Candidate for Next Baseline	Clear guidance	Would be nice	Explore / Experiment
	E-0003 If using MFA, signal with REFEDS MFA Profile 1.2		E-0007 Error handling with REFEDS MFA - part of E-0003
	E-0004 Standard Attribs		
	E-0006 distinguishes routable email from identifier attrib		
	E-0005 metadata validation behavior		
	E-0001 sign and validate		

E-0001

> “All entities will sign AuthnRequest and (Assertion and/or Response)”
>
> “All consumers will validate the signature of signed elements”
>
> What’s the cost of noncompliance?

- Let’s discuss AuthnRequest separate from the Response
- Not a lot of value to signing the AuthnRequest, the Response will go to a registered endpoint (not the spoofer)
- vs. the Response/Assertion signing is important
- Are there any specific drawbacks to only requiring a signed assertion instead of a (more thorough) signed Response?
    - It seems signing the Assertion is important

    - signing the Response doesn't seem too much more useful because only the Issuer and Status elements are outside the Assertion
- Denial-of-Service / confused deputy concerns:
    - What happens when someone spoofs an AuthnRequest to unwitting IdP, and it releases an assertion to a victim SP

Suggested text (when it's Baseline, it helps to be very precise in our wording):

When processing transactions among InCommon Federation registered entities:

- All entities MUST sign SAML Response - describe in further detail what that means
- All entities SHOULD sign SAML AuthnRequest
- All entities MUST validate the signature in a SAML Response
- When present, all entities SHOULD validate the signature of a SAML AuthnRequest

~~E-0002~~ withdrawn, replaced with E-0004

"use eduPerson"

Can we use a different profile? A more restricted one? A more broad one?

We cannot avoid talking about attrib release and attrib value interpretation

Contributing specs: 1. eduPerson, 2. Subject-id, pairwise-id OASIS spec, 3. Schac, 4. REFEDS Access Entity Categories 5. Ancient ldap attribs

Attribute bundle concerns (tagging/signalling, governance)

How will these attribs translate into an OIDC world?

How specific do we want to get? What's in scope and what's out of scope?

E-0003: " If/when using MFA, use REFEDS MFA Profile"

When processing transactions among InCommon Federation registered entities:

- REFEDS MFA Profile is the only supported MFA Signalling mechanism
- IdP must support REFEDS MFA Profile, even if no one performs MFA (needs to expand on this a bit)
- If an SP requires MFA, it must use REFEDS MFA Profile to request MFA.
- (added 2025-03-20) For the vendor, you will supply this feature and docs.

E-0004: Standard set of attributes, governance around them and about adding new ones – specifically the REFEDS Access Entity Categories

E-0005: Validate metadata by ignoring unknown extensions

E-0006: distinguish identifier from routable email address

2025-03-20:

Albert: mostly targeted at product vendors and institutions behind SPs

Good discussion on how to measure / enforce this; survey? metadata registration? tool measurement?

Do we give badges / trustmarks?

If they don't do it they can't claim they're "InCommon Compatible".

Do we need to specify this notion of "InCommon Compatibility" – is this in scope for this group?

E-0007: how to handle errors when using REFEDS MFA Profile

General points:

- Andy: "What would add the most value?"
- Albert: we need to tell people what's not allowed
- What's the cost of noncompliance? i.e. what if it were actually enforced?
- Focus on UX
- For each item:
  - What is it
  - Why it's important
  - How to accomplish

        - (to be fair this is already happening in separate WGs that are drafting deployment guidance)
        - (we're trying to make the process a bit more repeatable and systematic)
- 11 additional meetings before CommEx (a deadline reference point)
- Things we should definitely talk about:
    - measurement / metrics
    - what happens in cases of noncompliance
        - how we handle a bifurcated world where some people comply and some don't
    - the process of introducing, placing, evaluating new proposals
        - discussion internally
        - outside pressure by regulators
        - industry trends
        - (historically not done) surveying the broader community
    - the process of spinning up "working-groups" to eval and flesh out items
        - look at what's been done in the past
            - Shib
            - MACE-DIR
            - topic-based WGs spun up by TAC or CTAB

WGs and Advisory Groups should aim their deliverables to become expectations, when relevant

What's the public entrypoint / ingestion point for new proposals? (akin to feature request)

Intake

- How do we solicit ideas?
- How do we do evaluations?
- How do we do community consultations?
- If we chunk the intakes into known start times, we can galvanize participation at set times (e.g. once every quarter)
- We need some sort of intake form, then someone does an initial pass, and then they get assigned to committees
- How do we know which committee takes which items?
- For each item we need to think about how members would implement it (in every product)
- How do we flesh out community suggestions enough so that people understand what it actually is?

- Every item needs to have a sponsor
- What's expected of the sponsor?
    - Is it a Project Manager type role?
    - Are they expected to be involved in the resultant working group?
    - Are they expected to know the progression of an item through its evaluation stages?
- Maybe a Proposer can be distinct from a Sponsor?
- We need shorter iteration timelines.
- What's the MVP?
    - Online form to intake idea – "suggestion box"
    - At next meeting, a body (e.g. CTAB) responds to every suggestion that came in
    - On a periodic basis, e.g. every quarter, we brief the community on the stuff we've received
    - What is the communication mechanism with the community? IAM Online? Mailing List?
    - We used CTAB as an example but it'd be nice if it wasn't a closed group giving the first responses
    - How do we streamline the progression from "kernel of an idea" to "working group" (or the equivalent of a working group / debate chamber where the idea(s) can be fleshed out)
    - It sounds like we need a standing open body ("Expectations WG") to give the debate chamber to these ideas
    - Perhaps this body suggests which standing body takes it from there
    - The Expectations WG will spin up the targeted WG to develop an idea
        - Maybe also inform Steering
    - The SWAMID readiness check is quite nice (does not always translate nicely)
- How do we decide which group / venue is appropriate to articulate the "how" of each item, after this group articulates the "why" of each item?
    - Judgment call of this group per each item
        - Candidates: existing bodies (e.g. CACTI, TAC, CTAB, etc.) or new WG
        -

The Items/Guidelines themselves

- Let's base their structure off of AARC Guidelines
    - The granularity of the prescription varies per item
    - Judgment call for each of these how granular we want to get
    - Try to rely on third-party docs/specs where possible

Actually phrasing examples:

E-0003: Refeds MFA req:

- Proposal A: “InCommon Participants MUST use the REFEDS MFA Profile to signal that multi-factor authentication was used at the IdP.”
- Proposal B: “REFEDS MFA Profile is the only officially supported method of signalling the use of multi-factor authentication in the InCommon Federation.”
- Proposal C:
    - When processing transactions among InCommon Federation registered entities:
        - REFEDS MFA Profile is the only supported MFA Signalling mechanism
        - IdP must support REFEDS MFA Profile, even if no one performs MFA (needs to expand on this a bit)
        - If an SP requires MFA, it must use REFEDS MFA Profile to request MFA.

E-0001: Sign authN transactions and Validate signature

- Articulate rationale: prevent spoofing. Phrase in a sentence.
- Articulate how to accomplish using the supported protocols (SAML, OIDC).
- Proposal A:
    - When processing transactions among InCommon Federation registered entities:
        - The IdP, when communicating with the SP, MUST sign the response.
            - In SAML: SAML Response or SAML Assertion
            - In OIDC: must sign claim (consult the eduGAIN/REFEDS OpenID Federation profiling work)
        - The SP, when communicating with the IdP, SHOULD sign the authentication request
            - SAML AuthnRequest
        - All entities MUST validate signatures in the messages received during an authentication transaction.
        - When present, all entities SHOULD validate the signature of a SAML AuthnRequest
    - Which means:

        - All entities must register a valid signing key in SAML metadata

Taking stock

- How different does this need to be from the InC Baseline Exp adoption process that was undertaken in the past? Can we re-use the same ideas and learn from the experience?
    - At the time, there was an assessment made, from measurements/surveys, that the community isn't ready for some of them right now.
    - ~ 1 year of soliciting feedback
    - Differences and similarities, as recounted by Albert:
        - Difference: volume; baseline is on ~1 cycle/3 years; InCExp is more frequent and faster
        - Difference: degree of enforcement
            - The Baseline Community Consensus Process has surveys but the decision-making is very internal
                - We don't necessarily have to solve this, but we should point this out in our report
            - What happens in case of compliance or noncompliance with InC Exp?
                - Trustmarks?
    - The placement of items on the maturity model need to be based on measurements and metrics and a judgment call by the decision-making body
        - Measurements and metrics require an ongoing operational commitment (e.g. surveys, outreach, consultation, tooling, etc.)
        - We have not yet decided which body is the decision-making body… yet
        - Perhaps this body can recommend what the process of "adoption" looks like in the more general sense… across the board.
    - We should point out where this body has consensus, but also where this body was not able to achieve consensus. This output is useful for the intended audience of the report (i.e. CTAB and peer bodies or higher).
    - This body does not have to solve every problem, but it should point out areas where it has identified that future discussion and decision is required (e.g. how work moves through InCommon in general).

- Proposed frame / create cadence for InC's Federation Expectations:

Prep	Detail	Consensus	Implement	Report
Summer	Sept	Nov TechEx	Jan-May	May
CTAB + WG	TAC & Steering	participants	CTAB + I2	CTAB

- A yearly repeating process

Report drafting - details of the Named Items we want to highlight

Deliverables, per charter:

- Formulate an initial ordered list of at most 5 proposed initial new expectations to be evaluated and used to evaluate the process.
  - Each proposed expectation will address a specific need for enhanced scalability, trust, and/or interoperability among identity providers and service providers
  - InCommon expectations should be
    - where possible, independent of a specific messaging protocol (that is not formulated in terms of SAML, OIDC, or another protocol). And
    - make sense for adoption for “1 to 1” integrations not relying on the InCommon Federation as well as interactions among InCommon participants
  - The implementation of proposed expectations should specify
    - information required to meet the expectation
    - behaviors or processes required to meet the expectation
    - suggestions for implementation in specific protocol as and if appropriate
    - relation to or direct support of any of the 5 areas of work defined by InCommon Futures2
- Draft an ongoing process for CTAB to use in proposing and evaluating new expectations.
- Propose a "living" location where new expectations will be proposed and documented.
- Evaluate as wide a variety of opinions as possible and coalesce to a reasonable consensus.

In the most recent CTAB meeting (Mar 17, 2025), there’s developing an understanding that CTAB will apply a process similar to when Baseline Expectations was rolled out to serve as a forum where the ideas coming from intake can be fleshed out; i.e. there didn’t

appear to be any opposition to this quite yet, and folks seemed to be accepting of this outcome so far. The specifics of each item, i.e. the process by which we arrive at the bodies of recommendations may be spun out into individual WGs if needed.

The cadence doc (prepared and presented by CTAB leadership for the Mar 17 CTAB meeting) can be leveraged as a base for the deliverable to satisfy parts #2, 3, 4, partially.

We'll need to explain the 5-6 proposed deliverables as they would go through the cadence. Each of these represent different intake processes that have already played out. Let's narrate about these.

Discussion about specific items:

E-0001 "sign and validate SAML protocol messages": no existing WG working on this, but we all agree this should happen. I.e. organic de facto expectation, existing best practice; we will be codifying it formally. Eventually we'll run out of this type as we document best practices we want.

MG note: doing this for SAML responses is not just a federation MUST, it is a SAML2 protocol MUST.. it is a Baseline Expectation of using the SAML2 protocol, regardless of whether within the federation or not.

MG note: signing and validating of SAML requests is a should, and as the latest Shib IdP security issue announcement states, many of us consider a signed/validated SAML AuthnRequest to add very little real value (except in very rare edge cases)

--

E-0003: "If using MFA, signal with REFEDS MFA Profile 1.2": another de facto practice that is uncodified. "obvious" best practice if you're in the club, but hard for outsiders to understand the detail when unfamiliar.

Q: How specific do we need to be?

Answ suggestion: let's be vague and say what we want, but no impl. detail beyond referring to the actual spec.

Suggestion: For the vendor, you will supply this feature and docs.

Q: What onramp do we give folks who are using products that don't support this? Do we refer them to solutions that accomplish this?... How do we hold the vendors accountable?

E-0007: "E-0007 Error handling with REFEDS MFA"

a companion to E-0003; essentially about SAML error handling rather than this specific case

we think there should be an expectation about this but we're not sure what it is. Someone will have to investigate this, feasibility study, survey of community appetite, measurement of existing practices, etc.

E-0004: Standard Attribs: this is type #3 where there are existing workstreams from WGs, and want to recommend something / drive adoption, but maybe no measurement or no community appetite. This group can fill in the missing pieces, e.g. provide measurement, rank on maturity model, advance conversation cyclically.

Q: Is this supposed to be applied to just new integrations or also existing ones?

A: gradual option, with the hope that specifying this would drive adoption from both SPs and IdPs.

E-0006: distinguishes routable email from identifier attribute: type #4, where we would like something to be true but no one does it that way....

Salient points we want to highlight in the report

Here are the points we definitely want to cover in the report:

- The “cost of noncompliance”

    - These things are not intended to be things were noncompliance gets you kicked out, instead, these are for participants' benefit – and your benefit too, so you can assess where you are
- The significance of this InC Fed Expectations activity towards parts of InCommon business beyond CTAB
    - This effort proposes to formalize and provide a framework for topic intake and decision-making
    - Not intended to tread on existing processes, but augment them and fill the missing pieces
    - Intended to help community members attach and detach themselves from topics of interest as needed
    - Intended to enhance accountability, "keep ourselves honest", make sure work doesn't get forgotten about
    - Intended to demonstrate genuine community needs through measurement
    - Intended to write down expected behavior when it is already widespread community practice, therefore, this is partially a documentation effort!
    - Not intended to replace topic-specific finite-charter finite-end-date Working Groups, but to provide an attachment point before and/or after they complete their activity
- This InC Exp body, nor CTAB, does not have to solve every problem, but it should point out areas where it has identified that future discussion and decision is required (e.g. how work moves through InCommon in general, or where involvement from other steering committees or InC Leadership is needed to make decisions about goals, priorities, future direction)
- Albert's diagram: Federation Expectations Intake Journey
- Gabor:
    - Life stages of an intaken item:
    -
    - 1. Pre-intake
    - 2. Accepted for evaluation
    - 3. Initial eval
    - 4. Community consultation Type A
    - 5. Accept/reject into Workplan
    - 6. Workplan activities and status updates
    - 7. Community consultation Type B
    - 8. Loop to 6.
- 2024-04-03: Not every item has the same origin story. Some are aspirational, some codify existing best practice.
- Andy: How do we incentivize members to show compliance with Expectations?
    - The wording / requirements of each item needs to be specific and tailored towards the audience
- Discussion about SAML2Int – using it as a strawman of a spec that is aspirational and prescriptive but unclear adoption / conformance

- Michael: the federation is valid target for an expectation also, i.e we can recommend that the federation operator do a specific thing or enforce a specific thing too, e.g. we can propose that the federation should codify behaviors
-

2025-04-17 meeting

1. Gabor: what's the continuity of items in the running process? How do we ensure that some items will keep going for next year's cycle?

    David: continuity is assumed

2. Sprucing up E-0001, removing this stuff because it's not the final body:

    All entities will sign AuthnRequest and (Assertion and/or Response)

    - All consumers will validate the signature of signed elements
    - When present, all entities SHOULD validate the signature of a SAML AuthnRequest
    - Or in other words

Parking lot / idea board

- Device assurance / endpoint assurance enforced / signalled at assertion time
- MDQ (CTAB Feb 18)

Meeting 2025-04-24

1. What makes a good expectation? (This group should try to write down its knowledge for future parties)
    a.
2. How do we assign a steward for each expectation?
3. Who assigns these?