

eduPerson FAQ



Note: as of October 2018, stewardship and maintenance of eduPerson has been transferred from Internet2 to REFEDs (<https://refeds.org/>).

- What is eduPerson and how do organizations use it?
- What is the aim of defining eduPerson attributes?
- What is the relation of eduPerson to inetOrgPerson and other standards?
- How can use of eduPerson attributes protect users' privacy?
- Are eduPerson attributes intended or actually used (consumed) as LDAP attributes, or as attributes in SAML assertions?
- Are there canonical values of eduPersonAssurance that are or should be recognized by service providers?
- If eduPerson directory attributes are multi-valued, can one assume services will be able to properly consume corresponding multi-valued SAML attributes?
- Why does eduPerson include the eduPersonOrcid attribute but not eduPersonResearcherId? Won't this lead to new attributes for every kind of identifier?
- How are Identity Providers and Service Providers leveraging the eduPersonOrcid attribute?
- Is eduPersonTargetedID really an attribute in the usual sense? What is the relationship between it and SAML? Is it relevant in a context other than SAML?
- Why are there two seemingly similar identifiers eduPersonPrincipalName and eduPersonUniqueID?
- Under what circumstances would one use eduPersonEntitlement rather than LDAP group membership to indicate specific access privileges?

What is eduPerson and how do organizations use it?

eduPerson is an attribute schema that includes bindings to a Lightweight Directory Access Protocol (LDAP) schema and to SAML. It is designed to include and standardize widely-used person and organizational attributes in higher education and research that are not duplicated in other widely used objects such as inetOrgPerson.

What is the aim of defining eduPerson attributes?

The chief aim is to align practice across organizations around a common set of attributes for information specific to higher education and to IAM (Identity and Access Management) best practices promulgated by the [Internet2 Trust and Identity](#) community work.

What is the relation of eduPerson to inetOrgPerson and other standards?

eduPerson extends and profiles existing schema standards to avoid reinvention while adding attributes specific to and useful for higher education and research.

How can use of eduPerson attributes protect users' privacy?

(a) Services can rely on generic attributes that do not identify a specific person, such as eduPersonAffiliation or eduPersonEntitlement. (b) Services that need to maintain internal record (to manage preferences say) can use eduPersonTargetedID, providing a unique ID but not readily correlated to PII or to activity in other services (more in response to the question on eduPersonTargetedID).

Are eduPerson attributes intended or actually used (consumed) as LDAP attributes, or as attributes in SAML assertions?

That's a context-specific question. eduPerson may be used in both contexts (and in future ones). If you have no LDAP applications, then you may not find it useful or necessary to actually store attributes directly in LDAP and it may be simpler to just construct them as needed from within SAML software or in a database. However, if you do have the need or the ability to store them in LDAP, it will generally be easier to produce them in SAML too. The more your IDM infrastructure does, the less your SAML software has to do to compensate.

Are there canonical values of eduPersonAssurance that are or should be recognized by service providers?

The values of this attribute are generally specific to a community and there are none defined by the eduPerson specification (just as there are no values defined for eduPersonEntitlement). InCommon, for example, has defined assurance profiles and, more recently, an [MFA profile](#) that include values suitable for use with this attribute.

If eduPerson directory attributes are multi-valued, can one assume services will be able to properly consume corresponding multi-valued SAML attributes?

Attributes designed for searching, such as givenName, sn, or mail, are often not handled correctly if multiple values are supplied in a federated context. So in general, no, one can't assume that (and it may be necessary to release an alternative single-valued attribute to some services), but it is a good practice to report such bugs when they are identified. In terms of correctness, any multi-valued attribute is expected to be handled in that fashion in any context.

Why does eduPerson include the eduPersonOrcid attribute but not eduPersonResearcherId? Won't this lead to new attributes for every kind of identifier?

Yes, it will, very deliberately. Combining multiple types of data into a single attribute precludes use cases in which only a subset of that data may be relevant, unless the data is encoded in a way that allows the different types of data to be recovered. That in turn adds extra work to a consumer of the data.

Attributes are not "expensive" to create, and the more precise an attribute definition can be made, the more intelligent software can be when dealing with them. We should expect to see additional attributes created for any kind of identifier that gains adoption by the community.

How are Identity Providers and Service Providers leveraging the eduPersonOrcid attribute?

Some interesting use cases are found on the ORCID website here <http://orcid.org/organizations/institutions/usecases>

Is eduPersonTargetedID really an attribute in the usual sense? What is the relationship between it and SAML? Is it relevant in a context other than SAML?

The eduPersonTargetedID is an unusual attribute that does not map easily to an LDAP representation in the way that every other attribute in the schema does. Because its value is intended to be different for every "client application", it cannot easily be maintained in a typical LDAP directory and is not expected to be. That indeed makes it unusual.

The relationship with SAML has to do with the history of Shibboleth and its use of eduPerson as the "recommended" attribute vocabulary for the higher education community's use of SAML. The concept of a "directed" (pair-wise) identifier emerged from the work done on federated identity when SAML 2.0 was being standardized, and because Shibboleth was originally a SAML 1.1-based system, the Shibboleth community decided to develop an eduPerson attribute that had the characteristics of a concept in SAML 2.0 called a "persistent name identifier". That attribute was eduPersonTargetedID.

As the use of SAML 2.0 supplanted SAML 1.1, the need for an attribute distinct from the already-defined SAML NameID Format of "urn:oasis:names:tc:SAML:2.0:nameid-format:persistent" has waned, and the need for something called eduPersonTargetedID is now somewhat historical.

As to other contexts, that is unclear. The definition of eduPersonTargetedID is suitably generalized to be compatible with the SAML concepts it was copying, but may or may not be suitable as a way of describing similar concepts in other standards. This is an open question.

Why are there two seemingly similar identifiers eduPersonPrincipalName and eduPersonUniqueID?

eduPersonPrincipalName has the format of a name-based identifier, scoped to the domain of the Identity Provider; it will seem familiar to many users, but because it is name-based, the ePPN assigned to a given person is subject to change, which is a problem for services that maintain a user profile or record. In contrast, eduPersonUniqueID is intended never to change; it is more suitable as a permanent identifier of a specific user.

Under what circumstances would one use eduPersonEntitlement rather than LDAP group membership to indicate specific access privileges?

Answer to be developed. MACE-Dir community members are welcome to take a crack at answering this

See also

[MACE-Dir Working Group Space](#)

Practices in Directory Groups (2002) <http://doi.org/10.26869/TI.23.1>