



ScaleArc for SQL Server 3.11.0.4

Release Notes

Trademarks

© ScalArc, Inc. 2017

© ScaleArc 2017

ALL RIGHTS RESERVED.

THIS DOCUMENT IS PROPRIETARY TO SCALARC, INC., AND MAY NOT BE REVISED, MODIFIED OR OTHERWISE ALTERED WITHOUT SCALEARC'S AUTHORIZATION. ALL REPRODUCTIONS OF THIS DOCUMENT MUST NOT REMOVE OR REDUCE THE VISIBILITY OF ANY PROPRIETARY NOTICES.

ScaleArc®, ScaleArc iDB®, and ScalArc® are registered trademarks of ScalArc, Inc and/or its affiliates in the U.S, India and certain other countries. Other names and marks used herein are trademarks or service marks of their respective owners.

Disclaimer

The information disclosed in this document, including all designs and related materials, is the valuable property of ScalArc, Inc. ScaleArc reserves all copyrights and other proprietary rights to this document, including all design, manufacturing, reproduction, use, and sales rights thereto, except to the extent said rights have been expressly granted to others. The product(s) discussed in this document are licensed in accordance with the terms of ScaleArc's End User Software License Agreement. The information in this document is subject to change at any time, without notice.

Customer Support

Contact ScaleArc sales at sales@scalearc.com or call +1-408-837-2250. You can contact technical support at support@scalearc.com or call +1-855-800-SCAL(7225).

Contents

This release note covers the following topics:

- [Installation Information](#)—This section includes information for installation and upgrade of ScaleArc for SQL Server 3.11.0.4 release.
- [New Features](#) – This section lists the new features in this release.
- [Fixed Issues](#) – This section lists issues that have been fixed since the last release.
- [Known Issues](#) - This section lists existing known issues and provides solutions.
- [Highlight Behaviors](#) - This section highlights changes/ behaviors impacting end users in this release.
- [Limitations](#) – This section lists limitations in ScaleArc for SQL Server 3.11.0.4.
- [Troubleshooting Kerberos](#) – The section outlines troubleshooting techniques for Kerberos related issues.
- [Additional Resources](#) - This section lists additional resources for product training and documentation.

Installation Information

Getting started with ScaleArc for SQL Server 3.11.0.4 is fairly straight forward. Refer to the [Minimum/recommended system requirement to deploy ScaleArc server](#) article for more information on the prerequisites for deployment.

Refer to the [Upgrading ScaleArc](#) article for information on upgrading to ScaleArc v3.11.0.4 from a previous version of ScaleArc.

Visit ScaleArc support portal at <https://support.scalearc.com> for additional help articles.

New Features

This section lists the new features in ScaleArc for SQL Server 3.11.0.4

<i>Feature</i>	<i>Description</i>
Installation	ScaleArc now supports Google Cloud Platform deployments.

Fixed Issues

This section lists the issue that has been fixed since ScaleArc for SQL Server 3.11.0.3.

<i>Issue ID</i>	<i>Component</i>	<i>Description</i>
17150	Analytics	Delay in analytics is causing issues while trying to add cache patterns.

17249 17251	Analytics	A blank pop up screen is displayed in analytics unique query page even though data exists for the selected timeframe.
17344	Connection Management	For a cluster which has Kerberos authentication turned ON, ScaleArc does not put connections directly in passthrough. The authentication for the connection is offloaded by ScaleArc before putting the connection in passthrough.

Known Issues

This section provides a list of issues that are categorized as important.

<i>Issue ID</i>	<i>Component</i>	<i>Description</i>	<i>Solution</i>
5607	Caching	If the maximum query response size for cache is set to a large value, then the cache generation takes time.	Do not set the maximum query response size for cache to more than 5MB.
8550	Query Firewall Rules	ScaleArc may inadvertently delete DB firewall pattern while deleting DB Write Ignore pattern.	If a firewall pattern exists for a database, the last write ignore pattern should not be deleted. Deleting this pattern may cause inconsistency which could result in the deletion of the firewall pattern. Do not delete the pattern, but disable the rule.
8670	Analytics	Analytics not getting generated for existing clusters upon changing the time zone of ScaleArc machine.	Change time zone before putting ScaleArc into production.
8840	Services	If a user has configured SQL Server with dynamic ports, then ScaleArc will not be able to detect this change. When the port of a DB server which is added in a cluster with an instance name changes, it will be marked down by ScaleArc.	Delete the server marked down and add it again to the cluster.
9328	Install/ Upgrade	A session timeout may occur while upgrading from ScaleArc 3.0 to a current version of ScaleArc.	Before performing an upgrade, the session timeout should be increased to at least 120 minutes due to the large package format. Please note, even after increasing the session timeout, a pop message may still appear. Ignore the pop up message.
9342	UI-Cluster Configuration	When trying to create an Always On AG cluster with a windows user; if the AG Listener takes more than 5 seconds to respond to the login request from the UI, ScaleArc closes the underlying TCP connection to the AG Listener and displays "unable to fetch VNN server information" preventing cluster creation.	Click the "Fetch SQL Cluster Config" button multiple times to make the UI pull the AlwaysOn AG related info and create the cluster. This workaround may not work if

			response time always exceeds 5 sec.
10005	Network Settings	ScaleArc does not support changing the VLAN IP's in a NIC bond.	Currently there is no workaround for this issue.
10052	Caching	ScaleArc does not consider Cache and Firewall rules if the "use <DB name>" command within the query is terminated with newline characters ("\r\n"). ScaleArc fetches the response directly from DB server.	Currently there is no workaround for this issue.
10096	Network Settings	When the NIC bond is deleted; related VLAN interfaces do not get deleted.	Delete the related VLAN interface before deleting the bond.
10104	API	ScaleArc allows creating a clone of a cluster where the inbound IP is one which is not added in ScaleArc as well as not present on existing DB server. Ideally an error "VIP doesn't exists or give appropriate VIP as inbound port" should be indicated to the user.	Go back and assign a valid VIP as an inbound IP for the cloned cluster.
10438	Services/ Installation	Upon installation of Scalearc a few services such as failover, user_creds_monitor etc. do not work. The watchdog fails to start these services due to timezone/time mismatch.	This issue occurs only when there is timezone and time mismatch on server where ScaleArc is deployed. Please restart watchdog service once the ScaleArc initial setup is complete and timezone has been changed. Watchdog service is accessed by going to Settings -> System Settings -> from the Services tab select Watchdog and click on Restart Selected Services.
10657	Connection Management	Client read connections that are mapped to connections on a R/W DB server are candidates for migration to other healthy DB servers if the R/W DB server crashes. However, this connection migration is not consistent. Migration does not happen on connections on which a TCP RST is received when migration is yet to happen. A clean migration for all read connections mapped to the R/W server on a R/W server crash is not guaranteed.	Currently there is no workaround for this issue.
11105	Auto Failover	An error message is observed when a non-mirrored logical database is added as the first database in a cluster and SQL mirroring failover is selected.	The first database in a cluster must be in mirrored replication if SQL mirroring failover type is selected.

11214	Auto Failover	Manual failover fails if replication changes do not happen within the timeout period on the database server.	Increase the timeout period and trigger the manual failover again.
11957	Services	High RES memory utilization may be observed by Samba processes in case you have joined a domain through RODC. All of system RAM may be consumed by the Samba process.	To view memory utilization by Samba: Using top/htop on the machine, verify that the process named "samba" is consuming the most memory on the system. If yes, follow the resolution steps below. 1. Get shell access to ScaleArc. 2. Change to the super user (root), using "sudo su-" 3. Fire "/opt/idb/utills/smb restart" to restart the samba process.
12203	Core	The order of Query Routing rules will be ignored.	Currently there is no workaround for this issue.
12875	AlwaysOn/ DB Server Management	When a cluster is created with database servers configured in AlwaysOn Availability Groups, after the AG listener IP address is entered, and the cluster config is fetched, then the status is shown as 'down' if all the nodes are named instances. Once we click setup cluster, these instances appear green and available.	Currently there is no workaround for this issue.
13610	Prep Exec	If you turn QLLB and R/W split ON and set the following on the cluster: • Set Replay: SET IMPLICIT_TRANSACTIONS ON; • Write Ignore: SET IMPLICIT_TRANSACTIONS ON; Following behaviour can be seen on ScaleArc due to load balancing: Due to the setting, SQL Server will start a local transaction on the server that should be committed on the application end explicitly. Thus when we add this rule in set replay, ScaleArc would not know that the server has started a local transaction since we do not do response parsing. Once we lose the response, since QLLB is ON, we do active dissociation of the connection. So when the next query comes with this transaction descriptor, we load balanced it to a server, and the query then gives an error since this transaction is not valid on this connection. The client connection is thus terminated.	Do not add the SET IMPLICIT_TRANSACTIONS ON; in set replay or write ignore, as the server has the authority to start local transaction on the following set of queries: ALTER TABLE, FETCH, REVOKE, BEGIN TRANSACTION, GRANT, SELECT, CREATE, INSERT, TRUNCATE TABLE, DELETE, OPEN, UPDATE, DROP
13840	Auto Failover	While performing SQL Mirroring, if autofailover occurs for the second time before the suspended database server comes in replication (after first	Manually resume the replication for every DB which is in suspended mode after every auto failover.

		autofailover) this will result in data inconsistency on the DB server.	
13891	User Management	No error is displayed on Users and DB pop-up page when a user lacks "View server state' and 'View definition state" permissions	Currently there is no workaround for this issue.
13925	UI	ScaleArc is not able to load from Internet Explorer (v:8.0.7601.17514).	Update to the latest version of Internet Explorer and try logging into ScaleArc.
15054	UI	Verbose debug logs do not get downloaded when using Safari browser.	Safari browser requires the user to set the following settings. Click on Safari in the title bar in the browser -> Preferences -> Select Security -> and deselect the 'block pop-up windows setting'.
15055	UI	TCP dump file does not get downloaded when using Safari browser.	Safari browser requires the user to set the following settings. Click on Safari in the title bar in the browser -> Preferences -> Select Security -> and deselect the 'block pop-up windows setting'.
15364	UI	ScaleArc UI rendering issues are observed in Chrome version 52.0.2743.11.0.46 release 8.7 and earlier.	Refresh the UI and try the operation once more. Alternatively upgrade the chrome version to solve the UI rendering issues.
15641	Active Directory	In an Active Directory setup, if you unjoin and then join the same domain immediately without changing any settings on ScaleArc, then during the process of creating a new Kerberized cluster on ScaleArc, ScaleArc's machine account may not be able to access the SQL Server (even after assigning appropriate privileges).	Follow the steps below to successfully create a Kerberos cluster: <ul style="list-style-type: none"> • Login to the SQL Server using SSMS. Delete the machine account and add it again with the appropriate privileges. Alternatively, after the unjoin, wait for the AD to sync the machine account deletion change (might vary depending on the load on the AD server). • Perform the join again, login to the SQL Server using SSMS and add the ScaleArc's machine account. • Assign the ScaleArc account with correct

			<p>privileges on the SQL Server.</p> <ul style="list-style-type: none"> • Create the cluster once again.
15925	Kerberos	Kerberos delegation fails between sub-domains.	<p>While configuring subdomains in UI, enter information about the parent domain controller. For Kerberos mutual authentication between two subdomains, krb5.conf file in ScaleArc should have following information:</p> <ol style="list-style-type: none"> Sub-domain (to which ScaleArc has joined) Another sub-domain from where the client request is coming from Parent domain controller
16001	Active Directory	While trying to unjoin the AD domain (after adding secondary KDC in "/etc/krb5.conf" file for the joined realm), user may get the following error message "Could not leave the ADS: gss_init_sec_context failed with kinit succeeded but ads_sasl_spnego_gensec_bind(KRB5) failed: An internal error occurred. Failed to leave domain: failed to connect to AD: An internal error occurred."	<p>User (with root privileges) needs to do the following to unjoin the domain:</p> <ol style="list-style-type: none"> 1. Perform "kdestory" on each ScaleArc Machine. 2. Execute "kinit" to request a TGT.
16078	HA	It is observed that when a network restart or a DHCP lease renewal operation occurs, an alert is generated stating "ScaleArc recently encountered an issue that requires analysis". This alert is related to an abort operation performed by either Corosync or Pacemaker CRMd process.	This is expected and can be ignored. This issue does not have any side effects.
16197	Health Monitoring	Cluster landing page may be seen on the secondary machine when HA service is stopped. Eventually the system_monitor service will start the HA service and on refreshing the page, the cluster landing page should disappear.	Check the HA service status on the peer system to confirm which is the primary/secondary machine. Contact ScaleArc customer support if the problem still persists.
16321	HA	Once HA is configured, if a user changes the fencing type from cluster to external db server, the following event is generated "fencing is not configured".	Ignore the alert. Delete it from the UI to prevent it from appearing again.
16462	HA	On HA setups which have multiple network interfaces, when the interface on which HA is configured is brought down, failover does not occur as expected.	If an 'ifdown' operation for network maintenance is intended on the primary machine, then it is recommended to switch it to a secondary machine first. An

			'ifdown' operation for all the interfaces which owns VIP should be performed.
16494	Health Monitoring	If the port of a DB server in a cluster configured with an instance name is changed then subsequently the health of all the database servers in the ScaleArc cluster get marked down.	Delete and add the database server on the cluster whose instance name to port mapping has changed.
16524	HA	After performing a HA join, any one of the machines in the pair can behave as the primary machine.	Currently there is no workaround for this issue.
16601	HA	If the first 13 characters of the initial primary machine hostname in a HA setup is the same as the first 13 letters of the initial primary hostname of another HA setup, then you cannot use same witness server for fencing configuration. The witness storage will conflict and split brain resolution will be flawed.	Rename the hostname of the primary machine in a HA setup so that the first 13 characters are not identical to the hostname of another primary machine.
16670	Authentication Offload	When connection churn rate is high an intermittent drop in QPS can be observed when majority of these connections are requesting Kerberos authentication.	Click here for a detailed explanation of the issue and workaround.
16706	AD	In an Azure cloud environment, when the ScaleArc hostname is created using upper case letters an error is generated during AD join.	Configure ScaleArc's hostnames in lower case before launching ScaleArc instance in Azure cloud (if the instance is going to be used for Kerberos authentication).
16837	Upgrade	After upgrading from ScaleArc v3.11.0.1, connections requesting windows authentication on an existing cluster are running slower.	Turn OFF Replay cache from 'Cluster Settings > Debug > Replay cache' and turn it back ON.
16886	Firewall	When a user adds a firewall pattern from the security tab, the order of the new firewall pattern does not get automatically updated if there are multiple pages of rules.	Go to the last page of firewall patterns and add the new rule.
17406	Authentication Offload	An error is observed on SSMS 2016 on a Kerberised Always ON cluster where SSL connections with MARS is enabled. The error occurs even though the connections are successfully established in passthrough.	Currently there is no workaround for this issue.

Highlight Behaviors

This section highlights ScaleArc changes/ behaviors impacting end users for this release.

<i>ID</i>	<i>Component</i>	<i>Description</i>
	Upgrade	ScaleArc will automatically restart when upgrading from a previous version of ScaleArc (v.3.6.1 or earlier) to v3.11.0.4 because the kernel is also upgraded.
10805	Auto failover	While implementing auto failover with external API for SQL Mirroring, the following must be configured: <ol style="list-style-type: none"> 1. The database which is in replication on the DB server should be added in the cluster for that root user. 2. In case multiple logical DB's are added, then replicated DB should be first in that list. This is because external API picks the first database added for root user and performs operations as per that database. 3. The user should perform failover using ScaleArc 'manual failover' utility; otherwise ScaleArc does not reflect such role changes done by SQL servers in the back end. 4. With Asynchronous mirroring replication on DB servers, user is not recommended to use manual failover from ScaleArc manual failover utility. External API will error out as "Principal server is UP but some of the database server is configured as Async, so aborting the failover" with no changes in role of DB servers.
11183	Analytics	Analytics is unable to populate per minute analytic data from ScaleArc v3.4 to 3.11.0.4.
11254	Auto failover	<ol style="list-style-type: none"> 1. In SQL mirroring replication, the secondary database server must be set as Standby + No Traffic; otherwise, all queries served to such server turns out to be errors 2. If logical database on DB server is in async replication, then manual failover will fail with an error as "Principal server is UP but some of the database server is configured as async, so aborting the failover".
11425	Installation	During installation if the smallest disk has size less than 25GB, then installation fails with the following error: "Could not allocate requested partitions: not enough space for LVM requests." Please make sure the disk size is at a minimum 25GB.
11541	Core	ScaleArc core module terminates connections requesting data encryption for customers who are in enterprise license mode.
12312	Connection Management	When connection pooling is enabled and Idle Server Connection Time Out setting is more than 2 Minutes, ScaleArc internally creates/maintains a single connection as part of cluster health monitoring. This Connection times out when Idle Server Connection Time Out is reached.
12827	Core	The following error message 'Error while reading TDS-Prelogin packet from client' is observed in alert logs on the ScaleArc cluster that is added as a DB server on the super cluster. This message can be ignored.
13597	Auto failover	In SQL Mirroring there are only two servers are currently present in ScaleArc. If user makes R+W server unreachable from Scalearc using iptables rule, the server goes down in ScaleArc. That triggers Auto-failover in ScaleArc. However actually on the database server, R+W server is UP and running fine. Now, if we want to do the failover of individual logical databases then we need to fire an ALTER command for that database on the Secondary machine on ScaleArc. However,

		<p>those commands fail because Primary server is actually UP. So, attempt to do failover on individual databases do not succeed.</p> <p>Workaround:</p> <p>To perform Auto failover on individual logical databases, from the SQL Server management studio perform manual failover of logical databases configured as Principle on current Primary Server of Scalearc. This will trigger failover and failover service will make the UI role changes because it will not find any database configured as a Mirror on Secondary server.</p>
14323	Upgrade/ Downgrade	The entries added to the ignore_rpm_versions.conf file are valid only for that upgrade. For every new upgrade a fresh ignore_rpm_versions.conf file needs to be populated.
14893	Autofailover	User can perform a manual switchover even if autofailover is turned OFF.
15747	Connection Management	ScaleArc requires hostnames to be in lower case before joining as a machine account to the AD server. When upgrading from a previous version of ScaleArc to 3.11.0.4, ScaleArc will send an alert for the user to change the case of the hostname. When installing ScaleArc 3.11.0.4 for the first time, ScaleArc will take care of the hostname case internally.
15967	Active Directory	At times ScaleArc AD join may fail due to AD domain controller failure, try cleaning up stale data from AD domain and start the re-join process.
16042	UI	In HA, during an unjoin operation, if the secondary machine successfully unjoins from the domain and the primary unjoin fails, the error message that appears subsequently is misleading. The message, "ERROR 2418: Failed to un-join Scalearc as the Domain. Could not leave the ADS: Failed to leave domain: failed to connect to AD: No logon servers".
16232	HA	ScaleArc will not allow a cluster to be deleted or stopped from the cluster landing page if it is configured for fencing. To stop a cluster, change the fence type, stop the cluster, perform the required operation, start the cluster, and revert the fencing type.
16275	Authentication Offload	<p>When ScaleArc Authentication Offload is turned ON, the benchmark factory tool is unable to establish connection with ScaleArc. This issue occurs when TDS packet size is set to a size other than 4096 or 8000. The DB server setting should be changed to reflect a packet size of either 4096 or 8000 using the following SQL queries:</p> <pre>EXEC sp_configure 'show advanced options', 1; go EXEC sp_configure 'network packet size' go EXEC sp_configure 'network packet size',4096 go RECONFIGURE Go</pre>
16500	HA	A warning message pops up while configuring HA if hostname of both primary and secondary machines is identical. It is safe to ignore the warning; HA configuration will be successful despite the warning message.
16524	HA	It is expected that after an HA join, the Primary could be any one of the machines in the pair.
16563	AD	<p>Any RODC configuration with a ScaleArc hostname that contains uppercase characters should be unjoined prior to upgrading to 3.11.0.4. Refer to the following KB article for more information on unjoining a Windows AD domain:</p> <p>https://support.scalearc.com/kb/articles/4120</p>
16629	HA	Multiple failovers may occur when the network cable, used for communication with the witness server, is pulled. The network interface goes down and is later brought back up.
16661	HA	In a situation where:

		<ul style="list-style-type: none"> ▪ Secondary network interface is used for HA ▪ This interface doesn't have a gateway assigned ▪ Witness server and application are on another network and can be reached via the primary interface <p>If the HA interface is brought down by pulling the network cable, then an HA failover does not occur. Instead, the primary fences the secondary node. Any VIP's assigned to the secondary interface go down and will not failed over. This situation will not cause any disruption to the application traffic since the traffic is communicating via the VIP's on the primary interface which remains on the original node.</p>
16806	HA	Queries will not be processed for 60 seconds when the following scenario occurs. If a ScaleArc partition on primary is detected as 'read only', the Pacemaker will go into standby. The process of the primary machine moving in a standby mode while the secondary machine takes over takes time while the scalearc_traffic resource is stopped and restarted.
16915	License	While upgrading ScaleArc or installing a trial version, the trial license EULA must be accepted to complete the installation. Similarly, if a production license is installed, then the new EULA (specific to production license) must be accepted to complete the installation.

Limitations

This section describes limitations in ScaleArc for SQL Server 3.11.0.4. The following limitations have workarounds.

<i>Issue ID</i>	<i>Component</i>	<i>Description</i>	<i>Solution</i>
7612	Core	For traffic that fires RPC stored procedure with multiple OUTPUT parameters, please do not turn ON prepare cache handling.	Turn OFF prepare caching.
8842	Services	Replication monitoring service will not work for servers added as an instance name into a non-Always On cluster.	Configure servers with port number and not instance names.
10060	Authentication	When client adds SQL User for Windows Authentication, both formats of user specification "NetBIOS-domain-name\username" and "FQDN\username", should be added in ScaleArc	To overcome this situation Windows users on ScaleArc must be configured with DNS or NetBIOS domain names depending upon what the application is sending (e.g., us.scalearc.net\john, us\john).
10709	Auto Failover	Manual failover with external API continues even after API timeout has been reached. Because of which failover happens and DB server roles gets changed but logs says "ERROR Request did not complete successfully: Failed to do failover via external API."	Increase the API timeout value to avoid the issue from occurring.

11402	Auto Failover	External API failover operation is failing with syntax error if mirrored DB name has special characters in it.	The following are unsupported characters for DB names: ? > < * "
11780	Core	On creating a cluster with 2 DB servers having 100 db users and 100 DB's for each the user, with atleast one cache rule configured to each logical db, the cluster stops on its own after being started.	Limit the number of DB users, databases, and cache pattern combinations to less than 3,000.
14985	AD	If a Windows user (whose "password replication policy" is not allowed) is added manually after turning Auto Fetch User 'ON', then the user creds monitor deletes the user from the cluster.	Do not turn Auto Fetch User 'ON' if you want to manually add a Windows user. Or allow "password replication policy" to manually add a Windows user.
15988	Kerberos	Secondary KDC needs to be manually added to the realm.	User needs to manually add a secondary KDC in the "/etc/krb5.conf" file for each ScaleArc machine present in HA.

Following are limitations in ScaleArc for SQL Server 3.11.0.4 that do not have workarounds:

	Protocol Passthrough	ScaleArc does limited processing for TDS protocols 7.1 or lower, due to protocol limitations. ScaleArc by default puts all such traffic into passthrough. TDS 7.1 was introduced along with SQL 2000, now only seen on legacy drivers. Applications will need to upgrade their drivers to take full benefit of the ScaleArc. Updating drivers to the latest version typically does not require application changes
5608	Query Processing	MARS traffic is not supported by ScaleArc and it goes in pass-through. MARS is a feature from SQL server 2005 on-wards wherein there can be multiple outstanding requests being processed on a time-shared manner by the server-side-thread on a single connection.
5832	R/W Split	If a write query is received by ScaleArc as part of a multi-query packet and if the first query in the packet is a read query, then the entire query packet is considered as a read query and may be sent to a read-only server.
5887	Core	When a user connects to ScaleArc with a logical DB name that it does not have permissions on, the authentication will succeed through ScaleArc. But then when the client fires a query that causes ScaleArc to open a connection to the server the LOGIN fails and ScaleArc fails the query.
6140	Core	Multi-packet SET Replay is not supported i.e. if the combined packet size of SET queries exceeds 4k (negotiated packet size in case of Auth Offload ON), SET Replay is skipped. The amount of set queries that one can fit in a query packet depends on the query packet size that is negotiated during authentication. So, for a packet size of 4096: $(4096 - 8) / 2 = 2044$ characters concatenated (semicolon delimited) SET query can be accommodated in one packet.
6814	Logs	ScaleArc expects queries to be sent in UTF-16 encoding and English characters only. The behavior in case of 'non UTF-16' and in case of 'UTF-16 and non-English characters' is undefined.
8352	Analytics	ScaleArc Analytics does not detect and log stored procedures as different from normal query in the query log currently.
8514	Core	ScaleArc does not support for extended features of TDS like SESSION_RESTORE and FED_AUTH. When running traffic through ScaleArc, if a client requests extended features such as SESSION_RESTORE or FED_AUTH, during the login time, ScaleArc will

		disable these features, authenticate the client and then start processing the traffic as usual.
11196	Authentication Offload	Additional Certificate's common name (CN) validation is currently not performed by ScaleArc behaving as SQL client. SQL client that sets the "TrustServerCertificate" option does an additional validation along with the actual certificate validation. The client takes the Common Name from the certificate (CERT_COMMON_NAME) and tries to see that it matches the server's hostname that the client tried to connect to. If the client tried to connect to the server using an IP address, it would reverse lookup the hostname using the IP address and try to match the CERT_COMMON_NAME with this one. Such validation is currently not supported in ScaleArc.
11198	Caching	Cache pattern and the corresponding cached response are not persistent and on restart of a cluster the response is obtained from server rather than cache.
11388	Caching Core	With cluster having 100 users and 100 logical DBs assigned to each user, CPU utilization for save_cluster_cache () sometimes crossed 50% even in idle cluster (cluster without traffic).
11914	Authentication Offload	ScaleArc SSL does not support payload data larger than 4096 during the handshake process.
15889 15813	Core	TDS version 7.1 and above is required for Kerberos traffic. With older TDS versions, Kerberos connections cannot be processed in passthrough and would subsequently fail since ScaleArc does not offload authentication for such connections.
15946	Kerberos	"ALL IPs" during Kerberised cluster creation, is not supported for any of ScaleArc's HA configuration for ON-premise or On-cloud environment.
16102	UI	ScaleArc does not allow interface interrupt affinity to be set via the CPU distribution UI in case of cloud platform deployments. This feature is only available for non-cloud deployments like VM and bare metal.
16629	HA	Multiple failovers may occur when the network cable, used for communication with the witness server, is pulled. The network interface goes down and is later brought back up.
16707	HA	Primary DNS and search domain values change in Azure cloud, whenever machines are configured in HA.
16746	AD	ScaleArc is unable to join Active Directory if there is a single quote (') in the password.
16790	HA	Pure SSL only DB server as an external DB server for fence configuration is not supported.
16795	HA	Currently ScaleArc does offer support for a Kerberized user to configure HA fencing in case of external DB.

Troubleshooting Kerberos

If you unjoin and join the same domain immediately without changing any settings on ScaleArc, the scenario below may be encountered.

<i>Description</i>	<i>Solution</i>
While creating a new Kerberized cluster on ScaleArc, ScaleArc's machine account may not be able to access the SQL Server even after assigning correct privileges	<ul style="list-style-type: none"> • Login to the SQL Server machine using SSMS. • Delete the machine account and add it again with the correct privileges. OR • After the unjoin, wait for the AD to sync the machine account deletion change (might vary depending on the load on the AD server). • Perform the join again. • Login to the SQL Server using SSMS. • Add the ScaleArc's machine account. • Give the ScaleArc account the correct privileges on the SQL Server. • Create the cluster again.

Additional Resources

You can find news, articles, videos, webinars, and other useful information on [ScaleArc's web site](#).

To get the most out of the features in ScaleArc for SQL Server check out our [ScaleArc training videos](#).

Access [ScaleArc's Knowledge Base](#) for how-to articles, feature description, and troubleshooting information.

If you need further assistance with any ScaleArc product or service, please [contact us](#).