


WHAT ARE YOUR **GOALS**?



WITH

faircom **db** V12




One Core Technology — Endless Possibilities

faircom db

V12 HIGHLIGHTS

FairCom DB V12 (formerly known as c-treeACE) has a variety of new enhancements and features for the c-treeACE developer. This document is a quick overview of these exciting new additions.

You may also want to read this document online at docs.faircom.com/doc/v12update. It contains links to each feature in the FairCom DB V12 Update Guide where you can learn more about each feature.



Better Availability. Scalability. Productivity.

As a reminder, we've updated our product names and logos.

c-treeACE is now **faircom db**, c-treeEDGE is now **faircom edge**, c-treeRTG is still **c-treeRTG**



Create Big Solutions

- **Multitenant Cloud Apps**
- **Huge Transactions**
- **Big Data Analytics**




Go Bigger

- Millions of records per transaction
- Millions of open files
- 128 TB temp tables
- 4 GB transaction logs
- 64 K strings
- 2,500 columns per table

Speed Up Everything

- **Faster Applications**
- **More Concurrent Users**
- **Predictable Performance**
- **Lower Hardware Costs**
- **Reduced Tuning Costs**
- **Faster Software Development**



Go Faster


- Up to 3x faster overall performance
- Up to 4x faster indexes with improved key compression, locking, node pruning & sorting
- Up to 3x faster SQL for joins, filters, sorts, ORs and parameterized queries
- Insert many SQL rows in one statement
- Faster server-side bulk inserts, updates & deletes with up to millions of inserts per second
- New batch retrieval with server-side query expressions & index ranges
- Return specified fields instead of entire records
- Process data files in parallel forwards & backwards
- Faster file open, close, encryption & truncate
- Faster connections and app communications
- Faster transaction log checkpoints
- More concurrency with less lock contention
- Faster replication that runs in parallel
- Faster encryption at rest and in transit
- Low-cost OS compression on files for faster I/O
- Configure when data, index & logs flush to disk
- For extreme speed, run database in-process by dynamically or statically linking DB into app





Develop Faster

- **Lower Programming Costs**
- **Parallel Processing**
- **Distributed Processing**
- **New Capabilities**
- **New Languages**



Develop Easier

- Python NAV API
- Python SQLAlchemy
- JavaScript NAV API
- JSON & UTF-8
- Use plug-ins to run anything server-side
- Run SQL queries across databases
- Process data in parallel
- Process all data files forwards & backwards
- Tag, find, move & cache files more easily
- Schedule jobs
- Run distributed command queues
- Use C++ API to control replication
- Use JSON/HTTP API to control replication
- Use millisecond-resolution timestamps
- Leverage 50+ new functions

Low Administration

- **Automatic Self Tuning**
- **Built-in Browser Apps**
- **Simply Works**



Goal: Zero Administration

- Use new built-in browser apps to replicate, explore & monitor data
- Automatically replicate everything & anything
- Automatically deploy and resync replicas
- Automatically tune indexes
- Automatically manage data growth
- Automatically alert on low disk space
- Automatically reclaim deleted record space
- Automatically timestamp records
- New or enhanced CLI utilities to manage:
 - Data replication
 - Data partitions
 - Encryption
 - KEEPOPEN_LIST (data caches)
 - Recording and replaying transactions





Increase Uptime

- Stay Up Longer
- Recover Faster
- Upgrade Without Downtime



Faster Recovery

- Back up to STDOUT
- Restore from STDIN
- Faster restores from large backups
- Wildcards exclude & include files in backups

High Availability (Beta)

- Parallel Synchronous Replication
- Built-in Failover Clustering
- Linux Cluster Failover
- Windows Cluster Failover

For details see the Replication Overview Brochure.

Replicate Data For

- Cloud Computing
- Distributed Reads
- Disaster Recovery
- Real-time Analytics
- Machine Learning
- Global Scalability
- Microservices



Data Replication

Synchronous and Asynchronous Replication can now run with parallel threads:

- Replicate data at high speed across many servers
- Create multiple read-only servers for bulk reporting, ad hoc reporting, dashboards, intensive queries, real-time analytics, machine learning, etc.
- Create multiple disaster recovery servers in remote locations
- Shard global transaction data across many databases in the cloud
- Create up-to-date caches of data across microservices
- Create ACID-compliant, read-only replicas for an app to write to one server and read from many
- Create eventually consistent replicas for an app to read and write to any server





Run More Securely

- **Encrypt Everything:**
 - In Transit
 - At Rest
- **Monitor Security Events**
- **Audit Security in Logs**
- **Use LDAP Directories**
- **Use AWS Secrets Manager**
- **Secure Logins**



Enhanced Security

- Encrypt everything at rest and in transit. Encryption overhead reduced from 17% to 5%
- Authenticate using LDAP
- Log all login attempts to SYSLOG
- Automatically expire passwords
- Automatically lock out accounts
- Automatically enforce password strength
- Block & unblock connections
- Manage inherited permissions more easily
- Create read-only servers
- Manage Master Encryption Key in a custom callback

Diagnose Faster

- **Find Potential Problems**
- **Fix Problems Faster**
- **Report Issues Easier**



Diagnose Easier

New configuration options to:

- Detect memory buffer overruns
- Manage diagnostic crash dumps
- Log flush status
- Log index node queue failures


Monitor each logged-in account for:

- Disk reads & writes
- Data & index cache requests & hits
- Connection information
- Logical & physical file opens & closes
- File creates, renames, deletes
- Average log save time
- Cached data stats





Easily Integrate With

- **Web Applications**
 - **Enterprise Message Queues**
 - **Cloud IoT, iPaaS**
 - **Enterprise Service Bus**
 - **SOA**
 - **B2B**
- 

New APIs

REST API

- Create tables & indexes
- Query, insert, update & delete data

MQTT API

- Integrate database data with all major message queue solutions
- Integrate with cloud IoT solutions


JSON/HTTP API

- Use remote procedure calls to automate monitoring, SQL, ISAM & data replication

HTTP App Server

- Run custom browser-based apps from the database's built-in app server
- These apps can leverage the REST and JSON/HTTP APIs

Run More Places

- **ARM**
 - **Raspberry Pi**
 - **macOS**
 - **IBM AIX**
 - **System 390**
 - **32-bit & 64-bit**
- 

New Platforms

- 32-bit and 64-bit ARM CPUs
- Raspberry Pi OS
- Apple macOS 10.15 Catalina
- Apple macOS 10.14 Mojave
- IBM AIX for Power 9 (P9)
- IBM System 390 Linux
- Microsoft Visual Studio 2019

Warp speed – defined as faster than the speed of light, 299,792,458 meters per second.

As a species, we have always strived to ever increase our speed. Going even faster improves life for everyone by saving time, getting more done, or providing a thrill.

Today, the consensus of scientific knowledge is that humans traveling faster than the speed of light is impossible. Yet, history is full of scientifically “absurd” theories that became a reality through years of determination and research.

Our V12 cover honors the determination to explore where we all can GO. Motion and progress, not only in physics but in dreams and aspirations, make the impossible infinitely possible.

Crossing the galaxy may not be the goal today; nevertheless, the drive to improve and to progress is inherent in all of us. Join us in taking another step towards our goals and the impossible.

