

PeerReview: Practical Accountability for Distributed Systems

Andreas Haeberlen^{†‡}, Petr Kouznetsov[†], and Peter Druschel[†]
[†]Max Planck Institute for Software Systems, [‡]Rice University

ABSTRACT

We describe PeerReview, a system that provides accountability in distributed systems. PeerReview ensures that Byzantine faults whose effects are observed by a correct node are eventually detected and irrefutably linked to a faulty node. At the same time, PeerReview ensures that a correct node can always defend itself against false accusations. These guarantees are particularly important for systems that span multiple administrative domains, which may not trust each other.

PeerReview works by maintaining a secure record of the messages sent and received by each node. The record is used to automatically detect when a node's behavior deviates from that of a given reference implementation, thus exposing faulty nodes. PeerReview is widely applicable: it only requires that a correct node's actions are deterministic, that nodes can sign messages, and that each node is periodically checked by a correct node. We demonstrate that PeerReview is practical by applying it to three different types of distributed systems: a network filesystem, a peer-to-peer system, and an overlay multicast system.

Categories and Subject Descriptors

C.2.4 [Computer Systems Organization]: Computer-Communication Networks—*Distributed Systems*; D.4.5 [Software]: Operating Systems—*Reliability*

General Terms

Algorithms, Design, Reliability, Security

Keywords

Accountability, fault detection, distributed systems, Byzantine faults

Permission to make digital or hard copies of all or part of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. To copy otherwise, to republish, to post on servers or to redistribute to lists, requires prior specific permission and/or a fee.

SOSP'07, October 14–17, 2007, Stevenson, Washington, USA.
Copyright 2007 ACM 978-1-59593-591-5/07/0010 ...\$5.00.

1. INTRODUCTION

Nodes in distributed systems can fail for many reasons: a node can suffer a hardware or software failure; an attacker can compromise a node; or a node's operator can deliberately tamper with its software. Moreover, faulty nodes are not uncommon [44]. At large scale, it is increasingly likely that some nodes are accidentally misconfigured or have been compromised as a result of unpatched security vulnerabilities.

In systems that span multiple administrative domains, the lack of central administration tends to aggravate these problems. Moreover, multiple trust domains pose the additional threat of deliberate manipulation by node operators with different interests. Examples of systems with multiple administrative domains are network services such as DNS, NTP, NNTP and SMTP, federated information systems, computational Grids, Web services, peer-to-peer systems, and the Internet's inter-domain routing system.

We consider the use of *accountability* to detect and expose node faults in distributed systems [58]. For the purposes of this paper, an accountable system maintains a tamper-evident record that provides non-repudiable evidence of all nodes' actions. Based on this record, a faulty node whose observable behavior deviates from that of a correct node can be detected eventually. At the same time, a correct node can defend itself against any false accusations.

Clearly, accountability by itself is not sufficient for systems in which faults can have serious and irrecoverable effects, such as deleting all replicas of a file or dispensing cash from an automated teller machine. However, accountability offers several benefits, by itself and in combination with other techniques:

- **Deterring faults:** The mere presence of accountability can reduce the incidence of certain faults. For instance, the threat of exclusion may discourage freeloading, and the threat of public embarrassment may cause organizations to update and configure their software with increased diligence.

- **Detection in fault tolerant systems:** Accountability complements Byzantine fault tolerant techniques such as data or state machine replication [13]. Systems that use these techniques can tolerate a bounded number of faulty nodes, typically less than one-third of the system size. By enabling the timely recovery of faulty nodes, accountability can help these systems to maintain this bound, thus increasing their ability to withstand successive node faults.

- **Detection in best-effort systems:** Systems that provide best-effort services can naturally tolerate recoverable faults, because such faults “merely” degrade the system's service. However, too many unidentified faults can degrade

the system to the point of becoming unusable. Accountability enables recovery, which helps to maintain the system’s health.

- **Assigning blame:** In systems with multiple trust domains, accountability irrefutably pinpoints the party responsible for a problem, while allowing other principals to prove their innocence to customers, peering partners, and authorities.

This paper describes and evaluates PeerReview, a general and practical system that provides accountability for distributed systems. PeerReview creates a per-node secure log, which records the messages a node has sent and received, and the inputs and outputs of the application. Any node i can request the log of another node j and independently determine whether j has deviated from its expected behavior. To do this, i replays j ’s log using a reference implementation that defines j ’s expected behavior. By comparing the results of the replayed execution with those recorded in the log, PeerReview can detect Byzantine faults without requiring a formal specification of the system.

An accountable system requires strong node identities. Otherwise, an exposed faulty node could escape responsibility by assuming a different identity [21]. In this paper, we assume that each node is in possession of a cryptographic key pair that can be linked to a unique node identifier. We assume that correct nodes’ actions are deterministic and that communication between correct nodes eventually succeeds. These assumptions are common in the literature and reasonable in practice [3, 13, 30, 55].

PeerReview was designed for deployment in the Internet and is subject to certain limitations imposed by that environment. First, PeerReview cannot expose a faulty node that ignores some messages but never sends a message that a correct node would not send. The reason is that packet loss, or network or processing delays could make a correct node temporarily seem as though it were ignoring a message. With PeerReview, a node that does not respond to a message is eventually suspected by every correct node. A suspected node can exonerate itself by acknowledging the message. By restricting correct nodes to communicate with only non-suspected nodes, we can isolate a suspected node until it complies.

Second, PeerReview detects a fault after a correct node is causally affected by the fault, but not before. Since PeerReview does not assume the availability of trusted probes at each node, it can reliably observe only messages sent and received by correct nodes. Thus, PeerReview can detect only faults that manifest themselves through these messages. Detecting faults before they impact correct nodes would require PeerReview to make much stronger assumptions. We think that PeerReview strikes a reasonable balance between detection power and wide applicability in real-world systems.

We have applied PeerReview to three example systems: a network filesystem, a peer-to-peer email system, and an overlay multicast system. An experimental evaluation shows that PeerReview is applicable to a range of systems and that its overhead is reasonable in practice.

Specifically, signing messages imposes a fixed processing delay; this overhead is noticeable only in a LAN and when messages are short. The processing and message overheads depend on the fault assumptions and the desired strength of the detection guarantee. The processing overhead grows linearly with the number of nodes that need to inspect a given node’s actions to be sure at least one of the inspect-

ing nodes or the inspected node is correct. If we insist that every instance of misbehavior is eventually detected, then PeerReview’s message overhead grows with the square of the number of nodes in the system; experiments show that this limits PeerReview’s scalability to moderately sized systems of hundreds of nodes. However, if we settle for a probabilistic detection guarantee, then the message overhead becomes logarithmic in the number of nodes, and PeerReview can scale to much larger systems.

The rest of this paper is structured as follows. We discuss related work in Section 2. We provide a precise definition of the type of accountability and fault detection provided by PeerReview in Section 3. We describe the PeerReview algorithm in Section 4; our implementation in Section 5; and our three applications in Section 6. Section 7 presents the results of our evaluation, and Section 8 concludes this paper.

2. RELATED WORK

Accountability in distributed systems has been suggested as a means to achieve practical security [35], to create an incentive for cooperative behavior [20], to foster innovation and competition in the Internet [4, 36], and even as a general design goal for dependable networked systems [57]. However, it has been an open question whether accountability can be implemented in a general and efficient manner [58].

PeerReview offers strong accountability for any distributed system that can be modeled as a collection of deterministic state machines. To our knowledge, no prior work has achieved this level of generality. The type of accountability used in PeerReview, and a precise definition of the types of faults it can detect, has appeared in a prior workshop position paper [25]. This paper contributes the design and implementation of a practical system with these properties, as well as experimentally evaluates it with three different applications.

CATS [59] implements a network storage service with strong accountability properties. Like PeerReview, it maintains secure logs that record the messages sent and received by each node. Unlike PeerReview, however, CATS depends on a trusted publishing medium that ensures the integrity of these logs. CATS detects faults by checking logs against a set of rules that describes the correct behavior of a specific system (a network storage service). PeerReview does not require a formal specification of correct behavior; it verifies the nodes’ behaviors by replaying their logs using the system’s existing reference implementation.

The study of the Byzantine failure model originate in two papers by Lamport, Pease, and Shostak [34, 45]. State machine replication [33, 49] is a classic technique for masking a limited number of such Byzantine faults. Byzantine fault tolerance (BFT) protocols [13, 47, 56] can mask faults as long as less than one-third of the nodes are faulty [9]. Aiyer et al. [1] introduced the BAR model, which can tolerate a limited number of Byzantine nodes plus an unlimited number of “rational” nodes. PeerReview complements the existing work on BFT by providing strong accountability.

Generic simulations suggest a more general manner of handling Byzantine faults. Bracha [8] has described a protocol that hides the malicious effects of Byzantine faults by simulating more benign “identical Byzantine” faults on top of them. Simulations of even more restrictive classes of omission and crash faults in the Byzantine failure model were proposed by Srikanth and Toueg [52], Neiger and Toueg [43],

Coan [17], and Bazzi and Neiger [7]. Unlike PeerReview, these protocols are not intended to provide verifiable evidence of misbehavior. They are typically designed for broadcast-based algorithms and assume a synchronous system or a large fraction of correct nodes.

A trusted computing platform detects faults that involve modifications to a node’s software [23]. This approach requires special hardware features, a trusted OS kernel, and a software and hardware certification infrastructure.

The concept of failure detectors, abstract oracles that produce information about faults, was introduced by Chandra and Toueg [15] for the crash failure model. An extension of this work to a specific class of “muteness” failures was explored by Malkhi and Reiter [38] and Doudou et al. [22]. Kihlstrom et al. [31] proposed a consensus algorithm using a failure detector that produces suspicions of “detectable” misbehavior. Unlike the detector in PeerReview, [31] does not provide evidence of misbehavior and assumes algorithms that are based on reliable broadcast: every message is relayed to all nodes when received for the first time.

A technique for a statistical evaluation of the number of nodes with Byzantine faults was proposed by Alvisi et al. [2]. This technique is designed for specific replicated data services based on quorums and does not provide a means to identify which nodes are faulty.

A variety of techniques address specific types of misbehavior. These techniques include secure routing for structured overlays [11, 51], incentive mechanisms to prevent freeloading [18, 41], and content entanglement to prevent censorship [54]. These techniques can be more efficient than PeerReview, because they tend to be tailored to an application and a specific type of fault; however, they are difficult to reuse and offer no protection against unforeseen types of misbehavior. PeerReview offers a reusable, general detection mechanism for a large class of faults, including unanticipated faults.

Intrusion detection systems (IDS) can handle certain types of protocol violations [19, 27, 32]. However, they either are based on heuristics and require careful balancing between false positives and false negatives, or require a formal specification of the expected behavior, which can be difficult to write and maintain for a complex system. PeerReview avoids these problems by using a reference implementation of the system as an implicit specification.

Reputation systems such as EigenTrust [28] can be used to detect Byzantine faults, but they typically can detect only nodes that misbehave repeatedly. Also, a coalition of faulty nodes can denounce a correct node. PeerReview can detect even a single instance of detectable misbehavior, and it never exposes a correct node.

PeerReview’s secure logging and auditing techniques were inspired by secure timelines [39], as well as by earlier work on tamper-evident logs [50]. Secure histories have been used for other purposes [16]. The fork consistency model introduced with SUNDR [37] is similar to PeerReview’s model of observable behavior.

3. DETECTING FAULTY BEHAVIOR

Making a distributed system accountable involves two steps: the first creates a secure record of all nodes’ actions, the second inspects the recorded information and detects faulty behavior. In this section, we consider the problem of fault detection and discuss some of its limitations.

3.1 The problem of detection

We consider a distributed system of *nodes* in which each node is supposed to follow a given *protocol*. If it does, then we call the node *correct*; otherwise, we call the node *faulty*.

To illustrate the approach, we consider a simple toy protocol (Figure 1a). In this protocol, each node is responsible for providing and allocating 10 units of a resource, e.g. storage space. At any time, a node (client) may request resources from another node (server) by sending a message `REQUEST_k`. If the server has sufficient resources, it must allocate k units of the resource for the client and return a message `GRANT_k`. When the client is done with the resource, it sends a `RELEASE_k` to the server so that the server can release the k units and grant other requests.

There are two ways in which a node can be faulty in this protocol: it can either send a message that a correct node would never have sent in a given state (Figures 1c and 1d) or ignore a message a correct node would have accepted (Figure 1e).

Ideally, if a node misbehaves in either of these ways, each correct node should detect the misbehavior (we say that the correct node *exposes* the faulty node). An *ideal* fault detector should thus guarantee

- **Ideal completeness:** Whenever a node becomes faulty, it should be exposed by all correct nodes.

However, malicious nodes may try to disrupt the system by tricking some correct nodes into exposing other correct nodes. To prevent this, an ideal detector should also guarantee

- **Ideal accuracy:** No correct node is ever exposed by a correct node.

Of course, these idealistic requirements still lack important details. For example, it is often difficult to decide whether a message is being ignored or just delayed. It may also be difficult to identify misbehavior that cannot be observed by any correct node. Before we refine our requirements (Section 3.4), we consider some practical limitations of distributed detection mechanisms.

3.2 What can (or cannot) be detected?

Naturally, the power of any detection system is limited by the data that are (or are not) available to it. In this paper, we assume that a node must reason about faults of other nodes based on the messages it receives. This results in the following fundamental limitations for a practical detector.

First, we can only detect faults that directly or indirectly affect a message. For example, we cannot detect that a node’s CPU is overheating or that its display has failed. This would require a more powerful detector, e.g. one that has access to special hardware.

Second, it is difficult to verify whether a node correctly reports its external inputs. For example, a faulty weather station might report a light breeze as a heavy gale. Detecting faults of this type would require additional information, e.g. other weather stations in the same area or a satellite image.

Third, in an asynchronous system, it can be difficult to distinguish omission faults from messages with a particularly long delay. For example, in the absence of strong synchrony assumptions, if a node has been granted resources but does not release them for a long time, we can only *suspect* that it has become faulty, but there is no ‘smoking gun’ that would

```

client:
send(REQUEST_k)
wait_for(GRANT_k)
/* do work */
send(RELEASE_k)

server:
N := 10
repeat forever
  m := get next
  unprocessed Msg.
  if (m=REQUEST_k)
    if (N >= k)
      send(GRANT_k)
      N := N-k
    else buffer(m)
  elif (m=RELEASE_k)
    N := N+k

```




Figure 1: A simple example protocol (a) and five message exchanges: Correct behavior (b), incorrect response from node B that A can detect without C’s help (c), incorrect response from node B that A and C can detect only if they cooperate (d), missing response from node B (e), and faulty behavior by node C with B as an accomplice (f). In case (f), because node A cannot observe C’s behavior, it cannot expose the fault on C unless B cooperates.

allow us to expose it. The concept of suspected nodes is well known from the literature on unreliable fail-stop failure detectors [15]. A suspicion may sometimes turn out to be groundless, e.g. if the `RELEASE` message eventually arrives.

Finally, we can only detect faults that are *observable by a correct node*. The reason is that we cannot expect the faulty nodes to share any information about the messages they have observed. For example, they could be trying to cover the traces of other faulty nodes with whom they are colluding. Hence, if one faulty node sends a telltale message to another faulty node without changing its observable state, we cannot hope to expose the deviating node on the basis of this message only. Consider the example in Figure 1f. Here, node B can clearly observe that C is faulty, because it issues a second request without releasing its resources first. However, node A can only observe B’s `GRANT_8` message, which is correct. If node B does not share its observations with A, it can prevent C from being exposed, which makes B a *faulty accomplice* of C.

However, if none of the correct nodes can observe a particular fault, it is of little practical relevance because, from the correct nodes’ perspective, the faulty nodes are acting *as if they were correct*.

3.3 The letter versus the spirit of the law

It is tempting to think of fault detection as a panacea that works against all kinds of bad, or ‘faulty’, behavior. However, recall that our definition of ‘faulty’ depends on a specific protocol. Thus, an action may be against the *intent* behind the protocol (the spirit of the law), but may be perfectly legal with respect to the protocol’s specification (the letter of the law).

For instance, our example protocol in Figure 1a is not deadlock-free. Therefore, adding fault detection will neither make it deadlock-free nor enable nodes to detect or resolve deadlocks. However, if we augment the protocol to be deadlock-free, then a faulty node can create a deadlock only by violating the protocol. Fault detection will then be able to identify the fault, thus enabling the correct nodes to break the deadlock.

A similar argument applies to external input that is not verifiable. Consider a cooperative storage system that allows users to declare freely the amount of storage they wish to

contribute. If “zero” is a valid setting, fault detection cannot expose freeloaders who use this setting; their behavior is against the intent behind the protocol but does not violate its specification. However, if the protocol specifies a minimum storage capacity, fault detection can expose participants who refuse to store at least the minimal amount of data.

3.4 Problem statement

Now we are ready to relax and refine our “idealistic” requirements from Section 3.1. To capture the protocol dependence and the restriction to observable faults, we use the notions of *detectably faulty* and *detectably ignorant* nodes (formal definitions can be found in a technical report [24]).

Briefly, a node *i* is detectably faulty if it breaks the protocol in a way that causally affects a correct node. For instance, *i* sends an “incorrect” (with respect to its protocol) message *m* that, through a sequence of causally related messages, precedes an event observed by a correct node. The nodes that send messages caused by *m* are called *accomplices* of *i* (with respect to *m*). Note that a node can appear correct to each individual correct node and still be detectably faulty: this happens, for instance, if the node sends two observable messages that could never be sent by a correct node in the same execution (see the example in Figure 1d).

A node *i* is detectably ignorant if it never acknowledges that it received a message sent by a correct node (see the example in Figure 1e).

PeerReview produces two kinds of fault indications: *exposed* and *suspected*. Intuitively, we say that a node *i* is exposed by a node *j* if *j* has a proof of *i*’s misbehavior, and *i* is suspected by *j* if, according to *j*, *i* has not acknowledged a certain message sent to it. Note that *j* withdraws the “suspected” indication when it learns that *i* accepted the message. To account for long message delays, we require only that faults are detected *eventually*, after some delay.

Now we define the following requirements:

- **Completeness:** (1) Eventually, every detectably ignorant node is suspected forever by every correct node, and (2) if a node *i* is detectably faulty with respect to a message *m*, then eventually, some faulty accomplice of *i* (with respect to *m*) is exposed or forever suspected by every correct node.

- **Accuracy:** (1) No correct node is forever suspected by a correct node, and (2) no correct node is ever exposed by a correct node.

Although a system with these properties is weaker than our idealized detector from Section 3.1, it is still very strong in practice: every instance of detectably faulty behavior is eventually detected, and there are no false positives. As we will see later, relaxing completeness in favor of a probabilistic detection guarantee permits a highly scalable implementation, while still detecting faults with high probability and avoiding false positives.

4. DESIGN OF PEERREVIEW

In the previous section, we discussed fault detection and defined its properties. Next, we describe the design of PeerReview, an accountability system that satisfies these properties. A formal proof of these properties can be found in a technical report [24].

4.1 Overview

We first describe a simplified version of PeerReview, called FullReview. For FullReview, we make the (unrealistic) assumption that there is a trusted entity that can reliably and instantly communicate with all nodes in the system. The system’s membership is static, and each node knows the specification of the entire system.

FullReview works as follows: All messages are sent through the trusted entity, which ensures that all correct nodes observe the same set of messages in the same order. Furthermore, each node i maintains a log λ_{ij} for each other node j , and in this log it records all messages that were sent either from or to j . Periodically, i checks each of its logs against the system specification. If a node i finds that a node j has not yet sent the message it should have sent in its last observed state, then i *suspects* j until that message is sent. If j has sent a message it should not have sent according to the specification, then i *exposes* j .

It is easy to see that FullReview is both complete and accurate. On the one hand, when a node i sends an “incorrect” message, the trusted entity forwards the message to all nodes, and i is exposed by every correct node. Also i is suspected by every correct node as long as a message from i is missing. On the other hand, no correct node can be exposed or indefinitely suspected by any correct node.

However, FullReview is based on strong assumptions: a trusted, reliable communication medium and a formal system specification. Moreover, FullReview’s complexity is at least quadratic in the number of nodes, for messages, storage, and computation. In PeerReview, we refine this simple design to arrive at a practical system:

- Each node only keeps a full copy of its own log; it retrieves other logs when necessary. Nodes exchange just enough information to convince another node that a fault is, or is not, present.
- Tamper-evident logs and a commitment protocol ensure that each node keeps its log consistent with the set of messages it has exchanged with all correct nodes or else risk exposure.
- Each node is associated with a small set of other nodes, who act as its *witnesses*. The witnesses collect evidence about the node, check its correctness, and make the results available to the rest of the system.

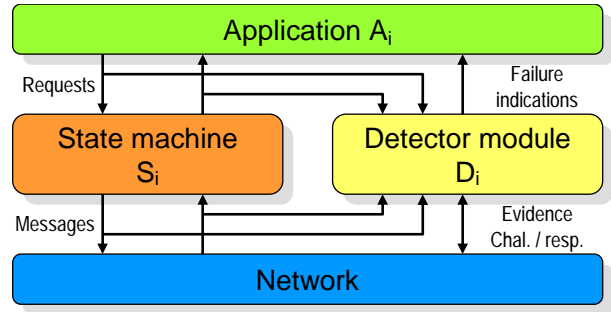


Figure 2: Information flow between application, state machine, and detector module on node i .

- PeerReview uses a reference implementation of the node software to check logs for faulty behavior. Thus, it does not require a formal system specification, which is difficult to obtain and maintain in practice.
- PeerReview uses a challenge/response protocol to deal with nodes that do not respond to some messages. This allows PeerReview to operate on an unreliable network that satisfies only weak synchrony assumptions.

We will describe each of these refinements in the following subsections.

4.2 System model

Each node i is modeled as a state machine S_i , a detector module D_i , and an application A_i (Figure 2). The state machine represents all the functionality that should be checked by PeerReview, whereas the application represents other functions that need not be checked, e.g. a GUI. The detector module D_i implements PeerReview; it can observe all inputs and outputs of S_i , and it can communicate with the detector modules on other nodes. We assume that a correct node implements S_i and D_i as specified, whereas a faulty node may behave arbitrarily.

The detector module issues *failure indications* about other nodes to its local application. Informally, *exposed(j)* is raised when i has obtained proof of j ’s misbehavior; *suspected(j)* says that i suspects that j does not send a message that it is supposed to send; *trusted(j)* is issued otherwise.

4.3 Assumptions

The design of PeerReview is based on the following assumptions:

1. The state machines S_i are deterministic.
2. A message sent from one correct node to another is eventually received, if retransmitted sufficiently often.
3. The nodes use a hash function $H(\cdot)$ that is pre-image resistant, second pre-image resistant, and collision resistant.

Assumptions 1–3 are common for techniques based on state machine replication [49], including BFT [13].

4. Each node has a public/private keypair bound to a unique node identifier. Nodes can sign messages, and faulty nodes cannot forge the signature of a correct node.




Figure 3: (a) A linear log and its hash chain, which is recursively defined on the log entries, and (b) a forked log with two branches.

- Each node has access to a reference implementation of all S_j . The implementation can create a snapshot of its state, and its state can be initialized according to a given snapshot.
- There is a function w that maps each node to its set of witnesses. It is assumed that for each node i , the set $\{i\} \cup w(i)$ contains at least one correct node; otherwise, PeerReview might lose completeness with respect to node i .

Assumption 4 can be met, for instance, by installing each node with a certificate that binds the node’s public key to its unique identifier. However, any type of name binding that avoids Sybil attacks [21] will work. In symmetric systems where all nodes run the same protocols, a node can simply use its own implementation as the reference implementation (Assumption 5). Otherwise, nodes can obtain a reference implementation for another node from a trusted source. The appropriate definition of w (Assumption 6) depends on the system configuration, which will be discussed in Section 5.3.

4.4 Tamper-evident logs

To enforce accountability, PeerReview must keep a secure record of the inputs and outputs of each node, and it must be able to detect if that record has been tampered with. PeerReview implements such a record using a technique inspired by secure histories [39].

A *log* is an append-only list that contains all the inputs and outputs of a particular node’s state machine in chronological order. The log also contains periodic state snapshots and some annotations from the detector module. Each log entry $e_k = (s_k, t_k, c_k)$ has a sequence number s_k , a type t_k , and some type-specific content c_k . The sequence numbers must be strictly increasing but may be non-contiguous; for example, a timestamp could be used. Additionally, each record includes a recursively defined hash value $h_k = H(h_{k-1} || s_k || t_k || H(c_k))$ (Figure 3a); $||$ stands for concatenation. The base hash h_{-1} is a well-known value.

The resultant hash chain, along with a set of *authenticators*, makes the log tamper-evident. An authenticator $\alpha_k^j = \sigma_j(s_k, h_k)$ is a signed statement by node j that its log entry e_k has hash value h_k ; $\sigma_j(\cdot)$ means that the argument is signed with j ’s private key.

By sending α_k^j to node i , a node j *commits* to having logged entry e_k and to the contents of its log before e_k . If j subsequently cannot produce a prefix of its log that matches

the hash value in α_k^j , then i has verifiable evidence that j has tampered with its log and is therefore faulty.

Moreover, i can use α_k^j as verifiable evidence to convince other nodes that an entry e_k exists in j ’s log. Any node can also use α_k^j to inspect e_k and the entries preceding it in j ’s log. To inspect x entries, i challenges j to return $e_{k-(x-1)}, \dots, e_k$ and h_{k-x} . If j responds, i calculates the hash value h_k from the response and compares it with the value in the authenticator. If j has not returned the correct log entries in the correct order, the hash values will differ. In this case, i has evidence that j is faulty. We discuss the case in which j does not respond to the challenge in Section 4.8.

Summary: *Logs and authenticators form a tamper-evident, append-only record of a node’s inputs and outputs.*

4.5 Commitment protocol

We must ensure that a node cannot add an entry to its log for a message it has never received. Also, we have to ensure that a node’s log is complete, i.e. that it contains an entry for each message sent or received by the node to or from a correct node.

When node i sends a message m to node j , i must commit to having sent m , and j must commit to having received m . They obtain an authenticator from the other node included in the message and its acknowledgment, respectively. This authenticator covers the corresponding log entry. A log entry for a received message must include a matching authenticator; therefore, a node cannot invent log entries for messages it never received.

When i is about to send m to j , it creates a log entry $(s_k, \text{SEND}, \{j, m\})$, attaches h_{k-1} , s_k and $\sigma_i(s_k || h_k)$ to m , and sends the result to j . Thus, recipient j has enough information to calculate h_k and to extract α_k^i . If the signature in α_k^i is not valid, j discards m . Otherwise, j creates its own log entry $(s_l, \text{RCV}, \{i, s_k, m\})$ and returns an acknowledgment with h_{l-1} , s_l , and $\sigma_j(s_l || h_l)$ to i . This allows i to extract and verify α_l^j . If i does not receive a valid acknowledgment, i sends a challenge to j ’s witnesses; we will explain the details in Section 4.8.

Summary: *The commitment protocol ensures that the sender (respectively the receiver) of each message m obtains verifiable evidence that the receiver (respectively the sender) of m has logged the transmission.*

4.6 Consistency protocol

A faulty node can attempt to escape detection by keeping more than one log or a log with multiple branches (Figure 3b). For example, this could be a promising strategy for node B in Figure 1d, who might keep one log for messages from A and another for messages from C. Both logs would show correct behavior, even though node B is clearly faulty. To avoid this attack, we exploit the fact that a node can produce a connecting log segment for each pair of authenticators it has ever signed if, and only if, it maintains a single, linear log.

If a node i receives authenticators from another node j , it must eventually forward these authenticators to the witness set $w(j)$. Thus, the witnesses obtain verifiable evidence of all the messages j has sent or received. Periodically, each witness $\omega \in w(j)$ picks the authenticators with the lowest and the highest sequence number and challenges j to return all log entries in this range. If j is correct, these log entries form a linear hash chain that contains the hash values in all the other authenticators. If they do not, ω has obtained

verifiable evidence that j is faulty. We discuss the case in which j does not respond to the challenge in Section 4.8.

Finally, each $\omega \in w(j)$ uses the log entries to extract all the authenticators that j has received from other nodes and sends them to the corresponding witness sets. This is necessary because j could be acting as a faulty accomplice of some node k ; it could forward k 's messages without sending the authenticators to $w(k)$. We note that this step has $O(|w(j)| \cdot |w(k)|)$ message complexity.

Summary: *The consistency protocol ensures that each node either maintains a single, linear log that is consistent with all the authenticators the node has issued, or it is exposed by at least one correct witness.*

4.7 Audit protocol

In the next step, we use a node's log to check whether the node's behavior conforms to that of its reference implementation. Each witness ω of a node i periodically looks up its most recent authenticator from i (say, α_k^i) and then challenges i to return¹ all log entries since its last audit, up to and including e_k . Then ω appends the new entries to its local copy $\lambda_{\omega i}$ of i 's log.

Next, ω locally creates an instance of i 's reference implementation and initializes it with a recent snapshot from $\lambda_{\omega i}$. Then, it replays all the inputs starting from that snapshot and compares S_i 's output with the output in the log. Since we require that S_i be deterministic, any discrepancy indicates that i is faulty. In this case, ω can use α_k^i and a suffix of $\lambda_{\omega i}$ as verifiable evidence against i , and this evidence can be checked by any correct node.

Summary: *The audit protocol ensures that, for each node i , either i 's actions are consistent with the reference implementation of i 's state machine, or i is exposed by at least one correct witness.*

4.8 Challenge/response protocol

The protocols described so far can expose faulty nodes if they respond to challenges. But what if a node does not respond to a challenge, or it fails to acknowledge a message that was sent to it? Unless we make stronger assumptions about synchrony, we cannot distinguish an uncooperative faulty node from a correct node that is slow or suffering from network problems.

When a node j concludes that another node i is refusing to cooperate, it indicates the suspected state for i and creates a *challenge* for i . The challenge must contain enough evidence to convince another correct node that if i were correct, it would be able to answer. The node j then sends the challenge to i 's witnesses, who forward it to i . If i does not send a response, the witnesses indicate that i is suspected. There are two types of challenges.

An *audit challenge* consists of two authenticators: a_k^i and a_l^i with $k < l$. After checking the signatures on a_k^i and a_l^i , any correct node has enough evidence to convince itself that either i is faulty or the entries e_k and e_l must exist in i 's log. If i is correct, it can answer the challenge with the corresponding log segment, whose hash values will form a hash chain connecting a_k^i to a_l^i .

A *send challenge* consists simply of a message m and the extra information appended by the consistency proto-

¹Because the audit protocol and the consistency protocol inspect the same log entries, our implementation retrieves them only once. Here, we separate the two protocols for clarity of presentation.

col (Section 4.6). After extracting and checking the authenticator from m , any correct node is convinced that i must acknowledge m . If i is correct and has not yet received m , i can accept m now and return an acknowledgment. If i has already received m , it can simply re-send the earlier acknowledgment.

Summary: *The challenge/response protocol ensures that, if a node i fails to respond to a challenge or does not acknowledge a message, it is eventually suspected by at least one correct witness. The suspicion persists until the node answers the challenges or acknowledges the message.*

4.9 Evidence transfer protocol

The mechanisms described so far ensure that at least one correct node obtains verifiable evidence of each fault or a challenge for each suspected node. We also need to make sure that all correct nodes eventually collect the same evidence (the same set of audit and send challenges) against faulty nodes. We achieve this by allowing every node i to periodically fetch the challenges collected by the witnesses of every other node j . Note that, in most practical settings, i may be interested only in accusations against the nodes that (directly or indirectly) communicate with it. In this case, i needs to contact only the witnesses of these nodes.

If a correct node i obtains a challenge for another node j , its detector indicates *suspected*(j). When i receives a message from j in this state, it challenges j . Once i has received valid answers to all pending challenges, its detector indicates *trusted*(j) again. If i obtains a proof of j 's misbehavior, its detector outputs *exposed*(j).

Summary: *The evidence transfer protocol ensures that all correct nodes eventually output a failure indication for each faulty node.*

4.10 Summary

The use of logs and authenticators helps to maintain a shared tamper-evident, append-only record of every node's activity. By periodically checking and replaying the logs, and comparing the results with the provided authenticators, the consistency, commitment and audit protocols ensure that a detectably faulty node is exposed or permanently suspected by at least one of its correct witnesses. The challenge/response protocol ensures that a node is suspected if it is unresponsive but allows a correct node to exonerate itself by responding. Finally, the evidence transfer protocol allows the correct nodes to share their evidence and eventually to come to the same conclusion about each detectably faulty node. To summarize, PeerReview has the completeness and accuracy properties described in Section 3.4. A proof of these properties can be found in a technical report [24].

4.11 Extension: Probabilistic guarantees

PeerReview provides strong guarantees under very conservative failure assumptions; this limits its scalability to moderately large systems. If we assume an upper bound φ on the fraction of faulty nodes in a system of N nodes, then we can ensure strong completeness only if we assign $\psi = \lceil \varphi N \rceil$ witnesses to each node. In this case, the message complexity is dominated by the consistency protocol's $O(\psi^2)$ complexity and hence grows with $O(N^2)$.

However, there are two extensions that can considerably improve PeerReview's scalability, at the expense of a slightly relaxed completeness guarantee. If we accept a small prob-

ability $P_f > 0$ that an all-faulty witness set exists, we need to assign only

$$\psi = \left\lceil \frac{\ln(1 - (1 - P_f)^{\frac{1}{N}})}{\ln \varphi} \right\rceil$$

witnesses to each node, which grows with $O(\log N)$. In addition, if we accept a small probability $P_m > 0$ that a given instance of misbehavior remains undetected, we can design a randomized consistency protocol in which a node sends an authenticator only with probability

$$\xi = \frac{1 - (1 - P_m)^{\frac{1}{\psi^2}}}{(1 - \varphi)^2}$$

This removes most of the redundant transmissions and reduces the message complexity of the consistency protocol to a constant. If both extensions are combined, the message complexity is dominated by the audit protocol and grows with $O(\log N)$.

Why does $P_m > 0$ lead to such a marked improvement? The reason is that, if $P_m = 0$, the consistency protocol must guarantee completeness even when both the sender’s and the receiver’s witness set contain just a single correct node. As the size of the witness set grows, this extreme case becomes more and more unlikely. A detailed probabilistic analysis of the randomized protocol can be found in a technical report [24].

5. IMPLEMENTATION

In this section, we describe our implementation of PeerReview. We discuss optimizations and practical engineering challenges, such as log truncation. Then, we discuss how the implementation can be adapted to specific environments through an appropriate choice of witnesses and configuration parameters.

5.1 The PeerReview library

Our PeerReview codebase is written as a C++ library, which includes all application-independent parts of the system. The library is used by all three example applications described in Section 6 and can be reused for other systems. It implements the tamper-evident logs and PeerReview’s five protocols. External libraries provide the cryptographic primitives and a transport layer for sending and receiving messages. The application system is expected to provide various callbacks, including one that instantiates a reference implementation of another node, which is used when PeerReview needs to check that node’s log.

Conceptually, the library is interposed between the transport layer and the application. All incoming and outgoing messages go through the library. The library may append headers, such as an authenticator for the consistency protocol, or send messages of its own, e.g. a challenge. Our PeerReview library contains 5,961 lines of code, counted by the number of newlines. It is available for download from the project homepage [46].

The library makes use of two simple optimizations to save bandwidth in the auditing protocol. First, when transferring log segments to a witness, it replaces all the outputs and state snapshots with hash values. The witness does not lose information that way, because it obtains the outputs during replay. It can compare the hash of these outputs with the hash values in the log. Second, if part of a message is never

examined by the state machine, the library hashes that part also, because it does not affect the replay. For example, a message handled by a routing protocol can be reduced to its routing header and a hash of its body. The library provides the full set of snapshots and outputs upon request, e.g. when a new witness needs to be initialized.

Another optimization applies to the consistency protocol. Rather than forwarding authenticators to the witnesses immediately, the library buffers them locally for some time T_{buf} and then sends them in batches. This increases the time to detection by up to T_{buf} but increases message efficiency because authenticators are very small. An authenticator with a SHA-1 hash and a 1024-bit RSA signature is just 156 bytes long.

The library also supports authenticated dictionaries [42], which applications can use to commit efficiently to large data structures, such as an entire disk image. This is useful, e.g. when a node generates evidence. The node can include only those parts of the data structure that the recipient needs to check the evidence, e.g. the disk blocks that have actually been accessed.

5.2 Log truncation

PeerReview’s append-only log must eventually be truncated. A simple solution is to allow each node to discard all log entries older than some time T_{trunc} . This approach requires very loosely synchronized clocks. Whether a node’s clock is sufficiently accurate can be checked by PeerReview.

Despite log truncation, exposed faulty nodes remain exposed forever, because the incriminating evidence can be verified without access to the corresponding log entries. A fault could remain undetected if legitimate evidence of the fault were to surface only after the log was truncated. This case can be avoided by keeping logs much longer than the expected duration of node and network outages. In practice, log entries can and should be kept at least on the order of months.

As a result of log truncation, a suspected node that was never exposed may become trusted again after T_{trunc} has elapsed. However, an attacker cannot gain much leeway from this exoneration, because a faulty node must remain silent for at least $T_{trunc} - T_{audit} \approx T_{trunc}$ (e.g. several months) after each misbehavior to avoid exposure. Moreover, a human operator could infrequently check for nodes that have been suspected for a long time and permanently revoke or refuse to renew such nodes’ certificates.

5.3 Configuring witnesses

Witness configuration depends on the type of system and the nature of the deployment in which PeerReview would be used. In a system that is overseen by a single organization (e.g. Planetlab), one can configure a dedicated set of machines as witnesses for all nodes. In a federated system (e.g. Internet inter-domain routing), each organization may wish to act as a witness for the organization’s peering partners. In a client-server system (e.g. a Web service), clients may act as witnesses for the servers they depend on. Alternatively, replicated servers can act as mutual witnesses. Lastly, in a peer-to-peer system (e.g. Skype), each node can be witnessed by a random set of other participating nodes.

Depending on the choice of witness configuration, an appropriate function w , which maps each node to its witness set, is defined. In the cases in which the system membership is relatively static, we can simply specify w in a configura-

tion file, signed by the appropriate authority and distributed throughout the system. If membership is dynamic, w must be dynamic as well; new witnesses can initialize their state by first obtaining some recent authenticators from the old witnesses and then performing an audit to obtain the latest checkpoint.

In peer-to-peer systems, we use consistent hashing [29] to map witnesses to nodes; in this case, each node acts as a witness for the k nodes whose node identifiers are closest to its own. This approach spreads the auditing overhead evenly across the nodes, and it minimizes the number of witness sets that are affected by a random fault. Nodes must not be allowed to choose their own identifiers, and they must be able to securely evaluate w at runtime, even in the presence of faulty nodes. Secure routing [11] ensures this.

In systems with dynamic witness sets, additional bandwidth is required to initialize new witnesses, which may limit PeerReview’s tolerance of high levels of churn. Note that *some* rate of churn among witnesses is actually desirable in configurations with small witness sets (Section 4.11): in this case, it ensures that the (unlikely) state in which some faulty node’s witnesses are all faulty does not persist indefinitely.

5.4 Choosing parameters

PeerReview’s most important parameter is the size of a node’s witness set. Consider, for instance, the case when witness sets of all nodes have the same size ψ . A higher ψ increases overhead. PeerReview’s storage requirement and CPU overhead grow with $O(\psi)$, and its message complexity with $O(\psi^2)$. To maintain PeerReview’s completeness guarantee, ψ must be chosen such that for each node i , the set $\{i\} \cup w(i)$ contains at least one correct node at any given time. In other words, either i or one of its ψ witnesses must be correct. In practice, ψ should be chosen as the minimal value that satisfies this criterion.

If a lower value is chosen for ψ , PeerReview might lose its strong completeness guarantee, i.e., it can miss some faults. However, it still does not suspect or expose any correct nodes. For more details, see Section 4.11.

The audit interval T_{audit} and the time limit T_{buf} for buffering authenticators determine the maximum time to detection. Shorter intervals result in quicker detection and thus reduce the damage a faulty node can do before it is evicted; however, short intervals also increase message overhead because logs are transferred in small pieces rather than in large segments.

It is also possible to perform audits on demand, rather than periodically. For instance, in a system providing best-effort service, witnesses may refrain from auditing until they observe a loss of service quality that indicates the presence of faults. With this policy, detection is not guaranteed during periods when the witnesses do not audit. During periods when the witnesses audit, PeerReview provides the usual guarantees.

The key length k of the nodes’ keys must be high enough to prevent an attacker from forging signatures; however, higher values also mean more CPU load for signing and verifying authenticators. For our experiments, we used RSA with the recommended key length of 1024 bits [5].

5.5 Alternatives to fault detection

PeerReview has two main components. One component provides a tamper-evident record of all observable actions; the other component detects faults by regularly checking that

record against a reference implementation. For some applications, it may be useful to replace the second component with another fault detection mechanism. For example, systems like BGP [48] have formal specifications of correct or incorrect behavior; in these cases, the auditors could directly check the record against the specification, e.g. using [6]. Another example is an electronic voting system, in which the secure record could be inspected in court in case an election is contested.

6. APPLICATIONS

In this section, we first discuss the general steps for applying PeerReview to an application, and then we briefly describe our three example applications. Additional details can be found in [24].

6.1 How to apply PeerReview

Applying PeerReview to a new system generally involves two steps. The first step is to decide which parts of the system should be included in the state machine S_i for each node i and thus checked by PeerReview. Including more subsystems potentially enables PeerReview to detect more types of faults. However, because PeerReview allows any node to audit past and current states of S_i , any subsystems that deal with sensitive information, e.g. with cryptographic keys or clear-text private information, should not be included.

The second step is to make sure S_i is deterministic. Given some initial state and a sequence of inputs, S_i must always produce the same outputs and finish in the same state. In practice, many systems have some sources of non-determinism, but usually these can be identified and dealt with. For example, if a system uses a pseudo-random number generator, we include its seed value in the state snapshot, so the same random sequence is generated during replay. If a protocol relies on real time, e.g. for timestamps or timeouts, we record the time of each event in the log; we use this information to update a virtual clock used during replay and to trigger timeouts.

If a system is concurrent and generates events in a non-deterministic order, we can use synchronization to make the order deterministic. For instance, we can add synchronization to ensure that jobs finish in the order in which they were started; alternatively, we can log the sequence in which the jobs actually finished and enforce this sequence during replay.

In general, adding PeerReview to a system presents the same challenges as adding state-machine replication (e.g. BFT), and the same solutions apply (see e.g. [10]). In addition to eliminating sources of non-determinism, PeerReview requires state snapshots to enable replay by witnesses, whereas BFT requires snapshots to initialize new replicas. To transfer snapshots between different implementations, BFT can use state abstraction [14]; this technique can be applied to PeerReview as well.

6.2 Application #1: Overlay multicast

Our first application delivers streaming content, such as audio or video, from a single source to a potentially large set of client nodes. Because the source may not have enough bandwidth to support all the clients simultaneously, the clients must contribute some of their own bandwidth and forward the content they receive to some other clients. We use a simple *multi-tree* multicast protocol to ensure that the forwarding load can be distributed well among the clients [12].

The content is striped across multiple trees, and each client is required to be an interior node in one of the trees and a leaf node in the others.

Overlay multicast provides a best-effort service. By applying PeerReview to this system, we gain the ability to detect and isolate misbehaving nodes, e.g. nodes that tamper with the content they are forwarding, or so-called freeloaders, who refuse to contribute resources to the system. The witnesses for a node are selected randomly among the remaining clients. If a client is suspected or exposed, its parents refuse to deliver to it any more content, and its children seek a different parent. This limits the damage a faulty node can cause, and it creates a disincentive for clients to misbehave. Therefore, PeerReview lends the system robustness in the face of malicious participants and freeloaders.

6.3 Application #2: Network filesystem

In this application, we added PeerReview to an existing NFSv2 server implementation [53]. PeerReview allows a set of NFS servers to check each other. Because NFS servers do not tolerate Byzantine faults, a faulty server may deliver incorrect data to a client. However, PeerReview allows us to detect quickly if an attacker has tampered with one of the file servers, e.g. by removing or corrupting data.

Each file server exports a different volume via NFS. In addition, each server is a witness for some of the other servers and audits them periodically. To save bandwidth, each witness maintains a full replica of the volumes it audits, rather than repeatedly transferring checkpoints. When a witness notices a fault, it uses the authenticated dictionary to reduce the size of the evidence, leaving only those parts of the volume that are relevant for checking. Then it distributes the evidence to the administrators and any interested clients. Thus, faults can be quickly repaired and cannot go unnoticed for long.

Our implementation uses the NFS server in the Linux 2.6.15 kernel in combination with a user-level wrapper process for PeerReview. Since the `ext2` file system is not completely deterministic, we had to apply a small 467-line kernel patch to remove sources of non-determinism². The patch also adds a mechanism for setting the file system time to a given value, which is necessary to get the same timestamps during replay.

6.4 Application #3: Peer-to-peer email

ePOST [40] is a peer-to-peer email service. Email and email folders are stored in a distributed hashtable (DHT), which is cooperatively implemented by all participating nodes. The DHT is replicated for availability and durability [26], and it can adapt to node failures and churn. To ensure confidentiality, all content is encrypted before it is added to the DHT. The key of an object in the DHT is the hash of its content, which makes the object's integrity easy to verify.

Given the existing security and fault tolerance mechanisms in ePOST, the remaining threat is denial of service. A faulty node can misroute or drop messages and thus prevent other nodes from retrieving content, or it can manipulate the topology of the peer-to-peer system to hide the presence of other nodes. We use PeerReview to identify misbehaving

²For example, `ext2` has a randomized block allocation policy, which matters because we use disk images as checkpoints. We made it deterministic by tying its random seed to the current time.

nodes and to eject them from the system. Thus, faulty nodes cannot degrade the (best-effort) service indefinitely.

To handle these threats, it is sufficient to include the DHT in the state machine. ePOST's cryptographic and IMAP mechanisms need not be checked by PeerReview, which preserves the confidentiality of email. We could not experiment with a live deployment because ePOST no longer has a significant user base; instead, our experiments are based on message traces from a past deployment provided by the ePOST authors.

6.5 Summary

Our three example applications are not the only, and may not be the most natural, applications for PeerReview. We chose them for two reasons. First, they represent different types of systems that PeerReview can be applied to. The filesystem is a small-scale client-server system for a LAN environment, whereas overlay multicast and ePOST are decentralized, cooperative systems for wide-area deployment. Second, the applications place stress on different aspects of PeerReview and explore its limitations. For instance, the fileservers have a large amount of state and a high rate of latency-sensitive requests. Overlay multicast has a high message rate and transmits a large volume of data. ePOST is a complex peer-to-peer system that includes distributed storage and a DHT.

7. EVALUATION

In this section, we present experimental results from our three example applications. We cover the systems' behavior under faults, the overheads of running PeerReview and its impact on application performance.

7.1 Methodology

Each of our applications places stress on a different aspect of PeerReview. For example, the network filesystem is sensitive to latency increases and is limited by throughput, whereas ePOST tolerates wide-area latencies and is typically lightly loaded. ePOST runs on nodes in residential or wireless networks and thus has limited bandwidth, whereas the network filesystem usually runs on a cluster with high-speed links. For this reason, we chose to evaluate each performance metric in the context of the application for which that metric is most critical.

We used Sun V20Z rack servers in a local cluster, consisting of dual 2.5 GHz Opteron CPUs connected with 1 Gbps switched Ethernet, running Linux 2.6.15. For the filesystem experiments, the workload was generated by a host in the same subnet. We compare the performance of the kernel-level NFS server with and without our user-level implementation of PeerReview.

For the multicast experiments, we generated a CBR stream of 300 Kbps, which is a common rate for streaming video to clients with broadband connections. We used a delay buffer of 10 seconds and counted blocks as lost if they were not received by this deadline. To obtain controlled conditions, we used a simple network emulator that forwards packets after a configurable delay. However, our implementation is complete and can also be run on a network testbed like PlanetLab.

Our ePOST experiment was driven by a one-day trace from a real ePOST deployment with 25 nodes/users, which was kindly provided by the authors of ePOST. The trace contains all messages sent on December 22, 2005, as well




Figure 4: Multicast with a single freeloader: Impact on an average correct node (left bar) and on the freeloader (right bar). The error bars show minimum and maximum values. The freeloader uses different strategies, which are described in the text.

as all DHT operations (get/put) and all churn events (online/offline). We could not obtain a list of the exact keys and sizes of the objects in the DHT, but we were able to estimate its overall size. The experiment used our network emulator, although our code also runs on real networks.

In all experiments, we applied consistent hashing to choose a set of ψ witnesses per node, and we configured a high audit frequency of $T_{audit} = 10$ secs. The buffer delay T_{buf} was 100 ms for overlay multicast and 5 secs otherwise.

7.2 Fault injection experiments

In our first experiment, we inject a few faults into our systems. In the first scenario, three³ faulty ePOST nodes F_1 – F_3 try to censor an object O . First, a node A joins and inserts O into the DHT. To make it easy for the three nodes, we choose O ’s key such that F_1 – F_3 store the replicas. Then A leaves the system and, 10 seconds later, another node B attempts to look up the object in the DHT. The request is routed to the faulty nodes, who reply that O does not exist.

Without PeerReview in place, lookups of O take a long time, because ePOST’s archival store has to retrieve the missing object each time. Moreover, this state can persist indefinitely, because the system cannot identify and remove the faulty nodes⁴. With PeerReview, however, the faulty nodes are exposed after their first incorrect response. The effect is the same as if F_i had left the system: ePOST repairs its routing tables, new replicas of O are restored from ePOST’s archival background store and subsequent lookups of object O succeed immediately. We note that this is possible even though A and B never communicate directly; they are never even online at the same time.

Our second experiment tests overlay multicast’s response to freeloaders. We set up three multicasts, each with one source, 100 clients, 10 trees, and $\psi = 2$ witnesses per

³ePOST stores three full replicas of each object in its DHT plus a number of erasure-coded fragments in its archival store. The archival store is accessed only if the object is not found in the DHT.

⁴Note that secure routing [11] would not solve this problem. The lookups are delivered to the right nodes, but these nodes do not respond correctly.

Run	Min	Avg	Max
nfs	66 μs	78 μs	99 μs
nfs+wrap	91 μs	113 μs	136 μs
peerreview-nosig	140 μs	144 μs	241 μs
peerreview-1024	1,591 μs	1,624 μs	1,938 μs

Table 1: Round-trip time for a NULL RPC.

node. One of the clients either behaved as expected (“Cooperative”) or freeloaded with a “Reluctant” strategy, forwarding content only when challenged by its witnesses, or freeloaded with a “Silent” strategy, refusing to forward any content. Without PeerReview, both strategies would enable the freeloader to receive all content without forwarding any content.

With PeerReview in place (Figure 4), the “Reluctant” strategy still enables the freeloader to receive most of the content in time but it must send more traffic than with the cooperative strategy because it cannot avoid forwarding each message eventually, and it must constantly answer challenges sent to it via its witnesses. The “Silent” strategy, by contrast, causes the freeloader to be exposed, so the other nodes stop sending content to it. Thus, both freeloading strategies are unattractive for the freeloader.

7.3 Message latency

PeerReview increases the latency of message transfers because the sender must sign each message before transmitting it, and the receiver must check the signature before accepting the message. Also, both sides must append an entry to their log.

To quantify the impact, we set up an experiment in which we measured the latency of an NFS NULL request (essentially a no-op). This is a worst-case scenario in terms of the impact on message latency: both the propagation delay and the processing time are very small, so the additional latency has a high impact. Table 1 shows our results. ‘nfs’ is the bare NFS server; ‘nfs+wrap’ adds the cost of invoking the user-level wrapper for each RPC request and each response; and ‘peerreview-nosig’ adds the PeerReview wrapper with dummy signatures. The final result is for the full PeerReview protocol with 1024-bit RSA keys.

Our results show that the overhead is dominated by the time needed to generate and verify cryptographic signatures. Each RPC requires two messages, and for each message, a signature must be generated by the sender and verified by the receiver. This adds a constant delay that depends on the key length and the cryptographic algorithm used. For RSA-1024, the delay is about 1.5 ms per RPC, which matters for local-area applications with many small requests but is insignificant for wide-area applications. Moreover, if we use the more efficient ESIGN algorithm with 2048-bit keys, two signatures can be generated and verified in less than 250 μs on the same hardware. As an additional optimization, small messages could be signed in batches.

7.4 Throughput

In configurations where nodes are witnesses for each other, a node must replay and check requests directed to ψ other nodes, as part of its responsibility as a witness. Hence, we expect the average throughput of each node to drop to $\frac{1}{\psi+1}$ of its capacity.




Figure 5: Throughput results for random 1 kB NFS read accesses over 10 GB of data.

To verify this, we deployed our network filesystem with four servers, using bare NFS in one experiment and adding PeerReview with $\psi = 2$ witnesses per server in another. Each server exported 10 GB worth of data. Then we sent random 1 kB read requests to each server at a constant rate, and we measured the average request latency. When the request rate reaches the server’s capacity, we expect the latency to increase sharply. Figure 5 shows our results for a single server. As expected, throughput drops to approximately one-third with PeerReview enabled.

However, if the workload is bursty, PeerReview can temporarily increase its throughput by deferring audits to periods of low load, at the expense of a slightly higher time to detection. We repeated the previous experiment with short request bursts (see “audits deferred” in Figure 5) and found that under these conditions, the PeerReview-enabled servers could almost reach the full throughput of the unmodified NFS server. This is important because, in practice, most systems experience bursty and/or diurnal workloads. With a sufficiently high T_{audit} , these systems can profit from accountability without having to pay for it with a lower throughput.

7.5 CPU overhead

If a PeerReview-enabled node handles many very short requests, the cryptographic operations required for each message can form a bottleneck. Suppose verifying one signature and creating another takes c seconds. Then a node cannot respond to more than $\frac{1}{c}$ incoming messages per second, even if the actual processing takes no time at all. However, concurrent signature generations are trivially parallelizable, so the problem can be overcome by using more CPUs or a CPU with multiple cores.

To quantify this, we saturated a PeerReview-enabled NFS server with NULL RPCs and found that it could handle 3,900 RPCs per second with one of the two CPUs disabled. When we enabled the second CPU, the throughput rose to 7,500 RPCs per second, a 92% increase. This shows that PeerReview can easily take advantage of additional CPUs. Because multi-core CPUs will become the standard, we expect that most applications will not be limited by the cost of asymmetric cryptography.

7.6 Network traffic and disk space

PeerReview requires network bandwidth for auditing and consistency checking, and disk space to store logs and evidence. As discussed in Section 4, the overhead depends mainly on the number ψ of witnesses per node. To quan-




Figure 6: Average network traffic in ePOST.

tify this, we used our ePOST application and the trace from a deployment with real users. We replayed the trace several times, varying the parameter ψ , and we measured the average number of bytes sent per node per hour. We also measured the average log size.

Figure 6 shows our results. With $\psi = 5$ witnesses, the average network traffic of an ePOST node increased from 18.0 Kbps to 80.5 Kbps. The overhead is initially dominated by the audits, but their share grows only with $O(\psi)$, whereas the consistency protocol’s share grows with $O(\psi^2)$. We also note that our optimizations from Section 5.1 are effective; without them, the audit traffic would be at least ψ times the payload.

PeerReview adds a considerable amount of network traffic. However, the bulk of the additional network traffic is caused by auditing and consistency. With the typical bursty or diurnal workloads that many systems experience, both of these tasks can be deferred to periods of low load. Thus, they do not necessarily reduce the peak throughput of the system, nor do they increase the peak network traffic.

On average, each node’s log grew at a rate of 14.4 MB per hour, independent of the parameter ψ . This means that a year’s log can be kept in ePOST in less than 125 GB of disk storage per node. In the multicast experiment, the log grew at a rate of 147 MB per hour. However, this includes 135 MB of streamed content, which could be safely discarded after a few minutes.

7.7 Scalability

In a system with a large number of nodes, it can be difficult to ensure an absolute bound on the number of faulty nodes. Instead, a bound on the fraction φ of faulty nodes is often assumed. To maintain PeerReview’s completeness guarantees, we must choose $\psi = \lceil \varphi N \rceil$ in this case; otherwise, all witnesses of a faulty node could also be faulty. Because the consistency protocol creates a per-node overhead that grows with $O(\psi^2)$, we arrive at an overhead of $O(N^2)$. Extrapolating from the results in Figure 6, and assuming that each node has a broadband connection with an upstream bandwidth of 512 Kbps, ePOST should be able to scale to a configuration with $\psi = 13$ witnesses. Assuming a bound of $\varphi = 10\%$ faulty nodes, the system would scale to about 130 nodes with broadband connections.

However, scalability can be improved considerably by relaxing PeerReview’s guarantees as described in Section 4.11. We implemented smaller witness sets and randomized consistency checking for both ePOST and overlay multicast. Figure 7 shows how the per-node traffic grows with the system size N if we assume a bound of $\varphi = 10\%$ faulty nodes and




Figure 7: Scalability for ePOST (left) and overlay multicast (right) with strong guarantees and with probabilistic guarantees, using smaller witness sets (SWS) and/or randomized consistency checking (RCC). The ePOST numbers are extrapolated from Figure 6, while the other results are from the multicast experiment.

$P_f = P_m = 10^{-6}$. When both modifications are combined, ePOST scales to over 10,000 nodes and overlay multicast to at least 1,000 nodes with broadband connections.

7.8 Additional benefits

While experimenting with PeerReview, we discovered some unexpected benefits. PeerReview turned out to be a useful tool for tracking down race conditions and other ‘heisenbugs’. These bugs are exposed by PeerReview because they typically occur either during execution or during replay, but not both. In this case, PeerReview conveniently produces evidence that can be used to reproduce the bug. We found several of these bugs in the original ePOST codebase, and these have been confirmed by the ePOST developers.

Another unexpected benefit was the ability to check for protocol conformance in a system that contains more than one implementation of the same protocol. Because each node uses its own implementation as a reference when witnessing another node, deviations from the protocol (e.g. in corner cases) can be detected. In this case, PeerReview provides evidence that can be used to reproduce the different behavior.

8. CONCLUSION

In this paper, we have described PeerReview, a general and practical system that provides accountability and fault detection for distributed systems. PeerReview guarantees the eventual detection of all Byzantine faults whose effects are observed by a correct node. Verifiable evidence of a fault is irrefutably linked to a faulty node, while correct nodes can defend themselves against false accusations. We applied PeerReview to three sample applications and experimentally evaluated the resultant systems. Our experiments indicate that PeerReview is practical and useful for a range of distributed applications. Despite its strong guarantees, PeerReview scales to moderately large systems; however, by relaxing completeness in favor of a probabilistic detection guarantee, PeerReview can scale to very large systems.

Acknowledgments

We thank the anonymous reviewers and our shepherd, Mike Schroeder, for their helpful comments. Thanks to Alan Misllove and James Stewart for their help with ePOST. This research was supported in part by US National Science Foundation grants ANI-0225660 and CNS-0509297. Andreas Haeberlen received an SOSP student travel scholarship, which was supported by Infosys.

9. REFERENCES

- [1] A. S. Aiyer, L. Alvisi, A. Clement, M. Dahlin, J.-P. Martin, and C. Porth. BAR fault tolerance for cooperative services. In *Proc. SOSP’05*, Oct 2005.
- [2] L. Alvisi, D. Malkhi, E. T. Pierce, and M. K. Reiter. Fault detection for Byzantine quorum systems. *IEEE Trans. Parallel Distrib. Syst.*, 12(9):996–1007, 2001.
- [3] R. Arends, R. Austein, M. Larson, D. Massey, and S. Rose. RFC 4033: DNS security introduction and requirements. <http://www.ietf.org/rfc/rfc4033.txt>, Mar 2005.
- [4] K. Argyraki, P. Maniatis, O. Irzak, and S. Shenker. An accountability interface for the Internet. In *Proc. 14th ICNP*, Oct 2007.
- [5] E. Barker, W. Barker, W. Burr, W. Polk, and M. Smid. Recommendation for key management. NIST Special Publication 800-57, revised Mar 2007.
- [6] H. Barringer, A. Goldberg, K. Havelund, and K. Sen. Rule-based runtime verification. In *Proc. VMCAI’04*, Jan 2004.
- [7] R. A. Bazzi and G. Neiger. Simplifying fault-tolerance: Providing the abstraction of crash failures. *J. ACM*, 48(3):499–554, 2001.
- [8] G. Bracha. Asynchronous Byzantine agreement protocols. *Information and Computation*, 75(2):130–143, Nov 1987.
- [9] G. Bracha and S. Toueg. Asynchronous consensus and broadcast protocols. *Journal of the ACM*, 32(4), 1985.
- [10] M. Castro. *Practical Byzantine Fault Tolerance*. PhD thesis, MIT LCS, Jan 2001.
- [11] M. Castro, P. Druschel, A. Ganesh, A. Rowstron, and D. S. Wallach. Secure routing for structured peer-to-peer overlay networks. In *Proceedings of OSDI’02*, Boston, MA, Dec 2002.
- [12] M. Castro, P. Druschel, A.-M. Kermarrec, A. Nandi, A. Rowstron, and A. Singh. SplitStream: High-bandwidth multicast in cooperative environments. In *Proceedings of SOSP’03*, Oct 2003.
- [13] M. Castro and B. Liskov. Practical Byzantine fault tolerance and proactive recovery. *ACM Trans. Comput. Syst.*, 20(4):398–461, 2002.
- [14] M. Castro, R. Rodrigues, and B. Liskov. BASE: Using abstraction to improve fault tolerance. *ACM Trans. Comput. Syst.*, 21(3):236–269, 2003.
- [15] T. D. Chandra and S. Toueg. Unreliable failure detectors for reliable distributed systems. *Journal of the ACM*, 43(2):225–267, Mar 1996.
- [16] B.-G. Chun, P. Maniatis, S. Shenker, and J. Kubiawicz. Attested append-only memory: Making adversaries stick to their word. In *Proc. SOSP*, Oct 2007.

- [17] B. A. Coan. A compiler that increases the fault tolerance of asynchronous protocols. *IEEE Trans. Computers*, 37(12):1541–1553, 1988.
- [18] L. P. Cox and B. D. Noble. Samsara: Honor among thieves in peer-to-peer storage. In *Proceedings of SOSP*, Oct 2003.
- [19] D. E. Denning. An intrusion-detection model. *IEEE Trans. on Software Engineering*, 13(2):222–232, 1987.
- [20] R. Dingledine, M. J. Freedman, and D. Molnar. *Peer-to-Peer: Harnessing the Power of Disruptive Technologies*, chapter Accountability. O’Reilly and Associates, 2001.
- [21] J. R. Douceur. The Sybil attack. In *Proceedings of IPTPS*, Mar 2002.
- [22] A. Doudou, B. Garbinato, R. Guerraoui, and A. Schiper. Muteness failure detectors: Specification and implementation. In *EDCC*, pages 71–87, 1999.
- [23] T. Garfinkel, M. Rosenblum, and D. Boneh. Flexible OS support and applications for trusted computing. In *Proc. HotOS’03*, May 2003.
- [24] A. Haeberlen, P. Kouznetsov, and P. Druschel. PeerReview: Practical accountability for distributed systems. Technical Report 2007-3, Max Planck Institute for Software Systems.
- [25] A. Haeberlen, P. Kouznetsov, and P. Druschel. The case for Byzantine fault detection. In *Proceedings of HotDep*, Nov 2006.
- [26] A. Haeberlen, A. Mislove, and P. Druschel. Glacier: Highly durable, decentralized storage despite massive correlated failures. In *Proc. NSDI*, May 2005.
- [27] M. Jelasity, A. Montresor, and O. Babaoglu. Detection and removal of malicious peers in gossip-based protocols. In *Proceedings of FuDiCo*, Jun 2004.
- [28] S. D. Kamvar, M. T. Schlosser, and H. Garcia-Molina. The EigenTrust algorithm for reputation management in p2p networks. In *Proc. 12th International WWW Conference*, May 2003.
- [29] D. R. Karger, E. Lehman, F. T. Leighton, R. Panigrahy, M. S. Levine, and D. Lewin. Consistent hashing and random trees: Distributed caching protocols for relieving hot spots on the world wide web. In *STOC*, 1997.
- [30] S. T. Kent, C. Lynn, J. Mikkelsen, and K. Seo. Secure border gateway protocol (S-BGP) - real world performance and deployment issues. In *NDSS*, 2000.
- [31] K. P. Kihlstrom, L. E. Moser, and P. M. Melliar-Smith. Byzantine fault detectors for solving consensus. *The Computer Journal*, 46(1):16–35, 2003.
- [32] C. Ko, G. Fink, and K. Levitt. Automated detection of vulnerabilities in privileged programs using execution monitoring. In *Proceedings of the 10th Computer Security Application Conference*, Dec 1994.
- [33] L. Lamport. Using time instead of timeout for fault-tolerant distributed systems. *ACM Trans. Prog. Lang. Syst.*, 6(2):254–280, 1984.
- [34] L. Lamport, R. Shostak, and M. Pease. The Byzantine generals problem. *ACM Trans. Program. Lang. Syst.*, 4(3):382–401, 1982.
- [35] B. W. Lampson. Computer security in the real world. In *Proc. Annual Computer Security Applications Conference*, Dec 2000.
- [36] P. Laskowski and J. Chuang. Network monitors and contracting systems: competition and innovation. In *Proc. SIGCOMM’06*, pages 183–194, Sep 2006.
- [37] J. Li, M. Krohn, D. Mazières, and D. Sasha. Secure untrusted data repository (SUNDR). In *Proceedings of OSDI 2004*, Dec 2004.
- [38] D. Malkhi and M. K. Reiter. Unreliable intrusion detection in distributed computations. In *CSFW*, pages 116–125, 1997.
- [39] P. Maniatis and M. Baker. Secure history preservation through timeline entanglement. In *Proc. of the 11th USENIX Security Symposium*, Jan 2002.
- [40] A. Mislove, A. Post, A. Haeberlen, and P. Druschel. Experiences in building and operating ePOST, a reliable peer-to-peer application. In *Proceedings of EuroSys 2006*, Apr 2006.
- [41] A. Nandi, T.-W. Ngan, A. Singh, P. Druschel, and D. Wallach. Scrivener: Providing incentives in cooperative content distribution systems. In *Proc. Middleware Conference*, Nov 2005.
- [42] M. Naor and K. Nissim. Certificate revocation and certificate update. In *Proc. USENIX Security Symposium*, pages 217–228, 1998.
- [43] G. Neiger and S. Toueg. Automatically increasing the fault-tolerance of distributed systems. In *PODC*, pages 248–262, 1988.
- [44] D. Oppenheimer, A. Ganapathi, and D. A. Patterson. Why do internet services fail, and what can be done about it? In *Proc. USITS’03*, Mar 2003.
- [45] M. C. Pease, R. E. Shostak, and L. Lamport. Reaching agreement in the presence of faults. *J. ACM*, 27(2):228–234, 1980.
- [46] PeerReview project homepage. <http://peerreview.mpi-sws.mpg.de/>.
- [47] H. V. Ramasamy, A. Agbaria, and W. H. Sanders. A parsimonious approach for obtaining resource-efficient and trustworthy execution. *IEEE Trans. on Dependable and Secure Comp.*, 4(1):1–17, Jan 2007.
- [48] P. Reynolds, O. Kennedy, E. G. Sirer, and F. B. Schneider. Securing BGP using external security monitors. Technical Report 2006-2065, Cornell University, Computing and Inform. Science, Dec 2006.
- [49] F. B. Schneider. Implementing fault-tolerant services using the state machine approach: a tutorial. *ACM Computing Surveys*, 22(4):299–319, 1990.
- [50] B. Schneier and J. Kelsey. Cryptographic support for secure logs on untrusted machines. In *Proc. 7th USENIX Security Symposium*, pages 53–62, Jan 1998.
- [51] A. Singh, T.-W. Ngan, P. Druschel, and D. Wallach. Eclipse attacks on overlay networks: Threats and defenses. In *Proceedings of INFOCOM*, Apr 2006.
- [52] T. K. Srikanth and S. Toueg. Simulating authenticated broadcasts to derive simple fault-tolerant algorithms. *Distributed Computing*, 2(2):80–94, 1987.
- [53] Sun Microsystems Inc. RFC 1094: Network file system protocol specification. <http://www.ietf.org/rfc/rfc1094.txt>, Mar 1989.
- [54] M. Waldman and D. Mazières. Tangler: a censorship-resistant publishing system based on document entanglements. In *Proc. 8th ACM conf. on Comp. and Comm. Security*, pages 126–135, 2001.
- [55] K. Walsh and E. G. Sirer. Experience with an object reputation system for peer-to-peer filesharing. In *Proc. NSDI*, May 2006.
- [56] J. Yin, J.-P. Martin, A. Venkataramani, L. Alvisi, and M. Dahlin. Separating agreement from execution for Byzantine fault tolerant services. In *Proc. SOSP*, pages 253–267, Oct 2003.
- [57] A. R. Yumerefendi and J. S. Chase. Trust but verify: Accountability for internet services. In *ACM SIGOPS European Workshop*, 2004.
- [58] A. R. Yumerefendi and J. S. Chase. The role of accountability in dependable distributed systems. In *Proceedings of HotDep*, Jun 2005.
- [59] A. R. Yumerefendi and J. S. Chase. Strong accountability for network storage. In *FAST*, 2007.