HowTo DR

Josh Berkus
PostgreSQL Experts
pgCon 2014

Disaster Recovery

"The process, policies and procedures that are related to preparing for recovery or continuation of technology infrastructure which are vital to an organization after a natural or human-induced disaster."

Wikipedia, February 2014

Disaster Recovery

Restoring services after the unexpected.

Disaster Recovery

Limiting:

1. Downtime
2. Data Loss

Do you have a DR Plan?

Is it fairly complete?

Have you *tested* it?

Our Disaster Recovery Plan Goes Something Like This...
HELP! HELP!
DILBERT
By Scott Adams

Dilbert is a copyright of Scott Adams. Used here as parody. All rights reserved to Scott Adams.

Threat Model

server
failure

**getting
hacked**

***natural
disaster***

server failure

storage failure

network failure

getting hacked

traffic spike

admin error

OS / VM problem

natural disaster

bad update

software bugs

server failure

storage failure

network failure

getting hacked

traffic spike

admin error

OS / VM problem

natural disaster

bad update

software bugs

Accepting Loss

The Nines

Nines	Down/Year
99.9%	9 hours
99.99%	1 hour
99.999%	5 minutes

The Nines

- Treats all downtime causes as identical
  - except the ones it ignores
- Doesn't address data loss
- Really “Business Continuity”
- also unrealistic

Disaster	Downtime	Data Loss	Detect
Server Failure			
Network Failure			
Admin Error			
Bad Update			
Storage Failure			
Getting Hacked			
Natural Disaster			

Disaster	Downtime	Data Loss	Detect
Server Failure	0	0	
Network Failure	0	0	
Admin Error	0	0	10 yrs
Bad Update	0	0	
Storage Failure	0	0	
Getting Hacked	0	0	10 yrs
Natural Disaster	0	0	

$$$$$$$$$$$$$$$$$$$$$$$

$$$$$$$$$$$$$$$$$$$$$$$

$$$$$$$$$$$$$$$$$$$$$$$

$$$$$$$$$$$$$$$$$$$$$$$

$$$$$$$$$$$$$$$$$$$$$$$

$$$$$$$$$$$$$$$$$$$$$$$

$$$$$$$$$$$$$$$$$$$$$$$

Disaster	Downtime	Data Loss	Detect
Server Failure	5min	1min	
Network Failure	3hrs	10min	
Admin Error	1hr	1hr	3 mo
Bad Update	1hr	1hr	
Storage Failure	5min	30min	
Getting Hacked	1hr	1hr	3 mo
Natural Disaster	6hrs	1hr	

$estimation

- Implementation
- Maintenance
- Storage
- Other Infrastructure

Imp / Maint

- Replication/backup setup
- Monitoring
- Troubleshooting
- Training
- Recovery tests

Storage

```
backups
X
( retention + 1 )
```

Storage

(800GB + 20GB)

X

(52 + 1)

= 43TB

Infrastructure

- Replica servers
- Hosting
- Networking & Bandwidth

Your DR Plan

Elements of a Plan

1. Backups/Replicas
2. Replacements
3. Procedures
4. People

Backups

- pg_dump - “logical”
- pg_basebackup - “binary”
- snapshot + PITR - “binary”

Backups++

- Periodic
- Portable
- Simple
- Recover point-in-time

Backups--

- Slow to restore
- Data loss interval

pg_dump

- very portable
  - across versions (to a degree)
- compressed
- can take a long time
  - both backup and restore

basebackup

- large file size
- not as portable
- faster for large databases
- can be used with PITR

Backups

- Good for:
  - natural disaster
  - admin error, bad update
  - software bugs
  - getting hacked
- Bad for everything else

Replication

- Streaming Replication
- Archive Replication
- Slony-I

Replication++

- Continuous
- Fast failover
- Low data loss

Replication--

- Extra hardware
- Complex
- High-maintenance
- Can hurt performance
- Can replicate failures

Replication

- Good For:
  - server, storage, network failure
- Bad For:
  - admin error, getting hacked
  - software bugs

Continuous Backup

- Also “PITR”
- Continuous like replication
- Partial recovery like backups
- Best of Replication & Backup
  - except slow restore times

THE REPLACEMENTS
I WILL DARE/COLOR ME IMPRESSED

… where you gonna restore those backups to?

Replacing Services

- servers
- network
- storage
- OS image
- software reversion

Procedures

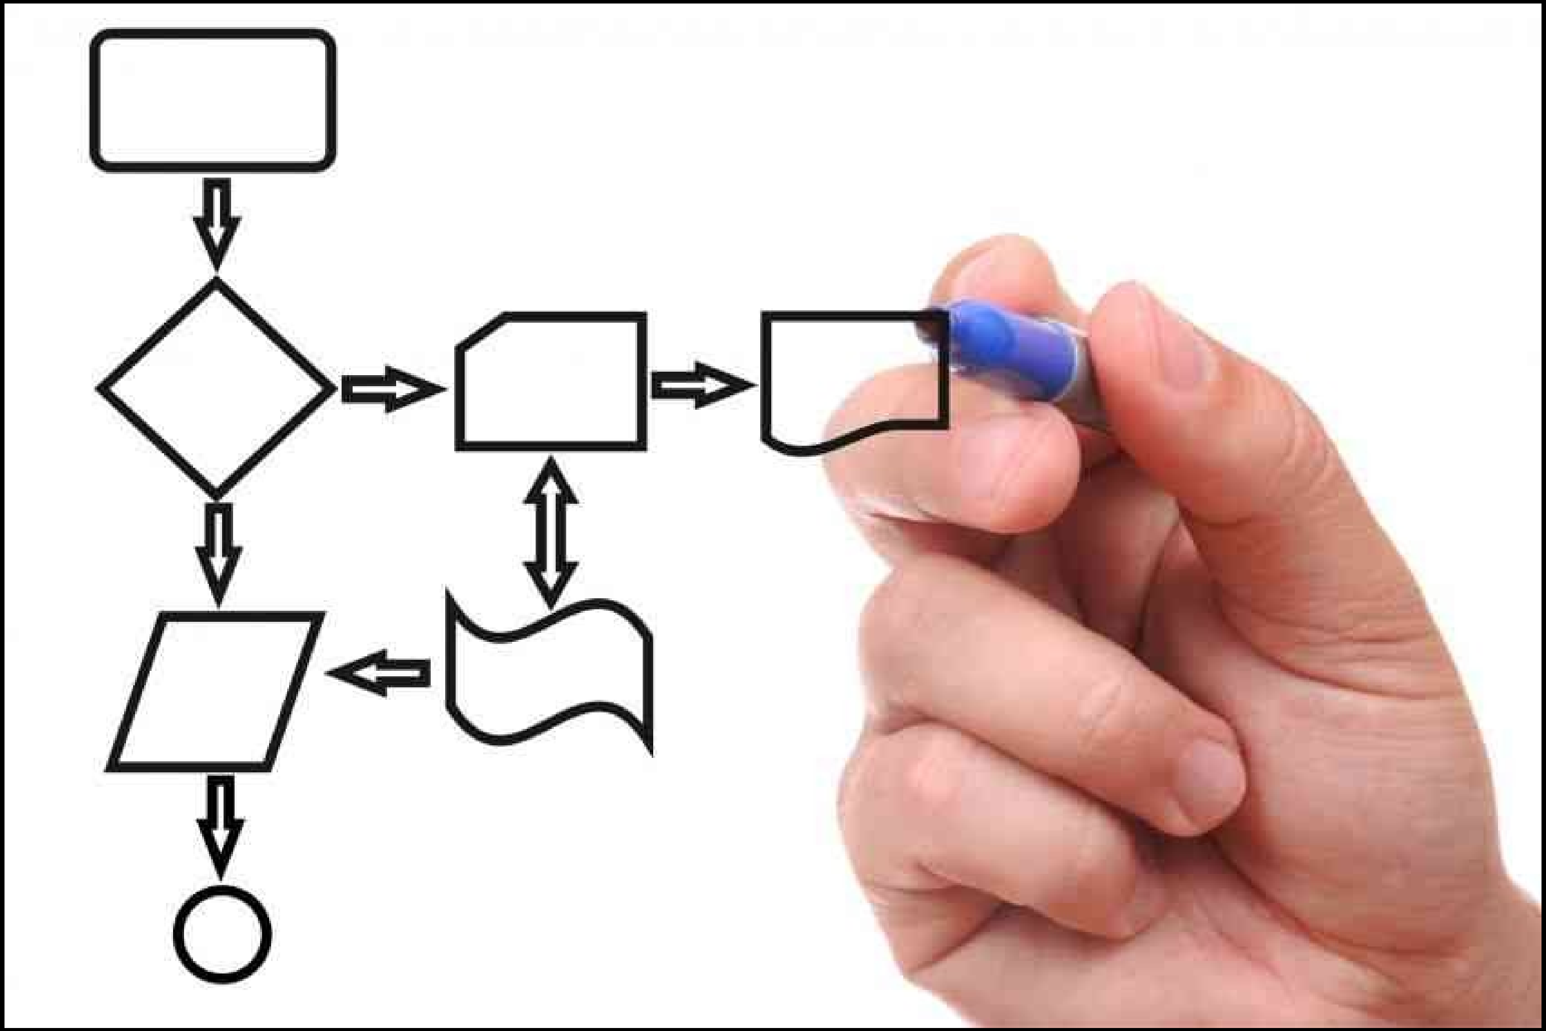

Written Procedures

3AM

is not the time

to improvise

Procedures

… for each recovery step

… for deciding what steps

Database Server Does Not Respond

1. Determine if physical server is down
   a. if network is down, use plan N1.
2. If not, try to restart database using command …
3. Still down? Fail over to replica using command …
4. Check replica.
5. Not working? Restore backup to test server 1 using command ...

Good: detailed written procedures

Better: written procedures with pastable commands

Best: tested single-command scripts

Fallback Procedures

- Sometimes recovery fails
- Have fallback procedures
- If the fallback fails

  … time for a meeting!

IMPROV
Where Comedy Begins

People

Who
You
Gonna
Call?

Know who to call

- on call staff
- experts in each service
- consultants/contractors
- vendors
- required authorizations

Contact Book

- Include as much contact information as possible
- Put copies in more than one place
  - including paper!
- Keep it up to date

Test Your DR

Good: when you create the procedure

Better: quarterly

Best: as part of daily/weekly provisioning

An untested backup
is one which
doesn't work.

DR in the Cloud

“It's a cloud, right?
That means it's
redundant, right?”

… not necessarily for your servers

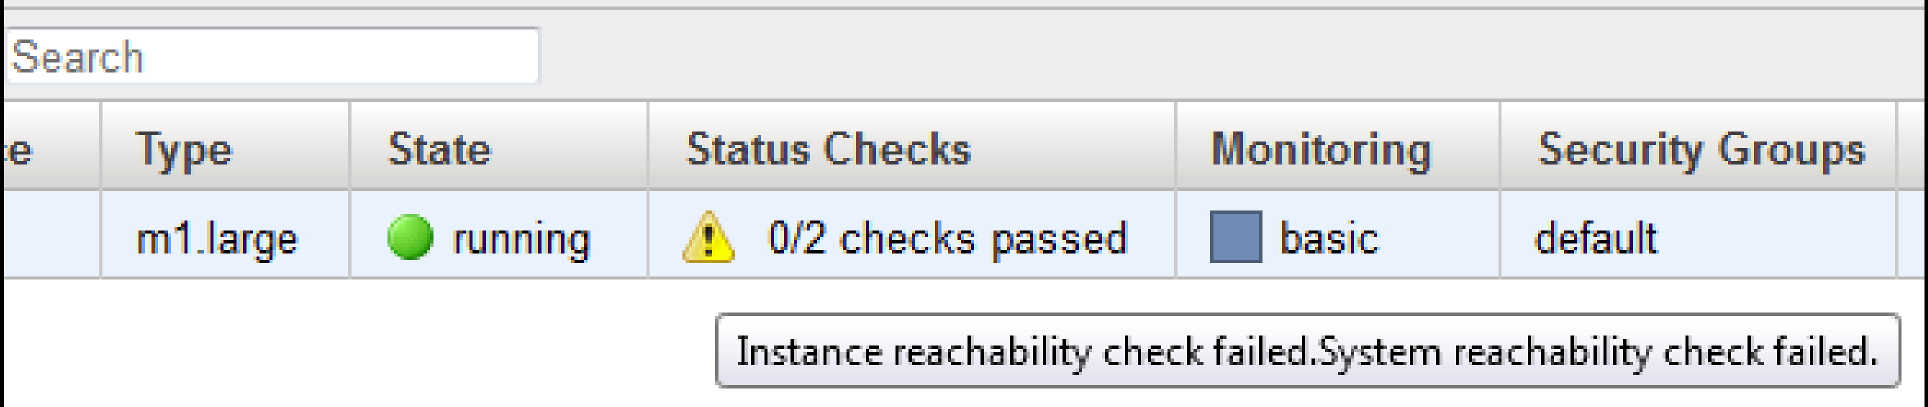

Search

e	Type	State	Status Checks	Monitoring	Security Groups
	m1.large	● running	⚠ 0/2 checks passed	■ basic	default

Instance reachability check failed.System reachability check failed.

unless you pay for it!

Some new problems

- Instance failure
- Resource overcommit
- Zone failures
- Admin error at scale

Some new solutions

- Redundant services
  - RDS, VIP, S3
- Rapid server deployment
- Cheap replicas

… otherwise pretty much the same.

backup locations

- shared instance storage (EBS)
  - fast failover for instance fail
- long-term storage API (S3)
  - redundant
  - large

Use your rapid deploy!

- Continuous backup to S3
- Deploy scripts + server images
  - Chef/Salt/Puppet/etc. helps here
- = fast recovery
  - with low running costs

DR Tips

- Have multiple copies of your plan
  - in multiple locations
- A SAN is not a DR solution
- One form of backup is seldom enough

Questions?

- Josh Berkus
  - www.databasesoup.com
  - www.pgexperts.com
- Coming up:
  - OSCON: AccidentalDBA
  - FOSS4G: Full day workshop

SOME RIGHTS RESERVED

Copyright 2013 PostgreSQL Experts Inc. Released under the Creative Commons Share-Alike 3.0 License. All images, logos and trademarks are the property of their respective owners and are used under principles of fair use unless otherwise noted.

PGX POSTGRESQL EXPERTS, INC.