ACTIAN
a division of HCLSoftware

NEWSLETTER

HCL Informix® User Newsletter

January 2026

Welcome to the latest edition of the HCL Informix® User Newsletter.

This issue is packed with updates on server enhancements, cloud capabilities, tooling, performance features, and migration tools. Whether you're a developer or DBA, you'll find key insights and resources to make the most of Informix in your environment.

In This Issue

1. Product Pulse
   - Recent Releases
     - HCL Informix 15
     - InformixHQ 3.0.3 Enhancements
     - Informix 4GL Enhancements
     - Looking to the Future: HCL Informix Roadmap
2. Dev and DBA Corner
   - Database Migration Tools Overview
     - DBEXPORT/DBIMPORT
     - High-Performance Loader (HPL)
3. Use Case Spotlight
   - Financial Controlling
4. Community and Events
   - Tech Talk: User Friendly External Smartblobs by Mary Schulte
5. Hidden Gem of the Month
   - TimeSeries in Informix – Underused and Powerful
6. Common Support Tickets and FAQ's
   - The authors of This edition

Product Pulse

Recent Releases: HCL Informix 15

In November 2024 we released HCL Informix 15 and we have been making incremental updates since. Updates include:

Marketplace Listings

Now, customers can deploy HCL Informix 15 directly on AWS and Microsoft Azure, taking advantage of their committed cloud spend and streamline procurement through a familiar and secure billing environment. This milestone marks a major step forward in delivering flexible, modern deployment options for organizations running HCL Informix at scale while making cloud adoption easier, faster, and more cost-effective.

AWS Listing >

Microsoft Azure Listing >

Marketplace Listing Blog >

ARM Platform Support

HCL Informix now runs natively on ARM64 platforms such as AWS Graviton and Azure Ampere Altra. ARM-based instances use up to **60% less energy** than comparable x86 instances, with significantly lower carbon intensity per workload. Additionally, organizations migrating workloads from x86 to ARM64 typically see **30% cost savings** with improved price-performance ratios. For large-scale deployments, this translates to millions in annual savings.

Product Launch Blog >

Informix is a trademark of IBM Corporation in at least one jurisdiction and is used under license.

Looking to the Future: HCL Informix Roadmap

As we look to the year ahead and beyond, the Actian team is working to improve upon our V15 release to make it easy for all users to adopt. Here are a few highlights that we are looking to deliver in our upcoming winter release:

New licensing solution incorporated in HCL Informix 15.0.1

- Actian's license management solution has been incorporated within many Actian portfolio products, and is being integrated with the HCL Informix server effective with the 15.0.1 release. This replaces the previously deprecated Flexnet licensing solution, and provides convenient access/tracking of licenses while simplifying the self-auditing process for customers.
- Starting in 15.0.1, valid license file keys will be required during installation of the HCL Informix server. All customers and partners working on upgrades to V15.0.1 or later should familiarize themselves with the licensing changes and ensure attention is allotted to this small additional task during the upgrade process.
- The Actian team will distribute a detailed communication about the licensing changes to all partners and customers in December concurrent with the release. The changes will be explained in the 15.0.1 release notes "What's New" section, as well as on the **Actian Community website**.
- Via the website we will provide instructions and an FAQ for partners and customers to learn more about how to use the Actian licensing tools to generate license file keys and validate license entitlement.
- Live Q&A sessions will be conducted several dates between December-March to review the new licensing solution and any questions you may have. You can contact us with licensing questions here: **HCLInformix_licensing@hcl-software.com**.

Performance Improvements in 15.0.1

Direct I/O Support for 4K Devices on Linux

Enables efficient use of 4K devices for 4K DBspace page sizes (or multiples thereof). An internal benchmark evaluation comparing the performance of KAIO + Direct I/O verses AIO shows a 3 to 8 times reduction in checkpoint flush duration for 4K devices.

- Direct I/O bypasses the kernel's cache, ensuring that pages (which may contain committed data) are written straight to disk for permanent storage, which is critical for DBspace chunks.)
- This enhancement required changes to memory alignment for ad hoc buffers to align on 4K boundaries, resulting in increased memory usage.
- Use the DIRECT_IO configuration parameter to control the use of direct I/O for cooked files and block devices used for DBspace chunks. This parameter enables direct I/O (bypassing file system buffering) on UNIX platforms or concurrent IO (bypassing both file system buffering and unnecessary write serialization) on AIX® operating systems.
- The DIRECT_IO configuration parameter is a bitmask allowing multiple feature-control flags to be combined.

`onconfig.std value`	`DIRECT_IO 0`
`flag values`	• `0x0` = Neither direct I/O nor concurrent I/O is used • `0x1` = Direct I/O, which bypasses file system buffering, is used if available • `0x2` = Concurrent I/O is enabled on AIX operating systems (The concurrent I/O option includes direct I/O and concurrent I/O.) • `0x4` = Apply any of the above flags to temporary DBspaces as well as permanent spaces.
`takes effect`	After you edit your onconfig file and restart the database server.

ACTIAN™ a division of HCLSoftware

Usage

Direct I/O can only be used for DBspace chunks whose file systems support direct I/O for the page size.

If direct I/O is enabled, KAIO (kernel asynchronous I/O) is used if the file system supports it. However, KAIO is not used if the environment variable KAIOOFF is set when the server is started. When direct IO and KAIO are both used, the number of AIO virtual processors can be reduced. If direct IO is used, but KAIO is not, the number of AIO virtual processors should not be reduced.

HCL Informix does not use direct or concurrent I/O for temporary dbspace chunks unless the DIRECT_IO value contains the 0x4 flag.

On AIX, if HCL Informix uses concurrent I/O for a chunk, another program (such as an online external backup program) must also use concurrent I/O. If not, the file open operation will fail.

If HCL Informix uses direct I/O for a chunk, and another program tries to open the chunk file without using direct I/O, the open operation will normally succeed, but there can be a performance penalty. The penalty can occur because the file system might attempt to ensure that each open operation views the same file data, either by not using direct I/O at all for the duration of the conflicting open operation, or by flushing the file system cache before each direct I/O and invalidating the file system cache after each direct write.

Direct I/O is used for DBspace chunks on Windows platforms regardless of the value of the DIRECT_IO configuration parameter.

Example values

`0x1` = Direct I/O is enabled for non-temporary dbspace chunks

`0x2` = (AIX only) Concurrent I/O is enabled for non-temp dbspace chunks

`0x5` = Direct I/O is enabled for all chunks that can support it

`0x6` = (AIX only) Concurrent I/O is enabled for all chunks that can support it

Device Block Sizes on Linux

Prior to 15.0.1.x, direct I/O on Linux was supported only on storage devices with a native or emulated 512-byte block size. Starting from 15.0.1.x, direct I/O is also supported on devices with 2k and 4k block sizes.

Parallelizing Checkpoints

- This new feature parallelizes the initial phase of a checkpoint, allowing multiple threads to concurrently process multiple buffer pools, which accelerates the flushing phase.
- The algorithm also sorts chunks based on the number of bytes to be written, prioritizing larger writes.
- The feature should result in faster checkpoint flush times especially in an instance with many chunks of diverse page sizes.

Parallel Shared Memory Dump

- Users can specify up to 16 directories in the `DUMPDIR` parameter, ideally on different disks, to split and write a shared memory dump file concurrently, reducing the time required.
- A dump size must exceed 256 MB to be split.
- The split files (e.g., SHM.000000, SHM.000001) must be copied into a single directory. The `onstat -O` command is then executed to merge the separate pieces into a single file named "shm_outfile".
- A new 'compress' attribute was added to the `DUMPCNT` configuration parameter for managed shared memory dumps.

  For example: `DUMPCNT thrdlimit=3,compress=9`

  When `DUMPDIR` contains multiple comma-separated directories, enabling parallel shared memory dumping, this 'compress' attribute in the `DUMPCNT` parameter value enables automatic compression of the separate dump file streams. Valid compression level values are 1-9, with 1 using the fastest compression algorithm and 9 using the most effective one. The `onstat -O` utility will automatically uncompress the resulting files when creating shm_outfile.

HCL Informix Roadmap >

Suggest New Capabilities >

ACTIAN™ a division of HCLSoftware

Use Case Spotlight

Retail & Supply Chain Management

Challenge: Rapid changes in product supply and demand make it difficult to understand how much product to make or keep on hand at any point in time. Additionally, external factors such as weather, global health crises, and natural disasters can create sudden shifts in the supply chain that can be difficult to manage.

Why HCL Informix: With its lightweight footprint, HCL Informix is easy to deploy at every stage of the supply chain, from the manufacturer to the point of sale. It supports multiple data types, including geospatial data for location tracking, JSON files for photo proof of delivery, and in-store customer transactions. HCL Informix effectively unifies diverse needs across the entire supply chain.

Benefit: Managing inventory with HCL Informix empowers organizations to quickly adapt during times of crisis or ensure that they're maximizing revenue by staying on top of the latest consumer trends and mitigating supply chain disruptions.

Supporting Capabilities:

- High-performance transactional processing
- Lightweight footprint
- Multi-model native
- Automated data management

Application Ecosystem:

- Point of sale (POS) systems
- Online stores • Supplier and customer relationship management (CRM)
- Warehouse management systems
- Order processing and tracking
- Smart sensors
- Accounts payable
- Taxation reporting software

> For examples of more use cases, check out the **Use Case Selection Guide.**
>
> **Get the Guide >**

Community and Events

Webinar: What's New in Genero v6 & HCL Informix 15

Get caught up on everything you missed in our recent Genero v6 and HCL Informix 15 webinar, now available on demand. In this joint session from Four Js and Actian, experts walk through the latest UI/UX upgrades, built-in Prometheus monitoring, modern security enhancements, and what's new in HCL Informix 15, along with practical guidance for smoother modernization and migration. Watch the full 35-minute session plus Q&A anytime.

> **Watch On-Demand >**

Tech Talk: What's New in HCL InformixHQ 3.1.0 By Ajay Satbhadre

Take a deep dive into the latest version of HCL InformixHQ, the modern web-based monitoring and administration console for HCL Informix databases. In this in-depth session, Ajay Satbhadre from the HCL InformixHQ development team showcases the new features and usability improvements in version 3.1.0, designed to simplify database management and enhance the user experience.

What you'll learn:

- **Role Management:** How to create and assign custom roles with fine-grained privileges for tighter access control.
- **H2 Database Backup:** New options for automatic and manual backups of HCL InformixHQ configuration data to keep your settings safe and recoverable.
- **Database Utilities in the UI:** Perform DB import, export, and schema operations directly in the web interface without using command-line tools.
- **Improved DB Space and Log Views:** See critical DB spaces at a glance with new visual indicators and filtering options.
- **PWA Support:** Install HCL InformixHQ as a Progressive Web App for faster access, offline capability, and mobile-friendly monitoring.
- **Enhanced Server Logs:** Filter by server or log level (info, error, warning, debug) to trace issues faster and with more context.

This tech talk is ideal for devs and DBAs looking to modernize their toolset and make day-to-day database operations more intuitive.

Watch On-Demand >

Tech Talk: Enterprise Replication Rehash By Paul Watson

Catch this practical refresher from a long-time HCL Informix expert. In under an hour, Paul walks through the real-world setup and care-and-feeding of Enterprise Replication (ER)—what's changed since his original talk, what still bites, and how to keep clusters humming.

What you'll learn:

- **Topologies that work:** Primary-target vs. update-anywhere, root/leaf catalogs, and when to mix ER with HDR/RSS.
- **Must-have prerequisites:** Clock sync, trusted connections, and dedicated CDR spaces (why syscdr placement matters).
- **Define → Start → Verify:** The essential cdr define server/replicat, cdr start, and cdr check flow—plus when to pick "repair" over "sync."
- **Keys & consistency:** Primary/unique/ER keys, the truth about "ripplecheck", and fast, low-outage ways to add both to big tables.
- **Throughput wins:** Avoiding XXL transactions, sizing CDRQMEM, reading onstat ER views, and using "delayed apply" windows.
- **Recent enhancement:** Cascading replication (A→B→C) in recent 14.x/15 releases.

This tech talk is perfect for DBAs and developers who want resilient DR, near-real-time downstream copies, or smoother warehouse feeds, without app rewrites.

Watch On-Demand >

Hidden Gems of the Month

Webinar: What's New in Genero v6 & HCL Informix 15

LOGBUF_INTVL

Use the LOGBUF_INTVL configuration parameter to specify the frequency, expressed in seconds, at which the database server checks to determine whether a logical log buffer flush is needed. Flushing periodically may benefit an instance that uses buffered logging and experiences idle periods during the day. Records will not be allowed to remain in the log buffer longer than LOGBUF_INTVL seconds before they are flushed to the logical log on disk.

HCL Informix Documentation >

DUMPSHMEM

In HCL Informix, DUMPSHMEM is a configuration parameter within the onconfig file that controls whether a shared memory dump file is created when an assertion failure occurs in the database server. This dump file is a snapshot of the database server's shared memory at the time of the failure, which can be crucial for diagnosing the root cause of issues. Values:

- `0` = Do not create a shared memory dump.
- `1` = Create a shared memory dump of all the shared memory that the database uses.
- `2` = Create a shared memory dump that excludes the buffer pools.
- `5` = Enables the managed shared memory dump feature. When permitted, create a shared memory dump of all the shared memory that the database uses.
- `6` = Enables the managed shared memory dump feature. When permitted, create a shared memory dump that excludes the buffer pools.

HCL Informix Documentation >

ACTIAN™ a division of HCLSoftware

Common Support Tickets and FAQ's

High Availability & Failover

Q: "How do I set up an HCL Informix Connection Manager for high availability?"

A: Follow the HCL documentation here: **Connection Manager Setup Guide**. Once configured, Connection Manager (CM) can automatically handle failover between primary and secondary servers based on your `HA_FOC_ORDER` setting in the primary's onconfig file (for example, HDR,RSS).

Q: "What happens when the primary server goes down in an HDR cluster?"

A: The CM detects the failure and initiates automatic failover. You'll see log messages such as:

- ALARM 2001 failover arbitrator automated failover in progress.
- ALARM 2002 failover arbitrator automated failover completed.

The CM promotes the HDR secondary to the new primary, then updates routing so clients reconnect automatically. When the original primary comes back online, it detects the new primary and safely shuts down. To reintroduce it as a secondary, restart with `oninit -PHY` followed by `onmode -d` secondary.

Q: "If I want to reverse roles and make the old primary active again, how do I do that?"

A: Run:

```
onmode -d make primary <old_primary>
```

Then restart the new primary as a secondary using:

```
oninit -PHY
onmode -d secondary
```

This safely restores the original configuration while maintaining data integrity and proper client routing through CM.

Q: "What's the best practice configuration for Connection Manager?"

A:

- Run CM in Proxy Mode, so application connections and CM connections share the same operational status when failover occurs.
- Configure redundant Connection Managers as group entries in `sqlhosts`. This ensures client reconnection continues seamlessly if one CM instance fails.

Q: "What are the differences between FULL_SYNC, NEAR_SYNC, and ASYNC HDR secondary modes?"

A:

- `FULL_SYNC`: Primary waits for the secondary to apply each transaction before committing. Guarantees zero data loss but can reduce performance on high-latency networks. Best when **data integrity** is critical.
- `NEAR_SYNC`: Primary waits only for the secondary to receive the data, not apply it. Balances integrity and performance; default and recommended for most workloads.
- `ASYNC`: Primary commits immediately without waiting for the secondary. Delivers maximum performance but carries risk of data loss on failover. Use when **speed > safety**.

	For `FULL_SYNC`, ensure `SEC_NONBLOCKING_CKPT=1` on the secondary to prevent blocking the primary during checkpoints.
Q: "How do RSS secondaries behave compared to HDR?""	A: Remote Standalone Secondary (RSS) servers are always asynchronous. They can lag behind the primary but support log staging if `DELAY_APPLY=-1` is configured. Key notes: • `RSS_FLOW_CONTROL` defaults to 0 (which means "on"), allowing some lag tolerance (up to 12× log buffer size). • Use staging (`DELAY_APPLY=-1`) to increase log delivery reliability but note that RSS cannot participate in cluster transaction coordination. • If RSS is your only replication and it lags behind the primary during a failure, some transactions may be lost.
Q: "Can transactions survive a failover?"	A: Only transactions started on an updatable secondary can survive a failover. To enable transaction survival, set `FAILOVER_TX_TIMEOUT` to a value greater than 0. Transactions initiated on the primary will always roll back if the primary fails mid-commit.

Stay in Touch

If you have any suggestions for future newsletter topics, we welcome your feedback on the newsletter. Please reach out our product management team at **Emily.Taylor@Actian.com**.

The Authors of this Edition

Nick Johnson	HCL Informix Product Marketing Actian, a division of HCLSoftware
Scott Norris	HCL Informix Sales Engineering Actian, a division of HCLSoftware
Emily Taylor	HCL Informix Product Management Actian, a division of HCLSoftware
Karen Qualley	HCL Informix Product Management Actian, a division of HCLSoftware

About Actian

Actian empowers enterprises to confidently manage and govern data at scale. Organizations trust Actian data management and data intelligence solutions to streamline complex data environments and accelerate the delivery of AI-ready data. Designed to be flexible, Actian solutions integrate seamlessly and perform reliably across on-premises, cloud and hybrid environments. Learn more about Actian, the data division of **HCLSoftware**, at **www.actian.com**.

About DB Solve Technologies Consulting Sdn Bhd

DB Solve Technologies Consulting enables organizations to maximize the value of Informix investments through end-to-end technology and consulting services. Founded in 2004 and based in Kuala Lumpur, DB Solve delivers expert support across the region, helping businesses modernize data infrastructure and streamline operations. With a focus on flexibility and deep product expertise, DB Solve is the best partner equipped with Informix technologies to provide reliable, scalable solutions tailored to client needs. Learn more about DB Solve at **mydbsolve.com**.

DB Solve Technologies Consulting Sdn. Bhd. E509, East Wing, Metropolitan Square, No 2, Jalan PJU 8/1, Bandar Damansara Perdana, 47820 Petaling Jaya, Selangor, Malaysia
Sales@mydbsolve.com +603-7625 5743 (Phone) +603-7625 5744 (Fax) **Mydbsolve.com**

Actian Australia Pty Limited Level 8, 815 Pacific Highway, Chatswood, NSW 2067 +61 (0)42-1577-815 **Actian.com**

© 2025 Actian Corporation. Actian is a trademark of Actian Corporation and its subsidiaries. All other trademarks, trade names, service marks, and logos referenced herein belong to their respective companies.

All Rights Reserved. V2-2026-01

ACTIAN
a division of HCLSoftware