

Reference Architecture

16 November 2018
vRealize Automation 7.4



vmware®

You can find the most up-to-date technical documentation on the VMware website at:

<https://docs.vmware.com/>

If you have comments about this documentation, submit your feedback to

docfeedback@vmware.com

VMware, Inc.
3401 Hillview Ave.
Palo Alto, CA 94304
www.vmware.com

Copyright © 2017–2018 VMware, Inc. All rights reserved. [Copyright and trademark information.](#)

Contents

vRealize Automation Reference Architecture Guide	4
Updated Information	5
1 Initial Deployment and Configuration Recommendations	6
2 vRealize Automation Deployment	7
3 vRealize Business for Cloud Deployment Considerations	10
4 vRealize Automation Scalability	11
Configure Manager Service for High Data Volume	13
Distributed Execution Manager Performance Analysis and Tuning	13
5 vRealize Business for Cloud Scalability	15
6 vRealize Automation High Availability Configuration Considerations	16
7 vRealize Business for Cloud High Availability Considerations	19
8 vRealize Automation Hardware Specifications and Capacity Maximums	20
9 vRealize Automation Small Deployment Requirements	23
10 vRealize Automation Medium Deployment Requirements	29
11 vRealize Automation Large Deployment Requirements	35
12 vRealize Automation Multi-Data Center Data Deployments	41

vRealize Automation Reference Architecture Guide

The *vRealize Automation Reference Architecture Guide* describes the structure and configuration of typical vRealize Automation deployments. In addition, it provides information about high availability, scalability and deployment profiles.

Intended Audience

This information is intended for anyone who wants to configure and manage vRealize Automation. The information is written for experienced Windows or Linux system users and administrators who are familiar with virtual machine technology and datacenter operations.

VMware Technical Publications Glossary

VMware Technical Publications provides a glossary of terms that might be unfamiliar to you. For definitions of terms as they are used in VMware technical documentation, go to <http://www.vmware.com/support/pubs>.

Updated Information

This *Reference Architecture* is updated with each release of the product or when necessary.

This table provides the update history of the *Reference Architecture*.

Revision	Description
16 NOV 2018	Minor edits.
05 OCT 2018	Minor updates, including updated graphics in the following topics: <ul style="list-style-type: none">■ Chapter 10 vRealize Automation Medium Deployment Requirements■ Chapter 11 vRealize Automation Large Deployment Requirements
15 JUN 2018	Minor updates.
12 APR 2018	Initial release.

Initial Deployment and Configuration Recommendations

1

Deploy and configure all VMware vRealize Automation components in accordance with VMware recommendations.

Keep your vRealize Automation, vRealize Business for Cloud, and vRealize Orchestrator in the same time zone with their clocks synchronized.

Install vRealize Automation, vRealize Business for Cloud, and vRealize Orchestrator on the same management cluster. Provision machines to a cluster that is separate from the management cluster so that user workload and server workload can be isolated.

Deploy Proxy Agents in the same data center as the Endpoint with which they communicate. VMware does not recommend placing DEM Workers in Remote Data Centers unless there is an express workflow skill based use case that requires it. All components except the Proxy Agents and DEM Workers must be deployed in the same Data Center or Data Centers within a Metro Area Network. Latency must be less than 5 milliseconds, and bandwidth must not be less than 1 GB/s between the Data Centers in the Metro Area Network.

For more information including a support statement, see the VMware Knowledge Base article *Installing the VMware vRealize Automation on a distributed multi-site instance* available at http://kb.vmware.com/selfservice/microsites/search.do?language=en_US&cmd=displayKC&externalId=213484.2.

vRealize Automation Deployment

2

Use the VMware resource recommendations as a starting point for vRealize Automation deployment planning.

After initial testing and deployment to production, continue to monitor performance and allocate additional resources if necessary, as described in [Chapter 4 vRealize Automation Scalability](#).

Authentication

When configuring vRealize Automation, you can use the default Directories Management connector for user authentication, or you can specify a pre-existing SAML based identity provider to support a single-sign on experience.

If two-factor authentication is required vRealize Automation supports integration with RSA SecurID. When this integration point is configured, users are prompted for their user ID and passcode.

Load Balancer Considerations

Use the Least Response Time or round-robin method to balance traffic to the vRealize Automation appliances and infrastructure Web servers. Enable session affinity or the sticky session feature to direct subsequent requests from each unique session to the same Web server in the load balancer pool.

You can use a load balancer to manage failover for the Manager Service, but do not use a load-balancing algorithm, because only one Manager Service is active at a time. Also, do not use session affinity when managing failover with a load balancer.

Use ports 443 and 8444 when load balancing the vRealize Automation Appliance. For the Infrastructure Website and Infrastructure Manager Service, only port 443 should be load balanced.

Although you can use other load balancers, NSX, F5 BIG-IP hardware, and F5 BIG-IP Virtual Edition are tested and are recommended for use.

See the vRealize Automation documentation for detailed information on configuring load balancers.

Database Deployment

vRealize Automation automatically clusters the appliance database in 7.0 and later releases. All new 7.0 and later deployments must use the internal appliance database. vRealize Automation instances which are upgrading to 7.1 or later must merge their external databases into the appliance database. See the vRealize Automation product documentation for more information on the upgrade process.

For production deployments of the Infrastructure components, use a dedicated database server to host the Microsoft SQL Server (MSSQL) databases. vRealize Automation requires machines that communicate with the database server to be configured to use Microsoft Distributed Transaction Coordinator (MSDTC). By default, MSDTC requires port 135 and ports 1024 through 65535.

For more information about changing the default MSDTC ports, see the Microsoft Knowledge Base article [Configuring Microsoft Distributed Transaction Coordinator \(DTC\) to work through a firewall](https://support.microsoft.com/en-us/kb/250367) available at <https://support.microsoft.com/en-us/kb/250367>.

The IaaS Manager Service host must be able to resolve the NETBIOS name of the IaaS SQL Server database host. If it cannot resolve the NETBIOS name, add the SQL Server NETBIOS name to the Manager Service machine `/etc/hosts` file and restart the Manager Service.

vRealize Automation supports SQL AlwaysON groups only with Microsoft SQL Server 2016. When installing SQL Server 2016, the database must be created in 100 mode. If you use an older version of Microsoft SQL Server, use a Failover Cluster instance with shared disks. For more information on configuring SQL AlwaysOn groups with MSDTC, see <https://msdn.microsoft.com/en-us/library/ms366279.aspx>.

Data Collection Configuration

The default data collection settings provide a good starting point for most implementations. After deploying to production, continue to monitor the performance of data collection to determine whether you must make any adjustments.

Proxy Agents

For maximum performance, deploy agents in the same data center as the endpoint to which they are associated. You can install additional agents to increase system throughput and concurrency. Distributed deployments can have multiple agent servers that are distributed around the globe.

When agents are installed in the same data center as their associated endpoint, you can see an increase in data collection performance of 200 percent, on average. The collection time measured includes only the time spent transferring data between the proxy agent and the manager service. It does not include the time it takes for the manager service to process the data.

For example, you currently deploy the product to a data center in Palo Alto and you have vSphere endpoints in Palo Alto, Boston, and London. In this configuration, the vSphere proxy agents are deployed in Palo Alto, Boston, and London for their respective endpoints. If instead, agents are deployed only in Palo Alto, you might see a 200 percent increase in data collection time for Boston and London.

Distributed Execution Manager Configuration

In general, locate distributed execution managers (DEMs) as close as possible to the model manager host. The DEM Orchestrator must have strong network connectivity to the model manager at all times. By default, the installer places DEM Orchestrators alongside the Manager Service. Create two DEM Orchestrator instances, one for failover, and two DEM Worker instances in your primary data center.

If a DEM Worker instance must run a location-specific workflow, install the instance in that location.

Assign skills to the relevant workflows and DEMs so that those workflows are always run by DEMs in the correct location. For information about assigning skills to workflows and DEMs by using the vRealize Automation designer console, see the vRealize Automation Extensibility documentation.

For the best performance, install DEMs and agents on separate machines. For additional information about installing vRealize Automation agents, see the vRealize Automation *Installing vRealize Automation* documentation.

vRealize Orchestrator

Use the internal vRealize Orchestrator instance for all new deployments. If necessary, legacy deployments can continue to use an external vRealize Orchestrator. See https://kb.vmware.com/selfservice/microsites/search.do?language=en_US&cmd=displayKC&externalId=2147109 for the procedure to increase the memory allocated to the internal vRealize Orchestrator instance.

For best product performance, review and implement configuration guidelines described in the *vRealize Orchestrator Coding Design Guide* prior to importing vRealize Orchestrator content into production deployments.

vRealize Business for Cloud Deployment Considerations

3

Deploy vRealize Business for Cloud, formerly known as vRealize Business Standard Edition, in accordance with VMware guidelines.

Load Balancer Considerations

Load balancing is not supported for data collection connections. For more information, see [Chapter 4 vRealize Automation Scalability](#). In the vRealize Business for Cloud appliance for user interface and API client connections, you can use the vRealize Automation load balancer.

4

vRealize Automation Scalability

Consider all applicable scalability factors when configuring your vRealize Automation system.

Users

The vRealize Automation appliance is configured for syncing less than 100,000 users. If your system contains more users, you may need to add memory to vRealize Automation Directories Management. For detailed information on adding memory to Directories Management, see "Add Memory to Directories Management" in *Configuring vRealize Automation*.

Concurrent Provisions Scalability

By default, vRealize Automation processes only eight concurrent provisions per endpoint. For information about increasing this limit, see *Configuring vRealize Automation*.

VMware recommends that all deployments start with at least two DEM-Workers. In 6.x each DEM-Worker could process 15 workflows concurrently. This was increased to 30 for vRealize Automation 7.0 and later.

If machines are being customized through Workflow Stubs, you should have 1 DEM-Worker per 20 Machines that will be provisioned concurrently. For example, a system supporting 100 concurrent provisions should have a minimum of 5 DEM-Workers.

For more information on DEM-Workers and scalability see [Distributed Execution Manager Performance Analysis and Tuning](#)

Data Collection Scalability

Data collection completion time depends on the compute resource capacity, the number of machines on the compute resource or endpoint, the current system, and network load, among other variables. The performance scales at a different rate for different types of data collection.

Each type of data collection has a default interval that you can override or modify. Infrastructure administrators can manually initiate data collection for infrastructure source endpoints. Fabric administrators can manually initiate data collection for compute resources. The following values are the default intervals for data collection.

Table 4-1. Data Collection Default Intervals

Data Collection Type	Default Interval
Inventory	Every 24 hours (daily)
State	Every 15 minutes
Performance	Every 24 hours (daily)

Performance Analysis and Tuning

As the number of resources collecting data increases, data collection completion times might become longer than the interval between data collection intervals, particularly for state data collection. To determine whether data collection for a compute resource or endpoint is completing in time or is being queued, see the Data Collection page. The Last Completed field value might show In queue or In progress instead of a timestamp when data collection last finished. If this problem occurs, you can increase the interval between data collections to decrease the data collection frequency.

Alternatively, you can increase the concurrent data collection limit per agent. By default, vRealize Automation limits concurrent data collection activities to two per agent and queues requests that exceed this limit. This limitation allows data collection activities to finish quickly without affecting overall performance. You can raise the limit to take advantage of concurrent data collection, but you must weigh this option against overall performance degradation.

If you increase the configured vRealize Automation per-agent limit, you might want to increase one or more of these execution timeout intervals. For more information about how to configure data collection concurrency and timeout intervals, see the vRealize Automation System Administration documentation. Manager Service data collection is CPU-intensive. Increasing the processing power of the Manager Service host can decrease the time required for overall data collection.

Data collection for Amazon Elastic Compute Cloud (Amazon AWS), in particular, can be CPU intensive, especially if your system collects data on multiple regions concurrently and if data was not previously collected on those regions. This type of data collection can cause an overall degradation in Web site performance. Decrease the frequency of Amazon AWS inventory data collection if it is having a noticeable effect on performance.

Workflow Processing Scalability

The average workflow processing time, from when the DEM Orchestrator starts preprocessing the workflow to when the workflow finishes executing, increases with the number of concurrent workflows. Workflow volume is a function of the amount of vRealize Automation activity, including machine requests and some data collection activities.

This chapter includes the following topics:

- [Configure Manager Service for High Data Volume](#)
- [Distributed Execution Manager Performance Analysis and Tuning](#)

Configure Manager Service for High Data Volume

If you expect to use a VMware vSphere cluster that contains a large number of objects, for example, 3000 or more virtual machines, modify the manager service config file with larger values. If you do not modify this setting, large inventory data collections might fail.

Modify the default value of the ProxyAgentServiceBinding and maxStringContentLength settings in the ManagerService.exe.config file.

Procedure

- 1 Open the ManagerService.exe.config file in a text editor.

Typically, this file resides at C:\Program Files (x86)\VMware\vCAC\Server.

- 2 Locate the binding name and readerQuotas lines in the file.

```
<binding name="ProxyAgentServiceBinding" maxReceivedMessageSize="13107200">
  <readerQuotas maxStringContentLength="13107200" />
```

Note Do not confuse these two lines with the similar lines that contain the following string: binding name = "ProvisionServiceBinding".

- 3 Replace the number values assigned to the maxReceivedMessageSize and maxStringContentLength attributes with a larger value.

The optimal size depends on how many more objects you expect your VMware vSphere cluster to contain in the future. For example, you can increase these numbers by a factor of 10 for testing.

- 4 Save your changes and close the file.
- 5 Restart the vRealize Automation manager service.

Distributed Execution Manager Performance Analysis and Tuning

You can view the total number of in progress or pending workflows at any time on the Distributed Execution Status page, and you can use the Workflow History page to determine how long it takes to run a given workflow.

If you have a large number of pending workflows, or if workflows are taking longer than expected to finish, add more Distributed Execution Manager (DEM) Worker instances to pick up the workflows. Each DEM Worker instance can process 30 concurrent workflows. Excess workflows are queued for execution.

You can adjust workflow schedules to minimize the number of workflows that start simultaneously. For example, rather than scheduling all hourly workflows to run at the beginning of the hour, you can stagger their run times so that they do not compete for DEM resources. For more information about workflows, see the vRealize Automation Extensibility documentation.

Some workflows, particularly certain custom workflows, can be CPU intensive. If the CPU load on the DEM Worker machines is high, consider increasing the processing power of the DEM machine or adding more DEM machines to your environment.

vRealize Business for Cloud Scalability

5

Configure your vRealize Business for Cloud installation for scalability in accordance with VMware guidelines.

vRealize Business for Cloud can scale up to 20,000 virtual machines across ten VMware vCenter Server instances. The first synchronization of the inventory data collection takes approximately three hours to synchronize 20,000 virtual machines across three VMware vCenter Server instances. Synchronization of statistics from VMware vCenter Server takes approximately one hour for 20,000 virtual machines. By default, the cost calculation job runs every day and takes approximately two hours for each run for 20,000 virtual machines.

Note In vRealize Business for Cloud 1.0, the default virtual appliance configuration can support up to 20,000 virtual machines. Increasing the limits of the virtual appliance beyond its default configuration does not increase the number of virtual machines that it can support.

vRealize Automation High Availability Configuration Considerations

6

If you require maximum system robustness, configure your vRealize Automation system for high availability in accordance with VMware guidelines.

vRealize Automation Appliance

The vRealize Automation appliance supports active-active high availability for all components except the appliance database. Starting with the 7.3 release, database failover is automatic if three nodes are deployed and synchronous replication is configured between two nodes. When vRealize Automation appliance detects database failure, it promotes a suitable database server to be the master. You can monitor and manage the appliance database on the Virtual Appliance Management Console **vRA Settings > Database** tab.

To enable high availability for these appliances, place them under a load balancer. For more information, see *Installing vRealize Automation*. Beginning with the 7.0 release, the appliance database, and vRealize Orchestrator are automatically clustered and available for use.

vRealize Automation Directories Management

Each vRealize Automation appliance includes a connector that supports user authentication, although only one connector is typically configured to perform directory synchronization. It does not matter which connector you choose to serve as the sync connector. To support Directories Management high availability, you must configure a second connector that corresponds to your second vRealize Automation appliance, which connects to your Identity Provider and points to the same Active Directory. With this configuration, if one appliance fails, the other takes over management of user authentication.

In a high availability environment, all nodes must serve the same set of Active Directories, users, authentication methods, etc. The most direct method to accomplish this is to promote the Identity Provider to the cluster by setting the load balancer host as the Identity Provider host. With this configuration, all authentication requests are directed to the load balancer, which forwards the request to either connector as appropriate.

For more information about configuring Directories Management for high availability, see *Configuring vRealize Automation*.

Infrastructure Web Server

The Infrastructure Web server components all support active-active high availability. To enable high availability for these components, place them under a load balancer.

Infrastructure Manager Service

The manager service component supports active-passive high availability. To enable high availability for this component, place two manager services under a load balancer. In vRealize Automation 7.3 and newer, failover is automatic.

If the active manager service fails, stop the Windows service, if it is not already stopped under the load balancer. Enable the passive manager service and restart the Windows service under the load balancer. See the *Installing vRealize Automation* documentation.

Agents

Agents support active-active high availability. For information about configuring agents for high availability, see the vRealize Automation configuration documentation. Check the target service for high availability.

Distributed Execution Manager Worker

A Distributed Execution Manager (DEM) running under the Worker role supports active-active high availability. If a DEM Worker instance fails, the DEM Orchestrator detects the failure and cancels workflows that the DEM Worker instance is running. When the DEM Worker instance comes back online, it detects that the DEM Orchestrator has canceled the workflows of the instance and stops running them. To prevent workflows from being canceled prematurely, leave a DEM Worker instance offline for several minutes before you cancel its workflows.

Distributed Execution Manager Orchestrator

DEMs running under the Orchestrator role support active-active high availability. When a DEM Orchestrator starts, it searches for another running DEM Orchestrator.

- If it finds no DEM Orchestrator instances running, it starts running as the primary DEM Orchestrator.
- If it does find another running DEM Orchestrator, it monitors the other primary DEM Orchestrator to detect an outage.
- If it detects an outage, it takes over as the primary instance.

When the previous primary instance comes online again, it detects that another DEM Orchestrator has taken over its role as primary and monitors for failure of the primary Orchestrator instance.

MSSQL Database Server for Infrastructure Components

vRealize Automation supports SQL AlwaysON groups only with Microsoft SQL Server 2016. When installing SQL Server 2016, the database must be created in 100 mode. If you use an older version of Microsoft SQL Server, use a Failover Cluster instance with shared disks. For more information on configuring SQL AlwaysOn groups with MSDTC, see Microsoft article <https://msdn.microsoft.com/en-us/library/ms366279.aspx>.

vRealize Orchestrator

An internal highly-available instance of vRealize Orchestrator is supplied as part of the vRealize Automation appliance.

vRealize Business for Cloud High Availability Considerations



Use the VMware vSphere HA feature for the vRealize Business for Cloud Edition appliance.

To configure the VMware vSphere HA feature on the VMware ESXi host, see the vCenter Server and Host Management documentation.

vRealize Automation Hardware Specifications and Capacity Maximums

8

Install appropriate components for your configuration and capacity needs on each vRealize Automation server profile in your environment.

Server Role	Components	Required Hardware Specifications	Recommended Hardware Specifications
vRealize Automation Appliance	vRealize Automation Services, vRealize Orchestrator, vRealize Automation Appliance Database	CPU: 4 vCPU RAM: 18 GB (See Chapter 4 vRealize Automation Scalability for more information.) Disk: 140 GB Network: 1 GB/s	Same as required hardware specifications.
Infrastructure Core Server	Web site, Manager Service, DEM Orchestrator, DEM Worker, Proxy Agent	CPU: 4 vCPU RAM: 8 GB Disk: 40 GB Network: 1 GB/s	Same as required hardware specifications.
Infrastructure Web Server	Web site	CPU: 2 vCPU RAM: 8 GB Disk: 40 GB Network: 1 GB/s	CPU: 2 vCPU RAM: 8 GB Disk: 40 GB Network: 1 GB/s
Infrastructure Manager Server	Manager Service, DEM Orchestrator	CPU: 2 vCPU RAM: 8 GB Disk: 40 GB Network: 1 GB/s	CPU: 2 vCPU RAM: 8 GB Disk: 40 GB Network: 1 GB/s
Infrastructure Web/Manager Server	Infrastructure Web/Manager Server	CPU: 2 vCPU RAM: 8 GB Disk: 40 GB Network: 1 GB/s	CPU: 2 vCPU RAM: 8 GB Disk: 40 GB Network: 1 GB/s
Infrastructure DEM Server	(One or more) DEM Workers	CPU: 2 vCPU RAM: 8 GB Disk: 40 GB Network: 1 GB/s Per DEM Worker	CPU: 2 vCPU RAM: 8 GB Disk: 40 GB Network: 1 GB/s Per DEM Worker

Server Role	Components	Required Hardware Specifications	Recommended Hardware Specifications
Infrastructure Agent Server	(One or more) Proxy Agent	CPU: 2 vCPU RAM: 8 GB Disk: 40 GB Network: 1 GB/s	CPU: 2 vCPU RAM: 8 GB Disk: 40 GB Network: 1 GB/s
MSSQL Database Server	Infrastructure Database	CPU: 2 vCPU RAM: 8 GB Disk: 40 GB Network: 1 GB/s	CPU: 8 vCPU RAM: 16 GB Disk: 80 GB Network: 1 GB/s
vRealize Business for Cloud Appliance	vRealize Business for Cloud Appliance services vRealize Business for Cloud Database Server	CPU: 2 vCPU RAM: 4 GB Disk: 50 GB Network: 1 GB/s	Same as required hardware specifications

vRealize Automation Recommended Capacity Maximums

The following resource capacity maximum values apply to the vRealize Automation large deployment profile.

Table 8-1. vRealize Automation Resource Capacity Maximums

Parameter	Maximum Value	
Tenant	100	
vSphere Endpoints	20	
Compute Resources	200	
Managed Machines	75,000	
Peak Concurrent Request		
	constant	50
	bursts	250
Peak requests per hour	400	
Business Groups	3000 (with 10 unique users per business group)	
Reservations	9000 (with 3 reservations per business group)	
Blueprints		
	CBP only	6000
	CBP + XaaS	8000
Catalog Items		
	across tenants	4000
	in a single tenant	6000

Table 8-1. vRealize Automation Resource Capacity Maximums (Continued)

Parameter	Maximum Value
User/Group sync with default 18 GB memory	
number of users	95027
number of groups	20403 (each group contains 4 users including one level of nesting)
User/Group with memory increased to 30 GB	
number of users	100,000
number of groups	750 (each group contains 4000 users and each user is in 30 groups)

vRealize Automation Small Deployment Requirements

9

A vRealize Automation small deployment comprises systems of 10,000 managed machines or fewer and includes the appropriate virtual machines, load balancers, and port configurations. The small deployment serves as a starting point for a vRealize Automation deployment that enables you to scale in a supported manner to a medium or large deployment.

When deploying vRealize Automation, use the Enterprise deployment process to provide a separate infrastructure Web site and Manager Service address.

Support

A small deployment can support the following items.

- 10,000 managed machines
- 500 catalog items
- 10 concurrent machine provisions

Requirements

A small deployment must be configured with the appropriate components.

- vRealize Automation appliance: vrava-1.ra.local
- Infrastructure Core server: inf-1.ra.local.
- MSSQL Database Server: mssql.ra.local
- vRealize Business for Cloud appliance: vrb.ra.local

DNS Entries

DNS Entry	Points To
vrava.ra.local	vrava-1.ra.local
web.ra.local	inf.ra.local
manager.ra.local	inf.ra.local

Certificates

The host names used in this table are examples only.

Server Role	CN or SAN
vRealize Automation appliance	SAN contains vra.va.sqa.local and vra.va-1.sqa.local
Infrastructure Core Server	SAN contains web.ra.local, managers.ra.local and inf-1.ra.local
vRealize Business for Cloud Server	CN = vrb.ra.local

Ports

Users require access to certain ports. All ports listed are default ports.

Server Role	Port
vRealize Automation appliance	443, 8444. Port 8444 is required for the Virtual Machine Remote Console. Port 8283 is required for access to the vRealize Orchestrator Control Center.

Administrators require access to certain ports, in addition to the ports that users require.

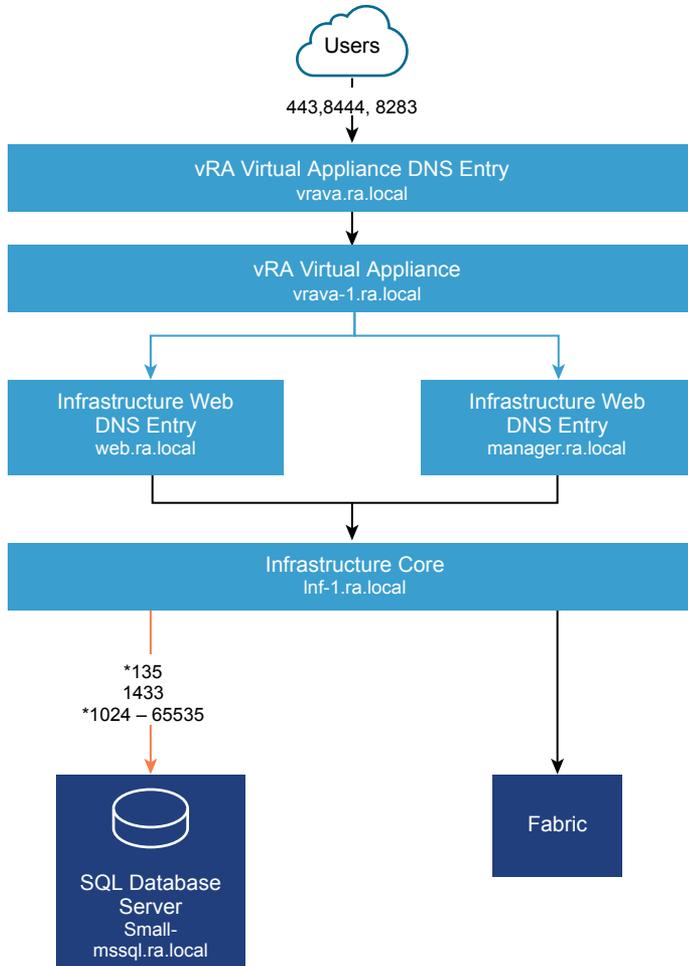
Server Role	Port
vRealize Automation appliance	5480, 8443. Port 8443 is used for advanced identity management configuration. VMware Identity Manager to Active Directory: 389, 636, 3268, 3269 VMware Identity Manager to Domain Controller: 88, 464, 135
vRealize Business for Cloud	5480

Server Role	Inbound Ports	Service/System Outbound Ports
vRealize Automation appliance	HTTPS: 443 Adapter Configuration: 8443 Remote Console Proxy: 8444 SSH: 22 Virtual Appliance Management Console: 5480	LDAP: 389 LDAPS:636 VMware ESXi: 902 Infrastructure Core requires access to vSphere endpoint port 443 to obtain a ticket for VMware Remote Console. The vRealize Automation appliance requires access to ESXi host port 902 to proxy traffic to the consumer. Infrastructure Core Server: 443 Kerberos Authentication: 88 Computer Object Password renewal: 464
Infrastructure Core Server	HTTPS: 443 MSDTC: 135, 1024 - 65535. For information about how to narrow this range, see the Database Deployment section of Chapter 2 vRealize Automation Deployment .	vRealize Automation virtual appliance: 443, 5480 vSphere Endpoint: 443 Infrastructure Core requires access to vSphere Endpoint Port 443 to obtain a ticket for VMware Remote Console. The vRealize Automation appliance requires access to ESXi host port 902 to proxy traffic to the consumer. MSSQL: 135, 1433, 1024 - 65535 MSDTC: 135, 1024 - 65535. For information about how to narrow this range, see the Database Deployment section of Chapter 2 vRealize Automation Deployment .

Server Role	Inbound Ports	Service/System Outbound Ports
MSSQL Database Server	MSSQL: 1433 MSDTC: 135, 1024 - 65535. For information about how to narrow this range, see the Database Deployment section of Chapter 2 vRealize Automation Deployment .	Infrastructure Core Server: 135, 1024 to 65535. For information about how to narrow this range, see the Database Deployment section of Chapter 2 vRealize Automation Deployment . MSDTC: 135, 1024 - 65535. For information about how to narrow this range, see the Database Deployment section of Chapter 2 vRealize Automation Deployment .
vRealize Business for Cloud Appliance	HTTPS: 443 SSH: 22 Virtual Appliance Management Console: 5480	vRealize Automation virtual appliance:443 Infrastructure Core:443
Global Catalog		Global Catalog: 3268, 3269

Minimum Footprints

Figure 9-1. Minimum footprint for small configuration of vRealize Automation




Not Shown:
All Infrastructure systems require access to Port 5480 of all vRealize Appliances for Log Collection (vRA Settings > Cluster > Collect Logs on Virtual Appliance:5480) to function.

For Virtual Machine Remote Console, vRealize Appliance requires access to VMware ESXi Port 902, and Infrastructure Core Server requires access to vSphere Endpoint Port 443.

*Please see Database Deployment section for information on how to narrow this range
In addition, bi-directional communication is required.

Figure 9-2. Minimum footprint for small configuration of vRealize Business for Cloud



vRealize Automation Medium Deployment Requirements

10

A vRealize Automation medium deployment comprises systems of 30,000 managed machines or fewer and includes the appropriate virtual machines, load balancers, and port configurations.

Support

A medium deployment can support the following items.

- 30,000 managed machines
- 1000 catalog items
- 50 machine provisions

Requirements

A medium deployment must meet the appropriate system configuration requirements.

Virtual Appliances

- vRealize Automation appliance 1: vrava-1.ra.local
- vRealize Automation appliance 2: vrava-2.ra.local
- vRealize Automation appliance 3: vrava-3.ra.local
- vRealize Business for Cloud Appliance: vrb.ra.local

Windows Server Virtual Machines

- Infrastructure Web/Manager Server 1 (Active Web or DEM-O, Active Manager): inf-1.ra.local
- Infrastructure Web/Manager Server 2 (Active Web or DEM-O, Passive Manager): inf-2.ra.local
- Infrastructure DEM Server 1: dem-1.ra.local
- Infrastructure DEM Server 2: dem-2.ra.local
- Infrastructure Agent Server 1: agent-1.ra.local
- Infrastructure Agent Server 2: agent-2.ra.local

Database Servers

- MSSQL Failover Cluster Instance: mssql.ra.local

Load Balancers

- vRealize Automation appliance Load Balancer: med-vrava.ra.local
- Infrastructure Web Load Balancer: med-web.ra.local
- Infrastructure Manager Service Load Balancer: med-manager.ra.local

Certificates

The host names that are used in this table are examples only.

Server Role	CN or SAN
vRealize Automation appliance	SAN contains the following host names: <ul style="list-style-type: none"> ▪ vrava.ra.local ▪ vrava-1.ra.local ▪ vrava-2.ra.local
Infrastructure Web or Manager Server	SAN contains the following host names: <ul style="list-style-type: none"> ▪ web.ra.local ▪ manager.ra.local ▪ inf-1.ra.local ▪ inf-2.ra.local
vRealize Business for Cloud Appliance	CN = vrb.ra.local

Ports

Users require access to certain ports. All ports listed are default ports.

Server Role	Port
vRealize Automation appliance Load Balancer	443, 8444. Port 8444 is required for the Virtual Machine Remote Console.

Administrators require access to certain ports, in addition to the ports that users require.

Server Role	Port
vRealize Automation appliance fVAMI	5480, 8443. Port 8443 is for advanced identity management configuration. VMware Identity Manager to Active Directory: 389, 636, 3268, 3269 VMware Identity Manager to Domain Controller: 88, 464, 135
vRealize Appliance Orchestrator Control Center	8283
vRealize Business for Cloud Server	5480

The following table shows inter-application communications.

Server Role	Inbound Ports	Outbound Ports for Service or System
vRealize Automation appliance	HTTPS: Adapter Configuration: 8443 Remote Console Proxy: 8444 Postgres: 5432 RabbitMQ: 4369, 25672, 5671, 5672 ElasticSearch: 9300, 40002, 40003 Stomp: 61613 SSH: 22	LDAP:389 LDAPS: 636 vRealize AutomationAppliance (All other): 5432, 4369, 25672, 5671, 5672, 9300, 40002, 40003 vRealize Automation Infrastructure Web Load Balancer: 443 VMware ESXi: 902. Infrastructure Web or Manager requires access to vSphere Endpoint port 443 to obtain a ticket for Virtual Machine Remote Console. The vRealize Automation appliance requires access to ESXi host port 902 to proxy console data to the user. Kerberos Authentication: 88 Computer Object Password renewal: 464
Infrastructure Web/Manager Server	HTTPS: 443 MSDTC: 135, 1024-65535. For information about how to narrow this range, see the Database Deployment section of Chapter 2 vRealize Automation Deployment .	vRealize Automation appliance Load Balancer: 443 vRealize Automation Infrastructure Web Load Balancer: 443 vRealize Automation Appliance (VA): 5480. vSphere Endpoint: 443. Infrastructure Web or Manager requires access to vSphere Endpoint port 443 to obtain a ticket for Virtual Machine Remote Console. The vRealize Automation appliance requires access to ESXi host port 902 to proxy console data to the user. MSSQL: 135, 1433, 1024 to 65535. For information about how to narrow this range, see the Database Deployment section of Chapter 2 vRealize Automation Deployment .
Infrastructure DEM Server	NA	vRealize Automation Appliance Load Balancer: 443 vRealize Automation Infrastructure Web Load Balancer: 443 vRealize Automation Infrastructure Manager Load Balancer: 443 vRealize Automation Appliance (VA): 5480.

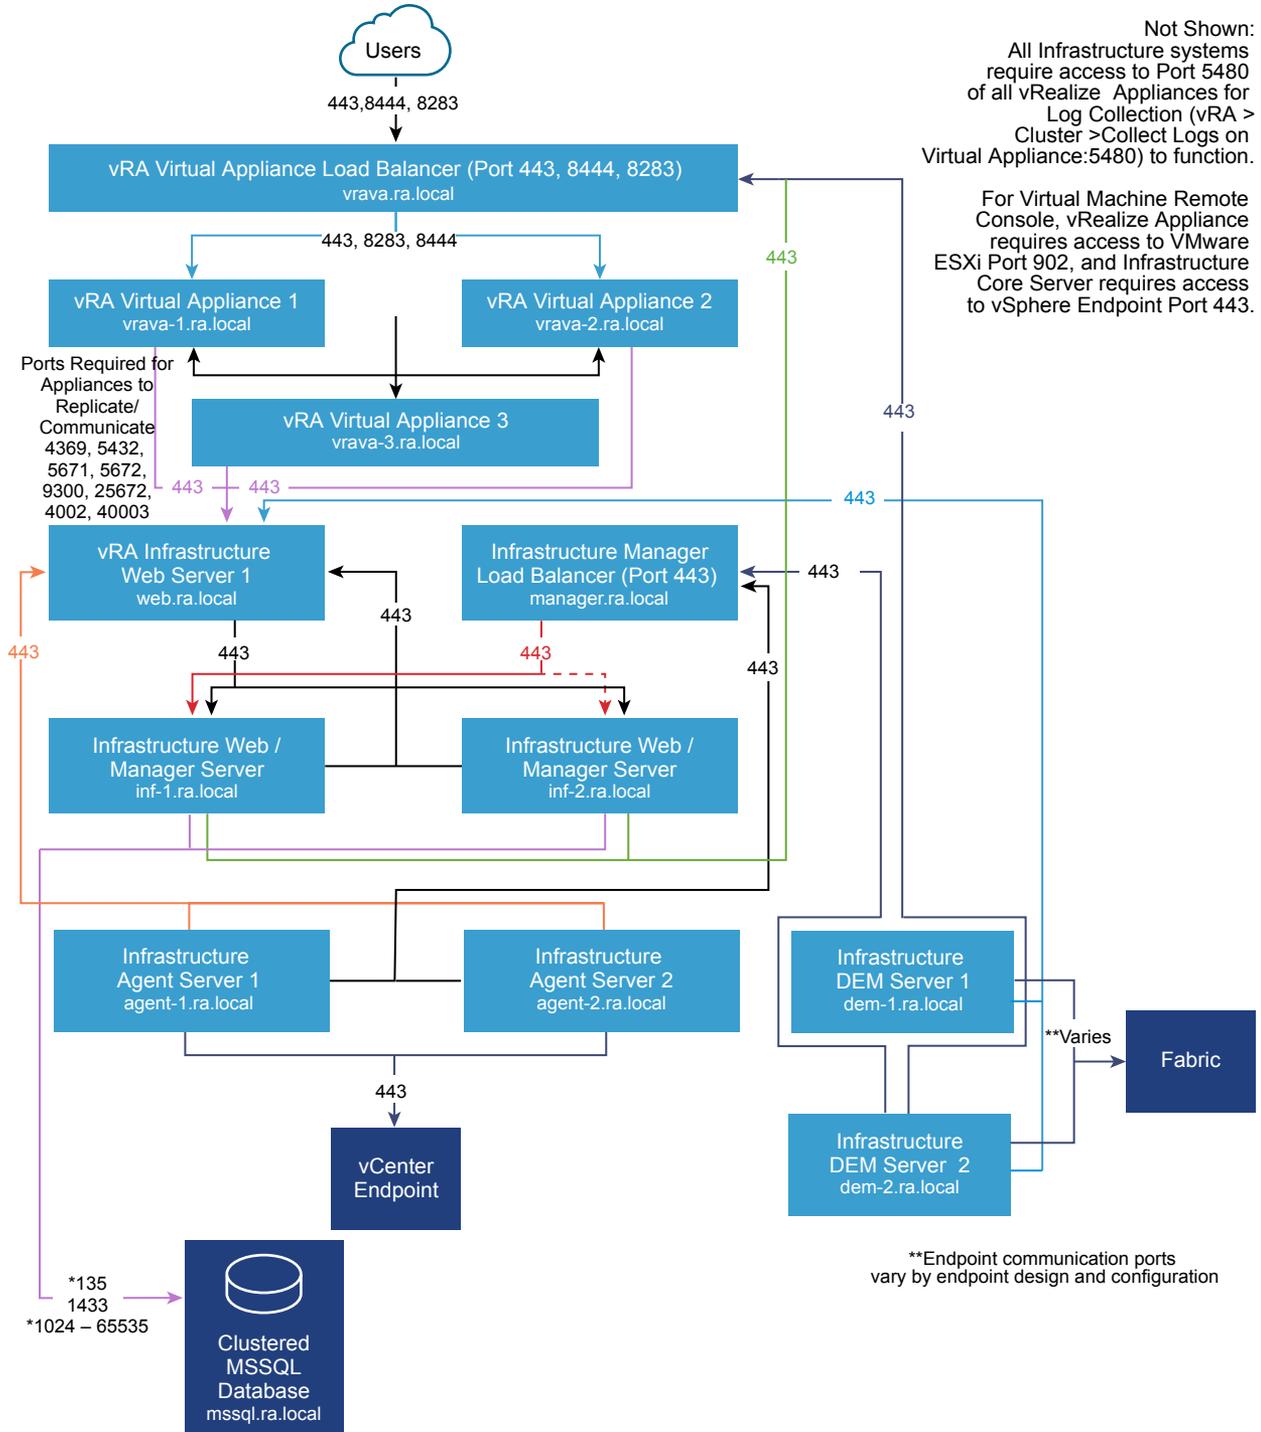
Server Role	Inbound Ports	Outbound Ports for Service or System
Infrastructure Agent Server	NA	vRealize Automation Infrastructure Web Load Balancer: 443 vRealize Automation Infrastructure Manager Load Balancer: 443 vRealize Automation Appliance (VA): 5480.
MSSQL Database Server	MSSQL: 1433 MSDTC: 135, 1024 - 65535. For information about how to narrow this range, see the Database Deployment section of Chapter 2 vRealize Automation Deployment .	Infrastructure Web/Manager Server: 135, 1024 - 65535. For information about how to narrow this range, see the Database Deployment section of Chapter 2 vRealize Automation Deployment .
vRealize Business for Cloud Server	HTTPS: 443 SSH: 22 Virtual Appliance Management Console: 5480	vRealize Automation Appliance Load Balancer: 443 vRealize Automation Infrastructure Web Load Balancer: 443
Global Catalog		Global Catalog: 3268, 3269

Load balancers require access through the following ports.

Load Balancer	Ports Balanced
vRealize Automation appliance Load Balancer	443, 8444
vRealize Automation Infrastructure Web Load Balancer	443
vRealize Automation Infrastructure Manager Service Load Balancer	443

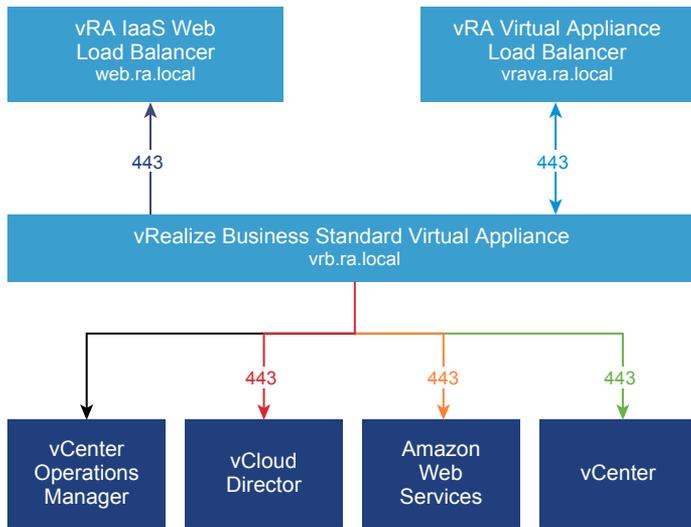
Graphics

Figure 10-1. Minimum footprint for vRealize Automation medium configuration



*Please see Database Deployment section for information on how to narrow this range
 In addition, bi-directional communication is required.

Figure 10-2. Minimum footprint for vRealize Business for Cloud medium deployment



vRealize Automation Large Deployment Requirements

11

A vRealize Automation large deployment comprises systems of 50,000 managed machines or fewer and includes the appropriate virtual machines, load balancers, and port configurations.

Support

A large deployment can support the following items.

- 50,000 managed machines
- 2500 catalog items
- 100 concurrent machine provisions

Requirements

A large deployment must meet the appropriate system configuration requirements.

Virtual Appliances

- vRealize Automation appliance 1: vrava-1.ra.local
- vRealize Automation appliance 2: vrava-2.ra.local
- vRealize Automation appliance 2: vrava-3.ra.local
- vRealize Automation appliance Appliance: vrb.ra.local

Windows Server Virtual Machines

- Infrastructure Web Server 1: web-1.ra.local
- Infrastructure Web Server 2: web-2.ra.local
- Infrastructure Manager Server 1: manager-1.ra.local
- Infrastructure Manager Server 2: manager-2.ra.local
- Infrastructure DEM Server 1: dem-1.ra.local
- Infrastructure DEM Server 2: dem-2.ra.local
- Infrastructure Agent Server 1: agent-1.ra.local
- Infrastructure Agent Server 2: agent-2.ra.local

- Clustered MSSQL Database: mssql.ra.local

Load Balancers

- vRealize Automation appliance load balancer: vrava.ra.local
- Infrastructure Web load balancer: web.ra.local
- Infrastructure manager service load balancer: manager.ra.local

Certificates

The host names used in this table are examples only.

Server Role	CN or SAN
vRealize Automation appliance	SAN contains the following host names: <ul style="list-style-type: none"> ▪ vrava.ra.local ▪ vrava-1.ra.local ▪ vrava-2.ra.local
Infrastructure Web server	SAN contains the following host names: <ul style="list-style-type: none"> ▪ web.ra.local ▪ web-1.ra.local ▪ web-2.ra.local
Infrastructure manager server	SAN contains the following host names: <ul style="list-style-type: none"> ▪ manager.ra.local ▪ manager-1.ra.local ▪ manager-2.ra.local
vRealize Business for Cloud appliance	CN = vrb.ra.local

Ports

Users require access to certain ports. All ports listed are default ports.

Server Role	Port
vRealize Automation appliance load balancer	443, 8444 Port 88444 is required for the VMware Remote Console.

Administrators require access to certain ports, in addition to the ports that users require.

Server Role	Port
vRealize Automation appliance	5480, 8443. Port 8443 is used for advanced identity management configuration. VMware Identity Manager to Active Directory: 389, 636, 3268, 3269 VMware Identity Manager to Domain Controller: 88, 464, 135
vRealize Business for Cloud server	5480

The system must support the appropriate inter-application communications.

Server Role	Inbound Ports	Outbound Ports for Service or System
vRealize Automation		
vRealize Automation appliance	HTTPS: 443 Adapter configuration: 8443 Remote console proxy: 8444 Postgres: 5432 Rabbit MQ: 4369, 25672, 5671, 5672 ElasticSearch: 9300, 40002, 40003 Stomp: 61613 SSH: 22 Control-Center: 8283	LDAP: 389 LDAPS: 636 vRealize Automation appliance: 5432, 4369, 25672, 5671, 5672, 9300, 40002, 40003. vRealize Automation infrastructure Web load balancer: 443 VMware ESXi: 902. Infrastructure Web requires access to vSphere endpoint port 443 to obtain a ticket for VMware Remote Console. The vRealize Automation appliance requires access to ESXi host Port 902 to proxy console data to the user. Kerberos Authentication: 88 Computer Object Password renewal: 464
Infrastructure Web server	HTTPS: 443 MSDTC: 443, 1024-65535. For information about how to narrow this range, see the database deployment section of Chapter 2 vRealize Automation Deployment .	vRealize Automation appliance load balancer: 443 vRealize Automation appliance virtual appliance: 5480. vSphere endpoint: 443. Infrastructure Web requires access to vSphere endpoint port 443 to obtain a ticket for VMware Remote Console. The vRealize Automation appliance requires access to ESXi host port 902 to proxy console data to the user. MSSQL: 135, 1433, 1024 to 65535. For information about how to narrow this range, see the database deployment section of Chapter 2 vRealize Automation Deployment .
Infrastructure manager server	HTTPS: 443 MSDTC: 135, 1024-65535. For information about how to narrow this range, see the database deployment section of Chapter 2 vRealize Automation Deployment .	vRealize Automation appliance load balancer: 443 vRealize Automation Infrastructure Web Load Balancer: 443 vRealize Automation appliance: 443, 5480 MSSQL: 135, 1433, 1024 to 65535. For information about how to narrow this range, see the database deployment section of Chapter 2 vRealize Automation Deployment .

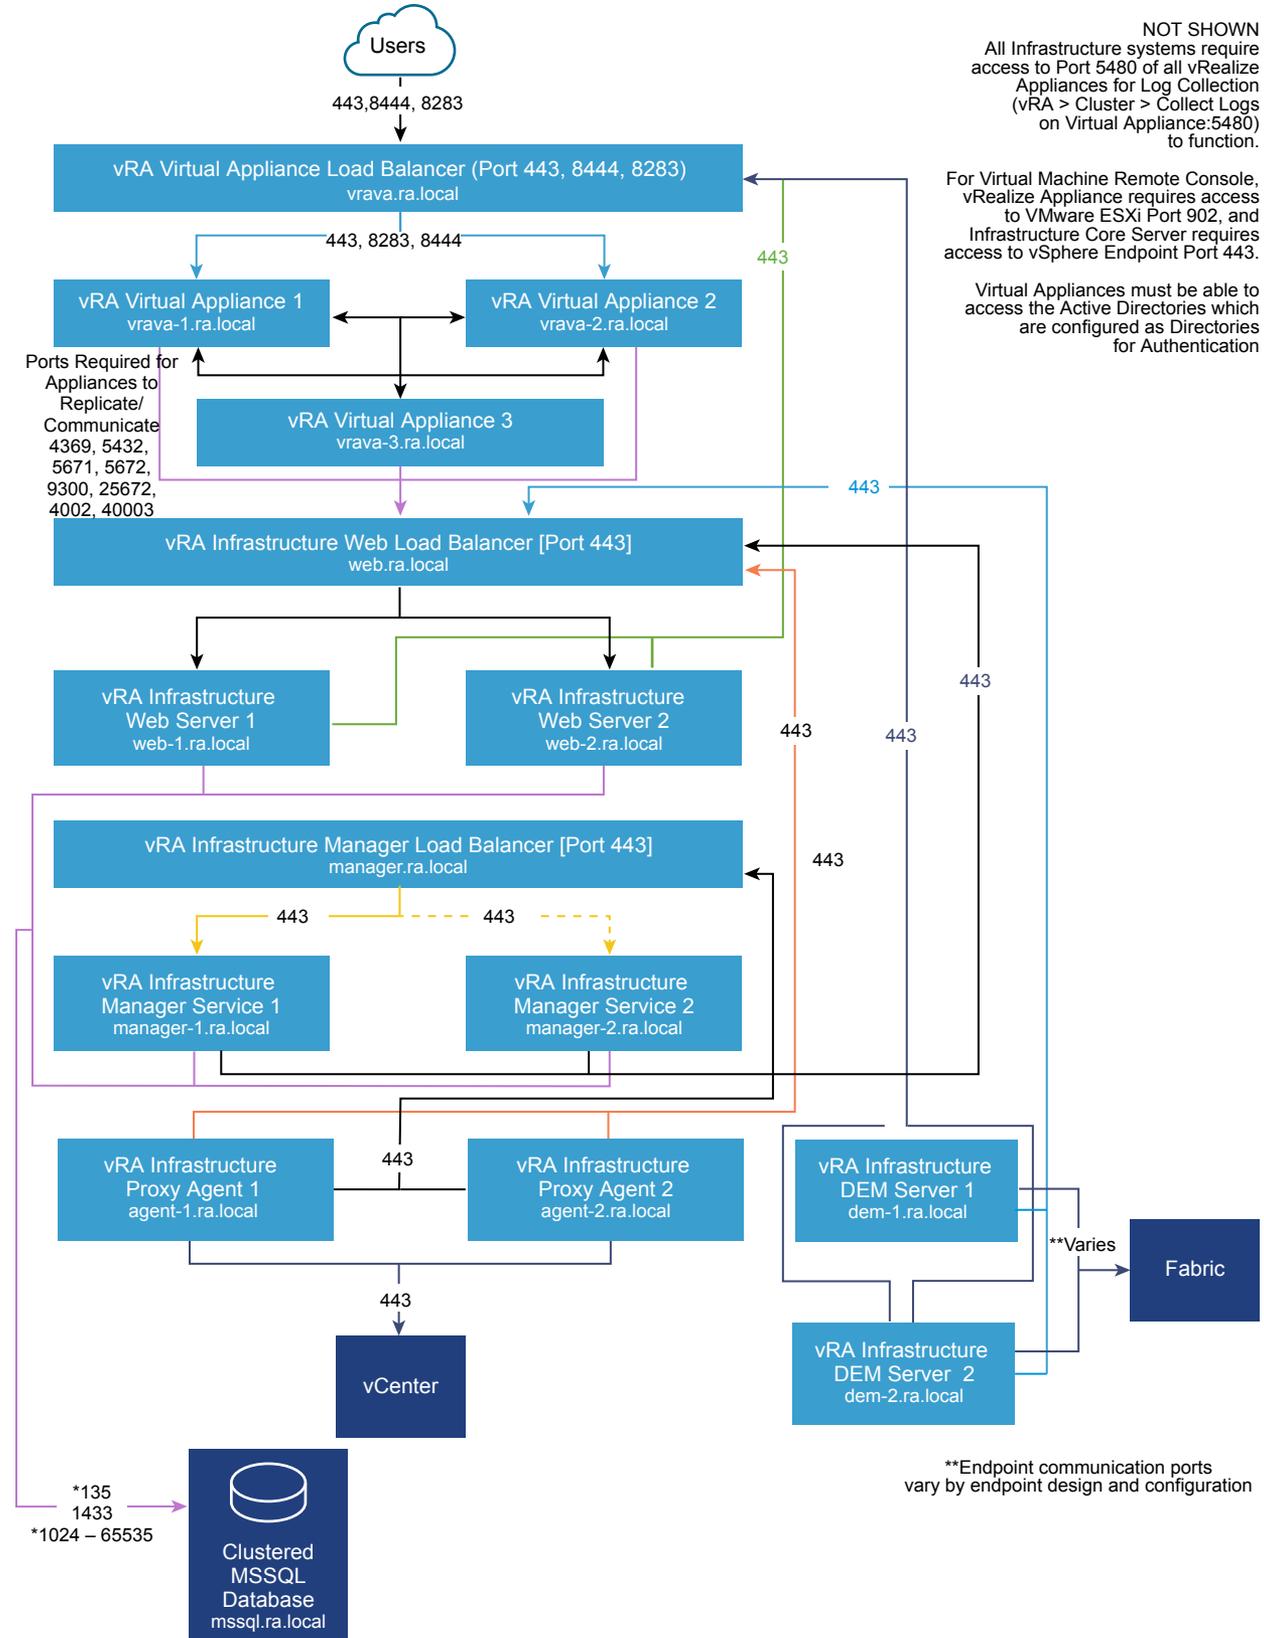
Server Role	Inbound Ports	Outbound Ports for Service or System
Infrastructure DEM server	NA	vRealize Automation appliance load balancer: 443 vRealize Automation infrastructure Web load Balancer: 443 vRealize Automation infrastructure manager load balancer: 443 vRealize Orchestrator load balancer: 8281 vRealize Automation appliance: 5480.
Infrastructure agent server	NA	vRealize Automation infrastructure Web load balancer: 443 vRealize Automation infrastructure manager load balancer: 443 vRealize Automation appliance: 5480.
MSSQL database server	MSSQL: 1433 MSDTC: 135, 1024-65535. For information about how to narrow this range, see the database deployment section of Chapter 2 vRealize Automation Deployment .	Infrastructure Web server: 135, 1024-65535. For information about how to narrow this range, see the database deployment section of Chapter 2 vRealize Automation Deployment . Infrastructure manager server: 135, 1024-65535. For information about how to narrow this range, see the database deployment section of Chapter 2 vRealize Automation Deployment .
vRealize Business for Cloud server	HTTPS: 443 SSH: 22 Virtual Appliance Management Console: 5480	vRealize Automation appliance load balancer: 443 vRealize Automation infrastructure Web load balancer: 443
Global Catalog		Global Catalog: 3268, 3269

Load balancers require access through the following ports.

Load Balancer	Ports Balanced
vRealize Automation appliance load balancer	443, 8444
vRealize Automation infrastructure Web load balancer	443
vRealize Automation manager server load balancer	443

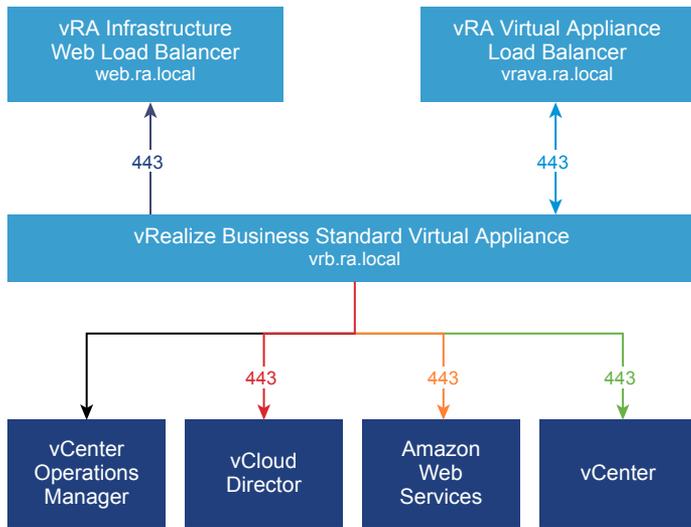
Graphics

Figure 11-1. Minimum footprint for vRealize Automation large configuration



*Please see Database Deployment section for information on how to narrow this range VMware, Inc.
In addition, bi-directional communication is required.

Figure 11-2. Minimum footprint for vRealize Business for Cloud large configuration



vRealize Automation Multi-Data Center Data Deployments

12

vRealize Automation supports managing resources in remote data centers.

To manage vSphere, HyperV, or Xen resources in remote data centers, deploy the proxy agent on a virtual machine in the remote data center.

Note The diagram below shows a vSphere deployment. Other endpoints require no additional configuration.

Because vRealize Orchestrator workflows will potentially communicate over a WAN, observe best practices as stated in the *vRealize Orchestrator Coding Design Guide*.

Table 12-1. Required Ports for WAN Communication

Role	Inbound Ports	Service/System Outbound Ports
vRealize Automation appliance - including embedded vRealize Orchestrator	N/A	vSphere endpoint: 443 ESXi Hosts: 903
vRealize Automation Infrastructure Load Balancer	vRealize Automation Infrastructure Proxy Agent: 443	N/A
vRealize Automation Infrastructure Web Server	N/A	vSphere endpoint: 443
vRealize Automation Infrastructure Manager Load Balancer	vRealize Automation Infrastructure Proxy Agent: 443	N/A
vRealize Automation Infrastructure DEM-worker Servers	N/A	Endpoint: **varies

* If DEM-workers are installed on the Manager Service machine or another server, these ports must be open between that machine and the target endpoint.

** The port required to a communicate with an external endpoint varies depending on the endpoint. By default for vSphere, this is port 443.

Figure 12-1. vRealize Automation Multi-Site Configuration

