Security Architectures and Algorithms for Publish-Subscribe Network Services

Mudhakar Srivatsa, James Caverlee and Ling Liu
College of Computing
Georgia Institute of Technology
{mudhakar, caverlee, lingliu}@cc.gatech.edu

Publish-Subscribe Services. A large number of emerging Internet applications requires information dissemination across different organizational boundaries, heterogeneous platforms, and a large, dynamic population of publishers and subscribers. A publish-subscribe (pub-sub) network service is a wide-area communication infrastructure that enables information dissemination across geographically scattered and potentially unlimited number of publishers and subscribers. A wide-area pub-sub system is often implemented as a collection of spatially disparate nodes communicating on top of a peer-to-peer overlay network. In such an environment, publishers publish information in the form of events and subscribers have the ability to express their interests in an event or a pattern of events by sending subscription filters to the pub-sub network. The pub-sub network uses content-based routing schemes to dynamically match each publication against all active subscriptions, and notifies the subscribers of an event if and only if the event matches their registered interest.

Publish-Subscribe Service Model. A pub-sub network service model allows an organization to outsource its physical resource management problems to a third-party pub-sub network. However, the ownership on published events still lies in the hands of the publisher. In essence, the pub-sub network service model separates resource management from ownership and access control. For example, a pub-sub network provides efficient and scalable delivery of events from a publisher to one or more subscribers (resource management). However, the publisher owns the content of a published event and is responsible for defining access control over the event (ownership and access control). The publishers may wish that the events are kept confidential from the pub-sub network nodes. Access control on a published event restricts the set of subscribers who are authorized to read a given event.

Security Issues. An important characteristic of pub-sub network services is the decoupling of publishers and subscribers combined with content-based routing protocols, enabling a many-to-many communication model. Such a model presents many inherent benefits as well as potential risks. On one hand, offloading the information dissemination task to the pub-sub network not only improves the scalability and the effectiveness of the pub-sub system, but also permits dynamic and fine-grained subscriptions. On the other hand, a pub-sub network model faces several security threats such as: denial of service (DoS) & host compromise attacks, authenticity, confidentiality and integrity of application data, and key distribution & management.

Denial of Service (DoS) Attacks. The pub-sub network service has to protect the application data routed by the pub-sub nodes from DoS and host compromise attacks. Protecting the pub-sub nodes from DoS and host compromise attacks improves service availability. In a pub-sub network service model, DoS attacks can target three different layers: (i) TCP/IP layer, (ii) pub-sub network layer, and (iii) application layer. The pub-sub network service has to develop solutions to mitigate *insider* DoS attacks, wherein a set of malicious pub-sub nodes attempt to launch a DoS attack on the applications hosted by the pub-sub network.

Authenticity Attacks. The pub-sub network service has to protect the applications data hosted by the pub-sub nodes from incorrect or fake (spoofed) application data. Protecting the pub-sub network nodes from incorrect or fake

application data guarantees the authenticity of application data hosted by the nodes. In a pub-sub network service model, authenticity attacks can be of two types: (i) an adversary may attempt to spoof the identity of a legitimate publisher and send incorrect or fake application data to the pub-sub network nodes, and (ii) an authentic publisher may flood the pub-sub network nodes with incorrect or inaccurate application data. The latter problem is prevalent in today's Internet wherein, we have multiple competitive web servers (with possibly conflicting interests) publish doctored information.

Confidentiality and Integrity Attacks. The pub-sub network service model has to protect the confidentiality and integrity from: (i) the pub-sub network nodes, and (ii) unauthorized users. The publisher may not trust the pub-sub network service with the confidentiality and integrity of the application data. The malicious pub-sub network nodes may be able to eavesdrop or corrupt the application data routed by them. In addition, malicious pub-sub nodes may collude with one another in their attempts to compromise the confidentiality and integrity of application data. The pub-sub network service model allows the publisher to specify access control rules on application data. These access control rules restrict the set subscribers that can access a given piece of application data hosted by the pub-sub network. However, malicious subscribers may be curious to access application data and services that they are not authorized to access. In addition, malicious subscribers may collude with one another and with the malicious nodes in the pub-sub network to compromise the confidentiality and integrity of application data.

Key Distribution and Management. A pub-sub network service model is faced with the challenge of having to meet the above security threats while preserving the performance and scalability of the application. Using cryptographic primitives to mitigate these security threats opens up new performance and scalability problems. Most cryptographic primitives assume an out-of-band distribution and management of cryptographic keys. In the pub-sub network service model, key distribution and management becomes a critical problem especially since the pub-sub network service typically employs tens of thousands of pub-sub nodes. Further, nodes can fail and leave the pub-sub network at a non-trivial rate; similarly, failed nodes can recover and join the pub-sub network at a non-trivial rate. Hence, the pub-sub network service model needs secure, efficient, and scalable key dissemination algorithms to handle a dynamic population of the pub-sub nodes, the publishers, and the subscribers.

Contributions. We have developed SGuard − a security architecture and a set of algorithms to secure wide-area pub-sub network services. Our design has been guided by the following two principles: (i) Cryptographic techniques need to be adapted using application specific knowledge in order to secure an application without compromising on its performance and scalability metrics. (ii) Using intrinsic properties such as the structure of the pub-sub network and the semantics of the application leads to powerful and effective security algorithms.

SGuard aims at developing a suite of security guards to: (i) protect the interfaces exported by the pub-sub network from denial of service (DoS) and host compromise attacks, (ii) protect the authenticity, confidentiality and integrity of application data as desired by the publisher, (iii) provide a secure key distribution & management algorithm for managing up to tens of thousands of pub-sub network nodes, and (iv) preserve the performance and scalability of the pub-sub network while meeting requirements (i), (ii) and (iii). SGuard comprises of a suite of security guards that can be seamlessly plugged into a pub-sub network service. We have also built prototype implementations of several security guards to show that SGuard is easily stackable on a pub-sub network service. Our experimental results so far indicate that secure a pub-sub network service while preserving its performance and scalability metrics.

Summary. In summary, the autonomous nature of the pub-sub network service model is very similar to that of the Internet itself, allowing multiple publishers to efficiently publish data and deliver services to a large population of geographically scattered subscribers. We believe that developing secure, efficient, and scalable techniques to guard pub-sub network services plays a very crucial role in making these services widely deployable.