

CHASM Skills Check

May 8, 2010

Abstract

This document contains a series of scenarios and questions designed to reveal the extent of your knowledge, but more importantly, how you approach problem solving. Try to answer each question to the best of your knowledge. If you don't feel you can adequately answer a problem, feel free to ask for help. If you ask for help by searching Google etc., make a note of it in your answer.

For most questions we have a list of things that we “must” or “should” see in a response. Most of the items we are looking for are in the “should” category.

1 Systems Coding

1. Describe the differences between blocking and non-blocking IO? Can you point out several uses for each in a project like CHASM?
2. What are the benefits and drawbacks for using threads in large-scale network applications?
3. A friend gives you the following piece of code and tells you that it's a UDP echo server that he is running on his brand-new x86_64 machine. He says that it will mirror all UDP packets received (that is, it will send them back to the client).

What's wrong with this (broken) UDP echo server (ignoring missing includes)? Point out all issues you see with it, and why they are problematic.

```
int main(int argc , char** argv)
{
    // Do some stuff to parse the user's input into a port to listen on.
    // Assume the presence of a 'port' short and a 'host' long that
    // contain the local address and port to bind to.

    int sockfd;
    struct sockaddr_in addr;
    int addrlen = sizeof(addr);
    char buf[512];

    // Create the socket.
    sockfd = socket(PF_INET, SOCK_DGRAM, 0);

    // Bind to the port.
    addr.sin_family = AF_INET;
    addr.sin_port = port;
    addr.sin_addr.s_addr = host;
    bind(sockfd , &addr , addrlen);

    // Listen for up to 5 unhandled connections.
    listen(sockfd , 5);
```

```

while (1)
{
    recvfrom(sockfd, buf, 512, 0, &addr, &addrlen);
    sendto(sockfd, buf, 512, 0, &addr, addrlen);
}
}

```

4. What is the difference between concurrency and parallelism?

2 Project Management

1. What's the difference between forwards compatible and backwards compatible network protocols?
2. You've been selected to start a new project. Your first task is to select a version control system. Unfortunately, none of the available systems are feature-complete. You have the option of picking five of the following seven features to be in the version control system you select for your project.
 - (a) Cheap branching.
 - (b) Speed.
 - (c) Compact checkouts.
 - (d) Has a staging/index area.
 - (e) Distributed.
 - (f) Flexible Workflow.
 - (g) Easy to learn.

Which of the features do you pick, and most importantly, why?

3. Spread throughout our codebase is the “Pimpl” or “Compiler-firewall Idiom” ¹. We selected this as a means of reducing potential ABI changes. What are some drawbacks (perceived or otherwise) that come with this design decision? What are some benefits?
4. In an ideal world, how would you keep a team of three developers in contact with each other when they are geographically distributed? Feel free to mention specific technology (or specific implementations of that technology) that you are most familiar with, and have the most access to.
5. Select two methods of automated testing of source code. Compare and contrast their benefits and deficiencies. Provide examples of tools that fall under each method.

3 Personal Background

1. Describe courses that you have taken that have given you background in any of the following areas (feel free to include areas not mentioned).
 - Network programming
 - Cryptography
 - Systems programming
 - Computer algorithms

A short, one-sentence description of each course (or a link to the syllabus) would be sufficient.

2. How much time will you be able to devote to CHASM this summer if you are selected?

¹<http://www.gotw.ca/publications/mill105.htm>

3. Do you have, or anticipate, any major commitments that would interrupt work this summer?
4. What motivates you to work on open source projects?
5. Do you have any code samples (preferably C or C++) that you can share with us?
6. Would you be available to meet with us on IRC to chat about your answers? If so, what times (specify timezone please) would you be free to meet. We're located UTC-04:00, but at least one of us is up most hours of the day.

Please keep your answers short if possible. We're not looking for essays, and we're not assigning numerical grades. The best way to gauge the experience of an individual is to work with them for several months. Our short time schedule does not allow us this luxury; instead, we hope that the above quiz gives you the opportunity to show us what you know.

University exams are most likely in progress around this time (we know this is the case in the US, but are unsure of elsewhere), so let us know if you're loaded with work, but still are looking to apply (we'll wait for your application).

Best of luck,
The CHASM Team
Robert
Ben
Joe