

This is your PostgreSQL on drugs

Aaron Thul

Electronic Medical Office Logistics

EMOL



Who am I?

- Computer & Database Geek, just like you
- Formerly a SysAdmin at Autoweb Communications
 - [PostgreSQL Build Your Car](#)
- Presently a IT manager at a EMOL
- PostgreSQL Evangelist
- Penguicon Organizer



With PostgreSQL and other Open Source software EMOL is

- Allowing Data collection from EMRs and other sources
- Aiding in Adherence to national standards
- Providing Physician and Practice level benchmarking
- Data Brokering
- Enabling Automation of National initiatives, such as the CMS PQRI



EMOL PostgreSQL Data

- Patient Records
- Billing Records
- Lab Results
- Clinical Records
- Inventory Management
- Patient Reported Data



Metadata

- Physicians Dictations
- Scanned Documents
- Images
 - XRAYs
 - MIRIs
 - CAT Scans



Metadata Storage

- ReiserFS with tail packing
- Each practice/doctor has a folder
- Files are stored with the same file name as the primary key in the database
- ReiserFS is a killer filesystem
- OpenSolaris & ZFS???
- Netapp Waffle???



EMOL Software Building Blocks

- Ubuntu Linux LTS (8.04)
- PostgreSQL (8.3)
- Perl (5.8.x)
- Windows Unified Data Storage Server 2003 (R2)
 - Yes Windows



EMOL Hardware Building Blocks

- HP ProCurve Switches
 - Support **considerably** cheaper than Smartnet
- SonicWall Firewalls
 - Support **considerably** cheaper than Smartnet
- Large number of SCSI and SATA Hard Drives
- Aberdeen Servers and DAS (Direct Attached Storage) Systems



Why PostgreSQL?

- Capable
- Required Features
- Database Team Experience
- Community
 - Documentation Project
 - Mailing Lists
 - IRC
 - Events Like This!



Why Perl?

- Practical **E**xtraction and **R**eport **L**anguage
- Development team experienced with Perl
- Unix-centric, and available for Windows
- Text parsing and normalizing
- I know it Perl is not sexy like
 - `INSERT 'new_popular_language' INTO languages;`
- Rapid prototyping
 - Weakly typed
 - Interpreted, though very fast
 - Supports objects



Who is Where?

- OS and PostgreSQL binaries on local disks
 - RAID 1 Mirror
 - 15k spindle drives
 - EXT3
- WAL Buffers on local disks
 - RAID 1 Mirror
 - 15k spindle speed
 - EXT2
- INDEXs
 - DAS (Direct Attached Storage) Units
 - RAID 6
 - 10 k spindle speed SCSI
 - EXT3
- TABLES
 - Multiple ISCSI Servers on SANS
 - 4 x 1 Gigabit Ethernet Interfaces Bonded
 - 8 x 1 Terabyte SATA drives per SAN Node RAID 6
 - EXT3



Data Daily

- Loading 10 GB data daily into PostgreSQL
- Loading 10 GB metadata daily



Data Size

```
SELECT relname, (relpages*8)/1024 as MB  
FROM pg_class  
ORDER BY relpages DESC;
```



How much data are we talking

- Largest Table: 1,844.73 GB
- Second Largest Table: 1,289.36 GB

- Largest Index: 411.91 GB
- Second Largest Index: 405.08 GB

- Total DB size on disk: 16,800.39 GB



Run it twice and make it faster

- Maintain a 1/500 set of random sample data
- ALL queries hit that data base first
- Only once query result is successful is the query moved onto production server



How do I sleep at night

- First Name
 - Last Names
 - Social Security Numbers
 - Birth Dates
-
- Needed to track people over time and geography



Protecting the Warehouse

- Simple processes that are followed
- Intrusion Prevention & Firewalls
- Security Monitoring & Management - MSSP
- Encrypted Communication
- Identity Management - Centralized management of users and groups – mitigates vulnerabilities that occur due to inconsistencies



Protecting the Warehouse

- Role-based security
- Functions everywhere we can
- Identity data symmetrically encrypted
- Data is anonymized in all but a few tables
- Role-based security
- All data is anonymized before it is sent out



Lessons Learned

- Server Ethernet Cards are not all made the same
- With 100+ drives be ready to RMA some disks
- You can never have too many DIMM slots
- You do get what you pay for with RAID controllers
- You can never have too many PCI-X Slots
- You can't have too big a cache on your RAID controller



More Lessons Learned

- Auto-vacuum is not always your friend
- Worry when a developer says “I have an idea”
- Some mistakes are just too much fun to make only once
- Never ask for directions from a two-headed tourist!

-Big Bird



Questions

- Web: <http://www.chasingnuts.com>
- Email: aaron@chasingnuts.com
- IRC: AaronThul on irc.freenode.org
- Jabber: apthul@gmail.com
- Twitter: @AaronThul
- AIM: AaronThul

