QlikTech

QlikDev SAMPA

Qlikview e otimização de hardware

Bob Craig

Technical Director Expert Services

America's

BUSINESS DISCOVERY

Safe Harbor Statement

This Presentation contains forward-looking statements, including, but not limited to, statements regarding the value and effectiveness of QlikTech's products, the introduction of product enhancements or additional products and QlikTech's growth, expansion and market leadership, that involve risks, uncertainties, assumptions and other factors which, if they do not materialize or prove correct, could cause QlikTech's results to differ materially from those expressed or implied by such forward-looking statements. All statements, other than statements of historical fact, are statements that could be deemed forward-looking statements, including statements containing the words "predicts," "plan," "expects," "anticipates," "believes," "goal," "target," "estimate," "potential," "may", "will," "might," "could," and similar words. QlikTech intends all such forward-looking statements to be covered by the safe harbor provisions for forward-looking statements contained in Section 21E of the Exchange Act and the Private Securities Litigation Reform Act of 1995. Actual results may differ materially from those projected in such statements due to various factors, including but not limited to: risks and uncertainties inherent in our business; our ability to attract new customers and retain existing customers; our ability to effectively sell, service and support our products; our ability to manage our international operations; our ability to compete effectively; our ability to develop and introduce new products and add-ons or enhancements to existing products; our ability to continue to promote and maintain our brand in a cost-effective manner; our ability to manage growth; our ability to attract and retain key personnel; the scope and validity of intellectual property rights applicable to our products; adverse economic conditions in general and adverse economic conditions specifically affecting the markets in which we operate; and other risks more fully described in QlikTech's publicly available filings with the Securities and Exchange Commission. Past performance is not necessarily indicative of future results. The forward-looking statements included in this presentation represent QlikTech's views as of the date of this presentation. QlikTech anticipates that subsequent events and developments will cause its views to change. QlikTech undertakes no intention or obligation to update or revise any forward-looking statements, whether as a result of new information, future events or otherwise. These forward-looking statements should not be relied upon as representing QlikTech's views as of any date subsequent to the date of this presentation.

This Presentation should be read in conjunction with QlikTech's periodic reports filed with the SEC (SEC Information), including the disclosures therein of certain factors which may affect QlikTech's future performance. Individual statements appearing in this Presentation are intended to be read in conjunction with and in the context of the complete SEC Information documents in which they appear, rather than as stand-alone statements.

Agenda

- Overview
- Hardware and QlikView
- QlikView Server Memory and CPU
- QlikView Optimization
- QlikView in a virtual environment
- What's new in QlikView 11

Before we start:

Regarding QlikView Server and QlikView Publisher

- Always use separate servers for QlikView Server and QlikView Publisher
- Use Publisher license and separate the load and Client/Server roles, else customer will have performance degrade.
- QVS and QVP uses memory and CPU totally differently
- ***”We only load at night”*** will seldom work in the long run. Loads are often still running when users arrive in the morning causing performance degrade.

This session will only cover QlikView Server performance optimization in Front End

Agenda

- Overview
- Hardware and QlikView
- QlikView Server Memory and CPU
- QlikView Optimization
- QlikView in a virtual environment
- What's new in QlikView 11

The Chain Is No Stronger Than Its Weakest Link

- All major components must be optimized for max performance
  - Hardware need to be optimized
  - QlikView Application need to be optimized
  - QlikView Server need to be optimized
  - Windows need to be optimized

- Only then is it possible to load 1600 simultaneous users as we did at the Swedbank technical tests.

Hardware Architecture Overview

- To understand how to optimize our hardware we need to understand the components and how they work together

- The primary “moving parts” in a modern hardware architecture are
  - Memory
  - CPU
  - Inter socket connection (QPI or Hyper-Transport)
  - Memory bus
  - Storage

Modern hardware architecture design overview

- Memory is connected to CPU one (4 cores)
- Via the memory bus (in this case thee channels) integrated in the CPU
- This relationship also applies to CPU two
- An inter socket connection (QPI or Hyper-Transport) connects the CPU's together and thereby also connecting the memory.
- Each CPU have it's own link to the IO Hub (IOH)

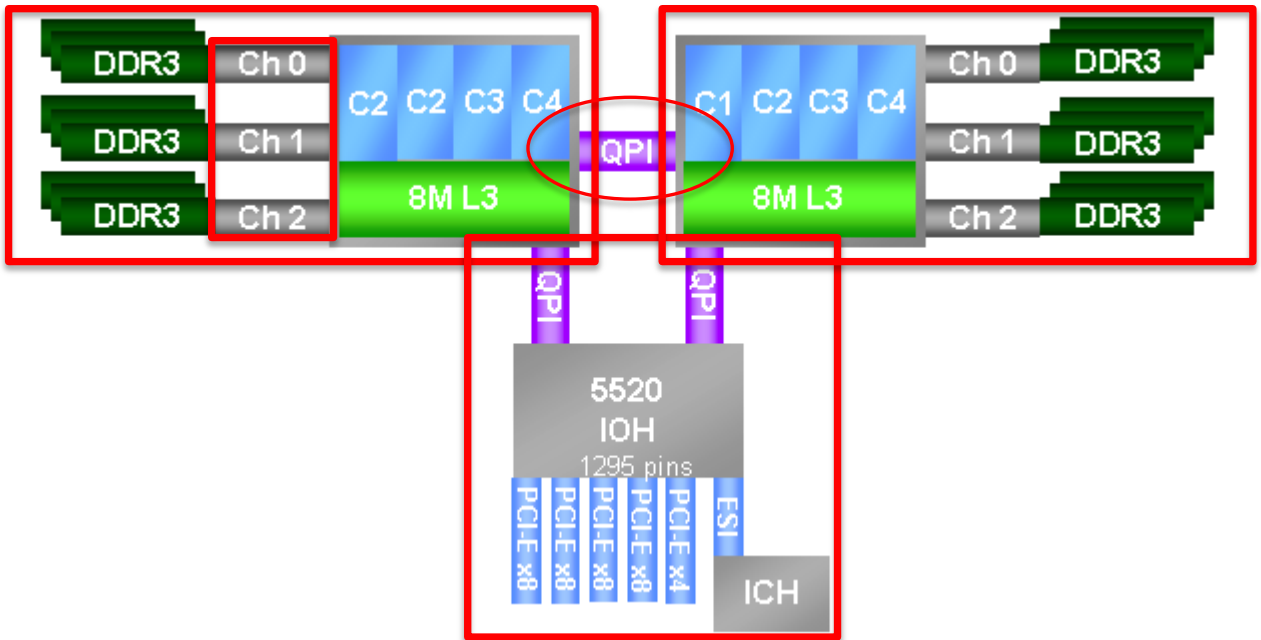

Modern hardware architecture design overview

- In a bigger system like this Nehalem EX the four CPU's need to communicate between each other and the inter socket connections then becomes complex

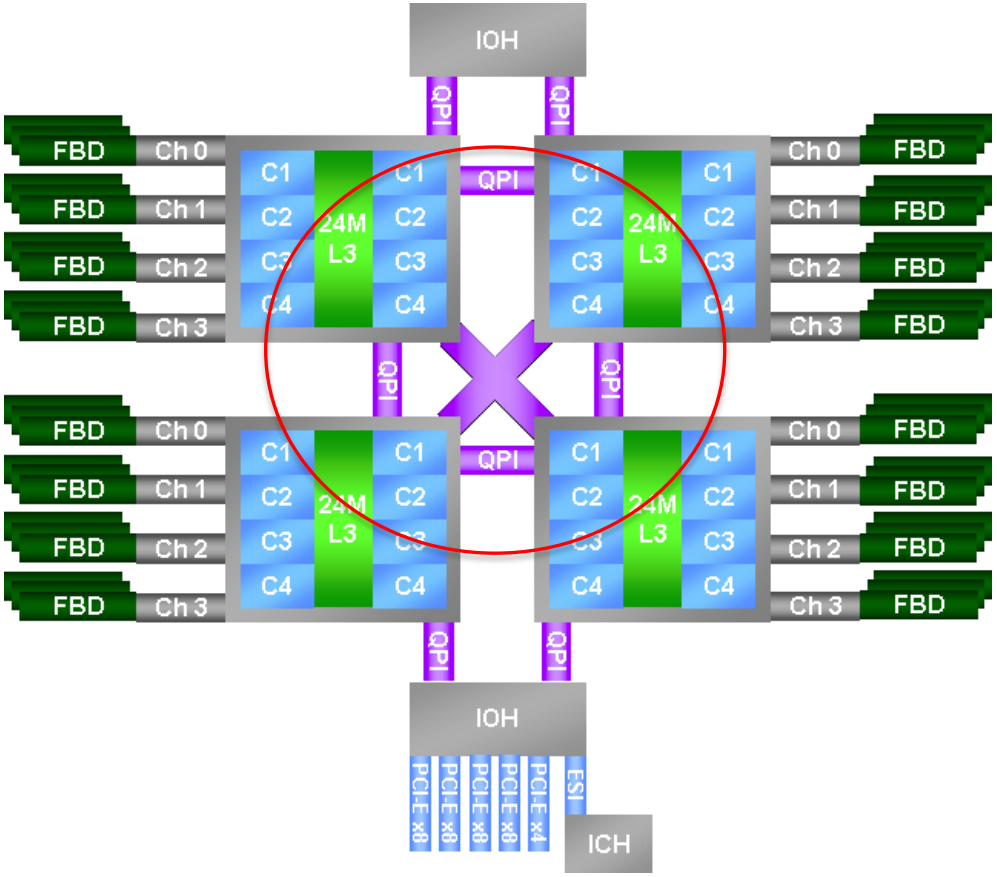

NUMA (Non-Uniform Memory Access), how does it work?

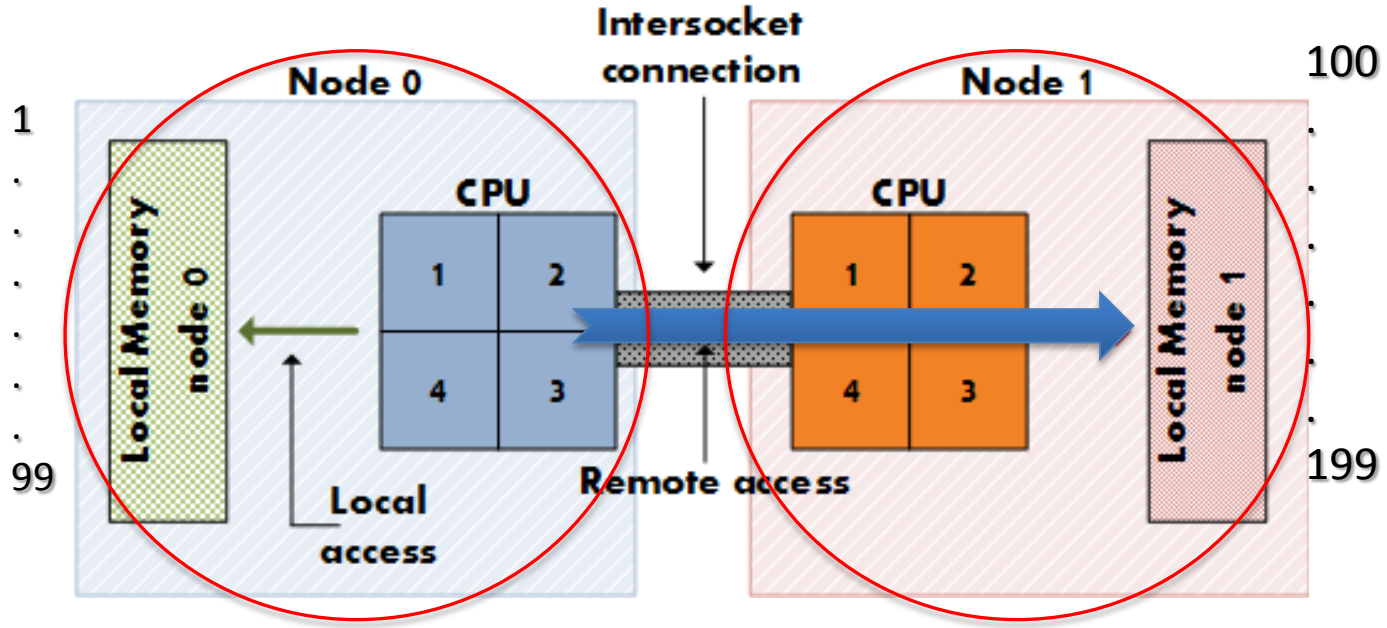

- Each processor have access its own local (and fast) memory, called a NUMA node
- Access to non-local memory (another NUMA node) will be slower
- Linear memory access (or NUMA) defines the memory on all nodes sequentially

NUMA and SUMA benchmarking

NUMA Enable
Max response time is 50 ms

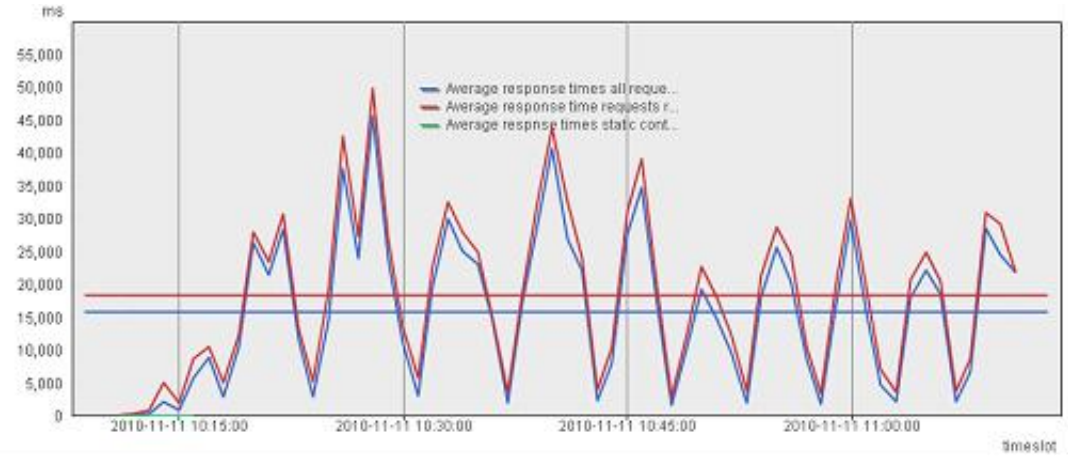

NUMA Disabled (SUMA), Max response time is now 1,8 ms

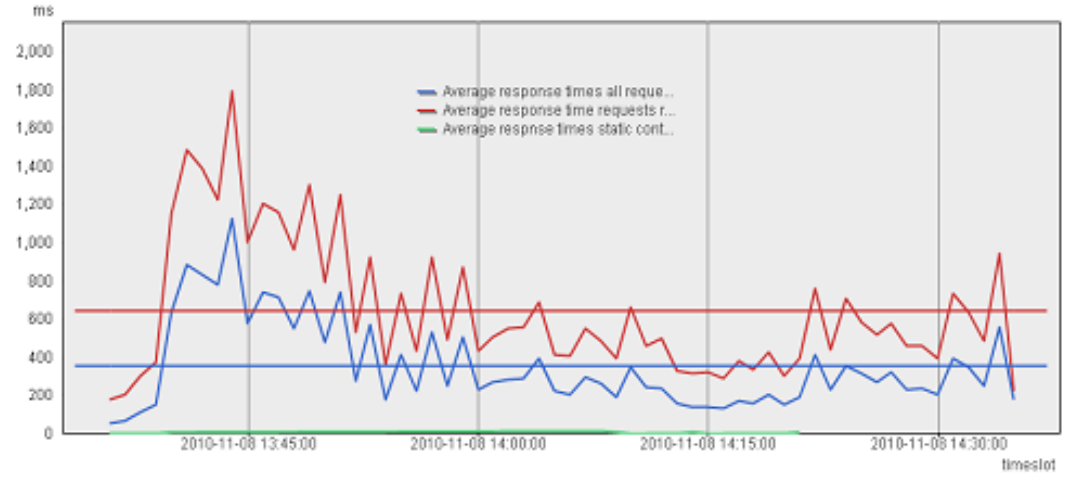

NUMA

- QlikView Server is not NUMA aware and performance will be degraded when NUMA is activated
- If NUMA is enable it's shown in Task Manager
- Windows operating system controls memory distribution between the NUMA nodes.
- This could explain the extremely bad performance results using NUMA!?

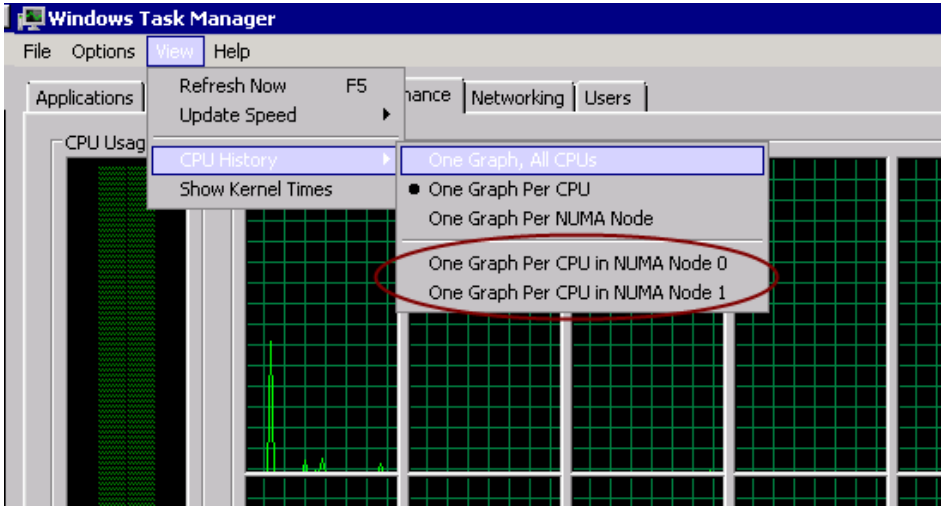

Disable NUMA

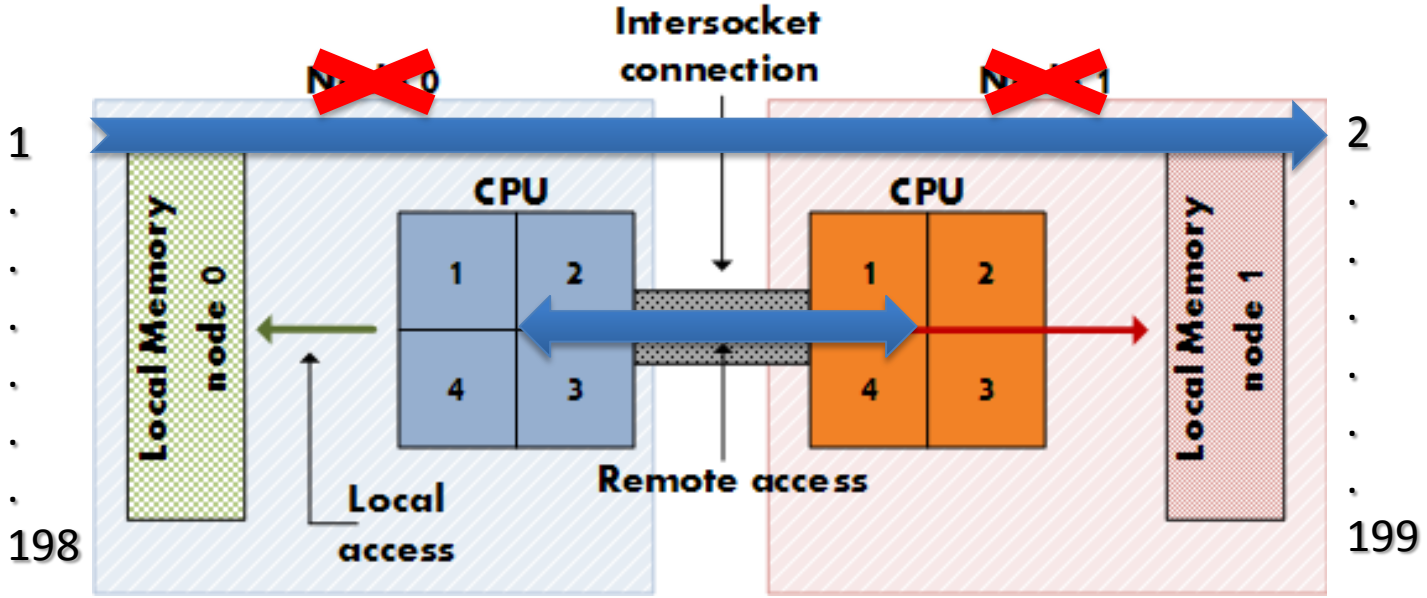

- Node interleaving (or SUMA) will address the memory across all the memory slots.
- Memory allocation and access is spread across the nodes.
- The isolated NUMA nodes will stop to exist
- Change in Bios to NUMA=Disable, Node interleaving=True or SUMA

Memory configuration

Fast memory access is extremely important for QlikView Server to get max performance

- Memory alignment: Align memory in the slots according to vendor instructions
- Memory frequency and size: Use DDR3,1333 MHz and on supported servers DDR3-1600 MHz or 1866 MHz
- Never mix different memory size and speed
- Memory bus speed: Faster CPU's will have a higher bus speed
- Hemisphere Mode: This mode allows interleaving between a processor's two memory controllers (Hemispheres) leading to improved performance on the memory buses.
  - Requires identical memory configuration across all memory controllers.
    Hemispheric mode is only available on CPU's with two memory controllers.
    Read server manual for detailed information

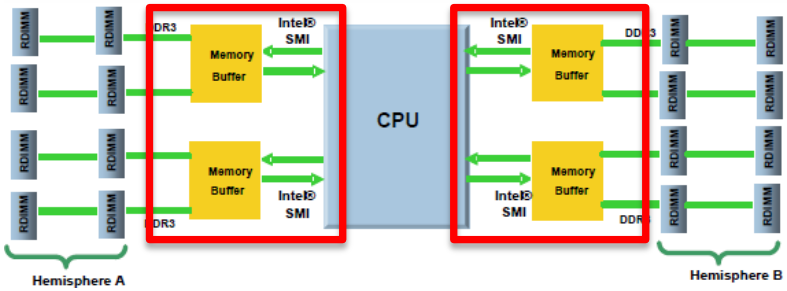

Energy saving settings

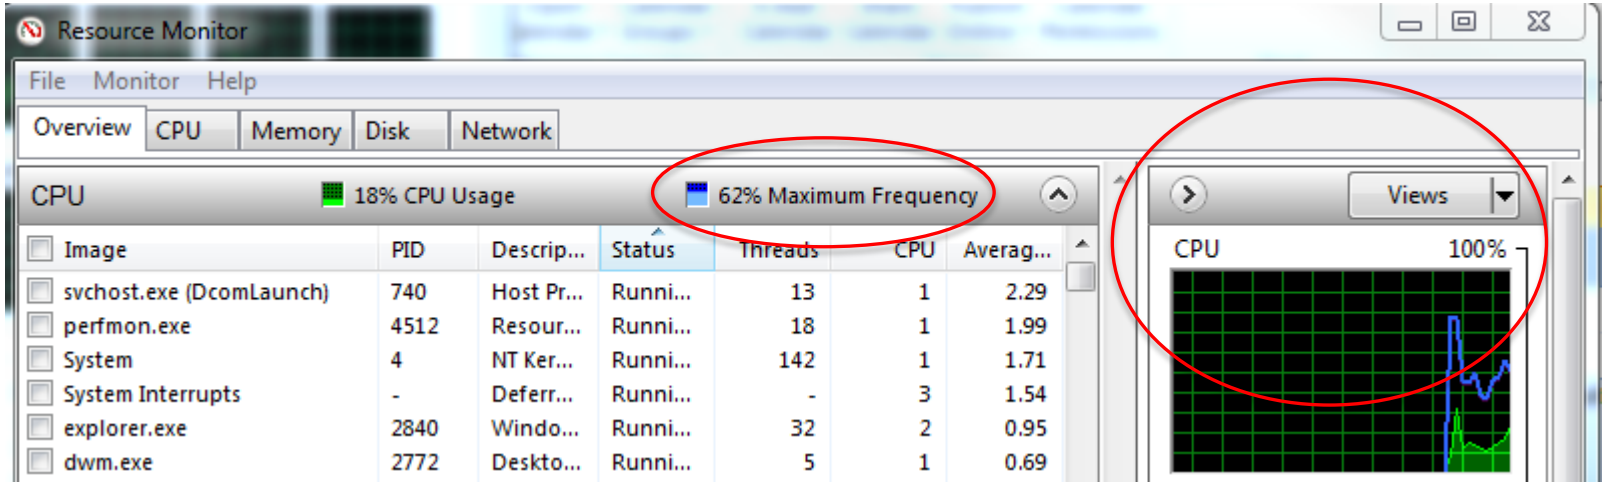

- Use Resource Monitor to identify if energy saving settings is active
- Modern processors only work at 60% efficiency when energy saving is turned on
- By activating Turbo Boost you will gain 3% to 9% extra CPU cycles
- Energy and Turbo boost settings are changed in Bios
- Disadvantage is the electric bills

Other possible hardware tweaks

- Hyper threading
  - Two threads on each core instead of one
  - Scalibility Center have noticed a QlikView Server performance degrading when using hyper threading
  - On smaller servers and laptops hyper threading could have a positive effect.
- Hardware prefetch
  - Processors is able to prefetch extra cache lines for every memory request
  - Have no measured performance benefit for QlikView Server
  - Could have a negative effect when QlikView Server fetching memory by loading the buss with unwanted requests.

Hardware Sizing

Performance Characteristics of Platforms

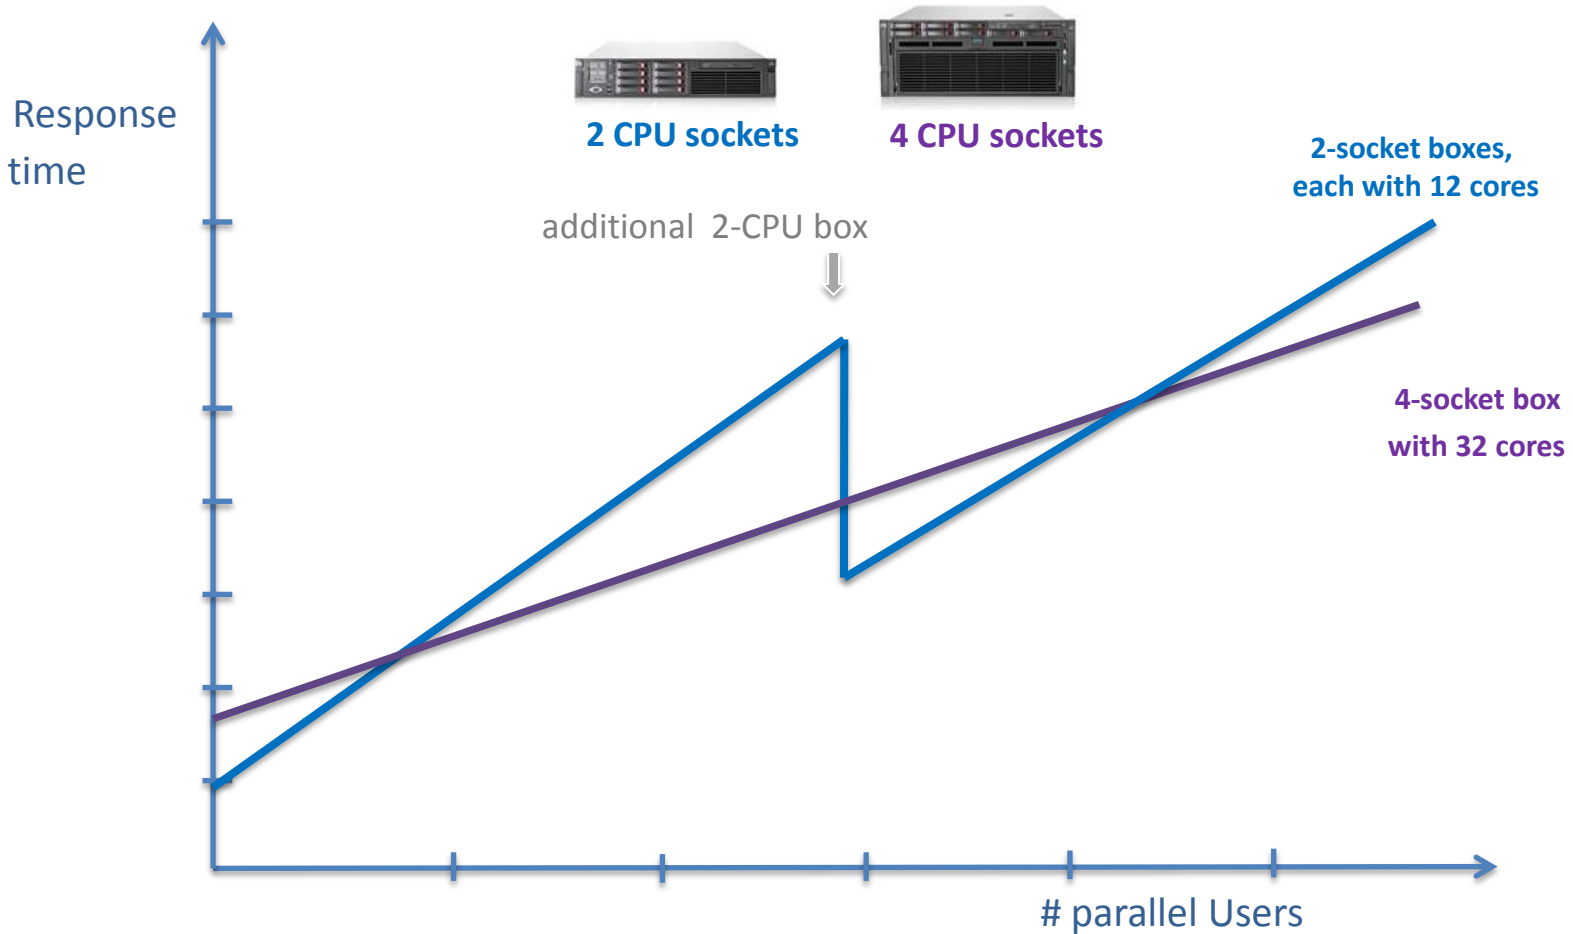

Hardware Sizing

Performance Characteristics of Platforms

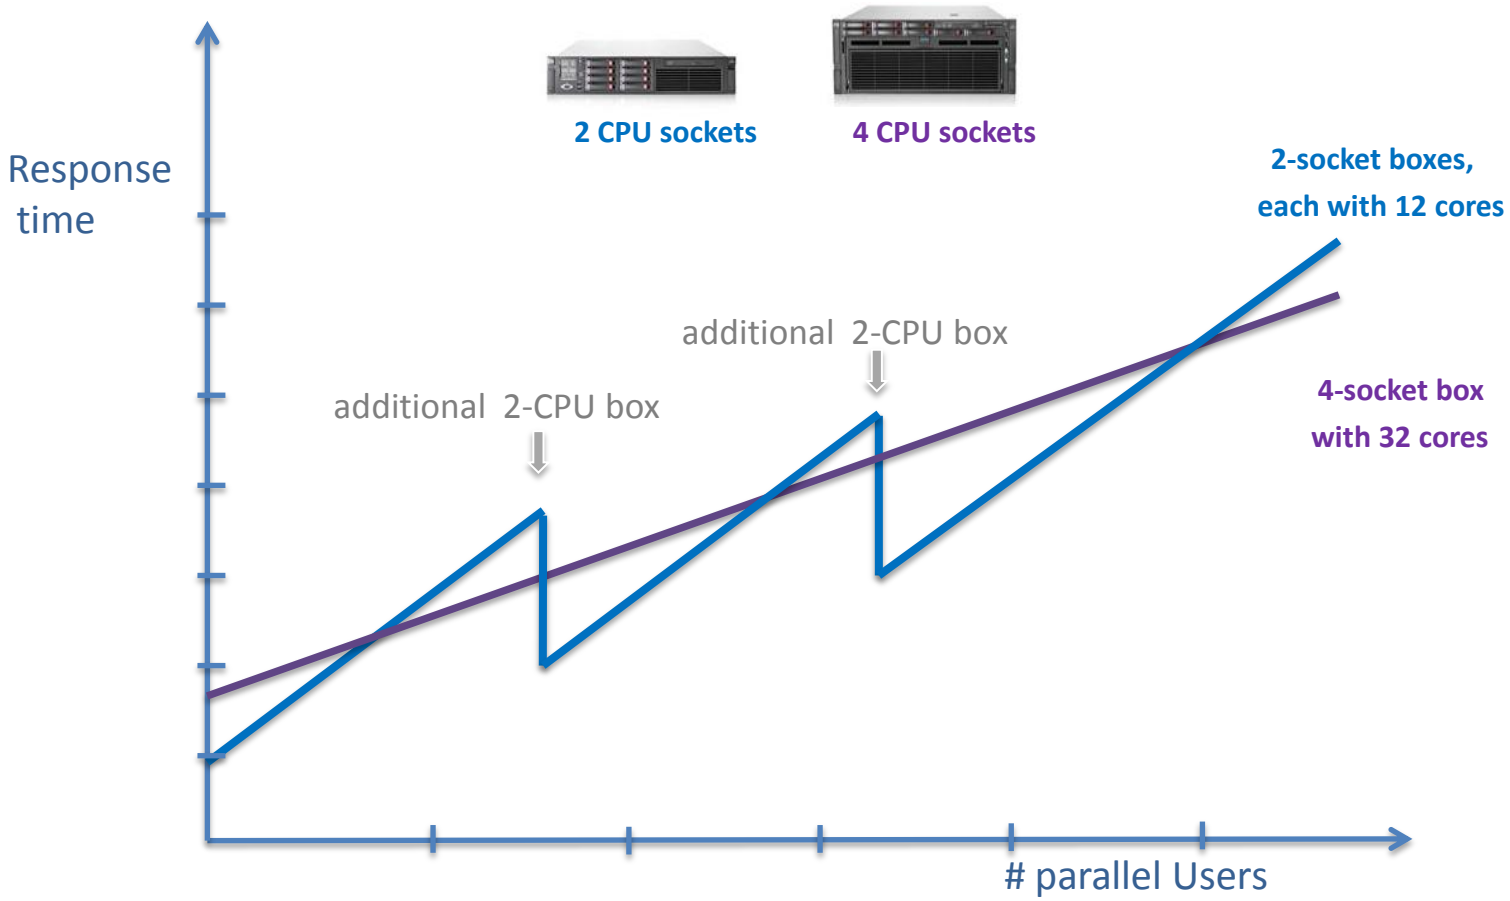

Hardware Sizing

2 –socket servers

- Advantages
  - Stepwise growth: small start invest, purchase newest technology when new box needed
  - Small user communities have best performance
  - Tasks with 1 or few threads are faster compared with 4-CPU-boxes
  - Low HW costs compared to big servers
  - More flexibility to distribute apps for analysis and reports users on different boxes
  - Redundancy for high availability

- Disadvantages
  - High SW costs (sometimes higher costs by service provider) when more servers are needed
  - Single box doesn't scale with very large user communities

Hardware Sizing

4/8 –socket servers

- Advantages
  - Due to the amount of total cores better handling of many user sessions - for large user communities
  - RAM- scalability (2 TB) => many applications , many users , big cache
  - Low SW costs (only 1 QVS license needed)
- Disadvantages
  - Expensive from the beginning
  - Lot of initial memory is needed for correct alignment (“Hemispheric Mode”)
  - No failover
  - Very big servers (8-socket via external interlink connections and mandatory NUMA mode)
    will have performance problems due to higher RAM access latencies

Servers tested by Scalibility Center

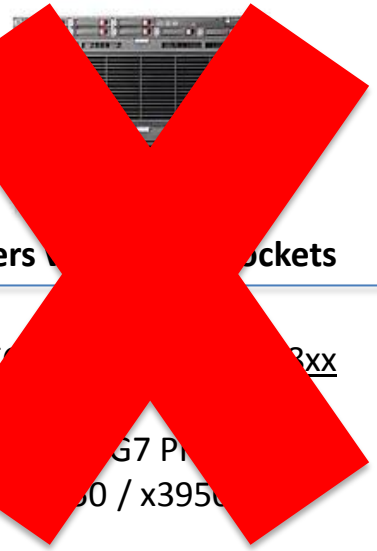

Small Servers with 2 CPU sockets

Intel XEONs X5675-5690

HP ProLiant DL380 G7
Dell R610 / R710
IBM x35nn / x36nn M3

Intel Xeon E5-2690
Sandy Bridge

HP ProLiant DL380p Gen8

~~A... ...tero...~~

HP Pr... ...5c G6

Big Servers with 4 CPU sockets

-Intel XEON X7560 and E-4870
Nehalem EX

HP DL580 G7
Dell R810 / R910 / DL980
IBM x3850 / x3950 X5

~~AMD ... 61x... ...2xx~~

~~H...~~

~~...3755 ...~~

Very Big Servers ... ...ockets

Intel XE... ...8xx

...G7 P...
IB... ...0 / x395...

Agenda

- Overview
- Hardware and QlikView
- QlikView Server Memory and CPU
- QlikView Optimization
- QlikView in a virtual environment
- What's new in QlikView 11

How does QlikView Applications use Memory?

Total RAM on the QlikView Server

RAM for Windows operating system
(approx 500 - 2000 MB)

RAM for QlikView Server Processes
(approx 30 - 100 MB)

RAM for other apps running on QlikView Server
(not recommended)

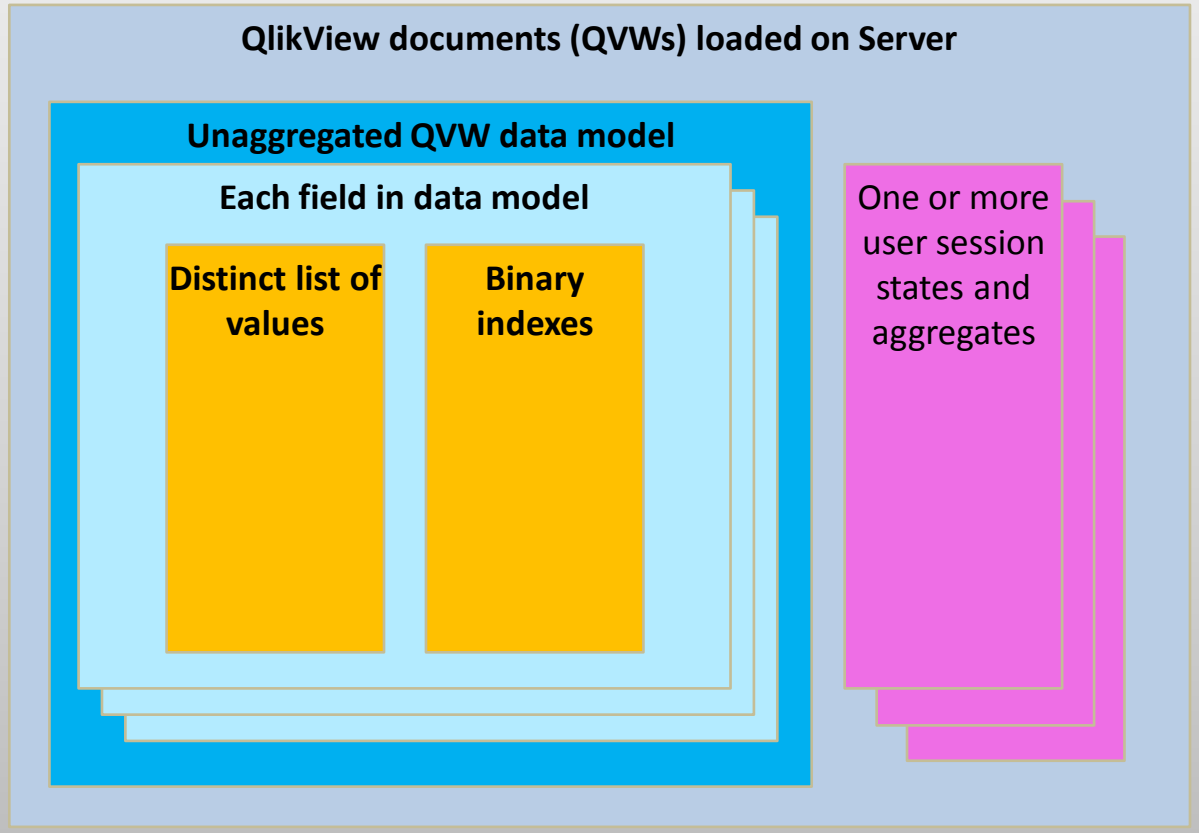

How does QlikView Server use Memory?

- **QlikView documents in memory**
  The **blue boxes** represent QlikView documents loaded into memory.
- **User Session State memory**
  Additional memory to keep user session information.
- **Cache memory**
  QlikView Server will use all available memory up to ***Working Set Low*** for calculation caching,
  this to reduce CPU usage and thereby gain speed.
- **Memory Swap point**
  When reaching memory swap point, ***Working set High*** plus having no cache left to swap out.
  QlikView will begin to swap memory on to disk resulting in performance degradation.

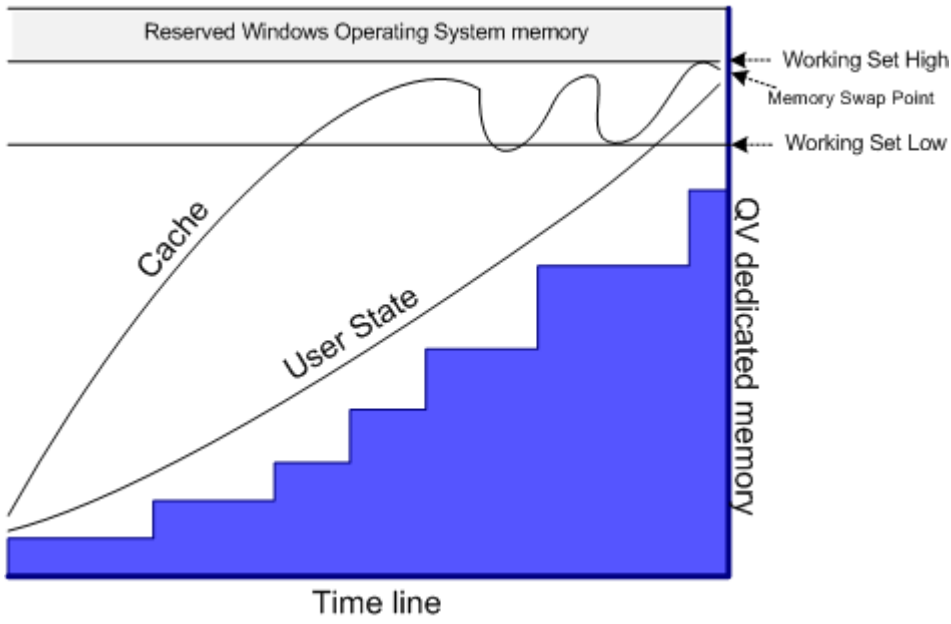

QlikView Server Cache Design

- QlikView uses an MRU (Most Recently Used) list classifying the cache based on recently used and complexity.
- When QlikView memory usage gets close to Working Set High the cache will be flushed based on the MRU list.
- The MRU list engine is single treaded
- This could be a bottleneck when using the cache in an extreme way
  like automated performance tests

How does QlikView Server use CPU?

- QlikView uses CPU for calculations and memory access
- 100% CPU utilization often means that the CPU is trying to retrieve data from memory. This is usually a symptom of slow memory access (or massive amount of data)
- Sometimes it looks like the CPU is at 30% to 60% but in reality the CPU is moving lightning fast between 0% and 100%.
  This when QlikView Server is collecting data from the memory and is a healthy behavior.
- A badly optimized function or calculation could utilize only one CPU (single threaded operation). Check task manager CPU usage tab during the test cycle.

Agenda

- Overview
- Hardware and QlikView
- QlikView Server Memory and CPU
- QlikView Optimization
- QlikView in a virtual environment
- What's new in QlikView 11

Impact of data model

Relations via link tables

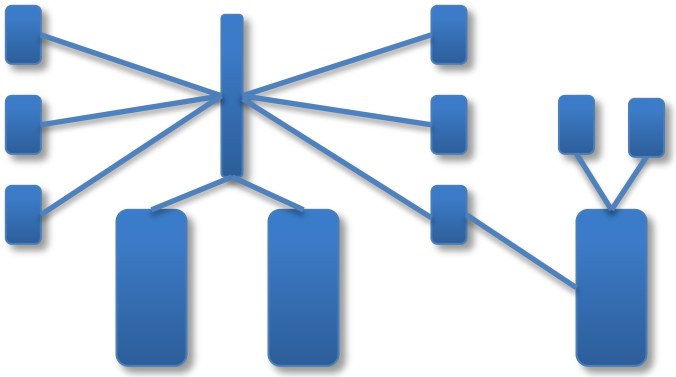

→ **Needs more CPU power**

VS.

Concatenated fact tables

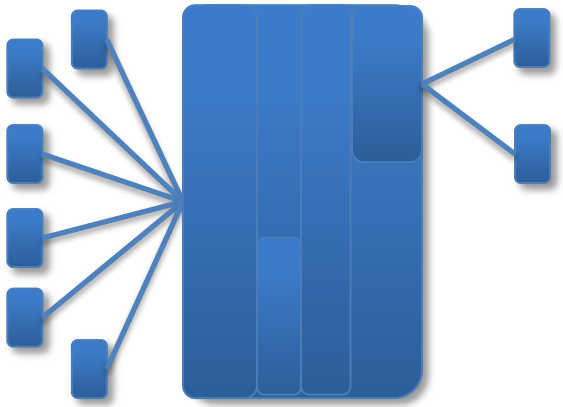

→ **Needs less CPU power but sometimes more RAM**

Impact of formula complexity

```
If....
sum( aggr(
 ( sum( {
        $<AGR.Datum={">=$(vStart_Datum_Z)<=$(vEnde_Datum)"},
          AGR.IFCNAM={$(=chr(39) & replace(replace( concat( aggr(                    if( rank(
          ( sum( {$<AGR.Datum={">=$(vStart_Datum_Z)<=$(vEnde_Datum)"} >}
AGR.ChangeRateDiff.kursv))      /
          ((
            sum( {$<AGR.Datum={">=$(vStart_Datum_N)<=$(vEnde_Datum)"} >} AGR.Stock.RateV)) /
($(#vChart_Month_Diff_N))
          )
                              ) <6, AGR.IFCNAM),AGR.IFCNAM),',',',','#,#'),'#',chr(39))&chr(39))}
                          >} AGR.ChangeRateDiff.kursv))
 ,AGR.InvObjective2))   /
 sum( aggr(
  ( sum( {
         $<AGR.Datum={">=$(vStart_Datum_N)<=$(vEnde_Datum)"},
          AGR.IFCNAM={$(=chr(39) & replace(replace( concat( aggr(                    if( rank(
          ( sum( {$<AGR.Datum={">=$(vStart_Datum_Z)<=$(vEnde_Datum)"} >}
AGR.ChangeRateDiff.kursv))
          /
          ((
            sum( {$<AGR.Datum={">=$(vStart_Datum_N)<=$(vEnde_Datum)"} >} AGR.Stock.RateV)) /
($(#vChart_Month_Diff_N))
          )
                              ) <6, AGR.IFCNAM),AGR.IFCNAM),',',',','#,#'),'#',chr(39))&chr(39))}
                          >} AGR.Stock.RateV)) / ($(#vChart_Month_Diff_N))
  )
 ,AGR.InvObjective2)
```

vs.

```
sum( {< Year={$(vPY)} >}   Sales.Amount  )
```

Impact of formula amount

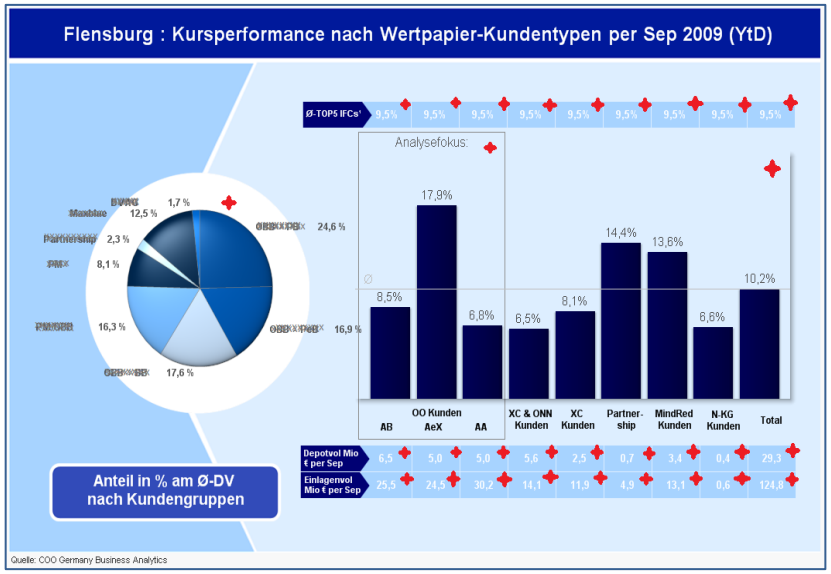

vs.

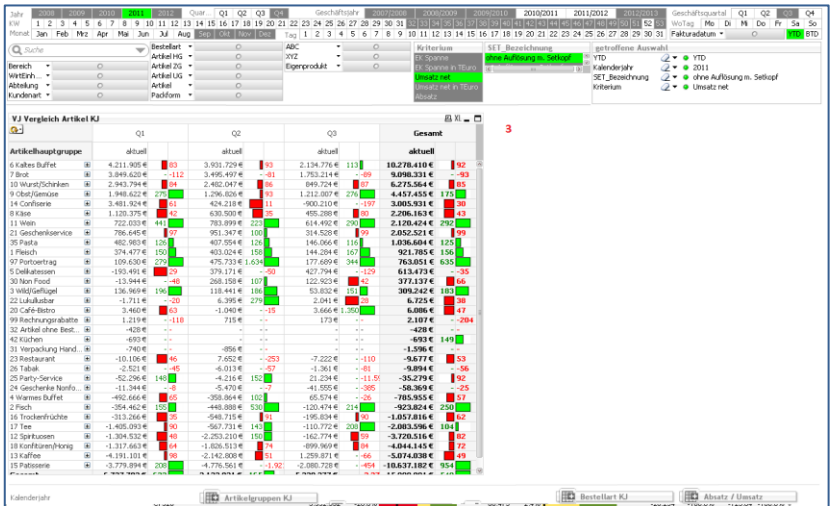

➡ **30** expressions in 30 objects

Max: 350 Expressions (incl. condition / color formulas) in one sheet**!**

➡ **3** expressions in 1 object

Impact of user types

Report users

- Direct access to their data section (the majority of the users have limited rights)
- Less analysis need
- Frequent, periodical access to the data of their business
- „one“-dimensional views

→ 10 – 30 clicks in one session → 0,5 clicks per minute within one hour

VS.

Analysis users

- Global, wider view at the data , top down , across many fields
- Often changes attributes, changes data vizalisation , create new filter combinations
- Intense but erratic access
- „multi“-dimensional views to data

→ 30-120 clicks in one session → 2 clicks per minute within one hour

Other influencing factors

- amount of dimensions in pivot tables
- cardinality of dimensions in diagramms and also in database
- Use of certain functions such as:
  - type conversions (date,num,dual)
  - count(distinct)
  - concat()
  - aggr()
  - firstsortedvalue()
  - rank()
  - …..

RAM Consumption Calculation Example

	.QVW File Size (in MB)	RAM Footprint	Per User Footprint	# Users Concurrent	Total RAM Consumed (MB)
QVW #1	500	2,000	200	50	12,000
QVW #2	1,000	4,000	400	100	44,000
				Total:	**56,000**

File Size is roughly 8% of original data size

File Size X 4 is estimated Footprint Size

Footprint Size X 10% is estimated Per User Footprint

For exact RAM Footprint save application with *Compression = None*

Application Architecture

Split big QVWs also in additional, different smaller ones!

Analysis Users

Few users with
heavy system load!

Report Users
Dashboard Users

Many users with
consistent system load!

Application Architecture

Scenario:
You have 800 million rows of data and a total user audience of 400 users.
A max concurrency of around 10%, gives you 40 max users at any given time.
1 QlikView application has been identified to meet the needs.

Option #1

800 M rows
4GB file size
16 GB footprint
1.6 GB per user
avg. resp. time ~5 sec

Avg. Resp. Time	Avg. Sessions per Day	Avg. Selections per Session	Total Wait Time per Day	# QVS Needed	CPU Cores	RAM
~5 sec	400	10	**5.6 hours**	**2**	**24**	**64**

Option #2

60 M rows
.5 GB file size
2 GB footprint
.2 GB per user
avg. resp. time <1 sec

800 M rows
4GB file size
16 GB footprint
1.6 GB per user
avg. resp. time ~5 sec

	Avg. Response Time	Avg. Sessions per Day	Avg. Selections per Session	Total Wait Time per Day	# QVS Needed	CPU Cores	RAM
App 1	< 1 sec	350	10	.95 hours	-	-	-
App 2	~5 sec	50	10	.7 hours	-	-	-
				1.7 hours	**1**	**24**	**48**

Option #2 is:

- The same solution
- Less than half the hardware needed
- Average 300% improvement in response times for users

Application data model

Schema alternatives for models with multiple source fact tables

	Alternative 1	Alternative 2	Alternative 3
	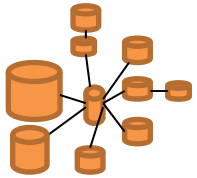	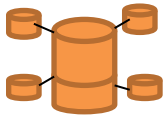	
Response Time	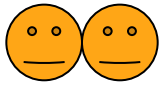	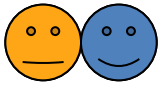	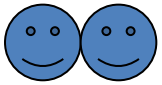
RAM consumption	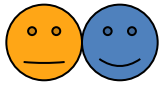	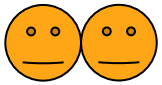	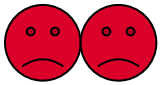
Script run time	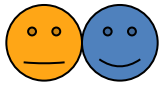	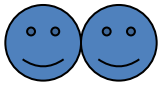	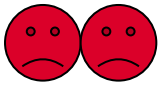
Flexibility Model		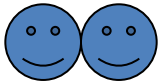	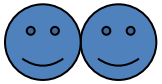
Complexity Script		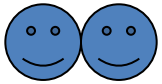	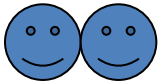

Application data model

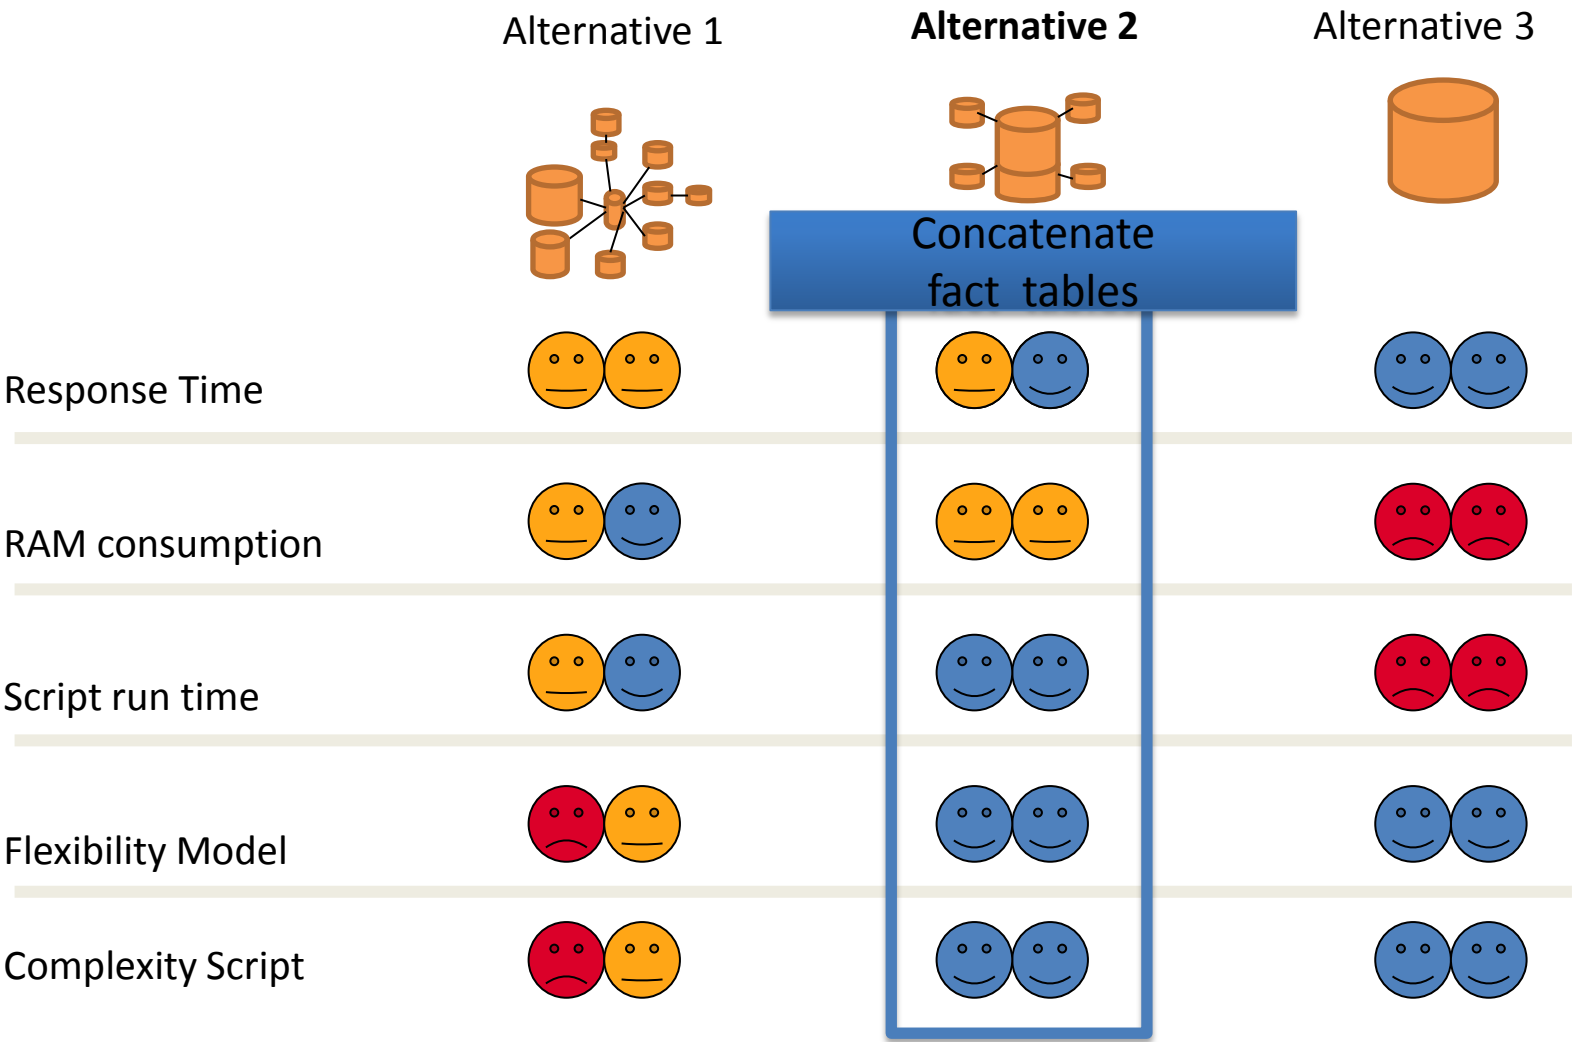

Design Alternative: Fact Table Concatenation

Sales

Region	Product	Date	Sales
Region A	P1	2009-01-31	100
Region A	P1	2009-02-28	120
Region A	P1	2009-03-31	140
Region A	P2	2009-01-31	500
Region A	P2	2009-02-28	550
Region A	P2	2009-03-31	600
Region B	P1	2009-01-31	50
Region B	P1	2009-02-28	55
Region B	P1	2009-03-31	60
Region B	P2	2009-01-31	200
Region B	P2	2009-02-28	180
Region B	P2	2009-03-31	160

Plan Yearly

Region	Date	Plan
Region A	2009-01-1	8000
Region B	2009-01-1	10000

Procurement Cost

Product	Date	Cost
P1	2009-01-31	130
P1	2009-02-28	1400
P1	2009-03-31	1600
P2	2009-01-31	500
P2	2009-02-28	650
P2	2009-03-31	600

Concatenated Facts

Region	Product	Date	Sales	Plan	Cost
Region A	P1	2009-01-31	100		
Region A	P1	2009-02-28	120		
Region A	P1	2009-03-31	140		
Region A	P2	2009-01-31	500		
Region A	P2	2009-02-28	550		
Region A	P2	2009-03-31	600		
Region B	P1	2009-01-31	50		
Region B	P1	2009-02-28	55		
Region B	P1	2009-03-31	60		
Region B	P2	2009-01-31	200		
Region B	P2	2009-02-28	180		
Region B	P2	2009-03-31	160		
Region A		2009-01-1		8000	
Region B		2009-01-1		10000	
	P1	2009-01-31			130
	P1	2009-02-28			1400
	P1	2009-03-31			1600
	P2	2009-01-31			500
	P2	2009-02-28			650
	P2	2009-03-31			600

Agenda

- Overview
- Hardware and QlikView
- QlikView Server Memory and CPU
- QlikView Optimization
- QlikView in a virtual environment
- What's new in QlikView 11

Does QlikView Server work in a virtual environment?

Yes and No

YES, QlikView Server is certified for Vm-Ware

- This means that QlikView Server works 100% in a virtual environment
- QlikView have the same support on VM as in physical environments
- New VM-Ware 5 release is better to handle memory and CPU
- QlikView Publisher works fine in a virtual environment
- Cloud computing is based on the same idea

NO, there are problems with QlikView Server in VM

- The VM host is a extra layer between Windows and hardware causing (10-40%) performance degrade.
- Memory access performance degrade, that is important for QlikView Server
- Memory ballooning will degrade QlikView Performance, Always turn off
- Sharing resources with other systems, try to get dedicated resources
- System Administrator can (and often will) reduce resources on the fly.

Agenda

- Overview
- Hardware and QlikView
- QlikView Server Memory and CPU
- QlikView Optimization
- QlikView in a virtual environment
- What's new in QlikView 11

What's new in QlikView 11 regarding performance?

- Memory allocation algorithm change
  QlikView now commits smaller memory blocks when allocating memory. Performance tests have shown that allocating bigger blocks takes time.
- Special NUMA disable tweak for systems that can't turn off NUMA.
  NumaRoundRobin=1 in c:\ProgramData\QlikTech\QlikView Server\Settings.ini
- Enhanced load balancing algorithm (CPU with RAM overload) for QlikView Server clusters
- QlikView Server performance has been overall improved in every QlikView version and SR, this is a result of ongoing Scalibility Center testing.

Questions?

QlikTech

III QlikDev SAMPA

Obrigado! Thank you!
Gracias! Tack!

BUSINESS DISCOVERY