

FHIR Infrastructure Minutes WGM 201905

Attendees

See [FHIR-I Attendance_bit.ly_fhiri - 201905.pdf](#)

Agenda for WGM

http://wiki.hl7.org/index.php?title=FHIR_Agenda_201905_WGM

Monday Q1

- General intro to FHIR-I for newcomers
- CapabilityStatement discussion
 - How to get this to scale? Current implementation can get quite large.
 - Options discussed
 - Breaking it up into smaller components with references
 - Using summary to retrieve smaller files
 - Terminology based solutions with code=value pairs
 - Smaller resources defining a single capability with a canonical URL, much like a profile (can ask a server if it has capability X via that URL)
 - May be sufficiently covered by the \$implements operation
 - Could end up with more capability resources than we have clinical resources
 - Metadata resource with the URL, etc, then a set of triples with context, property, value.
 - Would be like RDF
 - Using test suites to discover capabilities (i.e. try it see if it works)
 - Should Smart of FHIR, CDS Hooks, etc. capabilities be defined here or elsewhere? Probably elsewhere.
 - There are similar discovery-ish JSON files also in use
 - OpenAPI
 - Discussed OpenAPI authentication and authorization discovery
 - Sufficient to determine if a client is compatible
 - Consider just using this vs. having security stuff in CapabilityStatement?
 - Grahame likes this
 - Smart on FHIR
 - Currently duplicative (same info in both CapabilityStatement and their custom JSON file)
 - Consider moving everything to just the JSON file, but timeframe would be R5 or later (a couple years)
 - Motion: Grahame Grieve/Lloyd McKenzie:
 - Grahame moves that a group led by him create a strawman stripped down CapabilityStatement2, discuss on Zulip, and try out in Atlanta in Sept. Would be terminology driven. Will start by listing the problems we need to solve.
 - Discussion:
 - The strawman will be the changes we would like to make to CapabilityStatement, so there would not be 2 separate resources in R5.
 - Will try not to change the normative parts, but may deprecate some
 - Vote:
 - Abstain: 3
 - Against: 0
 - For: 39
 - Motion passes
- Bulk Data Discussion
 - Josh described the prefer=respond-async header.
 - See <http://hl7.org/fhir/async.html>
 - Currently the CapabilityStatement does not expose this, so now way to discover if a server supports an asynchronous response
 - Could be handled by the changes described above under the CapabilityStatement discussion
 - For now will likely do as an extension
 - Need a general solution to prefer headers at some point
 - Motion: Josh Mandel/Rick Geimer
 - Will not define a specific discovery protocol that a client could use to discover if a server supports async export. Will consider a general solution to prefer headers at some point. Though can currently declare that a given system supports a specific specification (i.e. supports Bulk Data). We will define in the Bulk Data spec a canonical URL that conforming servers can point to to declare conformance. Pointed to by CapabilityStatement.instantiates.
 - Discussion:
 - Discussion about whether the implementation guide is enough or if we need to be specific about what actor in the IG (i.e. client vs. server)
 - Decided that the actor CapabilityStatement is the way to go (so will leave as CapabilityStatement.instantiates vs. CapabilityStatement.implementationGuide).
 - Vote
 - Abstain: 9
 - Against: 0
 - For: 32
- Reviewed agenda
 - FHIR-I Q2: mustSupport discussion
- Adjournment: 10:27am

Monday Q2

- Discussion on mustSupport
 - Concept dating back to HL7v2, v3 – even if an element can be absent from specific instances, servers need to be able to support sending /receiving/using this element when relevant.
 - Do we need context to address "must support for what"?
 - There are some concerns that implementers ignore elements not marked as "must support," but it's not clear how widespread this behavior is
 - This concept also is not well understood by the implementer community
 - Testing of mustSupport: expectation is that in controlled testing, it's possible to expect systems to produce and/or consume data elements marked as mustSupport.
 - GF#20325
 - concern is with IGs that say "SHALL support" in narrative but don't set "mustSupport: true" in the StructureDefinitions of their profiles.
 - given different audiences using the same profiles, there's some nuance to the meaning... creeping toward "it MAY be the case that you mustSupport this element – please pay careful attention"
 - You could solve this with separate profiles for each context of use, but this is a lot of work / complexity, and leads to a proliferation of profiles
 - Motion: Lloyd Mckenzie / David Pyke:
 - solicit community feedback on a proposal to deprecate mustSupport and replace it with supportExpectations with type=url, definition of "points to the location that expresses the support expectations for the element."
 - Vote
 - Abstain: 8
 - Against: 0
 - For: 15
 - US core had some elements with required bindings and a data absent reason – but validation said: if required binding is present, there's no away around it. Alternative codes in the element, or extensions, aren't recognized.
 - how to deal with data absent reasons?
 - alternatively, require the use of an "other" or "unknown" type code in very required valueset.

Monday Q3

Agenda: Tracker Items

20688: Does Resource.meta.source point to a system, or a full resource URI, or something else?

- https://gforge.hl7.org/gf/project/fhir/tracker/?action=TrackerItemEdit&tracker_item_id=20688
- Discussion about what the URI should actually point to (the server, the actual resource on that server, etc.).
- Documentation mentions Provenance.entity.what[x] but that no longer exists (was a choice in STU3, but now just a resource reference in R4)
- Decided that Resource.meta.source can be the server, the actual resource, a URN (urn:oid:1.2.2...) etc. Will also update the text and examples in various places. Will remove the reference to what[x] in the definition, and do a global search/replace for what[x] to see what else needs to be updated.
- Persuasive
- Motion: Josh Mandel/Rick Geimer:
 - Abstain: 2
 - Against: 0
 - For: 19

20332: Create a \$restore operation

- https://gforge.hl7.org/gf/project/fhir/tracker/?action=TrackerItemEdit&tracker_item_id=20332&start=0
- Duplicate of 16198

20947: Clarification on which operations can be included in a batch/transaction request

- https://gforge.hl7.org/gf/project/fhir/tracker/?action=TrackerItemEdit&tracker_item_id=20947
- Discussed examples where you may not need a resource in the body when doing PUT/POST in a batch or transaction, such as POST-based search.
- A workaround would be to supply a Parameters resource with zero parameters inside.
- Not Persuasive
- Motion: Josh Mandel/Rick Geimer
 - Abstain: 0
 - Against: 0
 - For: 21

Monday Q4

FHIR-I w/CDS, II, MnM, OO

Agenda: FHIR Workflow, Objective FHIR

Notes:

- Record Status vs Clinical Status
 - FHIR is unable to represent all of the status combinations that can be expressed in CDA.

- Lloyd: Just because something CAN be done in CDA, doesn't mean that it is done in practice and so therefore FHIR doesn't necessarily need to support it. FHIR relies on specific use cases being brought forward to justify any modifications.
- Objective FHIR
 - Mark Kramer, MITRE Corporation
 - see also <https://chat.fhir.org/#narrow/stream/179166-implementers/topic/Objective.20FHIR>
 - MITRE has created a model expressing FHIR in an object oriented form
 - see <http://standardhealthrecord.org/guides/>
 - this model smoothes out some of the irregularities between FHIR resources (for example: createdBy, author, creator, etc are semantically equal)
 - implementation guides can be generated from the object oriented model
 - Lloyd: FHIR intentionally chose to NOT implement an object oriented approach to make it so that the FHIR resources in each domain could use terms and element names that were familiar to that domain. Sometimes an element that is named the same in different FHIR resources may have slightly different meanings in each domain
 - Action: MITRE to see if they can produce some reports that identify the irregularities and provide those to the FHIR Workflow work group to process
- FHIR Definitional Resources for Things
 - How do we help implementors know which resource to use?
 - FHIR has introduced a number of definitional resources and implementers are not clear on which to use.
 - Action: Lloyd will update the Creating FHIR Resources wiki to include updated design considerations.

Tuesday Q4

SDC Tracker items

19530: Document how to access variables created by the "variable" and "questionnaire-context"

- https://gforge.hl7.org/gf/project/fhir/tracker/?action=TrackerItemEdit&tracker_item_id=19530
- Discussion about whether % signs are necessary in all cases
- In FHIRPath, variables are required to have %
- Will require %
- Persuasive
- Motion: Brian Postlethwaite/Rick Geimer
 - Abstain: 0
 - Against: 0
 - For: 25

20396: Move choice/open-choice out of Questionnaire.item.type, and add Coding

- https://gforge.hl7.org/gf/project/fhir/tracker/?action=TrackerItemEdit&tracker_item_id=20396
- Will drop choice/open-choice, add Coding, and an additional element called choiceConstraint.
- Persuasive with Mod
- Motion: Rick Geimer/Brian Postlethwaite
 - Abstain: 3
 - Against: 0
 - For: 20

SIRB Project

- Anita Walden presented project overview
 - (Quick Google search found this, not sure if right link): <https://grants.nih.gov/grants/guide/notice-files/NOT-OD-16-094.html>
- Here because project wants to know if they should use SDC, and if so what is the best way to work with it.
- Want to use Questionnaire to obtain data for clinical trials
- Not a lot of back in forth for this use case, but could happen. FHIR-I has not done a lot of discussion around workflow with Questionnaire.
- Discussed a bit about what might be appropriate for this.
- Some discussion about existing SDC/Questionnaire implementations
- The project plans to ballot 1 or 2 implementation guides with different parts of the spec balloted at different times. Per Lloyd, FMG will likely recommend a single IG balloted multiple times. Also some concerns that some of the content may get more push back than others (another desire for 2 IGs).
- They were thinking of going with an informative document because they were worried about not having 2 committed implementers at ballot time, but Lloyd strongly recommended the normative track (STU).
- Anita asked FHIR-I if we would like to be a co-sponsor or interested party.
- Will review the PSS on next week's SDC call (Thursday at 5pm).

SDC Ballot Rec

21376: Should disabled data be visible? - SDC #65

- https://gforge.hl7.org/gf/project/fhir/tracker/?action=TrackerItemEdit&tracker_item_id=21376
- Discussion about whether it is appropriate to hide disabled data (enableWhen condition is false) vs. just disabling them but leaving them visible (i.e. maybe greyed out and uneditable).
- Will update the text to make it clear that non-enabled data SHALL not be displayed.
- Persuasive
- Motion: Brian Postlethwaite/Rick Geimer
 - Abstain: 1
 - Against: 0
 - For: 22

21377: Should non-enabled answers be removed? - SDC #66

- https://gforge.hl7.org/gf/project/fhir/tracker/?action=TrackerItemEdit&tracker_item_id=21377
- Will make this a SHALL, but should provide a warning of some kind to the user that they entered data that is now disabled and will be removed.
- Motion: Brian Postlethwaite/Rick Geimer
 - Abstain: 6
 - Against: 0
 - For: 16

Adjournment

Wednesday Q1

Co-chair: Josh Mandel

Minutes: Michael Donnelly

Pain points from the HL7 FHIR Connectathon this weekend:

- Hard to subscribe to exactly what you want.
- Knowing when something stops meeting your criteria (issue with expressivity)

Most real-world systems are broken down into **topics** and **subscriptions**; we want clear separation between those. One topic can have multiple subscriptions. Topics have been handled thus far as a Subscription linked to an EventDefinition

Good syntax and a framework will be better for R5 than what we have today.

Josh has a proposal to refactor FHIR subscriptions and test it out at the HL7 FHIR Connectathon before this fall's WGM. <https://docs.google.com/document/d/1rjOI49M-PBVD7DpLT9LolV4e3qNU1jZbBKdxTwtRE8/edit>

This won't use lots of extensions because there's a "developer ergonomics bias" against complicated structures like that.

We're looking at examples of subscription mechanisms that have topics with filters.

A trigger has:

- a human readable description
- a computable definition. native-FHIR systems (like HAPI) can use this; façade servers will have to write their own logic. The definition describes what resources are subscription candidates for this topic

A topic has a list of properties on which that topic can be filtered. These canFilterBy properties will be FHIR search criteria; but not all FHIR search criteria will filterable properties.

Once you create a topic, nothing happens until there's a subscription to that topic.

Topics can be server specific. If there are widely defined topics, a server will specify which of those it will support.

:checkbox: PARKING LOT: how do we align topics from server to server?

When a client wants to subscribe, it creates a subscription on the server. Much of this is like what we have today. Differences:

instead of search criteria, the client specifies the topic to which it wants to subscribe
it supplies a list of criteria for any filters defined in the topic
Topic has "canFilterBy" and Subscription has "filterBy"

(filterBy syntax can look a lot like FHIR search syntax)

How does a provider get notifications about patients he or she cares about?

Today you might be able to do something like "Encounter?patient:_has:Group:patient:_id=my-group"

With this proposal, you'd filterBy "matchType": "memberOfGroup" with "value": ["Group/my-high-risk-patients"]

:white_medium_square: PARKING LOT: Filter syntax

:checkbox: PARKING LOT: Scheduling use case

Previously, there was no clear way to learn about something falling out of your search set.

In the current system, when you have a REST hook channel, you specify a URL that the server treats as a FHIR base.

In this proposal, the server posts a bundle to that URL.

A FHIR server can parse the Bundle and use it. Other clients will know it'll get everything to the same place.

If a server only wants to batch up notifications (e.g. twice a minute), it can put more things in the Bundle.

Bundles can handle deletes.

A client can say they only want the IDs and not the full resources.

Some clients want nothing - just getting pinged on the endpoint with an empty payload will let them know something happened.

The last case could be useful if the message is going through somewhere unsecure.

How does a Bundle do a delete?

Bundle.entry.request.method=DELETE

This Bundle is the server telling the client, "This is what you should execute to look like me."

But many clients don't want to look like the server.

Should this be a transaction bundle or a history bundle?

A transaction Bundle only allows you to send the request part.

So history is probably what we want. (Josh updated his Google doc)

How can a server detect deletes?

Before: [Practitioner.org](#) = 123

After: [Practitioner.org](#):not=123

But the latter won't work because there's nothing to evaluate.

There's no resource to reason about.

Even with an implied search criterion of `_id` being the current context, nothing will ever be true (in FHIRpath) for it.

Maybe After: "\$not-found"

?

Looks like in addition to Before and After with FHIRpath, we need boolean's of `isDeletion` and `isCreation`.

:star: For later follow up: we need syntax to address detection of deletion and creation in native FHIR servers.

Scheduling use case

A client wants to know about schedule slots - when it can schedule appointments

They're looking at Slots as a proxy for time periods that a Practitioner is available.

Subscribe to Slot where `Practitioner=123`

They get an update when the Slot changes

The client does a couple of queries.

The problem is that there's no Practitioner on Slot.

Idea: instead of subscribing to Slot, the client could subscribe to changes in Schedule for Practitioner, but when there's a change the client would query for Slots for that Practitioner.

Canonical references

Josh would like an Argonaut IG to be able point to topics out there somewhere.

url - server declares its own canonical url

instantiatesUrl - declare alignment with an official topic

Should a `Subscription.topic` be to a shared URL or one specific to the server?

:star: This remains an open question.

Filter Syntax

Should we have filters split out into separate JSON elements or combined into comma separated lists?

Arbitrary. The group has a weak preference for JSON, and we'll try that out first, but we're open to considering the other way in the future.

Motion: @Eric Haas moved/Michael Donnelly seconded:

- Use Josh's proposal to draft a new Topic resource
- Update the Subscription resource
- make best effort to resolve the outstanding issues (see stars in this Zulip topic)
- get it into the CI build
- test this in an HL7 FHIR Connectathon track this September.

Vote: 13/0/0

Considering Michele's alternate proposal

Details at <https://github.com/CareEvolution/Public/blob/master/Subscriptions.md>

Use a List for subscriptions.

Entries are added to or removed from the List by the server. (The mechanism for this is up to the server.)

> Michele Mottini: I have an alternative proposal: <https://github.com/CareEvolution/Public/blob/master/Subscriptions.md> (that is backward compatible, so it can be used in existing servers)

There could be a list of "Admissions in the last 24 hours"

Josh: Since the items in the list are patients, the list should be called "Patients who were admitted in the last 24 hours".

A client can subscribe to the List and get notification of changes to it.

In the example of "patients admitted in the last 24 hours", the list would be updated when a new patient is admitted or when a patient on that list has been admitted for more than 24 hours.

Josh: In my opinion, this doesn't address the core concerns we have including scalability.

JM: the list could get longer and longer over time.

Michele Mottini: It works for us - and we have some fairly large deployments

Michele Mottini: (fairly large = tens of million patients, billion of data points)

Q1 ended.

Wednesday Q2

Chair: Josh Mandel

Minutes: Ewout Kramer

- Tracker item GF#19909 "clarity is needed for search on Period datatypes using date query parameter" - we still think all John wants is described. Non-persuasive. Chris Grenz/Christiaan Knaap: 7-0-0.
- Tracker item GF#21285 "Make lack of payload and re-query behavior ore visible" - persuasive: Christiaan Knaap/7-0-0
- Tracker item GF#21942 "Please provide an example of the error data file item response - BULKDATA #109" - we decide to rephrase the text so the confusing terms no longer appear (and there is already a good example): persuasive with mode: Ewout Kramer/Christiaan Knaap: 6-0-1.
- Tracker item GF#21003 "Backend services auth should use TLS 1.2 or better" - persuasive: Michael Donnelly/Nick Robinson: 7-0-0. Same vote and resolution for GF20996.
- Tracker item GF21068 "SMART Backend Services Auth Worked example should link to tagged artifacts" - persuasive: Ewout Kramer/Amy Ballard: 7-0-0.
- Tracker item GF21862 "Description of response/file/download Bundle - BULKDATA #77" - persuasive with mod: Nick Robinson/Christiaan Knaap: 7-0-0
- Tracker item GF21727 "Are JWS Algo's mandated or not? - BULKDATA #27" - persuasive: Michael Donnelly/Christiaan Knaap: 6-0-0

Wednesday Q4

Chair: Josh Mandel

Minutes: Michael Donnelly

How should JWK Set caching work?

https://gforge.hl7.org/gf/project/fhir/tracker/?action=TrackerItemEdit&tracker_item_id=21751

How should JWK Set caching work?

Michael Donnelly moved and @Nick Robison seconded that we will add:

- The client SHOULD return a "Cache-Control" header in its JWKS response
- The authorization server SHALL NOT cache a JWKS for longer than the client's cache-control header indicates.
- The authorization server SHOULD cache a client's JWK Set according to the client's cache-control header; it doesn't need to retrieve it anew every time.

Passed 6/0/2

Optionality of security layer weakens interoperability

https://gforge.hl7.org/gf/project/fhir/tracker/?action=TrackerItemEdit&tracker_item_id=21695&start=0

There are some implementations that use creative mechanisms to download the bulk data files once they're generated. The SMART Backend Services guide was developed with this use case in mind.

Currently we're only recommending this. If SMART Backend Services doesn't meet a systems needs, we should improve that guide.

CMS doesn't have a way currently to do Bulk FHIR with OAuth 2 that meets their security and engineering goals.

Would CMS provide a different security layer in addition to OAuth 2.0 or instead of it?

Instead.

Can CMS comment on the guide to try to enhance it to get it to meet their needs too?

Yes, that's why Nick is here this week.

Is Bulk Data different from the rest of FHIR? Is this more important here?

It's not more important, but this is an IG, so we need to get into more detail than the base spec needs.

Can we make a profile for the IG, where the profile defines the requirement to use OAuth 2.0 instead of the base IG?

Maybe? Seems confusing.

Could we say that servers either have to do OAuth 2.0 or CMS's delegated authorization scheme? And clients have to do whatever the server they're connecting to does?

What's CMS's issue?

The ACOs they work with aren't downloading the data directly; third parties are doing it for them, so they need a way to say who's allowed to get the data for them.

Could we say that Bulk FHIR requires SMART Backend Services, and later CMS can incorporate delegated authorization into SMART Backend Services?

People seem to feel okay about that.

CMS could come back later to change either Bulk FHIR or SMART.

Healthy discussion about how much this should be locked down.

ONC could specify a requirement to use the Bulk Data IG and the SMART Backend Services Guide.

In practice, the security layer will be negotiated between parties.

There was a motion to require SMART Backend Services that died for lack of a second.

@Nick Robison moved to find the ballot comment not persuasive, @Javier Espina seconded.

The motion passed with 3 votes for, 2 against, and 4 abstaining.

From *Robert's Rules Of Order Newly Revised In Brief* by Henry M. III Robert, William J. Evans, Daniel H. Honemann & Thomas J. Balch:

Do abstention votes count?

The phrase "abstention votes" is an oxymoron, an abstention being a refusal to vote. To abstain means to refrain from voting, and, as a consequence, there can be no such thing as an "abstention vote".

In the usual situation, where either a majority vote or a two-thirds vote is required, abstentions have absolutely no effect on the outcome of the vote since what is required is either a majority or two thirds of the votes cast. On the other hand, if the vote required is a majority or two thirds of the members present, or a majority or two thirds of the entire membership, an abstention will have the same effect as a "no" vote. Even in such a case, however, an abstention is not a vote.

Support encryption in bulk data

https://gforge.hl7.org/gf/project/fhir/tracker/?action=TrackerItemEdit&tracker_item_id=21051

CMS is working on per-client payload encryption.

Nobody else has had an issue with this yet; encryption in motion (TLS) is necessary but encryption at rest hasn't been needed for any other use cases.

Does FHIR support indicating that data are encrypted in the base FHIR spec?

No.

Should that be part of the Bulk FHIR spec?

That wouldn't drive implementers toward compatibility with the specification.

One thing that's tempting (but maybe not a good idea) would be to have an optional keyMap in the file to point at a descriptor about how to decrypt the file (e.g. a decryption algorithm and params).

CMS is happy to have conversation with the community about how to do this.

What was the threat assessment about this?

If someone gained access to the file system, unencrypted files could expose PHI.

If we're going to encrypt these files, we should have the opportunity to compress them first.

@Isaac Vetter moved to find the comment non-persuasive with mod and to revisit the topic when CMS brings back a proposal. @Chris Grenz seconded. The motion passed with 7 votes for, 0 against, and 2 abstaining.

Reconcile \$export vs \$everything

https://gforge.hl7.org/gf/project/fhir/tracker/?action=TrackerItemEdit&tracker_item_id=21050&start=0

\$export is more specific than \$everything (kind of)

\$export means

- async
- ndjson format
- a group of patients

\$everything means

- either one patient or **all** patients

One view: \$everything is for replicating a FHIR server. Converting from one system to another.

Does \$everything set an inaccurate expectation for clients? Would they be disappointed if they asked for "everything" and just got USCDI?

Michael moved to find this non-persuasive. @Adam Culbertson seconded.

The motion passed with 4 votes for, 1 against, and 4 abstaining.

We will keep the \$export name. Although the target data set is similar, the group sees different use cases for the operations.

Over time, \$everything has adopted aspects of the \$export operation. In the future, PA may want to more fully reconcile the \$everything operations with \$export.

Q4 ended.

Thursday Q2

Chair: Ewout Kramer (until 11:54) then Lloyd McKenzie

Notes: Rick Geimer (until 11:50) then Ewout Kramer

Agenda: Bulk Data Ballot Rec

[Combine group export overlap with the async workflow in R4 - BULKDATA #116](#)

[Combine patient export overlap with the async workflow in R4 - BULKDATA #115](#)

[Combine export overlap with the async workflow in R4 - BULKDATA #114](#)

- Proposal:
 - Update each operation to state, as part of the description: "The FHIR server MUST support invocation of this operation using the [FHIR Asynchronous Request Pattern](#)."
 - Update each operation parameters list. Continue to include the `_outputFormat` parameter inline to provide a holistic view, but condense the "Documentation," adding a reference to [FHIR async.html](#). "The format for the requested bulk data files to be generated, as per [FHIR Asynchronous Request Pattern](#). Defaults to `application/fhir+ndjson`."
- Motion: Josh Mandel/Michael Donnelly
 - Abstain: 1
 - Against 0
 - For: 10

[Clarify the scope of the authorization components and how it differs from the existing standard - BULKDATA #112](#)

- Proposal:
 - We note that the Backend Services authorization specification might eventually be hosted independently from the Bulk Data IG and is designed to work independently.
 - Before "Specifically, this profile describes the registration-time metadata required for a client to be pre-authorized, and the runtime process by which the client acquires an access token that can be used to retrieve FHIR resources.", add: "This specification handles use cases complementary to the SMART App Launch protocol[[link]]."
 - Include a note in <http://hl7.org/fhir/us/bulkdata/2019May/authorization/index.html> 1.0.2 "See Also"

The FHIR specification includes a set of [security considerations](#), including security, privacy, and access control. These considerations apply to diverse use cases and provide general guidance for choosing among security specifications for particular use cases.

- Motion: Josh Mandel/Rick Geimer
 - Abstain: 1
 - Against: 0
 - For: 12

[How do you specify the FHIR version to be used? - BULKDATA #58](#)

- Proposal:
 - Inform the commenter that clients should follow the regular pattern for determining a server's version as described at <http://build.fhir.org/versioning.html#cs-version>. We don't see special considerations for Bulk Data.
- Motion: Josh Mandel/Michael Donnelly
 - Abstain: 1
 - Against: 0
 - For: 12

[Incorrect input \(IN\) parameters are labeled as output \(OUT\) - BULKDATA #7](#)

- Proposal:
 - Update the table to document `_outputFormat` and `_type` as input parameters for each operation
- Motion: Josh Mandel/Chris Grentz
 - Abstain: 0
 - Against: 0
 - For: 13

[Clarify what date/time the `_since` flag uses. - BULKDATA #56](#)

- Proposal:
 - The `_since` flag is designed to filter based on when FHIR Resource state has changed, rather than clinical eventing times. While the export is based lastUpdated which is a less clinically meaningful time, it enables more complete data exchange. For servers pulling data in realtime from an external system, some evaluation would be required to determine when resource state had changed. We should add the following "Resources will be included in the response if their state has changed after the supplied time (e.g., if Resource.meta.lastUpdated is later than the supplied `_since` time)."
- Motion: Josh Mandel/Nick Robison
 - Abstain: 0
 - Against: 0
 - For: 14

[Can individual clinical provider credentials be used to initiate a bulk data export on all patients? - BULKDATA #55](#)

- Proposal:
 - We agree with the concern. The Bulk Data export specification recommends using Backend Services to protect bulk data, and that implementation guide does not allow end-user authorization of `$export`.
- Motion: Josh Mandel/Rick Geimer
 - Abstain: 0
 - Against: 0
 - For: 14

[A client is assigned an identifier. - BULKDATA #33](#)

- Proposal: We should write "SHALL register its public key" (i.e., delete "its client_id and" from this sentence).
- Motion: Issac Vetter/Josh Mandel
 - Abstain: 0
 - Against: 0
 - For: 14

[Healthcare orgs are already aware of their responsibility to protect PHI. Suggesting otherwise is silly. - BULKDATA #17](#)

- Proposal:
 - Adopt proposed wording: "Healthcare organizations have an imperative to protect PHI"
- Motion: Josh Mandel/Ewout Kramer
 - Abstain: 1
 - Against: 0
 - For: 13

[Transition to US Core FHIR IG - BULKDATA #51](#)

- Proposal:
 - Update "Common Clinical Data Set" to say "US Core Data for Interoperability"

- Update: "This use case exports all resources needed for the Common Clinical Data Set, as profiled by Argonaut. For a full list of these resources and profiles, see <http://www.fhir.org/guides/argonaut/r2/profiles.html>" to read: "This use case exports all resources needed for the US Core Data for Interoperability. For a full list of these resources and profiles, see <http://www.hl7.org/fhir/us/core/>."
- Motion: Josh Mandel/Rick Geimer
 - Abstain: 3
 - Against: 2
 - For: 9

Need to put actual ballot content into the ballot for github content - BULKDATA #50

- Proposal:
 - We should add a section at the bottom of the index page reproducing key items from the menu at the top (Exports, Backend Services, Operations, History)
- Motion: Josh Mandel/Michael Donnelly
 - Persuasive with mod
 - Abstain: 0
 - Against: 0
 - For: 12

IG should mandate that jku endpoint requires HTTPS, but not authnz. - BULKDATA #40

- Proposal:
 - We already state regarding the JWK Set URL, "This URL communicates the TLS-protected endpoint where the client's public JWK Set can be found." We should update this to add "This endpoint SHALL be accessible via TLS, without authentication or authorization." We should further update the Authentication JWT Header Values table, from "The URL to the JWK Set containing the public key(s)." to "The TLS-protected URL to the JWK Set containing the public key(s), accessible without authentication or authorization."
- Motion: Josh Mandel/Michael Donnelly
 - Persuasive
 - Abstain: 2
 - Against: 0
 - For: 10

Why is OAuth2 / mutual TLS required authorization - BULKDATA #96

- Proposal:
 - Replace: With each of the requests described herein, the client MUST provide proof of authorization. Implementers are encouraged to implement OAuth 2.0 access management in accordance with the [SMART Backend Services: Authorization Guide](#).
 - With: With each of the requests described herein, implementers are encouraged to implement OAuth 2.0 access management in accordance with the [SMART Backend Services: Authorization Guide](#).
- Motion: Josh Mandel/David Hay
 - Persuasive
 - Abstain: 1
 - Against: 0
 - For: 11

Bookmark to fine-grained permissions discussion

[Add fine grained requests and permissions for the _typeFilter - BULKDATA #85](#)

[Provide an example request with _typeFilter - BULKDATA #102](#)

[Does the inability to filter data potentially violate HIPAA's Minimum Necessary requirement? - BULKDATA #57](#)

[Provide an example request with _typeFilter - BULKDATA #102](#)

[Does _typeFilter support ALL FHIR query capabilities?](#)

[Suggest discussions with vendors trying to implement this feature and narrowing this down to what is viable prior to inclusion in IG. - BULKDATA #13](#)

[Drop _typeFilter parameter](#)

- Proposal:
 - We want to clarify that access control obligations can be metbet with a combination of in-band and out-of-band restrictions. (Example from CMS experience: some clients are authorized to access sensitive mental health information, and some aren't; this authorization is defined out of band, but when a client requests a full data set, filtering is automatically applied by the server.) While we wouldn't address HIPAA directly specification, we should add a note that "servers may limit the data returned to a specific client in accordance with local considerations (e.g., policies or regulations)."
 - The _typeFilter is labeled as experimental and optional, inside an IG at DSTU. Given these caveats
 - There ia rationale for keeping *something* in: a common question is "how would I filter this further?"
 - It's clearly labeled as experimental
 - We probably can't pin down the semantics further (e.g. "What happens with _include") -- so we should add a note in http://hl7.org/fhir/us/bulkdata/2019May/export/index.html#example-request-with-_typefilter specifically mentioning that we have not yet determined expectation for _include, _reinclude, or support for any specific search parameters.
- Motion: Ewout Kramer/David Hay
 - Abstain: 1
 - Against: 0
 - For: 10

Thursday Q3

Combine the async R4 details with this proposal due to mismatched parts of the specification - BULKDATA #97

- Proposal:
 - Agreed that there is substantial duplication with the FHIR async.html page. Today, the reason for this duplication is that content developed in the Bulk Data IG has been incorporated in FHIR, but there is not yet full alignment. We also have a number of requests for clarification, correction, and examples to these portions of the specification. We'd like to fix the issues direct in the Bulk Data IG, and then port them into the core FHIR Async.html page for future releases.
- Motion: Issac Vetter/Ewout Kramer
 - Abstain: 0
 - Against: 0
 - For: 11

Requesting dateTime instead of instant type for _since parameter - BULKDATA #4

- Proposal:
 - The _since parameter is aligned with the default FHIR spec (<https://www.hl7.org/fhir/http.html>) where _since requires an instant. Our motivation for this parameter in \$export is allowing a client to issue periodic data requests with a precise timestamp reflecting the time of the last fetch.
- Motion: Issac Vetter/Nick Robison
 - Abstain: 0
 - Against: 0
 - For: 11

Clarify if the use cases are using one or more specifications - BULKDATA #91

- Proposal:
 - The comment is about use cases, but the specification does not provide specific features or guidance at the use case level; as such, it's not clear how we could apply separate versioning.
 - Change: "These specifications are designed to support sharing any data that can be represented in FHIR. This means they should be useful for such diverse systems as:"
 - To: "This implementation guide is designed to support sharing any data that can be represented in FHIR. This means that the IG should be useful for such diverse systems as:"
- Motion: Rick Geimer/Abigail Watson
 - Abstain: 1
 - Against: 0
 - For: 10

Required access token lifetime is inconsistent in IG. - BULKDATA #36

- Proposal:
 - Strike "Each token issued under this profile MUST be short-lived, with an expiration time of no more than five minutes." since we say the same thing (but more clearly) two paragraphs above.
- Motion: Michael Donnelly/Rick Geimer
 - Abstain: 0
 - Against: 0
 - For: 11

Include an example of an extension in the example. - BULKDATA #111

- Proposal:
 - Add an example of an extension, like:
`{"extension": {"http://myserver.example.org/extra-property": true}}`
- Motion: Ewout Kramer/Rick Geimer
 - Abstain: 1
 - Against: 0
 - For: 10

Use a table format to specify content, type, optionality and description of the result body like CDS Hooks. - BULKDATA #107

- Proposal:
 - We should update our field definitions to use tables, like in CDS Hooks (e.g., <https://cds-hooks.org/specification/1.0/#suggestion>). The relevant columns are: Field, Optionality, Type, Description
- Motion: Rick Geimer/Nick Robison
 - Abstain: 0
 - Against: 0
 - For: 11

Is it mandatory behavior to support the bulk data files? - BULKDATA #106

- Proposal:
 - Change: "Note that the files-for-download MAY be served by a file server other than a FHIR-specific server."
 - To: "The files-for-download SHALL be accessible to the client at the URLs advertised. These URLs MAY be served by file servers other than a FHIR-specific server."
- Motion: Nick Robison/Michael Donnelly
 - Abstain: 0
 - Against: 0
 - For: 11

How do you know what subset of fhir resources are returned - BULKDATA #101

- Proposal:
 - As an answer to the question: the base expectations are identified as usual in FHIR, through the metadata conformance statement; additional constraints would be communicated out of band (e.g., at client registration time, based on client permissions).
 - Update the specification to state that "If the client explicitly asks for export of resources that the bulk data server doesn't support, the server SHOULD return details via OperationOutcome in an error response to the request,
- Motion: Nick Robison/Michael Donnelly
 - Abstain: 1
 - Against: 0
 - For: 10

Include the mime type application/fhir+json - BULKDATA #105

- Proposal:
 - Update the error response to say "Content-Type header of application/json" (just like the success response)
- Motion: Nick Robison/Michael Donnelly
 - Abstain: 0
 - Against: 0
 - For: 11

When a 200 OK is returned can the Response error be empty? - BULKDATA #104

- Proposal:
 - As the final bullet of <http://hl7.org/fhir/us/bulkdata/2019May/export/index.html#response---error-status-1>, add: "The choice of when to determine that an export job has failed in its entirety (error status), vs returning a partial success (complete status) is left up to the implementer."
- Motion: Nick Robison/Michael Donnelly
 - Abstain: 1
 - Against: 0
 - For: 10

Clarify if it is a "client-to-server" or "client-to-client" certificate - BULKDATA #47

Also: https://gforge.hl7.org/gf/project/fhir/tracker/?action=TrackerItemEdit&tracker_id=677&tracker_item_id=21739

- Proposal:
 - Change: Before a SMART client can run against a FHIR server, the client SHALL obtain a client-to-client digital certificate and SHALL register with that FHIR server's authorization service.
 - To: Before a SMART client can run against a FHIR server, the client SHALL generate or obtain an asymmetric keypair and SHALL register its public key set with that FHIR server's authorization service.
- Motion: Nick Robison/John Moerkhe
 - Abstain: 0
 - Against: 0
 - For: 9