

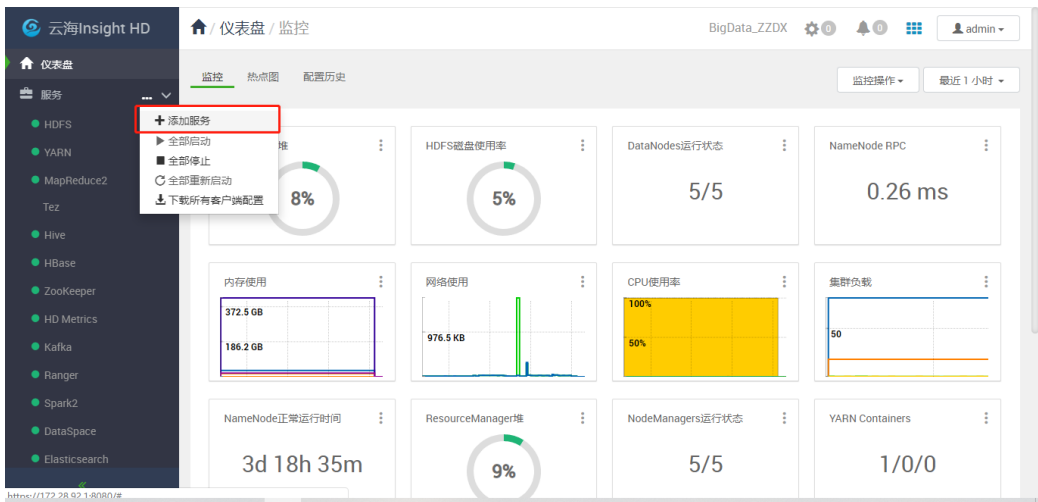
郑州大学电气工程学院公共安全大数据与人工智能平台设备采购项目 Insight HD集群功能测试报告

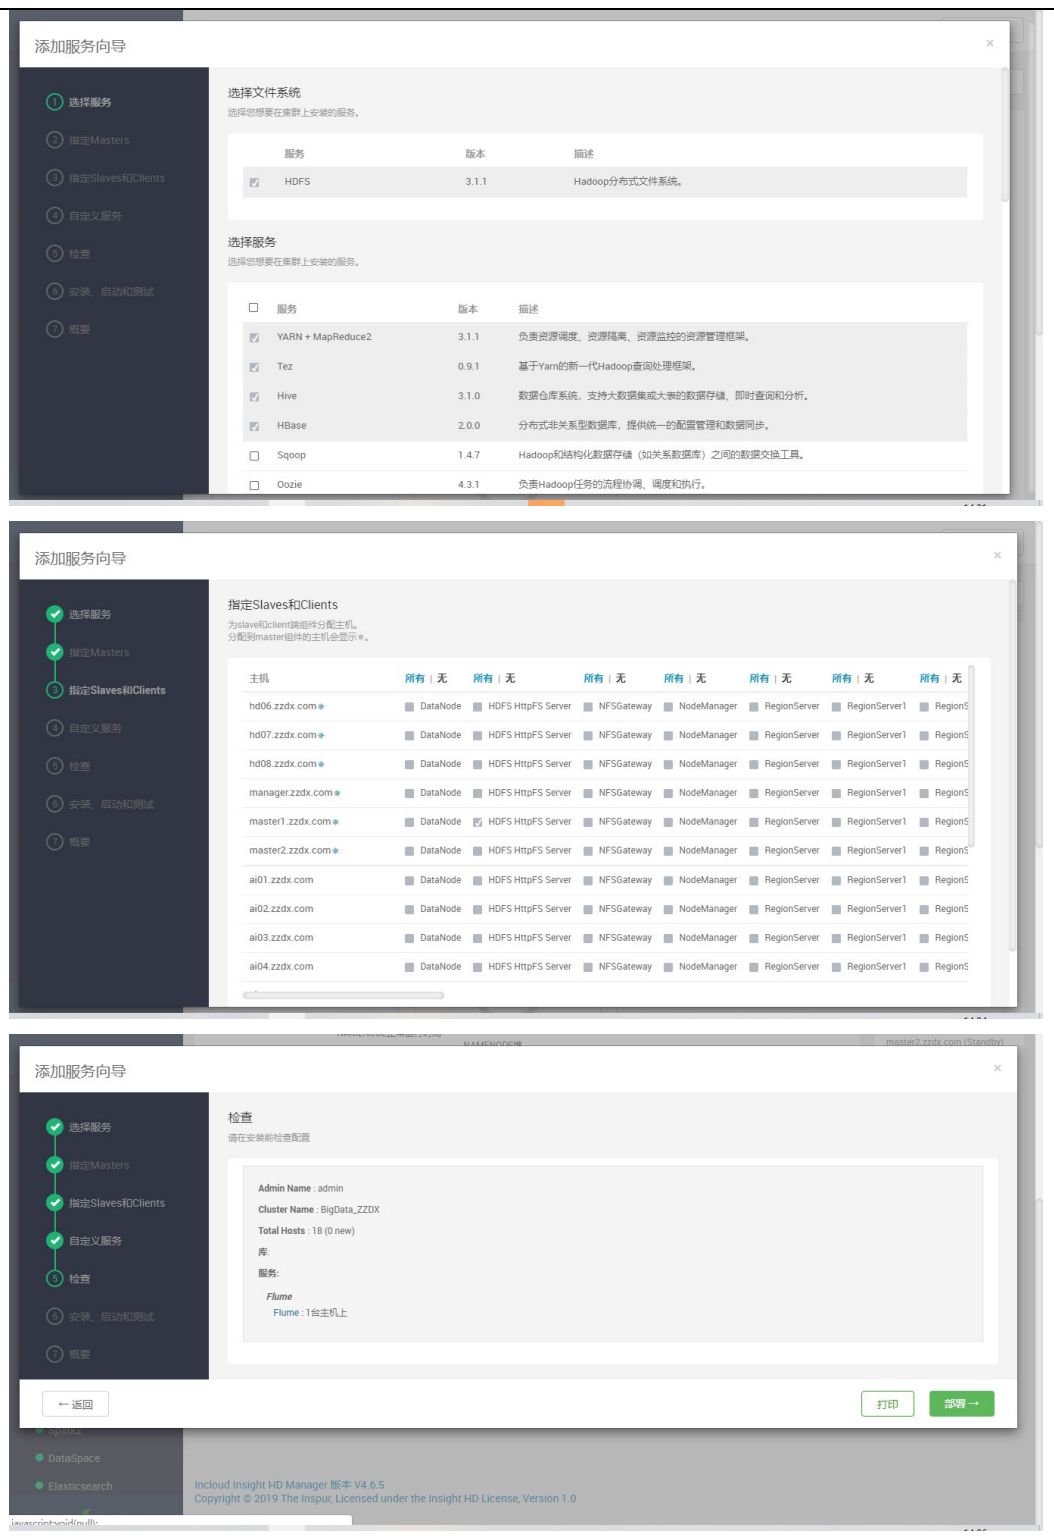
2021 年 5 月

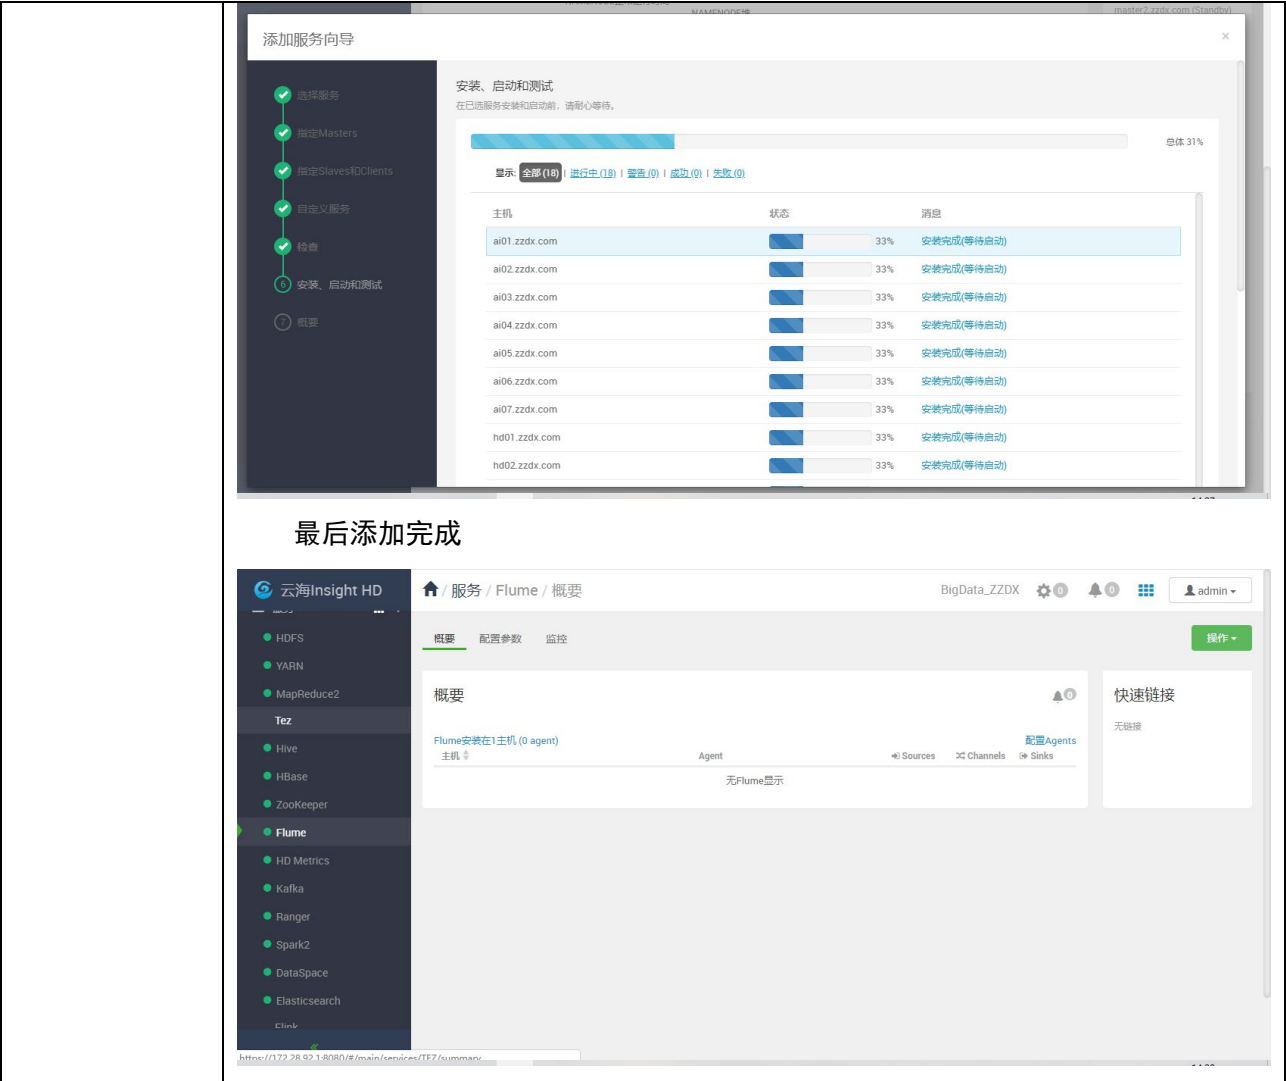
1. 测试内容

1.1 云海 Insight 运营功能测试

1.1.1 自动化部署

测试内容	自动化部署
预期结果	1、支持所有组件的自动化部署 2、安装部署操作以向导方式进行 3、是否显示安装进度 4、考察安装效率
测试结果	符合预期，测试通过 <input checked="" type="checkbox"/> 不符合预期，测试不通过 <input type="checkbox"/> 不具备测试条件，技术论证 <input type="checkbox"/>
测试说明	自动化部署（以 Flume 举例）  The screenshot shows the CloudInsight HD dashboard. On the left, there is a sidebar with a list of services: HDFS, YARN, MapReduce2, Tez, Hive, HBase, ZooKeeper, HD Metrics, Kafka, Ranger, Spark2, DataSpace, and Elasticsearch. A red box highlights the '+ 添加服务' (Add Service) button in the sidebar. A dropdown menu is open, showing options: '全部启动' (Start All), '全部停止' (Stop All), '全部重新启动' (Restart All), and '下载所有客户端配置' (Download all client configurations). The main dashboard area displays various monitoring widgets, including 'HDFS磁盘使用率' (HDFS disk usage) at 8%, 'HDFS磁盘使用率' (HDFS disk usage) at 5%, 'DataNodes运行状态' (DataNodes running status) at 5/5, 'NameNode RPC' at 0.26 ms, '内存使用' (Memory usage) showing 372.5 GB and 186.2 GB, '网络使用' (Network usage) at 976.5 KB, 'CPU使用率' (CPU usage) at 100% and 50%, '集群负载' (Cluster load) at 50, 'NameNode正常运行时间' (NameNode uptime) at 3d 18h 35m, 'ResourceManager堆' (ResourceManager heap) at 9%, 'NodeManagers运行状态' (NodeManagers running status) at 5/5, and 'YARN Containers' at 1/0/0.

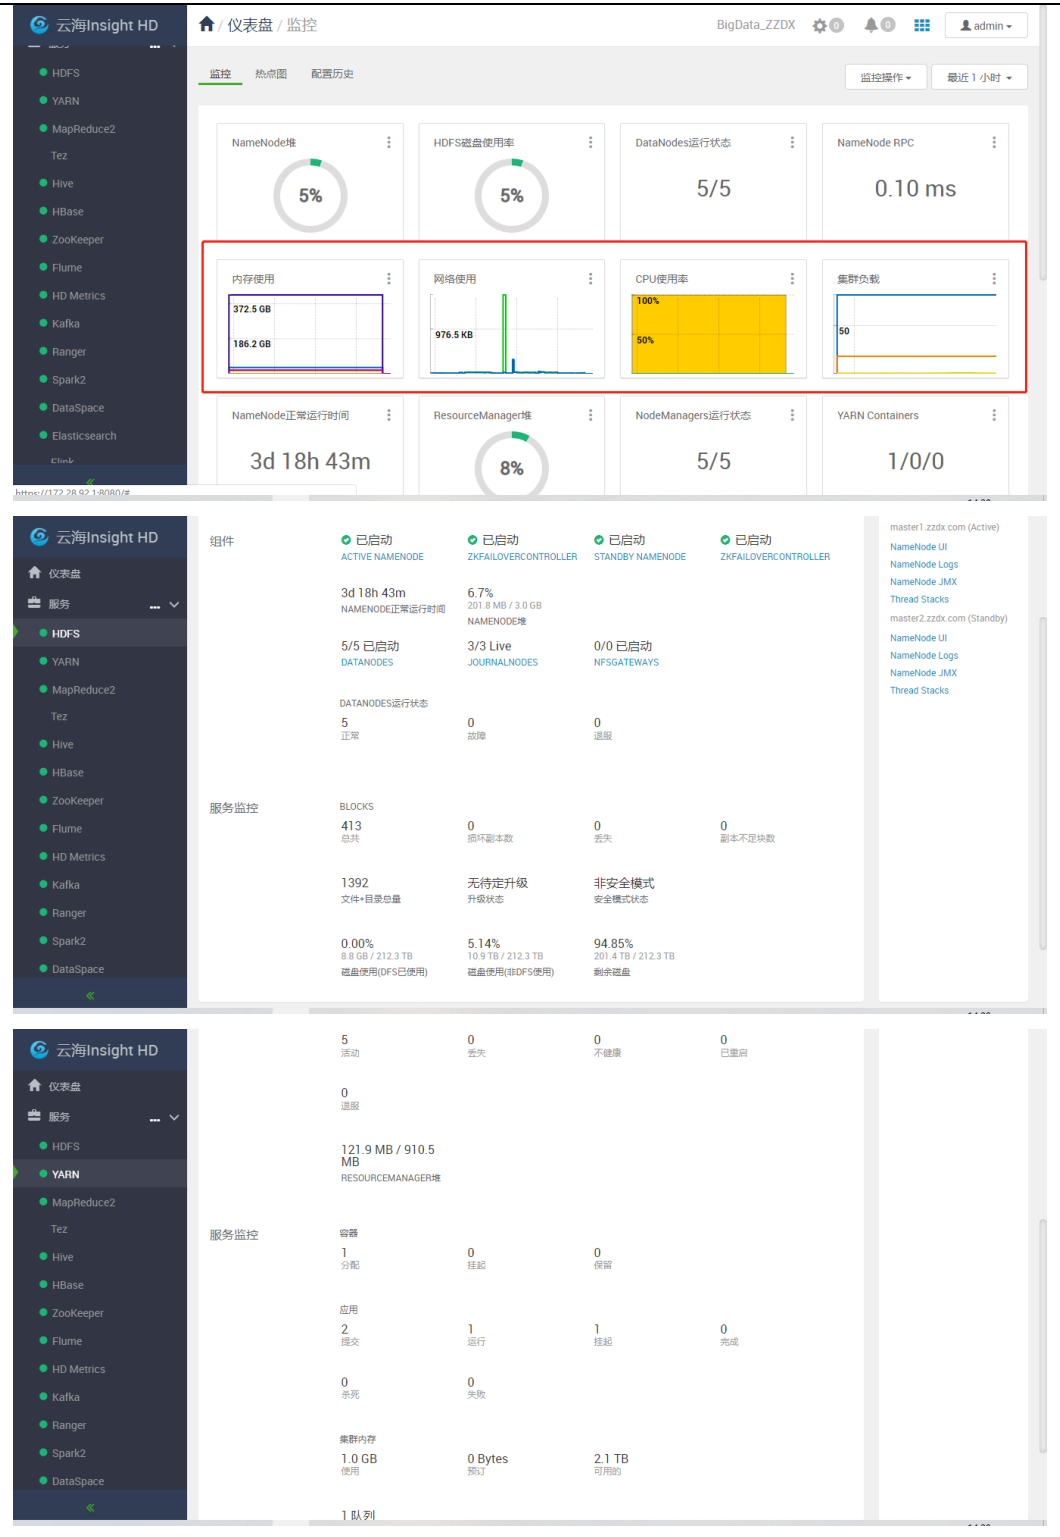




1.1.2 资源监控

测试内容	资源监控
预期结果	验证 Hadoop 管理系统对 Hadoop 系统性能进行图形化监控管理的功能 1、支持查看集群和单机的 CPU、内存、存储、网络负载 2、HDFS: HDFS 文件系统块总数、总大小、文件总数、剩余量、损坏块等 3、清单查询集群: RegionServer 的请求次数以及 RegionServer Regions 数目等 4、YARN: Current Applications、Finished Applications、NodeManagers
测试结果	符合预期，测试通过 <input checked="" type="checkbox"/> 不符合预期，测试不通过 <input type="checkbox"/> 不具备测试条件，技术论证 <input type="checkbox"/>

测试说明



1.1.3 集群操作

测试内容	集群操作
预期结果	1、通过管理平台界面一次性启停集群所有服务 2、在集群内通过管理了平台界面启停某些服务

测试结果	<div>符合预期，测试通过 <input checked="" type="checkbox"/></div> <div>不符合预期，测试不通过 <input type="checkbox"/></div> <div>不具备测试条件，技术论证 <input type="checkbox"/></div>
测试说明	<div></div> <div></div> <div></div>

1.1.4 日志管理

测试内容	测试大数据平台的日志管理功能
预期结果	集群管理平台能够收集日志，并提供日志预览和下载等功能
测试结果	符合预期，测试通过 <input checked="" type="checkbox"/> 不符合预期，测试不通过 <input type="checkbox"/> 不具备测试条件，技术论证 <input type="checkbox"/>
测试说明	<div><div><div><div>云海Insight HD</div><div>服务 / HDFS / 概要</div><div>BigData_ZZDX</div><div>admin</div></div><div><div>仪表盘</div><div>服务</div><div>HDFS</div><div>YARN</div><div>MapReduce2</div><div>Tez</div><div>Hive</div><div>HBase</div><div>ZooKeeper</div><div>Flume</div><div>HD Metrics</div><div>Kafka</div><div>Ranger</div><div>Spark2</div><div>DataSpace</div></div><div><div>概要</div><div>热点图</div><div>配置参数</div><div>监控</div></div><div><div>操作</div></div></div><div><div>组件</div><div>已启动</div><div>已启动</div><div>已启动</div><div>已启动</div><div>3d 18h 45m</div><div>10.6%</div><div>5/5 已启动</div><div>3/3 Live</div><div>0/0 已启动</div><div>DATANODES运行状态</div><div>5 正常</div><div>0 故障</div><div>0 退服</div><div>服务监控</div><div>BLOCKS</div><div>413 总共</div><div>0 损坏副本数</div><div>0 丢失</div><div>0 副本不足块数</div></div><div><div>快速链接</div><div>master1.zzdx.com (Active)</div><div>NameNode UI</div><div>NameNode Logs</div><div>NameNode JMX</div><div>Thread Stacks</div><div>master2.zzdx.com (Standby)</div><div>NameNode UI</div><div>NameNode Logs</div><div>NameNode JMX</div><div>Thread Stacks</div></div></div> <div><div>Directory: /logs/</div><div><div>SecurityAuth.audit</div><div>0 bytes</div><div>May 1, 2021 3:41:36 PM</div></div><div><div>gc.log-202105011541</div><div>2380 bytes</div><div>May 1, 2021 3:41:37 PM</div></div><div><div>gc.log-202105011543</div><div>1507 bytes</div><div>May 1, 2021 3:44:02 PM</div></div><div><div>gc.log-202105011544</div><div>6677 bytes</div><div>May 1, 2021 3:56:48 PM</div></div><div><div>gc.log-202105011558</div><div>2388 bytes</div><div>May 1, 2021 3:58:59 PM</div></div><div><div>gc.log-202105011559</div><div>4182 bytes</div><div>May 1, 2021 4:01:21 PM</div></div><div><div>gc.log-202105011601</div><div>6218 bytes</div><div>May 1, 2021 4:10:21 PM</div></div><div><div>gc.log-202105011610</div><div>4916 bytes</div><div>May 1, 2021 4:31:26 PM</div></div><div><div>gc.log-202105011631</div><div>5159 bytes</div><div>May 1, 2021 4:53:55 PM</div></div><div><div>gc.log-202105011653</div><div>5406 bytes</div><div>May 1, 2021 5:42:18 PM</div></div><div><div>gc.log-202105011742</div><div>4674 bytes</div><div>May 1, 2021 5:56:59 PM</div></div><div><div>gc.log-202105011819</div><div>40065 bytes</div><div>May 2, 2021 3:45:42 PM</div></div><div><div>gc.log-202105021545</div><div>11277 bytes</div><div>May 2, 2021 7:45:21 PM</div></div><div><div>gc.log-202105021945</div><div>162590 bytes</div><div>May 6, 2021 2:27:14 PM</div></div><div><div>hadoop-hdfs-journalnode-master1.zzdx.com.log</div><div>2974451 bytes</div><div>May 6, 2021 2:31:14 PM</div></div><div><div>hadoop-hdfs-journalnode-master1.zzdx.com.out</div><div>704 bytes</div><div>May 2, 2021 7:45:24 PM</div></div><div><div>hadoop-hdfs-journalnode-master1.zzdx.com.out.1</div><div>704 bytes</div><div>May 2, 2021 3:45:48 PM</div></div><div><div>hadoop-hdfs-journalnode-master1.zzdx.com.out.2</div><div>1122 bytes</div><div>May 1, 2021 6:19:36 PM</div></div><div><div>hadoop-hdfs-journalnode-master1.zzdx.com.out.3</div><div>704 bytes</div><div>May 1, 2021 5:42:23 PM</div></div><div><div>hadoop-hdfs-journalnode-master1.zzdx.com.out.4</div><div>704 bytes</div><div>May 1, 2021 4:53:23 PM</div></div><div><div>hadoop-hdfs-journalnode-master1.zzdx.com.out.5</div><div>704 bytes</div><div>May 1, 2021 4:31:19 PM</div></div></div> <div>预览和下载</div>

java HotSpot(TM) 64-Bit Server VM (25.121-b13) for linux-amd64_JRE (1.8.0_121-b13), built on Dec 12 2016 16:36:53 by "java_re" with gcc 4.3.0 20080428 (Red Hat 4.3.0-9)
Memory: 4k page, physical 527822268k (512197304k free), swap 16777212k (16777212k free)
Commandline flags: -XX:CMSInitiatingOccupancyFraction=70 -XX:ErrorFile=/var/log/hadoop/hdfs_err_pid%p.log -XX:InitialHeapSize=3221225472 -XX:MaxHeapSize=3221225472 -XX:MaxNewSize=402653184 -XX:MaxTenuringThreshold=4 -XX:NewSize=402653184 -XX:OldSpaceSize=16 -XX:ParallelGCThreads=8 -XX:ParallelGCThreads=8 -XX:ParallelGCThreads=8 -XX:PrintGC -XX:PrintGCDateStamps -XX:PrintGCDetails -XX:PrintGCClassStamps -XX:UseCMSInitiatingOccupancyOnly -XX:+UseCompressedClassPointers -XX:+UseCompressedOops -XX:+UseConcMarkSweepGC -XX:+UseFastMalloc
Heap
par new generation total 35392K, used 28092K [0x0000000700000000, 0x0000000718000000, 0x0000000718000000)
eden space 31462K, 82% used [0x0000000700000000, 0x000000070eece000, 0x0000000713400000)
from space 3929K, 0% used [0x0000000713400000, 0x0000000713400000, 0x00000007159a0000)
to space 3929K, 0% used [0x00000007159a0000, 0x00000007159a0000, 0x0000000718000000)
concurrent mark-sweep generation total 2762512K, used 0K [0x0000000718000000, 0x00000007c0000000, 0x00000007c0000000)
Metaspace used 15966K, capacity 18900K, committed 15900K, reserved 106700K
class space used 206K, capacity 218K, committed 2304K, reserved 1048678K

或者

```
[root@manager /]# cd /var/log/
[root@manager log]# ls
ambari-agent          ambari-metrics-monitor  dataspaces  hadoop-noproduce  hive          httpsflag.txt  insight_ntp.log  lastlog  ntpstats  spark2  yum.log
ambari-metrics-collector  ambari-server           flume       hadoop-yarn       hive-hcatalog  insight_http.log  insight_setup.log  livy2    ranger    sys-kernel  zookeeper
ambari-metrics-grafana  ansible.log             hadoop      hbase            http          insight_install.log  kafka
[root@manager log]# cd hadoop-yarn/
[root@manager hadoop-yarn]# ls
nodemanager  yarn
[root@manager hadoop-yarn]# cd yarn/
[root@manager yarn]# ls
hadoop-yarn-root-regisrtydns-manager.zdx.com.log  privileged-root-regisrtydns-manager.zdx.com.err  privileged-root-regisrtydns-manager.zdx.com.out.1
hadoop-yarn-root-regisrtydns-manager.zdx.com.out  privileged-root-regisrtydns-manager.zdx.com.err.1  privileged-root-regisrtydns-manager.zdx.com.out.2
hadoop-yarn-root-regisrtydns-manager.zdx.com.out.1  privileged-root-regisrtydns-manager.zdx.com.err.2  privileged-root-regisrtydns-manager.zdx.com.out.3
hadoop-yarn-root-regisrtydns-manager.zdx.com.out.2  privileged-root-regisrtydns-manager.zdx.com.err.3  privileged-root-regisrtydns-manager.zdx.com.out.4
hadoop-yarn-root-regisrtydns-manager.zdx.com.out.3  privileged-root-regisrtydns-manager.zdx.com.err.4  privileged-root-regisrtydns-manager.zdx.com.out.5
hadoop-yarn-root-regisrtydns-manager.zdx.com.out.4  privileged-root-regisrtydns-manager.zdx.com.err.5
hadoop-yarn-root-regisrtydns-manager.zdx.com.out.5  privileged-root-regisrtydns-manager.zdx.com.out
[root@manager yarn]#
```

1.1.5 用户管理

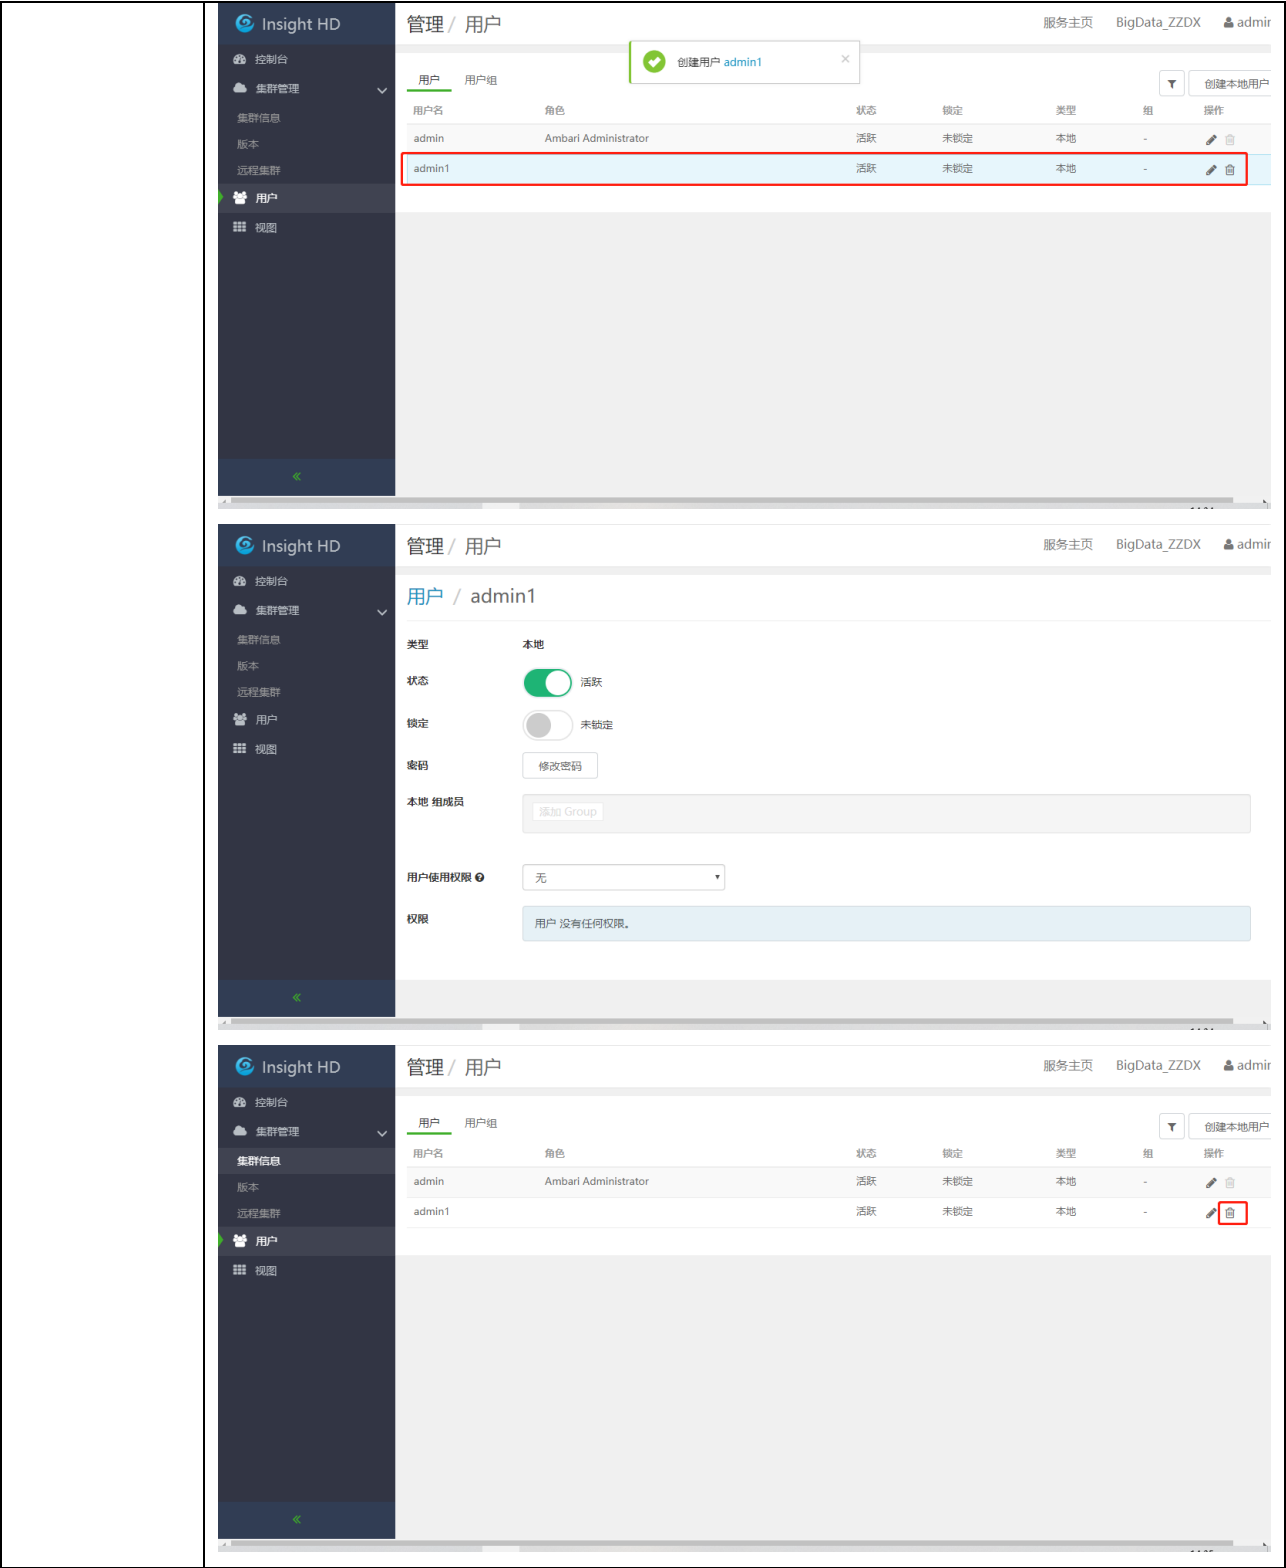
测试内容	用户管理（管理工具使用者）
预期结果	1、可以增加、修改、删除用户； 2、可以增加、修改、删除用户组，不同用户组可以设置不同的权限范围； 3、可以将不同用户指配到不同级别的用户组中，不同级别的用户所具有的权限范围不同； 4、支持修改账户的密码；
测试结果	符合预期，测试通过 <input checked="" type="checkbox"/> 不符合预期，测试不通过 <input type="checkbox"/> 不具备测试条件，技术论证 <input type="checkbox"/>

测试说明

The first screenshot shows the 'Overview' page for the HDFS service. The left sidebar lists various services, with HDFS selected. The main content area displays the status of HDFS components: ACTIVE NAMENODE, ZKFAILOVERCONTROLLER, STANDBY NAMENODE, and ZKFAILOVERCONTROLLER, all marked as '已启动' (Started). It also shows metrics for DATANODES, JOURNALNODES, and NFSGATEWAYS. The bottom section shows service monitoring metrics like BLOCKS, 副本数 (Number of Replicas), 丢失 (Lost), and 副本不足块数 (Number of Blocks with Insufficient Replicas).

The second screenshot shows the 'User Management' page. The left sidebar has '用户' (User) selected. The main content area shows a table of users, including 'admin' with the role 'Ambari Administrator'. A red box highlights the '创建本地用户' (Create Local User) button in the top right corner.

The third screenshot shows the 'Create Local User' dialog box. It contains fields for '用户名' (Username) with the value 'admin1', '密码' (Password), and '密码确认' (Confirm Password). A dropdown menu for '用户使用权限' (User Permissions) is open, showing options: '无' (None), '无' (None), 'Cluster User' (selected), 'Cluster Administrator', 'Cluster Operator', 'Service Administrator', and 'Service Operator'. '取消' (Cancel) and '保存' (Save) buttons are at the bottom right.



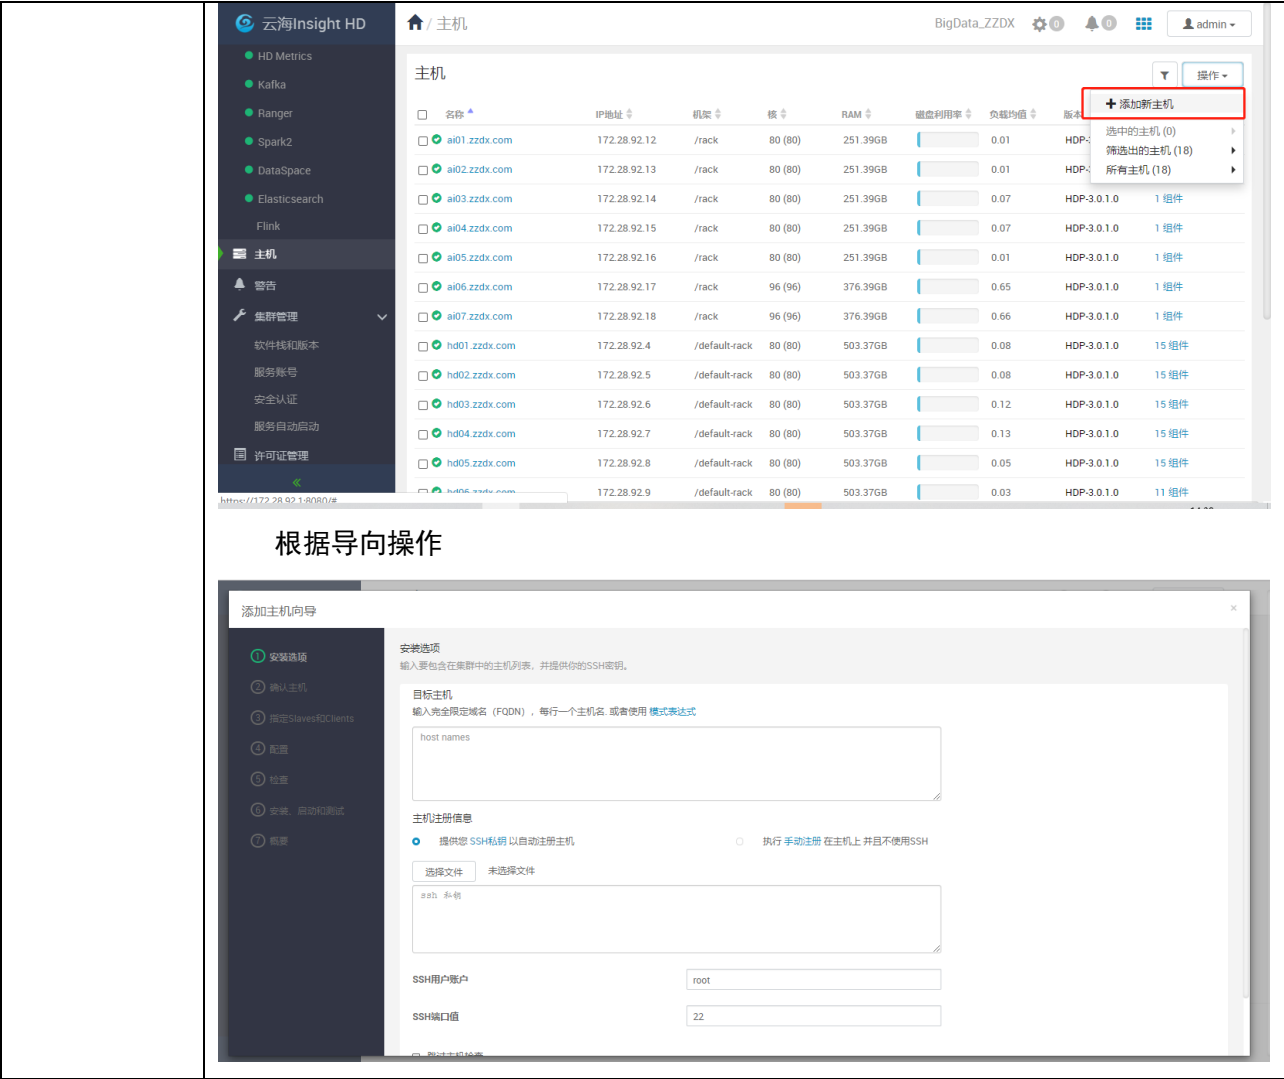
1.1.6 故障管理

测试内容	故障查看和管理
预期结果	1、告警级别（识别不同严重级别）

	2、告警阈值可配置 3、支持历史警报查询
测试结果	符合预期，测试通过 <input checked="" type="checkbox"/> 不符合预期，测试不通过 <input type="checkbox"/> 不具备测试条件，技术论证 <input type="checkbox"/>
测试说明	<div><div><div>云海Insight HD</div><div><div>HD Metrics</div><div>Kafka</div><div>Ranger</div><div>Spark2</div><div>DataSpace</div><div>Elasticsearch</div><div>Flink</div><div>主机</div><div>警告</div><div>集群管理</div><div>软件栈和版本</div><div>服务账号</div><div>安全认证</div><div>服务自动启动</div><div>许可证管理</div></div></div><div><div>警告</div><div>警告类型</div><div>警告名称定义</div><div>服务</div><div>最近状态改变</div><div>状态</div><div>正常 (10)</div><div>Metrics Monitor Status</div><div>HD Metrics</div><div>4 days ago</div><div>启用</div><div>正常 (10)</div><div>Ulimit for open files</div><div>Ambari</div><div>5 days ago</div><div>启用</div><div>正常 (10)</div><div>Ambari Agent心跳</div><div>Ambari</div><div>4 days ago</div><div>启用</div><div>正常 (5)</div><div>NodeManager Health</div><div>YARN</div><div>5 days ago</div><div>启用</div><div>正常 (5)</div><div>NodeManager Web UI</div><div>YARN</div><div>5 days ago</div><div>启用</div><div>正常 (5)</div><div>DataNode Heap Usage</div><div>HDFS</div><div>5 days ago</div><div>启用</div><div>正常 (5)</div><div>DataNode Unmounted Data Dir</div><div>HDFS</div><div>5 days ago</div><div>启用</div><div>正常 (5)</div><div>DataNode Process</div><div>HDFS</div><div>5 days ago</div><div>启用</div><div>正常 (5)</div><div>DataNode Web UI</div><div>HDFS</div><div>5 days ago</div><div>启用</div><div>正常 (5)</div><div>DataNode Storage</div><div>HDFS</div><div>5 days ago</div><div>启用</div></div></div> <div><div>云海Insight HD</div><div><div>HD Metrics</div><div>Kafka</div><div>Ranger</div><div>Spark2</div><div>DataSpace</div><div>Elasticsearch</div><div>Flink</div><div>主机</div><div>警告</div><div>集群管理</div><div>软件栈和版本</div><div>服务账号</div><div>安全认证</div><div>服务自动启动</div><div>许可证管理</div></div><div><div>警告</div><div>警告类型</div><div>警告名称定义</div><div>服务</div><div>最近状态改变</div><div>状态</div><div>通知 (0)</div><div>显示</div><div>所有 (0)</div><div>无警告报警</div><div>查看全部</div><div>正常 (10)</div><div>Metrics Monitor Status</div><div>HD Metrics</div><div>4 days ago</div><div>启用</div><div>正常 (10)</div><div>Ulimit for open files</div><div>Ambari</div><div>5 days ago</div><div>启用</div><div>正常 (10)</div><div>Ambari Agent心跳</div><div>Ambari</div><div>4 days ago</div><div>启用</div><div>正常 (5)</div><div>NodeManager Health</div><div>YARN</div><div>5 days ago</div><div>启用</div><div>正常 (5)</div><div>NodeManager Web UI</div><div>YARN</div><div>5 days ago</div><div>启用</div><div>正常 (5)</div><div>DataNode Heap Usage</div><div>HDFS</div><div>5 days ago</div><div>启用</div><div>正常 (5)</div><div>DataNode Unmounted Data Dir</div><div>HDFS</div><div>5 days ago</div><div>启用</div><div>正常 (5)</div><div>DataNode Process</div><div>HDFS</div><div>5 days ago</div><div>启用</div><div>正常 (5)</div><div>DataNode Web UI</div><div>HDFS</div><div>5 days ago</div><div>启用</div><div>正常 (5)</div><div>DataNode Storage</div><div>HDFS</div><div>5 days ago</div><div>启用</div></div></div>

1.1.7 扩展性

测试内容	集群的动态扩展
预期结果	通过管理工具对集群进行动态扩展
测试结果	符合预期，测试通过 <input checked="" type="checkbox"/> 不符合预期，测试不通过 <input type="checkbox"/> 不具备测试条件，技术论证 <input type="checkbox"/>
测试说明	扩容



1.2 云海 Insight 功能性测试

1.2.1 HDFS

测试内容	HDFS 的基本功能
预期结果	1、创建目录/查询/删除/下载/上传 2、从本地拷贝文件到 HDFS/从 HDFS 拷贝文件到本地
测试结果	符合预期，测试通过 <input checked="" type="checkbox"/> 不符合预期，测试不通过 <input type="checkbox"/> 不具备测试条件，技术论证 <input type="checkbox"/>
测试说明	1、创建目录 <code>hdfs dfs -mkdir /test</code>

```
[hdfs@manager yarn]$ hadoop fs -mkdir /test
[hdfs@manager yarn]$ hadoop fs -ls /
Found 15 items
drwxrwxrwt - yarn      hadoop      0 2021-05-01 16:30 /app-logs
drwxr-xr-x - hdfs      hdfs      0 2021-05-01 16:30 /apps
drwxr-xr-x - yarn      hadoop      0 2021-05-01 15:45 /ats
drwxr-xr-x - hdfs      hdfs      0 2021-05-01 15:45 /atsv2
drwxr-xr-x - hdfs      hdfs      0 2021-05-01 15:45 /hdp
drwx----- - livy      hdfs      0 2021-05-01 16:29 /livy2-recovery
drwxr-xr-x - mapred    hdfs      0 2021-05-01 15:45 /mapred
drwxrwxrwx - mapred    hadoop      0 2021-05-01 15:46 /mr-history
drwx----- - dataspace hdfs      0 2021-05-01 17:53 /public
drwxr-xr-x - hdfs      hdfs      0 2021-05-01 15:45 /services
drwxrwxrwx - spark     hadoop      0 2021-05-06 14:39 /spark2-history
drwxr-xr-x - hdfs      hdfs      0 2021-05-06 14:39 /test
drwxrwxrwx - hdfs      hdfs      0 2021-05-01 16:30 /tmp
drwxrwxrwx - hdfs      hdfs      0 2021-05-01 16:30 /user
drwxr-xr-x - hdfs      hdfs      0 2021-05-01 16:28 /warehouse
[hdfs@manager yarn]$
```

2、查询（查看）目录 `hdfs dfs -ls /` `hdfs dfs -du /`

```
[hdfs@manager yarn]$ hadoop fs -ls /
Found 15 items
drwxrwxrwt - yarn      hadoop      0 2021-05-01 16:30 /app-logs
drwxr-xr-x - hdfs      hdfs      0 2021-05-01 16:30 /apps
drwxr-xr-x - yarn      hadoop      0 2021-05-01 15:45 /ats
drwxr-xr-x - hdfs      hdfs      0 2021-05-01 15:45 /atsv2
drwxr-xr-x - hdfs      hdfs      0 2021-05-01 15:45 /hdp
drwx----- - livy      hdfs      0 2021-05-01 16:29 /livy2-recovery
drwxr-xr-x - mapred    hdfs      0 2021-05-01 15:45 /mapred
drwxrwxrwx - mapred    hadoop      0 2021-05-01 15:46 /mr-history
drwx----- - dataspace hdfs      0 2021-05-01 17:53 /public
drwxr-xr-x - hdfs      hdfs      0 2021-05-01 15:45 /services
drwxrwxrwx - spark     hadoop      0 2021-05-06 14:39 /spark2-history
drwxr-xr-x - hdfs      hdfs      0 2021-05-06 14:39 /test
drwxrwxrwx - hdfs      hdfs      0 2021-05-01 16:30 /tmp
drwxrwxrwx - hdfs      hdfs      0 2021-05-01 16:30 /user
drwxr-xr-x - hdfs      hdfs      0 2021-05-01 16:28 /warehouse
[hdfs@manager yarn]$ hadoop fs -du /
622482      1867446      /app-logs
162020      4832324268   /apps
694         2082         /ats
16794690    50384070     /atsv2
3040456097  9121368291   /hdp
0           0            /livy2-recovery
0           0            /mapred
254180      762540       /mr-history
0           0            /public
4490        13470        /services
2364297     402743904    /spark2-history
0           0            /test
1394556     4183668      /tmp
42377860    127133580    /user
2310992     6932976      /warehouse
```

3、删除文件 `hdfs dfs -rm /data.txt`

```
[hdfs@manager yarn]$ hadoop fs -put /test.txt /test/
[hdfs@manager yarn]$ hadoop fs -ls /test
Found 1 items
-rw-r--r--  3 hdfs hdfs      69 2021-05-06 14:42 /test/test.txt
[hdfs@manager yarn]$ hadoop fs -rm -r /test/test.txt
21/05/06 14:42:57 INFO fs.TrashPolicyDefault: Moved: 'hdfs://mycluster/test/test.txt' to trash at: hdfs://mycluster/user/hdfs/.Trash/Current/test/test.txt
[hdfs@manager yarn]$ hadoop fs -ls /test
[hdfs@manager yarn]$
```

4、上传 `hdfs dfs -put /hdfs/data.txt /test`

```
[hdfs@manager yarn]$ hadoop fs -put /test.txt /test/
[hdfs@manager yarn]$ hadoop fs -ls /test
Found 1 items
-rw-r--r--  3 hdfs hdfs      69 2021-05-06 14:42 /test/test.txt
[hdfs@manager yarn]$
```

5、下载 `hdfs dfs -get /test/BMGA.txt ./`

	<pre>[root@manager yarn]# hadoop fs -get /test/test.txt /opt/ [root@manager yarn]# ls /opt/ ambari.properties containerd disk_metric.sh gpu-monitor-Insight.tar InsightV4.6.5-release.tar MODEL-STUDIO rh ambari-server-2.7.1.0-0.x86_64.rpm disk_journal.sh disk_mysql.sh InsightHdInstall license-0.0.1-SNAPSHOT.jar nvdocker test.txt CentOS-7-x86_64-Everything-1708.iso disk_log.sh disk_zk.sh InsightV4.6.5-release.sha256sum license-0.0.1-SNAPSHOT.zip repoin.py</pre>
--	---

1.2.2 MapReduce

测试内容	Yarn 功能性测试
预期结果	提交/删除/查询 MR 任务
测试结果	符合预期，测试通过 <input checked="" type="checkbox"/> 不符合预期，测试不通过 <input type="checkbox"/> 不具备测试条件，技术论证 <input type="checkbox"/>
测试说明	<p>Wordcount 用例：</p> <p>/usr/hdp/3.0.1.0-187/hadoop-mapreduce/hadoop-mapreduce-examples.jar</p> <p>1、提交任务</p> <p><code>hadoop jar /usr/hdp/3.0.1.0-187/hadoop-mapreduce/hadoop-mapreduce-examples.jar wordcount /test/data.txt /output/wordcount01</code></p> <pre>[root@manager yarn]# su hdfs [hdfs@manager yarn]# hadoop jar /usr/hdp/3.0.1.0-187/hadoop-mapreduce/hadoop-mapreduce-examples.jar wordcount /test/test.txt /test/wordcount01 21/05/06 14:45:15 INFO client.AHSProxy: Connecting to Application History server at master1.zzdx.com/172.28.92.2:10200 21/05/06 14:45:15 INFO client.ConfiguredRMFailoverProxyProvider: Failing over to rm2 21/05/06 14:45:15 INFO mapreduce.JobResourceUploader: Disabling Erasure Coding for path: /user/hdfs/.staging/job_1619954590422_0001 21/05/06 14:45:15 INFO input.FileInputFormat: Total input files to process : 1 21/05/06 14:45:15 INFO mapreduce.JobSubmitter: number of splits:1 21/05/06 14:45:15 INFO mapreduce.JobSubmitter: Submitting tokens for job: job_1619954590422_0001 21/05/06 14:45:15 INFO mapreduce.JobSubmitter: Executing with tokens: [] 21/05/06 14:45:16 INFO conf.Configuration: Found resource resource-types.xml at file:/etc/hadoop/3.0.1.0-187/0/resource-types.xml 21/05/06 14:45:16 INFO impl.YarnClientImpl: Submitted application application_1619954590422_0001 21/05/06 14:45:16 INFO mapreduce.Job: The url to track the job: http://master2.zzdx.com:8088/proxy/application_1619954590422_0001/ 21/05/06 14:45:16 INFO mapreduce.Job: Running job: job_1619954590422_0001 21/05/06 14:45:24 INFO mapreduce.Job: Job job_1619954590422_0001 running in uber mode : false 21/05/06 14:45:24 INFO mapreduce.Job: map 0% reduce 0% 21/05/06 14:45:31 INFO mapreduce.Job: map 100% reduce 0% 21/05/06 14:45:39 INFO mapreduce.Job: map 100% reduce 100% 21/05/06 14:45:39 INFO mapreduce.Job: Job job_1619954590422_0001 completed successfully 21/05/06 14:45:39 INFO mapreduce.Job: Counters: 53 File System Counters FILE: Number of bytes read=77 FILE: Number of bytes written=475423 FILE: Number of read operations=0 FILE: Number of large read operations=0 FILE: Number of write operations=0 HDFS: Number of bytes read=164 HDFS: Number of bytes written=47 HDFS: Number of read operations=8 HDFS: Number of large read operations=0 HDFS: Number of write operations=2 Job Counters Launched map tasks=1</pre>

```

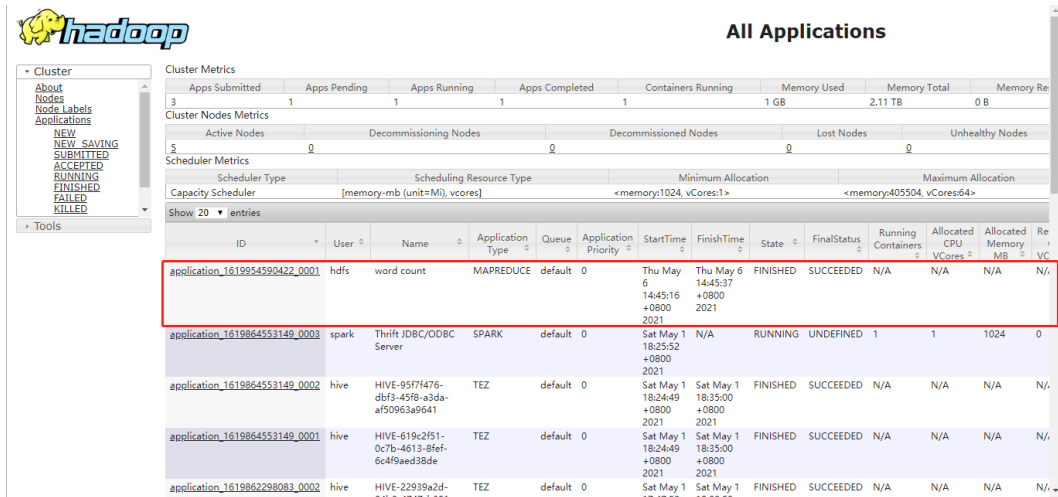
Map output materialized bytes=77
Input split bytes=95
Combine input records=12
Combine output records=6
Reduce input groups=6
Reduce shuffle bytes=77
Reduce input records=6
Reduce output records=6
Spilled Records=12
Shuffled Maps =1
Failed Shuffles=0
Merged Map outputs=1
GC time elapsed (ms)=127
CPU time spent (ms)=2250
Physical memory (bytes) snapshot=2846412800
Virtual memory (bytes) snapshot=62018621440
Total committed heap usage (bytes)=3233808384
Peak Map Physical memory (bytes)=2469511168
Peak Map Virtual memory (bytes)=21728722944
Peak Reduce Physical memory (bytes)=376901632
Peak Reduce Virtual memory (bytes)=40289898496

Shuffle Errors
BAD_ID=0
CONNECTION=0
IO_ERROR=0
WRONG_LENGTH=0
WRONG_MAP=0
WRONG_REDUCE=0

File Input Format Counters
Bytes Read=69
File Output Format Counters
Bytes Written=47
[hdfs@manager yarn]$

```

2、通过网页查看 MR 任务



The screenshot shows the Hadoop UI 'All Applications' page. A table lists various applications, with the first row highlighted in red. The table columns include ID, User, Name, Application Type, Queue, Application Priority, Start Time, Finish Time, State, Final Status, Running Containers, Allocated CPU, Allocated Memory, and Re...

ID	User	Name	Application Type	Queue	Application Priority	StartTime	FinishTime	State	FinalStatus	Running Containers	Allocated CPU	Allocated Memory	Re...
application_1619954590422_0001	hdfs	word count	MAPREDUCE	default	0	Thu May 6 14:45:16 +0800 2021	Thu May 6 14:45:37 +0800 2021	FINISHED	SUCCEEDED	N/A	N/A	N/A	N/A
application_1619864553149_0003	spark	Thrift JDBC/ODBC Server	SPARK	default	0	Sat May 1 18:25:52 +0800 2021	N/A	RUNNING	UNDEFINED	1	1	1024	0
application_1619864553149_0002	hive	HIVE-95f71476-dbf3-45f8-a3da-af50963a9641	TEZ	default	0	Sat May 1 18:24:49 +0800 2021	Sat May 1 18:35:00 +0800 2021	FINISHED	SUCCEEDED	N/A	N/A	N/A	N/A
application_1619864553149_0001	hive	HIVE-619c2f51-0c7b-4619-8fef-6c4f9aed38de	TEZ	default	0	Sat May 1 18:24:49 +0800 2021	Sat May 1 18:35:00 +0800 2021	FINISHED	SUCCEEDED	N/A	N/A	N/A	N/A
application_1619862298083_0002	hive	HIVE-22939a2d-...	TEZ	default	0	Sat May 1 18:24:49 +0800 2021	Sat May 1 18:35:00 +0800 2021	FINISHED	SUCCEEDED	N/A	N/A	N/A	N/A

3、命令行删除任务

yarn application -kill (任务名称)

运行成功后的结果

```

[hdfs@manager yarn]$ hadoop fs -cat /test/test.txt
java hello info info hello java spark hadoop java hadoop java nihao
[hdfs@manager yarn]$ hadoop fs -ls /test/wordcount01
Found 2 items
-rw-r--r--  3 hdfs hdfs          0 2021-05-06 14:45 /test/wordcount01/_SUCCESS
-rw-r--r--  3 hdfs hdfs        47 2021-05-06 14:45 /test/wordcount01/part-r-000000
^[[A[hdfs@manager yarn]$ hadoop fs -cat /test/wordcount01/part-r-000000
hadoop  2
hello   2
info    2
java    4
nihao   1
spark   1
[hdfs@manager yarn]$

```

1.2.3 Hbase

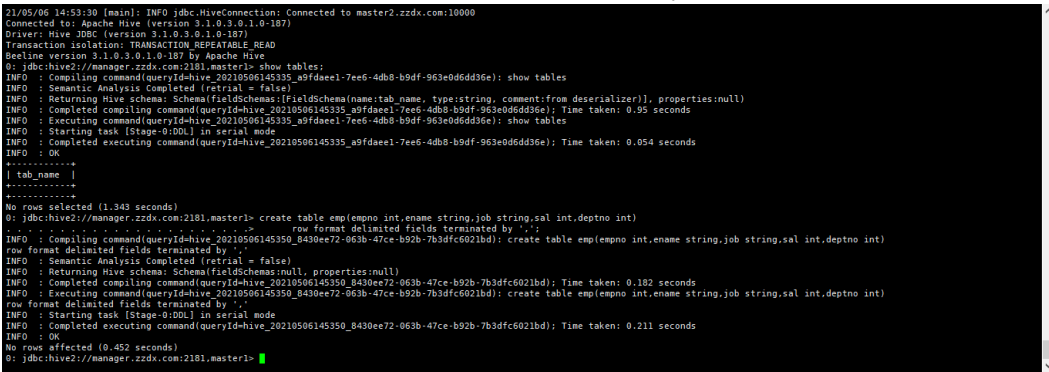
测试内容

Hbase 功能测试

预期结果	1、创建/插入/删除/启动/禁用表 2、查询/扫描表记录
测试结果	符合预期，测试通过 <input checked="" type="checkbox"/> 不符合预期，测试不通过 <input type="checkbox"/> 不具备测试条件，技术论证 <input type="checkbox"/>
测试说明	<p>进入 hbase shell: hbase shell</p> <p>1、list 列出表格</p> <pre>[root@manager yarn]# su hbase bash-4.2\$ hbase shell HBase Shell Use "help" to get list of supported commands. Use "exit" to quit this interactive shell. Version 2.0.0.3.0.1.0-187, r12d3ca6ec85ae2fclfa61f54f5c5a9ff6521cd81, Wed Sep 30 10:46:31 UTC 2020 Took 0.0012 seconds hbase(main):001:0> list TABLE 0 row(s) Took 0.3630 seconds => [] hbase(main):002:0></pre> <p>3、创建表 create 'person','info'</p> <pre>hbase(main):002:0> create 'person','info' Created table person Took 1.3465 seconds => Hbase::Table - person hbase(main):003:0> list TABLE person 1 row(s) Took 0.0140 seconds => ["person"] hbase(main):004:0></pre> <p>4、插入 put 'person','danny','info:sex','boy'</p> <pre>hbase(main):004:0> put 'person','danny','info:sex','boy' Took 0.2286 seconds hbase(main):005:0> scan 'person' ROW COLUMN+CELL danny column=info:sex, timestamp=1620283787648, value=boy 1 row(s) Took 0.0938 seconds hbase(main):006:0></pre> <p>5、删除 delete 'person','danny','info:sex'</p> <pre>hbase(main):009:0> scan 'person' ROW COLUMN+CELL David column=info:sex, timestamp=1620283838217, value=boy danny column=info:sex, timestamp=1620283787648, value=boy lisi column=info:sex, timestamp=1620283846266, value=boy zhangsan column=info:sex, timestamp=1620283857199, value=girl 4 row(s) Took 0.0124 seconds hbase(main):010:0> delete 'person','danny','info:sex' Took 0.0273 seconds hbase(main):011:0> scan 'person' ROW COLUMN+CELL David column=info:sex, timestamp=1620283838217, value=boy lisi column=info:sex, timestamp=1620283846266, value=boy zhangsan column=info:sex, timestamp=1620283857199, value=girl 3 row(s) Took 0.0093 seconds hbase(main):012:0></pre> <p>5、enable or disable</p>

	<pre>hbase(main):013:0> disable 'person' Took 0.7668 seconds hbase(main):014:0> enable 'person' Took 0.7418 seconds hbase(main):015:0> █</pre> <p>6、scan 'person'</p> <pre>hbase(main):015:0> scan 'person' COLUMN+CELL David column=info:sex, timestamp=1620283838217, value=boy lisi column=info:sex, timestamp=1620283846266, value=boy zhangsan column=info:sex, timestamp=1620283857199, value=girl 3 row(s) Took 0.0317 seconds hbase(main):016:0> █</pre> <p>7、get 'person','Lily'</p> <pre>hbase(main):018:0> get 'person','lisi' COLUMN CELL info:sex timestamp=1620283846266, value=boy 1 row(s) Took 0.0177 seconds hbase(main):019:0> █</pre>
--	---

1.2.4 Hive

测试内容	Hive 功能测试
预期结果	1、创建/更改/删除表 2、导入和导出数据 3、HQL 语句执行成功
测试结果	符合预期，测试通过 <input checked="" type="checkbox"/> 不符合预期，测试不通过 <input type="checkbox"/> 不具备测试条件，技术论证 <input type="checkbox"/>
测试说明	<p>进入 shell 行：hive</p> <p>1、创建表</p> <pre>create table emp(empno int,ename string,job string,sal int,deptno int) row format delimited fields terminated by ',' STORED AS TEXTFILE;</pre>  <p>2、插入数据（HIVE：支持插入，不支持删除和更新）</p> <pre>insert into emp values(1001,'danny','SA',5000,111);</pre>

```
0: jdbc:hive2://manager.zzdx.com:2181,master1> insert into emp values(1001,'danny','SA',5000,111);
INFO : Compiling command(queryId=hive_20210506145440_d7045024-2a61-4f36-b6c1-6ea70c5749a0): insert into emp values(1001,'danny','SA',5000,111)
INFO : Semantic Analysis Completed (retrial = false)
INFO : Returning Hive schema: Schema(fieldsSchemas:[FieldsSchema(name:col1, type:int, comment:null), FieldsSchema(name:col2, type:string, comment:null), FieldsSchema(name:col3, type:string, comment:null), FieldsSchema(name:col4, type:int, comment:null), FieldsSchema(name:col5, type:int, comment:null)], properties:null)
INFO : Completed compiling command(queryId=hive_20210506145440_d7045024-2a61-4f36-b6c1-6ea70c5749a0): Time taken: 2.668 seconds
INFO : Executing command(queryId=hive_20210506145440_d7045024-2a61-4f36-b6c1-6ea70c5749a0): insert into emp values(1001,'danny','SA',5000,111)
INFO : Query ID : hive_20210506145440_d7045024-2a61-4f36-b6c1-6ea70c5749a0
INFO : Total jobs = 1
INFO : Launching Job 1 out of 1
INFO : Starting task [Stage-1:MAPRED] in serial mode
INFO : Subscribed to counters: {} for queryId: hive_20210506145440_d7045024-2a61-4f36-b6c1-6ea70c5749a0
INFO : Tee session hasn't been created yet. Opening session
INFO : Dag name: insert into emp value...anny','SA',5000,111) (Stage-1)
INFO : Status: Running (Executing on YARN cluster with App id application_1619954590422_0002)

-----
VERTICES      MODE      STATUS TOTAL COMPLETED RUNNING PENDING FAILED KILLED
-----
Map 1 ..... container SUCCEEDED 1 1 0 0 0 0 0
Reducer 2 ..... container SUCCEEDED 1 1 0 0 0 0 0
-----
VERTICES: 02/02 (=====) 100% ELAPSED TIME: 6.63 s
-----
INFO : Status: DAG finished successfully in 6.02 seconds
INFO :
INFO : Query Execution Summary
INFO : -----
INFO : OPERATION DURATION
INFO : -----
INFO : Compile Query 2.67s
INFO : Prepare Plan 3.72s
```

```
INFO : INPUT_SPLIT_LENGTH_BYTES: 1
INFO : TaskCounter_Map_1_OUTPUT_Reducer_2:
INFO : ADDITIONAL_SPILLS_BYTES_READ: 0
INFO : ADDITIONAL_SPILLS_BYTES_WRITTEN: 0
INFO : ADDITIONAL_SPILL_COUNT: 0
INFO : OUTPUT_BYTES: 141
INFO : OUTPUT_BYTES_PHYSICAL: 161
INFO : OUTPUT_BYTES_WITH_OVERHEAD: 150
INFO : OUTPUT_LARGE_RECORDS: 0
INFO : OUTPUT_RECORDS: 1
INFO : SPILLED_RECORDS: 0
INFO : TaskCounter_Reducer_2_INPUT_Map_1:
INFO : FIRST_EVENT_RECEIVED: 66
INFO : INPUT_RECORDS_PROCESSED: 1
INFO : LAST_EVENT_RECEIVED: 66
INFO : NUM_FAILED_SHUFFLE_INPUTS: 0
INFO : NUM_SHUFFLED_INPUTS: 1
INFO : SHUFFLE_BYTES: 137
INFO : SHUFFLE_BYTES_DECOMPRESSED: 150
INFO : SHUFFLE_BYTES_DISK_DIRECT: 137
INFO : SHUFFLE_BYTES_TO_DISK: 0
INFO : SHUFFLE_BYTES_TO_MEM: 0
INFO : SHUFFLE_PHASE_TIME: 77
INFO : TaskCounter_Reducer_2_OUTPUT_out_Reducer_2:
INFO : OUTPUT_RECORDS: 0
INFO : Starting task [Stage-2:DEPENDENCY_COLLECTION] in serial mode
INFO : Starting task [Stage-0:MOVE] in serial mode
INFO : Loading data to table default.emp from hdfs://mycluster/warehouse/tablespace/managed/hive/emp/hive-staging_hive_2021-05-06_14-54-40_402_3774209835444939105-1/-ext-10000
INFO : Starting task [Stage-3:STATS] in serial mode
INFO : Completed executing command(queryId=hive_20210506145440_d7045024-2a61-4f36-b6c1-6ea70c5749a0): Time taken: 11.238 seconds
INFO : OK
No rows affected (13.993 seconds)
0: jdbc:hive2://manager.zzdx.com:2181,master1>
```

```
0: jdbc:hive2://manager.zzdx.com:2181,master1> select * from emp;
INFO : Compiling command(queryId=hive_20210506145543_f83ce2d2-2592-4e74-a23e-d0430285c32e): select * from emp
INFO : Semantic Analysis Completed (retrial = false)
INFO : Returning Hive schema: Schema(fieldsSchemas:[FieldsSchema(name:emp.empno, type:int, comment:null), FieldsSchema(name:emp.ename, type:string, comment:null), FieldsSchema(name:emp.job, type:string, comment:null), FieldsSchema(name:emp.sal, type:string, comment:null), FieldsSchema(name:emp.deptno, type:int, comment:null)], properties:null)
INFO : Completed compiling command(queryId=hive_20210506145543_f83ce2d2-2592-4e74-a23e-d0430285c32e): Time taken: 0.199 seconds
INFO : Executing command(queryId=hive_20210506145543_f83ce2d2-2592-4e74-a23e-d0430285c32e): select * from emp
INFO : Completed executing command(queryId=hive_20210506145543_f83ce2d2-2592-4e74-a23e-d0430285c32e): Time taken: 0.008 seconds
INFO : OK

+-----+
| emp.empno | emp.ename | emp.job | emp.sal | emp.deptno |
+-----+
| 1001      | danny    | SA      | 5000    | 111        |
+-----+
1 row selected (0.39 seconds)
0: jdbc:hive2://manager.zzdx.com:2181,master1>
```

4、删除表 drop table;

```
0: jdbc:hive2://manager.zzdx.com:2181,master1> show tables;
INFO : Compiling command(queryId=hive_20210506145604_838ecc15-356c-41ae-b565-4dca779083a5): show tables
INFO : Semantic Analysis Completed (retrial = false)
INFO : Returning Hive schema: Schema(fieldsSchemas:[FieldsSchema(name:tab_name, type:string, comment:from deserializer)], properties:null)
INFO : Completed compiling command(queryId=hive_20210506145604_838ecc15-356c-41ae-b565-4dca779083a5): Time taken: 0.025 seconds
INFO : Executing command(queryId=hive_20210506145604_838ecc15-356c-41ae-b565-4dca779083a5): show tables
INFO : Starting task [Stage-0:DDL] in serial mode
INFO : Completed executing command(queryId=hive_20210506145604_838ecc15-356c-41ae-b565-4dca779083a5): Time taken: 0.007 seconds
INFO : OK

+-----+
| tab_name |
+-----+
| emp      |
+-----+
1 row selected (0.069 seconds)
0: jdbc:hive2://manager.zzdx.com:2181,master1> drop table emp;
INFO : Compiling command(queryId=hive_20210506145608_2c64fb3b-5549-4c98-b68a-5d766cf82cf7): drop table emp
INFO : Semantic Analysis Completed (retrial = false)
INFO : Returning Hive schema: Schema(fieldsSchemas:null, properties:null)
INFO : Completed compiling command(queryId=hive_20210506145608_2c64fb3b-5549-4c98-b68a-5d766cf82cf7): Time taken: 0.025 seconds
INFO : Executing command(queryId=hive_20210506145608_2c64fb3b-5549-4c98-b68a-5d766cf82cf7): drop table emp
INFO : Starting task [Stage-0:DDL] in serial mode
INFO : Completed executing command(queryId=hive_20210506145608_2c64fb3b-5549-4c98-b68a-5d766cf82cf7): Time taken: 0.139 seconds
INFO : OK
No rows affected (0.184 seconds)
0: jdbc:hive2://manager.zzdx.com:2181,master1> show tables;
INFO : Compiling command(queryId=hive_20210506145613_53a63e86-0a8a-4d9c-8dbc-51b1ba56d014): show tables
INFO : Semantic Analysis Completed (retrial = false)
INFO : Returning Hive schema: Schema(fieldsSchemas:[FieldsSchema(name:tab_name, type:string, comment:from deserializer)], properties:null)
INFO : Completed compiling command(queryId=hive_20210506145613_53a63e86-0a8a-4d9c-8dbc-51b1ba56d014): Time taken: 0.021 seconds
```

```
0: jdbc:hive2://manager.zzdx.com:2181,master1> show tables;
INFO : Compiling command(queryId=hive_20210506145613_53a63e86-0a8a-4d9c-8dbc-51b1ba56d014): show tables
INFO : Semantic Analysis Completed (retrial = false)
INFO : Returning Hive schema: Schema(fieldsSchemas:[FieldsSchema(name:tab_name, type:string, comment:from deserializer)], properties:null)
INFO : Completed compiling command(queryId=hive_20210506145613_53a63e86-0a8a-4d9c-8dbc-51b1ba56d014): Time taken: 0.021 seconds
INFO : Executing command(queryId=hive_20210506145613_53a63e86-0a8a-4d9c-8dbc-51b1ba56d014): show tables
INFO : Starting task [Stage-0:DDL] in serial mode
INFO : Completed executing command(queryId=hive_20210506145613_53a63e86-0a8a-4d9c-8dbc-51b1ba56d014): Time taken: 0.007 seconds
INFO : OK

+-----+
| tab_name |
+-----+
+-----+
No rows selected (0.04 seconds)
0: jdbc:hive2://manager.zzdx.com:2181,master1>
```

5、导入

(HDFS) Load data inpath '/test/emp_data.txt' into table emp;

```
0: jdbc:hive2://manager.zzdx.com:2181,master1> create table emp(empno int,ename string,job string,sal int,deptno int)
row format delimited fields terminated by ',' STORED AS TEXTFILE;
INFO : Semantic Analysis Completed (retrial = false)
INFO : Returning Hive schema: Schema(fieldSchemas:null, properties:null)
INFO : Completed compiling command(queryId=hive_20210506152430_680e366f-1b89-4f27-ae5d-0c461d41da0c): create table emp(empno int,ename string,job string,sal int,deptno int)
INFO : Executing command(queryId=hive_20210506152430_680e366f-1b89-4f27-ae5d-0c461d41da0c): create table emp(empno int,ename string,job string,sal int,deptno int)
row format delimited fields terminated by ',' STORED AS TEXTFILE
INFO : Starting task [Stage:0:DDL] in serial mode
INFO : Completed executing command(queryId=hive_20210506152430_680e366f-1b89-4f27-ae5d-0c461d41da0c): Time taken: 0.105 seconds
INFO : OK
No rows affected (0.205 seconds)
0: jdbc:hive2://manager.zzdx.com:2181,master1> load data inpath '/test.txt' into table emp;
INFO : Compiling command(queryId=hive_20210506152500_de6fe36c-3e9f-4571-858b-3b8404ebd1e0): load data inpath '/test.txt' into table emp
INFO : Semantic Analysis Completed (retrial = false)
INFO : Returning Hive schema: Schema(fieldSchemas:null, properties:null)
INFO : Completed compiling command(queryId=hive_20210506152500_de6fe36c-3e9f-4571-858b-3b8404ebd1e0): Time taken: 0.115 seconds
INFO : Executing command(queryId=hive_20210506152500_de6fe36c-3e9f-4571-858b-3b8404ebd1e0): load data inpath '/test.txt' into table emp
INFO : Starting task [Stage:0:MOVE] in serial mode
INFO : Loading data to table default.emp from hdfs://mycluster/test.txt
INFO : Starting task [Stage:1:STATS] in serial mode
INFO : Completed executing command(queryId=hive_20210506152500_de6fe36c-3e9f-4571-858b-3b8404ebd1e0): Time taken: 0.213 seconds
INFO : OK
No rows affected (0.348 seconds)
0: jdbc:hive2://manager.zzdx.com:2181,master1>
```

6、导出

导出到本地

hive -e 'select * from default.emp' >> ./emp_export.txt

```
[root@manager ~]# hive -e 'select * from emp;' >> ./emp.txt
SLF4J: Class path contains multiple SLF4J bindings.
SLF4J: Found binding in [jar:file:/usr/hdp/3.0.1.0-187/hive/lib/log4j-slf4j-impl-2.10.0.jar!/org/slf4j/impl/StaticLoggerBinder.class]
SLF4J: Found binding in [jar:file:/usr/hdp/3.0.1.0-187/hadoop/lib/slf4j-log4j12-1.7.25.jar!/org/slf4j/impl/StaticLoggerBinder.class]
SLF4J: See http://www.slf4j.org/codes.html#multiple_bindings for an explanation.
SLF4J: Actual binding is of type [org.apache.hadoop.hive.log4j.Log4jLoggerFactory]
Connecting to jdbc:hive2://manager.zzdx.com:2181,master1.zzdx.com:2181,master2.zzdx.com:2181/default;serviceDiscoveryMode=zooKeeper;zooKeeperNamespace=hiveserver2
Enter username for jdbc:hive2://manager.zzdx.com:2181,master1.zzdx.com:2181,master2.zzdx.com:2181/default: hive
Enter password for jdbc:hive2://manager.zzdx.com:2181,master1.zzdx.com:2181,master2.zzdx.com:2181/default: *****
21/05/06 15:26:15 (main): INFO jdbc.HiveConnection: Connected to master2.zzdx.com:10000
Connected to: Apache Hive (version 3.1.0.3.0.1.0-187)
Driver: hive JDBC (version 3.1.0.3.0.1.0-187)
Transaction isolation: TRANSACTION_REPEATABLE_READ
INFO : Compiling command(queryId=hive_20210506152615_2d7cd946-54bd-460a-83bf-c74426018d3c): select * from emp
INFO : Semantic Analysis Completed (retrial = false)
INFO : Returning Hive schema: Schema(fieldSchemas:[FieldSchema(name:emp.empno, type:int, comment:null), FieldSchema(name:emp.ename, type:string, comment:null), FieldSchema(name:emp.job, type:string, comment:null), FieldSchema(name:emp.sal, type:int, comment:null), FieldSchema(name:emp.deptno, type:int, comment:null)], properties:null)
INFO : Completed compiling command(queryId=hive_20210506152615_2d7cd946-54bd-460a-83bf-c74426018d3c): select * from emp
INFO : Executing command(queryId=hive_20210506152615_2d7cd946-54bd-460a-83bf-c74426018d3c): select * from emp
INFO : Completed executing command(queryId=hive_20210506152615_2d7cd946-54bd-460a-83bf-c74426018d3c): Time taken: 0.007 seconds
INFO : OK
2 rows selected (0.266 seconds)
Beeline version 3.1.0.3.0.1.0-187 by Apache Hive
Closing: 0: jdbc:hive2://manager.zzdx.com:2181,master1.zzdx.com:2181,master2.zzdx.com:2181/default;serviceDiscoveryMode=zooKeeper;zooKeeperNamespace=hiveserver2
[root@manager ~]# cd /
[root@manager ~]# cat emp.txt
+-----+-----+-----+-----+-----+
| emp.empno | emp.ename | emp.job | emp.sal | emp.deptno |
+-----+-----+-----+-----+-----+
| 1001      | dummy     | SA      | 5000    | 11          |
| 1002      | lisi      | SA      | 500     | 10          |
+-----+-----+-----+-----+-----+
[root@manager ~]#
```

导出到 hdfs

export table emp to '/export_data';

```
0: jdbc:hive2://manager.zzdx.com:2181,master1> export table emp to '/export_data';
INFO : Compiling command(queryId=hive_20210506152743_9c2e00dc-ea3d-411b-a4e4-3b97bb8ae508): export table emp to '/export_data'
INFO : Semantic Analysis Completed (retrial = false)
INFO : Returning Hive schema: Schema(fieldSchemas:null, properties:null)
INFO : Completed compiling command(queryId=hive_20210506152743_9c2e00dc-ea3d-411b-a4e4-3b97bb8ae508): Time taken: 0.105 seconds
INFO : Executing command(queryId=hive_20210506152743_9c2e00dc-ea3d-411b-a4e4-3b97bb8ae508): export table emp to '/export_data'
INFO : Starting task [Stage:0:REPL_DUMP] in serial mode
INFO : Completed executing command(queryId=hive_20210506152743_9c2e00dc-ea3d-411b-a4e4-3b97bb8ae508): Time taken: 0.084 seconds
INFO : OK
No rows affected (0.208 seconds)
```

	<pre> [root@manager /]# hadoop fs -ls / Found 16 items drwxrwxrwt - yarn hadoop 0 2021-05-06 14:45 /app-logs drwxr-xr-x - hdfs hdfs 0 2021-05-01 16:30 /apps drwxr-xr-x - yarn hadoop 0 2021-05-01 15:45 /ats drwxr-xr-x - hdfs hdfs 0 2021-05-01 15:45 /atsv2 drwxr-xr-x - hive hdfs 0 2021-05-06 15:27 /export_data drwxr-xr-x - hdfs hdfs 0 2021-05-01 15:45 /hdp drwx----- - livy hdfs 0 2021-05-01 16:29 /livy2-recovery drwxr-xr-x - mapred hdfs 0 2021-05-01 15:45 /mapred drwxrwxrwx - mapred hadoop 0 2021-05-01 15:46 /mr-history drwx----- - dataspace hdfs 0 2021-05-01 17:53 /public drwxr-xr-x - hdfs hdfs 0 2021-05-01 15:45 /services drwxrwxrwx - spark hadoop 0 2021-05-06 15:28 /spark2-history drwxr-xr-x - hdfs hdfs 0 2021-05-06 14:45 /test drwxrwxrwx - hdfs hdfs 0 2021-05-01 16:30 /tmp drwxrwxrwx - hdfs hdfs 0 2021-05-06 14:42 /user drwxr-xr-x - hdfs hdfs 0 2021-05-01 16:28 /warehouse [root@manager /]# hadoop fs -ls /export_data Found 2 items -rw-r--r-- 3 hive hdfs 1485 2021-05-06 15:27 /export_data/_metadata drwxr-xr-x - hive hdfs 0 2021-05-06 15:27 /export_data/data [root@manager /]# hadoop fs -ls /export_data/data Found 1 items -rw-r--r-- 3 hive hdfs 42 2021-05-06 15:27 /export_data/data/test.txt [root@manager /]# hadoop fs -ls /export_data/data/test.txt -rw-r--r-- 3 hive hdfs 42 2021-05-06 15:27 /export_data/data/test.txt [root@manager /]# hadoop fs -cat /export_data/data/test.txt 1001,danny,SA,5000,11 1002,lisi,SA,500,10 [root@manager /]# </pre>
--	---

1.2.5 Flume

测试内容	Flume 功能测试
预期结果	实时采集文件目录
测试结果	<p>符合预期，测试通过 <input checked="" type="checkbox"/></p> <p>不符合预期，测试不通过 <input type="checkbox"/></p> <p>不具备测试条件，技术论证 <input type="checkbox"/></p>
测试说明	1、先编辑 agent 配置文件

```

[flume@manager myagent]$
[flume@manager myagent]$ more example.conf
# 配置Agent a1各个组件的名称
a1.sources = r1
a1.sinks = k1
a1.channels = c1

# 配置Agent a1的source r1的属性
a1.sources.r1.type = netcat
a1.sources.r1.bind = localhost
a1.sources.r1.port = 44444

# 配置Agent a1的sink k1的属性
a1.sinks.k1.type = logger

# 配置Agent a1的channel c1的属性, channel是用来缓冲Event数据的
a1.channels.c1.type = memory
a1.channels.c1.capacity = 1000
a1.channels.c1.transactionCapacity = 100

# 把source和sink绑定到channel上
a1.sources.r1.channels = c1
a1.sinks.k1.channel = c1
[flume@manager myagent]$

```

2、运行 agent 文件

```

/usr/hdp/3.0.1.0-187/flume/bin/flume-ng agent --name a1 --conf-file
/home/flume/myagent/example.conf --conf conf -Dflume.root.logger=INFO,console

```

2、先编辑 agent 配置文件

```

[root@manager myagent]# cat 4a.conf
a1.sources = r1
a1.sinks = k1
a1.channels = c1

a1.sources.r1.type = netcat
a1.sources.r1.bind = 172.28.92.1
a1.sources.r1.port = 44446

a1.sinks.k1.type = logger
a1.channels.c1.type = memory
a1.channels.c1.capacity=1000
a1.channels.c1.transactionCapacity=100

a1.sources.r1.channels=c1
a1.sinks.k1.channel=c1
[root@manager myagent]# █

```

2、运行 agent 文件

```

/usr/hdp/3.0.1.0-187/flume/bin/flume-ng agent --name a1 --conf-file
/home/flume/myagent/example.conf --conf conf -Dflume.root.logger=INFO,console

```

```

3.0.1.0-187/zookeeper/*:/conf:/lib/* -Djava.library.path=:/usr/hdp/3.0.1.0-187/hadoop/lib/native/Linux-amd64-64:/usr/hdp/3.0.1.0-187/hadoop/lib/native/Linux-amd64-64:/usr/hdp/3.0.1.0-187/hadoop/lib/native/Linux-amd64-64:/usr/hdp/3.0.1.0-187/hadoop/lib/native/org.apache.flume
ication --name a1 --conf-file /root/myagent/4a.conf
SLF4J: Class path contains multiple SLF4J bindings.
SLF4J: Found binding in [jar:file:/usr/hdp/3.0.1.0-187/flume/lib/slf4j-log4j12-1.6.1.jar!/org/slf4j/impl/StaticLoggerBinder.class]
SLF4J: See http://www.slf4j.org/codes.html#multiple_bindings for an explanation.
21/05/06 17:58:42 INFO node.PollingPropertiesFileConfigurationProvider: Configuration provider starting
21/05/06 17:58:42 INFO conf.FlumeConfiguration: Reloading configuration file:/root/myagent/4a.conf
21/05/06 17:58:42 INFO conf.FlumeConfiguration: Added sinks: k1 Agent: a1
21/05/06 17:58:42 INFO conf.FlumeConfiguration: Processing:k1
21/05/06 17:58:42 INFO conf.FlumeConfiguration: Processing:k1
21/05/06 17:58:42 INFO conf.FlumeConfiguration: Post-validation flume configuration contains configuration for agents: [a1]
21/05/06 17:58:42 INFO node.AbstractConfigurationProvider: Creating channels
21/05/06 17:58:42 INFO channel.DefaultChannelFactory: Creating instance of channel c1 type memory
21/05/06 17:58:42 INFO node.AbstractConfigurationProvider: Created channel c1
21/05/06 17:58:42 INFO source.DefaultSourceFactory: Creating instance of source r1, type netcat
21/05/06 17:58:42 INFO sink.DefaultSinkFactory: Creating instance of sink: k1, type: logger

```

3、启动监听并发送消息

```
Xshell 5 (Build 1005)
Copyright (c) 2002-2016 NetSarang Computer, Inc. All rights reserved.

Type 'help' to learn how to use Xshell prompt.
[ld:\~]$ telnet 172.28.92.1 44446

Connecting to 172.28.92.1:44446...
Connection established.
To escape to local shell, press 'Ctrl+Alt+J'.
sad
OK
asda
OK
dasdas
OK
sda
OK

Connection closed by foreign host.

Disconnected from remote host(172.28.92.1:44446) at 18:00:49.

Type 'help' to learn how to use Xshell prompt.
[ld:\~]$ █
```

1、查看 flume 采集记录

```
21/05/06 17:58:42 INFO source.DefaultSourceFactory: Creating instance of source r1, type netcat
21/05/06 17:58:42 INFO sink.DefaultSinkFactory: Creating instance of sink: k1, type: logger
21/05/06 17:58:42 INFO node.AbstractConfigurationProvider: Channel c1 connected to [r1, k1]
21/05/06 17:58:42 INFO node.Application: Starting new configuration: { sourceRunners: {r1=EventDrivenSourceRunner: { source:org.apache.flume.source.NetcatSource[name:r1,state:IDLE] }} sinkRunners: {k1=SinkRunner: { policy:org.apache.flume.sink.DefaultSinkProcessor@27315100 counterGroup: { name:null counters:{}} }} channels: {c1=org.apache.flume.channel.MemoryChannel[name: c1]} }
21/05/06 17:58:42 INFO node.Application: Starting Channel c1
21/05/06 17:58:42 INFO instrumentation.MonitoredCounterGroup: Monitored counter group for type: CHANNEL, name: c1: Successfully registered new MBean.
21/05/06 17:58:42 INFO instrumentation.MonitoredCounterGroup: Component type: CHANNEL, name: c1 started
21/05/06 17:58:42 INFO node.Application: Starting Sink k1
21/05/06 17:58:42 INFO node.Application: Starting Source r1
21/05/06 17:58:42 INFO source.NetcatSource: Source starting
21/05/06 17:58:42 INFO source.NetcatSource: Created serverSocket:sun.nio.ch.ServerSocketChannelImpl[/172.28.92.1:44446]
21/05/06 17:59:28 INFO sink.LoggerSink: Event: { headers:{} body: 73 61 64 00 sad. }
21/05/06 17:59:28 INFO sink.LoggerSink: Event: { headers:{} body: 61 73 64 61 00 asda. }
21/05/06 17:59:28 INFO sink.LoggerSink: Event: { headers:{} body: 64 63 73 64 61 73 00 dasdas. }
21/05/06 17:59:28 INFO sink.LoggerSink: Event: { headers:{} body: 73 64 61 00 sda. }
█
```

1.2.6 Spark

测试内容	spark 功能测试
预期结果	提交、执行任务
测试结果	符合预期，测试通过 <input checked="" type="checkbox"/> 不符合预期，测试不通过 <input type="checkbox"/> 不具备测试条件，技术论证 <input type="checkbox"/>
测试说明	Netcat Nc -l -p 1234 /usr/bin/run-example streaming.NetworkWordCount localhost 1234

```
[root@manager myagent]# nc -l -p 1234
sadasd
asd
asd
asd
as
das
d
sa
sf
ds
fsd
d
fsd
f
sdf

sdfsdf
s
dfs
dfsfsd
fsd
f
sd
fsd
fsd
f
sdf
[root@manager myagent]#
```

```
DD[7] at reduceByKey at NetworkWordCount.scala:55) (first 15 tasks are for partitions Vector(
1))
21/05/06 16:51:13 INFO YarnScheduler: Adding task set 6.0 with 1 tasks
21/05/06 16:51:13 INFO TaskSetManager: Starting task 0.0 in stage 6.0 (TID 73, hd02.zzdx.com,
executor 1, partition 1, NODE_LOCAL, 7660 bytes)
21/05/06 16:51:13 INFO BlockManagerInfo: Added broadcast_5_piece0 in memory on hd02.zzdx.com:
44471 (size: 1714.0 B, free: 408.9 MB)
21/05/06 16:51:13 INFO TaskSetManager: Finished task 0.0 in stage 6.0 (TID 73) in 49 ms on hd
02.zzdx.com (executor 1) (1/1)
21/05/06 16:51:13 INFO YarnScheduler: Removed TaskSet 6.0, whose tasks have all completed, fr
om pool
21/05/06 16:51:13 INFO DAGScheduler: ResultStage 6 (print at NetworkWordCount.scala:56) finis
hed in 0.055 s
21/05/06 16:51:13 INFO DAGScheduler: Job 3 finished: print at NetworkWordCount.scala:56, took
0.056983 s
-----
Time: 1620291070000 ms
-----
(d,2)
(asd,3)
(sdfsdf,1)
(as,1)
(,1)
(sadasd,1)
(dfsfsd,1)
(f,3)
(sa,1)
(das,1)
...

21/05/06 16:51:13 INFO JobScheduler: Finished job streaming job 1620291070000 ms.0 from job s
et of time 1620291070000 ms
21/05/06 16:51:13 INFO JobScheduler: Total delay: 3.916 s for time 1620291070000 ms (executio
```

1.2.7 Kafka

测试内容	kafka 功能测试
预期结果	新建 topic 并且正常生产、消费

测试结果	<p>符合预期，测试通过 <input checked="" type="checkbox"/></p> <p>不符合预期，测试不通过 <input type="checkbox"/></p> <p>不具备测试条件，技术论证 <input type="checkbox"/></p>
测试说明	<p>1、创建主题</p> <p>/usr/hdp/3.0.1.0-187/kafka/bin/kafka-topics.sh --create --zookeeper localhost:2181 --replication-factor 1 --partitions 1 --topic MyTopic</p> <p>--create 创建主题</p> <p>--topic 主题名称</p> <p>--zookeeper 集群地址</p> <p>--replication-factor 多少副本</p> <pre>[root@manager myagent]# /usr/hdp/3.0.1.0-187/kafka/bin/kafka-topics.sh --create --zookeeper localhost:2181 --replication-factor 1 --partitions 1 --topic MyTopic Created topic "MyTopic". [root@manager myagent]#</pre> <p>2、查看主题</p> <p>/usr/hdp/2.6.4.0-91/kafka/bin/kafka-topics.sh --zookeeper localhost:2181 --list</p> <pre>[root@manager myagent]# /usr/hdp/3.0.1.0-187/kafka/bin/kafka-topics.sh --zookeeper localhost:2181 --list MyTopic ambari_kafka_service_check [root@manager myagent]#</pre> <p>3、生产者</p> <p>./kafka-console-producer.sh --broker-list master.bigdata.com:6667,node01.bigdata.com:6667,node02.bigdata.com:6667 --topic MyTopic</p> <p>2、消费者</p> <p>/usr/hdp/2.6.4.0-91/kafka/bin/kafka-console-consumer.sh --zookeeper localhost:2181 --topic MyTopic --from-beginning</p> <div> <pre> -rwxr-xr-x. 1 root root 871 Sep 19 2018 kafka-delegation-tokens.sh -rwxr-xr-x. 1 root root 869 Sep 19 2018 kafka-delete-records.sh -rwxr-xr-x. 1 root root 863 Sep 19 2018 kafka-log-dirs.sh -rwxr-xr-x. 1 root root 862 Sep 19 2018 kafka-mirror-maker.sh -rwxr-xr-x. 1 root root 886 Sep 19 2018 kafka-preferred-replica-election.sh -rwxr-xr-x. 1 root root 1327 Sep 19 2018 kafka-producer-perf-test.sh -rwxr-xr-x. 1 root root 874 Sep 19 2018 kafka-reassign-partitions.sh -rwxr-xr-x. 1 root root 1234 Sep 19 2018 kafka-replay-log-producer.sh -rwxr-xr-x. 1 root root 874 Sep 19 2018 kafka-replica-verification.sh -rwxr-xr-x. 1 root root 9811 Sep 19 2018 kafka-run-class.sh -rwxr-xr-x. 1 root root 1376 Sep 19 2018 kafka-server-start.sh -rwxr-xr-x. 1 root root 987 Sep 19 2018 kafka-server-stop.sh -rwxr-xr-x. 1 root root 1235 Sep 19 2018 kafka-simple-consumer-shell.sh -rwxr-xr-x. 1 root root 945 Sep 19 2018 kafka-stream-application-reset.sh -rwxr-xr-x. 1 root root 863 Sep 19 2018 kafka-topics.sh -rwxr-xr-x. 1 root root 958 Sep 19 2018 kafka-verifiable-consumer.sh -rwxr-xr-x. 1 root root 958 Sep 19 2018 kafka-verifiable-producer.sh -rwxr-xr-x. 1 root root 4371 Sep 19 2018 kafka-zookeeper-run-class.sh -rwxr-xr-x. 1 root root 1722 Sep 19 2018 trogdor.sh drwxr-xr-x. 2 root root 4096 May 1 17:04 windows -rwxr-xr-x. 1 root root 867 Sep 19 2018 zookeeper-security-migration.sh -rwxr-xr-x. 1 root root 1401 Sep 19 2018 zookeeper-server-start.sh -rwxr-xr-x. 1 root root 1401 Sep 19 2018 zookeeper-server-stop.sh -rwxr-xr-x. 1 root root 958 Sep 19 2018 zookeeper-shell.sh [root@manager bin]# ./kafka-console-producer.sh --broker-list h096.zdx.com:6667,h07.zdx.com:6667 >hello >nihao >nihao >nihao >hello >java >inspur >hellpo </pre> <pre> down 21/05/06 16:51:15 INFO SchedulerExtensionServices: Stopping SchedulerExtensionServices (serviceOption=None, services=List(), started=false) 21/05/06 16:51:15 INFO YarnClientSchedulerBackend: Stopped 21/05/06 16:51:15 INFO MapOutputTrackerMasterEndpoint: MapOutputTrackerMasterEndpoint stopped 21/05/06 16:51:15 INFO MemoryStore: MemoryStore cleared 21/05/06 16:51:15 INFO BlockManager: BlockManager stopped 21/05/06 16:51:15 INFO BlockManagerMaster: BlockManagerMaster stopped 21/05/06 16:51:15 INFO OutputCommitCoordinatorOutputCommitCoordinatorEndpoint: OutputCommitCoordinator stopped 21/05/06 16:51:15 INFO SparkContext: Successfully stopped SparkContext 21/05/06 16:51:15 INFO ShutdownHookManager: Shutdown hook called 21/05/06 16:51:15 INFO ShutdownHookManager: Deleting directory /tmp/spark-01661525-fe6d-4a4d-bd41-12487f6ce893 21/05/06 16:51:15 INFO ShutdownHookManager: Deleting directory /tmp/spark-57dd9acc-0f18-482e-9716-44bd735c205f [root@manager ~]# "C [root@manager ~]# /usr/hdp/3.0.1.0-187/kafka/bin/kafka-console-consumer.sh --zookeeper localhost:2181 --topic MyTopic --from-beginning Using the ConsoleConsumer with old consumer is deprecated and will be removed in a future major release. Consider using the new consumer by passing [bootstrap-server] instead of [zookeeper] [1] hello nihao nihao hello java inspur hellpo </pre> </div>

1.2.8 ElasticSearch

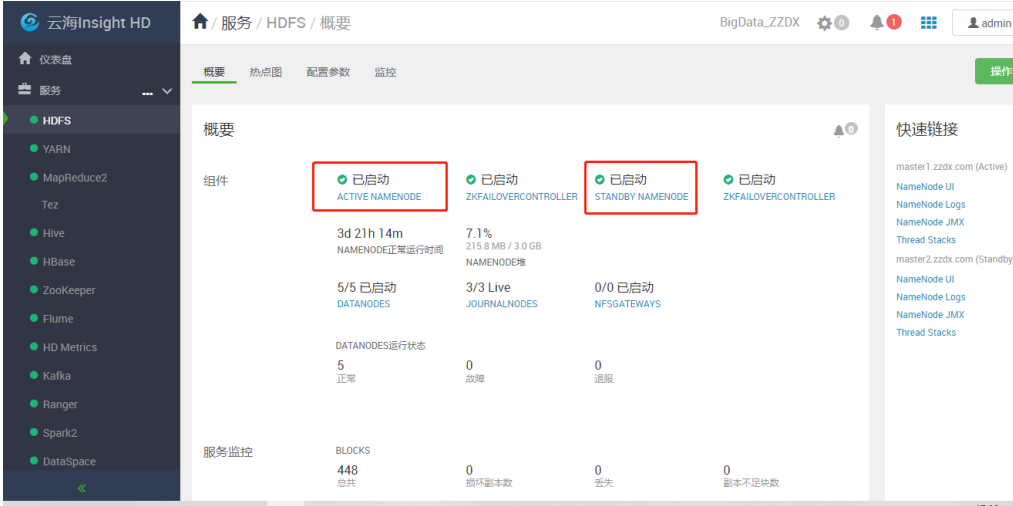
测试内容	Elasticsearch 功能测试
------	--------------------

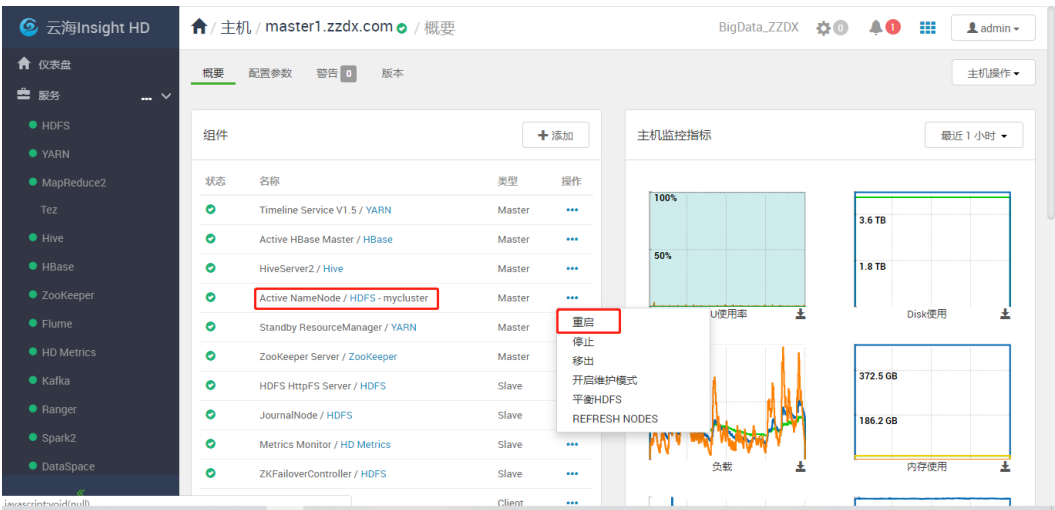
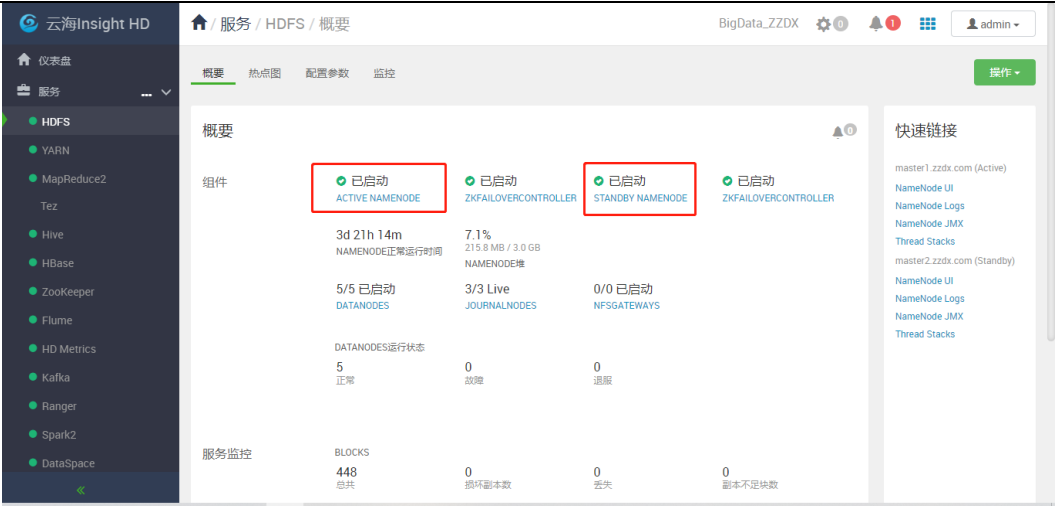
预期结果	查询上下文功能正常
测试结果	符合预期，测试通过 <input checked="" type="checkbox"/> 不符合预期，测试不通过 <input type="checkbox"/> 不具备测试条件，技术论证 <input type="checkbox"/>
测试说明	<p>1、创建 mapping</p> <pre>[root@manager bin]# curl -XPUT '172.28.92.9:9200/twitter?pretty' -H 'Content-Type: application/json' -d '{ "settings": { "index": { "number_of_shards": 3, "number_of_replicas": 2 } } }' { "acknowledged" : true, "shards_acknowledged" : true, "index" : "twitter" }</pre> <pre>[root@manager bin]# curl -XPUT '172.28.92.9:9200/twitter/_mapping?pretty' -H 'Content-Type: application/json' -d '{ "properties": { "user_name": { "type": "text", "fielddata": true, "analyzer": "standard", "search_analyzer": "whitespace" }, "content": { "type": "text" }, "status": { "type": "text" }, "publish_date": { "type": "text" } } }' { "acknowledged" : true }</pre> <p>2、插入数据</p> <pre>[root@manager bin]# curl -XPOST '172.28.92.9:9200/twitter/_doc/1?pretty' -H 'Content-Type: application/json' -d '{ "user_name": "test", "content": "test_content", "status": "test_status", "publish_date": "test_publish_date" }' { "_index" : "twitter", "_type" : "doc", "_id" : "1", "_version" : 1, "result" : "created", "_shards" : { "total" : 3, "successful" : 3, "failed" : 0 }, "_seq_no" : 0, "_primary_term" : 1 }</pre> <p>3、查询数据</p>

	<pre>[root@manager bin]# curl -XGET '172.28.92.9:9200/twitter/_search?pretty' { "took" : 1031, "timed_out" : false, "_shards" : { "total" : 3, "successful" : 3, "skipped" : 0, "failed" : 0 }, "hits" : { "total" : { "value" : 1, "relation" : "eq" }, "max_score" : 1.0, "hits" : [{ "_index" : "twitter", "_type" : "_doc", "_id" : "1", "_score" : 1.0, "_source" : { "user_name" : "test", "content" : "test_content", "status" : "test_status", "publish_date" : "test_publish_date" } }] } }</pre>

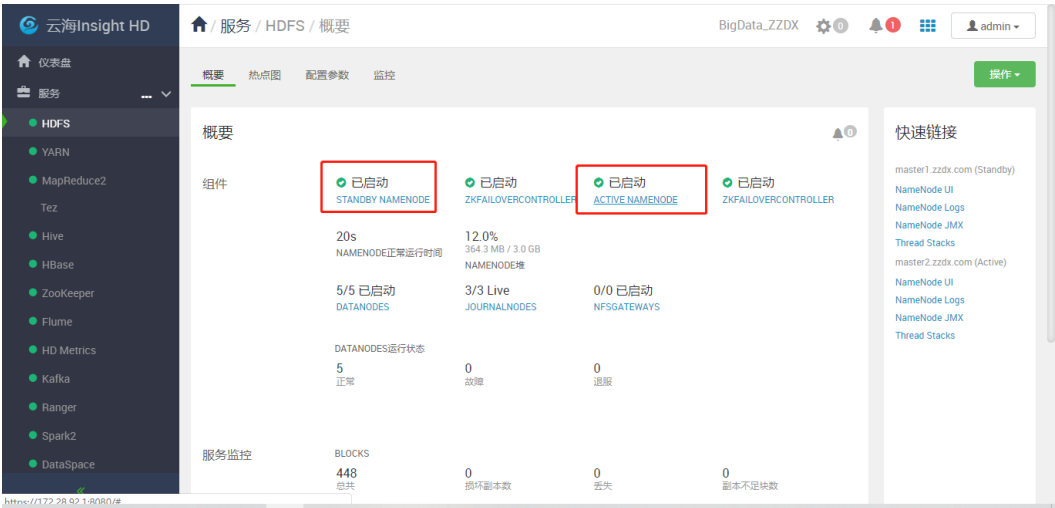
1.3 云海 Insight 高可用性测试

1.3.1 HDFS 的高可用测试

测试内容	HDFS 高可用实现
预期结果	Namenode 主节点失效后，standby 节点能够变更为 active namenode
测试结果	符合预期，测试通过 <input checked="" type="checkbox"/> 不符合预期，测试不通过 <input type="checkbox"/> 不具备测试条件，技术论证 <input type="checkbox"/>
测试说明	<p>开启 hdfs HA</p>  <p>功能测试：</p>



Standby namenode 会迅速变更为 active namenode 保证高可用



1.3.2 Hbase 的高可用测试

测试内容	Hbase 高可用实现
------	-------------

预期结果	HMaster 节点失效后，standby Hmaster 会变为 active Hmaster
测试结果	<div>符合预期，测试通过 <input checked="" type="checkbox"/></div> <div>不符合预期，测试不通过 <input type="checkbox"/></div> <div>不具备测试条件，技术论证 <input type="checkbox"/></div>
测试说明	<div><div><div>云海Insight HD</div><div>仪表盘</div><div>服务</div><div>HDFS</div><div>YARN</div><div>MapReduce2</div><div>Tez</div><div>Hive</div><div>HBase</div><div>ZooKeeper</div><div>Flume</div><div>HD Metrics</div><div>Kafka</div><div>Ranger</div><div>Spark2</div><div>DataSpace</div></div><div>服务 / HBase / 概要</div><div>BigData_ZZDX</div><div>admin</div><div>概要</div><div>组件</div><div>已启动</div><div>已启动</div><div>ACTIVE HBASE MASTER</div><div>STANDBY HBASE MASTER</div><div>5/5 Live</div><div>0/0 Live</div><div>0</div><div>REGIONSERVERS</div><div>PHOENIX QUERY 服务器</div><div>迁移中的REGION</div><div>0/0 运行状态</div><div>0/0 运行状态</div><div>0/0 运行状态</div><div>0/0 运行状态</div><div>REGIONSERVER1S</div><div>REGIONSERVER2S</div><div>REGIONSERVER3S</div><div>REGIONSERVER4S</div><div>0/0 运行状态</div><div>5/5 运行状态</div><div>5/5 运行状态</div><div>REGIONSERVERS</div><div>HBASE RESTSERVERS</div><div>HBASE THRIFTSERVERS</div><div>服务监控</div><div>4d 1h 8m</div><div>4d 1h 8m</div><div>每个RegionServer的</div><div>MASTER已启动</div><div>MASTER已激活</div><div>0.8个region</div><div>平均负载</div><div>4.8%</div><div>292.6 MB / 6.0 GB</div><div>MASTER堆</div><div>快速链接</div><div>master1.zzdx.com (Active)</div><div>HBase Master UI</div><div>HBase Logs</div><div>Zookeeper Info</div><div>HBase Master JMX</div><div>Debug Dump</div><div>Thread Stacks</div><div>master2.zzdx.com (Standby)</div><div>HBase Master UI</div><div>HBase Logs</div><div>Zookeeper Info</div><div>HBase Master JMX</div><div>Debug Dump</div><div>Thread Stacks</div></div>

云海Insight HD

仪表盘

服务

HDFS

YARN

MapReduce2

Tez

Hive

HBase

ZooKeeper

Flume

HD Metrics

Kafka

Ranger

Spark2

DataSpace

主机 master1.zzdx.com

BigData_ZZDX

admin

概要

配置参数

警告 0

版本

组件

添加

状态

名称

类型

操作

Timeline Service V1.5 / YARN

Master

...

Active HBase Master / HBase

Master

...

HiveServer2 / Hive

Master

...

Standby NameNode / HDFS - mycluster

Master

...

Standby ResourceManager / YARN

Master

...

ZooKeeper Server / ZooKeeper

Master

...

HDFS HttpFS Server / HDFS

Slave

...

JournalNode / HDFS

Slave

...

Metrics Monitor / HD Metrics

Slave

...

ZKFailoverController / HDFS

Slave

...

Client

主机监控指标

最近 1 小时

100%

3.6 TB

1.8 TB

1

0.5

负载

372.5 GB

186.2 GB

内存使用

重启

停止

开启维护模式

删除

云海Insight HD

仪表盘

服务

HDFS

YARN

MapReduce2

Tez

Hive

HBase

ZooKeeper

Flume

HD Metrics

Kafka

Ranger

Spark2

DataSpace

服务 / HBase / 概要

BigData_ZZDX

admin

概要

组件

已启动

已启动

STANDBY HBASE MASTER

ACTIVE HBASE MASTER

5/5 Live

0/0 Live

0

REGIONSERVERS

PHOENIX QUERY 服务器

迁移中的REGION

0/0 运行状态

0/0 运行状态

0/0 运行状态

0/0 运行状态

REGIONSERVER1S

REGIONSERVER2S

REGIONSERVER3S

REGIONSERVER4S

0/0 运行状态

5/5 运行状态

5/5 运行状态

REGIONSERVERS

HBASE RESTSERVERS

HBASE THRIFTSERVERS

服务监控

4d 1h 9m

4d 1h 9m

每个RegionServer的

MASTER已启动

MASTER已激活

0.8个region

平均负载

3.4%

206.0 MB / 6.0 GB

MASTER堆

快速链接

master1.zzdx.com

HBase Master UI

HBase Logs

Zookeeper Info

HBase Master JMX

Debug Dump

Thread Stacks

master2.zzdx.com (Active)

HBase Master UI

HBase Logs

Zookeeper Info

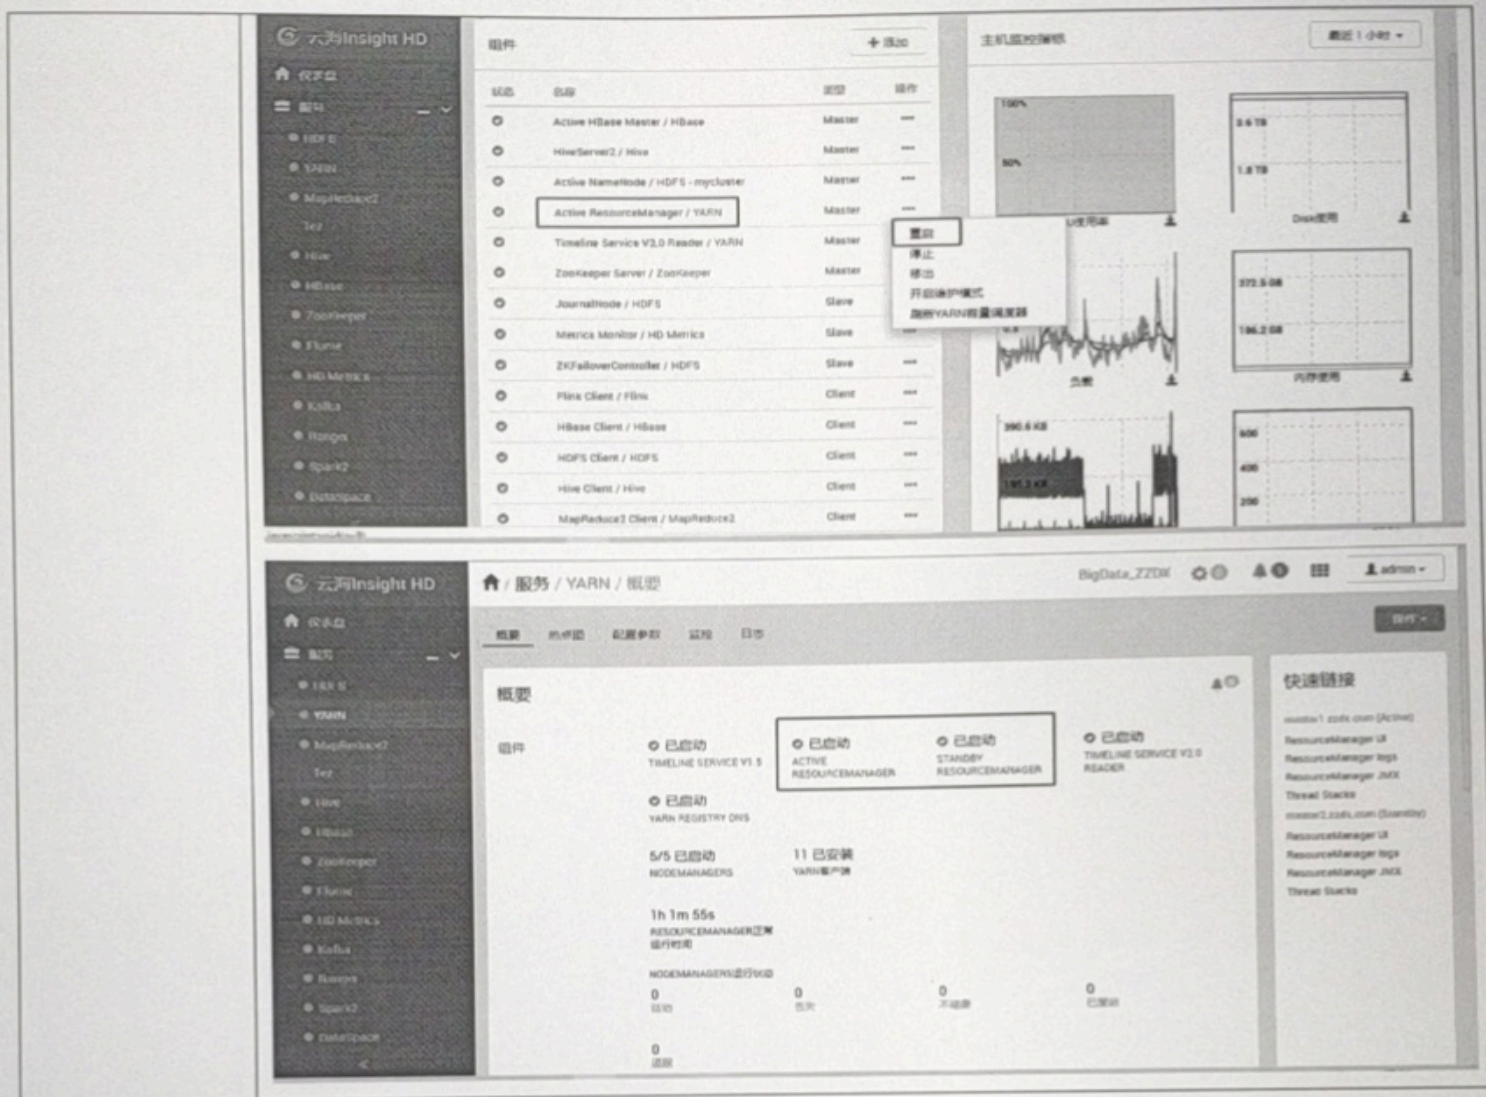
HBase Master JMX

Debug Dump

Thread Stacks

1.3.3 yarn的高可用测试

测试内容	yarn 高可用实现
预期结果	Resource Manager 失效及恢复测试
测试结果	符合预期，测试通过 <input checked="" type="checkbox"/> 不符合预期，测试不通过 <input type="checkbox"/> 不具备测试条件，技术论证 <input type="checkbox"/>
测试说明	<div>开启 HA</div> <div></div> <div>功能性测试</div> <div></div>



1.4 用户确认

测试结论	合格
测试人员签字	荆报
用户签字	任辉
日期	2021.5.6.