

Medical Center's Multiday Outage

February 2011

Recently, a major academic medical center lost almost all of its IT services for over three days, threatening the welfare of its patients. Most disturbing is the unnecessary sequence of events that led to the failure.

Healthcare systems dealing with keeping patients alive and with treating their illnesses are an important example of safety-critical systems. We would like these systems to have telephone system-like reliability – six 9s and beyond. But too often, we hear about major failures that can potentially cost lives or inflict serious injury on hospital patients.

A previous example is the failure of the U.S. Veterans Administration's new medical records system.¹ Wrong medical records for patients were delivered to doctors over a four-month period before the errors were uncovered.

In the case described below, it was not just medical records that were unavailable. Virtually all hospital services, including clinical analyses, were unavailable. This incident was reported in the *ACM Committee on Computers and Public Policy*.² The medical center was not identified.

Medical Center Down

Shortly before midnight on a recent Monday evening, a major system crash disabled virtually all of the IT functionality for the entire campus of a major academic institution. Included were the IT services for the institution's medical center. Financial records, patient medical records, clinical laboratory test ordering and reporting, imaging ordering and reporting, and pharmacy systems were all down. The systems were not restored to service until the following Friday, and it took even longer before the backlog of information that had accumulated during the outage could be entered into the systems.

The outage forced the cancellation of all elective procedures. Affected were 52 major procedures and numerous minor procedures such as colonoscopies.

All ambulance traffic was diverted to other hospitals – an estimated seventy diversions. There were substantial delays in ordering and obtaining laboratory and radiology reports for inpatients. Even with the ambulance diversions, it was difficult to clear out the hospital since doctors could not discharge patients until they could receive test reports.

¹ More Never Again II: Electronic Medical Records – Is the Cure Worse Than the Disease?, *Availability Digest*, February 2009.

² Robert L. Wears, MD, Health Information Technology Risks, *The Risks Digest (ACM Committee on Computers and Public Policy)*; March 20, 2010.

Fortunately, no patients were harmed, due in great part to ad hoc changes in procedures dreamed up on the spot by the medical-center staff.

What Caused This Outage?

The outage was caused by a chain of small events. It was aggravated by some underlying problems that had festered unknown for two years. The absence of any of the events or conditions in the failure chain could have prevented the outage.

The outage began with a hardware failure in a critical network component. The faulty component was quickly identified and repaired. However, the nature of the network fault was such that the data center's major servers had to be rebooted. No serious problem so far.

However, to the horror of the system operators, repeated attempts to reboot the servers failed with the servers all reporting critical errors. It took a while for the operators to determine that for some unknown reason, certain file permissions had been changed; and these changes prevented the systems from rebooting. The changes also prevented the operators from recovering the systems by reverting to previous versions of the systems.

It took the better part of two days to untangle the underlying causes of the outage and to correct them. The clinical systems were finally returned to service on Friday afternoon.

A Long Chain of Events

The failure chain started two years prior to the failure when the institution decided to upgrade its system with a "high availability" failover capability. The project was plagued with problems from the start and was eventually abandoned. However, unbeknown to the system managers, some of the changes that had been made to the system, including certain file permissions, never got rolled back.

In the two years that intervened, there was never a need to reboot the systems. It was only after the network failure two years later that what should have been a standard reboot was attempted. However, the latent changes that had been inadvertently left behind caused an emergency failover to be initiated; but the failover facility was nonexistent. This led to repeated reboot failures.

Once the problem was understood, the changes were successfully backed out. The databases underwent extensive integrity checks, and the clinical servers were then restarted. At this point, the clinical and financial transactions that had accumulated during the outage had to be entered into the system. This activity turned out to take more time than the outage itself, but no pre-existing data was lost. All services were finally restored on Friday afternoon, three and a half days after the initial outage.

The hospital staff was creative in responding to the problems created by the outage and for ensuring that critical services could still be performed, at least in some fashion. For instance, the financial accounting staff, now idle because their systems were down, were pressed into service as runners delivering materials, clinical orders and reports.

The Postmortem

Beyond the fact that system management was careless in backing out changes made during a failed project, a postmortem of the outage revealed several issues that contributed to the length of the outage:

- It was difficult for the frontline workers to convince help-desk personnel that there was a problem.
- The nature of the underlying fault – old changes that had not been backed out – was difficult to uncover.
- The medical center was slow in activating its internal disaster plan. The incident management group did not declare a disaster until sixteen hours into the outage.
- The system restart procedures had not been tested for two years. Such testing could have identified the problems by rebooting only one server rather than the entire data center and at a time when the problems might have been more easily recognized and corrected.

A previous risk analysis had estimated that a data-center outage would cost the medical center \$56,000 per hour. The total outage, therefore, represented a cost of close to \$4 million, not including lost revenue due to cancellations and diverted patients.

Lessons Learned

Sadly for this medical center, an attempt at high availability led to unavailability. Several lessons from this incident are obvious in retrospect:

- Why did a major project that was going to make significant modifications to existing systems impact the online systems? The project should have been implemented and tested on a development system and only migrated online when it was clear that it was operational. Perhaps management was trying to save the costs of a development system. That attempt at savings cost the medical center an estimated \$4 million.
- Change management is critical in achieving high availability. No change should be made to a system without both documentation and a plan to back out the change if it causes problems or is not needed.
- Failover testing (in this case, simple rebooting) must be carried out periodically to make sure that the procedures are correct and that they work. Testing them only when they are needed is a strategy fraught with danger.

The ACM report concludes with the following observation. “As more and more care delivery organizations with little experience in managing clinical, as opposed to business, systems install more and more advanced clinical HIT [health information technology] systems - systems that have not been developed from a safety-critical computing viewpoint, more frequent and potentially more consequential failures are likely.” This is a concern that we should take seriously.

Acknowledgement

Our thanks to our subscriber, Ron LaPedis, for bringing this incident to our attention.