IBM Informix Dynamic Server v10.0

Technical Overview

Carlton Doe
Consulting IT Specialist
Global Technical Partner Organization

cdoe@us.ibm.com

Trademarks and Disclaimers

References in this document to IBM products or services do not imply that IBM intends to make them available in every country. The following terms are trademarks or registered trademarks of International Business Machines Corporation in the United States, other countries, or both:

IBM			iSeries	OS/400	Informix	
	WebSphere					
IBM(logo)		pSeries	AIX		Cloudscape	MQSeries
e(logo)business		xSeries	DB2		DB2 Universal Database	CICS
Netfinity		zSeries	OS/390		IMS	

Lotus, Domino, Freelance Graphics, and Word Pro are trademarks of Lotus Development Corporation and/or IBM Corporation in the United States and/or other countries.

Java and all Java-based trademarks are trademarks of Sun Microsystems, Inc. in the United States, other countries, or both. Microsoft, Windows, Windows NT, and the Windows logo are trademarks of Microsoft Corporation in the United States, other countries, or both. ActionMedia, LANDesk, MMX, Pentium and ProShare are trademarks of Intel Corporation in the United States, other countries, or both. UNIX is a registered trademark of The Open Group in the United States and other countries. Other company, product and service names may be trademarks or service marks of others.

Information is provided "AS IS" without warranty of any kind.

All customer examples described are presented as illustrations of how those customers have used IBM products and the results they may have achieved. Actual environmental costs and performance characteristics may vary by customer.

Information in this presentation concerning non-IBM products was obtained from a supplier of these products, published announcement material, or other publicly available sources and does not constitute an endorsement of such products by IBM. Sources for non-IBM list prices and performance numbers are taken from publicly available information, including vendor announcements and vendor worldwide homepages. IBM has not tested these products and cannot confirm the accuracy of performance, capability, or any other claims related to non-IBM products. Questions on the capability of non-IBM products should be addressed to the supplier of those products.

All statements regarding IBM future direction and intent are subject to change or withdrawal without notice, and represent goals and objectives only. Contact your local IBM office or IBM authorized reseller for the full text of the specific Statement of Direction.

Some information in this presentation addresses anticipated future capabilities. Such information is not intended as a definitive statement of a commitment to specific levels of performance, function or delivery schedules with respect to any future products. Such commitments are only made in IBM product announcements. The information is presented here to communicate IBM's current investment and development activities as a good faith effort to help with our customers' future planning.

Performance is based on measurements and projections using standard IBM benchmarks in a controlled environment. The actual throughput or performance that any user will experience will vary depending upon considerations such as the amount of multiprogramming in the user's job stream, the I/O configuration, the storage configuration, and the workload processed. Therefore, no assurance can be given that an individual user will achieve throughput or performance improvements equivalent to the ratios stated here.

TOPIC LIST

Who I Am

- 12 + years of experience as DBA and engine admin mainly in retail environments
- Co-founded, presided over, and former member of the Board of Directors of the International Informix Users Group (IIUG)

- Wrote two InformixPress books on the engine:

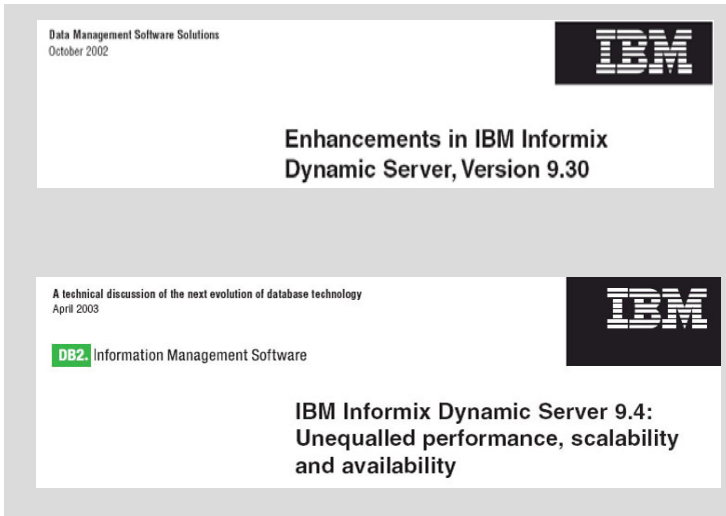

TOPIC LIST

Informix Dynamic Server 10.0 Technical Deep Dive

TOPIC LIST

Topic List

ADMIN & USABILITY	PERFORMANCE	REPLICATION	AVAILABILITY	SECURITY	APPLICATIONS
Rename dbspace	Configurable Page Sizes	DRAUTO	Online Index Build	PAM Authentication	JDBC 3.0 Support
Single-User mode	Memory Allocation to Non-PDQ Queries	Replicate Templates	Faster Recovery w/ Fuzzies	Secure Environment Check	.NET Support
HDR Setup w/ EBR	Dynamic OPTCOMPIND	Alter Table Support	Table Level Restore	Datablade Registration Restrictions	ESQL/C to DB2
Tablespace Tablespace Management	External Optimizer Directives	Replicate Resynch		Database Level Permissions	
SHMEM > 4G		Fixing Corrupt Secondary Indexes		Trigger Introspection	
Multiple Fragments in one dbspace					
ontape use of STDIO					
Misc Backup/Restore					

TOPIC LIST

IDS 10 Features
Performance

TOPIC LIST

Topics: Performance Enhancements

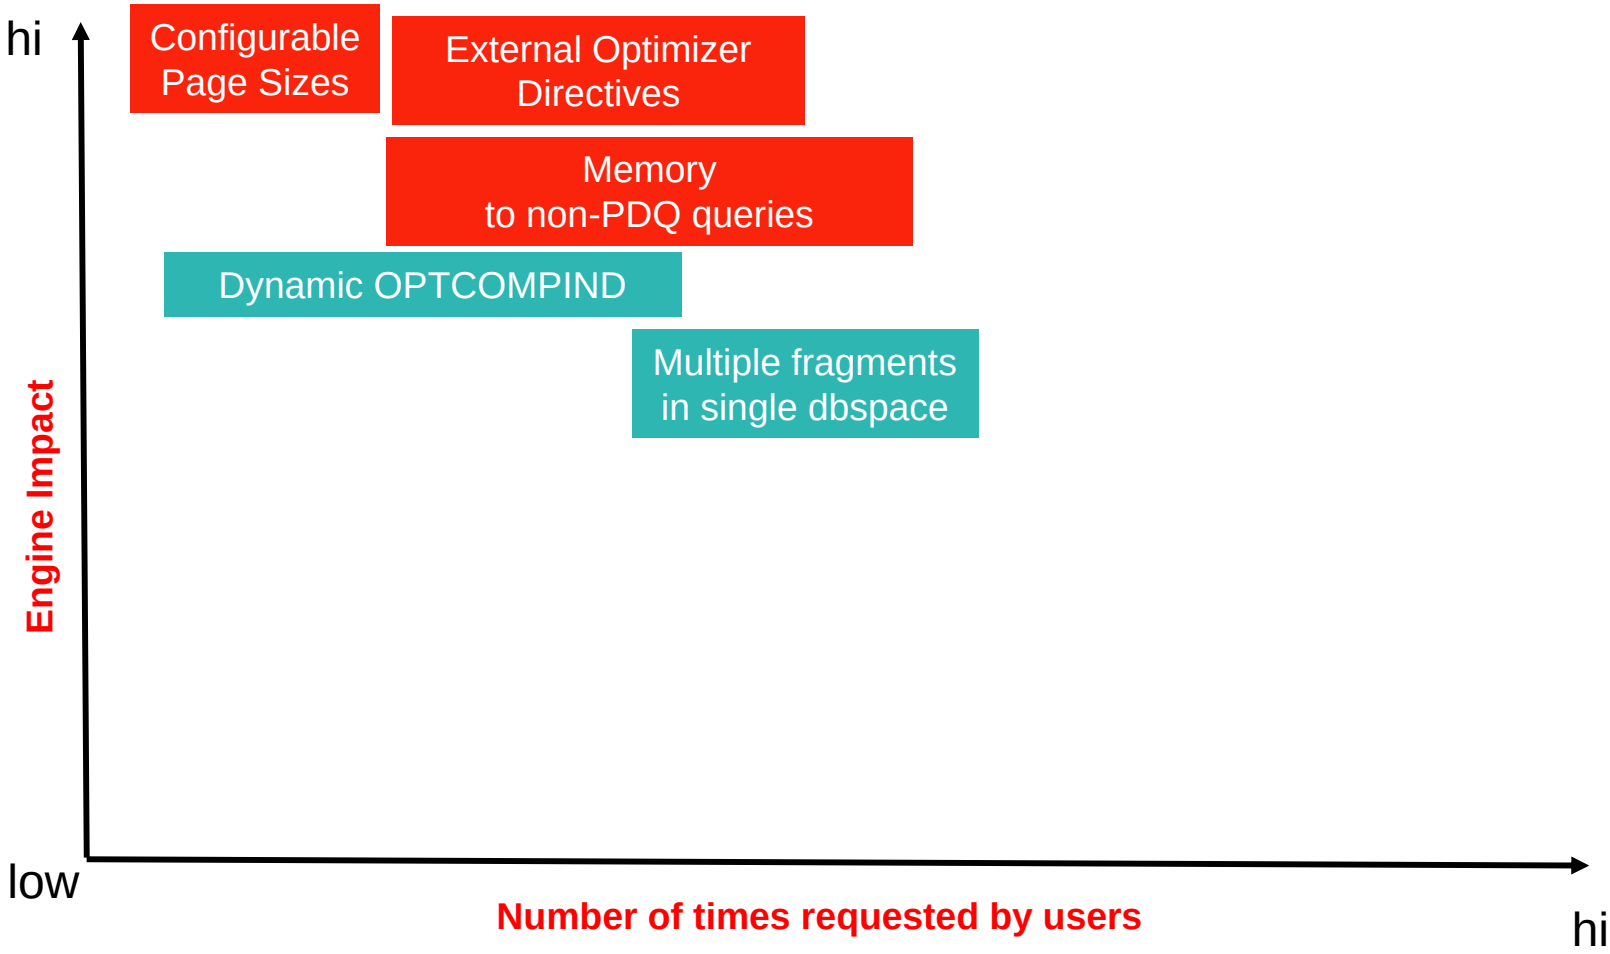

TOPIC LIST

IDS 10 Features
Performance

Configurable Page Sizes

TOPIC LIST

Configurable Page Sizes

- A feature that has been a “curiosity” to the customers.
- IDS XPS has a similar feature already
  - implemented differently – must define page size before building engine.
  - page size is the same for all structures.
- Similar to DB2
  - but not the same implementation.
  - and not for the same reason.

TOPIC LIST

Implementation – 2K Port

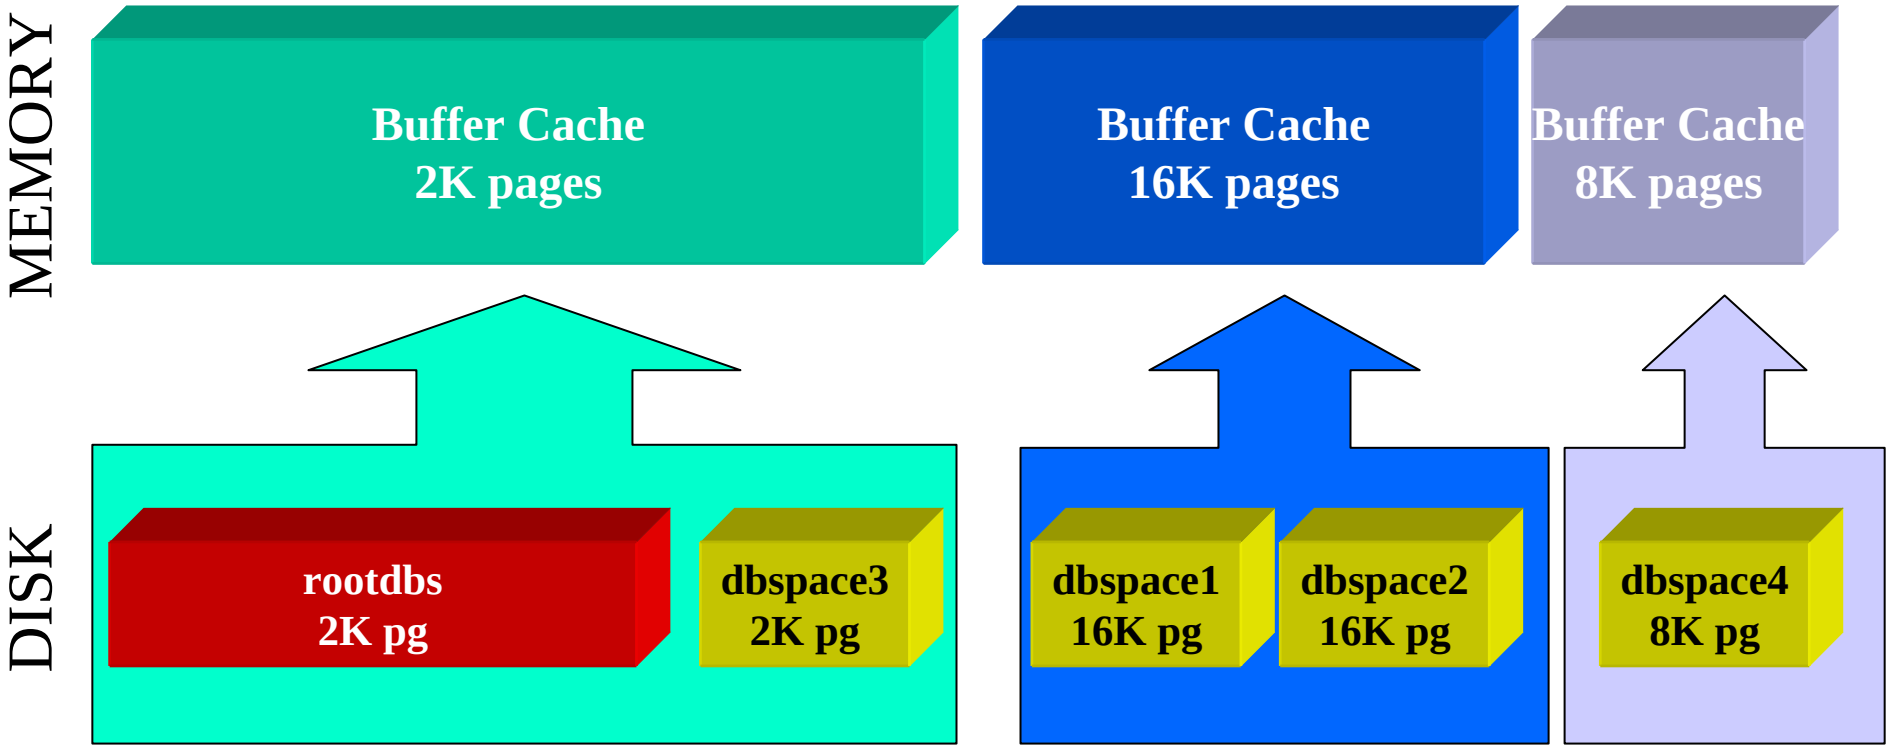

- each port still has a “default page size” – 2K or 4K
- some structure must be in a dbspace with the default page size
- the number of buffer pools depends on the “available page sizes”

TOPIC LIST

What did we solve with Configurable Pages?

- Objective #1: Space efficiency
  - larger pages up to **16K** bytes contiguous space
- Objective #2: Increased maximum key size
  - longer keys up to **3K** bytes
- Objective #3: Access efficiency
  - less I/O operations for data and indices

TOPIC LIST

First…some review on page/row overhead….

TOPIC LIST

Page Overhead Basics

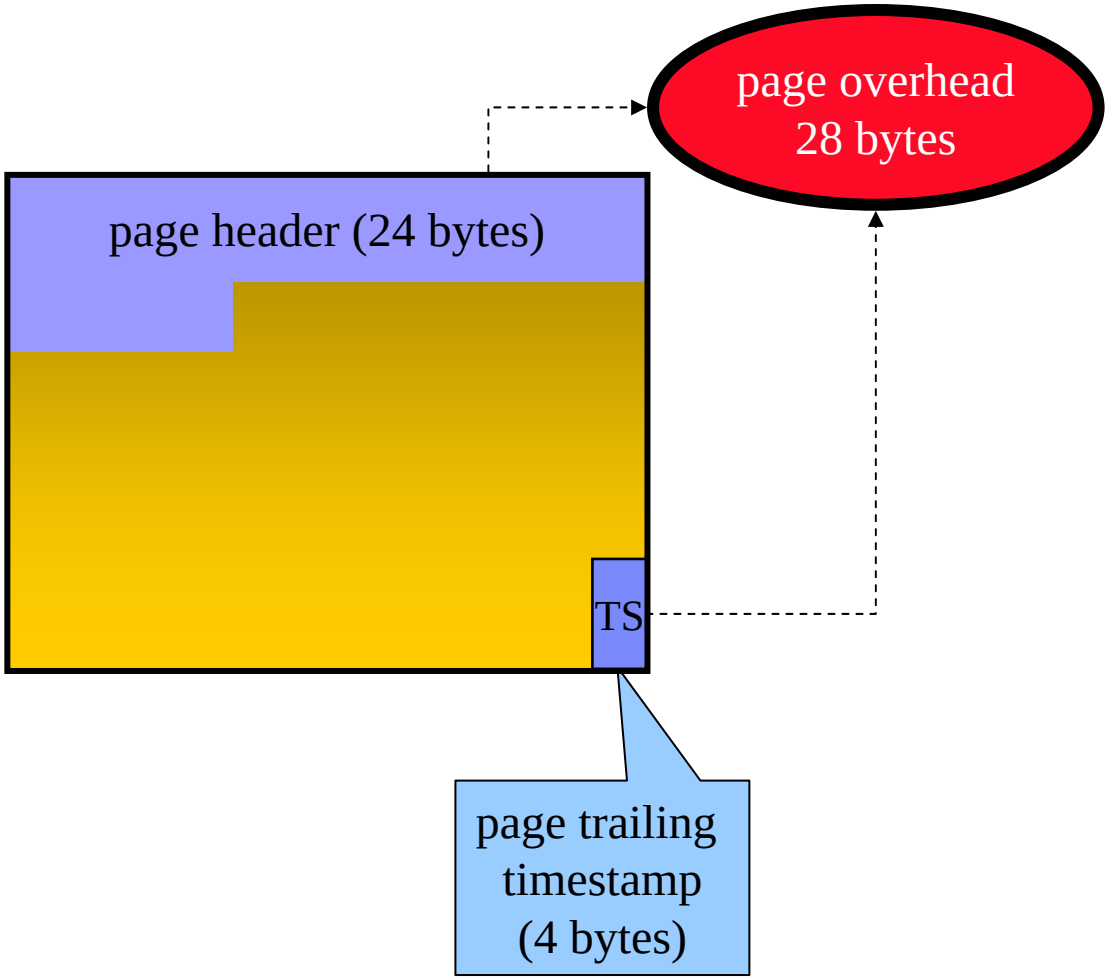

For example: a 2K page

page bytes	2048
page overhead	-28
available bytes	2020

TOPIC LIST

Row Overhead Basics – Index Pages

would point to data page or fragment…

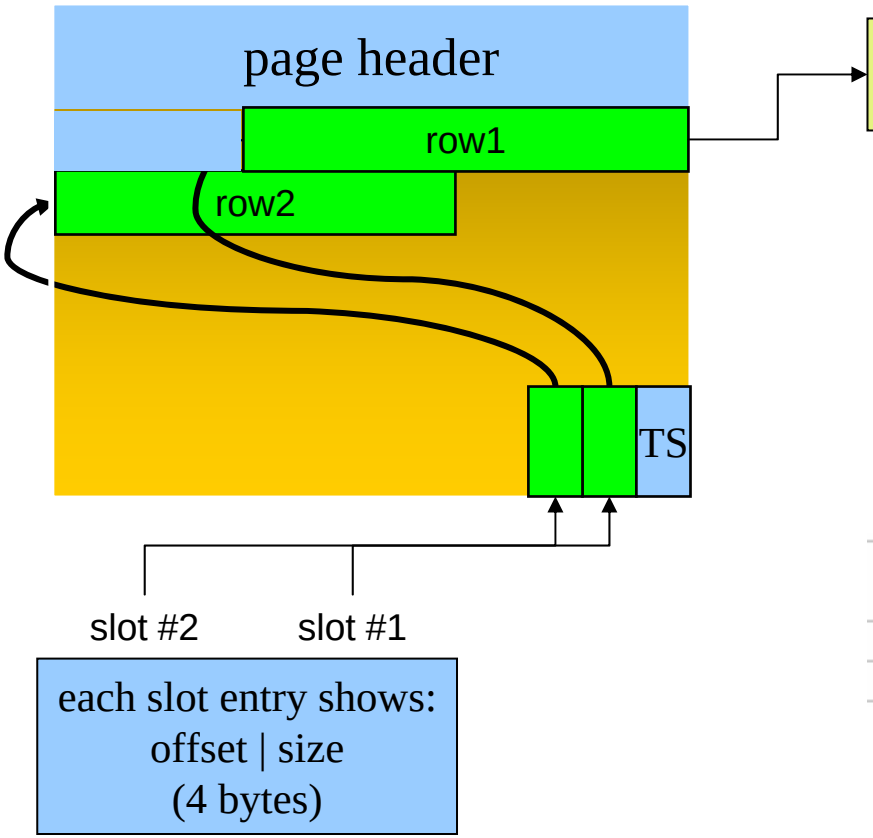

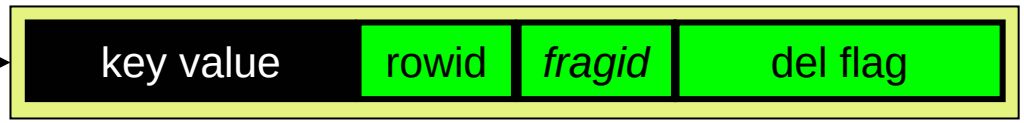

what is fragmented?

Data	Indexes	ROWID	FRAGID	DEL FLAG	Slot Table	Overhead Bytes
Yes	No	4	4	1	4	13
No	Yes	4	0	1	4	9

TOPIC LIST

Overhead Summary

- Page overhead — 28 bytes
- Row overhead — minimum - 4 bytes per row
  maximum – 13 bytes per row
- *When client tables approach 1B rows, this could be significant.*

TOPIC LIST

Objective #1: Space Efficiency

TOPIC LIST

So What About Space Efficiency?

- **Example**
  - A row size of 1200 bytes:
    - 1 row fits on a 2k page (6k every 3 rows).
    - 3 rows fit on a 4k page, a savings of %33
  - For thirty 1200-byte rows:
    - A 2k page size requires 60k.
    - A 4k page size would require only 40k.
    - A 6k page size requires just 36k, a 40% savings

TOPIC LIST

Proof – 2K dbspace

```
TBLspace Report for misc_db:informix.tenktup1

    Physical Address                3:516
    Creation date                   01/12/2005 01:24:21
    TBLspace Flags                  801         Page Locking
                                                TBLspace use 4 bit bit-
    Maximum row size                208
    Number of special columns       0
    Number of keys                  0
    Number of extents               1
    Current serial value            1
    Pagesize (k)                    2
    First extent size               8
    Next extent size                8
    Number of pages allocated       1120
    Number of pages used            1113
    Number of data pages            1112
    Number of rows                  10000
    Partition partnum               3145809
    Partition lockid                3145809

    Extents
         Logical Page      Physical Page         Size Physical Pages
                    0             3:1008         1120        1120
```

TOPIC LIST

Proof – 16K dbspace

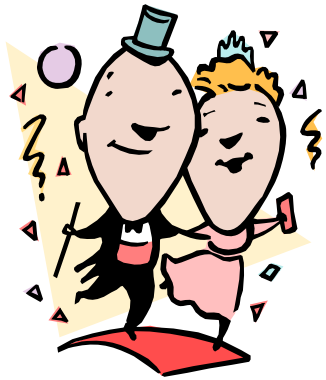

```
TBLspace Report for wisc_db:informix.tenktup2

    Physical Address                2:6904
    Creation date                   01/12/2005 01:24:22
    TBLspace Flags                  801         Page Locking
                                                TBLspace use 4 bit bit
    Maximum row size                208
    Number of special columns       0
    Number of keys                  0
    Number of extents               1
    Current serial value            1
    Pagesize (k)                    16
    First extent size               4
    Next extent size                4
    Number of pages allocated       132
    Number of pages used            131
    Number of data pages            130
    Number of rows                  10000
    Partition partnum               2097312
    Partition lockid                2097312

    Extents
         Logical Page      Physical Page        Size Physical Pages
                    0             2:7576         132        1056
```

2K pages … 1112
16K pages…. 130

88% reduction

TOPIC LIST

Objective #2: Larger Key Size

TOPIC LIST

Increased Maximum Key Size

- **Increased maximum key size - longer keys up to 3K bytes**
  - Placing more keys on a page, we support longer keys without drastically increasing index level depth.
    - this is very significant for duplicate indexes with a high number of dups or volatility.
    - this won't change customers that are already building "wider indexes" – they can't span pages anyway.
  - The pre-10.0 key size limit was also a roadblock to utilizing the UNICODE character set, which caused some key values to exceed the max length.

TOPIC LIST

Proof – unique index (INT – 4 bytes)

2K dbspace

```
Index Usage Report for index tk1 on wisc_db:informix.tenktup1

                    Average     Average
Level     Total  No. Keys  Free Bytes
-----  --------  --------  ----------
    1         1        72        1160
    2        72       138         214
-----  --------  --------  ----------
Total        73       137         227
```

```
TBLspace Usage Report for wisc_db:informix.tenktup1

Type                 Pages       Empty  Semi-Full
---------------- ---------- ---------- ----------
Free                     2
Bit-Map                  1
Index                   73
Data (Home)              0
                 ----------
Total Pages             76
```

16K dbspace

```
Index Usage Report for index tk2 on wisc_db:informix.tenktup2

                    Average     Average
Level     Total  No. Keys  Free Bytes
-----  --------  --------  ----------
    1         1         9       16252
    2         9      1111        1911
-----  --------  --------  ----------
Total        10      1000        3345
```

```
TBLspace Usage Report for wisc_db:informix.tenktup2

Type                 Pages       Empty  Semi-Full
---------------- ---------- ---------- ----------
Free                     1
Bit-Map                  1
Index                   10
Data (Home)              0
                 ----------
Total Pages             12
```

TOPIC LIST

Objective #3: Access Inefficiencies

TOPIC LIST

Access Inefficiencies : Forward Pointers

- Forward Pointers – data rows. On data pages - for rows larger than available page space on 2K or 4K ports, we utilize FORWARD POINTERS

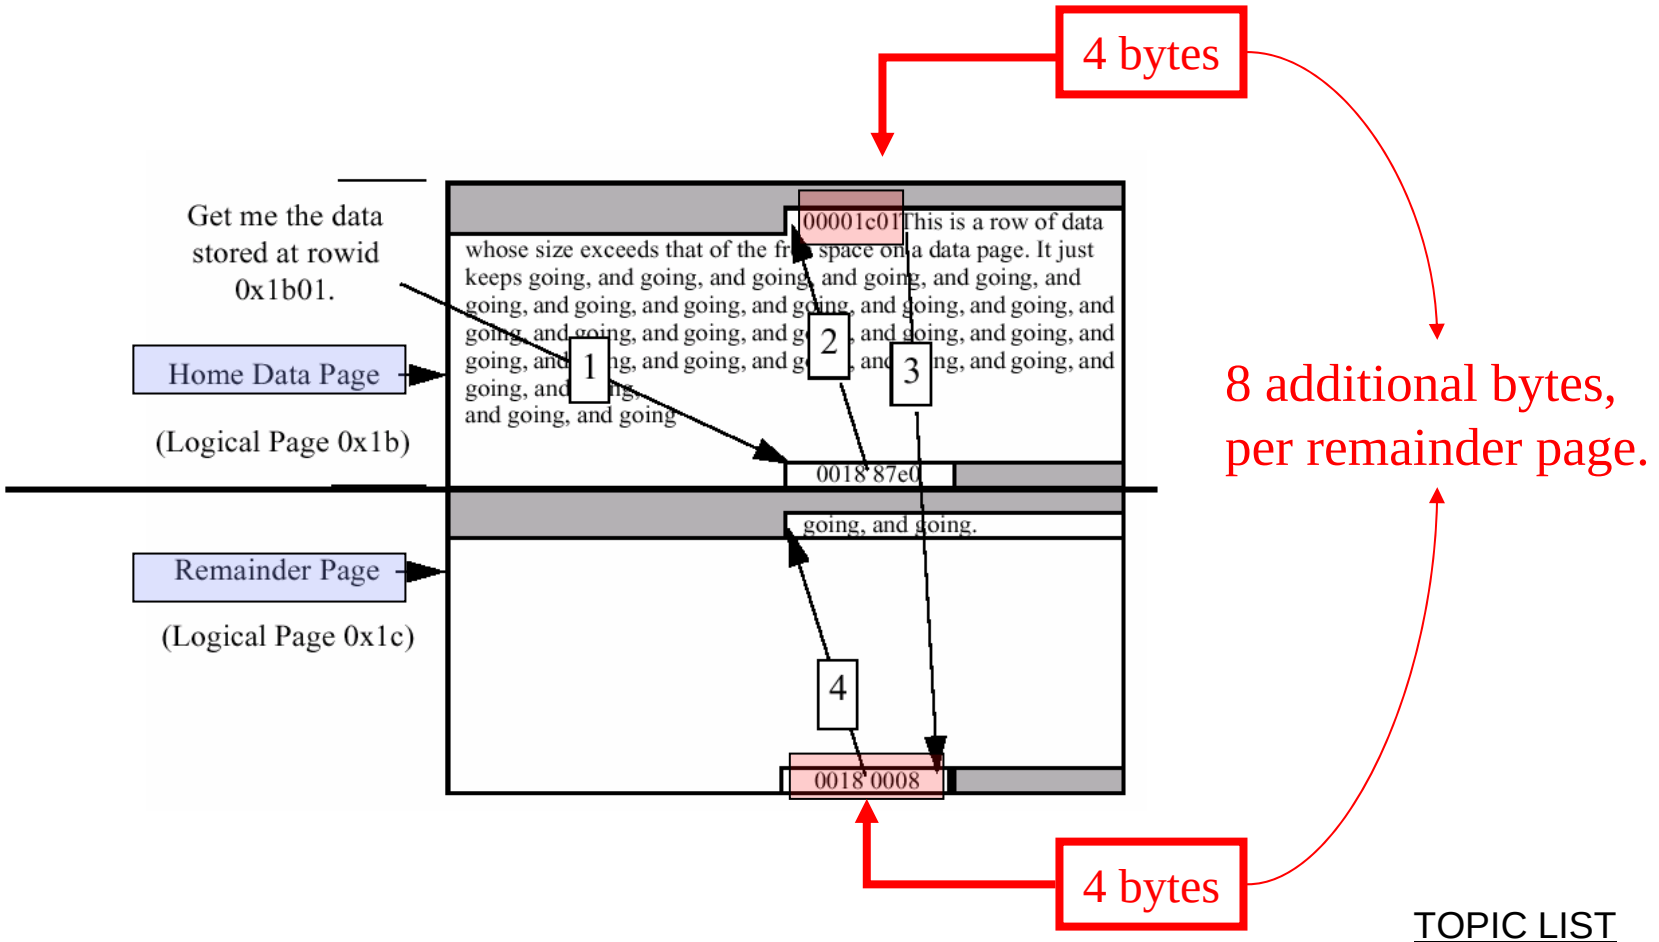

TOPIC LIST

Access Inefficiencies : Page Accesses

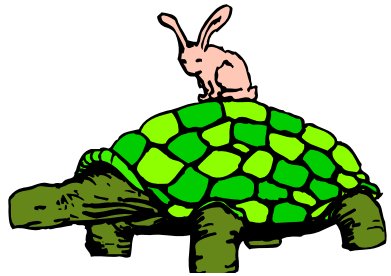

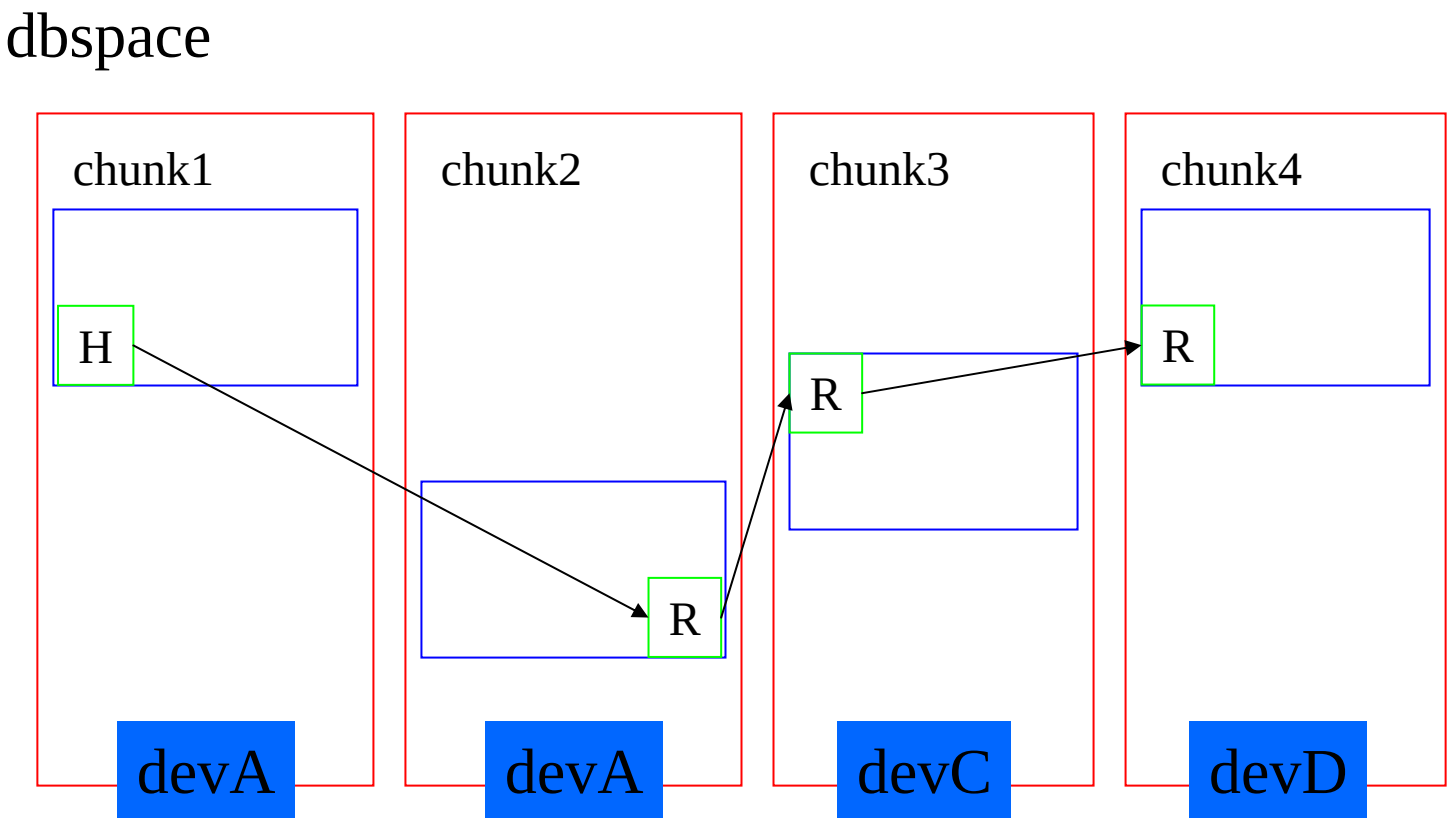

TOPIC LIST

Example: Creating a 16K page dbspace

onspaces –c –d sparky –k 16 –p <path> -o offset -s <size>

```
informix:/export/home/informix> onstat -d

IBM Informix Dynamic Server Version 9.50.UC1B5   -- Single-User -- Up 12:34:39 --

Dbspaces
address  number   flags      fchunk   nchunks  pgsize   flags    owner    name
bd277d8  1        0x40001    1        1        2048     N  B     informix rootdbs
d08b3f0  2        0x40001    2        1        16384    N  B     informix sparky
 2 active, 2047 maximum
```

```
informix:/export/home/informix> onstat -b

IBM Informix Dynamic Server Version 10.00.UC1      -- On-Line -- Up

Buffers
address  userthread flgs pagenum          memaddr  nslots pgflgs

Buffer pool page size: 2048
 0 modified, 500000 total, 524288 hash buckets, 2048 buffer size

Buffer pool page size: 16384
 0 modified, 1000 total, 1024 hash buckets, 16384 buffer size
```

TOPIC LIST

Miscellaneous

- Number of BUFFER pools is dependent upon default page size:
  - On a system with a 2k default page size, the maximum is 16
  - On a system with a 4k default page size, the maximum is 8.
- ONCONFIG parameter added

```
# The following parameters are related to the buffer pool
BUFFERPOOL      default,buffers=1000,lrus=8,lru_min_dirty=50.000000,lru_max_dirty=60.000000
BUFFERPOOL      size=2K,buffers=1000,lrus=8,lru_min_dirty=50.000000,lru_max_dirty=60.000000
BUFFERPOOL      size=16K,buffers=1000,lrus=8,lru_min_dirty=50.000000,lru_max_dirty=60.000000
~
~
```

TOPIC LIST

Page Concepts

Changing:

- home data pages
- remainder pages
- partition freemap pages
- partition partition pages
- chunk free list pages
- index pages
- partition blob pages - partition blobs will use the same page size as the large page size defined for the dbspace they reside in.

Not changing:

- maximum number of pages per partition still 16,777,216.
- maximum number of rows per page remains at 255.
- maximum number of parts per key remains 16.
- text or byte types in a blobspace are not affected.
- smartblobs are not affected.
- user-defined external spaces are not affected.
- r-tree indexes must be stored in a dbspace with the default page size.
- the rootdbs must be the default page size.
- the physical and logical log dbspaces must be the default page size.
- dynamically created logs must be in a default page size dbspace.

TOPIC LIST

Row Size / Number of Rows Considerations

- Primary purpose for CPS is to accommodate large rows (rows larger than the default page size).
- Be careful!
  - the “rows per data page limit” of 255 is STILL enforced with 10.0.
  - *regardless of row size, the max rows a data page can house is 255.*
  - this limitation is due to the ROWID format

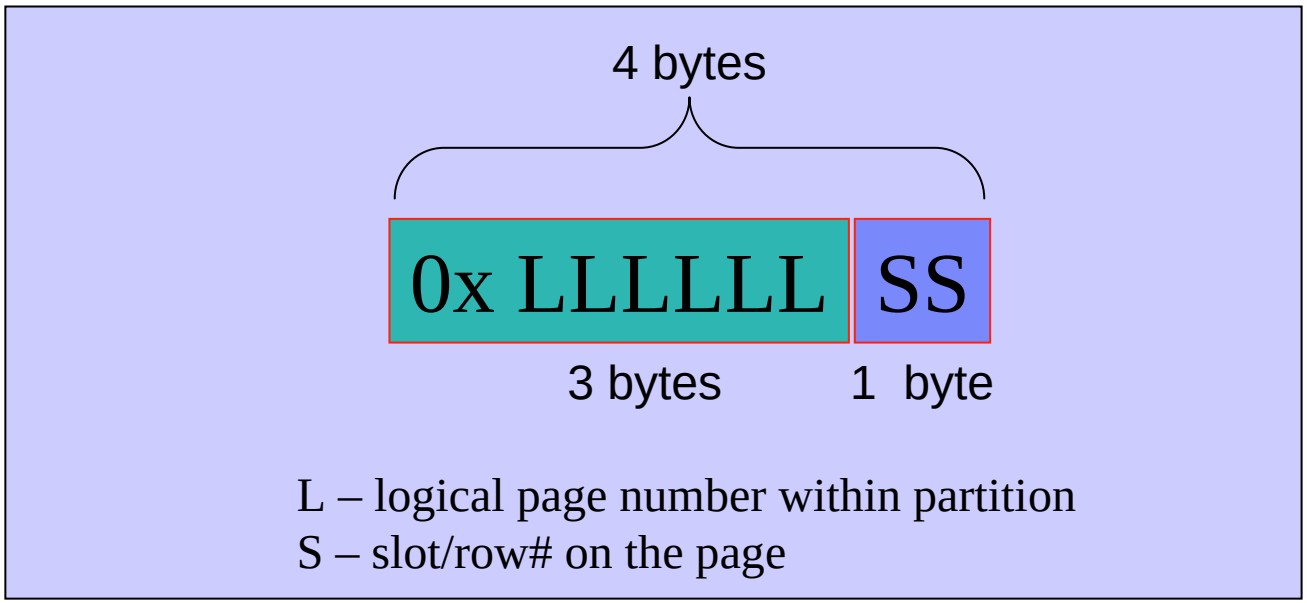

TOPIC LIST

onstat –g buf

```
informix:/export/home/informix> onstat -g buf |more

IBM Informix Dynamic Server Version 9.50.UC1B5   -- On-Line -- Up 10:14:56 -- 10137

Profile

Buffer pool page size: 2048
dskreads   pagreads   bufreads   %cached dskwrits   pagwrits   bufwrits   %cached
1346       7542       6244       78.44   527        7542       319         0.00
bufwrits_sinceckpt  bufwaits   ovbuff     flushes
0                   2          0          9

Fg Writes     LRU Writes    Chunk Writes
0             0             124


Buffer pool page size: 16384
dskreads   pagreads   bufreads   %cached dskwrits   pagwrits   bufwrits   %cached
164        1312       30379      99.46   492        1312       10972       95.52
bufwrits_sinceckpt  bufwaits   ovbuff     flushes
0                   0          0          5

Fg Writes     LRU Writes    Chunk Writes
0             0             492
```

TOPIC LIST

IDS 10 Features
Performance

Memory Allocation to non-PDQ Queries

TOPIC LIST

Performance : Non-PDQ queries - a problem….

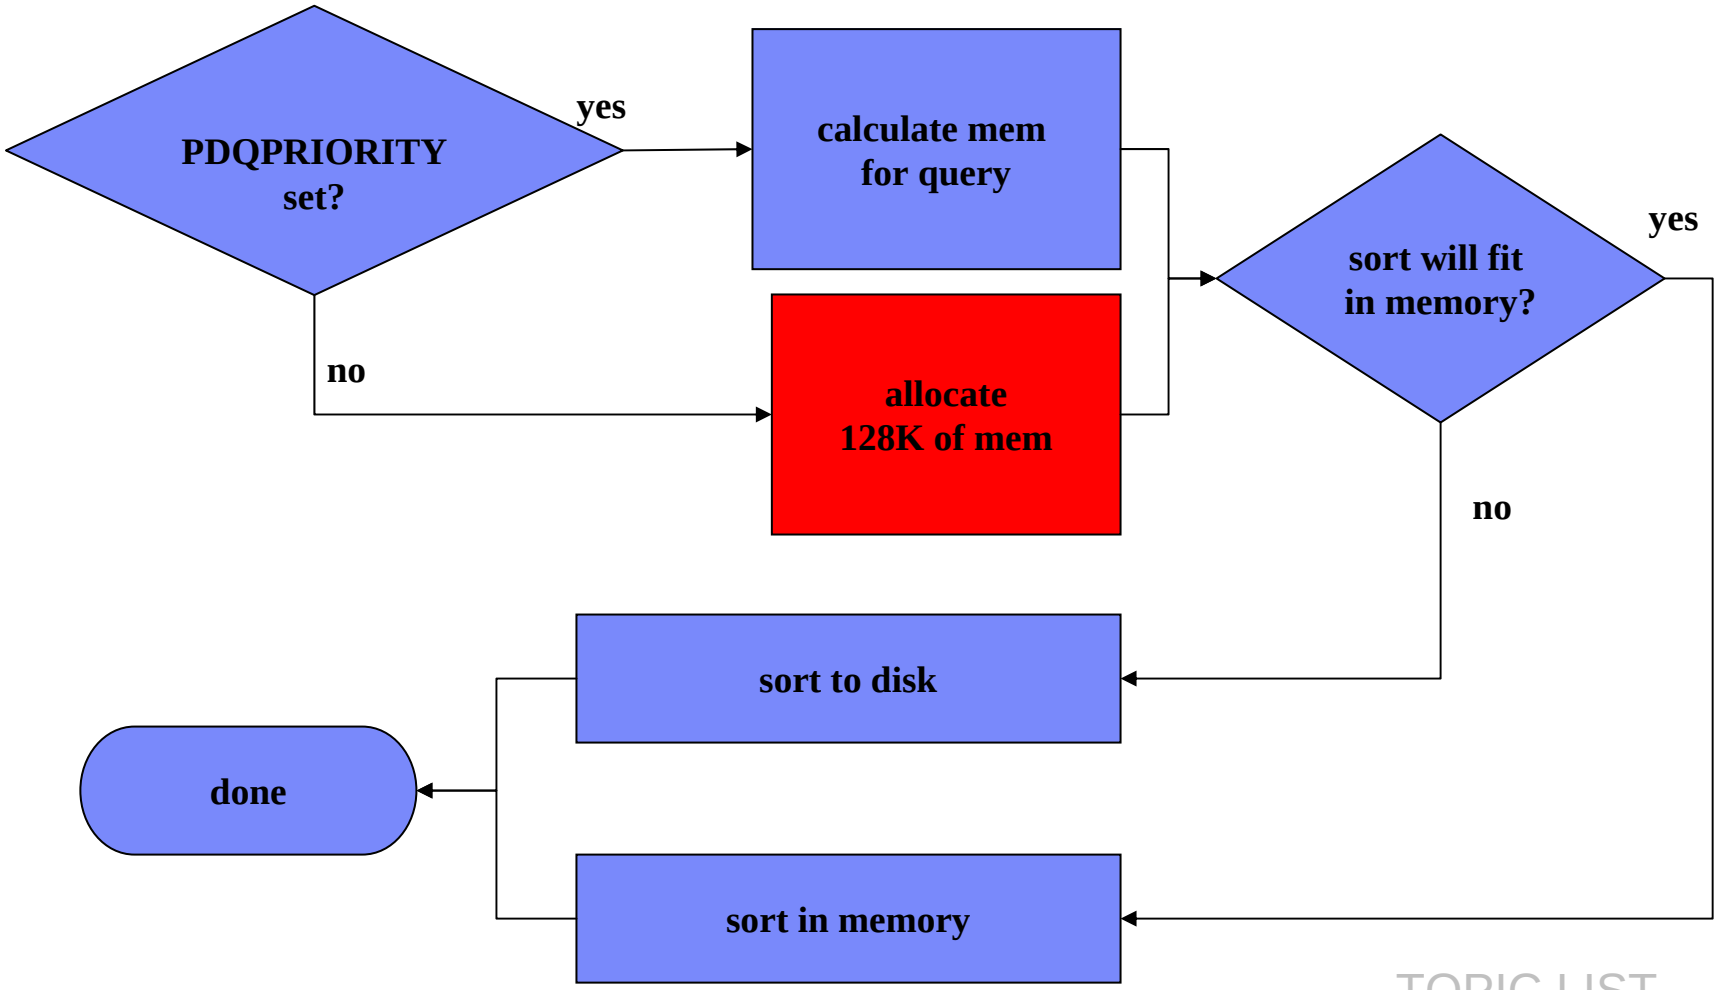

TOPIC LIST

Performance : Memory Allocation for non-PDQ Queries

- This feature was first available in 9.40.xC4.
- You can specify how much memory is allocated to non-PDQ queries.
  - The default of 128K can be insufficient for queries that specify ORDER BY, GROUP BY, hash joins, or other memory-intensive options.
- Use the new configuration parameter, DS_NONPDQ_QUERY_MEM, to specify more memory than the 128K that is allocated to non-PDQ queries by default.

TOPIC LIST

IDS 10 Features
Performance

External Optimizer Directives

TOPIC LIST

Performance : Storing/Applying External Directives

- You can create, save, and reuse external optimizer directives.
- External optimizer directives are useful when it is not feasible to rewrite a query for a short-term solution to a problem
  - for example, when a query starts to perform poorly.
- This feature is implemented as a new SQL statement, SAVE EXTERNAL DIRECTIVES
  - creates and registers external optimizer directives in a new **sysdirectives** table of the system catalog.
  - Use the new **IFX_EXTDIRECTIVES** environment variable or the **IFX_EXTDIRECTIVES** configuration parameter to enable this feature.

TOPIC LIST

IDS 10 Features
Performance

Dynamic OPTCOMPIND

TOPIC LIST

Performance: OPTCOMPIND

A Review

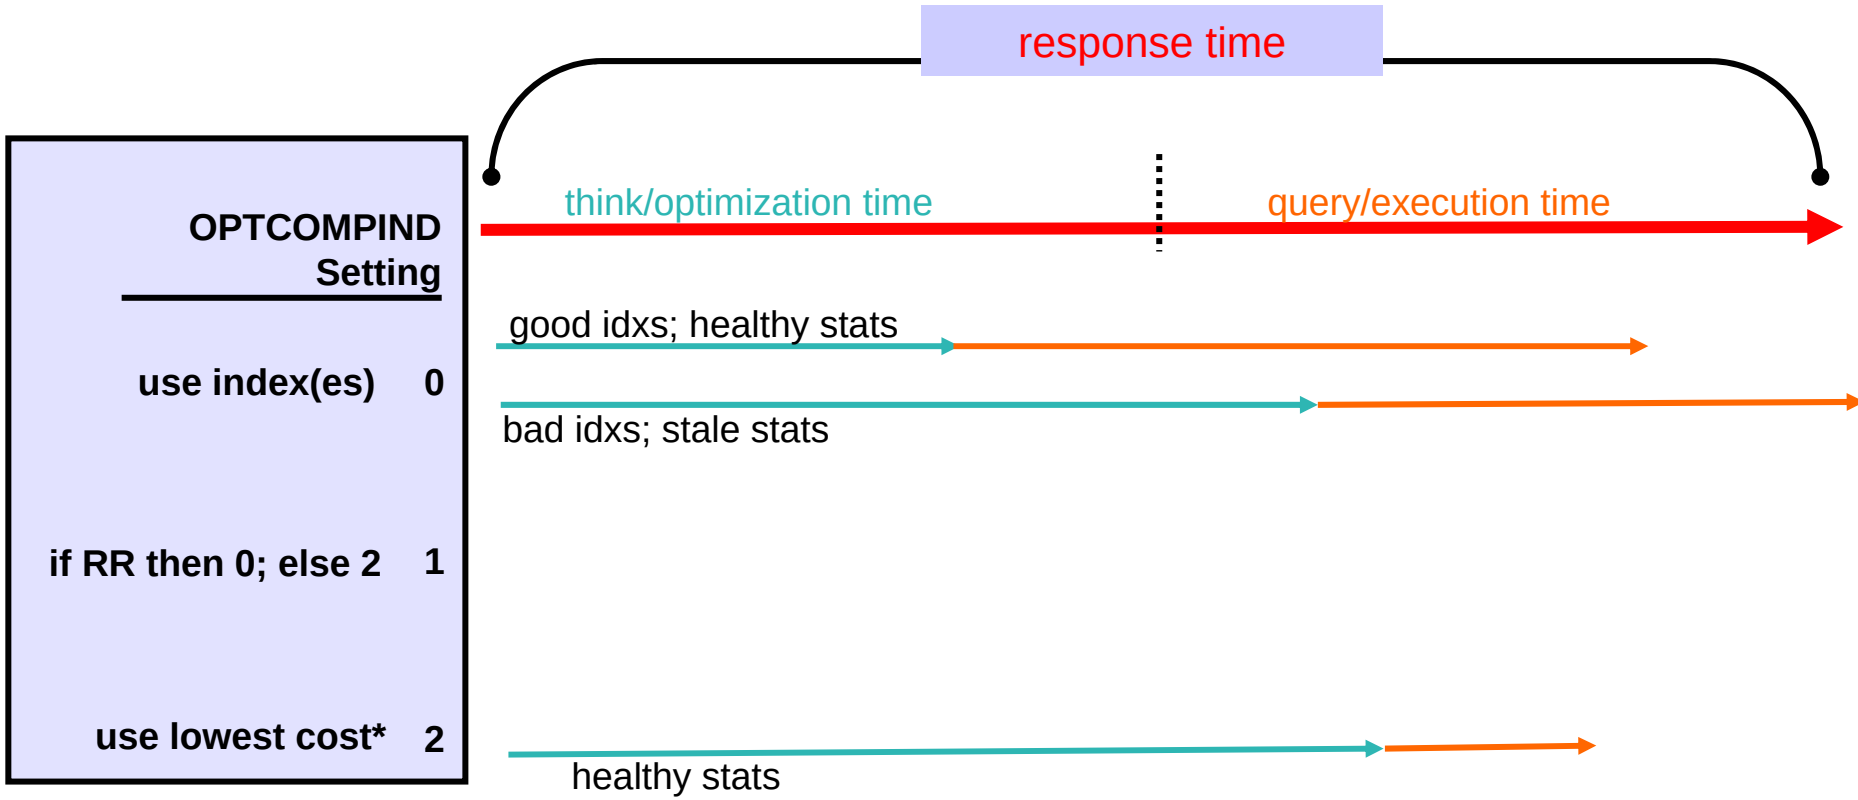

TOPIC LIST

* cost = I/O + (cpu *.03)

Performance : Dynamic OPTCOMPIND

- You can use SET ENVIRONMENT OPTCOMPIND to set **OPTCOMPIND** environment variable dynamically for the current session.
- The value that you enter using this statement takes precedence over the current setting specified in the ONCONFIG file.
- The default setting of the **OPTCOMPIND** environment variable is restored when your current session terminates.
- No other user sessions are affected by SET ENVIRONMENT OPTCOMPIND statements that you execute.

TOPIC LIST

IDS 10 Features
Performance

TOPIC LIST

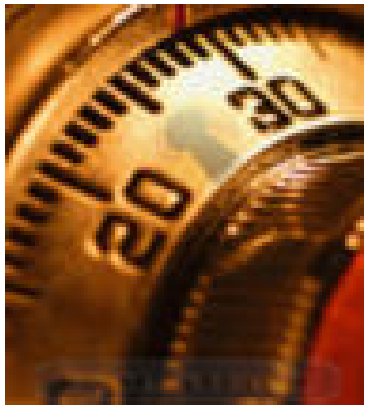

IDS 10 Features
Security

TOPIC LIST

Topics : Security Enhancements

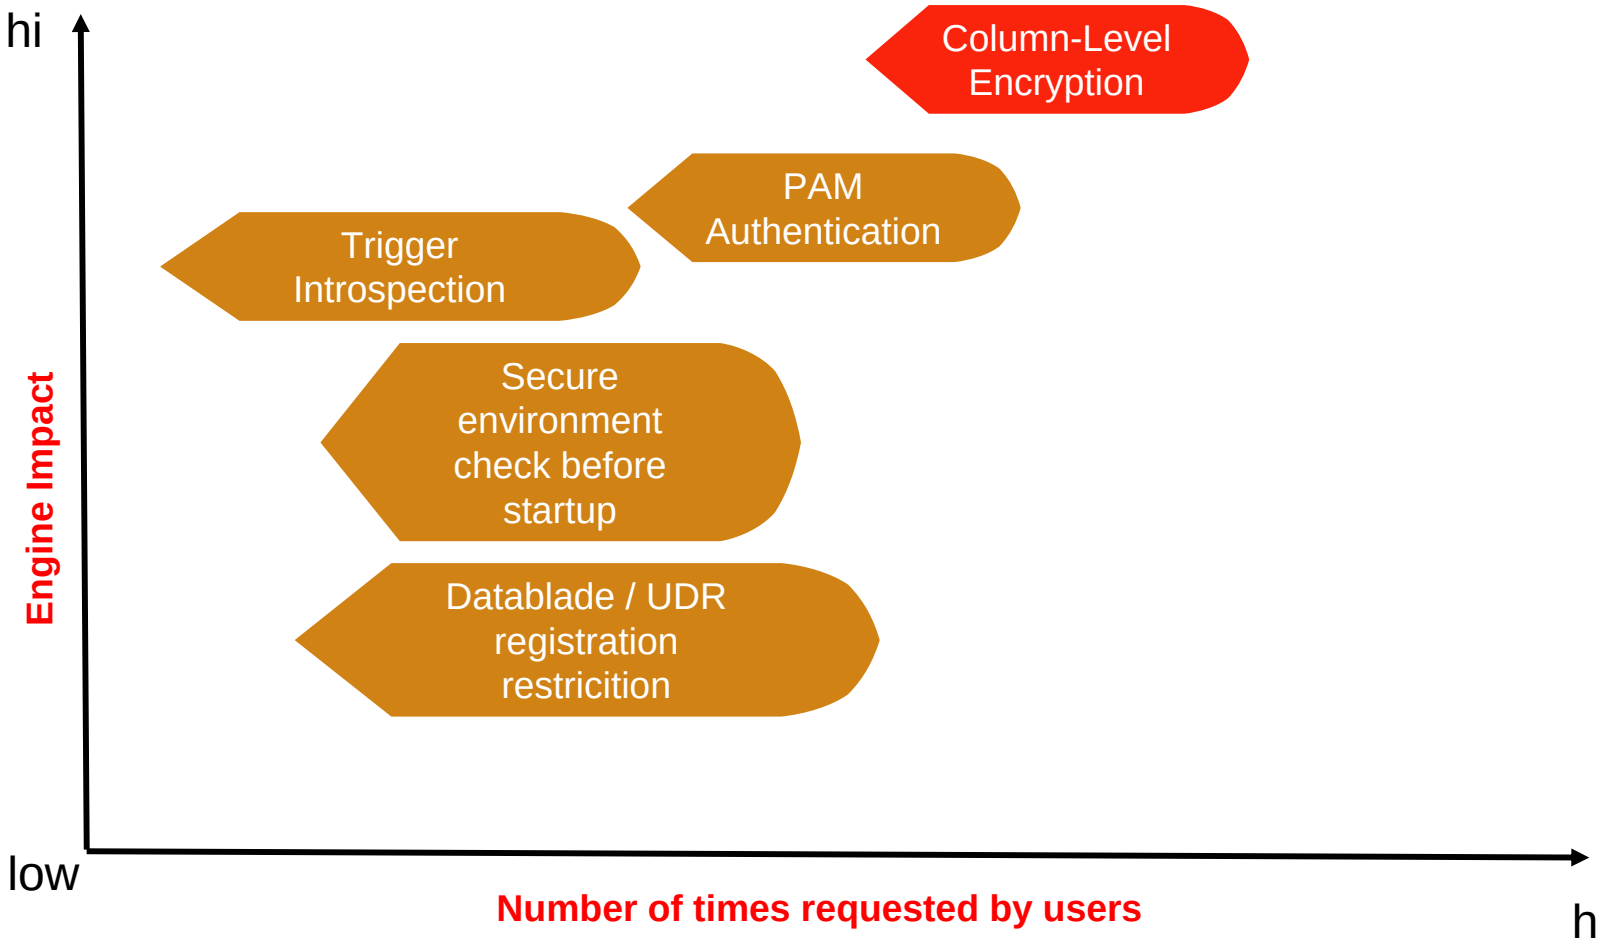

TOPIC LIST

IDS 10 Features
Security

Column-Level Encryption

TOPIC LIST

Security Enhancements: Column Level Encryption

- Highlights
  - Built-in SQL encryption functions to support data encryption
  - Latest cryptographic standards
  - 128 bit AES and Triple-DES
  - Passwords up to 128 characters
  - Session wide password management for easy programming
  - View/trigger/SPL support
  - Consistent with DB2 syntax

TOPIC LIST

Security Enhancements: Column Level Encryption

Usage examples

- INSERT INTO table VALUES (1, ENCRYPT_AES(data, password));
- UPDATE table SET column = ENCRYPT_TDES(data, password)) WHERE …;
- SELECT DECRYPT_CHAR(column, password) FROM table;
- EXECUTE FUNCTION ENCRYPT_AES(data, password, hint);

TOPIC LIST

Security Enhancements: Column Level Encryption

- Example – table creation / data insert
  - Create table

```
create table customer
        (customer_num           serial(101),
        fname                   char(43), -- encrypted char(15)
        lname                   char(43), -- encrypted char(15)
        company                 char(87), -- encrypted char(40)
        address1                char(67), -- encrypted char(30)
        address2                char(67), -- encrypted char(30)
        city                    char(15),
        state                   char(2),
        zipcode                 char(5),
        phone                   char(67), -- encrypted char(18)
        primary key (customer_num));
```

  - Insert data

```
set encryption password "my password";
insert into customer values (101, encrypt_aes("Ludwig"),
    encrypt_aes("Pauli"), encrypt_aes("All Sports Supplies"),
    encrypt_aes("213 Erstwild Court"), encrypt_aes(""), "Sunnyvale", "CA",
    "94086", encrypt_tdes("408-789-8075"));
```

TOPIC LIST

Security Enhancements: Column Level Encryption

- Example (Query)
  - Query

```
SELECT * FROM customer WHERE customer_num = 101

customer_num            101
fname       0tkn/AAAAEAndSyVxZpHZmA2/mCoJ6uZUHDSQ5I5u3l
lname       0V6j/AAAAEApupW+hft1mjw2CorFt7P9oWjO6cIjVzs
company     0Mc7/AAAAIAJbjI7Ul6N4oHMgIfh5wGo3559mnm8dIwznU2C+eivERQYUKse1WhoQ==
address1    0qVb/AAAAIAXoyVelIDxAWV8MqsX8mTdHDWCpG6A1bjGGgPdYioigDAz3/OdDinHw==
address2    0AP//AAAAAAe+M+wFglsPA=
city        Sunnyvale
state       CA
zipcode     94086
phone       1p1j/AAAAEAA1viBow2uu+7fLLJbEtuMucwiv9q0Mmx

SELECT customer_num, DECRYPT_CHAR(fname, 'my password') as fname,
      DECRYPT_CHAR(lname, 'my password') as lname FROM customer
      WHERE customer_num = 101

customer_num        101
fname                       Ludwig
lname                       Pauli
```

TOPIC LIST

Security Enhancements: Column Level Encryption

- Example (view)
  - Create a view

```
SET ENCRYPTION PASSWORD null;
CREATE VIEW custview (customer_num, firstname, lastname, company,
    city)
    AS SELECT customer_num, DECRYPT_CHAR(fname),
        DECRYPT_CHAR(lname), DECRYPT_CHAR(company), city
    FROM customer
```

  - Query from view

```
SET ENCRYPTION PASSWORD 'my password';
SELECT * FROM custview;
        customer_num            101
        firstname       Ludwig
        lastname        Pauli
        company         All Sports Supplies
        city            Sunnyvale
```

TOPIC LIST

Security Enhancements: Column Level Encryption

Estimating Space Requirements for encrypted data

Use the length() function:

Examples:

```
execute function
   length(encrypt_aes('1234-4321-2468-8642', 'aubergine', 'favorite vegetable'));

(expression)
        99
```

```
execute function length(encrypt_tdes('1234-4321-2468-8642', 'aubergine'));

(expression)
        55
```

It is critical you do this correctly!!!

TOPIC LIST

IDS 10 Features:
Security

Secure Environment Check

TOPIC LIST

Security: Secure Environment Check

- Server utilities on UNIX now check if the environment is secure by testing for the following:
  - The permissions on **$INFORMIXDIR** and some directories under it are correct. For each directory, check that the directory exists, is owned by user **informix** and the correct group, and that its permissions do not include write permissions for the group or other users.
  - The permissions on the ONCONFIG file are correct. The file must belong to the DBSA group. If the DBSA group is group informix (default), then the ONCONFIG file should be owned by user **informix** too; otherwise, the ownership is not constrained. The file must not have write permissions for others.
  - The permissions on the **sqlhosts** file are correct. Under a default configuration, the **sqlhosts** file is **$INFORMIXDIR/etc/sqlhosts**; the owner should be user **informix**, the group should be either the **informix** group or the DBSA group, and there should be no public write permissions. If the file is specified by setting the **INFORMIXSQLHOSTS** environment variable, then the owner and group are not checked, but public write permissions are not permitted.
  - The length of both the filenames **$INFORMIXDIR/etc/onconfig.std** and **$INFORMIXDIR/etc/$ONCONFIG** must be less than 256 characters.

TOPIC LIST

IDS 10 Features:
Security

Trigger Introspection

TOPIC LIST

Security: Trigger Introspection

- You can create user-defined routines that are invoked in trigger action statements
  - to obtain information about the triggers, triggering tables, views, statements, and the values of rows involved in the trigger actions.
  - Using the new DataBlade API routines, you can write a general purpose user-defined routine that can you can use to audit any table and any trigger event

TOPIC LIST

IDS 10 Features:
Security

Datablade Registration Restriction

TOPIC LIST

Security: Restricting Registration of DataBlade Modules and UDRs

- The DBSA (Database Server Administrator) can use a new built-in role, called EXTEND, to specify which users can register UDRs that include the EXTERNAL NAME clause.
- User-defined routines use shared-object files that are external to the database server and that could potentially contain harmful code.
- The DBSA can disable this feature by setting to “off” a new IDX_EXTEND_ROLE configuration parameter. This feature is intended to improve security and to control accessibility.

TOPIC LIST

IDS 10 Features
Security

PAM Authentication

TOPIC LIST

Security: PAM Authentication

- PAM – Pluggable Authentication Modules
  - It is a standardized system for allowing the OSA (Operating System Admininistrator) to configure how authentication is done.
  - Allows OSA to configure authentication methods.
  - Available on Linux, AIX, Solaris, HP-UX and others.
  - configured at the Operating System level.
  - APIs to write shared object (.so) at wws.sun.com/software/solaris/pam
  - PAM supports challenge-response protocols:
    - In response to initial authentication request,
    - PAM issues a challenge,
    - And waits for response from application.

TOPIC LIST

IDS 10 Features
Security

TOPIC LIST

IDS 10 Features
Administration & Manageability

TOPIC LIST

Topics : Administration & Manageability

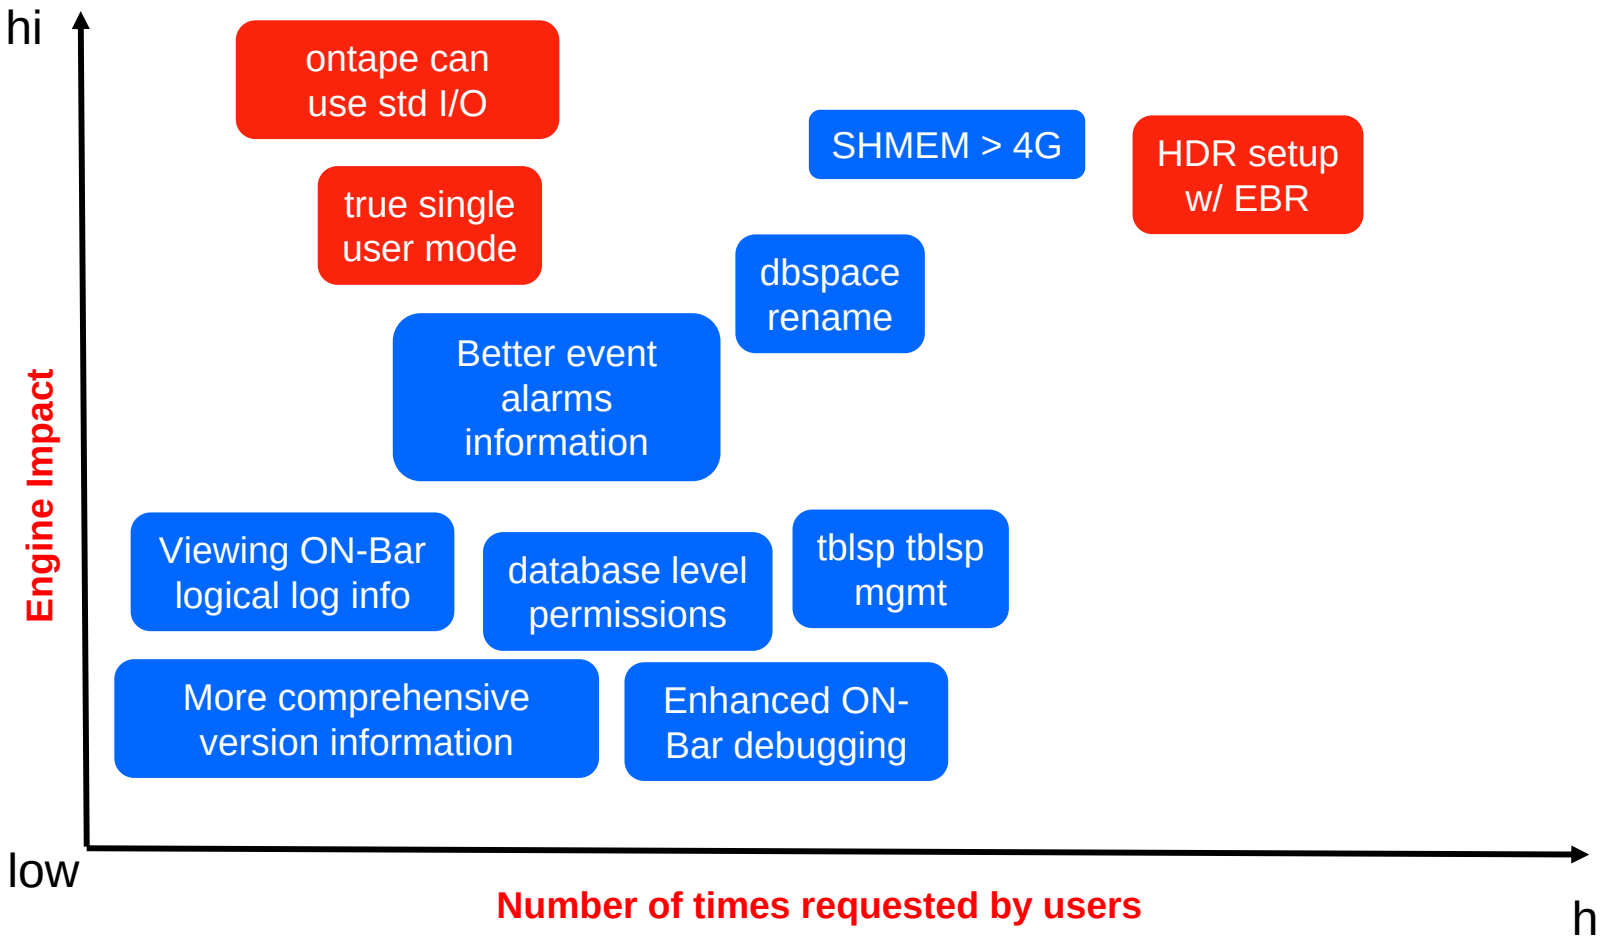

TOPIC LIST

IDS 10 Features
Administration

HDR Setup with EBR

TOPIC LIST

Administration: HDR Setup with EBR

HDR – High Availability Data Replication

- entire engine replicated.
- used for hot backups typically.
- used also for not-realtime report server.

EBR – External Backup and Restore

- a “non-IDS” backup to media of DBA’s choice
- added many releases ago to allow clients to, for example, “break off a mirror” for the archive/backup.

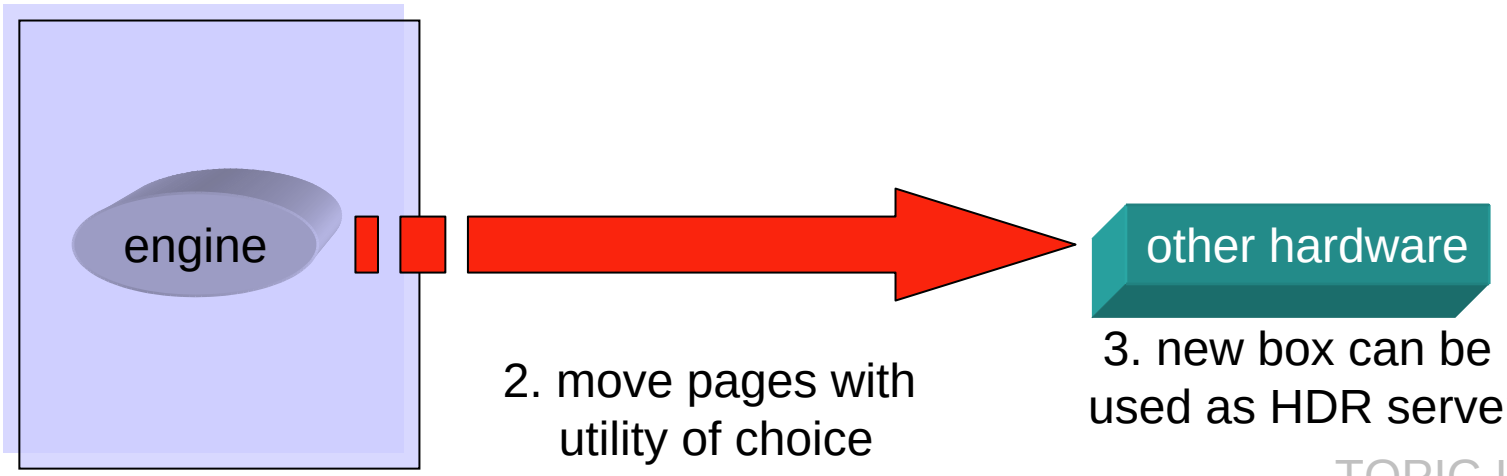

TOPIC LIST

IDS 10 Features
Administration

Single User Mode

TOPIC LIST

Administration : **Single User Mode**

Overview

- a mode intermediate between Quiescent mode & Online mode
- allows only the user ‘informix’ to connect and do any required maintenance

Utility Enhancements

- **oninit -j** … brings the engine from *offline* to *single-user mode*.
- **onmode -j** … brings the engine from *online* to *single-user mode*.

```
informix:/export/home/informix> onmode -j
This will change mode to single user. Only user informix can connect
in this mode.
Do you wish to continue (y/n)? y

All threads which are not owned by user informix will be killed.
Do you wish to continue (y/n)? y
informix:/export/home/informix> onstat -

IBM Informix Dynamic Server Version 9.50.UC1B5    -- Single-User -- Up

informix:/export/home/informix>
```

LIST

IDS 10 Features
Administration

Renaming Dbspaces

TOPIC LIST

Administration: Renaming Dbspaces

Description

The ability to change the name of a previously defined dbspace.

The Problem

- Existing customers who wish to reorganize their data are moving their data to a new dbspace, then reloading the data back into the original dbspace to regain the original dbspace name. The last reload step could be avoided by using a rename dbspace option. Helps in Recycling of Dbspaces.
- Time consuming operations as reorganizing the data in an existing dbspace can benefit from this feature.

- **The rename dbspace operation only changes the dbspace name; it does not reorganize data**

TOPIC LIST

Administration : Renaming Dbspaces

9.40.UC3

- Feature was first introduced with limitations:
  - Could not rename blobspaces, sbspaces, temporary, or external spaces.
  - Could not rename dbspaces that are referred to by the following configuration parameters:

    DBSPACETEMP, CDR_DBSPACE, SBSPACENAME, SBSPACETEMP, SYSSBSPACENEAME, CDR_QHDR_DBSPACE, CDR_QDATA_SBSPACE

  - Could not rename dbspaces if you are using High-Availability Data Replication.
  - Could not rename dbspaces when Enterprise Replication is active.

TOPIC LIST

Administration : Renaming Dbspaces

v10.0 Restrictions:

- Rename can not be done on critical spaces (i.e.) root dbspace, space containing physical log or logical logs.
- A dbspace with down chunks can not be renamed.
- Rename of spaces cannot be done with onmonitor.

TOPIC LIST

Administration : Renaming Dbspaces

Note:

- A level 0 archive of the renamed space and root dbspace need to be taken after renaming.
- Rename done on HDR primary will propagate to the secondary.

TOPIC LIST

Rename Dbspace Example

```
informix:/export/home/informix> onspaces -c -d skippy -k 16 -p /dev/rdsk/stu414B -o 0 -s 100000
Verifying physical disk space, please wait ...
Cannot build a new Space.
ISAM error: Cannot add dbspace of big page when Large Chunk support is disabled.

informix:/export/home/informix> onmode -BC 1
This command will enable creation of large chunks.
** WARNING **  This action cannot be undone.
** WARNING **  A level 0 archive of Root DBSpace will need to be done.
Do you wish to continue (y/n)? y

Expanded chunk capacity mode: enabled
informix:/export/home/informix> onmode -BC 2
This command will cause all chunks to be written in the new (big) format.
** WARNING **  This action cannot be undone.
** WARNING **  A level 0 archive of Root DBSpace will need to be done.
Do you wish to continue (y/n)? y

Expanded chunk capacity mode: always
informix:/export/home/informix> onspaces -c -d skippy -k 16 -p /dev/rdsk/stu414B -o 0 -s 100000
Verifying physical disk space, please wait ...
Space successfully added.

** WARNING **  A level 0 archive of Root DBSpace will need to be done.
informix:/export/home/informix>
```

TOPIC LIST

Administration : Rename Dbspace

```
informix:/export/home/informix> onstat -d

IBM Informix Dynamic Server Version 9.50.UC1B5   -- On-Line -- Up 00:02:17 -- 1013760 Kb

Dbspaces
address  number   flags      fchunk   nchunks  pgsize   flags    owner    name
ad277d8  1        0x60001    1        1        2048     N  B     informix rootdbs
c374410  2        0x60001    2        1        16384    N  B     informix skippy
 2 active, 2047 maximum

Chunks
address  chunk/dbs  offset     size       free       bpages     flags pathname
ad27928  1     1    0          50000      40212                 PO-B  /dev/rdsk/stu416A
c39e1a0  2     2    0          6250       6197                  PO-B  /dev/rdsk/stu414B
 2 active, 32766 maximum

NOTE: The values in the "size" and "free" columns for DBspace chunks are
      displayed in terms of "pgsize" of the DBspace to which they belong.

Expanded chunk capacity mode: always

informix:/export/home/informix>
```

TOPIC LIST

Rename Dbspace

```
informix:/export/home/informix> onmode -s

This will perform a GRACEFUL SHUTDOWN -
Do you wish to continue (y/n)? y
informix:/export/home/informix> onspaces -ren skippy -n sparky
Rename of Space completed successfully.

** WARNING **  A level 0 archive of Root DBSpace and the renamed
Space need to be done.
informix:/export/home/informix> onstat -d

IBM Informix Dynamic Server Version 9.50.UC1B5   -- Quiescent -- Up 00:04:07 -- 101376

Dbspaces
address  number   flags      fchunk   nchunks  pgsize   flags    owner    name
ad277d8  1        0x60001    1        1        2048     N  B     informix rootdbs
c374410  2        0x60001    2        1        16384    N  B     informix sparky
 2 active, 2047 maximum
```

TOPIC LIST

IDS 10 Features:
Administration

Database-Level Permissions

TOPIC LIST

Administration : Per Database Permissions

- You can create a default role and assign that role to individual users or to PUBLIC on a per-database level.
- Each user who is assigned to a default role receives the privileges of that role as well as whatever other privileges are granted to the user individually.
- The syntax of the GRANT, REVOKE, and SET ROLE statements support this feature.

TOPIC LIST

IDS 10 Features:
Administration

ontape using STDIO

TOPIC LIST

Administration: ontape use of standard I/O

- **ontape** can now use standard I/O instead of a tape device or disk file.
- Specifying stdout or stdin allows **ontape** to use pipes or programs for archives and restores.
  - For example, you can use compression to save media space, use cloning to duplicate the archive for safety reasons, or restore the data onto another server instance.
  - This feature is especially efficient for setting up High-Data Availability Replication by restoring the data to the secondary server while skipping the intermediary step of saving the data to a file or disk.

TOPIC LIST

IDS 10 Features
Administration

Tablespace Tablespace Management

TOPIC LIST

Administration : Managing the Tblspace Tblspace

- Overview
  - This feature changes the behavior of the sizing of the tblspace tblspace. It also changes the behavior of extending the tblspace tblspace.

- Description
  - A DBA will now be allowed to specify the first and next extent sizes for the tblspace tblspace.
  - When extending the tblspace tblspace preferences will be made when choosing the chunk for the next extent.

TOPIC LIST

Administration : Managing the Tblspace Tblspace

- Tblspace == partition
- TBLspace TBLspace
  - Special table that tracks other tables within a dbspace.
  - Every dbspace has its own tblspace tblspace that tracks tables (partitions) within its own dbspace.
  - There is only **one** tblspace tblspace per dbspace. But as with any tblspace it can have multiple extents. (see next slide)
  - Partnum is **0xDDD00001**, where DDD = dbspace number.

TOPIC LIST

Multiple Extents – 1 Tblspace Tblspace

Chunk Pathname	Size	Used	Free
2 /chunks/94/dbs1	10000	103	9897

Description	Offset	Size
RESERVED PAGES	0	2
CHUNK FREELIST PAGE	2	1
dbs1:'informix'.TBLSpace	**3**	**50**
FREE	53	37
dbs1:'informix'.TBLSpace	**90**	**50**
FREE	140	9860

Total Used: 103
Total Free: 9897

TOPIC LIST

Problem: Can’t drop chunk

Situation.

1. Customer adds a chunk temporarily to a dbspace to create a large table in the dbspace.
2. During processing the tblspace tblspace extends into the newly added chunk.
3. The customer drops the large table and now wants to drop the chunk.

TOPIC LIST

Resolution:

- Drop Dbspace
  - This causes the customer to have to unload the entire dbspace and drop all tables in the dbspace, then drop and recreate the dbspace.
  - This creates an outage

- Dialin To drop Chunk
  - Another possible solution was to have Advanced Support dial in and drop the chunk. They would also have to clean up the tblspace tblspace. This caused down time and is risky. (usually avoided)

TOPIC LIST

The feature

- Specify First/Next Extent sizes
  - When creating a dbspace you will now have the option to specify the first and next extent sizes.
  - When creating the instance (oninit –iy) you will be able to specify the first and next extent sizes of the root dbspace.

- New Algorithm for Extending Tblspace Tblspace
  - When creating an additional extent for a tblspace tblspace more emphasis is placed on the location of the extents.

TOPIC LIST

First/Next Extent Sizes – Root dbspace

- Root dbspace (oninit –iy)
  - TBLTBLFIRST
    - This **onconfig parameter** is used to specify the size of the first extent of the tblspace tblspace for the root dbspace.
  - TBLTBLNEXT
    - This **onconfig parameter** is used to specify the next extent size of the tblspace tblspace for the root dbspace.
  - If these parameters are not present, the defaults will be used. Defaults will be discussed later. **All sizes are in KB**, and must be a multiple of the page size.

TOPIC LIST

First/Next Extent Sizes – non root

- Non-Root dbspace (onspaces)
  - When creating a new dbspace with onspaces you will now be allowed to specify the first and next extent sizes for the tblspace tblspace.
    - -ef <first extent size> –en <next extent size>
  - If these options are not used the defaults will be used. Defaults will be discussed later. **All sizes are in KB**, and must be a multiple of the page size.

Example

```
onspaces –c –d dbs1 –p /spare2/dbs1.1 –o 0  \
        –s 10000 –ef 150 –en 150
```

TOPIC LIST

Reversion

- This feature will not prevent the reversion of any dbspaces.
- It will now be possible to revert back to a version prior to this feature, and now have extents for a tblspace tblspace that are not a default value or a doubling of that.
  - The next extent sizes will be adjust accordingly on the reversion process

TOPIC LIST

IDS 10 Features
Administration

Multiple Fragments in a Single DBSPACE

TOPIC LIST

Administration : Multiple Table Fragments in Single Dbspace

- You can create partitions within a dbspace that can each support a table fragment.
  - reduces the total number of dbspaces needed for a fragmented table.
- Storing multiple table fragments in a single dbspace improves query performance over storing each fragmented expression in a different dbspace.
- This feature improves performance and simplifies management of dbspaces.

TOPIC LIST

IDS 10 Features
Administration

SHMEM > 4G

TOPIC LIST

Administration: Shared Memory > 4 GB

- You can now specify that segments for shared memory be created as large as:
  - your operating system platform, or
  - the SHMMAX parameter allows.

TOPIC LIST

IDS 10 Features:
Administration

Misc Backup and Restore

TOPIC LIST

IDS v10.0 Backup and Restore Enhancements

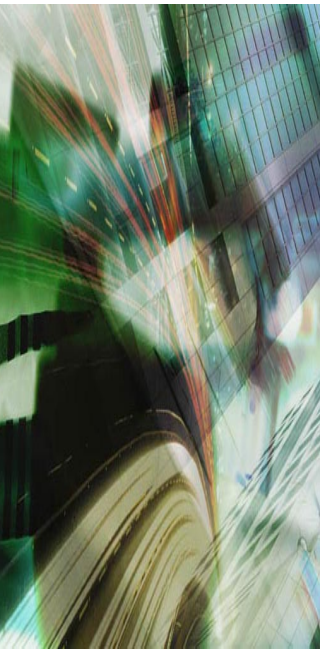

- View logical logs backed up by OnBAR
  - Similar to using OnLog utility
- OnBAR debugging level can be changed while OnBar is running
  - Saves time and disk space
  - Bar_Debug can be set as frequently as needed
- Tivoli Storage Manager XBSA is included with IDS

TOPIC LIST

IDS 10 Features
Administration & Usability

TOPIC LIST

IDS 10 Features
Availability

TOPIC LIST

Topics : Availability Enhancements

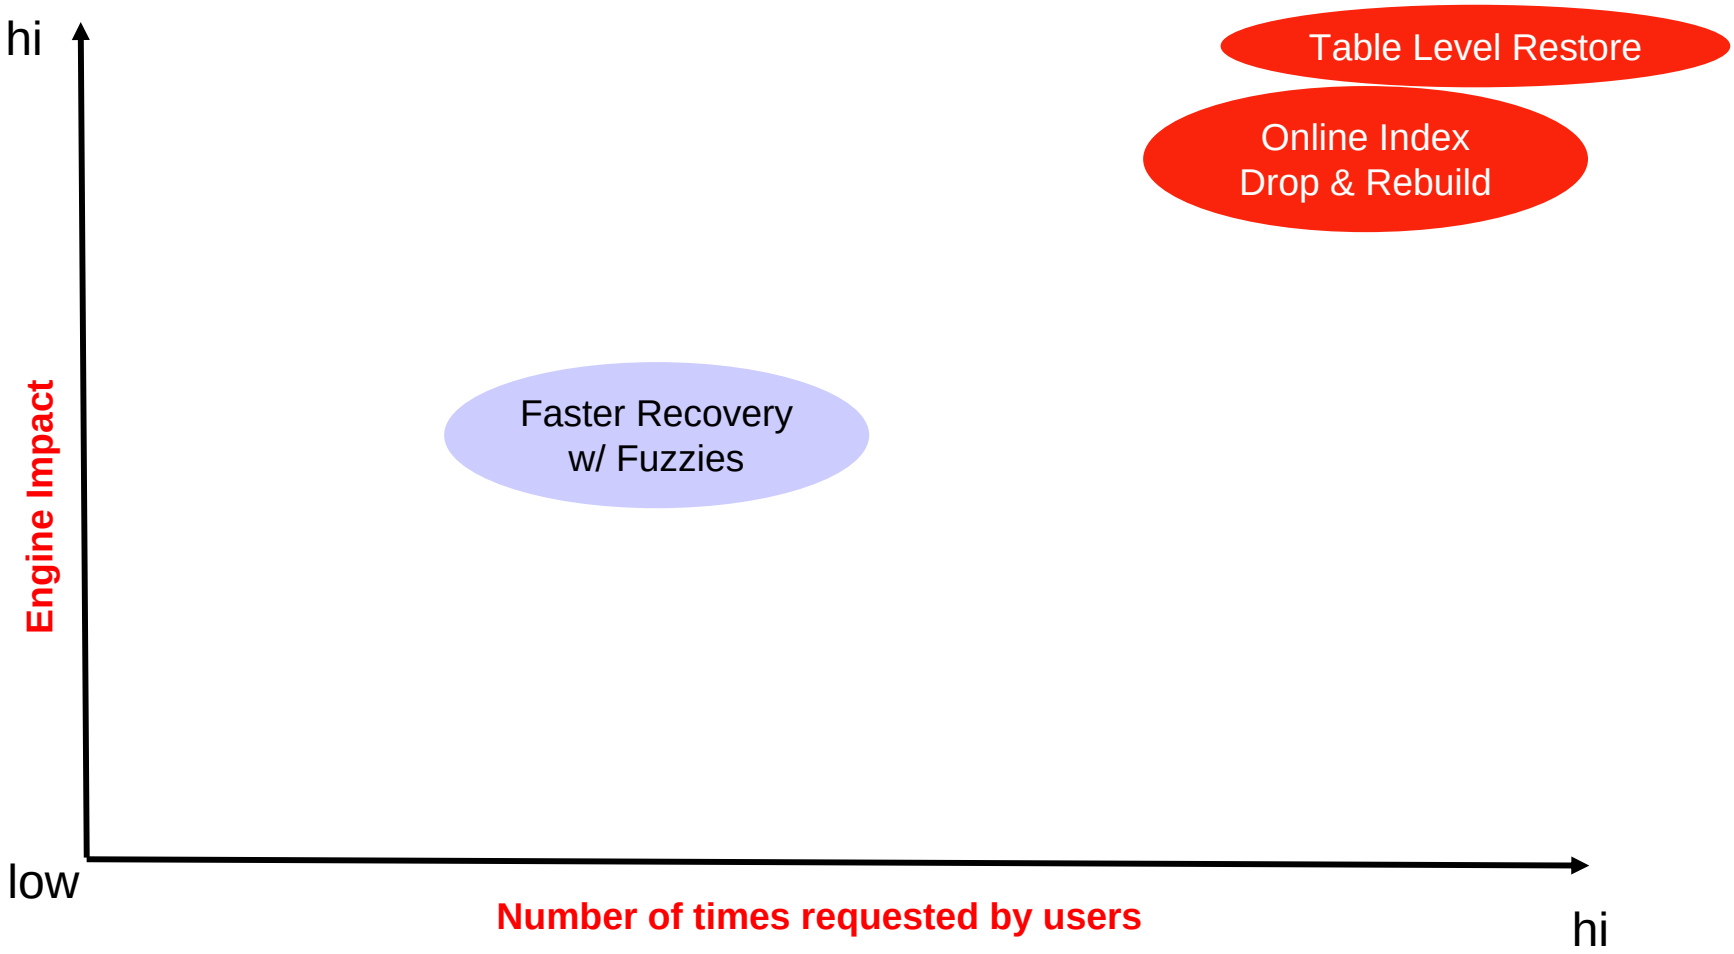

TOPIC LIST

IDS 10 Features
Availability

Online Index Rebuild

TOPIC LIST

Availability: Index Changes

- CREATE INDEX and DROP INDEX now supports DDL operations that apply no exclusive lock to the table on which the specified index is defined.
  - If you use this syntax to create an index on a table that other users are accessing, the index is not available until no user is updating the table.
  - After you issue the new syntax to drop an index, no one can reference the index, but concurrent DML operations can use the index until they terminate.
  - Dropping the index is deferred until no user is using the index.
  - This feature maintains the availability of the table within a production environment after an existing index has ceased to be efficient.

TOPIC LIST

IDS 10 Features
Availability

Table Level Restore

TOPIC LIST

Table Level Point-in-Time Restore (TLR)

- Purpose…to allow restoration of table data from a level 0 archive to a user specified point in time.
- Data may be restored to *any engine version or machine.*
- SQL Driven Distributed Restore (SDDR)
  - Extract a table or set of tables
  - You may apply a filter to the retrieved data
  - Retrieve just a subset of the columns
  - Repartitioning of the data

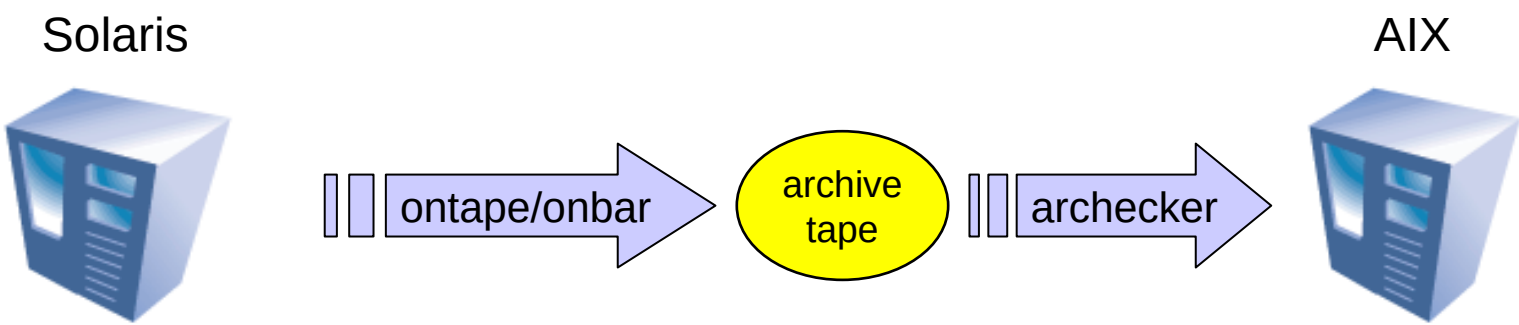

TOPIC LIST

Physical Restore

- Extracts requested data from a level 0 archive
- Two temporary tables create to store partial rows (i.e. rows not contained on a single page)

Physical Restore Flow

- The list of dbspaces is built from the create table statements in the schema command file
  - For onbar - only the dbspaces containing the table's data are scanned.
  - For ontape – currently all dbspaces scanned
- The data pages from the tables being restored are processed
  - Rows are extracted from the data pages
  - Converts to SQL Inserts or unloads format

TOPIC LIST

Logical Recovery

- Replays the logical logs required to bring the data from the level 0 archive point in time(s) to the DBA requested restore point in time.

- Two parts to logical recovery
  - LOG STAGER
    - Reads the logical logs file(s)
    - Filters out unwanted records
    - Insert the required log records into log staging tables
  - LOG APPLIER
    - Reads the log staging tables
    - Applies the records to the destination tables

TOPIC LIST

Schema Command File Example

```
database test1;
create table tlr     (
       a_serial    serial,
       b_char      char(20)
     ) in dbspace1;
create table tlr_dest     (
       a_serial    serial,
       b_char      char(20)
    ) in dbspace2;

insert into tlr_dest select * from tlr;

set workspace to dbs1,dbs2;

restore to '2003-01-01 01:01:01';
```

LIST

Data Types Not Support in the First Release

Extended Data Type
LISTS
MULTISET
SET
ROW
DISTINCT
OPAQUE
Built-in Data Type

- This mean these data types can not be in any of tables being restored
- If these data types do appear an error will be given when processing the command file

TOPIC LIST

Restore command

- Point in time to which the data is restored
  - A quoted time
  - To restore to the most recent time use the keyword *CURRENT*
- To extract data only from a level 0 archive use the "*NO LOG RESTORE*"
- If you omit the entire statement the default is "*RESTORE TO CURRENT*"

COMMAND	REASON
RESTORE [TO "timestamp" | CURRENT] [NO LOG RESTORE]	**Set the time the restore should use and if logical logs should be used.** **NO LOG RESTORE causes the data to be extracted from the level 0 archive only**

TOPIC LIST

Example Screen Output

```
archecker -bdvsX -f SETUP/setup_basic.cmd
AC_STORAGE              /tmp
AC_MSGPATH              /tmp/ac.log
AC_VERBOSE              on
AC_TAPEBLOCK             32 KB
AC_PAGESIZE             4096
AC_IXBAR                /vobs/tristarm/sqldist/etc/Ixbar_olympia.85
Dropping old log control tables
Extracting table d1:t1 into d2:t1n
Extracting table d1:t2 into d2:t2n
Scan PASSED
Control page checks PASSED
Table checks PASSED
Table extraction commands 2
Tables/fragments found on archive 2
LOADED: d2:t2n produced 37 rows.
LOADED: d2:t1n produced 36 rows.
Creating log control tables
Staging Log 1:14
Staging Log 1:15
Logically recovered d2:t1n Inserted 602 Deleted 19 Updated 0
Logically recovered d2:t2n Inserted 0 Deleted 9 Updated 0
```

Tables being extracted

Rows extracted from Physical Restore

Logical Logs Processed

Logical Recovery Actions

TOPIC LIST

Simple Schema Command File - Recovery of a Lost Table

- Extracts a table called *test1:tlr* from the most recent backup of *dbspace1* and places the data in table *test1:tlr*

```
database test1;
create table tlr     (
       a_serial    serial,
       b_integer   integer,
       c_char      char(20),
       d_decimal   decimal,
       ) in dbspace1;
insert into tlr
  select * from tlr;
```

TOPIC LIST

Restoring a Table from a Previous Backup

- Extracts a table called *test1:tlr* from the level 0 backup of *dbspace1* just prior to “2003-01-01 01:01:01” and places the data in table *test1:tlr.*
- The logical logs will be applied up to “2003-01-01 01:01:01”

```
database test1;
create table tlr    (
       a_serial   serial,
       b_integer  integer,
       c_char     char(20),
       d_decimal  decimal,
       ) in dbspace1;
insert into tlr
  select * from tlr;
restore to
    ‘2003-01-01 01:01:01’;
```

TOPIC LIST

Distributed Restore

- Extracts a table called *test1:tlr_1* from the most recent backup of *dbspace1* and places the data on the database server *rem_srv* in the table *rem_dbs:tlr_1*

```
database rem_dbs;
create table tlr_1 (
   ( columns  )

database test1;
create table tlr_1
  (  columns  ) in dbspace1;

insert into rem_dbs@rem_srv.tlr_1
  select * from tlr_1;
```

TOPIC LIST

Using External Tables

- Extracts a table called d1:source_tab from the most recent backup of dbspace1 and send the data in ASCII format with field delimited to the file named /tmp/pipe

```
database d1;
create table source_tab
   ( cola int) in dbspace1;


create external table target_tab
  (  cola int  )
  USING ('/tmp/pipe', delimited');


insert into target_tab
  select * from source_tab;
```

TOPIC LIST

Product Limitations

- The table schemas must be exact
- Since one of the goals of this feature is to recover accidentally dropped tables
  - Any drop of the fragment/table will be ignored during logical recovery
  - It will indicate the end of logical recovery for the table/fragment
- The tables/fragments being recovered must exist on the level 0 archive
- Stop logical recovery if an IPA has been performed since the level 0 restore
- A detached fragment is no longer part of the original table so recovery will be terminated on the detached fragment log record
- External tables are only physically restored
- Filters are only applied to tables which are physically restored

TOPIC LIST

IDS 10 Features
Availability

Faster Recovery w/ Fuzzy Checkpoints

TOPIC LIST

Availability: Recovering Quickly with Fuzzy Checkpoints

- Two new configuration parameters
  - FAST_RESTART_PHYSLOG
  - FAST_RESTART_CKPT_FUZZYLOG

- Reduces the time required for engine recovery. This supports high availability by improving recovery performance when using fuzzy checkpoints.

TOPIC LIST

IDS 10 Features
Replication

TOPIC LIST

Topics : Replication Enhancements

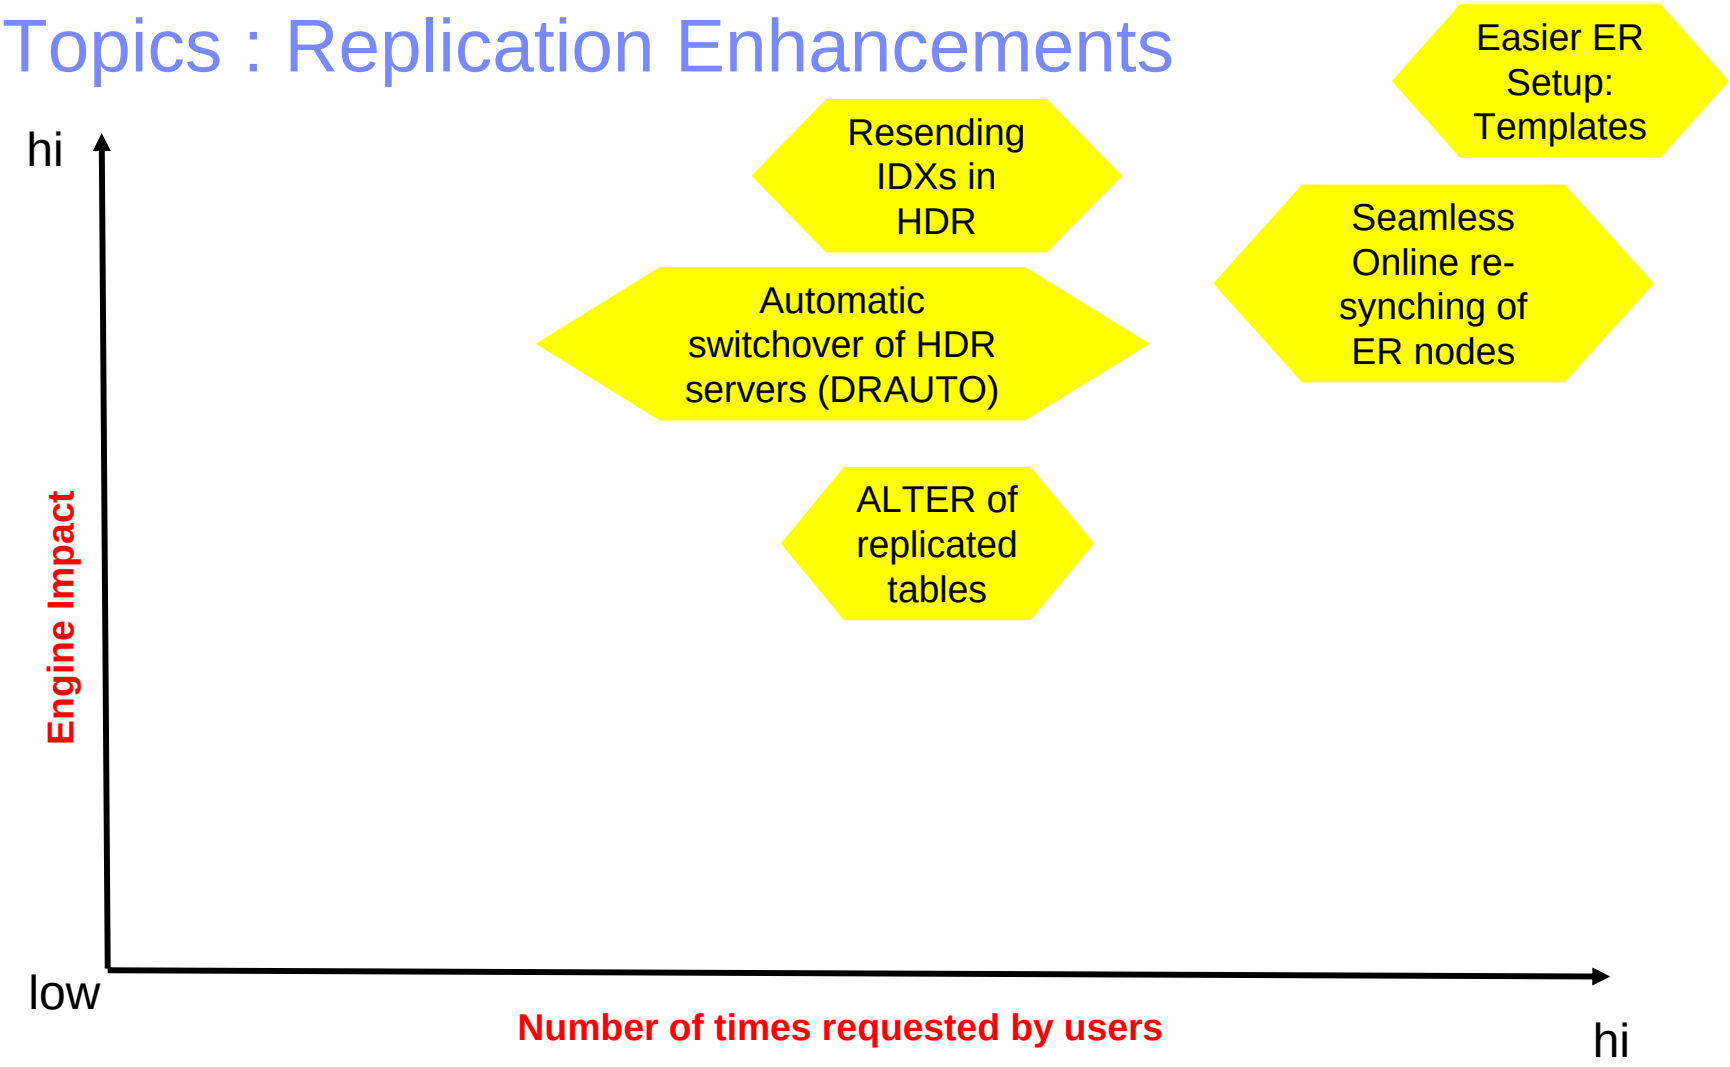

TOPIC LIST

IDS 10 Features
Replication

Fixing Corrupt Secondary Indexes

TOPIC LIST

HDR: Resending Indexes

- You can resend an index that became corrupt on the secondary server in an HDR pair.
- Resending an index is quicker than dropping and then rebuilding the index on the primary server.
- This feature increases the availability of the HDR primary server.

TOPIC LIST

IDS 10 Features
Replication

Alter Fragment Support

TOPIC LIST

ER Enhancements: Alter table/fragment support

- Overview
  - This feature provides alter support for tables being replicated via Enterprise Replication.
  - Currently, if the table schema needs to be altered or if the fragmentation strategy needs to be changed, then replication must be stopped, then alter is performed and then replication must be restarted. This is problematic as it makes it impossible to really consider ER in a 24X7 environment.
  - List of supported alter operations:
    - The ability to add/drop default values
    - The ability to add/drop SQL checks
    - The ability to add/drop fragments
    - The ability to attach/detach fragments
    - The ability to add/drop columns
    - The ability to recluster indexes
    - The ability to alter non replicated columns

TOPIC LIST

ER Enhancements: Alter table/fragment support

- Mastered Replicates
  - Currently no way to know if data types match between replicate nodes
  - Provides data type checking to eliminate possibility of corruption
  - New syscdr tables track data type information

TOPIC LIST

ER Enhancements: Alter table/fragment support

- Remastering process
  - Existing replicate can be redefined by “remastering the replicate”. Through remastering process, a new column can be added/dropped to/from a replicate definition.
  - Also an existing non-mastered replicate can be converted to a mastered replicate using remastering process.

TOPIC LIST

ER Enhancements: Alter table/fragment support

- Restrictions
  - ER must be in active state for altering a replicated table except in scenarios where adding/dropping check constraints and default values.
  - Alter operations are supported only on tables defined with mastered replicates.
  - Altering a replicated column size or data type is not supported. (Note: Modifying a non replicated column is supported).
  - Rename table operation is not supported
  - Rename column operation is not supported
  - Drop table operation is not supported

TOPIC LIST

IDS 10 Features
Replication

Replicate Templates

TOPIC LIST

ER Enhancements: Replicate Templates

- Overview
  - Ease of ER administration and setup
    - The entire Enterprise Replication domain can be setup using simple commands with options such as defining replicates on all tables within database, specified on the command line or using a input file.
    - A template can perform an initial data synchronization on new servers being added to a template..
    - A template can optionally create tables during realization if they do not exist on target servers during template realization.
  - Eliminates most of the table Schema related errors
    - Templates use the master dictionary from the Master node to create these tables to ensure consistent schemas between the nodes.

TOPIC LIST

IDS 10 Features:
Replication

DRAUTO

TOPIC LIST

HDR: DRAUTO

- You can automate switching servers for High-Availability Data Replication if the primary server fails by using the DRAUTO configuration parameter.

- If DRAUTO is set to either
  - RETAIN_TYPE or REVERSE_TYPE, the secondary database server switches to type standard automatically when an HDR failure is detected.
    - RETAIN_TYPE, the original secondary database server switches back to type secondary when the HDR connection is restored.
    - REVERSE_TYPE, the original secondary database server switches to type primary when the HDR connection is restored, and the original primary switches to type secondary.

TOPIC LIST

IDS 10 Features
Replication

Replicate Resynch

TOPIC LIST

ER Enhancements: Replicate Resynch

Overview

- Meant to be used:
  - Bring a newly participating table up-to-date with the ongoing replication
  - Repair a replicated table if replication was stopped or failed for some reason.

- Two ways to repair a table:
  - Quick way: By processing the ats/ris files.
  - Other way: By defining and running a ‘resynch job’.

TOPIC LIST

ER Enhancements: Replicate Resynch

Details of a resynch job:

- Creates a shadow replicate to carry the dummy updates as well as the forwarded data.
- Puts the replicate on target in ‘skip’ state – data coming from participants other than the
  source is skipped for the duration of the resynch.
- Puts the replicate on source in ‘forward’ state – data coming from other participants is forwarded to the target via ‘shadow replicate’.
- Creates three internal tables which are replicated between source and target.
- Generates stored procedures and triggers to:
  - Scan the source data
  - Handle the extra target row options – including cascade deletes if the option is delete
  - Cleanup the rows in the internal tables and populate the violations table for errors.
  - Do dummy updates on the source data to replicate to the target.

TOPIC LIST

ER Enhancements: Replicate Resynch

Listing the details of a repair job

$ cdr list resynch

```
RESYNCHJOB          REPLICATE/REPLSET                    STATE
-------------------------------------------------------------------------------------------
Repair_acct         Acct_repl                            Completed
Repair_txns         Txns_repl                            Defined
```

$ cdr list resynch Repair_acct

```
RESYNCHJOB          REPLICATE/REPLSET                    STATE
-------------------------------------------------------------------------------------------
Repair_acct         Acct_repl                            Completed

SOURCE
------
test@g_serv1:informix.account
      select * from 'informix'.account

TARGET
-------
test@g_serv2:informix.account
      select col1,col2 from 'informix'.account

BLOCK SIZE:                 10
TARGET ROW OPTION:          Delete
PROCESSED ROWS:             100
START TIME:                 2004-05-25 13:58:36
END TIME:                   2004-05-25 15:34:36
```

TOPIC LIST

IDS 10 Features
Application Development

TOPIC LIST

Topics : Application Development/Standards Enhancements

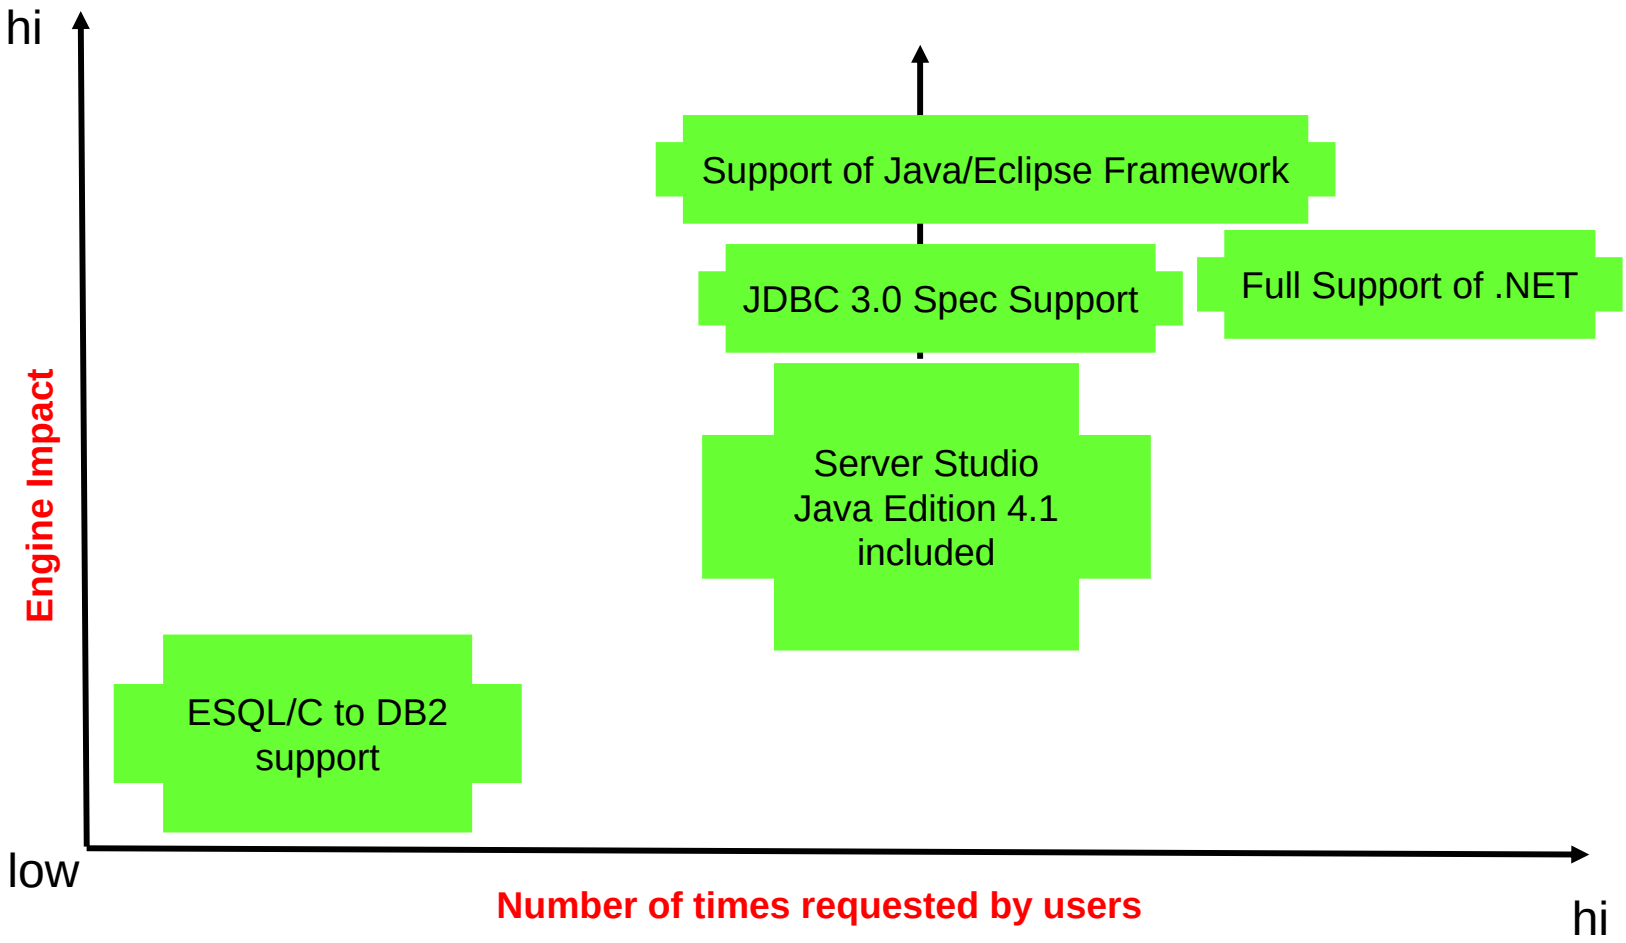

TOPIC LIST

IDS v10.0 Application Development

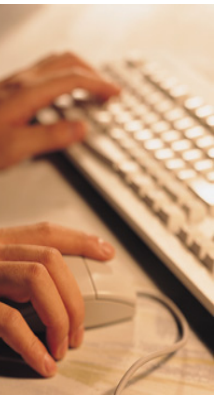

- Enterprise Generation Language (EGL) is a core technology in:
  - Rational Web Developer
  - Rational Application Developer
  - Rational Software Architect
  - WebSphere Application Server – Express
- EGL is ideal for Data-driven Rapid Application Development:
  - Web, Text, and Batch Applications
  - With IBM Databases (IDS v10.x and DB2 8.x)
- EGL generates Java at runtime
- Majority of Informix 4GL language constructs are included in the EGL
- Informix 4GL to EGL Conversion Utility is included with V6.0 products
  - GA with IDS v10, Q1 2005

TOPIC LIST

Applications: ESQL/C to DB2

- You can run Informix ESQL/C applications with DB2 servers and databases.
- The Informix ESQL/C product provides a new library that is called when you use the **esql** command to preprocess your files to work with DB2.
- Informix ESQL/C runs with DB2 Version 8.2, or later, running on Linux, UNIX, and Windows operating systems.

TOPIC LIST

Applications: JDBC 3.0 Support

- Version 3.0 of the IBM Informix JDBC Driver supports the following features in compliance with the Sun Microsystems JDBC 3.0 specification:
  - Internally update BLOB and CLOB data types using all methods introduced in the JDBC 3.0 specification.
  - Specify and control ResultSet holdability, leveraging on the Informix JDBC extension implementation.
  - Retrieve auto-generated keys from the database server.
  - Access multiple INOUT mode parameters in Dynamic Server through the CallableStatement interface.
  - Provide a valid large object descriptor and data to the JDBC client to send or retrieve BINARY data types as OUT parameters.
  - J/Foundation supports JRE Version 1.4 and the JDBC 3.0 specification.

TOPIC LIST

Applications: Full .NET Support

- The .NET Provider enables Windows .NET applications to access and manipulate data in IBM Informix databases.
  - The IBM Informix .NET Provider is a runtime library that encapsulates a data access API for use by Microsoft .NET applications.
  - It consists of a set of specialized classes that implement standard Microsoft ADO.NET interfaces and serves as a bridge between IBM Informix databases (data sources) and .NET applications.
- Windows client applications written in any .NET supported language can take advantage of the IBM Informix .NET Provider. Some examples of client applications are:
  - Visual BASIC .NET applications
  - Visual C# .NET applications
  - Visual J# .NET applications
  - ASP.NET web applications
- The IBM Informix .NET Provider a connection editor dialog box, a command editor dialog box, and a data adapter wizard as Microsoft Visual Studio add-ins.

TOPIC LIST

Helpful Links & Contacts

- Informix Instant Answers - E-mail: ibmifmx@us.ibm.com
- Informix Product Page - www.ibm.com/informix/
- Informix Connection Program (ICON) - www.informixconnect.com/icon/
- Informix Platform Roadmap - www.ibm.com/software/data/informix/pubs/roadmaps.html
- Informix Product Life Cycle (PLC) Roadmap - www.ibm.com/software/data/informix/support/plc/
- Informix International User Group (IIUG) - www.iiug.org
- DM Today Newsletter - www.ibm.com/software/data/db2infonews/ www.ibm.com/software/data/db2infonews/

TOPIC LIST

Helpful Links & Contacts

WW Informix Sales Support
Danilo Novelli
1-305-442-3604 T/L: 235-3604
spavlot@us.ibm.com

WW Informix Technical Enablement
Mark Scranton
1-303-395-3622 T/L: 273-3381
mark.scranton@us.ibm.com

IDS Brand Manager
Christine Normile
1-877-252-5399 T/L: 273-0981
Cnormile@us.ibm.com

IDS Product Manager
Frederick Ho
1-650-926-5824
hof@us.ibm.com

Informix Product Manager
Omkar Nimbalkar
1-650-926-1976 T/L: 630-1976
omkarn@us.ibm.com

Informix Program Manager
Bruce Weed
1-914-766-1717 T/L: 826-1717
bweed@us.ibm.com

Informix Direct Channel Manager
Christine Normile
1-877-252-5399 T/L: 273-0981
Cnormile@us.ibm.com

Informix Partner Channels Manager
Glenda Albertson
1-512-473-8123 T/L: 595-8123
gla@us.ibm.com

TOPIC LIST

Questions???

TOPIC LIST

TOPIC LIST