RTIR incident handling work-flow

JANET CSIRT

INTRODUCTION

Request Tracker for Incident Response (RTIR) is the incident handling and ticketing system used by JANET CSIRT, and builds upon the popular open source ticketing system Request Tracker (RT). RTIR was originally developed for JANET CERT by Best Practical, and further development since has been guided under the remit of the RTIR Working Group as part of TF-CSIRT.

This document is intended as an introduction to how RTIR works in practice. It closely follows the workflow used within JANET CSIRT and IRIS-CERT, whilst the content and structure of the document is closely based upon the internal workflow documentation used in JANET CSIRTs Quality Management System.

A small document such as this cannot pretend to give even an overview of every single function available to the user of RTIR and instead we have concentrated only on the features that an incident handler will come across on a daily basis. We hope that after reading this you have a clear idea of how RTIR should work, and also how the workflow can be adapted to fit your organization.

We have assumed that you have already installed and have access to an installation of RTIR available so you can read the text whilst following within the application. Although it has been avoided where possible, at some points we assume that you have some familiarity with programming in perl, and working with perl modules. This is necessary if you wish to develop, extend and change the way that RTIR works.

If you have any comments or feedback, or your organization is using RTIR in a way that you feel is worth including in this document please get in touch.

1 PURPOSE

This work instruction is to ensure that Computer Security Incident Response Team (CSIRT) members carry out incident handling duties consistently and effectively, and in particular that they follow an agreed work-flow pattern for the application Request Tracker for Incident Response (RTIR).

2 SCOPE

This work instruction is principally for the guidance of CSIRT team members, both at times when they are the person receiving all incoming messages and at other times when they may be working on ongoing incidents of their own or on related correspondence.

Some parts specify obligations for the CSIRT Manager.

3 PROCESS

3.1 Outline

The investigation of an issue reported to CSIRT may involve many different members of staff working on the same problem, possibly at different times. The RTIR software (see below) is intended to help team members to process and record information relating to the incident in a consistent way, and this work instruction indicates how the application should be used to best effect.

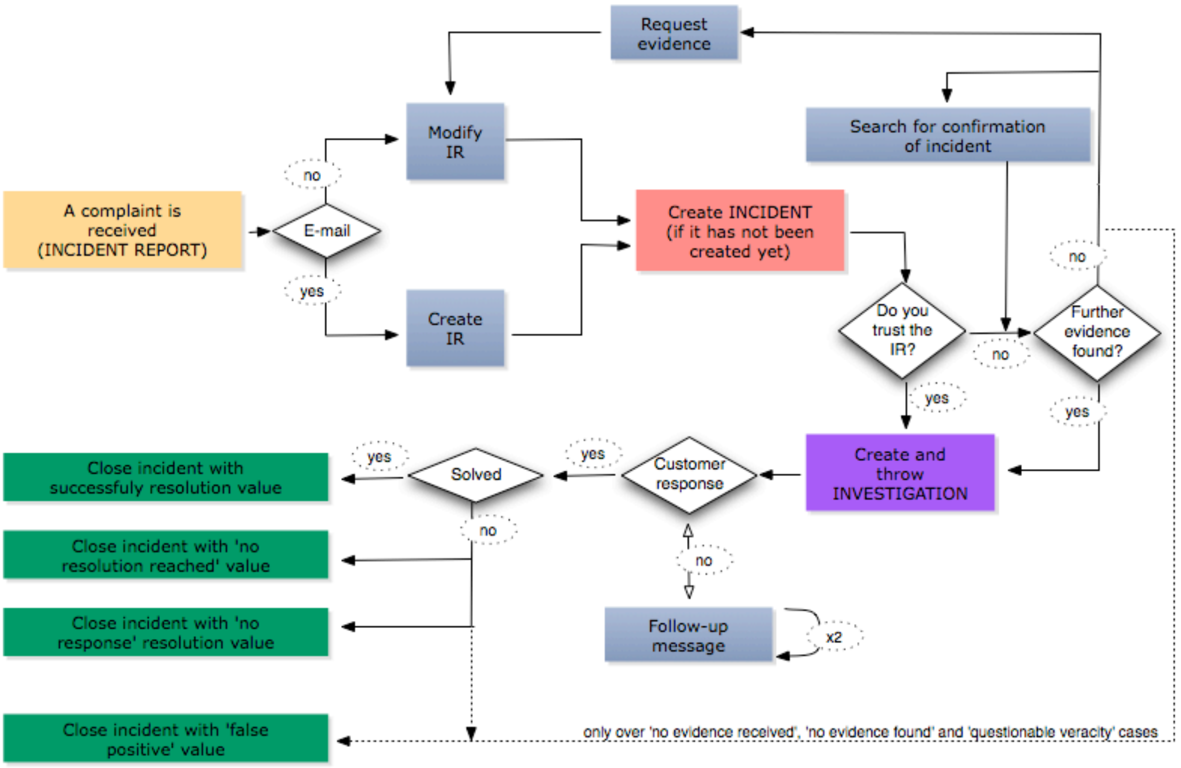
Request evidence
Modify IR
Search for confirmation of incident
no
A complaint is received (INCIDENT REPORT)
E-mail
Create INCIDENT (if it has not been created yet)
yes
no
Create IR
Do you trust the IR?
no
Further evidence found?
yes
yes
Close incident with successfuly resolution value
yes
Solved
yes
Customer response
Create and throw INVESTIGATION
no
Close incident with 'no resolution reached' value
no
Close incident with 'no response' resolution value
Follow-up message
x2
Close incident with 'false positive' value
only over 'no evidence received', 'no evidence found' and 'questionable veracity' cases

Figure 1 The RTIR Workflow

3.2 RTIR

RTIR is a customised user interface which sits on top of Request Tracker (RT), a popular ticketing system. Everyday use of RTIR is through a web interface and does not require any additional software to be installed on the user's machine.

Before a team member can access the RTIR interface they will need to login to the system with their user account and password. RTIR treats items created using any particular account as owned by the corresponding user, and attributes actions in the same way.

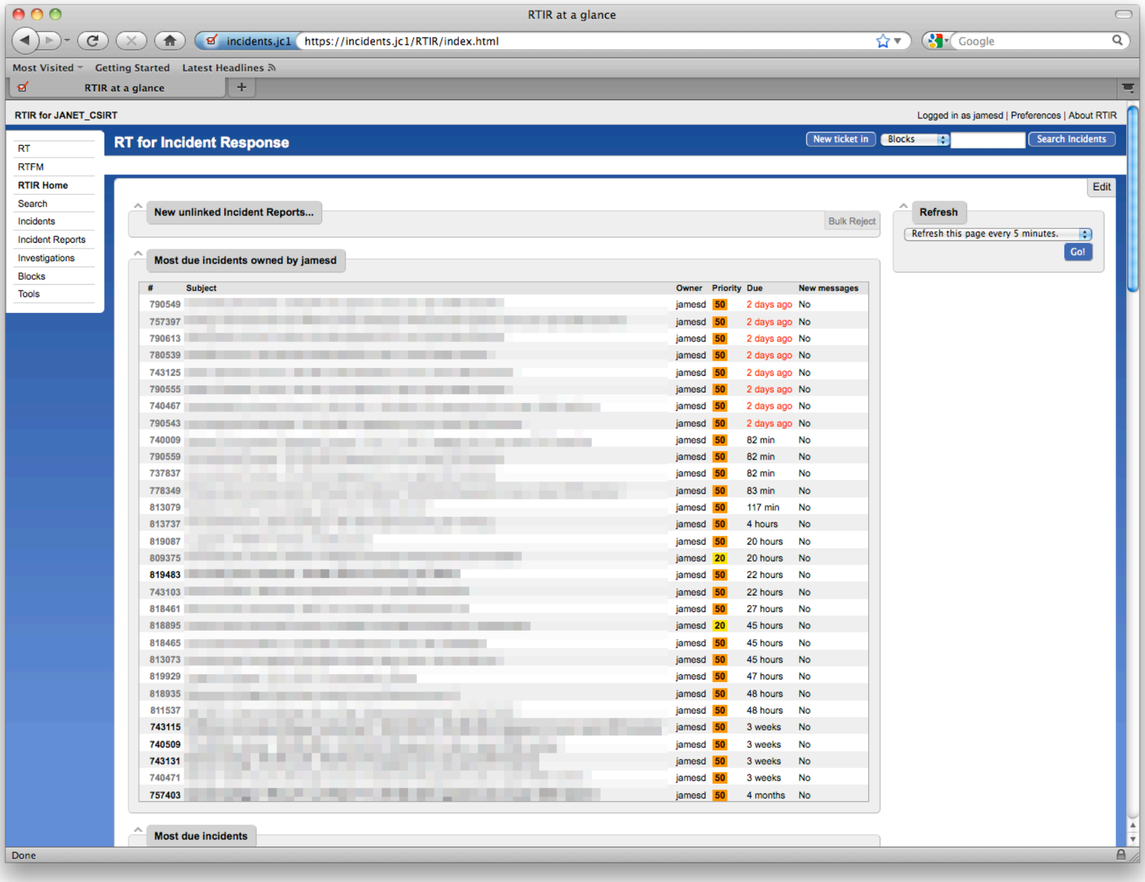

#	Subject	Owner	Priority	Due	New messages
790549	[redacted]	jamesd	50	2 days ago	No
757397	[redacted]	jamesd	50	2 days ago	No
790613	[redacted]	jamesd	50	2 days ago	No
780539	[redacted]	jamesd	50	2 days ago	No
743125	[redacted]	jamesd	50	2 days ago	No
790555	[redacted]	jamesd	50	2 days ago	No
740467	[redacted]	jamesd	50	2 days ago	No
790543	[redacted]	jamesd	50	2 days ago	No
740009	[redacted]	jamesd	50	82 min	No
790559	[redacted]	jamesd	50	82 min	No
737837	[redacted]	jamesd	50	82 min	No
778349	[redacted]	jamesd	50	83 min	No
813079	[redacted]	jamesd	50	117 min	No
813737	[redacted]	jamesd	50	4 hours	No
819087	[redacted]	jamesd	50	20 hours	No
809375	[redacted]	jamesd	20	20 hours	No
819483	[redacted]	jamesd	50	22 hours	No
743103	[redacted]	jamesd	50	22 hours	No
818461	[redacted]	jamesd	50	27 hours	No
818895	[redacted]	jamesd	20	45 hours	No
818465	[redacted]	jamesd	50	45 hours	No
813073	[redacted]	jamesd	50	45 hours	No
819929	[redacted]	jamesd	50	47 hours	No
818935	[redacted]	jamesd	50	48 hours	No
811537	[redacted]	jamesd	50	48 hours	No
743115	[redacted]	jamesd	50	3 weeks	No
740509	[redacted]	jamesd	50	3 weeks	No
743131	[redacted]	jamesd	50	3 weeks	No
740471	[redacted]	jamesd	50	3 weeks	No
757403	[redacted]	jamesd	50	4 months	No

Figure 2 - RTIR main page

3.3 Receiving an Incident Report via email

RTIR’s incident handling system relies primarily on e-mail.

E-mail messages reporting incidents (see 3.3.1) are sent to an address configured by CSIRT. Messages that constitute on-going correspondence in the handling of a ticket (see 3.16) include a number in the form [A CSIRT #34159] and are automatically appended to the corresponding RTIR ticket. All other messages are stored as new Incident Reports and appear in the *New unlinked Incident Reports* section of the RTIR homepage.

As the incident handling address has been in widespread use over many years, it receives a large amount of spam. Most unwanted messages are filtered out and do not reach RTIR, but those that do get through will also appear in this *New unlinked Incident Reports* section.

Some email messages will require decryption, validation or both using PGP software. It is possible to do this manually; however, RTIR can also provide some PGP functionality through integration with GnuPGP. For more information see ‘perldoc lib/RT/Crypt/GnuPG.pm’.

Some e-mail messages will arrive with MIME attachments, which RTIR stores with the message.

3.3.1 Sources of reports

E-mail messages to CSIRT can come from any source including:

- Technical and support staff of customer organizations
- Representatives of other ISPs and Communications Service Providers (CSPs), who may or may not be from Incident Response teams
- Customer Service Desk (usually forwarded messages)
- End users in customer organizations
- Customers of other ISPs and members of the public
- Officers of Law Enforcement Agencies
- Representatives of copyright holders
- Automated processes such as Intrusion Detection Systems
- Team members as a result of their own observations
- Other RT queues (although you cannot at this stage directly e-mail RTIR from another queue in the same RT system).

3.4 Incident Reports received other than by e-mail

Enter the details into RTIR by creating an Incident Report; this will appear in the unlinked Incident Reports section of the front page and can be dealt with in the same way as an Incident Report received via e-mail. For telephone calls, include notes of the call. For relevant application files, attach them to the Incident Report; text files including small log fragments should be entered inline. For material received by fax or post, include or attach an electronic record of as much of the information as practicable. Mark the original material with the date and the ticket number and store it in a secure filing area.

The preferred way to create a new Incident Report manually is to choose *[Incident Reports]* from the main page menu bar, then *[New Report]* from the secondary menu bar; complete the details and click *[Create]*.

For an Incident Report created as a result of an out of hours call-out, you may wish to set a different value for the Service Level Agreement (SLA) attribute. The SLA attribute governs how the timing of the response to this incident is reported in the statistics. Different types of incidents may require different levels of response according to the organization's requirements. To set the value for an existing Incident Report (for instance, when a call-out refers to an e-mail message already sent), use the *[Edit]* facility.

In practise we've found that Law Enforcement Agencies are the class of correspondents most likely to use media other than e-mail.

3.5 Initial Checks

Each message in the '*New unlinked Incident Reports*' section should then be dealt with, in ascending incident (and therefore, date) order unless it is clear from the message subjects that any are of particular urgency. First *Take* the report, clicking on the Take link, so that ownership changes from *Nobody* to *your user name*. This helps to avoid more than one incident response operative dealing with the same unlinked report, and ensures that all actions including rejection can be attributed to a particular team member. Once *Taken* (a ticket is *Taken* if it has an *Owner* other than *Nobody*) the email should be examined and it may be possible to *Reject* it (see 3.6).

3.6 Rejecting tickets

A number of legitimate incoming messages, such as reports not involving your own addresses or summary lists of IP addresses noted by firewall operators elsewhere, are for information only and once *Taken* and examined need no further attention. The *[Quick Reject]* button at the top of the Incident Report will change the report's state to *Rejected.* Rejected tickets are still searched for IP address matches, and can be linked to Incidents (see 3.12) although they will only be displayed if their state is *Open* or *Resolved*.

3.7 Bulk Reject

Unwanted e-mail such as spam or viruses and bounces where that address has been used as a spoofed From: address in either of those should also be rejected. At the beginning of a shift it is usually best to start by bulk rejecting tickets which you can identify as spam by seeing only the sender address and the subject. *[Bulk Reject]* will perform the *Take* and *Quick Reject* operations on one or more unlinked reports in a single action.

The bulk reject operation is performed by clicking *[Bulk Reject]* at the bottom of the list of *New unlinked Incident Reports*, selecting in the next screen the tickets you wish to close, then clicking *[Reject]*. After some seconds a confirmation screen appears, listing all the tickets rejected.

3.8 Checking whether the issue is already being dealt with

It is important to avoid generating multiple Incidents about the same issue. You may receive a number of reports about the same problem and these should be grouped together wherever possible. Parts of the messages matching regular expressions for IP addresses, hostnames and email addresses will be clickable links. Clicking on such a link will bring up a *Lookup* screen showing all other tickets which contain this text. It will also perform a whois lookup which returns extra information from external data sources. The absence of any existing tickets containing the same strings (typically IP address) is a good indication that this report is about a new issue, although the search results are limited to only the last two months. You can also change this search string to widen the search, for example removing the last octet from an IP address to see if there are any ongoing incidents about the same /24.

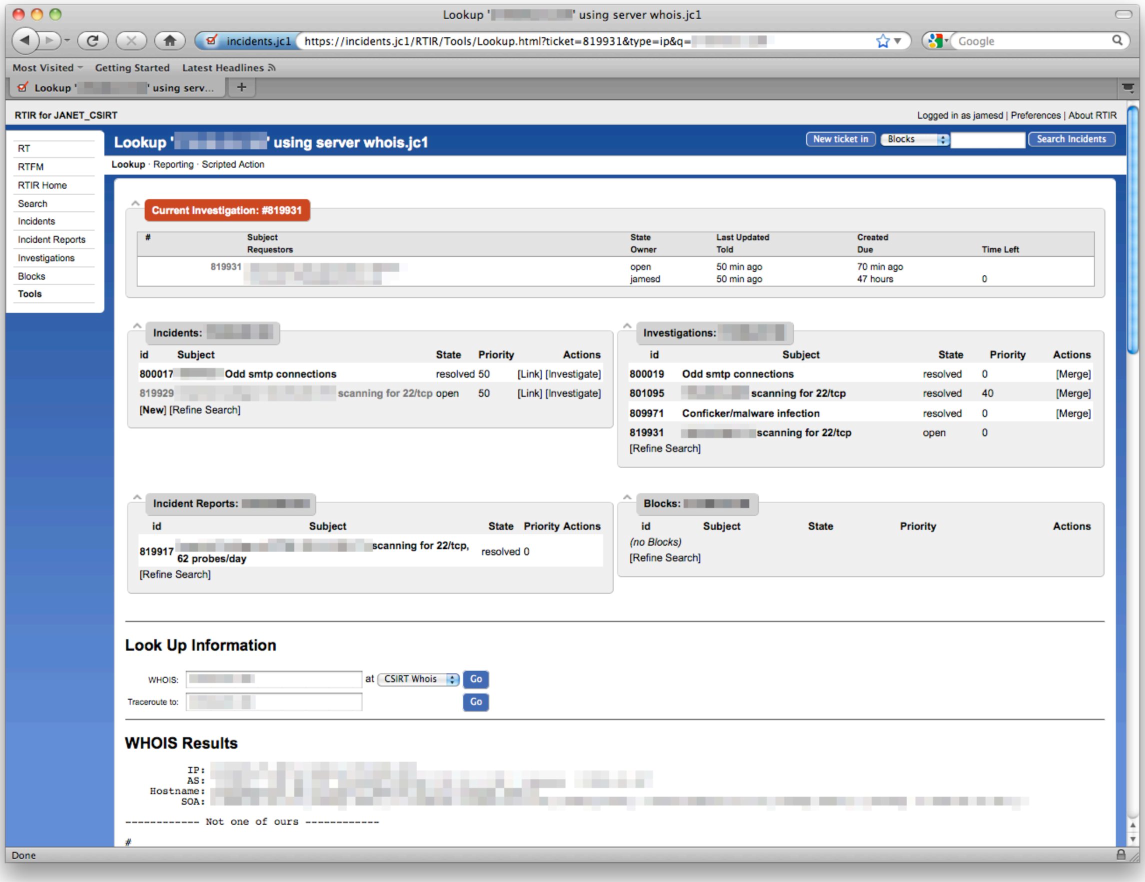

Lookup '[redacted]' using server whois.jc1
incidents.jc1
https://incidents.jc1/RTIR/Tools/Lookup.html?ticket=819931&type=ip&q=[redacted]
Google
Most Visited
Getting Started
Latest Headlines
Lookup '[redacted]' using serv...
RTIR for JANET_CSIRT
Logged in as jamesd | Preferences | About RTIR
RT
RTFM
RTIR Home
Search
Incidents
Incident Reports
Investigations
Blocks
Tools
New ticket in
Blocks
Search Incidents
Lookup · Reporting · Scripted Action
Done

Lookup '[redacted]' using server whois.jc1

Current Investigation: #819931

#	Subject Requestors	State Owner	Last Updated Told	Created Due	Time Left
819931	[redacted]	open jamesd	50 min ago 50 min ago	70 min ago 47 hours	0

Incidents: [redacted]

id	Subject	State	Priority	Actions
800017	[redacted] Odd smtp connections	resolved	50	[Link] [Investigate]
819929	[redacted] scanning for 22/tcp	open	50	[Link] [Investigate]

[New] [Refine Search]

Investigations: [redacted]

id	Subject	State	Priority	Actions
800019	Odd smtp connections	resolved	0	[Merge]
801095	[redacted] scanning for 22/tcp	resolved	40	[Merge]
809971	Conficker/malware infection	resolved	0	[Merge]
819931	[redacted] scanning for 22/tcp	open	0	

[Refine Search]

Incident Reports: [redacted]

id	Subject	State	Priority	Actions
819917	[redacted] scanning for 22/tcp, 62 probes/day	resolved	0	

[Refine Search]

Blocks: [redacted]

id	Subject	State	Priority	Actions

(no Blocks)
[Refine Search]

Look Up Information

WHOIS: [redacted] at CSIRT Whois Go
Traceroute to: [redacted] Go

WHOIS Results

```
      IP: [redacted]
      AS:
Hostname:
     SOA:
------------ Not one of ours ------------
#
```

Figure 3 IP Address Lookup

It is also advisable to look at the list of current Incidents in case any are about the same thing at the same organization but at a nearby IP address. It may then be appropriate to *Link* the report to the existing Incident.

3.9 Incident Reports containing multiple IP addresses

Although we discourage it, some Incident Reports may contain reports about a number of different issues. The method used to deal with such reports largely depends on the reason for the grouping.

For example, a site may report a compromise of a single machine which on further investigation shows that a number of other machines at other sites are compromised. As these are somewhat related it may make sense to group them in the same Incident by launching several Investigations from this single Incident.

Alternatively a correspondent may have just grouped reports of numerous unrelated port scans into a single Incident Report. Such an Incident Report should be split into numerous separate reports and a new Incident created for each event. It would also be wise to reply to the sender pointing out our preferred method (a single report in a single email).

Either of these methods are or will be valid for CSIRT teams. These examples above illustrate how RTIR can deal with them.

Although not originally possible, Incident Reports can now be linked to multiple incidents. (See 3.12, *Linking to an existing Incident*.)

3.10 Creating a new Incident

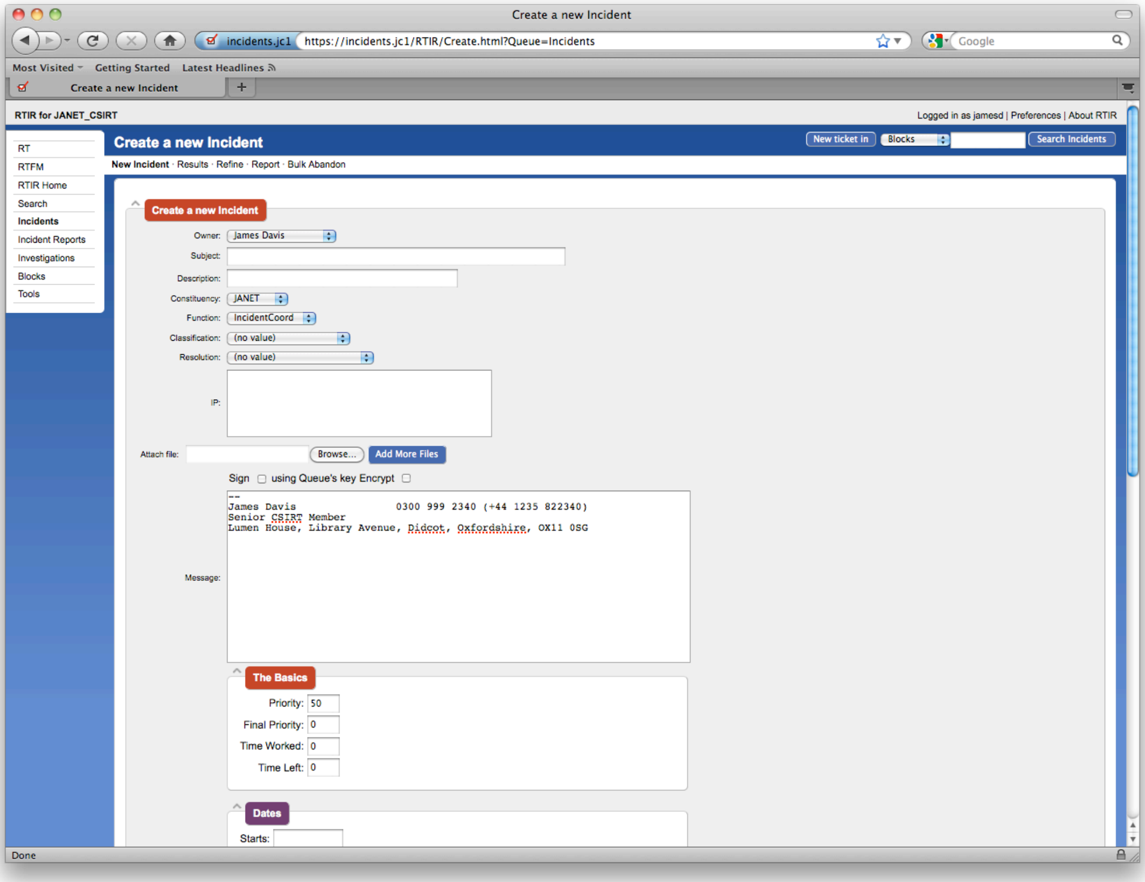

Owner: James Davis

Subject:

Description:

Constituency: JANET

Function: IncidentCoord

Classification: (no value)

Resolution: (no value)

IP:

Attach file:

Sign □ using Queue's key Encrypt □

Message:

```
--
James Davis                    0300 999 2340 (+44 1235 822340)
Senior CSIRT Member
Lumen House, Library Avenue, Didcot, Oxfordshire, OX11 0SG
```

Priority: 50

Final Priority: 0

Time Worked: 0

Time Left: 0

Starts:

Figure 4 Creating a new incident

If the Incident Report refers to an issue for which no Incident already exists, a new Incident must be created. There are two primary ways of doing so: from the Incident Report view screen click on the *[New]* link, or from the Lookup Tool click *[New]* in the Incident search results. Either action will bring up the New Incident creation page where various metadata about the Incident is kept. Some of this metadata is used for SLA reporting and it is important that this is entered correctly.

- Subject line – **MANDATORY**. A few words briefly outlining the incident. It is often useful to include something to identify the organization involved, and the IP address if there is only one (eg *"*Borchester 192.168.19.84 root compromise*"*). A default subject is copied from the Incident Report if there is one, but check as it may be best to edit it to the form above.
- Description (optional). A longer description.
- Constituency – **MANDATORY**. A CSIRT team may be responsible to two different groups of customers; for example it may also provide CSIRT services to another company. Not every team will need more than one option.
- Function (once mandatory, now optional). Some CSIRT teams make a distinction between the handling of security incidents and the handling of lesser forms of network abuse such as spam and copyright infringement. This can be used to mark incidents as belonging to different functions of the team.
- Classification – **MANDATORY** (no default). Closest match to type of incident for SLA reporting (see 3.11.1).
- Message (necessary but not enforced by RTIR). Brief evidence for the incident, which will form the initial basis for the message when an Investigation is launched. Likely to include log information. Default content is copied from the Incident Report if there is one, but it is often a good idea to remove extraneous or excessive content.
- Priority defaults to 50 which is normally acceptable. Higher values will set shorter alert times.

Clicking the *Create* button will create the Incident and if appropriate link the Incident Report to it. If any of the mandatory fields aren't populated then an error will be reported. You can also launch a new investigation direct from the Incident Creation page, but you should only do this if you are completely confident that there are no existing incidents or investigations related to the same issue. In practise this feature is rarely used and is not recommended for use until you are more experienced with the workflow.

The latest variant of RTIR allows you to launch the investigations from the Incident Creation page. However, it is recommended only for use if you are 100% sure that there is no other incident open or in existence related to the same issue.

3.10.1 Incident priority and classification

Incident priority is not necessarily an indication of urgency, severity or impact. Where for any reason it seems that a response to a message or Investigation is needed quickly, you should consider backing up the RTIR e-mail with a telephone call and at that stage there is no benefit from changing the priority from its default value. If you do increase the priority, the effect will be that Investigation and Incident Report correspondence becomes due sooner; but editing the Due Date in a ticket has a similar effect.

The values available for the Classification field match those required by the current SLA and may be adjusted from time to time. [SLA] gives explanatory text for each possible value, but in many cases assignment is arbitrary.

3.11 Linking to an existing Incident

If the Incident Report involves an issue which is already being dealt with, the Incident Report should be linked to the existing Incident. In the Incident Report view screen, click on the *Link* link in the *Basics* section, select the Incident from the search list then shown and press *Link*; or in the Lookup Tool click *Link* next to the Incident you wish to link to. Regardless of who owns this Incident, it is the responsibility of the team member linking in the new Incident Report either to answer this report promptly, checking whether it affects the rest of the ongoing Incident, or to ensure that it is brought to the attention of another team member who does so. They may also wish to consult with the Incident's owner (if they are available) before responding or otherwise becoming involved in work on the Incident.

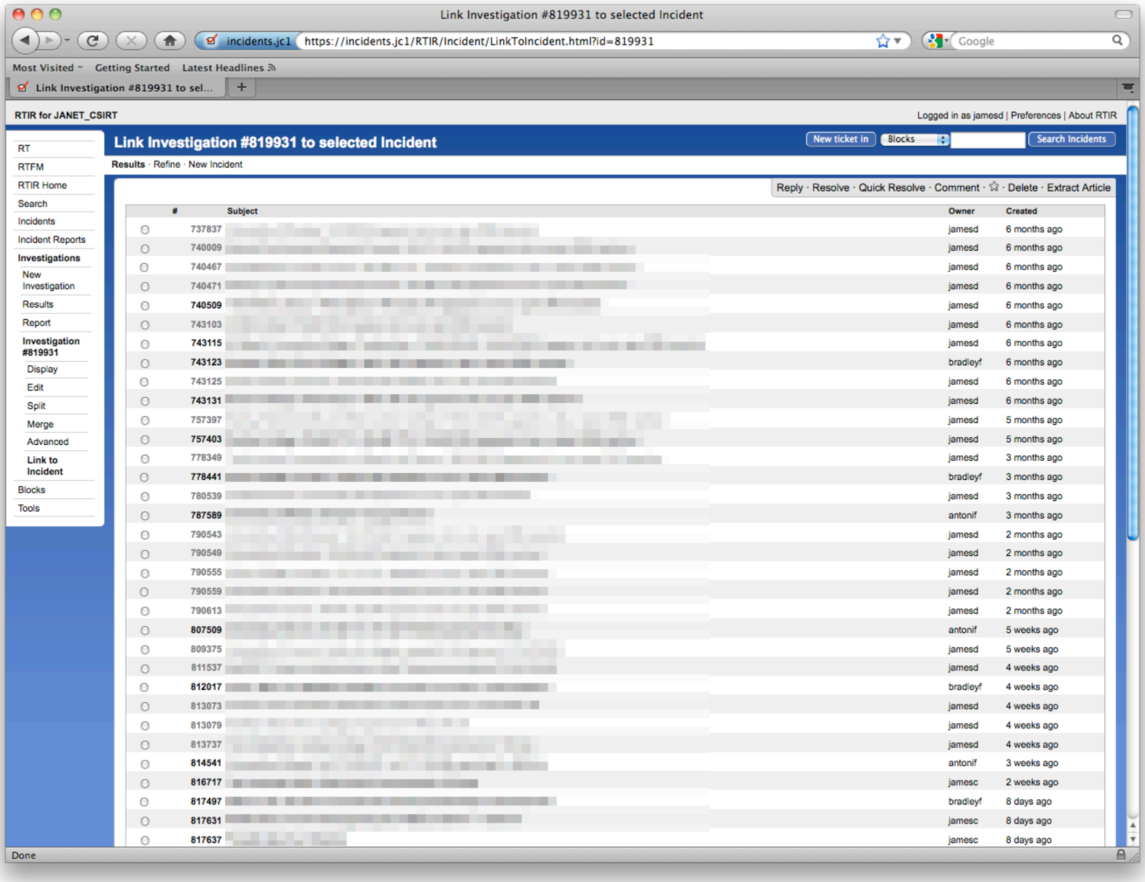

Link Investigation #819931 to selected Incident
incidents.jc1
https://incidents.jc1/RTIR/Incident/LinkToIncident.html?id=819931
Google
Most Visited
Getting Started
Latest Headlines
Link Investigation #819931 to sel...
RTIR for JANET_CSIRT
Logged in as jamesd | Preferences | About RTIR
RT
RTFM
RTIR Home
Search
Incidents
Incident Reports
Investigations
New Investigation
Results
Report
Investigation #819931
Display
Edit
Split
Merge
Advanced
Link to Incident
Blocks
Tools
Link Investigation #819931 to selected Incident
New ticket in
Blocks
Search Incidents
Results · Refine · New Incident
Reply · Resolve · Quick Resolve · Comment · ☆ · Delete · Extract Article
Done

	#	Subject	Owner	Created
○	737837	[redacted]	jamesd	6 months ago
○	740009	[redacted]	jamesd	6 months ago
○	740467	[redacted]	jamesd	6 months ago
○	740471	[redacted]	jamesd	6 months ago
○	740509	[redacted]	jamesd	6 months ago
○	743103	[redacted]	jamesd	6 months ago
○	743115	[redacted]	jamesd	6 months ago
○	743123	[redacted]	bradleyf	6 months ago
○	743125	[redacted]	jamesd	6 months ago
○	743131	[redacted]	jamesd	6 months ago
○	757397	[redacted]	jamesd	5 months ago
○	757403	[redacted]	jamesd	5 months ago
○	778349	[redacted]	jamesd	3 months ago
○	778441	[redacted]	bradleyf	3 months ago
○	780539	[redacted]	jamesd	3 months ago
○	787589	[redacted]	antonif	3 months ago
○	790543	[redacted]	jamesd	2 months ago
○	790549	[redacted]	jamesd	2 months ago
○	790555	[redacted]	jamesd	2 months ago
○	790559	[redacted]	jamesd	2 months ago
○	790613	[redacted]	jamesd	2 months ago
○	807509	[redacted]	antonif	5 weeks ago
○	809375	[redacted]	jamesd	5 weeks ago
○	811537	[redacted]	jamesd	4 weeks ago
○	812017	[redacted]	bradleyf	4 weeks ago
○	813073	[redacted]	jamesd	4 weeks ago
○	813079	[redacted]	jamesd	4 weeks ago
○	813737	[redacted]	jamesd	4 weeks ago
○	814541	[redacted]	antonif	3 weeks ago
○	816717	[redacted]	jamesc	2 weeks ago
○	817497	[redacted]	bradleyf	8 days ago
○	817631	[redacted]	jamesc	8 days ago
○	817637	[redacted]	jamesc	8 days ago

Figure 5 Linking an incident report to an existing incident

It is possible to link an Incident Report to an Incident that has been resolved by separately searching for the incident through the *Search* page, identifying and displaying the Incident on screen, and then re-opening the incident via the secondary menu bar.

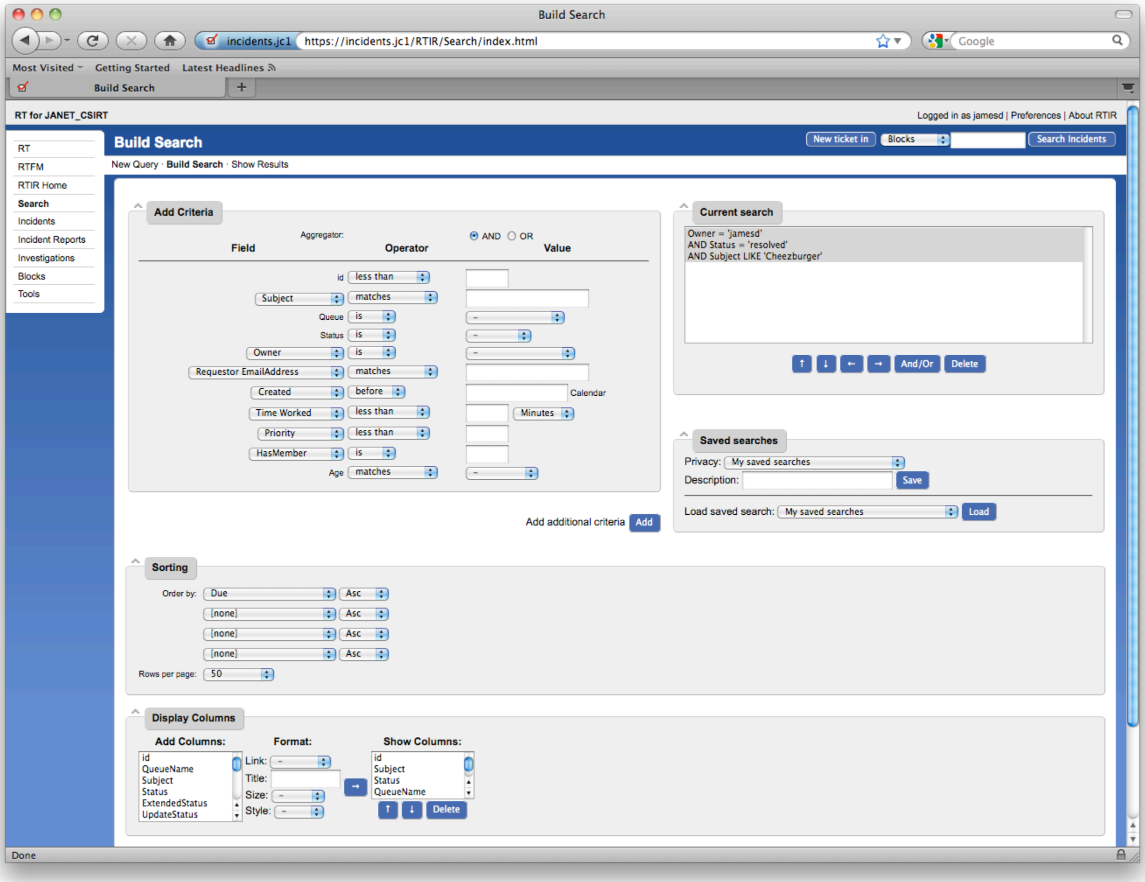

Owner = 'jamesd'
AND Status = 'resolved'
AND Subject LIKE 'Cheezburger'

Figure 6 Incident search screen

It is also possible to unlink an Incident Report from an Incident. From within the Incident Report, an *Unlink* option exists next to the entry for each Incident that the Incident Report is currently attached to.

3.12 Replying to the Incident Report

Once an Incident Report is linked to an Incident the next step should be to reply to it, or if no reply is appropriate (for instance, in the case of an automated report from your own monitoring systems) to resolve it.

SLA compliance can be measured by response times to Incident Reports. The software allows for the configuration of SLA timers which are customized by the RTIR administrator. This should be configured according to existing CSIRT procedures and policies. For more information on how your SLA may affect your use of RTIR, please see the documentation on the SLA.

Before replying make sure the *Requestors* (the senders of the incident report) are correct and do not include any unwanted e-mail address, particularly people CC'd on complaints. Requestors can be removed, added via the *Edit* tag in the side panel.

To reply to an Incident Report, first show it on the screen, then click on *Reply*. *Reply* within the body includes the message in the reply; *Reply* on the top title bar starts the reply with an empty message. Replying takes the form of a simple email. Stock replies can be added with the *Include RTFM article* option and then edited. RTFM articles can be created by navigating from *RTFM*, via Articles and selecting *New Article*, but the use of RTFM is outside of the scope of this document. If you do not require or expect any response from this correspondent you can then resolve the Incident Report immediately.

Replying to an Incident Report sets the Due date to a discrete number of days from the date of reply. The exact number of days can be configured by the RTIR administrator and should be set according to SLA policy and documentation. The Due date can also be adjusted in the *Edit* screen.

Some Incident Reports, such as automated complaints, may not require a reply. They should be resolved promptly to avoid any impact on response times in an SLA but the subsequent response will be entirely through Investigations.

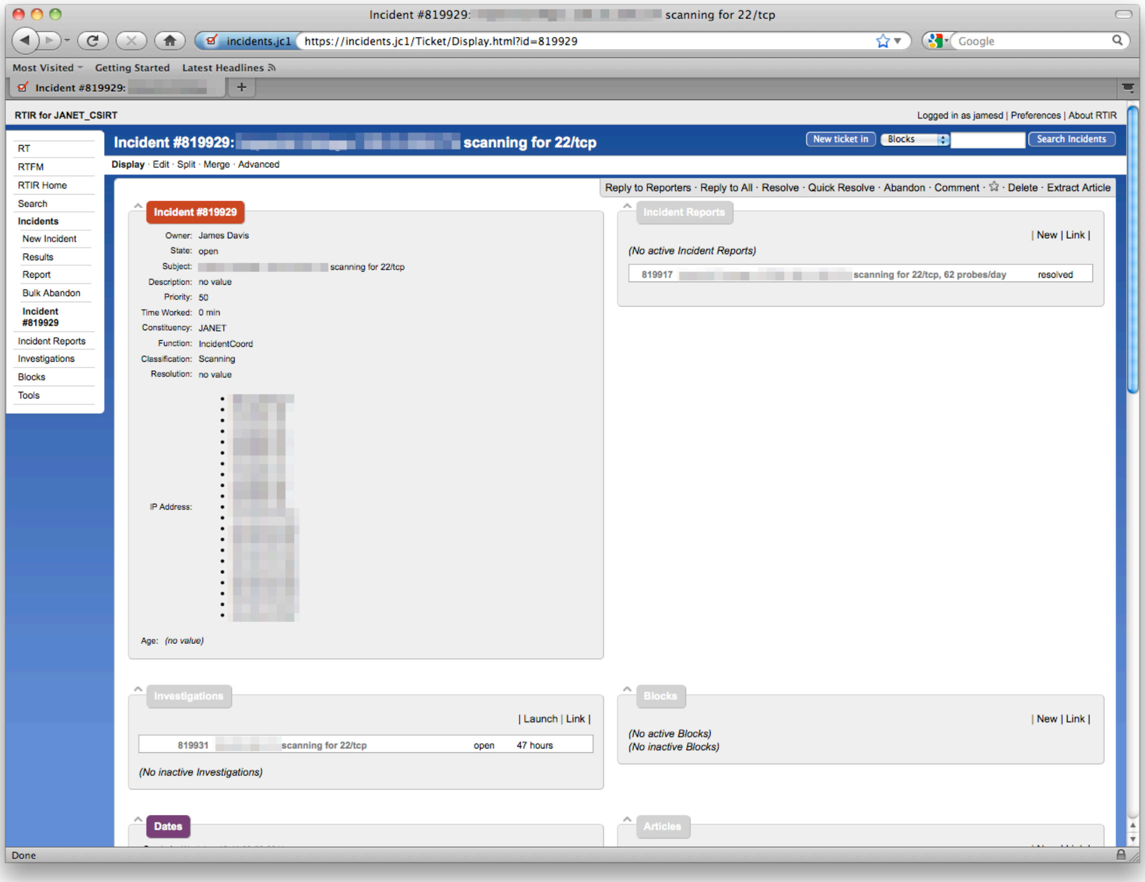

Incident #819929: [redacted] scanning for 22/tcp
incidents.jc1
https://incidents.jc1/Ticket/Display.html?id=819929
Google
Most Visited
Getting Started
Latest Headlines
Incident #819929: [redacted]
RTIR for JANET_CSIRT
Logged in as jamesd | Preferences | About RTIR
RT
RTFM
RTIR Home
Search
Incidents
New Incident
Results
Report
Bulk Abandon
Incident #819929
Incident Reports
Investigations
Blocks
Tools
New ticket in
Blocks
Search Incidents
Display · Edit · Split · Merge · Advanced
Reply to Reporters · Reply to All · Resolve · Quick Resolve · Abandon · Comment · ☆ · Delete · Extract Article
New	Link
Launch	Link
New	Link
Dates
Articles
Done

Incident #819929: [redacted] scanning for 22/tcp

Incident #819929

Owner: James Davis
State: open
Subject: [redacted] scanning for 22/tcp
Description: no value
Priority: 50
Time Worked: 0 min
Constituency: JANET
Function: IncidentCoord
Classification: Scanning
Resolution: no value
IP Address: [redacted]

Age: *(no value)*

Incident Reports

(No active Incident Reports)

819917	[redacted] scanning for 22/tcp, 62 probes/day		resolved

Investigations

819931	[redacted] scanning for 22/tcp	open	47 hours

(No inactive Investigations)

Blocks

(No active Blocks)
(No inactive Blocks)

Figure 7 An existing incident with incident report and investigation

3.13 Launching an Investigation

Once the Incident Report has been replied to, further investigative work is usually required. This will often be to a third party or customer organization, in response to a complaint received by CSIRT. Investigations can be launched from the Incident view, or from the *Lookup* tool; either route reaches the same screen. An Investigation launched from the *Lookup* tool can have the requestor field pre-populated with a single e-mail address from lookup results page - use *[Investigate To]*; alternatively, multiple email addresses can be selected via the check boxes on the lookup results page and the Investigation launched with the Investigate button. To initiate a lookup for a specific hostname or IP address, click on an active link in the history view of the Incident or its child Incident Reports.

By default a launched Investigation will take its Subject from the parent incident, and the initial message text from the first message of the parent incident (quoted with '>'). This is intended to make it easier to forward complaints to customer organizations. The subject line and message text will almost always need to be changed to make more sense to the customer receiving the message. In particular, do not depend too much on standard messages; say quite specifically and explicitly what you expect them to do (e.g. 'find the computer responsible, make it safe and let us know' or 'please confirm that you have taken action to apply your Acceptable Use policy to the user concerned').

In many cases it is appropriate to telephone the security contact as soon as you have sent the e-mail message, to ensure that they will find the message and are aware they will need to take some action in a timely manner. It may also turn out that contact details are defective, perhaps causing an e-mail to bounce, in which case this is the appropriate time to correct them and make sure that your database of contact details is correct. You may wish to link any bounces to the Incident as Incident Reports, or you may wish simply to *Quick Reject* them.

The default installation of RTIR does not automatically include references inserted into the subject line by a correspondent's ticketing system. This can cause multiple tickets to be opened on their ticketing system when the correspondence should have been added to an existing ticket. You can manually adjust the subject of your investigation to their required format through the *Edit* page. Extensions to automatically perform this action exist (see RT::Extension::ExtractSubjectTagOnTransaction on CPAN) or you can write your own customisations.

3.14 Initiating a block

It is appropriate to block traffic within your own network immediately if it is:

- a threat to the network from outside; or
- a threat to other networks from a customer organization.

If at all possible, help the customer organization which is the source of a problem to resolve it promptly within their own network. It is still sometimes useful to apply a block within the network to give the organization time to make their arrangements.

Every reasonable effort should be made to inform the organization or organizations directly affected by the block. However, your policies will normally dictate that the removal of a threat to the network takes precedence over the convenience of an individual organization.

To arrange for blocking, contact your Network Operations Centre (NOC) with details of the block that should be placed. You should normally ask for the action to be taken as close as possible to the source or destination of the unwanted traffic.

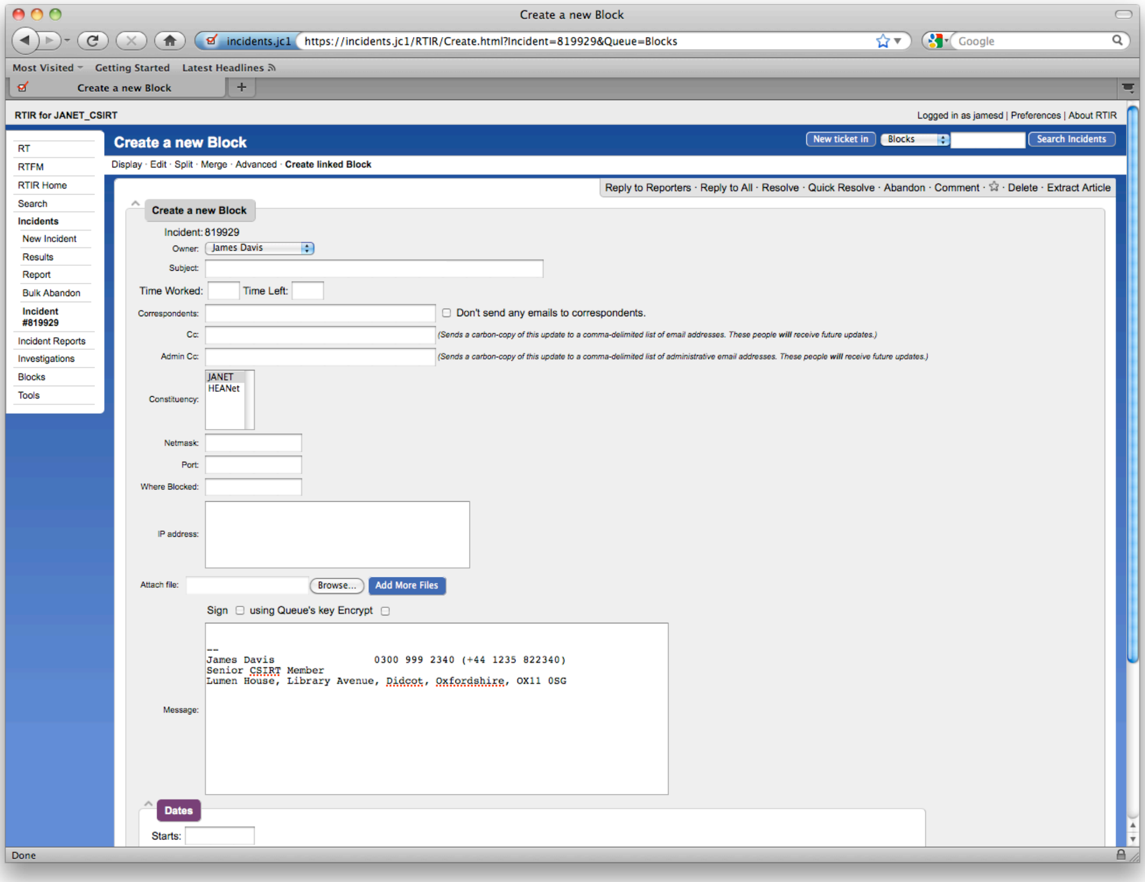

Create a new Block

Incident: 819929

Owner: James Davis

Subject:

Time Worked: Time Left:

Correspondents: □ Don't send any emails to correspondents.

Cc: *(Sends a carbon-copy of this update to a comma-delimited list of email addresses. These people **will** receive future updates.)*

Admin Cc: *(Sends a carbon-copy of this update to a comma-delimited list of administrative email addresses. These people **will** receive future updates.)*

Constituency: JANET
HEANet

Netmask:

Port:

Where Blocked:

IP address:

Attach file:

Sign □ using Queue's key Encrypt □

Message:

```
--
James Davis                 0300 999 2340 (+44 1235 822340)
Senior CSIRT Member
Lumen House, Library Avenue, Didcot, Oxfordshire, OX11 0SG
```

Starts:

Figure 8 Creating a new block

Create a child ticket of type *Block* from the Incident view. It will be in the ‘block pending’ state. The precise content and format of your message will depend on your documented procedures but will typically include information such as the IP address(es) to be blocked, particular services that are to be blocked, where the traffic is to be blocked, details of the affected customer, and the duration that the block is to be applied for. In urgent cases your procedure may be first to telephone your NOC and use the request sent by RTIR as confirmation and documentation of actions taken.

The block ticket will accumulate replies from the NOC. Set the Due date as appropriate.

Network blocks should be removed when the threat is no longer current, typically when the source organization has taken action to remove it. See 3.17 on removing a block. In many cases the communications involved in placing a block are highly regular and you can have RTIR change the status of a block automatically by configuring $RTIR_BlockAproveActionRegexp in RTIR_Config.pm.

3.15 Ongoing Incident handling

An Incident is ongoing once the Incident has been created; the Incident Reports replied to and the investigations launched; the team is awaiting action from an external party; all child Due dates are in the future; and this is reflected by the Due date of the parent Incident (which is the earliest child Due date). It will join the list of Incidents sorted by Due date listed on the RTIR homepage, owned by the team member who originally created it.

Further action will be required when either the Due date of a child ticket is reached or a third party replies to a child ticket (which sets the Due date to ‘now’). The parent Incident will then also have reached its Due date; it will appear at or near the top of the sorted list on the RTIR front page, and the Due date text will change from black to red.

Running time is measured from the ‘Started’ time which is reset each time the Incident Report is replied to.

The RTIR front page (RTIR Home) has two Incident lists, located beneath the unlinked Incident Reports list. The first Incident list shows Incidents owned by the currently logged-in team member. The second list is a list of all open Incidents, sorted by due date.

The aim for each Incident is to resolve all issues, complete all correspondence and ‘resolve’ the Incident as soon as practicable. No Incidents should appear in red (requiring action or follow up by a CSIRT member) for long periods, or exceed their Due dates by more than is permitted by your organization’s SLA.

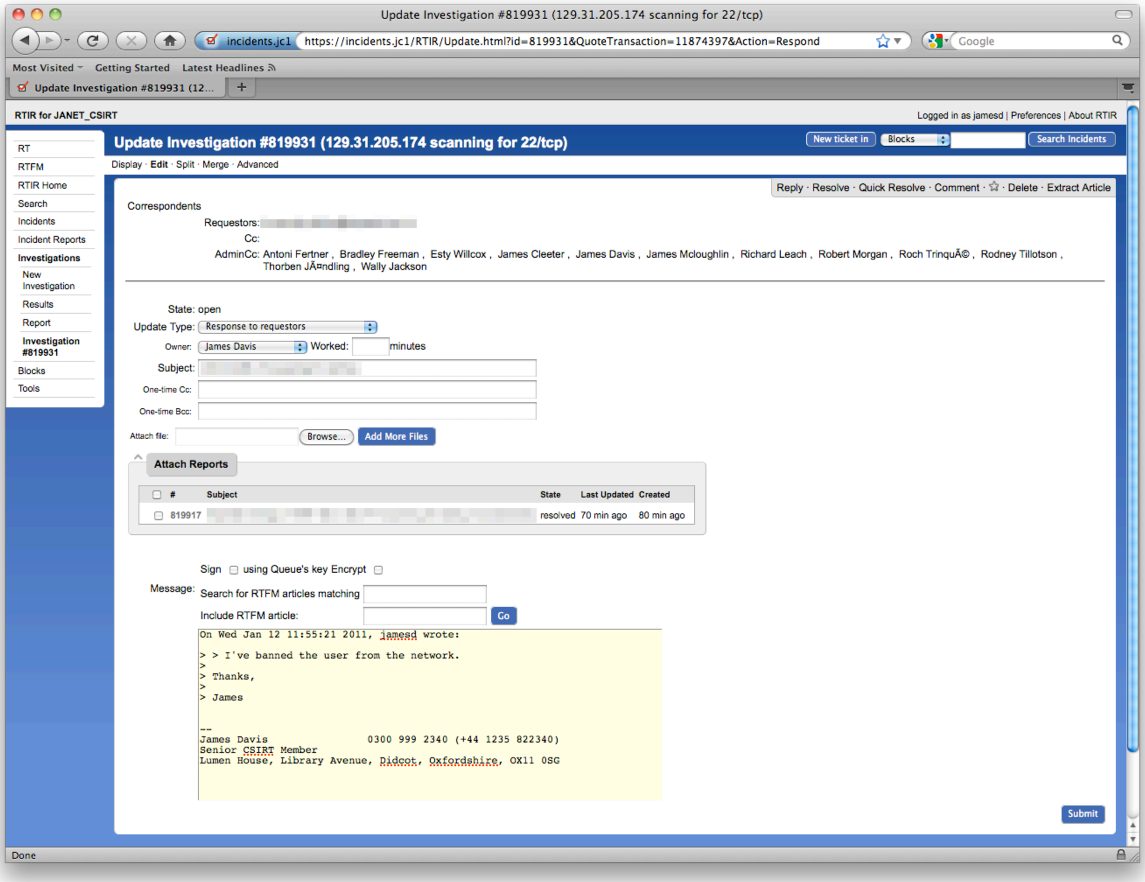
Update Investigation #819931 (129.31.205.174 scanning for 22/tcp)
incidents.jc1
https://incidents.jc1/RTIR/Update.html?id=819931&QuoteTransaction=11874397&Action=Respond
Google
Most Visited
Getting Started
Latest Headlines
Update Investigation #819931 (12...
RTIR for JANET_CSIRT
Logged in as jamesd | Preferences | About RTIR
New ticket in
Blocks
Search Incidents
RT
RTFM
RTIR Home
Search
Incidents
Incident Reports
Investigations
New Investigation
Results
Report
Investigation #819931
Blocks
Tools
Display · Edit · Split · Merge · Advanced
Reply · Resolve · Quick Resolve · Comment · ☆ · Delete · Extract Article
Browse...
Add More Files
Attach Reports
Go
Submit
Done

Update Investigation #819931 (129.31.205.174 scanning for 22/tcp)

Correspondents

Requestors: [redacted]

Cc:

AdminCc: Antoni Fertner , Bradley Freeman , Esty Willcox , James Cleeter , James Davis , James Mcloughlin , Richard Leach , Robert Morgan , Roch TrinquÃ© , Rodney Tillotson , Thorben JÃ¤ndling , Wally Jackson

State: open

Update Type: Response to requestors

Owner: James Davis Worked: minutes

Subject: [redacted]

One-time Cc:

One-time Bcc:

Attach file:

□	#	Subject	State	Last Updated	Created
□	819917	[redacted]	resolved	70 min ago	80 min ago

Sign □ using Queue's key Encrypt □

Message: Search for RTFM articles matching

Include RTFM article:

```
On Wed Jan 12 11:55:21 2011, jamesd wrote:

> > I've banned the user from the network.
>
> Thanks,
>
> James


--
James Davis                 0300 999 2340 (+44 1235 822340)
Senior CSIRT Member
Lumen House, Library Avenue, Didcot, Oxfordshire, OX11 0SG
```

Figure 9 Replying to an open investigation

Sometimes an e-mail Investigation can be left to time-out once, but after a period determined by the SLA and CSIRT procedures the follow-up should be by a telephone call. In some cases it is appropriate to telephone sooner than that; it may also be appropriate to re-send the original e-mail and possibly to include the Management contact as a copy recipient. Our experience is that a great many Incidents remain ‘open’ for weeks awaiting a meaningful response from correspondents, and that merely sending repeated e-mail reminders whenever RTIR flags that a reply is overdue achieves very little.

Record every telephone call made or received as a *Comment* in the appropriate ticket. If they are not obvious from the context, include the telephone number and name of the person you spoke to. If you were only able to leave a voice mail message, make sure you at least give your name, mention your CSIRT, and give your team’s telephone number and the RTIR ticket number. In your note of the call, state what message you left (e.g. “Asked her to call about this ticket”).

As is good practice for ordinary e-mail, avoid quoting the whole of the correspondence in a reply, and try to avoid multiple levels of quoting. Edit the text carefully to quote the previous message selectively, showing only the one or two lines to which you are responding or to which you want a response. Do not 'top post' by placing only a short reply followed by a quoting of the correspondent's entire previous message.

3.16 Resolving Incidents

When an Incident requires no further action it can be closed. Children of Incidents (Incident Reports, Investigations and Blocks) can be individually closed during the lifecycle of an Incident once each has run its course. When all children have reached this point (for example when a site reports the machine has been secured) we can close the Incident and all of its children in one action using the *Resolve* action in the Incident. This will close all children still open, and the Incident itself. (This no longer prompts for a comment to be added to each child ticket, unlike earlier behaviour.) If it is not obvious why the Incident is being closed then an explanatory comment should be added at this point.

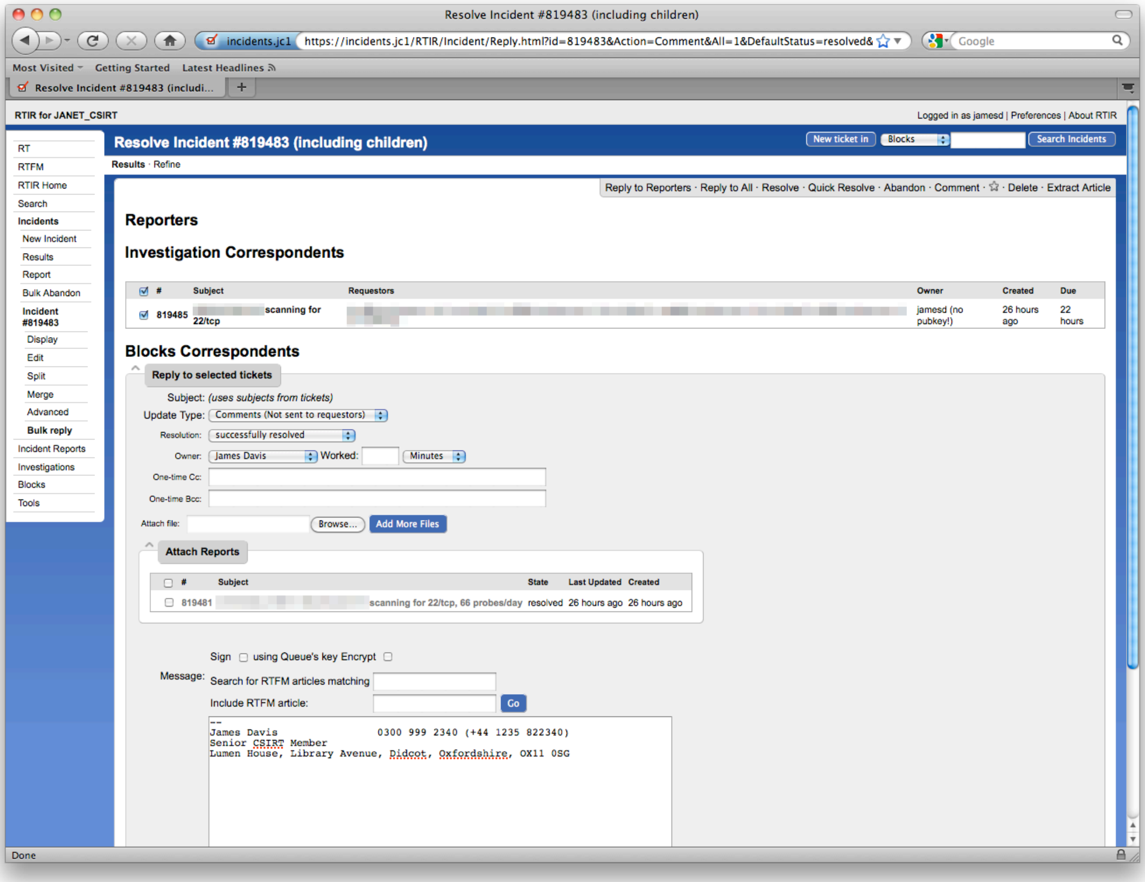
Resolve Incident #819483 (including children)
incidents.jc1
https://incidents.jc1/RTIR/Incident/Reply.html?id=819483&Action=Comment&All=1&DefaultStatus=resolved&
Google
Most Visited
Getting Started
Latest Headlines
Resolve Incident #819483 (includi...
RTIR for JANET_CSIRT
Logged in as jamesd | Preferences | About RTIR
New ticket in
Blocks
Search Incidents
RT
RTFM
RTIR Home
Search
Incidents
New Incident
Results
Report
Bulk Abandon
Incident #819483
Display
Edit
Split
Merge
Advanced
Bulk reply
Incident Reports
Investigations
Blocks
Tools
Results · Refine
Reply to Reporters · Reply to All · Resolve · Quick Resolve · Abandon · Comment · ☆ · Delete · Extract Article
Done

Resolve Incident #819483 (including children)

Reporters

Investigation Correspondents

☑	#	Subject	Requestors	Owner	Created	Due
☑	819485	[redacted] scanning for 22/tcp	[redacted]	jamesd (no pubkey!)	26 hours ago	22 hours

Blocks Correspondents

Reply to selected tickets

Subject: *(uses subjects from tickets)*
Update Type: Comments (Not sent to requestors)
Resolution: successfully resolved
Owner: James Davis Worked: Minutes
One-time Cc:
One-time Bcc:
Attach file: Browse... Add More Files

Attach Reports

□	#	Subject	State	Last Updated	Created
□	819481	[redacted] scanning for 22/tcp, 66 probes/day	resolved	26 hours ago	26 hours ago

Sign □ using Queue's key Encrypt □
Message: Search for RTFM articles matching
Include RTFM article: Go

```
--
James Davis                0300 999 2340 (+44 1235 822340)
Senior CSIRT Member
Lumen House, Library Avenue, Didcot, Oxfordshire, OX11 0SG
```

Figure 10 Resolving an incident

The 'Resolution' custom field is normally left as 'successfully resolved' but in some cases you may wish to set a different state which can later be reported upon.

In some cases, where actions are already documented, you may not wish to add any explanatory comments as to why an Incident, Investigation or Incident report is being resolved. In these cases you can use the *Quick Resolve* incident. This automatically resolves a ticket (and in the case of an Incident, its children) without giving any opportunity to set the Resolution custom field or add any explanatory comments.

(An alternative action is *Abandon* which was intended to indicate that CSIRT has given up on an Incident, perhaps because the Incident was created erroneously. Unfortunately abandoning an Incident will *Reject* its children so that they will no longer be displayed by default. Its use should be avoided until this behaviour is improved.)

If later an external party replies to a child ticket which has already been closed it will reopen itself and the parent incident. There is therefore no harm in closing an Incident which we believe is finished, but for which there is some chance of a follow up e-mail from a third party in the future (for example a site may send a write-up of the event or give details of malware).

If, instead, subsequent Incident Reports seem to indicate that the address remains unsafe, follow up the Incident in the usual way either by explicitly re-opening one of its child tickets or by replying to the earlier messages in the Investigation.

An Incident all of whose child tickets are ‘resolved’ will have no Due Date and will appear at the top of the ‘Most due Incidents’ list. Very occasionally it is useful to leave an Incident in this condition for reference, but normally it should be resolved.

3.17 Removing blocks

A special case when resolving an Incident is to make sure any associated blocks are removed. CSIRTs do not usually maintain long-term blocks; they only last as long as the incident. This is different to the other procedures for resolving incidents because it involves sending a request to the blocker (typically the NOC) to remove the block and waiting for confirmation before closing.

The stages are:

- Request removal (the corresponding RTIR button sends a message to the Service Desk and the Operations Team);
- Wait for confirmation (normally a message in response);
- Edit the state of the block ticket to ‘removed’ (this can be done automatically by configuring $RTIR_BlockAproveActionRegexp).

Block tickets in certain states can be set to ‘removed’ by resolving their parent Incident.

3.18 Dealing with heavy load

At certain periods (for example, on a Monday morning or during new worm outbreaks) the quantity of Incident Reports or ongoing Incidents awaiting action may be too much for a single team member to deal with. RTIR allows more than one person to work in parallel, although this requires some coordination. In general, it is best if only one person at a time works on any one ticket.

The working patterns needed to deal with such a situation will greatly depend on the type of heavy load situation. Simple coordination can be achieved by roughly splitting the work based on its topic (one team member deals with internally generated reports, another with spam reports, another everything else). More complex situations will need to be coordinated on an ad hoc basis, typically by the CSIRT Manager.

4 PERFORMANCE AND REPORTS

A Service Level Agreement specifies that initial response is to be made within certain periods. In addition, a CSIRT may seek to complete handling of routine incidents and the corresponding RTIR tickets within elapsed times given elsewhere which may depend on external events and influences.

Procedures for preparing the reports are outside the scope of this document, although RTIR has facilities in the *Tools* page for producing a simple monthly report including almost all the statistics and data required.

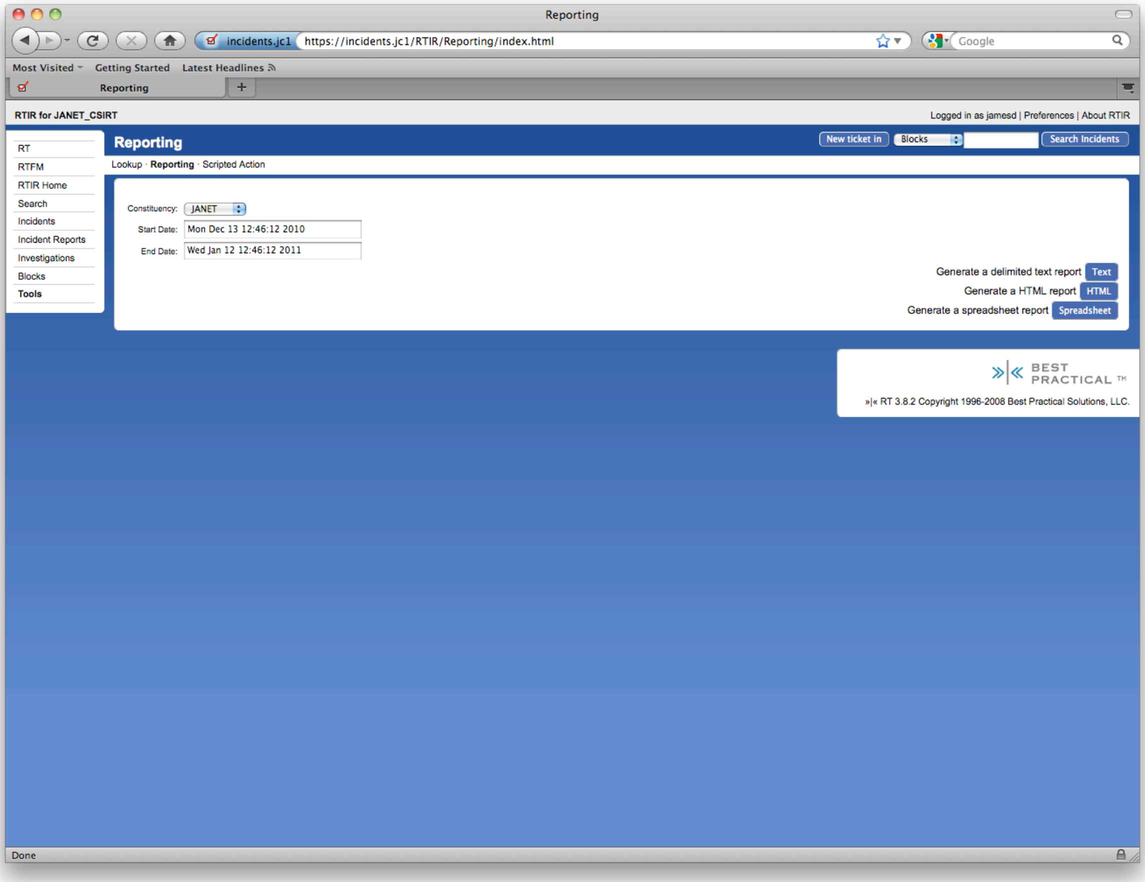

Constituency: JANET

Start Date: Mon Dec 13 12:46:12 2010

End Date: Wed Jan 12 12:46:12 2011

Generate a delimited text report

Generate a HTML report

Generate a spreadsheet report

»|« RT 3.8.2 Copyright 1996-2008 Best Practical Solutions, LLC.

Figure 11 The reporting tool

5 RECORDS

5.1 Record keeping, integrity and access

Most incident records exist only in electronic form and the RTIR database is the only repository. Records are the complete history of an Incident as constructed by RTIR including the correspondence in all its child tickets and the notes of action taken by team members and of other updates. Every item has a time stamp.

Access to incident records is through RTIR itself.

RTIR provides no facility for altering the historical record of the status of any ticket except by adding to it. More direct access to the underlying database (which might make it possible to alter existing data) is restricted by the configuration of the database servers.

5.2 Retention and availability

All the data should be backed up in accordance with policies.

Certain data which enters RTIR is not relevant to the work of the team, such as spam, or of only transient interest such as bulk lists of IP addresses sent for information by automatons. The resulting Incident Reports will all be left in the ‘Rejected’ state in accordance with the instructions in this document and separate retention rules apply to them. After a relatively short time they will be completely removed from the RTIR database and although they will be present on some backup media there is no automatic procedure for retrieving them.

The data relating to Incidents may also be removed from RTIR after some period, with an intermediate period during which statistical information about incidents may be available but not the complete correspondence.

Details of the processes and procedures for the ageing and removal of both wanted and unwanted material are still under discussion. At the time of writing a new version of RTIR is expected to be available soon with processes and procedures for the ageing and removal of both classes of material.