Attaching Genomic Files

For use cases in genomics where a deeper level of data than provided by the variant profile is warranted there are specifications such as VCF, BAM, CRAM, and MAF. Best practices in exchanging these files along with the metadata necessary to make use of them through a FHIR API remains a complicated and open issue.

While in future versions of FHIR other resources or options might be preferred, DocumentReference is a good candidate as a vehicle to provide access to genomic data formats such as VCF, BAM, CRAM and MAF. Attachments that are individually useful (such as clinical notes) are often sent in FHIR using the DocumentReference resource, which allows the file and its captured metadata (e.g. within the DocumentReference 'context' element) to be discovered via search queries. Importantly, DocumentReference can link to a DiagnosticReport or other structures in this IG as needed via `DocumentReference.context.related`.

When sending genomic files there are many considerations. For example, it is not unusual for files to be gigabytes in size. If embedded directly, servers receiving the files may have size constraints per resource or per transaction which may limit your options. Instead of sending a large data file, the file can be referenced by a URL and title using the `DocumentReference.content.attachment.url` element. This can point to an online resource that hosts the file or from where the file can be accessed. For genomic files, the host is likely not the FHIR server providing the DocumentReference data instance. Be aware that use of DocumentReference to provide access to files through URLs introduces authorization requirements that are out of scope of this Implementation Guide. See http://hl7.org/fhir/datatypes.html#Attachment for more information on the HL7 FHIR attachment data type.

Many facets of how the genomic files were generated will be needed for receivers to make use of these files, such as what pipelines, tools, and settings were used. The intended downstream use case must be carefully evaluated to ensure appropriate file preparation. The `DocumentReference.description` might be helpful for sending systems to provide guidance on how the file was generated. However, a fully computable approach has yet to be described.

In a future release, this IG may include a profile of `DocumentReference` or other artifacts to provide more specific guidance.