

Enterprise Architecture Overview

Appian's industry-leading application platform has been recognized by industry analysts and customers alike for its ability to deliver business value faster, thanks to its easy, powerful, and unified attributes. From business process management (BPM), intelligent automation, case management and enterprise mobility, Appian's key to success is its focus on customer satisfaction, no matter the project size. This is possible due to Appian's comprehensive, standards-based architecture and proven delivery approach.

This whitepaper describes Appian's platform architecture, redundancy and scalability considerations, security, integration capabilities, and Appian Cloud options.

APIAN ARCHITECTURE

Appian is an n-tiered, web-based enterprise application platform. It includes core components that work together with other systems to provide capabilities and services to users. Appian's highly scalable and extensible architecture is designed to operate in the most demanding enterprise environments, and it successfully does so today:

Appian customer deployments range from just 100 users to millions of active users around the world.

Appian is available as both a cloud and an on-premises solution. Appian Cloud allows customers to focus on delivering their solutions, while leaving the administration, monitoring, upgrades and architecture to Appian.



On-premise Appian solutions provide customers with total deployment and operation control.

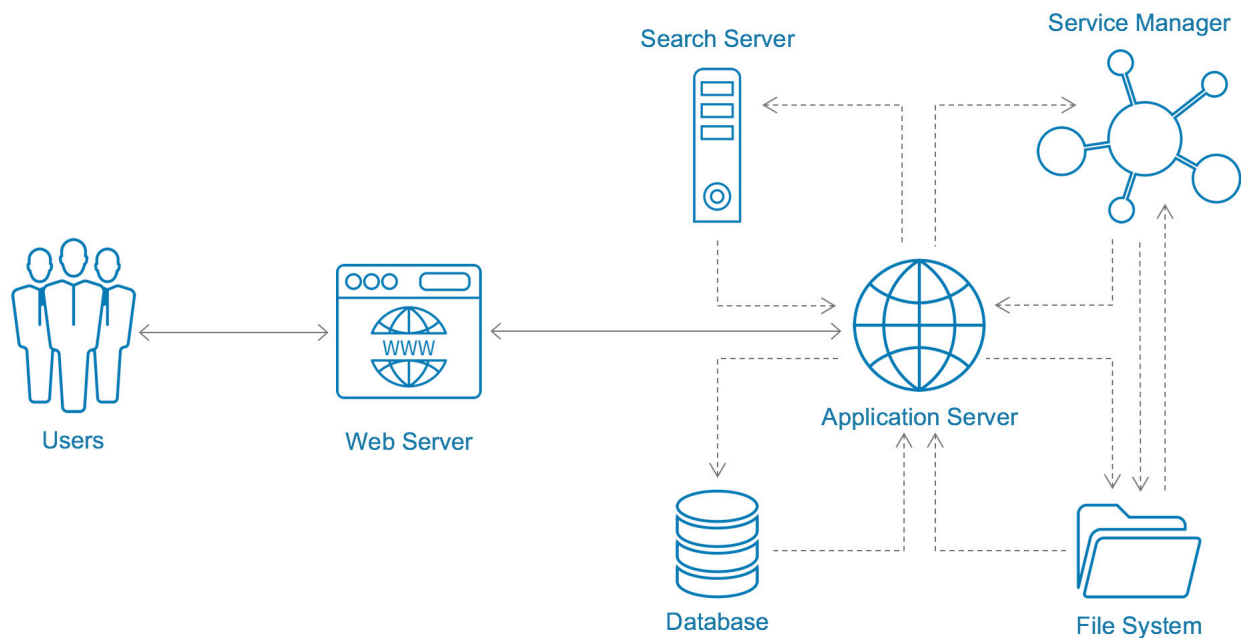
Applications built on Appian are entirely portable between cloud and on-premises deployments. The following sections walk through the Appian architecture and explain the function and interaction of each component.

Server Elements

Appian's server and software configuration will vary depending on the size and requirements of a deployment. Servers are generally distributed across the following components:

- Web server
- Application server
- Appian engines
- RDBMS (MS SQL Server, IBM DB2, Oracle, MySQL)
- Search server
- File store system (single or distributed)

These components can be operated individually or combined in servers, which can be replicated, distributed, and expanded as needed.



Appian can be deployed on a wide variety of enterprise server topologies, including physical and virtual, clustered or distributed, on-premise or on the cloud. Specific system requirements, deployment plans, and recommendations are available as part of the product delivery and they are updated with each release of Appian.

Software Elements

At a high level, Appian is composed of four main software components:

- The Appian web application
- Backend, Service Manager and in-memory Appian engines
- One or more search servers
- One or more relational databases hosting Appian internal (primary) and Appian application (business) databases

The Appian web application is a Java EE application that runs within an application server such as JBoss EAP or WebSphere. The web application serves requests from users' browsers or mobile apps and is primarily responsible for all end user, designer, and administrator web interface features.

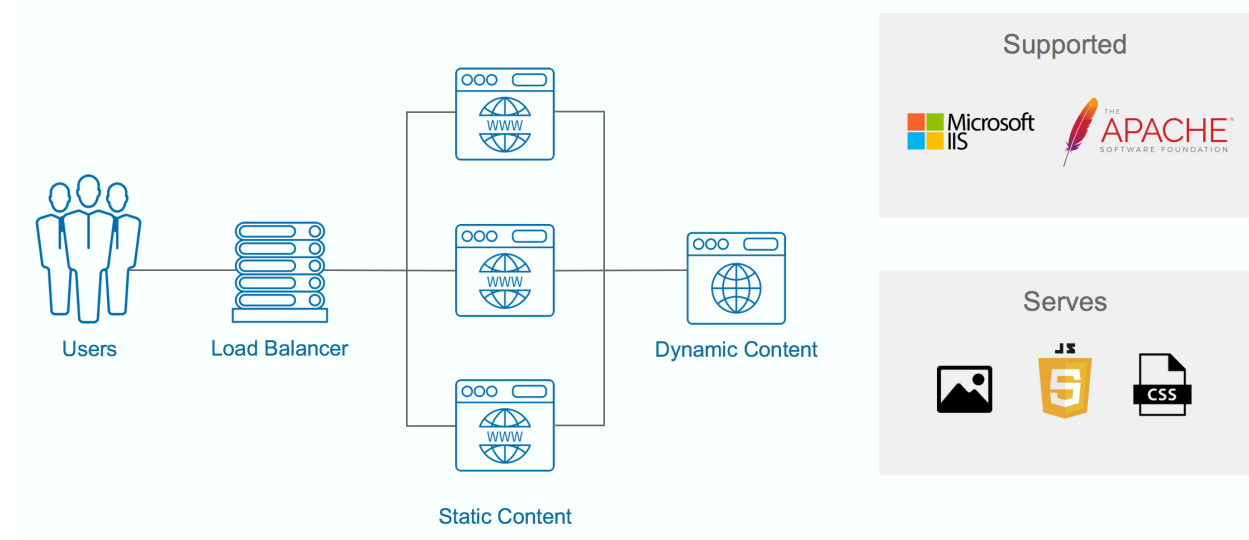
The Appian engines contain metadata for all the Appian objects created by the designers as well as runtime data created by users or processes. Data stored in the engines is accessed and updated by the web application.

The search server provides support for fast data retrieval capabilities used by many application features.

Appian requires a primary relational database such as MySQL, SQL Server, or Oracle to store runtime data for Appian interface features, as well as additional designer metadata not stored in the Appian engines. Appian can connect to additional business databases to access and store business data associated with individual applications.

Web Server

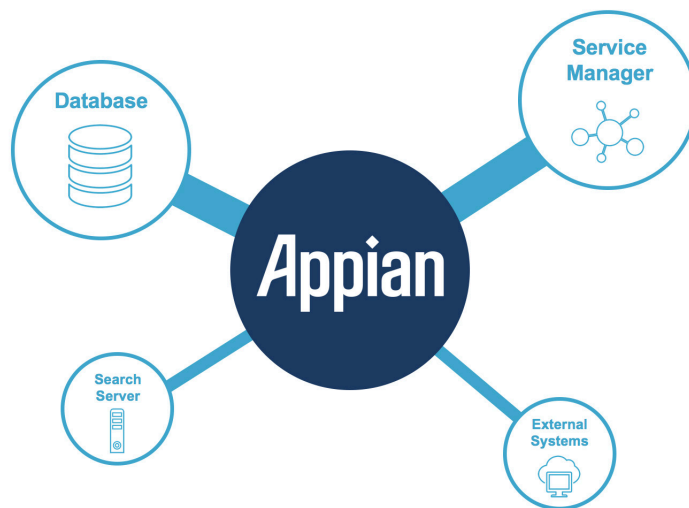
As with most web applications, Appian instances use a web server to handle client requests before passing traffic to the application server. The web server is typically used to handle static requests which enable that content to be cached by client browsers for improved performance. Certified Web Servers include Apache HTTP Server and Microsoft IIS.



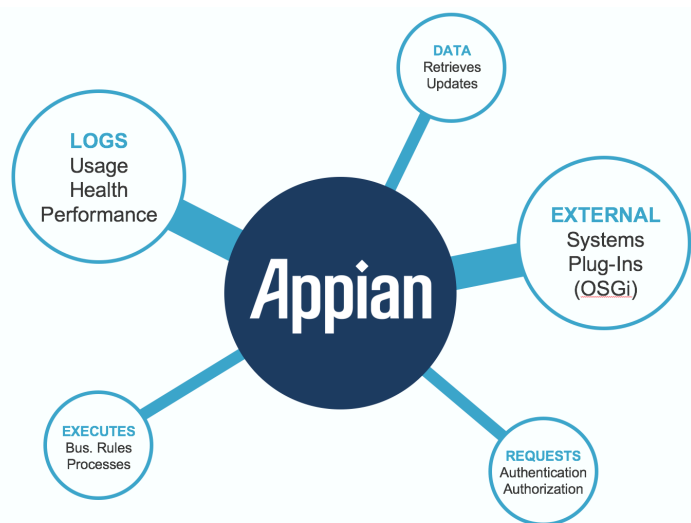
Requests for non-static content are passed through from the web server to the application server. The connection configuration between these components depends on which web and application server platforms are used.

Application Server

The application server provides built-in support for connecting to a wide range of related system components. The supported application server platforms (e.g. JBoss EAP) are implemented in Java.



The application server coordinates most of the interaction between system components and is responsible for a significant portion of Appian's functionality. It runs two Enterprise ARchive (EAR) modules: the Appian EAR and the search server EAR.



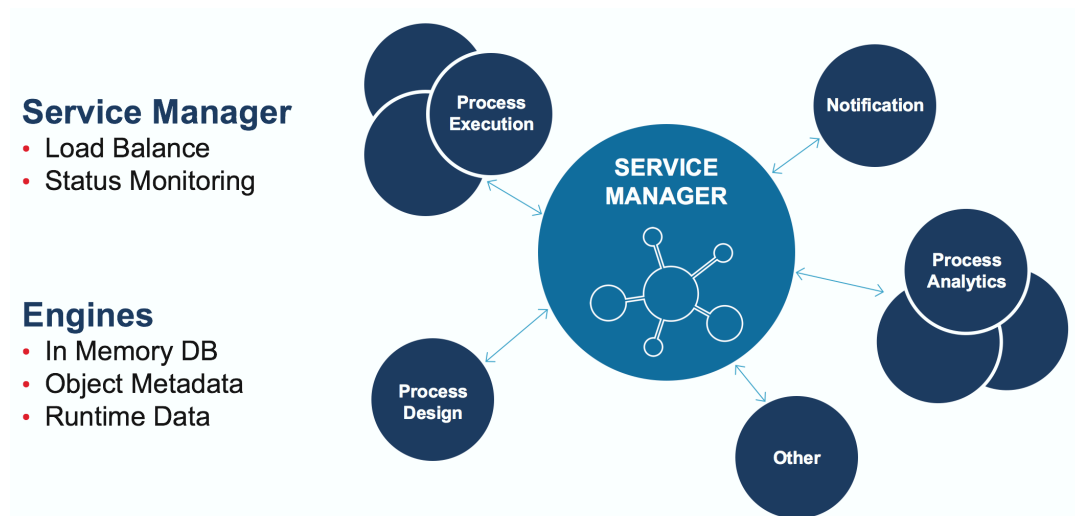
The Appian EAR represents the core of the Appian application.

- It handles end-user client, web or mobile, requests that are passed through from the web server (including authentication and authorization for those requests).
- It retrieves and updates data in the Appian engines, as well as the primary and business databases.
- It manages documents uploaded by users and generated by processes.
- It executes business rules, runs the activities defined in process models, and communicates with external systems.
- It can be extended using plug-ins deployed using the Java OSGi framework.
- It also is a central source of logging and other information about system usage, health, and performance.

Appian Engines

High performance, mission critical systems like Appian, require extremely high throughput and excellent response times to consistently deliver the proper user experiences. When dealing with potentially millions of transactions in a short period of time, choosing the right back-end infrastructure is critical.

Appian Engines are high performance in-memory databases that allow for very fast read and write access, as well as extremely high scalability, often performing at speeds that are orders of magnitude faster than typical relational databases. The in-memory data is also persisted to disk, allowing use of standard backup processes on all data. Real-time synchronous transaction logs capture all events that occur between in-memory checkpoints, allowing any recovery process to apply logged transactions on restart.



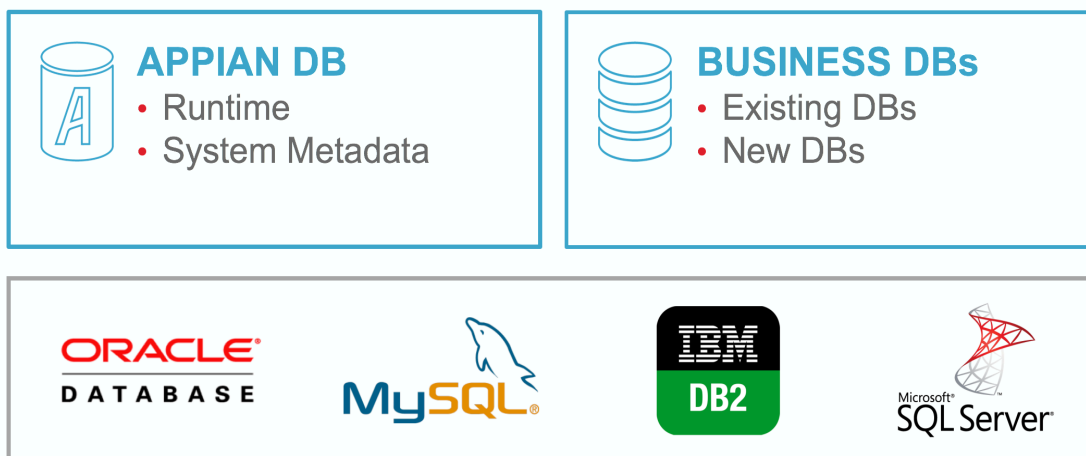
Each engine receives requests from the application server using TCP/IP protocol and manages its corresponding database, which stores data in memory. Engine transactions are also written to the file system. Finally, each engine may be replaced and distributed as needed to accommodate different scaling and redundancy/failover scenarios.

The search server contains an ElasticSearch server and aggregates data from the rest of the application to support features like tracking historical performance, viewing recent user activity, and analyzing design-time impacts/ dependencies. The search server runs as a stand-alone Java application, and multiple search servers can be configured to allow for both data redundancy and high availability.

Relational Database

Appian requires one relational database to be configured as a primary data source. This database is used to store system metadata like data types and security settings, as well as run-time data like News post content.

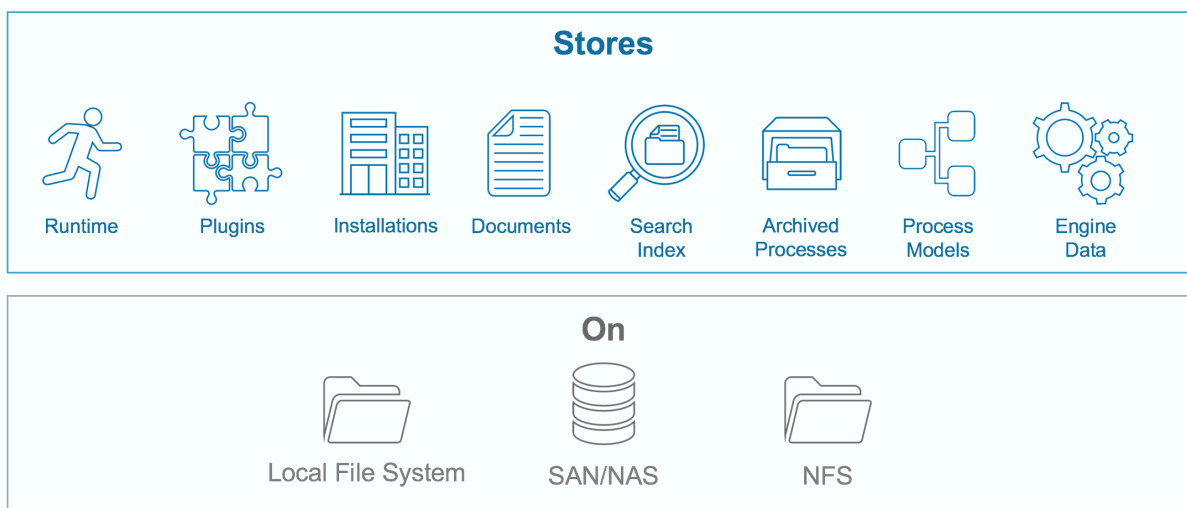
Additionally, Appian can connect to as many business databases as needed. These are used by applications built in Appian to store and retrieve business data specific to those applications. Data from these sources can be used in many places within an application, including in processes, on forms, as records, and in reports. Secondary databases may be created as part of either an Appian application or existing databases that Appian uses.



Database connections are managed using the application server's JDBC data source connection. These settings specify the database host, credentials, pooling, recovery, and other connection options. Certified platforms include MS SQL Server, IBM DB2, Oracle and MySQL.

File Storage

The application server and Appian engines are installed on—and store run-time content on—the file system. Each stores data in independent directories and does not use the file system to share data with one another. Appian can be deployed in a distributed environment with multiple application servers or multiple engine gateways on different servers.



Authentication Server

Appian uses Spring Security to provide a comprehensive and extensible authentication framework. Appian's native authentication is PCI compliant for all clients. Password complexity and expiration rules, as well as other PCI-Compliant configurations can also be set. User credentials are validated against Appian's internal account data in the personalization engine. The default configuration also includes features like "Remember Me" authentication and password complexity/expiration controls.


In addition, Appian supports LDAP/Active Directory, SAML, OAuth 2.0. All external authentication solutions require additional systems and configuration.



Clients: Web and Mobile client

Appian's end user application interfaces are supported on all major web browsers, and mobile native apps are available for the most popular platforms.


The design interface—used to build applications—and the system administration console are also 100% web-based.



All web and mobile clients access Appian using HTTP/S. Appian does not use browser extensions or plugins.

Integration Services

As an enterprise platform, Appian integrates with external systems to gather and display information to users, move data between systems, make decisions in business processes, and more. All of these integration capabilities are managed by the application server and can be extended using Java OSGi plug-ins.



Appian provides several built-in integration designers for standard HTTP and OpenAPI web services, as well as Artificial Intelligence services like Amazon, Google, Microsoft and others.

Appian also provides built in connectors that enable rapid integration with Content Management Interoperability Services (CMIS), Microsoft SharePoint, Microsoft Dynamics, Salesforce.com, SAP, and Oracle Siebel. Appian for Microsoft SharePoint provides additional support for exposing Appian collaborations, tasks and actions inside Microsoft SharePoint pages as SharePoint Web Parts.



In addition, Appian can connect using general purpose integration methods like JMS, file transfer protocols, email listeners, and consumption of web services. Both SOAP and REST web services are supported and can read and write data from user interfaces and transactional processes.



Appian also simplifies incoming integration by exposing process models as web services, as well as by allowing designers to create REST Web APIs. Web APIs provide a way to expose data that is stored inside of Appian or accessible by Appian to another system or to trigger Appian application processes from an external system.

Redundancy and Scalability

Appian supports a wide range of deployment topologies, from a single server to a fully clustered configuration. The flexibility in Appian's deployment topology allows for configurations that adapt to support of growing or changing demands. Redundancy and data replication can be handled at the hardware level (disk replication) or at the software level (real-time duplication via TCP/IP). Below is an overview of some of the most common configurations, as well as how each component can be scaled.

Minimum Infrastructure (Consolidated)

A consolidated environment contains all required software components in a single server and represents the minimum hardware required to run the full software stack.

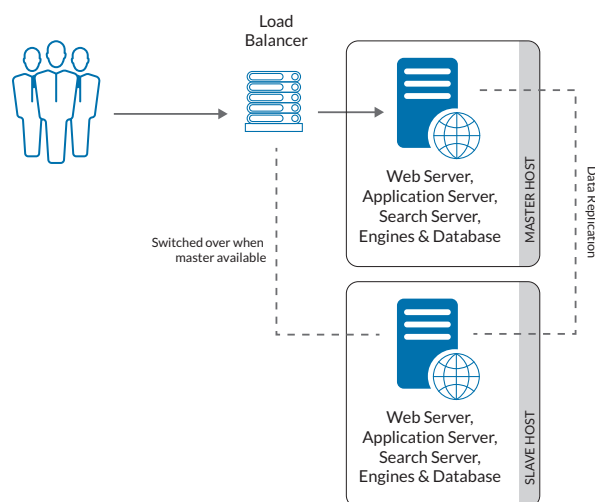
This is suited for development, testing, and low volume production environments that don't require high-availability.

- Simplest configuration to maintain
- Easy to snapshot for backup and restoration
- Limited scalability potential
- Does not provide high-availability or load balancing

Minimum Hardware High-Availability

Hardware-based high-availability relies on the hosting infrastructure to replicate data across the hosts, as well as the switch over when one becomes unavailable.

This is suited for low-volume environments that require high-availability and have two or more virtualization hosts available.



- Simplest configuration for high-availability environments
- Easy to snapshot for backup and restoration
- Limited scalability potential
- Provides high-availability
- Does not provide load balancing

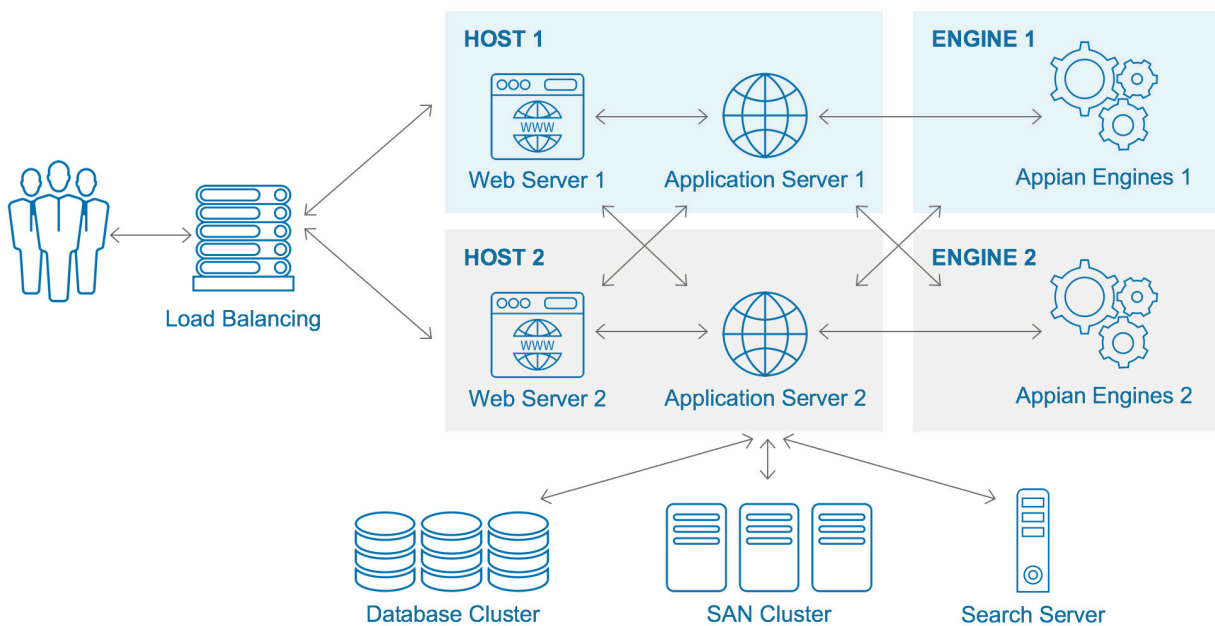
Distributing Components

As resource utilization increases, one or more components will compete for the available resources. A point will be reached where additional resources cannot be further added to the server, thus requiring the resource-intensive components to be distributed to their own servers that can be sized and managed individually according to the demand.

Fully Clustered

To handle very high volumes of concurrent users and processes, each software component can be hosted on a separate server and sized individually to meet extreme high-availability and load balancing scenarios. This allows for multiple nodes in the chain to be unavailable while still providing service.

Additional replication is required for this scenario to provide software-based load balancing and failover. This topology can be used when every component is heavily utilized and requires high-availability and load balancing.



- Most complex topology to configure and maintain due to the additional hardware elements
- Provides full scalability of every component for software-based, high-availability, and load balancing
- Requires more effort to coordinate backup schedules to create a valid restore point
- Requires a shared storage area to host the application data managed by the application servers
- Requires a load balancer to distribute user load across the web servers
- The SAN and database must be accessible by all application servers

“Appian leads the BPMS market in production deployments in the public cloud, and it is one of the few cloud platforms that can manage business outcomes as well as processes spanning on-premises and cloud environments.” – GARTNER

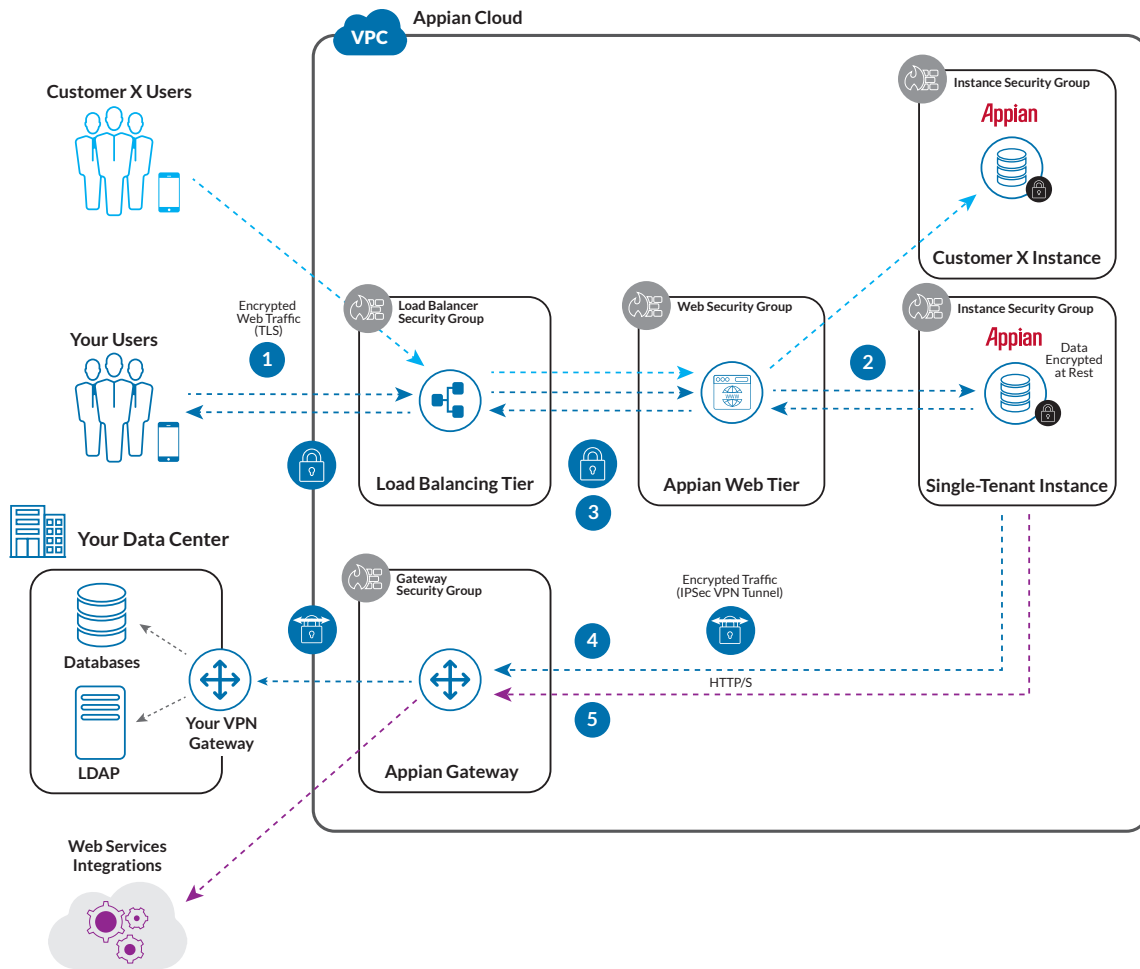
The market for enterprise corporate software is moving rapidly to the cloud. Cloud computing is an emerging way to deliver enterprise applications that are dynamically scalable, virtualized, and delivered as a service over the internet. Customers who employ Appian Cloud as their platform enjoy a number of benefits, including:

- Low startup costs • Fast deployment • Automatic upgrades
- No server maintenance • Fast return-on-investment • Predictable costs



Appian is the recognized market leader in cloud BPM, Case Management and Low-Code. What’s more, solutions built on Appian are 100% portable between on-premise and cloud computing deployment options in support of your unique deployment requirements.

In addition to Appian Cloud’s secure foundation, customers can choose their preferred geographic region to host their data and applications to meet their regulatory and security requirements. In all cases, Appian ensures all data contained in a customer’s Appian Cloud instance is protected and never copied outside the geographical regions designated by the customer.



- 1 Users access the Appian Cloud instance from their browsers or their mobile devices. All data in transit from the user to the Appian Web Tier is secured using industry standard Transport Layer Security (TLS) encryption. The Appian Web Tier is behind a Load Balancing Tier which forwards the traffic to one of multiple web servers processing the user request. The Appian Web Tier is also protected using AWS security groups and only allows HTTPS traffic originating from the Load Balancing Tier.
- 2 Intra-server communication in Appian Cloud is controlled using security groups. Security groups are default deny-all and configured with firewall rules to restrict access only from specific Appian Cloud services. For example, rules protecting your instance only allow traffic from the Appian Web Tier to the application server running on your instance. Traffic originating from other customers and all other unspecified traffic is blocked by the security group. For customers with specific compliance requirements, Appian supports the configuration of an IPSec VPN tunnel for communication between these two layers (i.e. Appian Web Tier and your Appian Instance).
- 3 Once the request is processed by your Appian Cloud instance, the security group will allow the response back for the TCP session created with the Appian Web Tier. Traffic sent back to the user will be encrypted in the same fashion as it was received.

- 4 IPsec VPN tunnels can be configured from your Appian instance to your on-premise resources such as databases or LDAP servers. AWS security groups are also configured to accept inbound IPsec traffic to your Appian Cloud instance only.
- 5 Customers can extend applications running on their Appian Cloud instance integrating with external systems, data repositories and web services. Customers are responsible for the security and encryption configuration when integrating with systems outside Appian Cloud.

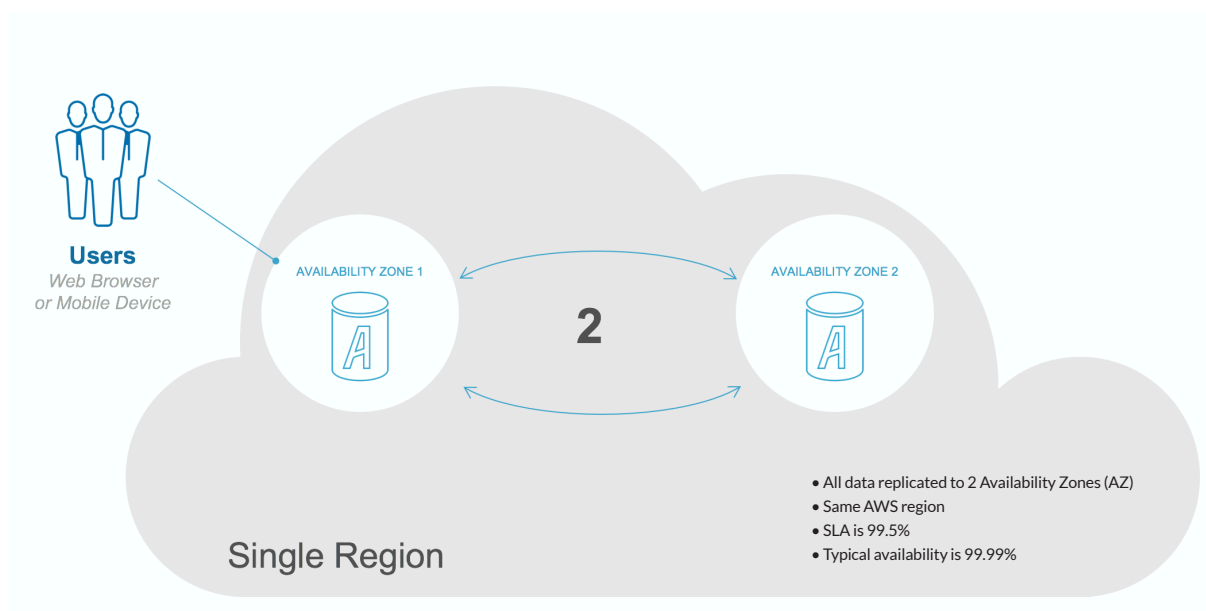
Reliability and Scalability

Appian Cloud provides reliability and security that can be tough to match even by internally managed environments. With a 99.95% SLA uptime guarantee, Appian can meet the needs of enterprise customers with the necessary guarantees, as well as historical performance in place.

The Appian Cloud architecture is designed to deliver maximum uptime and minimal to zero data loss but in case a major service disruption happens, Appian Cloud provides a formal Disaster Recovery plan to insure minimal-to-zero loss of production data.

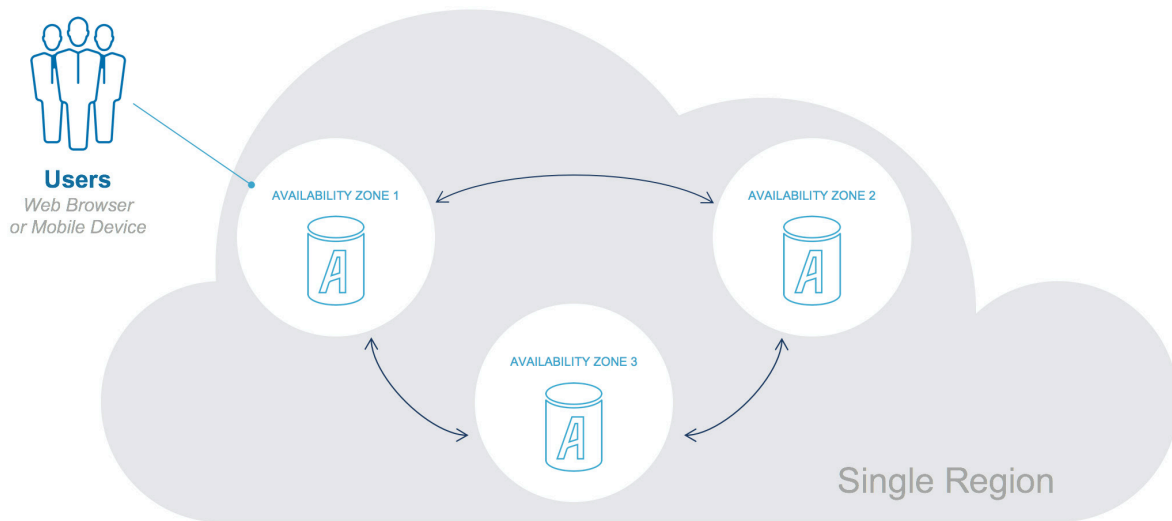
Appian Cloud Standard Availability

By default, Appian Cloud replicates all customer data across two physically separate data center locations within the same geographic region.



Appian Cloud High Availability

An optional High Availability offering extends replication to three data centers within the same region.




Individual systems in a given data center use mirroring to protect against hardware failures. In addition, all production data is included in nightly backups stored in a separate isolated location within the same geographical region.

Appian performs regular testing of its disaster recovery program to ensure readiness, viability and effectiveness.

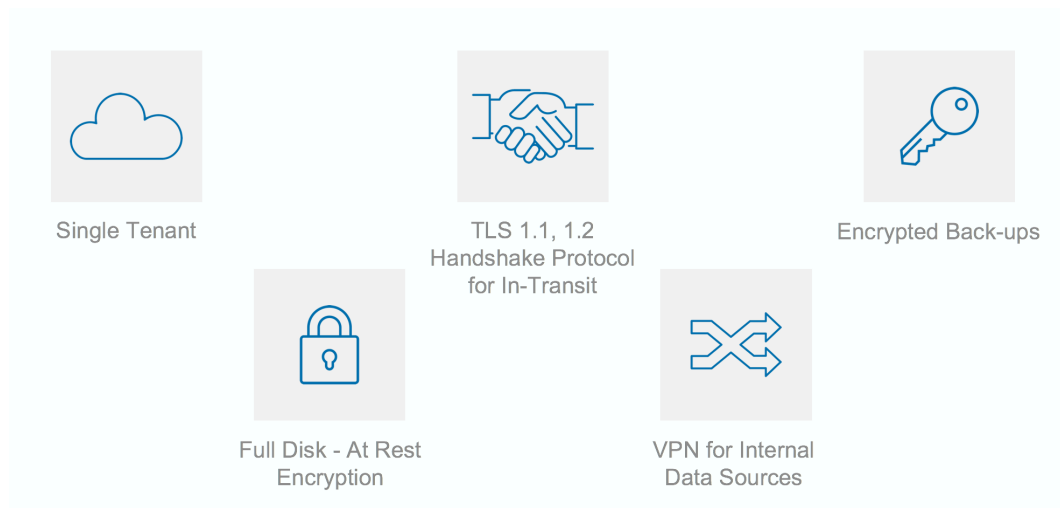
On-Demand Instances

Customers can expand server capacity with on-demand cloud development and test services, and benefit from automatic upgrades to the latest release. Instead of buying new equipment for short-term development or tests, customers can purchase a monthly cloud license to quickly provision capacity and increase business agility.



Security

Appian provides all the necessary security technology, personnel and procedures to meet the needs of enterprise customers.



Appian Cloud's security and compliance program follows the National Institute of Standards and Technology (NIST) 800-53 risk management framework, which directly aligns to HIPAA and other stringent compliance and security rules.

Appian Cloud implements the following NIST 800-53 controls:

- Access Controls and Authentication
- Audit and Accountability
- Contingency Planning
- Incident Response
- Personnel and Physical Security
- Risk Assessment
- System Acquisition and Integrity
- Systems Communication Protection

Authentication

- SAML, LDAP, Active Directory integration for authentication and single sign-on
- PCI DSS compliant login and password management features
- OAuth 2.0
- Virtual Private Network (VPN) for extending your data center

Communications

- TLS encryption for all communication

Storage

- Local geography hosting
- Data and application segmentation
- Encrypted data storage and backups

Security Certifications

FISMA Certified, FedRAMP, Gramm–Leach–Bliley Act, PCI-DSS, HIPAA, UK Data Protection Act, Mass Privacy Law, NIST Risk Management Framework, US ITAR.



Source: trust.appian.com

For up-to-date information on Appian's security and regulatory compliance program, please visit trust.appian.com

Portability

Avoid cloud lock-in with full portability of enterprise data and solutions. Appian customers enjoy the freedom to migrate between their cloud and on-premise deployments at any time, including all Process Models, Rules, Data, Reports, Content, Forms, Groups, and Collaboration capabilities. The bottom line? It's your data and your solutions. With Appian, you always have control of them.

Appian

Appian provides a leading low-code software development platform that enables organizations to rapidly develop powerful and unique applications. The applications created on Appian's platform help companies drive digital transformation and competitive differentiation.

For more information, visit www.appian.com