ECE590-03
Enterprise Storage Architecture

Fall 2016

Business Continuity: High Availability

Tyler Bletsch

Duke University

Includes material adapted from the course "Information Storage and Management v2" (modules 9-12), published by EMC corporation.

What is Business Continuity?

Business Continuity

It is a process that prepares for, responds to, and recovers from a system outage that can adversely affects business operations.

- An integrated and enterprise-wide process that includes set of activities to ensure "information availability"
- BC involves proactive measures and reactive countermeasures
- In a virtualized environment, BC solutions need to protect both physical and virtualized resources

Information Availability

Information Availability

It is the ability of an IT infrastructure to function according to business expectations, during its specified time of operation.

- Information availability can be defined with the help of:
  - **Accessibility**
    - Information should be accessible to the right user when required
  - **Reliability**
    - Information should be reliable and correct in all aspects
  - **Timeliness**
    - Defines the time window during which information must be accessible

Causes of Information Unavailability
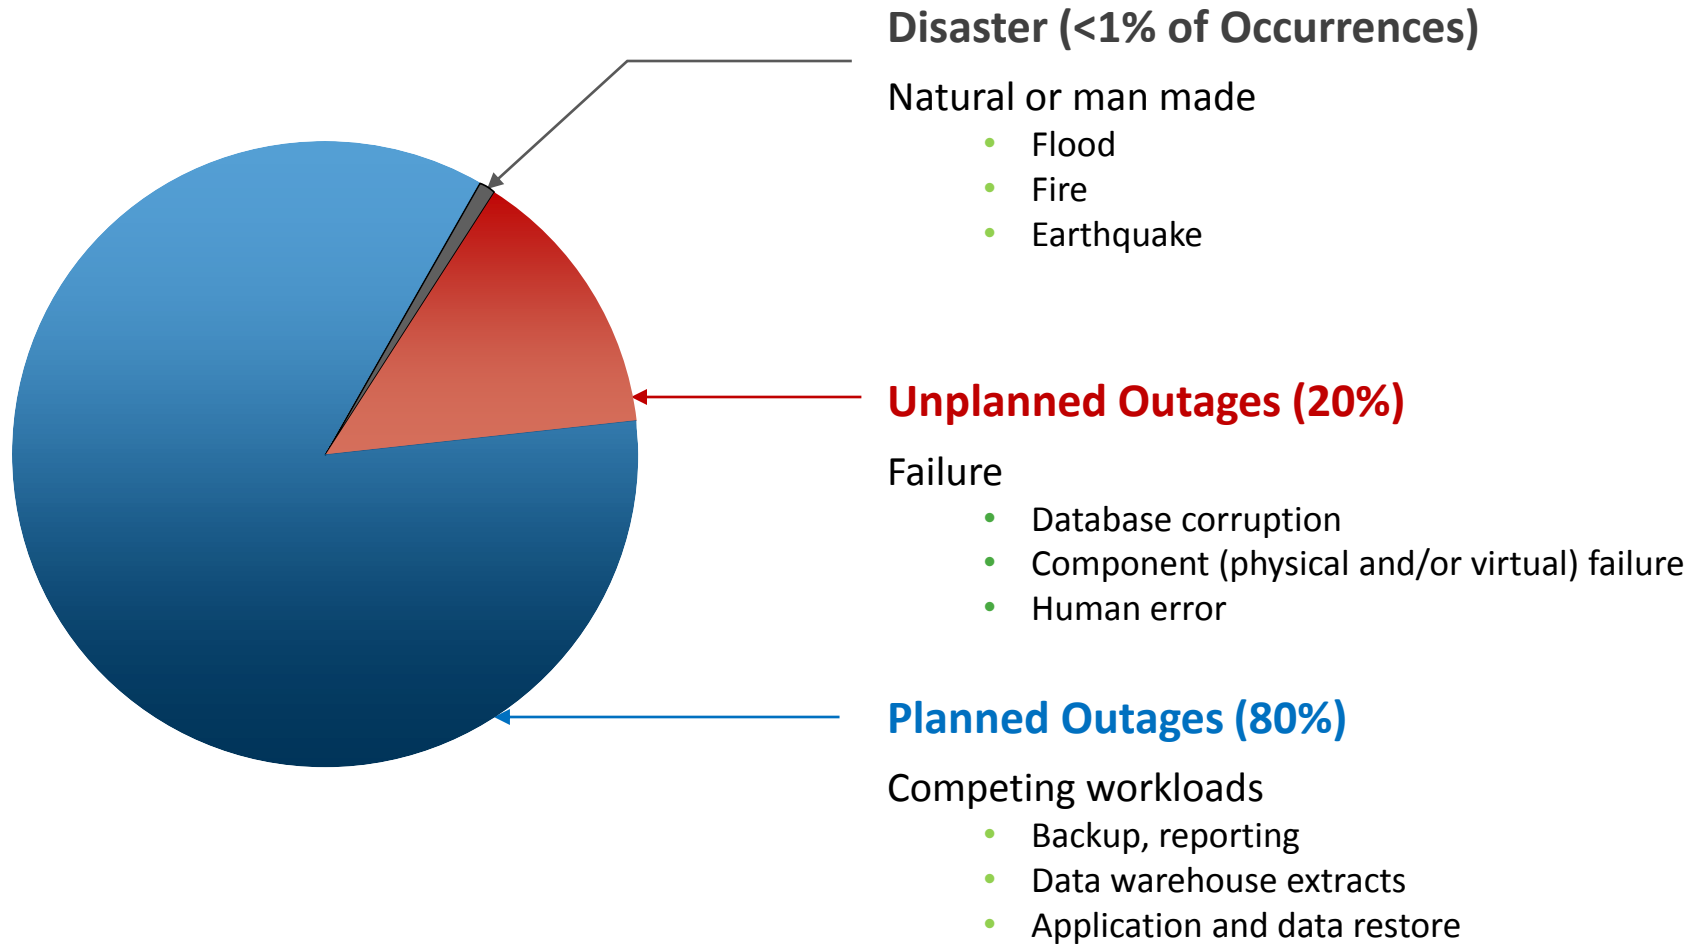
Disaster (<1% of Occurrences)

Natural or man made
- Flood
- Fire
- Earthquake

Unplanned Outages (20%)

Failure
- Database corruption
- Component (physical and/or virtual) failure
- Human error

Planned Outages (80%)

Competing workloads
- Backup, reporting
- Data warehouse extracts
- Application and data restore

Impact of Downtime

Lost Productivity

- Number of employees impacted x hours out x hourly rate

Know the downtime costs (per hour, day, two days, and so on.)

Lost Revenue

- Direct loss
- Compensatory payments
- Lost future revenue
- Billing losses
- Investment losses

Damaged Reputation

- Customers
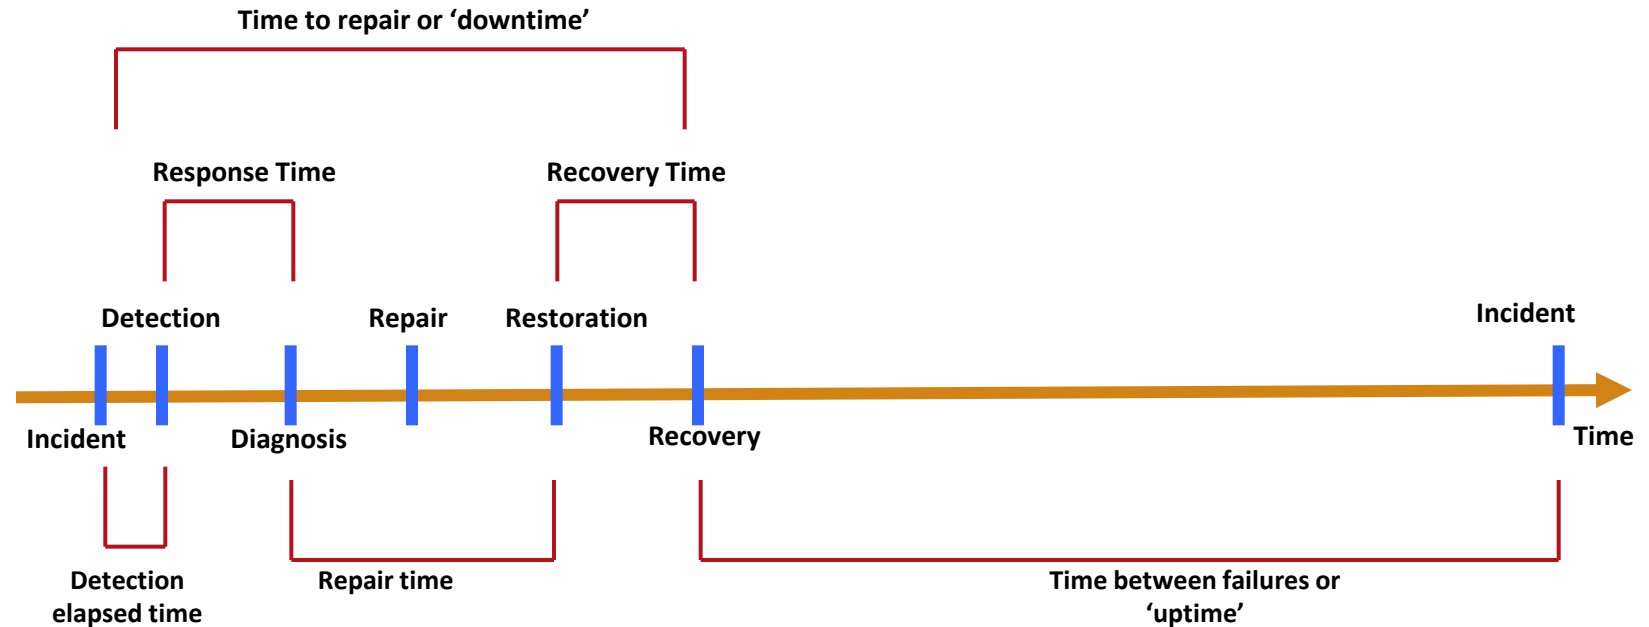- Suppliers
- Financial markets
- Banks
- Business partners

Financial Performance

- Revenue recognition
- Cash flow
- Lost discounts (A/P)
- Payment guarantees
- Credit rating
- Stock price

Other Expenses

- Temporary employees, equipment rental, overtime costs, extra shipping costs, travel expenses, and so on.

Measuring Information Availability

- MTBF: Average time available for a system or component to perform its normal operations between failures

  MTBF = Total uptime/Number of failures

- MTTR: Average time required to repair a failed component

  MTTR = Total downtime/Number of failures

  IA = MTBF/(MTBF + MTTR) or IA = uptime/(uptime + downtime)

Availability Measurement – Levels of '9s' Availability

Uptime (%)	Downtime (%)	Downtime per Year	Downtime per Week
98	2	7.3 days	3hrs, 22 minutes
99	1	3.65 days	1 hr, 41 minutes
99.8	0.2	17 hrs, 31 minutes	20 minutes, 10 secs
99.9	0.1	8 hrs, 45 minutes	10 minutes, 5 secs
99.99	0.01	52.5 minutes	1 minute
99.999	0.001	5.25 minutes	6 secs
99.9999	0.0001	31.5 secs	0.6 secs

Vendors love to brag about their nines

99,9 % UPTIME GUARANTEED

UPTIME 99.9% GUARANTEE

VPLS SOLUTIONS

99.999 PERCENT AVAILABILITY — Measured Across Entire Nimble Installed Base (July 2012 to November 2013)

Enterprise Hosting You Can Rely On

...ity Technology provides ...ality enterprise hosting ...lutions. We offer and stand ...our "five nines" up-time ...uarantees for corporations ...with high scalability and ...availability requirements.

...Our premium product offering ...ged hosting

UPTIME GUARANTEED 99.9% UPTIME GUARANTEED

SERIOUS HOSTING RELIABILITY · FIVE NINES 5X9 99.999% · SERIOUS SERVICE LEVEL AGREEMENTS

Comprehensive Technology Support Solutions

VPLS Solutions offers a full range of support solutions customized to fit your business needs. Whether you are in need of onsite technical support or remote network monitoring, our team of experienced technologists are here to help. VPLS Solutions offers a full range of network monitoring solutions from our three global VPLS Technical Assistance Centers (vTACs) located in Los ...es and Orange, CA and Bangkok, Thailand. VPLS Solutions will ...our network remotely 24×7 so that you can minimize staffing ...d the clock.

...99% uptime.

...out a 24×7 staff

UPTIME 99.9% GUARANTEE

99.999% uptime vs 99.9% uptime: the difference two extra "nines" makes

BY MITCHELL BERENSON - JANUARY 28, 2014
POSTED IN: ADMINS, HOSTED EXCHANGE, INFRASTRUCTURE, ...PROVIDERS

99 .999% UPTIME

99.9% Guaranteed Uptime

99.9% Uptime
Our all plans come with 99.9% uptime guarantee with no downtime.

99.9% vs. 99.999%
the reliability difference

...ime
...g server

...f at our data centers for ...ave collaborated with ...olutely dependable ...9.9% online availability ...Datacenters that

UPTIME GUARANTEE 99.9% NETWORK UPTIME ...E GUARA...

UPTIME GUARANTEE 99.9% UPTIME GUARANTEE

99,9 % UPTIME

99.999% UPTIME

We use one of the world's largest CDN ...Technologies, Inc.) to guarantee a 99.9... make sure your site never delays on th... communicate any scheduled downtim... site is unaffected by a slow server on... technology that excludes a visitor wh... fall-back has been reached.

Practical Use: No worries for the IT ... and client load-time guarantees.

We have a great infrastructure on ... Nothing less than 99.999% uptim...

Public announcement of schedule... downtime—ever!

Even when they don't actually know what that means...

Maximizing availability

- Two complementary approaches:

- **High Availability (HA)**
  - Design systems so that failures do not result in *any* downtime
  - Keys: redundancy and failover
  - "No single point of failure"

- **Disaster Recovery (DR)**
  - If the HA techniques are overwhelmed (e.g. due to a site failure, major human error, etc.), be able to recover data and restore functionality
  - Keys: replication and restore/failover
  - "Survive the inevitable multiple failure"

High Availability (HA)

Redundancy

- Core HA concept:
  - Identify single points of failure
  - Add redundancy to eliminate
  - Need policy for how to interface with redundant system

Active/active vs. active/passive

- **Active/active**: Both redundant components used in normal operation; symmetric design
  - + Higher utilization and capacity/performance
  - – Capacity/performance is reduced on failure

- **Active/passive**: A "primary" and "secondary" system; secondary only does work if primary fails; asymmetric design
  - + Failures don't affect capacity/performance
  - – Half the hardware is idle most of its life (low utilization)

The split brain problem

- Imagine an active/passive system
- What if the two redundant systems lose contact with each other, and each thinks its time to "take over"?
- Both are serving traffic and issuing commands!
- Result: chaos!

- Redundant computer systems must have protocol to govern takeover

SPECIFICITY WARNING

- Warning: a lot of the examples in this deck are particular implementations of general concepts.

- The storage system you use (or invent) will likely work a little different, but the core ideas will be the same

(Translation: I'm not trying to sell you NetApp, but because I happen to intimately know how they address HA issues, I'll use them as an example. I am not a shill.)

Layers on which to apply redundancy

Power

Cabling

Servers

Network switch

Storage controller

Disks
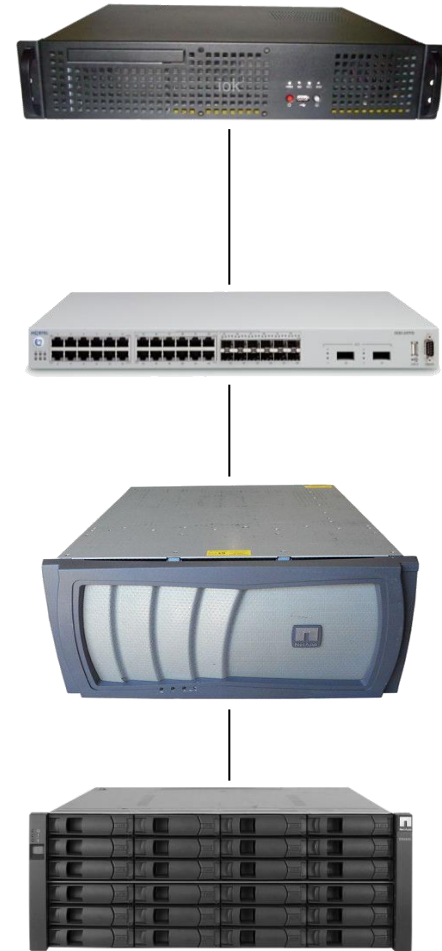
First, let's deal with power

- Everything has 2+ power supplies
  - Equipment can survive with half its power supplies dead
  - This protects against power supply failure
- Power comes from Power Distribution Units (PDUs) (basically rackmount power strips)
- HA power: Racks have two PDUs.
  - PDU 1 hooked to "left" power supply, PDU 2 hooked to "right" power supply
- Power supplies usually hot-swappable
  - Replace on fault without downtime

Figure from "Do dual-power supply servers increase redundancy?", TechTarget. Julius Neudorfer.
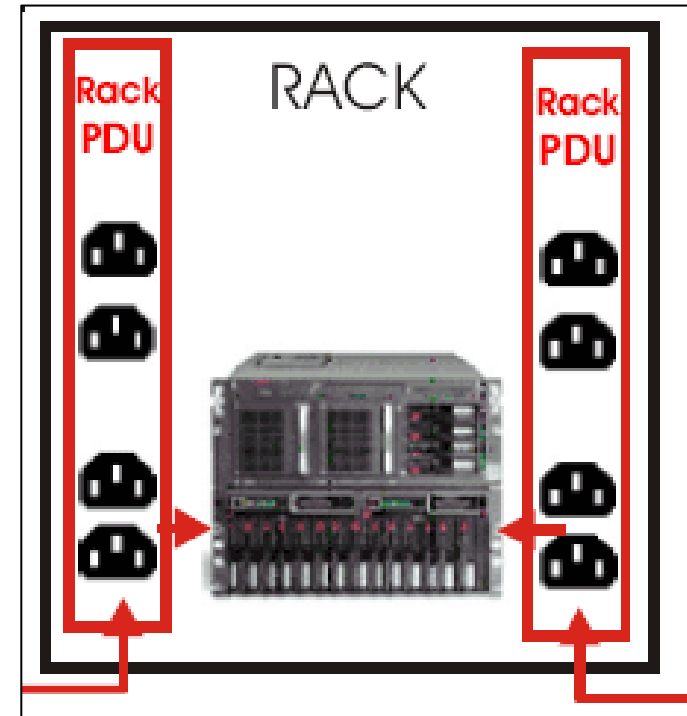
Utility power

- Single-feed environment: both strips get power from same utility
  - Power outage? All gear goes down.
  - Still protects against accidental disconnect, power supply failure, local tripped breaker
- Double-feed environment: two separate feeds from two separate power sub-stations
  - Also protects against utility power outage
  - Might even draw from two different power plants!
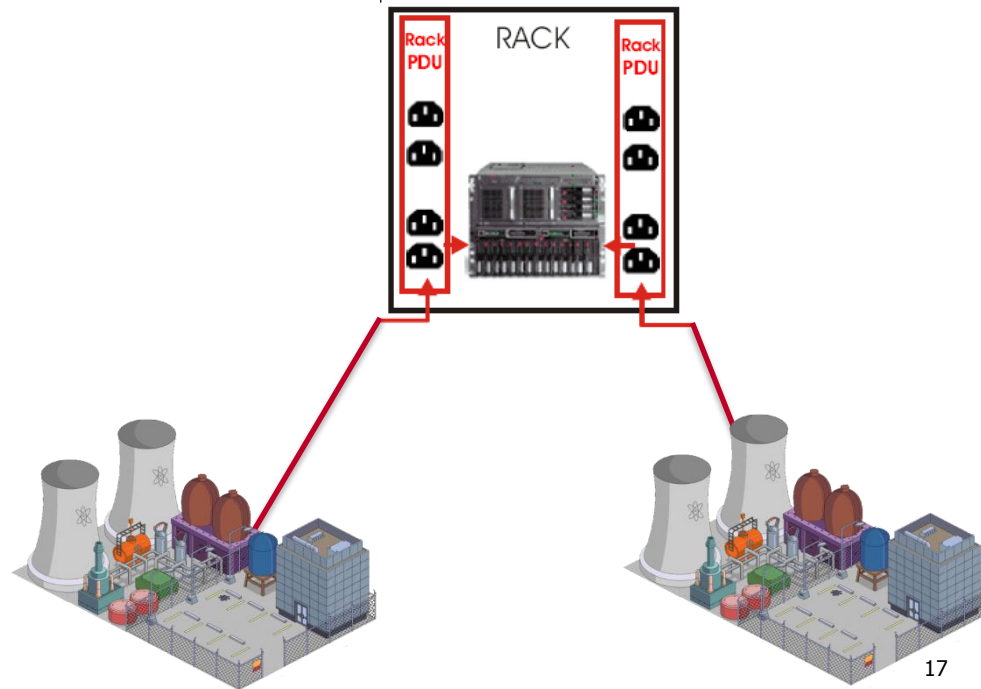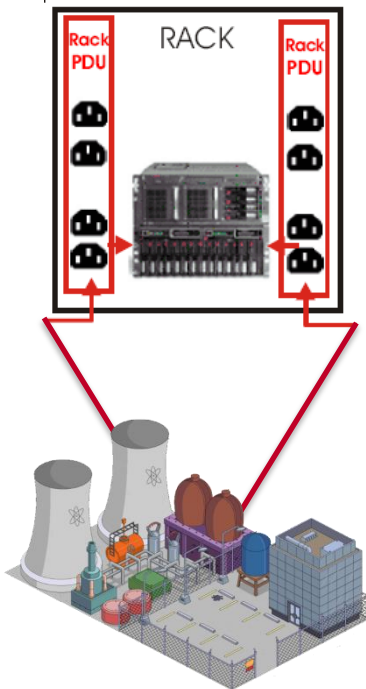
UPS: Uninterruptable Power Supply

- UPS: Uninterruptable Power Supply
  - Takes AC power in, gives AC power out
  - Keeps a big battery array charged
  - If AC power-in fails, AC power-out comes from battery array without interruption
    - DC power from batteries must be converted to AC with an inverter
  - Rated by **battery capacity (total energy)** and inverter **current capability (max power)**
- Use cases
  - Smooth power "**blips**" (momentary interruptions that would reboot everything)
  - Keep things running for a few **minutes**, long enough for graceful shutdown
  - Run things for a few **hours**
  - Keep things running long enough to start a **gasoline/diesel generator**

Small consumer UPS   Rackmount UPS   Building-scale battery array

Electric generators

- Need to survive long-term power loss?
- **Gasoline or diesel generator**
- Typically sized for whole or part of data center
- Large fuel tank on site, run time for days
- Can contract to have additional fuel brought during extended emergency
  - If fuel can't be brought, society is probably broken enough that it's okay your server isn't up…

Buy this Generator Online Today

Generac Commercial 150kW (Alum) NG 240V Single Phase

- Two Line LCD Tri-Lingual Digital Nexus Controller
- Isochronous Electronic Governor
- Sound Attenuated Enclosure
- Closed Coolant Recovery System
- Smart Battery Charger
- UV/Ozone Resistant Hoses
- ±1% Voltage Regulation
- Natural Gas or LP Operation
- 2 Year Limited Warranty
- UL 2200 Listed

Whoa! No review on this Product

Be the first to review this product and earn a chance for a $10 store credit when you place a review on this product post purchase.

Our Phone Lines Are Open!
Call Us Toll-Free: 855-453-4494

GENERAC

Model #
QT15068ANAC

MSRP: $29,999.00

NORWALL'S LOWEST PRICE $29,399.00

YOU SAVE: $600.00

Free Shipping (see terms)
Generac Referral Rebate (open)

Availability
Built To Order In 5-6 Weeks

QUANTITY 1

ADD TO CART

Add To Compare | Add to Wish List

Generac Power Systems QT15068ANAC Details and Specifications

≡ Description	✱ Warranty	🗎 Spec's	★ Reviews

Generac Commercial 150kW (Alum) NG 240V/Single Phase Liquid Cooled Standby Generator.

This premium-grade generator features a high-quality automotive style engine that runs at a low speed RPM mode (Quiet Test). This means a

Product Information PDFs

Total redundant power picture
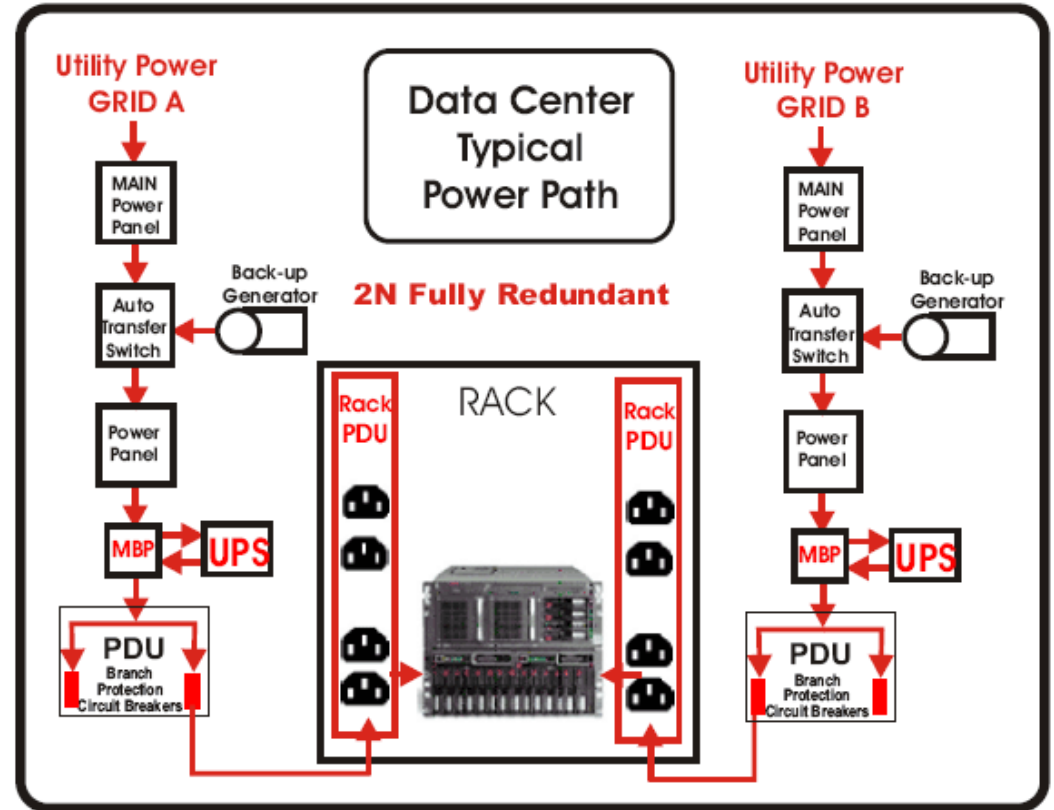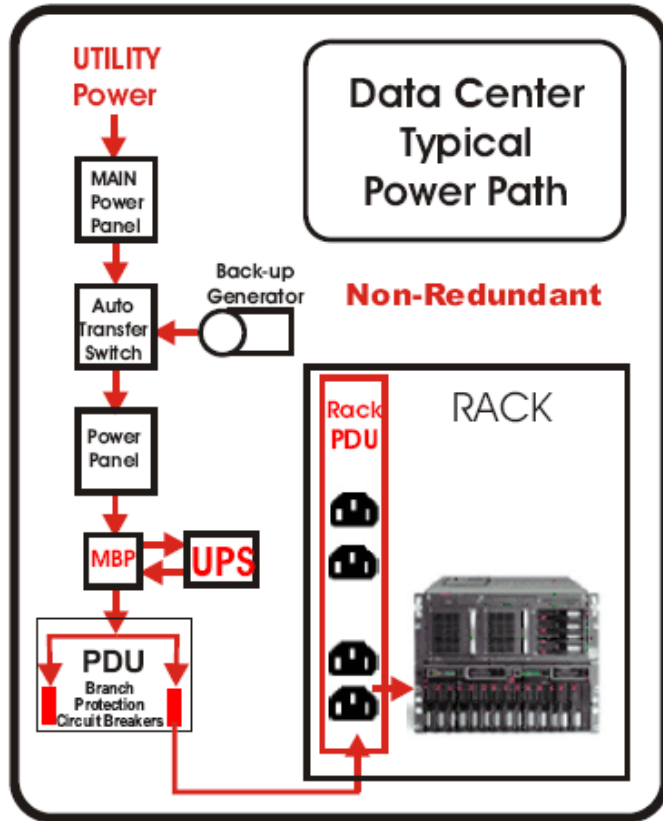
Figure from "Do dual-power supply servers increase redundancy?", TechTarget. Julius Neudorfer.

Layers on which to apply redundancy

Power

Cabling

Servers

Network switch

Storage controller
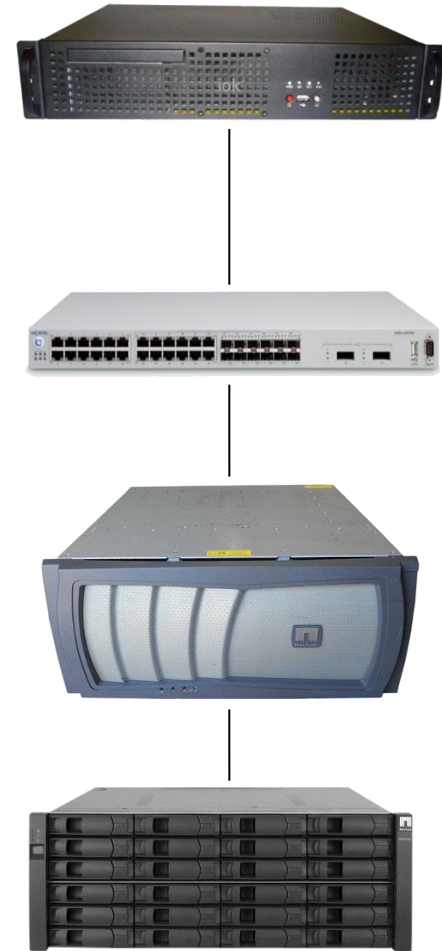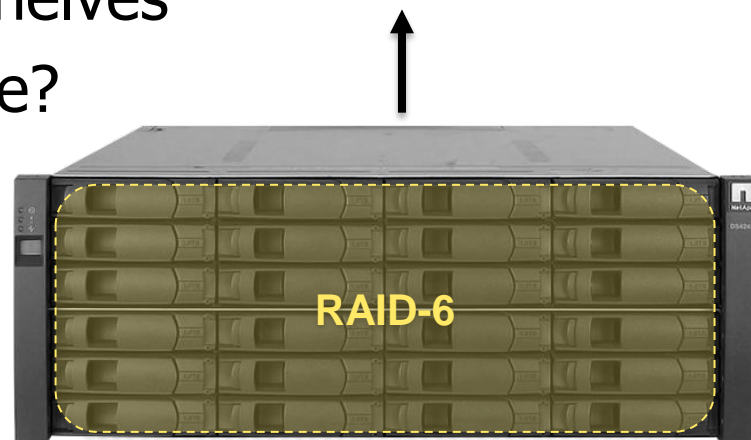
Disks

Redundancy: disks

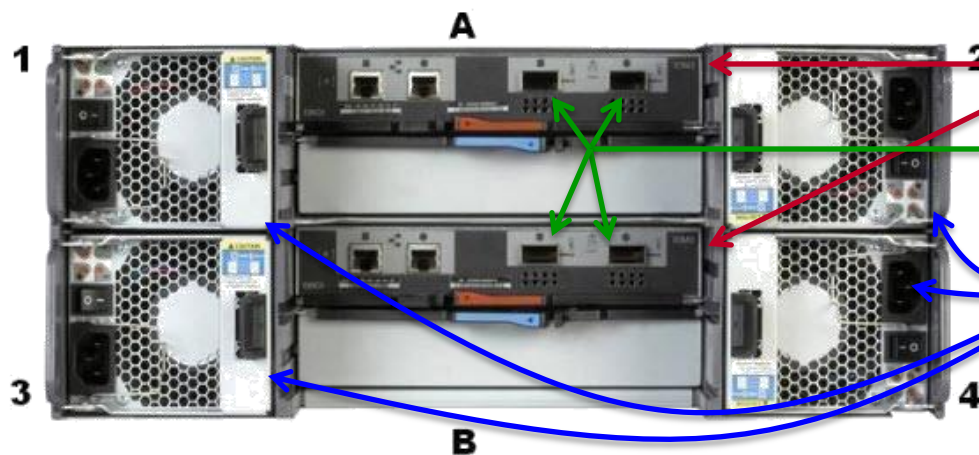- What technique do we know to provide redundancy against the failure of individual disks?

RAID

So we do that.

Redundancy: Disk shelves

- Disks physically installed into disk shelves
- Is there a single point of failure here?
  - Yes, the uplink!

- Actual back of this shelf:

RAID-6

- Two IO Modules
  - Each with two SAS ports
- Also two power supplies per IOM

Legend:
A - IOM Module A (upper)
B - IOM Module B (lower)
1 - Left power supply (upper)
2 - Right power supply (upper, and rotated)
3 - Left power supply (lower)
4 - Right power supply (lower, and rotated)

Redundancy: Disk shelves

- So now we did this. Any single points of failure?
  - Yes, the controller.

RAID-6

Layers on which to apply redundancy

Power

Cabling

Servers

Network switch

Storage controller

Disks
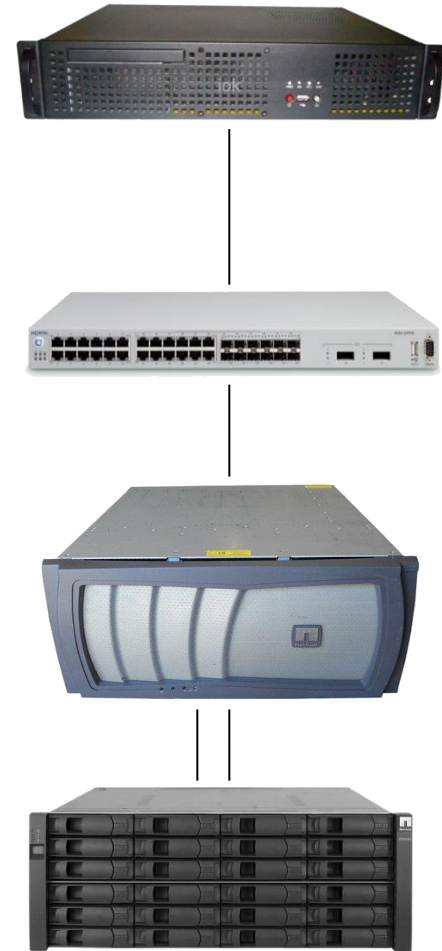
Redundancy: Storage controller

- ## So we want something like this.

Actual back of that storage controller and its connection to a disk shelf:

Controller A

Controller B

IOM A

IOM B

Note: you almost always want storage controllers to be highly available, so they're often sold as "two headed" units, where there's two of everything in one box. To keep it simple, we'll ignore that and use an example where there's literally two boxes.
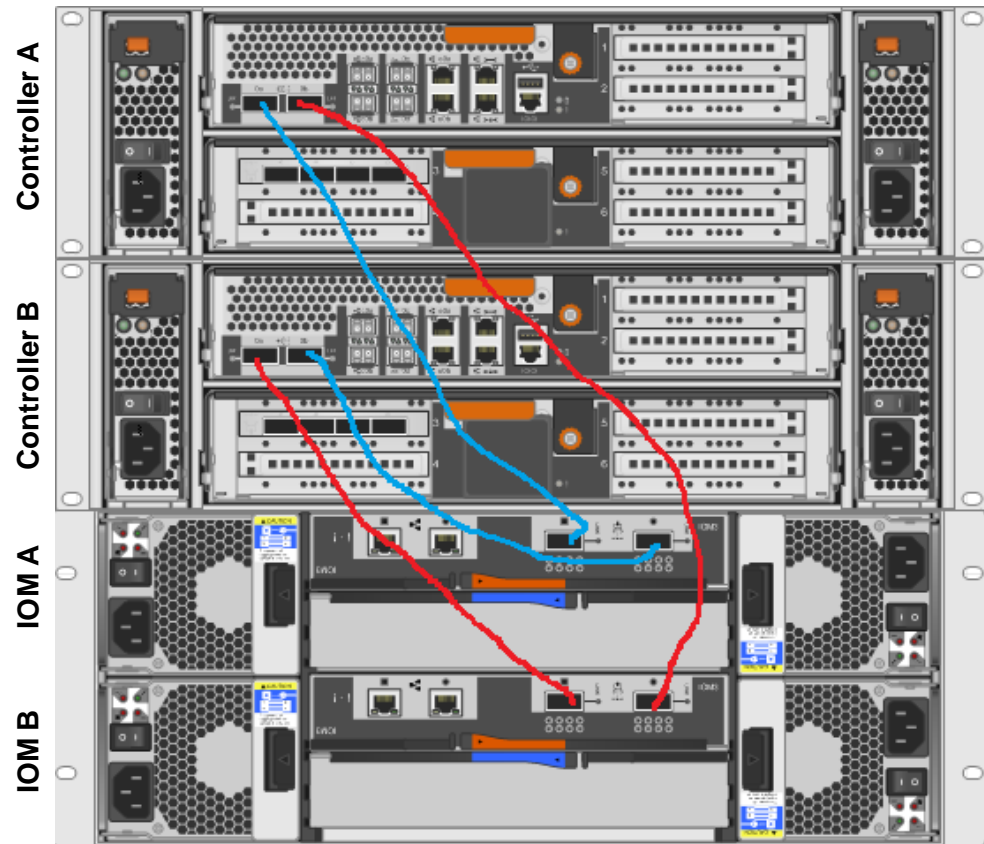
This is an example of what I call the "**universal HA topology**"

26

The universal HA topology

(a term I made up)

Server A — Inter-server link — Server B

Client A — Inter-client link — Client B

See the topology?

Server A — Inter-server link — Server B

Client A — Inter-client link — Client B

Controller A

Controller B

HA pair link

IOM A

IOM B
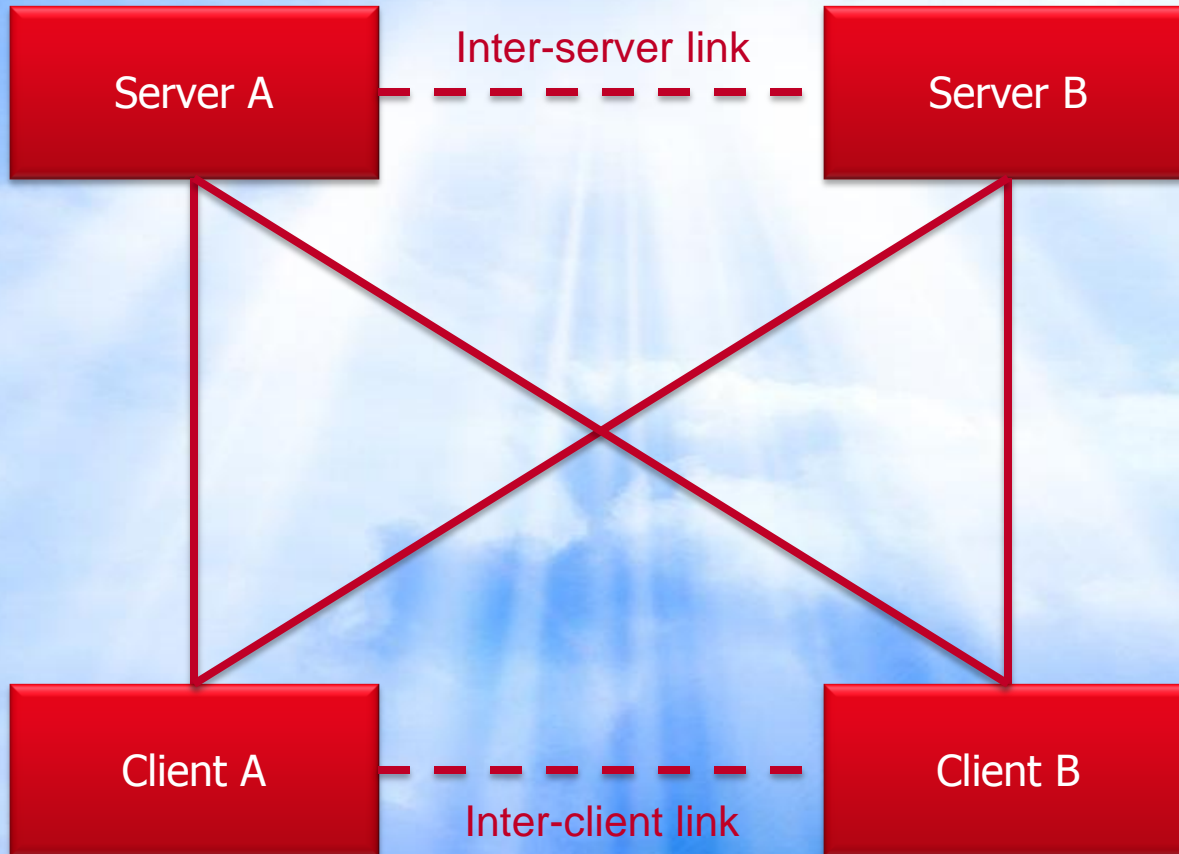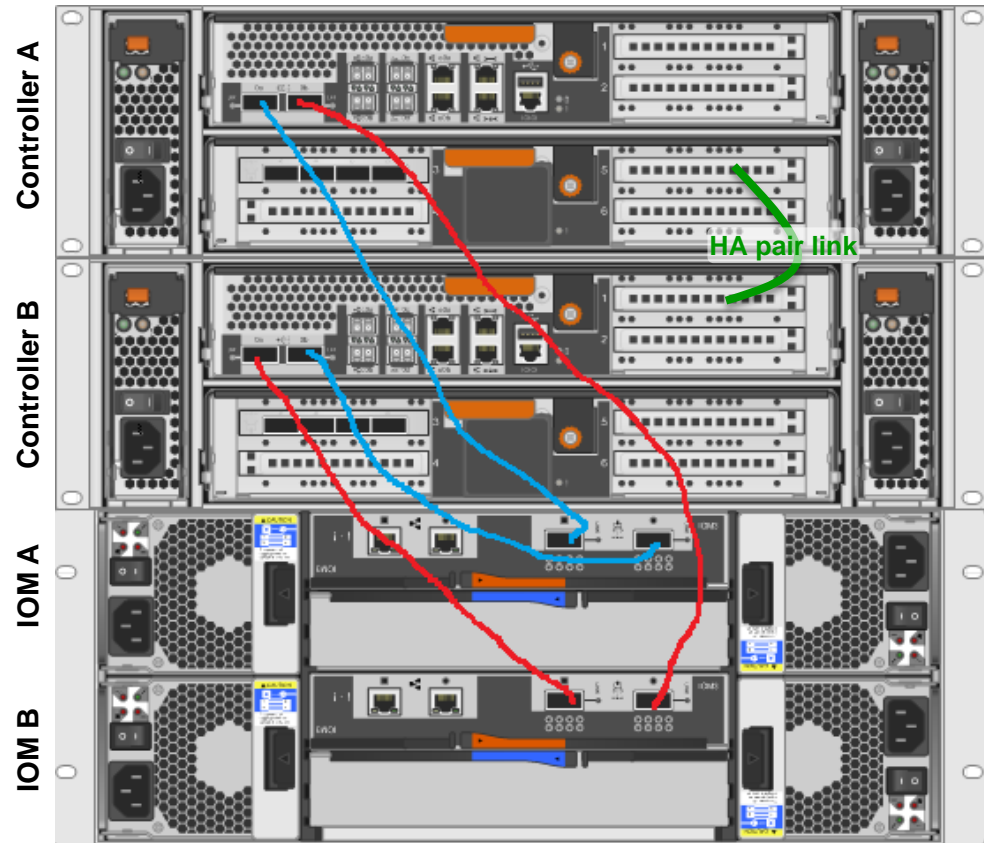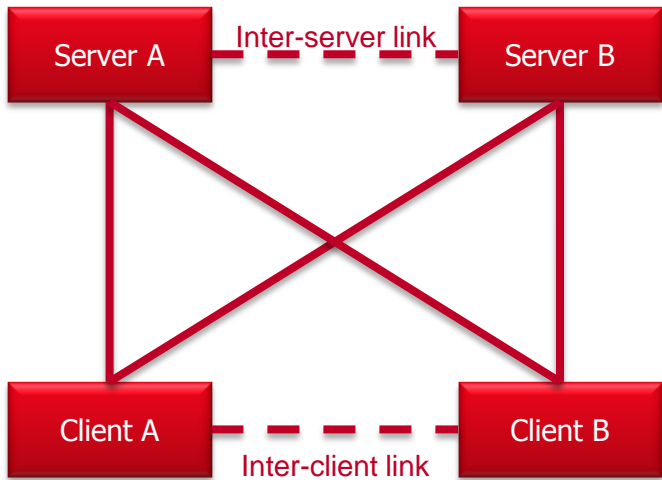
What does controller redundancy mean?

- Each controller can "see" all the disks, as well as each other
- Could use this to be active/passive:
  - All disks are "owned" by controller A
  - All storage is provisioned on controller A
  - If controller A goes down, controller B can "take over" and basically impersonate controller A without missing a beat
- Could use this to be active/active:
  - Half disks are "owned" by controller A, half by controller B
  - Provision workloads on each controller
  - If controller A goes down, controller B is responsible for both its own usual workload as well as controller A's
- In either case, when situation is resolved:
  - Controller B can "give back" to controller A the work it took over
  - This restores HA capability after a failure

Layers on which to apply redundancy

Power

Cabling

Servers

Network switch

Storage controller

Disks
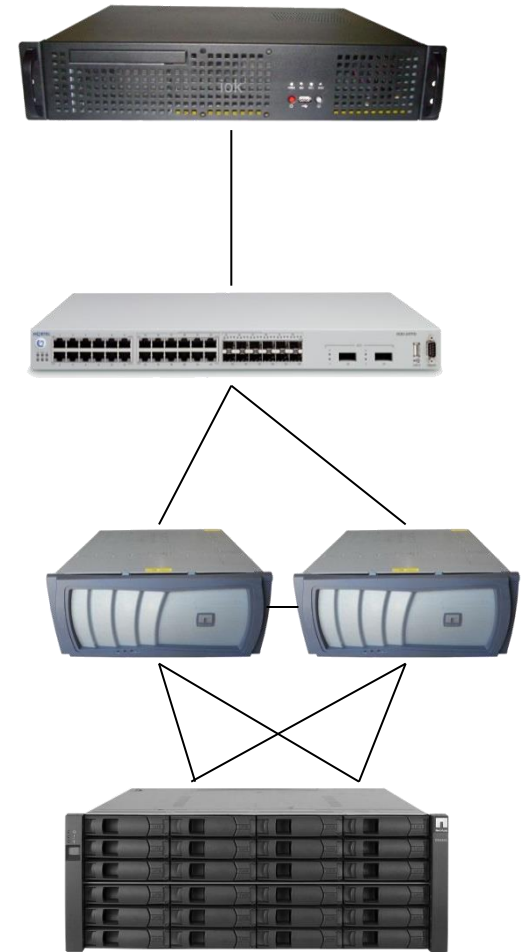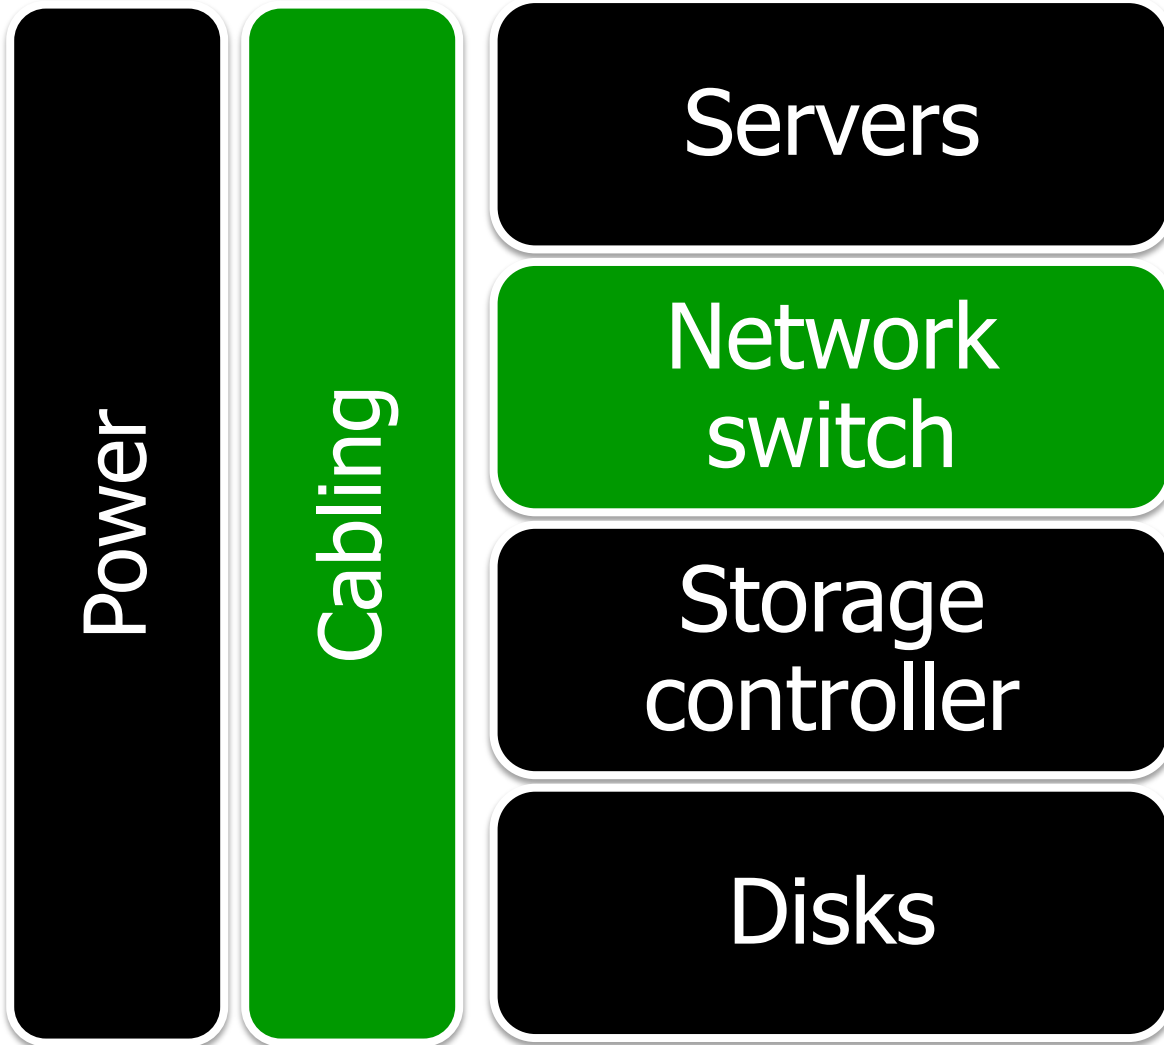
Network redundancy

- Apply the Universal HA topology
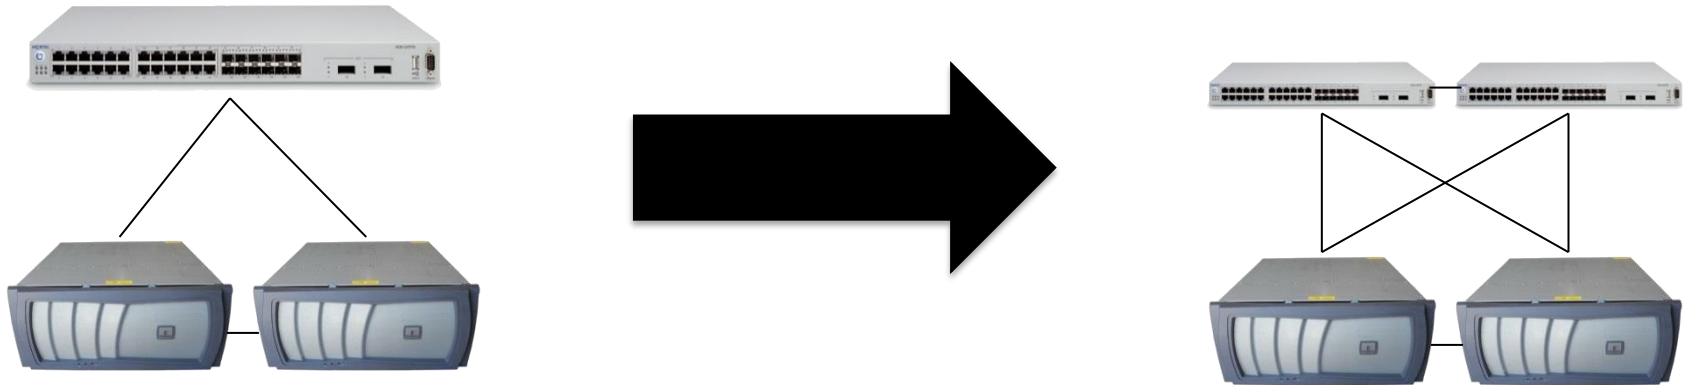
- In networking, this is known as **multipathing**
- Can be applied to Ethernet or Fibre Channel (FCP)

Ethernet Multipathing (1)

- Known as **Ethernet Link Aggregation**
- Inter-switch configuration:
  - Stackable switches
    - Two switches connected by a proprietary cable to extend the backplane
    - Both switches managed as single entity
  - Peer Link for "Virtual Port Channel" support (Cisco term)
    - Switches are separate, but can cooperate over an Ethernet link between them

Ethernet Multipathing (2)

- Downlinks:
  - Ports on each <u>switch</u> heading to the same target are logically combined (a "port channel" or "virtual port channel" in Cisco terminology)
  - Ports on each <u>endpoint</u> (e.g. storage controller) are similarly bonded, this is often called "link bonding"

- Port bonding situation can be configured statically on both switch and endpoint

- Alternative: Link Aggregation Control Protocol (LACP)
  - Standardized protocol to allow automatic negotiation of link bonding

Fibre Channel Multipathing

- In Ethernet, multipathing requires switch support and is somewhat of a special case

- In Fibre Channel, multipathing is common and doesn't require much special configuration
  - Just make sure the two switch configurations match
  - Vendors often have proprietary protocols or tools to ensure this

Layers on which to apply redundancy

Power

Cabling

Servers

Network switch
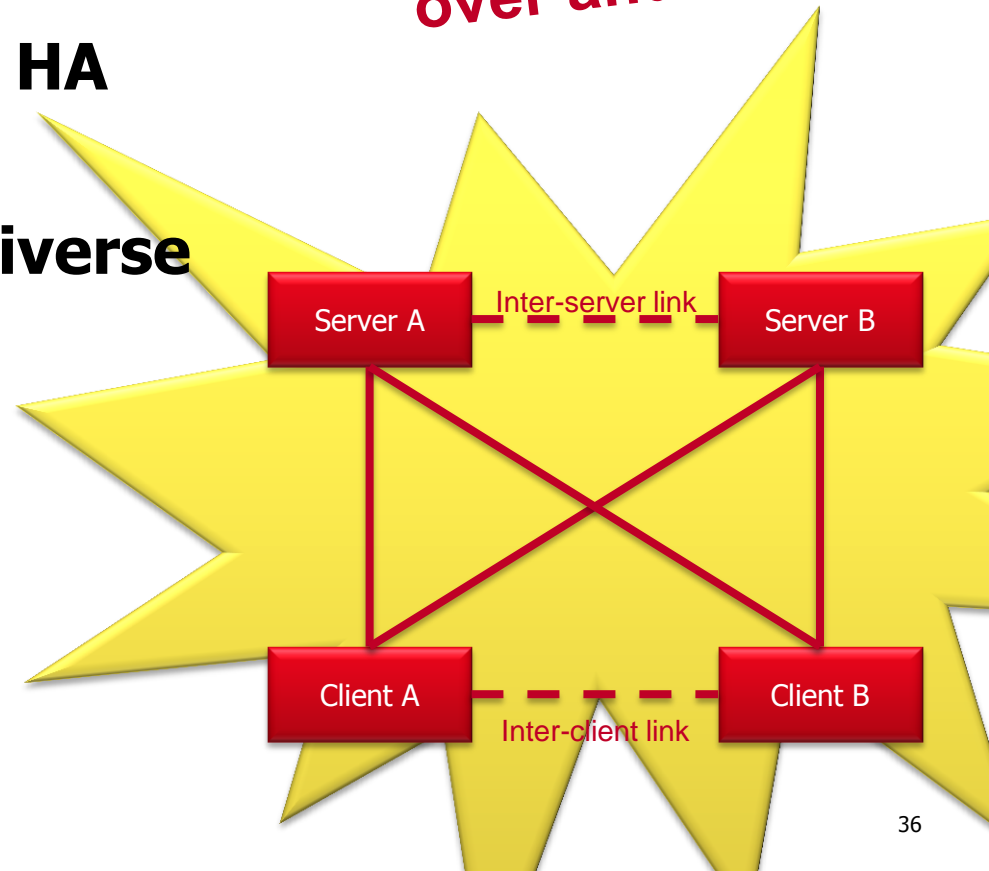
Storage controller

Disks
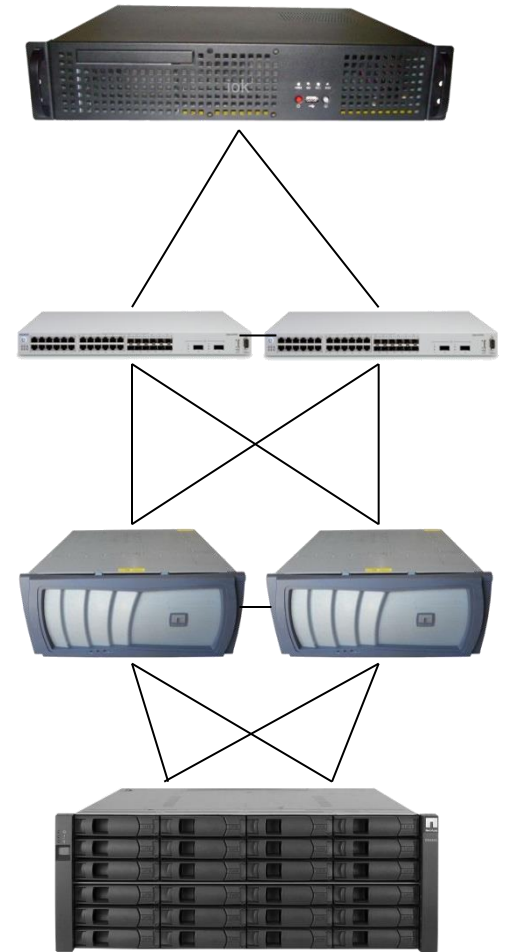
Server redundancy

- We need to make servers redundant

- We need a **topology** that does this

- The topology will make them **HA**

- This will take place in our **universe**

- what can we choose

- help

- **what** can we chooooooooooooooooooose???????

oh yeah, that thing we keep using over and over

| Server A | Inter-server link | Server B |

| Client A | Client B |

Inter-client link

Server redundancy

- Apply the Universal HA topology
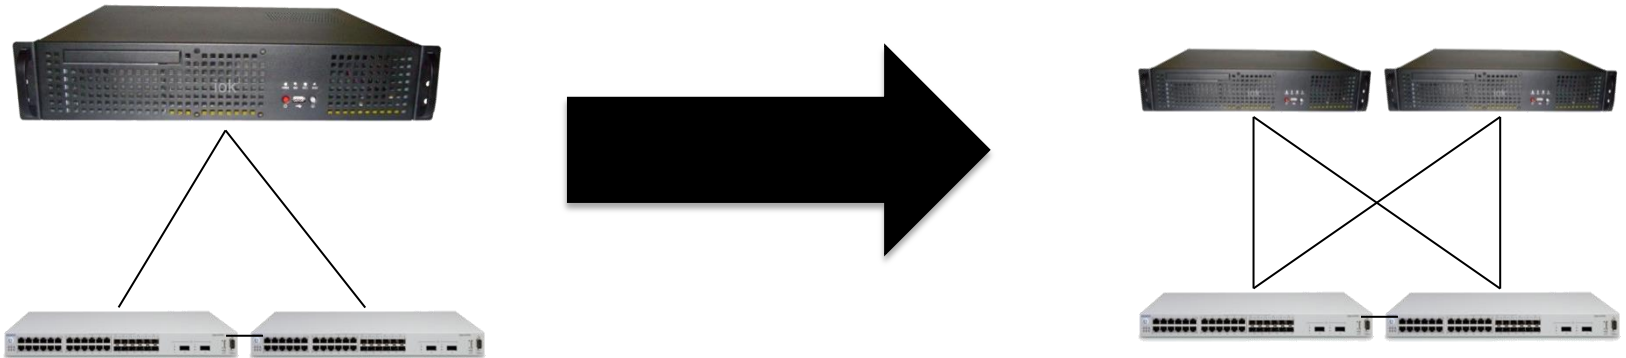
- However, typically have more than 2 servers – storage/network usually serves pool of many servers
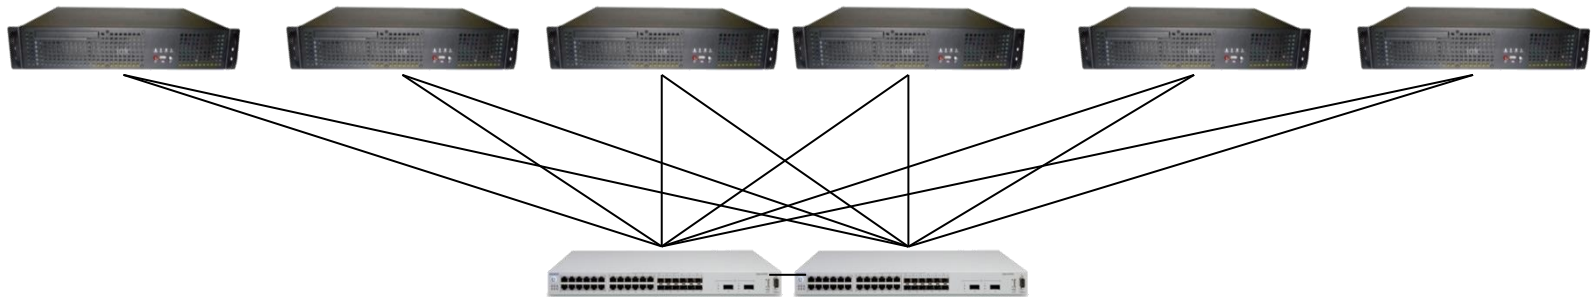
Software support for redundancy

- Physical connectivity is simple; software side is complex:
  - What is the effect of running two copies of the software?
    - Depends on the software!

- Many techniques/mechanisms to take advantage of redundancy:
  1. Truly redundant hardware
  2. Redundancy via hardware abstraction
  3. Redundancy via hypervisor abstraction
  4. Hypervisor-based virtual fault tolerance
  5. Application-based fault tolerance

Mechanisms for server redundancy

- **Truly redundant hardware**
  - Not common; very expensive and proprietary
  - Redundancy at the hardware level; software is oblivious
  - Requires very fine-grained synchronization, lots of bandwidth
  - **Result**: Hardware failure produces no measurable effect to the outside world; failed hardware can be hot-swapped afterward
  - **Example**: Tandem Systems shared-nothing message-passing system; now sold as "HP Integrity NonStop X"[1]

A Tandem T/16 memory board; it just has two of everything.

[1] The larger the company, the stupider the naming.

Mechanisms for server redundancy

- **Redundancy via hardware abstraction**
  - Whole servers can "fail over" to spares
  - All server-unique configuration is managed by a server manager
    - Ethernet MAC addresses
    - Fibre Channel World-Wide Names (WWNs)
    - BIOS/firmware settings (e.g. boot order)
  - Servers are of uniform configuration; some servers are "hot spares"
  - On failure:
    - Management system writes hardware config onto a spare server; boots it
    - The replacement server is indistinguishable from the original; software works the same
  - **Result**: Server hardware failure is just like a simple reboot
  - **Example**: Cisco Unified Computing System (UCS) servers

Mechanisms for server redundancy

- **Redundancy via _hypervisor_ abstraction**
  - Virtual machines can "fail over" to other physical hosts
  - Physical servers need not be uniform
  - Don't need entire "hot spare" hosts; just need "spare" compute/memory capacity on overall cluster
  - On physical server failure:
    - VM management system identifies affected VMs; boots them on another server
    - This is possible because VM virtual disks are on shared storage (SAN or NAS)
  - **Result**: Physical server failures act like VM reboot events
  - **Example**: The VMware HA Clustering feature

Mechanisms for server redundancy

- **Hypervisor-based virtual fault tolerance**
  - A VM is "run" on two separate physical hosts
  - On the "primary" host, the VM's system calls actually happen as normal
  - On the "secondary" host, the VM's system calls are faked – responses actually come from whatever the primary system got
  - Result: because computers are deterministic, both systems' computations proceed identically
  - On failure of the primary physical host, the secondary is simply made primary, and its system calls begin happening "for real"
  - Outside world cannot tell that changeover has happened
  - **Result**: Zero perceived downtime
  - **Example**: VMware Fault Tolerance feature

Mechanisms for server redundancy

- **Application-based fault tolerance**
  - The user application has built-in support for some kind of HA clustering
  - May work with performance-based clustering (i.e. scaling application performance by adding more servers) or be totally separate (e.g. an active/passive app)
  - Pro: Application does its own consistency, can achieve higher performance than the previous application-oblivious techniques
  - Con: Developers have to consciously design application with this in mind
  - **Result**: Depends on how app is built, but typically fault-tolerant apps allow server failure without measurable effect to outside world.
  - **Example**: Microsoft SQL Server Failover Clustering
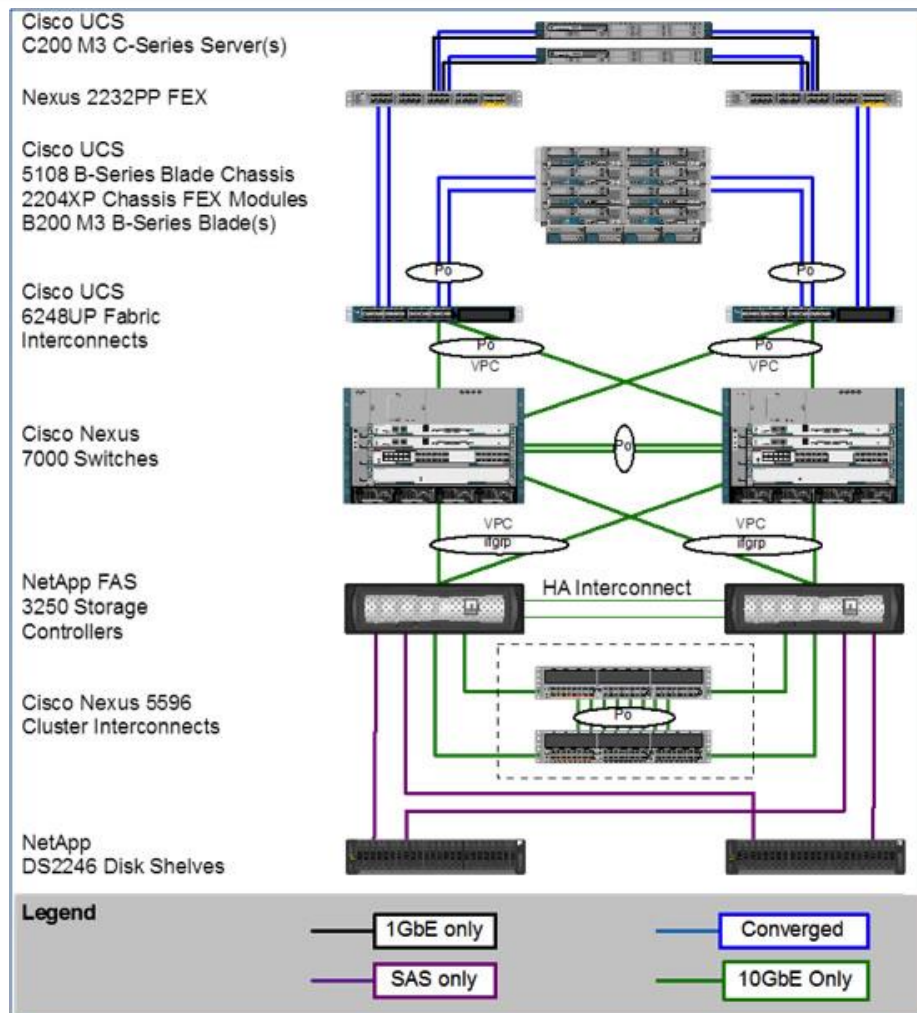
We did it!

Power

Cabling

Servers

Network switch

Storage controller

Disks
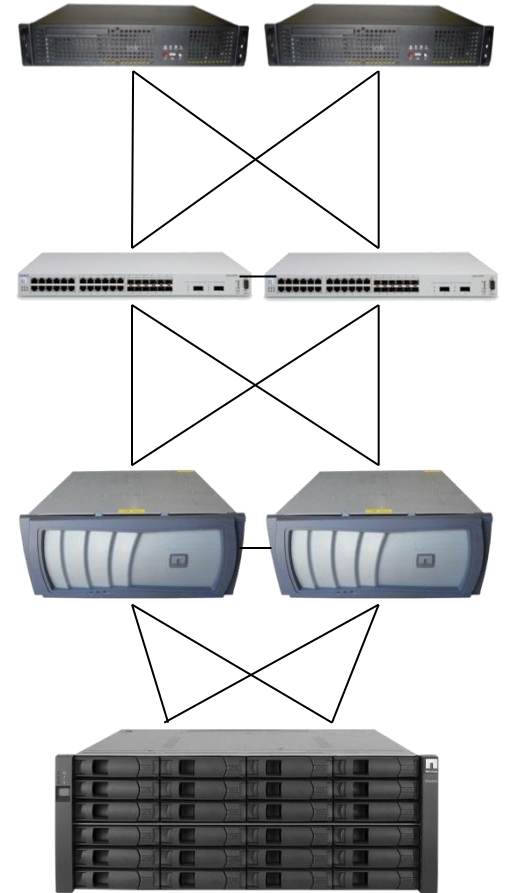
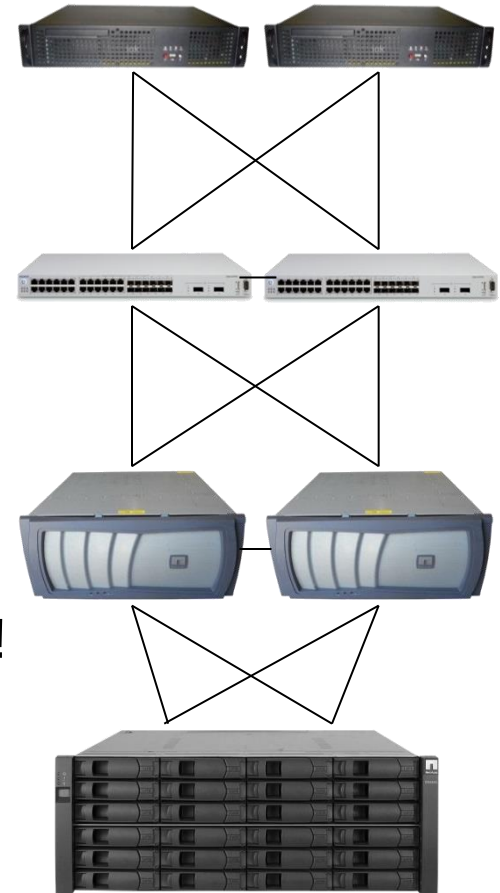# An example we saw in the course introduction

- From: http://community.netapp.com/t5/Tech-OnTap-Articles/FlexPod-Innovation-and-Evolution/ta-p/85156

But what if...?

- A meteor lands on our datacenter?
- Can we be HA against that?
  - Surprisingly, **yes** (for a small enough meteor)!
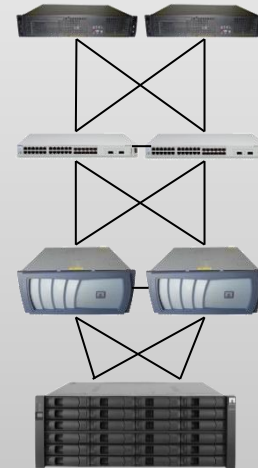
Site-level protection

Power | Cabling | Servers | Network switch | Storage controller | Disks

Thought experiment

- HA works if each redundant pair is 1 meter away
- Does it still work at 2 meters?
- Does it still work at 4 meters?
- Does it still work at 8 meters?
- …
- What's the limit?
- What affects the limit?
  - Latency
  - Ability of cable to carry data that far
- Practical answer: around 100km (depends on many things)
  - FYI: (100 kilometers) / the speed of light = 333 microseconds

- Result: Metro-scale clustering
  - "Metro-scale" = Around the size of a city

Max ~100km
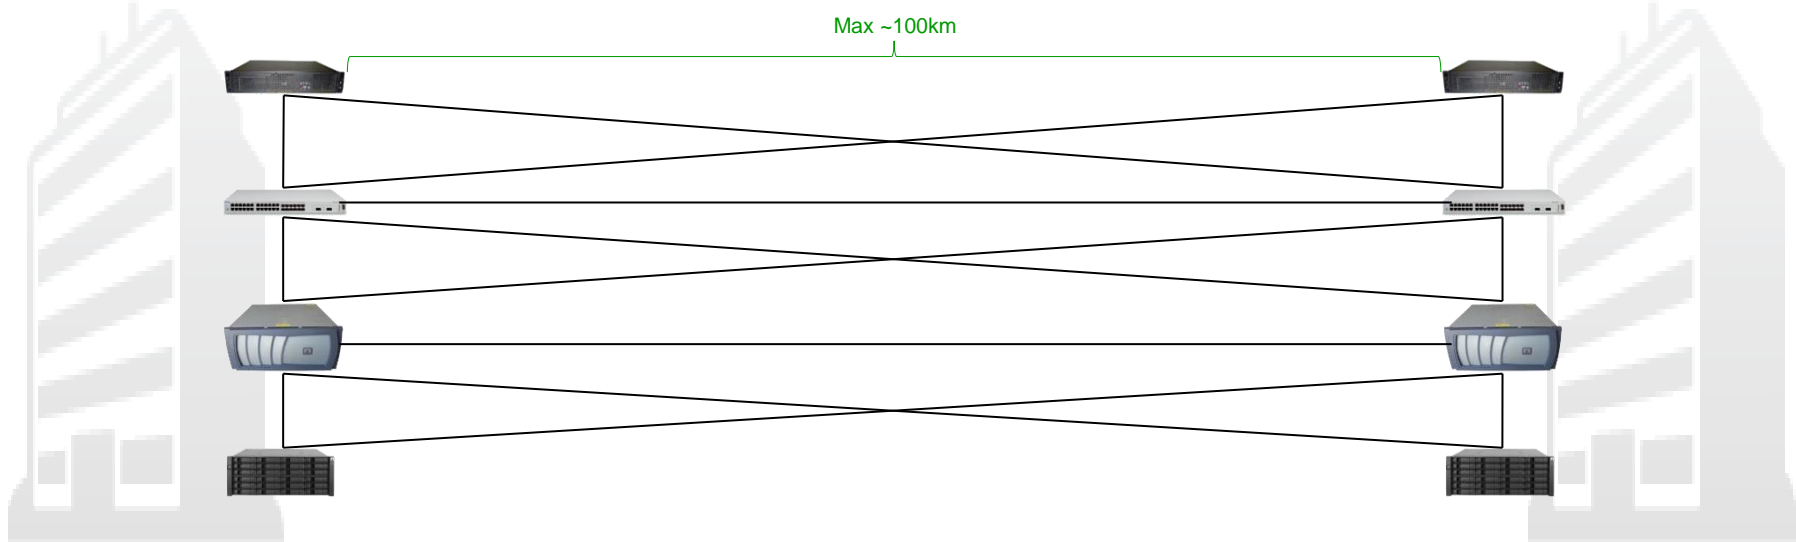
- Often deployed just at campus-scale (a few kilometers); sometimes deployed all the way between cities (especially in Europe, where cities are closer)

- Can also be applied to just storage: then it's a form of backup/replication, which we'll cover when we talk about disaster recovery

- **Result:** You can lose an ENTIRE DATACENTER and keep serving traffic with little to no interruption

- **Example**: NetApp Metrocluster plus VMware Stretch Cluster

Connectivity

- Connection between sites typically dedicated optical fiber
  - Fiber optics can run data faster over much longer distance than copper

- How to get?
  - Dark fiber: abandoned pre-existing line
  - New fiber: pay huge cost to run a new buried line
  - Leased line: pay for bandwidth on existing dedicated lines

- Can also tunnel over existing network or the internet
  - Performance penalty, or even unpredictable performance
  - Can be okay for iSCSI/NAS, not common for FCP

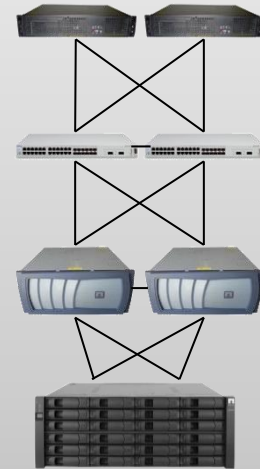## Site-level protection

Power

Cabling

Servers

Network switch

Storage controller

Disks

But what happens if something overwhelms these protections? Need ***disaster recovery*** (next).