The CDF-5 file format will look quite similar to the CDF-1 file format (see http://www.unidata.ucar.edu/software/netcdf/netcdf-4/newdocs/netcdf.html#NetCDF-Classic-Format)

- We define a new magic number version 'CDF-5'
- NON_NEG and STREAMING fields will be 8 bytes instead of four.
- we define a new 64 bit integer type (NC_INT64: a 64 bit signed integer)
- number of dimensions will be specified with a 64 bit field (Is this a ludicrous number of dimensions? Wanted to make sure we had all limits fully relaxed, but maybe 4 bytes is quite enough?)
- number of attributes and variables will be specified with an INT64 (not INT)
- 'padding' can be arbitrary -- not just to nearest 4-byte boundary but to possibly stat->st_blocksize or some other potentially multi-kilobyte boundary

We would make the following modifications to the formal BNF CDF-1 definition:

```
VERSION_BYTE:='\001'|    //File format version number
                         //for netCDF classic Format
              '\002'|    //File format version number
                         //for 64-bit offset Format
              '\005'     //File format version number  *new*
                         //for 64-bit data Format
NON_NEG := INT with non-negative value | // for classic format and
                                           64-bit offset format
           INT64 with non-negative value // for 64-bit data format *new*

OFFSET  := INT with non-negative value | // for classic format or
           INT64 with non-negative value // for 64-bit offset format
                                         // and for 64-bit data format
                                         // *new*
```

These file format changes must be accompanied by API changes:

- addressing data (start, count, stride) will have to be arrays of 64 bit values, not arrays of integers (which are often 32 bits).
- Maximum numbers of dimensions, attributes, and other limits will be relaxed.