

O'REILLY®
EUROPEAN
OPEN
SOURCE
CONVENTION
17-20 OCTOBER 2005

PostgreSQL in the Enterprise

11:35 St.John's Room II

Wed 19 Nov 2005

Simon Riggs

simon@2ndquadrant.com


Magnus Hagander

mha@sollentuna.net

PostgreSQL Global Development Group

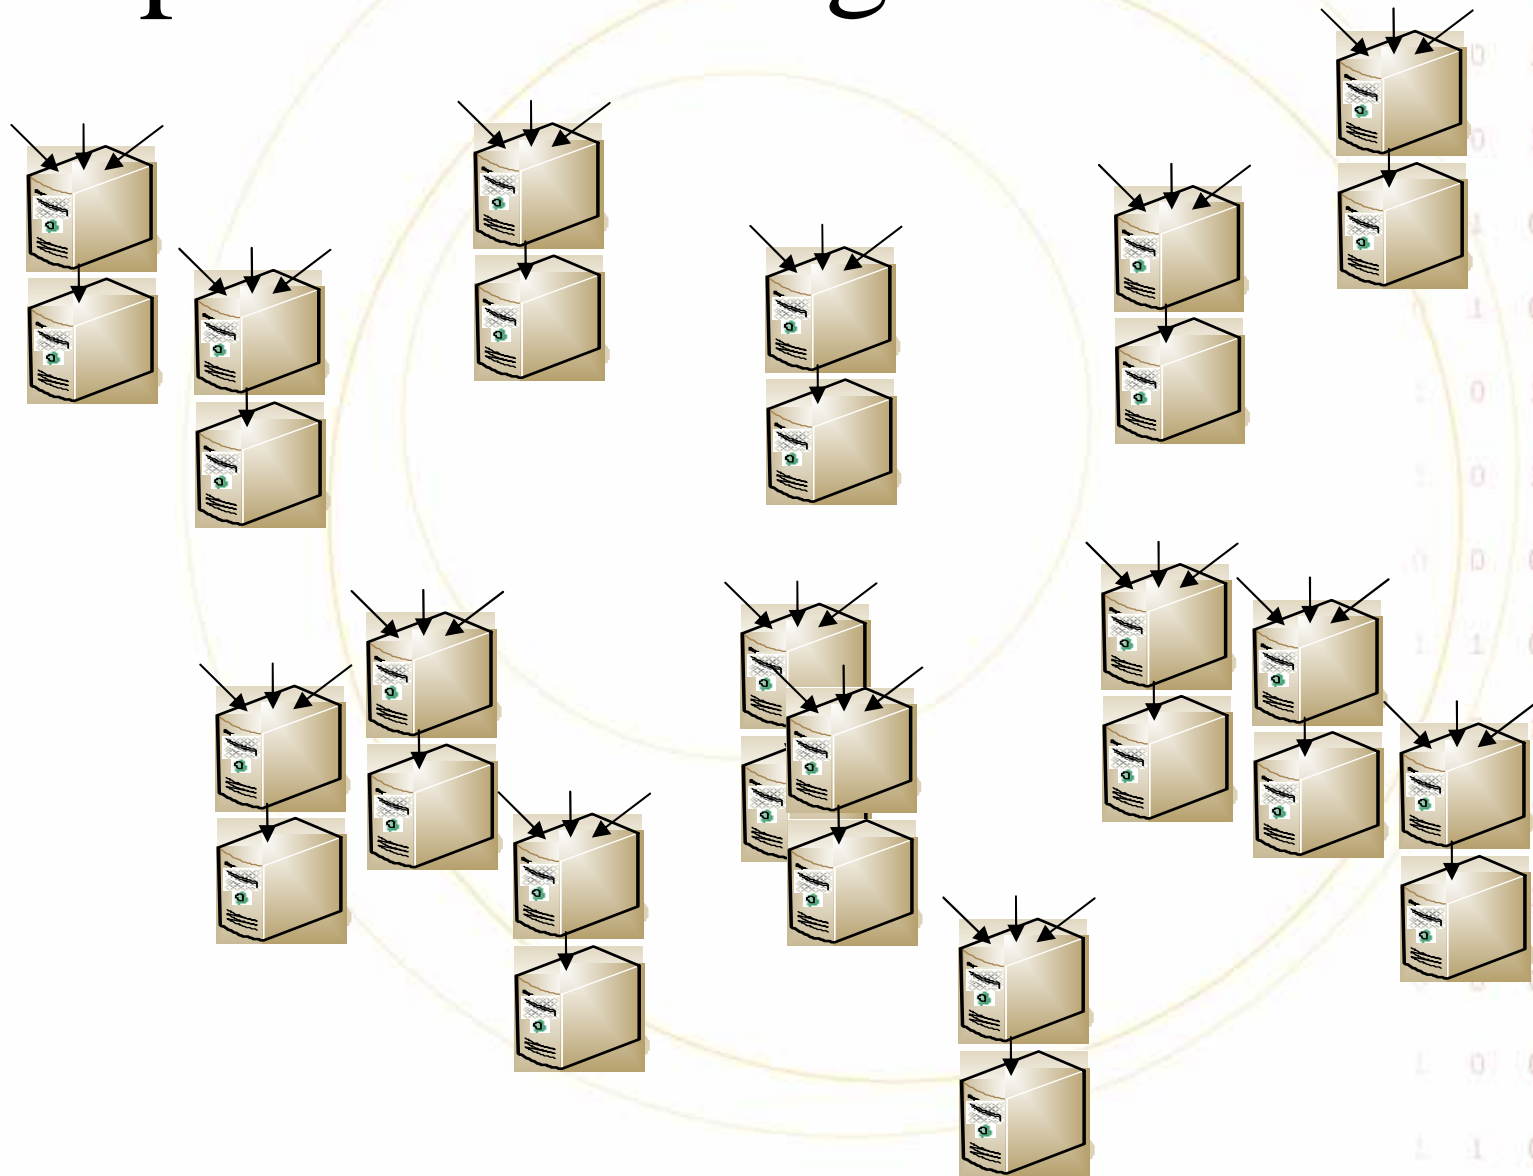
PostgreSQL Stack

- Mature platform for
 - LAMP
 - J2EE
 - BI/DW



1 1 0 0 0
1 0 0 1 0
0 0 1 0 1
0 0 1 1 0
0 0 1 0 1
0 1 0 0 0
1 1 0 0 1
1 0 1 0 0
1 0 1 1 0
0 0 0 1 1
1 1 0 0 0
0 1 1 0
1 0 1 0 1
1 0 1 0 1
0 0 0 1 0
1 0 0 1 1
1 1 0 0 0

Enterprise Challenge



Enterprise Requirements

- Reliability
- Availability
- Serviceability
- Manageability
- Usability
- Security
- Performance

1 1 0 0 0
1 0 0 1 0
0 0 1 0 1
0 0 1 1 0
0 0 1 0 1
0 1 0 0 0
0 1 0 0 1
1 0 1 0 0
1 0 1 1 0
0 0 0 1 1
1 1 0 0 0
0 1 1 0
1 0 1 0 1
1 0 1 0 1
0 0 0 1 0
1 0 0 1 1
1 1 0 0 0

Reliability

- “Coverity Inspected”
 - 1 defect every 39,000 lines code
 - all now fixed
- Mem Contexts
- Buildfarm
- Spike Source
- OSDL

PostgreSQL BuildFarm

Home Status Members Register PGFoundry PostgreSQL.org

PostgreSQL BuildFarm Status

Shown here is the latest status of each farm member for each branch it has reported on in the last 30 days.

Use the farm member link for history of that member on the relevant branch.

Legend

- = cassert
- = debug
- = integer-datetime
- = krb5
- = nls
- = openssl
- = pam
- = perl
- = python
- = tcl
- = thread-safety

Alias	System	Branch: HEAD	Status	Flags
tiretty	UnixWare 7.1.4 cc 4.2 i386	01:48:50 ago	OK Config	[c]
kookaburra	AIX 6.2 cc 6.0 powerpc	01:52:50 ago	OK Config	
echidna	FreeBSD 6-CURRENT gcc 3.4.2 x86	03:45:49 ago	OK Config	
wallaroo	Mac OS X 10.3.8 gcc 3.3 PowerPC G4	04:00:48 ago	OK Config	
platypus	FreeBSD 5.2.1-RELEASE-p5 gcc 3.3.3 amd64	04:10:49 ago	OK Config	
thruah	Fedora Core 1 gcc 3.3.2 i386	04:10:49 ago	OK Config	
herring	FreeBSD 6-CURRENT Intel(tm) CC 7.1 ia32	04:29:28 ago	OK Config	
asp	AIX 5.2 gcc 3.3.2 powerpc	04:52:50 ago	OK Config	
panda	Debian sid gcc 3.3.5 amd64	04:55:48 ago	OK Config	
viper	Fedora Core 3 gcc 3.4.2 x86_64	05:05:49 ago	OK Config	
shrew	Mandrake 10 gcc 3.3.2 AMD 2600x2	06:13:48 ago	OK Config	
osprey	NetBSD 2.0 gcc 3.3.3 m68k	06:15:34 ago	OK Config	

Availability

- System Availability
 - **Replication** →
 - Data Availability
 - Concurrent access
 - On-line maintenance
- *PostgreSQL Replication Solutions*
 - *Bruce Momjian, Core Team*
 - *14:20 Foyer Room*

1 1 0 0 0
1 0 0 1 0
0 0 1 0 1
0 0 1 1 0
0 0 1 0 1
0 1 0 0 0
0 1 0 0 1
0 1 0 0 0
1 0 1 1 0
1 0 1 1 0
0 0 0 1 1
1 1 0 0 0
0 1 1 0
1 0 1 0 1
1 0 1 0 1
0 0 0 1 0
1 0 0 1 1
1 1 0 0 0

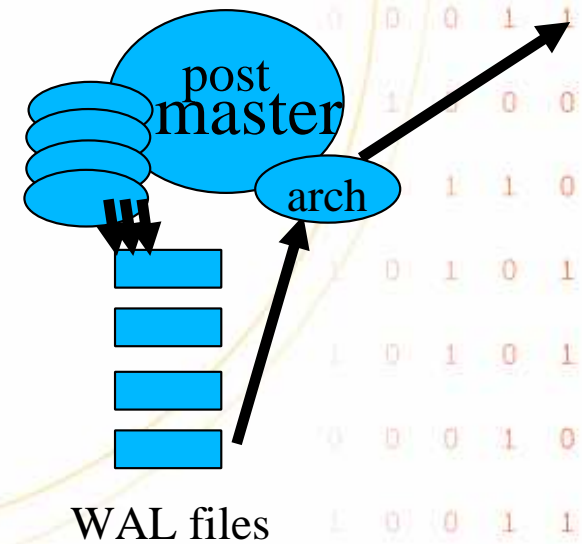
Serviceability

- If it goes down, will it come up?
- Programs
 - Good patches
 - Low time to fix
- Data
 - Recoverability

1 1 0 0 0
1 0 0 1 0
0 0 1 0 1
0 0 1 1 0
0 1 0 1
0 1 0 0 0
1 1 0 0 1
1 0 1 0 0
1 0 1 1 0
0 0 0 1 1
1 1 0 0 0
0 1 1 0
1 0 1 0 1
1 0 1 0 1
0 0 0 1 0
1 0 0 1 1
1 1 0 0 0

Hot Backup

- Enterprise Issue
 - Take backup at 04:00
 - System fails at 20:15
- HowTo
 - Set archive_command
 - pg_start_backup()
 - Take a base backup
 - pg_stop_backup()

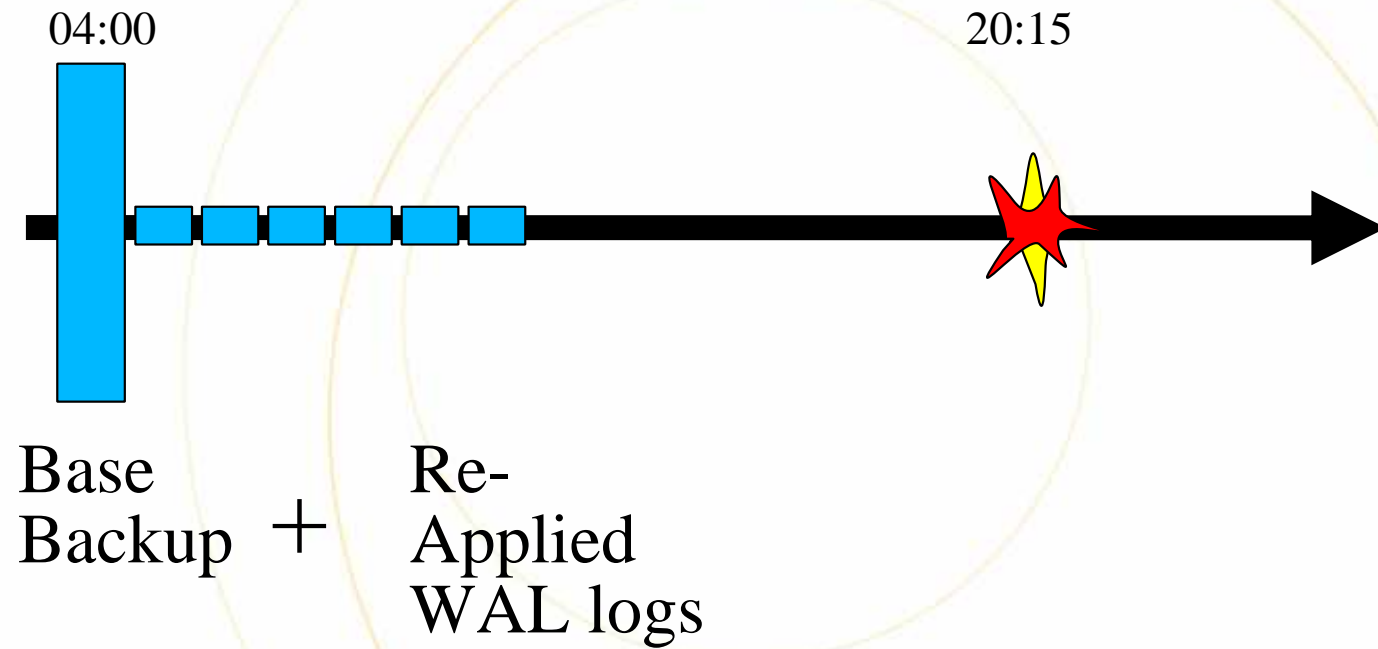


Archive Recovery




1 1 0 0 0
1 0 0 1 0
0 0 1 0 1
0 0 1 1 0
0 0 1 0 1
0 1 0 0 0
1 0 0 1
1 0 1 0 0
1 0 1 1 0
0 0 0 1 1
1 1 0 0 0
0 1 1 0
1 0 1 0 1
1 0 1 0 1
0 0 0 1 0
1 0 0 1 1
1 1 0 0 0

Archive Recovery




1	1	0	0	0
1	0	0	1	0
0	0	1	0	1
0	0	1	1	0
0	0	1	0	1
0	1	0	0	0
1	0	0	1	
1	0	1	1	0
0	0	0	1	1
1	1	0	0	0
0	1	1	0	
1	0	1	0	1
1	0	1	0	1
0	0	0	1	0
1	0	0	1	1
1	1	0	0	0

Archive Recovery




1 1 0 0 0
1 0 0 1 0
0 0 1 0 1
0 0 1 1 0
0 0 1 0 1
0 1 0 0 0
1 0 0 1
1 0 1 0 0
1 0 1 1 0
0 0 0 1 1
1 1 0 0 0
0 1 1 0
1 0 1 0 1
1 0 1 0 1
0 0 0 1 0
1 0 0 1 1
1 1 0 0 0

Point in Time Recovery



1	1	0	0	0
1	0	0	1	0
0	0	1	0	1
0	0	1	1	0
0	0	1	0	1
0	1	0	0	0
1	0	0	1	
1	0	1	0	0
1	0	1	1	0
0	0	0	1	1
1	1	0	0	0
0	1	1	0	
1	0	1	0	1
1	0	1	0	1
0	0	0	1	0
1	0	0	1	1
1	1	0	0	0

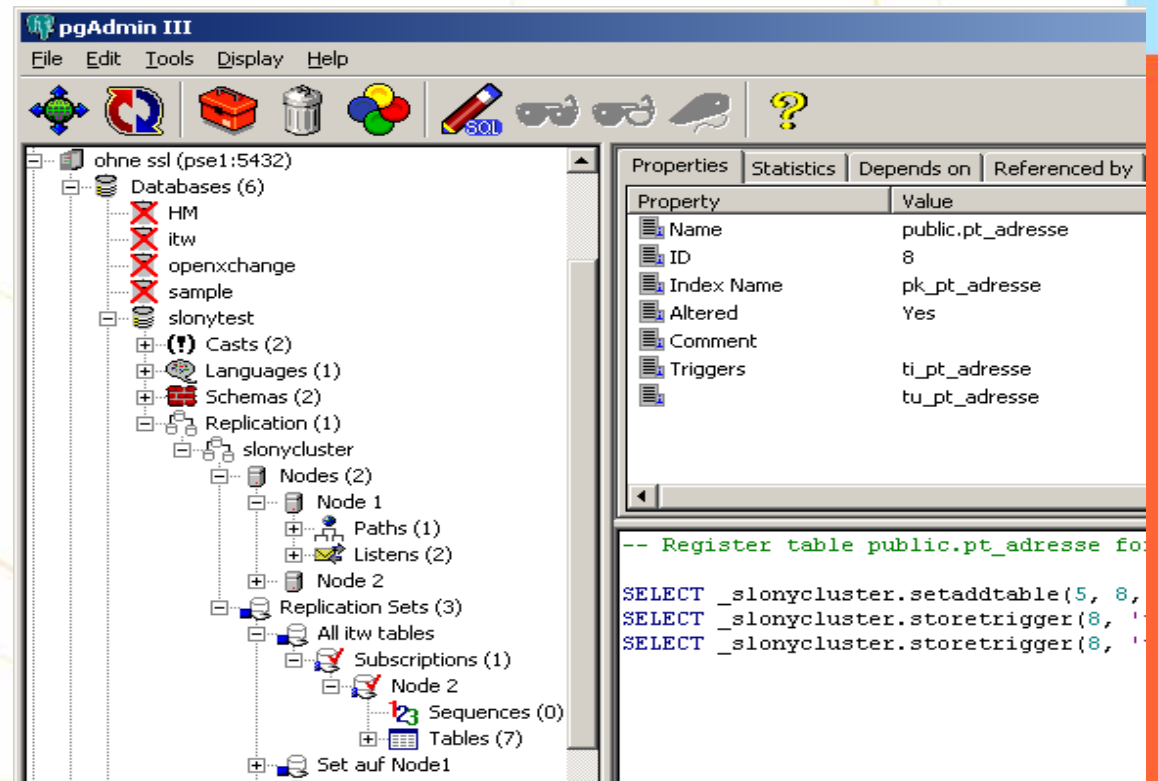
Multiple Complex Recoveries



1 1 0 0 0
1 0 0 1 0
0 1 0 1
0 0 1 1 0
0 0 1 0 1
0 1 0 0 0
1 0 0 1
1 0 1 0 0
1 0 1 1 0
0 0 0 1 1
1 1 0 0 0
0 1 1 0
1 0 1 0 1
1 0 1 0 1
0 0 0 1 0
1 0 0 1 1
1 1 0 0 0

Manageability

- How easy is it to drive?
- Linux, Windows, Solaris/AIX/HP-UX
- Monitoring
- Graphical remote admin
- Web remote admin



Usability

- Enterprise usability:
Does work with everything else?
 - SQL:2003
 - JDBC, ODBC
 - Extensibility: Compatibility
 - Tools
 - Applications
 - Staff

1 1 0 0 0
1 0 0 1 0
0 0 1 0 1
0 0 1 1 0
0 0 1 0 1
1 0 0 0
1 0 0 1
1 0 1 0 0
1 0 1 1 0
0 0 0 1 1
1 1 0 0 0
0 1 1 0
1 0 1 0 1
1 0 1 0 1
0 0 0 1 0
1 0 0 1 1
1 1 0 0 0

Security

- Can I keep control?
 - “PostgreSQL is by default a secure database...”
 - “...makes the database difficult to attack...”

The Database Hacker's Handbook

p.432

ISBN 0-7645-7801-4

1 1 0 0 0
1 0 0 1 0
0 0 1 0 1
0 0 1 1 0
0 0 1 0 1
0 1 0 0 0
1 1 0 0 1
1 0 1 0 0
1 0 1 1 0
0 0 0 1 1
1 1 0 0 0
0 1 1 0
1 0 1 0 1
1 0 1 0 1
0 0 0 1 0
1 0 0 1 1
1 1 0 0 0

Active Directory

- Authentication
 - Kerberos
- Directory information
 - LDAP
- ...

1 1 0 0 0
1 0 0 1 0
0 0 1 0 1
0 0 1 1 0
0 0 1 0 1
0 1 0 0 0
0 1 0 0 1
1 0 1 0 0
1 0 1 1 0
0 0 0 1 1
1 1 0 0 0
0 1 1 0
1 0 1 0 1
1 0 1 0 1
0 0 0 1 0
1 0 0 1 1
1 1 0 0 0

PG Kerberos: Build

- Prerequisites: MIT or Heimdal
 - On windows, MIT
- Enable Kerberos
 - ./configure --with-krb5
 - Default in many Linux distros
 - Included in pginstaller 8.1
- Set Service Principal Name
POSTGRES/host.domain@REALM
 - --with-krbsrvnam=POSTGRES
 - Runtime config in 8.1

1 1 0 0 0
1 0 0 1 0
0 0 1 0 1
0 0 1 1 0
0 0 1 0 1
0 1 0 0 0
1 1 0 0 1
1 0 1 0 0
1 0 1 1 0
0 0 0 1 1
1 1 0 0 0
0 1 1 0
1 0 1 0 1
1 0 1 0 1
0 0 0 1 0
1 0 0 1 1
1 1 0 0 0

PG Kerberos: Client

- Supported interfaces
 - libpq
 - OLEDB
 - ODBC 8.1+
 - Pretty much anything libpq based
- Natively unsupported interfaces
 - ODBC \leq 8.0
 - JDBC
 - Npgsql

1 1 0 0 0
1 0 0 1 0
0 0 1 0 1
0 0 1 1 0
0 0 1 0 1
0 1 0 0 0
0 1 0 0 1
1 0 1 0 0
1 0 1 1 0
0 0 0 1 1
1 1 0 0 0
0 1 1 0
1 0 1 0 1
1 0 1 0 1
0 0 0 1 0
1 0 0 1 1
1 1 0 0 0

PG Kerberos: Client

- **Service principal name (8.1)**

- Environment:
PG_KRB_SRVNAM=POSTGRES
- Connection string:
krbsrvname=POSTGRES

- **Registry key (MIT)**

- HKLM\SYSTEM\CurrentControlSet\Control\Lsa\Kerberos\Parameters
- HKLM\SYSTEM\CurrentControlSet\Control\Lsa\Kerberos (winXP SP2)
 - AllowTGTSessionKey = 0x01 (DWORD)

1 1 0 0 0
1 0 0 1 0
0 0 1 0 1
0 0 1 1 0
0 1 0 1
1 0 0 1
1 0 1 0 0
1 0 1 1 0
0 0 0 1 1
1 1 0 0 0
0 1 1 0
1 0 1 0 1
1 0 1 0 1
0 0 0 1 0
1 0 0 1 1
1 1 0 0 0

PG Kerberos: Join AD

- Create Service Principal
 - random pwd, never expire
- Export from DC

```
ktpass -princ  
  POSTGRES/host.domain.net@DOMA  
  IN.NET -crypto DES-CBC-MD5  
  -mapuser <spnuser> -pass  
  <password>  
  -out postgres.keytab
```

- Copy to PostgreSQL

1 1 0 0 0
1 0 0 1 0
0 0 1 0 1
0 0 1 1 0
0 0 1 0 1
0 1 0 0 0
1 1 0 0 1
1 0 1 0 0
1 0 1 1 0
0 0 0 1 1
1 1 0 0 0
0 1 1 0
1 0 1 0 1
1 0 1 0 1
0 0 0 1 0
1 0 0 1 1
1 1 0 0 0

PG Kerberos: Server

- Get Kerberos working first!
- postgresql.conf (8.1)
 - krb_srvname=POSTGRES
 - krb_caseinsens_users=yes
 - krb_server_keyfile=
/.../postgres.keytab
- pg_hba.conf

```
#type db user addr method
host all all 192.168.1.0/24 krb5
```

1 1 0 0 0
1 0 0 1 0
0 0 1 0 1
0 0 1 1 0
0 0 1 0 1
0 1 0 0 0
1 1 0 0 1
1 0 1 0 0
1 0 1 1 0
0 0 0 1 1
1 1 0 0 0
0 1 1 0
1 0 1 0 1
1 0 1 0 1
0 0 0 1 0
1 0 0 1 1
1 1 0 0 0

PG Kerberos: Connect

- Always specify server FQDN
- Never specify password
- Usually don't specify username

```
psql -h myserver.example.net mydb
```

DBLink-LDAP

- Read data from *any* LDAP directory
- Server-side user defined function
 - OpenLDAP on Unix
 - WinLDAP on Windows
 - Both talks to Active Directory

1 1 0 0 0
1 0 0 1 0
0 0 1 0 1
0 0 1 1 0
0 0 1 0 1
0 1 0 0 0
1 1 0 0 1
1 0 1 0 0
1 0 1 1 0
0 0 0 1 1
1 1 0 0 0
0 1 1 0
1 0 1 0 1
1 0 1 0 1
0 0 0 1 0
1 0 0 1 1
1 1 0 0 0

DBLink-LDAP

- List all users in the default container

```
CREATE VIEW all_my_users AS
SELECT * FROM
  ldap_search('dc.example.net',
    'CN=Users, DC=example, DC=net',
    'EXAMPLE\adreader',
    'TopSecret',
    '(&(objectClass=user)
      (objectCategory=person))',
    'cn,displayname')
  t(cn text, displayname text);
```

1 1 0 0 0
1 0 0 1 0
0 0 1 0 1
0 0 1 1 0
0 0 1 0 1
0 1 0 0 0
1 1 0 0 1
1 0 1 0 0
1 0 1 1 0
0 0 0 1 1
1 1 0 0 0
0 1 1 0
1 0 1 0 1
1 0 1 0 1
0 0 0 1 0
1 0 0 1 1
1 1 0 0 0

DBLink-LDAP

- Get the display name of a user


```
CREATE FUNCTION
  get_user_displayname(text)
  RETURNS text AS $$
SELECT displayname FROM
  ldap_search('dc.example.net',
    'DC=example, DC=net',
    'EXAMPLE\\adreader',
    'TopSecret',
    '(cn=' || $1 || ')',
    'displayname')
  t(displayname text)
$$ LANGUAGE 'sql';
```

```
1 1 0 0 0
1 0 0 1 0
0 0 1 0 1
0 0 1 1 0
0 0 1 0 1
0 1 0 0 0
1 1 0 0 1
1 0 1 0 0
1 0 1 1 0
0 0 0 1 1
1 1 0 0 0
0 1 1 0
1 0 1 0 1
1 0 1 0 1
0 0 0 1 0
1 0 0 1 1
1 1 0 0 0
```


Performance: 8.1

- Scalability improvements
- Optimizer & Executor enhanced
- Partitioning, Dynamic bit-map indexes
- Very large memory support
- Testing
 - DBT-2/TPC-C
 - DBT-3/TPC-H

1 1 0 0 0
1 0 0 1 0
0 0 1 0 1
0 0 1 1 0
0 0 1 0 1
0 1 0 0 0
1 0 0 1
1 0 0
1 0 1 1 0
0 0 0 1 1
1 1 0 0 0
0 1 1 0
1 0 1 0 1
1 0 1 0 1
0 0 0 1 0
1 0 0 1 1
1 1 0 0 0



- **PostgreSQL** <http://www.postgresql.org/>
 - Download 8.1 <http://www.postgresql.org/download/>
 - Press <http://www.postgresql.org/about/press/>
 - dblink-ldap <http://pgfoundry.org/projects/dblink-ldap/>
- **Speakers**
 - Simon Riggs simon@2ndquadrant.com
 - Magnus Hagander mha@sollentuna.net



PostgreSQL BOF

17:35

Seasons room

1 0 0 0 0 0
1 0 0 1 1 0
1 0 0 1 0 1
1 0 1 0 0 0
1 1 0 0 1
1 0 1 0 0
1 0 1 1 0
1 0 0 1 1
1 1 0 0 0
1 0 1 0 1
1 0 1 0 1
1 0 0 1 0
1 0 0 1 1
1 1 0 0 0