide Generated:

Utterances:

- ├ "Where's my order?"
- ├ "Track my package"
- ├ "Order status"
- └ "When will my order arrive?"

Intents:

- ├ Track Order
- ├ Delivery Status
- └ Tracking Number

Dialog Flow:

- ├ Ask for order number
- ├ Query order database
- ├ Provide shipping status
- ├ Offer further assistance
- └ Escalate if needed

Resolution Rate: 85-90% (self-service)

Use Case 2: Password Reset Process

Uploaded Process Document:

"Follow these steps for password reset:

1. Verify customer identity with security questions
2. Confirm email address on file
3. Send password reset link
4. Confirm reset completed
5. Offer additional support"

AI Guide Generated:

Intents:

- └ Password Reset Request
- └ Account Access Issue
- └ Security Verification

Dialog:

- └ "I'll help you reset your password"
- └ "Let me verify who you are"
- └ Customer answers security questions
- └ System confirms identity
- └ "Check your email for reset link"
- └ "Did you successfully reset?"
- └ Provide follow-up support

Resolution Rate: 92-95% (self-service)

Use Case 3: Appointment Scheduling

Prompt Input:

"Create a guide for customers to book service appointments. Must:

- Ask for service type
- Show available dates/times
- Confirm appointment
- Send confirmation"

AI Guide Generated:

Intents:

- └ Schedule Appointment
- └ Change Appointment
- └ Cancel Appointment

Dialog:

- └ "What service do you need?"
- └ "When works best for you?"
- └ Show available slots
- └ "Let me confirm..."
- └ Send calendar confirmation
- └ Offer reschedule option

Resolution Rate: 88-92% (self-service)

Use Case 4: Billing Question & Payment

Process Flow:

"Customers with billing questions:

1. Verify account
2. Explain charges
3. Offer payment options
4. Process if customer agrees
5. Send receipt/confirmation"

AI Guide Generated:

Intents:

- └ Check Bill Amount
- └ Explain Charges
- └ Make Payment
- └ Dispute Charge

Dialog:

- └ Confirm customer identity
- └ "Let me look up your account"
- └ Display current balance
- └ "Would you like to pay?"
- └ Secure payment processing
- └ Send receipt and confirmation
- └ Escalate disputes

Resolution Rate: 80-85% (self-service)

Modular Guides Within Virtual Agents

AI Guides are designed to create modular bot flows that can be combined within a single virtual agent - each guide typically addresses a specific use case like order tracking, password reset or appointment booking, and these flows can be added to your virtual agent's overall configuration

Multi-Guide Architecture Example

Virtual Agent: Customer Service Bot

└ Guide 1: Order Tracking

- | └ Intents: Track order, delivery status
- | └ Resolution: Order lookup + tracking info
- | └ Escalation: Complex shipment issues

└ Guide 2: Account Management

- | └ Intents: Password reset, update info
- | └ Resolution: Self-service account changes
- | └ Escalation: Security concerns

└ Guide 3: Billing & Payments

- | └ Intents: Check bill, make payment
- | └ Resolution: Payment processing
- | └ Escalation: Billing disputes

└ Guide 4: Appointment Booking

- | └ Intents: Schedule, reschedule, cancel
- | └ Resolution: Appointment management
- | └ Escalation: Complex scheduling needs

Router (AI determines which guide needed):

Customer: "Where's my order?"

→ Route to Guide 1: Order Tracking

→ Resolve: Self-service tracking info

Customer: "I want to schedule a service"

→ Route to Guide 4: Appointment Booking

→ Resolve: Appointment confirmation

Customer: "I have a billing question"

→ Route to Guide 3: Billing & Payments

→ Resolve or escalate accordingly

Knowledge Integration with AI Guides

In a future release, Genesys Cloud will allow AI Guides to use knowledge articles to answer customer questions while continuing through a defined process - guides can answer customer questions using approved knowledge content at any point in a conversation and then continue the task without losing context

Knowledge-Aware Guide Benefits

Before Knowledge Integration:

Customer: "What's your return policy?"

Guide: "Let me connect you with an agent
who can answer policy questions"

→ Escalation (unnecessary)

After Knowledge Integration:

Customer: "What's your return policy?"

Guide: [Accesses knowledge base]

"Here's our return policy...

Now, back to your order tracking,

I see your item was delivered..."

→ Continues conversation naturally

→ No escalation needed

Configuration Options

Guide authors can choose to inherit knowledge from the connected Virtual Agent or select a Guide-specific Knowledge source

Flexible knowledge sourcing:

- Use Virtual Agent's knowledge base
- Configure guide-specific knowledge source
- Support for knowledge Workbench v2
- Support for knowledge fabric configurations

Real Flow Scenario: AI Guide in Action

Customer Interaction Example:

Customer Calls: "I need to reset my password"

↓

Virtual Agent Answers (AI Guide: Account Management)

"Hi! I'm here to help. To reset your password, I'll need to verify your identity. What's your email address on file?"

↓

Customer: "john.smith@email.com"

↓

Guide Verifies: Customer email matches account

↓

Guide: "Thanks. What was your first pet's name?"

↓

Customer: "Fluffy"

↓

Guide Confirms: Security answer matches

↓

Guide: "Perfect! I've sent a password reset link to your email. Please check your inbox and click the link to set a new password."

↓

Customer: "Got it, thanks!"

↓

Guide: "You're welcome. Is there anything else
I can help with today?"

↓

Customer: "No, that's all"

↓

Guide: "Have a great day!"

↓

Interaction Complete:

├ Resolution: Self-service password reset

├ Tokens Used: 1-2 per interaction

├ Customer Satisfaction: High (immediate help)

└ Cost: ~\$0.05-0.15 per interaction

Licensing & Token Pricing

AI Experience Token Model

Organizations need Genesys Cloud AI Experience tokens to use AI Studio, with tokens based pricing model to monitor feature usage and consumption

Token Consumption

AI Guides consume tokens for each interaction session and for each guide created

Pricing Considerations

- **Per-interaction tokens** - Consumed when customers interact with guides
- **Per-guide tokens** - Consumed when creating new guides
- **Metered model** - Pay for what you use
- **Flexibility** - Tokens can scale up or down based on demand
- **Multiple LLM support** - AI Studio is compatible with proprietary, open source and Amazon Bedrock large language models (LLMs) and advanced frontier models from companies such as OpenAI, Anthropic and Google

Token Consumption Example

Small Organization:

- ┆ 3 AI Guides created (3 guides × tokens)
- ┆ 500 customer interactions/month (500 × tokens)
- ┆ Monthly token usage: ~750 tokens
- ┆ Estimated cost: \$XX-XXX/month

Mid-Market Organization:

- ┆ 10 AI Guides created (10 guides × tokens)
- ┆ 5,000 customer interactions/month (5,000 × tokens)
- ┆ Monthly token usage: ~5,100 tokens
- ┆ Estimated cost: \$XXX-XXXX/month

Enterprise Organization:

- ┆ 50 AI Guides created (50 guides × tokens)
- ┆ 50,000 customer interactions/month (50,000 × tokens)
- ┆ Monthly token usage: ~50,100 tokens
- ┆ Estimated cost: \$XXXX-XXXXX/month

Best Practices

Guide Design

- **Clear intent** - Define specific goal for each guide
- **Natural language** - Write instructions as you would for an employee
- **Process documentation** - Use existing playbooks and procedures
- **Modular approach** - Create focused guides for specific tasks
- **Escalation paths** - Define clear handoff scenarios
- **Testing** - Always test with real scenarios before production
- **Iteration** - Guides improve with updates and refinement

Guardrails & Governance

- Implement configurable, testable safety controls to help enable increased accuracy, appropriate tone and policy compliance to support responsible adoption of agentic AI
- Set boundaries for agent autonomy
- Define tone and brand voice
- Establish compliance requirements

- Monitor sensitive operations
- Implement audit trails
- Regular compliance reviews

Knowledge Management

- Author preview before publish - preview real knowledge responses during Guide configuration to confirm accuracy and behavior
- Keep knowledge articles current
- Ensure accurate information
- Update for policy changes
- Test knowledge accuracy
- Monitor article usage
- Gather feedback from guides using knowledge

Performance Optimization

- Monitor token usage closely
 - Track guide effectiveness metrics
 - Analyze escalation reasons
 - Gather customer feedback
 - Refine guides based on data
 - A/B test different approaches
 - Regular guide audits and updates
-

Common Implementation Scenarios

Scenario 1: Quick Start (Small Team)

Timeline: 2-3 weeks

Setup:

├ 2-3 simple guides

├ Basic integration (order lookup, FAQ)

├ Limited escalation paths

- └ Single channel deployment

Expected Results:

- └ 60-70% automation of routine inquiries
- └ Quick implementation, minimal training
- └ Fast ROI (4-6 weeks)
- └ Monthly token usage: ~300-500

Scenario 2: Comprehensive (Mid-Market)

Timeline: 6-8 weeks

Setup:

- └ 8-12 specialized guides
- └ Multiple backend integrations
- └ Complex escalation logic
- └ Omnichannel deployment
- └ Knowledge base integration

Expected Results:

- └ 50-65% automation of customer volume
- └ Significant cost savings
- └ 8-12 week ROI
- └ Monthly token usage: ~3,000-5,000

Scenario 3: Enterprise (Full Scale)

Timeline: 12-16 weeks

Setup:

- └ 30-50+ specialized guides
- └ Full system integration
- └ Advanced routing and logic
- └ Multi-language support
- └ Comprehensive knowledge integration
- └ Global channel deployment

Expected Results:

- └ 40-60% automation of global volume
- └ Major cost reduction

- └ Continuous optimization
- └ 6-10 week ROI
- └ Monthly token usage: ~20,000-50,000+

Troubleshooting Guide

Issue	Cause	Resolution
Guide doesn't understand customer intent	Insufficient training variations	Add more example utterances to training data
Escalations happening too frequently	Over-strict guardrails or incomplete logic	Expand guide capabilities and relax thresholds
Inaccurate information from knowledge	Stale knowledge articles	Update knowledge base and preview before publish
Slow response time	Token consumption or system latency	Optimize guide logic and integrations
Integration failures	Backend system connectivity issues	Verify API connections and data endpoints
Guides not deployed to queues	Missing Virtual Agent configuration	Check Architect flow configuration and deployment
Customer dissatisfaction	Poor conversation quality	Refine guide instructions and tone
Token overages	Guides consuming more than expected	Analyze usage patterns and optimize interactions
Editing feels clunky	Unfamiliar with new guide model	Review updated guide model documentation
Knowledge not accessible in guides	Knowledge not properly configured	Configure knowledge source and test access

AI Guides vs. Traditional Bot Flows

Feature	AI Guides	Traditional Flows
Creation Method	Natural language prompts	Manual design
Skill Required	Business knowledge	Technical/coding
Speed to Deploy	Days to weeks	Weeks to months
Iteration Time	Minutes to hours	Hours to days
Intent Recognition	AI-powered, flexible	Rule-based, rigid

Feature	AI Guides	Traditional Flows
Conversation Quality	Natural, contextual	Scripted, menu-driven
Maintenance	AI learns from data	Manual updates
Integration	Automatic via AI	Manual configuration
Cost to Build	Low (no developers)	High (developer time)
Scaling	Easy, rapid deployment	Time-consuming

Real-World Timeline Example

Week 1-2: Planning & Assessment

Day 1-3: Kickoff and planning

- └ Identify 3-5 use cases
- └ Gather process documentation
- └ Assess team readiness

Day 4-10: Documentation review

- └ Refine process flows
- └ Identify integration needs
- └ Plan escalation scenarios

Day 11-14: Setup and licensing

- └ Purchase AI Experience tokens
- └ Configure permissions
- └ Set up integrations

Week 3-4: Guide Creation

Day 15-20: First guide creation

- └ Upload process documentation
- └ Review AI-generated guide
- └ Customize and refine
- └ Connect to Virtual Agent

Day 21-28: Additional guides

- └ Create guides 2-4
- └ Refine based on feedback
- └ Configure escalation logic
- └ Integration testing

Week 5-6: Testing & Refinement

Day 29-35: Comprehensive testing

- └ Test all guide scenarios
- └ Validate integrations
- └ Test escalation paths
- └ Gather feedback

Day 36-42: Optimization

- └ Refine guide instructions
- └ Adjust confidence thresholds
- └ Optimize knowledge integration
- └ Performance tuning

Week 7-8: Deployment

Day 43-49: Pilot deployment

- └ Deploy to pilot queue
- └ Monitor closely
- └ Gather customer feedback
- └ Make adjustments

Day 50-56: Full rollout

- └ Expand to all queues
- └ Monitor metrics
- └ Provide agent support
- └ Celebrate success

Interview Cheat Sheet

Question	Answer
What is AI Studio?	Centralized workbench for building, managing, and deploying AI experiences
What is an AI Guide?	No-code feature that converts business instructions into Virtual Agents using AI
What licensing is required?	Genesys Cloud CX license + AI Experience tokens
How do AI Guides work?	Users describe goals in plain language, AI generates guide structure, user customizes
Can non-technical users create guides?	Yes, AI Guides are designed for business users without coding skills
What level of agentic AI is this?	Level 4 - semi-autonomous with defined boundaries and guardrails
How long to create a guide?	Minutes to hours depending on complexity
Can guides be combined?	Yes, multiple guides can be modular components in one Virtual Agent
How are guides customized?	Fully editable - users can change messages, logic, and integrations
What about knowledge integration?	Guides can access knowledge articles to answer questions within conversation
What's the pricing model?	Token-based - consumption tracked per interaction and per guide created
Can guides handle escalation?	Yes, with configurable escalation paths to human agents
What channels do guides support?	All Genesys channels - voice, chat, email, messaging, etc.
How do guides improve over time?	AI learns from interactions; user refines guides based on metrics
What's the expected ROI timeline?	4-12 weeks depending on implementation scope

Key Takeaways

- **No-Code Creation** - Natural language, no code required - easily build or refine virtual agents using plain language or existing documentation with no coding skills needed
- **Intelligent Automation** - AI-powered guides intelligently handle customer conversations with reasoning and planning
- **Enterprise Control** - Guardrails built in - implement configurable, testable safety controls to help enable increased accuracy, appropriate tone and policy compliance
- **Rapid Deployment** - Create and deploy guides in days, not months
- **Modular Design** - Build once, deploy anywhere across virtual agents, copilots and more to maintain consistency and reduce duplication of effort

- **Token-Based Pricing** - Pay only for what you use with flexible scaling
 - **Knowledge Integration** - Guides can answer customer questions using approved knowledge content during any step of a process
 - **Business User Friendly** - Designed for CX teams, not IT/developers
 - **Continuous Learning** - Guides improve as they handle more interactions
 - **Compliance Ready** - Built-in governance and audit capabilities
-

Customizable Summaries (Additional AI Studio Feature)

Customizable Summaries allows admins to shape interaction summaries to fit their business needs, enhancing consistency, compliance, and operation efficiency across human and AI agent interactions

Use Cases

- Tailor summaries to specific business requirements
 - Ensure compliance with regulatory standards
 - Align summaries with operational priorities
 - Support multiple languages and regions
 - Customize for different teams and use cases
-

Getting Started Checklist

Pre-Implementation

- Assess current contact center operations
- Identify 3-5 suitable use cases
- Gather process documentation and playbooks
- Determine success metrics for each guide
- Assess team readiness for AI automation
- Plan change management approach

Licensing & Setup

- Purchase Genesys Cloud AI Experience tokens
- Assign AI Studio permissions to team
- Configure role-based access controls
- Enable Virtual Agent module
- Set up backend system integrations
- Test knowledge base connectivity

Guide Development

- Create first AI Guide from process document
- Test guide functionality thoroughly
- Customize instructions and logic
- Configure escalation paths
- Connect to Virtual Agent
- Deploy to pilot queue

Monitoring & Optimization

- Monitor guide performance daily
 - Track customer satisfaction metrics
 - Analyze escalation patterns
 - Refine guides based on data
 - Scale to additional queues
 - Plan continuous improvements
-

Additional Resources

Official Documentation Links

- AI Studio Overview: help.genesys.cloud/articles/about-ai-studio/

- AI Guides Overview: help.genesys.cloud/articles/ai-guides-overview/
- Knowledge Integration: help.genesys.cloud/announcements/knowledge-integration-for-ai-guides/
- AI Studio Permissions: help.genesys.cloud/articles/ai-studio-permissions/

Support Contacts

- Genesys Sales: sales@genesys.com
 - Genesys Support: <https://support.genesys.com>
 - AI Services Team: ai-services@genesys.com
 - Community Forums: <https://community.genesys.com>
-

Document Version Info

Last Updated: March 2026

Source: Genesys PureCloud Official Documentation (help.genesys.cloud)

Validated: Current with January-March 2026 releases

Version: 1.0

Revision #1

Created 13 March 2026 19:19:40 by Cesar Gzz

Updated 14 March 2026 19:35:03 by Cesar Gzz