Minimum System Requirements For An AtScale Installation

This section describes the requirements for dedicated nodes for an AtScale installation, regardless of the data platform that you will access via AtScale.

Before You Begin

- Running Clustered AtScale requires the external load balancer and three AtScale machines, referred to as follows:
  - Coordinator Host - runs the AtScale Coordinator
  - Master Host (AtScale Application Host) - runs full AtScale (only designated "Master" during installation but not at run-time)
  - Standby Host (AtScale Application Host) - runs full AtScale (only designated "Standby" during installation but not at run-time)
- AtScale requires a number of ports to be available, and no other services should run on those ports. See below for the complete list of ports.

About This Task

Provision the AtScale hosts and open ports as described:

AtScale Application Host Specification

Property	AtScale Application Host System Requirements
64-bit Linux Operating Systems	Red Hat Enterprise Linux (RHEL) 8 Ubuntu 20.04 Ubuntu 22.04
Software and OS Utilities	curl, rsync, unzip, sed, awk, ping (use iputils-ping if Debian or Ubuntu), whoami, hostname, find, cp, mv, grep, sort, date, tail, timeout, ps, env, which, cat, libreadline-dev (required by Ubuntu 20.x)
Memory	128 GB recommended, 64 GB minimum
CPU	32 cores recommended, 16 cores minimum
Disk	200 GB (for AtScale installation and metadata catalog)
Network	1 Gbps reliable network connectivity between the AtScale server and the data warehouse Fully Qualified Domain Name (FQDN) host addresses should be used rather than IP addresses. Each FQDN must have a valid DNS entry and resolve to a valid IP address. Reverse DNS must be set up for your network. AtScale's master database will not start properly in a clustered installation without access to Reverse DNS. In the case of AtScale Clusters, elastic scaling is not supported, therefore your cluster configuration must consist of a static list of host names with working DNS entries. Changing the host membership of the cluster requires down-time.
OS Limits	Use ulimit to remove the atscale user limits for the following: open files - unlimited or 65536. cpu time - unlimited virtual memory - unlimited max user processes - unlimited or 65536. file size - unlimited Be sure to make the limit changes so that they survive a restart. Failure to raise these limits will prevent various AtScale services from starting.

AtScale Coordinator Host Specification

Property	AtScale Coordinator Node System Requirements
64-bit Linux Operating Systems	Red Hat Enterprise Linux (RHEL) 8 Ubuntu 20.04 Ubuntu 22.04
Software and OS Utilities	curl, rsync, unzip, sed, awk, ping (use iputils-ping if Debian or Ubuntu), whoami, hostname, find, cp , mv, grep, sort, date, tail, timeout, ps, env, which, cat
Memory	16 GB minimum
CPU	4 cores minimum
Disk	200 GB
Network	1 Gbps reliable network connectivity between the AtScale server and the data warehouse Fully Qualified Domain Name (FQDN) host addresses should be used rather than IP addresses. Each FQDN must have a valid DNS entry and resolve to a valid IP address. Reverse DNS must be set up for your network. AtScale's master database will not start properly in a clustered installation without access to Reverse DNS. In the case of AtScale Clusters, elastic scaling is not supported, therefore your cluster configuration must consist of a static list of host names with working DNS entries. Changing the host membership of the cluster requires down-time.
OS Limits	Use ulimit to remove the atscale user limits for the following: open files - unlimited or 65536. cpu time - unlimited virtual memory - unlimited max user processes - unlimited or 65536. file size - unlimited Be sure to make the limit changes so that they survive a restart. Failure to raise these limits will prevent various AtScale services from starting.

AtScale Virtualization Layer Host Specification

Property	AtScale Virtualization Layer Host System Requirements
64-bit Linux Operating Systems	Red Hat Enterprise Linux (RHEL) 8 Ubuntu 20.04 Ubuntu 22.04
Software and OS Utilities	curl, rsync, unzip, sed, awk, ping (use iputils-ping if Debian or Ubuntu), whoami, hostname, find, cp, mv, grep, sort, date, tail, timeout, ps, env, which, cat
Memory	32 GB minimum AtScale recommends allocated a maximum of 75% of memory for virtualization, leaving the remaining for the operating system and buffer cache. Recommended memory allocation depends on application use. To determine how much your application uses for a certain dataset size, use the following calculation: `Concurrent Users * Activity Rate * (2 * Average Row Size) * Average Rows Per Query Returned * 200% Join Factor * (200% of Average Query Complexity [1-2])` Average query complexity is relative to the use of features such as time relative calculations or calculations requiring MultiPass.
CPU	8-16 core recommended, 4 cores minimum Recommended memory allocation depends on application use. To determine how much your application uses for a certain dataset size, use the following calculation: `Concurrent Users X Activity Rate X 1 Cores Per Query X  Average Query Complexity [1-2]` Average query complexity is relative to the use of features such as Time Relative calculations, or calculations that require MultiPass.
Disk	200 GB minimum (for AtScale installation and metadata catalog)|br| AtScale recommends 4-8 disks per node configured without RAID as separate mount points. (Linux: Mount disks with the `noatime` Linux mount option.)
	1 GBPS reliable network connectivity minimum, 10 GBPS recommended Fully Qualified Domain Name (FQDN) host addresses should be used rather than IP addresses.

Network	Each FQDN must have a valid DNS entry and resolve to a valid IP address. Reverse DNS must be set up for your network. AtScale's master database will not start properly in a clustered installation without access to Reverse DNS. In the case of AtScale Clusters, elastic scaling is not supported, therefore your cluster configuration must consist of a static list of host names with working DNS entries. Changing the host membership of the cluster requires down-time.
OS Limits	Use ulimit to remove the atscale user limits for the following: open files - unlimited or 65536. cpu time - unlimited virtual memory - unlimited max user processes - unlimited or 65536. file size - unlimited Be sure to make the limit changes so that they survive a restart. Failure to raise these limits will prevent various AtScale services from starting.

- For more details on hardware provisioning for virtualization, see Virtualization Hardware Provisioning.

Reserve The Following Ports For AtScale Services:

Use this table to configure your Data Center's firewall.

- **IN:** Traffic originating from the end client/outside the cluster with a destination into the AtScale cluster.
- **OUT:** Traffic originating from AtScale with the destination outside the AtScale cluster.

Service	Description	Node Type(s) for Running Service	Direction	Default Port
engine	AtScale engine (HTTP, XMLA)	Application and Engine	IN	10502
engine	AtScale OLAP engine listener service (SQL)	Application and Engine	IN	11111-11119
engine	Communication to data sources	Application and Engine	OUT	JDBC Data Source-Specific

engine	Communication to LDAP	Application and Engine	OUT	LDAP-Specific
engine	Communication to KDC	Application and Engine	OUT	KDC-Specific
modeler	AtScale Design Center web application service	Application and Engine	IN	10500
modeler	Authorization messages	Application and Engine	IN	10503
data agent	Communication to data sources	Virtualization Listener, Virtualization Worker	OUT	JDBC Data Source-Specific
data agent	Communication to LDAP	Virtualization Listener, Virtualization Worker	OUT	LDAP-Specific
data agent	Communication to KDC	Virtualization Listener, Virtualization Worker	OUT	KDC-Specific
orchestrator	Node provisioning and monitoring		IN	10540
orchestrator UI	Node provisioning and monitoring		IN	10510
query-engine [1]	AtScale horizontal scaling	Application and Engine	IN	10541

[1] **License Protected Feature:** This service is not enabled by default and requires a special license to activate.

Internal Communications Within The AtScale Cluster:

Use the following table to configure internal Data Center firewalls, cross Data Center port forwarding, or validate that AtScale servers are operating with reasonable behavior.

- **IN:** Traffic originating from another AtScale node with a local destination.
- **LOCAL:** Traffic only within the local node.
- **OUT:** Traffic originating from the node to another AtScale node.
- **INTERNAL:** Should only be used for debugging purposes.

Services Running

Service	Description	Direction	Port
engine	Modeler requests	IN	10502
engine	Engine to engine communication	IN, OUT	10507
engine	AtScale authorization	LOCAL, OUT	10503
engine	Communicate with Virtualization Listener	OUT	10511
engine	Coordinator requests	LOCAL	10513
engine	ping	IN	10587
engine	ping_rtr	OUT	10588
engine	Egress HTTP	OUT	10582
engine	JMX	IN	13502
	Authorization		

modeler	messages	LOCAL, IN	10503
modeler	Authorization messages	LOCAL, OUT	10583
modeler	Modeler egress	LOCAL, OUT	10580
orchestrator	Node provisioning and monitoring, UI egress	OUT	10591
orchestrator	Node provisioning and monitoring, Auth egress	OUT	10583
servicecontrol	Web application for managing services	INTERNAL	10516
coordinator	Connect messages	IN, OUT	10505
coordinator	Election messages	IN, OUT	10506
coordinator	Client requests	IN	10513
database	Query requests	LOCAL, IN, OUT	10520
database	Leader detection service	LOCAL, IN, OUT	10519
egress	AtScale database leader service	IN, OUT	10518
egress	Database queries	LOCAL, OUT	10520
virtualization_supervisor [1]	Driver messages listener	IN	10531

virtualization_supervisor [1]	Manager port	IN	10532
virtualization_supervisor [1]	Master port	IN	10533
virtualization_supervisor [1]	Supervisor Web Port	INTERNAL	10534
agent	Configuration notification	LOCAL	Socket only
data_agent	Data segment listener	LOCAL	10502
data_agent	JMX	IN, INTERNAL	13502
data_agent	Remoting	IN, INTERNAL	10507
ingress	Configuration	IN	10525
ingress	Health	IN	10526
ingress	Coordinator messages	OUT	10513
virtualization_worker [1]	Worker port	IN	10535
virtualization_worker [1]	Worker web port	INTERNAL	10536
virtualzation_listener [1]	Listener service	IN	10511
virtualzation_listener [1]	Listener UI port	IN	10530
directory	Non-prod	IN	10514

	authorization		
directory	Non-prod authorization	OUT	10524
proxy stats	proxy stats	OUT	10517
query-engine [1]	AtScale horizontal scaling, jmx	IN	13541
query-engine [1]	AtScale horizontal scaling, remoting	IN	10507
gov_rules	Governance Rules Service	IN	10545
gov_rules	Governance Rules Service	OUT	10546
gov_enforcer	Governance Enforcer Service	IN	10547
gov_enforcer	Governance Enforcer Service	OUT	10548
service_registry	dns	IN	10554
service_registry	grpc	IN	10557
service_registry	http	IN	10555
service_registry	https	IN	10556
service_registry	sef_lan	IN	10551
service_registry	serf_wan	IN	10552

service_registry	server	IN	10550

[1] **License Protected Feature:** This service is not enabled by default and requires a special license to activate.

Hadoop Cluster Requirements For AtScale (Only Required When Connecting To A Hadoop Data Warehouse)

When AtScale is connected to a Hadoop cluster, the following permissions and cluster resources are required:

Resource	Requirement
Cluster disk space	Disk space in the Hive warehouse for AtScale smart aggregate tables. Allocate roughly 30% of the raw data size for storage. AtScale aggregate tables are created per cube in a designated Hive database.
Cluster permissions	The `atscale` user needs: read permissions to source data files; ownership of a dedicated user home directory in HDFS or MapR-FS; read and write permissions on the AtScale aggregate and UDAF schemas. If enabling AtScale impersonation then grant the read permission on the AtScale UDAF schema to all Active Directory users who will run queries through AtScale.
Cluster resources	Minimum open file limit = 5000