

Paying Lip Service to Business Continuity



Paul Koufalis, White Star Software
pk@wss.com



More than Disaster Planning

- 1000 little things can go wrong
 - So many moving parts: hardware, software, VMWare, network, fabric, SAN
- App can be offline for minutes, hours, days
 - Or maybe just painfully slow
- **Bottom line: is the business affected?**
 - Not about esoteric performance metric

Business Continuity Plan Requirements



- Clear SLA between business and I.T.
- Proper database and system administration
- Monitoring and alerting
- Continuous evolution and improvement

Paul Koufalis



- Progress DBA and UNIX admin since 1994
- Providing expert OpenEdge technical consulting
- Wide range of experience
 - Small 10 person offices to 3500+ concurrent users
 - AIX, HPUX, Linux, Windows...if Progress runs on it, I've worked on it
- Father to these two monkeys



pk@wss.com

Who is White Star Software?



- The oldest and most respected independent DBA consulting firm in the world
- Five of the world's top OpenEdge DBAs
- Author of ProTop, the #1 FREE OpenEdge Database Monitoring Tool
 - <http://protop.wss.com>



Today's Topics

- Recent business continuity examples
- Realistic service level agreements
- Common and avoidable problems
- Low hanging fruit
- Proactive or reactive?
- Finding the sweet spot along the cost-benefit curve



Recent Events

- Distributor: DB corruption
 - Some piece of hardware went URCKKK!!
 - DB was smashed
 - Sorta/kinda BCP plan – not usable
- **Down time: 12 hours**



Recent Events

- SAME CUSTOMER one month later
 - Progress executable corruption
- Down time: None
- **Pain time: 16 hours**
 - Running agents were fine but could not start new ones. Users and web suffered badly



Recent Events

- Manufacturer: VMware VMotion bug
 - Their products part of supply-chain of customers
 - Live VMotion (high availability !?! Riiiiight...) corrupted EVERYTHING
- **Down time: 30 hours**
 - Saving grace: Happened on Friday
 - Business impact was less severe



Recent Events

- Financial Services: FIRE !!!
 - Electrical panel caught fire
 - Data centre ok but 13 story building with no power
 - 2500 people with no computer/phone
 - Detailed BCP plan = 100 offsite workstations
 - Back to normal the next day at 4:00 AM
- **Down time: Officially zero**
 - The application was available



Dumb, Preventable Events

- Server hang due to full disk
 - FedEx log file
- Database crash
 - Locked files on Windows
 - Extent hit 2 GB “limit”
- All AI files full = DB stall or crash
 - This happens more often than you would think
 - Double bad: backups fail



Realistic SLA

- 24 years experience: businesses show little appetite for SLAs
- Too busy selling widgets
 - Concentrated on selling even more widgets
- Especially when everything is going *well*
 - No one wants to spend money <ahem> for nothing <ahem>



Realistic SLA

- Business assumes I.T. has I.T. stuff covered
 - W/out a written SLA, unlikely I.T. and business aligned
- With an SLA, it will be still be I.T.'s fault
 - But at least you have something to back you up
- Who has a clearly defined SLA to the business?



Realistic SLA

- Ask each business unit to assess impact of downtime
 - Manufacturing, shipping, finance...
 - What can you **not** do if app is down?
 - Do you have a manual workaround?
- Discuss outage scenarios: 1h, 4h, 24h down...
 - These things happen in the real world. They will happen to you
 - Enough work in the pipeline for an hour? A day?
- Discuss time-of-year outages
 - Spring (home improvement), Christmas (B2C), etc.



Realistic SLA

- Ask the business units the impact of bad performance
 - What if MRP isn't finished at 5:52 AM ?
- What about DB maintenance activity?
 - Backup = 27 hours
 - Corrupt index rebuild = 12h
 - Undo-redo processing after crash = 6h



Realistic SLA

- Clearly present impact to management
- Do NOT try to sugar-coat your findings
 - I.T. sometimes scared to tell management the truth
 - Don't want to look bad or be the bearer of bad news
- Get some guidelines from management
 - Maybe 15-minute SLA is overkill but losing a full day is out of the question
- Start devising a rough plan with cost estimates
- Go back to management for another round

Common & Avoidable Problems



I can't believe we went down because ... !

- Disk space: really? In 2017, you ran out of space?
- BI file grew to x GB and crashed
 - **DB start-up can take hours**
- AI files filled and locked
 - DB will crash or stall
- AppServer agents not available/locked
 - Spotty performance, eventual system hang-up
- Improper configuration
 - NOT a one-time task

Common & Avoidable Problems



Who is going to tell the CEO?

- Your backups haven't been valid for *how long !?!*
 - Restored backup from 2016-03-17. Uh-oh...
- We lost **how many hours** of data?
 - How do we get it back? What do you mean we can't!?!
- Performance is terrible
 - Suffering in silence
 - Users accept it = significant lost productivity
 - No one says anything because it's "normal"



Low Hanging Fruit

- Validate successful backups
 - Partial verify: block CRC check
 - Full verify: restore somewhere
- Enable after-imaging
 - Zero impact to the business
 - Ability to restore to an exact point in time
 - Protect against HUMAN ERROR
 - **Move the archives offsite**
- Configure DBs and other components properly



Low Hanging Fruit

- Upgrade to latest version of OpenEdge
- Professional health check of your environment
 - Easy and inexpensive
 - I am often surprised by what I find
- Monitoring and alerting
 - Roll your own or use existing tools

Proactive or Reactive Monitoring?



- Everyone already has a critical monitoring system
 - Your users!

- Reactive monitoring may be good enough for you
 - System crashes
 - Users call help desk
 - Help desk calls Mark
 - Mark does his magic
 - An hour later, everything is back to normal



Reactive Monitoring

- Mostly relies on luck
- You hope the issue will be minor
 - *“We’ve been running QAD for 20 years w/out a problem”*
- Problems often discovered accidentally
 - Ex.: restore backup in test environment and realize transactions are months old

Reactive Monitoring



- Monitoring is adjusted *after* each new type of event
- Your business processes may be resilient enough to absorb unplanned outages
 - Or not
- Do your customers live in a “6-8 weeks for delivery” world? Should you?



Proactive Monitoring

- Proactive approach is clearly better
- There is a cost associated
- Write your own tools
 - Costly to develop and maintain
 - Never comprehensive (reactive improvement)
 - Mish-mash of *stuff* accumulates over the years
- Use an established service like ProTop
 - Fixed cost
 - Comprehensive and constantly improving
 - Development/maintenance not your problem
 - Benefit from lessons learned by other users

Proactive Monitoring: Minimum Monitoring Points



- Database, UBrokers and other components up/down
- File system size
- Database
 - BI size
 - Extent sizes (WG limited to 2 GB)
 - Long transaction
 - AI and AI Archiver status
 - Replication status
 - Blocked users and deadlocks
 - Log file error messages
 - Backup Age
- Monitor the monitor



BCP Cost/Benefit

- At a minimum, implement low hanging fruit
 - Validated backups, after-imaging, health check, modern infrastructure
- “Next steps” examples in the next few slides

BCP Cost/Benefit Sweet Spot



- In-house monitoring
- Relative Cost: MEDIUM
- Complexity: HIGH
- Risk: HIGH (improve after incident)
- Down Time: N/A
- Data Loss: N/A

BCP Cost/Benefit Sweet Spot



- Professional Monitoring Service like ProTop
- Relative Cost: MEDIUM
- Complexity: LOW
- Risk: LOW
- Down Time: N/A
- Data Loss: N/A

BCP Cost/Benefit Sweet Spot



- Plan: Cold restore
 - Buy new HW (or provision VM)
 - Restore everything
- Relative Cost: LOW
- Complexity: VERY HIGH
- Risk: HIGH (unless procedure is well tested)
- Down Time: Hours to days (depends on HW avail)
- Data Loss: 15 - 30 minutes typical



BCP Cost/Benefit Sweet Spot

- Plan: Warm Spare
 - Provisioned fail-over HW up-and-running
 - Static data sync'd in near real-time
 - Backups and AI files sync'd in near real-time (ftp/scp)
- Relative Cost: MEDIUM-HIGH (licenses)
- Complexity: MEDIUM
- Risk: LOW
- Down Time: < 1 H
- Data Loss: 15 – 30 minutes typical

BCP Cost/Benefit Sweet Spot



- Plan: Hot Spare
 - HW/VM provisioned, equivalent to PROD
 - Static data sync'd in near real-time
 - Database changes sync'd in real-time
- Relative Cost: HIGH
- Complexity: HIGH
- Risk: MEDIUM
- Down Time: Minutes
- Data Loss: Zero-ish

BCP Cost/Benefit Sweet Spot



- Plan: Cluster + Hot Spare + DR site
 - Live PROD cluster box on same SAN
 - HW/VM provisioned for DR, equivalent to PROD
 - Static data sync'd in near real-time
 - Database changes sync'd in real-time
 - DR site for users
- Relative Cost: VERY HIGH
- Complexity: VERY HIGH
- Risk: HIGH (not tested adequately), otherwise MEDIUM
- Down Time: Minutes
- Data Loss: Zero-ish



Take Away Message

- Define an SLA with the business, no matter how simple
- Make sure you implement the basic monitoring recommendations (easy and cheap)
- Find your company's sweet spot along the BCP cost/benefit curve

Questions?



Thank You!



protop

#1 OpenEdge Database Monitoring Tool

<http://protop.wss.com>