SANDIA REPORT

SAND96–0536 • UC–900
Unlimited Release
Printed March 1996

Comprehensive Test Ban Treaty International Monitoring System Security Threats and Proposed Security Attributes

Timothy J. Draelos, Richard L. Craft

Prepared by
Sandia National Laboratories
Albuquerque, New Mexico 87185 *and Livermore, California 94550*
for the United States Department of Energy
under Contract DE-AC04-94AL85000

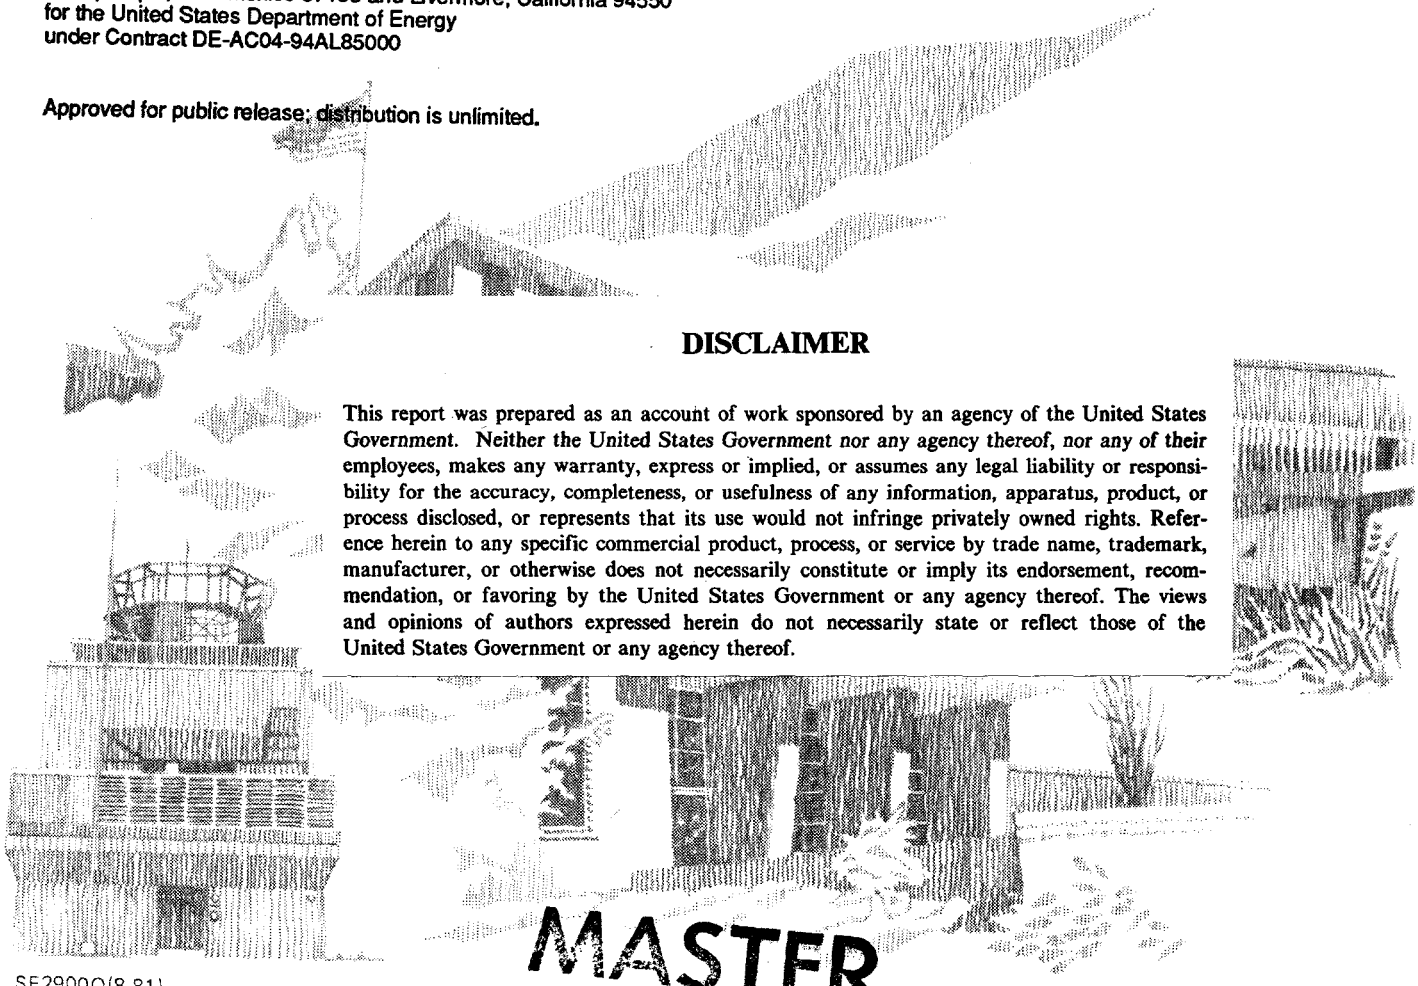

Issued by Sandia National Laboratories, operated for the United States Department of Energy by Sandia Corporation.

Available to DOE and DOE contractors from
Office of Scientific and Technical Information
PO Box 62
Oak Ridge, TN 37831

Prices available from (615) 576-8401, FTS 626-8401

Available to the public from
National Technical Information Service
US Department of Commerce
5285 Port Royal Rd
Springfield, VA 22161

NTIS price codes
Printed copy: A03
Microfiche copy: A01

SAND96-0536
Unlimited Release
Printed March 1996

Distribution
Category UC–900

Comprehensive Test Ban Treaty International Monitoring System Security Threats and Proposed Security Attributes

Timothy J. Draelos
Cooperative Monitoring Technologies Department

Richard L. Craft
Data Systems Security

Sandia National Laboratories
Albuquerque, NM 87185

Abstract

To monitor compliance with a Comprehensive Test Ban Treaty (CTBT), a sensing network, referred to as the International Monitoring System (IMS), is being deployed. Success of the IMS depends on both its ability to perform its function and the international community's confidence in the system. To ensure these goals, steps must be taken to secure the system against attacks that would undermine it; however, it is not clear that consensus exists with respect to the security requirements that should be levied on the IMS design. In addition, the CTBT has not clearly articulated what threats it wishes to address. This paper proposes four system-level threats that should drive IMS design considerations, identifies potential threat agents, and collects into one place the security requirements that have been suggested by various elements of the IMS community. For each such requirement, issues associated with the requirement are identified and rationale for the requirement is discussed.

Intentionally Left Blank

Contents

1. Introduction 1
- 1.1 Purpose 1
- 1.2 Scope 2
- 1.3 Document overview 2

2. Threats and Threat Agents 3
- 2.1 Four potential active threats to the IMS 3
- 2.2 Potential threat agents 4
  - IMS designers, implementors, and installers 5
  - Authorized IMS users 5
  - IDC insiders 5
  - Hackers and outside agents 6
  - Environment and system design weaknesses 6

3. Security Attributes 7
- 3.1 Integrity 7
- 3.2 Authenticity and nonrepudiation 8
- 3.3 Availability 11
- 3.4 Access control 12
  - Identification and authentication 12
  - Authorization 13
  - Access-control mechanisms 13
- 3.5 Accountability and intrusion detection 15
- 3.6 Visibility 17

4. Summary 19

5. Future Work 21

References 23

Intentionally Left Blank

Comprehensive Test Ban Treaty International Monitoring System Security Threats and Proposed Security Attributes

1. Introduction

The international community is preparing to sign a Comprehensive Test Ban Treaty (CTBT). To effectively monitor treaty compliance, the treaty will call for the implementation of an international monitoring system (IMS) capable of detecting and characterizing low-yield nuclear weapon tests. Much work has already been done to investigate what will be required to implement such a system. To date, less emphasis has been placed on considering how to secure such a system so that its results can be trusted by the community that depends on its products.

While a number of statements regarding desirable CTBT information security attributes have appeared in various papers produced by the CTBT community, it is not clear whether there is a consensus within the community as to what security requirements should be levied on the IMS design. The issues emphasized differ from one portion of the community to the next. A more fundamental concern is that no explicit statement exists regarding what threats the CTBT community will address. To make best use of the resources dedicated to the design, implementation, and operation of an IMS, the CTBT community needs to understand these issues and to have a clear plan for addressing them.

1.1 Purpose

This paper has three purposes:

- to detail the basic threats to the IMS that should concern the CTBT community,
- to collect into one place the information system security attributes that have been suggested by the various elements of the CTBT community, and
- to discuss the rationale behind and issues associated with each suggested attribute.

This paper is written to generate discussion within the CTBT monitoring community on those information security threats that should be given consideration, and the corresponding information security requirements that should be levied on the CTBT IMS design. The threat scenarios and threat agents discussed in this document have been proposed by Sandia National Laboratories. Most of the attributes in this document have been previously enumerated by the Advanced Research Projects Agency (ARPA) and its contractor, Trusted Information Systems,

as part of ARPA's GSETT-3[1] work; however, no documentation exists that addresses the implications of each suggested attribute.

Where a suggested security attribute has been drawn from a document, the attribute's source is referenced in brackets following the stated requirement (e.g., [Stated: Annex 5, Section 1.2.1 and Section 1.5]). The remaining attributes in this document (i.e., those marked with an empty set of reference brackets: [ ]) have been derived from discussions that Sandia has had with elements of the CTBT monitoring community.

1.2 Scope

The CTBT IMS will provide a means by which the international community can effectively monitor compliance with the Comprehensive Test Ban Treaty. The system is expected to include a global network of seismic, hydroacoustic, infrasound, and radionuclide sensors along with significant event detection, location, and parameterizing capabilities resident at a central international data center (IDC) and at various national data centers (NDC). The goal of this system is to detect, characterize, and report to the international monitoring community all events relevant to monitoring treaty compliance. The goal of information security in the CTBT is to provide the greatest assurance possible (within the constraints of funding) that human actions (intentional or otherwise), environmental factors, and inherent system characteristics (e.g., single-point dependencies) do not prevent the system from performing as it is intended.

1.3 Document overview

- Section 2 of this document proposes several potential threat scenarios and threat agents for the CTBT IMS.
- Section 3 enumerates and discusses each of the information security attributes that have been stated or implied.
- Section 4 summarizes the document.
- Section 5 discusses what steps will have to be taken to ensure that the fielded IMS is secure.

2. Threats and Threat Agents

In this section, four potential active threats to the IMS are identified and discussed. In addition to threats, this section suggests potential attackers of the IMS. It should be noted that since the focus of the IMS is the detection of nuclear weapon tests, the greatest threats to the monitoring system are likely to come from those parties not wanting the system to correctly detect events (e.g., countries wishing to test in violation of the treaty). For this reason, even though outside hackers are considered potential agents in each threat, this emphasis is secondary compared with the users of the system and those countries monitored by it.

2.1 Four potential active threats to the IMS

Threat 1: Hiding an event.

Discussion: The focus of this threat is an event that is, in fact, detectable by the IMS but that is somehow kept from appearing on an event bulletin produced by either the IMS's International Data Center (IDC) or by any of the National Data Centers (NDC) serviced by the IMS.

The most likely agent for this threat is a country desiring to conduct a test. For a sufficiently small yield, it is conceivable that the stations most likely to detect the test could be within the control of the testing country. Given this, reasonable attack scenarios include:

- Disabling the IMS elements (sensors, stations, NDCs, etc.) in the country for the duration of the event, and
- Leaving the system operational but overwriting event data with previously collected ambient data for that station.

It is also conceivable that an attack intended to hide the data could be launched from the IDC itself by means of an insider erasing source data, altering the processing subsystem in such a way as to avoid detection, or by altering the event bulletin after processing but before distribution. Unless special measures are put in place to prevent this, a credible case can be built for the viability of one-man attacks of this sort.

Finally, given the likelihood of using commercially available computers, open networking standards, and commercial networks in the implementation of the IMS, hackers from outside the CTBT community represent real threats to the integrity of the IMS's data. Source data queued up for processing might be destroyed, processing algorithms might be altered, event bulletins might be corrupted, etc.

Threat 2: Altering an event's characteristics.

Discussion: Rather than trying to totally hide the fact that an event occurred, an attacker may simply want to change the perceived characteristics of the event. Among other things, these changes could include making the event appear to be a different type of event (e.g., a mine blast rather than a low-yield test), modifying the event's apparent magnitude, or making the event appear to have

occurred in a location different from where it actually occurred.

As before, the most likely agent for this threat is a country desiring to conduct a test. Using the same techniques as described for "hiding an event" (the first threat), the country might launch an attack to modify the event's apparent characteristics. The motivation for this (as opposed to hiding an event) may be that the country's span-of-control within the IMS is such that it cannot control the data produced by all sensors capable of seeing the event.

As before, this kind of threat might be readily realized by a sufficiently skilled IDC insider working for the testing country. It might also be accomplished by a hacker wanting to exert some power by corrupting IMS data.

Threat 3: Creating an event.

Discussion: A number of different means might be used to realize this threat: First, old sensor data from previous events might be injected into the system, or new data might be completely synthesized and played into the system. Second, fictitious records might be placed in the incoming data queue at the central point of the system — the IDC. Third, an insider might also directly modify event bulletins after they are created; however, the source data alleged to have contributed would most likely have to be modified as well, in order to create a coherent picture. Finally, a hacker could achieve the same goals as an insider given the assumptions stated above.

Threat 4: Destroying trust in IMS integrity.

Discussion: While the first threat may be the most obvious, this final threat might be the most effective. If an adversary is able to create doubt as to the accuracy of the IMS, then it may be easier to initiate a test without consequence. "The system has reported a dozen false alarms already this year. So what justification do you have for wanting to come inspect my site this time?" an adversary might ask.

In practice, this threat might be realized in a number of ways: If the test to be conducted is likely to be detected by only a limited set of sensors and these are largely under that country's control, then the country might corrupt the sensor data on a regular basis in order to create a history of problems with those stations. Similarly, an IDC insider might accomplish the same thing by creating spurious events that cannot be corroborated by inspection or other means and by suppressing events that clearly should have been detected by the IMS. As before, an independent hacker might accomplish the same thing as an insider.

2.2 Potential threat agents

As might be deduced from these four threats, potential threat agents include:

- IMS designers, implementors, and installers.
- Authorized IMS users.
- IDC insiders.
- Hackers and outside agents.
- Environment and system design weaknesses.

The following paragraphs describe each of these agents in more detail.

IMS designers, implementors, and installers

IMS designers, implementors, and installers pose a special threat to the security of the IMS. No group has a better chance than this one to install malicious functionality into the system. For this reason, special care needs to be exercised during the design, implementation, and fielding stages to ensure that trustworthy components are being placed into operation.

Authorized IMS users

An authorized IMS user is any person authorized to access IMS information system resources or authorized to have physical access to IMS components. IMS users could include:

- Staff charged by a given country with monitoring events detected by the IMS.
- Operations staff responsible for day-to-day functioning of IMS components within a given facility.
- Personnel charged with maintenance of IMS components (sensors, computers, communications equipment, etc.).
- Other users who would find access to the IMS data useful (e.g., seismic researchers).

The primary concerns with this group center around its potential ability to use authorized access to subvert system operation. This could include, for example, removing power to IMS equipment, spoofing sensors by various means, or installing altered software that performs some undesirable function (e.g., corrupting data generated within a given time window). The most likely target of these attacks would be the data flowing from the sensors to the IDC; however, this group might also be able to use its authorized access to attack other elements of the IMS.

IDC insiders

An IDC insider is any person whose job relates to the operation of the IDC elements of the IMS. These could include:

- Analysts.
- System administrators.
- Computer maintenance personnel.

As with the IMS user group, the primary concern for this set of potential attackers is the use of authorized access to subvert system operation; however, in addition to raw sensor data, likely targets also include IDC products (e.g., event bulletins) and the processes used to produce them. Among other things, an IDC insider might alter stored records, change trustworthiness ratings for various stations, reveal IMS state-of-health testing schedules to countries wishing to covertly test (so that the countries know when specific sensors are not monitoring activities), or desensitize detection algorithms by altering key parameter files.

Hackers and outside agents

A hacker is characterized as having electronic, but not physical, access to certain IMS elements by virtue of the fact that these elements may be used by the IMS but not owned by it (e.g., Internet routers). This person has no authorization to IMS-specific assets. An outside agent faces the same constraints as the hacker but is differentiated by motivation and potential resources. The hacker is characterized as working alone and for the "thrill of it," and has relatively simple resources, whereas the outside agent may work as part of a team of professionals with very specific objectives in mind (e.g., intelligence operations) and with significant financial and technical resources. The primary concerns with this attacker are the subversion of IMS nodes connected to publicly accessible communications and the theft of computer resources (e.g., stealing CPU time or communications bandwidth, or using IMS disk space for sharing files with other hackers).

Environment and system design weaknesses

Environment includes all the natural phenomena associated with the locations in which the IMS elements operate. This could include lightning, flooding, high winds, wildlife, normal human activities in the area (e.g., fishing near hydro-acoustic sensors), etc. System design weaknesses include those features built into the system that render the IMS vulnerable to single-point failure. As opposed to the maliciously inserted functionality described earlier, this threat agent focuses on potential oversights and poor design and implementation practices of the IMS engineering community. The key issue to be considered with these weaknesses is whether or not the resulting failures are critical. If a sensor fails, all data from the region covered by that sensor is lost until repairs are made. If a communication link fails, all data must be buffered at the link's sending node until communications are reestablished. In the first case, data will be lost. In the second case, no data need be lost, assuming that the buffer's capacity is large enough to meet all reasonable repair time requirements.

3. Security Attributes

Within this document, security attributes are grouped into six categories:

- Integrity
- Authenticity and nonrepudiation
- Availability
- Access control
- Accountability and intrusion detection
- Visibility

In the sections that follow, each of these categories is defined in general terms and then the specific attributes for each category are enumerated and discussed. It should be noted that these attributes are statements about the security of the system as a whole and not about individual elements. How these attributes are implemented across the system and within individual system elements is left to a later discussion.

3.1 Integrity

The term "integrity" typically means data items have not been corrupted in some way. For example, in noisy communication environments, structures such as checksums are commonly used to determine whether or not sets of data have traversed the communications channels without being altered by noise on the channel (i.e., that the sets' "integrity" has been maintained). In the context of the CTBT IMS, the goal of integrity mechanisms is to ensure that, from the time that a given event generates a set of signals (e.g., seismic waveforms) until that event's data and characteristics are reported to IMS users, the data generated by the system is not corrupted. This goal recognizes that this corruption may be accidental — an artifact of natural phenomena or human error — or intentional.

Attribute 1: It must not be possible to alter or delete an IMS data product without the alteration or deletion being detected. [Stated: Annex 5, Section 1.2.1 and Section 1.5 of *GSETT-3 System Security Plan*, Ref. 1.]

Discussion: As illustrated by the GSETT-3 work, a number of data products are possible in a CTBT IMS. These include source data generated by the IMS sensing network as well as products (e.g., event bulletins) generated from this data by the various system processing nodes (e.g., the International Data Center or National Data Centers). In addition, this data may be "active" (i.e., newly produced) or archived.

Whether intentional or not, the effects of altering or deleting these products could include:

- completely hiding that a given event occurred, and
- changing the apparent characteristics of a detected event (event type, location, etc.).

In addition to this, alteration could create the appearance that an event has occurred, even though it never really did.

In considering this recommended security attribute, special attention should be paid to the qualifying statement that alteration or deletion should not occur "without detection." This acknowledges that there should be more emphasis placed on detection than

on prevention of alteration or deletion. This is significant in that it permits the development of security solutions that are sufficiently strong without being excessively expensive to implement.

Attribute 2: The source of an IMS data product must not be subject to being influenced to generate incorrect data without the influence being detected. [Stated: Annex 5, Section 1.5.][1]

Discussion: Whereas the first attribute in this section addresses integrity issues with respect to IMS products, this second attribute addresses the processes and equipment used to generate these products. Typical problems addressed by this attribute might include:

- Altering the data source's inputs.
- Altering other data that controls how the data of interest is created.
- Altering the processing done on the data.

Examples of the first case include augmenting or damping the signals arriving at sensors to alter the readings in some desired way. In this case, the equipment functions as designed but creates bogus data. This technique might be used on an on-going basis to tailor a site's ambient signals in such a way as to mask specific types of events. Similarly, if the IDC were somehow tricked into accepting bogus waveform data, then its event bulletins would be erroneous. In the second case, processing-related variables, such as threshold limits on sensing elements or "recipe files" in IDC "pipelines" could be altered such that fallacious data would be produced by the IMS. In the third case, processes could be replaced outright by invalid routines or modified in specific ways to produce invalid results.

In addition to the effects already cited for alteration or deletion of data, an attack like this might be intended to destroy user community confidence in the IMS's ability to execute by causing the system to alarm when no event has occurred or to do nothing when an event has clearly occurred.

3.2 Authenticity and non-repudiation

Authenticity typically refers to the ability to determine that a given transaction was initiated by a given source. In conjunction with integrity mechanisms, this establishes a user's trust in the usefulness and reliability of the transaction. Nonrepudiation mechanisms ensure that a transaction's origin can be uniquely traced back to its source by third parties.

While the mechanisms used to ensure integrity, authenticity, and non-repudiation are often very similar, it is worth noting that it is possible in a system like the IMS to have integrity without authenticity, authenticity with the ability to repudiate, and non-repudiation without integrity. It all comes down to how the system as a whole is designed.

For example, checksums are commonly used as an integrity mechanism in environments where corruption of data is possible (e.g., noisy communications). When a message is received in this kind of environment, the receiver computes a checksum over the message body and compares it with the

checksum received with the body. Matching checksums show that integrity has been maintained since the message left the transmitter, but they say nothing about the sender. In this kind of system, an adversary could create his own message with a correct checksum and inject it into the channel and no one would be the wiser. Similarly, if the system relies on a more secure mechanism, such as cryptographic checksums, to establish authenticity and integrity, an adversary may be able to record and then replay messages generated by a valid transmitter if the messages are not properly designed. As a final example of how these attributes are related and yet separate, even if the IMS design incorporated mechanisms that prevented replay and other attacks that might circumvent authenticity, designs are possible that permit a receiver to create a message like those coming from the system's transmitters. In this case, a third party could not prove the actual source of the message even though the messages are authentic (in as much as they are generated within the system boundary).

Within the CTBT IMS, there are two primary flows of information — raw data flowing from sensors to processing nodes (IDC and NDCs) and products (e.g., event bulletins) generated from this data and flowing back to system users. In both cases, CTBT users need to believe in the authenticity of the information. Both the processing nodes and the consumers of their products need to believe that the source data used in developing the products is not corrupted in any way. In addition, users of the products need to believe that the products have been correctly produced and remain unchanged since their production. At the same time, the individual countries or regions monitored by the IMS need to know that neither the system nor any of its users can generate authenticatable messages that implicate the countries or regions in events that never actually happened.

Given these needs, the suggested authenticity and nonrepudiation attributes for the CTBT IMS are as follows. (For more detail on these issues, see the SAND report *Authentication of Data for Monitoring a Comprehensive Test Ban Treaty.*)[3]

Attribute 1: An adversary must not be able to forge an IMS data product without the forgery being detected. [Stated: Annex 5, Section 1.2.1, Paragraph 2.][1]

Discussion: The goal here is to ensure that entities outside the IMS (including countries and regions making use of the IMS) are not able to create data products that are accepted as legitimate. The primary difference between this attribute and the integrity attributes is that this attribute acknowledges that there may be an active adversary, whereas the integrity attributes would suffice if there were only concerns about environmental factors corrupting IMS products.

Before ending this discussion, a couple of comments are worth making: First, the emphasis here is again on detection rather than prevention. The reason for stating it this way is that establishing authenticity in the IMS may not be a matter of applying technical measures but relying, instead, on system phenomenology. This leads us to the second point: there is a debate as to

whether the IMS requires any overt authentication capability at all. The argument is that events of interest to the CTBT community will fall into two categories: those that are of sufficient magnitude to be seen by a "large" number of IMS sensors and those that are so small as to be seen by only a few sensors.

In the case of the first, it has been suggested that the IMS is "self-authenticating" because an adversary would have difficulty altering the data streams from all of the sensors that picked up the large event's signature. Similarly, if a central processing node, like the IDC, were to falsify event data, the associated source data could still reside at the sensor stations or their NDCs, thus permitting the international monitoring community to discover the falsification at some point. In addition to these storage locations, a large number of non-IMS stations also operate throughout the world. The likelihood that some set of these will also see and independently report on an event lends credence to the notion that overt authentication measures may not be required in the IMS. The other side of the argument is that whether or not an adversary can successfully falsify an event really depends on the adversary's span-of-control within the IMS. If all data channels of interest pass through nodes controlled by the adversary, or if an adversary can extend his span-of-control by "hacking" nodes in the IMS, then the system may not be self-authenticating at all. In addition, if no coordinated reporting system exists for the non-IMS stations, there is some question as to how effective this "inherent safeguard" really is.

In the case of localized events, the span-of-control problem is much more likely to be a significant issue. In this case, using technical authentication means (e.g, digital signatures) and strong design features that inhibit loss of data because of single-point failures seems prudent.

Attribute 2: IMS product users must be able to authenticate the true source of all data (including that which has been archived). [Stated: Annex 5, Section 1.5.][1]

Discussion: This attribute elaborates on the previous by insisting that the recipient of a data product not only trust that the product is valid but also know that the product came from a given source. This guards against attacks where trusted sources try to impersonate other trusted sources to achieve the goals discussed in the previous section.

As an example of this problem, consider a system in which all sensors are authenticated with cryptographic signatures and all sensors use the same key. In this case, even if the secrecy of the keys was maintained, the receiver of the sensor messages could not be sure that a given message came from a given sensor. The only sure thing would be that it came from one of the sensors in the system.

Attribute 3: A source of data must not be able to repudiate origination of that data. [Stated: Annex 5, Section 1.5.][1]

Discussion: This attribute tightens the envelope even further by saying that not only does a receiver need to know that a data product came from a given source, but the receiver must also be able to prove to a third party that the product came from that source.

This attribute can be levied to protect both receiver and sender. In the first case, the receiver can use the data in good faith, confident that if a product based on that source data is shown to be erroneous because of the source data, then the receiver can prove his innocence in the matter. In the case of a sender, the attribute guards the sender against accusations of creating false data that the sender really did not create.

A typical example of a system in which this would fail is one in which a message's authenticator is created using private key cryptography. In this kind of system, the cryptographic key used by the sender to sign the message is the same key used by the receiver to authenticate the signature. A third party looking at this system might be convinced that a given message is authentic (i.e., was produced by a node in the system) but could not be convinced that it necessarily came from a given sender.

Attribute 4: Sensors and other IMS elements must respond only to authentic commands. [Stated: Annex 5, Section 1.5.][1]

Discussion: There are two issues here: (1) the element recognizing that the command came from an authorized user, and (2) the element recognizing that this is a unique message — distinct from any other previously generated by that user. The first issue acknowledges that IMS control functions should be limited to a subset of the IMS user population. The second addresses the notion that once executed, commands should not be replayable (lest authentic commands be captured by an adversary and then played at times convenient to his purposes). With respect to IMS sensors, this attribute is a necessary part of guarding against an adversary's ability to make the IMS "deaf" during the period that an event is occurring.

3.3 Availability

Two issues are typically addressed by availability:

- A system must not lose any data that it generates.
- The system design must ensure that data can be delivered fast enough to be useful.

From discussions with CTBT participants, it appears that these two aspects of availability are of roughly equal importance. Preservation of data is likely to be a key aspect of safeguarding against integrity and authenticity concerns. Timeliness of the system is important from a political viewpoint because the political officials served by the IMS would like to know about the occurrence of events before news reporters appear on their doorsteps asking for comments, as well as from a system effectiveness perspective, since timely data may affect the viability of subsequent activities (e.g., on-site inspection).

Attribute 1: A source of data must not lose data, even when responding to authentic commands. [Stated: Annex 5, Section 1.5.][1]

Discussion: There are two potential attacks here if the system is not designed correctly. First, an adversary might be able to mask detection of an event by placing the relevant sensors in a state (e.g., calibration) in which data is not

collected during the event. Second, an adversary (or even a system failure) might be able to swamp a sensor by continuously issuing authentic requests for service. Either way the sensor runs the potential of missing data that should have been captured.

Attribute 2: IMS product data must be universally available to authorized users. [Stated: Annex 5, Section 1.2.1, Paragraph 2.][1]

Discussion: Two aspects of availability are implied by this statement:

- The system has the ability to preserve the data it generates.
- The system has the ability to deliver data products within given time intervals to the required location.

With respect to the first issue, care must be taken with the IMS design to ensure that single-point failures cannot cause data produced by the IMS to be permanently lost. This means that data cannot exist in only one point in the system at any given time (e.g., on a single disk volume) and that protocols must accommodate the possibility of failures (e.g., data collected by a sensor site cannot be lost due to temporary loss of the communications channel fed by the sensor site). Multipoint failures should be of sufficiently low probability to ensure that data loss due to these failures does not degrade system performance.

With respect to the second issue, consideration needs to be given as to which, if any, of the IMS communication links and processes need to respond within a given amount of time. If any such links exist, steps need to be taken to prevent adversaries from being able to launch denial of service attacks. It should be noted that this attribute also implies that in specifying the final CTBT IMS design, maximum acceptable delays between specified system nodes, maximum processing times, acceptable data loss rates, etc. should be set.

Attribute 3: Product utility must be maintained [ ].

Discussion: Depending on the nature of the safeguards implemented, it may be possible for system users to lose the ability to use system products even though the products themselves are still available. For example, if products are cryptographically "signed" and the associated keys are stored only in a central location, loss of the key used to verify the signature could make authentication of the products impossible.

3.4 Access control

Access control refers to a system's ability to limit an entity's use of a system resource (i.e., data, processes, hardware, etc.). As such, access control has three constituent functions:

- Identification and authentication.
- Authorization.
- Access-control mechanisms.

Identification and authentication

In information security, identification and authentication is the process of establishing which user is requesting access to a system's data or resources. This user can be either a human or another computer. For human users, identity is commonly asserted (identification) by means of a user ID and

then verified (authentication) by means of a password; however, other techniques are used as well. These are typically categorized as:

- things that are known (passwords, pass phrases, challenge-response pairs, etc.),
- things that are possessed (tokens, keys, etc.), and
- attributes of the individual (retinal patterns, thumbprints, voice patterns, etc.).

For computers requesting access from other computers, different techniques exist. Among other things, these include cryptographic mechanisms (e.g., signatures based on shared private keys or certificate-based systems) and addressing techniques (i.e., providing service based on the network address information associated with the requesting computer).

Authorization

Authorization is the process by which a given user's right to access specified resources and data is established. Accessed resources can include both physical information system elements (e.g., disk drives or communications ports) and system programs/processes (editors, compilers, mail daemons, etc.). Depending on the system, the access-control policies embodied by authorization may be simple grant/deny schemes or may include more sophisticated measures, such as limiting rights to specified periods during the day.

Access-control mechanisms

Access-control mechanisms are those devices and procedures used to enforce the specified authorizations. Their effect may be to limit both physical access and logical access to resources. While constraints may be placed upon these mechanisms (e.g., "two-person control is not viable"), none have yet been identified for CTBT.

Under the CTBT, it is expected that users wishing to access system data and resources will fall into three categories:

- national and regional representatives charged by their respective governments with monitoring treaty compliance,
- those who can make use of IMS-generated data for other purposes (e.g., seismic research), and
- those who require access to system resources to keep the monitoring system operational.

There will likely be a need to limit access to those users authorized by system administrators (as opposed to leaving the system completely open to the world). In particular, anticipated communications capacity limitations may necessitate restricting the amount of data that any given user can request within a specified time interval. In addition, there will almost certainly be a need to limit "write privileges" for certain critical data (e.g., variables used to control the sensitivity of event-detection algorithms) to a closely monitored set of users. In addition, not every user will require access to every piece of data and to every system resource. For example, the average user charged with monitoring treaty compliance has little reason to inspect system audit logs. Similarly, a person authorized to change security-critical data should not have the ability to

also edit the logs regarding that data. In general, the philosophy of "least privilege" should be applied wherever possible (i.e., users should be given access to only those things needed to do their jobs and to nothing more).

The access control attributes that have been proposed for CTBT are as follows:

Attribute 1: The IMS must support the ability to ascertain the identity of users requesting access to system data or resources. [Implied: Volume 2, Part 3, Section 4.2.7 of *Operations Manual — Volume 2, Part 3: IDC Operations*, Ref. 2.]

Discussion: Each user must have some means of identifying himself to the system. This mechanism must provide the system with confidence (up to some as yet unspecified level) that the user is, in fact, who he claims to be. It should be noted that the level of trust required in identifying users may vary across the system. For example, it is certainly less of an issue if an unauthorized user reads system-generated data than if that same user is able to alter processes in such a way that system performance is degraded. Consequently, it may be desirable in design to apply stronger (and typically more costly) controls only where the consequences are more severe.

Traditional schemes, such as login ID and password, carry certain inherent weaknesses in distributed environments like the CTBT IMS. "Sniffing attacks" in recent years have pointed to the fact that it is quite easy for a relatively unsophisticated adversary to gain access to a system through capturing login IDs and passwords. As a result, techniques like one-time passwords have gained in popularity. On the other hand, these carry their own unique costs that may or may not be acceptable to the CTBT user community, the CTBT development community, or both. Irrespective, whatever scheme is used must provide sufficient uniqueness in identification to permit accountability requirements to be met (see Section 3.5).

Finally, note that identification can work in two directions: In the first, the IMS must be able to determine that a requester is, in fact, an authorized user. In the second, a user must be able to determine that the system to which he is connected is really the CTBT IMS.

Attribute 2: The IMS must have the ability to control which users have access to given system products and resources and to control the type of access permitted. [Implied: Volume 2, Part 3, Section 4.2.7.][2]

Discussion: Several issues are implied by this attribute: The first is that the CTBT community must completely identify the data and resources that constitute the IMS. These elements may be treated individually or in groups that are accessed as a whole. To a degree, how this is done may be determined by the inherent access-control mechanisms on the hardware selected for implementation.

Second, consideration needs to be given to how privileges will be divided such that no one user has sufficient access to corrupt IMS operation. For example, the staff from a region being monitored by a given sensor probably should not have the ability to configure the authentication capabilities

of that sensor; otherwise, that region may be able to falsify the data generated by the sensor. Similarly, information about IMS operation in a given region (e.g., when IMS sensors in that region will be "off-line" for self-test) should be kept from those IMS users with affiliations in that region.

Third, a decision will need to be made on how to grant access to users — as individuals or as groups. Granting rights to individual users creates more work for the person(s) charged with authorizing access but provides finer control granularity. Authorizing individuals as groups permits the authorizing agent to set privileges once for a group and to then simply enroll and disenroll individuals as group members. One additional consideration in this kind of scheme is whether or not individuals can be enrolled in multiple groups. If so, then care must be taken to not permit an individual to collect excessive privilege (per the second issue above).

Fourth, the question of who can grant what access privileges must be answered. The two obvious choices are centralized control (e.g., the IDC dictates what data and resources on what system nodes a given user can access) and distributed control (e.g., each node in the IMS controls access to data owned by that node). If control is distributed, some mechanism still needs to be put in place to ensure that no one individual collects excessive privilege. On the other hand, if centralizing control means that IDC operators have the ability to grant access to data and resources located at NDCs and stations, there may be some justifiable level of paranoia on the part of the owners of these nodes regarding their ability to protect their data and resources.

One authorization issue particularly worth noting is the confidentiality of subscription information. It has been suggested that some countries may not want others to know what they are monitoring. To truly safeguard this information will require protecting both the subscription information — both in transit between the requester and the IDC and while in storage at the IDC — and the products generated as a result of this information. In addition, traffic analysis issues may come into play here as well. For example, if an outsider can observe that every time an event is received from some portion of the world, a product is sent to a given subscriber, the outsider might be able to deduce something about the subscriber's monitoring interests.

One solution that has been suggested for this problem is to simply subscribe to everything that the IMS produces (or at least to some superset of the data of interest). The question that this approach naturally raises is whether the IMS has sufficient bandwidth to support this approach if every system user chooses to oversubscribe.

3.5 Accountability and intrusion detection

In this document, accountability refers to the ability to track security-critical actions taken by authorized users. Intrusion detection refers to the ability to assess whether unauthorized users have accessed IMS data and system resources.

As noted earlier, certain parts of the IMS are likely to be more security-critical than others. For these sections of the IMS, there is a real need to be able to monitor all security-relevant activities. The goal of this monitoring is two-fold: First, it serves as a safeguard to keep honest users honest and to cause other users to think twice about the possibility of being caught. Second, the information gained through this auditing serves as the basis for assessing the need for additional information safeguards once the system is deployed.

Before moving to the CTBT IMS accountability and intrusion detection attributes, it is worthwhile to comment on this last point. In safeguarding an information system against security threats, two radically different approaches might be taken:

(1) In the first, the system is analyzed in depth to determine any possible threats and the associated vulnerabilities that might permit these threats to be realized. Safeguards are then put in place to remove these vulnerabilities.

(2) In the second approach, a "bare" system is put into operation and all operations on the system are monitored and assessed. As certain operations deemed undesirable are detected, only the safeguards necessary to inhibit these operations are put into place.

While the first approach may field a strong system, it can be costly to realize and may provide solutions for problems that never arise. In the case of the second approach, "damage" may already be done by the time that audit logs reveal that it has occurred. While safeguards may be put into place to prevent future occurrences of the violation, they still haven't guarded against the original violation. This document proposes a middle-of-the-road approach where initial thought is given as to what the security goals of the IMS should be, the essential safeguards are installed, and then the system is monitored to ensure that an appropriate level of security is maintained. This approach values security while accepting "good enough" in place of "perfect."

Given this, the following accountability and intrusion detection attributes are suggested:

Attribute 1: IMS operators must be able to reconstruct all events relevant to elements critical to information security [ ].

Discussion: The goal here is to be able to detect that attacks are being launched so that the need for additional information security safeguards can be assessed. It should be noted that this kind of protection must begin during the system's development cycle (e.g., to guard against insertion of malicious code) and extend into its operation (to detect both internal and external attacks) and retirement (to guard against loss of confidentiality or authenticity if such requirements exist).

The primary issue here is determining what to audit. In this world, more is not always better. If too much data is collected, resources (e.g., disk space) can be overwhelmed. More importantly, too much information can make the auditor's job difficult by burying the critical information in a sea of data. On the other hand, if little thought is given as to what information is critical or if the system neglects to collect

this information, then personnel charged with maintaining system security may remain oblivious to very real threats.

Attribute 2: An adversary must not be able to subvert an information safeguard without detection. [Stated: Annex 5, Section 1.5.][1]

Discussion: It is difficult, if not impossible, to create systems that cannot be broken by a determined adversary; therefore, the goal of this attribute is to create a system that permits designated authorities to determine if information safeguards have been tampered with and that, to the degree possible, helps them track down who did the tampering.

Attribute 3: Audit data and processes shall be afforded at least the same integrity, authenticity, confidentiality, and availability considerations as primary product data. [Implied: Annex 5, Section 1.5.][1]

Discussion: If the system security philosophy relies as much or more on detection as it does on prevention, then the processes used to detect security violations and the data generated by these processes must be strictly protected. An adversary must not be able to corrupt audit logs. Auditors must be able to trust that the data they see in the logs has really been produced by the audit mechanisms. It may be desirable to ensure that audit data only be seen by designated authorities. Finally, the system must be able to preserve its audit data from loss and must be able to deliver the data to the designated authorities when required.

3.6 Visibility

As no part of the normal security parlance seems to apply to these requirements, the term "visibility" is used here to apply to a system's ability to ensure an openness or transparency to an operation or set of information. In a sense, this is the inverse of the access-control problem. As opposed to limiting disclosure to a specified set of entities, these requirements mandate full disclosure.

The visibility attributes suggested to date are as follows:

Attribute 1: IMS product users must be able to determine how each product was produced in order to independently reconstruct the processing events when desired [ ].

Discussion: It has been suggested that the users of IMS products (e.g., event bulletins) need to be able to determine how those products were produced. This includes knowing from what sensor data the product was derived, what algorithms were applied to that data and in what order, and what parameters were used to drive those algorithms.

Attribute 2: Countries or regions being monitored by IMS sensors shall be able to determine the exact information content of messages generated by those sensors [ ].

Discussion: The key concern here is that some safeguards that might be proposed would make it difficult (if not impossible) for the countries being monitored to determine what information is leaving their borders. For example, if a private key encryption algorithm were used to create a message authenticator for signing sensor data, the fact that the

cryptographic key would have to be kept secret from the monitored country would allow the authenticator bits to be used as a data channel. Solutions to this include the use of public key cryptography (which permits the host country to verify the signatures on all data) or the collection of signed messages by the host nation until such time as the private key (in a private key signature scheme) is made public and the country is able to verify the signatures.

A related problem is the existence of covert channels that arise because of protocol design or implementation or even just natural artifacts of the application domain. As an example, if a time stamp were placed in sensor-generated messages, then surreptitious code in the sensors could be used to modulate the interval between time stamps to pass covert information. For that matter, if the data itself contains a certain amount of noise as captured in the lowest N bits of a reading, these bits might successfully be used as a channel. Note that these same mechanisms could be used to control malicious code inside an IMS node. For example, even if unauthorized messages are rejected by an IMS node as invalid, the pattern of their arrival can be used to convey information to the malicious code.

4. Summary

This paper has addressed four active threats to the IMS:

- Hiding an event,
- Altering what the IMS detects about an event,
- Creating an event that never happened, and
- Destroying trust in the IMS's integrity.

This paper has also identified potential threat agents:

- IMS designers, implementors, and installers,
- Authorized IMS users,
- IDC insiders,
- Hackers and outside agents, and
- Environmental and system design weaknesses.

In addition, six sets of desirable security attributes for the IMS have been suggested. These are:

- **Integrity**
  1. It must not be possible to alter or delete an IMS data product without the alteration or deletion being detected. [Stated: Annex 5, Section 1.2.1 and Section 1.5.][1]
  2. The source of an IMS data product must not be subject to being influenced to generate incorrect data without the influence being detected. [Stated: Annex 5, Section 1.5.][1]
- **Authenticity and nonrepudiation**
  1. An adversary must not be able to forge an IMS data product without the forgery being detected. [Stated: Annex 5, Section 1.2.1, Paragraph 2.][1]
  2. IMS product users must be able to authenticate the true source of all data (including that which has been archived). [Stated: Annex 5, Section 1.5.][1]
  3. A source of data must not be able to repudiate origination of that data. [Stated: Annex 5, Section 1.5.][1]
  4. Sensors and other IMS elements must respond only to authentic commands. [Stated: Annex 5, Section 1.5.][1]
- **Availability**
  1. A source of data must not lose data, even when responding to authentic commands. [Stated: Annex 5, Section 1.5.][1]
  2. IMS product data must be universally available to authorized users. [Stated: Annex 5, Section 1.2.1, Paragraph 2.][1]
  3. Product utility must be maintained [ ].
- **Access control**
  1. The IMS must support the ability to ascertain the identity of users requesting access to system data or resources. [Implied: Volume 2, Part 3, Section 4.2.7.][2]

2. The IMS must have the ability to control which users have access to given system products and resources and to control the type of access permitted. [Implied: Volume 2, Part 3, Section 4.2.7.][2]

- **Accountability and intrusion detection**

  1. IMS operators must be able to reconstruct all events relevant to elements critical to information security [ ].
  2. An adversary must not be able to subvert an information safeguard without detection. [Stated: Annex 5, Section 1.5.][1]
  3. Audit data and processes shall be afforded at least the same integrity, authenticity, confidentiality, and availability considerations as primary product data. [Implied: Annex 5, Section 1.5.][1]

- **Visibility**

  1. IMS product users must be able to determine how each product was produced in order to independently reconstruct the processing events when desired [ ].
  2. Countries or regions being monitored by IMS sensors shall be able to determine the exact information content of messages generated by those sensors [ ].

5. Future Work

While most of the ideas in this paper are not new with this paper (or with CTBT, for that matter), the intent in producing this paper is to stimulate discussion in the U.S. CTBT community as to the validity of, and the issues surrounding, each suggested threat and security attribute. Given this discussion, sound decisions can be made with respect to the principles upon which the U.S. CTBT community wishes to base the IMS's system-level security architecture and component-level safeguards. Consensus with respect to these threats also serves as a basis for judging the relevance of various IMS attack scenarios and of assessing the relative value of potential safeguards.

Safeguards have already been proposed to address some of the security issues associated with the IMS and some of these have real value; however, one point should be clearly understood: **the IMS is a system and should be viewed as such in considering how to protect it**. In any system of reasonable size, an adversary usually has more than one way to achieve objectives. This fact demands that those charged with designing protection for a system understand the various ways in which the system might reasonably be attacked. At the same time, considering attacks and proposing safeguards — without having clearly identified what is to be protected (and *against* what protection is needed) — puts one at risk of squandering limited security resources on non-issues while inadequately safeguarding those things that genuinely need protection. This is the reason for stimulating discussion of IMS security within the U.S. CTBT community.

Great effort has been expended to lay a foundation of sound technical theory upon which an IMS can be built. In this research environment, realizing and extending functionality of the IMS has been paramount.

Consequently, system security, while important, has received less attention. Staff working some elements of the system, especially the sensing subsystems, have had little opportunity to consider security issues at all. Even so, the time left to affect treaty language with respect to system security is quickly running out. The U.S. government's goal is to close out changes to the language by March 29, 1996.

Between now and the time that the IMS is fielded, a number of steps will have to be taken in the process to secure the IMS. These include:

- system-level design activities.
  - — agreement upon the security objectives for the IMS and the likely threat agents.
  - — identification of viable attacks, given the system design.
  - — proposal of potential system-level safeguards and analysis of relative merits and costs.
- subsystem-/component-level design activities.
  - — assessment of risks associated with the detailed IMS design.
  - — proposal of potential component-level safeguards and analysis of relative merits and costs.

- implementation and installation activities.
  - — independent evaluation of implementations.
  - — secure installation of IMS elements.
- operational activities
  - — periodic auditing of IMS security elements.
  - — assessment of new threats discovered as technology evolves and installation of appropriate safeguards.

At each step, engineering tradeoffs will have to be made to balance the need for security against funding limitations and usability of the IMS.

This paper addresses the first task under system-level design activities. (A second Sandia-produced paper — *A Discussion of Attacks on Comprehensive Test Ban Treaty International Monitoring Systems* — that is nearing completion addresses the second and third tasks in the same category.) In addition to addressing specific attacks and potential safeguards for the system architecture, the paper proposes general principles to be followed in securing the IMS. Sandia has also been talking with sensor subsystem developers in the U.S., to begin the assessment portion of the component-level design activities. An additional goal of this work has been to sensitize component developers with respect to the ways in which their systems may be attacked and safeguards that might be applied to secure the components against these attacks. Sandia has also evaluated (see SAND report *Authentication of Data for Monitoring a Comprehensive Test Ban Treaty*[3]) the three architectural options for sensor data authentication that have been proposed by different elements of the U.S. CTBT community and has recommended a particular option. Finally, the authentication experiment being executed by ARPA, along with its recommendations for an IMS security management infrastructure, also addresses some of the needs of the system-level and component-level design tasks.

At the moment, much of the IMS work is best classified as exploratory. A general system design exists but various particulars are yet to be decided. Many of the specific component implementations that exist (both hardware and software) are for the purpose of evaluation and some even exist only on paper. Until it is clear that particular designs and implementations will be used in the actual IMS, the final two sets of tasks will remain "to be done." At this point in the system's development, the right kind of security work has been done by individual agencies. Now the U.S. CTBT community as a whole needs to discuss the issues and decide on a joint position with respect to the first two sets of tasks (system-level design and component-level design).

References

1. **GSETT-3 System Security Plan—Annex 5**, Conference on Disarmament Group of Scientific Experts, February 1995.

2. **Operations Manual—Volume 2, Part 3: IDC Operations**, Conference on Disarmament Group of Scientific Experts, February 1995.

3. **Authentication of Data for Monitoring a Comprehensive Test Ban Treaty**, Sandia National Laboratories, November 1995.

DISTRIBUTION

3 Air Force Technical Applications Center
Attn: Frank Pilotte, TT
David Russell, TTR
Bruce Varnum, TTD
1030 S. Highway A-1A
Patrick AFB, FL 32925-3002

1 DCI/ACIS
Attn: Larry Turnbull, 4W03
New Headquarters Building
Washington, D.C. 20505

2 Department of Energy
Attn: Leslie Casey, NN-20
Dave Watkins, NN-40
1000 Independence Avenue, SW
Washington, D.C. 20585

2 Lawrence Livermore National Laboratory
Attn: Dave Harris, L-205
Jay Zucca, L-205
P.O. Box 808
Livermore, CA 94551

3 Los Alamos National Laboratory
Attn: Wendee Brunish, MS F659
Mark Hodgson, MS D460
Rod Whitaker, MS F665
P.O. Box 1663
Los Alamos, NM 87545

2 Pacific Northwest Laboratory
Attn: Rich Hanlen, MS K6-48
Ray Warner, MS K6-40
P.O. Box 999
Richland, WA 99352

1 MS	0458	Laura Gilliom, 5133
1	0419	Bob Gough, 5336
1	0979	Dale Breding, 5704
1	0979	Larry Walker, 5704
5	0655	Tim Draelos, 5736
1	0655	Pres Herrington, 5736
5	0451	Rick Craft, 9415
1	0451	Judy Moore, 9415
1	1138	Ralph Keyser, 9432
1	0619	Tammy Locke, 12615
1	9018	Central Technical Files, 8523-2
5	0899	Technical Library, 4414
1	0619	Print Media, 12615
2	0100	Document Processing, 7613-2, for DOE/OSTI