

The Case for Opening Up

An Enterprise Strategy for Open Source

Derek Gottfrid
Jacob Harris
nytimes.com

FREE YOURSELF FROM BOOKING FEES.
Continental Airlines

The New York Times

Wednesday, March 26, 2008 Last Update: 1:22 PM ET

CONTINENTAL.COM
Continental Airlines



NYT Archive Since 1981 Search

Get Home Delivery | Personalize Your Weather

JOB
 REAL ESTATE
 AUTOS
 ALL CLASSIFIEDS

WORLD
 U.S.
 Politics
 Washington
 Education
 N.Y./REGION
 BUSINESS
 TECHNOLOGY
 SPORTS
 SCIENCE
 HEALTH
 OPINION
 ARTS
 Books
 Movies
 Music
 Television
 Theater
 STYLE
 Dining & Wine
 Fashion & Style
 Home & Garden
 Weddings/
 Celebrations
 TRAVEL

Maliki Gives Shiite Militias 72 Hours to Halt Fighting

By JAMES GLANZ and ANAHAD O'CONNOR 11:48 AM ET

Fighting in two of Iraq's largest cities threatens to destabilize a long-term truce that had helped reduce the level of violence in the five-year-old war.

• Photographs | Video

Paulson Reviewing Regulation of Investment Banks

By THE ASSOCIATED PRESS 11:20 AM ET

Treasury Secretary Henry M. Paulson Jr. said that the crash of Bear Stearns underscores the need for broad oversight of investment houses.

• CNBC Video: Paulson on Wall Street Regulation



Andy Wong/Associated Press

Arrests Swell in Tibet Mayhem

By DAVID BARBOZA 10:37 AM ET

China said that 660 people implicated in Tibetan protests and riots in western China had surrendered. Above, Chinese security officers on patrol in Lhasa.

• Unrest at Shattered Gateway to Tibet | Video | Times Topics

CITY ROOM
M.T.A. Votes to Sell West Side Rights to Tishman

TRAVEL » As Turmoil Subsides, Tourism in Nepal Surges

Katmandu's troubled decade did nothing to harm the city's age-old appeal.



On the Lam With Ken Kesey

A pilgrimage to Manzanillo in search of the man and the myth.



• Audio Slide Show

OPINION »

- Dowl: Hillary or Nobody?
- Op-Ed: McCain & the Media
- Freakonomics: Predictions
- Editorial: Taiwan & China
- Post a Comment

ACTIVE TRADERS »



MARKETS 1:24PM ET BigCharts

DOW	12,411.14	-121.46	-0.97%
NAS	2,312.96	-28.09	-1.20%
S&P	1,339.80	-13.19	-0.97%

My Portfolio » Stock Quotes: Go

• Tools: Alerts | Stocks | Sectors

• Earn up to 40,000 bonus points in one year – enough for one domestic round-trip airline ticket plus other business supplies
 • Earn 1 point for virtually every dollar you spend
 • Fly any airline anytime

0%

Open

NYTimes.com codebase 2 years ago

n %

Open*

* where $0 < n < 100$

NYTimes.com codebase today

Moving To N%

1. Context

- New Choices in Open Source
- Strategies For Switching

1. DBSlayer

- Building From Open Parts
- Embracing Open Standards

1. Times Archive

- The Compute Cloud Revolution
- Challenges For Licensed Software

1. Context

Context

- Custom web language built to serve pages on NYTimes.com
- Began as a dynamic ad serving platform
- Preprocessing model in HTML
- Compiles to bytecode executed by VM
- Custom CNX backend for DB access
- Built by 2 part-time employees
- Limited syntax for performance
- Limited mechanisms for extensions

Just For Comparison

Context

- 2 developers
- 30 users
- 0 books
- 0 libs / extensions
- Lagging support for new features

PHP

- 445 developers
- ~4 million users
- 1402 books
- 665 libs / extensions
- Quick addition of new features (PECL)

So why Context?

Build or Buy?

Six Years Ago

Build, Buy, or Download?

Today

Oh, The Things You'll Download!

- Web Servers
- Web Frameworks
- Load Balancers
- Proxies
- Template Languages
- CMSes
- Spiders
- Caches
- Message Queues
- Packet Analyzers
- Databases
- Data Warehouses
- XML parsers
- Search Engines
- Blogging Engines
- Image Processors

Don't
Bet Against
The Internet.

Eric Schmidt, Google

Switching To Open Source

- First we replaced components with open source equivalents:
 - Context → PHP
 - iPlanet → Apache
 - Oracle → MySQL
- Switching Works When:
 - Commodity Parts
 - Standards-compliant Interfaces
 - Loose Coupling

Phased Rollout

a → b

Much easier as

a → a' → a'' → b

With the Big Rewrite you will shoot yourself in the foot. Try aiming at each toe instead.

Belt and Suspenders

$a \rightarrow a' \rightarrow a'' \rightarrow b$

$a \longrightarrow b$

Dealing With Sticking Points

- Those cases specific to our system:
 - Cookie Authentication
 - Ad Serving
 - The CNX data backend
- Strategies For Adapting:
 - Native Language Bindings
 - Legacy Services
 - Testing Is Essential

Native Language Bindings

- Cookie Authentication
 - Identifies users to applications
 - C code within Context
- Native-language Adapters
 - PHP wrappers
 - Ruby native code
 - Java Native Interface (JNI), etc.
- Simplified Wrapper and Interface Generator (SWIG)

Legacy Services

- Ad serving
 - Tightly integrated into Context
 - Too costly to rewrite
- Using Context As Service:
 - Context still handles ad serving
 - Wrote PHP client stubs to fetch ads from Context
 - Will replace ad serving at a later date

Testing Is Essential

- Always Work With a Safety Net!
 - Unit tests
 - Regression tests
 - Integration tests
 - Phased Rollout
 - Belt and Suspenders

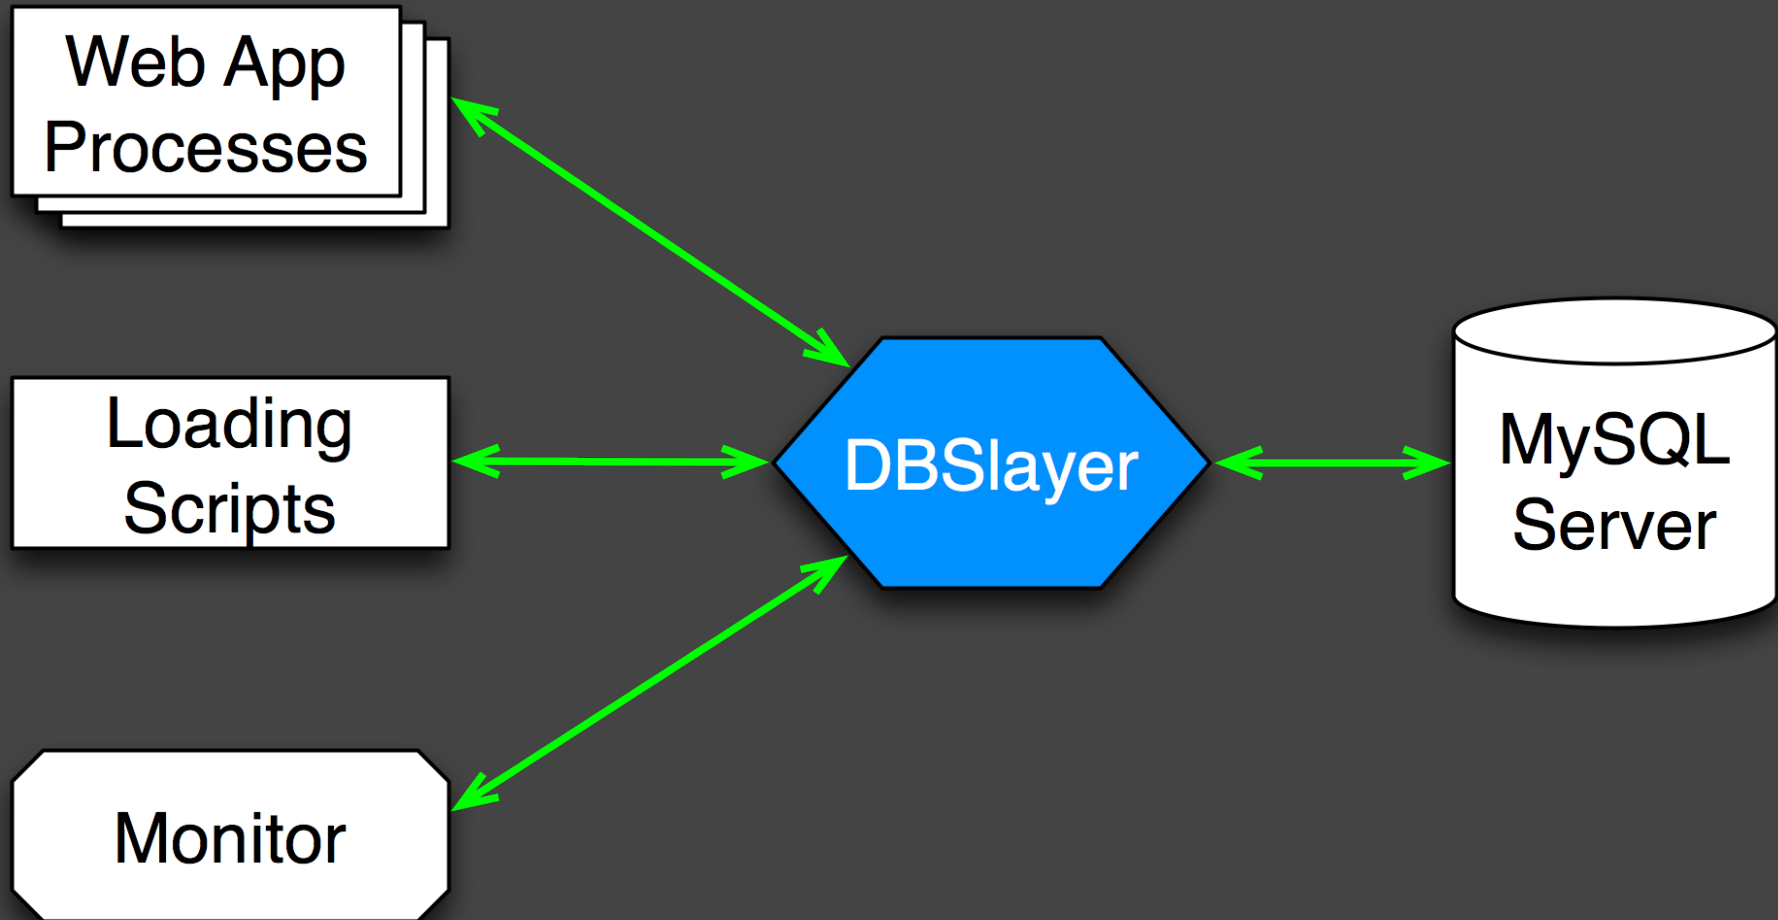
Now what about
that CNX data
backend?

2. DBSlayer

A Database Proxy Layer

- Replacement for the CNX backend
- Layer between applications and databases
- Connection pooling and proxying
- Stateless for scalability
- Speaks HTTP protocol
- Returns data in JSON format
- Absurdly Simple codebase

A Database Proxy Layer



Build It From Parts

Codebase =

1298 LOC, web server (*using libapr*)

+

520 LOC, database client (*calling libmysqlclient*)

+

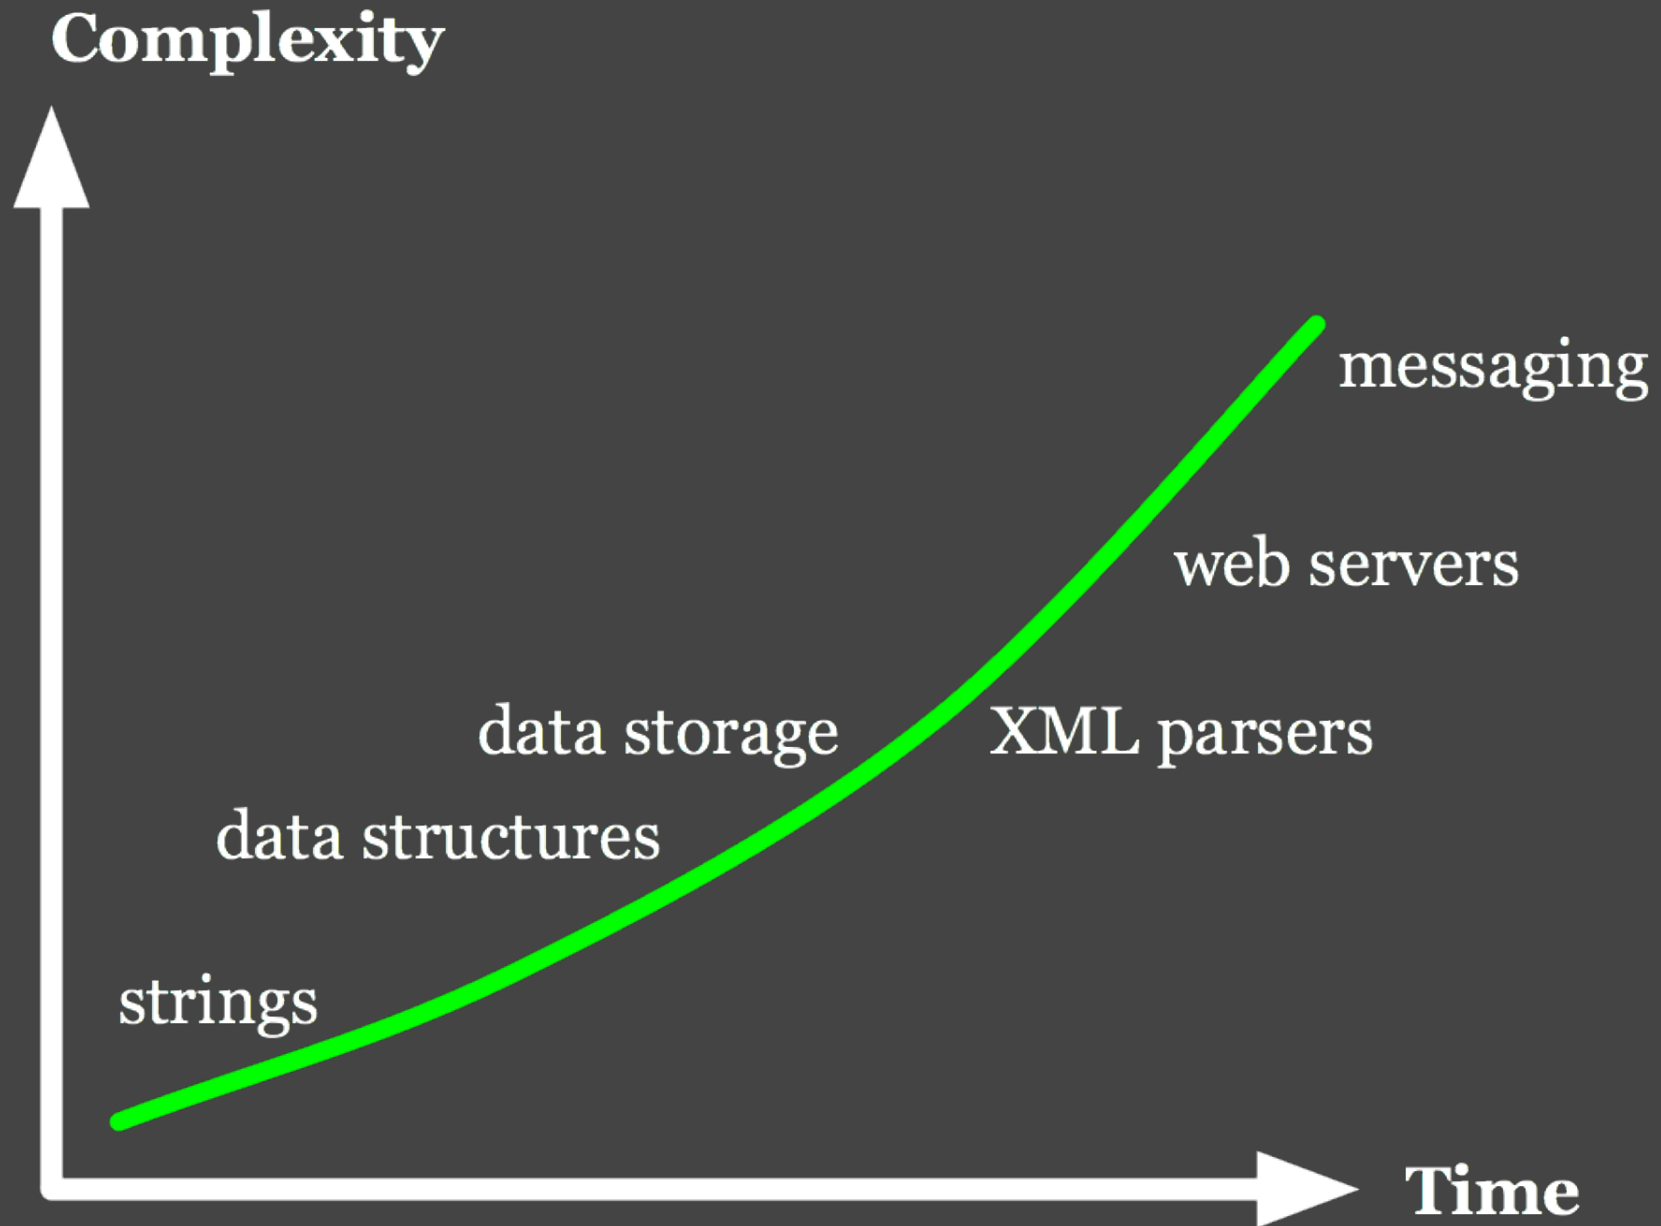
767 LOC, other utility classes (*JSON, etc.*)

=

2585 LOC total

With Open Source,
you don't just get
tools. You get the
libraries to make
your own tools.

Time Is On My Side



Open Protocols

- The CNX backend
 - Custom binary protocol
 - Custom data format
 - Usable only with Context
- DBSlayer
 - HTTP communication
 - JSON data format
 - Usable everywhere

Open Protocols

- Usable everywhere
 - Front-end Web Applications
 - Backend Loading Scripts
 - Monitoring Daemons
 - Web Caching Layers
 - Adapters in any language that can speak HTTP and JSON
 - Statistics Loggers
 - HTTP Packet Sniffers
 - Hand-coded Debugging Requests

Open Source is not just code.
Open Protocols and
Standards are where
Open Source really
kicks ass.

[code.nytimes.com/
projects/dbslayer](https://code.nytimes.com/projects/dbslayer)

Why Open Source It?

- Why Not?

- Giving something back
- We can't sell it
- When other people patch it, we win

3. Times Archive

What's Next

Previous Revolutions in Open Source:

- Compilers
- Applications
- Operating systems
- Component Libraries
- Services and Protocols

The New Frontier:

- **Server Virtualization**

Times Archive

- Opening up Times archives for free
- Millions of article PDFs available
- Each is dynamically created on demand
- Current servers woefully unable to handle projected surge in new traffic
- Traditional approach: buy more servers
- But...
 - Disk space is insanely cheap
 - These PDFs can be cached forever

Why Not Prebuild Everything?

- Static files outscale any demand.
- Generated PDFs can be indexed, linked, and used by other applications
- We live in the era of cheap storage
- All you need is to stitch each PDF:
 - 11 million articles **x**
 - ~12 article stitches / sec
 - = ~100 machine-days of work
- We're going to need some computers...

A COMPUTER WANTED.

WASHINGTON, May 1.—A civil service examination will be held May 18 in Washington, and, if necessary, in other cities, to secure eligibles for the position of computer in the Nautical Almanac Office, where two vacancies exist—one at \$1,000, the other at \$1,400..

The examination will include the subjects of algebra, geometry, trigonometry, and astronomy. Application blanks may be obtained of the United States Civil Service Commission.

Compute Clouds For All

- Amazon Web Services
 - S3: Scalable HTTP-networked storage at reasonable rates
 - EC2: Xen instances rentable on demand by the hour
- Scalable computing for the masses
 - All you need is an idea
 - And your boss' credit card...

Hadoop

- Open-source version of Google's Map-Reduce algorithm
- For parallelizable problems in two phases:
 - Map - perform some action on each data item
 - Reduce - optionally aggregate into a single computed result
- Abstracts away machine-machine coordination
- Also includes a distributed file system
- Error recovery for inevitable node failures



Putting It all Together

4TB of data on S3 (*TIFF images, metadata*)

+

Hadoop

+

JetS3t , iText PDF (*open source libraries*)

+

100 EC2 instances

+

< 1 KLOC glue code

Putting It All Together

11 million PDFs

1.5 TB of data


24 hours

\$240

“I don't kill my own chicken for Sunday dinner anymore, why should I provision my own server?”

BigJimmy's Blog

Now imagine doing
that with Windows



Microsoft Windows Vista Business FULL VERSION [DVD] [Old Version]

Other products by [Microsoft](#)

Platform: Windows XP / Vista / 2000

★★★★☆ (20 customer reviews)

List Price: ~~\$299.95~~

Price: **\$215.69**

You Save: **\$84.26 (28%)**

Availability: In Stock. Ships from and sold by [NothingButSoftware](#).

27 used & new available from **\$174.99**

[Share your own customer images](#)

Watch **the demo** [Flash]



Buy Now or Pre-Order with Service Pack 1

Buy this version now, then update to SP1 online--free of charge--when the new versions are released on March 18, 2008. Or pre-order [Windows Vista Business with SP1](#) now to get SP1 included in the box.

\$215.69 + \$4.99 shipping
In Stock. Ships from and sold by [NothingButSoftware](#)

Quantity:

Add to Shop

or [Sign in](#) to turn on 1...ing.

More Buying

J&R Music and Computer World

\$219.97 + \$6.09 shipping

In stock. Processing time additional 3 to 4 days

Amazon.com

\$279.99 & this item

with Super Saver Shipping
In Stock

ANTOnline

\$340.25 + \$6.58 shipping

In Stock

27 used & new from

But Amazon won't tell me an address to ship it to...



Price per License for Each Edition

Standard* \$

Enterprise* \$

Datacenter \$

Please use model A or B

A. Group servers by # of processors and average number of VMs

B. List each server individually

[Print](#) [Reset](#) [Need Help?](#)

* Standard and Enterprise editions include CALs. A single Standard Edition license does not grant rights to run both a host OS and guest OS, so an additional license is required for the host OS. Standard Edition is able to run on 1-4 processor servers and Enterprise Edition is able to run on 1-8 processor servers.

[Disclaimer](#)

A. Processor	# of Servers	Avg # of VMs per Server	Number of Licenses Needed			Price for Each Edition		
			Standard	Enterprise	Datacenter	Standard*	Enterprise*	Datacenter
1	<input type="text" value="0"/>	<input type="text" value="0"/>	0	0	0	\$ 0	\$ 0	\$ 0
2	<input type="text" value="0"/>	<input type="text" value="0"/>	0	0	0	\$ 0	\$ 0	\$ 0
4	<input type="text" value="0"/>	<input type="text" value="0"/>	0	0	0	\$ 0	\$ 0	\$ 0
8	<input type="text" value="0"/>	<input type="text" value="0"/>		0	0		\$ 0	\$ 0
16	<input type="text" value="0"/>	<input type="text" value="0"/>			0			\$ 0
32	<input type="text" value="0"/>	<input type="text" value="0"/>			0			\$ 0
64	<input type="text" value="0"/>	<input type="text" value="0"/>			0			\$ 0
			0	0	0	\$ 0	\$ 0	\$ 0

B. Server	# of Processors	# of VMs on Server	Number of Licenses Needed			Price for Each Edition		
			Standard	Enterprise	Datacenter	Standard*	Enterprise*	Datacenter
1	<input type="text" value="0"/>	<input type="text" value="0"/>	0	0	0	\$ 0	\$ 0	\$ 0
2	<input type="text" value="0"/>	<input type="text" value="0"/>	0	0	0	\$ 0	\$ 0	\$ 0
3	<input type="text" value="0"/>	<input type="text" value="0"/>	0	0	0	\$ 0	\$ 0	\$ 0
4	<input type="text" value="0"/>	<input type="text" value="0"/>	0	0	0	\$ 0	\$ 0	\$ 0
5	<input type="text" value="0"/>	<input type="text" value="0"/>	0	0	0	\$ 0	\$ 0	\$ 0
6	<input type="text" value="0"/>	<input type="text" value="0"/>	0	0	0	\$ 0	\$ 0	\$ 0
7	<input type="text" value="0"/>	<input type="text" value="0"/>	0	0	0	\$ 0	\$ 0	\$ 0
8	<input type="text" value="0"/>	<input type="text" value="0"/>	0	0	0	\$ 0	\$ 0	\$ 0
9	<input type="text" value="0"/>	<input type="text" value="0"/>	0	0	0	\$ 0	\$ 0	\$ 0
10	<input type="text" value="0"/>	<input type="text" value="0"/>	0	0	0	\$ 0	\$ 0	\$ 0
11	<input type="text" value="0"/>	<input type="text" value="0"/>	0	0	0	\$ 0	\$ 0	\$ 0
12	<input type="text" value="0"/>	<input type="text" value="0"/>	0	0	0	\$ 0	\$ 0	\$ 0
13	<input type="text" value="0"/>	<input type="text" value="0"/>	0	0	0	\$ 0	\$ 0	\$ 0
14	<input type="text" value="0"/>	<input type="text" value="0"/>	0	0	0	\$ 0	\$ 0	\$ 0
15	<input type="text" value="0"/>	<input type="text" value="0"/>	0	0	0	\$ 0	\$ 0	\$ 0
			0	0	0	\$ 0	\$ 0	\$ 0

Luckily, Microsoft is making it simpler...

Eventually, Microsoft
will figure it out.

And I'm sure it'll be
cheaply priced for
everybody. No doubt.

A Quick Recap

1. Context

- New Choices in Open Source
- Strategies For Switching

1. DBSlayer

- Building From Open Parts
- Embracing Open Standards

1. Times Archive

- The Compute Cloud Revolution
- Challenges For Licensed Software

Search



All the Code That's Fit to printf()

March 12, 2008, 1:55 pm

Join the 2008 Interns at NYTimes.com

By PLATFORM TECHNOLOGY

Do you thrive on difficult engineering challenges? Does your pulse quicken when you think about solving problems involving massive datasets or the technical hurdles in making search more relevant, personalization more powerful and easier to use, and delivering news and information instantly to millions of people around the world? Our platform team is responsible for the core software technologies that cut across all of our products and services. Internship opportunities are available for...

Search This Blog

All NYTimes.com Blogs »

About Open

A blog about open source technology at The New York Times, written by and primarily for developers. This includes our own projects, our work with open-source technologies at nytimes.com, and other interesting topics in the open source and Web 2.0 worlds.

DBSLayer

A simple database abstraction and connection pooling layer built on top of HTTP and JSON.

XSL Cache

A PHP extension that caches parsed XSLT stylesheets for faster XSL transforms.

Monthly Archives

Select Month

open.nytimes.com