1 Introduction

Apache Ranger provides centralized security for Enterprise Hadoop ecosystem, including fine-grained access control and centralized auditing. In 0.5 version, Apache Ranger introduced stack-model to make it easier for new components to use Apache Ranger authorization and auditing. Further, to enable extending/adapting Apache Ranger for new or deployment-specific authorization requirements, the stack model provides hooks like context-enrichers and policy-conditions.

In this document, we will see the details of extending Apache Ranger to support authorization based on the location from which the access is performed i.e. country/state/city from which the resource is accessed.

Here is the outline of the tasks to be done:

- Prepare a location data file containing IP address to location details mappings
- Register a context-enricher hook that adds the location details to the request
- Register a policy-condition to verify that the client location matches the locations specified in the policy
- Create/update Apache Ranger policies to specify the locations to allow/deny the access

2 IP location data file

IP location data file is a text file containing comma separated fields. Each line in the file contain the location details for a range of IP addresses. The format of the IP location data file is given below:

- Each line consists of comma separated fields
- The first line is treated as a header, containing names for each field
- Subsequent lines have location details for a range of IP addresses
- First field is the start IP-address of the range
- Second field is the end IP-address of the range
- Other fields have the location data for the IP-range specified in first two fields (inclusive)
- IP-addresses should be specified as long integers; but the context-enricher can read addresses in dot-notation when `IPInDotFormat` for the client IP-address `true`
- Example:

```
IP_FROM,IP_TO,COUNTRY_CODE,COUNTRY_NAME,REGION,CITY
"10.0.0.255","10.0.3.0","US","United States","California","Santa Clara"
"20.0.100.80","20.0.100.89","US","United States","Colorado","Broomfield"
"20.0.100.110","20.0.100.119","US","United States","Texas","Irving"
```

This data format is similar to commercially available data from providers like IP2Location. Depending upon the requirements, the data file can either be sourced from a commercial data provider (like IP2Location) or be created with the deployment specific details.

3 Register Context Enricher

When Apache Ranger plugin receives an authorization request, the request is passed through registered context-enrichers. The context enrichers have access to various request details – like user, resource accessed, access type, ip-address of the accessor, etc. The context enrichers can update the request context with additional information that can be used while evaluating Ranger policies.

Context enricher `RangerFileBasedGeolocationProvider`, available in tag-policy branch, adds geo-location data to the request context, based on the location details available in a data file. To register the context-enricher for a component (like HDFS/Hive/HBase/..), please update the component's service-def by including the following:

```
 "contextEnrichers": [
   {
     "enricher":
"org.apache.ranger.plugin.contextenricher.RangerFileBasedGeolocationProvider"
,
     "enricherOptions": {
       "FilePath": "/etc/ranger/geo/geo.txt",
       "IPInDotFormat": "true"
     }
   }
 ]
```

Please ensure that the data file is available to the components at the location specified in the above registration (`/etc/ranger/geo/geo.txt`, in this example).

When `RangerFileBasedGeolocationProvider` receives an authorization request, it locates the record in the IP location data for the client IP address specified in the request. If a record is found, each field in the record will be added to the request context.

Following example should help understand the details of context data get added by the enricher:

- Client IP address: `20.0.100.85`
- Matching record in IP location data:
  - `20.0.100.80,20.0.100.89,"US","United States","Colorado","Broomfield"`
- IP location data header:
  - `IP_FROM,IP_TO,COUNTRY_CODE,COUNTRY_NAME,REGION,CITY`
- Entries added to the request context:
  - `LOCATION_COUNTRY_CODE=US`
  - `LOCATION_COUNTRY_NAME=United States`
  - `LOCATION_REGION=Colorado`
  - `LOCATION_CITY=Broomfield`

Please note that the name of context entries will be the field name, prefixed with "`LOCATION_`".

4 Register Policy Condition

Apache Ranger provides policy-condition hooks to execute custom conditions while evaluating authorization requests. To determine the authorization result, Apache Ranger policy engine evaluates the policies that are applicable to the accessed resource. Only when various criteria like user/group, access-type, and policy-condition specified in the policy matches the request, the policy engine will use the policy to determine the result.

Policy condition `RangerContextAttributeValueNotInCondition`, available in tag-policy branch, returns true only when the specified request context value does not match the values specified in the policy. This can be used to check if the location in request context (which is populated by the context enricher detailed earlier) is outside the values specified in the policy - for example, to deny access to requests that originate outside of specified countries. To register the policy-condition for a component (like HDFS/Hive/HBase/..), please update the component's service-def by including the following:

```
  "policyConditions": [
    {
      "name": "location-outside",
      "label": "Accessed from outside of location?",
      "description": "Accessed from outside of location?",
      "evaluator":
"org.apache.ranger.plugin.conditionevaluator.RangerContextAttributeValueNotIn
Condition",
      "evaluatorOptions": {
        "attributeName": "LOCATION_COUNTRY_CODE"
      },
    }
  ]
```

Once this policy condition is registered with Ranger, the policy editing UI will prompt for condition values to be used during evaluation – as shown below:

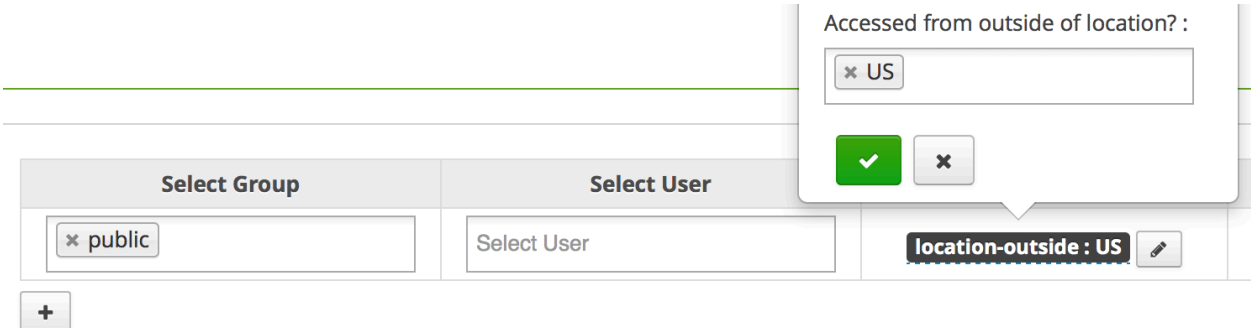

5 Example: Deny access to a Hive table from outside a specific country

In this section, we will see the details of an Apache Ranger policy that denies access to a specific Hive table when accessed from a location outside of a specific country. This example uses the context-enricher and policy-condition described earlier in this document.