

IAM Reference Architecture

By George B. Dobbs

© 2021 IDPro, George B. Dobbs

To comment on this article, please visit our [GitHub repository](#) and [submit an issue](#).

Table of Contents

ABSTRACT	3
INTRODUCTION	3
TERMINOLOGY	4
BASIC STRUCTURE OF THE MODEL	8
IDENTITY MANAGEMENT	9
RELYING PARTY	9
TRUST FRAMEWORK	10
TRUST ROOT	10
PROVISIONING	10
IDENTITY INFORMATION AUTHORITIES	11
GOVERNANCE	12
CREDENTIAL SERVICES & ENROLLMENT	12
ENROLLMENT	12
CREDENTIAL SERVICES	12
IDENTITY REGISTER	12
SERVICE PROVISIONING AGENT	13
JUST IN TIME PROVISIONING	13
AUDIT REPOSITORY	13
AUTHENTICATION AND SESSIONS	13
AUTHENTICATION	13
SESSIONS	14
AUTHORIZATION	15
LOCAL AUTHORIZATION	16
SHARED AUTHORIZATION	17
AUTHORIZATION MECHANISMS	17
ACCESS GOVERNANCE	17
CONTROL	18
OVERSIGHT	19
RISK CONTEXT	19
EXAMPLE: INFORMATION IN THE REQUEST	21

<i>Boundary control</i>	21
EXAMPLES: HISTORICAL USAGE	21
<i>Usage pattern match</i>	21
<i>Land speed violation</i>	21
EXAMPLE: THIRD PARTY	21
METADATA AND DISCOVERY	22
AUTHOR BIO	23
ACKNOWLEDGMENTS	23
REFERENCES	24

Abstract

This article provides a reference model to organize the presentation of technical details associated with various implementations of identity and access management (IAM) architectural concepts. The model is conceptual, as are the set of abstract components which it provides.

Additional articles will be made available in the IDPro Body of Knowledge that offer more specific technical use-cases based on the abstract concepts in this document.

Introduction

It has been said that all models are wrong, but some are useful.ⁱ This model attempts to find a level of generality that is broadly useful. Too general, and the model becomes untethered to reality and definitely not useful. Too specific, and the model will only work in some cases.

This IAM Reference Architecture leans more towards technical implementation and touches on some of the process, legal, and capability dimensions. This breadth of coverage is intended to give the reader a set of concepts that can be applied when thinking about IAM.

The principle behind this model assumes that the management of identities and access can (mostly) be separated from their use. This concept can apply to distributed systems as well as self-contained systems. So, when you see IAM working together with, say, an application, it may mean that these are separate physical systems. Alternatively, it could mean these parts are separate pieces of software running on a single system.

The main goal of this article is to allow consistent discussion of more specific use-cases by offering a common set of terms and concepts to be used across all IAM architectures.

While the model incorporates guidance from various standards and best practice documents, the primary structure for the model started with the ISO/IEC framing.ⁱⁱ The Unified Modeling Language (UML) detail was removed for simplicity, and the IAM model has been extended so that authorization, governance, and risk-control can be included.

Some of the ISO/IEC names have been changed to reflect more common usage. In some cases, the ISO names have been used in a more expansive way than their original definition.

In an attempt to adopt the most useful terminology, the model has been reviewed in conjunction with the FICAM,ⁱⁱⁱ Internet2,^{iv} NIST SP-800-63 definitions,^v NIST Zero Trust frameworks,^{vi} and with the Identity Stack as presented at Identiverse 2019.^{vii}

The model can be used to support varying levels of system complexity. For example:

- in a Distributed System environment, where the architecture may have a web-hosted application (the RP) that relies on a cloud identity service (the Identity Provider). The RP, in this case, could be a customer-facing application or a workforce-facing application;
- in a Single System model, where a computer's file system (the RP) provides access control based on the user information acquired at login (the IDP). In this case, both the file system and IAM function are encapsulated in an operating system.

Terminology

The terms are defined below. Those with a ✓ mark are the abstract components that comprise the model.

Two of the terms, IDM and Access Management, are used for a conceptual grouping of components. This is to aid understanding.

Item	Definition
Access Control	Various methods to limit access to data, systems, services, resources, locations by a user, a device or thing, or a service.
✓Access Governance (also known as Identity Governance and Administration (IGA))	Access Governance provides oversight and control over access rights implemented in multiple local or shared authorization systems. These rights may be controlled in a variety of ways, starting with the existence and validity of the digital identity. Other controls include various mechanisms such as policies, the mapping of roles, permissions, and identities. The abbreviation used is for Identity Governance and Administration and is commonly used in the commercial sector. This roughly corresponds to the Access Certification section of the first-class component Governance Systems in the FICAM model. IGA is not specifically addressed in the ISO/IEC model.
✓Access Management	The process and techniques used to control access to resources. This capability works together with identity management and the Relying Party to achieve this goal. The model shows access management as a conceptual grouping consisting of the Access

Governance function and the shared authorization component. However, access management impacts local authorization as well (through the governance function).

Assertion	A formal message or token that conveys information about a principal, typically including a level of assurance about an authentication event and sometimes additional attribute information. Sometimes this is called a Security Token.
Assurance Level	A category describing the strength of the identity proofing process and/or the authentication process. See NIST SP.800-63-3 for further information.
✓Attribute Provider	Sometimes the authority for attributes is distinguished from the authority for identities. In this case, the term Attribute Provider is sometimes used. It is a subset or type of an Identity Information Authority.
✓Audit Repository	A component that stores records about all sorts of events that may be useful later to determine if operations are according to policy, support forensic investigations, and allow for pattern analysis. Typically, this is highly controlled to prevent tampering. Audit Repository is the ISO name for this concept and is localized to the IDM. In this model, the term is generalized to indicate a service that supports event records from any part of the ecosystem.
✓Authentication (AuthN)	The act of determining that to a level of assurance, the principal/subject is authentic.
AuthN Assertion	A security token whereby the IDP provides identity and authentication information securely to the RP.
Authorization (AuthZ)	Authorization is how a decision is made at run-time to allow access to a resource. We break this down into two types: shared and local. The FICAM framework includes this as a subcomponent of the Access Management System. AuthZ is not included in the ISO or Internet2 models.
✓Shared AuthZ	Shared authorization is provided by a facility outside of the RP. It is shown here as part of the access management grouping.
✓Local AuthZ	Local authorization is handled by the RP.

Credential	A credential allows for authentication of an entity by binding an identity to an authenticator.
✓Credential Service Provider (CSP)	Following the guidance included in NIST 800-63-3, we include both the enrollment function and credential services together under the name Credential Services Provider.
Credential Services	Credential Services issue or register the subscriber authenticators, deliver the credential for use, and subsequently manage the credentials. We include PKI information for IAM architectures that must include system components that need certificates and private keys. This roughly corresponds to the FICAM component called Credential Management Systems.
Enforcement	The mechanism that ensures an individual cannot perform an action or access a system when prohibited by policy.
Enrollment	Also known as Registration. Enrollment is concerned with the proofing and lifecycle aspects of the principal (or subject). The entity that performs enrollment has sometimes been known as a Registration Authority, but we (following NIST SP.800-63-3) will use the term Credential Service Provider.
Entitlement	The artifact that allows access to a resource by a principal. This artifact is also known as a privilege, access right, permission, or an authorization. An entitlement can be implemented in a variety of ways.
✓Identity Information Authority (IIA)	This represents one or more data sources used by the IDM as the basis for the master set of principal/subject identity records. Each IIA may supply a subset of records and a subset of attributes. Sometimes the IIA is distinguished from the Identity Information Provider or IIP. We use IIA to include the service that actually provides the information as well as the root authority. This corresponds to Identity Information Source in ISO/IEC 24760-2 and Identity Sources in Internet2.
✓Identity Management (IDM)	A set of policies, procedures, technology, and other resources for maintaining identity information. The IDM contains information about principals/subjects, including credentials. It also includes other data such as metadata to enable interoperability with other components. The IDM is shown with a dotted line to indicate that

it is a conceptual grouping of components, not a full-fledged system in itself.

Identity Provider (IDP)	Identity Provider or IDP is a common term. We treat this as a subset of Identity Management. It consists of the service interfaces: AuthN/Assertion, Service Provisioning Agent, Session Management, Discovery Services, and Metadata Management.
✓Identity Register	This is the datastore that contains the enrolled entities and their attributes, including credentials. See the IDM section for elaboration. The terms Directory, Identity Repository, and Attribute Store are sometimes used as synonyms.
✓Metadata Management	The processes and techniques that allow the collection, use, and eventual deletion of control data used by the IDM to recognize and trust the Relying Party. This corresponds to Relying Party data in the Internet2 model.
✓Relying Party (RP)	A component, system, or application that uses the IDP to identify its users. The RP has its own resources and logic. Note that the term 'relying service' is used in the ISO/IEC standards to encompass all types of components that use identity services, including systems, sub-systems, and applications, independent of the domain or operator. We will use the more common Relying Party (or RP). An RP roughly corresponds to the Agency Endpoint in the FICAM model or to Identity Consumers in the Internet2 model.
✓Risk Context (RCTX)	Risk Context consists of additional facts that can be brought to bear to improve the overall security of the ecosystem. Internal or external events and facts can be applied to enable, limit, or terminate access. This is similar to the section Monitors and Sensors under FICAM's Governance Systems and to many of the inputs of the Policy Decision Point in the NIST Special Publication 800-207, a paper on Zero Trust.
Session	A period of time after an authentication event when an RP grants access to resources for the principal/subject. The duration of the session and the mechanism for enforcement vary by implementation.
✓Session Management	A coordinating function provided by an IDP to control sessions of subscribing RPs.

Trust Framework	This component represents the legal, organizational, and technical apparatus that enables trust between the IDM and the RPs.
✓Trust Root	A technical structure that provides the IDP and RP the ability to recognize each other with a high degree of certainty. This is similar to the concept of Trust Anchor (NIST SP.800-63-3), but we allow for a structure that relies on a mutually agreed-upon third party. A trust root derives from the operation of a Trust Framework.

Basic Structure of the Model

The most basic function of the identity system is to provide secure storage of the information about identities and a way for Relying Parties (RPs) to use that data to control access to resources. The following diagram shows the core components of an identity management system (IDM) that supports multiple RPs.

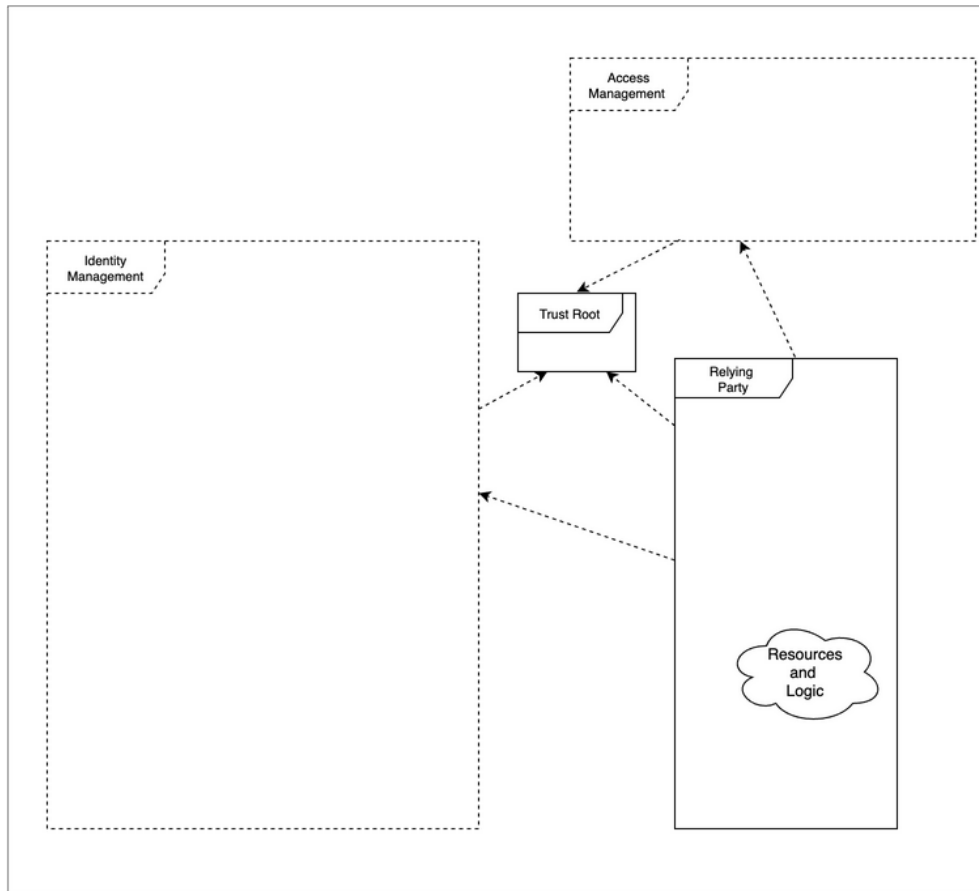


Figure 1: Basic Component Dependencies the IDM supports multiple relying parties. The core components of the IDM are shown. The dotted arrowed lines show dependencies

Identity Management

Identity Management (IDM) is a set of policies, procedures, technology, and other resources for maintaining identity information. In this model, it contains information about principals/subjects, including credentials. It also includes other data such as metadata to enable interoperability with other components. The IDM is shown with a dotted line to indicate that it is a conceptual grouping of components, not a full-fledged system in itself.

Relying Party

The Relying Party (RP) is a component, system, or application that uses the IDM to identify its users. The RP has its own resources and logic. It comes in many forms, all of which use identity services, including systems, sub-systems, and applications, independent of the domain or operator.

Trust Framework

This component represents the legal, organizational, and technical apparatus that enables trust between the IDM and the RPs.

When the IDM and the RP are not in the same organization, the Trust Framework takes on a salient aspect, resulting in multilateral or bilateral agreements. In simple cases, this may be a contract between two parties. In other cases, there may be a multilateral agreement. We will use the term federation loosely to cover both cases. These frameworks are described further in *Laws Governing Identity Systems (v2)*.^{viii}

These agreements, rules, and policies govern how the federation members operate and interact.^{ix} The parties of a federation establish mutual agreement upon an acceptable identity to be used between the parties in a federated relationship (for instance, the level of assurance used) in order to operate well. In addition, the definition and values of attributes of federated identities should be agreed upon. The parties should agree on the security/access policies of federated users between the parties in a federated relationship. For instance, whether there are duties to notify others in the event of security failures.

When IDM and the RP are in the same organization, the agreements may be more tacit.

When the IDM and RP are both built into a single system framework that allows for mutual trust may be completely opaque to the system operator, although the system developer is likely quite aware of it.

Trust Root

A trust root is a technical structure that provides the IDP and RP the ability to recognize each other with a high degree of certainty. This is similar to the concept of Trust Anchor (NIST SP.800-63-3), but we allow for a structure that relies on a mutually agreed-upon third party. A trust root derives from the operation of a Trust Framework. There is a need for a trust root so that the systems can operate without human involvement in every transaction. This may be done through a Public Key Infrastructure (PKI), where the parties agree to trust a common certificate authority that signs the certificates of all parties in the federation. This may be done through a set of several independent certificates that the parties agree to trust.

Provisioning

Provisioning is a term that encompasses the processes and methods that create, modify, and, eventually, delete the identity and profile information used by IT infrastructure and business applications. By these methods, records are created or updated in the identity register and removed from it. Often, provisioning needs to extend to applications to support authorization decisions. This is sometimes known as “downstream provisioning”.

The term “Onboarding” is sometimes used to refer to the sum of the initial provisioning activities in both the identity and access aspects.

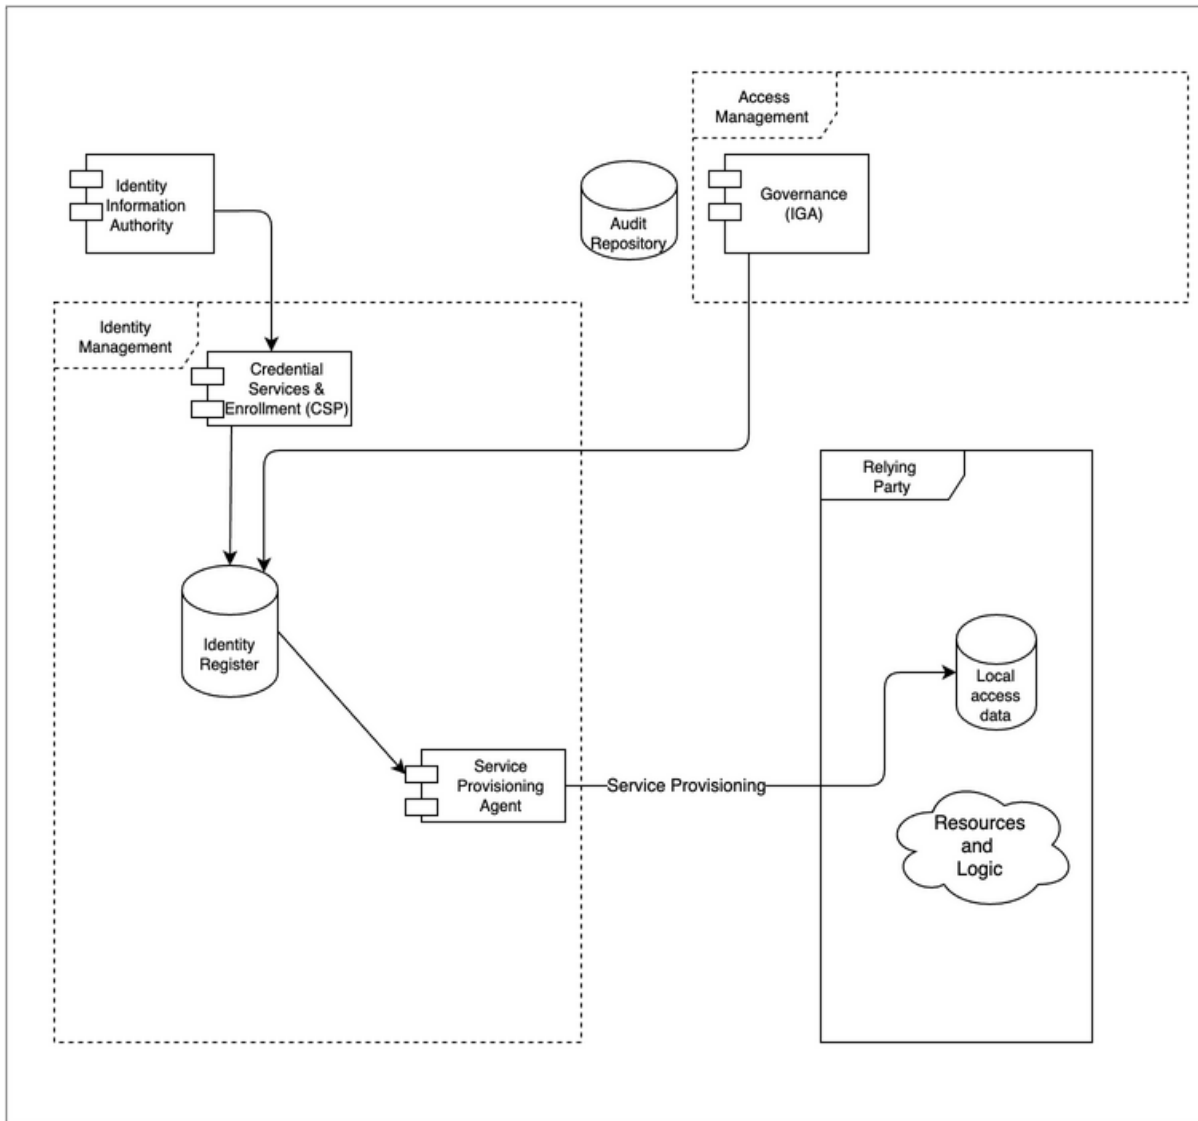


Figure 2: Provisioning: The Identity register receives updates from one or more external sources and administrative actions, passing the information on as needed.

Identity Information Authorities

While it is possible to have an IDM populated without attaching to an external data service, this is typically not the case. Usually, employee or customer data needs to be imported. This can be referred to as upstream provisioning.

Note that the authoritative sources for identity attributes transcend the HR system and may include email, phone, training certification system, etc. In some cases, a company may have more than one HR system.

Governance

The act of provisioning may include certain logic, best modeled as governance. In some cases, the IGA system takes on all the provisioning duties.

Credential Services & Enrollment

This function includes steps needed to originate and activate an identity. It is also concerned with ongoing maintenance such as password reset and key rotation. This function includes administrative activities and self-serve activities.

Enrollment

Also sometimes known as Registration. It involves such activities as proofing, verification or vetting, and recording sponsorship if needed. It also is responsible for the secure delivery of credentials. Enrollment ends when a user formally receives ownership of their digital identity and assumes control and ownership of their account's credentials.

Credential Services

Credential services include the creation and binding of passwords, cryptographic keys, and other authenticators. It is also concerned with ongoing maintenance such as password reset and key rotation. It also is in charge of revoking credentials as needed.

Identity Register

This is the datastore that contains the enrolled entities and their attributes, including credentials. In this model, we use the singular, as if it were one singular database. In practice, designs may store some attributes separately from identities. We also use this term to include the storage related to credentials, although all or some of the credentials may be stored in their own physical repository.

Identity Registers, by their nature, have high availability requirements. Often at the physical level, they contain multiple instances which are synchronized. The Identity Register could be implemented in several ways. Common methods include the use of general-purpose databases, optimized stores such as directories, either physical or virtual.

Importing data does not necessarily mean making a physical copy of data, although it often does. The notion also supports the idea of virtualization - where the import of identity information is done at run-time.

Service Provisioning Agent

Also noted is the function of propagating selected information further into the ecosystem. This typically occurs when an RP needs additional information about the users, e.g., for access control or personalization. The RP makes a copy of the identity data for future use in the application processes. A complete solution will support the full data lifecycle, including creation, update, and eventual deletion of the identity data stored locally.

Just in Time Provisioning

So far, the discussion of the provisioning function has been focused on “admin-time”. However, there are some cases where provisioning occurs at run time.

Not shown here, but sometimes implemented, are provisioning actions that occur on a just-in-time basis. This can happen when additional identity information is passed to an RP in real-time to support a specific application requirement, possibly including identity attributes (See Authentication and Sessions). A similar case involves the RP querying the IDM to acquire attributes (see Authorization later in this document)

Audit Repository

The audit repository is shown to indicate the accumulation of historical event data. To avoid clutter, we assume audit information is written but call that out with arrows in the diagram.

Authentication and Sessions

Authentication

Authentication is the process by which a subject’s credentials are used to verify their identity. The IDP checks and verifies credentials that are presented to it. There are multiple scenarios. Typically, the RP asks the Identity provider to gather the credentials from the user and receives an assessment from the IDP regarding the level of certainty that the user is authentic. Often the assessment (and more information about the user) is delivered to the RP via a security token, which is protected by cryptography. There are several varieties of security tokens.

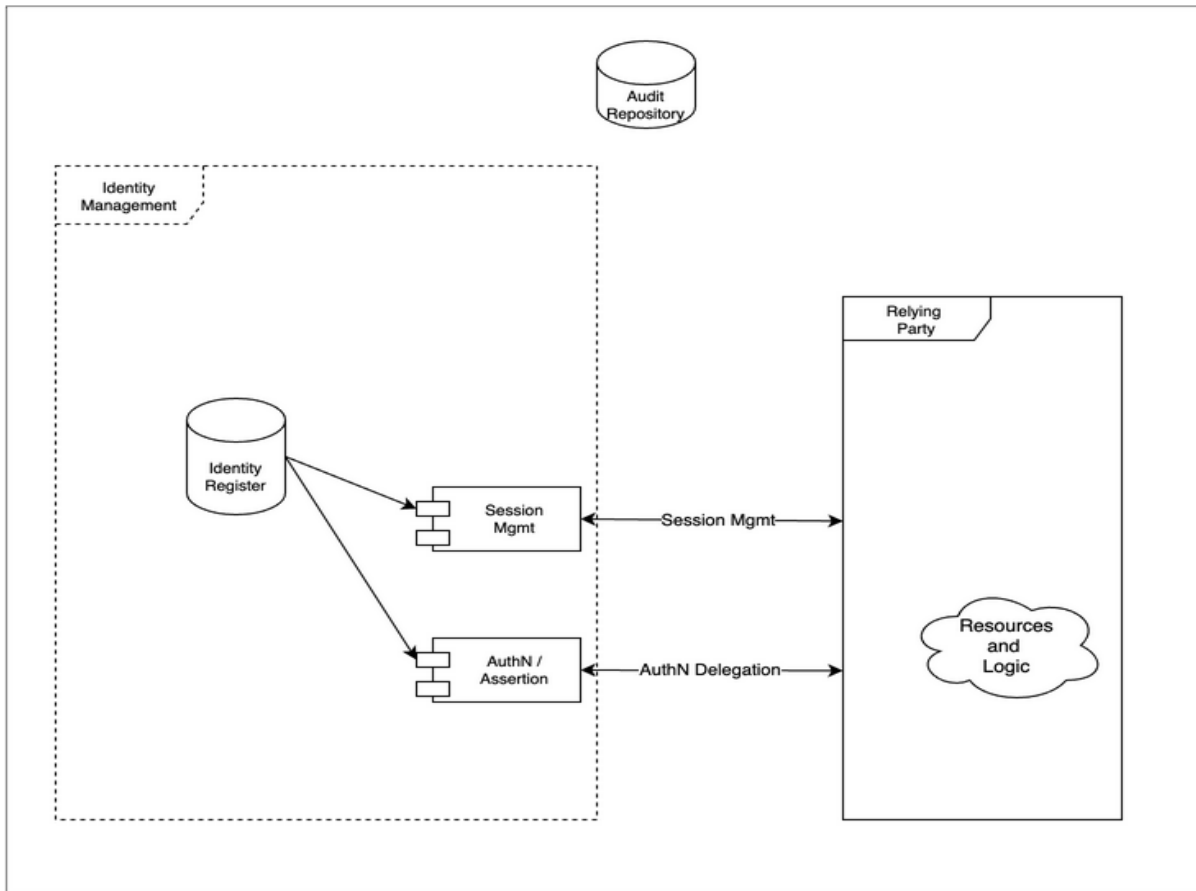


Figure 3: Authentication and Sessions: The Identity Register supports authentication scenarios. The IDP may monitor or participate in the full session lifecycle with the relying parties.

Sessions

A common pattern is to associate the authentication event with the start of a session. The session is mostly the concern of the RP. However, it is sometimes desirable to keep the sessions of several relying parties in synch. For instance, logging out of one session will terminate concurrent sessions. To do this, often the IDP will act to orchestrate sessions termination. In high-security environments, session management must support termination based on real-time identity data, such as when a user's entitlements have been modified.

The existence of a centralized point of view about sessions can be leveraged to support good security practices. For example, if the identity attributes of a user with an active session change and these new values then contravene an access control policy, the session should terminate. If session management becomes aware of a terminated account, it should end any active session that the user has. This could also occur in advanced scenarios that include facts presented by external risk monitors. See Risk Context below.

Sessions also support another important concept: step-up authentication. A session can keep track of the level of assurance of a particular authentication, so when a user requests access to a transaction or application requiring a higher level of identity assurance, the IDP can be prepared to determine the course of action, such as improving the certainty that the user is the right person by asking the user provide additional evidence. For example, maybe the password is good enough to review some information, but to withdraw money, the additional factor of a one-time password from a phone app is required. The detection of the assurance gap and subsequent action will logically be done at the RP, but to avoid a poor user experience in multiple RP scenarios, the step-up needs to be recorded in the session.

Authorization

Authorization models are many and diverse. The diagram illustrates two approaches for authorization: local and shared. As noted below, both approaches are subject to Access Governance.

Both approaches typically use subject attributes to help determine access, although some systems rely on direct enumerations mapping users to resources known as access control lists.

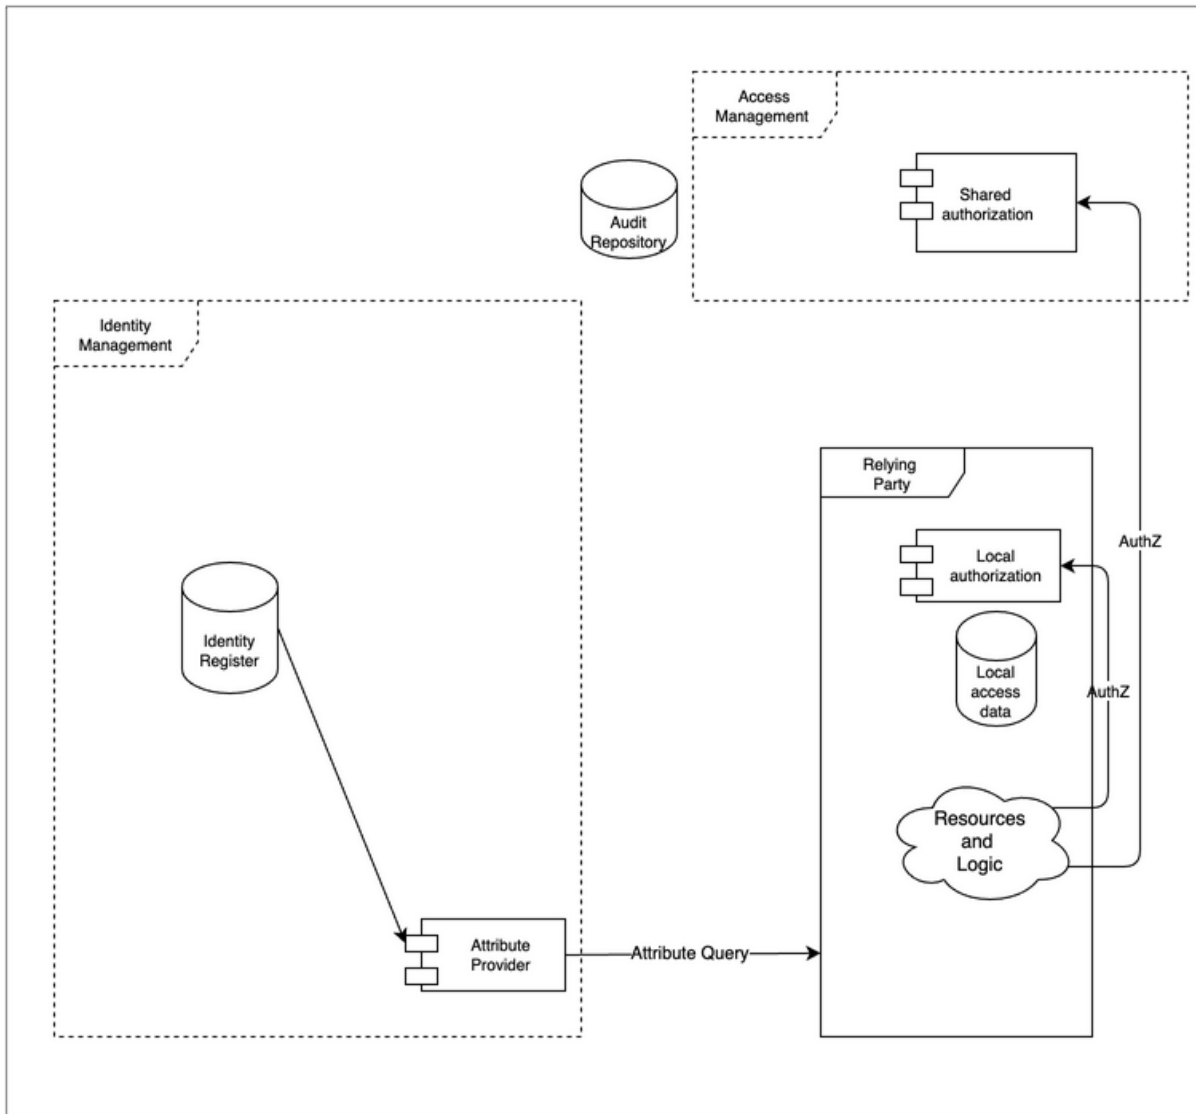


Figure 4: Authorization models: Some RPs perform authorization tasks internally. Sometimes authorization is a shared resource for many RPs.

Local Authorization

Many Relying parties perform authorization tasks internally. Often the fine-grained access control required by a protected resource makes this appealing. For instance, a financial management system may maintain a user's entitlements to specific functionality within the application. In this scenario, the application makes the authorization decision and implements (enforces) the result.

The controlling values may have been provisioned into the local access data store by the Provisioning process described above. Or the values can be acquired at run-time from the IDM as shown by the attribute query, which may provide the user's role or other attributes during the sign-on, perhaps as a value in the security token.

Shared Authorization

Sometimes authorization is a shared resource for many Relying parties. This design can improve the consistency of authorization decisions and supports organizations wishing to include advanced access decisions strategies such as those required by a “Zero Trust” access control approach. Shared authorization systems typically have a consistent approach to policy, such as a standardized policy language. In this scenario, the RP asks the shared authorization function to make the decision but implements (enforces) that itself.

Authorization Mechanisms

In either approach, the access rights may be established, maintained, and revoked in a variety of ways, starting with the existence and validity of the digital identity. Other controls include various mechanisms such as policies, roles, permissions, and identities. Some controls rely on user attributes, including group memberships or roles stored in an Identity Register. Some controls may depend on the properties of the accessed resource or the context of the request, such as time, device, or location.

Each mechanism relies on a particular logical data structure to implement the access control; that data structure becomes the focus of implementers. For instance, in role-based access control, there is some art involved in “Role Management” (defining and managing a useful set of roles) since too many roles become difficult to manage, whereas too few leads to users with access to things they don’t need. Similarly, in the case of policy-based access control, the set of policies (the Policy Rules) needs to be designed, stored, and managed.

Access Governance

Access Governance, also known as Identity Governance and Administration (IGA), provides control over access rights implemented in multiple local or shared authorization systems. This function is often broken into the administration of these rights and the oversight needed to ensure that these rights are in good order over time.

In enterprise systems, Access Governance focuses on managing staff (employee/contractor) entitlements. The concept can also apply to other scenarios, such as when business-to-business delegated administrative rights are required or to in business-to-customer scenarios where authorized third parties such as attorneys are required.

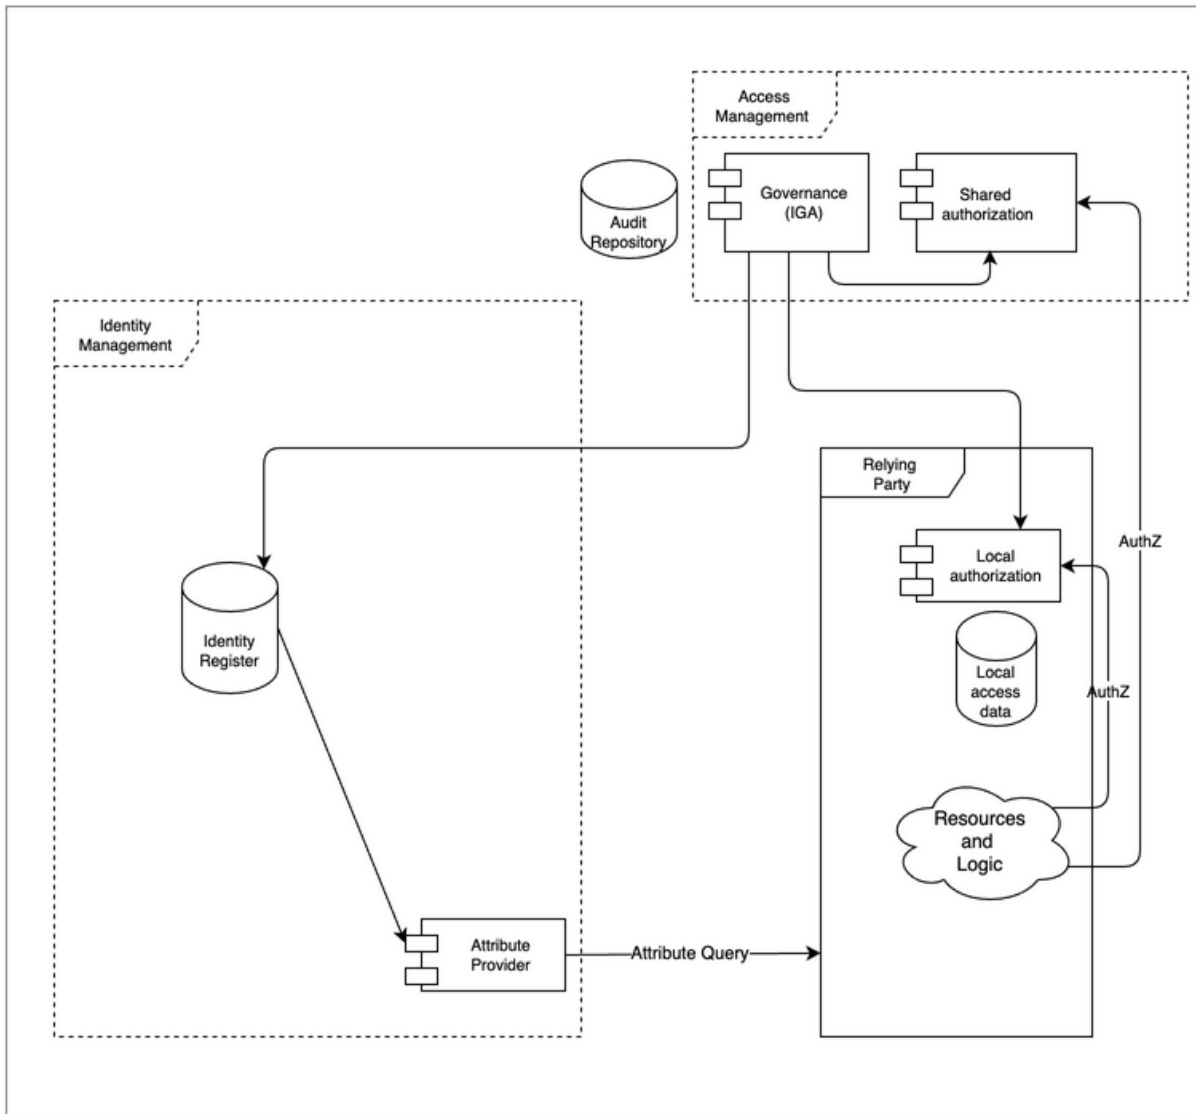


Figure 5: Access Governance provides oversight and control over access rights implemented in many Local authorization systems and, sometimes, in Shared authorization systems.

Control

The controls may also include methods such as procedures and workflows to ensure proper review. Typically, a request for access to resources is passed to one or more approvers and an audit trail is created.

Often deployed to prevent internal fraud is the “segregation of duties” control. The control defines groups of access rights that cannot be held by the same person. This control is best implemented in a location that has visibility to all the implicated access rights, i.e., the IGA system.

Oversight

Typically, governance activities review and potentially modify the data in one or more of the authorization components in order to effect a change in entitlements. Often organizations will have a formal process to review existing entitlements and may require a responsible party to certify or attest that the entitlements are in good order. Additional tools to ensure that IAM policies are effective at enforcing their stated controls include internal and external audits as well as analytic reports.

Risk Context

Risk Context (often abbreviated as RCTX) information can be valuable to improve the security of the relying service. Risk can be judged based on information in the request, information about the history of the user, or assertions/evidence from third parties.

The linkage from the Audit Repository illustrates that the Risk Context may consume the local historical data about events.




Figure 6: Risk Context: It is possible to use risk information in authentication decisions. For instance, if a stolen password is found on the dark web, don't allow login.

External events may be visible to the IDM operator through consortia or vendor packages. In some mutual-support scenarios, it may be possible for the IDM operator to also publish events for the benefit of others, supporting other operators' risk management requirements.

Events need to be delivered into the IDM so that they can selectively be used to modify the behavior of the authentication function. For example, armed with additional event data, the authentication function may request a step-up authentication or even plainly deny access.

In some severe scenarios, attaching the events to the session management function may be desirable so that current sessions can be reviewed and terminated if needed. The

OpenID Shared Signals and Events working group is developing standard ways to deliver these signals. ^x

As shown in the diagram, shared authorization systems may consume risk data as well. For example, an authorization might be denied if the subject's recent activity history is outside of normal bounds, possibly indicating a compromised credential. Logically this could happen with local authorization as well, but this is not shown.

Example: Information in the request

Boundary control

An authentication or authorization decision may be influenced by specific criteria, such as whether a request is coming from a known or unknown network. A more sophisticated version of this attempts to prohibit access from, say, certain countries.

Examples: Historical usage

Usage pattern match

Determine if this request is outside the normal usage patterns for a given individual. The reference to historical usage patterns allows for pattern detection and can help establish a metric for risk for a user, a specific transaction, or in general. Such activity can be called risk profiling.

Land speed violation

Amending the user's request and history with location information makes it possible to identify a likely compromised account because the user can't be in two places at once.

Such examples depend on signals from the local environment, but it is also possible to obtain signals from further afield.

Example: Third party

it is possible to determine commonly used passwords based on postings on the "dark web." Bad actors acquire these in the hope that users will use the same password at other sites. A countermeasure is for the IDM operator to require additional certainty if one of those passwords were presented.

Metadata and Discovery

Metadata refers to control data that allows the IDM and the Relying Parties to interoperate.




Figure 7: Metadata and discovery these two functions are involved with mutual recognition of the IDM and Relying Service.

One example is the registration of public-key certificates to enable mutual authentication. In some scenarios, this information is shared between the parties manually. At run-time for distributed systems, the technical root of trust is needed to validate the security channel (PKI)

Another example points out that configuration information is another form of metadata. OpenID Connect has a list of required, recommended, and optional values that describe a particular implementation aimed at providing a degree of automation during setup.

The metadata may include information that limits the types of interactions and scope of the data that is exchanged. It can also contain security information to allow the

counterparties to authenticate each other. For instance, public key components such as certificates with a common certificate authority may be used.

Discovery refers to protocols that facilitate automation. For instance, OpenID Connect defines a method for RPs to locate an end-point where a user's identity can be verified.^{xi} The concept is more supported by other methods such as SAML.^{xii} A Discovery service can advise where specific data can be accessed and which end-points are maintained to allow an RP to use the identity service.

Author Bio

George Dobbs manages architects at a major insurance company. He is also the chairman of the IDPro Body of Knowledge Committee. One of his interests is modernizing the use of Identity and Access Management techniques used by the firm. He is particularly interested in the area of customer-facing applications, including approaches to fraud prevention in call center and digital contexts. Related to this, he is interested in the evolution of distributed session management – notably distributed session termination. He is a founding member of IDPro and represented his firm in the Identity Ecosystem Steering Group (IDESG). Prior to his current position, he led the development of customer-facing identity for websites at three other insurers. He has led a local identity and access management user group since 2004. Prior to that, he was the chairman of the Network Applications Consortium.

Acknowledgments

The author would like to express gratitude to Ian Glazer, Graham Williamson, and Corey Scholefield for the detailed reviews of early drafts; Jon Lehtinen and Steve Hutchinson for some of the definitions from their unpublished Introduction to Identity Part 3 document; and Bertrand Carlier for his thorough and thoughtful review.

References

-
- ⁱ Wikipedia contributors, "All models are wrong," *Wikipedia, The Free Encyclopedia*, https://en.wikipedia.org/w/index.php?title=All_models_are_wrong&oldid=1032701969 (accessed 23 August 2021).
- ⁱⁱ ISO/IEC 24760-1 Second edition "IT Security and Privacy — A framework for identity management — Part 1: Terminology and concepts," <https://www.iso.org/standard/77582.html> and ISO/IEC 24760-2, 2015 "Information technology — Security techniques — A framework for identity management — Part 2: Reference architecture and requirements," <https://www.iso.org/standard/57915.html> (accessed 28 September 2021).
- ⁱⁱⁱ "FICAM Playbooks – FICAM Architecture – System Component Examples," Identity Assurance and Trusted Access Division in the GSA Office of Government-wide Policy, <https://playbooks.idmanagement.gov/arch/components/> (accessed 23 August 2021).
- ^{iv} Hazelton, Keith "The TAP Reference Architecture (RA)" <https://spaces.at.internet2.edu/pages/viewpage.action?pageId=98306902> (accessed 30 August 2021).
- ^v Grassi, Paul A., Michael E. Garcia, James L. Fenton, "NIST Special Publication 800-63-3 – Digital Identity Guidelines," National Institute of Standards and Technology, U.S. Department of Commerce, June 2017, <https://doi.org/10.6028/NIST.SP.800-63-3>.
- ^{vi} Rose, Scott, Oliver Borchert, Stu Mitchell, Sean Connelly, "NIST Special Publication 800-207 – Zero Trust Architecture," National Institute of Standards and Technology, U.S. Department of Commerce, August 2020, <https://doi.org/10.6028/NIST.SP.800-207>.
- ^{vii} Hutchinson, Steve, "Introduction to Identity Part 2 - June 25," Identiverse 2019, recording starting minute 27:39, <https://www.youtube.com/watch?v=zxKRUXmTLJs&list=PLpKq7xRilHaTDwAqpIU1UYpKZY03tfTMf&index=8>.
- ^{viii} Smedinghoff T. J., (2021) "Laws Governing Identity Systems (v2)," *IDPro Body of Knowledge* 1(5). <https://bok.idpro.org/article/id/8/>.
- ^{ix} Temoshak, David, Christine Abruzzi, "NISTIR 8149 - Developing Trust Frameworks to Support Identity Federations," National Institute of Standards and Technology, U.S. Department of Commerce, January 2018, <https://doi.org/10.6028/NIST.IR.8149>.
- ^x "Shared Signals and Events WG" <https://openid.net/wg/sse/> (Accessed 30 August 2021)
- ^{xi} Sakimura, N., J. Bradley, M. Jones, E., Jay, "OpenID Connect Discovery 1.0 incorporating errata set 1," OpenID Foundation, 8 November 2014, https://openid.net/specs/openid-connect-discovery-1_0.html (accessed 23 August 2021).
- ^{xii} Widdowson, Rod, Scott Cantor, "Identity Provider Discovery Service Protocol and Profile," OASIS Committee Specification 01, 27 March 2008, <https://docs.oasis-open.org/security/saml/Post2.0/sstc-saml-idp-discovery.pdf> (accessed 23 August 2021).