



DATASHEET

The Cloud-Native Enterprise Platform Purpose-Built for Scientific R&D

R&D is the driving force behind science-driven companies, and to stay ahead in today's competitive landscape, R&D teams need to accelerate their work. Fortunately, digital technologies like cloud computing, artificial intelligence, and data analytics provide an unprecedented advantage. But unlocking their full potential can be challenging for most R&D teams. That is where Enthought Edge comes in. Our purpose-built platform for scientific R&D enables you to easily access your complex scientific data, collaborate effectively on AI projects, and deploy your applications to users quickly and securely. All while keeping your IP secure and costs low. With Enthought Edge, you can focus on what you do best - solving complex scientific challenges - while we handle the rest.

Enthought Edge Powers Science at Every Level

Developed by scientists for scientists and backed by 20 years of leadership in advancing R&D innovation. Scientific data is uniquely complex, multimodal, and scattered across locations and formats. Researchers lose valuable time and productivity when they cannot quickly find data that is produced in the laboratories or access data that is managed in databases. They often face technical barriers when they need to share their work with collaborators or are burdened with setting up and maintaining their computational environments. Edge eliminates the friction and bridges the gap between what scientists need to be successful and what is available to them in enterprise settings. Edge powers science at every level - a single user, a lab or department, or across the enterprise. It can stand alone or be the R&D component of a larger corporate IT strategy.

Unlock Your Scientific Potential

Capabilities that accelerate discovery directly in the hands of R&D.

- Data App serves as a no-code interface that enables users to centralize, browse, and share data as well as connect to remote data sources using Data Connectors.
- Analysis App provides an enhanced JupyterLab with live connection to data, Enthought's enterprise-grade Python distribution with more than 100 million downloads, and the ability to share assets with a simple drag-and-drop.
- Compute Profiles enable users to access tiered computational resources (CPUs, GPUs, and RAM) in the cloud with one click.
- The Native Application Framework enables R&D teams to put digital solutions like dashboards, ML models, and web applications into production with no IT or DevOps learning curve.
- Cloud Watcher enables R&D organizations to selfmonitor their cloud spend and manage their budget proactively and as they wish.
- Time Machine eliminates the downtime risks involved in update rollouts and lets users always access all versions of their applications on Edge.

The Edge Scientific Workbench

One intuitive interface for R&D teams to access data, develop notebooks, and publish applications.

The Edge Scientific Workbench enables scientists and engineers to access data wherever it is stored, solve R&D problems in computational notebooks, and build web applications to rapidly put findings and solutions in production. Researchers can access their previously scattered data centrally through intuitive no-code/low-code interfaces. They can run analyses, build machine learning models, or prototype web apps all in a central JupyterLab environment that runs on Kubernetes for full scalability. And they can also self-serve previously-siloed IT capabilities to go from rapid prototypes to fully operational applications with zero IT learning curve—all in and from one environment, secure behind your organization's firewall.



All Edge applications—built-ins and custom-built by Enthought or users—can be accessed through the Scientific Workbench.

One Central Hub for R&D

Platform Overview

Enthought Edge has a modular two-layer architecture built to empower scientists and engineers while meeting the policy and security needs of the enterprise.

Edge consists of two layers: an application layer where users can access data, Jupyter notebooks, and their applications; and an infrastructure layer that manages cloud resources, scientific packages, and provides unified access to data. This modular architecture delivers analysis-ready data to the hands of R&D teams, gives them one-click access to scalable cloud compute, and expands capabilities that mirror the iterative nature of research all while allowing IT to ensure security and control.



The Edge Scientific Workbench application interface and underlying infrastructure.

Low to No Learning Curve

Scientists and engineers can easily access analysis-ready data and share assets using their favorite technologies in one place.



Ready-to-Use Familiar Tools

Computational notebooks are used by millions of users to do exploratory analysis and interactive programming. Edge makes them even more powerful. Edge has a builtin JupyterLab IDE that runs on a scalable Kubernetes infrastructure, allowing users to access fully flexible compute power in their notebooks while keeping their cloud costs low. Jupyter notebooks in Edge are supercharged with an enterprise-grade Python distribution to reduce the risk and inefficiencies associated with dependency management. In addition, Edge comes with hardwareoptimized builds of popular machine-learning packages like PyTorch and Tensorflow.





Unified Data Access

Edge provides a one-stop shop for data. Access and load all the data needed for analysis with Edge, irrespective of the data source—individual data stored in the JupyterLab, shared data files on Edge, and remote data can all be accessed from a single endpoint. All the notebooks launched in the JupyterLab on Edge come with a live connection to the Edge Data API. No need to write any hard-to-maintain glue code to send and parse HTTP requests, and no need to learn a dozen *vendor* APIs to get access to your *own* data. Zero overhead with one fully Pythonic API for all your R&D data needs.







Secure Asset Sharing with a Simple Drag-and-Drop

Collaboration is the foundation of high-throughput research. Edge makes sharing assets (data and code) with other authorized users in the organization as simple as drag-and-drop. Researchers can seamlessly access and share their data and scientific APIs with JupyterLab on Edge. Nothing will ever have to leave your secure and private environment. Build, run, and share your notebooks through one secure, central portal.





Tiered Cloud Resources with a Point-n-Click Interface

With Edge, each researcher can choose the just-right level of computational resources for the task at hand. Easily change as the work changes by selecting the optimal level through Edge *Compute Profiles*. A *Compute Profile* is a set of computational resources (CPUs, RAM, and GPUs) defined by the organization's manager that provides the equivalent of a dedicated virtual machine, giving the flexibility to scale the resources used and save on costs when not used. At the same time, managers can use the Edge *Cloud Watcher* to track how much cost is being incurred throughout the organization, set limits to contain costs, and be notified when cloud expenses pass certain milestones.



More than Data. More than Notebooks.

Bridge the technology gap so scientists can focus on the science.

While prototyping ideas is a necessary first step, your innovations must be put into production to unlock their full potential. Only by applying your ideas to solve real-world problems can your R&D efforts generate a competitive advantage, and Edge empowers you to do just that. Edge provides the digital infrastructure needed to efficiently transform your R&D knowledge into value-generating assets. With Edge, there is no need to set up and maintain one-off stacks for your digital solutions or deal with low-level and complex infrastructure components. Edge seamlessly integrates with existing corporate IT stacks, enabling R&D teams to build, deploy, and run digital solutions with zero downtime and minimal IT overhead.



Accelerate Time to Value, Mitigate Risks, and Reduce Total Cost of Ownership



Ready-to-Use Digital Infrastructure

Edge streamlines application deployment for R&D teams, allowing them to easily put their solutions into production without requiring IT expertise. This speeds up research and frees up valuable time for domain experts to focus on driving scientific innovation instead of troubleshooting technical issues. Meanwhile, Edge reduces the workload of IT professionals by providing a single, enterprise-grade stack to meet the infrastructure needs of R&D users. Managers can also monitor resource utilization through intuitive dashboards and proactively manage cloud costs using Edge's administration consoles.

	1
\bigcirc	

Full Visibility and Control Around Your Cloud Spend

With the flexibility of cloud computing comes the risk of overspending, but with Edge Cloud Watcher, you can actively monitor and manage cloud costs without disrupting important work. Edge monitors the use of data, compute resources, and applications and through detailed administration consoles, gives managers the ability to audit their team's usage with full granularity. The consoles also allow administrators to set caps on cloud costs and control the compute resources that are available in their organizations to maximize the efficiency of cloud budgets.

R&D Innovation Made Easy with the Edge Native Application Framework



Accelerate Your Development

Scientists and engineers working in R&D need to move fast. The Edge Native Application Framework enables R&D teams to rapidly turn prototypes to production-ready solutions with minimal IT, DevOps, or Cloud overhead. Our Kubernetesbased architecture eliminates the need to manage load balancing, identity management, SSL termination, reverse proxy management, and other infrastructure requirements for your applications. Best of all, you have full freedom to choose your preferred technology. Whether it is Dash, Panel, Streamlit, Voila, Shiny, Pynecone, Flet, or just React—if you can containerize it with Docker, you can productionize it with Edge.



Scientific apps have variable resource needs. To address this, Edge native applications allow each user to personalize the level of computational resources they use. Edge runs applications in a Kubernetes cluster, providing each user with their own dedicated backend instance of their apps. This design eliminates resource competition between users and the performance degradation and availability issues that commonly occur with traditional web apps. Additionally, it enables each user to select a different level of computational power (CPUs, RAM, and GPUs) with a single click. Meanwhile, Edge ensures cloud costs are under control and can be proactively managed through administration consoles.

Preserve Your Valuable IP

Scientific apps are constantly evolving, with new versions being regularly released as ML models are retrained, novel visualizations are added to dashboards, or scientific algorithms get extended. With Edge, you never have to worry about downtime risks or breakage issues during update rollouts. Previous versions of your native apps are preserved and adding a new version does not replace any of the older ones. The Edge Time Machine then lets you easily navigate through different versions of each app and launch the one you want with a point-and-click interface.



Scientific apps are critical in accelerating R&D, but they can limit innovation when they become black boxes. Tackling new challenges requires extending existing solutions. Once your custom applications are deployed on Edge, not only can users launch and use these apps, but they can also access the data and algorithms behind them. All they have to do is launch the built-in analysis environment in Edge to prototype new features, visualize data in new ways, or apply existing APIs on new use cases. Everything is integrated and ready-touse, providing a unique ecosystem that fosters continuous innovation.

Engineered for The Enterprise

Edge is designed for R&D and built for the enterprise.



Integration of the Best-in-Class Technologies

When building Edge, we made a deliberate choice to use standard technologies with wide adoption from the community. Not only can users get started immediately with familiar technologies, Edge comes with them already integrated for effortless implementation. Edge uses Kubernetes for running apps, as well as internal services like the Edge API. Istio service mesh is used for Kubernetes, to ensure secure isolation between apps and encryption of data in transit. We have standardized on the open-source KeyCloak application for authentication and SSO integration. Additionally, Edge also integrates easily with existing infrastructure using standard protocols (eg OAuth2, LDAP).



Single-Tenant and Self-Managed VPC Deployments

Edge provides flexible deployment models to deliver what suits the enterprise the best and fits within existing policies and processes. No more one-size-fits-none deployments. Use our central fully managed service, or we can set up private single-tenant deployments. The private deployments can be managed by Enthought or your in-house technical team.



No Lock-In at Any Level

Because Edge leverages widely-adopted technologies, there is no risk of "vendor lock-in." Users have the full benefits of Edge as a central hub for all their R&D data, analysis, and application needs but can still connect to their familiar tools like JupyterLab, Kubernetes, and Docker outside of Edge. And because Edge uses a straightforward file/ folder metaphor for data storage, there is no messy migration process or need to reformat data.

REQUEST A DEMO | Contact us info@enthought.com to see how Enthought Edge can drive innovation for R&D and your business.

Enthought Powers Scientific Computing for R&D

Enthought is a globally recognized leader in scientific computing, providing specialized solutions that accelerate scientific innovation across various industries. Our transformative solutions, from Al-assisted interpretations of subsurface seismic data to quantum simulations for material informatics and ML models for cancer therapeutics, have helped businesses achieve breakthrough discoveries in record time. With a unique blend of scientific expertise, technical proficiency, and business experience, Enthought is your trusted partner in accelerating your R&D organization. Enthought is headquartered in Austin, Texas, with additional offices in Houston, Texas; Cambridge, United Kingdom; Zürich, Switzerland; and Tokyo, Japan.



© 2023 Enthought, Inc. All Rights Reserved. enthought.com