One Identity Safeguard for Privileged Sessions 6.9.3

Scalability and High Availability in Safeguard

Contents

Introduction

This document describes the ways multiple appliances in the Safeguard product line can be deployed together.

Appliances of a Safeguard deployment: SPP and SPS

The backbone of a Safeguard deployment are two appliances: One Identity Safeguard for Privileged Passwords (SPP) and One Identity Safeguard for Privileged Sessions (SPS).

SPP and SPS appliances provide different functionality. You can use them together or independently.

- SPP provides asset and account discovery, password rotation and management, and access request workflow.

- SPS provides transparent or non-transparent interception of remote admin protocols, audit recording and video-like playback of sessions and analytics if One Identity Safeguard for Privileged Analytics (SPA) is licensed and enabled.

When used together, the two main operational modes are SPP-initiated (or Passwords-initated) and SPS-initiated (or Sessions-initiated).

- In SPP-initiated mode, users request access on the portal of SPP and when they are granted access, they are connected to the target account through SPS. See "Using SPS with SPP" in the Administration Guide.

- In SPS-initiated mode, users connect directly to a target server, SPS intercepts the traffic and fetches the required credentials from SPP.

SPP and SPS appliances solve scalability and high availability independently, but they can interoperate to ensure the correct operation of the entire deployment.

Most important terms around clustering

Clustering

The term clustering is often used with different meanings. SPP and SPS appliances can be clustered to provide:

- Shared configuration
- Scalability
- High Availability
- Disaster recovery
- Audit data replication
- Interoperation between SPP and SPS appliances

For clarity, we will use the more specific terms throughout this document where possible.

High Availability (HA)

Multiple SPP and SPS appliances can be connected to ensure high availability. This enables the continuation of vital technology infrastructure and systems.

Disaster recovery (DR)

SPP and SPS appliances can be connected to ensure immediate recovery after a natural or human-induced disaster. Disaster recovery reduces downtime and data loss.

Scalability

Connecting multiple appliances allows load distribution and scaling to loads beyond the serving capability of a single appliance, while ensuring that you can configure and operate the deployment as a single solution instead of multiple independent appliances. Both SPP and SPS clustering provide scalability features to reduce management and operational costs.

SPP-SPS Join

You can connect an SPP cluster to one or more SPS clusters to combine their functionality, for example, to provide password rotation and session recording for the same accounts.

Overview of clustering in SPP and SPS

One Identity Safeguard for Privileged Passwords (SPP)

SPP ensures shared configuration, scalability, high availability (HA), and disaster recovery through a single architecture. You can join 3 or 5 SPP appliances into a single cluster. All important information is replicated within the entire cluster and the cluster remains functional if some of the appliances fail. You can also distribute load between the appliances in the cluster.

Figure 1: Clustering in SPP
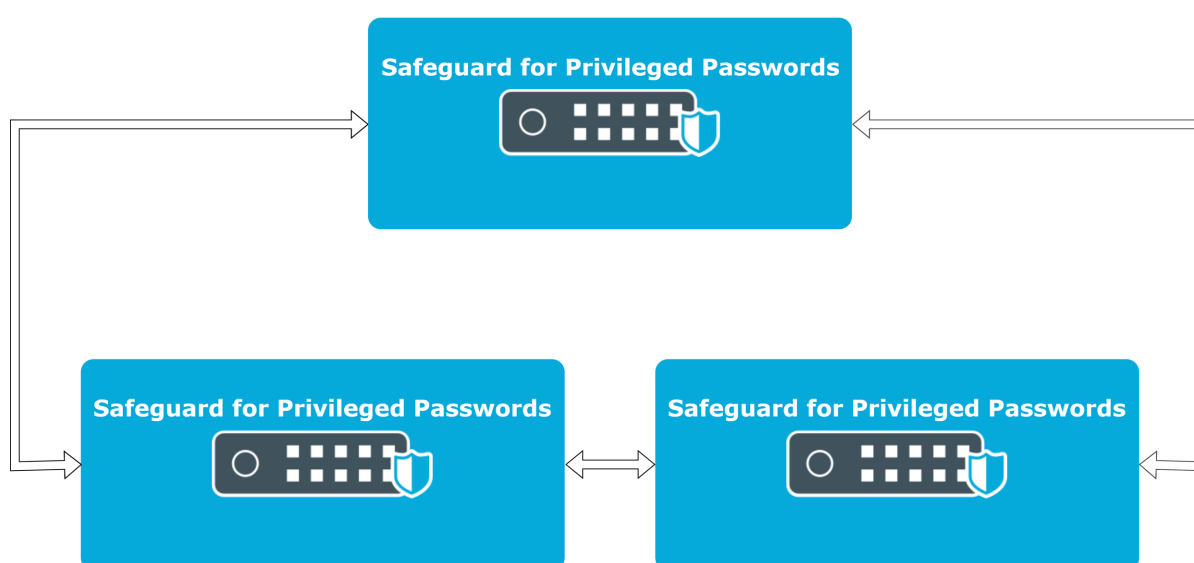
One Identity Safeguard for Privileged Sessions (SPS)

SPS follows a different approach and solves high availability and disaster recovery independently of shared configuration and scalability.

- Ensure high availability by adding a hot-spare pair to every SPS appliance that replicates all information from the first appliance and takes over all its functionality in case of a failure but serves no production traffic until the takeover occurs.

- Ensure shared configuration and scalability by clustering multiple SPS appliances (or HA pairs of appliances) together to control and monitor them from a single pane of glass.

To use HA and scalability at the same time, you need to configure them independently.

- A SPS HA pair always consists of two nodes: a primary and a secondary.
- An SPS scalability cluster consists of an arbitrary number of nodes with varying roles. For more information, see Scalability in One Identity Safeguard for Privileged Sessions (SPS)

Figure 2: Clustering in SPS
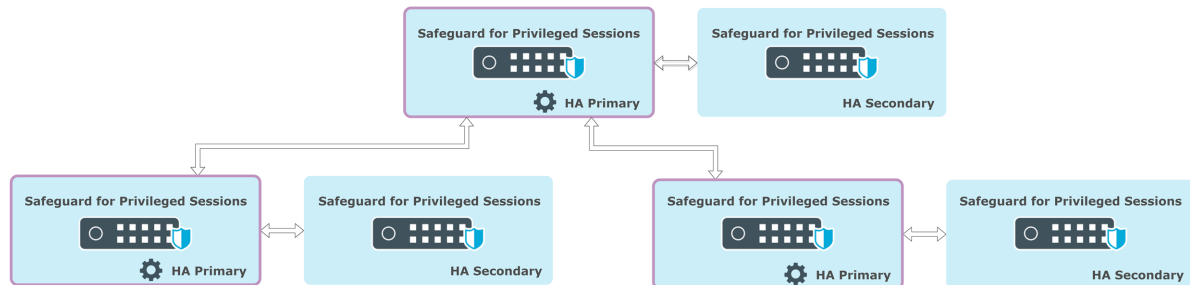
SPP and SPS clusters can work together and support each other's HA and scalability models through the SPP-SPS join.

The sections in this chapter describe how High Availability works in the Safeguard product line.

High Availability in One Identity Safeguard for Privileged Passwords (SPP)

In an SPP cluster, all vital data that is stored on the primary appliance is also stored on the replicas. The replicas provide a read-only view of the security policy configuration. You cannot add, delete, or modify the objects or security policy configuration on a replica appliance, but you can perform a password change, check operations and make password release and session access requests. Users can log in to replicas to request access, generate reports, or audit the data. You can request passwords and sessions from any appliance in a Safeguard cluster.

Figure 3: High Availability in SPP
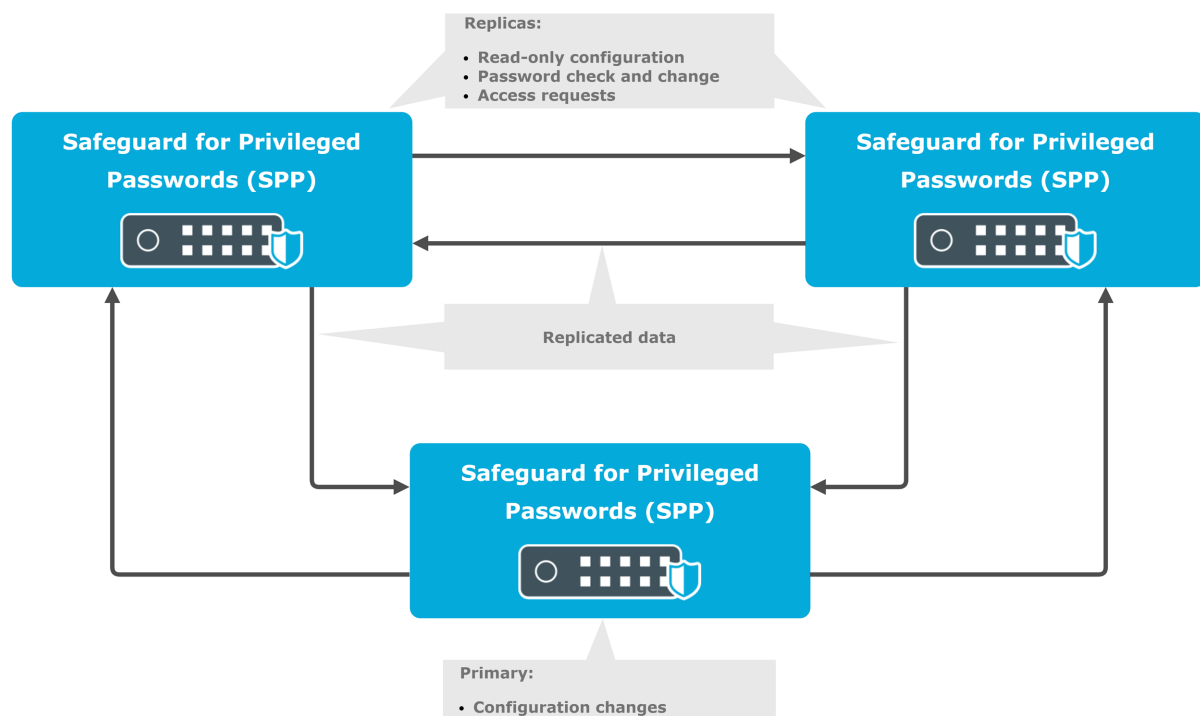
In the event of a disaster, where the primary appliance is no longer functioning, you can promote a replica to be the new primary appliance.

The full operation requires that the cluster has consensus (quorum). Consensus means that a majority of the clustered appliances is online and can communicate. If a cluster loses consensus, it goes into a read-only mode to prevent data inconsistencies and password check and it disables configuration changes. Offline Workflow Mode allows you to configure to either maintain or suspend access request workflows when consensus is lost.

For more information, see Disaster recovery and clusters in the One Identity Safeguard for Privileged Passwords Administration Guide.

High Availability in One Identity Safeguard for Privileged Sessions (SPS)

In case of physical appliances, add a hot-spare pair to every appliance to ensure high availability (HA). Place the two appliances ideally adjacent to each other and connect them with a direct cross cable. The secondary node in the pair replicates the entire disk contents of the primary node, including all configuration and audit recordings. The secondary node keeps all of its outside network connections in a disconnected state and serves no traffic until it takes over the role of the primary. You can initiate takeovers manually, or they are performed automatically if the primary node does not reply to the periodic heartbeat checks.

Figure 4: High Availability in SPS
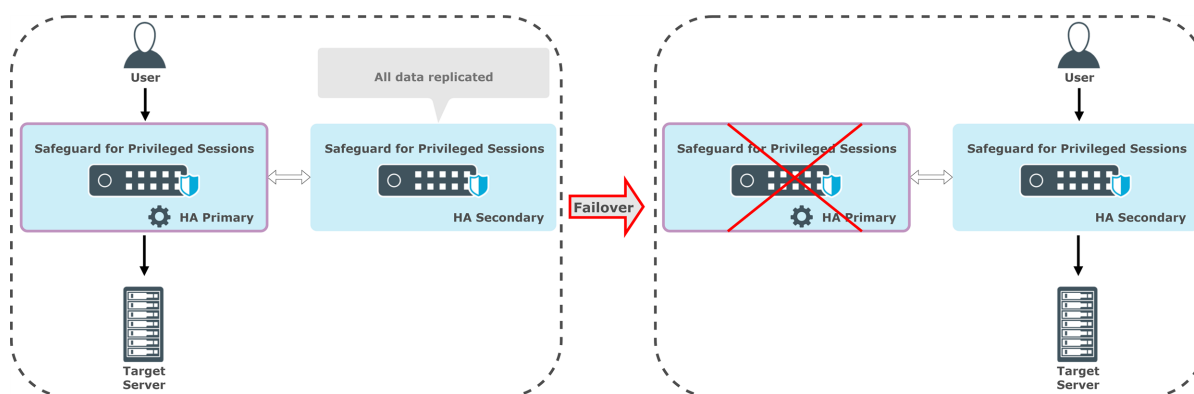
High availability is not available and not required for virtual appliances. All enterprise virtualization environments support high availability on the hypervisor level, which is more optimal compared to a host-based approach.

For more information, see "Managing a High Availability One Identity Safeguard for Privileged Sessions (SPS) cluster" in the Administration Guide.

High Availability in joint SPP and SPS deployments

To provide high availability (HA) in joint deployments, have at least 3 SPP nodes and add HA pairs to all SPS nodes. An SPP cluster of at least 3 nodes provides data redundancy and HA for the SPP functionality. In case a hot-spare SPS appliance needs to take over the primary role in a SPS HA pair, the SPP cluster automatically directs new connections to the new primary. This also solves the availability of all audit information in case of a hardware failure.

Figure 5: Joint SPP-SPS deployment
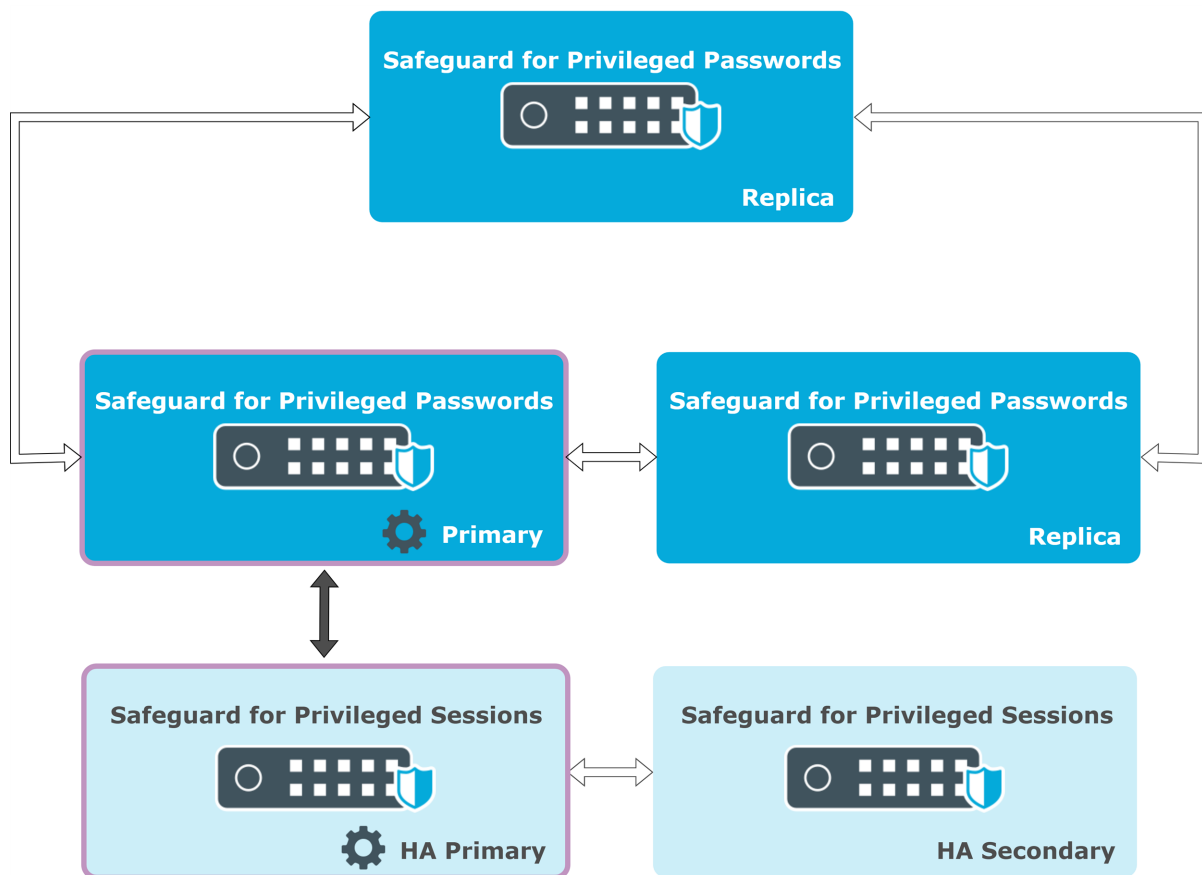
Backups

Although high availability (HA) protects against hardware failures, One Identity recommends enabling backups for both SPP and SPS appliances in both virtual and hardware deployments. Backups provide additional protection against:

- Software errors
- Mistakes that administrators make
- Large-scale disasters that affect many nodes of a cluster

However, backups alone do not provide a sufficient level of high availability because data during backup periods can be lost and a full restore from a backup may lead to a long period of service outage.

For more information on configuring backups, see the respective sections in the *Administration Guide*:

- SPP: Backup and restore in the One Identity Safeguard for Privileged Passwords Administration Guide
- SPS: "Data and configuration backups" in the Administration Guide

The sections in this chapter describe how scalability works in the Safeguard product line.

Scalability in One Identity Safeguard for Privileged Passwords (SPP)

The primary appliance in an SPP cluster automatically delegates platform management tasks such as password check and password change to appliances based on task load. Adding more appliances to the cluster allows performing more of these tasks.

You can customize load balancing through Managed Networks. **Managed Networks** are named lists of network segments serviced by a specific SPP appliance. Using **Managed Networks**, you can:

- Distribute the load so there is minimal cluster traffic.
- Use the appliances closest to the target asset to perform the task.

An SPP cluster has a default managed network that consists of all cluster members.

You can perform password request workflows through any appliance in the cluster if the cluster is healthy. For healthy clusters no automatic load balancing is performed.

Figure 6: SPP-managed networks
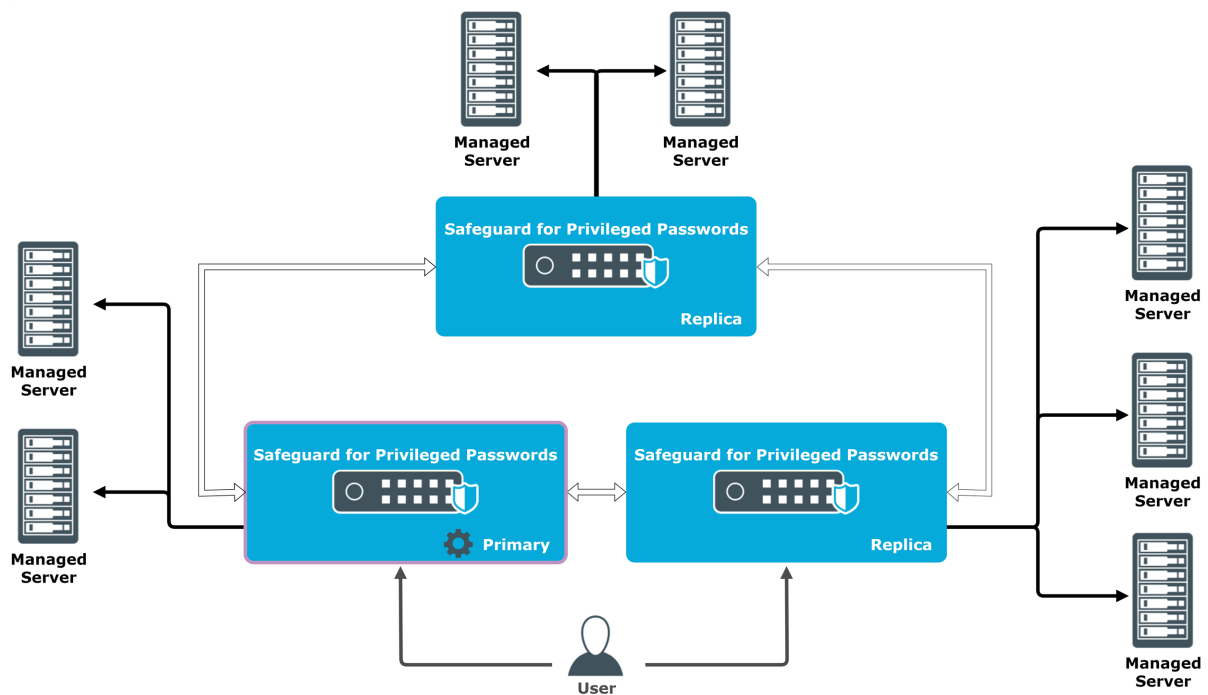
For more information on **Managed Networks**, see Managed Networks in the One Identity Safeguard for Privileged Passwords Administration Guide.

Scalability in One Identity Safeguard for Privileged Sessions (SPS)

You can join multiple SPS appliances or high availability (HA) pairs of appliances into a cluster and manage them from a single pane of glass.

The SPS cluster does not provide load balancing. You can set up network connections in a way that distributes the load among them. If SPP and SPS are used together, you can also use SPP to distribute the traffic. For more information, see Scalability in joint SPP and SPS deployments.

You can replicate the configuration of a primary node among the entire cluster.

For more information, see "Managing a cluster with configuration synchronization without central search" in the Administration Guide.

Figure 7: SPS-managed networks
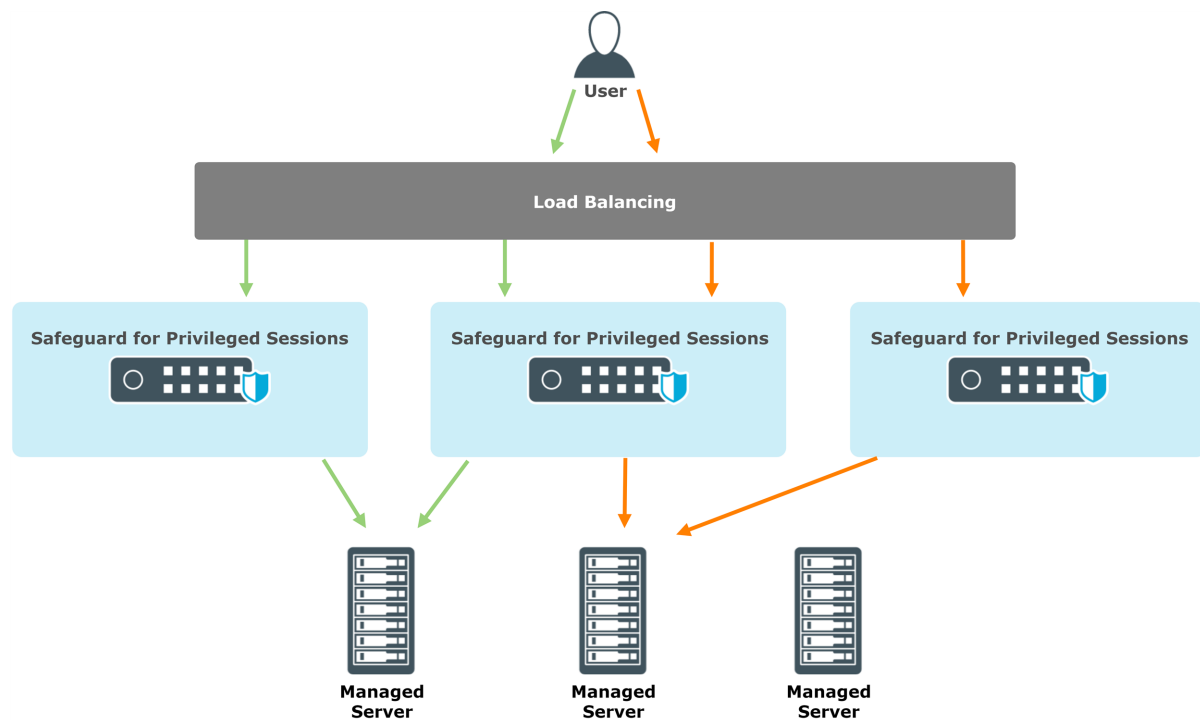
You can also make all audit information about the recorded sessions from all appliances available on a single search interface. This requires a dedicated search appliance or HA pair.

For more information, see "Managing a cluster with central search configuration and configuration synchronization" in the Administration Guide.

Scalability in joint SPP and SPS deployments

In case of SPP-initiated workflows, you can assign SPS scalability clusters to Managed Networks. The SPP cluster periodically checks the load on the members of the SPS cluster and assigns new connections to the best available appliance.

Figure 8: SPP-initiated workflow
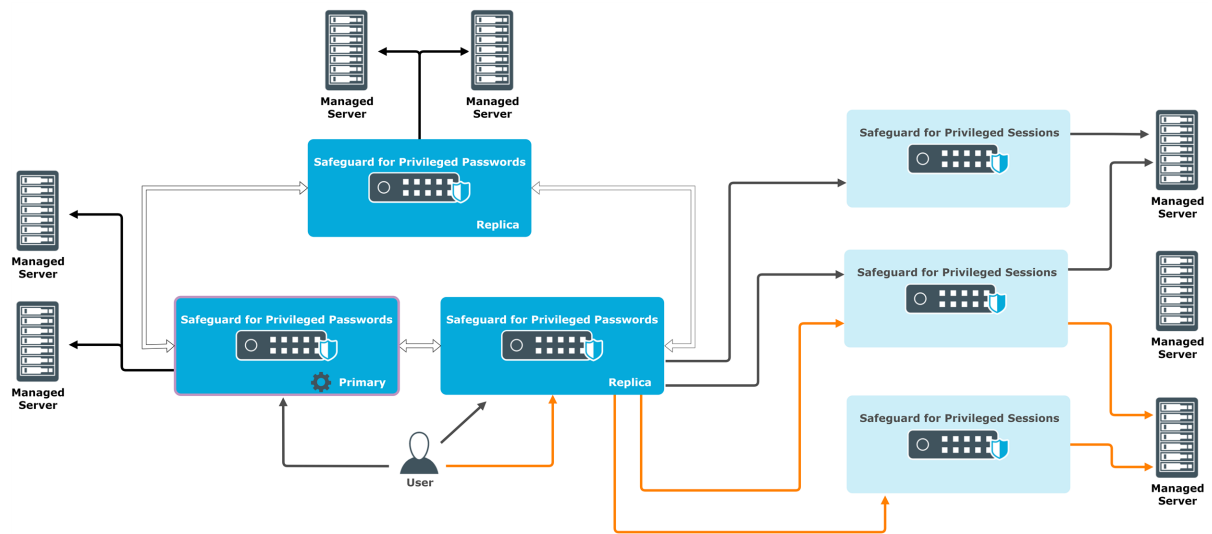
In case of SPS-initiated workflows, SPS appliances always target the primary appliance of the SPP cluster, but the queries do not usually require scaling out to multiple SPP appliances.

Figure 9: SPS-initiated workflow
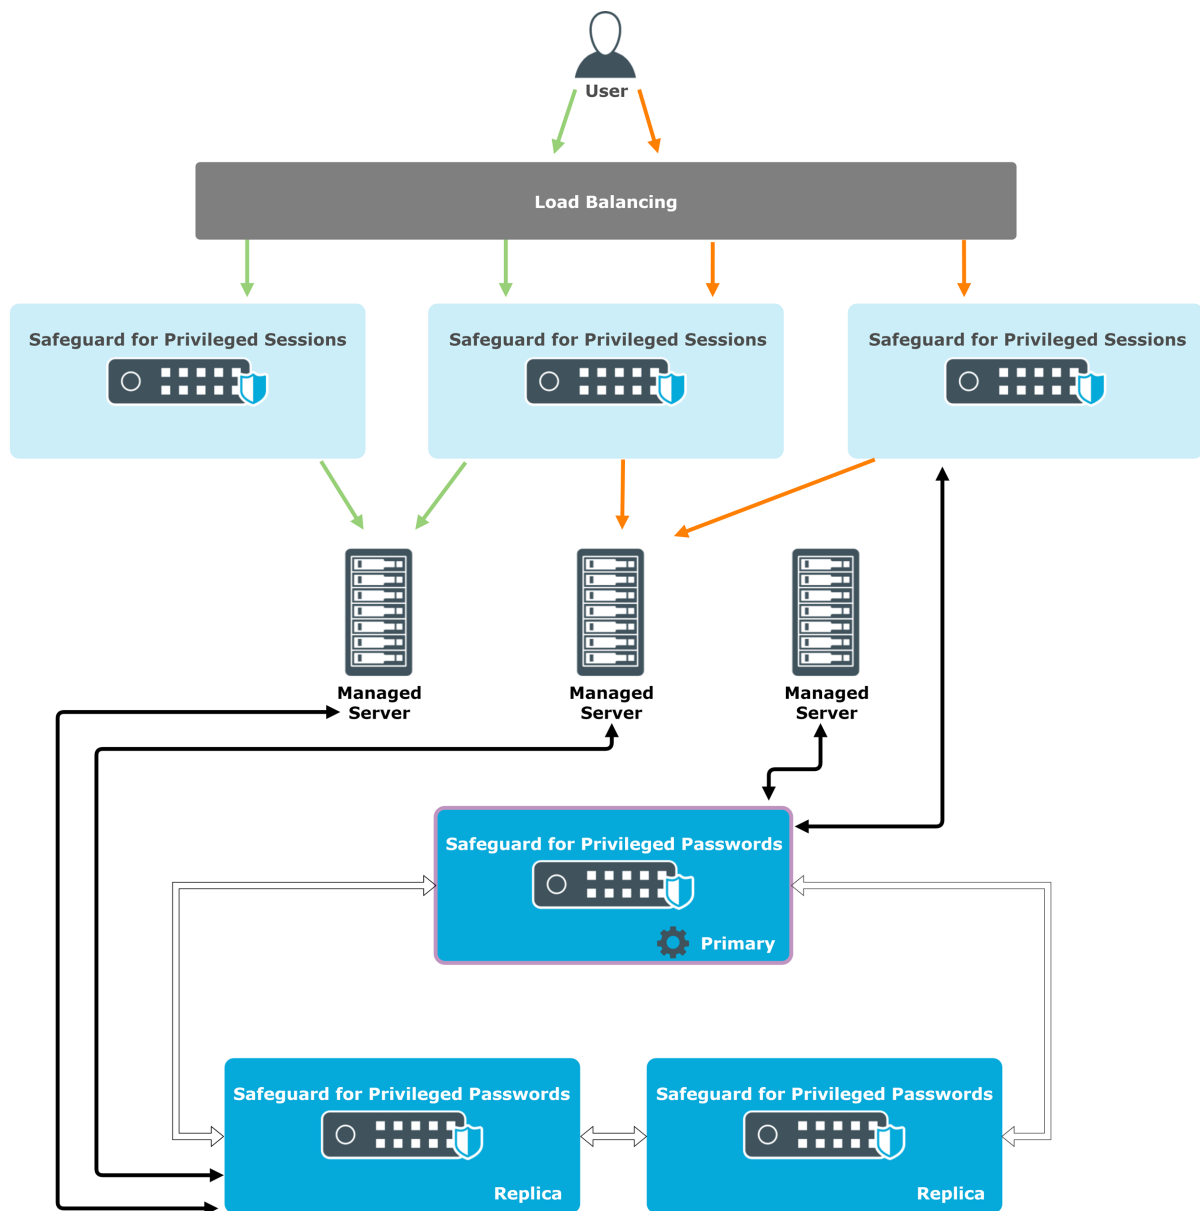
The sections in this chapter describe how Disaster Scenarios work in the Safeguard product line.

Disaster Scenarios in One Identity Safeguard for Privileged Passwords (SPP)

Failure of a replica node

If a replica node fails, the cluster detects that the node is out of the circulation and automatically redirects traffic. The cluster replicates all vital data between the nodes to prevent data loss.

One Identity recommends that managed networks contain multiple appliances so other nodes can take over the tasks of the failed node. You can reconfigure or disable managed networks to provide continuity of service.

Failure of the primary node

If the primary node fails, normal operation continues, but you cannot make changes to the configuration. You can promote any of the replica nodes manually to be the new primary node.

Failure of more than half of the cluster

If more than half of the cluster fails, the cluster switches into a read-only mode where you cannot make changes to the configuration and the cluster pauses password check and password change tasks. You can use **Offline Workflow** to manually or automatically restore the access request workflow. If the majority of the appliances have failed, you can use **Reset Cluster** to change to a new primary node without consensus. If all of the appliances in the cluster have failed, restore a backup.

Losing connectivity between appliances

The part of the cluster that becomes isolated and detects less than half of the original cluster switches into read-only mode, while the rest of the cluster remains active. In **Offline Workflow** you can configure the isolated nodes to continue serving access requests. When connection is re-established, the appliance state is automatically synchronized.

Disaster Scenarios in One Identity Safeguard for Privileged Sessions (SPS)

Failure of a node in a High Availability (HA) pair

If the failed node was the primary node, the hot-spare node automatically takes over the IP address and all traffic. Ongoing connections are disconnected. The cluster replicates all vital data between the pairs to prevent data loss. After replacing the failed node, perform a resynchronization.

NOTE: Resynchronization can last up to 24 hours.

In all of the following failure scenarios, if the failed node has a HA pair, the pair takes over all functionality automatically and the same recovery steps are required:

1. Replace the failed node
2. Resynchronize

Failure of a managed node (non-primary appliance) in the scalability cluster

If the managed node did not have a HA pair, traffic going through the managed node stops. The network configuration handles the outage and redirects traffic to another appliance in the cluster. In case of SPP-initiated workflows, SPP attempts to redirect the traffic towards a different SPS when the SPS configuration primary detects the outage. If central search is enabled, you can still perform searches, but video-like playback of sessions is not available.

Failure of the configuration primary in a scalability cluster

If the configuration primary did not have a HA pair, functioning nodes keep serving connections, but you cannot make any configuration changes in the cluster. You cannot move the configuration primary role to a different appliance, because the role must be restored from a backup.

Failure of the search master in a scalability cluster

If the search master did not have a HA pair, you cannot search in audit information, but all other functionality works. The other nodes buffer audit information until the search master node becomes available again. The nodes can survive approximately 24 hours of downtime when operating at full capacity, then they stop accepting new connections.

Losing connectivity between HA pairs

When connectivity is lost between the nodes of a HA pair, both appliances check whether they can detect the outside network. If a node only loses connection to the other node of the HA pair, both nodes start to operate as primary nodes. Two primary nodes cause

service outage that you have to recover manually. To prevent this, One Identity recommends you to configure redundant HA links between the nodes.

For more information, see "Redundant heartbeat interfaces" in the Administration Guide.

Losing connectivity between nodes in a scalability cluster

If connection is lost between nodes in a scalability cluster, individual nodes continue serving new connections, but some of the functionality is lost until an outage recovery, such as making configuration changes or searching in new audit information.

Disaster Scenarios in joint SPP and SPS deployments

The scenarios are handled the same way in the case of joint deployments and individual SPP and SPS clusters. Additional possible issues include:

SPP-initiated workflows

Losing the SPS configuration primary

Losing the SPP primary appliance: all of the functionality works until no configuration changes are required, SPP nodes reach out to SPS directly.

SPS-initiated workflows

Losing the SPS configuration primary

Losing the SPP primary appliance:

- SPS stores all the SPP cluster members and for password requests SPS connects to the first available node, starting with the primary node. If the primary node is not available, SPS makes a connection attempt to other nodes.

- SPS fetches information about the SPP cluster from the IP address that the SPS was joined with originally (the primary node at the time of the join). If an SPP is available at the IP address, which is not necessarily the primary node, the SPS cluster detects changes in the configuration.

- If a SPP cluster member can be replaced or reconnected on the same address, you do not need to do anything manually and traffic is not interrupted.

- If a SPP cluster member is permanently lost and cannot be replaced or reconnected, you need to use the new primary node to unjoin and rejoin the clusters.

One Identity solutions eliminate the complexities and time-consuming processes often required to govern identities, manage privileged accounts and control access. Our solutions enhance business agility while addressing your IAM challenges with on-premises, cloud and hybrid environments.

Contacting us

For sales and other inquiries, such as licensing, support, and renewals, visit https://www.oneidentity.com/company/contact-us.aspx.

Technical support resources

Technical support is available to One Identity customers with a valid maintenance contract and customers who have trial versions. You can access the Support Portal at https://support.oneidentity.com/.

The Support Portal provides self-help tools you can use to solve problems quickly and independently, 24 hours a day, 365 days a year. The Support Portal enables you to:

- Submit and manage a Service Request
- View Knowledge Base articles
- Sign up for product notifications
- Download software and technical documentation
- View how-to videos at www.YouTube.com/OneIdentity
- Engage in community discussions
- Chat with support engineers online
- View services to assist you with your product