




The Ultimate Guide to Atlassian Data Center

From technical overview
to installation

Introduction

Have you reached a point where your single server instance is just not quite enough to support your organization's level of growth or maturity? Or maybe you are just starting to notice that you are nearing that point and need help developing a plan – and fast – to get ahead of your organization's growth. We built Atlassian Data Center, our self-managed enterprise edition, to support our Server customers in their next stage of maturity and growth. This guide will give you everything you need to get up to speed on Data Center and how it can support you as you scale. You'll get an overview of its architecture, a deep dive into Data Center features by product, and a rundown on what it takes to prepare for and successfully upgrade to Data Center.



Contents

2	Introduction
4	What's the difference between Server and Data Center?
5	What features are in Data Center? <ul style="list-style-type: none">Core CapabilitiesProduct Specific Features<ul style="list-style-type: none">- Jira Software & Jira Service Desk- Confluence- Bitbucket- Crowd
8	What is required to set up Data Center? <ul style="list-style-type: none">Application nodesDatabaseLoad balancerFile system
10	How do I migrate to Data Center? <ul style="list-style-type: none">Step 1: Build a project teamStep 2: Review your existing Server environmentStep 3: Document current processesStep 4: Evaluate technology decisionsStep 5: Implement and test your deployment process
20	Resources & Help

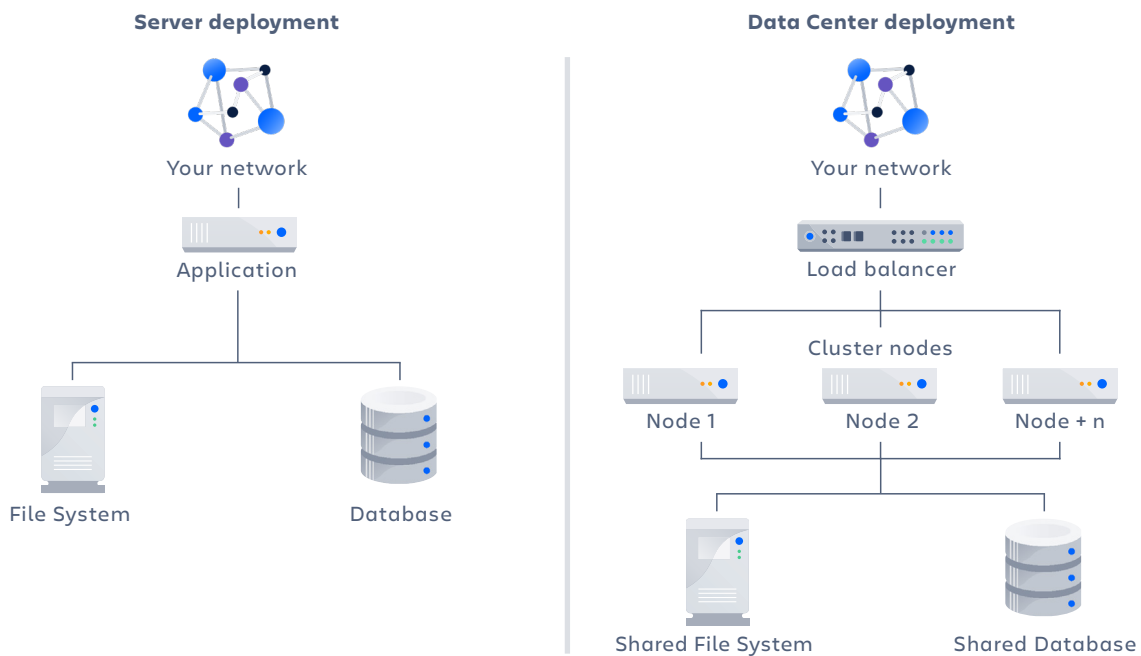
What's the difference between Server and Data Center?

Atlassian Data Center products are designed to support customers running in mission-critical environments and those with growing requirements, particularly around scaling administration and management of the tools.

COMPARING ARCHITECTURES

One of the key differences between Data Center and Server environments is their architecture. In a Server environment, there is a single node and if that node fails, your Atlassian application becomes unavailable to users.

Unlike Server products, Data Center products can be run in a multi-node environment, behind a load balancer, where each node in a cluster can essentially operate as an active instance. If a node fails, the load balancer will automatically failover sessions and direct new requests to any remaining active nodes in the cluster – minimizing downtime or completely avoiding it. By distributing load across multiple nodes, performance will improve. It also offers allowing for tactics like “traffic shaping”, which routes common, load-heavy requests through a single node – freeing up the remaining nodes for other functions. And, since there's no limit to the number of nodes you can run in a cluster, your environment can scale as your business grows and the demands on your Atlassian applications increase.



A deep dive into Data Center features

Beyond the difference in architecture, there are many other benefits that Data Center provides to meet the needs of our maturing Server customers.

The core capabilities

HIGH AVAILABILITY

Ensuring high availability is especially critical when your teams rely on Atlassian tools to get their work done. Incidents happen. Whether it's as a result of natural or human-induced circumstances, downtime causes productivity and revenue losses. Avoiding any potential downtime is critical to success. With Atlassian Data Center's active-active clustering model, you can ensure your teams always have access to the tools they need to do their job.

HORIZONTAL SCALABILITY

Beyond high availability, Data Center also allows you to scale. With Data Center, you can grow a cluster horizontally by adding more nodes. The load balancer routes and distributes incoming traffic across the nodes in the cluster depending on availability, allowing the instance to support more concurrent users.

SAML 2.0

A [SAML for Atlassian Data Center](#) A SAML for Atlassian Data Center app is embedded in each of the Data Center products. It allows each application to connect to an IdP to provide single sign-on (SSO). It's important to note that it only handles authentication. Application access and any required authorizations, such as ensuring that users belong to the appropriate groups/roles and have the necessary permissions, should be configured in the user directory and/or the application itself. We currently perform tests with the following identity providers:



INFRASTRUCTURE CHOICE

We are finding that more and more of our customers are choosing to deploy Atlassian tools using infrastructure as a service (IaaS) providers. In fact, 62% of Atlassian's self-hosted customers are choosing to deploy their applications on a virtual architecture. You have the option to deploy Atlassian Data Center using Amazon Web Services (AWS) or Microsoft Azure.

By deploying using one of these vendors, you can scale your deployment elastically by resizing and quickly launching additional nodes. These vendors also provide a number of managed services out-of-the-box with your Data Center instance and automatically handle all the configuration and maintenance.

DISASTER RECOVERY

Disaster recovery plans should be created for every tool an organization relies on. With the Data Center versions of Jira, Confluence and Bitbucket, your team can set up cold standby instances in the event of a disaster. As long as you have an active Data Center license, you will not need to purchase additional or separate licenses to start up Jira, Confluence or Bitbucket on a cold standby instance.

To get started, your team will need to create a standby system without starting it up. Then, you'll need to implement a data replication strategy. Replicating data to your standby location is crucial to a cold standby strategy. You don't want to failover to your standby instance and find that it's out of date or that it will take many hours to re-index your database(s).

DATA CENTER APPROVED APPS

Many of you have connected hundreds of Atlassian Marketplace apps to your various Atlassian instances. As critical as it is to ensure the availability and performance of your Data Center instances, the same holds true for your apps.

Before an app can be approved for Data Center use, Marketplace vendors are required to build and test their apps to a more rigorous, standard. This includes testing how they handle cache operations, whether they support required databases, if you can implement locking and availability in clustered environments, manage event handlers, and much more.

All test results must be submitted to Atlassian for a complete review and verification. Only the apps that have met these requirements will receive a ["Data Center"](#) listing on the Atlassian Marketplace.

Data Center features by product

Beyond Data Center's core capabilities, each product contains additional features that support the specific needs of their customers.

Jira Software and Jira Service Desk

PROJECT ARCHIVING

Reduce clutter in Jira and optimize performance by cleaning out irrelevant, outdated projects in just a few clicks.

ZERO DOWNTIME UPGRADES

Eliminate downtime by upgrading one node at a time - letting end users continue to work without interruption throughout the entire upgrade process.

CUSTOM FIELDS OPTIMIZER*

Scan your instance for custom fields that are hindering performance and remove their global contexts with just one click.

*The custom fields optimizer is only available in Jira Software currently. Coming soon to Jira Service Desk!

Confluence

READ-ONLY MODE

Minimize downtime and give your team access to content during maintenance or upgrades by turning on read-only mode.

Bitbucket

SMART MIRRORS

Improve Git clone speeds for distributed teams working with large repositories by setting up live mirror nodes, which are copies of repositories in remote locations.

INSTANCE CONSOLIDATION

Easily consolidate multiple Bitbucket instances or move from a Server instance to a Data Center instance with this migration tool.

ZERO DOWNTIME BACKUPS

Back up your instance as frequently as you need to without locking it for maintenance or inconveniencing your users and build agents with frequent downtimes.

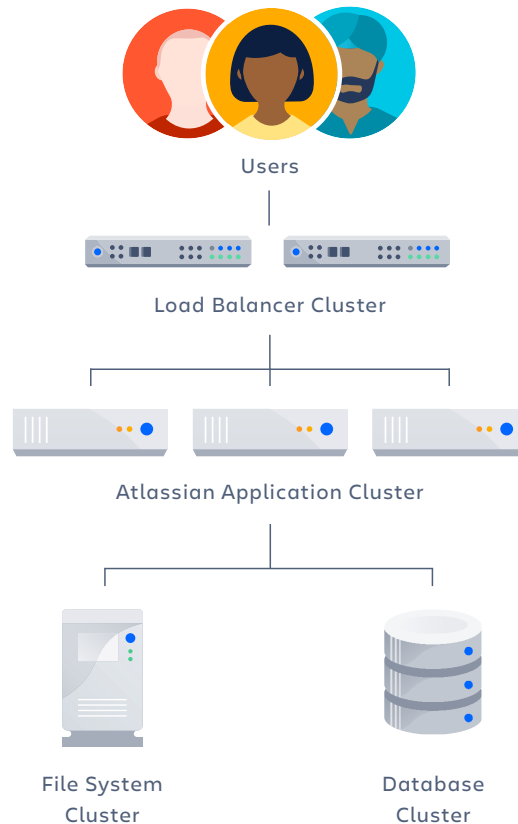
Crowd

DELEGATED ADMINISTRATION

Eliminate bottlenecks and save time by assigning group level administrators to add or remove users to groups.

What is required to set up Data Center?

You now know that Data Center is deployed in a multi-node environment, but what exactly does its architecture look like? There are four main components that make up a Data Center environment.



LOAD BALANCER

The load balancer is the first thing that your users' requests hit when they use an Atlassian Data Center application. Requests come into the load balancer and the load balancer then distributes each request to the application nodes. Data Center supports both hardware and software-based load balancers. For both software and hardware solutions, the load balancer should be connected to the application cluster using a high-speed LAN connection to ensure high bandwidth and low latency. All software load balancers should run on dedicated machines.

Data Center applications assume that each user's request will go to the same node during a session. If requests go to different nodes, users may be unexpectedly logged out, and they could even lose information stored in their session. Therefore, it is required to bind a session to the same node by enabling cookie-based "sticky sessions" (or session affinity) on the load balancer. When using cookie-based sticky sessions, you can use the cookie issued by the Atlassian application, or you can use a cookie generated by the load balancer.



Pro Tip: Add an extra layer of protection and prevent the load balancer from becoming a single point of failure by adding redundancy to your load balancing solution. You can do this by setting up two load balancers in an active-passive configuration, using a virtual IP address across both load balancers. If the active load balancer fails, it will fail over to the passive load balancer.

APPLICATION NODES

Application nodes are where the actual Atlassian application lives. Each node in your Data Center cluster must run on the same version of the application and be located in the same physical location. These nodes should be configured in a cluster, acting as one, to serve the application to your users. The number of nodes in your cluster depends on your needs and how you configure your application. Typically, we find that between 2-4 nodes is sufficient for most clusters, but use our [node sizing guides](#) to help you make the right decision.

Important Note: Confluence requires an additional application node specifically dedicated to Synchrony, the engine that powers collaborative editing in Confluence, and Bitbucket requires an additional application node specifically dedicated to Elasticsearch, which enables code search.

DATABASE

The database must be installed on its own node, separate from the application. Odds are you'd be able to use the same database configuration in Data Center that you're currently using with your Server instance. Databases that Atlassian supports include: Postgres, Oracle, MS SQL Server, and MySQL.

Important Note: MySQL is not supported for Bitbucket Data Center. Check the supported platforms page for each of your Data Center products for an up-to-date list of supported databases.

FILE SYSTEM

The shared file system is where all the underpinnings of the product are stored. This is where things like your attachments, icons, user information, apps, and source code lives.

In a Data Center environment, you need to set up your shared file system as its own node. You can use any NFS based NAS or SAN program for your shared file system. Just be sure to stay away from distributed protocols like DFS, as these are not supported.

How do I migrate to Data Center?

When you finally decide it's time to move to Data Center, you'll need to put together a migration plan. You'll want to front load as much staffing, planning, and decision-making as possible to manage problems as they arise throughout the process. A Data Center migration is one of the following:

- **In-place:** where a single Server application is deployed as a multi-node Data Center application
- **Consolidated:** where multiple Server applications running on isolated nodes (for example, different instances of Jira being used by different teams) are deployed as a single multi-node Data Center application

To make things easy, we've put together a reference of an in-place migration plan to help you develop your own.




Step 1: Build a project team

Enlist People

One of the most important parts in a Data Center migration is assembling the right team, as early as possible.

When assembling your Data Center migration team, it's important to keep in mind that every organization is like a snowflake: unique but with similar patterns. We see many different team configurations based on staff availability, skill sets, staff seniority, and organizational structure. You may find that one person can cover multiple areas of expertise, or that you need to extend your team. While there's no definitive answer to which roles and how many roles should be included, you can use this as a guide to help you determine the right people in your organization for the project.



KEY ROLES



Application Admin

The Application Admin handles the day-to-day administration. They have a deep 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.

Deployment responsibilities

- Verifies functionality and performance during testing to ensure that the Data Center application is operating properly
- When upgrading, determines which apps will be included, and decides whether or not to keep apps that are not Data Center certified
- Ensures that the users and permissions are maintained, or changed, correctly throughout the transition to Data Center



Sysadmin

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

Deployment responsibilities

- Ensures the Data Center application can connect to the database through multiple nodes, instead of a single server
- Gathers the hardware (physical or virtual) needed
- Performs the actual installation of the Data Center application
- Configures the load balancer to work with the Data Center cluster
- Creates, configures, and mounts the file system
- Pipes logs from on-disk to log aggregator



Project Lead

The Project Lead has a deep relationship with the business and knows how and why the application is used to meet company objectives. They also know how to make the right tradeoffs to maintain governance across applications.

Deployment 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 Data Center



Executive Sponsor

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

Deployment responsibilities

- Ultimate decision maker in the event of roadblocks
- Facilitates budget approvals for the proof of concept and rollout on production

AREAS OF EXPERTISE

In addition to these roles, there are also a number of areas of expertise you'll want to make sure each member of your team covers. Depending on what product you are migrating to Data Center, you may need some (or all) of the following:



Networking Engineering

Networking provides the connections between the different pieces of the Data Center architecture. Optimizing these connections will ensure you have a well-tuned deployment.



Database Management

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




Site Reliability

Site Reliability experts ensure that each application is running at peak performance. They help ensure uptime and bring systems back online in the event of an outage.



Security

It's advantageous, and often required, to test and audit your system from a security perspective. Those who have expertise in security practices will make sure that the system is safe and up to regulation.



Tip: Form a steering committee. Steering committees are made of key power-users and end-user stakeholders who can provide feedback on how the product is used at the individual and team level. They can give the thumbs-up on upgrades, sign off on governance changes, or help prioritize requests for new customizations, apps, and more.

Align your team to goals

Once your project team is assembled, it's important to have all members aligned on the business, functionality, and performance goals for the migration. Aligning stakeholders upfront ensures a smoother installation, testing, and release process so that your team can confidently deploy Data Center to production.

Don't forget to agree upon a target go-live date. You may have a business initiative that requires you to be live by a certain date, or a time sensitive budget request that requires you to make your purchase in a specific quarter. As you keep this in mind, communicate and build your timeline based on the agreed-upon target date.

Step 2: Review your existing Server environment

Before you move to Data Center, it's important to take a look at your existing Server environment. Not only should you understand existing conditions, but it's also important to look for ways to optimize the existing Server environment so those same improvements can be translated to your new Data Center environment.

BENCHMARK YOUR SERVER APPLICATION

Take a baseline measurement of your system's existing functionality and performance. That way, throughout testing, you can measure the improvements you should expect from migrating to Data Center in comparison to your original Server instance.

FINE-TUNE YOUR SERVER APPLICATION

Before your migrate, you should fine tune your Server application. Data Center does not inherently improve the performance of the individual nodes themselves. Any performance issues not related to concurrency, arising from suboptimal configurations or usage in the Server application, will likely persist or even worsen in Data Center. Spend the time identifying and finding ways to peel back usage issues so these issues don't come back to bite you. Although it may take 1-2 weeks to identify and make these performance optimizations, we've found that customers who don't do this early on tend to have a longer and more challenging deployment.

ASSESS AND UPDATE GOVERNANCE

How users interact with the application also affects application performance. Before deploying Data Center, asses these usage characteristics and determine whether you need to establish any restrictions on things like scripts that make REST calls or other integrations to protect performance.

Review installed marketplace apps

Using a large number of apps may degrade application performance. It's critical that you audit and remove apps that aren't crucial to system functionality to increase overall system performance. You'll also want to make sure all of your apps are compatible with Data Center and find alternatives when necessary.

Step 3: Document current processes

After application tuning, it's time to document Server environment. This documentation can help guide configuration decisions in your Data Center deployment, influence process modifications, and determine whether issues found after deployment are new or already existing.

Step 4: Evaluate technology decisions

Getting ahead of the technology decisions will speed up the design of a production-ready cluster environment tailored to your organization's needs. You'll need to start assessing the components we discussed earlier (a load balancer, database, a file system and your application nodes) and determine what you have available, what you may have to purchase, or any other key decisions around the supporting infrastructure, such as whether you are deploying on AWS, Azure, or on your own hardware.

Step 5: Implement and test your deployment process

We're done with the planning. Now it's time for action. Before you deploy, make sure you've evaluated and determined the team, technology and other key factors we've already mentioned.

BUILD YOUR CLUSTER

It's time to build your Data Center cluster. In addition to setting up each of various components in your cluster (application nodes, load balancer, database, file system), you also need to size the application nodes in your cluster based on your performance requirements. The number and size of the nodes in your cluster depends on your needs and how you configure your application. We have built documentation to help guide you on how to approach node sizing. Find out more about node sizing [here](#).

We've also put together some sample configurations you can reference. Atlassian does not endorse, approve, or recommend any specific vendors or configurations. These are provided as a reference. For more guidance on configuring your optimal environment, contact a [Technical Account Manager](#) or have [Premier Support](#) schedule a health check.

Jira Data Center on Amazon Web Services

Approximate User Count: 3,200 | [VISIT LINK](#)

Dimension	Production (AWS)
Application	Jira Data Center
Load Balancer	AWS ELB + nginx on each node
Node Size + Operating System	2x EC2 m4.large Nodes 4x CPU, 16GB RAM, 50GB Storage AWS EC2 Linux
Database	PostgreSQL 9.3 (AWS RDS)
Shared Storage	AWS EFS

Bitbucket Data Center

Approximate User Count: 4,700 | [VISIT LINK](#)

Dimension	Production (AWS)
Application	Bitbucket Data Center
Load Balancer	HAProxy
Node Size + Operating System	5x Application Nodes (VM) 8 CPU, 16GB RAM, 50GB RAM Ubuntu 14.04
Database	PostgreSQL 9.3
Shared Storage	NFS 500GB

Confluence Data Center

Approximate User Count: 4,200 | [VISIT LINK](#)

Dimension	Production (AWS)
Application	Confluence Data Center
Load Balancer	A10 Networks
Node Size + Operating System	3x Application Nodes (VM) 4 CPU, 32 GB RAM, 100 GB Storage CentOS 6.6
Database	PostgreSQL 9.3
Shared Storage	NFS 1TB

INSTALL THE PRODUCT

Are the hosts for your cluster all set up? It sounds like you're ready to install. Depending on what product you plan to install, we suggest you read one of the following guides:

[Bitbucket](#) [Confluence](#) [Crowd](#) [Jira Service Desk](#) [Jira Software](#)

MONITOR AND SCALE YOUR DATA CENTER ENVIRONMENT

Finally, in order to confidently deploy Data Center to production, you should run through an iterative set of functional tests, integration tests, and performance tests to vet your Data Center installation. Be prepared for the future. Also, make sure you are taking the right steps now to maintain and scale your environment down the line.

Testing

[Testing your application](#) helps determine whether your infrastructure is capable of supporting adequate performance levels under typical loads. You'll need to predict realistic usage patterns and load to conduct useful tests. We've created several performance testing frameworks to not only help you conduct your own tests, but also provide you insights on how we test our own Data Center applications.

- [Atlassian Performance Testing Framework](#)
 - [Performance Testing Framework for Bitbucket](#)
 - [Performance Testing Framework for Confluence](#)
- [Atlassian performance testing examples](#)
 - [Bitbucket Data Center sample performance test](#)
 - [Confluence Data Center sample performance test](#)
 - [Crowd Data Center sample performance test](#)
 - [Jira Data Center sample performance test](#)

Monitoring

To keep your instance healthy, you need to monitor it regularly. To do that, you need to know which metrics to track, create alarms for critical thresholds, and make plans for when you hit them. Do it right, and your Data Center application will continue performing at an optimal level.

Every deployment is different, so we've provided you with guidelines to help you develop a performance monitoring plan that best suits your environment.

- [Getting started with Data Center monitoring](#)
- [Tools for monitoring your Data Center application](#)
- [How Atlassians monitor their enterprise deployments](#)

Sizing

After your initial deployment, it's important to make sure your environment is set up correctly. Continue to assess your needs and configuration to determine if you need to adjust node sizes.

Need help? We've got you covered!

Think you need additional help preparing for or during your migration? Don't worry, we have a number of resources available to help you out!

TECHNICAL ACCOUNT MANAGERS

Think of a [Technical Account Manager](#) as your strategic partner for all things Atlassian. They'll help guide your journey by providing expertise and asking the questions that you wouldn't think to ask yourself.

PREMIER AND PRIORITY SUPPORT

[Atlassian Premier Support](#) provides fast 24/7 support for any issues the team may experience during the deployment. Premier Support can also help review the current Server installation and validate whether it's ready for a move to Data Center.

Atlassian Priority Support provides mission-critical coverage for support tickets, routed directly to our most senior Support Engineers.

SOLUTION PARTNERS

[Solution Partners](#) conduct hands-on system integrations, deployments, and upgrades. They'll also work with your team to customize Atlassian software to your specific needs. Solution Partners are a great option for organizations with complex requirements or are looking for onsite help. Visit our [Partner Directory](#) to find a Partner that's right for you.






ATLASSIAN COMMUNITY & USER GROUPS

The [Atlassian Community](#) has got your back! Stay connected by joining the enterprise community group, a place to share stories, tips, and best practices for using Atlassian products at scale. If you'd like to swap stories with other Atlassian customers in your area, the Atlassian User Group community has over 40,000 people in more than 30 countries! [Join a local user group today!](#)

Try it free!

Evaluate any of our
Data Center products.

Click the link to download.

-  [Crowd Data Center](#)
-  [Bitbucket Data Center](#)
-  [Jira Service Desk Data Center](#)
-  [Confluence Data Center](#)
-  [Jira Software Data Center](#)

