



SRE 4.1

Container Deployment Guide - Draft

Table of Contents

1	Hardware Requirements	1
1.1	CPU and Memory	1
1.2	Networking	2
1.2.1	Communication Matrix	2
1.3	Disk partitioning	3
2	Software Requirements	4
2.1	Database	4
2.1.1	Postgres	4
2.1.2	Influxdb	10
2.1.3	Mongodb (optional, required for CAC)	11
2.2	Container runtime	13
3	SRE Deployment	13
3.1	Importing SRE images	13
3.2	Extracting default configuration files	14
3.3	SRE Configuration	14
3.4	Deploying SRE container	14
3.4.1	create directory structure for SRE data directory	14
3.4.2	Start the container	14
3.4.3	Supported configuration options	14
3.4.4	Sample deploy of 2 EMs and 2 CPs	16
3.4.5	All-in-one deployment	16

1 Hardware Requirements

1.1 CPU and Memory

The minimum requirements for an SRE server are:

- vCPU: 4
- Memory Size: 8192 MB
- Hard Disk Size: 250GB (SSD)

Note

A more tailored dimensioning is needed on a project basis.

1.2 Networking

SRE should be equipped with a minimum of 1 Gigabit Ethernet link.

Element Managers can be configured with 1 or 2 NIC interfaces

- 1 for Management (GUI/SSH access)
- 1 for DB replication between all components (EM's and CP's) and for CP modules control

Note

They can be combined in a single interface

Call Processors can be configured with multiple NIC Interfaces

- 1 for Management (SSH access)
- 1 for DB Replication and access to the element Managers
- 1 or more for the Call Processing in case SRE needs to communicate with multiple voice networks.

Note

They can be combined in a single interface

1.2.1 Communication Matrix

Source	Destination	Interface	Protocol/DestinationPort	Description
terminal	EM	management	TCP/22	ssh/sftp
EM	CP	internal	TCP/22	ssh
browser	EM	management	TCP/8080 (*)	http/https GUI access
external provisioners	EM	management	TCP/5000 (*) (**)	REST APIs
EM	DNS server	management	UDP/53 (**)	dns resolution

Source	Destination	Interface	Protocol/DestinationPort	Description
EM/CP	NTP server	management	UDP/123	time synchronization
EM	EM	internal	TCP/5432	DB traffic
CP	EM	internal	TCP/5432	DB traffic
CP	CP	internal	TCP/5555 (*) (**)	kamailio-broker traffic for hitless update
CP	EM	internal	TCP/5000 (*)	db updates from service logic
CP	EM	internal	TCP/10000	SRE log and stats
CP	EM	internal	TCP/10001	SRE internal requests
CP	EM	internal	TCP/10002	Accounting data
EM	EM	internal	TCP/10003	Accounting synchronization
EM	SNMP managers	management	UDP/162 (*) (**)	SNMP traps
EM	SMTP server	management	TCP/587 (**)	mail server
EM	syslog server	management	UDP/514 (*) (**)	syslog data
EM	LDAP server	management	TCP/389/636 (*) (**)	GUI authentication
SIP endpoint(s)	CP	SIP	UDP/TCP/5060 (*) (**)	SIP traffic interface
HTTP endpoints	CP	HTTP	TCP/6000 (*) (**)	http traffic interface
ENUM endpoints	CP	ENUM	UDP/TCP/53 (*) (**)	ENUM traffic interface

(*) port can be customized (**) optional

1.3 Disk partitioning

Create the necessary partitions according to the table below.

The example below is sized for a total disk space of 250GB.

Partition	Size	Type	Description
/var/lib/pgsql	70 GB	Ext4 or XFS on LVM	PostgreSQL database
/data/sre/db/backups	50 GB	Ext4 or XFS on LVM	workspace for backups
/data/sre/db/wals	20 GB	Ext4 or XFS on LVM	work-ahead logs
/data/sre/provisioning	10 GB	Ext4 or XFS on LVM	provisioning data (EM only)
/data/sre/mongo	10GB	Ext4 or XFS on LVM	Mongo database (only if CAC or global caching is configured)

Note

Partition dimensions should be tailored to the specific requirements of each project.

2 Software Requirements

2.1 Database

2.1.1 Postgres

Postgres 14 and repmgr package are required and must be installed on all nodes.

{#discover-data-dir}

Location of the data directory can be retrieved with this command:

```
1 sudo -u postgres psql -c "show data_directory;"
2     data_directory
3 -----
4 /var/lib/postgresql/14/main
```

2.1.1.1 Primary Element Manager Node Configuration

The procedure in this section should be executed on the Primary Element Manager only.

Start postgresql with:

```
1 sudo systemctl start postgresql-14 (or postgresql depending on your
   ↪ distribution)
2 sudo systemctl enable postgresql-14
```

2.1.1.1.1 Configure the database access

Locate *pg_hba.conf* location depending on your distribution and edit it:

```
1 sudo vi pg_hba.conf
```

This file manages the access rights to the database. It lists all possible origins from where connections can be made to the database and which authentication method should be used. We will always apply 'trust' as an authentication method which means that we unconditionally trust the specified sources and allow them to connect to the database without any specific verification or authentication. The first line which must be present allows the localhost to initiate connections over TCP/IP connections to the database. In the default configuration file this line is configured with authentication method = 'ident'. You must assure it is configured as 'trust' for local host connections to all databases:

IPv4 local connections:

```
1 host all all 127.0.0.1/32 trust
```

Also, all other possible 'trusted' sources for database queries must be granted access to the databases previously created. Later in this document we will spin up a standby element manager and call processing nodes. All these nodes must be granted access to the sre and custom databases. For the sake of simplicity access can be granted to an entire range of IP addresses but for security reasons it is sometimes recommended that single hosts should be accepted by using a subnetmask of /32. Nevertheless, we should only accept database connections when using the database user 'sre'.

IPv4 remote connections:

```
1 host sre sre <local_subnet>/<subnetmask> trust
2 host postgres postgres <local_subnet>/<subnetmask> trust
```

2.1.1.1.2 Postgresql.conf tuning

Locate *postgresql.conf* location depending on your distribution and edit it:

```
1 sudo vi postgresql.conf
```

uncomment and set the parameter *listen_addresses* to "*". Increase the maximum number of connections to 1000:

```
1 listen_addresses = '*'
2 max_connections = 1000
```

Make sure that the destination directory is writable by the user postgres, on all nodes. Unless different directory is used, the configuration parameters to use for WAL archiving in */data/sre/db/wals/* are:

```
1 archive_mode = on
2 archive_command = 'test ! -f /data/sre/db/wals/%f && cp %p /data/sre/db/wals/%f
   ↪ '
```

Uncomment and set the following parameters (in this example the system keeps 200 WALs files of 16 MB each):

```
1 wal_level = replica
2 max_wal_senders = 10
3 wal_keep_size = 3200
```

Enable hot_standby:

```
1 hot_standby = on
```

For automatic EM switchover add the following configuration settings:

```
1 wal_log_hints = on
2 shared_preload_libraries = 'repmgr'
```

Your configuration file should look like the following (in red the deviations from default values):

```
1 # -----
2 # PostgreSQL configuration file
3 # -----
4
5 .....
6
7 #-----
8 # CONNECTIONS AND AUTHENTICATION
9 #-----
10
11 # - Connection Settings -
12
13 listen_addresses = '*'          # what IP address(es) to listen on;
14                                # comma-separated list of addresses;
15                                # defaults to 'localhost'; use '*' for
   ↪ all
16                                # (change requires restart)
17 #port = 5432                    # (change requires restart)
18 max_connections = 1000         # (change requires restart)
19 #superuser_reserved_connections = 3 # (change requires restart)
20 #unix_socket_directories = '/var/run/postgresql, /tmp' # comma-separated list
   ↪ of directories
21                                # (change requires restart)
22
23 .....
```

```

24
25
26 #-----
27 # WRITE AHEAD LOG
28 #-----
29
30 # - Settings -
31
32 wal_level = replica          # minimal, archive, hot_standby, or logical
33                               # (change requires restart)
34
35 .....
36
37 # - Archiving -
38
39 archive_mode = on           # enables archiving; off, on, or always
40                               # (change requires restart)
41 archive_command = 'test ! -f /data/sre/db/wals/%f && cp %p /data/sre/db/wals/%f
42     ↪ '                       # command to use to archive a logfile segment
43                               # placeholders: %p = path of file to archive
44                               #               %f = file name only
45                               # e.g. 'test ! -f /mnt/server/archivedir/%f &&
46     ↪ cp %p /mnt/server/archivedir/%f'
47 #archive_timeout = 0        # force a logfile segment switch after this
48                               # number of seconds; 0 disables
49 #-----
50 # REPLICATION
51 #-----
52
53 # - Sending Server(s) -
54
55 # Set these on the master and on any standby that will send replication data.
56
57 max_wal_senders = 10        # max number of walsender processes
58                               # (change requires restart)
59 wal_keep_size = 3200        # in logfile segments, 16MB each; 0 disables
60 #wal_sender_timeout = 60s   # in milliseconds; 0 disables
61 .....
62 # - Standby Servers -
63
64 # These settings are ignored on a master server.
65
66 hot_standby = on            # "on" allows queries during recovery
67                               # (change requires restart)

```



```
68 #max_standby_archive_delay = 30s           # max delay before canceling queries
69                                           # when reading WAL from archive;
70                                           # -1 allows indefinite delay
71 #max_standby_streaming_delay = 30s        # max delay before canceling queries
72                                           # when reading streaming WAL;
73                                           # -1 allows indefinite delay
74 #wal_receiver_status_interval = 10s       # send replies at least this often
75                                           # 0 disables
76 #hot_standby_feedback = off              # send info from standby to prevent
77                                           # query conflicts
78 #wal_receiver_timeout = 60s              # time that receiver waits for
79                                           # communication from master
80                                           # in milliseconds; 0 disables
81 #wal_retrieve_retry_interval = 5s        # time to wait before retrying to
82                                           # retrieve WAL after a failed attempt
83 wal_log_hints = on
84 shared_preload_libraries = 'repmgr'
```

Restart the database after changing the file.

2.1.1.1.3 Postgresql replication

Create a user for repmgr to setup replication:

```
1 su - postgres
2 createuser -s repmgr
3 createdb repmgr -O repmgr
4 exit
```

Edit the configuration file *pg_hba.conf* to change access rights for the replication:

```
1 ...
2 local replication repmgr trust
3 host replication repmgr 127.0.0.1/32 trust
4 host replication repmgr 10.0.11.30/32 trust
5 local repmgr repmgr trust
6 host repmgr repmgr 127.0.0.1/32 trust
7 host repmgr repmgr 10.0.11.30/32 trust
```

Restart the DB.

Edit the configuration file Postgresql cluster and set the parameters:

```
1 sudo vi repmgr.conf
```

```
1 node_id = <ID that should be unique in the SRE environment and greater than 0>
2 node_name=<name that should be unique in the SRE environment>
```

```

3 conninfo='host=<IP address of the server, this IP address should be accessible
    ↪ to all other nodes> dbname=repmgr user=repmgr'
4 data_directory='/var/lib/pgsql/14/data/'
5 ssh_options='-q -o ConnectTimeout=10'
6 failover='automatic'
7 reconnect_attempts=2
8 reconnect_interval=2
9 promote_command='/usr/pgsql-14/bin/repmgr standby promote -f /etc/repmgr/14/
    ↪ repmgr.conf --log-to-file; sudo docker restart sre'
10 follow_command='/usr/pgsql-14/bin/repmgr standby follow -f /etc/repmgr/14/
    ↪ repmgr.conf --log-to-file --upstream-node-id=%n'
11 repmgrd_pid_file='/run/repmgr/repmgrd-14.pid'
12 always_promote=true
13 service_start_command = 'sudo systemctl start postgresql-14'
14 service_stop_command = 'sudo systemctl stop postgresql-14'
15 service_restart_command = 'sudo systemctl restart postgresql-14'
  
```

2.1.1.2 Secondary Element Manager and Call Processor Node Configuration

Cloning the databases In order to clone the database from the master database, ensure that all the tablespaces directories are exactly the same as the Primary Element Manager node and that the access rights are identical. Check that the PostgreSQL is not running:

```

1 sudo systemctl stop postgresql-14 (or postgresql depending on your distribution
    ↪ )
  
```

Ensure that the data directory is empty. See [here](#) for the procedure to discover this path.

```

1 sudo rm -rf <data_dir>/*
  
```

Edit the configuration file *repmgr.conf* set the parameters:

```

1 node_id = <ID that should be unique in the SRE environment and greater than 0>
2 node_name=<name that should be unique in the SRE environment>
3 conninfo='host=<IP address of the server, this IP address should be accessible
    ↪ to all other nodes> dbname=repmgr user=repmgr'
4 data_directory=<postgres data dir>
5 ssh_options='-q -o ConnectTimeout=10'
6 failover='automatic'
7 reconnect_attempts=2
8 reconnect_interval=2
9 promote_command='/<path to repmgr bin>/repmgr standby promote -f /<path to
    ↪ repmgr conf>/repmgr.conf --log-to-file; sudo docker restart sre'
10 follow_command='/<path to repmgr bin>/repmgr standby follow -f /<path to repmgr
    ↪ conf>/repmgr.conf --log-to-file --upstream-node-id=%n'
11 repmgrd_pid_file='<full path to repmgrd pid file>'
  
```

```
12 always_promote=true
13 service_start_command = 'sudo systemctl start postgresql-14' # or postgresql
14 service_stop_command = 'sudo systemctl stop postgresql-14' # or postgresql
15 service_restart_command = 'sudo systemctl restart postgresql-14' # or
    ↪ postgresql
```

Connect as user postgres and clone the PostgreSQL data directory files from the master. Be sure to copy also *postgresql.conf* and *pg_hba.conf* from master.

```
1 su - postgres
2 /<path to repmgr bin>/repmgr -h <IP Mater EM> -U repmgr -d repmgr -f /<path to
    ↪ repmgr conf>/repmgr.conf standby clone
3 exit
```

As root, start PostgreSQL and enable it at boot:

```
1 sudo systemctl start postgresql-14 (or postgresql)
2 sudo systemctl enable postgresql-14 (or postgresql)
```

As postgres user, register the standby server:

```
1 su - postgres
2 /<path to repmgr bin>/repmgr -f /<path to repmgr conf>/repmgr.conf standby
    ↪ register
3 exit
```

2.1.1.3 Repmgrd

On all nodes start and enable repmgrd daemon:

```
1 sudo systemctl start repmgrd
2 sudo systemctl enable repmgrd
```

2.1.2 Influxdb

Support is provided for InfluxDB OSS versions 2.4 through 2.7. Please refer to [official site](#) for installation instructions.

Note

Influxdb is needed only on EM nodes.

2.1.2.1 Configuration

As root, start Influxdb and enable it at boot:

```
1 [root@em ~]# systemctl start influxd
2 [root@em ~]# systemctl enable influxd
```

Configure it with:

```
1 [root@em ~]# influx setup -u influxuser -p influxuser -t my-super-secret-token
   ↪ -o influxorg -b bucket -r 1h (it will ask for confirmation)
2 [root@em ~]# influx bucket delete -n bucket
```

2.1.3 MongoDB (optional, required for CAC)

MongoDB 5.0.x is required on all nodes. Please refer to [official site](#) for installation instructions.

2.1.3.1 Configuration

Note

The number of nodes for MongoDB must be an odd number (N+1) to minimize the possibility for an outage of CAC. Indeed, if less than N/2 nodes (example 1 node out of 3) are unavailable, there is still a quorum (majority of nodes) therefore a PRIMARY node is available and the CAC feature is generally available.

```
1 sudo cat /etc/mongod.conf
2 # Where and how to store data.
3 storage:
4     dbPath: /data/sre/mongodb
5     journal:
6         enabled: true
7
8 # network interfaces
9 net:
10     port: 27017
11     bindIp: 0.0.0.0
12
13 #security:
14
15 #operationProfiling:
16
17 replication:
18     replSetName: sre_location
```

Start and enable mongodb at boot:

```
1 sudo chown -R mongod.mongod /data/sre/mongodb
2 sudo systemctl restart mongod
3 sudo systemctl enable mongod
```

2.1.3.1.1 Add arbiter (optional)

Note

It is recommended to configure one of EM nodes as mongodb arbiter if total number of SRE CP hosts is even.

```
1 sudo cat /etc/mongod.conf
2 # Where and how to store data.
3 storage:
4     dbPath: /data/sre/mongodb/arb
5     journal:
6         enabled: true
7
8 # network interfaces
9 net:
10     port: 27017
11     bindIp: 0.0.0.0
12
13 #security:
14
15 #operationProfiling:
16
17 replication:
18     replSetName: sre_location
```

Start and enable mongodb at boot:

```
1 sudo mkdir /data/sre/mongodb/arb
2 sudo chown -R mongod.mongod /data/sre/mongodb
3 sudo systemctl restart mongod
4 sudo systemctl enable mongod
```

2.1.3.1.2 Cluster initialization

Connect to one SRE host and start the mongodb shell with:

```
1 mongo (or mongosh)
```

Issue the following command:

```
1 > rs.initiate({_id : "sre_location", members: [{_id: 0, host: "<address1>" }]}))
```

Add additional nodes to the cluster with (every node must have a unique id):

```
1 > rs.add({_id: 1, host: "<address2>" })
2 ...
```

If an arbiter has been configured, add it with:

```
1 > rs.addArb("<address3>")
```

If all went well, issue the following command to check where the PRIMARY node is located:

```
1 > rs.status()
```

2.1.3.1.3 Setting write concern

Connect to primary node and on that node launch mongo cli and insert the following command:

```
1 > db.adminCommand({"setDefaultRWConcern" : 1,"defaultWriteConcern" : {"w" :
↪ 1}})
```

2.2 Container runtime

Docker 20.10.7 or higher is required on all nodes. Please refer to [official site](#) for installation instructions.

Note

Using Docker installations from alternative sources, like Snap, is discouraged.

3 SRE Deployment

3.1 Importing SRE images

On EMs and CPs:

```
1 sudo docker import sre-oci.tar.gz netaxis/sre
```

3.2 Extracting default configuration files

```

1 sudo mkdir -p /opt/sre/
2 sudo mkdir -p /var/log/sre/
3 sudo mkdir -p /data/sre/
  
```

3.3 SRE Configuration

3.4 Deploying SRE container

3.4.1 create directory structure for SRE data directory

On all nodes run:

```

1 sudo mkdir -p /data/sre/accounting/state
2 sudo mkdir -p /data/sre/db/backups
3 sudo mkdir -p /data/sre/db/wals
4 sudo mkdir -p /data/sre/provisioning
  
```

3.4.2 Start the container

3.4.3 Supported configuration options

Option name	Description	Default Value
DB_USER	user for connecting to postgres database	sre
DB_PASSWORD	password for connecting to postgres database	sre
DB_HOST	postgres host to connect (hostname or address)	initialize and use embedded database
DB_NAME	name of the database to create	sre
DB_PORT	postgres port to connect	5432
REPMGR_NODE_ID	repmgr node id to use	1
INFLUXDB_HOSTS	list of comma separated influxdb hostnames or addresses	initialize and use embedded database
INFLUXDB_TOKENS	list of comma separated influxdb tokens for connection	generate token for embedded database

Option name	Description	Default Value
INFLUXDB_ORG	influxdb organization to use for connection	influxorg
INFLUXDB_USER	influxdb user to use for connection	influxuser
INFLUXDB_PASSWORD	influxdb password to use for connection	influxpassword
ENABLE_MONGODB	enable(1) or disable(0) embedded mongodb	1 (enabled)
MONGODB_HOSTS	list of comma separated mongodb hostnames or addresses	use embedded mongodb
ENABLE_MANAGER	make this instance a manager(1) or a client(0)	1 (enabled)
ENABLE_KAMAILIO	enable(1) or disable(0) embedded kamailio	1 (enabled)
KAMAILIO_PORT	use this port for sip traffic	5060
KAMAILIO_EXTRA_DEFINES	comma separated defines for kamailio	no extra defines
NUM_KAMAILIO_WORKERS	number of kamailio workers to start	8
NUM_CALL_PROCESSORS	number of call processor instances to start	1
ADMIN_PASSWORD	admin password	admin
ENABLE_HTTPS	enable(1) or disable(0) https for gui	0 (disabled)
GUI_PORT	listen port for gui	8080
API_PORT	listen port for rest apis	5000
ENABLE_ENUM_PROCESSOR	enable(1) or disable(0) enum processor	0 (disabled)
ENUM_PROCESSOR_PORT	listen port for enum traffic	53
NUM_ENUM_PROCESSORS	number of enum processor instances to start	1
ENABLE_HTTP_PROCESSOR	enable(1) or disable(0) http processor	0 (disabled)
HTTP_PROCESSOR_PORT	listen port for http traffic	6000
NUM_HTTP_PROCESSORS	number of http processor instances to start	1
MANAGER_HOSTS	list of comma separated of manager node addresses	127.0.0.1 (single node deployment)
IMPORT_PATH	path for importing datamodels and service logics at start	/import

3.4.4 Sample deploy of 2 EMs and 2 CPs

3.4.4.1 EMs

To start container on first EM:

```
1 sudo docker run -d --name sre -v /var/log/sre:/var/log/sre -v /data/sre:/data/  
↔ sre --network host \  
2 -e MANAGER_HOSTS=<em1 address>,<em2 address> \  
3 -e DB_HOST=<em1 address> \  
4 -e MONGODB_HOSTS=<em1 address>,<em2 address> \  
5 -e INFLUXDB_HOSTS=<em1 address>,<em2 address> \  
6 -e INFLUXDB_TOKENS=<influxdb token on em1>,<influxdb token on em2> \  
7 --restart=always netaxis/sre
```

To start container on second EM:

```
1 sudo docker run -d --name sre -v /var/log/sre:/var/log/sre -v /data/sre:/data/  
↔ sre --network host \  
2 -e MANAGER_HOSTS=<em1 address>,<em2 address> \  
3 -e DB_HOST=<em2 address> \  
4 -e MONGODB_HOSTS=<em1 address>,<em2 address> \  
5 -e INFLUXDB_HOSTS=<em1 address>,<em2 address> \  
6 -e INFLUXDB_TOKENS=<influxdb token on em1>,<influxdb token on em2> \  
7 --restart=always netaxis/sre
```

3.4.4.2 CPs

Start the container with the following options:

```
1 sudo docker run -d --name sre -v /var/log/sre:/var/log/sre -v /data/sre:/data/  
↔ sre --network host \  
2 -e ENABLE_MANAGER=0 \  
3 -e DB_HOST=<cp address> \  
4 -e MONGODB_HOSTS=<em1 address>,<em2 address> \  
5 --restart=always netaxis/sre
```

3.4.5 All-in-one deployment

For an all-in-one deployment only docker runtime is required on the host.

To start SRE run:

```
1 sudo docker run -d --name sre -v /var/log/sre:/var/log/sre -v /data/sre:/data/  
↔ sre --network host --restart=always netaxis/sre
```