

## High-Volume Transaction Processing

- Me
  - Chuck Boecking
  - 20+ year with iDempiere codebase
- Customer profile
  - 100K+ invoices per month
  - 1.5M+ GL Transaction
  - 60+ Organizations
  - ISP serving residential homes and commercial offices
- Monitoring - you need it
  - Zabbix
  - AWS Cloudwatch
- Heap - what is it and how to monitor it
  - Heap dump - leak vs usage
  - 3 levels
  - [Example script to extract free heap](#)
- Document Sequence
  - Native
  - UUID docno
- Table and Table Direct is a killer
  - [SQL performance utility](#)
    - update\_table\_dir\_view.sql
  - Would be nice:
    - Deferred drop-down lookup
    - Search window that auto populates
- Statement closed
  - MaxStatementsPerConnection =>  
/opt/idempiere-server/PostgreSQL/pool.properties
- Parallel multi-threaded
  - Ex: [zibraHashPasswords.java](#)
- Would be nice: Window => High volume => Search with no count
- High-speed export (copy command)
- Financial Reporting (60 organizations) - impacts window performance
- Auto allocation
- Payments into batch
- Delete record - dependency on FK
- Query for query effectiveness
  - [chuboe\\_debug\\_query\\_index\\_stats.sq](#)
- Materialized view
- Search using info window
- Zram
- When things go badly - contamination count is big
- Budget for optimization - bigger than you thing
- Activities will kill you system!

- 80 concurrent users per server limit - jetty connection pools?
- New FK in big/important tables (c\_invoiceline)
- Upgrading to a new version - performance testing...