
HBP Curation Documentation

Release 0.1.1-dev

Charles Tapley Hoyt

Oct 15, 2018

General Guidelines

1	Style Guide for BEL	3
1.1	Division of Content into Documents	3
1.2	Document Metadata	3
1.3	Organization of Terminologies	4
1.4	Usage of Short vs. Long Form	4
1.5	Usage of SET STATEMENT_GROUP	4
1.6	Proper Spacing	5
2	Cookbook	7
2.1	Post-Translational Modifications	7
3	Chemical Biology Curation Guidelines	9
3.1	Inhibitors	9
3.2	Receptor Binding	10
4	Encoding of Genetic Information	13
4.1	Epigenetics	13
4.2	Single Nucleotide Polymorphisms (SNPs)	13
5	Re-curation Guidelines	15
5.1	Normalizing Entities	15
5.2	Checking Correctness	16
6	Curation using INDRA	17
6.1	1. Identification of target entities	17
6.2	2. Gather statements from INDRA	17
6.3	3. Preparation of Curation Sheets	17
6.4	4. Curation	18
6.5	5. Re-Curation	18
7	Curation Guidelines for INDRA errors	19
7.1	Gene-centric	19
7.2	Relationship	19
7.3	Identification	19
7.4	Negation	20
7.5	Subject/object	20
7.6	Not Evidence	20

7.7	Site of Modification	20
7.8	Physical Contact Missing	20
8	Rational Enrichment	21
8.1	1. Assemble BEL from Relevant Sources	21
8.2	2. Identify Low-Information Nodes	21
8.3	3. Filter Low-Information Nodes	22
8.4	4. Curation	22
9	Curation with Git	23
9.1	1. Make a Branch	23
9.2	2. Commits with Incremental Progress	23
9.3	3. Merge	24
10	Indices and tables	25

This folder contains the BEL curation procedures and guidelines developed and used during the [Human Brain Pharmacome](#) project.

This document describes style guidelines for BEL. It was written with inspiration by the pragmatism and existence of the [PEP8](#) guidelines.

1.1 Division of Content into Documents

Each statement in BEL is an atomic piece of knowledge, and combine with annotations and provenance information makes a nano-publication. This header addresses the issue of how to organize that information into several .bel files.

Simply, each BEL document should represent the contents of one article. There may be reasons to include multiple articles in a single BEL document if there is crucial supporting information, but the task of assembling BEL for analysis is not the task of the curator.

One example where curation intervention was helpful in defining criteria was the use of the “Subgraph” annotation in NeuroMMSig, which sliced a large knowledgebase related to Alzheimer’s disease into several discrete subgraphs corresponding to biological pathways/mechanisms.

As an added benefit, the one-to-one correspondence of BEL scripts to citations makes the management of curators much easier since files will generally not conflict. This also encourages the externalization of list annotations for reuse.

1.2 Document Metadata

1.2.1 Versioning

BEL documents that are manually generated (as opposed to dumps of databases such as DrugBank) should use version numbers following [Semantic Versioning](#). A correct example, using the <MAJOR>.<MINOR>.<PATCH> format:

1.2.2 Authorship

Authors should be set comma-separated in alphabetical order by last name using

In the description, the contributions of each author can be listed. Some suggested roles are “curation”, “supervision”, “quality control”.

1.2.3 Contact Info

Consider that the authors of a BEL document and the responsible person for the integrity and correctness of the document might not be the same person. For example, this could be due to people moving to new projects. Only the person responsible for a given BEL document should list their contact information in the `SET DOCUMENT ContactInfo` field.

1.3 Organization of Terminologies

The term “terminologies” is used to refer to both BEL namespaces and BEL annotations in this section.

Terminologies’ keywords should use an uppercased version of their corresponding entry in Identifiers.org, when possible. Dots and dashes in resource names are removed for BEL, since they are not considered valid characters for use in keywords. Example: `ec-code` becomes `ECCODE`.

Namespaces should be listed first (interspersed URL and PATTERN definitions), then annotations (interspersed URL and PATTERN definitions), then annotations defined by lists. Within each group, all terminologies should be listed in alphabetical order by the keyword used.

Terminologies with multiple parts, like MeSH and GO, should **NOT** be split into multiple namespaces (e.g. `MESHD`, `MESHCS`, `MESHG`, `GOBP`, `GOCC`, `GOMF`). Update versions of these namespaces can be found at <https://github.com/pharmacometa/terminology/tree/master/external> and versioned using the git commit hashes. The following namespaces are already available:

Note, while `GFAM` is used for `hgnc.genefamily` for brevity, this isn’t *really* recommended.

1.4 Usage of Short vs. Long Form

All BEL functions (e.g., `proteinAbundance()`, `abundance()`, `pathology()`, etc.) should be abbreviated to the short forms (e.g., `p()`, `a()`, `path()`, etc.).

All BEL transformations (i.e., `activity()`, `translocation()`, and `reaction()`), as well as their specific arguments (i.e. `molecularActivity()`, `fromLocation()`, etc.) should be abbreviated to the short forms (i.e. `act()`, `tlloc()`, and `rxn()`).

All BEL relationships should be abbreviated with their short forms.

BEL is quite verbose - the theme is to always abbreviate when possible.

1.5 Usage of `SET STATEMENT_GROUP`

`STATEMENT_GROUP` is listed in the BEL specification as a privileged annotation - it does not need to be defined, and it can be set to anything without semantic validation.

Because it neither has inherent meaning, nor community practices ascribed to it, it is explicitly discouraged to use this annotation.

Some curators use the `STATEMENT_GROUP` to give information about who the curator was, or a certain “sprint” of curation, but these should already be addressed by the earlier point on the organization of BEL documents.

1.6 Proper Spacing

Ensure proper spacing. Without it, BEL is difficult to read and assess.

TODO develop a linter for continuous integration checking!

1.6.1 Spacing in BEL Terms

The following protein with a post-translational modification is difficult to read because there is no space between the comma following the identifier and the `pmod()` function:

The same, with proper spacing applied:

The same applies for all other variants (`sub()`, `frag()`, `loc()`, etc.) and other functions in which commas are applied. The following is another example in which the spacing between the comma following the identifier is correct, but the contents of the `pmod()` are not:

The same, with proper spacing applied:

1.6.2 Spacing in Annotations

The following single annotation is difficult to read because there are no spaces between 1) the annotation and the equals sign and 2) the equals sign and the value:

The same, with proper spacing applied:

The following multiple annotation is difficult to read, because there no spaces between 1) the annotation and the equals, 2) the equals and the open bracket, and 3) the entries within the brackets.

The same, with proper spacing applied:

1.6.3 Citation Information

Citations should be written succinctly when referring to databases like PubMed, PubMed Central and DOI. The remaining citation information can be looked up programatically after.

The same, with proper terseness:

2.1 Post-Translational Modifications

The control of post-translational modifications can be represented in two ways in BEL:

The qualitative representation is preferred, also with other modifications like truncations and fragmentations.

Chemical Biology Curation Guidelines

This document contains an initial (and as-of-yet incomplete) set of guidelines for representing quantitative information in BEL. They do not require any special extensions to its syntax, and should be compatible with any of the available parsers/compiler.

There is currently a BEL enhancement proposal for BEL 2.0.0+ for [native support of numeric annotations](#) that would benefit from public input, so that's highly encouraged!

3.1 Inhibitors

This example will focus on the ability for lovastatin (CHEMBL1487) to inhibit human HMG-CoA reductase (CHEMBL402, UniProt:P04035, HGNC:HMGCR, HGNC:5006).

3.1.1 Simple Representation

3.1.2 Medium-Granular Representation with BEL Default Namespace

3.1.3 Specific Representation

Using BEL 2.0, the molecular function HMGCR that lovastatin inhibits is its (hydroxymethylglutaryl-CoA reductase (NADPH) activity ([GO:0004420](#)).

In general, it might not be so obvious how specific of a GO term to choose. Additionally, a protein may have multiple functions. Complementary to GO is the ExPASy enzyme classification, which is also encoded in ChEBI and can be automatically added to BEL.

3.1.4 Assay Metadata

Taking inspiration from the ChEBML schema, several pieces of metadata can make inhibition experiments more useful:

1. Target Type (e.g., cell line, organism, single protein, complex, etc.)
2. Measurement Type (e.g., IC50, pIC50, EC50, pEC50, Ki)
3. Measurement Units (e.g., pM, nM, μ M, mM, M)
4. Measurement Relation (=, >, <, >=, <=, ~)
5. Measurement Value (floating point matching the regular expression: `^[+-]?[0-9]*\.[0-9]+([eE][+-]?[0-9]+)?`)
6. Assay Type (see below)
7. Cell line, target organism, and/or species

In some cases, the measurement value may be reported as a range. In these situations, use two complementary annotations for Measurement Range Lower and Measurement Range Upper.

Assay Type (adapted from ChEMBL)

- Binding (B) - Data measuring binding of compound to a molecular target, e.g. Ki, IC50, Kd.
- Functional (F) - Data measuring the biological effect of a compound, e.g. %cell death in a cell line, rat weight.
- ADMET (A) - ADME data e.g. t1/2, oral bioavailability.
- Toxicity (T) - Data measuring toxicity of a compound, e.g., cytotoxicity.
- Physicochemical (P) - Assays measuring physicochemical properties of the compounds in the absence of biological material e.g., chemical stability, solubility.
- Unclassified (U) - A small proportion of assays cannot be classified into one of the above categories e.g., ratio of binding vs efficacy.

After adding this metadata, we get:

3.1.5 Provenance

Using a valid citation that points to the original source of the information is preferred to using a reference to the database from which the relation comes. Example: it's better to use [PMID:2153213](#) in these examples referring to its original publication rather than citing ChEMBL.

However, data often comes in a table, and won't have a real evidence text. It's not *exactly* clear whether BEL requires evidences for each statement, so for now a placeholder string saying "Retrieved from X" and an additional annotation called `Database` set to X will allow forward-compatibility. Use an `identifiers.org` namespace whenever possible for X.

Finally, with both the assay metadata and provenance, we get:

3.2 Receptor Binding

This example will focus on the binding of zolpidem (ChEMBL911) to the GABA receptor alpha-5 subunit (ChEMBL5112, UniProt:A8K338, HGNC:GABRA5)

The binding of a chemical to a receptor is represented by the chemical causing a complex with the protein. Binding is typically measured with Ki.

Zolpidem is not a very strong binder to the GABA receptor alpha-5 subunit, so it is unlikely we'll find an annotation as to its binding type.

3.2.1 Binding Type

- *Full agonists* are able to activate the receptor and result in a strong biological response. The natural endogenous ligand with the greatest efficacy for a given receptor is by definition a full agonist (100% efficacy).
- *Partial agonists* do not activate receptors with maximal efficacy, even with maximal binding, causing partial responses compared to those of full agonists (efficacy between 0 and 100%).
- *Antagonists* bind to receptors but do not activate them. This results in a receptor blockade, inhibiting the binding of agonists and inverse agonists. Receptor antagonists can be competitive (or reversible), and compete with the agonist for the receptor, or they can be irreversible antagonists that form covalent bonds (or extremely high affinity non-covalent bonds) with the receptor and completely block it. The proton pump inhibitor omeprazole is an example of an irreversible antagonist. The effects of irreversible antagonism can only be reversed by synthesis of new receptors.
- *Inverse agonists* reduce the activity of receptors by inhibiting their constitutive activity (negative efficacy).

To Do:

- add full agonist example
- add partial agonist example
- add antagonist example

3.2.2 Allostery

In general, if allostery is not set, then it is assumed to be None.

Allosteric modulators do not bind to the agonist-binding site of the receptor but instead on specific allosteric binding sites, through which they modify the effect of the agonist. For example, benzodiazepines (BZDs) bind to the BZD site on the GABAA receptor and potentiate the effect of endogenous GABA.

- Positive allosteric modulator
- Negative allosteric modulator

Source: [https://en.wikipedia.org/wiki/Receptor_\(biochemistry\)](https://en.wikipedia.org/wiki/Receptor_(biochemistry))

When the binding type is set, we can also write a second statement with how the binding affects the activity of the receptor.

Basmisanil (ChEMBL3681419) is an inverse agonist of the GABA receptor alpha-5 subunit (UniProt:A8K338).

Encoding of Genetic Information

BEP6 introduces a proper syntax for representing epigenetic modifications, such as methylation for BEL 2.0.0+. The syntax follows the same style as the protein modification syntax, and can follow the identifier for a gene. A reference implementation has been included in PyBEL.

4.1 Epigenetics

The gene modification function, `gmod()`, as a syntax for encoding epigenetic modifications. Its usage mirrors the `pmod()` function for proteins and includes arguments for methylation.

The options for `gmod()` are currently:

- `M`, `Me`, and `methylation` all refer to methylation on the given gene
- `A`, `Ac`, and `acetylation` all refer to acetylation on the given gene

4.2 Single Nucleotide Polymorphisms (SNPs)

In general, a single nucleotide polymorphism (SNP) refers to a variant in a genetic sequence. The *de facto* identifiers for these variations are the RS numbers from dbSNP. A given identifier can point to two types of information: intragenic SNPs and intergenic SNPs.

4.2.1 Intragenic SNPs

A variation in the sequence of a protein-coding gene can have the consequence of an amino acid substitution or differential expression of the gene. In these cases, it's important to explicitly code a BEL statement linking a SNP to its gene with the variation to which it refers. Since BEL 2.0, variants can be encoded according to the HGVS nomenclature.

A variant can be named multiple ways depending on the “reference sequence” used. While it is possible to refer to a variation by the amino acid sequence or the chromosomal sequence, it is much easier to interpret biologically when

using the gene reference sequence. Given an `RS` identifier, dbSNP lists the reference sequence identifier and the HGVS string based on many of these sequences. More specifically, reference sequence identifiers starting with `NM` or `XM` refer to the genetic sequence (and should all have the same HGVS string), while the `NC` identifiers refer to the chromosome sequence and should be disregarded.

In the future, there might be a way to automate this procedure, but as a curator wants to encode intragenic SNPs, they should also make this equivalence explicit. These statements can be grouped together in a citation to dbSNP, the evidence can be dummy text, and the confidence level can be set with `SET Confidence = "Axiomatic"`.

First, dbSNP can be included using a regular expression definition of a namespace, since there are potentially billions of enumerated SNPs. This is a BEL 2.0.0+ feature that was proposed in [BEP5](#). Identifiers.org lists the database information at <https://www.ebi.ac.uk/miriam/main/datatypes/MIR:00000161>, and includes a regular expression that all accession numbers follow. This can be included as

The following is an example from https://www.ncbi.nlm.nih.gov/projects/SNP/snp_ref.cgi?rs=1235 that uses the `equivalentTo` relationship proposed by [BEP4](#) for BEL 2.0.0+ to link a dbSNP entry to the HGVS nomenclature applied to a gene.

A reference implementation is provided by PyBEL, in which the `hasVariant` relationship `g (HGNC:MDGA2) hasVariant g (dbSNP:rs1235)` is automatically added by the compiler.

An example from a genome-wide association study:

4.2.2 Intergenic SNPs

Some SNPs are not directly part of coding genes' regions. For these SNPs, it is not necessary to encode a relationship to a gene.

However, this also means that their functional consequences do not follow directly from the causal relationships connected to a particular gene. It must be kept in mind that these SNPs will need further qualifications to become useful, such as associations from LD-Block analysis or other studies from eQTL, etc.

4.2.3 LD-Block Information

Linkage disequilibrium (LD) block analysis find SNPs that co-occur together. These relationships can be inferred from data-driven approaches.

TODO

4.2.4 eQTL Information

Expression quantitative trait loci (eQTLs) connect variants to gene expression patterns.

TODO

Re-curation Guidelines

These guidelines were originally written for the re-curation of several NeuroMMSig subgraphs in the Alzheimer's Disease Knowledge Assembly during the Human Brain Pharmacome project, but may be generally applicable to other BEL scripts as well.

5.1 Normalizing Entities

5.1.1 Chemicals

1. Normalize chemical entities to preferred namespaces (*ChEBI*, *ChEBML*, *PubChem*) whenever possible. *MeSH* is explicitly discouraged because it is difficult to look up their structures as SMILES or InChI, even with resolving services like UniChem
2. Formalize knowledge about chemicals that have not yet been encoded in *ChEBI* (such as Selventa chemicals [*SCHEM*], the BELIEF chemical namespaces, etc.), drawing from other public resources (*PubChem*, *MeSH*, *CAS*, *ChEMBL*, UniChem, BridgeDB, ChemSpider, etc.) whenever possible

5.1.2 Protein Complexes and Families

FamPlex has emerged as a resource that maps families and protein complexes (including the Selventa mappings *SFAM* and *SCOMP* as well as other widely used namespaces like PFAM and InterPro).

1. Normalize all entities to FamPlex
2. Formalize knowledge about new families by making a pull request to FamPlex. See: <https://github.com/sorgerlab/famplex#contributing>

5.1.3 Other Entities

We are also building a terminology at <https://github.com/pharmacome/terminology>. This should not be done lightly, so see its [contribution guidelines](#) and [rules](#) before making a pull request.

5.2 Checking Correctness

- If statement can be asserted from the given evidence, add the annotation `SET Confidence = "High"`
- If the statement is wrong, fix it and add the annotation `SET Confidence = "Medium"`
- If it's not clear what BEL should represent the biology, add `SET Confidence = "Low"` for later discussion
- If the evidence string contains no reasonable biological knowledge/is nonsense, delete it and the related statements entirely. It's okay to remove BEL statements that are not supported.

5.2.1 Finalization

After all relevant statements have been checked for correctness, the curation leader will check all statements with `SET Confidence = "High"` or `SET Confidence = "Medium"` and change to `SET Confidence = "Very High"` if they agree. If they do not agree, they will fix it themselves.

Curation using INDRA

This document describes a procedure for using [INDRA](#) for acquiring automatically extracted relations in BEL for curation.

6.1 1. Identification of target entities

Curation begins from a seed list of entities. Often, this will be genes of interest to a project or based on a rational assessment of information density in a network.

6.2 2. Gather statements from INDRA

After acquiring statements from the [INDRA REST Database](#), several filters are made to the resulting

1. Statements containing chemicals are filtered. They are more difficult to curate because of the heterogeneity of chemical nomenclatures and identifiers, and many times can be replaced by the use of databases like ChEMBL, CTD, PubChem, etc.
2. Confidences are calculated for each statement based on INDRA's BELIEF engine, which takes into account high-quality material that has been previously curated in Pathway Commons, Selventa's Large BEL Corpus, and other sources. Statements of different levels of granularity are linked, and confidences can be calculated.
3. When several text-based evidences are available for a given statement, one is chosen randomly so as to avoid curator fatigue.

6.3 3. Preparation of Curation Sheets

The filtered content is exported to a CSV file containing the following columns:

1. INDRA UUID. A unique identifier assigned to each statement in the INDRA database
2. INDRA Confidence. The estimated percentage (0.0 - 1.0) for how correct a statement is

3. Text Reference. The PubMed identifier of the article from which the statement was automatically extracted
4. Text. The text from which the statement was extracted
5. INDRA API. The text mining system used by INDRA to extract the statement
6. BEL Subject
7. BEL Predicate
8. BEL Object
9. Checked
10. Correct
11. Changed
12. Annotations

6.4 4. Curation

Each statement should be read and assessed by the curator, then an “x” should be placed in the “Checked” column.

- If the statement was correct, an “x” should be placed in the “Correct” column.
- Otherwise, the statement should be fixed (assignment of entity types, relation, etc.) and an “x” should be placed in the “Changed” column.
- If the statement is total nonsense, then no checks should be placed in either the “Correct” or “Changed” columns. See the guidelines for curating information about errors on <https://github.com/pharmacome/curation/blob/master/indra-errors.rst>

If there are other BEL that can be extracted, make a new line with all of the same provenance information (uuid, reference, evidence, etc.) and just place an “x” in the “Changed” column.

If there are any annotations (cell type, species, cell line, tissue, experimental context, etc.) that are obvious from the evidence, then they can be denoted in the “Annotations” column using the BEL idiom of *SET ANNOTATION* ... statements.

6.5 5. Re-Curation

Finally, the statements should be exported to BEL and checked using the [re-curation procedure](#).

Curation Guidelines for INDRA errors

This document outlines the most common type of errors identified in the BEL statements extracted from INDRA.

When they're encountered:

1. Add a column with the label 'Error Type' to your document
2. Whenever you find a type of error mentioned in the table, please put the type of error in the column
3. If the error does not correspond to any of those categories, add a new error type to this table and give an example

7.1 Gene-centric

- **Knock Down.** A knock down of the gene is not labelled as such.
- **Gene skipping.** Exon/intron regulation is not labelled as such.
- **Promoter activity.** Missing promotor activity labels.

7.2 Relationship

- **Target.** Labelled as target but not clear from the evidence.
- **Modulate.** Labelled as modulation but not clear from the evidence.
- **Regulate.** Labelled as regulation but not clear from the evidence.
- **Mediate.** Labelled as mediation but not clear from the evidence.

7.3 Identification

The Named-entity recognition system labels an entity wrong. For example, confusing an abbreviation with a different meaning (e.g., FIP means USF2 but is also a recombinant protein (FasL Interfering Protein))

7.4 Negation

- **But not.** False positive negation.
- **Negative mediator.**

7.5 Subject/object

Subject and object are correctly labelled but there is not relationship between them.

7.6 Not Evidence

The evidence information is not sufficient to code the BEL statement.

7.7 Site of Modification

The entity was labelled with the wrong modification site.

7.8 Physical Contact Missing

The relationship was labelled as directed but the evidence does not show that.

Rational Enrichment

This document describes an approach to enrich knowledge assemblies based on topological novelty.

8.1 1. Assemble BEL from Relevant Sources

1. Assemble BEL from relevant sources
2. Re-curate content with questionable quality following the [re-curation procedure](#)
3. *Optional*: choose annotations that are relevant and filter the resulting network. During the initial stages of the Human Brain Pharmacome project, we used the “Subgraph” annotation from NeuroMMSig to select our [ten highest priority candidate mechanisms](#) of pathogenesis of Alzheimer’s disease.

8.2 2. Identify Low-Information Nodes

The information density of each node is assessed to provide an un-biased prioritization list of genes/proteins.

1. Chemicals and higher-order processes (e.g., biological processes, phenotypes, pathologies, clinical measurements, etc.) are filtered from the resulting network. From a molecular-biology point of view, these are the logical inputs and outputs to a biological system.
2. Proteins, mRNAs, miRNAs, and all variants are collapsed to their corresponding genes’ nodes.
3. The information density of each node is defined as the *degree*: the sum of the in-degree and out-degree.

Nodes with a degree of 0 necessarily had links to chemicals or higher-order processes that were removed in Step 1. They are interesting because there is not yet mechanistic information linking them to other genes and proteins that have been curated in the context of the network. Nodes with a degree of 1 or more have an increasing amount of information, and can be ranked such that the lowest degrees are prioritized for curation.

8.3 3. Filter Low-Information Nodes

Some of the nodes with low information have already been quantified mechanistically, but just not in the given knowledge assembly. In this step, several sources (Pathway Commons, Bio2BEL, etc.) are used to enrich these nodes to determine if they have already been curated. Those that are found in other sources are also excluded from curation.

8.4 4. Curation

Several steps can be taken to curate content after prioritizing genes.

1. [Curation with INDRA](#)
2. Manual search of literature for full-text. Eventually, storing the resulting PDFs in services like Mendeley will result in good automatic recommendations. Curation and [re-curation](#) can be done on prioritized papers.

This document describes a workflow for doing BEL curation using git for version control. It's inspired by the [Git Flow](#) philosophy.

The purpose of this document is not to explain what git is or how to use it. Lots of fantastic resources exist on the internet for this - here's a good place to start: <https://guides.github.com/activities/hello-world/>.

9.1 1. Make a Branch

The main/master branch of the git repository is protected in order to encourage assessment by multiple users.

1. Check out the master branch
2. Fetch and pull from the origin
3. Make a new branch with a descriptive (but succinct) name of what you'll be working on. If this branch corresponds to adding a new article, title the branch with the author's last name and publication year prefixed with `curation-`. Example: `curation-olsen2016`

9.2 2. Commits with Incremental Progress

Commit early and often. Each commit should describe what's being worked on, and follow best community practices for writing the commit. See: <https://chris.beams.io/posts/git-commit/>. In the future, pre-commit hooks could be used to enforce this policy.

Between commits, locally compiling the BEL document with a compiler like [PyBEL](#) to find errors is helpful as opposed to waiting until the whole curation task is finished. Other solutions using continuous integration to take the hassle out of installing and running a BEL compiler are also publicly available. For example, see: <https://github.com/cthoyt/pybel-git>.

9.3 3. Merge

If curation is being done on GitHub, make a [pull request](#). On GitLab, make a [merge request](#). This is a place where the final results can be checked one more time for syntactic/semantic correctness with PyBEL, and discussion between curators and managers can occur.

Each merge/pull request needs a name suitable for the public commit log, since the multiple commits will be combine (through a process known as squashing) before merging onto the master branch.

CHAPTER 10

Indices and tables

- `genindex`
- `modindex`
- `search`