Yellowbrick

How Yellowbrick Ensures Data Residency and Security Across All Environments

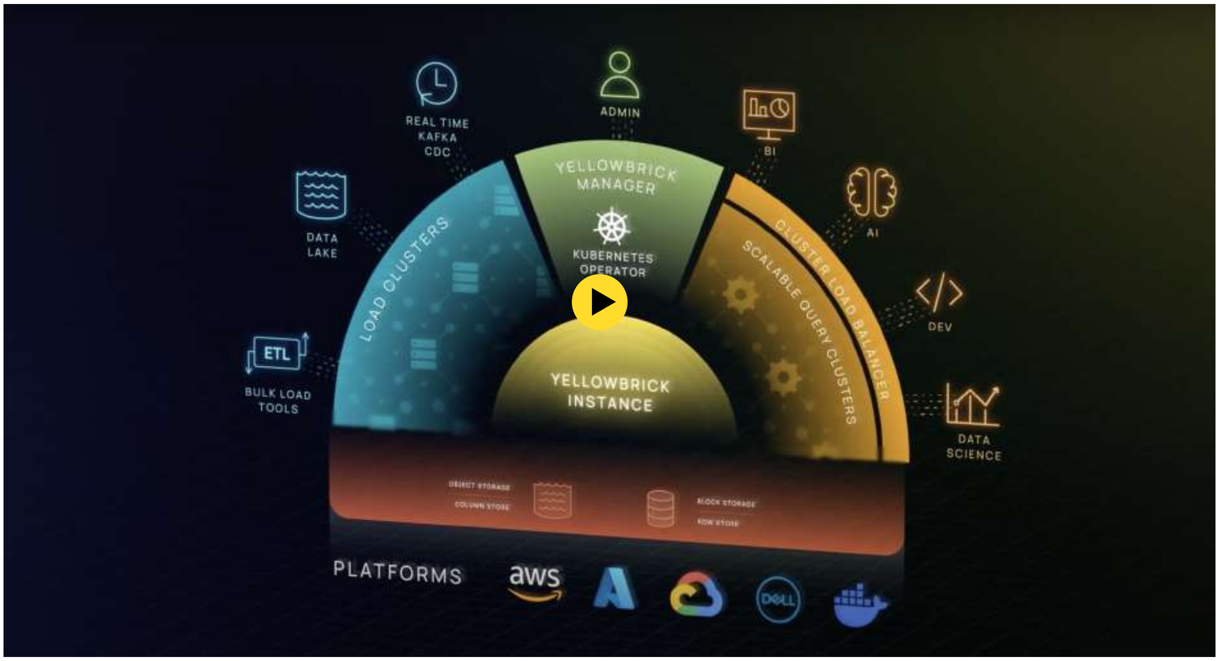

Yellowbrick's Private Data Cloud architecture addresses data residency concerns by offering availability in your cloud account in the public cloud of your choice, in any region, or on-premise, ensuring data processing occurs next to the data, keeping data under customer control, and providing secure methods for data transport and sharing.

PROOF POINTS

Keeping Data Local & Secure

Data residency concerns can arise from contractual, sovereign, and regulatory requirements, the practical realities of multinational organizations, risk management, or global application delivery. These lead to multi-region or even multi-cloud delivery needs. Your organization's risk appetite may even force sensitive data to reside on-premises, or perhaps your organization can deliver a private cloud at lower cost on-premises. Yellowbrick is the only true hybrid multi-cloud data platform that delivers in any cloud and on-premises with the same experience. Why suffer with multiple data solutions in different locations?

Data residency restricts sensitive data to the region or business where it originates, dictating both availability requirements and exportability. A continually increasing number of regulations around data privacy and protection are being formulated by governments and states globally. In addition to traditional data classifications such as PII and PHI, and broad classes of 'personal data' subject to GDPR, more classes of data are being considered as critical to national security.

Businesses must actively consider where data resides, manage its availability, and be sensitive to how it is exported. This is particularly challenging for businesses considering SaaS data platform vendors: Some SaaS platforms may be unavailable in certain geographic regions making a common data platform impossible to achieve. Platforms that require sending data to another vendors' cloud account for processing may not be acceptable, and even those that have spit data and control planes can still lead to leakage of sensitive data if there are multi-tenancy bugs.

Yellowbrick's Private Data Cloud architecture addresses these data residency concerns in a straightforward and easy to understand fashion:

- Yellowbrick is available in all regions and on-premise
- Processing is kept next to the data
- Data never leaves the customer's control
- Methods for securely transporting and sharing data are included

Innovative Architecture

Yellowbrick is the ideal data platform for businesses who care about data residency and security. By fully embracing Cloud Native architecture and Kubernetes, Yellowbrick's software runs in all major cloud providers and on-premises deployments yet runs largely on autopilot, requiring no specialist management skills or Kubernetes experience whatsoever.

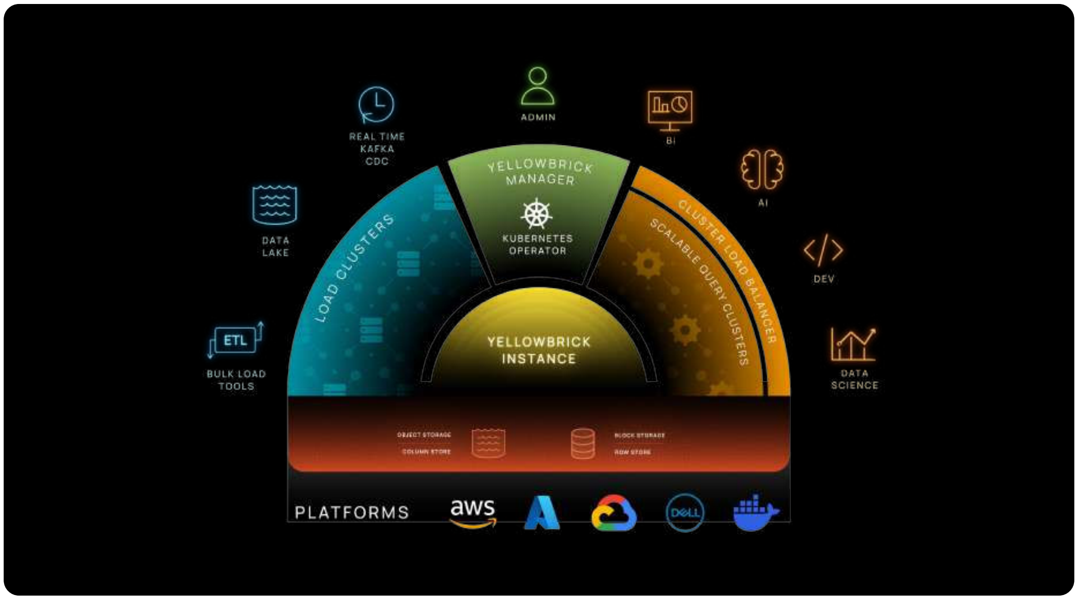

Yellowbrick is deployed into a customer's cloud account, secured in the same manner as all other applications. Data stored by Yellowbrick is persisted on object storage under the same cloud account. Similarly, Metadata is persisted in block storage under the same cloud account. The database software runs in the same cloud account, and even logs and diagnostics don't leave the cloud account. No queries or data are transmitted to Yellowbrick. For customers with very strict security requirements or sophisticated network topologies, the software can be deployed into a pre-formed OCI registry and VPC/VNET. No data egress is required.

In all cases, Kubernetes nodes are always running in private subnets with private IP addresses. The data path to object storage always happens over a VPC gateway endpoint, and customers can optionally provision other supported VPC service endpoints for access to cloud services if fully private data paths are needed. Service accounts (workload identity, IRSA) are used to access cloud resources such that no credentials or keys are passed around.

For customers running in on-premises data centers, Yellowbrick is delivered as a hyperconverged appliance which co-locates all compute, storage and networking with no dependency on outside resources or the public cloud.

Data Sharing When Needed

Although Yellowbrick defaults to a fully private deployment in customers' own networks, a rich set of functionality caters to advanced enterprise concerns.

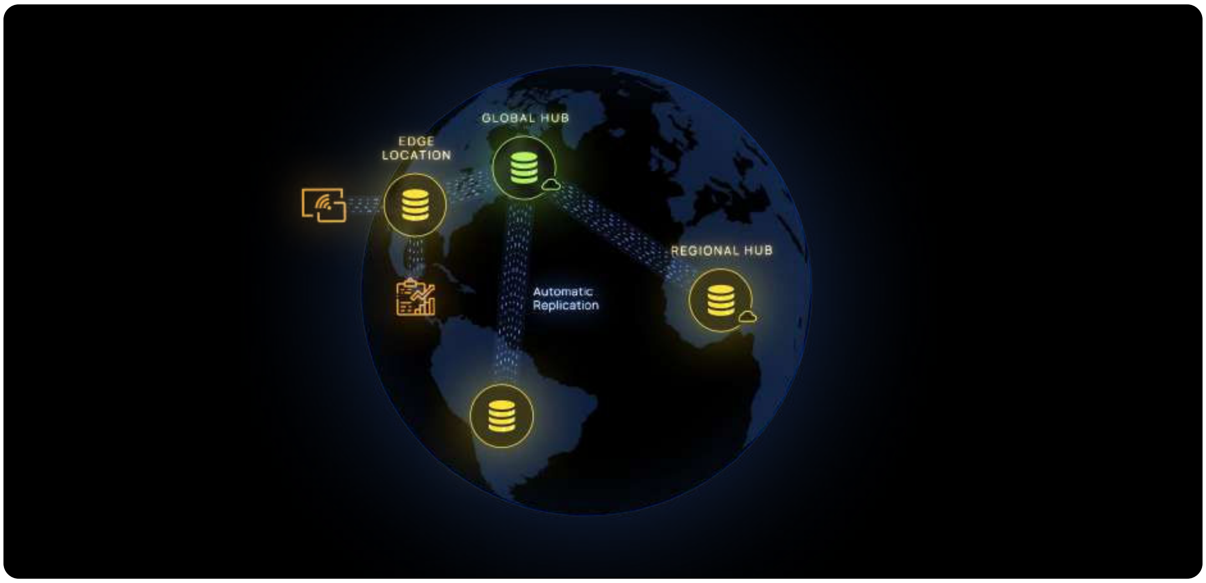

- Replication: Yellowbrick supports enterprise-grade replication for high availability and disaster recovery. Both schema, system metadata and data are replicated between regions, clouds, and data centers. Replication is asynchronous, supporting failover and failback, with configurable RPO down to a few minutes.
- Business continuity: Support for full, incremental and cumulative backups, both on-site and offsite, means unlimited backup retention.
- Bulk loading and unloading: Datasets can be loaded and unloaded at effective rates of 10s of terabytes per hour.
- Private data sharing (in preview): Point-in-time snapshots of databases can be curated and shared between users, across clouds and on-premises environments, minimizing copies of data.
- Standard integrations: From AirByte to Informatica, from Kafka to GoldenGate, Yellowbrick integrates with all the standard enterprise data movement tools.

Even though Yellowbrick runs within your own cloud account or data centre, it still ships with all the expected security features you'd expect. Authentication of users is accomplished by using database local accounts or external accounts authenticated using a choice of OAuth2 provider, Kerberos or a traditional enterprise LDAP integration. Role-based access controls (RBAC) allow fine-grained control of most database administration-type features without needing a super-user or 'root' account. Columnar masking and encryption allows columns containing sensitive data to be encrypted with customer-managed keys.

Although Yellowbrick's architecture is designed to minimize the use of insecure third party and open source components, all standard software packages have routine fixes for CVEs. Yellowbrick ships monthly software updates to remediate known vulnerabilities.

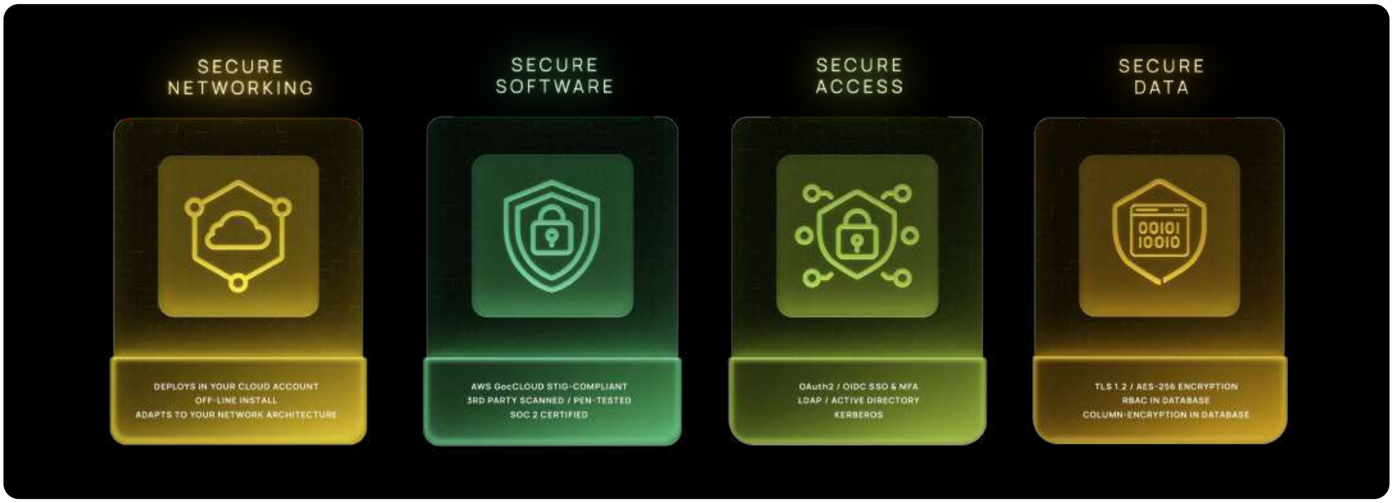
SECURE NETWORKING
DEPLOYS IN YOUR CLOUD ACCOUNT
OFF-LINE INSTALL
ADAPTS TO YOUR NETWORK ARCHITECTURE
SECURE SOFTWARE
AWS GocCLOUD STIG-COMPLIANT
3RD PARTY SCANNED / PEN-TESTED
SOC 2 CERTIFIED
SECURE ACCESS
OAuth2 / OIDC SSO & MFA
LDAP / ACTIVE DIRECTORY
KERBEROS
SECURE DATA
00101
10010
TLS 1.2 / AES-256 ENCRYPTION
RBAC IN DATABASE
COLUMN-ENCRYPTION IN DATABASE

Yellowbrick is FIPS-compliant and all external connections are TLS-encrypted. Regular updates for CVE remediation are provided.

As a mature enterprise offering, off-the-shelf integrations with many best-in-class enterprise data protection solutions are available and supported.