U32

e-Business Solutions Using DB2 UDB Data Links Technology on Unix & NT

Nagraj Alur

Orlando, Florida October 1 - 5, 2001

Topics

- e-business application characteristics
- Data Links technology
- Web Asset Integrity Solution demo -- LinkIntegrity+
- Conclusions
- Supporting foils

Please visit the Data Links website
www.ibm.com/software/data/db2/datalinks
for all kinds of collateral including
presentations, papers and brochures

"Red Book" : DB2 Data Links: Unstructured Data Support (SG24-6280)

e-business Application Characteristics

Business Critical Processes

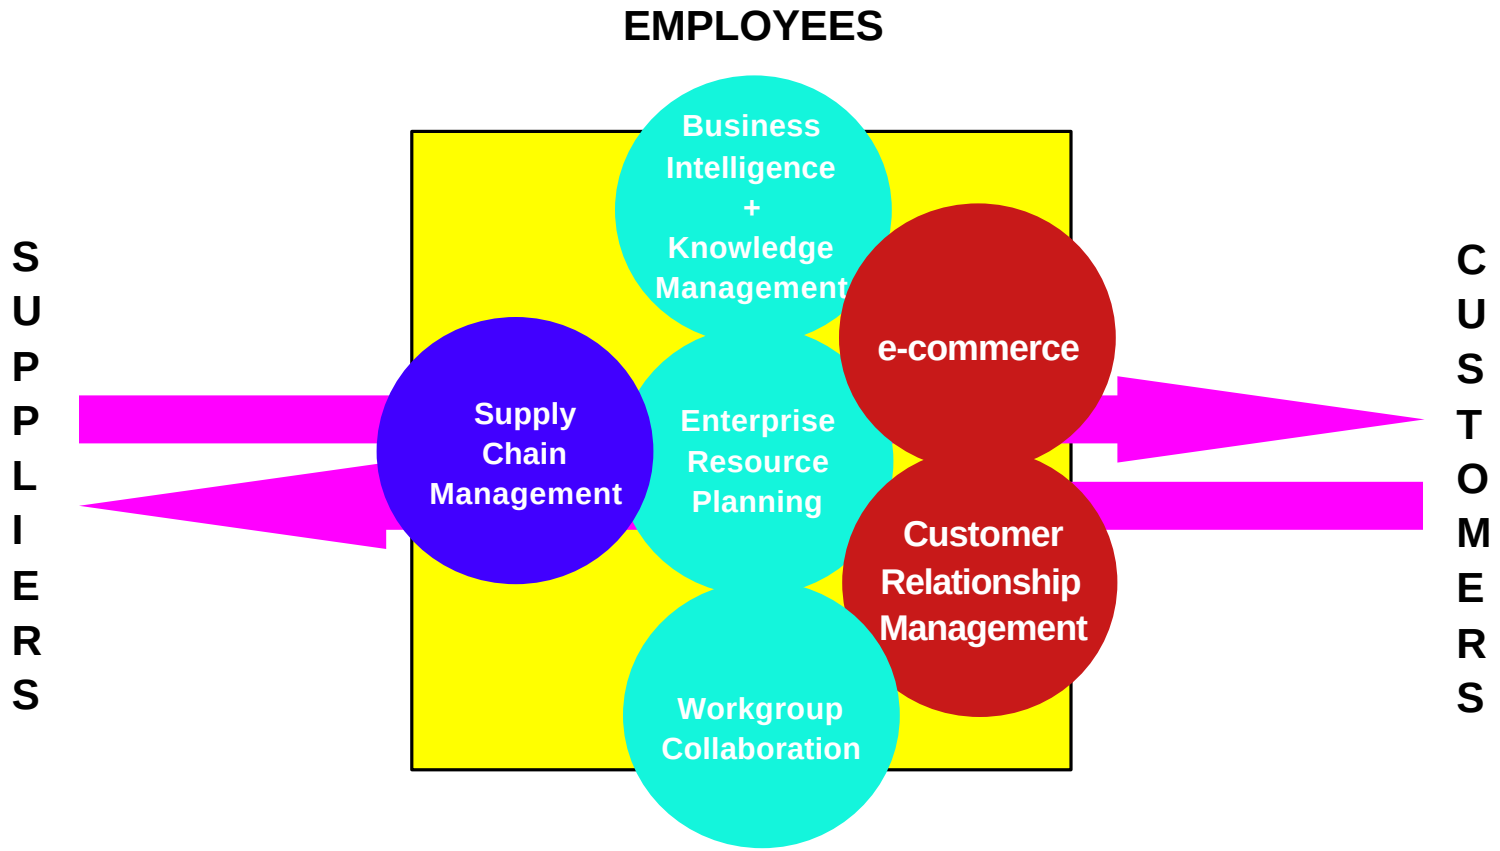

An e-business organization connects critical business systems directly to employees, customers, suppliers and distributors, via the internet, intranets and extranets to gain a competitive advantage

e-business Challenges

- Exacerbated infrastructure demands
- Technology demands
- "Internet time"
- Application demands

Content must be timely, accurate, consistent, and secure

e-business Demands on a Database

Why a DBMS

- Shared online access to information by hundreds and thousands of users
- Integrity
- Security
- Performance
- Optimal use of available resources

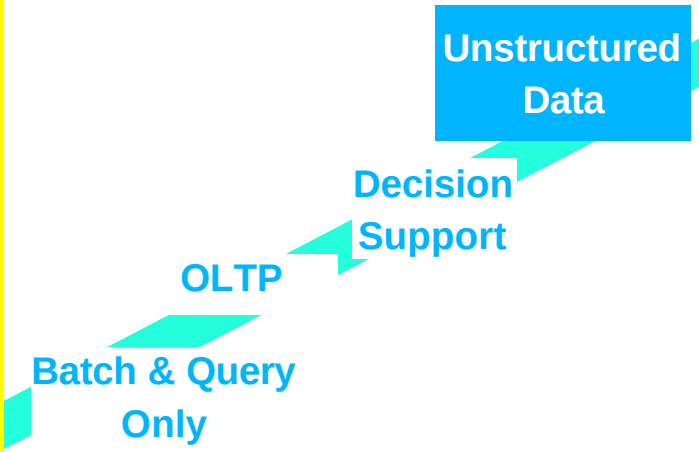

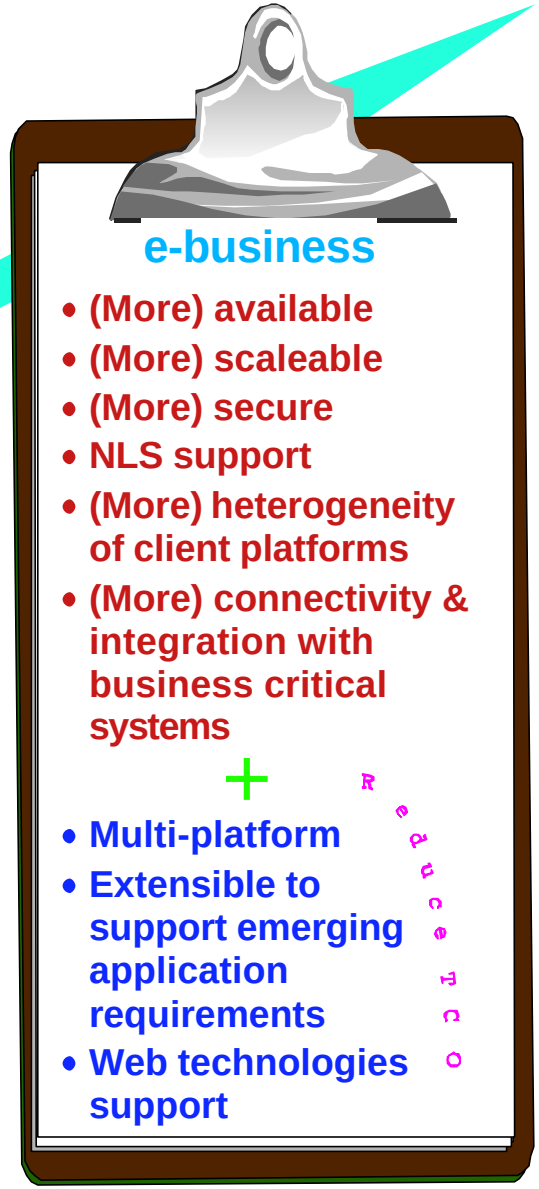

IBM's IMS is the first DBMS in mid 1960s
IBM invented the Relational Model in 1969
IBM DBMSs in 100% of Fortune 100 mission-critical applications
IBM leading visionary in final Gartner DBMS Magic Quadrant

Data Links Technology

Problem Domain Addressed

- 70%-95% (or more) of data resides in file systems as compared to databases
- Legacy applications using file systems abound, and new ones being developed involving unstructured data
- File system limitations
- New e-business applications and certain existing and emerging applications need to integrate database systems with existing and new file systems
- Storing files in BLOBs gives it DBMS capabilities

DB2 UDB Data Links

- Extends the following database management capabilities to file data in file systems
  - ➡ Referential Integrity
  - ➡ Value-based security
  - ➡ Coordinated backup & recovery
  - ➡ Replication

  All with transaction semantics
- Allows management of files **AS IF** it were stored in the database
- Retains file system performance, APIs and "look and feel"
- Leaves data where it is and manages it

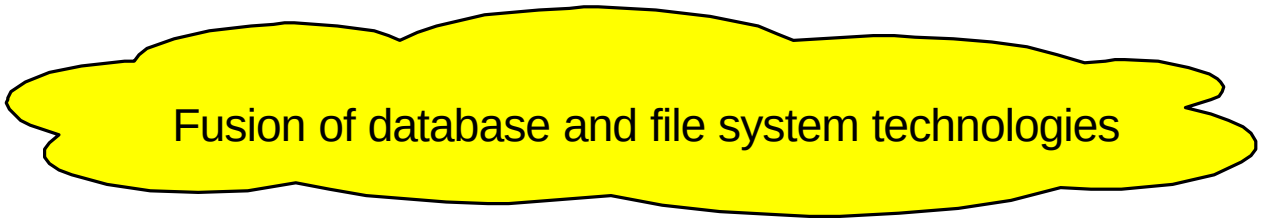

Data Links Programming Model

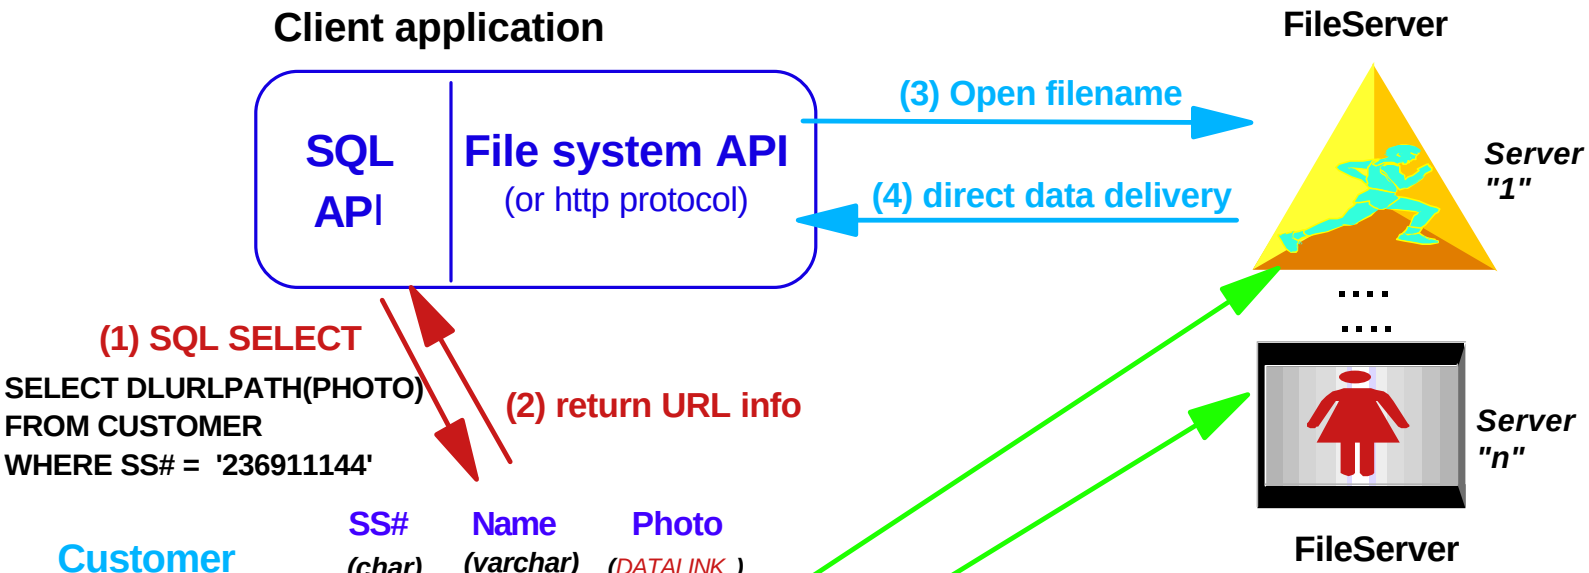


```
CREATE TABLE CUSTOMER
      (SS#      CHAR(8),
       NAME     VARCHAR (60),
       PHOTO    DATALINK LINKTYPE URL FILE LINK CONTROL
                INTEGRITY ALL READ PERMISSION  FS
                WRITE PERMISSION BLOCKED RECOVERY YES
                ON UNLINK RESTORE)
```

1. Leave the files as is
2. Establish the metadata for a set of files and define it as columns in a DB2 UDB table
3. Define a DATALINK (ISO/ANSI standard) column in the table that would contain the logical reference to the file in URL notation
4. Populate the table
5. Use SQL to search the metadata and find the files of interest -- get the URL of the file from the DATALINK column
6. Access the file using the native file system APIs

Data Links Architecture

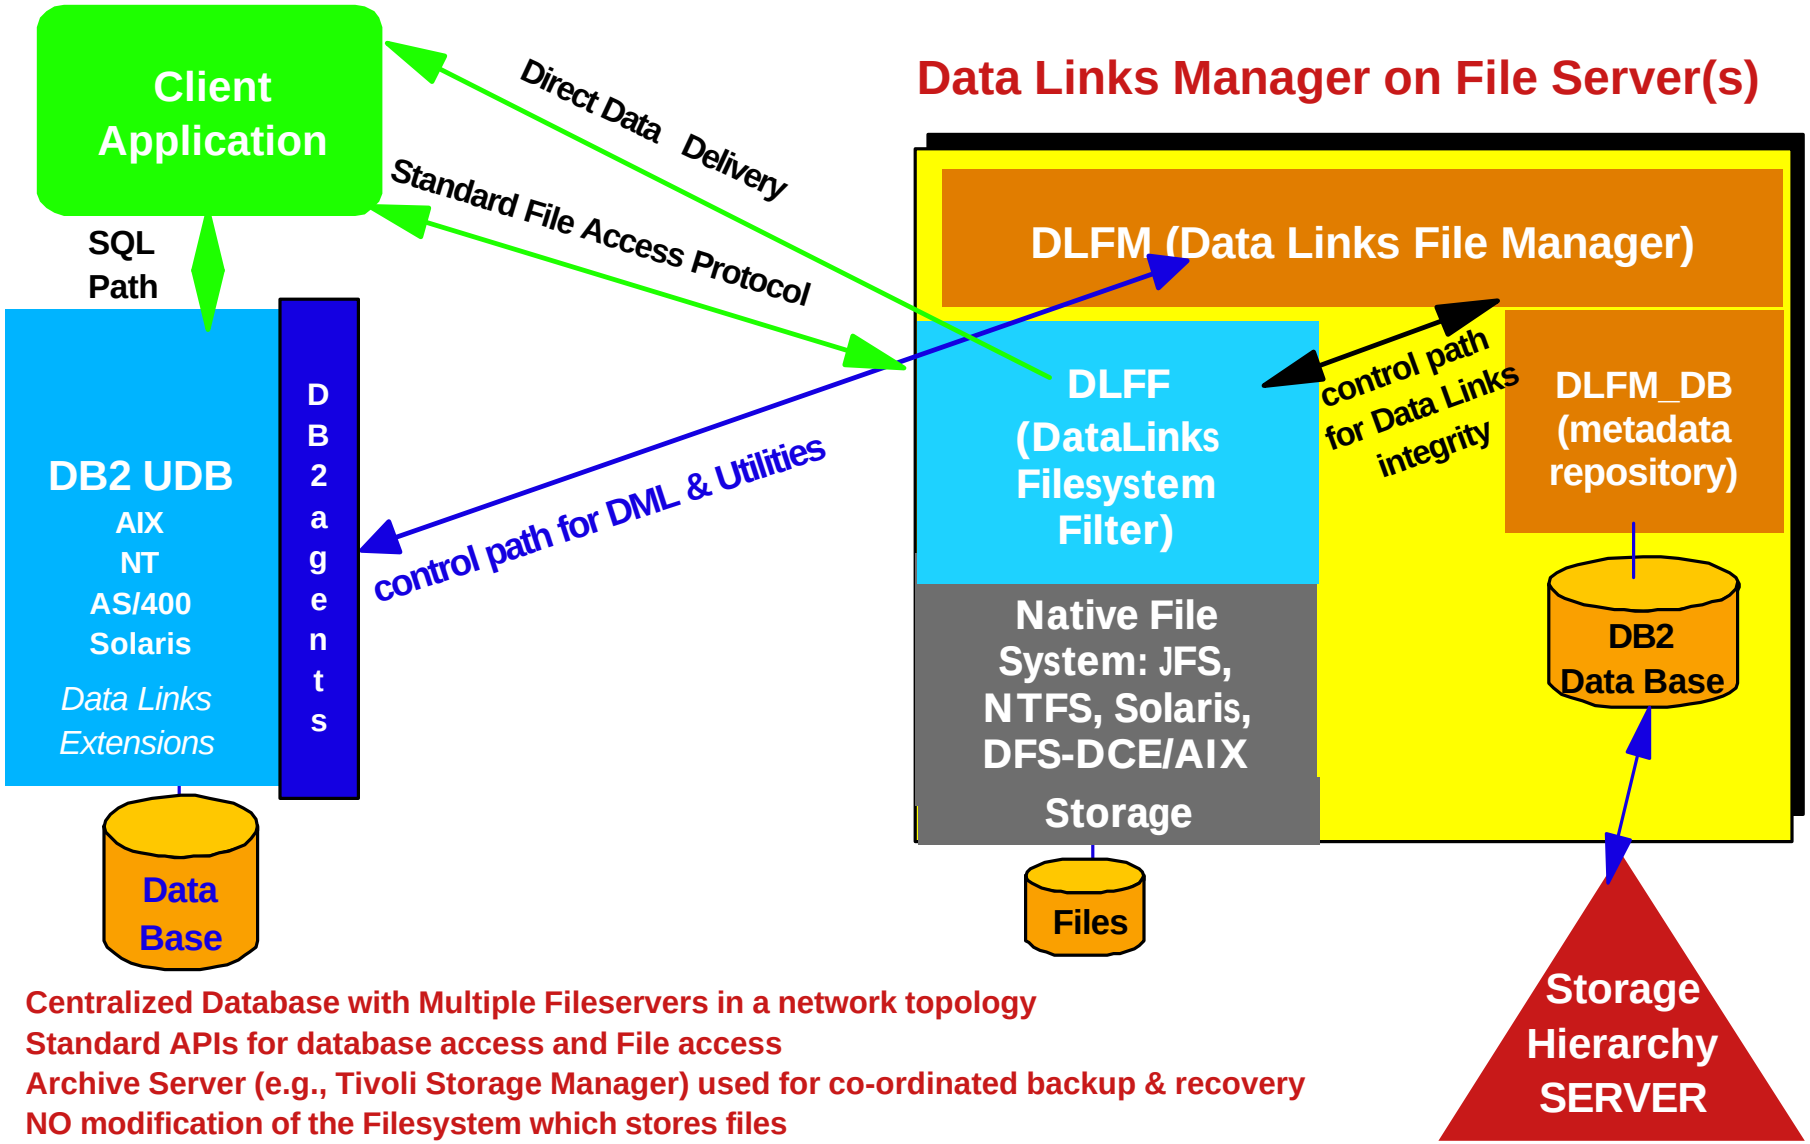

Centralized Database with Multiple Fileservers in a network topology
Standard APIs for database access and File access
Archive Server (e.g., Tivoli Storage Manager) used for co-ordinated backup & recovery
NO modification of the Filesystem which stores files

Data Links Process Model

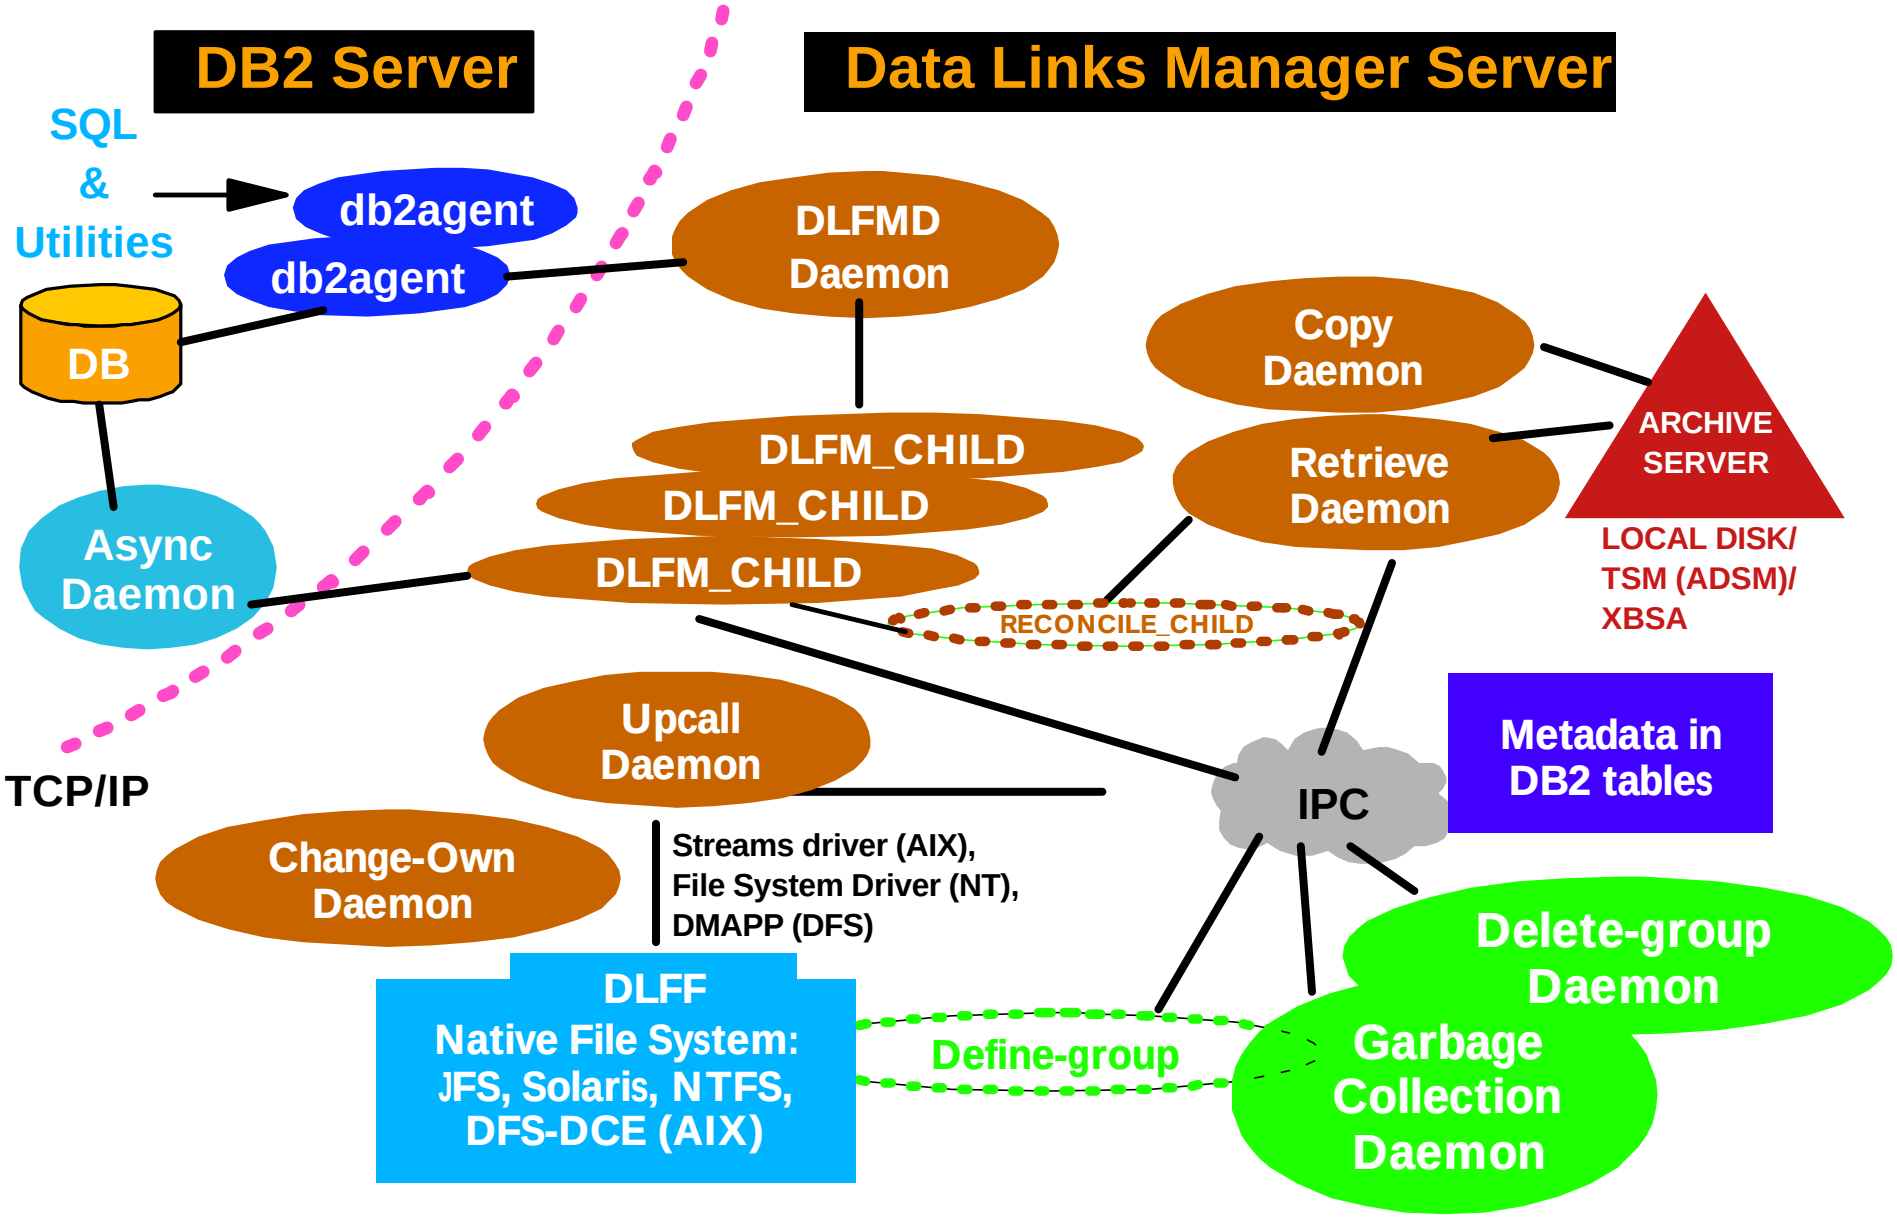

Referential Integrity

- Comes into play when rows are INSERTed, UPDATEd and DELETEd

```
INSERT INTO CUSTOMER VALUES ('123456789','BLANKETYBLANK',
                DLVALUE('HTTP://WWW.ALMADEN.IBM.COM/CDRIVE/BBPIC.GIF'))
UPDATE CUSTOMER
    SET PHOTO = DLVALUE('HTTP://WWW.ALMADEN.IBM.COM/CDRIVE/BB.GIF')
    WHERE SS# = '123456789'
DELETE FROM CUSTOMER  WHERE SS# = '123456789'
```

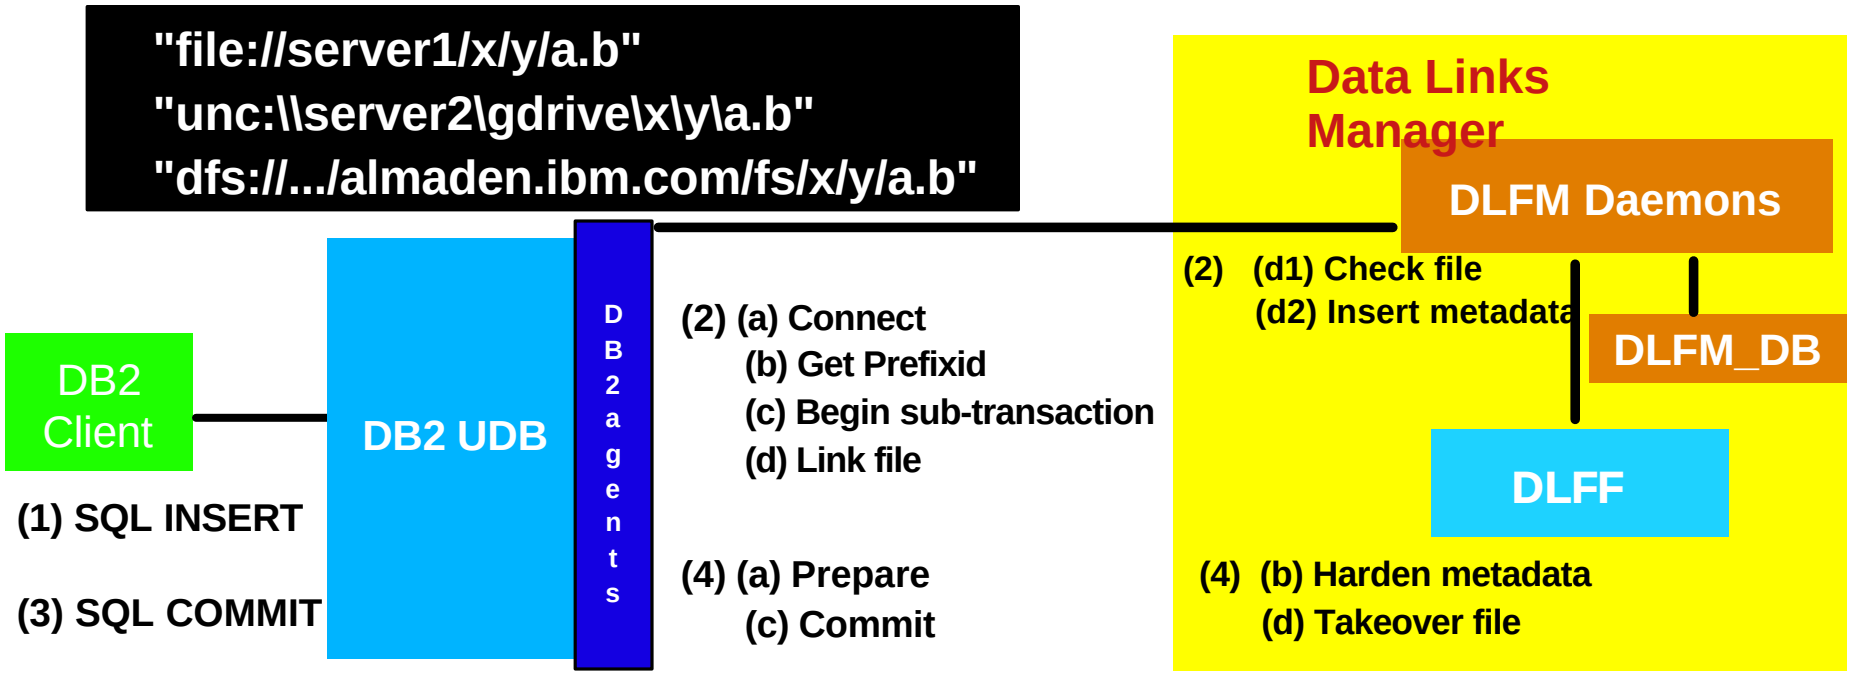

Coordinated Backup and Recovery

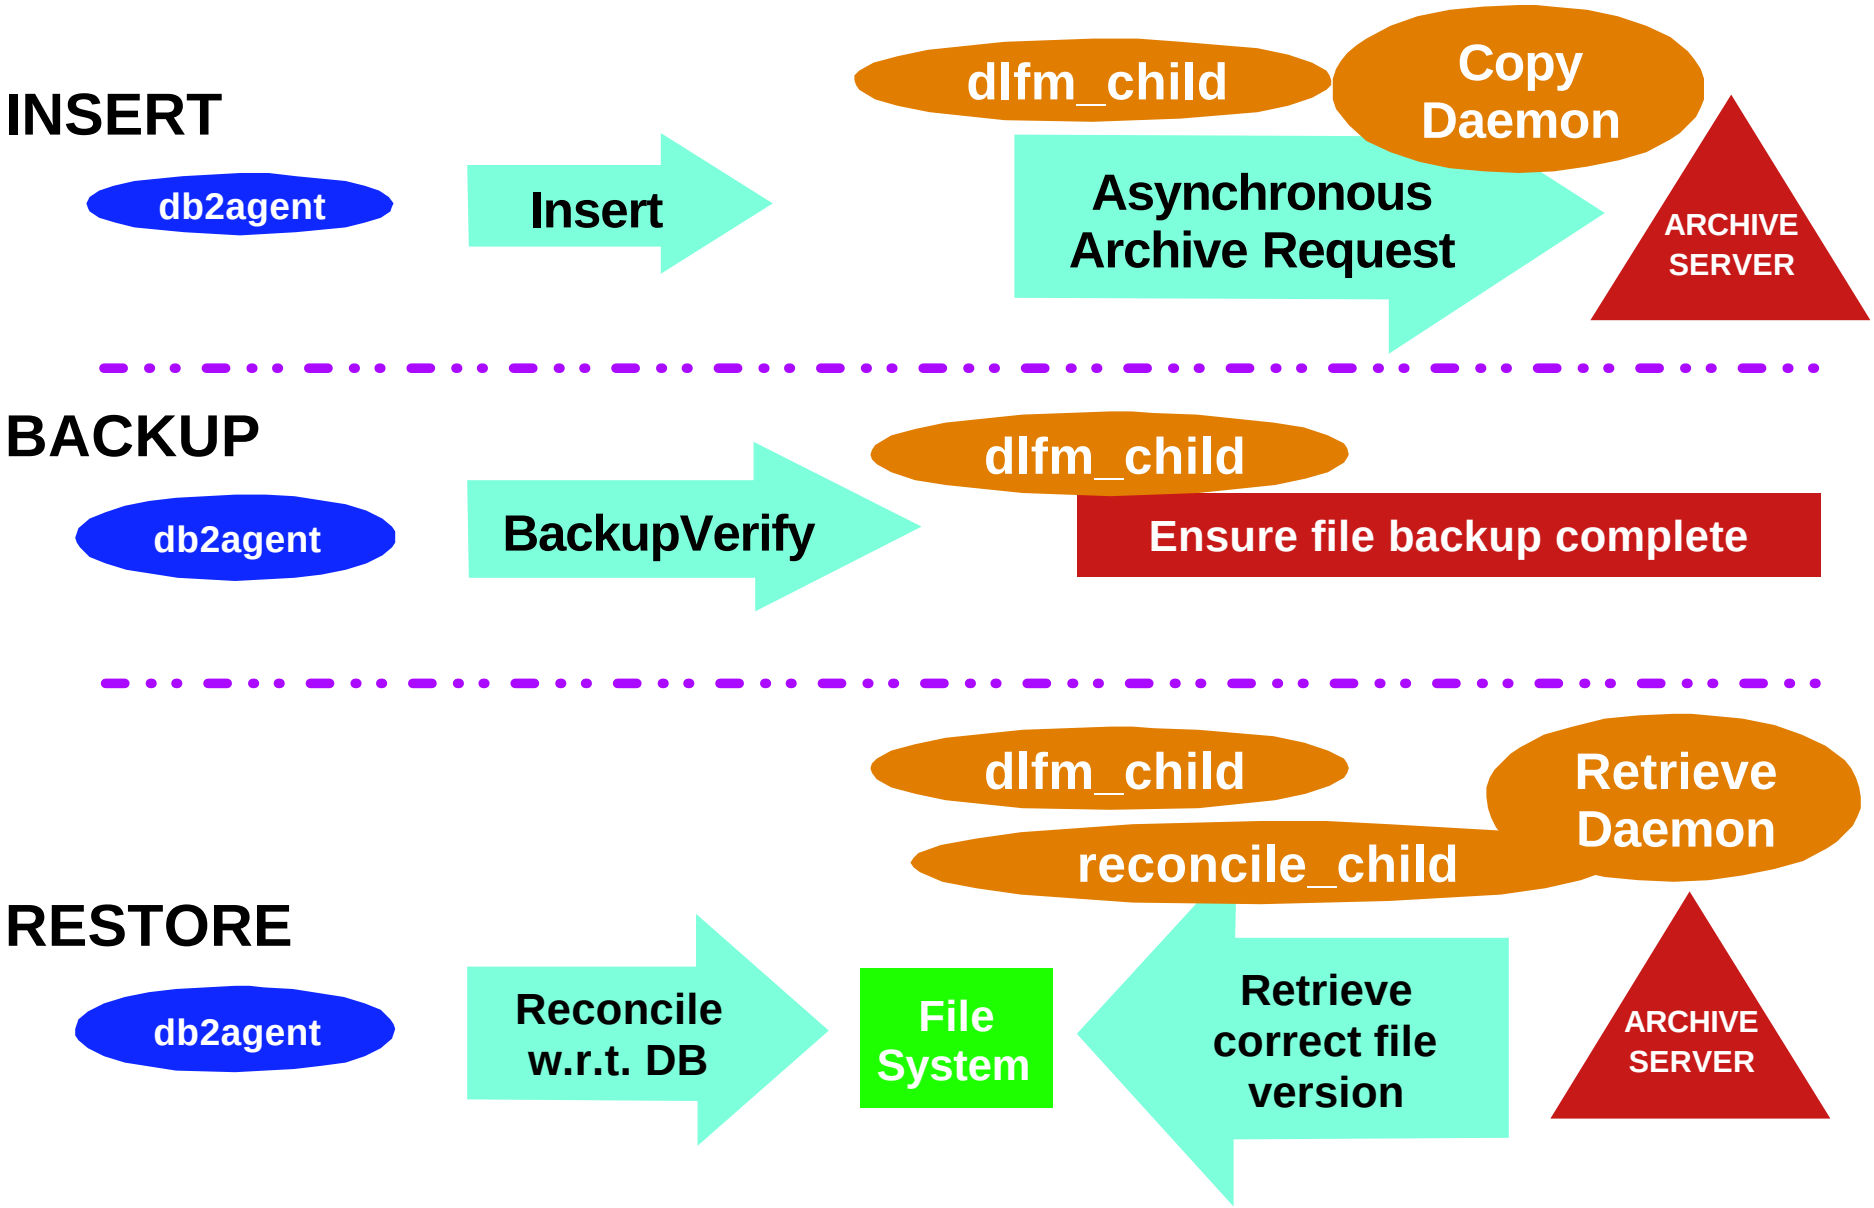

Access Security

- Access security type depends upon the READ PERMISSION attribute chosen when the DATALINK column is defined
  - ➡ READ PERMISSION FS specifies that existing filesystem permissions are to be honored
  - ➡ READ PERMISSION DB specifies that a database generated access token must be presented to DLFF before file access can be granted to the user
    - ✓ File ownership changed to database
    - ✓ Access token (25 or 30 characters in length) generated on query & embedded in the filename
    - ✓ Token validated by DLFF during filesystem open()
- Token generation and validation

Example: /videos/french.mpg ==> /videos/04E2_CS7Fo___biV4fhZ_0UM;french.mpg

  - ➡ Shared secret between DB2 and DLM (algorithm and key)
  - ➡ Two levels of security
    - ✓ MAC_0: encryption based on filename
    - ✓ MAC_1: encryption based on full path name

Access Performance

- DLFF is NOT in the read/write path
  - ➡ File access performance is not impacted compared to the native file system

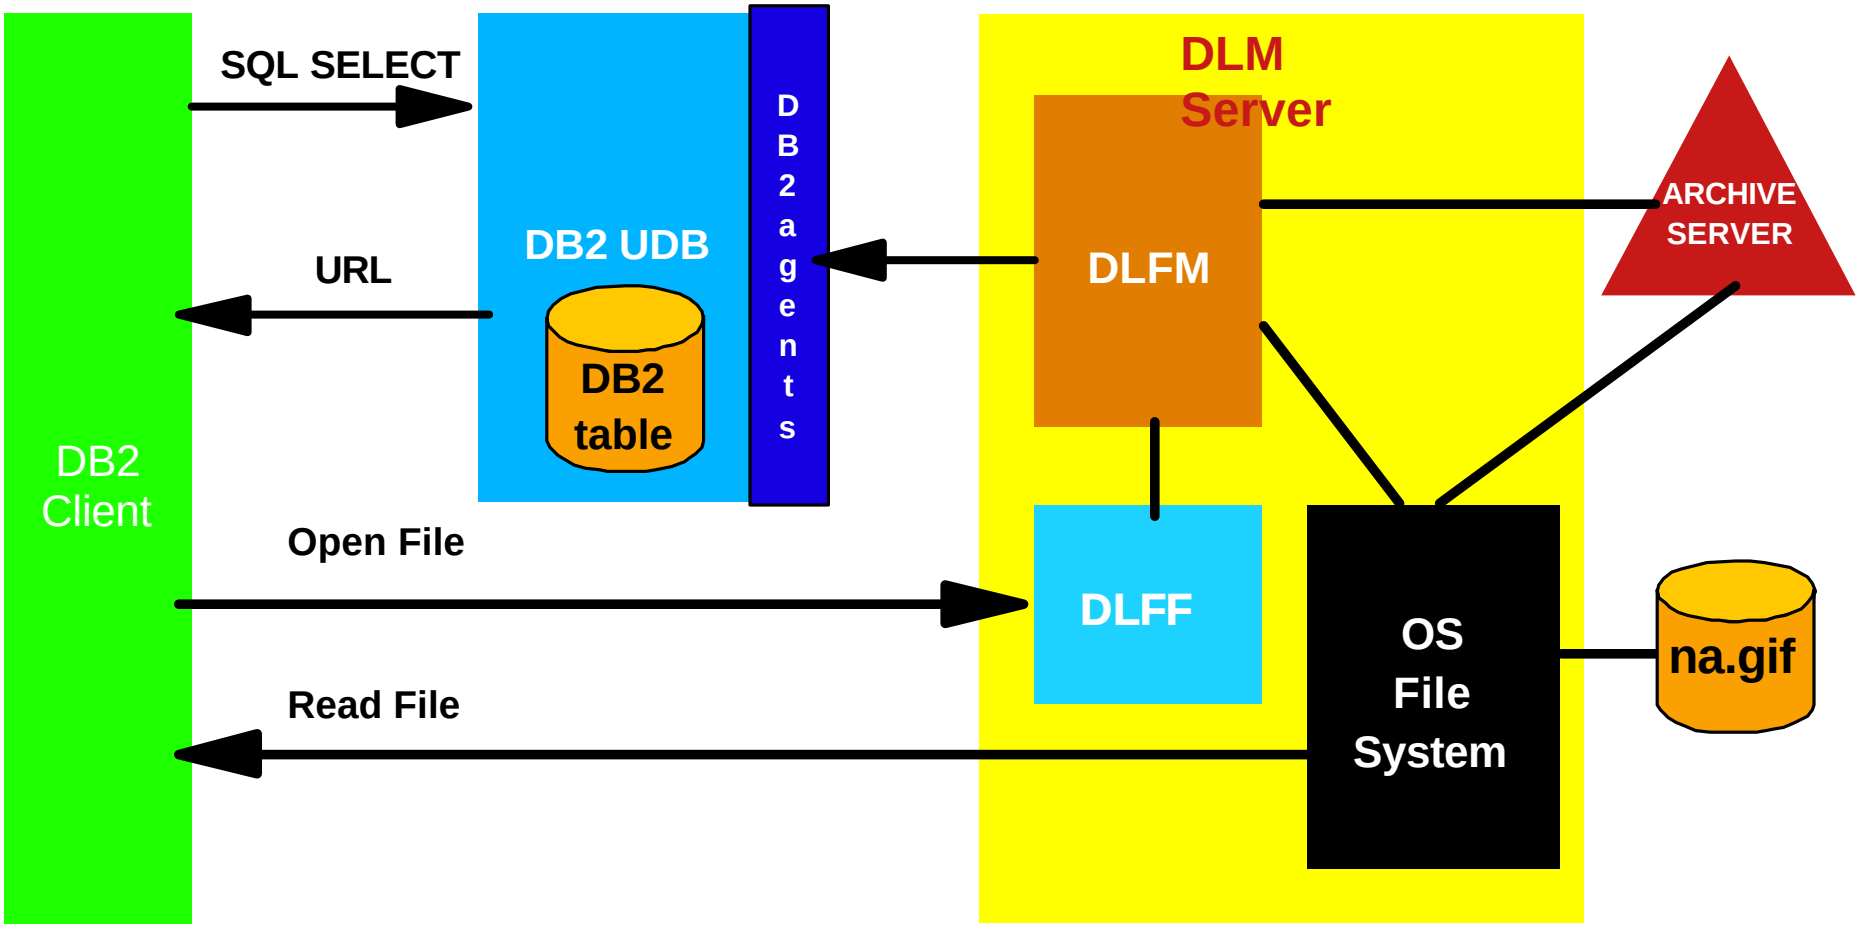

Utilities...

- All input formats to LOAD and IMPORT supported for tables containing DATALINK columns
  - ➡ DATALINK SPECIFICATION provides flexibility for transforming DATALINK values in data files
  - ➡ SAVECOUNT in LOAD causes consistency points for files linked in DLMs
  - ➡ Exceptions for DATALINK column(s) reported in exception table
  - ➡ LOAD COPY and LOAD REPLACE options not supported
- EXPORT
  - ➡ DB2 EXPORT command generates control file (TAR or ZIP) containing file references
  - ➡ dlfm_export generates a TAR (ZIP on NT) file based on control file
- IMPORT
  - ➡ dlfm-import uses the control file and TAR (ZIP on NT) file to materialize files prior to running IMPORT on DB2
- See DB2 Data Movements Guide for details

Utilities

- RECONCILE utility keeps DB2 table & DLFM meta data in sync
  - ➡ Operates at the table level
  - ➡ Table is scanned and a list of files (+version identifier) sent to DLM
  - ➡ DLM verifies and if required retrieves file from archive server
  - ➡ Unresolved references are recorded in an exception table
- db2_reconid_aid provides ability to run RECONCILE on all tables with DATALINK column(s)
- Fast RECONCILE at database without restore rollforward (internal)
  - ➡ Sync is done based on LSN

Data Links Replication

- Performance
  - ➡ Reduce network traffic by moving data closer to the application
  - ➡ Load balancing by providing multiple copies of a system image
- Availability
  - ➡ Standby or Failover capability in case of system failure
- Security
  - ➡ Isolate sensitive data

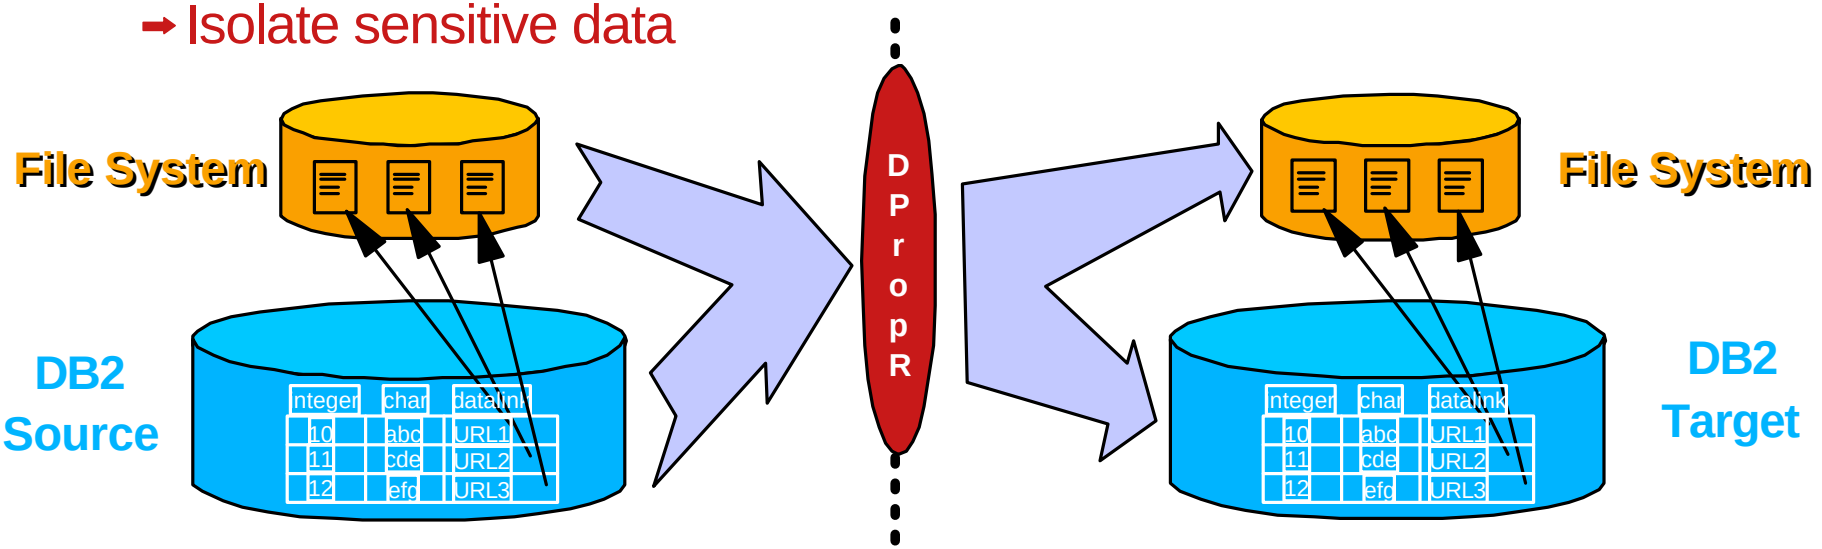

Both database data and external referenced files will be replicated together in an automatic and consistent way

DPROPR & Data Links

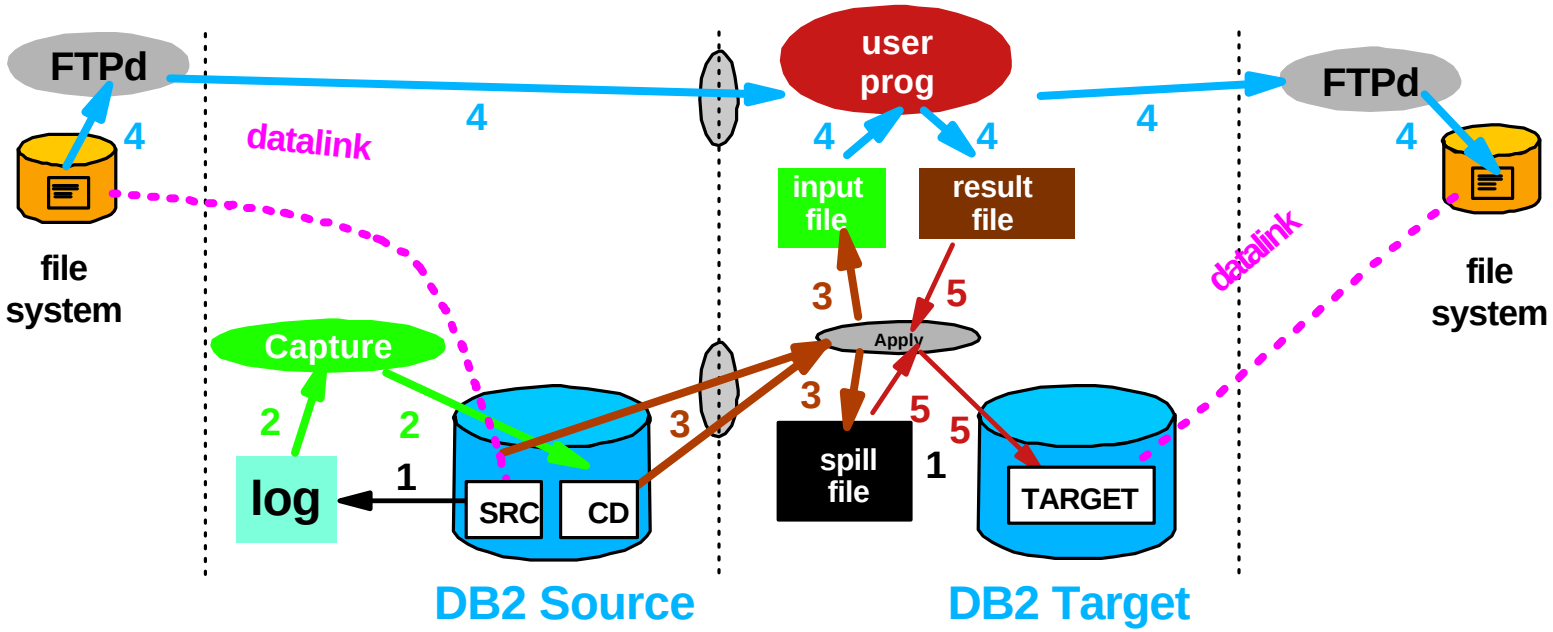

1. **Datalink changes are recorded in the database log.**
2. **Capture reads the database log and stores the changes to the CD table.**
3. **Apply copies the change data from the CD table to a spill file. At the same time, stores the Datalink file references in a separate file (input file).**
4. **User exit program maps the file references and copies them from the source file system to the target file system through FTP, and records file references in the "result" file**
5. **Apply propagates the metadata and the new DATALINK file reference from the "result file" to the target table**

Data Links Applications...

- e.Commerce
  - ➡ Product catalogs, price lists, brochures, thumbnail and full images, video, etc.
    - ✓ Integrity of file content
    - ✓ Integrity of file reference
- Supply Chain Management (SCM)
  - ➡ Common in automotive and aerospace industry for engineering designs
    - ✓ Large automotive manufacturer outsources 70% of a vehicle design
      - ◆ requires content sharing between different enterprises
      - ◆ needs replication of both the engineering drawings (files) and the metadata (database)
  - ➡ Customer support document system
    - ✓ Large airplane manufacturer needs to deliver maintenance documents in common format to relevant airlines

Data Links Applications

- Customer Relationship Management (CRM)
  - ➡ Holistic view of customer touchpoint interactions -- voice, e-mail, fax, web, database, etc.
    - ✓ Integrity of file content
    - ✓ Integrity of file reference
- ERP
  - ➡ Patient Information System where information is exchanged between hospitals and clinic -- Xrays, ECG charts, Doctor comments, medical history, etc.
  - ➡ Catalog distribution system -- catalogs include metadata & file data
  - ➡ Automotive insurance (vehicle damage pictures, claim forms, etc.)
- CAD/CAM
  - ➡ Engineering drawings
- Asset & Configuration Management
  - ➡ Content Management
    - ✓ Integrated Document Management
    - ✓ Media Access Management
    - ✓ Web Asset Management

BLOBs versus Data Links

- Storing files in BLOBs gives it DBMS capabilities
- DataLinks allows files to remain as is, while extending DBMS capabilities to them
- Use DataLinks when
  - ➡ Performance & scalability are of concern
  - ➡ Coexistence with existing and emerging applications that use the file system natively is required
- BLOBs appropriate when above issues not a concern

DB2 UDB is unique in the industry in offering the customer the choice to either implement BLOBs or Data Links

Lets the customer decide which option is most appropriate for their particular application requirement
(Single application may adopt both technologies)

Web Asset Integrity Solution Demo

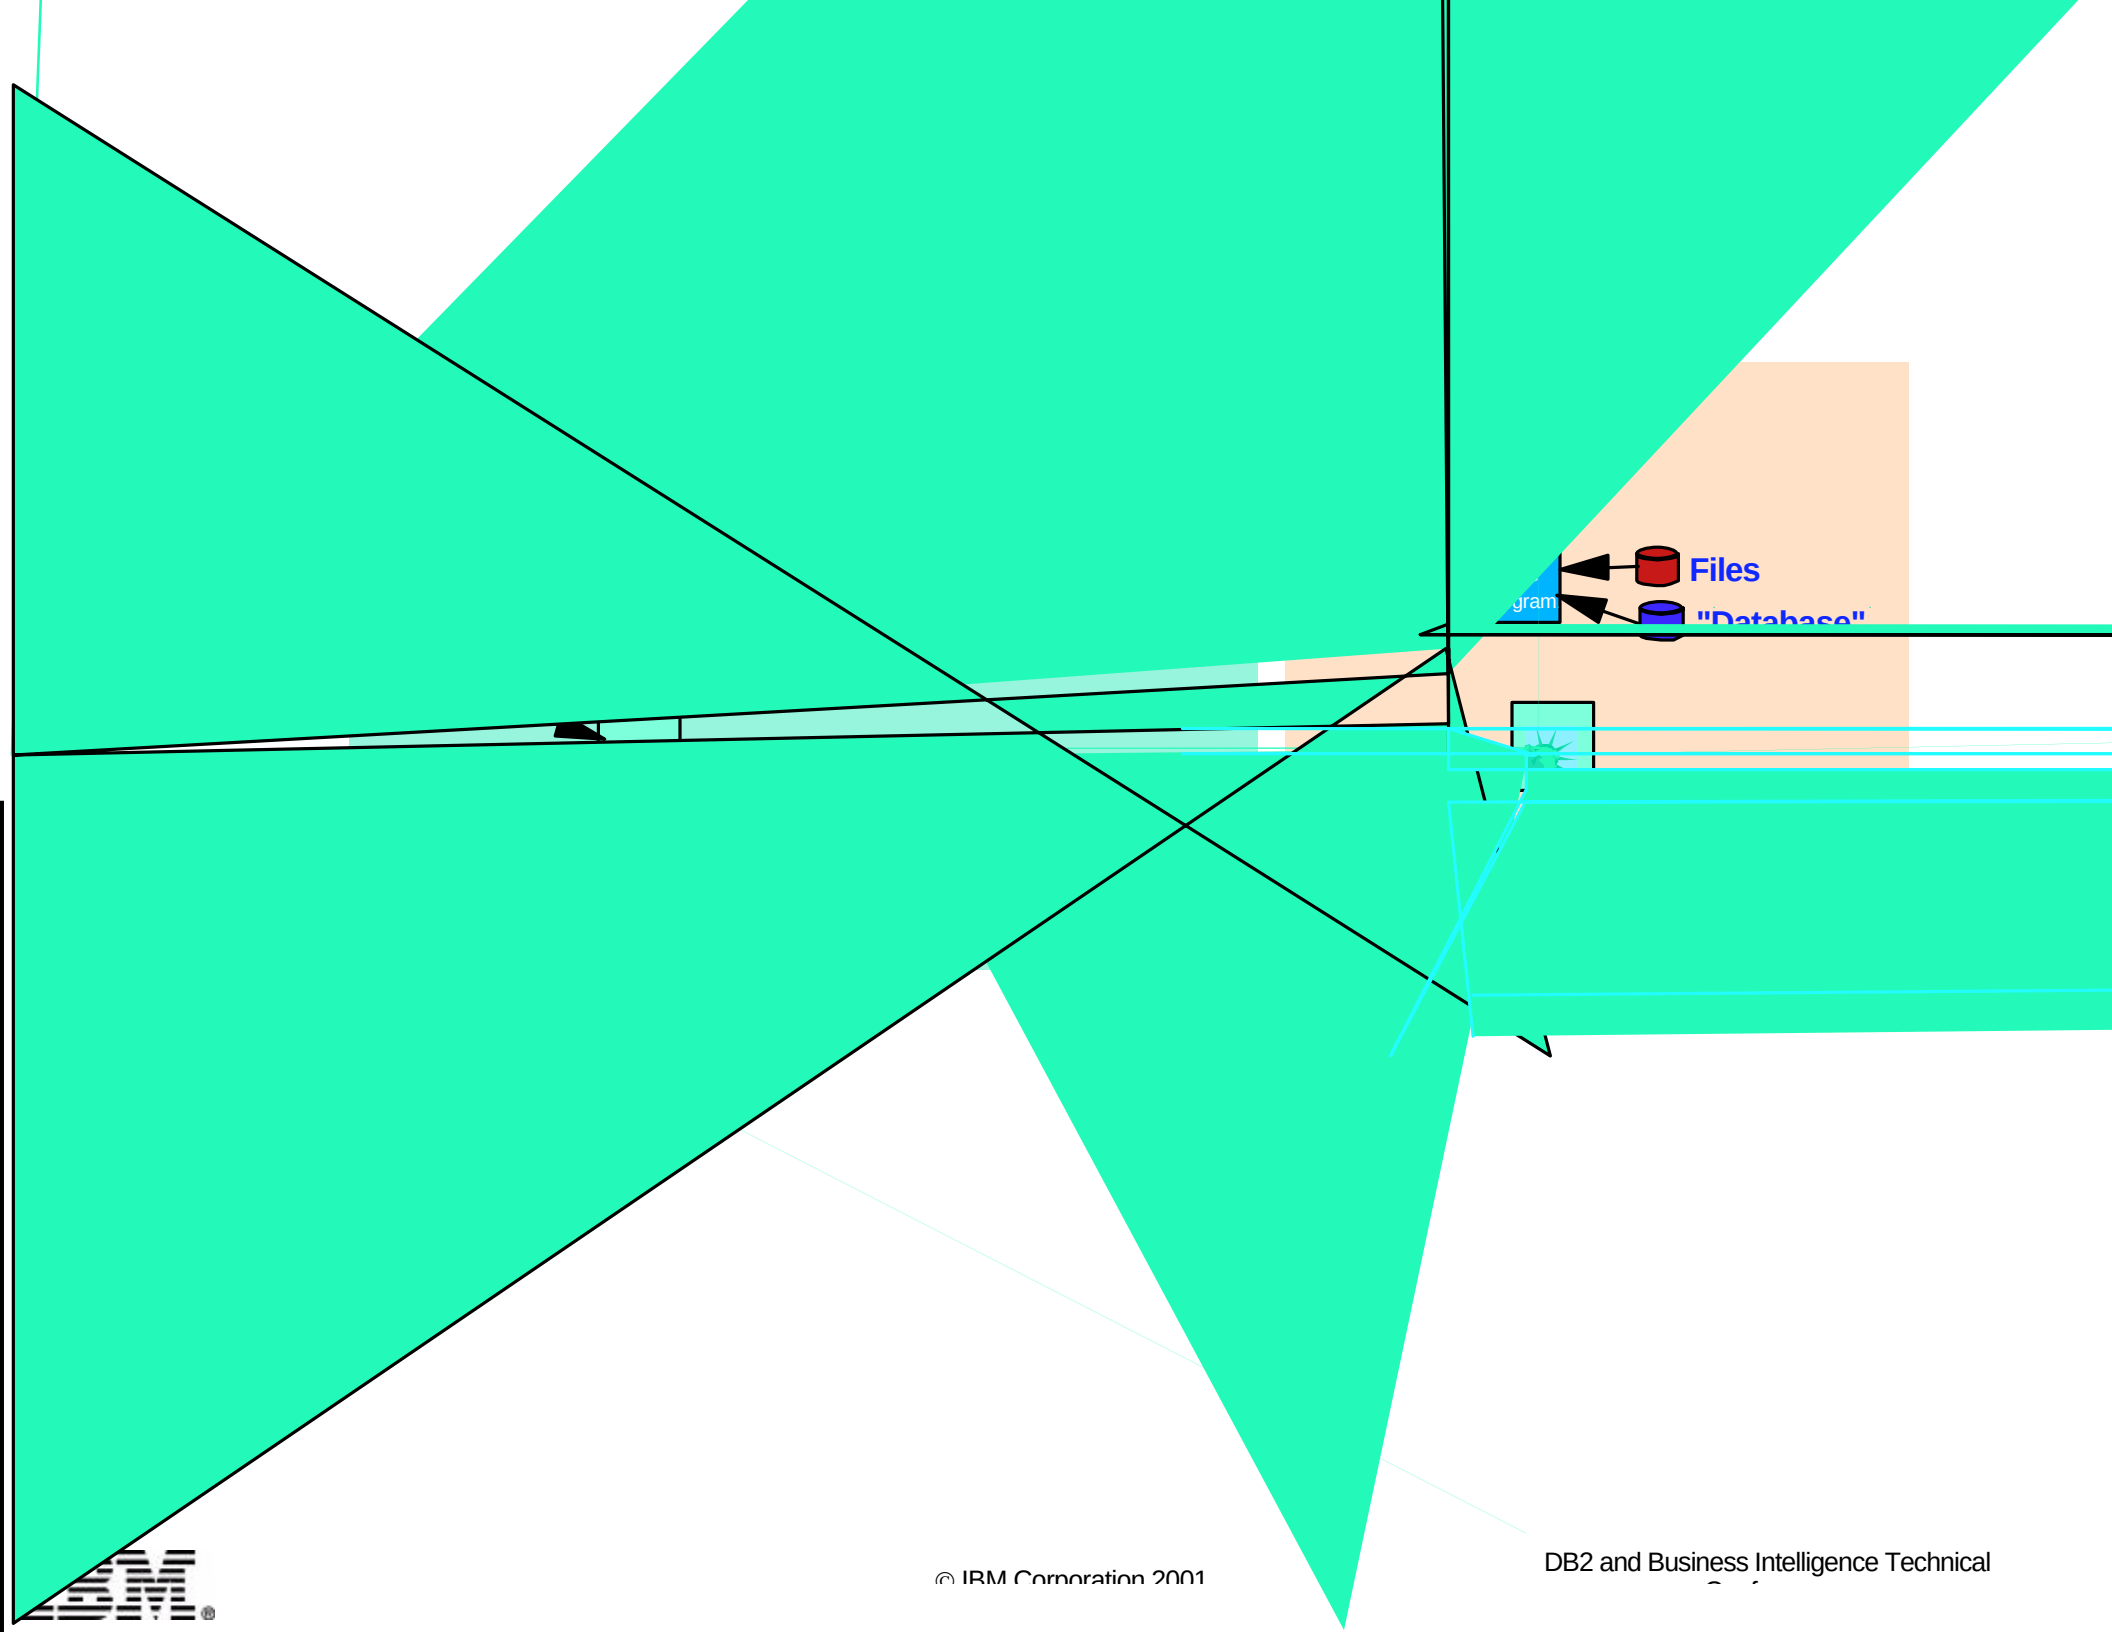

LinkIntegrity+ Overview

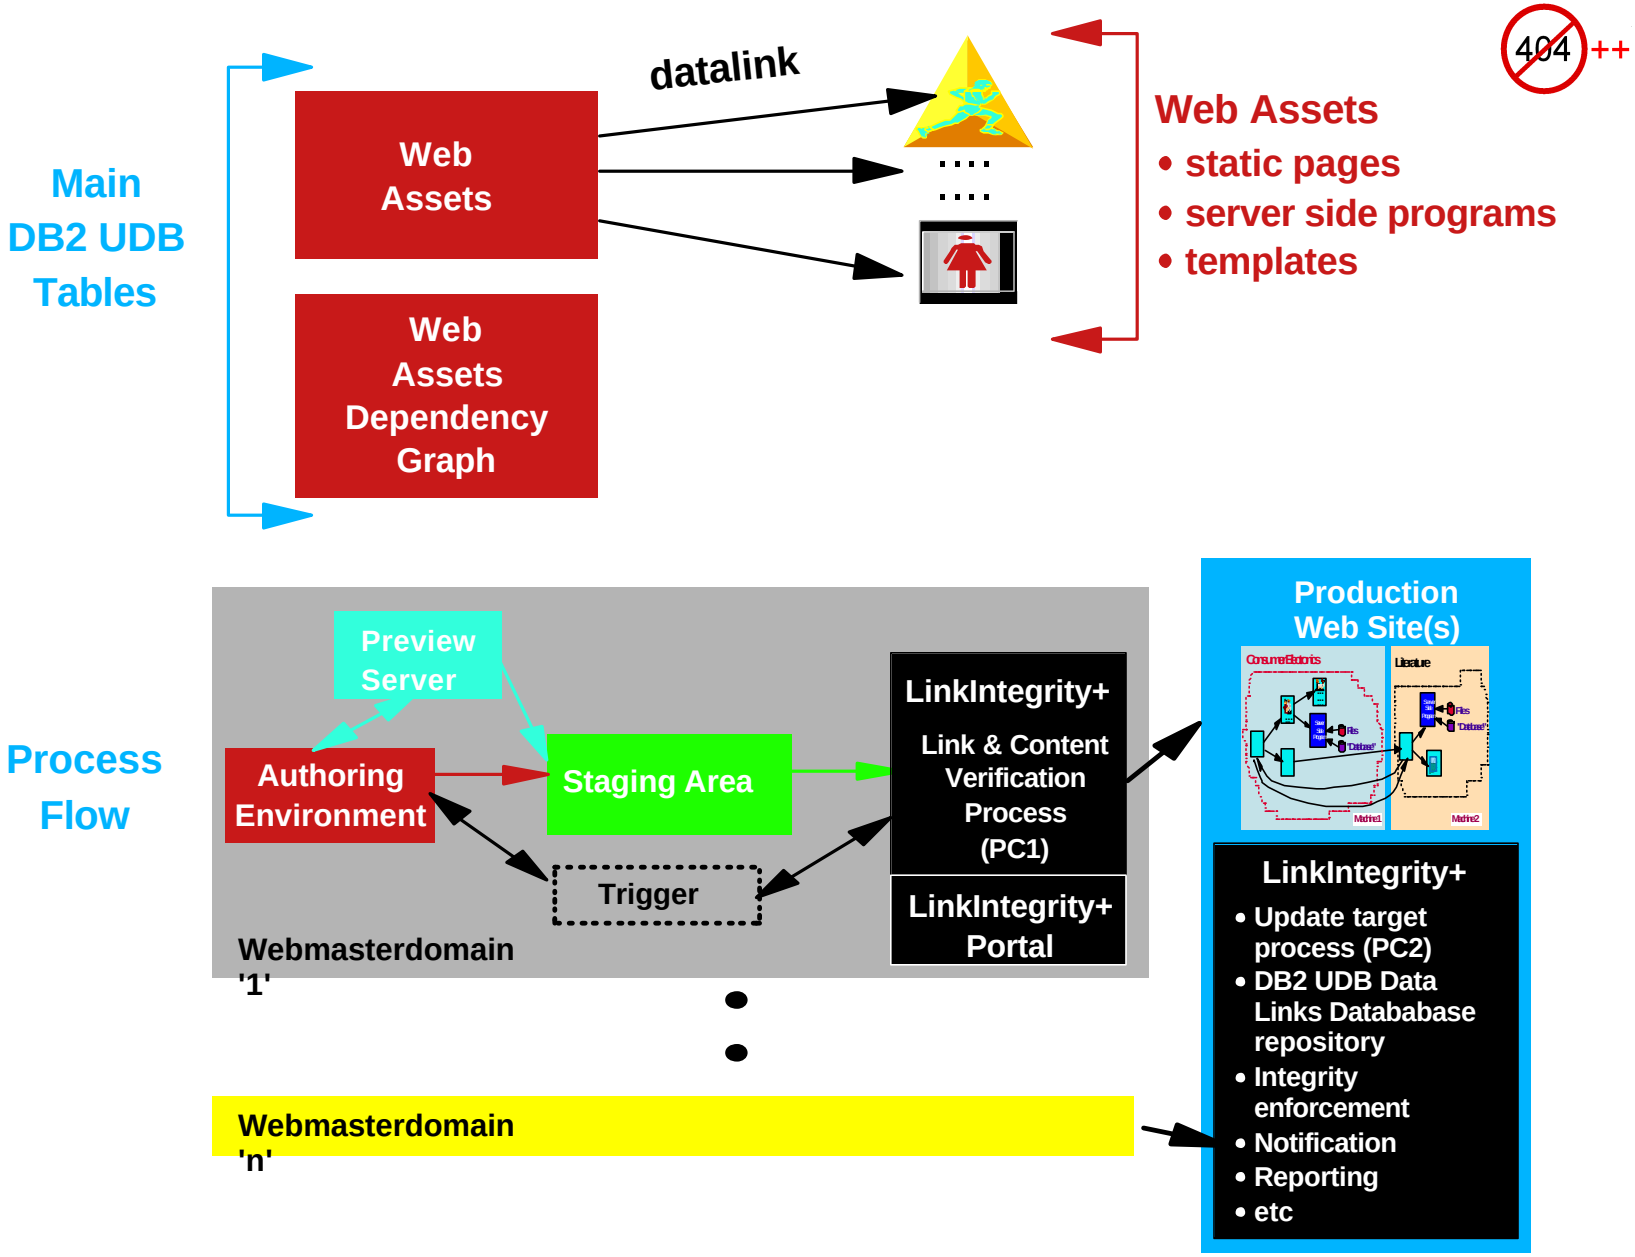

Conclusions

- Explosive growth in data stored in files critical to e-business
- e-business applications
  - ➡ Integrate structured and unstructured information from diverse sources
  - ➡ Co-exist with existing and emerging file system based applications
  - ➡ Demand mission critical capabilities of scalability, availability, security and integrity
- Data Links addresses e-business application demands by
  - ➡ Extending to file systems, the umbrella of mission-critical RDBMS capabilities of referential integrity, value-based security, transaction consistency and co-ordinated backup and recovery
  - ➡ Supporting coordinated database & file replication for load balancing, high availability and B2B requirements
  - ➡ Providing a scaleable multi-platform solution

Supporting Foils

DataLinks Terminology...

- Access Token

DataLinks Terminology

- DLFM
  - Data Links File Manager
  - Sub-component of DLM dealing with file metadata processing, user-process level daemons
  - Interacts with DLFF and DB2
- DLM
  - Data Link Manager
  - DataLinks application that is installed on the file server
- DPropR
  - IBM's replication technology
- Prefix
  - The mount point of the DLFF monitored filesystem

DATATYPE GRAMMAR

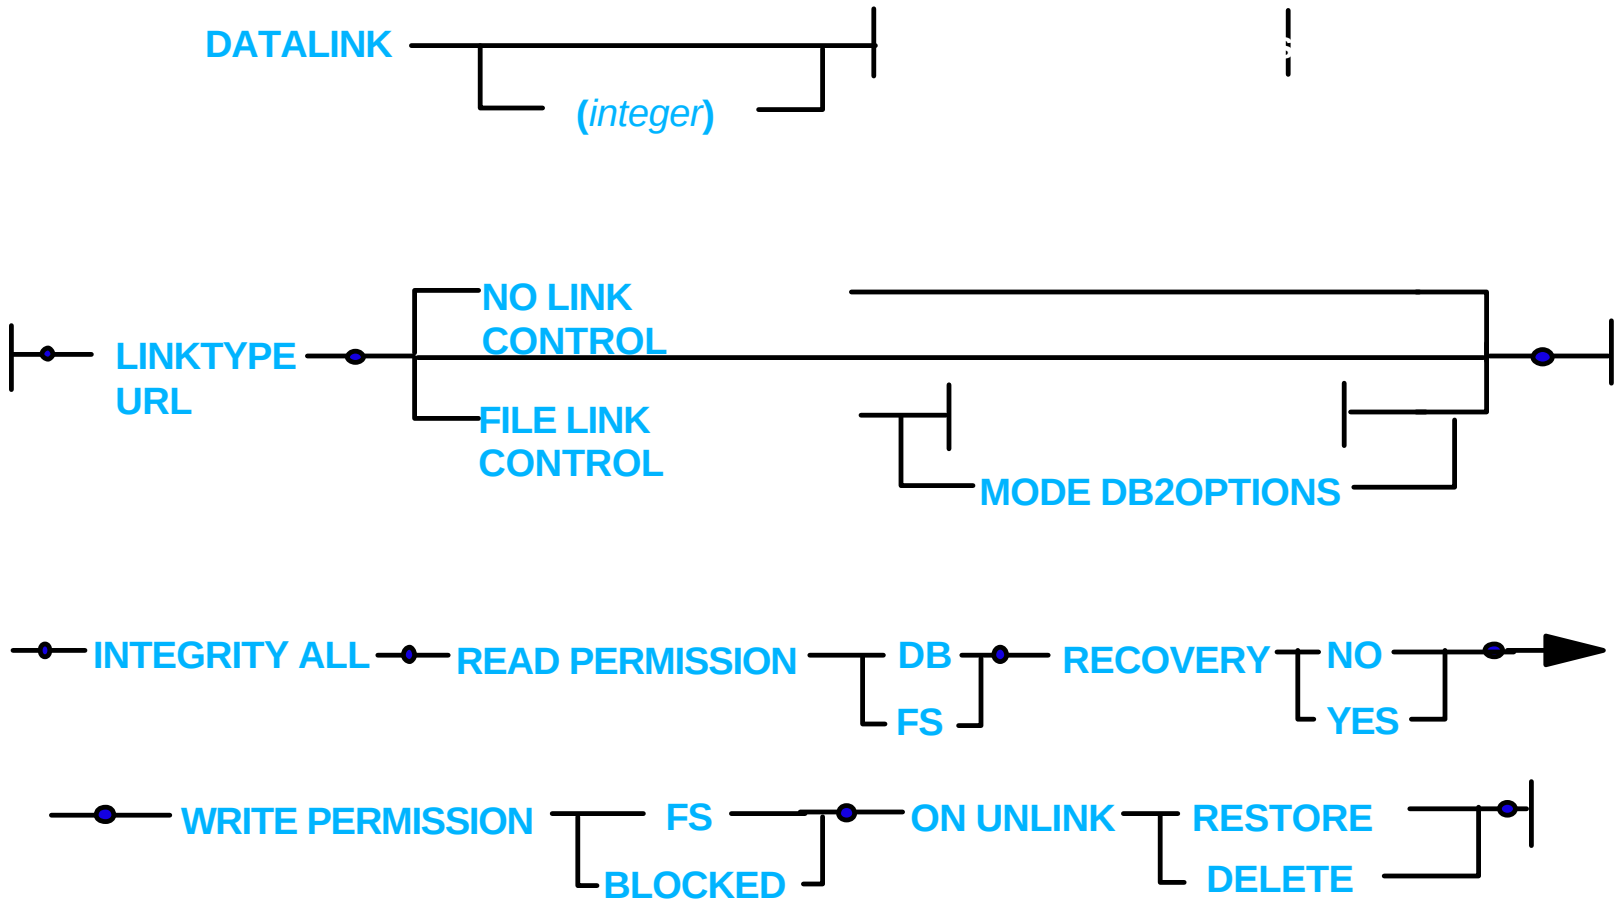

DATALINK Datatype Features

Opt #	Read	Write	Recovery	Unlink	Referential Integrity	DB Access
1	FS	FS	No	N/A	✓	
2	FS	Blocked	No	N/A	✓	
3	FS	Blocked	Yes	N/A	✓	
4	DB	Blocked	No	Delete	✓	✓
5	DB	Blocked	Yes	Delete	✓	✓
6	DB	Blocked	No	Restore	✓	✓
7	DB	Blocked	Yes	Restore	✓	✓

Valid Combinations for FILE LINK CONTROL Options

Scalar functions

- DLVALUE
- DLLINKTYPE
- DLURLSCHEME
- DLURLSERVER
- DLURLPATH
- DLURLPATHONLY
- DLURLCOMPLETE
- DLURLCOMMENT

CLI

- SQLBuildDataLink
- SQLGetDataLinkAttr

SQL Restrictions for DATALINK columns

- Cannot be part of an index
- cannot be part of a constraint
- cannot be compared

Table States

- DRP (DataLink Reconcile Pending)
  - ➡ Data Links Manager metadata is out-of-sync with table data
- DRNP (DataLink Reconcile Not Possible)
  - ➡ Data Links Manager metadata is missing for the table
- RESTORE and ROLLFORWARD utilities may set these states
- Table access is restricted in these states
  - ➡ SELECT is permitted
  - ➡ INSERT/DELETE is not permitted
  - ➡ UPDATE is permitted selectively in DRNP state
  - ➡ User may set state to DRNP is (s)he suspects integrity has been compromised
  - ➡ SELECT access may also be prohibited by setting the CHECK PENDING state in addition to the DRNP state
- RECONCILE utility should be run to bring the table out of DRP state
- See the SQL & DB2 Administration Guides for details

Some Configuration Parameters

- DL_EXPINT
  - ➡ Expiry time of the token generated by DB2
- DL_TOKEN
  - ➡ Algorithm choice for generating the token
- DL_UPPER
  - ➡ Whether the token generated can have only upper case or both upper and lower case characters in it
- DL_TIME_DROP
  - ➡ Number of days after a DROP for which the archive of unlinked files should be retained
- REC_HIS_RETENTN
  - ➡ Number of days entries are retained in the history file
- NUM_DB_BACKUP
  - ➡ Number of backups after which archive of unlinked files can be deleted

Architecture on DCE-DFS for AIX

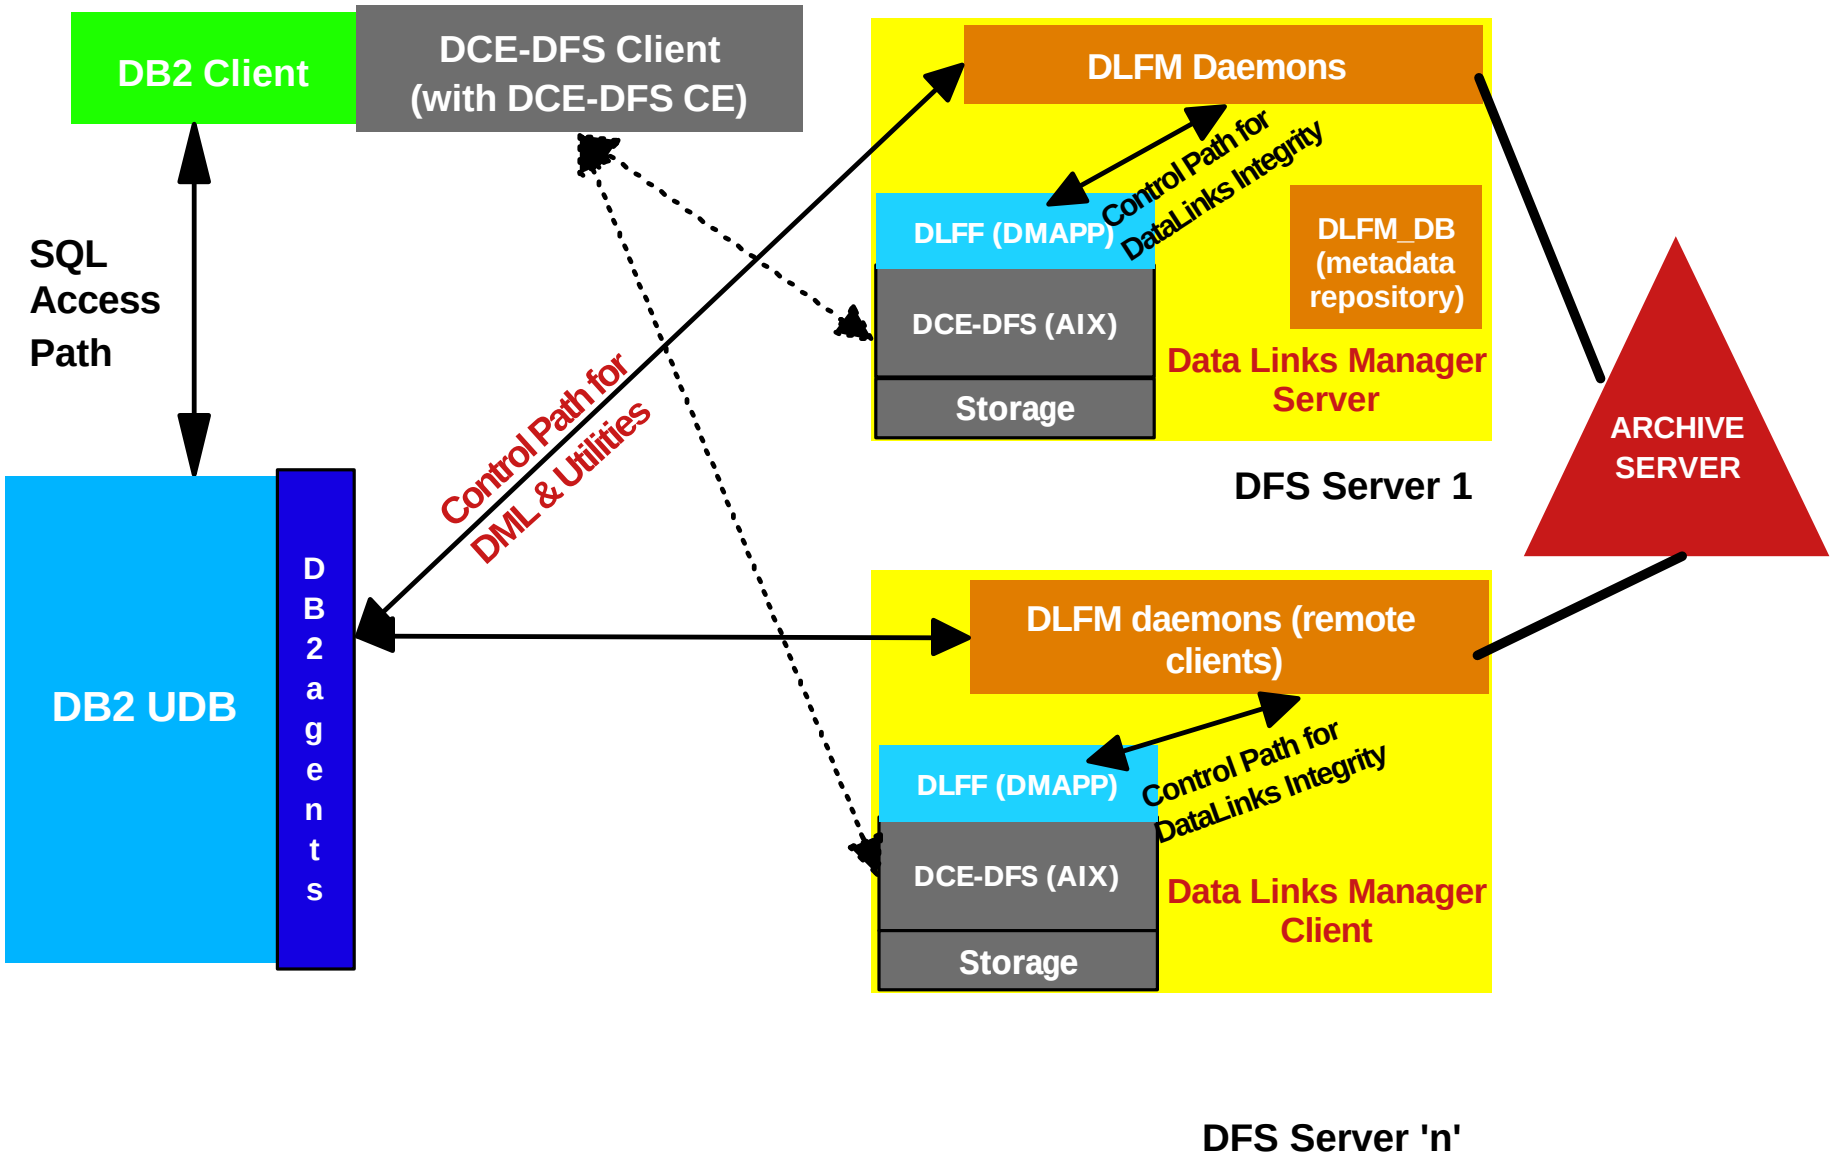

File Archive Optimization in DCE-DFS

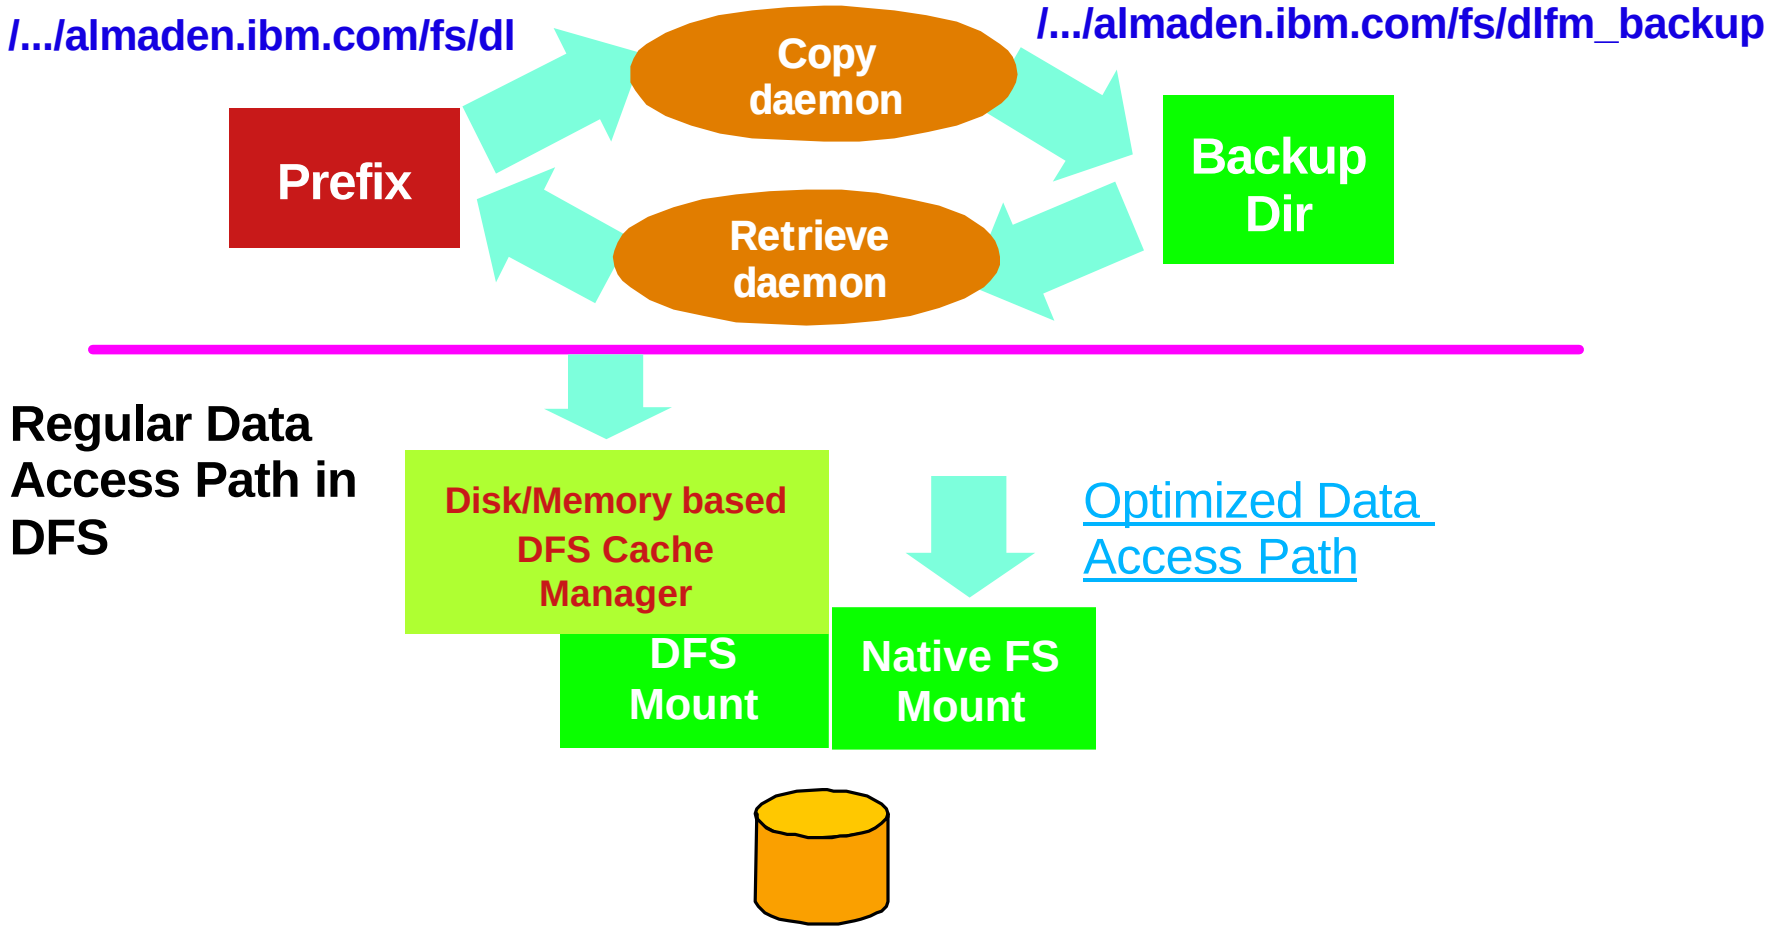

/.../almaden.ibm.com/fs/dl/kiran.pic <==> /localmount/dl/kiran.pic