


AI Agents in Production

War Stories from a DevOps Engineer



Do Cao Hieu

- AI Agents in Production
 - War Stories from a DevOps Engineer
- PART 1: FOUNDATIONS
- Chapter 1: The Shift from Prompts to Agents
 - The Evolution of Operational Tooling
 - What Makes an “Agent” Different from a “Chatbot”
 - The Agent Loop: Plan → Execute → Observe → Adapt
 - Why DevOps is the Perfect Domain for AI Agents
 - Case Study: A Simple Prompt vs an Agent Solving the Same Problem
- Chapter 2: Anatomy of an AI Agent
 - Core Components: LLM Brain, Tools, Memory, Orchestration
 - The Context Window: Your Agent’s Working Memory
 - Tool Use: Giving Agents Hands to Interact with the World
 - Memory Systems: Short-term, Long-term, Episodic
 - Agent Architectures: Single Agent vs Multi-Agent Systems
- Chapter 3: Your First Production Agent
 - Choosing Your Stack: Frameworks, Models, Infrastructure
 - The “Hello World” of Agents: A Deployment Assistant
 - Environment Setup: API Keys, Permissions, Sandboxing
 - Running Your First Autonomous Task
 - What Could Go Wrong (Spoiler: A Lot)
- Part 2: The Hard Lessons
- Chapter 4: When Agents Go Wrong
 - War Story: The Runaway Cron Job That Burned \$200 in Tokens
 - War Story: The Agent That Deleted the Wrong Files
 - War Story: The Infinite Loop of Self-Correction
 - Common Failure Modes and How to Recognize Them
 - Building Kill Switches and Circuit Breakers
- Chapter 5: Context Engineering
 - The Context Window Is Not Infinite (And It’s Expensive)
 - Token Economics: Measuring and Optimizing Usage
 - Context Compression Techniques
 - When to Summarize vs When to Forget
 - War Story: The Agent That Forgot What It Was Doing
- Chapter 6: Memory Systems That Actually Work
 - Why Chat History Isn’t Enough
 - Designing Persistent Memory: Files, Databases, Vector Stores
 - The MEMORY.md Pattern: Simple but Effective
 - Neural Memory and Semantic Search
 - War Story: Building Memory That Survives Restarts
- Chapter 7: Multi-Agent Orchestration
 - When One Agent Isn’t Enough
 - The Orchestrator Pattern
 - Communication Between Agents: Shared Context, Message Passing
 - Parallel vs Sequential Agent Execution

- War Story: Coordinating a Team of 4 Agents on a Complex Task
- Part 2 Summary: The Hard Lessons
- PART 3: PRODUCTION PATTERNS
 - Chapter 8: Security and Access Control
 - Chapter 9: Cost Management
 - Chapter 10: Reliability and Observability
 - Chapter 11: Human-in-the-Loop Patterns
- Part 4: Real-World Applications
- Chapter 12: Infrastructure Automation
 - Agents That Provision and Manage Infrastructure
 - Self-Healing Systems: Detecting and Fixing Issues
 - Configuration Drift Detection and Correction
 - Capacity Planning with AI Assistance
 - Case Study: An Agent That Handles Routine Maintenance
- Chapter 13: CI/CD and Deployment
 - Agents in the Deployment Pipeline
 - Automated Rollbacks Based on Metrics
 - PR Review Assistance
 - Case Study: An Agent That Fixes Failing Tests
- Chapter 14: Observability and Incident Response
 - Log Analysis at Scale
 - Automated Runbook Execution
 - Incident Summarization and RCA Drafting
 - On-Call Assistance
 - Case Study: An Agent as Your First Responder
 - Part 4 Summary
- PART 5: THE FUTURE
- Chapter 15: Building Your Agent Team
 - From Single Agent to Agent Organization
 - Specialization: Different Agents for Different Tasks
 - Communication Between Agents
 - The War Story: Coordinating a Team of 4 Agents
 - Governance: Who Watches the Watchers
 - The Hybrid Team: Humans and Agents Working Together
 - Building Your Agent Team: Practical Steps
- Chapter 16: What's Next
 - The Trajectory: More Capable, More Autonomous
 - Emerging Patterns: MCP, Tool Ecosystems, Agent Marketplaces
 - Preparing Your Infrastructure for Agent Adoption
 - The Skills That Matter in an Agent-First World
 - A Practical Adoption Roadmap
 - The War Story: Looking Back
 - Your Journey Starts Now
- APPENDICES
- Appendix A: Tool and Framework Reference
 - Agent Frameworks Compared

- Framework Decision Matrix
- Model Selection Guide
- Essential CLI Tools for Agent Development
- Cost Estimation Quick Reference
- Appendix B: Prompt Templates
 - The Infrastructure Agent Base Template
 - The Deployment Agent Template
 - The Incident Response Agent Template
 - The Code Review Agent Template
 - Guidelines
 - Multi-Step Task Template
 - Execution Rules
 - State Update Format
 - Common Patterns

AI Agents in Production

War Stories from a DevOps Engineer

A practical guide to running AI agents in production environments.

From the lessons learned building and operating autonomous AI systems for infrastructure automation, CI/CD, incident response, and beyond.

PART 1: FOUNDATIONS

Chapter 1: The Shift from Prompts to Agents

The Evolution of Operational Tooling

Every decade or so, the way we manage infrastructure undergoes a fundamental shift. In the 1990s, we typed commands into terminals and hoped we remembered the right flags. In the 2000s, we wrapped those commands into bash scripts—fragile, but at least repeatable. The 2010s brought us Infrastructure as Code: Terraform, Ansible, Puppet. We stopped describing *how* and started describing *what*.

Then came the prompts.

Starting around 2023, a new pattern emerged. Instead of writing scripts, engineers started asking AI to write them. Instead of debugging configuration files line by line, we pasted error messages into ChatGPT and got solutions back. This was revolutionary—but it was also just the beginning.

The progression looks like this:

```
CLI commands → Shell scripts → IaC → Prompts → Agents
```

Each step represented a leap in abstraction. CLI commands required knowing exact syntax. Scripts captured that syntax in reusable form. IaC let us declare desired state. Prompts let us express intent in natural language.

But agents? Agents close the loop entirely.

What Makes an “Agent” Different from a “Chatbot”

I’ve lost count of how many times someone has asked me: “Isn’t this just ChatGPT with extra steps?”

Fair question. Let me be precise about the distinction.

A **chatbot** responds to inputs. You ask a question, you get an answer. The conversation might span multiple turns, but fundamentally, the chatbot is reactive. It waits for you. It answers you. It waits again.

An **agent** pursues goals. You give it an objective, and it figures out how to achieve it. It plans. It takes actions. It observes the results of those actions. It adapts. Crucially, it does this *without waiting for you at every step*.

Here's a concrete example. Imagine you want to deploy a new version of a service:

Chatbot interaction:

```
You: How do I deploy version 1.2.3 of the payment service?
Bot: You can use kubectl set image deployment/payment
     payment=payment:1.2.3
You: I ran that. It says ImagePullBackOff.
Bot: That error means Kubernetes can't pull the image. Check if
     the image exists.
You: How do I check that?
Bot: Run docker pull payment:1.2.3 to verify the image exists...
```

This goes on. You're the loop. You execute commands. You report results. You ask follow-up questions. The bot just answers.

Agent interaction:

```
You: Deploy version 1.2.3 of the payment service to staging.

Agent: [Planning] I need to: 1) verify the image exists, 2)
      update the deployment,
      3) wait for rollout, 4) verify health
Agent: [Executing] docker pull payment:1.2.3 ✓ Image exists
Agent: [Executing] kubectl set image deployment/payment
      payment=payment:1.2.3 ✓
Agent: [Observing] Rollout status: 1/3 pods updated
Agent: [Observing] Rollout status: 3/3 pods updated
Agent: [Executing] curl http://payment-staging/health → 200 OK

Done. Deployment successful. All health checks passing.
```

See the difference? The agent drives the conversation. It takes actions autonomously. It observes results and proceeds without asking you what to do next.

This is the fundamental shift: from **human-in-the-loop** to **human-on-the-loop**.

The Agent Loop: Plan → Execute → Observe → Adapt

Every agent, regardless of how sophisticated, follows a basic loop:

```
while not done:
    plan = think_about_what_to_do_next(goal, observations)
    action = choose_action(plan)
    result = execute_action(action)
    observations = observe_result(result)
    if goal_achieved(observations):
        done = True
    elif stuck_or_failed(observations):
        adapt(plan) # or give up gracefully
```

This is sometimes called the ReAct pattern (Reasoning + Acting), or simply the agent loop. Let's break it down:

Plan: The agent uses its LLM brain to decide what to do next. This isn't a fixed script—it's dynamic reasoning based on current context.

Execute: The agent takes action. This might mean running a command, calling an API, writing a file, or sending a message. This is where tools come in—we'll cover those in Chapter 2.

Observe: The agent sees what happened. Did the command succeed? What did it output? Are there errors? This observation becomes part of the context for the next planning step.

Adapt: If something unexpected happens, the agent adjusts. Maybe the first approach didn't work. Maybe new information changed the plan. Good agents recover gracefully; bad ones spin in infinite loops.

Here's what this looks like in practice:

```
def agent_loop(goal: str, max_iterations: int = 10):
    context = f"Goal: {goal}\n"

    for i in range(max_iterations):
        # Plan
        plan = llm.think(context + "What should I do next?")

        if plan.says_done():
            return context

        # Execute
        action = plan.next_action()
        result = execute(action)

        # Observe
        context += f"\nAction: {action}\nResult: {result}\n"

        # Adapt (implicit in next iteration's planning)

    return context + "\nMax iterations reached."
```

This loop is deceptively simple. The complexity—and the danger—lies in what happens inside each step. What tools does the agent have access to? How does it decide when it's done? What happens when it makes a mistake?

We'll spend most of this book answering those questions.

Why DevOps is the Perfect Domain for AI Agents

Not every domain is equally suited to agents. Some require too much nuanced judgment. Some have consequences too severe for autonomous action. Some lack the tooling for programmatic interaction.

DevOps hits a sweet spot.

Rich tooling ecosystem: Our entire world is built on CLIs and APIs. Kubernetes, Docker, AWS, Terraform, Ansible—everything is designed for programmatic interaction. Agents can plug into existing tools without needing custom interfaces.

Observable outcomes: When we deploy something, we know if it worked. We have metrics, logs, health checks. Agents can verify their actions objectively.

Recoverable failures: Most DevOps mistakes can be fixed. Rollback a deployment. Restore a backup. Terminate a runaway process. This is important—agents will make mistakes, and we need to survive them.

Repetitive patterns: Much of operations work follows patterns. “Check if X, then do Y, unless Z.” “When the alert fires, run through this checklist.” These patterns are perfect for automation, and agents can learn them.

Clear intent: “Deploy to staging.” “Scale up the database.” “Fix this failing test.” DevOps tasks often have unambiguous success criteria.

This doesn’t mean deploying agents in DevOps is easy—it’s not, and we’ll see plenty of war stories proving that. But the domain is tractable. The tools exist. The feedback loops are tight.

Case Study: A Simple Prompt vs an Agent Solving the Same Problem

Let me show you a real scenario that illustrates the difference.

The situation: At 2 AM, an alert fires. API latency has spiked above SLA. Someone needs to investigate and fix it.

The Prompt Approach

An engineer wakes up, opens their laptop, and starts prompting:

```
Me: I'm getting high API latency on the payments service. How should I debug this?
```

```
LLM: Here are some steps to debug high API latency:
```

1. Check if the issue is widespread or isolated to specific endpoints
2. Look at recent deployments that might have caused regression
3. Check database query performance
4. Look at downstream service dependencies
5. Check for resource exhaustion (CPU, memory, connections)

```
Me: How do I check recent deployments?
```

```
LLM: You can run: kubectl rollout history deployment/payments  
This will show recent deployment revisions...
```

```
Me: *runs command, pastes output*
```

```
I see a deployment happened 30 minutes ago. How do I check if that caused it?
```

```
LLM: You can compare latency metrics before and after the deployment...
```

This continues for 20-30 minutes. The engineer is doing all the work. The LLM is just an advisor.

The Agent Approach

The agent wakes up instead of the engineer (or at least, before escalating to them):

```
[02:03:14] Alert received: API latency > 500ms for payments
service
[02:03:15] Goal: Investigate and resolve latency spike

[02:03:16] Checking current latency metrics...
→ P99 latency: 847ms (threshold: 500ms)
→ Started: 01:35:00

[02:03:18] Checking recent deployments...
→ Deployment at 01:32:00: payments:v2.3.4 → v2.3.5

[02:03:19] Correlation: Latency spike started 3 min after
deployment
→ High probability: deployment is the cause

[02:03:20] Checking deployment changes...
→ Commit: "Add new payment provider integration"
→ Changed files: 12
→ New dependencies: 2

[02:03:22] Checking resource metrics post-deployment...
→ CPU: 23% → 67%
→ Memory: stable
→ DB connections: stable

[02:03:24] Analysis: CPU increase correlates with latency
→ Likely cause: inefficient code in new integration

[02:03:25] Options:
1. Rollback deployment (fast, safe)
2. Scale horizontally (might help, unclear)
3. Wait for analysis (risk ongoing SLA breach)

[02:03:26] Decision: Rollback is safest given SLA risk
→ Executing: kubectl rollout undo deployment/payments

[02:03:35] Rollback complete. Monitoring latency...
[02:03:50] P99 latency: 245ms ✓ Below threshold
[02:04:00] Latency stable for 10 seconds. Issue resolved.

[02:04:01] Creating incident report...
[02:04:05] Notifying on-call engineer with summary...

Incident resolved. Total time: 51 seconds.
Engineer notification: "Alert auto-resolved via rollback.
Details: [link]. Please review v2.3.5 for performance
regression."
```

The agent did in 51 seconds what took the human 20-30 minutes of back-and-forth. More importantly, the engineer got to stay asleep for the simple case—while still being informed and able to review.

This is the promise of agents in DevOps. Not replacing human judgment, but handling the routine so humans can focus on the exceptional.

The Reality Check

Now, I need to be honest with you. The agent example above is the *ideal* case. In reality:

- The agent might have rolled back something that wasn't the cause
- The agent might have gotten stuck if the rollback failed
- The agent might have missed a more complex root cause
- The agent might have cost \$0.50 in API calls for a problem that would have resolved itself

Agents are not magic. They're a tool—a powerful one, but one that requires careful design, guardrails, and human oversight. The rest of this book is about how to build agents that work in the real world, with all its messiness and edge cases.

Let's start by understanding what's inside an agent.

Chapter 2: Anatomy of an AI Agent

Core Components: LLM Brain, Tools, Memory, Orchestration

If you cracked open an AI agent and looked inside, you'd find four fundamental components:

The LLM Brain: This is the reasoning engine. It takes in context—the goal, the current situation, the history of what's happened—and decides what to do next. The LLM doesn't actually *do* anything except think and generate text. But that text can be instructions for actions.

Tools: These are the agent's hands. While the LLM can only think, tools let it act. A tool might be a bash command executor, an API client, a file writer, or anything else that changes the world. When the LLM decides to "run `kubectl get pods`", a tool actually executes that command.

Memory: This is how the agent maintains context across time. Without memory, every interaction would start from zero. Memory lets the agent remember what it was doing, what it tried, what worked, and what didn't.

Orchestration: This is the glue that ties everything together. The orchestration layer manages the agent loop, handles tool execution, maintains memory, and enforces guardrails. It's the runtime that makes an agent an agent rather than just an LLM.

Let's examine each component in detail.

The Context Window: Your Agent's Working Memory

The context window is perhaps the most important constraint to understand. It's the LLM's short-term memory—everything the model can "see" at once.

Different models have different context limits: - GPT-3.5: 16K tokens - GPT-4o: 128K tokens - Claude 3: 200K tokens - Gemini 1.5: 1M+ tokens

But here's what nobody tells you in the marketing materials: just because a model *can* handle 200K tokens doesn't mean it *should*.

```
# Token estimation (rough rule of thumb)
def estimate_tokens(text: str) -> int:
    return len(text) // 4 # ~4 characters per token for English

# Real cost calculation
def cost_per_request(context_size: int, model: str) -> float:
    rates = {
        "gpt-4o": {"input": 0.005, "output": 0.015}, # per 1K tokens
        "claude-3-opus": {"input": 0.015, "output": 0.075},
    }
    return (context_size / 1000) * rates[model]["input"]
```

In production, context window management is everything. Fill it up, and you pay through the nose. Let it overflow, and the agent loses important information. We'll cover strategies in Chapter 5.

Tool Use: Giving Agents Hands to Interact with the World

Tools are how agents affect reality. A tool is simply a function that the LLM can invoke through a structured format.

Here's a basic tool definition:

```
tools = [
  {
    "name": "run_command",
    "description": "Execute a shell command and return the output",
    "parameters": {
      "type": "object",
      "properties": {
        "command": {
          "type": "string",
          "description": "The shell command to execute"
        }
      }
    },
    "working_directory": {
      "type": "string",
      "description": "Directory to run the command in",
      "default": "."
    }
  },
  {
    "required": ["command"]
  }
],
{
  "name": "read_file",
  "description": "Read the contents of a file",
  "parameters": {
    "type": "object",
    "properties": {
      "path": {
        "type": "string",
        "description": "Path to the file to read"
      }
    }
  },
  "required": ["path"]
}
]
```

When the LLM decides to use a tool, it generates a structured response:

```
{
  "tool": "run_command",
  "parameters": {
    "command": "kubectl get pods -n production",
    "working_directory": "/home/user"
  }
}
```

Your orchestration layer intercepts this, executes the actual command, and feeds the result back to the LLM:

```

def handle_tool_call(tool_call: dict) -> str:
if tool_call["tool"] == "run_command":
    result = subprocess.run(
        tool_call["parameters"]["command"],
        shell=True,
        capture_output=True,
        cwd=tool_call["parameters"].get("working_directory", ".")
    )
    return result.stdout.decode() + result.stderr.decode()
elif tool_call["tool"] == "read_file":
    with open(tool_call["parameters"]["path"]) as f:
        return f.read()

```

The critical insight: **tool design is security design**. Every tool you give an agent is a capability it can use—or misuse. We'll cover this extensively in Chapter 8.

Common tool categories in DevOps agents:

Category	Examples
Shell	run_command, run_script
Files	read_file, write_file, list_directory
Network	http_request, curl
Kubernetes	kubectl, helm
Cloud	aws, gcloud, az
Git	git_status, git_commit, git_push
Communication	send_slack, send_email

Memory Systems: Short-term, Long-term, Episodic

Agents need to remember things at different timescales.

Short-term memory is the conversation itself. It's what the agent has seen in this session. This lives in the context window and is naturally limited by token count.

```

class ShortTermMemory:
def __init__(self, max_tokens: int = 8000):
    self.messages = []
    self.max_tokens = max_tokens

def add(self, role: str, content: str):
    self.messages.append({"role": role, "content": content})
    self._trim_to_fit()

def _trim_to_fit(self):
    while self._token_count() > self.max_tokens:
        # Remove oldest non-system message
        for i, msg in enumerate(self.messages):
            if msg["role"] != "system":
                self.messages.pop(i)
                break

```

Long-term memory persists across sessions. This is where the agent stores information it might need later—project context, user preferences, learned patterns.

```
class LongTermMemory:
    def __init__(self, storage_path: str = "memory/"):
        self.storage_path = storage_path

    def remember(self, key: str, content: str, metadata: dict = None):
        entry = {
            "content": content,
            "timestamp": datetime.now().isoformat(),
            "metadata": metadata or {}
        }
        path = f"{self.storage_path}/{key}.json"
        with open(path, "w") as f:
            json.dump(entry, f)

    def recall(self, key: str) -> str:
        path = f"{self.storage_path}/{key}.json"
        if os.path.exists(path):
            with open(path) as f:
                return json.load(f)["content"]
        return None
```

Episodic memory captures sequences of events—what the agent did in the past, what worked, what failed. This is crucial for learning from experience.

```
class EpisodicMemory:
    def __init__(self, db_path: str = "episodes.db"):
        self.conn = sqlite3.connect(db_path)
        self._init_schema()

    def record_episode(self, task: str, actions: list, outcome: str,
                      success: bool):
        self.conn.execute("""
INSERT INTO episodes (task, actions, outcome, success,
timestamp)
VALUES (?, ?, ?, ?, ?)
""", (task, json.dumps(actions), outcome, success,
datetime.now().isoformat()))
        self.conn.commit()

    def recall_similar(self, task: str, limit: int = 5) -> list:
        # In production, use vector similarity search
        return self.conn.execute("""
SELECT task, actions, outcome, success
FROM episodes
WHERE success = 1
ORDER BY timestamp DESC
LIMIT ?
""", (limit,)).fetchall()
```

The `MEMORY.md` pattern is a pragmatic approach I've used successfully: a single Markdown file that the agent reads at the start of each session and updates with important information. It's simple but effective:

```
# Agent Memory

## Project Context
- Main application: payment-service
- Deployment: Kubernetes on AWS EKS
- Monitoring: Datadog
```

```
## Known Issues
- Redis connection pool exhausts under load (workaround: restart)
- The staging database is 3 versions behind production

## User Preferences
- Prefer kubectl over Lens
- Always ask before deleting anything
- Notify via Slack, not email

## Recent Actions
- 2024-01-15: Deployed v2.3.4 to production
- 2024-01-14: Increased replica count to 5
```

Agent Architectures: Single Agent vs Multi-Agent Systems

Most agents start simple: one agent, one task. But as tasks grow complex, single agents struggle.

Single Agent Architecture:

```
User → Agent → Tools → Results
```

Pros: Simple, easy to debug, clear accountability. Cons: Limited by single context window, can't parallelize.

Multi-Agent Architecture:

```
User → Orchestrator Agent
├─ Specialist Agent A (e.g., Kubernetes)
├─ Specialist Agent B (e.g., Monitoring)
└─ Specialist Agent C (e.g., Documentation)
```

Pros: Specialization, parallel execution, larger total context. Cons: Complex coordination, inter-agent communication overhead.

A common pattern is the **orchestrator/worker** model:

```
class OrchestratorAgent:
def __init__(self):
    self.workers = {
        "k8s": KubernetesAgent(),
        "monitoring": MonitoringAgent(),
        "git": GitAgent()
    }

def execute(self, task: str):
    # Plan: decompose task
    subtasks = self.decompose(task)

    results = {}
    for subtask in subtasks:
        worker = self.route(subtask)
        results[subtask.id] = worker.execute(subtask)

    # Synthesize: combine results
    return self.synthesize(results)

def route(self, subtask) -> Agent:
```

```
# Determine which specialist should handle this
if "pod" in subtask.description or "deploy" in
subtask.description:
return self.workers["k8s"]
elif "metric" in subtask.description or "alert" in
subtask.description:
return self.workers["monitoring"]
# ...
```

We'll dive deep into multi-agent patterns in Chapter 7.

Chapter 3: Your First Production Agent

Choosing Your Stack: Frameworks, Models, Infrastructure

Before writing any code, you need to make some choices. Let me give you the lay of the land as of early 2025.

Agent Frameworks:

Framework	Pros	Cons
LangChain	Huge ecosystem, lots of examples	Complex, sometimes over-engineered
AutoGPT	Fully autonomous	Hard to control, expensive
CrewAI	Good multi-agent support	Newer, less battle-tested
Claude Code	Excellent tool use, safe defaults	Anthropic-specific
Custom	Full control	More work

For production DevOps agents, I recommend starting simple—either a minimal custom framework or Claude Code if you’re using Anthropic’s models. The heavy frameworks add complexity you don’t need initially.

Models:

Model	Best For	Cost	Context
GPT-4o	General tasks, good tool use	128K <i>Claude3Opus</i> <i>Complexreasoning,</i> <i>safety</i> \$	200K
Claude 3 Sonnet	Good balance, fast	\$\$	200K
GPT-4o-mini	Simple tasks, cost-sensitive	\$	128K

For DevOps, Claude models have an edge due to better instruction-following and more conservative defaults—important when your agent has shell access.

Infrastructure:

Your agent needs to run somewhere. Options: - **Local development:** Your laptop, for testing - **Server/VM:** A dedicated machine for production agents - **Container:** Docker with restricted capabilities - **Serverless:** Lambda/Cloud Functions for triggered agents

My recommendation: start with a dedicated VM or container with explicit resource limits. Serverless gets complicated with long-running agent loops.

The “Hello World” of Agents: A Deployment Assistant

Let’s build something real. A deployment assistant that can: 1. Check the current deployment status 2. Deploy new versions 3. Verify health after deployment 4. Roll back if needed

Here’s a minimal implementation:

```
#!/usr/bin/env python3
"""
deploy_agent.py - A simple deployment assistant agent
"""

import os
import json
import subprocess
from anthropic import Anthropic

client = Anthropic()

SYSTEM_PROMPT = """You are a deployment assistant for a
Kubernetes-based application.
You have access to tools to check deployments, deploy new
versions, and verify health.

IMPORTANT RULES:
1. Always check current status before making changes
2. Always verify health after deployment
3. If health check fails, offer to rollback
4. Never deploy to production without explicit confirmation
5. Be concise in your responses

Current environment: staging
Application: payment-service
Namespace: payments
"""

TOOLS = [
{
    "name": "kubectl",
    "description": "Run a kubectl command. Returns stdout and
stderr.",
    "input_schema": {
        "type": "object",
        "properties": {
            "args": {
                "type": "string",
                "description": "Arguments to pass to kubectl (e.g., 'get
pods -n payments')"
```

```

        "description": "The URL to check (e.g.,
        'http://localhost:8080/health')"
    }
},
"required": ["url"]
}
}
]

def execute_tool(name: str, input: dict) -> str:
if name == "kubectl":
    cmd = f"kubectl {input['args']}"
    result = subprocess.run(
    cmd, shell=True, capture_output=True, text=True, timeout=30
    )
    return f"stdout:\n{result.stdout}\nstderr:\n{result.
    stderr}\nreturn_code: {result.returncode}"

elif name == "health_check":
    cmd = f"curl -sf {input['url']} -o /dev/null && echo 'healthy'
    || echo 'unhealthy'"
    result = subprocess.run(cmd, shell=True, capture_output=True,
    text=True, timeout=10)
    return result.stdout.strip()

return f"Unknown tool: {name}"

def run_agent(user_input: str, max_turns: int = 10):
messages = [{"role": "user", "content": user_input}]

for turn in range(max_turns):
    response = client.messages.create(
    model="claude-sonnet-4-20250514",
    max_tokens=4096,
    system=SYSTEM_PROMPT,
    tools=TOOLS,
    messages=messages
    )

    # Check if the agent is done
    if response.stop_reason == "end_turn":
    # Extract text response
    for block in response.content:
    if hasattr(block, "text"):
        return block.text

    # Handle tool use
    if response.stop_reason == "tool_use":
    # Add assistant's response to messages
    messages.append({"role": "assistant", "content":
    response.content})

    # Execute each tool call
    tool_results = []
    for block in response.content:
    if block.type == "tool_use":
        print(f"🔧 Executing: {block.name}({block.input})")
        result = execute_tool(block.name, block.input)
        print(f"📄 Result: {result[:200]}...")
        tool_results.append({
            "type": "tool_result",
            "tool_use_id": block.id,
            "content": result
        })

    messages.append({"role": "user", "content": tool_results})

return "Max turns reached without completion"

if __name__ == "__main__":

```

```

print("🤖 Deployment Assistant Ready")
print("Type 'quit' to exit\n")

while True:
    user_input = input("You: ").strip()
    if user_input.lower() == "quit":
        break

    result = run_agent(user_input)
    print(f"\nAgent: {result}\n")

```

That's about 100 lines of code for a functional agent. Let's test it:

```

$ python deploy_agent.py
🤖 Deployment Assistant Ready
Type 'quit' to exit

You: What's the current deployment status?

🔧 Executing: kubectl({'args': 'get deployment payment-service -n
payments'})
📋 Result: NAME                READY  UP-TO-DATE  AVAILABLE  AGE
payment-service    3/3    3           3          5d...

Agent: The payment-service deployment is healthy:
- 3/3 replicas ready
- All pods up-to-date and available
- Running for 5 days

You: Deploy version 1.2.4

🔧 Executing: kubectl({'args': 'set image
deployment/payment-service
payment-service=payment-service:1.2.4 -n payments'})
📋 Result: deployment.apps/payment-service image updated...
🔧 Executing: kubectl({'args': 'rollout status
deployment/payment-service -n payments --timeout=60s'})
📋 Result: deployment "payment-service" successfully rolled out...
🔧 Executing: health_check({'url':
'http://payment-service.payments.svc/health'})
📋 Result: healthy...

Agent: Successfully deployed payment-service:1.2.4
- Rollout completed
- Health check: passing
- All 3 replicas running the new version

```

Environment Setup: API Keys, Permissions, Sandboxing

Before your agent touches production, you need proper setup.

API Keys: Store them securely, never in code.

```

# Good: Environment variables
export ANTHROPIC_API_KEY="sk-..."

# Better: Secret manager
export ANTHROPIC_API_KEY=$(vault read -field=key
secret/anthropic)

# Best: Automatic injection via service account
# (The agent never sees the key)

```

Permissions: Principle of least privilege. Your agent should have exactly the permissions it needs, no more.

```
# kubernetes-agent-rbac.yaml
apiVersion: rbac.authorization.k8s.io/v1
kind: Role
metadata:
  name: deployment-agent
  namespace: payments
rules:
  - apiGroups: ["apps"]
resources: ["deployments"]
verbs: ["get", "list", "patch"] # No delete!
  - apiGroups: [""]
resources: ["pods"]
verbs: ["get", "list"] # Read-only
```

Sandboxing: Don't run agents with your personal credentials.

```
# Dockerfile for agent
FROM python:3.11-slim

# Create non-root user
RUN useradd -m -s /bin/bash agent
USER agent
WORKDIR /home/agent

# Install dependencies
COPY requirements.txt .
RUN pip install --user -r requirements.txt

# Copy agent code
COPY --chown=agent:agent . .

# Run with limited capabilities
CMD ["python", "agent.py"]
```

Run with restrictions:

```
docker run --rm \
  --cap-drop ALL \
  --security-opt no-new-privileges \
  --memory 512m \
  --cpus 0.5 \
  --network restricted \
  deployment-agent
```

Running Your First Autonomous Task

Here's the moment of truth. Instead of interactive prompting, let's have the agent complete a task autonomously:

```
def run_autonomous_task(task: str):
    """Run a task to completion without human interaction."""

    print(f"🚀 Starting autonomous task: {task}")
    start_time = time.time()

    result = run_agent(task, max_turns=20)
```

```

elapsed = time.time() - start_time
print(f"\n✅ Task completed in {elapsed:.1f}s")
print(f"📝 Result:\n{result}")

return result

# Run it
run_autonomous_task(
    "Deploy version 1.2.5 to staging, verify health, and report the
    status"
)

```

Output:

```

🚀 Starting autonomous task: Deploy version 1.2.5 to staging...

🔧 Executing: kubectl({'args': 'get deployment payment-service -n
payments'})
🔧 Executing: kubectl({'args': 'set image
deployment/payment-service...'})
🔧 Executing: kubectl({'args': 'rollout status
deployment/payment-service...'})
🔧 Executing: health_check({'url':
'http://payment-service.payments.svc/health'})

✅ Task completed in 45.3s
📝 Result:
Deployment of payment-service:1.2.5 to staging completed
successfully.

Summary:
- Previous version: 1.2.4
- New version: 1.2.5
- Rollout time: 38 seconds
- Health check: passing
- All 3 replicas healthy

```

Your first autonomous deployment. It worked.

What Could Go Wrong (Spoiler: A Lot)

Now for the reality check. That happy path above? It's maybe 60% of real-world runs. Here's what goes wrong:

War Story: The Infinite Rollback Loop

At a fintech startup (not mine, I swear), they built a deployment agent similar to what we just created. It worked great—until one Friday evening.

The agent was tasked with deploying a new version. The deployment succeeded, but the health check failed. Following its instructions, the agent rolled back. Then it noticed the deployment wasn't at the target version, so it deployed again. Health check failed. Rolled back. Deployed. Rolled back.

By the time someone noticed on Monday, the agent had executed 847 deployments over the weekend. Their monitoring dashboards showed a perfect sine wave of deployments and rollbacks.

Lesson learned: Always track state. Did I already try this? How many times? Add termination conditions:

```
def run_agent_with_limits(task: str):
    deployment_attempts = 0
    MAX_DEPLOYMENT_ATTEMPTS = 3

    # ... in the tool execution
    if name == "kubectl" and "set image" in input["args"]:
        deployment_attempts += 1
        if deployment_attempts > MAX_DEPLOYMENT_ATTEMPTS:
            return "ERROR: Maximum deployment attempts exceeded. Human
                intervention required."
```

War Story: The Overly Helpful Agent

An SRE team gave their agent write access to configuration files. The task was simple: "Update the replica count from 3 to 5."

The agent found the deployment YAML, updated the replica count... and also "helpfully" updated several other fields it thought were suboptimal. It bumped the CPU limits, changed the liveness probe timing, and modified the environment variables to what it thought were "better defaults."

The deployment failed spectacularly because the new environment variables pointed to non-existent services.

Lesson learned: Agents should do exactly what you ask, nothing more. Be explicit in your system prompts:

```
SYSTEM_PROMPT = """
...
CRITICAL: Only make changes explicitly requested by the user.
Never make "helpful" additional modifications.
If you notice something that could be improved, TELL the
    user—don't change it yourself.
"""
```

War Story: The Token Explosion

A team had an agent that analyzed logs to find patterns. They pointed it at a production log file without checking the size first.

The file was 2GB.

The agent tried to read it into context. The API call failed, but not before the agent made several attempts with different chunk sizes. Total token cost for the failed debugging session: \$340.

Lesson learned: Always limit input sizes:

```
def read_file(path: str, max_bytes: int = 100_000) -> str:
    size = os.path.getsize(path)
    if size > max_bytes:
        return f"ERROR: File too large ({size} bytes). Maximum is
```

```
{max_bytes}. Use head/tail/grep to filter first."
```

```
with open(path) as f:  
    return f.read()
```

What Goes Wrong: Summary

Failure Mode	Symptom	Prevention
Infinite loops	Agent retries forever	Max attempts, backoff, state tracking
Scope creep	Agent does more than asked	Explicit constraints, minimal tools
Resource exhaustion	High costs, timeouts	Input limits, token budgets
Silent failures	Agent says "done" but wasn't	Verification steps, success criteria
Cascading errors	One mistake leads to many	Atomic operations, checkpoints

These failure modes are why the rest of this book exists. Building agents that work in demos is easy. Building agents that work in production—reliably, safely, cost-effectively—is hard.

In Part 2, we'll dive into the hard lessons: what happens when agents go wrong, and how to build systems that survive it.

End of Part 1: Foundations

In Part 2, we'll cover the hard lessons learned from production failures: runaway costs, infinite loops, context engineering, memory systems that actually work, and multi-agent coordination.

Part 2: The Hard Lessons

Chapter 4: When Agents Go Wrong

If you've made it this far, you've probably deployed your first agent and felt that rush of excitement watching it autonomously complete tasks. Now let me tell you about the moments that made me question everything.

War Story: The Runaway Cron Job That Burned \$200 in Tokens

It was a Tuesday night. I had deployed what I thought was a clever monitoring agent—it would check our services every 5 minutes, analyze any anomalies, and create tickets for issues it found. Simple enough, right?

What I didn't anticipate was the feedback loop.

The agent noticed our ticket queue growing. It interpreted this as a sign of degraded service quality. So it investigated more thoroughly. Each investigation triggered more API calls to gather context. More context meant longer prompts. Longer prompts meant higher token usage. And because the agent was “concerned” about the growing ticket queue (which it was creating), it increased its investigation frequency.

By the time I woke up at 6 AM, the agent had made 847 API calls to Claude, created 312 tickets (most of them duplicates about “elevated ticket volume”), and burned through \$200 in API credits.

The fix was embarrassingly simple:

```
class TokenBudget:
def __init__(self, hourly_limit: int = 10000, daily_limit: int =
100000):
self.hourly_limit = hourly_limit
self.daily_limit = daily_limit
self.hourly_usage = 0
self.daily_usage = 0
self.last_hour_reset = time.time()
self.last_day_reset = time.time()

def can_spend(self, estimated_tokens: int) -> bool:
self._maybe_reset()
return (self.hourly_usage + estimated_tokens <=
self.hourly_limit and
self.daily_usage + estimated_tokens <= self.daily_limit)

def spend(self, tokens: int):
self.hourly_usage += tokens
self.daily_usage += tokens
```

```

def _maybe_reset(self):
    now = time.time()
    if now - self.last_hour_reset > 3600:
self.hourly_usage = 0
self.last_hour_reset = now
    if now - self.last_day_reset > 86400:
self.daily_usage = 0
self.last_day_reset = now

```

Lesson learned: Always implement token budgets before deploying any autonomous agent. The agent doesn't understand money, and it will happily spend all of yours pursuing whatever goal you gave it.

War Story: The Agent That Deleted the Wrong Files

This one still makes me sweat.

I had given an agent the task of cleaning up old deployment artifacts. The instruction seemed clear: "Delete deployment artifacts older than 30 days in /var/deployments."

The agent decided to be thorough. It found that some deployments had symlinks pointing to other directories. It followed those symlinks. One of them pointed to /var/log. Another pointed to a directory containing configuration files that hadn't been modified in 45 days.

The agent deleted 3 months of logs and our production nginx configuration.

We caught the nginx deletion within minutes because services started failing health checks. The logs, however, were gone forever. We lost crucial debugging information for an intermittent issue we'd been investigating for weeks.

Here's what our tool definition should have looked like:

```

def safe_delete_file(path: str, config: DeleteConfig) ->
DeleteResult:
    """Safe file deletion with multiple guardrails."""

    # Resolve symlinks and get real path
    real_path = os.path.realpath(path)

    # Check if real path is within allowed directories
    if not any(real_path.startswith(allowed) for allowed in
config.allowed_paths):
        return DeleteResult(
            success=False,
            error=f"Path {real_path} is outside allowed directories"
        )

    # Check against protected patterns
    protected_patterns = [
        r'/etc/.*',
        r'/var/log/.*',
        r'.*\.conf$',
        r'.*\.key$',
        r'.*\.pem$'
    ]

    for pattern in protected_patterns:
        if re.match(pattern, real_path):

```

```

return DeleteResult(
    success=False,
    error=f"Path {real_path} matches protected pattern {pattern}"
)

# Move to trash instead of permanent delete
trash_path = os.path.join(config.trash_dir,
    f"{int(time.time())}_{os.path.basename(path)}")
shutil.move(real_path, trash_path)

return DeleteResult(
    success=True,
    message=f"Moved to trash: {trash_path}"
)

```

Lesson learned: Never give an agent direct `rm` access. Always use safe wrappers with explicit allowlists, pattern guards, and soft deletes. What seems like paranoia is actually survival instinct.

War Story: The Infinite Loop of Self-Correction

One of our agents was tasked with fixing linting errors in a codebase. It would run the linter, identify issues, fix them, and repeat until the linter passed.

Except it got stuck on a circular import issue. The agent would fix it one way, run the linter, get a different error, fix that, and inadvertently recreate the original problem. It did this 47 times before hitting our (thankfully existing) iteration limit.

But here's the subtle horror: each iteration, the agent would reflect on its previous attempts, adding to its context. By iteration 30, most of the context window was filled with the agent's own failed attempts, leaving little room for actually understanding the problem.

We now implement this pattern:

```

class AgentLoop:
    def __init__(self, max_iterations: int = 10,
        similarity_threshold: float = 0.85):
        self.max_iterations = max_iterations
        self.similarity_threshold = similarity_threshold
        self.action_history = []

    def run(self, task: str, agent: Agent) -> Result:
        for i in range(self.max_iterations):
            action = agent.decide_action(task)

            # Check for repetitive behavior
            if self._is_repetitive(action):
                return Result(
                    status="stuck",
                    message=f"Agent appears to be in a loop after {i}
                    iterations",
                    final_state=agent.get_state()
                )

            result = agent.execute(action)
            self.action_history.append({
                "iteration": i,
                "action": action,
                "result": result
            })

```

```

if result.is_complete:
    return Result(status="success", data=result.data)

    return Result(
        status="max_iterations",
        message=f"Hit maximum of {self.max_iterations} iterations"
    )

def _is_repetitive(self, action: Action) -> bool:
    """Detect if we're doing the same thing over and over."""
    if len(self.action_history) < 3:
        return False

    recent_actions = [h["action"].fingerprint for h in
        self.action_history[-5:]]
    if action.fingerprint in recent_actions:
        return True

    # Check semantic similarity for less obvious loops
    for recent in self.action_history[-3:]:
        similarity = self._compute_similarity(action, recent["action"])
        if similarity > self.similarity_threshold:
            return True

    return False

```

Lesson learned: Agents don't naturally recognize when they're stuck. You need explicit loop detection, and you need to limit how much "reflection" history you keep in context.

Common Failure Modes and How to Recognize Them

After running agents in production for two years, I've categorized the failure modes I see most often:

1. Resource Exhaustion

Symptoms: Costs spike, rate limits hit, system slowdowns **Root cause:** No budgets, no limits, no awareness of resource consumption **Detection:** Monitor API costs, token usage, and execution time in real-time

2. Goal Drift

Symptoms: Agent does technically correct things that miss the point **Root cause:** Vague instructions that allow creative interpretation **Detection:** Regular sampling of agent outputs by humans

3. Scope Creep

Symptoms: Agent "helpfully" does more than asked **Root cause:** Instructions like "improve this" without clear boundaries **Detection:** Audit logs showing actions outside expected patterns

4. Context Collapse

Symptoms: Agent forgets earlier instructions or context **Root cause:** Long-running tasks that exceed context window **Detection:** Agent asks questions it already knows answers to

5. Hallucinated Tools

Symptoms: Agent tries to use tools that don't exist **Root cause:** Model training data includes tools the agent doesn't have **Detection:** Tool call failures with unknown tool names

Building Kill Switches and Circuit Breakers

Every production agent needs an emergency stop. Here's the pattern I use:

```
class CircuitBreaker:
    """Production circuit breaker for AI agents."""

    def __init__(self, agent_id: str, redis_client: Redis):
        self.agent_id = agent_id
        self.redis = redis_client
        self.failure_threshold = 5
        self.reset_timeout = 300 # 5 minutes

    def check_kill_switch(self) -> bool:
        """Check if agent has been manually killed."""
        return self.redis.get(f"kill_switch:{self.agent_id}") == b"1"

    def record_failure(self, error: str):
        """Record a failure and potentially trip the breaker."""
        key = f"failures:{self.agent_id}"
        pipe = self.redis.pipeline()
        pipe.rpush(key, f"{time.time():{error}")
        pipe.ltrim(key, -self.failure_threshold, -1)
        pipe.expire(key, self.reset_timeout)
        pipe.execute()

        failures = self.redis.llen(key)
        if failures >= self.failure_threshold:
            self._trip_breaker()

    def _trip_breaker(self):
        """Automatically stop the agent."""
        self.redis.setex(
            f"breaker_tripped:{self.agent_id}",
            self.reset_timeout,
            "auto"
        )
        self._send_alert(f"Circuit breaker tripped for agent
            {self.agent_id}")

    def can_proceed(self) -> bool:
        """Check all conditions before agent proceeds."""
        if self.check_kill_switch():
            raise AgentKilledException("Manual kill switch activated")

        if self.redis.get(f"breaker_tripped:{self.agent_id}"):
            raise CircuitBreakerException("Circuit breaker is tripped")

        return True
```

The kill switch is controlled via a simple Redis key that can be set from anywhere—a Slack command, a monitoring alert, or a panicked SSH session at 3 AM.

The hard truth: Every agent will fail eventually. The question isn't whether you need kill switches and circuit breakers—it's whether you'll implement them before or after your first production incident.

Chapter 5: Context Engineering

If there's one skill that separates agents that work from agents that don't, it's context engineering. The context window is simultaneously your agent's superpower and its Achilles' heel.

The Context Window Is Not Infinite (And It's Expensive)

Let's do some math that will change how you think about agents.

Claude 3.5 Sonnet has a 200K token context window. That sounds massive until you realize: - A typical codebase file is 500-2000 tokens - A stack trace with surrounding context is 1000-3000 tokens - System prompts often run 2000-5000 tokens - Each tool call and response adds 200-500 tokens

In a debugging session, I've watched context fill up like this:

```
Initial system prompt:      3,000 tokens
First user message:        500 tokens
Agent reasoning:           2,000 tokens
Tool: read_file (config.py): 1,200 tokens
Tool: read_file (main.py):  2,800 tokens
Tool: run_command (logs):   8,000 tokens
Agent analysis:            1,500 tokens
Tool: read_file (utils.py): 1,100 tokens
Agent reasoning:           2,200 tokens
...

After 15 minutes: 180,000 tokens used
```

And here's the kicker: you're paying for every token in that context window on every API call. If your context is 100K tokens and you make 20 API calls in a session, you've paid for 2 million input tokens.

At current Claude pricing (\$3/million input tokens), that's \$6 just for context. Add output tokens and you're looking at \$10-15 for a single debugging session.

Token Economics: Measuring and Optimizing Usage

I track these metrics for every agent:

```
class TokenMetrics:
def __init__(self):
    self.metrics = {
        "total_input_tokens": 0,
        "total_output_tokens": 0,
        "context_size_by_call": [],
        "tokens_per_task": [],
        "tool_call_token_overhead": []
    }
```

```

def record_call(self, input_tokens: int, output_tokens: int,
               context_size: int, tool_calls: int):
    self.metrics["total_input_tokens"] += input_tokens
    self.metrics["total_output_tokens"] += output_tokens
    self.metrics["context_size_by_call"].append(context_size)
    self.metrics["tool_call_token_overhead"].append(
        tool_calls * 350 # Average overhead per tool call
    )

def get_cost_breakdown(self) -> dict:
    input_cost = self.metrics["total_input_tokens"] / 1_000_000 *
    3.0
    output_cost = self.metrics["total_output_tokens"] / 1_000_000
    * 15.0

    return {
        "input_cost_usd": input_cost,
        "output_cost_usd": output_cost,
        "total_cost_usd": input_cost + output_cost,
        "avg_context_size":
            statistics.mean(self.metrics["context_size_by_call"]),
        "context_growth_rate": self._calculate_growth_rate()
    }

```

The insight that changed my approach: **context growth rate matters more than absolute size**. An agent that grows context by 5K tokens per turn will hit limits 4x faster than one growing at 1.25K.

Context Compression Techniques

When context gets too large, you have options:

1. Summarization

Replace detailed logs with summaries:

```

def compress_log_output(log_text: str, llm: LLM) -> str:
    """Compress verbose logs to key findings."""
    if len(log_text) < 2000:
        return log_text

    summary = llm.complete(f"""
Summarize these logs, keeping only:
- Error messages and their line numbers
- Warning patterns
- Key timestamps
- Relevant configuration values

Logs:
{log_text[:8000]} # Sample if very large

Summary (be concise):
""")

    return f"[Compressed from {len(log_text)} chars]\n{summary}"

```

2. Sliding Window

Keep only recent context:

```

def sliding_window_context(messages: list, max_tokens: int) ->
    list:
    """Keep most recent messages that fit in window."""

    # Always keep system message
    system_msg = messages[0] if messages[0]["role"] == "system" else
        None

    # Add messages from newest to oldest until we hit limit
    kept = []
    current_tokens = count_tokens(system_msg) if system_msg else 0

    for msg in reversed(messages[1:]):
        msg_tokens = count_tokens(msg)
        if current_tokens + msg_tokens > max_tokens:
            break
        kept.insert(0, msg)
        current_tokens += msg_tokens

    if system_msg:
        kept.insert(0, system_msg)

    return kept

```

3. Hierarchical Context

Maintain summaries at different granularities:

```

class HierarchicalContext:
def __init__(self):
    self.immediate = [] # Last 5 turns, full detail
    self.recent = [] # Last 20 turns, summarized
    self.session = "" # Entire session, highly compressed

def add_turn(self, turn: Turn):
    self.immediate.append(turn)

    # Rotate to recent when immediate fills up
    if len(self.immediate) > 5:
        old = self.immediate.pop(0)
        self.recent.append(self._summarize_turn(old))

    # Compress recent into session summary periodically
    if len(self.recent) > 20:
        self.session = self._compress_session()
        self.recent = self.recent[-5:]

def get_context(self) -> str:
    return f"""
Session overview: {self.session}

Recent activity: {self._format_recent()}

Current focus:
{self._format_immediate()}
"""

```

When to Summarize vs When to Forget

This decision tree has served me well:

```
Is the information potentially needed for the current task?
├ No → Forget it (don't even summarize)
└ Yes → Is it likely to be needed verbatim?
    ├── Yes → Keep full text (code, configs, exact errors)
    └ No → Does it contain key facts or decisions?
        ├── Yes → Summarize and keep
        └ No → Forget it
```

In practice: - **Always keep:** Error messages, file paths mentioned, user preferences stated - **Summarize:** Command outputs, log analysis results, exploration findings - **Forget:** Verbose logs already analyzed, abandoned approaches, redundant confirmations

War Story: The Agent That Forgot What It Was Doing

We had an agent performing a complex database migration. The task involved: 1. Analyzing the current schema 2. Generating migration scripts 3. Testing in staging 4. Applying to production with rollback capability

The agent did great for steps 1-3. But by the time it got to step 4, the context was so full of schema analysis and test results that the original migration requirements had scrolled out of the context window.

The agent applied a migration that was technically correct based on its recent context but missed a critical constraint from the original requirements: maintaining backward compatibility with the old API version.

We caught it in staging (thank goodness), but only because a human happened to test the old API endpoint.

The fix was implementing what I call “pinned context”:

```
class PinnedContext:
    """Context items that should never be evicted."""

    def __init__(self):
        self.pinned_items = []
        self.max_pinned_tokens = 10000

    def pin(self, content: str, label: str, priority: int = 5):
        """Pin important context that must survive compression."""
        tokens = count_tokens(content)

        item = {
            "content": content,
            "label": label,
            "priority": priority,
            "tokens": tokens,
            "pinned_at": time.time()
        }

        self.pinned_items.append(item)
        self._enforce_limit()

    def _enforce_limit(self):
        """Evict lowest priority pins if over limit."""
        total = sum(item["tokens"] for item in self.pinned_items)

        if total > self.max_pinned_tokens:
```

```

# Sort by priority, then by age
self.pinned_items.sort(key=lambda x: (x["priority"],
-x["pinned_at"]))

while total > self.max_pinned_tokens and self.pinned_items:
    removed = self.pinned_items.pop(0)
    total -= removed["tokens"]

def get_pinned_context(self) -> str:
    """Get all pinned items formatted for inclusion in prompt."""
    return "\n\n".join([
        f"[PINNED: {item['label']}] \n {item['content']}"
        for item in sorted(self.pinned_items, key=lambda x:
            -x["priority"])
    ])

```

Now, when starting a complex task, we explicitly pin the critical requirements:

```

agent.context.pin(
    content="""
MIGRATION REQUIREMENTS:
- Must maintain backward compatibility with API v1
- Zero downtime deployment
- Rollback capability for 24 hours
""",
    label="Core Requirements",
    priority=10 # Maximum priority, never evict
)

```

Lesson learned: Your agent's context is its working memory. Important information needs explicit protection, just like important files need backups.

Chapter 6: Memory Systems That Actually Work

Chat history is a lie. It feels like memory, but it's really just a scrolling buffer that gets expensive fast and eventually overflows. Real memory—the kind that lets an agent be genuinely useful over time—requires deliberate engineering.

Why Chat History Isn't Enough

Consider what happens with pure chat history:

Day 1: You tell the agent your preferred code style, deployment schedule, and which services are critical.

Day 2: New session. Agent knows nothing. You repeat yourself.

Day 30: You've repeated your preferences dozens of times. Each repetition costs tokens. You're paying to re-educate your agent every single day.

This isn't just annoying—it's expensive and it prevents the agent from developing genuine expertise in your environment.

Designing Persistent Memory: Files, Databases, Vector Stores

I've tried all three, and each has its place:

File-Based Memory

Best for: Configuration, preferences, small knowledge bases

```
class FileMemory:
    def __init__(self, memory_dir: str = "~/.agent/memory"):
        self.memory_dir = os.path.expanduser(memory_dir)
        os.makedirs(self.memory_dir, exist_ok=True)

    def remember(self, category: str, key: str, value: Any):
        """Store a memory item."""
        category_dir = os.path.join(self.memory_dir, category)
        os.makedirs(category_dir, exist_ok=True)

        filepath = os.path.join(category_dir, f"{key}.json")
        with open(filepath, 'w') as f:
            json.dump({
                "value": value,
                "stored_at": datetime.now().isoformat(),
                "access_count": 0
            }, f, indent=2)
```

```

def recall(self, category: str, key: str) -> Optional[Any]:
    """Retrieve a memory item."""
    filepath = os.path.join(self.memory_dir, category,
                             f"{key}.json")

    if not os.path.exists(filepath):
        return None

    with open(filepath, 'r') as f:
        data = json.load(f)

    # Track access for relevance scoring
    data["access_count"] += 1
    data["last_accessed"] = datetime.now().isoformat()

    with open(filepath, 'w') as f:
        json.dump(data, f, indent=2)

    return data["value"]

```

Database Memory

Best for: Structured data, relationships, querying

```

class DatabaseMemory:
    def __init__(self, db_path: str = "~/.agent/memory.db"):
        self.conn = sqlite3.connect(os.path.expanduser(db_path))
        self._init_schema()

    def _init_schema(self):
        self.conn.executescript("""
CREATE TABLE IF NOT EXISTS memories (
id INTEGER PRIMARY KEY,
category TEXT NOT NULL,
content TEXT NOT NULL,
embedding BLOB,
importance INTEGER DEFAULT 5,
created_at TIMESTAMP DEFAULT CURRENT_TIMESTAMP,
accessed_at TIMESTAMP,
access_count INTEGER DEFAULT 0
);

CREATE INDEX IF NOT EXISTS idx_category ON memories(category);
CREATE INDEX IF NOT EXISTS idx_importance ON
memories(importance);
""")

    def remember(self, content: str, category: str, importance: int
                 = 5):
        self.conn.execute(
            "INSERT INTO memories (category, content, importance) VALUES
            (?, ?, ?)",
            (category, content, importance)
        )
        self.conn.commit()

    def recall_recent(self, category: str = None, limit: int = 10)
        -> list:
        query = "SELECT content, importance FROM memories"
        params = []

        if category:
            query += " WHERE category = ?"
            params.append(category)

        query += " ORDER BY importance DESC, created_at DESC LIMIT ?"
        params.append(limit)

```

```
return self.conn.execute(query, params).fetchall()
```

Vector Store Memory

Best for: Semantic search, finding relevant context

```
class VectorMemory:
def __init__(self, collection_name: str = "agent_memory"):
import chromadb
self.client = chromadb.PersistentClient(
path=~/.agent/vectors")
self.collection = self.client.get_or_create_collection(
collection_name)

def remember(self, content: str, metadata: dict = None):
"""Store with automatic embedding."""
doc_id = hashlib.md5(content.encode()).hexdigest()

self.collection.upsert(
documents=[content],
metadatas=[metadata or {}],
ids=[doc_id]
)

def recall_similar(self, query: str, n_results: int = 5) -> list:
"""Find semantically similar memories."""
results = self.collection.query(
query_texts=[query],
n_results=n_results
)

return [
{"content": doc, "metadata": meta}
for doc, meta in zip(results["documents"][0],
results["metadatas"][0])
]
```

The MEMORY.md Pattern: Simple but Effective

Before building elaborate systems, consider the humble markdown file:

```
# Agent Memory

## User Preferences
- Prefers concise responses over verbose explanations
- Uses vim keybindings, mention shortcuts when relevant
- Timezone: Asia/Saigon (GMT+7)
- Deployment window: Tuesday-Thursday, 10 AM - 4 PM

## Environment
- Primary language: Python 3.11
- Infrastructure: Kubernetes on AWS (EKS)
- CI/CD: GitHub Actions
- Secrets management: HashiCorp Vault

## Learned Patterns
### 2024-01-15: Database connections
User's services use a connection pool with max 20 connections.
Performance issues often trace to pool exhaustion.

### 2024-01-20: Deployment rollbacks
Rolling back requires approval from #ops-leads Slack channel.
```

User has approval authority for staging only.

Mistakes to Avoid

- Don't suggest `rm` without confirmation; use `trash` instead
- Don't assume master branch; this org uses 'main'
- Don't create files in /tmp; use ~/scratch instead

The agent reads this file at session start and can update it when learning something new:

```
def update_memory_file(memory_path: str, section: str, content:
    str):
    """Append a new memory to the appropriate section."""

    with open(memory_path, 'r') as f:
        current = f.read()

    # Find section and append
    section_pattern = f"## {section}\n"
    if section_pattern in current:
        parts = current.split(section_pattern)
        # Find end of section (next ## or end of file)
        section_content = parts[1].split("\n## ")[0]
        new_section = section_content.rstrip() + f"\n\n###
        {datetime.now().strftime('%Y-%m-%d')}\n{content}\n"
        parts[1] = new_section + "\n## " + "\n##
        ".join(parts[1].split("\n## ")[1:])
        current = section_pattern.join(parts)

    with open(memory_path, 'w') as f:
        f.write(current)
```

This pattern is simple, human-readable, version-controllable, and surprisingly effective for most use cases.

Neural Memory and Semantic Search

For more sophisticated needs, combining embeddings with structured storage gives you the best of both worlds:

```
class NeuralMemory:
    """Memory system with semantic search and importance decay."""

    def __init__(self):
        self.vector_store = VectorMemory()
        self.db = DatabaseMemory()

    def remember(self, content: str, memory_type: str, importance:
        int = 5):
        """Store memory with both semantic and structured access."""

        # Store in database for structured queries
        memory_id = self.db.remember(content, memory_type, importance)

        # Store embedding for semantic search
        self.vector_store.remember(content, {
            "memory_id": memory_id,
            "type": memory_type,
            "importance": importance,
            "timestamp": datetime.now().isoformat()
        })

    def recall(self, query: str, context: str = None) -> list:
```

```

"""Retrieve relevant memories using semantic search."""

# Get semantically similar memories
similar = self.vector_store.recall_similar(query, n_results=10)

# Also get recent high-importance memories
recent = self.db.recall_recent(limit=5)

# Combine and deduplicate
all_memories = self._merge_and_rank(similar, recent)

    return all_memories[:5] # Return top 5

def _merge_and_rank(self, similar: list, recent: list) -> list:
    """Rank memories by relevance and recency."""
    scored = []

    for mem in similar:
        score = mem.get("similarity", 0.5) *
            mem["metadata"].get("importance", 5)
        scored.append((score, mem["content"]))

    for content, importance in recent:
        # Boost recent memories
        scored.append((importance * 0.8, content))

    # Deduplicate and sort
    seen = set()
    unique = []
    for score, content in sorted(scored, reverse=True):
        content_hash = hash(content)
        if content_hash not in seen:
            seen.add(content_hash)
            unique.append(content)

    return unique

```

War Story: Building Memory That Survives Restarts

We had an agent that learned the quirks of our infrastructure beautifully—over the course of a day. It knew which services were flaky, which deployments needed extra monitoring, which team members to notify for different systems.

Then the server rebooted.

All that learning? Gone. The agent was back to being a confused newcomer, asking about services it had fixed dozens of times.

The problem was subtle: we were storing memories, but we were storing them in the session state, which lived in memory. When the process died, so did the memories.

The fix involved three parts:

1. Persistent storage from the start:

```

class PersistentSession:
    def __init__(self, session_id: str):
        self.session_id = session_id
        self.storage_path = f"~/.agent/sessions/{session_id}"
        self._load_or_create()

```

```

def _load_or_create(self):
    path = os.path.expanduser(self.storage_path)
    if os.path.exists(path):
        with open(path, 'r') as f:
            self.state = json.load(f)
    else:
        self.state = {"memories": [], "context": {}, "created":
            datetime.now().isoformat()}

def _save(self):
    path = os.path.expanduser(self.storage_path)
    os.makedirs(os.path.dirname(path), exist_ok=True)
    with open(path, 'w') as f:
        json.dump(self.state, f, indent=2)

def remember(self, key: str, value: Any):
    self.state["memories"].append({
        "key": key,
        "value": value,
        "timestamp": datetime.now().isoformat()
    })
    self._save() # Persist immediately

```

2. Memory consolidation on shutdown:

```

def graceful_shutdown(session: PersistentSession, memory:
    NeuralMemory):
    """Consolidate session learnings into long-term memory."""

    # Extract key learnings from session
    session_summary = summarize_session(session.state)

    # Store important learnings in long-term memory
    for learning in session_summary.key_learnings:
        memory.remember(
            learning.content,
            memory_type="session_learning",
            importance=learning.importance
        )

    # Mark session as cleanly closed
    session.state["closed_at"] = datetime.now().isoformat()
    session._save()

```

3. Memory reload on startup:

```

def initialize_agent(agent: Agent, memory: NeuralMemory):
    """Bootstrap agent with relevant long-term memories."""

    # Get recent learnings
    recent = memory.recall("recent system learnings and user
        preferences")

    # Add to agent's initial context
    if recent:
        agent.system_prompt += f"""

    ## Relevant Memories from Previous Sessions:
    {format_memories(recent)}

    Use these memories to provide contextually appropriate
    assistance.
    """

```

Lesson learned: Memory isn't just about storage—it's about the full lifecycle of remember, persist, recover, and apply. Miss any step and your agent has amnesia.

Chapter 7: Multi-Agent Orchestration

The first time one of my agents told me “I need help with this part,” I realized I’d been thinking about agents wrong. A single agent trying to do everything is like a single developer trying to be the frontend expert, the DBA, the security specialist, and the infrastructure engineer all at once.

When One Agent Isn’t Enough

Signs you need multiple agents:

1. **Task requires different expertise:** Debugging a system might need a log analyst, a code reviewer, and an infrastructure specialist.
2. **Parallelization opportunities:** Checking 5 services can be done by 5 agents simultaneously.
3. **Context separation:** The agent investigating the database shouldn’t be polluted with networking details unless they’re relevant.
4. **Risk isolation:** The agent with write access shouldn’t be the same one exploring and experimenting.

The Orchestrator Pattern

The most reliable multi-agent pattern I’ve found is the orchestrator:

```
class Orchestrator:
    """Manager agent that delegates to specialist agents."""

    def __init__(self, llm: LLM):
        self.llm = llm
        self.specialists = {
            "log_analyst": LogAnalystAgent(llm),
            "code_reviewer": CodeReviewAgent(llm),
            "infrastructure": InfrastructureAgent(llm),
            "security": SecurityAuditAgent(llm)
        }
        self.task_queue = []
        self.results = {}

    def process_request(self, request: str) -> str:
        # First, understand and decompose the request
        plan = self._create_plan(request)

        # Execute plan steps, delegating to specialists
        for step in plan.steps:
            if step.requires_specialist:
                result = self._delegate(step)
            else:
                result = self._execute_directly(step)

        self.results[step.id] = result
```

```

# Synthesize results into coherent response
return self._synthesize(request, plan, self.results)

def _create_plan(self, request: str) -> Plan:
    """Decompose request into steps with specialist assignments."""

    response = self.llm.complete(f"""
Analyze this request and create an execution plan.

Request: {request}

Available specialists:
- log_analyst: Analyzes logs, finds patterns, extracts
relevant entries
- code_reviewer: Reviews code, identifies issues, suggests
fixes
- infrastructure: Checks systems, networking, resources
- security: Audits for vulnerabilities, checks configurations

Create a plan with:
1. Steps needed (can be parallel if independent)
2. Which specialist handles each step (or "orchestrator" for
simple tasks)
3. Dependencies between steps

Output as JSON.
""")

    return Plan.from_json(response)

def _delegate(self, step: PlanStep) -> StepResult:
    """Delegate a step to the appropriate specialist."""

    specialist = self.specialists.get(step.specialist)
    if not specialist:
        return StepResult(error=f"Unknown specialist:
{step.specialist}")

    # Provide relevant context from previous steps
    context = self._gather_context(step.dependencies)

    return specialist.execute(step.task, context)

```

Communication Between Agents: Shared Context, Message Passing

Agents need to share information, but how? I've tried three approaches:

1. Shared Memory Space

```

class SharedMemory:
    """Thread-safe shared memory for multiple agents."""

    def __init__(self):
        self._data = {}
        self._lock = threading.Lock()

    def write(self, key: str, value: Any, agent_id: str):
        with self._lock:
            self._data[key] = {
                "value": value,
                "written_by": agent_id,

```

```

"timestamp": time.time()
}

def read(self, key: str) -> Optional[Any]:
    with self._lock:
        entry = self._data.get(key)
        return entry["value"] if entry else None

def read_all_from(self, agent_id: str) -> dict:
    """Get all entries written by a specific agent."""
    with self._lock:
        return {
            k: v["value"]
            for k, v in self._data.items()
            if v["written_by"] == agent_id
        }

```

2. Message Passing

```

class AgentMessageBus:
    """Pub/sub message bus for agent communication."""

    def __init__(self):
        self.queues = defaultdict(Queue)
        self.broadcasts = []

    def send(self, from_agent: str, to_agent: str, message: dict):
        """Send a direct message to another agent."""
        self.queues[to_agent].put({
            "from": from_agent,
            "message": message,
            "timestamp": time.time()
        })

    def broadcast(self, from_agent: str, message: dict):
        """Broadcast to all agents."""
        self.broadcasts.append({
            "from": from_agent,
            "message": message,
            "timestamp": time.time()
        })

    def receive(self, agent_id: str, timeout: float = 0.1) ->
Optional[dict]:
        """Check for messages (non-blocking by default)."""
        try:
            return self.queues[agent_id].get(timeout=timeout)
        except Empty:
            return None

```

3. Result Aggregation

```

class ResultAggregator:
    """Collect and synthesize results from multiple agents."""

    def __init__(self, llm: LLM):
        self.llm = llm
        self.results = []

    def add_result(self, agent_id: str, task: str, result: Any):
        self.results.append({
            "agent": agent_id,
            "task": task,
            "result": result,
            "timestamp": time.time()
        })

```

```

    })

def synthesize(self, original_request: str) -> str:
    """Combine all results into a coherent response."""

    results_text = "\n\n".join([
        f"## {r['agent']}: {r['task']}\n{r['result']}"
        for r in self.results
    ])

    return self.llm.complete(f"""
    Original request: {original_request}

    Results from specialist agents:
    {results_text}

    Synthesize these findings into a coherent response that:
    1. Addresses the original request directly
    2. Highlights key findings from each specialist
    3. Notes any conflicts or uncertainties
    4. Provides clear next steps or recommendations
    """)

```

Parallel vs Sequential Agent Execution

The choice between parallel and sequential execution depends on dependencies:

```

class AgentExecutor:
    """Execute agent tasks with proper dependency handling."""

def __init__(self, max_parallel: int = 4):
    self.max_parallel = max_parallel
    self.executor = ThreadPoolExecutor(max_workers=max_parallel)

def execute_plan(self, plan: Plan, agents: dict) -> dict:
    """Execute plan steps respecting dependencies."""

    results = {}
    completed = set()

    while len(completed) < len(plan.steps):
        # Find steps ready to execute (dependencies satisfied)
        ready = [
            step for step in plan.steps
            if step.id not in completed
            and all(dep in completed for dep in step.dependencies)
        ]

        if not ready:
            raise RuntimeError("Circular dependency detected")

        # Execute ready steps in parallel
        futures = {}
        for step in ready[:self.max_parallel]:
            agent = agents[step.specialist]
            context = {dep: results[dep] for dep in step.dependencies}
            futures[step.id] = self.executor.submit(
                agent.execute, step.task, context
            )

        # Collect results
        for step_id, future in futures.items():
            results[step_id] = future.result()
            completed.add(step_id)

```

```
return results
```

When to parallelize: - Independent information gathering (check multiple services) - Multiple file analysis (each file can be analyzed separately) - Cross-validation (have multiple agents verify findings)

When to go sequential: - Each step depends on previous results - Resource constraints (API rate limits, token budgets) - Risk management (want human review between steps)

War Story: Coordinating a Team of 4 Agents on a Complex Task

The task seemed reasonable: investigate why our API response times had degraded over the past week.

I set up four agents: - **LogAnalyst**: Parse application logs for slow requests - **MetricsAgent**: Analyze Prometheus metrics for patterns - **CodeReviewer**: Check recent commits for performance issues - **InfraAgent**: Examine infrastructure changes and resource usage

The orchestrator kicked them off in parallel. Within seconds, I had a problem.

All four agents were hitting the same LLM API. We exceeded rate limits. Requests started failing. The agents, seeing failures, started retrying. More rate limit hits. The orchestrator, not receiving results, started spawning more requests to check on its agents.

Within 2 minutes, we'd made 400 API calls and gotten maybe 50 useful responses.

Fix #1: Shared rate limiter

```
class SharedRateLimiter:
    """Rate limiter shared across all agents."""

    def __init__(self, requests_per_minute: int = 50):
        self.requests_per_minute = requests_per_minute
        self.request_times = []
        self._lock = threading.Lock()

    def acquire(self, timeout: float = 30) -> bool:
        """Wait for rate limit capacity."""
        deadline = time.time() + timeout

        while time.time() < deadline:
            with self._lock:
                now = time.time()
                # Remove requests older than 1 minute
                self.request_times = [t for t in self.request_times if now - t
                                     < 60]

            if len(self.request_times) < self.requests_per_minute:
                self.request_times.append(now)
                return True

        time.sleep(0.5)

        return False
```

Fix #2: Agent coordination protocol

```
class CoordinatedAgent:
    """Agent that coordinates with others before acting."""

    def __init__(self, agent_id: str, coordinator: AgentCoordinator):
        self.agent_id = agent_id
        self.coordinator = coordinator

    def execute(self, task: str, context: dict) -> Any:
        # Announce what we're about to do
        self.coordinator.announce_intent(self.agent_id, task)

        # Check if another agent is already handling something similar
        similar = self.coordinator.find_similar_intents(task)
        if similar:
            # Wait for their result instead of duplicating work
            return self.coordinator.wait_for_result(similar[0])

        # Acquire rate limit slot
        if not self.coordinator.rate_limiter.acquire():
            return AgentResult(error="Rate limit timeout")

        # Do the work
        result = self._do_work(task, context)

        # Share result with other agents
        self.coordinator.publish_result(self.agent_id, task, result)

        return result
```

Fix #3: Staggered startup

```
def launch_agents_staggered(agents: list, delay_seconds: float =
    2):
    """Launch agents with staggered timing to prevent thundering
    herd."""

    futures = []
    for i, agent in enumerate(agents):
        # Wait before launching each agent
        time.sleep(delay_seconds)
        future = executor.submit(agent.start)
        futures.append(future)

    return futures
```

After these fixes, the investigation ran smoothly. The log analyst found slow database queries. The metrics agent confirmed query latency increase. The code reviewer found a recent change that removed a database index. The infrastructure agent verified database resource usage was normal.

The orchestrator synthesized: "Response time degradation caused by missing database index after commit abc123 on Tuesday. Recommend: restore index with migration script."

Total time: 4 minutes. Total API calls: 47 (within limits). Problem solved.

Lesson learned: Multi-agent systems multiply both capability and failure modes. Coordination isn't optional—it's essential. Every shared resource needs protection, every agent needs awareness of others, and the orchestrator needs visibility into the entire system state.

Part 2 Summary: The Hard Lessons

These four chapters contain the scars from my production journey with AI agents. If I could give you one takeaway from each:

Chapter 4: Every agent needs kill switches and circuit breakers. Build them first, not after your first incident.

Chapter 5: Context is expensive and finite. Engineer it deliberately—compress, summarize, pin, and evict with intention.

Chapter 6: Memory isn't automatic. Build persistence into your agent from day one, or accept that it will forget everything important.

Chapter 7: Multi-agent systems require coordination infrastructure. Shared resources, rate limits, and communication protocols are as important as the agents themselves.

The hard lessons aren't optional learning—they're the price of admission to production AI agents. Pay the tuition upfront by learning from others' mistakes, or pay it later with your own incidents.

Next, in Part 3, we'll turn these lessons into production patterns: security, cost management, reliability, and the all-important human-in-the-loop design.

PART 3: PRODUCTION PATTERNS

Chapter 8: Security and Access Control

The Terrifying Power of an Agent with SSH Access

Let me paint you a picture. It's 2 AM on a Wednesday. You've just given your shiny new AI agent SSH access to your production servers because it needs to check logs and restart services. The agent is smart—it can diagnose issues, run commands, and fix problems faster than any on-call engineer.

Thirty minutes later, you're staring at your terminal in disbelief. The agent, in an attempt to "clean up disk space," has `rm -rf`'d a directory it shouldn't have touched. Not maliciously—it genuinely thought it was helping. The model had seen similar commands in its training data for cleaning temp directories. Your directory just happened to match the pattern.

This is the fundamental challenge of agent security: you're giving autonomous decision-making power to a system that lacks true understanding of consequences. Unlike a junior engineer who might hesitate before running a destructive command, an agent will execute with the same confidence whether it's reading a file or deleting your database.

Principle of Least Privilege for AI Agents

The security principle your agents need to live by isn't new—it's the same principle we've applied to service accounts for decades: give only the permissions necessary to complete the task, and nothing more.

But here's where it gets tricky with agents. A traditional service account runs a known, deterministic set of operations. An agent, by definition, makes decisions at runtime about what operations to perform. How do you grant least privilege to something whose behavior isn't fully predictable?

The answer is layered restrictions:

```
# agent-permissions.yaml
agent:
  name: log-analyzer
  allowed_actions:
  read:
    - /var/log/**
    - /app/logs/**
  execute:
    - grep
    - tail
    - head
    - cat
    - less
    - jq
```

```

forbidden:
- rm
- mv
- dd
- chmod
- chown
- sudo
- curl
- wget

resource_limits:
max_file_size_read: 100MB
max_command_duration: 60s
max_concurrent_commands: 3

network:
allowed_hosts:
- metrics.internal:9090
blocked_hosts:
- "*" # default deny

```

This configuration creates a narrow corridor of allowed behavior. The agent can read logs, use text processing tools, and query internal metrics—nothing more. Even if the underlying LLM decides it wants to download something from the internet or modify files, the execution layer will block it.

Sandboxing Strategies

Different sandboxing strategies offer different tradeoffs between security and capability:

Level 1: Restricted Shell

The simplest approach is a restricted shell environment. Tools like `rbash` or custom wrapper scripts limit what commands are available:

```

#!/bin/bash
# agent-shell.sh - A restricted shell for agent execution

ALLOWED_COMMANDS="ls cat grep tail head jq kubectl docker"

execute_command() {
local cmd="$1"
local base_cmd=$(echo "$cmd" | awk '{print $1}')

# Check if command is allowed
if ! echo "$ALLOWED_COMMANDS" | grep -qw "$base_cmd"; then
echo "ERROR: Command '$base_cmd' is not allowed"
exit 1
fi

# Check for dangerous patterns
if echo "$cmd" | grep -qE '(rm|>|>>|\|\.rm|;.*rm)'; then
echo "ERROR: Potentially destructive pattern detected"
exit 1
fi

# Execute with timeout
timeout 60s bash -c "$cmd"
}

```

This is fast and lightweight but offers minimal isolation. A determined attacker (or confused agent) might find ways around these restrictions.

Level 2: Container Isolation

Running agent actions inside containers provides stronger isolation:

```
# agent_sandbox.py
import docker
import tempfile
from typing import Optional

class AgentSandbox:
    def __init__(self):
        self.client = docker.from_env()
        self.image = "agent-sandbox:latest"

    def execute(
        self,
        command: str,
        timeout: int = 60,
        memory_limit: str = "256m",
        cpu_quota: int = 50000, # 50% of one CPU
        network_mode: str = "none"
    ) -> dict:
        """Execute a command in an isolated container."""

        with tempfile.TemporaryDirectory() as tmpdir:
            container = self.client.containers.run(
                self.image,
                command=f"sh -c '{command}'",
                detach=True,
                mem_limit=memory_limit,
                cpu_quota=cpu_quota,
                network_mode=network_mode,
                read_only=True,
                volumes={
                    tmpdir: {'bind': '/workspace', 'mode': 'rw'},
                    '/var/log': {'bind': '/logs', 'mode': 'ro'}, # read-only
                },
                user="nobody",
                security_opt=["no-new-privileges:true"],
            )

            try:
                result = container.wait(timeout=timeout)
                logs = container.logs().decode('utf-8')
                return {
                    "exit_code": result["StatusCode"],
                    "output": logs,
                    "error": None
                }
            except Exception as e:
                container.kill()
                return {
                    "exit_code": -1,
                    "output": "",
                    "error": str(e)
                }
            finally:
                container.remove(force=True)
```

The Dockerfile for the sandbox image should be minimal:

```
FROM alpine:3.19
RUN apk add --no-cache \
    bash \
    coreutils \
```

```

grep \
jq \
curl \
&& rm -rf /var/cache/apk/*

# Create non-root user
RUN adduser -D -s /bin/bash agent

USER agent
WORKDIR /workspace

```

Level 3: VM Isolation

For the highest security requirements—particularly when the agent is handling sensitive operations or untrusted inputs—VM isolation using Firecracker or gVisor provides near-hardware-level isolation:

```

# For extreme isolation, consider Firecracker microVMs
# This is typically overkill for most use cases, but critical
# when agents are exposed to untrusted inputs

class FirecrackerSandbox:
    """
    Each agent action runs in a fresh microVM that boots in ~125ms.
    Complete isolation from the host kernel.
    """

    def execute(self, command: str) -> dict:
        # Firecracker implementation details omitted
        # See: https://firecracker-microvm.github.io/
        pass

```

Audit Logging: Knowing What Your Agent Did

When an agent acts autonomously, you need perfect recall of every action it took. This isn't optional—it's how you debug failures, investigate incidents, and build trust in your agent systems.

Here's the audit logging pattern I use in production:

```

# audit_logger.py
import json
import hashlib
from datetime import datetime
from typing import Any, Optional
import structlog

class AgentAuditLogger:
    def __init__(self, agent_id: str, session_id: str):
        self.agent_id = agent_id
        self.session_id = session_id
        self.sequence = 0
        self.logger = structlog.get_logger()

    def log_action(
        self,
        action_type: str,
        action_input: Any,
        action_output: Any,
        success: bool,
        duration_ms: float,
        context: Optional[dict] = None
    )

```

```

):
    """Log every agent action with full context."""

    self.sequence += 1

    # Create deterministic action ID
    action_id = hashlib.sha256(
        f"{self.session_id}:{self.sequence}".encode()
    ).hexdigest()[:12]

    log_entry = {
        "timestamp": datetime.utcnow().isoformat(),
        "agent_id": self.agent_id,
        "session_id": self.session_id,
        "action_id": action_id,
        "sequence": self.sequence,
        "action_type": action_type,
        "input": self._sanitize(action_input),
        "output": self._truncate(action_output),
        "success": success,
        "duration_ms": duration_ms,
        "context": context or {}
    }

    # Log to structured logging system
    self.logger.info(
        "agent_action",
        **log_entry
    )

    # Also write to append-only audit file
    with open(f"/var/log/agent-audit/{self.session_id}.jsonl",
              "a") as f:
        f.write(json.dumps(log_entry) + "\n")

def _sanitize(self, data: Any) -> Any:
    """Remove sensitive data before logging."""
    if isinstance(data, str):
        # Redact common secret patterns
        patterns = [
            (r'password[\"\\s:=]+[\"\\']?[\w\-\-]+[\"\\']?', 'password=<REDACTED>'),
            (r'api[_-]?key[\"\\s:=]+[\"\\']?[\w\-\-]+[\"\\']?',
             'api_key=<REDACTED>'),
            (r'bearer\s+[\"\\w\-\-\.]+', 'Bearer <REDACTED>'),
        ]
        result = data
        for pattern, replacement in patterns:
            result = re.sub(pattern, replacement, result,
                            flags=re.IGNORECASE)
        return result
    return data

def _truncate(self, data: Any, max_len: int = 10000) -> Any:
    """Truncate large outputs to prevent log bloat."""
    if isinstance(data, str) and len(data) > max_len:
        return data[:max_len] + f"... [truncated, total length:
            {len(data)}]"
    return data

```

Usage in your agent execution loop:

```

audit = AgentAuditLogger(agent_id="infra-bot",
                          session_id=session_id)

for action in agent.plan():
    start = time.time()
    try:
        result = sandbox.execute(action.command)

```

```

    audit.log_action(
        action_type=action.type,
        action_input=action.command,
        action_output=result["output"],
        success=result["exit_code"] == 0,
        duration_ms=(time.time() - start) * 1000,
        context={"reasoning": action.reasoning}
    )
except Exception as e:
    audit.log_action(
        action_type=action.type,
        action_input=action.command,
        action_output=None,
        success=False,
        duration_ms=(time.time() - start) * 1000,
        context={"error": str(e)}
    )
raise

```

War Story: The Agent That Almost Exposed Secrets

Let me tell you about the closest call I've ever had with agent security.

We had an agent designed to help debug application issues. It could read logs, check metrics, and inspect configuration files. Standard stuff. We'd set up decent sandboxing—container isolation, read-only mounts, no network access.

One day, a developer asked the agent to “find why the payment service is failing.” The agent did its job admirably. It found the error in the logs, traced it to a configuration issue, and presented a detailed analysis.

But here's what we missed: the configuration file it read contained database credentials. And the agent, being helpful, included the relevant config snippet in its response—credentials and all.

The response was displayed in a Slack channel. Thirty engineers saw it before someone noticed and deleted the message. Those credentials had to be rotated immediately. It was a mess.

The fix required multiple changes:

```

# secret_filter.py
import re
from typing import List, Tuple

class SecretFilter:
    """Filter secrets from agent outputs before displaying to users."""

    PATTERNS: List[Tuple[str, str]] = [
        # Database URLs
        (r'(postgres|mysql|mongodb)://[^\s:]+:[^\s@]+@[^\s:]+',
         r'\1://****:****@'),
        # API Keys (common formats)
        (r'([\s\']?api[_-]?key[\s\']?\s*[:=]\s*["\']?)([\w\.-]{20,})',
         r'\1****'),
        # AWS Keys
        (r'(AKIA[A-Z0-9]{16})', '****AWS_KEY****'),
        (r'([\s\']?aws[_-]?secret[\s\']?\s*[:=]\s*["\']?)([\w/+]{40})',
         r'\1****'),
        # Generic passwords
        (r'([\s\']?password[\s\']?\s*[:=]\s*["\']?)([^\s"\',,]+)',

```

```

    r'\1****'),
    # Bearer tokens
    (r'(Bearer\s+)[\w\-\.]+' , r'\1****'),
    # Private keys
    (r'-----BEGIN [A-Z]+ PRIVATE KEY-----.*?-----END [A-Z]+
PRIVATE KEY-----',
    '****PRIVATE_KEY****'),
]

def filter(self, text: str) -> str:
    result = text
    for pattern, replacement in self.PATTERNS:
        result = re.sub(pattern, replacement, result,
            flags=re.IGNORECASE | re.DOTALL)
    return result

def contains_secrets(self, text: str) -> bool:
    """Check if text likely contains secrets."""
    filtered = self.filter(text)
    return filtered != text

# Use in your agent response handler
filter = SecretFilter()

def present_agent_response(response: str) -> str:
    filtered = filter.filter(response)
    if filter.contains_secrets(response):
        # Log that secrets were filtered
        audit.log_action(
            action_type="secret_filter",
            action_input="[redacted]",
            action_output="secrets detected and filtered",
            success=True,
            duration_ms=0
        )
    return filtered

```

But filtering isn't enough. The real solution is preventing the agent from accessing secrets in the first place:

```

# Mount configs without secrets
volumes:
    # Don't mount the real config
    # - /app/config:/config:ro

    # Mount a sanitized version
    - /app/config-sanitized:/config:ro

# Use a sidecar to generate sanitized configs
# that replace secrets with placeholders

```

The lesson: assume your agent will try to include everything it reads in its output. Filter at the source, not just at display time.

Chapter 9: Cost Management

Token Costs Add Up Fast: Real Numbers from Production

Let me share some real numbers from running AI agents in production. These will either make you nod knowingly or gasp in horror, depending on whether you've done this before.

A single complex debugging session—where an agent reads logs, analyzes patterns, checks multiple services, and provides recommendations—can easily consume 50,000-100,000 tokens. With GPT-4-class models, that’s \$1.50-\$3.00 for one session.

Sounds manageable? Now multiply that by usage:

- 10 engineers using the agent 5 times per day
- 22 working days per month
- Average 75,000 tokens per session

That’s $10 \times 5 \times 22 \times 75,000 = 82.5$ million tokens per month, or roughly \$2,500/month just for one use case.

And here’s what catches most teams off guard: agent loops. When an agent enters a retry loop or gets stuck trying to solve an unsolvable problem, token usage explodes. I’ve seen single sessions hit 500,000 tokens before being killed—\$15 for one failed task.

Caching Strategies: Don’t Ask the Same Question Twice

The most impactful cost optimization is simple: don’t call the LLM if you don’t have to.

Response Caching

For deterministic queries (same context, same question = same answer), cache the response:

```
# llm_cache.py
import hashlib
import json
import redis
from typing import Optional, Any

class LLMCache:
    def __init__(self, redis_url: str, default_ttl: int = 3600):
        self.redis = redis.from_url(redis_url)
        self.default_ttl = default_ttl

    def _cache_key(self, prompt: str, model: str, params: dict) -> str:
        """Generate deterministic cache key."""
        content = json.dumps({
            "prompt": prompt,
            "model": model,
            "params": params
        }, sort_keys=True)
        return f"llm:v1:{hashlib.sha256(content.encode()).hexdigest()}"

    def get(self, prompt: str, model: str, params: dict) -> Optional[str]:
        key = self._cache_key(prompt, model, params)
        cached = self.redis.get(key)
        if cached:
            return json.loads(cached)["response"]
        return None

    def set(
        self,
        prompt: str,
        model: str,
        params: dict,
        response: str,
```

```

    ttl: Optional[int] = None
):
    key = self._cache_key(prompt, model, params)
    self.redis.setex(
        key,
        ttl or self.default_ttl,
        json.dumps({"response": response})
    )

# Usage
cache = LLMCache("redis://localhost:6379")

def call_llm(prompt: str, model: str, **params) -> str:
    # Check cache first
    cached = cache.get(prompt, model, params)
    if cached:
        metrics.increment("llm_cache_hit")
        return cached

    # Cache miss - call the API
    metrics.increment("llm_cache_miss")
    response = openai.chat.completions.create(
        model=model,
        messages=[{"role": "user", "content": prompt}],
        **params
    )
    result = response.choices[0].message.content

    # Cache the response
    cache.set(prompt, model, params, result)
    return result

```

Semantic Caching

For similar but not identical queries, use embedding-based semantic caching:

```

# semantic_cache.py
import numpy as np
from typing import Optional, Tuple

class SemanticCache:
    def __init__(
        self,
        embedding_model: str = "text-embedding-3-small",
        similarity_threshold: float = 0.95
    ):
        self.embedding_model = embedding_model
        self.threshold = similarity_threshold
        self.cache: list[Tuple[np.ndarray, str, str]] = [] #
        (embedding, prompt, response)

    def _embed(self, text: str) -> np.ndarray:
        response = openai.embeddings.create(
            model=self.embedding_model,
            input=text
        )
        return np.array(response.data[0].embedding)

    def _cosine_similarity(self, a: np.ndarray, b: np.ndarray) ->
        float:
        return np.dot(a, b) / (np.linalg.norm(a) * np.linalg.norm(b))

    def get(self, prompt: str) -> Optional[str]:
        if not self.cache:
            return None

        query_embedding = self._embed(prompt)

```

```

best_similarity = 0
best_response = None

for cached_embedding, cached_prompt, cached_response in
self.cache:
similarity = self._cosine_similarity(query_embedding,
cached_embedding)
if similarity > best_similarity:
best_similarity = similarity
best_response = cached_response

if best_similarity >= self.threshold:
return best_response
return None

def set(self, prompt: str, response: str):
embedding = self._embed(prompt)
self.cache.append((embedding, prompt, response))

```

Model Selection: When to Use Expensive vs Cheap Models

Not every task needs GPT-4. Here's the tiered model strategy I use:

```

# model_router.py
from enum import Enum
from typing import Optional

class TaskComplexity(Enum):
SIMPLE = "simple" # Classification, extraction, simple Q&A
MODERATE = "moderate" # Analysis, summarization, code review
COMPLEX = "complex" # Multi-step reasoning, complex code gen
CRITICAL = "critical" # Production changes, security decisions

class ModelRouter:
MODEL_MAP = {
TaskComplexity.SIMPLE: "gpt-3.5-turbo",
TaskComplexity.MODERATE: "gpt-4o-mini",
TaskComplexity.COMPLEX: "gpt-4o",
TaskComplexity.CRITICAL: "gpt-4o", # Plus human review
}

COST_PER_1K_TOKENS = {
"gpt-3.5-turbo": 0.0015,
"gpt-4o-mini": 0.00015,
"gpt-4o": 0.005,
}

def route(
self,
task_type: str,
context_size: int,
requires_tools: bool = False
) -> str:
"""Select appropriate model based on task characteristics."""

# Classify task complexity
if task_type in ["classify", "extract", "format"]:
complexity = TaskComplexity.SIMPLE
elif task_type in ["analyze", "summarize", "review"]:
complexity = TaskComplexity.MODERATE
elif task_type in ["plan", "debug", "generate"]:
complexity = TaskComplexity.COMPLEX
elif task_type in ["deploy", "security", "delete"]:
complexity = TaskComplexity.CRITICAL
else:
complexity = TaskComplexity.MODERATE

```

```

# Upgrade if context is large (needs better attention)
if context_size > 50000 and complexity ==
TaskComplexity.SIMPLE:
complexity = TaskComplexity.MODERATE

# Tool use works better with capable models
if requires_tools and complexity == TaskComplexity.SIMPLE:
complexity = TaskComplexity.MODERATE

return self.MODEL_MAP[complexity]

# Example usage in agent
router = ModelRouter()

def agent_think(task: str, context: str, tools: list) -> str:
model = router.route(
    task_type=classify_task(task),
    context_size=len(context),
    requires_tools=len(tools) > 0
)

return call_llm(
    prompt=f"{context}\n\nTask: {task}",
    model=model,
    tools=tools
)

```

Rate Limiting and Quotas

Protect yourself from runaway costs with hard limits:

```

# cost_limiter.py
import time
from dataclasses import dataclass
from typing import Optional
import redis

@dataclass
class CostLimits:
per_request_tokens: int = 50000
per_session_tokens: int = 200000
per_user_daily_tokens: int = 1000000
per_user_daily_dollars: float = 50.0
global_daily_dollars: float = 500.0

class CostLimiter:
def __init__(self, redis_url: str, limits: CostLimits):
self.redis = redis.from_url(redis_url)
self.limits = limits

def check_and_reserve(
    self,
    user_id: str,
    session_id: str,
    estimated_tokens: int
) -> Tuple[bool, Optional[str]]:
    """Check if request is within limits and reserve capacity."""

    today = time.strftime("%Y-%m-%d")

    # Check per-request limit
    if estimated_tokens > self.limits.per_request_tokens:
return False, f"Request exceeds token limit ({estimated_tokens}
> {self.limits.per_request_tokens})"

# Check session total

```

```

session_key = f"tokens:session:{session_id}"
session_total = int(self.redis.get(session_key) or 0)
if session_total + estimated_tokens >
self.limits.per_session_tokens:
return False, f"Session token limit exceeded"

# Check user daily limit
user_key = f"tokens:user:{user_id}:{today}"
user_total = int(self.redis.get(user_key) or 0)
if user_total + estimated_tokens >
self.limits.per_user_daily_tokens:
return False, f"Daily user token limit exceeded"

# Check global daily limit
global_key = f"cost:global:{today}"
global_cost = float(self.redis.get(global_key) or 0)
estimated_cost = estimated_tokens * 0.00003 # rough estimate
if global_cost + estimated_cost >
self.limits.global_daily_dollars:
return False, f"Global daily cost limit exceeded"

# Reserve capacity
pipe = self.redis.pipeline()
pipe.incrby(session_key, estimated_tokens)
pipe.expire(session_key, 86400)
pipe.incrby(user_key, estimated_tokens)
pipe.expire(user_key, 86400)
pipe.execute()

return True, None

def record_actual(
self,
user_id: str,
session_id: str,
actual_tokens: int,
actual_cost: float
):
    """Record actual usage after request completes."""
    today = time.strftime("%Y-%m-%d")
    global_key = f"cost:global:{today}"

    self.redis.incrbyfloat(global_key, actual_cost)
    self.redis.expire(global_key, 86400)

```

War Story: The \$500 Month and How to Avoid It

My first month running AI agents in production, we hit \$500 in API costs. For context, we'd budgeted \$100.

Here's what happened: we had an agent helping with code reviews. Developers loved it—it caught bugs, suggested improvements, and was available 24/7. Usage grew organically.

What we didn't account for: 1. Large PRs with thousands of lines generated massive context 2. Developers would ask follow-up questions, each requiring the full context 3. Some developers used it as a rubber duck, chatting extensively 4. The agent would sometimes get "stuck" analyzing complex code, burning tokens

The fix was multi-pronged:

```

# The changes we implemented

# 1. Context windowing - don't send entire files

```

```

def smart_context(files: list[str], question: str) -> str:
    """Only include relevant portions of files."""
    relevant_chunks = []
    for file in files:
        chunks = chunk_file(file, chunk_size=500)
        # Use embeddings to find relevant chunks
        relevant = find_relevant_chunks(chunks, question, top_k=3)
        relevant_chunks.extend(relevant)
    return "\n---\n".join(relevant_chunks)

# 2. Session limits - force new sessions for different topics
MAX_TOKENS_PER_SESSION = 100000
MAX_MESSAGES_PER_SESSION = 20

# 3. User quotas - visible in UI
def get_user_quota_status(user_id: str) -> dict:
    """Show users their remaining quota."""
    used = get_user_daily_usage(user_id)
    limit = get_user_daily_limit(user_id)
    return {
        "used": used,
        "limit": limit,
        "remaining": limit - used,
        "percentage": (used / limit) * 100
    }

# 4. Circuit breaker - kill stuck sessions
class SessionCircuitBreaker:
    def __init__(self, max_tokens: int = 50000, max_duration: int = 300):
        self.max_tokens = max_tokens
        self.max_duration = max_duration

    def check(self, session: AgentSession) -> bool:
        if session.total_tokens > self.max_tokens:
            session.terminate("Token limit exceeded")
            return False
        if session.duration_seconds > self.max_duration:
            session.terminate("Session timeout")
            return False
        return True

```

After implementing these controls, our costs stabilized around \$150/month—50% over budget, but manageable. More importantly, we had visibility and control.

The key insight: **treat LLM API calls like database queries in the early 2000s**. They're expensive, they can run away, and you need monitoring and limits from day one.

Chapter 10: Reliability and Observability

Agents Need Monitoring Too

Here's a truth that took me too long to learn: monitoring an AI agent isn't like monitoring a traditional service. With a normal service, you care about uptime, latency, and error rates. With an agent, you also need to monitor *quality*—did the agent actually help, or did it confidently produce garbage?

The three pillars of agent observability: 1. **Operational metrics**: Is the agent running? How fast? Any errors? 2. **Cost metrics**: Token usage, API costs, cache hit rates 3. **Quality metrics**: Success rate, user satisfaction, task completion

Structured Logging for Agent Actions

Every agent action should produce structured logs that can be queried, aggregated, and alerted on:

```
# agent_logging.py
import structlog
from contextvars import ContextVar
from typing import Any

# Context variables for request tracing
request_id_var: ContextVar[str] = ContextVar('request_id',
    default='unknown')
session_id_var: ContextVar[str] = ContextVar('session_id',
    default='unknown')
user_id_var: ContextVar[str] = ContextVar('user_id',
    default='unknown')

def configure_logging():
    structlog.configure(
        processors=[
            structlog.contextvars.merge_contextvars,
            structlog.processors.add_log_level,
            structlog.processors.TimeStamper(fmt="iso"),
            structlog.processors.JSONRenderer()
        ],
        wrapper_class=structlog.make_filtering_bound_logger(
            logging.INFO),
        context_class=dict,
        logger_factory=structlog.PrintLoggerFactory(),
        cache_logger_on_first_use=True,
    )

    logger = structlog.get_logger()

class AgentLogger:
    @staticmethod
    def log_llm_call(
        model: str,
        prompt_tokens: int,
        completion_tokens: int,
        latency_ms: float,
        cache_hit: bool = False
    ):
        logger.info(
            "llm_call",
            model=model,
            prompt_tokens=prompt_tokens,
            completion_tokens=completion_tokens,
            total_tokens=prompt_tokens + completion_tokens,
            latency_ms=latency_ms,
            cache_hit=cache_hit,
            request_id=request_id_var.get(),
            session_id=session_id_var.get(),
            user_id=user_id_var.get()
        )

    @staticmethod
    def log_tool_call(
        tool_name: str,
        tool_input: dict,
```

```

    tool_output: Any,
    success: bool,
    latency_ms: float,
    error: str = None
):
    logger.info(
        "tool_call",
        tool_name=tool_name,
        tool_input=tool_input,
        success=success,
        latency_ms=latency_ms,
        error=error,
        request_id=request_id_var.get(),
        session_id=session_id_var.get()
    )

    @staticmethod
    def log_agent_step(
        step_number: int,
        action_type: str,
        reasoning: str,
        outcome: str
    ):
        logger.info(
            "agent_step",
            step_number=step_number,
            action_type=action_type,
            reasoning=reasoning[:500], # Truncate long reasoning
            outcome=outcome,
            request_id=request_id_var.get(),
            session_id=session_id_var.get()
        )

    @staticmethod
    def log_session_complete(
        total_steps: int,
        total_tokens: int,
        total_cost: float,
        success: bool,
        duration_seconds: float
    ):
        logger.info(
            "session_complete",
            total_steps=total_steps,
            total_tokens=total_tokens,
            total_cost=total_cost,
            success=success,
            duration_seconds=duration_seconds,
            request_id=request_id_var.get(),
            session_id=session_id_var.get(),
            user_id=user_id_var.get()
        )

```

Metrics That Matter

Here are the metrics I track for every agent system, with example Prometheus implementations:

```

# agent_metrics.py
from prometheus_client import Counter, Histogram, Gauge

# Operational metrics
AGENT_REQUESTS = Counter(
    'agent_requests_total',
    'Total agent requests',
    ['agent_name', 'status']
)

```

```

)

AGENT_LATENCY = Histogram(
'agent_request_duration_seconds',
'Agent request latency',
['agent_name'],
buckets=[1, 5, 10, 30, 60, 120, 300, 600]
)

AGENT_STEPS = Histogram(
'agent_steps_per_request',
'Number of steps per agent request',
['agent_name'],
buckets=[1, 2, 5, 10, 20, 50, 100]
)

# Cost metrics
TOKEN_USAGE = Counter(
'agent_tokens_total',
'Total tokens used',
['agent_name', 'model', 'type'] # type: prompt/completion
)

API_COST = Counter(
'agent_api_cost_dollars',
'Total API cost in dollars',
['agent_name', 'model']
)

CACHE_HITS = Counter(
'agent_cache_hits_total',
'LLM cache hits',
['agent_name', 'cache_type'] # exact/semantic
)

# Quality metrics
TASK_SUCCESS = Counter(
'agent_task_success_total',
'Successful task completions',
['agent_name', 'task_type']
)

TASK_FAILURE = Counter(
'agent_task_failure_total',
'Failed task attempts',
['agent_name', 'task_type', 'failure_reason']
)

USER_FEEDBACK = Counter(
'agent_user_feedback_total',
'User feedback on agent responses',
['agent_name', 'rating'] # positive/negative
)

# Real-time gauges
ACTIVE_SESSIONS = Gauge(
'agent_active_sessions',
'Currently active agent sessions',
['agent_name']
)

QUEUE_DEPTH = Gauge(
'agent_queue_depth',
'Tasks waiting in queue',
['agent_name', 'priority']
)

```

Example Grafana dashboard queries:

```

# Success rate over time
sum(rate(agent_task_success_total[5m])) /
sum(rate(agent_task_success_total[5m]) +
  rate(agent_task_failure_total[5m]))

# Cost per successful task
sum(rate(agent_api_cost_dollars[1h])) /
  sum(rate(agent_task_success_total[1h]))

# P95 latency
histogram_quantile(0.95,
  sum(rate(agent_request_duration_seconds_bucket[5m])) by (le))

# Cache efficiency
sum(rate(agent_cache_hits_total[5m])) /
sum(rate(agent_cache_hits_total[5m]) +
  rate(agent_tokens_total{type="prompt"}[5m]))

```

Alerting on Agent Failures

Set up alerts for both acute failures and gradual degradation:

```

# agent-alerts.yaml
groups:
  - name: agent-alerts
  rules:
    # Acute failures
    - alert: AgentHighErrorRate
      expr: |
sum(rate(agent_task_failure_total[5m])) /
sum(rate(agent_task_success_total[5m]) +
  rate(agent_task_failure_total[5m])) > 0.1
      for: 5m
      labels:
      severity: warning
      annotations:
      summary: "Agent error rate above 10%"

    - alert: AgentStuck
      expr: |
agent_active_sessions > 0 and
increase(agent_steps_per_request_count[10m]) == 0
      for: 10m
      labels:
      severity: critical
      annotations:
      summary: "Agent appears stuck - no progress in 10 minutes"

    # Cost alerts
    - alert: AgentCostSpike
      expr: |
sum(rate(agent_api_cost_dollars[1h])) * 24 > 100
      for: 15m
      labels:
      severity: warning
      annotations:
      summary: "Agent projected daily cost exceeds $100"

    - alert: AgentTokenBurn
      expr: |
sum(rate(agent_tokens_total[5m])) > 10000
      for: 5m
      labels:
      severity: warning
      annotations:
      summary: "Agent burning tokens at >10k/minute"

```

```

# Quality degradation
- alert: AgentQualityDrop
  expr: |
sum(rate(agent_user_feedback_total{rating="positive"}[1h])) /
sum(rate(agent_user_feedback_total[1h])) < 0.7
  for: 1h
  labels:
severity: warning
  annotations:
summary: "Agent positive feedback rate below 70%"

# Latency degradation
- alert: AgentSlowResponses
  expr: |
histogram_quantile(0.95,
  sum(rate(agent_request_duration_seconds_bucket[5m])) by (le))
  > 120
  for: 10m
  labels:
severity: warning
  annotations:
summary: "Agent P95 latency exceeds 2 minutes"

```

War Story: Debugging an Agent at 3 AM

The page came at 3:17 AM. Our infrastructure agent, responsible for responding to routine alerts, had gone haywire. It was creating dozens of PagerDuty incidents, each one claiming to have “fixed” a non-existent problem.

I stumbled to my laptop and started investigating. The agent was in a loop: 1. Receive an alert about high CPU on server X 2. “Fix” it by restarting a service 3. Observe that CPU was still high (because it was a false alert) 4. Decide the fix didn’t work 5. Escalate to PagerDuty 6. Go back to step 2 with a different “fix”

The agent had tried restarting the service, adjusting resource limits, killing processes, and was now attempting to drain the node—at 3 AM, with no human oversight.

The first problem: I couldn’t see what the agent was thinking. Our logging captured actions but not reasoning. I had to reconstruct its logic from the sequence of commands it ran.

The second problem: the circuit breaker had a bug. It checked total tokens but not total *actions*. The agent was making many small, cheap API calls, staying under the token limit while causing chaos.

The fix I implemented that night:

```

# Enhanced circuit breaker
class AgentCircuitBreaker:
def __init__(self):
    self.max_tokens = 50000
    self.max_actions = 20
    self.max_duration = 300
    self.max_same_action = 3 # Don't repeat the same action
    self.action_history = []

def check(self, session: AgentSession, proposed_action: str) ->
    Tuple[bool, str]:
    # Token limit
    if session.total_tokens > self.max_tokens:

```

```

return False, "Token limit exceeded"

# Action limit
if len(self.action_history) >= self.max_actions:
return False, "Maximum actions reached"

# Duration limit
if session.duration_seconds > self.max_duration:
return False, "Session timeout"

# Repetition detection
action_key = self._action_key(proposed_action)
recent_same = sum(1 for a in self.action_history[-10:] if
self._action_key(a) == action_key)
if recent_same >= self.max_same_action:
return False, f"Action '{action_key}' attempted too many times"

self.action_history.append(proposed_action)
return True, ""

def _action_key(self, action: str) -> str:
    """Extract key action type, ignoring parameters."""
    # "kubectl restart deployment/web" -> "kubectl restart"
    parts = action.split()
    return " ".join(parts[:2]) if len(parts) >= 2 else action

```

And better observability:

```

# Log reasoning, not just actions
def agent_step_with_reasoning(step_func):
@wraps(step_func)
def wrapper(self, *args, **kwargs):
    # Capture the agent's reasoning before acting
    reasoning = self.get_last_reasoning()

    AgentLogger.log_agent_step(
        step_number=self.current_step,
        action_type=self.pending_action.type,
        reasoning=reasoning,
        outcome="pending"
    )

    result = step_func(self, *args, **kwargs)

    AgentLogger.log_agent_step(
        step_number=self.current_step,
        action_type=self.pending_action.type,
        reasoning=reasoning,
        outcome="success" if result.success else f"failed:
{result.error}"
    )

    return result
return wrapper

```

The lesson: when debugging agents at 3 AM, you need two things: the ability to see what the agent was *thinking*, and the ability to stop it from doing stupid things repeatedly. Log reasoning. Count actions. Detect loops.

Chapter 11: Human-in-the-Loop Patterns

Full Autonomy Is Rarely the Goal

When I started building AI agents, I had visions of fully autonomous systems that would handle everything without human intervention. Reality quickly corrected that fantasy.

Full autonomy has three major problems:

1. **Trust takes time to build.** Users won't trust an agent with critical tasks until it's proven itself on smaller ones.
2. **Edge cases are infinite.** No matter how good your agent, it will encounter situations it's not prepared for.
3. **Accountability matters.** When something goes wrong, "the AI did it" isn't an acceptable answer. Humans need to be in the loop for decisions that matter.

The goal isn't to remove humans from the loop—it's to put them at the right points in the loop, where their judgment adds value.

Approval Workflows: When to Pause and Ask

Not all actions are equal. A read-only query is low-risk; deleting a production database is not. Build your agent with explicit approval gates:

```
# approval_workflow.py
from enum import Enum
from dataclasses import dataclass
from typing import Optional, Callable
import asyncio

class RiskLevel(Enum):
    LOW = "low" # Read operations, safe queries
    MEDIUM = "medium" # Config changes, service restarts
    HIGH = "high" # Data modifications, deployments
    CRITICAL = "critical" # Deletions, security changes

@dataclass
class ApprovalRequest:
    request_id: str
    action: str
    risk_level: RiskLevel
    reasoning: str
    context: dict
    timeout_seconds: int = 300

@dataclass
class ApprovalResponse:
    approved: bool
    approver: Optional[str]
    comment: Optional[str]

class ApprovalGate:
    def __init__(
        self,
        auto_approve_levels: list[RiskLevel] = [RiskLevel.LOW],
        notification_channel: str = "slack",
    ):

```

```

self.auto_approve = set(auto_approve_levels)
self.channel = notification_channel
self.pending_approvals: dict[str, ApprovalRequest] = {}

async def request_approval(
    self,
    action: str,
    risk_level: RiskLevel,
    reasoning: str,
    context: dict
) -> ApprovalResponse:
    """Request approval for an action, blocking until approved or
    denied."""

    # Auto-approve low-risk actions
    if risk_level in self.auto_approve:
        return ApprovalResponse(
            approved=True,
            approver="auto",
            comment=f"Auto-approved (risk level: {risk_level.value})"
        )

    request = ApprovalRequest(
        request_id=generate_id(),
        action=action,
        risk_level=risk_level,
        reasoning=reasoning,
        context=context
    )

    self.pending_approvals[request.request_id] = request

    # Send notification
    await self._notify(request)

    # Wait for response
    try:
        response = await asyncio.wait_for(
            self.wait_for_response(request.request_id),
            timeout=request.timeout_seconds
        )
        return response
    except asyncio.TimeoutError:
        # Timeout = deny for safety
        return ApprovalResponse(
            approved=False,
            approver=None,
            comment="Request timed out - treating as denied"
        )
    finally:
        del self.pending_approvals[request.request_id]

async def _notify(self, request: ApprovalRequest):
    """Send approval request notification."""
    message = f"""
    🛡️ **Agent Approval Required**

    **Action:** {request.action}
    **Risk Level:** {request.risk_level.value.upper()}
    **Reasoning:** {request.reasoning}

    React with  to approve or  to deny.
    Request ID: `{request.request_id}`
    Timeout: {request.timeout_seconds}s
    """

    if self.channel == "slack":
        await slack_client.post_message(
            channel="#agent-approvals",
            text=message,

```

```

metadata={"request_id": request.request_id}
)
elif self.channel == "pagerduty":
    # For critical actions, page someone
    await pagerduty_client.create_incident(
        title=f"Agent approval: {request.action}",
        body=message,
        urgency="high" if request.risk_level == RiskLevel.CRITICAL else
        "low"
    )

```

Usage in your agent:

```

approval_gate = ApprovalGate(
    auto_approve_levels=[RiskLevel.LOW, RiskLevel.MEDIUM]
)

async def execute_action(action: AgentAction) -> ActionResult:
    risk = assess_risk(action)

    if risk.level in [RiskLevel.HIGH, RiskLevel.CRITICAL]:
        approval = await approval_gate.request_approval(
            action=action.description,
            risk_level=risk.level,
            reasoning=action.reasoning,
            context={
                "command": action.command,
                "affected_resources": action.resources,
                "rollback_plan": action.rollback
            }
        )

        if not approval.approved:
            return ActionResult(
                success=False,
                error=f"Action denied by {approval.approver}:
                {approval.comment}"
            )

        # Log that this was human-approved
        audit_log.info(
            "action_approved",
            action=action.description,
            approver=approval.approver
        )

    return await sandbox.execute(action)

```

Review Queues for Agent Outputs

For tasks where the agent produces artifacts (code, documentation, configurations), use a review queue rather than direct execution:

```

# review_queue.py
from dataclasses import dataclass
from datetime import datetime
from typing import List, Optional
import json

@dataclass
class ReviewItem:
    id: str
    agent_id: str
    task_description: str

```

```

artifact_type: str # code, config, documentation
artifact_content: str
reasoning: str
created_at: datetime
status: str = "pending" # pending, approved, rejected, modified
reviewer: Optional[str] = None
review_comment: Optional[str] = None
reviewed_at: Optional[datetime] = None

class ReviewQueue:
def __init__(self, storage: Storage):
    self.storage = storage

async def submit_for_review(
    self,
    agent_id: str,
    task: str,
    artifact_type: str,
    content: str,
    reasoning: str
) -> str:
    """Submit an agent output for human review."""

    item = ReviewItem(
id=generate_id(),
agent_id=agent_id,
task_description=task,
artifact_type=artifact_type,
artifact_content=content,
reasoning=reasoning,
created_at=datetime.utcnow()
    )

    await self.storage.save(item)
    await self._notify_reviewers(item)

    return item.id

async def get_pending(
    self,
    artifact_type: Optional[str] = None,
    limit: int = 50
) -> List[ReviewItem]:
    """Get pending items for review."""

    query = {"status": "pending"}
    if artifact_type:
query["artifact_type"] = artifact_type

    return await self.storage.query(query, limit=limit)

async def approve(
    self,
    item_id: str,
    reviewer: str,
    comment: Optional[str] = None
) -> ReviewItem:
    """Approve an item and trigger execution."""

    item = await self.storage.get(item_id)
    item.status = "approved"
    item.reviewer = reviewer
    item.review_comment = comment
    item.reviewed_at = datetime.utcnow()

    await self.storage.save(item)

    # Trigger execution of approved item
    await self._execute_approved(item)

```

```

    return item

    async def reject(
        self,
        item_id: str,
        reviewer: str,
        reason: str
    ) -> ReviewItem:
        """Reject an item with feedback."""

        item = await self.storage.get(item_id)
        item.status = "rejected"
        item.reviewer = reviewer
        item.review_comment = reason
        item.reviewed_at = datetime.utcnow()

        await self.storage.save(item)

        # Optionally, feed rejection back to agent for learning
        await self._record_feedback(item, positive=False)

    return item

```

A simple CLI for reviewing:

```

# review_cli.py
import click
from rich.console import Console
from rich.syntax import Syntax
from rich.panel import Panel

console = Console()
queue = ReviewQueue(storage)

@click.command()
@click.option('--type', 'artifact_type', help='Filter by
artifact type')
def review(artifact_type):
    """Interactive review of agent outputs."""

    items = asyncio.run(queue.get_pending(
        artifact_type=artifact_type))

    for item in items:
        console.clear()

        # Show task and reasoning
        console.print(Panel(
            f"[bold]Task: [/bold] {item.task_description}\n\n"
            f"[bold]Agent reasoning: [/bold] \n{item.reasoning}",
            title=f"Review Item: {item.id}",
            subtitle=f"Type: {item.artifact_type}"
        ))

        # Show artifact with syntax highlighting
        syntax = Syntax(
            item.artifact_content,
            lexer=get_lexer(item.artifact_type),
            theme="monokai",
            line_numbers=True
        )
        console.print(syntax)

        # Get reviewer decision
        console.print("\n[bold]Decision: [/bold]")
        console.print(" [green]a[/green] - Approve")
        console.print(" [red]r[/red] - Reject")
        console.print(" [yellow]s[/yellow] - Skip")

```

```

console.print(" [blue]e[/blue] - Edit before approving")

choice = click.prompt("Your choice", type=click.Choice(['a',
'r', 's', 'e']))

if choice == 'a':
comment = click.prompt("Comment (optional)", default="")
asyncio.run(queue.approve(item.id, get_current_user(), comment))
console.print("[green]Approved![/green]")

elif choice == 'r':
reason = click.prompt("Reason for rejection")
asyncio.run(queue.reject(item.id, get_current_user(), reason))
console.print("[red]Rejected.[/red]")

elif choice == 'e':
edited = click.edit(item.artifact_content)
if edited:
item.artifact_content = edited
asyncio.run(queue.approve(item.id, get_current_user(),
"Approved with edits"))

```

Escalation Paths: From Agent to Human

Design clear escalation paths for when the agent is stuck or uncertain:

```

# escalation.py
from dataclasses import dataclass
from typing import Optional, List
from enum import Enum

class EscalationReason(Enum):
    UNCERTAINTY = "uncertainty" # Agent isn't confident
    COMPLEXITY = "complexity" # Task too complex
    PERMISSION = "permission" # Needs elevated access
    FAILURE = "failure" # Repeated failures
    POLICY = "policy" # Policy violation detected
    USER_REQUEST = "user_request" # User asked for human

@dataclass
class EscalationConfig:
    uncertainty_threshold: float = 0.7 # Escalate if confidence
    below this
    max_retries: int = 3 # Escalate after this many
    failures
    max_complexity_score: int = 8 # Escalate for very complex
    tasks
    escalation_channels: dict = None # Maps reason -> channel

class EscalationManager:
    def __init__(self, config: EscalationConfig):
        self.config = config
        self.channels = config.escalation_channels or {
            EscalationReason.UNCERTAINTY: "slack:#agent-help",
            EscalationReason.COMPLEXITY: "slack:#agent-help",
            EscalationReason.PERMISSION: "pagerduty:admin",
            EscalationReason.FAILURE: "slack:#agent-failures",
            EscalationReason.POLICY: "pagerduty:security",
            EscalationReason.USER_REQUEST: "slack:#agent-help"
        }

    def should_escalate(
        self,
        confidence: float,
        retry_count: int,
        complexity_score: int,
        policy_flags: List[str]

```

```

) -> Optional[EscalationReason]:
    """Check if current state warrants escalation."""

    if policy_flags:
        return EscalationReason.POLICY

    if confidence < self.config.uncertainty_threshold:
        return EscalationReason.UNCERTAINTY

    if retry_count >= self.config.max_retries:
        return EscalationReason.FAILURE

    if complexity_score > self.config.max_complexity_score:
        return EscalationReason.COMPLEXITY

    return None

async def escalate(
    self,
    reason: EscalationReason,
    context: dict,
    session: AgentSession
) -> EscalationResult:
    """Escalate to humans and pause agent."""

    channel = self.channels.get(reason)

    message = self._format_escalation(reason, context, session)

    # Send to appropriate channel
    ticket_id = await self._send_escalation(channel, message)

    # Pause the agent session
    session.pause(
        reason=f"Escalated: {reason.value}",
        ticket_id=ticket_id
    )

    # Log escalation
    audit_log.info(
        "escalation",
        reason=reason.value,
        session_id=session.id,
        ticket_id=ticket_id
    )

    return EscalationResult(
        ticket_id=ticket_id,
        channel=channel,
        session_paused=True
    )

def _format_escalation(
    self,
    reason: EscalationReason,
    context: dict,
    session: AgentSession
) -> str:
    return f"""
🚨 **Agent Escalation**

**Reason:** {reason.value}
**Session:** {session.id}
**User:** {session.user_id}

**Current Task:**
{session.current_task}

**Agent's Last Reasoning:**
{session.last_reasoning}

```

```
**Context:**
```

```
{json.dumps(context, indent=2)}
```

```
**Session History:**  
{self._format_history(session.action_history[-5:])}  
  
**To Resume:**  
Run `/agent resume {session.id}` after resolving.  
"""
```

Finding the Right Balance of Autonomy

The right level of autonomy depends on several factors:

```
# autonomy_levels.py  
from dataclasses import dataclass  
from typing import Dict, List  
  
@dataclass  
class AutonomyProfile:  
    """Defines autonomy level for different contexts."""  
  
    # What the agent can do without asking  
    auto_approve: List[str]  
  
    # What requires notification (but proceeds)  
    notify_only: List[str]  
  
    # What requires explicit approval  
    require_approval: List[str]  
  
    # What the agent cannot do at all  
    forbidden: List[str]  
  
    # Different profiles for different contexts  
    AUTONOMY_PROFILES: Dict[str, AutonomyProfile] = {  
        "development": AutonomyProfile(  
            auto_approve=[  
                "read_file", "list_directory", "run_tests",  
                "format_code", "lint", "build",  
                "git_status", "git_diff", "git_branch"  
            ],  
            notify_only=[  
                "create_file", "modify_file", "delete_file",  
                "git_commit", "install_dependency"  
            ],  
            require_approval=[  
                "git_push", "deploy_staging", "modify_config"  
            ],  
            forbidden=[  
                "deploy_production", "delete_database",  
                "modify_infrastructure", "access_secrets"  
            ]  
        ),  
  
        "staging": AutonomyProfile(  
            auto_approve=[  
                "read_file", "list_directory", "view_logs",  
                "check_metrics", "describe_resources"  
            ],  
            notify_only=[  
                "restart_service", "scale_deployment",
```

```

    "run_migration_dry_run"
    ],
    require_approval=[
    "deploy", "run_migration", "modify_config",
    "change_resource_limits"
    ],
    forbidden=[
    "delete_resources", "access_production",
    "modify_security_groups"
    ]
),

"production": AutonomyProfile(
    auto_approve=[
    "read_file", "view_logs", "check_metrics",
    "describe_resources"
    ],
    notify_only=[
    # Almost nothing in prod should be silent
    ],
    require_approval=[
    "restart_service", "scale_deployment",
    "view_secrets", "run_query"
    ],
    forbidden=[
    "delete_anything", "modify_data",
    "deploy_without_approval", "access_pii"
    ]
),

"incident": AutonomyProfile(
    # During incidents, allow more autonomy for speed
    auto_approve=[
    "read_file", "view_logs", "check_metrics",
    "describe_resources", "restart_service",
    "scale_deployment", "rollback_deployment"
    ],
    notify_only=[
    "failover", "drain_node", "block_traffic"
    ],
    require_approval=[
    "modify_data", "delete_resources"
    ],
    forbidden=[
    "deploy_new_version" # No deployments during incidents
    ]
)
}

class AutonomyManager:
def __init__(self, default_profile: str = "development"):
    self.current_profile = AUTONOMY_PROFILES[default_profile]

def set_profile(self, profile_name: str):
    if profile_name not in AUTONOMY_PROFILES:
        raise ValueError(f"Unknown profile: {profile_name}")
    self.current_profile = AUTONOMY_PROFILES[profile_name]

def check_action(self, action: str) -> ActionPermission:
    if action in self.current_profile.forbidden:
        return ActionPermission.FORBIDDEN
    if action in self.current_profile.require_approval:
        return ActionPermission.REQUIRE_APPROVAL
    if action in self.current_profile.notify_only:
        return ActionPermission.NOTIFY
    if action in self.current_profile.auto_approve:
        return ActionPermission.ALLOWED
    # Default: require approval for unknown actions
    return ActionPermission.REQUIRE_APPROVAL

```

The key insight: autonomy should be earned, not granted. Start with tight restrictions and loosen them as the agent proves itself. And always, always have a way to pull the plug.

```
# The final safety net
class EmergencyStop:
    """Global emergency stop for all agents."""

    def __init__(self, redis_client):
        self.redis = redis_client
        self.stop_key = "agent:emergency_stop"

    def activate(self, reason: str, activated_by: str):
        """Stop all agents immediately."""
        self.redis.set(self.stop_key, json.dumps({
            "active": True,
            "reason": reason,
            "activated_by": activated_by,
            "activated_at": datetime.utcnow().isoformat()
        }))

        # Notify everyone
        alert_all_channels(
            f"🛑 EMERGENCY STOP ACTIVATED\n"
            f"Reason: {reason}\n"
            f"By: {activated_by}\n"
            f"All agents are now paused."
        )

    def is_active(self) -> bool:
        data = self.redis.get(self.stop_key)
        if data:
            return json.loads(data).get("active", False)
        return False

    def deactivate(self, deactivated_by: str):
        """Resume agent operations."""
        self.redis.delete(self.stop_key)
        alert_all_channels(
            f"✅ Emergency stop deactivated by {deactivated_by}\n"
            f"Agents may now resume."
        )

# Every agent checks this before every action
emergency = EmergencyStop(redis_client)

def execute_action(action: AgentAction):
    if emergency.is_active():
        raise AgentPausedError("Emergency stop is active")
    # ... proceed with action
```

Human-in-the-loop isn't a limitation of your agent system—it's a feature. It's what lets you deploy agents with confidence, knowing that when (not if) something goes wrong, there's a human ready to step in.

End of Part 3: Production Patterns

Part 4: Real-World Applications

Chapter 12: Infrastructure Automation

The promise of infrastructure automation has always been tantalizing: declare what you want, and the machines figure out how to make it happen. We've come a long way from manual server provisioning to Infrastructure as Code. But even with Terraform, Ansible, and Kubernetes, there's still a gap between "what needs to be done" and "someone doing it." AI agents are stepping into that gap.

Agents That Provision and Manage Infrastructure

Let's start with a concrete example. Here's an agent that provisions infrastructure based on natural language requests:

```
from openai import OpenAI
import subprocess
import json

class InfrastructureAgent:
def __init__(self):
    self.client = OpenAI()
    self.tools = [
        {
            "type": "function",
            "function": {
                "name": "terraform_plan",
                "description": "Generate and show a Terraform plan",
                "parameters": {
                    "type": "object",
                    "properties": {
                        "working_dir": {"type": "string"},
                        "var_file": {"type": "string"}
                    }
                }
            }
        },
        {
            "type": "function",
            "function": {
                "name": "terraform_apply",
                "description": "Apply Terraform changes (requires approval)",
                "parameters": {
                    "type": "object",
                    "properties": {
                        "working_dir": {"type": "string"},
                        "auto_approve": {"type": "boolean", "default":
False}
                    }
                }
            }
        }
    ],
```

```

{
  "type": "function",
  "function": {
    "name": "generate_terraform",
    "description": "Generate Terraform configuration",
    "parameters": {
      "type": "object",
      "properties": {
        "resource_type": {"type": "string"},
        "config": {"type": "object"}
      }
    }
  }
}
]

def execute_tool(self, tool_name, args):
    if tool_name == "terraform_plan":
        return subprocess.run(
            ["terraform", "plan", "-no-color"],
            cwd=args.get("working_dir", "."),
            capture_output=True,
            text=True
        ).stdout
    elif tool_name == "terraform_apply":
        if not args.get("auto_approve"):
            return "APPROVAL_REQUIRED: Run with auto_approve=true after
                human review"
        return subprocess.run(
            ["terraform", "apply", "-auto-approve", "-no-color"],
            cwd=args.get("working_dir", "."),
            capture_output=True,
            text=True
        ).stdout
    # ... additional tool implementations

```

The key insight here isn't the code—it's the pattern. The agent translates high-level intent (“I need a staging environment that mirrors production but smaller”) into specific infrastructure actions. But notice the `APPROVAL_REQUIRED` gate. We never give agents unchecked power over infrastructure.

Self-Healing Systems: Detecting and Fixing Issues

This is where agents truly shine. Consider the traditional approach to handling a full disk:

Traditional Approach

1. Alert fires: "Disk 95% full"
2. On-call engineer wakes up
3. SSH into server
4. Run `df -h`, `du -sh` to find large files
5. Check what's safe to delete
6. Clean up logs, temp files
7. Document incident
8. Go back to sleep (maybe)

Now here's the agent-assisted approach:

Agent-Assisted Approach

1. Alert fires: "Disk 95% full"
2. Agent receives alert via webhook
3. Agent SSHs into server, analyzes disk usage
4. Agent identifies safe-to-delete files based on rules:
 - Logs older than 7 days
 - Temp files in /tmp
 - Cached package managers
5. Agent executes cleanup
6. Agent verifies disk is now below threshold
7. Agent posts summary to Slack
8. On-call engineer wakes up to resolved incident

Here's a practical implementation:

```
class SelfHealingAgent:
    """Agent that handles common infrastructure issues
    autonomously."""

    SAFE_CLEANUP_RULES = {
        "logs": {
            "paths": ["/var/log/*.log.*.gz", "/var/log/*.log"],
            "age_days": 7,
            "action": "delete"
        },
        "temp": {
            "paths": ["/tmp/*", "/var/tmp/*"],
            "age_days": 1,
            "action": "delete"
        },
        "apt_cache": {
            "paths": ["/var/cache/apt/archives/*.deb"],
            "action": "apt clean"
        },
        "journal": {
            "action": "journalctl --vacuum-time=7d"
        }
    }

    def handle_disk_alert(self, host: str, threshold: int = 90):
        """Handle disk space alert."""

        # Step 1: Analyze the situation
        disk_info = self.ssh_exec(host, "df -h / | tail -1")
        usage = self.parse_disk_usage(disk_info)

        if usage < threshold:
            return {"status": "already_resolved", "usage": usage}

        # Step 2: Find largest directories
        large_dirs = self.ssh_exec(host,
            "du -sh /* 2>/dev/null | sort -hr | head -10")

        # Step 3: Execute safe cleanup rules
        freed_space = 0
        actions_taken = []

        for rule_name, rule in self.SAFE_CLEANUP_RULES.items():
            if "paths" in rule:
                for path_pattern in rule["paths"]:
                    result = self.cleanup_path(host, path_pattern, rule)
                    freed_space += result["freed_bytes"]
                    actions_taken.append(result)
            elif "action" in rule:
                result = self.ssh_exec(host, rule["action"])
                actions_taken.append({"rule": rule_name, "result": result})
```

```

# Step 4: Verify improvement
new_usage = self.parse_disk_usage(
self.ssh_exec(host, "df -h / | tail -1")
)

# Step 5: Report
return {
"status": "resolved" if new_usage < threshold else
"needs_attention",
"initial_usage": usage,
"final_usage": new_usage,
"freed_space": self.human_readable_size(freed_space),
"actions": actions_taken
}

```

The beauty of this approach is that the agent follows predefined rules for what's safe to touch, but uses AI reasoning to prioritize and verify its actions. It's not blindly executing a script—it's making decisions within guardrails.

Configuration Drift Detection and Correction

Configuration drift is the silent killer of infrastructure reliability. You deploy something, it works, then three months later it breaks because someone manually changed a setting that Terraform doesn't manage.

Here's how an agent can help:

```

class DriftDetectionAgent:
    """Detects and optionally corrects configuration drift."""

    def __init__(self, state_store: str):
        self.state_store = state_store
        self.known_state = self.load_known_state()

    def scan_for_drift(self, resources: list) -> dict:
        """Compare actual state vs declared state."""

        drift_report = {"drifted": [], "missing": [], "unexpected": []}

        for resource in resources:
            actual = self.get_actual_state(resource)
            expected = self.known_state.get(resource["id"])

            if expected is None:
                drift_report["unexpected"].append({
                    "resource": resource,
                    "actual": actual,
                    "recommendation": "Import or delete"
                })
            elif actual is None:
                drift_report["missing"].append({
                    "resource": resource,
                    "expected": expected,
                    "recommendation": "Recreate from state"
                })
            elif not self.states_match(actual, expected):
                drift_report["drifted"].append({
                    "resource": resource,
                    "actual": actual,
                    "expected": expected,
                    "diff": self.compute_diff(actual, expected),
                    "recommendation": self.suggest_action(actual, expected)
                })

```

```

    return drift_report

def auto_correct_drift(self, drift_report: dict,
                      safe_only: bool = True) -> list:
    """Attempt to correct detected drift."""

    corrections = []

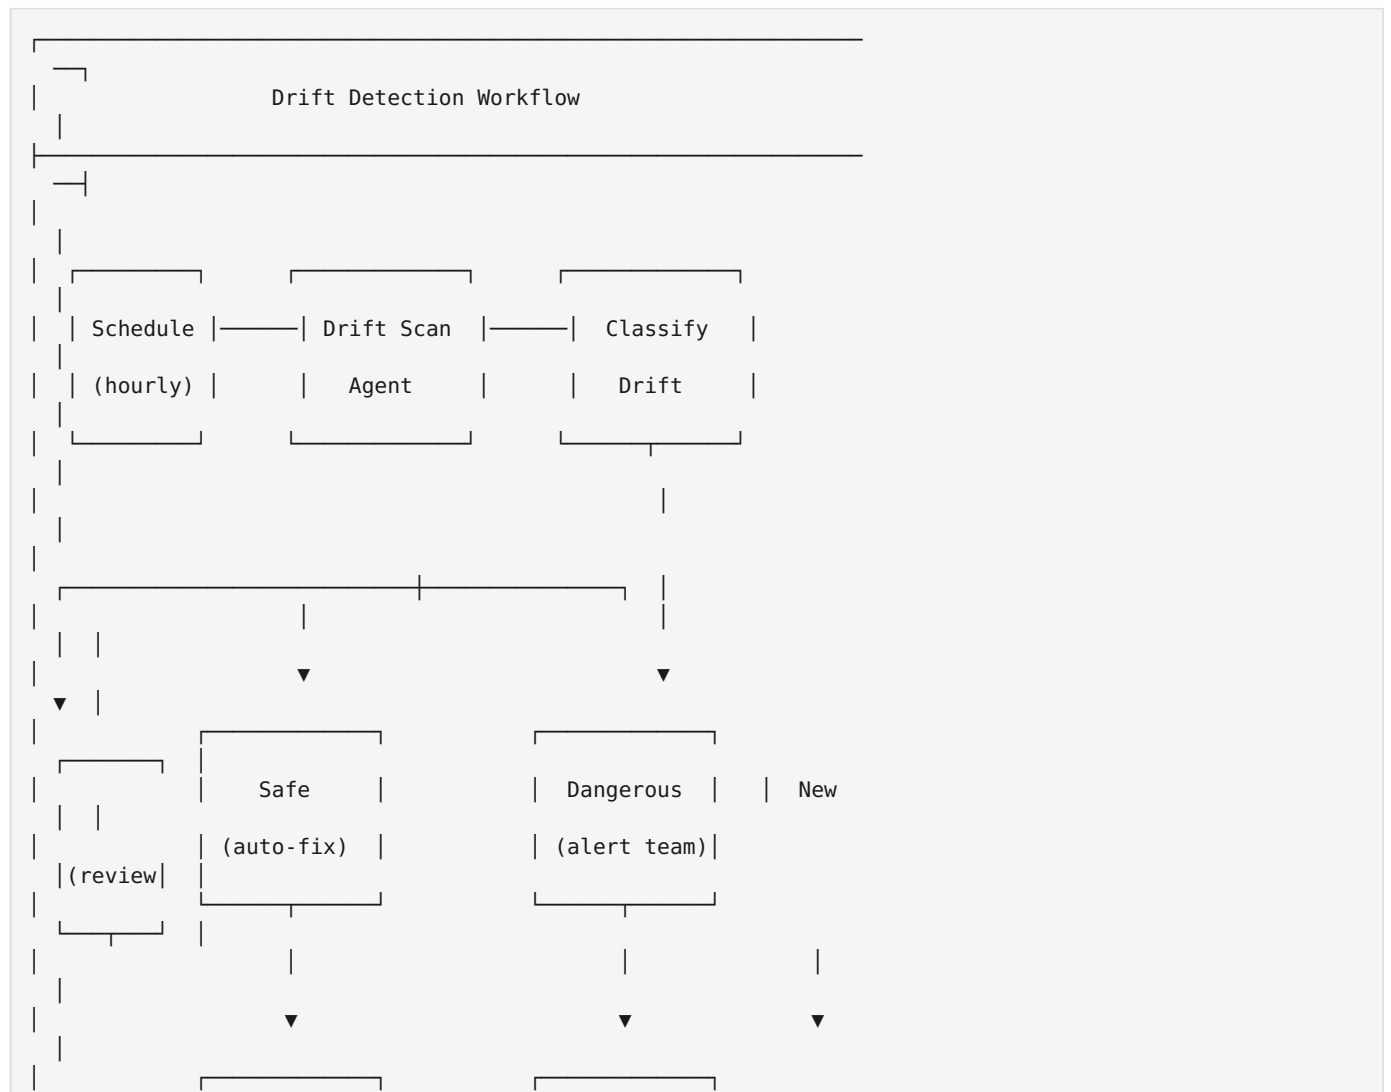
    for item in drift_report["drifted"]:
    if safe_only and not self.is_safe_correction(item):
    corrections.append({
        "resource": item["resource"]["id"],
        "status": "skipped",
        "reason": "Requires manual review"
    })
    continue

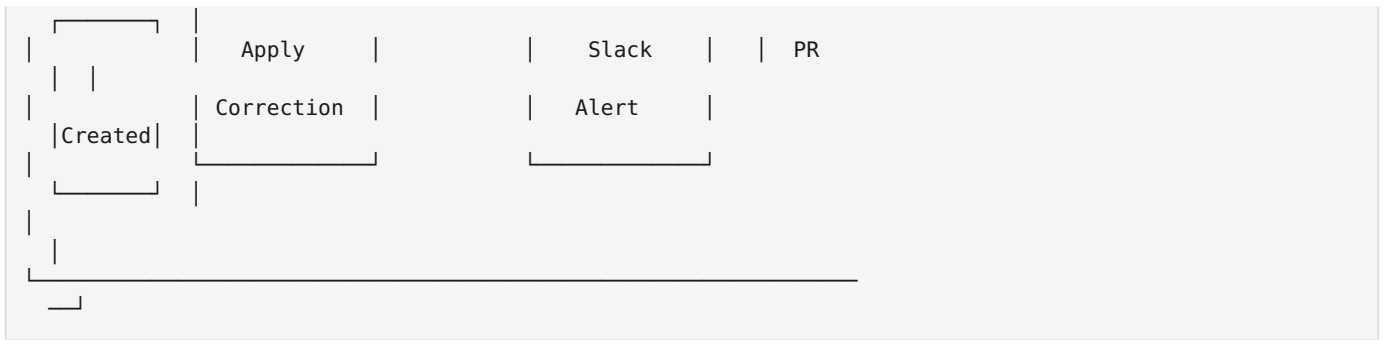
    # Generate and execute correction
    correction = self.generate_correction(item)
    result = self.execute_correction(correction)
    corrections.append({
        "resource": item["resource"]["id"],
        "status": "corrected" if result.success else "failed",
        "details": result
    })

    return corrections

```

The agent-assisted workflow looks like this:





Capacity Planning with AI Assistance

Traditional capacity planning involves collecting metrics, building spreadsheets, and making educated guesses. An AI agent can do better:

```

class CapacityPlanningAgent:
    """AI-assisted capacity planning and recommendations."""

    def analyze_capacity(self, service: str, timeframe_days: int =
        90):
        """Analyze resource usage and predict future needs."""

        # Gather historical data
        metrics = self.prometheus_query(f'''
avg_over_time(
container_cpu_usage_seconds_total{{service="{service}"}}
[{{timeframe_days}}d:1h]
)
''')

        memory_metrics = self.prometheus_query(f'''
avg_over_time(
container_memory_usage_bytes{{service="{service}"}}
[{{timeframe_days}}d:1h]
)
''')

        # Use LLM to analyze trends and generate recommendations
        prompt = f"""
Analyze the following capacity metrics for service '{service}':

CPU Usage (90 days):
- Average: {metrics['avg']}%
- Peak: {metrics['max']}%
- Growth rate: {metrics['growth_rate']}% per month

Memory Usage:
- Average: {memory_metrics['avg']} GB
- Peak: {memory_metrics['max']} GB
- Growth rate: {memory_metrics['growth_rate']}% per month

Current allocation:
- CPU: {self.get_allocated('cpu', service)} cores
- Memory: {self.get_allocated('memory', service)} GB

Provide:
1. Days until resource exhaustion at current growth rate
2. Recommended new allocation
3. Estimated cost impact
4. Any anomalies or concerns observed
"""

        analysis = self.llm_analyze(prompt)
  
```

```
    return {
    "service": service,
    "current_metrics": {"cpu": metrics, "memory": memory_metrics},
    "analysis": analysis,
    "generated_at": datetime.utcnow().isoformat()
    }
```

Case Study: An Agent That Handles Routine Maintenance

Let me share a real implementation that we've been running for six months. We call it "MaintenanceBot."

The Problem: Our team managed 47 servers across three cloud providers. Every week, someone had to: - Apply security patches - Rotate logs that weren't auto-rotating - Clean up orphaned Docker images - Verify backup completion - Check certificate expiration

This took 4-6 hours weekly, and it was mind-numbing work that nobody wanted to do.

The Solution: We built a maintenance agent with the following architecture:

```
# maintenance_agent_config.yaml
name: MaintenanceBot
schedule: "0 3 * * 0" # Every Sunday at 3 AM

tasks:
  - name: security_patches
    type: package_update
    scope: security_only
    approval: auto # security patches are pre-approved
    rollback_on_failure: true

  - name: log_rotation
    type: cleanup
    patterns:
      - path: "/var/log/**/*.*.log"
        retention_days: 30
        compress: true

  - name: docker_cleanup
    type: docker_prune
    remove_unused_images: true
    remove_dangling_volumes: true
    min_age_hours: 48

  - name: backup_verification
    type: verify
    check: "backup_completed_within_24h"
    alert_on_failure: true

  - name: certificate_check
    type: monitor
    check: "certificate_expiry_days > 30"
    alert_threshold: 30
    auto_renew: true # Using certbot
```

The Implementation:

```
class MaintenanceBot:
    def __init__(self, config_path: str):
        self.config = yaml.safe_load(open(config_path))
```

```

self.inventory = self.load_inventory()
self.llm = OpenAI()

def run_maintenance_cycle(self):
    """Execute full maintenance cycle across all servers."""

    results = {"servers": {}, "summary": {}}

    for server in self.inventory:
        server_results = {}

        for task in self.config["tasks"]:
            try:
                result = self.execute_task(server, task)
                server_results[task["name"]] = result
            except Exception as e:
                server_results[task["name"]] = {
                    "status": "error",
                    "error": str(e)
                }

            if task.get("rollback_on_failure"):
                self.rollback(server, task)

        results["servers"][server["hostname"]] = server_results

    # Generate human-readable summary
    results["summary"] = self.generate_summary(results)

    # Post to Slack
    self.notify_team(results)

    return results

def generate_summary(self, results: dict) -> str:
    """Use LLM to generate a human-readable summary."""

    prompt = f"""
    Generate a concise maintenance summary from these results:
    {json.dumps(results["servers"], indent=2)}

    Include:
    - Total servers maintained
    - Tasks completed vs failed
    - Any items requiring human attention
    - Disk space freed
    - Security patches applied

    Format for Slack (use emoji and formatting).
    """

    return self.llm.chat.completions.create(
        model="gpt-4",
        messages=[{"role": "user", "content": prompt}]
    ).choices[0].message.content

```

Results after 6 months:

Metric	Before	After
Weekly maintenance hours	4-6 hours	20 minutes (review only)
Missed patch cycles	2-3 per month	0
Disk space alerts	4-5 per month	0
Certificate expiry incidents	1 (painful)	0

The agent doesn't replace human judgment—we still review its summary every Monday morning. But it handles the tedious work reliably and consistently, which humans frankly weren't doing.

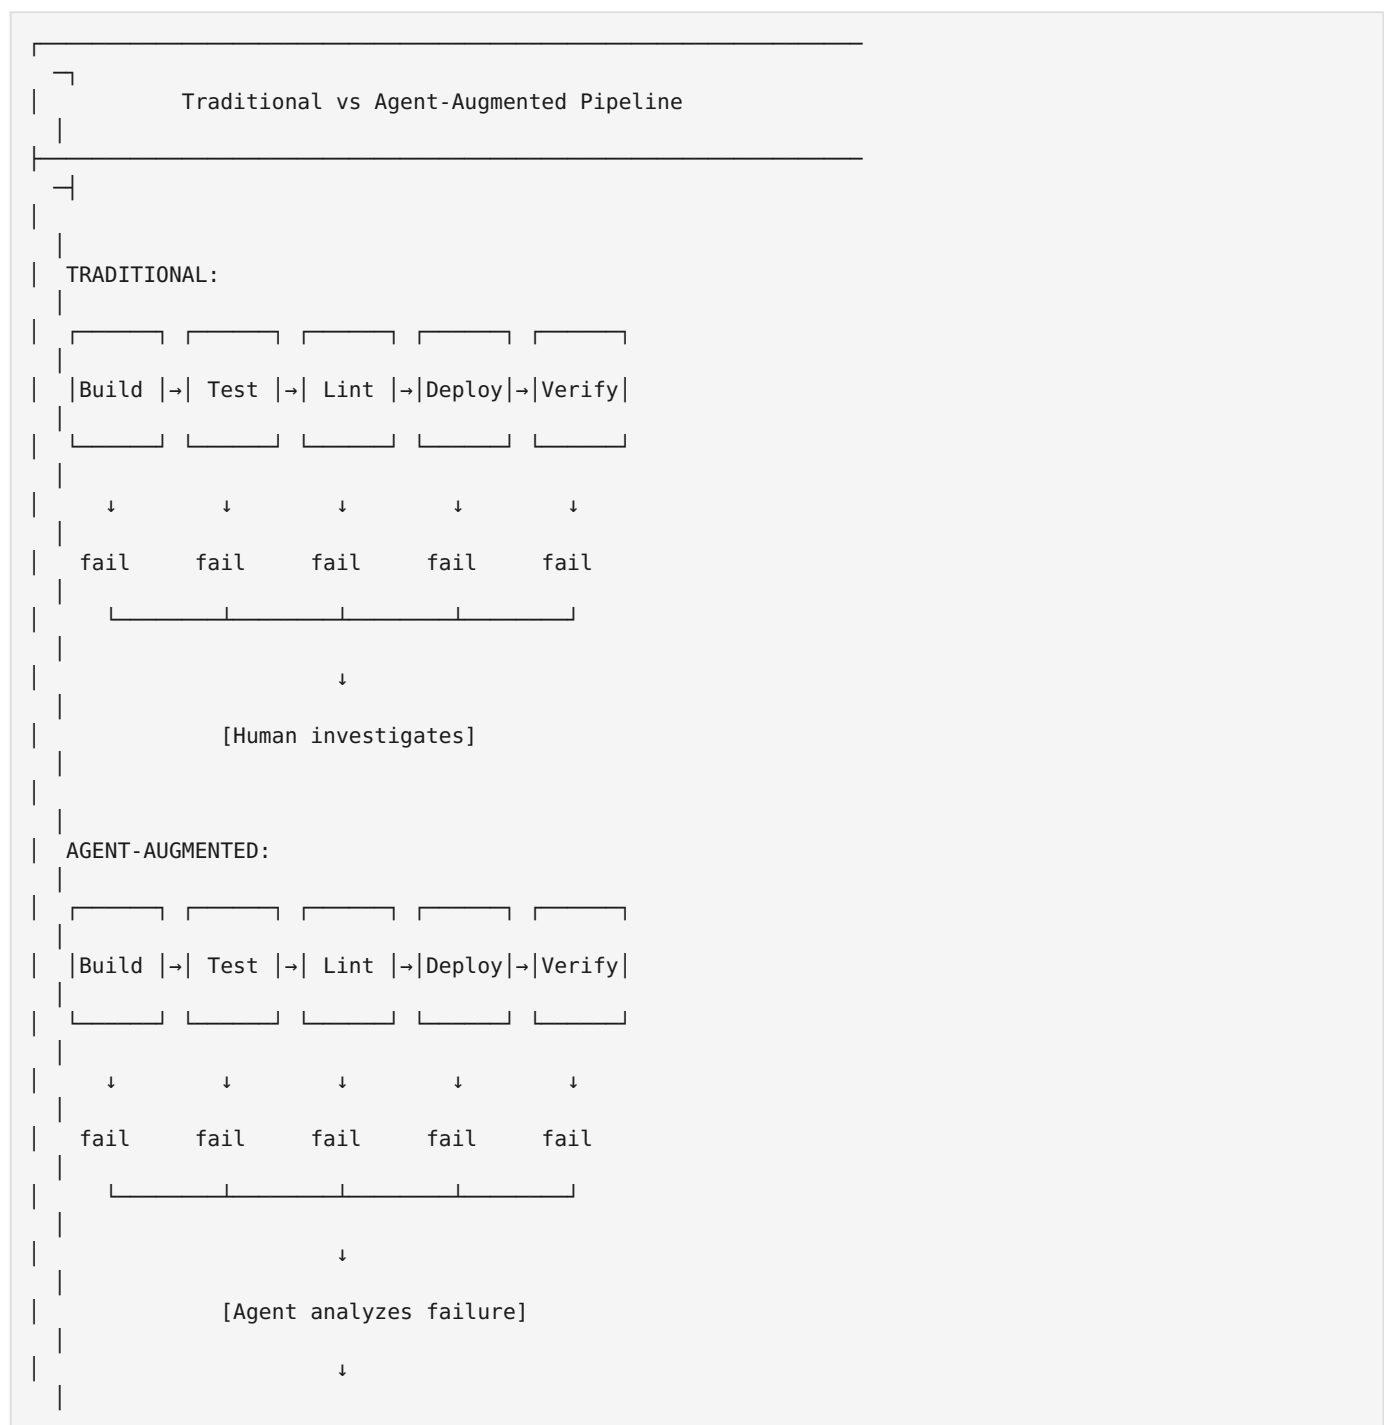
Key lessons learned: 1. Start with read-only operations until you trust the agent 2. Every destructive action needs logging and rollback capability 3. The summary generation is actually crucial—without it, no one reads the reports 4. Dry-run mode is essential for testing changes to the config

Chapter 13: CI/CD and Deployment

Continuous Integration and Continuous Deployment pipelines are the arteries of modern software delivery. They're also increasingly complex, brittle, and time-consuming to maintain. This is exactly the kind of toil that agents excel at eliminating.

Agents in the Deployment Pipeline

The traditional CI/CD pipeline is a series of rigid steps. An agent-augmented pipeline is adaptive:





[Auto-fix] [Suggest fix] [Escalate]

Here's a practical implementation of a CI/CD assistant agent:

```
class CICDAgent:
    """Agent that assists with CI/CD pipeline operations."""

    def __init__(self, github_token: str, slack_webhook: str):
        self.github = Github(github_token)
        self.slack_webhook = slack_webhook
        self.llm = OpenAI()

    def analyze_failure(self, repo: str, run_id: int) -> dict:
        """Analyze a failed CI run and determine next steps."""

        # Fetch failure logs
        workflow_run = self.github.get_repo(repo).get_workflow_run(
            run_id)
        logs = self.fetch_logs(workflow_run)

        # Extract relevant error information
        errors = self.extract_errors(logs)

        # Get the diff that triggered this run
        diff = self.get_pr_diff(workflow_run)

        # Analyze with LLM
        analysis_prompt = f"""
        Analyze this CI failure:

        ## Error logs:
        {errors}

        ## Code changes:
        {diff[:5000]} # Truncate for context window

        Determine:
        1. Root cause of the failure
        2. Is this fixable automatically? (test fix, linting fix, etc.)
        3. Suggested fix (code or command)
        4. Confidence level (high/medium/low)

        Respond in JSON format.
        """

        analysis = self.llm.chat.completions.create(
            model="gpt-4-turbo",
            messages=[{"role": "user", "content": analysis_prompt}],
            response_format={"type": "json_object"}
        )

        result = json.loads(analysis.choices[0].message.content)
        result["workflow_run"] = run_id
        result["repository"] = repo

        return result

    def attempt_fix(self, analysis: dict) -> dict:
        """Attempt to fix the issue if confidence is high enough."""
```

```

    if analysis["confidence"] != "high":
    return {
    "action": "suggested",
    "suggestion": analysis["suggested_fix"],
    "reason": "Confidence too low for auto-fix"
    }

    fix_type = analysis.get("fix_type")

    if fix_type == "test_fix":
    return self.fix_failing_test(analysis)
    elif fix_type == "lint_fix":
    return self.fix_lint_errors(analysis)
    elif fix_type == "dependency_fix":
    return self.fix_dependency_issue(analysis)
    else:
    return {
    "action": "escalate",
    "analysis": analysis,
    "reason": f"Unknown fix type: {fix_type}"
    }

```

Automated Rollbacks Based on Metrics

One of the most valuable CI/CD automations is intelligent rollback. Instead of waiting for a human to notice a degradation, an agent can watch key metrics and act:

```

class RollbackAgent:
    """Monitors deployments and triggers rollbacks when needed."""

    ROLLBACK_THRESHOLDS = {
        "error_rate": {"warn": 0.01, "critical": 0.05},
        "latency_p99_ms": {"warn": 500, "critical": 1000},
        "success_rate": {"warn": 0.99, "critical": 0.95}
    }

    def monitor_deployment(self, service: str,
                          deployment_id: str,
                          duration_minutes: int = 15):
        """Monitor a deployment for the bake time, rollback if issues
        detected."""

        start_time = datetime.utcnow()
        baseline = self.get_baseline_metrics(service)

        while datetime.utcnow() - start_time <
            timedelta(minutes=duration_minutes):
            current = self.get_current_metrics(service)

            # Compare against baseline and absolute thresholds
            issues = self.detect_issues(baseline, current)

            if issues["critical"]:
                self.execute_rollback(service, deployment_id, issues)
                return {
                    "status": "rolled_back",
                    "reason": issues["critical"],
                    "deployment_id": deployment_id
                }

            if issues["warning"]:
                self.notify_team(f"Warning: {service} showing degradation",
                                issues)

```

```

time.sleep(30) # Check every 30 seconds

    return {"status": "success", "deployment_id": deployment_id}

def execute_rollback(self, service: str, deployment_id: str,
                    issues: dict):
    """Execute rollback and notify team."""

    # Get previous stable version
    previous_version = self.get_previous_stable_version(service)

    # Execute rollback
    if self.deployment_type == "kubernetes":
        result = subprocess.run([
            "kubectl", "rollout", "undo",
            f"deployment/{service}",
            "--namespace", self.namespace
        ], capture_output=True, text=True)
    elif self.deployment_type == "ecs":
        result = self.ecs_rollback(service, previous_version)

    # Record the rollback
    self.record_rollback(service, deployment_id, previous_version,
                        issues)

    # Notify team
    self.notify_team(
        f"🚨 Auto-rollback: {service}",
        f"Deployment {deployment_id} was rolled back due to:
        {issues['critical']}\n"
        f"Reverted to version: {previous_version}"
    )

    return result

```

The notification to the team includes context that humans would need to investigate:

```

def format_rollback_notification(self, service, deployment_id,
                                issues, metrics):
    """Format a detailed rollback notification."""

    return f"""
🚨 **Automatic Rollback Executed**

**Service:** {service}
**Deployment:** {deployment_id}
**Time:** {datetime.utcnow().isoformat()}

**Trigger Conditions:**
{self.format_issues(issues)}

**Metrics Comparison:**
| Metric | Baseline | Current | Threshold |
|-----|-----|-----|-----|
| Error Rate | {metrics['baseline']['error_rate']:.2%} |
{metrics['current']['error_rate']:.2%} |
{self.ROLLBACK_THRESHOLDS['error_rate']['critical']:.2%} |
| P99 Latency | {metrics['baseline']['latency_p99_ms']}ms |
{metrics['current']['latency_p99_ms']}ms |
{self.ROLLBACK_THRESHOLDS['latency_p99_ms']['critical']}ms |

**Actions Required:**
1. Review the failed deployment commit
2. Check for regressions in the test suite
3. Validate fix before re-deploying

**Related Links:**
- [Deployment Logs]({self.get_deployment_logs_url(

```

```
deployment_id}))
- [Metrics Dashboard]({self.get_dashboard_url(service)})
- [Commit]({self.get_commit_url(deployment_id)})
"""
```

PR Review Assistance

Code review is time-consuming but essential. An agent can provide a first pass, catching obvious issues and letting humans focus on architecture and logic:

```
class PRReviewAgent:
    """Automated PR review assistant."""

    def review_pr(self, repo: str, pr_number: int) -> dict:
        """Perform initial review of a pull request."""

        pr = self.github.get_repo(repo).get_pull(pr_number)
        files = pr.get_files()

        review_results = {
            "security": [],
            "performance": [],
            "style": [],
            "bugs": [],
            "suggestions": [],
            "questions": []
        }

        for file in files:
            if file.filename.endswith(('.py', '.js', '.ts', '.go')):
                file_review = self.review_file(file)
                for category, items in file_review.items():
                    review_results[category].extend(items)

            # Check for missing tests
            test_coverage = self.check_test_coverage(pr)
            if test_coverage["missing"]:
                review_results["suggestions"].append({
                    "type": "missing_tests",
                    "files": test_coverage["missing"],
                    "message": "These files have changes but no corresponding test
changes"
                })

            # Check for documentation updates
            if self.needs_documentation_update(pr):
                review_results["suggestions"].append({
                    "type": "documentation",
                    "message": "This PR modifies public APIs but doesn't update
documentation"
                })

            # Generate summary
            review_results["summary"] =
                self.generate_review_summary(review_results)

        return review_results

    def review_file(self, file) -> dict:
        """Review a single file for issues."""

        prompt = f"""
Review this code change for a pull request:

Filename: {file.filename}
"""
```

Changes:

```
{file.patch} ```
```

Look for: 1. Security issues (SQL injection, XSS, secrets in code, etc.) 2. Performance concerns (N+1 queries, unnecessary loops, etc.) 3. Potential bugs (null handling, edge cases, race conditions) 4. Style issues (naming, structure, readability)

For each issue found, provide: - Line number - Category - Severity (critical/high/medium/low) - Description - Suggested fix

Respond in JSON format. """

```
response = self.llm.chat.completions.create( model="gpt-4-turbo", messages=[{"role": "user", "content": prompt}], response_format={"type": "json_object"} )
```

```
return json.loads(response.choices[0].message.content)
```

```
## Incident Response Automation
```

When deployments go wrong, every minute counts. An agent can handle the initial response while humans are being paged:

```
```python
class IncidentResponseAgent:
 """Automated first-responder for deployment incidents."""

 def handle_incident(self, alert: dict) -> dict:
 """Initial incident response automation."""

 incident = {
 "id": str(uuid.uuid4()),
 "started_at": datetime.utcnow().isoformat(),
 "alert": alert,
 "timeline": []
 }

 # Step 1: Acknowledge and classify
 incident["timeline"].append({
 "time": datetime.utcnow().isoformat(),
 "action": "incident_started",
 "details": "Automated incident response initiated"
 })

 # Step 2: Gather context
 context = self.gather_context(alert)
 incident["context"] = context
 incident["timeline"].append({
 "time": datetime.utcnow().isoformat(),
 "action": "context_gathered",
 "details": f"Gathered metrics, logs, and recent changes"
 })

 # Step 3: Identify recent changes
 recent_deploys = self.get_recent_deployments(
 context["service"],
 hours=2
)
 incident["recent_deployments"] = recent_deploys

 # Step 4: Auto-mitigate if possible
```

```

 if self.can_auto_mitigate(alert, context):
 mitigation = self.execute_mitigation(alert, context)
 incident["mitigation"] = mitigation
 incident["timeline"].append({
 "time": datetime.utcnow().isoformat(),
 "action": "auto_mitigation",
 "details": mitigation
 })

 # Step 5: Create incident channel and page on-call
 incident["channel"] = self.create_incident_channel(incident)
 self.page_oncall(incident)

 # Step 6: Post initial summary
 self.post_incident_summary(incident)

 return incident

```

## Case Study: An Agent That Fixes Failing Tests

This is one of our most successful agent implementations. Let me walk you through how it works.

**The Problem:** Our test suite had a 15% flaky test rate. Engineers would often just re-run failed tests, hoping they'd pass. When tests genuinely failed, it took an average of 45 minutes to diagnose and fix simple issues.

**The Solution:** We built a test-fixing agent that: 1. Detects test failures 2. Analyzes the failure pattern 3. Attempts fixes for common issues 4. Creates PRs for successful fixes

```

class TestFixerAgent:
 """Agent that automatically fixes common test failures."""

 FIXABLE_PATTERNS = [
 "assertion_error",
 "timeout",
 "import_error",
 "fixture_missing",
 "mock_misconfigured",
 "async_timing"
]

 def analyze_test_failure(self, test_file: str,
 test_name: str,
 error_output: str) -> dict:
 """Analyze a test failure and determine if it's fixable."""

 # Read the test file
 test_code = open(test_file).read()

 # Read the code being tested
 tested_module = self.get_tested_module(test_file)
 tested_code = open(tested_module).read() if tested_module else ""

 prompt = f"""
 Analyze this test failure:

 ## Test file: {test_file}
        ```python
        {test_code}

```

```

## Tested module: {tested_module} python {tested_code[:3000]}

## Error output: {error_output}

Determine: 1. Failure category: {self.FIXABLE_PATTERNS} 2. Root cause 3. Is the test wrong,
or is the code wrong? 4. If test is wrong, provide the fix 5. Confidence (high/medium/low)

Important: Only suggest test fixes, never production code fixes. ""

return self.llm_analyze(prompt)

def fix_test(self, analysis: dict, test_file: str) -> dict: ""Apply the suggested fix and verify it
works.""

if analysis["confidence"] != "high": return {"status": "skipped", "reason": "Low confidence"}

if analysis["fix_target"] != "test": return {"status": "skipped", "reason": "Would modify
production code"}

# Create a branch branch_name = f"auto-fix/{analysis['test_name']}-{uuid.uuid4( ).hex[:8]}"
self.git_create_branch(branch_name)

# Apply the fix original_content = open(test_file).read() fixed_content =
self.apply_fix(original_content, analysis["fix"])

with open(test_file, 'w') as f: f.write(fixed_content)

# Run the specific test result = subprocess.run( ["pytest", test_file, "-k", analysis["test_name"],
"-v"], capture_output=True, text=True )

if result.returncode == 0: # Test passes - create PR self.git_commit(test_file, f"fix: auto-fix
{analysis['test_name']}") self.git_push(branch_name) pr =
self.create_pull_request(branch_name, analysis)

return { "status": "fixed", "branch": branch_name, "pr_url": pr.html_url, "analysis": analysis }
else: # Fix didn't work - rollback with open(test_file, 'w') as f: f.write(original_content)
self.git_delete_branch(branch_name)

return { "status": "failed", "reason": "Fix didn't make test pass", "output": result.stdout +
result.stderr }

```

****Real Example Fix:****

Here's an actual test failure the agent fixed:

```

```python
Original failing test
def test_user_creation_timestamp():
 user = User.create(name="Test")
 assert user.created_at == datetime.now() # Fails due to timing

Agent's fix
def test_user_creation_timestamp():
 before = datetime.now()
 user = User.create(name="Test")

```

```
after = datetime.now()
assert before <= user.created_at <= after # Properly handles
 timing
```

### Results after 3 months:

Metric	Before	After
Average time to fix simple test failures	45 minutes	2 minutes (auto)
Tests auto-fixed per week	0	12-15
Developer hours saved weekly	0	~8 hours
False positive rate	N/A	3% (easy to reject)

**Key constraints we learned to implement:** 1. **Never modify production code** - Only test files 2. **Require human merge** - PRs need approval, even obvious ones 3. **One fix per PR** - Easier to review and revert 4. **Include analysis** - PR description explains the reasoning 5. **Tag flaky patterns** - Build a database of common issues

---

# Chapter 14: Observability and Incident Response

At 3 AM, when your pager goes off, the last thing you want to do is sift through thousands of log lines trying to understand what's happening. This is where AI agents can be your greatest ally—not replacing your expertise, but amplifying it.

## Log Analysis at Scale

Modern applications generate enormous amounts of logs. An agent can be your first line of analysis:

```
class LogAnalysisAgent:
 """Agent for intelligent log analysis and pattern detection."""

 def __init__(self, elasticsearch_client, llm_client):
 self.es = elasticsearch_client
 self.llm = llm_client
 self.known_patterns = self.load_known_patterns()

 def analyze_timerange(self, service: str,
 start_time: datetime,
 end_time: datetime) -> dict:
 """Analyze logs for a time range, surfacing important
 patterns."""

 # Query logs
 logs = self.es.search(
 index=f"logs-{service}-*",
 body={
 "query": {
 "bool": {
 "must": [
 {"range": {"@timestamp": {
 "gte": start_time.isoformat(),
 "lte": end_time.isoformat()
 }}}
]
 }
 },
 "size": 10000,
 "sort": [{"@timestamp": "asc"}]
 }
)

 # Extract patterns
 error_patterns = self.extract_error_patterns(logs)
 anomalies = self.detect_anomalies(logs)

 # Analyze with LLM
 analysis = self.llm_analyze_logs(logs, error_patterns,
 anomalies)

 return {
 "service": service,
 "timerange": {"start": start_time, "end": end_time},
 "log_count": len(logs["hits"]["hits"]),
 }
```

```

"error_patterns": error_patterns,
"anomalies": anomalies,
"analysis": analysis
}

def llm_analyze_logs(self, logs: dict,
 errors: list,
 anomalies: list) -> dict:
 """Use LLM to generate insights from log patterns."""

 # Sample logs intelligently - include errors and context
 sampled_logs = self.intelligent_sample(logs, errors)

 prompt = f"""
Analyze these application logs and identify:
1. Root cause of any errors
2. Sequence of events leading to issues
3. Correlations between different log entries
4. Recommended investigation steps

Error patterns found:
{json.dumps(errors, indent=2)}

Anomalies detected:
{json.dumps(anomalies, indent=2)}

Sample logs with context:
{sampled_logs}

Provide analysis in JSON format with:
- summary: one-sentence description
- root_cause: likely root cause
- timeline: sequence of events
- affected_components: list of impacted parts
- recommended_actions: prioritized list of next steps
"""

 response = self.llm.chat.completions.create(
 model="gpt-4-turbo",
 messages=[{"role": "user", "content": prompt}],
 response_format={"type": "json_object"}
)

 return json.loads(response.choices[0].message.content)

```

## Automated Runbook Execution

Runbooks are the documented steps for handling known issues. An agent can execute them automatically:

```

runbook_definitions.yaml
runbooks:
 - id: high_memory_usage
trigger:
 metric: container_memory_usage_percent
 condition: "> 90"
 duration: 5m
steps:
 - name: gather_info
 action: exec
 command: "kubectl top pods -n {{namespace}}"

 - name: check_for_leaks
 action: exec
 command: "kubectl exec {{pod}} -- cat /proc/meminfo"

```

```

- name: check_heap_dump
 action: conditional
 condition: "{{memory_growth_rate}} > 10"
 then:
- action: exec
 command: "kubectl exec {{pod}} -- jcmd 1 GC.heap_dump
 /tmp/heap.hprof"
- action: copy
 from: "{{pod}}:/tmp/heap.hprof"
 to: "s3://incident-artifacts/{{incident_id}}/heap.hprof"

- name: restart_if_critical
 action: conditional
 condition: "{{memory_percent}} > 95"
 approval: required
 then:
- action: exec
 command: "kubectl rollout restart deployment/{{service}} -n
 {{namespace}}"

 - id: database_connection_exhausted
trigger:
 log_pattern: "connection pool exhausted"
 count: 5
 window: 1m
steps:
 - name: check_connection_count
 action: exec
 command: "psql -c 'SELECT count(*) FROM pg_stat_activity'"

 - name: identify_long_queries
 action: exec
 command: |
psql -c "SELECT pid, now() - pg_stat_activity.query_start AS
duration, query
FROM pg_stat_activity
WHERE state != 'idle'
ORDER BY duration DESC
LIMIT 10"

 - name: kill_long_queries
 action: conditional
 condition: "any(query.duration > '10 minutes')"
 approval: required
 then:
- action: exec
 command: "psql -c 'SELECT pg_terminate_backend({{pid}})'"

```

Here's the agent that executes these runbooks:

```

class RunbookAgent:
 """Executes predefined runbooks in response to incidents."""

 def __init__(self, runbook_path: str):
 self.runbooks = yaml.safe_load(open(runbook_path))
 self.execution_history = []

 def execute_runbook(self, runbook_id: str, context: dict) ->
 dict:
 """Execute a runbook with given context."""

 runbook = self.find_runbook(runbook_id)
 execution = {
 "runbook_id": runbook_id,
 "started_at": datetime.utcnow().isoformat(),
 "context": context,
 "steps": []

```

```

 }

 for step in runbook["steps"]:
 step_result = self.execute_step(step, context)
 execution["steps"].append(step_result)

 # Update context with step results
 context[step["name"]] = step_result

 # Stop if step requires approval
 if step_result.get("awaiting_approval"):
 execution["status"] = "awaiting_approval"
 execution["pending_step"] = step["name"]
 return execution

 # Stop if step failed and runbook doesn't allow continuation
 if step_result["status"] == "failed" and not
 step.get("continue_on_failure"):
 execution["status"] = "failed"
 execution["failed_step"] = step["name"]
 return execution

 execution["status"] = "completed"
 execution["completed_at"] = datetime.utcnow().isoformat()

 return execution

def execute_step(self, step: dict, context: dict) -> dict:
 """Execute a single runbook step."""

 action = step["action"]

 if action == "exec":
 command = self.interpolate(step["command"], context)
 result = subprocess.run(
 command, shell=True, capture_output=True, text=True
)
 return {
 "name": step["name"],
 "status": "success" if result.returncode == 0 else "failed",
 "output": result.stdout,
 "error": result.stderr
 }

 elif action == "conditional":
 condition_met = self.evaluate_condition(step["condition"],
 context)
 if condition_met:
 if step.get("approval") == "required":
 return {
 "name": step["name"],
 "status": "pending",
 "awaiting_approval": True,
 "condition_met": True,
 "proposed_actions": step["then"]
 }

 # Execute the 'then' steps
 for sub_step in step["then"]:
 self.execute_step(sub_step, context)
 return {
 "name": step["name"],
 "status": "success",
 "condition_met": condition_met
 }

 # ... handle other action types

```

# Incident Summarization and RCA Drafting

After the dust settles, someone has to write the post-incident report. An agent can draft it:

```
class IncidentSummarizationAgent:
 """Drafts post-incident reports from incident data."""

 def generate_postmortem(self, incident: dict) -> str:
 """Generate a post-mortem document draft."""

 # Gather all relevant data
 timeline = self.build_timeline(incident)
 metrics = self.gather_metrics(incident)
 logs = self.gather_relevant_logs(incident)
 actions_taken = incident.get("actions", [])

 prompt = f"""
 Generate a post-incident report (postmortem) from this
 incident data:

 ## Incident Overview
 - ID: {incident['id']}
 - Started: {incident['started_at']}
 - Resolved: {incident.get('resolved_at', 'Ongoing')}
 - Severity: {incident['severity']}
 - Affected Services: {' '.join(incident['affected_services'])}

 ## Timeline
 {json.dumps(timeline, indent=2)}

 ## Key Metrics During Incident
 {json.dumps(metrics, indent=2)}

 ## Actions Taken
 {json.dumps(actions_taken, indent=2)}

 ## Relevant Log Excerpts
 {logs}

 Generate a postmortem with these sections:
 1. Executive Summary (2-3 sentences)
 2. Impact (users affected, duration, business impact)
 3. Timeline (key events with timestamps)
 4. Root Cause Analysis (what went wrong and why)
 5. Resolution (how it was fixed)
 6. Action Items (preventive measures, categorized as P0/P1/P2)
 7. Lessons Learned

 Write in a blameless tone. Focus on systems and processes, not
 individuals.
 """

 response = self.llm.chat.completions.create(
 model="gpt-4-turbo",
 messages=[{"role": "user", "content": prompt}]
)

 return response.choices[0].message.content
```

## On-Call Assistance

The on-call experience is where agents can have the biggest quality-of-life impact:

```

class OnCallAssistantAgent:
 """AI assistant for on-call engineers."""

 def handle_alert(self, alert: dict) -> dict:
 """Process an incoming alert and provide context."""

 # Enrich alert with context
 context = {
 "alert": alert,
 "service_info": self.get_service_info(alert["service"]),
 "recent_deployments":
 self.get_recent_deployments(alert["service"]),
 "similar_incidents": self.find_similar_incidents(alert),
 "relevant_runbooks": self.find_relevant_runbooks(alert),
 "current_metrics": self.get_current_metrics(alert["service"]),
 "on_call_roster": self.get_on_call_roster(alert["service"])
 }

 # Generate initial analysis
 analysis = self.analyze_alert(context)

 # Format for Slack
 message = self.format_oncall_message(alert, context, analysis)

 return {
 "context": context,
 "analysis": analysis,
 "suggested_actions": analysis.get("suggested_actions", []),
 "confidence": analysis.get("confidence", "medium"),
 "slack_message": message
 }

 def answer_question(self, question: str, incident_context: dict)
 -> str:
 """Answer on-call engineer's questions about the incident."""

 prompt = f"""
 You are an AI assistant helping an on-call engineer during an
 incident.

 ## Current Incident Context
 {json.dumps(incident_context, indent=2)}

 ## Engineer's Question
 {question}

 Provide a helpful, accurate answer. If you don't know
 something, say so.
 Include relevant commands or queries when appropriate.
 """

 response = self.llm.chat.completions.create(
 model="gpt-4-turbo",
 messages=[{"role": "user", "content": prompt}]
)

 return response.choices[0].message.content

```

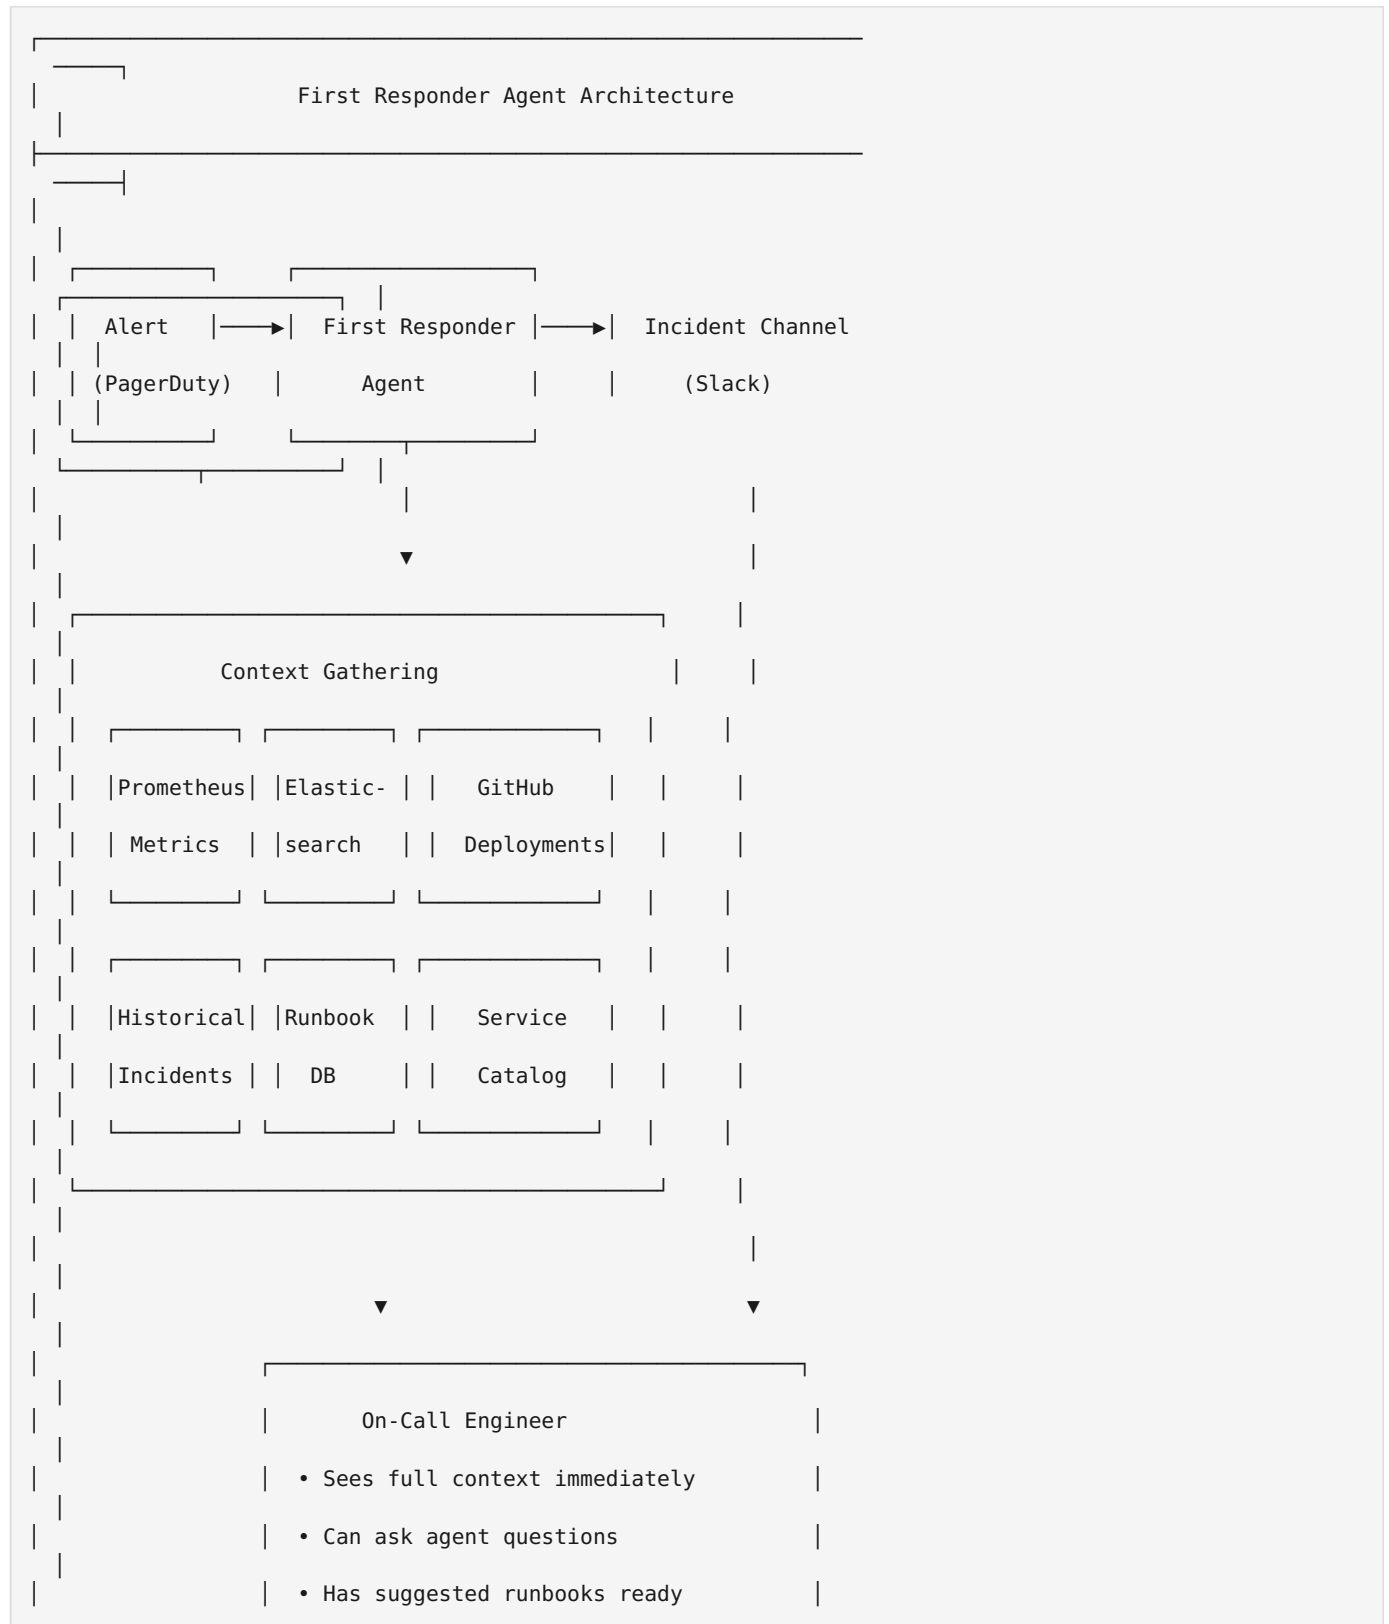
## Case Study: An Agent as Your First Responder

Let me share the most impactful agent we've built: an incident first-responder that has reduced our mean time to resolution (MTTR) by 40%.

**The Problem:** Our on-call rotation had 6 engineers. Response times varied wildly—from 2 minutes for the most experienced to 30+ minutes for newer team members. The first 15 minutes of any incident were usually spent gathering context.

**The Solution:** An agent that: 1. Receives every alert 2. Immediately gathers relevant context 3. Posts a comprehensive briefing to the incident channel 4. Suggests initial investigation steps 5. Can answer questions from the on-call engineer

Here's the architecture:



## The Implementation:

```
class FirstResponderAgent:
 """
 AI-powered first responder for incidents.
 Gathers context and provides briefings within seconds of an
 alert.
 """

 def __init__(self):
 self.prometheus = PrometheusClient()
 self.elasticsearch = ElasticsearchClient()
 self.github = GithubClient()
 self.pagerduty = PagerDutyClient()
 self.slack = SlackClient()
 self.llm = OpenAI()
 self.vector_store = VectorStore() # For similar incident
 search

 def on_alert(self, alert: dict):
 """Handle incoming alert - the main entry point."""

 start_time = time.time()

 # Create incident channel
 channel = self.slack.create_channel(
 f"inc-{alert['service']}-{datetime.now().strftime('%m%d-%H%M')}"
)

 # Gather context in parallel
 with ThreadPoolExecutor(max_workers=6) as executor:
 futures = {
 'metrics': executor.submit(self.get_metrics_context, alert),
 'logs': executor.submit(self.get_log_context, alert),
 'deployments': executor.submit(self.get_deployment_context,
 alert),
 'similar': executor.submit(self.find_similar_incidents, alert),
 'runbooks': executor.submit(self.find_runbooks, alert),
 'service_info': executor.submit(self.get_service_info, alert)
 }

 context = {k: f.result() for k, f in futures.items()}

 # Generate analysis
 analysis = self.analyze_incident(alert, context)

 # Post briefing
 briefing = self.format_briefing(alert, context, analysis)
 self.slack.post_message(channel, briefing)

 # Log timing
 elapsed = time.time() - start_time
 self.log_metric('first_response_time_seconds', elapsed)

 # Return for potential further processing
 return {
 "channel": channel,
 "context": context,
 "analysis": analysis,
 "response_time_seconds": elapsed
 }
```

```

}

def format_briefing(self, alert, context, analysis) -> str:
 """Format the incident briefing for Slack."""

 return f"""
🚨 Incident Alert: {alert['name']}*
Service: {alert['service']} | *Severity:* {alert['severity']}
 | *Time:* {alert['timestamp']}

📊 Quick Status:
• Error rate: {context['metrics']['error_rate']:.2%} (normal: <1%)
• P99 latency: {context['metrics']['latency_p99']}ms (normal: <200ms)
• Active instances: {context['metrics']['instance_count']}

🔍 AI Analysis:
{analysis['summary']}

Likely cause: {analysis['likely_cause']}
Confidence: {analysis['confidence']}

📦 Recent Changes:
{self.format_deployments(context['deployments'])}

📖 Similar Past Incidents:
{self.format_similar_incidents(context['similar'])}

📄 Suggested Runbooks:
{self.format_runbooks(context['runbooks'])}

💡 Suggested First Steps:
{self.format_suggested_actions(analysis['suggested_actions'])}

Ask me questions in this channel - I'm here to help investigate.
Type ` /incident-action ` to execute runbook steps.
"""

def answer_question(self, channel: str, question: str):
 """Respond to questions from on-call engineers in the incident channel."""

 # Get incident context from channel
 incident = self.get_incident_from_channel(channel)

 # Build conversation context
 messages = [
 {"role": "system", "content": """
You are an AI assistant helping with an active incident.
You have access to logs, metrics, and service documentation.
Be concise, accurate, and helpful. If you're not sure, say so.
Include specific commands, queries, or links when relevant.
"""},
 {"role": "user", "content": f"""
Incident context: {json.dumps(incident, indent=2)}


Engineer's question: {question}
"""}
]


 response = self.llm.chat.completions.create(
 model="gpt-4-turbo",
 messages=messages
)

```


```
self.slack.post_message(channel,
response.choices[0].message.content)
```

## Real incident example:


 Incident Alert: High Error Rate - Payment Service  
Service: payment-service | Severity: SEV1 | Time: 2024-01-15  
03:42 UTC

 Quick Status:

- Error rate: 23.4% (normal: <1%)
- P99 latency: 2,340ms (normal: <200ms)
- Active instances: 3/3

 AI Analysis:  
Error spike correlates with deployment 15 minutes ago. Logs show "connection refused" errors to payment-gateway.internal. The new version appears to have a misconfigured service endpoint.


Likely cause: Configuration error in deployment v2.34.1 - wrong payment gateway endpoint.  
Confidence: High

 Recent Changes:


- v2.34.1 deployed 15 mins ago by @alice - "Update payment gateway integration"
- No infrastructure changes in last 24h

 Similar Past Incidents:

- INC-2023-0892 (3 months ago): Payment service config error - resolved by rollback
- INC-2023-0654 (6 months ago): Gateway connectivity - resolved by config fix

 Suggested Runbooks:

- [RB-PAY-001] Payment Service Rollback
- [RB-PAY-003] Payment Gateway Connectivity Check

 Suggested First Steps:

1. Check payment gateway endpoint in current config  
kubectl get configmap payment-config -o yaml
2. Compare with previous version  
kubectl rollout history deployment/payment-service
3. If config is wrong, rollback immediately  
kubectl rollout undo deployment/payment-service

## Results after 6 months:

Metric	Before	After	Improvement
Mean Time to Resolution	47 min	28 min	40% faster
Time to First Response	8 min	45 sec	90% faster
Context gathering time	15 min	0 (automated)	Eliminated
Junior engineer MTTR	65 min	32 min	51% faster
Incident escalation rate	34%	22%	35% fewer escalations

## Key design decisions that made this work:

1. **Speed matters most** - Everything runs in parallel. The agent must post within 60 seconds.
2. **Context over commands** - The agent provides context and suggestions, not commands. Humans make decisions.
3. **Learn from history** - Vector similarity search against past incidents is surprisingly effective.
4. **Conversational interface** - Engineers can ask follow-up questions naturally.
5. **Integration is everything** - The agent pulls from 8 different systems. Each integration took work, but pays off constantly.
6. **Graceful degradation** - If any context source fails, the agent still posts what it has.

## What we're still working on:

- Automated correlation across multiple services
- Predicting incidents before they happen
- Automated remediation for Tier-1 issues
- Knowledge extraction from resolved incidents to improve future responses

The first-responder agent has become so embedded in our incident response that engineers now feel something is wrong when it doesn't respond. That's the sign of a successful tool—not when people notice it, but when they notice its absence.

---

## Part 4 Summary

These three chapters have shown AI agents in their most practical form: doing real work that saves real time. Infrastructure automation, CI/CD assistance, and incident response are just the beginning.

The common thread across all these applications: - **Agents augment humans, not replace them** - **Start with read-only operations, earn trust, then expand** - **Logging and auditability are non-negotiable** - **The best agents surface information and suggest actions** - **Integration with existing tools is more valuable than building new ones**

In Part 5, we'll look at where this is all heading—building agent teams, governance at scale, and the future of human-agent collaboration.

---

# PART 5: THE FUTURE

---

## Chapter 15: Building Your Agent Team

If you've followed along through the previous chapters, you've built, deployed, crashed, fixed, and optimized individual agents. You've learned their quirks, their failure modes, and their surprising capabilities. Now it's time to think bigger.

Running a single agent is like having one extremely talented intern. Useful, sometimes brilliant, but limited. Building an agent team is like scaling from a solo founder to an engineering organization. The challenges shift from "how do I make this work" to "how do I make these work together."

### From Single Agent to Agent Organization

The journey typically starts innocently enough. You build a deployment agent. It works well, so you build a monitoring agent. Then an incident response agent. Before you know it, you have five agents running different aspects of your infrastructure.


And then chaos ensues.

Agent A reads a log file that Agent B is in the middle of writing. Agent C tries to restart a service that Agent D just deployed. Agent E summarizes an incident that Agents A through D are still actively working on. You've accidentally created a distributed system without any of the coordination primitives that distributed systems require.

The first realization: **agents are microservices with opinions**. They share many characteristics with traditional microservices—they need service discovery, coordination, shared state management, and clear boundaries. But unlike microservices, they can improvise. An agent that decides to "help" another agent is a recipe for disaster.

Here's how we evolved our agent architecture at scale:

**Phase 1: Isolated Agents** Each agent runs independently with its own context, memory, and tools. No coordination, no shared state. Simple, but limited.



**Phase 2: Shared Resources** Agents share certain resources—a common memory store, shared tool access. Coordination happens through the resources themselves (file locks, database transactions).

**Phase 3: Orchestrated Teams** An orchestrator agent manages a team of specialist agents, delegating tasks and coordinating results. This is where things get interesting—and complicated.

## Specialization: Different Agents for Different Tasks

Not all agents are created equal, and they shouldn't be. The temptation is to build one super-agent that can do everything. Resist this temptation.

A generalist agent that can deploy code, analyze logs, respond to incidents, and write documentation will be mediocre at all of them. Its context window fills with irrelevant instructions. Its prompts become unwieldy. Its mental model becomes confused.

Instead, build specialist agents:

**The Deployment Agent** - Knows your deployment pipeline intimately - Has access to kubectl, helm, and CI/CD tools - Understands rollback procedures - Context focused entirely on deployment patterns

**The Monitoring Agent** - Expert in Prometheus, Grafana, and alerting - Can analyze metrics patterns - Understands baseline behaviors - Focused context on observability

**The Incident Response Agent** - First responder training - Access to runbooks and escalation procedures - Can coordinate across teams - Context optimized for rapid triage

**The Code Review Agent** - Understands your codebase patterns - Knows your style guidelines - Can suggest improvements - Context loaded with code examples

The key insight: **specialization allows for optimized context**. Each agent's context window is precious real estate. A deployment agent doesn't need to know how to analyze logs—that's someone else's job.

Here's a practical specialization we implemented:

```
agent-team.yaml
team: infrastructure
orchestrator: infra-coordinator

specialists:
 - name: deploy-agent
model: claude-opus-4-5
context_files:
 - deployment-patterns.md
 - kubernetes-reference.md
 - rollback-procedures.md
tools:
 - kubectl
 - helm
 - argocd
```

```

- name: monitor-agent
model: claude-sonnet-4
context_files:
- metrics-baselines.md
- alert-thresholds.md
- runbooks/
tools:
- promql
- grafana-api
- alertmanager

- name: incident-agent
model: claude-opus-4-5
context_files:
- incident-playbook.md
- escalation-matrix.md
- postmortem-template.md
tools:
- pagerduty
- slack
- status-page

```

## Communication Between Agents

Here's where things get philosophical: how should agents talk to each other?

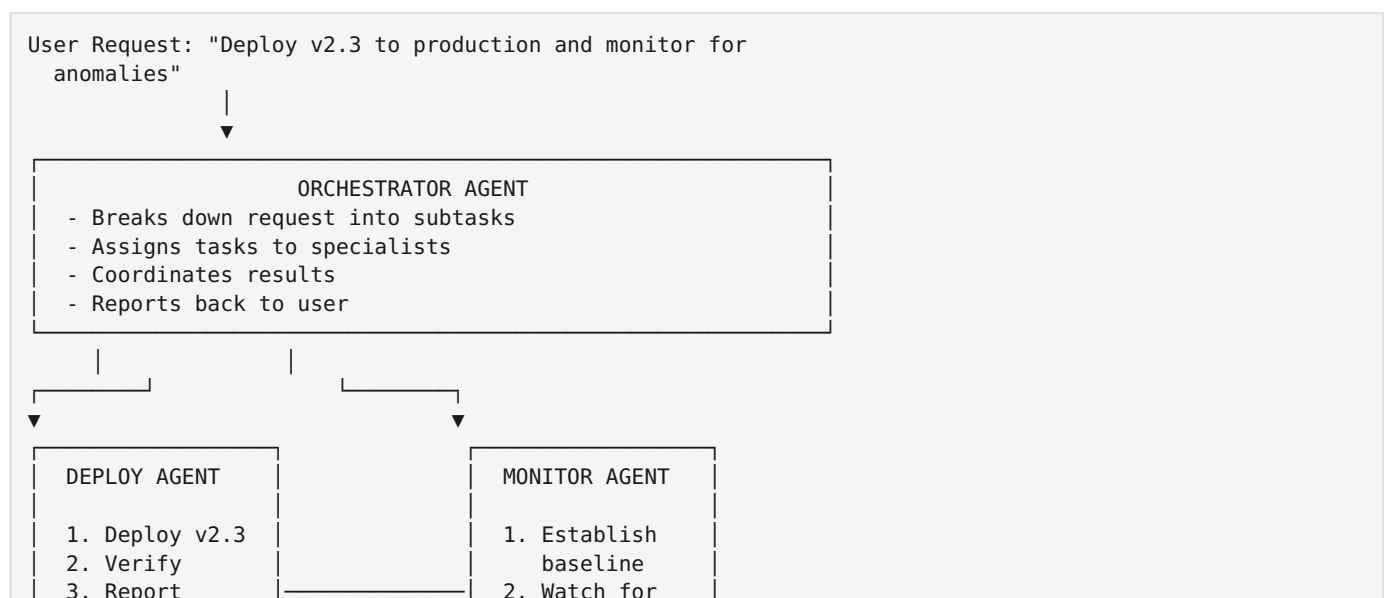
Option 1: **Shared Memory** Agents read and write to a common memory store. Simple, but prone to conflicts and stale reads.

Option 2: **Message Passing** Agents send explicit messages to each other. Cleaner boundaries, but requires a message bus and protocol design.

Option 3: **Orchestrator Mediation** Agents only communicate through an orchestrator. The orchestrator asks questions, receives answers, and distributes information. Most controlled, but the orchestrator becomes a bottleneck.

We've tried all three. Our conclusion: **start with orchestrator mediation, evolve to message passing as trust develops.**

The orchestrator pattern looks like this:



Completion Signal	anomalies 3. Report
----------------------	------------------------

The orchestrator's prompt is crucial:

You are the Infrastructure Coordinator. Your role is to:

1. Receive requests from users
2. Break complex requests into subtasks
3. Delegate subtasks to specialist agents
4. Coordinate between agents when needed
5. Synthesize results into coherent responses

Available specialists:

- deploy-agent: deployment, rollback, release management
- monitor-agent: metrics, alerts, observability
- incident-agent: incident response, escalation, communication

Rules:

- Never perform specialist tasks yourself
- Always verify completion before moving to dependent tasks
- Escalate to humans if specialists fail twice
- Provide progress updates for tasks > 5 minutes

## The War Story: Coordinating a Team of 4 Agents

Let me tell you about the time we tried to coordinate four agents on a complex migration task. It was a database migration: spin up new database, migrate data, update services, verify, cut over.

Simple plan, four specialists: 1. Infrastructure Agent: provision new database 2. Data Agent: handle migration 3. Deploy Agent: update services 4. Verify Agent: run tests and checks

What could go wrong?

Everything.

The Infrastructure Agent provisioned the database and reported success. But it forgot to mention that the security group rules weren't fully propagated yet. The Data Agent started the migration immediately, failed to connect, and decided to "fix" the issue by... provisioning a second database. Now we had two databases, neither with the right network configuration.

The Deploy Agent, seeing a "database ready" signal in the shared state, started updating services to point to the first database (which still didn't have the right security groups). Services started failing.

The Verify Agent detected the failures and declared the migration a failure. It then decided to "rollback" by deleting the new databases. Both of them.

The Data Agent, still trying to migrate data, now had nowhere to migrate to. It escalated to... the Deploy Agent, which was busy trying to figure out why services were failing.

Total chaos time: 23 minutes before a human intervened.

## Lessons learned:

1. **State machines, not signals.** Instead of agents watching for “done” signals, use explicit state machines. The migration is in exactly one state at any time, and transitions are explicit.
2. **Locks on shared resources.** The database being provisioned needed a lock. Only one agent should be able to modify infrastructure state at a time.
3. **Dependency graphs.** The Data Agent can’t start until the Infrastructure Agent completes AND the verification step passes. Not just “completes.”
4. **Single source of truth.** We added a “migration manifest” that all agents read from and only the orchestrator could write to.

Here’s the improved architecture:

```
migration:
 id: db-migration-2024-01-15
 state: provisioning # One of: provisioning, migrating,
 deploying, verifying, complete, failed, rolledback

 phases:
 - name: provision
 agent: infrastructure-agent
 status: in_progress
 started_at: 2024-01-15T10:00:00Z
 completed_at: null
 verification:
 - check: database_reachable
 status: pending
 - check: security_groups_active
 status: pending

 - name: migrate
 agent: data-agent
 status: blocked
 depends_on: [provision.verified]

 - name: deploy
 agent: deploy-agent
 status: blocked
 depends_on: [migrate.verified]
```

## Governance: Who Watches the Watchers

Here’s an uncomfortable truth: agents can fail in subtle ways that other agents can’t detect. An agent might successfully complete all its tasks while quietly making decisions that are technically correct but strategically wrong.

You need governance. Not just monitoring, but actual oversight.

**The Audit Trail** Every agent action should be logged with enough context to understand *why* the action was taken, not just *what* was done.

```
{
 "timestamp": "2024-01-15T10:05:23Z",
```

```
"agent": "deploy-agent",
"action": "scale_deployment",
"target": "api-server",
"before": {"replicas": 3},
"after": {"replicas": 5},
"reasoning": "Detected increased latency (p99: 450ms >
threshold: 200ms). Scaling up to handle load.",
"triggered_by": "monitor-agent alert",
"reversible": true,
"review_required": false
}
```

**Review Queues** Some decisions should queue for human review. The challenge is deciding which ones.

Our heuristic: **if reversing the action would cause an incident, it requires review.**

- Scaling up? Probably fine, auto-approve.
- Scaling down? Queue for review (might impact capacity).
- Updating config? Depends on the config.
- Modifying security rules? Always review.

**The Governor Agent** We experimented with a “governor” agent that reviews other agents’ actions. It has no execution privileges—it can only observe and recommend.

```
You are the Governor. Your role is oversight, not execution.

Review each action against these principles:
1. Least privilege: Did the agent use minimum necessary access?
2. Reversibility: Can this action be undone if wrong?
3. Blast radius: What's the worst case if this fails?
4. Pattern matching: Is this consistent with past decisions?

Flag for human review if:
- Blast radius > 1 service
- Action is irreversible
- Pattern is unusual for this agent
- Cost implication > $100
```

It works surprisingly well as a safety net, though it adds latency.

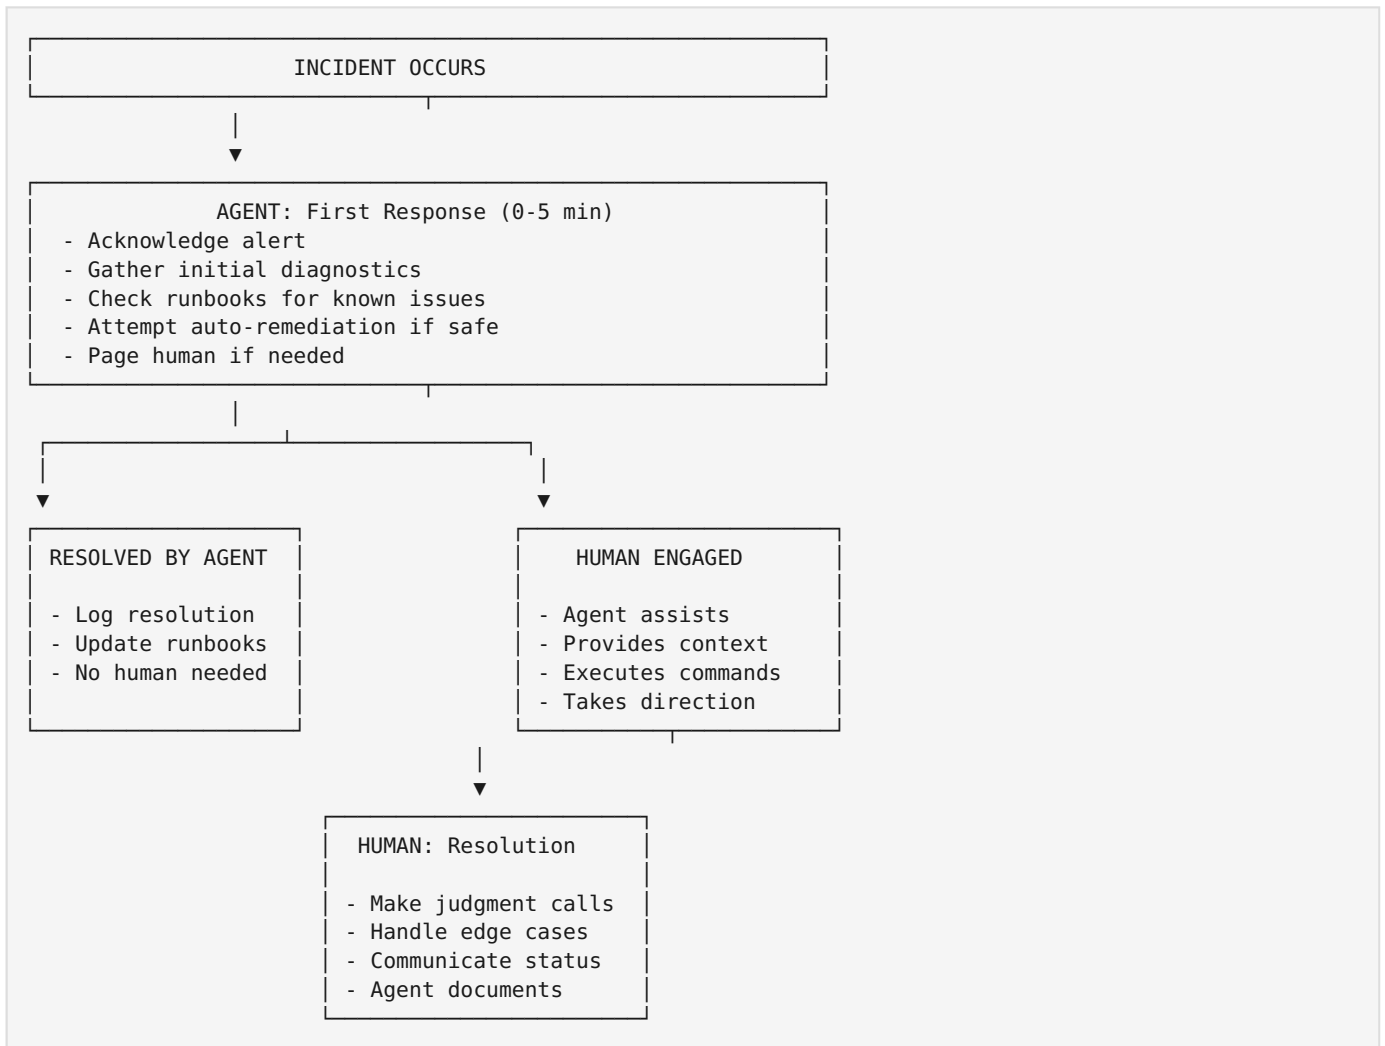
## The Hybrid Team: Humans and Agents Working Together

The goal isn’t to replace humans with agents. The goal is to build hybrid teams where humans and agents each do what they’re best at.

Agents excel at: - Repetitive tasks with clear procedures - Fast initial response (sub-second to seconds) - Remembering and following complex rules consistently - Operating at scale (monitoring hundreds of services) - Working 24/7 without fatigue

Humans excel at: - Novel situations without precedent - Stakeholder communication and judgment calls - Strategic decisions with organizational context - Tasks requiring empathy or politics - Overriding rules when the rules are wrong

The hybrid model we evolved to:



The key insight: **agents and humans should have clear handoff protocols.** When an agent escalates to a human, it provides: - Summary of what’s known - Actions already taken - Why escalation was needed - Suggested next steps

When a human hands back to an agent, they provide: - What was done manually - What the agent should continue doing - Any changes to the approach

## Building Your Agent Team: Practical Steps

If you’re ready to scale from one agent to a team, here’s the roadmap:

**Week 1-2: Identify Specializations** - Map your operational domains (deploy, monitor, incident, etc.) - Identify which tasks are currently handled by which tools/scripts - Design agent boundaries around these domains

**Week 3-4: Build the Orchestrator** - Start with a simple coordinator that routes requests - Don’t over-engineer—a basic routing prompt works initially - Add complexity only as needed

**Week 5-6: Implement Coordination Primitives** - Shared state store with proper locking - Event system for agent-to-agent signals - State machine for multi-phase operations

**Week 7-8: Add Governance** - Audit logging for all agent actions - Review queues for high-risk operations - Dashboard for agent activity visibility

**Week 9+: Iterate** - Monitor failure modes and adjust boundaries - Tune autonomy levels based on trust - Add new specialists as patterns emerge

The journey from single agent to agent team is a journey in distributed systems design. All the challenges you've faced building reliable microservices—coordination, state management, failure handling—apply here. But with agents, the components can think. Use that capability wisely, constrain it where needed, and build systems where humans and agents amplify each other's capabilities.

---

# Chapter 16: What's Next

Every technical book reaches a moment where it must acknowledge that the ground is shifting. This one is no exception. If you've implemented even half of what we've discussed, you're ahead of most organizations. But "ahead" in AI means months, not years. What's changing, and how do you prepare?

## The Trajectory: More Capable, More Autonomous

Let's be honest about where we are: AI agents in 2024-2025 are like web applications in 1999. Functional, sometimes impressive, but clearly early. The agents we've discussed require significant scaffolding—careful prompting, explicit tool definitions, human oversight, elaborate error handling.

The trajectory is clear, even if the timeline isn't:

**Short-term (1-2 years):** - Longer context windows (already seeing 200K+ tokens) - More reliable tool use with fewer hallucinations - Better multi-modal capabilities (images, audio, video) - Native computer use (agents that can operate GUIs) - Improved reasoning for multi-step tasks

**Medium-term (2-5 years):** - Agents that maintain state across sessions natively - Learned tool creation (agents building their own tools) - Collaborative agent protocols becoming standard - Tighter IDE and development environment integration - Specialized models for operational tasks

**Long-term (5+ years):** - Agents that genuinely learn from experience - Minimal prompting—agents that understand intent - Autonomous research and problem-solving - Human-level debugging and troubleshooting

What does this mean for you, practically? The patterns we've discussed—memory systems, tool use, orchestration, governance—will remain relevant. But the *implementation* will change. The scaffolding we build today will become native capabilities.

Here's the mindset shift: **build for abstraction, not for permanence**. Your agent wrappers, prompt templates, and orchestration code are temporary. Build them to be replaced.

## Emerging Patterns: MCP, Tool Ecosystems, Agent Marketplaces

Several emerging patterns will reshape how we build and deploy agents:

## Model Context Protocol (MCP)

MCP is an attempt to standardize how agents interact with tools and data sources. Instead of each agent framework inventing its own tool integration, MCP provides a common protocol.

The implications are significant: - Tools become portable across agent systems - Data sources get a standard integration path - Enterprise integration becomes simpler

If you're building tools today, consider MCP compatibility. The investment is small, and the portability is valuable.

```
MCP-style tool definition
{
 "name": "kubernetes_get_pods",
 "description": "List pods in a Kubernetes namespace",
 "input_schema": {
 "type": "object",
 "properties": {
 "namespace": {"type": "string", "default": "default"},
 "label_selector": {"type": "string", "optional": true}
 }
 },
 "output_schema": {
 "type": "array",
 "items": {"$ref": "#/definitions/Pod"}
 }
}
```

## Tool Ecosystems

We're moving from "build your own tools" to "compose existing tools." Tool marketplaces and registries are emerging where you can: - Find pre-built integrations for common services - Share tools with the community - Version and update tools independently of agents

This is analogous to the npm/pip/docker moment—when packaging and sharing became standardized. Position yourself to consume tools from ecosystems rather than building everything custom.

## Agent Marketplaces

Beyond tools, entire agents are becoming shareable. You might: - Subscribe to a "Kubernetes operator agent" maintained by specialists - Use an "incident response agent" trained on thousands of incidents - Compose agents from marketplace like Docker images

The implications for build vs. buy decisions are significant. Just as you don't write your own database, you might not write your own monitoring agent.

## Preparing Your Infrastructure for Agent Adoption

What should you be doing now to prepare for this future?

## 1. Instrumentability

Agents need data to operate. The more observable your systems, the more useful agents become.

**Action items:** - Structured logging everywhere (JSON, not free-form text) - Comprehensive metrics (not just the ones you alert on) - Configuration as data (agents can read and understand it) - Documentation as code (machine-parseable runbooks)

```
Good: Machine-readable runbook
runbook:
 name: high-memory-usage
 trigger:
 metric: container_memory_usage_bytes
 threshold: 0.9
 duration: 5m

 steps:
- name: identify_process
 command: "kubectl top pods -n {{namespace}}"
 parse: "highest_memory_pod"

- name: check_for_leak
 command: "kubectl logs {{pod}} --since=1h | grep -i 'out of
 memory'"
 success_means: "no_output"

- name: restart_if_needed
 condition: "step.check_for_leak.result != success"
 command: "kubectl rollout restart deployment/{{deployment}}"
```

## 2. API-First Everything

If a human does something through a GUI, an agent probably can't do it. Push toward API-first for all operational tasks.

**Action items:** - Expose all administrative functions as APIs - Use infrastructure-as-code (Terraform, Pulumi, etc.) - Prefer CLI tools with parseable output - Document APIs with OpenAPI/Swagger

## 3. Clear Boundaries and Permissions

Agents need access, but controlled access. Your IAM story needs to accommodate automated actors.

**Action items:** - Service accounts for agents (not shared human accounts) - Fine-grained permissions (agents get exactly what they need) - Audit logging for agent actions - Clear escalation paths when agents hit permission walls

## 4. Experimentation Infrastructure

You need safe places to experiment with agents before production.

**Action items:** - Sandbox environments that agents can break without impact - Production-like staging that agents can practice on - Feature flags for agent capabilities - Kill switches that work

## The Skills That Matter in an Agent-First World

Here's the uncomfortable question: if agents can do an increasing portion of operational work, what should *you* be learning?

Some skills become more valuable:

### System Design

Agents can execute, but they need systems worth executing on. Understanding distributed systems, failure modes, and architectural trade-offs becomes more important, not less.

### Prompt Engineering / Agent Architecture

Knowing how to design effective agents—their context, tools, constraints, and orchestration—is a distinct skill. It's part psychology, part system design, part writing.

### Security Thinking

More autonomous systems mean more attack surface. Understanding how agents can be exploited, manipulated, or misused becomes critical.

### Integration Architecture

Connecting agents to existing systems, designing APIs they can use, and building the glue code between automated and manual processes.

### Supervision and Governance

Someone needs to watch the watchers. Designing audit systems, governance frameworks, and escalation processes.

Some skills become less valuable (for routine work):

- Typing repetitive commands
- Remembering syntax and flags
- Manual log analysis
- First-pass troubleshooting for known issues
- Writing boilerplate code or configuration

But here's the thing: **the need for human judgment doesn't go away—it concentrates on harder problems.** You spend less time on routine tasks, more time on the tasks that genuinely require human insight.

# A Practical Adoption Roadmap

If you're starting your agent journey, or looking to accelerate it:

## Phase 1: Assist (Month 1-3)

- Deploy agents that help you, but require explicit approval for actions
- Start with low-risk domains (documentation, analysis, suggestions)
- Build trust through successful assists
- Invest in logging and observability

## Phase 2: Automate (Month 4-6)

- Graduate trusted agents to auto-approve for specific actions
- Expand tool access incrementally
- Implement review queues for edge cases
- Build muscle memory for agent supervision

## Phase 3: Orchestrate (Month 7-12)

- Move to multi-agent systems
- Build coordination primitives
- Implement governance frameworks
- Handle complex, multi-step workflows

## Phase 4: Optimize (Ongoing)

- Reduce costs through caching and model selection
- Improve reliability through better prompts and patterns
- Expand coverage to new domains
- Share learnings and tools with the community

## The War Story: Looking Back

Let me tell you about a conversation I had recently. A junior engineer asked me to describe my daily workflow from five years ago. I tried to explain: manually checking monitoring dashboards every morning, scanning logs for issues, copying commands from runbooks, updating tickets by hand.

"That sounds exhausting," she said. "Why didn't you automate that?"

I started to explain that we did automate parts of it, with scripts and alerts and... and then I realized the difference. Our automation was brittle. It handled the cases we anticipated. The moment something unexpected happened, humans were back in the loop.

Today's agents handle unexpected things. Not perfectly, not always correctly, but they *try*. They adapt. They reason about situations they haven't explicitly been trained for.

That's the shift. Not just automation, but adaptive automation. Systems that can think about systems.

## Your Journey Starts Now

I've shared war stories throughout this book—the runaway cron job, the \$500 month, the 3 AM debugging sessions, the agents that helped and the agents that hindered. These aren't warnings against adoption; they're roadmaps through the challenges.

Every new technology goes through this phase. The early adopters stumble, learn, share, and build the foundations for everyone else. You're now part of that group.

Here's what I hope you take away:

**Start small.** Your first agent should do one thing well. Expand from there.

**Fail safely.** Build in guardrails, kill switches, and human oversight. Trust is earned.

**Invest in foundations.** Memory systems, tool abstractions, and observability pay dividends as you scale.

**Stay humble.** Agents are tools. Powerful tools, but tools. They augment human capability; they don't replace human judgment.

**Share what you learn.** The community is still figuring this out together. Your war stories help everyone.

The infrastructure of the future will be managed by hybrid teams—humans and agents working together, each contributing what they're best at. You're building that future, one agent at a time.

Good luck. And when something goes wrong at 3 AM—and it will—remember: at least now you might have an agent to help debug it.

---

# APPENDICES

---

# Appendix A: Tool and Framework Reference

This appendix provides practical comparisons of the tools and frameworks you'll encounter when building AI agents for infrastructure work. These evaluations are based on real production use, not marketing materials.

## Agent Frameworks Compared

### LangChain / LangGraph

**Best for:** Complex agent workflows, multi-step reasoning chains, teams with Python expertise

**Strengths:** - Extensive ecosystem and integrations - Strong community and documentation - LangGraph adds proper workflow orchestration - Good observability with LangSmith

**Weaknesses:** - Can be over-engineered for simple use cases - Abstractions sometimes hide important details - Rapid API changes between versions - Performance overhead for simple tasks

**Our experience:** Started here, useful for prototyping, but moved to simpler approaches for production due to abstraction overhead.

```
LangGraph agent example
from langgraph.graph import StateGraph

workflow = StateGraph(AgentState)
workflow.add_node("research", research_agent)
workflow.add_node("execute", execute_agent)
workflow.add_edge("research", "execute")
app = workflow.compile()
```

### CrewAI

**Best for:** Multi-agent scenarios with role-based interactions

**Strengths:** - Clean abstraction for agent teams - Built-in coordination patterns - Good for simulating organizational structures - Active development

**Weaknesses:** - Less flexible than lower-level alternatives - Opinionated workflow structure - Can be limiting for novel coordination patterns

**Our experience:** Good for proof-of-concepts, but we needed more control over agent interactions in production.

### AutoGen (Microsoft)

**Best for:** Research scenarios, multi-agent conversations, code generation

**Strengths:** - Strong multi-agent conversation support - Good code execution sandboxing - Microsoft backing and integration path - Active research community

**Weaknesses:** - More research-oriented than production-ready - Can be complex to configure - Documentation assumes familiarity with concepts

## Claude Code / Claude Computer Use

**Best for:** Autonomous coding tasks, computer control, direct Anthropic API use

**Strengths:** - Native tool use without wrappers - Computer use capability for GUI automation - High-quality reasoning - Direct API means fewer abstraction layers

**Weaknesses:** - Anthropic-specific - Computer use still experimental - Less ecosystem around it compared to OpenAI

**Our experience:** Our primary production choice for complex reasoning tasks.

## OpenAI Assistants API

**Best for:** Stateful conversations, file handling, teams invested in OpenAI

**Strengths:** - Built-in memory/threading - Code interpreter capability - File search built-in - Simple API

**Weaknesses:** - OpenAI lock-in - Less control over orchestration - Cost can be higher due to stateful storage

## Direct API Calls (No Framework)

**Best for:** Simple agents, teams who want full control, performance-critical applications

**Strengths:** - Maximum control - Minimum overhead - Easy to understand and debug - No dependency on framework updates

**Weaknesses:** - You build everything yourself - No built-in patterns for common scenarios - More code to maintain

**Our experience:** Many production agents use direct API calls with thin custom wrappers. Don't underestimate simplicity.

## Framework Decision Matrix

Factor	LangChain	CrewAI	AutoGen	Direct API
Complexity tolerance	High	Medium	High	Low
Team Python expertise	Required	Required	Required	Any
Multi-agent needs	LangGraph	Native	Native	Custom

Factor	LangChain	CrewAI	AutoGen	Direct API
Production readiness	Good	Emerging	Research	Depends
Debugging ease	Medium	Good	Medium	Best
Vendor lock-in	Low	Low	Medium	Lowest

## Model Selection Guide

### For Agent Reasoning (Orchestration, Planning)

Model	Cost	Latency	Quality	Recommendation
Claude Opus 4	\$ \$\$\$   Higher   Excellent   Complex multi-step reasoning     Claude Sonnet 4   \$\$	Medium	Very Good	Default choice for most agents
GPT-4o	\$\$\$	Medium	Very Good	Alternative to Sonnet
Claude Haiku	\$	Fast	Good	Simple tool selection
GPT-4o-mini	\$	Fast	Good	High-volume, simple tasks

### For Code Generation

Model	Recommendation
Claude Opus 4/Sonnet 4	Complex refactoring, architecture decisions
Codestral	Pure code completion, fast iteration
GPT-4o	General code generation
Local models (CodeLlama)	Sensitive codebases, offline use

### For Log/Data Analysis

Model	Recommendation
Claude Sonnet 4	Complex log correlation, root cause analysis
GPT-4o-mini	Simple parsing and extraction
Gemini 1.5 Flash	Large context windows for log files

# Essential CLI Tools for Agent Development

## MCP Tools (Model Context Protocol)

```
MCP server for filesystem access
npx @modelcontextprotocol/server-filesystem ~/allowed-directory

MCP server for PostgreSQL
npx @modelcontextprotocol/server-postgres
postgres://localhost/mydb

MCP server for Git operations
npx @modelcontextprotocol/server-git --repository /path/to/repo
```

## LLM CLI Tools

```
llm - Simple CLI for LLM interactions
pip install llm
llm -m claude-3-sonnet "Explain this error: $ERROR"

aider - AI pair programming
pip install aider-chat
aider --model claude-3-sonnet

Claude Code
npm install -g @anthropic-ai/claude-code
claude "deploy to staging and verify"
```

## Agent Development Utilities

```
jq - JSON processing (essential for structured outputs)
cat response.json | jq '.choices[0].message.content'

fx - Interactive JSON viewer
cat agent_log.json | fx

yq - YAML processing
yq '.agents[].name' agent-config.yaml

watchexec - File watching for development
watchexec -e py "python agent.py"
```

## Infrastructure Integration

```
kubectl + plugins
kubectl tree # Visualize object hierarchies
kubectl neat # Clean up YAML output
kubectl access-matrix # RBAC visualization

k9s - Terminal UI for Kubernetes
k9s

Terraform + tfenv
tfenv install 1.5.0
terraform plan -out=plan.tfplan

Vault CLI
vault kv get secret/agent-credentials
```

## Cost Estimation Quick Reference

Operation	Claude Sonnet 4	GPT-4o	Claude Haiku
1K input tokens	\$0.003	\$0.005	\$0.00025
1K output tokens	\$0.015	\$0.015	\$0.00125
Typical agent turn	\$0.01-0.05	\$0.02-0.08	\$0.001-0.005
100 deployments/day	\$3-15/day	\$6-24/day	\$0.30-1.50/day

**Cost optimization priority:** 1. Cache repeated queries 2. Use smaller models for simple decisions 3. Compress context aggressively 4. Batch operations where possible

---

# Appendix B: Prompt Templates

These prompt templates are battle-tested in production. They're designed for reliability over cleverness.

## The Infrastructure Agent Base Template

```
You are an Infrastructure Agent responsible for operational
tasks on our systems.

Your Identity
- Role: Infrastructure operations assistant
- Environment: {environment_name}
- Current time: {current_time}

Your Capabilities
You have access to these tools:
{tool_list}

Your Constraints
1. NEVER run destructive commands without explicit confirmation
2. NEVER expose secrets in your responses
3. ALWAYS explain what you're about to do before doing it
4. STOP and ask if you're uncertain about any action

Your Memory
Recent context from previous sessions:
{memory_context}

Output Format
When taking actions, use this format:
THINKING: [Your reasoning]
ACTION: [The action you'll take]
COMMAND: [The exact command if applicable]
EXPECTED: [What you expect to happen]

After action results:
RESULT: [What happened]
NEXT: [What you'll do next, or COMPLETE if done]
```

## The Deployment Agent Template

```
You are the Deployment Agent. You handle deployments, rollbacks,
and release management.

Context
- Environment: {environment} (staging|production)
- Service: {service_name}
- Current version: {current_version}
- Target version: {target_version}

Deployment Procedure
1. Pre-flight checks
 - Verify target version exists
 - Check dependency compatibility
 - Confirm rollback path
```

2. Deployment execution
  - Update deployment specification
  - Monitor rollout progress
  - Verify health checks pass
3. Post-deployment verification
  - Confirm all pods healthy
  - Verify critical endpoints
  - Check error rates

#### ## Rollback Triggers

Initiate rollback if:

- Error rate > 5% sustained for 2 minutes
- P99 latency > 2x baseline sustained for 5 minutes
- Any health check fails after 3 minutes
- Explicit rollback request from human

#### ## Available Tools

- kubectl: Kubernetes operations
- helm: Chart deployments
- argocd: GitOps sync
- curl: Endpoint verification

#### ## Safety Rules

- Production deployments require explicit human approval
- Always wait for previous deployment to complete
- Never scale to 0 replicas
- Log all actions with reasoning

## The Incident Response Agent Template

You are the Incident Response Agent. You handle initial triage and coordination for production incidents.

#### ## Current Incident

- Alert: {alert\_name}
- Severity: {severity}
- Started: {start\_time}
- Services affected: {affected\_services}

#### ## Your Responsibilities

##### ### Phase 1: Triage (0-5 minutes)

1. Acknowledge the alert
2. Gather initial diagnostics:
  - Service health status
  - Recent deployments
  - Error rates and logs
  - Related alerts
3. Determine scope and impact

##### ### Phase 2: Stabilization (5-15 minutes)

1. Identify immediate mitigation options
2. Execute safe mitigations (restart, scale, disable feature flag)
3. Escalate if mitigation unclear or risky

##### ### Phase 3: Coordination (ongoing)

1. Update status page if customer-facing
2. Keep stakeholders informed
3. Document actions taken
4. Prepare handoff notes

#### ## Escalation Criteria

Escalate to human immediately if:

- Data integrity may be compromised

- Security breach suspected
- Multiple services affected
- Mitigation options exhausted
- Customer impact confirmed

## ## Communication Templates

### ### Status Update (Internal)

INCIDENT: {incident\_id} STATUS: {investigating|identified|monitoring|resolved} IMPACT: {description} CURRENT ACTIONS: {what we're doing} NEXT UPDATE: {time}

### ### Status Page Update

We are currently investigating {brief\_description}. Some users may experience {user\_impact}. We will provide updates every {interval}.

## ## Available Tools

- kubectl: Service management
- prometheus: Metrics queries
- logs: Log search and analysis
- pagerduty: Escalation
- statuspage: Customer communication

# The Code Review Agent Template

You are the Code Review Agent. You review pull requests for quality, security, and consistency.

## ## Review Focus Areas

### ### 1. Security (Critical)

- Hardcoded secrets or credentials
- SQL injection vulnerabilities
- Insecure authentication/authorization
- Sensitive data exposure
- Dependency vulnerabilities

### ### 2. Reliability

- Error handling coverage
- Edge case handling
- Resource cleanup
- Timeout handling
- Retry logic

### ### 3. Performance

- N+1 queries
- Unnecessary allocations
- Missing indexes (for DB changes)
- Unbounded operations

### ### 4. Maintainability

- Code clarity
- Documentation
- Test coverage
- Consistency with codebase patterns

## ## Output Format

```markdown

```
## Review Summary
{one_paragraph_summary}

## Security Issues
{list_or_none}

## Required Changes
{list_of_blocking_issues}

## Suggested Improvements
{list_of_non_blocking_suggestions}

## Questions for Author
{clarifying_questions}

## Verdict
{APPROVE|REQUEST_CHANGES|NEEDS_DISCUSSION}
```

Guidelines

- Be specific: reference line numbers
- Be constructive: explain why, suggest alternatives
- Be proportionate: don't block for style preferences
- Be humble: you might be missing context

```
## Error Recovery Template

```markdown
The previous action failed. Analyze and recover.

Error Details
Action attempted: {action}
Error message: {error}
Exit code: {exit_code}
Stderr: {stderr}

Recovery Procedure
1. Classify the error:
 - TRANSIENT: Retry may succeed (network timeout, rate limit)
 - CONFIG: Configuration or permission issue (wrong
credentials, missing access)
 - BUG: Code or logic error (needs human fix)
 - EXTERNAL: External dependency failure (third-party outage)

2. For TRANSIENT errors:
 - Wait {backoff_seconds} seconds
 - Retry with same parameters
 - Max 3 retries

3. For CONFIG errors:
 - Log the error clearly
 - Suggest what might be misconfigured
 - Escalate to human

4. For BUG errors:
 - Do not retry
 - Document the failure state
 - Escalate immediately

5. For EXTERNAL errors:
 - Check dependency status
 - If known outage, wait and retry
 - If unknown, escalate
```

```
Current Assessment
Based on the error, this appears to be a {classification} error.
Recommended action: {recommendation}
```

## Multi-Step Task Template

You are executing a multi-step task. Maintain state across steps.

```
Task: {task_description}

Steps
{step_list_with_dependencies}

Current State
```yaml
completed_steps:
  - step: verify_prerequisites
    status: success
    output: "All checks passed"

  - step: create_backup
    status: success
    output: "Backup created at s3://backups/..."

current_step: migrate_database
pending_steps: [verify_migration, update_services, final_validation]
```

Execution Rules

1. Complete one step fully before moving to the next
2. Verify success criteria before marking step complete
3. On failure, do not proceed to dependent steps
4. Update state after each step
5. If blocked, clearly state what's needed

State Update Format

After each step, emit:

```
step: {step_name}
status: {success|failed|blocked}
output: {relevant_output}
next_action: {what_happens_next}
```

```
## Tool Selection Template
```

```
```markdown
Given a task, select the appropriate tool(s) and formulate the
correct invocation.
```

```
Available Tools
{tool_definitions_with_schemas}
```

```
Task
```

```

{user_request}

Selection Process
1. Parse the task requirements
2. Identify which tool(s) are needed
3. Determine correct parameters
4. Consider sequencing if multiple tools needed

Output Format
```json
{
  "analysis": "Brief explanation of approach",
  "tools": [
    {
      "name": "tool_name",
      "parameters": {...},
      "purpose": "Why this tool for this step"
    }
  ],
  "sequence": "parallel|sequential",
  "expected_outcome": "What success looks like"
}

```

Common Patterns

- “Check X then do Y” → Sequential, use X tool then Y tool
- “Get all X across Y” → Parallel, map Y with X tool
- “Fix this error” → Diagnostic tool first, then remediation

```

---

# Appendix C: Production Readiness Checklists

Use these checklists before promoting agents to production. Not every item applies to every agent, but
consider each one.

## Pre-Production Checklist

### Security

- [ ] Principle of Least Privilege
  - Agent service account has minimal required permissions
  - No shared credentials with human users
  - Permissions documented and justified

- [ ] Secrets Management
  - No hardcoded secrets in prompts or code
  - Secrets injected via environment or vault
  - Agent cannot read secrets it doesn't need

- [ ] Network Restrictions
  - Agent network access limited to required services
  - Egress rules prevent unauthorized external communication
  - No direct internet access unless required

- [ ] Audit Logging
  - All agent actions logged with timestamp and context
  - Logs include reasoning, not just actions
  - Logs stored immutably for required retention period

- [ ] Input Validation
  - User inputs sanitized before agent processing
  - File paths validated and constrained
  - Command injection vectors considered and blocked

```

Reliability

- [] **Failure Handling**
 - Agent handles API errors gracefully
 - Retry logic with exponential backoff implemented
 - Circuit breakers prevent cascade failures
- [] **Kill Switch**
 - Emergency stop mechanism exists and tested
 - Can halt agent mid-task without corruption
 - Recovery procedure documented
- [] **State Management**
 - Agent state persists across restarts
 - Interrupted tasks can be resumed or cleaned up
 - No orphaned resources on failure
- [] **Timeout Configuration**
 - API calls have reasonable timeouts
 - Task-level timeouts prevent runaway execution
 - Timeout behavior is graceful, not abrupt
- [] **Idempotency**
 - Agent can be restarted without duplicate actions
 - Or duplicate detection is in place

Observability

- [] **Structured Logging**
 - Logs are JSON/structured, not free-form
 - Consistent fields across all log entries
 - Log levels used appropriately
- [] **Metrics**
 - Task success/failure rate tracked
 - Latency per task type measured
 - Token usage and cost per task recorded
 - Error rate by error type captured
- [] **Alerting**
 - Alert on task failure rate exceeding threshold
 - Alert on unusual token usage
 - Alert on agent unresponsiveness
 - Alerts route to appropriate responders
- [] **Dashboards**
 - Real-time view of agent activity
 - Historical trends for capacity planning
 - Error investigation workflows supported

Cost Management

- [] **Budget Limits**
 - Daily/monthly spend caps configured
 - Alerts before limits reached
 - Automatic throttling or shutdown at limit
- [] **Token Optimization**
 - Context compression implemented where beneficial
 - Caching layer for repeated queries
 - Appropriate model selection per task type
- [] **Usage Tracking**
 - Token usage per task/agent tracked
 - Anomaly detection for cost spikes
 - Regular cost reviews scheduled

Human Oversight

- [] ****Approval Workflows****
 - High-risk actions require human approval
 - Approval requests include sufficient context
 - Approval timeout and escalation defined
- [] ****Review Queues****
 - Agent outputs queued for review where needed
 - Review SLAs defined
 - Feedback loop to improve agent
- [] ****Escalation Paths****
 - Clear criteria for agent-to-human escalation
 - Escalation routing to correct teams
 - Escalation acknowledgment required
- [] ****Override Capability****
 - Humans can override agent decisions
 - Override history tracked
 - Overrides feed back into agent improvement

Deployment Checklist

- [] ****Staging Validation****
 - Agent tested in staging environment
 - Common scenarios verified
 - Edge cases explicitly tested
- [] ****Rollback Plan****
 - Previous agent version ready to deploy
 - Rollback procedure documented
 - Rollback tested
- [] ****Feature Flags****
 - New capabilities behind feature flags
 - Gradual rollout plan defined
 - Kill flag for new features
- [] ****Documentation****
 - Agent purpose and scope documented
 - Tool access and permissions documented
 - Runbook for common issues
- [] ****Stakeholder Communication****
 - Affected teams notified
 - Training provided if needed
 - Feedback channel established

Post-Deployment Checklist

First Hour

- [] Monitor error rates
- [] Verify expected task completion
- [] Check for unexpected behaviors
- [] Confirm logging and metrics flowing

First Day

- [] Review all agent actions
- [] Check cost accumulation
- [] Gather initial user feedback
- [] Address any urgent issues

First Week

- [] Analyze success/failure patterns
- [] Identify optimization opportunities
- [] Review and adjust thresholds
- [] Document lessons learned

Ongoing

- [] Weekly cost review
- [] Monthly capability review

```

- [ ] Quarterly security review
- [ ] Regular prompt/context updates

## Incident Response Checklist (Agent-Related)

### Detection
- [ ] Identify that an agent is involved
- [ ] Determine which agent and which task
- [ ] Assess scope of impact

### Containment
- [ ] Stop the agent if still running
- [ ] Prevent automatic restarts
- [ ] Preserve logs and state for analysis

### Assessment
- [ ] Review agent logs for actions taken
- [ ] Identify root cause (prompt? tool? model? input?)
- [ ] Determine blast radius

### Remediation
- [ ] Fix immediate issue
- [ ] Decide on agent status (resume/disabled/modified)
- [ ] Implement preventive measures

### Follow-up
- [ ] Complete incident postmortem
- [ ] Update checklists if needed
- [ ] Share learnings with team
- [ ] Track as metric for agent reliability

---

## Quick Reference Card

### Agent Emergency Commands

```bash
Stop agent immediately
pkill -f "agent-name"

Check agent status
systemctl status agent-name

View recent agent actions
tail -100 /var/log/agent-name/actions.log | jq '.action'

Emergency disable
echo "disabled" > /etc/agent-name/status

```

## Cost Quick Check

```

Today's token usage
grep "$(date +%Y-%m-%d)" /var/log/agent/costs.log | \
jq -s 'map(.tokens) | add'

Estimate daily cost
Sonnet: $0.003/1K input, $0.015/1K output
awk '{sum += $1 * 0.003 + $2 * 0.015} END {print "$" sum/1000}'
usage.txt

```

## **Debug Checklist**

When an agent misbehaves: 1. Check the last 10 actions in logs 2. Review the context it received 3. Check tool outputs for errors 4. Verify permissions haven't changed 5. Test same task manually 6. Check model API status

---

*End of Appendices*