GT.M Update

K.S. Bhaskar

Development Director, FIS GT.M™

Empowering
the Financial World

FIS

Overview

- **Business**
- **Technical**

Business

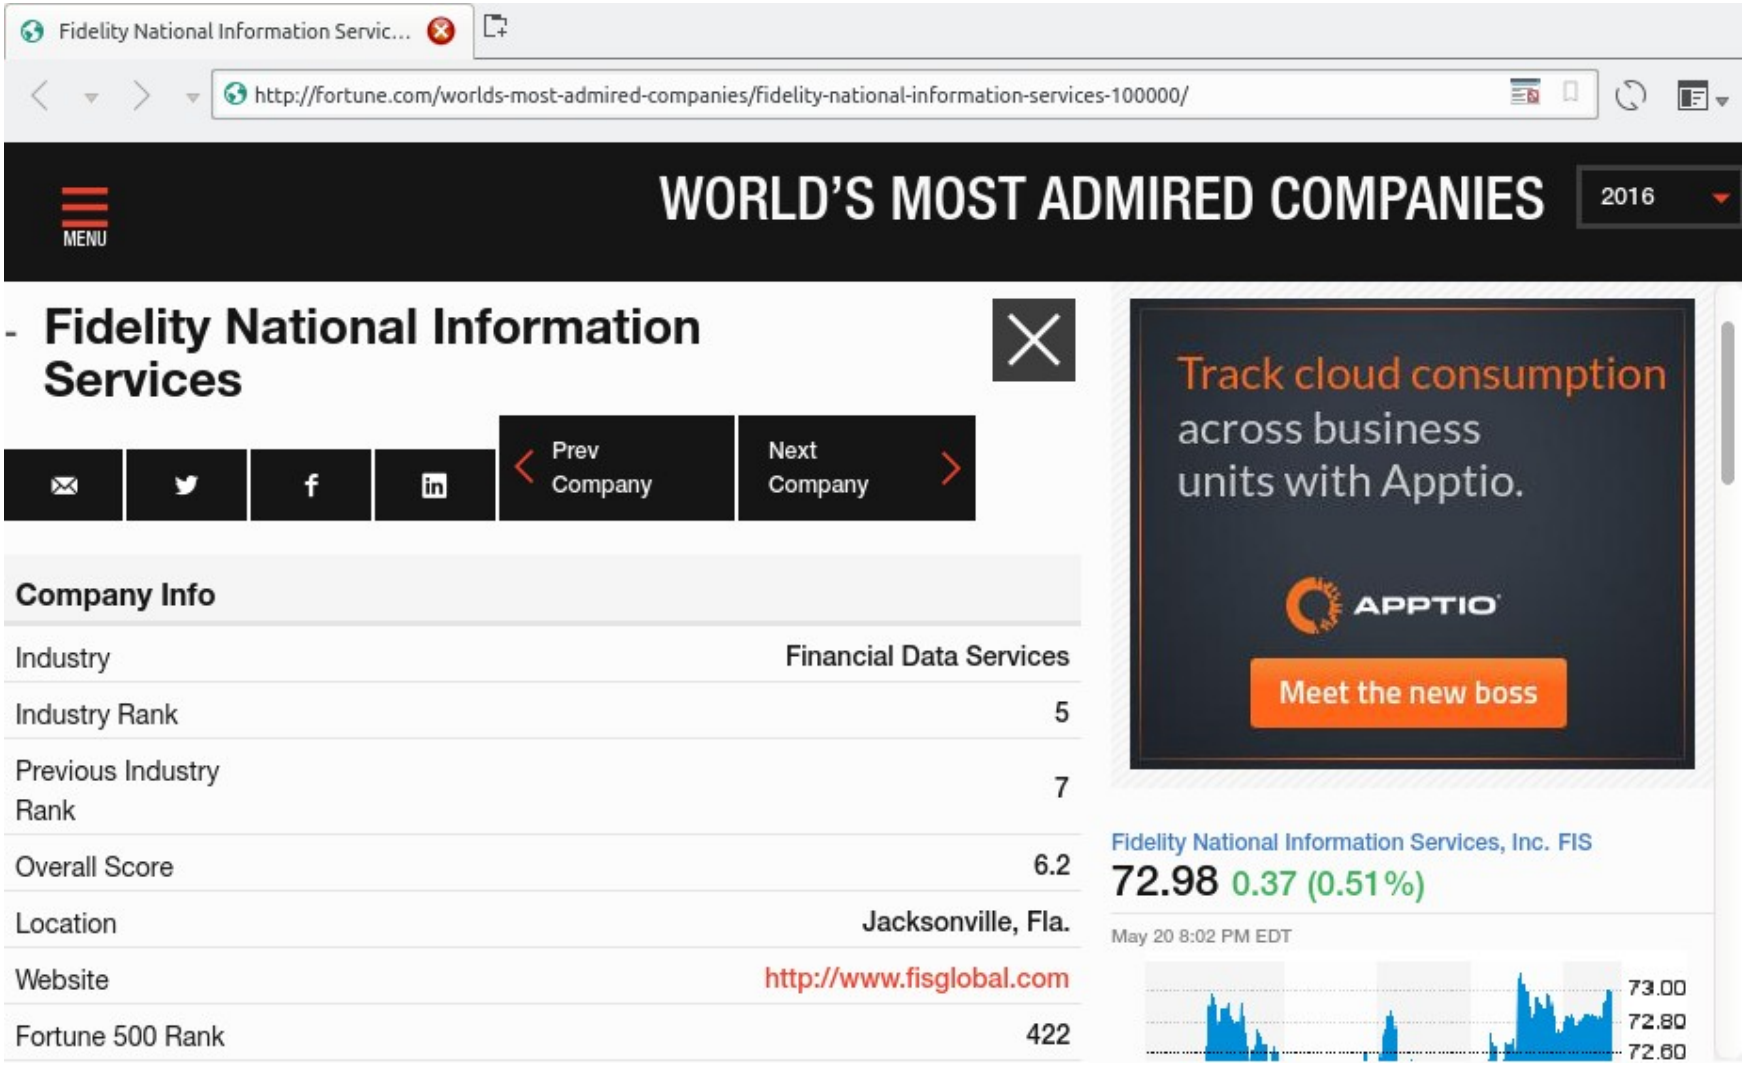
Fidelity National Information Servic...
http://fortune.com/worlds-most-admired-companies/fidelity-national-information-services-100000/
MENU
WORLD'S MOST ADMIRED COMPANIES
2016
- Fidelity National Information Services
Prev Company
Next Company
Company Info
Industry
Financial Data Services
Industry Rank
5
Previous Industry Rank
7
Overall Score
6.2
Location
Jacksonville, Fla.
Website
http://www.fisglobal.com
Fortune 500 Rank
422
Track cloud consumption across business units with Apptio.
APPTIO
Meet the new boss
Fidelity National Information Services, Inc. FIS
72.98 0.37 (0.51%)
May 20 8:02 PM EDT
73.00
72.80
72.60

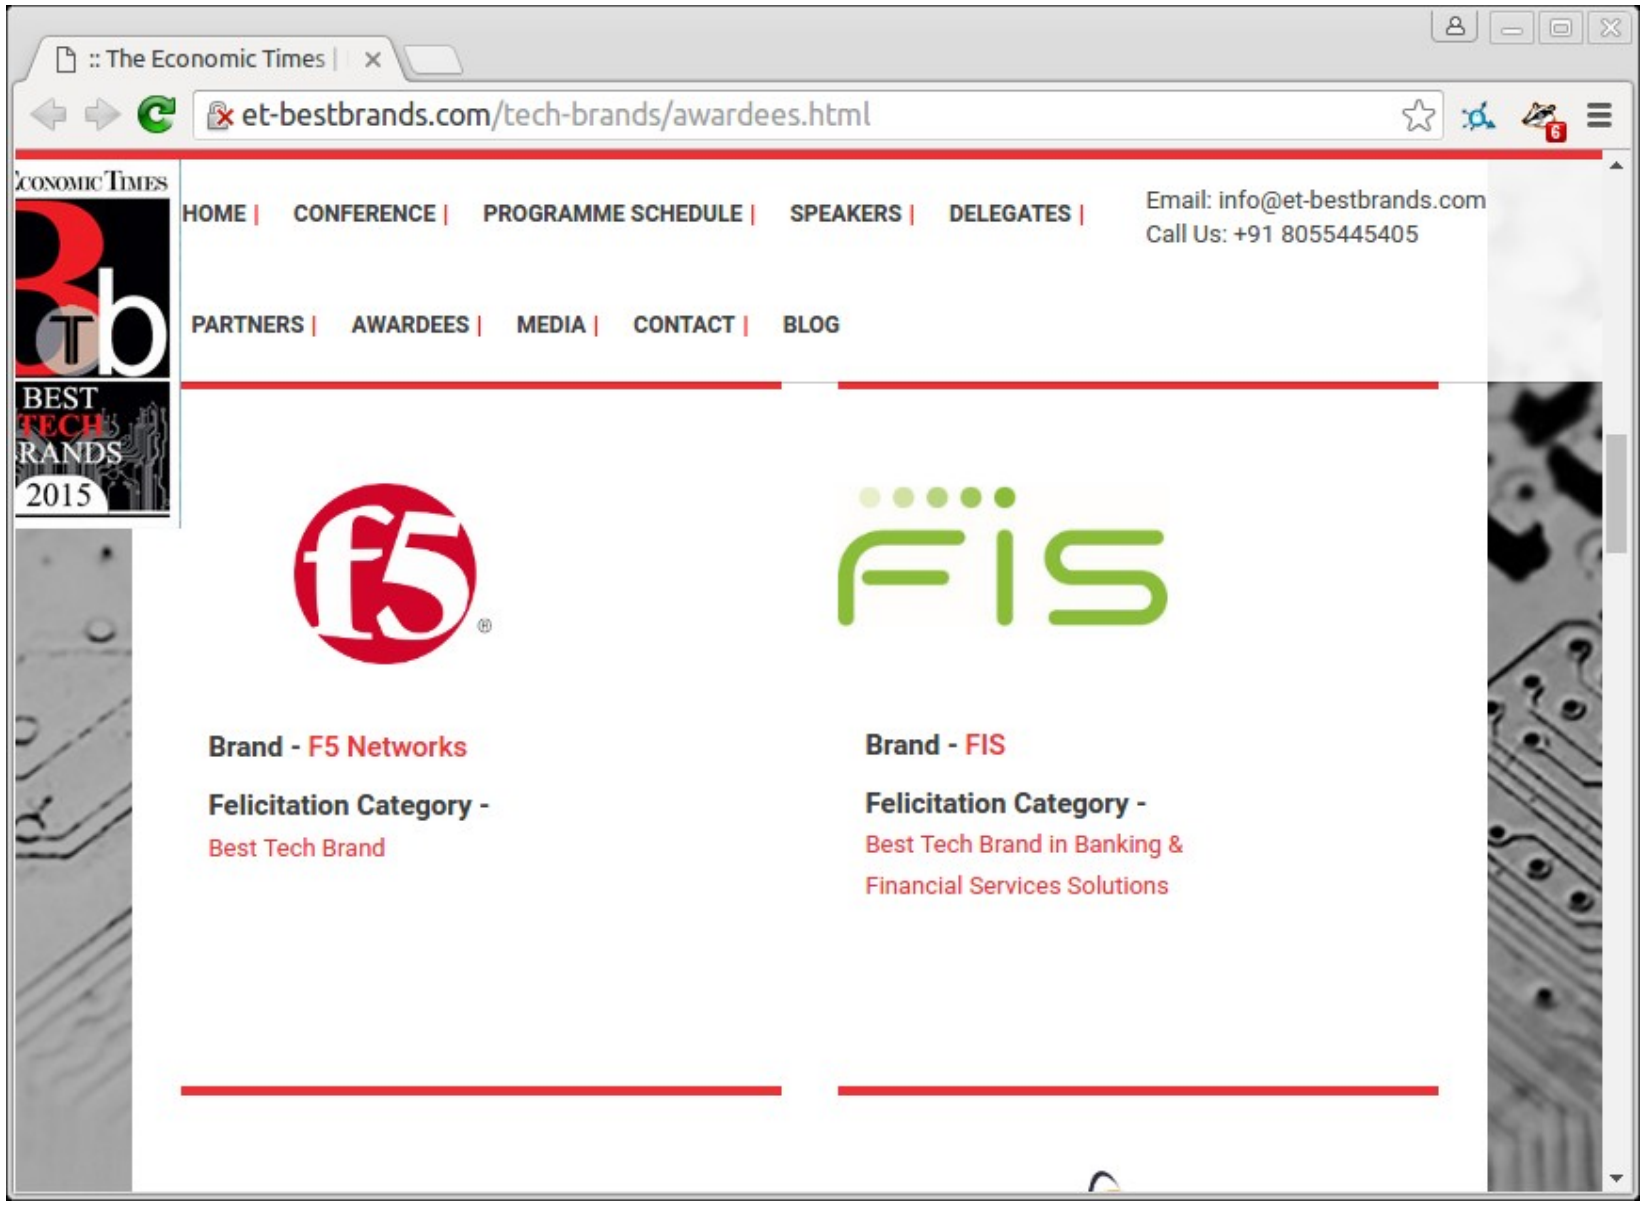
:: The Economic Times
et-bestbrands.com/tech-brands/awardees.html
HOME | CONFERENCE | PROGRAMME SCHEDULE | SPEAKERS | DELEGATES |
Email: info@et-bestbrands.com
Call Us: +91 8055445405
PARTNERS | AWARDEES | MEDIA | CONTACT | BLOG
BEST TECH BRANDS 2015
Brand - F5 Networks
Felicitation Category -
Best Tech Brand
Brand - FIS
Felicitation Category -
Best Tech Brand in Banking & Financial Services Solutions

JT E-health programme to

www.jordantimes.com/news/local/e-health-programme-be-implemented-nationwide

THE JORDAN TIMES

WEATHER, AMMAN **Today,** May 22, 2016

30 c° / Sunny

Home » Local » E-health programme to be implemented nationwide

E-health programme to be implemented nationwide

By **Khetam Malkawi** - Nov 16,2015 - Last updated at Nov 16,2015

AMMAN — The government on Monday signed a five-year framework agreement with the Electronic Health Solutions Company (EHS) to regulate the implementation of the e-health programme "Hakeem" at the national level.

The agreement — signed by Planning and International Cooperation Minister Imad Fakhoury and EHS Chairman Rami Farraj — aims at coordinating cooperation between EHS — the implementing party of Hakeem programme — and the concerned

Planning and International Cooperation Minister Imad Fakhoury and EHS Chairman

Technical

VistA on GT.M on Windows

Hats off to Christopher Edwards, Sam Habiel, and Cygwin

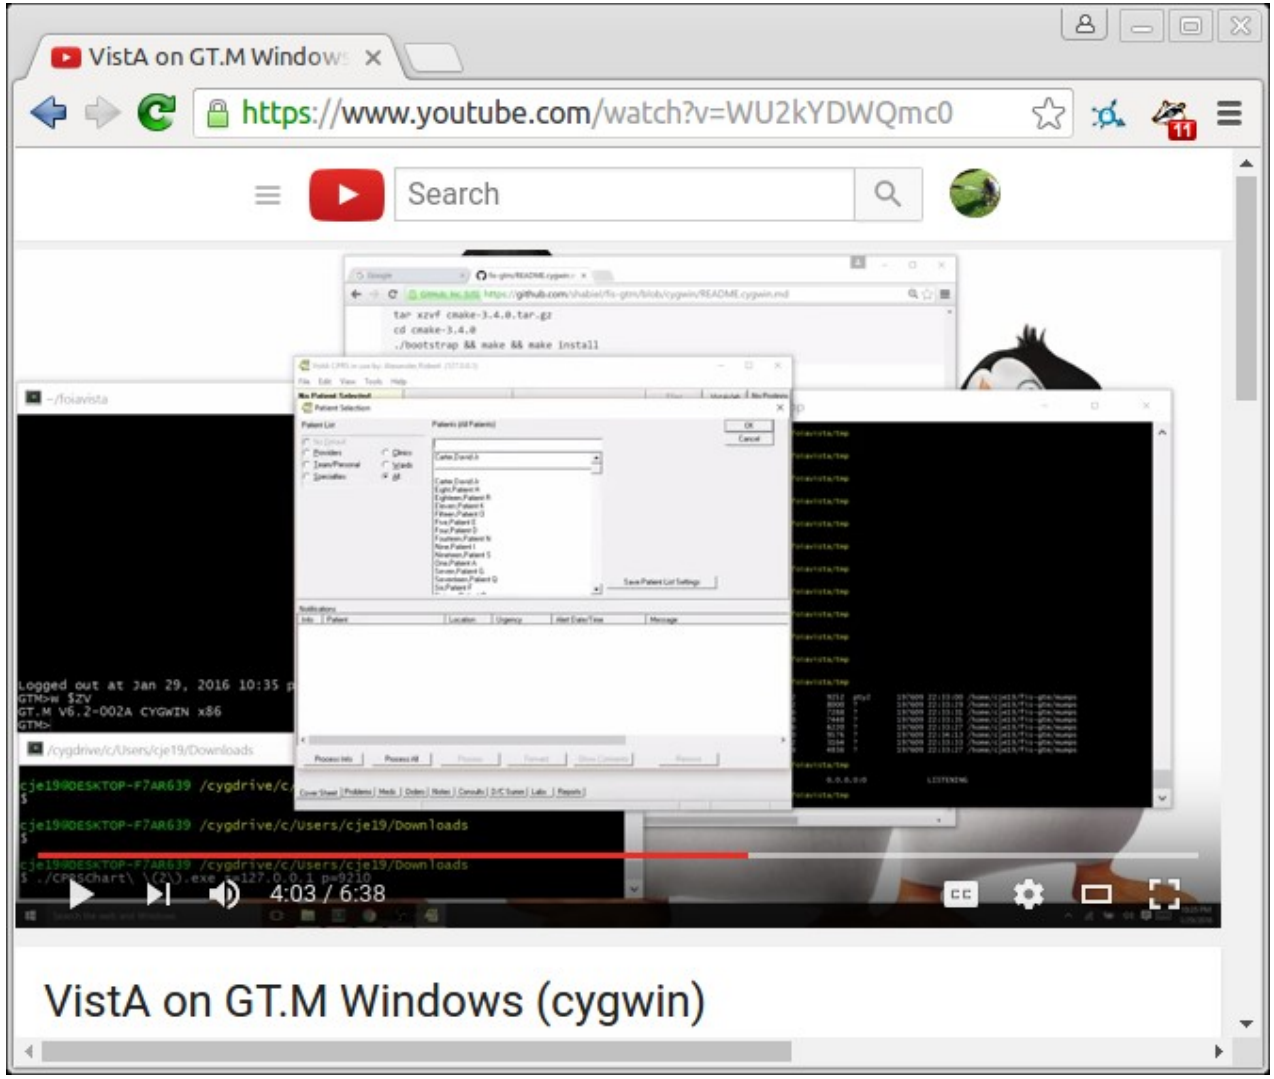

Releases

- GT.M
  - V6.3-000 – major new enhancements
  - V6.3-000A – timely availability of fixes
- Plugins
  - Peer replication

GT.M Functionality

- **Security – Reduce “attack surface”**
- **Performance**
- **And more**

Change database encryption password “on the fly”

- Database files vulnerable as a consequence of large data volumes and long lived keys
- Defense is to change keys periodically
  - Pre-V6.3-000: extract and load into new database
  - V6.3-000: Define two keys in configuration file, and run MUPIP REORG ENCRYPT to change encryption key while database is in use
- Can also encrypt unencrypted database files

Other Security Enhancements

- TLS key renegotiation
  - Protection against attacker who records a session for future
  - Effect with WRITE /TLS("renegotiate"[„[tlsid][„options]])
- Initialization vectors for database encryption – different for each block
  - Better protection against attacks seeking to exploit large volumes of data
  - Automatic for databases created with V6.3-000
  - Zero IV retained for databases created with older releases – changing the encryption key switches to using initialization vectors

Performance and Scalability

- Faster process termination
  - Especially when large numbers of processes and many regions both exist
- Option to allow more than 32Ki concurrent processes accessing a database region
- String expressions concatenation, $[Z]ASCII(), $[Z]EXTRACT(), $[Z]PIECE(), and $ZSUBSTR() on literals computed at compile time
  - In UTF-8 mode can result in BADCHAR warnings at compile time (errors at run time)
- Performance optimization for certain common UTF-8 mode use cases
- MUPIP JOURNAL commands have PARALLEL[=n] qualifier
- Enhanced replication throughput

And more...

- %PEEKBYNAME() gives symbolic access to internal statistics provided by $ZPEEK()
  - LISTALL^%PEEKBYNAME / LIST^%PEEKBYNAME(.var) list all the fields, but sheer number is overwhelming
  - Fields not documented because internals can change from release to release, even if most may not – work with GT.M support on specific needs
- MUPIP JOURNAL ROLLBACK FORWARD
  - Better / faster recovery from certain catastrophic operational failures, e.g., affecting multiple locations

Plugins

- Extend GT.M with useful functionality, but not part of core GT.M
  - Different release cycles from GT.M, and one another
- Usable “out of the box”, but intended to be extensible
  - “Copy and extend” is a common programming model
- Released under the same license as GT.M
- Unmodified plugins supported as part of GT.M support, modified plugins not supported by FIS
- Existing plugins – zlib, POSIX, GTMJI
- New – Peer Replication

Serialization of Updates / Transactions

- **Business logic that needs serialization**
  - **Order is required, e.g., each transaction on an account depends on the result of the preceding transaction on that account; a diagnosis must be made or at least suspected before a treatment can be prescribed**
  - **Unserialized updates cannot be aggregated**
- **Business logic that doesn't need serialization**
  - **Order can be relaxed, e.g., a balance inquiry on an account does not depend on the result of the previous balance inquiry (involves updates because each access to an account must be recorded)**
  - **Unserialized updates can be aggregated**

BC / SI Replication vs. Peer Replication

- Centralized update serialization performs better than distributed update serialization
  - As long as central decision maker is not resource constrained (CPU, RAM, IO)
  - Systems like VistA (even EHS Jordan's Hakeem) are nowhere near limits of today's computer systems
  - BC / SI replication to distribute serialized updates from central decision maker
  - Unidirectional, implemented by GT.M, which maintains order, and restores it on recovery from a partitioning event
- Distributed updates can be closer to client, are more resilient to network outages, can use more economical computing resources, and can have greater elasticity with varying load
  - Peer replication to aggregate updates from multiple service points
  - Uni- or bidirectional, implemented by application code, which recovers aggregation (but not order) on recovery from a partitioning event

Sample configuration

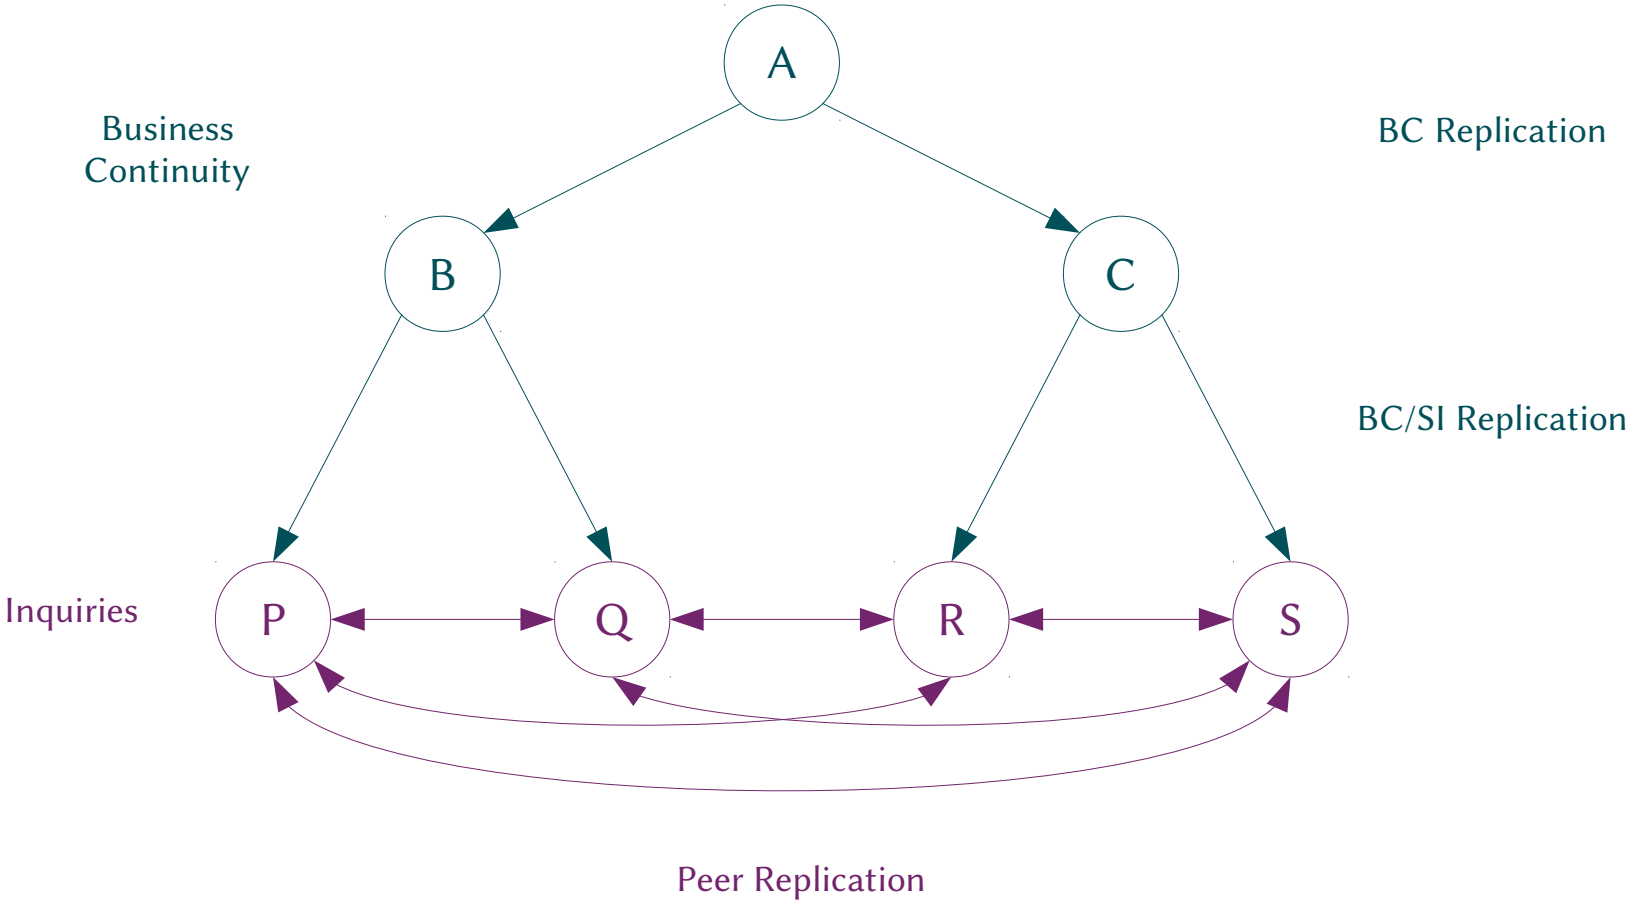

Peer Replication

- Implemented with pure application level code
- Triggers to capture updates to be replicated using peer replication
- Push / pull processes to monitor & replicate captured updates between instances

Peer Replication Schematic

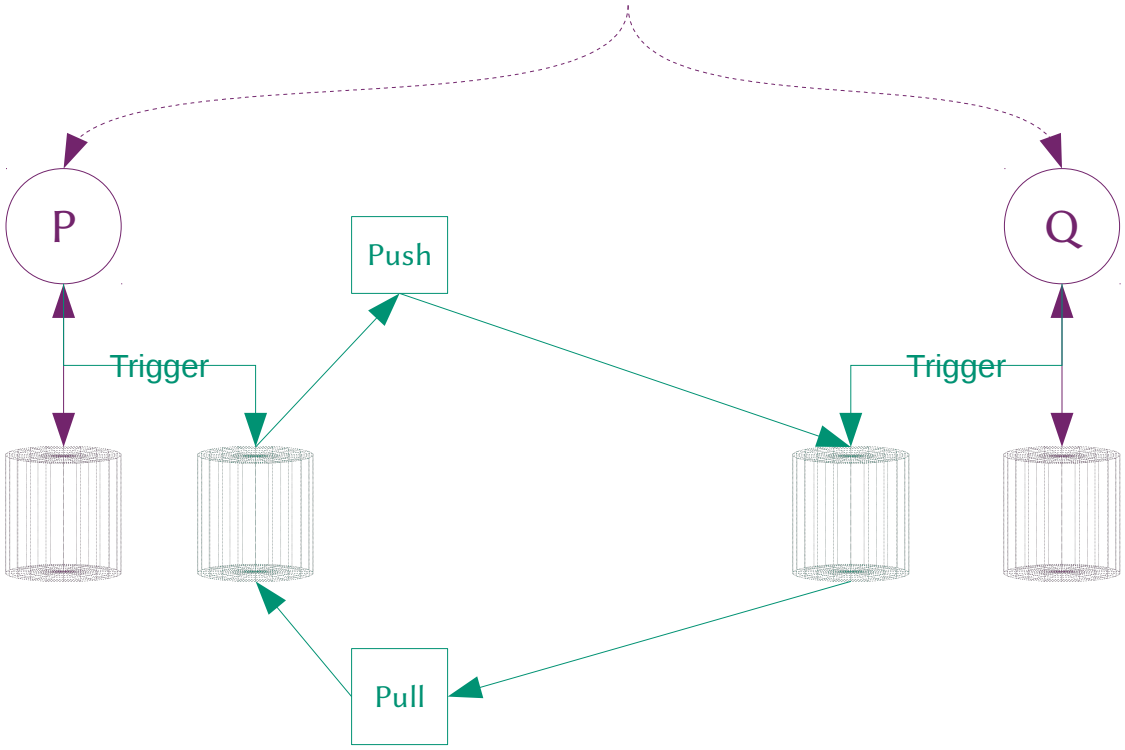

Peer Replication Distribution

- **Reference implementation of plugin is one routine**
- **Packaged with small sample application and scripts to demonstrate / test software**
- **Documented in readme and with help:** `mumps -run GTMPeerRepl –help`

Questions and Answers

Contact – note new address and office landline!

K.S. Bhaskar
Development Director, FIS GT.M
200 Campus Drive
Collegeville, PA 19426, USA
ks.bhaskar@fisglobal.com
+1 (484) 302-3174
http://fis-gtm.com