

IBM Informix Dynamic Server provides flexible business continuity for your agile business

Continuous availability solutions



Continuous availability custom-built for you

Continuous access to business information can rarely be achieved with a “one-size-fits-all” solution. Your business has its own unique structure and needs, and the design of your information systems and delivery should also be unique to meet them.

A business employing 10 people requires information availability on a much different scale than a corporation of thousands. At the same time, small businesses need to plan for growth, and design systems that are adaptable for the future. As a business grows and expands across geographies, its information infrastructure needs the agility to maneuver to meet new business opportunities without missing a beat. Whether information systems power the local grocery store or the largest mega-store in the world, business survival depends on delivering the right information to the right people at the right place and the right time.

*In some industries, downtime costs can amount to up to 16 percent of revenue.**

Finding the right mix of technology

Successful businesses weave together the right technologies for their organization to create a virtual “fabric of availability” for their information. In other words, their data systems ensure global availability of all business data and applications at all times to make business itself continuously available. Achieving this level of flexibility and availability often requires multiple approaches and layers of technology.

Finding the right product and combination of technologies is essential to implementing the most effective and cost-efficient solution.

IBM Informix Dynamic Server— Flexible, scalable solutions

While such a delicate mix may seem complicated, IBM Informix® Dynamic Server (IDS) has a long legacy of providing such flexible, customizable solutions. The ease of use and reliability

Continuous availability solutions

of IDS simplifies IT management and database administration. IDS offers technology that translates into critical business advantages such as:

- *Having better access to information to make insightful business decisions and improve processes.*
- *Gaining a more holistic and accurate picture of customers and their needs to improve service for bottom-line results.*
- *Eliminating time wasted searching for information or waiting for data transmissions, helping to improve employee productivity.*
- *Rapidly detecting information threats or fraud, reducing risk and streamlining compliance with government and industry regulations.*

IDS 11 offers significant new technology for businesses like yours to weave a flexible availability fabric that's just right for your needs. It begins with the basics—backup and restore utilities with the granularity and scalability to work with all or the smallest subset of your data. Additional texture is added with the ability to create and maintain near-line copies of the production environment, or to have an exact copy of the environment standing by, maintained in real-time and able to immediately take over if the primary can no longer function. The options explode with the ability to have not just one, but as many continually maintained copies as you'd like—either nearby or as far away as needed to build global availability. Finally, with the new Continuous Availability Feature, you can share a single network-mounted copy of data among a cluster of servers, broadening your reach while conserving expenses.

These availability solutions bring even greater value by allowing you to load-balance query workloads across any or all of them without interfering with or compromising their part of the availability fabric. You can weave these solutions together to protect yourself against any technology failure or disaster—be it manmade or natural. With IDS, these solutions work together seamlessly and easily.

Scalable and flexible backup and restore for protection against data loss

The first step in building a data availability solution is protecting your data against loss. The IDS backup and restore functionality permit you to backup at the level of granularity you need—from the entire environment to a single table (or portion thereof) using any number and type of backup devices. When restoring you can choose, to the second, the moment in time you want the system to reflect when the operation completes. Recent enhancements give you even more flexibility; you can restore subsets of a table's rows or just certain attributes of rows. You can also restore to different physical servers, running the same or different operating systems.

The IDS backup and restore functionality sits on top of the Dynamically Scalable Architecture (DSA) which, unlike other data servers on the market today, efficiently uses and maximizes the resources provided by the physical server hosting the IDS data server. This architecture enables you to access and use more sophisticated data server technology when it's needed—quickly, easily and reliably.




Multiple options = The ability to make the best choice for your business

As you look beyond just protecting your data, to an environment where data is replicated for business or availability needs, IDS offers an unmatched technology portfolio that's both wide and deep and can be as simple or as sophisticated as you want it to be. They include:

Disk mirroring

The next step for protecting your data, beyond backup and restore, is using disk mirroring. With other data servers you need to purchase and use additional hardware or software to create the mirrors. IDS doesn't; you can define mirror pairs within the database environment. When both devices are available, the data server will use both for I/O operations increasing throughput. If one fails, the other seamlessly assumes the complete I/O load. See Figure 1.


Figure 1.



Building near-line copies with Continuous Log Restore

The next step is to move beyond basic data protection to start creating a data availability fabric. Many businesses look into creating an exact copy of their environment. For cost, technological, or business reasons, the copy may not need to be immediately available. The IDS Continuous Log Restore functionality is the perfect response to this requirement. As logs containing database changes fill on the primary server, they are sent to secondary servers that can be local or geographically distant. Once there, the logs can either be applied or not, as business needs dictate. The amount of time required to bring these secondary servers to active mode is reduced from that required to execute a full backup and restore operation, to just the amount of time it takes to roll the latest log(s) forward. The frequency of log shipment is completely flexible to suit the needs of the application. This solution is ideal for customers using IDS Express Edition, who don't require more complex availability solutions, but need a more immediate availability option than a complete restore. See Figure 2.

Figure 2.




Continuous availability solutions

Creating a hot standby with High Availability Data Replication

After creating a copy of the environment, the next element to weave into the availability fabric is having the secondary server become a true, online copy of the production server with the ability to immediately take over should the primary become unavailable. The High Availability Data Replication (HDR) functionality, available in IDS since 1993, has kept businesses, small and large, running in the face of a worst-case disaster—losing their production server.

With HDR, you create a complete copy of the database environment (server and data), usually located within close proximity to the production server. The HDR secondary is literally kept in sync with the production server as transactions occur. It can immediately take over the processing without losing a single transaction should the production server become unavailable. Unlike other data servers, the HDR secondary is more than fail-over copy; it can also be used to support user operations. The HDR secondary supports query SQL operations without additional modification or configuration, while still performing its primary role of supporting the production server. This is an ideal choice to support report generation, online queries and other activities to distribute workload and improve performance. See Figure 3.


Figure 3.



Create global redundancy with Remote Secondary servers

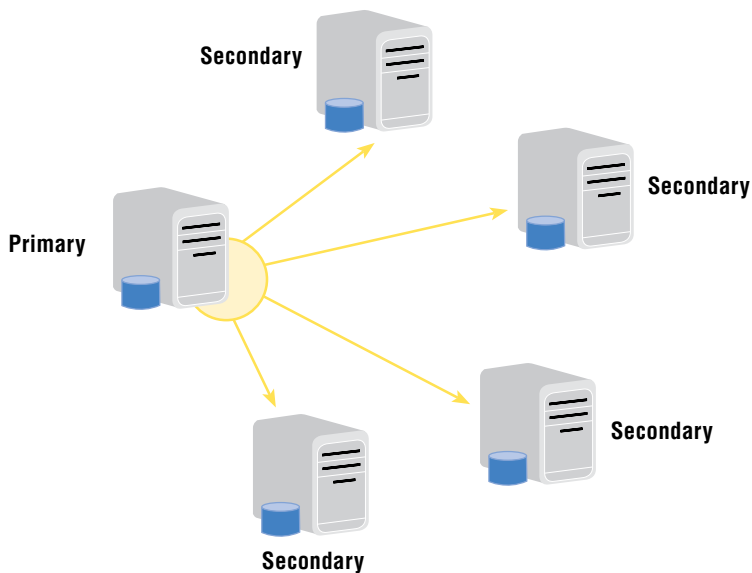
Business is never local and your availability solution shouldn't be either. Remote Secondary servers extend HDR secondary functionality to provide N+1 redundancy on a truly global scale.

Like an HDR secondary, Remote Secondary servers are independent and complete copies of the production server. The secondary databases are kept in sync with the primary as transactions occur. Unlike an HDR secondary, the communication from the primary to these copies uses a full duplex protocol allowing Remote Secondary servers to literally be located anywhere in the world—without affecting throughput and performance on the production server. You can use these servers to create a “bunker backup” in a geographically remote location from your primary production location. They are also intended to be a “backup of the backup” should the HDR secondary need to assume full processing responsibilities. Any one of these secondary servers can be promoted to HDR secondary, ensuring there is always a hot standby server.




Continuous availability solutions

All servers can be used for query operations, allowing you to distribute the application processing load. For example, in an Internet-based transaction system, your application can connect customers to the closest copy of the data anywhere in the world while they browse for details, pricing and other information. Once they decide to buy, the application executes the transaction against the production server located in your primary site.



Create enterprise-wide redundancy with the Continuous Availability Feature

With the ubiquity of network-mounted storage solutions, it is only natural that they should be used to enrich your availability fabric. The IDS Continuous Availability Feature, using Shared Disk Secondary servers, extends your investment in storage functionality, allowing additional IDS servers to access the same physical copy of data on disk. Since the IDS servers can be located anywhere on your network, they too can provide global redundancy while helping reduce your cost to implement—since they don't need their own set of disks. Shared Disk Secondary servers bring additional functionality beyond just reflecting what's on the disk; these servers also replicate the content of the production server's shared memory structures. This makes Shared Disk Secondary servers the best choice for hot site failover in the event of a production server failure. Servers from the IBM BladeCenter® family of products are ideally suited for hosting this type of IDS instances.



Blade Server



Continuous availability solutions

Weaving it all together using the options that best meet your needs

What kinds of failures are you most concerned about? What do you need to feel comfortable and ensure the availability of your data services in the event of any kind of failure? The availability fabric you use needs to adapt to you, not the other way around. IDS gives you the ability to seamlessly protect against server, disk, network and even location failures without losing a transaction. There are options for networks of any size and speed—and even no network at all. Each option is interoperable with all others and can be interchanged at will. Every one of these servers natively supports SQL query operations, adding even more value and ROI by balancing your transaction load. Finally they are completely compatible with Enterprise Replication, the IDS technology used to build global clusters of fully active/active servers.

Simplicity is the key

It is one thing to design an answer to a problem, but all too often the implementation is worse than what you're trying to solve. That is never the case with IDS, known throughout the IT industry as an easy data server to manage and maintain. All of the options available to you to craft your availability fabric can be instantiated in literally just a few minutes with, at worst, a single five-word command. Expensive services aren't required. Even during a downtime situation (planned or unplanned), you can modify the environment while still supporting normal operations. Once created, you are not locked into your initial design; you can change it at any time and reconfigure the instances without tearing the fabric apart.

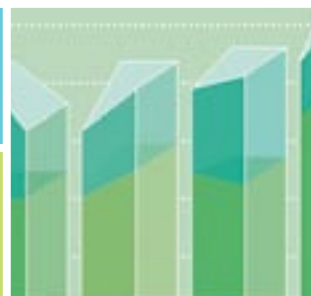
In fact, in almost all cases, the data server instances will automatically adjust on their own to the changes that are occurring without any direction from you. Finally, the technology is part of the data server's core, not an add-on requiring additional hardware and software. As such, it benefits from the IDS Dynamic Scalable Architecture, resulting in the fastest, most scalable and reliable solution available.

A complete solution from IBM

IDS is just one part of a complete portfolio of IBM Information Management solutions. IDS product adoption is growing faster than the data server market itself, due to the flexible, advanced technology it provides and the product's on-going legacy of success. IBM continues to make improvements in this technology that will serve our customers into the future. Finally, IBM and IBM Business Partners offer a spectrum of implementation and customization services, rounding out a full portfolio of hardware, software and services that can keep your critical business information—and your business itself—continuously available.

For more information

To learn more about IBM Informix Dynamic Server, visit: ibm.com/informix/ids.





© Copyright IBM Corporation 2007

IBM Software Group
Route 100
Somers, NY 10589
U.S.A.

Produced in the United States of America
07-07
All Rights Reserved

BladeCenter, IBM, the IBM logo and Informix
are trademarks of International Business Machines
Corporation in the United States, other countries,
or both.

Other company, product or service names
may be trademarks or service marks of others.

*The Costs of Enterprise Downtime:
North American Vertical Markets 2005
Information Research

TAKE BACK CONTROL WITH **Information Management**