


Database-as-a-Service In Your Own Datacenter

OVERVIEW

04	CENTRALIZED MANAGEMENT
08	HIGH AVAILABILITY & DISASTER RECOVERY
09	CONNECTION POOLING
10	MONITORING & ALERTS
13	PERFORMANCE MANAGEMENT
18	REPORTS
19	BACKUPS & RESTORES
22	ADMINISTRATION
25	SECURITY
28	SUPPORT



“ Database-Platform-as-a-Service (DBPaaS) is the fastest growing segment in the Platform-as-a-Service (PaaS) sector for its ability to break customers out of cloud vendor lock-in and support their need for agility. ”



Gartner

CENTRALIZED MANAGEMENT

Control Sprawl

Creating Clusters

Provisioning

Deprovisioning

Hybrid Clouds

Admin Access

OPTIMIZE YOUR ORGANIZATIONAL AGILITY

With the unstoppable rise of open-source databases, businesses are adopting multiple database types to power their applications. This creates the complex challenge of ensuring all components of your infrastructure are organized, compatible, and accessible to avoid mission-critical application failures through avoidable issues. While adopting new databases and technologies, DBA's are expected to do more with less resources at the ready. At ScaleGrid, we give you the tools to control database sprawl across your IT infrastructure.

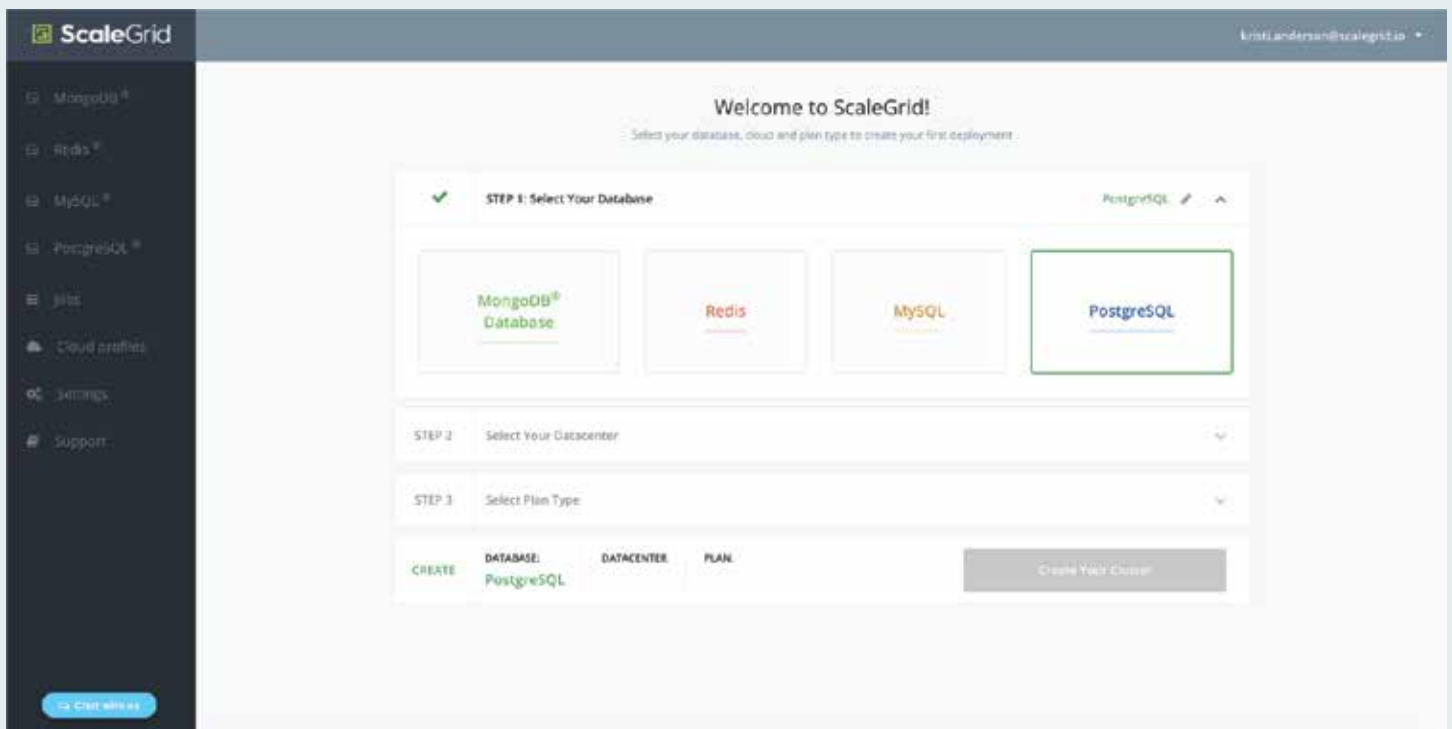
ScaleGrid Database-as-a-Service (DBaaS) platform allows developers, DevOps engineers, and database administrators (DBA) to automate daily tasks, optimize their infrastructure, and refocus their efforts on new innovations for their applications with speed and agility.

Understand your database versions, backup status, and performance history for all of the clusters in your organization through one single, centralized management console.

CONTROL DATABASE SPRAWL

ScaleGrid supports relational databases PostgreSQL and MySQL, as well as NoSQL databases, MongoDB and Redis, on AWS, Azure, DigitalOcean, VMware, Openstack and bare metal. Capture all of your key IT performance indicators for ultimate application performance visibility, and easily manage your infrastructure, resources, and information through one powerful platform.

DBaaS simplifies the administration of users, stakeholders, and permissions, implementing consistency and standardization into your deployments. With DBaaS automation, you can make tasks simpler and more effective, which is vital in the era of Big Data and agile demand-driven service based applications. Organizations can control server and database sprawl by applying a single source of truth (SSOT) to manage all of their database clusters. DBaaS also helps reduce IT staff time from performing redundant, time consuming activities such as provisioning, management, administration and monitoring, and enables your in-house Database Administrators (DBAs) to focus on more core, value-added activities such as logical database administration and application architecture rationalization.



CREATING POSTGRES SQL CLUSTERS

Easily create new PostgreSQL clusters in minutes. The intuitive creation wizard walks you through all the steps to customize your deployment setup in a few simple clicks. Visit our [PostgreSQL Help Docs](#) for step-by-step instructions to get started:

VMware
Deployments

Bare Metal
Deployments

Docker
Deployments

PROVISIONING POSTGRES SQL

Completely customize your new PostgreSQL cluster at creation. All configurations can be modified at any time in your PostgreSQL cluster profile.

VM Size

Customize your virtual machine size, disk size, disk IOPS, etc to meet your deployment needs.

PostgreSQL Version

Use the PostgreSQL version you need and upgrade at any time YOU choose with the help of ScaleGrid support.

Replication

Customize your replica setup as a primary-standby-standby or primary-standby-quorum, and distribute across locations.

Enable SSL

Secure your PostgreSQL clusters by enabling SSL to create a secure connection between your server and client.

Data Disk Encryption

Enable disk encryption to encrypt with Linux Unified Key Setup (LUKS) to securely encrypt storage media volumes.

Enable PgBouncer

Leverage the number one PostgreSQL connection pooler to manage your database connections and improve performance.

Firewall Rules

Lock down access to the intended clients of your deployments by creating firewall rules at the cluster and account level.

Data Imports

Enter your server name, port, DB name, username and password, authentication DB, and SSL strategy to import your data.

DEPROVISIONING POSTGRESQL

Save hours deleting clusters that are no longer needed by permanently removing them in two clicks. Delete directly from your PostgreSQL cluster page, and confirm by simply entering the cluster name to verify your intent.

HYBRID CLOUD

Manage both your PostgreSQL cloud and on-premise deployments run in your private datacenter(s) through one centralized console. Easily distribute your data across both public and private clouds to strengthen your high availability, increase your security, and leverage the flexibility of the cloud.

ADMIN ACCESS

You also maintain full PostgreSQL superuser admin access so you're never locked into a DBaaS, and can migrate out as easily as you came in. Adopt a vendor-agnostic approach to your infrastructure to maintain an agile strategy, and eliminate the financial risk of committing to a single provider.

HIGH AVAILABILITY DISASTER RECOVERY

Configure your high availability framework so you can instantaneously recover from potential failures that could impact in any part of your system. We help you implement redundancy in your infrastructure and data, failure detection and correction mechanism, a failover mechanism, as well as an application and user redirection mechanism to ensure that all of the necessary procedures are able to carry forward to the new master.

HIGH AVAILABILITY

At ScaleGrid, we enable you to deploy your PostgreSQL replica sets across multiple datacenters or Availability Zones (AZ's) so you can minimize the impact of losing any particular AZ. If the AZ of the master goes down for any reason, one of the replicas is automatically elected as the new master.

DISASTER RECOVERY

We also enable you to distribute your PostgreSQL replica sets across regions or data centers. This way, even if an entire region goes down, it does not affect the availability of your replica set.

Automate failovers to standby servers using industry standard high availability (HA) frameworks for PostgreSQL on-premise or in the cloud.

CONNECTION POOLING

Manage your database connections with PgBouncer, the most popular PostgreSQL Connection Pooler, built-in to your database management platform. PgBouncer is a middleware that sits between the database and clients and creates a pool that clients can connect to. We allow you to enable PgBouncer in the cluster creation process, and customize these settings to meet the needs of your deployment:

Pooling Mode: Determines how soon connections return to the pool.

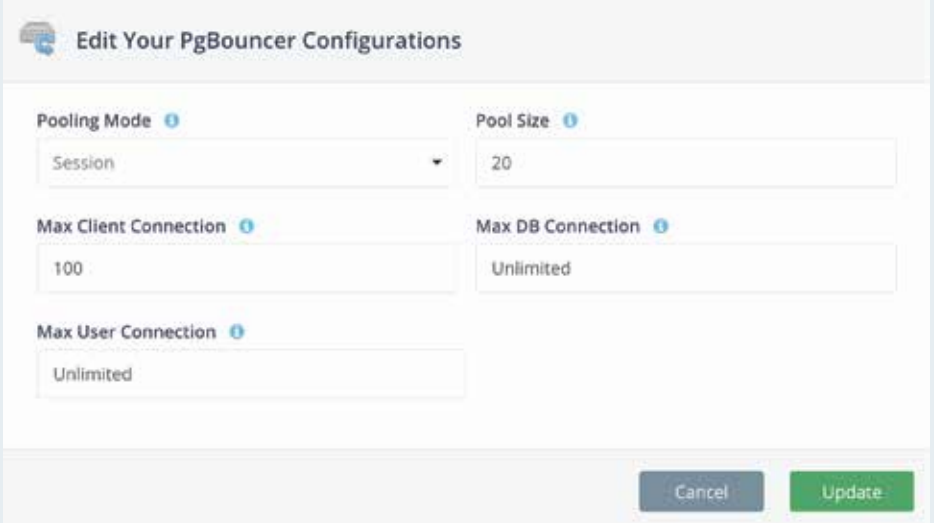
Pool Size: Max cached connections per pool.

Max Client Connection: Max client connections PgBouncer will allow.

Max DB Connection: Max connections to a single database PgBouncer will allow.

Max User Connection: Max connections by a single user PgBouncer will allow.

The ScaleGrid platform handles all of the necessary PgBouncer maintenance, backend configurations, upgrades and security patches to keep the connection pooler optimized for your deployment.



The screenshot shows a web interface titled "Edit Your PgBouncer Configurations". It contains five input fields with informational icons (i) next to their labels:

- Pooling Mode:** A dropdown menu currently set to "Session".
- Pool Size:** A text input field containing the value "20".
- Max Client Connection:** A text input field containing the value "100".
- Max DB Connection:** A text input field containing the value "Unlimited".
- Max User Connection:** A text input field containing the value "Unlimited".

At the bottom right of the form are two buttons: a grey "Cancel" button and a green "Update" button.

MONITORING

ALERTS









MONITORING

All PostgreSQL solutions include advanced monitoring to ensure the continuous health of your deployments. Each cluster comes with its own unique Monitoring Console that allows you to analyze your PostgreSQL and OS metrics, including CPU, memory, disk, operations, connections, and locks. Here are all the supported metrics:

CPU - User (%)	Locks - Row Share
CPU - System (%)	Locks - Share Update Exclusive
CPU - Nice (%)	Locks - Share
CPU - IO Wait (%)	Locks - Row Exclusive
CPU - Total (%)	Locks - Exclusive
Memory - Total (KB)	Locks - Access Exclusive
Memory - Used (KB)	Locks - Total
Memory - Free (KB)	Scheduled Checkpoints
Memory - Buffers (KB)	Requested Checkpoints
Memory - Cached (KB)	Temporary Bytes Written to Disk (Bytes)
Disk - Read (KB/sec)	Replication Lag (secs)
Disk - Write (KB/sec)	Cache Hit Ratio
Operation - Select (per second)	Deadlocks
Operation - Update (per second)	Conflicts
Operation - Insert (per second)	Write Ahead Log Size (Bytes)
Operation - Commit (per second)	Memory Used by Postgres Processes (KB)
Operation - Delete (per second)	Network In (KB/sec)
Operation - Rollback (per second)	Network Out (KB/sec)
Active Connections	Load Avg - 1 min
Idle Connections	Load Avg - 5 mins
Locks - Access Share	Load Avg - 15 mins

Automate your PostgreSQL monitoring and alerts so you can ensure the health of your deployments.


POSTGRES MONITORING CONSOLE



CUSTOM ALERTS & RULES

Customize alerts on any PostgreSQL database metric you'd like to monitor so you can receive notifications on any potential issues with your deployment. Each cluster is established with default alerts to notify you once any dangerous threshold has been reached, and you can also override the global rules by configuring custom alerts at the cluster and account-level so your alerts and rules meet the needs of your deployment.

We also allow you to receive alerts through your notification method of choice. Get your alerts through your email or SMS text messages, or leverage one of our alert integrations, including PagerDuty, Opsgenie, and Slack.

 **Create Cluster-Level Alert Rule**

Pick a metric

Notifications

Create rules on Operating System and Database Metrics.

Alert type

Metric

Metric

CPU - User (%)

Condition

Greater than

Threshold (%)

Condition lasts for

2 mins


Cancel

Next


PERFORMANCE MANAGEMENT

EASY TROUBLESHOOTING WITH SLOW QUERY ANALYSIS

Generate slow query heatmaps to identify queries that need indexes across any time period. By default, slow queries are defined as queries that take over 100ms, presented by type through color-coded bubbles in your heatmap for easy pattern recognition. Click into a bubble to see the underlying query, the time it happened, and its duration:



Visualize the top queries causing read and write-load on your system. Hover over a query to identify the queries that are typically not indexed:



At ScaleGrid, we give you several tools to continuously track, benchmark, and identify issues that might be affecting your performance, and quickly optimize to improve your speed through proper allocation and application of resources.

SLOW QUERY ANALYZER CONTINUED

You can also analyze your PostgreSQL slow queries in a tabular format:

Slow Query Details

Show 10 entries

Search:

#	<input checked="" type="checkbox"/>	Type	Users	Query	Database	Count	Plan Node types (Visualize)	Duration (ms)	Rows Returned (Avg)	Rows scanned	Rows modified
1	<input checked="" type="checkbox"/>		sgpos:gres	INSERT INTO TestTabXXX VALUES (XXX,XXX, XXX, XXX)	sq	11924	ModifyTable, Result	0	0	0	1
10	<input checked="" type="checkbox"/>		sgpos:gres	SELECT max(age(datfrozenxid)) FROM pg_catalog.pg_databases	postgres	29	Aggregate, Seq Scan	0	1	6	0
11	<input checked="" type="checkbox"/>		sgpos:gres	SELECT pg_catalog.pg_current_xlog_location()	postgres	29	Result	0	1	0	0
12	<input checked="" type="checkbox"/>		sgpos:gres	SELECT pg_catalog.pg_xlog_location_diff(CAST(XXX AS pg_lsn), CAST(XXX AS pg_lsn))	postgres	29	Result	0	1	0	0
13	<input checked="" type="checkbox"/>		sgpos:gres	select count(*) * XXX * XXX * XXX as total_size from pg_catalog.pg_ls_dir(XXX)	postgres	29	Aggregate, Function Scan	2	1	18	0
14	<input checked="" type="checkbox"/>		sgpos:gres	SELECT buffers_backend as Backend, buffers_clean as BGWriter, buffers_checkpoint as Checkpoint FROM pg_catalog.pg_stat_bgwriter	postgres	29	Result	0	1	0	0
15	<input checked="" type="checkbox"/>		sgpos:gres	SELECT SUM(tup_inserted) AS Inserted, SUM(tup_updated) AS Updated, SUM(tup_returned) AS Returned, SUM(tup_fetched) AS Fetched, SUM(xact_commit) AS Commit, SUM(xact_rollback) AS Rollback, SUM(temp_bytes) AS	postgres	29	Aggregate, Seq Scan	13	1	6	0
16	<input checked="" type="checkbox"/>		sgpos:gres	SELECT mode, COUNT(graded) FROM pg_catalog.pg_locks WHERE graded=XXX GROUP BY mode	postgres	29	Aggregate, Function Scan	2	2	3	0
17	<input checked="" type="checkbox"/>		sgpos:gres	select extract(epoch from current_timestamp - pg_catalog.pg_postmaster_start_time()) as uptime	postgres	29	Result	0	1	0	0
18	<input checked="" type="checkbox"/>		sgpos:gres	select count(*) as total_wal_files from pg_catalog.pg_ls_dir(XXX)	postgres	29	Aggregate, Function Scan	0	1	18	0

Showing 1 to 10 of 24 entries

< Previous

1

2

3

Next >

Type	Users	Query	Database	Count (default)
Plan Node Types (Visualize)	Duration (ms)	Rows Returned (Avg)	Rows Scanned	Rows Modified

QUERY PLAN VISUALIZER

For complicated queries, ScaleGrid provides a Query Plan Visualizer so you can easily visualize the ordered set of steps picked by the query plan. Visually examine the optimal plan that was selected and is expected to run the fastest, and click into the step to examine the various key and value details so you can optimize your PostgreSQL slow queries.

ScaleGrid PostgreSQL Plan Visualizer

UPDATE pgbench_branches SET balance = balance + 100 WHERE bid = 100;

ModifyTable
Seq Scan

ID	Node Type	Startup Cost	Total Cost
0	ModifyTable	0	3.25
1	Seq Scan	0	3.25

Plan Details

Key	Value
Alias	pgbench_branches
Id	0
Operation	Update
Parallel Aware	false
Plan Rows	1
Plan Width	370
Relation Name	pgbench_branches
Schema	public

Cancel

Actual Rows	Actual Loops
0	1
1	1

ACTIVITY STATS

Instead of analyzing PostgreSQL stats by running and deciphering individual commands, save hours with ScaleGrid’s PostgreSQL Activity Stats which allows you to visualize your key server activity statistics, including:

Size Stats

Tables	Total number of tables.
Indexes	Total number of indexes.
Size	Size of database on disk.
Table Size	Largest tables by size on disk.

Bloat & VACUUM Stats

Dead Rows	Tables in this deployment with most dead rows. Higher number of dead rows can impact performance, and indicate bloated tables.
Table Bloat Estimate	An estimate of table bloat computed from usage and size statistics. This is an estimate only - please contact ScaleGrid Support to calculate actual bloat metrics
Index Bloat Estimate	An estimate of the bloat on your BTree Indexes computed from usage and size statistics. This is an estimate only - please contact ScaleGrid Support to calculate actual bloat metrics
Auto-Vacuum - Most Frequent	Most frequently vacuumed tables in this deployment. These are the tables where auto-vacuum has been called most often. This typically means there is a lot of updates/deletes going on in this table.

ACTIVITY STATS CONTINUED

Read & Write Stats

Read Throughput	Largest tables in this deployment, by number of rows fetched from database.
Update Throughput	Tables with most rows updated. This graph shows how many rows were updated in the normal manner, (Number of Updates), vs how many updated rows took advantage of Heap-Only-Update (Number of HOT Updates). HOT Updates are designed to be faster, but work only if no indexed column is being updated. If your total updates are very high for a table, and very few of them are HOT updates, your update performance will degrade.
Activity Select	Top tables by select activity.
Activity - Read Index	Top indices by scan activity.
Activity - Insert	Top tables by insert activity.
Activity - Update	Top tables by update activity.
Cache Effectiveness	Top tables, by cache effectiveness. This graph shows the ratio of (Blocks read from Buffer Cache) to (Total Blocks Read). If this ratio is very low, your database performance might degrade. Using a machine with more RAM would help improve performance.
Scan	Top tables, by number of sequential scans initiated. If the number of sequential scans is much higher than number of index scans, it might

REPORTS

MONTHLY REPORTS

Get monthly reports sent right to your inbox every month. Broken down by server, quickly analyze the key metrics of your database performance:

Cluster Activity (per second)

- Rows Selected
- Rows Updated
- Rows Inserted
- Rows Deleted
- Commit Ops
- Rollback Ops

Database Statistics

- Databases Added
- Databases Removed
- Databases
- Storage Size
- Rows
- Index Size

Database Size Growth (Top 10)

Table Details

- Tables Added
- Tables Removed
- Indexes Added
- Indexes Removed
- Tables
- Storage Size
- Rows
- Index Count
- Index Size

Table Growth Stats By Size (Top 10)

Table Growth Stats By Rows Count (Top 10)

Activity Stats

- Tables
- Indexes
- Size
- Table Size
- Index Size
- Dead Rows
- Table Bloat Estimate
- Index Bloat Estimate
- Auto-vacuum - Most Frequent
- Read Throughput
- Update Throughput
- Activity - Select
- Activity - Read Index
- Activity - Insert
- Activity - Update
- Cache Effectiveness
- Pgbench.pgbench_branches
- Cache Effectiveness
- Scan

Stay on top of your database performance with cluster-level monthly reports emailed directly to you each month, and saved to your Postgres dashboard.

BACKUPS & RESTORES


Stay proactive on disaster recovery and business continuity planning to protect your mission-critical applications. Leverage tools and best practices to solidify a comprehensive strategy and safeguard your data through regular and consistent backup procedures.

ON-DEMAND & RECURRING BACKUPS

Automate recurring backups for your PostgreSQL standalone instances and replicas so your production clusters are backed up at regular intervals. ScaleGrid also allows you to backup on-demand so you have the most current data available before performing any application changes.

The easiest way to recover your PostgreSQL data. ScaleGrid automates your restore jobs so you can quickly recover a backup of your PostgreSQL instance or replica set in a single click.

Ensure your business is able to recover from a detrimental disaster event, while implementing the most powerful safeguards to protect against it.

 **Backup - AWSTest**

On-Demand

Scheduled

☒ Enable Scheduled Backups?

Every

1 day

Start time (UTC)

14:00

Max backups to retain

7

Cancel


Schedule

CONTINUOUS BACKUPS

Backup your data in real-time with continuous backups so all of your data is backed up automatically every few seconds so you always have the most recent data available to recover when you need it.

ONE-CLICK RESTORES

Recover your data in seconds with one-click restores. The management console allows you to select from your scheduled, on-demand backups, or continuous backups from your selected point in time, and encourages you to perform an on-demand backup before restoring so you can preserve your existing data.

 **Restore - AWSTest**

WARNING: Restoring will erase all your existing data. Please create a new Backup if you would like to preserve the existing data.

Backup List

Scheduled - 05/09/19 14:00 UTC-9uMUvg


Name	Scheduled - 05/09/19 14:00 UTC-9uMUvg
Taken on	2019-05-09 02:11:43 +00:00

Cancel

Restore

SELECTIVE RESTORES / PEEK AT BACKUP

Peek at your backups through our selective restores so you easily recover a single row or document from a table or collection without having to impact your entire production deployment. The peek at backup feature creates a new standalone instance with the data from the selected backup, and then allows you to perform a selected partial restore for a single collection.

 **Peek - TestCluster2**

Create a new standalone instance with the data from the selected backup

Cluster Name
P-TestCluster2-5234

Backup List
Scheduled - 10/21/19 13:00 UTC-W5QBUE

Name	Scheduled - 10/21/19 13:00 UTC-W5QBUE
Taken on	2019-10-21 01:26:32 +00:00 UTC

CancelPeek

RPO & RTO

Establish a Recovery Time Objective (RTO) to determine the amount of time it should take to do a full recovery (dependent on the amount of data in the database), and a Recovery Point Objective (RPO) to address these critical questions:

- How much data will you lose if all your servers melt down today?
- Have you setup continuous backup?
- Can you recover to any point in time or just to certain points?

ADMINISTRATION

CLONE

The PostgreSQL Clone feature allows you to take a look at the contents of either a scheduled or on-demand a backup without impacting your main deployment, and carry over the configurations and contents of the backup to your new cluster. Clone is a useful tool for testing the effects any change would have on your production deployment in a separate environment, which has the exact setup as that of production.

SCALE

Dynamically scale your PostgreSQL deployments in a single click. The ScaleGrid platform performs a 'rolling operation' that upgrades one node at a time with zero downtime. The only visible impact to the application is a switchover of the master from one node to another.

Automate your database administrators time-consuming tasks so they can focus on proactive optimizations and development.



The image shows a screenshot of the 'Scale - TestCluster2' dialog box in the ScaleGrid platform. The dialog has a title bar with a blue icon and the text 'Scale - TestCluster2'. Below the title bar, there is a descriptive text: 'Scale allows you to dynamically scale up your CPU, memory and disk size. If you have a replica set, the scale operation performs a rolling upgrade resulting in almost no downtime'. Underneath this, there is a section titled 'New size' with a dropdown menu currently showing 'X4XLarge'. Below the dropdown, there is a list of specifications for the selected size: 'r4.2xlarge - 64 GB of memory', '8 virtual cores with 27 EC2 Compute Units', '64 bit', '700 GB', and '4000 EBS Provisioned IOPS'. At the bottom right of the dialog, there are two buttons: 'Cancel' and 'Scale'.

Scale - TestCluster2

Scale allows you to dynamically scale up your CPU, memory and disk size. If you have a replica set, the scale operation performs a rolling upgrade resulting in almost no downtime

New size

X4XLarge

r4.2xlarge - 64 GB of memory
8 virtual cores with 27 EC2 Compute Units
64 bit
700 GB
4000 EBS Provisioned IOPS

Cancel Scale

IMPORT

Import your existing PostgreSQL data into your PostgreSQL deployment at ScaleGrid. The Import wizard walks you through the steps to migrate your data over, and ensure that the data is compatible with your newly created cluster.

FOLLOWER CLUSTERS

Easily import data on a scheduled basis from an existing PostgreSQL cluster managed through ScaleGrid into another cluster in your account by setting up a "follower cluster". This is very useful when you have a development or testing setup that you wish to sync with the production server on a regular schedule. Once you create a follower relationship, each day at your chosen time, the latest backup from the source cluster will be restored on to the destination cluster.

here'. The main content area has two sections: 'Select source cluster' with a dropdown menu showing 'ProdCluster', and 'Start time (UTC)' with a dropdown menu showing '0:00'. At the bottom right are two buttons: 'Cancel' (gray) and 'Schedule' (green)." data-bbox="47 529 635 794"/>

 **Setup as Follower - TestCluster2**

For more information please refer [here](#)

Select source cluster

ProdCluster

Start time (UTC)

0:00

Cancel Schedule

POSTGRESQL CONFIGURATION MANAGEMENT

The ScaleGrid platform allows you to easily manage the configuration of all of your PostgreSQL clusters. Update your configuration settings, and allow ScaleGrid to handle the necessary adjustments behind the scenes. Any dynamic adjustments require a restart, and ScaleGrid will automatically perform a rolling upgrade to apply the configuration to all your nodes.

Parameter Name	Value	Description
autovacuum	on	Starts the autovacuum subprocess
autovacuum_analyze_scale_factor	0.1	Number of tuple inserts, updates, or deletes prior to analyze as a fraction of ...
autovacuum_analyze_threshold	50	Minimum number of tuple inserts, updates, or deletes prior to analyze
autovacuum_freeze_max_age	200000000	Age at which to autovacuum a table to prevent transaction ID wraparound
autovacuum_max_workers	3	Sets the maximum number of simultaneously running autovacuum worker p...
autovacuum_multixact_freeze_max_age	400000000	Multixact age at which to autovacuum a table to prevent multixact wraparou...
autovacuum_naptime (s)	60	Time to sleep between autovacuum runs
autovacuum_vacuum_cost_delay (ms)	20	Vacuum cost delay in milliseconds, for autovacuum
autovacuum_vacuum_cost_limit	-1	Vacuum cost amount available before napping, for autovacuum

MANAGE DATABASE VERSIONS

Use the PostgreSQL version of your choice, and never be forced into an upgrade before you're ready. We support a minimum of three (3) versions of each database, and will only discontinue a version after all of our users have upgraded to a newer version.

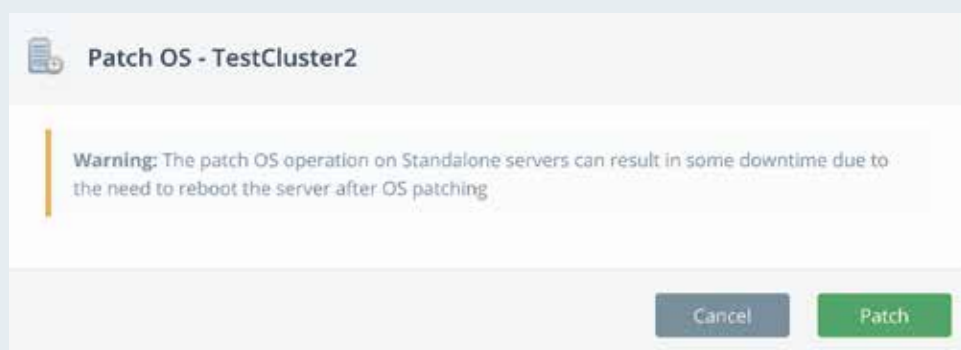
UPGRADE DATABASE VERSIONS

Save time and headaches with fully managed database upgrades. At ScaleGrid, our expert support team performs your upgrades on your schedule, and works with you to ensure all the necessary preparations are made before executing your upgrade.

PATCH OPERATING SYSTEM

Database clusters managed and hosted by ScaleGrid are patched on a monthly basis. The clusters are typically patched on the last weekend of the month, and only security related patches are installed. A custom maintenance time window can also be specified for the patching operation.

You can also trigger on-demand OS patch jobs from the console. The OS patching is performed on a “rolling” basis – one server at a time is patched and brought back into operation. This allows the entire cluster to be



SECURITY

Optimizing for security has never been more necessary, as organizations are only as successful as the protection of their customers and proprietary data. Data is the lifeline of an organization, and companies must take every action necessary to fortify their infrastructure to ensure the safety of the information stored and leveraged within their applications.

BEST PRACTICE DEPLOYMENT

Wondering if your infrastructure is designed for a best practice deployment? Our experts understand and design the best practices for deploying, managing, and securing your database applications, and are deeply knowledgeable on the complex intricacies of both NoSQL and SQL in the cloud and on-premise. At every stage, starting with migration to the ScaleGrid platform, we optimize your deployments to the best possible configuration for the unique needs of your application so you face the lowest possible security threats and are backed by a team of experts who will quickly address any security challenge your infrastructure ever encounters.

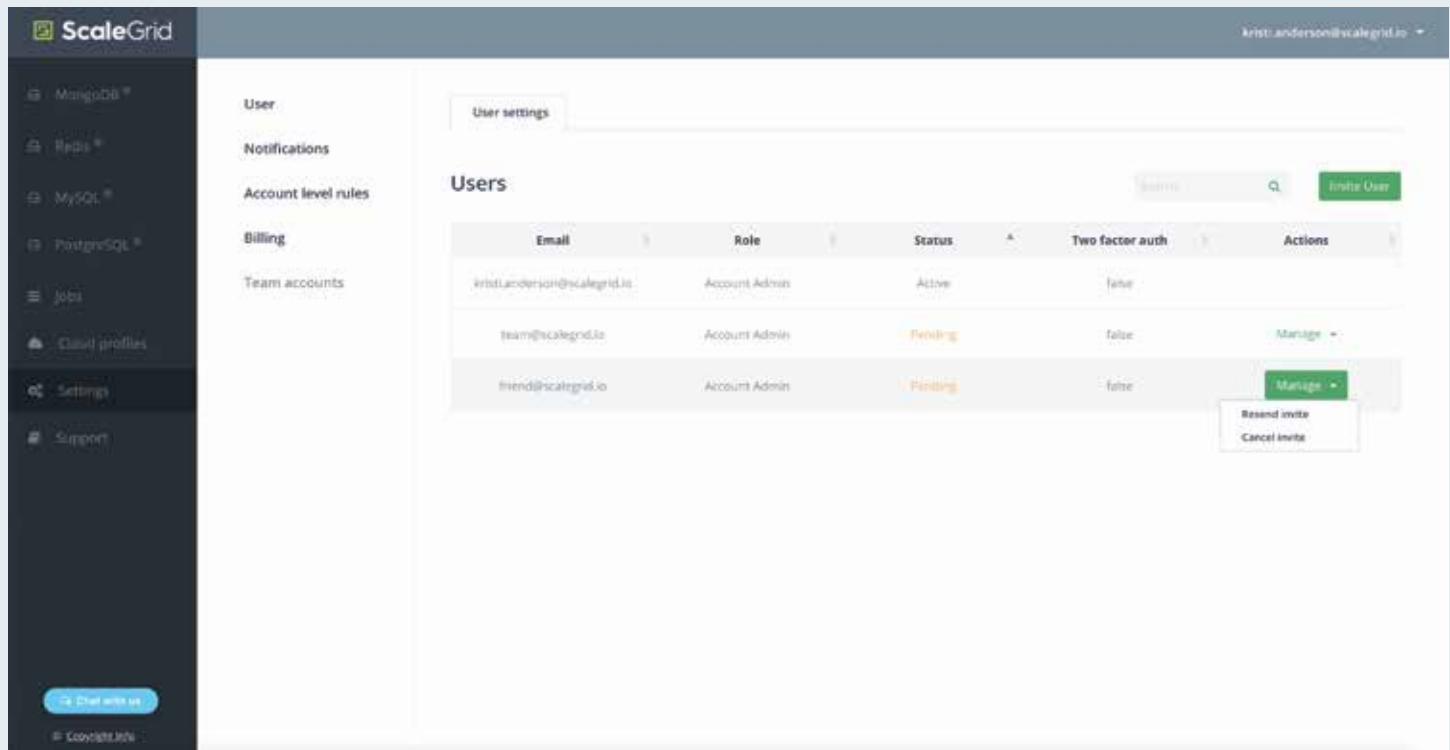
TWO-FACTOR AUTHENTICATION

Implement an additional layer of security over your ScaleGrid console account username and password with two-factor authentication. Any application that supports time-based, One-Time Password (OTP) protocol may be used, and Google Authenticator App is the most popular used by our customers.

Implement our industry-leading best practices for your PostgreSQL cloud deployments to develop a solid infrastructure that's positioned for growth. Fortify your mission critical applications with advanced security tools to protect your data.

TEAM MANAGEMENT

Invite an unlimited number of users to your ScaleGrid account so your data-base team can collaborate and manage shared resources through the PostgreSQL management platform. Easily invite new users in a few simple clicks, and control their privileges through a desired role assignment. You can also enforce two-factor authentication for your users, and delete users anytime they no longer need or are allowed access to the platform.



The screenshot displays the ScaleGrid user management interface. On the left is a dark sidebar with navigation links: MongoDB, Redis, MySQL, PostgreSQL, Jobs, Cloud profiles, Settings, and Support. The main content area is titled 'Users' and includes a search bar and an 'Invite User' button. Below the header is a table with columns: Email, Role, Status, Two factor auth, and Actions. The table lists three users: kristi.anderson@scalegrid.io (Account Admin, Active), team@scalegrid.io (Account Admin, Pending), and friend@scalegrid.io (Account Admin, Pending). A dropdown menu for the 'friend@scalegrid.io' user shows options: 'Manage', 'Resend invite', and 'Cancel invite'.

Email	Role	Status	Two factor auth	Actions
kristi.anderson@scalegrid.io	Account Admin	Active	false	
team@scalegrid.io	Account Admin	Pending	false	Manage
friend@scalegrid.io	Account Admin	Pending	false	Manage Resend invite Cancel invite

ENCRYPTION AT REST

Implement encryption at rest using Linux Unified Key storage (LUKS), the standard for volume based encryption in Linux. The LUKS encrypting keys are stored both on the ScaleGrid controller servers and on the database servers.

On the database servers, the keys are stored under root credentials and are used to mount the volumes when the server restarts. This ensures the data volumes are encrypted and can only be mounted on servers of this cluster. When encryption at rest is enabled, backups are also encrypted to ensure they can only be restored to the servers of this cluster.

ENCRYPTION IN TRANSIT (SSL)

Encrypt your PostgreSQL data “in transit” using SSL to secure communications over network. Configured at cluster creation, PostgreSQL is deployed in requireSSL mode, requiring SSL on all inbound connections and replication traffic between the PostgreSQL replicas.

By default, self-signed certificates are installed during cluster creation, and you can additionally bring your own wildcard certificates and install them on the servers.

SUPPORT

Regardless of your deployment size, our experts are here to provide the technical expertise necessary to address any issue that may arise with your database operations.

24/7/365 SUPPORT

At ScaleGrid, 24/7 expert support is included with all of our Cloud DBaaS plans, and premium enterprise-level support packages for our on-premise solutions. From the moment you begin, we help you migrate all of our deployments, customize your settings, deploy your clusters through our advanced Management Console, and continue that support throughout your entire lifespan with our company.

NICHE TROUBLESHOOTING EXPERTISE

Get support from the actual database developers who designed the ScaleGrid platform to help you quickly identify and resolve any issue your infrastructure faces. Our entire support team has niche expertise in PostgreSQL on both public cloud setups (AWS, Azure, DigitalOcean), and on-premise. Trust the experts who improve the infrastructure security, availability, and performance for thousands of deployments across the world to save your team countless hours troubleshooting minor to mission-critical issues.

“I have to come to rely on the stability and support that ScaleGrid has provided. ScaleGrid has always been there to support any needs we have and the software just works. We have several collections with millions of records and there is never an issue.”



CONTACT US

Get in touch to learn more about our Enterprise On-Premise PostgreSQL database management and support solutions to help optimize your private datacenter environments.

sales@scalegrid.io

support@scalegrid.io

[On-Premise Solutions](#)

[PostgreSQL Solutions](#)

[Free 30-Day Trial](#)

[Documentation](#)