



Agentic Work Transformation: From Prototype to Production with Gemini Enterprise

A Strategic Webinar Series for Operationalizing AI



Meet Your Presenters



Brandon Carter

Marketing Director



Kylan Thomson

Solutions Architect



Brian Kinlaw

Sr. Solutions Architect



Webinar Series Schedule



Webinar 1

July 7 • 1:00 PM

THE ART OF POSSIBLE

Define business outcomes and identify where AI can create measurable impact.



Webinar 2

July 9 • 1:00 PM

THE FOUNDATION

Ground agents in trusted enterprise data architectures for reliable outputs.



Webinar 3

July 14 • 1:00 PM

THE WORKSHOP

Build your first functional, no-code virtual agent prototype step-by-step.



Webinar 4

July 16 • 1:00 PM

THE ASSEMBLY LINE

Scale your virtual agent for robust, enterprise-grade production environments.



The Strategic Shift

The Big Picture

AI is moving beyond sandbox experimentation.

Organizations are now focused on moving from one-off prompts to reliable, scalable, and secure AI workflows that create measurable business impact.

Key shift:

From AI experimentation → to AI operationalization

The Brain & Body Advantage

Gemini Enterprise is the brain.

It grounds AI in secure, proprietary enterprise data.

Gemini for Workspace is the body.

It brings that intelligence into the tools employees already use every day — Gmail, Docs, Sheets, and Chat.



+35% Operational Velocity



Our Journey Map: From Prototype to Production

This series walks through the practical path from AI experimentation to production-ready agents.



Session 1: Define

The Art of Possible

Define business outcomes and identify where AI can create measurable impact.



Session 2: Ground

The Foundation

Ground agents in trusted enterprise data.



Session 3: Build

The Workshop

Build your first no-code prototype.



Session 4: Scale

The Assembly Line

Scale your agent for production.




Session 1: The Art of Possible



The Art of Possible

Reimagine the way work gets done with agents that bring together Gemini's advanced reasoning, Google-quality search, and enterprise data





The “Why”: Defining Business Outcomes

Successful AI initiatives start with the business problem — not the technology.



The Business Gap

Teams are slowed down by siloed data, manual processes, and disconnected workflows.



The Agentic Solution

Gemini Enterprise connects trusted data with AI agents that can reason, retrieve, and assist across workflows.



Measurable Business Impact

The goal is to create outcomes the business can track — like revenue growth, operational efficiency, increased ROI, and reduced churn.

Shift the question from: “What can Gemini do?” to **“What business problem can we solve?”**



Showcase Customer Outcomes

The End State: Prove Production-Grade Impact



HCLS: Clinical Trial Data Assistant

Helps teams analyze siloed clinical trial data, improve consistency, and prepare submission-ready materials faster.

-50% Reduced siloed trial analysis time



Banking: Compliance Boundaries Agent

Reviews customer or loan scenarios against defined policy boundaries, flags exceptions, and supports auditability.

98% Accuracy on qualifying transactions



Manufacturing: Procurement Agent


Automates low-risk procurement decisions, routes exceptions for review, and improves operational efficiency.

+35% Operational efficiency improvement




The Integrated Experience: Invocation Within Everyday Tools


Gemini brings trusted enterprise intelligence into the tools teams already use.




How We Search
Find answers faster.



How We Create
Generate useful content.



How We Write
Draft in the flow of work.



How We Learn
Understand context quickly.



Gemini Enterprise

Securely Linking to Enterprise Data Sources

Gemini Enterprise helps agents connect to the systems and knowledge your teams already rely on.





The “Brain and Body” Analogy

Agentic Work Transformation happens when secure enterprise intelligence is connected to the tools employees already use.



Gemini Enterprise = The Brain

Grounds AI in your organization's proprietary data, business context, and governance.



Gemini for Workspace = The Body

Brings that intelligence into everyday work through Gmail, Docs, Sheets, and Chat.

Together, they turn AI from a separate tool into an integrated way of working.



Governance as a Foundation: Security from Day One

Production-grade AI requires governance to be built into the process from the start — not added at the end.

Promevo Trust Acceleration Framework (TAF)



Define the Roadmap

Align on the business outcome, data sources, and production path.



Establish Principles

Set expectations for security, compliance, accuracy, and responsible use.



Apply Access Controls

Use role-based permissions so users only access authorized data.



Keep Humans in the Loop

Review, evaluate, and improve agent outputs before scaling.

Governance is what makes AI secure, auditable, and ready for production.



The Executive Dilemma

The Problem: **The Data Swamp**

Organizations are moving quickly to deploy generative AI, but many are building on fragmented, unclassified data.

Shadow Data

You can't secure what you can't see.

Compliance Risk

Sensitive data must be protected before AI can safely use it.

Hallucination Risk

Conflicting data silos can produce confident, costly errors.

The Solution: **The 21-Day Agentic Readiness Lab (ARL)**

A focused three-week sprint designed to assess whether your environment is ready for secure, compliant GenAI.

What it Delivers:

- A clearer view of your data landscape for your "Human Firewall"
- Security and governance recommendations enable AI Guardians
- Production-ready readiness artifacts feed directly into the Trust Acceleration Framework
- A practical path toward enterprise AI deployment

Session 1 Completed: The Roadmap Begins

Webinar Series Tracker

1 of 4 Sessions Complete

1

Session 1: The Art of Possible

Completed today. Defined outcomes & aligned business vision.

2

Session 2: The Foundation

Up next: July 9. Grounding agents in secure data structures.

3

Session 3: The Workshop

Coming July 14. Hands-on no-code assistant builder.

4

Session 4: The Assembly Line

Coming July 16. Scaling for enterprise production.

Bridge Today's Gap: Start Mapping Outcomes

Moving from general concepts to real workflows requires structured planning. Align your stakeholders before our next technical deep dive.

Contact us to learn about our Gemini Enterprise Accelerator program.

promevo.com/contact-us



Session 2: The Foundation



The AI Paradox

Visually contrasting sheer brilliance with enterprise disconnectedness.

Sheer Brilliance



The Brilliant LLM (Creative Genius)

Sonnet to a Metallic Toaster

*"Thou silent sentinel of morning's grace,
Whose elements of fire begin to glow,
To heat the bread within thy warm embrace,
And crisp the slice with steady heat below."*

- Flawless, creative, yet entirely disconnected

Utterly Useless



The Enterprise Reality (Siloed & Blind)

"Analyze our Q2 regional pipeline performance against budget."

System Access Restricted

"I'm sorry, as an AI, I do not have access to your internal corporate databases or secure local business systems."



The Foundation: Grounding Agents in Your Data

An agent is only useful when it can access the right data, from the right systems, with the right permissions.



Goal: Move from a promising prototype to a reliable, secure, production-ready agent.




The “How”: Grounding in Reality

An agent is only as good as the data it can access.

To move from generic AI responses to useful business insights, agents need to be grounded in trusted enterprise data.

- Connect agents to approved business systems
- Ground responses in real company knowledge
- Improve trust, accuracy, and relevance
- Create answers teams can actually use




The average company
interacts with

254

distinct AI applications

At work, our capabilities are limited by
an abundance of siloed applications
that do single tasks

1. Harmonic Security






Connecting Your Data Universe

Application Connectors

Work happens across dozens of applications. Gemini Enterprise helps bring that information together so agents can search, reason, and respond with context from trusted sources.

- Connect to the systems teams already use
- Reduce the friction of siloed applications
- Ground agent responses in approved data sources
- Create a more connected knowledge layer





From Data to Insight

Grounding turns disconnected business data into specific, trustworthy answers.

Generic AI



What was our Q1 revenue for the Project Phoenix account?



I don't have access to your company's internal sales data.



Grounded Agent



What was our Q1 revenue for the Project Phoenix account?



Project Phoenix generated \$1.2M in Q1 revenue, based on the approved sales report and finance summary.



From Insight to Actionable Workflow

The value of a grounded agent doesn't stop at answering a question. It can help move the next step of work forward.

- Turn insights into drafts, summaries, or recommendations
- Push information into Docs, Sheets, Gmail, or Chat
- Reduce manual handoffs between tools
- Keep a human in the loop before action is finalized






Generic AI




What was our Q1 revenue for the Project Phoenix account?



Grounded Agent



What was our Q1 revenue for the Project Phoenix account?



Moving the next step of work forward.



Security and Access Control

Grounded agents should only access the data each user is authorized to see.

- Apply role-based access controls across connected sources
- Respect existing permissions from enterprise systems
- Protect sensitive information from unauthorized access
- Keep governance tied to the workflow

User Role	Drive	Salesforce	BigQuery	HR Docs
Sales	✓	✓		
Finance	✓		✓	
HR	✓			✓




The Human-in-the-Loop Principle

Production-ready agents still need human oversight.

Human review helps teams validate outputs, catch gaps, and improve the agent before it is scaled across the business.

- Review responses for accuracy and context
- Provide feedback when outputs miss the mark
- Refine instructions and data sources over time
- Approve key actions before they move forward

Human oversight is not a blocker. It is how agents become trustworthy.



Session 2 Completed: The Data Secured

Webinar Series Tracker

2 of 4 Sessions Complete



Session 1: The Art of Possible

Completed today. Defined outcomes & aligned business vision.

2

Session 2: The Foundation

Up next: July 9. Grounding agents in secure data structures.

3

Session 3: The Workshop

Coming July 14. Hands-on no-code assistant builder.

4

Session 4: The Assembly Line

Coming July 16. Scaling for enterprise production.

Solve the Executive Dilemma & Avoid the Swamp

Now that we have covered grounding architectures, let's assess your actual infrastructure for secure enterprise Gemini integration.

Contact us to learn about our Gemini Enterprise Accelerator program.

promevo.com/contact-us



Session 3: The Workshop



Meet Your Presenters



Christine Page
Sr. Marketing Manager



Lauryn Thomas
Associate Solutions Architect



Brian Kinlaw
Sr. Solutions Architect

Today's Agenda

Webinar Series Schedule



Webinar 1

July 7 • 1:00 PM

THE ART OF POSSIBLE

Define business outcomes and identify where AI can create measurable impact.



Webinar 2

July 9 • 1:00 PM

THE FOUNDATION

Ground agents in trusted enterprise data architectures for reliable outputs.



Webinar 3

July 14 • 1:00 PM

THE WORKSHOP

Build your first functional, no-code virtual agent prototype step-by-step.

- Part 1: The Builder Interface (Touring the no-code environment)
- Part 2: Build-Along (Creating your first Sales Analyst agent)
- Part 3: Testing & Refinement (Validating outputs and iterative design)
- Part 4: Wrap up.



Webinar 4


July 16 • 1:00 PM

THE ASSEMBLY LINE

Scale your virtual agent for robust, enterprise-grade production environments.

➤ The Illusion of Viral Demos


Moving from perception (demos) to reality (enterprise production requirements).



Insulating Complexity with GEAP

The Gemini Enterprise Agent Platform abstracts away infrastructure burdens, letting you focus on outcomes.

◆ Gemini Enterprise Agent Platform



AI Integration Strategies

Understanding the increasing intelligence & complexity of AI tools.

Direct API Call

WHAT

A direct API call is a simple request-response: send a prompt, get an answer. No memory or built-in tools.

WHEN

Use a direct API call for simple, single requests where you want full control and speed

AI Application (Assistant)

WHAT

An AI application is a ready-made tool with a UI that guides users through specific AI-powered tasks (e.g. chatbot)

WHEN

Use an AI application for predictable, human-supervised tasks where consistency matters

AI Agent (Agentic AI)

WHAT


An AI agent is an autonomous system that pursues goals independently using planning, tools and self-correction.

WHEN

Use an AI agent for tasks that require independent thinking, multi-step orchestration and on-the-fly decisions

What is an Agent?

An agent combines a generative AI model with tools, memory, and orchestration logic to solve complex problems independently.



Key Components

- **Model:** Used to reason over goals, determine the plan and generate a response
- **Tools:** Fetch data, perform actions or transactions by calling other APIs or services
- **Orchestration:** Maintain memory and state (including the approach used to plan), tools, data provided/fetched, etc

▶ A Governed, Two-Track Strategy

Successful enterprises run two parallel lanes to maximize velocity and differentiation.

Two Powerful Entry Points

Build your agents on Google Cloud

Benefit from integrated tooling without the burden of maintaining infrastructure.

Runtime

Sandbox

Sessions

Memory Bank

Gateway

Identity

Registry

Model Armor

Anomaly Detection

Custom Agents (GEAP 'Build' Lane)

Gemini Enterprise App (The 'Empower' Lane)

What

Sophisticated agentic workflows for complex, mission-critical scenarios

Platform where non-technical staff can securely find, use, and even build their own agents for broad productivity & self-service automation

Who

Software Engineering & Data Scientists (Pro-Code)

Non-technical users, e.g. Analysts, HR, Sales (No-Code/Low Code)

Sample Use-Cases

- Core operations (e.g., automating multi-step loan underwriting workflows)
- Customer experience (e.g., 'super agents' to modify orders & process refunds)

- Functional expert (e.g., "HR Benefits Advisor" or "IT Helpdesk Assistant")
- Role-Based Accelerators (e.g., "Sales Outreach Generator" for personalized email generation)

Values

Strategic differentiation, high ROI on core processes.

Scale, speed, reduced IT backlog, and employee empowerment.

■ The No-Code Path: Agent Designer


Empowering teams through no-code build. Helps teams move from AI strategy to a working prototype.

Agent Capabilities:

- **Build:** Create customizable custom agents
- **Orchestrate:** Manage complex workflows
- **Define & Instruct:** Control agent behavior
- **Enrich:** Connect data source integrations
- **Preview & Test:** Real-time execution testing


To help you:

- **Complete Repetitive Tasks:** Automate daily routines
- **Compare Documents:** Analyze side-by-side files
- **Compose Replies:** Generate instant email drafts
- **Generate Images:** Create stunning visual content
- **And More:** Adaptable to endless requirements



■ The No-Code Path: Prototyping & Testing

A strong prototype starts with a clear use case, the right context, and simple testing.



Instructional Design Agent Designer

Reliable Prompting Foundation

Writing clear, unambiguous instructions for repeatable results.

Key Principles

1. Clarity & Precision (Avoid Ambiguity)
2. Structured Inputs & Outputs
3. Defined Evaluation Metrics (Accuracy, Consistency)

Iterative Refinement

Continuous evaluation and feedback for optimal performance.

"You don't need to be a developer to build a powerful AI assistant."

4-Step Testing Continuum

- ✓ Validate accuracy against expected answers.
- ✓ Check whether responses use the right sources.
- ✓ Find edge cases or unclear instructions.
- ✓ Fix issues before expanding access.



Building Your First No-Code Prototype

Create the initial agent

1. Click the Overview button in the left menu.
2. Click the URL for the Gemini Enterprise webapp. The application will open in a new window.
3. In the application, click New agent in the left menu.
4. Using the prompt box, input a prompt like the following. The Create an agent component ensures that Gemini knows to create an agent from the prompt:
Create an agent named the Supply Chain Agent that is an expert at supply chain processes including planning materials, ordering supplies, and logistics for a marathon race or event. Include any subtasks and research tasks needed
5. Click the Submit paper airplane.

Purpose: Define the agent's purpose and goals.



Context: Connect relevant data and sources.



Instructions: Add clear instructions and guardrails.



Prototype: Preview, test, and refine the output.



The Hands-On Workshop

Building Your First No-Code Prototype

Modify the agent

Agents can be modified manually in the *Flow* view, or by using prompts.

1. In the prompt pane, ask Gemini to create subagents by prompting *add sub-agents as needed*.
2. You may need to provide an additional affirmation to create sub-agents.
3. The *Flow* view will update.
4. Click Create in the upper right to save the agent.

Prompt: plan the supplies needed for a 10,000 runner race in Las Vegas in April.

Share the agent

Agents that are created in Gemini Enterprise are automatically registered in Agent Registry.

Share agents in Gemini Enterprise

You can share agents within Gemini Enterprise, and from Agent Registry.

1. To share the agent in Gemini Enterprise, click Agents in the left menu.
2. Click the 3 ellipsis on the card representing your agent.
3. Add any user names, and click Done.

View agents in Agent Registry

1. Go to the [Agent Registry](#) page.
2. In the Runtime filter, click the filter, then select *Gemini Enterprise*.
3. The agent you created will be displayed.
4. Click the agent's name.
5. You can view information about the agent including the agent's identity. The identity can be granted additional or conditional access as needed.

Live Demonstration

Outreach Specialist Agent and Supply Chain Agent

- Moving beyond the front-end UI to inspect backend telemetry.
- 100% transparent logging natively within the Gemini Enterprise Agent Platform.
 - Diagnosing system timeouts and logic paths in real-time.

▶ The Pro-Code Path: PCIA Use Case

How Promevo Solves Problems with Custom Agents


Marketing's Problem

- Spends too much time manually sifting through industry news.
- Challenging to find current event-related blog topics.
- Difficult and slow to Identify available Subject Matter Experts (SMEs)

The Solution:

Promevo Content Intelligence Agent (PCIA) automates discovery and SME matching.

LLM: Gemini 3.0 Flash
Agent Framework: Google ADK
Deployment: Agent Runtime





New chat

Search

Starred

Library

Agents



Deep Research



Generate Client Meeting Sum...



Generate Pipeline Report



Idea Generation

Preview



NotebookLM



+ New agent

Chats

Latest news updates

HubSpot Pipeline Export

Professional Service Deals

HubSpot pipeline export

HubSpot Professional Service



Plus

Agents

Made by Google



Deep Research

Get in-depth answers grounded in web research.

By Google



Preview



Idea Generation

Generate and refine innovative solutions through a competitive, multi-agent...

By Google



NotebookLM

Quickly summarize and take notes for research with AI.

By Google

From your organization



PCIA

Promevo Content Intelligence Agent (PCIA). An Agentic AI solution serves as a Content...



sales-research-tool

Its primary function is to gather intelligence on target companies from both interna...

P PCIA

Google Cloud & Security

- Quantum frontiers may be closer than they appear
 - Summary: Google is accelerating its timeline for post-quantum cryptography migration to ensure long-term data security against emerging quantum computing capabilities.
 - Author: Sophie Schmieg
 - Tags: Google Research, Safety & Security
 - Source: [The Keyword](#)

- OpenAI shutting down Sora video generation app without explanation

- **Summary:** In a surprising move, OpenAI has announced the shutdown of its Sora AI video generation application. No specific reason for the closure has been provided yet.
- **Author:** Ben Schoon
- **Tags:** AI, News
- **Source:** [9to5Google](#)

Google Ads & Business Solutions

Generative AI may display inaccurate information, including about people, so double-check its responses.

Code Blueprint: The Orchestrator

```
root_agent = LlmAgent(
    name='orchestrator_agent',
    model=llm,
    instruction=
        "You are a helpful assistant that aids Promevo marketing team in creating timely and engaging content. " \
        "You will help the team by sifting through the latest news and updates in the tech industry, identifying relevant topics, \
        and suggesting content ideas. " \
        "Use the news agent to provide the latest news updates related to the user's query. " \
        "If the user wants to find SME experts in the company based on a specific topic, " \
        "use the SME finder agent to find the relevant SME experts and use the SME scheduler agent to check their \
        availability and schedule meetings with them. " \
        "The current date and time is: " + datetime.datetime.now().strftime("%Y-%m-%d %H:%M:%S"),
    sub_agents=[
        sme_finder,
        sme_scheduler,
        News_agent ] )
```

Trace Log: Orchestration & Delegation in Action

pcia_agent Service Configuration

Dashboard Traces Sessions **Playground** Memories Evaluation

The Playground requires the following permissions to show information in the Trace and Events tabs.

Project: agentspace-promo-internal

- cloudtrace.traces.get

To acquire access, contact your organization administrator.

Event 2 of 6

Event

```
id: "7980027761131520000"
invocationId: "e-7aa66999-9f6c-4d3f-9a11-88c1f9bd1146"
author: "orchestrator_agent"
content:
  role: "model"
  parts:
    0:
      functionCall:
        name: "transfer_to_agent"
        args:
          agent_name: "news_agent"
timestamp: 1774452028117
actions:
```

Playground Preview

SESSION ID 4590245129368895488

#1

#2

#3

#4

Type a Message...

Event 2 of 6

Author: `orchestrator_agent`

Action: `transfer_to_agent`

Args: { agent_name: "news_agent" }

The Orchestrator evaluates its plan and delegates the initial search task to the specialized News Agent.

Trace Log: The Secure Tool Call (Event 4 of 6)

pcia_agent Service Configuration

Dashboard Traces Sessions **Playground** Memories Evaluation

The Playground requires the following permissions to show information in the Trace and Events tabs.

Project: agentspace-promevo-internal

- cloudtrace.traces.get

To acquire access, contact your organization administrator.

Event 4 of 6

Event

```
id: "4737436029424762880"
invocationId: "e-7aa66999-9f6c-4d3f-9a11-88cf9bd1146"
author: "news_agent"
content:
  role: "model"
  parts:
    0:
      functionCall:
        name: "fetch_and_return_news_from_rss_feeds"
        args:
          date_range_days: 0
          target_date: "2026-03-25"
timestamp: 1774452030241
```

Playground Preview

SESSION ID 4590245129368895488

#2 transfer_to_agent

#3 ✓ transfer_to_agent

#4 fetch_and_return_news_from_rss_feeds

#5 ✓ fetch_and_return_news_from_rss_feeds

Type a Message...

Execution Context

Author: news_agent

Action: **FunctionCall**

Tool: fetch_and_return_news...

Args: target_date: "2026-03-25"

Security Boundary

The specialized News Agent executes its specific, isolated function.

It is restricted from accessing directories or calendars.

Trace Log: Grounded Generation (Event 5 of 6)

pcia_agent Service Configuration

Dashboard Traces Sessions **Playground** Memories Evaluation

The Playground requires the following permissions to show information in the Trace and Events tabs.

Project: agentspace-promevo-internal

- cloudtrace.traces.get

To acquire access, contact your organization administrator.

Event 5 of 6

Event

```
id: "7286473418516463616"
invocationId: "e-7aa66999-9f6c-4d3f-9a11-88c1f9bd1146"
author: "news_agent"
content:
  role: "user"
  parts:
    0:
      functionResponse:
        name: "fetch_and_return_news_from_rss_feeds"
        response:
          result: '{"message": "Found 9 articles for 2026-03-25.", "date_queried": "2026-03-25", "articles": [{"title": "Quantum frontiers may be closer than they appear", "link": "https://blog.google/innovation-and-ai/technology/safety-security/cryptography-"}'}
```

Playground Preview

SESSION ID 4590245129368895488

#3 ☒ transfer_to_agent

#4 ☐ fetch_and_return_news_from_rss_feeds

#5 ☒ fetch_and_return_news_from_rss_feeds

Here are the latest news updates from Google and the broader tech ecosystem relevant to Google Cloud Workspace and AI integration:

Type a Message...

Event Metadata

Author: news_agent

Action: **FunctionResponse**

Result: Structured JSON data from verified news sources.

"Grounded Generation: Factual, Structured, Verifiable"

The agent relies strictly on facts rather than hallucinating.

Trace Log: Synthesis & Security (Event 6 of 6)

← pcia_agent Service Configuration

Dashboard Traces Sessions **Playground** Memories Evaluation

The Playground requires the following permissions to show information in the Trace and Events tabs.

Project: agentspace-promevo-internal

- cloudtrace.traces.get

To acquire access, contact your organization administrator.

Event 6 of 6

Event

```
id: "2575708208286924800"
invocationId: "e-7aa66999-9f6c-4d3f-9a11-88c1f9bd1146"
author: "news_agent"
content:
  role: "model"
  parts:
    0:
      text: "Here are the latest news updates from Google and the broader tech ecosystem for March 25, 2026, with a focus on topics relevant to Google Cloud, Workspace, and AI integration. ### **Top Security & Infrastructure News** * * * [Quantum frontiers may be closer than they appear] (https://blog.google/innovation-and-ai/technology/safety-security/cryptography-
```

Playground Preview

SESSION ID 4590245129368895488

+ New Session

#4 fetch_and_return_news_from_rss_feeds

#5 ✓ fetch_and_return_news_from_rss_feeds

Here are the latest news updates from Google and the broader tech ecosystem for March 25, 2026, with a focus on topics relevant to Google Cloud, Workspace, and AI integration.

Top Security & Infrastructure News

- Quantum frontiers may be closer than they appear

Type a Message...

Trace Log: Zero-Trust Security

●●● GEAP PLAYGROUND | TRACE_LOG_EVENT_06

[2026-06-14 23:46:44] INFO: Initializing Event 6 of 6...

Author: news_agent

Action: Final text synthesis summarizing the tech news.

> Attempting unauthorized action: access_restricted_calendar

ERROR: Permission Denied

[SUCCESS] Agent synthesis complete within IAM boundaries.

The news summary has been generated based on authorized public web sources only.

Security Guardrail

Identity and Access
Management (IAM)
Enforcement


The agent strictly inherits
the human user's IAM
permissions.

*"Security check ensures no
escalation or lateral
movement beyond the user's
defined scope."*

Zero-Trust Verified

■ The Agent Lifecycle

A continuous, iterative framework for building agents at scale



Session 3 Completed: The Prototype Built

Webinar Series Tracker

3 of 4 Sessions Complete



Session 1: The Art of Possible

Outcome alignment & Brain/Body strategy.



Session 2: The Foundation

Data connectors & security boundaries established.



Session 3: The Workshop

Completed today. Live agent builder walkthrough.



Session 4: The Assembly Line

Up next: July 16. Bridging from prototype to enterprise scale.

Take Your Prototype out of the Sandbox

You have seen how to build a no-code agent. The next step is to distribute it safely to a focused departmental user group to evaluate performance.

Contact us to discuss our Gemini Enterprise Accelerators.

promevo.com/contact-us



Session 4: From Ideation to Prototype

Today's Agenda

Webinar Series Schedule



Webinar 1

July 7 • 1:00 PM

THE ART OF POSSIBLE

Define business outcomes and identify where AI can create measurable impact.



Webinar 2

July 9 • 1:00 PM

THE FOUNDATION

Ground agents in trusted enterprise data architectures for reliable outputs.



Webinar 3

July 14 • 1:00 PM

THE WORKSHOP

Build your first functional, no-code virtual agent prototype step-by-step.



Webinar 4

July 16 • 1:00 PM

THE ASSEMBLY LINE

Scale your virtual agent for robust, enterprise-grade production environments.

- Part 1: Assembly Line Intro
- Part 2: Meta-Agents
- Part 3: Scaling Considerations
- Part 4: Wrap up.



Reviewing the Journey: From Ideation to Prototype

Before scaling an agent for production, it's important to revisit how we got here.



Defined the Business Outcome

Identified the problem the agent needs to solve.



Grounded the Agent in Data

Connected trusted enterprise sources and permissions.



Built the Prototype

Created a no-code agent with clear instructions and business context.



Tested and Refined

Validated outputs before moving toward production.

A strong production agent starts with a clearly defined, well-tested prototype.



Advanced Agent Builder Techniques for Real-World Impact

Once the prototype is working, advanced techniques help make the agent more useful, reliable, and scalable.

Better Instructions

Refine how the agent understands its role, task, and response format.

Smarter Grounding

Connect the agent to the right sources for more relevant answers.



Workflow Actions

Move beyond answers into drafts, updates, summaries, or next steps.

Quality Feedback

Use testing and review to improve performance over time.

Gemini Enterprise Agent Library



✓ 100+ Industry Use Cases

✓ System-Specific Filtering


✓ Workflow Visualization

✓ Design Specs



The Assembly Line


Scaling your Agent for Production Transforming prototypes to reliable enterprise assets





The Meta-Agent: Strategic Discovery

Building an agent to discover *new* agentic use cases




Enterprise Data
Connectors



Automating End-to-End Workflows with the “Generic Agent”

The “Generic Agent” Demo: Multi-Tool Execution from a Single Directive





Scaling Consideration: Reliability

A production agent needs to deliver consistent, trustworthy results across real users, real workflows, and real edge cases.

Monitoring

Track usage, performance, and recurring issues over time.

Testing

Validate responses against known questions, scenarios, and expected outcomes.

Evaluation Metrics

Measure accuracy, source quality, consistency, and user confidence.



Scaling Consideration: Rollout Strategy

A successful prototype needs a rollout strategy before it can support broader teams, workflows, and user groups.

Infrastructure Readiness

Prepare the systems, permissions, and connected data sources needed to support more users.

Change Management

Give teams the training, guidance, and expectations they need to adopt the agent confidently.

Rollout Strategy

Start with focused user groups, learn from feedback, and expand access in phases.



Scaling Consideration: ROI Definition

Estimating Annualized and Per-Run ROI Based on Rate, Time Savings, and Frequency

Average Hourly Rate

Establish the average hourly rate of the employees who would typically perform the process or task manually.

Time Saved Per Run

Estimate the hours saved per individual run (e.g., 1.5 hours) by delegating the process to the no-code agent.

Execution Frequency

Determine how often the agent runs (daily, monthly, quarterly) to calculate per-run metrics and roll them up into annualized ROI.



Scaling Consideration: Calculating Workspace ROI

Measuring Impact: Time Saved and Productivity Gains by Eliminating Context-Switching

Time Saved

Reduce the time spent searching for information, switching tools, and recreating work.

Workflow Efficiency

Move from insight to action inside Gmail, Docs, Sheets, and Chat.

Adoption Potential

Make AI easier to use by embedding it into familiar daily workflows.



Conclusion & Q&A

From Vision to Production: Recapping the prototype-to-production journey

Over the course of this series, we walked through how to move from a business vision to a scalable, production-ready agent.



Define the Vision

Start with a measurable business outcome.



Build the Foundation

Ground the agent in trusted enterprise data.



Create the Prototype

Use no-code tools to build, test, and refine the first version.



Scale for Production

Add reliability, governance, security, and ROI measurement.



Webinar Series Completed: The Road to Scale

Webinar Series Tracker

Series Fully Complete



Session 1: The Art of Possible

Operationalized AI outcome framework aligned.



Session 2: The Foundation

Avoided the Data Swamp with secure grounding.



Session 3: The Workshop

Completed hands-on sandbox agent build.



Session 4: The Assembly Line

Completed today. Production architecture masterclass.

Securely Hardening Your Multi-Agent Fleet

Now that you know the 12 critical infrastructure layers required for real enterprise production, let's turn your blueprints into high-ROI operational reality.

Contact us to discuss our Gemini Enterprise Accelerators.

promevo.com/contact-us



Thank You!