CSci 4271W (010 Section) Project Instructions

Project 3 **Due: May 1st and May 5th, 2025**

1 Project: Badly Coded Bookmark Manager (BCBM)

As you know from projects 1 and 2, Badly Coded, Inc., (BCI) is a software vendor responsible for keeping security students at UMN busily employed for many years. (Ask your friends! They're known for such incredible products as the "Badly Coded Versioning System (bcvs)" the "Badly Coded Compression System (bczip)", the "Badly Coded Print Daemon (bcpd)", and the "Badly Coded Calendar Alert System (bccal)" to name but a few...) The latest product under development at BCI is the Badly Coded Bookmark Manager (`bcbm`, for short), and the team responsible for this project has approached your group for a security assessment of (the pre-alpha version of) this code.

bcbm is intended to be a cross-browser and cross-device bookmark manager system: it can be run once on a user's device, and accessed from any browser on the device. Using the browser interface, users can add bookmarks, delete them, and order them in a "top 10" list of favorite bookmarks. bcbm can also synchronize a user's bookmarks across multiple devices, by uploading and downloading encrypted copies of the user's bookmarks file to a back-end bcbm cloud server. This synchronization between devices is managed through an email registration process: users enter an email address and password in the client that is submitted to the cloud server when uploading and downloading the bookmarks file; to set up an account, users visit a page on the cloud server, enter their email address, and are sent a confirmation link that can be used to set or reset the password.

Functionally, bcbm itself consists of two codebases:

- The bcbm "client", `bcbmc` is actually a small HTTP application server program (built using the CrowCPP framework) that runs on a user's device. Users connect to the client from a web browser, and the client writes web pages that accept user inputs and contain the results of user actions, which are then displayed in the browser. The client also handles the client-side portion of synchronizing between multiple devices and communicates with...
- A cloud server, `bcbms` for facilitating synchronization between devices. The cloud server program has functionality to transfer encrypted bookmark files, verify user credentials, and register or update credentials using registration codes.

However, a number of other entities are potentially involved: `bcbmc` connects to third-party websites to download website icons and web page titles; the user's interactions are sent through one or more web browsers, and the browser (typically) displays pages from other web sites that could potentially try to issue commands to the client through included resources. As you progress through the stages of the project, you'll see more and more of the code for bcbm, but these other entities, and the communication with them, should also be considered as you think about the potential threats to the client and server.

2 Project 3: The everything bagel

For the purpose of the project, we're considering the bcbm code in three stages. For project 3, we'll finally consider the full client code and server code.

To install the full `bcbmc` client and the full `bcbms` server on your VM (please don't install this on a real machine - it's really buggy code) follow these steps:

1. Clone the source repos using:

```
git clone https://github.umn.edu/badly-coded-alpha/bcbmc-sync.git
git clone https://github.umn.edu/badly-coded-alpha/bcbms-reg.git
```

2. Build the client by changing directory into `bcbmc-sync` and running `./install.sh`; then build the server by changing directory back to `bcbms-reg` and running `./install.sh`.

This will build and install the full `bcbmc` executable, as well as the full `bcbms` server. `bcbms` is installed as a "systemd service", which means the operating system monitors the process, restarting it if it crashes, and it runs as `root`. To start `bcbms` on your VM, type `sudo systemctl start bcbms`, and to stop it run `sudo systemctl stop bcbms`. You can view the output logs of `bcbms` with `sudo journalctl -u bcbms`. The install.sh script also configures your VM (but only your VM) to route the hostname `bcbm.badlycoded.net` back to itself, so only a `bcbmc` instance running on your VM will be able to connect to the server running on your VM.

To interact with `bcbmc`, type `bcbmc` on the command line, and then from another terminal **on your VM** (connected with X forwarding), open a browser (chromium or firefox, say) and navigate to `http://localhost:8888/`. (Note: `localhost` is a name for whatever computer you are currently using; it's a way to tell a program that expects to connect to other computers to connect instead to the computer the tool is running on. So the `localhost` URL will only work to connect to bcbmc when you run the browser on the VM.)

The code for bcbmc is split across several files:

- The main code for the application server is in `bcbmc.cpp`, which uses the CrowCPP framework. The framework itself (the code in `crow_all.h`) should be excluded from your analysis, but you should at least skim the documentation for CrowCPP to help you read and understand the code in `bcbmc.cpp`. The `main` function is in `bcbmc.cpp` and uses Crow to "route" requests to the other modules in the system. It also calls functions to set the six-digit "authentication" code that (attempts to?) prevent(s) other pages/programs from calling the bookmark management functions of `bcbmc`.
- The file `acode.c` contains the code for choosing a random authentication code.
- The file `manage.c` contains most of the code for responding to user requests. For most interactions, a function is called from `bcbmc.cpp` with the inputs provided by the user, and the function writes an html file with results to be displayed. `bcmbc.cpp` then returns this file to the browser for display.
- The file `sync.c` includes all of the code for synchronizing bookmarks across devices. This includes connecting to the cloud server to download the latest archive (consisting of an encrypted bookmark file, public encryption keys for each device connected to the user's account, and the icons for each bookmarked site) of the user's bookmark file; decrypting the encrypted bookmark file using the local decryption key for the device, applying any changes in the local transaction file (`bcbm.tx`) to the cloud version of the bookmark file; re-encrypting the patched bookmark file under a randomly-chosen symmetric key, then encrypting that symmetric key under each device public key, making a new archive, and uploading the new archive to the cloud server.

The code for `bcbms` is also split across several files:

- The main code for the cloud server is in `bcbms.cpp`. Like the client, this is a CrowCPP application that "routes" requests to various back-end functions.
- The `cloudsync.cpp` file handles the cloud side of synchronization, verifying uploaded archives then moving them to a common directory; and returning archives for downloading on request.
- The `regdb.cpp` file handles checking of username/password requests, as well as using an email registration code system to allow users to set/re-set their passwords.

- A `sqlite3` database that records user and site registration information, managed by a series of scripts so that the server process doesn't dynamically construct SQL statements.

Some other notes about interacting with/testing bcbm and bcbms for project two:

- **Email Recipients**: Because of the network setup for the VMs, a server running on your VM can only send emails to accounts on other VMs, of the form `account@csel-xsme-s25-csci4271-NNN.cselabs.umn.edu`. You can access emails sent to, e.g. the `student` account on your VM using the `mail` command – type `man mail` on your VM for more information. (If you need more accounts, the `useradd` command can be used to add extra user accounts, e.g. `student2`, `student3`, ..., on your VM.) Mails sent to external email addresses will silently fail.

- A possible threat that the developers clearly considered is tampering/spoofing of the user's commands by other websites displayed in the browser. This is a possibility because any website can tell the browser to load certain "resources," commonly images, from another website. A simple example of this is the file:

```
<html><img src="http://localhost:8888/666666/add/http%3A%2F%2Fevil.com%2F"></html>
```

  which would cause the browser to send an "add" command for `http://evil.com/` to their local bcbmc client. (Thanks to the randomized authentication string, the command will almost certainly be rejected!) You can experiment with similar files as examples of this type of threat: local files can be loaded in the browser. For testing, you can also use the `curl` command line tool (on your VM) to send a request to bcbmc and observe the response:

```
% curl "http://localhost:8888/666666/add/http%3A%2F%2Fevil.com%2F"
```

  (Since bcbmc itself uses `curl` you might benefit from spending some time reading its man pages)

- bcbmc also interacts with web pages when adding them as bookmarks. For testing purposes, you can place pages in your CSELabs user homepage and add/delete them as bookmarks as you see fit. A PoC exploit using a page from one of these directories should include the code for the page, and a curl command to call the corresponding action in the client.

- **NO SOCIAL ENGINEERING!** The purpose of this project is to find problems inherent in the code of bcbm; so attacks that involve asking the user to run `bcbmc` with odd arguments, or paste strange values into `bcbmc` are out of scope. (However, it is OK to assume that the user might bookmark an attacker-controlled page or load an attacker-controlled page in their browser.)

- While there are several Elevation of Privilege attacks planted in this code, and many more may be possible due to the fact that it is Badly Coded, don't forget to look for other possible STRIDE threats, such as potential spoofing, tampering, information disclosure and denial of service against the client server. Your final report, after all, should include at least two vulnerabilities that are not in the EoP category.

- You can also interact (through a browser) with a bcbmc "client" running on your VM from a laptop connected directly to the UMN network, a CSELabs machine/vole session, or a laptop/desktop connected to the UofM VPN, using "ssh port forwarding." This feature of ssh intercepts network connections sent to a local application and "forwards" them to an application running on the destination host. To do this from, e.g. a CSELabs machine you would connect to your VM as follows:

```
$ ssh -L 8888:csel-xsme-s25-csci4271-NNN:8888 student@csel-xsme-s25-csci4271-NNN
csel-xsme-s25-csci4271-NNN % bcmbc
```

Then start chromium or firefox on the lab machine and navigate to `http://localhost:8888/`.

- You can access the human-usable components of the server running on your from any machine connected to the UMN network, e.g. `https://csel-xsme-s25-csci4271-NNN.cselabs.umn.edu/reg/start`; however, the TLS certificate used by the server is self-signed and for `bcbm.badlycoded.net`, so you'll need to click through the certificate warnings issued by your browser when doing so.
- The use of encryption on the bookmarks file indicates the developers were also concerned with leaking potentially sensitive information about users' bookmarks to the cloud server (for example, imagine if a user had a bookmark to `alcoholicsanonymous.com` or `plannedparenthood.org`). It would certainly be worth investigating whether the measures taken here are effective.
- **Note about DoS:** Since bcbms is a network server, availability is clearly an important feature, and so denial of service is certainly a consideration. However, `bcbms` is installed and started with a script that detects when the server process crashes and restarts the server. So a crash bug that doesn't impact the ability of the server to restart normally and resume functioning is more of a bug than a DoS vulnerability. Also keep in mind that there's nothing a programmer can do to prevent network-level DoS or DDoS: please exclude these from your threat model and vulnerability reports. (The U's OITSec will get quite irate with us if someone tries to DDoS the network for their course project, so don't do that.)

If you have questions about getting `bcbmc` or `bcbms` to work, or about the wording of the assignment, or about finding someone to work with, use the `project` folder on Piazza so that others can benefit from the answers to your questions. If you have any other questions about what is allowed for the project or other details, please ask the instructor or a TA as early as possible.

3 Assignment and Timeline

As also explained on the Project overview page, your assignment is to review and produce a security assessment of the BCBM code so far. For Project 3, you should produce a report that includes (1) a system design section that answers the question "what are we building?", (2) a threat model using STRIDE that outlines the types of threats you evaluated the system against, and (3) a summary of findings that documents and explains at least 6 vulnerabilities, at least one of which should be in the code released for project 3. You will also produce "proof of concept" programs showing that your vulnerabilities can be exploited, and a `README.md` file explaining how to run your proof of concept programs and what to expect. Finally, your report should have a separate section entitled "Addendum: Revisions and Group Accountability" which details which sections were primarily drafted or revised by each group member along with how you have revised the portions of your report based on the feedback you received for Projects 1 and 2 (of course if you're working on individually, this will just cover revisions). To demonstrate the separate efforts of each group member, you will need to individually submit drafts of a section you have been working on before each full group submission. Recall that each student in a group should individually be the primary author, over the course of the semester, of 2 vulnerability subsections, and one design section or revision, and one threat model section or revision, so you should rotate responsibiliies for drafts compared to projects 1 and 2.

3.1 Grading

- As noted on the project information page,

  - 18 points are assigned for including the required elements. Check that you have included: design section, DFD, all DFD elements, threat model section, listing of security goals in the threat model, STRIDE by each element in the DFD, summary of findings, 6 vulnerabilities, mitigation discussions for all vulnerabilities;
  - 5 points are assigned for the per-student section draft;
  - 5 points are assigned for the "Revisions and Group Accountability" addendum; did you address the feedback you received from projects 1 and 2?
  - 28 points are assigned for high-quality writing, as described in a separate writing rubric web page.
  - 18 points are split between the six vulnerabilities in your project report: for each vulnerability, did you clearly describe an actual vulnerability, correctly explain how it could be mitigated, clearly distinguish it from the other vulnerabilities?
  - 8 points are available for the clarity and correctness of the `README.md` file in your proof of concept repo.
  - 18 points are available for the readability, reproducibility, and correctness of all of the proof-of-concept programs in your repo
  - **Extra Credit**: two points of extra credit is available for each additional, correct, vulnerability plus proof of concept (limit 4) above the six required vulnerabilities.

- To re-emphasize an important point from the project information page: **Claiming a vulnerability without confirming your claims about it is academic and professional dishonesty and may result in a grade of 0 for the project**. (An example would be to claim that a stack buffer in the function `gen_auth_code` can be overflowed, resulting in a tampering attack, without giving any evidence that an input from an untrusted source, of the appropriate length, could actually reach the `gen_auth_code` function. If all you are sure of is that the code looks risky but you're not sure an attack is possible, you can get partial credit for saying that honestly.)

3.2 Timeline:

- **By 11:59pm, Thursday, May 1st**: Each member of a group will submit a PDF draft of one section of the report that they have been responsible for. This submission isn't for detailed grading or feedback, but to confirm that each member of the group has been working on part of the report.

- **By 11:59pm, Monday, May 5th**: You should submit your PDF report to Gradescope. The report should also include a clearly labeled reference to a private **UMN** github repo (so, on `github.umn.edu`) that is accessible to o `smccaman` and `foiso001`, where we can find your six (or more) proof of concept programs and `"README.md"` file. Note that since May 5th is the last day of instruction this semester, there will be no extensions for project 3.

Happy Hacking!