



# Intelligent Network Operations and Management – It's about the Data

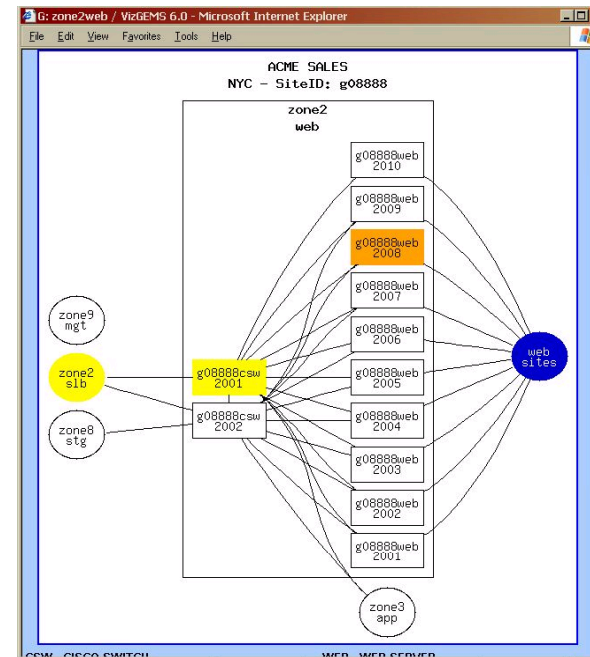
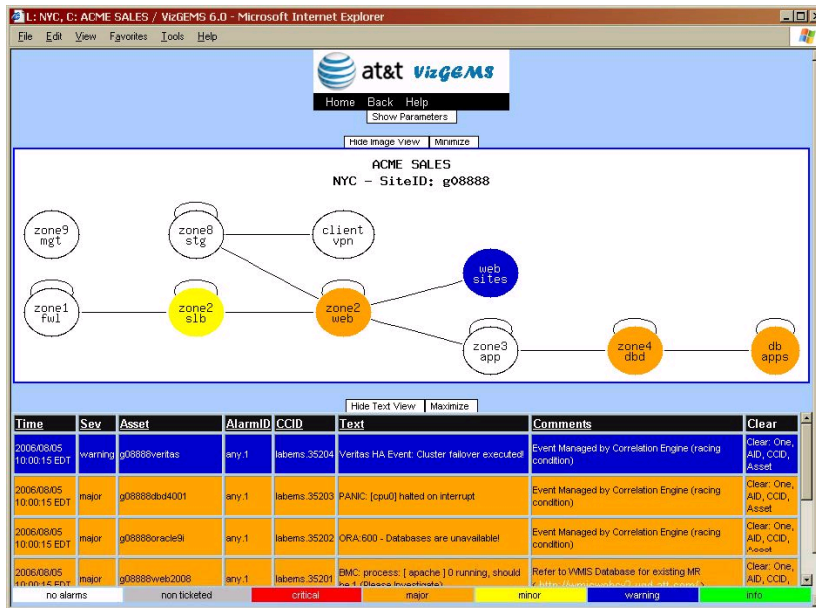
Keynote Talk: June 2011  
IEEE ISCC 2011, Corfu, Greece

**Mahmoud Daneshmand, PhD**  
**AT&T Labs Research**



# Operations – Give me all your data

Operations is increasingly about End-to-End control across multiple networks/services; across multiple layers in network, computing, and software stacks; and, across a variety of time-frames. It is, therefore, a problem of integration of huge amounts of very heterogeneous data in real time.



# **How Much Information? 2009**

## **Report on American Consumers**

*Roger E. Bohn & James E. Short*

**Global Information Industry Center, UCSD**

### **Industry–University Collaborations**

**AT&T**

**Cisco Systems**

**IBM**

**Intel Corporations**

**LSI**

**Oracle**

**Seagate Technology**



# Scale – Still the Heart of the Challenge



## The Answer:

**3,600,000,000,000,000,000,000**

Scale includes volume, volatility, complexity, reliability, and security



# How Much Information? 2009

- The goal of HMI? Project is to create a census of the world's data and information
- First Report: Information at the US Consumer Level
- How Much Information was Consumed by Individuals in the U.S. in 2008?

*3.6 Zettabytes (ZB =  $10^{21}$  bytes)*

- *How much is 3.6 Zettabytes?* If we printed 3.6 zettabytes of text in books, and stacked them as tightly as possible across the United States including Alaska, the pile would be 7 feet high



# How Much Information? 2009

- Measures of information include all data delivered to people, whether for personal consumption, for communication or for any other reason: **Cable TV, Broadcast TV, Radio, Telephone Line, Internet, Wireless, etc.**

## *Network Data*

### *The AT&T Network*

- *Currently carries 18.7 Petabytes of data traffic on an average business day (PB =  $10^{15}$  bytes)*
- **Nearly 5 Billion calls per day**

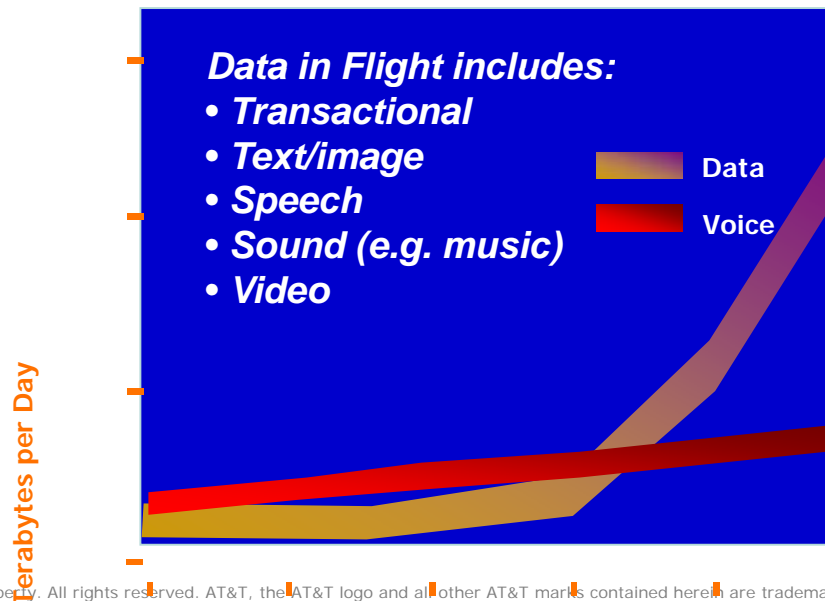
## *Information in Flight*



# Information in Flight





Existing delivered information is estimated to be O(Zettabytes) per year. Most of the new information, and much of the existing is, in any year, on a network – i.e. in flight. Video is responsible for the enormous growth in total amount of information in flight.



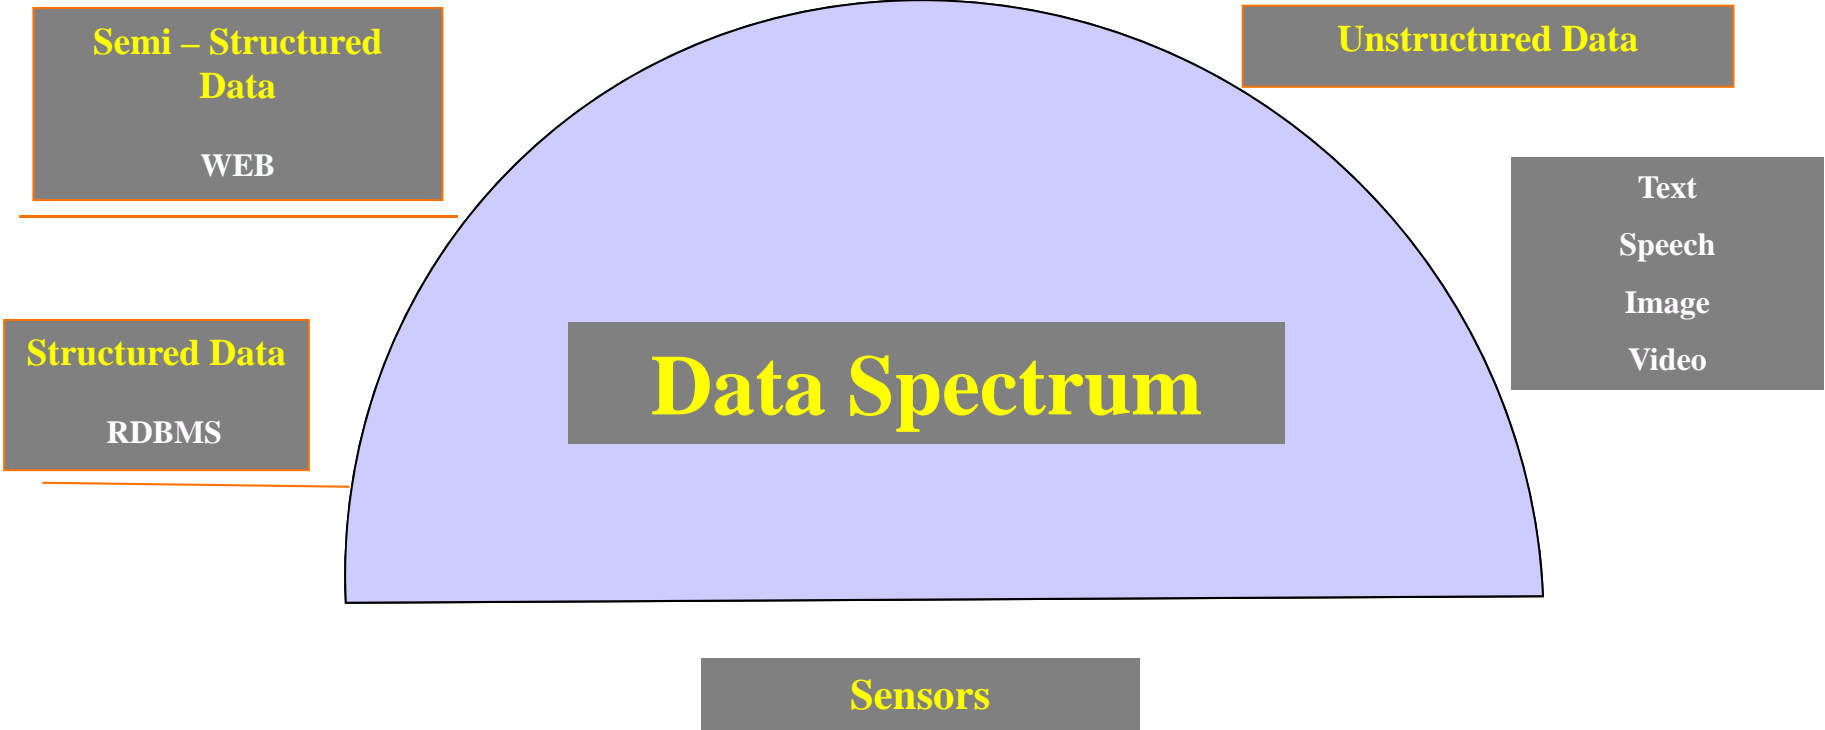
# Operations – Give me all your data

Operations is increasingly about End-to-End control across multiple networks/services; across multiple layers in network, computing, and software stacks; and, across a variety of time-frames. It is, therefore, a problem of integration of huge amounts of very heterogeneous data in real time.





# AT&T InfoLab



# Operations Today

## Operations



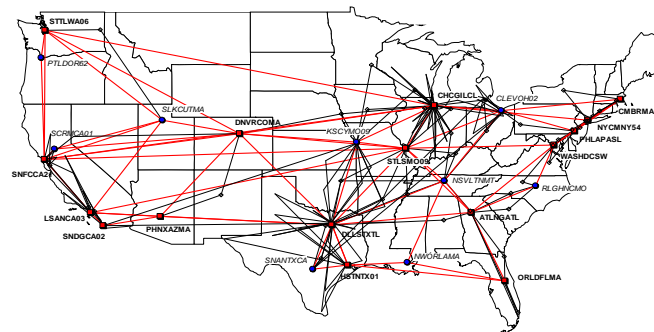
## Services



## Applications



## Network



# Network Scale and Scope

**3,824**

MPLS nodes serving 163 countries

**928K**

Route Miles of fiber

**9152**

Major network buildings;  
211K small locations

**125K+**

Wi-Fi hotspots

**50K+**

Cell sites covering >70 countries

**38**

Internet Data Centers on 4 continents

**24x7**

Network monitoring & management

**100%**

Fortune 1000 companies are AT&T customers

**87M**

Wireless customers

**17.5M**

Consumer Broadband connections

**2.3M**

IPTV customers in >40 markets

**20+**

Petabytes of data per day

**1.6B**


Text messages/day

***“Setting the industry standard for network reliability”***

© 2010 AT&T Intellectual Property. All rights reserved. AT&T, the AT&T logo and all other AT&T marks contained herein are trademarks of AT&T Intellectual Property and/or AT&T affiliated companies. All other marks contained herein are the property of their respective owners.



# Operations Cycle



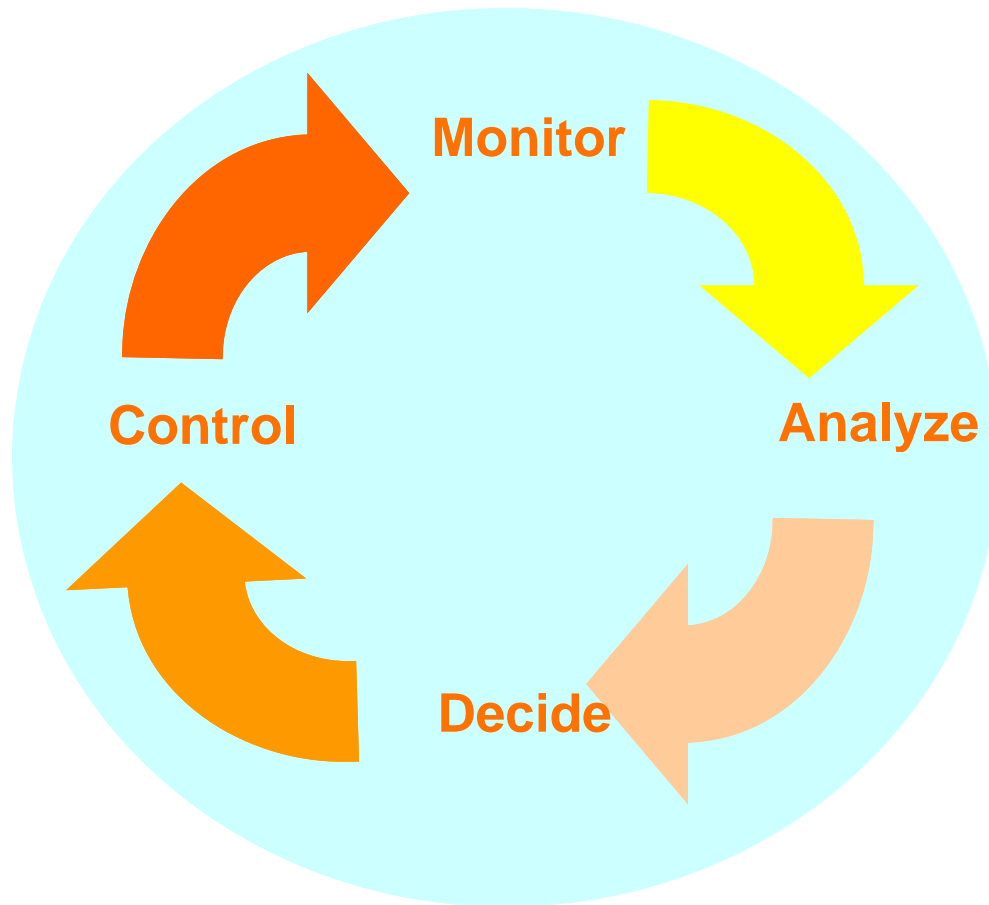
## Each phase can be complex in large networked systems

- Monitoring involves data across multiple hosts and multiple sources.
- Analyzing may involve heuristics or evaluation over time.
- Decision may involve evaluating tradeoffs or distributed algorithms.
- Control may involve distributed coordination across multiple hosts.

**All done in a running system and an environment that continues to change.**



# Operations Cycle



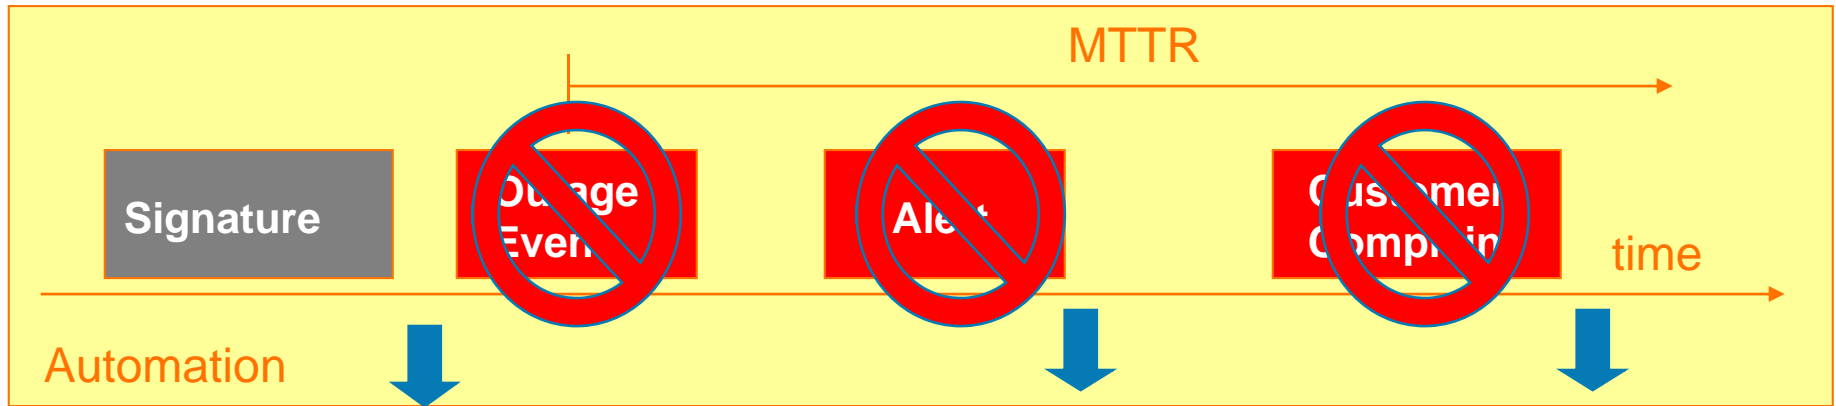
This often involves huge amounts, and many types, of data:

- **Examples may include:**
  - **Circuit Switched Networks – SS7, CDR, ...**
  - **IP Networks - Netflow, SNMP, ...**
  - **Mobile Networks – UE, Physical & Location, RNC/MS, ...**
  - **Systems – Logs, utilization, virtualization, ...**
  - **Sensors – SmartGrid, TeleHealth, RFID, ...**

**Data flows are very widely distributed, and often measure in the Giga/Tera/Peta Bytes per day.**



# Network Based Computing: Goal - Adaptive Maintenance



**Predictive maintenance**

- Eliminates MOOs and outages
- Further reduces costs, including overtime pay
- Tradeoff of maintenance and outage costs
- Delighted customers

**Proactive maintenance**

- Reduces MOOs and MTTR
- Reduces operational costs
- Customer care costs reduced
- Automated customer status
- Happier customers

**Reactive maintenance**

- Highest MOOs
- SLA risk
- High care costs
- Disappointed customers

**Adaptive maintenance**  
§ Reduces false positives and unnecessary costs



# What's Changed

- |   |  |
|---|--|
| <b>1. Volumes</b>                         | <b>Larger Volumes</b>                  |
| <b>2. Network per Service</b>             | <b>Multiple Services over IP</b>       |
| <b>3. Vertical Ecosystem</b>              | <b>Horizontal Ecosystem</b>            |
| <b>4. Sarnoff/Metcalf Net Value Laws</b>  | <b>Reed's Law, Col, Social N.</b>      |
| <b>5. Dial Access</b>                     | <b>BB &amp; Wireless BB Access</b>     |
| <b>6. VoIP, VXML</b>                      | <b>SIP Services, IMS Services, WSI</b> |
| <b>7. Video Broadcast</b>                 | <b>Video Multicast, IPTV</b>           |
| <b>8. Fraud</b>                           | <b>Security/Privacy</b>                |
| <b>9. Structured/OO Programming</b>       | <b>Web Services</b>                    |
| <b>10. Structured, Transactional Data</b> | <b>Data Streams, Spectrum of Data</b>  |
| <b>11. Research Directions:</b>           |  |



# Scale

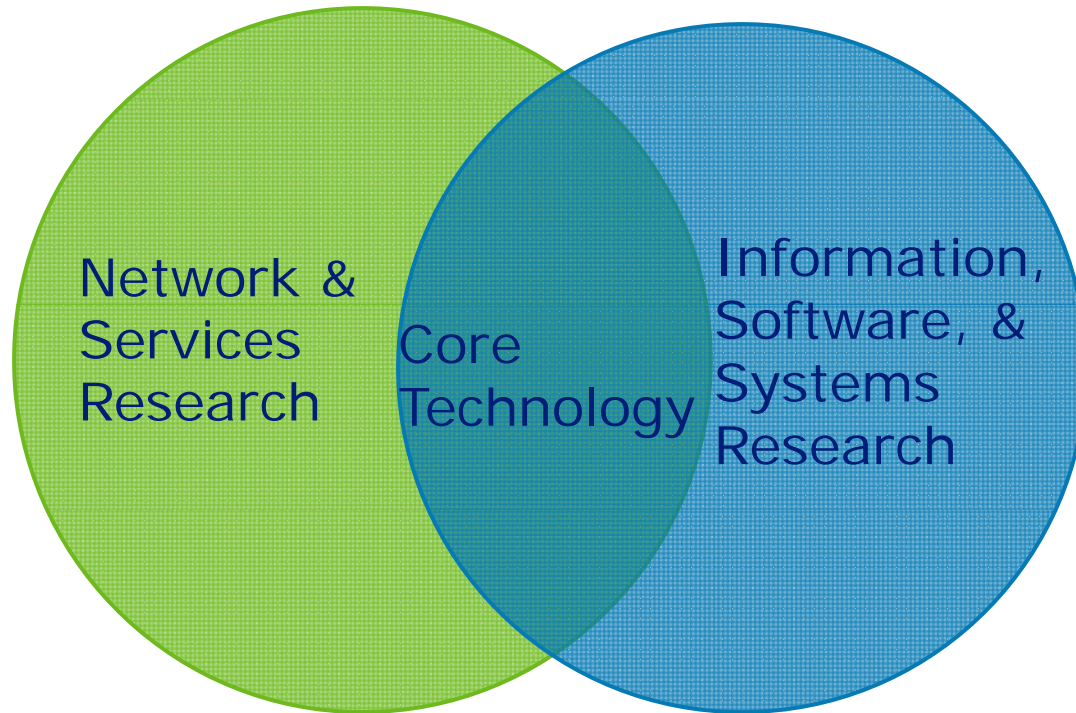
## What makes it challenging!

- **Lots of records**
  - 5.22 Billion records per day**
  - 3 Trillion events in the DB**
  - 2 PB of data overall**
- **Lots of calling/called parties**
  - 400 Million unique users per month**
- **Lots of diversity in calling behavior**
  - Every call, no matter how strange it may seem, can be perfectly legitimate**





# AT&T Labs - Research



**Research**

**Impact**

**Influence**



# Information, Software, & Systems Research

Information/ Statistics / ML / Data Mining Research

Software Research

AI

Scale

Compression

Tools, Libraries, & Algorithms

Security

Specification & Test

Systems Research

Mobile/Nomadic/AdHoc

Distributed, Network Based

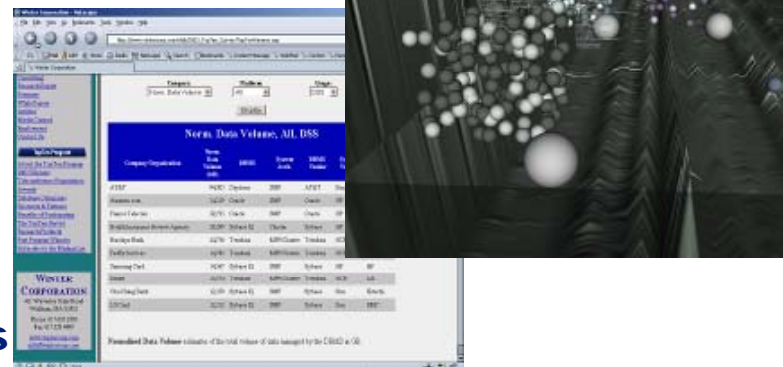

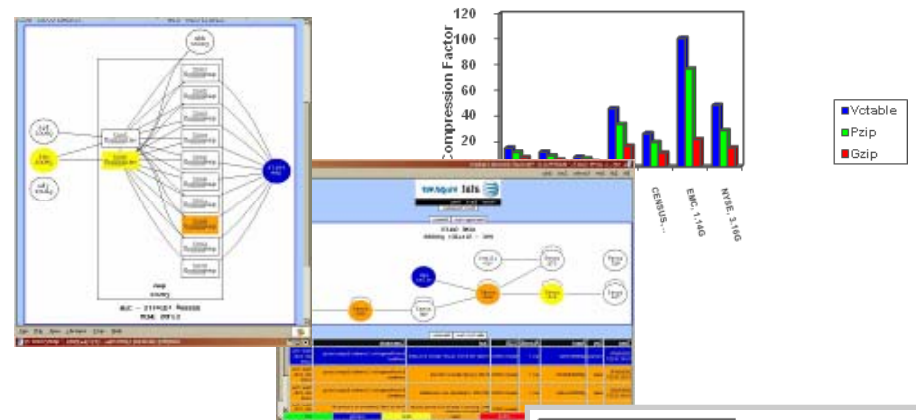
Data Base, Data Streams

Visualization & Data Integration

Monitor, Analyze, Control, Adapt


Information- Based Services

Incubation & Professional Services



# Data Lifecycle

*Can we completely automate the staging of data for analysis (From origination to analysis ready)?*



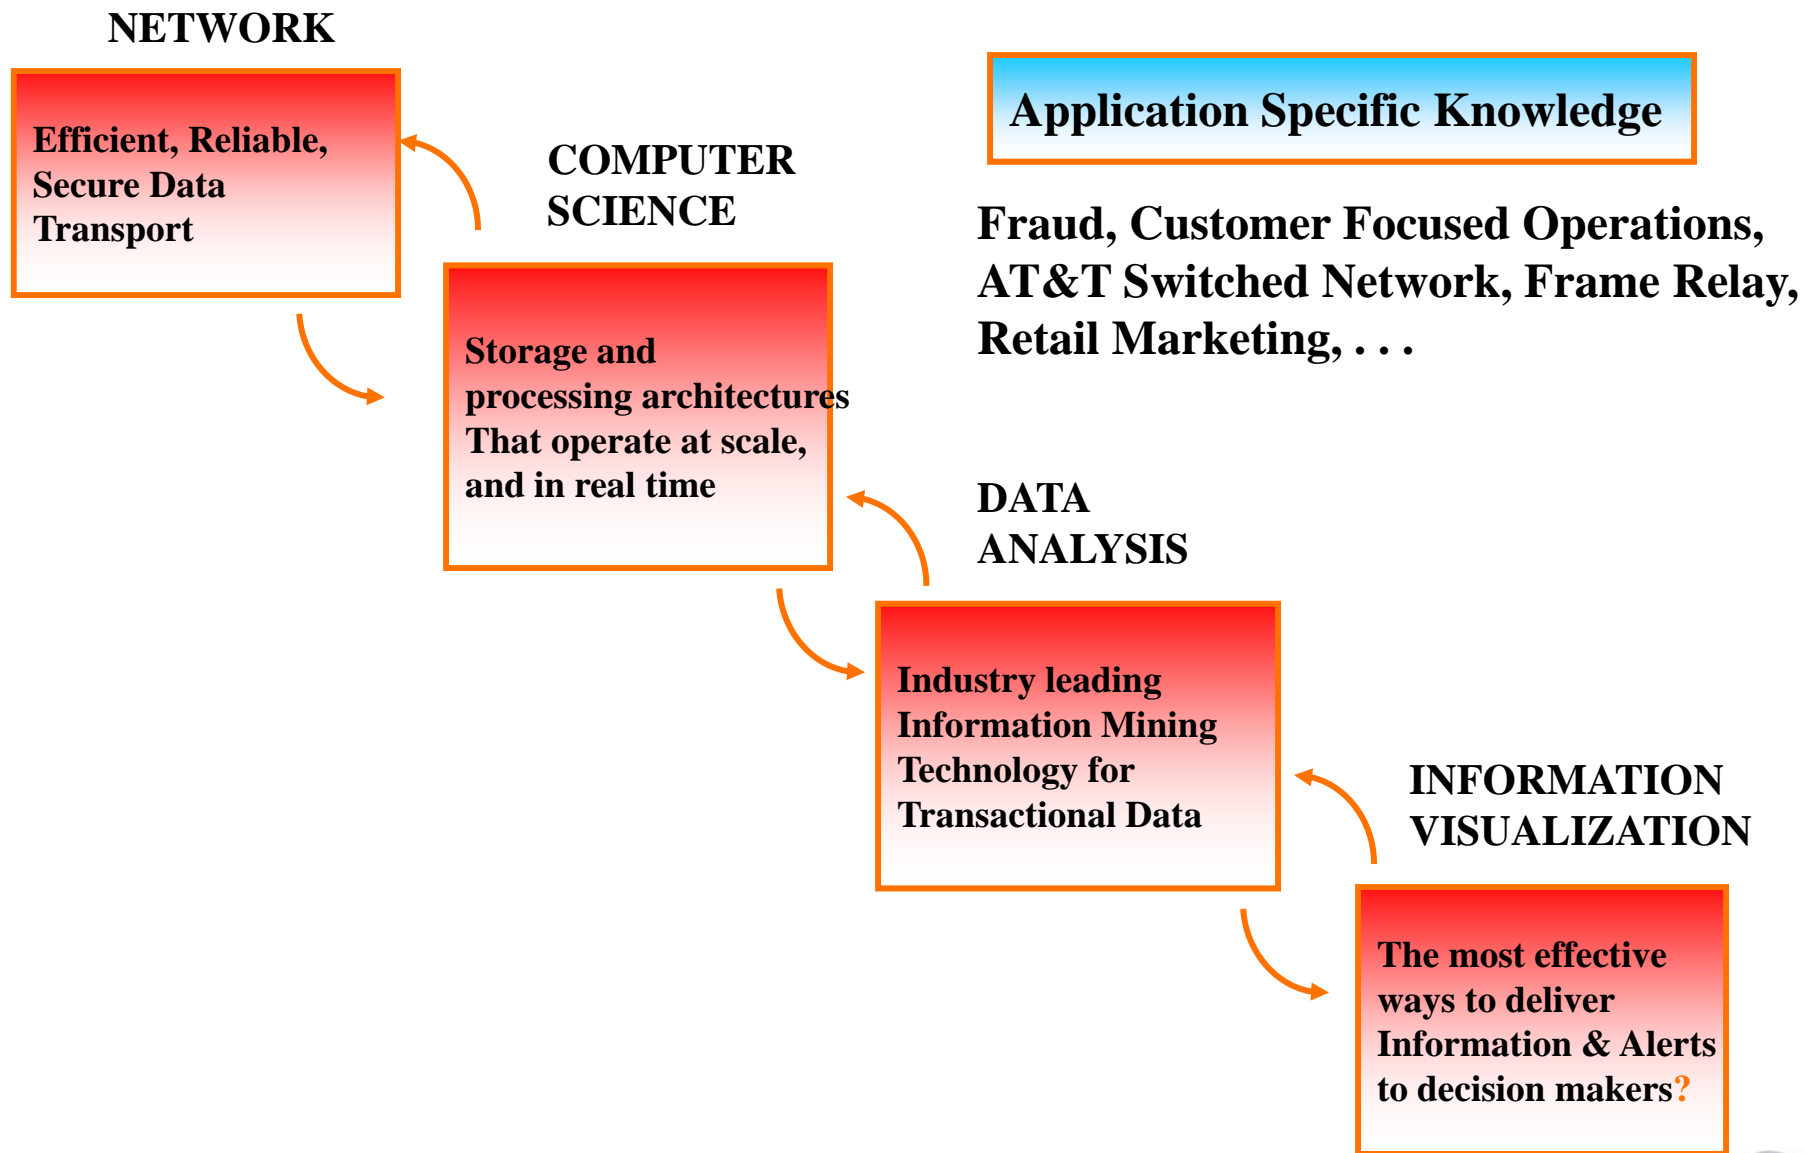
# Information & Software Systems Research

## *A few very high level questions*

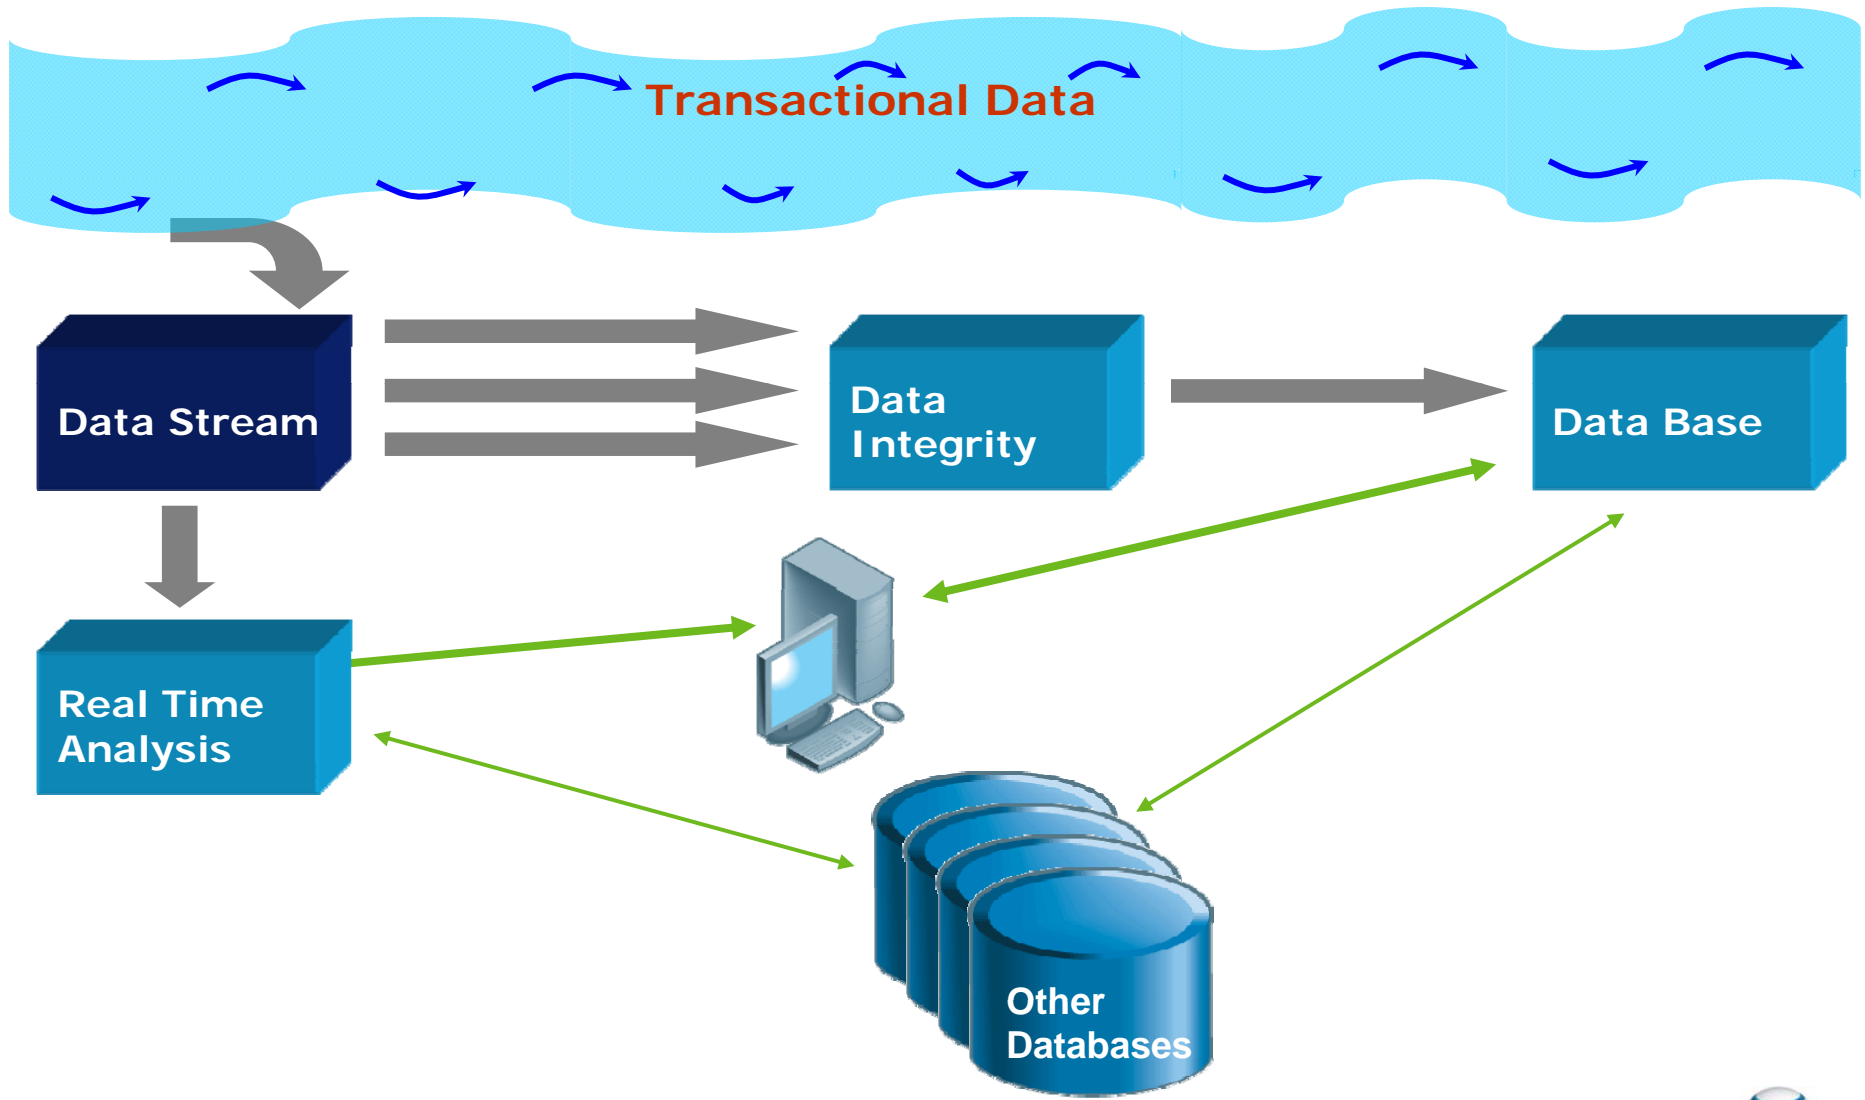
- *Can we make the End to End, Top to Bottom, operations completely transparent?*
- *What happens when we have direct access to all videos ever made – Anywhere, Anytime, Anyway?*
- *Can we completely automate the staging of data for analysis (From origination to analysis ready)?*
- *Can we predict actions of our customers based on behavior (e.g. churn, retention)?*
- *What will the next great analytic computing environment look like?*
- *Tools: What systems tasks can we make disappear?*
- *Can we run a global IP network in the dark?*
- *How do we best design, build, test, & protect systems as distribution & mobility become integral parts of software/data systems?*
- *How should we think about Web and Data Privacy, Security?*



# AT&T InfoLab Approach



# Large Scale Data Management

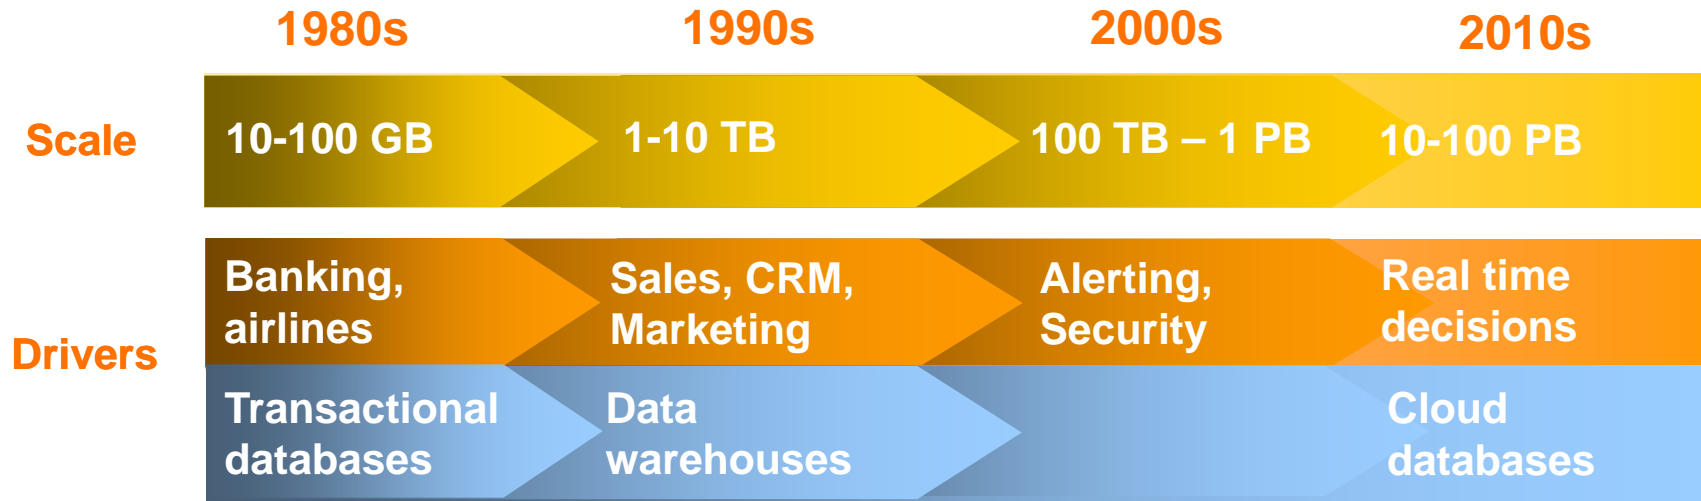


# Challenges in Data Management @ Scale

- **Scale:** billions of records generated daily
- **Analysis:** desire to support (near) real-time queries and analyses
- **Access:** rapidly integrate multiple databases
- **Quality:** pinpoint potential data quality issues
- **Security:** authorized users, authorized uses



# Evolution of Data Management: Scale





# Daytona: Managing Data at AT&T Scale

Norm. Data Volume, Unix, DW							
Company/Organization	Norm. Data Volume (GB)	DBMS	Platform	Architecture	DBMS Vendor	System Vendor	Storage Vendor
AT&T	330,644	Daytona	UNIX	Federated/SMP	AT&T	HP	HP
AT&T	93,468	Daytona	UNIX	Federated/SMP	AT&T	Sun	Sun
Nielsen Media Research	17,969	Sybase IQ	UNIX	Centralized/SMP	Sybase	Sun	EMC
Yahoo!	17,014	Oracle	UNIX	Centralized/SMP	Oracle	Fujitsu Siemens	EMC
UBS AG	14,177	Oracle	UNIX	Centralized/SMP	Oracle	Sun	EMC
China Telecom Corporation Co.,Ltd. GuangZhou Research Institute	13,241	Sybase IQ	UNIX	Centralized/SMP	Sybase	Sun	Sun
Reliance Infocomm Ltd	11,500	Oracle	UNIX	Centralized/SMP	Oracle	Sun	EMC
Cellcom	10,345	Oracle RAC	UNIX	Centralized/Cluster	Oracle	HP	EMC
Turkcell	9,504	Oracle	UNIX	Centralized/SMP	Oracle	Sun	Hitachi
JPMorganChase	8,875	DB2	UNIX	Centralized/MPP	IBM	IBM	IBM

Copyright 2005 Winter Corporation

- **Massive amounts of data can be collected, but hard to manage in commercial DBs**
- **Daytona enables scalable data management**
  - organizes and stores ~1 PB of data + indices on disk, with a data dictionary
  - uses compression, horizontal partitioning
  - enables concise, natural expression of sophisticated queries
  - provides answers to those queries quickly
  - manages data in concurrent environment
  - has proven reliability
- **Sample applications across AT&T**
  - call detail (largest warehouse)
  - STORM/FLOOD: network security
  - Darkstar/TAS: IP data analysis

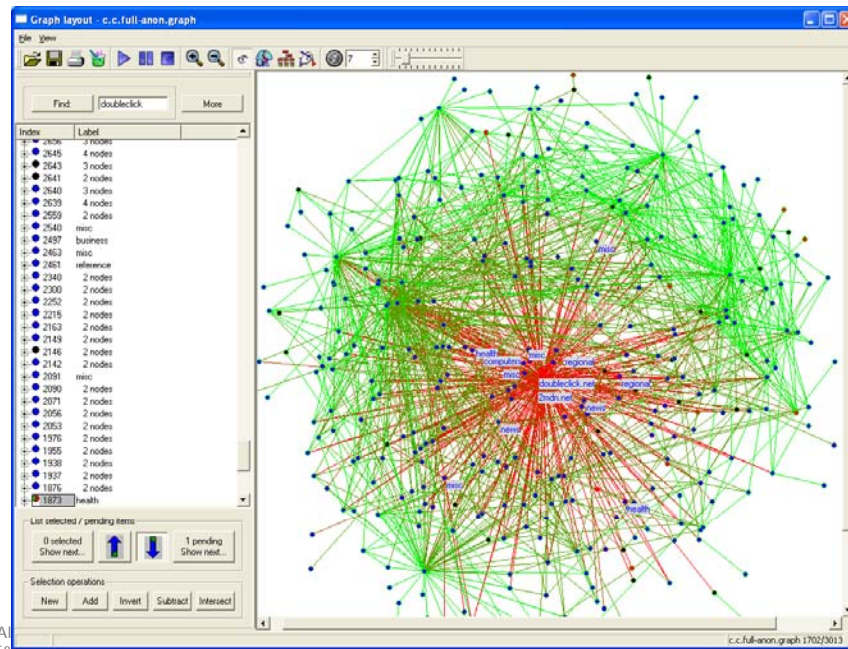


# Data: Movement, Management, Analysis, Research

## Visualization

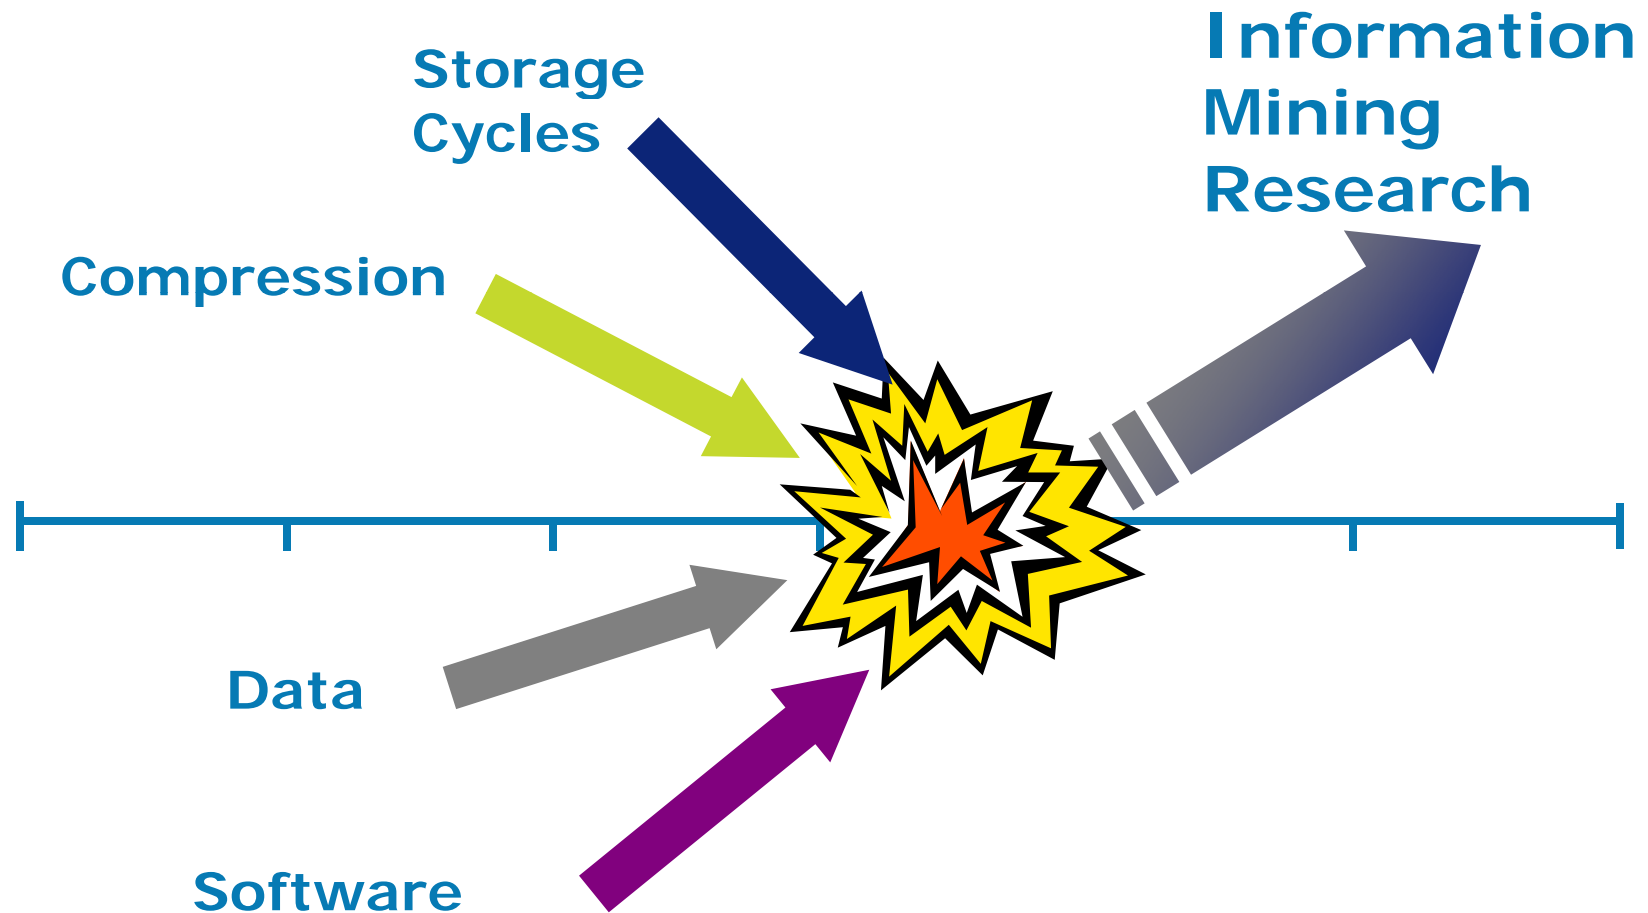
The scale and power of data related technology continues to increase dramatically. Volume and complexity that would have seemed impossible just a few years ago is now somewhat routine.

### Analysis/Visualization on a Large Graph

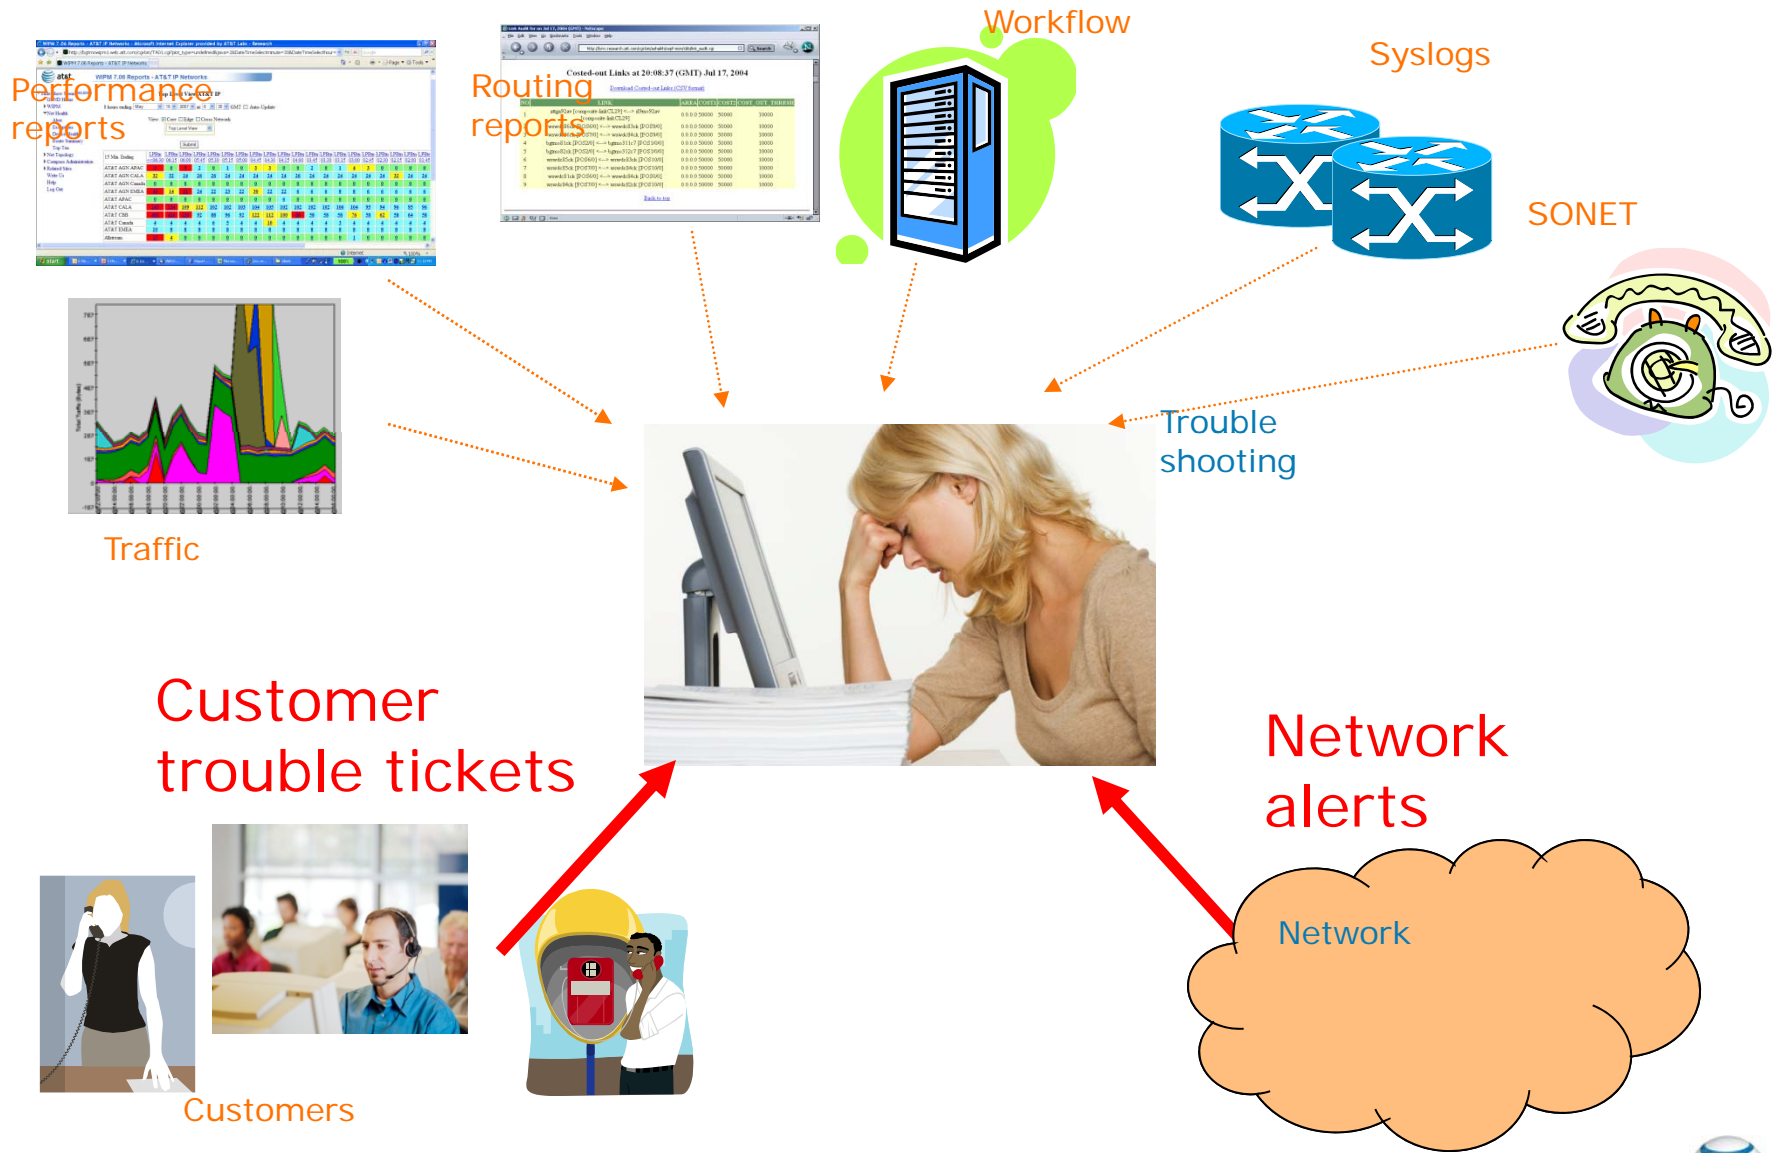


# DRIVING INNOVATION

## - Technology Drivers and Trends



# Operations Challenge: Troubleshooting

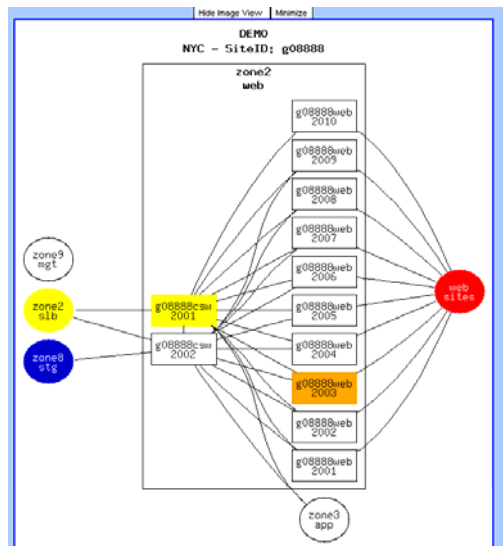


# Decide – Analysis & Visualization



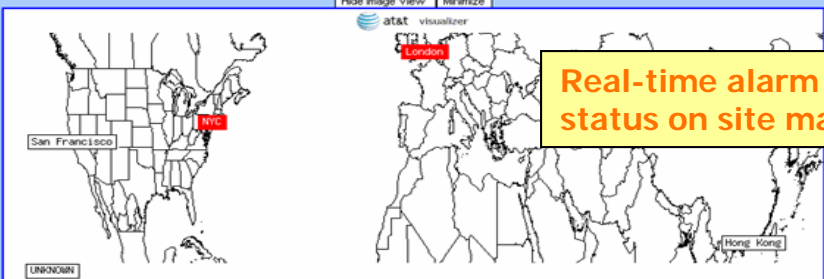
The amount and variety of data are far too great for all but the most powerful tools. Robust, automated analysis is essential to control in real time. Humans are still the best visual pattern recognition engines available.

## An example of a small system



# Visualizer

## Managing Systems Distributed Worldwide



**Real-time alarm status on site map**

ABC Company  
NYC Area, NY - SiteID: 009999

WWW 00 01 02 03 06

RTR 81

IDC srvo

FwX 1a 1b

DBD 00 01 02

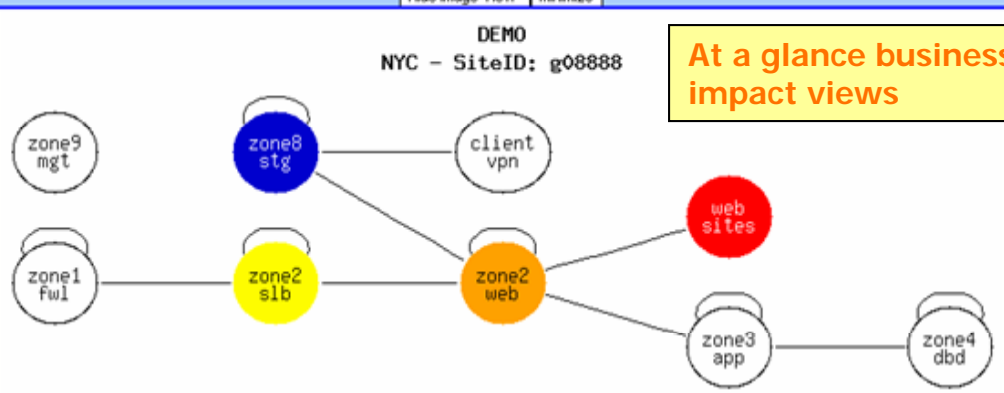
CSW 01 02

BI 41nyc 42nyc

ASW 01 02

**High level equipment inventory**

Time	Severity	Object	CCID	Msg Text	Comments
2006-05-16 09:00:02 EDT	minor	g0888slb2001	ldchev1.17375	PORTMON: [ port 80 ] [ host=g0888web2003 ] [ failcount=>3 ] [ portscan failed ]	Event Managed by Correlation Engine (racing condition)
2006-05-16 09:00:02 EDT	minor	g0888csw2001	ldchev1.17376	SYSLOG: %LINK-3-DOWN: Link on interface GigabitEthernet4/11, changed state to down	Event Managed by Correlation Engine (racing condition)
2006-05-16 09:00:02 EDT	major	g0888web2003	ldchev1.17377	BMC process: [ apache ] 0 running, should be 1 (Please investigate).	Refer to VMS Database for existing MR ( http://www.webbox2.ugd.all.com/ )
2006-05-16 09:00:02 EDT	critical	w0888app3002	ldchev1.17378	CLERT: perl[cpu0]/Thread=0x502c7ba0 [AFT1] errD 0x00179024.ca80da7 UE Error(s)	Event Managed by Correlation Engine (racing condition)

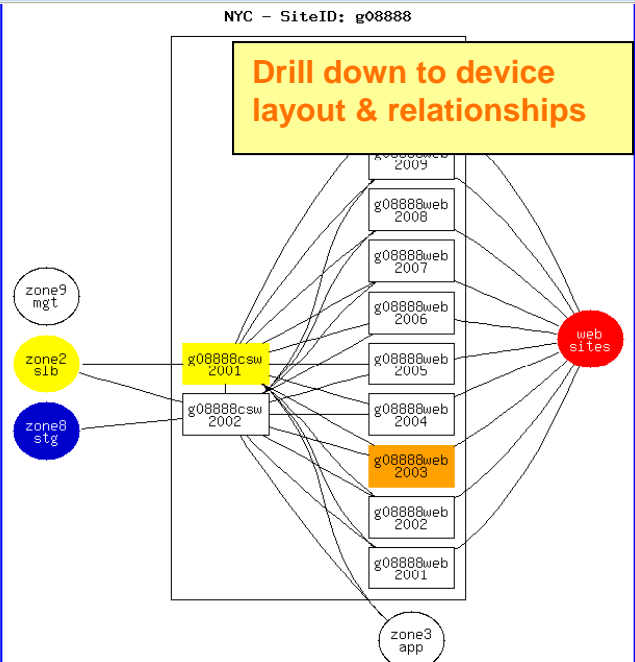


**At a glance business impact views**

ABC Company  
NYC Area, NY - SiteID: 009999

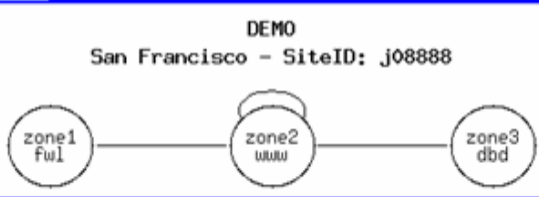
no alarms non-ticketed normal warning minor major critical

ASW - LOAD BALANCER  
CSW - CISCO SWITCH  
DBD - DATABASE SERVER  
FWX - FIREWALL DEVICE  
RTR - ROUTER  
WWW - WEB SERVER

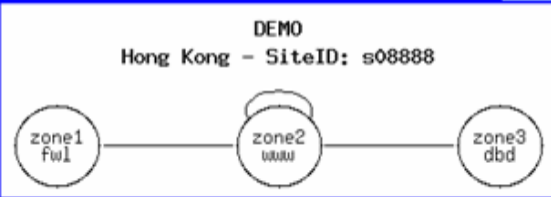


**Drill down to device layout & relationships**

San Francisco - SiteID: j08888



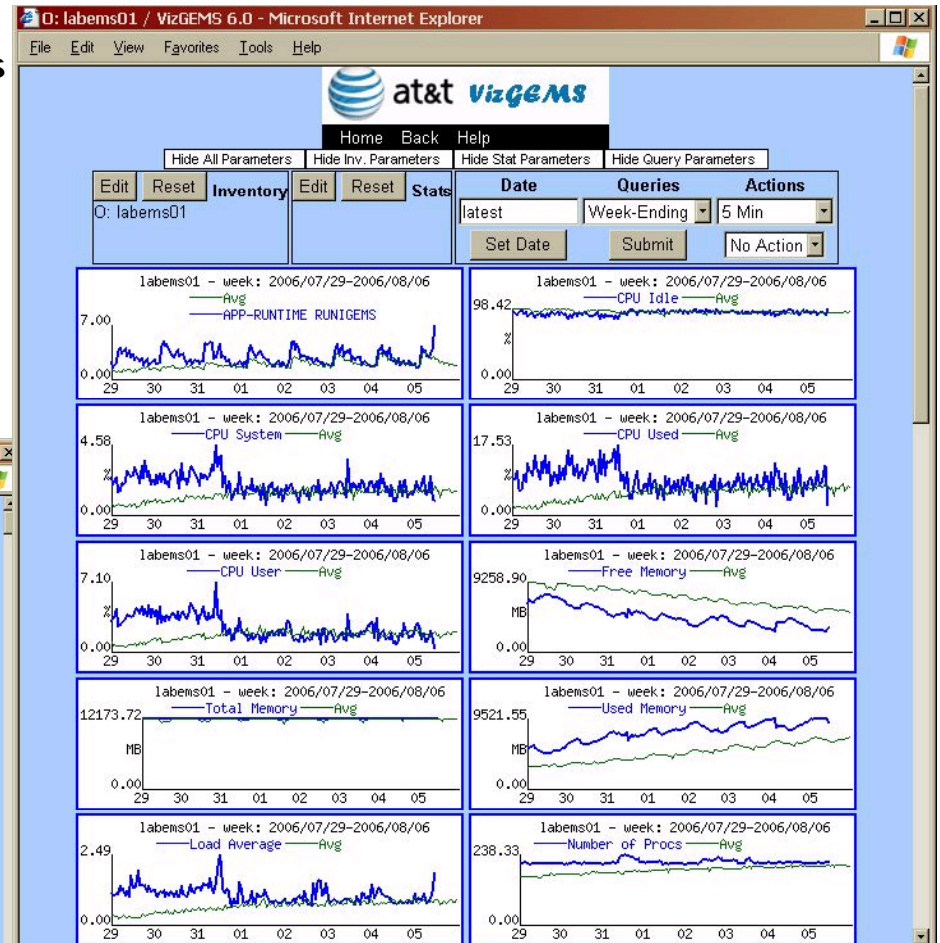
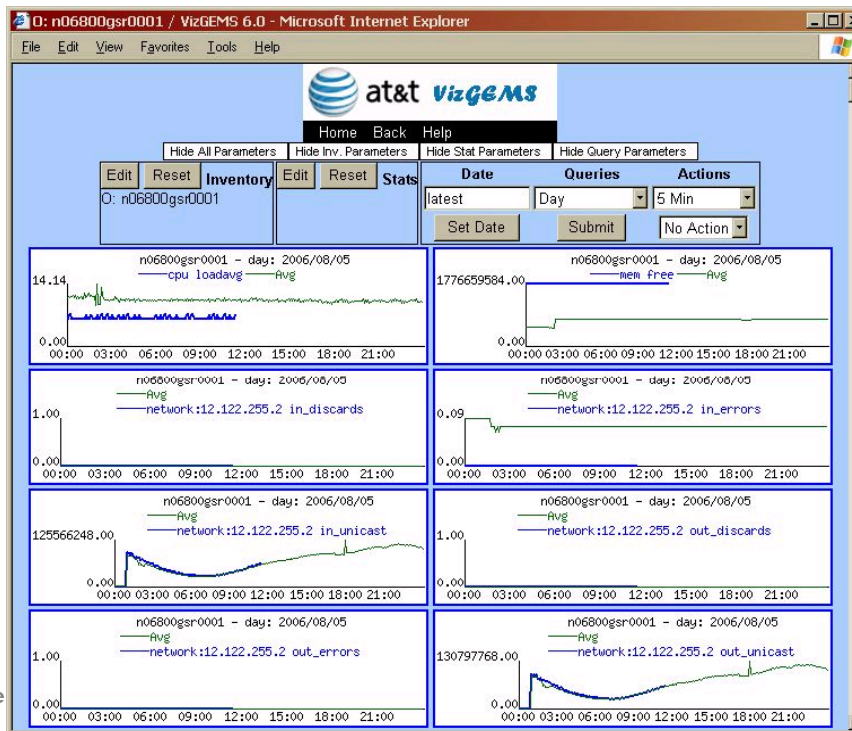
Hong Kong - SiteID: s08888



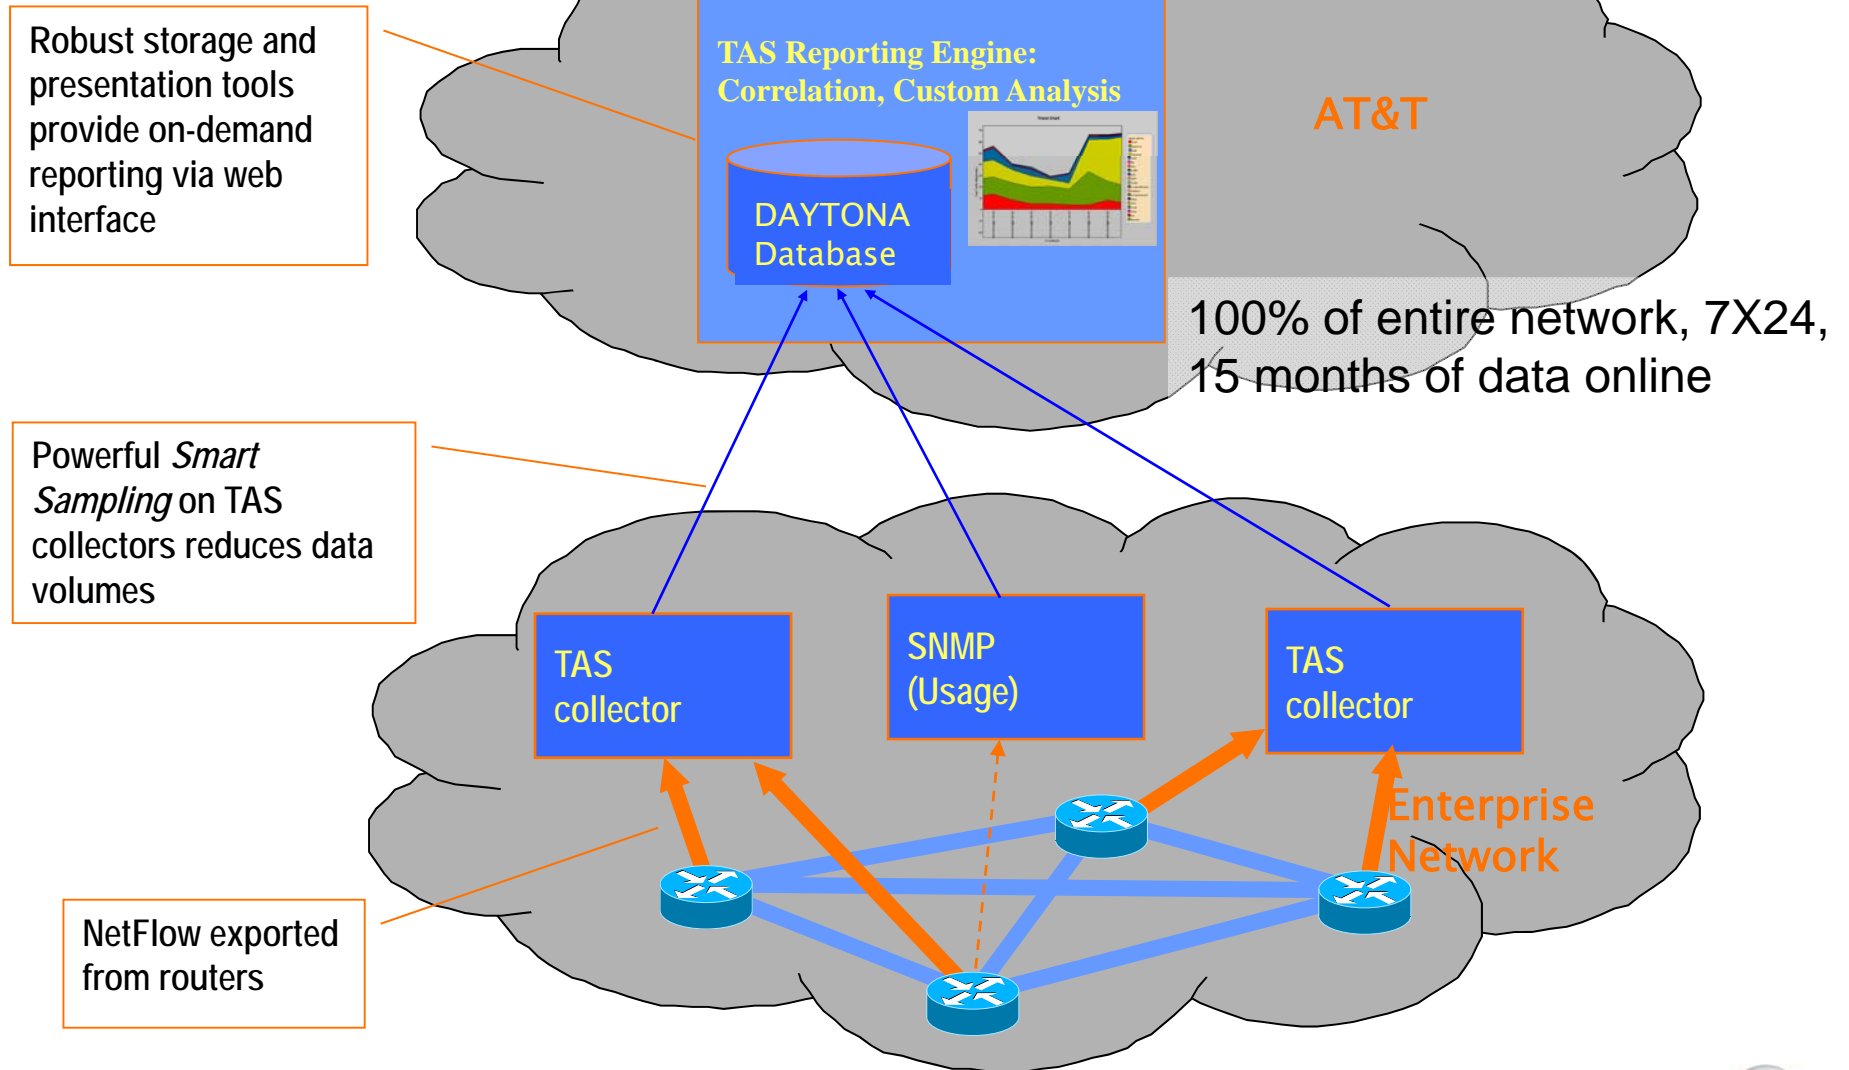


# Enhanced Control of Networks and Servers

- Early identification of trends
- Control charts identify unusual conditions
- Dramatic reduction of unnecessary alarms
- Complete drill down for fast resolution of problems
- Visual query capability
- All views, e.g. physical/logical, related




# Providing Traffic Analysis for Large IP Networks





# Netflow



## IP flow abstraction

- set of packets identified with same address, ports, etc.
- packets that are close together in time
- possible protocol-based flow demarcation; e.g., TCP FIN

## IP flow summaries

- reports of measured flows exported periodically from routers
  - flow identifiers, total packets/bytes, router state
- Several flow definitions in commercial use

## Eight keys define a Netflow:

- Source Address, Destination Address, Source Port, Destination Port, Layer 3 Protocol, TOS Byte (DSCP), Input Interface, Output Interface



# This Takes You From ... To ...

**From networks consisting of numerous, uncoordinated, error-prone systems**

- e.g., manual grappling with the changing state of the art in packet filters, route maps, IP, MPLS routing, layer 1-3 interworking, ...

**To networks where *operators* leverage automated network-wide views to assure performance**

- e.g., “assure negligible customer impact from planned cable intrusion scheduled tonight in New Mexico at midnight, mountain time”

**To networks where *designers* leverage automated mechanisms for real-time network response**

- e.g., “survive any single fiber or router failure with all link utilizations < 70%”

**To networks where *VPN customers* can leverage detailed views of their traffic and flexible, policy-driven routing**

- e.g., “traffic from east coast customer X sites should prefer X’s Dallas data center and X’s Phoenix Internet Gateway”



# Recommender Systems

The screenshot shows the Amazon.com homepage for user Chris Volinsky. At the top, the Amazon logo is on the left, and a personalized greeting reads: "Hello, Chris Volinsky. We have [recommendations](#) for you. (Not [Chris?](#))". Below this is a navigation bar with links for "Chris's Amazon.com", "Today's Deals", "Gifts & Wish Lists", and "Gift Cards". A search bar contains "Amazon.com". Below the search bar are links for "Shop All Departments", "Chris's Amazon.com", "Your Browsing History", "Recommended For You", and "Rate These Items".

The main content area is titled "Today's Recommendations For You". Below the title, it says: "Here's a daily sample of items recommended for you. Click here to [see all recommendations](#)".

Three items are displayed in a row:

- Tribute (Audio CD)**: An audio CD cover with a green background and the word "Tribute" in large, stylized letters.
- Java Concurrency in Practice (Paperback)**: A book cover with a blue and white design, featuring the title "JAVA CONCURRENCY IN PRACTICE" and the author "BRIAN GOETZ".
- Car Talk (Audio CD)**: An audio CD cover featuring a man playing a guitar, with the title "CAR TALK" and "BORN NOT BORN" visible.

Below the items is a horizontal menu of categories: "Action & Adventure", "All Categories", "Applied", "Book Characters", "Classics", "Classics", "ECM Jazz & World", "Electronica", "Fiction", "Humorous", "Indie Rock", "Jazz", "Literature", "Rock", "Shakespeare", "Shakespeare, William", "Singer-Songwriters", "Software Development", and "Sports".



# Recommender Systems

**Data consists of users (typically customers of a service) and items (typically products).**

**Which users like which items**

- Purchase history**
- Viewing history**
- Rating information**

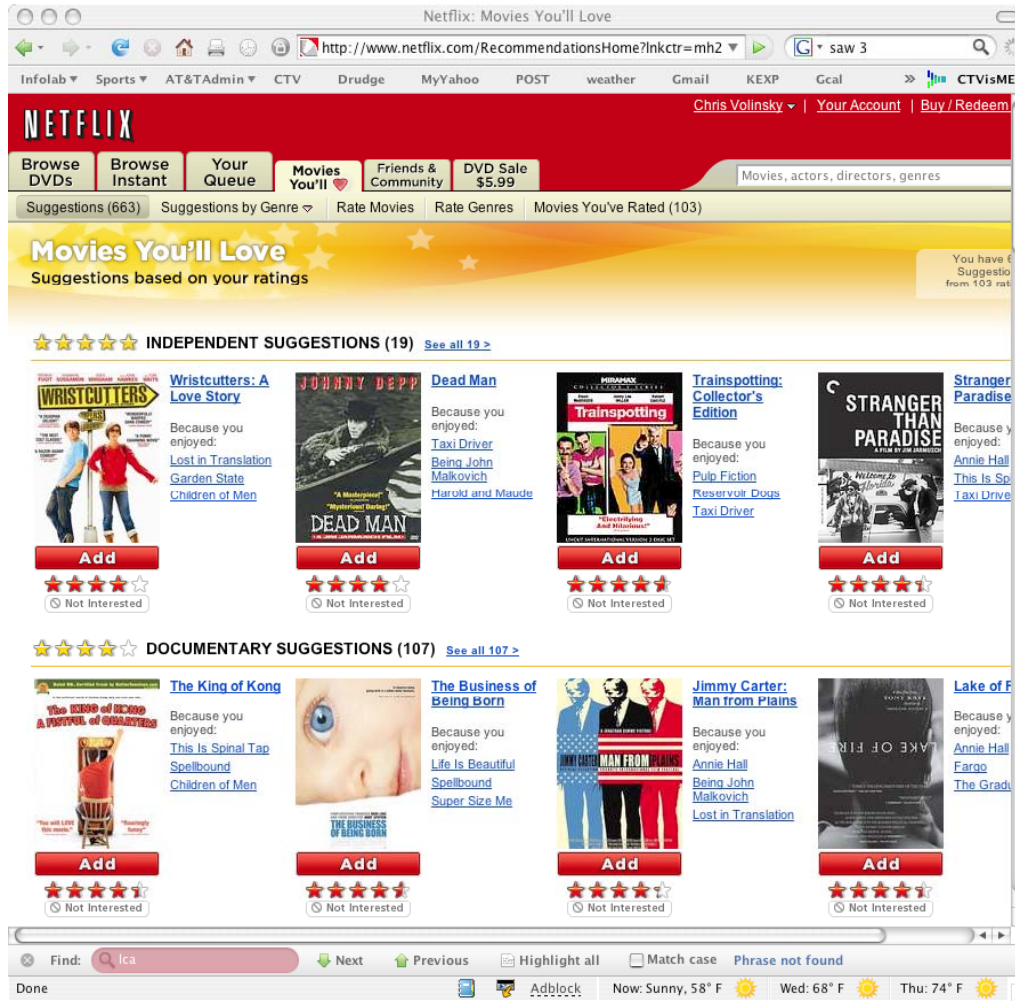
**Goal: recommend new items for users**



# Netflix

A US-based DVD rental-by mail company

8.2M customers, 100K titles, ships 1.9M DVDs per day



Good recommendations = happy customers



# Netflix Prize

October, 2006:

- Offers **\$1,000,000** for an improved recommender algorithm

Training data

- 100 million ratings
- 480,000 users
- 17,770 movies
- 6 years of data: 2000-2005

Test data

- Last few ratings of each user (2.8 million)
- Evaluation criterion: root mean squared error (RMSE)
- Netflix Cinematch RMSE: 0.9514
- results submitted by email

Competition

- 31K teams, 173 countries, 3,600 submissions
- \$1 million grand prize for 10% improvement on Cinematch
- If 10% not met, \$50,000 annual "Progress Prize" for best improvement

user	movie	score	date
1	21	?	2002-01-03
1	123	?	2002-04-04
2	125	?	2002-05-05
2	873	?	2002-05-05
2	8	?	2003-05-03
3	116	?	2003-10-10
4	245	?	2004-10-11
5	232	?	2004-10-11
5	998	?	2004-10-11
5	1012	?	2004-12-12
6	64	?	2005-01-02
6	156	?	2005-01-31

© 2010 AT&T Intellectual Property. All rights reserved. AT&T, the AT&T logo and all other AT&T marks contained herein are trademarks of AT&T Intellectual Property and/or AT&T affiliated companies. All other marks contained herein are the property of their respective owners.



# Nearest Neighbors

		users											
		1	2	3	4	5	6	7	8	9	10	11	12
movies	1	1		3			5			5		4	
	2			5	4			4			2	1	3
	3	2	4		1	2		3		4	3	5	
	4		2	4		5			4			2	
	5			4	3	4	2					2	5
	6	1		3		3			2			4	



- unknown rating




- rating between 1 to 5



# Nearest Neighbors

		users											
		1	2	3	4	5	6	7	8	9	10	11	12
movies	1	1		3		?	5			5		4	
	2			5	4			4			2	1	3
	3	2	4		1	2		3		4	3	5	
	4		2	4		5			4			2	
	5			4	3	4	2					2	5
	6	1		3		3			2			4	



- estimate rating of movie 1 by user 5





# Recommender Systems – Netflix Prize

- Cinematch RMSE was beaten in 2 weeks, by 10% in 3 years
- 51,051 contestants on 41,000+ teams from 169 different countries.
- 23428 valid submissions from 3133 different teams
- Great drama leading up to the Progress & Grand Prizes....
- Won by team led by AT&T Shannon Labs!

**Recommending new Products/Services/Processes based on customer behavior, preferences, characteristics.**



# BellKor solution...

**Our solution is a mixture of many different models.**

**Mostly variants of two main classes of collaborative filtering models**

- Nearest Neighbors**
- Latent Factor (SVD)**



# Nearest Neighbors

		users											
		1	2	3	4	5	6	7	8	9	10	11	12
movies	1	1		3		?	5			5		4	
	2			5	4			4			2	1	3
	<u>3</u>	2	4		1	2		3		4	3	5	
	4		2	4		5			4			2	
	5			4	3	4	2					2	5
	<u>6</u>	1		3		3			2			4	

**Neighbor selection:**  
**Identify movies similar to 1, rated by user 5**



# Nearest Neighbors

		users											
		1	2	3	4	5	6	7	8	9	10	11	12
movies	1	1		3		?	5			5		4	
	2			5	4			4			2	1	3
	<u>3</u>	2	4		1	2		3		4	3	5	
	4		2	4		5			4			2	
	5			4	3	4	2					2	5
	<u>6</u>	1		3		3			2			4	

Compute similarity weights:

$$s_{13}=0.2 \quad s_{16}=0.3$$



# Nearest Neighbors

		users											
		1	2	3	4	5	6	7	8	9	10	11	12
movies	1	1		3		<b>2.6</b>	5			5		4	
	2			5	4			4			2	1	3
	<u>3</u>	2	4		1	<b>2</b>		3		4	3	5	
	4		2	4		5			4			2	
	5			4	3	4	2					2	5
	<u>6</u>	1		3		<b>3</b>			2			4	

Predict by taking weighted average:

$$(0.2*2+0.3*3)/(0.2+0.3)=2.6$$



# Selected Research Directions

## The Next 5 Years

### Pervasive:

- Scale
- Security
- Mobility
- Operations
- Reliability

- **Network Based Computing:** Corporate grade “cloud” computing will be routinely employed in critical applications. Location and Presence.
- **Rich Media:** Mission critical, interactive applications will employ multimedia and move seamlessly between all 3 canonical screens.
- **Networks (including Internet) of everything:** *Billions of devices, many mobile, will interact as computing, sensing, and communications platforms.*
- **Information Leverage:** Collection, Analysis, Visualization, & Distribution will include all forms of data (e.g. Relational, Semi-Structured, Text, Speech, Video, Image) integrated, near real time, and at huge scale.
- **Communities of Interest:** Collaboration, Social Networks, Group Oriented Services.



# Networking & Services Research

## Traffic analysis & network management

*GSTool, DPI*

*Darkstar: predictive network mgt*

## Network design and performance analysis

*CBB, AGN, IPAG network design*

## Optical and wireless technology

*ULH and GPON, advanced radio*

## Speech and natural language understanding

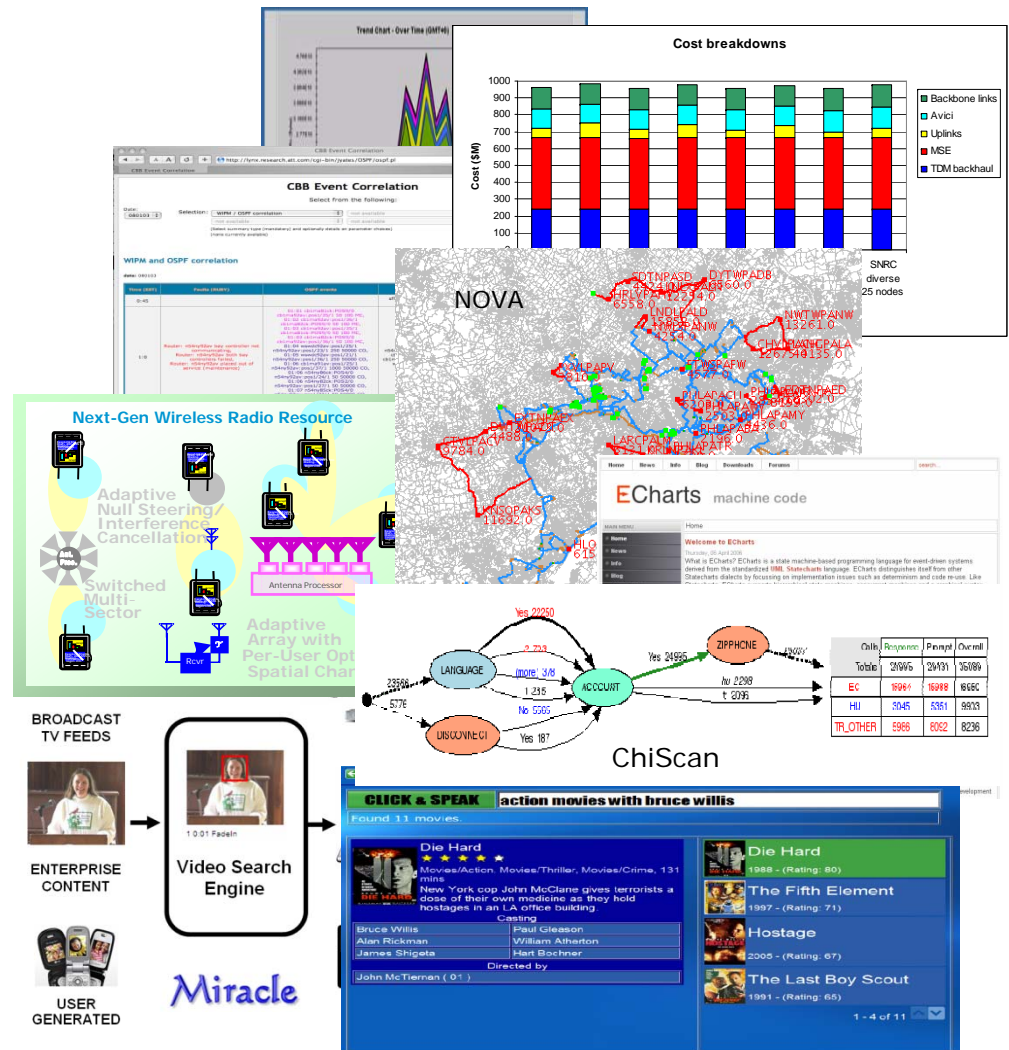
*Watson ASR & Natural Voices TTS*

## Speech, web, and email data mining

*Search: Local business & IPTV*

## Converged & multimedia applications

*Miracle multimedia content processing*



*Rethink Possible*

