
Lightning Documentation

Release 0.1.0

Curoverse

October 05, 2015

1	Getting started	3
1.1	Lightning at a Glance	3
1.2	Importing a Genome	5
1.3	Lightning Design Doc	5
1.4	Installation	7
2	Tiling Overview	9
3	Compact Genome File (CGF) Format	11
4	Lantern Specifications	13
5	Tile Library	15
6	Annotile: Annotating Tile Variants	17
6.1	Adding an Annotation Pipeline to Annotile	17
6.2	How does Annotile Work?	17
6.3	Storing Annotation Details	18
6.4	Future Annotile Functionality	18
7	Data Structures Specifications	19
7.1	Data Structures Specifications, v0.1.0	19
7.2	Data Structures Specifications, v0.1.1	26
8	Lightning API Specifications	35
8.1	Alternate Lightning v0.1.0 API Specifications	35
8.2	Batch Processing	50
8.3	Reasoning Behind the API	53
8.4	Lightning Errors	53
8.5	Lightning v0.1.1 API Specifications	54
8.6	Versioning	63
8.7	Paging	63
9	Software Development Kits	65
10	Sprite	67
11	Indices and tables	69

Contents:

GETTING STARTED

Welcome to Lightning! This documentation section is most relevant for first time users. [Lightning at a Glance](#) provides an overview of Lightning: its uses and components. How-to guides and tutorials are expected to be placed here as the project grows.

1.1 Lightning at a Glance

Lightning is software designed to enable fast queries and machine learning on genomic data. The genomic data we currently focus on are human whole genomes that are aligned and are called by external software, then imported into Lightning. From there, Lightning:

- Stores quality information
- Stores phased and unphased genomes
- Allows fast retrieval of called sequences from regions of interest
- Defines flexible queries:
 - Filters by subsets of the population
 - And/or by specific regions of interest
- Normalizes standard called genome files (such as VCF and gVCF), such that each variant is expressed in the same way
- Incorporates new data fast and painlessly
- Stores annotations from ClinVar and annotation pipelines, such as CAVA

Lightning is made possible by the process of *tiling*, which takes advantage of the high degree of redundancy in a population of genomes. Tiling partitions genomes into *tiles*: overlapping, variable-length sequences that begin and end with unique *k*-mers, termed *tags*. Once a genome has been tiled, the sequences for each tile are stored in a *tile library*. These sequences may be annotated by using *Annotile*. The tiled genomes are stored in *Compact Genome Format* (CGF) files. Genomes stored as CGF files are loaded into *Lantern*, which is our in-memory database designed to respond to queries quickly. Finally, *Sprite* is a web browser application for interacting with Lightning.

Note: Stated another way, Lightning's basic method is to consider short snippets of genomic sequences as the basic building block of genomes. These short snippets are of variable length, but are mostly in the range of 250 base pairs long. Splitting genomes into short segments allows for savings by only storing a single copy of redundant sequences.

Each genome is partitioned into these these short read segments. From all *tiles* in a population, a *tile library* can be constructed. Tiles are chosen to have 24mer tags on either end that overlap with neighboring tiles. *Tags* are chosen with with some uniqueness constraint on them and provide convenient anchor points to differentiate tiles from one another.

Currently, all tags are chosen to be at least 2 edit distance away from each other. The tag set is fixed and acts as anchor points to partition future sample genomic sequences wishing to be analyzed.

The hope is that tiles, along with information on the population used to generate them, can also be used to aid in read placement.

Because most genomic sequences are redundant, duplicate tiles need not be stored in a population of genomic sequences. At each tile position, multiple tile variants are stored representing the variation in a population for that tile. Given a partitioned genomic sequence and a tile library, a compact representation of a genome can be constructed by storing the variant numbers contiguously.

1.1.1 Components

- Tiling
- Compact Genome Format (CGF)
- Lantern
- Tile Library
- Annotile
- Sprite
- Software development kits
- API

1.1.2 Motivation

We developed Lightning in response to the difficulty and time-consuming nature of merging VCFs, querying subsets of a population, finding poorly sequenced regions, and similar issues. After using various ad-hoc solutions, we eventually stepped back and committed time and effort to developing a more sensible and sustainable solution. We hope it will be useful to the broader research community and welcome your feedback.

1.1.3 Usage

We host a Lightning instance for public whole genomes, including 502 genomes from the 1000 Genomes Project and 178 genomes from the Harvard Personal Genome Project. Here, the public can interact with Lightning using Sprite.

Todo

Decide where the public lightning instance will be

Currently, Lightning assumes it has access to an Arvados cluster. For a user to run their own instance of Lightning on their data, they must set up an Arvados cluster and import their data into it. From there, they should follow the instructions in [Importing Genomes](#) to add new genomes. To add new annotations, they should go to [Annotile](#).

1.1.4 Further research

Lightning is currently in its infancy. This state means Lightning has many use cases and future directions that can be explored, some of which are listed here.

- Implement Lightning Servers for multiple species

- Include RNA sequences in Lightning

Todo

Keep filling out list

1.1.5 Contributors

- Abram Connelly
- Sarah (Sally) Guthrie
- Jiayong Li
- Nancy Ouyang
- Alexander (Sasha) Wait Zaranek

1.1.6 Contact

lightning@curoverse.com

Todo

Make sure this email is live

1.2 Importing a Genome

Importing a genome is difficult. We'll probably do it using pipelines. Have one pipeline in pipeline_templates that is run when someone uploads a genome.

We acknowledge that phenotypes or meta-information about each genome is vital for interpreting results. However, the breadth of available phenotype databases and the number of possible pieces of information makes using Lightning as a phenotype database infeasible. Sprite provides a preliminary phenotype database, which includes information about whether each genome (or set of calls, also known as a callset) is a reference. Additionally, it stores whether the phases are well known, as well as the sex and the ethnicity.

Todo

Implement a better process for importing a genome Document the process

1.3 Lightning Design Doc

Goal: Specify Lightning v0.1.0 in enough detail that it can be implemented and tested.

1.3.1 Lightning v0.1.0 Features

Features here are expected to only work for well-sequenced tile variants. Poorly sequenced tile variants/partial tile variants should return a null value.

- Stores 1 byte of quality information for all bases (well sequenced versus poorly sequenced)

- Stores human genomes assuming phasing from called genome input
- Retrieves well-sequenced subsequences from population (subsequences identified by tile position)
 - Subsequences may be in the form of bases
 - Subsequences may be in the form of tiles
- Allows filtering
 - by population subsets (named by id)
 - by regions of interest (named by tile)
 - by tile variant (named by md5sum, only well-sequenced tiles)
 - by combinations of the above
- Provides translation between tile positions and reference locations (GRCh37, GRCh38)
- Tiles genomes provided in GFF and gVCF formats
- Can convert tiled genomes to GFF and/or gVCF
- One pipeline is required to add a new genome
- Stores annotations from ClinVar and annotation pipelines, such as CAVA
- Retrieves annotations for one or multiple tile variants (identified by tile variant identifier)
- Retrieves tile variants with one or multiple annotations (identified by annotation identifier)

1.3.2 Lightning v0.1.0 Benchmarking Goals

- Uses up to 70 MB per human genome
- Retrieves the sequences of 50 contiguous tiles from 100 human genomes in 5 seconds
- Retrieves the 50 contiguous tiles from 100 human genomes in less than a second
- Filtering on above query with any option does not change the observable time
- Tiling genomes from GFF/gVCF takes 1 hour per genome
- Annotates the tile library for 100 human genomes in 2 hours
- Filtering on annotations or tile variants for tile library of 100 genomes takes less than 1 second

1.3.3 Lightning v0.1.0 Usage

Lightning v0.1.0 should be run by Curoverse and should include 178 whole genomes from the Personal Genome Project and 502 whole genomes from the 1000 Genomes Project. The tile library should be annotated with CAVA, ClinVar, and ExAC data. The public should be able to retrieve well-sequenced subsequences from this population, filter on them by population subsets, by regions of interest, by tile variant, and/or by annotations.

Users should be able to upload whole genomes, with the understanding that this Lightning instance is public. These newly added genomes should be added to the tile library, new sequences should be annotated with CAVA, ClinVar, and ExAC annotations, and these genomes should be added to those we can filter on.

1.3.4 Lightning v0.1.0 API

See [Lightning v0.1.0 API Specifications](#) for API details.

1.3.5 Future Versions Lightning Features

- Support partial tile variants for features detailed in v0.1.0
- Stores unphased and phased human genomes, including phase groups for features detailed in v0.1.0
- Provides translation between tiles and alternate reference locations
- Compare multiple genomes with each other

1.3.6 Future Versions Lightning Usage

Users should set up an Arvados cluster. From there, they should be able to run Sprite, which will allow them to import their data without making it public. From there, users should be able to use Sprite identically to the v0.1.0 Lightning instance.

1.4 Installation

Since Lightning is in development, our installation process has not yet been determined, implemented, or documented.

Todo

Create a good installation process and document it

TILING OVERVIEW

Lightning is made possible by the process of `tiling`, which takes advantage of the high degree of redundancy in a population of genomes. Tiling partitions genomes into `tiles`: overlapping, variable-length sequences that begin and end with unique `k`-mers, termed `tags`. This document is an overview of our tiling implementation, along with the crunch scripts and pipeline templates available for tiling genome inputs.

We currently accept GFF files and Complete Genomics CGI-var files as inputs.

Todo

- Document PASTA/FASTA
 - Write basic instructions for writing one's own conversion of file to a tiling (As I understand it, one must write a conversion between that file format and PASTA, then probably add a pipeline template component).
-

COMPACT GENOME FILE (CGF) FORMAT

Compact Genome File Format is our preliminary format that stores tiled genomes compactly.

Todo

Document details about CGF format

LANTERN SPECIFICATIONS

Lantern is our in-memory database which stores individual tiled genome sequences. It allows us to make very fast comparisons between genomes and filter on genome sequences. Since it is contained in random access memory (RAM), queries are very fast.

Todo

Document Lantern and the REST APIs used to interact with Lantern

TILE LIBRARY

The tile library is our in-memory database that stores tile variants and their information. Tile variants are associated with their sequence, MD5 hash digest, tile position, and other meta-information. The tile library stores this information compactly while enabling queries on it.

Todo

Describe/specify the queries the Tile Library supports

ANNOTILE: ANNOTATING TILE VARIANTS

6.1 Adding an Annotation Pipeline to Annotile

Build a component for the *annotate_tile_variants* pipeline. This is likely to require implementing a crunch script and creating a docker image.

6.1.1 Automated Annotation Pipelines

If the annotation pipeline you wish to port cannot use a VCF file (aligned against GRCh37) as input, you will also need to implement a component to generate the required input for your annotation pipeline. Use the *create-vcf-per-tile* component as an example.

The components you implement must output a collection with a directory named *annotile_input*.

Todo

After implementing annotile, document the output format requirements

6.1.2 User-added Annotations

User added annotations can be converted into a component in the *annotate_tile_variants* pipeline by hard-coding annotations applying to specific tile variants. A component which does this already is linked to a Sprite user-added annotations app. Use this component and application as an example.

Annotating genomes is an essential piece of bioinformatic workflows. Annotations are produced by many sources, including programs and users, and we want to ensure Lightning supports the addition of all these annotations to tile variants, can add software-produced annotations to new tile variants automatically, and can respond to queries about which tile variants have which annotations or which annotations apply to particular tile variants.

Despite the importance of annotations in bioinformatics, Lightning is not meant to be an annotation database - it is not meant to support complex queries on types of annotations, source of annotations, date of generation or modification, or keywords. These types of queries and annotation storage should be done by add-ons, which can be added to Sprite using django apps or can be added to one's own application interacting with Lightning.

6.2 How does Annotile Work?

Lightning associates each tile variant with its `variant` value, which is our term for the MD5 hash digest of its sequence. These `variant` values can be associated with metadata about that tile variant. Annotile is simply a many-to-many database, which associates variant values to user-specified annotations identifiers.

Warning: Lightning v0.1.0 does not support annotating tiles with poorly sequenced regions. These require more complex tile variant representations and are thus left to future versions.

Annotile is populated by running the pipeline template *annotate_tile_variants*. Each component in *annotate_tile_variants* runs annotation software on each tile variant, producing detailed annotation information for the annotation app and information for loading into Annotile.

For details about porting an annotation pipeline or adding user-written annotations, see [Adding an Annotation Pipeline to Annotile](#).

Annotile also provides querying capabilities, allowing a user to find which annotations are associated with which tile variants and vice versa. For specific details about annotation querying, see *../api/index*.

6.3 Storing Annotation Details

[Sprite](#) has a basic app for storing and visualizing specific annotations, but it is not the only option for storing and querying annotations, since we are designing Lightning to be able to support new annotation databases. A user can build a django app to plug into Sprite on their Lightning instance, or a user can build their own application in their language of choice to work similarly to Sprite, then plug their annotation database into that.

6.4 Future Annotile Functionality

Future functionality is expected to focus on Annotile's querying capabilities, and will probably include a direct querying function from genomes to annotations and vice versa.

DATA STRUCTURES SPECIFICATIONS

Contents:

7.1 Data Structures Specifications, v0.1.0

7.1.1 ArvadosUUID

The UUID of an Arvados object.

```
'<string>-<string>-<string>'
```

7.1.2 TilePosition

A compact representation of a tile position. A string of 3 period-separated integers (in base 16). The first integer is the tag set version integer, the second is the path number, the third is the step.

```
'<int>.<int>.<int>'
```

```
Version.Path.Step
```

7.1.3 TilePositionRange

A compact representation of a range of tile positions. A string of 3 period-separated integers (in base 16) followed by a hyphen and another integer in base 16. The first integer is the tag set version integer, the second is the path integer, the third is the first step integer, which is the step to start retrieving from (inclusive and 0-indexed). The integer following the hyphen indicates the step to stop retrieving from (**exclusive** and 0-indexed).

```
'<int>.<int>.<int>-<int>'
```

```
Version.Path.StartStep-EndStep  
  StartStep: inclusive, 0-indexed  
  EndStep: exclusive, 0-indexed
```

7.1.4 TileVariant

A compact representation of a tile variant. A string of 3 period-separated integers (in base 16), followed by another period and one string. The first integer is the tag set version integer, the second is the path integer, the third is the step integer, and the fourth is the MD5 hash digest of the tile variant sequence.

```
'<int>.<int>.<int>.<string>'
```

```
Version.Path.Step.VariantMD5SUM
```

7.1.5 NotTileVariant

Only used to build *TileVariantClause*. When used, it indicates a selection on the specimens that do not have that tile variant. A *TileVariant* (in base 16), preceded by a tilda.

```
'~<int>.<int>.<int>.<string>'
```

```
Version.Path.Step.VariantMD5SUM
```

7.1.6 ClauseEntry

Used as a builder for *TileVariantClause*'s. Items may be of type *TileVariant* or *NotTileVariant*. If the item is of type *TileVariant*, the *ClauseEntry* evaluates to true if that *TileVariant* exists in the population. If the item is of type *NotTileVariant*, the *ClauseEntry* evaluates to true if the *TileVariant* following the tilda does *not* exist in the population.

```
<TileVariant> | <NotTileVariant>
```

7.1.7 TileVariantClause

The 'OR' list of SAT (Boolean Satisfiability Problem). Ensuring the validity of the clause is a task placed on the client. A list of length 1 or more. Each item in the list is of type *ClauseEntry* and is checked against each phase independently of the other items. The *TileVariantClause* evaluates to true if *any* of the clauses evaluate to true.

```
[ <ClauseEntry>, <ClauseEntry>, ... ]
```

Example using default behavior:

[0.0.0. md5sum1 , ~0.0.1. md5sum2] will evaluate to True for any genomes in the population that have tile variant md5sum1 at *TilePosition* 0.0.0 on at least one of their phases and/or do not have tile variant md5sum2 at *TilePosition* 0.0.1 on at least one of their phases.

7.1.8 TileVariantLogic

The 'AND' list of SAT (Boolean Satisfiability problem). Ensuring the validity of the clause is a task placed on the client. A list of one or more *TileVariantClause*'s. By default, each *TileVariantClause* is evaluated against each phase independently of the other *TileVariantClause*'s.

```
[<TileVariantClause>, <TileVariantClause>, ... ]
```

Examples:

[[0.0.0. md5sum1 , ~0.0.1. md5sum2]] will evaluate to True for any genomes in the population that have tile variant md5sum1 at *TilePosition* 0.0.0 on at least one of their phases and/or do not have variant md5sum2 at *TilePosition* 0.0.1 on at least one of their phases.

[[0.0.0. md5sum1], [0.0.1. md5sum2]] will evaluate to True for any genomes in the population that have tile variant md5sum1 at *TilePosition* 0.0.0 on at least one of their phases *and* do not have variant md5sum2 at *TilePosition* 0.0.1 on at least one of their phases.

7.1.9 TileVariantDetail

The metadata information associated with one well sequenced tile variant. Dictionary containing the keys:

- `tile-variant`: The *TileVariant* identifier.
- `tag-length`: The length of the tags.
- `start-tag`: The sequence of the start tag. Must be of length 0 (if the tile is at the start of the path) or the length specified by `tag-length`. Cannot have n's.
- `end-tag`: The sequence of the end tag. Must be of length 0 (if the tile is at the end of the path) or the length specified by `tag-length`. Cannot have n's.
- `is-start-of-path`: A boolean indicating whether the tile is at the start of the path.
- `is-end-of-path`: A boolean indicating whether the tile is at the end of the path.
- `sequence`: The sequence of the tile. Cannot include n's (since the tile variant must be well-sequenced).
- `md5sum`: The md5sum of the tile sequence.
- `length`: The length of the tile sequence.
- `number-of-positions-spanned`: The number of tile positions this tile variant spans. Must be greater or equal to 1.
- `population-frequency`: The percentage of the population that contains this tile variant. Each well-sequenced phase counts as 1 entry in the population.
- `population-count`: The number of well-sequenced phases that contain this tile variant.
- `population-total`: The number of phases that contain a well sequenced tile at this tile position. Each phase counts as 1 entry in the population.

```
{
  'tag-length': <int>,
  'start-tag': <string>,
  'end-tag': <string>,
  'is-start-of-path': <boolean>,
  'is-end-of-path': <boolean>,
  'sequence' : <string>,
  'md5sum': <string>,
  'length': <int>,
  'number-of-positions-spanned': <int>,
  'population-frequency': <float>,
  'population-count': <int>,
  'population-total': <int>
}
```

Validations for each key are as follows:

```
TileVariantDetail['tag-length'] >= 1
TAG_LENGTH = TileVariantDetail['tag-length']
TileVariantDetail['start-tag'] matches '^[acgt]{TAG_LENGTH}$|^$'
TileVariantDetail['end-tag'] matches '^[acgt]{TAG_LENGTH}$|^$'
TileVariantDetail['sequence'] matches '[acgt]'
```

```
TileVariantDetail['md5sum'] == MD5_hash_digest(TileVariantDetail['sequence'])
TileVariantDetail['length'] == len(TileVariantDetail['sequence'])
TileVariantDetail['number-of-positions-spanned'] >= 1
0 <= TileVariantDetail['population-frequency'] <= 1
TileVariantDetail['population-total'] >= 0
```

7.1.10 Assembly

A description of an assembly (JSON-formatted). Keys:

- `assembly-name`: the assembly name (string)
- `assembly-pdh`: the portable data hash referencing the collection of FASTA files for this assembly (string)

```
{
  'assembly-name': <string>,
  'assembly-pdh': <string>,
}
```

7.1.11 Locus

A description of an assembly locus (JSON-formatted). Keys:

- `assembly-name`: the assembly name (string)
- `assembly-pdh`: the portable data hash referencing the collection of FASTA files for this assembly (string)
- `chromosome-name`: the chromosome name (string)
- `indexing`: Indicates the indexing of `start-position` and `end-position`. (hard-coded to 0).
- `start-position`: start position; the inclusive beginning of the loci range. Must be greater than 0 and less than `end-position`. Inclusive. (Integer)
- `end-position`: end position; the exclusive end of the loci range. Must be greater than `start-position` and less than or equal to the length of the chromosome in the specified assembly.

```
{
  'assembly-name': <string>,
  'assembly-pdh': <string>,
  'chromosome-name': <string>,
  'indexing': 0,
  'start-position': <int>,
  'end-position': <int>
}
```

7.1.12 CMPFunction

List of supported comparison functions.

```
ENUM('eq', 'lt', 'lte', 'gt', 'gte')
```

7.1.13 CMPTuple

Representation of an integer comparison - used for filtering queries. Also includes a range comparison. For the range comparison, the first integer is the inclusive beginning of the range, the second int is the exclusive end of the range.

```
( <CMP-fn>, <float>)  
OR  
( 'range', <float>, <float>)
```

7.1.14 VCF Data Structures Specifications

gVCFBlock

Genotype quality ranges used for banding. List of integers (length greater than or equal to 2). Each pair of integers creates a minGQ (inclusive) and maxGQ (exclusive) pair for a gVCFBlock. Must be strictly increasing, the first entry must be 0, and the last entry must be 2147483647 (the largest unsigned integer representable in 32 bits).

```
[0, <int>, ..., 2147483647]
```

gVCFMetaData

Representation of a gVCF Header. JSON-formatted with keys:

- `fileformat`, whose value indicates the format of the VCFLines returned by the Lightning server
- `fileDate`, the current date, format (YYYYMMDD)
- `source`, the Lightning server producing the VCF MetaData
- `assembly`, *Assembly* data type containing the location of the reference FASTA file used to generate the VCF lines.
- `info`, used to indicate the end of a VCFLine.
- `format`, for genotype fields
- `alt`, used to refer to non-reference alternate alleles.
- [optional] `gvcfblock`, for information about splitting gVCF blocks

```
{  
  'fileformat': 'VCFc4.2',  
  'fileDate': <int>,  
  'source': 'Lightningv0.1.0',  
  'assembly': <Assembly>,  
  'info': [  
    {  
      'ID': 'END',  
      'Number': 1,  
      'Type': Integer,  
      'Description': 'Stop position of the interval'  
    }  
  ],  
  'format': [  
    {  
      'ID': 'GT',
```

```
        'Number':1,
        'Type':'String',
        'Description':'Genotype'
    }
],
'alt': [
    {
        'ID': 'NOT_REF',
        'Description':'Represents any possible alternative allele at this location'
    }
],
'gvcfblock': gVCFBlock_
}
```

VCFMetaData

Representation of a VCF Header. JSON-formatted with keys:

- `fileformat`, whose value indicates the format of the VCFLines returned by the Lightning server
- `fileDate`, the current date, format (YYYYMMDD)
- `source`, the Lightning server producing the VCF MetaData
- `assembly`, *Assembly* data type containing the location of the reference FASTA file used to generate the VCF lines.
- `format`, for genotype fields

```
{
  'fileformat':'VCFc4.2',
  'fileDate':<int>,
  'source':'Lightningv0.1.0',
  'assembly':<Assembly>,
  'format': [
    {
      'ID': 'GT',
      'Number':1,
      'Type':'String',
      'Description':'Genotype'
    }
  ]
}
```

VCFSampleFormatData

JSON-formatted format field format for a VCF or gVCF file. Used for Genotype fields. Keys:

- `sample-name`, the name of the sample, normally defined by the column header
- `GT`, the genotype of the sample

```
{
  'sample-name': <string>,
  'GT':<string>
}
```

VCFLine

Representation of VCF line. Keys:

- `chrom`: Chromosome. An identifier from the reference genome or an angle-bracketed ID String (<ID>) pointing to a contig in the assembly file (the file pointed to by `assembly` in *VCFMetaData*). The colon symbol (:) must be absent from all chromosome names to avoid parsing errors when dealing with breakends. (String, no white-space permitted).
- `pos`: Position. The reference position, 1-indexed. Telomeres are indicated by using positions 0 or N+1, where N is the length of the corresponding chromosome or contig. (Integer).
- `ref`: Reference base(s). Each base must be one of A,C,G,T,N (case insensitive). Multiple bases are permitted. The value matching the `pos` key refers to the position of the first base in this string. For simple insertions and deletions in which either the `ref` or one of the `alt` alleles would otherwise be null/empty, the `ref` and `alt` values must include the base before the event (which must be reflected in the `pos` field), unless the event occurs at position 1 on the contig in which case it must include the base after the event; this padding base is not required (although it is permitted) for variations such as complex substitutions or other events where all alleles have at least one base represented in their strings. If any of the `alt` alleles is a symbolic allele (an angle-bracketed ID String <ID>), then the padding base is required and `pos` denotes the coordinate of the base preceding the polymorphism. (String)
- `alt`: Alternate base(s). List of alternate non-reference alleles called on at least one of the samples. Options are strings made up of the bases A,C,G,T,N,* (case insensitive) or an angle-bracketed ID String (<ID>) or a breakend replacement string as described in the VCFv4.2 section on breakends. The * allele is reserved to indicate that the allele is missing due to an upstream deletion. If there are no alternative alleles, the list should be empty. (String; no whitespace, commas, or angle-brackets are permitted in the ID String itself).
- `filter`: Filter status. A list of length 1 with value *PASS* if this position has passed all filters (if a call is made at this position). Otherwise, if the site has not passed all filters, a list of codes for filters that fail. [*q10*, *s50*] might indicate that at this site the quality is below 10 and the number of samples with data is below 50% of the total number of samples. 0 is reserved and should not be used as a filter string. If filters have not been applied, then the list should be empty. (List of strings with no white-space or semi-colons permitted).
- `format`: Genotype information (JSON-formatted). Key is 'GT'. Values are a list of strings with no white-space, semi-colons, commas or equals-signs.

```
{
  'chrom':<string>,
  'pos':<int>,
  'ref':<string>,
  'alt': [<string>, ...],
  'filter': [<string>, ...],
  'format': [VCFSampleFormatData_, ...]
}
```

gVCFLine

Representation of gVCF line. Keys are the same as *VCFLine* with one additional key:

- `info`: Additional information (dictionary, optional) to indicate when a gVCF block ends. Only 1 valid key exists in this version: "END". Values are a list of integers. The list must be of length 1.

```
{
  'chrom':<string>,
  'pos':<int>,
  'ref':<string>,
  'alt': [<string>, ...],
  'filter': [<string>, ...],
  'format': [VCFSampleFormatData_, ...],
  'info': {"END": [<int>]}
}
```

7.2 Data Structures Specifications, v0.1.1

7.2.1 TilePosition

(No changes, see *TilePosition*.)

7.2.2 TilePositionRange

(No changes, see *TilePositionRange*.)

7.2.3 TileVariant

(No changes, see *TileVariant*.)

7.2.4 NotTileVariant

(No changes, see *NotTileVariant*.)

7.2.5 ClauseEntry

(No changes, see *ClauseEntry*.)

7.2.6 TileVariantClause

(No changes, see *TileVariantClause*.)

7.2.7 TileVariantLogic

(No changes, see *TileVariantLogic*.)

7.2.8 TileVariantDetail

Todo

Specify md5sum for poorly sequenced tiles Different frequencies for more information - how many are well sequenced? How many are there, regardless of sequencing? Etc

The metadata information associated with one tile variant. Dictionary containing the keys:

- `tag-length`: The length of the tags.
- `start-tag`: The sequence of the start tag. Must be of length 0 (if the tile is at the start of the path) or the length specified by `tag-length`. Cannot have n's.
- `end-tag`: The sequence of the end tag. Must be of length 0 (if the tile is at the end of the path) or the length specified by `tag-length`. Cannot have n's.
- `is-start-of-path`: A boolean indicating whether the tile is at the start of the path.
- `is-end-of-path`: A boolean indicating whether the tile is at the end of the path.
- `sequence`: The sequence of the tile. May include n's.
- `md5sum`: TODO.
- `length`: The length of the tile sequence.
- `number-of-positions-spanned`: The number of tile positions this tile variant spans. Must be greater or equal to 1.
- `population-frequency`: The percentage of the population that contains this tile variant. Each phase counts as 1 entry in the population.
- `population-total`: The number of phases that contain a tile at this tile position. Each phase counts as 1 entry in the population.

```
{
  'tag-length': <int>,
  'start-tag': <string>,
  'end-tag': <string>,
  'is-start-of-path': <boolean>,
  'is-end-of-path': <boolean>,
  'sequence' : <string>,
  'md5sum': <string>,
  'length': <int>,
  'number-of-positions-spanned': <int>,
  'population-frequency': <float>,
  'population-total': <int>
}
```

Validations for each key are as follows:

```
TileVariantDetail['tag-length'] >= 1
TAG_LENGTH = TileVariantDetail['tag-length']
TileVariantDetail['start-tag'] matches '^[acgt]{TAG_LENGTH}$|^$'
TileVariantDetail['end-tag'] matches '^[acgt]{TAG_LENGTH}$|^$'
TileVariantDetail['sequence'] matches '[acgtn]'
TileVariantDetail['md5sum'] == MD5_hash_digest(TileVariantDetail['sequence'])
TileVariantDetail['length'] == len(TileVariantDetail['sequence'])
TileVariantDetail['number-of-positions-spanned'] >= 1
0 <= TileVariantDetail['population-frequency'] <= 1
TileVariantDetail['population-total'] >= 0
```

7.2.9 Locus

(No changes, see *Locus*)

7.2.10 CMPTuple

(No changes, see *CMPTuple*)

7.2.11 CMPFunction

(No changes, see *CMPFunction*)

7.2.12 VCF Data Structures Specifications

VCFInfoField

JSON-formatted information field format for a VCF or gVCF file. Keys that are not defined here may be included. Required keys are:

- 'ID', the name of the info field.
- 'Number', indicates the number of values that can be included with the INFO field. If the field has one value per alternate allele, the value should be 'A'. If the field has one value for each possible allele (including the reference), the value should be 'R'. If the field has one value for each possible genotype, the value should be 'G'. If the number of possible values varies, is unknown, or is unbounded, the value should be '.'.
- 'Type', options are 'Integer', 'Float', 'Flag', 'Character', and 'String'.
- 'Description', the description of the info field.
- [optional] 'Source'
- [optional] 'Version'

If the 'Type' is equal to 'Flag', the INFO field does not contain a Value entry and the number should be 0.

```
{
  'ID': <string>,
  'Number':<int>|'A'|'R'|'G'|'.',
  'Type':<string>,
  'Description':<string>,
  'Source':<string>,
  'Version':<string>
}
```

VCFFilterField

JSON-formatted filter field format for a VCF or gVCF file. Keys that are not defined here may be included. Required keys are 'ID' and 'Description'.

```
{
  'ID': <string>,
  'Description':<string>
}
```

VCFFormatField

JSON-formatted format field format for a VCF or gVCF file. Used for Genotype fields. Keys that are not defined here may be included. Required keys are:

- 'ID'
- 'Number', indicates the number of values that can be included with the INFO field. If the field has one value per alternate allele, the value should be 'A'. If the field has one value for each possible allele (including the reference), the value should be 'R'. If the field has one value for each possible genotype, the value should be 'G'. If the number of possible values varies, is unknown, or is unbounded, the value should be '.'.
- 'Type', options are 'Integer', 'Float', 'Character', and 'String'
- 'Description'

```
{
  'ID': <string>,
  'Number':<int>|'A'|"R"|"G"|'.'.',
  'Type':<string>,
  'Description':<string>
}
```

VCFAAlternativeAlleleField

JSON-formatted alternative allele field format for a VCF or gVCF file. Used to build symbols for alternate alleles (used for imprecise structural variants). Keys that are not defined here may be included. Required keys are 'ID' and 'Description'. The ID field indicates the type of structural variant and can be a colon-separated list of types and subtypes. The ID values are case sensitive and may not contain whitespace or angle brackets. The first level type must be one of the following:

- DEL, deletion relative to the reference
- INS, insertion of novel sequence relative to the reference
- DUP, region of elevated copy number relative to the reference
- INV, inversion of reference sequence
- CNV, copy number variable region (may be both deletion and duplication; should not be used when a more specific category may be applied)

Reserved subtypes include:

- DUP:TANDEM, tandem duplication
- DEL:ME, deletion of a mobile element relative to the reference
- INS:ME, insertion of a mobile element relative to the reference

```
{
  'ID': <string>,
  'Description':<string>
}
```

VCFContigField

JSON-formatted contig field format for a VCF or gVCF file. Keys that are not defined here may be included. Required keys are 'ID' and 'Description', and 'URL'. 'URL' points to the location of the

contig.

```
{
  'ID': <string>,
  'Description':<string>,
  'URL':<string>
}
```

VCFSampleField

JSON-formatted sample field format for a VCF or gVCF file. Used to define sample to genome mappings. This is the only info given by the VCFv4.2 specifications. I believe the length of 'Genomes', 'Mixture', and 'Description' must be the same.

```
{
  'ID': <string>,
  'Genomes': [<string>, ...],
  'Mixture': [<string>, ...],
  'Description': [<string>, ...]
}
```

VCFMetaData

Representation of a VCF Header. JSON-formatted with keys:

- `fileformat`, whose value indicates the format of the VCFLines returned by the Lightning server
- `fileDate`, the current date, format (YYYYMMDD)
- `source`, the Lightning server producing the VCF MetaData
- `reference`, the location of the reference fasta file used to generate the VCFLines
- `assembly`, same as `reference`
- [optional] `info`, for information field formats
- [optional] `filter`, for filters that have been applied to the data
- [optional] `format`, for genotype fields
- [optional] `alt`, for symbolizing imprecise structural variants
- [optional] `contig`, for pointing to sequence contigs
- [optional] `sample`, for defining sample to genome mappings
- [optional] `pedigree`, for defining relationships between genomes. Can be a list of `name:genome` or a url pointing to a pedigree database

```
{
  'fileformat':<string>,
  'fileDate':<int>,
  'source':<string>,
  'reference':<string>,
  'assembly':<string>,
  'info': [VCFInfoField_, ...],
  'filter': [VCFFilterField_, ...],
  'format': [VCFFormatField_, ...],
  'alt': [VCFAlternativeAlleleField_, ...],
  'contig': [VCFContigField_, ...],
}
```

```
'sample': [VCFSampleField_, ...],
'pedigree': [{<string>:<string>}, ...] OR <string>
}
```

Example:

```
{
  'fileformat': 'VCFv4.2',
  'fileDate': 20150921,
  'source': 'Lightningv0.1.0',
  'reference': '1adbd1bd00358fe6ff2303ec8f3169ce+83454',
  'assembly': '1adbd1bd00358fe6ff2303ec8f3169ce+83454'
}
```

VCFSampleFormatData

JSON-formatted format field entry for a VCF or gVCF line. Used for Genotype fields of specific samples (identified by the `sample-name` key below). Further keys are the values associated with the 'ID' of *VCFFormatField*'s returned by *VCFMetaData* ['format']. Values are a list of strings with no white-space, semi-colons, commas or equals-signs. Useful predefined keys include:

- **GT** : genotype. The values associated with this key are strings of allele values for the specified sample in `sample-name`. The alleles are separated by / (for unphased) or | (for phased). The allele values are 0 for the reference allele (which is in the `ref` field of the *VCFLine*), 1 for the first allele listed in the `alt` field of the *VCFLine*, 2 for the second allele in the `alt` field of the *VCFLine* and so on. If a call cannot be made for a sample at a given locus, . should be specified for each missing allele in the GT field.
- **PS** : phase set. The values associated with this key are non-negative 32-bit integers indicating the phase set this genotype belongs to. A phase set is a set of phased genotypes. Phased genotypes for an individual sample that are on the same chromosome and have the same PS value are in the same phased set. A phase set specifies multi-marker haplotypes for the phased genotypes in the set. All phased genotypes that do not contain a PS subfield are assumed to belong to the same phased set. If the genotype in the GT field is unphased, the corresponding PS field is ignored.

One additional key (not defined in *VCFMetaData* ['format']) is required:

- `sample-name`, the name of the sample, which is normally defined by the column header in a VCF file.

```
{
  'sample-name': <string>,
  <string>:[<string>, ...],
  <string>:[<string>, ...],
  ...
}
```

VCFLine

Representation of VCF line. The info field maps the IDs from *VCFInfoField*'s defined in *VCFMetaData*. Required keys:

- **chrom**: Chromosome. An identifier from the reference genome or an angle-bracketed ID String (<ID>) pointing to a contig in the assembly file (the file pointed to by `assembly` in *VCFMetaData*). The colon symbol (:) must be absent from all chromosome names to avoid parsing errors when dealing with breakends. (String, no white-space permitted).

- `pos`: Position. The reference position, 1-indexed. Telomeres are indicated by using positions 0 or N+1, where N is the length of the corresponding chromosome or contig. (Integer).
- `id`: Identifier. List of unique identifiers if available. If there is no identifier available, then the list is empty. (List of strings, no white-space or semi-colons permitted)
- `ref`: Reference base(s). Each base must be one of A,C,G,T,N (case insensitive). Multiple bases are permitted. The value matching the `pos` key refers to the position of the first base in this string. For simple insertions and deletions in which either the `ref` or one of the `alt` alleles would otherwise be null/empty, the `ref` and `alt` values must include the base before the event (which must be reflected in the `pos` field), unless the event occurs at position 1 on the contig in which case it must include the base after the event; this padding base is not required (although it is permitted) for variations such as complex substitutions or other events where all alleles have at least one base represented in their strings. If any of the `alt` alleles is a symbolic allele (an angle-bracketed ID String <ID>), then the padding base is required and `pos` denotes the coordinate of the base preceding the polymorphism. (String)
- `alt`: Alternate base(s). List of alternate non-reference alleles called on at least one of the samples. Options are strings made up of the bases A,C,G,T,N,*, (case insensitive) or an angle-bracketed ID String (<ID>) or a breakend replacement string as described in the VCFv4.2 section on breakends. The * allele is reserved to indicate that the allele is missing due to an upstream deletion. If there are no alternative alleles, the list should be empty. (String; no whitespace, commas, or angle-brackets are permitted in the ID String itself).
- `qual`: Quality. Phred-scaled quality score for the assertion made in ALT. $-10 \cdot \log_{10}(\text{prob}[\text{call in alt is wrong}])$. If `alt` is empty, then this is $-10 \cdot \log_{10}(\text{prob}[\text{variant}])$, otherwise, this is $-10 \cdot \log_{10}(\text{prob}[\text{no variant}])$. If this is unknown, return None.
- `filter`: Filter status. A list of length 1 with value *PASS* if this position has passed all filters (if a call is made at this position). Otherwise, if the site has not passed all filters, a list of codes for filters that fail. [*q10*, *s50*] might indicate that at this site the quality is below 10 and the number of samples with data is below 50% of the total number of samples. *0* is reserved and should not be used as a filter string. If filters have not been applied, then the list should be empty. (List of strings with no white-space or semi-colons permitted)
- `info`: Additional information (JSON-formatted). Keys are the values associated with the 'ID' key in the *VCFInfoField*'s of *VCFMetaData*['info']. Values are a list of strings with no white-space, semi-colons, commas or equals-signs. List may be empty for Flag info keys.
- [optional] `format`: Genotype information (List of *VCFSampleFormatData*). Given if 'format' is defined in *VCFMetaData*.

```
{
  'chrom':<string>,
  'pos':<int>,
  'id':[<string>, ...],
  'ref':<string>,
  'alt':[<string>, ...],
  'qual':<float>|None,
  'filter':[<string>, ...],
  'info': {
    <string>:[<string>, ...],
    <string>:[<string>, ...],
    ...
  },
  'format': [VCFSampleFormatData, ...]
}
```

gVCFBlock

Genotype quality ranges used for banding. List of integers (length greater than or equal to 2). Each pair of integers creates a minGQ (inclusive) and maxGQ (exclusive) pair for a gVCFBlock. Must be strictly increasing, the first entry must be 0, and the last entry must be 2147483647 (the largest unsigned integer representable in 32 bits).

```
[0, <int>, ..., 2147483647]
```

gVCFMetaData

Representation of a gVCF Header. JSON-formatted with keys:

- `fileformat`, whose value indicates the format of the VCFLines returned by the Lightning server
- `fileDate`, the current date, format (YYYYMMDD)
- `source`, the Lightning server producing the VCF MetaData
- `reference`, the location of the reference fasta file used to generate the VCFLines
- `assembly`, same as `reference`
- `alt`, for symbolizing imprecise structural variants. Includes 'NOT_REF' for the non-reference alternate allele.
- `info`, for information field formats. Includes 'END', which is used to indicate the end of a VCFLine.
- [optional] `gvcfblock`, for information about splitting gVCF blocks
- [optional] `filter`, for filters that have been applied to the data
- [optional] `format`, for genotype fields
- [optional] `contig`, for pointing to sequence contigs
- [optional] `sample`, for defining sample to genome mappings
- [optional] `pedigree`, for defining relationships between genomes. Can be a list of name:genome or a url pointing to a pedigree database

```
{
  'fileformat':<string>,
  'fileDate':<int>,
  'source':<string>,
  'reference':<string>,
  'assembly':<string>,
  'info': [
    {
      'ID': 'END',
      'Number':1,
      'Type':Integer,
      'Description':'Stop position of the interval'
    },
    VCFInfoField_,
    ...
  ],
  'alt': [
    {
      'ID': 'NOT_REF',
      'Description':'Represents any possible alternative allele at this location'
```

```
    },
    VCFAlternativeAlleleField_,
    ...
  ],
  'gvcfblock': gVCFBlock_,
  'filter': [VCFFilterField_, ...],
  'format': [VCFFormatField_, ...],
  'contig': [VCFContigField_, ...],
  'sample': [VCFSampleField_, ...],
  'pedigree': [{<string>:<string>}, ...] OR <string>
}
```

LIGHTNING API SPECIFICATIONS

We highly recommend reading the appropriate version of [Data Structures Specifications](#) before reading the API specifications you are interested in.

Components for each API specification:

- Server Namespace
- API Calls
- API Examples

Contents:

8.1 Alternate Lightning v0.1.0 API Specifications

We highly recommend reading [Data Structures Specifications, v0.1.0](#) before diving into this API. These API only support well-sequenced tile variants

8.1.1 Lightning Server Namespace

```
/status : returns the API version running on the server
/tile-library
  /tag-sets : returns the tag set version information for all versions
              supported by this Lightning server instance.
  /{tag-set-identifier} : given the tag set version identifier, returns
                        information about that tag set.
  /paths : given the tag set version integer, returns the paths in that
           tag set.
  /{path-int} : given the tag set version identifier and the path integer,
               returns information about that path.
  /tile-positions : given the tag set version identifier, returns the tile
                  positions in that tag set.
  /{tile-position-id} : given the tag set version identifier and tile
                       position identifier, returns information about
                       that tile position.
  /locus : given the tag set version identifier, tile position identifier,
           and optional query parameters containing assembly information,
           returns locus information about the tile position.
  /tile-variants : given a tag set version identifier, returns the tile
                  variants in that tag set in this Lightning server instance.
  /{tile-variant-id} : given the tag set version identifier and tile variant
                       identifier, returns details about the tile variant.
```

```
/locus : given the tag set version identifier, tile variant identifier,
         and optional query parameters containing assembly information,
         returns locus information about the tile variant.
/subsequence: given the tag set version identifier, tile variant
              identifier, and query parameters containing locus
              information, returns the subsequence of the tile variant
/annotations: given the tag set version identifier and tile variant
              identifier, returns the annotation identifiers applying
              to that tile variant.
/annotations : returns a list of annotation identifiers loaded into the Lightning
              instance.
/{annotation-id} : given an annotation id, returns the tile variants associated
                  with that annotation.
/callsets : returns a list of all genome names, termed callsets, loaded into this
           Lightning server instance.
/{callset-name} : given the callset name, returns details about the callset.
/gvcf : given the callset name and locus query parameters, returns a list
       of gVCF lines.
/vcf : given the callset name and locus query parameters, returns a list of
      VCF lines.
/tile-variants : given the callset name and tile position query parameters,
                returns the tile variants the callset has at the given tile
                position.
/assemblies : returns the available assemblies
/{assembly-id} : returns the details about the assembly, including a list of
                loaded loci (valid locations) on the Lightning server instance.
/searches : returns a list of searches that have been performed
/{search-id} : returns the specific search and the answer of the search
```

8.1.2 GET /status

Request used to obtain the status of the Lightning server instance, which currently only includes the API version the server is running. Does not require any query parameters.

Response body:

```
{
  'api-version': <int>.<int>.<int>
}
```

Example response:

```
{ 'api-version' : 0.1.0 }
```

8.1.3 GET /tile-library/tag-sets

Request used to get the available tag sets on this Lightning instance. Does not require any query parameters. Returns a list of tag set unique identifiers (portable data hashes of the collection containing the tag set). This collection contains information about the tag set and the path dividers.

Response body:

```
[<string>, ...]
```

Example response:

```
['d87075c41962489cb9ce7d63da1d7841', '047ae54fba97385716acd2c552fae763']
```

8.1.4 GET /tile-library/tag-sets/{tag-set-identifier}

Request used to get information about the given tag set identifier. Does not require any query parameters.

Response body:

```
{
  'tag-set-identifier' : <string>,
  'tag-set-integer': <int> (base 16)
}
```

Example response:

```
{
  'tag-set-identifier' : 'd87075c41962489cb9ce7d63da1d7841',
  'tag-set-integer' : 00
}
```

8.1.5 GET /tile-library/tag-sets/{tag-set-identifier}/paths

Request used to get the available paths for a specific tag set on this Lightning server instance. Tag set is specified using the tag set identifier in the uri. Does not require any query parameters. Returns a list of path integers (in base 16).

Response body:

```
[<int>, ...]
```

Example response for a BRCA Lightning Instance:

```
[247, 2c5]
```

8.1.6 GET /tile-library/tag-sets/{tag-set-identifier}/paths/{path-int}

Request used to get information about a specific path for a specific tag set on this Lightning server instance. Tag set is specified using the tag set identifier in the uri. The path is identified using the path integer, written in base 16 in the uri. Does not require any query parameters. Returns a dictionary with path information for the specified path and tag set.

Response body:

```
{
  'path' : <int> (base 16),
  'num-positions' : <int> (base 10)
}
```

Example response:

```
{
  'path' : 2c5,
  'num-positions' : 12462
}
```

8.1.7 GET /tile-library/tag-sets/{tag-set-identifier}/tile-positions

Request used to get the available tile positions for a specific tag set on this Lightning server instance. Tag set is specified using the tag set identifier in the uri. Does not require any query parameters. In the future, might support query parameters filtering on information about the tile (like its path). Returns a list of tile position identifiers (*TilePosition*).

Response body:

```
[<TilePosition>, ...]
```

Example response for a BRCA Lightning Instance:

```
['00.247.0000', '00.247.0001', ..., '00.247.1bfb', '00.2c5.0000', ..., '00.2c5.30ae']
```

8.1.8 GET /tile-library/tag-sets/{tag-set-identifier}/tile-positions/{tile-position-id}

Request used to get information about a specific tile position for a specific tag set on this Lightning server instance. The tag set is specified using the tag set identifier in the uri. The tile position is identified using the *TilePosition* in the uri. Does not require any query parameters. Returns a dictionary with tile position information for the specified tag set and tile position.

Response body:

```
{
  'tile-position': <TilePosition>,
  'total-tile-variants': <int>, (base 10)
  'well-sequenced-tile-variants': <int>, (base 10)
  'num-genomes': <int> (base 10)
}
```

Example Response:

```
{
  'tile-position': '00.247.0000',
  'total-tile-variants': 25,
  'well-sequenced-tile-variants': 0,
  'num-genomes': 680
}
```

8.1.9 GET /tile-library/tag-sets/{tag-set-identifier}/tile-positions/{tile-position-id}/locus

Request used to get locus information about a specific tile position for a specific tag set on this Lightning server instance. The tag set is specified using the tag set identifier in the uri. The tile position is identified using the *TilePosition* in the uri. Does not require any query parameters, but an assembly identifier may be used to get information about a specific assembly. Returns a list of *Locus*'s. If no query parameters are specified, the list returned contains the loci for all assemblies in the Lightning server instance.

GET Query Parameters:

Parameter name	Type	Notes
assembly-name	<string>	Optional
assembly-pdh	<string>	Optional

Response body:

```
[<Locus>, ...]
```

Example Query Parameters:

Parameter name	Value
assembly-name	'hg19'
assembly-pdh	'dad94936d4144f5e0a289244d8be93e9+5735'

Example response body:

```
[
  {
    'assembly-name': 'hg19',
    'assembly-pdh': 'dad94936d4144f5e0a289244d8be93e9+5735',
    'chromosome-name': '13',
    'indexing': 0,
    'start-position': 32199976,
    'end-position': 32200225
  }
]
```

8.1.10 GET /tile-library/tag-sets/{tag-set-identifier}/tile-variants

Request used to get the available tile variants for a specific tag set on this Lightning server instance. Tag set is specified using the tag set identifier in the uri. Does not require any query parameters. In the future, might support query parameters filtering on information about the tile variant (like its path). Returns a list of tile variant identifiers (*TileVariant*).

Response body:

```
[<TileVariant>, ...]
```

Example response for a BRCA Lightning Instance:

```
[
  '00.247.0000.830003ac103a97d8f7992e09594ac68e',
  '00.247.0000.455577ff6b0d38188477ee2bfb2f0ea8',
  ...,
  '00.247.1bfb.c95325c08a449529143776e18561db71',
  '00.2c5.0000.1948117b4a56e4ad73d36dce185110fd',
  ...,
  '00.2c5.30ae.bc952f709d7419f7e103daa2b7e469a9'
]
```

8.1.11 GET /tile-library/tag-sets/{tag-set-identifier}/tile-variants/{tile-variant-id}

Request used to get information about a specific tile variant for a specific tag set on this Lightning server instance. The tag set is specified using the tag set identifier in the uri. The tile variant is identified using the *TileVariant* in the uri. Does not require any query parameters. Returns details about the specified *TileVariant* as a *TileVariantDetail*.

Response body:

```
<TileVariantDetail>
```

Example response:

```

{
  'tile-variant': '00.2c5.30ae.bc952f709d7419f7e103daa2b7e469a9',
  'tag-length': 24,
  'start-tag': 'gccaggagttttaaaactactga',
  'end-tag': '',
  'is-start-of-path': False,
  'is-end-of-path': True,
  'sequence' : 'gccaggagttttaaaactactgatgccacctcccacacccaaaagtctgattaattgatctagggtatggcctgag
  'md5sum': 'bc952f709d7419f7e103daa2b7e469a9',
  'length': 394,
  'number-of-positions-spanned': 1,
  'population-frequency': 0.5,
  'population-count': 150,
  'population-total': 300
}

```

8.1.12 GET `/tile-library/tag-sets/{tag-set-identifier}/tile-variants/{tile-variant-id}/locus`

Request used to get locus information about a specific tile variant for a specific tag set on this Lightning server instance. The tag set is specified using the tag set identifier in the uri. The tile variant is identified using the *TileVariant* in the uri. Does not require any query parameters, but an assembly identifier may be used to get information about a specific assembly. Returns a list of *Locus*'s. If no query parameters are specified, the list returned contains the loci for all assemblies in the Lightning server instance.

GET Query Parameters:

Parameter name	Type	Notes
assembly-name	<string>	Optional
assembly-pdh	<string>	Optional

Response body:

```
[<Locus>, ...]
```

Example Query Parameters:

Parameter name	Value
assembly-name	'hg19'
assembly-pdh	'dad94936d4144f5e0a289244d8be93e9+5735'

Example response body:

```

[
  {
    'assembly-name': 'hg19',
    'assembly-pdh': 'dad94936d4144f5e0a289244d8be93e9+5735',
    'chromosome-name': '13',
    'indexing': 0,
    'start-position': 32199976,
    'end-position': 32200225
  }
]

```

8.1.13 GET `/tile-library/tag-sets/{tag-set-identifier}/tile-variants/{tile-variant-id}/subsequence`

Request used to get a subsequence of a specific tile variant for a specific tag set on this Lightning server instance. The tag set is specified using the tag set identifier in the uri. The tile variant is identified using the *TileVariant* in the uri. Query parameters follow the *Locus* data structure. These query parameters are required. The response is a dictionary with one key: `sequence`.

If the *Locus* provided touches a locus outside of the tile variant, the API should return an error.

GET Query Parameters:

Parameter name	Type
<code>assembly-name</code>	<string>
<code>assembly-pdh</code>	<string>
<code>chromosome-name</code>	<string>
<code>indexing</code>	<int>
<code>start-position</code>	<int>
<code>end-position</code>	<int>

Response body:

```
{
  'sequence': <string>
}
```

Example Query Parameters:

Parameter name	Value
<code>assembly-name</code>	'hg19'
<code>assembly-pdh</code>	'dad94936d4144f5e0a289244d8be93e9+5735'
<code>chromosome-name</code>	'13'
<code>indexing</code>	0
<code>start-position</code>	32199976
<code>end-position</code>	32199983

Example response body:

```
{
  'sequence': 'gggtac'
}
```

8.1.14 GET `/tile-library/tag-sets/{tag-set-identifier}/tile-variants/{tile-variant-id}/annotations`

Request used to get the annotations associated with a specific tile variant for a specific tag set on this Lightning server instance. The tag set is specified using the tag set identifier in the uri. The tile variant is identified using the *TileVariant* in the uri. Does not require any query parameters. The response is a list of annotation identifiers. If no annotation identifiers exist that are associated with that tile variant, returns an empty list.

Response body:

```
[ <string>, <string>, ... ]
```

Example response body:

```
[ 'annotation1', 'annotation3' ]
```

8.1.15 GET /annotiles

Request used to get all the annotations loaded into this Lightning server instance. Does not require any query parameters.

Response body:

```
[ 'annotation0', 'annotation1', 'annotation2', 'annotation3' ]
```

8.1.16 GET /annotiles/{annotation-id}

Request used to determine which tile variants (*TileVariant*) are associated with a particular annotation identifier. Does not require any query parameters.

Response body:

```
[<TileVariant>, ...]
```

Example response body:

```
[ '00.2c5.30ae.bc952f709d7419f7e103daa2b7e469a9' ]
```

8.1.17 GET /callsets

Request used to determine which callsets (called genomes) are loaded into this Lightning server instance. Returns a list of the names of the callsets. These names are expected to be identical to the names used in the phenotype database the user chooses. Does not require any query parameters.

Response body:

```
[<string>, <string>, ... ]
```

Example response body:

```
[ 'human1-complete-genomics', 'human1-illumina', 'human2', 'assembly-hg19' ]
```

8.1.18 POST /callsets

Possibly a v0.1.1 functionality.

Creates a callset (called genomes) in this Lightning server instance. Returns an Arvados UUID pointing to the pipeline instance adding the callset. The callset will appear on the server once the pipeline finishes successfully.

Request body:

```
{
  'callset-name': <string>,
  'callset-format': <string>,
  'callset-collection-pdh': <string>
}
```

Response body:

```
<ArvadosUUID>
```

8.1.19 GET /callsets/{callset-name}

Request used to find details about a specific callset. Returns dictionary containing details about the callset. Phenotypic details may be passed by querying the phenotype database specified by the user, but this phenotype database is not part of Lightning.

Response body:

```
{
  'callset-name' : <string>,
  'callset-locator': <string>
}
```

Example response body:

```
GET /callsets/assembly-hg19
Response body: {
  'callset-name': 'assembly-hg19',
  'callset-locator': '1cf491c1ea99543da01c5a8f6b8a6dba+228008/hg19'
}
```

8.1.20 GET /callsets/{callset-name}/gvcf-header

Request used to obtain a valid gVCF header. Currently should not change based on the callset name specified in the uri. Requires query parameters specifying an assembly and an optional parameter specifying the gVCF blocks (*gVCFBlock*) to use. Responds with dictionary of type *gVCFMetaData*.

GET Query Parameters:

Parameter name	Type	Notes
assembly-name	<string>	
assembly-pdh	<string>	
gvcf-block	<list of ints>	Optional

Response body:

```
<gVCFMetaData>
```

Example Query Parameters:

Parameter name	Value
assembly-name	'hg19'
assembly-pdh	'dad94936d4144f5e0a289244d8be93e9+5735'

Example response body:

```
{
  'fileformat': 'VCFv4.2',
  'fileDate': '20150928',
  'source': 'Lightningv0.1.0',
  'assembly': 'dad94936d4144f5e0a289244d8be93e9+5735/hg19',
  'info': [
    {
      'ID': 'END',
      'Number': 1,

```

```

        'Type':Integer,
        'Description':'Stop position of the interval'
    }
],
'format': [
    {
        'ID': 'GT',
        'Number':1,
        'Type':'String',
        'Description':'Genotype'
    }
],
'alt': [
    {
        'ID': 'NOT_REF',
        'Description':'Represents any possible alternative allele at this location'
    }
],
'gvcfblock': [0, 2147483647]
}

```

8.1.21 GET /callsets/{callset-name}/gvcf

Request used to obtain a list of valid gVCF lines for specified callset in the uri. Requires query parameters specifying a *Locus* to retrieve lines for. The query parameters may include an optional parameter to specify the gVCF blocks (*gVCFBlock*) to use. Responds with a list of dictionaries of type *gVCFLine*. These represent the gVCF lines of the specified callset.

GET Query Parameters: ::

Parameter name	Type	Notes
assembly-name	<string>	
assembly-pdh	<string>	
chromosome-name	<string>	
indexing	<int>	
start-position	<int>	
end-position	<int>	
gvcf-block	<list of ints>	Optional

Response body:

```
[<gVCFLine>, ...]
```

Example query parameters:

Parameter name	Value
assembly-name	'hg19'
assembly-pdh	'dad94936d4144f5e0a289244d8be93e9+5735'
chromosome-name	'13'
indexing	0
start-position	32199976
end-position	32200225

Example response body:

```
[
  {
    'chrom':'13',
    'pos':32199977,

```

```

    'ref':G,
    'alt':['<NON_REF>'],
    'filter':[],
    'format': [
      {'sample-name':'human1-illumina', 'GT':'0/0'}
    ],
    'info':{'END':[32200122]}
  },
  {
    'chrom':'13',
    'pos':32200123,
    'ref':T,
    'alt':['A','<NON_REF>'],
    'filter':[],
    'format': [
      {'sample-name':'human1-illumina', 'GT':'0/1'}
    ]
  },
  {
    'chrom':'13',
    'pos':32200124,
    'ref':G,
    'alt':['<NON_REF>'],
    'filter':[],
    'format': [
      {'sample-name':'human1-illumina', 'GT':'0/0'}
    ],
    'info':{'END':[32200225]}
  }
]

```

8.1.22 GET /callsets/{callset-name}/vcf-header

Request used to obtain a valid VCF header. Currently should not change based on the callset name specified in the uri. Requires query parameters specifying an assembly. Responds with dictionary of type *VCFMetaData*.

GET Query Parameters:

Parameter name	Type
assembly-name	<string>
assembly-pdh	<string>

Response body:

```
<VCFMetaData>
```

Example Query Parameters:

Parameter name	Value
assembly-name	'hg19'
assembly-pdh	'dad94936d4144f5e0a289244d8be93e9+5735'

Example response body:

```

{
  'fileformat':'VCFc4.2',
  'fileDate':20150928,

```

```

'source': 'Lightningv0.1.0',
'assembly': 'dad94936d4144f5e0a289244d8be93e9+5735/hg19',
'format': [
  {
    'ID': 'GT',
    'Number': 1,
    'Type': 'String',
    'Description': 'Genotype'
  }
]
}

```

8.1.23 GET /callsets/{callset-name}/vcf

Request used to obtain a list of valid VCF lines for specified callset in the uri. Requires query parameters specifying a *Locus* to retrieve lines for. Responds with a list of dictionaries of type *gVCFLine*. These represent the VCF lines of the specified callset.

Parameter name	Type
assembly-name	<string>
assembly-pdh	<string>
chromosome-name	<string>
indexing	<int>
start-position	<int>
end-position	<int>

GET Query Parameters: ::

Response body:

```
[<VCFLine>, ...]
```

Example query parameters:

Parameter name	Value
assembly-name	'hg19'
assembly-pdh	'dad94936d4144f5e0a289244d8be93e9+5735'
chromosome-name	'13'
indexing	0
start-position	32199976
end-position	32200225

Example response body:

```

[
  {
    'chrom': '13',
    'pos': 32200123,
    'ref': 'T',
    'alt': ['A'],
    'filter': [],
    'format': [
      { 'sample-name': 'human1-illumina', 'GT': '0/1' }
    ]
  }
]

```

8.1.24 GET /callsets/{callset-name}/tile-variants

Request used to obtain a list of tile variants for each phase for the callset specified in the uri. A list of *TilePosition* or *TilePositionRange* may be specified using optional query parameters. If no query parameters are specified, all tile variants are returned. Responds with a dictionary with two keys: `callset-name` and `tile-variants`. The value associated with tile variants is a list of lists. Each list represents a phase and contains objects of type *TileVariant*. Tile variants are returned if they intersect at with the given positions, even if they span outside the tile positions given.

GET Query Parameters: ::	Parameter name	Type	Notes
	<code>tile-positions</code>	<code>[<TilePosition><TilePositionRange>, ...]</code>	Optional

Response body:

```
{
  'callset-name': <string>,
  'tile-variants': [
    [ <TileVariant>, ... ],
    [ <TileVariant>, ... ],
    ...
  ]
}
```

Example query parameters:

Parameter name	Type
<code>tile-positions</code>	<code>['00.247.0000-0003']</code>

Example response body:

```
{
  'callset-name': 'human1-illumina',
  'tile-variants': [
    [
      '00.247.0000.830003ac103a97d8f7992e09594ac68e',
      '00.247.0001.a31fd29383d072a5ccf7027ec37df093',
      '00.247.0002.a42a3e835440e21dda2cfd65162e85f0'
    ],
    [
      '00.247.0000.455577ff6b0d38188477ee2bfb2f0ea8',
      '00.247.0001.30c792a4fc1f0bd88dcc10907e6f27e6'
    ],
  ]
}
```

8.1.25 GET /assemblies

Request used to obtain a list of assemblies (*Assembly*) available in this Lightning server instance. Does not require any query parameters.

Response body:

```
[ <Assembly>, ... ]
```

Example response body:

```
[
  {
    'assembly-name': 'hg19',
```

```

    'assembly-pdh': 'dad94936d4144f5e0a289244d8be93e9+5735'
  },
  {
    'assembly-name': 'GRCh38',
    'assembly-pdh': '047ae54fba97385716acd2c552fae763+5735'
  }
]

```

8.1.26 GET /assemblies/{assembly-id}

Request used to obtain details about a specified assembly, given by the portable data hash in the uri, available in this Lightning server instance. Does not require any query parameters. Details are returned using a list of *Locus*'s.

Response body:

```
[ <Locus>, ... ]
```

Example response body:

```

[
  {
    'assembly-name': 'hg19',
    'assembly-pdh': 'dad94936d4144f5e0a289244d8be93e9+5735',
    'chromosome-name': '13',
    'indexing': 0,
    'start-position': 32199976,
    'end-position': 34000000
  },
  {
    'assembly-name': 'hg19',
    'assembly-pdh': 'dad94936d4144f5e0a289244d8be93e9+5735',
    'chromosome-name': '17',
    'indexing': 0,
    'start-position': 40899976,
    'end-position': 44900000
  },
]

```

8.1.27 GET /searches/help

Returns the filters this Lightning server instance supports and the options available for each filter. If the filter 'tile-variants' is available, options will be an empty list for that filter, since those may be obtained by GET /tile-library/tag-sets/{tag-set-identifier}/tile-variants. However, the key 'tile-variants' is included for completeness. The filter 'callsets' is treated identically: options will be an empty list for that filter, the key is included for completeness, and the options for 'callsets' may be obtained by GET /callsets. Does not require any query parameters. Details are returned in a dictionary with the keys matching the available filters. The values matching these keys are the available options for those filters.

Response body:

```

{
  <string> : [ <string>, <string>, ... ],
  <string> : [ <string>, <string>, ... ],

```

```
    ...
  }
```

Example response body:

```
{
  'tile-variants' : [],
  'callsets': [],
  'phasing': ['any-phase', 'all-phases']
}
```

8.1.28 GET /searches

Returns the searches that have been performed on this Lightning server instance. Does not require any query parameters. Details are returned as a list of pipeline uuids (*ArvadosUUID*).

Response body:

```
[<ArvadosUUID>, ... ]
```

Example response body:

```
['su921-d1hrv-lcrzu0qrs1iiu03', 'su921-d1hrv-vg62va2y75c5wxm']
```

8.1.29 GET /searches/{search-id}

Returns the details about the search, given by the Arvados UUID (*ArvadosUUID*) in the uri, that has been performed on this Lightning server instance. Does not require any query parameters. Details are returned as a dictionary. The value associated with 'response' will be None if the pipeline is still running or failed.

Response body:

```
{
  'search-id': <ArvadosUUID>,
  'parameters': {
    <string>:[<string>],
    ...
  },
  'response': None | [<string>, ...]
}
```

Example response body:

```
{
  'search-id': 'su921-d1hrv-lcrzu0qrs1iiu03',
  'filters': {
    'tile-variants': [['00.247.0000.830003ac103a97d8f7992e09594ac68e']],
    'phasing': ['any-phase']
  },
  'response': ['human1-complete-genomics', 'human1-illumina', 'assembly-hg19']
}
```

8.1.30 POST /searches

Creates a search with the specified parameters. The parameters should be available in the response GET /searches/help. If the parameter 'tile-variants' is used, the value associated with that key should be

of type (*TileVariantLogic*). Returns the pipeline uuid (format *ArvadosUUID*) of that search. The output of that pipeline will be a list of callsets matching the search.

Request body:

```
{
  <string> : [ <string>, <string>, ...],
  ...
}
```

Response body:

```
<ArvadosUUID>
```

Example request body:

```
{
  'tile-variants': [['00.247.0000.830003ac103a97d8f7992e09594ac68e']],
  'phasing': ['any-phase']
}
```

Example response body:

```
'su921-dlhrv-1crzu0qrs1iiu03'
```

8.2 Batch Processing

The API calls are excellent for one-off queries. However, many users may wish to use these calls many times for many samples and/or over large parts of the genome. To speed up these queries and reduce load on the server, Lightning supports batch processing.

Todo

Design and document batch processing

/tile-variants	/filter [returns information about the requested tile variants, accepts request] containing a list of tile variants
	/loci [returns locus information about the requested tile variants, accepts] request containing a list of tile variants
/assemblies	[returns a list of loaded loci (valid locations) for all assemblies]
/tiles	[returns tile variants subsequences and list of positions touching the] loci, requires request containing loci
/annotiles	/filter-annotations [given a request of tile variants and filters, returns] the annotations applying to those tile variants according to the filters.
	/filter-tile-variants [given a request of annotation identifiers and filters,] returns the tile variants with those annotations that satisfy the filters.

8.2.1 POST /tile-variants/filter

The only reason this is a POST is because the request arguments could be too big to fit in the URL.

Requires a request body with a dictionary of filters to apply and a list of tile variants (*TileVariant*), tile positions (*TilePosition*), and/or ranges of tile positions (*TilePositionRange*). If a tile position is given, all variants at that position

are evaluated against the filters. If a tile position range is given, all variants for all positions in the range are evaluated against the filters.

Returns a dictionary of (*TileVariant*) as keys with (*TileVariantDetail*) as the values for the tile variants in the tile library that satisfy all the filters.

Request body:

```
{
  'tile-variants' : [<TileVariant> or <TilePosition> or <TilePositionRange>, ... ],
  'filter': {
    'sequence': <REGEX-string>,
    'start-tag': <REGEX-string>,
    'end-tag': <REGEX-string>,
    'length' : <CMP-tuple>,
    'number-of-positions-spanned': <CMP-tuple>,
    'population-frequency': <CMP-tuple>,
    'population-total' : <CMP-tuple>,
    'is-start-of-path': <boolean>,
    'is-end-of-path': <boolean>
  }
}
```

Response body:

```
{
  <TileVariant> : <TileVariantDetail>,
  <TileVariant> : <TileVariantDetail>,
  ...
}
```

Example request body:

```
{ 'tile-variants': [ '0.0.1.d87075c41962489cb9ce7d63da1d7841' ] }
```

Example response body:

```
{
  '0.0.1.d87075c41962489cb9ce7d63da1d7841':
    {
      'tag-length': 24,
      'start-tag': 'acggcagtagtttggccgctcggt',
      'end-tag': 'tcagaatgttggagggcggtacg',
      'is-start-of-path': false,
      'is-end-of-path': false,
      'sequence' : 'acggcagtagtttggccgctcgntaaatcagaatgttggagggcggnnnn',
      'md5sum': 'd87075c41962489cb9ce7d63da1d7841'
      'length': 51,
      'number-of-positions-spanned': 1,
      'population-frequency': 15,
      'population-total': 15,
    }
}
```

Todo

Fix example response

8.2.2 POST /tile-variants/loci

The only reason this is a POST is because the request arguments could be too big to fit in the URL.

Requires a request body with a list of *TileVariant*'s, *TilePosition*'s, and/or *TilePositionRange*'s, and a list of assemblies (as strings) to return loci for. If a *TilePosition* is given, all variants at that position are returned. If a *TilePositionRange* is given, all variants for all positions in the range are returned.

Returns a dictionary of *TileVariant* with a list of *Locus*'s.

Request body:

```
{
  'tile-variants' : [<TileVariant> or <TilePosition> or <TilePositionRange>, ... ],
  'assembly': [<assembly-name>, <assembly-name>, ... ]
}
```

Response body:

```
{
  <TileVariant> : [ <Locus>, <Locus>, ... ],
  <TileVariant> : [ <Locus>, <Locus>, ... ],
  ...
}
```

Example request body:

```
{
  'tile-variants': [ '0.0.1.d87075c41962489cb9ce7d63da1d7841' ]},
  'assembly': ['hg19']
}
```

Example response body:

```
{
  '0.0.1.d87075c41962489cb9ce7d63da1d7841' : [ ['hg19', 'chr1', 0, 3000, 3248] ]
}
```

8.2.3 POST /assembly-loci/tiles

Returns a list of tile-positions touching the given loci, a dictionary of tile-variants touching those loci with the sequence of the tile variant, cut off if appropriate, and the tag-length.

The only reason this is a POST is because the request arguments could be too big to fit in the URL.

Examples for the INDEL behavior can be found at lightning-dev4.curoverse.com/pad/p/lightning-indel-behavior

Request body:

```
[ <Locus>, <Locus>, ... ] OR { <string>: <Locus>, <string>: <Locus>, ... }
```

Response body:

```
{
  'tile-positions':
    {
      (<string>, <Locus>) OR <Locus> : [ <TilePosition> OR <TilePositionRange>, <TilePosition>
      (<string>, <Locus>) OR <Locus> : [ <TilePosition> OR <TilePositionRange>, <TilePosition>
      ...
    },
}
```

```

'tile-variants':
  {
    (<string>, <Locus>) OR <Locus> : { <TileVariant>: <string>, <TileVariant>:<string>, ...
    (<string>, <Locus>) OR <Locus> : { <TileVariant>: <string>, <TileVariant>:<string>, ...
    ...
  },
'tag-length' : <int>
}

```

Example request body:

```

{
  'exon1' : [ hg38', 'chr19', 1, 44905870, 44905924 ],
  'exon2' : [ 'hg38', 'chr19', 1, 44906603, 44906668 ]
}

```

Example response body:

```

{
  'tile-positions' : {
    ( 'exon1', [ 'hg38', 'chr19', 1, 44905870, 44905924 ] ) : [ '0.756.742-744' ],
    ( 'exon2', [ 'hg38', 'chr19', 1, 44906603, 44906668 ] ) : [ '0.756.745-748' ],
  },
  'tile-variants': {
    ( 'exon1', [ 'hg38', 'chr19', 1, 44905870, 44905924 ] ): {
      '0.756.741.5' : "ctaccgtttaggcggatatacgcgtacggcagctactcaacgcaccggaacttgtgtttgtgtgtgtggtcgc
      '0.756.742.0' : "acggcagtagtnttgccgctcggtaaacgtcagaatgtttgaggggcggtagcgctagagatatcacctct
      '0.756.742.1' : "acggcagtagtnttgccgctcggtaaacgtcagaatgtttgaggggcggtagcgctagagatatcacctct
      ...
      '0.756.743.0' : "gtcgcccactacgcacgttatatgagagagctggcagatgccttatggaagtgactgctaccgtttgttga
    },
    ( 'exon2', [ 'hg38', 'chr19', 1, 44906603, 44906668 ] ): { ... }
  },
  'tag-length' : 24
}

```

8.3 Reasoning Behind the API

API design! Here are some explanations!

Todo

Document API explanations

8.4 Lightning Errors

8.4.1 UnknownAssembly

Used when user requests an assembly that is not loaded into the Lightning instance

8.5 Lightning v0.1.1 API Specifications

We highly recommend reading [Data Structures Specifications, v0.1.1](#) before diving into this API.

8.5.1 Lightning Server Namespace

```

/status : returns the API version running on the server, current load level, etc.
/tile-library
  /tag-version : returns the tag set version information for all versions supported
                by the server.
  /{version-int} : given the tag set version integer, returns a dictionary keyed
                  by path integer with a value of the number of steps
  /{path-int} : given the tag set version integer and the path integers, returns
               a dictionary keyed by step integer with a value of the number
               of variants
  /{step-int} : given the tag set version integer, path integer, and step
               integers, returns a list of tuples:
               (md5sum of sequence, frequency, total population size)
  /{variant-md5sum} : given the tag set version integer, path integer,
                     step integer, and the md5sum of the tile variant,
                     returns details about the tile variant
/tile-variants
  /filter : returns information about the requested tile variants, accepts request
           containing a list of tile variants
  /loci : returns locus information about the requested tile variants, accepts
         request containing a list of tile variants
/callsets : returns a list of all genome names (md5sums of the input callset file)
  /{callset-name} : returns details about the callset
  /gvcf : returns a list of gvcf lines, requires request containing locus or loci
         to return. This request is expected to be extended to enable annotation placement
  /vcf : returns a list of vcf lines, requires request containing locus or loci
        to return. This request is expected to be extended to enable annotation placement
  /filters : returns valid filters (which may optionally include phenotypical
           filters)
  /filter-phase-insensitive : returns a list of callsets matching filters, accepts
                             request containing filters and subsection of the callsets to run query
  /filter-phase-sensitive : returns a dictionary of callsets mapped to the integers
                            of their phases matching filters, accepts request containing filters and
                            of the callsets to run query on
  /tile-variants : returns tile variants for subsection of callsets over positions
                 given, requires request containing callsets and positions
/assembly-loci : returns a list of loaded loci (valid locations) for all assemblies
/tiles : returns tile variants subsequences and list of positions touching the
        loci, requires request containing loci

```

8.5.2 GET /status

Returns the status of the server, including the api-version running on the server, possibly the current load level, etc.

Request body:

```
None
```

Response body:

```
{
  'api-version': <int>[.<int>]*,
  ...
}
```

8.5.3 GET /tile-library/tag-version

Probably will be unchanged from v0.1.0

8.5.4 GET /tile-library/{version-int}

Probably will be unchanged from v0.1.0

8.5.5 GET /tile-library/{version-int}/{path-int}

Probably will be unchanged from v0.1.0

8.5.6 GET /tile-library/{version-int}/{path-int}/{step-int}

Probably will be unchanged from v0.1.0

8.5.7 GET /tile-library/{version-int}/{path-int}/{step-int}/{variant-md5sum}

Specifications will probably will be unchanged from v0.1.0. But the response will change because the data structure changes.

8.5.8 POST /tile-variants/filter

The only reason this is a POST is because the request arguments could be too big to fit in the URL.

Requires a request body with a dictionary of filters to apply and a list of tile variants (*TileVariant*), tile positions (*TilePosition*), and/or ranges of tile positions (*TilePositionRange*). If a tile position is given, all variants at that position are evaluated against the filters. If a tile position range is given, all variants for all positions in the range are evaluated against the filters.

Returns a dictionary of (*TileVariant*) as keys with (*TileVariantDetail*) as the values for the tile variants in the tile library that satisfy all the filters.

Request body:

```
{
  'tile-variants' : [<TileVariant> or <TilePosition> or <TilePositionRange>, ... ],
  'filter': {
    'sequence': <REGEX-string>,
    'start-tag': <REGEX-string>,
    'end-tag': <REGEX-string>,
    'length' : <CMP-tuple>,
    'number-of-positions-spanned': <CMP-tuple>,
    'population-frequency': <CMP-tuple>,
    'population-total' : <CMP-tuple>,
    'is-start-of-path': <boolean>,
  }
}
```

```
    'is-end-of-path': <boolean>
  }
}
```

Response body:

```
{
  <TileVariant> : <TileVariantDetail>,
  <TileVariant> : <TileVariantDetail>,
  ...
}
```

Example request body:

```
{ 'tile-variants': [ '0.0.1.d87075c41962489cb9ce7d63da1d7841' ] }
```

Example response body:

```
{
  '0.0.1.d87075c41962489cb9ce7d63da1d7841':
    {
      'tag-length': 24,
      'start-tag': 'acggcagtagttttgccgctcggg',
      'end-tag': 'tcagaatgtttgaggggcggtacg',
      'is-start-of-path': false,
      'is-end-of-path': false,
      'sequence': 'acggcagtagttttgccgctcgntaaatcagaatgtttgaggggcgnnnnn',
      'md5sum': 'd87075c41962489cb9ce7d63da1d7841',
      'length': 51,
      'number-of-positions-spanned': 1,
      'population-frequency': 15,
      'population-total': 15,
    }
}
```

Todo

Fix example response

8.5.9 POST /tile-variants/loci

The only reason this is a POST is because the request arguments could be too big to fit in the URL.

Requires a request body with a list of *TileVariant*'s, *TilePosition*'s, and/or *TilePositionRange*'s, and a list of assemblies (as strings) to return loci for. If a *TilePosition* is given, all variants at that position are returned. If a *TilePositionRange* is given, all variants for all positions in the range are returned.

Returns a dictionary of *TileVariant* with a list of *Locus*'s.

Request body:

```
{
  'tile-variants' : [<TileVariant> or <TilePosition> or <TilePositionRange>, ... ],
  'assembly': [<assembly-name>, <assembly-name>, ... ]
}
```

Response body:

```
{
  <TileVariant> : [ <Locus>, <Locus>, ... ],
  <TileVariant> : [ <Locus>, <Locus>, ... ],
  ...
}
```

Example request body:

```
{
  'tile-variants': [ '0.0.1.d87075c41962489cb9ce7d63da1d7841' ],
  'assembly': [ 'hg19' ]
}
```

Example response body:

```
{
  '0.0.1.d87075c41962489cb9ce7d63da1d7841' : [ [ 'hg19', 'chr1', 0, 3000, 3248 ] ]
}
```

8.5.10 GET /callsets

Returns a list of the names of the callsets held in this lightning server.

Request body:

```
None
```

Response body:

```
[<string>, <string>, ... ]
```

Example response body:

```
[ 'human1-complete-genomics', 'human1-illumina', 'human2', 'assembly-hg19' ]
```

8.5.11 GET /callsets/{callset-name}

Returns dictionary containing details about the callset. If the callset name is the name of an assembly or reference, the key 'is-reference' is in the response body. If the callset is a non-reference sequence, the keys will not include 'is-reference', but other keys indicating phenotype may be present.

Request body:

```
None
```

Reference/assembly response body:

```
{
  'name' : <string>,
  'is-reference': <string>
}
```

Non-reference callset response body:

```
{
  'name': <string>,
  <string> : <string>,
  <string> : <string>,
}
```

```
} ...
```

Example 1:

```
GET /callsets/assembly-hg19
Response body: {
  'name' : 'assembly-hg19',
  'is-reference': 'hg19'
}
```

Example 2:

```
GET /callsets/human1-illumina
Response body: {
  'name' : 'human1-illumina',
  'phases-known': 'false',
  'sex': 'female',
  'ethnicity': 'Caucasian'
}
```

8.5.12 GET /callsets/{callset-name}/gvcf (acts like a POST)

Todo

Is this a POST or a GET? Stopped intensive editing here

Requires a request containing information about the loci to return. Returns a list of gVCF lines for the callset at these loci.

Request body:

```
{
  'loci' : [ <Locus>, <Locus>, ... ] OR { <string> : <Locus>, <string> : <Locus>, ... }
}
```

Response body:

```
{
  (<string>, <Locus>) OR <Locus> : [
    <gVCF-line>,
    <gVCF-line>,
    ...
  ],
  (<string>, <Locus>) OR <Locus> : [
    <gVCF-line>,
    <gVCF-line>,
    ...
  ],
  ...
}
```

8.5.13 GET /callsets/{callset-name}/vcf (acts like a POST)

Todo

Is this a POST or a GET? Stopped intensive editing here

Requires a request containing information about the loci to return. Returns a list of VCF lines for the callset at these loci.

Request body:

```
{
  'loci' : [ <Locus>, <Locus>, ... ] OR { <string> : <Locus>, <string> : <Locus>, ... },
}
```

Response body:

```
{
  (<string>, <Locus>) OR <Locus> : [
    <VCF-line>,
    <VCF-line>,
    ...
  ],
  (<string>, <Locus>) OR <Locus> : [
    <VCF-line>,
    <VCF-line>,
    ...
  ],
  ...
}
```

8.5.14 GET /callsets/filters

Returns the filters the lightning server supports and the options available for each filter (options not shown for filter 'tile-variants', but key 'tile-variants' is included for completeness). None of the filters can be named 'tile-variants', 'is-reference', or 'callsets'.

Request body:

None

Response body:

```
{
  'tile-variants' : [ ],
  'is-reference': [ <string>, <string>, ... ],
  <string> : [ <string>, <string>, ... ],
  <string> : [ <string>, <string>, ... ],
  ...
}
```

Example response body:

```
{
  'tile-variants' : [ ],
  'is-reference': [ 'hg19', 'hg38' ],
  'sex': [ 'male', 'female', 'other' ]
}
```

8.5.15 POST /callsets/filter-phase-insensitive

Returns a list of callsets that match the request filters on at least one phase. Besides the tile-variants logic, all filters are AND'd. If the value associated with 'callsets' is an empty list, the query is done on all callsets in the lightning

server.

Request body:

```
{
  'tile-variants-sat' : <TileVariantLogic>,
  'callsets' : [ <string>, <string>, ... ],
  'is-reference': [ <string>, <string>, ... ],
  <string> : [ <string>, <string>, ...],
  <string> : [ <string>, <string>, ...],
  ...
}
```

Response body:

```
[ <string>, <string>, ... ]
```

Todo

example

8.5.16 POST /callsets/filter-phase-sensitive

Returns a dictionary of callsets with a list of phases that match the request filters. Besides the tile-variants logic, all filters are AND'd. If the value associated with 'callsets' is an empty list, the query is done on all callsets in the lightning server.

Request body:

```
{
  'tile-variants-sat' : <TileVariantLogic>,
  'callsets' : [ <string>, <string>, ... ],
  'is-reference': [ <string>, <string>, ... ],
  <string> : [ <string>, <string>, ...],
  <string> : [ <string>, <string>, ...],
  ...
}
```

Response body:

```
{
  <string> : {
    'phase-groups-known' : <boolean>,
    'phases': [ <int>, <int>, ... ]
  },
  <string> : {
    'phase-groups-known' : <boolean>,
    'phases' : [ <int>, <int>, ... ]
  },
  ...
}
```

Todo

example

8.5.17 POST /callsets/tile-variants

Returns a dictionary of callsets with a list of variants for each phase. The variants are returned only for the tile positions given by the tile-positions list. If the value of 'tile-positions' is an empty list, all tile-variants are returned. If the value associated with 'callsets' is an empty list, the query is done on all callsets in the lightning server. The only reason this is a POST is because the request arguments could be too big to fit in the URL.

Request body:

```
{
  'tile-positions': [ <TilePosition> or <TilePositionRange>, ... ],
  'callsets': [ <string>, <string>, ... ]
}
```

Response body:

```
{
  <string>: {
    'phase-groups' : None OR [ <int>-<int>, <int>-<int>, ... ],
    'sequence' : [
      [ <TileVariant>, <TileVariant>, ... ],
      [ <TileVariant>, <TileVariant>, ... ],
      ...
    ]
  },
  <string>: {
    'phase-groups' : None OR [ <int>-<int>, <int>-<int>, ... ],
    'sequence' : [
      [ <TileVariant>, <TileVariant>, ... ],
      [ <TileVariant>, <TileVariant>, ... ],
      ...
    ],
  },
  ...
}
```

Example request body:

```
{
  'tile-positions': [ '0.0.1', '0.0.2' ],
  'callsets': ['assembly-hg19', 'human2' ]
}
```

Example response body:

```
{
  'assembly-hg19': {
    'phase-groups' : [ 0-2 ],
    'sequence' : [ [ '0.0.1.bab3...', '0.0.2.1f35...' ] ]
  },
  'human2': {
    'phase-groups' : None,
    'sequence': [
      [ '0.0.0.1afe...', '0.0.2.1f35...' ],
      [ '0.0.1.bab3...', '0.0.2.1f35...' ]
    ]
  }
}
```

```

    ]
  }
}

```

8.5.18 GET /assembly-loci

Returns a list of loci in the server

Request body:

```
None
```

Response body:

```
[ <Locus>, <Locus>, ... ]
```

Example response body:

```

[
  ['hg19', 'chr1', 0, 0, 249250620],
  ['hg19', 'chr2', 0, 0, 243199372],
  ...,
  ['hg38', 'chr1', 0, 0, 248956421],
  ...
]

```

8.5.19 POST /assembly-loci/tiles

Returns a list of tile-positions touching the given loci, a dictionary of tile-variants touching those loci with the sequence of the tile variant, cut off if appropriate, and the tag-length.

The only reason this is a POST is because the request arguments could be too big to fit in the URL.

Examples for the INDEL behavior can be found at lightning-dev4.curoverse.com/pad/p/lightning-indel-behavior

Request body:

```
[ <Locus>, <Locus>, ... ] OR { <string>: <Locus>, <string>: <Locus>, ... }
```

Response body:

```

{
  'tile-positions':
    {
      (<string>, <Locus>) OR <Locus> : [ <TilePosition> OR <TilePositionRange>, <TilePosition>
      (<string>, <Locus>) OR <Locus> : [ <TilePosition> OR <TilePositionRange>, <TilePosition>
      ...
    },
  'tile-variants':
    {
      (<string>, <Locus>) OR <Locus> : { <TileVariant>: <string>, <TileVariant>:<string>, ...
      (<string>, <Locus>) OR <Locus> : { <TileVariant>: <string>, <TileVariant>:<string>, ...
      ...
    },
  'tag-length' : <int>
}

```

Example request body:

```
{
  'exon1' : [ 'hg38', 'chr19', 1, 44905870, 44905924 ],
  'exon2' : [ 'hg38', 'chr19', 1, 44906603, 44906668 ]
}
```

Example response body:

```
{
  'tile-positions' : {
    ( 'exon1', [ 'hg38', 'chr19', 1, 44905870, 44905924 ] ) : [ '0.756.742-744' ],
    ( 'exon2', [ 'hg38', 'chr19', 1, 44906603, 44906668 ] ) : [ '0.756.745-748' ],
  },
  'tile-variants': {
    ( 'exon1', [ 'hg38', 'chr19', 1, 44905870, 44905924 ] ): {
      '0.756.741.5' : "ctaccgtttaggcggatatcgcgtagcggcagctactcaacgcaccggaacttgtgtttgtgtgtgtggtcg",
      '0.756.742.0' : "acggcagtagttttgccgctcggtaacgctcagaatgtttgaggggcggtacggctagagatatcacctc",
      '0.756.742.1' : "acggcagtagtnntgccgctcggtaacgctcagaatgtttgaggggcggtacggctagagatatcacctc",
      ...
      '0.756.743.0' : "gtcgcccactacgcacggttatatgagagagctggcagatgccttatggaagtgactgctaccgtttgttga",
    },
    ( 'exon2', [ 'hg38', 'chr19', 1, 44906603, 44906668 ] ): { ... }
  },
  'tag-length' : 24
}
```

8.6 Versioning

URL Parameter Versioning: The client specifies the version as part of the URL path:

```
GET /v0.1.0/status HTTP/1.1
Host: lightning.curoverse.com
Accept: application/json
```

For more information and implementation details, reference <http://www.django-rest-framework.org/api-guide/versioning/>.

8.7 Paging

Limit Offset Pagination: The client specifies the limit and offset using request query parameters

limit: indicates the maximum number of items to return. It's not required. Default limit is 100. Maximum limit is 1000.

offset: indicates the starting position of the query in relation to the complete set of unpaginated items

For information and implementation details, reference <http://www.django-rest-framework.org/api-guide/pagination/>.

SOFTWARE DEVELOPMENT KITS

This section is here for completeness: we currently do not have sdks available for Lightning.

SPRITE

Sprite is a django web application that provides visualizations and easy interaction with Lightning, as well as preliminary phenotype and tile variant annotation databases.

Todo

- Specify phenotype database
 - Specify annotation database(s)
-

Note: Please keep in mind Lightning is under development, as is its documentation. Feel free to file bugs or documentation errors at <https://dev.arvados.org/projects/lightning>.

Todo

Make sure <https://dev.arvados.org/projects/lightning> is public and a well-behaved redmine account.

For an introduction, basic information about using Lightning, and Lightning's design document, see [Getting started](#).

For a description of the process of tiling (the abstraction of genomic sequences that makes Lightning possible) and the functions and pipelines we provide for tiling genomes, see [Tiling Overview](#).

For a description of our representation of tiled genomes (Compact Genome Format), see [Compact Genome File \(CGF\) Format](#).

For a description of our in-memory database for tiled genomes (Lantern), including the REST API it supports, see [Lantern Specifications](#).

For a description of our in-memory database for tile variants (the tile library), see [Tile Library](#).

For a description of Annotile, the way we support annotations of tile variants, how to import annotation software, and how to add human-generated annotations, see [Annotile: Annotating Tile Variants](#).

For a description of Lightning Data Structures, used for interaction with the Lightning APIs, see [Data Structures Specifications](#).

For a description of Lightning APIs, see [Lightning API Specifications](#).

For a description of Lightning Software Development Kits, see [Software Development Kits](#).

Finally, for a description of the web browser application that runs on Lightning (Sprite), see [Sprite](#).

INDICES AND TABLES

- genindex
- modindex
- search