

# Capacity Analytics for System z

**Dr Rajkumar Munusamy**

Tivoli System z WW Product Manager

**Eric Chicha**

z Business Analytics Technical Leader WW

**Alan Place**

Tivoli TDSz and TADz Dev Manager

29 May 2013



# IBM Capacity Analytics on System z is about predicting business demand and preparing to meet it successfully

## Key Takeaways



1. Addressing **Cloud, Big Data and Mobile** business requirements requires on-going capacity management
2. Good capacity management will **reduce risks and decrease costs** by making sure no unexpected workloads impact business and SLAs
3. IBM capacity management solution on System z using business analytics provides increased **flexibility and productivity**

## What is Capacity Management and what new technologies make it important to focus on today?

Capacity Management performs business focused analysis of IT services and infrastructure to drive agreed service levels in cost effective and timely manner

### Cloud computing

- Cloud value primarily about maximizing efficiency and cost savings
- Capacity management required to understand workload impacts

### Mobile

- Capacity patterns harder to predict when mobile driving workloads 24/7

### Big Data

- Broke process of human experts analyzing information
- Fuels the need for better capacity management process



## Capacity Analytics on System z will provide savings today and going forward

### ▪ Reduce costs

- HW, SW and labor costs
- Fewer physical servers required to run workloads
- Reduce number of required licenses

### ▪ Ensure application availability


- Resources overloaded?
- Significant changes in environment over time?
- Can supply meet demand?
- Are business policies met?

### ▪ Optimize resources

- Right size virtual machines
- Identify trends for workload balancing



## Capacity Management includes different components that provide detailed targeted information



- Translates business needs and plans into requirements for services and IT infrastructure
- Management, control and prediction of end-to-end performance and capacity of workloads
- Management, control and prediction of performance, utilization and capacity of IT technology

## How Capacity Management keeps Business and IT aligned

**Executives want to know how a change in their business will impact IT demand and if they have sufficient capacity to satisfy the demand:**

- *We are bringing in four more hospitals over the next 6 months. Do we have enough capacity in our infrastructure?*

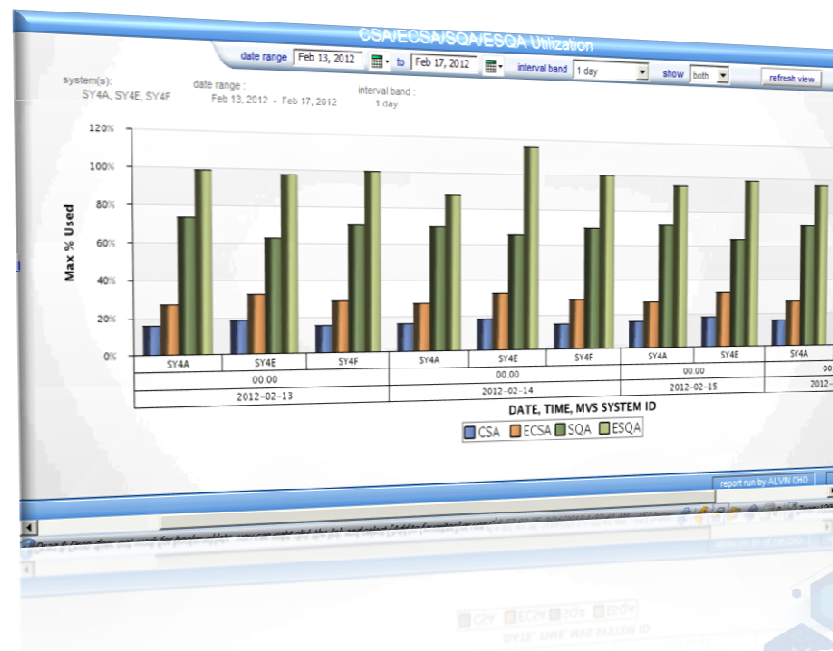
### **Language difference between business and IT**

- *When business launches a marketing campaign, the IT organization should be informed to take appropriate measures, for instance to increase the number of application servers because more customers will connect with the internet environment.*
- *Does business understand impact of adding Mobile access to banking transactions, and how does IT predict increased usage and IT overhead ?*

# Capacity Analytics: Understand how current system is running

## *Immediate Insights to System Performance*

- Scorecards
- Dashboards
- Reports



# Capacity Analytics: Use data to figure out future usage

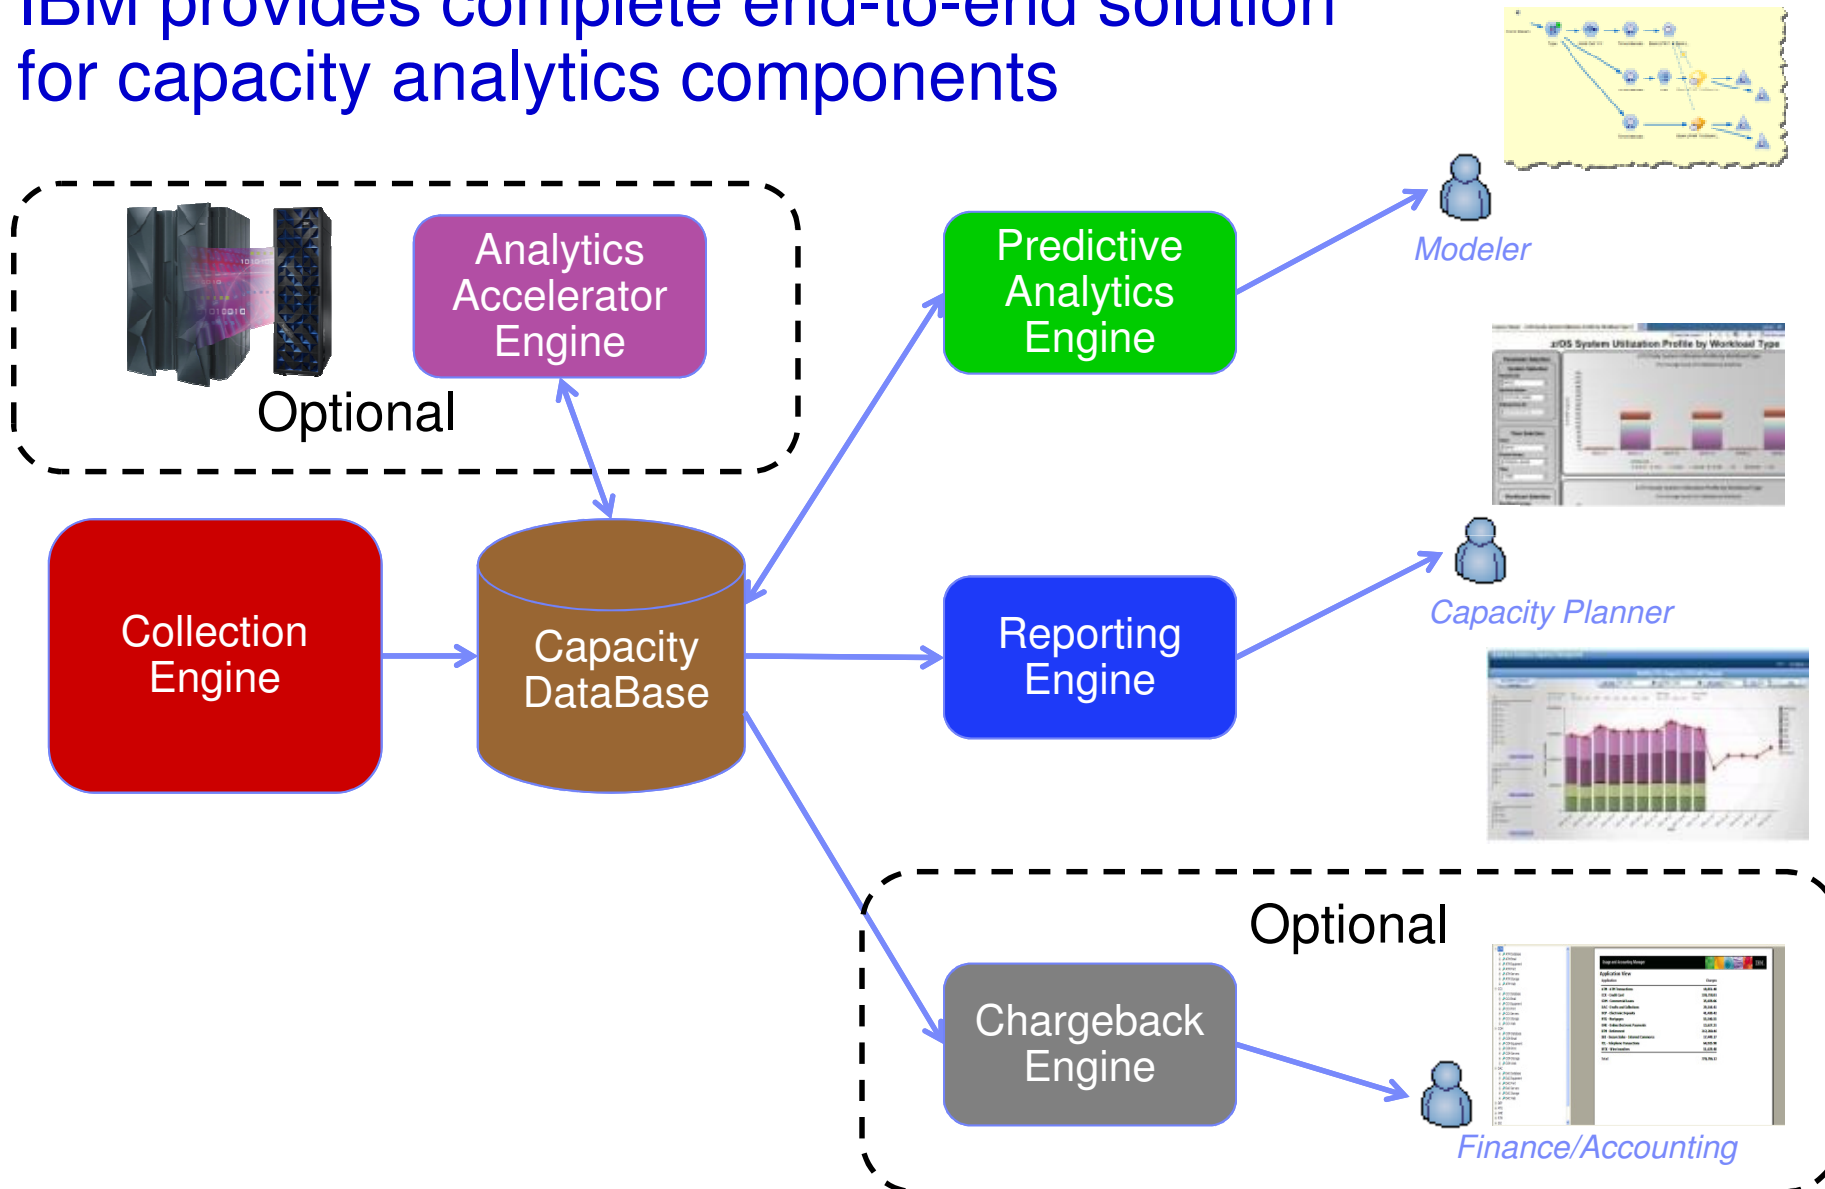
## *Forecast to Plan & Allocate Resources*

- What-If Analysis
- Predictive Analysis



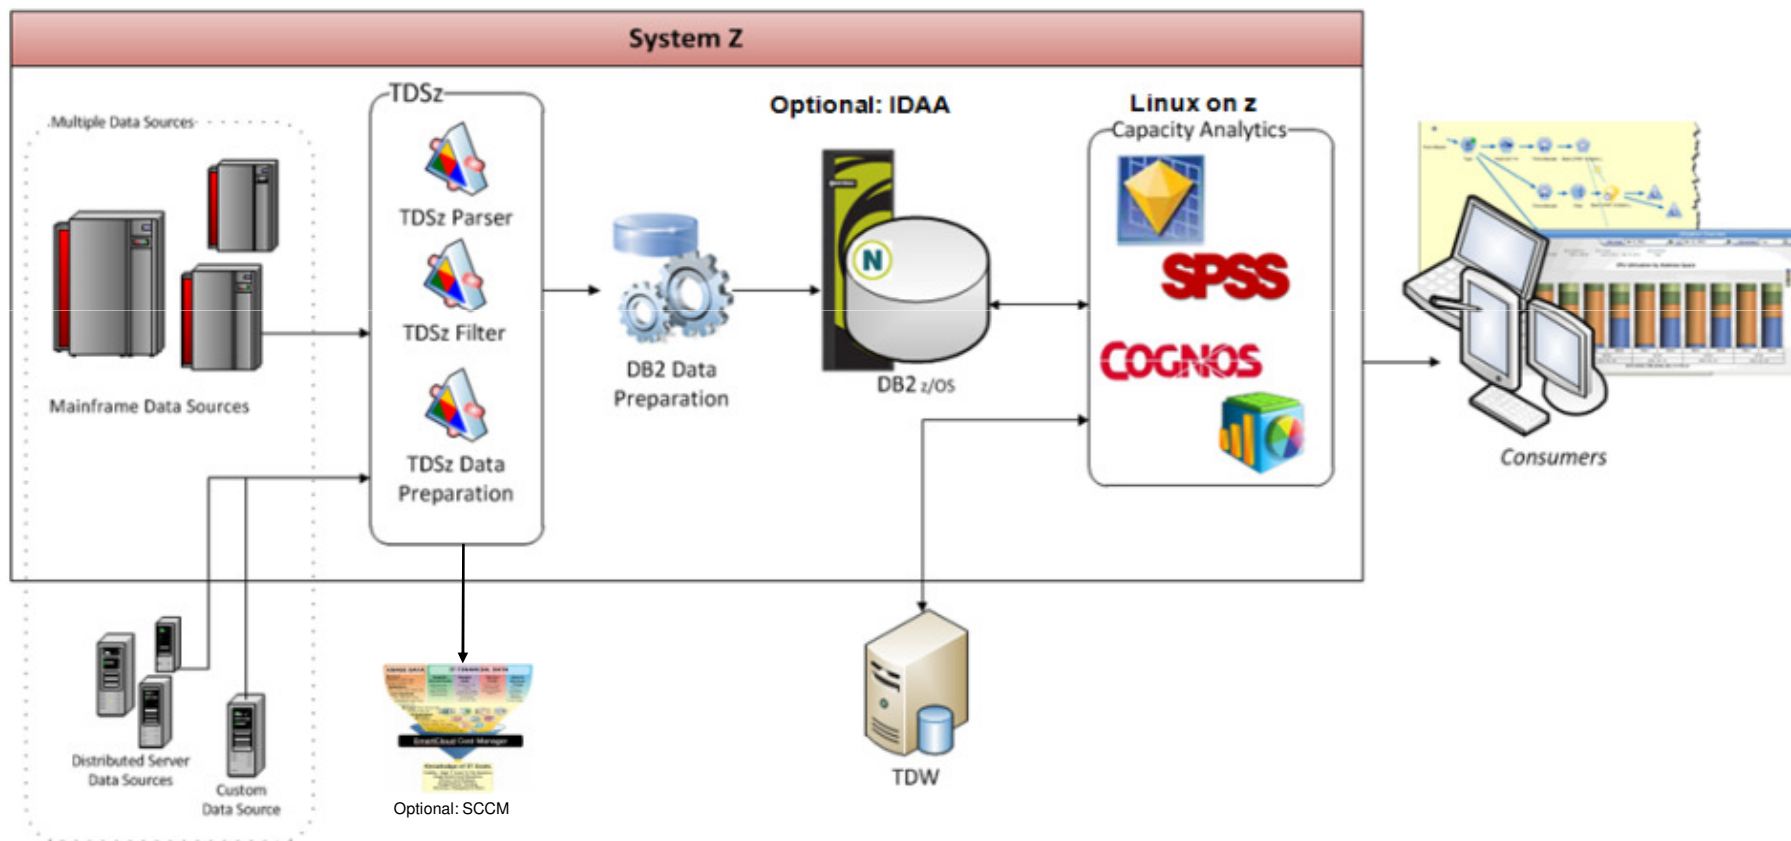


# IBM provides complete end-to-end solution for capacity analytics components



# Core Architecture

## IBM Capacity Analytics – Core Architecture



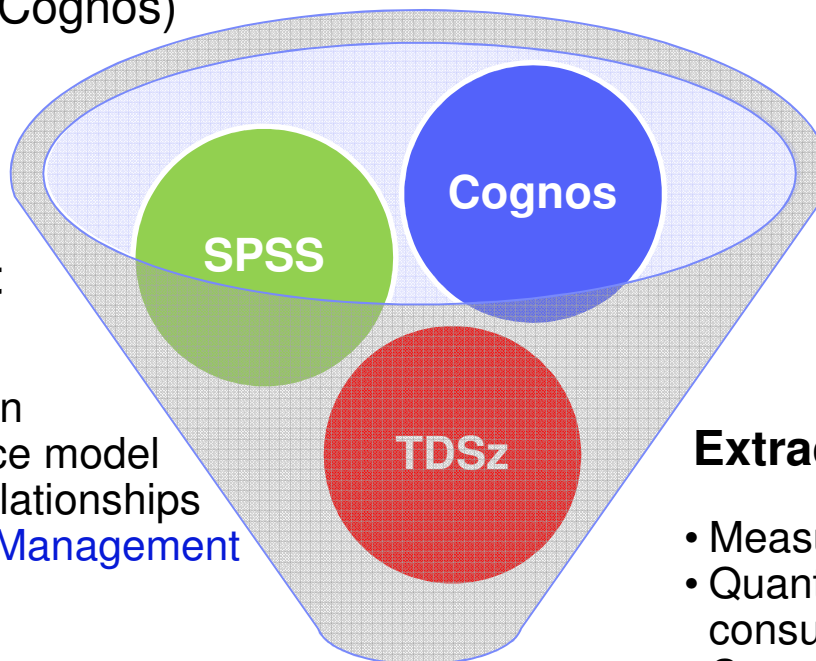
# IBM solution capabilities work together



- Collection Engine (TDSz)
- Predictive Engine (SPSS)
- Reporting Engine (Cognos)

## Correlate & Forecast

- Granularity / Statistical
- Forecasting / Prediction
- Application performance model
- Correlation of data / relationships
- Use beyond Capacity Management



Capacity Management

## Reporting

- Report / Presentation layer
- Federated data model
- Use beyond Capacity Management

## Extract, Categorize, Store

- Measure SLA compliance
- Quantify increased IT resource consumption or abnormal spikes
- Compare trends to pinpoint where consumption increased
- Converts raw data into business-relevant information
- Basis for mainframe accounting

## Reporting Starter Kit

- Reduce report working set size
- Users can drill-down to lower level detail
- Customize frequency of data updates
- User selectable data window – view 1 day or 1 year of data
- Moves the paradigm away from static reporting to a self service model
- Provides exec level dashboard on delivery against SLAs

Standard Predictive Model:  
Inventory: LPAR Configuration

CEC: Processor Complex(s)  
CEC: Processor Complex(s) with LPAR information  
CICS: File Usage  
CICS: Program Usage  
CICS: Subsystem Overview  
CPU: CPU Utilization - CEC Level  
CPU: CPU Utilization - LPAR/System Level  
CPU: CPU Utilization - Service Class Level  
CPU: Daily CPU Usage by LPAR with FORECAST  
CPU: Monthly Usage by LPAR with FORECAST

DB2 - CPU Utilization Details  
DB2 - Package(s) Overview  
DB2 - Plan(s) Overview  
DB2 - Subsystem(s) Details  
DB2 - Utilization Overview

I/O: Channel Utilization  
I/O: DASD IO Performance S

Storage: CSA/ECSA/SQA/ESQA Utilization

TDSz: Data Collection by System  
TDSz: Data Collection Currency  
TDSz: Installed Components

zLinux – CPU Usage by System  
zLinux – Memory by System  
zLinux – Paging by System  
zLinux – # Processes by System  
zLinux – # Users by System

Exception Detection  
LPAR MIPS Usage Control Chart  
Simple Exception Detection Chart

# Reporting with Cognos BI can show today and help predict future usage

**Easy to use interface**

**LPAR Forecast**

**Current usage by LPAR**

IBM Cognos Viewer - CPU Usage Report (in the wild v0.2)

Processor: LPAR Daily CPU Forecast

Utilization Overview

Max % Used

DATE, TIME, MVS SYSTEM ID

CSA ECSA SQA ESQA

DB1A DB1M

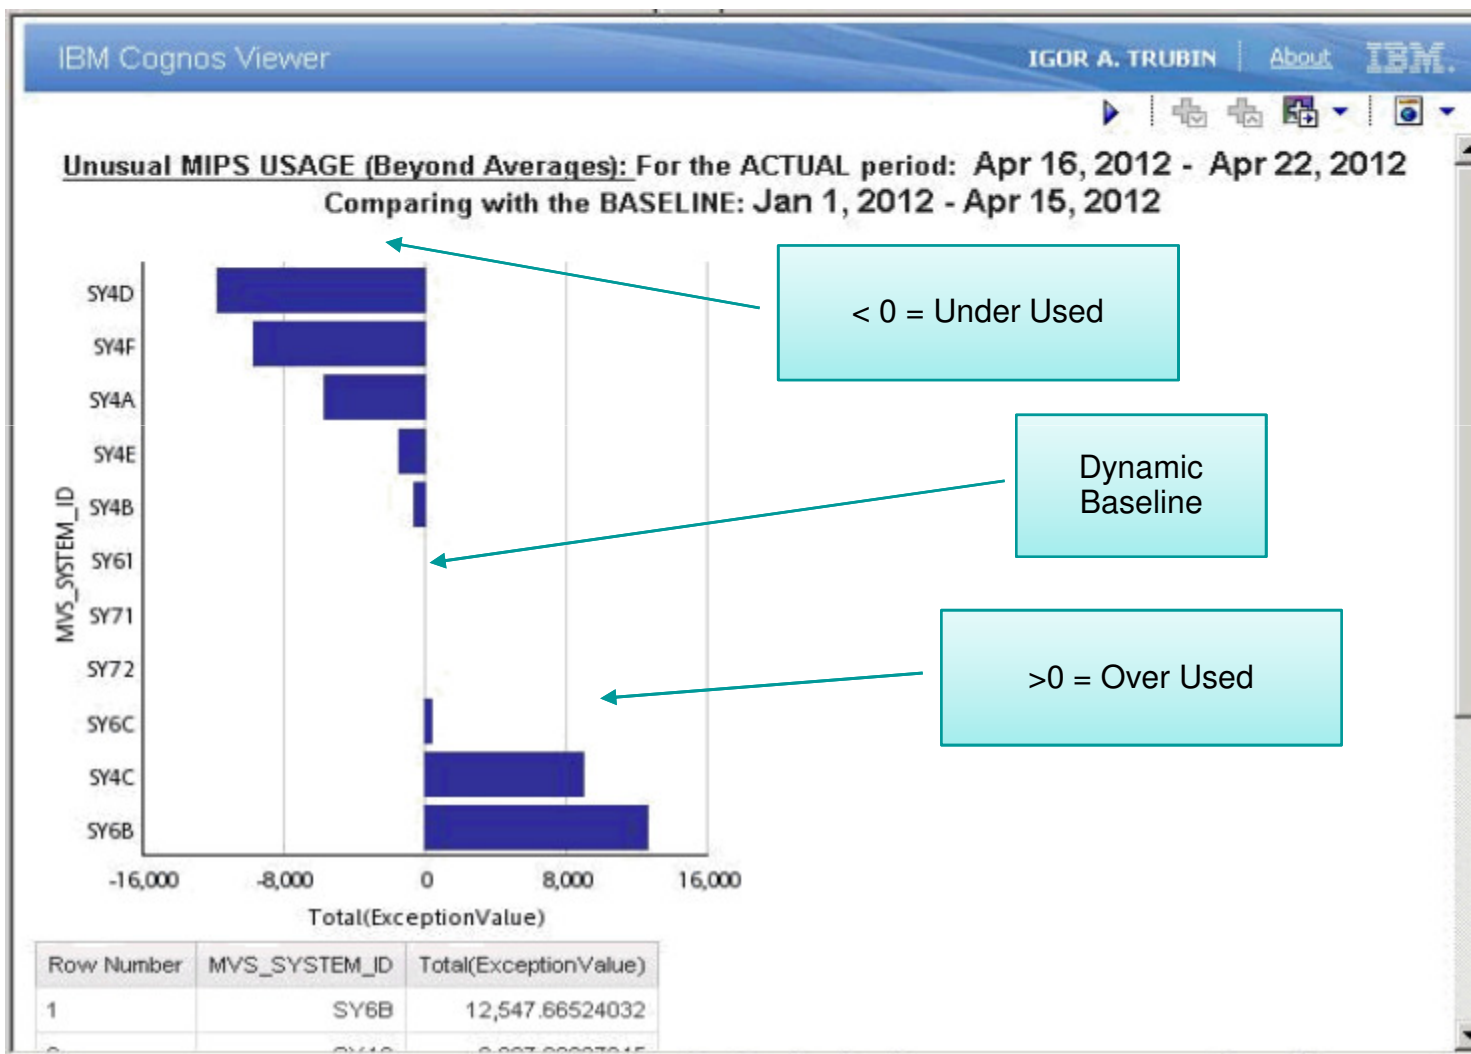


# Reporting with Cognos BI – Advanced Filtering


The screenshot illustrates the IBM Cognos BI interface for reporting on CPU usage. It is divided into several key sections:

- Filtering Section (Top Left):** Contains a 'date range' field set to 'Apr 1, 2011' to 'Aug 24, 2011' and an 'interval band' dropdown set to '1 day'. Both are highlighted with red boxes.
- View Options (Top Right):** Includes a 'show chart' dropdown (highlighted with a red box) and a 'refresh view' button.
- Filter Panel (Bottom Left):** A sidebar for filtering by 'cpu / processor / lpar'. It includes:
  - cpu serial no.:** A dropdown menu with '31CE' selected.
  - processor type(s):** A list with checkboxes for 'CP', 'IFA', and 'IIP', all of which are checked.
  - lpar name(s):** A list of lpar names with checkboxes, including '\*PHYSICAL', 'COGVMLNX', 'COGVMLN2', 'COGVMLN3', 'COGVMLN4', 'ICF0F', and 'ICF01', all of which are checked.
- Main Report Area (Center):** Displays a bar chart showing CPU usage over time. Above the chart, it lists the selected filters: 'date range: Apr 1, 2011 - Aug 24, 2011', 'interval band: 1 day', 'cpu serial no.: 31CE', 'processor type(s): CP, IFA, IIP', and 'lpar name(s): \*PHYSICAL, COGVMLNX, COGVMLN2, COGVMLN3, COGVMLN4, ICF0F, ICF01, ICF1D, ICF1E, ICF1F, ICF3A, SPSVMLNX, SPSVMLN2, STLAB7, STLABF6, STLAB4A, STLAB4B, STLAB4C, STLAB4D, STLAB4E, STLAB4F, STLAB51, STLAB52, STLAB53, STLAB54, STLAB6B, STLAB6C, STLAB6D, STLAB60, STLAB61, SVLXCOD1, SVLXCOD4, SVLXCOD5, SVLXCOD6, SVLXCOD7, SVLXCOD8, SVLXCOD9'. Below the chart is a legend for 'lpar name, processor type' with various colored squares corresponding to the data series.

# Other Reports: Statistical Exception & Trend Detection System




# Other Reports: Statistical Exception & Trend Detection System





# Other Reports: Statistical Exception & Trend Detection System

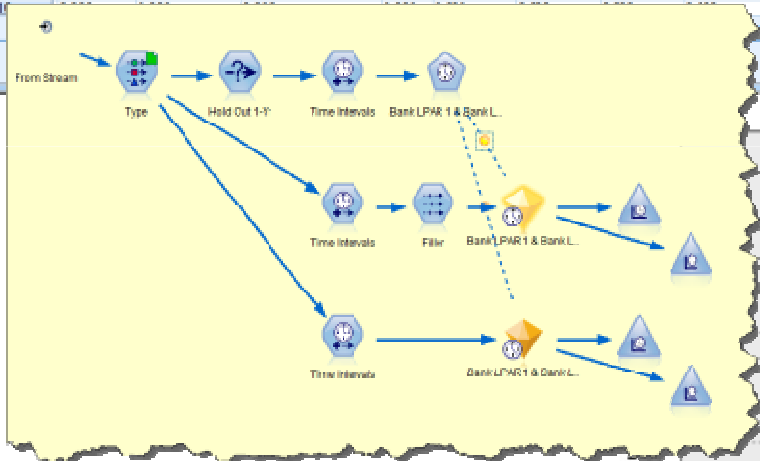


# Powerful Model with SPSS – See impact over time

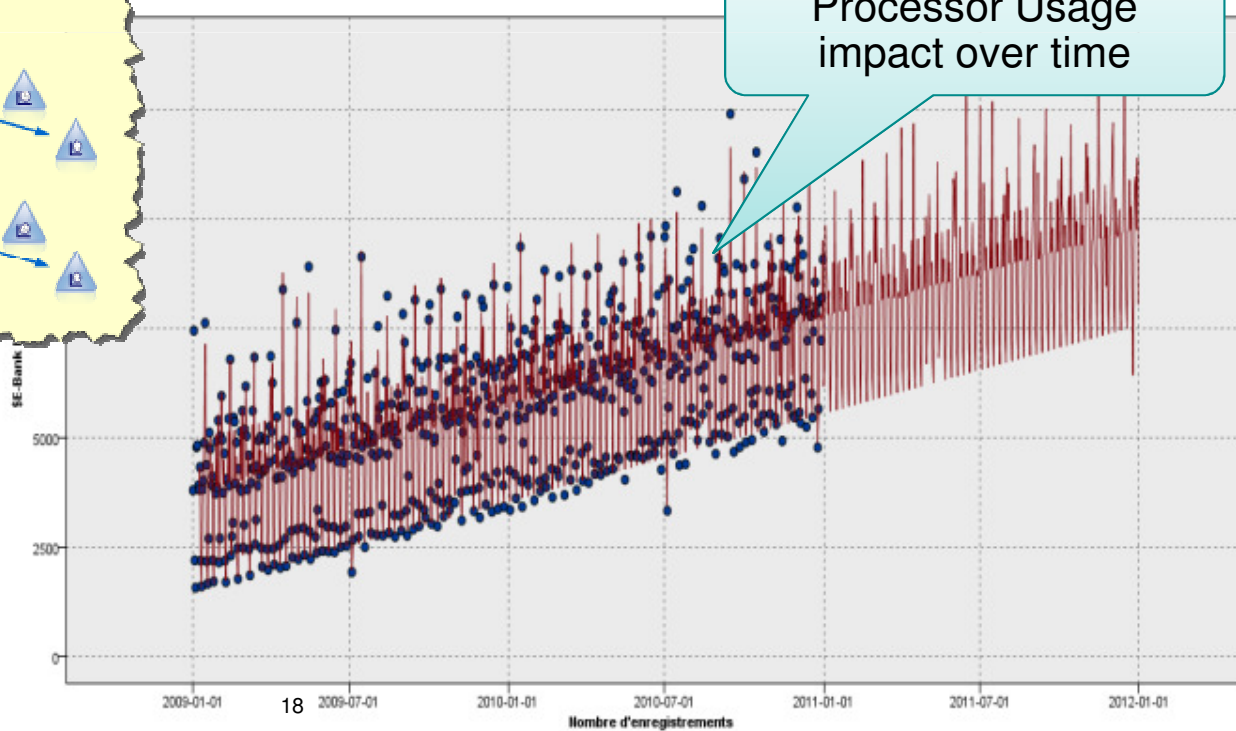
Preview from Causal Factors Node (25 fields, 13 records)

	Ides of Mnth	Fifteenth of Mnth	Prior Day Was Holiday	Holiday	SS Payday 1	SS Payday 2	SS Payday 3	SS Payday 4	Second Bus Day of Mnth	Third Bus Day of Mnth	Day After Thanksgiving
1	0.000	0.000	0.000	1.000	0.000	0.000	0.000	0.000	0.000	0.000	0.000
2	0.000	0.000	1.000	0.000	1.000	0.000	0.000	1.000	0.000	0.000	0.000
3	0.000	0.000	0.000	0.000	0.000	0.000	0.000	0.000	0.000	0.000	0.000
4	0.000	0.000	0.000	0.000	0.000	0.000	0.000	0.000	0.000	0.000	0.000
5	0.000	0.000	0.000	0.000	0.000	0.000	0.000	1.000	0.000	0.000	0.000
6	0.000	0.000	0.000	0.000	0.000	0.000	0.000	0.000	1.000	0.000	0.000
7	0.000	0.000	0.000	0.000	0.000	0.000	0.000	0.000	0.000	0.000	0.000
8	0.000	0.000	0.000	0.000	0.000	0.000	0.000	0.000	0.000	0.000	0.000
9	0.000	0.000	0.000	0.000	0.000	0.000	0.000	0.000	0.000	0.000	0.000
10	0.000	0.000	0.000	0.000	0.000	0.000	0.000	0.000	0.000	0.000	0.000

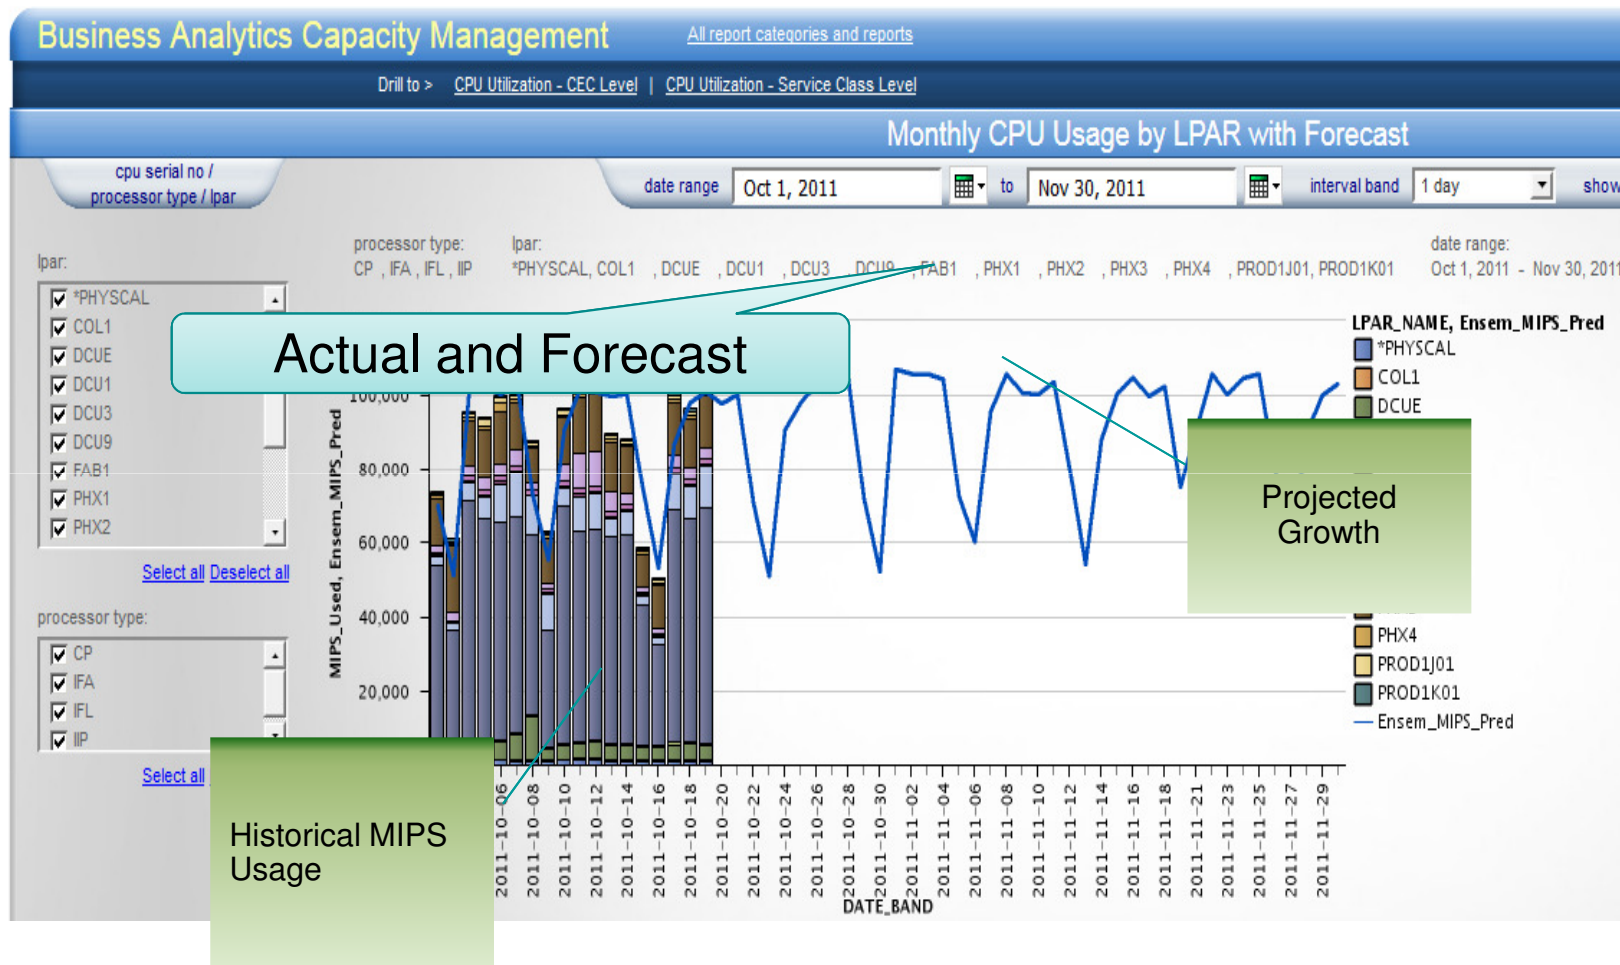
Inputs for time series model



Processor Usage impact over time



# Show business impact of growth over time



\* Actual and Prediction on same chart

# Multiple Distribution Methods



Reply | Reply to All | Forward | Subject

cognosD1    Rapport : CPU: CPU Utiliz  
cognosD1    Rapport : Data Collection  
notifications    Report: TDSz: Data Collec  
cognosD1    Report: Held Batch Jobs

21/06/2012 16:44    40K

**Rapport : CPU: CPU Utilization - LPAR/System Level**  
cognosD1 to: Marc Soumahoro Amadou    22/06/2012 09:10  
Custom expiration date: 22/06/2013    [Show Details](#)

1 attachment

System Level.mht

Bonjour Mr AMADOU,

Veuillez trouver ci-joint le rapport "Utilisation CPU par LPAR" demandé

Email

## Predictive Analytics provides



*Industry Leader – Predictive Analytics*

- **Flexibility**

- Supports wide range of forecasting techniques and forecasting requirements
- i.e. MSU to MIPS ratios change, so models change

- **Accuracy**

- From simple to rich models – select best practical fit

- **Productivity**

- Automated creation and selection of best fitting model
  - For each individual forecasted item (target)
- Scheduled automatic forecasting and reports – users see forecast results without going into each LPAR
- Process unlimited number of targets within one stream

**Accurate and Productive Modeling Requires Flexibility and Automation**

## IBM Capacity Analytics - Solution value

Analytics that listen, measure and analyze IT system performance to more effectively:

### Leverage Assets



Lower Total Cost of Ownership for Capacity Management  
Leveraging current products and capacity

### Balance Workload



Plan future based on historical performance  
Understand what-if scenarios  
Balance workload based on business targets

### Improve Customer Experience



Respond quickly with accurate, timely and relevant insight into **system requests** to ensure a **consistent experience**

## IBM Capacity Analytics – Extended Enterprise Vision

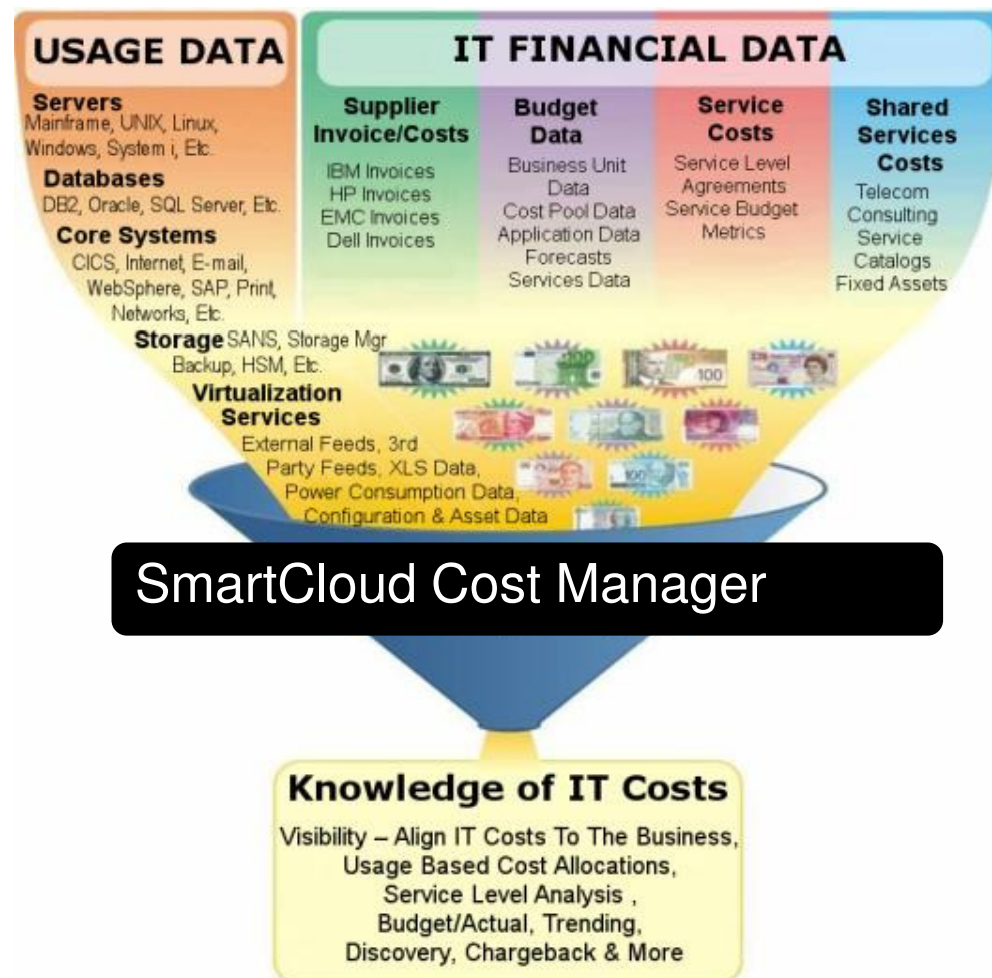
Capacity Analytics is one of the building blocks of Business Analytics at Enterprise level





# Exploit accounting to see cost impact from Capacity Management activities

Know what IT Costs with TDSz and SmartCloud Cost Management for System z





## IBM can help today



- IBM Experts visit to understand requirements
- However big or small IT, DON'T ignore capacity management!
- IBM Capacity Analytics fundamental building block to meet business demand cost effectively
- Let IBM help in System z journey to build:
  - Robust Collection Engine (TDSz)
  - Industry leader Predictive Engine (SPSS)
  - Rich, Flexible Reporting Engine (Cognos)
- IBM strength does not stop there
  - Breadth of Solution unmatched by competitors
  - Adds more business value to System z

## IBM investing in Capacity Analytics and can assist and support customers

<b>Nigel Bland</b>	<b>zTivoli Sales</b>	<a href="mailto:nigel_bland@au1.ibm.com">nigel_bland@au1.ibm.com</a>
<b>Steve Talbot-Walsh</b>	<b>z Client Architect</b>	<a href="mailto:stwalsh@au1.ibm.com">stwalsh@au1.ibm.com</a>
<b>Catherine Hawkins</b>	<b>z Client Technical Manager</b>	<a href="mailto:chawkins@au1.ibm.com">chawkins@au1.ibm.com</a>
<b>David Rintoul</b>	<b>zTivoli Client Tech Specialist</b>	<a href="mailto:david.rintoul@au1.ibm.com">david.rintoul@au1.ibm.com</a>
<b>Alan Morgan</b>	<b>BA Sales Manager</b>	<a href="mailto:alanmorgan@au1.ibm.com">alanmorgan@au1.ibm.com</a>
<b>Jason Burke</b>	<b>BA Client Technical Manager</b>	<a href="mailto:jason.burke@au1.ibm.com">jason.burke@au1.ibm.com</a>
<b>Raj Munusamy</b>	<b>zTivoli WW Product Manager</b>	<a href="mailto:rajkumar.m@uk.ibm.com">rajkumar.m@uk.ibm.com</a>
<b>Eric Chicha</b>	<b>zBA WW CTP Leader</b>	<a href="mailto:eric.chicha@fr.ibm.com">eric.chicha@fr.ibm.com</a>
<b>Anita Cox</b>	<b>z Demand Programs Professional</b>	<a href="mailto:anita.cox@au1.ibm.com">anita.cox@au1.ibm.com</a>

---

## Further Information

IBM developerWorks portal TDS for z/OS

<http://www.ibm.com/developerworks/wikis/display/tivolidoccentral/Tivoli+Decision+Support+for+zOS>

IBM developerWorks portal TCR (articles, video's, message boards,...)

<https://www.ibm.com/developerworks/community/groups/service/html/communityview?communityUuid=9caf63c9-15a1-4a03-96b3-8fc700f3a364>

IBM Cognos 8 portal (Guides for admin, report creation, framework...)

<http://publib.boulder.ibm.com/infocenter/c8bi/v8r4m0/index.jsp?topic=/com.ibm.swg.im.cognos.c8bi.doc/welcome.html>

IBM Cognos 10 support/documentation site

[http://www-947.ibm.com/support/entry/portal/documentation/software/cognos/cognos\\_business\\_intelligence](http://www-947.ibm.com/support/entry/portal/documentation/software/cognos/cognos_business_intelligence)

IBM SPSS predictive analytics: Optimizing decisions at the point of impact

<http://w3.itso.ibm.com/abstracts/redp4710.html?Open>

TCR support site

[http://www-947.ibm.com/support/entry/portal/Overview/Software/Tivoli/Tivoli\\_Common\\_Reporting](http://www-947.ibm.com/support/entry/portal/Overview/Software/Tivoli/Tivoli_Common_Reporting)

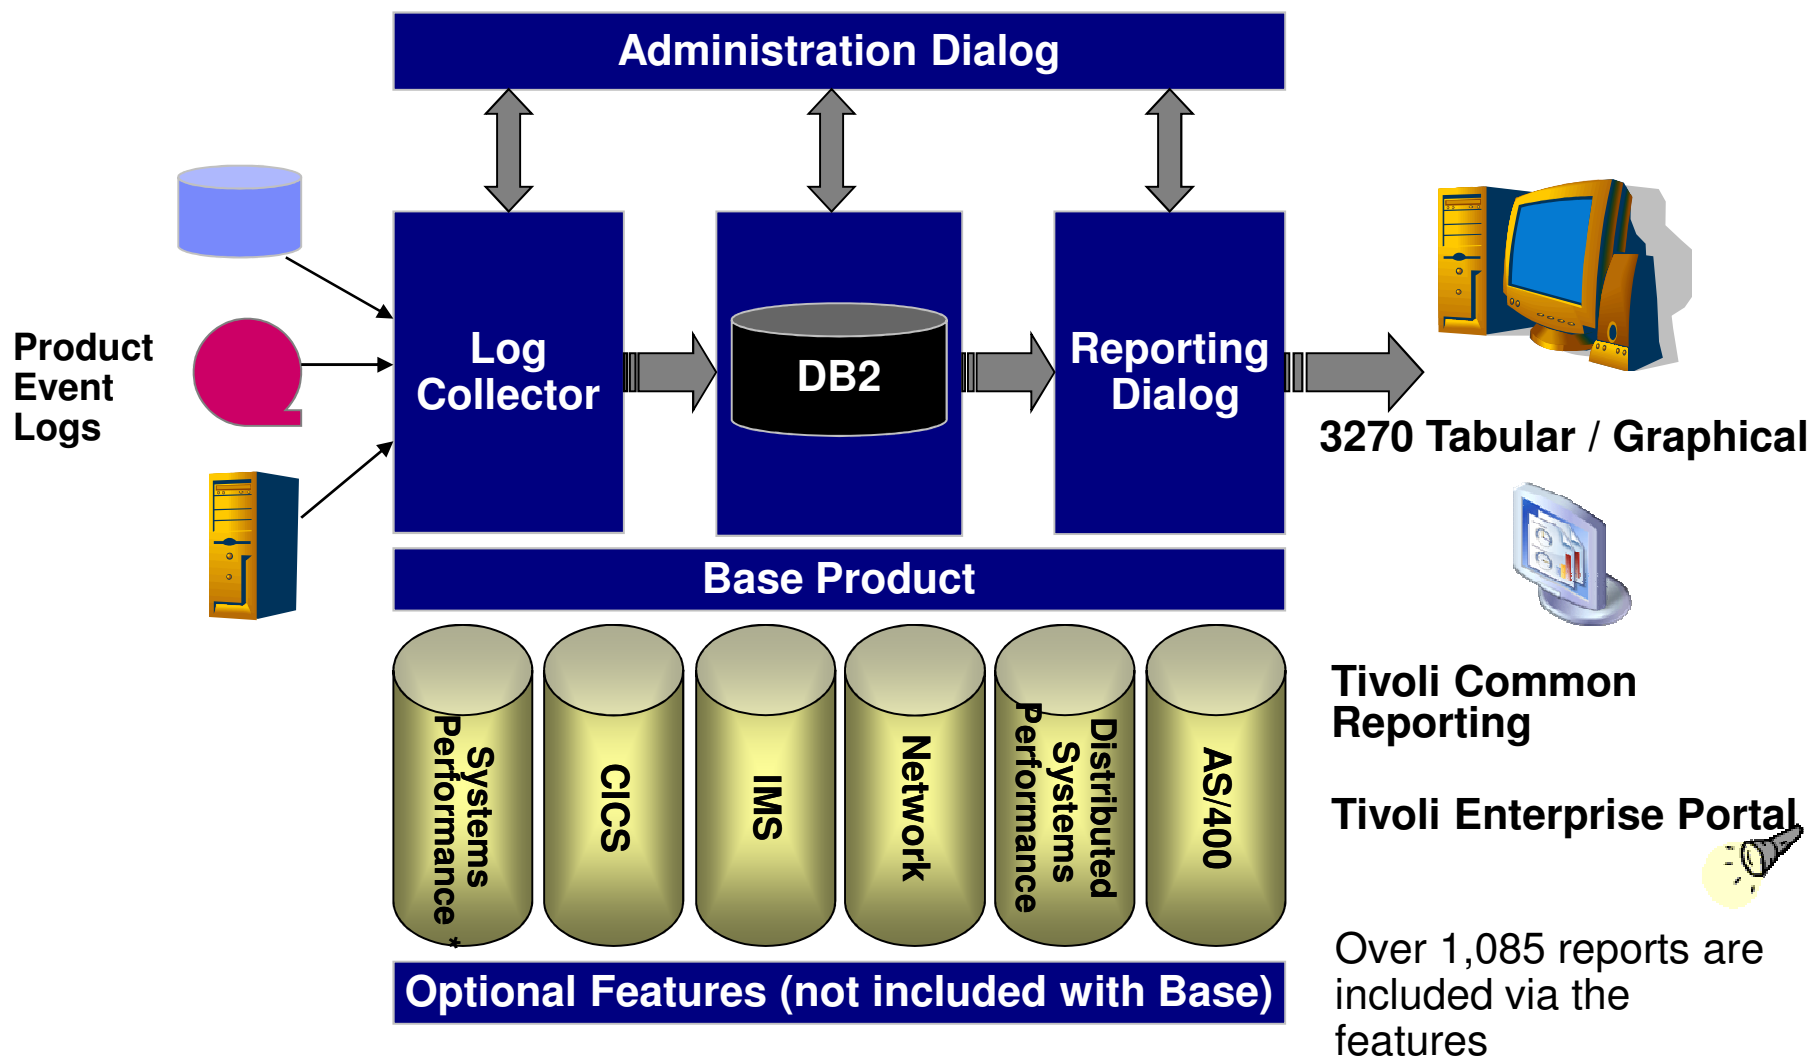
# DEMO



## Capacity Management questions:

- **System/Workload Characteristics, Performance and Trending**
  - How is my environment performing overall?
  - Which are my most used servers/LPARs for a given resource type?
  - Are there any bottlenecks in my current environment and where?
  - Am I reaching capacity on resources and which resource? When will I exhaust capacity?
  - Which is my top resource consumers for a given resource type?
  - Which are my least used servers/LPARs for a given resource type?
  - Which are my bottom resource consumers for a given resource type?
  - Do I have any outstanding abnormal behavior this week compared to last week (other periods can be used)?
  - Are my systems/workloads balanced or unbalanced?
  
- **System/Workload Estimation and Optimization (optimize and keep optimized – what if)**
  - How many more VMs can I add to a cluster/server based on usage history?
  - How much more resources do I need to add additional VMs to environment?
  - How, where do I add capacity if existing systems are not enough for future growth for optimized capacity usage?
  - Where do I place new workloads? Do I really need to add more resources?
  - How can I optimize the VM/LPARs placement to maximize usage and minimize costs?
  - How can I optimize the app placement to maximize usage and minimize costs?

# Tivoli Decision Support for z/OS



# An example: Processors view / Drill through

IBM Cognos Viewer - Processor Complexes

Business Analytics Capacity Management  
Cognos software

DATE: **Aug 23, 2011**

CPU SERIAL #	CPU MODEL #	# OF ACTIVATED ENGINES	# OF CPs (TOTAL)	# OF CPs (DED)	# OF CPs (NON-DED)	# OF zAAPs(TOTAL)	# OF zAAPs (DED)	# OF zAAPs (NON-DED)	# OF zIIPs (TOTAL)	# OF zIIPs (DED)	# OF zIIPs (NON-DED)
441E	2097	64	64	28	36	0	0	0	0	0	0
<b>31CE</b>	2097	64	62	43	31	1	0	1	1	0	1

DED = Dedicated    NON-DED = Non-Dedicated

All systems:

- Serial number
- Engines
- CPs, ZIIP, ZAAP dedicated or not
- ...

IBM Cognos Viewer - LPAR Configuration

Business Analytics Capacity Management  
Cognos software

PROCESSOR: **31CE**

LPAR NAME	LPAR #	PROCESSOR TYPE	CAPPED?	WAIT COMPLETION?	# OF LOGICAL PROCESSORS	DEDICATED PROCESSORS?	LPAR WEIGHT	CENTRAL STORAGE (MB)	EXPAN STORAGE (MB)
ICF01	1	CP	No	No	1	No	10	4,096	
STLABB7	2		No	No	10	No	5	40,960	
STLAB4C	3		No	Yes	16	Yes		227,328	
STLAB72	4		No	No	8	No	10	51,200	
COGVMLN2	5		No	No	12	No	10	92,160	
STLABF6	6		No	No	2	No	10	4,096	
SVLXCOC9	7		No	No	10	No	10	12,288	
SVLXCOT7	8		No	No	8	No	10	24,576	
COGVMLN3	9		No	Yes	8	Yes		81,920	
SVLXCOGC	10		No	Yes	8	Yes		51,200	
ICF3A	11		No	Yes	1	Yes		4,096	
STLAB6B	12		No	Yes	8	Yes		83,968	
STLAB6C	13		No	Yes	2	Yes		16,384	
STLAB6D	14		No	No	2	No	10	51,200	
<b>CP - Total</b>							<b>75</b>	<b>745472</b>	<b>1</b>
STLABF6	6	zAAP	No	No	1	No	10	4,096	

Go into more details for one system:

- LPARs
- Capping information
- Logical CPs
- ...



# Other Reports

Held Batch Jobs - IBM Cognos Viewer - Windows Internet Explorer

http://svlxcod3.svl.ibm.com:8080/cognos10/cgi-bin/cognos.cgi?b\_action=cognosViewer&ui.action=run&ui.object=%2fcontent%2ffolder%5b%40nam...

IBM Cognos Viewer - Held Batch Jobs

Date: **BATCH JOBS WITH HOLD TIMES GREATER THAN ZERO**

From: Jun 15, 2011

To: Jun 16, 2011

SYSPLEX	SYSID	JOB	JOB_NUMBER	JOB_TIMESTAMP	JOB_START_TIME	JOB_END_TIME	TOTAL ELAPSED SECS	TOTAL JOB HELD SECS
PLX4AB	SY4E	MANSEFEL1	JO802681	Jun 15, 2011 12:00:17 PM	Jun 15, 2011 12:22:01 PM	Jun 15, 2011 12:22:07 PM	1,310.57	1,296.14

1

Done

Batches

IBM Cognos Viewer - Data Collection by System

Business Analytics Capacity Management

Cognos software

Select a SYSTEM\_ID

\* STLAB4A

SYSTEM	SMF LOG DATASET NAME	SUCCESSFULLY COLLECTED?	ID OF USER THAT RAN THE COLLECT	TIMESTAMP OF DATA COLLECTION
STLAB4A	SMF.STLAB4A.G2045V00	Y	OMVSKERN	Aug 24, 2011 10:56:22 AM
STLAB4A	SMF.STLAB4A.G2044V00	Y	OMVSKERN	Aug 24, 2011 9:47:30 AM
STLAB4A	SMF.STLAB4A.G2043V00	Y	OMVSKERN	Aug 24, 2011 8:31:41 AM
STLAB4A	SMF.STLAB4A.G2042V00	Y	OMVSKERN	Aug 24, 2011 7:08:27 AM
STLAB4A	SMF.STLAB4A.G2041V00	Y	OMVSKERN	Aug 24, 2011 3:37:54 AM
STLAB4A	SMF.STLAB4A.G2040V00	Y	OMVSKERN	Aug 24, 2011 12:11:33 AM
STLAB4A	SMF.STLAB4A.G2039V00	Y	OMVSKERN	Aug 23, 2011 10:47:03 PM
STLAB4A	SMF.STLAB4A.G2038V00	Y	OMVSKERN	Aug 23, 2011 9:43:40 PM
STLAB4A	SMF.STLAB4A.G2037V00	Y	OMVSKERN	Aug 23, 2011 8:34:33 PM
STLAB4A	SMF.STLAB4A.G2036V00	Y	OMVSKERN	Aug 23, 2011 7:32:08 PM
STLAB4A	SMF.STLAB4A.G2035V00	Y	OMVSKERN	Aug 23, 2011 6:35:04 PM
STLAB4A	SMF.STLAB4A.G2034V00	Y	OMVSKERN	Aug 23, 2011 4:37:16 PM
STLAB4A	SMF.STLAB4A.G2033V00	Y	OMVSKERN	Aug 23, 2011 3:08:05 PM
STLAB4A	SMF.STLAB4A.G2032V00	Y	OMVSKERN	Aug 23, 2011 2:04:26 PM
STLAB4A	SMF.STLAB4A.G2031V00	Y	OMVSKERN	Aug 23, 2011 1:25:48 PM

Top Page up Page down Bottom

TDS data monitoring

IBM Cognos Viewer - Installed Components

Business Analytics Capacity Management

Cognos software

pReports

Component Name:

COMPONENT\_NAME PART\_NAME DESCRIPTION TIME\_INSTALLED STATUS OBJECT\_NAME OBJECT\_TYPE

CICSMON		CICS Monitoring Component	Mar 7, 2011 2:45:54 PM	I	CICS001	REPORT
CICSMON		CICS Monitoring Component	Mar 7, 2011 2:45:54 PM	I	CICS002	REPORT
14 TRAN RES		CICS Monitoring Component	Mar 7, 2011 2:45:54 PM	I	CICST01	REPORT
15 EJB		CICS Monitoring Component	Mar 7, 2011 2:45:54 PM	I	CICSJ01	REPORT
4 CMF BASIC APPL		CICS Monitoring Component	Mar 7, 2011 2:45:54 PM	I	CICS401	REPORT
		CICS Monitoring Component	Mar 7, 2011 2:45:54 PM	I	CICS402	REPORT
		CICS Monitoring Component	Mar 7, 2011 2:45:54 PM	I	CICS403	REPORT
		CICS Monitoring Component	Mar 7, 2011 2:45:54 PM	I	CICS404	REPORT
		CICS Monitoring Component	Mar 7, 2011 2:45:54 PM	I	CICS405	REPORT
		CICS Monitoring Component	Mar 7, 2011 2:45:54 PM	I	CICS406	REPORT
		CICS Monitoring Component	Mar 7, 2011 2:45:54 PM	I	CICS407	REPORT
		CICS Monitoring Component	Mar 7, 2011 2:45:54 PM	I	CICS408	REPORT
5 CMF DBCTL APPL		CICS Monitoring Component	Mar 7, 2011 2:45:54 PM	I	CICSA51	REPORT

Status: \* I

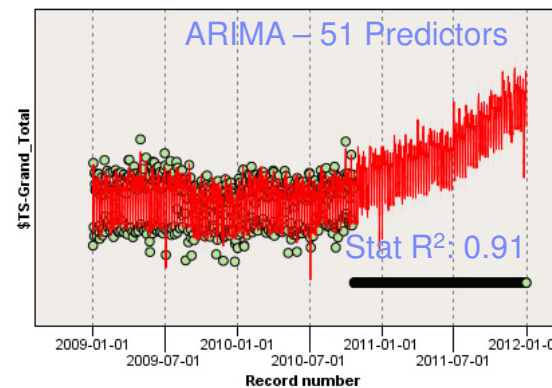

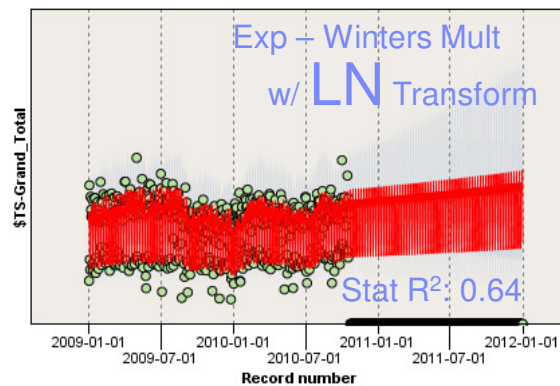



Object Type: \* REPORT

Finish

TDS components monitoring



# SPSS Various Models



**And in addition:**

► **Expert Modeller – SPSS selects the best model automatically**



# Tivoli Asset Discovery for z/OS

*Software Management for the Mainframe*

IBM Tivoli Asset Discovery for z/OS provides discovery, monitoring and reporting to understand z/OS product and application usage.

R



**Get the most for your budget**

O



**Financial accountability**

I



**Protect against losing money**

*“With Tivoli software, we can better align software usage with end-user needs while reducing costs.”*

*–Bob Becker  
Principal Information  
Technology Consultant, Farmers  
Group Inc., a subsidiary of  
Zurich Financial Services Group*

## TADz - Key value to z/OS Operations Management

- z/OS products and applications are SHARED by many users and business units.
  - **Managing this shared environment relies on educated guesswork unless you have automated tools and up-to-date knowledge bases** to continually understand z/OS software usage
  - TADz helps customers **avoid possible large revenue loses due to unexpected outages**
    - TADz shows **exactly where products & applications are deployed and which jobs/userids are using them.** This enables better software upgrade planning, change control and reduced support
    - Plan and verify **Disaster Recovery** systems have the necessary product libraries replicated to support business critical applications.
- Many z/OS customers have older SW versions and inherited systems
  - In order to **effectively manage inherited/merged environments** it is critical to understand product usage
  - **Reduce support costs and software license costs by consolidating environments**
    - Planning and smoothing migrations is greatly assisted by inspecting product usage with TADz's interactive web reporting.



## IBM DB2 Analytics Accelerator

*Do things you could never do before!*

### ■ What is it?

- A high performance appliance that integrates Netezza technology with zEnterprise technology, to deliver dramatically faster business analysis

### ■ What does it do?

- Speeds complex queries
- Lowers the cost of long term storage
- Minimizes latency
- Improves security and reduces risk
- Complements existing investments





## What's new in DB2 Analytics Accelerator V3

- High Performance Storage Saver
  - Significantly reduced the cost of storage resources
  - Option to store data only once: in the accelerator
- Incremental Update
  - Data changes are propagated to the accelerator as they happen
  - Uses Change Data Capture technology
  - Extends the accelerator use to reporting on operational data
- New optimizations
  - Tables or partitions refresh much faster and less resource intensive
  - Optimized unloading data from DB2

