

Disaster Recovery and Business Continuity Planning BoF

Internet2/ESnet Joint Techs, Minneapolis
5:45PM-7:45, February 13th, 2007

Radisson University Ballroom B

Joe St Sauver, Ph.D. (joe@uoregon.edu)

<http://www.uoregon.edu/~joe/dr-bcp-bof/>

Thanks For Joining Us Today for This BoF!

- Let's begin by going around the room, and have everyone briefly **introduce themselves**. Please give **your name**, and **the name of the institution you're with**.
- I'd also encourage you to **sign in on the sheet that's going around**.
- I'll then show a few introductory slides to get things started
- Finally we'll open things up for the rest of this session's time slot so that attendees can share what they're thinking about when it comes to Disaster Recovery and Business Continuity Planning.

DR/BCP Mailing List

- Before I forget, let me also mention that if you're interested in discussing disaster recovery and business continuity planning-related issues, there's a (relatively quiet) mailing list available for that purpose:

BC List Subscribe/Unsubscribe

<http://listserv.educause.edu/cgi-bin/wa.exe?SUBED1=bc&A=1>

A Little Background

- This BoF follows the well-attended Disaster Planning and Recovery BoF which took place at the Fall Internet2 Member Meeting in Chicago, and is meant to provide an opportunity for attendees to discuss contemporary approaches to disaster recovery and business continuity planning
- Kenneth Green's Campus Computing Project is an annual survey of pressing computing and network issues in higher education. A summary for the most recent year is available at <http://www.campuscomputing.net/summaries/2006/index.html> Checking that summary, the critical graph for this BoF is <http://www.campuscomputing.net/summaries/2006/4-disaster.html> which is captioned "Little Progress on IT Disaster Planning?"
Only a little over sixty percent of public universities are doing disaster planning.

Are There Threats That We Should Be Worrying About?

- **Hurricane Katrina** made business continuity requirements tangible for many folks for the first time.

Directnic's "Survival of New Orleans Weblog" provides a nice summary of what they saw

<http://interdictor.livejournal.com/2005/08/29/>

Ironically, UO directly saw some parts of the issues associated with that event since we provided secondary DNS for one New Orleans site that was hit. [lesson learned there? name server TTLs matter!]

- There are plenty of other examples, too.

Earthquakes

- “The Indian subcontinent was missing 70 percent of capacity," Barney observed. “China was missing at least 50 to 60 percent. Taiwan was almost 100 percent down.””

Quake-driven crash shows vulnerability of Asia's link to Net
<http://www.mercurynews.com/mld/mercurynews/news/16570731.htm>

January 29, 2007

- The United States has earthquake exposure as well:
 - everyone thinks about California and Alaska
 - but don't forget the New Madrid Fault Zone in Missouri
http://en.wikipedia.org/wiki/New_Madrid_Fault_Zone
 - and the Pacific Northwest's Cascadia Subduction Zone
http://www.ess.washington.edu/SEIS/PNSN/HAZARDS/CASCADIA/cascadia_zone.html

Facilities Fire

- Fire Devastates Dutch Internet Hub

[http://www.theregister.co.uk/2002/11/20/
fire_devastates_dutch_internet_hub/](http://www.theregister.co.uk/2002/11/20/fire_devastates_dutch_internet_hub/)

“A fire at the University of Twente in the Netherlands today has destroyed one of the fastest computer networks in Europe. The fire, the cause of which is currently unknown, has gutted a building housing the vast majority of the University's computer servers and networking equipment. While the fire is still burning, fire crews attending the scene have brought the conflagration under control. Staff at the University told us damage from the fire could cost the University €10 million or above. Although the university's network is (obviously) down, technicians at the University expressed optimism that the network could be restored in a matter of days rather than weeks.”

Widespread Loss of Power

- Major power outage hits New York, other large cities
<http://www.cnn.com/2003/US/08/14/power.outage/>

Power began to flicker on late Thursday evening, hours after a major power outage struck simultaneously across dozens of cities in the eastern United States and Canada. By 11 p.m. in New Jersey, power had been restored to all but 250,000 of the nearly 1 million customers who had been in the dark since just after 4 p.m. [...]

Power was being restored in Pennsylvania and Ohio, too. In New York City, however, Con Edison backed off previous predictions that power for most of the metropolitan area would be restored by 1 a.m. Friday. The power company had predicted that residents closer to Niagara Falls in upstate New York would have to wait until 8 a.m. [...] In just three minutes, starting at 4:10 p.m., 21 power plants shut down, according to Genscape, a company that monitors the output of power plants.

Loss of Facilities Access

- **February 12, 2007 4:46PM**

Quarantine lifted on building shuttered after anthrax attack

Health officials on Monday lifted a quarantine of a building once occupied by a tabloid newspaper but vacated after an anthrax attack killed a photo editor. [...] Bob Stevens, a photo editor for American Media Inc., died in **October 2001** after being exposed to anthrax in an envelope mailed to the building, which housed offices of the National Enquirer. Stevens' diagnosis brought to light widespread anthrax attacks that paralyzed the nation with bioterrorism fears shortly after the Sept. 11 attacks on New York and Washington. The publisher later moved from the building, and the case remains unsolved. **The cleanup began in July 2004**. The building was fumigated with chlorine dioxide.

<http://www.heraldtribune.com/apps/pbcs.dll/article?AID=/20070212/BREAKING/70212012&start=1>

Distributed Denial of Service Attacks

- Another possible “disaster” is a distributed denial of service attack. You might still have connectivity, but that connectivity will be overly full to the point of being unusable. See “Explaining Distributed Denial of Service Attacks to Campus Leaders,” <http://www.uoregon.edu/~joe/ddos-exec/ddos-exec.ppt> (or .pdf)
- If your backup site is connected via the same link that’s being DDoS’d, your backup site may lose synchronization with the primary site, and your backup site may not be accessible when needed.
- Should you have separate wide area connectivity for your backup site that is sheltered from “normal” network traffic, either as a separate packet connection, or as a circuit like connection?

And The List Goes On...

- Misadventures are a fact of life...
- The question is, “How do we **deal with** those risks? Are we at least **reasonably ready** to overcome foreseeable issues?”
- You carry a spare tire, you’ve got fire extinguishers, you buy insurance for your home and car, etc.
- For most sites, part of minimizing risk exposure in a business environment is having a disaster recovery and business continuity plan.

The Old Disaster Recovery Paradigm

- Reciprocal shared space at a partner site
- Data archived to tape
- Just-in-time delivery of replacement hardware
- Small number of key applications (typically enterprise ERP system)
- At least some down time is acceptable
- Proforma/low probability of occurring
- **Is that still a realistic paradigm? NO.**

What's Mission Critical?

- Domain name system?
- Enterprise SAN/NAS (data storage)?
- Enterprise Identity Management System?
- ERP System?
- Voice over IP?
- Teaching and Learning System?
- Institutional Web Presence?
- Email and Calendaring?
- Building control and access systems (smart build HVAC, elevators, door controls, alarm systems, etc.)
- The network itself?
- All of the above and more?

What Are Today's Restoration/Recovery Time Frames

- Hitless/non-interruptible?
- Restoration on the order of seconds?
- Minutes?
- Hours? <== I suspect this is what we need
- Days?
- Weeks? <== Is this where we are?
- Longer?

- **Assertion: time to recover is a key driver.**

Key Driver? Total Data Volume

- How many GB/TB/PB worth of data needs to be available post-event?
- If that data needed to be transferred over a network or restored from archival media post-event, **how long would it take to do that?**
- What about failing back over to a primary system once the crisis is over (including moving all the data that's been modified during the outage)

Key Driver? Data Change Rate

- If restoration has to occur from a checkpoint/periodically archived media, how much data would be at risk of loss since that snapshot?
- Are the transactions which occurred since that time securely journal'd, and can they be replayed if need be? Or would those transactions simply be lost?

Key Driver? Required Lower Level Infrastructure

- Secure space with rackage
- Power and cooling
- Local loop and wide area connectivity
- System and network hardware
- How long would it take to get/install/configure that lower level infrastructure from scratch, if it isn't already there?
- Office space for staff?

Key Driver? System Complexity

- Today's systems are complex.
- Replicating complex systems takes time and may require specialized expertise
- Specialized expertise may not be available during a crisis
- Detailed system documentation may not be available during a crisis.
- Debugging a specialized system may take time...
- Not going to want to try rebuilding everything on a crash basis.

Strawman Proposal/Suggestion

- **Doing disaster recovery/business continuity today requires a hot/spinning off site facility with synchronized data.**

Key Constraint? Cost

- Facilities themselves? (NOT cheap)
- Hardware? (commodity PCs are cheap, but enterprise-class SAN/NAS boxes are NOT)
- Software? (ERP licenses are NOT cheap!)
- Staff? (Personnel costs often dominate IT budgets -- what would staff impacts be?)
- Network connectivity? (Function of facility separation distance, bandwidth required, and redundancy demands)

Key Constraint: Distance (and Direction!)

- You need to be far enough away that a given disaster doesn't hit both your primary and backup sites (Katrina lesson: backup site for hurricanes should not also be coastal!).
- Some disasters can impact a surprisingly large region (e.g., power grid issues)
- Azimuth can be as important as distance traveled when attempting to get clear of a disaster zone. A backup site that's a hundred miles away (but right on the same fault line) is not a good backup site.
- But distance comes at a cost:
 - direct network connectivity may be milage sensitive
 - latency and bandwidth delay problems may complicate use of truly remote backup sites
 - staff may need to travel to the remote site; if a backup site is very remote, on-site staff (or remote hands) may be needed

Can Advanced Networks Help Universities Get Better Geographic Separation?

- Current default/de-facto financial limit for remote sites is often the physical extent of a statewide or regional network.
- Unfortunately, disasters are also often statewide or regional in scope -- you need greater separation!
- Keeping large volumes of data synchronized between two sites requires high throughput and minimum latency on a point-to-point basis -- both characteristics of today's advanced networks.
- Inter-filer synchronization traffic may not be encrypted -- does that help make this a perfect application for static lambda-based connections?
- Are there other potential roles for advanced networks when it comes to helping sites become more resilient and robust?

Or Should We Be Moving Away From The Hot Site Model Entirely?

- Should we achieve resilience via outsourcing/managed hosting?
(for example, a growing number of sites outsource their teaching and learning systems)
- Should we be trying to virtualize everything that's important in an inherently distributed way?
 - We've got federated authentication systems such as Shibboleth...
<http://shibboleth.internet2.edu/>
 - We've got common replicated and survivable file systems
(particularly impressed by the work at UTK on LoCI, see
<http://loci.cs.utk.edu/> or commercial services such as the Amazon Simple Storage Service, S3, <http://aws.amazon.com/s3>)
- Or should we be thinking about something else entirely?

Testing

- Whatever you do end up doing, be SURE to test it...
- When it comes to disaster recovery and business continuity planning, the key to making this real is going to be testing the plan. Give yourself an **intentional** disaster!
- Hypothesis: many sites do not and will not test (and probably for dang good reasons)
- That's probably a sign we have a lot of work to do!
- Where would folks like to go from here? What's on your mind?