

Tableau Server 10.0

High Availability: Delivering mission-critical analytics at scale

Kitty Chou, Product Manager
Mike Klaczynski, Product Marketing

Table of Contents

- Self-Service Analytics is Mission-Critical 3
- Understanding High Availability..... 3
- Tableau Server Scalability 4
- Out-of-the-Box High Availability 4
- Understanding Tableau Server High Availability.....5
 - General 5
 - Gateway..... 5
 - Application Server 6
 - Coordination Service..... 6
 - Cluster Controller7
 - Repository7
 - Backgrounder 9
 - Data Server 9
 - Cache Server 9
 - Data Engine 9
 - File Store10
 - Search & Browse 11
 - VizQL Server 11
- Integrating with Third-party Monitoring Tools.....11
- Failover of Primary Server Node.....12
- Configure a Backup Primary13
- Monitoring Cluster State 14
- Architectural Considerations15
- Select the Best Configuration.....15
 - Minimal Three-Node HA Deployment16
 - 3+ Node Deployments 17
- Disaster Recovery with Tableau Server 18
- The Simplest DR Strategy18
- Third-party DR Solutions18
- Tableau DR Capabilities.....18
- Beyond High Availability19

Self-Service Analytics is Mission-Critical

Today more than ever, self-service analytics and data-driven decision-making are becoming the norm in organizations worldwide. Users and decision makers have come to depend on immediate access to data and self-service tools to answer their questions in real time. Executives understand the importance of data-driven decisions at their companies, and rely on these systems daily. This reliance on data requires a high degree of availability to the underlying systems. A platform's capabilities need to be more accessible and easily configurable by existing teams and enterprise tools.

Tableau Server 10.0 delivers the future of mission-critical self-service analytics. It enables rapid self-service data exploration, promotes trust in content and data through robust governance, and is easy to deploy, manage, and scale across any enterprise. In this paper we will explore how Tableau Server 10.0 delivers self-service analytics at scale with high availability.

Understanding High Availability

The goal of highly available systems is to minimize downtime of the system. Availability is commonly expressed as the “number of nines” and effectively measured as the percentage of actual uptime versus expected uptime. The table below shows how the number of nines corresponds to annual downtime.

Number of 9's	Availability Percentage	Total Annual Downtime
1	90%	36.5 days
2	99%	3 days, 15 hours
3	99.9%	8 hours, 45 minutes
4	99.99%	52 minutes, 34 seconds
5	99.999%	5 minutes, 15 seconds

Figure 1: Typical availability metrics and corresponding annual downtimes.

System administrators often have service level agreements (SLA) with their business users to define an acceptable threshold for downtime. Based on that SLA, they will choose deployment architectures to meet those goals. Most system administrators plan downtime for maintenance, upgrades, and patching. In addition, there is some likelihood of unexpected failures, which is referred to as unplanned downtime. Of course, administrators need to conduct planned maintenance for hardware or software updates; the goal is to minimize unplanned downtime.

We understand how important it is for users to readily see and understand their data. We also realize there will always be events that threaten the availability of business intelligence systems, whether

related to hardware, software, networks, or even human error. At a minimum, Tableau Server 10.0 processes will automatically restart to keep your system running in the event of component failure. A properly configured multi-node deployment uses redundant processes to achieve server high availability (HA). However, unlike most systems, Tableau makes it easy to set up and configure your analytics environment for HA.

Tableau Server Scalability

Tableau Server is architected to scale up and scale out. It provides large organizations with enterprise-grade deployment stability while still retaining the simple and easy-to-use qualities that make it appealing for smaller teams. Depending on your environment, Tableau Server can run on one or more computers — and run one or more component processes on the same node—in order to best serve both your user demands and your HA requirements.

Internally, Tableau manages several cloud-scale Tableau Server deployments, these include Tableau Public and Tableau Online. Tableau Public is a customized Tableau Server deployment that supports millions of public views each week across the world. As part of our engineering and release process, we deploy beta versions of Tableau Server software to Tableau Public to fine-tune stability and quality before we release to our corporate customers.

Out-of-the-Box High Availability

Installing and configuring high availability for Tableau Server is easy. A default installation takes just minutes. You can install Tableau Server on the primary computer as well as all additional worker nodes in the cluster, then configure and form a highly available cluster using the configuration utility.

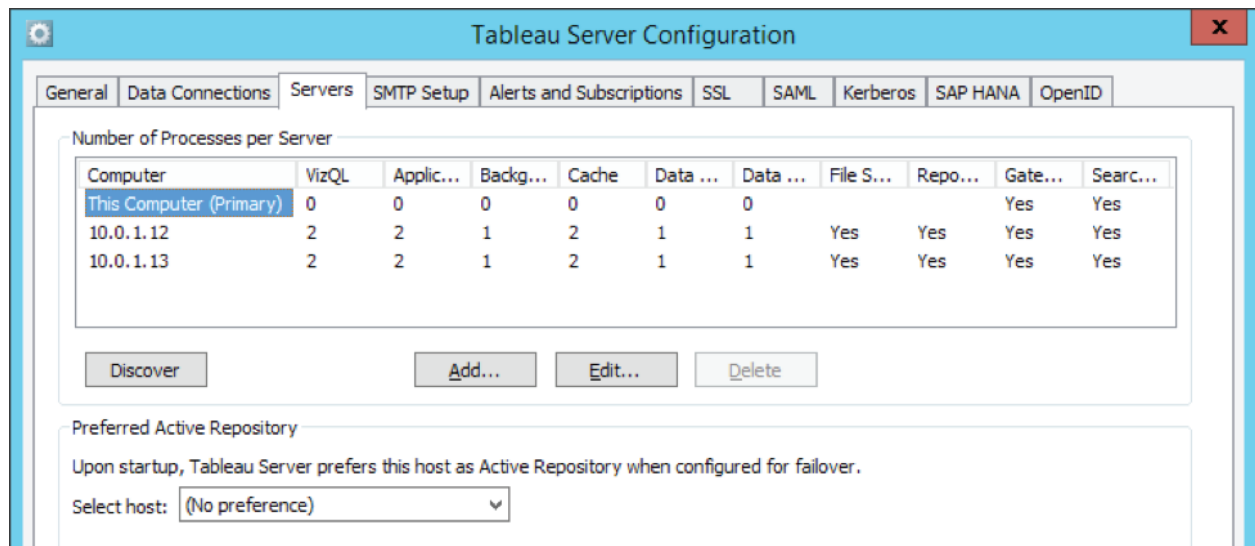


Figure 2: The Tableau Server configuration utility provides flexibility in the count and configuration of nodes, and processes per node in a cluster.

Understanding Tableau Server High Availability

Tableau Server has several dedicated processes that ensure the entire system is accessible by end users. This section assumes familiarity with Tableau Server components and what they do. If you are not familiar with them, we suggest first reading the [Tableau Server Administration Guide](#).

Understanding how to make Tableau Server highly available is really the same as understanding how to make each of its components highly available. In other words, high availability for every component must be ensured in order for the entire Tableau Server cluster to be highly available, providing redundancy for single points of failure. Let us consider each component in turn.

General

Tableau Server has built-in automation for restarting all its server processes. This automation ensures high availability by automatically restarting server processes that fail and notifying administrators. The hardware or virtual machine that houses Tableau Server must be healthy for this automation to occur.

To protect against failure of an entire node, it is important to configure the unique server processes so that they are redundant across different nodes in the cluster. This redundancy is possible for all processes except the Licensing service, which can only run on the primary node. We discuss how to address this process limitation in more detail within the *Failover of Primary Server Node* section.

Gateway

Starting with Tableau Server 8.1, the Gateway process can run on any and all nodes of the Tableau Server cluster, not just the primary Tableau Server node. This has important implications for HA.

Prior to Tableau Server 8.1, only one node could be configured to run a Gateway process. And when that one process failed, Tableau Server couldn't communicate with anyone outside the cluster, and also caused disruption of many internal communications. In Tableau Server 10.0, when multiple nodes are configured to run a Gateway process, communication with the server remains intact as long as at least one Gateway process is running.

Mitigating Risk of Gateway Failures

The key to Gateway high availability is to have more than one node in a Tableau Server cluster and to configure more than just one to run the Gateway process. In fact, we recommend that you configure a

Gateway process on each node. This mitigates the risk of the Gateway process being a single point of failure and leading to service unavailability.

So what happens when a Gateway process fails? As mentioned previously, if no Gateway processes are running, the entire Tableau Server cluster will be unavailable. If other Gateway processes remain running, requests made to those working Gateways will be processed normally. However, any requests received by the failed Gateway will not be redirected and will continue to fail, despite the presence of other functioning Gateways. Failed Gateway processes automatically restart; so as long as the computer itself is functioning, the failed Gateway process will relaunch and resume serving requests.

To make your installation even more robust to Gateway failures, we recommend placing your Tableau Server cluster behind an external load balancer (ELB) to ensure that requests only get routed to functioning Gateway processes. The [Tableau Server Administrator Guide](#) provides some additional guidance on setting this up.

Application Server

Achieving high availability with Application Server is easy. You just need to configure Application Server instances on each node in the Tableau Server cluster.

What happens when an Application Server process fails? Requests being handled by that instance will fail, but subsequent requests will be routed to other running Application Server processes. Assuming the node containing the failed Application Server is still running, the failed process should automatically restart within seconds.

Coordination Service

The Coordination Service is part of the base installation of Tableau Server. The number of nodes installed with the Coordination Service is dependent on the number of nodes in the cluster as outlined in the following table:

# of Nodes in the Cluster	# of Nodes with Coordination Services
1 - 2	1
3 - 4	3
5+	5

The Coordination Service process will be installed sequentially on the first n nodes of the cluster (including the primary node) where n is the number of nodes with Coordination Service – based on the table above.

Understanding Quorum

A complete Tableau Server outage will occur if the number of running Coordination Service processes does not constitute a quorum, which is based on the total number of configured Coordination Services processes. A quorum is just another way of saying an absolute majority. So a cluster with either three or four computers is able to tolerate the loss of, at most, one node (one instance of the Coordination Service). A cluster with either five or six computers can tolerate losing up to two Coordination Service processes.

Notably, a cluster with only two nodes cannot tolerate the loss of even a single Coordination Service process, and this is the reason why full HA, including automated failover, requires a minimum of three nodes.

What happens when a Coordination Service process fails? Nothing, as long as the number of remaining Coordination Service processes still constitutes a quorum. If the number of still functioning Coordination Service processes number less than a quorum, the entire Tableau Server cluster becomes unavailable in order to protect the referential integrity of the underlying Postgres database.

Failed Coordination Service processes automatically restart as long as the computer itself is okay.

Cluster Controller

Cluster Controller is also part of the Tableau Server 10.0 base installation. One Cluster Controller process is configured on each node of the cluster during installation, no explicit configuration is required.

What happens when a Cluster Controller process fails? All other Tableau Server components on the same node will become unavailable and display as “unavailable” on the Tableau Server status page. Any Repository process running on that node will also be unavailable. (To fully understand this effect, read the *Repository* section below.) In order to protect against a Cluster Controller process failure, ensure that each unique server component has redundancy and is running on at least two different nodes in the cluster.

Failed Cluster Controller processes automatically restart as long as the computer itself is okay.

When the Cluster Controller restarts, it also restarts any Repository process configured on that node.

Repository

The Repository is a necessary database for Tableau Server to function. If there is no fully-functioning Repository, the entire Tableau Server cluster is unavailable. At any point in time, there can be only one fully-functioning Repository in use, and this Repository is referred to as the “active” Repository. All operations involving the Repository run on this active Repository. All operations involving the Repository run on this active Repository.

To improve availability, you can configure Tableau Server with an additional, “passive” Repository on a different node of the cluster. The contents of the active Repository constantly stream to the passive Repository. In the event of a failure of the active Repository, a cluster configured for high availability will automatically promote the passive Repository to active status for continued availability of the server. Customers who want high availability will certainly want to configure a passive Repository. In an entire cluster, there can be a maximum of only two Repositories—with one active and one passive—they cannot be located on the same node.

The Cluster Controller manages Repository startup, shutdown, and any failover from active to passive. This is why a failure in the Cluster Controller process can cause a problem with the Repository. If the Cluster Controller process that started a Repository fails, then the Repository will also fail.

What happens when a Repository process fails? That depends. The following list summarizes the various cases:

- If the passive Repository fails, then users should experience no impact. Everything will continue to work since the active Repository is still functioning. In the background, the passive Repository will be restarted, and data replication will resume, though there may be some delay before the passive Repository is again fully synchronized with the active Repository.
- If the active Repository fails, and if there is no fully-synchronized passive Repository, then Tableau Server is unavailable until the active Repository can be restarted. The system will attempt to do this automatically, but, depending on the reason for the failure, this may not be possible.
Note: This means – if only the active Repository is currently usable, then Tableau Server cannot be considered highly available. The active Repository is a single point of failure for the entire system if there is no synchronized passive repository
- If the active Repository fails, and there is a fully-synchronized passive Repository available, and the cluster is configured for HA, then a failover to the passive Repository will be automatically triggered. After the failover, the previously passive Repository will be the new active Repository. The system will restart the failed previously active Repository as the new passive Repository and begin synchronization. It will also restart other relevant processes automatically so they are aware of the newly-promoted active Repository and can reconnect. During this short window of restarts, users will experience a service interruption. However, administrators do not have to intervene to ensure continued availability as this is an automated sequence. Should you wish to manually promote the passive Repository to become the active Repository, you can do so using the `tabadmin failoverrepository` command.

Backgrounder

To make the Backgrounder service highly available, you should configure more than one Backgrounder processes to run on multiple nodes in the cluster. When deciding where and how many Backgrounders to run, consider each machine's available capacity based on other server processes.

What happens if a Backgrounder process goes down? The jobs that the Backgrounder is working on fail and are not retried. However, most background jobs are scheduled to run periodically, and the same background task will be picked up and performed normally at the next scheduled time by a functioning Backgrounder process.

Failed Backgrounder processes automatically restart as long as the computer itself is okay, but they will not retry the jobs that failed.

Data Server

To make Data Server highly available, configure one or more Data Server processes to run on multiple nodes of the cluster.

What happens if a Data Server process fails? Queries running via proxy through the Data Server process will fail, resulting in a failed view rendering. Subsequent requests, including a retry of the failed operation, should succeed as long as a working Data Servers exist that can accept rerouted requests.

Tableau Server is not dependent on Data Server to function; however, without a running Data Server, the cluster loses its ability for workbooks to proxy through to external data sources. Any view that does not use Data Server for one of its data sources should still function correctly.

Cache Server

The Cache Server provides a shared external query cache. It's a cache of key/value pairs that hold information from previous queries and speed up future requests. This means that if a Cache Server process becomes unavailable—or even if all of them do—the consequences are relatively mild. Tableau Server will still work but actions may take longer as they do not have pre-cached results available. As the queries re-run, the restarted Cache Server gets re-populated, eventually speeding things up for end users. In effect, Cache Server does not have an availability impact; however, it does have an impact on various end user performance scenarios. Just like all the other processes, the Cache Server will also restart itself as long as the machine itself is healthy.

Data Engine

The Data Engine component loads and queries data extracts when using in-memory analytics. To make the Data Engine highly available, simply configure one or more Data Engine processes to run

on multiple nodes of the cluster. All Data Engines run in active/active mode and perform exactly the same functions. It is important to note that any node configured to run a Data Engine process will also be configured to run the File Store process. The File Store process manages storage and replication of extract files and is described next.

What happens when a Data Engine process goes down? Queries currently running on that Data Engine process will fail, resulting in a failed view rendering or failed extract refresh. Running the same operation again will automatically be reassigned to a different Data Engine that is functioning.

Failed Data Engine processes are automatically restarted; as long as the computer itself is okay, the Data Engine process will relaunch.

File Store

As mentioned previously, a File Store process will automatically instantiate on any node that runs one or more Data Engine processes. The File Store process manages storage and replication of extract files between nodes.

How File Store Works

An extract file is created in the system when a user first publishes it to Tableau Server or an extract refresh of it occurs. Immediately after one of these events, the extract exists on a single File Store on a single node. The specific extract file cannot be said to be highly available yet as it does not have redundancy and is therefore a single point of failure. The File Store processes communicate with each other to quickly replicate local extracts to all other File Store nodes in the cluster. The File Store process is designed to copy the files as quickly as the cluster network resources allow for, but can take a variable amount of time depending on the size of the extract. Once a copy is available on multiple nodes within the clusters, the extract file is considered highly available.

What happens if a File Store process fails? There are two consequences:

- Copying of extract files to and from the affected node stops.
- Removal of no-longer-needed extract files on the affected node is suspended. This removal process is usually referred to as “extract reaping.”

The suspension of extract reaping has no immediate effect. It simply means that a buildup of unwanted extract files will begin to consume disk space on that node. Eventually this will cause a problem, but correctly sized nodes should have a comfortable buffer of disk space.

The lack of file replication means that new extract files added to the failed File Store node will not be available on any other File Store nodes in the cluster, and vice-versa. After the File Store process restarts, the system corrects itself by ensuring synchronization between all File Stores on all nodes.

Failed File Store processes automatically restart as long as the computer itself is okay. The File Store process should quickly resume its function, including synchronization of all files, regardless if they were added during the period of failure or after.

Search & Browse

Achieving high availability for the Search & Browse process is easy. The system should be configured so that the Search & Browse process is running on multiple computers.

What happens if a Search & Browse process fails? Tableau Server is largely unusable, and though users can still log into the system, workbook content will appear to be missing. The content is not actually missing, it is simply not being returned in the search results, it will be redisplayed after the Search & Browse process restarts. If more than one Search & Browse process is configured and running on multiple nodes when the failure occurs, requests made to a failed Search & Browse process will also fail, but subsequent requests will be routed to working Search & Browse processes. Each Search & Browse process indexes across all nodes in the cluster, therefore if all but one Search & Browse process fails, results will still be returned across all nodes.

VizQL Server

To achieve high availability for the VizQL Server process, simply configure one or more instances to run on multiple computers.

What happens if a VizQL Server process fails? If there is only one VizQL Server process and it fails then Tableau Server will no longer be able to render any views. High availability requires configuring redundant VizQL processes. A fairly typical configuration consists of two to four VizQL Server processes on each node. This simultaneously serves the need for high availability and scalability. If multiple VizQL Server processes are running, then the failure of a single process will result in the failure of requests and loss of session data at the time of that processes failure. Any future requests will be routed to the other working VizQL Server processes across the Tableau Server cluster.

So far, we have seen how each server process behaves on failure and how you can mitigate risks of failure to ensure that the Tableau Server cluster as a whole is highly available. In addition to planning for each failure scenario, you should also proactively monitor the cluster for past process failures.

Integrating with Third-party Monitoring Tools

In addition to monitoring system health using Tableau Server's built-in mechanisms, you can also remotely receive a machine-readable XML version of each processes status. You will need to enable remote access to the server and use the appropriate URL for your instance: `http://<my_tableau_server>/admin/systeminfo.xml`

This will return a `status.xml` file, which can be parsed to identify the server's status and integrated into other systems or monitoring tools:

Sample machine-readable XML from the `status.xml` file provided by remote access to the server status URL

```
<systeminfo xmlns:xsi="http://www.w3.org/2001/XMLSchema-instance">
  <machines>
    <machine name="my_tableau_server">
      <repository worker="my_tableau_server:8060" status="Active"
preferred="false"/>
      <dataengine worker="my_tableau_server:27042" status="Active"/>
      <applicationserver worker="my_tableau_server:8600" status="Active"/>
      <apiserver worker="my_tableau_server:8000" status="Active"/>
      <vizqlserver worker="my_tableau_server:9100" status="Active"/>
      <dataserver worker="my_tableau_server:9700" status="Active"/>
      <backgrounder worker="my_tableau_server:8250" status="Active"/>
      <gateway worker="my_tableau_server:80" status="Active"/>
      <searchandbrowse worker="my_tableau_server:11000" status="Active"/>
      <cacheserver worker="my_tableau_server:6379" status="Active"/>
      <filestore worker="my_tableau_server:9345" status="Active"
pendingTransfers="0" failedTransfers="0" syncTimestamp="2015-02-
27T20:30:48.564Z"/>
      <clustercontroller worker="my_tableau_server:12012"
status="Active"/>
      <coordination worker="my_tableau_server:12000" status="Active"/>
    </machine>
  </machines>
  <service status="Active"/>
</systeminfo>
```

Failover of Primary Server Node

The primary Tableau Server node is identified as the server where the first Tableau Server installation occurred. The primary server is unique in that it includes a unique license management process and other administrative functions, in addition to potentially being a fully-fledged server installation. In larger deployments, you can deploy the primary server with only the unique barebones administrative functions, omitting processes that are redundant on worker nodes. While Tableau Server gives you a lot of deployment flexibility, it is important to remember that its processes poll the primary server every seventy-two hours to perform licensing checks. If the primary server is not reachable during a license check, it will fail and your Tableau Server deployment will become “unlicensed” and be disabled.

However, in the real world, your primary server could fail for any number of reasons outside the realm of the Tableau Server software—for example a hardware, OS, or network failure. In such cases, if you have deployed the cluster behind an external load balancer and have Gateways on all nodes in the

cluster, the rest of the cluster will remain available to service requests. If however, the primary node is still disabled when the licensing check window expires, the entire cluster will become unlicensed. To prepare for such failures, be sure to have a warm backup of the primary node on standby.

Configure a Backup Primary

If the primary server fails, you will experience downtime of administrative tasks until the issue causing the server malfunction is resolved. For this reason, you should allocate a dedicated backup machine as the “backup” primary in the event the primary server node goes down. Once configured and ready, the backup primary should not be turned on or connected to the cluster. This ensures that licensing and administrative functions continue to work against the primary server node. Learn how to setup and stage a backup primary server by following directions in our [Tableau Server Admin Guide](#).

If the primary server fails, we have also documented the steps for switching to the backup primary server in our [Tableau Server Admin Guide](#). These steps alert the rest of the cluster about the change and automatically reroutes their administrative and licensing requests to the now new primary server, these steps can also be automated.

Monitoring Cluster State

System administrators can monitor the state of Tableau Server 10.0 processes on the Tableau Server status page. This provides administrators with an understanding of the health of the overall cluster by displaying the server processes across all nodes as well as the status of the active and passive Repositories.

Server Status

Process Status
The real-time status of processes running in Tableau Server.

Process	Primary 10.0.1.11	Worker 1 10.0.1.12	Worker 2 10.0.1.13
Cluster Controller	✓	✓	✓
Gateway	✓	✓	✓
Application Server		✓ ✓	✓ ✓
VizQL Server		✓ ✓	✓ ✓
Cache Server		✓ ✓	✓ ✓
Search & Browse	✓	✓	✓
Backgrounder		✓	✓
Data Server		✓	✓
Data Engine		✓	✓
File Store		🔄 Synchronizing	🔄 Synchronizing
Repository		✓	✓

 Active Busy Passive Unlicensed Down Status unavailable

Figure 3. The Server Status admin page shows the status of each process on each node.

Tableau Server can also be configured to send email alerts to server administrators when there is a system failure via the Tableau Server Configuration utility. This utility can also be used to provide early warning of disk space capacity issues.

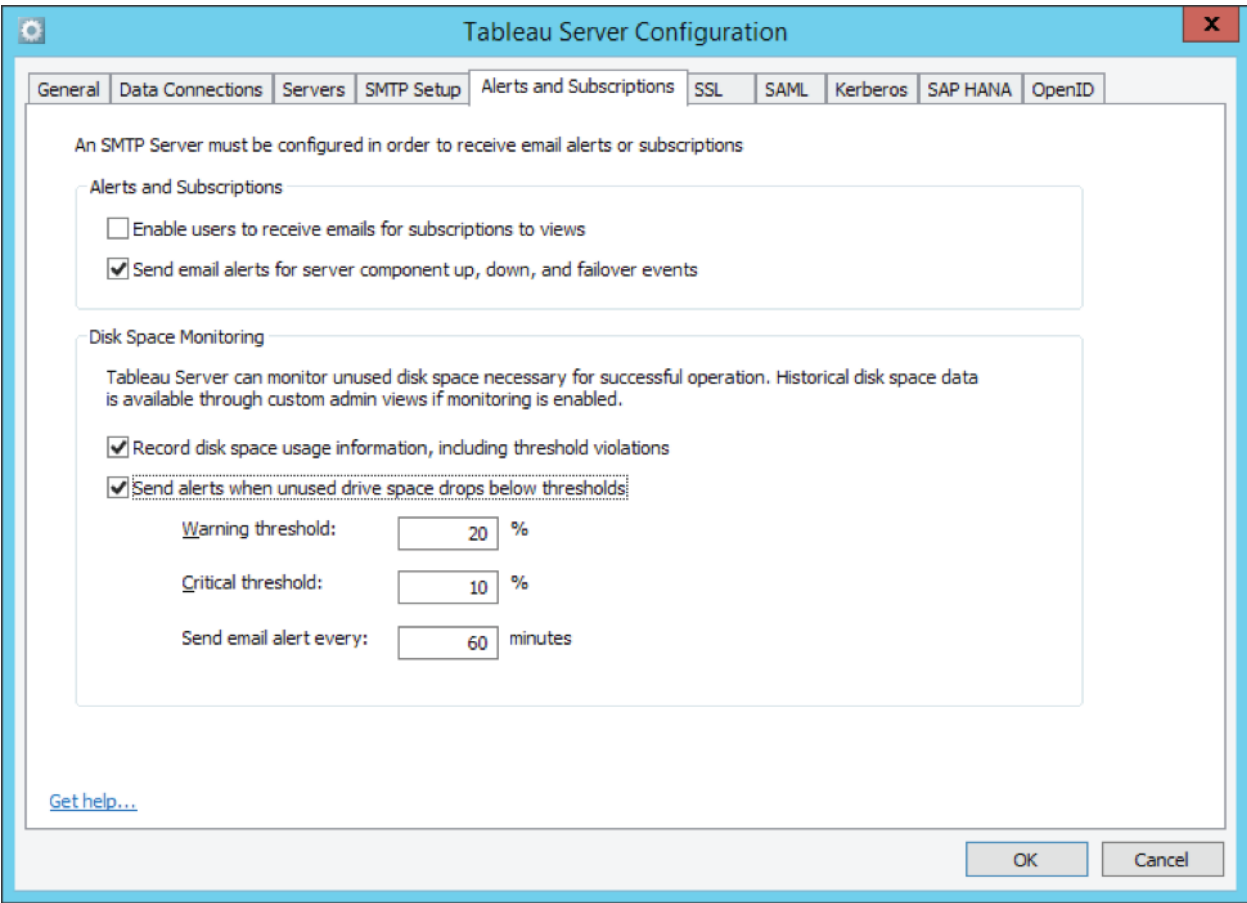


Figure 4. The Tableau Server Configuration utility provides the ability to set alerts on system failures.

Architectural Considerations

For a successful implementation of high availability, you should understand the uptime objectives and expected service levels of your organization, then plan for redundancy accordingly.

Select the Best Configuration

Determining your optimal cluster configuration will depend on both the needs of your business and the resources available in your organization. Although there are several possible successful configurations for high availability, it's important to understand the tradeoffs between clusters of various sizes so you can make the best choice for your environment. Additionally, you may choose between symmetrical (each worker node has all of the same components) or asymmetrical computer configurations (each node has a varying number and type of components). Often times, having a symmetrical configuration for worker nodes simplifies cloning and appending additional nodes to the cluster. However, keep in mind the constraints of a single active and single passive Repository as you plan your configuration.

Minimal Three-Node HA Deployment

As we discussed earlier, you must have at least three nodes in the cluster to be running fully in HA mode. A three-node cluster is also a great starting point for larger deployments. Should one node fail, you will still have a quorum with the remaining active two nodes. Figure 5 below, shows the responsibilities of the primary node for administrative and licensing functions. It also illustrates the two additional worker nodes that power the data and visualization functionality of Tableau Server. The exact mix of these components may vary depending on your scalability needs in addition to your HA needs. For example, while the graphic shows only one VizQL process per worker node, in reality you could have two or more VizQL Server processes running on a worker to support HA and your scalability needs. In addition, if you have more core capacity on your primary node, you may choose to include additional server processes such as VizQL Server or the Data Engine/File Store combination for added scalability.

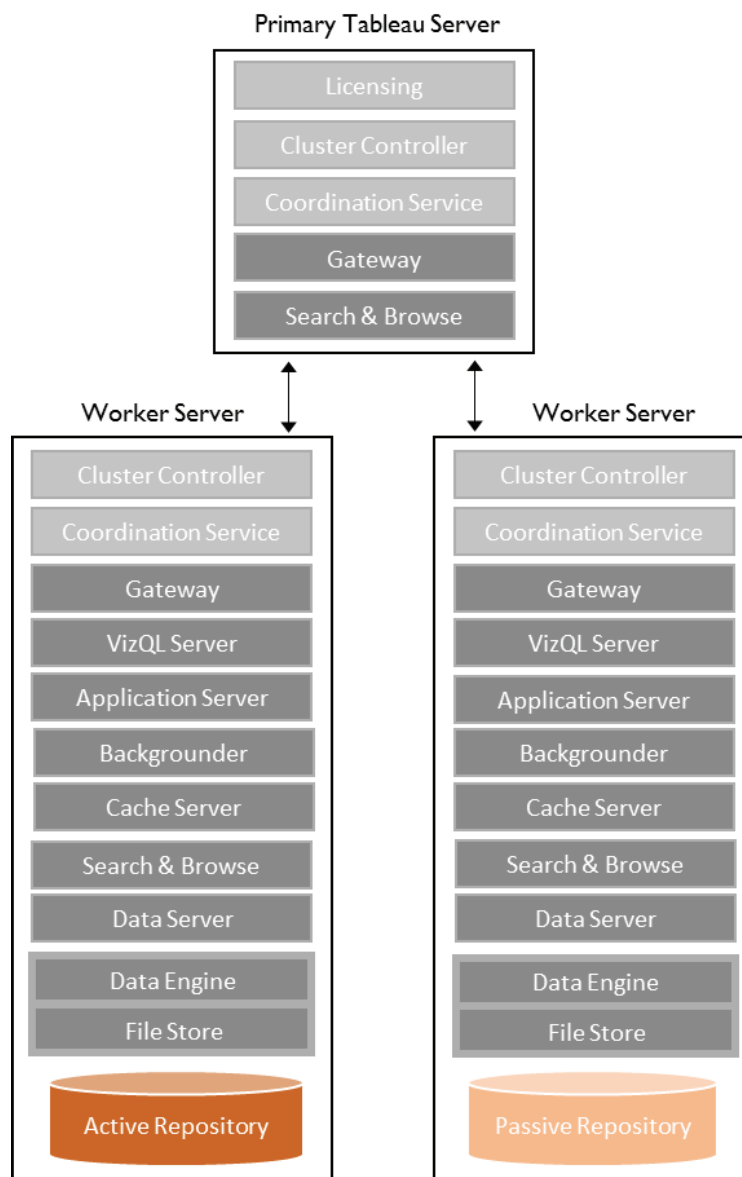


Figure 5. Example of a minimal 3-node HA Deployment

Having Data Engines installed on different computers gives you added redundancy in case of hardware failure and rapid backup recovery. If you still need additional redundancy, you can deploy an additional Data Engine on each worker node.

The processes on your primary server will not be affected if either worker goes down. However, the workers are dependent on the primary. In this configuration, if a worker with the active repository fails, the worker with the passive repository automatically becomes active. However, as discussed above, it is best to have a backup primary configured and staged to ensure quick recovery from a primary server failure. In addition, if your organization has access to an external load balancer, you should deploy the cluster behind that ELB for improved end-user availability. If you are just getting started with high availability, this is a good configuration to start with.

3+ Node Deployments

When using three or more nodes, deploying an odd number of nodes in a cluster is a preferred topology for improved high availability. Clusters with an even node count have the same quorum ability as clusters with one fewer node. For example, should a two-node outage occur, both clusters with three and four nodes will be equally affected, losing quorum and resulting in a shutdown.

Adding a 4th, even node to a three-node cluster will spread the workload across more nodes helping with scalability, but also reduce the risk of complete cluster failure by providing additional redundant processes.

Setting aside other considerations and just focusing on the assumption that you want the least downtime, you should consider a deployment architecture that is five or more nodes. While we recommend the three-node architecture in this paper as a great starting point, large enterprise-wide deployments with mission-critical high availability requirements should consider using five or more nodes as a baseline.

For organizations that rely heavily on extracts and extract refreshes, a commonly used deployment architecture includes deploying Backgrounder processes onto dedicated worker nodes. Extract refresh workloads can sometimes interfere with the visualization workloads supported by VizQL Server processes, so by partitioning Backgrounder processes to dedicated worker machines, you can avoid resource contention between these two different workloads. To ensure that you have process redundancy across multiple nodes in the cluster, Backgrounder-only workers are often deployed in pairs.

As we mentioned, high availability reduces downtime due to critical failures. In addition to fulfilling service-level requirements and deploying a topology that supports high availability, you may still encounter failure situations that demand rapid recovery. In case of disaster and/or hardware failures, plan for system recovery based on your disaster recovery goals and objectives.

Disaster Recovery with Tableau Server

When planning for disaster recovery (DR) in your Tableau environment, there are two main factors to consider: (1) your business tolerance for downtime as part of the recovery process, and (2) your business tolerance for lost data. This is defined in terms of Recovery Time Objective (RTO), a measure of how much time can elapse before a full recovery, and Recovery Point Objective (RPO), a measure of tolerance to data loss.

When determining your RTO and RPO thresholds it is important to understand the tradeoffs, which are mostly financial. Regardless of the specifics, it is costly to support a very low recovery time objective when you take into account the infrastructure and staff necessary to bring the systems back online rapidly. If you choose to spend less on the infrastructure and responsiveness of staff, you should be able to tolerate much longer RTO (longer downtimes). This tradeoff is not specific to Tableau Server and is valid for any server software deployment.

Your tolerance for data loss will be determined by the individual application and the rate of data change in that application, among other things. For example, a 24-hour RPO means that at most, you will only lose the data from the last 24 hours.

The Simplest DR Strategy

In many cases, the simplest and most effective DR strategy is running nightly backups and having a standby server machine, warm or hot, staged and ready in the recovery data center. Compared to traditional BI tools, Tableau Server is quick to install, configure, and get running, and in most cases you may meet your business objectives by installing a clean server on a stand-by machine and restoring from a backup. The amount of time needed to install the server and restore a backup depends on the size of the backup and the server's complexity, but could easily be accomplished in a matter of hours.

Third-party DR Solutions

There are many third-party solutions available that provide 15-minute or longer RPOs that your IT teams may already be using. Although third-party backup recovery and DR solutions are outside the scope of this whitepaper, when used with such recovery systems, Tableau Server should be treated just like any other mission-critical application on the system, and the third-party backup recovery technology should be configured accordingly.

Tableau DR Capabilities

In case you can't use third-party DR technologies, Tableau Server includes a built-in native backup and recovery capability. Tableau Server can be backed up while still running, so you can create full,

hot backups. In addition, these hot backups are as smart as Tableau Server, and they only back up the extract files, the Repository databases, and the relevant configuration. You can also run backups while Tableau Server is not running—during a planned outage, for instance.

You can choose to back up your server every 24 hours if you have a 24-hour RPO, or you can change the backup schedule per your RPO. Your recovery time will depend on whether you have a software or hardware failure as well as your ability to provision new hardware on site. If you have another node standing by with the same instance of Tableau Server already installed and configured, restoring Tableau Server to production is essentially a quick command line admin operation using a saved backup file.

For most enterprises, disaster recovery policies also impose distance requirements between the data center where the server is primarily running and the data center that would be the recovery target (e.g. in the United States, DR sites must be at least 250 miles apart). In these situations, Tableau can be treated as any other business-critical application being managed for disaster recovery. Leverage Tableau's built-in backup and recovery technologies to meet your DR needs.

Beyond High Availability

Achieving high availability requires more than just the features and functionality built into Tableau Server to deliver Self-Service Analytics at scale. Failures can occur outside the scope of the Tableau Server software, such as hardware or network failures. You can achieve improved high availability by deploying Tableau Server on infrastructure that is more resilient to these types of failures such as virtualized HW in public or private clouds. Some vendors even provide live migration capabilities such as VMware vMotion.

True high availability comes from understanding what users need and following best practices and processes to meet those needs. Although Tableau makes high availability easier to achieve with its robust features, these should not replace best practices such as regular maintenance and scheduled backups. Tableau realizes the importance of business intelligence applications to help your business make decisions. Tableau Server is ready for your organizations most critical analytical questions.

About Tableau

Tableau helps people transform data into actionable insights that make an impact. Easily connect to data stored anywhere, in any format. Quickly perform ad hoc analyses that reveal hidden opportunities. Drag and drop to create interactive dashboards with advanced visual analytics. Then share across your organization and empower teammates to explore their perspective on data. From global enterprises to early-stage startups and small businesses, people everywhere use Tableau's analytics platform to see and understand their data.

Resources

[Tableau for the Enterprise: An IT overview](#)

[Tableau Server Scalability: A Technical Deployment Guide for Server Administrators](#)

[Server Admin Guide](#)

[Tableau on Amazon Web Services](#)

