Configuring Storage Area Networks Using Mandatory Security

Benjamin Aziz[l*]   Simon N. Foley[†]   John Herbert[†]

Garret Swart[‡]

[l]e-Science Centre, Science and Technology Facilities Council, Didcot OX11 0QX, UK
b.aziz@rl.ac.uk
[†]Department of Computer Science, University College Cork, Cork, Ireland
{s.foley,herbert}@cs.ucc.ie
[‡]Oracle Corporation, 500 Oracle Parkway, Redwood Shores, CA 94065, USA
garret.swart@oracle.com

Abstract

Storage-area networks are a popular and efficient way of building large storage systems both in an enterprise environment and for multi-domain storage service providers. In both environments the network and the storage has to be configured to ensure that the data is maintained securely and can be delivered efficiently. In this paper, we describe a model of mandatory security for SAN services that incorporates the notion of risk as a measure of the robustness of the SAN's configuration and that formally defines a vulnerability common in systems with mandatory security, i.e. cascaded threats. Our abstract SAN model is flexible enough to reflect the data requirements, tractable for the administrator, and can be implemented as part of an automatic configuration system. The implementation is given as part of a prototype written in OPL.

Keywords: Storage Area Networks (SAN), formal validation methods, security, configuration analysis

1   Introduction

Storage-Area Networks (SANs) constitute an important element in modern IT infrastructures due to their efficient management of the underlying storage capabilities in environments shared by several servers, applications, users or even organisations. The purpose of any SAN service is to provide virtual storage services, called a *datasets*, to its clients, typically file systems or database servers.

---

*Corresponding author's address: e-Science Centre, Science and Technology Facilities Council, Didcot OX11 0QX, UK. Tel: +44 (0)1235 778840. Fax: +44 (0)1235 445945.

Each dataset is constructed from physical disks but unlike a physical disk, which has a fixed set of properties that were set when the disk was designed, a dataset has a set of requirements that are specified when the dataset is created and these properties may change over the life of the dataset. These properties may include current capacity, availability in the face of component failures, reliability of the data stored on the dataset in the face of media failures, and any real time performance requirements. These requirements could be in excess of what is provided by a physical device, however, it is up to the administrators to configure the SAN with enough resources so that datasets meeting the requirements can be provisioned.

Like any other multi-user system, a SAN also has security requirements. These can include data privacy – protecting data from unauthorized readers, data integrity – protecting data from unauthorized updates, and additional privacy of the data traffic. One approach to data privacy is to encrypt the data before the application writes it to the dataset. This provides very good privacy but does not help with integrity or traffic analysis attacks. Another, more typical approach yielding all three security properties, is to secure the devices used to provide the SAN service. In an enterprise configuration, this is often done by a combination of security kernels in the SAN devices themselves, as well as firewalls that restrict access of the SAN service to those client machines authorised to use the service. In such an environment, both the firewall and the SAN's own internal security system must be breached before unauthorized access can take place.

The kind of security threat that we are concerned with in this paper is an attack by an *authorized* user of the SAN. It is more likely to take place in a large enterprise SAN environment that is shared by a wide range of organizations, or in a service provider that is providing storage services to many different (and possibly competing) customers. In such an environment, the administrator cannot use the firewall to prevent the attack because it is initiated from a valid user of the SAN that must be allowed to use the service. The administrator must rely on the configuration of the SAN's internal security mechanisms only to determine which datasets can be accessed, and which ones cannot.

Using a single SAN system for storing multiple datasets, accessed by many different authorized clients, means that the implementation of the SAN must be trusted to provide the correct semantics. This reliance on a single layer of software to provide data separation is a potential soft spot in the security of the system. One solution to this security problem, which we adopt in this paper, is to specify carefully the data separation requirements on devices making up the SAN. Once such requirements are specified, this will lead to a search for an acceptable SAN configuration that will maintain a low risk level associated with losing the privacy or integrity of the data being stored on the SAN.

In carrying out an attack, an authorized user attempts to 'copy' data from one dataset to another, in violation of the data-separation/security policy. The attack has consequence for both confidentiality (leaking of dataset) and integrity (unauthorized modification of dataset). The attack may be intentional, where the user deliberately attempts to violate the data-separation policy, or uninten-

tional, where the software that the user executes contains malicious code that attempts to violate the policy unknown to the user. This data-separation security policy is not unlike the requirements for a Chinese Wall security policy [5], and the model that we develop in this paper also provides a novel risk-based interpretation for Chinese Walls.

In this paper we do not consider how underlying security services for SANs might be implemented [14], rather, we take a modeling approach and provide an abstract definition of what is meant by (data-separation) security in a SAN. We are concerned with threats from insiders, and in particular, the primary contribution in this paper is the development of a framework that can be used to determine a secure configuration of a SAN. In addition, the configuration is chosen so as to be within an acceptable level of risk as specified by the security administrator.

The rest of the paper is structured as follows. In Section 2, we define a simple model of SANs. In Section 3, we define the notion of a secure SAN; essentially a SAN equipped with generalized form of mandatory security. In Section 4, we define a configuration analysis for our secure SANs that aims at minimising the risk level and in Section 5, we discuss the issue of cascaded vulnerabilities in secure SANs. In Section 6, we define our notion of an optimal configuration and in Section 7, we review an implementation of the configuration analysis in OPL and finally, in Section 8, we conclude the paper and discuss directions for future work.

2   Storage-Area Networks

According to the model of SANs defined in [20, 1], a SAN is composed from the following elements, as illustrated in Figure 1, where the terminal connectors express many-one and many-many relations among the different elements of the model. These elements are described in the following paragraphs.

Disks   These are the physical storage units that may include tapes, hard disks, optical and solid state devices. We shall write the set of all disks as $DISK = \{disk, disk', \ldots\}$.

Disk Controllers   Usually, disks are arranged into arrays which are connected to one or more controllers. Each controller has a direct access to its attached arrays and, in general, the controller will retain a few spare disks inside the array for use in the event of a failure. We refer to the set of disk controllers as $CNTR = \{ctr, ctr', \ldots\}$.

Logical Volumes   A controller usually builds one or more large virtual disks, which we call logical volumes, out of the physical storage units that it is connected to using RAID techniques [17]. For availability reasons, logical volumes
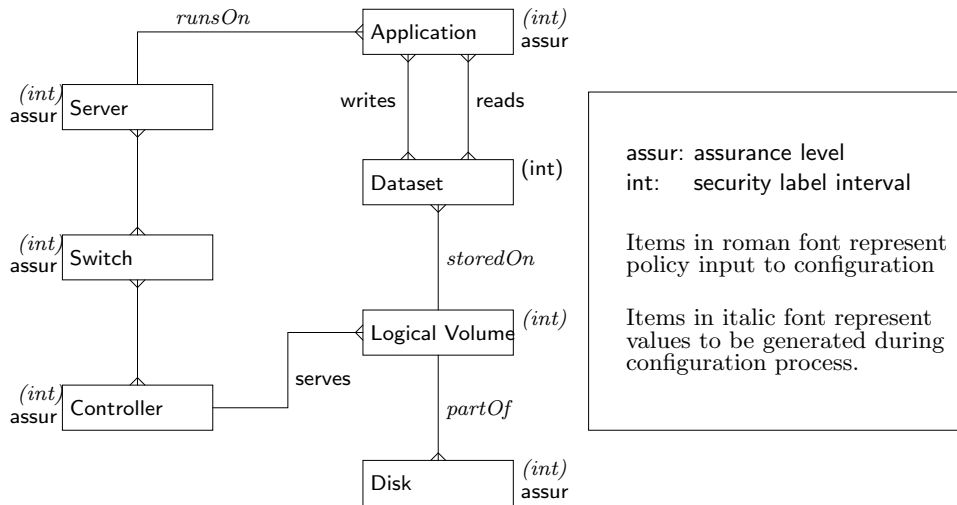
3

Figure 1: The Structure of a SAN.

are not typically tied to a single controller and the failure of the primary controller will trigger its back-up to controllers to start serving the logical volumes it was serving.

We define the set of logical volumes as $LV = \{lv, lv', \ldots\}$. Furthermore, we define the following two functions, which relate logical volumes to disks and controllers, respectively:

$$partOf : DISK \rightarrow LV$$
$$serves : CNTR \rightarrow \wp(LV)$$

Hence, $partOf(disk) = lv$ denotes that $disk$ is part of the logical volume lv, and $serves(ctr) = \{lv_1, \ldots, lv_n\}$ denotes that logical volumes $lv_1, \ldots, lv_n$ are served by controller ctr.

Datasets[1]   In our simplified model, a dataset is a collection of related data files that are accessed by applications as a unit. Controllers implement the datasets by constructing RAID-based logical volumes with the appropriate properties out of a set of identical physical disks and then partitioning them to datasets of the appropriate size. Examples of datasets include the set of web files comprising a website, the set of files holding a database and a user's email files. We write the set of datasets as $DATASET = \{data, data', \ldots\}$.

The following function relates datasets to logical volumes:

$$storedOn : DATASET \rightarrow LV$$

---

[1]In SANS, these are more specifically called Logical Units or LUNs.

such that $storedOn(data) = lv$ signifies the fact that the dataset, $data$, is stored on the logical unit, lv.

Applications, Streams, and Servers   Applications, which constitute a set, $APP = \{app, app', \ldots\}$, are active entities that read and write information stored in one or more datasets. In effect, datasets are regarded as virtual disks that are accessed by applications using *streams*. A stream is regarded as a triple:

$$(app, op, data) \in STREAM$$

where an application, app, accesses some dataset, $data$, using operation, $op \subseteq \{R, W\}$, which is a subset of the Read and Write capabilities.

At any one time, an application runs on a particular *application server*, which may be running more than one application. We write the set of application servers as $SERVER = \{srv, srv', \ldots\}$, and we define the following function, relating applications to their application servers:

$$runsOn : APP \rightarrow SERVER$$

such that $runsOn(app) = srv$ expresses the fact that application, app, is currently running on the application server, srv.

Switches   SAN switches, much like switches in a Local Area Network (LAN), are used to connect the components of a SAN (i.e. its controllers, other switches and application servers) and to connect a SAN to a set of LANs. For a SAN application to be able to access a SAN, the application must either be connected to the LAN which is connected to the SAN, or must be directly connected to the SAN via a network controller specific to the SAN fabric being used.

Assuming that $DEVICE = CNTR \cup SWITCH \cup SERVER$ is the set of devices of a SAN, ranged over by $dev, dev', \ldots$, then we can define the following function:

$$connects : SWITCH \rightarrow \wp(DEVICE)$$

such that $connects(swt) = \{dev_1, \ldots, dev_n\}$ denotes the fact that swt currently connects the devices, $dev_1, \ldots, dev_n$. It is also possible to refer to the transitive closure of $connects$ as $connects^*$.

3   A Mandatory Security Model

In this section we propose a basic mandatory security model that will be used later to characterize security in SANs. The model builds on and extends the label-based model of Multilevel Security with partially trusted subjects [4, 2, 7]. It has been shown that a wide variety of application security policies can be encoded in terms of label-based policies [10].

3.1 Label-based Security Policies

A *label-based security policy* is a lattice of security labels (classifications), SC, with partial ordering, $\leq$, and the least-upper and greatest-lower bound operators, $\sqcup, \sqcap$, respectively. An example is the multilevel security policy with labels, $SC = \{unclassified, secret, topsecret\}$. In this paper, we shall adopt the lattice resulting from the powerset of the set of all organisation names ordered by subset inclusion with the empty set and the full set being the bottom and top elements, respectively. More formally, assume that $ORG = \{org_1, org_2, \ldots\}$ is the set of organisations, $\wp(ORG)$ is the powerset of ORG, then our security lattice is defined in the standard manner as:

$$(\wp(ORG), \subseteq, \cap, \cup, \emptyset, ORG)$$

Hence, $\forall x, y \in SC : x \leq y \Leftrightarrow x \subseteq y$. Intuitively, this means that label y is more important than label x if the set of organisations associated to y is a superset of that associated to x. This association could have different semantics depending on the type of entities x and y are being used to label. In what follows, we are interested mainly in *read/write* semantics.

Let $ENTITIES$ represent the set of all components that can source and/or sink information (disks, controllers, applications, etc.). Every entity, e, is bound to an interval of the policy lattice, where $int(e) = (x, y) \in SC \times SC$, and $x \leq y$, is interpreted to mean that entity e may sink information at class y or lower and may source information at class x or higher. We also write $int(e) = [int_\perp(e), int_\top(e)]$.

If entity e is a subject (in the traditional sense) then $int(e) = [x, y]$ corresponds to a partially trusted subject that may view/read information at class y and lower (vmax) and may write may write/alter information at class x and higher (amin). Conventional objects may be interpreted within this model as entities that are bound to an interval $[x, x]$ with a single level. Informally, $int(e) = [x, y]$ means that an entity, e, can be trusted to properly manage multilevel information within the security interval, $[x, y]$.

Within this model, the definition of a secure system is simply a generalization of the Bell-La Padula axioms (the simple security condition and star property). A system is secure if for all entities, A and B, such that information can flow from A to B then $int_\perp(A) \leq int_\top(B)$ holds. Other possible properties such as the integrity property can also be modeled using, for example, the Biba integrity axioms (where the flow of information is dictated by $int_\top(B) \leq int_\perp(A)$). In this paper, we shall only deal with the Bell-La Padula axioms.

EXAMPLE 1 A SAN is to be configured to manage IBM, HP and $Exxon$ information. The powerset lattice resulting from the set $\{IBM, HP, Exxon\}$ is shown in Figure 2. According to our labeling policy, an application server that is being used to securely process IBM and Exxon data will be assigned the security interval $[\{\}, \{IBM, Exxon\}]$ and a dataset belonging to IBM will have the security interval $[\{IBM\}, \{IBM\}]$. $\triangle$
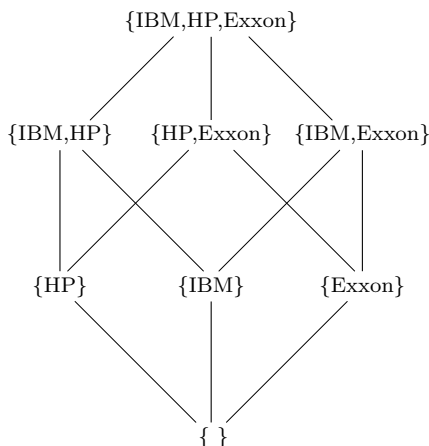
Figure 2: The powerset lattice of $\{\text{IBM}, \text{HP}, \text{Exxon}\}$.

In a MAC model subjects do not necessarily correspond to users. For example, a subject might be a process that is created and/or owned by a user. In our model, a user is *cleared* to an interval and may create/own a subject with an interval that is a subset of the user's clearance. Thus, a user that is cleared to interval [{Exxon}, {IBM,Exxon}] might create processes A and B with classification intervals [{Exxon}, {Exxon}] and [{Exxon,IBM}, {IBM,Exxon}], to process Exxon information and aggregates of Exxon and IBM information (and process A may send data to B, but not vice-versa). Thus, while the user is permitted to read IBM data, the user is not permitted to write IBM-only information. A wide variety of MAC policies can be enforced with these simple policies and the reader is referred to [10] for further examples.

3.2   Assurance Levels

An *assurance level* expresses the level of confidence in the capability of an entity to properly meet its security requirements. An assurance policy is a lattice of assurance levels, A, with partial ordering, $\leq$. The assurance of a device may depend on the complexity of the device's function, the amount of testing that has been applied to the device, the frequency of use, and the development methodology that has been used. Traditionally, assurance ratings are given as a total ordering, for example, $A1 > B3 > B2 > B1 > \ldots$ is the assurance ratings from [21]. In this paper, following [11], we generalize this to a lattice structure as it allows for incomparable assurance levels. Every entity, e, has an associated assurance rating, $assur(e) \in A$. An off-the-shelf physical disk might have a low assurance rating; a multilevel secure application system that has been formally evaluated might have a high assurance rating, while a less formally developed embedded device with very limited functionality might also

have a high assurance rating.

3.3 Risk Functions

The relationship between assurance levels and security intervals can provide an indication of how much an entity should be relied upon. A low-assurance off-the-shelf physical disk configured with security interval $[\{IBM\}, \{IBM\}]$ can be relied upon to manage single-level IBM data. However, the same disk should not be configured with an interval $[\{IBM\}, \{IBM, HP\}]$ as there is not sufficient assurance that it will reliably manage/separate the data between these competing organizations. On the other hand, a disk configured with the interval $[\{\}, \{IBM, Exxon\}]$ is sufficient to manage/separate the non-competing IBM and Exxon data.

We use a *risk function* to encode the relationship between assurance levels and security intervals. Given an entity, e, with assurance level, a, and a security interval, $[x, y]$, then the risk that the entity can be compromised is defined as $risk([x, y], a, e) \in \mathbb{N}$ along with the standard "+" operator to aggregate risk values and the "<" operator to compare them.

It would be interesting in the future to develop algorithms that would compute the risk level of an entity, given its assurance level and security interval. An early approach that adopted the notion of *risk* was introduced in the "Red Book" [21] and the "Yellow Book" [6], which also discussed the issue of what constitutes acceptable risk levels when evaluating the security of computing systems.

Note that our choice of natural numbers as a quantification of risk is arbitrary, however this choice is simple and useful. Other choices (other algebras) could have been equally valid, however, the exploration of these is outside the scope of the paper.

EXAMPLE 2 Consider the security policy from the previous example and an assurance lattice, $lo \leq hi$. There is a low security risk to using an off-the-shelf disk, $disk$, for *single* level data:

$$
\begin{aligned}
risk([\{IBM\}, \{IBM\}], lo, disk) &= risk([\{HP\}, \{HP\}], lo, disk) \\
&= risk([\{Exxon\}, \{Exxon\}], lo, disk) \\
&= risk([\{HP, IBM\}, \{HP, IBM\}], lo, disk) \\
&= 1
\end{aligned}
$$

There is a high security risk when using the same disk to manage/store multilevel data from competing organizations:

$$
\begin{aligned}
risk([\{\}, \{IBM, HP\}], lo, disk) &= risk([\{\}, \{IBM, HP, Exxon\}], lo, disk) \\
&= 40
\end{aligned}
$$

However, there is less of a risk using the off-the-shelf disk to store multilevel

data from non-competing organizations:

$$risk([\{\}, \{IBM, Exxon\}], lo, disk) = risk([\{\}, \{HP, Exxon\}], lo, disk)$$
$$= 10$$

If we assume that a specialized high assurance disk, $disk'$, will properly manage/partition data at different security classifications, then there is less risk when using this disk to manage data from competing organizations:

$$risk([\{\}, \{IBM, HP\}], hi, disk') = risk([\{\}, \{IBM, HP, Exxon\}], hi, disk')$$
$$= 10$$

$$\triangle$$

In providing a relationship between assurance levels and security intervals, the risk function provides a novel approach to characterizing aggregation problems. This contrasts with the lattice based strategies for Chinese Walls that are described in [9, 16, 19] which can be thought of as defining a ((un)acceptable aggregation) binary risk relation.

4   Configuring Secure SANs

A *secure SAN* is a SAN extended with the mandatory security Model defined in the previous section. In the context of our security model, configuring a SAN means searching for a configuration of the SAN devices that meets the specified security policy, the applications' data requirements and any service level agreement (SLA) that may have been agreed with the customers of the data, and that has the least amount of risk possible.

More precisely, before the configuration process can commence we require:

- The security policy, that is, the lattice of security classes, SC, and the risk function, $risk$,

- the application requirements, that is, the security intervals of all datasets, $int(data)$, and the set of streams that relate each application to the datasets that it reads and writes.

- a set of risk limitations in the form of a security class and a maximum risk threshold. These limitations correspond to a customer requirements for an upper bound on the risk in storing one a particular security point class.

- the device specifications, that is, the set of servers, controllers, switches and disks that the SAN is to be configured from. For each such device, e, we need its assurance level, $assur(e)$.

Solving the configuration problem will result in finding values for the $partOf$, $serves$, $storedOn$, $runsOn$ and $connects$ functions that define a particular instance of a SAN system.

4.1 Defining Security Intervals

A dataset, $data$, is initially assigned a point interval, $int(data) = [x, x]$, representing the sensitivity of that dataset. This is a reasonable assumption since datasets are passive entities that can only be manipulated and will never themselves manipulate other datasets. Using the point intervals of a set of datasets, it is possible to compute the interval of an application that will access those datasets through streams, using the lattice *meet*, $\sqcap$, and *join*, $\sqcup$, operations:

$int(app) = [(\sqcap set_\perp), (\sqcup set_\top)]$
where,
$set_\perp = \{int_\perp(data) \mid (app, op, data) \in STREAM \ \wedge \ op = W\}$
$set_\top = \{int_\top(data') \mid (app, op, data') \in STREAM \ \wedge \ op = R\}$

Now, for a particular setting of the *storedOn* function, we can define the security intervals of the logical volumes:

$int(lv) = [(\sqcap set_\perp), (\sqcup set_\top)]$
where,
$set_\perp = \{int_\perp(data) \mid data \in DATASET \ \wedge \ storedOn(data) = lv\}$
$set_\top = \{int_\top(data) \mid data \in DATASET \ \wedge \ storedOn(data) = lv\}$

From the security intervals of logical volumes and given a particular definition of the *serves* function, we can define the security intervals of controllers:

$int(ctr) = [(\sqcap set_\perp), (\sqcup set_\top)]$
where,
$set_\perp = \{int_\perp(lv) \mid lv \in LV \ \wedge \ lv \in serves(ctr)\}$
$set_\top = \{int_\top(lv) \mid lv \in LV \ \wedge \ lv \in serves(ctr)\}$

Similarly, security intervals of disks may be defined based on the security intervals of the logical volumes and a definition of the *partOf* function:

$int(disk) = int(partOf(disk))$

On the other hand, the security interval of an application server is defined based on the security intervals of its applications running and a definition of the *runsOn* function:

$int(srv) = [(\sqcap set_\perp), (\sqcup set_\top)]$
where,
$set_\perp = \{int_\perp(app) \mid app \in APP \wedge runsOn(app) = srv\}$
$set_\top = \{int_\top(app) \mid app \in APP \wedge runsOn(app) = srv\}$

Finally, intervals of switches are computed from intervals of the devices they connect (i.e. other switches, controllers, disks and application servers), given a

definition of the *connects* function:

$int(swt) = [(\sqcap set_\perp), (\sqcup set_\top)]$
where,
$set_\perp = \mu \; swt'.$
$\qquad (\{int_\perp(srv) \mid srv \in SERVER \wedge srv \in connects(swt)\} \cup$
$\qquad \{int_\perp(ctr) \mid srv \in CNTR \wedge ctr \in connects(swt)\} \cup$
$\qquad \{int_\perp(swt') \mid swt' \in SWITCH \wedge swt' \in connects(swt)\})$
$set_\top = \mu \; swt'.$
$\qquad (\{int_\top(srv) \mid srv \in SERVER \wedge srv \in connects(swt)\} \cup$
$\qquad \{int_\top(ctr) \mid srv \in CNTR \wedge ctr \in connects(swt)\} \cup$
$\qquad \{int_\top(swt') \mid swt' \in SWITCH \wedge swt' \in connects(swt)\})$

The usage of the least-fixed point operator, μ, is required since the definition of *connects* is recursive.

EXAMPLE 3 Given the following dataset intervals:

$int(data_1) = [\{IBM, Exxon, foo\}, \{IBM, Exxon, foo\}]$
$int(data_2) = [\{IBM, foo\}, \{IBM, foo\}]$
$int(data_3) = [\{foo\}, \{foo\}],$

and an application, *app*, using the following streams

$(app, \{R\}, data_1)$
$(app, \{R, W\}, data_2)$
$(app, \{W\}, data_3),$

then then if *app* is classified with the interval, $[\{foo\}, \{IBM, Exxon, foo\}]$, it can handle the above data using the streams indicated. $\triangle$

5 Cascade Vulnerabilities

Cascade vulnerabilities were first identified in [21] in networks of systems with multilevel security classifications and assurance levels. In such systems, the problem arises in scenarios similar to that shown in Figure 3.

Given the ordering on security classifications, $C \leq S \leq TS$, and ordering, $B1 \leq B2 \leq B3$, on levels of assurance, then this system exhibits a case of the cascade vulnerability problem arising from the fact that an attacker classified at security level, C, and possessing the capabilities to break assurance level, $B2$ (and below), can attack and break $Sys.A$ and $Sys.C$. As a result, the attacker can go through the following scenario: First, it downgrades information labeled TS on $Sys.A$ to level S. Second, it communicates this information to $Sys.C$. Finally, it downgrades the information to level C, which is the level of the attacker itself, i.e. it can read the information perfectly legally now. This

Sys.A Sys.B

TS		TS
S		S
		C

| S |
| C |

Sys.C

| Assurance Levels	
System	Level
Sys.A	B2
Sys.B	B3
Sys.C	B1

Figure 3: The Cascade Vulnerability Problem.

scenario reflects a vulnerability, which undermines the assurance requirement stating that for information to be down-classified from TS to C, an attacker needs to be able to have capability at level $B3$ or higher.

Cascade vulnerability attacks are possible in any network model with security and assurance classifications. In the following sections, we present an analysis of the vulnerability in our model of secure SANs.

5.1 Introductory Definitions

First, we need to clarify what we mean by a network.

Definition 5.1 (Network) *Assuming a set of nodes, $\mathcal{N} = \{n_1, n_2, \ldots\}$, a set of edges, $\mathcal{E} = \{(n, n', [x, x]), \ldots\}$, where only data classified at security level, $[x, x]^2$ (or lower), is allowed to flow from n to n', then a network is defined as the following quadruple:*

$$(\mathcal{N}, \mathcal{E}, int, assur)$$

Where $int : \mathcal{N} \rightarrow (SC \times SC)$ and $assur : \mathcal{N} \rightarrow A$ are the usual mappings from nodes to their security intervals and to their assurance levels, respectively.

A network then contains nodes, edges and security and assurance classification mappings. In a certain sense, a network is a directed graph with the extra information of assurance and security level mappings. Based on this definition of a network, we define a chain as follows.

Definition 5.2 (Chain) *Given a network, $(\mathcal{N}, \mathcal{E}, int, assur)$, then a chain is defined as any list of nodes, $[n_1, \ldots, n_m]$, which satisfies the property:*

$$\forall n_i, n_{i+1} \in [n_1, \ldots, n_m], \exists [x, x] \in SC \times SC : (n_i, n_{i+1}, [x, x]) \in \mathcal{E}$$

Hence, in Figure 3, we have that $[Sys.B, Sys.A, Sys.D, Sys.C]$ is a chain, and so is $[Sys.D, Sys.C]$. Furthermore, a bad chain is a chain that can be compromised by some attacker.

---

[2] We follow the convention adopted in Section 4.1 that data have point security intervals.

Definition 5.3 (Bad Chain) *Given an attacker, I, with a capability to break a maximum assurance level, $a_I \in A$, then we say that a chain, $[n_1, \ldots, n_m]$, is a bad chain with respect to I, if the following is true:*

$$\forall n \in [n_1, \ldots, n_m] : assur(n) \leq a_I$$

This implies that every element in the chain can be broken by I. Using these definitions, we can now define a cascade vulnerability formally as follows.

Definition 5.4 (Cascade Vulnerability) *We say that an attacker, I, with security level, $int(I) \in SC \times SC$, and capability, $a_I \in A$, poses a cascaded threat using a bad chain, $[n_1, \ldots, n_m]$, to data classified at level $[x, x] \in SC \times SC$ on some node, $n \notin [n_1, \ldots, n_m]$, if the following condition holds true:*

$$\exists n_i, n_j \in [n_1, \ldots, n_m] : (n, n_i, [x, x]) \in \mathcal{E} \ \wedge \ int_\top(I) \geq int_\perp(n_j)$$

The condition in Definition 5.4 highlights the role of two important nodes in a bad chain: the first, n_i, which is linked at security level, $[x, x]$, with the victim node, n, acts as an indirect leakage point since data classified at $[x, x]$ can legally pass from n to n_i. The second node, n_j, acts as a direct point of leakage of data to the attacker since the latter can legally read any data n_j is allowed to write. This is true because the attacker has a read-level of $int_\top(I)$, which is more than or equal to the write-level, $int_\perp(n_j)$. In other words, a cascade vulnerability arises whenever an attacker becomes capable of manipulating a bad chain to obtain sensitive data from a node, which has some sort of connection to the bad chain, and then downgrade and read the sensitive data.

In the following sections, we give concrete examples of the cascade vulnerability problem in the cases of networks of logical volumes and networks of application servers.

5.2 Networks of Logical Volumes

Using the model of SANs as defined in Section 2, we define a network of logical volumes as follows.

Definition 5.5 (Network of Logical Volumes) *Given some definition of the function, $storedOn : DATASET \rightarrow LV$, and a set of streams, $STREAM$, then we can define a network of logical volumes as follows:*

$(\mathcal{N} = \{lv_1, \ldots, lv_m\},$
$\mathcal{E} = \{(lv_i, lv_j, [x, x]) \mid lv_i, lv_j \in \{lv_1, \ldots, lv_m\} \ \wedge$
$\qquad\qquad\qquad \exists (app, R, data_i), (app, W, data_j) \in STREAM \ \wedge$
$\qquad\qquad\qquad storedOn(data_i) = lv_i \ \wedge \ storedOn(data_j) = lv_j \ \wedge$
$\qquad\qquad\qquad [x, x] = int(data_i)\},$
$int : \{lv_1, \ldots, lv_m\} \rightarrow (SC \times SC),$
$assur : \{lv_1, \ldots, lv_m\} \rightarrow A))$

Hence, the flow of information from lv_i to lv_j is modeled as the presence of an application, app, which reads from a dataset, $data_i$, stored on lv_i and then writes to another dataset, $data_j$, stored on lv_j. Note that since it is $data_i$ which is *flowing* from lv_i to lv_j, then $[x, x]$ is taken as the interval of $data_i$.

By adopting this definition of a network of logical volumes, we can reason about the presence of a cascade threat towards a logical volume, $lv \in \mathcal{N}$, at data interval, $[x, x]$, from an attacker, I, capable of breaking a bad chain of logical volumes, $[lv_1, \ldots, lv_m]$, as follows:

$$cascadeLv(lv, [x,x], I) \overset{\text{def}}{=}$$
$$\exists lv_i, lv_j \in [lv_1, \ldots, lv_m] : (lv, lv_i, [x, x]) \in \mathcal{E} \ \wedge \ int_\top(I) \geq int_\bot(lv_j)$$

This definition of cascade vulnerability in a network of logical volumes is a straightforward adaptation of Definition (5.4).

5.3 Networks of Application Servers

The other network that we identify within our model of secure SANs is the network of application servers.

Definition 5.6 (Network of Application Servers) *Given some definition of the function, $runsOn : APP \to SERVER$, and a set of streams, $STREAM$, then we define a network of application servers by the following quadruple:*

$$(\mathcal{N} = \{srv_1, \ldots, srv_m\},$$
$$\mathcal{E} = \{(srv_i, srv_j, [x, x]) \mid srv_i, srv_j \in \{srv_1, \ldots, srv_m\} \ \wedge$$
$$\exists (app_i, W, data), (app_j, R, data) \in STREAM \ \wedge$$
$$runsOn(app_i) = srv_i \ \wedge \ runsOn(app_j) = srv_j \ \wedge$$
$$[x, x] = int(data)\},$$
$$int : \{srv_1, \ldots, srv_m\} \to (SC \times SC),$$
$$assur : \{srv_1, \ldots, srv_m\} \to A))$$

In this definition, information flows from srv_i to srv_j if there exists a dataset, $data$, which acts as a shared space to which an application, app_i, running on srv_i writes and from which another application, app_j, running on srv_j reads from. The security interval, $[x, x]$, is taken directly as the interval of ata. Note that this manipulation of streams is quite different from that of the previous section. However, the definition of the cascade vulnerability is quite similar: A cascaded threat towards an application server, srv, is possible at security level, $[x, x]$, given some intruder, I, and a bad chain of application servers, $[srv_1, \ldots, srv_m]$, if the following predicate holds true:

$$cascadeSrv(srv, [x,x], I) \overset{\text{def}}{=}$$
$$\exists srv_i, srv_j \in [srv_1, \ldots, srv_m] : (srv, srv_i, [x, x]) \in \mathcal{E} \ \wedge \ int_\top(I) \geq int_\bot(srv_j)$$

6 Optimal Configurations

After defining the security intervals, $int(e)$, of every SAN entity, e, as in Section 4.1, and given that entities have fixed assurance levels, $assur(e)$, then the problem of finding optimal configurations consists in searching for definitions of one or more of the functions, $partOf$, $serves$, $runsOn$, $storedOn$ and $connects$, such that the total risk level is at a minimum value:

$$\sum_{\forall e \in ENTITIES} risk(int(e), assur(e), e) \tag{1}$$

Individual customers may insist on a Service Level Agreement (SLA) that limits the risk that their data is compromised. For example, a customer may provide a security interval int_{SLA} and a limit ν_{SLA} and require that:

$$\sum_{\forall e \in ENTITIES:\ int(e)\ \cap\ int_{SLA}\ \neq\ \{\}} risk(int(e), assur(e), e) \leq \nu_{SLA} \tag{2}$$

The SLA assures the customer that the risk in the configuration for storing their data is low enough.

Furthermore, the customer may require that no instances exist of the two cascade vulnerabilities as defined in Section 5 given some attacker model I:

$$\forall lv \in LV, x \in SC \quad : \quad \neg cascadeLv(lv,[x,x],I) \tag{3}$$

$$\forall srv \in SERVER, x \in SC \quad : \quad \neg cascadeSrv(srv,[x,x],I) \tag{4}$$

EXAMPLE 4 In a particular incomplete configuration of a secure SAN, assume that the aggregate risk levels of applications, logical volumes, controllers, servers and switches are as follows:

$$\sum_{\forall app \in APP} risk(int(app), assur(app), app) = 70$$

$$\sum_{\forall lv \in LV} risk(int(lv), assur(lv), lv) = 50$$

$$\sum_{\forall ctr \in CNTR} risk(int(ctr), assur(ctr), ctr) = 45$$

$$\sum_{\forall srv \in SERVER} risk(int(srv), assur(srv), srv) = 45$$

$$\sum_{\forall swt \in SWITCH} risk(int(swt), assur(swt), swt) = 30$$

Then, we would like to find out the optimal definition of the $partOf$ function (i.e. i.e. the assignment distribution of physical disks to logical volumes), for each of the following:

a. $\sum_{\forall e \in ENTITIES} risk(int(e), assur(e), e) = 300$

b. $\nu_{SLA} = 280$

To solve case (a), we need to find out the maximum risk level disks may have:

$$\sum_{\forall disk \in DISK} risk(int(disk), assur(disk), disk) = 300 - (70 + 50 + 45 + 45 + 30) = 60$$

From the definition of disk security intervals, we obtain:

$$\sum_{\forall disk \in DISK} risk(int(partOf(disk)), assur(disk), disk) = 60$$

Assuming we already have definitions of *serves*, *storedOn*, *runsOn* and *connects*, then we can solve the above equation to find out the most suitable definition(s) of *partOf* function and thereby complete the configuration of the SAN.

In the case of (b), the procedure is similar, except that now we have:

$$\sum_{\forall disk \in DISK} risk(int(disk), assur(disk), disk) = 280 - (70 + 50 + 45 + 45 + 30) = 40$$

$\triangle$

7 OPL Implementation

To test our understanding of this security and configuration model and to test its usefulness, we implemented the model and used it to generate the lowest risk configuration that meets the requirements. We decided to use OPL [13] for the implementation language because of its built-in logic and search capabilities.

The OPL program consists of five pieces:

- Input Data Model: Describes all the data that must be supplied to define a particular instance of the problem to be solved. The input data is generally validated to make sure that the request is not obviously inconsistent.

- Variable Data Model: Describes the data that the program is to determine values for.

- Constraints: A set of relations that must hold between the variables and the input data. The number of these constraints can depend on the input provided. If all the constraints hold for a particular assignment to the variable data, that assignment is called a feasible solution.

- Objective function: A function that is maximized or minimized from among the feasible solutions. OPL reports new maxima or minima as they are determined during the search process, finally terminating when the search space of variable data has been exhausted.

- Search procedure: An optional plan for how to find the optimal solution. Typically this involves carefully choosing the order in which the variable data is examined and noticing when further changes will be ineffective.

The OPL input data consists of an instantiation of the Input Data Model. The output of the OPL program is a sequence of successively improving feasible solutions. For this application, the Input Data Model is used to represent all needed input: the security policy, the application requirements, any SLA requirements, and the device specifications. We do validation of the input data to ensure that the security class forms a lattice and that the risk function is consistent with the lattice, and also to ensure that any static requirements, e.g. requirements on the applications themselves, are met.

The Variable Data Model is used to represent the interval for each device and the SAN configuration functions, that is, *storedOn*, *runsOn*, *serves*, *partOf* and *connects*. In the worst case, finding the optimal configuration means examining every combination of values in the Variable Data Model, so it is very important to make sure that there is a minimum of redundancy or over specification in the model. The constraints fall into several categories:

- Device interval constraints implement the formulae defined in the previous section.

- Configuration consistency constraints make sure that the configuration meets the basic requirements, for example, that each logical volume is assigned enough disks to store the assigned datasets, that servers and controllers are all connected to switches, and switches to each other.

- Canonicalization constraints prune all but one equivalent configuration from the configuration space. This is important for reducing the search space as discussed in [12].

- SLA constraints ensure that the risk for a particular security class is limited to the agreed value.

The final piece of the OPL program is the search method. In this case it simply makes judicious choices about which part of the variable data space to explore first. The primary issue in the search is to make sure that the intervals are evaluated once the needed bits of the configuration have been generated. Quick elimination of infeasible or less optimal alternatives is the key to a fast running OPL application.

As a tool for exploring the space of SAN risk models, OPL worked well for us. As part of this research we went through several iterations of the risk model, expressing each iteration in OPL. This process was straightforward and helped uncover holes in our thinking. By using a logic language the semantic distance between the model and the implementation was small enough that the implementation actually acted as a useful means of expression, not just a hurdle to be jumped. The discipline of implementing and running the model also provided a check, keeping the model simpler than it might have been.

However the OPL implementation that we used was designed as a programming environment not as a deployment tool. Embedding it into a larger n-tier system would be difficult. In addition there are inputs where the search procedure we used could require exponential time, this is hard to avoid when looking

for an exact solution, because the optimization problem itself is NP-complete by reduction to bin packing. To address both issues, a custom solver is the best answer, both for efficiency, encapsulation and packaging. Once the model is fixed, a custom solver allows for more efficient evaluation of the constraints and the objective functions, and for the use of approximation techniques, such as simulated annealing or genetic algorithms, that have been tuned to the model, Maintaining such a solver increases the cost of changing the model, but when looked at as part of the total cost of maintaining a complex SAN management product, it is not the highest order term.

8 Conclusion and Future Work

In this paper, we presented a model of mandatory security for SANs. The primary contribution in this paper was the development of a framework that can be used to determine a secure data-separation configuration of a SAN. The model incorporated an evaluation of the total risk incurred when data stored on a SAN are compromised. We also outlined an implementation of the configuration process in OPL.

A label-based model is used to represent security in SANs. While conceptually simple, lattice/label-based models can be used to characterize mechanisms that support a wide range of mandatory security requirements [3, 21, 10, 9, 7, 15, 16, 19]. Therefore, we conjecture that the results in this paper can be usefully applied to other mandatory protection models such as Role Based Access Control; this is a topic of ongoing research.

The SAN security model extends the dual-label/partially trusted subject lattice model with the addition of a risk function. This function is used to encode the level of risk associated with storing/managing combinations of information on entities evaluated to certain degrees of assurance. This is more flexible than the conventional assurance/evaluation criteria approach [21]; the risk function is used to guide the generation of a secure configuration within an acceptable degree/measure of risk. The data-separation security policy is not unlike the requirements for a Chinese Wall security policy [5], and the model proposed in this paper provides a novel risk-based interpretation for Chinese Walls that can be used in a non-SAN situation. Given the duality between Chinese wall (aggregation) polices and separation of duty policies, we argue that the proposed model can also be used to capture dynamic separation of duty policies that are risk-based interpretations of those described in [8]. This is a topic for future work.

One limitation of our configuration technique is the possibility that it might be hard to apply it to newer shared storage models like those provided by Amazon's S3 (aws.amazon.com/s3). Primarily because a dataset on S3 would typically be spread out over the entire data center with portions stored on every single logical volume. This reduces their cost of administration and increases their parallelism but at the cost of not being able to bound the security foot print of the dataset.

Future work would explore the definition of risk functions and how they behave towards changes in assurance levels and security intervals. Also, we plan to investigate more deeply other security properties, such as integrity and availability, as well as other models of security, such as the continuous Usage Control (UCON$_{\mathrm{ABC}}$) model [18]. Other future work includes extending the OPL implementation to be able to capture instances of the cascade vulnerability problem.

References

[1] Benjamin Aziz, Simon N. Foley, John Herbert, and Garret Swart. Configuring storage area networks for mandatory security. In Csilla Farkas and Pierangela Samarati, editors, *Proceedings of the* 18$^{\mathrm{th}}$ *IFIP Annual Conference on Data and Applications Security*, pages 357–370, Sitges, Catalonia, Spain, July 2004. Kluwer.

[2] D.E. Bell. Concerning modelling of computer security. In *Proceedings of the* 1988 *IEEE Symposium on Security and Privacy*, pages 8–13, Oakland, California, USA, April 1988. IEEE Computer Society Press.

[3] D.E. Bell and L.J. La Padula. Secure computer systems: Unified exposition and multics interpretation. Technical Report ESD–TR–75–306, Mitre Corporation, Bedford, M.A., U.S.A., July 1975.

[4] M. Branstad et al. Trusted mach design issues. In *Proceedings of the* 3rd *AIAA/ASIS/DODCI Aerospace Computer Security Conference*, Orlando, Florida, USA, December 1987. IEEE Press.

[5] D.F.C. Brewer and M.J. Nash. The chinese wall security policy. In *Proceedings of the* 1989 *IEEE Symposium on Security and Privacy*, pages 206–214, Oakland, California, USA, May 1989. IEEE Computer Society Press.

[6] DoD. Computer security requirements guidance for applying the department of defence trusted computer system evaluation criteria in specific environments. Technical Report CSC-STD-003-85, DoD Computer Security Center, 1985. Yellow Book.

[7] Simon N. Foley. A model for secure information flow. In *Proceedings of the* 1989 *IEEE Symposium on Security and Privacy*, pages 248–258, Oakland, California, USA, May 1989. IEEE Computer Society Press.

[8] Simon N. Foley. Secure information flow using security groups. In *Proceedings of the* 3rd *IEEE Computer Security Foundations Workshop*, pages 62–72, Franconia, New Hampshire, USA, June 1990. IEEE Computer Society Press.

[9] Simon N. Foley. Aggregation and separation as noninterference properties. *Journal of Computer Security*, 1(2):159–188, 1992.

[10] Simon N. Foley. The specification and implementation of commercial security requirements including dynamic segregation of duties. In *Proceedings of the* 4th *ACM Conference on Computer and Communications Security*, pages 125–134, Zurich, Switzerland, April 1997. ACM Press.

[11] Simon N. Foley. Conduit cascades and secure synchronization. In *Proceedings of the ACM New Security Paradigms Workshop*, pages 141–150, Cork, Ireland, September 2000. ACM Press.

[12] Eugene Freuder. Eliminating interchangeable values in constraint satisfaction problems. In *Proceedings of the* 9th *National Conference on Artificial Intelligence*, volume 1, pages 227–233, Anaheim, California, USA, July 1991. MIT Press.

[13] Pascal Van Hentenryck. *The OPL Optimization Programming Language.* MIT Press, Cambridge, Massachusetts, USA, January 1999.

[14] Yongdae Kim, Maithili Narasimha, Fabio Maino, and Gene Tsudik. Secure group services for storage area networks. In *Proceedings of the First International IEEE Security in Storage Workshop*, pages 80–93, Greenbelt, Maryland, USA, December 2002. IEEE Computer Society.

[15] T.M.P. Lee. Using mandatory integrity to enforce 'commerical' security. In *Proceedings of the Symposium on Security and Privacy*, pages 140–146, Oakland, C.A., U.S.A., 1988. IEEE Press.

[16] Catherine Meadows. Extending the brewer-nash model to a multilevel context. In *Proceedings of the IEEE Symposium on Security and Privacy*, pages 95–102, Oakland, California, USA, May 1990. IEEE Computer Society Press.

[17] D.A. Patterson, G.A. Gibson, and R.H. Katz. A case for redundant arrays of inexpensive disks (raid). In *Proceedings of the* 1988 *ACM Conference on the Management of Data (SIGMOD)*, pages 109–116, Chicago, USA, June 1988. ACM Press.

[18] R. Sandhu and J. Park. The $UCON_{ABC}$ usage control model. *ACM Transactions on Information and System Security*, 7(1):128–174, 2004.

[19] Ravi S. Sandhu. Lattice-based access control models. *Computer*, 26(11):9–19, November 1993.

[20] Garret Swart. Storage management by constraint satisfaction. In *Proceedings of the Workshop on Immediate Applications of Constraint Programming*, Kinsale, Cork, Ireland, September 2003.

[21] TNI. Trusted computer system evaluation criteria: Trusted network interpretation. Technical report, National Computer Security Center, 1987. Red Book.