

IntelliMagic



# Machines are for Answers, Humans are for Questions

Session 222b Tuesday Nov 7  
11:35 – 12:05 Beauregard

**Dr. Gilbert Houtekamer** – IntelliMagic Founder, CEO

Availability Intelligence



# Session Abstract

Big data is the big thing, and AI and predictive analytics will all put us out of job, or so some think.

We all know that computers are way better than humans for processing large amounts of data. They also, if artificially intelligent enough, are especially good at doing the things you tell them to do that would otherwise require human intelligence and unrealistic amounts of time.

So instead of human vs. machine, it can become human and machine vs. problem.

This is especially needed for predicting and diagnosing IT infrastructure performance and cost-efficiency problems. These problems are simply too time consuming for the human analysts to proactively search the vast and complex data sources for answers that machines can provide in seconds.

In this presentation, we will show how enabling the machine to utilize domain knowledge and workload information, also known as human intelligence, can be used in modern IT Operations Analytics (ITOA) to automatically highlight the areas that are the most important for a human eye and brain to spend time on. This discussion applies to any IT environment, our examples will be from z/OS and VMware SAN.



# Agenda

1. Questions vs. Answers
2. Human vs. Machines
3. How to make Computers Give Better Answers



## Questions vs. Answers

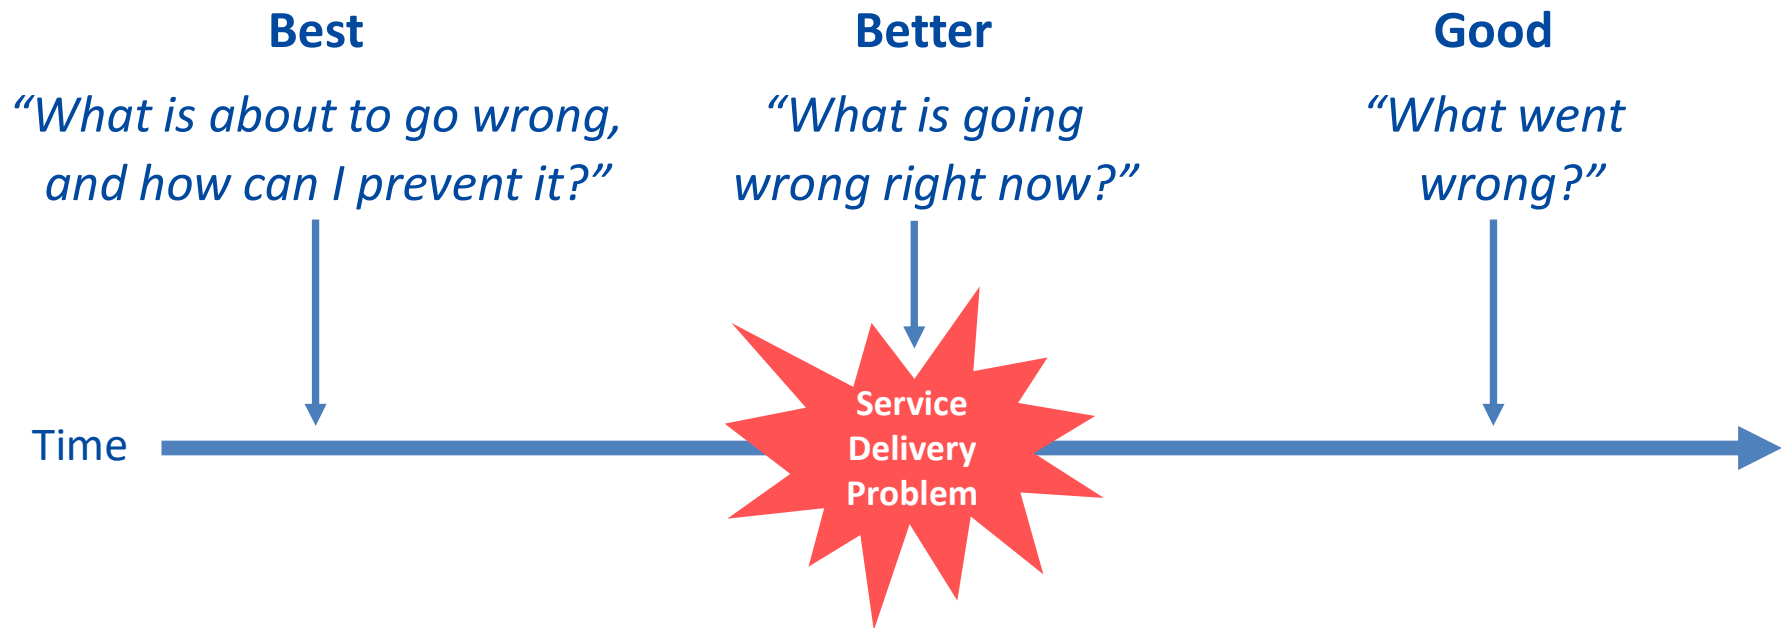
- *"Computers are useless. They can only give you answers."*  
- Pablo Picasso
- *"My greatest strength as a consultant is to be ignorant and ask a few questions."* - Peter Drucker
- *"The art and science of asking questions is the source of all knowledge."* - Thomas Berger
- *"Everyone gets so much information all day long they lose their common sense."* - Gertrude Stein



## "Machines are for answers; humans are for Questions"

- Attributed to Kevin Kelley <https://www.edge.org/conversation/the-technium>
  - *"Machines are for answers; humans are for questions."*
  - *"The world that Google is constructing—a world of cheap and free answers—having answers is not going to be very significant or important. Having a really great question will be where all the value is."*
- True for performance and capacity planning as well
  - The value of humans is in asking great questions
    - NOT in spending time searching many data points for answers
  - The value of computers is in deriving answers (intelligence)

# Some Questions are more Valuable than Others



# Why the Best Questions are Hard to Answer: Challenges from the Environment

- **Size of Environment**
  - MIPS / TB / Users
  - Huge decrease in MF Ops headcount / resource...
- **Complexity of Environment**
  - More virtualization means less visibility into utilizations
  - More networking, cross-platform collaboration
- **Workload (Un)Predictability**
  - 24 x 7 access, mobile devices driving back-end transactions

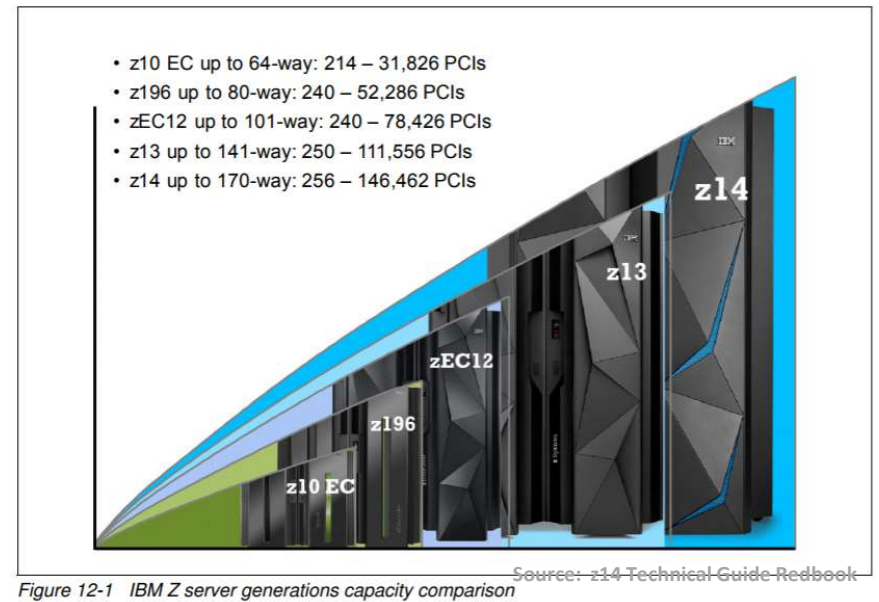


Figure 12-1 IBM Z server generations capacity comparison



## Why the Best Questions are Hard to Answer: Challenges due to the nature of the RMF/SMF Data

- **Interpretation of the Data**

- What does the data mean for my application workloads on my infrastructure?
  - Is it good or bad? Risky or plenty of performance horsepower?
  - Specialized expertise required for each element of the infrastructure
- Abstraction of logical and physical infrastructure
- Shared physical resources amongst multiple users/applications

- **Quantity of Data**

- GB's of data per day – this is good:
  - **The richer the data source, the smarter the analytics can be if designed right!**
- And even more measurement data to support has been recently created ...





# Humans vs. Machines



## Computers are Taking Jobs, but what Kind?

- *"Artificial intelligence is the science of making machines do things that would require intelligence if done by men"*
  - Marvin Minsky 1968
- For enterprise computing infrastructure perf/cap jobs:
  - The jobs that humans don't have time to do anyway
  - Repetitive, process-driven jobs
  - Jobs that require the same kind of expert-knowledge to assess metrics over and over
  - Important jobs that are not being done today due to lack of time.



## What defines an human expert?

- Understanding of a discipline as a whole
- Insight in what matters
- Ability to combine information from multiple sources
- Experience



## Example: Human Expert's Post-Mortem Approach

Incidents can be investigated with logging data to find the **root cause**.

What do the Experts do:

- Apply subject matter knowledge and experience
- Look for specific metrics and relations they know to matter
- Do back of the envelope sanity checks
- Zoom in on likely culprits until root cause is found
- Check for configuration problems and errors



## What defines a computer AI system?

- Ability to process certain types of data
- Ability to 'understand' that data by applying best-practice rules
- Ability to apply statistical analysis to data
- Ability to report findings such that humans can act upon them



## Example: Computer AI looking for symptoms

Logging data can be analyzed looking for unusual or undesirable patterns

What do the computer can do do:

- Apply subject matter knowledge to the data
- Look for specific metrics to validate their range
- Check response time with workload and hardware using model
- Check for configuration problems and errors



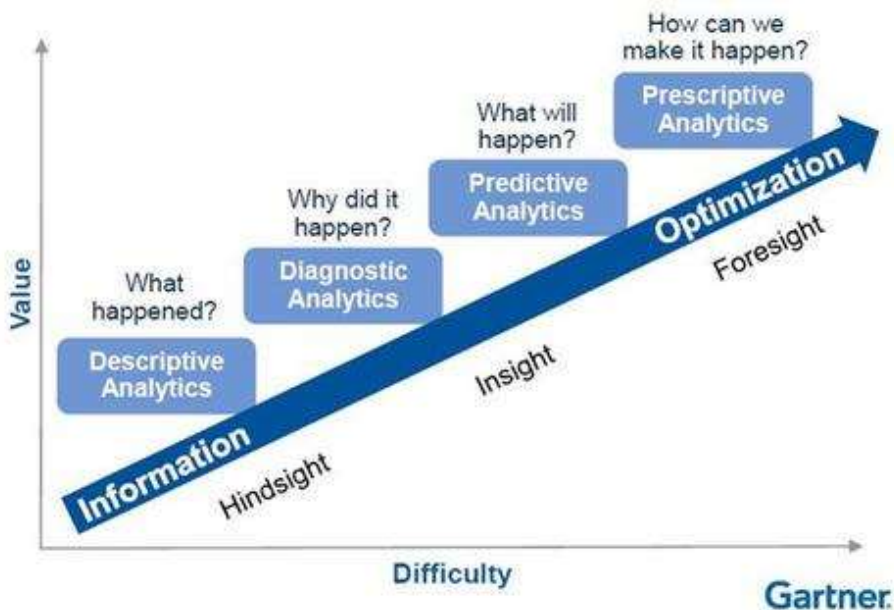
## Computer – Human Relationship

- Not “humans vs. computers” but “humans and computers vs. problem”
- IT performance and capacity teams should:  
*“Invite AI to be part of the team”*
- Artificial Intelligence needs Real (Human) Intelligence to be actionable
- Human Intelligence needs Artificial Intelligence to be indispensable

# Force Multiplier Technology for the Perf/Cap Team

Gartner's analytics maturity model:


- Even large sites are still using shovel-era technology (not even on chart)
- You don't need a new "reporting tool".
  - Already have more reports than the team can manually look at / interpret
- *"Humans are for questions, machines are for answers"* (K. Kelly)
  - The same built-in, automated expert knowledge enables answers at all 4 analytics maturity levels
  - Automate the tedious things your experts have to do – give them intelligence, not data or static reports!





# Force Multiplier Technology for The Perf/Cap Team

- Technology that enables more to be accomplished with less human effort
- Smart Strategy for:
  - Current experts: Empower them with a force multiplier
  - New staff: Accelerate time to productivity with a force multiplier
- Where is your fulcrum?
  - If your RMF/SMF solution relies on a catalogue of static reports, your fulcrum is far right, requiring inordinate amounts of effort
  - Predictive work simply cannot be done with antiquated technology
    - It is not the team's fault!





# How to Make Computers Give Better Answers



## Humans vs Computers

	Strength	Weakness
Human	Insight Questions	Looking at millions of charts Gets bored & Retires
Computer	Number crunching Applying (smart) assessments	Understanding



## Humans and Computers

	What?	When
Human	Define questions to be addressed and assessments to be done during (log) data analysis	Once
Computer	Run software programs with millions of data points and present findings	At least daily
Human	Analyze findings via deep-dive ideas	Based on findings or phone call
Computer	Provide deep-dive repository to support investigation	




## What happened?

- Human instructs the computer what to look for
- Computer (log crunching) highlights findings
  - Massive data reduction
- Human interprets these findings, and likely wants root-cause or context
- Computer (data repository with user interface) provides the answers

=> Bionic Performance Analyst

# What happened?



=> Bionic Performance Analyst



## What this really does

Apply hundred of questions to millions of fields to provide the answers we were looking for

- The purpose of the questions should be to **predict** potential availability issues (based on the performance data base)
- If 'it' knows **why** an issue is identified, 'it' can **explain** it.
- Performance data base can answer the follow-on questions to **resolve** the problems

Integration of rules with structured data repository enables experts to investigate and validate findings





## 5 Perf/Cap Problem Areas Mature Analytics Solves

### 1. **Speed: Fast and Current**

- Problem: Finding answers about what happened takes too long. Supporting new infrastructure and data types is hard. Proactive and predictive analysis not feasible.

### 2. **Usability: Visual and Interactive**

- Problem: Physical, virtual, logical infrastructure shared resources are very difficult to tie together and navigate to understand application performance problems.

### 3. **Intelligence: Predictive and Contextual**

- Problem: Understanding what the measurement data means in the context of the infrastructure running it takes great expertise and time from different silos

### 4. **Versatility: Expanded Application**

- Problem: it takes so much time and expertise that the data is not used for important purposes like predictive analysis, DevOps performance testing, etc.

### 5. **Agility: Cloud and Collaborative** – Problem: Difficult access, maintenance, sharing





## (1) Speed: Fast and Current

- Intuitive, GUI-based interface makes data easily accessible
  - Enables analysis to begin quickly, without the need to write programs or other methods requiring extended time to master
  - Easy access to explore data expedites learning
- Support for emerging data types provides insights required to manage the associated technologies
  - Eliminates need to invest resources developing own programs to process and report on new data sources
  - Challenges finding skills or time often result in significant gaps in visibility into data required to support newer technologies



## (2) Usability: Visual and Interactive

- Easy visibility to data boosts learning and productivity and surfaces issues otherwise difficult to identify
  - Gain insights made apparent by the data
  - View relationships between metrics
  - Determine if observed values represent “normal” or “anomaly”
  - Identify patterns and trends
  - Integrate separate data sources to present new insights



## (2) Usability: Visual and Interactive (cont.)

- Flexible drilldown capability expedites learning curve
  - Can select alternative analytical paths based on current display
  - Though inexperienced analyst may not know precisely what to look for, he or she can begin exploring with minimal effort
  - Cost of going down "wrong path" is minimal
- Immense difference in analytical effectiveness between:
  - Leveraging drilldowns to logically connected metrics accessible with a click, and
  - Analysis done through accessing a catalog of static reports



### (3) Intelligence: Predictive and Contextual

- Identify key metrics to answer “where should I start looking?”
  - Out of dozens or more available for most areas
  - Consolidate into a single view or limited number of views
- Automatically assess “good” and “bad” conditions for those key metrics and rate the severity of the assessed condition.
  - Indicates issues that may warrant further investigation
- Identify current or potential risks to production application availability due to thousands of z/OS infrastructure elements

## (4) Versatility: Expanded Application

- Deliver new, meaningful metrics derived from raw SMF data
- Enable application view of infrastructure resources (e.g., CF)



## (5) Agility: Cloud and Collaborative



- Cloud solution offers rapid startup and no ongoing product support effort
- Single integrated interface enables collaboration with other mainframe disciplines



# Questions?