





Document Title: Framing a Discussion to Foster SP Middlething Deployments

Document Repository ID: TI.168.1

DOI: 10.26869/TI.168.1

Persistent URL: <http://doi.org/10.26869/TI.168.1>

Authors:

- Tom Barton, independent consultant,  <https://orcid.org/0000-0003-1878-344>
- Ken Klingenstein, Internet2,  <https://orcid.org/0000-0003-3182-4566>
- Mark Rank, Cirrus Identity,  <https://orcid.org/0000-0001-8930-9247>
- David Walker, independent consultant,  <https://orcid.org/0000-0003-2540-0644>
- Albert Wu, Internet2  <https://orcid.org/0000-0001-7570-0923>

Publication Date: December 4, 2022

Sponsor: InCommon Technical Advisory Committee

© 2023 Internet2 This work is licensed under a Creative Commons Attribution 4.0 International License.

Framing a Discussion to Foster SP Middlething Deployments

(December 4, 2022)

How to Use This Document	1
Background	1
Considerations for Federation	3
Federation Trust	3
Effort Required for Participation and Operation	4
User Experience	4
Implementation Guidance	4
Observations and Questions for Community Consideration	5
General	5
Federation Trust	5
Effort Required for Participation and Operation	5
User Experience	6
Implementation Guidance	6
Appendix A: Use Cases - Details	7
Use Case - EDUCAUSE's federated access gateway	7
Use Case - CILogon	9
Use Case - Academic Journal Publishing Platforms (e.g., Elsevier, Highwire Press, SilverChair)	10
Use Case - Domain-specific research gateways	11
Comparison of Use Cases	12
Appendix B: Typical Functions of SP Middlethings	14

How to Use This Document

This document is intended to start a community discussion of issues related to SP Middlethings, as described below. We ask questions but provide no answers; the answers are for the community to decide. We have explored a few representative middlething use cases that are in place today to identify their distinguishing characteristics, as well as the middlethings' impacts on various stakeholders in the areas of federation trust, security, privacy, ease of participation and operation, and user experience. Note: No attempt was made to study the entire space of middlething use cases, only enough of it to identify significant [Considerations for Federation](#). We finish with a collection of [Observations and Questions](#) for future exploration.

Middlethings are being deployed now, establishing operating principles as they go, and it may be difficult and/or expensive to accommodate those principles in the future if discovered to be in conflict with each other or with existing federation practices. The time is right to explore these issues and resolve them, according to their importance and urgency. We present this to the community to start that process.

Background

The architecture of today's R&E federations presumes a secure end-to-end communication channel between Identity Providers (IdPs) and Service Providers (SPs). Over time, multiple use cases have arisen requiring (automated) mediation of that communication. This mediator breaks the assumption of the end-to-end channel. Reasons for this mediation include protocol translation, enhancement and/or transformation of the information exchanged, managing the complexity of interacting within a multilateral federation when doing so within the SP is not possible or undermines its function, or aggregation of common applications and data sets into a single service for commonality of the user interface or the technical architecture.

In the Summer of 2022, the InCommon Technical Advisory Committee formed an *ad hoc* group to study the potential impacts of this mediation on federation policy, privacy, transparency, usability, and technical architecture. This is that group's report.

We have chosen to call these mediators *middlethings*. "Middlething" is a deliberately ambiguous term, potentially referring to anything that exists along the path between two communicating things. Without context, it could refer to proxy servers, browsers, routers, even fiber optic cable. This report, however, concerns itself with middlethings that actively translate, transform, filter, or enhance the information exchanged between identity providers and *mediated service providers*, primarily for the benefit of the mediated service providers¹. Browsers, routers, cable, *etc.* are out of scope.

What is a mediated service provider?

¹ There are also use cases where the mediation is primarily for the benefit of the IdP that should be investigated. Our focus on mediated SPs is merely to control scope, given the time available.

We use this term to describe the resource a user actually wants to access when signing in. This resource is generally not directly registered in the federation because it is mediated through the “middlething”. For example, to access NIH’s eRA application, a federated user signs in through the NIH Login Gateway. The NIH Login Gateway is the middlething. It is the entity registered in InCommon. eRA, on the other hand, does not appear in the federation. Mediated service providers, may in fact, not speak the federation’s supported protocol(s) at all.

Further, our focus is on use cases that support research and education, particularly research collaboration and scholarly publishing. Examples of middlethings that are within our focus include CILogon, the NIH Login Gateway, EDUCAUSE, ezProxy, and scholarly journal publishing platforms such as Silverchair, Highwire Press, Elsevier, etc.

In their classic paper on the dynamics of the Internet ecosystem, “Tussle in Cyberspace: Defining Tomorrow’s Internet” (<https://david.choffnes.com/classes/cs4700fa14/papers/tussle.pdf>) David Clark, *et al*, describe the forces and tensions among the participants in cyberspace as tussles. A similar analysis of the federated identity space helps understanding of how to incorporate middlethings into the model. Different stakeholders have different tussles:

- **For the user**, the dynamic between ease of use and protecting privacy
- **For service providers**, facilitating users while balancing their investment and risk
- **For the identity provider**, facilitating users while protecting their privacy and the institution’s risk
- **For the federation operator**, extending the reach of federation trust for all types of entities

We build on the work of others. In particular, the Scalable Negotiator for a Community Trust Framework in Federated Infrastructures (Snctfi) (<https://www.igtf.net/snctfi/>), sponsored by the Interoperable Global Trust Federation (IGTF) (<https://www.igtf.net>) and the Authentication and Authorisation for Research Communities (AARC) project (<http://aarc-project.eu/>), provides significant guidance from the research cyberinfrastructure community for middlething² operators and mediated service providers³ to enable their inclusion in identity federations in an interoperable and trustworthy manner. Also, Federated Identity Management for Research Collaborations (FIM4Rv2) (<http://doi.org/10.5281/zenodo.1296031>) identifies considerations for various constituencies, including Research Community Proxies (*i.e.*, middlethings).

For more information about the functions provided by SP middlethings, see [Appendix B: Typical Functions of SP Middlethings](#).

² Called “Service Provider to Identity Provider (SP-IdP) Proxies” by Snctfi.

³ Called “Constituents” by Snctfi.

Considerations for Federation

Middlethings have potential impact along multiple dimensions. In order to study this, we examined use cases that were chosen to shed light on the middlethings used to mediate access to services supporting a range of academic activities; what we learned is presented in [Appendix A: Use Cases - Details](#).

We have grouped these characteristics along the dimensions of Federation Trust, Effort Required for Participation and Operation, User Experience, and Implementation Guidance.

Federation Trust

As stated above, middlethings break the assumption of end-to-end communication channels among federation entities. Their impact on trust involves a few issues, including:

- **Who has managerial responsibility for the middlething and its associated SPs?** Federation trust, being among participating organizations as opposed to software systems like SPs and middlethings, is not affected when a middlething and its mediated SPs are under the managerial control of the same organization.

When different organizations have managerial responsibility for the middlething and a mediated SP, then the following issues come into play.

- **Are modifications made to identity assertions from IdPs and is this known by the mediated SPs?** The degree to which modifications made to identity assertions are deemed as appropriate by the mediated SPs' management, as well as the alignment of those modifications with federation policy, determine impact on trust. When the modifications are deemed appropriate and comply with federation policy, the impact is low. Otherwise, the impact is indeterminate, depending on the risk profile of the mediated SPs and IdPs federation-wide.
- **Do middlethings protect the privacy of identity assertions?** Middlethings have access to identity information that otherwise would be protected in transit by strong encryption. The security provided by the middlething, and the purposes it makes of the information affect trust.
- **The transparency of middlethings' operation.** Can end users learn which third parties are responsible for handling their identity information? Where do end users go for help? Do IdPs and the federation operator know who to contact for incident response? Is the overall topology of the federation known, or are some portions masked?
- **Do middlethings' and mediated SPs' policies and practices align with the federation's?** Since federation architecture implies that an organization managing a middlething will have joined a federation, that organization will be bound by federation policies. Assuming compliance, there is no impact on trust. If the SPs it mediates are also required to join the federation, there is again no impact on trust. If they are not

required to join the federation, however, there is an indeterminate impact on trust. How might any inconsistencies between SP policies and federation policies be identified and addressed?

Effort Required for Participation and Operation

There are multiple potential impacts of middlethings on federation participation and federation operation:

- **Support** - Middlethings may introduce additional parties into support processes, affecting end user support and incident response.
- **Participation Practices** - Depending on a middlething's business model, its mediated SPs may not have joined the federation. Also, are policies implemented by the middlething administration consistent with those of the federation? Are clarifications and/or changes to federation policy warranted?
- **Federation Tooling** - If clarifications or changes are made to federation to federation policy to support middlethings, does that create a need for modifications or additions to the federation operator's tool set?
- **Deployment Assistance** - Some middlething operators may require assistance in various forms from the federation operator, such as architectural patterns, ready to use solutions and software, or operational guidance

User Experience

Middlethings may impact users' experience. Here are some examples:

- **Service Discovery** - Middlethings have the potential to help or hinder the discovery of services within the federation, depending on the availability of service discovery tools within the federation, and whether the middlething uses those tools.
- **Consistent User Experience** - As with Service Discovery, middlethings may help or hinder consistency in users' experience of using the federation.
- **Expected Behavior** - Middlethings have the potential to modify the behavior of IdPs, discovery services, service catalogs, *etc*, for good or ill.
- **Finding Help** - Middlethings have the potential to help or hinder a user's search for help.

Implementation Guidance

Expectations of federation participants may be different for middlething operators than for operators of SPs and IdPs.

Observations and Questions for Community Consideration

General

- The focus of this report is on middlethings that primarily benefit mediated service providers. Middlethings that primarily benefit mediated identity providers also warrant attention.
- There is a lot that could be done to foster middlething deployment. How urgent and important are the various issues raised in this report?
- Are there emerging technologies that may affect the relationship between middlethings and the rest of the federation?
- Do privacy frameworks like GDPR impact the relationship between middlethings and the rest of the federation?
- Should these issues be addressed by each individual federation, or should some issues be addressed internationally?
- Is there a role for funding agencies to foster appropriate deployment of middlethings?

Federation Trust

- Much of a middlething's impact on trust depends on whether it and its mediated SPs are under common management, perhaps involving outsourcing from one to the other.
- When middlethings are under the management of third parties, should the community establish expectations for the relationship between middlethings and their mediated SPs, addressing such things as modification of identity assertions, privacy protection, transparency of operation, or other policies and practices of middlething operators?

Effort Required for Participation and Operation

- When and how should a middlething's support personnel be included in end user support, as well as incident response?
- Do we have the right participation model for SP middlething operators and the relying parties behind them? Should a middlething's mediated SPs be registered federation participants? If so, who performs registration and other administrative tasks? If not, what expectations, if any, should the federation have of a middlething's onboarding practices, including reporting of information about the middlething's mediated SPs?
- What additional technology tools, if any, are needed by a federation operator to support middlethings?

User Experience

- How can middlethings help (or even improve) users' intuitive understanding of federated access activities?
- Middlethings have the potential to help users discover services that are available via the federation. Are there specific actions that should be undertaken to foster this?
- Middlethings have the potential to increase consistency in user experience.
- consistent use of terms/ key visuals when describing federated access activities. Are there specific actions that should be undertaken to foster this?
- How should a middlething hand a user off to the IdP or mediated SP for support? Do IdPs and mediated SPs need to make accommodations for middlethings?

Implementation Guidance

- Are the expectations for middlething operators different from those for SP and IdP operators?
- Is there a generalizable “right” architecture for an SP middlething? Should we advocate for such a thing (as we do for IdPs) so that there is consistency and therefore easier adoption?
- Do federation operators have a role in providing standardized implementation and deployment guidance for middlethings within the federation?

Appendix A: Use Cases - Details

These use cases highlight various aspects of SP middlethings that deserve further thought as they are incorporated into the federation. We have identified the following use cases. (Note that, due to the constrained schedule, our group was not able to learn all details of all use cases. When this is the case, we have indicated potential future work with “<td>.”

Use Case - EDUCAUSE’s federated access gateway

The EDUCAUSE "middlething" (<https://www.educause.edu/Login.ashx?returnUrl=/>) is a "Coke Classic" SAML authentication proxy with limited assertion modification, business logic processing, and policy enforcement capabilities. It is registered in InCommon as a service provider (<https://sso.educause.edu/sp>) by the EDUCAUSE organization and carries R&S tagging.

Business Model

The EDUCAUSE “middlething” presents a single SP in InCommon. This single SP gates access to a dozen resources, all operated by EDUCAUSE or by a vendor contracted by EDUCAUSE to provide services to its community members.

EDUCAUSE enters into business arrangements with all of the downstream SPs. We assume that they all have traditional data sharing contract language to address liability and risk.

Services and Functions

The EDUCAUSE "middlething" performs the following functions:

1. **Architectural Abstraction Layer** - provides an abstraction layer between federation IdPs (~200) and about a dozen back-end service providers so that changes to participating IdPs do not cause cascading changes to service providers and vice versa. The upstream IdPs see the middlething as an SP, and the mediated SPs see the middlething as an IdP.
2. **SP Bridge to Federation** - domesticates many of the mediated service providers so that they can function in a multilateral federation, by providing metadata processing, per SP attribute transformation, and addressing SAML implementation gaps. EDUCAUSE maintains third-party service agreements (contracts) with these SPs. The mediated SPs are a mix of custom software, traditional enterprise solutions (for example NetForum), and HigherEd solutions such as Instructure Canvas.
3. **Custom "internal IdP" integration** -- EDUCAUSE uses a CRM called NetForum and historically EDUCAUSE participants that didn't have an InCommon IdP would create credentials in NetForum. As part of the deployment, Cirrus built a custom bridge to NetForum using APIs to present NetForum accounts as a SAML IdP that could be a peer to InCommon IdPs at the Proxy. The NetForum IdP is NOT registered in InCommon and

can only be used with the EDUCAUSE proxy (EDUCAUSE does have an IdP in InCommon for EDUCAUSE staff).

4. **Business logic execution** -- When end users log in with federation IDPs, they are put through a linking flow to connect their assertion to records in NetForum. If a link is not found, they are put through a CRM registration to onboard. The linkage does rely on the member organization's IdP entityId which is mapped to the organization's membership records.
5. **Attribute Pixie Dusting** (assertion decoration for downstream SP access) -- The middlething has what is effectively an attribute authority call to NetForum during login to add attributes from the CRM to be used for the service providers. I believe all of these attributes are from the end user's NetForum profile. Data release consent is handled as part of the end user's agreement to use EDUCAUSE services.
6. **Policy enforcement** -- The proxy does a limited amount of access control by allowing or preventing assertions to some SPs based on registration flags (thus limited ABAC). End users that don't have access to certain SPs receive a "Not Authorized" message at the proxy.

Risks/Compliance/Regulatory Environment

<tbd>

Technology

The Educause Middlething is a deployment of the Cirrus Proxy (hosted, SimpleSAMLphp based).

The proxy has been in operation for over 2 years.

Needs and Desires

EDUCAUSE has stated that maintaining a proxy solution between the IdPs and SPs allows it to scale out the number of IdPs used for authentication, while maintaining flexibility to select and deploy backend service providers.

Use Case - CIlogon

CIlogon provides an integrated open source identity and access management platform for research collaborations, combining federated identity management (Shibboleth, InCommon) with collaborative organization management (COmanage). Federated identity management enables researchers to use their home organization identities to access research applications, rather than requiring yet another username and password to log on. Collaborative organization management enables research projects to define user groups for authorization to collaboration platforms (e.g., wikis, mailing lists, and domain applications). CIlogon implements the AARC Blueprint Architecture and the REFEDS Assurance Framework.

Business model

CIlogon is an open source project, with source code in GitHub. Research collaborations, including major science research gateways, use CIlogon to connect with federations. Adopters include: 2i2c, ACCESS, Apache Airavata Test Drive, Ask.CI, ATLAS Connect, Australian BioCommons, BNL Quantum Astrometry, Brainlife.io, CADRE, CERN PanDA, Chem Compute, ClassTranscribe, CloudBank, Clowder, CMS Connect, Connect.ci, Custos, CyberGISX, CyVerse, DataCite, Duke CI Connect, Einstein Toolkit, FABRIC, Fermilab, Flywheel, GeoChemSim, Globus, GW-Astronomy, HubICL, HTRC, ImPACT, LIGO, LROSE, LS-CAT, LSST, Mass Open Cloud, MIT Engaging OnDemand, MSU HPCC OnDemand, MyGeoHub, NCAR PRESTO, NEON, NIH ClinOmics, NIH KnowEnG, Ocean Observatories Initiative, Open Science Chain, OSC OnDemand, OSG Connect, Pacific Research Platform, QUBES, SciGaP, SCiMMA, SEAGrid, SeedMeLab, SimVascular, Social Media Macroscopic, UCLA JupyterHub, and Vanderbilt JupyterHub.

Services and Functions

See <https://doi.org/f6dqgk>.

Risks/Compliance/Regulatory Environment

<tbd>

Technology

See <https://doi.org/f6dqgk>.

Needs and Desires

<tbd>

Use Case - Academic Journal Publishing Platforms (e.g., Elsevier, Highwire Press, SilverChair)

These are businesses who operate online academic journal hosting/publishing platforms. Among their other functions, these platforms perform “middlething” IAM functions. The services they provide can include traditional academic journals publication; collaboration services for editing, refereeing, *etc.*; management interfaces for librarians; deep linking; and more.

Business Model

Academic Journal Publishing Platforms charge a fee to journal creators to use their platform to host/publish their journals. They also charge a subscription fee from institutions who wish to subscribe to the journal(s).

Services and Functions

<td>

Risks/Compliance/Regulatory Environment

<td>

Technology

varies

Needs and Desires

- Better user experience
- Easier linking / access to articles leading to higher on-demand purchase/subscription hits.

Use Case - Domain-specific research gateways

These gateways serve primarily to aggregate and integrate research resources. They are common in the sciences (NSF supports their creation - see https://www.sdsc.edu/News%20Items/PR20220906_science_gateways_center.html).

Examples:

- GIS Sandbox - <http://www.gisandbox.org/>
- Hydroshare - <https://www.hydroshare.org/>
- Science Gateways Community Institute (SGCI) - <https://sciencegateways.org/> for gateways as a service

Business Model

Most are spun up by grant funds awarded to researchers and educators. There is no formal business model; the focus is on the mission, not sustainability.

Services and Functions

They provide a number of important services to dedicated communities of interest, including:

- Aggregation of computing, data and other resources for a specific research domain
- Easy entry to expand a research community and provide outreach and education
- Democratize access, to both data and supercomputers
- Provide some local groups and permissions for access control

Risks/Compliance/Regulatory Environment

Little understanding of regulations about security, privacy, etc.

Technology

Varies

Needs and Desires

- Ability to add resources specific to the community.
- Solutions to sustainability challenges.

Comparison of Use Cases

	HE online community (Educause)	CILogon	Academic Journal Publishing Platforms	Domain-specific research gateway
Function / Service				
Protocol translation (e.g., SAML to OIDC)	✓	✓	✓	✓
Present single SP in federation	✓	✓	✗	✓
Identity linking / merging	✓	?	?	✓
User attribute minting and transformation	Not directly – attribute minting is handled by EDUCAUSE’s CRM system	✓	?	✓
Business Model / Org Structure				
Same organization operates Middlething and all resources behind it	yes	no	no	yes
All resource organizations are Federation Participants	yes	yes	no	yes
Middlething relays user information to resource	yes	yes	?	yes
Middlething operator discloses measures taken to safeguard user information shared with external resources	n/a	U Illinois web privacy policy applies	no	yes
Federation Trust				
Resources relying on the middlething have clear understanding of how identity assertions from IdPs are relayed/transformed/enhanced by the middlething	?	?	?	?
Middlething discloses how it protects the privacy of identity assertions in a way consistent with Federation requirements	?	?	?	?
An IdP or user can readily discover which resources live and/or are receiving information behind the middlething?	?	?	?	?
Middlethings’ and resources’ policies and practices align with the federation’s	?	?	?	?

	HE online community (Educause)	CILogon	Academic Journal Publishing Platforms	Domain-specific research gateway
Other?				
Federation Participation and Operation				
End-to-end user support processes and hand-off is clear	?	?	?	?
Resources receiving IdP-asserted information via Middlething are Federation Participants	?	?	?	?
Current Federation tooling fully supports middlethings registration and operation	no	no	no	no
Others?				
User Experience				
Service Discovery - User intuitively understands what services are available behind the middlething	✓	?	✓	?
Consistent User Experience - middlething promotes federated sign-in experience, i.e, home institution discovery	No - because Federation has not advocated for such consistent experience	No - because Federation has not advocated for such consistent experience	No - because Federation has not advocated for such consistent experience	No - because Federation has not advocated for such consistent experience
Finding Help -	?	?	?	?

Appendix B: Typical Functions of SP Middlethings

Middlethings are deployed to provide services that address a number of issues. We have observed the following:

- **Protocol normalization.** SPs often are packaged with minimal SAML implementations that do not interoperate well (or at all) within R&E multilateral federations. Middlethings are often deployed to facilitate federation interoperation. Educause is an example of a middlething that does this.
- **Protocol translation.** SPs are often packaged with no SAML implementation. Middlethings can provide translation services to/from other protocols, often OIDC. CILogon is an example of such a middlething.
- **Enhancement of identity information.** Middlethings can store and assert to mediated SPs information about users that is not supported by home institutions. Educause and CILogon are examples of middlethings that do this.
- **Enforcement of access control and other policies.** Being in the middle, middlethings can act as gatekeepers to enforce policies, such as those governing access.
- **Integration of multiple SPs.** Middlethings can act as portals, integrating multiple mediated SPs, probably also providing one or more of protocol normalization, protocol translation, and enhance of identity information. Educause, journal aggregators, and NIH are examples of this.