Secure routing for structured peer-to-peer overlay networks

M. Castro, P. Druschel, A. Ganesch, A. Rowstron, D.S. Wallach
5th Unix Symposium on Operating Systems Design and Implementation (OSDI), December 2002

Seminar of Distributed Computing
Anna Wojtas

Security in Peer-to-Peer networks

- Peer-to-Peer networks are meant to be open and autonomous
  - →availability
  - →authenticity of documents
  - →anonymity
  - →access control
- Possible attacks:
  - →denial of service
  - →poisoning attack
  - →insertion of viruses to carried data

Agenda

- Definition: Overlay network
- Motivation
- Model
- Secure node assignment
- Secure routing table maintenance
- Secure message forwarding
- Self-certifying data
- Conclusions

Definition: Overlay network

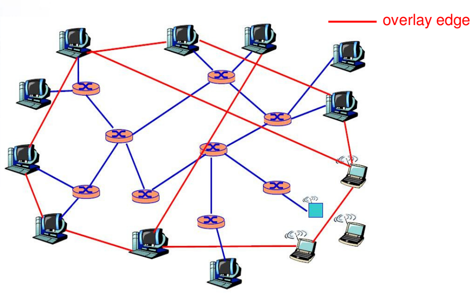

Agenda

- Definition: Overlay network
- **Motivation**
- Model
- Secure node assignment
- Secure routing table maintenance
- Secure message forwarding
- Self-certifying data
- Conclusions

Motivation

Status quo (2002):
- self-organizing
- scalable
- fault-tolerant
- provide effective load balancing

Support for open environments:
- robustness against malicious nodes

Agenda

- Definition: Overlay network
- Motivation
- **Model**
- Secure node assignment
- Secure routing table maintenance
- Secure message forwarding
- Self-certifying data
- Conclusions

Model: routing overlay

- Large Id space (128-bit)
- Node identifiers → nodeIds
- Application-specific objects → keys
- Mapping key x nodeId → key's root
- nodeIds x IP addresses → routing table
- Closest nodeIds → neighbor set
- Key → replica keys → replica roots → replica function

Model: Pastry

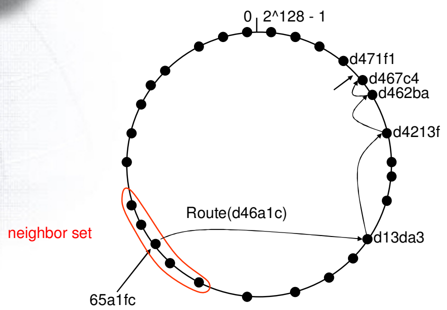

Pastry cont.

0x	1x	2x	3x	4x	5x		7x	8x	9x	ax	bx	cx	dx	ex	fx	6x
60x	61x	62x	63x	64x		66x	67x	68x	69x	6ax	6bx	6cx	6dx	6ex	6fx	65x
650x	651x	652x	653x	654x	655x	656x	657x	658x	659x		65bx	65cx	65dx	65ex	65fx	65ax
65a0x		65a2x	65a3x	65a4x	65a5x	65a6x	65a7x	65a8x	65a9x	65aax	65abx	65acx	65adx	65aex	65afx	65a1x

nodeId 65a1x

Model: system

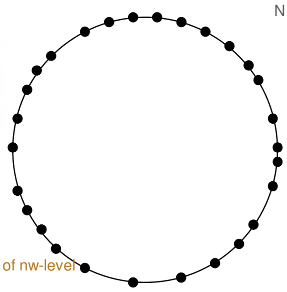

static IP

Communication:
- network-level,
- overlay-level

Adversary:
- complete control of nw-level communication
- delay messages between correct nodes

Model: secure routing

Routing primitive:
- best-effort service to deliver a message to a replica root associated with a given key

Cannot be used to construct secure applications:
- corrupt, delete, deny access to or supply stale copies of replicas

Model: secure routing cont.

Secure routing primitive:

- ensures that when a non-faulty node sends a message to a key k, the message reaches all non-faulty members in the set of replica roots with a very high probability

Requires solution for:

- securely assigning nodeIds to nodes
- securely maintaining the routing tables
- securely forwarding messages

Agenda

- Definition: Overlay network
- Motivation
- Model
- Secure node assignment
- Secure routing table maintenance
- Secure message forwarding
- Self-certifying data
- Conclusions

Secure node assignment

Attacks:

- network partitioning
- DoS on single nodes / objects

→Attacker cannot choose the value of the nodeId assigned to the node she controls

Solution:

- certified nodeIds

Secure assignment cont.

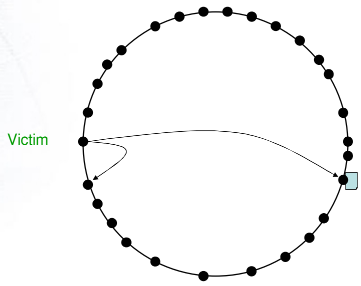

→Victim's access to the overlay completely mediated by the attacker

→Control of other nodes accessing a victim's file

Secure assignment cont.

More attacks:

- delete, corrupt or deny access to objects

→ attacker cannot choose the value of the nodeId assigned to the node she controls

Solution:

- certified nodeIds

Secure node assignment

Certified nodeIds:

- CAs assign nodeId certificates
- binding of a random nodeId to the public key for a IP address → nodeId swapping attacks harder

→only for static IP addresses

→works well only for fixed nodeIds

→doesn't solve all problems…

Secure assignment cont.

Sybil attacks:

- peer impersonates multiple virtual peers
- →destroy cohesion of the overlay
- →observe network status
- →slow down, destroy overlay
- →DoS

→attacker cannot easily obtain a large number of nodeId certificates

Secure assignment cont.

Solution:

- pay for certificates
  - cost $20, controlling 10% of
    - 1000 nodes → $2,000
    - 1,000,000 nodes → $2,000,000
- bind nodeIds to real-world identities
  - for overlays run by an organization

Secure assignment cont.

Distributed nodeId generation:

- CA is point of failure
- techniques to moderate the rate at which attackers can acquire nodeIds
- crypto puzzles

Agenda

- Definition: Overlay network
- Motivation
- Model
- Secure node assignment
- Secure routing table maintenance
- Secure message forwarding
- Self-certifying data
- Conclusions

Secure routing table maintenance

Goal:

- create routing table, neighbor sets for joining nodes and maintaining them
- secure nodeId assignment necessary but not sufficient

→ Attacks…

Secure routing table cont.

- Routing algorithms using network proximity information:

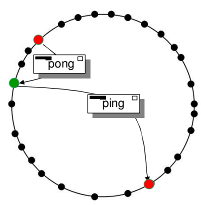

→ Increased probability that faulty nodes are used for routing

Secure routing table cont.

- Systems with weak constraints on routing updates
  - updates received during joining
  - periodical fetch of routing table entries

→attackers can easily supply updates pointing to faulty nodes
  - probability of routing table entry is faulty after update (1-f)*f +f*1 > f
  - fraction of faulty entries → 1

Secure routing table cont.

Theoretical solution:

- strong constraints on the set of nodeIds that can fill each slot of the routing table
- e.g. closest nodeId to some point in id space

→can be verified

→independent of network proximity information

Secure routing table cont.

Practical solution (Pastry):

- 2 routing tables
- locality-aware routing table exploits network proximity information for efficient routing
  - used to forward messages to achieve good performance
  - prefix D whatever
- additional table constraints routing table entries
  - used when the efficient routing technique fails
  - prefix D suffix

Agenda

- Definition: Overlay network
- Motivation
- Model
- Secure node assignment
- Secure routing table maintenance
- Secure message forwarding
- Self-certifying data
- Conclusions

Secure message forwarding

- Certified IDs & secure routing table maintenance

→guarantees that each constraint routing table has an average fraction f of entries pointing to faulty nodes

→attacker can reduce probability of successful delivery by not forwarding according to the algorithm

Secure message forwarding cont.

Attacks:

- drop the message
- route the message to the wrong place
- pretend to be the key's root

→Probability of routing successfully to a replica root is $(1-f)^h$

- h is the number of average hops for delivering a message
- h depends on the overlay

Secure message forwarding cont.

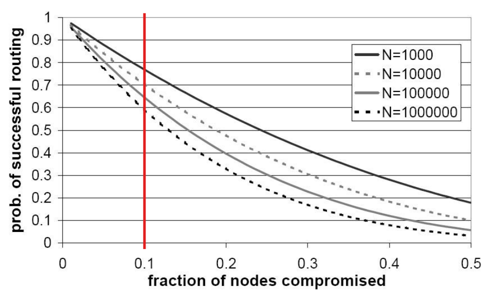

→ it is important to have a mechanism to route securely

Secure message forwarding cont.

Theoretical solution:

- route a message efficiently
- apply failure test to determine if routing has worked
- upon failure of the test use redundant routing

Secure message forwarding cont.

Practical solution (Pastry):

- use locality-aware routing table for efficient routing
- collect the prospective set of replica roots from the prospective root node
- apply failure test to the set
- if test negative, accept the replica roots as correct
- if test positive, send message copies over diverse routes towards various replica roots

Secure message forwarding cont.

Failure test:

- average density of nodeIds per unit of "volume" in the id space is greater than the average density of faulty nodes

→compare densities

replica roots = subset of key's root neighbor set

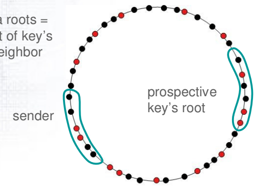

μ_{sender}
average numerical distance between consecutive nodes in sender's neighbor set

$rn = id_0, \ldots, id_{l+1}$
prospective root neighbor set

μ_{rn}
average numerical distance between consecutive nodes in rn

Test:

- all nodes in rn have a valid nodeId certificate
- $\mu_{rn} < \mu_{sender} * \gamma$

Secure message forwarding cont.

Problems

- false positives (α), false negatives (β)

→ γ controls tradeoff between α and β

Attacker can

- collect nodeId certificates of node that have left the overlay

→ increase density of a prospective root neighbor set

- include nodeId it controls and nodeIds of correct nodes

Solution

- sender has to contact all neighbors to find out if they are alive and have the same nodeId certificate

Secure message forwarding cont.

NodeId suppression attack

- suppress nodeIds close to the sender

→increase false negatives (β)

- suppress nodeIds in the root's neighbor set

→increases false positives (α)

- combination of both

→routing test is not very accurate

→tradeoff increased α to achieve targeted β

→β=0.001, c=f ≤ 0.3 → α_{no_attack}=0.12, α_{attack}=0.77

Secure message forwarding cont.

Redundant routing

- use multiple routes
- neighbor set anycast

probability of reaching all correct replica roots ~ probability that at least one of the anycast messages is forwarded over a route with no faults

for 100,000 nodes, $l = 32 \rightarrow 0.999$

s collects in a set N $l/2+1$ numerically closest to x on the left and on the right

only certificates with valid signed nonces are added to N and marked pending

after timeout or after all replies received, s sends a list with nodeIds in N to each node marked pending in N and marks the nodes done

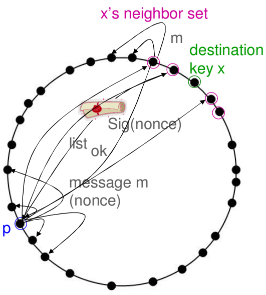

Agenda

- Definition: Overlay network
- Motivation
- Model
- Secure node assignment
- Secure routing table maintenance
- Secure message forwarding
- Self-certifying data
- Conclusions

Self-certifying data

- minimize use of secure routing by storing self-certifying data in the overlay
- clients use efficient routing to request a copy of an object
- client performs integrity check and use secure routing only upon failure
- does not help when inserting new objects
- node joining requires secure routing

→self-certifying data can eliminate the overhead of secure routing in common cases

Conclusions

- The authors analyzed various approaches for the problems
- Weak performance evaluation
- Paper cited in ~40 other papers

Questions?