

CS312 Homework #4 Answer Key

March 2, 2016

Questions

1. Ansible primarily uses a push driven model, however it can also support a pull driven model.
 - (a) **True**
 - (b) False
2. Nagios plugins keep track of core monitoring logic including service states.
 - (a) True
 - (b) **False**
3. NRPE checks are initiated by Nagios
 - (a) **True**
 - (b) False
4. Horizontal scaling means adding more resources to a node
 - (a) True
 - (b) **False**

5. Vertical scaling is easier and more simple than horizontal scaling.
 - (a) **True**
 - (b) False
6. Hosts in Nagios must have at least once service defined
 - (a) **True**
 - (b) False
7. Which Nagios configuration component defines what hosts and services to monitor?
 - (a) Main configuration file
 - (b) Resource files
 - (c) **Objection definition files**
 - (d) CGI configuration file
8. Which consensus algorithm does CoreOS use for `etcd`?
 - (a) Paxos
 - (b) **Raft**
 - (c) Ripple
 - (d) RAID5
9. Which configuration management system is the oldest platform that is still in use today?
 - (a) Chef
 - (b) Puppet
 - (c) Ansible
 - (d) **CFEngine**

10. Describe two resources you might want to monitor on a server. Elaborate on how you might decide on acceptable thresholds for each resource.

This answer can vary but some good resources to check include: server load, disk usage, memory, swap usage, services such as http and ssh, network traffic, and server ping-able to name a few. The thresholds depend on the resource, however generally its a value that allows you to respond before the service goes completely down. You also need to consider rate of change sometimes because of things like log rotation.

11. Name two primary differences between active checks vs. passive checks for monitoring.
 - (a) **Active checks are initiated from a central server while passive checks are initiated from the remote host.**
 - (b) **Active checks require a daemon listening on a port to run on the remote server, while passive checks only require a daemon listening on a port on the central server.**
 - (c) **Active requires a central server to schedule checks to the remote hosts.**
 - (d) **Passive checks can scale much better than active checks.**
 - (e) **Passive checks work behind a firewall easier than Active checks.**

12. Describe what the following command does in detail:

```
/usr/lib64/nagios/plugins/check_nrpe -H foo.example.org \  
-c check_cpu
```

The `check_nrpe` contacts an NRPE agent running on a remote server. In this case it's contacting the server `foo.example.org` and requesting to check the command that is defined as `check_cpu` on the NRPE agent on the remote host. This is an example of an active check.

13. Which contact configuration directive in Nagios configures whether or not a contact will receive notifications about service problems and recoveries?

service_notifications_enabled

14. What are specificity and sensitivity? Give a 1–2 sentence description of each.

Specificity is the true-negative rate. In other words, it is the rate of true negatives compared to total reported negatives. Sensitivity is the true-positive rate, and is the ratio of true positives over total positives.

15. Give an example of a highly sensitive test that has low specificity.

A highly sensitive test has a high true-positive rate. A test with low specificity has a high rate of false-negatives. In other words, the test correctly identifies positives but incorrectly identifies negatives. Specific examples will vary.

16. What is time-series data? Give an explanation of why it is important.

Time series data is a set of vectors consisting of a value and a timestamp. This data is important because it is a common format for displaying and analyzing metrics.

17. Explain the CoreOS update process. What is it modeled after? Do you think this is a good or bad update model? Explain.

The CoreOS update process is a dual-partition scheme based on Omaha, the update engine for Chrome and ChromeOS. Essentially, there are two partitions, A and B. The machine runs on partition A until an update comes along, and it updates on partition B, and reboots to use partition B. This provides many advantages, the most important of which is atomic (and therefore safe) rollbacks.

18. Name three scenarios that virtual IP's are good for increasing redundancy.

Virtual IPs can provide redundancy in routers (and other networking gear), databases that have some form of replication (for providing a single IP for everyone to talk to), and redundant VPNs.

19. Name three ways that horizontal scaling can add complexity. Briefly explain why the complexity is necessary to scale horizontally.

Horizontal scaling can add complexity in the following ways:

- (a) **Load Balancing**
- (b) **Latency**
- (c) **Consensus**

Horizontal Scaling almost always requires some increased form of complexity because it often introduces new problems into the domain, such as keeping multiple nodes in the same state.

20. Suppose you have a system running with a single load balancer, three web nodes, and two database nodes. What is the single point of failure? How could you get rid of this single point of failure? After fixing it, what other points of failure might be singular?

The load balancer is the single point of failure. You could fix it by introducing a second load balancer backed by a virtual IP for fail-over. After fixing the load balancer, other single points of failure could include things like: power; networking; cooling.