



SCALITY RING

AN EDGENEXUS ADC DEPLOYMENT GUIDE



Contents

Document Properties	2
Document Disclaimer	2
Copyrights	2
Trademarks	2
Edgenexus Support	2
Introduction	3
Document Intention	3
Scality Ring versions supported	3
Acronyms	3
VIPs, Ports, and Other Bits	4
HTTP VIP Services	4
HTTPS VIP Services	4
Port Requirements	4
Sizing the EdgeADC for Scality	4
Deployment Scenarios	5
Scality Ring Virtual Service Methods	5
Creating the Scality Real Server Monitors	6
Scality S3	6
VIP – Scality S3 – SSL Bridging	7
VIP – Scality S3 – SSL Offload	9
Layer 7 Summary	11
High-Speed Transactions and Layer 4 with DSR	12
Configuring the EdgeADC Layer 4 DSR for Scality	13

Document Properties

Document Number: 2.0.6.24.21.12.06

Document Creation Date: May 20, 2021

Document Last Edited: June 24, 2021

Document Author: Jay Savoor

Document Last Edited by:

Document Referral: EdgeADC - Version All

Document Disclaimer

Screenshots and graphics in this manual may differ slightly from your product due to differences in your product release version. Edgenexus ensures that they make every reasonable effort to ensure that the information in this document is complete and accurate. Edgenexus assumes no liability for any errors. Edgenexus makes changes and corrections to the information in this document in future releases when the need arises.

Copyrights

© 2021 All rights reserved.

Information in this document is subject to change without prior notice and does not represent a commitment on the manufacturer's part. No part of this guide may be reproduced or transmitted in any form or means, electronic or mechanical, including photocopying and recording, for any purpose, without the express written permission of the manufacturer. Registered trademarks are properties of their respective owners. Every effort is made to make this guide as complete and as accurate as possible, but no warranty of fitness is implied. The authors and the publisher shall have neither responsibility nor liability to any person or entity for loss or damages arising from using the information contained in this guide.

Trademarks

The Edgenexus logo, Edgenexus, EdgeADC, EdgWAF, EdgeGSLB, EdgeDNS are all trademarks or registered trademarks of Edgenexus Limited. All other trademarks are the properties of their respective owners and are acknowledged.

Edgenexus Support

If you have any technical questions regarding this product, please raise a support ticket at: support@edgenexus.io

Introduction

Scality Ring is one of the leading object-based storage solutions on the market that uses Amazon S3 as its specialist foundation. Scality's revolutionary technology means businesses small, large, and very large can take advantage of object-based storage in their data centers or cloud located S3 using the native Amazon S3-API.

High availability is supported in Scality architecture using load balancers that are placed in advanced positions. The EdgeADC, with its advanced health monitoring and flightpath rules technology, makes the availability of the storage infrastructure assured.

The EdgeADC is capable of load-balancing Scality Ring, and this guide explains how to set this up.

Document Intention

This document is aimed at administrators who need to load-balance their Scality storage nodes efficiently and quickly.

The quickest way to do this is to use the Scality Ring jetPACK, in which we have done all the work for you. However, we will also show each item that is configured, so you have a better understanding.

Scality Ring versions supported

This document and support are valid for Scality Ring v7.44 and above.

Acronyms

VIP – Virtual IP

VS – Virtual Service

VIPs, Ports, and Other Bits

When load balancing Scality Ring, the following VIPs will be needed for Scality operations.

HTTP VIP Services

- S3 (HTTP) for handling HTTP requests from S3 client applications

HTTPS VIP Services

- S3 (HTTPS) for handling HTTPS requests from S3 client applications

Port Requirements

The following are the port requirements for the Scality Ring platform.

Port	Protocol	Service Type	Explanation
80	TCP	L4-TCP or L7 HTTP	This port is used to handle all HTTP requests from S3 client applications. You can use Layer 4 TCP with SSL Passthrough or Layer 7 with SSL Offload or SSL Bridging.
443	TCP	L4-TCP or L7 HTTPS	This port is used to handle all HTTPS requests from S3 client applications. You can use Layer 4 TCP with SSL Passthrough or Layer 7 with SSL Offload or SSL Bridging.

Sizing the EdgeADC for Scality

The ADC can operate in either physical or virtual deployments. The reverse proxy engine within the ADC is optimized for speed and efficiency. The ADC will use all available threads automatically.

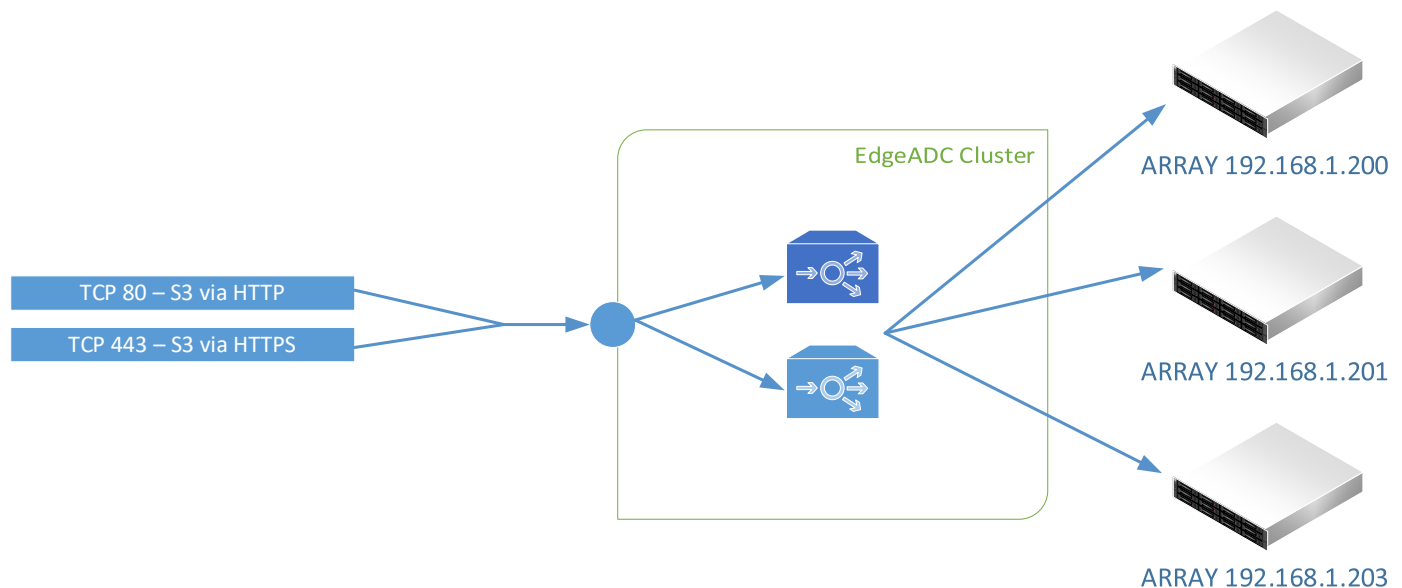
In virtualized environments, we recommend that you set the ADC to 4 vCPU with 8GB RAM, to begin with, and scale up when you need to.

We recommend that you utilize the hardware platforms from our partners in physical environments, with the base system being a quad-core Intel Xeon with 8GB RAM.

In both cases, 50GB of disk storage space should be sufficient.

Deployment Scenarios

Connections to the Scality Ring system occur by clients connecting to the VIP or Virtual IP service created on the ADC. The ADC then load-balances the connections to the Scality nodes configured within the ADC and linked to the VIP. An example diagram is shown below.



Scality Ring Virtual Service Methods

There are four different methods of configuring the ADC for use with Scality Ring.

Non-Encrypted Port 80	In this mode, the traffic will enter the ADC using an un-encrypted VIP using port 80. It will then be sent onto the nodes using the same means. Traffic will not be encrypted when using this mode and is not recommended for best practices. ADC service type Layer 4 TCP is used.
SSL Passthrough	In this mode, the traffic enters the ADC on port 443 using SSL. The traffic is sent onto the nodes without inspection. ADC service type Layer 4 TCP is used.
SSL Bridging	In this mode, the SSL traffic is terminated in the ADC and then re-encrypted before passing to the nodes. When this mode is chosen, you will need to have the SSL certificate installed on the nodes and install it in the ADC. This mode is the recommended best practice method for security reasons. ADC service type HTTP is used.
SSL Termination	This mode allows SSL traffic to be received by the ADC, which then terminates the SSL encryption internally before passing it to the nodes using unencrypted SSL. This mode may not be the desired choice for security reasons. ADC service type HTTP is used.

The following pages will take you through the VIP configuration. Please take care to configure correctly to avoid issues in operations.

Creating the Scality Real Server Monitors

A vital requirement within the ADC infrastructure for Scality, and indeed for many other installations, is the definition of real server monitors that will check on the back-end nodes' health and then use the data to determine if traffic should be sent to them.


- Navigate to Library > Real Server Monitors using the navigation pane.
- Click Add Monitor.
- A new blank entry will be created.
- Create the monitors using the guides below.

Scality S3

Name	Description	Monitoring Method	Page Location	Required Content	User	Password
Scality-S3	As needed	HTTP 200OK	/_/healthcheck/deep/	200 OK	Username	Password

VIP – Scality S3 – SSL Bridging

The first VIP and VS we are going to create is the one that handles the S3 traffic. We will be showing the creation of an HTTPS VIP, but the detail for making the HTTP VIP/VS is almost the same. The method being used here is SSL Bridging. In this method, the SSL traffic enters the ADC, is then terminated internally, any inspection required is carried out, and the traffic is then re-encrypted and sent to the nodes.



- The first step is to create the VIP and initial VS
- Log into the ADC and go to IP Services. This location should be the default entry point.
- Click Add Service
- You will see an empty row into which you will add values similar to the one below. The field values we provide are examples for your reference.

IP Address	Subnet Mask	Port	Service Name	Service Type
192.168.1.222	255.255.255.0	443	S3 Scality	HTTP

So this has now created the initial VIP with the entry IP address of 192.168.1.222. In this example, we show a NAT IP address, and the assumption is that there is a firewall between the ADC and the public Internet. You can, of course, have a public IP address as the VIP entry point.

- Now we will define the Real Servers (RS) section.
- Click on the Servers tab to display the Real Servers listing.
- There is a ready-created blank entry to aid you in adding the RS entries.
- Please enter the details relevant to your infrastructure following the examples we have provided below. In our case, we have three array nodes, but you may have more.

Address	Port	Weight	Calculated Weight	Notes	ID
192.168.1.201	443	100	100	Array Node 1	

- Click Update to save.
- Click the Copy Server button and make changes for the second array node.

Address	Port	Weight	Calculated Weight	Notes	ID
192.168.1.202	443	100	100	Array Node 2	

- Click Update to save.
- Click the Copy Server button and make changes for the third array node.

Address	Port	Weight	Calculated Weight	Notes	ID
192.168.1.203	443	100	100	Array Node 2	

- Click Update to save.

You can add a name for the server group if you wish.

We have now defined our first VIP, and its two connected Real Server nodes. We have to do some more work yet to do.

The next stage is to configure the Basic tab.

- Click on the Basic tab within the Real Servers section.
- Make changes as follows:

Field	Value
Load Balancing Policy	Least Connections
Server Monitoring	Scality-S3
Caching Strategy	Off
Acceleration	Compression
Virtual Service SSL Cert	Your SSL certificate
Real Server SSL Cert	Your SSL certificate

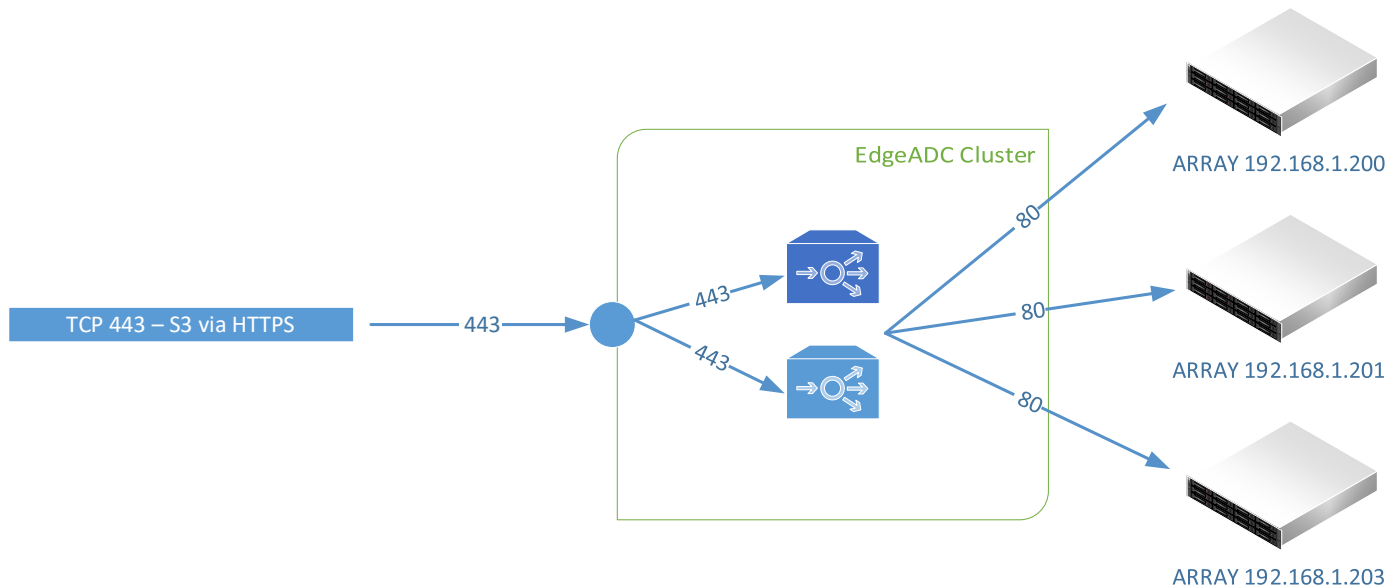
- Click Update when done.

There are no configurations to be done within the Advanced tab.

Note: To add your SSL certificate, please consult the EdgeADC Administration Guide

VIP – Scality S3 – SSL Offload

The VIP and VS we are going to create is the one that handles the S3 traffic. We will be showing the creation of an HTTPS VIP, with the traffic being passed onto the Scality node-set using HTTP 80.



- The first step is to create the VIP and initial VS
- Log into the ADC and go to IP Services. This location should be the default entry point.
- Click Add Service
- You will see an empty row into which you will add values similar to the one below. The field values we provide are examples for your reference.

IP Address	Subnet Mask	Port	Service Name	Service Type
192.168.1.222	255.255.255.0	443	S3 Scality	HTTP

So this has now created the initial VIP with the entry IP address of 192.168.1.222. In this example, we show a NAT IP address, and the assumption is that there is a firewall between the ADC and the public Internet. You can, of course, have a public IP address as the VIP entry point.

- Now we will define the Real Servers (RS) section.
- Click on the Servers tab to display the Real Servers listing.
- There is a ready-created blank entry to aid you in adding the RS entries.
- Please enter the details relevant to your infrastructure following the examples we have provided below. In our case, we have three array nodes, but you may have more.

Address	Port	Weight	Calculated Weight	Notes	ID
192.168.1.201	80	100	100	Array Node 1	

- Click Update to save.
- Click the Copy Server button and make changes for the second array node.

Address	Port	Weight	Calculated Weight	Notes	ID
192.168.1.202	80	100	100	Array Node 2	

- Click Update to save.
- Click the Copy Server button and make changes for the third array node.

Address	Port	Weight	Calculated Weight	Notes	ID
192.168.1.203	80	100	100	Array Node 2	

- Click Update to save.

You can add a name for the server group if you wish.

We have now defined our first VIP and its connected Real Server nodes. We have to do some more work yet to do.

The next stage is to configure the Basic tab.

- Click on the Basic tab within the Real Servers section.
- Make changes as follows:

Field	Value
Load Balancing Policy	Least Connections
Server Monitoring	Scality-S3
Caching Strategy	Off
Acceleration	Compression
Virtual Service SSL Cert	Your SSL certificate
Real Server SSL Cert	None

The above configuration will ensure SSL termination/offload.

- Click Update when done.

There are no configurations to be done within the Advanced tab.

Note: To add your SSL certificate, please consult the EdgeADC Administration Guide

Layer 7 Summary

The Layer 7 traffic load balancing configurations are now complete, and the ADC should look something like the example below.

The screenshot displays the EDGE NEXUS GUI interface. At the top, there is a status bar with 'GUI Status' (green checkmark), 'Home' (house icon), 'Help' (question mark icon), and a user dropdown 'admin'. The left sidebar contains a 'NAVIGATION' menu with 'Services' (selected), 'App Store', and 'IP-Services'. The main content area is divided into two sections: 'Virtual Services' and 'Real Servers'.

Virtual Services Section:

- Buttons: Copy Service, Add Service, Remove Service.
- Search bar.
- Table with columns: Mode, VIP, VS, Enabled, IP Address, SubNet Mask / Prefix, Port, Service Name, Service Type.

Mode	VIP	VS	Enabled	IP Address	SubNet Mask / Prefix	Port	Service Name	Service Type
Active			<input checked="" type="checkbox"/>	192.168.1.222	255.255.255.0	443	Scalify Ring S3	HTTP

Real Servers Section:

- Tabs: Server (selected), Basic, Advanced, flightPATH.
- Group Name: Server Group.
- Buttons: Copy Server, Add Server, Remove Server.
- Table with columns: Status, Activity, Address, Port, Weight, Calculated Weight, Notes, ID.


Status	Activity	Address	Port	Weight	Calculated Weight	Notes	ID
	Online	192.168.1.200	443	100	100	Array Node 1	
	Online	192.168.1.201	443	100	100	Array Node 2	
	Online	192.168.1.203	443	100	100	Array Node 3	

High-Speed Transactions and Layer 4 with DSR

The reverse proxy within the EdgeADC is an exceptionally high speed and has been built for high-speed transactional environments. However, there may be occasions when your needs demand even higher throughputs, and in such cases, we would recommend you switch to Layer 4 TCP load balancing, with Direct Server Return rather than Reverse Proxy.

DSR or Direct Server Return is often used when the Request is small, but the Request-Reply is large. Examples could include large image files, video playback and streaming, and large data file retrieval. In such cases, the traversal of data through the Reverse Proxy would probably delay its return to the originator, and passing it directly to the source is highly advantageous.

An example of DSR is shown in the illustration below.



In the example shown, the Request traverses through the EdgeADC and is distributed to the Scality nodes using the load balancing policy.

When the Request-Reply is sent back to the user, it bypasses the EdgeADC and is returned directly to the user, a Direct Server Return, or DSR.

Configuring the EdgeADC Layer 4 DSR for Scalify

The configuration of the EdgeADC for Layer 4 and Direct Server Return is very similar to the Layer 7 method. There are essentially two fields that need changing.

- Service Type – for each virtual service defined, and should be changed to Layer 4 TCP.
- Connectivity – This is found under Real Servers > Advanced > Connectivity and should be changed to Direct Server Return. This change needs to be done for EACH VIP.

Please see the screenshots below.

EDGE NEXUS

GUI Status Home Help admin

IP-Services

NAVIGATION

- Services
- App Store
- IP-Services

Mode	VIP	VS	Enabled	IP Address	SubNet Mask / Prefix	Port	Service Name	Service Type
Active			<input checked="" type="checkbox"/>	192.168.1.222	255.255.255.0	443	Scalify Ring S3	Layer 4 TCP

Real Servers

Server Basic Advanced flightPATH

Group Name: Server Group

Copy Server Add Server Remove Server

Status	Activity	Address	Port	Weight	Calculated Weight	Notes	ID
Online	Online	192.168.1.200	443	100	34	Array Node 1	
Online	Online	192.168.1.201	443	100	34	Array Node 2	
Online	Online	192.168.1.203	443	100	50	Array Node 3	

Library View System Advanced Help

Real Servers

Server Basic Advanced flightPATH

Connectivity: Direct Server Return

Connection Timeout (sec): 600

Cipher Options: Defaults

Monitoring Interval (sec): 1

Client SSL Renegotiation: ☒

Monitoring Timeout (sec): 10

Client SSL Resumption: ☒

Monitoring In Count: 2

SNI Default Certificate: None

Monitoring Out Count: 3

Security Log: On

Max. Connections (Per Real Server):

Update

VERY IMPORTANT

For Direct Server Return to work, the following **mandatory conditions must be met**.

- a. The EdgeADC and the Scality Nodes must be on the same network segment / switching fabric. This requirement is due to the load balancing method working by rewriting MAC ID by operating at Layer 2 of OSI.
- b. Each Scality Ring node must take ownership of the VIP address so they can all accept requests and send back responses. The address will need to be assigned to a loopback adapter.
- c. Each Scality Ring node must be configured to not respond to ARP requests for the VIP address or advertise they own the VIP address.