



Running the Atlassian Stack on Amazon Web Services

Scandio GmbH © 2018



scandio
TECHNOLOGY & CONSULTING

sales@scandio.de

Table of Contents

1. Using the Atlassian Stack in the AWS Cloud.....	4
2. The Atlassian Stack.....	4
2.1. Jira Software.....	5
2.2. Confluence.....	5
2.3. Bitbucket Server.....	5
2.4. Bamboo.....	5
2.5. Jira Service Desk.....	5
2.6. Crowd.....	5
3. Comparison to Atlassian Cloud.....	5
4. Complementary Tools.....	6
4.1. JFrog Artifactory.....	6
4.2. Jenkins CI Server.....	7
4.3. Sonarqube Code Quality.....	7
4.4. Matomo Web Analytics.....	7
5. Important AWS Services for the Atlassian Stack.....	7
5.1. Amazon Elastic Compute Cloud (EC2).....	7
5.2. Amazon Relational Database Service (RDS).....	7
5.3. Elastic Load Balancing (ELB).....	8
5.4. Amazon Virtual Private Cloud (VPC).....	8
5.5. Amazon Simple Storage Service (S3).....	8
5.6. Route 53.....	8
5.7. AWS Certificate Manager (ACM).....	8
5.8. Key Management Service (KMS).....	8
5.9. Identity and Access Management (IAM).....	8
5.10. CloudWatch Monitoring.....	8
5.11. Other Services.....	8
6. High Availability and Scalability.....	9
7. Network Architecture.....	10
8. Installation and Migration.....	10
9. Monitoring, Alerting, Capacity Planning.....	12
10. Security & Privacy.....	13
11. Cost optimization through relief of infrastructure and capacities.....	13
12. Time saving and increased efficiency.....	14
13. Supplier Information.....	15
13.1. Amazon Web Services (AWS).....	15
13.2. Atlassian.....	15
13.3. Scandio GmbH.....	16

Executive Summary

The decision to host their system in a self-administered on-premise or in a service provider operated environment is often a challenge for many IT-decisionmakers. This whitepaper deals with operating the Atlassian stack in the AWS Cloud and it is intended to provide a possible scenario for all those facing this decision.

There are recurring important aspects of hosting Atlassian products in the AWS Cloud. We suggest the following course of action:

- **High Availability and Scalability** – Scalability is no longer a problem thanks to self-service, as new infrastructure is available at the push of a button if required. "Elastic Load Balancing" distributes incoming application traffic to multiple destination server instances. It can route the variable load of traffic to multiple availability zones.
- **Security** – Working in the environment of AWS services does not mean that all components are available on the Internet. Complete isolation and partially controlled opening is easy to configure using virtual private cloud (VPC) networks. Systems can also be isolated from each other in accordance with company guidelines.
- **Cost optimization** – The costs of the AWS services can be easily determined. It eliminates up-front costs and establishes a comfortable pay-as-you-go model based on actual expenses and not on capital investments. Through further optimization measures such as "auto scaling" or "reserved instances", costs can also be saved during operation.
- **Time saving and increased efficiency** – go live can be realized faster than with on-premise solutions, since coordination with internal infrastructure is no longer necessary and is now entirely the responsibility of the cloud provider. Efficiency gains continue to be made during ongoing operations and can be applied to all maintenance and upgrade work that now takes place externally.

The scenario presented in this white paper is based on the use of server versions of Atlassian applications. Even though Atlassian itself provides a cloud version of it, the server version has some advantages:

- **Storage space** – The Atlassian Cloud offers up to 100 GB per application. A server version installed in the AWS cloud is not limited in storage space and can be increased at any time.
- **Number of users** – The maximum number of total users is 5,000 in Atlassian Cloud. Since the products are also very popular in non-software teams this can be a limiting factor.
- **Data backup** – Atlassian does not currently offer automatic backups to restore the system in an emergency. This must be done manually by the customer. The use of AWS Services facilitates zero-downtime backups and enables the fully automated orchestrated creation of test systems.
- **App Availability** – The Atlassian Marketplace (<https://marketplace.atlassian.com/>) offers over 3,000 apps for its products like Confluence or Jira. The apps are essential for integration into the company and for mapping complex processes and interfaces. Atlassian offers only 1,200 extensions by now for their cloud version. Since the installation in the AWS Cloud is in principle a server version, the apps can be used to their full extent.
- **Customization** – The Connect API for Atlassian Cloud makes the products extensible. However, the possibilities are significantly limited by the complex cloud security concept.

Of course, the possibilities have to be weighed up individually, cleanly and with difficulty in each case. This white paper can be used to make a well-founded decision.

1. Using the Atlassian Stack in the AWS Cloud

This white paper presents the basic concepts of using the Atlassian stack in the AWS cloud and discusses which components and toolchains make sense. To further establish the context, the main advantages of working with the cloud compared to a local environment are presented. In addition, this white paper covers many of the unique benefits of an AWS hosted service. The clarity of the content is supported by the initial definitions of terms.

For more than a decade, Australian software manufacturer Atlassian's products have been widely used in enterprises. The triumphal march began in IT departments and in software development. With the spread of agility and the DevOps movement throughout the enterprise world, large and midsize companies are recognizing further challenges:

- How can one ensure that the systems perform well and are permanently available with strongly increasing user bases?
- How can one guarantee the highest quality of operation at the lowest possible cost without burdening already scarce internal capacities?
- How can a company enable external partners to use the tools together with its employees to create the necessary transparency for innovation projects?

A trend can be clearly deduced from the experience of recent years. The so-called agile toolchain can be operated optimally, with high availability, securely and in compliance with data protection regulations in the AWS cloud. The services provided by AWS not only accelerate the commissioning or relocation, but also increase the availability and scalability of the platform without any special project expenses.

In the following sections, the Atlassian Stack and the AWS technologies, which enable its highly available and scalable operation, are described in detail. In addition to information on the extensibility of the toolchain through other applications and on security and data protection within the cloud, you will also find technical details on infrastructure, architecture and monitoring in AWS.

2. The Atlassian Stack

An agile tool chain describes an application landscape that is required for the successful execution of agile projects. Atlassian products play an increasingly important role here.

The essential products for this are:

- Jira Software
- Confluence
- Bitbucket Server
- Bamboo
- Jira Service Desk
- Crowd

Interlocked, the products show their full strength and can be used profitably for the entire company in agile project management, team collaboration, knowledge management and in the area of continuous integration or continuous delivery.

2.1. Jira Software

Jira Software is the heart of agile project management. Jira Software can be used for all types of projects, free of usage requirements. Scrum and Kanban, but also waterfall driven projects are possible. Extensions such as Jira Portfolio enable multi-project management and support large companies in scaling their newly gained agility. Jira software can not only be used for temporary IT projects, but also enables the handling of the entire task, error and test case management. Process types, workflows, notifications and permissions are freely definable and customizable for each project scenario.

2.2. Confluence

With Confluence, Atlassian's second major product, all unstructured information from projects or the entire company is usually managed. With Confluence, a company is able to create its own communities, manage requirements engineering for projects and create the single point of truth for knowledge management. Confluence and Jira complement each other very well and have a deep functional integration. Confluence also makes it possible to manage documentation in an audit-compliant manner.

2.3. Bitbucket Server

Bitbucket Server is Atlassian's source code management system. Based on the Git standard, Bitbucket enables all general Git functions on an enterprise level. Authorizations and user management are fully integrated and the link to Jira software provides a completely transparent overview. Git extensions such as pull requests or forks are also standard features of Bitbucket.

2.4. Bamboo

Bamboo is Atlassian's Continuous Integration Server. With Bamboo, software can be created from source code management, automatically tested and delivered to the stage and production environment. Bamboo can also be optimally integrated with Jira software, Bitbucket and Fisheye.

2.5. Jira Service Desk

With Jira Service Desk it is possible to extend Jira so that internal and external customers can use Jira either as an ITIL compliant service desk or as a simple help desk. By linking Jira Service Desk with Jira Software, IT and development teams can work together on a common platform to resolve incidents faster and implement changes more reliably.

2.6. Crowd

All Atlassian products can be controlled via the central user management Crowd. With Crowd, companies have the possibility to use different user administrations like Office 365, LDAP/Active Directory or a local user management at the same time. Crowd also allows the implementation of various single sign-on solutions, and since version 3.3 it has been possible to delegate group management to the project managers with Crowd.

3. Comparison to Atlassian Cloud

Atlassian is not only a manufacturer of on-premise software. Many products can be used directly on Atlassian's cloud platform. The offer is extensive and convenient and is developing rapidly. Many innovations of the core products are only available in the Atlassian cloud and are also visible in the on-premise products after successful mass testing.

But what are the differences between this solution and a self-organized or externally controlled operation in the AWS Cloud? There are several criteria to consider here:

- **Storage space** - Depending on the number of users, the Atlassian Cloud offers a maximum of 100 GB per application. This is not sufficient for all companies and must be taken into account when making a decision. If installed in the AWS cloud, there is no limit to the application size and this can be increased at any time.
- **Number of users** -The maximum number of total users (all products) was increased from 2,000 to 5,000 users in September 2018, but the products are also very popular in teams far away from software development. Therefore, user numbers beyond the 5,000 users in a company are no longer uncommon. For installations in the AWS cloud, there is no limit to the number of users.
- **Data backup** - Atlassian does not currently offer automatic backups to restore the system in an emergency. This must be done manually by the customer. Atlassian does not currently provide any automatic mechanism. The use of AWS Services facilitates zero-downtime backups and enables the fully automated orchestrated creation of test systems.
- **App availability** - The Atlassian Marketplace offers more than 3,000 apps for the products. The apps are essential for integration into the company and for mapping complex processes and interfaces, but not all apps are currently available for the Atlassian Cloud. The offering in this category has been reduced to just under 1,200 extensions. Since the installation in the AWS Cloud is in principle an on-premise version, the apps can be used to their full extent.
- **Adaptability** - The Connect API for Atlassian Cloud makes the products extensible. However, the possibilities are significantly limited by the complex cloud security concept.

From the point of view of European data protection, the Atlassian Cloud cannot yet be used to its full extent. Atlassian is currently working on a GDPR/EU-DSGVO compliant solution. This is planned for 2019. A computer center that ensures the unique storage of data in the European Union is currently only visible on paper. Therefore, deadlines for the technical implementation cannot yet be set.

In this respect, the Atlassian software can be used in the AWS cloud, as in this environment all conditions of the European data protection regulation are fulfilled. The AWS services are also certified according to ISO 27001, ISO 27017 and ISO 27018.

Another important point concerns the regions and end points of the respective services in order to ensure optimal feasibility.

To reduce data latency in the applications, most Amazon Web Services offer a regional endpoint. An endpoint is a URL that acts as an entry point for a Web service. For example, apigateway.eu-central-1.amazonaws.com is an endpoint for the Amazon API Gateway based in Frankfurt, Germany.

4. Complementary Tools

In addition to the Atlassian Tools, the following services are available in the area of software development.

4.1. JFrog Artifactory

JFrog's Artifactory allows you to store a variety of different build artifact types that arise during development.

4.2. Jenkins CI Server

The Jenkins CI Server is the leading Open Source Build Server, with which almost all tasks in the CI/CD environment can be implemented. Due to the large selection of extensions, various software projects can be built and easily integrated with various services.

4.3. Sonarqube Code Quality

Sonarqube is a popular source code quality management tool that can detect potential errors during the development process using static source code analysis techniques.

4.4. Matomo Web Analytics

Formerly known as Piwik, it can be configured to be DSGVO and works council compliant and enables page view analysis within web applications. For example, page views can be tracked in Confluence and, based on this, targeted investments can be made in topics that receive special attention within the organization.

5. Important AWS Services for the Atlassian Stack

Operating the Atlassian stack requires some of the more than one hundred AWS services. They form the backbone of a simple and cost-optimized operation. These are:

- Amazon Elastic Compute Cloud (EC2)
- Amazon Relational Database Service (RDS)
- Elastic Load Balancing (ELB)
- Amazon Virtual Private Cloud (VPC)
- Amazon Simple Storage Service (S3)
- Route 53
- AWS Certificate Manager (ACM)
- Key Management Service (KMS)
- Identity and Access Management (IAM)
- CloudWatch Monitoring

Each AWS service offers exactly the right performance and the deep integration of the services with each other simplifies the configuration. The following description of the AWS services gives an overview.

5.1. Amazon Elastic Compute Cloud (EC2)

The web service Amazon Elastic Compute Cloud (Amazon EC2) provides secure, scalable computing capacities in the cloud. With EC2 it is possible to create all classic server services for Linux and Windows. A selection of ready-made images simplifies use.

5.2. Amazon Relational Database Service (RDS)

Amazon Relational Database Service (Amazon RDS) quickly and easily provides relational databases in the cloud that automate time-consuming administrative tasks. The RDS databases are highly available out of the box and can be scaled almost arbitrarily.

5.3. Elastic Load Balancing (ELB)

Elastic Load Balancing distributes incoming application traffic to multiple targets such as Amazon EC2 instances, Docker Containers, and IP addresses. It can route the variable load of application traffic into a single Availability Zone or into multiple Availability Zones. ELB replaces all reverse proxy systems with the highest availability.

5.4. Amazon Virtual Private Cloud (VPC)

The Amazon Virtual Private Cloud (Amazon VPC) provides a logically isolated area of the AWS cloud where AWS resources run in a defined virtual network. VPC also enables the use of stateless and stateful firewalls (security groups) and subnets to selectively allow and prohibit connections. A VPN connection to the local data center can also be guaranteed via this service.

5.5. Amazon Simple Storage Service (S3)

Amazon S3 is a 99.9999999999% durable object storage device designed to store and retrieve any amount of data from any location: websites and mobile apps, enterprise applications, and data from IoT sensors or devices.

5.6. Route 53

Amazon Route 53 is a highly available and scalable Domain Name System (DNS) web service for the cloud. This makes it easy for the organization to extend DNS routing.

5.7. AWS Certificate Manager (ACM)

AWS Certificate Manager is a service that easily deploys and manages public and private Secure Sockets Layer (SSL) and Transport Layer Security (TLS) certificates for use with AWS services and internal associated resources.

5.8. Key Management Service (KMS)

The AWS Key Management Service (KMS) is a managed service that facilitates the creation and control of encryption keys used for data encryption and uses FIPS 140-2 validated hardware security modules to protect the security of keys.

5.9. Identity and Access Management (IAM)

AWS Identity and Access Management (IAM) enables secure management of access to AWS services and resources. With IAM, AWS users and groups can be created and managed, and access to AWS resources can be allowed or denied with permissions.

5.10. CloudWatch Monitoring

Amazon CloudWatch is a monitoring and management service for developers, system operators, site reliability engineers (SRE) and IT managers. This can be used to directly read out metrics via AWS services such as EC2 and RDS, send alarms via SNS as e-mails or trigger other actions such as auto scaling.

5.11. Other Services

"Behind the Scenes" is a wide range of other AWS services, but their equivalents pose significant challenges to the local data center. The most important are listed below:

- **CodeDeploy** - AWS CodeDeploy is a managed service that automates and documents deployments, including on EC2 instances. With CodeDeploy, new features can be released more easily and quickly, downtime during application deployment can be avoided, and complex application updates can be handled more easily. AWS CodeDeploy can also be used to automate software deployments.
- **Lambda** - With AWS Lambda, code can be executed without the need to deploy and manage servers. You only have to pay for the computing time that is actually consumed - there is no cost if the code is not executed.
- **SES** - Amazon Simple Email Service (Amazon SES) is an arbitrarily scalable cloud-based email sending service that provides an SMTP endpoint.
- **WorkMail** - Amazon WorkMail is a secure, managed enterprise email and calendar service that supports existing desktop and mobile email applications and Atlassian products through its IMAP connectivity.
- **SNS** - Amazon Simple Notification Service (SNS) is a long-lived, secure, fully managed pub/sub messaging service that enables the decoupling of microservices, distributed systems and serverless applications. Amazon SNS supports a wide range of subscriber endpoints including SQS, lambda capabilities, HTTP/S webhooks and mail.
- **CloudFront** - Amazon CloudFront is a fast content delivery service (CDN) that delivers data, videos, applications and APIs securely, with low latency, high transfer speeds and within a developer-friendly environment to customers around the world.

6. High Availability and Scalability

Atlassian's software products can be implemented as high-performance, fail-safe data center variants - with the exception of Bamboo. Each application can thus be run in different data center sections simultaneously on any number of nodes. With load balancers and fail-safe databases and storage systems, an availability of over 99% can be achieved. Updates and backups can be carried out with the data center variant without planned downtimes. This enables a simple and highly efficient operation.

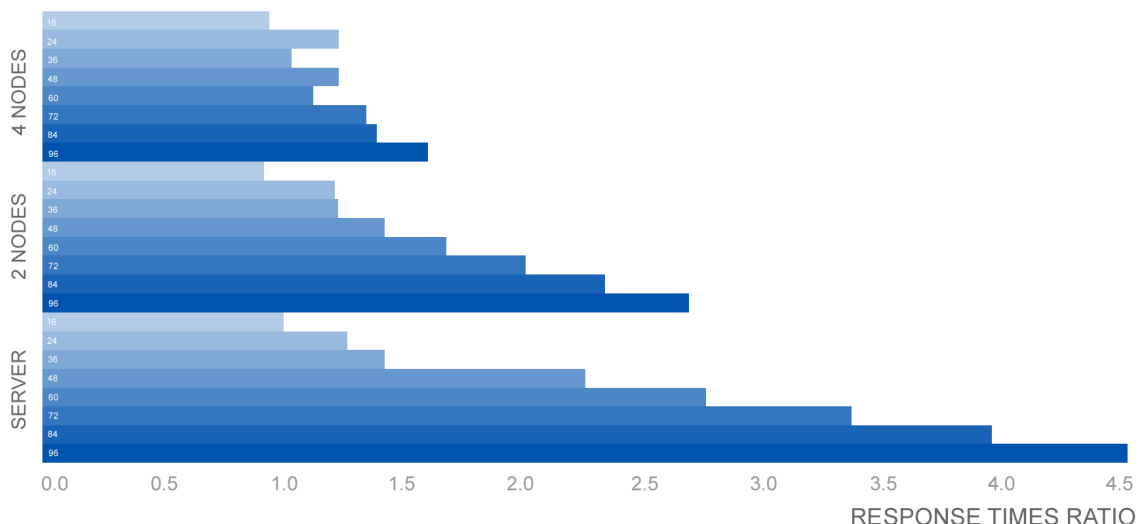


Figure 1 Relative response times of different Confluence installations. Access was tested with 16, 24, 36, 48, 60, 72, 84 and 96 browsers on a single server and on two data center installations with two and four nodes each.

(Source: Atlassian (<https://confluence.atlassian.com/enterprise/confluence-data-center-performance-656408756.html>))

Figure 1 shows a comparison of the response times of a server installation, a data center variant with two and one with four nodes. The loading time of the server installations was defined with a reference value of 1.0 for 16 browsers. It is clear that the response times increase with the number of simultaneous accesses. This increase can be easily absorbed by several nodes and thus there is hardly any difference for the user when using four nodes.

If the Atlassian products are operated in the AWS Cloud, the AWS services for databases (Relational Database Service RDS), storage systems (Elastic Block Store EBS, Elastic Files System EFS), load balancers (Elastic Load Balancer ELB) and automatic scaling (Auto Scaling Group ASG) can be fully used. Whether a data center deployment is worthwhile for an enterprise can be determined by the core requirements of

- Number of active (concurrent) users
- general availability in different time zones
- System performance at high load
- complete reliability

These are recognized, evaluated and decided on. The available AWS services for scaling and high availability reduce the effort of implementation enormously.

7. Network Architecture

The installation of an application environment including the integration of the required services and their connectivity is often underestimated because it is very costly in practice. The considerations usually begin with the basic network architecture: Should the services only be accessible via the company network or a VPN connection, or is free availability via the Internet desired in order to enable simple onboarding of new employees and external partners?

In the first case, a VPN connection between an isolated AWS network area and the company network is the best option, in the second case the services are configured for general availability via the Internet. Another central issue is user administration, which can be connected to an existing LDAP or ActiveDirectory system as required or implemented using a Single Sign-On (SSO) solution. Questions about DNS, certificates for transport encryption and the connection for e-mail services for sending notification e-mails and, in the case of Jira ServiceDesk, receiving replies from users are often forgotten.

Working in the environment of AWS services does not mean that all components are available on the Internet. Complete isolation and partially controlled opening is easy to configure using virtual networks. Systems can also be isolated from each other in accordance with company guidelines.

Due to the variety of AWS services, these requirements can be completely implemented in the cloud, but can also be provided by the company data center as required. Experience has shown, however, that additional effort can then be expected through communication and the implementation of connectivity between the local data center and AWS. This ultimately results in increased implementation time.

8. Installation and Migration

By today's standards, a manual installation is considered obsolete. Instead, systems are provisioned as code (IaC) using CloudFormation or Terraform infrastructure, for example.

Automation with CodeDeploy and Ansible can also be used for application installation and configuration. To accelerate the installation process, the creation of individual EC2 instance images (AMIs) can also be done by e.g. packers and the installation and configuration part can be minimized when starting a new instance. This Immutable Infrastructure not only speeds up the provisioning process considerably, but also enables a versioning of the server installation that enables traceability and a rollback in the event of an error. A significant increase in stability is achieved by reducing the human error source.

From the experience of past migration projects some topics must not be forgotten.

- Early reduction of the time-to-live of the Domain Name System (DNS) is essential, otherwise propagation can take several hours or even days. In the worst case, some users would not be able to access the service during this time.
- When making backups from the old installation, it is particularly important to note that it is not sufficient to create a backup via the user interface. Instead, a copy of the database contents and hard disk should be created, since not all elements are contained in the native backup, or the backup can even fail in large installations.
- Closely linked to this topic is the goal of downtime minimization: for installations with large space requirements on a hard disk (several hundred GB), it is advisable to start an incremental copy in advance to prevent downtimes of many hours or days.

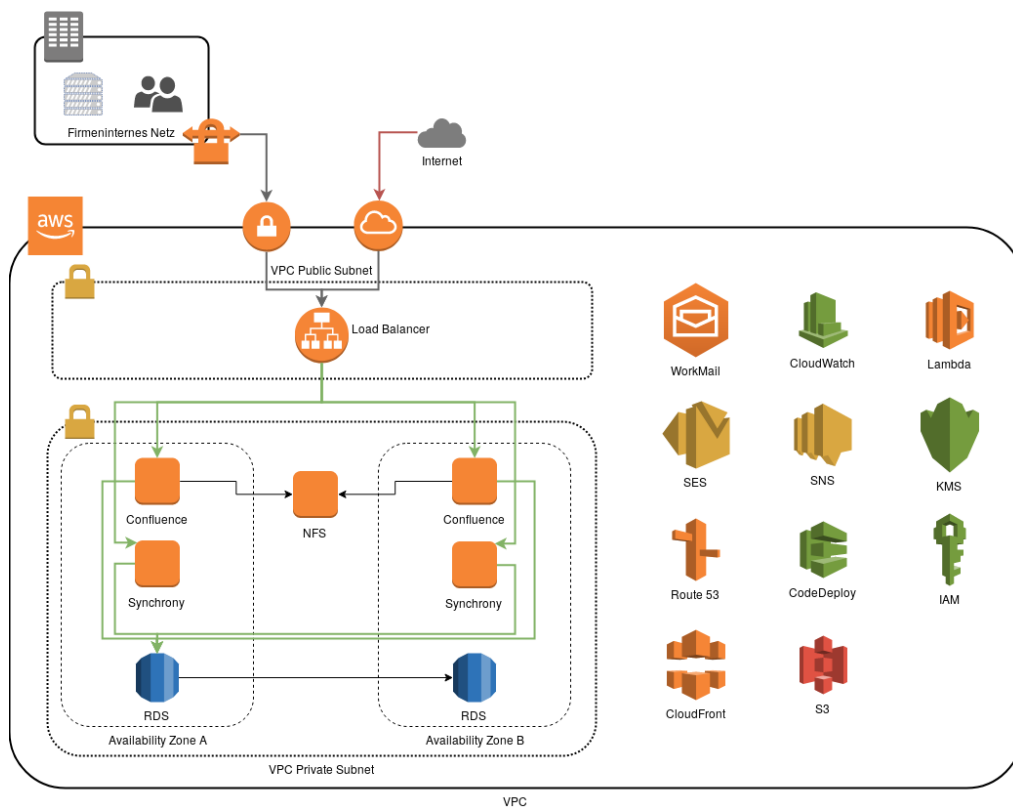


Figure 2 Overview of the AWS architecture for operating a Confluence in the data center variant with Synchrony outsourced to two separate servers. The application instances, the RDS databases as well as the load balancer are highly available and fail-safe distributed across different availability zones and can thus withstand a failure. On the right, additional services are listed that guarantee deployment, monitoring, sending and receiving mail as well as backup and security.

(Source: Scandio GmbH auf der Grundlage von AWS)

Figure 2 shows all the services used at a glance. In addition, the architecture and the interaction of EC2, ELB and RDS as well as the VPN connection to a local data center are shown schematically. As an example, a data center installation of Confluence with a separate Synchrony (Synchrony is the solution for collaborative processing in Confluence) is selected. Users can only access the application via the Elastic Load Balancer. In a single-server installation, the load balancer serves as a reverse proxy. In both cases, traffic is forwarded to the otherwise isolated private subnet. The same load balancer is also used to forward requests to Synchrony.

The RDS database is also located in the private subnet and can only be accessed from the application instances. All these resources exist at least twice in different availability zones to ensure the greatest possible reliability. For the database, this must be selected directly in the configuration of RDS ("Multi-AZ"), the various EC2 instances are provisioned separately. Another instance also serves as a central NFS server for data that has to be shared between the application instances. AWS Elastic File System (EFS) would also be an option, but for large installations, the lack of snapshot capability leads to complications during backup and test system generation, as copying the entire installation can take too much time. Using the snapshot feature for EBS volumes and RDS databases, zero-downtime backups can be created quickly and easily. The recovery of the complete system as a basis for a test system is thus possible, even with very large installations with several hundred gigabytes, without problems in less than 20 minutes and can easily be automated.

With sufficient automation, for example through Ansible and CodeDeploy, this design can also create an auto-scaling environment, which enables automatic addition and removal of new EC2 instances at high or low load through AWS-internal mechanisms. The limits within which this scaling moves are individually configurable.

Amazon WorkMail and SES are required for receiving and sending mail, and CloudWatch alerts can be sent directly as mail via SNS, for example. The consistent use of IAM and careful handling of the assignment of rights to individual persons guarantee the highest possible level of security. The Key Management Service facilitates the strict encryption of data at all points and lambda functions are used to create regular backups of the hard disks. To reduce worldwide access times, CloudFront can be used as a CDN. Depending on the application, however, it may make more sense to maintain separate installations in different regions and to use internal application synchronization.

9. Monitoring, Alerting, Capacity Planning

With CloudWatch, AWS offers a central service that collects infrastructure, system and application metrics as well as log files at the operating system and application level.

The metrics of AWS services are provided automatically. EC2 instances also automatically record the data visible to the hypervisor. Fine granular information on system level, as well as data of the application must be transferred via custom metrics to CloudWatch. In addition to Java process data, JMX Beans are particularly suitable for this purpose. Especially the latter metrics are indispensable to get an insight into the application state and to react early to possible bottlenecks and problems and to be able to initiate countermeasures.

In addition to the metrics, log files of all applications are indispensable for error diagnosis. These can be transferred using the CloudWatch agent and are therefore immediately available for defined periods of time, even without direct access to the respective instance being necessary.

Based on the collected CloudWatch metrics, certain actions can be initiated. Events can be used to send a notification, for example, but Atlassian Data Center installations can also be implemented in an auto-scaling environment.

The long-term recording of CPU, RAM and memory requirements is the basis for achieving an optimal ratio between the required performance and the cost of the individual application.

10. Security & Privacy

But how is the topic data protection for an enterprise of the Atlassian products on AWS to be evaluated? In Europe, the topic of data protection (GDPR/EU-DSGVO) has played a very important role since 25 May 2018, but also before that. Companies take the regulations in the meantime very seriously and check the correctness of the implementation continuously. A toolchain based on Atlassian products is relevant for data protection for various reasons. Personal data of employees and external partners are stored on the platform. AWS offers a very comprehensive package of services around the topic GDPR/EU-DSGVO and can prove that all services of AWS are data protection compliant. Many Scandio customers use the security of a European AWS data center for this purpose, in particular the locations in Germany and Ireland.

Even more important than the formal confirmation of the data protection conformity is the actual technical implementation of the data security and the access protection based on it. Amazon Web Services offers deeply integrated functions for encryption, monitoring, logging and access control that go far beyond the standard in a medium-sized company. Partial access to applications from the Internet can be controlled just as easily. Multi-factor authentication can also be implemented with simple means.

In the AWS Shared Responsibility Model concept, safety and compliance are shared responsibilities between AWS and the customer. This common model makes operations easier for customers because AWS operates, manages, and controls the host operating system and virtualization layer components, and also provides physical security for the sites where the services run. Customer assumes responsibility for the guest operating system and its management (including updates and security patches), for other related application software, and for configuring the firewall provided by AWS for the security group.

11. Cost optimization through relief of infrastructure and capacities

Traditional OnPremise solutions require a hardware infrastructure with servers, storage and networks to operate. All devices must be housed in secure and expensive data centers. In addition, all hardware needs to be maintained, updated and regularly replaced when it becomes obsolete. These tasks are not necessarily within the scope of the customer's professional competence and represent high - often difficult to calculate - entry costs before they can be worked with effectively at all.

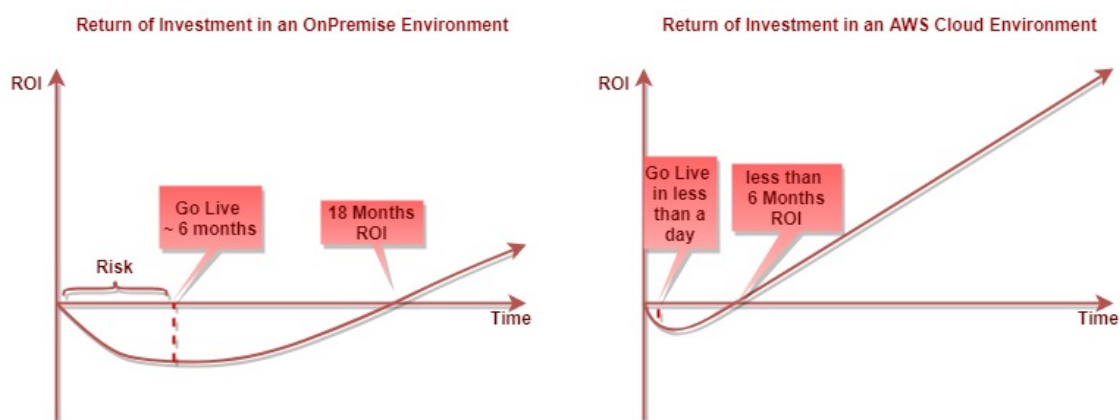


Figure 3 Comparison of the return on investment in an on-premise installation and a variant installed in the AWS Cloud

(Source: Scandio GmbH)

In Software-as-a-Service, on the other hand, all of the above considerations are addressed and managed by the vendor. Customers do not have to buy hardware or invest in high-availability or security solutions. All infrastructure and performance requirements are defined in Service Level Agreements (SLAs) with the AWS infrastructure provider, whose business and core competence is the provision of software services.

The cost of AWS services can be easily determined in advance with the AWS Calculator. It eliminates up-front costs and establishes a convenient pay-as-you-go model based on actual expenses rather than capital investments. By using AWS, companies do not buy hardware or virtual servers, but rent computing time. This consideration is enormously important. RAM, number of CPU cores and storage sizes can be adjusted at any time. This gives companies the opportunity to book and pay only for the capacity actually used. By means of several optimization measures, the costs of the infrastructure can be further reduced during operation. These are:

- **Reserved Instances** - By making fixed bookings for a period of 12 months or 36 months, costs can be reduced by up to 60%, depending on the service.
- **Spot Instances** - This can reduce the cost of the EC2 service by up to 90%. The only difference between on-demand instances and spot instances is that spot instances can be interrupted by EC2 with a two-minute notification lead time when EC2 needs capacity.
- **Continuous Monitoring** - When organizations continuously review actual computing power and storage needs, overcapacity is permanently reduced and is assumed to be at least 50% in the local data center.
- **Auto Scaling** - Using Atlassian's data center option, the smallest possible server nodes can be used. Additional nodes are only added when actually needed. Assuming two peak phases of two hours each per day, the costs can also be significantly reduced here, depending on time zone usage.

In general, it can be assumed that the costs of AWS services do not represent a significant increase in costs compared to the local data center, but they do represent an increase in flexibility. In addition, previously occupied capacities in the company are freed up, which can now concentrate on the core competence of the company and thus profitably contribute to the company.

12. Time saving and increased efficiency

Not only the monetary savings described above should be considered, but also the time saved when the applications are installed in the AWS cloud. Our own experience has shown that a live visit can be implemented faster than with on-premise solutions, since coordination with the internal infrastructure is no longer necessary and is now entirely the responsibility of the cloud provider.

On premise	AWS
Setup and maintain database cluster	Use the Relational Database Service and configure in detail with a few clicks
Setup and maintain a highly available Load Balancer	Use the Elastic Load Balancing Service and configure in detail with a few clicks
Run through internal processes to modify virtual machines	Easily change the instance type of your EC2 machines and RDS instances

On premise	AWS
A certain percentage of your infrastructure has to be idle to be available in case of emergencies	Only pay for what you use
Setup and maintain mail servers	Use the Simple Email Service for sending and WorkMail for receiving mails
Coordinate and maintain domains and certificates	Use self-renewing Amazon certificates together with Route53

Table 1 Comparison of selected resources in an on-premise environment and in the AWS cloud

These efficiencies continue throughout operations and can be applied to all maintenance and upgrades. Since the cloud application and the entire infrastructure are outside the enterprise, companies are not involved in maintenance or upgrades because they are performed exclusively by the cloud providers in their data centers.

The applications can be offered across the enterprise without regional commitment or additional costs. Because users only need one Internet connection to access the cloud application, the location becomes irrelevant. This is a big advantage for increasingly mobile employees and potential partner companies.

13. Supplier Information

13.1. Amazon Web Services (AWS)

In 2006, Amazon Web Services (AWS) began offering enterprise IT infrastructure services as Web services (now commonly known as cloud computing). Today, AWS offers a highly reliable, scalable, and cost-effective infrastructure platform in the cloud that runs hundreds of thousands of businesses in 190 countries around the world.



The AWS Cloud offers a wide range of infrastructure services, including computing power, storage options, networks and databases, delivered as utilities: On-demand, available in seconds, with pay-as-you-go pricing. From data warehousing to delivery tools, directories and content delivery, over 100 AWS services are available.

13.2. Atlassian

Atlassian was founded in 2002 with a mission to unleash the potential in every team. Atlassian’s suite of team collaboration software - Jira Software, Confluence, Bitbucket, Jira Service Desk and Portfolio for Jira - makes it easy for teams to organize, discuss and complete work. Today, Atlassian products help teams of all types and sizes in virtually every industry - from software and engineering teams to sales and marketing teams. For many years, Atlassian has supported large companies in introducing agile methods.



13.3.Scandio GmbH

Scandio is a software forge full of technology-enthusiastic consultants, developers and problem solvers. Based on our extensive know-how and many years of experience, Scandio has been successfully developing tailor-made solutions for its customers for 15 years. They also support companies in the areas of cloud strategy, UI/UX consulting and DevOps engineering, as well as automation, orchestration and agile development. Scandio Atlassian has been a Platinum Solution Partner for more than 10 years. In 2018, Scandio was voted Partner of the Year by Atlassian in the "Enterprise" category. Thanks to the long-standing partnership with AWS, Scandio's consultants were able to build up massive knowledge in both areas. Scandio's customers include BMW, BSH, Infineon and Siemens.

