

Exam Questions CCA-500

Cloudera Certified Administrator for Apache Hadoop (CCA-H)

<https://www.2passeasy.com/dumps/CCA-500/>



NEW QUESTION 1

You are migrating a cluster from MAppReduce version 1 (MRv1) to MapReduce version 2 (MRv2) on YARN. You want to maintain your MRv1 TaskTracker slot capacities when you migrate. What should you do/

- A. Configure yarn.applicationmaster.resource.memory-mb and yarn.applicationmaster.resource.cpu-vcores so that ApplicationMaster container allocations match the capacity you require.
- B. You don't need to configure or balance these properties in YARN as YARN dynamically balances resource management capabilities on your cluster
- C. Configure mapred.tasktracker.map.tasks.maximum and mapred.tasktracker.reduce.tasks.maximum in yarn-site.xml to match your cluster's capacity set by the yarn-scheduler.minimum-allocation
- D. Configure yarn.nodemanager.resource.memory-mb and yarn.nodemanager.resource.cpu-vcores to match the capacity you require under YARN for each NodeManager

Answer: D

NEW QUESTION 2

Given:

```
[user1@host1 ~] yarn application -list
```

Total Applications: 3

Application ID	Application-Name	Application-Type	User	Queue	State	Final-State	Progress	Tracking
Application_1374638600275_0109	Sleep Job	MAPREDUCE	user1	KILLED	KILLED	KILLED	100%	host1:54059
Application_1374638600275_0121	Sleep Job	MAPREDUCE	user1	FINISHED	SUCCEEDED	SUCCEEDED	100%	host1:19888/Jobhistory/Job_1374638600275_0121
Application_1374638600275_0020	Sleep Job	MAPREDUCE	user1	FINISHED	SUCCEEDED	SUCCEEDED	100%	host1:19888/Jobhistory/Job_1374638600275_0020

You want to clean up this list by removing jobs where the State is KILLED. What command you enter?

- A. Yarn application --refreshJobHistory
- B. Yarn application --kill application_1374638600275_0109
- C. Yarn rmadmin --refreshQueue
- D. Yarn rmadmin --kill application_1374638600275_0109

Answer: B

Explanation: Reference:http://docs.hortonworks.com/HDPDocuments/HDP2/HDP-2.1-latest/bk_using-apache-hadoop/content/common_mrv2_commands.html

NEW QUESTION 3

You are running a Hadoop cluster with MapReduce version 2 (MRv2) on YARN. You consistently see that MapReduce map tasks on your cluster are running slowly because of excessive garbage collection of JVM, how do you increase JVM heap size property to 3GB to optimize performance?

- A. yarn.application.child.java.opts=-Xsx3072m
- B. yarn.application.child.java.opts=-Xmx3072m
- C. mapreduce.map.java.opts=-Xms3072m
- D. mapreduce.map.java.opts=-Xmx3072m

Answer: C

Explanation: Reference:<http://hortonworks.com/blog/how-to-plan-and-configure-yarn-in-hdp-2-0/>

NEW QUESTION 4

You observed that the number of spilled records from Map tasks far exceeds the number of map output records. Your child heap size is 1GB and your io.sort.mb value is set to 1000MB. How would you tune your io.sort.mb value to achieve maximum memory to disk I/O ratio?

- A. For a 1GB child heap size an io.sort.mb of 128 MB will always maximize memory to disk I/O
- B. Increase the io.sort.mb to 1GB
- C. Decrease the io.sort.mb value to 0
- D. Tune the io.sort.mb value until you observe that the number of spilled records equals (or is as close to equals) the number of map output records.

Answer: D

NEW QUESTION 5

You are planning a Hadoop cluster and considering implementing 10 Gigabit Ethernet as the network fabric. Which workloads benefit the most from faster network fabric?

- A. When your workload generates a large amount of output data, significantly larger than the amount of intermediate data
- B. When your workload consumes a large amount of input data, relative to the entire capacity of HDFS
- C. When your workload consists of processor-intensive tasks
- D. When your workload generates a large amount of intermediate data, on the order of the input data itself

Answer: A

NEW QUESTION 6

You have just run a MapReduce job to filter user messages to only those of a selected geographical region. The output for this job is in a directory named westUsers, located just below your home directory in HDFS. Which command gathers these into a single file on your local file system?

- A. Hadoop fs --getmerge --R westUsers.txt

- B. Hadoop fs –getmerge westUsers westUsers.txt
- C. Hadoop fs –cp westUsers/* westUsers.txt
- D. Hadoop fs –get westUsers westUsers.txt

Answer: B

NEW QUESTION 7

You are running a Hadoop cluster with a NameNode on host mynamenode, a secondary NameNode on host mysecondarynamenode and several DataNodes. Which best describes how you determine when the last checkpoint happened?

- A. Execute hdfs namenode –report on the command line and look at the Last Checkpoint information
- B. Execute hdfs dfsadmin –saveNamespace on the command line which returns to you the last checkpoint value in fstime file
- C. Connect to the web UI of the Secondary NameNode (<http://mysecondary:50090/>) and look at the “Last Checkpoint” information
- D. Connect to the web UI of the NameNode (<http://mynamenode:50070/>) and look at the “Last Checkpoint” information

Answer: C

Explanation: Reference:<https://www.inkling.com/read/hadoop-definitive-guide-tom-white-3rd/chapter-10/hdfs>

NEW QUESTION 8

On a cluster running MapReduce v2 (MRv2) on YARN, a MapReduce job is given a directory of 10 plain text files as its input directory. Each file is made up of 3 HDFS blocks. How many Mappers will run?

- A. We cannot say; the number of Mappers is determined by the ResourceManager
- B. We cannot say; the number of Mappers is determined by the developer
- C. 30
- D. 3
- E. 10
- F. We cannot say; the number of mappers is determined by the ApplicationMaster

Answer: E

NEW QUESTION 9

Your Hadoop cluster contains nodes in three racks. You have not configured the dfs.hosts property in the NameNode’s configuration file. What results?

- A. The NameNode will update the dfs.hosts property to include machines running the DataNode daemon on the next NameNode reboot or with the command dfsadmin–refreshNodes
- B. No new nodes can be added to the cluster until you specify them in the dfs.hosts file
- C. Any machine running the DataNode daemon can immediately join the cluster
- D. Presented with a blank dfs.hosts property, the NameNode will permit DataNodes specified in mapred.hosts to join the cluster

Answer: C

NEW QUESTION 10

Which two features does Kerberos security add to a Hadoop cluster?(Choose two)

- A. User authentication on all remote procedure calls (RPCs)
- B. Encryption for data during transfer between the Mappers and Reducers
- C. Encryption for data on disk (“at rest”)
- D. Authentication for user access to the cluster against a central server
- E. Root access to the cluster for users hdfs and mapred but non-root access for clients

Answer: AD

NEW QUESTION 10

Which YARN daemon or service monitors a Controller’s per-application resource using (e.g., memory CPU)?

- A. ApplicationMaster
- B. NodeManager
- C. ApplicationManagerService
- D. ResourceManager

Answer: A

NEW QUESTION 11

Identify two features/issues that YARN is designated to address:(Choose two)

- A. Standardize on a single MapReduce API
- B. Single point of failure in the NameNode
- C. Reduce complexity of the MapReduce APIs
- D. Resource pressure on the JobTracker
- E. Ability to run framework other than MapReduce, such as MPI
- F. HDFS latency

Answer: DE

Explanation: Reference:[http://www.revelytix.com/?q=content/hadoop-ecosystem\(YARN, first para\)](http://www.revelytix.com/?q=content/hadoop-ecosystem(YARN, first para))

NEW QUESTION 16

Assume you have a file named foo.txt in your local directory. You issue the following three commands:

Hadoop fs -mkdir input

Hadoop fs -put foo.txt input/foo.txt

Hadoop fs -put foo.txt input

What happens when you issue the third command?

- A. The write succeeds, overwriting foo.txt in HDFS with no warning
- B. The file is uploaded and stored as a plain file named input
- C. You get a warning that foo.txt is being overwritten
- D. You get an error message telling you that foo.txt already exists, and asking you if you would like to overwrite it.
- E. You get a error message telling you that foo.txt already exist
- F. The file is not written to HDFS
- G. You get an error message telling you that input is not a directory
- H. The write silently fails

Answer: CE

NEW QUESTION 17

You are working on a project where you need to chain together MapReduce, Pig jobs. You also need the ability to use forks, decision points, and path joins. Which ecosystem project should you use to perform these actions?

- A. Oozie
- B. ZooKeeper
- C. HBase
- D. Sqoop
- E. HUE

Answer: A

NEW QUESTION 21

You decide to create a cluster which runs HDFS in High Availability mode with automatic failover, using Quorum Storage. What is the purpose of ZooKeeper in such a configuration?

- A. It only keeps track of which NameNode is Active at any given time
- B. It monitors an NFS mount point and reports if the mount point disappears
- C. It both keeps track of which NameNode is Active at any given time, and manages the Edits fil
- D. Which is a log of changes to the HDFS filesystem
- E. If only manages the Edits file, which is log of changes to the HDFS filesystem
- F. Clients connect to ZooKeeper to determine which NameNode is Active

Answer: A

Explanation: Reference: Reference:[http://www.cloudera.com/content/cloudera-content/cloudera-docs/CDH4/latest/PDF/CDH4-High-Availability-Guide.pdf\(page 15\)](http://www.cloudera.com/content/cloudera-content/cloudera-docs/CDH4/latest/PDF/CDH4-High-Availability-Guide.pdf(page 15))

NEW QUESTION 24

Which is the default scheduler in YARN?

- A. YARN doesn't configure a default scheduler, you must first assign an appropriate scheduler class in yarn-site.xml
- B. Capacity Scheduler
- C. Fair Scheduler
- D. FIFO Scheduler

Answer: B

Explanation: Reference:<http://hadoop.apache.org/docs/r2.4.1/hadoop-yarn/hadoop-yarn-site/CapacityScheduler.html>

NEW QUESTION 27

On a cluster running CDH 5.0 or above, you use the `hadoop fs -put` command to write a 300MB file into a previously empty directory using an HDFS block size of 64 MB. Just after this command has finished writing 200 MB of this file, what would another use see when they look in directory?

- A. The directory will appear to be empty until the entire file write is completed on the cluster
- B. They will see the file with a `._COPYING_` extension on its nam
- C. If they view the file, they will see contents of the file up to the last completed block (as each 64MB block is written, that block becomes available)
- D. They will see the file with a `._COPYING_` extension on its nam
- E. If they attempt to view the file, they will get a `ConcurrentFileAccessException` until the entire file write is completed on the cluster
- F. They will see the file with its original nam
- G. If they attempt to view the file, they will get a `ConcurrentFileAccessException` until the entire file write is completed on the cluster

Answer: B

NEW QUESTION 28

Which scheduler would you deploy to ensure that your cluster allows short jobs to finish within a reasonable time without starting long-running jobs?

- A. Complexity Fair Scheduler (CFS)
- B. Capacity Scheduler
- C. Fair Scheduler
- D. FIFO Scheduler

Answer: C

Explanation: Reference:http://hadoop.apache.org/docs/r1.2.1/fair_scheduler.html

NEW QUESTION 32

For each YARN job, the Hadoop framework generates task log file. Where are Hadoop task log files stored?

- A. Cached by the NodeManager managing the job containers, then written to a log directory on the NameNode
- B. Cached in the YARN container running the task, then copied into HDFS on job completion
- C. In HDFS, in the directory of the user who generates the job
- D. On the local disk of the slave node running the task

Answer: D

NEW QUESTION 37

What two processes must you do if you are running a Hadoop cluster with a single NameNode and six DataNodes, and you want to change a configuration parameter so that it affects all six DataNodes.(Choose two)

- A. You must modify the configuration files on the NameNode only
- B. DataNodes read their configuration from the master nodes
- C. You must modify the configuration files on each of the six DataNodes machines
- D. You don't need to restart any daemon, as they will pick up changes automatically
- E. You must restart the NameNode daemon to apply the changes to the cluster
- F. You must restart all six DataNode daemon to apply the changes to the cluster

Answer: BD

NEW QUESTION 38

Your cluster's mapred-start.xml includes the following parameters

```
<name>mapreduce.map.memory.mb</name>
```

```
<value>4096</value>
```

```
<name>mapreduce.reduce.memory.mb</name>
```

```
<value>8192</value>
```

And any cluster's yarn-site.xml includes the following parameters

```
<name>yarn.nodemanager.vmem-pmem-ratio</name>
```

```
<value>2.1</value>
```

What is the maximum amount of virtual memory allocated for each map task before YARN will kill its Container?

- A. 4 GB
- B. 17.2 GB
- C. 8.9 GB
- D. 8.2 GB
- E. 24.6 GB

Answer: D

NEW QUESTION 43

You have recently converted your Hadoop cluster from a MapReduce 1 (MRv1) architecture to MapReduce 2 (MRv2) on YARN architecture. Your developers are accustomed to specifying map and reduce tasks (resource allocation) tasks when they run jobs: A developer wants to know how specify to reduce tasks when a specific job runs. Which method should you tell that developers to implement?

- A. MapReduce version 2 (MRv2) on YARN abstracts resource allocation away from the idea of "tasks" into memory and virtual cores, thus eliminating the need for a developer to specify the number of reduce tasks, and indeed preventing the developer from specifying the number of reduce tasks.
- B. In YARN, resource allocations is a function of megabytes of memory in multiples of 1024m
- C. Thus, they should specify the amount of memory resource they need by executing `-D mapreduce-reduces.memory-mb-2048`
- D. In YARN, the ApplicationMaster is responsible for requesting the resource required for a specific launch
- E. Thus, executing `-D yarn.applicationmaster.reduce.tasks=2` will specify that the ApplicationMaster launch two task containers on the worker nodes.
- F. Developers specify reduce tasks in the exact same way for both MapReduce version 1 (MRv1) and MapReduce version 2 (MRv2) on YARN
- G. Thus, executing `-D mapreduce.job.reduces=2` will specify reduce tasks.
- H. In YARN, resource allocation is function of virtual cores specified by the ApplicationMaster making requests to the NodeManager where a reduce task is handled by a single container (and thus a single virtual core). Thus, the developer needs to specify the number of virtual cores to the NodeManager by executing `-p yarn.nodemanager.cpu-vcores=2`

Answer: D

NEW QUESTION 48

Your cluster has the following characteristics:

? A rack aware topology is configured and on

? Replication is set to 3

? Cluster block size is set to 64MB

Which describes the file read process when a client application connects into the cluster and requests a 50MB file?

- A. The client queries the NameNode for the locations of the block, and reads all three copie
- B. The first copy to complete transfer to the client is the one the client reads as part of hadoop's speculative execution framework.
- C. The client queries the NameNode for the locations of the block, and reads from the first location in the list it receives.
- D. The client queries the NameNode for the locations of the block, and reads from a random location in the list it receives to eliminate network I/O loads by balancing which nodes it retrieves data from any given time.
- E. The client queries the NameNode which retrieves the block from the nearest DataNode to the client then passes that block back to the client.

Answer: B

NEW QUESTION 49

You have a cluster running with a FIFO scheduler enabled. You submit a large job A to the cluster, which you expect to run for one hour. Then, you submit job B to the cluster, which you expect to run a couple of minutes only.

You submit both jobs with the same priority.

Which two best describes how FIFO Scheduler arbitrates the cluster resources for job and its tasks?(Choose two)

- A. Because there is a more than a single job on the cluster, the FIFO Scheduler will enforce a limit on the percentage of resources allocated to a particular job at any given time
- B. Tasks are scheduled on the order of their job submission
- C. The order of execution of job may vary
- D. Given job A and submitted in that order, all tasks from job A are guaranteed to finish before all tasks from job B
- E. The FIFO Scheduler will give, on average, and equal share of the cluster resources over the job lifecycle
- F. The FIFO Scheduler will pass an exception back to the client when Job B is submitted, since all slots on the cluster are use

Answer: AD

NEW QUESTION 50

Assuming you're not running HDFS Federation, what is the maximum number of NameNode daemons you should run on your cluster in order to avoid a "split-brain" scenario with your NameNode when running HDFS High Availability (HA) using Quorum- based storage?

- A. Two active NameNodes and two Standby NameNodes
- B. One active NameNode and one Standby NameNode
- C. Two active NameNodes and on Standby NameNode
- D. Unlimite
- E. HDFS High Availability (HA) is designed to overcome limitations on the number of NameNodes you can deploy

Answer: B

NEW QUESTION 55

Which command does Hadoop offer to discover missing or corrupt HDFS data?

- A. Hdfs fs -du
- B. Hdfs fsck
- C. Dskchk
- D. The map-only checksum
- E. Hadoop does not provide any tools to discover missing or corrupt data; there is not need because three replicas are kept for each data block

Answer: B

Explanation: Reference:<https://twiki.grid.iu.edu/bin/view/Storage/HadoopRecovery>

NEW QUESTION 60

Your cluster is configured with HDFS and MapReduce version 2 (MRv2) on YARN. What is the result when you execute: `hadoop jar SampleJar MyClass` on a client machine?

- A. SampleJar.Jar is sent to the ApplicationMaster which allocates a container for SampleJar.Jar
- B. Sample.jar is placed in a temporary directory in HDFS
- C. SampleJar.jar is sent directly to the ResourceManager
- D. SampleJar.jar is serialized into an XML file which is submitted to the ApplicatoionMaster

Answer: A

NEW QUESTION 62

You are running a Hadoop cluster with a NameNode on host mynamenode. What are two ways to determine available HDFS space in your cluster?

- A. Run `hdfs fs -du /` and locate the DFS Remaining value
- B. Run `hdfs dfsadmin -report` and locate the DFS Remaining value
- C. Run `hdfs dfs /` and subtract NDFS Used from configured Capacity
- D. Connect to `http://mynamenode:50070/dfshealth.jsp` and locate the DFS remaining value

Answer: B

NEW QUESTION 66

Your cluster is running MapReduce version 2 (MRv2) on YARN. Your ResourceManager is configured to use the FairScheduler. Now you want to configure your

scheduler such that a new user on the cluster can submit jobs into their own queue application submission. Which configuration should you set?

- A. You can specify new queue name when user submits a job and new queue can be created dynamically if the property `yarn.scheduler.fair.allow-undecleared-pools = true`
- B. `Yarn.scheduler.fair.user.fair-as-default-queue = false` and `yarn.scheduler.fair.allow-undecleared-pools = true`
- C. You can specify new queue name when user submits a job and new queue can be created dynamically if `yarn.scheduler.fair.user-as-default-queue = false`
- D. You can specify new queue name per application in `allocations.xml` file and have new jobs automatically assigned to the application queue

Answer: A

NEW QUESTION 71

Which YARN process run as “container 0” of a submitted job and is responsible for resource requests?

- A. ApplicationManager
- B. JobTracker
- C. ApplicationMaster
- D. JobHistoryServer
- E. ResoureManager
- F. NodeManager

Answer: C

NEW QUESTION 75

During the execution of a MapReduce v2 (MRv2) job on YARN, where does the Mapper place the intermediate data of each Map Task?

- A. The Mapper stores the intermediate data on the node running the Job’s ApplicationMaster so that it is available to YARN ShuffleService before the data is presented to the Reducer
- B. The Mapper stores the intermediate data in HDFS on the node where the Map tasks ran in the `HDFS /usercache/&(user)/apache/application_&(appid)` directory for the user who ran the job
- C. The Mapper transfers the intermediate data immediately to the reducers as it is generated by the Map Task
- D. YARN holds the intermediate data in the NodeManager’s memory (a container) until it is transferred to the Reducer
- E. The Mapper stores the intermediate data on the underlying filesystem of the local disk in the directories `yarn.nodemanager.local-DIFS`

Answer: E

NEW QUESTION 78

Choose three reasons why should you run the HDFS balancer periodically?(Choose three)

- A. To ensure that there is capacity in HDFS for additional data
- B. To ensure that all blocks in the cluster are 128MB in size
- C. To help HDFS deliver consistent performance under heavy loads
- D. To ensure that there is consistent disk utilization across the DataNodes
- E. To improve data locality MapReduce

Answer: CDE

Explanation: <http://www.quora.com/Apache-Hadoop/It-is-recommended-that-you-run-the-HDFS-balancer-periodically-Why-Choose-3>

NEW QUESTION 80

Your Hadoop cluster is configuring with HDFS and MapReduce version 2 (MRv2) on YARN. Can you configure a worker node to run a NodeManager daemon but not a DataNode daemon and still have a functional cluster?

- A. Ye
- B. The daemon will receive data from the NameNode to run Map tasks
- C. Ye
- D. The daemon will get data from another (non-local) DataNode to run Map tasks
- E. Ye
- F. The daemon will receive Map tasks only
- G. Ye
- H. The daemon will receive Reducer tasks only

Answer: B

NEW QUESTION 85

You want to node to only swap Hadoop daemon data from RAM to disk when absolutely necessary. What should you do?

- A. Delete the `/dev/vmswap` file on the node
- B. Delete the `/etc/swape` file on the node
- C. Set the `ram.swap` parameter to 0 in `core-site.xml`
- D. Set `vm.swapfile` file on the node
- E. Delete the `/swapfile` file on the node

Answer: D

NEW QUESTION 90

You need to analyze 60,000,000 images stored in JPEG format, each of which is approximately 25 KB. Because you Hadoop cluster isn’t optimized for storing

and processing many small files, you decide to do the following actions:

1. Group the individual images into a set of larger files
2. Use the set of larger files as input for a MapReduce job that processes them directly with python using Hadoop streaming.

Which data serialization system gives the flexibility to do this?

- A. CSV
- B. XML
- C. HTML
- D. Avro
- E. SequenceFiles
- F. JSON

Answer: E

Explanation: Sequence files are block-compressed and provide direct serialization and deserialization of several arbitrary data types (not just text). Sequence files can be generated as the output of other MapReduce tasks and are an efficient intermediate representation for data that is passing from one MapReduce job to another.

NEW QUESTION 93

A user comes to you, complaining that when she attempts to submit a Hadoop job, it fails. There is a Directory in HDFS named /data/input. The Jar is named j.jar, and the driver class is named DriverClass.

She runs the command:

Hadoop jar j.jar DriverClass /data/input/data/output The error message returned includes the line:

PrivilegedActionException as:training (auth:SIMPLE) cause:org.apache.hadoop.mapreduce.lib.input.InvalidInputException:

Input path does not exist: file:/data/input What is the cause of the error?

- A. The user is not authorized to run the job on the cluster
- B. The output directory already exists
- C. The name of the driver has been spelled incorrectly on the command line
- D. The directory name is misspelled in HDFS
- E. The Hadoop configuration files on the client do not point to the cluster

Answer: A

NEW QUESTION 94

Which YARN daemon or service negotiates map and reduce Containers from the Scheduler, tracking their status and monitoring progress?

- A. NodeManager
- B. ApplicationMaster
- C. ApplicationManager
- D. ResourceManager

Answer: B

Explanation: Reference:<http://www.devx.com/opensource/intro-to-apache-mapreduce-2-yarn.html>(See resource manager)

NEW QUESTION 96

You have installed a cluster HDFS and MapReduce version 2 (MRv2) on YARN. You have no dfs.hosts entry(ies) in your hdfs-site.xml configuration file. You configure a new worker node by setting fs.default.name in its configuration files to point to the NameNode on your cluster, and you start the DataNode daemon on that worker node. What do you have to do on the cluster to allow the worker node to join, and start sorting HDFS blocks?

- A. Without creating a dfs.hosts file or making any entries, run the commands `hadoop dfsadmin -refreshNodes` on the NameNode
- B. Restart the NameNode
- C. Creating a dfs.hosts file on the NameNode, add the worker Node's name to it, then issue the command `hadoop dfsadmin -refreshNodes` on the Namenode
- D. Nothing; the worker node will automatically join the cluster when NameNode daemon is started

Answer: A

NEW QUESTION 98

Which process instantiates user code, and executes map and reduce tasks on a cluster running MapReduce v2 (MRv2) on YARN?

- A. NodeManager
- B. ApplicationMaster
- C. TaskTracker
- D. JobTracker
- E. NameNode
- F. DataNode
- G. ResourceManager

Answer: A

NEW QUESTION 103

.....

THANKS FOR TRYING THE DEMO OF OUR PRODUCT

Visit Our Site to Purchase the Full Set of Actual CCA-500 Exam Questions With Answers.

We Also Provide Practice Exam Software That Simulates Real Exam Environment And Has Many Self-Assessment Features. Order the CCA-500 Product From:

<https://www.2passeasy.com/dumps/CCA-500/>

Money Back Guarantee

CCA-500 Practice Exam Features:

- * CCA-500 Questions and Answers Updated Frequently
- * CCA-500 Practice Questions Verified by Expert Senior Certified Staff
- * CCA-500 Most Realistic Questions that Guarantee you a Pass on Your FirstTry
- * CCA-500 Practice Test Questions in Multiple Choice Formats and Updatesfor 1 Year