Markscheme

November 2019

Computer science

Higher level

Paper 3

9 pages

Subject details: **Computer science HL paper 3 markscheme**

Mark allocation

Candidates are required to answer **all** questions. Total 30 marks.

General

A markscheme often has more specific points worthy of a mark than the total allows. This is intentional. Do not award more than the maximum marks allowed for that part of a question.

When deciding upon alternative answers by candidates to those given in the markscheme, consider the following points:

- Each statement worth one point has a separate line and the end is signified by means of a semi-colon (;).

- An alternative answer or wording is indicated in the markscheme by a "/"; either wording can be accepted.

- Words in (…) in the markscheme are not necessary to gain the mark.

- If the candidate's answer has the same meaning or can be clearly interpreted as being the same as that in the markscheme then award the mark.

- Mark positively. Give candidates credit for what they have achieved and for what they have got correct, rather than penalizing them for what they have not achieved or what they have got wrong.

- Remember that many candidates are writing in a second language; be forgiving of minor linguistic slips. In this subject effective communication is more important than grammatical accuracy.

- Occasionally, a part of a question may require a calculation whose answer is required for subsequent parts. If an error is made in the first part then it should be penalized. However, if the incorrect answer is used correctly in subsequent parts then **follow through** marks should be awarded. Indicate this with "**FT**".

- Question 4 is marked against markbands. The markbands represent a single holistic criterion applied to the piece of work. Each markband level descriptor corresponds to a number of marks. When assessing with markbands, a "best fit" approach is used, with markers making a judgment about which particular mark to award from the possible range for each level descriptor, according to how well the candidate's work fits that descriptor.

General guidance

Issue	Guidance
Answering more than the quantity of responses prescribed in the questions	• In the case of an "identify" question read all answers and mark positively up to the maximum marks. Disregard incorrect answers. • In the case of a "describe" question, which asks for a certain number of facts *eg* "describe two kinds", mark the first two correct answers. This could include two descriptions, one description and one identification, or two identifications. • In the case of an "explain" question, which asks for a specified number of explanations *eg* "explain two reasons …", mark the first two correct answers. This could include two full explanations, one explanation, one partial explanation *etc*.

1. (a)   *Award **[2 max]**.*
    *Award **[1]** for each smartphone resource identified up to **[2 max]**.*

   Screen size;
   Memory;
   Processor speed;
   Available primary or secondary storage;
   Battery life;
   GPS Receiver;

   Do not accept any resources that are obviously on the server side
   Do not accept software as a resource.   **[2]**

   (b)   *Award **[2 max]**.*
    *Award **[1]** for identifying why an API may be useful and **[1]** for a brief elaboration*
    *up to **[2 max]**.*

   An API ensures abstraction;
   as the user only how the service is requested, how the service is implemented;

   An API acts as a "contract" or definition of the services of the CAD system;
   which makes it clearer to client application developers and back-end developers;

   The API would ensure arguments (parameters) are passed to the system;
   in the correct order and type);

   An API would make sure that the app is expecting a specific data type;
   which will be returned from a CAD request/service;

   An API provides a consistent method/protocol/interface for building a resource;
   and thus, reduces the amount of programming required/provides structure.

   An API would allow the smartphone app and the CAD system to be "compatible";
   can be easily updated if the situation changes;

   An API ensures authentication/authorisation/security checks are carried out;
   before allowing operations to be done;

   Using an API would allow server to be more secure;
   as unused ports and binary/executables/programs/scripts could be hidden from
   users. (*i.e.* Only the API handler would be accessible, not the program modules);

   Using an API would allow changes to be coordinated between all users;
   and backward compatibility when the app is upgraded;

   An API standardises the operations available across different clients/devices;
   and the level of service would not be as device dependent;

   An API would allow the server to track usage of individuals (through keys);
   To identify when certain users were overusing services;   **[2]**

2. (a) *Award [4 max].*
Award [1] for identifying part of the process that a smartphone uses to obtain its location up to [4 max].

Trilateration;
Uses signal strength and latency based on very accurate atomic clock;
Complex formula taking into consideration the curvature of the earth to arrive at latitude, longitude and altitude;
The smartphone uses the distance between satellites to create overlapping "spheres" that intersect in a circle (which is the phones location);
When fewer satellites are available the accuracy diminishes and may be specified to within a radius only;
Assisted/augmented GPS can help accuracy by adding location information received from or about WIFI networks and the triangulation of cell towers;
The smartphone depends on data received from GPS satellites, but that can be obstructed or distorted by obstruction (*eg* tall buildings), weather or water; **[4]**

(b) The answer could include the following:

A multi-tier approach for developers provides:
Abstraction of layers;
Less complexity;
Faster development time;
Easier testing;
Work can be divided better amongst teams;
Performance – system can be divided and allocated to many processors/nodes;
Easier addition of clients, databases or logic which does not "break" other layers;

The use of standards and protocols provides:
Future proofing;
Interoperability;
Communication with legacy devices;
Compatibility with best of breed 3rd party products (*eg* firewalls, load balancers,
Quality of Service and network monitoring software *etc*).
Quick patching of discovered errors, exploits *etc*;
Standards and protocols may already be implemented and have available code references or libraries which would reduce development time *etc*;
Security updates and upgrades from the community; **[4]**

[1–2 marks]
A limited or superficial response that indicates a basic understanding of one or both approaches. Uses little appropriate subject specific terminology.

[3–4 marks]
A competent explanation of the importance of one or both approaches. There is appropriate use of subject specific terminology throughout the response.

3. *The answer could include the following:*

Redundancy:
- There is another network card in the server which is waiting to take the role of the first.
- The user session information is replicated onto another server so that the other server can take over the operation of that client, should the first one be completely unavailable (*eg* If there is no redundant network card).

Failover:
- At the point failure, the OS realises that the network card is not working and automatically switches to the redundant network card to continue conversation with the client.
- At the point of failure, the load balancer/application server/partner server realises that the server is no longer reachable and from that point onwards all requests by the owner of that session are directed to the server which contains a replicated copy of the original session. The session is then replicated to a third machine.

Replication of session information:
- The information contained in the session of a user can either be stored only in one node of the cluster (sticky sessions), or it can be stored in shared memory or secondary storage accessible by the whole cluster. There are performance overheads, because if the session information is constantly changing, or it has to be accessed with each request then it would usually be much faster to keep it in primary memory in one server (sticky sessions) to which the user is always directed by the load balancer / application server. **[6]**

[1–2 marks]
A limited response that indicates very little understanding of the topic or the reason is not clear. Uses little or no appropriate subject specific terminology. No reference is made to the scenario in the stimulus material. The response is theoretical and descriptive.

[3–4 marks]
A superficial explanation of why this hardware failure would not seriously affect the users who are currently connected to that particular server. There is some use of appropriate subject specific terminology in the response.

[5–6 marks]
A thorough explanation of why this hardware failure would not seriously affect the users who are currently connected to that particular server. Explicit and relevant references are made to the scenario in the stimulus material. There is appropriate use of subject specific terminology throughout the response.

4.   *The answer could include the following:*

- It is very difficult to guarantee anonymity as so many technologies and levels of logging are involved.
- Encryption can be used to make the contents of the communication inaccessible (or at least much more difficult to access).
- Information such as IP address, can be used to trace somebody but would need the cooperation of an ISP *etc*.
- IP Addresses are logged at various different levels: Client OS, ISP, routers on the internet (many of them), host webserver (in a weblog file). The server could be set up to not log this information.
- Logging of IP addresses by the ISP, local router at client's point of origin and routers all along the internet cannot necessarily be controlled or secured to prevent them logging accesses.
- MAC addresses that identify the specific device used are logged by switches.
- VPNs or proxy servers could be used to make the users' origin difficult to obtain.
- IP addresses can be shared between many different computers in one area/organization, using NAT (network address translation) so might not uniquely identify a person.
- The client device itself could be compromised with key loggers, spyware *etc* and therefore even if all transmission risks were removed, someone could still get access to the information and source.
- Various products/services/protocols exist which try to make anonymity easy (such as Tor browser) but this is still an area that is developing
- If the servers are vulnerable to physical attack, then people could forcefully enter the hosting facility and install monitoring software or steal data. Sufficient security, biometric access locks etc. could help to prevent this.
- If the data is backed-up, then this historic data should be anonymised or destroyed.
- How ethically correct is it for the providers of the emergency services to receive what could be hoax calls and have no way of knowing themselves who has made them. Could this be counterproductive and therefore allow people to drain emergency resources and potentially send police to unnecessary operations? The equivalent to this online service would be the "confidential telephone" services available by the police in several countries, but these are not anonymous either (though they don't necessarily pass through an open communications platform like the Internet).

Marks	Level descriptor
No marks	• No knowledge or understanding of the relevant issues and concepts. • No use of appropriate terminology.
Basic 1–3 marks	• Minimal knowledge and understanding of the relevant issues or concepts. • Minimal use of appropriate terminology. • The answer may be little more than a list. • No reference is made to the information in the case study or independent research.
Adequate 4–6 marks	• A descriptive response with limited knowledge and/or understanding of the relevant issues or concepts. • A limited use of appropriate terminology. • There is limited evidence of analysis. • There is evidence that limited research has been undertaken.
Competent 7–9 marks	• A response with knowledge and understanding of the related issues and/or concepts. • A response that uses terminology appropriately in places. • There is some evidence of analysis. • There is evidence that research has been undertaken.
Proficient 10–12 marks	• A response with a detailed knowledge and clear understanding of the computer science. • A response that uses terminology appropriately throughout. • There is competent and balanced analysis. • Conclusions are drawn that are linked to the analysis. • There is clear evidence that extensive research has been undertaken.

[12]

Total: [30]