

Architecture Reference Diagrams

Administrator Guide

Version: 2023.x

Publication Date: 12/19/2024

Architecture Reference Diagrams Administrator Guide

Version: 2023.x, Publication Date: 12/19/2024

© Delinea, 2024

Warranty Disclaimer

DELINEA AND ITS AFFILIATES, AND/OR ITS AND THEIR RESPECTIVE SUPPLIERS, MAKE NO REPRESENTATIONS ABOUT THE SUITABILITY OF THE INFORMATION CONTAINED IN THE DOCUMENTS AND RELATED GRAPHICS, THE SOFTWARE AND SERVICES, AND OTHER MATERIAL PUBLISHED ON OR ACCESSIBLE THROUGH THIS SITE FOR ANY PURPOSE. ALL SUCH MATERIAL IS PROVIDED "AS IS" WITHOUT WARRANTY OF ANY KIND. DELINEA AND ITS AFFILIATES, AND/OR ITS AND THEIR RESPECTIVE SUPPLIERS, HEREBY DISCLAIM ALL WARRANTIES AND CONDITIONS WITH REGARD TO SUCH MATERIAL, INCLUDING ALL IMPLIED WARRANTIES AND CONDITIONS OF MERCHANTABILITY, FITNESS FOR A PARTICULAR PURPOSE, TITLE AND NON-INFRINGEMENT.

THE MATERIAL PUBLISHED ON THIS SITE COULD INCLUDE TECHNICAL INACCURACIES OR TYPOGRAPHICAL ERRORS. CHANGES ARE PERIODICALLY ADDED TO THE INFORMATION HEREIN. DELINEA AND ITS AFFILIATES, AND/OR ITS AND THEIR RESPECTIVE SUPPLIERS, MAY MAKE IMPROVEMENTS AND/OR CHANGES IN THE MATERIAL DESCRIBED HEREIN AT ANY TIME.

Disclaimer of Liability

IN NO EVENT SHALL DELINEA AND ITS AFFILIATES, AND/OR ITS AND THEIR RESPECTIVE SUPPLIERS, BE LIABLE FOR ANY SPECIAL, INDIRECT, OR CONSEQUENTIAL DAMAGES (INCLUDING LOSS OF USE, DATA, PROFITS OR OTHER ECONOMIC ADVANTAGE) OR ANY DAMAGES WHATSOEVER, WHETHER IN AN ACTION OF CONTRACT, NEGLIGENCE, OR OTHER TORTIOUS ACTION, ARISING OUT OF OR IN CONNECTION WITH THE USE OR PERFORMANCE OF SOFTWARE, DOCUMENTS, PROVISION OF OR FAILURE TO PROVIDE SERVICES, OR MATERIAL AVAILABLE FROM THIS SITE.

Table of Contents

| | |
|--|----------|
| Administrator Guide | i |
| Delinea Architecture Reference Diagrams | 1 |
| Account Lifecycle Manager Architecture | 1 |
| Account Lifecycle Manager Cloud Architecture | 1 |
| Component Definitions | 1 |
| Cloud Deployment Network Configuration | 2 |
| Account Lifecycle Manager On-Prem Architecture | 8 |
| Component Definition | 8 |
| On-Prem Deployment Network Configuration | 9 |
| Password Reset Server | 13 |
| Architecture Component Definitions | 13 |
| Single Site, Single Server | 14 |
| Single Site, Multi-Server | 15 |
| Multi-Site Design A - AlwaysOn AG | 16 |
| Multi-Site Design B - AlwaysOn AG | 18 |
| Delinea Platform Architecture | 20 |
| Delinea Platform: High-Level Overview | 20 |
| Privileged Remote Access | 21 |
| Delinea Connector | 22 |
| Privileged Access Service Architecture | 23 |
| Architecture Component Definitions | 23 |
| Hyper-Scalable Privileged Access Service: High Availability Multi-Availability Zone in Amazon AWS | 24 |
| Definitions | 24 |
| System Requirements | 25 |
| Diagram | 25 |
| Hyper-Scalable Privileged Access Service: High Availability Multi-Availability Zone in Microsoft Azure | 26 |
| Definitions | 26 |
| System Requirements | 27 |
| Diagram | 27 |
| Hyper-Scalable Privileged Access Service: High Availability Disaster Recovery | 28 |
| Definitions | 28 |
| System Requirements | 29 |
| Diagram | 29 |
| Hyper-Scalable Privileged Access Service: Three-Box Single Site | 31 |
| Definitions | 31 |
| System Requirements | 31 |
| Diagram | 31 |
| Hyper-Scalable Privileged Access Service: High Availability Single Site | 33 |
| Definitions | 33 |
| System Requirements | 33 |
| Diagram | 33 |
| Hyper-Scalable Privileged Access Service: Scalable Single Site | 35 |

Table of Contents

| | |
|---|----|
| Definitions | 35 |
| System Requirements | 35 |
| Diagram | 35 |
| Privilege Manager Architecture | 37 |
| Privilege Manager Cloud Architecture | 37 |
| Diagram | 37 |
| Web Application Firewall (WAF) | 37 |
| Requirements for the Architecture | 39 |
| Privilege Manager Single-Tenant Cloud Customer Example Architecture | 39 |
| Diagram | 39 |
| Privilege Manager Reference Architecture Diagrams | 40 |
| Component Definitions | 40 |
| Single Site with Minimum HA | 41 |
| Overview | 41 |
| Requirements | 41 |
| Virtual IP/Virtual Computer Object Requirements | 41 |
| Diagram | 41 |
| Single Site Minimum HA (Reverse Proxy/Azure Bus) | 42 |
| Overview | 42 |
| Requirements | 43 |
| Virtual IP or Computer Object Requirements | 43 |
| Diagram | 43 |
| Multi-Site Minimum HA/DR - Lower Cost, Manual Failover | 44 |
| Overview | 44 |
| Requirements | 45 |
| Virtual IP or Computer Object Requirements | 45 |
| Diagram | 45 |
| Multi-Site Average HA/DR - Average Cost, Manual Failover | 46 |
| Overview | 46 |
| Requirements | 47 |
| Virtual IP or Computer Object Requirements | 47 |
| Diagram | 47 |
| Best HA/DR - Highest Cost/Manual Failover | 48 |
| Overview | 48 |
| Requirements | 49 |
| Virtual IP/Virtual Computer Object Requirements | 49 |
| Diagram | 50 |
| Secret Server Architectures | 50 |
| Secret Server Cloud and On-Premises | 51 |
| Secret Server Cloud Only | 51 |
| Distributed Engine Example Architectures | 51 |
| General | 51 |
| Virtual IP or Computer | 52 |
| General | 54 |

Table of Contents

| | |
|--|-----|
| Virtual IP or Computer | 54 |
| General | 56 |
| Virtual IP or Computer | 56 |
| General | 58 |
| Virtual IP or Computer | 58 |
| General | 60 |
| Virtual IP or Computer | 60 |
| General | 62 |
| Virtual IP or Computer | 62 |
| General | 64 |
| Virtual IP or Computer | 64 |
| RabbitMQ Helper Example Architectures | 66 |
| Secret Server Example Architectures | 74 |
| General | 75 |
| Virtual IP or Computer | 75 |
| General | 77 |
| Virtual IP or Computer | 77 |
| General | 79 |
| Virtual IP or Computer | 79 |
| General | 81 |
| Virtual IP or Computer | 81 |
| General | 83 |
| Virtual IP or Computer | 83 |
| General | 85 |
| Virtual IP or Computer | 86 |
| General | 88 |
| Virtual IP or Computer | 88 |
| General | 91 |
| Virtual IP or Computer | 91 |
| Secret Server and DevOps Secrets Vault Example Architectures | 93 |
| Secret Server and Privilege Manager Example Architectures | 97 |
| Session Recording Example Architectures | 112 |
| Scenario A | 117 |
| Scenario B | 118 |
| Secret Server Resilient Secrets Architecture | 125 |
| Secret Server Cloud Customer Example Architectures | 127 |
| secretservercloud.com | 129 |
| secretservercloud.co.uk | 129 |
| secretservercloud.ca | 130 |
| secretservercloud.eu | 130 |
| secretservercloud.com.sg | 131 |
| secretservercloud.com.au | 131 |
| Secret Server Hybrid Multi-Tenant Cloud Architecture | 132 |
| secretservercloud.com | 134 |

Table of Contents

secretservercloud.co.uk 134
secretservercloud.ca 135
secretservercloud.eu 135
secretservercloud.com.sg 136
secretservercloud.com.au 136

Delinea Architecture Reference Diagrams

This documentation area provides a collection of all product Architecture Reference Diagrams currently available.

- [Account Lifecycle Manager Architecture](#)
- [Password Reset Server Architecture](#)
- [Platform Architecture](#)
- [Privileged Access Service Architecture](#)
- [Privilege Manager Architecture](#)
- [Secret Server Architecture](#)

Account Lifecycle Manager Architecture

This section contains example architecture diagrams for Account Lifecycle Manager and related technologies.

Account Lifecycle Manager Cloud Architecture

Component Definitions

Account Lifecycle Manager: Provides a cloud app design and intuitive UI as a front end to Active Directory allowing enterprise IT users to more easily and efficiently request, approve, privilege, manage, and retire service accounts by delegating the Active Directory intricacies to the cloud app.

Thycotic One Identity:Accounts that are used to provision User access to your ALM instance. Thycotic One accounts are used in some of our other products as well. This is not pictured in the architectural diagram as it is contained as part of the ALM sign on process.

ALM Remote Engine: A Windows Service that runs on your organization's hardware. It manages interactions between the ALM cloud service and your Active Directory installation. It also supports ALMs integration with your organizations Secret Server/DSV Instance. This includes support for integration with:

- Secret Server Cloud
- DevOps Secrets Vault

App Services: These are shared resources between multiple customers

Databases: These are customer independent for each ALM instance

Active Directory Server: This is the active Directory Server you intend to integrate the ALM Remote Engine with. Please note that relationships between a ALM Remote Engine to a domain within your environment is a 1:1 mapping. This implies that ALM Remote Engines can only manage one domain at a time.

Integrations: Currently there are additional integrations for Azure AD and ServiceNow that are pictured in the reference architecture

Cloud Deployment Network Configuration

1. Web Application Firewall (WAF): IP Address whitelisting is not necessary unless outbound firewall rules are in place. Public IP is based on geographical location.
All regions: 45.60.38.37, 45.60.40.37, 45.60.32.37, 45.60.34.37, 45.60.36.37, 45.60.104.37
2. Content Delivery Network (CDN): IP Address whitelisting is not necessary unless outbound firewall rules are in place. Public IP is based on geographical location.
All regions: <https://docs.microsoft.com/rest/api/cdn/edgenodes/list> (type=Standard_Verizon)
3. ALM Engine: IP Address whitelisting is not necessary unless outbound firewall rules are in place. If outbound firewall rules are in place, the ALM Engine should be allowed to the WAF IP address ranges listed above.
4. Active Directory Server: Must allow outbound communication from the ALM Remote Engine over TCP 636 (LDAPS) to your Active Directory Server.
5. Secret Server / DSV: Must allow outbound communication from the ALM Remote Engine over TCP 443 (HTTPS) to one of the following respective credential stores: Secret Server, Secret Server Cloud, or DSV. Please be mindful that if you are integrating with Secret Server Cloud, the ALM Remote Engine must also be able to communicate with the WAF IP address ranges (above) for Secret Server Cloud. DSV uses API Gateway regional endpoints with custom domain names in AWS. If you are integrating with DSV and have outbound restrictions from your ALM Remote Engine, it would be best to whitelist based on the tenant specific DSV custom domain name URL. If there is a hard requirement for outbound filtering based on IP address ranges, your rules will be dependent on this list:
<https://ip-ranges.amazonaws.com/ip-ranges.json> (you can filter out EC2 ranges).
6. Certificate CRLs: Whitelisting is not necessary unless outbound firewall rules are in place. If whitelisting is necessary, access to CRL distribution points is necessary

| | |
|-------------------------------|--|
| accountlifecyclecloud.com: | http://crl.godaddy.com/gdig2s1-1019.crl (web server) |
| | http://crl.godaddy.com/gdroot-g2.crl (web server) |
| accountlifecyclecloud.eu: | (unknown) |
| accountlifecyclecloud.com.au: | (unknown) |
7. Azure AD Integration: This is optional and extends its directory service support to include Azure AD. This allows ALM to manage accounts located in Azure AD. As this communication comes from the ALM Remote Engine, access would be outbound to the customer Azure AD environment over TCP 443 (HTTPS)

ALM Single Domain Design Example: Minimal Footprint or Cost

Requirements for Reference Architecture

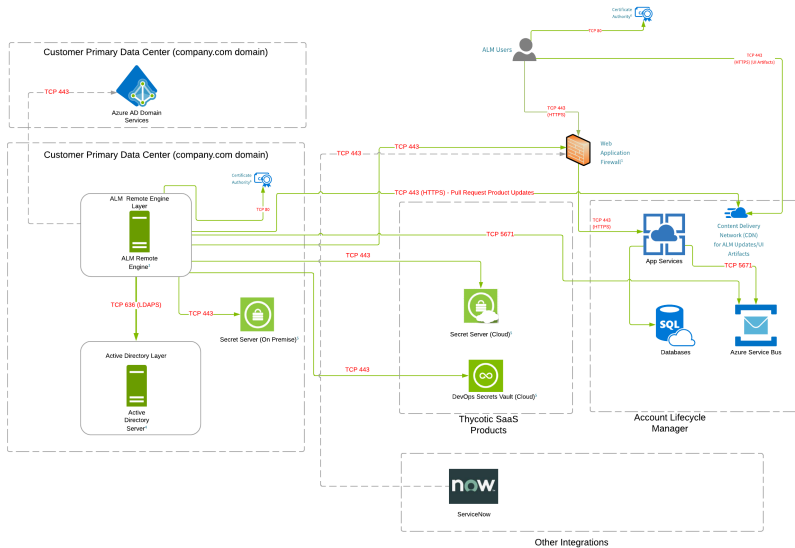
- Communication lines in green are required. Lines that are gray and dotted may or may not be required dependent on individual customer requirements.
- This design is fully supported by Delinea.
- All ALM Remote Engine servers to be running on Windows Server 2012 or later with .NET 4.7.1 or greater.

Delinea Architecture Reference Diagrams

- ALM Remote Engine
 - Minimum System Requirements: 2 Cores, 2 GB RAM
 - Recommended Requirements: 4 Cores, 4 GB RAM.



Arrows indicate direction of initial connection

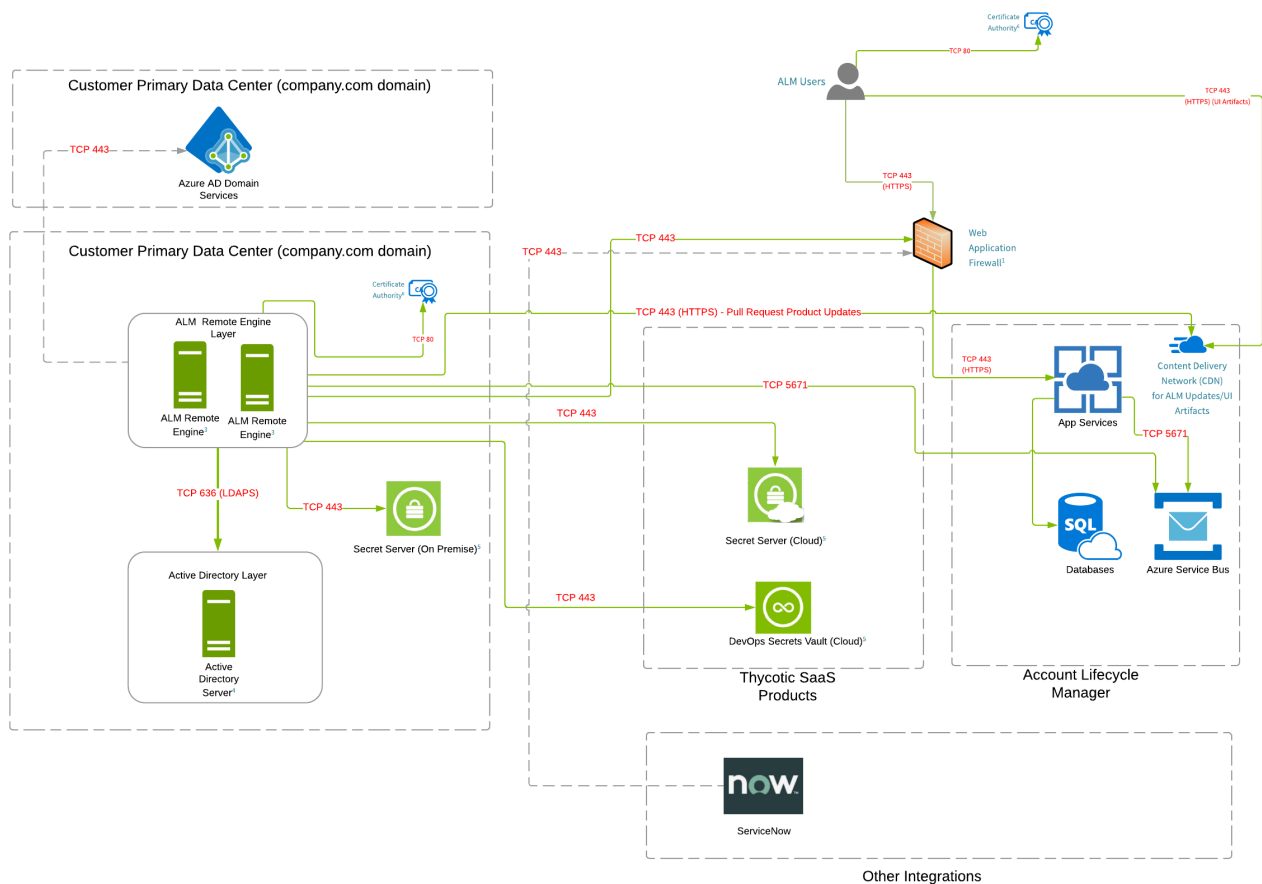


ALM Single Domain Design Example: HA

Requirements for Reference Architecture

- Communication lines in green are required. Lines that are gray and dotted may or may not be required dependent on individual customer requirements.
- This design is fully supported by Delinea.
- All ALM Remote Engine servers to be running on Windows Server 2012 or later with .NET 4.7.1 or greater.
- ALM Remote Engine
 - Minimum System Requirements: 2 Cores, 2 GB RAM

- Recommended Requirements: 4 Cores, 4 GB RAM.



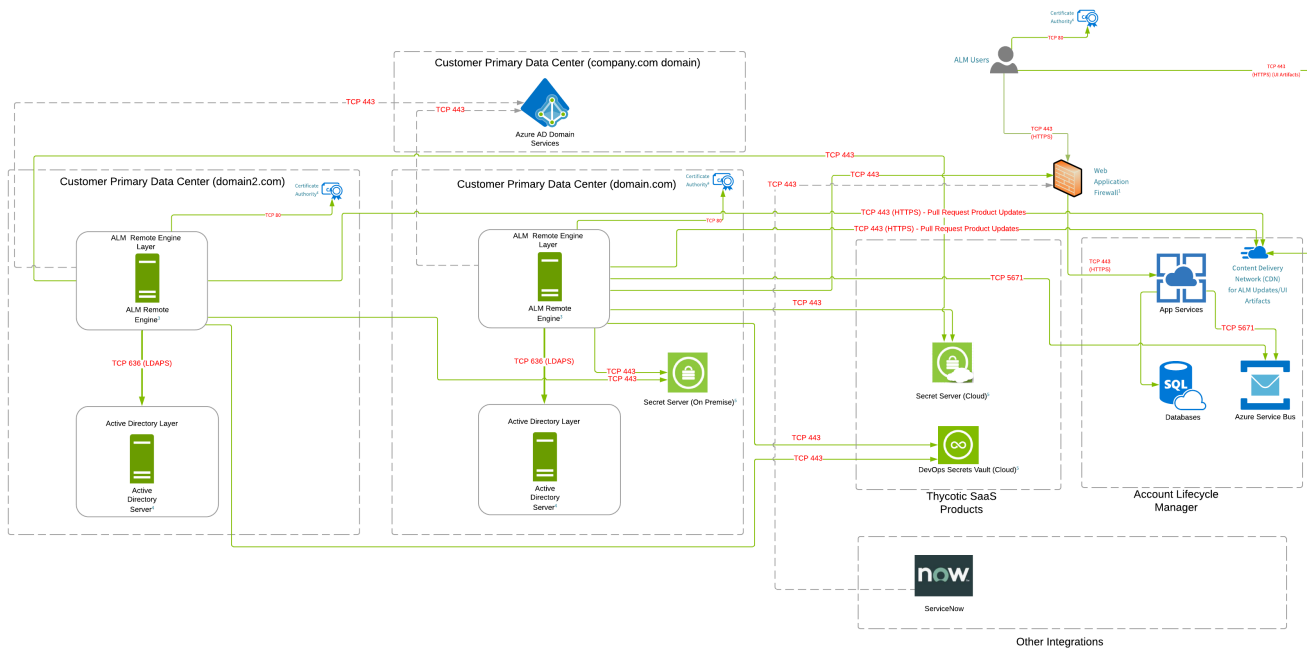
 Arrows indicate direction of initial connection

ALM Multiple Domains Design Example: Minimal Footprint or Cost

Requirements for Reference Architecture

- Communication lines in green are required. Lines that are gray and dotted may or may not be required dependent on individual customer requirements.
- This design is fully supported by Delinea
- All ALM Remote Engine servers to be running on Windows Server 2012 or later with .NET 4.7.1 or greater.
- All ALM Remote Engine servers require communication back to the Web (WAF) ranges over TCP 443. This has been pictured in the company.com domain but has not been pictured for other domains to make it easier to interpret.
- ALM Remote Engine
 - Minimum System Requirements: 2 Cores, 2 GB RAM
 - Recommended Requirements: 4 Cores, 4 GB RAM.

Delinea Architecture Reference Diagrams

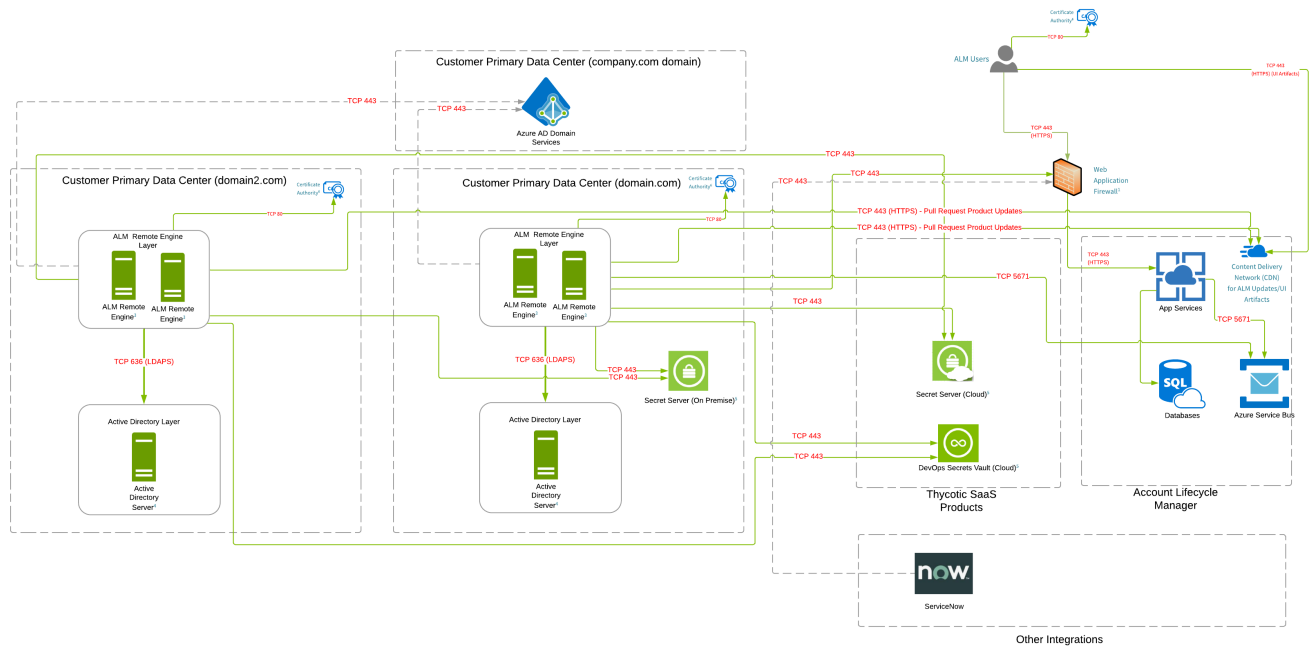


ALM Multiple Domains Design Example: HA

Requirements for Reference Architecture

- Communication lines in green are required. Lines that are gray and dotted may or may not be required dependent on individual customer requirements.
- This design is fully supported by Delinea.
- All ALM Remote Engine servers to be running on Windows Server 2012 or later with .NET 4.7.1 or greater.
- All ALM Remote Engine servers require communication back to the Web (WAF) ranges over TCP 443. This has been pictured in the company.com domain but has not been pictured for other domains to make it easier to interpret.
- ALM Remote Engine
 - Minimum System Requirements: 2 Cores, 2 GB RAM
 - Recommended Requirements: 4 Cores, 4 GB RAM.

Delinea Architecture Reference Diagrams

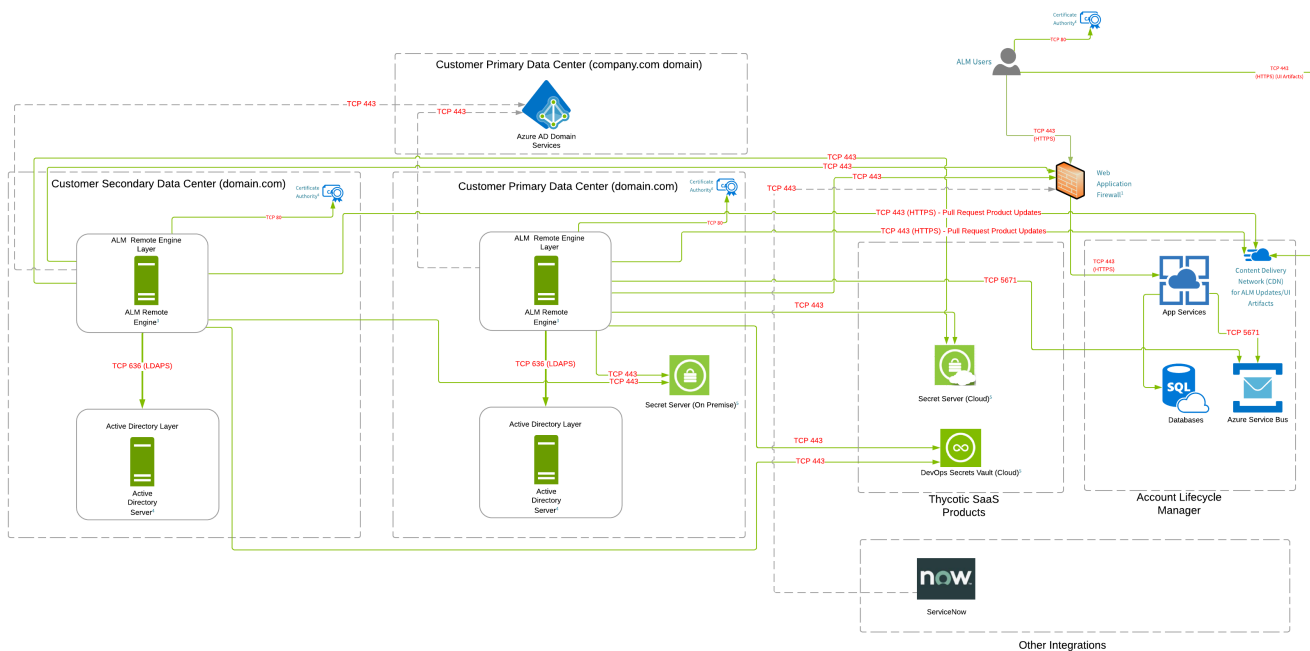


ALM Single Domain with Multiple Data Centers Design Example: HA/DR

Requirements for Reference Architecture

- Communication lines in green are required. Lines that are gray and dotted may or may not be required dependent on individual customer requirements.
- This design is fully supported by Delinea.
- All ALM Remote Engine servers to be running on Windows Server 2012 or later with .NET 4.7.1 or greater.
- ALM Remote Engine
 - Minimum System Requirements: 2 Cores, 2 GB RAM
 - Recommended Requirements: 4 Cores, 4 GB RAM.

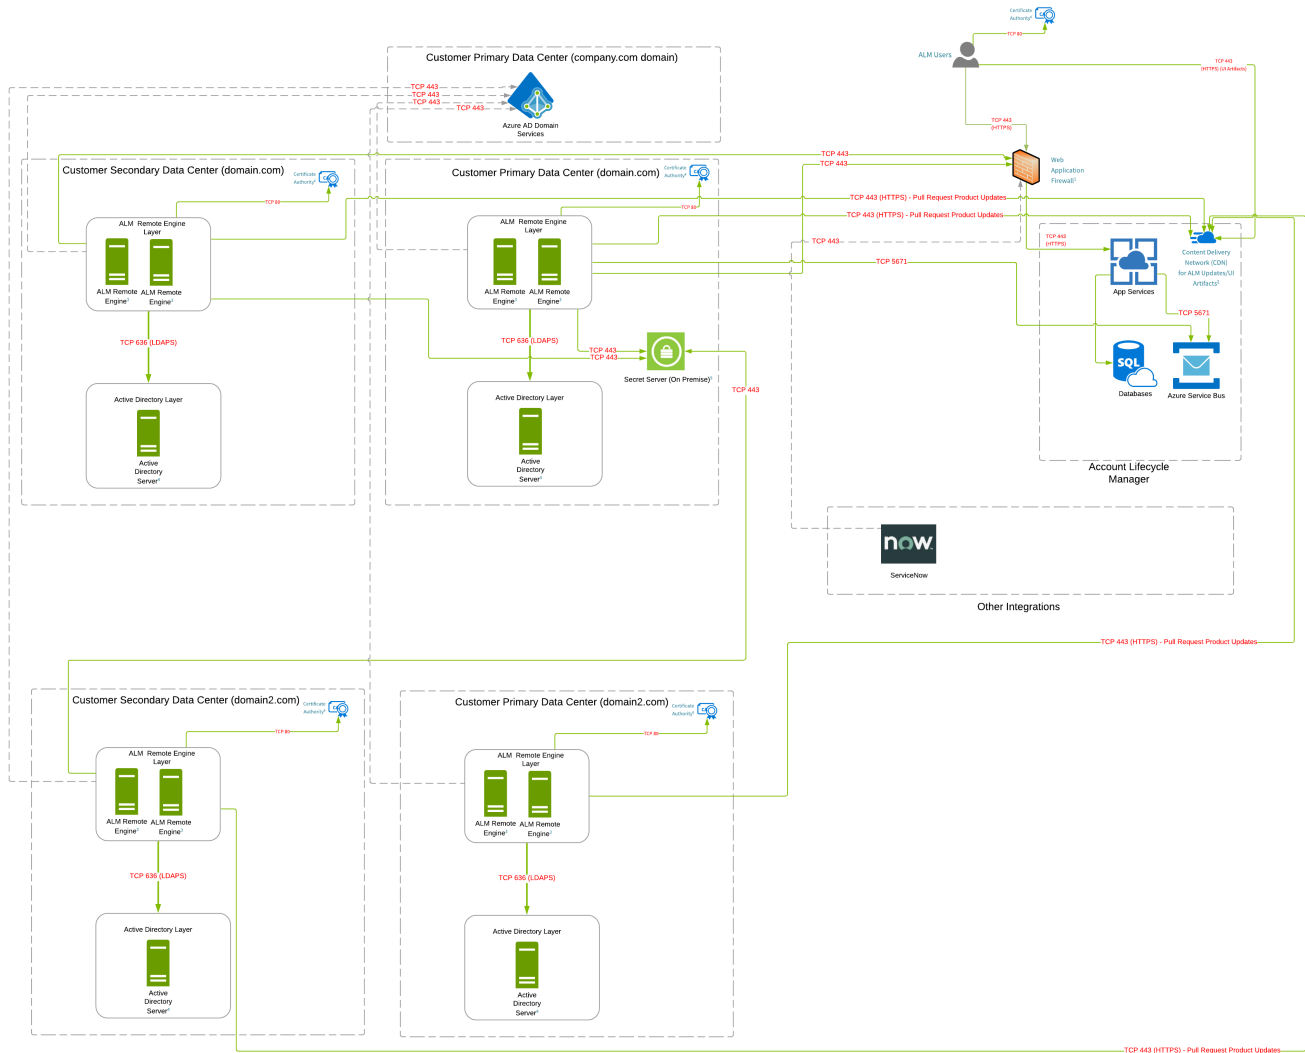
Delinea Architecture Reference Diagrams



ALM Multiple Domains and Data Centers Design Example: HA/DR

Requirements for Reference Architecture

- Communication lines in green are required. Lines that are gray and dotted may or may not be required dependent on individual customer requirements.
- This design is fully supported by Delinea.
- All ALM Remote Engine servers to be running on Windows Server 2012 or later with .NET 4.7.1 or greater.
- All ALM Remote Engine servers require communication back to the Web (WAF) ranges over TCP 443. This has been pictured in the company.com domain but has not been pictured for other domains to make it easier to interpret.
- ALM Remote Engine
 - Minimum System Requirements: 2 Cores, 2 GB RAM
 - Recommended Requirements: 4 Cores, 4 GB RAM.



Account Lifecycle Manager On-Prem Architecture

Component Definition

Account Lifecycle Manager - Provide a cloud app design and intuitive UI as a front end to Active Directory allowing enterprise IT users to more easily and efficiently request, approve, privilege, manage, and retire service accounts by delegating the Active Directory intricacies to the cloud app.

Thycotic One Identity - These are the accounts that are used to provision User access to your ALM instance. Thycotic One accounts are used in some of our other products as well. This is not pictured in the architectural diagram as it is contained as part of the ALM sign on process.

ALM Remote Engine - The ALM engine is a Windows Service that runs on your organization's hardware. It manages interactions between the ALM cloud service and your Active Directory installation. It also supports ALMs integration with your organizations Secret Server/DSV Instance. This includes support for integration with:

Delinea Architecture Reference Diagrams

- Secret Server On Prem Installation
- Secret Server Cloud
- DevOps Secrets Vault

App Services - These are shared resources between multiple customers

Databases - These are customer independent for each ALM instance

Active Directory Server - This is the active Directory Server you intend to integrate the ALM Remote Engine with. Please note that relationships between a ALM Remote Engine to a domain within your environment is a 1:1 mapping. This implies that ALM Remote Engines can only manage one domain at a time.

Integrations - Currently there are additional integrations for Azure AD and ServiceNow that are pictured in the reference architecture

On-Prem Deployment Network Configuration

1. Docker Production Images: IP Address allowlisting is not necessary unless outbound firewall rules are in place. If outbound firewall restrictions are in place, please allow:
 - *.docker.io
 - *.docker.com
2. Content Delivery Network (CDN): IP Address allowlisting is not necessary unless outbound firewall rules are in place. Public IP is based on geographical location. All regions:
[https://docs.microsoft.com/rest/api/cdn/edgenodes/list \(type=Standard_Verizon\)](https://docs.microsoft.com/rest/api/cdn/edgenodes/list?type=Standard_Verizon) The installation script is pulled from - <https://alc-cdn01.azureedge.net>
3. Active Directory Server: Must allow outbound communication from the ALM Remote Engine over TCP 636 (LDAPS) to your Active Directory Server.
4. Secret Server / DSV: Must allow outbound communication from the ALM Remote Engine over TCP 443 (HTTPS) to one of the following respective credential stores: Secret Server, Secret Server Cloud, or DSV. Please be mindful that if you are integrating with Secret Server Cloud, the ALM Remote Engine must also be able to communicate with the WAF IP address ranges (above) for Secret Server Cloud. DSV uses API Gateway regional endpoints with custom domain names in AWS. If you are integrating DSV and have outbound restrictions from your ALM Remote Engine, it would be best to whitelist based on the tenant specific DSV custom domain name URL. If there is a hard requirement for outbound filtering based on IP address ranges, your rules will be dependent on this list -<https://ip-ranges.amazonaws.com/ip-ranges.json> (you can filter out EC2 ranges).
5. **Optional Integrations**
 - Azure AD Integration (Optional): Extends directory service support to include Azure AD. This allows ALM to manage accounts located in Azure AD. As this communication comes from the ALM Remote Engine, access would be outbound to the customer Azure AD environment over TCP 443 (HTTPS)
 - LetsEncrypt (Optional): If leveraging LetsEncrypt during installation, please review this article for any potential outbound firewall requirements <https://letsencrypt.org/docs/integration-guide/#firewall-configuration>

Delinea Architecture Reference Diagrams

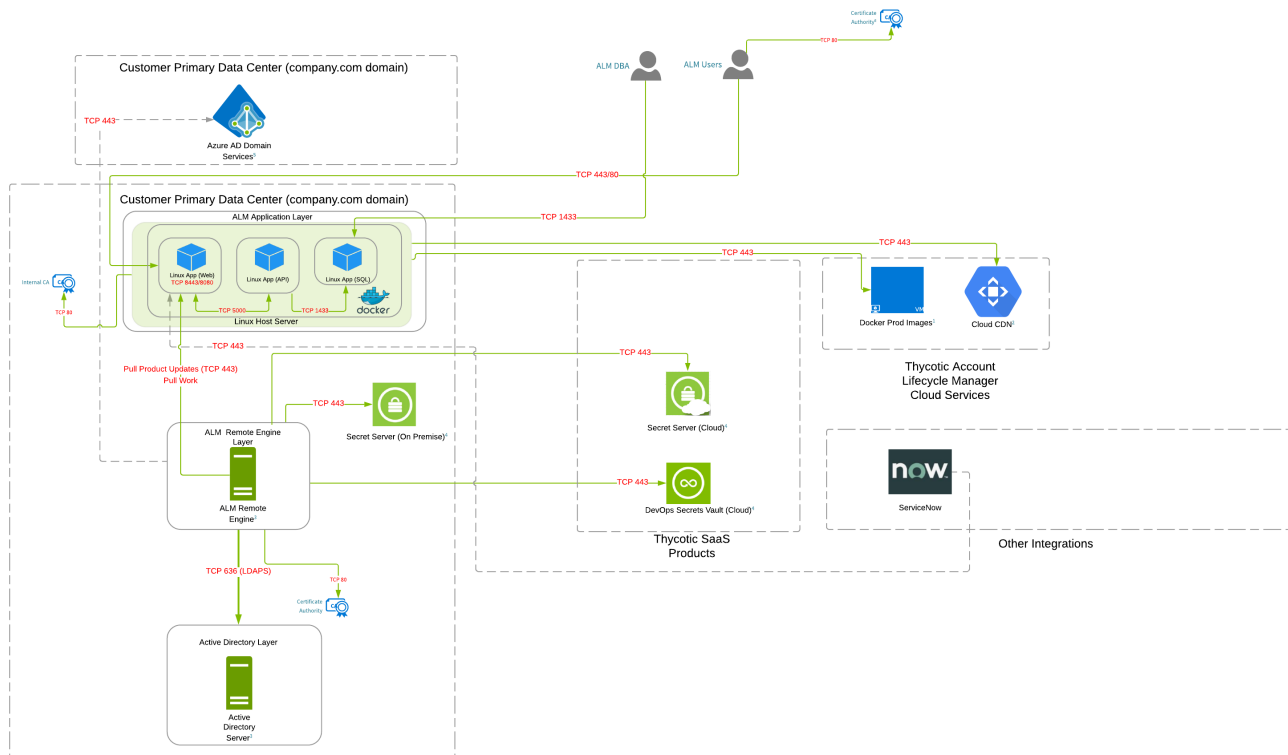
- Additional Allowlisting (Optional): When installing pre-requisite software (i.e Docker Compose), it may be desirable to temporarily allow communication from the ALM Application Layer to the resources below. Alternatively this can be downloaded and transferred directly to the Linux Host Server.

- *.github.com
- *.amazonaws.com (<https://github-production-release-asset-2e65be.s3.amazonaws.com>)

ALM Single Domain Design (On Prem) Example - Minimal Footprint/Cost

Requirements for Reference Architecture

- Communication lines in green are required. Lines that are gray and dotted may or may not be required dependent on individual customer requirements.
- This design is fully supported by Delinea.
- All ALM Remote Engine servers to be running on Windows Server 2012 or later with .NET 4.7.1 or greater.
- It is possible for customers to use an existing SQL server instance that is not part of a docker container. This means the database can reside on a Windows MSSQL instance or a Linux instance running MSSQL.
- ALM Remote Engine
 - Minimum System Requirements - 2 Cores 2GB RAM
 - Recommended Requirements - 4 Cores 4 GB RAM.



ALM Multi Domain Design (On Prem) Example - HA/DR (Engines + SQL)

Requirements for Reference Architecture

- Communication lines in green are required. Lines that are gray and dotted may or may not be required dependent on individual customer requirements.
- This design is fully supported by Delinea.
- All ALM Remote Engine servers to be running on Windows Server 2012 or later with .NET 4.7.1 or greater.
- All ALM Remote Engine servers require communication back to the Web container over TCP 443. This has been pictured in the company.com domain but has not been pictured for other domains to make it easier to interpret.
- It is possible for customers to use an existing SQL server instance that is not part of a docker container. This means the database can reside on a Windows MSSQL instance or a Linux instance running MSSQL.
- ALM Remote Engine
 - Minimum System Requirements - 2 Cores 2GB RAM
 - Recommended Requirements - 4 Cores 4 GB RAM.

Delinea Architecture Reference Diagrams

Database Server is a primary component of the solution. SQL Server hosts the PRS database. We are compatible with SQL Server 2005 or newer running on Windows Server 2008–Windows Server 2019. The PRS database can be put on a stand-alone server, a FCI, or preferably using an AlwaysOn AG for clustered environments. The database can be added to an existing production SQL cluster or instance, but it is important to retain proper sizing of the environment. Windows authentication only is advised.

Single Site, Single Server

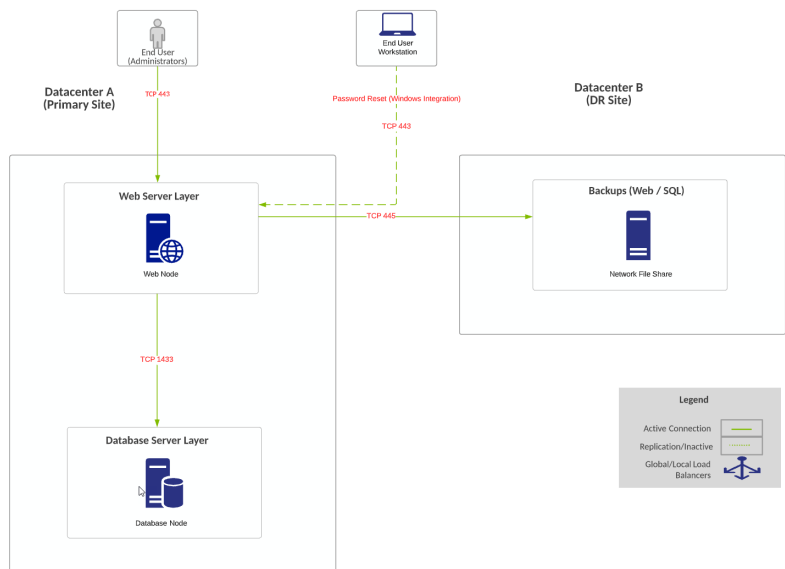
This is an example of a standard Password Reset Server installation. Password Reset Server Standard and Microsoft SQL Server Standard are installed on a single Windows Server (Production). A Disaster Recovery Plan for this configuration would consist of the **Manual** Web Application Backup and **Manual** Application Database Backup procedures. User-managed strategies can also be employed, such as a **Manual** File Backup and **Manual** Database Backup. If the Windows Server is virtualized, strategies such as making scheduled Snapshots or having a Hot/Cold Site would provide additional layers of redundancy. The recovery time for reverting to Password Reset Server-generated backup files is roughly 30-60 minutes plus any time needed to prepare the server for Password Reset Server installation in the user's environment.

Diagram



The reference number for this diagram is #11-A-1.

Figure: Standard Installation on a Single Windows Server with Minimal Footprint



Definitions

Minimum footprint design

Requirements

- Web Node - Minimum 2 core 2 Ghz or higher per core, 4 GB RAM
- Web Node - Recommended 4 core 2 Ghz or higher per core, 8 GB RAM
- IIS 7, 8
- Database - 1000 users=300 MB DB
- SQL Standard Edition (SQL Server 2005 - 2016)
- All port requirements are listed on diagram

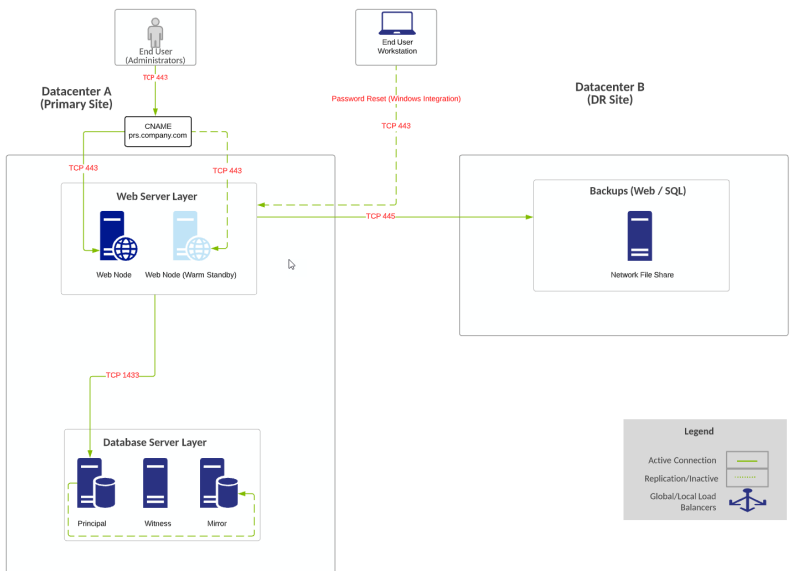
Single Site, Multi-Server

This is an example of a common Password Reset Server installation. Password Reset Server and Microsoft SQL Server Standard are installed on a number of Windows Servers, with the number depending on the requirements of the organization. The diagram shows Microsoft SQL Mirroring enabled with a Witness Server, but a Witness Server is not required for this scenario. A Disaster Recovery Plan for the above configuration would consist of failover for Microsoft SQL Server issues. If the failover members were to themselves fail, then **Automated** Web Application Backups and **Automated** Application Database Backups can be used to restore functionality to the Inactive server. If these Servers are virtualized, strategies such as making scheduled Snapshots or having a Hot/Cold Site would provide additional layers of redundancy. The recovery time for reverting to Password Reset Server-generated backup files is roughly 30-60 minutes plus any time needed to prepare the server for Password Reset Server installation in the user's environment.

Diagram

 The reference number for this diagram is #11-A-2.

Figure: Hot Standby Mode



Definitions

- Hot standby node (requires manual intervention to bring online)
- SQL Mirroring in primary location configured

Requirements

- Web Node - Minimum 2 core 2 Ghz or higher per core, 4 GB RAM
- Web Node - Recommended 4 core 2 Ghz or higher per core 8 GB RAM
- IIS 7, 8
- Database - 1000 users=300MB DB
- SQL Standard Edition (SQL Server 2005 - 2016)
- All port requirements are listed on diagram

Multi-Site Design A - AlwaysOn AG

This is an example of a Password Reset Server installation that leverages a hot standby node in another data center. In the event of a disaster, manual intervention is required to bring the web node server in that location online. Load Balancers are used to help minimize DNS-related changes during fail over events. This design variation leverages **Basic** Availability Groups as part of SQL Standard licensing, for a low cost, new design option to provide HA/DR for the PRS back-end database. Synchronous replication configured between data centers is advisable only when data centers are in close physical proximity to one another or latency is less than 30ms.

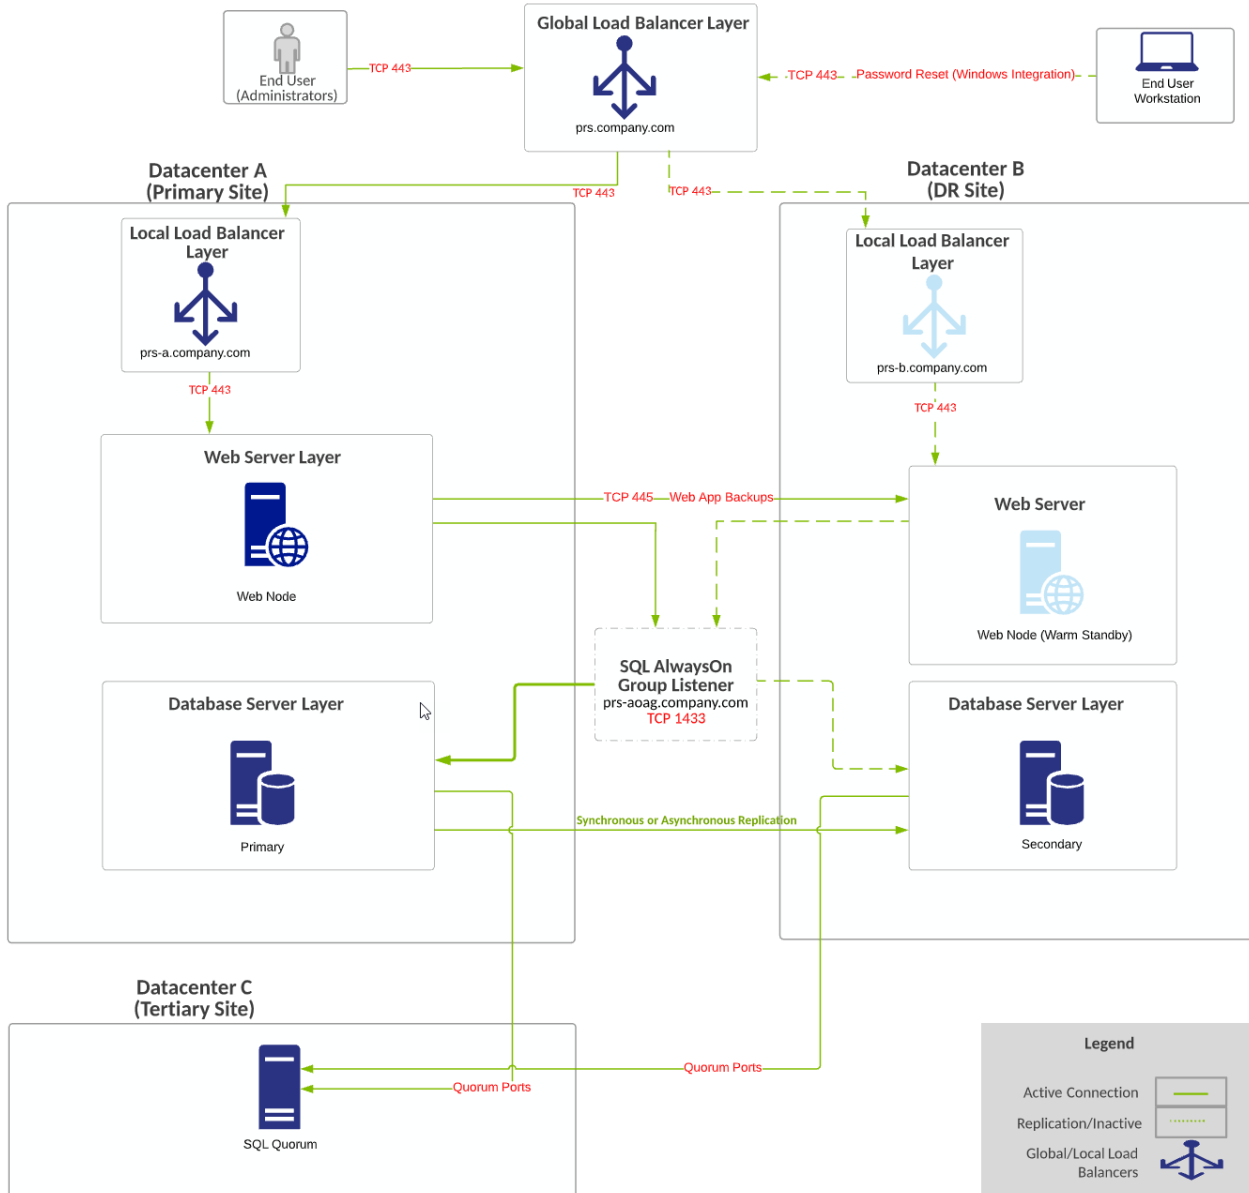
Diagram



The reference number for this diagram is #11-B-1.

Figure: Hot Standby Node in Another Data Center

Delinea Architecture Reference Diagrams



Definitions

- Hot standby node in another data center (requires manual intervention to bring online)
- Traffic to DR site in GLB configuration can be included but should be explicitly disabled in the pool until Web Server in DR is brought online
- Basic Availability Groups used as part of the solution

Requirements

- Web Node - Minimum 2 core 2 Ghz or higher per core, 4 GB RAM
- Web Node - Recommended 4 core 2 Ghz or higher per core 8 GB RAM
- IIS7, 8
- Database - 1000 users = 300 MB DB
- SQL Standard Edition (SQL Server 2016)
- All port requirements are listed on diagram
- Synchronous or asynchronous replication is dependent on latency. Recommend asynchronous replication between data centers with more than 30ms latency

Virtual IP or Computer Object Requirements

- prs.company.com:443 (1 virtual IPs - Global Load Balancer)
- prs-a.company.com:443. prs-b.company.com:443 (2 virtual IPs - Local Load Balancer)
- prs-aoag.company.com:1433 (created as part of SQL AlwaysOn Configuration)
 - prs-aoag.company.com computer object/Virtual IP
 - 2 virtual Ip addresses may be required as part of this configuration
- Windows Failover Cluster Object (created as part of Windows Failover Clustering Configuration)
 - computer object/Virtual IP
 - 2 additional virtual IP addresses may be required as part of Windows Failover Cluster for single site design for the network configuration of the Failover Cluster representing both networks at each respective site

Multi-Site Design B - AlwaysOn AG

This is an example of a Password Reset Server installation that leverages multiple hot standby nodes. One is set up locally and another is set up in another data center. In the event of a disaster, manual intervention is required to bring the web node server in that location online. Load Balancers are used to help minimize DNS-related changes during fail over events. This design variation leverages **Enterprise** Availability Groups as part of SQL **Enterprise** licensing, for a low cost, new design option to provide HA/DR for the PRS back-end database. A local node is available for patching scenarios to fail the database over to another node within the same data center. Synchronous replication configured between data centers is advisable only when data centers are in close physical proximity to one another or latency is less than 30ms.

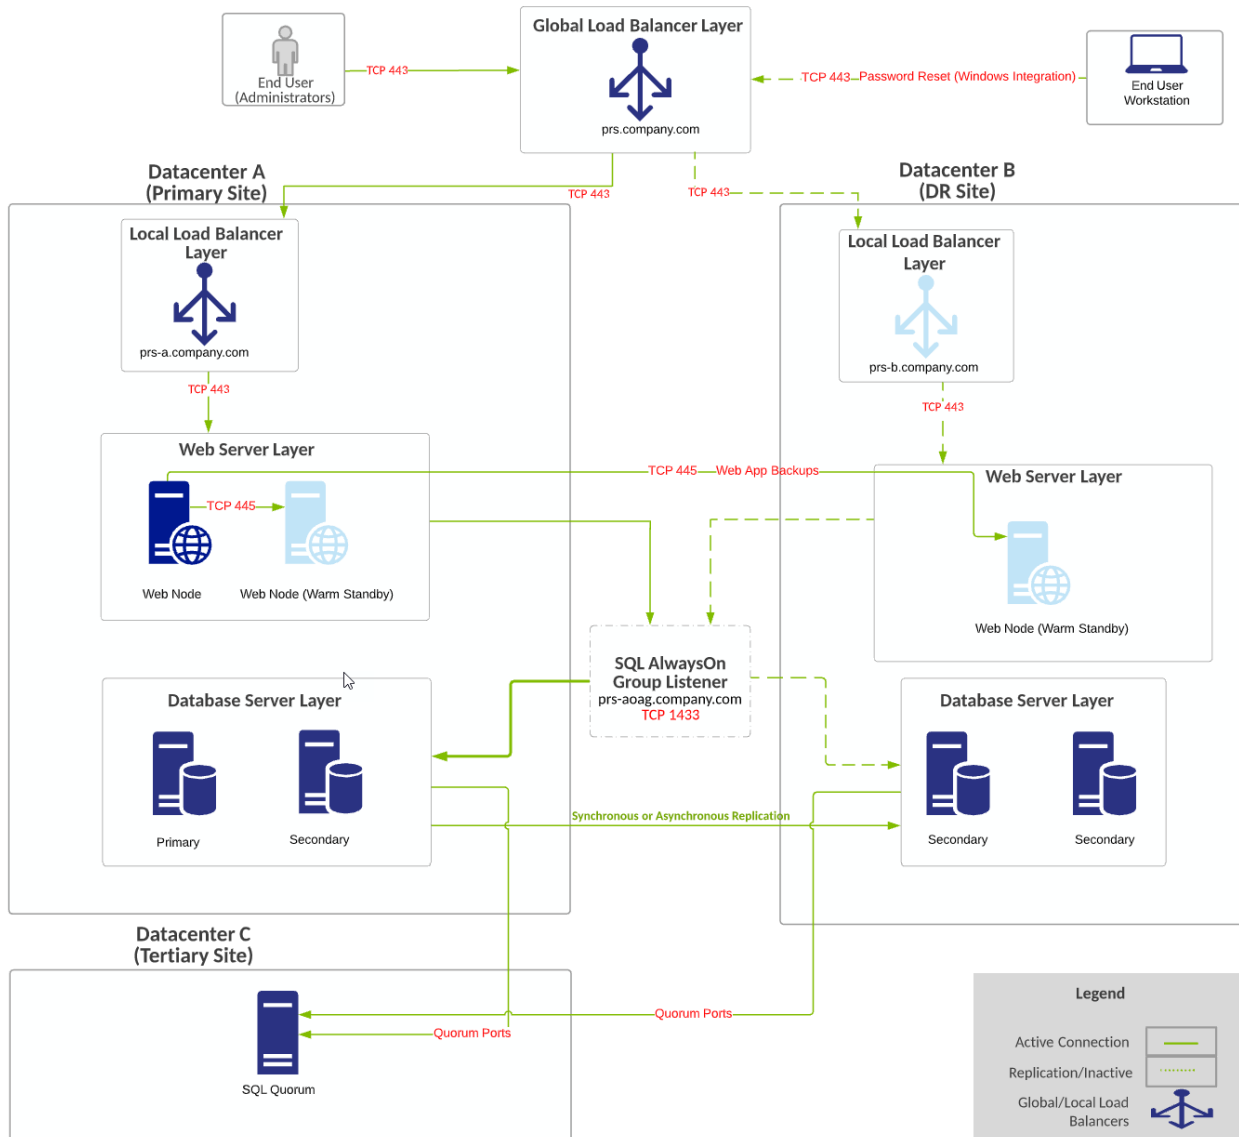
Diagram



The reference number for this diagram is #11-B-2.

Figure: Hot Standby Node Locally and in Another Data Center

Delinea Architecture Reference Diagrams



Definitions

- Hot standby node locally and in another data center (requires manual intervention to bring online)
- Traffic to DR site in GLB configuration can be included but should be explicitly disabled in the pool until Web Server in DR is brought online
- AlwaysOn Enterprise Availability Groups used as part of the solution

Requirements

- Web Node - Minimum 2 core 2Ghz or higher per core, 4GB RAM
- Web Node - Recommended 4 core 2Ghz or higher per core 8GB RAM
- IIS 7, 8

Delinea Architecture Reference Diagrams

- Database - 1000 users = 300MB DB
- SQL Enterprise Edition (SQL Server 2012 - 2016)
- All port requirements listed on diagram
- Synchronous or asynchronous replication is dependent on latency. Recommend asynchronous replication between data centers with more than 30 ms latency.


Virtual IP or Computer Object Requirements

- prs.company.com:443 (1 virtual IPs- Global Load Balancer)
- prs-a.company.com:443. prs-b.company.com:443 (2 virtual IPs- Local Load Balancer)
- prs-aoag.company.com:1433 (created as part of SQL AlwaysOn Configuration)
 - prs-aoag.company.comcomputer object/Virtual IP
 - 2 virtual IP addresses may be required as part of this configuration
- Windows Failover Cluster Object (created as part of Windows Failover Clustering Configuration)
 - computer object/Virtual IP
 - 2 additional virtual IP addresses may be required as part of Windows Failover Cluster for single site design for the network configuration of the Failover Cluster representing both networks at each respective site.

Delinea Platform Architecture

These architectural diagrams are provided to help you gain a high-level understanding of the underlying infrastructure and technology stack that supports the Delinea Platform. Additionally, you can leverage this material if you are interested in allowing access to the Delinea Platform and its related services in your firewall.


We are continuously improving and optimizing our architecture to ensure that our service is scalable, secure, and efficient.

 **Note:** The suggested list of ports in this document shows all of the default port numbers. These default ports may differ based on your environment and your own unique requirements. In all cases, the ports and addresses listed below should be excluded from packet inspection to allow for normal service operation.

Please see [Delinea Platform Architecture and Topology](#) for additional information.

Delinea Platform: High-Level Overview

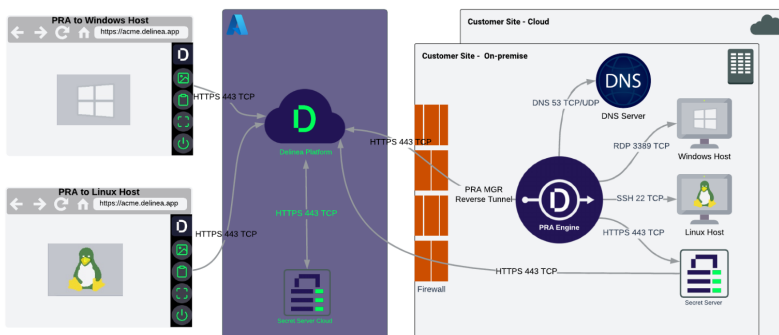
- The diagram below highlights the overall architecture of the Delinea Platform.
 - Shared services are foundational services that provide infrastructure and other common resources that are designed to be consumed by various applications such as authentication, notification, and audit.
 - Application services are built on top of the platform shared services, and are designed to provide functionality that is unique to the application such as vaulting and remote access.

 **Note:** The Delinea Platform is evolving with every new release. The overview diagram below may be forward-looking from that perspective.

Delinea Architecture Reference Diagrams

- No internet-facing ingress ports are required for the PRA Engine
- Outbound access on port 443 TCP from PRA Engine to the Delinea Platform via Impreva ingress
- Internal access on port 53 TCP/UDP from PRA Engine to DNS server for name resolution of target machines
- Internal access on port 3389 TCP from PRA Engine to Windows-based target machines for RDP access
- Internal access on port 22 TCP from PRA Engine to Linux-based target machines for SSH access
- Internal access on port 443 TCP from PRA Engine to Secret Server (on-premise) to enable integration with Delinea Platform and leverage secret access. Only required if Secret Server (on-premise) is in use.


Delinea Remote Access Service



Delinea Connector

The Delinea Connector enables secure communication between the Delinea Platform and AD directories. Typically, the Delinea Connector is installed on-premises and requires access to an Active Directory Domain Controller.

- No internet-facing ingress ports are required for the Connector
- Outbound access on port 443 TCP from the Connector to the Delinea Platform via Impreva WAF

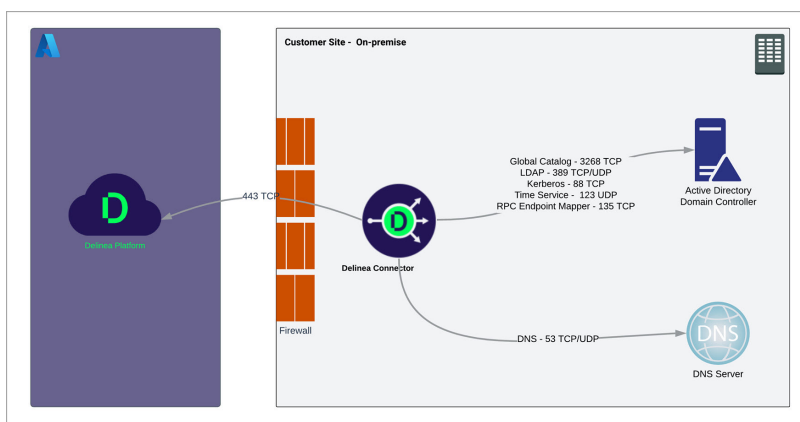
 **Note:** Requests from the Delinea Platform to the Delinea Connector are made via the TCP Relay hosts. Such requests for instance include querying for AD user's details. All data is encrypted.

| Region | TCP Relay Hosts IP Address Range |
|----------------|---|
| Canada | 20.104.14.80 - 20.104.14.87 |
| Australia | 20.211.60.240 - 20.211.60.247 |
| United Kingdom | 20.49.210.72 - 20.49.210.79 |
| United States | 20.242.252.136 - 20.242.252.143; 52.148.145.72 - 52.148.145.79; 20.85.110.128 - 20.85.110.135 |
| Southeast Asia | 20.195.89.80 - 20.195.89.87 |
| Europe | 20.8.3.112 - 20.8.3.119 |

Delinea Architecture Reference Diagrams

- Internal access on port 3268 TCP from Delinea Connector to AD Domain Controller for Global Catalog access
- Internal access on port 123 UDP from Delinea Connector to AD Domain Controller for time synchronization
- Internal access on port 389 TCP/UDP from Delinea Connector to AD Domain Controller for handling normal authentication queries
- Internal access on port 88 TCP from Delinea Connector to AD Domain Controller used for Kerberos authentication
- Internal access on port 135 TCP from Delinea Connector to AD Domain Controller for remote procedure call (RPC) endpoint mapping
- Internal access on port 53 TCP/UDP from Delinea Connector to DNS server for name resolution (this might be the DC itself depending on your environment)

Delinea Connector



Privileged Access Service Architecture

This section contains example architecture diagrams for on-premise Privileged Access Service deployments.

Architecture Component Definitions

Background node: contains Hyper-scalable PAS software and manages background jobs such as regularly rotating passwords, re-syncing with the Domain Controller and running reports. Background nodes communicate with Web nodes, Cache (Redis), and the Database (PostgreSQL) servers. The Jobs dashboard (*/jobs*) provides a view of the Background node workload. You can add more Background nodes to scale up your architecture if you notice delays or jobs are queued for extended periods of time.

The Relay node allows the Delinea Privileged Access Service to communicate with the Connector. Connectors are used to enable Active Directory integration, RDP access, and other integrations with the infrastructure. All TCP Relay nodes receive requests to forward data from Web nodes and/or Background nodes. If a request is Connector-bound (instead of, logging), it is forwarded along the Connector-initiated pipe.

Delinea Architecture Reference Diagrams

The Logging node centralizes the logs onto a single system for easier diagnostics, as well as allowing the logs to be watched on the Management node. The command `Centrify-Pas-WatchLogs.ps1` will not work without a logging node.

Cache (Redis) server: caches repeat operations to improve database performance. The cache (Redis) never originates requests; it only receives and answers requests.

Connector: Integrates your Active Directory/LDAP service with Privileged Access Service. The connector allows you to, among other things, specify groups whose members can register and manage devices. It also monitors Active Directory/LDAP for group policy changes, which it sends to Privileged Access Service to update registered devices.

Database (PostgreSQL) server: external database that is only used for Hyper-scalable PAS. The database (PostgreSQL) never originates requests; it only receives and answers requests.

End User: These are the users connecting to your Delinea website. These users may be performing administrative tasks (admins), or just using the solution.

Load balancer: load balances traffic to multiple servers (for Web node and connector traffic). The load balancer must have a static IP address, with an appropriate entry connecting the name (URI) to the address in the DNS.

Management node: scripts are executed from the Management node to manage the cluster. The Management node is not part of the cluster itself, however It does need to be able to reach Web, Background and TCP Relay nodes and have full database access. While the management node needs full database access, it doesn't directly communicate with any other nodes beyond the initial installation.

TCP Relay node (Relay and Logging): Relay and Logging TCP Relay nodes contain Hyper-scalable PAS software and bridge between other technologies such as Active Directory, RDP hosts, log aggregation, and the Hyper-scalable PAS deployment. Although a separate Logging node is not mandatory, Delinea suggests you deploy a separate Logging node.

Web node: contains Hyper-scalable PAS software and manages incoming web requests and provides REST endpoints (provides web API functionality). Web nodes communicate with Background nodes, TCP Relay nodes, Cache (Redis), and the Database (PostgreSQL) servers. All user-access to Hyper-scalable PAS is through the Web nodes, which are reached at the host address through the load balancer. The Web nodes do not typically perform long-running or scheduled tasks; their job is to respond quickly to user requests. Only active Web nodes, those with the current active Deployment ID, respond to requests, and therefore only Web nodes receive traffic from the load balancer. You can add more Web nodes to scale up your architecture.

Hyper-Scalable Privileged Access Service: High Availability Multi-Availability Zone in Amazon AWS



If you are a current customer with support hours for Delinea Professional Services, you can discuss any of these diagrams in detail with one of our Professional Services support architects.

Definitions

- Recommended for large production environments
- Major Application components split primarily across many servers with each major layer offering HA capability
- Connectors installed in one or more locations

Delinea Architecture Reference Diagrams

- Flows depict connectivity for web based connection between user and destination systems and using native SSH or RDP clients to connect through connector to destination systems (proxied use cases)
- Servers in other Availability Zone are located within the same region. **Not all communication lines have been pictured for this active/active design between all AZs to make the diagram comprehensible. Communication within the Web, Background, Relay, and Database Layers may be required between Availability Zones using the ports described. As an example, the Web Server's in AZ1 will require communication to the Background and Relay Nodes in AZ2 based on the ports pictured in AZ1 location and vice versa. For any customers with tight firewalls between AZs, please inspect closely the ports that may be required across AZs by referencing - <https://docs.centrixy.com/Content/managed/NetworkTopo.htm>.**
- Designs that are spread across regions will require similar failover process as depicted in [HSPAS High availability single site](#).
- Customers may leverage PaaS offerings for PostgreSQL and Redis for AWS. This reference architecture pictures multi-AZ support for these layers. **Be aware that during a failover to a replica, communication/port requirements to the replica will be required.**

System Requirements

- 8 Core, 8 GB RAM for the Application, Web, Logging, and Relay layers
- 4 Core, 16 GB RAM for the Connector
- 8 Core, 32 GB RAM for Postgres
- 8 Core, 32 GB RAM for Redis Cache
- Management Server does not need to be a net-new system and can have minimal specs

Diagram

Figure: High availability multi-availability zone in Amazon AWS (large deployment)

Delinea Architecture Reference Diagrams

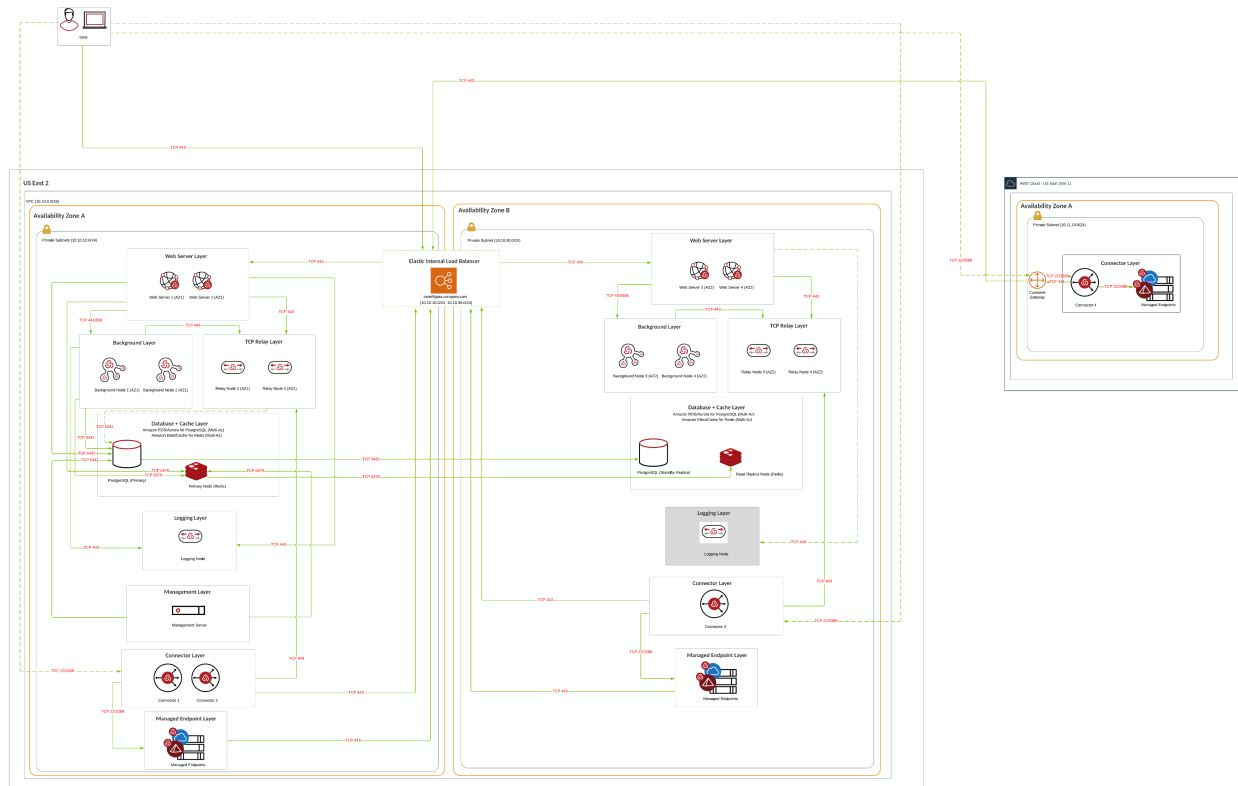



Figure: Diagram legend



Hyper-Scalable Privileged Access Service: High Availability Multi-Availability Zone in Microsoft Azure



If you are a current customer with support hours for Delinea Professional Services, you can discuss any of these diagrams in detail with one of our Professional Services support architects.

Definitions

- Recommended for large production environments
- Major Application components split primarily across many servers with each major layer offering HA capability
- Connectors installed in one or more locations

Delinea Architecture Reference Diagrams

- Flows depict connectivity for web based connection between user and destination systems and using native SSH or RDP clients to connect through connector to destination systems (proxied use cases)
- Servers in other Availability Zone are located within the same region. **Not all communication lines have been pictured for this active/active design between all AZs to make the diagram comprehensible. Communication within the Web, Background, Relay, and Database Layers may be required between Availability Zones using the ports described. As an example, the Web Server's in AZ1 will require communication to the Background and Relay Nodes in AZ2 based on the ports pictured in AZ1 location and vice versa. For any customers with tight firewalls between AZs, please inspect closely the ports that may be required across AZs by referencing - <https://docs.delinea.com/online-help/cloud-suite/before-deploy/firewall.htm?Highlight=firewall%20rules>.**
- Designs that are spread across regions will require similar failover process as depicted in [HSPAS High availability single site](#)
- Customers may leverage PaaS offerings for PostgreSQL and Redis for Azure. This reference architecture pictures multi-AZ support for these layers. **Be aware that during a failover to a replica, communication/port requirements to the replica will be required.**

System Requirements

- 8 Core, 8 GB RAM for the Application, Web, Logging, and Relay layers
- 4 Core, 16 GB RAM for the Connector
- 8 Core, 32 GB RAM for Postgres
- 8 Core, 32 GB RAM for Redis Cache
- Management Server does not need to be a net-new system and can have minimal specs

Diagram

Figure: High availability multi-availability zone in Microsoft Azure (large deployment)

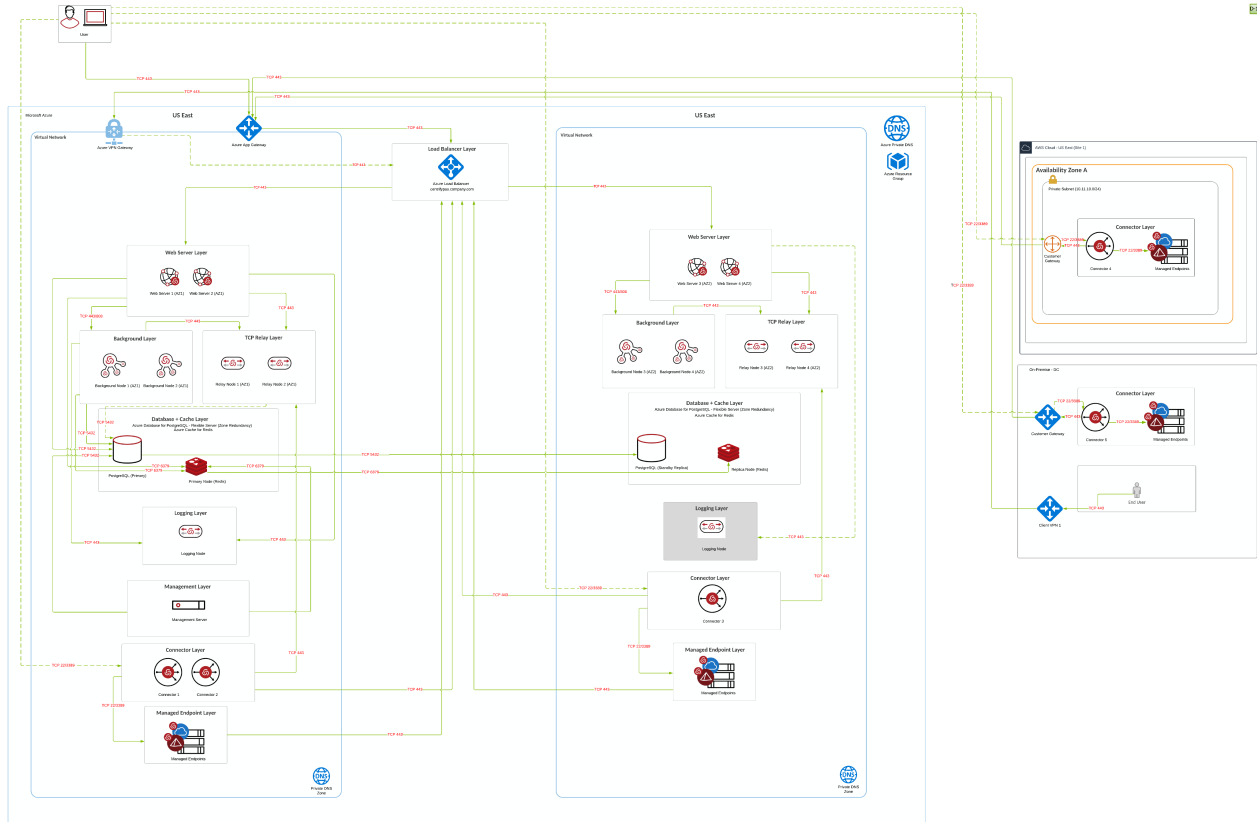




Figure: Diagram legend



Hyper-Scalable Privileged Access Service: High Availability Disaster Recovery

 If you are a current customer with support hours for Delinea Professional Services, you can discuss any of these diagrams in detail with one of our Professional Services support architects.

Definitions

- Recommended for large production environments
- Major Application components split primarily across many servers with each major layer offering HA capability
- Connectors installed in one or more locations

Delinea Architecture Reference Diagrams

- Flows depict connectivity for web based connection between user and destination systems and using native SSH or RDP clients to connect through connector to destination systems (proxied use cases)
- Servers in Disaster Recovery Location can be built but should be considered "cold" standby servers for disaster recovery event. Layers highlighted in grey should be brought online manually. Load Balancer configuration should only be marked as active once all components have been brought online, then traffic can be redirected to DR location.
- **There are constraints between HSPAS and Redis as it relates to multi-site designs. A manual failover may require a partial rebuild by modifying the installation and redeployment of nodes. Servers can be prepared and "ready" but this may impact your RTO and RPO for failover.**

System Requirements

- 8 Core, 8 GB RAM for the Application, Web, Logging, and Relay layers
- 4 Core, 16 GB RAM for the Connector
- 8 Core, 32 GB RAM for Postgres
- 8 Core, 32 GB RAM for Redis Cache
- Management Server does not need to be a net-new system and can have minimal specs

Diagram

Figure: Hyper-scalable Privileged Access Service: High availability disaster recovery

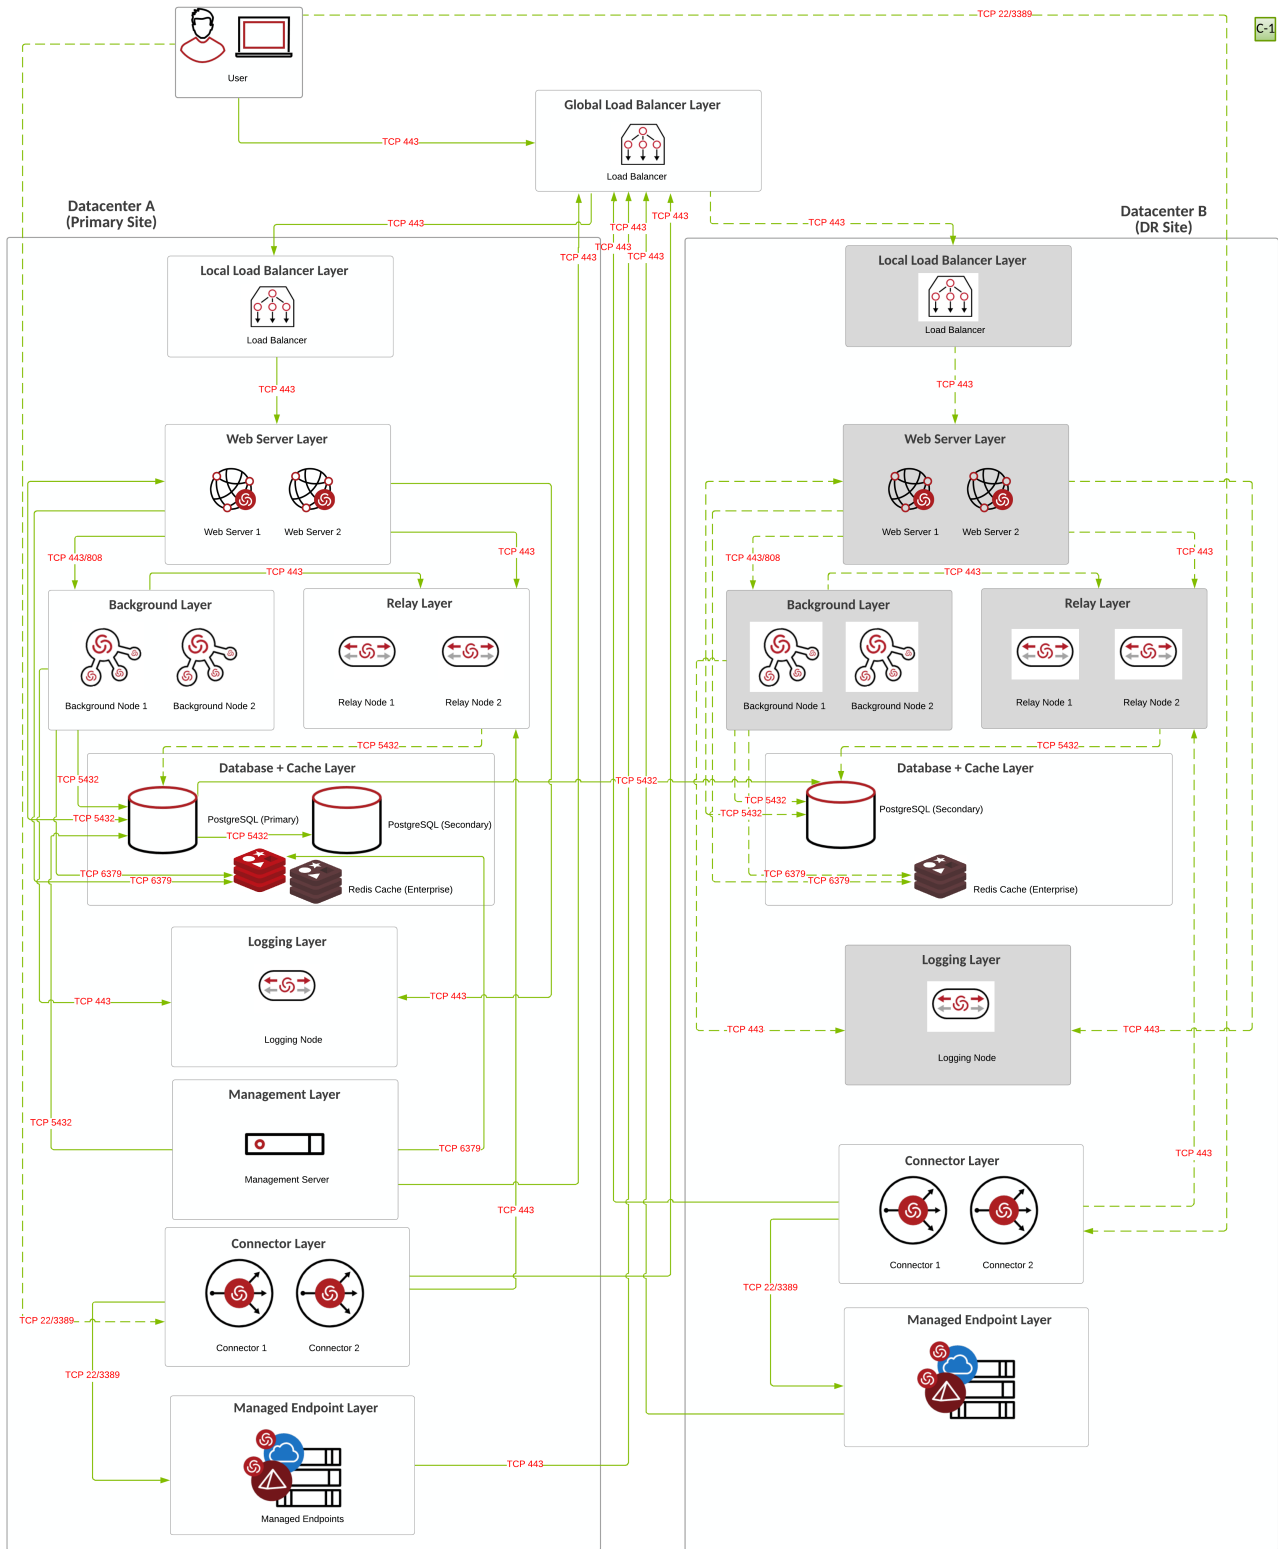



Figure: Diagram legend



Hyper-Scalable Privileged Access Service: Three-Box Single Site



If you are a current customer with support hours for Delinea Professional Services, you can discuss any of these diagrams in detail with one of our Professional Services support architects.

Definitions

- Recommended for small footprint production environments
- Major Application components split primarily across two servers
- Connectors installed in one or more locations. For minimal footprint, the Connector can also be installed on the management + relay node, keeping the solution entirely to 3 systems
- Flows depict connectivity for web based connection between user and destination systems and using native SSH or RDP clients to connect through connector to destination systems (proxied use cases)

System Requirements

- 8 Core, 16 GB RAM for the shared Application + Database layer
- 4 Core, 16 GB RAM for the Connector
- 8 Core, 32 GB RAM for Postgres
- 8 Core, 32 GB RAM for Redis Cache
- Management Server is shared with the Relay node
- An additional node (node not pictured) for the Logging component is optional and can be recommended or not based on feedback from Professional Services

Diagram

Figure: Hyper-scalable Privileged Access Service: 3 box single site

B-1

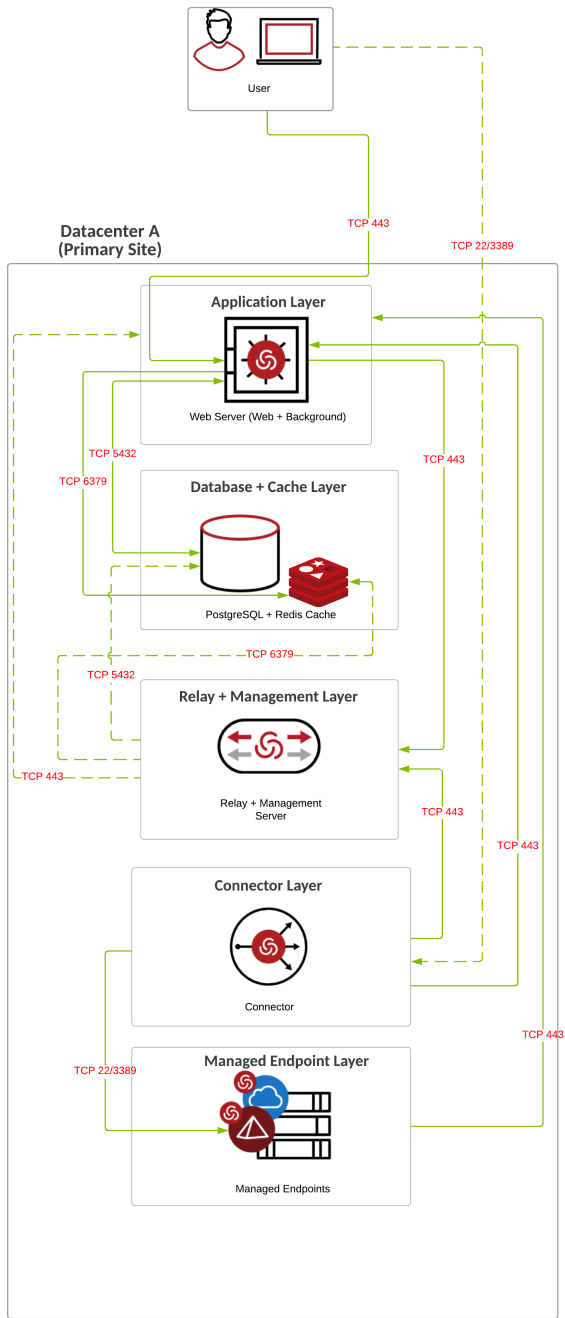



Figure: Diagram legend



Hyper-Scalable Privileged Access Service: High Availability Single Site



If you are a current customer with support hours for Delinea Professional Services, you can discuss any of these diagrams in detail with one of our Professional Services support architects.

Definitions

- Recommended for large production environments
- Application components split across independent servers, additional nodes to provide HA added for most layers
- Connectors installed in one or more locations
- Flows depict connectivity for web based connection between user and destination systems and using native SSH or RDP clients to connect through connector to destination systems (proxied use cases)

System Requirements

- 8 Core, 8 GB RAM for the Application, Web, Logging, and Relay layers
- 4 Core, 16 GB RAM for the Connector
- 8 Core, 32 GB RAM for Postgres
- 8 Core, 32 GB RAM for Redis Cache
- Management Server does not need to be a net-new system and can have minimal specs

Diagram

Figure: Hyper-scalable Privileged Access Service: High availability single site

Delinea Architecture Reference Diagrams

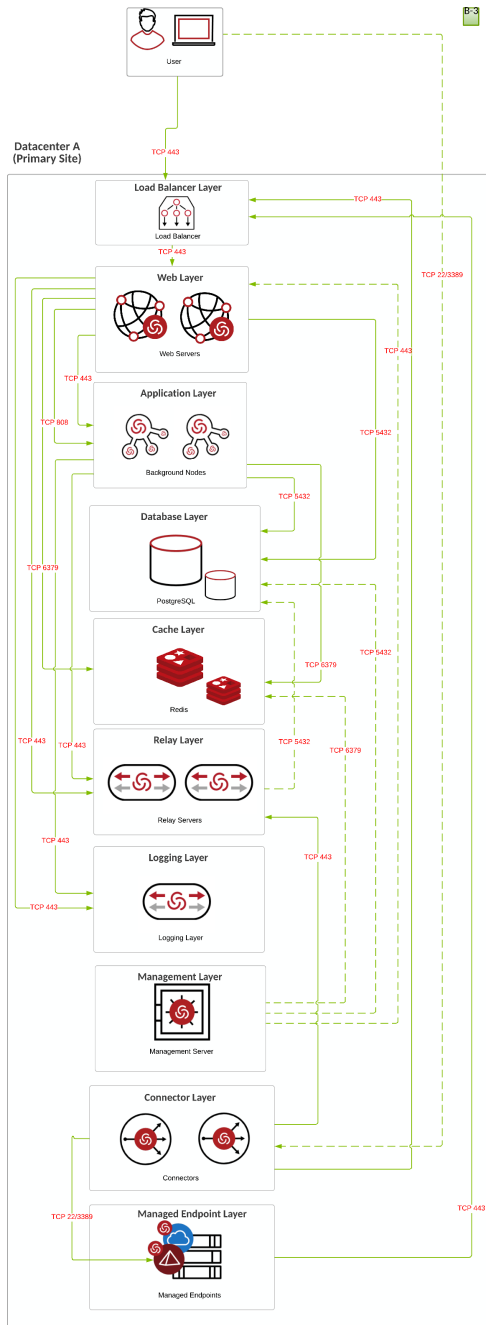



Figure: Diagram legend



Hyper-Scalable Privileged Access Service: Scalable Single Site



If you are a current customer with support hours for Delinea Professional Services, you can discuss any of these diagrams in detail with one of our Professional Services support architects.

Definitions

- Recommended for medium sized production environments
- Application components split across independent servers
- Connectors installed in one or more locations.
- Flows depict connectivity for web based connection between user and destination systems and using native SSH or RDP clients to connect through connector to destination systems (proxied use cases)

System Requirements

- 8 Core, 8 GB RAM for the Application, Web, Logging, and Relay layers
- 4 Core, 16 GB RAM for the Connector
- 8 Core, 32 GB RAM for Postgres
- 8 Core, 32 GB RAM for Redis Cache
- Management Server does not need to be a net-new system and can have minimal specs

Diagram

Figure: Scalable single site

Delinea Architecture Reference Diagrams

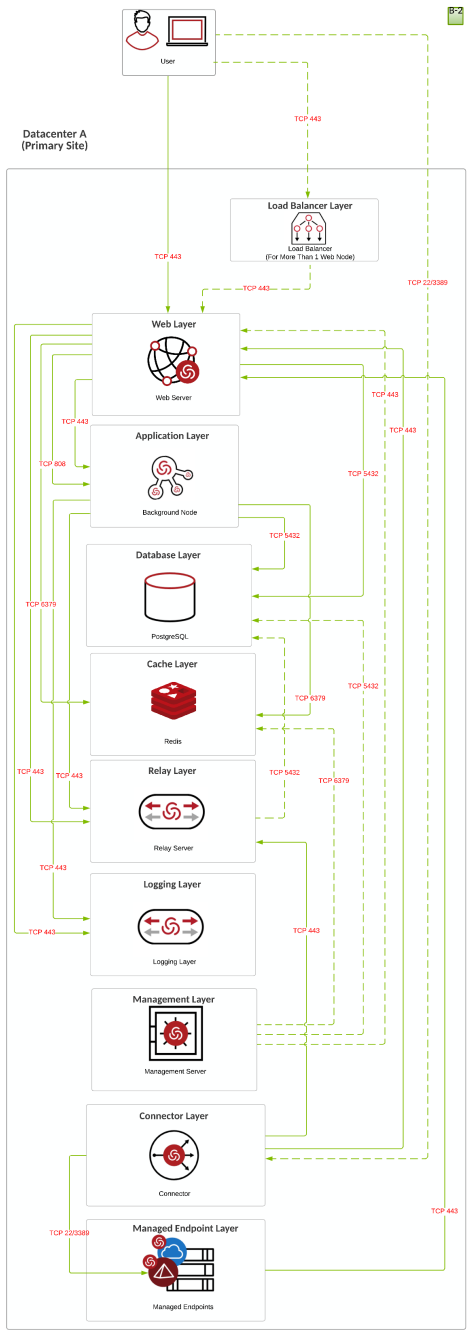



Figure: Diagram legend



Privilege Manager Architecture

This section contains example architecture diagrams for Privilege Manager and related technologies.

Privilege Manager Cloud Architecture

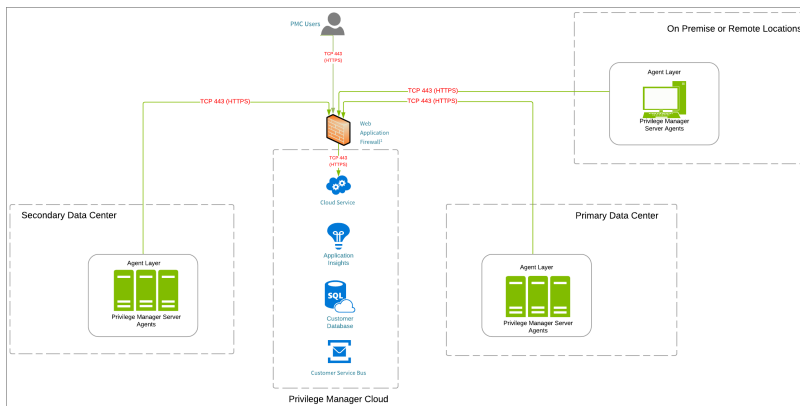


If you are a current customer with support hours for Delinea Professional Services, you can discuss any of these diagrams in detail with one of our Professional Services support architects.

The information provided is a Privilege Manager Cloud agent connectivity example.

Diagram

Figure: Privilege Manager Cloud Architecture



Arrows indicate the direction of initial connection.

Web Application Firewall (WAF)

IP Address allowlisting is not necessary unless outbound firewall rules are in place. Generally, the public IP the hostname resolves to is based on geographical location of the request source.

US

- 13.92.238.189
- 20.62.140.151
- 20.72.149.142
- 20.81.14.178
- 20.81.21.130
- 20.81.65.234
- 20.81.68.147
- 20.81.123.91

Delinea Architecture Reference Diagrams

- 20.84.34.132
- 20.85.201.237
- 20.88.176.22
- 20.119.38.215
- 20.163.166.200
- 20.172.242.207
- 20.228.249.169
- 20.231.58.68
- 20.232.195.171
- 23.101.142.103
- 40.76.159.46
- 40.88.250.211
- 52.136.127.49
- 52.146.66.102
- 52.146.66.244
- 52.149.238.230
- 52.149.205.98
- 52.154.69.186
- 52.186.102.0
- 52.191.103.234
- 52.191.235.73
- 52.224.76.213
- 52.224.254.56
- 52.255.201.247
- 54.154.73.179
- 172.174.100.176

EU

- 20.56.240.207
- 20.61.137.98
- 20.76.59.64
- 20.86.109.175
- 20.93.154.191

Delinea Architecture Reference Diagrams

- 52.157.86.0
- 104.45.72.251

AU

- 20.37.5.19
- 20.39.64.207

Canada

- 20.63.51.11
- 20.116.80.217

SEA

- 20.43.149.152

Requirements for the Architecture

The agents on endpoints may be installed anywhere around the world. The diagram can be adjusted to include specific locations relevant to our customers.

Agents may be installed on Windows, macOS, or Unix/Linux endpoints.

- Windows 7 or newer
- Windows Server 2012+
- MacOS 10.11 (El Capitan or newer)
- Unix/Linux: CentOS 7+, RHEL 7+, Oracle Linux 7+, Ubuntu 18.04+

Also refer to:

- [Agent Software](#)
- [Agent Installation](#)

Privilege Manager Single-Tenant Cloud Customer Example Architecture



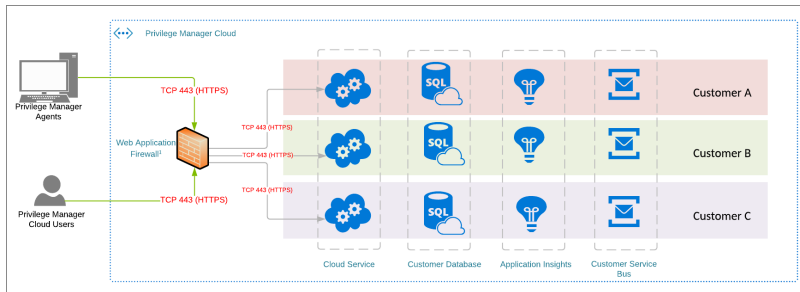
If you are a current customer with support hours for Thycotic Professional Services, you can discuss any of these diagrams in detail with one of our Professional Services support architects.

The information provided is a Privilege Manager Cloud reference example.

Diagram

Figure: Privilege Manager Cloud Customer Architecture

Delinea Architecture Reference Diagrams



Arrows indicate the direction of initial connection.

Privilege Manager Reference Architecture Diagrams

Component Definitions

App User - These are the users connecting to your Privilege Manager websites. These users will be limited to the users that perform administrative tasks (admins), to use the solution in a helpdesk role, or to perform approvals or audits.

Privilege Manager Agents - These is used for application control and local user/group management.

Load Balancers - Load balancers are often involved in the solution to help distribute web traffic to more than one web server. Local and Global load balancers, if available, may be used in the solution to further lower potential application downtime during upgrades, patching, and single site failures.

Web Server - This is a primary component of the solution. Our web servers use IIS 7 and newer and will only work on Windows Server 2008 R2 - Windows Server 2016. For multiple web server (clustered) solutions, the web application itself can be made cluster aware and does not require being built as part of an IIS farm. Each web server acts as its own stand alone web server.

Database Server - This is a primary component of the solution. Microsoft SQL Server hosts the Privilege Manager databases. We are compatible only with SQL Server 2012 or newer running on Windows Server 2012 R2 - Windows Server 2016. The Delinea databases can be put on a stand alone server, a FCI, or preferably using an AlwaysOn AG for clustered environments. The databases can be added to an existing production SQL cluster or instance, but proper sizing of the environment should be done. Windows authentication only is advised.

Reverse Proxy / Azure Service Bus - A properly configured Reverse Proxy will act as a buffer between Privilege Manager agents and the Privilege Manager server(s) to limit server exposure. Use nginx, F5, or Windows Application Request Routing 3.0 and URL Rewrite in IIS on a DMZ Server, to prevent a direct connection between Agent endpoints and your Privilege Manager web server(s). Alternatively, Azure Bus can be used, to prevent Agent endpoints connecting directly to your Privilege Manager web server(s).

Secret Server - Optionally, Secret Server can be installed with Privilege Manager to use an authentication source and a storage vault for Privilege Manager credentials. Using Secret Server as the authentication source for Privilege Manager allows MFA options for login. Also, application role assignment can be assigned in Secret Server. If using Secret Server features (beyond authentication and vault storage for Privilege Manager), Secret Server should be on separate servers - for this, see [Secret Server + Privilege Manager Architectures](#).

Note: Every component of Privilege Manager can be made highly available to ensure a redundant architecture and to scale for future growth.

Privilege Manager can be setup for various types of authentication methods:

- Azure Active Directory Authentication
- NTLM for local webserver authentication
- ThycoticOne for Cloud Instances

Single Site with Minimum HA

Overview

- Minimum Cost HA Configuration.
- Single Site design, no native DR capacity. DR can be provided by means of VM replication if subnets are spanning locations, otherwise re-ip + DNS changes may be necessary.
- Due to SQL Basic Availability groups with Standard Edition, you will need to have multiple instances of SQL and a dedicated AlwaysOn availability group configuration.
- Some customers may choose to use a separate web reverse proxy or azure service bus configuration for Privilege Manager agent TCP 443 communication.
- Optionally, Privilege Manager can be integrated with Secret Server (installed on the same web servers or, preferably, on dedicated servers) for authentication and as a storage vault for Privilege Manager credentials. If using Secret Server features (beyond authentication and vault storage for Privilege Manager), Secret Server should be on separate servers.

Requirements

- SQL Standard Edition - Basic Availability Group Configuration.
- Local load balancers can be utilized for all web server nodes.
- Configuring a file share witness for SQL quorum voting is required for SQL to stay online during single node unplanned failures.
- SQL Quorum Ports - <http://dsfnet.blogspot.com/2013/04/windows-server-clustering-sql-server.html>

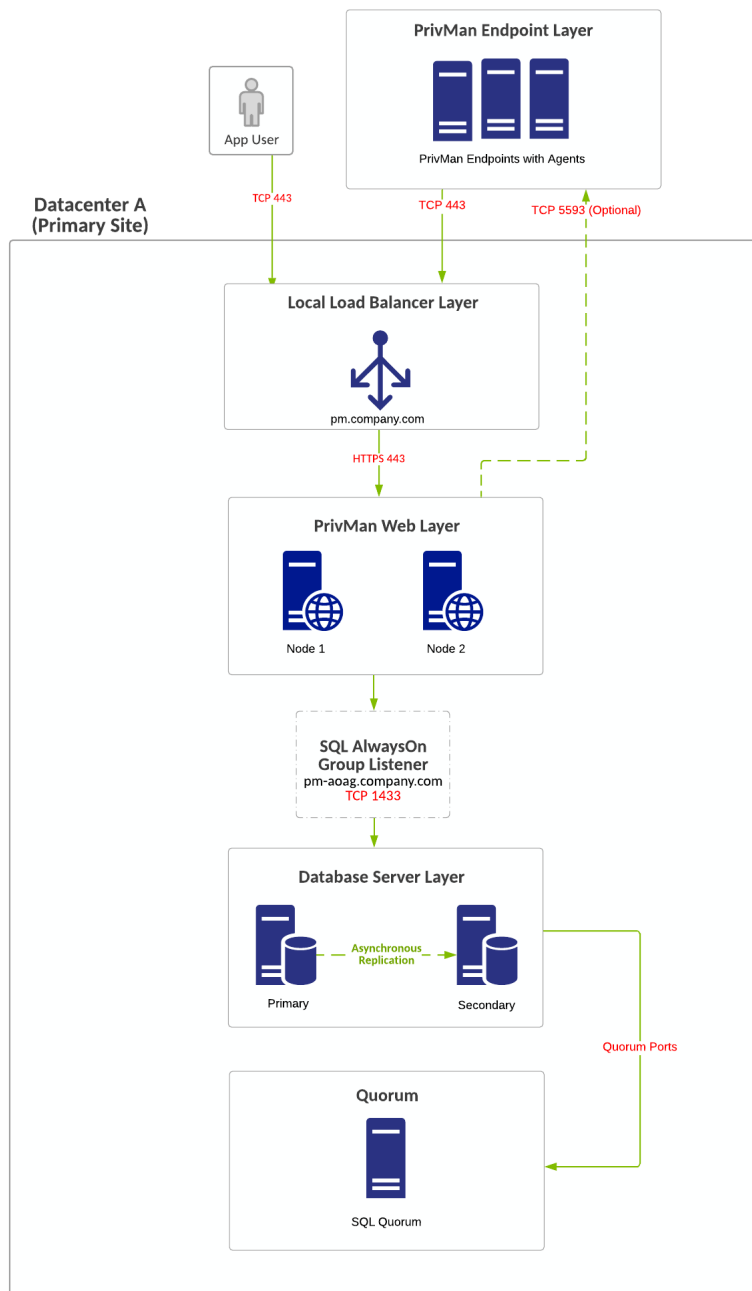
Virtual IP/Virtual Computer Object Requirements

- pm.company.com:443 (Load Balancer).
- pm-aoag.company.com:1433 (created as part of SQL AlwaysOn Configuration).
 - pm-aoag.company.com computer object/Virtual IP.
- Windows Failover Cluster Object (created as part of Windows Failover Clustering Configuration).
 - computer object/Virtual IP.
 - 1 additional virtual IP addresses may be required as part of Windows Failover Cluster for single site design for the network configuration of the Failover Cluster.

Diagram



The reference for this diagram is A-1.



Single Site Minimum HA (Reverse Proxy/Azure Bus)

Overview

- Minimum Cost HA Configuration.
- Single Site design, no native DR capacity. DR can be provided by means of VM replication if subnets are spanning locations, otherwise re-ip + DNS changes may be necessary.

Delinea Architecture Reference Diagrams

- Due to SQL Basic Availability groups with Standard Edition, you will need to have multiple instances of SQL and a dedicated AlwaysOn availability group configuration.
- Some customers may choose to use a separate web reverse proxy or azure service bus configuration for Privilege Manager agent TCP 443 communication.
- Optionally, Privilege Manager can be integrated with Secret Server (installed on the same web servers or, preferably, on dedicated servers) for authentication and as a storage vault for Privilege Manager credentials. If using Secret Server features (beyond authentication and vault storage for Privilege Manager), Secret Server should be on separate servers.

Requirements

- SQL Standard Edition - Basic Availability Group Configuration.
- Local load balancers can be utilized for all web server nodes.
- Configuring a file share witness for SQL quorum voting is required for SQL to stay online during single node unplanned failures.
- SQL Quorum Ports - <http://dsfnet.blogspot.com/2013/04/windows-server-clustering-sql-server.html>

Virtual IP or Computer Object Requirements

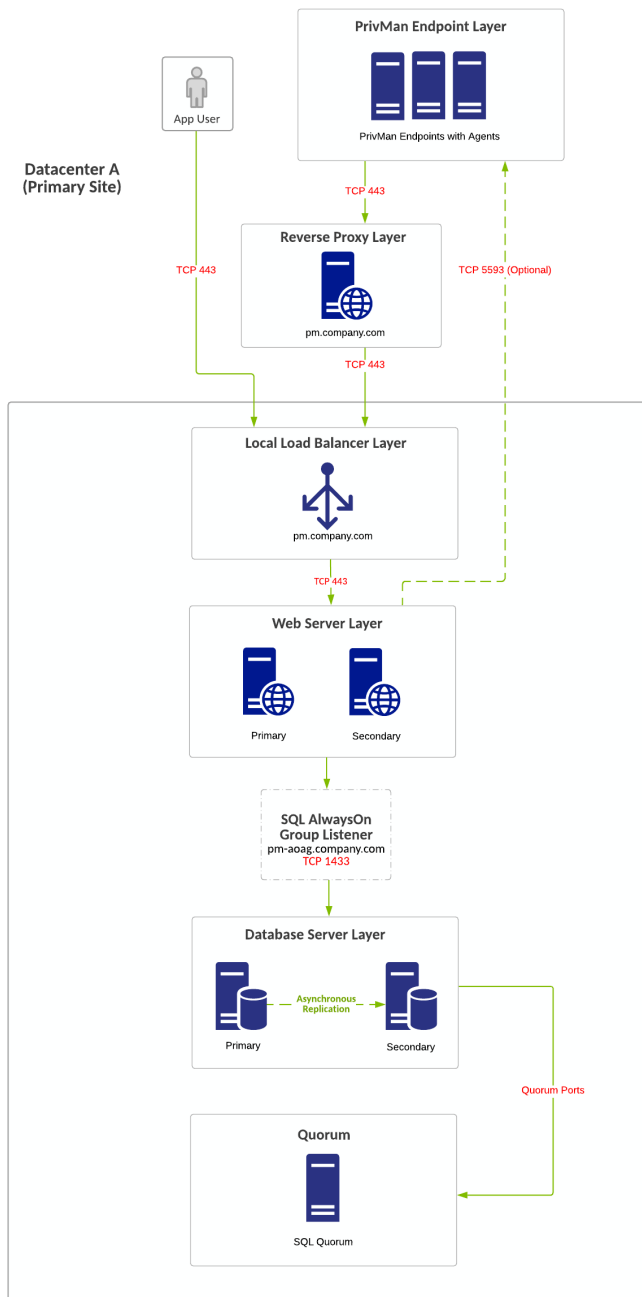
- pm.company.com:443 (Load Balancer).
- pm-aoag.company.com:1433 (created as part of SQL AlwaysOn Configuration).
 - pm-aoag.company.com computer object/Virtual IP.
- Windows Failover Cluster Object (created as part of Windows Failover Clustering Configuration).
 - computer object/Virtual IP.
 - 1 additional virtual IP addresses may be required as part of Windows Failover Cluster for single site design for the network configuration of the Failover Cluster.

Diagram



The reference for this diagram is A-2.

Delinea Architecture Reference Diagrams



Multi-Site Minimum HA/DR - Lower Cost, Manual Failover

Overview

- Minimum Cost HA Multi-Site Configuration - Lower Infrastructure Footprint for DR..
- Multi-Site Design. SQL AlwaysOn configurations will be asynchronous for Privilege Manager database.

Delinea Architecture Reference Diagrams

- DR site acts at temporary site only with no intention for long-term usage. Services in DR site being down can incur downtime.
- If a Global Load Balancer is not used, the Reverse Proxy and App Users will need to be directed to the DR Web Node, typically via DNS or IP updates.
- Due to SQL Basic Availability groups with Standard Edition, you will need to have multiple instances of SQL and a dedicated AlwaysOn availability group configuration.
- Privilege Manager can reside on the same database servers as Secret Server or separate database servers, but Secret Server and Privilege Manager should not share the same database itself. Due to SQL Basic Availability groups with Standard Edition, you will need to have multiple instances of SQL and a separate AlwaysOn availability group configuration.

Requirements

- SQL Standard Edition - Basic Availability Group Configuration.
- If no Global Load Balancers Exist due to costs/infrastructure missing, local load balancers can be utilized for all web server nodes but DNS change may be required if primary location goes offline.
- Configuring a file share witness for SQL quorum voting is required for SQL to stay online during single node unplanned failures.
- SQL Quorum Ports - <http://dsfnet.blogspot.com/2013/04/windows-server-clustering-sql-server.html>

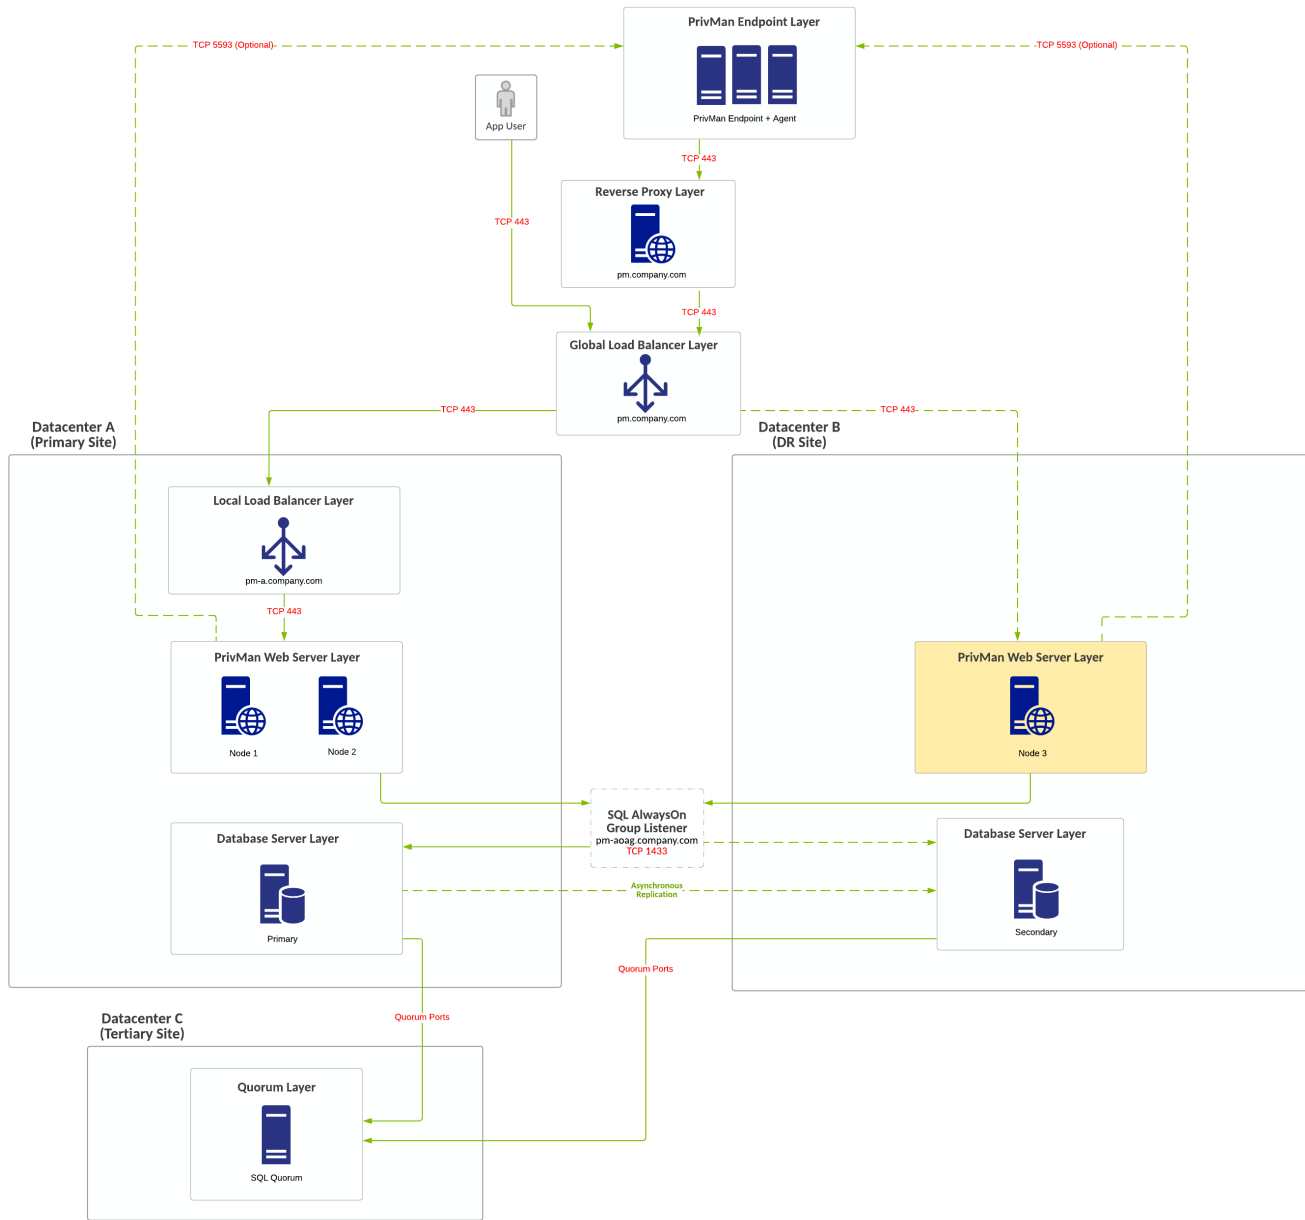
Virtual IP or Computer Object Requirements

- pm.company.com:443 (Load Balancer).
 - pm-a.company.com (Local Load Balancer).
- pm-aoag.company.com:1433 (created as part of SQL AlwaysOn Configuration).
 - pm-aoag.company.com computer object/Virtual IP.
- Windows Failover Cluster Object (created as part of Windows Failover Clustering Configuration).
 - computer object/Virtual IP.
 - 1 additional virtual IP addresses may be required as part of Windows Failover Cluster for single site design for the network configuration of the Failover Cluster.

Diagram



The reference for this diagram is B.



Multi-Site Average HA/DR - Average Cost, Manual Failover

Overview

- Average Cost HA Multi-Site Configuration - Lower Infrastructure Footprint for DR.
- Multi-Site Design. SQL AlwaysOn configurations will be asynchronous for Privilege Manager database.
- Secondary SQL Node at Primary Site for Planned Failover "Patching", Secondary SQL Node in DR Site for Unplanned Failover.
- DR site can act as permanent secondary site for long term use.

Delinea Architecture Reference Diagrams

- Database requires manual failover at primary or DR location.
- Global Load Balancers are configured to force all traffic to go to primary site unless primary site is down (priority group activation).
- Web node in DR Site is inactive and manually activated when failover is needed.
- If the data centers have low latency between networks, it may be possible to leave the PrivMan web server in DR online, active, and processing work.
- Some customers may choose to use a separate web reverse proxy, as shown, or azure service bus configuration for Privilege Manager agent TCP 443 communication.
- Optionally, Privilege Manager can be integrated with Secret Server (installed on the same web servers or, preferably, on dedicated servers) for authentication and as a storage vault for Privilege Manager credentials. If using Secret Server features (beyond authentication and vault storage for Privilege Manager), Secret Server should be on separate servers.

Requirements

- SQL Standard Edition.
- Global and Local Load Balancers.
- Configuring a file share witness for SQL quorum voting is recommended. A cloud witness or DFSR share is recommended for witness configuration. Simultaneous failure of both SQL nodes in the primary location can cause the failover cluster to not survive.
- SQL Quorum Ports - <http://dsfnet.blogspot.com/2013/04/windows-server-clustering-sql-server.html>

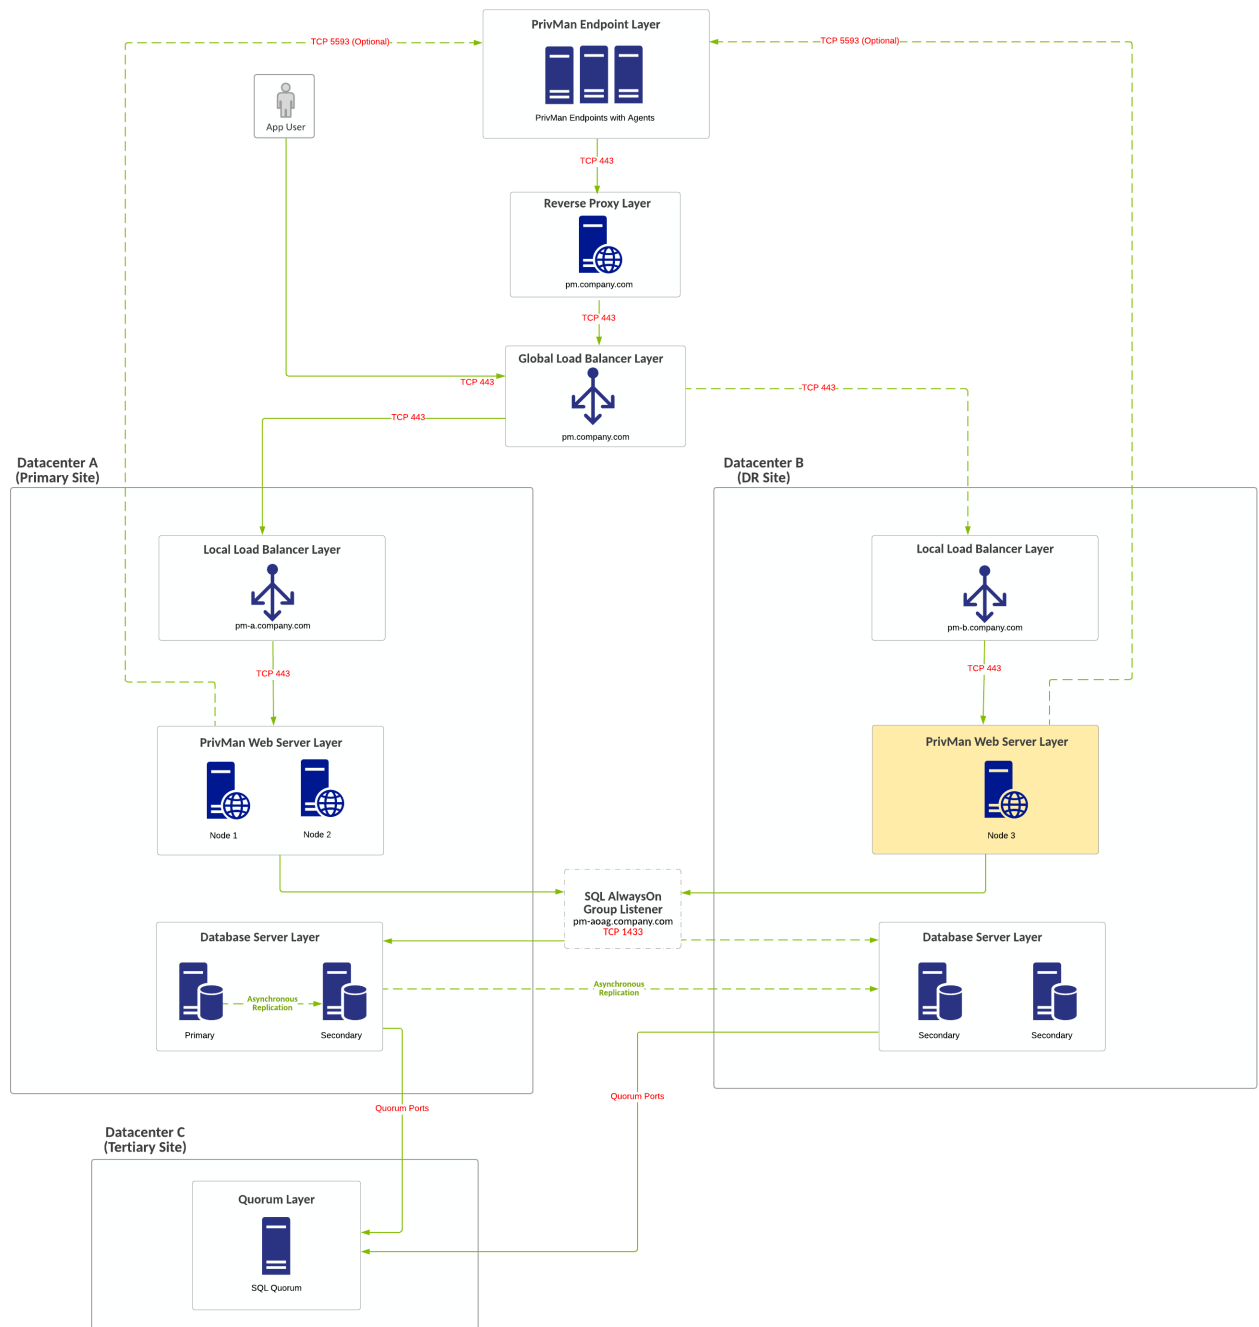
Virtual IP or Computer Object Requirements

- pm.company.com:443 (Global Load Balancer).
 - pm-a.company.com:443 (Local Load Balancer).
 - pm-b.company.com:443 (Local Load Balancer).
- pm-aoag.company.com:1433 (created as part of SQL AlwaysOn Configuration).
 - pm-aoag.company.com computer object/Virtual IP.
 - 2 virtual IP addresses may be required as part of this configuration.
- Windows Failover Cluster Object (created as part of Windows Failover Clustering Configuration).
 - computer object/Virtual IP.
 - 2 additional virtual IP addresses may be required as part of Windows Failover Cluster for single site design for the network configuration of the Failover Cluster representing both networks at each respective site.

Diagram



The reference for this diagram is C.



Best HA/DR - Highest Cost/Manual Failover

Overview

- Highest Cost HA Multi-Site Configuration - Higher Infrastructure Footprint in DR.
- Multi-Site Design. SQL AlwaysOn configurations will be asynchronous for Privilege Manager database.

Delinea Architecture Reference Diagrams

- Secondary SQL Node at Primary Site for Planned Failover "Patching", Secondary SQL Node in DR Site for Unplanned Failover.
- DR site can act as permanent secondary site for long term use.
- Database requires manual failover at primary or DR location.
- Global Load Balancers are configured to force all traffic to go to primary site unless primary site is down (priority group activation).
- Web nodes in DR Site is inactive and manually activated when failover is needed.
- If the data centers have low latency between networks, it may be possible to leave the PrivMan web server in DR online, active, and processing work.
- Some customers may choose to use a separate web reverse proxy, as shown, or azure service bus configuration for Privilege Manager agent TCP 443 communication.
- Optionally, Privilege Manager can be integrated with Secret Server (installed on the same web servers or, preferably, on dedicated servers) for authentication and as a storage vault for Privilege Manager credentials. If using Secret Server features (beyond authentication and vault storage for Privilege Manager), Secret Server should be on separate servers.

Requirements

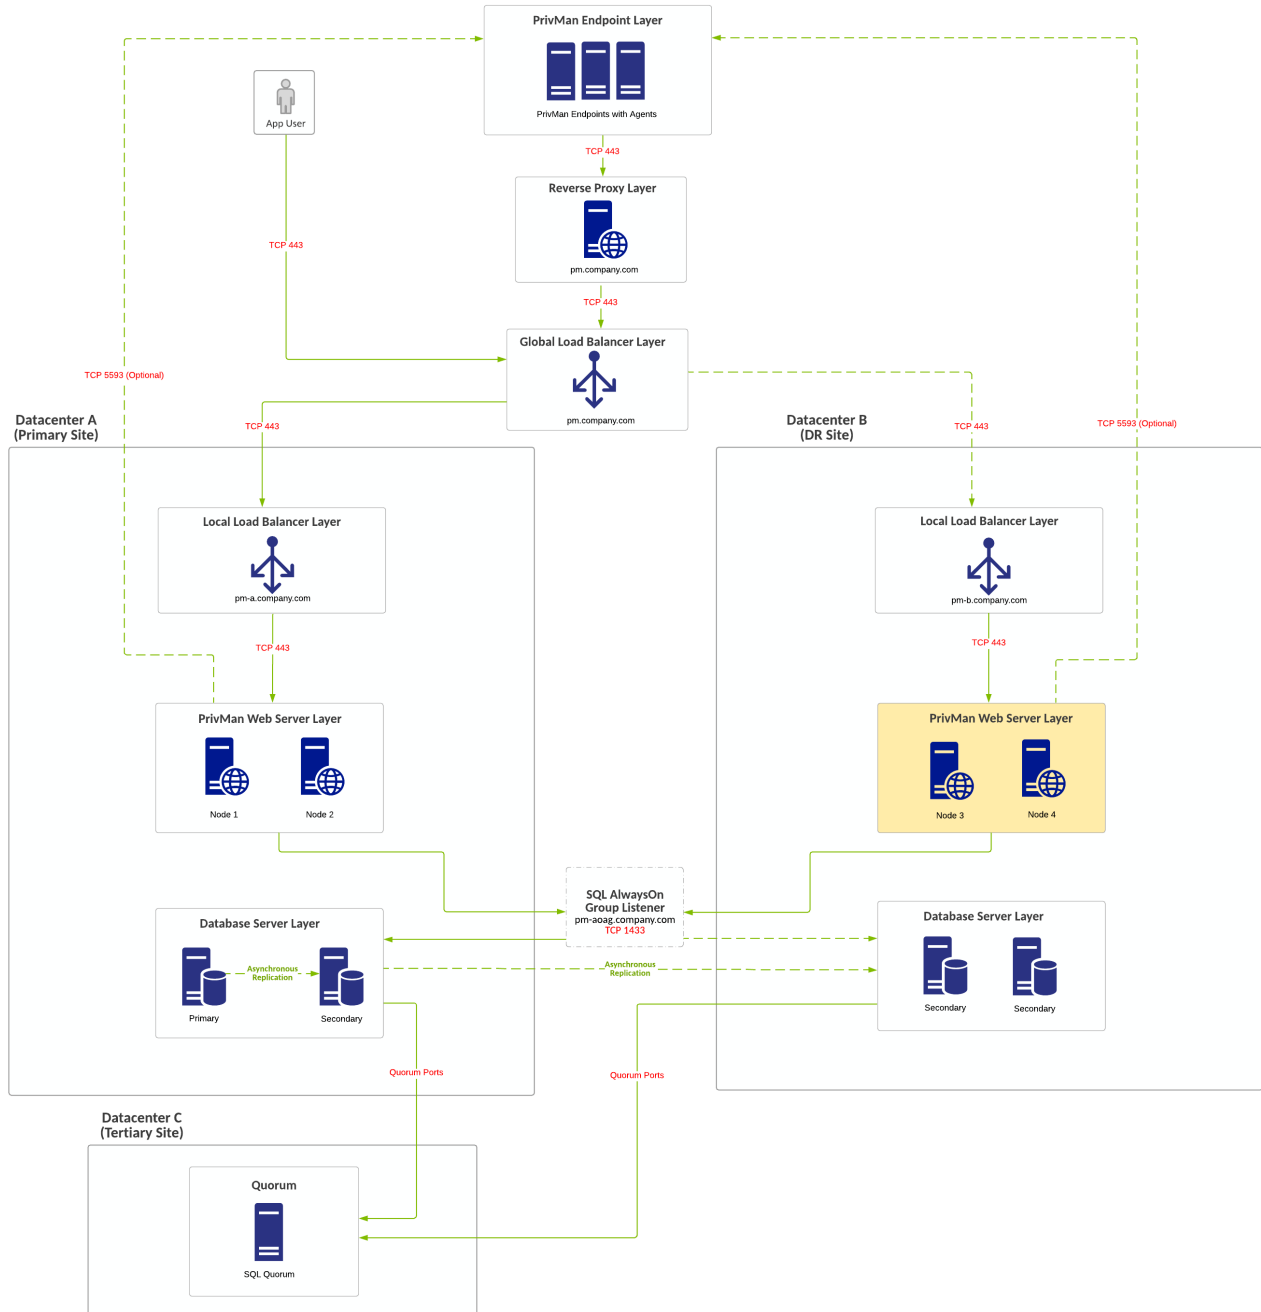
- SQL Standard Edition.
- Global and Local Load Balancers.
- Configuring a file share witness for SQL quorum voting is recommended. A cloud witness or DFSR share is recommended for witness configuration. Simultaneous failure of both SQL nodes in the primary location can cause the failover cluster to not survive.
- SQL Quorum Ports - <http://dsfnet.blogspot.com/2013/04/windows-server-clustering-sql-server.html>

Virtual IP/Virtual Computer Object Requirements

- pm.company.com:443 (Global Load Balancer).
 - pm-a.company.com:443 (Local Load Balancer).
 - pm-b.company.com:443 (Local Load Balancer).
- pm-aoag.company.com:1433 (created as part of SQL AlwaysOn Configuration).
 - pm-aoag.company.com computer object/Virtual IP.
 - 2 virtual IP addresses may be required as part of this configuration.
- Windows Failover Cluster Object (created as part of Windows Failover Clustering Configuration).
 - computer object/Virtual IP.
 - 2 additional virtual IP addresses may be required as part of Windows Failover Cluster for single site design for the network configuration of the Failover Cluster representing both networks at each respective site.

Diagram

 The reference for this diagram is D.



Secret Server Architectures

This section contains example architecture diagrams for Secret Server and related technologies.

Secret Server Cloud and On-Premises

- [Secret Server Architecture](#)
- [Distributed Engine Example Architectures](#)
- [RabbitMQ Helper Example Architectures](#)
- [Secret Server and DevOps Secrets Vault Example Architectures](#)
- [Secret Server and Privilege Manager Example Architectures](#)
- [Session Recording Example Architectures](#)

Secret Server Cloud Only

- [Secret Server Cloud Customer Example Architectures](#)
- [Secret Server Hybrid Multi-Tenant Cloud Architecture](#)

Distributed Engine Example Architectures



If you are a current customer with support hours for Delinea Professional Services, you can discuss any of these diagrams in detail with one of our Professional Services Solutions Architects.

This reference architecture is our best practice architecture for Secret Server Distributed Engines (DEs). The two most common use cases are:

- Distributing work across fire-walled networks using the fewest ports possible, to help ensure a better network security model.
- Separating work tasks away from the Web servers and placing the processing work on other dedicated servers within the Delinea infrastructure to improve overall performance.

Both of these use cases are covered with minimal and best high-availability solutions. The final reference architecture in this collection combines both uses cases with a high-availability solution.

Minimal HA Single-Site Deployment with No Distributed Engines

Overview



Please see [Secret Server Example Architectures](#) for additional design variations.

Requirements

General

- SQL Server Standard Edition with basic availability group configuration.
- SQL Server 2012 R2+.
- Use Windows authentication for SQL Server.
- You can use local load balancers for Web server nodes.

Delinea Architecture Reference Diagrams

- We require a file-share witness for SQL quorum voting for SQL to stay online during single-node unplanned failures.
- [Distributed Engine Ports](#)
- [Distributed Engine Proxy Configuration](#)
- [SQL Quorum Ports](#)

Virtual IP or Computer

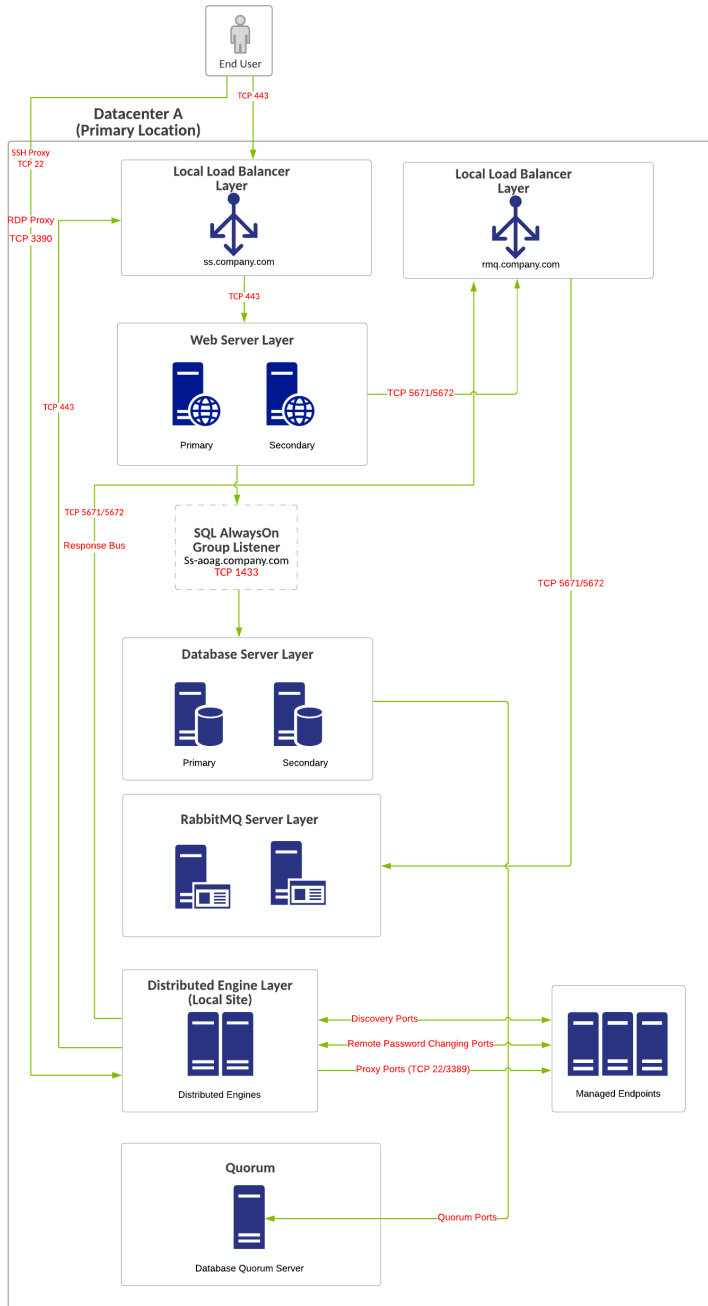
- ss.company.com: 443 (load balancer)
- ss.company.com: 5671 or 5672 (load balancer)
- ss-aoag.company.com: 1433 (created as part of SQL AlwaysOn configuration). Computer object or virtual IP.
- Windows Failover Cluster Object (created as part of Windows failover clustering configuration):
 - Computer object or virtual IP
 - One additional virtual IP address may be required as part of Windows failover cluster for single-site design for the network configuration of the failover cluster.

Diagram



Note: The reference number for this diagram is A1.

Figure: Minimal HA Single-Site Deployment with No Distributed Engines



Minimal HA Single-Site Deployment with Distributed Engines for Additional Data Centers

Overview

- Minimum-cost HA configuration.
- No shared storage requirement.
- RabbitMQ Helper installed on separate dedicated servers.

Delinea Architecture Reference Diagrams

- Two DEs for HA of local site, which is included with all licensing models.
- Distributed Engine licenses required for this design:
 - Three DE site licenses added (for DMZ, secondary, and cloud Locations), one DE included per site.
 - One DE per site license added, which allows for second DE in each DE site for HA.
- All DEs require callback communication to Web servers (TCP 443) and to the RabbitMQ Helper response bus (TCP 5672 or 5671). This is pictured with one set of distributed engines (local site) but is not pictured for other DEs to keep the diagram easier to interpret.



Please see [Secret Server Example Architectures](#) for additional design variations.

Requirements

General

- SQL Server Standard Edition with basic availability group configuration.
- SQL Server 2012 R2+.
- Use Windows authentication for SQL Server.
- You can use local load balancers for Web server nodes.
- We require a file-share witness for SQL quorum voting for SQL to stay online during single-node unplanned failures.
- [Distributed Engine Ports](#)
- [Distributed Engine Proxy Configuration](#)
- [SQL Quorum Ports](#)

Virtual IP or Computer

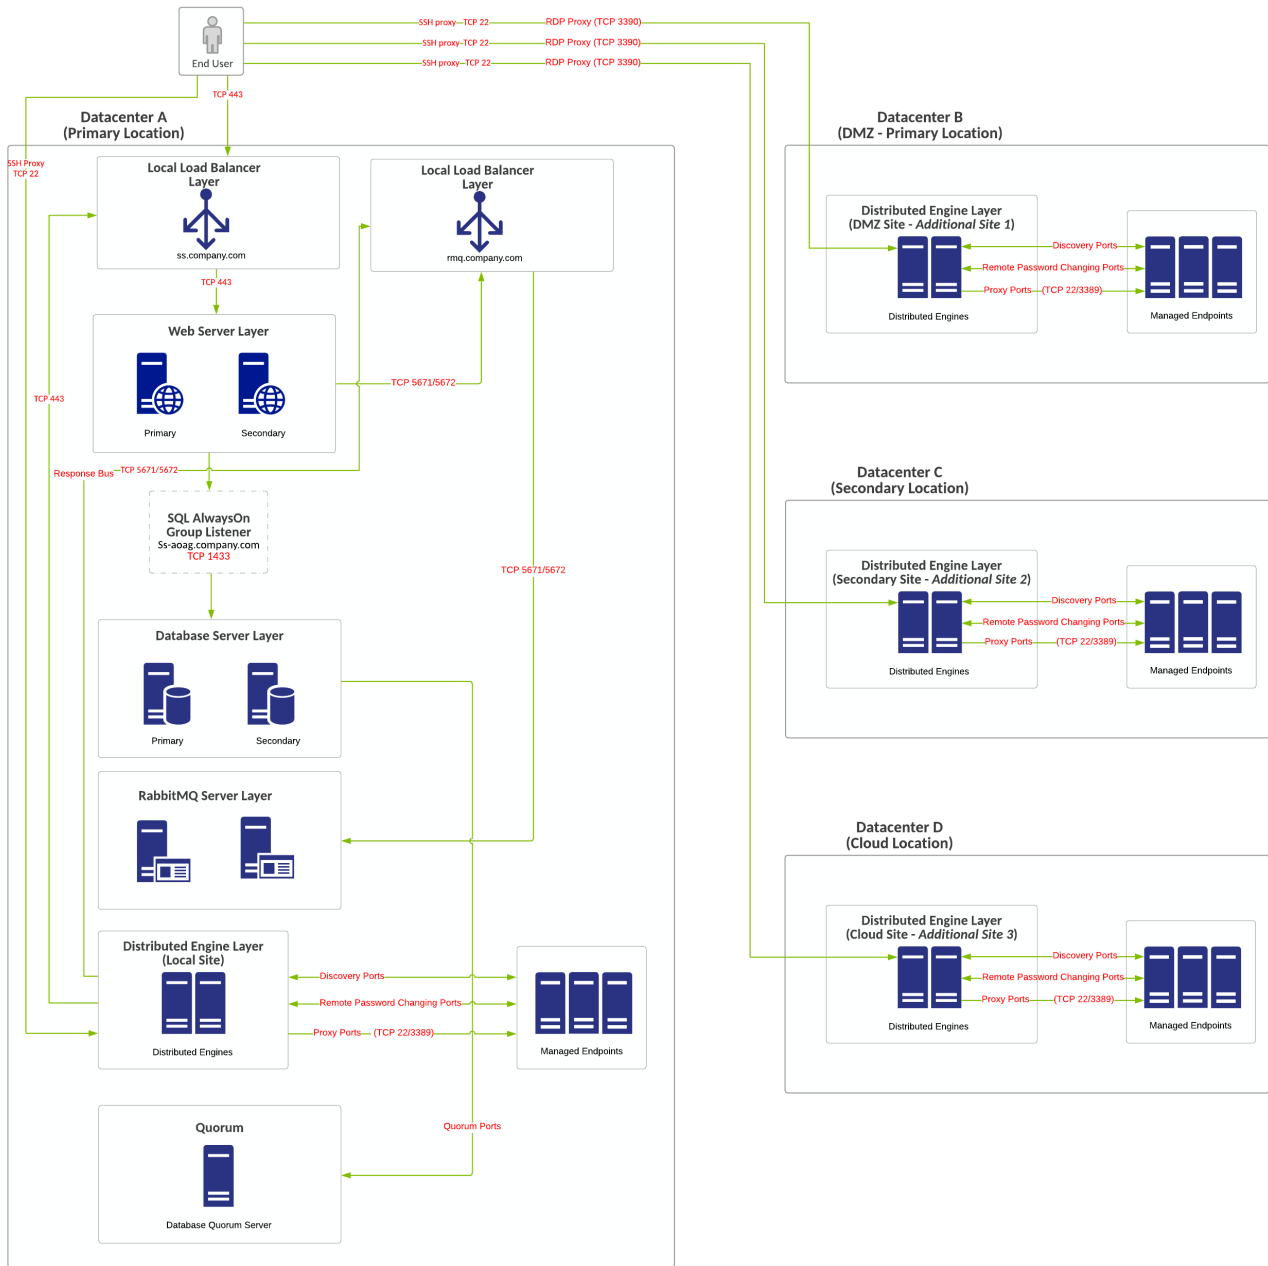
- ss.company.com: 443 (load balancer)
- ss.company.com: 5671 or 5672 (load balancer)
- ss-aoag.company.com: 1433 (created as part of SQL AlwaysOn configuration). Computer object or virtual IP.
- Windows Failover Cluster Object (created as part of Windows failover clustering configuration):
 - Computer object or virtual IP
 - One additional virtual IP address may be required as part of Windows failover cluster for single-site design for the network configuration of the failover cluster.

Diagram



The reference number for this diagram is A2.

Figure: Minimal HA Single-Site Deployment with Distributed Engines for Additional Data Centers.



Minimal HA Single-Site Deployment with Distribute Engines for Separate Work Tasks

Overview

- Minimum-cost HA configuration.
- No shared storage requirement.
- RabbitMQ Helper installed on separate dedicated servers.
- Two DEs for HA of local site, which is included with all licensing models.

Delinea Architecture Reference Diagrams

- Local site for AD or LDAP, SMTP, SIEM, or RADIUS integration.
- Distributed Engine licenses required for this design:
 - Two DE site licenses added (for secret and discovery tasks), one DE included per site.
 - One DE per site license added, which allows for second DE in each DE site for HA.
- Single-site design with no native DR capacity. DR can be provided by VM replication if subnets are spanning locations, otherwise re-IP + DNS changes may be necessary.
- All DEs require callback communication to Web servers (TCP 443) and to the RabbitMQ Helper response bus (TCP 5672 or 5671). This is pictured with one set of distributed engines (local site) but is not pictured for other DEs to keep the diagram easier to interpret.

 **Note:** Please see [Secret Server Example Architectures](#) for additional design variations.

Requirements

General

- SQL Server Standard Edition with basic availability group configuration.
- SQL Server 2012 R2+.
- Use Windows authentication for SQL Server.
- You can use local load balancers for Web server nodes.
- We require a file-share witness for SQL quorum voting for SQL to stay online during single-node unplanned failures.
- [Distributed Engine Ports](#)
- [Distributed Engine Proxy Configuration](#)
- [SQL Quorum Ports](#)

Virtual IP or Computer

- ss.company.com: 443 (load balancer)
- ss.company.com: 5671 or 5672 (load balancer)
- ss-aoag.company.com: 1433 (created as part of SQL AlwaysOn configuration). Computer object or virtual IP.
- Windows Failover Cluster Object (created as part of Windows failover clustering configuration):
 - Computer object or virtual IP
 - One additional virtual IP address may be required as part of Windows failover cluster for single-site design for the network configuration of the failover cluster.

Diagram


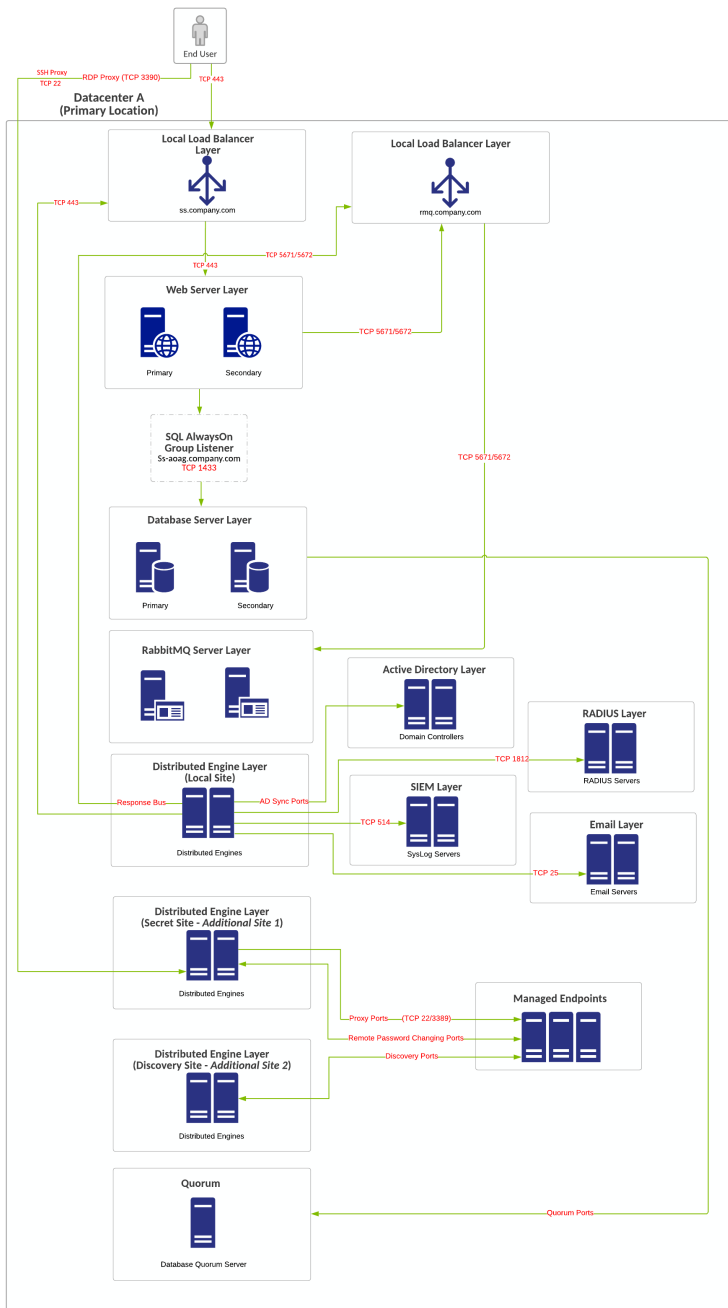
 The reference number for this diagram is A3.

Figure: Minimal HA Single-Site Deployment with Distribute Engines for Separate Work Tasks

Delinea Architecture Reference Diagrams



Best HA Multi-Site Deployment with No Distributed Engines

Overview

- All DEs require callback communication to Web servers (TCP 443) and to the RabbitMQ Helper response bus (TCP 5672 or 5671). This is pictured with one set of distributed engines (local site) but is not pictured for other DEs to keep the diagram easier to interpret.

 **Note:** Please see [Secret Server Example Architectures](#) for additional design variations.

Requirements

General

- SQL Server Standard Edition with basic availability group configuration.
- SQL Server 2012 R2+.
- Use Windows authentication for SQL Server.
- Global and local load balancers.
- We require a file-share witness for SQL quorum voting for SQL to stay online during single-node unplanned failures.
- [Distributed Engine Ports](#)
- [Distributed Engine Proxy Configuration](#)
- [SQL Quorum Ports](#)

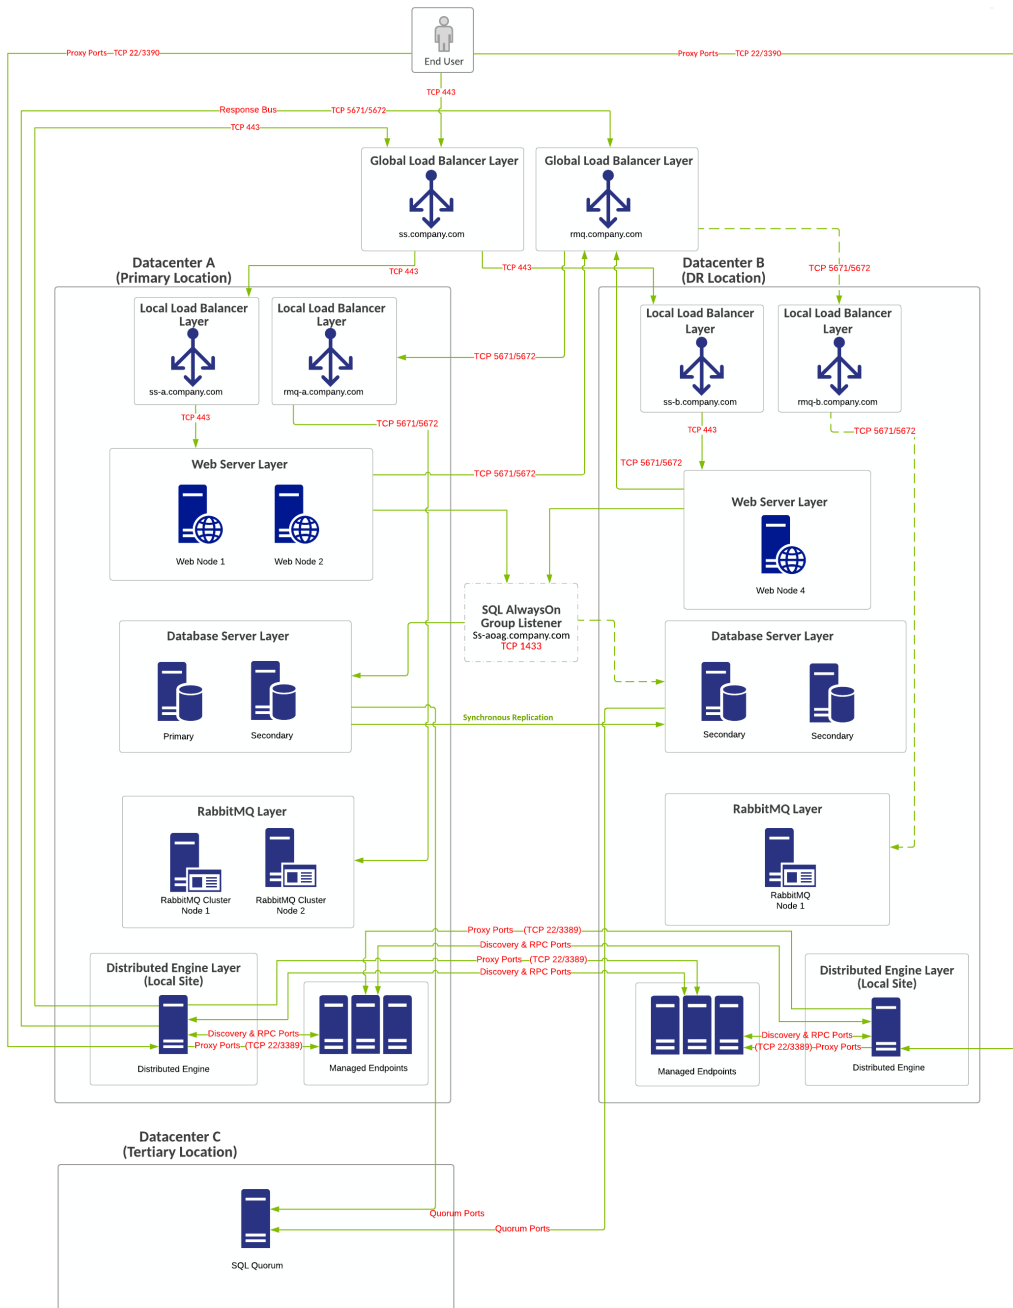
Virtual IP or Computer

- ss.company.com: 443 and rmq.company.com: 5671 or 5672 (two virtual IPs—global load balancer).
- ss-a.company.com: 443 and ss-b.company.com:443 (two virtual IPs—local load balancer).
- rmq-a.company.com: 5671 or 5672 (load balancer) and rmq-b.company.com: 5671 or 5672 (two virtual IPs—local load balancer).
- ss-a.company.com: 443 and ss-b.company.com:443 (two virtual IPs—local load balancer).
- Windows Failover Cluster Object (created as part of Windows failover clustering configuration):
 - Computer object or virtual IP
 - One additional virtual IP address may be required as part of Windows failover cluster for single-site design for the network configuration of the failover cluster.

Diagram



Figure: Best HA Multi-Site Deployment with No Distributed Engines



Best HA Multi-Site Deployment with Distributed Engines for Additional Datacenters

Overview

- Two DE site licenses added (for DMZ and cloud locations), one DE included per site.
- One DE per site license added, which allows for second DE in each DE site for HA and a third one for the local site (added to the primary location).

Delinea Architecture Reference Diagrams

- All DEs require callback communication to Web servers (TCP 443) and to the RabbitMQ Helper response bus (TCP 5672 or 5671). This is pictured with one set of distributed engines (local site) but is not pictured for other DEs to keep the diagram easier to interpret.

 **Note:** Please see [Secret Server Example Architectures](#) for additional design variations.

Requirements

General

- SQL Server Enterprise Edition with availability group configuration.
- SQL Server 2012 R2+.
- Use Windows authentication for SQL Server.
- Global and local load balancers.
- We require a file-share witness for SQL quorum voting for SQL to stay online during single-node unplanned failures.
- [Distributed Engine Ports](#)
- [Distributed Engine Proxy Configuration](#)
- [SQL Quorum Ports](#)

Virtual IP or Computer

- ss.company.com: 443 and rmq.company.com: 5671 or 5672 (two virtual IPs—global load balancer).
- ss-a.company.com: 443 and ss-b.company.com: 443 (two virtual IPs—local load balancer).
- rmq-a.company.com: 5671 or 5672 (load balancer) and rmq-b.company.com: 5671 or 5672 (two virtual IPs—local load balancer).
- ss-a.company.com: 443 and ss-b.company.com: 443 (two virtual IPs—local load balancer).
- Windows Failover Cluster Object (created as part of Windows failover clustering configuration):
 - Computer object or virtual IP
 - Two additional virtual IP addresses may be required as part of Windows failover cluster for single-site design for the network configuration of the failover cluster, representing both networks at each site.

Diagram


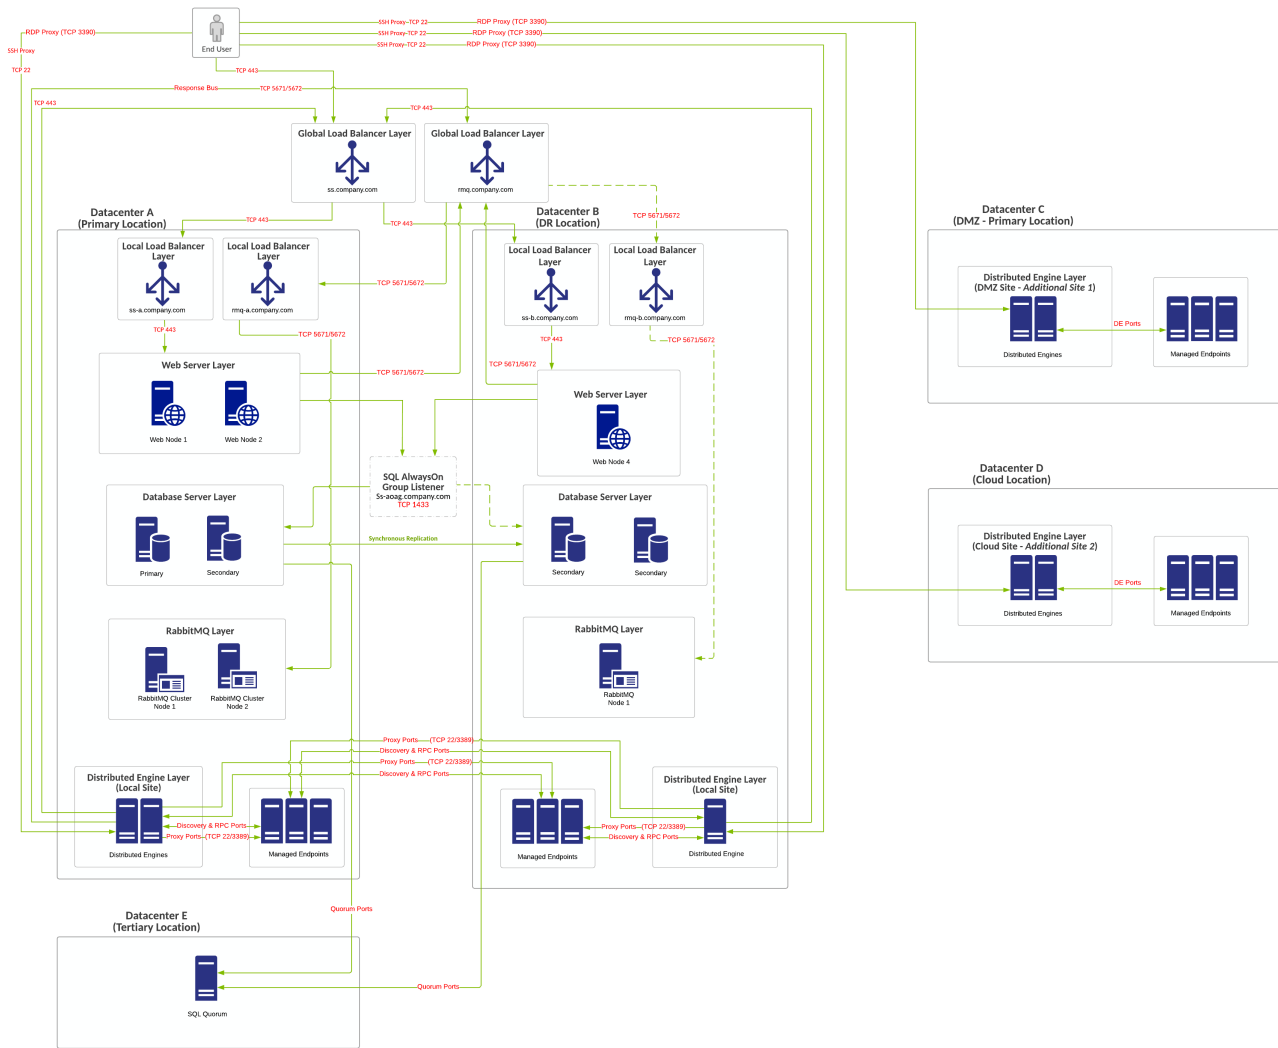
 The reference number for this diagram is B2.

Figure: Best HA Multi-Site Deployment with Distributed Engines for Additional Data Centers.

Delinea Architecture Reference Diagrams



Best HA Multi-Site Deployment with Distributed Engines for Separate Work Tasks

Overview

- Distributed Engine licenses required for this design:
 - Two DE site licenses added (for secret and discovery tasks), one DE included per site.
 - One DE per site license added, which allows for second DE in each DE site for HA and a third one for the local site (added to the primary location).
- All DEs require callback communication to Web servers (TCP 443) and to the RabbitMQ Helper response bus (TCP 5672 or 5671). This is pictured with one set of distributed engines (local site) but is not pictured for other DEs to keep the diagram easier to interpret.

Please see [Secret Server Example Architectures](#) for additional design variations.

Requirements

General

- SQL Server Enterprise Edition with availability group configuration.
- SQL Server 2012 R2+.
- Use Windows authentication for SQL Server.
- Global and local load balancers.
- We recommend a file-share witness for SQL quorum voting. We recommend a cloud witness or DFSR share for witness configuration. This can handle the failure of both SQL Server nodes in the primary location.
- [Distributed Engine Ports](#)
- [Distributed Engine Proxy Configuration](#)
- [SQL Quorum Ports](#)

Virtual IP or Computer

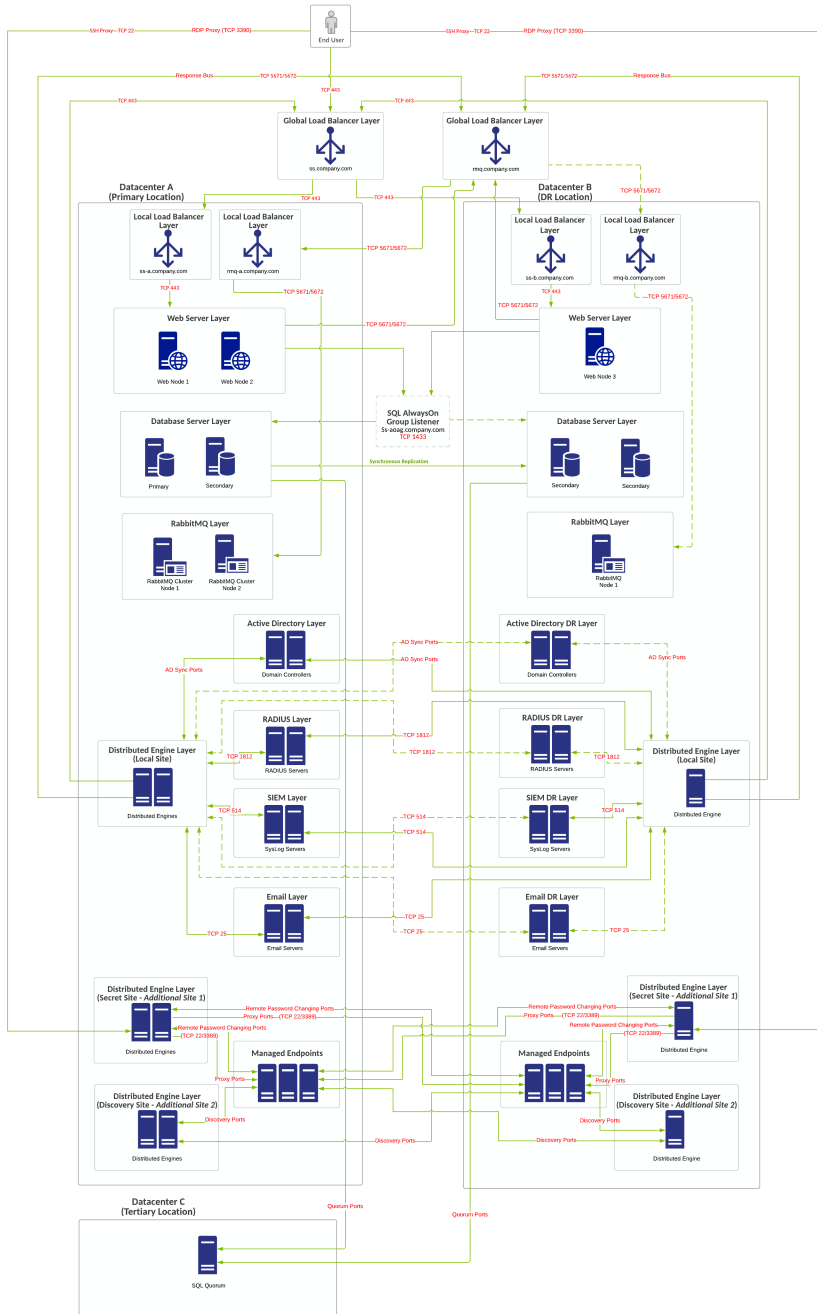
- ss.company.com: 443 and rmq.company.com: 5671 or 5672 (two virtual IPs—global load balancer).
- ss-a.company.com: 443 and ss-b.company.com: 443 (two virtual IPs—local load balancer).
- rmq-a.company.com: 5671 or 5672 (load balancer) and rmq-b.company.com: 5671 or 5672 (two virtual IPs—local load balancer).
- ss-a.company.com: 443 and ss-b.company.com: 443 (two virtual IPs—local load balancer).
- Windows Failover Cluster Object (created as part of Windows failover clustering configuration):
 - Computer object or virtual IP
 - Two additional virtual IP addresses may be required as part of Windows failover cluster for single-site design for the network configuration of the failover cluster, representing both networks at each site.

Diagram



Figure: Best HA Multi-Site Deployment with Distributed Engines for Separate Work Tasks

Delinea Architecture Reference Diagrams



Best HA Multi-Site Deployment with Distributed Engines for Additional Datacenters with Separate Work Tasks

Overview

- Distributed Engine licenses required for this design:
 - Five DE site licenses added (for primary secret and discovery tasks, DMZ site, and cloud secret and discovery tasks), one DE included per site.
 - Two DE per site licenses added, which allows for second DE in each DE site for HA and a third one for the local site (added to the primary and DR locations).
- All DEs require callback communication to Web servers (TCP 443) and to the RabbitMQ Helper response bus (TCP 5672 or 5671). This is pictured with one set of distributed engines (local site) but is not pictured for other DEs to keep the diagram easier to interpret.



Please see [Secret Server Example Architectures](#) for additional design variations.

Requirements

General

- SQL Server Enterprise Edition with availability group configuration.
- SQL Server 2012 R2+.
- Use Windows authentication for SQL Server.
- Global and local load balancers.
- We recommend a file-share witness for SQL quorum voting. We recommend a cloud witness or DFSR share for witness configuration. This can handle the failure of both SQL Server nodes in the primary location.
- [Distributed Engine Ports](#)
- [Distributed Engine Proxy Configuration](#)
- [SQL Quorum Ports](#)

Virtual IP or Computer

- ss.company.com: 443 and rmq.company.com: 5671 or 5672 (two virtual IPs—global load balancer).
- ss-a.company.com: 443 and ss-b.company.com: 443 (two virtual IPs—local load balancer).
- rmq-a.company.com: 5671 or 5672 (load balancer) and rmq-b.company.com: 5671 or 5672 (two virtual IPs—local load balancer).
- ss-aoag.company.com: 1433 (created as part of SQL AlwaysOn configuration). Computer object or virtual IP. May require two virtual IP addresses.
- Windows Failover Cluster Object (created as part of Windows failover clustering configuration):

Delinea Architecture Reference Diagrams

- Computer object or virtual IP
- Two additional virtual IP addresses may be required as part of Windows failover cluster for single-site design for the network configuration of the failover cluster, representing both networks at each site.

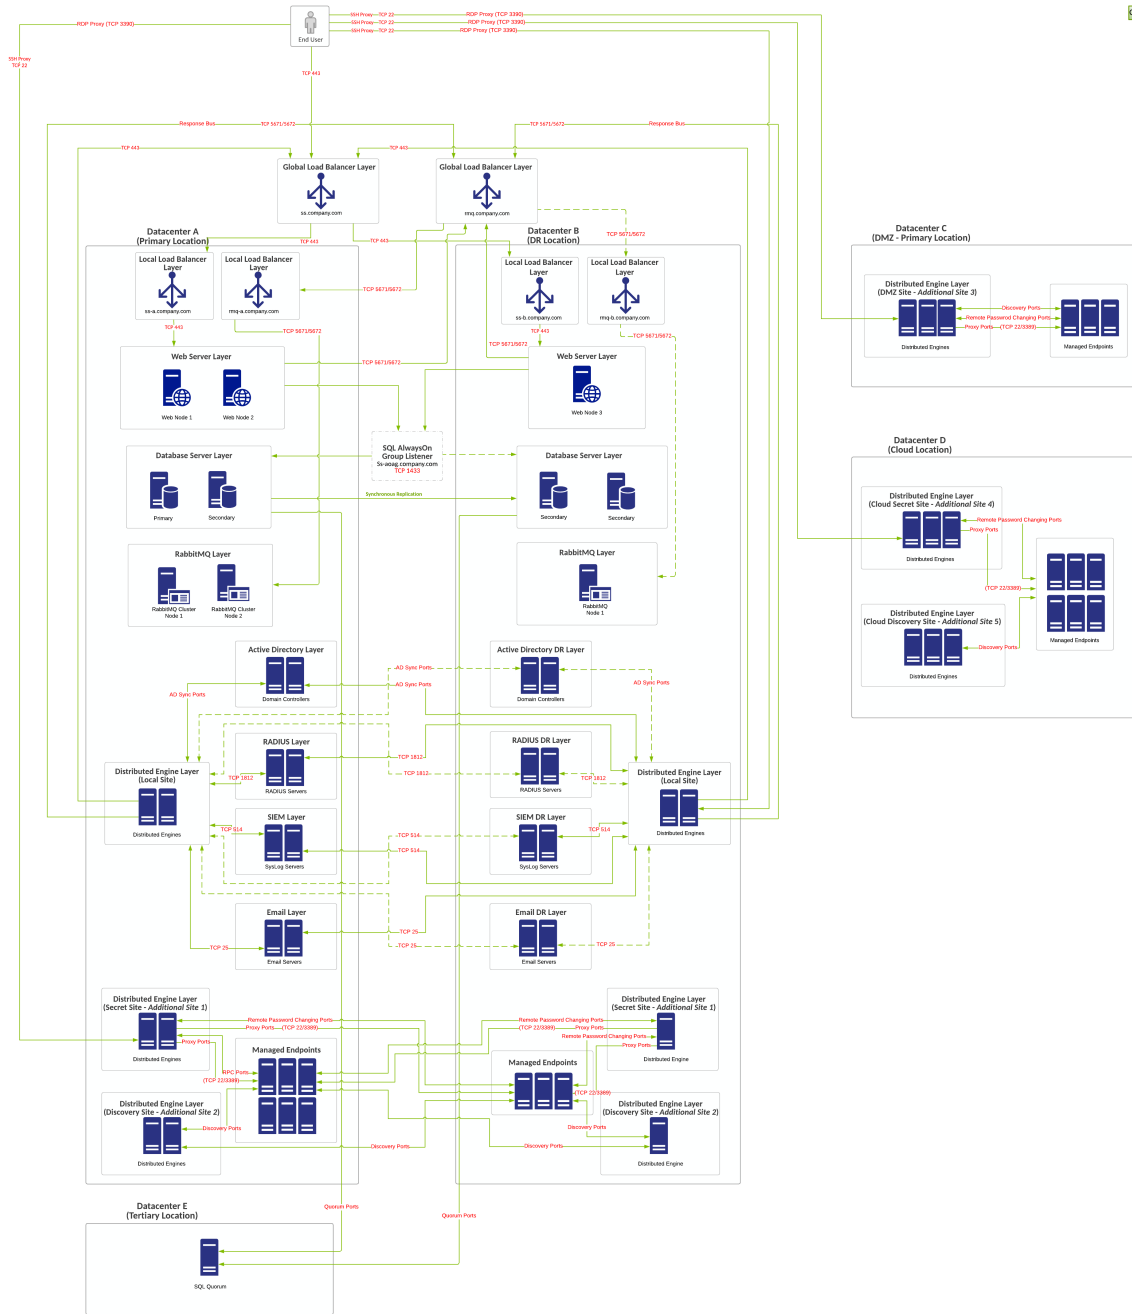
Diagram




The reference number for this diagram is C1.

Figure: Best HA Multi-Site Deployment with Distributed Engines for Additional Data Centers with Separate Work Tasks

Delinea Architecture Reference Diagrams



RabbitMQ Helper Example Architectures

 If you are a current customer with support hours for Delinea Professional Services, you can discuss any of these diagrams in detail with one of our Professional Services Solutions Architects.

Basic Single-Site Deployment

Overview

- This is the lowest cost and most simple option.
- No high-availability (HA) design.
- No failover capability of RabbitMQ Helper functions in a DR site.
- Distributed engines communicate to SS over callback ports (typically TCP 443, which is not shown in the diagram).
- RabbitMQ Helper traffic is set via static DNS entry and communicates by TCP port 5671 or 5672 (do not use if using TLS).
- All Sites connect to one site connector in SS.

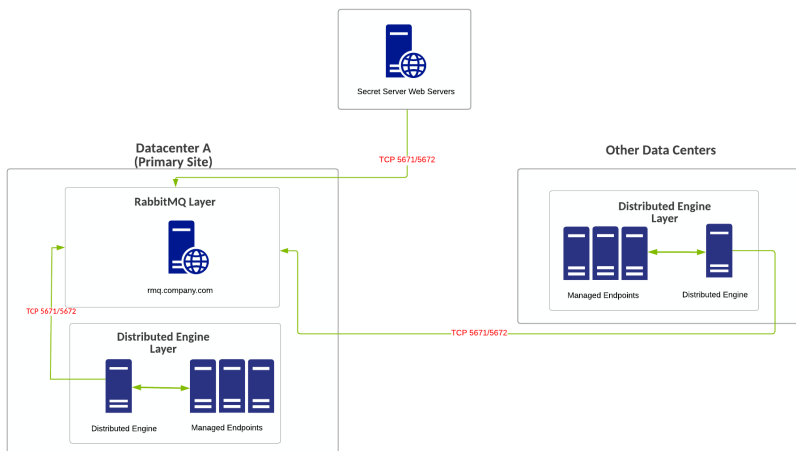
Requirements

- Cross-data-center communication between DEs and RabbitMQ Helper
- Does not require the creation of RabbitMQ Helper policies, which are based on input from Professional Services.

Diagram

 The reference number for this diagram is A.

Figure: Basic Single-Site Deployment



Basic Multi-Site Deployment

Overview

- This is the basic single-site deployment plus a hot standby.
- No HA design.
- Hot standby of RabbitMQ Helper for DR.
- Distributed engines communicate to SS over callback ports (typically TCP 443, which is not shown in the diagram).
- RabbitMQ Helper traffic is set via static DNS entry and communicates by TCP port 5671 or 5672 (do not use if using TLS). You can use a CNAME record and update the record during failover.
- All Sites connect to one site connector in SS.

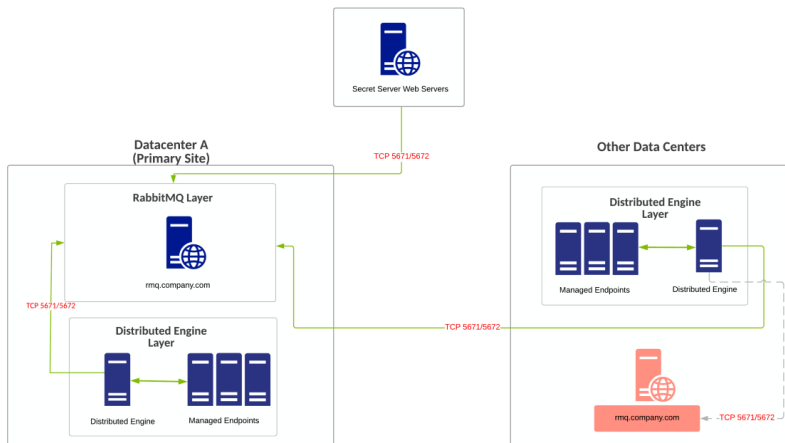
Requirements

- Cross-data-center communication between DEs and RabbitMQ Helper.
- Does not require the creation of RabbitMQ Helper policies, which are based on input from Professional Services.

Diagram

 The reference number for this diagram is B.

Figure: Basic Multi-Site Deployment



Minimum HA Multi-Site Deployment with Independent Nodes

Overview

- Minimum HA RabbitMQ Helper design.
- Independent RabbitMQ Helper nodes under load balancer configuration.
- No failover capability of RabbitMQ Helper functions in DR site.
- Distributed engines communicate to SS over callback ports (typically TCP 443, which is not shown in the diagram).
- RabbitMQ Helper load balancer configuration only sends traffic to only node one unless it is down, then then it sends traffic to node two. Communicates by TCP port 5671 or 5672 (do not use if using TLS) . This provides localized high availability for patching and more.
- Singular nodes built as part of a load balancer configuration between multiple data centers is also a design possibility with the same type of load balancer configuration mentioned above. Do so will lose localized HA of RabbitMQ Helper.
- All Sites connect to one site connector in Secret Server (RabbitMQ Helper FQDN).

Requirements

- Cross-data-center communication between DEs and RabbitMQ Helper.
- Does not require the creation of RabbitMQ Helper policies, which are based on input from Professional Services.
- Local load balancer.

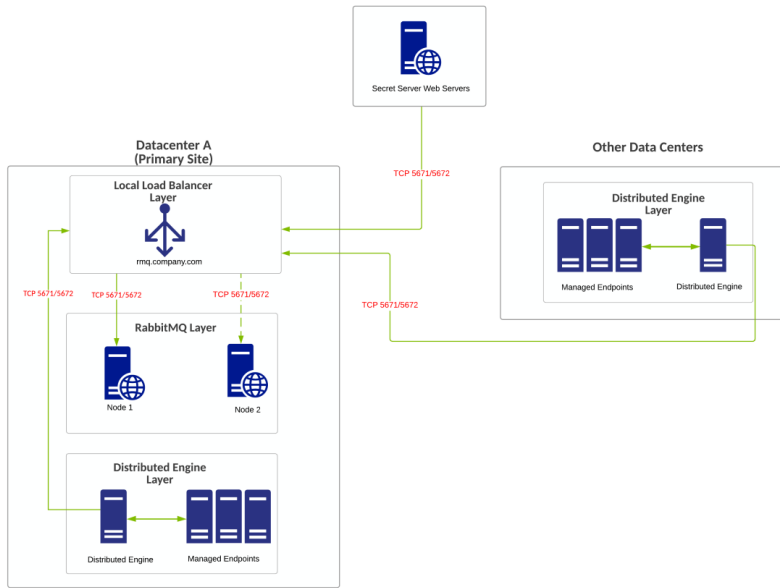
Diagram



The reference number for this diagram is C.

Figure: Minimum HA Multi-Site Deployment with Independent Nodes

Delinea Architecture Reference Diagrams



Minimum HA Multi-Site Deployment with Cluster

Overview

- Minimum HA RabbitMQ Helper design
- No failover capability of RabbitMQ Helper functions in DR site.
- Distributed engines communicate to SS over callback ports (typically TCP 443, which is not shown in the diagram).
- RabbitMQ Helper traffic is set to round-robin through load balancers to RabbitMQ Helper nodes in cluster. Communicates by TCP port 5671 or 5672 (do not use if using TLS).
- All Sites connect to one site connector in Secret Server (RabbitMQ Helper FQDN).

Requirements

- Cross-data-center communication between DEs and RabbitMQ Helper.
- Requires the creation of RabbitMQ Helper policies, which are based on input from Professional Services.
- Local load balancer.

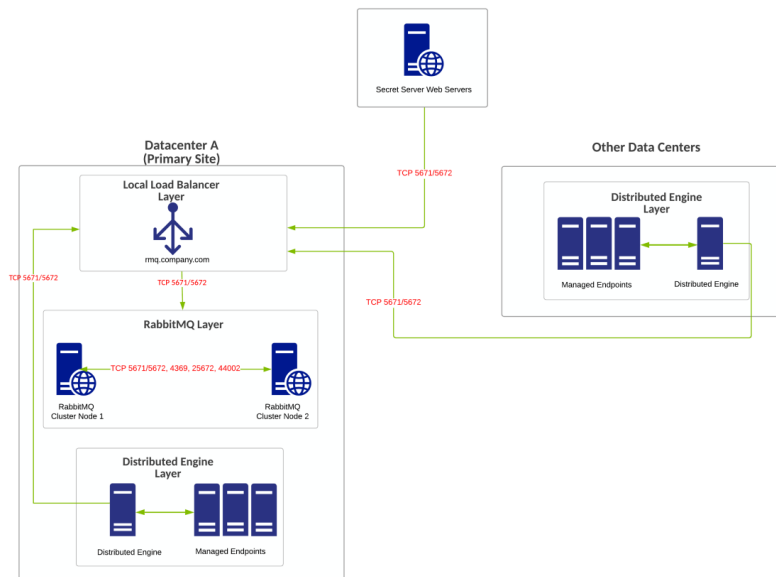
Diagram



The reference number for this diagram is D.

Figure: Minimum HA Multi-Site Deployment with Cluster

Delinea Architecture Reference Diagrams



Average HA/DR Multi-Site Deployment

Overview

- RabbitMQ Helper clusters in multiple locations (typically primary and secondary DR site).
- Distributed engines communicate to SS over callback ports (typically TCP 443, which is not shown in the diagram).
- RabbitMQ Helper traffic for each local load balancer configuration is set to round-robin traffic through load balancer to RabbitMQ Helper nodes in each respective cluster. Communicates by TCP port 5671 or 5672 (do not use if using TLS).
- Two site connectors configured in SS. This is a good design for manual failover between the primary and disaster recovery site but will require manual changes within SS for full functionality to resume in DR. The distributed engines global configuration has only the option for one response bus. This implies that traffic may be sent to one RabbitMQ Helper cluster but will be processed via response by the primary data center RabbitMQ Helper cluster.
- The dotted line traffic for the DR distributed engines back to the RabbitMQ Helper cluster indicates the down primary data center being down and the DR data center being online. In this case, this traffic would become active if the primary site is down. Any functions or secrets in SS assigned explicitly to that site connector will not function without reassigning those functions to the other site connector and RabbitMQ Helper cluster.

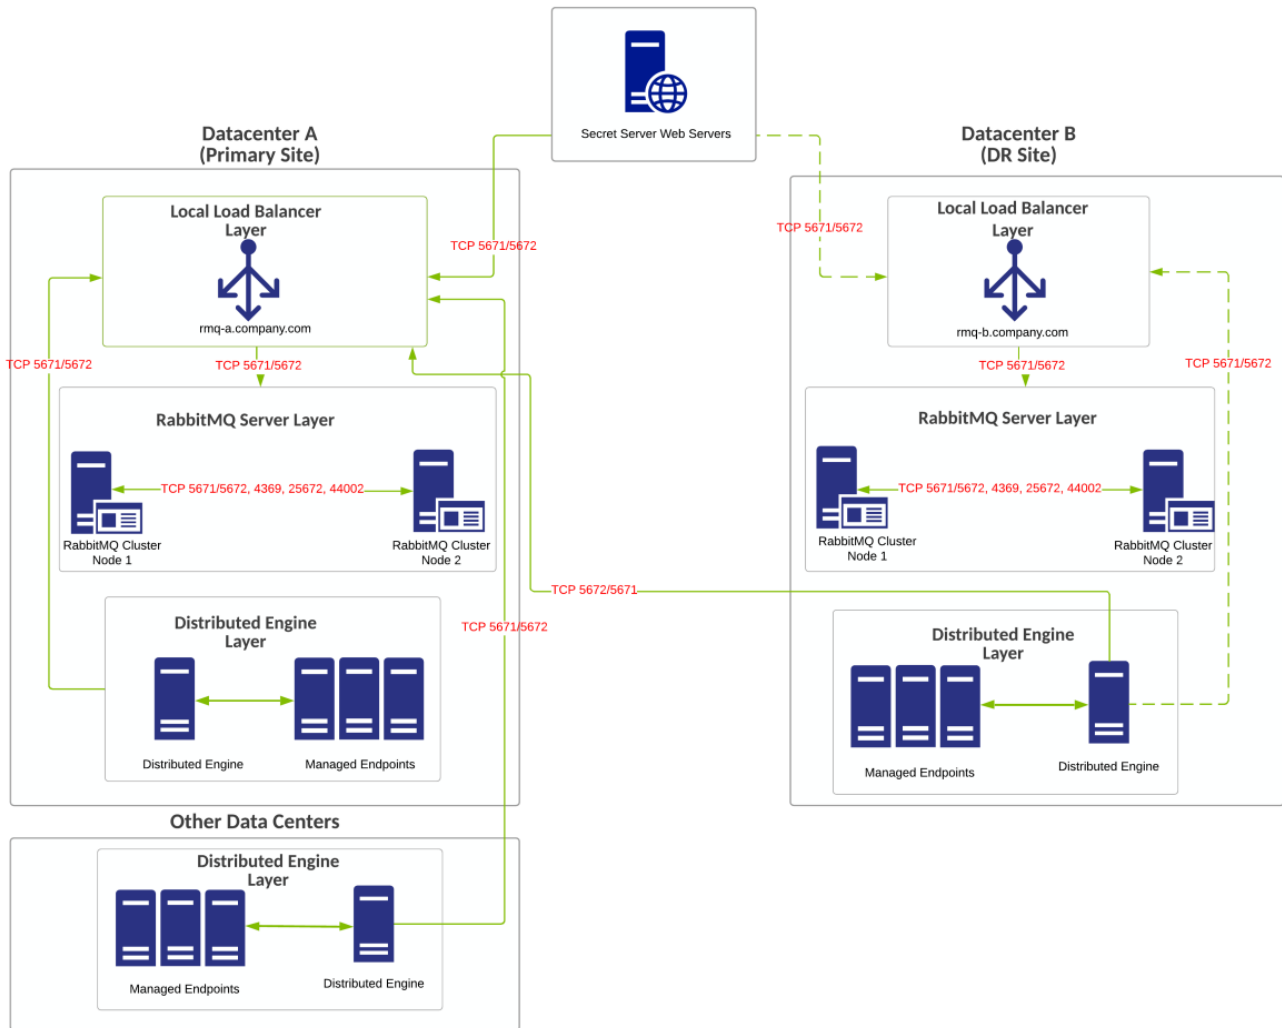
Requirements

- Requires the creation of RabbitMQ Helper policies, which are based on input from Professional Services.
- Multiple load balancers.

Diagram

 The reference number for this diagram is E.

Figure: Average HA/DR Multi-Site Deployment



Best HA/DR Multi-Site Deployment

Overview

- RabbitMQ Helper clusters are in multiple locations (typically primary and secondary DR site).
- Distributed engines communicate to SS over callback ports (typically TCP 443, which is not shown in the diagram).

Delinea Architecture Reference Diagrams

- RabbitMQ Helper traffic is set to force all traffic to the primary cluster in the primary site, unless the primary site is down. Communication through load balancer is via TCP port 5671 or 5672 (do not use if using TLS).
- Options available for site connector creation or use:
 - Sites and site connectors are intended to be primarily location based.
 - This design is best for situations where you want to localize RabbitMQ Helper traffic and control traffic directly through the load balancer configurations.
- This design is best suited for true DR situations where we assume when the primary location is down or offline that the entire data center is also.
- If the primary location is only partially down (RabbitMQ Helper only is offline), it is possible secrets or features in SS assigned to that specific site connector will not function correctly.

Site connector options:

- Three site connectors and three or more sites configured in SS to localize RabbitMQ Helper traffic for each respective location-based RabbitMQ Helper cluster:
 - One site connector for the global load balancer URL (Used for all other locations/sites).
 - One site connector for primary site.
 - One site connector for secondary (DR) site.
- Other data center's distributed engines typically connect to the primary site RabbitMQ Helper cluster global load balancer FQDN/VIP.
- Alternatively, if traffic isolation is not as important, you can have one site connector and all sites point to it. Cross-data-center communication between DEs and all RabbitMQ Helper clusters is required. A one-site connector design will eliminate application-side changes during failover.
- Traffic should be forced to RabbitMQ Helper cluster or nodes in the primary site for proper functionality.

Requirements

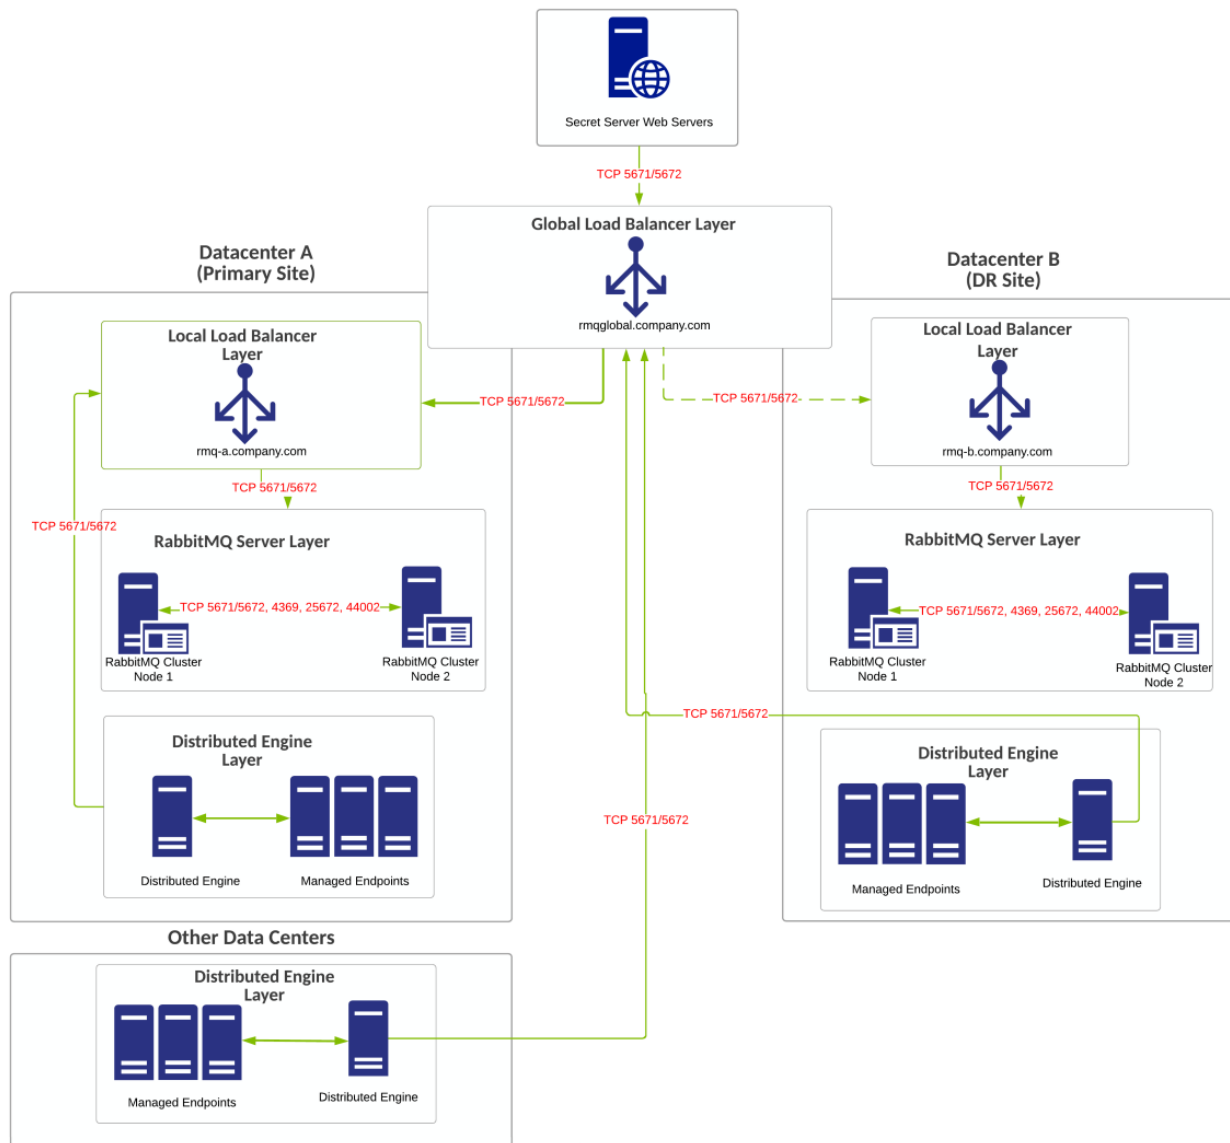
- Require the creation of RabbitMQ Helper policies, which are based on input from Professional Services.
- Multiple load balancers.
- Global load balancers.

Diagram



The reference number for this diagram is F.

Figure: Best HA/DR Multi-Site Deployment



Secret Server Example Architectures



If you are a current customer with support hours for Thycotic Professional Services, you can discuss any of these diagrams in detail with one of our Professional Services Solutions Architects.

Minimal Single-Site Enterprise Deployment

Overview

- Single-site minimum-cost HA configuration.
- No shared storage requirement.

Delinea Architecture Reference Diagrams

- RabbitMQ Helper Installed on SS Web servers (typically in a cluster on a primary + secondary node).
- Single-site design with no native DR capacity. DR can be provided by VM replication if subnets are spanning locations, otherwise re-IP + DNS changes may be necessary.

Requirements

General

- SQL Standard Edition: Basic availability group configuration. Local load balancers can be used for all Web server nodes.
- Configuring a file share witness for SQL quorum voting is required for SQL to stay online during single node unplanned failures.
- [Distributed Engine Ports](#).
- [SQL Quorum Ports](#).

Virtual IP or Computer

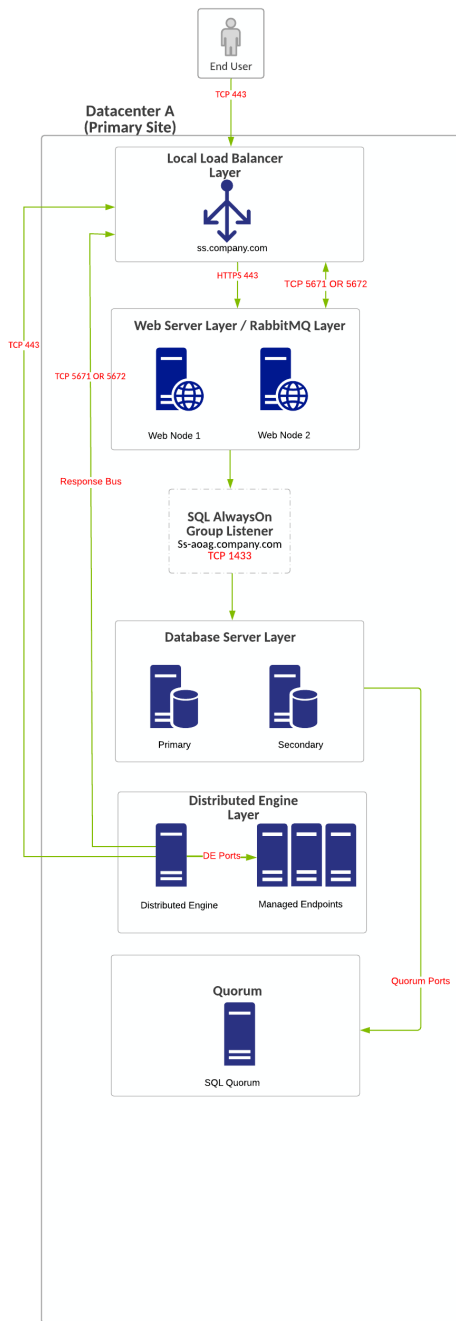
- ss.company.com: 443 (load balancer)
- ss.company.com: 5671 or 5672 (load balancer)
- ss-aoag.company.com: 1433 (created as part of SQL AlwaysOn configuration). Computer object or virtual IP.
- Windows Failover Cluster Object (created as part of Windows failover clustering configuration):
 - Computer object or virtual IP
 - One additional virtual IP address may be required as part of Windows failover cluster for single-site design for the network configuration of the failover cluster.

Diagram



The reference number for this diagram is A1.

Figure: Minimal Single-Site Enterprise Deployment



Average Single-Site Enterprise Deployment

Overview

- Single-site minimum-cost HA configuration.
- No shared storage requirement.
- RabbitMQ Helper Installed on separate, dedicated servers.

Delinea Architecture Reference Diagrams

- Single-site design with no native DR capacity. DR can be provided by VM replication if subnets are spanning locations, otherwise re-IP + DNS changes may be necessary.

Requirements

General

- SQL Standard Edition: Basic availability group configuration. Local load balancers can be used for all Web server nodes.
- Configuring a file share witness for SQL quorum voting is required for SQL to stay online during single node unplanned failures.
- [Distributed Engine Ports](#).
- [SQL Quorum Ports](#).

Virtual IP or Computer

- ss.company.com: 443 (load balancer)
- ss.company.com: 5671 or 5672 (load balancer)
- ss-aoag.company.com: 1433 (created as part of SQL AlwaysOn configuration). Computer object or virtual IP.
- Windows Failover Cluster Object (created as part of Windows failover clustering configuration):
 - Computer object or virtual IP
 - One additional virtual IP address may be required as part of Windows failover cluster for single-site design for the network configuration of the failover cluster.

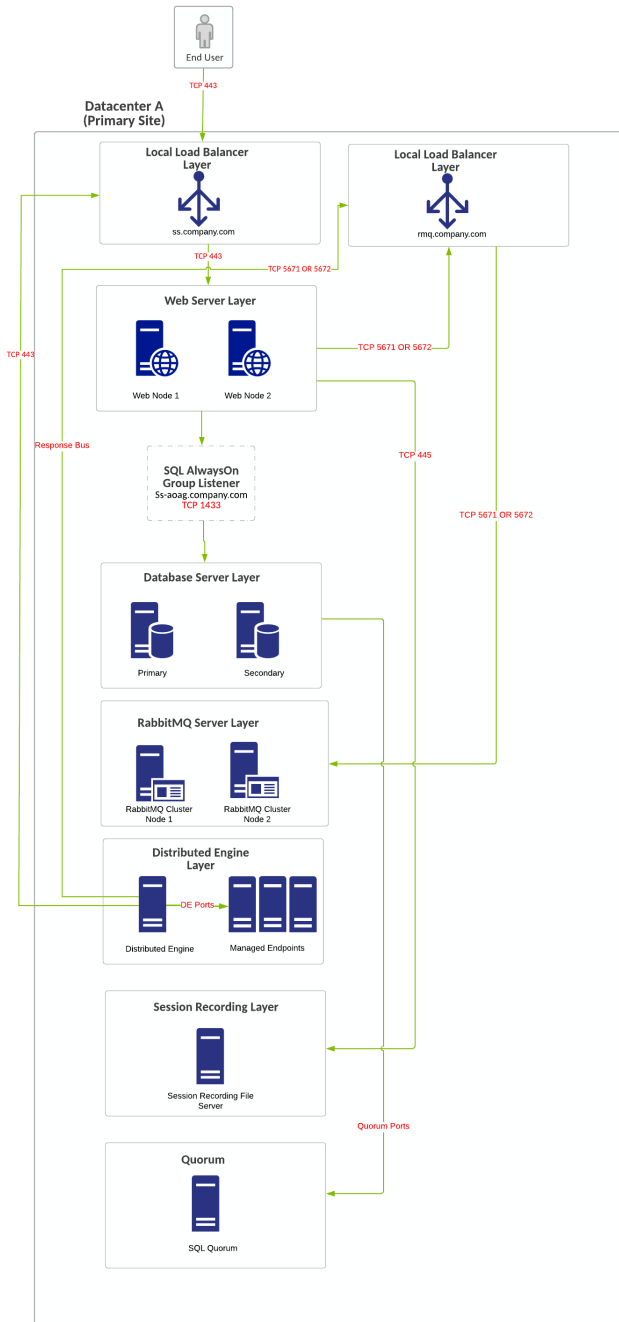
Diagram



The reference number for this diagram is A2.

Figure: Average Single-Site Enterprise Deployment

Delinea Architecture Reference Diagrams



Minimal Multi-Site Enterprise Deployment

Overview

- Minimum cost HA/DR configuration
- No shared storage requirement
- Lower infrastructure footprint for DR

Delinea Architecture Reference Diagrams

- DR site acts as temporary site only with no intention for long-term usage or becoming primary site.
- No secondary SQL node at primary site for "planned" failover.
- Secondary SQL node in DR site for planned or unplanned failover.
- Uses single RabbitMQ Helper site connector design.
- Can accommodate automatic failover with synchronous replication (30 ms or less latency between SQL DB nodes).
- Global load balancers for Web and RMQ configurations are configured to force all traffic to primary site unless primary site is down (priority group activation)

Requirements

General

- SQL Standard Edition (basic availability group configuration).
- If no global load balancers exist due to costs or infrastructure missing, local load balancers can be used for all Web server nodes, but DNS change may be required if primary location goes offline.
- Configuring a file share witness for SQL quorum voting is required for SQL to stay online during single-node unplanned failures. A cloud witness is recommended.
- [SQL Quorum Ports](#).

Virtual IP or Computer

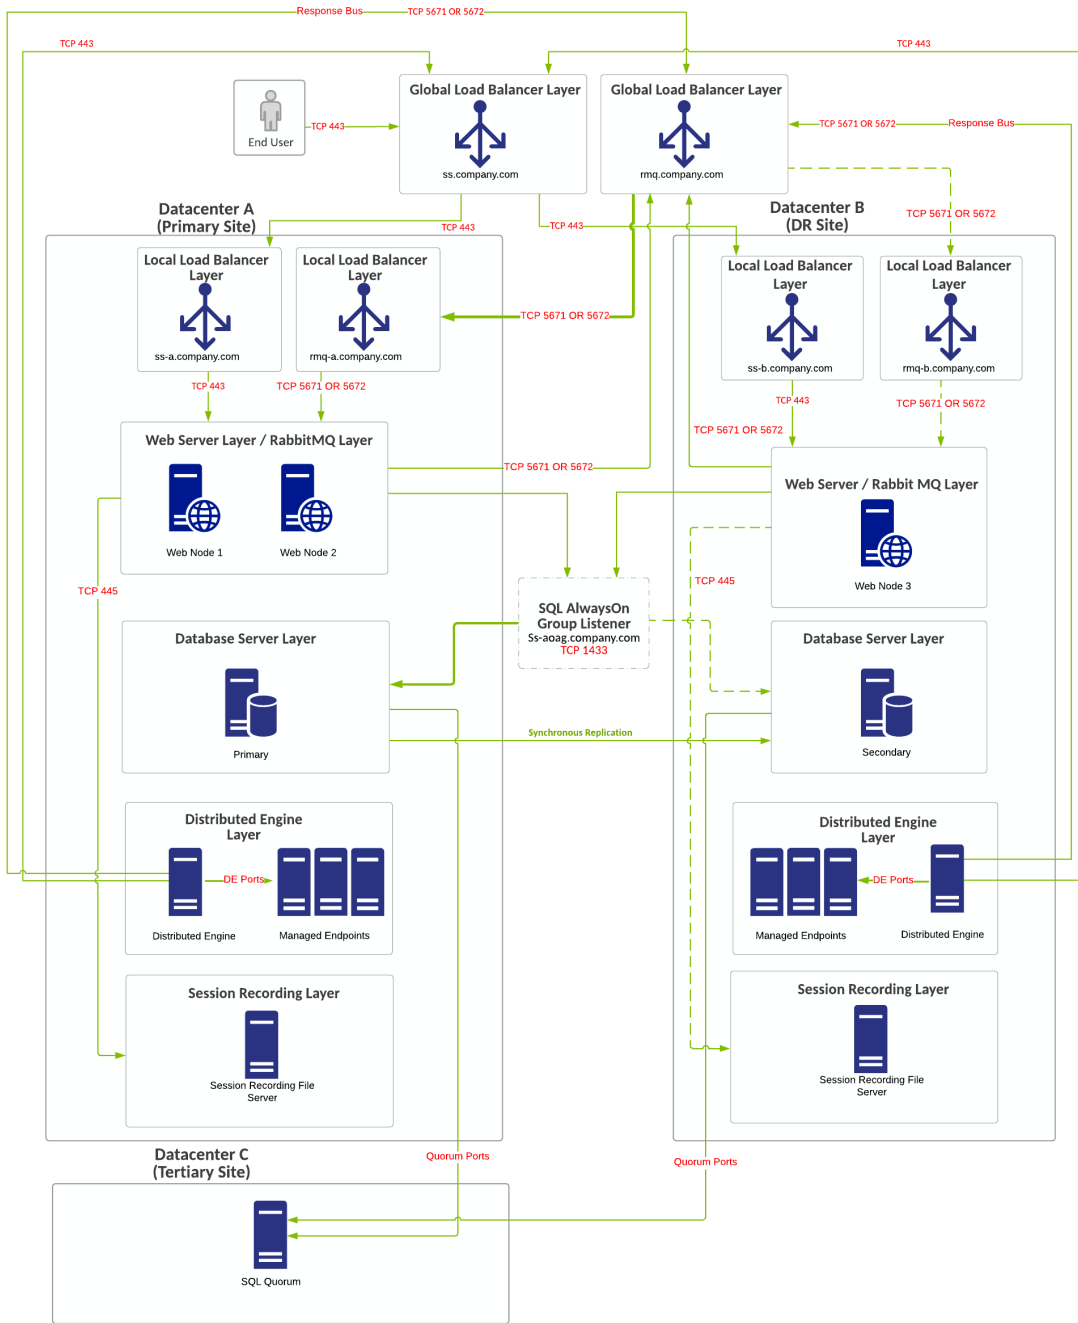
- ss.company.com: 443, rmq.company.com: 5671 or 5672 (2 virtual IPs, global load balancer)
- ss-a.company.com: 443, ss-b.company.com: 443 (2 virtual IPs, local load balancer)
- rmq-a.company.com: 5671 or 5672 (load balancer), rmq-b.company.com: 5671 or 5672 (2 virtual IPs, local load balancer)
- ss-aoag.company.com:1433 (created as part of SQL AlwaysOn configuration):
- ss-aoag.company.com computer object or virtual IP.
- Two virtual IP addresses may be required as part of this configuration .
- Windows failover cluster object (created as part of Windows failover clustering configuration):
 - Computer object or virtual IP.
 - Two additional virtual IP addresses may be required as part of Windows failover cluster for single site design for the network configuration of the failover cluster representing both networks at each respective site.

Diagram



The reference number for this diagram is B1.

Figure: Minimal Multi-Site Enterprise Deployment



Average Multi-Site Enterprise Deployment

Overview

- Minimum cost HA/DR configuration with no shared storage requirement, resulting in a smaller infrastructure footprint for DR.
- RabbitMQ Helper installed on SS Web servers (typically in a cluster on a primary+secondary node).

Delinea Architecture Reference Diagrams

- DR site acts at temporary site only with no intention for long-term usage or becoming primary site.
- Multiple site connector design intended for RabbitMQ Helper. All distributed engines communicate to one local load balancer/RMQ cluster for their response bus. There are two site connectors, one for each location.
- Global load balancers are unavailable, thus requiring a manual failover process/DNS change for Web traffic to the DR site. Additional application-specific changes will be needed for full functionality to resume in the DR site:
 - Change the internal site connector to the DR site connector.
 - Change the response bus to the DR site connector.
- Design assumes the Web server in the DR location is joined to the cluster and is online so that work is also being generated from Web servers in the DR location. If server is online and has roles configured, cross-data-center RabbitMQ Helper communication may occur. Work may be generated at either location's Web servers and is placed on either site connector, based on how secrets are configured.

Requirements

General

- SQL Enterprise Edition.
- DNS change may be required if primary location is offline.
- Configuring a file share witness for SQL quorum voting is required for SQL to stay online during single-node unplanned failures. A cloud witness is recommended.
- [Distributed Engine Ports](#).
- [SQL Quorum Ports](#).

Virtual IP or Computer

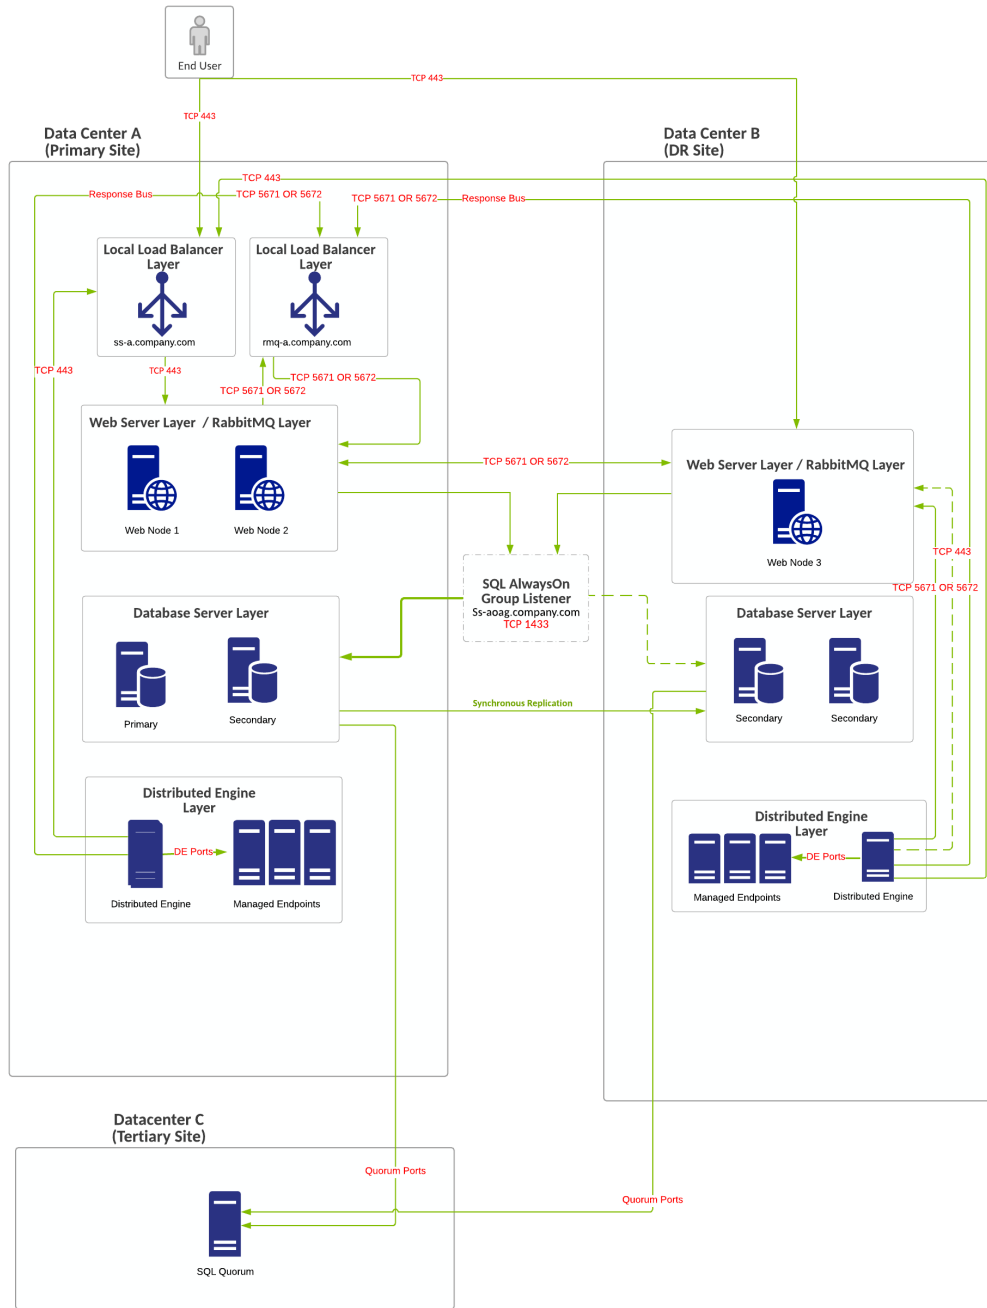
- ss.company.com: 443 (load balancer).
- rmq.company.com: 5671 or 5672 (load balancer).
- ss-aoag.company.com: 1433 (created as part of SQL AlwaysOn configuration). ss-aoag.company.com computer object/virtual IP.
- Windows failover cluster object (created as part of Windows failover clustering configuration):
 - Computer object/virtual IP.
 - One additional virtual IP addresses may be required as part of Windows failover cluster for single-site design for the network configuration of the failover cluster.

Diagram



The reference number for this diagram is B2.

Figure: Average Multi-Site Enterprise Deployment



Best Multi-Site Enterprise Deployment (C1)

Overview

- Higher Cost HA/DR Configuration with no shared storage requirement. Smaller Infrastructure footprint for DR RabbitMQ Helper clusters installed on dedicated systems.
- RabbitMQ Helper clusters installed on dedicated systems,

Delinea Architecture Reference Diagrams

- DR site acts as temporary site only when there is no intention of long-term use or becoming primary site. Services in DR site being down can incur downtime.
- Secondary SQL node at primary site for planned failover "patching." Secondary SQL Node in DR Site for unplanned failover.
- Uses a single RabbitMQ Helper site-connector design.
- Can accommodate automatic failover with synchronous replication (30 ms or less latency between SQL DB nodes).
- Global load balancers for Web and RMQ configurations are configured to force all traffic to go to primary site unless primary site is down (priority group activation).
- RabbitMQ Helper three-node cluster minimum for production location in keeping with RabbitMQ Helper design best practices. Classic mirrored queues will be retired in the future. Until then, two-node cluster designs may still be acceptable for customers who want less infrastructure footprint.

Requirements

General

- SQL Enterprise Edition.
- Global and local load balancers.
- Configuring a file share witness for SQL quorum voting is recommended. A cloud witness or DFSR share is recommended for witness configuration. Simultaneous failure of both SQL nodes in the primary location will still allow cluster to survive.
- [Distributed Engine Ports](#).
- [SQL Quorum Ports](#).

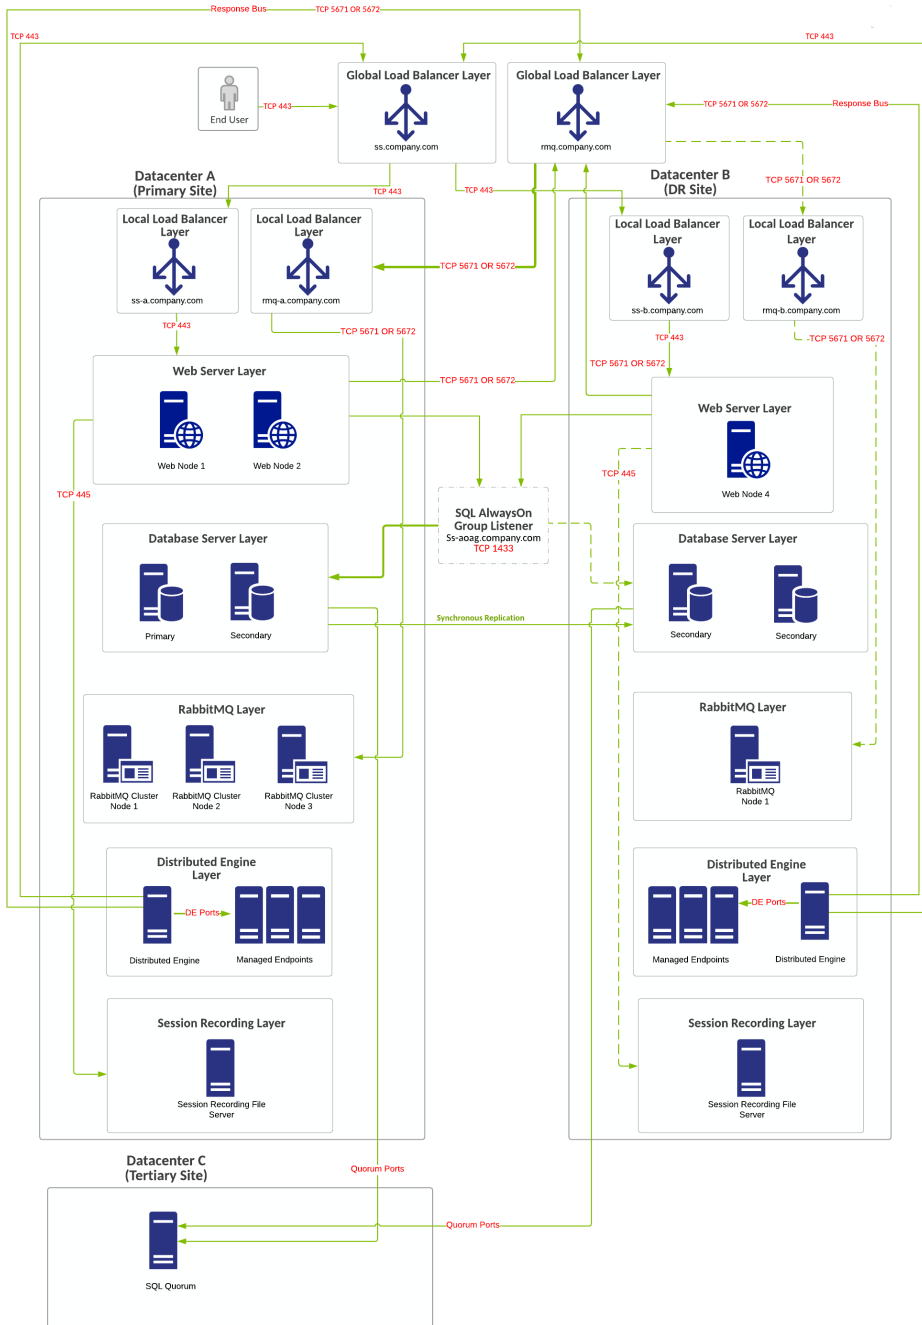
Virtual IP or Computer

- ss.company.com: 443, rmq.company.com: 5671 or 5672 (two virtual IPs, global load balancer).
- ss-a.company.com: 443, ss-b.company.com: 443 (two virtual IPs, local load balancer).
- rmq-a.company.com: 5671 or 5672 (Load Balancer), rmq-b.company.com: 5671 or 5672 (two virtual IPs, local load balancer).
- ss-aoag.company.com: 1433 (created as part of SQL AlwaysOn configuration):
 - ss-aoag.company.com computer object/virtual IP.
 - Two virtual IP addresses may be required as part of this configuration.
- Windows failover cluster object (created as part of Windows failover clustering configuration):
 - Computer object/virtual IP.
 - Two additional virtual IP addresses may be required as part of Windows failover cluster for single-site design for the network configuration of the failover cluster representing both networks at each site.

Diagram

 The reference number for this diagram is C1.

Figure: Best Multi-Site Enterprise Deployment (C1)



Best Multi-Site Enterprise Deployment (C2)

Overview

- Higher cost HA/DR configuration with no shared storage requirement. Smaller Infrastructure footprint for DR.
- RabbitMQ Helper clusters installed on dedicated systems.
- DR site acts as temporary site only when there is no intention of long-term use or becoming primary site. Services in DR site being down can incur downtime.
- Secondary SQL node at primary site for planned failover "patching." Secondary SQL Node in DR Site for unplanned failover.
- Can accommodate manual failover only with asynchronous replication (30 ms or more latency between SQL DB nodes).
- Singular site connector using only the local site. The local site comes with two distributed engines now. The distributed engine layer can be online/active but cross-data-center communication may occur (as depicted). It can alternatively be marked as yellow as other layers.
- Global load balancers are unavailable, thus requiring a manual failover process/DNS change for Web traffic to the DR site. Additional application-specific changes will be needed for full functionality to resume in DR site:
 - Change the internal site connector to the DR site connector.
 - Change the response bus to the DR site connector.
- Design assumes the Web server in DR location is joined to the cluster but is shut down (marked in yellow) so that no work is being generated from the Web servers in the DR location. If server is online and has roles configured, cross-data-center RabbitMQ Helper communication may occur. The RabbitMQ Helper layer is also marked as yellow, indicative of the server being built/online but the RabbitMQ Helper service is shut down.
- In this design, a manual failover may proceed in the following order:
 - Bring database node only and force it to become the new primary.
 - Bring online the Web server layer.
 - Update DNS record to point to the DR Web server node
 - Bring RabbitMQ Helper layer online for additional feature functionality. Switch the response bus, internal site connector, and custom URL if necessary.
- If necessary while in DR, session recordings should be temporarily switched to database storage while in DR.
- RabbitMQ Helper three-node cluster minimum for production location in keeping with with RabbitMQ Helper design best practices. Classic mirrored queues will be retired in the future. Until then, two-node cluster designs may still be acceptable for customers who want less infrastructure footprint.

Requirements

General

- SQL Enterprise Edition.
- DNS change may be required if primary location goes offline.

Delinea Architecture Reference Diagrams

- Configuring a file share witness for SQL quorum voting is recommended. A cloud witness or DFSR share is recommended for witness configuration. Simultaneous failure of both SQL nodes in the primary location will still allow cluster to survive.
- [Distributed Engine Ports](#).
- [SQL Quorum Ports](#).

Virtual IP or Computer

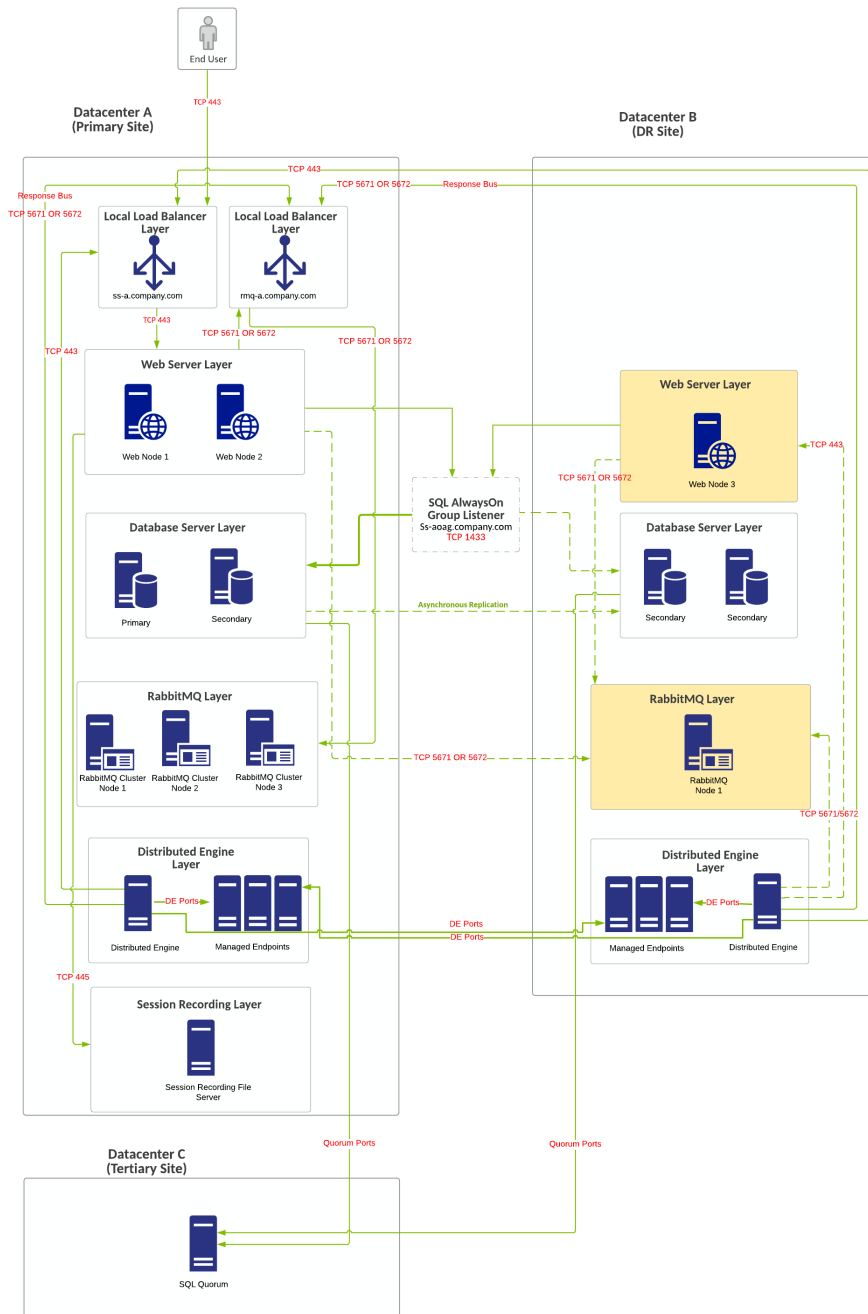
- ss.company.com: 443 (load balancer).
- rmq.company.com: 5671 or 5672 (load balancer).
- ss-aoag.company.com: 1433 (created as part of SQL AlwaysOn configuration). ss-aoag.company.com computer object/virtual IP.
- Windows failover cluster object (created as part of Windows failover clustering configuration):
 - Computer object/virtual IP.
 - One additional virtual IP addresses may be required as part of Windows failover cluster for single site design for the network configuration of the failover cluster.

Diagram



The reference number for this diagram is C2.

Figure: Best Multi-Site Enterprise Deployment (C2)



Best Multi-Site Enterprise Deployment (C3) with Manual Failover

Overview

- Higher cost HA/DR configuration with no shared storage requirement.
- Smaller Infrastructure footprint.
- RabbitMQ Helper cluster installed on dedicated systems.

Delinea Architecture Reference Diagrams

- DR site acts as temporary site only when there is no intention of long-term use or becoming the primary.
- Secondary SQL node at primary site for planned failover "patching." Secondary SQL Node in DR Site for unplanned failover.
- Can accommodate manual failover only with asynchronous replication (30 ms or more latency between SQL DB nodes).
- Singular site connector using only the local site. The local site comes with two distributed engines.
- The distributed engine layer can be online/active. Cross-data-center communication may occur (as depicted). It can alternatively be marked as yellow as with other layers.
- Design assumes the Web server in DR location is joined to the cluster but is shut down (marked in yellow), so no work is being generated from Web servers in the DR location. Other layers also have its services shut down. The database is the only layer not marked as yellow due to services needing to run for the asynchronous replica to function. The Local Load balancer configurations are simply disabled in the global load balancer configuration and can be manually brought online when traffic is ready to flow to DR.
- In this design, a manual failover may proceed in the following order:
 1. Bring database node only and force it to become the new primary.
 2. Bring online the Web server layer.
 3. Enable the load balancer configurations to direct traffic to the Web node.
 4. Bring distributed engine and RabbitMQ Helper layer online for additional feature functionality.
- Session Recordings, if necessary while in DR, should be temporarily switched to database storage while in DR.
- RabbitMQ Helper three-node cluster minimum for production location in keeping with RabbitMQ Helper design best practices.
- Classic mirrored queues will be retired in the future. Until then, two-node cluster designs may still be acceptable for customers who want less infrastructure footprint.

Requirements

General

- SQL Enterprise Edition.
- Global and local load balancers.
- Configuring a file share witness for SQL quorum voting is recommended. A cloud witness or DFSR share is recommended for witness configuration. Simultaneous failure of both SQL nodes in the primary location will still allow the cluster to survive.
- [Distributed Engine Ports](#).
- [SQL Quorum Ports](#).

Virtual IP or Computer

- ss.company.com: 443, rmq.company.com: 5671 or 5672 (two virtual IPs, global load balancer).
- ss-a.company.com: 443, ss-b.company.com: 443 (two virtual IPs, local load balancer).

Delinea Architecture Reference Diagrams

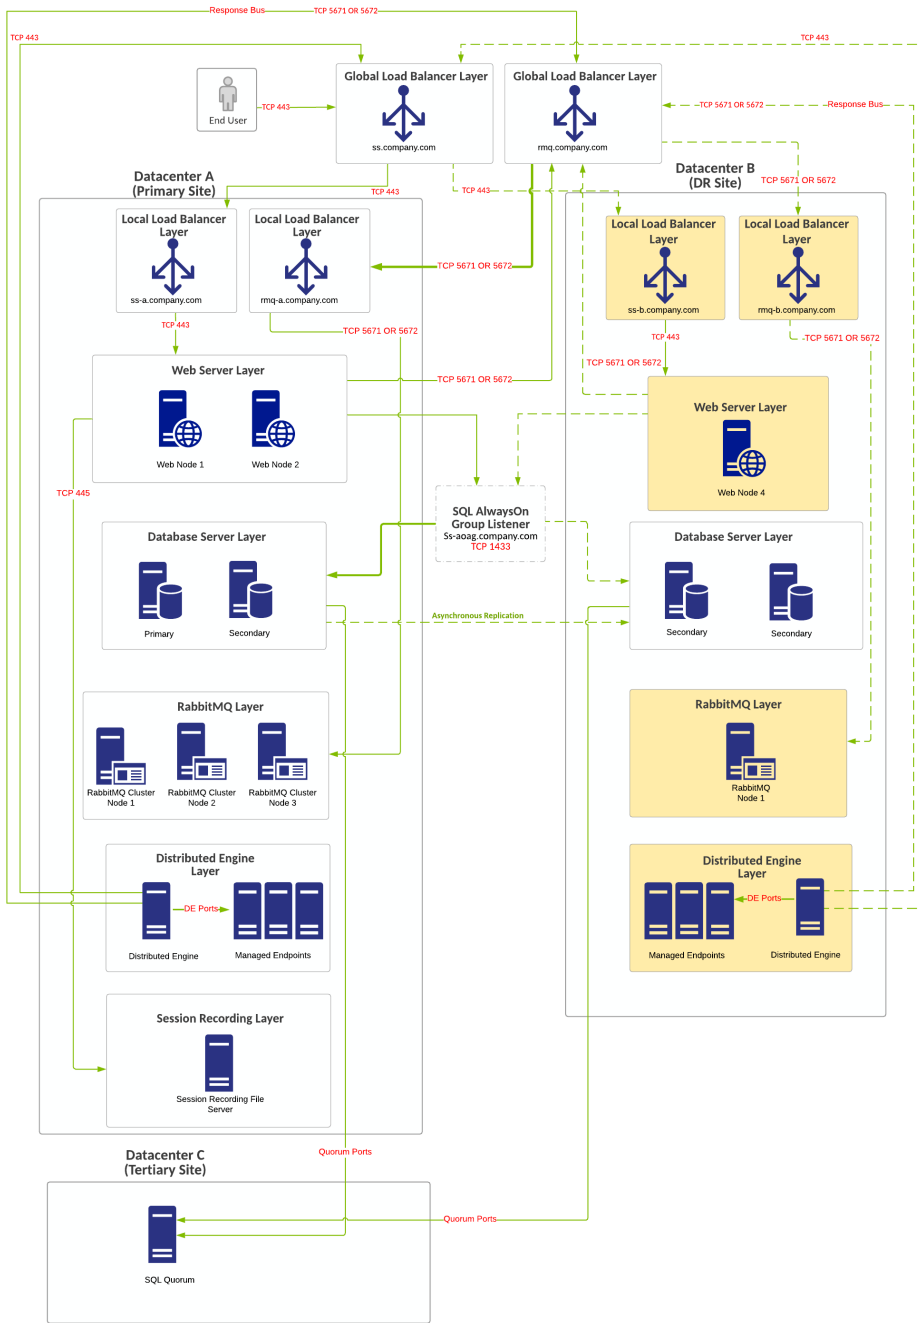
- rmq-a.company.com: 5671 or 5672 (Load Balancer), rmq-b.company.com: 5671 or 5672 (two virtual IPs, local load balancer).
- ss-aoag.company.com: 1433 (created as part of SQL AlwaysOn configuration):
 - ss-aoag.company.com computer object/virtual IP.
 - Two virtual IP addresses may be required as part of this configuration.
- Windows failover cluster object (created as part of Windows failover clustering configuration):
 - Computer object/virtual IP.
 - Two additional virtual IP addresses may be required as part of Windows failover cluster for single-site design for the network configuration of the failover cluster representing both networks at each site.

Diagram



The reference number for this diagram is C3.

Figure: Best Multi-Site Enterprise Deployment with Manual Failover (C3)



Best Multi-Site Enterprise Deployment (C4) with Automatic Failover

Overview

- Higher cost HA/DR configuration with no shared storage requirement.
- Smaller Infrastructure footprint.
- RabbitMQ Helper cluster installed on dedicated systems.

Delinea Architecture Reference Diagrams

- Equal infrastructure at two sites—DR site can act as permanent secondary site for long-term use.
- Secondary SQL node at primary site for planned failover "patching." Secondary SQL Node in DR Site for unplanned failover.
- Uses a single RabbitMQ Helper connector design.
- Global load balancers for Web and RMQ configurations are configured to force all traffic to go to primary site unless primary site is down (priority group activation).
- RabbitMQ Helper three-node cluster minimum for production location in keeping with RabbitMQ Helper design best practices.
- Classic mirrored queues will be retired in the future. Until then, two-node cluster designs may still be acceptable for customers who want less infrastructure footprint.
- Multiple session recording servers in each location should consider leveraging DFSR or other methods to provide near real-time synchronization of data storage between data centers for active/active designs. This can impact design cost.

Requirements

General

- SQL Enterprise Edition.
- Global and local load balancers.
- Configuring a file share witness for SQL quorum voting is recommended. A cloud witness or DFSR share is recommended for witness configuration. Simultaneous failure of both SQL nodes in the primary location will still allow the cluster to survive.
- [Distributed Engine Ports](#).
- [SQL Quorum Ports](#).

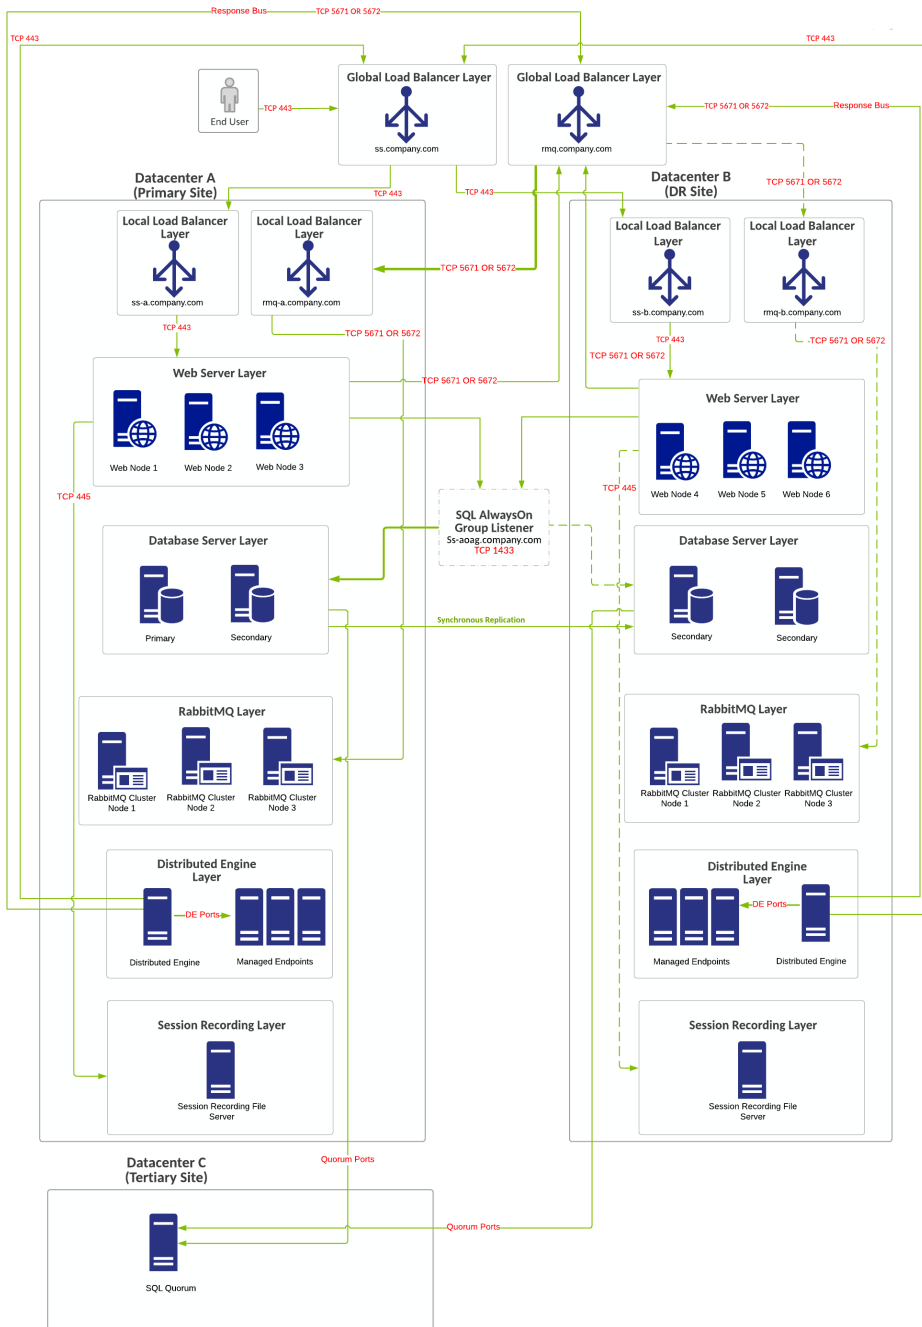
Virtual IP or Computer

- ss.company.com: 443, rmq.company.com: 5671 or 5672 (two virtual IPs, global load balancer).
- ss-a.company.com: 443, ss-b.company.com: 443 (two virtual IPs, local load balancer).
- rmq-a.company.com: 5671 or 5672 (Load Balancer), rmq-b.company.com: 5671 or 5672 (two virtual IPs, local load balancer).
- ss-aoag.company.com: 1433 (created as part of SQL AlwaysOn configuration):
 - ss-aoag.company.com computer object/virtual IP.
 - Two virtual IP addresses may be required as part of this configuration.
- Windows failover cluster object (created as part of Windows failover clustering configuration):
 - Computer object/virtual IP.
 - Two additional virtual IP addresses may be required as part of Windows failover cluster for single-site design for the network configuration of the failover cluster representing both networks at each site.

Diagram

 The reference number for this diagram is C4.

Figure: Best Multi-Site Enterprise Deployment with Automatic Failover (C4)



Secret Server and DevOps Secrets Vault Example Architectures



If you are a current customer with support hours for Delinea Professional Services, you can discuss any of these diagrams in detail with one of our Professional Services Solutions Architects.

This reference architecture is our best practices for using Delinea Secret Server (SS) with DevOps Secrets Vault (DSV). You can use DSV with Secret Server On-Premises or Cloud. DevOps Secrets Vault (DSV) is nimble, multi-platform privileged access manager specifically for developers and other IT professionals. So why might you want to integrate it with SS? Customers can use DSV as a native "backup" of critical SS secrets, effectively extending their on-premise instance into the cloud during SS outages. Similarly Secret Server Cloud (SSC) customers can back up their critical secrets to another cloud, effectively extending their SSC instance into another vendor's cloud during SS outages. SSC is hosted in Azure, and DSV is hosted in Amazon Web Services, which has a 99.999% uptime.



We suggest only doing this for critical secrets and not every secret in your environment. Back up critical IT infrastructure secrets, such as those needed if your data centers went down—routers, firewalls, and the like.

We discuss the following scenarios:

- SS on-premises with a single DSV instance (A-1)
- SSC with a single DSV instance (A-2)
- SS on-premises with a single DSV instance in a separate region (B-1)
- SS on-premises with multiple DSV instances in the same or a separate region (C-1)

Secret Server On-Premises with a Single DSV Instance

Overview

- Customer is using an on-premise installation of SS installed in physical or private cloud environment and is leveraging a free instance of DSV (Limited to 250 secrets, 2,500 API calls per month).
- When SS is down, user-to-jump-host connectivity becomes active and uses a break-the-glass account to connect to the jump host. The jump host then has the DSV executable available and can retrieve credentials when SS is down.
- DSV SLA is 99.999%.

Requirements

- SS Premium/Professional/Platinum licensing and DSV Free Edition.
- Each tenant in the SS instance has a settable interval to check if secrets need to be pushed DSV again. Once that interval hits (or they use an event pipeline to trigger it), SS checks all the secrets associated to that tenant. Any secret that has had a change since the last time that tenant had the secret pushed to it is pooled. SS authenticates and for each secret it POSTs or PUTs the secret (depending on creation or update).

Delinea Architecture Reference Diagrams

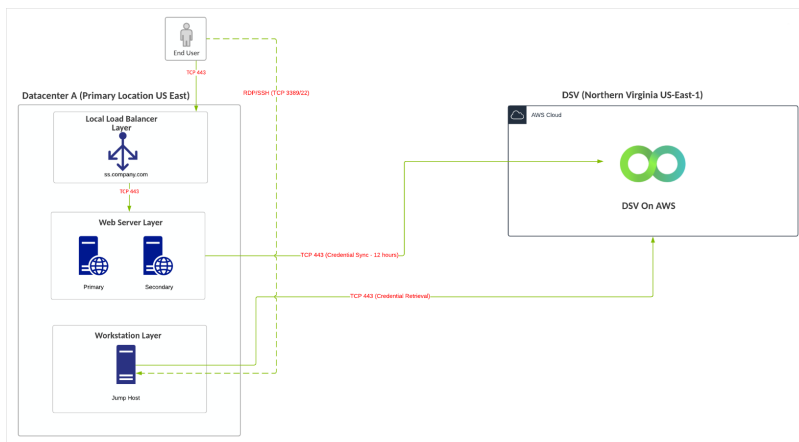
- The number of requests equals the number of updated tenants (authenticating for each tenant) plus the number of secrets that need updating updated per tenant (updating the secret in DSV).
- Compare the total 2,500 free API calls per month to your number of requests to determine if this fits the "free" DSV licensing model.

Diagram



The reference for this diagram is A-1.

Figure: Secret Server On-Premises with a Single DSV Instance



Secret Server Cloud with a Single DSV Instance

Overview

- Customer using an on-premise installation of SS installed in physical or private cloud environment and is leveraging a free instance of DSV (Limited to 250 secrets, 2,500 API calls per month).
- When SS is down, the DSV executable is available and can retrieve credentials when SS is down. This is more convenient but less secure than other options.
- DSV SLA is 99.999%.
- SSC SLA is 99.9%.

Requirements

- SS Premium/Professional/Platinum licensing and DSV Free Edition.
- Each tenant in the SS instance has a settable interval to check if secrets need to be pushed to DSV again. Once that interval hits (or they use an event pipeline to trigger it), SS checks all the secrets associated to that tenant. Any secret that has had a change since the last time that tenant had the secret pushed to it is pooled. SS authenticates and for each secret it POSTs or PUTs the secret (depending on creation or update).
- The number of requests equals the number of updated tenants (authenticating for each tenant) plus the number of secrets that need updating updated per tenant (updating the secret in DSV).

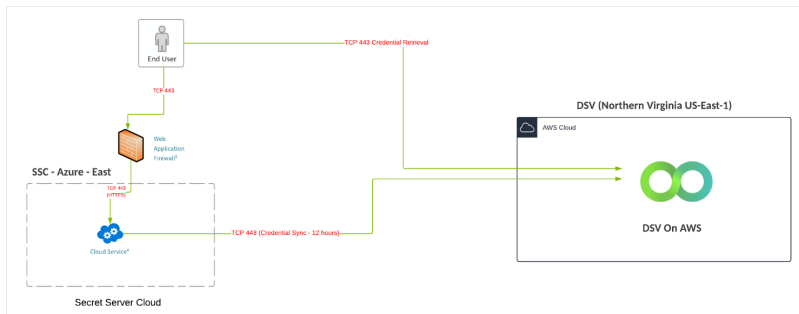
Delinea Architecture Reference Diagrams

- Compare the total 2,500 free API calls per month to your number of requests to determine if this fits the "free" DSV licensing model.

Diagram

 **Note:** The reference for this diagram is A-2.

Figure: Secret Server Cloud with a Single DSV Instance



Secret Server On-Premises with a Single DSV Instance in a Separate Region

Overview

- Customer is using an on-premise installation of SS installed in a physical or private cloud environment and is using a free instance of DSV (Limited to 250 secrets, 2,500 API calls per month).
- When SS is down, user-to-jump-host connectivity becomes active and uses a break-the-glass account to connect to the jump host. The jump host then has the DSV executable available and can retrieve credentials when SS is down.
- Multiple jump hosts are provisioned in case the primary site is down.
- DSV SLA is 99.999%.

Requirements

- SS Premium/Professional/Platinum licensing and DSV Free Edition.
- Each tenant in the SS instance has a settable interval to check if secrets need to be pushed to DSV again. Once that interval hits (or they use an event pipeline to trigger it), SS checks all the secrets associated to that tenant. Any secret that has had a change since the last time that tenant had the secret pushed to it is pooled. SS authenticates and for each secret it POSTs or PUTs the secret (depending on creation or update).
- The number of requests equals the number of updated tenants (authenticating for each tenant) plus the number of secrets that need updating updated per tenant (updating the secret in DSV).
- Compare the total 2,500 free API calls per month to your number of requests to determine if this fits the "free" DSV licensing model.

Diagram


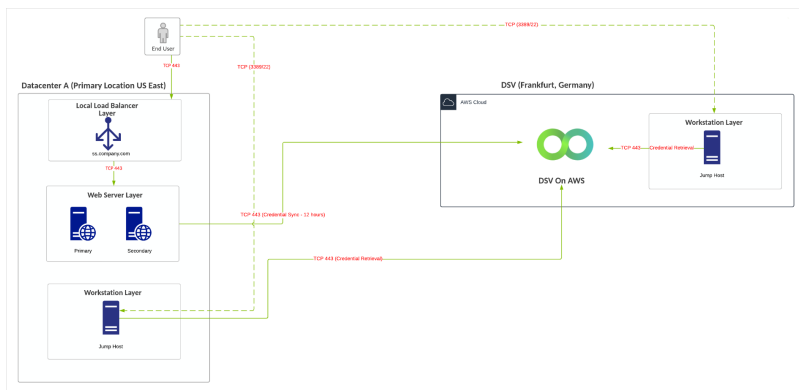
 The reference for this diagram is B-1.

Figure: Secret Server On-Premises with a Single DSV Instance in a Separate Region



Secret Server On-Premises with a Multiple DSV Instances in the Same or a Separate Region

Overview

- Customer is using an on-premise installation of SS installed in a physical or private cloud environment and is using a free instance of DSV (Limited to 250 secrets, 2,500 API calls per month).
- When SS is down, user-to-jump-host connectivity becomes active and uses a break-the-glass account to connect to the jump host. The jump host then has the DSV executable available and can retrieve credentials when SS is down.
- Multiple jump hosts are provisioned in case the primary site is down.
- Users from different regions can access their own regional DSV instance.
- DSV SLA is 99.999%.

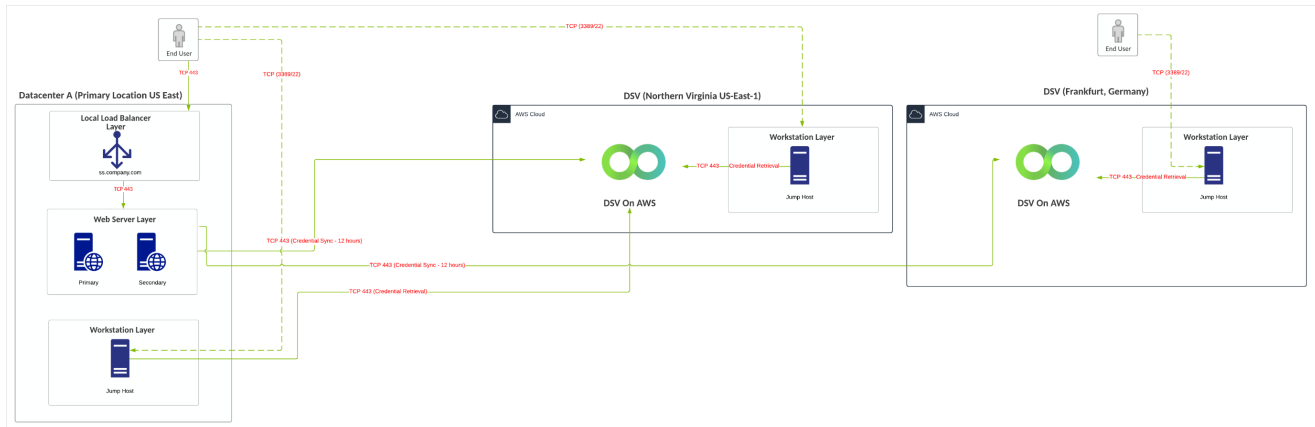
Requirements

- SS Premium/Professional/Platinum licensing and DSV Free Edition.
- Each tenant in the SS instance has a settable interval to check if secrets need to be pushed to DSV again. Once that interval hits (or they use an event pipeline to trigger it), SS checks all the secrets associated to that tenant. Any secret that has had a change since the last time that tenant had the secret pushed to it is pooled. SS authenticates and for each secret it POSTs or PUTs the secret (depending on creation or update).
- The number of requests equals the number of updated tenants (authenticating for each tenant) plus the number of secrets that need updating updated per tenant (updating the secret in DSV).
- Each tenant in the SS instance has a settable interval to check if secrets need to be pushed to DSV again. Once that interval hits (or they use an event pipeline to trigger it), SS checks all the secrets associated to that tenant. Any secret that has had a change since the last time that tenant had the secret pushed to it is pooled. SS authenticates and for each secret it POSTs or PUTs the secret (depending on creation or update).
- In this model, the DSV instance is provisioned in two regions. Some secrets synchronize to one region while other secrets synchronize to another. This may be ideal for large global deployments.


Diagram

 The reference for this diagram is C-1.

Figure: Secret Server On-Premises with a Multiple DSV Instances in the Same or a Separate Region



Secret Server and Privilege Manager Example Architectures


 If you are a current customer with support hours for Thycotic Professional Services, you can discuss any of these diagrams in detail with one of our Professional Services Solutions Architects.

Secret Server and Privilege Manager Integration

The benefits of Privilege Manager (PM) integration with Secret Server (SS) include:

- SS can be the authentication source for PM. This:
 - Adds SS MFA login options to PM log ons.
 - Gives one place for role assignments for both products.
- PM can use SS as a storage container.

When you use SS as the authentication source for PM, role permissions assigned in SS apply and determine user-access levels in PM.

 Without SS integration, the authentication sources for PM can be NTLM for the Web server or Azure AD.

When using SS as a storage container for PM credentials:

- PM creates secrets for each local credential managed by PM.
- PM creates secrets for each configuration credential stored in PM. This includes the credentials PM uses for foreign system integration, such as AD sync and ServiceNow.

Delinea Architecture Reference Diagrams

- PM pulls any changes from secrets. PM only stores the credentials in SS to use SS workflow options and for users to view them.

This integration is supported when the two applications are installed on the same server or separate servers, as long as PM can communicate with SS via the SS REST API.

Single Site with Minimum HA

Overview

- Minimum-cost configuration with no shared storage requirement.
- RabbitMQ Helper (for SS) is installed on the SS Web servers (typically in a cluster).
- Single-site design with no native DR capacity. DR can be provided by VM replication if subnets are spanning locations. Otherwise Re-IP + DNS changes may be necessary.
- PM is installed on separate Web servers.
- PM can integrate with SS for authentication and credential storage.
- PM can reside on the same database servers as SS or on separate ones, but SS and PM should not share the same database.



Due to SQL basic availability groups with the Standard Edition, you need to have multiple SQL instances and a separate AlwaysOn availability group configuration.

Requirements

- SQL Standard Edition with a basic availability group configuration.
- You can use local load balancers for all Web server nodes.
- For SQL to stay online during single-node unplanned failures, you must configure a file-share witness for SQL quorum voting.

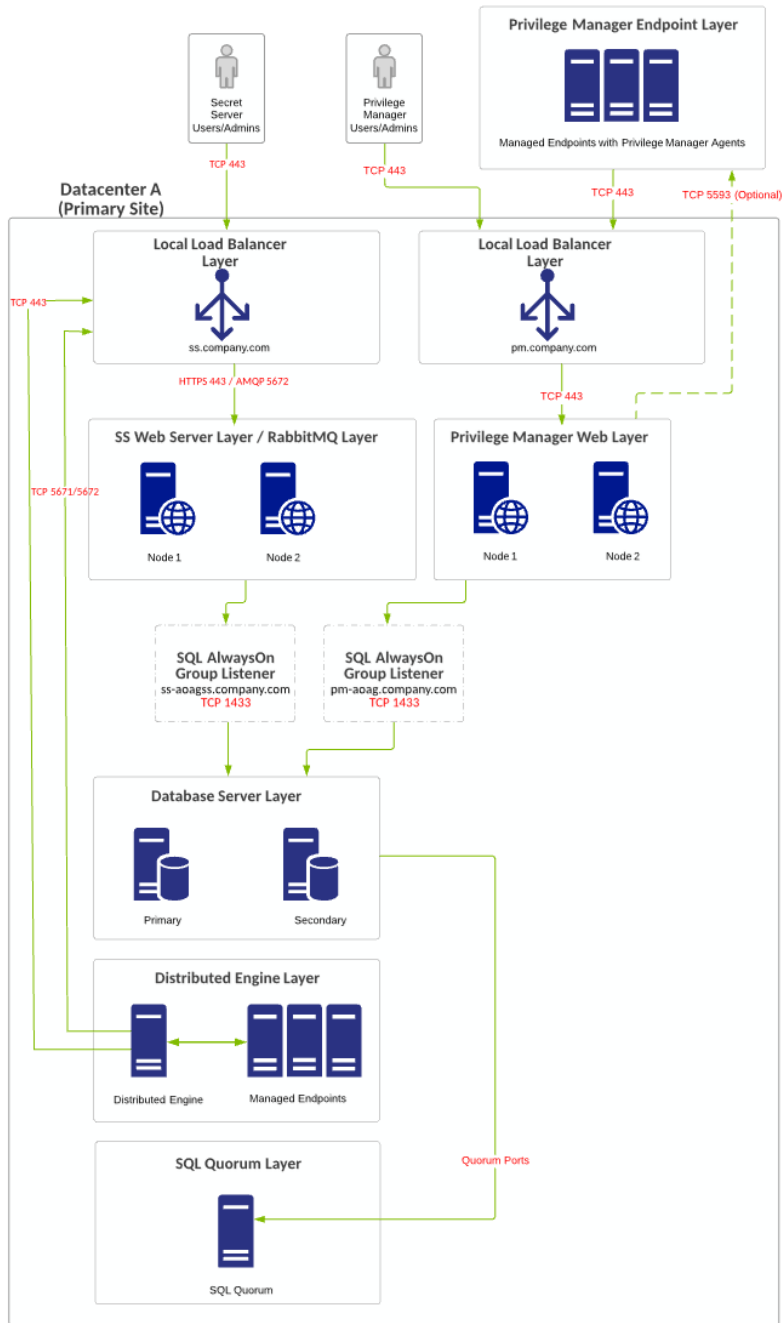
Diagram



The reference for this diagram is A-1.

Figure: Single Site with Minimum HA

Delinea Architecture Reference Diagrams



Single Site with Minimum HA and Separate RabbitMQ Helper

Overview

- Minimum-cost HA configuration with no shared storage requirement.
- RabbitMQ Helper (for SS) is installed on the SS Web servers (typically in a cluster).

Delinea Architecture Reference Diagrams

- Single-site design with no native DR capacity. DR can be provided by VM replication if subnets are spanning locations. Otherwise Re-IP + DNS changes may be necessary.
- PM is installed on separate Web servers.
- PM can integrate with SS for authentication and credential storage.
- PM can reside on the same database servers as SS or on separate ones, but SS and PM should not share the same database.



Due to SQL basic availability groups with the Standard Edition, you need to have multiple SQL instances and a separate AlwaysOn availability group configuration.

- You can use a separate Web reverse proxy or Azure service bus configuration for Privilege Manager agent TCP 443 communication.

Requirements

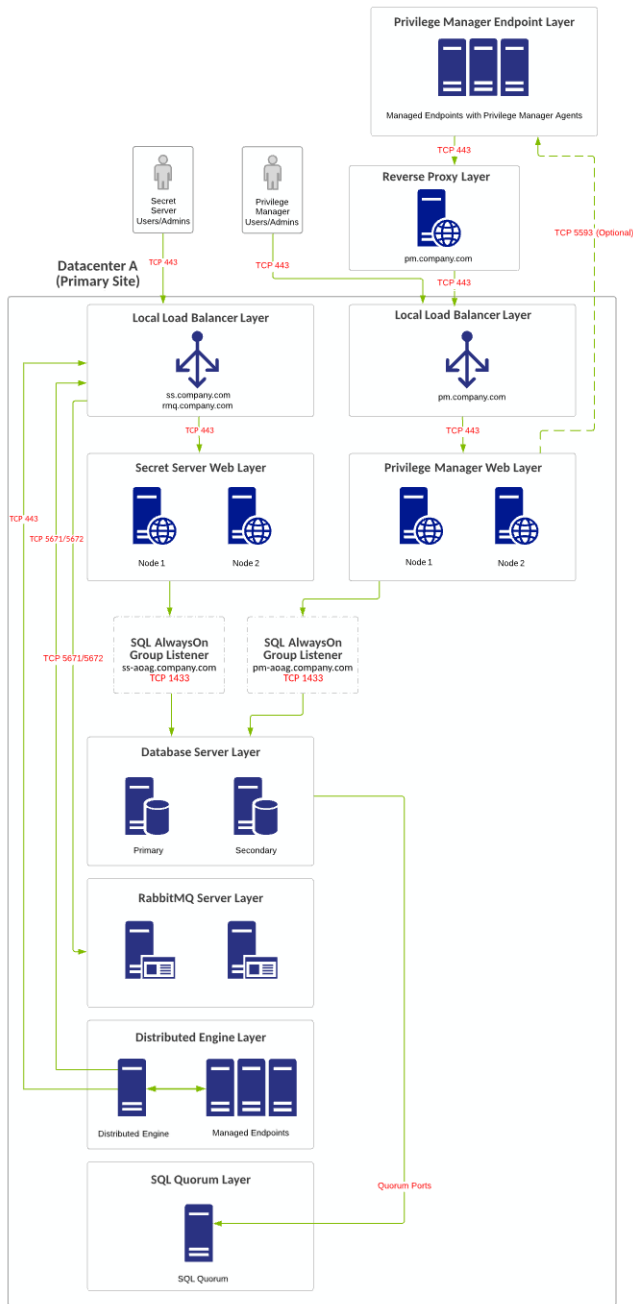
- SQL Standard Edition with a basic availability group configuration.
- You can use local load balancers for all Web server nodes.
- For SQL to stay online during single-node unplanned failures, you must configure a file-share witness for SQL quorum voting.

Diagram



The reference for this diagram is A-2.

Figure: Single Site with Minimum HA and Separate RabbitMQ Helper



Multiple Site with Manual Failover

Overview

- Minimum-cost HA configuration with no shared storage requirement.
- RabbitMQ Helper (for SS) is installed on the SS Web servers (typically in a cluster).

Delinea Architecture Reference Diagrams

- SQL AlwaysOn configurations are either synchronous or asynchronous for the SS database and asynchronous only for the PM database.
- DR site acts as a temporary site only with no long-term use. Services in DR site being down can incur downtime.
- PM is installed on separate Web servers.
- PM can integrate with SS for authentication and credential storage.
- PM can reside on the same database servers as SS or on separate ones, but SS and PM should not share the same database.



Due to SQL basic availability groups with the Standard Edition, you need to have multiple SQL instances and a separate AlwaysOn availability group configuration.

- You can use a separate Web reverse proxy or Azure service bus configuration for Privilege Manager agent TCP 443 communication.

Requirements

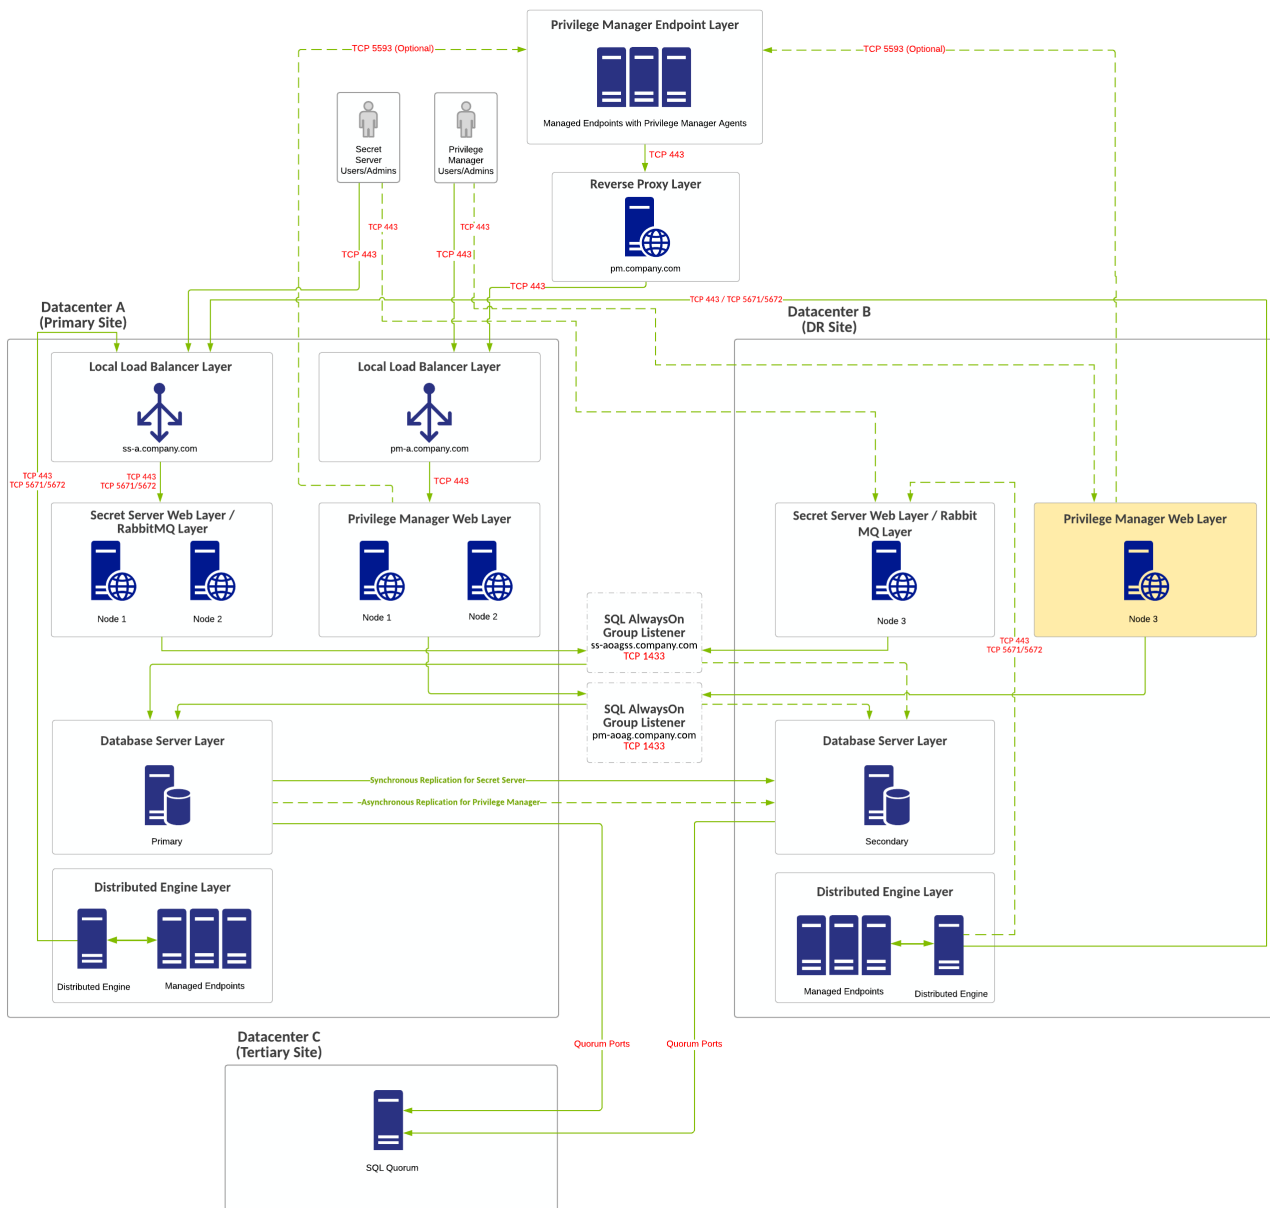
- SQL Standard Edition with a basic availability group configuration.
- If global load balancers are not available due to cost or limited infrastructure, you can use local load balancers for all Web server nodes, but DNS change may be required if primary location goes offline.
- For SQL to stay online during single-node unplanned failures, you must configure a file-share witness for SQL quorum voting.

Diagram



The reference for this diagram is B-1.

Figure: Multiple Site with Manual Failover



Multiple Site with Manual Failover and Separate RabbitMQ Helper

Overview

- Minimum-cost HA configuration with no shared storage requirement.
- RabbitMQ Helper (for SS) is installed on the SS Web servers (typically in a cluster).
- SQL AlwaysOn configurations are either synchronous or asynchronous for the SS database and asynchronous only for the PM database.
- DR site acts as a temporary site only with no long-term use. Services in DR site being down can incur downtime.

Delinea Architecture Reference Diagrams

- PM is installed on separate Web servers.
- PM can integrate with SS for authentication and credential storage.
- PM can reside on the same database servers as SS or on separate ones, but SS and PM should not share the same database.



Due to SQL basic availability groups with the Standard Edition, you need to have multiple SQL instances and a separate AlwaysOn availability group configuration.

- You can use a separate Web reverse proxy or Azure service bus configuration for Privilege Manager agent TCP 443 communication.

Requirements

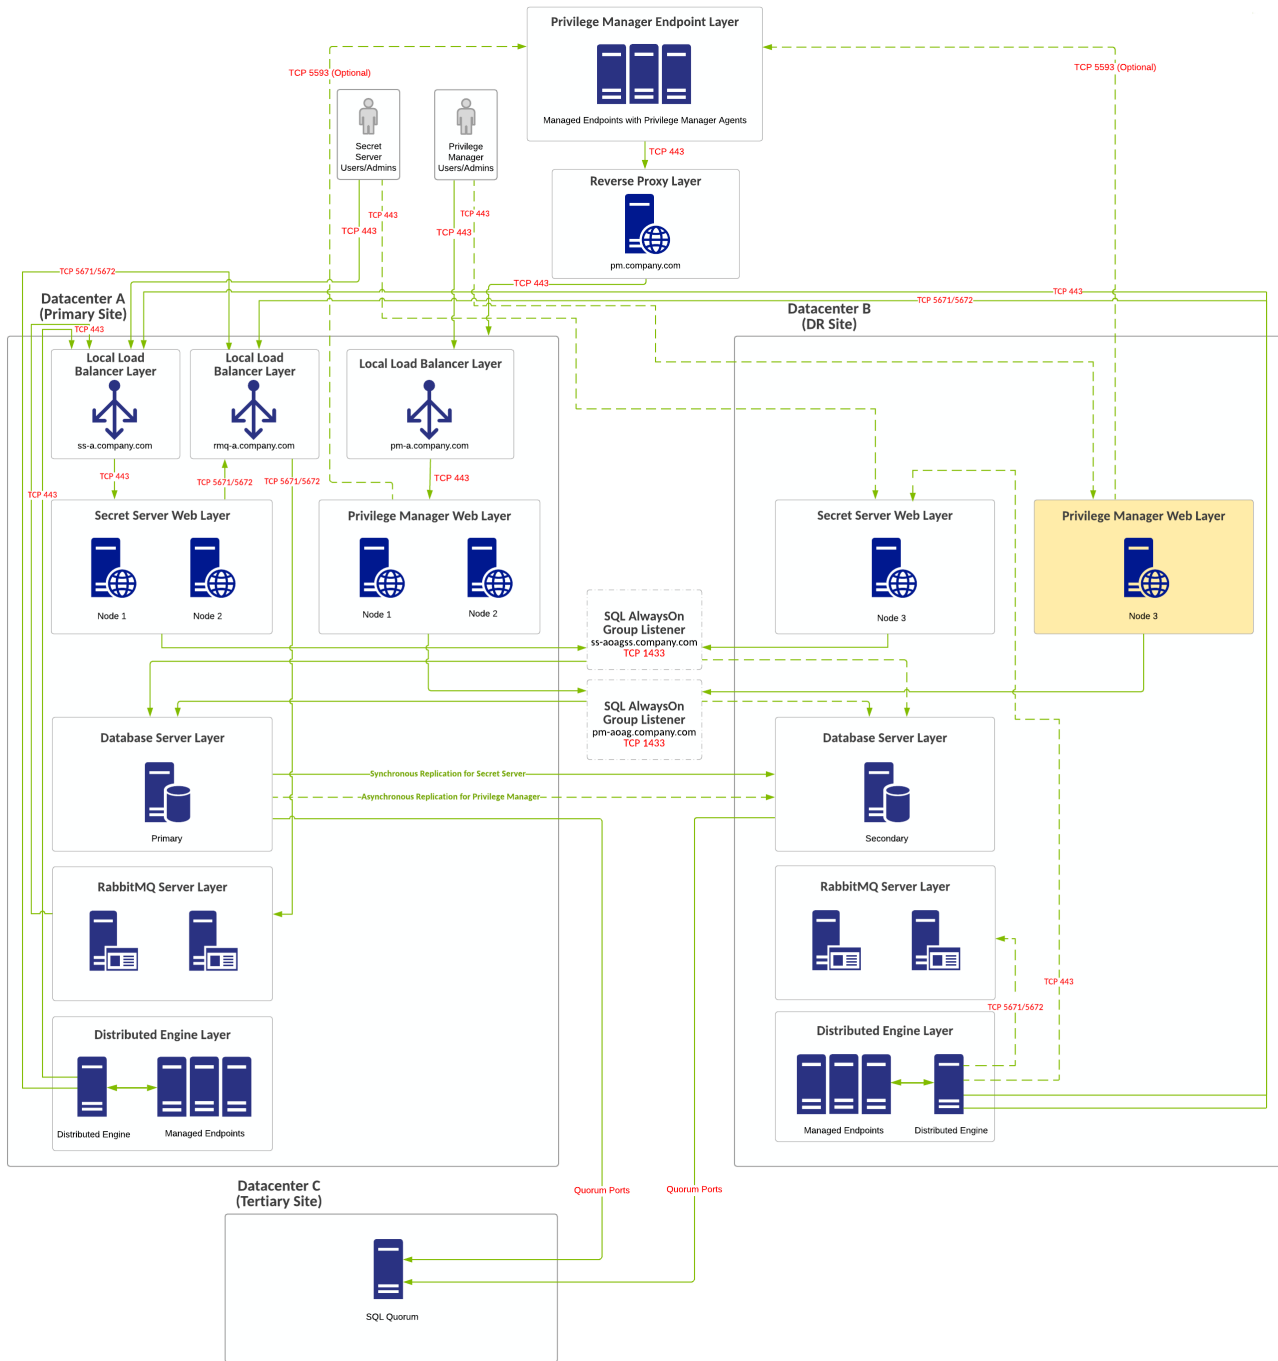
- SQL Standard Edition with a basic availability group configuration.
- If global load balancers are not available due to cost or limited infrastructure, you can use local load balancers for all Web server nodes, but DNS change may be required if primary location goes offline.
- For SQL to stay online during single-node unplanned failures, you must configure a file-share witness for SQL quorum voting. We recommend a cloud witness.

Diagram



The reference for this diagram is B-2.

Figure: Multiple Site with Manual Failover and Separate RabbitMQ Helper



Multiple Site with Automatic Failover

Overview

- Improved HA configuration with no shared storage requirement.
- RabbitMQ Helper (for SS) is installed on the SS Web servers (typically in a cluster).

Delinea Architecture Reference Diagrams

- SQL AlwaysOn configurations are either synchronous or asynchronous for the SS database and asynchronous only for the PM database.
- DR site acts as a temporary site only with no long-term use. Services in DR site being down can incur downtime.
- PM is installed on separate Web servers.
- PM can integrate with SS for authentication and credential storage.
- PM can reside on the same database servers as SS or on separate ones, but SS and PM should not share the same database.



Due to SQL basic availability groups with the Standard Edition, you need to have multiple SQL instances and a separate AlwaysOn availability group configuration.

- You can use a separate Web reverse proxy or Azure service bus configuration for Privilege Manager agent TCP 443 communication.

Requirements

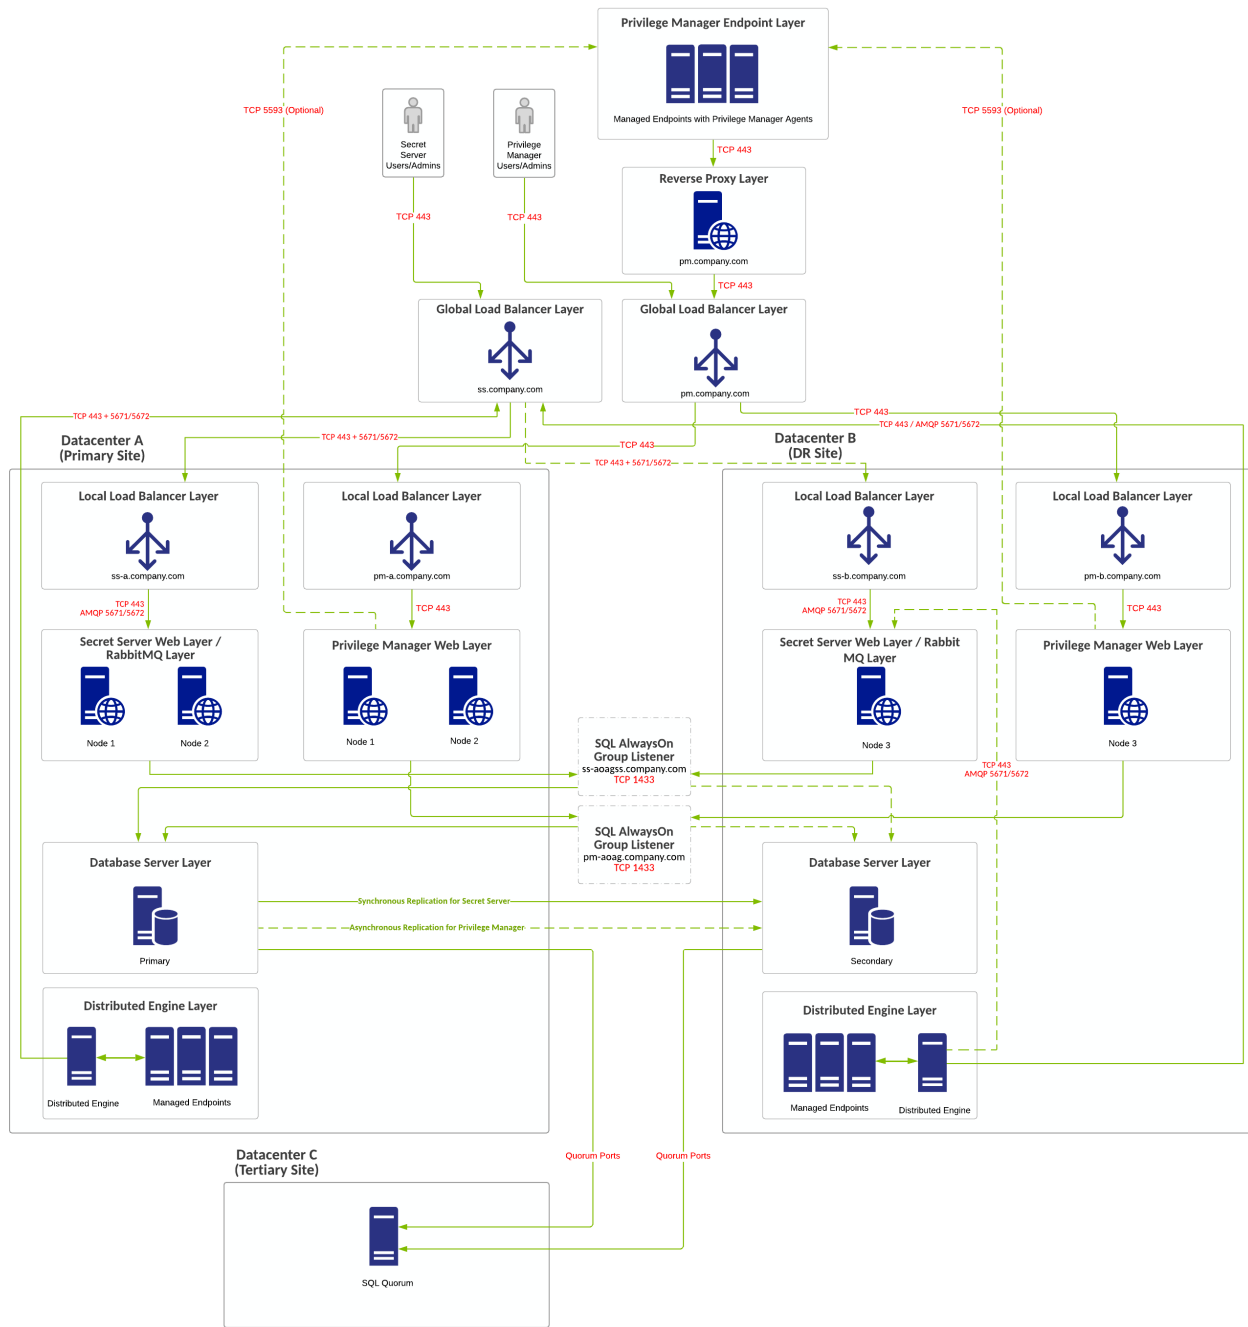
- SQL Standard Edition with a basic availability group configuration.
- If global load balancers are not available due to cost or limited infrastructure, you can use local load balancers for all Web server nodes, but DNS change may be required if primary location goes offline.
- For SQL to stay online during single-node unplanned failures, you must configure a file-share witness for SQL quorum voting. We recommend a cloud witness.

Diagram



The reference for this diagram is C-1.

Figure: Multiple Site with Automatic Failover



Multiple Site with Automatic Failover and Separate RabbitMQ Helper

Overview

- Improved HA configuration with no shared storage requirement.
- RabbitMQ Helper (for SS) is installed on dedicated servers (typically in a cluster).

Delinea Architecture Reference Diagrams

- SQL AlwaysOn configurations are either synchronous or asynchronous for the SS database and asynchronous only for the PM database.
- DR site acts as a temporary site only with no long-term use. Services in DR site being down can incur downtime.
- PM is installed on separate Web servers.
- PM can integrate with SS for authentication and credential storage.
- PM can reside on the same database servers as SS or on separate ones, but SS and PM should not share the same database.



Due to SQL basic availability groups with the Standard Edition, you need to have multiple SQL instances and a separate AlwaysOn availability group configuration.

- You can use a separate Web reverse proxy or Azure service bus configuration for Privilege Manager agent TCP 443 communication.

Requirements

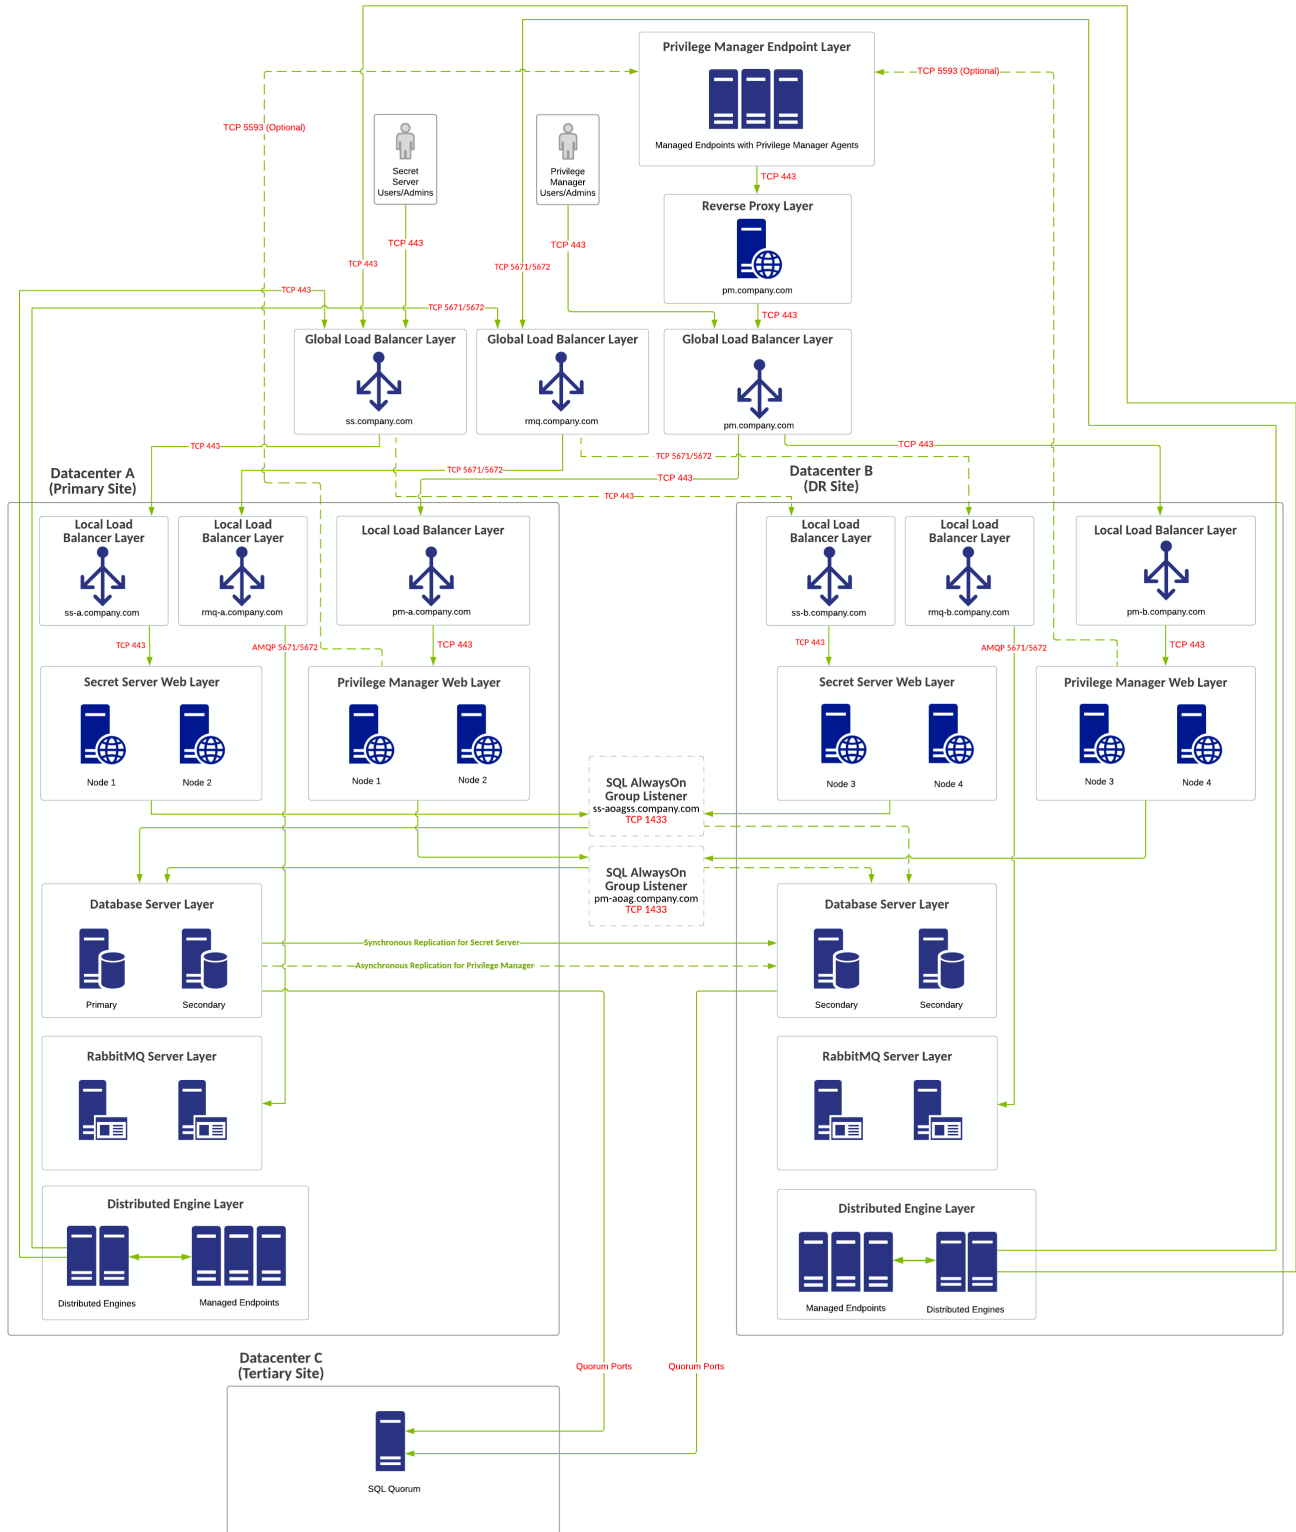
- SQL Enterprise Edition.
- Global and local load balancers.
- If global load balancers are not available due to cost or limited infrastructure, you can use local load balancers for all Web server nodes, but DNS change may be required if primary location goes offline.
- For SQL to stay online during single-node unplanned failures, you must configure a file-share witness for SQL quorum voting. We recommend a cloud witness.
- [Distributed Engine Ports](#).
- [SQL Quorum Ports](#).

Diagram



The reference for this diagram is C-2.

Figure: Multiple Site with Automatic Failover and Separate RabbitMQ Helper



Best Multiple Site with Automatic Failover and Separate RabbitMQ Helper

Overview

- Best HA configuration with no shared storage requirement.
- RabbitMQ Helper (for SS) is installed on dedicated servers (typically in a cluster).
- SQL AlwaysOn configurations are either synchronous or asynchronous for the SS database and asynchronous only for the PM database.
- DR site acts as a temporary site only with no long-term use. Services in DR site being down can incur downtime.
- PM is installed on separate Web servers.
- PM can integrate with SS for authentication and credential storage.
- PM can reside on the same database servers as SS or on separate ones, but SS and PM should not share the same database.



Due to SQL basic availability groups with the Standard Edition, you need to have multiple SQL instances and a separate AlwaysOn availability group configuration.

- You can use a separate Web reverse proxy or Azure service bus configuration for Privilege Manager agent TCP 443 communication.

Requirements

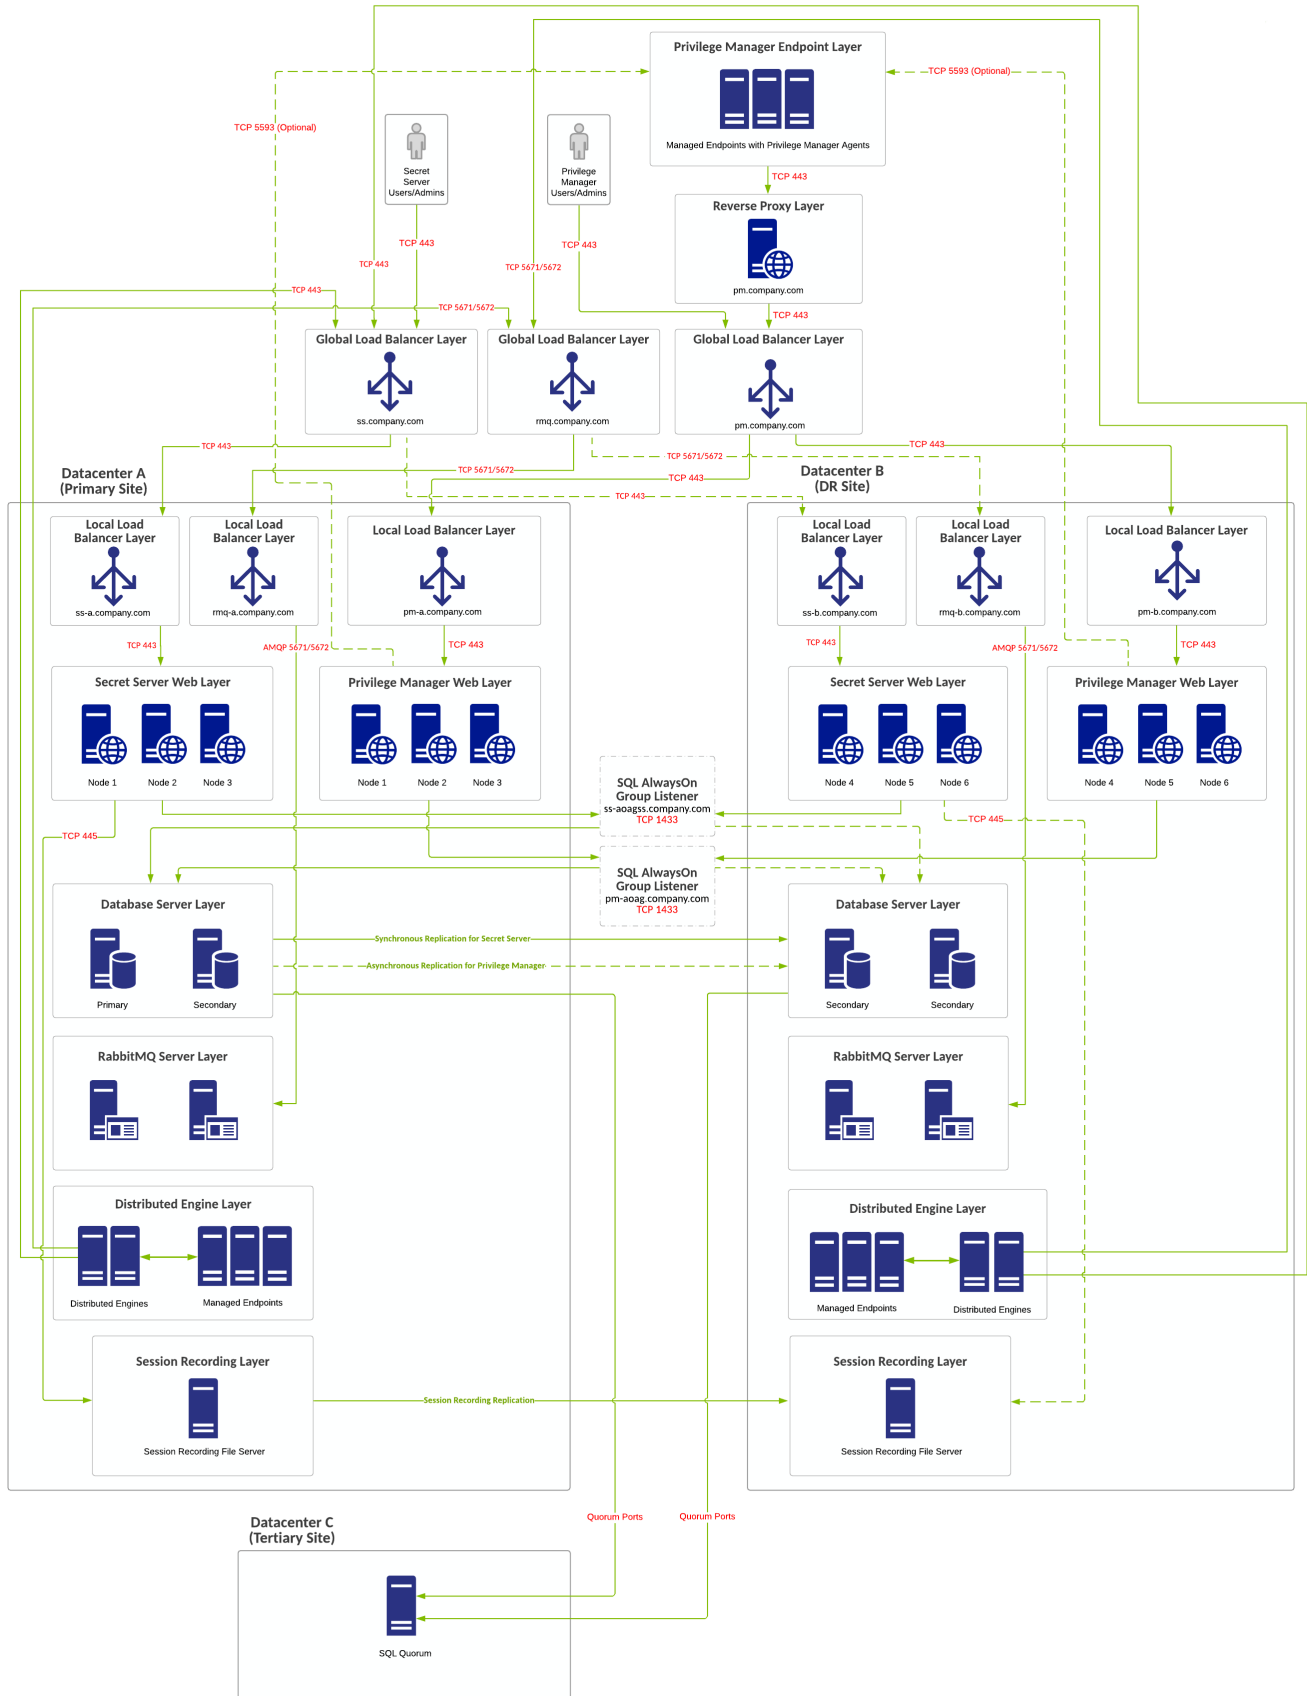
- SQL Enterprise Edition.
- Global and local load balancers.
- If global load balancers are not available due to cost or limited infrastructure, you can use local load balancers for all Web server nodes, but DNS change may be required if primary location goes offline.
- For SQL to stay online during single-node unplanned failures, you must configure a file-share witness for SQL quorum voting. We recommend a cloud witness.
- [Distributed Engine Ports](#).
- [SQL Quorum Ports](#).

Diagram




The reference for this diagram is C-3.

Figure: Best Multiple Site with Automatic Failover and Separate RabbitMQ Helper



Session Recording Example Architectures

 **Note:** If you are a current customer with support hours for Delinea Professional Services, you can discuss any of these diagrams in detail with one of our Professional Services Solutions Architects.

Endpoint via Secret Server Launcher

Overview

User connects to SS and clicks the launcher for a secret with a non-proxied session. This is the default. Session recording is enabled on the secret. This scenario applies to both SS and SSC.


Simplified flow:

User > SS > Protocol Launcher > Endpoint System

Data Gathered

Visual, video is created for video playback.

Flow

 **Note:** These numbers also appear on the diagram.

1. User logs on SS. User clicks icon to launch a session for a secret within SS with session recording enabled.
2. SS launches protocol handler.
3. Protocol handler requests instructions from SS.
4. SS sends secret and the target remote server.
5. Protocol launcher starts RDP session to target server. Recording begins.
6. End user system via the protocol handler does one of two things:
 - RDP/SSH launchers: Uploads either recorded video segments every second.
 - Mac Launchers: Takes screenshots every second. This information is uploaded via HTTPS (configurable) and stored in the database.
7. The Web servers' session recording role does any encoding, transcoding, or re-composition of videos to ensure videos can be played back through the session monitoring page.
8. The final video recording is stored in the database or file-share server (preferable) after the video has been processed. For Secret Server Cloud customers, this is stored in customer-specific BLOB. Playback can only occur through the Session Monitoring section within SS.

Diagram


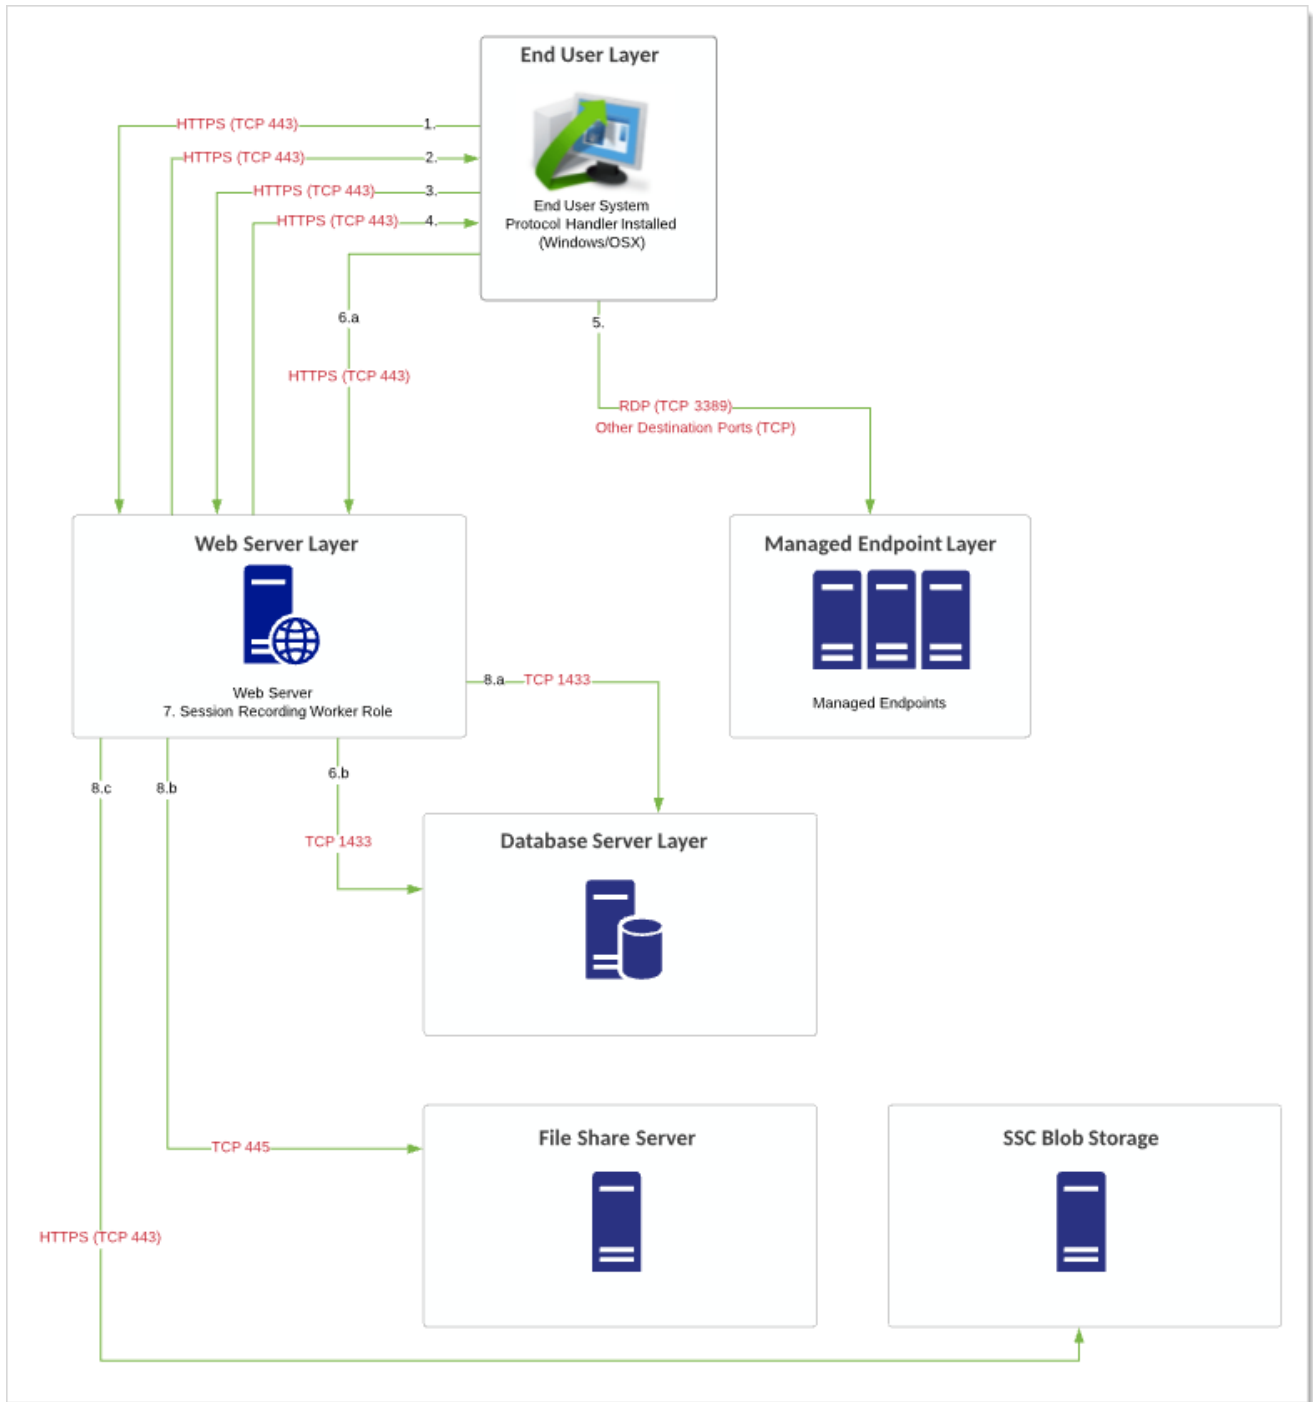
 **Note:** The reference for this diagram is Flow 1.

Figure: Endpoint via Secret Server Launcher



Notes

- Many of these ports are configurable. These diagrams reflect best practice ports. We strongly recommend using RabbitMQ Helper for your message bus. The ports attached to RabbitMQ Helper/MemoryMQ only need to be singular (SSL OR non-SSL). If using RabbitMQ Helper for example, most customers will only need to utilize port

TCP 5672.

- Session recording storage can be configured per site (separate file share servers), but the processing is still finalized from a Web server with session recording enabled. This means that data written from the Web server to the file share server may happen across different physical locations, which may add to the complexity of your networking requirements and may cause network saturation.
- If your client connection cannot support the needed bandwidth, the session data is still transmitted, but it takes longer to process each session. The protocol handler retries sending recordings to the Web server five times over the course of an hour. If that fails, it stops until the machine or service is restarted. The protocol handler itself terminates its launched session after 10 seconds if SS becomes unreachable. In a scenario where the protocol handler has terminated because SS is down, it stores the video recording in session monitoring up to the point when SS went down. Please see [Basic Session Recording](#) and [Advanced Session Recording](#) for more information.

Endpoint via Secret Server Launcher and Proxy

User connects to SS and clicks the launcher for a secret with a proxied session. Session recording is enabled on the secret. This scenario applies to both SS and SSC.

Simplified flow:

User > Secret Server > Launcher > Secret Server/DE Proxy > Endpoint System


Data Gathered

- For SSH secrets: Visual and Textual (keystroke/terminal output). Video is created for video playback and enhanced with textual data from the proxy.
- RDP secrets With RDP tunneling enabled: - Visual only. Video is created for video playback (same as the first scenario).
- For RDP proxy (new in 10.8): Video and Keystrokes

 **Note:** As of the December 10th 2019 release, you can now choose from:

- Record keystrokes only
- Record video only
- Do not record

Flow

 These numbers also appear on the diagram.

1. User logs on SS. User clicks icon to launch a session for a secret within SS with session recording enabled, and the secret proxies the request through either the Web server or a DE.
2. SS launches protocol handler.
3. Protocol handler requests instructions from SS.
4. SS sends connection details for Web server or DE proxy back to the protocol handler.

Delinea Architecture Reference Diagrams

5. Protocol launcher starts tunneling session to Web server or DE.
6. DE or Web server proxy requests connection parameters for destination from SS.
7. SS sends connection parameters to an individual Web server or DE proxy.
8. Proxy starts remote session to target server. Recording begins.
9. End user system via the protocol handler does one of two things:
 - RDP/SSH launchers: Uploads either recorded video segments every second.
 - Mac Launchers: Takes screenshots every second.

This information is uploaded via HTTPS (configurable) and stored in the database.

10. The Web servers' session recording role does any encoding, transcoding, or re-composition of videos to ensure videos can be played back through the session monitoring page.
11. The final video recording is stored in the database or file-share server (preferable) after the video has been processed. For Secret Server Cloud customers, this is stored in customer-specific BLOB. Playback can only occur through the Session Monitoring section within SS.

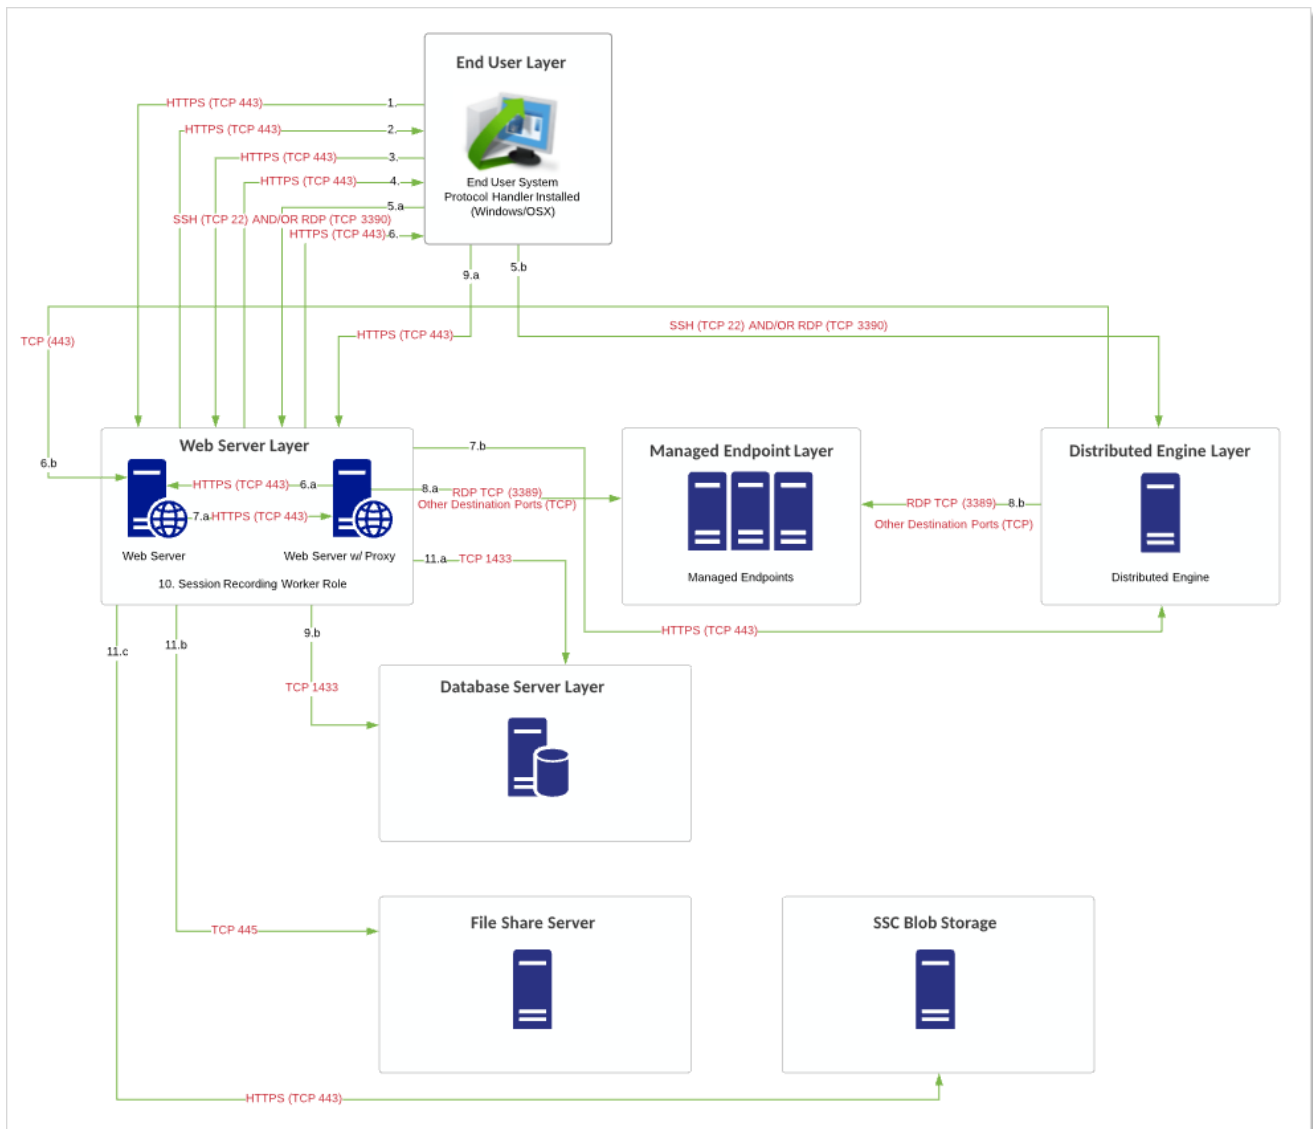
Diagram



The reference for this diagram is Flow 2.

Figure: Endpoint via Secret Server Launcher and Proxy

Delinea Architecture Reference Diagrams



Notes

- Many of these ports are configurable. These diagrams reflect best practice ports. We strongly recommend using RabbitMQ Helper for your message bus. The ports attached to RabbitMQ Helper/MemoryMQ only need to be singular (SSL OR non-SSL). If using RabbitMQ Helper for example, most customers will only need to utilize port TCP 5672.
- Session recording storage can be configured per site (separate file share servers), but the processing is still finalized from a Web server with session recording enabled. This means that data written from the Web server to the file share server may happen across different physical locations, which may add to the complexity of your networking requirements and may cause network saturation.
- If your client connection cannot support the needed bandwidth, the session data is still transmitted, but it takes longer to process each session. The protocol handler retries sending recordings to the Web server five times

over the course of an hour. If that fails, it stops until the machine or service is restarted. The protocol handler itself terminates its launched session after 10 seconds if SS becomes unreachable. In a scenario where the protocol handler has terminated because SS is down, it stores the video recording in session monitoring up to the point when SS went down.

- Please see [Basic Session Recording](#) and [Advanced Session Recording](#) for more information.

Endpoint via Secret Server Credentials and Proxy

Overview

Scenario A: User connects to SS and generates proxy credentials for a specific secret with session recording enabled. The user launches a terminal (PuTTY) session or RDP proxy session outside of SS, connecting to either the SS Web server or a DE.

Scenario B: User connects via a PuTTY session to the SS Web server or a DE proxy using SSH Terminal. User then runs TTY commands to manually initiate a connection to a destination system with a retrieved secret via the command line.

This scenarios both apply to both SS and SSC.

Simplified flow:

Scenario A:

User > Secret Server > Generate Proxy Credentials > Secret Server/DE Proxy > Endpoint System

Scenario B:

User > SSH Terminal> Secret Server > Provide Proxy Credentials > Secret Server/DE Proxy > Endpoint System

Data Gathered

Textual data that passed through the proxy (client to server and server to client) is recorded and can be viewed without a video recording.

Flow



These numbers also appear on the diagrams.

Scenario A

1. User logs on SS. User retrieves the proxy credential username and password for the configured secret.
2. User launches a PuTTY terminal session or RDP proxy session to a proxied Web server or DE, supplying the generated proxy credentials.
3. The Web server or DE matches the proxy credentials to the correct secret and uses the real credentials to connect to the destination system. Thus, the real credentials are never exposed to you or your machine.
4. Terminal output (server data) and client data (keystrokes) are recorded and are either written to the database by the Web server directly (Web server proxy) or is recorded by the DE proxy where the data is sent back in periodic 30-second chunks through RabbitMQ Helper/MemoryMQ bus. The Web servers pull this work off the

Delinea Architecture Reference Diagrams

RabbitMQ Helper/MemoryMQ bus for the engine worker role to store it in the database.

5. SS indexes the textual information using its background worker role for presentation in the session monitoring page.

Scenario B

1. User launches a PuTTY terminal session or RDP proxy session to a proxied Web server or DE, supplying the user's credentials. This requires the user to access to the relevant secrets.
2. The Web server or DE validates the username and password. If using a DE, a message is sent over the message bus to use SS to validate the credentials. The Web server provides a response to the DE. In some cases, SS may request additional information, such as a 2FA PIN code or a password re-do.
3. Terminal output (server data) and client data (keystrokes) are recorded and are either written to the database by the Web server directly (Web server proxy) or is recorded by the DE proxy where the data is sent back in periodic 30-second chunks through RabbitMQ Helper/MemoryMQ bus. The Web servers pull this work off the RabbitMQ Helper/MemoryMQ bus for the engine worker role to store it in the database.
4. SS indexes the textual information using its background worker role for presentation in the session monitoring page.

Diagram



The reference for these diagrams is Flow 3.

Figure: Endpoint via Secret Server Credentials and Proxy (Scenario A)

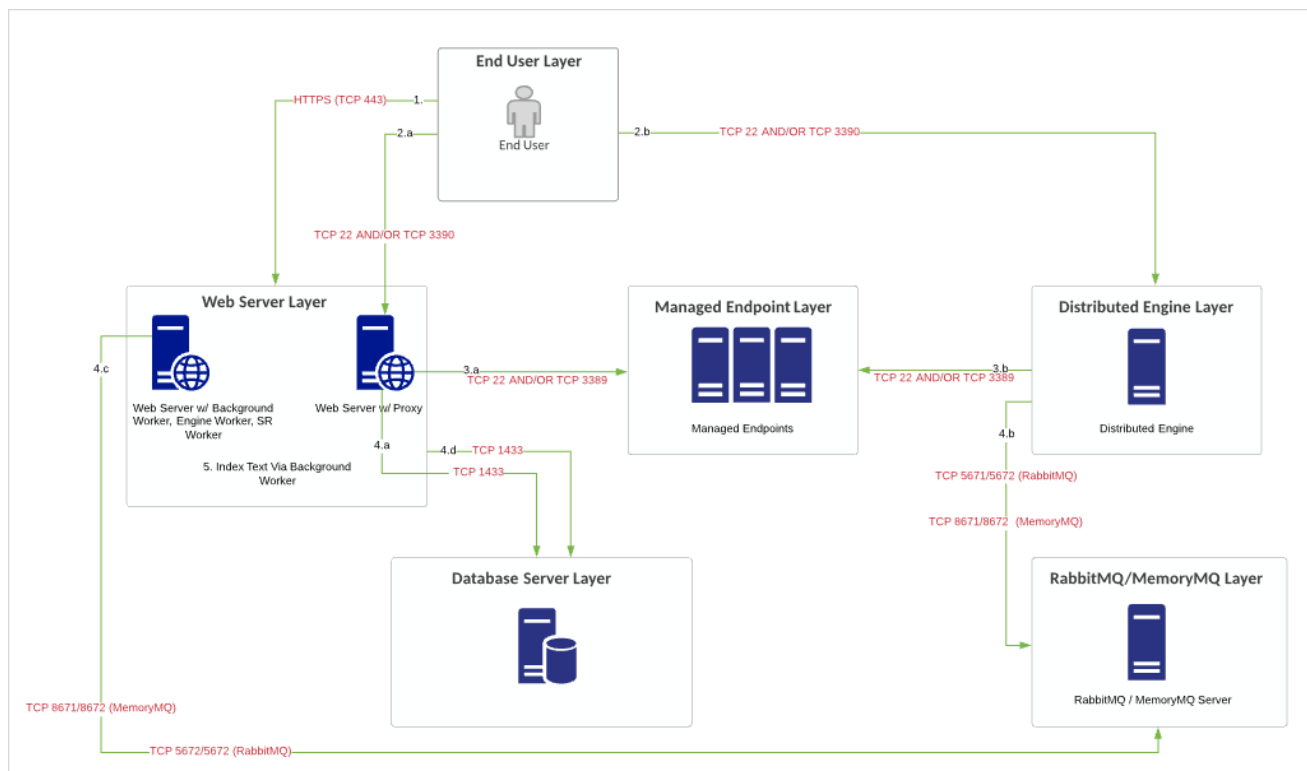
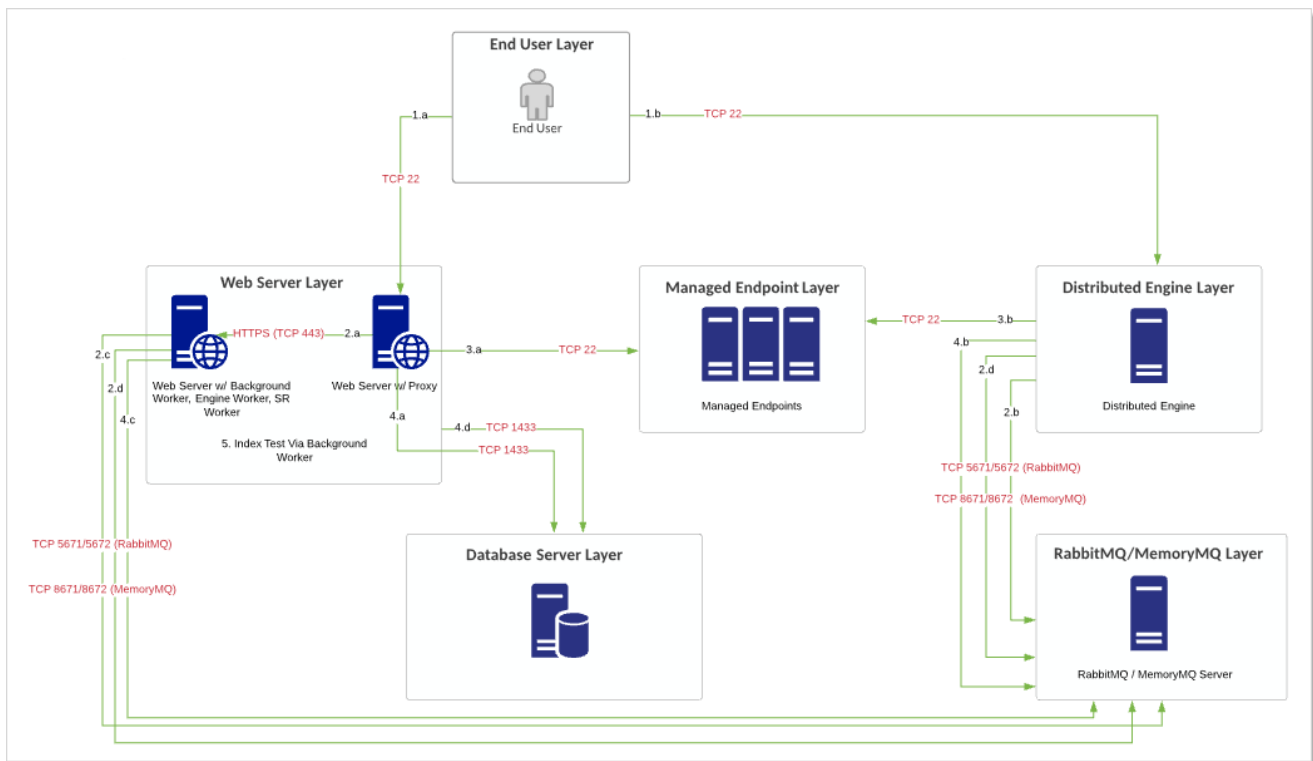


Figure: Endpoint via Secret Server Credentials and Proxy (Scenario B)



Notes

- Many of these ports are configurable. These diagrams reflect best practice ports. We strongly recommend using RabbitMQ Helper for your message bus. The ports attached to RabbitMQ Helper/MemoryMQ only need to be singular (SSL OR non-SSL). If using RabbitMQ Helper for example, most customers will only need to utilize port TCP 5672.
- This flow offers two different scenarios and is most relevant to Linux administrators and SSH-related secrets.
- Please see [Basic Session Recording](#) and [Advanced Session Recording](#) for more information.

Endpoint via Secret Server Launcher and ASRA

Overview

User connects to SS and clicks the launcher for a secret with a non-proxied session. This is the default. Session recording is enabled on the secret. The Advanced Session Recording Agent (ASRA) is installed on the endpoint. This scenario applies to both SS and SSC.

Simplified flow:

User > Secret Server > Launcher > Endpoint System + Advanced Session Recording Agent


Data Gathered

- Data Gathered: Visual, process, keyboard. Video is created for video playback and enhanced with process and keyboard data from the ASRA. The data includes visual data provided during the connection sequence with the destination system. Recorded session for playback is based on secret name.

 **Note:** As of the December 10th 2019 release, you can now choose from:

- Record keystrokes only
- Record video only
- Do not record


Flow

 **Note:** These numbers also appear on the diagram.

1. User logs on SS. User clicks icon to launch a session for a secret within SS with session recording enabled.
2. SS launches protocol handler.
3. Protocol handler requests instructions from SS.
4. SS sends the secret and the remote target server.
5. Proxy starts remote session to target server. Recording begins.
6. End user system via the protocol handler does one of two things:
 - RDP/SSH launchers: Uploads either recorded video segments every second.
 - Mac Launchers: Takes screenshots every second.

This information is uploaded via HTTPS (configurable) and stored in the database.

7. The destination system, via the ASRA, uploads keystroke and process (metadata) data by sending this information to the Web servers via HTTPS (TCP 443).

 Legacy ASR agents for on-premise deployments (prior to 10.7.000059) upload keystroke/process (metadata) by way of the response bus (RabbitMQ Helper/MemoryMQ). This is no longer pictured in this diagram below.

8. The Web servers' session recording role does any encoding, transcoding, or re-composition of videos to ensure videos can be played back through the session monitoring page.
9. The final video recording is stored in the database or file-share server (preferable) after the video has been processed. For SSC customers, this is stored in customer-specific BLOB. Playback can only occur through the Session Monitoring section within SS.

Diagram


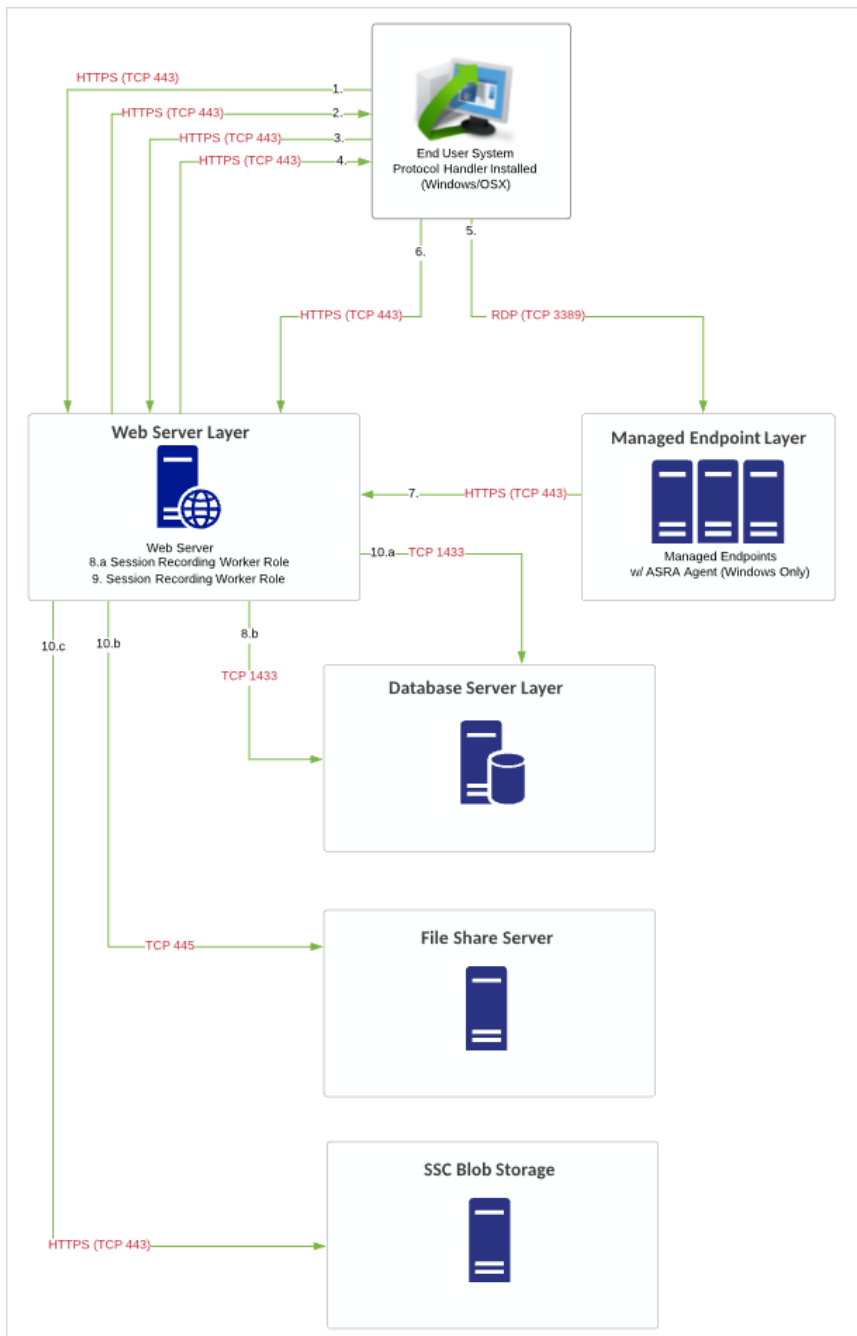
 The reference for this diagram is Flow 4.

Figure: Endpoint via Secret Server Launcher and ASRA

Delinea Architecture Reference Diagrams



Notes

- Many of these ports are configurable. These diagrams reflect best practice ports. We strongly recommend using RabbitMQ Helper for your message bus. The ports attached to RabbitMQ Helper/MemoryMQ only need to be singular (SSL OR non-SSL). If using RabbitMQ Helper for example, most customers will only need to utilize port TCP 5672.
- This flow requires installing ASRA on the destination endpoint. ASRA is only available for Windows clients.

Delinea Architecture Reference Diagrams

- This flow can be combined with [SSH proxying](#).
- Session recording storage can be configured per site (separate file share servers), but the processing is still finalized from a Web server with session recording enabled. This means that data written from the Web server to the file share server may happen across different physical locations, which may add to the complexity of your networking requirements and may cause network saturation.
- If your client connection cannot support the needed bandwidth, the session data is still transmitted, but it takes longer to process each session. The protocol handler retries sending recordings to the Web server five times over the course of an hour. If that fails, it stops until the machine or service is restarted. The protocol handler itself terminates its launched session after 10 seconds if SS becomes unreachable. In a scenario where the protocol handler has terminated because SS is down, it stores the video recording in session monitoring up to the point when SS went down.
- Please see [Basic Session Recording](#) and [Advanced Session Recording](#) for more information.

Endpoint and ASRA

Overview

User connects directly to an endpoint without accessing SS at any point. The Advanced Session Recording Agent (ASRA) is installed on the endpoint. This scenario applies to both SS and SSC.

Simplified flow:

User > Endpoint System + Advanced Session Recording Agent

Data Gathered

- Data Gathered: Visual, process, keyboard. Video is created for video playback and enhanced with process and keyboard data from the ASRA.



As of the December 10th 2019 release, you can now choose from:

- Record keystrokes only
- Record video only
- Do not record

Flow



These numbers also appear on the diagram.

1. User logs on the destination system. SS is not accessed. This is sometimes called "headless recording."
2. The destination system, via the ASRA, uploads keystroke and process (metadata) data by sending this information to the Web servers via HTTPS (TCP 443).



Legacy ASR agents for on-premise deployments (prior to 10.7.000059) upload keystroke/process (metadata) by way of the response bus (RabbitMQ Helper/MemoryMQ). This is no longer pictured in this diagram below.

3. The Web server's session recording worker role pulls the keystroke, process (metadata), and video data off the external bus. This information is uploaded via HTTPS (configurable) and stored in the database.
4. The Web servers' session recording role does any encoding, transcoding, or re-composition of videos to ensure videos can be played back through the session monitoring page.
5. The final video recording is stored in the database or file-share server (preferable) after the video has been processed. For SSC customers, this is stored in customer-specific BLOB. Playback can only occur through the Session Monitoring section within SS.

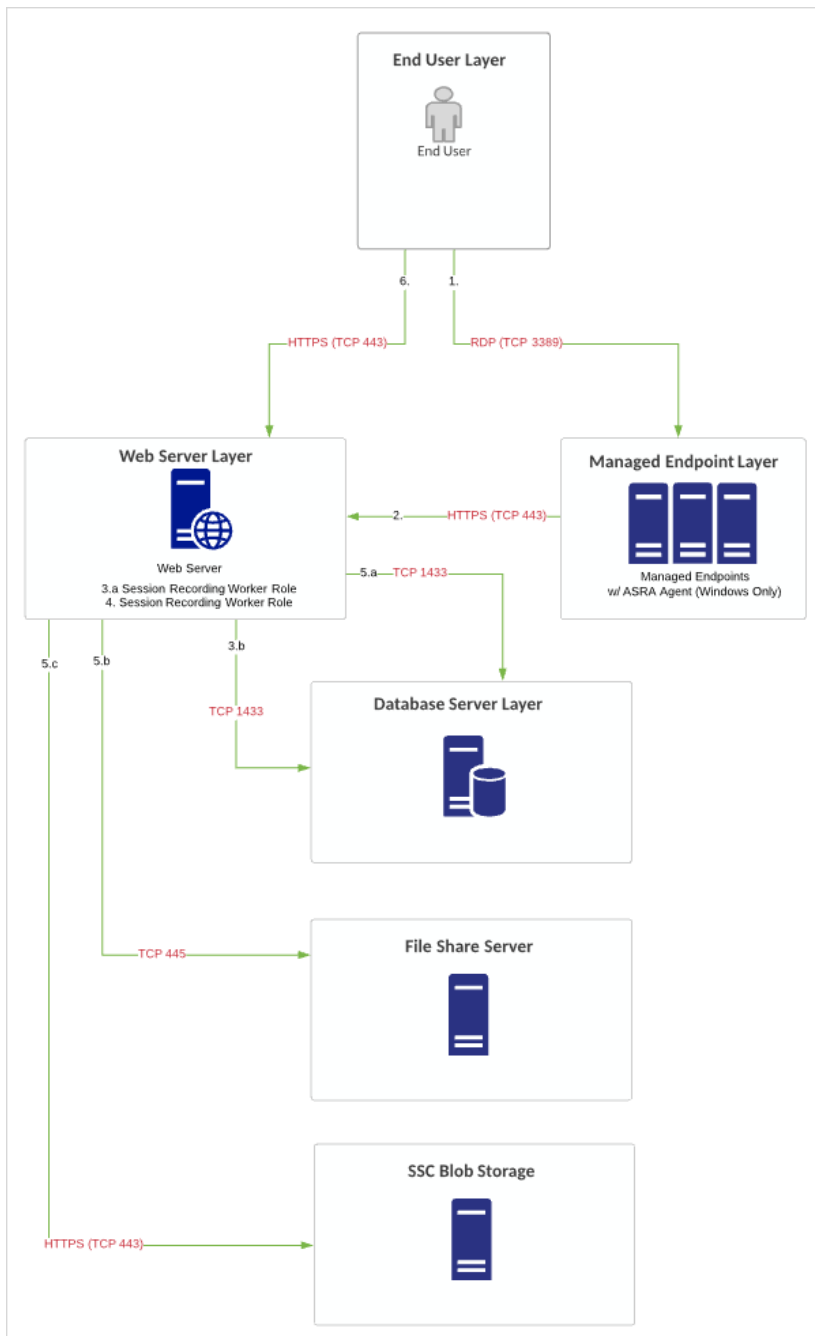
Diagram



The reference for this diagram is Flow 5.

Figure: Endpoint and ASRA

Delinea Architecture Reference Diagrams



Notes

- Many of these ports are configurable. These diagrams reflect best practice ports. We strongly recommend using RabbitMQ Helper for your message bus. The ports attached to RabbitMQ Helper/MemoryMQ only need to be singular (SSL OR non-SSL). If using RabbitMQ Helper for example, most customers will only need to utilize port TCP 5672.
- This flow requires installing ASRA on the destination endpoint. ASRA is only available for Windows clients.

Delinea Architecture Reference Diagrams

- This flow can be combined with [SSH proxying](#).
- Session recording storage can be configured per site (separate file share servers), but the processing is still finalized from a Web server with session recording enabled. This means that data written from the Web server to the file share server may happen across different physical locations, which may add to the complexity of your networking requirements and may cause network saturation.
- If your client connection cannot support the needed bandwidth, the session data is still transmitted, but it takes longer to process each session. The protocol handler retries sending recordings to the Web server six times over the course of three hours. If that fails, it stops until the machine or service is restarted.
- Please see [Advanced Session Recording](#) and [Basic Session Recording](#) for more information.

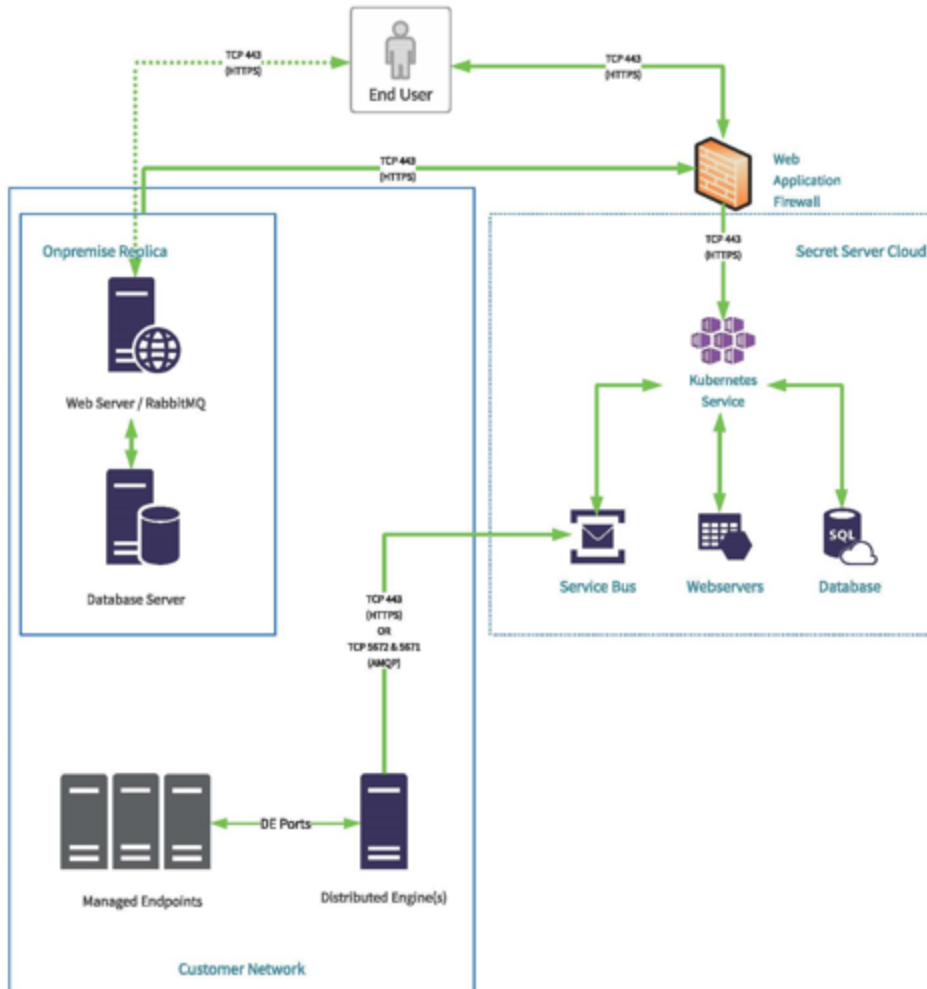
Secret Server Resilient Secrets Architecture

Diagram

Figure: Secret Server Resilient Secrets Architecture



| | | |
|--|---|----------------|
| Secret Server Cloud Disaster Recovery Onpremise Replica Reference Architecture | | Delinea |
| Version 1.1 | Created On: 6/29/2023 Updated: 6/30/2023 | |



Note: While the image above shows an architecture with a Secret Server Cloud source instance and an On-Premises replica instance, the architecture would be the same for an On-Premises source and Cloud replica setup.

To learn more about using Resilient Secrets with the Delinea Platform, please see additional [Delinea Platform documentation](#).

Requirements for Secret Server Cloud Replica Instances

- If you have On-Premises source instance and Cloud replica instance, you need to whitelist the same inbound IP addresses as the RADIUS authentication incoming from Secret Server Cloud. ([Learn More](#))
- The externally facing source server needs a valid certificate signed by a trusted CA . You cannot upload your own CA cert to Secret Server Cloud.

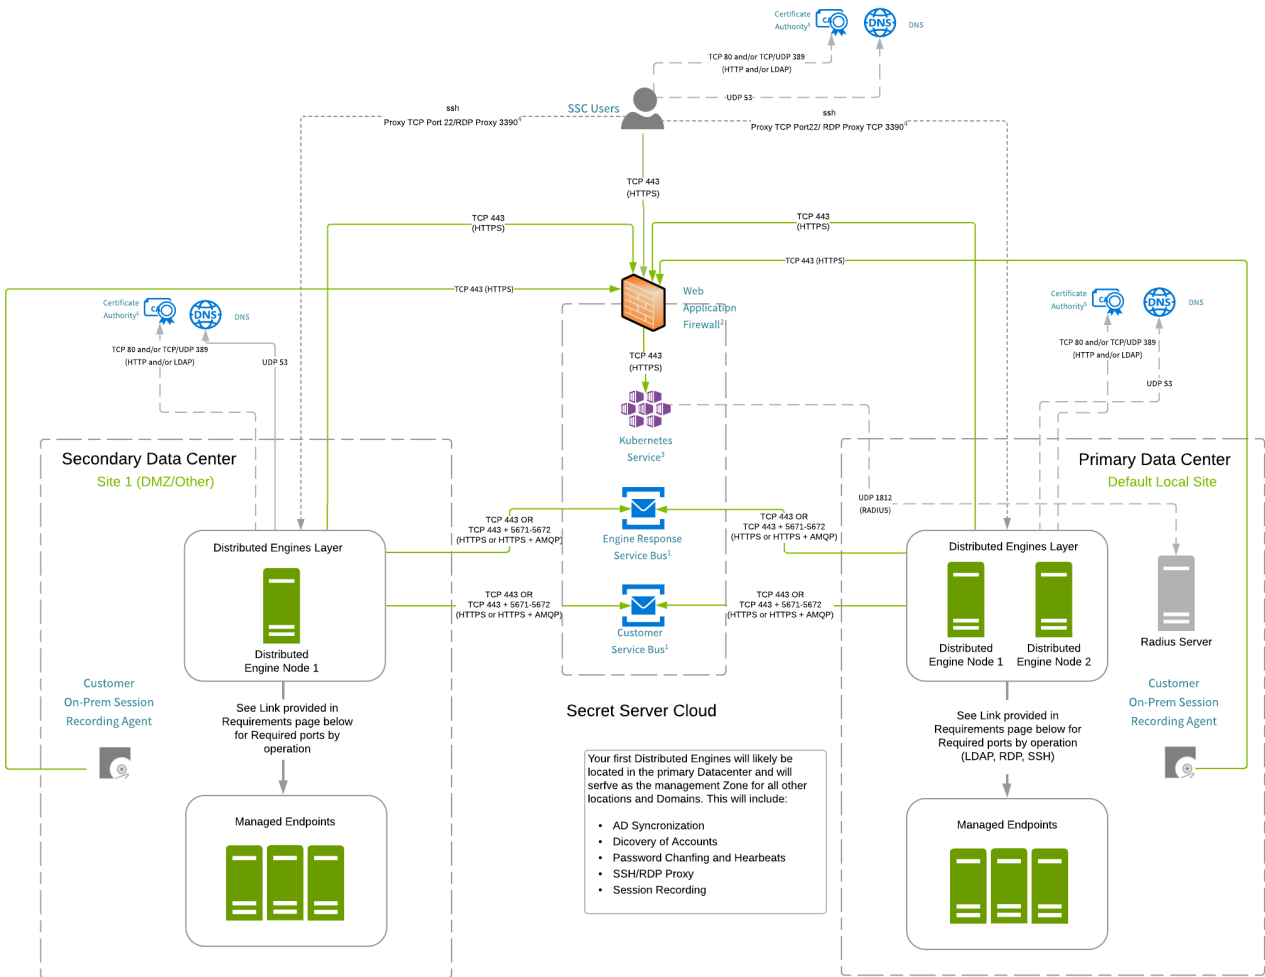
Secret Server Cloud Customer Example Architectures


Note: If you are a current customer with support hours for Delinea Professional Services, you can discuss any of these diagrams in detail with one of our Professional Services support architects.


Note: See "Delinea Platform Architecture" on page 20 for additional cloud architecture.

Multi-site with ASR Agents Example Architecture

Figure: Multi-site with ASR Agents Example Architecture




 **Note:** This design is fully supported by Delinea.

 **Note:** Arrows indicate the direction of initial connection.

 Reference architecture requirements:

- Ports for accessing, managing and discovering end-points must have the required ports opened between the site Distributed Engines and the appropriate devices. Please see [Ports Used by Secret Server](#).
- All Distributed Engines servers must run on Windows Server 2016 or higher. Windows Server 2022 is supported by Secret Server 11.0 or later. It may work with earlier versions, but that has not been officially confirmed.
- Distributed Engines servers must have 4 cores and 4 GB RAM. We encourage increasing CPUs before RAM to improve DE efficiency.

 Your first distributed engines will likely be located in the primary data center and will serve as the management zone for all other locations and domains. This includes:

- AD synchronization
- Account discovery
- Password changing and heartbeats
- SSH and RDP proxy
- Session recording

Details for All Architectures

1: *Service Buses*

IP Address allow-listing is not necessary unless outbound firewall rules are in place. If IP allow-listing is necessary, please contact [Delinea Support](#) to obtain the shared engine response service bus and your dedicated customer service bus hostnames. The TCP port requirement is based on the transport type configured in the distributed engine settings. The default is Web sockets, which requires TCP 443. If the AMQP option is selected within the application, TCP 5671/5672 ports are also required.

2: *Web Application Firewall (WAF)*

IP Address allow-listing is not necessary unless outbound firewall rules are in place. Generally, the public IP the hostname resolves to is based on geographical location of the request source. All IPs below should be allow-listed to ensure uninterrupted connectivity.

All regions:

- 45.60.32.37
- 45.60.34.37
- 45.60.36.37

Delinea Architecture Reference Diagrams

- 45.60.38.37
- 45.60.40.37
- 45.60.104.37

3: RADIUS

Inbound allow-listing is necessary if RADIUS authentication is configured. IP addresses for RADIUS authentication configuration:



For convenience, the new additional ranges (as of October 2024) are consolidated into a single CIDR (Classless Inter-Domain Routing) block per region.

secretservercloud.com

- 20.65.118.12 (Primary)
- 23.102.107.104 (Primary)
- 23.102.107.220 (Primary)
- 23.102.106.185 (Primary)
- 23.102.108.55 (Primary)
- 52.224.253.7 (Primary)
- 52.224.253.4 (Primary)
- 52.151.206.73 (Primary)
- 52.151.206.77 (Primary)
- 52.151.206.35 (Primary)
- 20.228.138.112/29 (Primary)
- 52.160.67.39 (DR)
- 52.160.67.38 (DR)
- 104.40.25.170 (DR)
- 138.91.163.99 (DR)
- 137.135.51.234 (DR)
- 52.190.184.16/29 (DR)

secretservercloud.co.uk

- 20.0.46.111 (Primary)
- 51.142.243.172 (Primary)
- 20.0.46.112 (Primary)
- 20.0.46.123 (Primary)
- 20.0.46.124 (Primary)

Delinea Architecture Reference Diagrams

- 20.162.162.64/29 (Primary)
- 51.104.62.220 (Secondary)
- 51.104.62.213 (Secondary)
- 51.104.63.38 (Secondary)
- 51.104.62.185 (Secondary)
- 51.104.62.252 (Secondary)
- 20.117.16.40/29 (Secondary)

secretservercloud.ca

- 52.228.117.246 (Primary)
- 52.228.113.119 (Primary)
- 52.139.7.40 (Primary)
- 52.139.7.137 (Primary)
- 52.139.7.197 (Primary)
- 40.85.220.216/29 (Primary)
- 52.229.119.193 (DR)
- 52.229.119.89 (DR)
- 52.235.39.79 (DR)
- 52.235.39.125 (DR)
- 52.235.39.5 (DR)
- 20.220.90.80/29 (DR)

secretservercloud.eu

- 20.79.64.213 (Primary)
- 20.79.65.3 (Primary)
- 20.79.226.78 (Primary)
- 20.79.226.180 (Primary)
- 20.79.226.116 (Primary)
- 51.116.178.152/29 (Primary)
- 20.50.180.242 (DR)
- 20.50.180.187 (DR)
- 20.50.154.28 (DR)
- 20.50.176.86 (DR)

- 20.50.156.219 (DR)
- 20.16.113.88.144/29 (DR)

secretservercloud.com.sg

- 20.195.97.220 (Primary)
- 20.195.98.154 (Primary)
- 20.212.128.73 (Primary)
- 20.212.128.75 (Primary)
- 20.212.128.74 (Primary)
- 52.237.113.56/29 (Primary)
- 65.52.165.108 (DR)
- 65.52.160.251 (DR)
- 52.184.100.188 (DR)
- 52.184.101.189 (DR)
- 52.184.101.213 (DR)
- 23.100.88.144/29 (DR)

secretservercloud.com.au

- 20.37.251.37 (Primary)
- 20.37.251.120 (Primary)
- 20.37.5.233 (Primary)
- 20.37.5.227 (Primary)
- 20.37.5.48 (Primary)
- 20.37.1.16/29 (Primary)
- 20.53.142.34 (DR)
- 20.53.142.37 (DR)
- 20.53.80.77 (DR)
- 20.53.81.216 (DR)
- 20.53.82.77 (DR)
- 23.101.211.80/29 (DR)

4: *Distributed Engine (DE)*

If external clients must be able to connect to internal SSH or RDP endpoints, an SSH proxy can be configured on the DE. Additionally, TCP port 22 needs to be open for inbound connections on the DE server, as well as have an appropriate configuration to allow inbound connections from the public Internet.

5: Certificate CRLs

Allow-listing is not necessary unless outbound firewall rules are in place. If it is necessary, access to CRLs or OSCP endpoints may be required. CRL and OSCP endpoints may differ from customer to customer. To determine the endpoints, review the certificates presented by the:

- Web application firewall
- Customer service bus
- Engine response service bus
- CDN for DE updates



Obtaining and reviewing certificates is not within the scope of this document, but you can find resources online, such as [OCSP & CRL and Revoked SSL Certificates](#), which is not owned or maintained by Delinea.



As of October 2024, we switched from one-year-validity Sectigo certificates to 90-day-validity Let's Encrypt certificates. Let's Encrypt certificates are widely trusted and have effectively become the standard for cloud-native operations. See [About Let's Encrypt](#) for more information.

Secret Server Hybrid Multi-Tenant Cloud Architecture



Note: If you are a current customer with support hours for Delinea Professional Services, you can discuss any of these diagrams in detail with one of our Professional Services support architects.



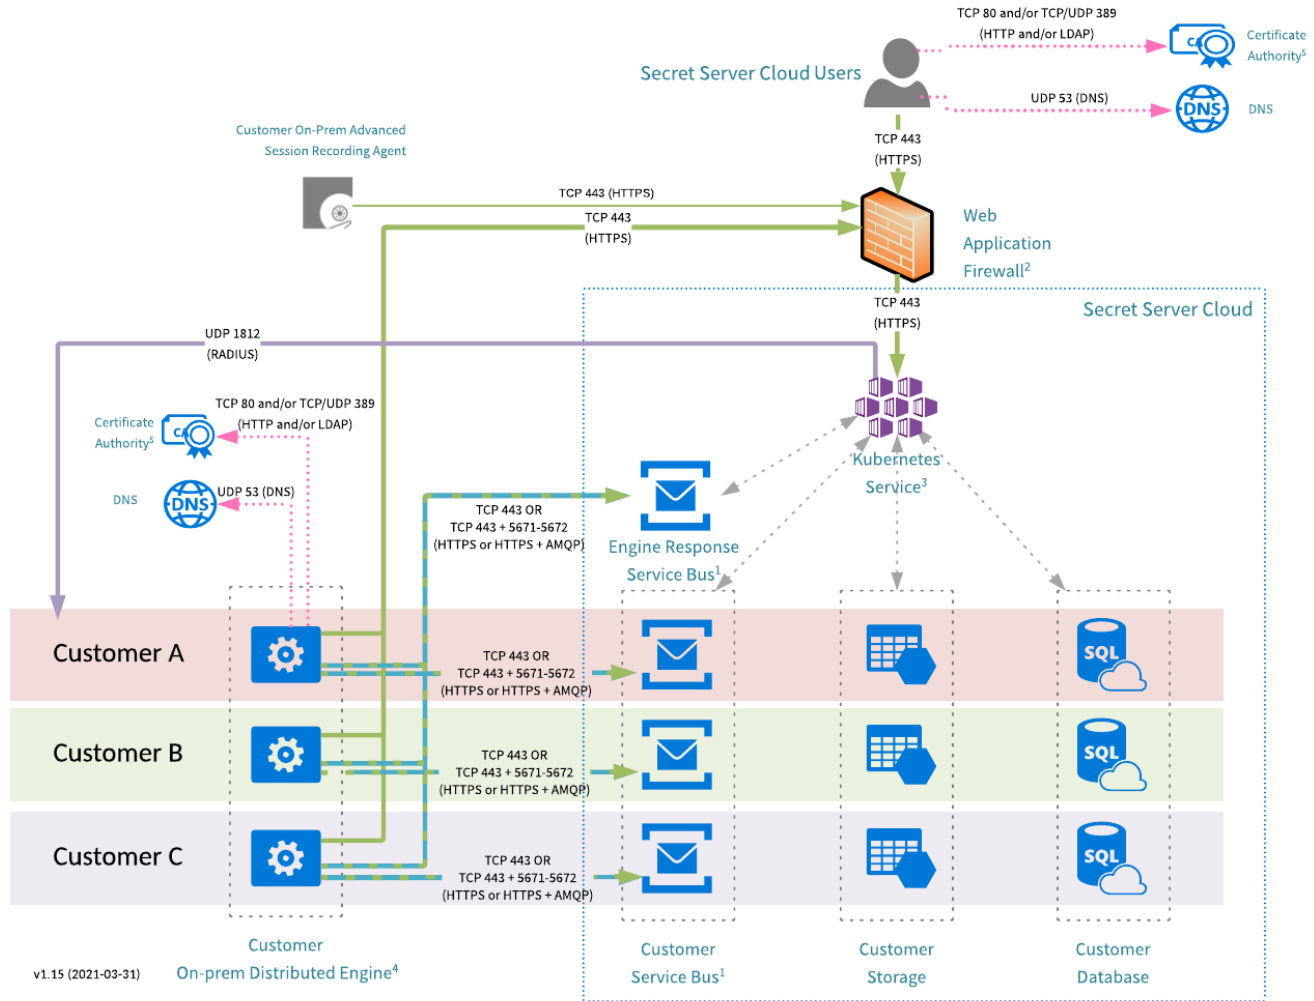
Note: This, the standard Secret Server Cloud architecture, is "hybrid multi-tenant" because only the front-end is multi-tenant, it is shared with other customers. The databases, service buses, and storage accounts are single-tenant (dedicated to you).



Note: See "Delinea Platform Architecture" on page 20 for additional cloud architecture.

Diagram

Figure: Secret Server Cloud Architecture



Note: Arrows indicate the direction of initial connection.

Details

1: Service Buses

IP Address allowlisting is not necessary unless outbound firewall rules are in place. If IP allowlisting is necessary, you can find your customer-specific service bus IP addresses by navigating to `customer.secretservercloud.com/Admin/Diagnostics.aspx`. The TCP port requirement is based on the transport type configured in the distributed engine settings. The default is Web sockets, which requires TCP 443. If the AMQP option is selected within the application, TCP 5671/5672 ports are also required.

The shared engine-response-bus hostnames are:

- `thycotic-ssc-eu-er-sb-01-prod-g.servicebus.windows.net`
- `thycotic-ssc-eu-er-sb-01-prod-b.servicebus.windows.net` (Active)

2: Web Application Firewall (WAF)

IP Address allowlisting is not necessary unless outbound firewall rules are in place. Generally, the public IP the hostname resolves to is based on geographical location of the request source. All IPs below should be allowlisted to ensure uninterrupted connectivity.

All regions:

- 45.60.32.37
- 45.60.34.37
- 45.60.36.37
- 45.60.38.37
- 45.60.40.37
- 45.60.104.37

3: RADIUS

Inbound allowlisting is necessary if RADIUS authentication is configured. IP addresses:

secretservercloud.com

- 20.65.118.12 (Primary)
- 23.102.107.104 (Primary)
- 23.102.107.220 (Primary)
- 23.102.106.185 (Primary)
- 23.102.108.55 (Primary)
- 52.224.253.7 (Primary)
- 52.224.253.4 (Primary)
- 52.151.206.73 (Primary)
- 52.151.206.77 (Primary)
- 52.151.206.35 (Primary)
- 52.160.67.39 (DR)
- 52.160.67.38 (DR)
- 104.40.25.170 (DR)
- 138.91.163.99 (DR)
- 137.135.51.234 (DR)

secretservercloud.co.uk

- 20.0.46.111 (Primary)
- 51.142.243.172 (Primary)

Delinea Architecture Reference Diagrams

- 20.0.46.112 (Primary)
- 20.0.46.123 (Primary)
- 20.0.46.124 (Primary)
- 51.104.62.220 (Secondary)
- 51.104.62.213 (Secondary)
- 51.104.63.38 (Secondary)
- 51.104.62.185 (Secondary)
- 51.104.62.252 (Secondary)

secretservercloud.ca

- 52.228.117.246 (Primary)
- 52.228.113.119 (Primary)
- 52.139.7.40 (Primary)
- 52.139.7.137 (Primary)
- 52.139.7.197 (Primary)
- 52.229.119.193 (DR)
- 52.229.119.89 (DR)
- 52.235.39.79 (DR)
- 52.235.39.125 (DR)
- 52.235.39.5 (DR)

secretservercloud.eu

- 20.79.64.213 (Primary)
- 20.79.65.3 (Primary)
- 20.79.226.78 (Primary)
- 20.79.226.180 (Primary)
- 20.79.226.116 (Primary)
- 20.50.180.242 (DR)
- 20.50.180.187 (DR)
- 20.50.154.28 (DR)
- 20.50.176.86 (DR)
- 20.50.156.219 (DR)

secretservercloud.com.sg

- 20.195.97.220 (Primary)
- 20.195.98.154 (Primary)
- 20.212.128.73 (Primary)
- 20.212.128.75 (Primary)
- 20.212.128.74 (Primary)
- 65.52.165.108 (DR)
- 65.52.160.251 (DR)
- 52.184.100.188 (DR)
- 52.184.101.189 (DR)
- 52.184.101.213 (DR)

secretservercloud.com.au

- 20.37.251.37 (Primary)
- 20.37.251.120 (Primary)
- 20.37.5.233 (Primary)
- 20.37.5.227 (Primary)
- 20.37.5.48 (Primary)
- 20.53.142.34 (DR)
- 20.53.142.37 (DR)
- 20.53.80.77 (DR)
- 20.53.81.216 (DR)
- 20.53.82.77 (DR)

4: Distributed Engine (DE)

If external clients must be able to connect to internal SSH or RDP endpoints, an SSH proxy can be configured on the DE. Additionally, TCP port 22 needs to be open for inbound connections on the DE server, as well as have an appropriate configuration to allow inbound connections from the public Internet.

5: Certificate CRLs

Allowlisting is not necessary unless outbound firewall rules are in place. If it is necessary, access to CRLs or OSCP endpoints may be required. CRL and OSCP endpoints may differ from customer to customer. To determine the endpoints, review the certificates presented by the:

- Web application firewall
- Customer service bus

Delinea Architecture Reference Diagrams

- Engine response service bus
- CDN for DE updates



Obtaining and reviewing certificates is not within the scope of this document, but you can find resources online, such as [OCSP & CRL and Revoked SSL Certificates](#), which is not owned or maintained by Delinea.