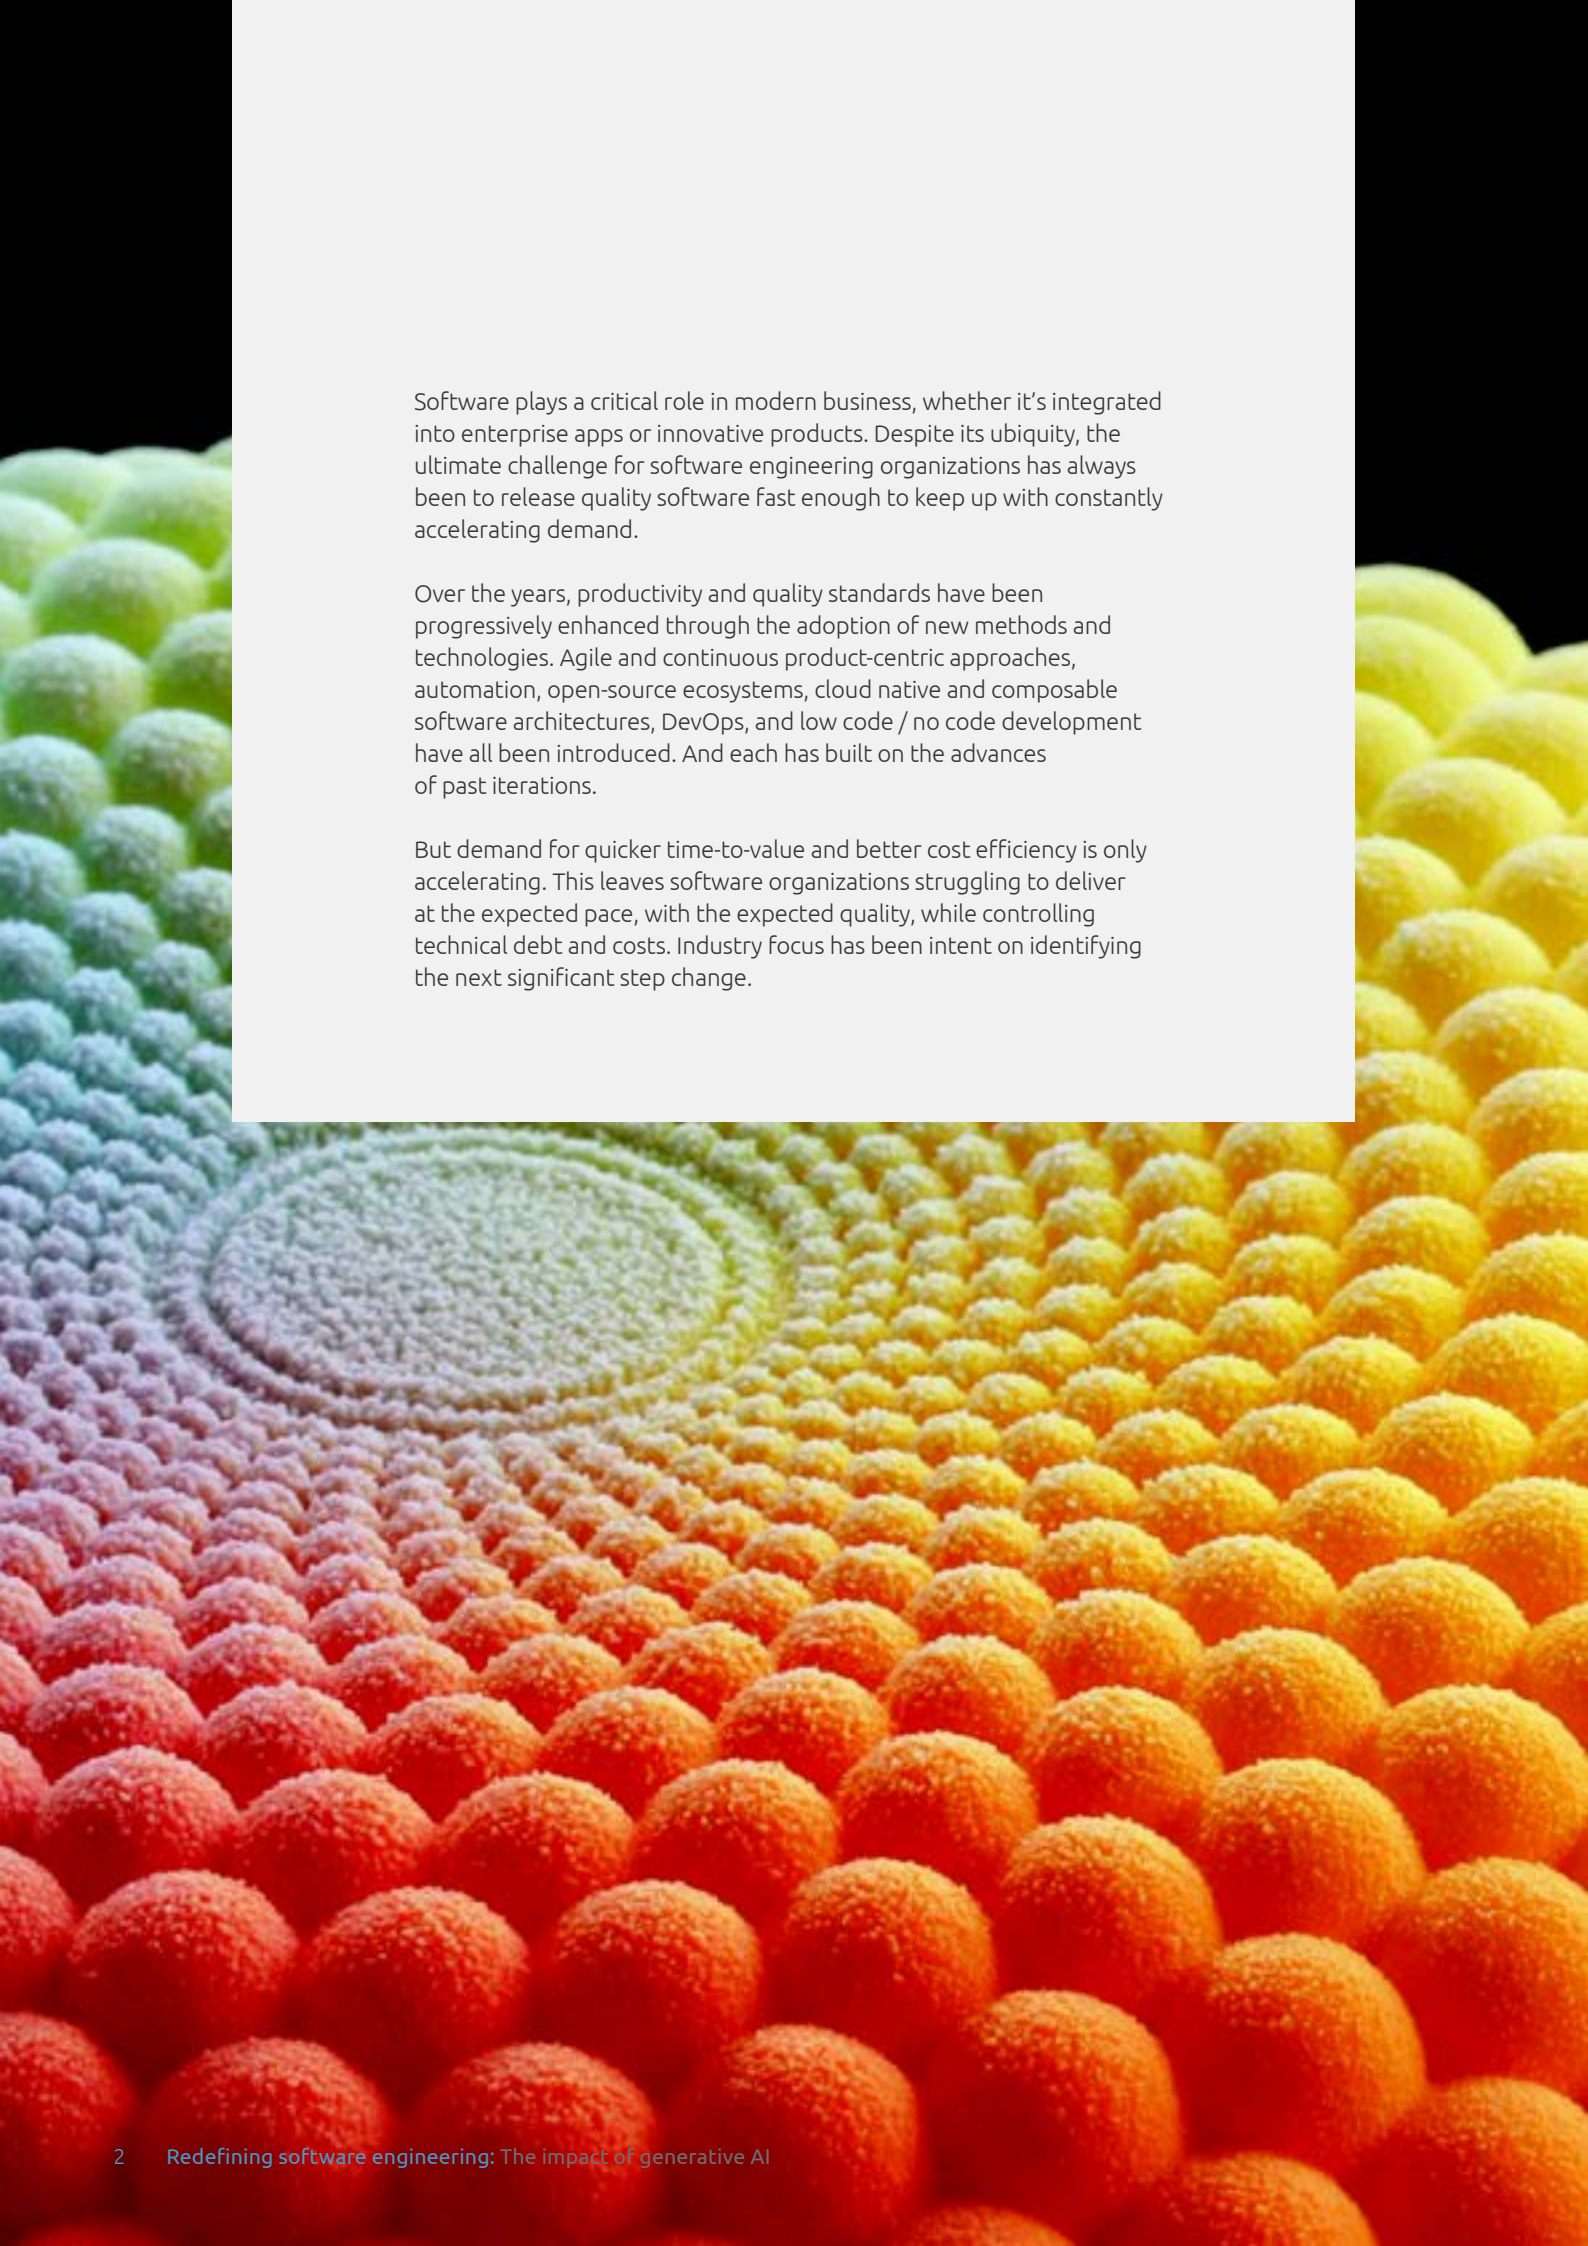


Redefining software engineering

The impact of generative AI



Software plays a critical role in modern business, whether it's integrated into enterprise apps or innovative products. Despite its ubiquity, the ultimate challenge for software engineering organizations has always been to release quality software fast enough to keep up with constantly accelerating demand.

Over the years, productivity and quality standards have been progressively enhanced through the adoption of new methods and technologies. Agile and continuous product-centric approaches, automation, open-source ecosystems, cloud native and composable software architectures, DevOps, and low code / no code development have all been introduced. And each has built on the advances of past iterations.

But demand for quicker time-to-value and better cost efficiency is only accelerating. This leaves software organizations struggling to deliver at the expected pace, with the expected quality, while controlling technical debt and costs. Industry focus has been intent on identifying the next significant step change.

Generative AI is software engineering's latest and greatest evolution

Generative AI now stirs wonder and excitement across many application fields, not least software engineering. Bill Gates is calling its latest advancement the most revolutionary technological achievement in over 40 years,¹ with the potential to dramatically improve the way organizations meet business and IT challenges.

The intersection of business and technology will be at the heart of generative AI's impact. Developments on one side will inevitably influence the other. For example, better software engineering can open doors to more automated routine tasks, such as data entry and document processing, giving leading organizations both a software and business-oriented edge.

Our latest Capgemini Research Institute report shows that 61% of organizations see enabling more innovative work, such as developing new software features and services, as the leading benefit generative AI. Close behind are improving software quality (49%) and increasing productivity (40%). Organizations are utilizing these productivity gains on innovative work such as developing new software features (50%) and upskilling (47%). Very few aim to reduce headcount (4%).²

According to a 2023 Capgemini Research Institute survey across 800 organizations, 67% of executives see the most potential for generative AI in the IT function to drive innovation and create value.² Moreover, according to Forrester research, "Off-the-shelf and custom AI software spend will double from \$33 billion in 2021 to \$64 billion in 2025 and will grow 50% faster than the overall software market, with an annual growth rate of 18%."³

Applying generative AI to software engineering significantly assists software teams in the multitude of tasks they perform across the traditional software development life cycle (SDLC). From business needs analysis and writing agile user stories to software design, coding, (retro) documentation, packaging, deployment, testing, and monitoring. In both legacy software modernization and new software development contexts, generative AI gives back time to software engineers. They are freed up to focus on business demand, software quality, security, and the advanced features required by new software.

Generative AI will supplement bandwidth - using the same capacity - for higher productivity and efficiency, along with faster time to market. But only if organizations and their employees commit to getting on board, keeping in mind that early adopters will take a leading position in the field.

In addition, Generative AI presents an opportunity, not only to transform, but also to standardize and enhance the delivery of digital projects. The benefits brought by using generative AI in software engineering can be reused to create more value, reduce IT costs, and minimize technical debt.

Generative AI also has a positive impact on software professionals' job satisfaction. 69% of senior software professionals and 55% of junior software professionals report high levels of satisfaction from using generative AI for software. 78% of software professionals are optimistic about generative AI's potential to enhance collaboration between business and technology teams.⁵



^[1] <https://www.gatesnotes.com/The-Age-of-AI-Has-Begun>

^[2] [Capgemini Research Institute "Turbocharging software", June 2024](#)

^[2] [Capgemini Research Institute "Generative AI in Organizations", July 2023](#)

^[3] [Global AI Software Forecast, Forrester Research, Inc. September 29th, 2022](#)

^[5] [Capgemini Research Institute "Turbocharging software", June 2024](#)

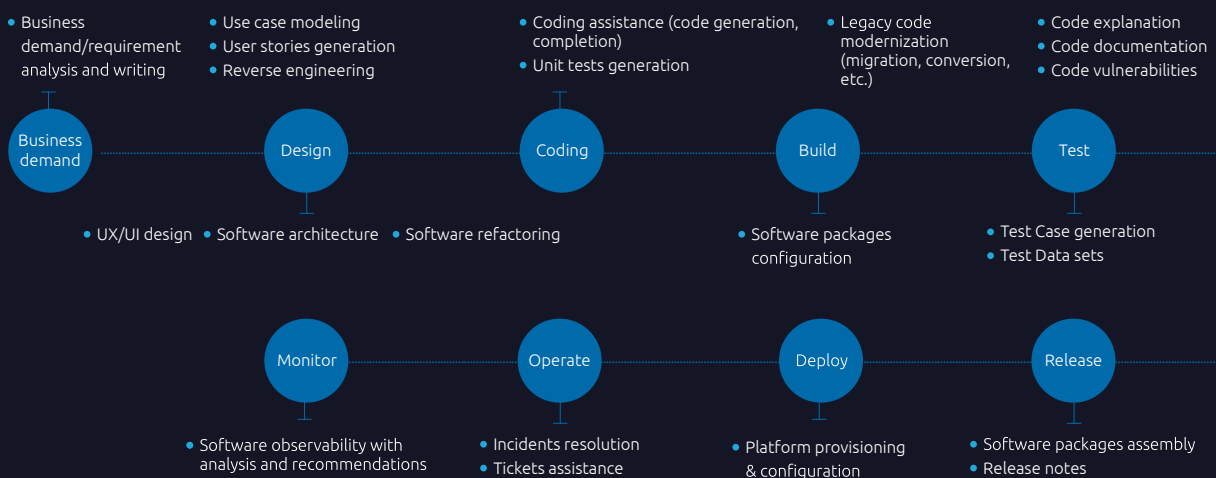
Exploring generative AI's potential across the software development life cycle

As cited in our research, Generative AI has the potential to yield benefits across the SDLC. It can be integrated at any stage – from business needs analysis and writing agile user stories to software design, coding, documentation, packaging, deployment, testing, and augmenting the work of software engineers and helping increase efficiency, improve quality, and enhance job satisfaction.⁴



Potential application areas of generative AI in the SDLC

(DevOps) Software Lifecycle



(Agile) Product Teams / (Waterfall) Development Teams

- Backlog and roadmap planning
- Effort estimations
- Team communication and collaboration
- Product value stream performance recommendations
- Team effectiveness analysis and improvement
- Process facilitation (plannings, retrospective, burndown, etc.)

Industrialized Software Engineering Platform

- Agile Process Management/ALM
- Developer workplace (IDE)
- DevOps automation
- Tests automation
- Generative AI Foundations

^[4] Capgemini Research Institute "Turbocharging software", June 2024



Capgemini is experimenting with multiple at-scale use cases, in real world contexts. The aim is to see where and how generative AI can augment the many tasks software teams typically perform. Here is a look at some important areas for generative AI integration:

Assisting business demand analysis: from epics to features to user stories

Business analysts add value by interacting with business stakeholders, such as product owners and business users. Generative AI and large language models (LLMs) are instrumental in helping them both understand business needs and turn them into actionable user stories. For example, by identifying and resolving inconsistencies or ambiguities in business needs. By prioritizing those needs and creating a roadmap for their implementation. By effectively communicating them to stakeholders, while spotting patterns and trends in their analysis. And by supporting the generation of natural language descriptions of epics, features, and user stories.

Accelerating software design and coding

Software engineers use generative AI when transforming backlog stories into software, through design and coding activities. Generative AI can be harnessed to create design outputs such as UI mock-ups, entity models, and microservice APIs. This can lead to a significant productivity improvement without compromising quality, as design outputs are always reviewed, updated, and validated by software engineers or technology leads.

Although adoption of generative AI for software engineering is still in its early stages with 9 in 10 organizations yet to scale, 27% of organizations are running generative AI pilots, and 11% have started leveraging generative AI in their software function. Generative AI is expected to play a key role in augmenting software workforce with better experience, tools and platforms, and governance, assisting in more than 25% of software design, development, and testing work by 2026.⁵

Generative AI is also poised to redefine conventional programming practices by shifting the focus from coding to prompt engineering and code proofreading. Andrej Karpathy, an OpenAI computer scientist, recently said: “the hottest new programming language is English.”⁶ As an example, using plain language, software engineers can describe the intended functionality of a software feature, then review, update, and validate the generated output. There are many other examples, such as auto-completion of code, generating code for unit testing, (retro) documentation, and code migration from one language to another.

^[5] [Capgemini Research Institute “Turbocharging software”, June 2024](#)

^[6] <https://twitter.com/karpathy/status/1617979122625712128?lang=en-GB>

Developers value generative AI because it already supports them during coding. It can either suggest clean code directly or evaluate existing code to improve software quality if it finds issues.

Making the build easier

Generative AI can assist software engineers in developing packages more easily, automating release notes, and building infrastructure as code.

Improving testing

Simple errors are not the only cause of software deficiencies. Often, subpar software quality can be traced back to the testing phase, where the test cases and/or related test data sets fail to include the full spectrum of possible user inputs and scenarios.

Generative AI can assist developers in writing more complete test cases, in which user stories serve as prompt engineering context for maximum relevance. Generative AI not only helps search, manage, and maintain test data, while detecting and solving security vulnerabilities. It can also generate a massive amount of synthetic information that closely resembles real world data to ensure tests cover a broader range of use cases. According to Gartner by 2025, 20% of all test data will be synthetically generated.⁷

However, the adoption of most use cases is at a nascent stage, and few organizations are implementing or piloting them. Coding assistance, the highest adopted use case, is at 39%, while test case generation stands at 26%. Globally less than one in four organizations are focusing on any use case on average.

Lessons learned

Our current experimentations cover multiple technology combinations across two dimensions. First, vendor packaged solutions that include generative AI extensions to developer tools and environments. And second, foundational LLMs with advanced techniques for contextualized prompt engineering.

In the case of vendor packaged solutions, we establish expertise in complementing LLMs with an organization's specific knowledge. We use techniques such as retrieval augmented generation or fine tuning, which considers existing epics and user stories, as well as industry / business domain documentation, public and/or internal design patterns, coding guidelines, and templates. To tap the full potential of an LLM, developers must acquire the skill and practice of good prompt engineering, which requires proper training, playbooks, and prompt libraries.



^[7] [Gartner Inc: "Gartner Identifies Three Technology Trends Gaining Traction in Banking and Investment Services in 2022", May 24, 2022.](#)

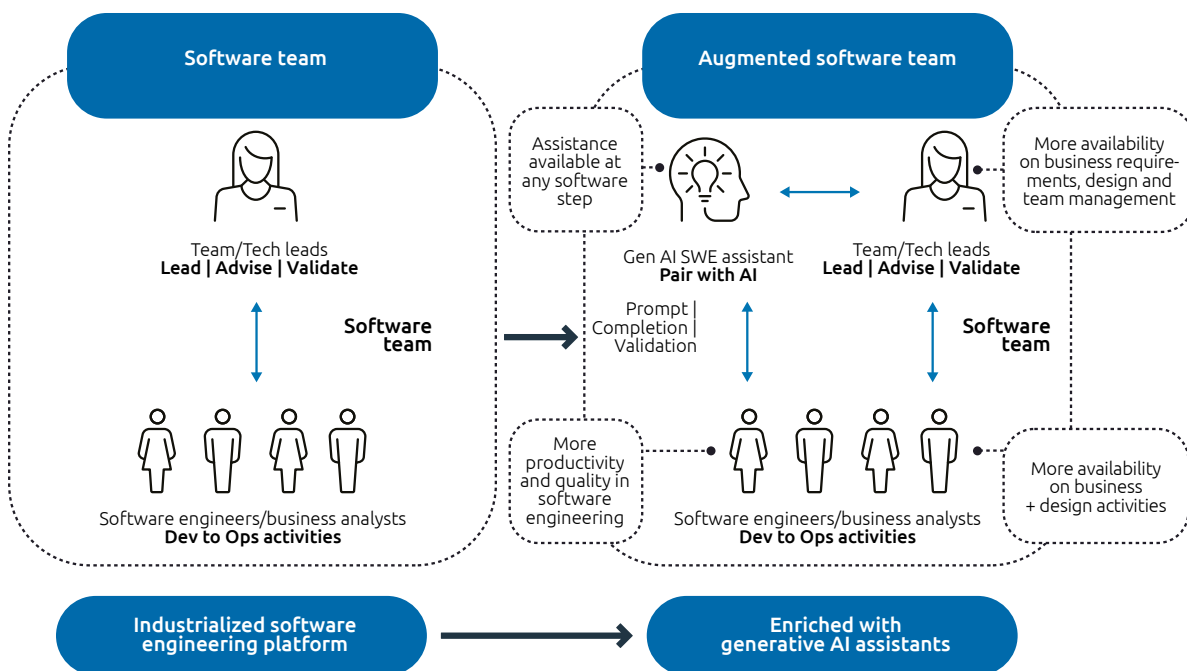
Bringing tomorrow's software engineering paradigm into view: augmented software teams

We believe generative AI will gradually transform the way software is developed. The advent of LLMs has introduced a compelling rationale for a paradigm shift to more AI-assisted (augmented) software teams.

The integration of generative AI into software engineering offers the promise of significantly elevated productivity and enhanced quality. The foundational principles and methods of Agile and DevOps are retained, including collaboration, adaptability, time to value, product centricity, continuous feedback loops.



How generative AI is impacting software teams



By working in concert with AI assistants powered by LLMs, augmented software teams can automate mundane tasks, expedite debugging, and promote more nuanced, data driven decision making. This optimizes the software life cycle and helps to achieve more milestones along the way.

We call this approach conversational software engineering. This is where software teams “talk” to the AI assistant,

asking it to generate snippets of code, troubleshoot issues, or even help design the architecture of software. The idea is to make the software engineering workflow more dynamic and interactive, through continuous conversations with the AI assistant in the development environment. It’s a way to streamline tasks and solve problems more efficiently, enabling software teams to focus on more complex and creative aspects that boost overall productivity.

The modern evolution of pair programming with generative AI

Augmented software teams can effectively benefit from and optimize usage of generative AI through a revised organization and new ways of working, guided by the following principles:

- **Augmented pair programming:** The basic unit of an augmented software team is one or multiple pairs of software engineers working intermittently with each other or individually with the AI assistant.
- **AI-human collaboration:** Team members work individually with the AI assistant to automate repetitive tasks, understand and solve problems, or brainstorm ideas, leveraging an LLM's speed and knowledge, while ensuring human creativity.
- **Human-human collaboration:** After conversations with an AI assistant, humans review the generated output and enhance its quality.
- **Senior coordination:** A senior lead oversees the augmented team, coordinating efforts and resolving conflicts, while ensuring control and validation to promote a smooth workflow and healthy team dynamic.

Of course, organizations need to prepare for generative AI use by delivering technology prerequisites, building a repository of platforms and tools for a seamless and augmented software engineering experience.

According to Capgemini Research Institute, the biggest gap in essential prerequisites is usually access to platforms and tools, including IDEs, automation and testing tools, and collaboration tools. Only 27% of organizations claim to have above-average availability of these. Of testing domain professionals, 24% say they have access to these tools, compared to 19% of project and program management professionals, further highlighting gaps within the software engineering function.⁸



^[8] [Capgemini Research Institute "Turbocharging software", June 2024](#)

Mismanaged AI integration: risks and consequences

Organizations that neglect to harness the power of generative AI to transform their software engineering paradigm will relegate their business to the rear ranks in terms of market standards. But there is undoubtedly a right and wrong approach to implement the technology. And the wrong approach can be just as debilitating as a defiant anti-AI stance. Both will deliver a devastating blow to an organization's chances of basic survival in today's dynamic software landscape.

Here are some scenarios that might unfold if organizations choose a strategy lacking the necessary precautions:

- 1. Falling behind market standards.** Organizations will see a decline in productivity, quality, and business agility compared to the market. They will struggle to keep pace with competitors who have successfully adopted a plan that combines human intelligence with technological innovation.
- 2. Shadow IT threat.** Developers, and even business users, might resort to using unvetted generative AI tools, leading to a shadow IT phenomenon. This poses a substantial hazard as it could compromise intellectual property, resulting in the inadvertent leakage of confidential data. In turn, this harms the trust customers have in the organization.
- 3. Uncontrolled costs.** Ineffective management of generative AI use could lead to spiraling costs that can negate the automation and efficiency gains.
- 4. Risk management challenges.** Organizations may find themselves ill-equipped to manage the multifaceted risks associated with generative AI, including legal, intellectual property, cybersecurity, and ethical concerns. This could expose the business to litigation and reputational damage.
- 5. Employee concerns.** An absence of official communication regarding the organization's stance on generative AI and the future roles of its software engineers also brings risk. This could lead to employees voicing their legitimate fears and concerns across the organization. Or simply leaving the company for a competitor with a more transparent generative AI view.
- 6. Talent attraction.** During recruitment interviews, the best software engineers ask whether they will benefit from generative AI tools. A company with no plans for the technology may miss out on highly skilled candidates.
- 7. Loss of sustainability objectives.** Since the rising carbon footprint of AI systems is a real concern, neglecting sustainability considerations could result in long term environmental and societal repercussions.

As an example, let's consider the scenario of employees using unvetted generative AI tools. According to Capgemini Research Institute, only 27% and 32% of organizations have the platforms and tools, and talent prerequisites in place, respectively, to implement generative AI for software engineering. Over 60% lack governance and upskilling programs for generative AI for software. As a result, of those software professionals who use generative AI, 63% use unauthorized tools and solutions. Nearly a third of the workforce is self-training on generative AI for software as only 40% of the organizations are upskilling employees on these skills. This is an extremely high risk. Using unauthorized tools without proper governance and oversight exposes organizations to functional, security and legal risks like hallucinated code, code leakage and IP issues.⁹

^[9] [Capgemini Research Institute "Turbocharging software", June 2024](#)





A structured strategy for a software engineering transformation

To start this journey with generative AI, organizations will reap huge benefits from partnering with an experienced and trusted advisor. In addition to having considerable AI expertise, this partner should be ready to oversee pilots and form collaborative project teams that eventually flourish into joint software houses.

A variety of factors, such as company size, staffing pyramid, and industry-specific goals, mean that each organizational unit progresses at a different pace and with a different model. Therefore, each will have distinct needs and challenges to consider before generative AI can become a part of daily operations. Organizations should reflect on the following:



1. Assess the organization's maturity and develop a roadmap with clear objectives

Within a company, there is usually not only one SDLC, but many. Typically, every program, project, and product has its own SDLC.

A meticulous assessment and deep understanding of the SDLC of an IT domain will reveal the current maturity of its software engineering processes. And whether they align with industry best practices. The evaluation will also identify areas for improvement, particularly where bottlenecks or inefficiencies occur.

Based on this assessment, objectives can be defined by selecting the most promising transformation enablers, as well as identifying their associated risks and challenges. Some enablers will be powered by generative AI while, for others, it will be more a matter of applying software engineering best practices. It's important to define key metrics to measure the outcomes of the future transformation and involve stakeholders early to foster a collaborative environment.

Finally, a comprehensive value, accessibility, and risk analysis will serve as a basis for establishing a roadmap for change. All options will be weighted and prioritized to make informed decisions about resource allocation as different domains, teams, and roles are defined for each option.

2. Run pilots/projects and measure the impact of generative AI

With the priorities set, it's time to select the best generative AI tools to achieve them. And put up guardrails to manage any legal and cybersecurity risks, while controlling costs.

Measurement and continuous improvement will be pivotal. As generative AI is inserted into more SDLC processes, organizations need to gauge its impact. This means measuring critical aspects such as productivity enhancements, software quality, time-to-market, and developer experience. A feedback loop should be put in place so the pace and scope of the deployment can be adapted to account for inefficiencies. Or respond to evolving needs and specific objectives.

But to be truly competitive, organizations need more than internal measurement. At Capgemini, we have developed an industrialized value measurement protocol that evaluates the objective impact of generative AI across an organization's many SDLCs. It is used to measure and compare an organization's metrics against our benchmark, which factors in all our related internal and external generative AI projects. This gives organizations a clear view of how they stack up against their peers.

To put this into perspective, 75% of organizations with more than \$20 billion of annual revenue have adopted (piloted or scaled) generative AI, compared to 23% of their smaller counterparts, with an annual revenue between \$1 - 5 billion.¹⁰

^[10] [Capgemini Research Institute "Turbocharging software", June 2024](#)

Measuring and benchmarking generative AI's impact at scale

There are two questions we hear most often in the software industry. First, how will generative AI-based assistance affect software productivity and quality? And, second, how will software teams' organization and ways of working change because of it?

At the beginning of 2023, Capgemini started a large-scale global program to experiment with use cases and measure generative AI's impact through pilots, both internally and jointly with clients. We deployed a reliable and consistent measurement protocol and now we are progressively consolidating our pilot results into a repository for internal and external usage / benchmarking.

Our unique and proven protocol measures generative AI's real-world impact for various use cases and technologies across the SDLC. The aim is to uncover how productivity, quality, security, time-to-market, and developer experience have changed because of generative AI.



A typical pilot includes one to multiple teams, for a minimum duration of six sprints, but preferably nine, with a significant and representative backlog of user stories and complexity mix to work with. We can implement different team configurations:

- Existing teams transition to using generative AI.
- New teams start without generative AI then later integrate it into their workflows.
- Shadow teams work alongside existing teams, using generative AI to perform the same tasks.

All setups can include different variations in terms of structure and seniority mix. And in all cases, a solid baselining is a fundamental prerequisite. This enables robust and representative comparisons, reports, and insights across different metrics:

1. **Velocity:** aggregated (team), per story complexity, per developer expertise
2. **Quality:** quality indexes, code duplicates and smells, code efficiency, unit test coverage
3. **Security:** vulnerabilities and potential risks
4. **Time-to-market:** from a change in the software code to its release in production
5. **Developer experience:** assessed through comprehensive surveys

Our pilots are delivering quantified results for coding, design, and business analysis use cases. We have seen significant productivity improvements, with better-than-expected quality and security, for both junior and senior developers tackling complex and very complex user stories.

We are learning, documenting, and packaging best practices for generative AI adoption through various technologies and architecture solutions, prompt engineering handbooks, and prompt libraries. We are also discovering proven and applicable augmented software team patterns for a conversational software engineering model that will make organizations perform more efficiently.

First results from internal and client pilots are progressively confirming a range of 10% - 30% workload gains across the whole life cycle for custom software development, without degrading software quality neither security, with potential peaks of 20% to 40% specifically for coding phases.

3. Deploy generative AI at scale

After experimenting with generative AI through real world pilots and projects, many organizations will want to broaden their applications and possibly involve hundreds or even thousands of developers.

For such large deployments, careful consideration must be given to the organizational and HR implications. Since various roles, including business analysts, software architects and engineers, and software testers may need to change, it's important to introduce a generative AI upskilling program. This will help shape the new software

engineering pyramid and lay the foundation for a strategic plan that gradually integrates generative AI into the different stages of the software life cycle. A further step may be to offer coaching and assistance to steer employees as they prepare to work alongside a host of new generative AI tools.

Finally, a dedicated team, or value realization office, should be set up to define business cases, measure progress, and ensure the outcomes meet expectations. This is essential, as the pace and scope of each deployment will invariably need adjustment.

The dawning of a new era

As a global leader in software engineering, Capgemini has a proven history of advising clients on how to elevate what is possible. Always with human creativity accenting machine efficiency. We've orchestrated many step changes in this industry, and none has stirred as much excitement and innovative promise as generative AI. It's a giant leap forward that has fundamentally changed the way software is created.

Embracing generative AI is not merely a choice, but a strategic imperative. It is predicted that 70% of professional developers will be using the technology by 2027.¹¹ And so the faster organizations get this transition gets underway, the further ahead of the competition they will be. Ultimately, the question is whether they will be an early adopter and transform at scale, leading the way. Or a late bloomer, playing catch-up.

Although the path to full adoption of generative AI may seem daunting, it's essential that organizations harness its unprecedented transformative power now.

Are you ready to reshape your software engineering?

^[11] <https://www.gartner.com/en/articles/set-up-now-for-ai-to-augment-software-development#:~:text=By%202027%2C%2070%25%20of%20professional,from%20less%20than%2010%25%20today.>



About Capgemini

Capgemini is a global business and technology transformation partner, helping organizations to accelerate their dual transition to a digital and sustainable world, while creating tangible impact for enterprises and society. It is a responsible and diverse group of 340,000 team members in more than 50 countries. With its strong over 55-year heritage, Capgemini is trusted by its clients to unlock the value of technology to address the entire breadth of their business needs. It delivers end-to-end services and solutions leveraging strengths from strategy and design to engineering, all fueled by its market leading capabilities in AI, cloud and data, combined with its deep industry expertise and partner ecosystem. The Group reported 2023 global revenues of €22.5 billion.

www.capgemini.com

For more information on Capgemini's Generative AI for Software Engineering offer, visit:

www.capgemini.com/solutions/generative-ai-for-software-engineering/

Authors:

Pierre-Yves Glever

Global Head of C&CA, Capgemini

Arnaud Balssa

Global Head of Business Technology, Capgemini Invent

Stephane Girard

CTO, Global C&CA, Capgemini

Thilo Hermann

Deputy CTO, Global C&CA, Capgemini

Serge Baccou

Offer Leader – Cloud & Generative AI, Capgemini Invent

Lorna Neville

Marketing Director, Global C&CA, Capgemini



Get the future you want