Configuration

Configuration

Node.xml configuration properties

Node.xml contains miscellaneous node-wide configuration information, such as the node's role in the network. The configuration identifies if the node performs the following tasks:

- Authenticates users
- Broadcasts queries
- Aggregates responses
- Performs queries
- Routes incoming and outgoing messages

The order of the elements in the configuration file is important. Properties are required or optional.

Required Node Configurat n Propertie	Description s	Default	
<version></version>	Describes the human-readable version of SPIN on which this node is running. (This field is deprecated and will go away in future SPIN versions.)		
<nodenam< th=""><th>Describes the human-readable node nam</th><th></th><th></th></nodenam<>	Describes the human-readable node nam		
Optional elements	If these elements are omitted, the defaults apply.	Default:	
<isauthent ator></isauthent 	c Can this node authenticate users?	False	
<isbroadca er></isbroadca 	st Should this node propagate the queries it receives to other nodes.	True	
<isaggrega or></isaggrega 	t Identifies if this node should aggregate responses from other nodes.	False, if the node is not the root of a heierarchy, otherwise true.	
<isqueryal< th=""><th>Identifies if this node should respond to queries.</th><th colspan="2">True. Set to false if the node should only broadcast queries and aggregate results, for example.</th></isqueryal<>	Identifies if this node should respond to queries.	True. Set to false if the node should only broadcast queries and aggregate results, for example.	
<identityse iceClass></identityse 	rv Identifies the fully-qualified name of the JVM class used to authenticate users based on credentials. Only required if <isauthenticator> is true. The supplied class, if present, must implement org.spin.query.message.identity. IdentityService.</isauthenticator>	Null	
<queries></queries>	dentifies the queries, other than the basic queries, to which nodes respond. Each maps a queryType string to the fully-qualified name of a JVM class that implements that query. For each <queries> element, the <querytype> and <classname> sub-elements are required. May be present zero or more times. Note that all nodes respond to the following queries, at a minimum: Spin.Echo Spin.Audit Spin.Discovery</classname></querytype></queries>	no additional queries	
<certificati TTL></certificati 	When clients query a SPIN network, the queries must be digitally signed. Nodes examine queries when they are received and respond only if the digital signature is valid and not expired. The <certificationttl> element describes, in milliseconds, how old a signature can be before being considered invalid due to expiration.</certificationttl>	3600000, or one hour)	
<cachettl< th=""><th>In some circumstances, SPIN caches the results of queries to allow clients to retrieve them at a later time. However, results that are unclaimed after some period are purged automatically. Specified, in milliseconds, by the <cachettl> element.</cachettl></th><th>3600000, or one hour</th><th></th></cachettl<>	In some circumstances, SPIN caches the results of queries to allow clients to retrieve them at a later time. However, results that are unclaimed after some period are purged automatically. Specified, in milliseconds, by the <cachettl> element.</cachettl>	3600000, or one hour	
<resultsto reType></resultsto 	The type of underlying result store used to cache query results. Allowed values are 'SimpleInMemory' and 'Ehcache'. The SimpleInMemory result store is generally faster, but fails un-gracefully when memory runs out. The EHCache result store is generally slower, but fails more gracefully.	SimpleInMemory	
<queryacti onMapCla ssName></queryacti 	Defines the fully-qualified name of a JVM class that implements org.spin.node.QueryActionMap and provides QueryAction instances by name. Useful for creating SPIN using Spring, Guice, or another dependency-injection framework, or for instantiating QueryActions lazily, instead of eagerly at node-startup time, which is the default behavior.	Null. Eagerly instantate QueryActions defined by <queries> elements)</queries>	
<broadcas tTimeoutP eriod></broadcas 	SPIN nodes attempt to broadcast queries they receive to other nodes defined as being <i>children of</i> , or <i>downstream from</i> the current node. A timeout period for these broadcast attempts is defined by the broadcastTimeoutPeriod> element. The <duration></duration> sub-element defines the amount of the timeout. The <untycert< b=""> sub element defines the unit for the value of the <duration> field. enum constants in java.util.concurrent.TimeUnit define values for this field. Note the following examples:</duration></untycert<>		 stTimeou tPeriod>

<duration>5</duration>

<unit>SECONDS</unit>

</broadcastTimeoutPeriod> <broadcastTimeoutPeriod>

<duration>1</duration>

<unit>HOURS</unit>

</broadcastTimeoutPeriod> <broadcastTimeoutPeriod>

<duration>500</duration>

<unit>MILLISECONDS</unit>

</broadcastTimeoutPeriod> | 5 seconds |

routingtable.xml configuration properties

A SPIN node's place in a network topology is defined in routingtable.xml. SPIN nodes may participate in different overlay networks, fulfilling different roles in a topology, by joining different peer groups. A peer group defines a group of nodes and how they fit into a network topology. In routingtable.xml, nodes specify their place in different network topologies for zero or more peer groups.

Nodes always belong to the peer group, LOCAL, which includes only the current node. This is useful for querying only one node, without triggering query propagation.

When clients query a SPIN network, they send, along with their query, the peer group for which that the query is intended. Nodes use that peer group to look up how to propagate the query and where to return the results.

For a given peer group, nodes may have zero or one parents, and zero or more children. If a parent node is specified, query results from the current node, and those downstream, will be sent to the parent for aggregation. If no parent is specified, results are stored at the current node until the client retrieves them. It is also possible for a client to specify an arbitrary node as the point at which results will be aggregated and stored for retrieval.

If any child nodes are specified, the current node broadcasts all queries that it receives to the child nodes. Note that node, node B, can be listed as the child of node A, but node B does not have to list node A as its parent.

routingtable.xml contains zero or more <peerGroup> elements. The structure of these elements is as follows:

Property	Description	Optional /Required
peerGroup	LOCAL peer group named with no parent or children added implicitly	optional
<groupname></groupname>	Unique name of the peer group	required
<parent></parent>	The current node's parent (default null - no parent)	optional
<children></children>	Zero or more child nodes that to which the current node should broadcast queries, if it receives a query destined for this peer group. The structure of <parent> and <children> elements are identical.</children></parent>	
<endpointtype></endpointtype>	Values: Only supports SOAP. (Local used only for testing.)	required
<address></address>	Identifies the URL of the referenced node's WSDL, such as http://my-node.med.harvard.edu:8080/node-server/node?wsdl	require

keystore.xml configuration properties

A keystore is a repository of security certificates, either Certification Authority Certificates or Public key certificates. Keystore.xml contains information about the keystore file that contains the node's cryptographic keypair. The keypair includes the keystore file password and the keystore private key alias that you should use, if more than one private key exists.

The serial number of the node private key identifies the SPIN node.

The serial number of their private key supports the PKCS12 and JKS keystore formats.

Nodes are definitively identified on a network by the serial number of their private key. This file contains information about the keystore file, which contains the Node's cryptographic keypair, including the password of the keystore file and the alias of the private key in the keystore to use---if there is more than one.

The keystore file referenced by keystore.xml will also contain the certificates of authenticating and querying agents that the Node trusts.

PKCS12 and JKS keystore formats are supported

You can locate keystore.xml at http://spin.org/xml/res/keystore. Note that the order of elements in this file is important.

Required Elements	Description	Default Value
<file></file>	The file containing the keystore <file>/opt/spin/main/conf/spin.keystore</file> Path, either absolute or relative, to the keystore file containing cryptographic certificates.	

<password></password>	Password for the file referenced by the <file> element.<password>spinkeystore</password></file>	
Optional elements	The alias of the node's public/private keypair is optional. alias of the keypair is used. If no keyAlias is specified but multiple pub/priv keypairs exist, PKITool's constructor throws an exception.	
<keyalias></keyalias>	Identifies the alias of the certificate containing this node's private key. If more than one public/private keypair is contained in the referenced keystore, this element can be used to choose one. If only one public/private keypair exists in the referenced keystore, Spin choses that keypair, and this element is optional.	Null. Node tries to detect the public /private keypair.
<capublicke yAlias></capublicke 	May be present 0 or more times. The aliases of the public keys of certificate authorities that this node trusts. If a node trusts a particular certificate authority by importing the CA's public cert into the node's keystore and referencing it with the <capublickeyalias> element, the node will implicitly trust all queriers for which the CA vouches.</capublickeyalias>	Null. No CAs are trusted.
<attachcerti ficateToSign ature></attachcerti 	Relevant only if a <capublickeyalias> element (or elements) is present. If set to true, then the public cert of the CA that vouched for a querier is attached to outbound queries.</capublickeyalias>	False
<setsystem Properties></setsystem 	Work around some failures with establishing SSL connections by modifying JVM-wide settings. Set to true as a last resort when troubleshooting SSL/HTTPS communication. This element is deprecated and will likely go away in a future release of SPIN.	False
<keystorety pe></keystorety 	The format of the referenced keystore file; See org.spin.tools.config.KeyStoreType for allowed values (currently 'JKS' and 'PKCS12') Example keystore.xml: http://scm.chip.org/svn/repos/spin/base/trunk/examples/src/test/resources/keystore.xml	JKS (Default) or PKCS12.