ATLASSIAN

Summary of internal security incidents

July 2021 - June 2022

Table of contents

3 **Introduction**

4 **Recent developments in our incident response program**

5 **More information about Atlassian's incident response processes**

6 Incidents by type

9 Incidents by severity

10 Summary of incidents by severity

13 Method of detection

14 Time to incident detection

16 Time to incident containment

This paper provides a summary of the security incidents at Atlassian over the last 12 months (July 2021-June 2022). The purpose of this summary is to provide additional transparency – via a range of metrics - around the security incidents we have experienced, and the performance of our incident management processes. Atlassian maintains a significant amount of both qualitative and quantitative data internally for the purpose of tracking security incidents from the time they are identified through to resolution – this paper summarizes much of that data in aggregated form. Our intention is to publish updated incident data for each financial year moving forward, consistent with our core company values and in particular, continuing to be an open company with no bullshit.

[1]Atlassian catalogues its security incidents according to the Verizon VERIS framework

Recent developments in our incident response program

Like any company with a good security program, we don't rest on our laurels when it comes to our approach to handling incidents. As part of our ongoing efforts to evolve and improve our security incident management processes, there have been some notable developments in the last 12 months that we also want to highlight:

1. We've implemented enhancements to our forensic analysis capabilities by introducing our internally developed Cloud Forensics Kit (CFK). The CFK is a bundle of server-less automations that do parts of forensic analysis within AWS Cloud, plus a command-line interface tool for our analysts to create standard forensic instances that can be used during the investigation and collection of forensic evidence. The introduction of the CFK will also help us to contain incidents more effectively and reduce the average time it takes us to contain an incident.

2. Our Security Intelligence team leveraged a new Security Orchestration and Response (SOAR) platform to assist with investigations of incidents and to help with managing tasks during an incident. This included:

    - Automated response and tracking of phishing incident reports;
    - Automated enrichment of incident alerts so that our team has as much information available to them as possible about an incident; and
    - Automated containment of hosts involved in an incident with an Endpoint Detection and Response (EDR) solution.

While not a recent development, it is important to note we have a detailed post incident review process involving multiple stakeholders across Atlassian, and we continue to conduct these reviews across all incidents in order to identify any preventative measures we can implement in order to reduce the likelihood of a re-occurrence of similar incidents in future.

More information about Atlassian's incident response processes

While protective security measures continue to be a crucial cornerstone of our security program, we also know Atlassian's ability to detect and respond to security incidents effectively is just as important to deal with the threat environment of both now and the future. While this paper does not discuss the specifics of our approach to incident management, you can find detailed information about this on our **Trust Center**:

- **Our approach to managing security incidents**
- **Our security detections program**

What we mean by a security incident

Consistent with the definition of security incident we have adopted for our incident management processes, this paper's focus is on incidents where there was an existing or impending potential negative impact to the confidentiality, integrity or availability of our customers' data, Atlassian's data, or Atlassian's services.

An important caveat is that where a potential incident was identified, but subsequent investigations indicated the incident occurred a result of customer error, it is not included in the data we have provided in this paper.

Scope of this report

This report details Security Incidents identified and investigated in Atlassian Cloud Platform, Atlassian Corporate Infrastructure, and endpoints and mobile devices during the time frame July 2021 – June 2022, which is Atlassian Fiscal Year (FY) 22.

This report does not include information about Security Incidents at customers who are hosting Server or Data Center versions of our products.

Total incidents

In FY 22, Atlassian had **34** confirmed security incidents. This is a significant 60% reduction as compared to FY 21, where **84** incidents were identified.

Incidents by type

In the graph below, we break down the 34 incidents for FY 22 according to type. The numerical data is also provided in the table under the graph. Just over 22% of incidents during the financial year were classified as malware infections. The next highest category were related to incidents involving a product vulnerability (6 incidents, or 17%).

Incidents by type (FY 22)

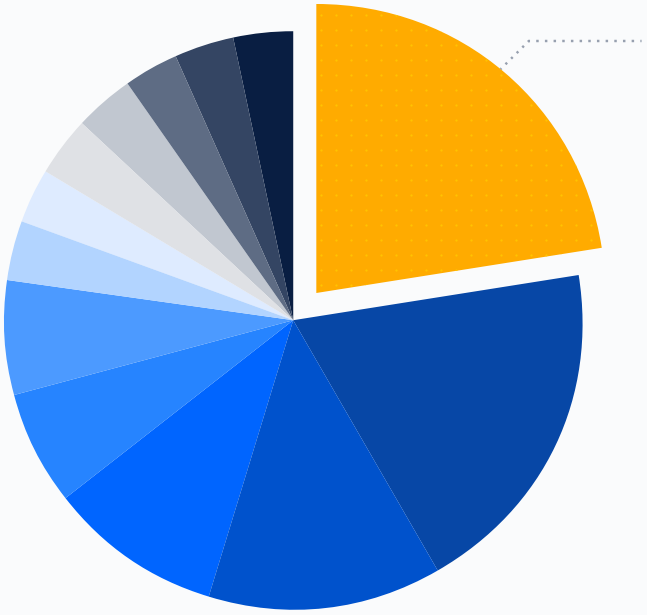

Just over 22% of incidents during the financial year were classified as malware infections

- Malware infection
- False positives
- Compromised account
- Data exposure
- Human error
- Social engineering
- Lost or stolen asset
- Bug
- Platform abuse
- Non-product related vulnerablility
- Third party incident
- Employee relations
- Product vulnerability

Incident Type	Number of incidents (FY 22)
Malware infection	7
Product vulnerability	6
Data exposure	4
Platform abuse	4
Compromised account	2
Human error	2
Non-product related vulnerability	2
Bug	1
Employee relations	1
Lost or stolen asset	1
False positives	1
Third party incident	1
Social engineering	1
Other	0

For clarity, we include the definitions for our various incident types below:

- **Malware infection** – the root cause of the incident was malware being downloaded onto an Atlassian IT asset. For FY 22, all malware related incidents occurred in our workstation fleet – other environments (e.g. our production environments) were not affected;

- **Product vulnerability** - an incident where the root cause was an exploitation of a vulnerability in an Atlassian product;
- **Data exposure** - root cause of an incident was some type of data being exposed accidentally (e.g. data was published to an online service, inadvertently sent to logging services, or accidentally shared with a third party);
- **Platform abuse** - involves an incident where the root cause was our products being used in a way that violated Atlassian's terms of service;
- **Human error** – incidents where the root cause is predominantly attributable to a human error of some kind;
- **Non-product related vulnerability** - an incident caused by exploitation of a vulnerability in software that is not an Atlassian product (e.g. operating systems or platform tools);
- **Bug** - instances where the root cause of an incident was a bug introduced into code that wasn't raised as a vulnerability (e.g. access control breaking due to regression);
- **Compromised account** – incidents where an internal Atlassian user account is compromised;
- **Employee relations** – incidents relating to actions that were possible violations of company policies by staff;
- **Lost or stolen asset** – root cause of the incident involved a lost or stolen Atlassian asset;
- **Third Party Incident** – root cause was a third party trusted by Atlassian being compromised; and
- **Social Engineering** – root cause related to social engineering of some kind (e.g. via a Phishing scam).

Incidents by severity

Atlassian designates one of four severity levels to an incident. These are:

Severity level	Description
Severity Level 0 (Highest)	Crisis incident with maximum impact
Severity Level 1	Critical incident with very high impact
Severity Level 2	Major incident with significant impact
Severity Level 3 (Lowest)	Minor incident with low impact

We use a variety of indicators to determine the severity of an incident – these vary depending on the product involved but will include consideration of whether there is a total service outage (and the number of customers affected), whether core functionality is broken, and whether there has been any data loss.

In the table below, we provide the number of incidents for each severity level for FY 22.

Severity level	Number of incidents (FY 22)
Severity Level 0 (Highest)	0
Severity Level 1	4
Severity Level 2	4
Severity Level 3 (Lowest)	26

As can be seen from this data, in FY 22 there were no incidents that were classified by us as Level 0 (the highest level of severity). For FY 22 there were 8 incidents classified as either Critical (Level 1) or Major (Level 2).

Summary of incidents by severity

We have provided a summary of the nature of each of the incidents for the Level 1, Level 2, and Level 3 categories for FY 22 below.

Severity Level 1 incidents

There were four incidents in this category for FY 22:

- One incident involved the investigation of suspicious activity within Atlassian's environment (specifically a jump box associated with accessing the Trello environment), as reported in one of our Splunk alerts. The investigation determined this alert was a false positive, and the incident was closed.
- In a second incident, Atlassian was alerted to suspicious activity within one of its AWS accounts which involved the creation of a suspicious Lambda function by an external threat actor. The function included the downloading of a script and executing it. The incident was in part attributable to a security breach of a trusted third party – a detailed investigation was undertaken by Atlassian that included the reset of two accounts identified as compromised, remediation of API keys for these accounts, and the removal of the suspicious function. Upon investigating, the threat actor was attributed to be our internal Red Team performing a simulated attack.
- The third incident involved the disclosure to Atlassian via its support portal of a critical remote code execution vulnerability in all Confluence Server and Data Center versions. A working exploit has not been confirmed for Confluence Cloud. Investigations are ongoing and a patch has been developed by Atlassian to address the vulnerability.
- The fourth incident was to coordinate response to Confluence Server and Data Center vulnerability CVE-2022-26134, which was actively being exploited on internet-facing Confluence instances, and resulted in **Confluence Security advisory 2022-06-02**.

Severity Level 2 incidents

There were four incidents in this category for FY 22:

- One incident involved an external party (a security researcher) gaining access to internal AWS credentials by exploiting an Apache vulnerability in a specific EC2 instance. Our vulnerability scanners did not detect this instance was subject to this vulnerability. Once identified, the vulnerability was promptly patched and the incident closed when it was confirmed no additional servers were vulnerable.
- In another incident, our Security Intelligence team was alerted via Splunk that some Atlassian-owned EC2 instances were performing crypto-mining. Subsequent investigations revealed that the instances were compromised by an external threat actor who had exploited a vulnerability in Confluence. Atlassian immediately isolated relevant hosts from being internet-accessible, took snapshots of each system's state and obtained memory captures, and stopped compromised instances from continuing to run wherever possible. Where instances could not be stopped, the vulnerability was immediately patched. Most systems that were compromised were test instances, with no customer data. No evidence showed any data was exfiltrated by a threat actor.
- Atlassian was informed of an issue in Trello which meant that in some circumstances, users within a workspace could, without being granted authorization, access private boards in the workspace. A patch was created so that this issue could be addressed.
- Atlassian detected suspicious processes running on a number of hosts in a data center AWS account. The hosts were subsequently shut down and we are continuing investigations into how the hosts were exploited, although it is likely attributable to an unpatched Confluence vulnerability.

Severity 3 incidents

There were twenty-six incidents in this category for FY 22. These incidents are classified as low in severity and given the quantity, we have limited our summary below to some of the more notable in this category:

- One of our staff received a phone call from scammers purporting to be from a telecommunications company support team. A remote desktop session was briefly initiated with the staff member, the user's laptop was re-imaged and their account credentials reset.
- Account credentials for one staff member were leaked in the form of an API token in a GitHub code repository. The leaked tokens were subsequently revoked.
- Two user accounts in Trello were identified as being used to facilitate command and control activity for a specific malware campaign. The accounts were banned and a new detection algorithm was created to monitor for similar activity in future.
- A database containing test data attached to a test Jira instance was discovered by a third party to be publicly accessible from the internet using default credentials. The instance was promptly shut down and remedial actions were taken to avoid the occurrence of similar incidents in future.
- There were two separate incidents involving hardware assets, one relating to a malware compromise of an Atlassian IT asset and another concerning a lost / stolen IT asset which were also addressed as part of our standard incident response processes.

Method of detection

Below is a breakdown of the method of detection for each identified incident for FY 22. Detection methods have been split into three different categories:

1. **Staff** – Incident identified by an Atlassian staff member noticing and reporting something suspicious, or through our Threat Modelling or our Security Detections capabilities
2. **Monitoring** – Incident was identified through our standard detection and monitoring capabilities
3. **Other** – Incident was reported to us via another avenue such as a customer or external vendor

In FY 22, there was a relatively even split between detection methods, with the Staff category and the Monitoring category almost equal (11 and 13), with 7 incidents reported via another avenue.

Incidents by detection method

Incident Detection Method	Number of incidents (FY 22)
Staff	13
Monitoring	13
Other	8

Time to incident detection

In this section, we provide information on how long it took Atlassian to detect incidents for FY 22. This refers to the time it took us to initially identify the occurrence of an incident, but does not include the time it took us to contain the incident. The first graph and table below provide the breakdown for time to detection across all incidents. More than 50% of all incidents for the financial year were detected within 8 hours.

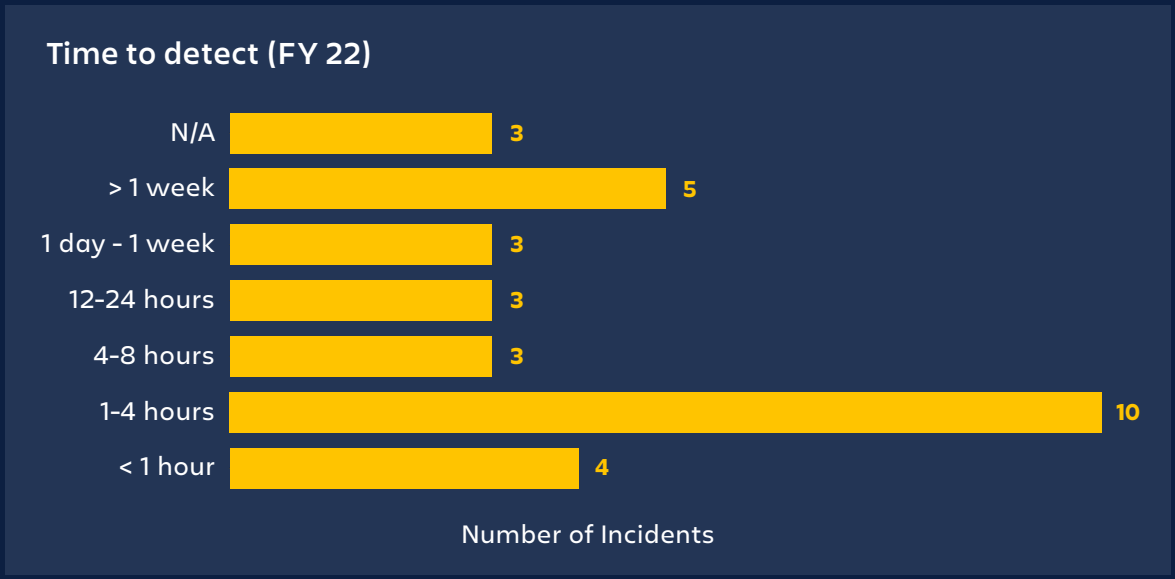

Time to detect	Number of incidents (FY 22)
Under 1 hour	4
1-4 hours	10
4-8 hours	3
12-24 hours	3
1 day - 1 week	3
More than 1 week	5
Not applicable	6

Note: Not applicable refers to incidents that were not detected, and were reported to us via another avenue such as a customer or external vendor.

The second graph and table break this down further by providing the time to detection based on the severity level of the incident (noting that there were no Level 0 incidents for the FY 22 period). For the 5 incidents that took more than 1 week to detect, 4 were of the lowest severity type (Level 3) and one was a Level 2 incident. More than 60% of our incidents were detected within one day.

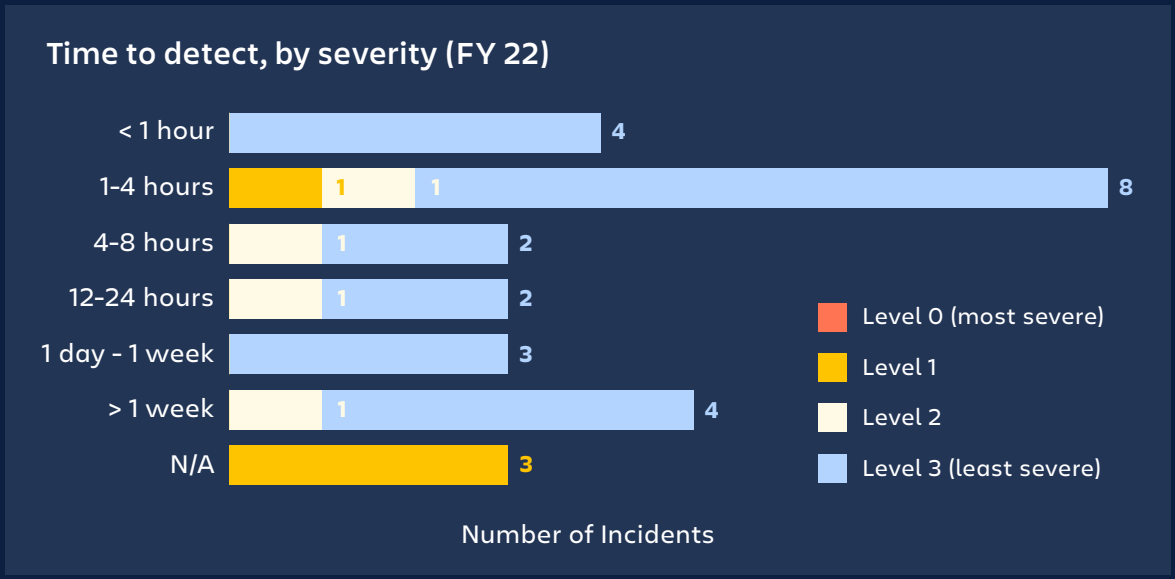

Time to detect	Level 0 (most severe)	Level 1	Level 2	Level 3 (least severe)
Under 1 hour	0	0	0	4
1-4 hours	0	1	1	8
4-8 hours	0	0	1	2
12-24 hours	0	0	1	2
1 day - 1 week	0	0	0	3
More than 1 week	0	0	1	4
Not applicable	0	3	0	0

Note: Not applicable refers to incidents that were not detected, and were reported to us via another avenue such as a customer or external vendor.

Time to incident containment

In this section, we provide information on how long it took Atlassian to contain incidents for FY 22, once they were detected. The containment time refers to the time it took Atlassian to ensure detected incidents no longer presented a security risk.

The first graph and table below provide the breakdown for time to containment across all incidents.

Time to contain (FY 22)

Time to contain	Number of incidents (FY 22)
Under 8 hours	15
8-24 hours	2
1-3 days	8
More than 10 days	6

The second graph and table break this down further by providing the time to containment, based on the severity level of the incident (noting that there were no Level 0 incidents for the FY 22 period). Our three Level 1 incidents for the year were all contained within 8 hours. A total of 15 out of the 34 incidents for the year (44%) were contained within 8 hours.

Time to contain (FY 22)

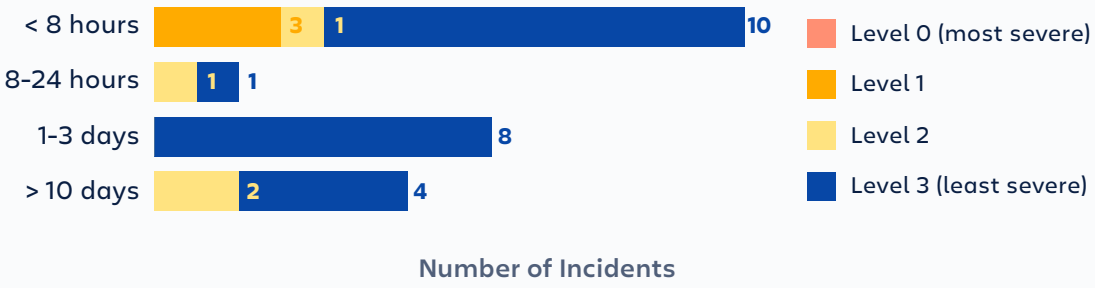

Time to detect	Level 0 (most severe)	Level 1	Level 2	Level 3 (least severe)
Under 8 hours	0	3	1	10
8-24 hours	0	0	1	1
1-3 days	0	0	0	8
More than 10 days	0	0	2	4

Note:
It is important to take into account when looking at this data that some incidents take longer to contain because they require our security intelligence team to work with other teams across the organisation to support investigation and containment efforts.

For FY 22 and beyond, Atlassian has had a specific focus on reducing time to containment, and we are continuing to implement strategies to facilitate this.

These include:

- Developing capabilities to support automated acquisition and time-lining of forensic artefacts;
- Developing capabilities to supporting containing endpoints infected with malware;
- Conducting breach readiness analyses with various products to understand the most effective way to respond to different types of security incidents; and
- Ensuring we have adequate logging in place to support our investigations process.

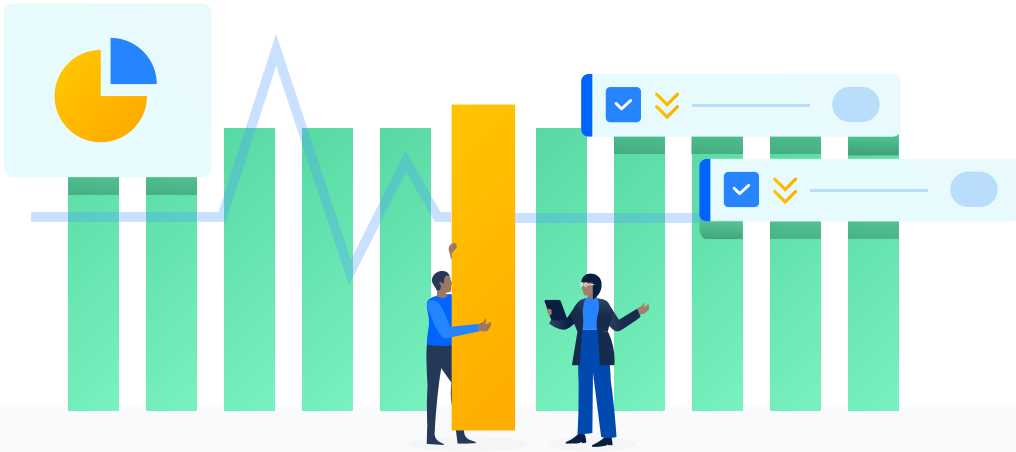

Have more questions?

While we hope that the data in this paper has provided a helpful insight into our incident response program for the last financial year, if you would like more information please send an email to security@atlassian.com, or alternatively visit our support portal and lodge a request at support.atlassian.com/contact.

ATLASSIAN