


# What a Chief AI Officer Actually Does

Turning fragmented AI activity into enterprise capability and measurable value

Ed Keisling | SVP, Chief AI Officer, Progress Software | 30 minutes



# About me...



Scan to connect

## About Progress

Progress Software products provide the context and control organizations need to reliably extract value from AI—**context** drawn from an organization's data, content, workflows and business rules, and **control** over the security, governance and cost of their AI initiatives.

## My Career

Pegasystems, Vecna Technologies, AthenaHealth, Progress Software


## My approach

Servant leadership, decentralized decision making, collective ownership of results.


# Investing and Innovating in AI



**Enable organizations to develop, deploy and manage responsible AI-powered applications and digital experiences**



**AI-power our products to better serve customers and drive revenue**



**Accelerate our business operations through AI**

# Five Focus Areas



## Strategy

Market narrative,  
product roadmap,  
executive priorities




## Enablement

Training, champions,  
communities, hands-  
on adoption



## Governance

AI Council, policy,  
risk, intake, approvals



## Platform

Reusable patterns,  
shared components,  
discovery



## Economics

Tokens, model  
routing, budgets,  
chargeback, ROI

# Hub and Spoke

## Mission

Establish Progress as an AI-native product and engineering organization, and a thought leader at the forefront of AI-enabled product delivery.

### Lean central hub — 7 people

- Product and AI product strategy
- AI strategy and advocacy (product + engineering)
- AI engineering, program and delivery
- AI intel and automation


### Distributed AI Champions network

- Drives adoption inside the BUs
- Localizes guidance and best practice
- Builds social proof and follow-through
- Vanguard runs the frontier experiments

---

# Strategy


Market narrative, product roadmap, executive priorities




# Market Dynamics have Shifted




The frontier labs are pushing toward autonomous engineering




An idea can become a working app in minutes



Software can now be built by describing the outcome



Teams can move from idea to product in days



*One signal I watch: new Stack Overflow questions per month*

# Move from AI programs to an AI-native operating model

## AGENT FLEETS

Move from pair programming toward governed autonomous workflows.

## AI AS A SERVICE

Codify shared components and patterns for product teams.

## VALUE TARGETS

Set baseline-to-target metrics for capabilities and engineering throughput.

## M&A ACCELERATOR

Use the playbook to assess and onboard acquisitions without scaling cost linearly.

**The next constraint is not access to better models. It is our ability to redesign work and institutionalize what works.**

# The metrics are maturing

Progress started with access and activity, then moved toward behavior, impact and economics

STARTED WITH

## Access & activity

- Licenses and requests
- Training participation
- Satisfaction surveys
- Initial utilization



THEN ADDED

## Adoption & maturity

- Active and repeat use
- Crawl / Walk / Run
- Team maturity progression
- Surface and workflow usage



MOVING TOWARD

## Impact & economics & ROI


- Quality and task success
- Cycle time and throughput
- Cost per output
- Revenue, retention, capacity

**Key shift: from “Is AI being used?” to “Is AI changing the economics or performance of the workflow?”**

---

# Enablement

Training, champions, communities, hands-on adoption



# Satisfaction Metrics



How would you feel if you could no longer use GitHub Copilot?


|                          |      |
|--------------------------|------|
| Not Dissapointed         | 7.5% |
| Somewhat Disappointed    | 49%  |
| Very Disappointed        | 41%  |
| Not Currently Using G... | 2.5% |

How would you feel if you could no longer use GitHub Copilot?

|                           |       |
|---------------------------|-------|
| Not Dissapointed          | 4.2%  |
| Somewhat Disappoin...     | 29.5% |
| Very Disappointed         | 65.3% |
| Not Currently Using GH... | 1.1%  |

"Using GitHub Copilot, I am able to complete my coding tasks more..."

|                          |       |
|--------------------------|-------|
| Strongly Disagree        | 1.1%  |
| Disagree                 | 1.1%  |
| Neither Agree nor Dis... | 11.6% |
| Agree                    | 56.8% |
| Strongly Agree           | 29.5% |




# AI Utilization within Engineering

- Measured by number of requests to the models


Premium Request Growth: Actual vs Projected

Bars = monthly requests • Pink line = daily run rate • \* = projected full month



Premium Requests — All

Each bar = one user, sorted high



# GitHub Copilot Maturity Phases

## What the Phases Mean & Typical Tooling

Use this as the interpretation key for maturity cohorts and phase-based PR/cycle-time comparisons.

### Phase 0 · Not Engaged

Not yet engaged with Copilot in a way that meets cohort activity thresholds.

**Typical tools:** Mostly manual IDE workflows; little or no Copilot Chat/CLI/agent usage.

### Phase 1 · Code First

Code-first usage centered on inline suggestions and basic IDE assistance.

**Typical tools:** Code completion, inline chat, and editor-integrated Copilot basics.

### Phase 2 · Agent First

Agent-first behavior with at least one advanced agent surface in regular use.

**Typical tools:** Copilot CLI, cloud coding agent, or Copilot Code Review as primary workflow helpers.


### Phase 3 · Multi-Agent

Multi-agent adoption across multiple advanced Copilot surfaces.

**Typical tools:** Combination of CLI + cloud agent + code review (and related agent-capable surfaces).

## Cohort Mix Over Time

Stacked bars = seat count per phase · Pink line = % in advanced cohorts (right axis)



# Engineering AI maturity model (Crawl / Walk / Run)



## Overall maturity



Crawl

Crawl

Walk

Crawl


## How we use it

- Checkpoints enablement impact over time
- Highlights where to invest vs. amplify
- Now calibrated with Copilot usage data

---

# Governance

AI Council, policy, risk, intake, approvals



# CAIO team with Other teams




**CAIO**

**Product AI Strategy and  
Enablement**



**IT AI Ops**

**Enterprise Tools  
& Automation**




**AI CoE - G&A**

**G&A AI Automation**

**AI Council**

**Oversight, Risk and Governance**

# The CAIO connects accountabilities that otherwise remain separate





# AGENT RISK CLASSIFICATION MODEL

A simple framework for classifying AI agents based on data, permissions, autonomy, impact, and authority.

|                                                                                                                                      | ✓ LOW                                                                     | ★ MEDIUM                                     | ⚠ HIGH                                                                                     | 🔥 CRITICAL                                                                                        |
|--------------------------------------------------------------------------------------------------------------------------------------|---------------------------------------------------------------------------|----------------------------------------------|--------------------------------------------------------------------------------------------|---------------------------------------------------------------------------------------------------|
| 1  What data does it use?                           | Public, user-provided, or non-sensitive internal data                     | Routine internal business data               | Confidential, customer, proprietary, or sensitive data                                     | Restricted, regulated, PII/PHI/PCI, or highly sensitive data                                      |
| 2  What permissions does it have?                   | Read-only within an individual's existing privileges                      | Limited write in standard work tools         | Broad read or write access to important business or operational systems                    | Privileged or admin access, or write access to financial, regulated, or security-critical systems |
| 3  How autonomous is it?                            | User initiates each action                                                | Bounded multi-step tasks with user oversight | Acts independently within defined workflows                                                | Operates continuously or makes high-impact decisions without prior approval                       |
| 4  How broad is its deployment or potential impact? | Single individual; impact limited to that user                            | Team or project                              | Multiple teams, business processes, or customers                                           | Enterprises-wide, external, regulatory, or material business impact                               |
| 5  What consequential authority does it have?      | No authority to approve, commit, delete, or trigger consequential actions | Can initiate or modify routine workflows     | Can trigger important workflows, publish externally, or create material operational impact | Can approve or execute financial, legal, regulated, security, deletion, or irreversible actions   |



## SIMPLE RULE



Overall risk equals the **highest** application rating across the five rows.



## ALWAYS REQUIRED

- ✓ Using approved AI tools
- ✓ In alignment with Progress policies
- ✓ Human accountability for outcomes



## RISK LEVEL GUIDE

- |            |                                                        |
|------------|--------------------------------------------------------|
| ● LOW      | No AI request needed with approved tooling             |
| ● MEDIUM   | AI request needed                                      |
| ● HIGH     | Security review request required before production use |
| ● CRITICAL | C-Staff approval needed                                |




The risk level inherits all previous risk tier actions

---

# Platform

Reusable patterns, shared components, discovery




# The Reuse Imperative

We analyzed 569 AI roadmap items across all BUs. Here are the results.

68% fall into shared AI focus areas:

- Building Agents: 83%
- RAG Embeddings & Pipelines: 25%
- Publishing APIs via MCP: 18%
- LLM & Agent Observability: 13%

Top 3 AI Focus Areas in our Backlog



# 4 measurement dimensions + MVS

| Business Impact & Value                                                                                                                           | User Experience & Engagement                                                                                                                                                               | Operational & Trust                                                                                                                                                                                               | Model Performance & Quality                                                                                                                                                                                               |
|---------------------------------------------------------------------------------------------------------------------------------------------------|--------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|-------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|---------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <ul style="list-style-type: none"><li>• How much value and monetization the AI feature creates</li><li>• Attach Rate or Conversion Lift</li></ul> | <ul style="list-style-type: none"><li>• Whether users discover, adopt, and get to first value</li><li>• AI Activation Rate</li><li>• Time to First Value or Task Completion Rate</li></ul> | <ul style="list-style-type: none"><li>• Whether the feature is scalable, cost-effective, and reliable</li><li>• <b>Cost per AI Task / Inference</b></li><li>• <b>End-to-End Latency or Availability</b></li></ul> | <ul style="list-style-type: none"><li>• Whether outputs are good enough, safe enough, and improving</li><li>• <b>Task Success Rate</b></li><li>• <b>Critical Error Rate, Grounded-ness, or Format Adherence</b></li></ul> |

Proof the metrics translate to value

52%

AI Summary (SecOps) reduction in event-prioritization time — 2.5 hrs saved per analyst per day, 160+ analyst-hours/month


50%

AI Threat Briefings feature adoption — 2 to 3x the typical 15-25% SaaS/enterprise adoption rate

---

# Economics

Tokens, model routing, budgets, chargeback, ROI



## Four lenses prevent one-dimensional decisions

### ADOPTION

Activation  
Repeat use  
Workflow penetration

### IMPACT

Cycle time  
Throughput / capacity  
Quality or risk

### ECONOMICS

Spend attribution  
Cost per output  
Incremental cost vs. value

### SCALE

Reuse  
Duplication avoided  
Time to subsequent value


## QUESTION 1 · ENTERPRISE ROI

ROI on a spectrum




- Marketing automations
- Support deflections
- Pull Requests automatically generated by AI
- Time saved on a task
- Meeting summaries
- Faster delivery
- More agents
- More PR's
- Production adoption/stickiness
- Expense Reduction
- Top line growth

# GitHub copilot Cost Visibility and Governance




Cumulative AIC — Month-over-Month Tracking



- Established policy and governance for Copilot budgets post pricing change
- Training and Enablement for ICs and Managers on how to optimize model usage
- Given ICs dashboard to get near-real time visibility into their usage with drilldowns & Managers to review their team's usage
- Worked with IT and BU Prod/Engg leaders to setup AI Credit limits and limit extension process

# PRs created/merged example



Web-UI — PRs per Sprint (detailed)

| Period         | Created | Merged | +Additions | -Deletions | Files | Avg Turn |
|----------------|---------|--------|------------|------------|-------|----------|
| Pre-PI Block 3 | 12      | 6      | +2,910     | -1,055     | 52    | 4.0h     |
| Pre-PI Block 4 | 27      | 23     | +31,230    | -1,382     | 210   | 4.4h     |
| Pre-PI Block 5 | 26      | 20     | +34,922    | -8,246     | 169   | 10.5h    |
| Pre-PI Block 6 | 15      | 12     | +11,871    | -14,232    | 124   | 14.0h    |
| Sprint 1       | 3       | 3      | +5,020     | -28        | 36    | 0.5h     |
| Sprint 2       | 14      | 9      | +1,350     | -449       | 53    | 4.5h     |
| Sprint 3       | 52      | 45     | +49,446    | -13,697    | 360   | 7.7h     |
| Sprint 4       | 74      | 56     | +33,911    | -16,336    | 424   | 3.4h     |
| Sprint 5       | 79      | 61     | +173,529   | -4,932,270 | 1,357 | 4.2h     |
| Sprint 6*      | 44      | 41     | +32,903    | -4,961     | 437   | ~3.5h    |

*Sprint 5 deletion count inflated by SDK widget migration (PR #265 removed ~4.9M lines of legacy package code)*

Per-Developer Correlation

Each developer's own monthly AI usage vs their own monthly PR outcomes (up to 6 monthly observations per developer, 2025-10 → 2026-03) — unlike the table above, this compares each developer only against themselves over time. Every cell shows both scores and signals. Select a developer to expand the full timeframe, monthly evidence, formulas, and factual interpretation.

| DEVELOPER    | MERGED PRS                   | CYCLE TIME                   | MERGE RATE                   | LINEs                        | FILES                        | COMMITs                      |
|--------------|------------------------------|------------------------------|------------------------------|------------------------------|------------------------------|------------------------------|
| Thierry Ciot | P 0.451<br>↑ moderate        | P 0.434<br>↑ moderate        | P 0.629<br>↑ strong          | P 0.184<br>↑ weak            | P 0.040<br>→ negligible      | P -0.241<br>↓ weak           |
|              | S 0.486<br>↑ moderate        | S 0.600<br>↑ strong          | S 0.700<br>↑ very strong     | S -0.700<br>↓ very strong    | S -0.500<br>↓ strong         | S -0.500<br>↓ strong         |
|              | n=6                          | n=5                          | n=5                          | n=5                          | n=5                          | n=5                          |
|              | P N/A<br>→ insufficient data | P N/A<br>→ insufficient data | P N/A<br>→ insufficient data | P N/A<br>→ insufficient data | P N/A<br>→ insufficient data | P N/A<br>→ insufficient data |
| Thierry Ciot | S N/A<br>→ insufficient data | S N/A<br>→ insufficient data | S N/A<br>→ insufficient data | S N/A<br>→ insufficient data | S N/A<br>→ insufficient data | S N/A<br>→ insufficient data |
|              | n=6                          | n=0                          | n=1                          | n=0                          | n=0                          | n=0                          |
|              | P 0.027<br>→ negligible      | P 0.828<br>↑ very strong     | P -0.857<br>↓ very strong    | P -0.650<br>↓ strong         | P -0.494<br>↓ moderate       | P 0.925<br>↑ very strong     |
|              | S 0.086<br>↑ negligible      | S 0.600<br>↑ strong          | S -0.941<br>↓ very strong    | S -0.257<br>↓ weak           | S -0.543<br>↓ strong         | S 0.812<br>↑ very strong     |
| Thierry Ciot | n=6                          | n=6                          | n=6                          | n=6                          | n=6                          | n=6                          |
|              | P 0.334<br>↑ moderate        | P -0.034<br>→ negligible     | P -0.344<br>↓ moderate       | P 0.586<br>↑ strong          | P 0.804<br>↑ very strong     | P 0.941<br>↑ very strong     |
|              | S 0.600<br>↑ strong          | S -0.087<br>↓ negligible     | S -0.464<br>↓ moderate       | S 0.943<br>↑ very strong     | S 0.943<br>↑ very strong     | S 0.899<br>↑ very strong     |
|              | n=6                          | n=6                          | n=6                          | n=6                          | n=6                          | n=6                          |

DEVELOPER STORY · 2025-10 → 2026-03

From 2025-10 to 2026-03, Thierry Ciot's AI usage and Merged PRs have a Pearson r of 0.334 (positive, moderate) and a Spearman rho of 0.600 (positive, strong). A positive relationship means higher AI usage aligned with higher merged prs; a negative relationship means it aligned with a lower value.

View raw PRs

Merged PRs

Evaluated 2025-10 → 2026-03 · 6 valid monthly observations

PEARSON R

0.334

↑ positive · moderate

SPEARMAN P

0.600

↑ positive · strong


Across 6 valid months from 2025-10 to 2026-03, Thierry Ciot's AI usage and Merged PRs have a Pearson r of 0.334 (positive, moderate) and a Spearman rho of 0.600 (positive, strong). A positive relationship means higher AI usage aligned with higher merged prs; a negative relationship means it aligned with a lower value.

| MONTH   | AI REQUESTS | METRIC VALUE (MERGED PRS) |
|---------|-------------|---------------------------|
| 2025-10 | 71          | 11                        |
| 2025-11 | 439         | 24                        |
| 2025-12 | 107         | 9                         |
| 2026-01 | 350         | 17                        |
| 2026-02 | 1,005       | 47                        |
| 2026-03 | 1,772       | 15                        |

► Formula, filters, and caveats

---

# What we've learned to date



# What we have learned

1

**Not everyone needs to be a citizen developer**

Building the wrong thing faster creates a drag on the good innovation that is occurring within the org. Identify the builders and have everyone else “consume”

2

**You need to recapture/reallocate savings from automations**

You need to be intentional about reallocating team members after workflows are accelerated, otherwise they will just pull in the next batch of work that may not be critical to the business.

3

**More AI usage is not better**

Zero or near zero adoption is a red flag, virtually everything else needs an understanding of the project, person and problems they need to solve. High usage can also be a sign of inefficient usage.

4

**We need to be intentional about connecting value streams together with AI**

Optimizing in silos creates bottlenecks in other departments. Focus instead on “How can we reduce the cycle time for customers to renew”

5


**Build Governance in**

To use AI well, safely, and securely is complex. Simplify this by building the guardrails in so people do the right thing and use the tools in a way that make sense.


**The enterprise question is whether value and learning are growing faster than AI cost, complexity and risk.**

# A short list of resources I have found useful


## READ



Practical ideas for building organizations that can adapt faster



A useful framework for explaining why the change matters



A simple reminder that people experience change differently

## LISTEN

### Artificial Intelligence Podcast

Long-form interviews with many of the people building the models

### Networking

Often the most practical source of what is actually happening

## TRACK

### MIT AI Risk Initiative

[airisk.mit.edu](https://airisk.mit.edu)

### Frontier lab threat and safety reports

Useful input when reviewing risk and policy

THE JOB IN ONE SENTENCE

# Make AI compound.

Across products.

Across teams.

Across acquisitions.

Across time.

Value and learning must grow faster than cost,  
complexity and risk.



