

Byzantine Fault-Tolerant Confidentiality

Jian Yin Arun Venkataramani Jean-Philippe Martin Lorenzo Alvisi Mike Dahlin

1 Introduction

As the world becomes increasingly interconnected, more and more important services such as business transactions are deployed as *access anywhere services* – services that are accessible by remote devices through the Internet and mobile networks. Such services often must access confidential data to provide service. For example, an online bank service must access a user’s checking account to process an online transfer request. In such a scenario, guarantees of availability, integrity, and confidentiality are essential. By availability, we mean that services must provide service 24/7 without interruption. By integrity, we mean that services must process clients’ requests correctly. By confidentiality, we mean that services must restrict who sees what data.

Unfortunately, software for providing access anywhere services may contain bugs, and hackers may exploit these bugs to disrupt service or steal confidential data. One solution is to eliminate bugs by using better software engineering practice or formal methods [10]. However, the complexity of real world systems and application software makes it difficult to produce bug-free implementations with existing techniques.

A promising approach is to use redundancy to harden services against bugs. A traditional Byzantine fault-tolerant (BFT) system runs different implementations of the same service on several replicas and ensures that correct computation is performed by enough correct replicas to mask incorrect replicas [1, 2, 12, 13, 15]. Recent research has shown that BFT systems can be practical for several important services as they can be implemented with low overheads compared to the unreplicated services [3, 4, 14]. However, although traditional BFT systems improve availability and integrity through redundancy, existing BFT architectures make such systems more prone to compromising confidentiality. In a traditional BFT system, replicas send all replies directly to clients. So if a hacker manages to compromise even one replica, he can steal confidential data. Thus, there is a fundamental tradeoff between increasing availability and integrity on one hand and strengthening confidentiality on the other: the former is typically achieved by increasing the number of replicas and the latter is typically achieved by reducing the number of machines that have access to confidential data.

This paper discuss how to use redundancy to simultaneously improve availability, integrity, and confidentiality. We propose an architecture for such a system - Confidential BFT

(CBFT). In CBFT, service replicas connect to a firewall and can send messages to the outside world only through it. The firewall runs a majority voting algorithm just like the clients of a traditional BFT system. It filters out faulty messages that may contain confidential data.

The firewall system has to be correct to provide confidentiality. Even though the firewall is simple, building a formally verified bug-free firewall may not be feasible. However, redundancy can be used to improve the robustness of the firewall. Such a firewall system consists of a group of nodes that are interconnected such that any path from a service replica to the outside world is longer than a threshold, f . Thus, any communication from any service replica to the outside world must go through at least one correct node as long as there are fewer than $f + 1$ faulty firewall nodes. Moreover, a correct node in a firewall chain can independently ensure that a unique sequence of replies results from a sequence of requests just as if this sequence of requests were processed by a single correct server. Thus, faulty machines are prevented from using steganography to leak confidential data.

In summary, this paper investigates how to build available, high-integrity, and confidential access anywhere services by using redundancy and outlines the architecture of one implementation of such an architecture, CBFT.

2 Background

There is much previous work on how to use redundancy to tolerate Byzantine faults. This work falls into two categories: i) using client voting to improve availability and integrity without considering confidentiality and ii) using secret sharing to improve confidentiality for a limited class of services that do not manipulate confidential data.

Redundancy techniques to improve availability and integrity without considering confidentiality have been extensively studied. Typically in such systems, all service replicas send replies to clients, and voting takes place at clients. There is a significant body of work on BFT quorum systems [1, 12, 13] in which a group of service replicas use intersection. properties of quorum sets to guarantee that clients retrieve values of variables written previously. Byzantine fault tolerance for arbitrary services captured by state machines have been studied in both theoretical and practical settings [2, 3, 4, 14, 15].

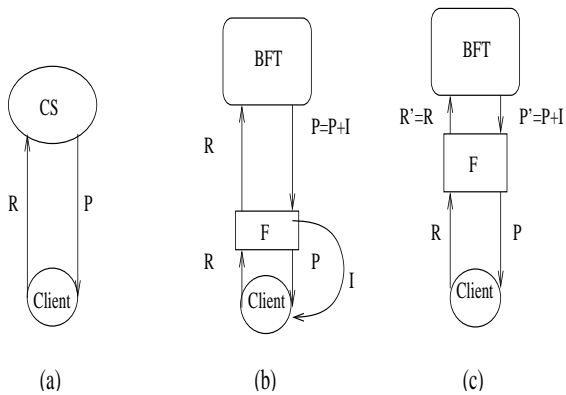


Figure 1: Fault-tolerance models

In particular, Castro and Liskov [3, 4, 14] have shown that BFT systems can be practical because i) the performance of several BFT services implemented by them are comparable to their unreplicated counterparts [3] and ii) multiple implementations of the same legacy services by different vendors can be used to achieve failure independence of replicas [14].

There are also several studies [5, 6, 7, 9, 11, 16, 17] on sharing secrets among several service replicas that prevent fewer than some threshold number of service replicas from compromising secrecy. However, hiding the unencrypted value of confidential data from servers effectively prevents servers from manipulating confidential data. Thus, such techniques can only be applied to a limited class of services in which the servers only retrieve, use, or transmit opaque data such as storage systems, file systems, distributed certification authority, and message transmissions.

3 A CBFT system

The traditional state machine approach for implementing Byzantine fault-tolerant services delegates the responsibility of combining the outputs of the ensemble of state machine replicas to the client. To quote Schneider [15], *“the voter - a part of the client - is faulty exactly when the client is, so the fact that an incorrect output is read by the client due to a faulty voter is irrelevant.”* Though such a model is appropriate for services deployed locally, it fails to address confidentiality for access anywhere services. In particular, if a hacker manages to compromise one replica in such a system, the system has no mechanisms to prevent the compromised replicas from sending confidential data back to the hacker.

As shown in Figure 1(b), traditional approaches to tolerate Byzantine faults in replicas attempt to emulate the abstract notion of a single correct server (CS) as in (a). Here R is a request sent by the client and P the response the client receives, and P' the set of responses filtered by the voting

component F to eliminate incorrect information I received by F from faulty replicas in the BFT. I can be used by a malicious replica as a channel to transmit confidential data. The obvious solution is to move the voter, F , away from clients into the service provider’s domain as shown in (c). Unfortunately, this is challenging since F is a potential point of vulnerability for availability, integrity, and confidentiality. In particular, Schneider’s “fault sharing” argument above no longer applies – faults in the voter F can hurt availability and integrity of correct clients. Furthermore, a fault in the voter can reintroduce a covert channel that allows the information I to escape to clients.

One solution is to appeal to software engineering and formal methods to construct a perfect voter. Another approach is to use redundancy to construct a highly reliable voter out of imperfect components. To understand the interplay of redundancy and confidentiality, our goal is to design a voting component that acts as a “firewall” to filter responses. This component is comprised of individual units, each of which is susceptible to Byzantine faults, but the voter as a whole protects confidentiality without compromising availability. This is a hard problem, especially in the face of sophisticated attacks that could compromise confidential information by delaying certain replies, omitting requests, spoofing requests not generated by clients, etc.

In order to tolerate faults in the voter, the voter itself must be replicated, otherwise a single failure – of the voter – can prevent the system from producing the correct output or protecting confidentiality. This additional level of replication threatens to mire one recursively in the original problem. In our firewall system, we solve this problem by i) ensuring that every response from the server replica set is verified by at least one correct voter and ii) ensuring that each voter can independently verify the correctness of any reply.

To provide intuition, we first describe a simplified firewall that filters some incorrect information I , but is vulnerable to poor availability and in-band signaling that can leak confidential data. We then outline how to address these limitations. The system consists of two components as illustrated in Figure 2(a):

- a) A backend BFT consisting of set of replicated servers U to tolerate s Byzantine faults.
- b) A firewall network F , organized as a chain of $f + 1$ machines, at most f of which are faulty. Voters can communicate directly only with their neighbors. Only f_c at the bottom of the chain can communicate directly with clients, and the topmost machine f_s alone can communicate with the server replicas.

Clients send encrypted, signed requests to f_c . Each machine in the chain forwards the request up the chain towards f_s which broadcasts it to all servers in U . Servers in U are




Figure 2: Confidential BFT

expected to respond with encrypted, signed replies to f_s . On receiving $s + 1$ identical replies with verifiable correct signatures for a given request, f_s forwards the reply and the set of $s + 1$ signatures to the machine just below it in the chain. Each correct firewall thereafter independently verifies the $s + 1$ signatures and forwards them further down the chain, until finally f_c sends them to the client. If any firewall other than f_s detects that the replies do not match, it discards the reply and triggers a fault-detection alarm.

P1: Since every response is verified by $f + 1$ machines in the firewall (at least one of which is correct) before reaching the client, the client only receives correct replies.

The system described above, though simple has the following drawbacks:

1. The system has no availability guarantee. Even if one of the machines in the firewall is compromised, it could drop requests or replies.
2. A faulty firewall machine might leak confidential information by altering the membership of the set of correct signatures vouching for a reply.
3. A faulty firewall machine might leak confidential information by inserting, omitting, or reordering requests or replies.
4. The system is susceptible to timing attacks, i.e., the malicious servers and/or the machines in the firewall could use the delay in sending responses to encode confidential information.

Availability of the system can be improved simply by augmenting the firewall with multiple independent chains each $f + 1$ in length as shown in Fig. 2(b). Assuming at most f faulty machines, $f + 1$ such chains are sufficient to guarantee 100% availability. Since P1 still holds, the client can choose any reply that it receives. In general if v machines are susceptible to crash failures, it is a straightforward graph-theoretical problem to prove that a minimum of $(f + 1)(v + 1)$

machines are necessary to protect confidentiality and ensure availability.

The second problem can be easily rectified by using threshold signatures [8] with a threshold of $s + 1$. Each of the replicas sends its reply along with its partial signature to f_s . f_s waits for $s + 1$ matching replies, combines them to form just one message and one signature, and forwards them down the chain. Each machine now independently verifies the reply before sending it further down.

One solution to (3) is for each firewall machine to ensure that it transmits replies in exactly the order that it receives requests. If there exists at least one correct firewall CF on a data path, then clients or firewall machines nearer to clients than CF can gain no confidential information from request ordering, insertion, or omission.

Timing attacks are difficult to tackle in a airtight manner without affecting the performance of the system. Timing attacks can be countered in two ways: i) we can reduce the ability of faulty nodes to affect latency, for example by having correct nodes insert delays deliberately, ii) we can examine unusual latency fluctuations at some machines to detect possible intrusion and recover proactively.

With the above additions, it can be shown that the CBFT system emulates the single correct server abstraction as depicted in Fig. 1. We show the equivalence in a single chain firewall by considering a sequence of requests seen at the first correct machine from the bottom in a chain and by showing that it produces the unique sequence of replies that the abstract single correct server would generate. By restricting the flexibility of the information sent out by a correct machine in the firewall we remove the means for malicious machines to hide confidential information therein, thereby preventing steganography. It is straightforward to extend the argument to multiple chains.

4 Conclusion

In this abstract we argue that availability, integrity, and confidentiality are essential for access anywhere services and that traditional assumptions behind Byzantine Fault Tolerant state machine replication do not protect confidentiality when clients are not trusted. A Confidential BFT system (CBFT) differs from traditional BFT systems in that it increases confidentiality guarantees instead of reducing them, while maintaining the availability and integrity guarantees provided by the redundancy. The approach we propose in this paper is general because it applies to any state machine. We are currently in the process of implementing a prototype of CBFT to evaluate its performance.

References

- [1] R. Bazzi. Synchronous Byzantine quorum systems. In *Proceedings of the 16th Annual ACM Symposium on Principles of Distributed Computing (PODC '97)*, August 1997.
- [2] G. Bracha and S. Toueg. Asynchronous consensus and broadcast protocols. In *Journal of the Association for Computing Machinery*, October 1995.
- [3] M. Castro and B. Liskov. Practical Byzantine fault tolerance. In *Proceedings of the Third Symposium on Operating Systems Design and Implementation*, February 1999.
- [4] M. Castro and B. Liskov. Proactive recovery in a Byzantine-fault-tolerant system. In *Proceedings of the 4th Symposium on Operating Systems Design and Implementation*, October 2000.
- [5] Y. Deswarte, L. Blain, and J. Fabre. Intrusion tolerance in distributed systems. In *Proc. Symp. on Research in Security and Privacy*, pages 110–121, Oakland, CA, USA, 1991. IEEE Computer Society Press.
- [6] M. Franklin and R. N. Wright. Secure communication in minimal connectivity models. In *Advances in Cryptology – EUROCRYPT '98*, pages 346–360, 1998.
- [7] M. Franklin and M. Yung. Secure hypergraphs: Privacy from partial broadcast (extended abstract). In *Proceedings of the Twenty-Seventh Annual ACM Symposium on the Theory of Computing*, pages 36–44, 29 May–1 June 1995.
- [8] R. Gennaro, S. Jarecki, H. Krawczyk, and T. Rabin. Robust threshold DSS signatures. In *Theory and Application of Cryptographic Techniques*, pages 354–371, 1996.
- [9] R. Gennaro, T. Rabin, S. Jarecki, and H. Krawczyk. Robust and efficient sharing of RSA functions. *Journal of Cryptology*, 13(2):273–300, Spring 2000.
- [10] M. Kaufmann, P. Manolios, and J. Strother Moore. *Computer-Aided Reasoning: An Approach*. Kluwer Academic Publishers, June 2000.
- [11] J. Kubiawicz, D. Bindel, P. Eaton, Y. Chen, D. Geels, R. Gummedi, S. Rhea, W. Weimer, C. Wells, H. Weather- spoon, and B. Zhao. OceanStore: An architecture for global-scale persistent storage. *ACM SIGPLAN Notices*, 35(11):190–201, November 2000.
- [12] D. Malkhi and M. Reiter. Byzantine quorum systems. In *Distributed Computing*, 1998.
- [13] D. Malkhi, M. Reiter, and A. Wool. The load and availability of Byzantine quorum systems. In *SIAM Journal on Computing*, December 2000.
- [14] R. Rodrigues, M. Castro, and B. Liskov. BASE: Using abstraction to improve fault tolerance. In *Proceedings of the 18th ACM Symposium on Operating Systems Principles*, 2001 October.
- [15] F. Schneider. Implementing Fault-tolerant Services Using the State Machine Approach: A tutorial. *Computing Surveys*, 22(3):299–319, September 1990.
- [16] T. Wu, M. Malkin, and D. Boneh. Building intrusion-tolerant applications. In *Proceedings of the 8th USENIX Security Symposium (SECURITY-99)*, pages 79–92, Berkeley, CA, August 23–26 1999. Usenix Association.
- [17] L. Zhou, F. Schneider, and R. Renesse. COCA: A secure distributed on-line certification authority. In *Technical Report TR2000-1828*, Computer Science Department, Cornell University, 2000.