Fortanix®

Whitepaper

Resiliency in Fortanix DSM SaaS

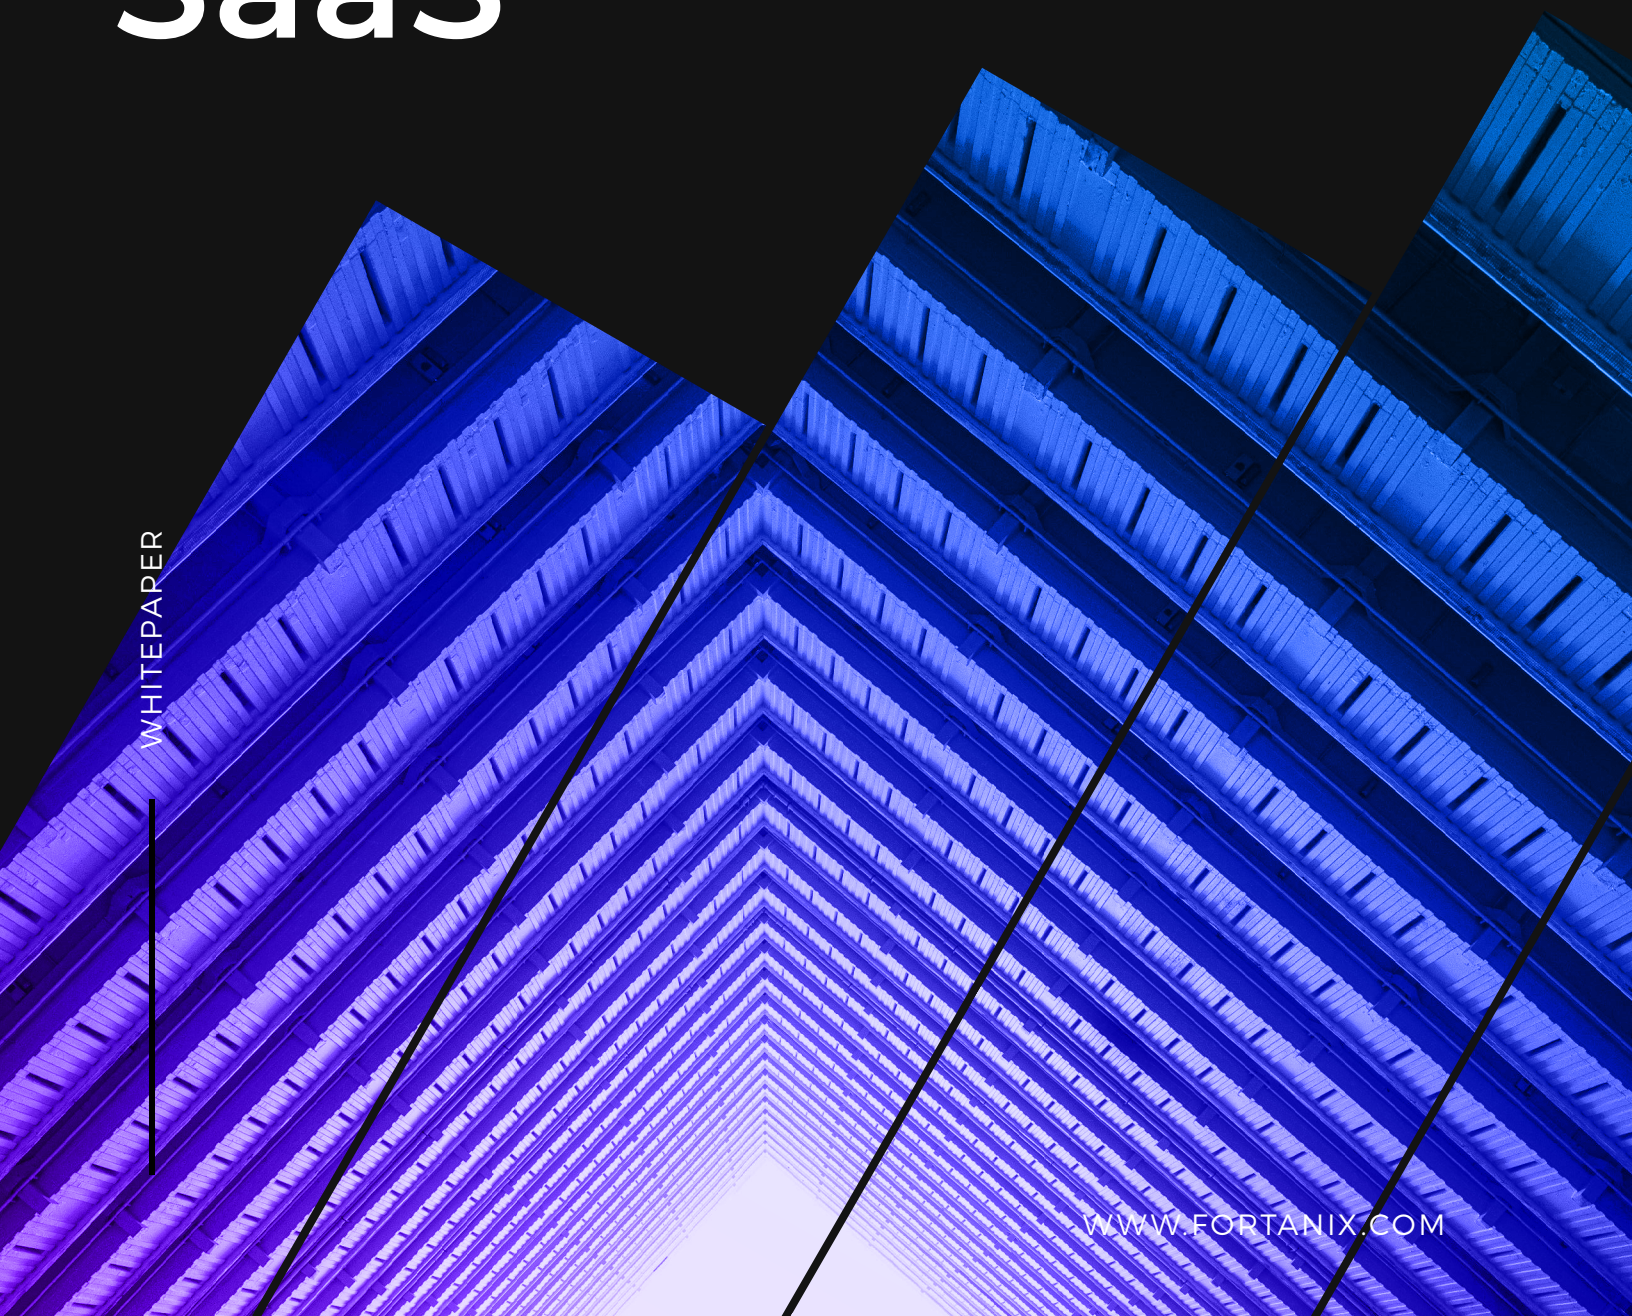

WHITEPAPER

WWW.FORTANIX.COM

Introduction

Fortanix Data Security Manager (DSM) SaaS is a global service available over the Internet and offers security services, such as cryptographic key lifecycle management, encryption, tokenization, and secrets management. DSM SaaS is often used by enterprises in business-critical applications, which makes it necessary for the service to be extremely reliable and resilient.

In this white paper, we first state the objectives for resiliency in DSM SaaS and then describe the architecture and the operational procedures which help us achieve those objectives. In the end, we discuss further enhancements that are being worked on. This paper does not cover the security objectives and security properties of DSM SaaS which are covered in a different white paper.

Resiliency Objectives

We have defined the following objectives for resiliency in DSM SaaS:

High Availability (HA)	Capacity	Latency	Fairness	Disaster Recovery (DR)
DSM SaaS offers an SLA uptime of 99.95%. However, the system is de-signed to achieve continuous availability.	DSM SaaS maintains capacity to offer predictable performance even at peak load and to withstand DDoS attacks. Capacity can be easily scaled up by adding more hardware.	The latency of an API request to DSM SaaS should not be impacted by an increase in load or congestion when operating under capacity.	DSM SaaS is a multi-tenant service, and the objective of fairness is to ensure that a noisy tenant does not starve other tenants.	The RTO and RPO for DSM SaaS is zero except for some catastrophic scenarios.

The resiliency objectives are met by three factors – the design of the DSM SaaS architecture, our Software Development Lifecycle and the DSM SaaS operations.

info@fortanix.com | 3910 Freedom Circle Suite 104, Santa Clara, CA 95054, United States

DSM SaaS Architecture

DSM SaaS is globally deployed across 6 independent "regions," each consisting of a cluster of DSM "nodes" deployed across 3 "sites" or data centers (DCs). Each DSM node is a FIPS 140-2 Level 3 validated Fortanix FX2200 HSM appliance. Each region is labeled based on its geographic location and offers the service through a unique URL as listed below. There are no network links between these regions, and no data gets exchanged between them. In the future, there is a plan to use a common login to authenticate to each region through a common and global Identity and Access Management (IAM) service. Even with a global IAM, the regional cluster of nodes will continue to operate regionally only exchanging data and communication within their cluster.

The architecture can be conceptualized in the form of the network topology (how an incoming request is routed to a DSM node), the compute cluster (how a request is processed) and the data and storage layer (how data is stored and replicated).

Network Topology

Each region of DSM SaaS has a unique URL and has 2 open ports – 443 and 5696, on which it receives requests over the Internet using REST API (https://support.fortanix.com/apidocs) and KMIP API (https://docs.oasis-open.org/kmip/spec/v1.2/os/kmip-spec-v1.2-os.html) respectively.

A (primary) network provider in each region is responsible for routing requests coming from the Internet to the router at the nearest data center. This is done using Anycast routing. The primary network provider also implements a DDoS protection layer that filters out any DDoS attempts that are detected.

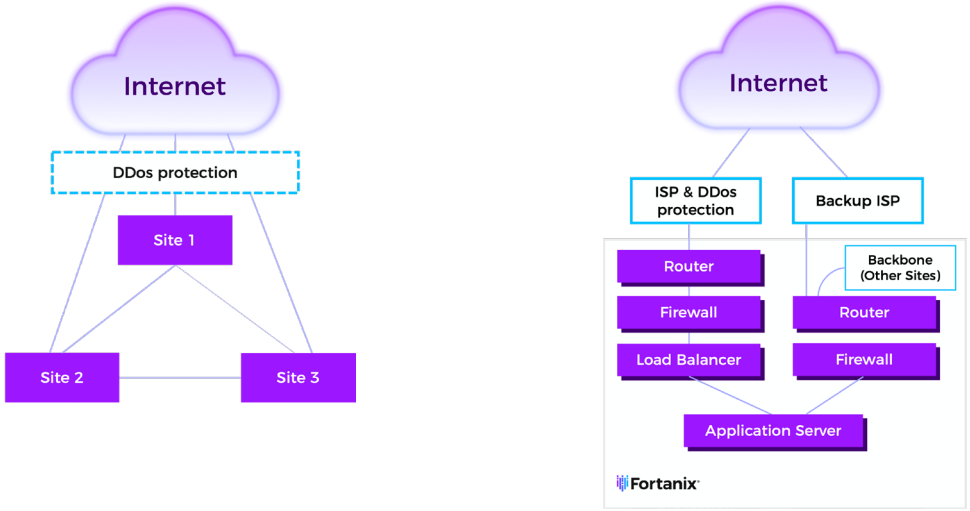

To ensure reliable connectivity, a backup network provider is also available with an automatic failover from primary to backup built in. The backup network provider operates over its own independent circuit, and the backup transit traffic is routed through a dedicated backup edge device. The backup network provider does not have DDoS protection.

Each site has two Layer 3 routers which operate in active-active configuration. The routers support Anycast routing for traffic coming from the Internet. The routers also provide connectivity to the other sites which is required for intra-region clustering. This is done over the Equinix fabric via two diverse, independent connections. There is also a backup edge device which is connected to the backup network provider. This configuration allows for device failures to be handled with N+1 redundancy (i.e., a single component failure should have no impact on availability).

Once an API request reaches a site, it is directed to one of the multiple DSM nodes deployed in that site by a Layer 4 load balancer. The load balancing mechanism currently used is to route requests to the least loaded node, measured by the number of open TCP connections.

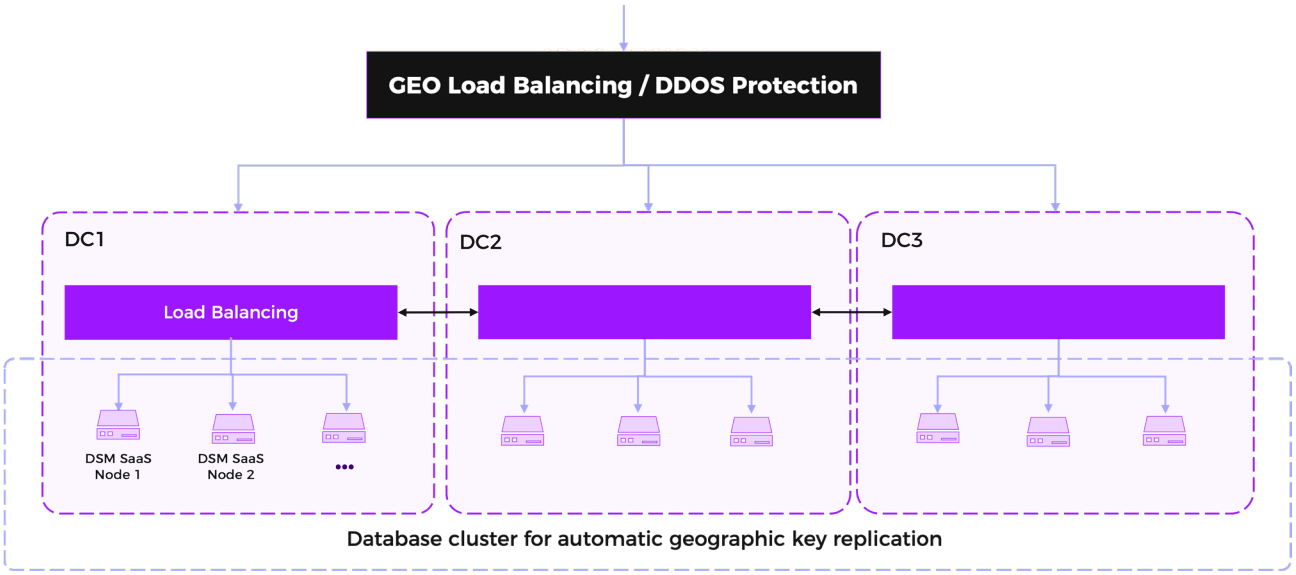

Compute Cluster

The DSM cluster operates as an Active-Active cluster, which enables every node in a DSM cluster to respond to any incoming request. This is possible because every node runs the same software and has access to the same shared database.

Every site in a DSM region is served using several nodes depending on the capacity requirements of the cluster. The capacity of the cluster increases linearly with the size of the cluster, so the capacity of the cluster can be scaled up by adding additional DSM nodes to the cluster.

Every node implements a "quality of service" or QoS framework to achieve fairness. This is done by implementing a weighted priority queuing scheme. Any incoming request is inserted into a queue based on their entitlements (for example, paid customer vs customer on free trial), and their current usage pattern (to throttle a sudden burst). The nodes in a cluster synchronize information to prevent a burst of traffic from propagating across the nodes. This scheme also prevents starvation of requests from any source.

Data and Storage Layer

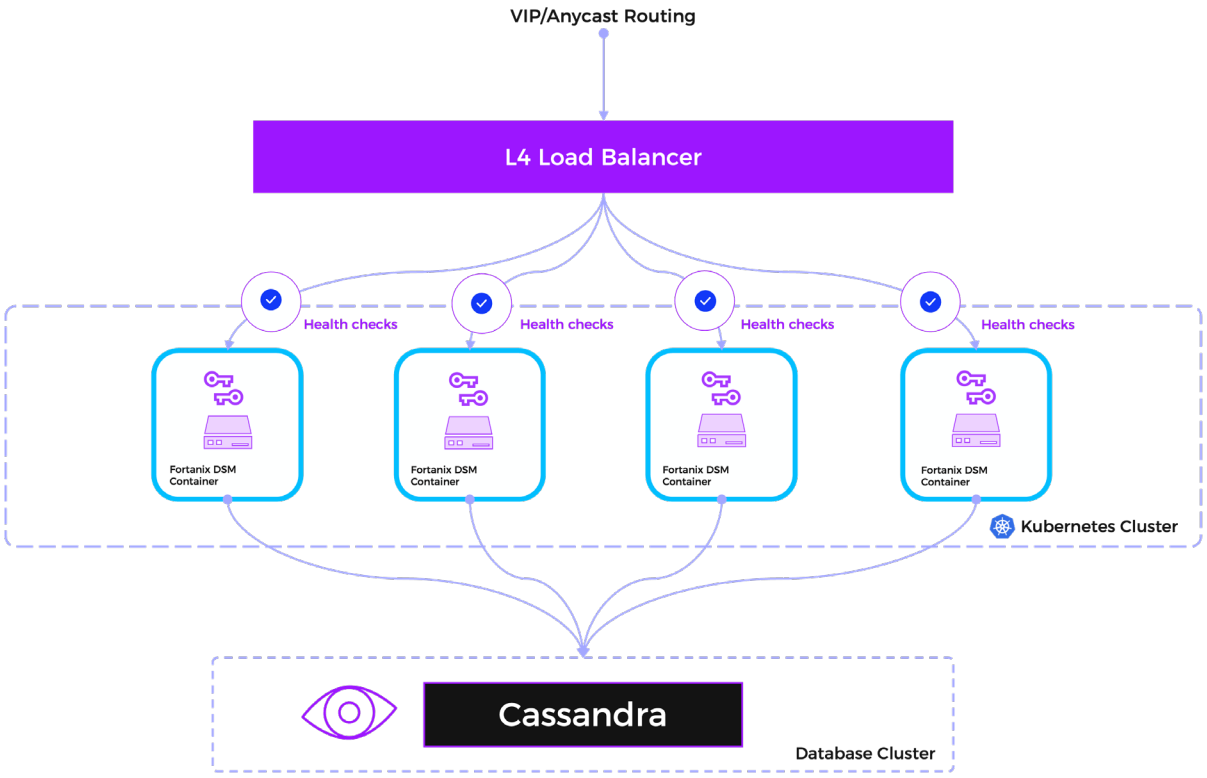

Every node in a regional cluster is enabled to respond to any incoming request. This is achieved by having every node have an identical and unified view of all the data stored in a DSM cluster – key material, metadata, configuration, credentials, logs, etc. Data is stored in a single, distributed, and clustered instance of Cassandra database that runs across all nodes in a DSM cluster. Data is stored using a "strong consistency" paradigm, which ensures consistency of API operations regardless of where a request lands in a cluster.

To achieve strong consistency when writing into the database, a majority of nodes in the cluster (or quorum) should be available. Since each site in a region typically has similar number of nodes, it means that writes can be performed with strong consistency even when a site or data center is completely cut-off from the other two sites.

Daily backups of the database are taken and stored in a remote location. These backups can be used to rebuild a DSM SaaS cluster in case of a catastrophic failure where the entire DSM SaaS cluster in a region is lost.

How the Architecture Helps Achieve Resiliency Objectives

The architecture described in this section helps achieve the resiliency objectives in the following ways:

HA	Capacity	Latency	Fairness
DSM SaaS has redundancy built into the network layer, the compute layer, as well as the data layer. As long as a majority of nodes in the cluster are available (e.g., even if an entire site is unavailable), and the incoming requests can be routed over the network to a healthy node in a healthy site, all operations will succeed.	The DDoS protection layer which detects and filters out any DDoS attacks enables DSM SaaS to maintain predictable capacity. The capacity can be scaled up by simply adding more DSM nodes as the compute cluster is linearly scalable.	Anycast routing routes all incoming requests to the nearest site, thus reducing network latency. The DDoS protection layer ensures the site has enough capacity thus reducing network contention or congestion which would impact latency. The load balancing in each site sends the request to the least loaded node, which ensures the processing latency for the request is minimized.	Fairness objectives are achieved using the QoS mechanism.

info@fortanix.com | 3910 Freedom Circle Suite 104, Santa Clara, CA 95054, United States

Software Development Lifecycle (SDLC)

The Fortanix SDLC focuses on producing and deploying software which can help achieve the resiliency objectives of DSM SaaS. Some of the steps of SDLC particularly relevant to resiliency are described in this section.

Quality Assurance (QA)

Fortanix performs extensive and rigorous testing before any release which consists of unit and integration testing, performance and stress testing, and regression testing. A dedicated QA team is responsible for qualifying a build before it can be approved for release.

When new APIs are added in a software version, extensive testing is done to ensure that the new and the related old API can work together during the software update rollout as software updates are meant to be non-disruptive. For certain new features deemed high risk, we roll them out across two successive software updates so that we can observe the feature in its intermediate state and further reduce the risk of disruption.

Staging Environment

Fortanix maintains a scaled down replica of a region of DSM SaaS as the staging environment. Any new release of DSM SaaS is first deployed in the staging environment. After the software update, a full set of test cases are run, and a soaking period is used to further qualify the software version for production release.

Software Updates

Software updates on DSM SaaS are done about once a month. Software updates are designed to be transparent to end users and not service impacting.

Software updates are done in a staged fashion where different regions of the SaaS are updated one after the other. This allows us to pause the rollout across regions and roll back if necessary if an issue is found.

Inside a region, software updates are done progressively across the nodes in the cluster. This ensures that a majority of nodes are always available in the cluster, and the cluster continues to operate normally during the update.

How the SDLC Helps Achieve Resiliency Objectives

Fortanix SDLC helps achieve resiliency in the following ways:

High Availability (HA)	Capacity
Extensive testing is done to ensure that a new software release does not impact the availability of DSM SaaS, either during the software update process or post-update.	A bad software update has the potential to bring down an entire cluster, but extensive tesing and the deployment in the staging cluster is done to minimize the possibility of that happening. However, if a buggy software update still makes it to the production update process, the update process is closely monitored and can be rolled back if a problem is found.

info@fortanix.com | 3910 Freedom Circle Suite 104, Santa Clara, CA 95054, United States

SaaS Operations

Fortanix has a dedicated SaaS Operations (SaaS ops) team which works across the globe and has 24*7 coverage. The SaaS ops team is responsible for all operational aspects of DSM SaaS – monitoring, responding to issues, performing change procedures such as software updates, performing periodic disaster drills, capacity planning, providing customer notifications, etc. While the SaaS ops team is there to support the operations of the service, the adherence to customer SLAs is achieved through the architecture and the automation that we have built.

Monitoring and Alerting

Fortanix uses multiple observability tools to monitor the health of DSM SaaS. This includes monitoring of the network infrastructure, the compute infrastructure, the database and various services running inside DSM SaaS, and the health of external services that DSM SaaS relies on, such as the DDoS protection service.

If a fault is observed in any of the systems being monitored, alerts are sent immediately to the on-call members of the team, which get escalated quickly to higher levels if not responded to quickly.

Disaster Drills

The SaaS ops team maintains a runbook for various disaster scenarios and performs periodic drills to practice and train on the procedures outlined in the runbook. These drills include scenarios, such as having an entire site may go down, but the service should not be degraded, and we should be able to restore the site as quickly as possible.

The results of the disaster drills performed annually become part of our SOC2 Type 2 compliance report. The latest report can be obtained upon request.

Capacity Planning

The monitoring framework for DSM SaaS captures data on usage trends which the SaaS ops team uses to determine future capacity needs. This allows us to proactively provision more capacity by adding hardware nodes to a region before the load reaches peak capacity.

Communication and Notification

Fortanix SaaS ops team maintains https://status.fortanix.com to provide notification and updates to customers. Any change procedure or software update is done in a maintenance window, and the customers are notified ahead of time about them. Any outage or degradation of service is also mentioned on this webpage, and the subscribers are alerted about it.

info@fortanix.com | 3910 Freedom Circle Suite 104, Santa Clara, CA 95054, United States

Change Procedure

Any change to production environment of DSM SaaS is heavily controlled and follows a rigorous planning and approval cycle. Where possible, automation is used to avoid human error. An internal ticketing system is used to request any change where the requester is expected to fill out a form describing the change in detail and identifying any perceived risks. This ticket is then reviewed and approved. A notification is sent to customers informing them about a maintenance window when the change would be implemented, and then the change is performed during that window.

How Operations Help Achieve Resiliency Objectives

Fortanix SaaS Operations help achieve the resiliency objectives in the following ways:

HA	DR	Capacity	Latency
By actively always monitoring the health of DSM SaaS and having a robust alerting mechanism, Fortanix SaaS ops team is able to act quickly if there is a degradation in service and restore full service before availability is impacted.	Regular disaster drills keep the SaaS ops team trained to be able to react quickly when an actual disaster may happen.	Through active capacity planning, the SaaS ops team ensures that DSM SaaS has enough capacity to satisfy peak and average load.	Latency may go up when there is a site failure, and the request may get redirected to a site which is farther from the source. The SaaS ops team works actively to restore the site to ensure network latency can go back to normal after the incident.

Future Enhancements

One feature we plan to add in DSM SaaS that will further improve resiliency is support for Virtual Private Cloud (VPC). This would allow tenants to create a mechanism to establish a direct and dedicated connection between their infrastructure and DSM SaaS, thus bypassing the Internet. The VPC connectivity provides improved resiliency at the network layer. A direct network would have predictable, and generally lower latency and would be less susceptible to disruption that might happen over the Internet.

info@fortanix.com | 3910 Freedom Circle Suite 104, Santa Clara, CA 95054, United States

Catastrophic Scenarios

Fortanix DSM SaaS offers best-in-class resiliency, but there are some extremely unlikely catastrophic events which could still impact availability, or impact our ability to keep RPO and RTO to zero:

All sites in a region becoming unreachable: We currently maintain a primary and a backup network provider to route traffic from the Internet to DSM SaaS. Loss of connectivity to the service would require both network providers to fail simultaneously. In case of this unlikely event, the resolution would be either through the restoration of one of the network providers, or by enabling another network provider at the time of the outage.

All sites in a region becoming simultaneously unavailable to serve requests: All sites in a region becoming simultaneously unavailable to serve requests – For this to happen, there has to be a nationwide power outage, or a national disater knocking off data centers in different parts of a region simultaneously. The resolution would be to wait for the DSM nodes and networking equipment to be restored in at least two out of three sites in a region.

Sites in a region get disconnected from each other: For this to happen, there needs to be multiple simultaneous network failures in intra-cluster connectivity. While this is extremely unlikely, as there are multiple layers of redundancy built into the network fabric we rely on for intra-cluster connectivity, the resolution would be to wait for restoration of connectivity. In this failure event, it's likely that individual sites will continue to support a read-only mode which would allow connected applications to be unaffected.

Loss or corruption of data or software in a region: In case of an extremely unlikely event of the data or code in a cluster getting wiped out or becoming unrecoverable, daily backups ensure the cluster can be rebuilt and restored to an earlier point in time.

One or more sites in a region become unreachable: This could happen when there is a fault with the primary ISP, for example, during a DDoS attack which impacts the ISP operations. This could cause a partial outage in the DSM SaaS cluster, and the automatic site-to-site failover would not kick in because the fault is external to the DSM SaaS cluster. If this scenario persists, the Fortanix SaaS ops team needs to intervene, and divert traffic to one or more sites in the region which are healthy. This scenario could lead to a temporary and partial outage on DSM SaaS. Future enhancements would support a private connectivity options to protect against this.

Conclusion

In conclusion, Fortanix DSM SaaS is designed to be highly resilient. The resiliency objectives of HA, DR, predictable capacity, low latency, and fairness, is achieved using a combination of a highly-redundant architecture, an extensive resiliency-focused SDLC, and a strong SaaS ops team following rigorous operational procedures. Fortanix DSM SaaS customers should rest assured that Fortanix uses best-in-class methods in all these areas to stand behind its SLA commitments. Future enhancements will continue to make DSM SaaS more resilient.

info@fortanix.com | 3910 Freedom Circle Suite 104, Santa Clara, CA 95054, United States