OpenID

Enabling Large-Scale Multi-Party Federations with OpenID Connect

January 24, 2020

Michael B. Jones

Identity Standards Architect – Microsoft

with significant contributions by Roland Hedberg

Working Together

OpenID

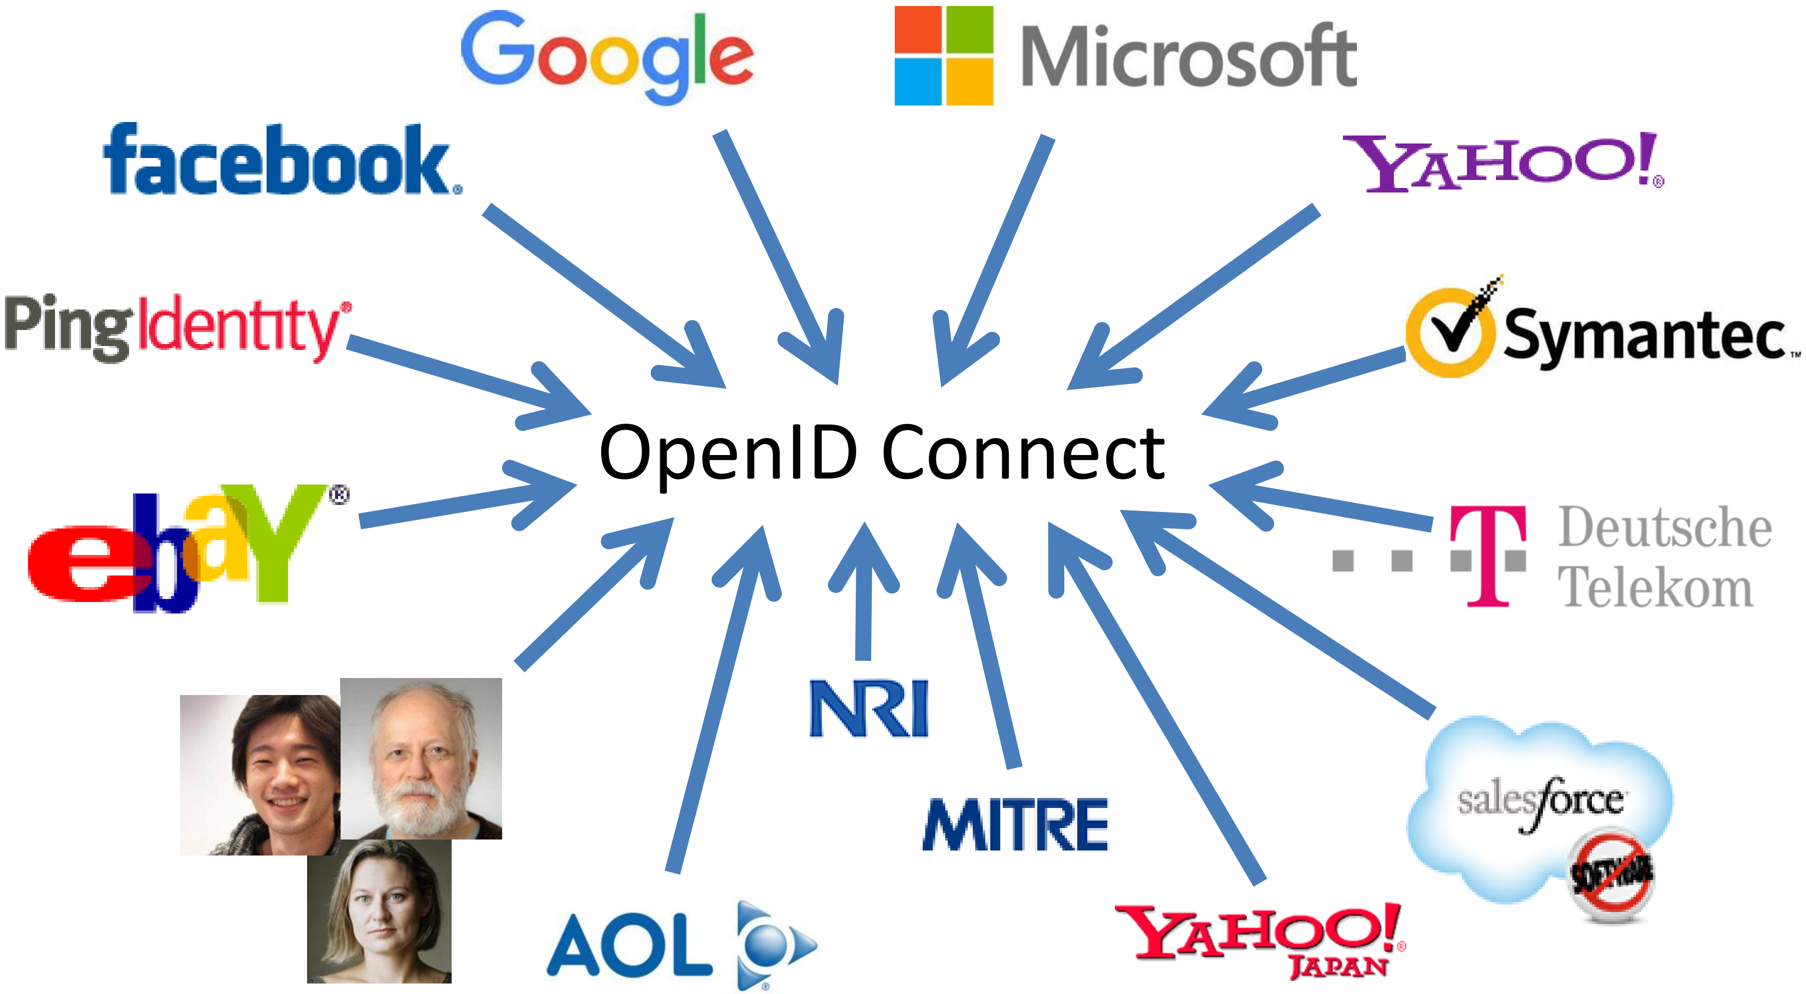

OpenID Connect

You’re Probably Already Using OpenID Connect!

- If you log in with Android, Apple, AOL, Deutsche Telekom, eBay, Gigya, GSMA, Google, Janrain, KDDI, Microsoft, NEC, NTT, Okta, PayPal, Salesforce, Softbank, Symantec, Verizon, or Yahoo! Japan, you’re already using OpenID Connect
  - Many other sites and apps large and small also use OpenID Connect
- Not a consumer brand
  - Rather, very widely deployed, simple, secure identity infrastructure

What has OpenID Connect Achieved? OpenID

- Widely used for
  - Web apps
  - Native apps
  - Enterprise apps
  - Cloud apps
  - Financial apps
- Over 100 certified deployments at https://openid.net/certification
- Available for essentially all modern development platforms
- Increasingly preferred by developers over SAML

Numerous Awards

- OpenID Connect won 2012 European Identity Award for Best Innovation/New Standard
  - http://openid.net/2012/04/18/openid-connect-wins-2012-european-identity-and-cloud-award/
- OAuth 2.0 won in 2013
- JSON Web Token (JWT) & JOSE won in 2014
- OpenID Certification program won 2018 Identity Innovation Award
- OpenID Certification program won 2018 European Identity Award

A Microsoft Perspective

- At Identiverse in 2019, Alex Simons – Microsoft's VP of Identity Program Management, reported:
  - ***Over 95% of all Azure Active Directory (AAD) authentications use OpenID Connect***
  - We're doing over 20 billion authentications per day

But what about Research and Education?

- Research and Education sector has numerous large-scale identity federations
  - – Many national and regional federations
    - Such as SWAMID in Sweden and InCommon in the United States
    - Some have thousands of sites
  - – Inter-federations among dozes of federations, such as eduGAIN
- These allow identities from any federation member to be used at relying parties from any federation member
  - – For instance, using a University of Washington account at CERN
- BUT… today these are nearly all based on SAML 2
  - – Mostly using Shibboleth software

Significant OpenID Connect Interest in Research and Education Sector

- Research and Education OpenID Working Group
  - https://openid.net/wg/rande/
  - Profiling OpenID Connect for use in R&E applications
  - Including mapping EduPerson schema to OpenID Connect claims
- Multiple OpenID Connect implementations for R&E world:
  - University of Chicago Shibboleth Plug-in was an early implementation
  - GÉANT OpenID Connect Shibboleth Plug-In
    - Now supported and distributed with Shibboleth software

Federation using OpenID Connect

- OpenID Connect Federation specification
  - https://openid.net/specs/openid-connect-federation-1_0.html
  - Enables establishment and maintenance of multi-lateral federations using OpenID Connect
- Incorporates lessons learned from SAML-based federations
  - Defines hierarchical JSON-based metadata structures for federation participants
- Second Implementer's Draft just approved
- Rest of this presentation describes how Federation is being achieved with OpenID Connect

Establishing Trust within a Federation OpenID

- How do a Relying Party and an Identity Provider know that they're in the same federation?
  - Important for trust, liability, accountability, and reliability
- Shibboleth/SAML approach:
  - Federation Operator polls participants for their metadata, concatenates it into a huge flat file, and distributes it to all nightly
  - In production use, but brittle and not scalable
    - SAML world developing Metadata Query protocol to try to move away from this
- New OpenID Connect Federation approach:
  - Hierarchical metadata, where organizations publish metadata about themselves and Federation Operators publish statements about orgs
  - Scalable, maintainable

Use of Hierarchical Metadata

OpenID

- Each leaf member publishes self-signed metadata about itself
  - Relying Parties
  - Identity Providers
- Organizations publish signed metadata about the members that belong to them
- Federation operators publish signed metadata about orgs
- Inter-federations publish signed metadata about federations
- Hierarchical metadata is an online graph data structure

Trust Chains

- Participants follow metadata trust chains from leaves up to common roots, verifying signatures
- Both participants are members of a federation if a common trusted root is found
- Participants can be members of multiple federations

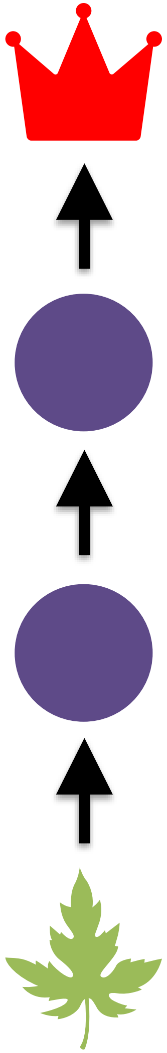

Metadata Representation

- Each metadata statement is a signed JSON Web Token (JWT)
  - These are called Entity Statements
- They make statements about
  - The entity itself
  - Keys used by the entity
  - Policies of the entity
  - Other entities up the trust chain that they are willing to trust
    - This is how trust chains can be followed to federation roots

Example Entity Statement

```
{
 "iss": "https://feide.no",
 "sub": "https://ntnu.no",
 "iat": 1516239022,
 "exp": 1516298022,
 "jti": "7l2lncFdY6SlhNia",
 "metadata_policy": {
    "openid_provider": {
      "issuer": {"value": "https://ntnu.no"},
      "organization_name": {"value": "NTNU"},
      "id_token_signing_alg_values_supported":
        {"subset_of": ["RS256", "RS384", "RS512"]},
   }
 },
  "jwks": {
    "keys": [
      {
        "e": "AQAB",
        "kid": "key1",
        "kty": "RSA",
        "n": "pnXBOusEANuug6ewezb9J_...",
        "use": "sig"
      }
    ]
  },
  "authority_hints": [
    "https://edugain.org/federation"
  ]
}
```

Collecting a Trust Chain

Self-signed entity statement. First entity in trust chain.

1. From the claim *authority_hints*, pick superior entity.
2. Grab superior’s self-signed entity statement (using .well-known)
3. Request superior’s view of subordinate (federation API). Add to the trust chain.
4. GOTO 1

Repeat until superior is a trusted trust anchor

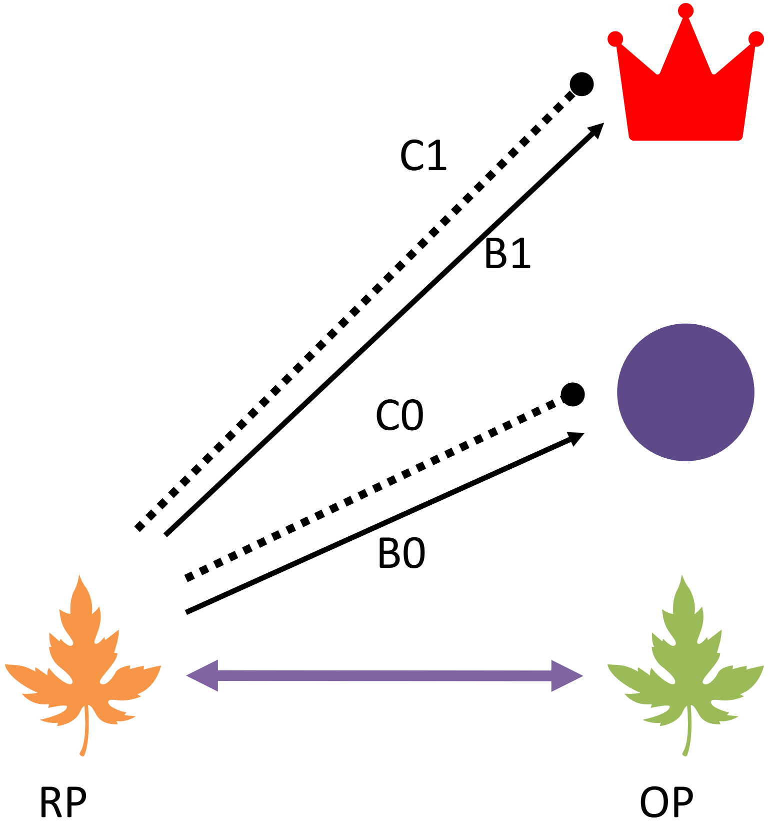

SAML vs. OpenID Connect

SAML

- Appearing in a metadata file means you are part of a federation

OpenID Connect

- Entities with trust chains up to the same trust anchor belong to the same federation

SAML vs. OpenID Connect

SAML

- An entity's complete metadata must be accepted by the federation operator for the entity to be allowed into the federation

OpenID Connect

- The federation operator sets the boundaries of what is acceptable

Praise for OpenID Connect Approach

Shibboleth author Scott Cantor publicly said at a federation conference:

- *“Given all my experience, if I were to redo the metadata handling today, I would do it along the lines in the OpenID Connect Federation specification.”*

Policy Language for Entity Statements

- subset_of
- one_of
- superset_of
- add
- value
- default
- essential
- Path length/name restrictions
- Trust/certification marks

Applying Metadata Policies

- Policies applied top-down from root to leaves of trust chain
- Policies higher in the chain override those lower in the chain

- For instance, a Federation Operator might specify that only a particular set of signing algorithms may be used
  - Policies are applied to all entities in the federation

SAML vs. OpenID Connect

SAML

- It is rare that an entity belongs to more then one federation. Given eduGAIN, it is actually recommended that an entity only belong to one.

OpenID Connect

- There is no drawback to belonging to multiple federations

SAML vs. OpenID Connect

SAML

- The is no metadata negotiation

OpenID Connect

- The RP proposes and the OP decides, subject to applicable policies from the trust chain

Client Registration Methods

- Automatic
  - The client preforms no client registration. Instead, it sends an authorization request with *client_id* == *entity_id* and client authentication method *private_key_jwt.*
  - The OP fetches the RP's self-signed entity statement.
- Explicit
  - The client performs dynamic client registration. The OP responds with an entity statement about the RP with metadata policy.
  - The RP provides the OP with its self-signed entity statement in the body of the client registration request.

OpenID Connect Federation Past

- Second Implementer’s Draft Approved in January 2020
- Spec refined from discussions at multiple federation events
  - NORDUnet 2017
  - SURFnet 2018
  - TNC/REFEDS 2019
  - Internet2/REFEDS 2019
  - Now OpenID Japan Workshop 2020
- Hackathon with interop among multiple implementations
  - Internet2/REFEDS 2019

OpenID Connect Federation Future

OpenID

- OpenID Foundation holding three interop events in 2020
  - Much like five interops were held for OpenID Connect
  - Interop results will be used to improve the specification
  - Contact Roland Hedberg roland@catalogix.se to participate
  - Join OpenID Federation Interop mailing list
    - https://groups.google.com/forum/#!forum/openid-federation-interop
- It’s time for feedback from developers and early deployers
  - ***Will you be one?***
  - Please read (and implement!) the spec and give us your feedback!

OpenID Connect Federation Resources

- OpenID Connect Federation Specification
  - https://openid.net/specs/openid-connect-federation-1_0.html
- OpenID Connect Page
  - https://openid.net/connect/
- OpenID Connect Working Group Mailing List
  - https://lists.openid.net/mailman/listinfo/openid-specs-ab
- OpenID Blog
  - https://openid.net/
- Mike Jones' Blog
  - https://self-issued.info/

Open Conversation

- Where would you like to see OpenID Connect Federation used?
- What would you like the working group to know or do?