

The Bitbucket Data Center guide for success



Table of Contents

3	Bitbucket Data Center: built for teams, optimized for scale
5	Understanding what's inside Bitbucket Data Center
9	Before you get started
12	Planning your deployment
22	Support & resources



01

Bitbucket Data Center:
built for teams, optimized for scale

Bitbucket Data Center: built for teams, optimized for scale

Bitbucket helps the world's best performing teams find bugs faster, collaborate with ease, and merge with confidence. But as organizations evolve in size and complexity, so does the need for their software. If Bitbucket is the center of your team's work and you need to ensure that performance and reliability scale efficiently across any distance, it may be time to upgrade to Atlassian's enterprise edition of Bitbucket, Bitbucket Data Center.

In this guide, we'll run through everything you need to know about making the switch.

Built for strategic administration

Data Center is a self-managed enterprise edition that gives you the reliability and compliance your organization demands, with the infrastructure you need to stay ahead of team growth.

In addition to all the Server features your teams know and love, Data Center offers additional benefits in the areas of:



Security & compliance



Performance & reliability



Infrastructure & automation

Data Center paves the way for organizations

INNOVATE AT SCALE

High-performing teams depend on Atlassian to get stuff done, at scale (and fast).

Data Center empowers teams to drive success with the horizontal scalability required for seamless collaboration.

GROW FAST WITH CONFIDENCE

Scale with confidence and build a long-term strategy for scaling teams, applications, and infrastructure.

TRUST YOUR TOOLS

In the face of tough compliance, usage, and geographic constraints, your business can depend on our tooling and practices to get stuff done, at scale.

At the end of this guide, we share the story of a fictional company, Streaming+, on their journey to upgrading to Bitbucket Data Center. This story is intended to highlight how the tools shared in this guide can be applied to any organization.



02

Understanding what's inside Bitbucket Data Center

Understanding what's inside Bitbucket Data Center

While Bitbucket Server fits many development teams' needs, as organizations continue to grow, take on greater complexity, hire remote developers, and encounter new compliance requirements, the software needed to support those teams goes beyond the traditional server solution.

Enter Bitbucket Data Center. Built with large organizations in mind, Data Center has the security, reliability, and performance that enterprises depend on to move work forward. Here's a brief overview of the features that differentiate Bitbucket Data Center from our Server solution:

Security & Compliance

Integrating security into the software development lifecycle is often an unrealized ideal. While security, compliance, and regulatory standards are often top of mind for many business leaders, it can be tough to embed good security practices into the way individual teams work. To address some of the manual procedures and lack of visibility that admins often deal with, we've added a variety of solutions that are built into the fabric of Bitbucket Data Center to ensure teams align with the internal and external demands for security.

USER MANAGEMENT

Comply with security standards and streamline user login with built-in SAML or OpenID Connect single sign-on support. Data Center supports these authentication protocols as well as several IdPs:

- Windows Azure
- Windows
- Onelogin
- Okta
- PingIdentity
- Bitium

Don't see your IP in the list? Don't worry, you can also set up a [custom SAML or OpenID connection](#).

ADVANCE AUDITING

Track and log actions in an instance, developing a security-relevant chronological record that documents your organization's digital footprint. Advanced Auditing includes a range of features that streamline an organization's ability to demonstrate compliance, security strength, and manage risk while maintaining the right level of visibility to monitor the state of the business.

Performance & reliability

While every company is at a different stage of their digital transformation, there is consistently a need for tools that strengthen collaboration and provide reliable performance. In an era of distributed teams and remote work, that need is amplified, with teams depending on their software tools to keep work moving forward and productivity on track. In Bitbucket Data Center we offer high availability, CDN support, and Smart Mirrors to increase performance and deliver a consistent experience to teams wherever they're located. Pair that with rate limiting and disaster recovery that implement preventive measures to improve the stability of your instance, Data Center offers enterprise-level performance, not available on Server.

HIGH AVAILABILITY

Active clustering ensures users have uninterrupted access to critical applications. Data Center uses industry-standard load-balancing, database clustering, and a shared file system to prevent downtime in the event of unexpected hardware failure.

RATE LIMITING

Protect your instance from threats that can flood your REST API and cause performance degradation and downtime. With rate limiting, you can control the volume of requests from automated integrations and scripts, making sure that your Bitbucket instance remains stable.

SUPPORT FOR CDN

Deliver faster load times for geographically distributed offices by using a Content Delivery Network (CDN), a globally distributed network of servers that chase static resources locally. CDNs speed up response times by reducing the physical distance between a server and your team.

SMART MIRRORS

Improve git clone speeds and increase CI/CD capacity for distributed teams by setting up one or more live mirror nodes that host copies of repositories in remote locations. Teams can also cluster mirrors into “farms” grouped behind a load balancer to reduce time spent waiting for builds.

GIT LFS AND STORAGE SHARDING

For teams with large repositories, Git LFS and storage sharding can improve how large files are handled to streamline updates, improve clone speed, and facilitates improvements to the performance of your instance.

DISASTER RECOVERY

In the event of a disaster, Bitbucket Data Center has you covered with a [disaster recovery strategy](#) that allows you to failover to a standby system. This means your Bitbucket instance, and more importantly your Bitbucket managed data (source code), will remain online even as your primary system is unavailable.

Infrastructure & automation

To support organizations that are on the path to cloud or would like to reduce the administrative overhead of an on-premise solution, Data Center gives teams the ability to deploy Bitbucket on infrastructure as a service (IaaS). Moving to an IaaS environment can increase the speed, productivity, and innovation on your team by redeploying resources that were previously focused on managing your servers to other parts of your business. Bitbucket Data Center also streamlines the deployment process by offering Quickstart templates that include preset configurations and recommended settings for the leading IaaS providers, Amazon Web Services (AWS) and Microsoft's Azure.

AMAZON WEB SERVICES

Leverage our partnership with AWS to deploy with confidence and unlock added capabilities, like our support for Amazon Aurora. Aurora boosts stability and reduces downtime with a fault-tolerant relational database built for cloud.

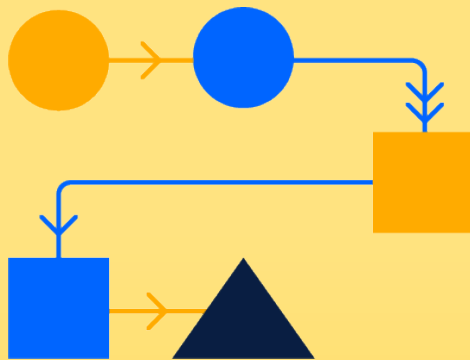
AZURE

Deploy Data Center on Azure to scale quickly with minimal overhead and reduce infrastructure maintenance and costs.

DOCKER

As enterprises grow and become more geo-distributed, they inevitably encounter dissonance between setups. Reduce administrative overhead by leveraging our Docker container images for agility and consistency in your deployment, regardless of what hardware you're using.





03

Before you get started

Before you get started

As with any new investment the team is embarking on, you want to start by creating a plan - a path to upgrade that is as painless as possible. We hope this guide is the first step in that mission!

We'll help you evaluate how you want to deploy, what type of infrastructure you want to use, and an understanding of the necessary requirements. As the project lead, you'll also want to bake in any unexpected interruptions or roadblocks to success, anything from red tape to budget constraints that may come up as part of this process.

We're confident that once you take the leap, Bitbucket Data Center will give you, your team, and organization all the tools to be successful.

Get to know the latest Bitbucket release

Before you decide to download Data Center you should familiarize yourself with the latest Bitbucket release. Beyond the new features, you'll access through Data Center, you'll want to prepare your team for all the changes that have been made between your current version and the new one you're moving to.

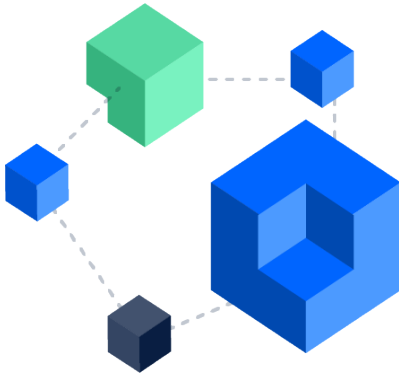
For larger, more complex instances or for organizations that prefer to upgrade only 1-2 times a year, we suggest upgrading to Long Term Support releases (formally known as Enterprise releases). These releases will continue to receive backported security updates and critical bug fixes throughout the standard 2-year support window and are intended for customers that prioritize stability over the latest functionality.

Performance recommendations and requirements

Depending on the size of your instance and the performance you're looking for, you may need to invest in additional hardware to maximize all the benefits of Bitbucket Data Center. For example, if achieving true availability is a top priority for your team you'll want to move to a clustered approach, which requires additional hardware for the application nodes, shared file system, and load balancer.



You can find a full list of system requirements on our documentation [site](#).



On-premise vs. infrastructure as a service

Data Center gives teams the flexibility to choose the infrastructure that best suits their needs. Your team can choose to deploy on-premise or leverage the cloud computing features and capabilities available through infrastructure as a service (IaaS) providers. With 42% of Data Center instances deployed in the cloud, we've built additional capabilities, like our Quickstart templates, to support teams as they move to the cloud.

Review your disaster recovery strategy

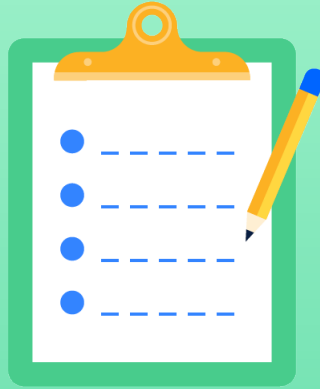
A crucial step in ensuring a successful implementation and that your business will continue running smoothly is reviewing (or creating) a disaster recovery plan. With disaster recovery for Atlassian's Data Center products, you can configure index replication, attachment updates, and database synchronization to comply with your business's specific requirements.

We also recommend creating a runbook that is detailed enough that anyone on your team can complete the steps and recover your service, regardless of prior knowledge.



A well prepared runbook should include:

- Detection of the problem
- Isolation of the current production environment and bringing it down gracefully
- Synchronization of data between failed production and intended recovery point
- Warm up instructions for the recovery instance
- Documentation, communication, and escalation guidelines



04

Planning your deployment

Planning your deployment

Now that you've decided that Bitbucket Data Center is right for your team, it's time to get into the details. We've put together a step by step process that will make the transition to Bitbucket Data Center as seamless as possible.

1. Build a team

Enlist the right people to support this project and get their buy-in early and regularly throughout the process. An ideal Data Center deployment is a fully-fledged project with defined roles and responsibilities that span multiple teams.



Key roles involved in making this move successful

EXECUTIVE SPONSOR

Given the scope of migrating to Data Center, it will be useful to have executive buy-in for the project.

Responsibilities:

- Ultimate decision-maker in the event of roadblocks
- Facilitates all budget approvals

PROJECT LEAD

Project lead has a deep understanding of the business and how Data Center will align with those objectives. They also know how to make the right tradeoffs to maintain governance across applications.

Responsibilities:

- Keeps the project on track with key milestones and estimated dates to achieve them
- Owns the schedule, ensures task completion, and resolves cross-functional issues
- Communicates project updates to stakeholders and announcements to end-users
- Works with a procurement representative who performs the purchasing of Bitbucket Data Center



SYSADMIN

The Sysadmin handles everything from the infrastructure to the application's interface. They are concerned with backups, storage, network, and performance.

Responsibilities:

- Ensures Bitbucket can connect to the database through multiple nodes, instead of a single server
- Addresses L1 support issues that may arise during migration
- Gathers the hardware (physical or virtual) needed
- Performs the installation of Bitbucket Data Center
- Ensures that if you choose to deploy in a cluster, that all the components are working properly. This includes configuring the load balancer, making sure that the database can connect to multiple nodes, and setting up the file system.

APPLICATION ADMIN

The Application Admin handles the day-to-day administration. They have in-depth knowledge of the product, care about performance, reliability, and evaluate and maintain Marketplace apps. They may also work closely with end-users to understand their needs and provide assistance or training.

Responsibilities:

- Verifies functionality and performance during testing to ensure that the Data Center application is operating properly
- Manages all the Marketplace apps
- Ensures that the users and permissions are maintained, or changed, correctly throughout the transition

In addition to these roles, you may need support from:



DATABASE MANAGEMENT

Database experts ensure the database is running smoothly. They understand the complexities and tuning required to operate Bitbucket Data Center at scale.



NETWORK ENGINEERING

They analyze connectivity and update networking requirements for Data Center. This team can also provide expertise and guidance on selecting and configuring the load balancer if the business decides on a clustered environment.



SITE RELIABILITY

Site reliability experts ensure that Data Center properly operates and scales with the rest of the organization's technology infrastructure. They also help ensure uptime and bring systems back online in the event of an outage.



SECURITY

This team is crucial in ensuring that Data Center is set up for success. They can review the plan, offer recommendations, and provide expertise in making sure that Bitbucket Data Center is safe and complies with all the internal and external regulatory requirements.

Once you have your team in place, you'll want to make sure that all stakeholders align on the business, functionality, and performance goals for the migration. Proper upfront alignment ensures a smoother installation, testing, and release process.



2. Review your current installation and document everything

Just like with any project, before you can mark this deployment as a success, you'll have to evaluate the impact and improvements Data Center has made for the business. To do that confidently, you'll want to take a baseline measurement of your system's functionality and performance before implementation.

As part of this assessment, you'll also want to make sure you've made all the optimizations possible to your server instance. This extra due diligence will not only simplify the upgrade process, but this will also signal to your leadership team that the decision to move to Data Center was well thought out and necessary for the business to continue to scale effectively.

Next, you'll want to assess how users interact with Bitbucket and if there are any usage characteristics that should be addressed before migrating. For example, users that run frequent reports can place strain on the system, causing performance degradation. It may be necessary to place certain restrictions on scripts that make REST calls and other kinds of integrations. Specifically, scripts that execute many write operations can impede cluster performance as these write operations will sync across all nodes.

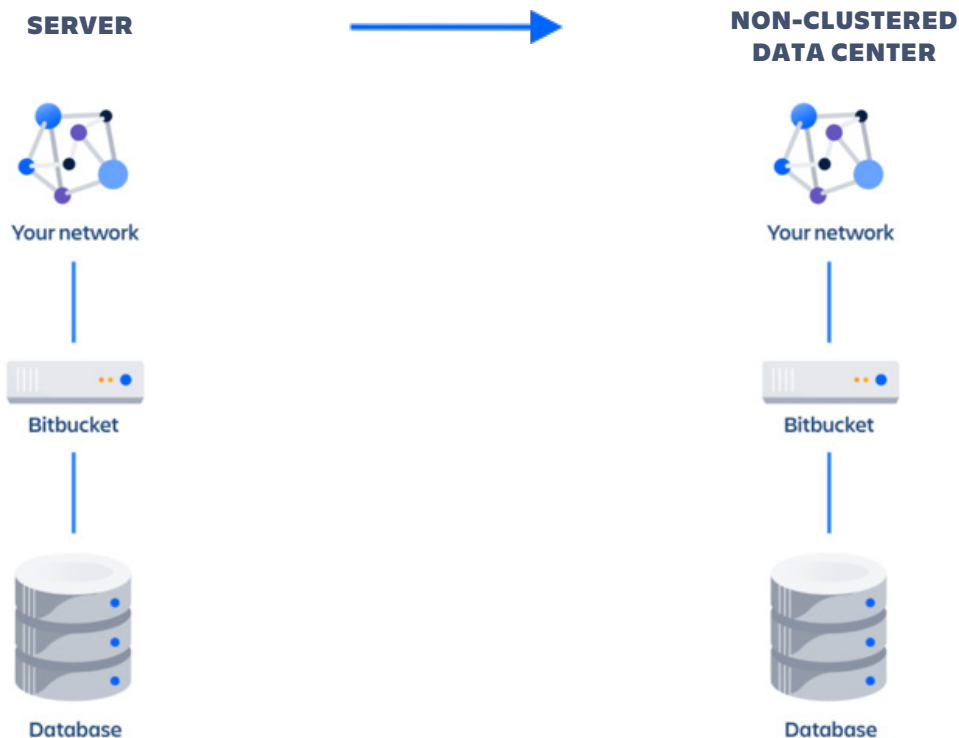
Last but not least, you'll want to evaluate all the [Marketplace apps](#) that are being used by the team. A thorough analysis includes reviewing the number of apps that are being used, identifying redundancies, and finding the Data Center approved versions of your essential apps. Data Center approved apps undergo a rigorous technical review process (run by Atlassian). During this process, an app gets tested for performance, stability, and overall suitability for Data Center environments.

And throughout this entire process, you'll want to create a runbook and document everything you've learned about your existing instance and the needs of the team. The runbook will help guide your configuration choices, influence decisions, and identify whether issues observed after deployment are new or existed before deployment. Some items to note include general system behavior benchmarks and API access patterns, as well as backup, reporting and monitoring processes, and any maintenance or disaster recovery plans.

3. Evaluate technology decisions

One of the biggest decisions that teams have to make on their path to upgrading is the type of hardware and infrastructure they'll want their new Data Center instance to run on. Bitbucket Data Center allows you to choose from either a clustered or non-clustered environment.

For customers where high availability isn't their top priority and would prefer to retain the same architecture they have in their Server instance, the non-clustered approach is the way to go. It can take as little as a few minutes to upgrade to Data Center with a non-clustered environment and provides many of the enterprise-grade features to teams that were not available on their server instance (ex. Rate Limiting, SAML, OpenID support, and Advanced Auditing).

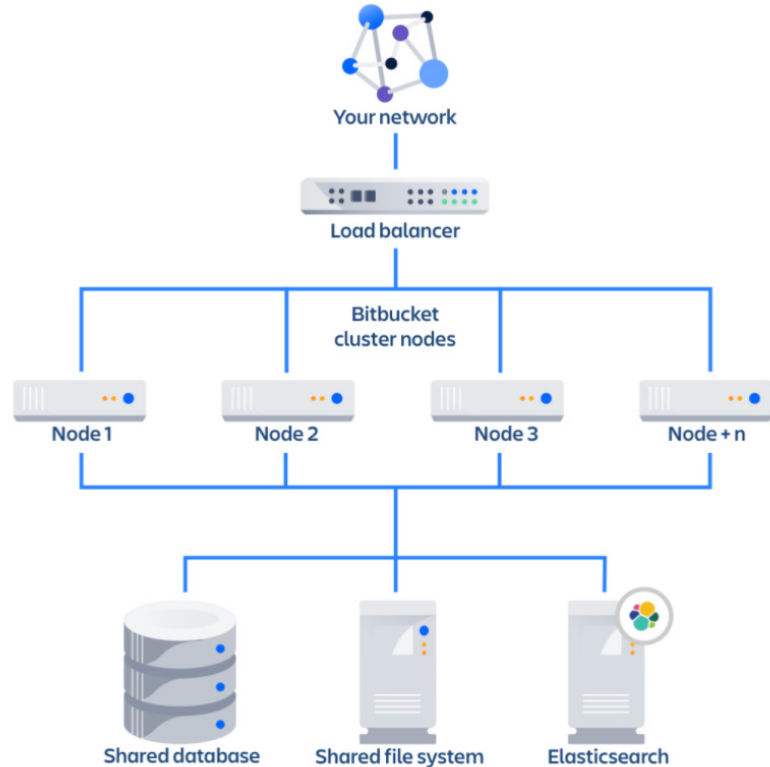


For teams, where high availability is one of the main reasons for switching to Data Center, you can set up a clustered environment where Data Center sits behind a load balancer. Each node in your cluster can operate as an active instance of your Bitbucket application. If a node fails, the load balancer will automatically direct new requests to any remaining active nodes in the cluster; this can minimize downtime or potentially avoid it altogether. Smart Mirrors and Smart Mirror Farms are only available in a clustered environment.

NON-CLUSTERED DATA CENTER



CLUSTERED DATA CENTER



Key Components of a Clustered Data Center

LOAD BALANCER

A load balancer distributes incoming user requests across your cluster to minimize response time and avoid overloading any single node. The load balancer also returns the response from the selected server to the user. For Data Center applications, the load balancer serves three essential functions:

- Distributes traffic efficiently across multiple nodes
- Ensures high availability by sending traffic only to nodes that are online (requires health check monitoring)
- Enables the ability to add and remove nodes

APPLICATION NODES

Cluster nodes share the workload of incoming requests, and if a node fails, the load balancer will send requests to the remaining nodes in the cluster.

SHARED DATABASE

When setting up your shared database, make sure that it allows at least the number of maximum connections across all nodes. For example, if you have three nodes that can each support a maximum of 50 connections, then your shared database must allow at least 150 connections. Bitbucket Data Center also requires you to use an external database. You can not use Bitbucket Server's internal HSQL or H2 database with Bitbucket Data Center.



Tip: All the usual database vendors in Bitbucket Server's supported platforms are supported by Bitbucket Data Center, with one exception: we do not recommend MySQL due to inherent deadlocks that can occur in this database engine at high load.

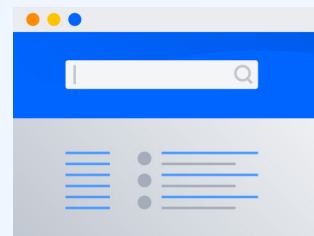
SHARED FILE SYSTEM

All clustered deployments require a shared file system, such as network-attached storage via a file system protocol, such as NFS. All application nodes should have access to a shared directory in the same path. Bitbucket Data Center requires a high performance shared file system such as a SAN, NAS, RAID server, or high-performance file server optimized for I/O. For more information on setting up Bitbucket Data Center's shared file server, check out this [page](#).

ELASTICSEARCH NOSE

Bitbucket Data Center requires a remote Elasticsearch installation to enable code search. The easiest way to set up and deploy an Elasticsearch cluster for Bitbucket Data Center is to use the Amazon's Elasticsearch service, but you can also set up a remote Elasticsearch instance on your own hardware.

For a detailed list of requirements, see [our documentation site](#).



4. Assess your requirements

If your team chooses the clustered approach, the best way to realize the full potential of Bitbucket Data Center is to tailor your set up based on the needs of your organization. Here are a few examples of things that have come up when reviewing some of our customers' installations:

- Evaluate if a cluster with fewer larger nodes or a cluster with more smaller nodes is a better fit for your environment.
- Remember to account for cascading failure when determining node size. This occurs when an instance fails, and the remaining instances cannot handle the incoming load.
- Consider traffic segmentation so that specified nodes exclusively handle specific kinds of traffic.
- Investigate virtualization and automation tools such as Docker, Chef, and Puppet to better manage scaling the cluster.
- Consider investing in additional monitoring tools such as New Relic, Splunk, and ELK.
- Ensure existing monitoring tools are reconfigured to handle multi-node environments.

To help you get started, we've provided a Bitbucket Data Center sample deployment and monitoring strategy that shares how we use and monitor our Data Center application.

Install and configure Bitbucket Data Center on a test environment

Now that you've reviewed all the requirements and made the appropriate decisions, it's time to install and configure Bitbucket Data Center on a test environment.

HOW YOU INSTALL DEPENDS ON YOUR ENVIRONMENT

Option A If you're installing on your own hardware and prefer to take the non-clustered approach, the upgrade is as easy as applying your Data Center license. If you are [new to Bitbucket](#), you can choose to download using an installer from a zip or archive file or using a Docker container. For those that want to install Bitbucket Data Center in a cluster for optimal scalability and performance, you can start by applying your new license and installing the application on your first node. Once you've installed the first Bitbucket application node, you will then provision the shared database, shared filesystem, and elastic search node to use with Bitbucket Data Center. Then you'll add more Bitbucket nodes and configure your load balancer. Once the setup is complete, you can start Bitbucket one node at a time before testing your cluster.

Option B If you're using an IaaS platform, Bitbucket Data Center currently supports deployments on AWS and Azure. Use our Quick Start guides to launch, configure, and run the AWS compute, network, storage, and other services required to deploy a specific workload on AWS, using AWS best practices for security and availability. Or deploy in Azure with the Atlassian template for Azure Resource Manager. This platform allows you to scale your deployment elastically by resizing and quickly launching additional nodes. It also provides a number of managed services that work out of the box with Bitbucket Data Center.



Infrastructure recommendations for AWS can be found on our [documentation site](#).

Test your Data Center installation, then launch to production

Whether you're deploying on your own hardware or using an IaaS platform, you should run through an iterative set of functional tests, integration tests, and performance tests in order to deploy to production confidently. Here is where the team you've set up plays a significant role, not only will they help you define the metrics for success, but they will help implement each of the tests and determine if Bitbucket is production-ready. Through the process of testing, you and your team will determine where you need to refine your configurations and harden procedures governing application usage. Once the team is in agreement and you've tested, refined, and documented the entire process, you are ready to launch Bitbucket Data Center on production!



Pro Tip: Organizations typically push to production over a weekend or some other time of low usage. Whenever you decide to make the switch, make sure that notifications are sent to end-users throughout this process. As one of your biggest stakeholders, it's important they stay informed and engaged throughout the process.

Maintenance

To keep your instance in tip-top shape, you should plan to do regular health checks, review how users are experiencing Bitbucket Data Center, and assess if there are opportunities to optimize or upgrade certain elements of your environment. Regular testing and monitoring of performance is vital to ensuring the integrity and reliability of your instance. We've compiled a [list of tools](#) we see many of our customers using.



05

Support & Resources

Support & Resources

At Atlassian, we believe that great teamwork requires more than just great tools. In addition to our solutions and proven practices, we also offer great support. Teams that help ensure customer success at every step of their Atlassian journey.

CUSTOMER SUCCESS MANAGER

As a new Data Center customer, you have access to a Customer Success Manager as an ongoing resource throughout your first year. They can discuss your organization's goals, help define success criteria, and provide you with the right guidance to help you get up and running quickly.

ATLASSIAN COMMUNITY

Find answers, support, and inspiration from other Atlassian users. We recommend joining the [Enterprise community group](#) for stories, tips, and best practices for using Atlassian products at scale.

TECHNICAL ACCOUNT MANAGEMENT

Think of a Technical Account Manager as your strategic advisor for all things Atlassian. Backed by the resources of the entire Atlassian organization and armed with best practices, Technical Account Managers help you get the most out of Bitbucket Data Center, and your broader Atlassian investment.

PREMIER SUPPORT

Atlassian Premier Support offers our highest level of support, with 24/7 access to a team of dedicated senior support engineers. Premier Support can work hand-in-hand with Technical Account Managers and Solution Partners to save your team valuable time.

The Premier Support team can also perform health checks by analyzing your application and logs, to help prevent outages and ensure that your deployment fully meets the needs of your users.

SOLUTION PARTNERS

Atlassian Solution Partners offer a range of services for enterprises, including system integrations, deployments, migrations, and version upgrades. They're a great option if your organization has complex requirements, or if you need dedicated, hands-on help with your Data Center deployment.



To chat with a representative to learn more about our services or Bitbucket Data Center, contact the team [here](#).

Unlock benefits across the Atlassian suite

At Atlassian, we believe in the power of teamwork. Behind every great human achievement, there is a team, and we want to help unleash the potential of every team. To do that, we don't only create amazing products, but we want to create amazing products that streamline collaboration across teams. Here are two products that have built-in integrations with Bitbucket:

Jira

Jira Software helps the world's best agile teams plan, track, and release great software at scale. Teams that integrate Jira Software with Bitbucket release versions 14% more often than those who don't. You can interact with Jira issues in Bitbucket, automatically transition issues as work is completed in Bitbucket and consolidate user management across these tools. To learn more about the benefits of this integration, check out our [documentation site](#).

Bamboo

Bamboo is the choice of professional teams for continuous integration, deployment, and delivery. When you integrate Bamboo with Bitbucket, commit, branch, build, and deployment information is shared for users of both applications. These solutions streamline the software development workflow and make it easier to ship quality code faster. To learn more about the benefits of this integration, check out our [documentation site](#).

If you're ready to start the conversation about whether Bitbucket Data Center is the right fit for you, reach out and [set up a chat today](#).

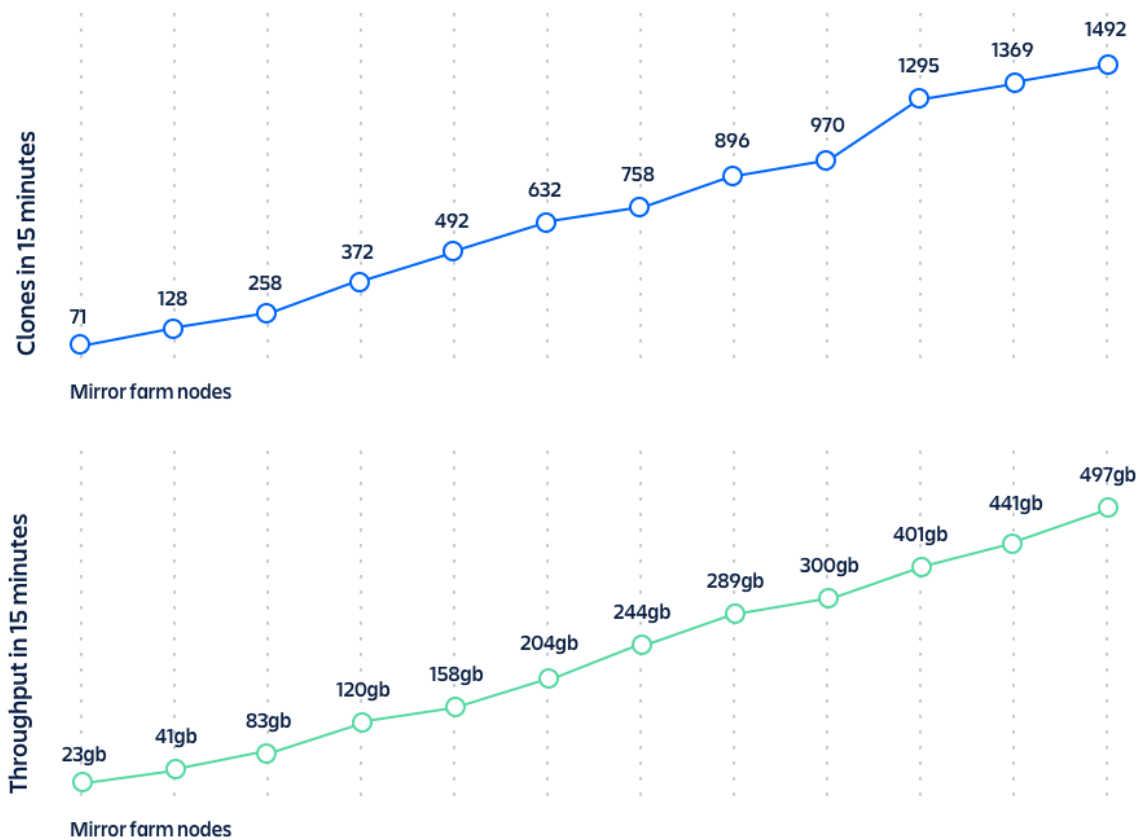
How Streaming+ Stays Focused on Innovation with Bitbucket Data Center

Streaming+ is a media-services provider that offers an online video library of content from all over the world. In the last few years, they've seen an explosion of growth with over 15 million subscribers and a development team that has gone from 150 to 400+ employees. During this period of hyper-growth, they've begun to notice a decrease in performance, more incidents of downtime, and complaints of unreliable experiences from teams in different locations.

Streaming+ has relied on Atlassian products for several years now, first adopting Jira and then choosing Bitbucket to retain an integrated toolset. As these challenges began to occur more regularly, they searched for a solution that could carry the burden of a larger team.

Bitbucket Data Center was ultimately chosen because it provided the high availability, scale, and strategic administration that their server solution lacked. Additionally, choosing Data Center made moving to an IaaS an obvious deployment choice. Amazon Web Services (AWS) not only offered greater cost savings but unlocked the added benefit of support for Aurora, a fault-tolerant relational database that could boost stability and reduce downtime.

The Streaming+ team also decided on a clustered approach, setting up multiple nodes to limit incidences of downtime and increase performance across the team. They were also able to leverage Smart Mirror Farms, to improve their git clone speeds, and increase their CI/CD capacity. They noticed that as they increased the number of mirror farm nodes, the load balancer throughput also increased linearly.



Beyond the added capacity and performance, the Streaming+ team implemented new user management and advanced auditing features to improve security, demonstrate compliance, and make it easier to monitor potential risks. “Thanks to Bitbucket Data Center, our team has the tools to focus on innovation without worrying about scale”, says Mark Nielsen, Head of Engineering for Streaming+.

To date, the Streaming+ team has had 75,000+ commits and 432+ contributors, across 30 countries, using Bitbucket Data Center. As the team has grown, the Atlassian products have continuously adapted to meet the needs of their modern, distributed teams.



Try it for free

[Learn more](#)