ITRoadmap
CONFERENCE & EXPO

LATIN
SCHOOL of CHICAGO

Lessons Learned about Consolidation, Virtualization and Disaster Recovery

Shandor Simon
Director, Networking Services
Latin School of Chicago

Agenda

- The Latin School of Chicago
- Why redundancy and resiliency is important to us
- How it was supposed to work
- What went wrong?
- How did we respond? What have we learned?
- Results
- Q&A

Latin School of Chicago

- Founded in 1888, The Latin School of Chicago is a private independent day school located on the Near North Side of Chicago
- 1,100 students in JK through 12th Grade
- 3 Buildings, 300,000 sq ft
- 250+ Faculty/Staff
- 80+ servers, 2 Data Centers
- 700+ school-owned computers

ITRoadmap CONFERENCE & EXPO

LATIN SCHOOL of CHICAGO

Latin School of Chicago – Technology Environment

- Academic Technology (projectors & interactive whiteboards in every classroom, class websites, blogs, wikis, etc.)
- Administrative Technology (accounting, payroll, development, attendance, grades, etc.)
- Operational Technology (building controls, access control, surveillance, etc.)
- Technology Infrastructure (DNS, DHCP, Active Directory, etc.)

IT Roadmap CONFERENCE & EXPO

LATIN SCHOOL of CHICAGO

The Challenge

- Our school has become dependent on our technology infrastructure functioning. In our new building, the network must work for the classroom doors to unlock.
- We have invested in redundant server, storage and network virtualization, multiple server rooms.
- Yet we still had a failure.
- What happened, and more importantly what can we learn?

Virtual Server Infrastructure

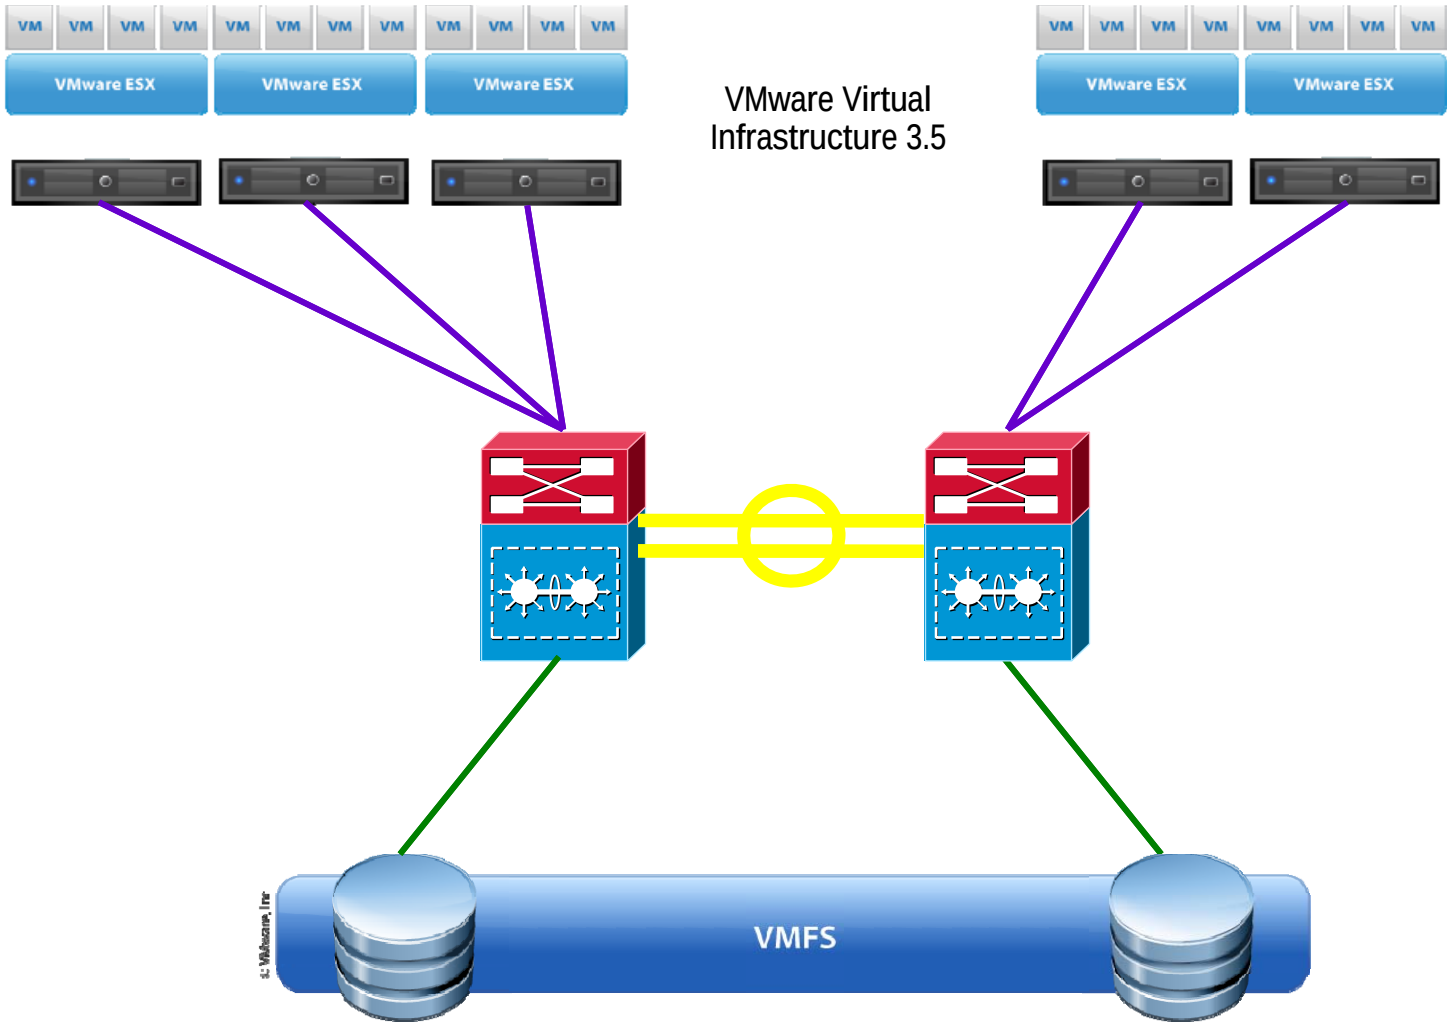

How VM failure is supposed to work

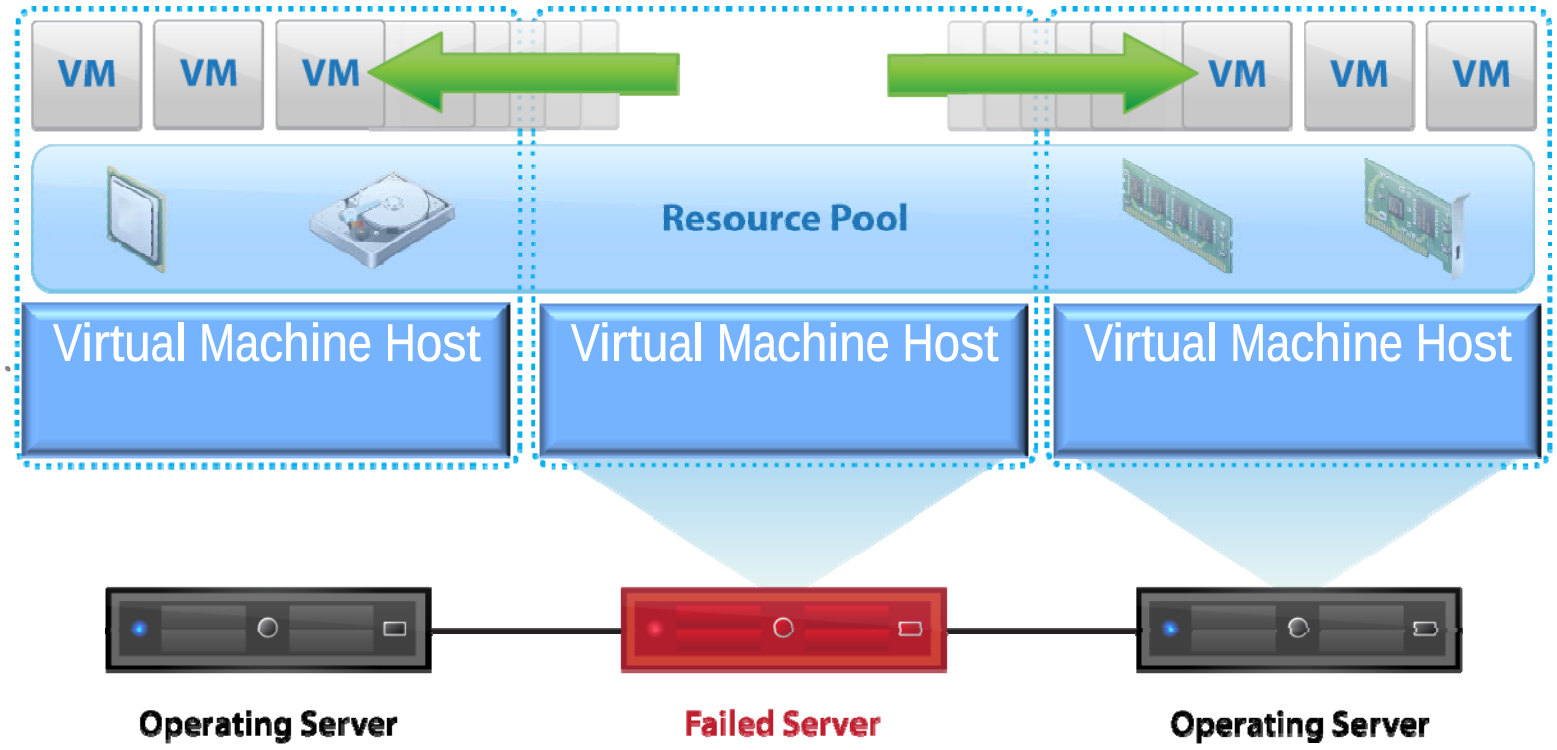

How our iSCSI SAN is supposed to work

- Our SAN is a software based iSCSI solution that runs on commodity hardware. We have four servers, each with 6TB of RAID 5 storage.
- The SAN writes the same data to at least two different servers, and spreads the data across all four servers.
- The SAN is configured to write one copy of the data to an “even” server, the other to an “odd” server.

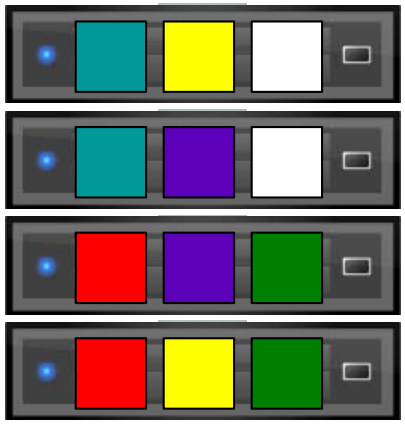

How our iSCSI SAN is supposed to work

- By placing all the even servers in one building, and the odd servers in another, we will not lose data with the loss of one location.
- If the link between the sites fail, it is important to allow changes to happen only on one side of the SAN.
- Our vendor makes use of a quorum where one side has the most “votes.” We have a special tie-breaking server.

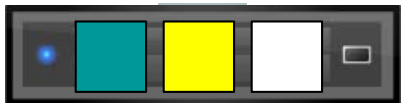
1

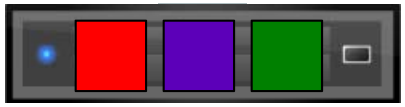
3

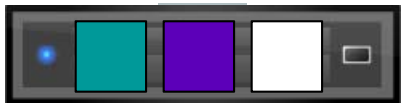
2

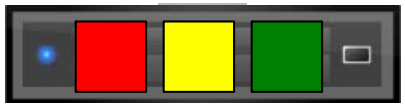
4

Network Redundancy

- Our core virtual switch / router physically spans two locations. Two 10 Gb links connect the two physical boxes.
- The VSS switch appears as a single logical switch to the rest of the network.
- All of our access switches are wired back to both locations.
- Either physical box can fail without impacting the core network functions.
- We have redundant Internet connections to both locations, and redundant firewalls

Facilities

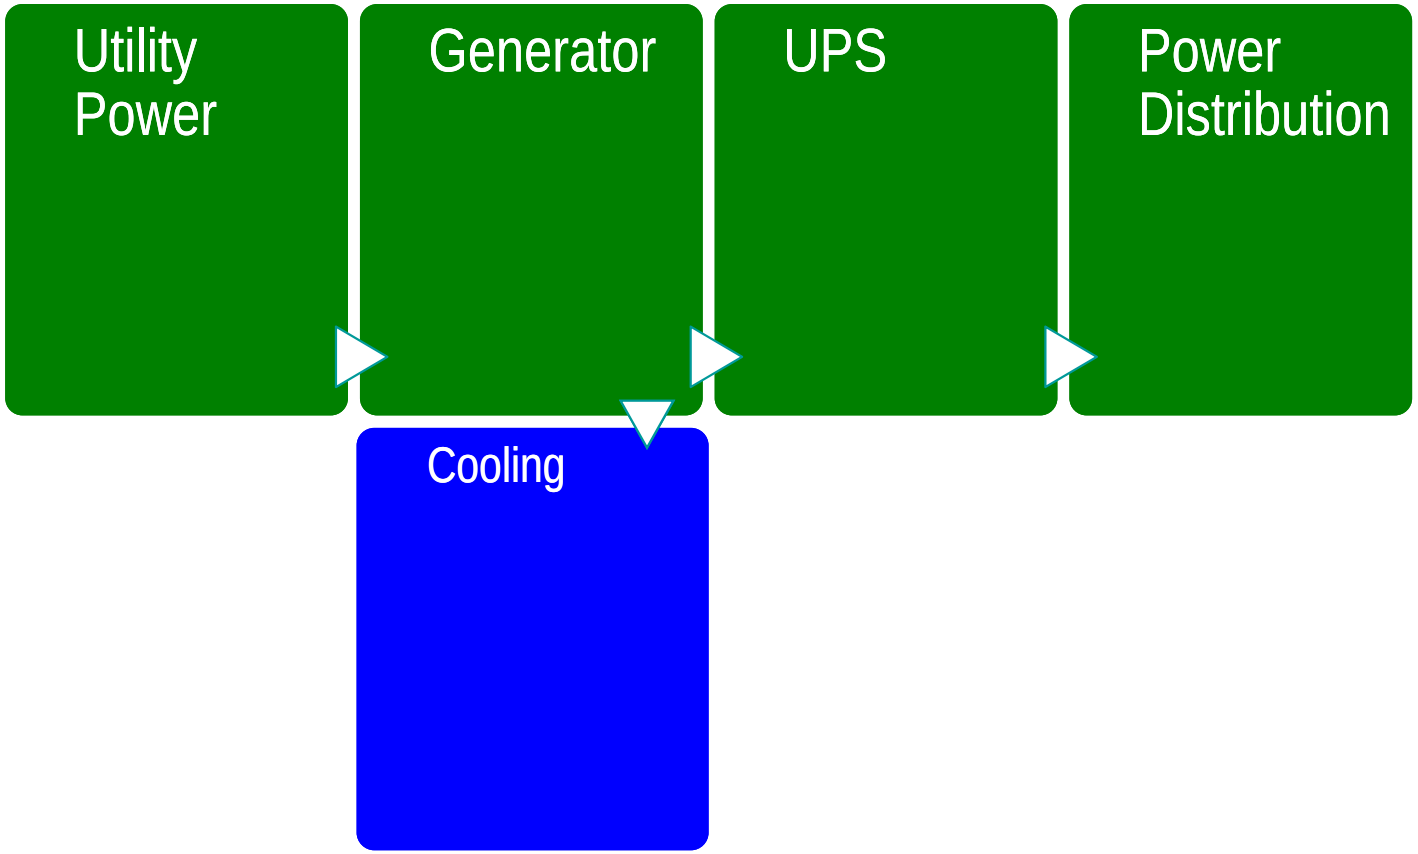

Facilities

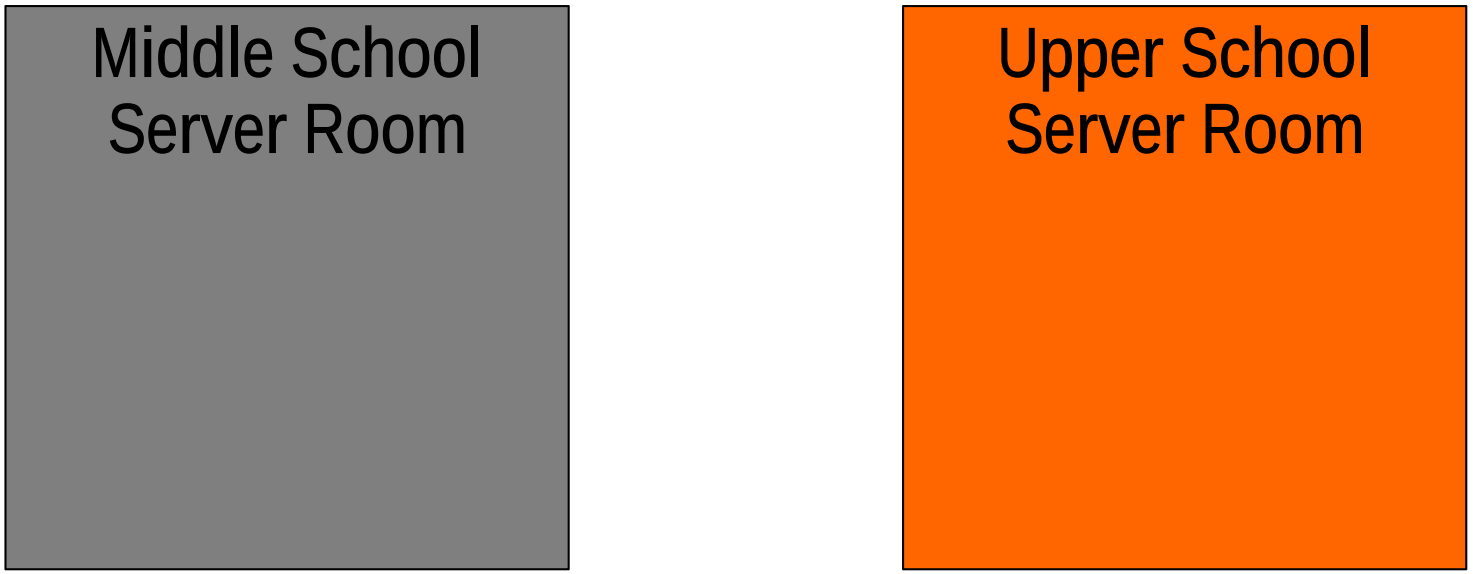

Putting it together

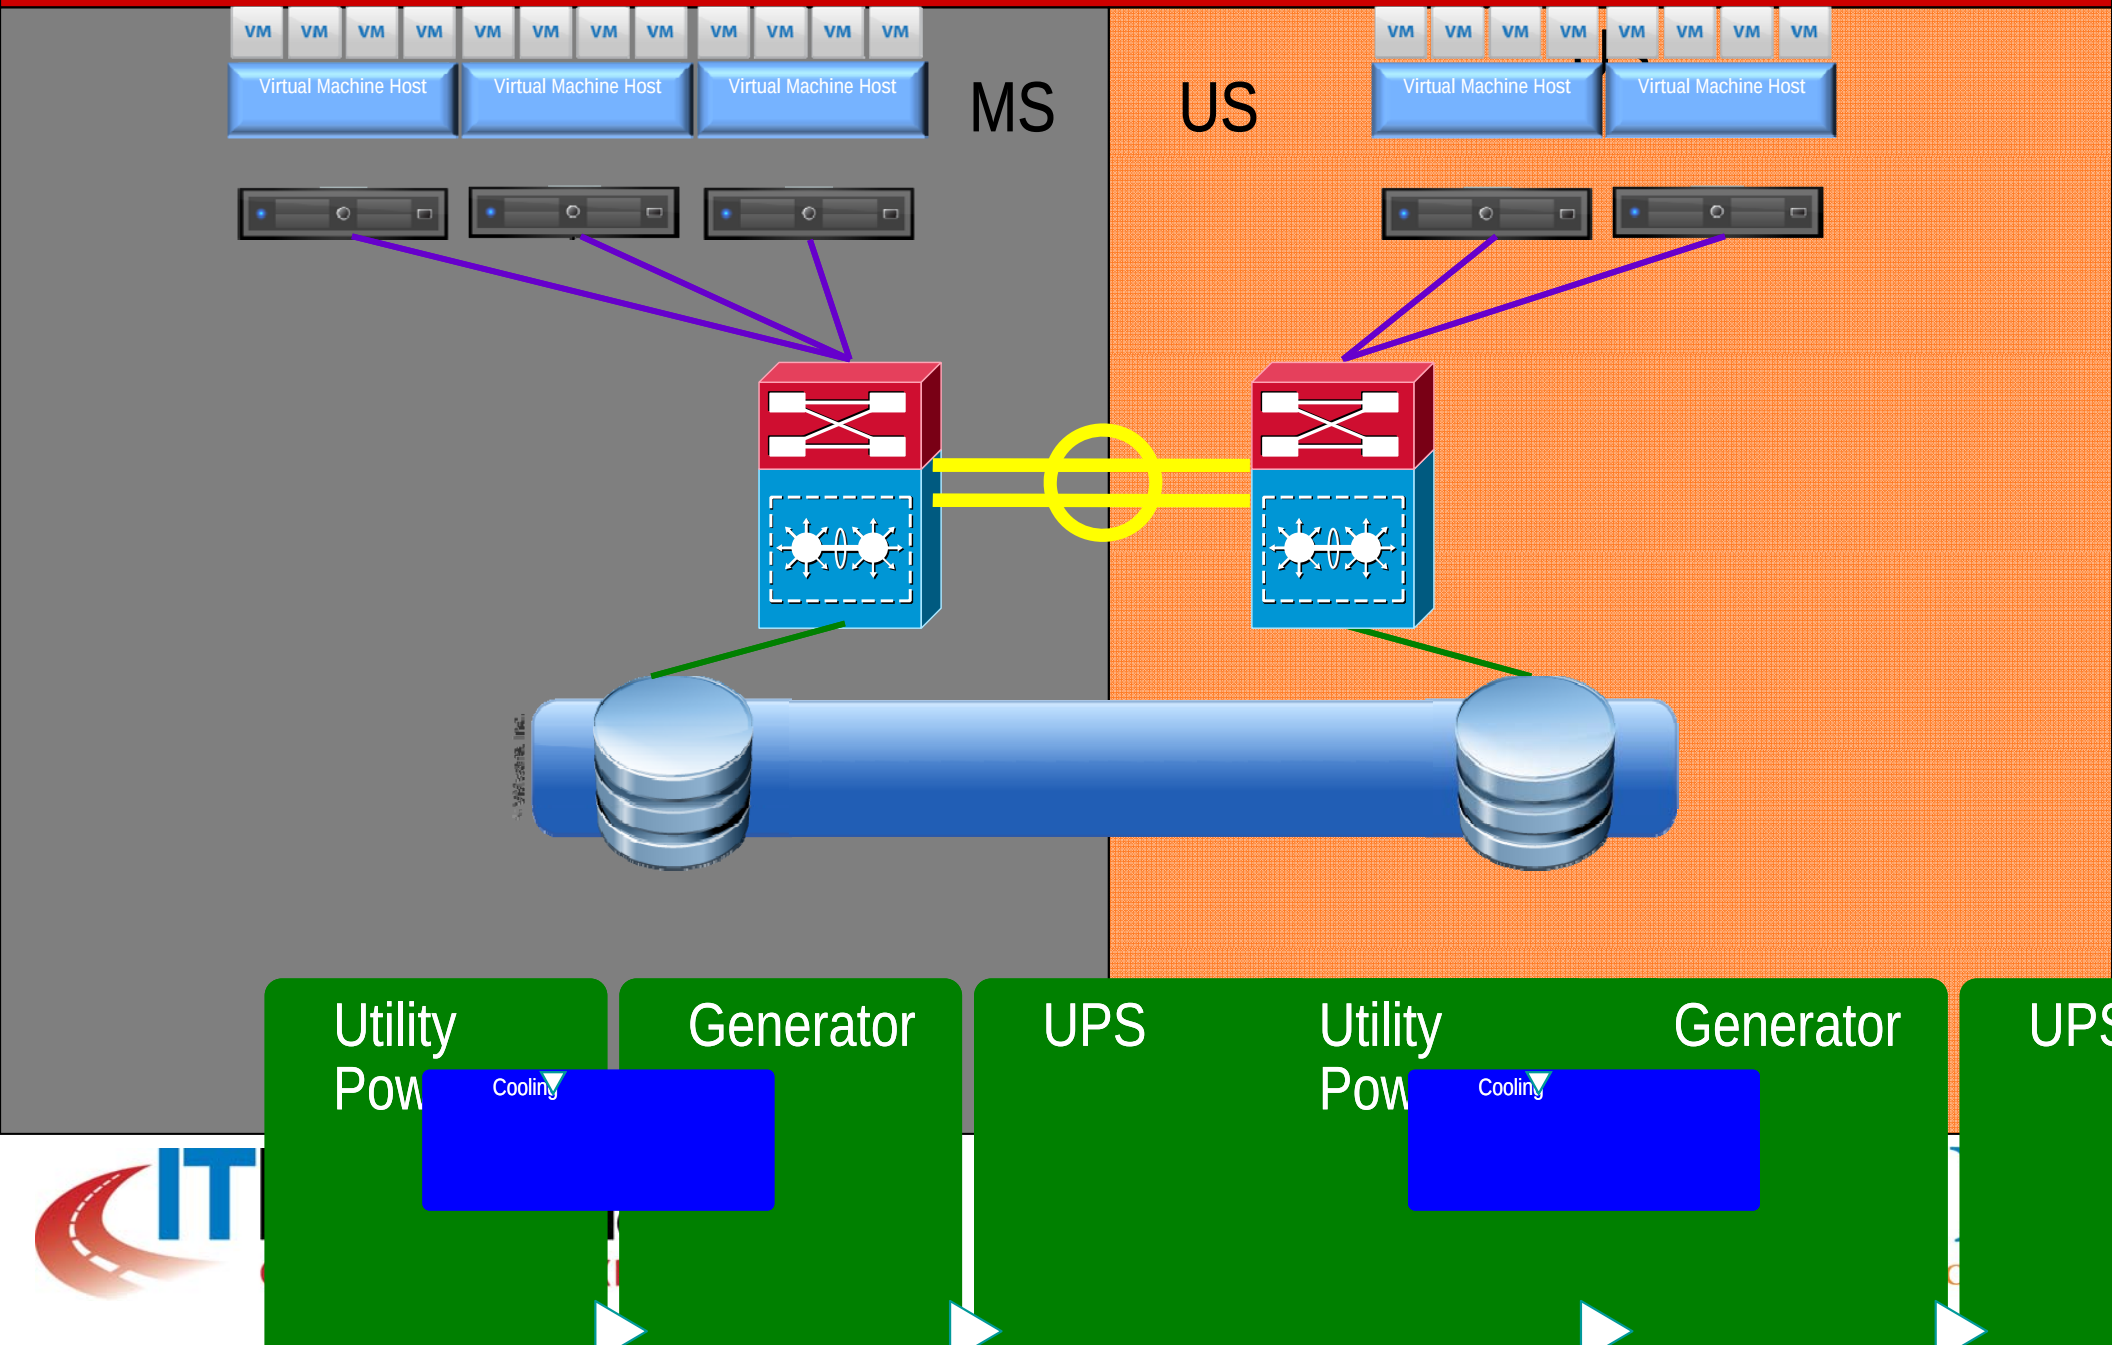

What we expected…

- We could survive failure of either data center.
- Servers in a failed location would restart in the other location.
- That the overall reliability of the system was sufficiently high.
- That our setup could handle failure of an entire data center location, or of a component on one side.

Expectations

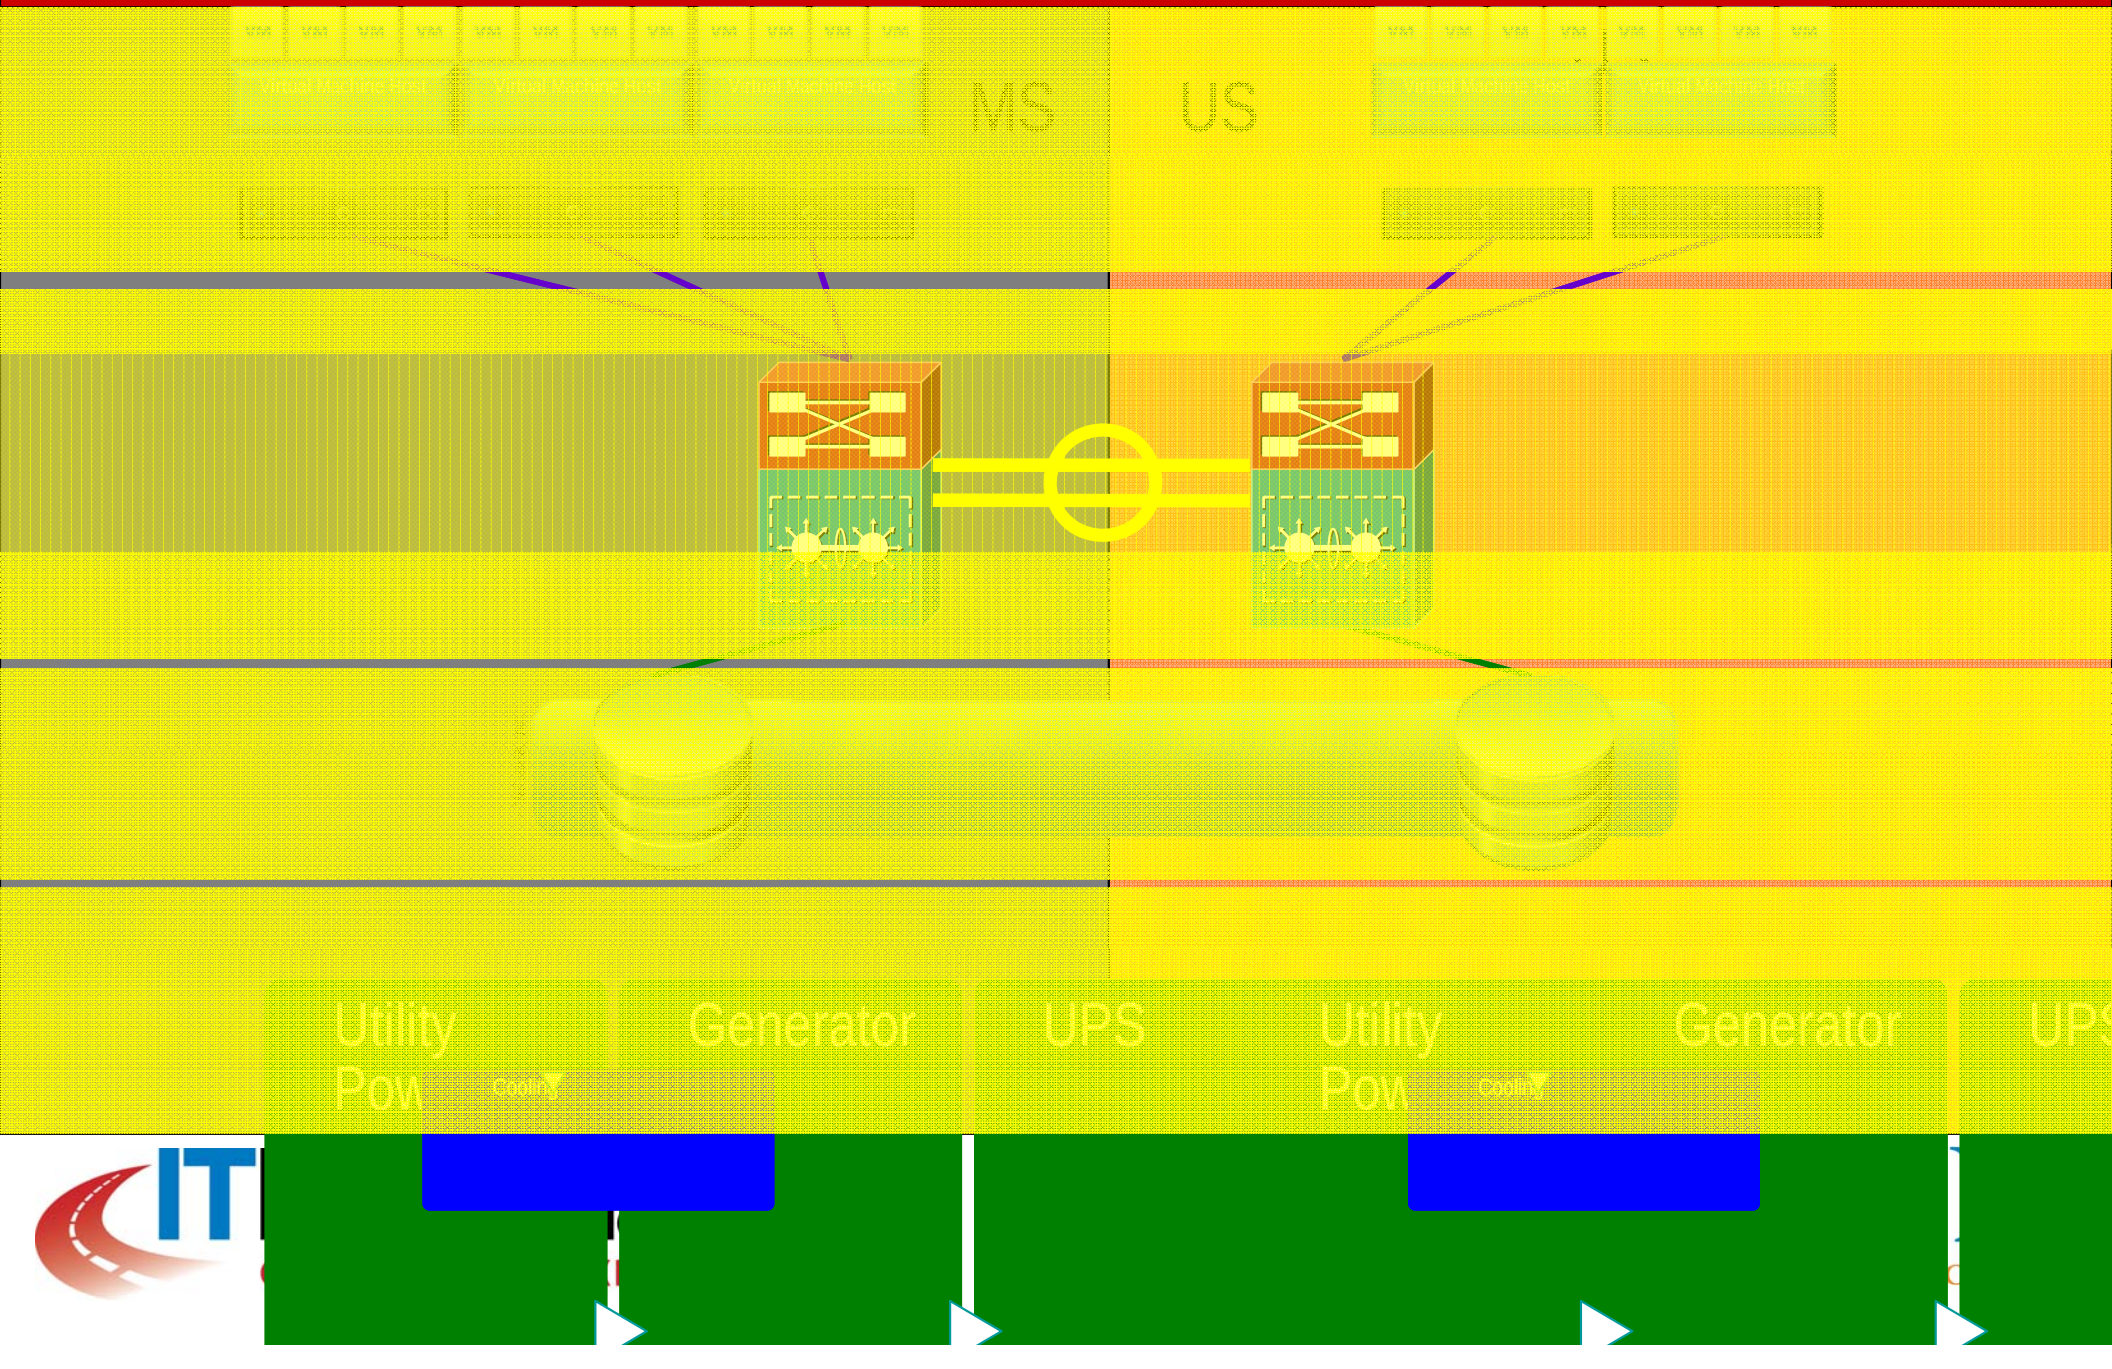

Actual Failure

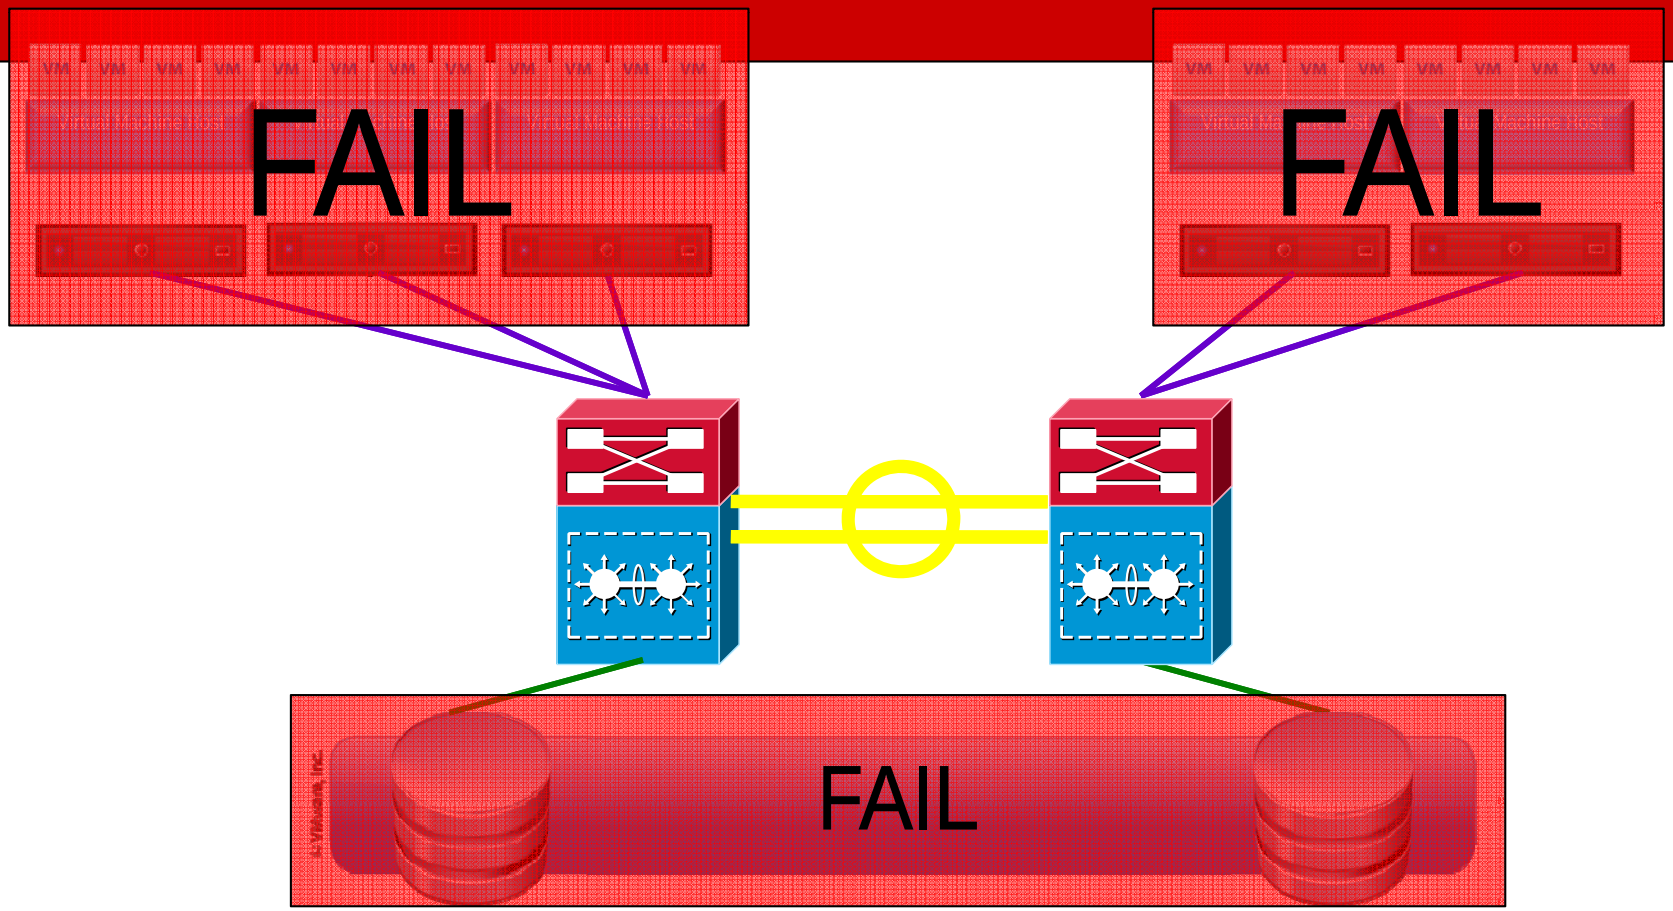

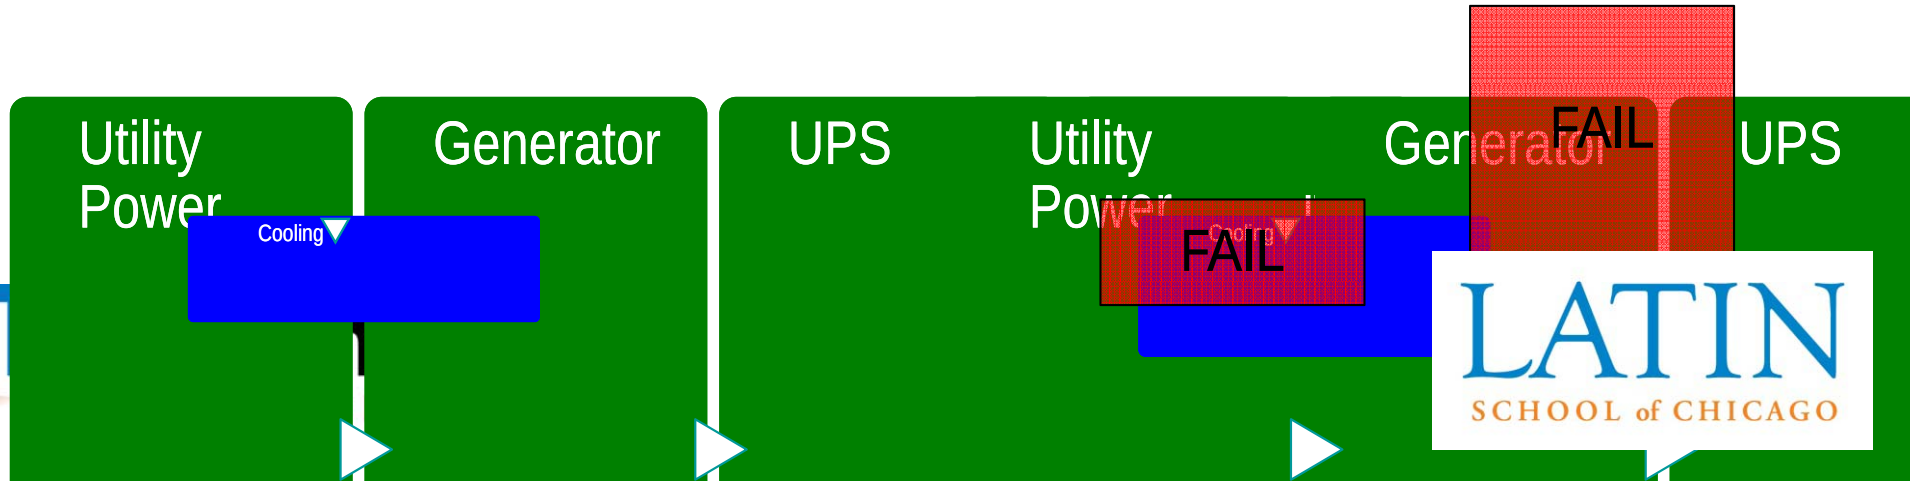

LATIN SCHOOL of CHICAGO

What went wrong?

- We didn't properly account for load on our PDUs under full power, nor account for a cascading failure caused by servers with multiple power supplies shifting load.
- Since our tie-breaking server was improperly placed, our SAN did not achieve a quorum at the site that survived, and went read-only.
- Functioning virtual servers crashed at the good site when their disks went read-only, and failed servers were unable to be restarted.
- Active Directory, Doman Name Service, and DHCP all went down. AD and DNS should have survived, but the backup server that hosted the VMs that did not use the SAN was down for maintenance.

Our response

- Switched to DNS/DHCP appliances that can failover.
- Moved database and mail servers off of SAN (performance win), and replicate the data between data centers twice a day.
- Added additional PDUs and spread the load across them.
- Enhanced monitoring of data center facilities to a service that notifies via phone and uses a call list.
- Moved SAN tie-breaking server outside of either data center.
- During scheduled downtime, we test that our failover works the way we expect it to.

Recovery Time vs. Recovery Point

- Recovery Time: how long it takes to recover from a failure.
- Recovery Point: how much data is lost.
- In our initial design we focused on preventing failure and data loss.
- We did not plan on recovery time, and it showed (took way too long to get back on-line).
- Our users would prefer to lose a few hours of mail than a few hours of work.
- We now assume everything will fail, and run recovery exercises to see how long it takes us to recover.

Results

- Our air conditioner failed again, but this time did not cause additional failures. With better notification, we remotely forced our infrastructure to failover to the other site, and shut down our equipment, dealing with the problem in the morning.
- In another case a PDU failed, causing half of our SAN to go down. The SAN continued to function out of the other site.
- Recently we were able to recover from a server failure by reverting to a replicated copy, causing only a 15-minute outage and minimal data loss.

Avoiding our mistakes

- Make sure you understand the limitations of your facilities infrastructure (cooling, power, etc.)
- Test how your infrastructure reacts to failure, considering how different systems interact.
- Understand that recovery time might be as or more important then recovery point.
- Watch out for redundant parts of your infrastructure having the same dependences (ex: our “redundant” servers all using the same SAN).

ITRoadmap
CONFERENCE & EXPO

Shandor Simon
Director, Networking Services
Latin School of Chicago