2011 Sixth IEEE International Conference on Networking, Architecture, and Storage

Secure Fragment Allocation in a Distributed Storage System with Heterogeneous Vulnerabilities

Yun Tian, Shu Yin, Jiong Xie
Ji Zhang, Xiao Qin
Department of Computer Science
and Software Engineering
Auburn University
Auburn, Alabama 36849-5347
http://www.eng.auburn.edu/∼xqin

Mohammed I. Alghamdi
Al-Baha University
Al-Baha City
Kingdom of Saudi Arabia
Email: mialmushilah@bu.edu.sa

Meikang Qiu
Department of Electrical
and Computer Engineering
University of Kentucky
Lexington, Kentucky
Email: mqiu@engr.uky.edu

Yiming Yang
Intel Corporation
NM 87124
Email: yiming.yang@intel.com

I. Abstract

***Abstract*—There is a growing demand for large-scale distributed storage systems to support resource sharing and fault tolerance. Although heterogeneity issues of distributed systems have been widely investigated, little attention has yet been paid to security solutions designed for distributed storage systems with heterogeneous vulnerabilities. This fact motivates us to investigate a fragment allocation scheme called S-FAS to improve security of a distributed system where storage sites have a wide variety of vulnerabilities. In the S-FAS approach, we integrate file fragmentation with the secret sharing technique in a distributed storage system with heterogeneous vulnerabilities. Storage sites in a distributed systems are classified into a variety of different server types based on vulnerability characteristics. Given a file and a distributed system, S-FAS allocates fragments of the file to as many different types of nodes as possible in the system. Data confidentiality is preserved because fragments of a file are allocated to multiple storage nodes. We develop storage assurance and dynamic assurance models to evaluate the quality of security offered by S-FAS. Analysis results show that fragment allocations made by S-FAS lead to enhanced security because of the consideration of heterogeneous vulnerabilities in distributed storage systems.**

II. Introduction

A. Security Problems in Distributed Systems

There is an increasing demand to develop large-scale distributed storage systems supporting data-intensive services that provide resource sharing and fault tolerance. The confidentiality of security-sensitive files must be preserved in modern distributed storage systems, because distributed systems are exposed to an increasing number of attacks from malicious users [9].

Although there exist many security techniques and mechanisms (for example, [6] and [17]), it is quite challenging to secure data stored in distributed systems. In general, security mechanisms need to be built for each component in a distributed system, then a secure way of integrating all the components in the system must be implemented. It is critical and important to maintain the confidentiality of files stored in a distributed storage system when malicious programs and users compromise some storage nodes in the system.

In addition to cryptographic systems, secret sharing is an approach to providing data confidentiality by distributing a file among a group of n storage nodes, to each of which a fragment of the file is allocated. The file can be reconstructed only when a sufficient number (e.g., more than k) of the fragments are available to legitimate users. Attackers are unable to reconstruct a file using the compromised fragments, if a group of servers are compromised and fewer than k fragments are disclosed.

B. Heterogeneous Vulnerabilities

Although heterogeneity issues of distributed systems have been widely investigated, little attention has been paid to security solutions designed for distributed storage systems with heterogeneous vulnerabilities. This problem motivates us to focus on heterogeneity issues concerning security mechanisms of distributed storage systems.

In a large-scale distributed system, different storage sites have a variety of ways to protect data. The same security policy may be implemented in various mechanisms. Data encryption schemes may vary; even with the same encryption scheme, key lengths may vary across the distributed system. The above mentioned factors can contribute to different vulnerabilities among storage sites. Although security mechanisms deployed in multiple storage sites can be implemented in a homogeneous way, different vulnerabilities may exist due to heterogeneities in computational units.

We start to address security heterogeneity issues by dividing storage servers into different server-type groups. Each server type represents a level of security vulnerability. In a server-type group, storage servers with the same vulnerability share the same weakness that allows attackers to reduce the servers' information assurance. Although it may be difficult to classify all servers in a system into a large number of groups, a practical way of identifying server types is to organize these with similar vulnerabilities into one group.

In light of the server types and heterogeneous vulnerabilities, we investigate, in this study, a fragment allocation scheme called S-FAS in order to improve security of a distributed system where storage sites have a wide variety of vulnerabilities.

978-0-7695-4509-7/11 $26.00 © 2011 IEEE
DOI 10.1109/NAS.2011.14

C. File Fragmentation and Allocation

The file fragmentation technique is often used in many distributed and parallel systems to improve availability and performance. Several file fragmentation schemes have been proposed to achieve high assurance and availability in a large distributed system [7][15]. In real-world distributed systems, the fragmentation technique is usually combined with replication to achieve better performance at the cost of increased security risk to data stored in the systems. A practical distributed system normally contains multiple heterogeneous servers providing services with various vulnerabilities. Unfortunately, the existing fragmentation algorithms do not take the heterogeneity issues into account.

To address the above mentioned limitations, we focus on the development of a file fragmentation and allocation approach to improving the assurance and scalability of a heterogeneous distributed system. If one or more fragments of a file have been compromised, it is still very hard for a malicious user to reconstruct the file from the compromised fragments. Our solution is different from those previously explored, because our approach utilizes heterogeneous features regarding vulnerabilities among servers.

To evaluate our method for fragment allocations, we develop static and dynamic assurance models to quantify the assurance of a heterogeneous distributed storage system handling data fragments. Experimental results show that increasing heterogeneity levels can improve file assurance in a distributed storage system.

D. Main Contributions

The following are four main contributions that we have made with this study:

- We address the heterogeneous vulnerability issue by dividing storage nodes of a distributed system into different server-type groups based on their vulnerabilities. Each server-type group - representing a level of vulnerability - contains storage nodes with the same security vulnerability.
- We propose a secure fragmentation allocation scheme called S-FAS to improve security of a distributed system where storage nodes have a wide variety of vulnerabilities.
- We develop storage assurance and dynamic assurance models to quantify information assurance and to evaluate the proposed S-FAS scheme.
- We discover principles to improve assurance levels of heterogeneous distributed storage systems. The principles are general guidelines to help designers achieve a secure fragment allocation solution for distributed systems.

E. The Organization of this Paper

The rest of the paper is organized as follows: In Section III, we review the related work. Section IV presents the system and threat models of this study. Section V describes S-FAS - a secure fragmentation allocation scheme. In Section VI, we develop an assurance model and a dynamic assurance model for distributed storage systems. In Sections VII, we quantitatively evaluate the proposed S-FAS scheme in the context of distributed systems. Section VIII summarizes this paper and outlines our future work.

III. Related Work

A. Security Techniques for Distributed Systems

Much research has been performed to improve security of distributed and high-performance computing systems such as Grids. For example, Pourzandi *et al.* proposed a structured security approach that incorporates both distributed authentication and distributed access control mechanisms [9].

Intrusion detection techniques have been widely used to provide basic assurance of security in distributed systems. However, most intrusion detection techniques are inadequate to protect data stored in distributed systems [3]. One of the most effective approaches to improving information assurance in distributed systems is intrusion tolerance [2] [13] [15]. To enhance security assurance, researchers have developed a range of intrusion-tolerant tools and mechanisms. The fragmentation technique summarized below is one of the intrusion tolerance methods that can be used in combination with intrusion detection techniques.

B. Fragmentation Techniques

A fragmentation technique partitions a security sensitive file into multiple fragments that are distributed across different storage servers in a distributed system. A lot of fragmentation schemes have been proven to be valuable tools to improve security of data stored in distributed systems (see, for example, [14][4][5][7][10][11]).

Many fragmentation approaches aim to improve availability and performance of distributed systems by applying data replication methods. For example, Dabek *et al.* developed a wide-area cooperative storage system in which they implemented a fragmentation scheme to improve availability and to facilitate load balancing [1].

Although combining a fragmentation scheme and a replication scheme can enhance performance and availability, data replications may impose security risks due to an increasing number of file fragments handled by distributed storage servers. A file is more likely to be compromised when more replications of the file are stored in a distributed storage system.

All existing file fragmentation technologies are inadequate to address the issue of heterogeneous vulnerabilities in large-scale distributed systems. Our preliminary results show that security can be improved in a distributed storage system when a fragmentation scheme incorporates the heterogeneous-vulnerability feature.

C. Secret Sharing

Secret sharing - independently invented by Shamir and Blakley - is a method of distributing a secret among a group of participants, each of which is allocated a share of the secret. The secret can be successfully reconstructed only

when a sufficient number of shares are given and combined together [8][11].

Shamir proposed the (k, n) secret sharing scheme that divides data D into n pieces in such a way that D can be easily reconstructed from any k pieces. If fewer than k pieces are disclosed, no one can reconstruct D from the revealed pieces.

The (k, n) secret sharing scheme was proposed as a robust key management approach with $\mathrm{n} = 2\mathrm{k} - 1$. A key can be recovered even when $\lfloor n/2 \rfloor = k - 1$ of the n pieces that are destroyed. Attackers cannot reconstruct the key even when security breaches expose $\lfloor n/2 \rfloor = k - 1$ of the remaining k pieces [11].

The secret sharing scheme has been extended and employed in different application domains [12]. For example, Bigrigg *et al.* proposed an architecture called PASIS for secure storage systems. The PASIS architecture integrates the secret sharing scheme with information dispersal to improve security, integrity and availability [16][18]. In a storage system with PASIS, the confidentiality of data stored in the system is still preserved, even if an attacker compromises a limited (i.e., fewer than the threshold) subsets of storage nodes. The aforementioned secret-sharing solutions designed for distributed storage systems ignore the issue of heterogeneous vulnerabilities. This fact motivates us to extend the secret sharing scheme by considering heterogeneity in vulnerabilities, in the context of distributed storage systems.

D. Comparison of Our Work with Existing Solutions

Our fragment allocation solution we describe in this paper is entirely different from the existing fragment allocation schemes found in the literature. Our approach aims to incorporate the vulnerability heterogeneity feature of distributed storage systems into file fragment allocation. Our solution captures heterogeneous features regarding vulnerabilities of the nodes in order to improve the security level of the data stored in a distributed system. In this study, the data replication technique is not considered, because fragment allocation and data replication are independent of each other. Thus, the reliability and performance of fragment allocation schemes can be improved when data replication modules are integrated.

IV. System and Threat Model

We outline, in this section, the system and threat models that capture main characteristics of distributed storage systems. The system model is used as a basis to design the S-FAS fragmentation allocation scheme, whereas the threat model helps us identify vulnerabilities and certain potential attacks in distributed storage systems.

A. System Model

The S-FAS fragmentation allocation scheme was designed for a distributed storage system (see Fig. 1) where each storage site is a cluster storage subsystem. Different cluster storage subsystems may be connected within some subnetworks to form a larger scale distributed storage sysytem.

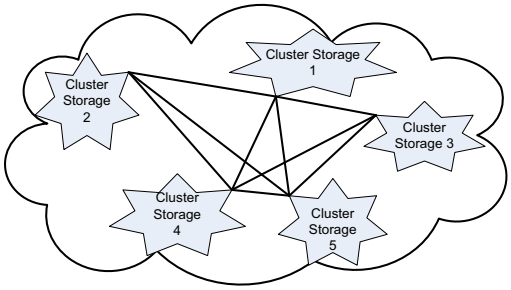

Fig. 1. A distributed storage system is comprised of a set of cluster storage subsystems. Multiple fragments of a file can be stored either in storage nodes within a single cluster storage subsystem or in nodes across multiple cluster storage subsystems. See Fig. 2 for details on a cluster storage subsystem.

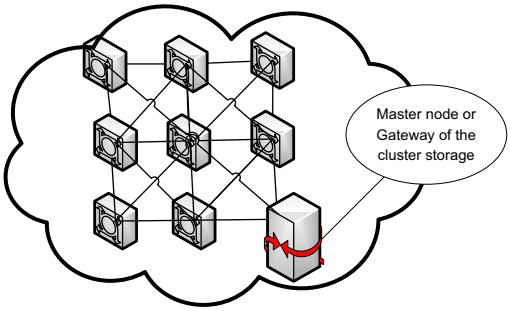

Fig. 2. A cluster storage subsystem consists of a number of storage nodes and a gateway. Storage nodes are divided into different server-type groups, each of which represents a level of security vulnerability.

Fig. 2 depicts a cluster storage subsystem, which consists of a number of storage nodes and a gateway. Considering heterogeneous vulnerability in large-scale storage systems, we divide storage nodes into different server-type groups, each of which represents a level of security vulnerability .

Before presenting details on the system model, let us summarize all notations used throughout this paper in Table I.

In this study, we consider a distributed storage system containing L cluster storage subsystems, i.e., $R_1, R_2, ..., R_L$. Cluster storage subsystems R_i consists of H_i storage nodes, i.e., $R_i = \{r_{i1}, r_{i2}, ..., r_{iH_i}\}$. All the storage nodes connected in cluster R_i have heterogeneous vulnerabilities.

Since all the nodes, including a master node, are fully connected in a cluster storage subsystem, we model the topology of a cluster storage system as a general graph. Cluster storage subsystem R_i has a gateway, which hides the cluster's internal architecture from users by forwarding file requests to storage nodes.

Data in cluster storage subsystem R_i can be accessed

TABLE I
NOTATION USED IN THE SYSTEM AND THREAT MODELS.

Notations	Meaning
N	Number of server nodes in the system
U	The whole system considered
L	Number of subsystems in the whole system
H_i	Number of server nodes in the subsystem i
F	A file stored in the system
F_i	Fragment i of file F
T_j	Server type j in the system
K	The total number of server types
S_j	The size of a certain server type in a cluster
m	Threshold for the secret sharing scheme
n	The total number of fragments for each file in the secret sharing scheme
R_i	Cluster storage subsystem
r_{ij}	Node j in subsystem i
X	The event that a set of storage nodes is chosen to be attacked
Y	The event that if X occurs, at least m fragments can be compromised using the same attack method.
Z	The event of a successful attack to a certain fragment of a file
V	The event file F is compromised under one attack method
P_N	The successful probability of an attack on a node
P_f	The successful probability to compromise a fragment in a compromised node
$P(X)$	The probability of event X occurring
$P(Y)$	The probability of event Y occurring
$P(Z)$	The probability of event Z occurring
$P(V)$	The probability of event V occurring
α	An allocation mapping of file F
$SA(\alpha)$	The storage assurance of an allocation mapping α of file F
$\mathrm{DA}(\alpha)$	The dynamic assurance of an allocation mapping α of file F
q	Number of fragments needed to reconstruct a file transmitted from outside of the subsystem
g	Number of fragments compromised out of the q fragments transmitted from outside of the subsystem
P_L	The probability that a fragment is intercepted during its transmission
P_D	The probability that a file F is intercepted because of the compromised transmitted fragments

through its master node. When a read request is submitted to cluster R_i, the master node is responsible for reconstructing file fragments and returning the file to users. When a write request of a file is issued, the master node updates all the fragments of the file.

Legitimate users access cluster storage subsystems through master nodes; malicious users may bypass the master nodes to access storage nodes without being authorized. See Section IV-B below for details on the threat model.

B. Threat Model

It is not reasonable to assume that if a malicious user breaks into a storage node, fragments of a file stored on the node are thereby compromised. Normally, a malicious user needs two steps to compromise fragments of a file stored on a server. First, the malicious user must successfully attack the server. Second, fragments are retrieved by the malicious user.

Let P_N be the probability that a storage server is successfully attacked; let P_f be the probability that authorized users retrieve fragments stored on the server, provided that the server has been compromised. We define event Z as a successful attack on a fragment (i.e., unauthorized disclosure of the fragment). Since the above two consecutive attack steps are independently, the probability that event Z occurs is a product of probability P_N and probability P_f. Thus, the probability that a fragment is disclosed to an unauthorized attacker can be expressed as:

$$\mathrm{P}\,(\mathrm{Z}) = P_N * P_f. \tag{1}$$

In a dynamic allocation environment, a malicious user can use a compromised node to collect other needed fragments of the file when the fragments are passing through the compromised node.

If encryption keys are disclosed to attackers, unauthorized interceptions of encrypted files stored on the attacked node may occur. Given two storage nodes with different vulnerabilities, successful attacks of the nodes are not correlated. This statement is true for many potential threats, because compromising one storage node does not necessarily lead to the successful attack of the second one.

V. S-FAS: A Secure Fragment Allocation Scheme

In this section, we first outline the motivation for addressing the heterogeneity issues in the vulnerability of distributed storage systems. Next, we describe a security problem addressed in this study. Last, we present a secure fragment allocation scheme called S-FAS for distributed storage systems.

A. Heterogeneity in the Vulnerability of Data Storage

Since the existing security techniques (see Section III) developed for distributed systems are inadequate for distributed systems with heterogeneity in vulnerabilities, the focus of this study is heterogeneous vulnerabilities in large-scale distributed storage systems. Vulnerabilities of storage nodes in a distributed system are heterogeneous in nature due to the following four main reasons. First, storage nodes have different ways to protect data. Second, a security policy can be implemented in a variety of mechanisms. Third, the key length of an encryption scheme may vary across multiple storage nodes. Fourth, heterogeneities exist in computational units of storage sites. We believe that future security mechanisms for distributed systems must be aware of vulnerability heterogeneities.

B. A Motivational Example

If the above heterogeneous vulnerability features are not incorporated into fragment allocation schemes for distributed storage systems, a seemingly secure fragment allocation decision can lead to a breach of data confidentiality. The following motivational example illustrates a security problem caused by ignoring vulnerability heterogeneities.

Let us consider a file F with three partitioned fragments: f_a, f_b, and f_c, and a distributed storage system (see Fig. 3) that contains 16 storage nodes divided into 4 server-type groups (or server groups for short), i.e., T_1, T_2, T_3, and T_4. Storage nodes in each server group offer similar services with the same level

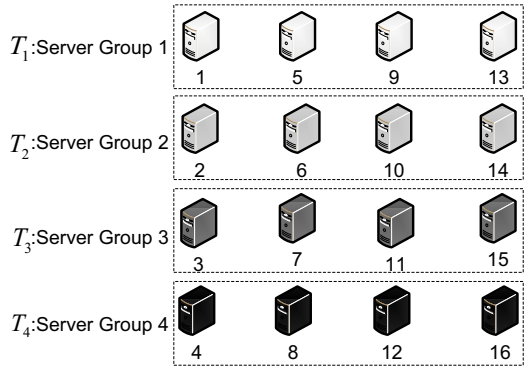

Fig. 3. A distributed storage system contains 16 storage nodes, which are divided into 4 server-type groups (or server groups for short), i.e., T_1, T_2, T_3, and T_4. Servers in each group have the same level of security vulnerability.

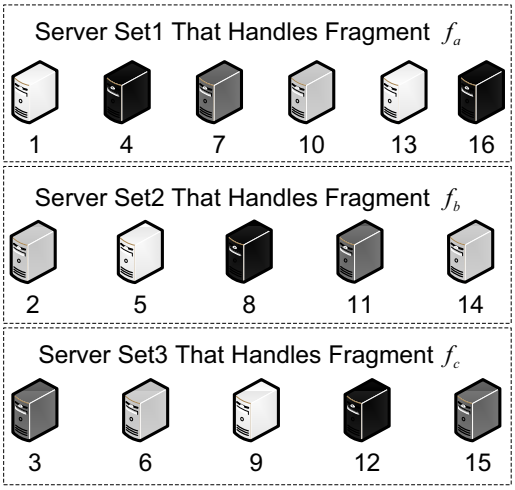

Fig. 4. Possible insecure file fragment allocation decision made using a hashing function (see Eq. 11 in [7]): Server set 1 handles fragment f_a, server set 2 handles fragment f_b, and server set 3 handles fragment f_c. Server set 1 contains storage nodes r_1, r_4, r_7, r_10, r_13, and r_16; server set 2 contains storage nodes r_2, r_5, r_8, r_{11}, and r_{14}; and server set 3 contains storage nodes r_3, r_6, r_9, r_{12}, and r_{15}. It is possible that fragments f_a, f_b, and f_c may be allocated to storage nodes that belong to the same server-type group. For example, the three fragments are respectively stored on nodes r_4, r_8, and r_{12}, which share the same vulnerability in server group T_4. Rather than three attacks, one successful attack against server group T_4 allows unauthorized users to access the three fragments of file F.

of vulnerability. In this example, server group T_1 consists of nodes r_1, r_5, r_9, r_{13}, i.e., $T_1 = \{r_1, r_5, r_9, r_{13}\}$. Similarly, we define the other three server groups as: $T_2 = \{r_2, r_6, r_{10}, r_{14}\}$, $T_3 = \{r_3, r_7, r_{11}, r_{15}\}$, and $T_4 = \{r_4, r_8, r_{12}, r_{16}\}$.

Fig. 4 shows that it is possible to make insecure fragment allocation decisions that do not take vulnerability heterogeneity into account. The decision made using a hashing function (see Eq. 11 in [7]) randomly allocates the three fragments of file F to three different nodes, each of which belongs to one of the three server sets illustrated in Fig. 4. For example, the three fragments f_a, f_b, and f_c are stored on nodes r_1, r_6, and r_8, respectively. This fragment allocation happens to be a good solution, because r_1, r_6, and r_8 have different vulnerabilities as the three nodes belong to different server groups (i.e., T_1, T_2, and T_4). A malicious user must launch three successful attacks (one for each server group) in order to compromise all three fragments.

The above fragment allocation scheme fails to address the threat described in Section IV. This is because an attacker can first retrieve one fragment of F by compromising a single node, then the attacker simply waits for the other two fragments to be passed through the compromised node. To solve this security problem, Zanin *et al.* developed a static algorithm to decide whether a particular storage node is authorized to handle a file fragment of F [18]. Zanin's algorithm can generate an insecure fragment allocation because heterogeneous vulnerabilities are not considered. For example, the three fragments are respectively stored on nodes r_4, r_8, and r_{12}, which share the same vulnerability in server group T_4 (see Fig. 4). Rather than three attacks, one successful attack against server group T_4 allows unauthorized users to access the three fragments of file F. Two other insecure fragment allocations are: (1) allocating f_a, f_b, f_c to nodes r_1, r_5, and r_9, respectively; and (2) allocating f_a, f_b, f_c to nodes r_7, r_{11} and r_{15}, respectively. These three fragment allocation decisions are unacceptable, because the fragments are assigned to a group of storage nodes with the same vulnerability, meaning that an attacker who comprised one node within a group can easily compromise the other nodes in the group. The attacker can reconstruct F from f_a, f_b, and f_c stored on the comprised server group.

C. Design of the S-FAS Scheme

To solve the above security problem, we have to incorporate vulnerability heterogeneities into fragment allocation schemes. Specifically, we design a simple yet efficient approach to allocating fragments of a file to storage nodes with various vulnerabilities. Since allocating fragments of a file into different storage clusters can degrade performance, our S-FAS scheme attempts to allocate fragments to storage nodes within a cluster. If the number of nodes with different vulnerabilities cannot meet the aforementioned criterion, file fragments must be allocated across multiple clusters. To improve the assurance of a distributed storage system while maintaining high I/O performance, each cluster storage subsystem has to be built with high vulnerability heterogeneity. This causes the fragments of a file to be less likely distributed across multiple storage clusters.

Because of the following two reasons, the S-FAS scheme can significantly improve data security when fragments are stored in a large-scale distributed storage system. First, S-FAS integrates the fragmentation technique with secret sharing. Second, S-FAS addresses the issue of heterogeneous vulnerabilities when file fragments are allocated to a distributed storage system.

The S-FAS scheme makes fragment allocation decisions by following the four policies below:

- **Policy 1:** All the storage nodes in a distributed storage system are classified into multiple server-type groups (server group for short) based upon their various vulnerabilities. Each server group consists of storage nodes with the same vulnerability level.
- **Policy 2:** To improve security of a distributed storage system, S-FAS allocates fragments of a file to storage nodes belonging to as many different server groups as possible. In doing so, it is impossible to compromise the file's fragments using a single successful attack method.
- **Policy 3:** The fragments of a file are trying to be allocated to nodes with a wide range of vulnerability levels all within a single cluster storage subsystem. The goal of this policy is to improve performance of the storage system by making the fragments less likely to be distributed across multiple clusters.
- **Policy 4:** The (m, n) secret sharing scheme is integrated with the S-FAS allocation mechanism.

If a file's fragment-allocation decisions are guided by the above four policies, successful attacks against less than m server groups have little chance to gain unauthorized accesses of files stored in a distributed system. In other words, if the number of compromised fragments of a file is less than m, attackers are unable to reconstruct the file from the fragments that are accessed by the unauthorized attackers. The S-FAS scheme can improve information assurance of files stored in a distributed storage system without enhancing confidentiality services deployed in cluster storage subsystems of the distributed system, because S-FAS is orthogonal to security mechanisms that provide confidentiality for each server group in a distributed storage system. Thus, S-FAS can be seamlessly integrated with any confidentiality service employed in distributed storage systems in order to offer enhanced security services.

VI. Assurance Models

We developed assurance models to quantitatively evaluate the security of a heterogeneous distributed storage system in which S-FAS handles fragment allocations.

A. Storage Assurance Model

For encrypted files, their encryption keys are partitioned and allocated using the same strategy that handle file fragments. Once a storage node in set U is compromised, file fragments and encryption key fragments stored on the node are both breached. If a malicious user wants to crack a file, at least m nodes within U must be successfully attacked.

We first investigate the probability that a file is compromised using one attack method. Let X be the event that a set of storage nodes is chosen to be attacked. Let Y be the event that if X occurs, at least m fragments can be compromised using the same attack method. As we already defined in Section IV, event Z represents a successful attack to a certain fragment of a file. Applying the multiplication principle, we calculate the probability that V - an event that file F is compromised under one attack - occurs as:

$$\mathrm{P(V)} = \sum_{\mathrm{j=1}}^{\mathrm{k}} \mathrm{P(X)P(Y)P(Z)} \tag{2}$$

where $P(X)$, $P(Y)$ and $P(Z)$ are probabilities that events X, Y and Z occur when the total number of different server-type groups (server group for short) is K. The probability $P(V)$ is proportional to probability $P(Z)$, which largely depends on the quality of security mechanisms deployed in the storage system, as well as the attacking skills of hackers.

Note that when k equals 1, there is no vulnerability difference among storage nodes. Supposing that all the fragments of a file can be compromised using one successful attack method, the probability that Y occurs becomes 1. Then, we can express $P(V)$ as:

$$\mathrm{P(V)} = \sum_{\mathrm{j=1}}^{\mathrm{k}} \mathrm{P(X)P(Z)} \tag{3}$$

Let S_j be the number of storage nodes in $server\ type\ T_j$ set and N be the total number of nodes in a distributed system. The probability that nodes in set T_j are randomly attacked can be derived as $\mathrm{P(X)} = \frac{\mathrm{S_j}}{\mathrm{N}}$.

Probability $P(Y)$ in Eq. 2 can be calculated as follows:

$$\mathrm{P\,(Y)} = \sum_{\mathrm{i=m}}^{\mathrm{n}} \frac{C_{S_j}^{i} C_{N-S_j}^{n-i}}{C_N^n}, (j = 2, \ldots K) \tag{4}$$

where C_N^n is the total number of possibilities of allocating fragments of a file, and the product of $C_{S_j}^i$ and $C_{N-S_j}^{n-i}$ is the number of possibilities that a file is compromised using a successful attack method which means at least m (It may be $m+1$, $m+2$, ..., n) fragments of the file are compromised.

To simplify the model, one may assume that security mechanisms and attacking skills have no significant impacts on information assurance of the entire distributed storage system. This assumption is reasonable because of two factors. First, S-FAS is independent of security mechanisms that provide confidentiality for server groups in a distributed storage system. Second, if empirical studies can provide values for probability $P(Z)$, the probability $P(V)$ can be derived from $P(Z)$ and the model (see Eq. 4) that calculates $P(Y)$. Since the study of the distribution of $P(Z)$ is not within the range of this work, in Section VII the impact of probability $P(Z)$ on $P(V)$ is ignored by setting the value of $P(Z)$ to 1.

Now we can derive Eq. 2 from Eq. 4 as below:

$$\mathrm{P(V)} = \sum_{\mathrm{j=1}}^{\mathrm{K}} \left(\frac{\mathrm{S_j}}{\mathrm{N}} \mathrm{P(Z)} \sum_{\mathrm{i=m}}^{\mathrm{n}} \frac{\mathrm{C_{S_j}^{i}\, C_{N-S_j}^{n-i}}}{\mathrm{C_N^n}} \right) \tag{5}$$

The confidentiality of file F is assured if F is not compromised. Thus, we can derive the assurance $SA(\alpha)$ of the storage system from Eq. 5 as:

$$\begin{aligned} \mathrm{SA\,(\alpha)} &= 1 - \mathrm{P\,(V)} \\ &= 1 - \sum_{\mathrm{j=1}}^{\mathrm{K}} \left(\frac{\mathrm{S_j}}{\mathrm{N}} \mathrm{P(Z)} \sum_{\mathrm{i=m}}^{\mathrm{n}} \frac{\mathrm{C_{S_j}^{i}\, C_{N-S_j}^{n-i}}}{\mathrm{C_N^n}} \right) \end{aligned} \tag{6}$$

B. Dynamic Assurance Model.

During read and write operations, some fragments of a file may be transmitted among different storage clusters or subnetworks. We assume that data transmissions within a cluster are secure, while connections among clusters and subnetworks may be insecure. Let P_L be the probability that a fragment is intercepted during its transmission on an insecure link. We consider a common case in which some fragments of file F are allocated outside a cluster. The probability P_D that a fragment of F is intercepted during its transmission can be expressed as:

$$\mathrm{P_D} = \mu_1\mu_2\mathrm{P_L} + \mu_3\left[1-\mathrm{P_L}\right]\mathrm{P_L} \tag{7}$$

where $\mu_1 = 1$ indicates that connections among storage clusters are insecure and $\mu_1 = 0$ means the connections are secure. $\mu_2 = 1$ indicates that fragments are transferred among different clusters, otherwise $\mu_2 = 0$. Similarly, $\mu_3 = 1$ means that fragments are transmitted across different subnetworks. When μ_1, μ_2, and μ_3 equal to 0, there is no fragment transmission risk. If q fragments need to be collected outside a cluster processing read/write operations, then probability $\mathrm{P_q(g)}$ that g out of q fragments are intercepted can be expressed as:

$$\mathrm{P_q\left(g\right)} = \mathrm{C_q^g P_D}^{\mathrm{g}}(1-\mathrm{P_D})^{\mathrm{q-g}} \tag{8}$$

Now we model the dynamic assurance of an allocation mapping α of file F. For simplicity, let us focus on a time period during which there is only one attempt to attack storage nodes where F is stored. During this time period, we assume that only one read or write operation is issued to access F. There are two cases where file F can be compromised. First, a malicious user can reconstruct F from m compromised fragments using the same attack method. Second, although less than m fragments are compromised, other g fragments are intercepted during their transmissions. Hence, we can derive the dynamic assurance $\mathrm{DA}(\alpha)$ from the storage risk (see Eq. 5) and the transmission risk (see Eq. 8), as shown here:

$$DA(\alpha) = 1- \left(P\left(V\right) + \left(\sum_{g=(m-i)}^{q} P_q(g)\right)\sum_{j=1}^{K}\left(\frac{S_j}{N}\times\sum_{i=0}^{m-1}\frac{C_{S_j}^{i}\,C_{N-S_j}^{n-i}}{C_N^n}\right)\right) \tag{9}$$

VII. Evaluation of System Assurance

The assurance models described in Section VI indicate that system assurance is affected by the number K of storage types, the number N of storage nodes in the system, and the number S_j of nodes in the jth storage type. In addition, threshold m and the number n of fragments in a file also have an impact on system assurance. Now, we quantitatively evaluate the impacts of these factors on the information assurance of distributed storage systems. We first obtain a comprehensive evaluation of S-FAS in terms of data storage assurance (see Sections VII-A to VII-E). Then, we consider dynamic assurance of S-FAS (see Sections VII-F and VII-G).

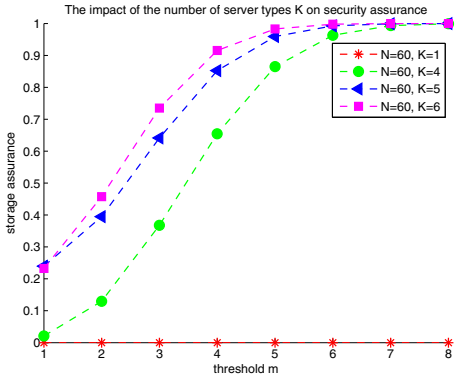

Fig. 5. *Heterogeneous system and homogeneous system using secret sharing scheme.* In all the four test cases, N is set to 60. K is set to 1, 4, 5, and 6, respectively. When K is 1, there is only one server group in the system.

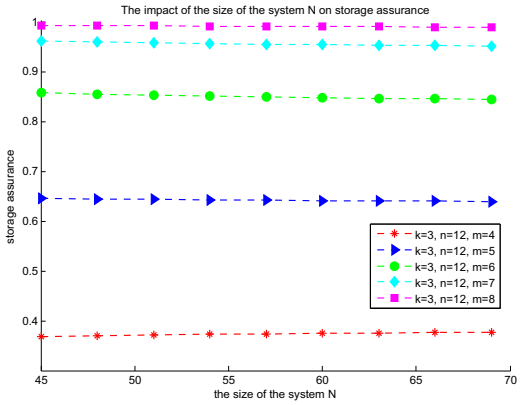

Fig. 6. *The impact of the system size N on storage assurance.*

We compare our approach with a traditional fragment allocation scheme that does not consider vulnerability heterogeneities. We evaluated a distributed storage system with the threshold value m. The default number n of fragments of a file is set to 12 and S_j=$\frac{N}{K}$ for all j from 1 to K.

A. Impact of Heterogeneity on Storage Assurance

If all storage nodes in the evaluated distributed system are identical in terms of vulnerability, the probability that fragments of a file can be compromised using one successful attack method is 1. Fig. 5 shows the impact of the number K of storage types on system assurance. Results plotted in Fig. 5 suggest that for a distributed system with homogeneous vulnerability, threshold m has no impact on system assurance. When it comes to a distributed system with heterogeneous vulnerabilities, the system assurance increases significantly with the increasing values of K and threshold m (see Fig. 5). Such a trend implies that a high heterogeneity level of vulnerability gives rise to high confidentiality assurance.

B. Impact of System Size on Storage Assurance

To quantify the impact of system size N on data assurance of a file stored in the system, we gradually increase system

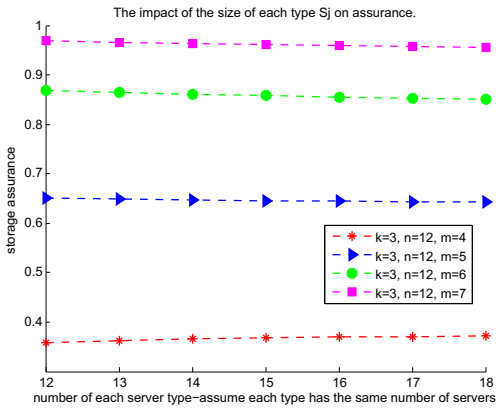

Fig. 7. The impact of server-group size on data storage assurance. The server-group size means the number of storage nodes in a server-type group. Note that the storage nodes within a server group share the same level of vulnerability. The server-group size varies from 12 to 18 with an increment of 1.

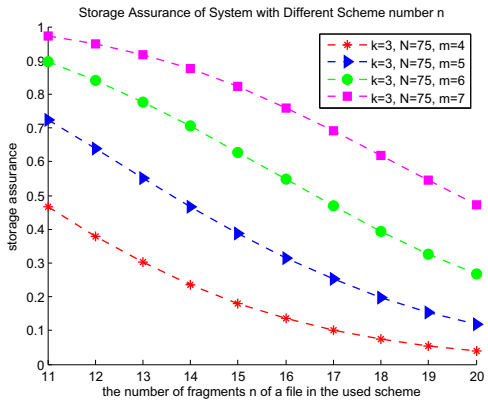

Fig. 8. The impact of the number n of fragments of a file on storage assurance. The number n of fragments increases from 11 to 20. The parameters k and N are set to 3 and 75, respectively.

size from 45 to 70 by increments of 5. We keep k at 3 and also vary m from 4 to 8. Fig. 6 reveals that the storage assurance of the system is not very sensitive to the system size, indicating that storage assurance largely depends on the vulnerability heterogeneity level rather than system size. Thus, large-scale distributed storage systems with low levels of vulnerability heterogeneities may not have higher assurance than small-scale distributed systems. These results suggest that one way to improve system assurance is to increase vulnerability heterogeneity while increasing the scale of a distributed storage system. A high heterogeneity level in vulnerability helps in increasing threshold m, making it harder for attackers to compromise multiple server groups and reconstruct files.

C. Impact of Size of Server Groups on Storage Assurance

Fig. 7 illustrates the impact of server-group size on data storage assurance. Note that the server-group size is the number of storage nodes in a server-type group, in which all the storage nodes share the same level of vulnerability. We vary the server-group size from 12 to 18 with an increment of 1. We observe from Fig. 7 that when threshold m is small (e.g, $m = 4$), the assurance of systems with large server-group sizes is slightly higher than that of systems with small server-group sizes. Interestingly, the opposite is true when the threshold m is large (e.g, $m > 4$). Given a fixed number of storage nodes in a distributed storage system, increasing the server-group size can decrease the number of server groups, which in turn tends to reduce vulnerability heterogeneity. The results shown in Fig. 7 match the results in the previous experiments in which a low level of vulnerability heterogeneity (or larger server-group sizes) results in degraded storage assurance.

D. Impact of Number n of File Fragments on Storage Assurance

Fig. 8 illustrates the impact of the number n of fragments of a file on storage assurance. In this experiment, we increase the number n of fragments from 11 to 20 and measured data storage assurance using our model. The parameters k and N are set to 3 and 75, respectively. We also vary threshold m from 4 to 7. Results depicted in Fig. 8 confirm that the system assurance is reduced with the increasing value of fragment number n. The results indicate that a large number of file fragments leads to low data storage assurance of the file. This assurance trend is reasonable because more fragments are likely to be allocated to storage nodes with the same vulnerability. If one storage node is compromised by an attacker, fragments stored on nodes with the same vulnerability can also be collected by the attacker, who is more likely to be able to reconstruct the file from the disclosed fragments.

In addition, Fig. 8 shows that increasing the value of threshold m can improve storage assurance. This pattern is consistent with the results obtained in the previous experiments.

E. Impact of Threshold m on Storage Assurance

Figs. 5-8 clearly show the impact of threshold m on storage assurance of a distributed system. More specifically, regardless of other system parameters, the storage assurance always goes up with the increasing threshold value m. The results indicate that the more fragments an attacker needs in order to reconstruct a file, the higher data storage assurance can be preserved for the file in distributed storage systems. These results suggest that to improve data storage assurance of a file, one needs to partition the file and allocate fragments in such a way that an attacker must compromise more server groups (the best case is m server groups) in order to reconstruct the file.

F. Impact of P_L on Dynamic Assurance

Now we are in a position to evaluate dynamic assurance of distributed storage systems. The three parameters μ_1, μ_2, and μ_3 in Eq. 7 have an important impact on dynamic assurance because these parameters indicate whether there is risk during fragment transmissions. Please refer to Sections VII-A to VII-E for details on the impacts of a set of parameters on data storage assurance.

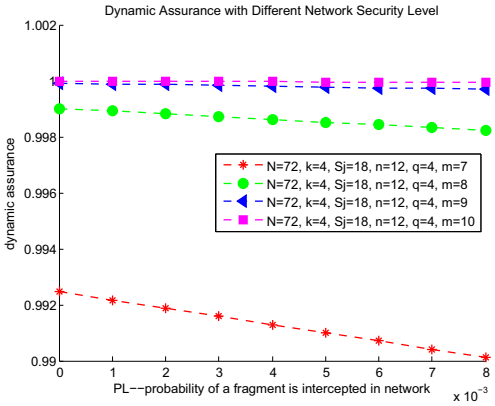

Fig. 9. Impact of P_L - the probability that a fragment might be intercepted by an attacker during the fragment's transmission through an insecure link. P_L is varied from 0 to $8 * 10^{-3}$ by increments of $1 * 10^{-3}$. Threshold m is varied from 7 to 10

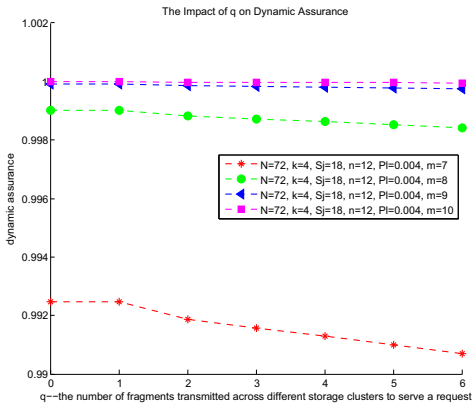

Fig. 10. Impact of q - the number q of fragments transmitted to and from a storage cluster. q is chosen from 0 to 6 with an increment of 1. Threshold m is set from 7 to 10)

P_L - the probability that a fragment might be intercepted by an attacker during the fragment's transmission through an insecure link - has a noticeable impact on dynamic assurance of a distributed storage system provided that threshold m is small (e.g., smaller than 9). Fig. 9 shows the dynamic assurance of a distributed system when P_L is varied from 0 to $8 * 10^{-3}$ by increments of $1 * 10^{-3}$. We also vary threshold m (i.e., m is varied from 7 to 10) to evaluate the sensitivity of dynamic assurance on parameter P_L under different threshold m.

Fig. 9 demonstratively confirms that when threshold m is equal to or smaller than 8, a large value of P_L results in low dynamic assurance of the system. The results are expected since a high value of P_L means that the transmitted fragments are likely to be intercepted by an attacker. Once the attacker has collected enough fragments of a security-sensitive file, the file could be reconstructed. When threshold m is larger than 8, the dynamic assurance is not noticeably sensitive to the probability P_L that a fragment is compromised during its network transfer.

G. Impact of q on Dynamic Assurance

Like parameter P_L, the number q of fragments transmitted to and from a storage cluster also has an impact on the dynamic assurance of a distributed storage system. Intuitively, Fig. 10 shows that when the number of fragments of a file that must be transmitted through insecure links is increasing, the dynamic assurance of the file drops. Interestingly, when threshold m is larger than 8, the dynamic assurance becomes very insensitive to the number q of fragments. This observation suggests that when the threshold is small, the S-FAS fragment allocation scheme must pay particular attention to lower the value of q in order to maintain a high dynamic assurance level.

In addition, we observe from Fig. 10 that dynamic assurance is always lower than the corresponding storage assurance (where q=0 in Fig. 10). This trend is always true because in a dynamic environment, file fragments have to be transmitted through insecure network links where malicious users may intercept the fragments in order to reconstruct files.

VIII. Conclusion and Future Work

It is critical to maintain the confidentiality of files stored in a distributed storage system, even when some storage nodes in the system are compromised by attackers. Secret sharing is an efficient way to preserve data confidentiality by distributing a file among a group of n storage nodes, to each of which a fragment of the file is stored. The file can be reconstructed from at least k fragments. If fewer than k fragments are disclosed to attackers, the file's confidentiality can still be preserved. In recognizing that storage nodes in a distributed system have heterogeneous vulnerabilities, we investigate a secure fragment allocation scheme by incorporating secret sharing and heterogeneous vulnerability to improve security of distributed storage systems.

We addressed the security heterogeneity issue by classifying storage servers into different server-type groups (or server group for short), each of which represents a level of security vulnerability. With heterogeneous vulnerabilities in place, we developed a fragment allocation scheme called S-FAS to improve security of a heterogeneous distributed system. S-FAS allocates fragments of a file in such a way that even if attackers compromised a number of server groups and fewer than k fragments are disclosed, the file cannot be reconstructed by the attackers from the compromised fragments.

To evaluate the S-FAS scheme, we built the static and dynamic assurance models in order to quantify the assurance of a heterogeneous distributed storage system processing file fragments. We demonstrate that S-FAS incorporates the vulnerability heterogeneity feature into file fragment allocation for distributed storage systems. Experimental results show that increasing heterogeneity levels can improve file assurance in a distributed storage system.

There are three future research directions of this study. First, we will make an effort to improve the performance of the S-FAS fragment allocation scheme in a heterogeneous distributed

system. Second, we will integrate the data replication technique with S-FAS to enhance reliability and performance of the fragment allocation scheme for distributed systems. Third, we will implement a distributed storage system prototype where S-FAS is deployed. In this prototype, we will evaluate performance of S-FAS in a real-world system.

ACKNOWLEDGMENT

The work reported in this paper was supported by the U.S. National Science Foundation under Grants CCF-0845257 (CAREER), CNS-0917137 (CSR), CNS-0757778 (CSR), CCF-0742187 (CPA), CNS-0831502 (CyberTrust), CNS-0855251 (CRI), OCI-0753305 (CI-TEAM), DUE-0837341 (CCLI), and DUE-0830831 (SFS), as well as Auburn University under a startup grant and a gift (Number 2005-04-070) from the Intel Corporation.

REFERENCES

[1] F. Dabek, M. F. Kaashoek, D. Karger, R. Morris, and I. Stoica. Wide-area cooperative storage with cfs. In *SOSP '01: Proceedings of the eighteenth ACM symposium on Operating systems principles*, pages 202–215, New York, NY, USA, 2001. ACM.

[2] Y. Deswarte, L. Blain, and J.-C. Fabre. Intrusion tolerance in distributed computing systems. In *Research in Security and Privacy, 1991. Proceedings., 1991 IEEE Computer Society Symposium on*, pages 110 –121, May 1991.

[3] D. L. Kewley and J. F. Bouchard. Darpa information assurance program dynamic defense experiment summary. *IEEE Trans. on Systems, Man and Cybernetics, Part A: Systems and Humans*, 31(4):331 –336, Jul 2001.

[4] J. Kubiatowicz, D. Bindel, Y. Chen, S. Czerwinski, P. Eaton, D. Geels, R. Gummadi, S. Rhea, H. Weatherspoon, W. Weimer, C. Wells, and B. Zhao. Oceanstore: an architecture for global-scale persistent storage. *SIGPLAN Not.*, 35:190–201, November 2000.

[5] S. Lakshmanan, M. Ahamad, and H. Venkateswaran. Responsive security for stored data. *IEEE Trans. on Parallel and Distributed Systems*, 14(9):818 – 828, 2003.

[6] H. Mantel. On the composition of secure systems. In *2002 IEEE Symposium on Security and Privacy.*, 2002.

[7] A. Mei, L. V. Mancini, and S. Jajodia. Secure dynamic fragment and replica allocation in large-scale distributed file systems. *IEEE Trans. on Parallel and Distributed Systems*, 14(9):885 – 896, Sept. 2003.

[8] T. P. Pedersen. Non-interactive and information-theoretic secure verifiable secret sharing. In *Proceedings of the 11th Annual International Cryptology Conference on Advances in Cryptology*, CRYPTO '91, pages 129–140, London, UK, 1992. Springer-Verlag.

[9] M. Pourzandi, D. Gordon, W. Yurcik, and G. A. Koenig. Clusters and security: distributed security for distributed systems. In *Cluster Computing and the Grid, 2005. CCGrid 2005. IEEE International Symposium on*, volume 1, pages 96 – 104 Vol. 1, May 2005.

[10] M. O. Rabin. Efficient dispersal of information for security, load balancing, and fault tolerance. *J. ACM*, 36:335–348, April 1989.

[11] A. Shamir. How to share a secret. *Commun. ACM*, 22(11):612–613, 1979.

[12] G. J. Simmons. How to (really) share a secret. In *CRYPTO '88: Proceedings of the 8th Annual International Cryptology Conference on Advances in Cryptology*, pages 390–448, London, UK, 1990. Springer-Verlag.

[13] B. M. Thuraisingham and J. A. Maurer. Information survivability for evolvable and adaptable real-time command and control systems. *Knowledge and Data Engineering, IEEE Transactions on*, 11(1):228 – 238, 1999.

[14] M. Tu, P. Li, I-Ling Yen, B. M. Thuraisingham, and L. Khan. Secure data objects replication in data grid. *IEEE Trans. on Dependable and Secure Computing*, 7(1):50 –64, 2010.

[15] T. Wu, M. Malkin, and D. Boneh. Building intrusion tolerant applications. In *SSYM'99: Proceedings of the 8th conference on USENIX Security Symposium*, pages 7–7, Berkeley, CA, USA, 1999. USENIX Association.

[16] J. J. Wylie, M. W. Bigrigg, J. D. Strunk, G. R. Ganger, H. Kiliccote, and P. K. Khosla. Survivable information storage systems. *Computer*, 33(8):61 –68, Aug 2000.

[17] W. Yurcik, G. A. Koenig, X. Meng, and J. Greenseid. Cluster security as a unique problem with emergent properties: Issues and techniques. In *5th LCI International Conference on Linux Clusters: The HPC Revolution 2004*, pages 18–20, 2004.

[18] G. Zanin, A. Mei, and L. V. Mancini. Towards a secure dynamic allocation of files in large scale distributed file systems. In *HOT-P2P '04: Proceedings of the 2004 International Workshop on Hot Topics in Peer-to-Peer Systems*, pages 102–107, Washington, DC, USA, 2004. IEEE Computer Society.