



Enterprise Database Inside
or:
There's an elephant in my Solaris!

**Josh Berkus, PostgreSQL Lead
Database Technology Group**



What is this elephant doing here?



This elephant showed up in Solaris 10 Update 2 in June 2006

Did you notice it?

Why would we add an elephant to Solaris?



What is this elephant doing here?




Since Solaris is the leading operating system for data centers, we simply wanted to make sure that every Solaris user had access to an enterprise database system -- right out of the box.



Introduction to PostgreSQL for Solaris

- What is PostgreSQL?
- Why use PostgreSQL?
- Getting Started with PostgreSQL
- Solaris Features and PostgreSQL



What is PostgreSQL?



What is PostgreSQL?

- Traditional relational database system
 - > client-server
 - > intended for dedicated database servers
 - > aimed at professional DBAs
 - > used for data-centric projects and enterprise data stores



What is PostgreSQL?

- Object-relational database system
 - > Incorporates advanced ideas from database research
 - > Designed to be extensible by the DBA
 - > Database objects can be re-designed or added



What is PostgreSQL?


- Worldwide, thriving open source community

Estimated direct users	more than 300,000
Estimated "embedded" users	40 million
Developers	about 250
Developer time zones	14
Active online community	about 50,000
Contributing companies	more than 40
Non-profits orgs	5
Japanese member DBAs	more than 2500
Languages	7 full / 5 partial



21-year-old Database Project

- **1986:** POSTGRES at the University of California, Berkeley
 - > Michael Stonebraker project
 - > Successor to INGRES
- **1994:** first commercialized
 - > as Illustra (later merged into Informix)
- **1995:** open-sourced
 - > Ported to SQL
 - > PostgreSQL Global Development Group formed
- **1998:** ported to Japanese, supported in Japan
- **1999:** first full-time developers & corporate support
- **2005:** Multiple PostgreSQL start-ups
- **2006:** supported by Sun



Why PostgreSQL



Why PostgreSQL: Key Features

- High-Concurrency ACID Transactions
- Complex Queries
- HA Features
- SMP Scaling
- Extensibility
- Standards Compliance
- Security
- License



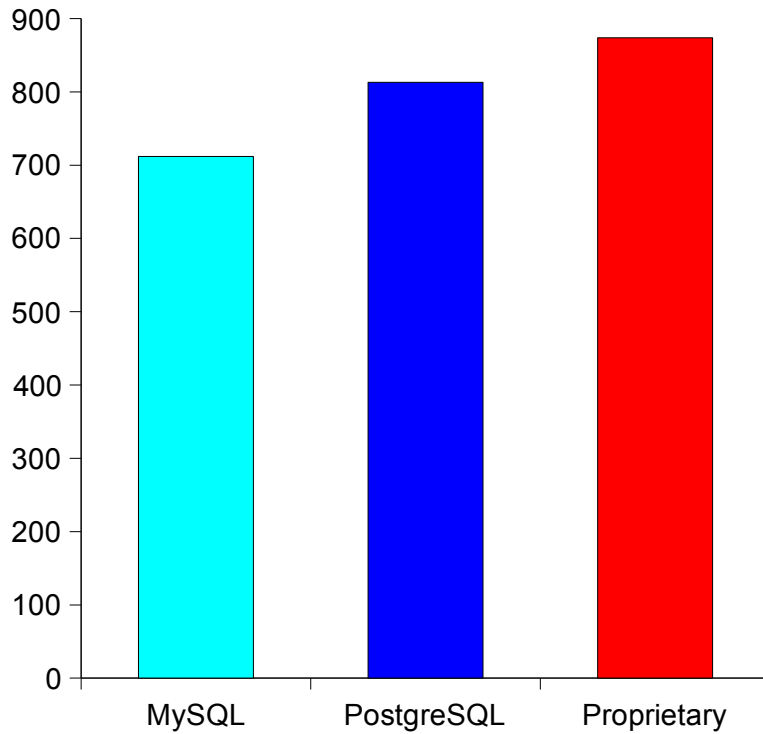
High-Volume ACID Transactions

- Atomic, Consistent, Isolated & Durable (ACID) transactions
 - > never lose data if your hardware is good
 - > bulletproof crash recovery
 - > transaction-safe table and schema changes (DDL)
 - > all isolation levels
 - > SQL SAVEPOINTS (subtransactions)
- High-performance concurrency
 - > effortless isolation
 - > fast rollbacks

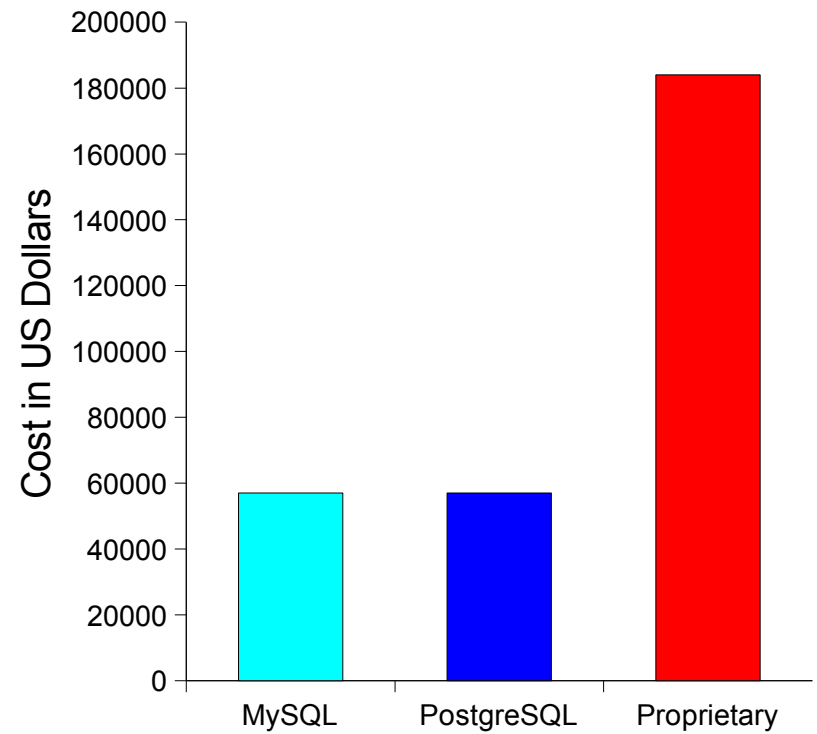


OLTP Performance & Cost

J2EE Througput



Acquisition Cost Comparison





High Availability Features

- Point-In-Time Recovery (PITR)
 - > 2 modes:
 - > Continuous Warm Standby
 - > Recover-to-point-in-time (recover from DBA error)
- Hot Backup
 - > backup remotely, as a regular SQL user
- Transactional DDL
 - > Concurrent Index Build: index without locking your tables
 - > Modify database objects without risk or database-wide locks



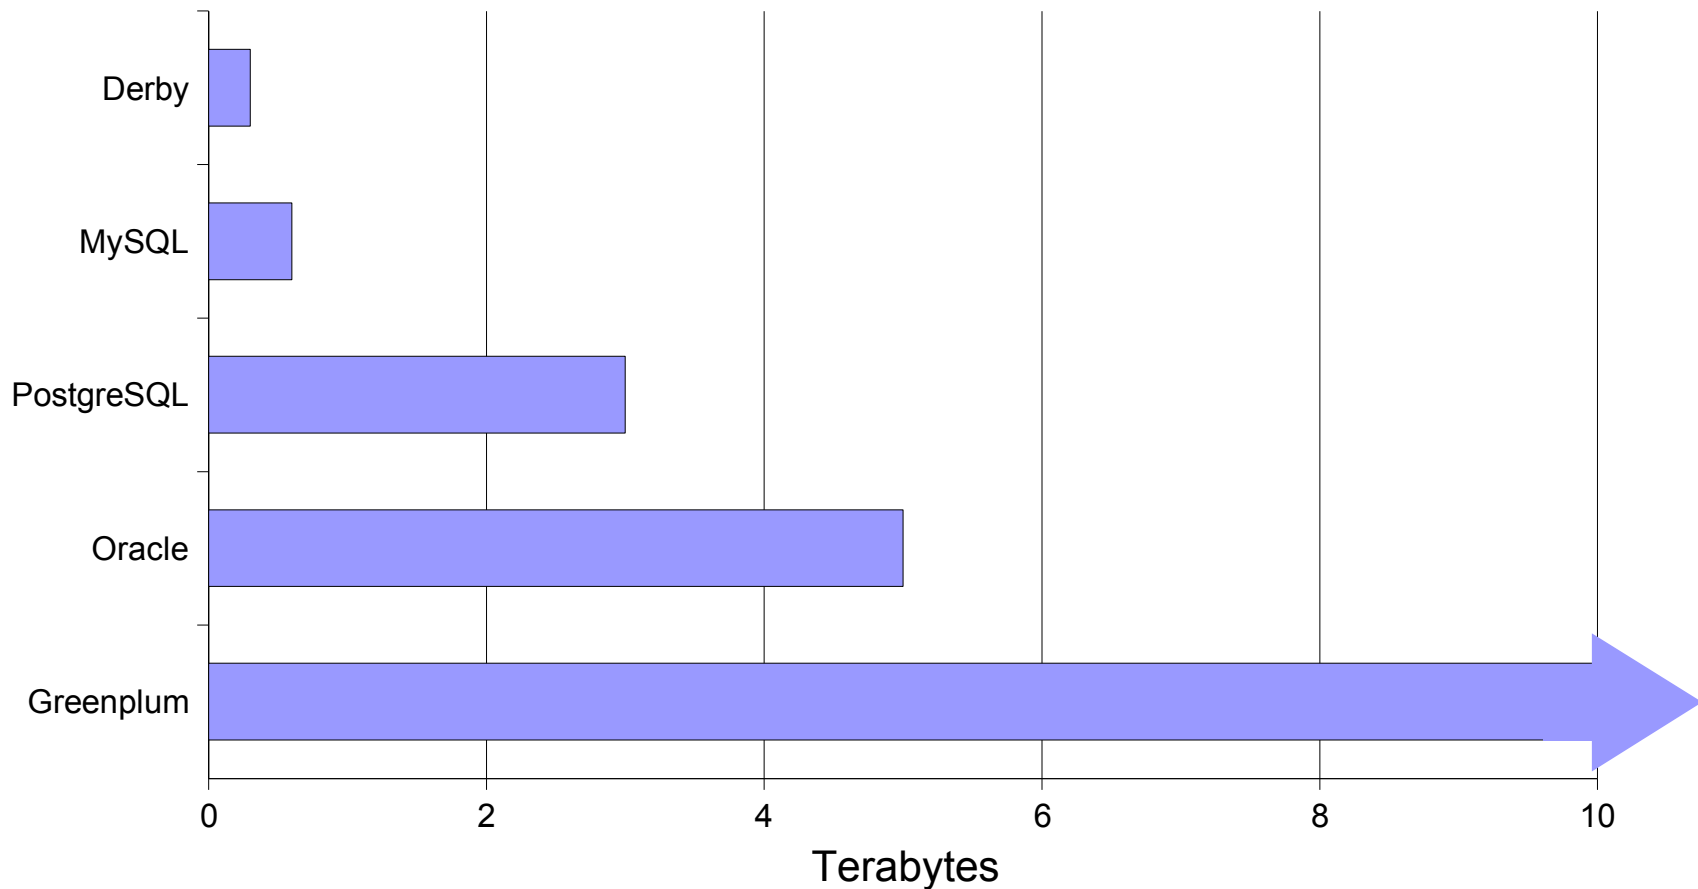
Complex Queries

- Optimizes and Executes Large, Complex Queries
 - > 50+ line queries
 - > 40-way joins with genetic optimization
 - > All join types supported
 - > outer joins
 - > subqueries
 - > UNION, EXCEPT, INTERSECT
 - > Partitioning
 - > Perfect for small data warehouses
 - > 1 TB to 7 TB



Data Warehousing Comparison

- General data warehouse size





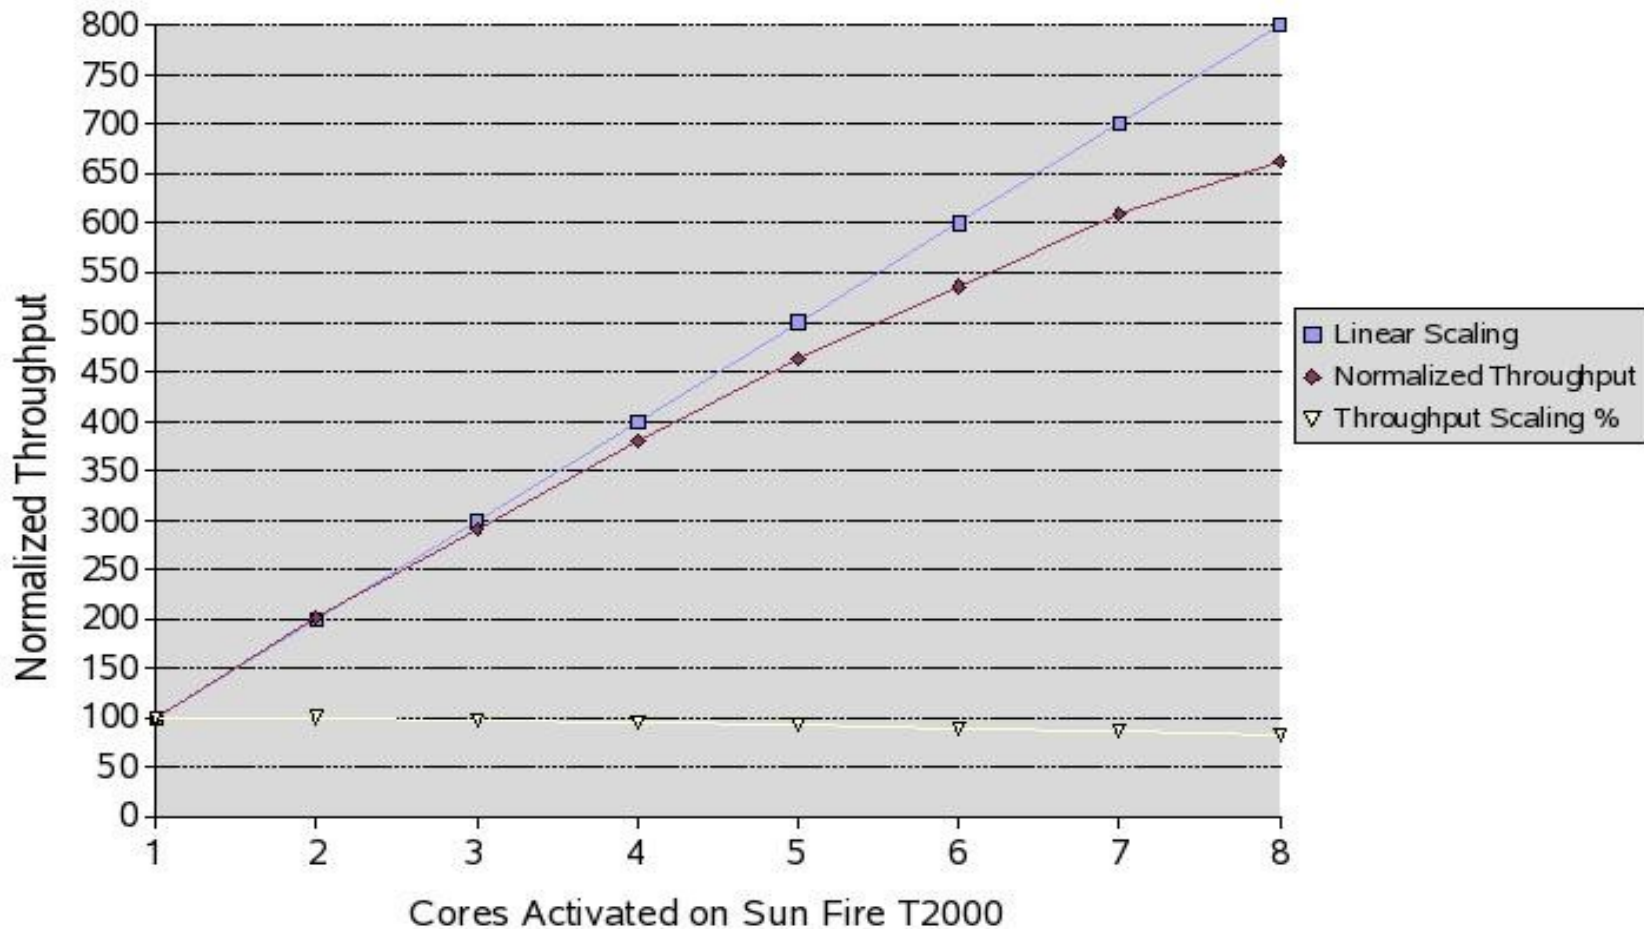
SMP Scaling

- PostgreSQL is the most scalable multi-core open source database.
 - > Scales linearly to 8 to 30 cores, depending on workload
 - > Great CPU utilization on 8+ core machines with many concurrent connections
 - > Good for Niagara!



SMP Scalability on Niagara

PostgreSQL 8.2.4 Scaling on Sun Fire T2000 with Solaris 10 11/06





Extensibility

- *Object-relational* extensibility
 - > Add your own data types, operators and index types
 - > advanced full-text search engine
 - > Geographic Information Service
 - > genomics
 - > cryptography
 - > Write stored procedures in 12 different languages
 - > C, SQL, PL/pgSQL, Ruby, Java, Perl, Python, R, C++ and more
 - > Easily port stored procedures from Oracle and DB2




Standards Compliance

- Broad SQL compliance
 - > Implements almost all of SQL 92
 - > views, domains, schema, procedures ...
 - > better syntax compliance than most DBMSes
 - > Many features from SQL 99 and 03
 - > composite types, information schema, savepoints ...
- Commitment to standards compliance
 - > Certified JDBC3 compliant
 - > ODBC and XA support
 - > Support most language encodings
 - > UTF8, Euc-JP, GBK, LATIN1, BIG5, more



Security



“By default, PostgreSQL is possibly the most security-aware database available ...”

-- Database Hacker's Handbook,
2006



Security Features

- Access Control List and multiple authentication methods
 - > MD5, SHA, PAM, Kerberos, OpenSSL
- Full support for SQL ROLES
 - > including inherited roles and security definer functions
- Huge array of logging options
 - > connections, transactions, DDL, offsite logging
- Fewest bugs of any major SQL database
 - > according to Coverity reviews



License

- All of the PostgreSQL core is under the permissive “BSD” license
 - > allows combining with proprietary software
 - > does not require reporting or open-sourcing
 - > very popular with OEMs and ISVs


PostgreSQL is released under the BSD license.

PostgreSQL Database Management System

Portions Copyright (c) 1996-2007, The PostgreSQL Global Development Group

Portions Copyright (c) 1994, The Regents of the University of California

Permission to use, copy, modify, and distribute this software and its documentation for any purpose, without fee, and without a written agreement is hereby granted, provided that the above copyright notice and this paragraph and the following two paragraphs appear in all copies. ...



Getting Started



Where to Get It

- Solaris 10
 - > PostgreSQL 8.1 came with Update 2
 - > Update 4 now has PostgreSQL 8.2
 - > recommended
- Solaris Express Developer Edition
 - > SXDE 9/07 is performance-optimized
- No need for additional installs
 - > PostgreSQL is just there



Getting Started: Activation

Three Easy Steps (on Solaris 10u4 and SXDE)

1. Log in as the user “postgres”

```
su - postgres
```

1. Initialize the database directory (as postgres user)

```
/usr/postgres/8.2/bin/initdb -D  
/var/postgres/8.2/data
```

2. Start the database service (as root)

```
svcadm enable postgresql:version_82
```

demo: showing the two-step process of activation.



Getting Started: SMF

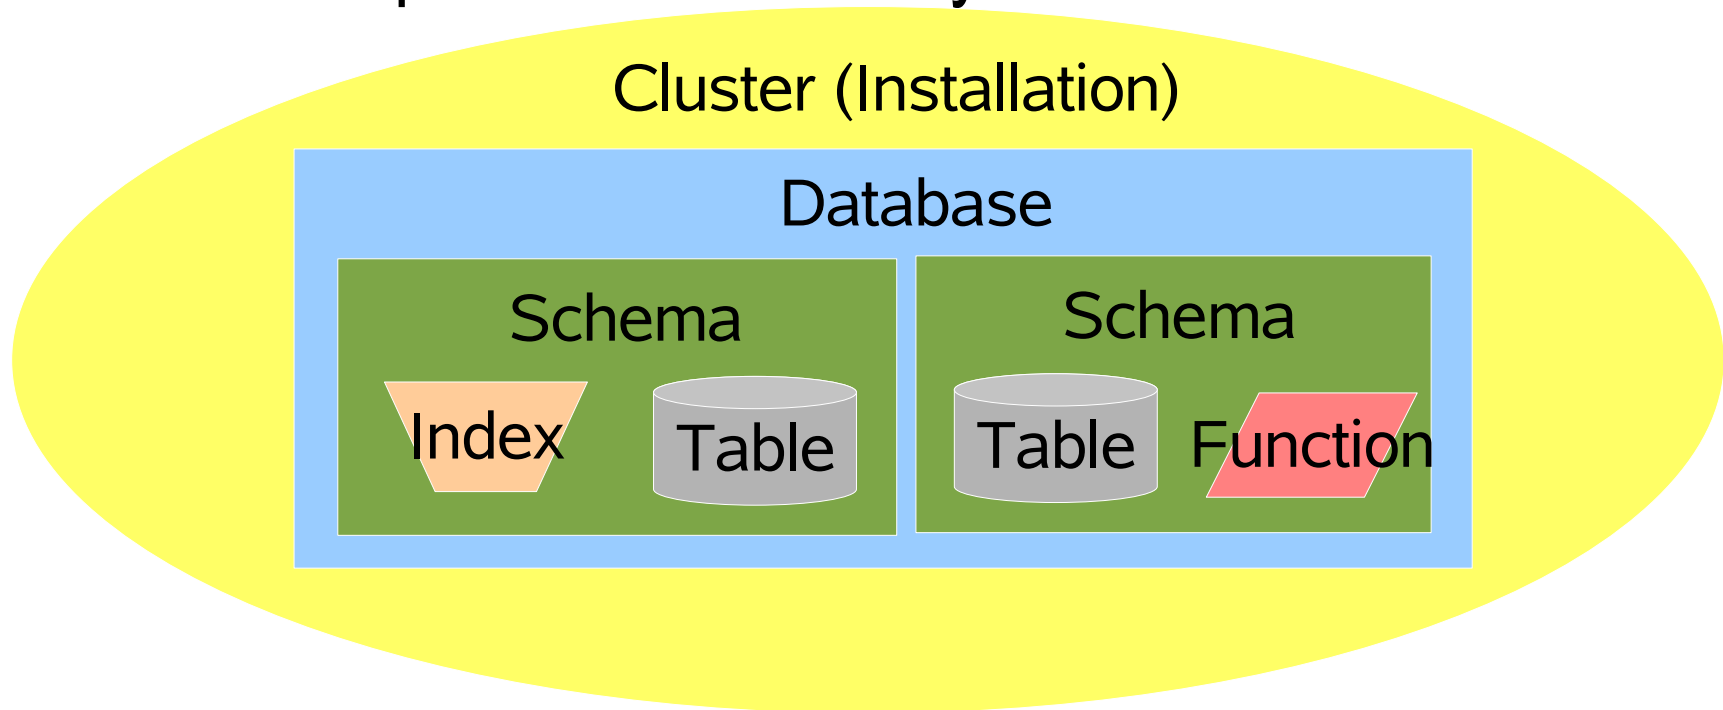
- Use SMF to manage PostgreSQL
 - > better for administering whole server rooms
 - > easily accomodates varying configurations
 - > auto-restarts PostgreSQL in the event of a DOS attack or system fault

demo: run PostgreSQL 8.2.0. Crash it using known bug. Show how SMF auto-restarts it, preventing DOS. **How to change SMF Properties**



Getting Started: First Database

- Data encapsulation hierarchy:




- Create your first database

demo: creating a database. Show how to use CREATEDB, psql, and how to load a SQL file. Show pgAdmin.



Performance Configuration

- Configuration is required
 - > Like Oracle, PostgreSQL requires some manual performance configuration for production use.
 - > postgresql.conf database configuration
 - > Solaris tuning
 - > filesystem and disks
 - > There is too much information to cover here:
 - > <http://wikis.sun.com/display/DBonSolaris/PostgreSQL>
- Autoconfiguration is planned
 - > Sun engineers are working on it



Solaris Features & PostgreSQL



Solaris + PostgreSQL = Enterprise

- More Highly Available
- Additional Security
- Better Scalability
- Enterprise Support



More Highly Available

- Sun Cluster support
 - > Sun Cluster agent available
 - > Fully supported
- ZFS + Zones for Hot Standby
 - > ZFS cloning & replication lets you “clone” your database
 - > Zones allows you to “clone” your DB server
 - > OmniTI is using this technique to support a 6TB data warehouse on PostgreSQL

demo: zones + zfs for hot standby



Add Security with Zones

- Zones are Perfect for “hosted” PostgreSQL
 - > allows 100% isolation of PostgreSQL instances per customer/application
 - > set CPU & RAM limits for each Zone (in SXDE)
 - > PostgreSQL can't do this natively
 - > no per-user limits, no ability to hide the system tables
 - > Being used by Joyent.com to host 100's of customers.



Better Scalability with DTrace

- The first database with DTrace probes
 - > MySQL, Derby, more later
 - > only 14 probes now, more to be added by PostgreSQL community
- Real-time monitoring of SMP performance
 - > locks and lock conflicts
 - > transactions
 - > kernel resources
- Part of future performance monitoring tools
 - > planned for 2007-2008

demo: show PostgreSQL+ **Chime** or DLight demo.




Enterprise Support

- Enterprise & Standard Support available for Solaris
 - > 10x to 50x cheaper than proprietary DBs
 - > 24x7 support
 - > The only truly worldwide PostgreSQL support
- Expert Support Team
 - > PostgreSQL Core Team member and several contributors
 - > Sun performance experts with 10+ years of experience tuning databases



PostgreSQL on Solaris: Summary

- Included with Solaris 10 / SXDE
- Enterprise-grade
 - > reliable
 - > high-performance
 - > supported
- Full of cool features
- Try it out today!





More Information

- Questions: postgresql-feedback@sun.com
- Blogs on PostgreSQL
 - > Josh Berkus: blogs.ittoolbox.com/database/soup
 - > Jignesh Shah: blogs.sun.com/jkshah
 - > Tom Daly: blogs.sun.com/tomdaly
 - > Robert Lor: blogs.sun.com/robertlor/
- PostgreSQL for Solaris: <http://sun.com/postgresql>
- PostgreSQL: www.postgresql.org
- OpenSolaris databases community:
www.opensolaris.org/os/community/databases/