Distributed Systems Unassessed Coursework: Publish-Subscribe

Lecturer: M. Sloman

1. Compare pub-sub interaction with multicast communication
2. Briefly describe Topic, content and type based subscription
3. In a large-scale system when is it better to broadcast subscriptions to all brokers and when should publications be broadcast to all brokers?
4. Describe the Siena event Routing system

Distributed Systems Coursework Solution: Publish-Subscribe

Lecturer: M. Sloman

1. Multicast communication uses a multicast address in a message on a LAN which identifies a set of hosts to which the message will go. All Hosts which are part of members of the multicast group address, receive the message and decide how to process it. On the internet, the sender has to know the set of members of the multicast group and send unicast messages to each of them.

   In a pub-sub system, consumers subscribe to a message based on a predicate related to the message contents. Publishers generate events without knowing who the subscribers are. In general subscribers may not know who the publishers are and multiple publishers can generate events matched by a subscription. Brokers perform the matching between publishers and subscribers and hence there are at least 2 hops for and event to go from a publisher to subscriber. Sophisticated routing can be used to prevent multiple copies of events traversing the same path when destined for different subscribers.

2. Topics define a subject relevant to events and are typically arranged in a hierarchy so a subscriber to Topic T1.1 gets event related to topics T1.1.1, T1.1.2, T1.1.3 etc. as well as any topic below these.

   Content based subscription defines a predicate which applies to all fields in a message so allows for much finer granularity in selecting which events are required but requires more processing in brokers to match subscriptions with published events.

   Type based subscription users a Type language for defining events. A subscription refers to a subtree root and all subtypes in the tree are delivered to the subscription. Content filtering can be used on the attributes of the type. This is similar to a combination of topic and content subscription.

3. Broadcast subscriptions to all brokers can be used in a system with large numbers of publishers and comparatively few subscriptions. Each publisher will send events to its local broker which will be able to match with subscribers and route the event appropriately.

   If the system has large number of subscriber but comparatively few publishers then it is more efficient to subscribe to your local broker which stores the subscriptions and performs matches against all events. The publisher's broker, broadcasts the event to all other brokers.

4. Advertisements about events that will be published are propagated to form an advertisement tree rooted at the publisher's broker. When a subscription arrives at a node that is a member of the tree, it sends it up the tree so avoiding broadcasting it to all neighbouring nodes. Events from a publisher are sent down the tree following the reverse path set up by matching subscriptions. When a broker B1 forwards a subscription s to B2, B2 stores s and notes that s came from B. When B2 receives an event e that matches s, it forwards e to B1, so events are only forwarded to brokers that have subscribed to receive them. When a new subscription s' is received, by a broker, it checks whether s' is subsumed by (i.e. is a subset of) previous subscriptions, and if so s' is not forwarded up the tree as the broker is already receiving events which match s'.