

Coding Practices for Programming

Section A: Linux kernel style based coding

Coding style is based on Linux Kernel coding style. The following excerpts in this section are mostly taken *as is* from an article written by Greg Kroah-Hartman titled, “*Documentation/Coding Style and Beyond*” (available on Internet and on the ICP website) and <http://pantransit.reptiles.org/prog/CodingStyle.html>. It is required that the relevant parts of the articles are clearly understood. Given below is the condensed set of rules:

Indentation

Tabs are 8 characters, and thus indentations are also 8 characters.

Rationale: The whole idea behind indentation is to clearly define where a block of control starts and ends. Especially when you've been looking at your screen for 20 straight hours, you'll find it a lot easier to see how the indentation works if you have large indentations.

Now, some people will claim that having 8-character indentations makes the code move too far to the right, and makes it hard to read on a 80-character terminal screen. The answer to that is that if you need more than 3 levels of indentation, you most likely made a mistake and should fix your program.

In short, 8-char indents make things easier to read, and have the added benefit of warning you when you're nesting your functions too deep. Heed that warning.

Placing Braces

All blocks (including `if`, `for` and `while`) except functions are in K&R style as:

```
if (x == y) {  
    ..  
} else if (x > y) {  
    ...  
} else {  
    ....  
}
```

Functions are written as:

```
int function(int x)  
{  
    body of function  
}
```

Naming

Mixed-case names (like `ThisVariableIsATemporaryCounter`) are frowned upon, descriptive names for global variables are a must. To call a global function "foo" is a shooting offense.

Global variables (to be used only if you really need them) need to have descriptive names, as do global functions. If you have a function that counts the number of active users, you should call that "`count_active_users()`" or similar, you should not call it "`cntusr()`".

Encoding the type of a function into the name (so-called Hungarian notation) is forbidden - the compiler knows the types anyway and can check those, and it only confuses the programmer.

Local variable names should be short, and to the point. If you have some random integer loop counter, it should probably be called "i". Calling it "`loop_counter`" is non-productive, if there is no chance of it being misunderstood. Similarly, "`tmp`" can be just about any type of variable that is used to hold a temporary value.

Functions

Functions should be short and sweet, and do just one thing. They should fit on one or two screenfuls of text (the ISO/ANSI screen size is 80x24, as we all know), and do one thing and do that well.

The maximum length of a function is inversely proportional to the complexity and indentation level of that function. So, if you have a conceptually simple function that is just one long (but simple) case-statement, where you have to do lots of small things for a lot of different cases, it's ok to have a longer function.

However, if you have a complex function, and you suspect that a less-than-gifted first-year high-school student might not even understand what the function is all about, you should adhere to the maximum limits all the more closely. Use helper functions with descriptive names (you can ask the compiler to in-line them if you think it's performance-critical, and it will probably do a better job of it that you would have done).

Another measure of the function is the number of local variables. They shouldn't exceed 5-10, or you're doing something wrong. Re-think the function, and split it into smaller pieces. A human brain can generally easily keep track of about 7 different things, anything more and it gets confused. You know you're brilliant, but maybe you'd like to understand what you did 2 weeks from now.

Commenting

Comments are good, but there is also a danger of over-commenting. **Never** try to explain **how** your code works in a comment: it's much better to write the code so that the working is obvious, and it's a waste of time to explain badly written code.

Generally, you want your comments to tell **what** your code does, not **how**. Also, try to avoid putting comments inside a function body: if the function is so complex that you need to separately comment parts of it, you should probably go back to **Functions** section for a while. You can make small comments to note or warn about something particularly clever (or ugly), but try to avoid excess. Instead, put the comments at the head of the function, telling people what it does, and possibly **why** it does it.

The short function description cannot run multiple lines, but the other descriptions can (and they can contain blank lines). All further descriptive text can contain the following markups:

```
funcname(): function
$ENVVAR: environment variable
&struct_name: name of a structure (up to two words, including struct)
@parameter: name of a parameter
%CONST: name of a constant
```

So a simple example of a function comment with a single argument would be:

```
/**
 * my_function - does my stuff
 * @my_arg: my argument
 *
 * Does my stuff explained.
 */
```

Comments can and should be written for structures, unions and enums. The format for them is much like the function format:

```
/**
 * struct my_struct - short description
 * @a: first member
 * @b: second member
 *
 * Longer description
 */
struct my_struct {
    int a;
    int b;
};
```

typedef

typedef should **not** be used in naming any of your structures. Most main kernel structures have no typedef. Using a typedef only hides the real type of a variable. There are records of some kernel code using typedefs nested up to four layers deep, preventing the programmer from easily telling what type of variable they are actually using. If the programmer does not realize the size of the structure it can cause very large structures.

Section B: Building and testing of a program

Never begin to start keying-in the program unless pseudocode is written.

If possible, make a dry run over the pseudocode and see if it at least **seems** to do what is required of it.

Divide your program into units that can be compiled or checked **while** building the code to ensure things going smoothly. It is a good practice to check that at the end of each unit, one can verify if **progress** in the **correct direction** is being made. The identification of units, hence, themselves is a subjective matter and can be done easily when one has complete idea about the entire problem in the form of a pseudocode.

An example of units for the construction of a program to print a number in Fibonacci series is given below:

1. Build the minimal valid program (*ie*, a program that returns '0').
2. Declare the function template, say `Fib(int i)`, and the function `Fib` that returns '0'.
3. Write the function and check it with a value, say '3'. It is expected to return '2'.
4. Take a number from user and verify the output using the function.

Section C: Sample code

```
/*
 * Program to:
 * (a) add two numbers
 * (b) add marks in various subjects to obtain total
 */

#include <stdio.h>

/**
 * struct rec - stores marks obtained in various subjects
 * @icp: marks obtained in ICP
 * @dfs: marks obtained in DFS
 * @total: total marks
 */
struct rec {
    int icp;
    int dfs;
    int total;
};

int add_num(int *, int *);
void add_marks(struct rec *);

int main()
{
    int b;
    int c;
    struct rec b2020020;

    b=100;
    c=200;
    b2020020.icp=60;    /* marks in ICP */
    b2020020.dfs=70;   /* marks in DFS */

    printf("add_num: %d + %d = %d\n\n", b, c, add_num(&b, &c));

    add_marks(&b2020020);

    printf("rec.total = %d\n\n", b2020020.total);

    return 0;
}

/**
 * add_num - returns sum of two integers obtained through pointers
 * @x: pointer to first number
 * @y: pointer to second number
 *
 * We assume that the sum of the two integers given by the user is a valid
 * and correct integer.
 */
int add_num(int *x, int *y)
{
    return ( *x + *y );
}

/**
 * add_marks - finds rec.total
 * @record: &struct rec
 *
 * We take the student record and add the marks to total
 */
void add_marks(struct rec *record)
{
    record->total = record->icp + record->dfs;
}
```