Cloud architecture for enterprise-grade applications

Introduction: Continually evolving to meet enterprise needs

Pega is powering mission-critical processes at the world's leading organizations – from collecting data from [300+ million citizens](#) to delivering real-time next best actions for [70+ million customers](#) to advancing [R&D for life-saving drugs](#).

The applications that do this work require enterprise-grade performance, scale, flexibility, and resilience. The technologies and architectures that deliver on these requirements are continually advancing, and so is the architecture of [Pega Infinity™](#) the low-code platform for AI-powered decisioning and workflow automation that enterprises rely on to personalize customer engagement, accelerate acquisition and onboarding, automate customer service, streamline operations, and resolve exceptions.

For 35+ years, Pega has incorporated new technologies and approaches into the architecture of our low-code platform. Core languages have evolved from CICS to PL/1 to C++ to Java, JavaScript, and Kotlin, while typical infrastructure has moved from mainframes to bare metal servers to public clouds – leveraging containerization.

In 2018, Pega announced Project fnx – our program to evolve Pega's underlying architecture, focused on:

- Adopting a **microservices architecture** backed by leading-edge cloud technologies for increased agility, innovation, and resiliency

- Standardizing on **cloud-native technologies,** such as Kubernetes and Docker, for deployment repeatability and elastic auto-scaling

- Deploying a **new, modern front-end architecture** for applications, which incorporates technologies like React and web components for responsive end user experiences

- Introducing new products, as exemplified by the recent announcement of **Pega Launchpad**

This paper details Pega Infinity's cloud-native architecture – enabled by Project fnx – and illustrates how that architecture delivers on the unique performance, scale, flexibility, and resilience needs of the world's leading organizations.

Please note: While many of these of these architectures are consistent across Pega Cloud® and self-managed "client cloud" environments, the focus of this paper is on Pega Cloud – which both increases project ROI and saves costs for Pega projects.
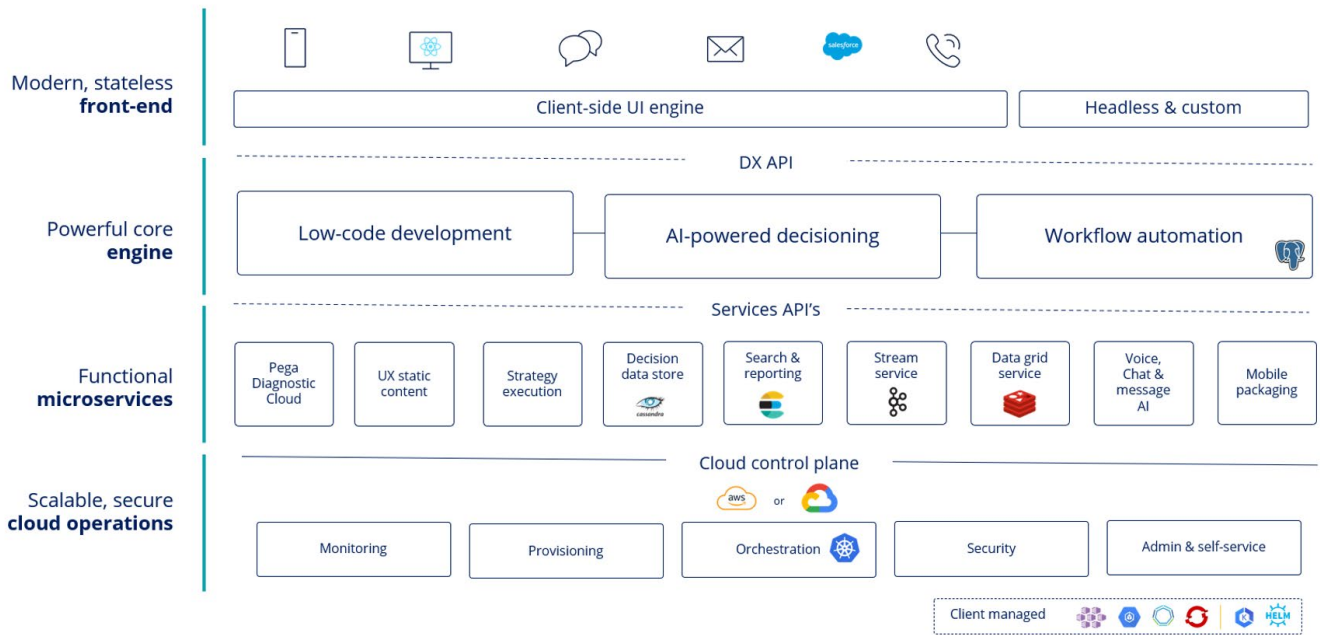
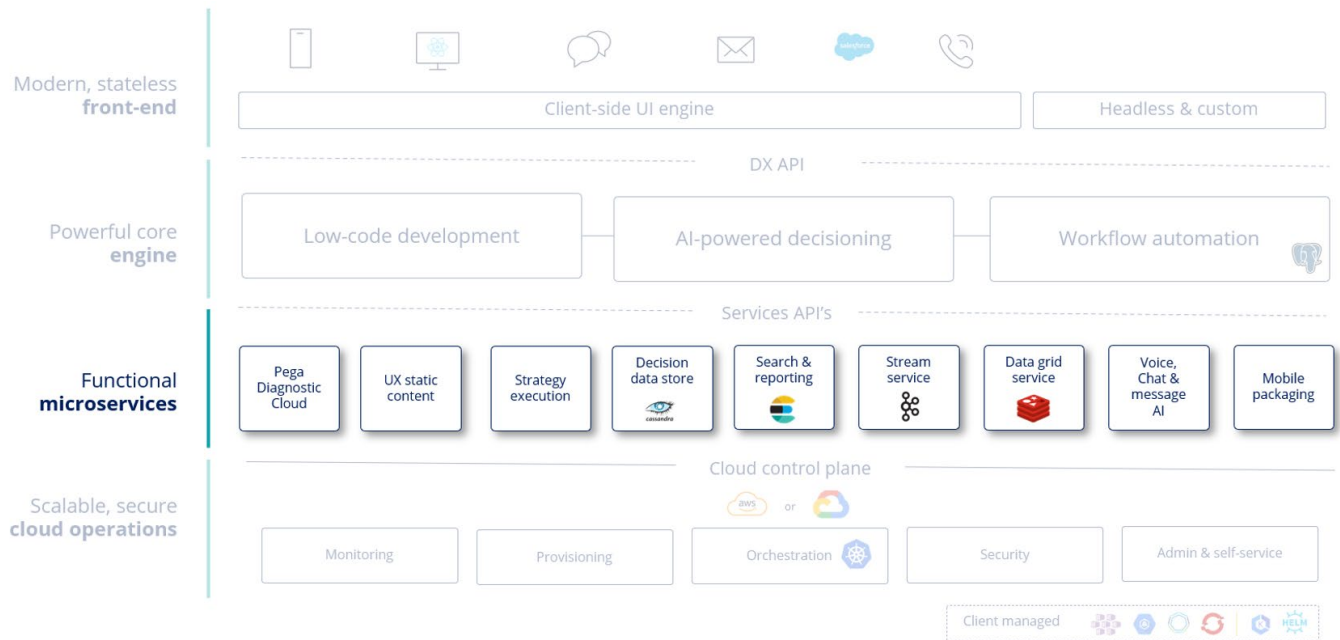High level depiction of the Pega Platform's cloud-native architecture

"Pega's cloud-native, multitenant, microservices-based architecture can scale up and scale down, helping clients address use cases more cost-effectively."

Microservices for scale & resiliency

Pega's standard deployment includes core workflow and AI engines – securely provisioned per account – along with a number of microservices that power supporting platform functionality. Pega chose a microservices architecture because it delivers on the uptime, scale, and resiliency needs of leading enterprises.
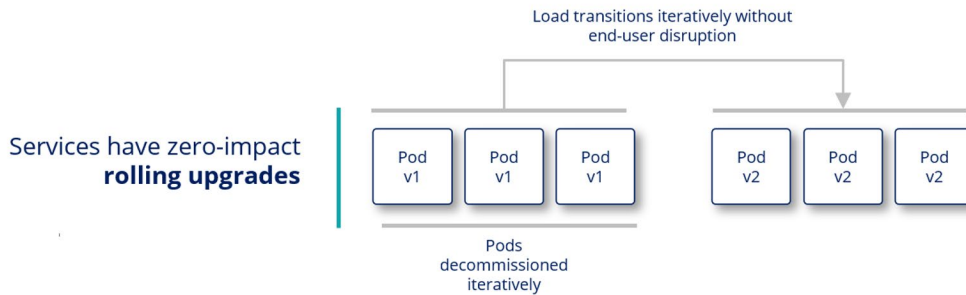

High level, non-exhaustive depiction of the Pega Platform's cloud-native architecture focusing on functional microservices

Pega's microservices have a few shared characteristics:

- **They are deployed independently:** Run, get updated, and scale separately from other services and the Platform at large.

- **They have well-defined APIs:** Clear inputs and outputs get invoked by the Platform behind the scenes to power capabilities.

- **They typically include a processing and technology layer:** The technology layer is typically storage provided by the cloud infrastructure provider (e.g. AWS Elasticsearch for Search & Reporting capabilities). The processing layer is comprised of a horizontally scalable service-broker, which abstracts the technology layer from the platform-at-large.

Seamless, rolling global updates = faster features, security patches, and fixes.

Users expect applications to be available all the time, and developers are expected to deploy new versions of them several times a day. Pega standardized on Kubernetes, which enables rolling updates. Rolling updates allow deployment updates to take place with zero downtime – all by incrementally updating pods (deployable units of Docker containers) instances with new ones. System load is slowly transitioned to the new pods, while the old pods are decommissioned.

Load transitions iteratively without
end-user disruption

Services have zero-impact
rolling upgrades

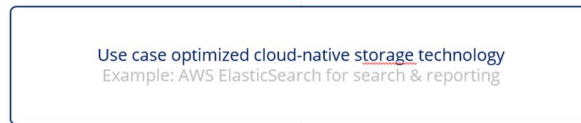| Pod v1 | Pod v1 | Pod v1 | | Pod v2 | Pod v2 | Pod v2 |

Pods
decommissioned
iteratively

Services provide consistent performance under load through horizontal and vertical autoscaling.

Pega's cloud control plane monitors all service utilization – and automatically provisions and routes load to service pods as processing volume increases. For additional storage, the backing technology volume is monitored – and scaled up and down to meet usage. Leveraging these approaches, Pega Cloud® is proven to power the customer service operations of more than 20,000 representatives and millions of customers with no degradation in performance.

Service layer scales
out to meet
processing demand

| Pod | Pod | Pod | Pod | Pod | Pod | →

Backing technology
scales up to meet
Storage demand

Use case optimized cloud-native storage technology
Example: AWS ElasticSearch for search & reporting

Seamless response to disruption = higher availability.

Similar to rolling updates, Kubernetes also provides the ability to monitor and replace unresponsive or corrupt pods seamlessly. When a node is corrupted or goes down, it will automatically be identified and replaced – zero disruption to service availability.

The results?

Hundreds of seamless updates per week to Pega's services – meaning new features, fixes, and security remedies more often. Since adopting microservices, Pega has seen a 10X decrease in the rate of client-reported issues. And when an issue does get reported and fixed, it gets fixed for every client instantly (no patches, hotfixes, or upgrades needed).

"Increasingly, cloud is the means with which they pursue digital transformation both quickly and while maintaining their customer-obsessed strategies. The same rule applies to business processes powered by Pega applications **to empower employees and to enable business transformation, organizations are operating in Pega Cloud."**

Forrester Research, The Total Economic Impact™ Of Pega Cloud, May 2021

Cloud-native technologies for flexibility and self-service

Pega is committed to adopting cloud engineering best practices and cloud-native technologies throughout the stack.
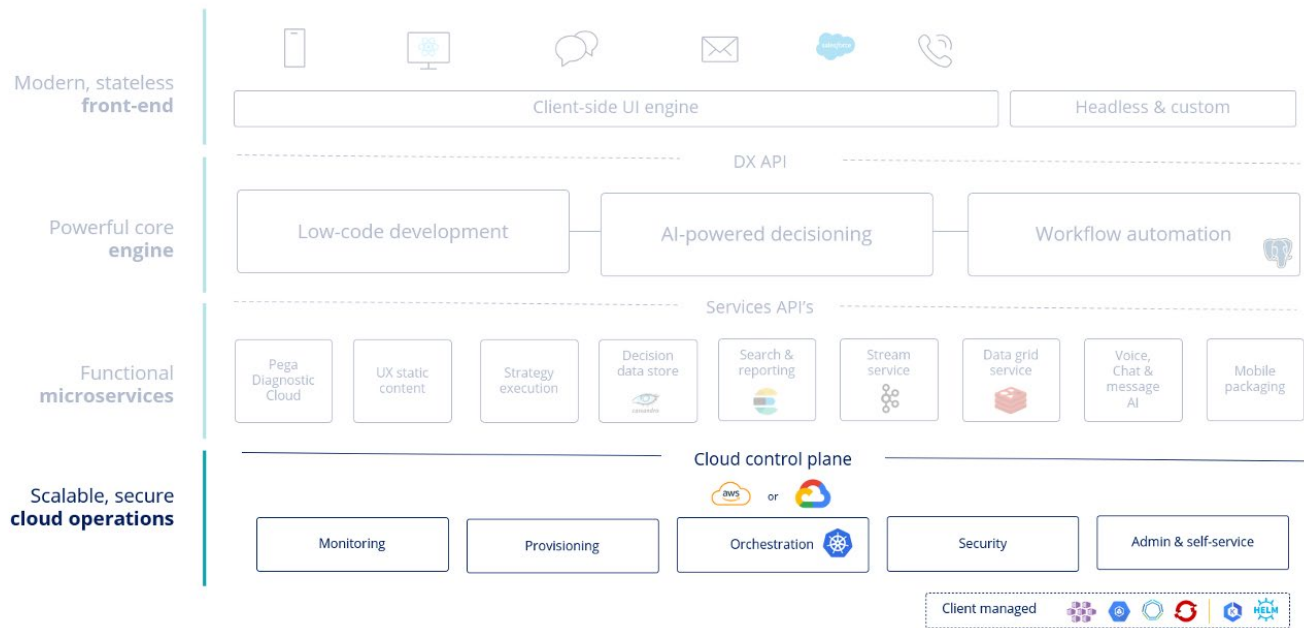
High level depiction of the Pega Platform's cloud-native architecture, focusing on the cloud operations control plane

Standard deployment pattern = more flexibility and streamlined troubleshooting.

Docker and Kubernetes is the recommended deployment approach across both Pega Cloud and client-managed environments. Pega supports and provides Helm charts for commonly managed Kubernetes providers (e.g. AWS EKS, Google Cloud GKE, and more). As discussed above, this provides a number of runtime benefits – in scale, availability, and agility. Beyond that, adopting a standard deployment pattern enabled Pega to seamlessly expand Pega Cloud managed services to the Google Cloud Platform (where previously AWS was the sole option). In addition, standard deployment enables IT admins and support to accelerate troubleshooting by eliminating infrastructure as a factory.

More connected and automated = more self-service.

Adoption of cloud operations technologies as part of the Pega Cloud control plane – such as AWS Cloud Watch and New Relic – enables additional insights, controls, and automations which can be offered in a self-service manner to administrators of Pega environments.

For example, Pega Predictive Diagnostic Cloud™ is a central dashboard for IT system admins. Predictive Diagnostic Cloud monitors real-time logs from Pega applications and analyzes log data to provide actionable insights and alerts to users on the performance, health, and usage of their active applications.
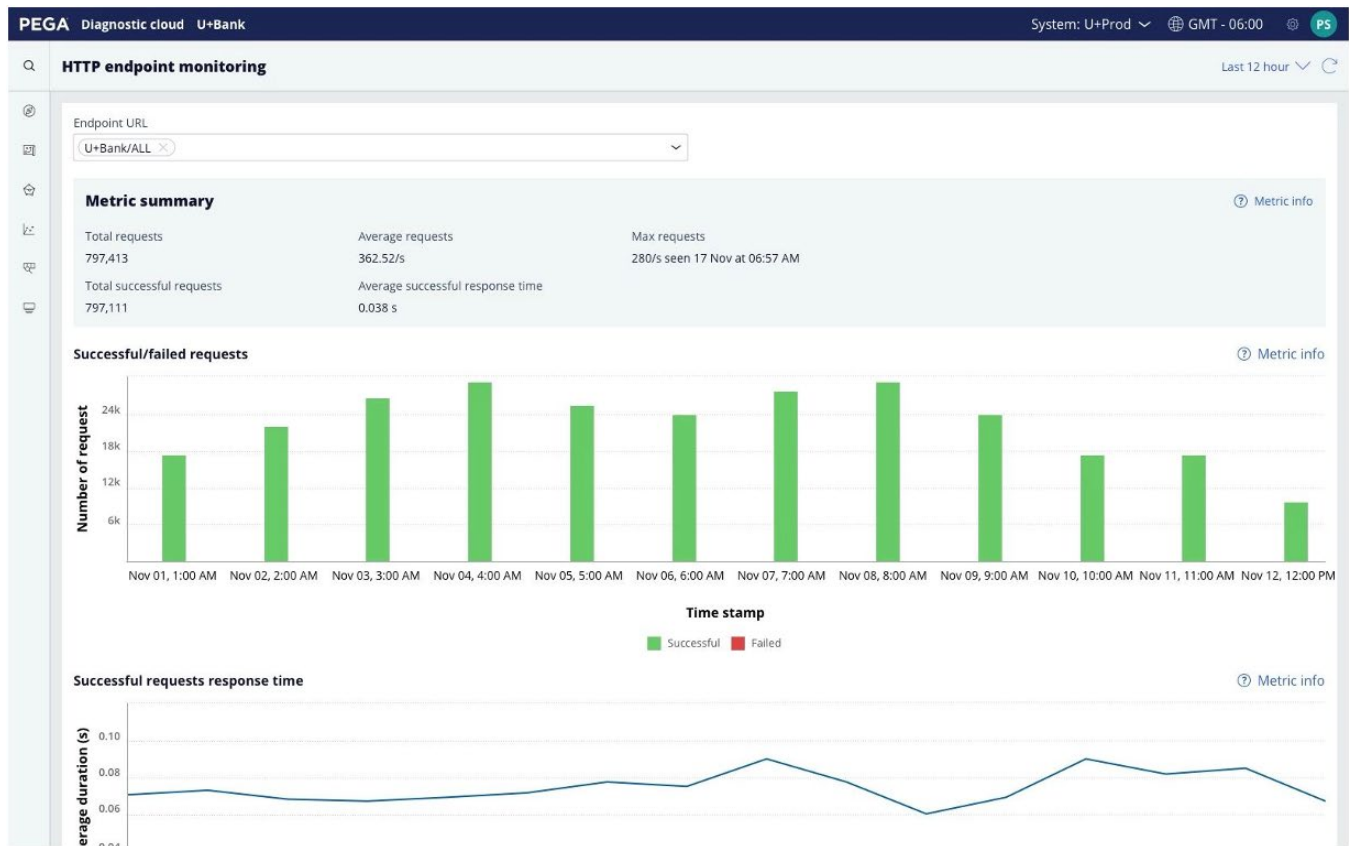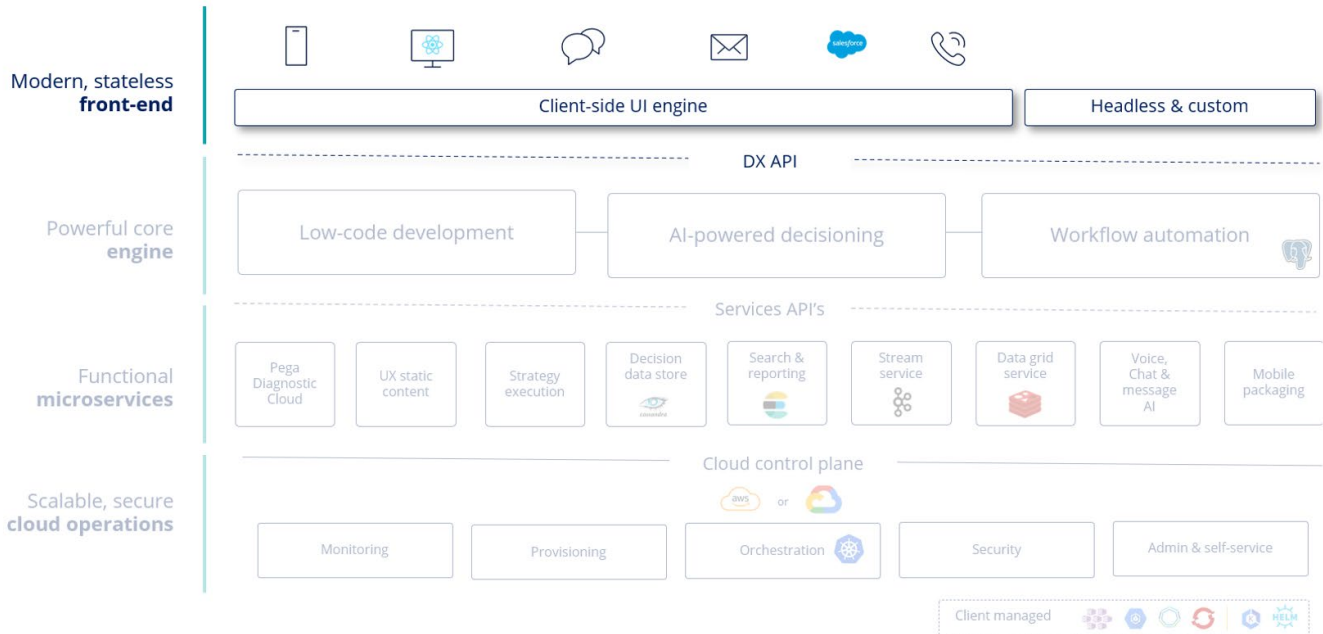

Example of monitoring insights provided by Pega Diagnostic Cloud

Modern front-end architecture for responsiveness and openness
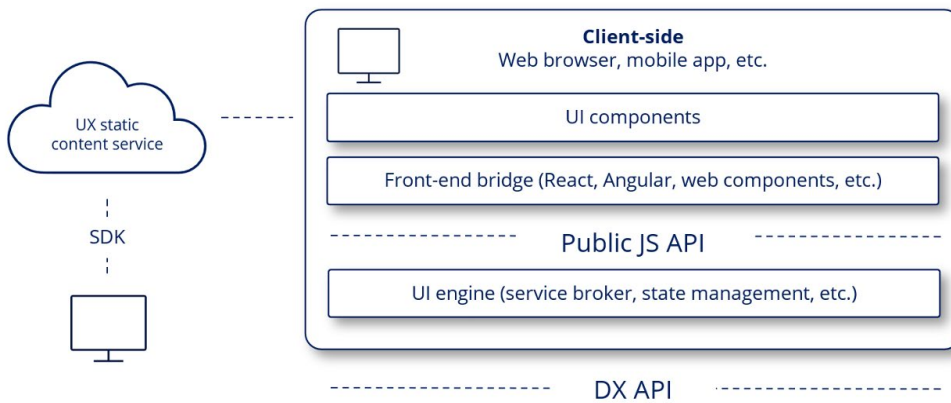
In addition to an evolved cloud architecture, Project fnx focuses on rethinking Pega's approach to the front end.

High level depiction of the Pega Platform's cloud-native architecture focused on modern front-end architecture

Core to this architecture are:

- **Digital Experience (DX) API**: Automatically-generated, model-driven REST API endpoints enable a programmatic read, create, and update cases, data, and assignments. Unlike traditional process APIs, DX API provides UI metadata in the response – so consistent front-end experiences can be interpreted in any channel.

- **UX static content service:** Serve a library of UX static content (React components, JavaScript code, web components, etc.) to the browser. By default, the UX static content service contains all React components needed to power the Pega Cosmos Design System. Pega also makes a software development kits (SDK) available, enabling clients to develop and publish React components that extend Pega Cosmos.

- **Client-side components:** A new client-side engine called Constellation acts as a service broker between UI components and Pega's core services through DX API. It manages API calls, state, caching, and encryption – and abstracts it to a public JavaScript API that powers the UI componentry. Constellation drives improved user experiences via the Pega Cosmos Design System, but also provides a framework for clients to build their own design systems directly into the low-code workflow applications they build in Pega.

Open architecture = endless flexibility.

This approach unlocks new options for seamlessly embedding Pega into an enterprise's UI landscape – taking into account existing systems, portals, design systems, and standards. Developers who use their favorite front-end frameworks typically need to build their own components to connect them into back-end systems and maintain the website's distinct look and feel. Pega provides an orchestration layer with pre-built components to accelerate the integration of user experiences developed on popular front-end frameworks, including new and updated SDKs for Angular, React, and web components available via Pega Marketplace.

Optimized design system = increased productivity.

In addition to evolved architecture, the Pega application look and feel has advanced to adopt a new design system: Pega Cosmos. Pega Cosmos is purpose-built for workflow automation applications; application details, actions, and utilities are consistently and prescriptively presented to users. Pega Cosmos is shown to increase end-user productivity while accelerating developer productivity.
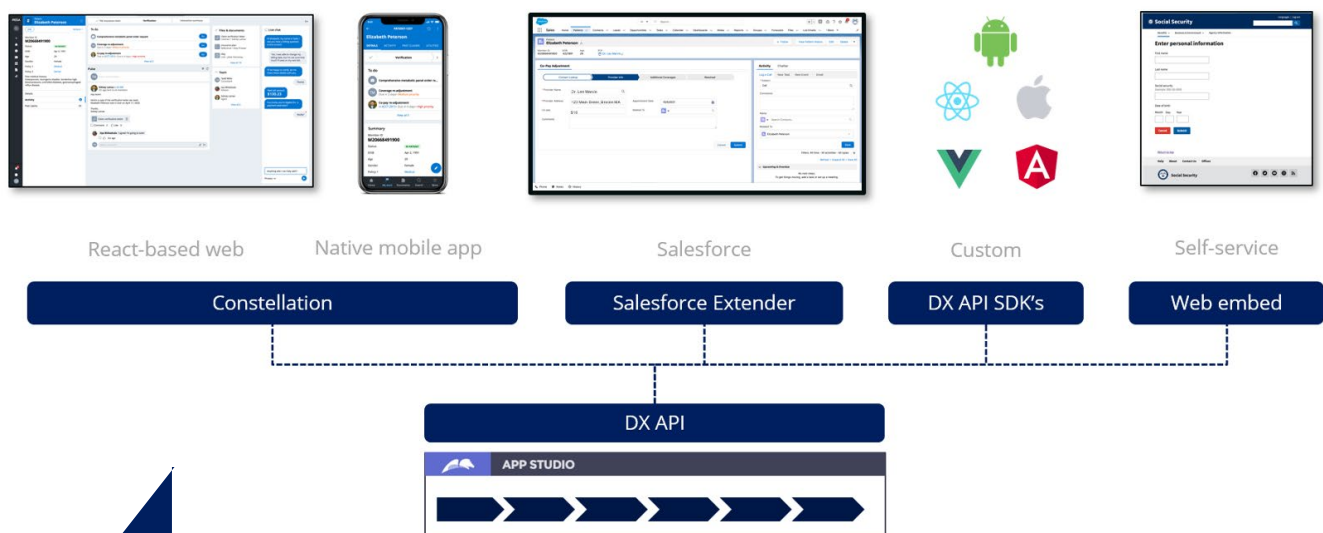
A view of the Pega Cosmos Design System
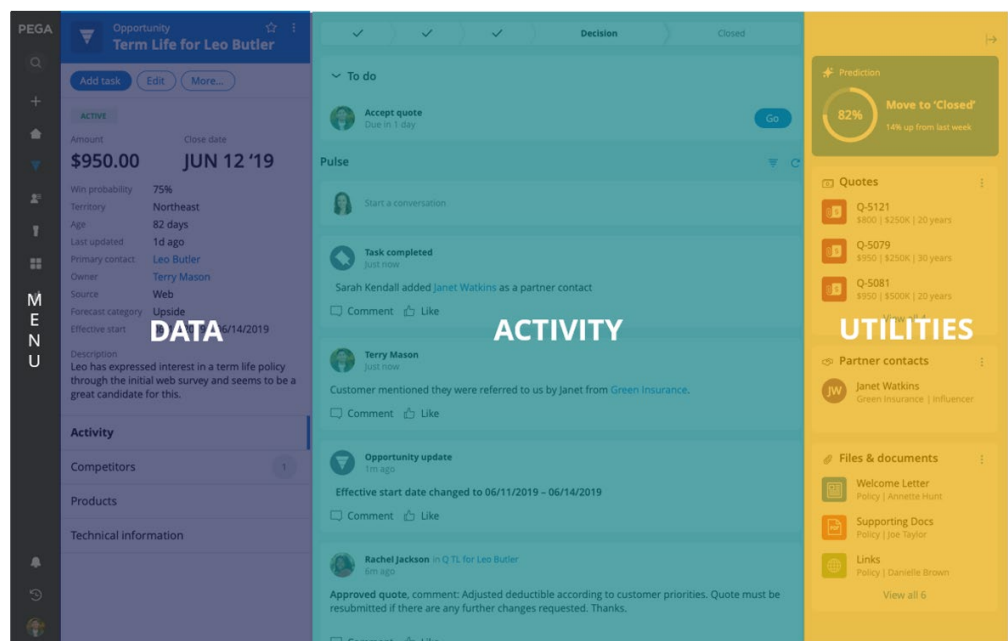
Optimized architecture = optimized performance.

The introduction of a client-side UI engine enables smaller, less frequent payloads between client and server. As a result, the overall experience is available to users 3X faster and the overall network payload is 7X smaller when compared to server-generated UI.

Center-out workflows = consistency across any channel.

Because all channels are powered by a unified API – web, mobile, chat, self-service, headless, and even Salesforce Lightning – all end-user experiences are ensured to be consistent. In addition, this approach both accelerates development, as workflows can quickly be deployed to new channels, and system maintenance costs, as any changes made to workflows are instantly reflected across all channels.

React-based web   Native mobile app   Salesforce   Custom   Self-service

| Constellation | Salesforce Extender | DX API SDK's | Web embed |

DX API

APP STUDIO

What's available?

The DX API, front-end technology SDKs and starter kits, and Pega Process Extender for Salesforce Lightning are all available now. As of Pega 8.7, Cosmos React applications (which include the new client-side engine) are available for evaluation for new build-from-scratch Pega projects.

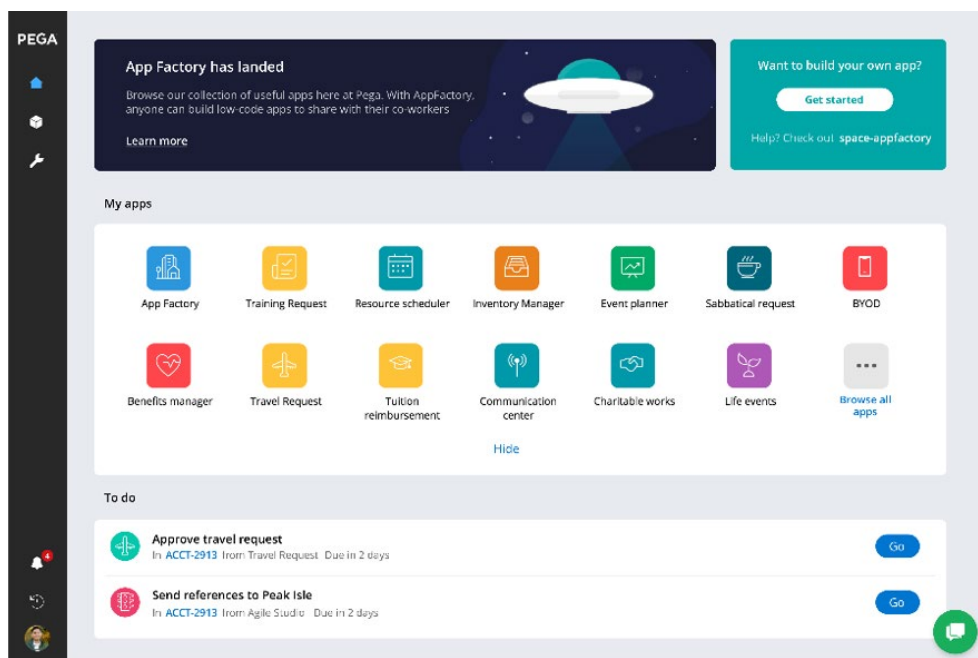# End-to-end visibility and orchestration with Process Fabric

Gartner's vision of a [composable business](#) recommends enterprises move from large, horizontal functions to develop modular, verticalized business capabilities – leading to:

- More speed through discovery
- Greater agility through modularity
- Better leadership through orchestration
- Resilience through autonomy

Through new architectural approaches enabled by Project fnx, Pega is uniquely set up to enable enterprises to deliver on this vision. Pega Infinity's cloud architecture allows organizations to deploy multiple, smaller applications that each contain a small number of case types or workflows. This allows for easier deployment of changes and isolation of issues. And it's increasingly necessary in a global, distributed world where data residency requirements enforce keeping certain application data within jurisdictional boards. But as your applications become more distributed, Pega has delivered capabilities to ensure you can manage both the distributed development of low-code workflow applications and weave distributed workflows together into holistic experiences for employees and customers.

Automated SDLC = streamlined development of business capabilities.

Pega App Factory allows organizations to effectively scale business capability application development by enabling IT automation of developer enablement, coaching, reuse, and governance. Pega App Factory is a Pega application that gives enterprise IT a framework to configure the standard software development lifecycle of low-code applications. This includes the ability to discover and set up reusable components for new projects, templatize and subsequently set up DevOps pipelines (automated test execution, security validations, approvals, and more), and enable developer collaboration, enablement, and coaching.

Pega App Factory dashboard

## DX API = unified view across applications.

As applications are deployed, Process Fabric uses the DX API to access workflow data across workflows. This provides all users with a single, unified experience ¬– personalized to them across all apps. For example, the Process Fabric Hub provides an open API that listens for new and updated assignments across applications (Pega and third party). This saves employees time swapping between multiple systems to get their work done – all by providing them with a unified worklist, ability to run cross-application reports, and a single view of all the distributed work they are stakeholders of.
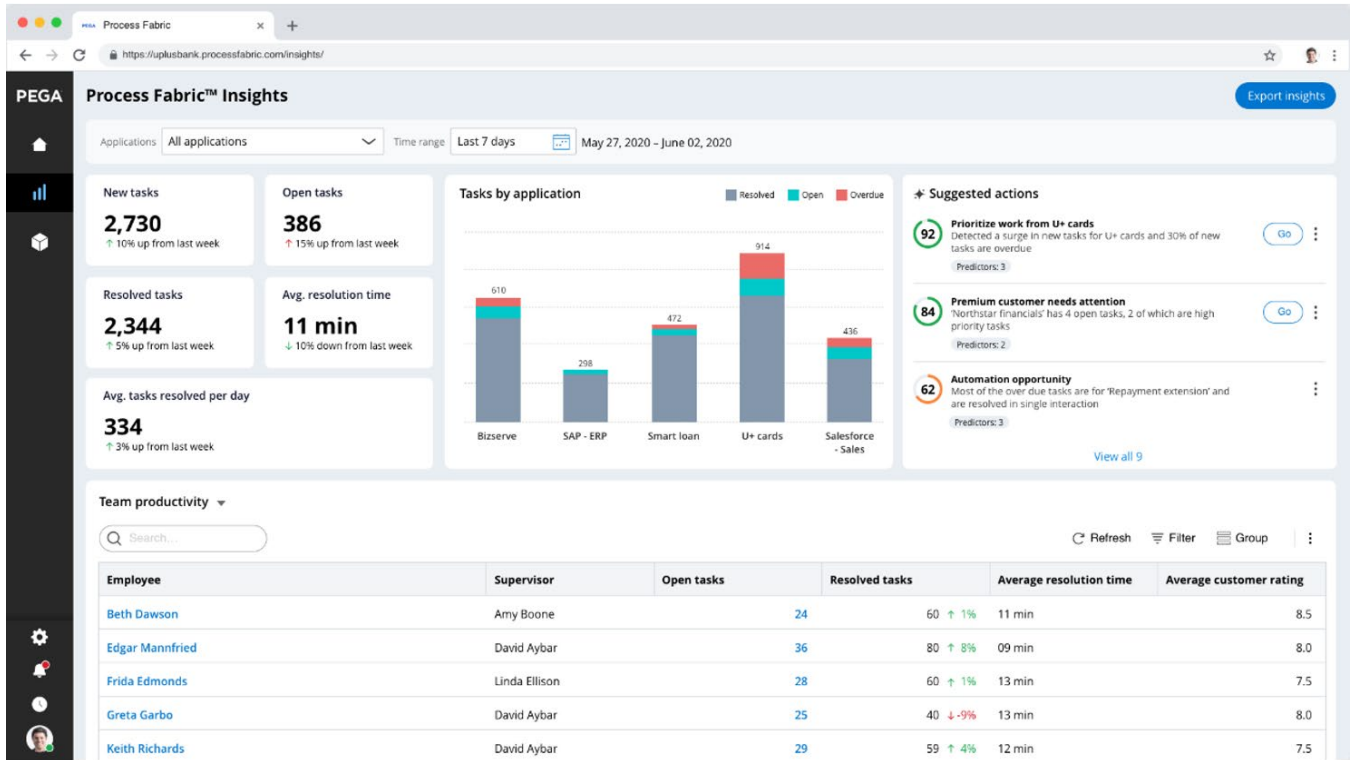
Pega Process Fabric reporting dashboard
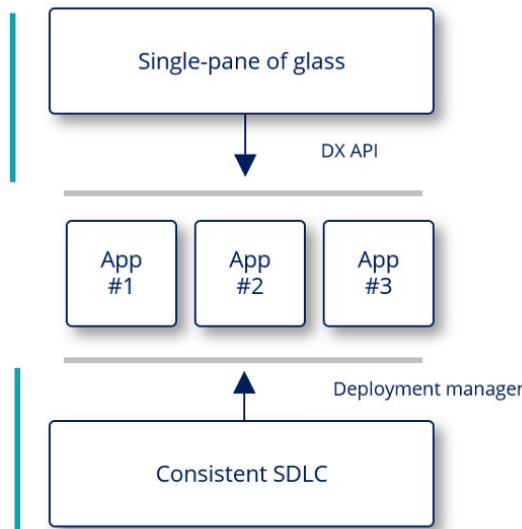
Depiction of the relationship between Pega App Factory and Pega Process Fabric

Conclusion: The evolution never stops

So long as the expectations for enterprise-grade performance, scale, flexibility, and resilience advance, so will Project fnx and Pega's investment in the adoption of new technologies and architectures.

To date, through Project fnx, Pega's cloud-native architecture has evolved to:

- Adopt a **microservices architecture** backed by leading-edge cloud technologies for increased agility, innovation, and resiliency

- Standardize on **cloud-native technologies**, including Kubernetes and Docker, for deployment repeatability and elastic auto-scaling

- Leverage a new, **modern front-end architecture** that incorporates technologies like React and web components for flexible, consistent, responsive end-user experiences

- Enable enterprises to deliver visibility and orchestration for end-to-end processes that cut across distributed workflow systems with **Pega Process Fabric®**

It all adds up to one thing: Powering mission-critical applications and operations for the world's leading enterprises – today and into the future.

About Pegasystems

Pega delivers innovative software that crushes business complexity. From maximizing customer lifetime value to streamlining service to boosting efficiency, we help the world's leading brands solve problems fast and transform for tomorrow. Pega clients make better decisions and get work done with real-time AI and intelligent automation. And, since 1983, we've built our scalable architecture and low-code platform to stay ahead of rapid change. Our solutions save people time, so our clients' employees and customers can get back to what matters most.

For more information, please visit us at **pega.com**

PEGA®