

RadiantOne FID: A Federated Identity and Directory Service Based on Virtualization

Faster Authentication, Smarter Authorization, and a Common Identity for WAM and Federation

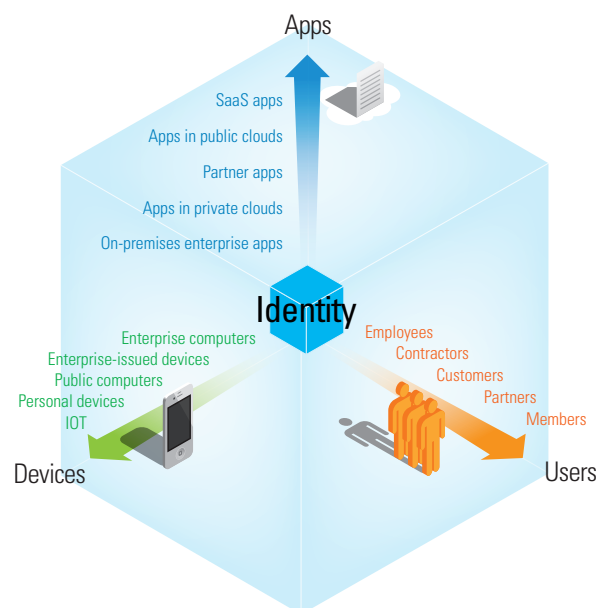
RadiantOne provides a federated identity and directory service based on identity and context virtualization, combined with a high-capacity/high performance directory based on big data technology. Working with your existing identity stores, RadiantOne FID delivers faster deployments, lower integration costs, and the flexibility you need to navigate changing business requirements. With RadiantOne, virtual directory technology has evolved into an easy-to-use, enterprise-grade solution for stronger authentication and faster Single Sign-On (SSO), fine-grained authorization, customized user experience, and the speed and reliability that only a storage layer based on big data technology can provide.

Integrating Identity within a Fragmented Infrastructure

Corporate identity systems are often the result of years of technological accumulation. They feature a conglomeration of disparate data silos made up of multiple directories, data sources, and protocols (such as LDAP, SQL, and web services). Identities are spread across many disparate data silos and the same identity often exists in more than one source.

Compounding this problem is the rapid expansion and evolution of applications, access devices, and user populations (and their identity sources). This is not a one-time phenomenon. In fact, companies and organizations are facing a permanent challenge when it comes to security and IAM.

To enable authentication, SSO, security policies, and smart authorization, the system must identify and authenticate users and gather their attributes. Performing these tasks across multiple heterogeneous data silos constitutes **a major integration task** that requires costly, time-consuming customization and synchronization. This leaves enterprises less able to adapt to changing business requirements, whether they're federating with partners, adding new users following a merger or acquisition, or adopting critical new cloud-based applications.



Rapid growth and evolution of applications, access devices, and user populations are creating a tough challenge for security and IAM.

The Secret to Fully Leveraging your Federation IdP: Integrating your Internal Identity Sources

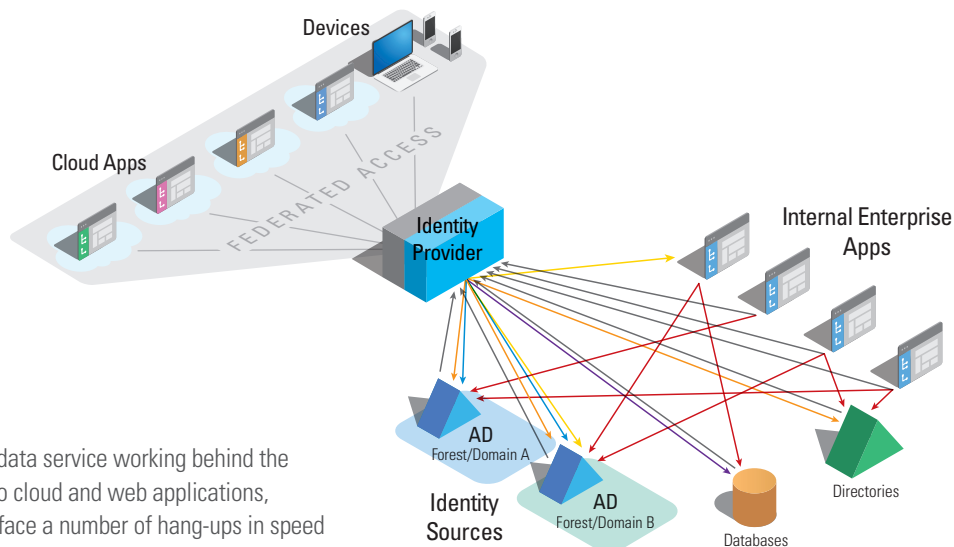
To address those challenges, many companies and organizations have adopted SSO based on federation standards. However most of their existing installations or newer deployments overlook another challenge—the fragmentation of their internal identity sources.

When you access an application secured by SAML or OpenID Connect (for instance, a cloud-based application), the authentication requests are routed to an Identity Provider (IdP). The IdP is then responsible for ultimately authenticating the user. Deploying an IdP assumes that there is a unique or unified identity source acting as an authoritative source of identity behind your IdP.

The reality for most sizeable enterprises or organizations is that there are multiple silos of identity, composed of various data sources and protocols (such as LDAP, SQL, and web services). And as a result, the lack of integration hampers the scope of their deployment (not every identity source can participate in the SSO service) and /or slows down the overall time it takes for authentication (too many scattered sources to look through to identify a user).

The identity source fragmentation challenge—Popular IdPs aren't designed to sort through internal heterogeneous data stores. Without a common source of identity and attributes, incorporating or reorganizing populations is costly, time consuming, and requires extensive customization.

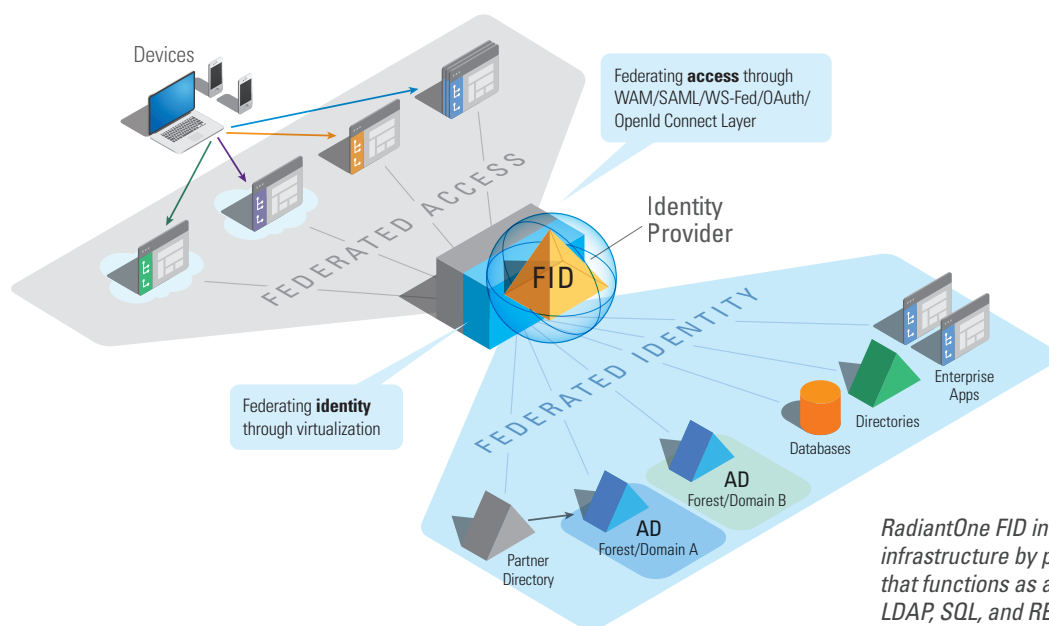
Without some form of integrated identity data service working behind the scenes, functions such as providing SSO to cloud and web applications, password management, and provisioning face a number of hang-ups in speed and resilience, not to mention security risks.



RadiantOne: Fast Authentication and SSO Through Identity Federation

For any security and identity team, it takes hard work to ensure that all users—including employees, members, providers, and vendors—have secure access to a set of applications. Aside from the different methods and limitations imposed by the applications themselves, with some supporting SSO using federation standards and others supporting only proprietary methods, there is the challenge of multiple identity sources with their different formats and divisions, such as Active Directory domains and forests, LDAP, SQL, and more.

Optimized for the enterprise environment, RadiantOne hides the heterogeneity of existing identity sources, providing simple, logical, standards-based access to all the identities within your organization—no matter where or how they're stored.

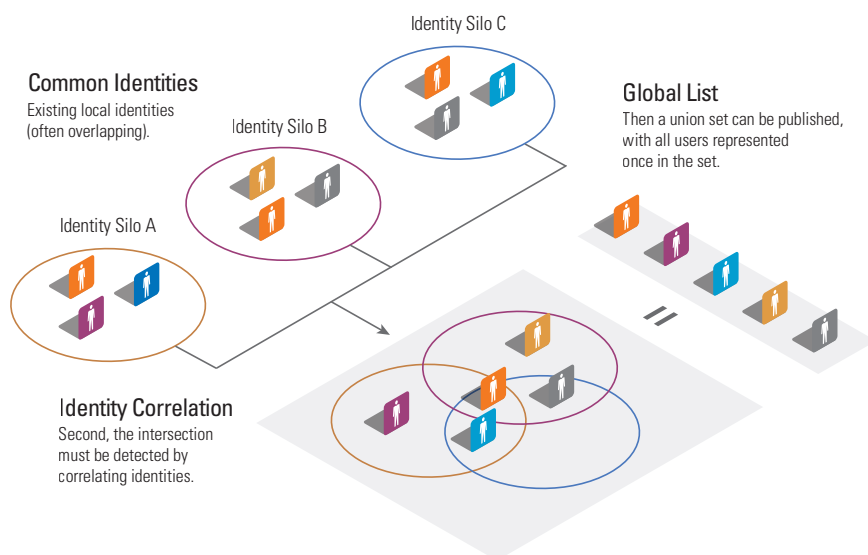


RadiantOne FID integrates your identity infrastructure by providing a central hub that functions as an access point for all LDAP, SQL, and REST-enabled applications across security domains and AD forests.

This makes deployments much faster, whether WAM or Federation IdP. And because it provides a common source of identity and group information, it greatly simplifies the process of reorganizing existing populations and incorporating new populations that come from mergers, acquisitions, partners, etc.

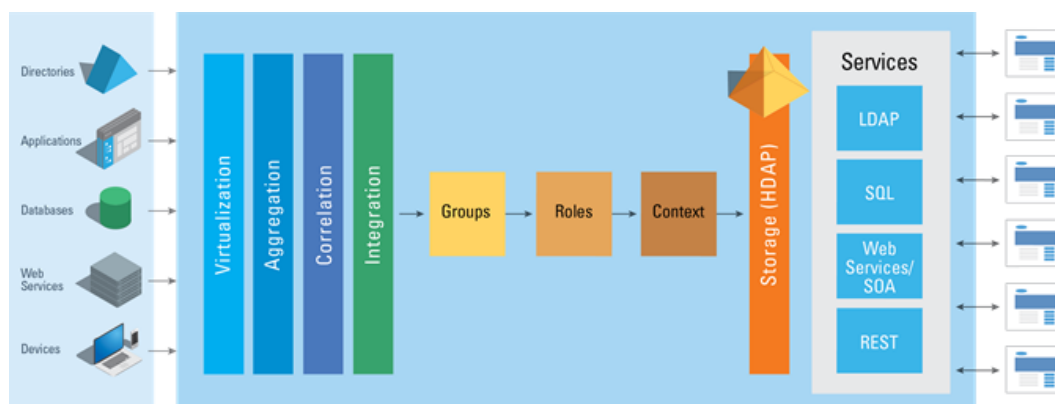
Virtualization Shields Applications from the Complexity of Backend Data Stores

The key capabilities of RadiantOne FID are data virtualization, union through identity integration, and creation of the global profile through join. It discovers and extracts schemas and data models from backend sources, then translates the schemas into an XML-based data format in the virtualization layer.



RadiantOne creates a global list where all duplicate accounts have been removed—the process of correlation links same-accounts so that a union-compatible list of all identities from across all data sources can be made.

If there is user overlap, RadiantOne FID correlates the accounts of same-users across data sources to each other to create a unified global list, and then joins those accounts to provide complete profiles of users.



Architecture: RadiantOne acts as an abstraction layer between applications and the underlying identity silos. Virtualization protects applications from the complexity of backends.

Application-specific attributes can be stored at the virtualization layer, without requiring schema extensions to the underlying data stores. Then, RadiantOne FID builds custom hierarchical views containing the complete profiles to meet the needs of each application, so that all your applications get the identity they need, in a format they can understand. And RadiantOne's persistent cache feature adds reliability and scalability to your identity infrastructure, ensuring that profile information is always available and always up to date.

LDAP

SQL

Web Services/SOA

REST

RadiantOne supports any LDAP v3 service.

The JDBC driver for FID allows virtualized views of identity data in SQL.

RadiantOne supports SOAP and WSDL, both XML-based technologies.

RadiantOne supports ADAP (Adaptive Directory Access Protocol), a REST interface for LDAP directory services for web and mobile applications.

Because RadiantOne FID actually tailors data from across data sources to meet the individual needs of each application, it's easy to extend access to new services as they are deployed—without worrying about what is contained in the underlying data stores, or what format the data is in.

Smarter Authorization and Dynamic Group Capabilities

With a virtualized central access point, authentication and authorization are much faster and much easier. Credential checking is delegated to the authoritative underlying data source. RadiantOne FID feeds your policy server user attributes from a variety of identity stores via standard-based protocols, so your policy server can perform richer, attribute-based authorization based on a more complete identity picture.

When basing authorization on group membership, RadiantOne FID can be used to rationalize and aggregate existing groups, flatten nested groups if needed, and even compute dynamic groups with members from multiple sources.


RadiantOne:

- ▲ allows you the flexibility to create groups and define membership across data sources
- ▲ acts as a central access and translation point in cases where all necessary groups exist, but you need a way to re-use them across applications
- ▲ leverages and remaps existing groups by aggregating them into FID. To find groups and members, applications only need to search against FID to check for group membership
- ▲ virtualizes your existing groups and the translation and Distinguished Name (DN) remapping happen automatically

Powerful Integration Layer Supported by HDAP, the RadiantOne Big Data Directory

RadiantOne FID is made up of two main parts, an integration layer and a storage layer. The integration layer is based on virtualization and used for identity aggregation and correlation, group rationalization, and modeling application-specific virtual views.

The storage layer, **HDAP**, is the world's most elastically scalable, cluster-based directory. It is fully **LDAP v3 compliant** with a modern architecture based on Big Data (Hadoop, Lucene) that is much more scalable and fault tolerant than legacy LDAP directories.



The two parts of the RadiantOne FID—an integration layer based on virtualization, and a storage layer (HDAP) based on big data technology—work together to provide an unparalleled federated identity and directory service.

When combined with the RadiantOne virtualization layer, HDAP acts as a highly-scalable cache, delivering high-speed look-ups as well as high-volume storage. As a cache, HDAP stores user attributes and other contextual information joined from across disparate sources—such as SQL, LDAP, web services—and “materialized” for faster access into HDAP. So you get all the richness that comes from accessing a variety of data sources, without having to pay the high cost in time and talent of dynamic and distributed joins.

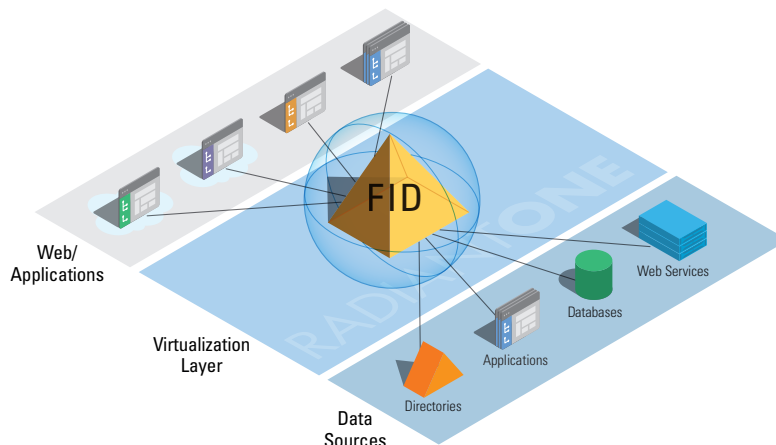
And when demand increases, the cluster-based architecture of RadiantOne is highly available and extremely scalable—in fact, you can bring new nodes online in minutes.

Persistent Caching Means the Directory is Always Up To Date

RadiantOne FID uses virtualization to create an on-demand central directory based on existing identity sources. To do this, it caches and stores a fully indexed and complete image in an LDAP directory, which is automatically updated, creating an always on, always up-to-date persistent cache.

In this way, RadiantOne offers a guaranteed level of performance because the cache is ready without having to be primed with an initial set of queries. And there's no need to worry about how quickly the underlying data sources can respond because they don't need to be queried.

If the directory receives an update for information that's stored in the cache, the underlying data sources will be updated and the persistent cache will also be refreshed automatically. RadiantOne also offers the option of configuring real-time cache refreshes, which will automatically update the persistent cache image when data changes on the backend sources.



The persistent cache functions as a materialized view stored as a full LDAP directory, allowing RadiantOne to deliver data at speeds expected from a standard LDAP directory.

RadiantOne FID: A Complete Federated Identity and Directory Service

A single point of access for all applications: Get one view of all your users and rationalize duplicates. With RadiantOne, your applications have one source for the identity they need—presented in the exact format they require.

Reduced costs and shorter timelines: Updating your identity infrastructure to meet changing business requirements can be costly, time-consuming, and may open you up to security risks. With RadiantOne, you get a reusable identity service that virtualizes all your existing identity information—for quick, risk-free deployments that don't break the bank.

Fast identity retrieval: RadiantOne can be accessed to quickly search/retrieve a unique user profile that will be used for authentication and authorization. Because users are identified against the FID global list of identities, applications do not need to perform lengthy searches of multiple data stores and can authenticate quickly.

Massive scalability: FID scales to millions of users without sacrificing speed. FID enables clients to access data in SQL databases and other non-directory sources at the speed of a directory and also reduces the load on backends by forwarding queries to only the relevant underlying data stores.

Persistent cache: FID creates an on-demand central directory based on existing identity sources and then caches and stores a fully indexed and complete image of it. This image is automatically updated to create an always on, always up-to-date persistent cache that greatly increases speed and reliability. The stored data is available even in the event of a failure in an underlying identity data store.

Cache for global groups: The LDAP-compliant persistent cache can be used for the storage of global groups that include members from multiple data sources being aggregated by FID. It can also store application-specific attributes for new applications so you don't have to extend your existing AD schema.

A 360° view of each user: FID builds a complete user profile, bringing together all attributes, regardless of where or how they're stored. So it's easy to enforce more flexible, finer-grained authorization policies without costly custom coding, and to deliver more personalized services.

Contextual search and management: RadiantOne extracts and aggregates the contextual relationships between identities stored across different silos. With this information, applications and businesses can create complex, context-driven views that can be consumed by authorization policy engines and/or used to build rich customer profiles to enhance user experience.

RadiantOne 7.2 Features

- ▲ Fully encrypt data at rest with HDAP, the RadiantOne Big Data Directory. With new attribute encryption functionality, RadiantOne allows directory administrators to store entries in an encrypted format as an extra layer of security.
- ▲ Enable multi-factor authentication to strengthen security for every app. With plug-in support for multi-factor authentication, RadiantOne makes it quick and easy to leverage third-party security methods, from token cards and certificates to biometrics.
- ▲ Protect key accounts with Privileged ID/Access Management Integration. RadiantOne 7.2 includes an embedded framework for integrating with Privileged Identity/Access Management (PIM/PAM) software. This adds an additional layer of security for key accounts and enables better data governance.
- ▲ Easily provision cloud apps such as Azure AD and Salesforce. With new cloud provisioning capabilities, RadiantOne uses a simple drag and drop interface to build the identity image for each application and then push that customized image to cloud providers.
- ▲ Server management enhancements with improved logging and monitoring. Now it is easier for customers with cluster-based deployments to monitor the state of the services across all nodes, and for those with multiple clusters to monitor the state of inter-cluster replication.

Specifications

See the *Hardware Sizing Guide* to get the best performance with RadiantOne. For supported platforms and system requirements, please see the System Requirements at https://www.radiantlogic.com/website/media/documents/RadiantOne_System_Requirements.pdf