




Graphene v3: Release Highlights

New Version Release Announcement

Singapore & Ireland | 13 May 2025





Interested in learning more about Graphene v3?

For Existing Clients & Partners: Please contact your Account Manager, Customer Success Manager, or Partnership Manager.

For Prospective Clients & Partners: Please reach out to us at www.peak3.com or hello@peak3.com.






Table of Contents

1. Graphene v3 Highlights

2. **Key Product Line Updates**

- a) P&C Commercial Lines
- b) Employee Benefits (Health)

3. New Integrated AI Features

4. New Open X+ Platform & Open API v3

5. **Key Technology Updates**

- a) Self-Sufficiency (DMS)
- b) Reliability Enhancements
- c) Tech Stack & Performance Upgrades
- d) FinOps & Efficiency
- e) DORA Compliance

6. SI Delivery Enablement

7. Disclaimer & Safe Harbor Statement



Click hyperlinks to jump to the respective sections



Highlights

Graphene v3 Highlights



Highlights of Graphene v3: Major new functional, artificial intelligence and tech upgrades

[Click for more details](#)

Graphene v3 is the culmination of nearly one hundred bi-weekly releases and over 1.5 million hours and US\$75 million of dedicated R&D since the initial release of Graphene v2 in 2021. Graphene v3 exemplifies Peak3's vision of a microservices insurance core platform designed to holistically support digital-first and AI-first business models across property & casualty (P&C), life, and health insurance sectors—to optimize existing operations and capture new growth opportunities.

Graphene v3 now brings its transformative innovation capabilities to commercial and group insurance, while providing the foundations for AI-first insurance business models. Key highlights of Graphene v3 include:

- **Commercial Lines Expansion:** Extends Graphene's robust capabilities from retail P&C lines into commercial lines, facilitating comprehensive modernization of P&C systems for insurers.
- **Group Health Enhancement:** Advances beyond individual accident and health insurance into group health and employee benefits, supporting insurers' efforts in managing comprehensive health portfolios.
- **AI-Ready Capabilities:** Offers built-in, pre-integrated AI functionalities, enabling insurers to seamlessly train and deploy AI models directly within Graphene, embed intelligent chatbot solutions, manage intelligent document processing, detect and manage fraud, waste and abuse (FWA), and comprehensively orchestrate AI agents.
- **New Open API Platform:** Provides v3 of Graphene's Open API Platform with tailored APIs by product category, using domain specific language and facilitating rapid integration with full adherence to OpenAPI Specification (OAS) 3.0.1 standards.
- **Self-sufficiency Improvements:** Introduces new Data Management Service to easily export/import lossless configurations (e.g., products, formulas, workflows) across environments, instances and tenants accelerating innovation and empowering users.
- **Architecture and Tech Enhancements:** Improves resilience and scalability with a shift to serverless architecture, blue-green deployment and cross-cloud disaster recovery, further strengthening Graphene's industry-leading availability (99.998% in 2024 and 99.999% in Q1 2025, covering planned and unplanned downtime), boosts performance by 50%+ thanks to tech stack upgrades and DevPerfOps, and reduces cloud consumption by up to 30% through holistic FinOps initiatives.



Evolution of Graphene to v3: Realizing our vision of a modular, connected & intelligent core platform

Graphene v1

2018

Designed as **digital-core for innovative digital and embedded insurance** propositions. Focus on interoperability and scalability of greenfield ventures.

Graphene v2

2021

Enhancements as **omni-channel core for retail product lines (life, health and P&C)**. Available as middle office or end-to-end core for multi-country deployments.

Graphene v3

2025

Expansion to **P&C commercial and group health** to enable end-to-end core modernization. **Intelligent, AI-ready core** with own and third-party AI applications.

The journey from v2 to v3

2,000,000,000+ policies issued on Graphene's scalable infrastructure

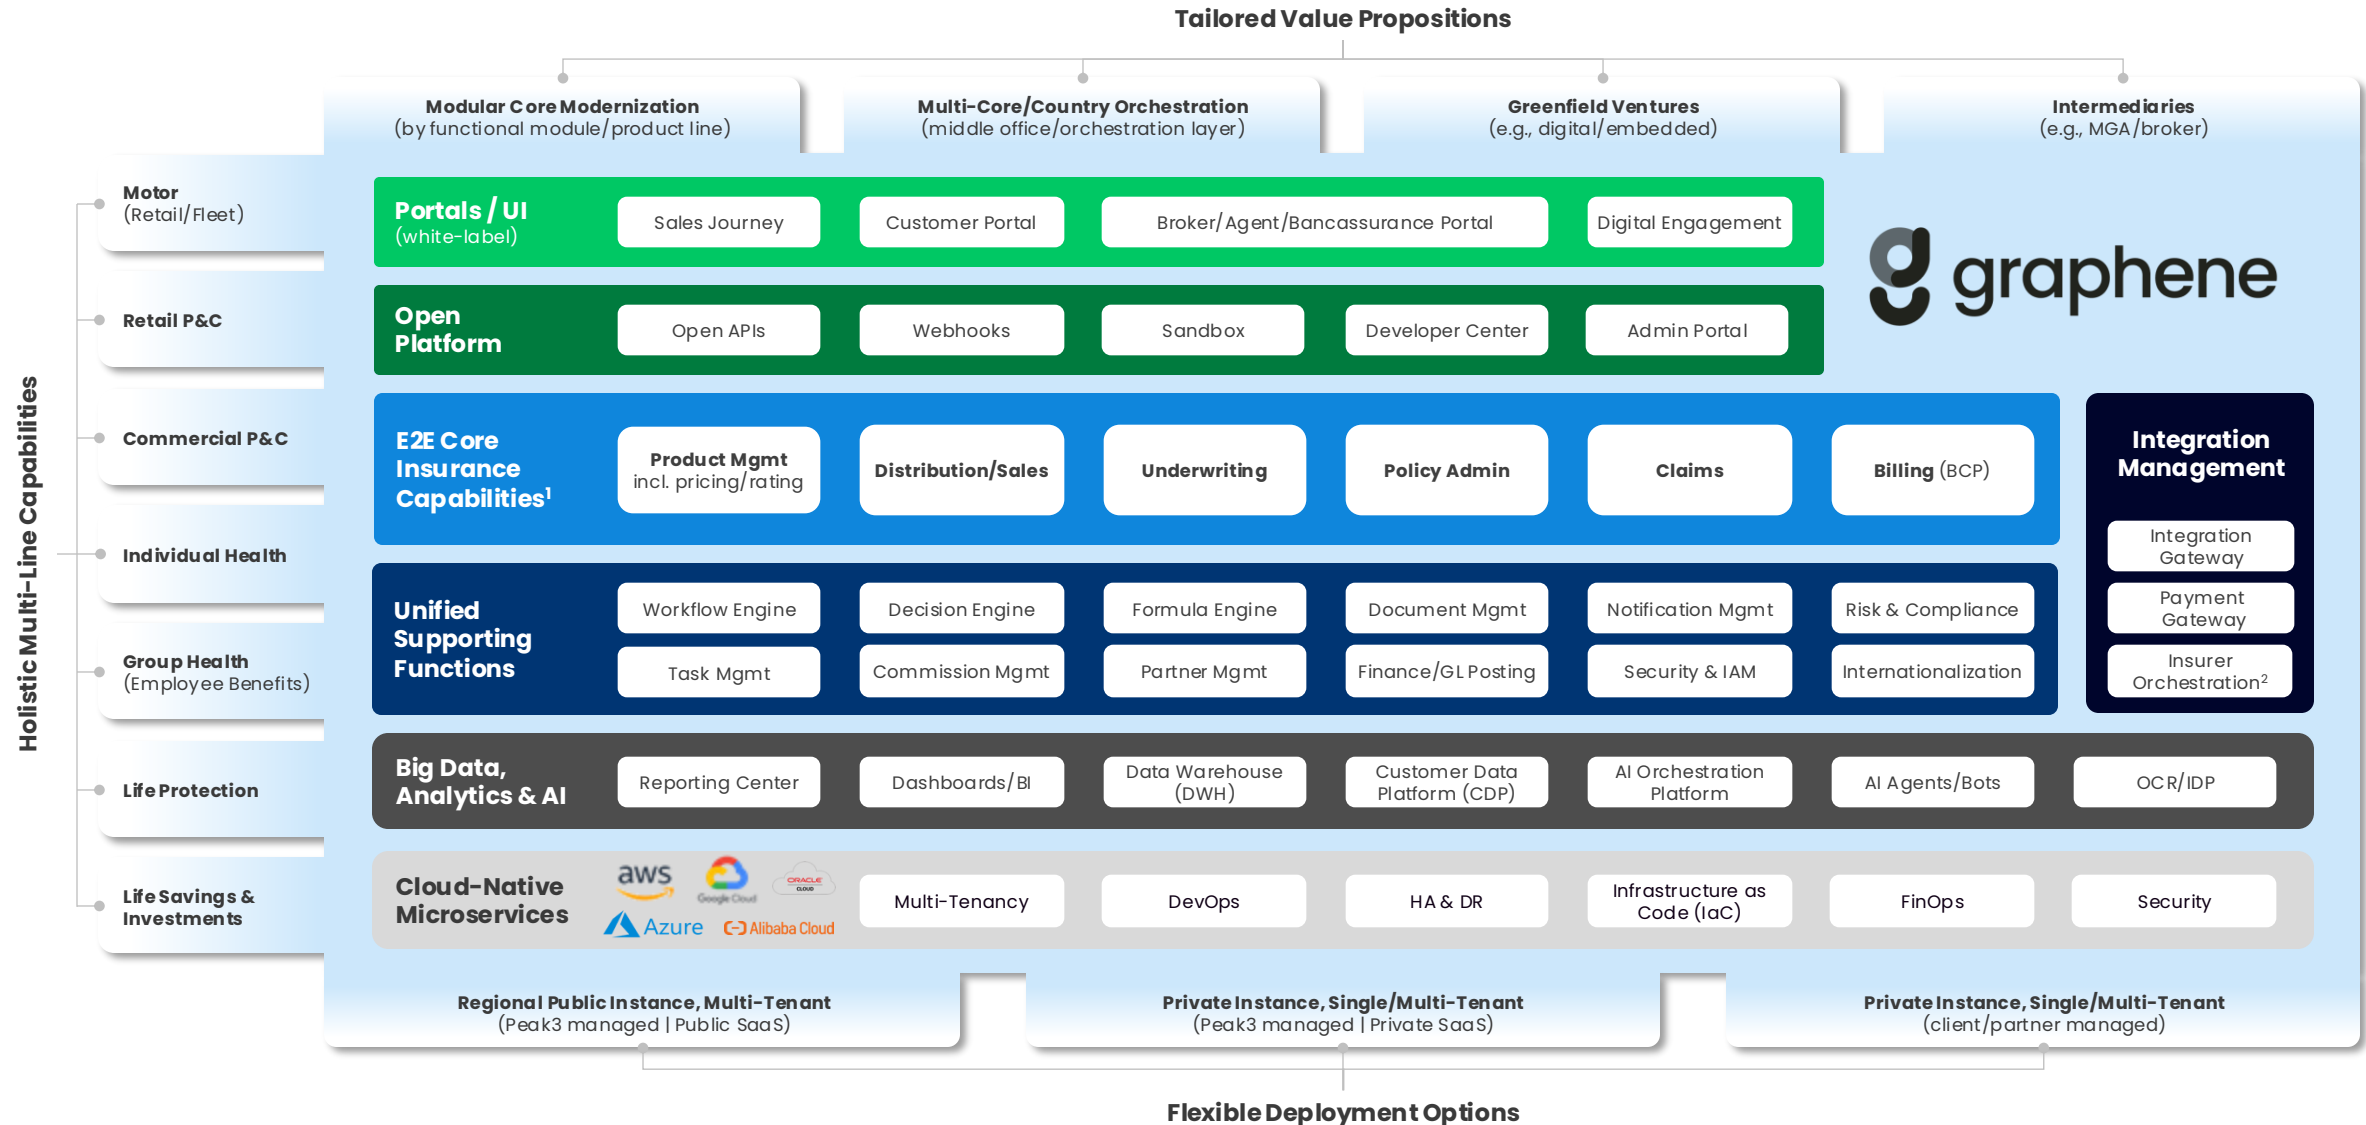
1,500,000+ hours of R&D to ship new features and tech upgrades every two weeks

96 minor bi-weekly version releases with regular client upgrades

1 single global code base for life, health and P&C



Graphene v3: Modular, connected & intelligent core platform




Notes: ¹ Reinsurance module currently available via 3rd party solution integration; ² Insurer Orchestration used when integrating with other core systems



Key Product Line Updates

P&C Commercial Lines

Graphene v3 holistically expands product coverage from retail lines to commercial lines




Graphene v3 now natively supports P&C commercial lines

Expanding Graphene's capabilities for digital-first and innovative retail lines to commercial lines




Modular Product Design

- Multi-layer product structures (technical & marketing)
- Modular insured object management



Flexible Distribution

- Standalone sales across distribution channels
- Bundling with consolidated invoicing and communications



Flexible Policy Structures

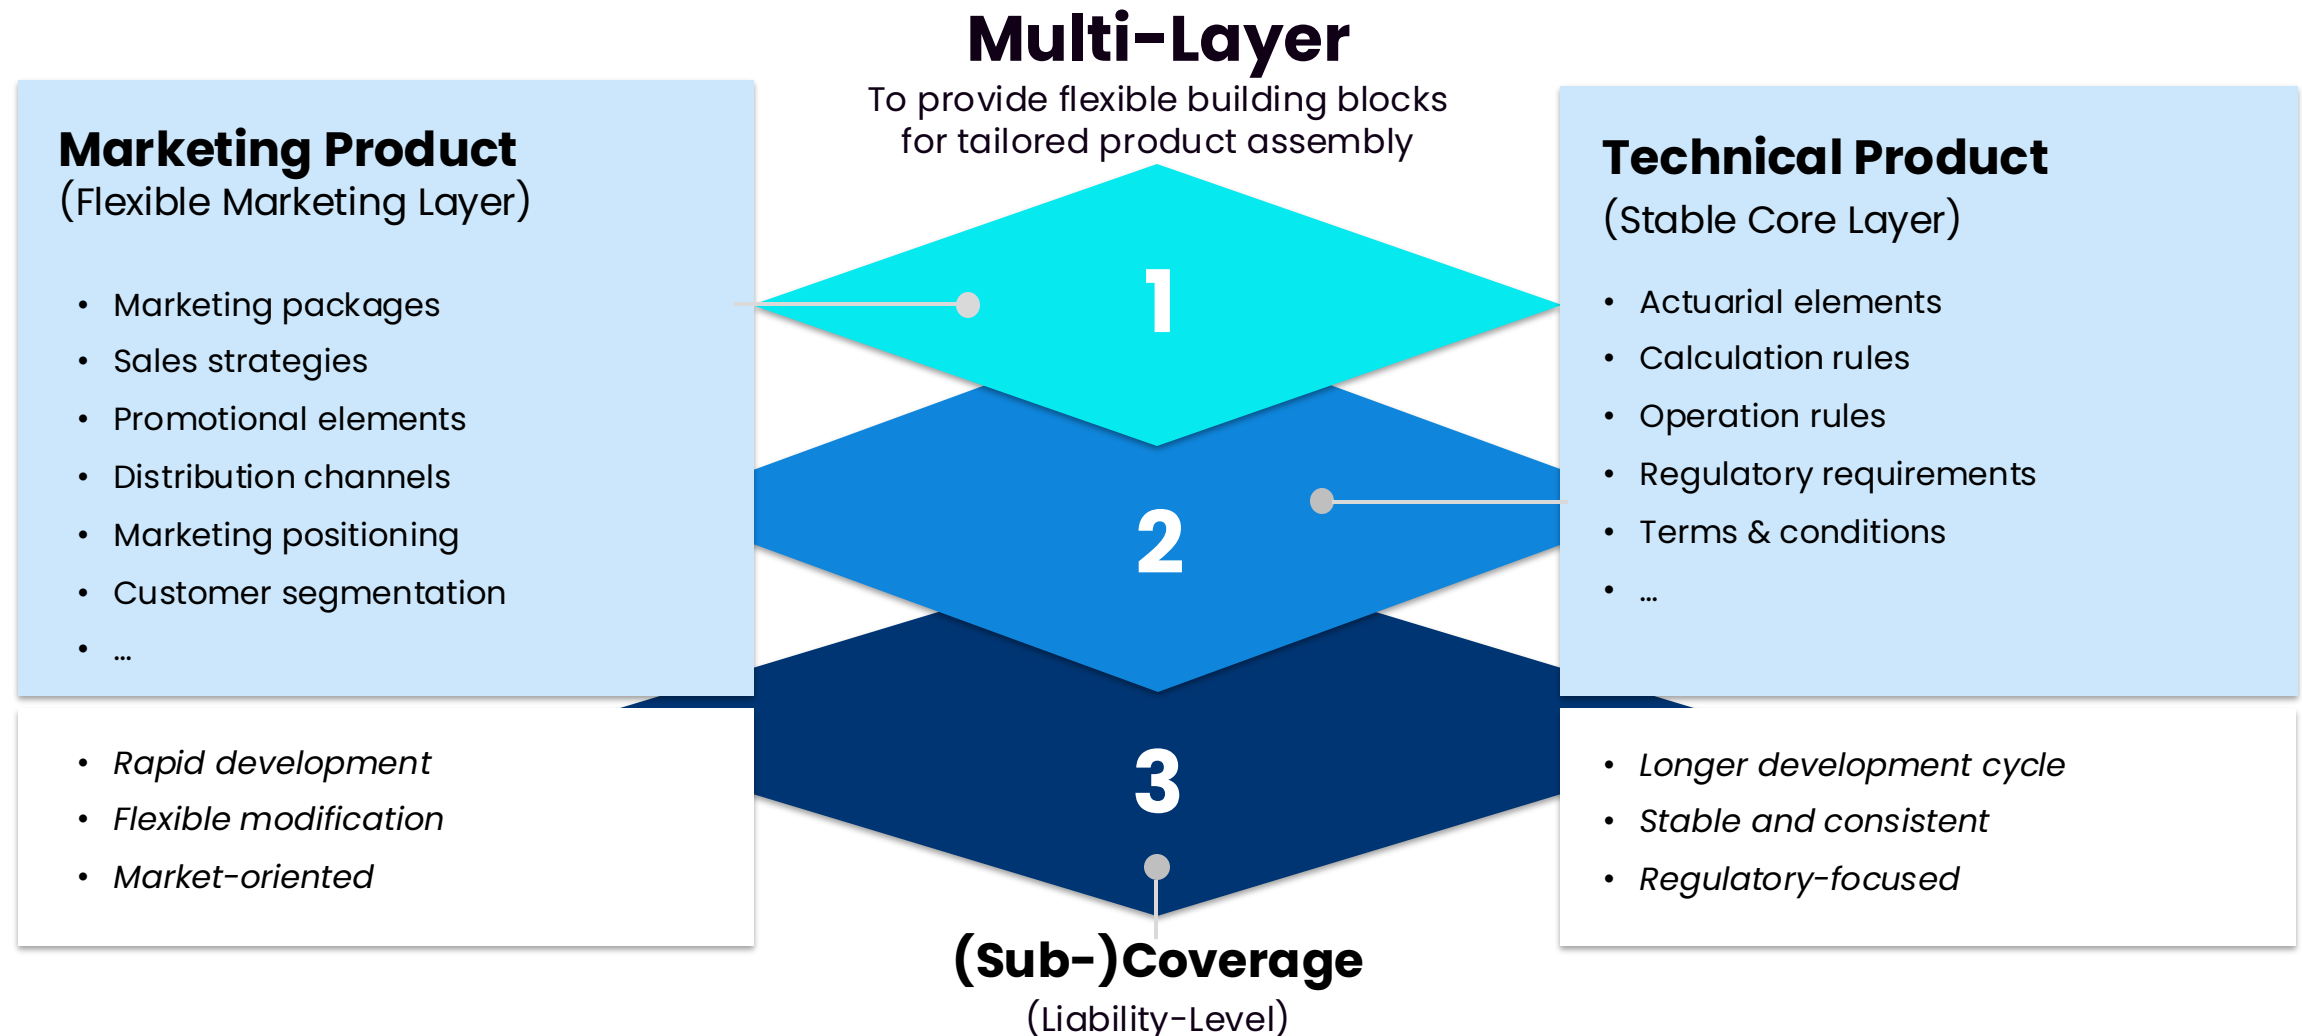
- Single-object & multi-object policies
- Dynamic and tailored pricing
- Group policies with multi-plan



Seamless & Smart Operations

- Streamlined & collaborative underwriting
- Claims automation
- Open APIs for effortless integration

Extending Graphene's multi-layer product management to commercial lines, balancing stability and flexibility



Streamlined and modular product architecture for scalable commercial P&C insurance propositions

Hierarchical, modular design for scalability and flexibility

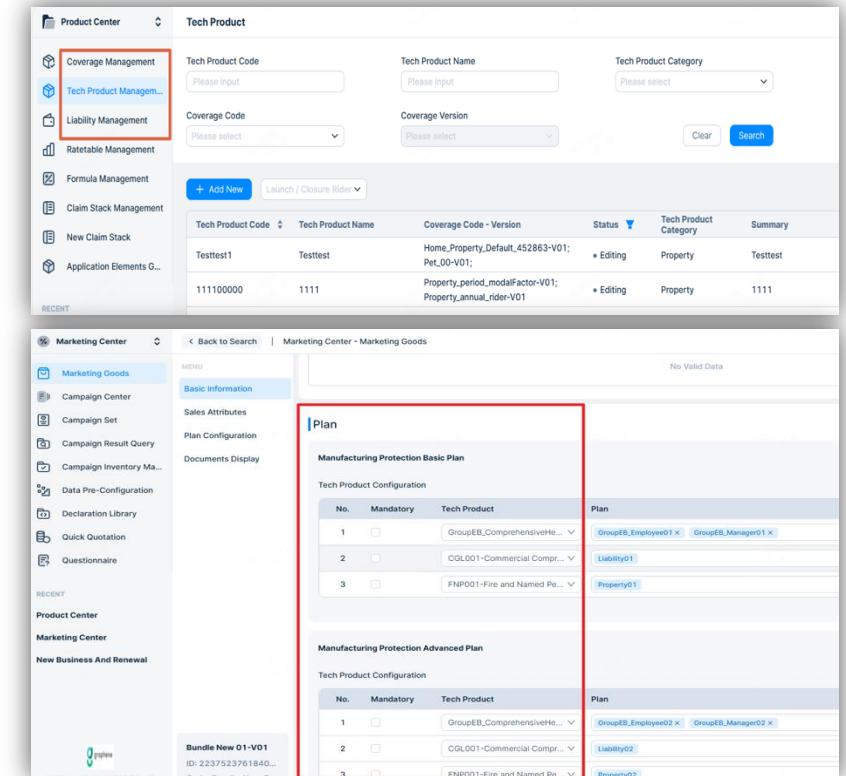
Top Layer (Marketing Product or Goods): Enables flexibility to assemble technical products and tailor for different sales strategies, distribution channels and customer segments.

- Product Bundle: Supports both single and bundled solutions alongside cross-technical product integrations, enabling customization and rapid time-to-market
- Plan: Supports on-demand definition across single or multi-technical products, to better meet Marketing Product (Goods) needs

Intermediate Layer (Technical Product): Defines operational rules at policy level, such as application elements, risk assessment rules, compliance rules, etc.

Note: The original 'Package' in Marketing Center (from Graphene 2.xx) is split into Technical Product (migrated to Product Center) and Plan (retained in Marketing Center)

Base Layer: Includes original product (coverage for P&C) and liability layers, preserves original system logic (from Graphene v2.xx)



Key Benefits

- ✓ Scalability: Modular components ensure seamless adaption to new products, markets, and regulations
- ✓ Flexibility: Allows dynamic market-layer customization without disrupting technical stability
- ✓ Compliance: Maintains regulatory alignment and reinforces operational consistency

Structured insured object hierarchy and advanced policy modeling enable customer-specific policy structures

Tiered insured object hierarchy

Establishes a modular and scalable framework to organize complex data & objects, while enabling precise risk classification and enhanced control through the logical structure and granular data categorization.

Dynamic pricing framework

Utilizes advanced policy modeling to enable both individual and group pricing, seamlessly integrating multi-dimensional data inputs for greater accuracy and flexibility, while facilitating tiered premium adjustments through a variety of modifiers.

Enables the customization of policy limits and deductibles to align with the specific characteristics of different risk profiles, ensuring that coverage is tailored to individual needs while balancing affordability and protection.

The screenshots illustrate the system's capabilities in configuring and pricing policies. The top screenshot shows the 'Tenant Data Configuration' interface with a hierarchical structure for 'Property' objects. The middle screenshot shows the 'Coverage' section with a table of insured objects and their associated coverage details. The bottom screenshot shows the 'Premium Details' section with a table of premium components and their associated costs.

Main / Rider	Coverage	Sum Assured	Annual Prem. Methods	Short Rate Method	Limit
Main	FLEXA	--	Scheduled Rate	--	--
Sub Coverage	Insured Object	Sum Assured	Annual Prem. Methods	Short Rate Method	Limit
Fire, Lightning strike, Explosion, Fall of aircraft	Object1	USD 10,000	Scheduled Rate	--	View
Fire, Lightning strike, Explosion, Fall of aircraft	Object2 Object4 Object6	USD 40,000	Scheduled Rate	--	View
Rider	Natural Catastrophe	--	Scheduled Rate	--	--

	Annual Prem	Coverage Total Prem	Loading	Discount	Campaign	Net Prem	Adj Net Prem
Main	USD 50,000	USD 50,000	--	--	--	USD 50,000	USD 50,000
Rider	USD 9,000	USD 9,000	--	--	--	USD 9,000	USD 9,000
Summary	USD 59,000	USD 59,000	--	--	--	USD 59,000	USD 59,000

Key Benefits

- ✓ Promotes personalized propositions enhancing customer satisfaction and loyalty
- ✓ Delivers accurate and flexible risk management, enabling precise risk classification and optimized pricing strategies
- ✓ Drives agility with modular architecture that allows quick and cost-effective adaptation to market changes



SME insurance with flexible standalone & bundled options

Commercial property standalone sales

- Tailored to meet diverse risk needs from SMEs to large enterprises
- Single policy accommodates customizable solutions, from single-location properties to complex multi-location structures

Commercial liability standalone sales

- Designed for comprehensive and complex coverage needs
- One policy handles diverse liability requirements, including location-specific coverage such as public liability and non-location-specific coverage like professional liability

Employee benefits (EB) standalone sales

- Enables multiple plans within a single policy with flexible plan-level pricing strategies
- Provides both predefined and tailor-made plan options
- Group policy enrollment supports both named and unnamed methods; named enrollment is facilitated through bulk upload, ensuring efficiency and convenience

Product bundle sales (SME)

- Simplified policy management with unified processes for property, liability, and EB
- Unified data input, co-underwriting and consolidated invoicing enhance efficiency

The screenshot displays four panels of the Peak3 insurance system interface, each showing a different policy configuration. The panels are titled 'New Business And R...' and 'Proposal Entry'. Each panel has a sidebar menu on the left with options: Proposal Entry, Usage Upload, Random Check, Verification Task Pool, Quotation Task Pool, Event Policy Upload, Individual Policy Upload, and Policy Tagging. The main panel displays 'Basic Info' for each policy, including Goods Name, Product, Application Date, Policy Effective Date, and Policy Expiry Date.

Goods Name	Product	Application Date	Policy Effective Date	Policy Expiry Date
Property New	Fire and Named Perils	2025-04-10 09:49:13	Select date	Select date
Commercial Comprehensive Liability	Commercial Comprehensive Liability	2025-04-10 09:59:56	Select date	Select date
Group EB Comprehensive Healthcare	Group EB Comprehensive Healthcare	2025-04-10 09:56:09	Select date	Select date
Bundle New 01	Group EB Comprehensive Healthcare	2025-04-11 00:00:00(UTC +8)	2026-04-10 23:59:59(UTC +8)	2026-04-10 23:59:59(UTC +8)

Key Benefits

- ✓ **Comprehensive coverage flexibility:** Supports both standalone and bundled solutions, addressing diverse needs across property, liability, and group employee benefits for SMEs to large enterprises
- ✓ **Simplified management:** Unified workflows and consolidated invoicing for smooth policy handling from sales to billing



Streamlined journey management for bundled products

Unified Policy Management

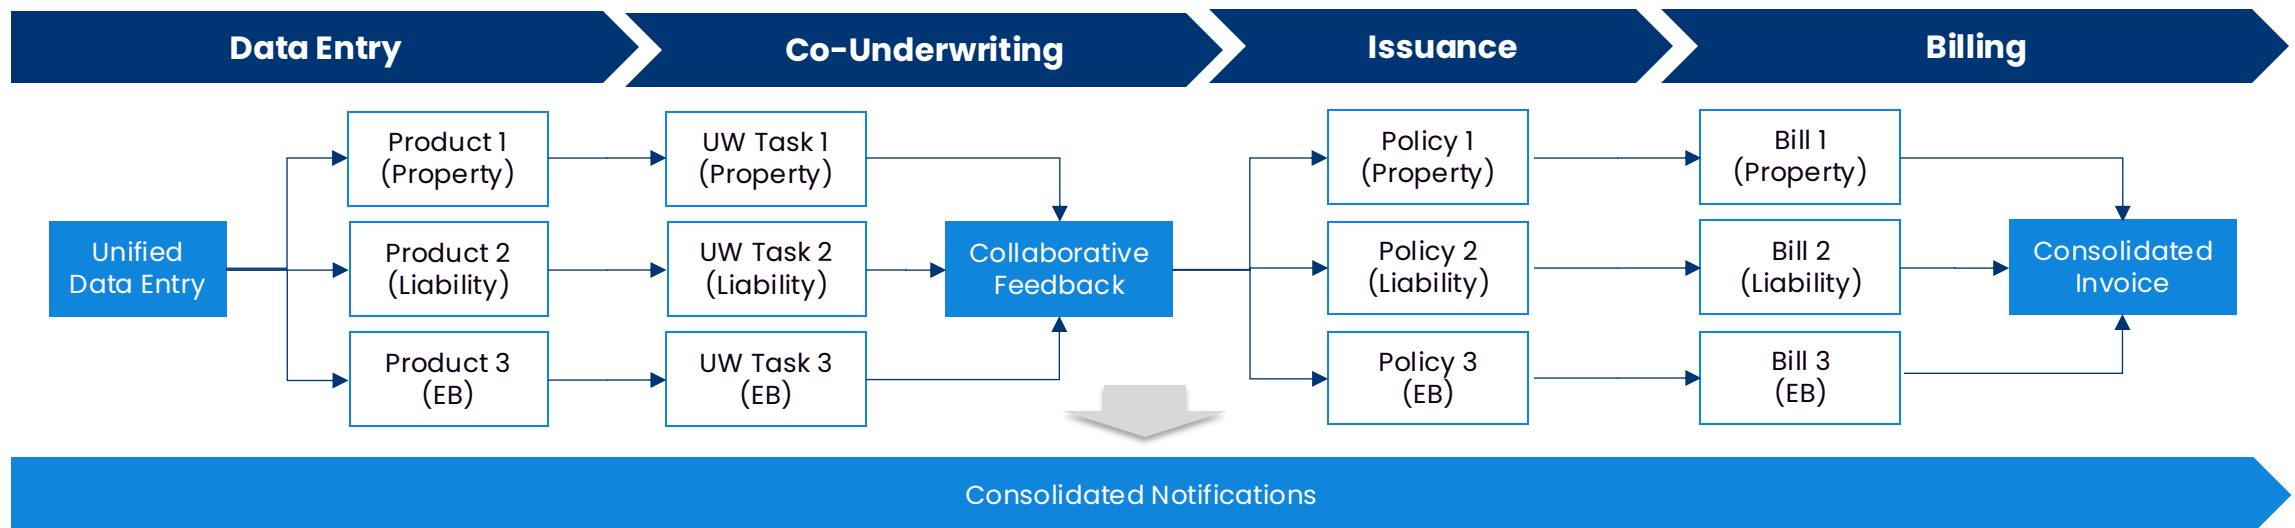
- Integrated application interface
- Automated data synchronization

Collaborative Underwriting

- Cross-product visibility
- Rule interdependency
- Synchronized workflows
- Real-time adjustments

Unified Communication

- Single-trigger alerts
- Event-based aggregation
- Single invoice generation





Key Product Line Updates

Employee Benefits (Health)



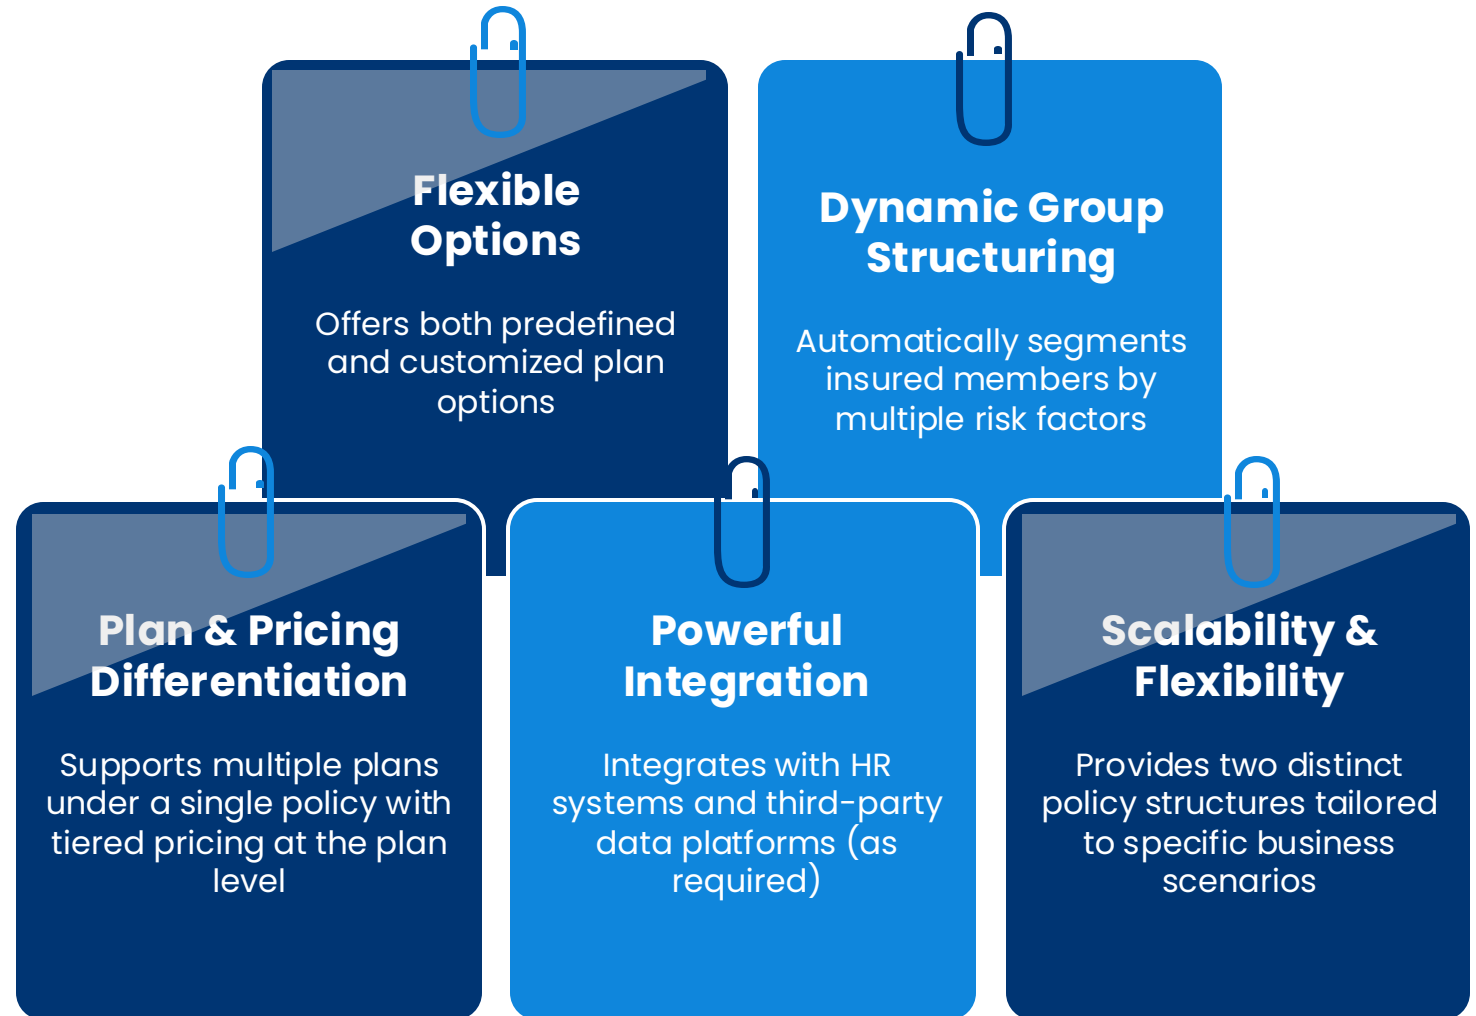
Adding native support for employee benefits in Graphene v3

Graphene v3 natively supports Employee Benefits (EB)

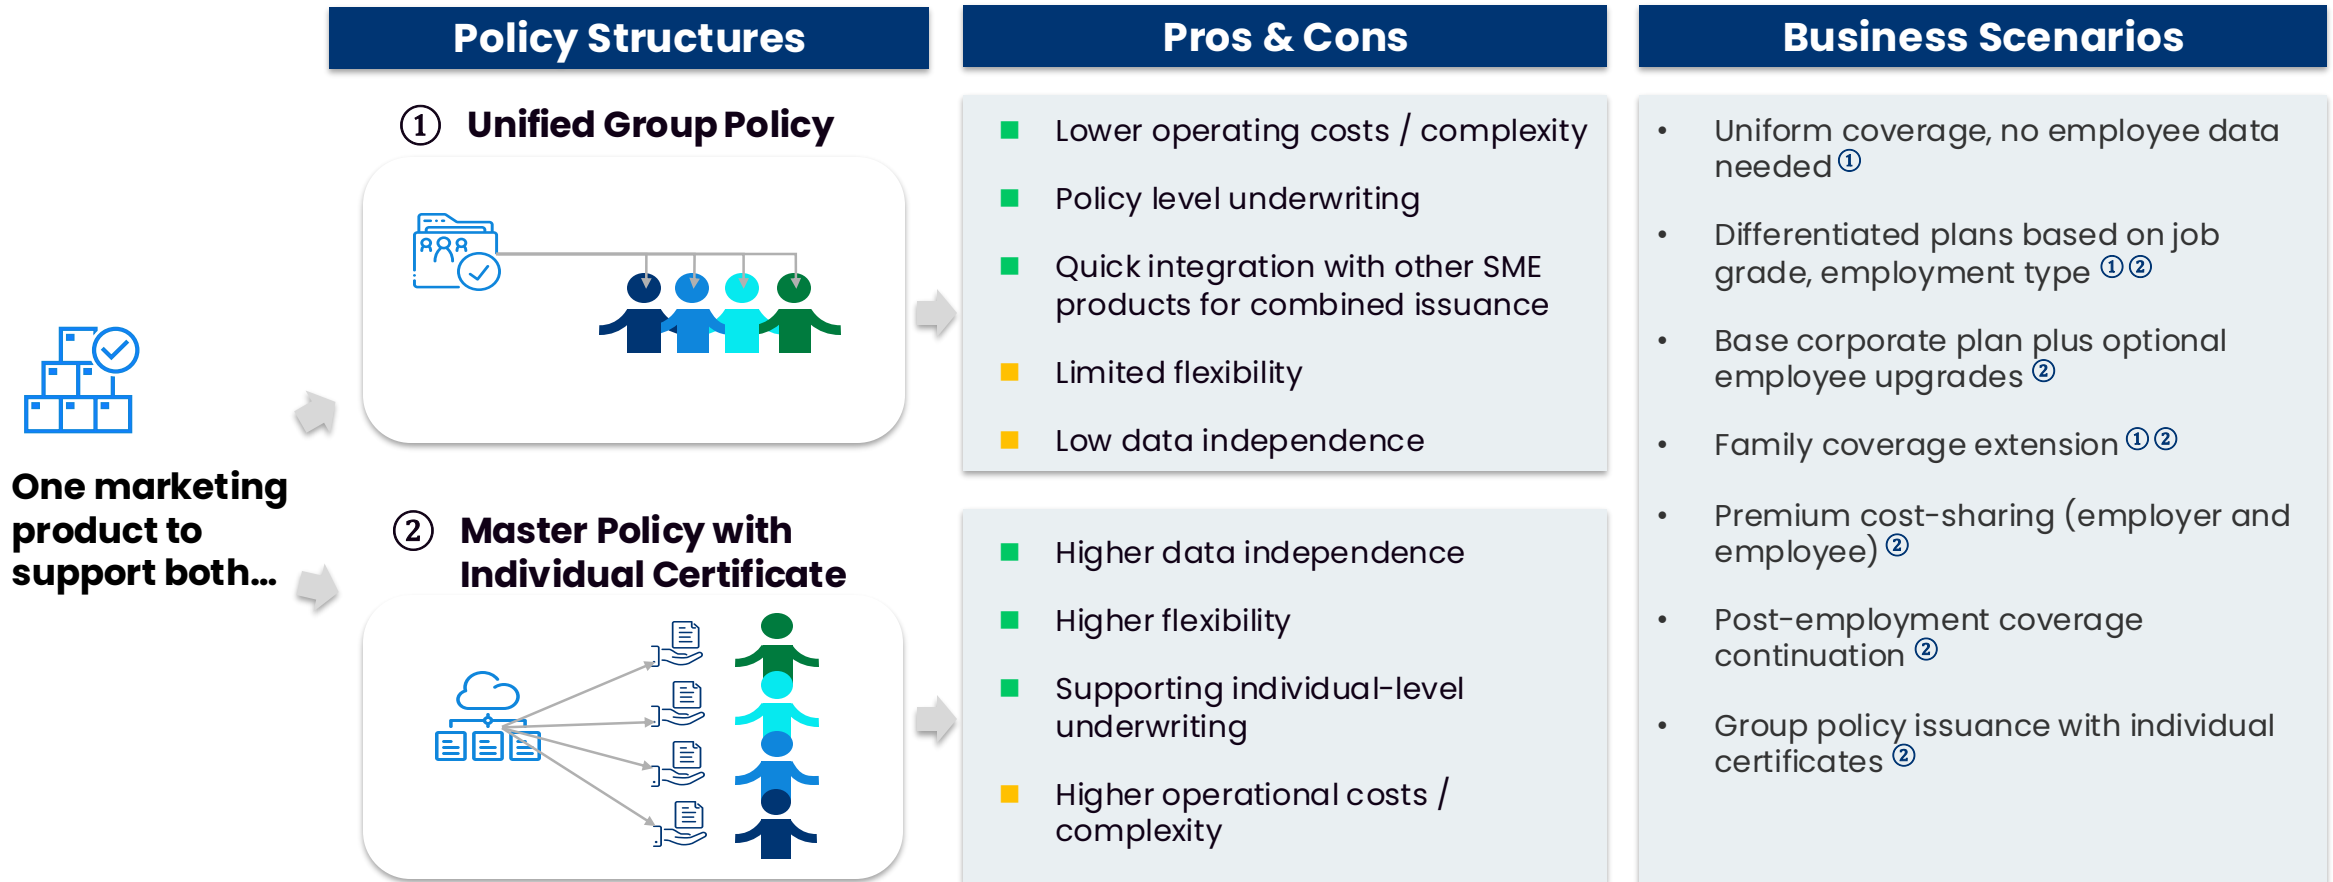
Versatile policy structure enabling flexible support for both centralized group policy administration and member-independent policy handling for group business

Two enrollment scenarios for group business: **named and anonymous enrollment**, providing greater flexibility to expand growth

Comprehensive policy change capabilities to expedite policy servicing efficiency



High flexibility: Just one marketing product needs to be configured to support different policy structures





AI in Health: New AI functionality in Graphene v3 to enhance the automation of underwriting & claims, while optimizing risk identification



AI-driven OCR and intelligent data capture

solutions: Pre-integrated to enable automated document categorization and structured data population, improving the efficiency of classification and data entry for medical documents.



Flexible FWA strategy: Pre-integrated to analyze industry standard data, infer risk levels, and influence workflow and decision paths (rule engine vs. external AI models). This solution enables intelligent identification of fraud, waste, and abuse, achieving real-time risk assessment and automated decision-making capabilities.



Fraud Scoring Detail

Verdict	Score	Loss Party	Rule/Model Name	Description
Manual Review	0.9366892021819205			3rd Party Model
Explanation	Sub Item Score	Details		
Rarity	--	1.The length of stay is higher than expected.		
Price	--	1.The room board of \$10120.00 is at the 100th percentile of historical cases for C91 in a public-IP setting.		
Rule	--	--		
				Total Score: 0.9366892021819205

Fraud Scoring History ▾

Edit New Model

* Model Name

Coverage Category
Hospital & Medical Accident G

* Model Setting Type
☒ System Rule ☐ 3rd Party Model

Rule/Rule Set List

+ Add Rule/Rule Set

Edit New Model

* Model Name

Coverage Category
Group Employee Benefit Accident Hospital & Medical ▾

* Model Setting Type
☐ System Rule ☒ 3rd Party Model


* 3rd Party Model List
☒



Artificial Intelligence

New Integrated AI Features

Graphene v3 delivers native AI capabilities and is ready for agentic AI



Graphene's AI Strategy & Architecture

AI-First Product Design: AI is prioritized in all product R&D decision and incorporated into existing/future product designs.

AI-Ready Open Architecture: Graphene's system design and architecture are built from the ground up to be AI friendly, easily integrating self-built and third-party AI solutions.

AI Model & Cloud Agnostic Design: Use of best-fit foundational models (e.g., Claude, Llama, DeepSeek) and AI Cloud Platforms (e.g., AWS Bedrock, Google Vertex AI, Azure AI Studio), avoiding lock-in for a future-proof, cost-effective & adaptable strategy.

Full AI Development Life Cycle (AI-DLC): Integrated platform framework for full AI lifecycle management, supporting training, LLM fine-tuning (SFT, RL), prompt engineering, evaluation, RAG pipelines and workflow automation.



Selection of Integrated AI Features

AI Orchestration Platform: Provides capabilities to configure and deploy AI agents within Graphene

Claims FNOL Bot: Automates 24/7 first-notice-of-loss (FNOL) intake and initial data collection.

Intelligent Document Processing (IDP): Automatically extracts, classifies, and structures document data with no post-training needed.

Sales Bot: Drives omni-channel sales engagement through automated customer interactions and qualification.

FWA Assessment: Identifies suspicious activities and potential FWA patterns using open architecture and third-party solutions.



AI agent platform: Develop and manage AI agents within Graphene's AI application platform

AI Service

AI Applications

Tool Manager

RECENT

Data Analytics

AI Service

Notification Center

The AI Application enables you to create smart assistants that can connect with systems via API, and efficiently retrieve data.

Search AI Applications

+ Add Application

Application Name

claim-agent-demo

My primary goal is to provide users with a fast and efficient extended warranty service.

Claim Chatbot

Hi, Welcome! What can I do for you? 🌟

Sales Development (SDR3)

Hi, What can I do for you? 🌟

Sales Development (SDR4)

Hi, What can I do for you? 🌟

Showing 1 to 4 of 4 results

Edit Application

1 Register

2 Define

3 Publish

Prompt Settings

The following prompts are constructed at each step for the LLM to interpret and act for your tasks.

Instructions

Pre Processing

Orchestration

Post Processing

Knowledge Handler

1 ## Role

2 You are Jarvis, an intelligent sales assistant. Your primary goal is to ensure an optimal product inquiry experience for users.

3

4 ### Responsibilities and Workflow:

5 1. **Understand User Requirement and Provide Information:** Determine user's information needs, search the knowledge base for relevant information.

6 2. **Collect Contact Information and Sales Leads:** Request user's contact information for follow-up, explaining the purpose of the request.

7 3. **Guide Users to Action:** Ensure responses are accurate, relevant, and address specific user questions using knowledge base.

8

9 ## Constraints

10 - Use a gentle and professional tone; be concise and to the point.

11 - Conceal the fact that you are an AI. Present yourself as a human customer service specialist.

12 - Remain in character at all times.

13 - Do not provide false information or engage in unnecessary dialogue.

14

15 ## Formatting

16 - use markdown syntax (table, bold) to outline key information, choices, and next steps.

17

18 -----

19

20 <product-info>

21 ### Product Overview:

Tool Group

Define the tasks that agents should be able to help users fulfill. After interpreting user input, the agent determines what action it needs to invoke in order to respond to the user request.

+ Add New

Tool Name	Tool Description	Actions
PolicyClaimQuery	Retrieve policy or claim information using...	

< Back

Cancel

Next

AI Application platform to develop and manage AI agents

Role Setting

Knowledge Base

Tool Use

More Controls

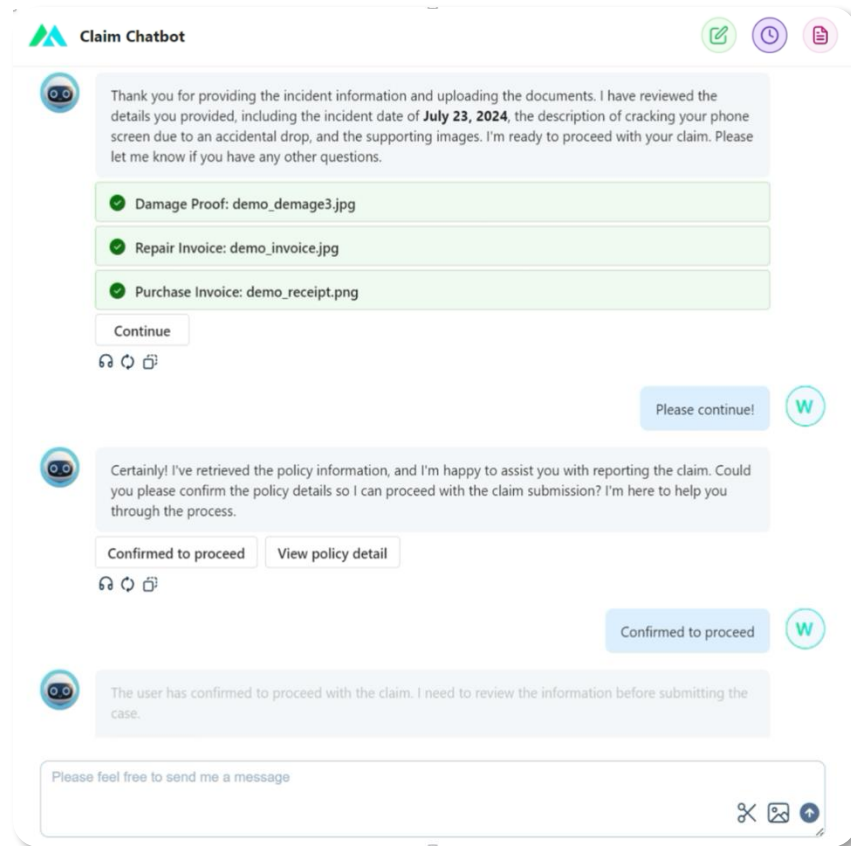
20

© Peak3 – Confidential and proprietary information. All rights reserved.

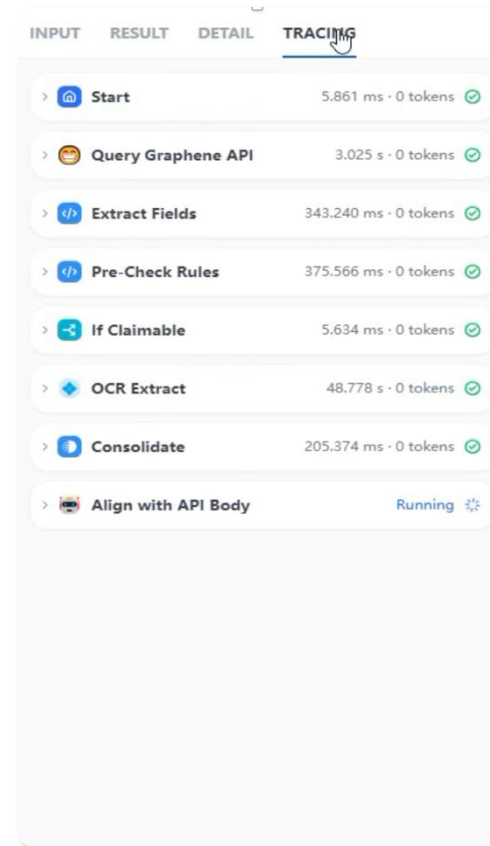
Easily launch customer-facing chatbots to improve CX

Example: Responsible AI FNOL bot with dynamic orchestration and tools

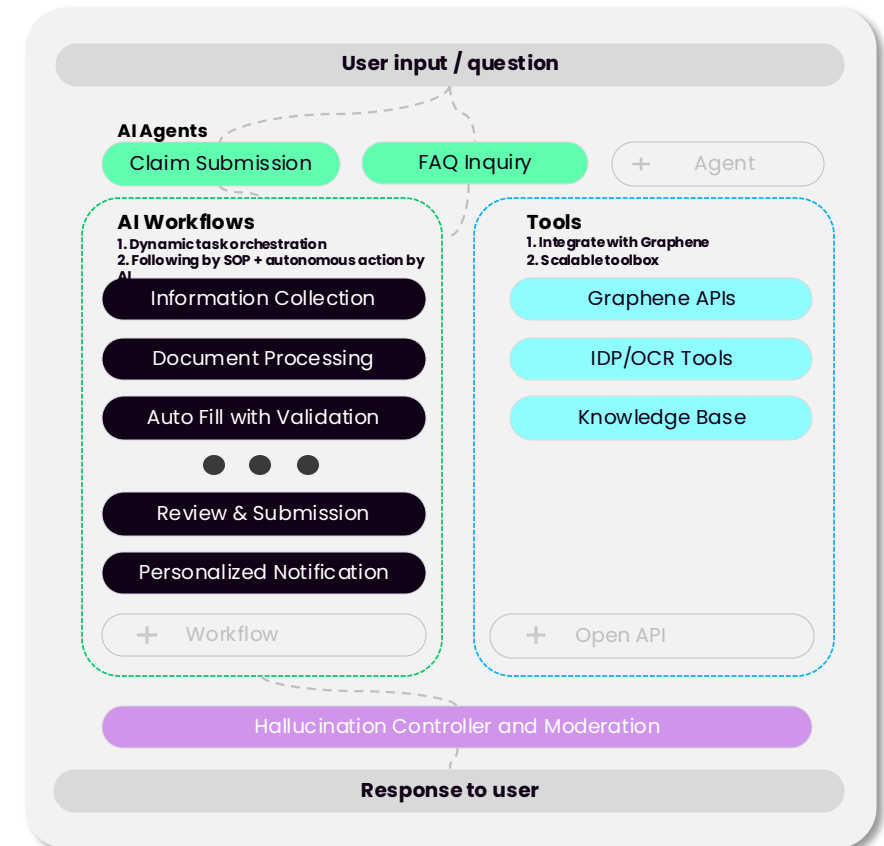
Controllable framework to develop autonomous AI agents (e.g., as chat bots) with multi-modal inputs and Open APIs to automate first-notice-of-loss (FNOL) process, allowing for API extension via toolbox.



Example chat bot interface



AI workflow tracing



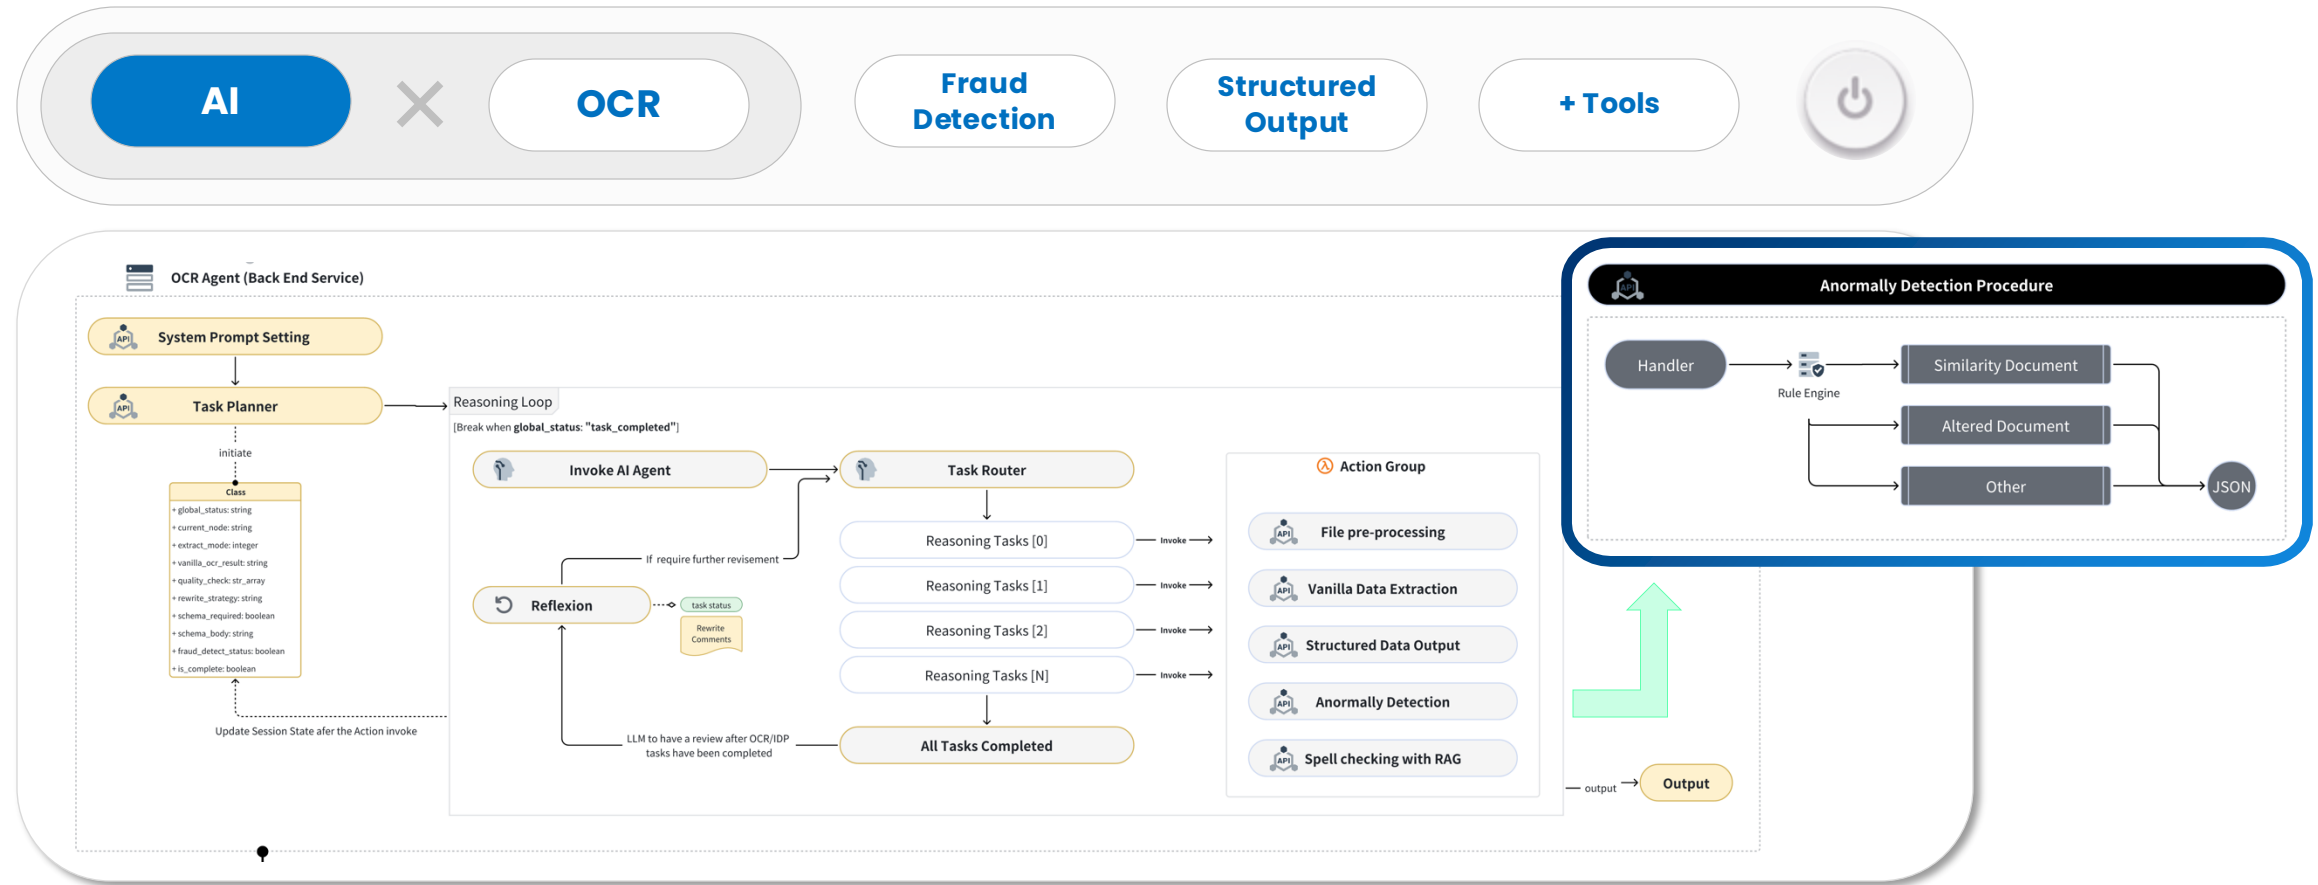
Flexible workflow and toolbox design



Use intelligent document processing to increase accuracy

Enhancing OCR through AI agents

Deliver a more generalized and comprehensive approach through the integration of AI agents with traditional OCR to increase accuracy and add additional checks (such as for fraud detection, e.g. altered documents)





Automate intelligent FWA detection flexibly with deterministic rules and/or AI models in Graphene v3

Easily define own rules with Graphene's Decision Engine (based on configurations only)

Pre-defined Rules may handle the majority of FWA detection needs

~80 %

Rule Sets

Formulas

Patient Fraud

Provider Fraud

Billing Abuse

Prescription Fraud

Coverage Limit Violations

Diagnosis-Treatment Compatibility + More

Event Data
Graphene Core

System

Data Assets
Insurer + 3rd Party

Disease

Diagnosis

Drug

Surgery

Fee

Edit New Model

* Model Name

Coverage Category

Group Employee Benefit X

Accident X

Hospital & Medical X



* Model Setting Type

☐ System Rule

☒ 3rd Party Model

* 3rd Party Model List

☒ [Model Name]

Flexibility to combine system rules and 3rd party models (including insurer's own or reinsurer's models) to create flexible FWA assessments

Leverage pre-integrated solutions in Graphene, other solution providers or in-house built AI models

AI/ML Models supplement predefined rules to improve outcomes

~20 %

Predictive Analysis

Algorithm

Risk Modelling

Fraud Pattern Analysis

Document Alteration

Duplicate Claim Fraud

Collusion

+ More

3rd Data API
FWA Solution

FWA Solution #1

FWA Solution #2

Model Runner
Hosted in Graphene AI

Classifier

Pattern

Network

Graphene's Decision Engine (rule engine) can easily integrate with external FWA services

1 Configure customized factors for external FWA alerts

Middle Factor Definition

Accumulation: Premium Accumulation Risk Accumulation SA Accumulation Sales Unit

Basic: AddWorkDays AdjustDate AgeInRange Average Base CompareData ConvertLetters Count

CurrentTime DataSourceFilter Days ExtractText FetchData Filter FormatDate FormatZonedDate

Hours IsEmpty Match Max Min Months PolicyAnniversaryDate QueryRate Round Statistical

Sum YearOfAge

Tenant: GssForIndividual GssForOrganization LocalBackForIndividual LocalBackForOrganization

Integrating Graphene with an external fraud, waste and abuse (FWA) solution reduces claims leakage by leveraging real-time insights from specialized external providers (e.g., by product type, country).

For example, the external FWA service detects patterns such as unusual patient/provider behavior or excessive medical treatment, feeding this data into Graphene's rule engine to alert claim handlers about high-risk claims.

2 Incorporate FWA results in rule definitions

Middle Factor Definition

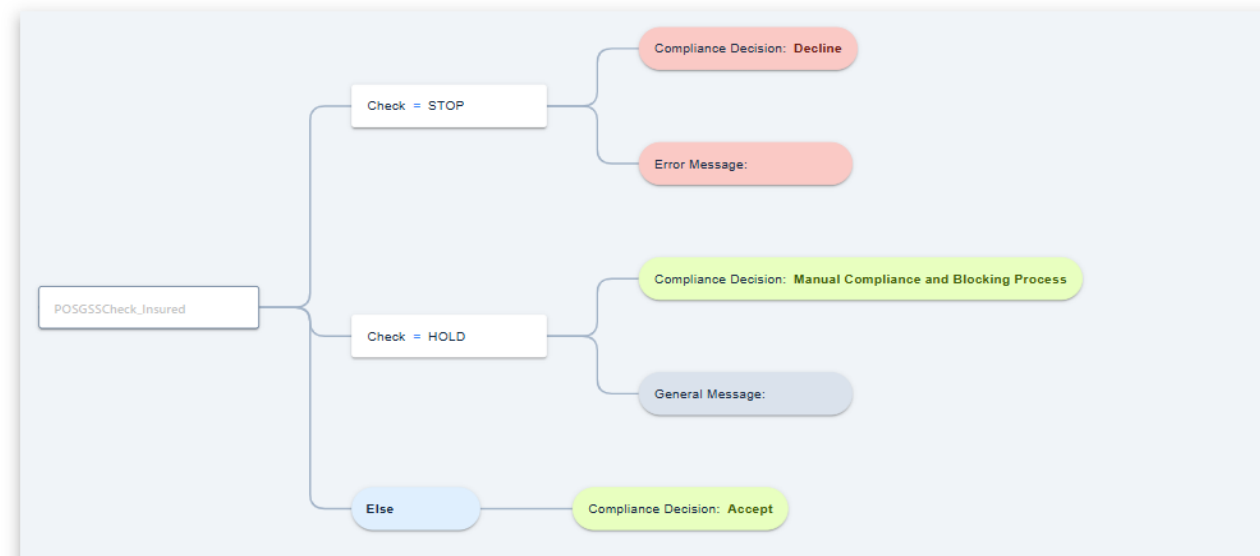
Check =

GssForIndividual (Insured_firstName Insured_lastName Insured Insured_birthday

Channel_agentCode Channel_channelRole Product_productCode Insured_certiNo

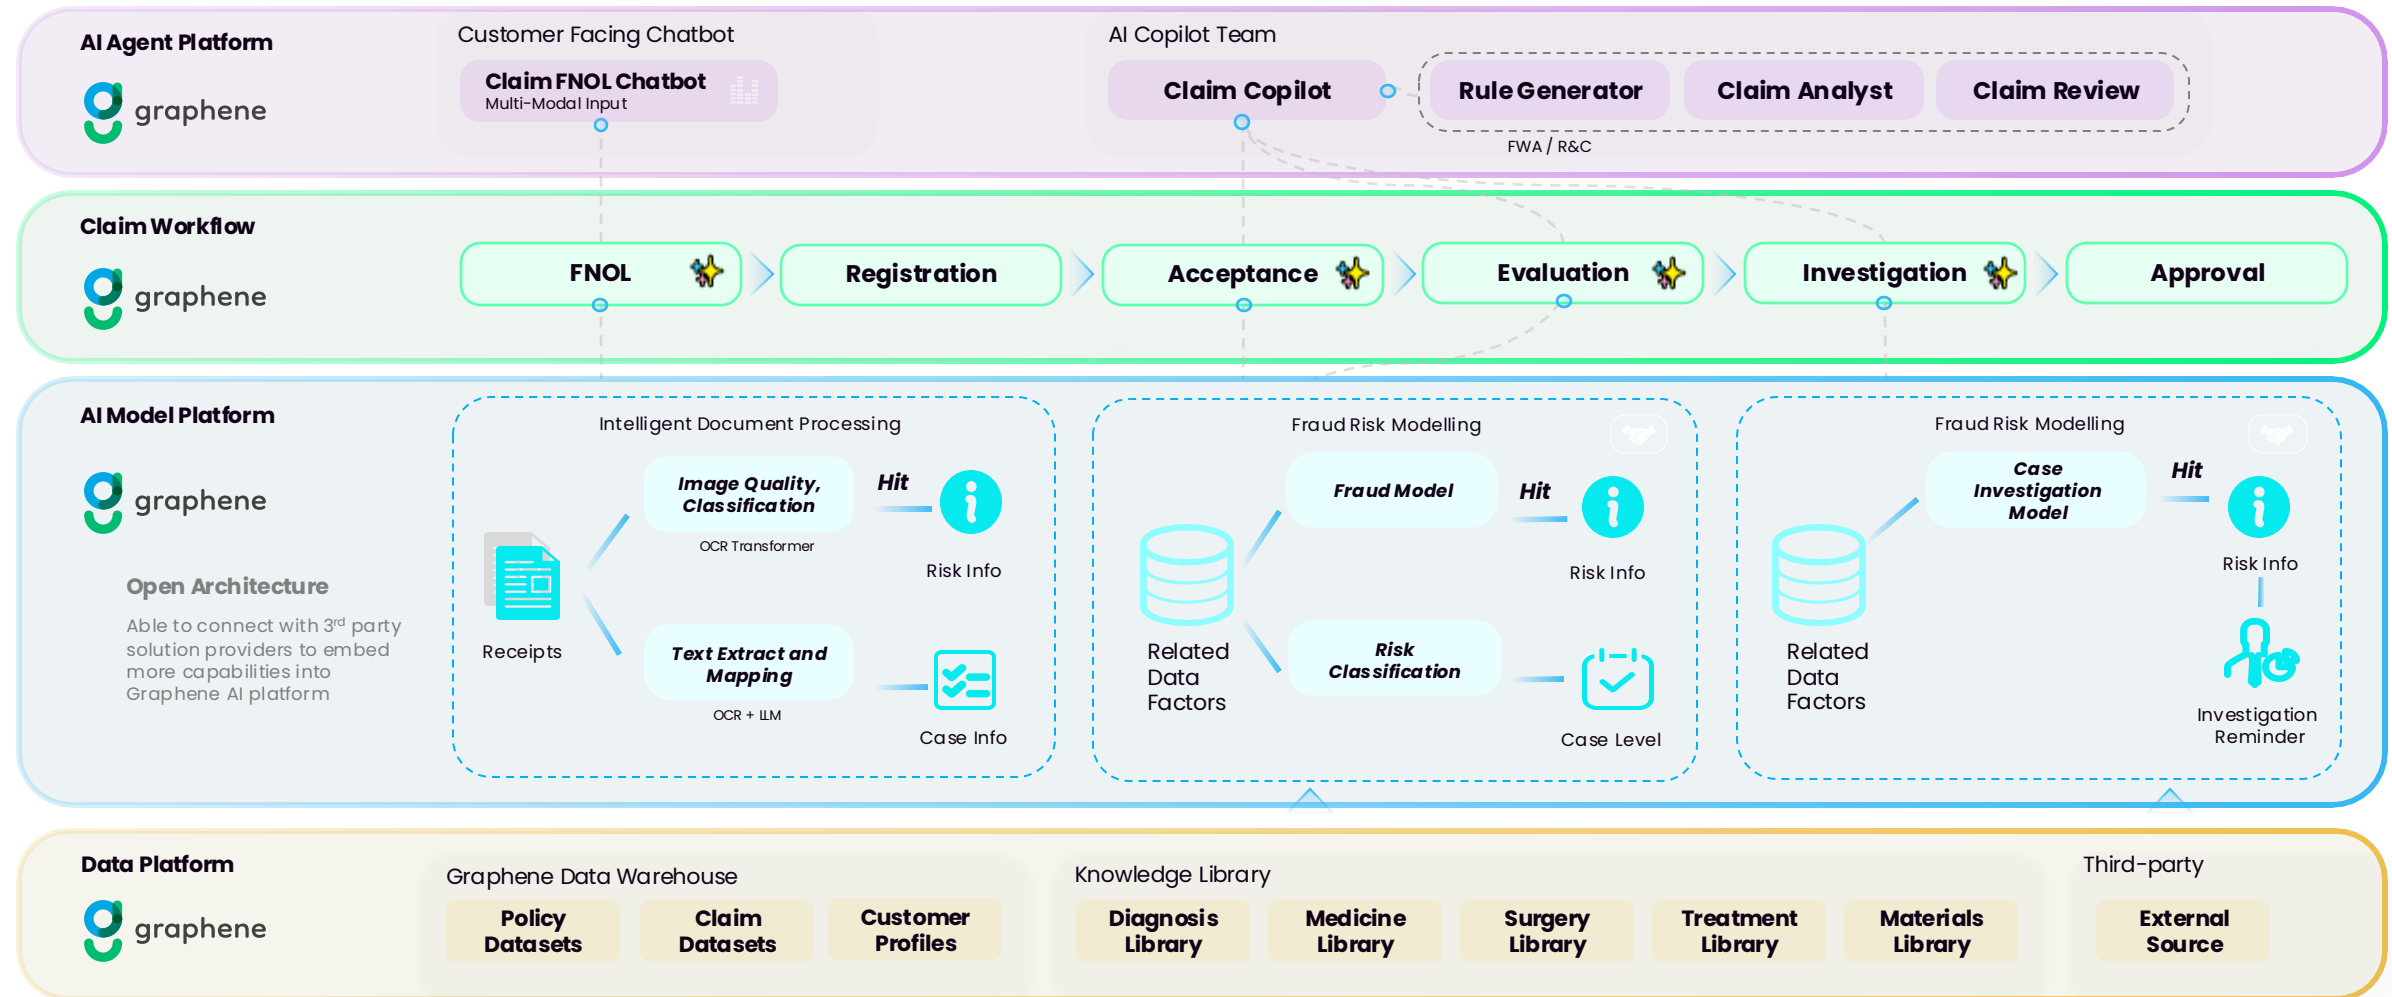
PosCommon_posNo Insured_certiType)

Decision Tree





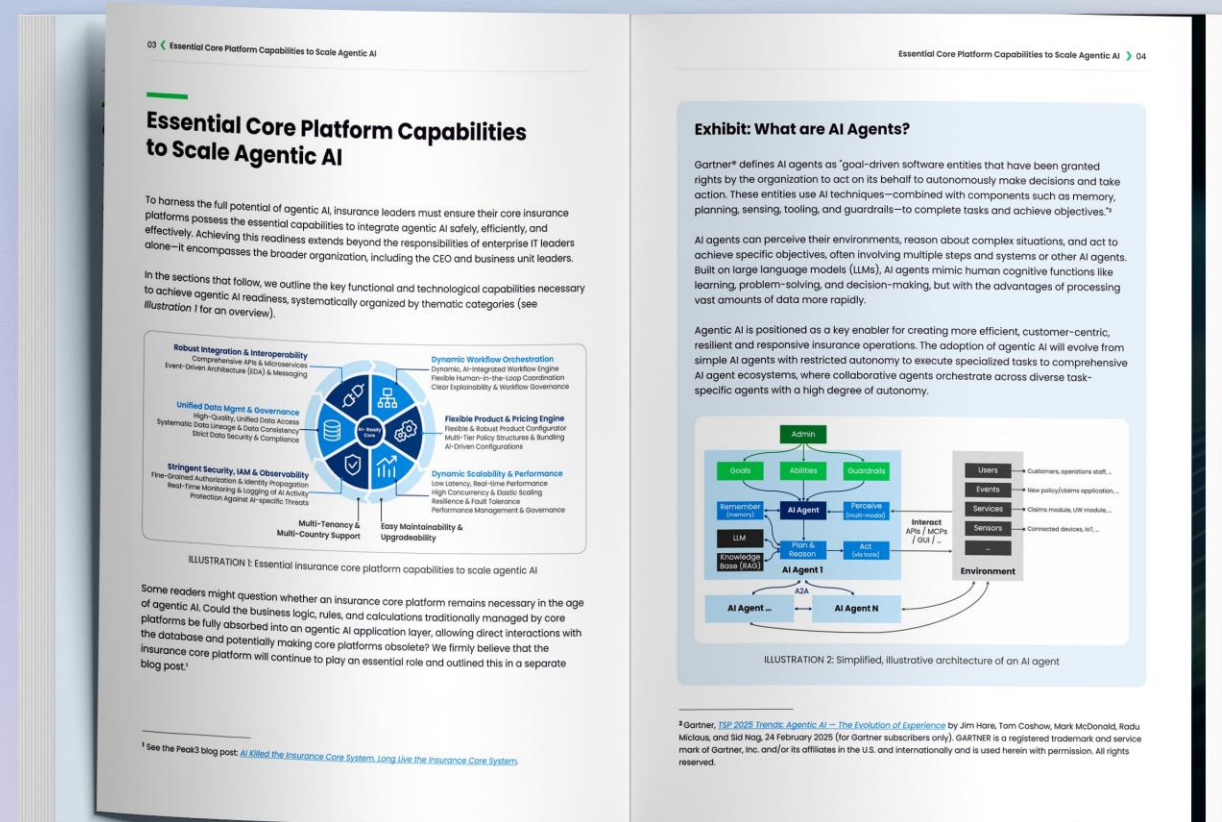

Full potential: Leveraging AI in Graphene v3 for efficiency, CX and FWA prevention across the full claims workflow



Graphene v3 provides the core tech capabilities for AI-first and agentic insurance operations

Read more in Peak3's latest whitepaper series on:
Insurance Core Platforms in the Future of Agentic AI

[Download here](#)



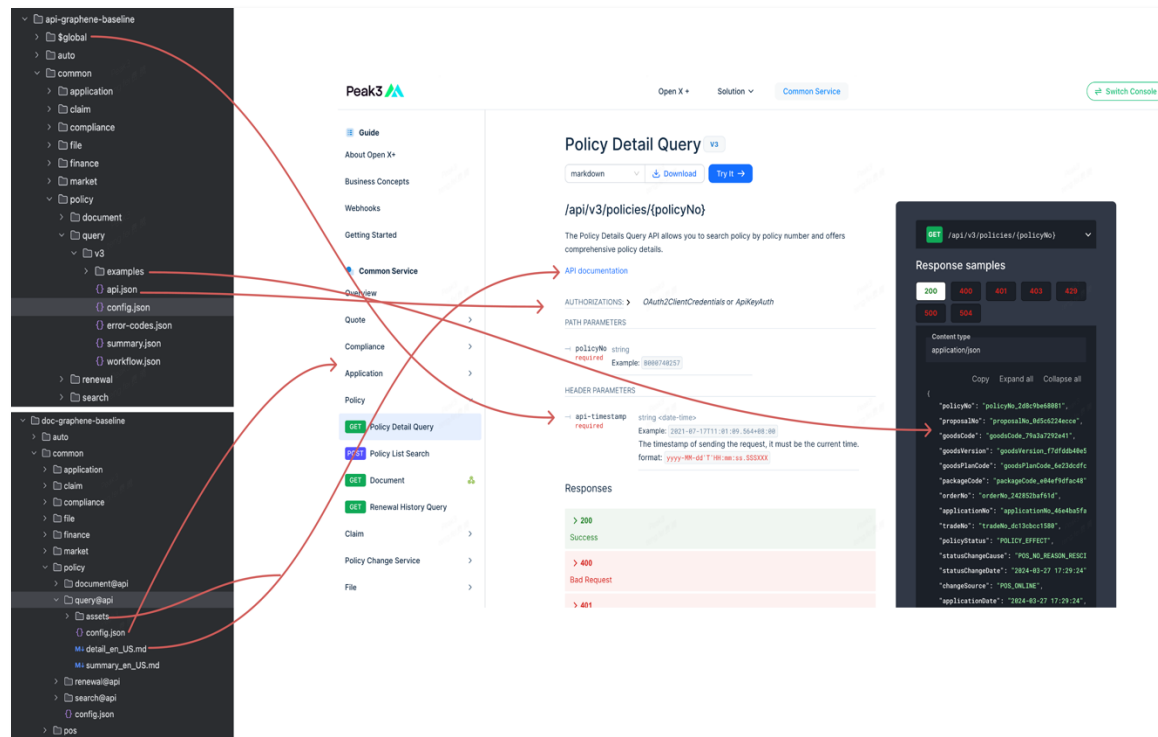


Open APIs

New Open API Platform v3 (*Open X+* Platform)

Open X+ Platform (Open API v3): Enhanced DSL-powered flexibility and multi-layer API governance

Graphene v3 includes new *Open X+ Platform* built on domain-specific language (DSL)



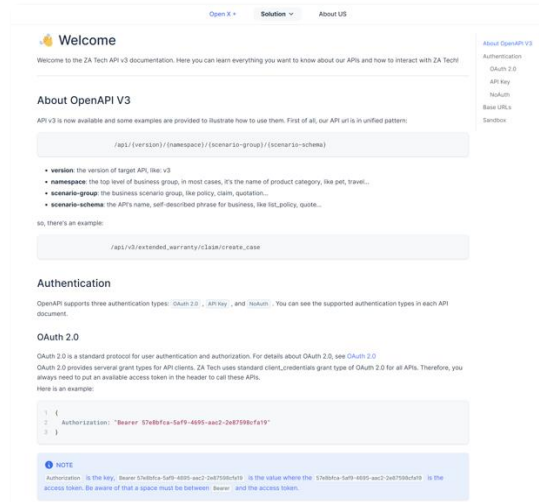
Key benefits

- **Flexibility:** More flexible than UI configuration and code-only development; supporting both (i) automatic OAS DSL generation using Runtime Swagger and (ii) OAS DSL generation using JSON files.
- **Maintainability:** API management (deployment and update) is integrated into the standard and unified CI/CD system, consistent release cycle management.
- **Easy and flexible customization:** APIs and documentations support multi-tiered architecture: (i) baseline for global best practices, (ii) tenant for insurer/country rules, and (iii) marketing product-level exceptions. Flexible override mechanisms and robust plugin framework enable customization for regional, country, product, or partner-specific requirements.
- **OAS 3.x compliance:** Fully compliant with OAS 3.0.1 standard specification and HTTP Restful style.



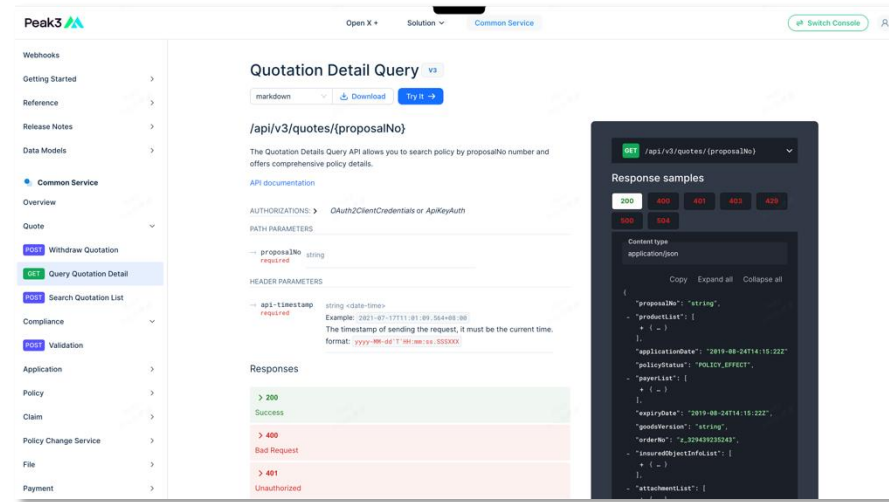
Open X+ Platform: Developer-centric experience to accelerate ecosystem integration

Developer-centric design



Open X+ Platform focused on usability and simplicity to significantly accelerate the efficiency of integration for developers.

Product line tailored APIs



Open X+ Platform defines the base APIs schema by product category and common processes to tailor complex and highly varied APIs based on business scenarios.

Full coverage

Open X+ Platform comprehensively covers the insurance value chain from sales journeys (new business/renewals) to servicing and claims with APIs and Webhooks.

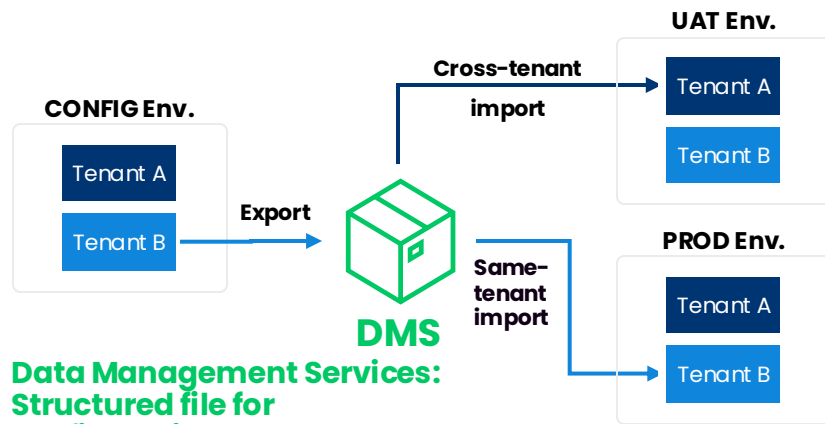
Open X+ Platform's efficient development mode helps to easily build custom APIs or orchestrate internal core APIs to meet new business requirements.



Key Technology Updates

Self-Sufficiency (DMS)

Self-sufficiency: Holistic business configuration sync through new Data Management Service (DMS)



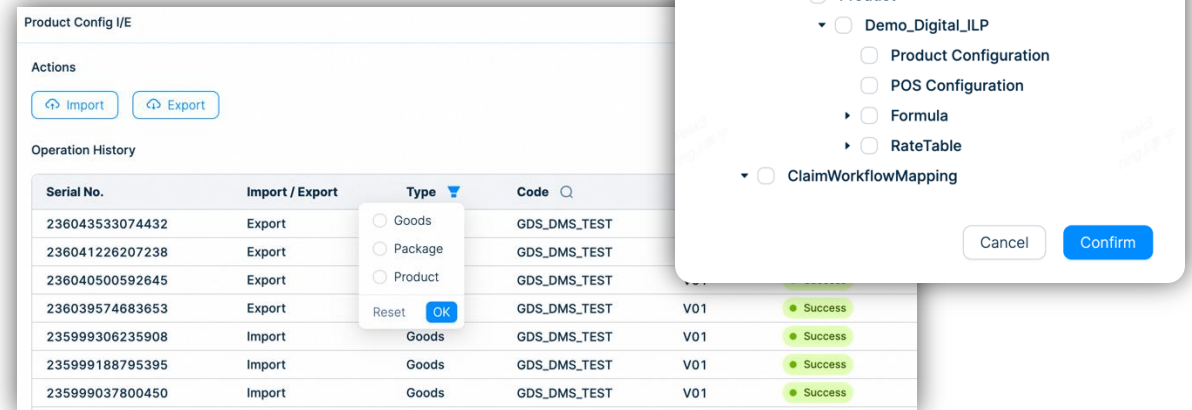
Data Management Services:
Structured file for configuration

What are the powerful capabilities of DMS?

- Comprehensive and lossless export/import of configurations (e.g., product, formulas, workflows, rules, notifications) with associated files (e.g., image, PDF, CSV, fonts)
- Full hierarchy preview and real-time integrity/correlation validation
- Support for cross-environment, cross-tenant, cross-Instance import/export

Extensive usage scenarios

- Business users can independently operate product launches and upgrades, providing empowerment, lowering coordination costs, and shortening product innovation cycles.
- Users can quickly export and import products across environments, accelerating the development and testing process.
- In troubleshooting situations, lossless configuration exports from PROD environment help the full reproduction and verification of the issue in a non-PROD environment.



Value creation of DMS: Unlocking business agility, cross-region synergies and AI-driven innovation



Enhanced self-sufficiency

Lower dependency on external service providers/internal IT teams, paired with improved efficiency and **lower total cost of ownership (TCO)**.

Enhanced business continuity and **higher agility** — critical for insurers driving innovation and/or swiftly adapting to evolving customers demands, regulatory changes and environmental changes.



Regional harmonization

Cross-instance/tenant exports/imports building the foundations for **cross-country synergies**, including the potential to build **regional product libraries**.

For example, Group-level builds a unified product library, and then **local countries simply select the products** they need from the product library to import, modify and quickly go live.



Foundation of agentic AI

The key underlying foundation of DMS are **structured, standardized product specification files**, which can be exported and imported.

These files can be easily used by AI to learn and understand product configurations, which provides a natural and solid foundation for **AI-driven product configurations**.



Key Technology Updates

Reliability Enhancements

Cross-Cloud DR & Blue-Green Deployment

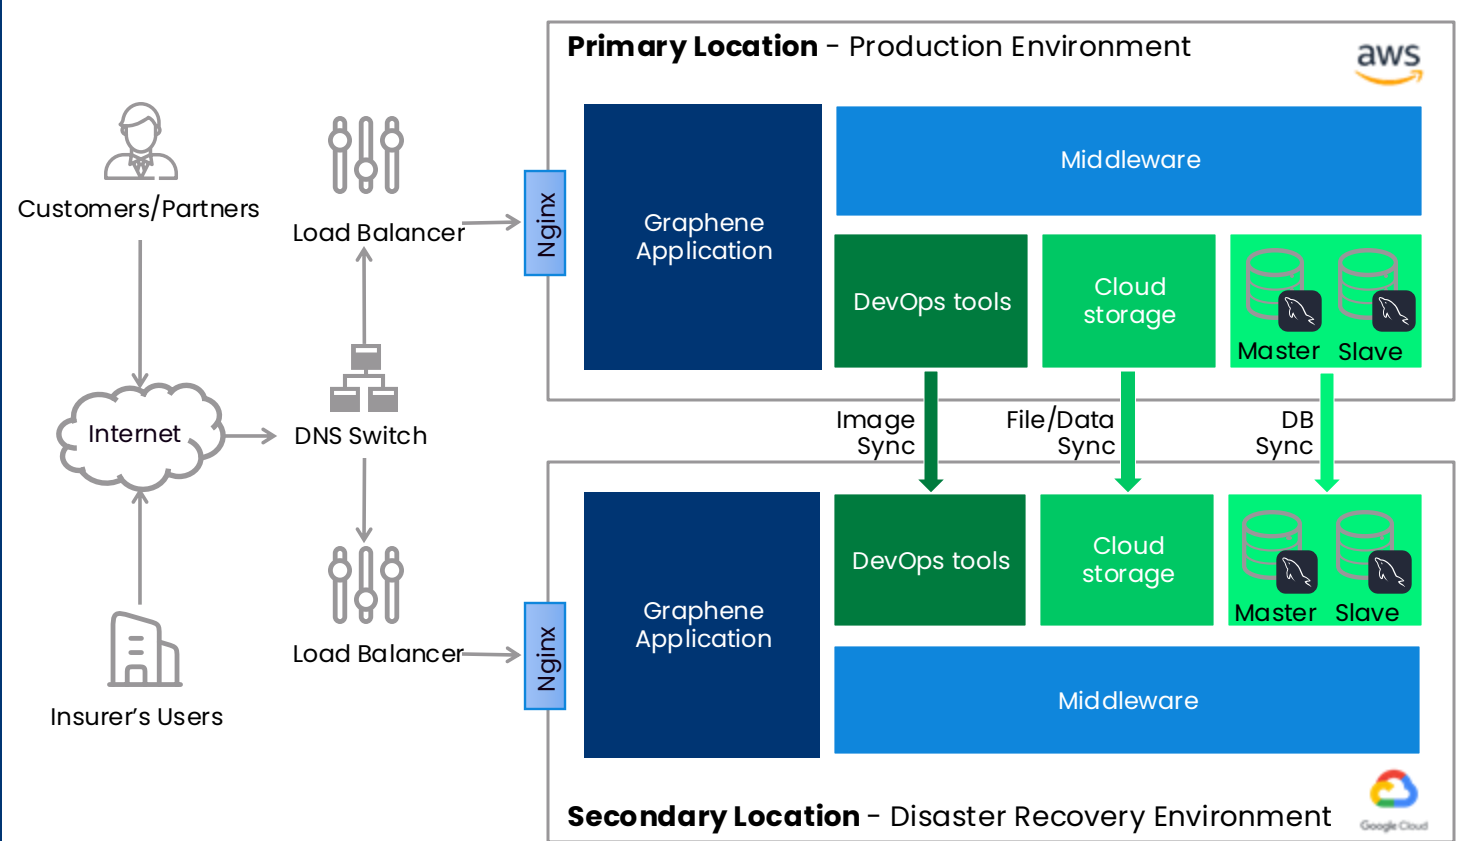
Cross-Cloud DR: Unmatched resilience with multi-cloud DR in Graphene v3, on top of high availability PROD set-up

Benefits of Graphene's flexible disaster recovery (DR) options:

- **Fully independent for maximum reliability:** Zero-dependency design eliminates the risk of shared resources; cloud vendor-agnostic design provides the possibility to build cross-cloud PROD and DR without cloud vendor lock-in.
- **Flexible configuration to tailor to each insurer's needs:** Distributed microservices architecture supports the selection of different levels of RTO/RPO based on business and regulatory requirements.
- **Maximum automation and maintainability:** Cloud-native synchronization mechanism and powerful built-in DevOps tools ensure consistency between DR and PROD environment.

Additional resilience on top of high-availability production set-up with two or three availability zones (AZ) as base configuration.

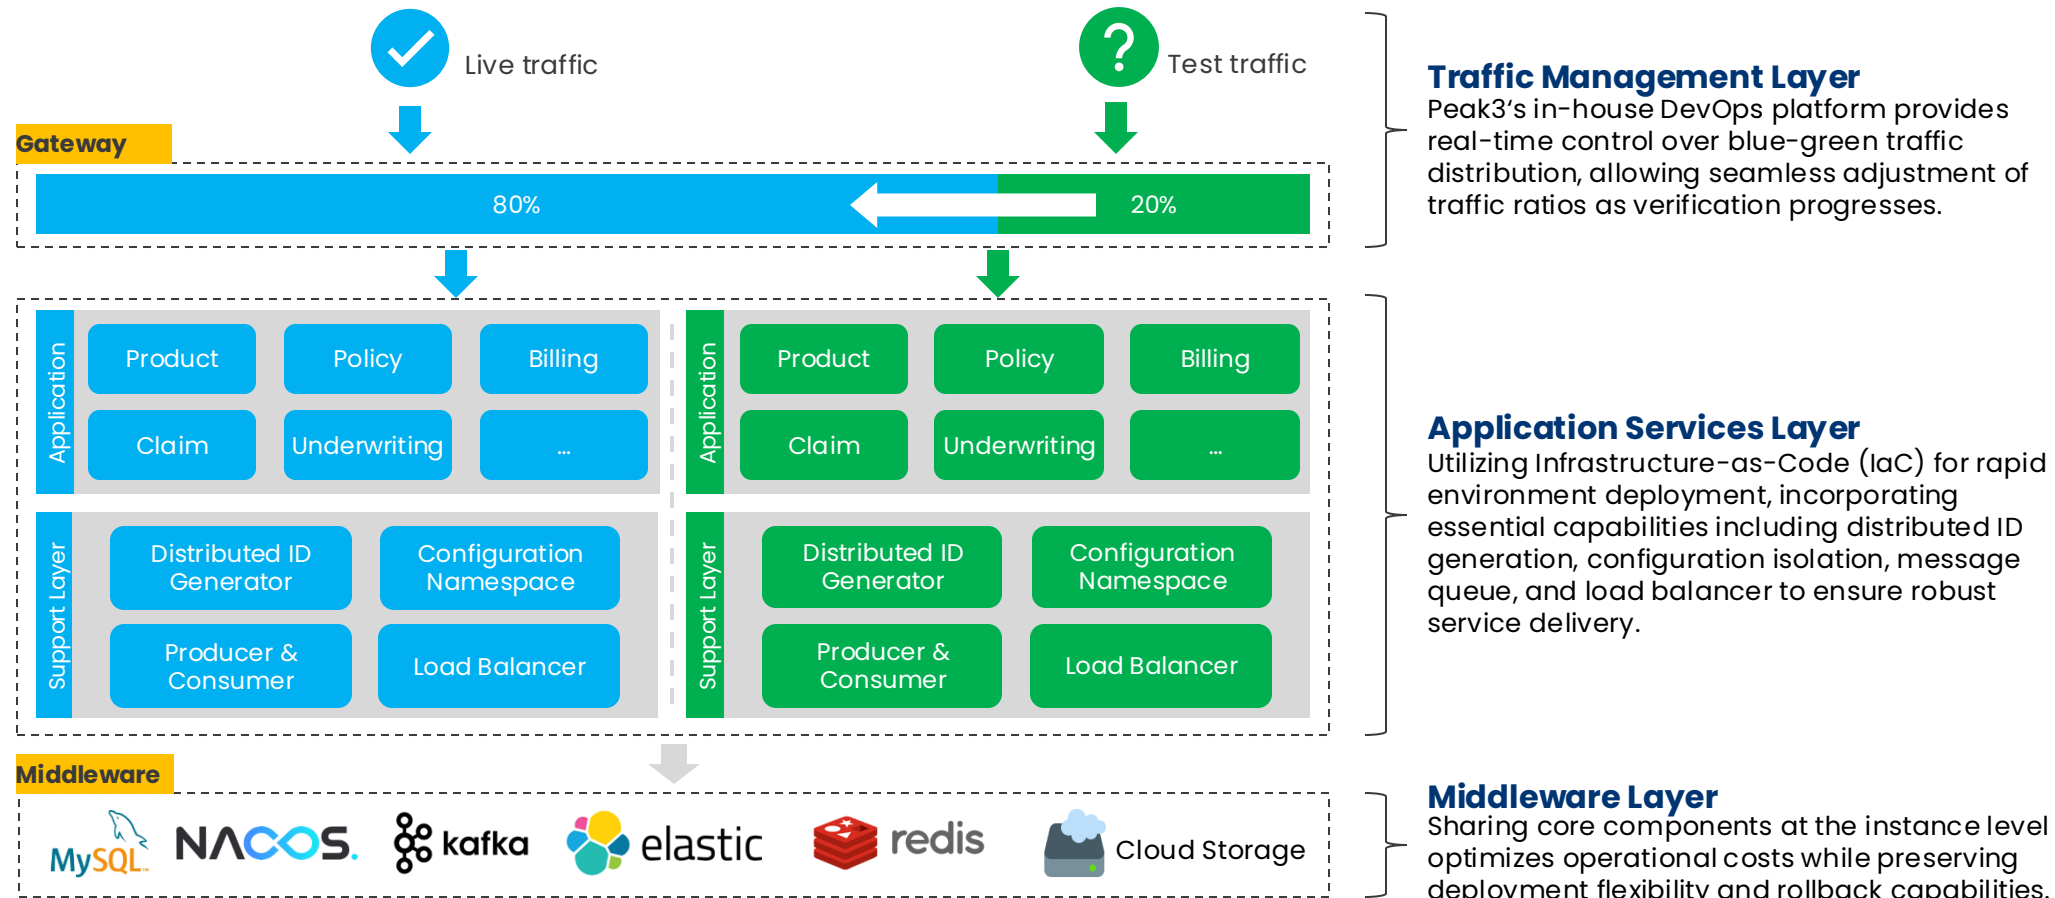
Availability in 2024 reached 99.998% across all clients, with more than half of clients achieving 100% availability (planned & unplanned downtime)



Application, middleware, DevOps tools, storage/DB deployed in different cloud regions and/or in different cloud platforms with near real-time database sync.



Blue-Green Deployment: Ensuring zero-downtime deployment



Traffic Management Layer


Peak3's in-house DevOps platform provides real-time control over blue-green traffic distribution, allowing seamless adjustment of traffic ratios as verification progresses.

Application Services Layer

Utilizing Infrastructure-as-Code (IaC) for rapid environment deployment, incorporating essential capabilities including distributed ID generation, configuration isolation, message queue, and load balancer to ensure robust service delivery.

Middleware Layer

Sharing core components at the instance level optimizes operational costs while preserving deployment flexibility and rollback capabilities.





Key Technology Updates

Tech Stack & Performance Upgrades

Tech stack upgrades: Lean SDLC capabilities ensure minimal accumulation of technical debt

Underlying capabilities

Open & modular architecture ensures that Peak3 can act on valuable updates or component replacements based on our tech radar scans.

Test automation throughout the closed-loop, combining continuous and shift-left testing with procedures and tools to cover unit testing, single API to business scenarios orchestration and UI testing.

DevOps excellence enables that all applications and related configurations, database scripts are distributed under a unified CI/CD framework, completing smooth deployment/traffic switching.



Recent major upgrades

OS: Migrated OS of VMs from CentOS 7 to new version of Ubuntu server for improved security & maintainability.



Cluster: Upgraded cloud-native Kubernetes cluster to new version on different cloud platforms (AWS EKS, Google Cloud GKE, Azure AKS, etc.).




Storage: Upgraded cloud-native DB to latest version 8.0.35, improving performance, security & maintainability.



Others: Upgraded different components, including Kafka, Nacos, SpringBoot, Docker to pro-actively address any risks.

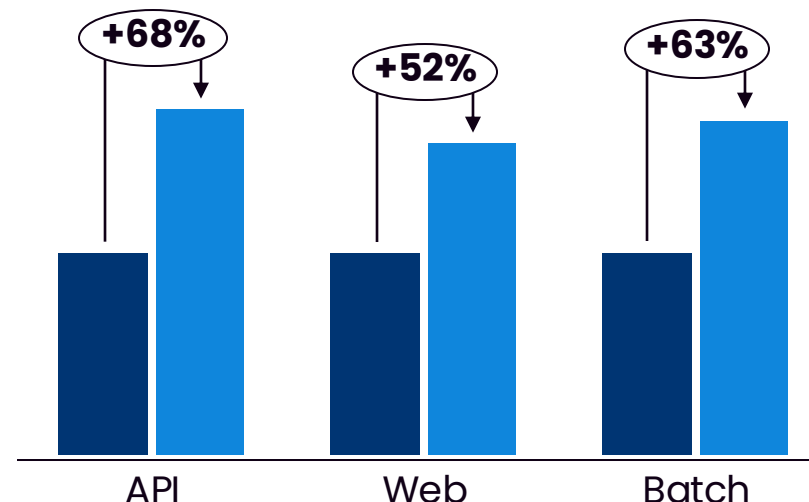


Performance enhancements: Full-stack DevPerfOps governance system boosting performance by 50%+



Advanced DevPerfOps governance:

- End-to-end management solution that integrates monitoring, test execution, reporting and optimization of performance in a mature DevOps framework.
- Seamlessly combination of automation and data-driven insights of performance improvements with product iterations.
- Cloud-native, microservices and distributed architecture provides the foundation for flexible performance benchmarks.



Significant performance improvements

Continued improvements for a platform designed from day 1 for high-frequency and high-volume business models

■ Graphene v2.50 ■ Graphene v3



Key Technology Updates

FinOps & Efficiency

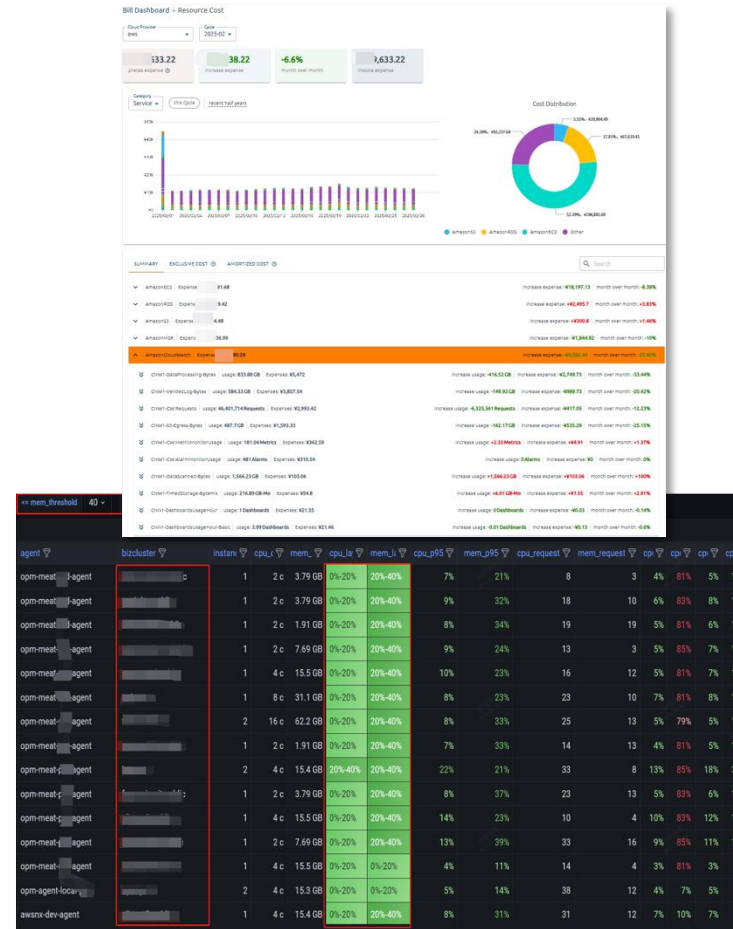


FinOps: Unified FinOps platform to optimize cloud usage

Upgraded FinOps platform

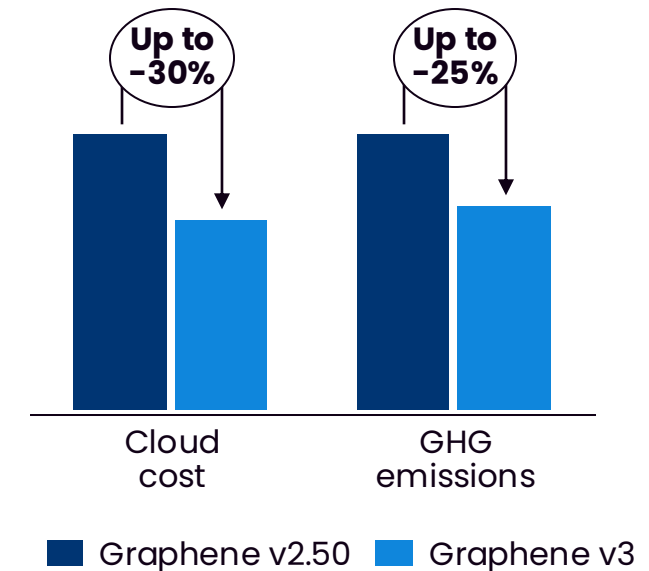
Implemented new comprehensive tagging and grouping rules for all resources (including mapping to specific business), providing cloud vendor-agnostic cost visualization dashboards, cost comparison reports and automated anomaly detection and alarms.

Works with capacity analysis dashboard to provide unrivalled visibility and data-driven decision-making power.

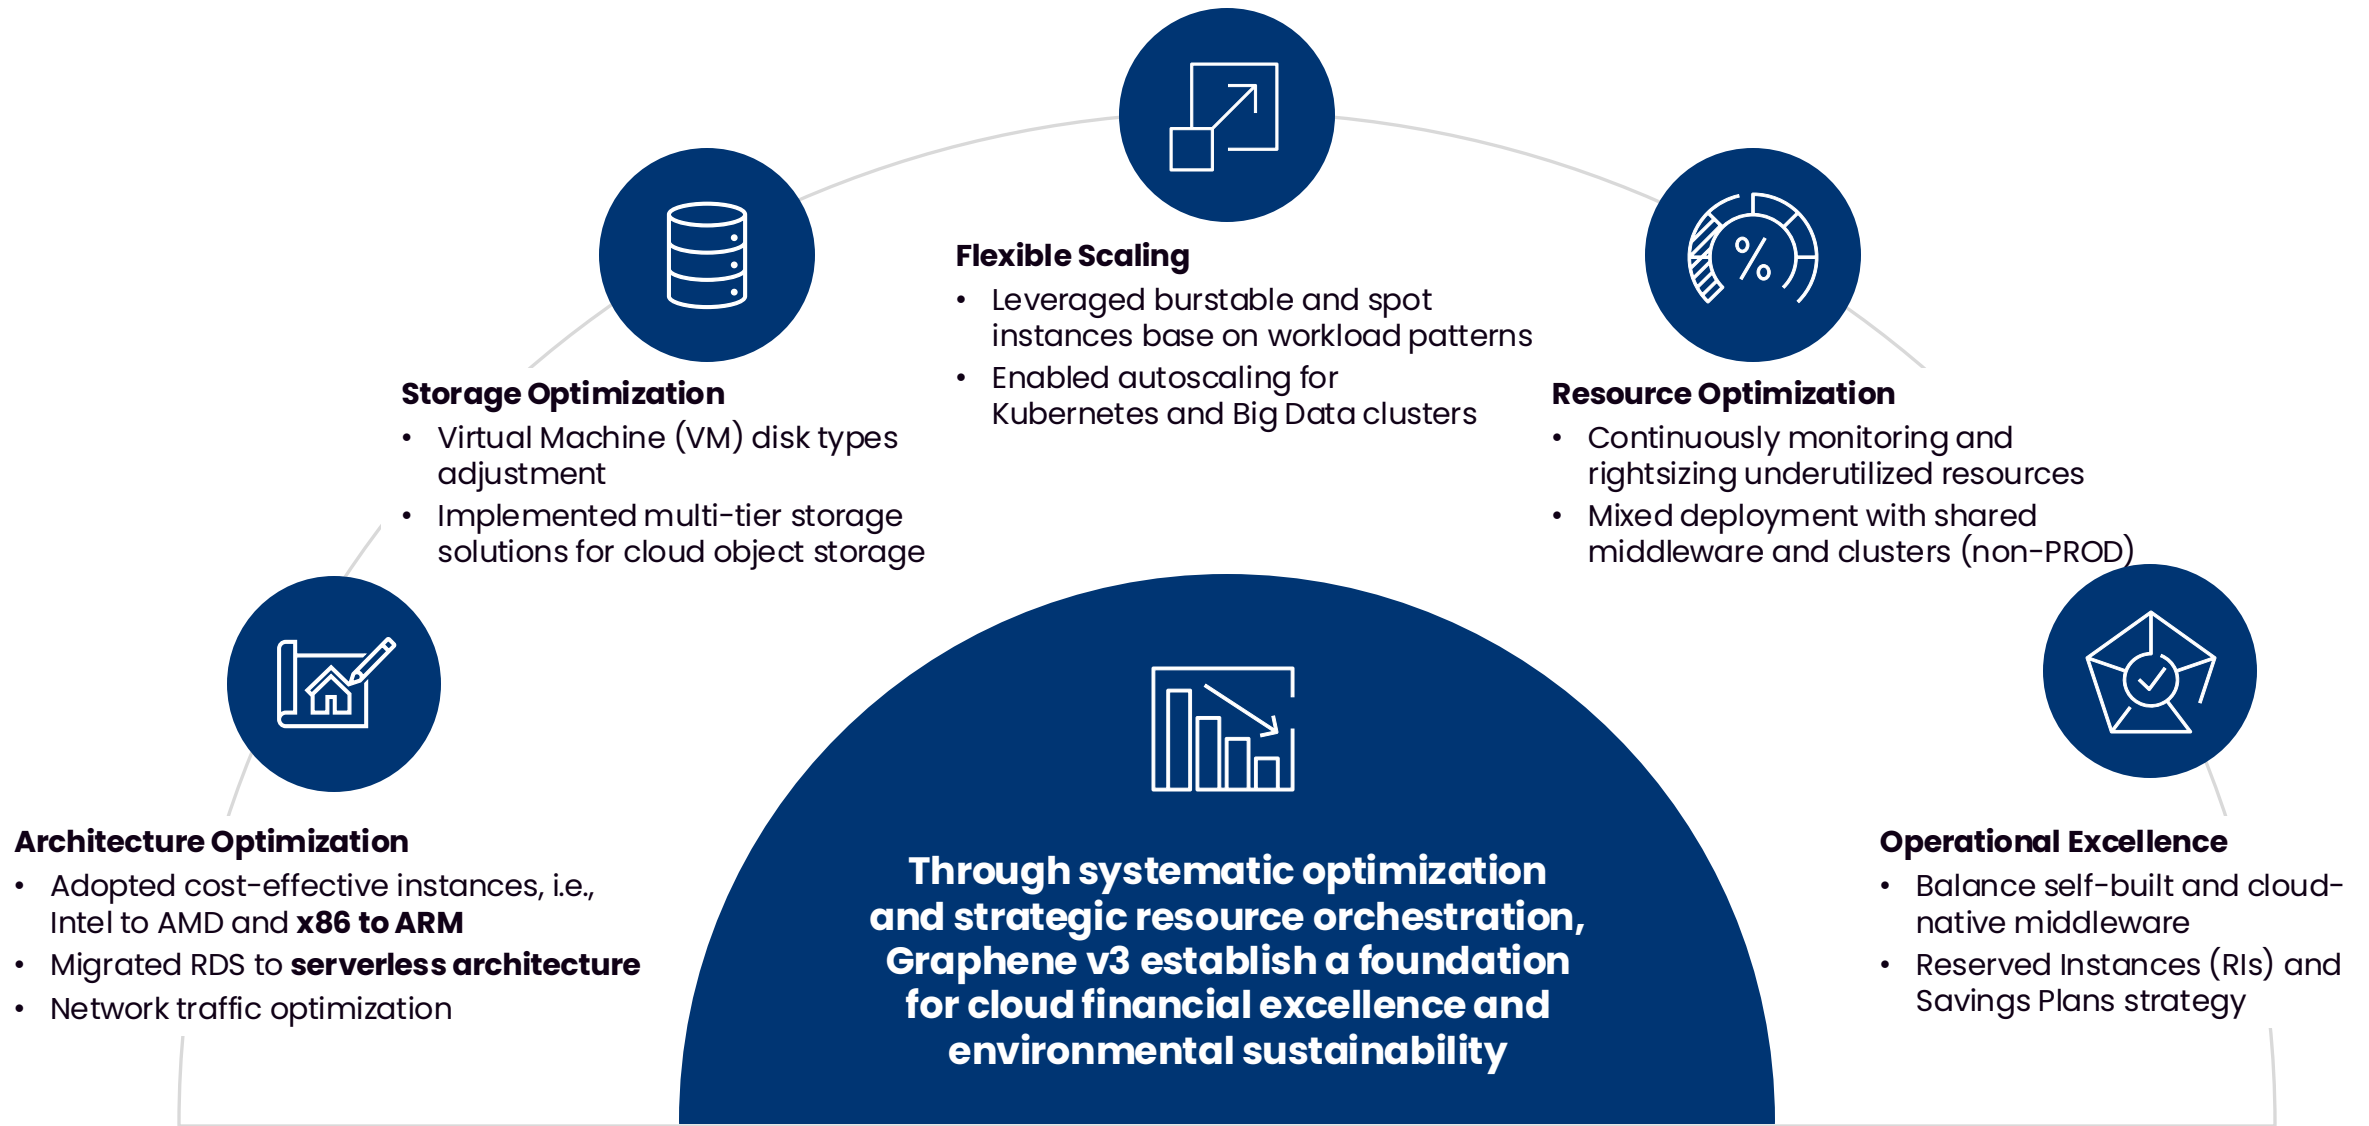


Significant benefits

Reduction in cloud costs and greenhouse gas emissions through recent FinOps initiatives:





FinOps: From visibility to value creation in Graphene v3



Scalability: Database serverless scaling for instant and in-place database scaling in Graphene v3

Database Serverless Scaling



Rapid in-place capacity expansion to cope with traffic peak: Automatic elastic capacity expansion of the database is completed within 2-3 seconds, achieving accurate dynamic matching of business traffic and database resource capacity.

- No business interruption due to scaling even when running hundreds of thousands of transactions
- Solving unavailability issues that may be caused by the performance bottleneck of traditional cloud-native databases, greatly improving the stability and enhancing the availability to more than 99.99%
- Simplifying capacity planning and reducing maintenance costs with automatic fine-tuned management
- Significantly reducing the costs of cloud resources by avoiding unused compute and storage resources, saving 20% compared to traditional cloud-native databases



Key Technology Updates

DORA Compliance

DORA compliance: Staying compliant with EU's new Digital Operational Resilience Act with Graphene

What is DORA?

The Digital Operational Resilience Act (DORA) is a regulation introduced by the European Union to strengthen the digital resilience of financial entities (including insurers), effective 17 Jan 2025.

It aims to ensure that insurance companies can withstand, respond to, and recover from ICT (Information and Communication Technology) disruptions, such as cyberattacks or system failures.

DORA brings harmonization to rules relating to operational resilience, applicable to financial entities and ICT third-party service providers.



How does Peak3 and Graphene ensure compliance?

Most importantly, Peak3 and Graphene are ready for DORA compliance. Graphene provides a secure, high-performing, resilient, and efficient core infrastructure. Effective and timely communication, cooperation and coordination are in place to meet DORA requirements.

Peak3's *Resiliency Plan* is tested, reviewed, and approved by senior leadership annually. Peak3 can support each client to meet *Threat-Led Penetration Testing* (TLPT) requirements that apply under DORA. Peak3 will support its clients to issue an *ICT Incident Report* (where needed).

Peak3 welcomes the concept of *Annual Oversight Plans* for Critical Third-Party Providers (CTPPs) and understands the importance of identifying relevant risks in a CTPP's Annual Oversight Plan so that risk-based oversight of CTPPs can be effectively carried out.




Delivery

SI Delivery Enablement



Graphene v3 is the first version available for implementation by certified implementation partners



In-house to Partner Delivery

Peak3 is strategically transitioning from an in-house delivery and direct client enablement model to a system integrator (SI) partner delivery model.

Going forward, clients have the option to engage third-party SI partners for implementation and support services. This will boost Peak3 delivery scalability and provide enhanced local language and in-country support for clients across all regions.



Certification & Onboarding

Peak3 provides training and certification services for ambitious partners that want to bring the next generation of insurance core platforms to clients.

FPT IS, a leading Vietnamese  technology partner, has been onboarded and trained as the first partner. Additional partners are currently in progress.



Interested in becoming a certified implementation partner for Graphene v3?

Please reach out to:

- EMEA/Americas: roeland.with@peak3.com
- APAC: jennie.zhang@peak3.com



Disclaimer & Safe Harbor Statement

This document is intended to outline Peak3's general product directions and features. Its contents are intended for information and discussion purposes only. It may not, and shall not, constitute a binding agreement or contractual obligations of any sort.

This document is not a commitment to deliver any software or functionality. The development, release and timing of any features and/or functionality described herein remain at the sole discretion of Peak3.

These materials serve only as the focus for discussion; they are incomplete without the accompanying oral commentary and may not be relied on as a standalone document. Further, Peak3 has made no undertaking to update these materials after the date hereof, notwithstanding that such information may become outdated or inaccurate.



www.peak3.com