ATLASSIAN

Atlassian Data Center: advanced auditing

How to monitor security events in
Data Center with third-party integrations

Executive summary

Security, compliance, and governance – three things that are top of mind for any modern-day enterprise.

One of the biggest challenges an IT department faces is what to do with their data. Not only is your IT department responsible for finding a cost effective way to retain your organization's data, but they also need to understand what the data means. The number of users accessing your applications isn't just about how much data they generate. Rather, the data paints a picture of your organization's overall application security. Who has admin permissions? How many failed login attempts have been made? Answering these types of question can help you quickly identify security concerns and let your IT team take action.

To help your IT department to start building those insights, we've added advanced auditing to our Atlassian Data Center products – our self-managed enterprise edition. Advanced auditing enables you to track the different events within your instance, share and store you data with a third-party monitoring tool, and gain additional insights into your instance using those tools.

Finding a solution that can help your IT teams manage these changing requirements can be challenging, which is why we've built our Atlassian Data Center products. Data Center, our self-managed enterprise edition, is purpose built to meet the growing demands of the enterprise.

Advanced auditing: What are the core capabilities?

Advanced auditing in Atlassian Data Center helps organizations increase security, meet compliance demands, and improve visibility and workflow across their instance(s). The advanced auditing feature in Data Center expands upon some of the core auditing capabilities in Server to offer an enterprise-grade auditing solution. These enhancements include extended event coverage and capabilities to scale and better leverage your auditing data for strategic decisions.

ADVANCED AUDITING CAPABILITIES

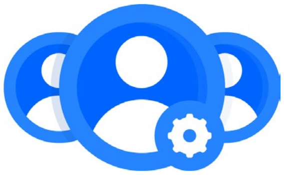

Configurable log level

Customize what your instance is tracking

The configurable log level provides organizations the ability to customize the volume of event coverage with four different levels (off, base, advanced, full) on a per "coverage area" basis. Coverage areas include global configuration and administration, user management, permissions, local configuration and administration, security, end-user activity, and apps (Jira Software Data Center, Bitbucket Data Center, Confluence Data Center). This allows organizations to set up their auditing capabilities to meet their unique demands. In addition to customizing log levels, organizations have the ability to set their database retention period (this will be addressed in further detail below).

File externalization

Integrate third-party integration tools

File externalization provides the ability to integrate with best in breed third-party monitoring tools for secure long term storage of event data, improved data analysis capabilities, and to help meet specific security and compliance demands your organization may have.

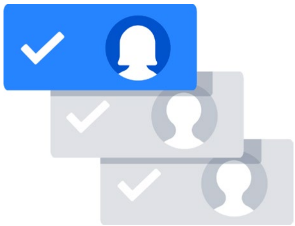

Audit delegation

Unblock teams and improve agility

Audit delegation enables admins to scale their capacity by delegating auditing capabilities, reducing time spent on simple requests, and unblocking teams by giving them the visibility and agility needed to get their work done, improving overall workflow.

Audit UI

Quick and easy in-product visibility

The UI provides a quick and simple way to review the instance without having to spend unnecessary time digging through page after page.

What does this all mean for your enterprise?

Organizations will now have a security-relevant digital record of what is happening in their instance and the ability to securely store and analyze data to help them meet their unique needs whether that is improved security, demonstrating compliance, improving workflow, and/or better visibility.

Nuts and bolts of advanced auditing

Coverage levels

The logs generated by advanced auditing show all the global events in your instance. Depending on the needs of your organization, you may want to change the level of information that you log.

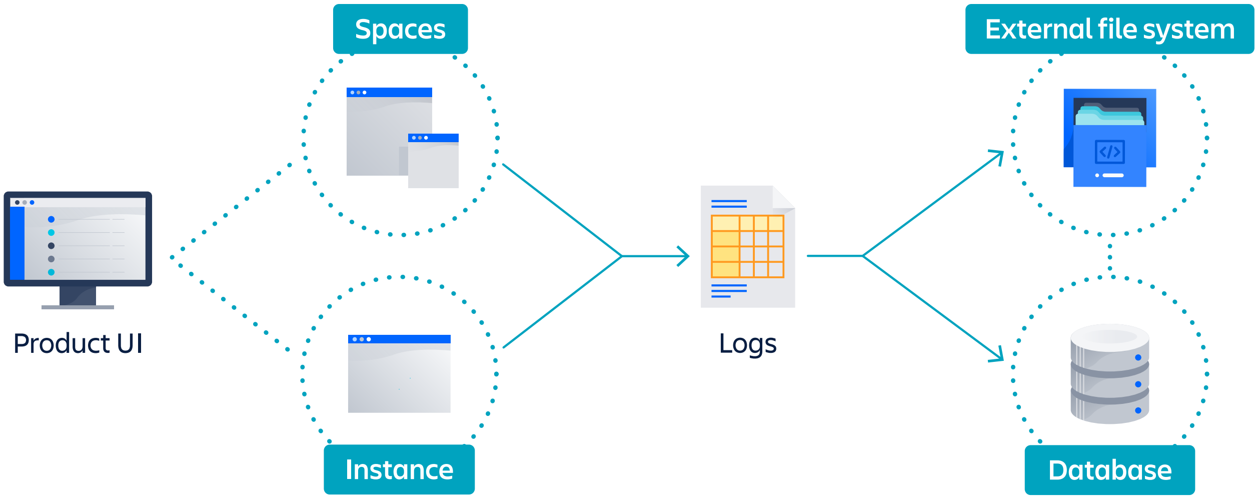

There are four types of global coverage available for advanced auditing:

- **Base:** Logs low-frequency and some of the high-frequency core events from selected coverage areas in your instance, project, or space. By default, advanced auditing is set to the base logging level.

> Note: Events that occurred prior to enabling advanced auditing are classified as base logging events. We migrate the most recent 10 million events once this capability is enabled. Any additional events are still stored on the database if you want to back it up.

- **Advanced:** Logs the core events as well as the low and medium frequency events from your instance, project, or space's activity.

- **Full:** Logs all the events available with the base and advanced coverage level, plus additional events for a comprehensive record of your instance activity.

- **Off:** Turns off logging events in your instance.

Here are some things you'll want to consider when selecting a logging level

- What is the level of compliance required by your organization?
- How much information do you need to ensure that your instance remains secure?
- How much insight do you want into your how your instance is being used?
- What level of logging will help you take action if something goes wrong in your instance?

Retention period

Based on the coverage type that you select, all of the raw event data is recorded in the audit logs and stored on your database until you've reached the retention period, which defaults to 20 years.

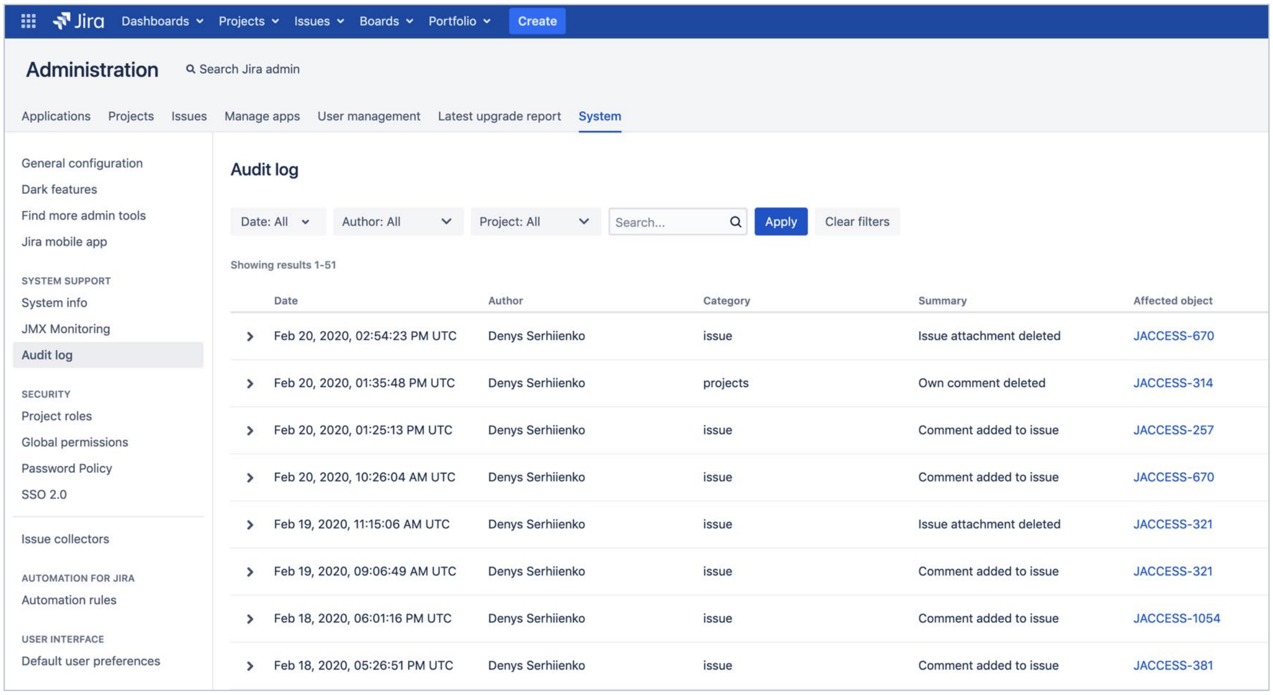

Showing results 1-51

Date	Author	Category	Summary	Affected object
Feb 20, 2020, 02:54:23 PM UTC	Denys Serhiienko	issue	Issue attachment deleted	JACCESS-670
Feb 20, 2020, 01:35:48 PM UTC	Denys Serhiienko	projects	Own comment deleted	JACCESS-314
Feb 20, 2020, 01:25:13 PM UTC	Denys Serhiienko	issue	Comment added to issue	JACCESS-257
Feb 20, 2020, 10:26:04 AM UTC	Denys Serhiienko	issue	Comment added to issue	JACCESS-670
Feb 19, 2020, 11:15:06 AM UTC	Denys Serhiienko	issue	Issue attachment deleted	JACCESS-321
Feb 19, 2020, 09:06:49 AM UTC	Denys Serhiienko	issue	Comment added to issue	JACCESS-321
Feb 18, 2020, 06:01:16 PM UTC	Denys Serhiienko	issue	Comment added to issue	JACCESS-1054
Feb 18, 2020, 05:26:51 PM UTC	Denys Serhiienko	issue	Comment added to issue	JACCESS-381

You do have the flexibility to define your retention period, but decreasing the period will delete all the logs that exceed the time period you set and you will no longer see them in your audit log view. When deciding the right retention period for your instance, there are several factors to consider, such as number of active users, type of event coverage you select, and auditing and legal constraints.
For most organizations, we recommend setting your retention period at three years to start. However, there is a limit of 10 million records, so, if you're a larger organization, you may hit that number quicker. Leveraging the file externalization capability of advanced auditing will allow you to use your third-party monitoring tools to store your audit logs for longer periods of time.

> Pro tip: If you have a smaller data set, you can leverage REST and query the audit logs for the information that you need, or you can implement AuditConsumer, which will receive the same information as the file receiver. For larger data sets, leveraging the external file system will help with your data storage.

Admin usage and permissions

Most traditional audit logging features tend to look at only global events in your instance. With our advanced auditing feature, not only can you audit your entire instance, but you can also audit individual spaces and projects. This gives space and project admins the flexibility to audit their individual spaces without having to filter through all of the global events logged in the instance.

If you're a space or project admin, you can specify the name of your project or space to show only the logs that are relevant to you in the product UI.

Check out our advanced auditing documentation to learn more.

Technical aspects of advanced auditing

Because of its architecture, getting advanced auditing up and running for your Data Center applications is simple.

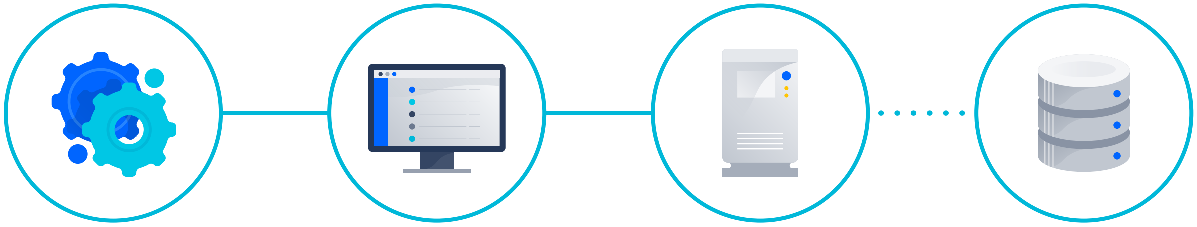

When you download or upgrade to Jira Software 8.8, Jira Service Desk 4.10, Bitbucket 7.0, and Confluence 7.5, advanced auditing is automatically available for your to start using.

As we continue to add functionality to advanced auditing, you'll be able to get access to these capabilities by downloading the latest version from the Atlassian Marketplace.

You can even use your apps to log audit events too. These apps will show up in the apps event logging category.

Components of advanced auditing

There are two main architectural components to advanced auditing:

- Product UI
- External file system

> Note: While implementing an external file system is a component of advanced auditing, you don't have to immediately setup your monitoring tools to consume the log files. You can still use the product UI to review event changes in your instance and files will be stored on your database.

Product UI

From the Administration page of your Data Center product, you can immediately start logging events. All you need to do is specify the logging level you need and easily filter the types of events depending on your organizations requirements.

External file system

One of the key components of advanced auditing is the external file system, or file externalization. Rather than just storing your audit logs on your database, you can leverage your existing third-party monitoring tools to store your audit logs.

> Note: If you're considering deploying your Data Center applications on Amazon Web Services (AWS), advanced auditing is already integrated into our AWS Quick Start templates for Jira Data Center, Confluence Data Center, and Bitbucket Data Center. That means that you can immediately start using AWS CloudWatch to monitor and store your audit logs (learn more later in the whitepaper).

Third-party monitoring tools

The following section covers best practices, tips, and recommendations for popular Data Center integrations

AppDynamics

AWS CloudWatch

Datadog

Elastic Stack

Splunk

Sumo Logic

APPDYNAMICS™
part of Cisco

Integrating AppDynamics with Data Center

ⓘ You need to configure log monitoring on each node in your DC cluster.

1. Log in to your Data Center instance as an admin, and click through to the audit integration settings (**Admin Console > Audit Log › "…" Menu › Integration**). Once there, you may configure the file directory path where you want to store your logs.

ⓘ Remember this path because you'll need to point to it in your monitoring tool.

2. Install the Analytics Agent on each server that is hosting your Data Center product.
3. Configure the Centralized Log Management settings and connect to your Controller.
4. Now that your agents are installed and connected, use the Centralized Log Management UI to configure your log data sources using source rules.
    - Define the source rule by pointing to the file directory path entered in your Data Center product.
    - Upload a recent log file as a "sample log" to preview and customize your Field Extraction rules.
5. Wait up to 5 minutes for the agent to pickup the latest configuration.

APPDYNAMICS™
part of Cisco

Tips, tricks, and recommendations

AppDynamics new (4.3+) Centralized Log Management UI simplifies configuration.

You can parse the log file using pattern matches and define criteria that will build the patterns better.

The Analytics Agent checks the logs and sends new events to AppDynamics once a minute.

By default, the controller communicates new configuration information to the Analytics Agent every five minutes. So, it could be up to five minutes before the agent starts tailing your log file using the new source rule configuration.

If you're interested in learning more about AppDynamics, check out the AppDynamics documentation.

Integrating AWS CloudWatch with Data Center

ⓘ You need to configure log monitoring on each node in your DC cluster.

1. Log in to your Data Center instance as an admin, and click through to the audit integration settings (**Admin Console > Audit Log › "…" Menu › Integration**). Once there, you may configure the file directory path where you want to store your logs.

ⓘ Remember the path you set in Data Center. You'll need it for the CloudWatch agent configuration.

2. Install and configure AWS CloudWatch agent. Your configuration might look similar to the following:

```
{
  “logs”:
   {
       “logs_collected”: {
           “files”: {
               “collect_list”: [
                   {
                       “file_path”: “/opt/atlassian/data/
                       jira/log/audit/*.log”,
                       “log_group_name”: “jira-audit.log”,
                       “log_stream_name”: “enterprise-demo-tis”
                   }
               ]
           }
       },
       “log_stream_name”: “default_log_stream_name”
      }
}
```

3. Confirm the events are visible in CloudWatch.

Tips, tricks, and recommendations

Review the AWS CloudWatch Logs concepts with your team. It's important that everyone understands how CloudWatch works and what certain terminology means.

Make sure you are following AWS security best practices by assigning the appropriate IAM instance roles and only granting people in those roles access to CloudWatch.

The Cloudwatch @timestamp is based on ingest time, which may not always match the actual event timestamp. This is not configurable.

But don't worry, to filter or search by the correct time, you can use you can use:

```
fromMillis(timestamp.epochSecond*1000)
```

You can adjust the max delay for events being flushed to CloudWatch by adjusting the `force_flush_interval` in your agent configuration. By default, the delay is set to 5 second.

If you're interested in learning more about CloudWatch, check out the CloudWatch documentation.

Integrating Datadog with Data Center

You need to configure log monitoring on each node in your DC cluster.

1. Log in to your Data Center instance as an admin, and click through to the audit integration settings (**Admin Console > Audit Log › "…" Menu › Integration**). Once there, you may configure the file directory path where you want to store your logs.

Remember this path because you'll need to point to it in your monitoring tool.

2. Install the Datadog agent.

   Note: You need to enable log collection for each server that your running Data Center on.

   Your `/etc/datadog-agent/conf.d/jira.d/conf.yaml` might look something like this:

```
Nlogs:
  - type: file
    path: “/opt/atlassian/data/jira/log/audit” #path
    configured in Jira UI.
    service: “jira” #match specific app_name used in APM
    source: “jira” #general technology name
```

3. Give the Datadog agent read and execute permissons on the log file and parent directories.

4. Confirm that the agent has access by checking the status of the Datadog agent.

5. View events in Datadog. You can also choose to create facets for the area you are auditing or add other attributes you want to make sure are monitored

6. You're all set. Confirm that you can see your events in Datadog.

Tips, tricks, and recommendations

Setup email notifications to alert your team and help troubleshoot events that have been logged in your instance.

Set up metric monitors to notify you when anomalies are detected in your data.

Audit logs are generated in real-time, so your events are pulled into Datadog every 15 seconds.

With Datadog Notebooks, tell a story with your data. Notebooks combines graphs of your data with accompanying text to help you and your team reflect on your data over time.

If you're interested in learning more about Datadog, check out the Datadog documentation.

Integrating Elastic Stack with Data Center

You can use either Filebeat or Logstash to send events to Elastic Cloud. In our setup, we've used Filebeat to send events to our instance. Filebeat is a lightweight alternative to logstash and provides timestamps and message information for events. You can also use Logstash in tandem with Filebeat for further enrichment or processing of your events before sending them to Elastic Stack.

ⓘ You need to configure log monitoring on each node in your DC cluster.

1. Log in to your Data Center instance as an admin, and click through to the audit integration settings (**Admin Console > Audit Log › "…" Menu › Integration**). Once there, you may configure the file directory path where you want to store your logs.

ⓘ Remember this path because you'll need to enter it in Elastic Stack.

2. Install Filebeat and connect it to your Elastic Search instance.
3. Modify `thefilebeat.yml` file to point to the log directory you configured in your Data Center application.

   Note: You should mark the input as a series of JSON lines rather than using plain text.

4. Start Filebeat and confirm you can see recent auditing events in ElasticSearch.

elastic

Example `filebeat.yml` (partial)

```
filebeat.inputs:
- type: log # use generic log input type
  paths:
    # match the path from Atlassian Audit settings.
    - /opt/atlassian/data/jira/log/audit/*.log
  # add atleast 1 json config.
  # WARNING: placing under root will cause conflicts
  # on source, system and other fields.
  json.keys_under_root: false
  json.overwrite_keys: true
  close_inactive: 24h # ensure this valuye exceeds longest
possible del;ay in events.
  tags: [“advanced-audit”] # optional
```

5. You’re all set. For more info, check out Elastic Stack documentation.

Tips, tricks, and recommendations

To make events searchable, configure Elastic’s Filebeat to parse the file as JSON objects, allowing filtering or querying on their value.

Setting `keys_under_root:` true will cause field conflicts.

If you want the audit details to be under your root of the message you must use additional processors in filebeat.yml to modify event fields names.

Filebeat will stop monitoring files that don’t change often. You may need to adjust `close_inactive` to account for any idle periods in your instance (default is 5 minutes).

Filebeat will check for new files every 10 seconds, and send updates for existing files every 1 second. Both of these intervals are adjustable.

splunk>

Integrating Splunk with Data Center

ⓘ You need to configure log monitoring on each node in your DC cluster.

1. Log in to your Data Center instance as an admin, and click through to the audit integration settings (Admin Console > Audit Log > "..." Menu > Integration). Once there, you may configure the file directory path where you want to store your logs.

ⓘ Remember this path because you'll need to enter it in Splunk.

2. Download and install Splunk's Universal Forwarder and Universal Forwarder credentials.
3. In Splunk, configure a new data source.
   - From **Settings**, choose **Add Data**.
   - Select **Forward**.
   - Choose the server matching your Data Center instance and give it a name.
   - Select **Files and Directories**.
   - Enter the path you used in your Data Center application.
   - For **Source type**, choose your newly defined type (see tips on left), and an appropriate index.
4. In a few moments Splunk's universal forwarder should poll for the updated configuration, and start sending data.
5. You're all set. Confirm that you can see your events in Splunk

splunk>

Tips, tricks, and recommendations

Splunk can extract the audit data into fields. This is easily achieved with a custom source type with the `json` extraction type. You should define this before adding your logs.

To properly define the timestamp, define the source type with an extracted timestamp value as follows:

- Indexed extractions: json
- Timestamp Extraction: Custom
- Format: %s (seconds since epoch)
- Timestamp fields: timestamp.epochSecond

If you're interested in learning more about Splunk, check out the Splunk documentation.

Integrating Sumo Logic with Data Center

ⓘ You need to configure log monitoring on each node in your DC cluster.

1. On the Administration page of your application, click Audit Integration, and enter the file directory path where you want to store your logs.

ⓘ Remember this path because you'll need to provide it during the Collector configuration.

2. In Sumo Logic, download and install the connector. You can either use the Collector's online wizard, or install it manually.
3. Enter the name of the audit log directory from your Data Center application in the Collector UI or the configuration files.
4. The Collector will collect and analyze any existing files and make them available for query.
5. You're all set. Confirm that you can see your events in Sumo Logic.

Tips, tricks, and recommendations

Here are some best practices on how to deploy Sumo Logic so that you're able to get the most from your logs

- Good Source Category, Bad Source Category
- Local and Centralized Data Collection

Find events of interest using Sumo Logic's enhanced search capabilities. Here's an example:

```
_sourceCategory="prod/atlassian/jira/audit" | json
"author.name","auditType.area","auditType.action"
```

- Make sure you optimize your search to find all of the relevant information you need faster.
- Specify a Field Extraction Rule to parse JSON fields at the time the log files are ingested.

Detect patterns in your log files and group similar events together with the LogReduce algorithm. You can influence LogReduce to further collapse or divide events as desired.

If you're interested in learning more about Sumo Logic, check out the Sumo Logic documentation.

Ready to get started?

If you’re still interested in learning more about advanced auditing, check out our advanced auditing webinar that talks about how advanced auditing gives you the visibility to make informed business decisions, safeguard your data, and meet internal and external regulations.

Download these Data Center offerings and try them free

Bitbucket Data Center

Confluence Data Center

Jira Software Data Center

Want to learn more?

atlassian.com/datacenter

ATLASSIAN