Forming an Incident Response Team

Danny Smith
Australian Computer Emergency Response Team
c/- Prentice Centre
The University of Queensland
Brisbane, Qld, 4072.
Australia

auscert@auscert.org.au
24 Hour Hotline: +61 7 365 4417
Fax: +61 7 365 4477

Abstract

Forming an Incident Response Team (IRT) in the 1990s can be a daunting task. Many people forming an IRT have no experience with doing this. This paper examines the role an IRT may play in the community, and the issues that should be addressed both during the formation and after commencement of operations. It may be of benefit to existing IRTs as it may raise awareness of issues not previously addressed.

1. Introduction

On 8 March 1993, the Security Emergency Response Team (SERT) commenced Incident Response operations in Australia. Prior to this, discussions had been held with other Incident Response Teams (IRTs) to discuss the establishment of the team, and what would be required. A significant amount of work was performed just before, and immediately after, commencement to establish operations and tools. Further communication with other IRTs assisted SERT to establish policies and helped SERT to grow in its own constituency, and the computer security community at large.

Since that time, SERT has undergone many changes, and these transitions could not have been effected as smoothly as they were without the work that had been achieved earlier.

This paper looks at the topic of what it takes to form an IRT. It examines what issues need to be addressed and resolved prior to, and after, forming an IRT. It looks at the constituency, policies, relationships, information, equipment, tools, and interaction with the wider community.

Much of the information in this paper is not new. It has been steadily collected from a number of sources over time, and various amounts of it have been applied by SERT with varying success. The overwhelming message throughout this paper is:

"You are not alone!".

2. How did SERT start?

Dateline 1992: The Australian Academic and Research Network (AARNet) has been running for two years. It connects all the academic and research institutions in Australia, which is now the third largest country on the Internet in terms of connected hosts[1]. During this year, many institutions started experiencing a dramatic increase in the number of computer security intrusions, particularly network based attacks.

This was a new problem for Australia to face. In the past, most attacks originated locally, and were dealt with by local institution statutes. At first, the attacks had nuisance value, but they soon started to reach plague proportions. Australia was used as a launch pad for attacks to overseas sites.

One particular group of individuals[2] concentrated on the South-East Queensland corner, and used three Universities in particular. From here, they launched attacks to overseas institutions, which ultimately threatened a large amount of research funding coming into the country.

Cooperation between these three Universities was always extremely good, and only a coordinated response to this problem resulted in the apprehension of the intruders. It was during these attacks that a decision was made that Australia was large enough that it must fend for itself in the international arena.

The decision was that an Incident Response Team was required. Much talk was generated on the topic, but no progress was made. The three Brisbane-based Universities, Queensland University of Technology, Griffith University, and The University of Queensland, combined their efforts and applied to the Federal government for funds to establish a response team. Late in 1992, this application was rejected by the government.

The Universities then made a crucial decision: the IRT was essential, so they decided to just start it anyway, and fund it themselves. During February, a number of staff members worked hard to get a team to operational readiness as quickly as possible. This included developing crude tools for incident tracking, and establishing a secure premises to operate from. On 8-Mar-1993, the SERT team was announced to its constituency.

It was also during this time, that SERT communicated heavily with the CERT Coordination Centre in Pittsburgh. SERT outlined their intentions to commence operations, and received an enormous amount of assistance from CERT. Much electronic mail was exchanged, which finally culminated in a conference telephone call between the two teams. During this phone call, the two teams exchanged ideas on the issues that the SERT team would need to address to become operational.

Subsequent communications between CERT and SERT clarified and defined how the two teams would interoperate. Many issues relating to the international nature of the interaction required resolution, even the little things such as date format. Is 1/8/94 the 1st of August or the 8th of January? It was during this time that CERT was forming a new relationship with the DFNCERT team in Germany as well.

Subsequent incidents highlighted shortcomings in the operation, which were addressed and rectified as time went on. Many of SERT's problems stemmed from the way it was formed: it had no authority to act, it just existed. Convincing the community that the SERT team was essential was a hard and long task. This was achieved through steady and constant communications, dedication to assisting sites with security problems, and as large a public exposure as could be achieved without burning out staff, or destroying the travel budget. Acceptance was finally realised when most of the Computer Centre Directors in the Australian Universities contributed funds towards SERT's operation.

---

[1]Since relegated to fourth position!

[2]One of this group has since pleaded guilty to 166 charges of unauthorised access of a computer system. Other members of the group are still awaiting trial.

Since that time, SERT has been transformed into AUSCERT with a formal contract signed between AUSCERT and AARNet. AUSCERT now acts with the authority of AARNet, and is seeking to extend its constituency to the whole of Australia and beyond. The operation commenced on 1 April 1994.

The rest of this paper looks at the decisions that were made, or were advised should be made. At times, comparisons are drawn between SERT and CERT, to highlight some fundamental differences in the two operations. This comparison highlights advantages and disadvantages of the two types of teams.

3. Pre-establishment Tasks

Having decided (or been directed) to form an Incident Response Team, there are a number of tasks that can be completed before commencing operation. The ultimate success of the ongoing team may be the direct result of how well some of these tasks have been completed. This list of tasks is not exhaustive, and cannot cater for the myriad of local issues. These will need to be addressed by each individual team.

3.1. Reason for existence

"Why should there be an Incident Response Team?". This question, although obvious, is crucial. There may be many answers, all of which are equally valid. Ultimately, it is the answer to this question that will earn the respect and cooperation of the constituency.

Possible answers include:

- a local team that understands local issues;
- a team that operates in the same time zone as the constituency;
- separate security services from the network providers;
- to increase the security of the constituent's computer systems;
- to educate system administrators in their roles;
- to coordinate incident response at a central point;
- to scope the size of the security problem;
- to determine trends in attacks.

The lack of a clear reason for the existence of the IRT will ultimately result in a lack of support, both financially and administratively, which will lead to the demise of the team. If the constituency does not want the team, then its effectiveness will be minimised. This may lead to funding cuts, and eventual closure.

In Australia, it was determined that at the time it was the third largest nation in the world in terms of the number of registered Internet hosts. That fact, coupled with the comparatively low population makes Australia one of the highest Internet users per capita. It determined that Australia should take responsibility for its own security problems rather than relying on the limited resources of the United States. In addition, the timezone difference made cooperating with the United States difficult to perform effectively; a local team that understood local issues was required. This was an important issue in the justification for SERT.

3.2. Goals

Forming an Incident Response Team without a goal is like implementing computer security measures without a policy. If the goal of what needs to be achieved is unclear, then any efforts by the IRT will always be performed on an "ad-hoc" basis, without a clear picture in mind. This may cause precious team resources to be fruitlessly expended on ventures that yield limited results.

One thing that is consistent across all Incident Response Teams is that they do not have sufficient resources to do their job to the ability they would like. Their working day is a continual compromise of priorities. The lack of clearly defined goals makes priority decisions arbitrary at best, opening the possibility for error resulting in mistrust from the community.

Deciding goals generally follows immediately from answering the question about the reason for the IRT's existence. Once the goals are defined, they should be communicated to the community being served. Many misunderstandings between an IRT and its community have occurred because members of that community misunderstood the role and goals of the IRT. Clear and well defined communication of the goals of the IRT is essential if the community is to work with the IRT; not against it.

The expression of the goals may be made in the form of a *mission statement* to the constituency. The day to day operation of the IRT is then measured against the question, "Does this situation and action fit within the mission statement of the team?". A measure of success of the team's operations may be determined through some empirical measurement of how well these goals are being met.

Some examples of goals may include:

- raising the floor of Internet security;
- assisting sites in proactive security ventures;
- increasing the awareness of security incidents;
- determining the scope of the security problem;
- assisting the community in applying the best security practices available.

3.3. Constituency

When a team is formed, it must have a clearly defined scope of operations. The people it serves must know that they have an Incident Response Team, and the team must know who is, and who isn't, in the constituency. The scope of the constituency is usually defined by the community that is funding the IRT (either directly or indirectly). This may be based on the network provider, geographical considerations, or organisational considerations.

When decisions on the boundaries of the constituency are made, they should be communicated not only to those members that form the constituency, but also any members that do not fall within the boundaries. This might be done through other IRTs. Other IRTs also need to know where the boundaries of the constituency are defined so that they can direct appropriate queries to the correct team.

At times, it is possible that a site that is not within the defined constituency will request assistance. If that site falls under the defined constituency of another IRT, it is in the best interests of the IRTs and the site in question to have them contact the local IRT for assistance. If the site does not wish to do this, then it is polite to request permission to advise the local IRT that the incident will be dealt with internally, at the request of the site.

If permission is not given, then assistance should still be given to the site, with an attempt to resolve the issue of constituency as soon as possible. It is the experience of SERT that it almost never gets to this stage.

3.3.1. Defining the Constituency

Defining a constituency is not as trivial a task as it first seems. Constituencies may be defined by a number of constraints:

- geographical boundaries;
- network provider;

- organisational dependencies.

Existing IRTs are defined by a selection of all the above. In some cases, a site may be contained within the constituency of two or more IRTs. In many cases, there are sites that do not have an IRT. By default, the CERT Coordination Centre will always provide assistance to those sites, on an incident priority basis.

Some IRTs are defined by their network providers, which may or may not cover the entire country. If a site is in one country, but connected by the network provider of another country (which may or may not have an IRT), this has the potential for much confusion. The United States has a large number of IRTs covering a range of constituencies, with each team being established to meet the specific needs of their constituency.

3.3.2. Advertising

Having established the boundaries of the constituency, it is essential to advertise the existence of the newly formed IRT. This can be done in many ways:

- mailouts;
- electronic mail to network and site contacts;
- USEnet news;
- conferences;
- press releases.

It is important that the constituency learns about the existence of the IRT, and then establishes communication with that team to learn about their goals, mission, and policies. The mechanisms above are useful for advising the wider community of the existence of the team.

A mechanism for communicating the goals, mission, and policies could be through a "registration" procedure. By asking each site to register a 24-hour contact point with the IRT so that the site could be contacted after hours, a database of constituent sites can be established. At the same time, communication lines are opened with that site to provide information on the goals and policies of the IRT. If possible, establish a mechanism for rapidly contacting all members within the constituency (such as an electronic mailing list).

In Australia, this initially met with limited success. After a few incidents that resulted in some sites being uncontactable, the number of registered sites has risen steadily. This is a fundamental difference between the SERT and CERT operations. CERT by nature of its constituency can never establish a one-to-one relationship with all of its constituents; there are too many and they are too diverse. SERT has a well defined constituent list, and has worked to establish the ability to rapidly contact any constituent site on a 24 hour basis.

3.3.3. Identifying Trusted Contacts

If the IRT is to communicate security information with a site, then it needs to know whom that information is going to. If the constituency is relatively small and well defined, it is possible to establish a database of "registered site security contacts" in advance, rather than establishing a security contact for each incident as it occurs. This register should be independently verified.

Initial thoughts may be to solicit this contact information by asking each site to nominate their contact. This can be easily achieved using electronic mailing lists that already exist for the operation of the network. Any contact information received should be independently verified for correctness.

This method however, registers "contacts" for a site, not the "nominated security contact". These contacts may not be the appointed security personnel of the institution. The appointed security personnel may not be technically minded, but they might contain the authority to make decisions and contact the correct staff during an emergency. It is the responsibility of each constituent site to nominate the most appropriate site security contact.

Therefore, if the IRT is intending to form a register of trusted security contacts, it is strongly recommended that these contacts be determined by approaching the Chief Executive of the organisation, and asking that person to indicate who their appointed security contact is.

Another concern for some countries might be that the collection and storage of this information may contravene local laws (such as Data Protection and Privacy). This must be addressed on a case by case basis.

3.3.4. Information Releases

Incident investigation may require that certain items of information such as machine names and contacts be released to other parties. Rather than seek permission to release this information on a case by case basis, it may be easier to seek permission prior to any incident. Many sites do not mind their site name, contact information, or affected machine names communicated with any necessary parties to assist in the resolution of incidents. Seeking this permission in advance may reduce the time taken to resolve an incident, especially an international one where timezones become an issue, and any delay may be crucial.

3.3.5. Trusted Communications Paths

Once the community is identified, the IRT needs to be able to communicate with that community in a secure way. Many people think that this means that it should be impossible for an intruder to read the electronic mail that is issued from the IRT to the sites, and mail sent from the site to the IRT. This is only one aspect of this complex topic.

Electronic mail is by far the easiest form of communications for an IRT to deal with. Automated tools can be used to process the information, thus reducing the load on IRT staff. However, if a site has been compromised, then it may not be possible for them to send electronic mail (for example, if they have disconnected from the network). Other forms of communication will be required (such as phone, pager, and fax).

Working in the international community and with other IRTs sometimes requires the exchange of sensitive data. "Sensitive" may merely be a copy of a draft advisory that is still being verified for correctness. Early release of this information may result in further damage to the community.

Data encryption is another method of exchanging sensitive data securely. This relies on the end points of the communication being secure. If any of these end points is not secure, then the encrypted data should not be stored in plain text, and the encryption keys should be kept offline. The use of data encryption should be determined by the classification of the data.

The release of any public information from the IRT should be done in such a way that if any false information is released by a third party pretending to be the IRT, the fraudulent message will be detected. This may involve the use of digital signatures, certificates, or encryption.

The final topic in this area is the ability to access the secured systems within the IRT from outside the normal base of operations. This may occur for example if staff are travelling, or are operating after hours. These communications channels should also be secured against network sniffing.

3.4. Scope of Operation

What types of incidents will be handled by the team? What types of incidents will the team *not* handle? These questions must be answered and those answers communicated to the community. For instance, the types of incidents that may or may not be handled could include:

- intrusions;
- software vulnerabilities;
- requests for security information;
- requests to speak at conferences;
- requests to perform on-site training;
- requests to perform on-site security audits;
- requests to investigate suspected staff;
- viruses;
- international incidents;
- illegal activities such as software piracy;
- requests to undertake keystroke monitoring.

In addition, a decision must be made on what level of assistance will be provided. Will the team merely forward notification of security incidents onto the affected sites, or will they work completely with the site to determine the extent of the intrusion and help them to better secure their sites?

3.5. Identify Savings to the Community

Part of the justification for forming an Incident Response Team is to identify the savings to the community. This is typical of any risk analysis situation, where the costs of reducing the risk should not exceed the costs of the potential loss. Possible savings could include:

- real money costs in staff time handling incidents;
- costs of staff gathering and verifying security information;
- lost opportunity costs;
- loss of reputation (or gaining a reputation!);
- threat to "sensitive" data.

3.6. Scope of Expertise

Small teams in particular cannot have a complete set of skills required in today's complex and diverse array of computer hardware and software. There is no shame in admitting to the constituency that the team does not possess the necessary skills to tackle a certain problem. If the team finds itself in this situation, they could cultivate contacts within the community that do possess the required skills. Develop a level of trust with these contacts over time and use them from time to time when the team's skills are inadequate.

Be careful of always making use of the same people as they become less reluctant to help over time (due to other work commitments), and risking the wrath of their management. In general, people are willing to assist in true emergency situations, but are more reluctant to devote time to more mundane situations, or bolster the ranks of the IRT for free if the team is inadequately staffed.

3.7. Staff size and Makeup

About the only common attributes between existing Incident Response Teams are that they are under-funded, under-staffed, and overworked. Determining the appropriate number of staff to employ is a fine balance between the expected (and probably as yet, unknown) workload, and the budget constraints.

It is SERT's experience that one full-time technical person can comfortably handle one new incident per day, with 20 incidents that are still open and being investigated. Anything over this rate does not allow for any other involvement than incident response. This may have many negative aspects.

Besides the technical team, there must be management, administrative, and clerical support. These services may be contracted from other organisations, or people may be employed to fulfil these roles.

The biggest issue facing staffing levels is staff burnout. It is a problem that if staff are continually placed under stress by being on 24 hour callout, and working long hours on complex incidents, their mental and physical health may begin to suffer. It is highly recommended that to operate on a 24 hour callout basis, a *minimum* of three full time staff are required.

Staff should be rotated through the high stress positions, and when they are rostered off, they should be given the opportunity to pursue other less stressful activities such as tool and course development. However, staff should always be available to assist when the emergency load becomes excessive.

The incident *rate* is not a constant. There will be quiet times, and there will be busy times. The success of an IRT is usually measured in how they perform during the busy times, as this is when most members of the constituency are exposed to the IRT. There must be sufficient capacity in the team to effectively deal with large and complex incidents. Failure to do this will result in dissatisfaction from the constituency.

There are other duties for team members to perform when the incident load is light. Seminar preparation, tool development, policy writing, and most importantly, looking to the team's own security (which is often forgotten).

It is an unfortunate fact of life that incidents do not occur at a steady rate. What may initially be a quiet moment in the office can be shattered through a single electronic mail message. Incidents can, and do, occur in bursts. This is particularly true immediately after information on how to exploit vulnerabilities is made public. The posting of an exploitation script is usually a recipe for long hours within Incident Response Teams.

Possible solutions to this problem may involve the ability to recall staff at a moment's notice to assist with the higher than normal incident rate. This has negative implications of staff burnout. Another solution is to have trusted staff from other institutions on standby who could lend technical assistance in times of emergencies.

Not all incidents are created equal! This paper discusses incident load in terms of numbers of incidents. One incident may involve a single system and be dealt with in five minutes, whilst another incident may involve a large number of systems over many sites and continents, requiring an enormous amount of coordinating and analysis. Long running incidents are partially covered in the "open and investigating" incident category detailed above. This does not take into account the amount of effort required to resolve the incident, or the severity and priority of the incident.

3.8. Identify Technology Dealt With

Given that it is not possible for an IRT to have all the necessary experience to deal with every platform and system, a decision should be taken as to what technology will be dealt with, and what incidents may need to be referred to other groups or other IRTs for assistance. The choices could include:

- hardware platforms;
- operating systems and revisions;
- vendor packages;
- third party packages;

- public domain packages;
- viruses;
- worms;
- Trojan horses.

This information should be communicated to the other IRTs.

3.9. Identify Depth of Analysis

When investigating incidents and vulnerabilities, the depth of analysis may vary, depending on the size, experience, and spare capacity of the IRT. In general, the more time spent on analysis, the faster the problem will be resolved. However, some problems take an enormous amount of time to resolve, and may be beyond the experience of the team. A decision should be made as to what level of analysis will be applied to vulnerabilities and incidents.

Some IRTs merely act as a clearing house for security information, providing no assistance to the affected site to become more secure. Others will examine a vulnerability in depth, and determine not only a workaround and fix, but also an explanation as to why the vulnerability occurred, and examine other packages for similar problems. Most teams do not have this level of resource available.

Possible actions of the IRT when examining incidents or vulnerabilities may include:

- sending information on, but providing no further assistance;
- assisting sites to resolve the problem;
- assisting sites by examining their security and providing suggestions;
- examining source code to find the vulnerability;
- providing workarounds and example fixes to vulnerabilities;
- assisting vendors in patching vulnerabilities and testing solutions;
- detailed examination of vulnerabilities to determine why they occurred;
- examining other packages for similar vulnerabilities.

3.10. Budget

When submitting a budget for funding, the budget should contain a significant component for staff travel. This travel is used to attend conferences and workshops, meetings with constituent members, meetings with other IRTs, and meetings with the funding providers. Once the IRT starts up, it will be called upon to present papers at a variety of conferences and workshops, and this requires a large amount of travel.

3.11. Authority and Reporting

Each IRT has a management structure controlling their activities, and monitoring their progress. This management requires regular reports from time to time. The management also may exercise some level of authority over the IRT (such as demanding to know information like affected sites, or vulnerability details).

In addition, it is often misconstrued by the constituency that the IRT has some form of "authority" over them, and can direct other sites to "get their act together!". In general, this is not true. The IRT usually acts as an "advisory service", rather than an enforcing agency. Sites are more willing to report failures of security to someone that is in a position to help, rather than someone that is in a position to discipline.

The authority over the constituent members needs to be clearly defined, and communicated regularly to the constituency. Mistrust in the IRT will prevent security incidents being reported, resulting in incomplete information and an inability to assist sites with security. If the IRT has no authority over the constituency, then the constituency should be left in no doubt about this situation. The lower the authority by the IRT over the constituency, the more chance there is that the constituency will be reporting security incidents and seeking assistance.

In addition, any authority that may be exercised by the management over the IRT should be clearly communicated to the constituency. If the management may request access to any information, then the constituency should be aware of this, and accept it. Any reports that are generated for management should contain only the minimum of detail required for management to perform its duties. This level of reporting should also be communicated to the constituency. In general, the constituency is provided a summary of this reported information as a form of "statistics" or report of the progress of security within its constituency.

3.12. Policies

There is no point in advising the constituency that they need to have security policies if the IRT does not have one itself.

Policies are very important to establish early, so that all staff members take appropriate action in the majority of situations encountered. Policies should begin with a policy "framework" that shows how the various policies relate together.

Policy statements contain directives of a general nature, that may be implemented using the most appropriate techniques available. For example, the statement:

> "Data will be transmitted using DES in ECB mode."

is not a policy statement, as technology may change. A policy statement is better worded as;

> "Data will be transmitted encrypted using the best available technology at the time that ensures message content confidentiality.".

Many of the policies of an IRT will need to be communicated with the constituency so that they understand the role, goals, and intentions of the IRT. This helps to build trust in the IRT as the constituency fully understands what will happen with any information sent to the IRT, and what assistance they can expect from the IRT.

Some of the policies may not be considered to be public knowledge. In particular, policies relating to the internal workings of the IRT are probably best kept internal to the IRT as the constituents do not need to know this information. Determine which policies are public knowledge, communicate those policies to the constituency, and any other persons requesting them from time to time.

One of the major policies to develop is how to handle the release of information to various aspects of the community. These policies will need to deal with issues about what information is public, and who is authorised to communicate that information. Example situations include:

Press	The press has a job to do in getting the latest story that makes headlines and sells papers. As such, it is the experience of some people that they are not always accurately or completely reported, with some words being taken out of context. It is the policy of some of the IRTs that operational staff will not communicate with the press, but pass them to a nominated "press officer" that is briefed with the information that is public release only.
Incoming Calls	When a call is received into the IRT, the way it is handled may depend on the type of request. Determine which information is public, and only release that to unsolicited callers. For example, a caller may indicate they

	are from site X, and ask for an update on the status of incident Y. This caller may be the intruder attempting to determine what is known about their activities. If in doubt, call the person back using the contact information that has been registered for that site. If the caller is seeking public information, then there is no problem in just releasing that.
Sites	When communicating with sites, it is important to decide what they should be told in relation to their incident. For example, if other sites had already reported compromises as a result of a vulnerability, should this information be released to the caller? Should the current state of knowledge on the vulnerability be released?
Law Enforcement	In some countries, it is a legal requirement to advise law enforcement agencies of any knowledge of illegal activities. This must be resolved prior to commencing IRT operations. The detail of information passed to law enforcement should be determined and communicated back to the constituency.
Other IRTs	Resolving incidents will most likely involve the use of other IRTs, especially ones located in other countries. The level of information communicated with other IRTs should be determined and the constituency advised, either initially, or on a case by case basis. In general, it is almost impossible to resolve an incident without revealing the names of the source and target machines involved.

In general, it is important to identify what the IRT will do in terms of its operation. It is just as important to determine what the IRT won't do. For example, the IRT won't:

- investigate individuals;
- communicate vulnerability information without a fix;
- release site names and contacts without permission;
- advise law enforcement without permission;
- fix a constituent's security problems for them, but will offer advice.

3.13. Enforcement

Once policies are determined and enacted, there must be a mechanism to determine that they are being adhered to. Failure to adhere to stated policies can lead to a breach of trust in the IRT, finally resulting in termination of its services. It is vitally important that all staff members understand the policies and undertake to adhere to them.

Policies should not be overbearing. They should be implementable, acceptable, and testable. If the staff do not accept the policies, they will ultimately be forgotten. Some metric of compliance may need to be developed, to ensure that any steady relaxing of the adherence to policies does not go unnoticed.

3.14. Incident Response

Prior to commencement of operations, the IRT needs to decide how it will deal with incidents as they are reported. In many cases, the IRT will lack the necessary experience to know best how to deal with incidents initially. This experience will come with time. In the meantime, some communication with other IRTs to seek information on how to handle "fictitious" situations may provide some guidelines on where to start.

Make up an incident. Have someone communicate it to the IRT, and determine internally how this incident should be handled. Role playing and scenario analysis will assist the team in making rational decisions under pressure.

It is at this time that contact should be commenced with other IRTs. Trust will take some time to build with these teams, so it is important to be patient. Communicate the team's policies to the other IRTs, and let them provide some response on their experiences. Requesting information about current incidents and vulnerabilities will almost certainly be met with stony silence.

There are a number of other useful groups that can be contacted at this time, other than IRTs. These groups may be doing research and development into computer security tools and products, or may be experts in areas that the newly formed IRT does not have any experience in. Security research groups will be able to educate the IRT members on the latest advances in computer security. Contacts with local vendors should also be established so that rapid comments on vulnerabilities can be achieved.

One area that some teams decide not to develop expertise in is combating computer viruses. There are many vendors of anti-virus software, and a number of groups doing virus research. Contacts with these groups should be made, to allow for expert opinion when dealing with virus incidents.

Contacts with law enforcement should be established as many computer security incidents involve a breach of local laws. Whilst it may not be the role of the IRT to investigate criminal activity, they may be required to liaise with the law enforcement officers to provide expert assistance. Policies should be developed between these two groups as to how they will operate with each other.

3.15. Legal Issues

Local laws and conventions may affect how the IRT operates. These legal issues will require resolution prior to commencement of operations. In general, different countries will have different laws governing the various aspects outlined below. It is impossible to give a general guideline, and local legal counsel should be sought by the IRT.

If an IRT gives advice on security issues, and the site is further compromised, there may be a liability issue. In general, this is not the case, provided the IRT provides the best advice possible, based upon the knowledge that was available at the time. The IRT must undertake to obtain the most up-to-date advice possible at all times. Staff should be trained in security issues, and that training regularly updated from time to time. This issue may be reflected as a "duty of care" to the constituency.

Many countries now have enacted "freedom of information" (FOI) legislation that allows individuals to request access to varying amounts of data, particularly personal data, and have that data corrected if it is in error. If the laws allow individuals to request access to any data, then sensitive vulnerability data may be at risk. The law may require the appointment of an FOI officer.

Any information that is stored within the IRT should remain private, unless permission is granted by the constituent site to release it. There may be certain types of information that must be kept confidential according to certain laws. As well, the storing of information that identifies individuals may contravene local laws on the use of computer databases to store personal information.

If an IRT is to become involved in investigating computer security incidents, it may require monitoring network communications to determine the actions of intruders. In many countries, monitoring keystrokes may constitute a breach of privacy. For many companies, any data stored or transmitted internally is deemed to belong to the company for its official use, and therefore, is not private data. Any company data may be viewed by designated company officials, under policy guidelines.

Many intruders make use of the telephone system and modems for their initial connection into the computer networks. In many (most?) countries, monitoring a telephone line is illegal, and capturing the calling telephone number may also be a breach of privacy. In the cases where the telephone line is used, it is often illegal to tap the telephone line, but not illegal to monitor the connection once the data is within the organisation's boundaries on their networking equipment.

4. Equipment

Prior to commencing operations, the Incident Response team will require a number of items of equipment. The choice of equipment will vary, depending on the chosen constituency, the scope of analysis work, the types of incidents being investigated, the size of the team, the physical and geographical location, and approximately two thousand other related issues.

4.1. Phones

The IRT will require telephone access for contacting constituent sites, other IRTs, vendors, management, and other external contacts. For convenience, this phone access must be able to perform a number of basic functions. These might include:

- call pick from any other extension, while still maintaining the security that external personnel cannot pick up the calls;
- a central phone point that acts as the main contact point for the team. This point should be able to be answered by any other team member at their desk;
- the ability to switch calls to another party to answer calls when the team is unavailable (perhaps after hours, or during a team meeting);
- access to long distance and international direct dialling. The majority of the team's work will be communicating with people who are based some distance away;
- compatibility with existing infrastructure equipment.

The telecommunications equipment will require maintenance by other parties. The IRT may need to be mindful that the phone lines may not possess the desired security. Whilst there are a number of analogue speech scramblers on the market, many of these of not all that secure. The security of the telephone will vary from country to country, according to local laws, equipment, and telecommunication authorities.

4.2. Answering Services (24 hour contact)

An unfortunate part of the IRT's work is that the Internet is a 24-hour operation that spans the globe. To this end, the team must be able to be contacted on a 24 hour basis by constituents and other national and international IRTs. This may be done in several ways:

- "registering" an after hours contact with any person that needs to contact the team on a 24 hour basis. This is usually a team member's private phone number. This has obvious implications for privacy, and is not very satisfactory as the only other point of contact is when that team member is at home;
- the use of pagers. This may have negative aspects as an intruder may launch a "denial of service" attack by continually paging the team after hours. There may be a number of techniques to combat this threat, many of which can be implemented by the local PTT. Many alphanumeric pagers have a number of ways of being accessed, including a data dialup service. This opens the way for electronic mail to pager access. This can then be access controlled based upon the address of the sender.
- call forwarding of the central number to an answering service. This service could ask a few basic questions, and then issue pager or telephone calls to the necessary team members. This option has the highest security if a form of dial-back can be established. An intruder could make a nuisance of themselves by calling the answering service, and supplying random numbers for the call back. This is especially antisocial if it is done out of hours, with calls directed to innocent bystanders.

4.3. Fax

Some constituents in certain situations may not wish to send details to the IRT through electronic mail if there is a concern that the network or other central system that controls the mail has been compromised. The facsimile machine is another possibility for data transfer in this situation. The fax machine should be physically secure, and the security of the fax transmission will be as good as that for a normal phone conversation.

This adds an extra burden on the IRT as the fax must be associated with a particular incident when tracking that incident. Some suggestions on mechanisms to do this are:

- retype the fax into the incident tracking database. This has implications of typing errors;
- use a fax modem and software, and store incoming faxes in electronic form (for example, bit mapped Postscript);
- maintain a paper file of each incident. This will soon mount up to be unmanageable.

There is no one correct method. The desired method used to associate incident information that is not received in electronic format will vary, depending on the structure if the incident database, the type of information received, and the mechanism used to send that information.

4.4. Systems and Networks

One of the roles of an IRT may be to analyse incidents to determine trends and intelligence of future attacks. To do this, some form of incident analysis and database tools must be used. Since most of the information supplied to the IRT is already in machine readable form, a computer system is the obvious choice of tool. The team must be able to be reached via the Internet so that information can be sent to it, and other forms of information (such as Advisories) can be sent back to the constituency.

4.4.1. IP Address Range

Careful planning prior to the commencement of the team will save an amount of restructuring in the future. Since the IRT must be connected to the Internet, it must use a range of IP addresses. These may be "borrowed" from the organisation that provides the team's infrastructure (for example, being assigned a subnet for use). A better recommendation is to apply to the Network Information Centre for a separate IP address range. This has no immediate benefits, but will have significant benefits should the team be required to relocate its base of operations to some other administrative or geographical location.

4.4.2. Domain Name

The team will be required to register a domain name with the Network Information Centre and the network providers. It is important to place the team under the correct higher level domain from the outset. Both the CERT Coordination Centre and SERT originally started under one domain, and have subsequently moved to a more appropriate domain. This has implications of having to maintain backward compatibility with old names for many years.

Originally, the SERT team was placed under the `.edu.au` domain (`sert.edu.au`). This was mainly due to the way that this team was formed. It was quickly pointed out that SERT's constituency covered more than educational institutions. A number of research, government, and commercial organisations were contained within the constituency definition. Ultimately, this caused confusion and mistrust (some constituents thought that SERT would only operate for educational institutions).

The migration to AUSCERT has allowed the new team to move under its correct parent domain as `auscert.org.au`. Since AUSCERT is a non-profit organisation without direct association with any particular form of organisation, and since it may be contracted by more than one network provider, the logical conclusion for AUSCERT was that it was an "organisation". The CERT Coordination Centre is also now addressed as `cert.org`.

Careful choice of a domain in the initial stages will remove the drama of changing names at a future point in time, requiring backwards compatibility.

4.4.3. Subnetting

It is a good idea to be allocated a complete subnet from a larger network address space or a complete network address space, rather than be allocated a range of addresses within another organisation's network. This allows the possibility of subnetting the address space further to form a number of different networks. The separate networks can then be protected using different security policies.

Example subnetworks may include:

public	this network contains public access machines such as anonymous ftp, gopher, and world wide web servers. Information stored on this subnet is deemed to be public release;
test	it may be desirable to have a testing subnet. This network may or may not be secured, and any testing on this network will minimise the impact on production machines. The nature of testing vulnerabilities often leaves a machine open to attack. It would be desirable to make this network secure from outside connections (although, other IRTs may require access when cooperating on a vulnerability analysis). Should the test machines be compromised, they should not have access to the secure network, and they should not contain any sensitive information;
secure	the IRT will require a network that is secure against intrusion. This network will hold sensitive information such as ongoing investigations, site contacts, site names, and vulnerability information;
highly secure	the highly secure network may be used to store the most sensitive of information. It should not allow *any* connections into it, but may allow connections out of it. These outgoing connections must be carefully audited to prevent the accidental "down-classification" of data, by moving it to another network.

There may be other requirements for separate networks. Splitting the network into four subnets should provide reasonable flexibility for future plans. For example, a complete Class C address space may be split into four separate subnets, allowing 62 hosts on each (not including the network and broadcast addresses).

4.4.4. Test Equipment

If the team is to be involved in vulnerability analysis (proactive operations), then a range of test equipment will be required. This test equipment should be chosen to best serve the needs of the constituency. There will generally be insufficient funds to get one of every platform running all software. It is under these circumstances that other IRTs will be able to contribute test platforms and expertise.

The test equipment should not contain any sensitive data, and should not be required in the day to day operations of the team. It is possible that testing security vulnerabilities will reduce the security of this system, or even cause it to fail.

4.4.5. Routers/Firewalls

Once the availability of a security team is announced, it is likely to become a target for all sorts of reasons. As with plumbers who always have leaky taps at home, carpenters whose kitchens require repair, builders whose doors need adjusting, it is possible that in the rush to assist other sites, the IRT fails to attend to its own security.

Nothing makes an intruder look better than to break into a computer run by an Incident Response Team. Nothing destroys the constituency's trust faster than if the IRT's machines are compromised.

Sensible security starts at home. There must be dedicated hardware and software designed to increase the security of the internal systems. The router and its filters must be under the administrative control of the IRT (or appointed staff), and should be reviewed regularly for effectiveness.

Solutions may involve managed bridges, routers, or software firewalls. The decision is based on expertise to establish an effective filtering mechanism, and budget constraints.

4.4.6. Non-replayable Authentication

Incident Response Team staff will be required to operate from outside the secure environment from time to time. This may be as a result of visiting another site to assist them, attending a conference or workshop, or operating after hours. If access to the secured network is to be granted to team members, then they must be made aware of the possibility of trojan horses and network sniffers operating in the network.

Some form of non-replayable authentication sequence is required. This may take the form of one-time password generators, software systems such as S/Key, or some other locally developed mechanisms. These systems should be secure, such that no matter how many password "tokens" are captured, the next password in the series cannot be guessed or determined.

4.5. Shredder

This piece of equipment is generally quite cheap, but may be necessary. There has been an amount of literature that discusses intruders "trashing": searching through waste paper bins for snippets of information.

An Incident Response Team will be given many pieces of sensitive information. This information may not necessarily be how to break into a computer, but it might mention a sensitive site by name. The negative press generated by such a leak of information could spell the end for an IRT. Confidentiality also means destroying information when it is no longer required; hard copy included.

4.6. Safe

Much attention is given to the logical security of the IRT, but what about the physical security? The budget should contain provisions for a fire-proof safe. This safe could store for example:

- encryption keys. If all data on the disks is stored encrypted, then theft of the disks will not reveal any sensitive information;
- backup media. This prevents a thief from stealing the backup media and analysing the information stored on it. It also provides a mechanism for recovering quickly from a disaster such as a fire. Such disasters may be the result of a malicious act directed at the IRT as a result of their activities.

4.7. Backup Media

It is important to establish and maintain a system backup strategy. This requirement is not unique to IRTs, but should be practiced by any organisation that cannot afford extended down time or loss of data. Backup media could consist of tapes, shadowed disks, or other removable media.

Provisions should be made for the secure storage of this media. If all backup media is stored on-site, then a disaster may result in the total loss of all information. Some form of secure off-site storage is required. This could include:

- a fire-proof safe on the premises as discussed;
- a trusted organisation that provides such a service;
- encrypting all data prior to backup.

Note that any data sent off-site for storage should be afforded the same level of security as the on-line data. It should be protected from unauthorised disclosure, modification, or loss.

4.8. Information Security

The information stored on the secure systems requires extreme protection from unauthorised disclosure and modification. Information may be in transit on a network or phone link, stored on a disk or tape, stored on paper, or in the memories of IRT staff.

A number of requirements may be placed upon the IRT for the security of that data. The requirements may stem from classification issues, legal issues, or a need for privacy of affected sites (policies of the IRT). Regardless of the specialist requirements for security, a number of common elements of security are required.

The first is physical security. The premises that house the IRT must be physically secured from intrusion by unauthorised personnel. This could include mechanisms like physical and electronic locks, intrusion detectors and alarms, security guards, and security badges. The premises must be able to be accessed by IRT staff on a 24 hour, 7 day a week basis.

As indicated previously, some form of network router will be required to connect in to the Internet. In addition, traffic filtering will be essential to prevent sensitive IRT data from being sniffed on other parts of the network. This filtering should ideally be performed The filter should only allow connections into the secured subnet from a small subset of trusted hosts.

Provision of a filtering mechanism to prevent unauthorised connections does not diminish the responsibility of the IRT for their own host security. Careful attention must be paid to the security of hosts on the secure subnet in case access is gained to that network by an intruder. This includes items such as integrity checks, log file and system audits, password use, security enhancement and assessment tools, and data encryption. The secured systems must be able to exist in an environment where any connections may be made from the Internet to them. It is important to ensure that resources are devoted to this task, and the task of effective system administration. It is very easy to become complacent about security when dealing with it daily.

It may be advantageous to "classify" the data stored within the IRT according to some set criteria, and then define how each category of data should be handled. Possible examples are:

public	may be transmitted in plain text, and released to any person;
private	information that is private between a constituent and the IRT. This may include incident data, site contacts, and equipment lists. This data will only

	be sent to registered security contacts for that organisation. No indication of the presence of this data will be made to other people. Data may be transmitted in plain text if acceptable to the constituent. Encryption may be used otherwise.
sensitive	this may include information on how to exploit an old vulnerability. Whilst some of this information is in the public arena, it may not be widely known. Release of this information will result in increased incident loads. This data may be shared with trusted constituents and other IRTs on a needs basis. It must be transmitted encrypted.
highly sensitive	this may include information about sensitive constituent sites (such as the military), or exploitation information about current vulnerabilities that have no solution yet. This data must be stored and transmitted encrypted. It may be shared with other IRTs on a needs basis.
classified	this may include data that is otherwise classified by other organisations such as the military or law enforcement. Storage and handling of this data will only be performed by security cleared personnel according to the requirements of local laws.

Mechanisms must be established for communicating with other organisations (such as constituents, law enforcement, or other IRTs) using data encryption. The tools and procedures will vary depending on local conditions at each end of the communication. As a baseline, DES encrypted, uuencoded text is acceptable to most places in the world. The issue of key management should be addressed and resolved satisfactorily for encryption to work successfully.

There may be requirements for speech encryption or scramblers with some constituents or law enforcement. These issues are local to the community and should be considered as part of the overall data security policy.

4.8.1. Data Origin Authentication and Integrity

An issue that is developing on the network is the ability of anyone to forge news and mail articles. This may throw doubt on the integrity of information released from the IRT. In general, any constituent can verify the origin and content of the information by contacting the IRT through some other mechanism (such as the telephone). This may become unworkable if the number of constituents is large. Mechanisms should be investigated which allows each constituent to verify the content and origin of the information released from the IRT. These techniques may include Certificates and Digital Signatures, PGP, and PEM. Current ITAR regulations on the exportability of encryption software make the interoperability of the United States and Canada with the rest of the world extremely difficult.

4.8.2. Trusted Staff

A major issue for all Incident Response Teams is the selection of staff. It may be felt that one of the most important attributes of a staff member is their experience in computer security. However, ultimately the success of the team could be undermined if that team member exhibits behaviours that undermined the trust of the constituency in the team.

The author's personal opinions are that the following attributes are required in any team members, and are placed in priority ordering:

- integrity. A lack of staff integrity will result in the ultimate demise of the IRT through mistrust. This integrity may require a staff member to abide by the policies of the team, even if they do not agree with them;

- operating system administration experience. The team members must have significant experience in managing computer systems, and preferably ones that are in a large network. This is the type of person that the team is trying to assist, and experience in this area will help the team understand the problems faced by the constituency;
- programming experience. If the team member cannot program quickly and effectively, and cannot read source code quickly and gain an understanding of a program, then their ability to analyse new security incidents will be limited. In many circumstances, the analysis must be done quickly and effectively. There is little room for learning, and little room for error;
- communication skills. The team member will be required to present talks and write papers in their role of educating the constituency. If this cannot be achieved effectively, then the security incidents will continue to occur. When dealing with an incident, sometimes all that is required is a "friendly ear" and some offer of advice. Effective verbal communication skills are essential when assisting sites that have experienced a compromise;
- security experience. Knowledge of computer security is desirable, but can ultimately be learned "on the job". If the team member is very experienced in this area, but lacks skills in the others, then their usefulness as a team member over time will rapidly diminish.

These days of "equal opportunity", "freedom of information", and accountability for all actions makes staff selection a complex topic. Staff selection criteria, job advertisements, and interview and screening techniques must be carefully addressed before employing new staff.

Ultimately, the staff members are accountable for their actions. The IRT may make a requirement of their staff to sign some form of "non-disclosure" agreement that binds the staff member to their responsibilities; even after they have left the team. These responsibilities may include maintaining the confidentiality of vulnerability information, site contacts, incident details, and information relating to the IRT itself. Whilst disciplinary action (such as dismissal) could be taken against staff while they are still employed for breaching this agreement, generally the only possible action that can be taken after the staff member has left the team is some form of legal action.

Depending on the constituency of the IRT and the type of information that the IRT will be required to deal with, team members may be required to undergo some form of civil or military security clearance. This opens up a range of problems should the IRT wish to employ foreign nationals, or people with criminal histories. Where possible, the decision to obtain a security clearance should be left to the staff member, with no requirement being placed upon them to do so by the team. This may split the team in two however; equipment and staff that are security cleared, and those that are not. This must be handled on an individual basis.

If the team is to store classified data (such as court evidence, military data, incident data) then there may be issues dictated by law or other convention on the storage and access to this data. Such issues may be the use of certain types of safes and locks, right through to the choice of the colour of the folders the data is stored in. These issues should be addressed on an individual basis.

Other staff issues involve access to the premises by non-IRT staff such as cleaners, security personnel, network and system administrators, electricians, window cleaners, pest control, management, and the general public. Whilst it is usual to deny access to the public, it may be a requirement of the physical location of the team to allow access to a range of other personnel. This may be done during hours or after hours. Good practices by team members (such as locking away any sensitive data each night or when the office is unattended) will reduce the risk of allowing access to other personnel. If the team is holding classified data, there may be a requirement to obtain security clearances for other personnel, or at least have a team member present at all times that other personnel are on the premises.

5. Commencing Operations

At some point, the newly formed team will transition from the development stage to the operational stage. Ideally, when this transition arrives, there should be very little for the IRT to do except commence operations. More than likely however, the IRT will still be preparing their database, acquiring staff and equipment, training staff, and establishing telecommunications. Despite the obvious chaos behind the scenes, the IRT must present a professional and educated front to its constituents from the outset.

A set time must be determined to commence IRT operations. When this time arrives, an announcement should be made in the form of a press release, and that announcement should be transmitted as widely as possible. If the IRT then waits at this point for the calls to start coming in, then they have immediately failed.

The IRT needs to identify its constituency and then go out and "sell" its services. Make the constituency aware of what the IRT intends to do, and why it is doing it. Educate the constituency on the policies, goals, and mission of the IRT. Tell them why the IRT was formed, and why it is important that the IRT coordinate incidents for the community.

It is at this point that the IRT should solicit trusted contact information from its constituents. This is discussed later.

An IRT cannot exist in isolation from the rest of the world. The IRT should at this point establish communications with other existing IRTs. A form of trusted communications should be established between the IRT and its constituents, and the IRT and other IRTs.

The IRT should then consider applying to become a member of FIRST - the Forum of Incident Response and Security Teams. This application will take some time to achieve, and the IRT should communicate with other IRTs about the benefits of becoming a member of FIRST, and the procedures required to join.

6. Operations - Learning

Commencing the operation of an Incident Response Team is a major task. Initially, the staff that are selected to operate the team may have little experience in incident response, or even security issues in general. A steep learning curve is experienced, with obvious drawbacks along the path. The constituency will have an expectation that they are dealing with "experts" in the field. There are few people in existence who can claim that they know about every operating system, every hardware platform, every piece of third party software, every security device, every book published, and all the security implications that go along with this.

It is important to recognise from the outset that the team members cannot be expected to know everything. The constituency should be advised when some particular query or situation falls outside the expertise of the IRT. In this way, there are no false expectations formed by the constituency, which ultimately leads to disappointment and mistrust in the IRT.

The solution to the problem is to identify trusted members of the constituency who are experts in one field or another. Establish and foster those contacts, and use them when required. Be careful not to use a small set of contacts all the time, or that resource may be withdrawn from the IRT. These trusted contacts need not be given all the information relating to a query (for example, there is no need to disclose the originating site name). They should be supplied with enough information to assist with your query, and should be prepared to accept that they will not always be told the complete story. The amount of information released will depend on the type and sensitivity of the query, and the level of trust with the contact.

6.1. Report incidents

Get the sites within the constituency to report all security incidents, and then analyse these incidents. Initially, this is a large task, and requires many hours. As the experience of the team increases, many incidents fall into the same generic "class" of attack or vulnerability, and analysis will proceed more quickly.

Reporting all security incidents, no matter how minor, allows a central reporting facility such as the IRT to determine the "bigger picture" of security within the constituency. For example, a site detects four `tftp` probes that were unsuccessful. In isolation, it is rather an innocuous and poor attempt to exploit a vulnerability. However, if more than 50% of the constituency reports four `tftp` probes from the same site, then this represents a coordinated and determined attack. It may be that upon further analysis of this incident, other sites are identified in which this very same attack was successful without detection.

Determining the scope of the security incidents within the community assists not only the IRT in justifying staff and equipment funding, but assists the constituency in justifying their security staffing levels. In general, only a small number of security incidents are ever reported. Many of the poorer attempts are repelled by a site, and ignored. These incidents may have only been repelled due to the ability of experienced staff to configure the systems appropriately. Therefore, the existence of those staff members can be justified since the attack was not successful.

6.2. Contact other IRTs

Other IRTs in the world may be able to assist with the learning curve. In general, each IRT will have a collection of security related literature and tools that they are more than willing to disclose to new teams.

Additional information that may be of use is an idea of the "state of the world" in security. How many incidents per 1,000 hosts? What platforms, operating systems, versions are experiencing the most incidents? Why? Information like this may help to channel learning efforts into more productive areas initially.

6.3. Documents and Tools

The Internet has a wealth of documents and tools relating to computer security. Where does one start? A number of sources could be explored initially:

- `archie` searching;
- USEnet news groups like `alt.security`, `comp.security.misc`, and `comp.security.unix`;
- mailing lists such as `bugtraq` and `firewalls;`
- information supplied by other IRTs and constituent members;
- anonymous ftp areas like `ftp.cert.org` and `ftp.auscert.org.au`;
- World Wide Web - start at `www.first.org` and see where that leads!

6.4. Library of Reference Material

There are many books written on the subjects of computer, network, and data security. Over time, the team should build up a library of reference material that can be referred to during incident and vulnerability investigation.

Some books are better than others. Many books are reviewed and these reviews are published either in USEnet news, or in popular computing magazines and journals.

6.5. Journal Subscriptions

There are also a number of journals that are devoted to security, and networking issues in general. A subscription for regular issues will ensure that no important articles are missed. It may be that many articles are of no immediate benefit to the team, but the background knowledge gained will assist the team members in keeping up to date with current trends and technology in the security world.

6.6. Staff Training

A number of courses in security are run as a commercial enterprise. These may or may not address the issues that affect an Incident Response Team. Many of these courses are aimed at the commercial and government infrastructures, covering topics such as security policies, viruses, and data encryption.

If the staff lack the necessary system administration skills, then they should attend a course on this topic. Solid system administration skills, coupled with the latest knowledge in security, is a good recipe for protecting computer systems. If the system administration skills are poor, then the security knowledge is wasted.

Ultimately, it is the experience of many IRTs that they are presenting the courses, rather than attending them!

6.7. Visits to Existing Response Teams

A common belief in many new teams is that all the other IRTs will simply hand over all of their information. This is simply not true, due to many considerations. The major consideration is the policies under which this information was obtained in the first place. There is a moral obligation for IRTs to protect their data, and they cannot release it to any new team that announces itself.

In order to build up trust within the IRT community, face to face visits are required. Without knowing the members of a team on a personal basis and understanding their policies and level of integrity, it is almost impossible to exchange any information of a sensitive nature. Trust takes time to build, and can be destroyed in seconds. If it is the intention of the new IRT to contribute to the cooperation of IRTs around the globe, some effort is required to establish credibility over time.

7. Operations - Reactive

The highest priority task for an IRT is to respond to incidents as they occur. This may involve working with the affected site to determine the cause of the incident and help them become secure again, or it may involve finding a solution to a vulnerability that is being actively exploited to compromise many sites. Reactive response is always done on a priority basis; where are the team's resources most effectively utilised?

7.1. Routine

The day to day operations of an IRT are very difficult to define as much of their work is event driven. A single telephone call or electronic mail message can change the structure of an entire week! There should always be a number of background activities occurring and these should be scheduled for attention from time to time by all staff members. This may include reading journals, papers, and books, or auditing the security of the systems and networks within the team.

7.2. Operations Manual

Many of the team's operations should be standardised and documented, so that team members can make informed and appropriate decisions in the majority of cases. All standard operations should be documented, and reviewed from time to time. The operations manual needs to address at least the following issues:

- handling vulnerabilities;
- creating advisories;
- handling difficult contacts;
- handling unauthenticated callers;
- information disclosure;
- coordinating with other IRTs;
- systems management;
- backup strategy;
- disaster recovery;
- off-site operations.

7.3. Administrative

There are many day to day administrative issues that must be addressed. These may include period reporting to management and constituency about incident levels and intelligence on attacks, plus some form of measurement of progress and success. It is not possible to base a measure of success on incident levels; overall, success can be measured by the severity and number of incidents, and the type of incidents being observed.

IRT staff must not be left on incident response for long periods or they will become detached from the changing world of security and system administration. Staff should be given the opportunity from time to time to pursue other interests to allow for a break from incident response, as well as providing a growth path for their professional development.

Classified data and access to it by security cleared staff may require a period activity such as auditing the procedures used to access data, submitting period report forms to an agency, or changing locks and keys. These should be adhered to as instructed by local law and convention.

If staff have been required to sign a non-disclosure agreement or acceptable use policy document, then they should be reminded of their responsibilities from time to time. It is the role of management to ensure that all staff abide by these agreements. If a staff member leaves the team, procedures should be in place to terminate their system access, change passwords, change encryption keys, retrieve physical and logical access devices, and to debrief staff on their continued obligation for confidentiality and privacy of the information.

Policies should be put into place to increase the physical security of the premises. This could include such items as storing encryption keys in a safe, locking all documents away at night or when the office is unattended, locking computer screens to prevent access to the systems from accounts left logged in, and locking backup media away.

Ideally, encryption keys, physical keys, passwords, safe access, and so on should be granted on a "need to know" basis. This may prevent non-security cleared staff being exposed to classified material, and will increase the auditability of all staff actions. Passwords should be changed regularly, and not reused.

If the team is to be available on a 24 hour basis, then there must be some form of roster to rotate the responsibility of answering emergency calls. If a call is received and responded to after hours, a decision must be made whether there is assistance that can be provided immediately, or whether the caller can wait and contact the IRT during office hours. Often, if a caller has placed an emergency call to the IRT out of business hours, they perceive it to be an emergency. In the majority of these calls, the caller is seeking some assistance to give them confidence in their immediate actions. The caller should be reassured as to what course of action is best to take, and then left to perform those actions. They can then be followed up during business hours with more detailed information about a solution to their problem.

It is rare that an emergency call cannot be dealt with reasonably quickly, and deferred to the next day. There is a small subset of emergency calls that require detailed analysis of a problem, development of a solution, and communication of that solution back to the constituency after hours. In these cases, there will always be another IRT somewhere on the globe that has team members awake and alert, and able to provide assistance.

Many smaller teams may find that their management structure is actually larger than the team itself. When there are several managers involved, it is important to have a clear understanding of the chain of command, which manager is responsible for which decisions, and what the correct reporting structure is if a manager is unavailable. The concept of "too many chiefs" may cause the team to become fragmented and confused. Just as the team must present a unified front to its constituency, so must the management present a unified front to the IRT.

The management structure should be based on the size of the team, and the structure of the constituency it serves. There may be cause to establish a management structure that is based upon the types of service provided to the constituency. This may vary depending on the service, so that the management needs of that service are best met.

7.4. Contacts

If the IRT has a reasonably small and well defined constituency, it is highly advantageous to build a database of contacts within each of the constituent sites. These contacts should be verified independently as the database is formed so that trust can be placed in the integrity of that information. When soliciting this contact information, it may be best to approach the CEO or head of the constituents and ask them to nominate their appointed security contact. In this way, the IRT will be dealing with contacts who have the knowledge and authority to act on situations.

Useful contact information includes:

- names (get more than one in case the first one is unavailable);
- addresses;
- organisations;
- main switch telephone numbers;
- contact office telephone numbers;
- electronic mail addresses;
- IP address ranges;
- list of hardware and software in use;
- after hours contact points such as home phones, pagers, mobiles.

For after hours contact information, it is best to make this optional as some sites may choose not to include this information. If the IRT does have contact information, then they are able to contact sites after hours and warn them of potential vulnerabilities and threats rather than waiting until the next working day.

If the IRT does not have registered contact information for a site, then they can use public information to track down a contact. This may include NIC databases, telephone books, electronic mail lists, or word of mouth.

It is also advantageous to seek the necessary permissions for information disclosure prior to any incidents occurring. This can be done as part of the registration process. When investigating incidents, it may be difficult not to reveal the affected system's name to the other party. As well, it may be necessary to work through another IRT to achieve resolution. In some countries, it may be legally binding that if an individual has knowledge of a crime being committed, they must advise the law enforcement authorities. It is advantageous to seek prior approval from constituents to pass only the necessary data on to third parties when assisting with the investigation of incidents. These third parties include other sites, other IRTs, or law enforcement. The information may include hostnames and connection records, site contact information (not after hours), and site names.

7.5. Unsolicited/Unauthenticated Calls

If the IRT receives a call from a person claiming to belong to a site, requesting information about a particular incident, what will the appropriate action to take be? The caller may be the intruder, seeking information about how much is known about the incident.

It is important to have clearly defined policies and procedures for dealing with unauthenticated calls and messages. Messages are easier to handle as they need not be responded to immediately, and further confirmation may be sought through some other means.

Phone conversations are more difficult to defer. It is possible to authenticate the user by either knowing them personally, asking them for particular information that only they could know (perhaps part of a previous phone conversation), or the call could be terminated, and then reestablished by calling the registered contact point for that person. Establishing the caller's identity is important if information is not to be leaked to unauthorised people.

Another mechanism for dealing with unauthenticated calls is to determine in advance exactly what information is deemed to be public. In that case, all staff will be able to decide if the caller's request can be satisfied by releasing public information (such as a request for an Advisory, or information about the team).

Information that is not public should not be released to any unauthenticated person. The presence of this information should not be revealed, nor any indication that the team will likely be given such information.

7.6. Point

It is desirable to establish a roster of staff that act as the focus for all information flow into and out of the IRT. This prevents confusion resulting from misunderstandings as to which team member is performing which function. All information into the team must be logged and replied to. All calls into the team should be answered by a single person. This way, a professional and unified front can be shown to the constituency. Failure to establish this regime may result in messages being lost and going unanswered, calls being taken and not logged, updated information on incidents not being passed back to the constituents, or messages being answered differently by two different team members.

Some IRTs call this position "point" or "point duty" (adapted from the original military role of being the first person in a patrol). Staff who are not rostered on point duty should be given the ability to stay away from the high interrupt load that this task entails, and be left to concentrate on other issues such as education and tool development. This allows staff to develop professionally, and gives them a rest from the pressure of dealing with security incidents.

The staff not on point duty may be called upon in times of emergency to assist with problems if required. The tasks of point duty may include:

- answering and logging all incoming telephone calls and faxes;
- answering and logging all incoming electronic mail;
- reviewing all outstanding incidents for action;
- updating the incident database with new information;
- administrative duties such as backups;

When answering calls or electronic mail, a number of items of information should be sought immediately. This will reduce the overhead of dealing with the incident at later times. This information is:

- primary contact details: name, telephone, fax, electronic mail;
- secondary contacts if the primary is not available;
- affected machine names and addresses;
- how the incident was discovered;
- the source of the incident if known;
- action taken to resolve the incident, including other sites contacted;
- action the site wishes the IRT to perform.

Messages should be responded to in real time wherever possible. This may mean sending a short reply thanking the person for the message, and advising that it will be looked at within a set period of time. This gives that person confidence that the IRT will address the issue, and indicates how long they should wait before taking further action.

Communicating expectations to the constituency is extremely important. They must understand exactly what will be done with information, when, and by whom. This removes any misunderstandings, and establishes more trust in the professionalism of the IRT. The expectations that should be made clear are:

- what will the IRT do with this information;
- when will it be done;
- will it be forgotten;
- who will followup the information;
- will the information be passed to other people;
- what should the reporting site do at this point.

Whenever any action is taken on an incident, the reporting site and any other affected sites should be kept informed of any progress or changes in the incident status. In this way, the sites are reassured that their needs are being addressed by the IRT, and that reporting information to the IRT has received attention. Sites are more likely to continue reporting information in this case, which is important to the IRT's operations.

7.7. Incident Numbers and Database

When incidents are reported, they should be logged into an incident database. This database should be used to collate all information relating to the incident in one place for all IRT staff to view and act upon. This information should include all electronic mail, telephone conversations, facsimile transmissions, and IRT staff notes. All staff must be able to update the incident database as new information is received and the status of an incident changes. They must be able to determine the status of an incident with a high degree of confidence, and each member must be able to arrive at the same decision as to the required action to be taken as a result of stored incident information. The integrity of the database should be protected from multiple, simultaneous updates.

The database should be able to be searched based upon a number of criteria including site, dates, vulnerable software, methods of intrusion, geographical location, IP address, incident status, duration of incident, and other criteria required by the IRT.

To distinguish individual incidents uniquely, some form of incident identification is required. Protecting the privacy of affected sites is paramount, so some form of numbering is appropriate. These incident numbers should not contain any information about the affected sites, nor the severity of the incident. Some teams use random numbers. The SERT team uses numbers calculated from the date and time the incident was first logged. The incident number is a 10 digit number that consists of YYMMDDHHMM. The structure of these numbers allows for automated analysis and summary reporting. It may however reveal some information about the incident if it is known that a site experienced an incident at a particular time.

The database should be used to generate statistics on the incidents. These statistics should be periodically reported back to the constituency and the IRT management structure. Statistics could include:

- number of incidents, open and closed;
- number of calls for help;
- number of queries received;
- number of phone calls received;
- number of electronic mail messages processed;
- time an incident requires before closure;
- breakdown of severity of incidents;
- analysis of incident trends;
- other statistics as required by the constituency or management.

7.8. Hot Lists and Refer Again

As incidents are logged, certain actions are required to be performed such as contacting affected sites, contacting vendors, or seeking further information from the affected site. This may take some time to retrieve. A useful tool is the ability to place an incident on "hold" for a set period, and have the incident brought to the attention of IRT staff (the person on point duty) after that period expires. In this way, incidents are not forgotten, and constituent sites are followed up to ensure that they also do not forget to perform any actions requested of them.

8. Operations - Proactive

What does taking a "proactive" role mean? An Incident Response Team may find that it does not have sufficient resources to deal with any more activity than reacting to incidents as they occur. They spend all of their time communicating with affected sites, assisting them to recover, and collating the data.

What is required is a way of analysing the incidents, identifying patterns and trends, determining intelligence on the likely next wave of attacks, and working to prevent these attacks before they reach large proportions. Ideally, it is good if a security vulnerability can be identified and fixed prior to any exploitation of it. This has an enormously positive effect of reducing the incident load, and increasing the security of the constituency, and the Internet as a whole.

Vulnerability analysis may be a difficult task, and should not be undertaken lightly. It requires the skill to be able to read source code very quickly, gain an understanding of the problem, and the appreciate the complicated subtleties of possible solutions. Many solutions are less than optimal due to the large number of platforms and operating systems that they must be compatible with. Experience in multiple platforms is a major bonus when examining vulnerabilities that affect more than one type of system.

Some vulnerabilities occur in vendor controlled modules. Many of these modules are shipped to customers (IRTs included) in binary form only. This may prevent the IRT from examining and testing the vulnerability. In this case, effective communications must be established with the correct team within the vendor to work towards a solution. Vendors have a responsibility to test their solutions, and distribute them to *all* of their customers (not just the ones connected to the Internet). Documentation must be produced, media copied and distributed, and customer support centres advised and trained. This takes time!

Some vendors are now releasing their security patches to the Internet as soon as they become available. This has a positive effect in that sites that are connected to the Internet may fetch the patches quickly. Many vendors now make their security related patches available to anyone for free (without a software maintenance contract).

If the vulnerability affects more than one vendor (since many of the vendor's operating systems have been taken from the same original source code tree), then the problem of coordinating fixes becomes extremely complex. Releasing information about a vulnerability and solution for only a subset of vendors may reveal information about the vulnerability existing for other vendors that do not have a solution yet. Withholding information about the vulnerability until all vendors have a solution allows more time for the vulnerability to be exploited.

Ultimately, it may be in the vendors best interests to make a small subset of their source code available to (at least) the IRTs. This source code should include the modules that run privileged on their operating systems. This assists the vendors as they now have access to a large number of security specialists, all working to benefit the vendor! The intruders often already have source code that they are examining.

8.1. Proactive Roles to Prevent Incidents

Having decided to take a proactive role with vulnerabilities, the IRT must decide on what activities they will expend their resources on. Vulnerabilities are usually discovered by a constituent site, either when analysing an incident, or uncovering it by accident.

If the discovery of the vulnerability was due to the analysis of an incident, then the intruders must already know the information, and are exploiting it. This reduces the amount of time that the IRTs and vendors can work on the problem to determine an optimal solution. If the discovery was by accident, then so long as the IRT can rely on the integrity of the constituent, then there may be sufficient time to analyse the problem and ensure a total solution. However, any vulnerability that can be discovered by accident at one site, can easily be discovered by another site.

Solutions created under pressure have an extreme potential to contain other related vulnerabilities, or cause some other functionality to fail. If a section of software is complex enough that the original programmer made a mistake, then it is still complex enough that a code maintainer will also make a mistake.

When a vulnerability is reported, it is important to determine a number of basic facts:

- is this vulnerability easy to reproduce;
- does this vulnerability affect different versions of this software;
- what previous level of access is required to exploit this vulnerability;
- does this vulnerability grant privileged access;
- does this vulnerability affect other vendor versions;
- is this vulnerability being actively exploited;
- can this (or other) vulnerabilities be further exploited to gain privileged access;
- how many systems within the constituency and the Internet are affected.

This gives some form of metric as to the severity of this vulnerability, and the resources required to effectively deal with it.

A number of courses of action are open to the IRT:

- the IRT should report the vulnerability to the vendor or vendors. The responses will require coordination to ensure uniform release of information to the community;
- the IRT may actively examine the source code to assist with understanding and fixing the vulnerability. The IRT members will require exceptional programming skills to perform this task effectively;
- the IRT may become involved with testing patched software to determine that the solution removes the vulnerability, does not introduce new vulnerabilities, and does not cause any functionality to fail.

Many IRTs do not have sufficient resources to pursue this type of activity on a full time basis. Forming trust relationships with other IRTs and requesting their assistance is one mechanism for combining many skilled personnel onto one problem.

Many solutions are not determined within a few hours. Some may take several days. Since the Internet is a global network, it may be advantageous to establish relationships with other international IRTs that are in different timezones. The vulnerability and its current state of analysis can then be passed from one IRT to the next in a global chain, following the daylight and office hours around the world.

8.2. Education and Training

An extremely important role for the IRT is to educate the community on issues relating to security. This may be done in several ways, depending on the requirements of the constituency.

8.2.1. Advisories

Advisories generally are a document that raises a single issue about computer security. They are usually a long living document, and may be referred to from time to time. Examples of content may be the announcement of a vulnerability and solution, a suggestion relating to some administrative matter (such as the use of login banners), or the announcement of tool kits.

Advisories that announce a vulnerability should contain information on the scope of the vulnerability (the versions and platforms affected), a description of the severity of the problem (including any exploitation), and one or more solutions. It is then up to the constituent to decide the most appropriate solution to apply in their situation.

8.2.2. Conference Presentations

Conference presentations are a mechanism to discuss the latest research or latest trends in computer security. This is a good forum to relate back to the constituency some information that affects them directly, such as the number and severity of the incidents, and general trends that have been determined.

8.2.3. Workshop Presentations

Workshops may take the form of a conference style presentation, or may be more hands on. Hands on security workshops are an effective teaching aid to assist new system administrators in the techniques require to monitor and audit their systems. This requires a lot of preparation time and resources (a laboratory full of systems).

8.2.4. Panel Sessions

This type of session allows several people of differing experience and focus to come together and provide a session of much broader content. This is usually an interactive session with comments and questions invited from the audience. It places the security professionals within reach of the constituency. This is important as the IRT must always maintain contact with its constituency.

8.2.5. Journal Articles

Formal papers may be written and published in journals. Less formal papers may be published in magazines and editorials. These papers should always be made available to the Internet, provided it does not breach copyright.

8.2.6. Exercises

This is a concept that was developed by SERT, but has not been actively employed yet. A "security exercise" was designed to be a short 10 to 15 minute activity that increased the security of the computer systems by a small amount. It was felt that the constituency contained a wide range of experience and expertise among the system administrators, and some of the basic skills of system administration and security auditing could be steadily improved. Many system administrators are too busy to attend lengthy courses. The security exercise was in essence a correspondence course without assessment, and without lengthy study. An example of a security exercise might be to request system administrators to examine one day's system log files. Any lines contained in that log file that are not understood should be investigated and researched. These exercises would be issued regularly.

8.2.7. Book Reviews

Security is a rapidly growing topic. Many books are appearing, some better than others. It is impossible for each member of an IRT to read all the books and understand their content. The IRT must choose a subset of available literature for its library, and the chosen books must suit the needs of the IRT and the constituency. This can be determined prior to purchase by reading book reviews.

If the IRT reads a book that they feel is of benefit to the wider community, then they should make a book review available to the constituency and the Internet.

8.2.8. Courses

Traditional education involves classrooms, lectures, and tutorials. This is still an effective form of educating the constituency. Courses may be developed, and run at regular intervals, either at the base of the IRT, or within the constituent sites.

These courses may be presented as a paid service, which covers the cost of preparation, staff time, and travel.

As further advances are made in multi-media, it will not be long before courses that are accessed through the Internet start to appear.

8.2.9. Security Audits and On-site Consulting

Many IRTs are requested to provide on-site consulting and security audits. This may involve policy formulation, examining procedures and suggesting improvements, to acting as a "tiger team" by trying to actively break into the site. Tiger teams in general are not a good idea, as there are legal implications of actively trying to break into a computer system, and it may reveal sensitive exploitation details to the general public.

8.2.10. Goals

It is the goal of the education process to raise the community's awareness to security. It is the author's experience that the majority of incidents occur due to poor system configurations and poor system management. A competent, educated, and diligent system administrator has a much better chance of defending against intruders and detecting them quickly if there is an intrusion, thereby reducing the severity and scope of the incident.

The education role must give sufficient information to all system administrators to raise their awareness of security issues. It may involve discussing new tools and techniques, highlighting when new versions of software fix vulnerabilities, describing methods of attack used by intruders, or assist in resolving local legal issues.

The community must be made aware that security is a total community response. One vulnerable site may put the entire community at risk. "I don't need a good password because all I ever do is word process". This attitude requires modification. Once this account is compromised, it provides a stepping stone into the community. Step by step, the intruder may steadily compromise systems. Denying the intruder the initial foothold into the network prevents these attacks.

8.3. Research and Development

The IRT may choose to perform active research and development with the aim of providing tools and techniques that improve security. This is especially true of IRTs that are based at research or educational institutions.

Many excellent tools have been developed which are now in common use. Without these tools, many more systems would be compromised. Research and development is extremely important, but must be adequately funded to achieve any results. Many more tools are required that assist users who are not computer literate. As the cost of computing decreases, this allows more inexperienced system administrators to be connected to the network with their own machines. Configuration tools should make decisions on behalf of the system administrator, and set up sensible default configurations that are secure as well as useable.

Research may also take the form of analysing coding structures, developing tool kits for programmers to use, writing educational material, or developing new ways for information to be processed, presented, or configured.

9. Operations - Off-site

As indicated before, Incident Response Team staff will be required to operate from outside the secure environment from time to time. This may be as a result of visiting another site to assist them, attending a conference or workshop, or operating after hours. If access to the secured network is to be granted to team members, then they must be made aware of the possibility of trojan horses and network sniffers operating in the network. This may be the result of using equipment that is administered by people other than the IRT.

Some form of non-replayable authentication sequence is required. This may take the form of one-time password generators, software systems such as S/Key, or some other locally developed mechanisms. These systems should be secure, such that no matter how many password "tokens" are captured, the next password in the series cannot be guessed or determined.

Since computer incidents may occur 24 hours a day, 7 days a week, it is important that team members be able to operate from a number of bases, including their private home. This reduces the impact of incidents on the team's private lives by not requiring them to be physically located on the premises during the investigation. This may require extra equipment such as secondary telephone lines (allowing access to the systems simultaneously as voice access), terminal equipment, modems, pagers, mobile phones, and so on.

If staff are to be on call 24 hours a day, then they require mechanisms for making long distance telephone calls without incurring a charge to the premises they are calling from. This alleviates the problem of being at a friend's place when required to make several international phone calls. A mobile phone removes this requirement, but a mobile phone may not possess adequate security. All team members should be able to be contacted during emergency situations. This may require home telephone numbers or the use of pagers. If a team member knows they cannot be contacted (for example, on a boat fishing!), the other team members should be made aware of this.

On-call staff members must be able to contacted by the constituency and other IRTs, without invading on the privacy of those members. In addition, it should be possible to rotate the on-call status among staff members without the adjusting the way the community contacts the team. This may be achieved through call forwarding, pagers, or staff to answer the central phone 24 hours a day.

One issue often overlooked is the ability to travel into the work premises should it be required. It is not possible to ask team members to dedicate their lives to the IRT 24 hours a day. People may for example be attending a celebration in which an amount of alcohol might have been consumed. If an incident occurs at this point, the team member may not be able to drive into the office. Mechanisms should be made available to allow for the use of taxis or some other arrangements in unusual circumstances.

During large conferences (particularly ones hosted by the sponsoring organisation), it may be required that a significant number of the team attend the conference. If the team is small, this could easily account for all members. Plans and equipment should be put into place to allow the *entire* operations to be moved between cities. This may involve telephone access, the ability for the team to be contacted, and access to the secure computer systems. This setup could also be used in the case of emergencies where the office is inaccessible (for example, a bomb threat during an incident).

Response time to incidents may be critical. Careful thought given to off-site operations may significantly reduce the response time to an incident, and allow many team members to contribute effort. This is especially important when the incident is large and complex.

10. Working with the Larger Community

Ultimately, the aim of most Incident Response Teams is to reduce the number and severity of incidents[3]. This cannot be effectively achieved by sitting in the office and waiting for the phone calls to advise of a new incident. Only through education and understanding of security issues can a reduction of incidents be achieved.

The education role is one of the most important, and can take a significant amount of resources from the IRT. However, successful ventures in this area will ultimately have a positive effect on the rest of the IRT by reducing the number and severity of incidents to respond to.

Education may be achieved in several ways. With each incident, some small amount of extra effort should be spent in increasing the knowledge of the affected system administrator. Introduce them to a new security tool, or work with them so they completely understand why this incident occurred and how to prevent it happening again. This helps one system administrator.

Analyse the incident. Why did this incident occur? Inexperience, or is it a general problem to the constituency? If the wider community may benefit, then spend more effort in designing an "education

[3]Despite the feelings of some users of USEnet!

package" that can be given to the rest of the constituency. This package need not explain who was affected by this vulnerability, nor even how to actively exploit it. If the constituency trusts the IRT, then they will act on the information and seek independent verification later. The package should contain a number of items:

- A description of where the problem lies. This should include affected version if possible as not all versions may be affected;
- A description of the severity of the problem. If this problem can be used to gain privileged access, then it should be acted upon quickly;
- An idea of how widely this information is distributed. If it is well known, and currently being actively exploited widely, then the constituency should act quickly to resolve it;
- A solution to the problem. Sometimes the solution is not optimal due to the vulnerability affecting more than one platform. Solutions such as "disable the service" may be the only option if no adequate solution can be found quickly. Include a description of the impact of applying each solution. This decision of which solution to adopt should be made by the constituents; not the IRT. Each site knows their own risks and will act accordingly. Solutions such as "disconnect from the network" are far more severe as in general, the final solution will be distributed through the network. A better solution in these circumstances might be to filter all but trusted sites and the IRT until further notice. Make it a policy of the team to *never* post information without also posting a solution! This helps noone.

In general, the information package that is released will become public information. It may be challenged in the future, and the team must be able to defend it. Check each statement for truthfulness, and act according to the best of the team's ability in the present situation.

Many teams already release this type of information in a document called an "Advisory". These advisories assist sites to increase their security, thus preventing compromises utilising the same vulnerability.

Many sites will wish to know the extent of the security problem so that they can justify the required level of security staffing. Attendance at conferences and presenting papers containing statistics, trends, and future predictions provides good public relations to the constituency, as well as feedback of the situation. It is very easy for a site to become complacent about security if they believe that no incidents are occurring. It may be that many sites around experiencing security incidents continually.

Conferences, workshops, "birds-of-a-feather" sessions, rump sessions, panel sessions, and so on are an ideal forum for providing education on security. The number of topics that could be covered are almost limitless including security policies, secure programming practices, good system administration skills, disaster recovery, and tool analysis. Well-presented papers and sessions will increase the respect of the constituency for the professionalism of the IRT.

Another forum could be a security training workshop, dedicated to only security issues. This is a lot of work, and needs to be well organised. If the equipment can be obtained, this is the best place to organise hands-on training of configuring systems, and making them more secure. Many basic system administration skills must be learned on a running system, and production systems are not always the best platform to do this. Once the basic system administration skills are covered, security tools could be installed, and a demonstration of their effectiveness explored. Hands-on training is more effective than conference proceedings, advisories, or telephone calls. These workshops could be performed as a charged service, helping to recover the cost of preparing them, and the use of the equipment.

As the operating systems, third party packages, and configuring them becomes more complex, it is becoming increasing difficult to state with certainty that a system is configured correctly. IRTs are well placed to contribute to the pool of available security tools. In particular, it will become more important to assist novice system administrators with basic system administration skills. Configuration tools, security assessment and enhancement tools, and a number of "wrappers" to make their use easier for a system administrator will little or no knowledge will ensure that these tools are at least applied in some minimal form, thereby increasing the security of those machines. If the tools are too difficult to use, or contain too many options, they will not be used at all.

For example, SERT in conjunction with Sun Microsystems developed the Megapatch. One problem with applying security patches to SunOS was that it was difficult to determine which patches should be applied and in which order[4]. The Megapatch is a tool that is applied to a newly installed SunOS system, and applies all known security patches in the correct order. In addition, it installs and configures a number of security assessment and enhancement tools such as COPS, Tripwire, and TCP Wrapper, and enables C2 security. These security enhancements are provided with conservative initial configurations that protect the system from unauthorised intrusion. The tool is designed to make it easy to apply by the novice system administrator.

Ultimately, it is the role of the IRT to become the trusted source of security information in the community. The constituents should be given the opportunity to learn that the IRT has competent and diligent staff. If the IRT indicates that a security vulnerability exists, then the constituents should be confident that the IRT has either tested the vulnerability and solutions, or has a high degree of confidence that the information is correct. In addition, the constituents should learn that the IRT has integrity and honours the privacy of each institution.

11. Working with FIRST

The Forum of Incident Response and Security Teams (FIRST) is a collection of IRTs, vendors, and other interested parties that are working together to improve computer security. Since many of the IRTs are formed to cater to the particular requirements of their constituencies, they cannot effectively deal with other constituents. This is the reason that there are many IRTs. FIRST is designed to improve the communication and cooperation between the IRTs and registered vendors.

FIRST basically supplies an umbrella secretariat to assist communication between all of its members. Much of the work within FIRST is done on a volunteer basis, and supplied from within the various FIRST members.

FIRST provides a forum for IRTs and other security experts to discuss security vulnerabilities, and cooperate to find an acceptable solution. Other information that is shared may be intelligence on methods used by intruders, warnings of security situations to be aware of, draft advisories for review, and ensuring that all members see publicly released information from the wide range of sources. The benefits to be gained from membership in FIRST are directly proportional to the amount of effort that the IRT is willing to supply.

12. Conclusion

Forming an Incident Response Team in the 90s is a difficult task. Fortunately, there are many willing individuals that are able to provide guidance that will help the newly formed team avoid many pitfalls. It is possible, and highly desirable, to perform much of the establishment work prior to commencement of the team. Once operations start, then the time available to formulating the new team will become limited.

[4] This problem is not unique to Sun Microsystems!

Policies, procedures, equipment, premises, contacts, and staff should be established before commencing operations. More likely however, is that many of these items will be missing or inadequate. The team must struggle on as best as it can while it is forming. Clear communication with the constituency will alleviate the startup problems and any confusion that might be caused by them.

Obtaining good advice from other established teams and establishing good practices will make the startup of the new team far less difficult, and will take the team from strength to strength in their operations. Once the incident load increases, there will be little resource to "redo" some aspect of the operation of the team. Getting it right the first time will remove the need to expend precious resources on fixing a problem, as well as converting over the existing procedures and data to the new operation.

12.1. Acknowledgments

What makes me an expert on this topic? Simple - I had to do it once! I could not have achieved the formation of a security IRT in Australia without the tremendous support from many individuals.

Firstly, Tom Longstaff for assisting with the ideas that are contained within this paper. We met in Pittsburgh in August 1993 and were discussing the various issues that require resolution when forming an IRT. Next thing, Tom had captured all of our ideas in a set of notes that formed the basis of this paper.

Moira West: If ever there was a heroine in the security field (in my opinion), it is Moira. She has weathered my abuse, my triumphs, my disappointments, my anger, my frustration, and my humour throughout the time that SERT has been operating. Through all of this, she has provided enormous support and guidance, and for that I will always be grateful.

Barbara Fraser: Barbara visited Australia for a conference just prior to our learning that our government funding request was unsuccessful. Barbara paved the way for forming the IRT in Australia. She firmly placed the idea into the minds of the people that had the power to make this happen. She showed Australia what an IRT was all about, and why Australia needed to have its own. Without that visit, Australia may not have an IRT today.

IR Group in CERT: These people are modern day heroes. They wear abuse, scorn, derision, lies, and back stabbing, and still keep trying to help the very people that do this to them. Being at the forefront of this technology and procedures means that mistakes are made. The Internet community is not forgiving of mistakes. Keep your chin up guys - there are more people out there that appreciate your efforts than there are who fight you!

Klaus-Peter Kossakowski: Just as SERT was commencing, Germany formed its DFNCERT team. I didn't learn of this until August. This was the start of more work to resolve international cooperation issues. Peter has had to fight just as hard for funding as we have. He has supported the SERT team absolutely, and I look forward to further cementing our relationship with DFNCERT. I asked Peter to contribute to this panel as they have recently also been through the exercise of forming an IRT, so they are also experts!

Georgia Killcrece: Although Georgia is part of CERT's IR group, I have singled her out for agreeing to participate in this panel. Georgia knows what it is like to try and operate in a hostile environment, and her experience has helped us face our constituents with more confidence.

Sandy Sparks: Sandy also agreed (was coerced) to be on this panel of presenters. SERT has had only minor dealings with CIAC, but has been impressed on all occasions with their integrity and professionalism. This cannot be achieved in a team without strong management, which Sandy will now educate me in!

Alan Coulter, Geoffrey Dengate, John Noad: The Directors of the Computer Centres of the three cooperating Brisbane Universities. Without their vision and support, the SERT team would still be a part of people's imaginations.

Graham Rees: My immediate manager and good friend. He has had to tolerate and soothe my ruffled feathers when the going got tough, the budget was lean, and there was no more resource to apply to the problem - "Just do the best you can!". Graham was always willing to help and support, and that is a great boost when times get tough.

Finally, Rob McMillan. Rob and I created the initial stages of SERT. Every IRT could use a person like Rob: intelligent, talented, full of integrity, full of great ideas, trustworthy, and grossly underpaid!

13. Information Sources

This section contains a number of papers, articles, security tools, and general information sources. These are not the sources of information used to create this paper, but are sources of security information that a newly forming IRT may find useful to obtain and peruse. These references have been used at different times by the author in other papers.

13.1. Papers

[Alv90] De Alvare A. M., *How Crackers Crack Passwords or What Passwords to Avoid*, Proceedings of the UNIX Security Workshop II, Portland, August 1990.

[BB91] den Boer B. and Bosselaers A., *An Attack on the Last Two Rounds of MD4*, Proceedings of the Crypto'91 conference, Santa Barbara, August 1991.

[BB93] den Boer B. and Bosselaers A., *Collisions for the compression function of MD5*, Pre-proceedings of the EUROCRYPT 93 conference, Lofthus, May 1993.

[Bis87] Bishop M., *How to Write a Setuid Program*, ;login, Volume 12, Number 1, January/February 1987.

[Bis92a] Bishop M., *Proactive Password Checking*, Proceedings of the 4th Workshop on Computer Security Incident Handling, Denver, August 1992.

[BKS90] Baran F., Kaye H., and Suarez M., *Security Breaches: Five Recent Incidents at Columbia University*, Proceedings of the UNIX Security Workshop II, Portland, August 1990.

[BM91] Bellovin S. and Merritt, M., *Limitations of the Kerberos Authentication System*, Proceedings of the USENIX Winter 1991.

[Bra90] Brand R., *Coping with the Threat of Computer Security Incidents: A Primer from Prevention through Recovery*. CERT 0.6, June 1990.

[Bro93] Brown L, *On Implementing Security Extensions to the TCP Transport Layer*, Proceedings of the 16th Australian Computer Science Conference (ASCS-16), Brisbane, February 1993.

[Che92] Cheswick W.. *An evening with Berferd in which a Cracker is Lured, Endured, and Studied*, Proceedings of the Winter USENIX Conference, San Francisco, January 1992.

[Cly93] Clyde R., *DECnet Security (Not Necessarily an Oxymoron)*, Computers and Security, March 1993.

[Coh92] Cohen F., *A Formal Definition of Computer Worms and Some Related Results*, Computers and Security, Volume 11, Number 7, November 1992.

[Cov90] Covert J., *Functional Specification for Callouts for LOGINOUT & DECnet Session*, Version T1.0.0, Digital Equipment Corporation, July 1990.

[Cur90] Curry D., *Improving the Security of your UNIX System*, ITSTD-721-FR-90-21, SRI International, April 1990.

[Din90] Dinkel C., *Secure Data Network System (SDNS) Network, Transport and Message Security Protocols*, NIST, NISTIR-90/4250, March 1990.

[Edw90] Edwards B., *How to Survive a Computer Disaster*, Proceedings of the DECUS Symposium, August 1990.

[FIP77] Federal Information Processing Standards Publication 46, *Data Encryption Standard*, National Bureau of Standards, U.S. Department of Commerce, January 1977.

[HY92] Harn L. and Yang S., *Group Oriented Undeniable Signature Schemes without the Assistance of a Mutually Trusted Party*, Proceedings AUSCRYPT '92, Gold Coast, December 1992.

[JM91] Janson P and Molva R., *Security in Open Networks and Distributed Systems*, Computer Networks and ISDN Systems, Volume 22, Number 5, October 1991.

[KC90] Kaplan R., and Clyde R., *Viruses, Worms, and Trojan Horses - Part VI: The War Continues*, Proceedings DECUS Fall 1990, Las Vegas, 1990.

[KCS90] Kohl J., Neuman B., and Steiner J., *The Kerberos Network Authentication Service*, MIT Project Athena, Version 5 Draft 3, October 1990.

[KK92] Koblas D. and Koblas M., *SOCKS*, Proceedings of the USENIX Security Symposium, 1992.

[Kle90] Klein D., *"Foiling the Cracker": A Survey of, and Improvements to, Password Security*, Proceedings of the UNIX Security Workshop II, Portland, August 1990.

[Kur90] Kuras J., *An Expert Systems Approach to Security Inspection of UNIX*, Proceedings of the UNIX Security Workshop II, Portland, August 1990.

[LAB92] Lampson B., Abadi M., Burrows M., and Wobber E., *Authentication in Distributed Systems: Theory and Practice*, acm Transactions on Computer Systems, November 1992.

[Lau92] Laun, R., *Asymmetric User Authentication*, Computers and Security, Volume 11, Number 2, April 1992.

[Law93] Lawrence L., *Digital Signatures - Explanation and Usage*, Computers and Security, Volume 12, Number 3, May 1993.

[LS93] Longstaff T. and Schultz E., *Beyond Preliminary Analysis of the WANK and OILZ Worms: A Case Study of Malicious Code*, Computers and Security, Volume 12, Number 1, February 1993.

[Mor90] Moraes M., *YP is not secure*, Security Digest, Volume 3, Issue 12, May 1990.

[RSA78] Rivest R., Shamir A., and Adleman L., *A Method for Obtaining Digital Signatures and Public-key Cryptosystems*, Communications of the ACM, February 1978.

[Spa88] Spafford, E., *The Internet Work Program: An Analysis*, Technical Report CSD-TR-823, Department of Computer Science, Purdue University, November 1988.

[Spa92] Spafford E., *OPUS: Preventing Weak Password Choices*, Computers and Security, May 1992.

[TAP90] Tardo J., Alagappan K., and Pitkin R., *Public Key Authentication using Internet Certificates*, Proceedings of the UNIX Security Workshop II, Portland, August 1990.

13.2. Books

[Arn93] Arnold N., *UNIX Security: A Practical Tutorial*, McGraw-Hill Inc., 1993.

[Bha93] Bhaskar K., *Computer Security: Threats and Countermeasures*, NCC Blackwell, 1993.

[CLS91] Caelli W., Longley D., and Shain M., *Information Security Handbook*, Stockton Press, 1991.

[DEC88a] *Guide to DECnet-VAX Networking Version 5.0*, Digital Equipment Corporation, April 1988.

[DEC88b] *VMS Access Control List Editor Manual Version 5.0*, Digital Equipment Corporation, April 1988.

[DEC89a] *Guide to VMS System Security Version 5.2*, Digital Equipment Corporation, June 1989.

[DEC89b] *VAX C Run-Time Library Reference Manual Version 3.1*, Digital Equipment Corporation, December 1989.

[DEC90] *VMS Authorize Utility Manual Version 5.4*, Digital Equipment Corporation, August 1990.

[Far91b] Farrow R., *Unix System Security: How to Protect your Data and Prevent Intruders*, Addison-Wesley, April 1991.

[Gro93] Grottola M., *The UNIX Audit: Using UNIX to Audit UNIX*, McGraw-Hill Inc., 1993.

[GS91] Garfinkel S. and Spafford G., *Practical UNIX Security*, O'Reilly and Associates, Inc., 1991.

[IBM89] *Virtual Machine/Directory Maintenance - Operation and Use, Release 4*, International Business Machines,1989.

[MP92] Mui L. and Pearce E., *X Window System Administrator's Guide*, O'Reilleys & Associates Inc., 1992.

[OSI92] *The OSI Security Package, OSISEC Users Manual V0.2*, July 1992.

[SS94] Shaffer S. and Simon A., *Network Security*, AP Professional, 1994.

[Ste90] Stevens W., *UNIX Network Programming*, Prentice Hall, 1990.

[Sto89] Stoll C., *The Cuckoo's Egg*, Doubleday, 1989.

[Sun90a] System and Network Administration, SUN Microsystems, Revision A, March 1990.

[Sun90b] SunOS Reference Manual, Volume 1, SUN Microsystems, Revision A, March 1990.

[Sun90c] SunOS Reference Manual, Volume 2, SUN Microsystems, Revision A, March 1990.

[Tan89] Tanenbaum A., *Computer Networks*, Prentice-Hall International Inc.1989.

13.3. Security Tools

[Bis92b] Bishop M., *README file for passwd+*, anonymous ftp from dartmouth.edu, June 1992.

[Far91a] Farmer D., *README.1 file from COPS system*, anonymous ftp from cert.org, November 1991.

[Goa92] Goatley H., *Supervisor Reference Guide*, anonymous ftp from ftp.spc.edu, October 1992.

[Hei90] Heirtzler J., *shadow.howto file from shadow system*, anonymous ftp from csc2.anu.edu.au, April 1990.

[Hoo90] Hoover C., *README file from npasswd system*, anonymous ftp from ftp.cc.utexas.edu, March 1990.

[KHW93] Karn P., Haller N., and Walden J., *S/Key One Time Password System*, anonymous ftp from thumper.bellcore.com, July 1993.

[KS92] Kim G. and Spafford E., *README file from Tripwire system*, anonymous ftp from cert.org, November 1992.

[LeF92] LeFebvre W., *Restricting Network Access to System Daemons under SunOS*, securelib system, anonymous ftp from eecs.nwu.edu, 1992.

[MLJ92] McCanne S., Leres C., and Jacobson V., *README file from tcpdump system*, anonymous ftp from ftp.ee.lbl.gov, May 1992.

[Muf92] Muffett A., *"Crack Version 4.1" A Sensible Password Checker for Unix*, anonymous ftp from cert.org, March 1992.

[Ney92] Ney S., *README file from TAP system*, anonymous ftp from ftp.cs.tu-berlin.de, March 1992.

[SSH93] Safford D., Schales D., and Hess D., *Texas A & M Network Security Package Overview*, anonymous ftp from sc.tamu.edu, July 1993.

[Ven92] Venema W., *BLURB file from TCP Wrapper system*, anonymous ftp from cert.org, June 1992.

[Zim92] Zimmermann P., *README file from PGP system*, anonymous ftp from ghost.dsi.unimi.it, November 1992.

13.4. Articles

[CER92] Computer Emergency Response Team, *Internet Security for UNIX System Administrators*, Presented at AARNet Networkshop, December 1992.

[CER93] Computer Emergency Response Team Advisory 93:14, *Internet Security Scanner (ISS)*, September 1993.

[Hey93a] Van Heyningen M., *RIPEM Frequently Asked Questions*, USEnet newsgroup alt.security.ripem, 31 March 1993.

[Hey93b] Van Heyningen M., *RIPEM Frequently Noted Vulnerabilities*, USEnet newsgroup alt.security.ripem, 31 March 1993.

[SER93] Security Emergency Response Team Advisory 93.04, *Guidelines for Developing a Sensible Password Policy*, June 1993.

[TIS93] *TIS/PEM FAQ (Frequently Asked Questions)*, anonymous ftp from ftp.tis.com, June 1993.

13.5. Standards

[ISO92] International Standards Organisation ISO 9594-8: *The Directory: Authentication Framework*, 1992 (also known as CCITT Recommendation X.509).

[RFC783] Sollins K., *The TFTP Protocol (Revision 2)*, Network Working Group, RFC783, June 1981.

[RFC1094] Sun Microsystems Inc., Network File System Protocol Specification, Network Working Group, RFC1094, March 1989.

[RFC1319] Kaliski B., *The MD2 Message-Digest Algorithm*, Network Working Group, RFC1319, April 1992.

[RFC1320] Rivest R., *The MD4 Message-Digest Algorithm*, Network Working Group, RFC1320, April 1992.

[RFC1321] Rivest R., *The MD5 Message-Digest Algorithm*, Network Working Group, RFC1321, April 1992.

[RFC1421] Linn J., *Privacy Enhancement for Internet Electronic Mail: Part I: Message Encryption and Authentication Procedures*, Network Working Group, RFC1421, February 1993.

[RFC1422] Kent S., *Privacy Enhancement for Internet Electronic Mail: Part II: Certificate-Based Key Management*, Network Working Group, RFC1422, February 1993.

[RFC1423] Balenson D., *Privacy Enhancement for Internet Electronic Mail: Part III: Algorithms, Modes, and Identifiers*, Network Working Group, RFC1423, February 1993.

[RFC1424] Kaliski B., *Privacy Enhancement for Internet Electronic Mail: Part IV: Key Certification and Related Services*, Network Working Group, RFC1424, February 1993.