Enterprise Identity Federation: Mitigating Risks and Liabilities

Author: Guy Huntington, President, Huntington Ventures Ltd.
Date: September 2017

Table of Contents

Introduction

There is a growing trend for enterprises to use identity federation with their customers and between enterprises offering single sign on, authorization, and seamless integration with many social media apps.

When I go into enterprises deploying identity federation, I frequently tell my teams that I have four letters stenciled across my forehead: R I S K. To mitigate risk from federation requires an enterprise view of the risk from Legal, Governance, Business, and IT.

It has been my experience that Business and Legal don't fully understand the risks involved, and instead trust their IT department to "handle it." This is why I have written this paper. It is aimed at Business, Legal, and IT leaders within an enterprise that is either embarking on identity federation and/or expanding their use of it. By reading this paper, you will learn the types of things your enterprise should be doing to mitigate federation risks and potential liabilities.

Governance – Legal Agreement

The place to start is the governance of any federation arrangement. A federation is first and foremost a legal agreement between two or many federation partners. The legal agreement needs to: specify the risks; clearly explain risk mitigation business processes, technical standards, and processes that will be adhered to; and where possible, assign legal liabilities.

Federation Legal Agreement

The following areas need to be addressed in a federation agreement:

Assurance	Identity
	Credential
Identity Management	Identity creation
	Identity modification
	Identity termination
	Archival

Federation Protocols	The exact standards to be used (in the following Federation Risk section, I will explain why this is VERY important)
Environments	Production
	Other environments
	Endpoints used per environment
	The business processes used to migrate between environments
	Service Level Agreements
	Performance specifications
Security	Encryption used
	Digital Certificate authorities used
	Digital Certificate revocation processes
	Digital Certificate update processes
	Denial of Service and Distributed Denial of Service attack mitigation measures
	Third party penetration tests
	Describe each attack vector and describe how the risk will be either mitigated or where the liability is assigned
Monitoring, Incident Management, and Reporting	Synthetic transactions once per minute (or whatever the agreed upon time frame)
	Incident management escalation processes
	Integrated incident management ticketing systems
	Types of reports to be generated and automatically exchanged between partners
Privacy Breach	Business processes that occur when a privacy breach is detected
	Technical processes for resolving a privacy breach
Disaster Recovery	Business and technical processes used when one of the federation parties goes down
Identity Data Standards	Identity attributes to be exchanged
	Data standards for each attribute
	Format for the subject identifier
	Ensure that subject identifier never changes for the identity during their lifetime
Session Management	Define how session management will be done
Test Accounts	Define all test accounts used in all environments

Risks and Liabilities

The risks and potential liabilities of each of the above areas are examined in further detail in the following subsections:

Assurance

Assurance is a measure of risk. There are three types of assurance:

- Identity
- Credential
- Session

Identity Assurance

How does an enterprise know that a person claiming to be Jane Doe is actually Jane Doe?

Governments and enterprises around the world typically use an agreed upon rating scale from 1 to 4 where 1 indicates low assurance and 4 indicates high assurance. For each rating level, the underlying identity documentation type and the verification processes are clearly specified.

Depending on whether your enterprise's role as either the identity provider (OpenID provider) or the enterprise relying upon it (Relying Party), there will be varying risks associated with how the identity was originally "proofed". Thus the legal contract needs to specify exactly what the identity assurance is.

Credential Assurance

How will your enterprise know that an electronic transaction with Jane Doe is actually being carried out with Jane Doe on the other end of the internet connection? Similar to identity assurance, governments and enterprises around the world agree on levels of risk for the credentials used. Credentials are either: something you know, something you have, or something you are. Various combinations of these factors will have a higher assurance than those that do not.

If you are the Relying Party in a federation agreement, you are dependent upon the credential assurance the OpenID Provider uses. If the credential assurance is low, then you might have to either accept the risk that the entity on the other end might not be who they claim to be or, you might try to assign some of the legal liability to the OpenID Provider.

Session Assurance

For a given user's session, there is an identity assurance level AND the user might have more than one types of credentials they can use to log on with. If your enterprise is a Relying Party, you may or may not care about assuring yourself that the identity on the other end is who they claim to be (identity and credential assurance) for a given session. Session assurance is a measure of the lower of the two values.

As an example, Jane Doe has an identity assurance level of 3. She also has two types of authentication: one at a credential assurance level of 2 and the other at level 3. She might choose to use the weaker form of credential assurance to log on with. Your application/service might require a session assurance level of 3. Since Jane is at a session assurance level of 2, she would be stopped from accessing your service until she uses her stronger credential assurance.

So, in summary, identity, credential, and session assurance levels need to be specified in the federation legal agreement.

Identity Management

There are three general types of risks associated with identity federation regarding identity management:

1. How the identity is created, modified, terminated, and archived in the authoritative source (usually but not always the OpenID Provider)
2. How fast the change occurs between the OpenID Provider and the Relying Party
3. How fast the Relying Party implements the change

The first type of risk (creating, modifying, terminating and archiving an identity) is something that occurs in the business and technical processes used by the OpenID Provider and/or the authorization source.

The second type of risk (notification of a change between the federation partners) is determined by how the federation is configured as well as the time when the authoritative source enters in the information to be sent to the Relying Party.

The third type of risk (Relying Party implementing the change) is determined on the Relying Party's internal systems and the time it takes to update the systems within.

Each one of the above has potential risks and liabilities for the federation partners. Depending upon the degree of risk, the federation legal agreement needs to specify the risk and then assign legal liability to one or more of the federation partners wherever possible.

Federation Protocols

The two most commonly used federation protocols, OpenID Connect (OIDC) and Secure Assertion Markup Language (SAML) are not one "thing". Instead, both use a variety of different protocols upon which they are built. To help Business individuals understand what the implications are, let us use the analogy of Lego building blocks and OpenID Connect as the example.

Imagine that you are constructing a strong structure using the blocks. At the bottom of the structure you use six large blocks. This is the foundational piece. However, there are many different ways to assemble each of the foundational blocks, i.e. each of the large blocks is composed of smaller blocks that can be assembled differently to make the larger foundational one.

On top of the large foundational blocks, you then assemble five mid tier blocks. Like the foundational ones supporting them, they too have different ways for each of them to be configured, i.e. each mid-tier block is made up of smaller blocks.

On top of this you then assemble five top tier blocks. Like the foundational and mid-tier blocks supporting them, they too have different ways they can be configured.

The overall "strength" of the structure is dependent upon each of the building blocks and how they are configured. If the ones at the bottom have been poorly configured, then the structure might fall down.

This analogy can be applied to the picture below:

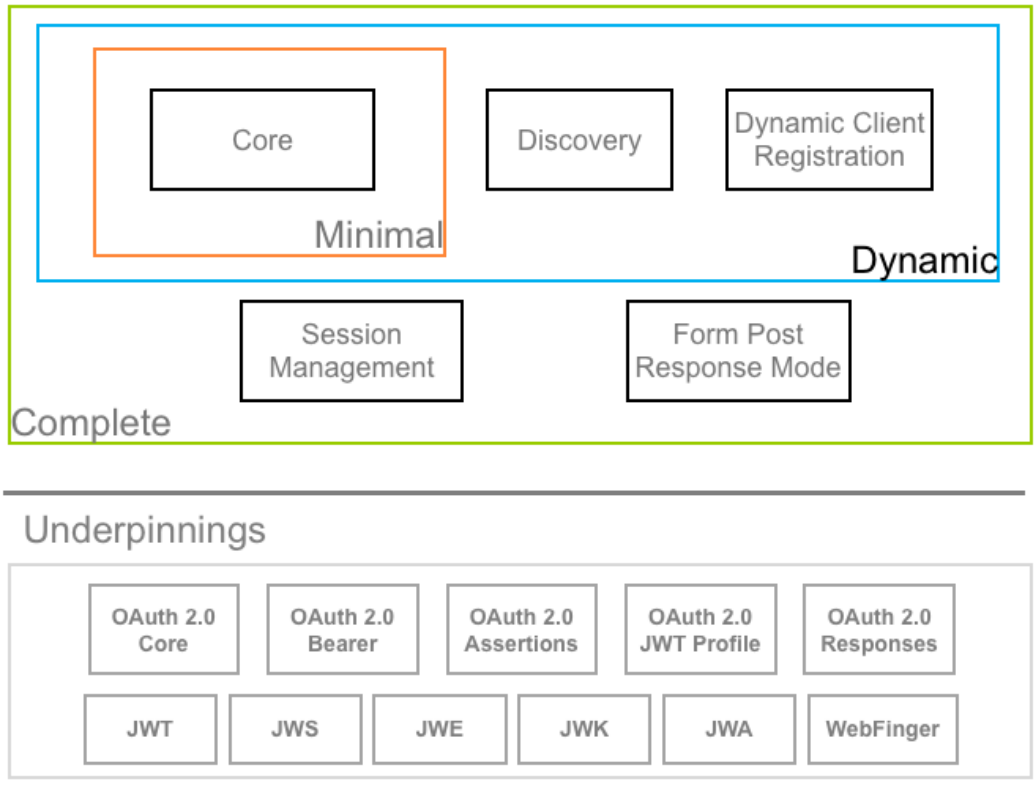

- Foundational blocks are the JSON components and WebFinger
- Mid-tier blocks are the OAuth components
- OIDC is the top tier blocks comprised of Core, Discovery, Dynamic Client Registration, Session Management, and Form Post Response Mode

When I work with enterprise teams, I usually print all this information out and create one to two very large binders. I then tell the teams that this is where most, but not all, of the risk lies. Defining the exact way the various protocols will be configured is critical to the overall security of the deployment (see the security section for more details). Therefore, the legal agreement must specify the EXACT configuration for each of the components.

Environments

A typical federation deployment uses at least two environments, i.e. Production and UAT (User Acceptance Test). Sometimes more are used. Regardless, in federation one partner, the identity provider, is usually hosting the environments and the other parties, i.e. the Relying Parties, are integrating with these environments.

The environments are internet facing, which means increased external attack vectors. Federation involves end-points (i.e. URL addresses) where the federation interaction occurs.

The federation partners need to have a very clear understanding of the business processes used to migrate into these environments, the service level agreements, hours of support, performance specifications, key contact information, etc. As well, when change management is done on the host environment, the other federation partners need to be notified in advance and then automatic testing needs to be done to ensure everything still works as expected. When an emergency patch is done, the legal contract must specify the exact notification and testing processes.

All of the above needs to be specified in an appendix to the legal federation agreement. This way there are no unexpected surprises.

Security

One of the first things I do when working with an enterprise doing federation is to get my team to assemble an attack vector spreadsheet outlining all the potential attack vectors with their likelihood of happening, the risks, and suggested mitigation strategies. Many of these attack vectors need to be in an appendix to the legal agreement stating exactly how the federation partners will mitigate their risk and, wherever possible, assign potential liability. Thus the Business and Legal teams need to go through a risk spreadsheet to understand each risk and then sign off internally and also with their federation partners.

Another thing to consider is the use of digital certificates. A typical federation transaction encrypts the internet to create a secure tunnel, the tokens are digitally signed, and the actual information within the token is also encrypted. For each single federation transaction, a digital certificate revocation check is usually done, which involves going out from the federation partner to see if the digital certificate issued by a Certificate Authority (CA) is still valid.

There is risk and potential liability associated with digital certificates. The legal agreement needs to specify what CAs are acceptable, the type of encryption used, revocation checking and protocols used for it, as well as the business and technical processes used when a digital certificate is about to expire and a replacement one is then installed.

An area of growing concern is Denial of Service and Distributed Denial of Service attacks. Over the last year:

- The size of these attacks has increased to several hundred gigabits per second
- The cost for renting out a botnet to make these attacks ranges from $5-40/hour

Recently a teen was arrested after making over $300,000 from conducting these types of attacks as well as renting out their botnet. Enterprises should review with their Internet Service Provider (ISP) how durable they are to detecting these attacks and withstanding them. I usually take my teams and we do a close examination to see how we would quickly detect a Denial of Service attack through the monitoring systems.

I always encourage my clients to build a standard target federation architecture. This is followed by an independent third-party firm review the architecture. For each federation deployment external penetration tests are conducted.

Monitoring, Incident Management and Reporting

Depending on the degree of risk associated with federation transactions, enterprises need to pay close attention to monitoring, incident management, and reporting between the federation partners.

At each of my federation engagements, I get my team to “follow the electrons” between the enterprise and its federation partners all the way through to the end application/service. This means understanding the external firewalls/load balancers, web servers in the demilitarized zone (DMZ), internal firewalls/load balancers, internal identity management systems and on to the application servers and supporting databases. A typical federation transaction will take somewhere between 40-70 steps.

Next, the enterprise realizes it is joined at the hip to the federation partners via the federation legal agreement. Service level agreements are required to be met, which in turn usually means running a synthetic transaction once per minute between the identity provider and the Relying Party.

All of the 40-70 federation transaction steps need to be monitored, with log files and metadata rapidly available for troubleshooting. When an incident occurs between the federation partners, it usually requires an integrated incident management ticketing system to track the issue through the partners along with an agreed upon incident management escalation pathway. This is accompanied by automatic reports.

All of the above needs to be clearly specified in a legal agreement appendix to avoid any confusion that might adversely affect the enterprise's customers and/or generate potential legal liability.

Privacy Breach

Next year in Europe the GDPR (General Data Protection Regulation) comes into effect. This means that greater attention must be paid to the handling of identity data. Jurisdictions like Canada are just implementing mandatory privacy breach reporting. The bottom line is that in a federation agreement, the business processes are clearly identified for notification of a privacy breach as well as the technical processes for resolving them.

Disaster Recovery

All federation partners need to agree on notification processes for when a disaster brings down the enterprise's services and how the partners will then be notified when the systems are brought up to quickly conduct tests. This also needs to be stated in the federation legal agreement.

Identity Data Standards

For every identity attribute exchanged in a claim/token, the actual data standard for each one must be noted. I've found in some enterprises that some international data standards for identity attributes are not adhered to.

Of particular importance in the subject identifier (name identifier in SAML). The federation agreement should state that it is unique for the identity for their lifetime. Sometimes email addresses are used as the subject identifier. This is not unique for the identity for their life.

The legal agreement should also state the business and technical processes used for making any change to an existing attribute. This way no unexpected surprises may occur.

Session Management

The internet is stateless. One of the appeals of using OIDC is the use of tokens which means that federating partners don't have to maintain lists of users and their state. HOWEVER, it all depends on degree of risk. Some federation applications/services might be low risk while others are high risk. Typically, high risk applications require stronger identity and credential assurance as well as shorter user session times. A major concern is often regarding idle session times, and the risk of a malicious person taking control of a user's computer.

The federation legal agreement should state what the session management practices are. Depending on the federation partner, they may or may not want to assign legal liability to the other partners dependent upon how the session is managed by the OIDC identity provider.

Test Accounts

In a federation agreement, there is a wide number of different test accounts required:

- Synthetic transactions between federation partners
- Help desk test accounts
- Incident management test accounts
- Performance testing

If there is more than one environment, test accounts need to be recreated for each environment. Often, it is the OIDC identity provider who creates these accounts with the Relying Parties consuming them.

Therefore, this should be specified in an appendix to the federation agreement, with agreed upon business processes when making a change to any test account.

Summary

With the advent of OpenID Connect, it is now possible to achieve identity federation on the fly via a discovery service. This will work for some enterprises where the risks are low and the benefits outweigh the risks. I can see a day when an enterprise might not only do this for some of their customers and services but also operate a more traditional identity federation where the risks are higher and stronger risk mitigation is required. It's what I call a "hybrid identity federation model".

Before your enterprise enters into a federation agreement with any federation parties, regardless of the type of identity federation used, your Business, Legal, and IT leaders need to be aware of all the potential risks and liabilities. As this paper shows, there are many to consider.

Each one needs to be considered, decisions need to be reached on possible impact and liabilities to the enterprise and then risk mitigation measures need to be identified to offset the risk. Some of the risks can be mitigated by technology and configuration settings; others cannot. If possible, in the legal agreement, liability should be assigned to one of the partners.

All business and technical processes used in federation need to be clearly defined and stated in a legal agreement appendix. This eliminates any potential confusion possibly impacting your user's experience between your enterprise and your federation partners.

In the world at large, there are very few lawyers specializing in identity federation contracts as well as identity architects/program managers who have successfully deployed and maintained them. As enterprises rush into identity federation, there is a growing risk of poorly deploying the systems resulting in data breaches, malware/ransomware attacks, etc. This can adversely affect your enterprise's brand.

I've frequently found large enterprises take months to deploy federation agreements at high cost. When I left the enterprises, they were able to integrate in a matter of weeks to a month with a federation partner. Achieving this requires:

- Template legal agreements
- Target federation architecture
- Standard automated testing tools
- End to end synthetic monitoring
- Enterprise risk spreadsheets
- Attack vector spreadsheets for the federation protocol you are deploying
- Standard business processes
- Integrated incident management/escalation
- Third party penetration testing
- Excellent Relying Party and identity provider guides answering all the questions before integration
- Standardized reports
- Training for IT teams on federation protocols and risk mitigation

- Educational management decks on subjects like the federation protocols, risks, liabilities, and assurance
- Automated ways to rapidly configure a new federation environment
- Checklists for the legal, business and IT teams to follow

I am taking my federation deployment experience and working with partners, like Nulli, to create a standardized consulting model to rapidly customize identity federations and deploy. If you'd like to learn more about this, then please contact myself.

About the Author

Guy Huntington is a veteran identity architect, program and project manager who's lead as well as rescued many large identity projects with many of them involving identity federation. His past clients include Boeing, Capital One, Kaiser Permanente, WestJet, Government of Alberta and Alberta Blue Cross. As one of his past clients said "He is a great find, because he is able to do high quality strategic work, but is also well-versed in project management and technical details, so he can traverse easily from wide to deep. With Guy, you get skills that would typically be encompassed in a small team of people."