Introduction to Dynamic Infrastructure

By Alan G. Labouseur

MARIST SCHOOL OF COMPUTER SCIENCE & MATHEMATICS

MARIST SCHOOL OF COMPUTER SCIENCE & MATHEMATICS

IBM

MODULE THREE

INFORMATION INFRASTRUCTURE

Introduction to Dynamic Infrastructure

Designed, compiled, written, and edited by
Alan G. Labouseur
www.Labouseur.com / Alan.Labouseur@Marist.edu

Contents

Remarks

I liked the discussions last week especially of the virtues of the name "Dynamic Infrastructure".

- I was leaning toward "Smart Business Infrastructure" myself, but you guys convinced me that DI is the way to go.
- I've even changed the labels on my web site for this.

The case studies were very good too. Random thoughts:

- I worry a little bit about Brian throwing darts at his screen, but better his than mine.
- Mike, I wonder how much Muamalat is liking their Sun hardware now that it's really Oracle that's runnign things.
- On being green, two things:
  - There's a difference between really doing it and "green-washing". That might be a good discussion topic for this week.
  - I agree about the solar panels... were are they? Why isn't my iPad solar charged?
- As Volkswagen discovered, virtualization is rilly rilly cool. We'll get into that in the second-to-last class.
- Brian makes a great point about IBM Rochester not talking about how much it cost to implement all that instrumentation. I wonder what the break-even point was (or is).
- The University of Bari is not alone in cloud computing: our z/Series here at Marist support over 600 virtual Linux servers for various and sundry student and faculty projects. (Or it once did. To be honest, I don't hear that much about it any more.)
- About the Energy-efficient Data Center: IBM had a "Green Data Center" in building 701 in their Poughkeepsie lab. Within it lies a heat monitor device which is basically an Apple Macbook Air (running Windows) bolted to a metal

Review

Let's review. Last time our intrepid (and dynamic) adventurers . . .

- learned that the three primary goals of a Dynamic Infrastructure are to
  - Improve Service
    - with cloud based services and solutions, service management industry solutions, and application management and hosting.
  - Reduce Costs
    - with energy efficient servers, storage, and facilities, virtualization and consolidation, effective information infrastructure, and standardization and automation.
  - Manage Risk
    - with pervasive and preventive security solutions, global resiliency and security centers, comprehensive resiliency solutions, and compliance and long- term information retention policies.
- and that we have to "Think Different" to do so.
  - Trivia: what's the origin of the words "Think" and "different" in the graphic on slide 13 of last week's module?
- read about the 007 pillars of a Dynamic Infrastructure:
  - Information Infrastructure, Service Management, Asset Management, Energy Efficiency, Business Resiliency, Security, and Virtualization
- discovered that Dynamic Infrastructure is, in a rather clichéd sense, not a destination but rather a **Journey**

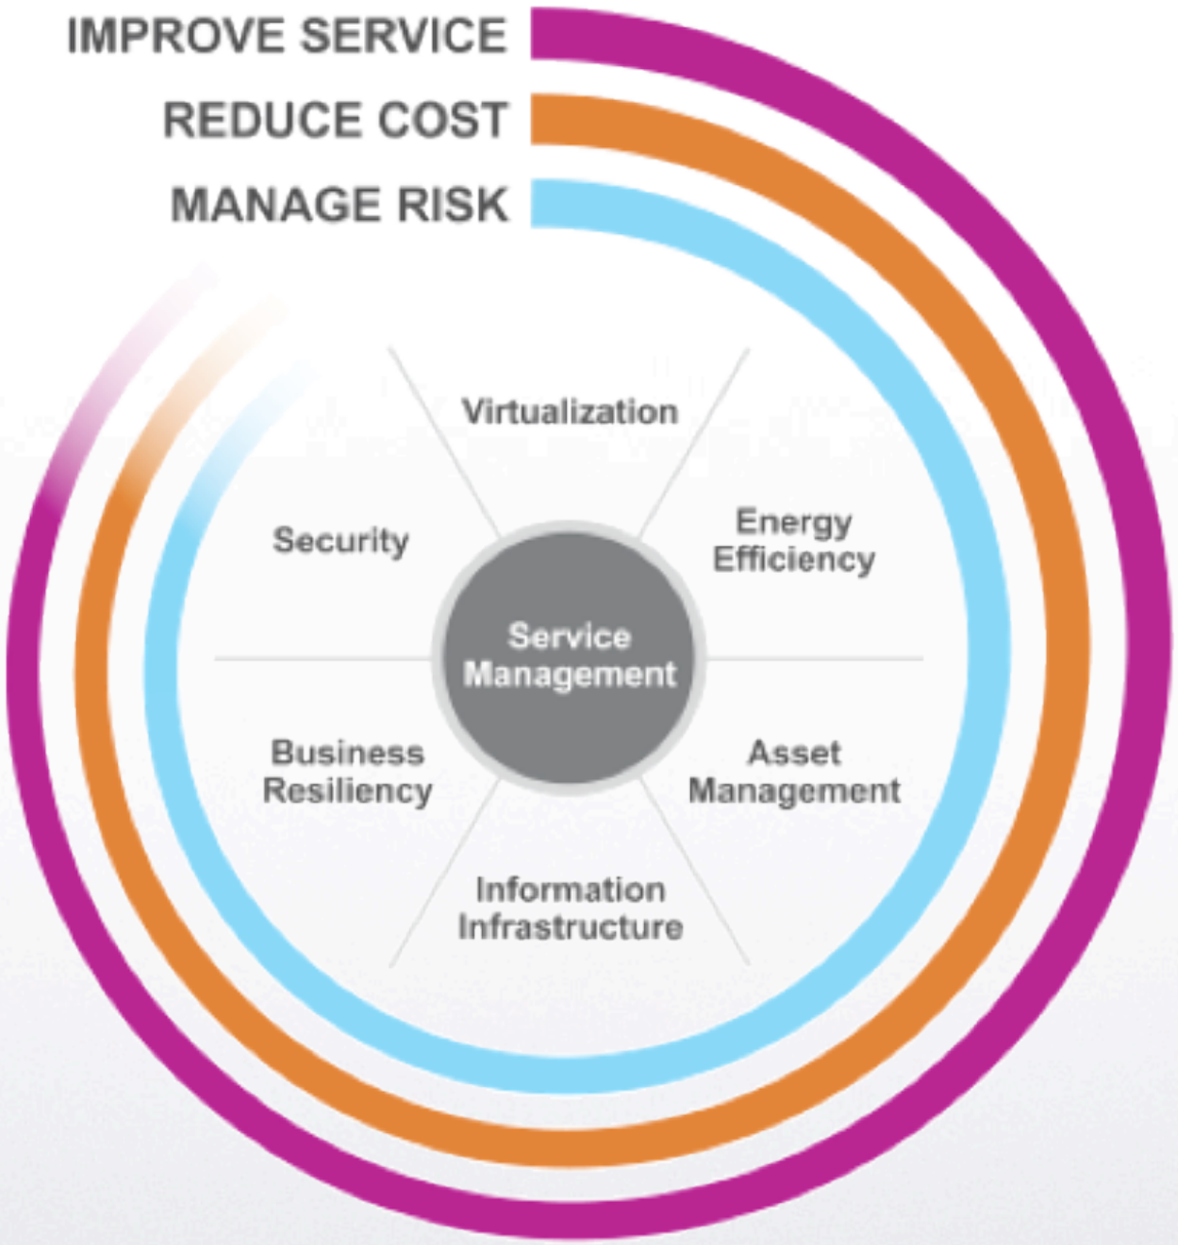

Information Infrastructure

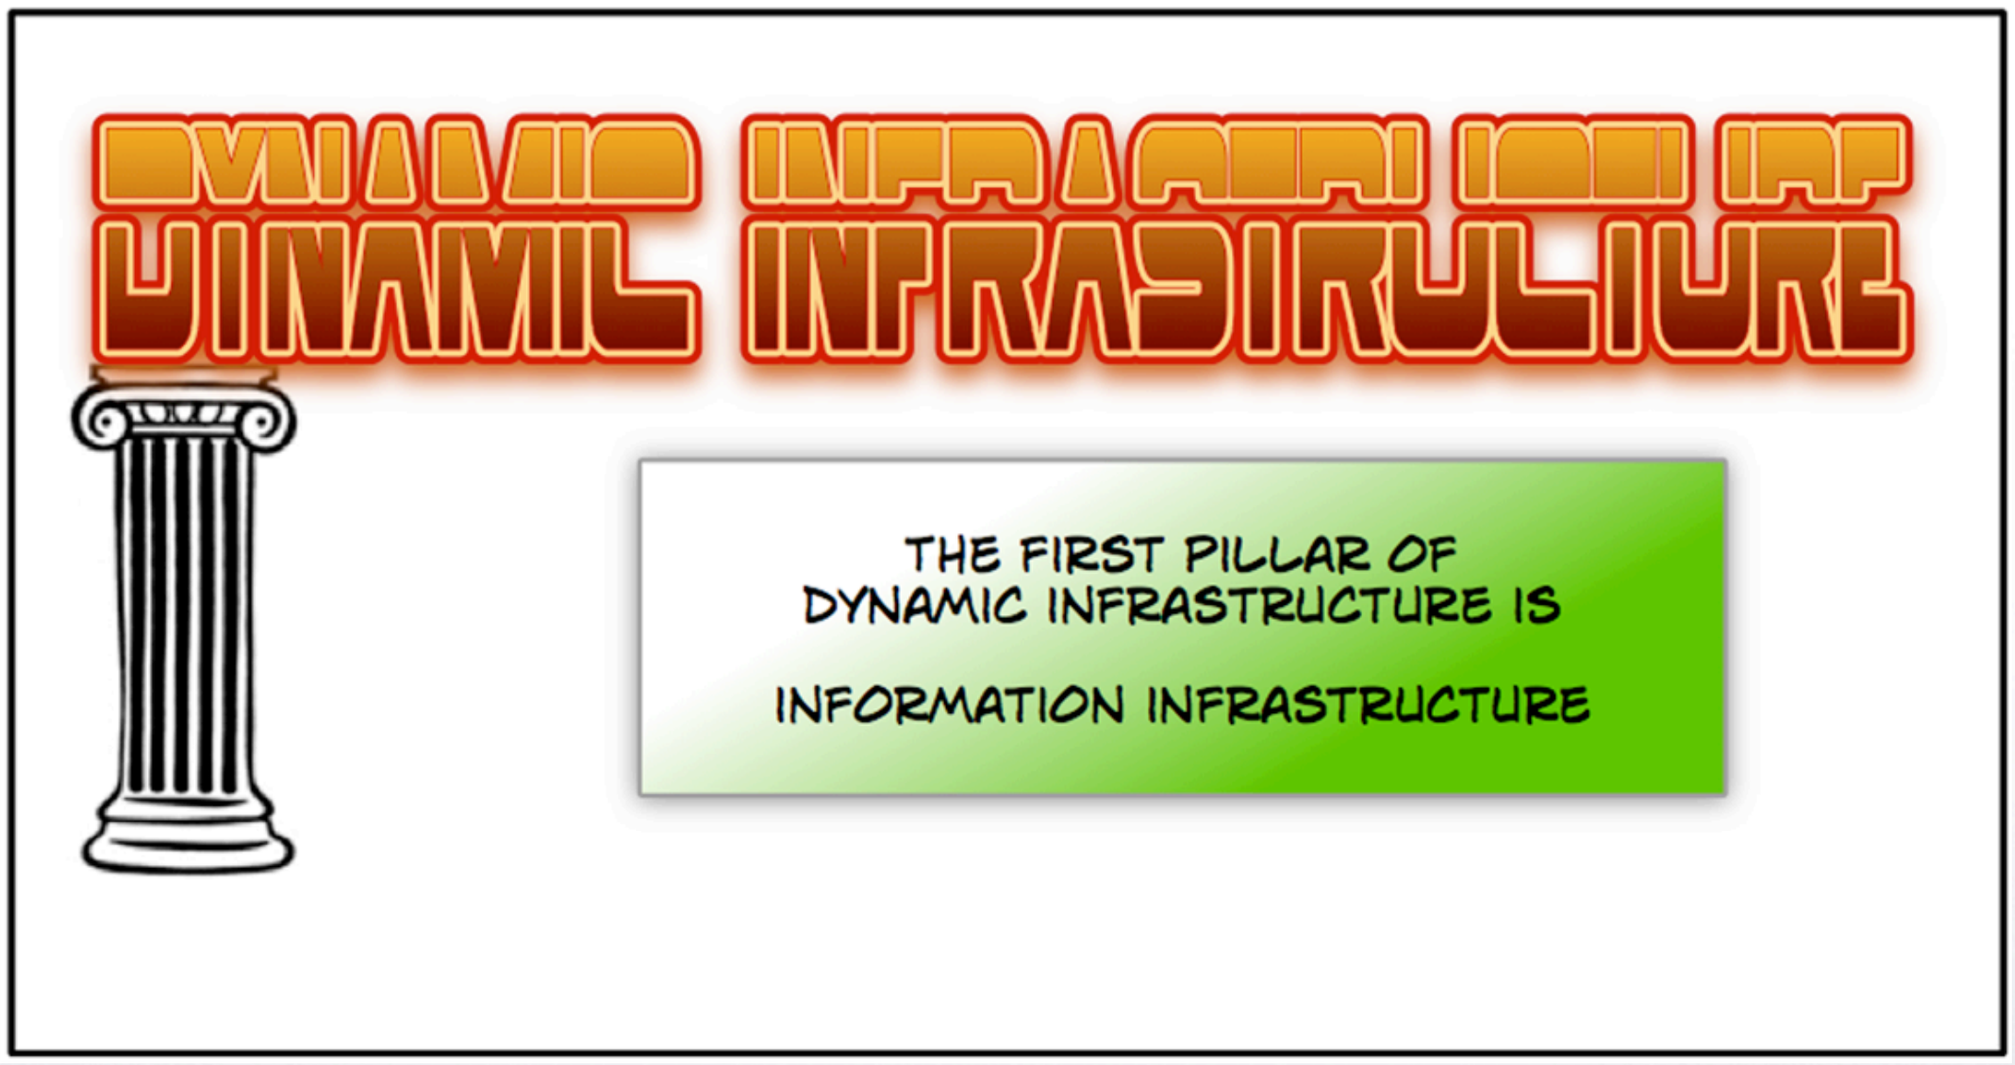

Information Infrastructure

Connecting everything is a cohesive Information Infrastructure which provides data to each requester above it in a timely manner.

This infrastructure must manage the **plethora of data** that is growing in every business.

"Plethora of data". But we're talking about information, right?

Or are we?

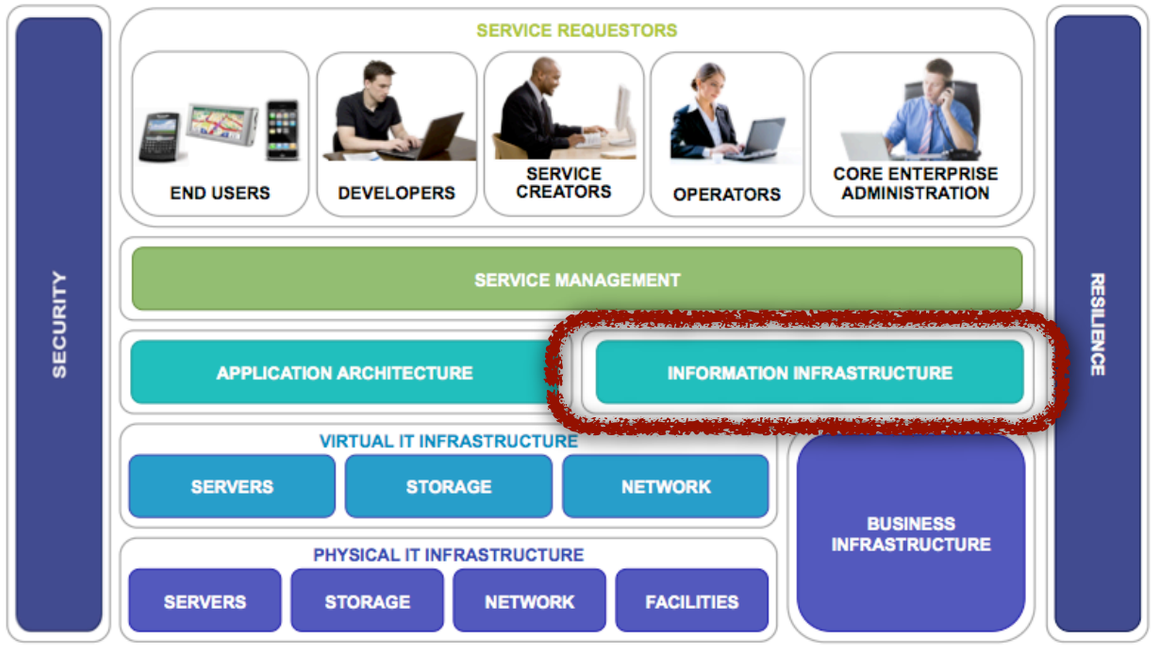

Information Infrastructure

Are we talking about **data** or about **information**?

Is there a difference?

Information Infrastructure

We are talking about turning **data** into **information**.

There is a huge difference.

Data consists of facts and figures, letters, numbers, and points of dubious distinction. That's easy. But few possess information. **Information** is data given context and meaning. Information leads to insight. Piling up data is easy. Gaining insight from the data pile is hard. Keep this in mind as we explore this topic.

Information Infrastructure

Questions to ask yourself about Information Infrastructure

- Do any areas of your business have growing storage needs that seem out of control?
- Do you know the life cycle of your data?
- Are you storing lots of information to comply with regulations?
- Are your users over-provisioning storage?
- Can you prove that the data has not been altered?
- How hard is it to do capacity planning?

You want to have good answers to all of these.

In order to achieve this, we need to consider four (4) areas:

compliance, **availability**, **retention**, and **security**.

Information Infrastructure

Information Infrastructure addresses concerns around four challenges: data **compliance**, **availability**, **retention**, and **security**. Why is this hard? Because there's a lot of data. Like a whole lot. (Rilly.) And we have several data issues:

- Data Volume
  - The codified information base of the world is expected to double every 11 hours. Data is exploding, and its nature is changing to machine-generated – coming from sensors, RFID, meters, GPS systems and more.
- Data Variety
  - With the expansion of data comes large variances in the complexion of the available data – 80% of data is now unstructured, contributed largely by email, documents, images and video.
- Data Risk
  - With more than 30 new compliance regulations worldwide, and more each year, the amount of data that is subject to regulation is growing at 64% per year. More than 60% of IT executives now rate compliance with regulations a top challenge.
- Data Storage
  - Many industries require certain data to be stored for more than 50 years. But on average 37% of a firm's data is inactive or expired. Storing and archiving this data unnecessarily increases business risk, energy consumption and IT costs.

We need to address these data issues in the context of **compliance**, **availability**, **retention**, and **security** in our journey toward a Dynamic Infrastructure.

The Compliance Challenge

- In the current environment, driven by strict regulations, ensuring the availability, retention, and the access to business information needs to be supported reliably and cost-effectively.
- The compliance burden will only continue to grow as we provide more web-based services to a diverse set of users, including employees, customers and business partners.
- The public relations damage that comes with the forced publicity of non-compliance has created the most powerful incentives for stronger information security and compliance. “Stay off the front page of the newspaper” is a common mantra.
- It is important to define what and when to archive information but **it is just as important to define when to purge the information**. Not managing the disposal of information can cause an explosion of retained data and make management and searching the archive difficult.
- Companies are now expressing a pain around the ridiculously high cost of compliance. They know they have to deal with e-discovery and are dealing with it, but per Gartner, most Fortune 1000 companies have 86 cases at any one time and are spending on average $1.5M per case
- Information privacy rules are not the only things that impact information infrastructures. Organizations must also abide by record retention regulations that force them to save information for specified periods of time. Over half of organizations surveyed by ESG* are saving e-mails to comply with such mandates.
- Requests to quickly respond to regulatory agency information management edicts, legal e-discovery demands, and internal audit inquiries are becoming daily occurrences in many organizations. In these types of situations, organizations need to know they can find the information quickly. It is also imperative that they are able to document who accessed and changed that information over its lifetime

* ESG = Enterprise Strategy Group. See acknowledgements.

Information Infrastructure > Challenges

The Availability Challenge

- From managing day-to-day operations to safeguarding intellectual property on which the viability of the business itself may depend, availability remains a cornerstone of our IT lives. Information availability is directly related to IT resource performance, and is one of the highest priorities for assuring agile, responsive IT.
- Today's infrastructure was not designed to efficiently manage the estimated two billion people on web, especially considering that users today are already uploading more than 10 hours of video to YouTube every minute.
- Cost effective, scalable technologies are needed to enable Internet scale and speed for the management of vast amounts of online information – coupled with continuous access.
- Businesses must be prepared for a new phase of cloud computing: giving consumers access to data and systems remotely, from any device, anywhere.
- Data must be available when and where it is needed.
  - Employees rely on it to do their jobs, business processes are designed around it, and now, more than ever, external constituents such as partners and customers need it to execute their own operations.
  - Sixty percent of enterprise organization respondents to ESG research (1000 employees or more) indicated that if their mission critical applications were unavailable for more than four hours, they would experience significant revenue loss or other adverse impacts.

The Retention Challenge

- Today's information infrastructure suffers from massive inefficiency in both duplicate sources of the same data and excessive energy costs.
- Analysts have stated that 50 percent of data centers will run out of power or space for their data centers sometime this year.
- We need retention systems that eliminate duplicate data, coupled with green, cost-efficient archival solutions for long-term retention.
  - On average, companies need to keep content for 7 years, and more for some industries. In healthcare, this is 50+ years and the content is increasing exponentially in HCLS.
  - Clients also want to keep and preserve certain information for cultural heritage reasons.
- How do you future proof your information?
- How do you optimize your infrastructure for retention?
- Companies save copies of information in the event of corruption or data loss. Oftentimes, copies are also created for disaster recovery reasons as well as for test and development, compliance, and many other purposes. In fact, 51% of organizations responding to an ESG survey create between two and ten copies of their primary database instances for the purpose of testing enterprise applications in user environments.
- In data protection scenarios, companies have reasonable data expiration policies. However, copies of data saved for compliance, testing, and other purposes are rarely deleted. A lack of consistent information disposal processes and policies often means that large numbers of files are retained much longer than required by business needs or compliance and audit mandates. The net result is excessive use of costly storage resources.

The Security Challenge

- Ensuring that the data in a data center is secure and being accessed by those authorized has become a top concern for all data centers large and small.
- A recent data center hack cost one company more than $60 million dollars in damages through stealing of data and unauthorized use of credit card information of consumers.
- With several governments and industries having information privacy laws and significant financial penalties, organizations have to worry about making information available to the right people. In the United States alone, 656 data breaches were reported during 2008, a 47% increase over 2007.
- Many organizations find that significant amounts of data are not consistently backed up due to a lack of standardized processes and policies, as well as the use of disparate tools.
- Among enterprise and medium-sized IT decision makers, ESG's research indicates only 28% are complete confident that their organization's information is fully protected.
- Encryption is a common security tactic, but it has it's own issues (which we'll address later) and raises several questions:
  - Can data availability be assured against the effective loss of data due to poorly managed encryption or the loss or damage of decryption keys?
  - Can keys be both adequately secured and adequately available without exposing the entire strategy to risk?
  - Can the cost and performance impact of encryption be reduced or eliminated?
  - Can it assure security without further complicating already complex storage and data center environments?

Information Infrastructure

Compliance

Availability

Retention

Security

That's a lot of challenges.

How about some strategies?

Conquering Compliance

- Volume
  - Archive data.
    - Archiving is the migration of data from online application-accessible storage to an offline long-term storage repository which addresses the following issues:
      - ✓ the “Save Everything” attitude
      - ✓ Excessive storage costs (physical and electronic)
      - ✓ Litigation (underway or impending)
      - ✓ Government audit / investigation (underway or impending)
      - ✓ Excess time to locate materials/records
      - ✓ System Changes (especially to electronic document/records system)
      - ✓ Vital e-mail and Records Protection
  - Pay attention to Information Lifecycle Management (ILM).
    - ILM manages the disposition of data from its creation to its placement on initial online storage to its migration to archive long- term storage to its final deletion.
    - In today’s environments, most organizations leverage tiering of online storage to optimize the efficiency of their environment. The illustration below depicts how ILM manages the process of data mobility from inception to final deletion.

Conquering Compliance

Required Reading

- Variety
  - Handling structured and unstructured data with search systems.
- Risk
  - It's critical that data is archived in accordance with laws and regulations and is stored in a way that allows it to be retrieved effectively should it be needed.
  - Compliance involves, amongst other things, asset management, identity access, storage, server, file systems and database monitoring, and especially change management and configuration management. All of the compliance requirements are addressed by enterprise archive solutions.
  - Archiving is an intelligent process for placing on the right tier inactive or infrequently accessed data that still has value while providing the ability to preserve, search, and retrieve the data during a specified retention period.
- Storage - getting the right data on the right tiers

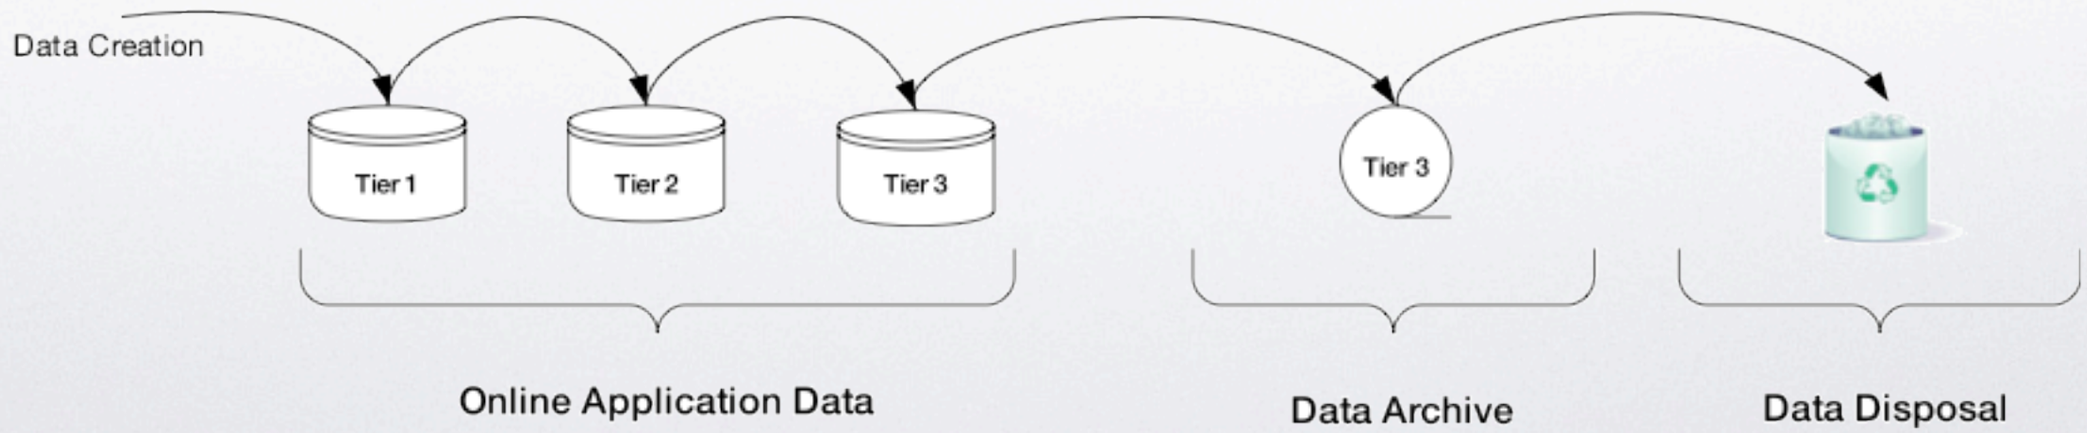

Accomplishing Availability

- Risk
  - Businesses are at risk when their data is no available to them. Data must be available where and when it's needed, even in the face of disaster.
  - Availability usually means that there is little to no *unplanned* downtime. You need planned downtime for routine maintenance, upgrades, etc. It's unplanned downtime that we want to minimize.
  - Disaster Recovery (DR) is one of the key factors in accomplishing data availability.
    - You need to calculate your Recovery Time Objective (RTO) and Recovery Point Objective (RPO)
- Storage
  - Duplicate backup data (replication) is another
    - Logical Replication . . .
      - ✓ maintains a hot backup copy of the environment that is fully available to the business.
      - ✓ uses journal-based data resiliency. Every change to the data is logged, sent to the backup system, and immediately applied there.
    - Physical Replication (sometimes called Hardware High Availability) . . .
      - ✓ maintains an offline backup copy of the environment that is ready and waiting to reconstruct the environment with the latest possible copy of the data.
      - ✓ uses disk-based replication native to the storage systems which employs sector-based replication capability between two or more storage environments.

Information Infrastructure > Strategies

Accomplishing Availability

- Volume - There are several availability solutions, um... available.
  - Tape backup and archiving
    - Old school, reliable, but slow: long RPO and RTOs, like days.
  - Disk backup and practical availability - Backup to secondary disks or partitions
    - More modern, fairly reliable, and a little faster: RPO and RTOs in the range of hours
  - Continuous data protection - Real-time backup and mirroring on active disks
    - Modern, somewhat reliable, fast: RPO and ROTs in minutes and seconds.
  - High Availability - a backup server with a current replica of your application environment is always available
    - Modern, somewhat reliable, fast: RTO of seconds to minutes and an RPO of zero.
- Variety - how do you handle all kinds of structured and unstructured data in different environments?
  - Use one of the above on each platform
  - Search appliances
  - Real-time Database Sharing
    - Goal: multi-way replication among all your platforms: Postgres, DB2, Oracle, Microsoft SQL Server, etc.
    - Strategy: Extraction, integration, Translation, encoding, and Loading (EiTeL).

Required Reading

Required Reading

Required Reading

Resolving Retention

- Volume - We can retain huge volumes of data with a few different storage technologies
  - Direct Attached Storage (DAS) - digital storage directly attached to a server or workstation
  - Storage Area Networks (SAN) - networks designed to connect storage devices to servers.
    - Normally built on a specialized network infrastructure specifically designed to handle storage communications.
    - usually a Fibre Channel network using the SCSI command set or structured as a network using TPC/IP
    - FCP, FC-IP, iFCP, and SAS are common protocols used in a SAN.
    - connections generally include one or more servers (hosts) and one or more disk arrays, tape libraries, or other storage devices.
  - Network Attached Storage (NAS) - a remote server presents its storage to other systems and allows it to be “mounted” or “mapped” to the target server’s existing file system giving the appearance of additional local storage being available.
  - Tape libraries (sometimes called “tape silos” or “tape jukeboxes”) are large storage devices containing tape drives, slots to hold tape cartridges, a barcode reader to identify tape cartridges, and a robot for loading tapes. (Robots are cool.)
    - Magnetic tape, magneto-optical disks, and optical tape storage use many of the same concepts of storage.
    - Optical recording media is used primarily for Write-Once-Read-Maybe (WORM) capabilities.

Resolving Retention

- Risk - We can mitigate some storage risk with RAID: Redundant Arrays of Independent Disks.
  - RAID uses multiple hard drives to appear as a single logical disk, combining multiple hard drives with a parity error correction mechanism to protect data from individual disk failures.
  - Six levels or RAID were originally specified. Today there are well over a dozen different combinations of the original six levels, modifications of the original levels, and vendor proprietary RAID definitions.
  - One or more user-definable RAID levels is at the heart of most modern storage subsystems.
  - The most commonly implemented RAID levels include:
    - RAID 0: Striped without data protection (JBOD – "Just a Bunch of Disks")
    - RAID 1: Mirrored (100% redundancy)
    - RAID 3: Striped data (dedicated parity disk)
    - RAID 5: Striped (parity evenly distributed across disks)
  - Common nested RAID levels:
    - RAID 6: Striped data with parity blocks distributed across two disks
    - RAID 01: A mirrored set of striped disks
    - RAID 10: A striped set of mirrored disks
    - RAID 30: A stripe across dedicated parity RAID systems
    - RAID 100: A stripe of a stripe of mirrors

Solving Security

- Volume, variety, risk, and storage concerns can all be addresses with encryption.
- Encryption can . . .
  - secure data on storage devices when the devices leave the owners' control, such as when devices are repurposed, lost, or returned for an expired lease, repair, or for warranty replacement.
  - protect data from unauthorized or unintentional discovery, lowering risks of exposure while preserving retention requirements and assuring information security and compliance.
- On way to implement this is to encrypt what's stored on your disks.
  - Disk drive-level encryption can be very effective because all data can be encrypted directly in storage media.
  - It can alleviate performance concerns by embedding the encryption engine directly into the individual drives. Even cooler, encryption performance scales linearly as each drive includes its own cryptographic engine.
  - Today's self-encrypting disk and tape drives can deliver security benefits with a high degree of transparency to applications and application servers, as well as to storage systems themselves. They must, however, be accompanied by effective key management policies to deliver on their full potential.
- Speaking of key management . . .

Solving Security

- For successful data security coupled with real-time availability, we must be careful with the keys.
  - Keep encryption keys secure.
    - Encrypted data is only as secure as the keys used to authenticate, encrypt and—more importantly—to decrypt or “unlock” data. The security of these keys is, therefore, paramount.
    - Lost or exposed keys can compromise the security of the data they are supposed to protect. Effective key management must accommodate a variety of risks, use cases, application security, and related IT management systems. Assuring secure and reliable key management is one of the most critical factors that can make or break a successful deployment.
  - Keep encryption keys available.
    - Key availability is essential to the availability of the data itself. High availability requires redundant copies of critical information, and while cryptographic keys are indeed critical information, precautions must be taken to ensure backup copies are also secure to mitigate risk to any copies.
    - The availability of keys in real-time operations is even more critical. One might think that this is a typical if not critical requirement of key management, but you may be surprised to learn that serving keys to the appropriate systems, applications and services is not always an integral feature of a key management products. We need to assure that these capabilities are part of real-time key management system.

Information Infrastructure

In summary, you build the frame of strong Information Infrastructure with the CARS approach:

Conquering compliance

Accomplishing availability

Resolving retention

Solving security

This is the engine of Dynamic Infrastructure, and if you do it right, your business can react to changes in the marketplace with speed, faster than fast, quicker than quick. Maybe even . . . *Lightning*.

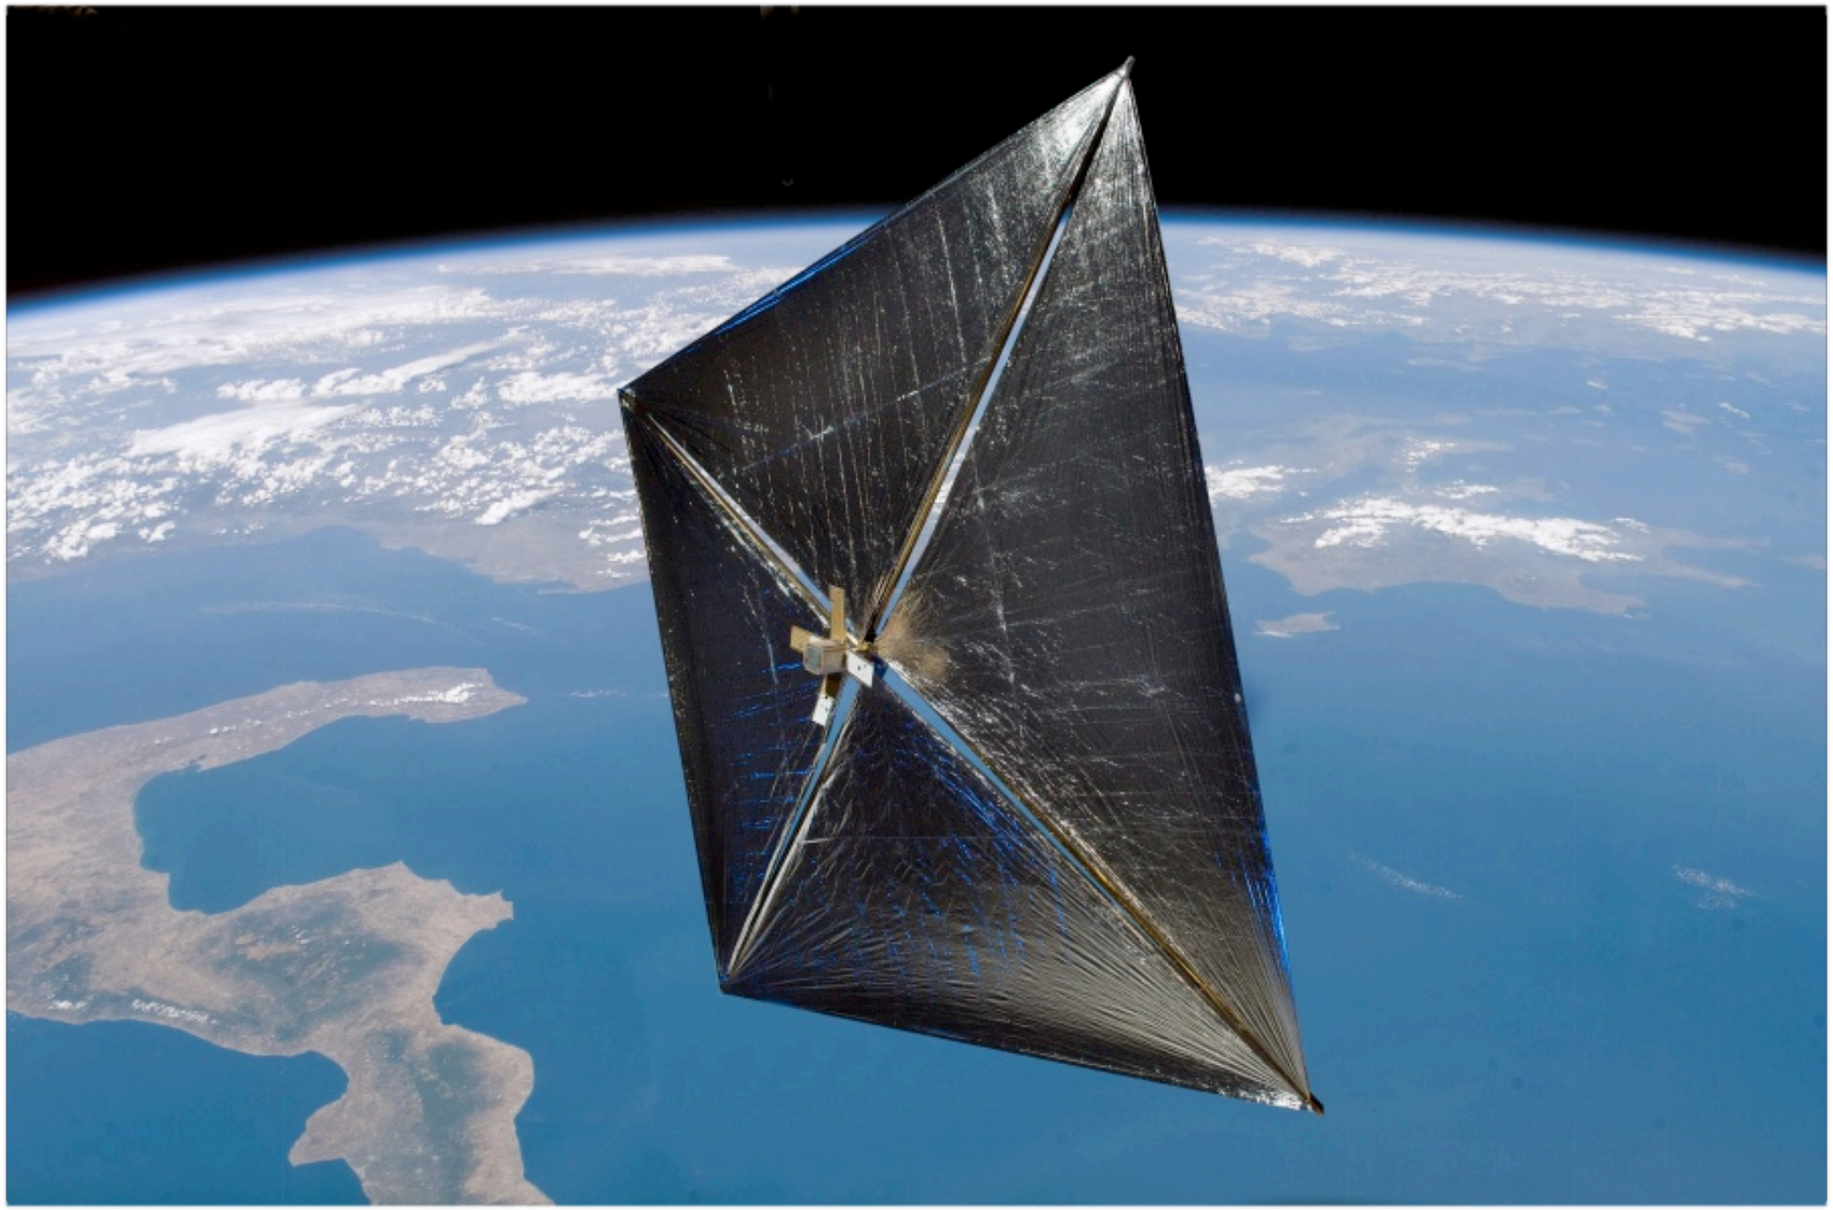

Illustration Credit: courtesy NASA

Required Readings and Videos

Websites

- RAID in Linux - http://www.ibm.com/developerworks/linux/library/l-raid1/index.html
- More on DR - http://findarticles.com/p/articles/mi_m0BRZ/is_8_23/ai_109665173/
- The Hacker Threat - http://www.bloomberg.com/apps/news?pid=newsarchive&sid=a8xQO5KRZil0

Papers (already linked from the prior slides)

- Taming the Information Explosion [*pdf*]
- Data Archiving for Compliance and Cost [*pdf*]
- High Availability [*pdf*]
- Essential Guide to Disaster Recovery [*pdf*]

Videos

- Data Security at the US Census Bureau.
- CIO Pain points: Improving data security for mobile devices.

Optional Readings

Papers

- RAID - the early history [*pdf*]
- Real-time Database Sharing [*pdf*]
- EMC-storage-management [*pdf*]
- Security for Data at Rest [*pdf*]
- Availability Management- http://www.itilnews.com/availability_management.html
- Inside Intel's Security Organization http://www.networkworld.com/news/2010/091610-intel-security.html
- Information System Security Management in the new millennium- http://tols17.oulu.fi/~jhyvonen/Tietoturvan%20hallinta%20tenttimateriaalit/Luento1.pdf

Self-test

- What are the four (4) categories or challenges for Information Architecture?
- What are several data issues that need to be addressed with these for challenges?
- Why do you need to plan for downtime?
- What is the difference between data and information?
- What does the ILM (Information Lifecycle Management) do?
- Name a key factor in data availability?
- What two things should you calculate in disaster recovery?
- What are goals and strategies of Real-time Database Sharing?

Self-test Answers

- What are the four (4) categories or challenges for Information Architecture?
  - Compliance, Availability, Retention, and Security
- What are several data issues that need to be addressed with these for challenges?
  - Data volume, data variety, data risk, and data storage.
- Why do you need to plan for downtime?
  - Planned downtime is essential for routine maintenance, upgrades, etc.
- What is the difference between data and information?
  - Data are facts numbers and figures where information is that data given context and meaning.
- What does the ILM (Information Lifecycle Management) do?
  - It manages the disposition of data from its creation to its placement onto online storage then to its migration to archive long-term storage to its final deletion.
- Name a key factor in data availability?
  - Disaster Recovery
- What two things should you calculate in disaster recovery?
  - Recovery Time Objective and Recovery Point Objective.
- What are goals and strategies of Real-time Database Sharing?
  - The goal is to have multi-way replication amount all platforms. The strategy is extraction, integration, translation, encoding and loading.

Discussions

1. Going Green: How do we tell if it's real or just "green washing"?
   - What's "green washing"?
   - Why do people do it?
   - How can we tell when it's going on?
2. Information Security vs. Availability
   - Define both.
   - What are the tradeoffs?
3. Did you notice how many Dynamic Infrastructure concepts there were in that "CIO Pain Points" video?
   - Discuss.

Remember our discussion expectations and guidelines.

Acknowledgements

Some of the source material and a few of the graphics in this module came from the IBM World Wide Client Technology Centers's very own Frank De Gilio.

Some additional source material was provided by David Graves and Paul Kontogiorgis of IBM in 2006.

More additional material from:

- IBM's "Data Archiving: Foundation Capabilities for Compliance and Cost Optimization" April 2009
- Enterprise Strategy Group's white paper: "IBM Information Infrastructure Initiative Tames the Information Explosion" by Brian Babineau, 2009.
- "Security for Data at Rest" An ENTERPRISE MANAGEMENT ASSOCIATES® (EMATM) White Paper Prepared for IBM in September 2008
- "High Availability for IBM i" prepared for IBM by Vision Solutions in Irvine, CA in 2010
- "Essential Guide to Disaster Recovery" prepared for IBM by Vision Solutions in Irvine, CA in 2010

Acknowledgements

Still more additional material from:

- iStockphoto.com
- Elsewhere and else**when**, but only when the stars are right.

Alan thanks his diligent (and patient) student Carley Keefe for her contributions to this work.

Colophon

This work was authored in Keynote by Alan G. Labouseur in June 2010 from his home in Pleasant Valley, NY.

This is a photo from Beihai park in Beijing, one of Alan's favorite places in China, or the world.

Distractions that made writing go slower:

- The This Week in Venture Capital podcast
- "Medium Raw" by Anthony Bourdain
- Stackoverflow.com

Music that made writing go faster:

- iTunes Genius Mixes: Blues, Funk
- Specific artists: Marcus Miller / Stanley Clarke / Aerosmith / Steely Dan / Lionel Hampton / Stu Hamm