FUJITSU Enterprise Postgres 14 for Kubernetes

Reference Guide

Linux

FUJITSU

J2UL-RF14-05ENZ0(00)
June-2023

Preface

Purpose of this document

This document is a reference, and explains parameter.

Intended readers

This document is aimed at people who manage and operate.

Readers of this document are also assumed to have general knowledge of:

- Linux
- Kubernetes
- Containers
- Operators

Structure of this document

This document is structured as follows:

Abbreviations

The following abbreviations are used in this manual:

Full Name	Abbreviations
FUJITSU Software Enterprise Postgres for Kubernetes FUJITSU Software Enterprise Postgres	FEP or FUJITSU Enterprise Postgres
Transparent Data Encryption	TDE
Custom Resource	CR
Custom Resource Definition	CRD
Persistent Volume	PV

Trademarks

- Linux is a registered trademark or trademark of Mr. Linus Torvalds in the U.S. and other countries.
- Red Hat and all Red Hat-based trademarks and logos are trademarks or registered trademarks of Red Hat, Inc. in the United States and other countries.
- S/390 is a registered trademark of International Business Machines Corporation ("IBM") in the U.S. and other countries.

Other product and company names mentioned in this manual are the trademarks or registered trademarks of their respective owners.

Export restrictions

If this document is to be exported or provided overseas, confirm legal requirements for the Foreign Exchange and Foreign Trade Act as well as other laws and regulations, including U.S. Export Administration Regulations, and follow the required procedures.

Issue date and version

```
Edition 5.0: June 2023
Edition 4.0: October 2022
Edition 3.0: September 2022
Edition 2.0: June 2022
Edition 1.0: March 2022
```

Copyright

Contents

Chapter 1 Custom Resource Parameters

This chapter explains the parameter.

1.1 FEPCluster Parameter

Equivalent Kubernetes command: kubectl apply -f FEPClusterCR.yaml

This operation will create a FEPCluster with supplied information in FEPClusterCR.yaml.

Initial configuration and subsequent changes to FEP Cluster are done through FEP Cluster CR.

Field	Default	Details
metadata.name	new-fep	Name for the Cluster. FEP server container will use this value for Patroni scope. e.g. new-fep
spec.fep.autoPodRestart	<omitted>	Optional This parameter affects the behaviour when value(s) of CPU, memory and/or image for FEP and/or optional Backup container are updated in FEPCluster CR. If it is NOT defined and set to True, operator will automatically create an action CR to make values effective by restarting all pods in an orderly fashion to minmise outage. If is set to False, automatic restart of PoDs will NOT happen. To make the changes effective, user must restart pods by creating action CR with type ‘pod_restart’ and arguments ‘ALL’
spec.fep.fepVersion	<omitted>	Optional When deploying a new FEP cluster, this parameter controls which FEP major version will be used for the deployment. If not specified, Operator will use latest FEP version supported by the Operator. When fepVersion is defined but not spec.fep.image.image, Operator will deploy the specific version of FEP. When both fepVersion and image are defined, Operator will use the image and discard the value of fepVersion. Current support value: 12, 13, 14 Note: Changing fepVersion from one version to another version is not supported after deployment.
spec.fep.customAnnotation.allDeployments	{} (*)	Contents under this are optional. User can remove {} and add multiple key-value pairs. All of these pair will be

Field	Default	Details
		added to annotations of FEP statefulSet and FEP Pods. If left at default, no annotation is added to Pods and statefulSets
spec.fep.image.image	<omitted>	FEP server container image to be used quay.io/fujitsu/fujitsu-enterprise-postgres-14-server:ubi8-14-0.0 It is optional Image line is omitted by default. This key has a higher precedence than fepVersion. If both fepVersion and image are omitted, Operator will use the latest FEP version that it supports. If both fepVersion and image are specified, Operator will use the specified image and ignore the value in fepVersion.
spec.fep.image.pullPolicy	IfNotPresent	
spec.fep.mcSpec.limits	cpu: 500m memory: 700Mi	
spec.fep.mcSpec.requests	cpu: 200m memory: 512Mi	
spec.fep.sysExtraLogging	false	To turn extra debugging on, set value to true It can be turned on/off at any time
spec.fep.instances	1	Number of nodes in the cluster, including both Master and Replicas. In Example CR, it is kept at 1 for certification. However, user can change it to 3 for 1 master and 2 replicas.
spec.fep.servicePort	27500	TCP port for FEP master service
spec.fep.syncMode	off	Replication Mode: off - async replication on - sync replication
spec.fep.forceSsl	true	Controls that the communication to the server should only be via SSL. Changes are reflected in pg_hba.conf
spec.fep.locale	<omitted> (*)	Optional Can only be specified when creating a FEPCluster. Database Cluster Locale Settings: ja_JP - Japanese locale Default - C

Field	Default	Details
spec.fep.monitoring		This is an Optional section. This defines whether monitoring enabled(true) or disabled(false) , MTLS enabled or disabled & Basic authentication enabled or not
spec.fep.monitoring.enable	false	If set true, the operator will create FEPExporter with given spec
spec.fep.monitoring.fepExporter		This is Optional section. Exporter spec section applied only if enable: true
spec.fep.monitoring.fepExporter.authSecret		This is Optional section. Base Authentication secret to provide username & encrypted password of user
spec.fep.monitoring.fepExporter.authSecret.secretName	(created by user)	Mandatory Name of secret that contains username and password
spec.fep.monitoring.fepExporter.authSecret.userKey	(created by user)	Mandatory Key of username in specified secret
spec.fep.monitoring.fepExporter.authSecret.passwordKey	(created by user)	Mandatory Key of password in specified secret
spec.fep.monitoring.fepExporter.tls		This is optional section. FEPExporter MTLS specs. Mandatory if tls specs defined for Prometheus specs
spec.fep.monitoring.fepExporter.tls.certificateName	(created by user)	Mandatory.This points to Kubernetes TLS secret that contains the certificate of FepExporter. Prometheus will use this for certificate authentication. The certificate itself is stored in the key tls.crt.
spec.fep.monitoring.fepExporter.tls.caName	(created by user)	Mandatory This points to Kubernetes configmap that contains additional CA the client use to verify a server certificate. The CA is stored in the key ca.crt.
spec.fep.monitoring.fepExporter.customLabel		Optional List of key value pair to be added to Prometheus ServiceMonitor label. The following label will always be added to ServiceMonitor, regardless if a value is specified here or not. fepsmgrp: sm-fep-exporter
spec.fep.monitoring.prometheus		This is Optional section. Prometheus specs are mandatory if tls specs defined for FEPExporter
spec.fep.monitoring.prometheus.tls		Prometheus MTLS specs
spec.fep.monitoring.prometheus.tls.certificateName	(created by user)	This is an Optional parameter. These points to Kubernetes TLS secret that contains the certificate of Prometheus. FEPExporter will use this for certificate

Field	Default	Details
		authentication. The certificate itself is stored in the key tls.crt.
spec.fep.monitoring.prometheus.tls.caName	(created by user)	This is an Optional parameter. This point to Kubernetes configmap that contains additional CA the client use to verify a server certificate. The CA is stored in the key ca.crt.
spec.fep.podAntiAffinity	false	Defines that all the pods should not run on same worker node
spec.fep.podDisruptionBudget	false	Allows to maintain minimum number of pods of an application even when some nodes are voluntarily drained for say, maintenance
spec.fep.replicationSlots		List of Patroni permanent replication slots.
spec.fep.replicationSlots.demo_subscription1		The 'demo_subscription1' is the slot name. This name cannot be same as any pod name (e.g., new-fep-sts-01) in the cluster. Otherwise, the slot will not be created.
spec.fep.replicationSlots.type	logical	Must be 'logical' for logical replication
spec.fep.replicationSlots.database	postgres	Specify the database name for logical replication
spec.fep.replicationSlots.plugin	pgoutput	FEP supports 'pgoutput' by default.
spec.fep.usePodName		Optional Setting this key to true will make internal POD communication, both Patroni and Postgres to use hostname, instead of IP address. This is important for TLS as the hostname of the POD is predictable and can be used to create Server Certificate, whereas IP address is unpredictable and cannot be used to create Certificate. There is no negative effect setting this key to true even if TLS (i.e. Server Certificate) is not used.
spec.fep.patroni.tls.certificateName	(created by user)	Optional This point to Kubernetes TLS secret that contains the certificate for Patroni. The certificate itself is stored in the key tls.crt. This field is optional. When this key is set, the Operator will ignore the value in systemCertificates
spec.fep.patroni.tls.caName	(created by user)	Optional This points to Kubernetes configmap that contains additional CA for Patroni to verify client. The CA is stored in the key ca.crt. This field is optional.
spec.fep.postgres.tls.certificateName	(created by user)	Optional This points to Kubernetes TLS secret that contains the certificate for Postgres.

Field	Default	Details
		The certificate itself is stored in the key tls.crt. This field is optional. When this key is set, Operator will ignore the value in systemCertificates
spec.fep.postgres.tls.caName	(created by user)	Optional This point to Kubernetes configmap that contains additional CA for Postgres to verify client. The CA is stored in the key ca.crt. This field is optional.
spec.fep.postgres.tls.privateKeyPassword	(created by user)	Optional This points to Kubernetes secret that contains the password for the above private key. This field is optional.
spec.fep.pgAuditLog.auditLogPath		Directory where auditlog files are stored. This should match the one defined in fepChildCrVal.customPgAudit.log_directory
spec.fep.pgAuditLog.schedules		Schedule to upload auditlog
spec.fep.pgAuditLog.schedules.upload		Upload schedule in crontab format
spec.fep.pgAuditLog.endpoint.protocol		Web server protocol to upload auditlog file. Supported type is http or https
spec.fep.pgAuditLog.endpoint.url		Webserver URL to upload the auditlog files
spec.fep.pgAuditLog.endpoint.customCertificateName		Optional Secret that contains the certificate to setup communication with Web server
spec.fep.pgAuditLog.endpoint.insecure	false	Optional equivalent to curl -insecure option
spec.fep.pgAuditLog.endpoint.authentication		Optional Name of secret that contains the username and password for Basic authentication to Web server
spec.fep.pgAuditLog.endpoint.fileUploadParameter	file	Optional The file upload parameter defined by the web server
spec.fep.pgBadger.schedules.create		The 'create' schedule to create report and upload it to endpoint
spec.fep.pgBadger.schedules.cleanup		The 'cleanup' schedule to delete the report left in container
spec.fep.pgBadger.options.incremental	false	Default: false; When set to True: create incremental report in pgbadger
spec.fep.pgBadger.endpoint.authentication		a secret to contain authentication info to access endpoint support basic auth only
spec.fep.pgBadger.endpoint.customCertificateName		Client certitificate reference in customCertificate CR

Field	Default	Details
spec.fep.pgBadger.endpoint.fileUpload Parameter	file	The file upload parameter defined by the web server
spec.fep.pgBadger.endpoint.insecure	false	equivalent to curl -insecure option
spec.fep.pgBadger.endpoint.url		Web server url to upload the report file
spec.fepChildCrVal.customCertificates		Optional This is an optional parameter, which comprises of the parameters mentioned below. It is an array of elements to define certificates. Used to setup SSL connection between publisher and subscriber clusters for logical replication
spec.fepChildCrVal.customCertificate s.userName		Optional This should be the username of the publisher database. When this parameter is specified, an empty folder is created under FEP Server Container- /tmp/ custom_certs/<username>. The custom certificates are mounted in this empty folder. However, if this parameter is not specified, the section is ignored and folder is not created; hence the certificates are not mounted without it.
spec.fepChildCrVal.customCertificate s.certificateName	(created by user)	Optional This points to Kubernetes TLS secret that contains the custom certificate. The certificate itself is stored in the key tls.crt.
spec.fepChildCrVal.customCertificate s.caName	(created by user)	Optional This points to Kubernetes configmap that contains CA certificate to verify server. The CA is stored in the key ca.crt.
spec.fepChildCrVal.backup		Optional This section is defined to enable fepbackup sidecar for cluster backup feature.
spec.fepChildCrVal.backup.image.ima ge	<omitted>	FEP backup container image to be used quay.io/fujitsu/fujitsu-enterprise-postgres-14-backup:ubi8-14-0.0 It is optional. Image line is omitted by default. In such a case, it will pick up URL of image from operator container environment. If you specify the image, Operator will take that image to deploy backup container
spec.fepChildCrVal.backup.image.pull Policy	IfNotPresent	
spec.fepChildCrVal.backup.mcSpec.li mits	cpu: 0.2	

Field	Default	Details
	memory: "300Mi"	
spec.fepChildCrVal.backup.mcSpec.requests	cpu: 0.1 memory: "200Mi"	
spec.fepChildCrVal.backup.pgbackrestParams	[global] repo1-retention-full=7 repo1-retention-full-type=time log-path=/database/log/backup	"|" When nothing is specified, and the parameter set in pgbackrest.conf is described from the line below.
spec.fepChildCrVal.backup.pgbackrestKeyParams		Optional "|" is fixed, and the following line describes the parameters to be set in pgbackrest.conf. The value described by this parameter is masked with *****.
spec.fepChildCrVal.backup.caName		Optional Set to use a CA file other than the system default. Specifies the name of the Configmap you created.
spec.fepChildCrVal.backup.repoKeySecretName		Optional Specifies the name of the Kubernetes Secret generated from the object storage key file. Specify in array format.
spec.fepChildCrVal.backup.schedule.num	0	Number of schedules to set The maximum number of backup schedules is 5.
spec.fepChildCrVal.backup.scheduleN.schedule	"|"	Backup schedule in cron format. The date and time is UTC time.
spec.fepChildCrVal.backup.scheduleN.type	"|"	full: Perform a full backup (Back up the contents of the database cluster). incr ー Perform an incremental backup (Back up only the database cluster files that were changed to the last backup migration).
spec.fepChildCrVal.backup.scheduleN.repo	1	Optional Gets a backup in the specified repository. The range is 1 to 256.
spec.fepChildCrVal.customPgAudit	[output] logger = 'auditlog' log_directory = '/database/log/audit' log_truncate_on_rotation = on log_filename = 'pgaudit-%a.log' log_rotation_age = 1d log_rotation_size = 0	PgAudit file content

Field	Default	Details
	[rule]	
spec.fepChildCrVal.customPgHba	# define pg_hba custom rules here to be merged with default rules. # TYPE DATABASE USER ADDRESS METHOD	Entries to be inserted into pg_hba.conf

spec.fepChildCrVal.customPgParams	# define custom postgresql.conf parameters below to override defaults.
	# Current values are as per default FEP deployment
	shared_preload_libraries='pgx_datamasking,pg_prewarm,pg_stat_statements,fsep_operator_security'
	session_preload_libraries='pg_prewarm'
	max_prepared_transactions = 100
	max_worker_processes = 30
	max_connections = 100
	work_mem = 1MB
	maintenance_work_mem = 12MB
	shared_buffers = 128MB
	effective_cache_size = 384MB
	checkpoint_completion_target = 0.8
	# tcp parameters
	tcp_keepalives_idle = 30
	tcp_keepalives_interval = 10
	tcp_keepalives_count = 3
	# logging parameters in default fep installation
	# if log volume is not defined, log_directory should be
	# changed to '/database/userdata/data/log'
	log_directory = '/database/log'
	log_filename = 'logfile-%a.log'
	log_file_mode = 0600
	log_truncate_on_rotation = on
	log_rotation_age = 1d
	log_rotation_size = 0
	log_checkpoints = on
	log_line_prefix = '%e %t [%p]: [%l-1] user=%u,db=%d,app=%a,client=%h'
	log_lock_waits = on
	log_autovacuum_min_duration = 60s
	logging_collector = on
		Postgres configuration in postgresql.conf
		If the FEP server container utilizes images with a FEPBaseVersion less than 15, exclude fsep_operator_security from the configuration.

Field	Default	Details

	pgaudit.config_file='/opt/app-root/src/pgaudit-cfg/pgaudit.conf'
	log_replication_commands = on
	log_min_messages = WARNING
	log_destination = stderr
	# wal_archive parameters in default fep installation
	archive_mode = on
	archive_command = 'pgbackrest --stanza=backupstanza --config=/database/userdata/pgbackrest.conf archive-push %p'
	wal_level = replica
	max_wal_senders = 12
	wal_keep_segments = 64
	track_activities = on
	track_counts = on
	password_encryption = 'md5'	

spec.fepChildCrVal.storage.dataVol		Mandatory volume
spec.fepChildCrVal.storage.dataVol.size	2Gi (**)	Size of data volume. Data volume must be specified
spec.fepChildCrVal.storage.dataVol.storageClass	<omitted> (*)	StorageClass for data volume: When this line is omitted, the PV created will use default storage class in the Kubernetes cluster
spec.fepChildCrVal.storage.dataVol.accessModes	<omitted> (*)	accessModes for data volume: Specified as an array of accessModes e.g. [ReadWriteMany] If omitted, it will be treated as [ReadWriteOnce]
spec.fepChildCrVal.storage.walVol		Mandatory volume
spec.fepChildCrVal.storage.walVol.size	1200Mi (**)	Size of WAL volume. WAL volume must be specified
spec.fepChildCrVal.storage.walVol.storageClass	<omitted> (*)	StorageClass for WAL volume: When this line is omitted, the PV created will use default storage class in the Kubernetes cluster
spec.fepChildCrVal.storage.walVol.accessModes	<omitted> (*)	accessModes for WAL volume: Specified as an array of accessModes e.g. [ReadWriteMany] If omitted, it will be treated as [ReadWriteOnce]

Field	Default	Details
spec.fepChildCrVal.storage.tablespaceVol		Optional volume
spec.fepChildCrVal.storage.tablespaceVol.size	512Mi (**)	Size of tablespace volume. This volume is optional and can be omitted
spec.fepChildCrVal.storage.tablespaceVol.storageClass	<omitted> (*)	StorageClass for tablespace volume: When this line is omitted, the PV created will use default storage class in the Kubernetes cluster
spec.fepChildCrVal.storage.tablespaceVol.accessModes	<omitted> (*)	accessModes for tablespace volume: Specified as an array of accessModes e.g. [ReadWriteMany] If omitted, it will be treated as [ReadWriteOnce]
spec.fepChildCrVal.storage.archivewalVol		Mandatory if backup section is defined. Optional otherwise
spec.fepChildCrVal.storage.archivewalVol.size	1Gi (**)	Size of archivewal volume. This volume is optional and can be omitted
spec.fepChildCrVal.storage.archivewalVol.storageClass	<omitted> (*)	StorageClass for Archived WAL volume: When this line is omitted, the PV created will use default storage class in the Kubernetes cluster When the number of instance is more than 1 and backup is not done on S3, both archivewalVol and backupVol must be hosted on Shared storage such as NFS with respective storageClass
spec.fepChildCrVal.storage.archivewalVol.accessModes	<omitted> (*)	accessModes for Archived WAL volume: Specified as an array of accessModes e.g. [ReadWriteMany] If omitted, it will be treated as [ReadWriteOnce] When the number of instance is more than 1 and backup is not done on S3, both archivewalVol and backupVol must be hosted on Shared storage such as NFS with accessMode set to [ReadWriteMany]
spec.fepChildCrVal.storage.logVol		Optional volume
spec.fepChildCrVal.storage.logVol.size	1Gi (**)	Size of log volume. This volume is optional and can be omitted

Field	Default	Details
spec.fepChildCrVal.storage.logVol.storageClass	<omitted> (*)	StorageClass for log volume: When this line is omitted, the PV created will use default storage class in the Kubernetes cluster
spec.fepChildCrVal.storage.logVol.accessModes	<omitted> (*)	accessModes for log volume: Specified as an array of accessModes e.g. [ReadWriteMany] If omitted, it will be treated as [ReadWriteOnce]
spec.fepChildCrVal.storage.backupVol		Mandatory if backup section is defined. Optional otherwise
spec.fepChildCrVal.storage.backupVol.size	2Gi (**)	Size of backup volume. This volume is optional and can be omitted
spec.fepChildCrVal.storage.backupVol.storageClass	<omitted> (*)	StorageClass for backup volume: When this line is omitted, the PV created will use default storage class in the Kubernetes cluster When the number of instance is more than 1 and backup is not done on S3, both archivewalVol and backupVol must be hosted on Shared storage such as NFS with respective storageClass
spec.fepChildCrVal.storage.backupVol.accessModes	<omitted> (*)	accessModes for backup volume: Specified as an array of accessModes e.g. [ReadWriteMany] If omitted, it will be treated as [ReadWriteOnce] When the number of instance is more than 1 and backup is not done on S3, both archivewalVol and backupVol must be hosted on Shared storage such as NFS with accessMode set to [ReadWriteMany]
spec.fepChildCrVal.sysUsers.pgAdminPassword	<omitted>	Password for user "postgres" Available character types Alphanumeric characters (A-Z, a-z), numbers (0 -9), symbols (~! @ # $^ & * () - = < >,.? ; : |/+) If this parameter is omitted, the Operator automatically generates a password. If the FEP server container uses an image with a FEPBaseVersion less than 15, be sure to specify this parameter.
spec.fepChildCrVal.sysUsers.pgdb	mydb (*)	Database to be created during provisioning

Field	Default	Details
		Available character types Alphanumeric characters (A-Z, a-z), numbers (0 -9), and underscores (_) However, you cannot start with a number. Upper case letters are treated as lower case letters. Maximum string length 63 characters
spec.fepChildCrVal.sysUsers.pguser	mydbuser (*)	Database user to be created during provisioning Available character types Alphanumeric characters (A-Z, a-z), numbers (0 -9), and underscores (_) However, you cannot start with a number. Upper case letters are treated as lower case letters. Maximum string length 63 characters
spec.fepChildCrVal.sysUsers.pgpassword	mydbpassword	Password for database user pguser Available character types Alphanumeric characters (A-Z, a-z), numbers (0 -9), symbols (~! @ # $^ & * () - = < >,.? ; : |/+)
spec.fepChildCrVal.sysUsers.pgrepluser	repluser (*)	Database user for replication Available character types Alphanumeric characters (A-Z, a-z), numbers (0 -9), and underscores (_) However, you cannot start with a number. Maximum string length 63 characters
spec.fepChildCrVal.sysUsers.pgreplpassword	repluserpwd	Password for database user repluser Available character types Alphanumeric characters (A-Z, a-z), numbers (0 -9), symbols (~! @ # $^ & * () - = < >,.? ; : |/+)
spec.fepChildCrVal.sysUsers.tdepassphrase	tde-passphrase	TDE keystore passphrase
spec.fepChildCrVal.sysUsers.pgRewindUser	rewind_user	Database user for Rewind Available character types Alphanumeric characters (A-Z, a-z), numbers (0 -9), and underscores (_)

Field	Default	Details
		However, you cannot start with a number. Maximum string length 63 characters
spec.fepChildCrVal.sysUsers.pgRewindUserPassword	rewind_password	Password for database user rewinduser Available character types Alphanumeric characters (A-Z, a-z), numbers (0 -9), symbols (~! @ # $^ & * () - = < >,.? ; : |/+)
spec.fepChildCrVal.sysUsers.pgMetricsUser		Optional user for FEPExporter connection. Can be defined afterwards Available character types Alphanumeric characters (A-Z, a-z), numbers (0 -9), and underscores (_) However, you cannot start with a number. Upper case letters are treated as lower case letters. Maximum string length 63 characters
spec.fepChildCrVal.sysUsers.pgMetricsUserPassword		Optional Password for metrics user. Can be defined afterwards Available character types Alphanumeric characters (A-Z, a-z), numbers (0 -9), symbols (~! @ # $^ & * () - = < >,.? ; : |/+)
spec.fepChildCrVal.sysUsers.pgAdminTls.certificateName		This points to Kubernetes TLS secret that contains the certificate of Postgres user “postgres”. Patroni will use this for certificate authentication. The certificate itself is stored in the key tls.crt. This field is optional.
spec.fepChildCrVal.sysUsers.pgAdminTls.caName		This points to Kubernetes configmap that contains additional CA the client use to verify a server certificate. The CA is stored in the key ca.crt. This field is optional.
spec.fepChildCrVal.sysUsers.pgAdminTls.sslMode	prefer	Specify the type of TLS negotiation with the server. - disable - allow - prefer - require

Field	Default	Details
		- verify-ca - verify-full
spec.fepChildCrVal.sysUsers.pgreplUserTls.certificateName		This points to Kubernetes TLS secret that contains the certificate of Postgres user “repluser”. Patroni will use this for certificate authentication. The certificate itself is stored in the key tls.crt. This field is optional.
spec.fepChildCrVal.sysUsers.pgreplUserTls.caName		This points to Kubernetes configmap that contains additional CA the client use to verify a server certificate. The CA is stored in the key ca.crt. This field is optional.
spec.fepChildCrVal.sysUsers.pgreplUserTls.sslMode	prefer	Specify the type of TLS negotiation with the server. - disable - allow - prefer - require - verify-ca - verify-full
spec.fepChildCrVal.sysUsers.pgRewindUserTls.certificateName		This points to Kubernetes TLS secret that contains the certificate of Postgres user “rewinduser”. Patroni will use this for certificate authentication. The certificate itself is stored in the key tls.crt. This field is optional.
spec.fepChildCrVal.sysUsers.pgRewindUserTls.caName		This points to Kubernetes configmap that contains additional CA the client use to verify a server certificate. The CA is stored in the key ca.crt. This field is optional.
spec.fepChildCrVal.sysUsers.pgRewindUserTls.sslMode	prefer	Specify the type of TLS negotiation with the server. - disable - allow - prefer - require - verify-ca - verify-full
spec.fepChildCrVal.sysUsers.pgMetricsUserTls.certificateName		Optional This points to Kubernetes TLS secret that contains the certificate of Postgres user defined by pgMetricsUser. FEPExporter will use this for certificate

Field	Default	Details
		authentication. The certificate itself is stored in the key tls.crt.
spec.fepChildCrVal.sysUsers.pgMetricsUserTls.caName		Optional This points to Kubernetes configmap that contains additional CA the client use to verify a server certificate. The CA is stored in the key ca.crt.
spec.fepChildCrVal.sysUsers.pgMetricsUserTls.sslMode	prefer	Optional Specify the type of TLS negotiation when FEPExporter connects to FEP server. - disable - allow - prefer - require - verify-ca - verify-full
spec.fepChildCrVal.sysTde	(*)	Optional If the user selects a file-based TDE, you do not need to define it. Required when implementing TDE with a key management system (KMS).
spec.fepChildCrVal.sysTde.tdeType	(*)	Optional The parameter itself is optional, but required when spec.fepChildCrVal.sysTde is defined. Specify tdek.
spec.fepChildCrVal.sysTde.tdek		Optional Defines the connection information to the KMS. Required when tdek is specified for spec.fepChildCrVal.sysTde.tdeType.
spec.fepChildCrVal.sysTde.tdek.targetKmsName	(*)	Specifies the KMS name. It has the same value as name in sysTde.tdek.kmsDefinition.
spec.fepChildCrVal.sysTde.tdek.targetKeyId		Specifies the key ID (Identifier attribute in KMIP) attached to the encryption key in KMS. When you update this parameter, the Operator automatically updates the master key.
spec.fepChildCrVal.sysTde.tdek.kmsDefinition		Specifies KMS connection information. Specify in array format.
spec.fepChildCrVal.sysTde.tdek.kmsDefinition[0].name	(＊)	The name given to the KMS (key management system name) specified in spec.fepChildCrVal.sysTde.tdek.targetKmsName. The KMS name must be a string of no more than 63 characters beginning with

Field	Default	Details
		a-z, consisting of a-z, numbers (0 -9), and underscores. Upper and lower case letters are the same.
spec.fepChildCrVal.sysTde.tdek.kmsDefinition[0].type	(*)	Specifies the type of KMS. Currently, only kmip is available.
spec.fepChildCrVal.sysTde.tdek.kmsDefinition[0].address	(*)	Specifies the host name or IP address of the KMS.
spec.fepChildCrVal.sysTde.tdek.kmsDefinition[0].port	(*)	Specifies the KMS port.
spec.fepChildCrVal.sysTde.tdek.kmsDefinition[0].authMethod	(*)	Specifies the authentication method in KMS. Currently, the only possible value is cert.
spec.fepChildCrVal.sysTde.tdek.kmsDefinition[0].sslpassphrase		Optional Specifies the passphrase of the client certificate private key file when connecting to KMS. This can be omitted if no passphrase is set in the private key file.
spec.fepChildCrVal.sysTde.tdek.kmsDefinition[0].cert		Optional Specifies the name of the Secret/ ConfigMap containing the certificate file, etc., when cert is specified as authMethod.
spec.fepChildCrVal.sysTde.tdek.kmsDefinition[0].cert.certificateName	(*)	Specifies the TLS Secret name that contains the client certificate and private key for TLS communication with KMS.
spec.fepChildCrVal.sysTde.tdek.kmsDefinition[0].cert.caName	(*)	Specifies the ConfigMap name that contains the file name of the SSL Certificate Authority certificate. Used to verify the server certificate of the connection destination.
spec.fepChildCrVal.systemCertificates.key		Use spec.fep.postgres.tls specification instead.
spec.fepChildCrVal.systemCertificates.crt		Use spec.fep.postgres.tls specification instead.
spec.fepChildCrVal.systemCertificates.cacrt		Use spec.fep.postgres.tls specification instead.
spec.fepChildCrVal.autoscale.scaleout.policy	off	Specifies whether to use the automatic scale out feature and the metric to base on. Specify one of the following: - cpu_utilization (if based on CPU utilization) - connection_number (if based on number of connections) - off (without automatic scale out) If omitted, off is assumed.
spec.fepChildCrVal.autoscale.scaleout.threshold	40	Specifies an integer as the threshold for performing scale out.

Field	Default	Details
		- When cpu_utilization is specified for policy Specifies the average CPU utilization as a percentage for the threshold. If this option is omitted, 40 (40%) is assumed. - When connection_number is specified for policy Specifies the average value of the number of connections as a threshold. If you omit this option, 40 is assumed.
spec.fepChildCrVal.autoscale.scaleout.metricName	pg_capacity_connection_average	Specify this parameter if policy is connection_number. Ignored if policy is cpu_utilization. The custom metrics server must publish the average number of connections in the FEP cluster under this name. If omitted, pg_capacity_connection_average is assumed.
spec.fepChildCrVal.autoscale.scaleout.stabilizationWindowSeconds	0	This parameter controls the stability of scaling (variation in the number of replicas). Scale out is not performed unless the metric exceeds the threshold for more than the number of seconds specified for this parameter. If omitted, 0 is assumed.
spec.fepChildCrVal.autoscale.limits.maxReplicas	2	Maximum number of replicas (0 to 15) (Value out of range) Do not perform auto scale out
spec.fepChildCrVal.restore		Optional Defines to restore specified backup data stored in object storage.
spec.fepChildCrVal.restore.pgbackrestParams		Optional "|" is fixed, and the following line describes the parameters to be set in pgbackrest.conf. Specifies the object storage where the backup data is stored. If you need to use a root certificate other than the default, specify the following: repo1-storage-ca-path =/pgbackrest/storage-certs/filename The CA file is registered in ConfigMap and the ConfigMap name is listed in spec.fepChildCrVal.restore.caName.
spec.fepChildCrVal.restore.pgbackrestKeyParams		Optional "|" is fixed, and the following line describes the parameters to be set in

Field	Default	Details
		pgbackrest.conf. The value described by this parameter is masked with *****. Specify the parameter you want to mask, such as a password.
spec.fepChildCrVal.restore.caName		Optional Set to use a CA file other than the system default. Specifies the name of the ConfigMap created, in list format. The ConfigMap specified is mounted in /pgbackrest/storage-certs.
spec.fepChildCrVal.restore.repoKeySecretName		Optional Specifies the name of the Kubernetes Secret generated from the object storage key file. Specify in array format. The specified Secret will be mounted in /pgbackrest/storage-key.
spec.fepChildCrVal.restore.mcSpec.limits	cpu: 200m memory: 300Mi	Optional CPU and memory allocated to the container performing the restore
spec.fepChildCrVal.restore.mcSpec.requests	cpu: 100m memory: 200Mi	Optional CPU and memory allocated to the container performing the restore
spec.fepChildCrVal.restore.restoretype	latest	Optional Select the type of restore (latest or PITR).
spec.fepChildCrVal.restore.restoredate		Optional Specifies the date to restore when spec.fepChildCrVal.restore.restoretype is "PITR".
spec.fepChildCrVal.restore.restoretime		Optional Specifies the time to restore when spec.fepChildCrVal.restore.restoretype is "PITR".
spec.fepChildCrVal.restore.image		Optional Image of the container to perform the restore It is omitted by default. In this case, the URL for image is obtained from the operator container environment.
spec.fepChildCrVal.restore.imagePullPolicy	IfNotPresent	Optional
spec.fepChildCrVal.upgrade		Optional When this field is defined, a major version upgrade is performed.

Field	Default	Details
		However, if spec.fepChildCrVal.restore is defined, the FEPCluster build stops.
spec.fepChildCrVal.upgrade.sourceCluster		Specifies the FEPClusterCR name from which to migrate data. Required if spec.fepChildCrVal.upgrade is defined.
spec.fepChildCrVal.upgrade.mcSpec.limits	cpu: 200m memory: 300Mi	Optional Specifies the maximum number of resources to allocate to the upgrade execution container.
spec.fepChildCrVal.upgrade.mcSpec.requests	cpu: 100m memory: 200Mi	Optional Specifies the lower limit of resources allocated to the upgrade execution container.
spec.fepChildCrVal.upgrade.image		Optional By default, the URL of image is obtained from the operator container environment.
spec.fepChildCrVal.upgrade.imagePullPolicy	IfNotPresent	Optional Specifies the pull policy for the container image. - Always - IfNotPresent - Never
spec.fepChildCrVal.upgrade.source.pgAdminTls.certificateName		Optional If you do not define spec.fepChildCrVal.sysUsers.pgAdminTls.certificateName for the data source, it points to the Kubernetes TLS secret that contains the certificate for the Postgres user "postgres" in the data source. If the data source FEP has set the authentication method for the upgrade execution container to "cert", then the upgrade execution container uses the certificate defined as secret.
spec.fepChildCrVal.upgrade.destination.pgAdminTls.certificateName		Optional If you have not defined the spec.fepChildCrVal.sysUsers.pgAdminTls.certificateName of the newly created FEPCluster, it points to the Kubernetes TLS secret that contains the certificate of the Postgres user “postgres ” in the data source. If you create a new FEP with the "cert" authentication method for the upgrade execution container, the upgrade

Field	Default	Details
		execution container uses the certificate defined as secret.
spec.fepChildCrVal.upgrade.storage		Optional Defines the storage for storing dump files.
spec.fepChildCrVal.upgrade.storage.storageClass		Optional If omitted, the default storage class for your environment is used.
spec.fepChildCrVal.upgrade.storage.size	2Gi	Optional Specifies the size of the storage to store the dump file.
spec.fepChildCrVal.upgrade.storage.accessModes	ReadWriteOnce	Optional accessModes for store the dump file Specified as an array of accessModes e.g. [ReadWriteMany] If omitted, it will be treated as [ReadWriteOnce]
spec.fep.remoteLogging.enable		Set to true to forward logs from fluentbit to fluentd
spec.fep.remoteLogging.image		Optional Fluentbit image to be used. If not specified, Operator will use the latest version that is supported by the Operator.
spec.fep.remoteLogging.pullPolicy	IfNotPresent	Optional

Note

- (*) - These parameters can be specified only at creation time and should not be changed. Any change to these parameters will be ignored and will not have any effect on FEP cluster functioning.
- (**) - The storage volumes size can be increased provided underlying storage supports the operation. Optional volumes can be specified only at initial FEP cluster creation. If an optional volume is added later, operator will ignore it and no action will be taken.
- User should do or remove unsupported CR changes manually.
- spec.fep.postgres.tls CR specification should be used instead of spec.fepChildCrVal.systemCertificates. The lateral spec can still be used, however spec.fep.postgres.tls gives better flexibility to control MTLS access of the cluster.
- Either spec.fep.postgres.tls specification (old specification) or spec.fepChildCrVal.systemCertificates should be used. They should not be used interchangeable.
- Server certificate specified under spec.fep.postgres.tls can be rotated by changing the secret and executing reload (e.g. using FEPAction); however for others specified in the CR, it is required to do restart of the PoDs

While in running state - following value will dynamically appear in the FEPCluster to reflect the cluster status

Field name	Details
status.fepStatus.fepClusterReady	Will be true or false to reflect if the whole cluster is ready. Kubernetes cluster information is fetched to check number of instances 'READY' & 'RUNNING' is equal to number of Configured instances.

Note

“fepClusterReady" flag will be set at first FEPCluster creation time only. fepClusterReady flag does not participate in the next reconcillation loop)

1.2 Custom Resource Parameters

This section explains the Custom Resource Parameters.

1.2.1 FEPCluster Custom Resource Parameters

Category	Details
CRD Name	`FEPCluster`
`Definition`	`///`

Operations	`Create: kubectl create -f fepcluster.yaml`
	`Delete: kubectl delete fepcluster <clusername>`
	`Update: kubectl apply -f fepcluster.yaml`
	`List: kubectl get fepcluster`

FEPCluster CR Example

```
apiVersion: fep.fujitsu.io/v2
kind: FEPCluster
metadata:
  name: new-fep
  namespace: new-fep
spec:
  fep:
///
        wuC4
        -----END CERTIFICATE-----
```

It should also be noted that all the passwords / passphrase and certificates will be masked after the creation of the CR. This includes

- Also, inital pgAdminPassword: admin-password
- pgpassword: mydbpassword
- pgreplpassword: repluserpwd
- tdepassphrase: tde-passphrase
- pgRewindPassword: rewind_passoword (Optional - if defined)
- pgMetricsPassword: metrics_password (Optional - if defined)
- sslpassphrase under sysTde.tdek.kmsDefinition (if defined)
- certificate.key

- certificate.crt
- certificate.cacrt

Values of child CRs at the time of initial deployment of cluster, are stored in FEPCluster under fepChildCrVals, e.g. for Server certificates, Configuration of FEP, User details.

All fields for FEPCluster CR and its child CRs should be managed through FEPCluster CR only. Operator will reflect the changes to respective child CR to be processed. The fields that not allowed to change will not be reflected from parent to child CR and hence will not have any affect.

1.2.2 FEP Cluster Configuration

Configuration of all aspects of FEP Cluster is done through FEPCluster CR only.

All fields for FEPCluster CR and its child CRs should be managed through FEPCluster CR only. Operator will reflect the changes to respective child CR to be processed. The fields that not allowed to change will not be reflected from parent to child CR and hence will not have any affect. Refer to "1.1 FEPCluster Parameter" for details.

All child CRs are marked as internal objects in RedHat OCP and will not appear on console. However, it can be checked on command line using oc or kubectl commands.

Following table shows Child CRs of FEPCluster CR and respective sections in parent CR related to given child CR.

Configuration changes are made in these sections will update allowable fields only in corresponding child CR.

Child CR Name	Relevant sections in FEP Cluster CR
FEPBackup	spec.fepChildCrVal.backup
FEPCert	spec.fepChildCrVal.systemCertificates
FEPConfig	spec.fepChildCrVal.customPgAudit spec.fepChildCrVal.customPgHba spec.fepChildCrVal.customPgParams
FEPUser	spec.fepChildCrVal.sysUsers
FEPVolume	spec.fepChildCrVal.storage

1.2.3 FEPConfig Child Custom Resource Parameters

Field	Default	Details
metadata.name	<same-as-in-FEPCluster>	This value is inherited from parent FEPCluster CR
metadata.namespace	<same-as-in-FEPCluster>	This value is inherited from parent FEPCluster CR
spec.customPgAudit	All line specified in spec.fepChildCrVal.customPgAudit of FEPCluster CR	Audit rules can be updated in this section. Requires restart. Note: initial values inherited once only at start. Changes to FEPConfig directly
spec.customPgHba	All line specified in spec.fepChildCrVal.customPgHba of FEPCluster CR	pg_hba rules can be added in this section Note: Inherited once at start. Changes to FEPConfig directly
spec.customPgParams	All line specified in spec.fepChildCrVal.customPgParams of FEPCluster CR	All postgres parameters are listed here to overwrite defaults. Note: Inherited once at start. Changes to FEPConfig directly

Field	Default	Details
spec.replicationSlots		Optional: Details of replication slots if defined in FEPCluster

Example of FEPConfig CR created

```
apiVersion: fep.fujitsu.io/v1
kind: FEPConfig
metadata:
  name: new-fep-19ncfg
  namespace: cfg-expt
spec:
  sysExtraLogging: false
  customPgAudit: |
    # define pg audit custom params here to override defaults.
    # if log volume is not defined, log_directory should be
    # changed to '/database/userdata/data/log'
    [output]
    logger = 'auditlog'
    log_directory = '/database/log/audit'
    log_truncate_on_rotation = on
    log_filename = 'pgaudit-%a.log'
    log_rotation_age = 1d
    log_rotation_size = 0
    [rule]

  customPgHba: |
    # define pg_hba custom rules here to be merged with default rules.
    # TYPE     DATABASE        USER        ADDRESS        METHOD
  customPgParams: |+
    # define custom postgresql.conf parameters below to override defaults.
    # Current values are as per default FEP deployment
    shared_preload_libraries='pgx_datamasking,pgaudit,pg_prewarm,pg_stat_statements'
    session_preload_libraries='pg_prewarm'
    max_prepared_transactions = 100
    max_worker_processes = 20
    max_connections = 100
    work_mem = 1MB
    maintenance_work_mem = 20MB
    shared_buffers = 128MB
    effective_cache_size = 384MB
    checkpoint_completion_target = 0.8
    pgx_global_metacache = 10MB
    temp_buffers = 10MB

    # tcp parameters
    tcp_keepalives_idle = 30
    tcp_keepalives_interval = 10
    tcp_keepalives_count = 3

    # logging parameters in default fep installation
    # if log volume is not defined, log_directory should be
    # changed to '/database/userdata/data/log'    log_directory = '/database/log'
    log_filename = 'logfile-%a.log'
    log_file_mode = 0600
    log_truncate_on_rotation = on
    log_rotation_age = 1d
    log_rotation_size = 0
    log_checkpoints = on
    log_line_prefix = '%e %t [%p]: [%l-1] user=%u,db=%d,app=%a,client=%h'
    log_lock_waits = on
```

```
log_autovacuum_min_duration = 60s
logging_collector = on
pgaudit.config_file= '/opt/app-root/src/pgaudit-cfg/pgaudit.conf'
log_replication_commands = on
log_min_messages = WARNING
log_destination = stderr

# wal_archive parameters in default fep installation
archive_mode = on
wal_level = replica
max_wal_senders = 10
wal_keep_segments = 64
wal_sender_timeout = 60s
track_activities = on
track_counts = on
```

1.2.4 FEPUser Child Custom Resource Parameters

Field	Default	Details
metadata.name	<same-as-in-FEPCluster>	This value is inherited from parent FEPCluster CR
metadata.namespace	<same-as-in-FEPCluster>	This value is inherited from parent FEPCluster CR
spec.pgAdminPassword	spec.fepChildCrVal.users.pgAdminPassword of FEPCluster CR	postgres superuser password. Masked once secret is created/changed Note: initial values inherited once only at start. Changes to FEPUser directly
spec.pgdb	spec.fepChildCrVal.users.pgdb of FEPCluster CR	Name of a user database Note: Created once only at start. Cannot be changed
spec.pgpassword	spec.fepChildCrVal.users.pgpassword of FEPCluster CR	Password for superuser for user database pgdb. Masked once secret is created/changed Note: initial values inherited once only at start. Changes to FEPUser directly
spec.pguser	spec.fepChildCrVal.users.pguser of FEPCluster CR	Name of a user database Note: Created once only at start. Cannot be changed
spec.pgrepluser	spec.fepChildCrVal.users.pgrepluser of FEPCluster CR	Name of a database user for replication
spec.pgreplpassword	spec.fepChildCrVal.users.pgreplpassword of FEPCluster CR	Password for pgrepluser
spec.tdepassphrase	spec.fepChildCrVal.users.tdepassphrase of FEPCluster CR	Passphrase for encrypting/decrypting keystore file which contains the TDE encryption key
spec.pgRewindUser	rewind_user	Database user for Rewind
spec.pgRewindUserPassword	rewind_password	Password for database user rewinduser
spec.pgMetricsUser	spec.fepChildCrVal.sysUsers.pgMetricsUser	Optional See details in FEPCluster CR

Field	Default	Details
spec.pgMetricsPassword	spec.fepChildCrVal.sysUsers.pgMetricsPassword	Optional See details in FEPCluster CR
spec.pgAdminTls	spec.fepChildCrVal.sysUsers.pgAdminTls	Optional section See details in FEPCluster CR
spec.pgrepluserTls	spec.fepChildCrVal.sysUsers.pgrepluserTls	Optional section See details in FEPCluster CR
spec.pgRewindUserTls	spec.fepChildCrVal.sysUsers.pgRewindUserTls	Optional section See details in FEPCluster CR
spec.pgMetricsUserTls	spec.fepChildCrVal.sysUsers.pgMetricsUserTls	Optional section See details in FEPCluster CR

Example of FEPUser CR created

```
apiVersion: fep.fujitsu.io/v1
kind: FEPUser
metadata:
  name: new-fep-19n
  namespace: testswatiproject
spec:
  pgAdminPassword: '***************'
  pgdb: mydb
  pgpassword: '***************'
  pgreplpassword: '***************'
  pgrepluser: repluser
  pguser: mydbuser
  tdepassphrase: '***************'
  sysExtraLogging: false
  pgRewindUser: rewind_user
  pgRewindUserPassword: rewind_password
  pgAdminTls:
    certificateName: admin-client-certs-secret
    caName: admin-ssl-rootcert-configmap
    sslMode: prefer
  pgrepluserTls:
    certificateName: repluser-client-certs-secret
    caName: repluser-ca-name-configmap
    sslMode: prefer
  pgRewindUserTls:
    certificateName: rewinduser-client-certs-secret
    caName: rewinduser-ca-name-configmap
    sslMode: prefer
```

Note

- Password and Passphrase are masked in output from CR. The original values can still be found in the respective Kubernetes secrets and configmaps.
- TDE is enabled by default with given tdepassphrase and must have a value.
- TDE is enabled by using the key tdepassphrase with the desired passphrase.Do not remove this key once TDE is enabled. Otherwise, the database may go into a crash loop. If the Cluster is running on Async Replication and a failover/switchover occured during the crash loop, there could be data lost. The team is looking at preventing the deletion of this passphrase from Operator even if customer tries to remove it in customer resource.

- Database users and their passwords managed by the FEPUser CR should not be changed in the SQL interface. Inconsistencies with the information managed by the operator can cause problems with operator operation. If you make changes in the SQL interface, use the SQL interface again to restore the original state.

1.2.5 FEPVolume Child Custom Resource Parameters

1.2.5.1 Create Volumes

Volumes for the cluster nodes(pods) are initially created in accordance with the values set in fepChildCrVal’ storage section of the parent FEPCluster CR.

The parent FEPCluster CR creates a child FEPVolume CR with the respective startup values and the relevant controller(FEPColume Controller) takes care of creating the required volumes.

Note

- After you create the FEPCluster for the first time, you cannot add new volumes later or modify the storageClass or accessModes.
- You can resize the initially created volume only if the underlying storageClass supports dynamic resizing.

Below is the schema of the FEPVolume CR:

Field	Mandatory	Sub-Field	Default	Description
archivewalVol	No	size	1Gi	Volume size of the archive log. Refer to "Estimating Database Disk Space Requirements" in the FUJITSU Enterprise Postgres Installation and Setup Guide for Server to help you design the size.
		storageClass	Defaults to platform default if omitted	SC is only set at start
		accessModes	Defaults to ReadWriteOnce if omitted	Access mode is only set at start
backupVol	No	size	2Gi	Volume size of the backup. Estimate based on the following formula: (full backup generations + incr backup generations + 1) * dataVol size
		storageClass	Defaults to platform default if omitted	SC is only set at start
		accessModes	Defaults to ReadWriteOnce if omitted	Access mode is only set at start
dataVol	Yes	size	2Gi	Volume size of the data. Refer to "Estimating Database Disk Space Requirements" in the FUJITSU Enterprise Postgres Installation and Setup Guide for Server and base the design on table/index size.
		storageClass	Defaults to platform default if omitted	SC is only set at start

Field	Mandatory	Sub-Field	Default	Description
		accessModes	Defaults to ReadWriteOnce if omitted	Access mode is only set at start
logVol	No	size	1Gi	Volume size of the log. If you change the log output level (default: WARNING), measure the actual amount of log output in a test environment.
		storageClass	Defaults to platform default if omitted	SC is only set at start
		accessModes	Defaults to ReadWriteOnce if omitted	Access mode is only set at start
tablespaceVol	No	size	512Mi	Volume size of the tablespace. When using tablespaces, as with dataVol, you should refer to "Estimating Database Disk Space Requirements" in the FUJITSU Enterprise Postgres Installation and Setup Guide for Server for information on sizing.
		storageClass	Defaults to platform default if omitted	SC is only set at start
		accessModes	Defaults to ReadWriteOnce if omitted	Access mode is only set at start
walVol	Yes	size	1200Mi	Volume size of the transaction log. Refer to "Estimating Database Disk Space Requirements" in the FUJITSU Enterprise Postgres Installation and Setup Guide for Server to help you design the size. Note that the default value for max_wal_size is 1 GB.
		storageClass	Defaults to platform default if omitted	SC is only set at start
		accessModes	Defaults to ReadWriteOnce if omitted	Access mode is only set at start

The 'accessMode' is been incorporated for the inclusion of pgBadger layer later. Giving it a shared volume capability will allow pgBadger Container to read logs from multiple server instance (master / replica) and expose it via a WebServer.

1.2.5.2 Delete Volumes

Equivalent Kubernetes command: kubectl delete FEPVolume <cr_name>

This operation will remove all the PVCs and possibly PVs depending on the default reclaimPolicy of the storageclass used per volume.

With right backup and restore integration by customer, they may not need volumes to be persisted.

Note

Do not delete this CR unless the Cluster has been removed.

Example of FEPVolume CR created

```
apiVersion: fep.fujitsu.io/v1
kind: FEPVolume
metadata:
  name: new-fep-19n
  namespace: testswatiproject
spec:
  archivewalVol:
    size: 1Gi
  backupVol:
    size: 2Gi
  dataVol:
    size: 2Gi
  logVol:
    size: 1Gi
  tablespaceVol:
    size: 512Mi
  walVol:
    size: 1Gi
  selectedVolList:
  - name: data
  - name: tablespace
  - name: wal
  - name: log
  sysExtraLogging: false
```

1.2.6 FEPCert Child Custom Resource Parameters

1.2.6.1 Create/ Update Certificates

Certificate secret for the FEP cluster is initially created in accordance with the values set in fepChildCrVal’ certs section of the parent FEPCluster CR.

Below is the schema of the FEPCert CR:

Field	Default	Description
cacrt	Defaults to dummy self signed crt from parent FEPCluster CR	Can be replaced with customer’ s own CA cert
crt	Defaults to dummy self signed crt from parent FEPCluster CR	Can be replaced with customer’ s own trusted cert
key	Defaults to dummy key from parent FEPCluster CR	Can be replaced with customer’ s own key

By default, Operator will create Kubernetes secrets to store the CA Cert, Server Cert and Key file. These files are exposed under the mount point /fep-certs in the container. The default FEPCluster template will also set the following postgres parameters in postgresql.conf.

```
ssl = on
ssl_cert_file = '/fep-certs/fep.crt'
ssl_key_file = '/fep-certs/fep.key'
ssl_ca_file = '/fep-certs/ca.crt'
```

It should also be possible to change the certificates by end user, by changing ALL key, crt and cacrt. However, user will need to restart the cluster to let change take effect.

1.2.6.2 Delete Certificates

Equivalent Kubernetes command: kubectl delete FEPCert <cr_name>

This operation will remove the secret containing the TLS Certificates and keys for the cluster.

Below is an example CR for certificates to be used by FEP server container

```
apiVersion: fep.fujitsu.io/v1
kind: FEPCert
metadata:
  name: new-fep
  namespace: ansible-operator-poc
spec:
  key: |-
    -----BEGIN RSA PRIVATE KEY-----
    MIIEowIBAAKCAQEA4AI33yvHZws+jta6qpV6wzJqF8odIfTIpCfbrVcUUtLFKJ1I
    2e4SceTKi6O3C/I1XuvWlpng5IO65+fQQLO06z1/AuQT78YUn/Wlm9x1aHVsv4AN
    B5JWWqDOjrRT3o7nRPGXfilabP0rGE2mJJcVR9nExJ3IeaktgT3sb8YlXvtchyYp
    mjdbfxabTz07ig0+6/cwKoRRxOK8Uf7f5euE0cI/490J6r5Rs4lgD8sIQNCUFlTF
    YvmAH7gcdssSFBt8NPlUATHEsoFmlW0DKCJWNhTLOht+s6L/1zwTHLjPG2pdkG6W
    dgmu5H2pDml8CDNLDv98Aj7i+I5SRKKcVPlnuQIDAQABAoIBAFPQYKlOzw/+BA0b
    yMIUpdctIMb/54CR/xR0mVw1DbSjigNVPjHUQvB8Y1B2FAITQObgJO06bAv0QdWN
    Rb0/v/yYiNJDFjaLjaIAHlO/2+oWrXbFaZqgpVDJhB+e1xaZr2x7XGxm+p925k30
    l6pvIRY+I8JRKvZiV1VZHwL/R3JOtPr++xMZtLVjVOI+f+ySqJ+TZHuAjm49EKxj
    cEmmJ28b7QcziXsvKy00f+zbqLIBKXQdZAFU5eEr1BsDRXdRW+Kf0XIvftuy4BJZ
    voKT+VGhEvF/qysswL4+6IAO6tpuYnnM0Y2d3sOGoWPkTcQK0MekYKzL/WmtCjNs
    9hodJtECgYEA5EWyhEOf4uOKe5TDp697UCUvXLoOR58FDe/S8XNvScn29jjOkqIg
    OMoqo9xAkJTNTzqn5UUdt1x/pgM2NxlPLFijrc0zQlX3SoOO2ryDd9WNi7YKtN16
    KJqa536WeZu2OEbuAZ+S3GALVy1RPeTNPnUOmKnF06DjDUGzLNCZy10CgYEA+zfw
    952DWuz1U0Z4wvAEqqcgUKXPKrkTXV/iUnjkDkrLYVr0ZofDNTXrdHl+UedFmaOC
    cieZn6DNhcdz5tKtyysGMH3g/qs9PfoGUngvcXsy0Egk04l3x1jc8TTCLqXZXYaQ
    HMsx51n+R58oncPtzYSUOr9qQ6PbC2CstTbFJA0CgYEAjGEsUliAB/jknfEzjXjG
    PdhQUxb8VyE864Az2lah9t/kJzFyIAziAeqZ5GE7t247AGFTBRTHHI8e1Qoemi3P
    Wbc9GVIbFs1lIYbcIDpUIyrKPEP8O5QEXtoNLxXTFgAjRGKiVY87spjCAJ+W2ZhO
    e/1it5GYXfgQCYQA2yuBmOUCgYANRkR2YR1axaCk+NlSu6oTdmdPu6M5x7PNQE7O
    OtMaKjua9lppvIzFGAdMDUtueoEEAE7ZR1xnwfB6PDLUpJdIYAqgr1YfPt8qkjaZ
    Tv56yZ7CwL0pbF8m6nwqRrZoDp1wwraEvvvxFKFKGY/k3kCHlpTakdjEoDjn3gDi
    RnWeVQKBgCEneMSzucei5LRppRtRaJw/Btll8qlPMlX3W7dxQ3cLwpmLOn0m51Fp
    PIZ44zYK8R6fu4+/sSrlfaIg86Ugeufp6YNxyNROKxUGza5vDIu5OftwWtBeg+UK
    Z8lLWNdX6pp7WMujmF3H1DrkBbauYMUKZ4UxUYtelgHERMePIxwb
    -----END RSA PRIVATE KEY-----
  crt: |-
    -----BEGIN CERTIFICATE-----
    MIIDUTCCAjmgAwIBAgIRAMocW3qMoHrD6qRvMPppMkMwDQYJKoZIhvcNAQELBQAw
    NzEQMA4GA1UEChwHRnVqaXRzdTEjMCEGA1UEAwwaRkVQIFJvb3QgQ0EgZm9yIEt1
    YmVybmV0ZXMwHhcNMjEwMjA2MDQzMjM2WhcNMjYwMjA1MDQzMjM2WjA/MRAwDgYD
    VQQKEwdGdWppdHN1MSswKQYDVQQDEyJGVUpJVFNVIEVudGVycHJpc2UgUG9zdGdy
    ZXMgU2VydmVyMIIBIjANBgkqhkiG9w0BAQEFAAOCAQ8AMIIBCgKCAQEA4AI33yvH
    Zws+jta6qpV6wzJqF8odIfTIpCfbrVcUUtLFKJ1I2e4SceTKi6O3C/I1XuvWlpng
    5IO65+fQQLO06z1/AuQT78YUn/Wlm9x1aHVsv4ANB5JWWqDOjrRT3o7nRPGXfila
    bP0rGE2mJJcVR9nExJ3IeaktgT3sb8YlXvtchyYpmjdbfxabTz07ig0+6/cwKoRR
    xOK8Uf7f5euE0cI/490J6r5Rs4lgD8sIQNCUFlTFYvmAH7gcdssSFBt8NPlUATHE
    soFmlW0DKCJWNhTLOht+s6L/1zwTHLjPG2pdkG6Wdgmu5H2pDml8CDNLDv98Aj7i
    +I5SRKKcVPlnuQIDAQABo1AwTjAdBgNVHSUEFjAUBggrBgEFBQcDAQYIKwYBBQUH
    AwIwDAYDVR0TAQH/BAIwADAfBgNVHSMEGDAWgBQcwrrUO0u+FhIUuVdrDRCQRsi6
    ZjANBgkqhkiG9w0BAQsFAAOCAQEAm5dxBoI9pScOCvRAchg4CprdRDSJb9K6yB3O
    nCAxnM47iHeXnY3WlnI388kHu8DU7O4ba1tJbGs3KY9KzioPk43pU12jWkO1onoF
    +mTDjx/Ef1cYWA9r5q/LtgTa6Q2sxV4O2x67QW82aAnaxO34dV5zWCPIvAoovZBV
    HRT+BgCg3r2vD1RGKK2nl1aYJtWhO1SZubam+VttdZ/vbM9oOJctxmImsEtBXjkY
    KteePdQtLL5o03JhyXWyRshCq+HMmKf2KgyY8gvydGcP4eLQdBWcW40LcnVq6UjT
    0kJycJEKngMVademq1ZWHGaiYB7hyT6GhgIcHUJ2cKrPgbEh1Q==
    -----END CERTIFICATE-----
  cacrt: |-
    -----BEGIN CERTIFICATE-----
    MIIDTzCCAjegAwIBAgIUYssQ8I74US5g+1+Z7CHuaDgkZnEwDQYJKoZIhvcNAQEL
    BQAwNzEQMA4GA1UEChwHRnVqaXRzdTEjMCEGA1UEAwwaRkVQIFJvb3QgQ0EgZm9y
    IEt1YmVybmV0ZXMwHhcNMjEwMjA2MDM1MjI4WhcNMzEwMjA0MDM1MjI4WjA3MRAw
```

```
    DgYDVQQKDAdGdWppdHN1MSMwIQYDVQQDDBpGRVAgUm9vdCBDQSBmb3IgS3ViZXJu
    ZXRlczCCASIwDQYJKoZIhvcNAQEBBQADggEPADCCAQoCggEBAMs97gUF0xkUzCgL
    7MiiDju9ySr/ziwjvcYU7jA9ML+SLmftMs3HtcYbAmSntqI+MDBSR/FAJTOoytuT
    pV+mCFcGj2YAjDpliHPeNcUpbryy4YMChF3+MovkIwGCksxo5rhiWhGmoBYpA48P
    4Xe8SPlzqMzhFvNeKzyiUhvjutS2Y1Ss38lsTaurFPx64vQ2PaC54XzdwMptXtpb
    tYmWSzCpJWwxZ6lF3vitdA2w0tnBWNyctAd0+RIM/fvArxiIqseAux9t0uogm5to
    lRIhvekuxOpXBPEqtIYQ4j9XUW2JH8vUDnzPkPvjrq+A3Ug8OyyfGVrW7+VYXozu
    c4aP7P0CAwEAAaNTMFEwHQYDVR0OBBYEFBzCutQ7S74WEhS5V2sNEJBGyLpmMB8G
    A1UdIwQYMBaAFBzCutQ7S74WEhS5V2sNEJBGyLpmMA8GA1UdEwEB/wQFMAMBAf8w
    DQYJKoZIhvcNAQELBQADggEBAMDwD85RAaWEBptFgLzKw+9xEUy1vcZaonAuA1qc
    T342XTueyAugxkC11HwdCGgGS34VyctfMGqj4AW6pA2ez4tLrbOps4DmV4sw8uBL
    8pgRDgfly3ob9FEg2wa0hmrwX9jH5Bt4vySUE2785uPAqaspT2UNtTBxS85BUi1T
    sKId2Rtil6an281Z81wyWVI6Jm2D4MG0mbsiGcTPlCtdg/UljvDYymXlAvd4vNhl
    k9hDa13TgDqJKgKdTIcmZoNQdpEVgFcO0h9AEUy5AuLqxHq60dLfZ6ESGPlMI7Lm
    i4PzYbCnBmOe+7TnHcPSyrnehs66Ik+oifRd82eYS7vKjFw=
    -----END CERTIFICATE-----
```

Note

This approach of specifying FEPCerts is getting deprecated. Should follow Secrets as referred in section to configure Certs for Server, Patroni and Users.

1.2.7 FEPBackup Child Custom Resource Parameters

Field	Default	Details
apiVersion	fep.fujitsu.io/v1	Fixed
kind	FEPBackup	Fixed
metadata.name	<clustername>	Enter the CR name.
spec.pgbackrestParams	"|"	"|" It is fixed, and the parameter set in pgbackrest.conf is described from the line below.
spec.schedule.num	Integer	Number of schedules to set The maximum number of backup schedules is 5.
spec.scheduleN.schedule	-	Write the date and time of the Nth schedule in cron format. The date and time is UTC time.
spec.scheduleN.type	full/incr	full: Perform a full backup (Back up the contents of the database cluster). incr ー Perform an incremental backup (Back up only the database cluster files that were changed to the last backup migration).
spec.preScript	" "	This parameter must specify a default value.
spec.postScript	" "	This parameter must specify a default value.

Example of FEPBackup CR created

```
apiVersion: fep.fujitsu.io/v1
kind: FEPBackup
metadata:
  name: fepcluster-backup
spec:
  schedule:
    num : 2
  schedule1:
```

```
    schedule : "0 0 1 * *"
    type : "full"
  schedule2:
    schedule : "0 0 1-6 * *"
    type : "incr"
  preScript: " "
  postScript: " "
  pgbackrestParams: |
    # define custom pgbackrest.conf parameters below to override defaults.
    [global]
    repo1-retention-full = 30
    repo1-retention-full-type = time
...
```

1.2.8 FEPRestore Custom Resource Parameters

Field	Default	Details
apiVersion	fep.fujitsu.io/v1	Fixed
kind	FEPRestore	Fixed
metadata.name	-	Enter the CR name.
spec.fepVersion		Optional To use FEPRestore image of given version. Possible values: 12, 13 & 14
spec.image	<current-released-image>	FEP restore container image to be used quay.io/fujitsu/fujitsu-enterprise-postgres-14-restore:ubi8-14-0.0 It is optional. Image is left blank by default. In such a case, it will pick up URL of image from operator container environment. If you specify the image, Operator will take that image to deploy container
spec.imagePullPolicy	IfNotPresent	
spec.mcSpec.limits	cpu: 0.2 memory: "300Mi"	
spec.mcSpec.requests	cpu: 0.1 memory: "200Mi"	
spec.fromFEPcluster	<from_clustername>	The name of the FEPcluster from which to restore
spec.toFEPcluster	<to_clustername>	Specifies the name of the FEP cluster to restore to. When restoring to an existing cluster, do not specify the line of this parameter.
spec.restoretype	latest/PITR	latest - Restore Latest State PITR - Date-Time Restore
spec.restoredate	-	If spec.restoretype is PITR, specify the day of PITR (UTC) in YYYY-MM-DD format Be sure to use single quotes. Example) '2020-11-25'

Field	Default	Details
spec.restoretime	-	If spec.restoretype is PITR, specifies the PITR time (UTC) in HH: MM: SS format Be sure to use single quotes. Example) '02:50:43'
spec.restoreTargetRepo		Optional If you are using multiple repositories, specify the repository from which to restore. If not specified, "1" is substituted.
spec.changeParams.fepChildCrVal.backup.pgbackrestParams		Optional Specify this to change the spec.fepChildCrVal.backup.pgbackrestParams setting in FEPClusterCR when restoring to a new DB cluster.
spec.changeParams.fepChildCrVal.backup.pgbackrestKeyParams		Optional Specify this to change the spec.fepChildCrVal.backup.pgbackrestKeyParams setting in FEPClusterCR when restoring to a new DB cluster.
spec.changeParams.fepChildCrVal.backup.caName		Optional Specify if you want to change the spec.fepChildCrVal.backup.caName setting of FEPClusterCR when restoring to a new DB cluster.
spec.changeParams.fepChildCrVal.backup.repoKeySecretName		Optional Specify if you want to change the spec.fepChildCrVal.backup.repoKeySecretName setting of FEPClusterCR when restoring to a new DB cluster.
spec.changeParams.fepChildCrVal.storage.backupVol		Optional Specify this to change the spec.fepChildCrVal.storage.backupVol setting in FEPClusterCR when restoring to a new DB cluster.
spec.changeParams.fepChildCrVal.storage.archivewalVol		Optional Specify this option to change the spec.fepChildCrVal.storage.archivewalVol setting for FEPClusterCR when restoring to a new DB cluster.
spec.changeParams.fepChildCrVal.storage.dataVol		Optional Specify this to change the spec.fepChildCrVal.storage.dataVol setting for FEPClusterCR when restoring to a new DB cluster.
spec.changeParams.fepChildCrVal.storage.walVol		Optional Specify this to change the spec.fepChildCrVal.storage.walVol setting for FEPClusterCR when restoring to a new DB cluster.
spec.changeParams.fepChildCrVal.storage.logVol		Optional Specify this to change the spec.fepChildCrVal.storage.logVol setting for FEPClusterCR when restoring to a new DB cluster.

Field	Default	Details
spec.changeParams.fepChildCrVal.storage.tablespaceVol		Optional Specify this to change the spec.fepChildCrVal.storage.tablespaceVol setting for FEPClusterCR when restoring to a new DB cluster.

Example of FEPRestore CR created

```
apiVersion: fep.fujitsu.io/v1
kind: FEPRestore
metadata:
  name: feprestore
spec:
  mcSpec:
    limits:
      cpu: 200m
      memory: 300Mi
    requests:
      cpu: 100m
      memory: 200Mi
  fromFEPcluster: fepcluster1
  toFEPcluster: fepcluster2
  restoretype: latest
  imagePullPolicy: IfNotPresent
```

Example of Point-In-Time-Recovery using FEPRestore CR

```
apiVersion: fep.fujitsu.io/v1
kind: FEPRestore
metadata:
  name: feprestore
spec:
  mcSpec:
    limits:
      cpu: 300m
      memory: 700Mi
    requests:
      cpu: 200m
      memory: 512Mi
  fromFEPcluster: fepclusterA
  toFEPcluster: fepclusterB
  restoretype: PITR
  restoredate: 2020-11-25
  restoretime: 02:50:43
  imagePullPolicy: IfNotPresent
```

Note

Upon successful completion, custom resources in FEPRestore are automatically deleted.

1.2.9 FEPPgpool2 Custom Resource Parameters

Equivalent Kubernetes command: kubectl create FEPPgpool2

This operation will create a PGPool2 with supplied information.

Field	Default	Details
apiVersion	fep.fujitsu.io/v1	Fixed
kind	FEPPgpool2	Fixed
metadata.name	-	List the name of the FEP Pgpool2 container.
metadata.namespace	-	Specify the namespace of the environment where you want to deploy the operator.
spec.fepVersion		Optional To use FEPPgpool2 image of given version. Possible values: 12, 13 & 14
spec.image	<current-released-image>	FEPPgpool2 container image to be used quay.io/fujitsu/fujitsu-enterprise-postgres-14-pgpool2:ubi8-14-0.0 It is optional. Image is left blank by default. In such a case, it will pick up URL of image from operator container environment. If you specify the image, Operator will take that image to deploy container.
spec.count	2	List the number of FEP Pgpool2 containers to create.
spec.serviceport	9999	Describes the TCP port for connecting to the FEP Pgpool2 container.
spec.statusport	9898	Identifies the TCP port for connecting to the PCP process.
spec.limits.cpu	400m	List the number of CPUs (restriction) to allocate to resources.limits.cpu.
spec.limits.memory	512Mi	Specifies the memory size (restriction) to allocate to resources.limits.memory.
spec.requests.cpu	200m	List the number of CPUs (request) to allocate to resources.requests.cpu.
spec.requests.memory	256Mi	Specifies the memory size (request) to allocate to resources.requests.memory
spec.fepclustername	new-fep	Enter the FEPCluster name to connect to.
spec.customhba		If you want to use pool _ hba.conf, describe what pool _ hba.conf should contain from the line below.
spec.customparams	listen_addresses = '*' pcp_listen_addresses = '*' num_init_children = 32 reserved_connections = 0 enable_pool_hba = off allow_clear_text_frontend_auth = off authentication_timeout = 80 backend_weight0 = 1 backend_weight1 = 1 backend_flag0 = 'ALWAYS_PRIMARY'	"|" and the Pgpool-II parameters. Refer to "Pgpool-II parameters" for detail.

Field	Default	Details
	backend_flag1 = 'DISALLOW_TO_FAILOVER' connection_cache = on max_pool = 4 listen_backlog_multiplier = 2 serialize_accept = off child_life_time = 300 client_idle_limit = 0 child_max_connections = 0 connection_life_time = 0 reset_query_list = 'ABORT; DISCARD ALL' client_min_messages = info log_min_messages = debug1 log_statement = on log_per_node_statement = on log_client_messages = on log_hostname = on log_connections = on log_line_prefix = '%t: pid %p: ' load_balance_mode = on ignore_leading_white_space = on white_function_list = '' black_function_list = 'currval,lastval,nextval,setval' black_query_pattern_list = '' database_redirect_preference_list = '' app_name_redirect_preference_list = '' allow_sql_comments = off disable_load_balance_on_write = 'transaction' statement_level_load_balance = on sr_check_period = 0 sr_check_user = 'postgres' delay_threshold = 0 log_standby_delay = 'none' ssl = on	

Field	Default	Details
	ssl_ciphers = 'HIGH:MEDIUM:+3DES:!aNULL' ssl_prefer_server_ciphers = off ssl_ecdh_curve = 'prime256v1' ssl_dh_params_file = '' relcache_expire = 0 relcache_size = 256 check_temp_table = catalog check_unlogged_table = on enable_shared_relcache = off relcache_query_target = primary wd_port0 = 9000 failover_on_backend_error = off	
spec.custompcp	"|"	If you use the pcp command, "|" and the contents of pcp.conf from the line below.
spec.customsslkey	"|"	If you want to do it, "|" and the Beethoven key content in the line below.
spec.customsslcert	"|"	If you want to do it, "|" and the contents of the public x 509 certificate from the line below.
spec.customsslcacert	"|"	If you want to do it, "|" and the following lines describe the contents of the CA root certificate in PEM format.
spec.customlogsize	100 Mi	Specifies the persistent volume size for log output.
spec.storageclassname		Specifies the storage class for log output.

Pgpool-II parameters

The parameters that can be specified are shown in the table below. For details on the parameters, refer to the Pgpool-II manual.

Category	Parameter name (Specified format)	Restart required after change
Connection settings	listen_addresses (string)	Y
	pcp_listen_addresses (string)	Y
	num_init_children (integer)	Y
	reserved_connections (integer)	Y
Authentication settings	enable_pool_hba (boolean)	
	allow_clear_text_frontend_auth (boolean)	
	authentication_timeout (integer)	
Backend settings	backend_weight0 (floating point)	
	backend_weight1 (floating point)	
	backend_flag0	
	backend_flag1	

Category	Parameter name (Specified format)	Restart required after change
Connection pooling	connection_cache (boolean)	Y
	max_pool (integer)	Y
	listen_backlog_multiplier (integer)	Y
	serialize_accept (boolean)	Y
	child_life_time (integer)	Y
	client_idle_limit (integer)	
	child_max_connections (integer)	Y
	connection_life_time (integer)	Y
	reset_query_list (string)	
Error reporting and log acquisition	client_min_messages (enum)	
	log_min_messages (enum)	
	log_statement (boolean)	
	log_per_node_statement (boolean)	
	log_client_messages (boolean)	
	log_hostname (boolean)	
	log_connections (boolean)	
	log_error_verbosity (enum)	
	log_line_prefix (string)	
Load sharing settings	load_balance_mode (boolean)	Y
	ignore_leading_white_space (boolean)	
	white_function_list (string)	
	black_function_list (string)	
	black_query_pattern_list (string)	
	database_redirect_preference_list (string)	
	app_name_redirect_preference_list (string)	
	allow_sql_comments (boolean)	
	disable_load_balance_on_write (string)	Y
	statement_level_load_balance (boolean)	
Health check	connect_timeout (integer)	
Streaming replication check	sr_check_period (integer)	
	sr_check_user (string)	
	sr_check_password (string)	
	sr_check_database (string)	
	delay_threshold (integer)	
	log_standby_delay (string)	
Secure Sockect Layer (SSL)	ssl (boolean)	Y
	ssl_ciphers (string)	Y
	ssl_prefer_server_ciphers (boolean)	Y
	ssl_ecdh_curve (string)	Y

Category	Parameter name (Specified format)	Restart required after change
	ssl_dh_params_file (string)	Y
Other parameters	relcache_expire (integer)	Y
	relcache_size (integer)	Y
	enable_shared_relcache (boolean)	Y
	relcache_query_target (enum)	
	check_temp_table (enum)	
	check_unlogged_table (boolean)	

1.2.10 FEPAction Custom Resource Parameters

Specify parameters in the format described below.

Custom resource spec	Default	Change effect
.spec.targetClusterName		Must specify target FEP Cluster name within namespace mentioned in metadata.
.spec.targetPgpool2Name		Must specify target FEPPgpool2 name within namespace mentioned in metadata when using pgpool2_restart.
.spec.fepAction.type		Must specify action type. Supported action types are: restart pod_restart reload list switchover failover pgpool2_restart backup open_tde_masterkey
.spec.fepAction.args		Must specify arguments needed for given action. For details of args corresponding to each action refer to "1.2.10.1 FEPAction Specific Operation Details".
.spec.fepAction.backupType	full	Options If you specify backup for fepAction.type, the type of backup is used. full : Performs a full backup (backs up the contents of the database cluster). incr : Perform an incremental backup (Back up only the database cluster files that were changed during the last backup migration).
.spec.fepAction.backupRepo	1	Options Gets a backup in the specified repository. The range is 1 to 256.
.spec.sysExtraLogging		To turn extra debugging on, set value to true. It can be turned on/off at any time.

After execution of FEPAction CR, status is reflected in fepStatus field that is dynamically inserted in current FEPAction CR as needed.

fepStatus field used for FEPAction CR are described here

fepStatus (with possible values)	Remarks
fepActionStatus:	fepStatus is inserted at the top of FEPAction CR
fepActionCondition: Success | Failure	This flag is inserted in fepAction CR to reflect success or failure of requested action
fepActionResult: > “details”	The result contains verbose details corresponding to the specific action been executed. Should be noted that it is either plain text of HTTP output.
processedTimestamp: <time stamp>	Denotes time of action execution by the Operator

```
apiVersion: fep.fujitsu.io/v1
kind: FEPAction
fepActionStatus:
 fepActionCondition: Success
metadata:
 name: new-fep-reload-action
 namespace: myns
spec:
 fepAction:
  args:
  - new-fep-sts-0
  - new-fep-sts-1
  type: reload
 sysExtraLogging: false
 targetClusterName: new-fep
```

Note

- Please do not use the FEPAction to perform a switchover or restart while executing backup. Failed to get the backup.
- You must create a new FEPAction custom resource for each operation.

1.2.10.1 FEPAction Specific Operation Details

Action type - reload

The reload action will manually reload the FEP database on the targeted FEPCluster.

“reload” action type expects users to specify the name of individual FEP pods that they want to run the database reload operation on. They specify that in the args section under the FEPAction CR spec as below :

```
spec:
fepAction:
 args:
 - nf-131851-sts-0
 - nf-131851-sts-1
 type: reload
 targetClusterName: nf-131851
```

Action type - restart

The restart action will manually restart the FEP database on the targeted FEPCluster.

“restart” action type expects users to specify the name of individual FEP pods that they want to run the database restart operation on. They specify that in the args section under the FEPAction CR spec as below:

```
spec:
fepAction:
 args:
 - nf-131851-sts-0
 - nf-131851-sts-1
 type: restart
 targetClusterName: nf-131851
```

Action type - pod_restart

The pod_restart action will restart specified list of POD for given target cluster. User can specify key word ‘ALL’ under ‘args’ section to restart all pods in target cluster. Alternatively, user can give the list of pods to be started in target cluster. User should either give ALL or the list of the pods.

This action restarts the replica pods first. Once all replicas have been restarted, it switches over the mastership to one of the replica before restarting old master pod. If it is a single node cluster, master will be restarted in its current state. This action is automatically created to restart pods when image or machine specs are changed for fep or backup container depending on autoPodRestart flag in FEPCluster CR (see more details in FEPCluster CR section):

```
spec:
fepAction:
args:
 - nf-131851-sts-0
 - nf-131851-sts-1
 type: pod_restart
 targetClusterName: nf-131851
```

Action type - list

The list action will return the status of the targeted FEPCluster.

“list” action type expects users to specify just the target cluster name to list the details of the same. Looks like below:

```
spec:
fepAction:
 type: list
 targetClusterName: nf-131851
```

Action type - switchover

The switchover action performs a manually switchover of the current leader/primary database from one pod to another pod of the targeted FEPCluster.

"switchover" action type expects users to specify the name of the target cluster that they want to perform switchover. args section is not required for switchover as FEPAction operator code will internally find it and promote new master. FEPAction CR spec as below:

```
spec:
fepAction:
 type: switchover
 targetClusterName: nf-131851
```

Action type - failover

The failover action performs a manually failover of the current primary database from one pod to another pod of the targeted FEPCluster. The difference between switchover and failover is that, switchover expects the primary database is running at the time whereas failover can force switchover of primary role from a non-responding pod to another pod. Note that failover is a disruptive action and may cause data lost.

“failover” action type expects users to specify the names of the candidate pods that they want to failover to. They specify that in the args section under the FEPAction CR spec as below:

```
spec:
fepAction:
 args:
```

```
 - nf-131851-sts-1
 - nf-131851-sts-2
 type: failover
 targetClusterName: nf-131851
```

Here, nf-131851-sts-1 and nf-131851-2 are the candidate pods to failover to. In this example, the current primary pod would be nf-131851-sts-0.

Action type - pgpool2_restart

“pgpool2_restart” action type expects users to specify the name of individual FEPPgpool2 resource that they want to restart operation on. They specify that in the targetPgpool2Name section under the FEPAction CR spec as below:

```
spec:
fepAction:
 type: pgpool2_restart
 targetPgpool2Name: nf-131851-pgpool2
```

Action type - backup

The "backup" action performs a backup on the target FEPCluster.

The "backup" action type requires you to specify the type of backup and the repository in which to store the data.

In the fepAction section of the FEPAction custom resource specification, specify the following:.

```
spec:
  targetClusterName: new-fep
  fepAction:
    type: backup
    backupType: full
    backupRepo: 1
```

- Regardless of how the backup was performed (scheduled or FEPAction), if backups were performed at the same time by the same FEPCluster, subsequent backups will fail.
- If the backup repository Retention Option is specified in the FEPCluster custom resource spec.fepChildCrVal.backup.pgbackrestParams, the backup files obtained by the FEPAction are also deleted as specified by the option.

Action type - open_tde_masterkey

The open_tde_masterkey action opens a keystore for a TDE-enabled target cluster.

The "open_tde_masterkey" action type requires the user to specify the name of the target cluster on which the keystore will be opened. The args section is not required.

Specify the following:

```
spec:
  targetClusterName: nf-131851
  fepAction:
    type: open_tde_masterkey
```

1.2.11 FEPExporter Custom Resource

Field	Default	Details
apiVersion	fep.fujitsu.io/v1	Mandatory as it is
kind	FEPExporter	Mandatory as it is

Field	Default	Details
metadata.name	fep-monitor	Name of FEPExporter CR - must be unique in namespace
metadata.namespace	fep-ns	Namespace - OCP populates it as current
spec.prometheus		Optional Prometheus MTLS spec section
spec.prometheus.tls		
spec.prometheus.tls.certificateName		Optional This points to Kubernetes TLS secret that contains the certificate of Prometheus ServiceMonitor. FEPExporter will use this for certificate authentication. The certificate itself is stored in the key tls.crt.
spec.prometheus.tls.caName		Optional This points to Kubernetes configmap that contains additional CA the client use to verify a server certificate. The CA is stored in the key ca.crt.
spec.fep.remoteLogging.enable		Set to true to forward logs from fluentbit to fluentd
spec.fep.remoteLogging.image		Optional Fluentbit image to be used. If not specified, Operator will use the latest version that is supported by the Operator.
spec.fep.remoteLogging.pullPolicy	IfNotPresent	Optional
spec.fepExporter.		Exporter spec section
spec.fepExporter.authSecret		Optional Base Authentication secret to provide username & encrypted password of user
spec.fepExporter.authSecret.secretName		Secret name
spec.fepExporter.authSecret.userKey		Key of username in specified secret
spec.fepExporter.authSecret.passwordKey		Key of password in specified secret
spec.fepExporter.customLabel		Custom label to be added to Prometheus ServiceMonitor
spec.fepExporter.tls		FEPExporter MTLS specs
spec.fepExporter.tls.certificateName		Optional This point to Kubernetes TLS secret that contains the certificate of FepExporter. Prometheus will use this for certificate authentication. The certificate itself is stored in the key tls.crt.
spec.fepExporter.tls.caName		Optional This points to Kubernetes configmap that contains additional CA the client use to verify a server certificate. The CA is stored in the key ca.crt.
spec.fepExporter.disableDefaultQueries	false	Optional Not defined or set to false => Create default queries Defined and set to true => Do not create default queries.

Field	Default	Details
spec.fepExporter.disableDefaultAlertRules	false	Optional Not defined or set to false => Create default alert rules Defined and set to true => Do not create default alert rules. If Default queries are disabled => Do not create default alert rule.
spec.fepExporter.exporterLogLevel	error	Set logging level: one of debug, info, warn, error
spec.fepExporter.fepClusterList		Array of FEPCluster to monitor
spec.fepExporter.image.image		quay.io/fujitsu/fujitsu-enterprise-postgres-exporter:ubi8-14-0.0 Optional If not specified; image name is picked up from operator environment variable
spec.fepExporter.image.pullPolicy	IfNotPresent	Always or IfNotPresent
spec.fepExporter.mcSpec.limits	cpu: 500m memory: 700Mi	Max CPU allocated to exporter container Max memory allocated to exporter container
spec.fepExporter.mcSpec.requests	cpu: 200m memory: 512Mi	CPU allocation at start for exporter container memory allocation at start for exporter container
spec.fepExporter.scrapeInterval	30s	Optional This parameter may be specified to change statistics scraping frequency. If specified, Prometheus will poll FEPExporter at given interval. CHANGE THIS PARAMETER ONLY IF REALLY REQUIRED
spec.fepExporter.scrapeTimeout	30s	Optional This parameter may be specified to change statistics scraping timeout. If specified, Prometheus will wait for FEPExporter for maximum this given period to return statistics. CHANGE THIS PARAMETER ONLY IF REALLY REQUIRED
spec.fepExporter.sysExtraLogging	true	To turn on extra debugging messages for operator, set value to true *It can be turned on/off at any time*
spec.fepExporter.restartRequired	false	True: To restart FEPExporter, when there is any change found in CR or FEPCluster False: Will not restart FEPExporter
spec.fepExporter.userCustomQueries		Optional Section Example user's custom query to extract additional metrics.

```
    usr_example:
      query: "SELECT EXTRACT(EPOCH FROM (now() - pg_last_xact_replay_timestamp())) as lag"
      master: true
      metrics:
        - lag:
```

```
            usage: "GAUGE"
            description: "Replication lag behind master in seconds"
```

1.2.12 FEPAutoscale Custom Resource

When FEPClusterCR is defined, FEPAutoscaleCR is defined.

The parameters are as follows:

Configuration changes are made in FEPClusterCR.

Field	Default	Details
apiVersion	fep.fujitsu.io/v1	Fixed
kind	FEPAutoscale	Fixed
metadata.name	Same as FEPClusterCR	Fixed
metadata.namespace	Same as FEPClusterCR	Fixed
spec.scaleout.policy	off	[cpu_utilization/connection_number/off]
spec.scaleout.threshold	cpu_utilization: 40 connection_number: 40	Threshold
spec.scaleout.metricName	pg_capacity_connection_average	Specify this parameter if policy is connection_number. The custom metrics server must publish the average number of connections in the FEP cluster under this name.
spec.scaleout.stabilizationWindowSeconds	0	If the duration (seconds) threshold of this parameter has been exceeded continuously, a scale out is performed.
spec.limits.maxReplicas	2	Maximum number of replicas (0 to 15) If the value is out of range, no automatic scale out is performed.

1.2.13 FEPUpgrade Custom Resource

If "spec.fepChildCrVal.upgrade" is defined for the FEPCluster custom resource, the FEPUpgrade custom resource is defined.

The parameters are as follows:

Field	Default	Details
apiVersion	fep.fujitsu.io/v1	Fixed
kind	FEPUpgrade	Fixed
metadata.name	Same as FEPClusterCR	Fixed
metadata.namespace	Same as FEPClusterCR	Fixed
spec.upgrade		
spec.upgrade.sourceCluster		Specifies the FEPClusterCR name from which to migrate data. Required.
spec.upgrade.mcSpec.limits	cpu: 200m memory: 300Mi	Optional Specifies the maximum number of resources to allocate to the upgrade execution container.
spec.upgrade.mcSpec.requests	cpu: 100m memory: 200Mi	Optional Specifies the lower limit of resources allocated to the upgrade execution container.

Field	Default	Details
spec.upgrade.image		Optional If omitted, the URL for image is obtained from the operator container environment.
spec.upgrade.imagePullPolicy	IfNotPresent	Optional Specifies the pull policy for the container image. - Always - IfNotPresent - Never
spec.upgrade.source.pgAdminTls.certificateName		Optional If you do not define spec.fepChildCrVal.sysUsers.pgAdminTls.certificateName for the data source, it points to the Kubernetes TLS secret that contains the certificate for the Postgres user "postgres" in the data source. If the data source FEP has set the authentication method for the upgrade execution container to "cert", then the upgrade execution container uses the certificate defined as secret.
spec.upgrade.destination.pgAdminTls.certificateName		Optional If you have not defined the spec.fepChildCrVal.sysUsers.pgAdminTls.certificateName of the newly created FEPCluster, it points to the Kubernetes TLS secret that contains the certificate of the Postgres user “postgres ” in the data source. If you create a new FEP with the "cert" authentication method for the upgrade execution container, the upgrade execution container uses the certificate defined as secret.
spec.upgrade.storage		Optional Defines the storage for storing dump files.
spec.upgrade.storage.storageClass		Optional If omitted, the default storage class for your environment is used.
spec.upgrade.storage.size	2Gi	Optional Specifies the size of the storage to store the dump file.
spec.upgrade.storage.accessModes	ReadWriteOnce	Optional accessModes for store the dump file Specified as an array of accessModes e.g. [ReadWriteMany] If omitted, it will be treated as [ReadWriteOnce]

1.2.14 FEPLogging Custom Resources

The fepLogging section needs to be added under spec to define required parameters for FEPLogging configuration.

Following is a sample template :

```
spec:
  fepLogging:
    elastic:
      authSecret:
        secretName: elastic-auth
        passwordKey: password
        userKey: username
      host: elastic-passthrough.apps.openshift.com
      logstashPrefix: postgres
      port: 443
      scheme: https
      sslVerify: true
      tls:
        certificateName: elastic-cert
        caName: elastic-cacert
    image:
      pullPolicy: IfNotPresent
    mcSpec:
      limits:
        cpu: 500m
        memory: 700Mi
      requests:
        cpu: 200m
        memory: 512Mi
    restartRequired: false
    sysExtraLogging: false
    scrapeInterval: 30s
    scrapeTimeout: 30s
    tls:
      certificateName: fluentd-cert
      caName: cacert
  prometheus:
    …
```

Below is the list of all parameters defined in the fepLogging section, along with their brief description

Custom Resource spec	Required/ Optional	Change Effect	Updating value allowed
spec.fepLogging.image.image	Optional	Fluentd Image of FEPLogging	Yes
spec.fepLogging.image.pullPolicy	Required	Fluentd Image pull policy of FEPLogging	Yes
spec.fepLogging.mcSpec.limits.cpu	Required	Max CPU allocated to fluentd container	Yes
spec.fepLogging.mcSpec.limits.memory	Required	Max memory allocated to fluentd container	Yes
spec.fepLogging.mcSpec.requests.cpu	Required	CPU allocation at start for fluentd container	Yes
spec.fepLogging.mcSpec.requests.memory	Required	Memory allocation at start for fluentd container	Yes
spec.fepLogging.sysExtraLogging	Required	To turn on extra debugging messages for operator, set value to true. It can be turned on/off at any time	Yes
spec.fepLogging.restartRequired	Required	To restart FEPLogging instance for applying any new configuration for example after certificate rotation	Yes
spec.fepLogging.scrapeInterval	Optional	Scrape interval for Prometheus to fetch metrics from FEPLogging instance	Yes

Custom Resource spec	Required/ Optional	Change Effect	Updating value allowed
spec.fepLogging.scrapeTimeout	Optional	Scrape Timeout for Prometheus to fetch metrics from FEPLogging instance	Yes
spec.fepLogging.elastic.host	Optional	Target Elasticsearch host name	Yes
spec.fepLogging.elastic.port	Optional	Target Elasticsearch port number	Yes
spec.fepLogging.elastic.authSecret.secretName	Optional	Secret name which contains Elasticsearch authentication username & password	Yes
spec.fepLogging.elastic.authSecret.userKey	Optional	Username key specified in Elasticsearch authentication secret	Yes
spec.fepLogging.elastic.authSecret.passwordKey	Optional	Password key specified in Elasticsearch authentication secret	Yes
spec.fepLogging.elastic.logstashPrefix	Optional	Logstash prefix to differentiate index pattern in elastic search. Default value is postgres	Yes
spec.fepLogging.elastic.scheme	Optional	Connection scheme between FEPLogging & Elasticsearch. Possible options http & https	Yes
spec.fepLogging.elastic.sslVerify	Optional	Set to true if you want to verify ssl certificate. If set to false then will not consider TLS crtificate	Yes
spec.fepLogging.elastic.tls.certificateName	Optional	Kubernetes secret name which holds fluentd certificate	Yes
spec.fepLogging.elastic.tls.caName	Optional	Kubernetes configmap which holds cacert of Elasticsearch to verify Elasticsearch TLS connection	Yes
spec.fepLogging.tls.certificateName	Optional	Kubernetes secret name which holds Fluentd certificate	Yes
spec.fepLogging.tls.caName	Optional	Kubernetes configmap which holds cacert of Fluentd to configure MTLS between FEPLogging & Prometheus	Yes
spec.prometheus.tls.certificateName	Optional	Kubernetes secret name which holds Prometheus certificate	Yes
spec.prometheus.tls.caName	Optional	Kubernetes configmap which holds cacert of Fluentd to configure MTLS between FEPLogging & Prometheus	Yes

1.2.15 FEP Custom Resources - spec.fep.pgBadger

Custom Resource spec	Change Effect
pgBadger.schedules.create	The 'create' schedule to create report and upload it to endpoint
pgBadger.schedules.cleanup	The 'cleanup' schedule to delete the report left in container
pgBadger.options.incremental	Default: false; When set to True: create incremental report in pgbadger
pgBadger.endpoint.authentication	a secret to contain authentication info to access endpoint support basic auth only
pgBadger.endpoint.customCertificateName	Client certitificate reference in customCertificate CR
pgBadger.endpoint.fileUploadParameter	The file upload parameter defined by the web server Default: 'file'

Custom Resource spec	Change Effect
pgBadger.endpoint.insecure	equivalent to curl -insecure option, default to false
pgBadger.endpoint.url	Web server url to upload the report file

Appendix A Default Metrics Queries

```
pg_capacity_connection:
    query: |
      select sys, idle, idleintx, idleintx10min, idleintx1hour, idleintx1day, idleintx1week,
(curr.idle + curr.idleintx + curr.active) total,  s.setting "max" from
      (
       select
       count(CASE WHEN a.state is null THEN 1 END) sys,
       count(CASE WHEN a.state='idle' THEN 1 END) idle,
       count(CASE WHEN a.state='idle in transaction' OR a.state='idle in transaction (aborted)' THEN
1 END) idleintx,
       count(CASE WHEN (a.state='idle in transaction' OR a.state='idle in transaction (aborted)') AND
age(now(), state_change) > interval '10 min'  THEN 1 END) idleintx10min,
       count(CASE WHEN (a.state='idle in transaction' OR a.state='idle in transaction (aborted)') AND
age(now(),state_change) > interval '1 hour'  THEN 1 END) idleintx1hour,
       count(CASE WHEN (a.state='idle in transaction' OR a.state='idle in transaction (aborted)') AND
age(now(),state_change) > interval '1 day'  THEN 1 END) idleintx1day,
       count(CASE WHEN (a.state='idle in transaction' OR a.state='idle in transaction (aborted)') AND
age(now(),state_change) > interval '1 week'  THEN 1 END) idleintx1week,
       count(CASE WHEN a.state='active' THEN 1 END) active
       from  pg_stat_activity a
      ) curr, pg_settings s where name = 'max_connections'
    master: true
    metrics:
      - sys:
          usage: 'GAUGE'
          description: 'Number of system connections.'
      - idle:
          usage: 'GAUGE'
          description: 'Number of idle connections.'
      - idleintx:
          usage: 'GAUGE'
          description: 'Number of idle in transaction connections.'
      - idleintx10min:
          usage: 'GAUGE'
          description: 'Number of idle in transaction connections running longer than 10 min.'
      - idleintx1hour:
          usage: 'GAUGE'
          description: 'Number of idle in transaction connections running longer than 1 hour.'
      - idleintx1day:
          usage: 'GAUGE'
          description: 'Number of idle in transaction connections running longer than 1 day.'
      - idleintx1week:
          usage: 'GAUGE'
          description: 'Number of idle in transaction connections running longer than 1 week.'
      - total:
          usage: 'GAUGE'
          description: 'Number of total connections.'
      - max:
          usage: 'GAUGE'
          description: 'Max number of connections.'

  pg_capacity_schema:
    query: |
      SELECT current_database() AS database_name, table_schema,
COALESCE(SUM(pg_total_relation_size('"'||table_schema||'"."'||table_name||'"')), 0) AS size
      FROM information_schema.tables GROUP BY table_schema
    master: true
    metrics:
      - database_name:
          usage: 'LABEL'
```

```
          description: 'Database name.'
      - table_schema:
          usage: 'LABEL'
          description: 'Table schema name.'
      - size:
          usage: 'GAUGE'
          description: 'Disk space of schema.'

  pg_capacity_tblspace:
    query: |
      SELECT pg_tablespace.spcname AS tablespace_name, pg_tablespace_size(pg_tablespace.spcname) AS
tablespace_size FROM pg_tablespace
    master: true
    metrics:
      - tablespace_name:
          usage: 'LABEL'
          description: 'Table space name.'
      - tablespace_size:
          usage: 'GAUGE'
          description: 'Disk space of table space.'

  pg_capacity_tblvacuum:
    query: |
      SELECT current_database() datname, t.table_schema,  count(t.table_name) table_count
      FROM information_schema.tables t
      INNER JOIN pg_catalog.pg_stat_user_tables tu on t.table_schema::text=tu.schemaname::text and
t.table_name::text=tu.relname::text
      and
      age(now(),greatest(COALESCE(last_vacuum, '1970-01-01Z'), COALESCE(last_autovacuum,
'1970-01-01Z'))) > interval '1 day'
      GROUP BY t.table_schema
    master: true
    metrics:
      - datname:
          usage: 'LABEL'
          description: 'Database name.'
      - table_schema:
          usage: 'LABEL'
          description: 'Table schema name.'
      - table_count:
          usage: 'GAUGE'
          description: 'Number of tables without vacuum for more than a day.'

  pg_capacity_longtx:
    query: |
      with xact_count as (
       SELECT COALESCE(datname, '') datname, count(1)
       FROM pg_stat_activity
       where backend_type='client backend' and age(now(), COALESCE(xact_start, '1970-01-01Z')) >
interval '5 minutes'
       group by datname
      )
      select d.datname, coalesce(xc.count, 0) as count from pg_database d left join xact_count xc on
d.datname=xc.datname
    master: true
    metrics:
      - datname:
          usage: 'LABEL'
          description: 'Database name.'
      - count:
          usage: 'GAUGE'
          description: 'Number of transactions running longer than 5 minutes.'
```

```
  pg_capacity_tblbloat:
    query: |
      SELECT DISTINCT
        current_database() as datname, schemaname, tablename as relname, /*reltuples::bigint,
relpages::bigint, otta,*/
        CASE WHEN relpages < otta THEN 0 ELSE bs*(sml.relpages-otta)::BIGINT END AS wastedbytes
      FROM (
        SELECT
          schemaname, tablename, cc.reltuples, cc.relpages, bs,
          CEIL((cc.reltuples*((datahdr+ma-
           (CASE WHEN datahdr%ma=0 THEN ma ELSE datahdr%ma END))+nullhdr2+4))/(bs-20::float)) AS otta,
         COALESCE(c2.relname,'?') AS iname, COALESCE(c2.reltuples,0) AS ituples, COALESCE(c2.relpages,
0) AS ipages,
          COALESCE(CEIL((c2.reltuples*(datahdr-12))/(bs-20::float)),0) AS iotta -- very rough
approximation, assumes all cols

        FROM (
          SELECT
            ma,bs,schemaname,tablename,
            (datawidth+(hdr+ma-(case when hdr%ma=0 THEN ma ELSE hdr%ma END)))::numeric AS datahdr,
           (maxfracsum*(nullhdr+ma-(case when nullhdr%ma=0 THEN ma ELSE nullhdr%ma END))) AS nullhdr2
          FROM (
            SELECT
              schemaname, tablename, hdr, ma, bs,
              SUM((1-null_frac)*avg_width) AS datawidth,
              MAX(null_frac) AS maxfracsum,
              hdr+(
                SELECT 1+count(*)/8
                FROM pg_stats s2
                WHERE null_frac<>0 AND s2.schemaname = s.schemaname AND s2.tablename = s.tablename
              ) AS nullhdr
            FROM pg_stats s, (
              SELECT
                (SELECT current_setting('block_size')::numeric) AS bs,
                CASE WHEN substring(v,12,3) IN ('8.0','8.1','8.2') THEN 27 ELSE 23 END AS hdr,
                CASE WHEN v ~ 'mingw32' THEN 8 ELSE 4 END AS ma
              FROM (SELECT version() AS v) AS foo
            ) AS constants
            GROUP BY 1,2,3,4,5
          ) AS foo
        ) AS rs
        JOIN pg_class cc ON cc.relname = rs.tablename
       JOIN pg_namespace nn ON cc.relnamespace = nn.oid AND nn.nspname = rs.schemaname AND nn.nspname
<> 'information_schema'
        LEFT JOIN pg_index i ON indrelid = cc.oid
        LEFT JOIN pg_class c2 ON c2.oid = i.indexrelid
      ) AS sml
      ORDER BY wastedbytes DESC
    master: true
    metrics:
      - datname:
          usage: 'LABEL'
          description: 'Database name.'
      - schemaname:
          usage: 'LABEL'
          description: 'Schema name.'
      - relname:
          usage: 'LABEL'
          description: 'Name of this table.'
      - wastedbytes:
          usage: 'GAUGE'
          description: 'Number of bytes wasted for table.'
```

```
pg_performance_locking_detail:
  query: |
    SELECT blocked_locks.pid AS blocked_pid,
    blocked_activity.usename AS blocked_user,
    blocking_locks.pid AS blocking_pid,
    blocking_activity.usename AS blocking_user,
    blocked_activity.query AS blocked_statement,
    1 locks
    FROM pg_catalog.pg_locks blocked_locks
    JOIN pg_catalog.pg_stat_activity blocked_activity ON blocked_activity.pid = blocked_locks.pid
    JOIN pg_catalog.pg_locks blocking_locks
    ON blocking_locks.locktype = blocked_locks.locktype
    AND blocking_locks.DATABASE IS NOT DISTINCT FROM blocked_locks.DATABASE
    AND blocking_locks.relation IS NOT DISTINCT FROM blocked_locks.relation
    AND blocking_locks.page IS NOT DISTINCT FROM blocked_locks.page
    AND blocking_locks.tuple IS NOT DISTINCT FROM blocked_locks.tuple
    AND blocking_locks.virtualxid IS NOT DISTINCT FROM blocked_locks.virtualxid
    AND blocking_locks.transactionid IS NOT DISTINCT FROM blocked_locks.transactionid
    AND blocking_locks.classid IS NOT DISTINCT FROM blocked_locks.classid
    AND blocking_locks.objid IS NOT DISTINCT FROM blocked_locks.objid
    AND blocking_locks.objsubid IS NOT DISTINCT FROM blocked_locks.objsubid
    AND blocking_locks.pid != blocked_locks.pid
   JOIN pg_catalog.pg_stat_activity blocking_activity ON blocking_activity.pid = blocking_locks.pid
    WHERE NOT blocked_locks.GRANTED
  master: true
  metrics:
    - blocked_pid:
        usage: 'LABEL'
        description: 'Blocked process id.'
    - blocked_user:
        usage: 'LABEL'
        description: 'Blocked user.'
    - blocking_pid:
        usage: 'LABEL'
        description: 'Blocking process id.'
    - blocking_user:
        usage: 'LABEL'
        description: 'Blocking user.'
    - blocked_statement:
        usage: 'LABEL'
        description: 'Blocked statement.'
    - locks:
        usage: 'GAUGE'
        description: 'Number of processes in blocked state.'

pg_performance_locking:
  query: |
    WITH
    locks as (
     SELECT blocked_locks.DATABASE, count(blocked_locks.pid) locks
     FROM pg_catalog.pg_locks blocked_locks
     JOIN pg_catalog.pg_stat_activity blocked_activity ON blocked_activity.pid = blocked_locks.pid
     JOIN pg_catalog.pg_locks blocking_locks
     ON blocking_locks.locktype = blocked_locks.locktype
     AND blocking_locks.DATABASE IS NOT DISTINCT FROM blocked_locks.DATABASE
     AND blocking_locks.relation IS NOT DISTINCT FROM blocked_locks.relation
     AND blocking_locks.page IS NOT DISTINCT FROM blocked_locks.page
     AND blocking_locks.tuple IS NOT DISTINCT FROM blocked_locks.tuple
     AND blocking_locks.virtualxid IS NOT DISTINCT FROM blocked_locks.virtualxid
     AND blocking_locks.transactionid IS NOT DISTINCT FROM blocked_locks.transactionid
     AND blocking_locks.classid IS NOT DISTINCT FROM blocked_locks.classid
     AND blocking_locks.objid IS NOT DISTINCT FROM blocked_locks.objid
     AND blocking_locks.objsubid IS NOT DISTINCT FROM blocked_locks.objsubid
```

```
       AND blocking_locks.pid != blocked_locks.pid
       JOIN pg_catalog.pg_stat_activity blocking_activity ON blocking_activity.pid =
blocking_locks.pid
       WHERE NOT blocked_locks.GRANTED group by blocked_locks.DATABASE
      ),
      dbs as (
       select * from pg_catalog.pg_database
      )
      select dbs.datname, coalesce(locks.locks, 0) locks from dbs left join locks on dbs.oid=DATABASE
    master: true
    metrics:
      - datname:
          usage: 'LABEL'
          description: 'Database name'
      - locks:
          usage: 'GAUGE'
          description: 'Number of processes in blocked state.'

  pg_replication:
    query: |
      SELECT CASE WHEN pg_last_wal_receive_lsn() = pg_last_wal_replay_lsn() THEN 0 ELSE GREATEST (0,
EXTRACT(EPOCH FROM (now() - pg_last_xact_replay_timestamp()))) END AS lag
    master: true
    metrics:
      - lag:
          usage: "GAUGE"
          description: "Replication lag behind master in seconds"

  pg_postmaster:
    query: |

      SELECT pg_postmaster_start_time as start_time_seconds from pg_postmaster_start_time()
    master: true
    metrics:
      - start_time_seconds:
          usage: "GAUGE"
          description: "Time at which postmaster started"

  pg_stat_user_tables:
    query: |
      SELECT
        current_database() datname,
        schemaname,
        relname,
        seq_scan,
        seq_tup_read,
        idx_scan,
        idx_tup_fetch,
        n_tup_ins,
        n_tup_upd,
        n_tup_del,
        n_tup_hot_upd,
        n_live_tup,
        n_dead_tup,
        n_mod_since_analyze,
        last_vacuum,
        last_autovacuum,
        last_analyze,
        last_autoanalyze,
        vacuum_count,
        autovacuum_count,
        analyze_count,
        autoanalyze_count
```

```
      FROM
        pg_stat_user_tables
    master: true
    metrics:
      - datname:
          usage: "LABEL"
          description: "Name of current database"
      - schemaname:
          usage: "LABEL"
          description: "Name of the schema that this table is in"
      - relname:
          usage: "LABEL"
          description: "Name of this table"
      - seq_scan:
          usage: "COUNTER"
          description: "Number of sequential scans initiated on this table"
      - seq_tup_read:
          usage: "COUNTER"
          description: "Number of live rows fetched by sequential scans"
      - idx_scan:
          usage: "COUNTER"
          description: "Number of index scans initiated on this table"
      - idx_tup_fetch:
          usage: "COUNTER"
          description: "Number of live rows fetched by index scans"
      - n_tup_ins:
          usage: "COUNTER"
          description: "Number of rows inserted"
      - n_tup_upd:
          usage: "COUNTER"
          description: "Number of rows updated"
      - n_tup_del:
          usage: "COUNTER"
          description: "Number of rows deleted"
      - n_tup_hot_upd:
          usage: "COUNTER"
          description: "Number of rows HOT updated (i.e., with no separate index update required)"
      - n_live_tup:
          usage: "GAUGE"
          description: "Estimated number of live rows"
      - n_dead_tup:
          usage: "GAUGE"
          description: "Estimated number of dead rows"
      - n_mod_since_analyze:
          usage: "GAUGE"
          description: "Estimated number of rows changed since last analyze"
      - last_vacuum:
          usage: "GAUGE"
         description: "Last time at which this table was manually vacuumed (not counting VACUUM FULL)"
      - last_autovacuum:
          usage: "GAUGE"
          description: "Last time at which this table was vacuumed by the autovacuum daemon"
      - last_analyze:
          usage: "GAUGE"
          description: "Last time at which this table was manually analyzed"
      - last_autoanalyze:
          usage: "GAUGE"
          description: "Last time at which this table was analyzed by the autovacuum daemon"
      - vacuum_count:
          usage: "COUNTER"
          description: "Number of times this table has been manually vacuumed (not counting VACUUM
FULL)"
      - autovacuum_count:
```

```
          usage: "COUNTER"
          description: "Number of times this table has been vacuumed by the autovacuum daemon"
      - analyze_count:
          usage: "COUNTER"
          description: "Number of times this table has been manually analyzed"
      - autoanalyze_count:
          usage: "COUNTER"
          description: "Number of times this table has been analyzed by the autovacuum daemon"

  pg_statio_user_tables:
    query: |
      SELECT current_database() datname, schemaname, relname, heap_blks_read, heap_blks_hit,
idx_blks_read, idx_blks_hit, toast_blks_read, toast_blks_hit, tidx_blks_read, tidx_blks_hit FROM
pg_statio_user_tables
    metrics:
      - datname:
          usage: "LABEL"
          description: "Name of current database"
      - schemaname:
          usage: "LABEL"
          description: "Name of the schema that this table is in"
      - relname:
          usage: "LABEL"
          description: "Name of this table"
      - heap_blks_read:
          usage: "COUNTER"
          description: "Number of disk blocks read from this table"
      - heap_blks_hit:
          usage: "COUNTER"
          description: "Number of buffer hits in this table"
      - idx_blks_read:
          usage: "COUNTER"
          description: "Number of disk blocks read from all indexes on this table"
      - idx_blks_hit:
          usage: "COUNTER"
          description: "Number of buffer hits in all indexes on this table"
      - toast_blks_read:
          usage: "COUNTER"
          description: "Number of disk blocks read from this table's TOAST table (if any)"
      - toast_blks_hit:
          usage: "COUNTER"
          description: "Number of buffer hits in this table's TOAST table (if any)"
      - tidx_blks_read:
          usage: "COUNTER"
          description: "Number of disk blocks read from this table's TOAST table indexes (if any)"
      - tidx_blks_hit:
          usage: "COUNTER"
          description: "Number of buffer hits in this table's TOAST table indexes (if any)"

  pg_database:
    query: |

      SELECT pg_database.datname, pg_database_size(pg_database.datname) as size_bytes FROM pg_database
    master: true
    cache_seconds: 30
    metrics:
      - datname:
          usage: "LABEL"
          description: "Name of the database"
      - size_bytes:
          usage: "GAUGE"
          description: "Disk space used by the database"
```

```
  pg_stat_statements:
    query: |
      SELECT t2.rolname, t3.datname, queryid, calls, total_plan_time / 1000 as
total_plan_time_seconds, total_exec_time / 1000 as total_exec_time_seconds, min_plan_time / 1000 as
min_plan_time_seconds, min_exec_time / 1000 as min_exec_time_seconds, max_plan_time / 1000 as
max_plan_time_seconds, max_exec_time / 1000 as max_exec_time_seconds, mean_plan_time / 1000 as
mean_plan_time_seconds, mean_exec_time / 1000 as mean_exec_time_seconds, stddev_plan_time / 1000 as
stddev_plan_time_seconds, stddev_exec_time / 1000 as stddev_exec_time_seconds, rows, shared_blks_hit,
shared_blks_read, shared_blks_dirtied, shared_blks_written, local_blks_hit, local_blks_read,
local_blks_dirtied, local_blks_written, temp_blks_read, temp_blks_written, blk_read_time / 1000 as
blk_read_time_seconds, blk_write_time / 1000 as blk_write_time_seconds FROM pg_stat_statements t1
JOIN pg_roles t2 ON (t1.userid=t2.oid) JOIN pg_database t3 ON (t1.dbid=t3.oid) WHERE t2.rolname !=
'rdsadmin'
    master: true
    metrics:
      - rolname:
          usage: "LABEL"
          description: "Name of user"
      - datname:
          usage: "LABEL"
          description: "Name of database"
      - queryid:
          usage: "LABEL"
          description: "Query ID"
      - calls:
          usage: "COUNTER"
          description: "Number of times executed"
      - total_plan_time_seconds:
          usage: "COUNTER"
          description: "Total plan time spent in the statement, in milliseconds"
      - total_exec_time_seconds:
          usage: "COUNTER"
          description: "Total exec time spent in the statement, in milliseconds"
      - min_plan_time_seconds:
          usage: "GAUGE"
          description: "Minimum plan time spent in the statement, in milliseconds"
      - min_exec_time_seconds:
          usage: "GAUGE"
          description: "Minimum exec time spent in the statement, in milliseconds"
      - max_plan_time_seconds:
          usage: "GAUGE"
          description: "Maximum plan time spent in the statement, in milliseconds"
      - max_exec_time_seconds:
          usage: "GAUGE"
          description: "Maximum exec time spent in the statement, in milliseconds"
      - mean_plan_time_seconds:
          usage: "GAUGE"
          description: "Mean plan time spent in the statement, in milliseconds"
      - mean_exec_time_seconds:
          usage: "GAUGE"
          description: "Mean exec time spent in the statement, in milliseconds"
      - stddev_plan_time_seconds:
          usage: "GAUGE"
          description: "Population standard deviation of plan time spent in the statement, in
milliseconds"
      - stddev_exec_time_seconds:
          usage: "GAUGE"
          description: "Population standard deviation of exec time spent in the statement, in
milliseconds"
      - rows:
          usage: "COUNTER"
          description: "Total number of rows retrieved or affected by the statement"
      - shared_blks_hit:
```

```
          usage: "COUNTER"
          description: "Total number of shared block cache hits by the statement"
      - shared_blks_read:
          usage: "COUNTER"
          description: "Total number of shared blocks read by the statement"
      - shared_blks_dirtied:
          usage: "COUNTER"
          description: "Total number of shared blocks dirtied by the statement"
      - shared_blks_written:
          usage: "COUNTER"
          description: "Total number of shared blocks written by the statement"
      - local_blks_hit:
          usage: "COUNTER"
          description: "Total number of local block cache hits by the statement"
      - local_blks_read:
          usage: "COUNTER"
          description: "Total number of local blocks read by the statement"
      - local_blks_dirtied:
          usage: "COUNTER"
          description: "Total number of local blocks dirtied by the statement"
      - local_blks_written:
          usage: "COUNTER"
          description: "Total number of local blocks written by the statement"
      - temp_blks_read:
          usage: "COUNTER"
          description: "Total number of temp blocks read by the statement"
      - temp_blks_written:
          usage: "COUNTER"
          description: "Total number of temp blocks written by the statement"
      - blk_read_time_seconds:
          usage: "COUNTER"
          description: "Total time the statement spent reading blocks, in milliseconds (if
track_io_timing is enabled, otherwise zero)"
      - blk_write_time_seconds:
          usage: "COUNTER"
          description: "Total time the statement spent writing blocks, in milliseconds (if
track_io_timing is enabled, otherwise zero)"
```

Appendix B Default Alert Rules

```
apiVersion: monitoring.coreos.com/v1
kind: PrometheusRule
metadata:
  name: {{ ansible_operator_meta.name }}-{{ item.name }}-alertrules
  namespace: {{ ansible_operator_meta.namespace }}
  labels:
    app: prometheus-postgres-exporter-alertrules
    name: {{ ansible_operator_meta.name }}-{{ item.name }}-alertrules
spec:
  groups:
    - name: fep-container
      rules:
        - alert: ContainerDisappeared
          annotations:
            description: {{ 'Container {{$labels.container}}/{{$labels.pod}} from
{{$labels.namespace}} has been disappeared' }}
            summary: Container Pod disappeared.
          expr: time() -
            container_last_seen{ container="fep-patroni",
namespace="{{ ansible_operator_meta.namespace }}", pod=~"^{{ item.name }}-sts-.*" } >  60
          labels:
            severity: warning
        - alert: ContainerHighCPUUsage
          annotations:
            description: {{ 'Container {{$labels.container}}/{{$labels.pod}} from
{{$labels.namespace}} has been high on CPU usage(>80%) for 5 mins' }}
            summary: High Container CPU usage.
          expr:
(sum(node_namespace_pod_container:container_cpu_usage_seconds_total:sum_rate{pod=~"{{ item.name }}-
sts.*", namespace="{{ ansible_operator_meta.namespace }}", container="fep-patroni"}) by
(pod,namespace,container)/sum(kube_pod_container_resource_limits_cpu_cores) by
(pod,namespace,container))*100 > 80
          for: 5m
          labels:
            severity: warning
        - alert: ContainerHighRAMUsage
          annotations:
            description: {{ 'Container {{$labels.container}}/{{$labels.pod}} from
{{$labels.namespace}} has been high on RAM usage(>80%) since 30 mins' }}
            summary: High container memory usage.
          expr: sum(container_memory_working_set_bytes{pod=~"{{ item.name }}-sts.*",
namespace="{{ ansible_operator_meta.namespace }}", container="fep-patroni"}  /
container_spec_memory_limit_bytes * 100) by (pod, container, instance) > 80
          for: 30m
          labels:
            severity: warning
        - alert: PVCLowDiskSpace
          annotations:
            description: {{ 'Found low disk space on {{$labels.persistentvolumeclaim}} in
{{$labels.namespace}} namespace.' }}
            summary: {{ 'Found low disk space on {{$labels.persistentvolumeclaim}} in
{{$labels.namespace}} namespace.' }}
          expr:
kubelet_volume_stats_available_bytes{namespace="{{ ansible_operator_meta.namespace }}",
persistentvolumeclaim=~"fep.*{{ item.name }}.*"}/ (kubelet_volume_stats_capacity_bytes) * 100 < 10
          for: 5m
          labels:
            severity: warning
    - name: postgres
      rules:
```

```
        - alert: PostgresqlDown
          annotations:
            description:  "Postgresql one or more instances are down in FEPCluster {{ item.name }} in
{{ ansible_operator_meta.namespace }} namespace. Please check the FEP pods in this cluster"
            summary: "Postgresql FEPCluster {{ item.name }} in {{ ansible_operator_meta.namespace }}
namespace is degraded"
          expr: count(pg_static{ namespace="{{ ansible_operator_meta.namespace }}",
service="{{ ansible_operator_meta.name }}-service", server=~"{{item.name}}-sts.*" }) <
{{item.instances | length}}
          labels:
            severity: error
        - alert: PostgresqlTooManyConnections
          annotations:
            description: {{ 'PostgreSQL instance has too many connections on server
{{ $labels.server }} in {{ $labels.namespace }} namespace.' }}
            summary: {{ 'Postgresql too many connections (FEPCluster server {{ $labels.server }})' }}
          expr: pg_capacity_connection_total{namespace="{{ ansible_operator_meta.namespace }}",
service="{{ ansible_operator_meta.name }}-service", server=~"{{ item.name }}-sts.*"}/
pg_settings_max_connections > 0.9
          labels:
            severity: warning
```