zEnterprise Hybrid Smarter-Planet Application Architecture and Smarter-Server Design

Monte Bauman
IBM Columbus
mbauman@us.ibm.com

Monday, August 08, 2011

9557

Topics

- Requirements
- Design
  - zEnterprise
  - z10 vs. z196
  - zManager
- Use Cases
- Conclusion
- Next Steps

Topics

Requirements

Topics

It's About the Business…
(and the Business Process that we automate)

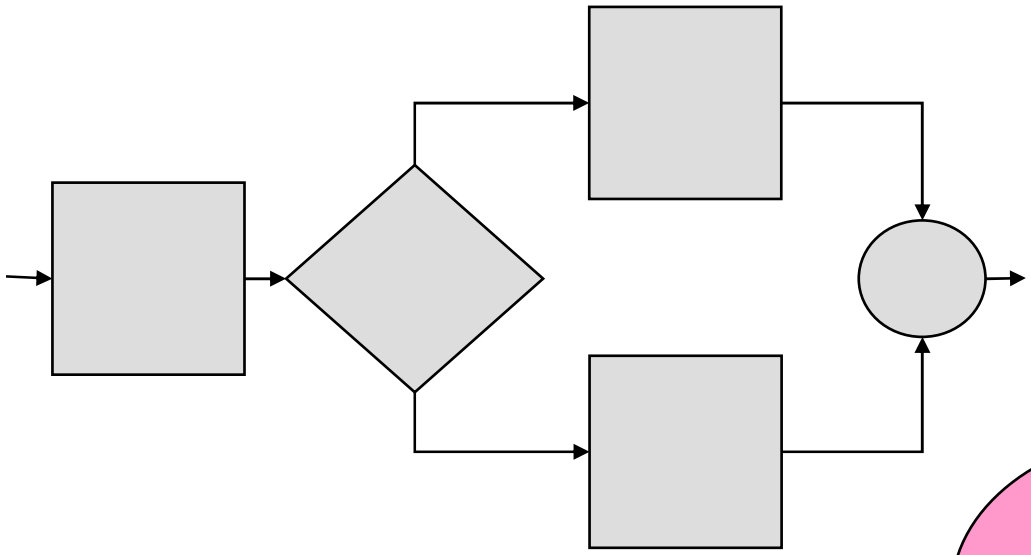

It is ALL about the transaction!

Topics

It's About the Business Process…
(indeed … lots of business processes)

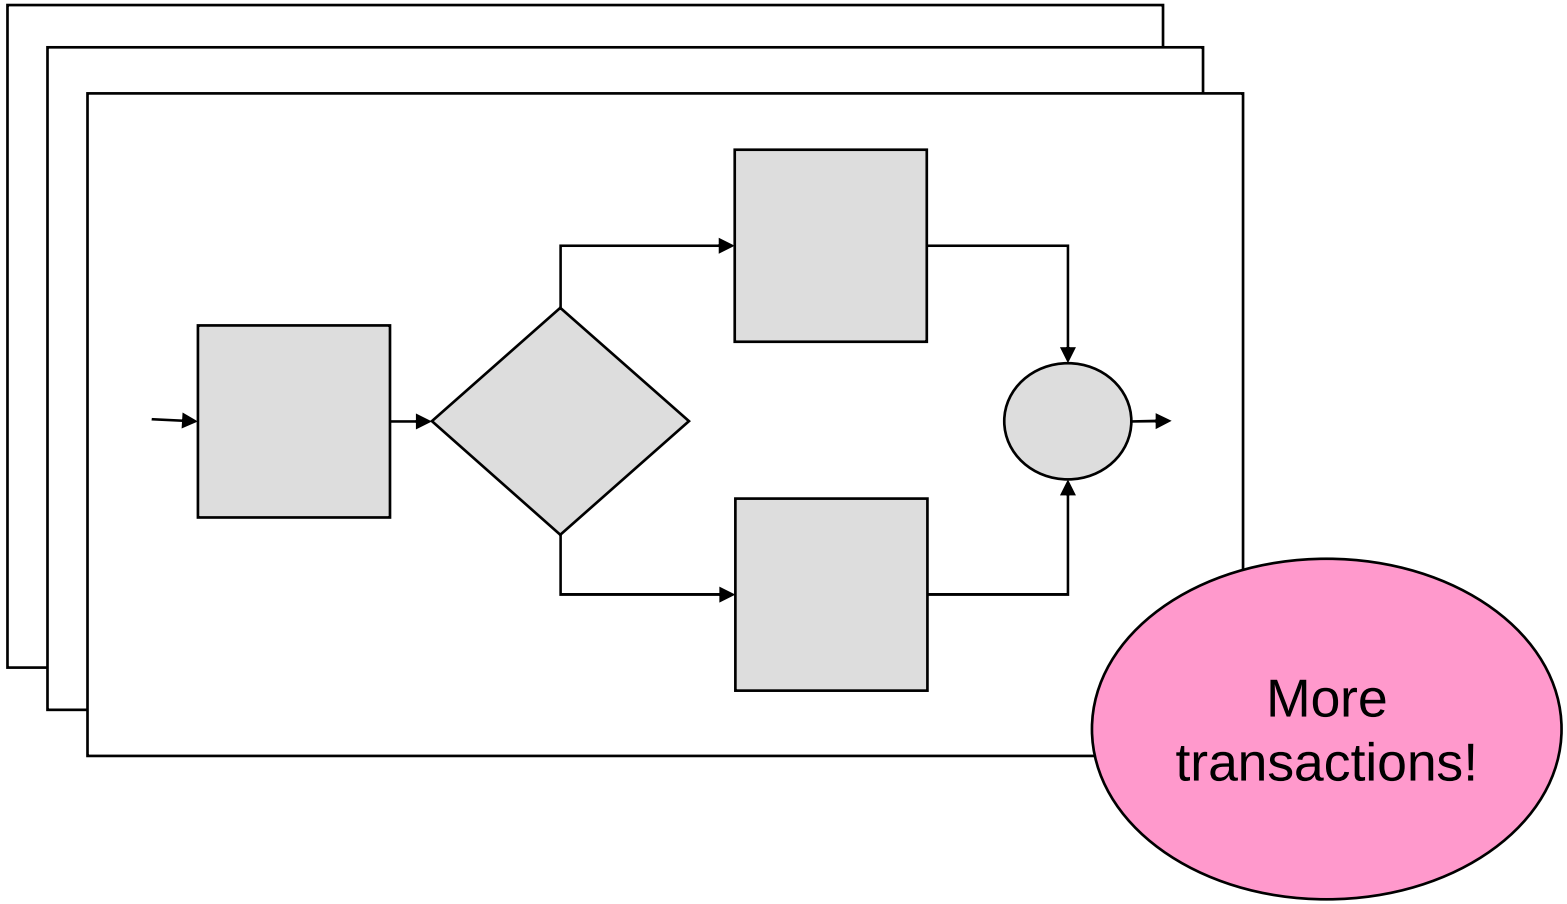

It's About the Business Process

(indeed … lots of inter-related (SOA-ish) business processes)

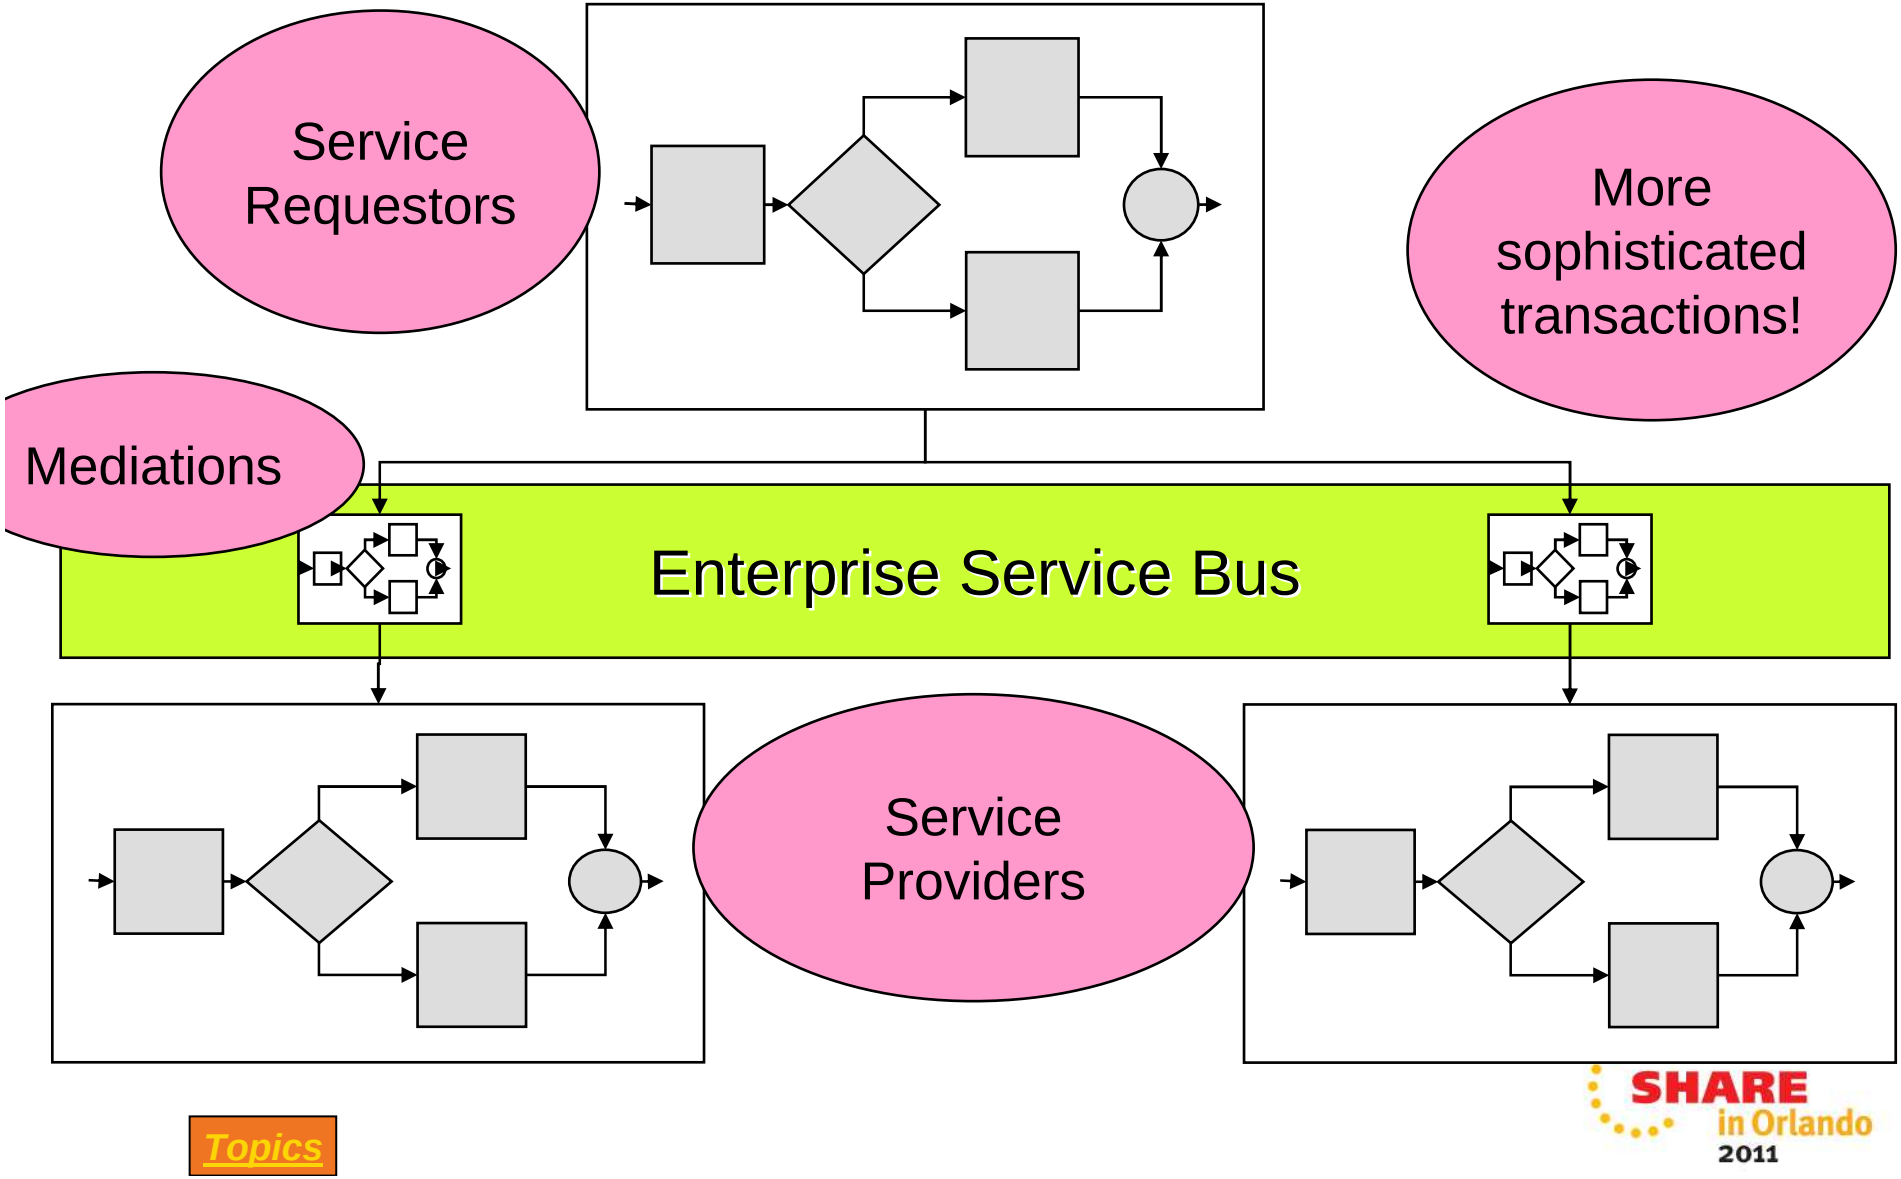

Topics

It’s About Automating/Enabling the Business Process (using repeatable architectural patterns)

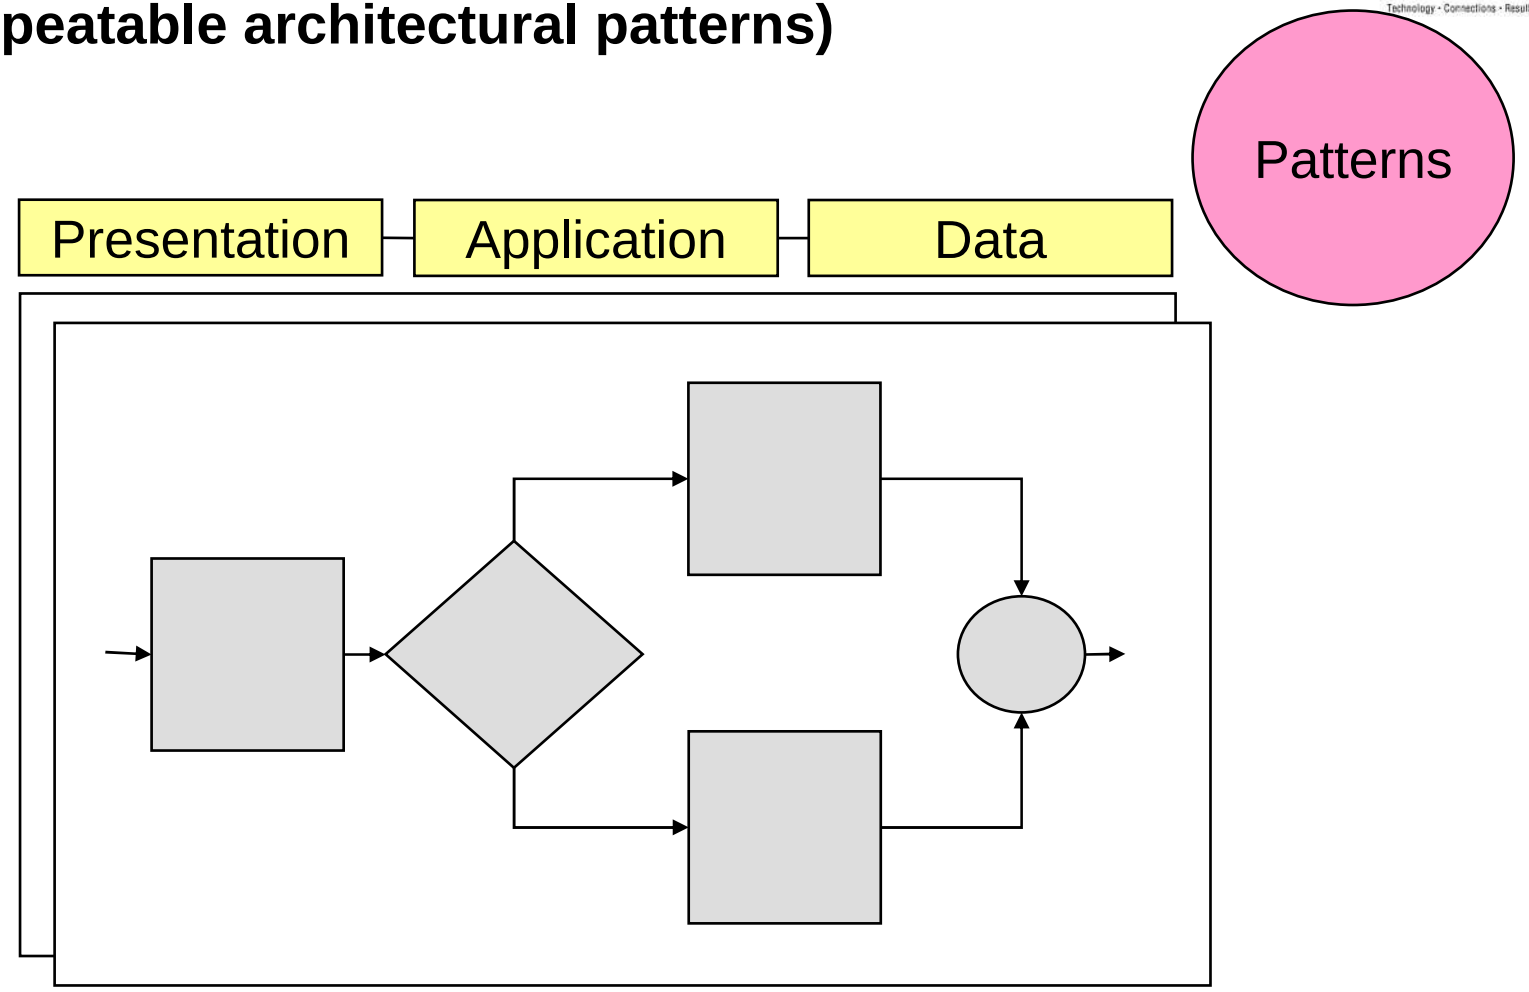

Topics

It's About Building/Deploying the Business Process (the application development lifecycle)

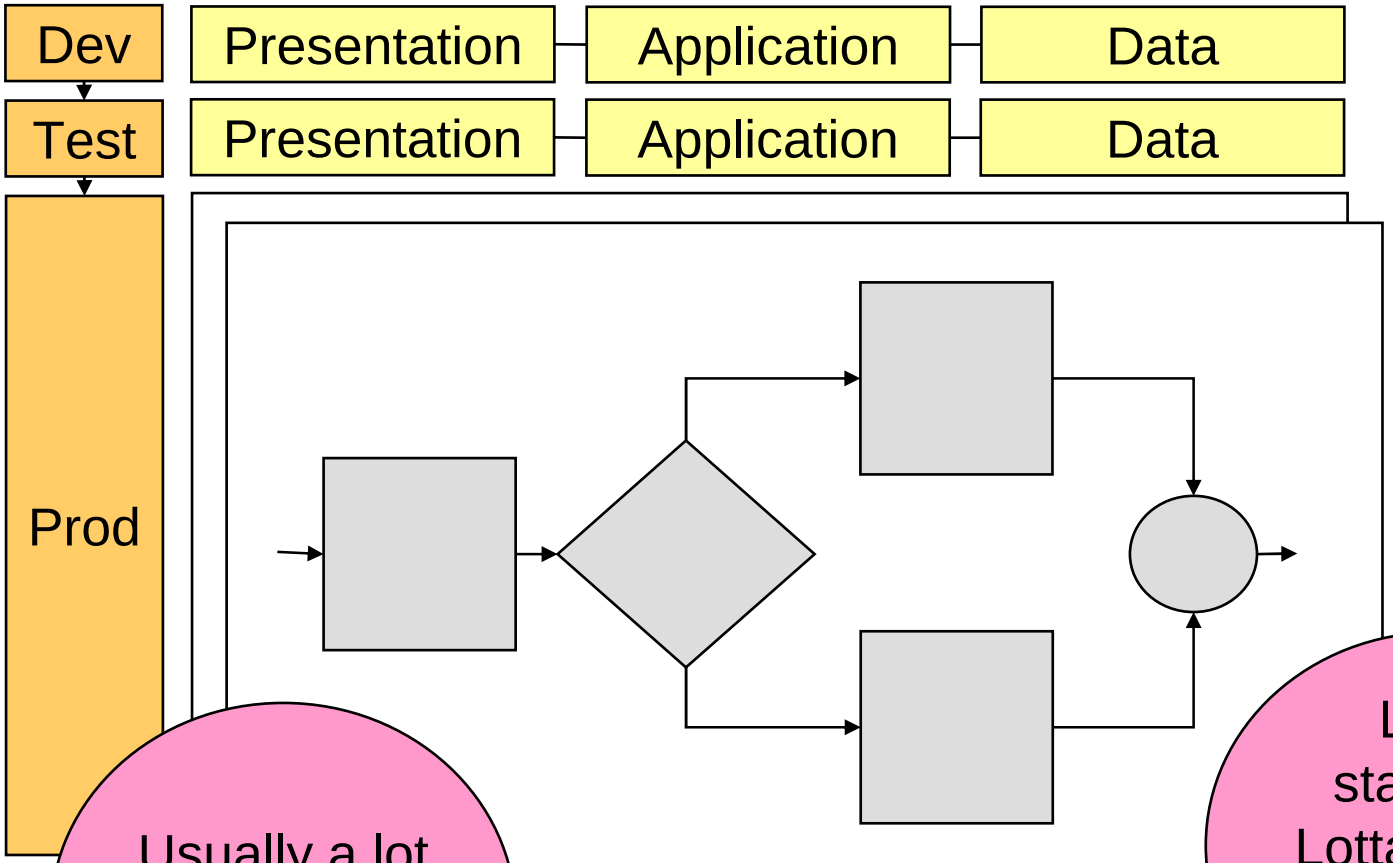

It’s About the Data

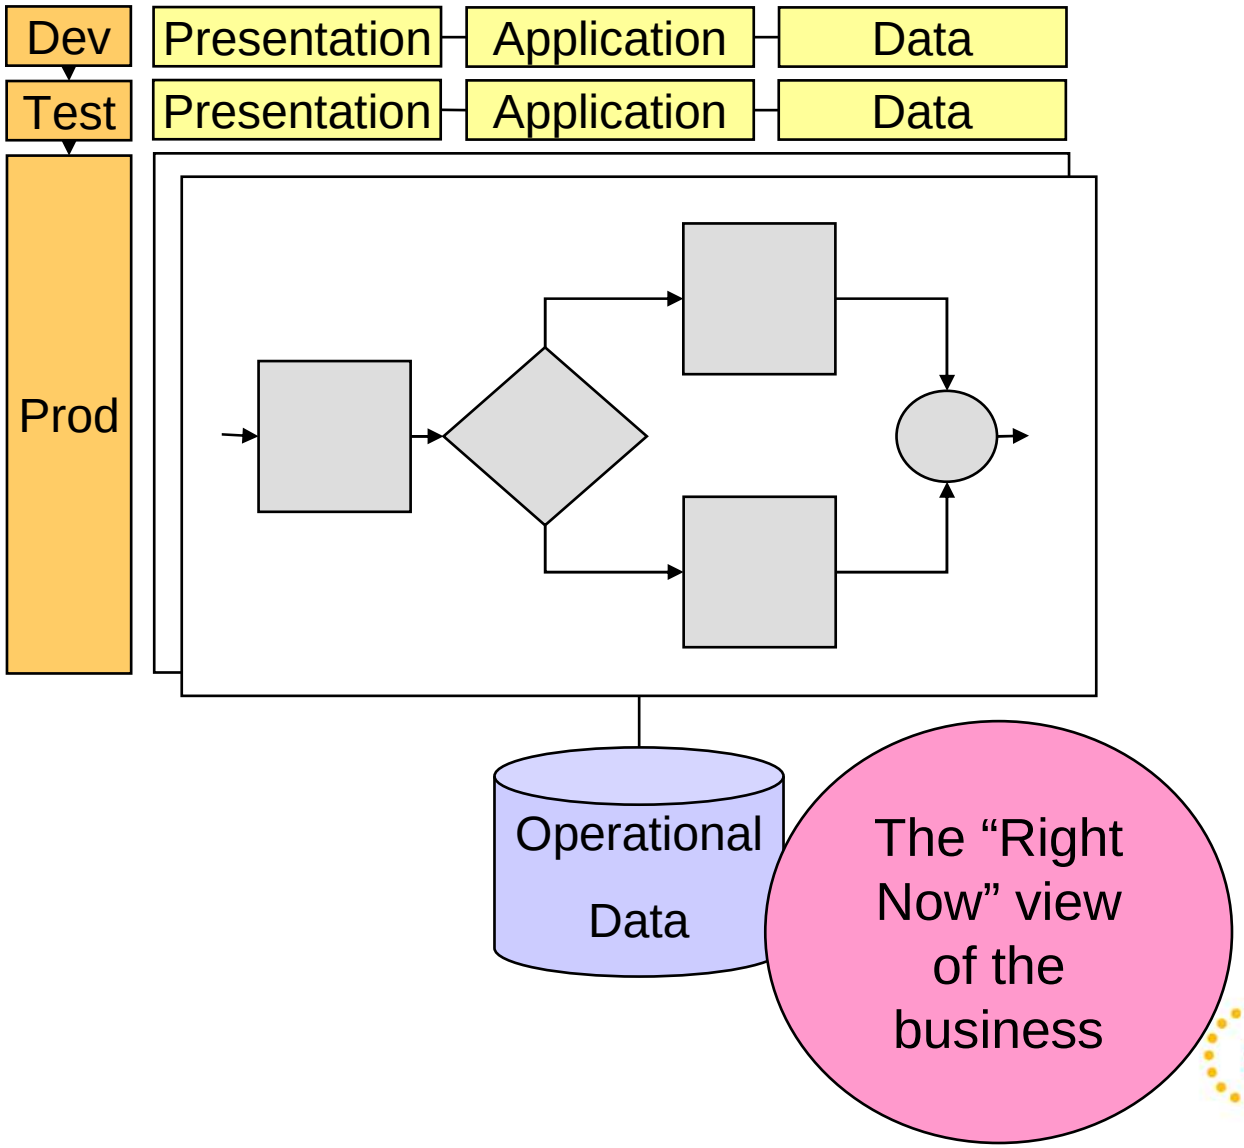

Topics

It's About the Data … It's About the Information (in the data)
(Business Intelligence is about turning data into information)

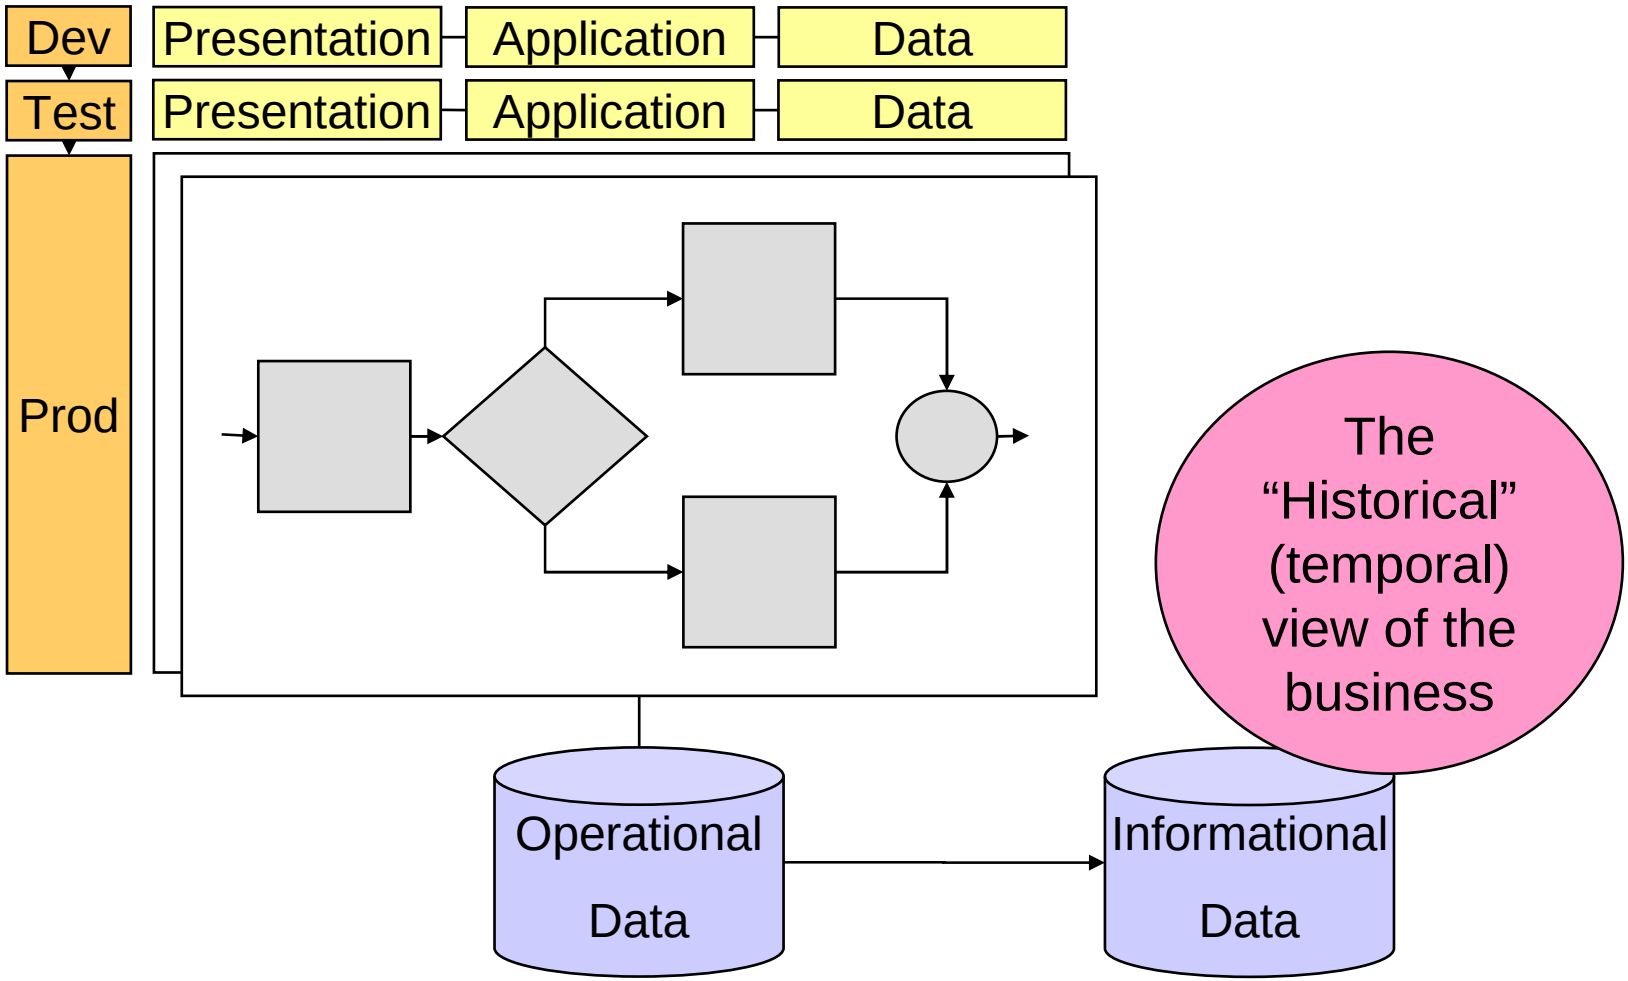

It’s About Making the Business (Process) Better

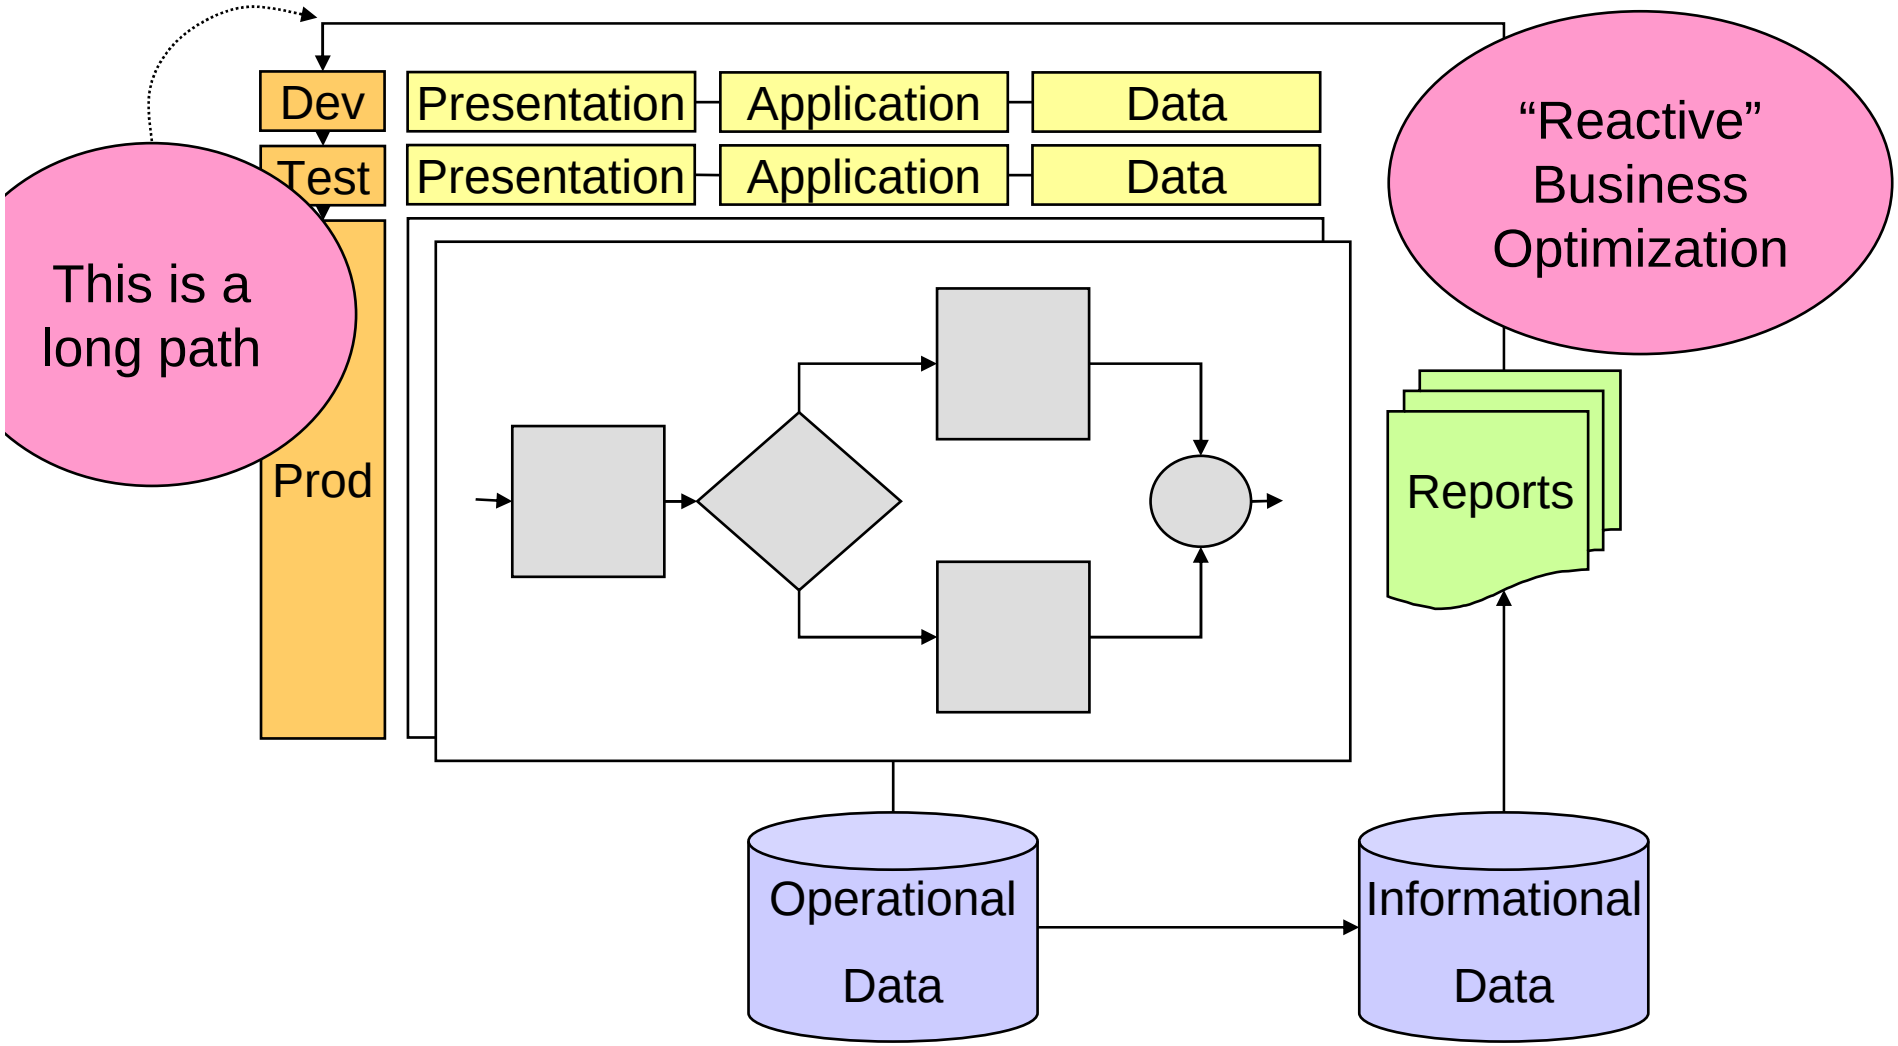

Topics

And then there was the “Smarter Planet” (starting with the “Connections”)

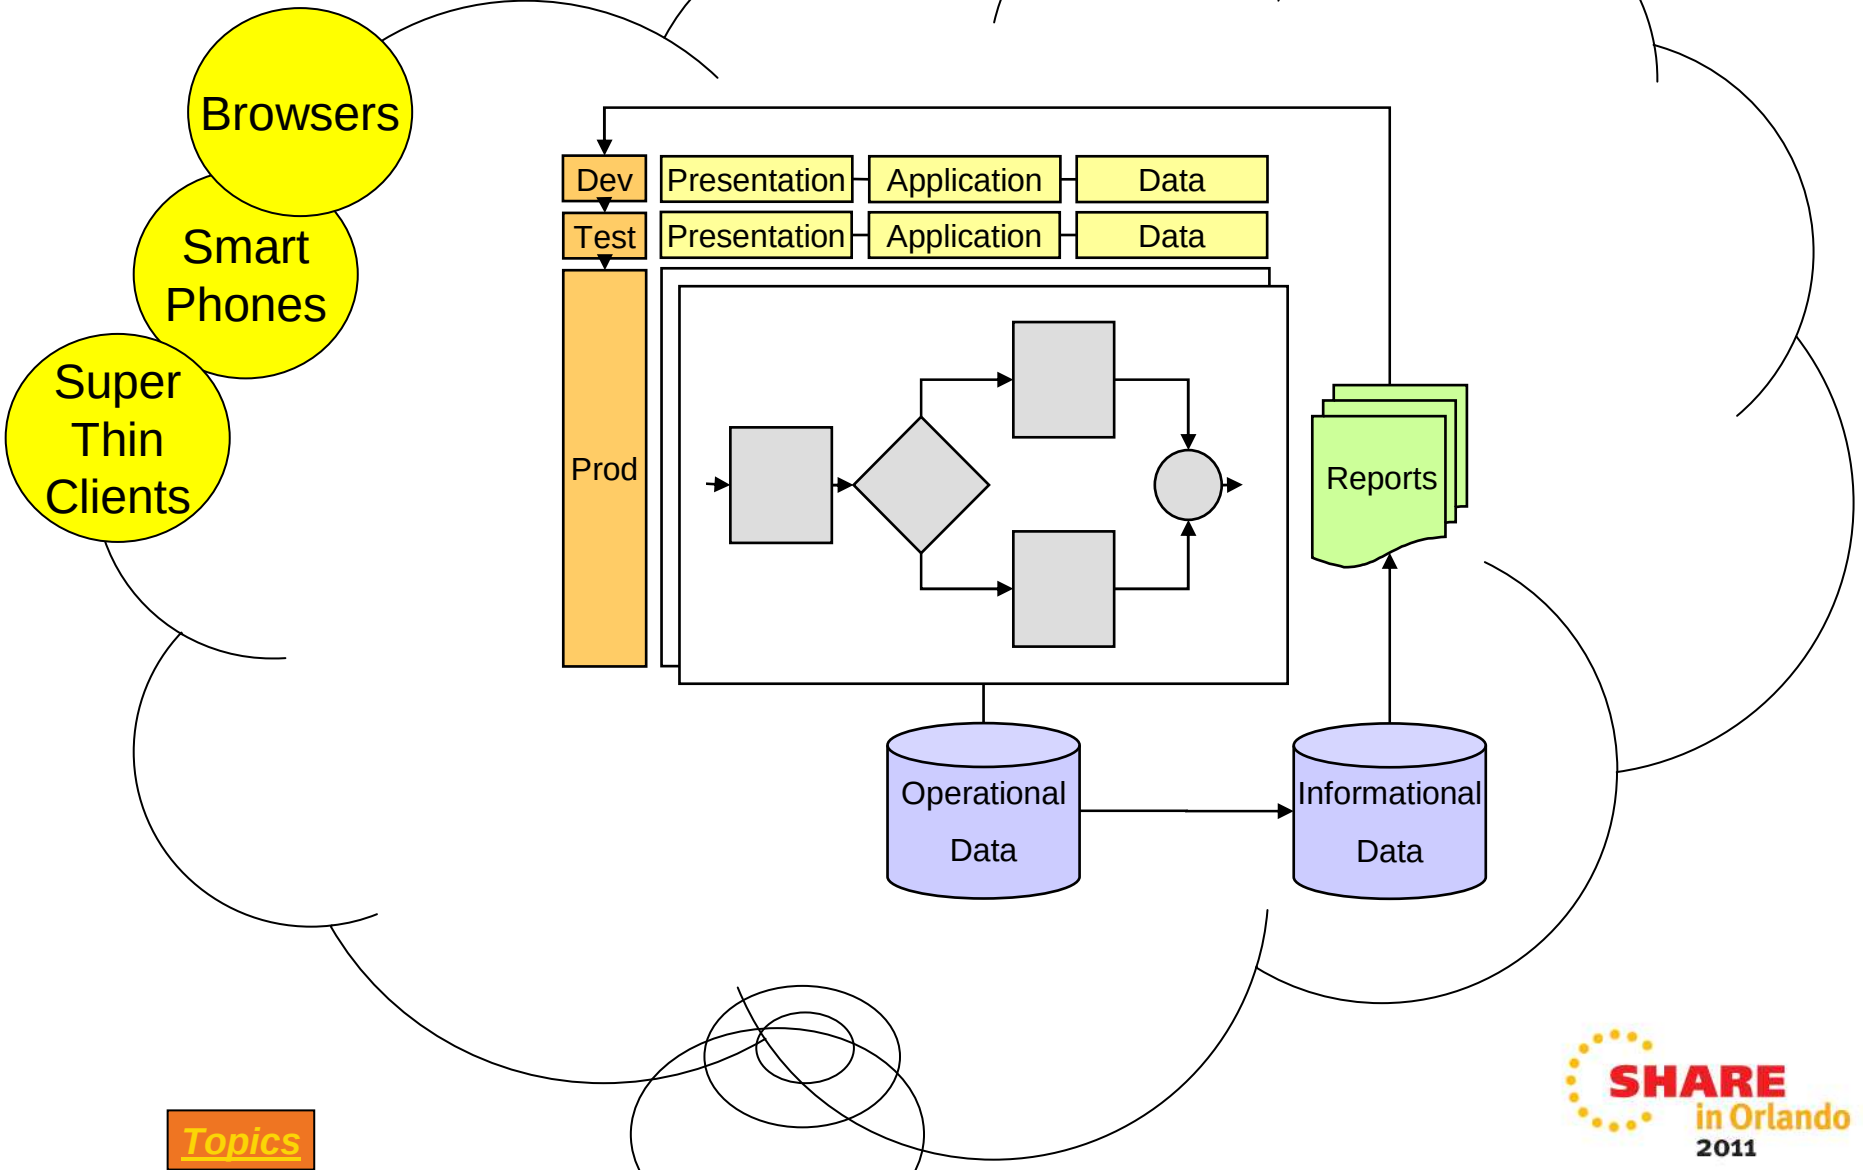

Topics

And then there was the “Smarter Planet” (leading to the “Internet of Things”)

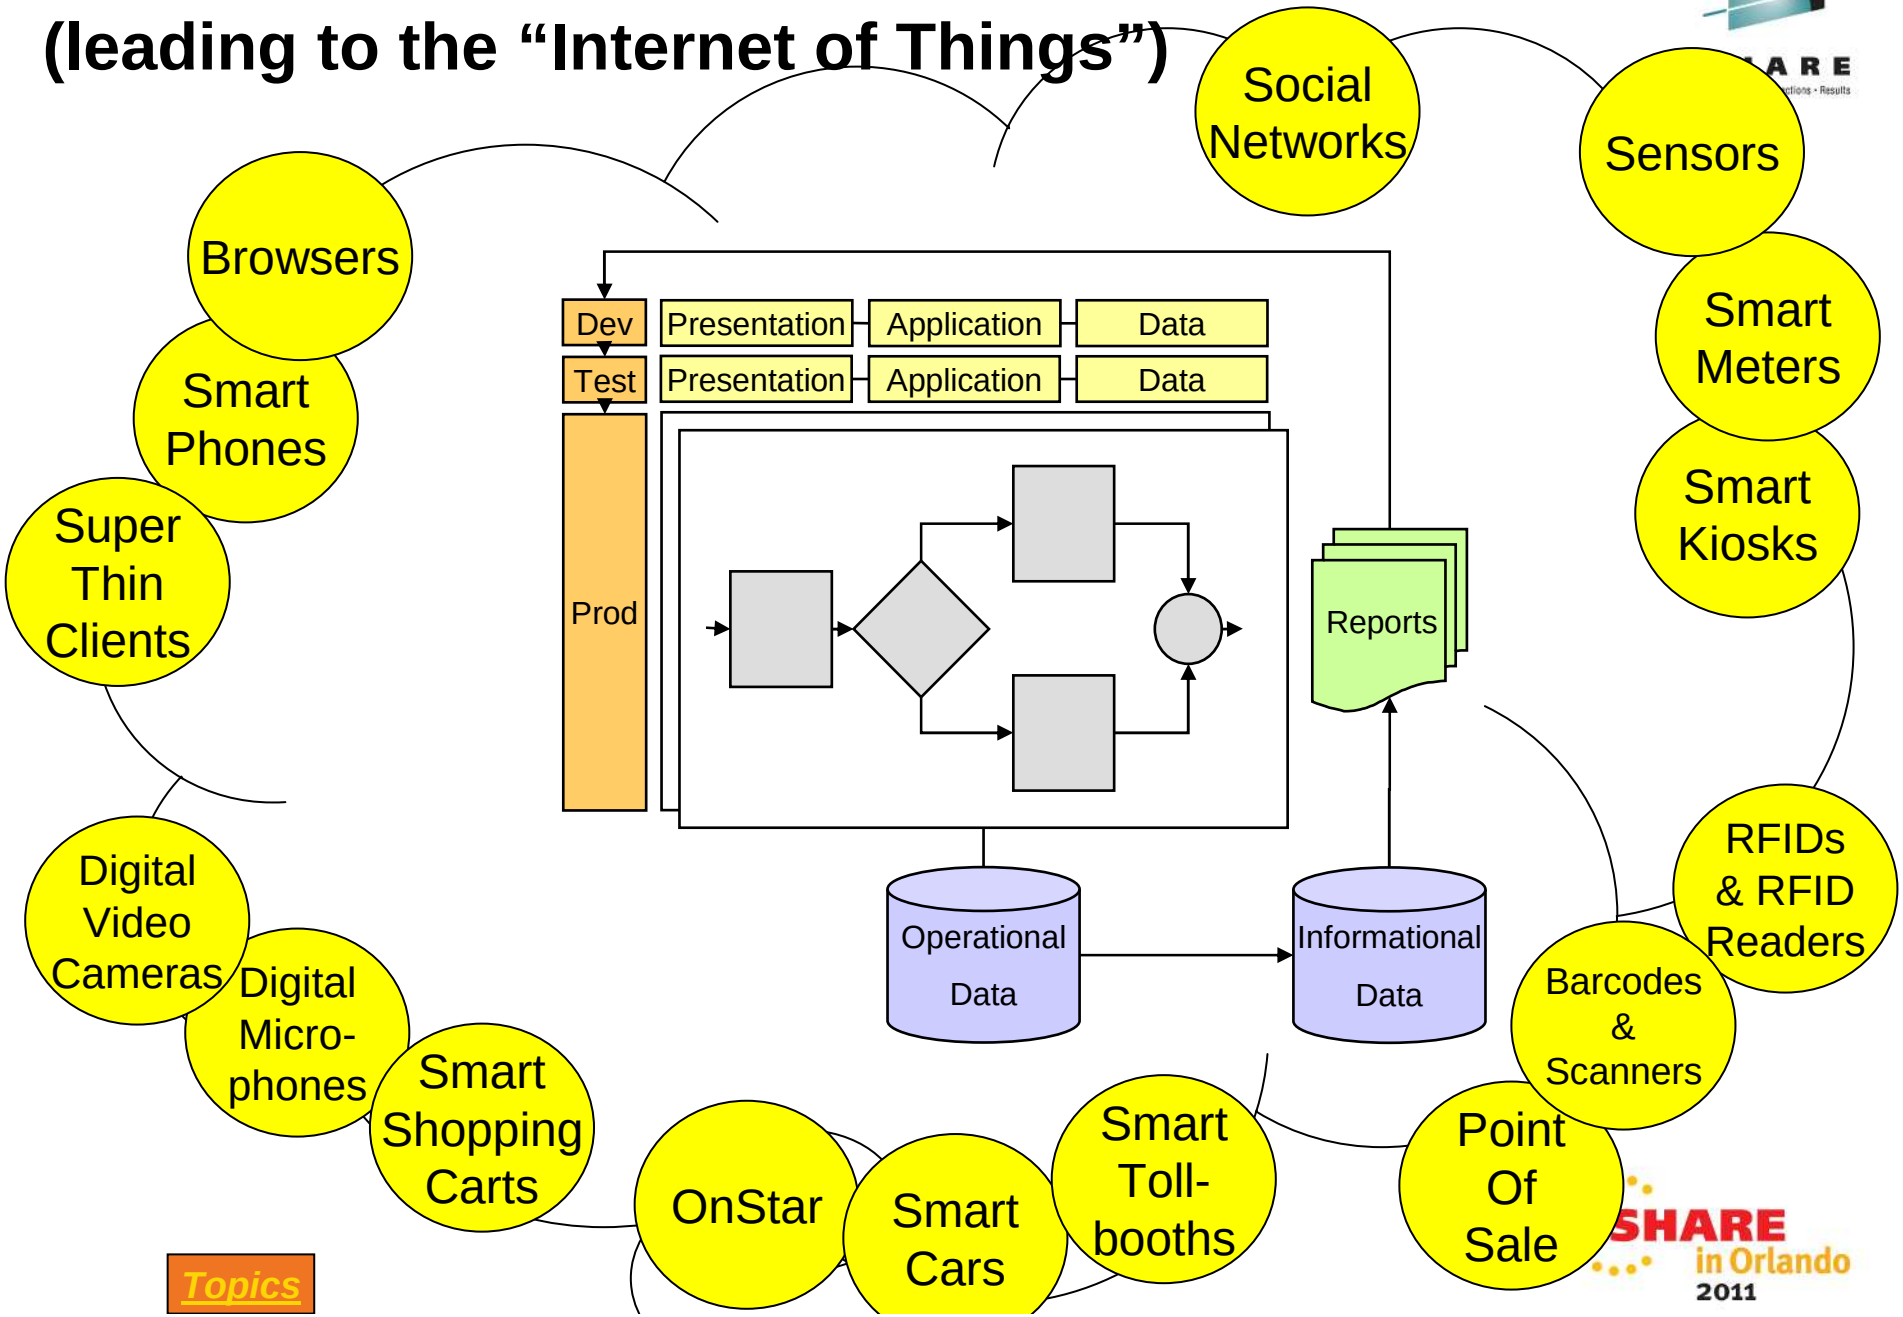

Topics

And then there was the "Smarter Planet" (creating "Real World Data" … i.e. BIG DATA)

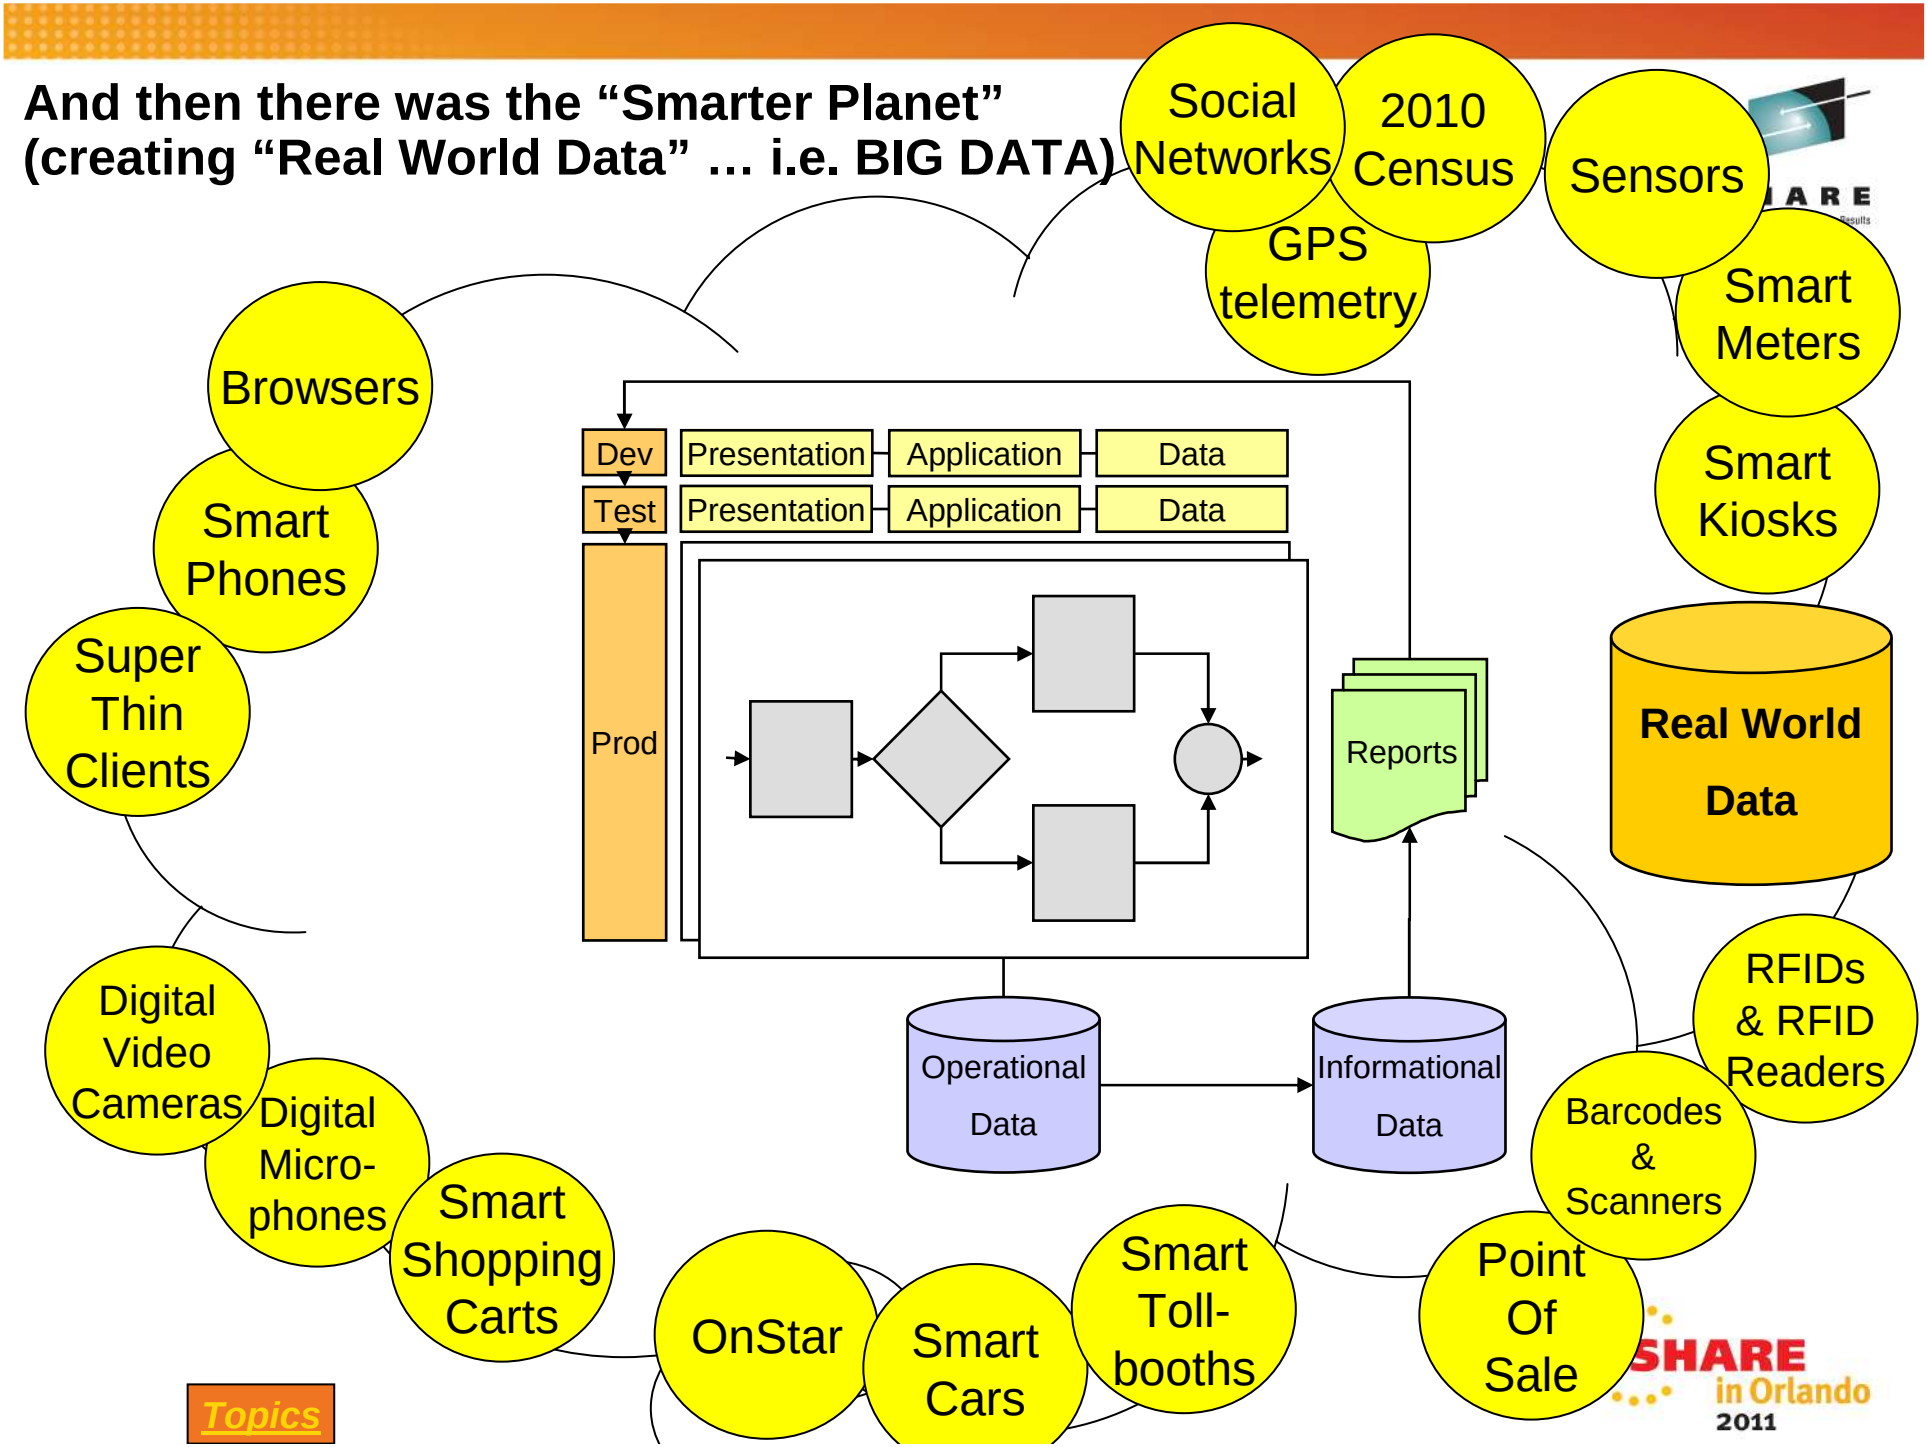

Topics

And then there was the "Smarter Planet" (and extending the reach of IT)

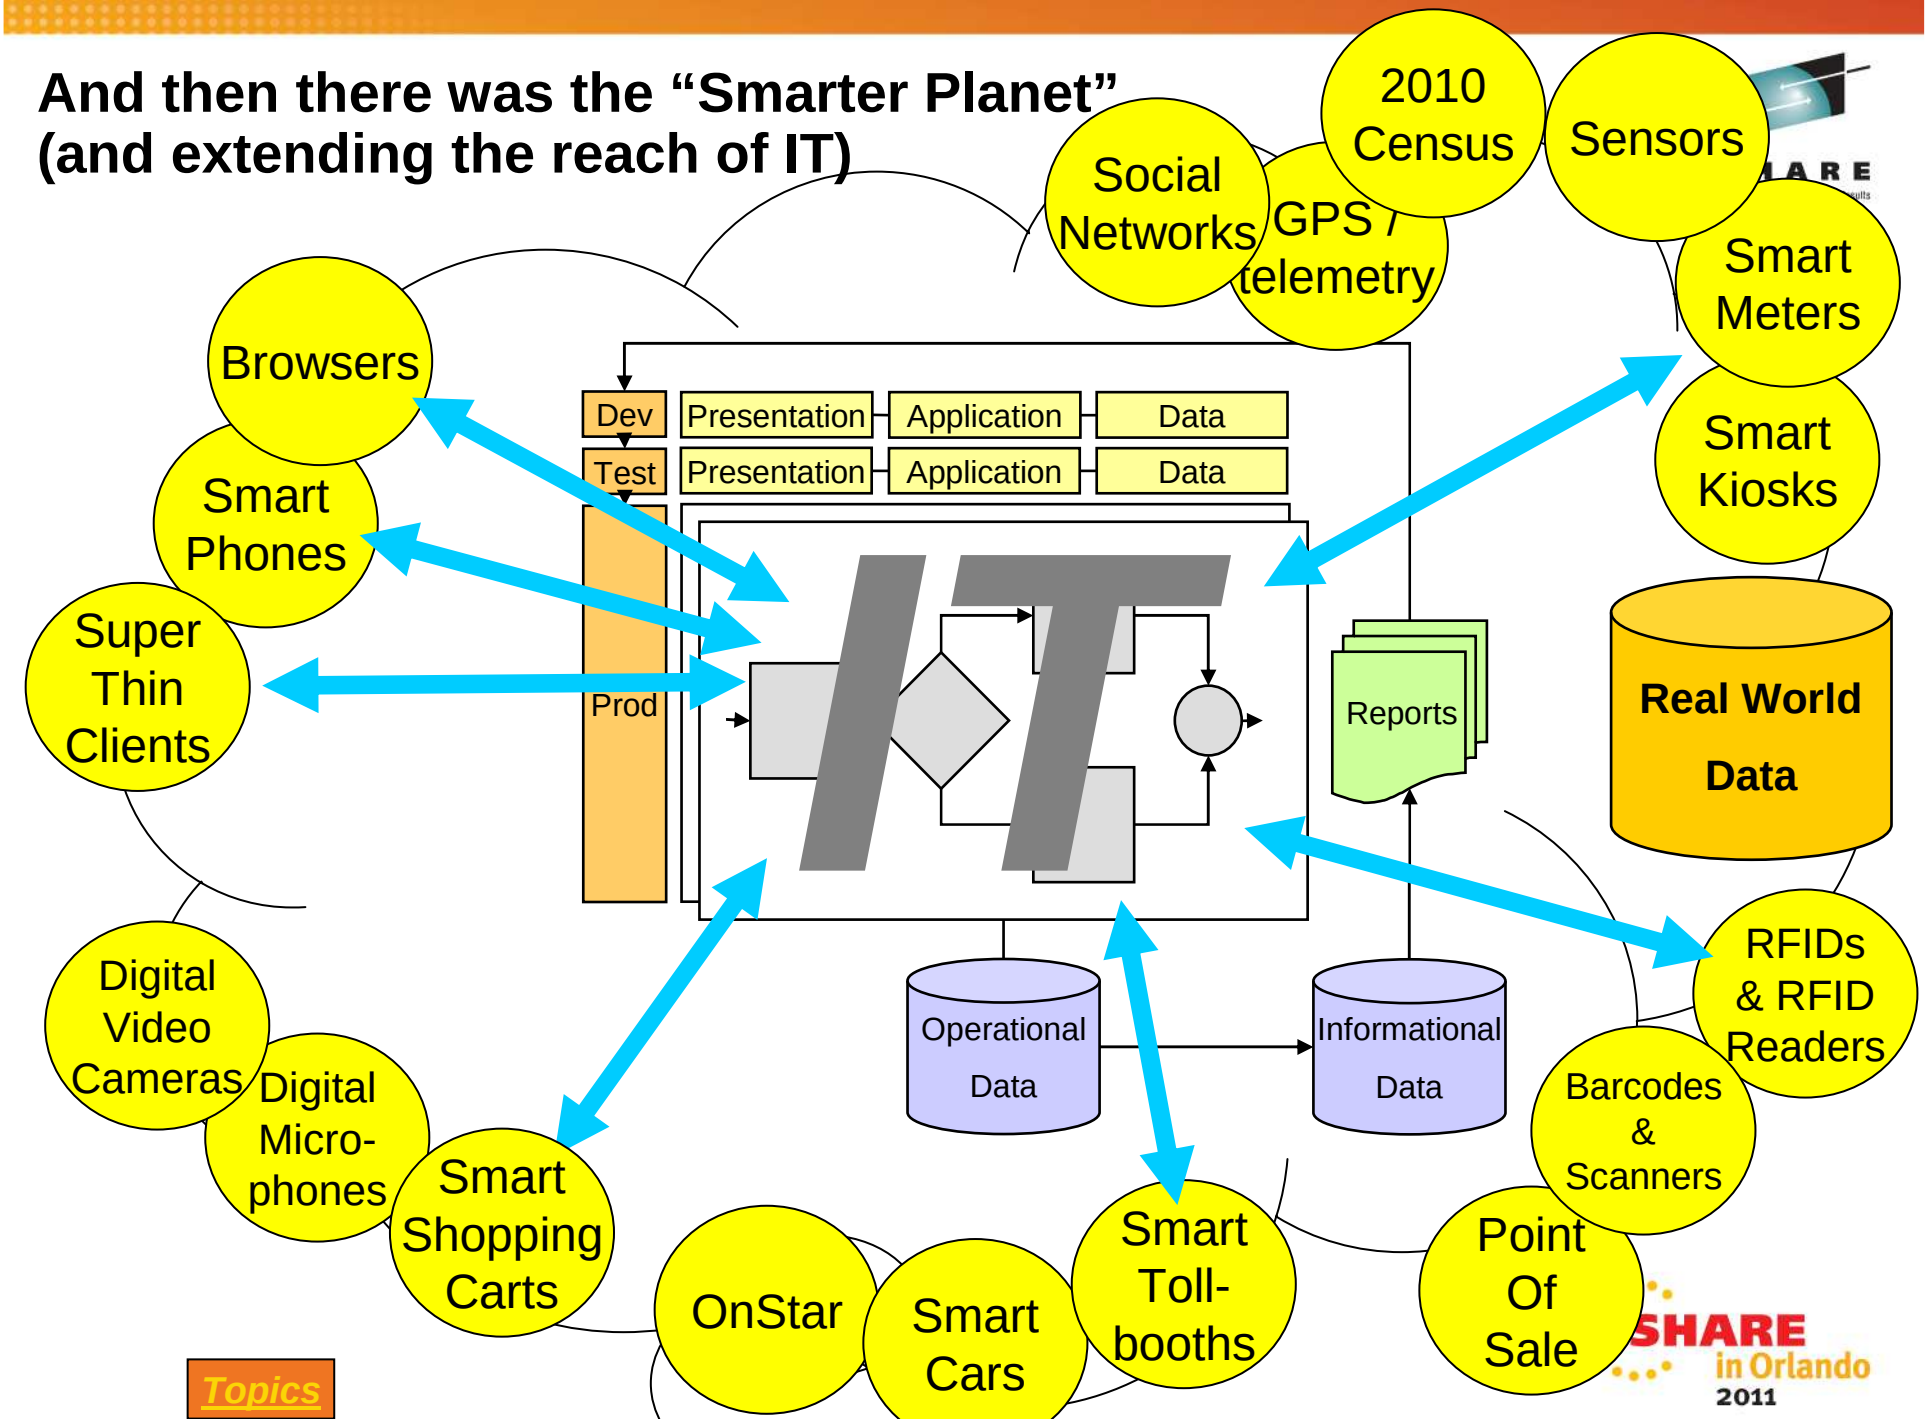

Topics

And then there was the “Smarter Application” (offering a “Smarter” way to evolve the business)

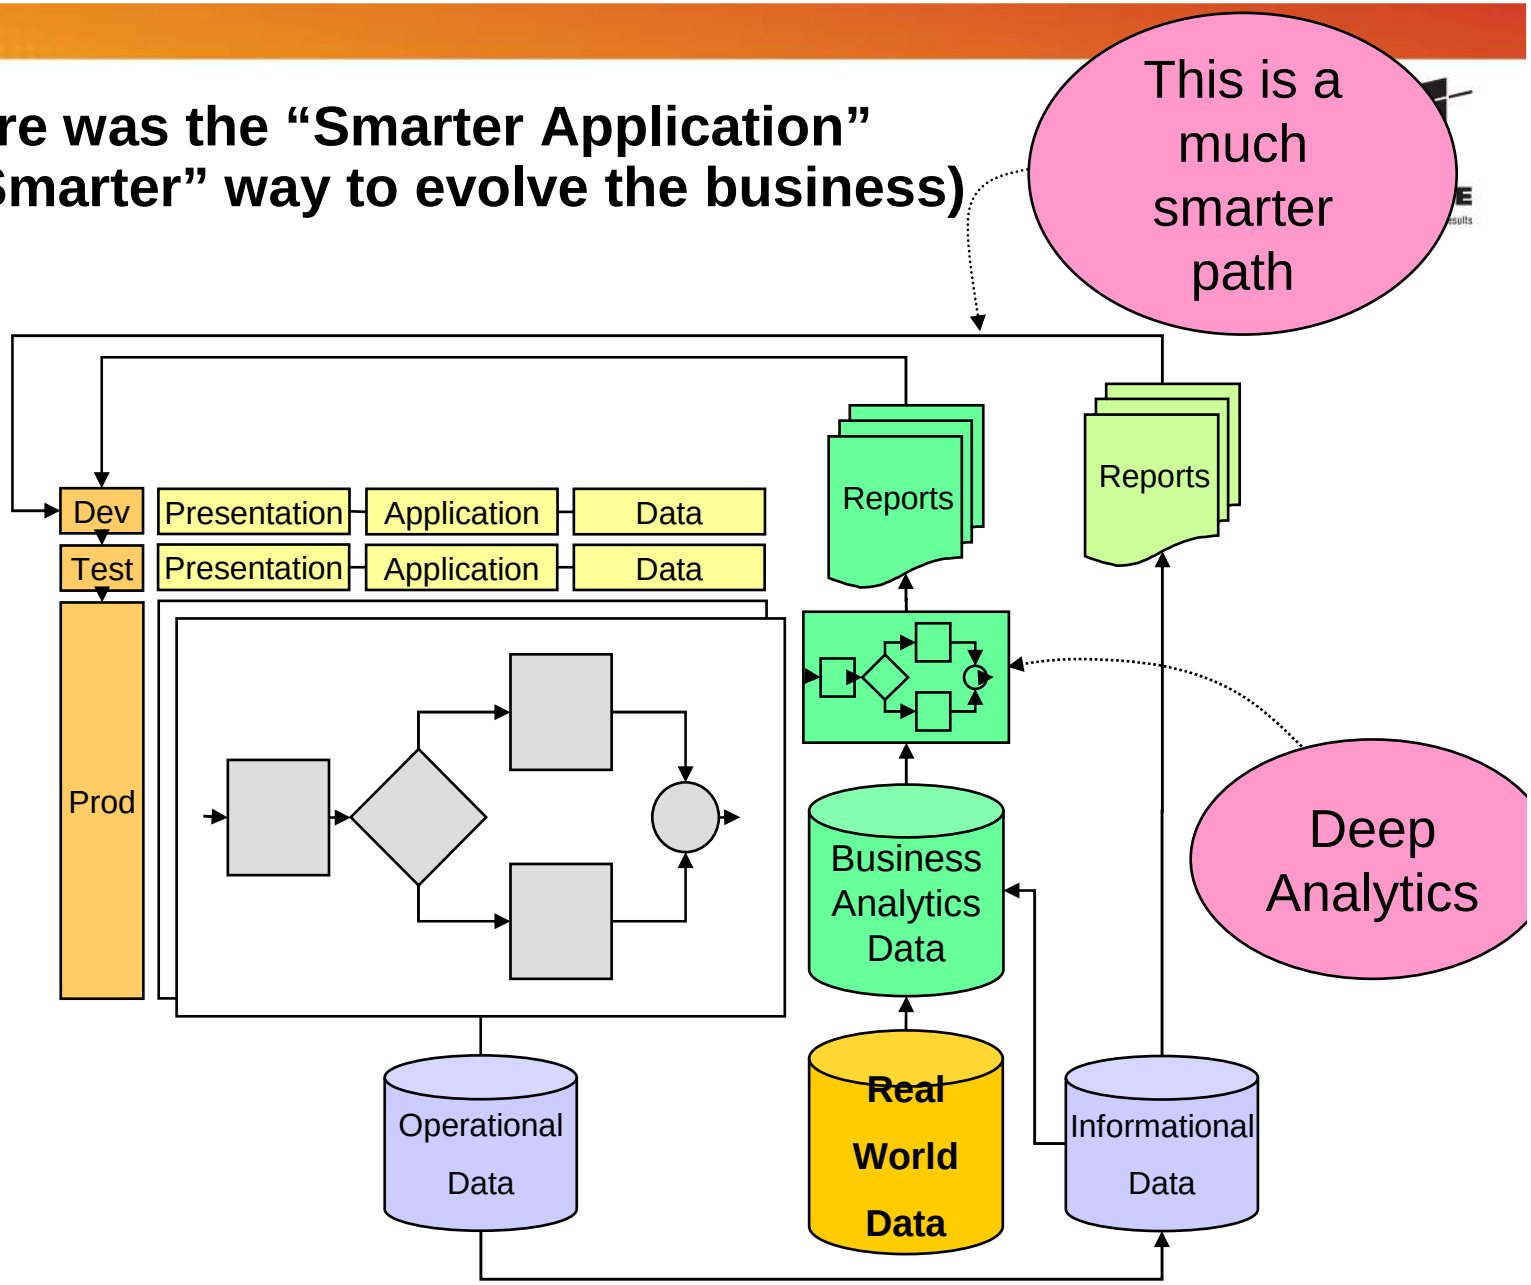

Topics

And then there was the “Smarter Application”
(offering a “Smarter” and “Faster” way to evolve the business)

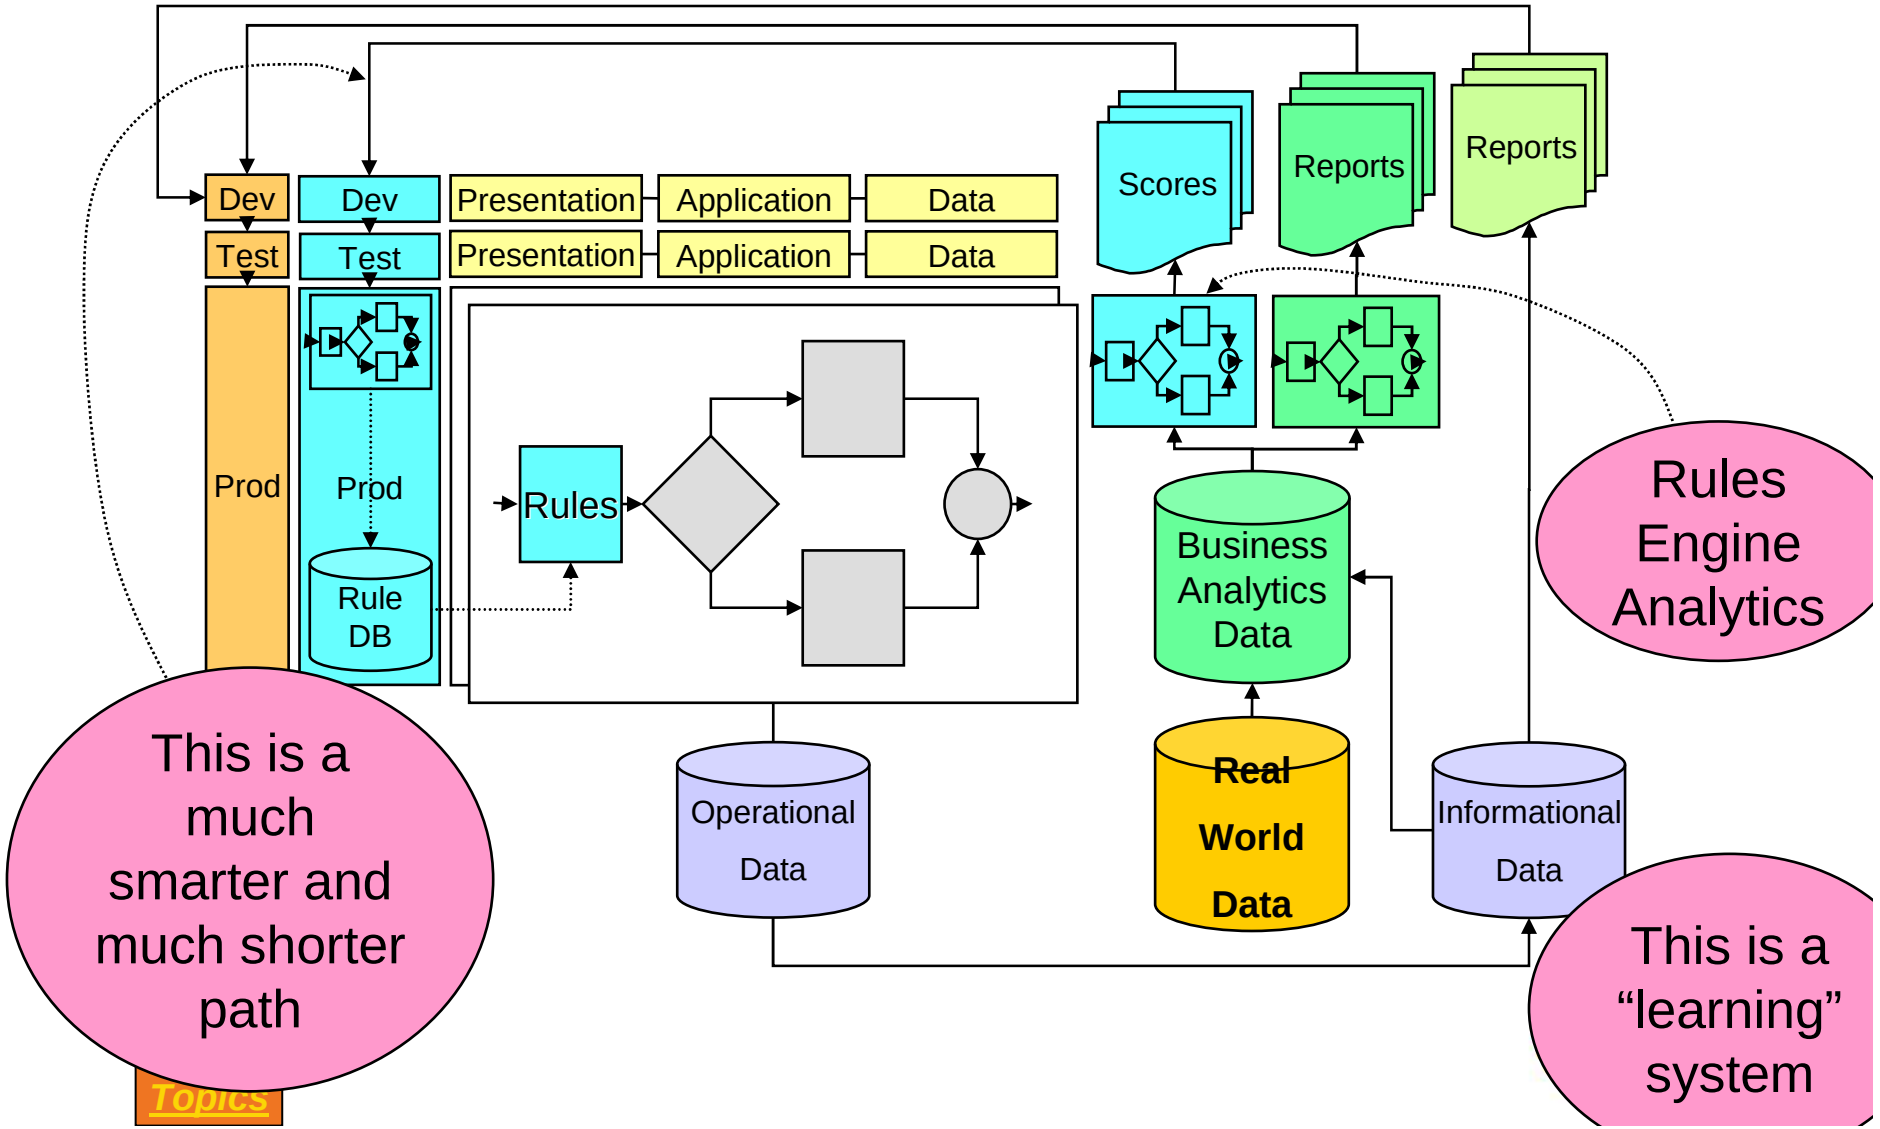

And then there was the “Smarter Application”
(offering a “Smarter” TRANSACTION!!!!)

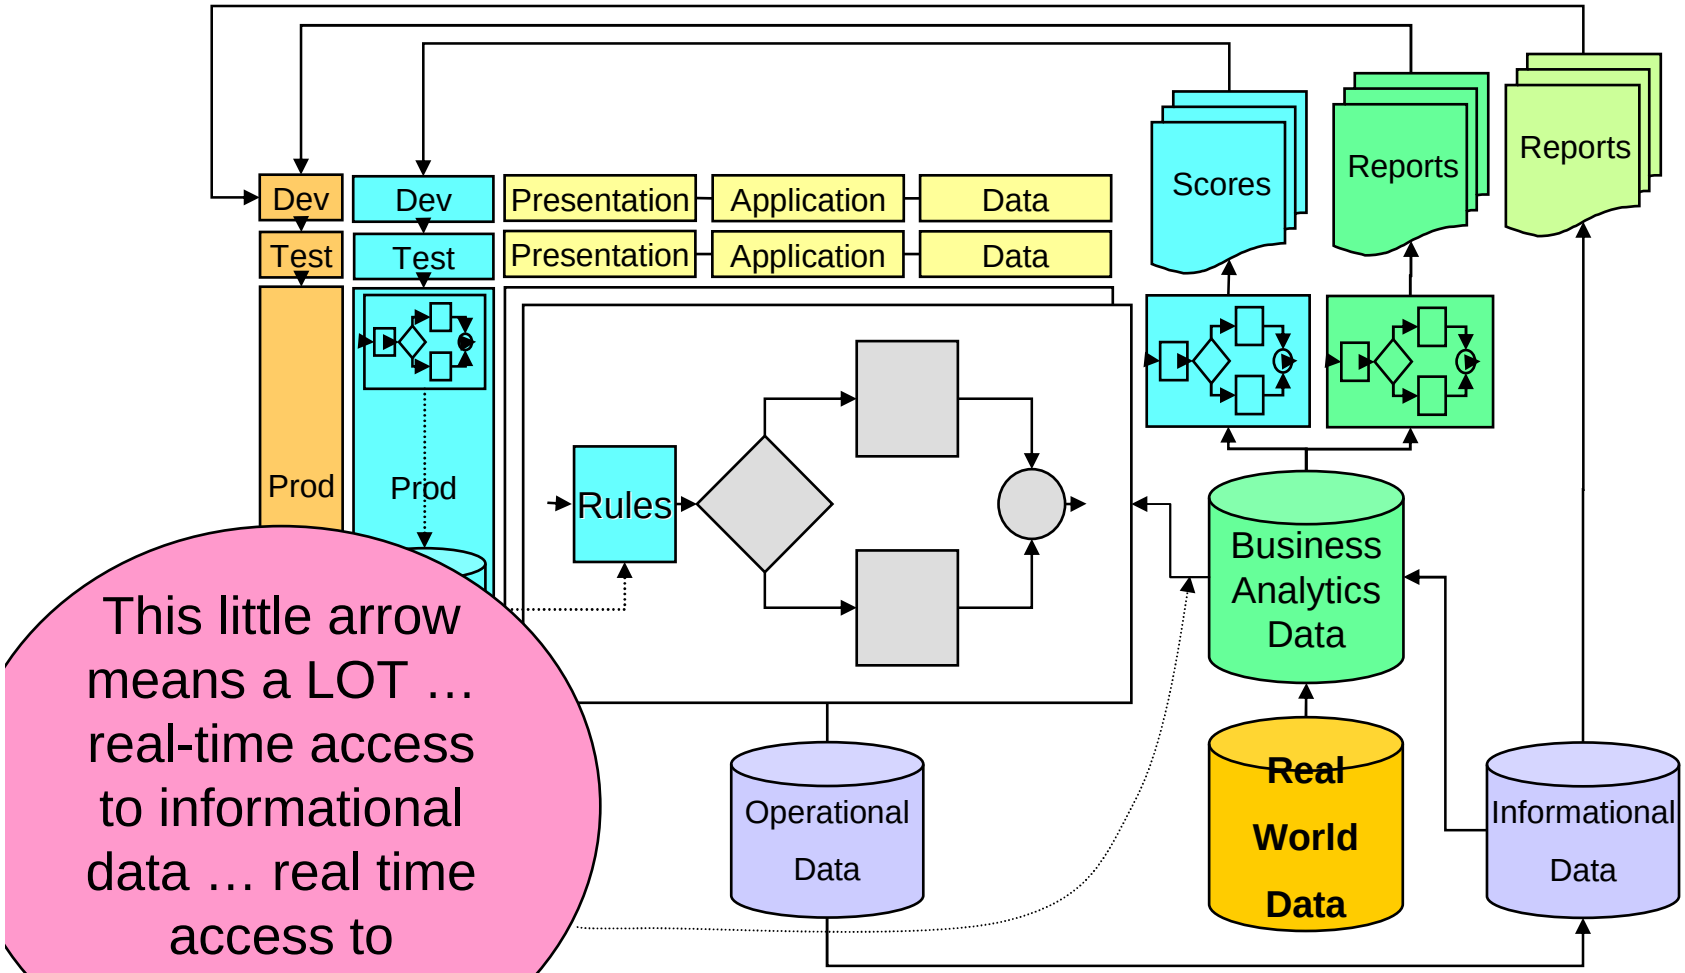

And then there was the “Smarter Application” (driving requirements for “Smarter” servers)

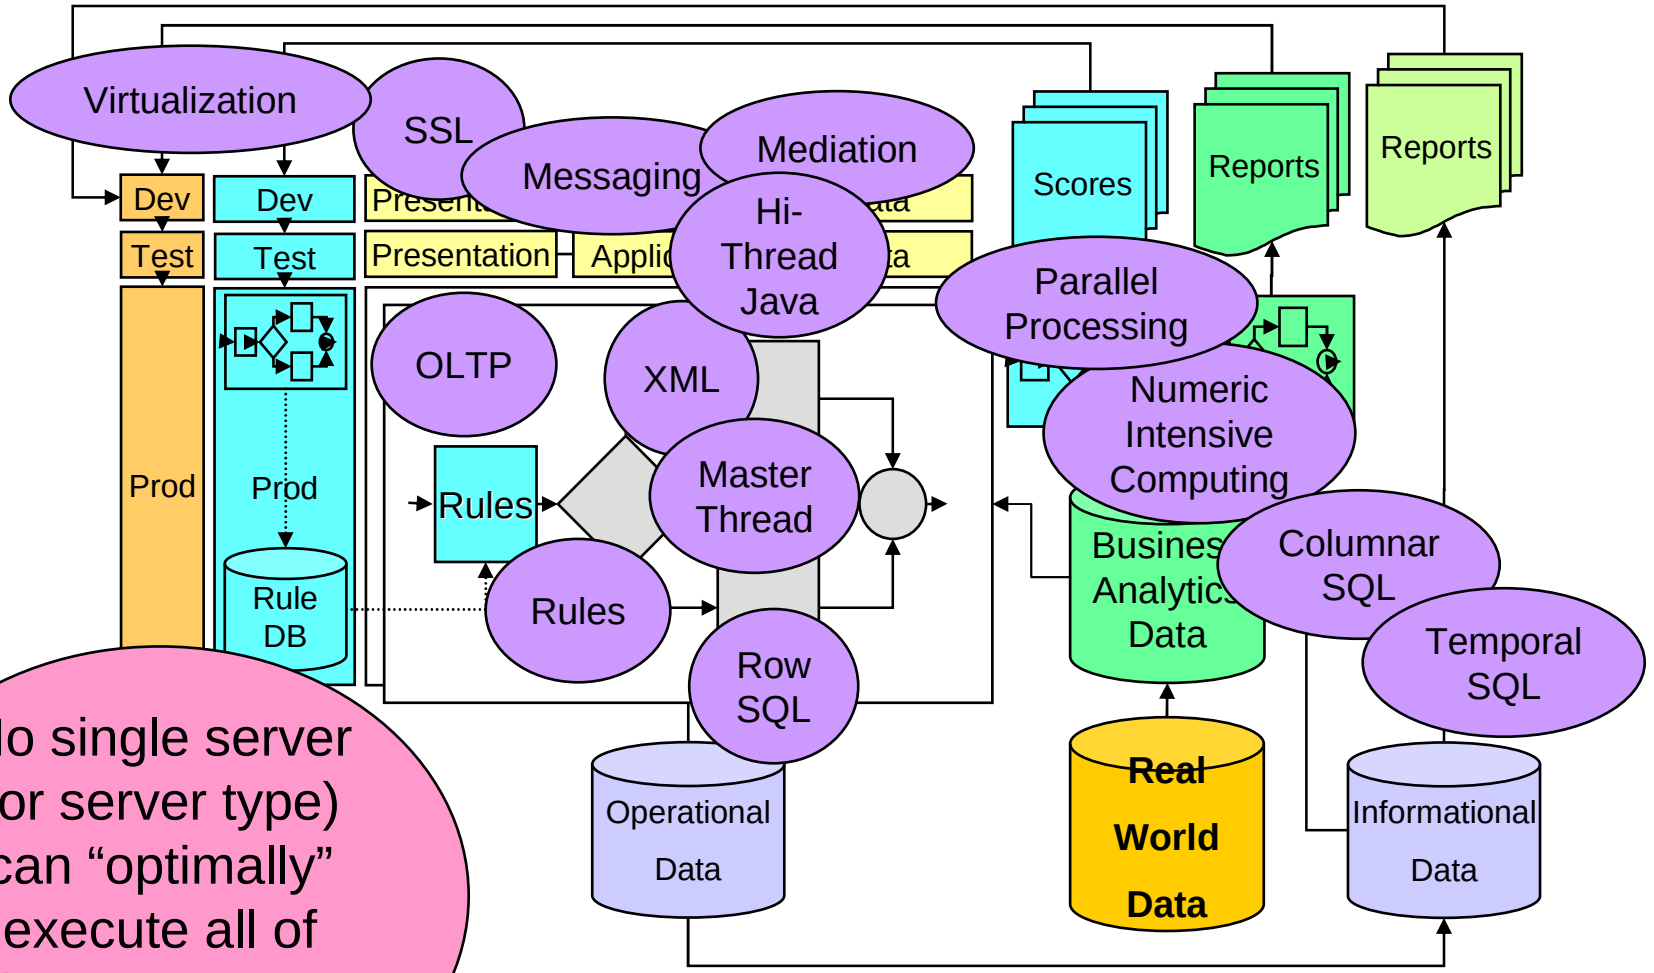

No single server (or server type) can “optimally” execute all of these compute models

In Summary

(smarter planet > smarter apps > smarter servers)

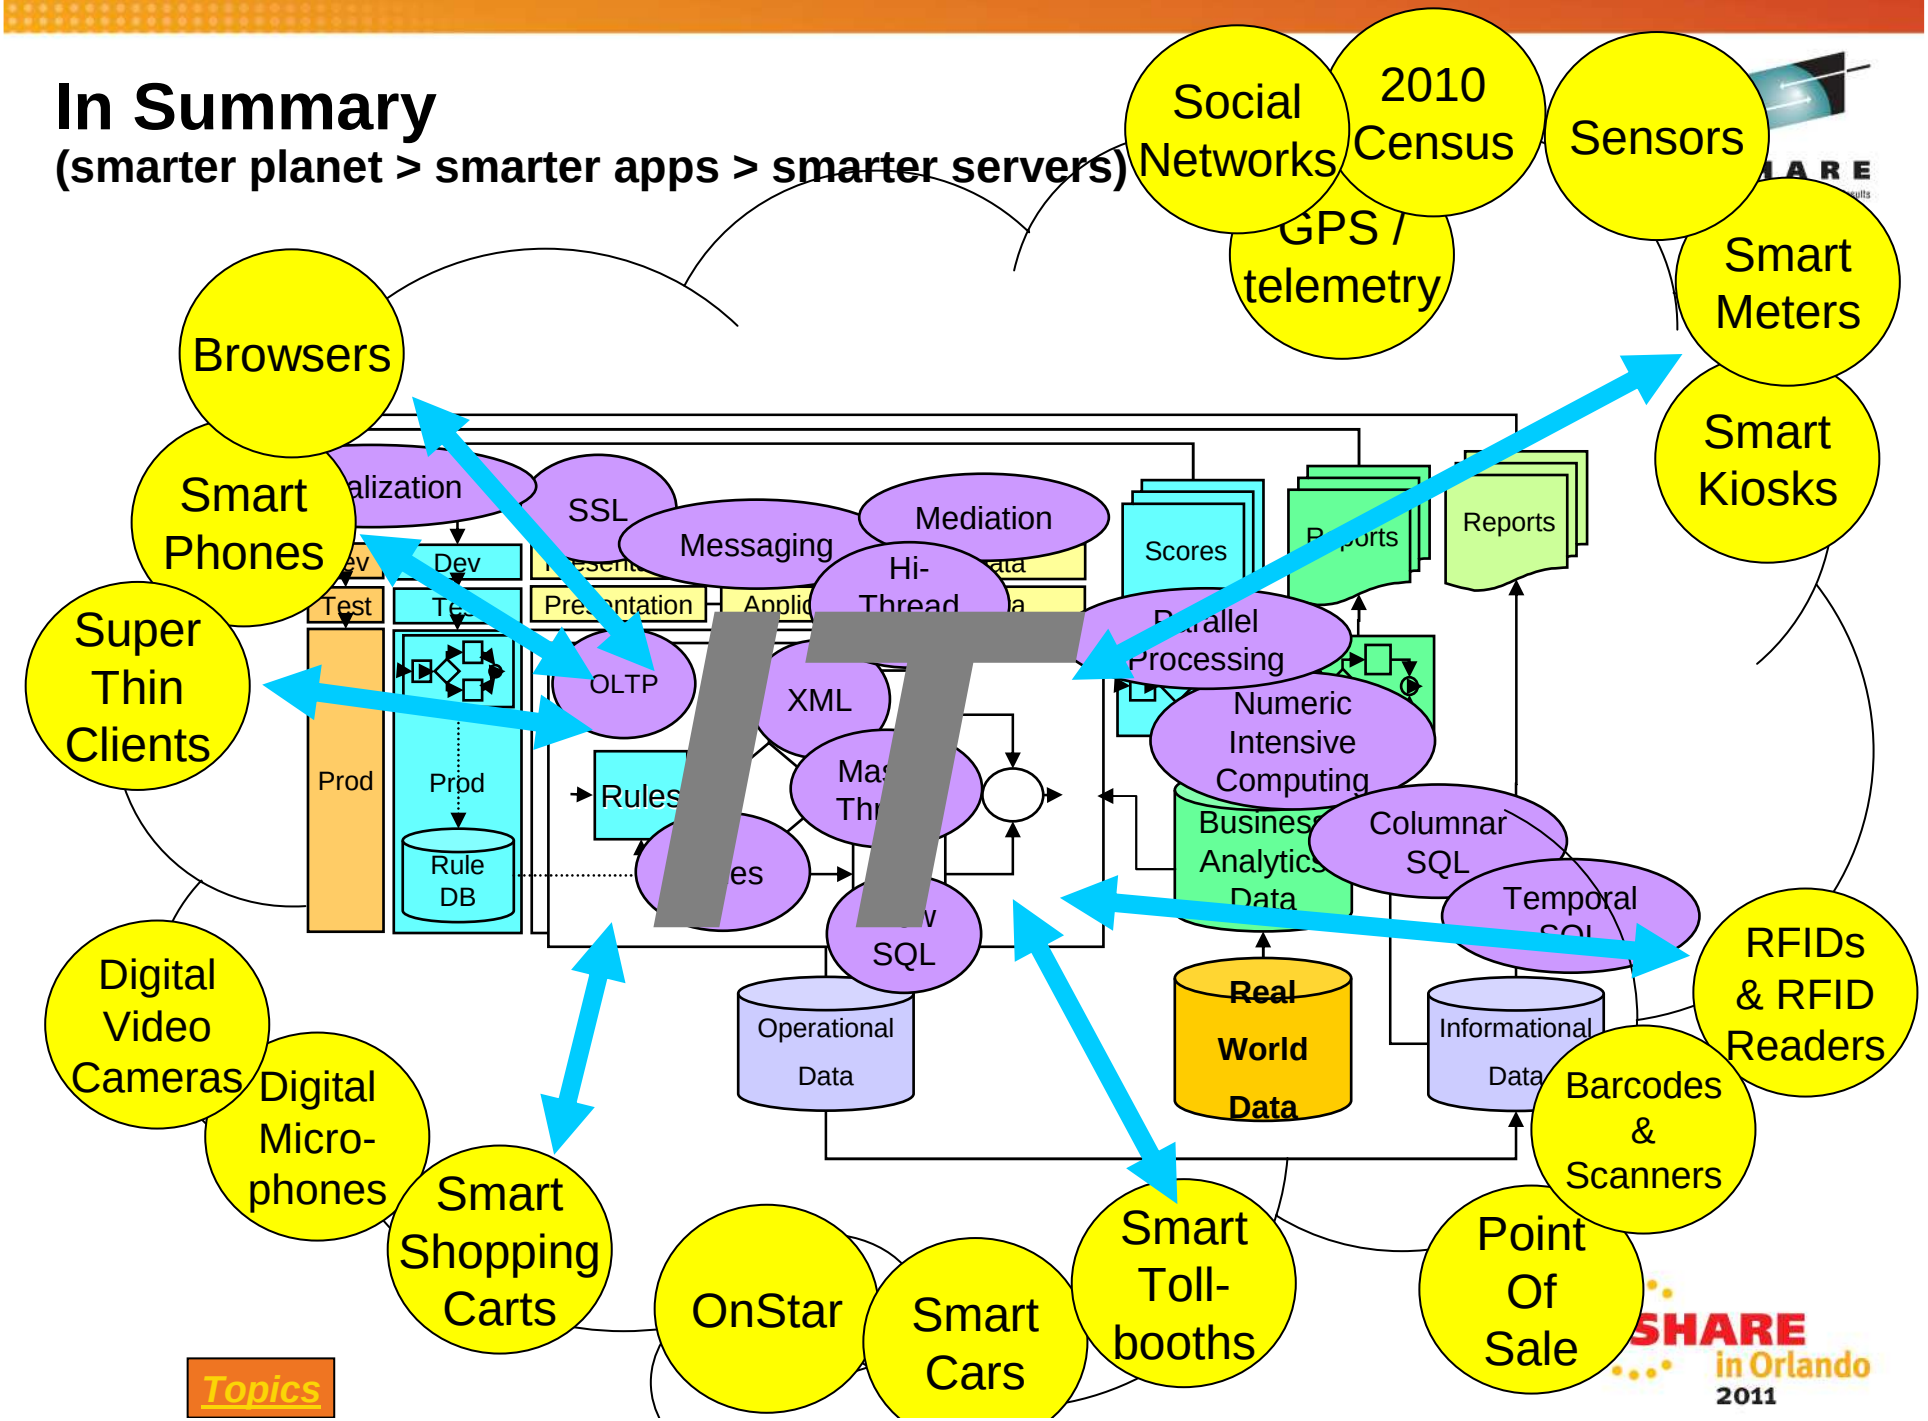

Topics

End of Section

Design Principles

Topics

General Purpose Computing

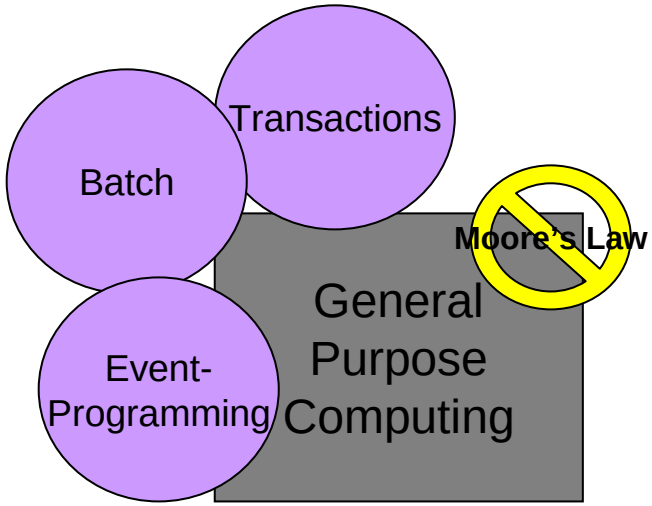

General Purpose Computing…

Continues to be the cornerstone of business process automation.

New event-based programming models (a natural fit with real-world extension of IT)

If Moore's Law is not broke …
it is at least getting bent!

What's your definition?
(2X density improvement every 18mo)
(2X speed-up every 2yr)
(2X price/performance every 2yr)

CMOS technology progress is slowing.

Logic density growth is slowing.

Frequency growth is slowing.
(some may say "reversing")

Cost/performance improvement is???

Many technology providers are taking non-traditional approaches to afford continued price/performance progress.

Topics

High Performance Computing

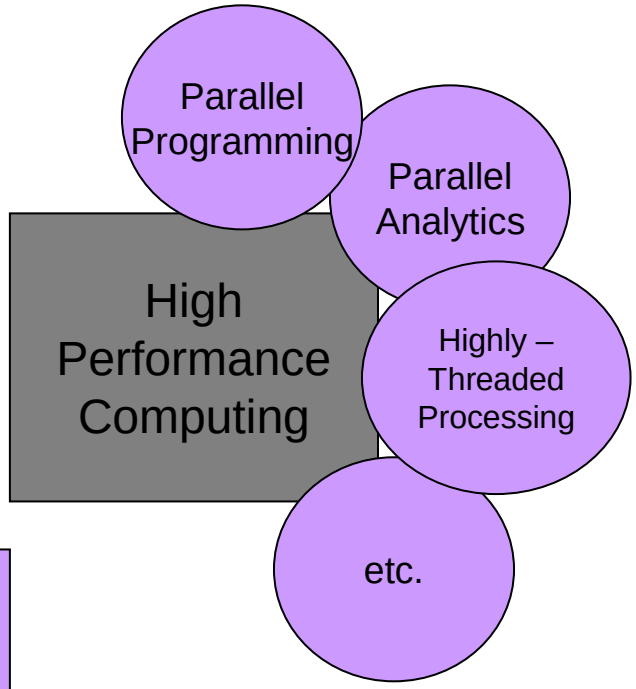

High Performance Computing…

Parallel algorithms
Vector processing
Analytic parallel-data algorithms
Predictive analytics
(patterns and correlations used to predict outcomes in real time)

Topics

Application Domain Computing

Application Domain Computing…

Specialized hardware & algorithms
Single-purpose
Appliance form factor

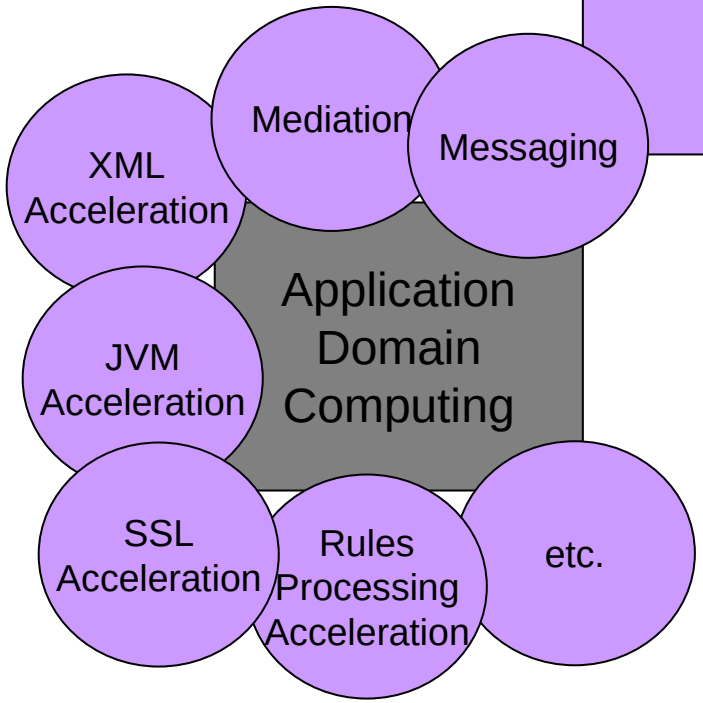

Topics

Network-Speed Computing

Network-Speed Computing…

Real-world data feeds
Processing into actionable form
Processing into reduced/stored form
Processing into patterns
Processing into events (triggers)

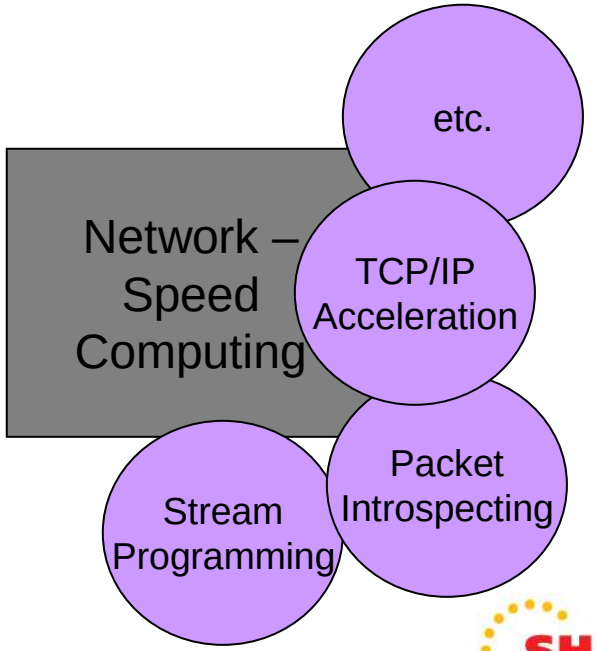

Topics

Hybrid Systems

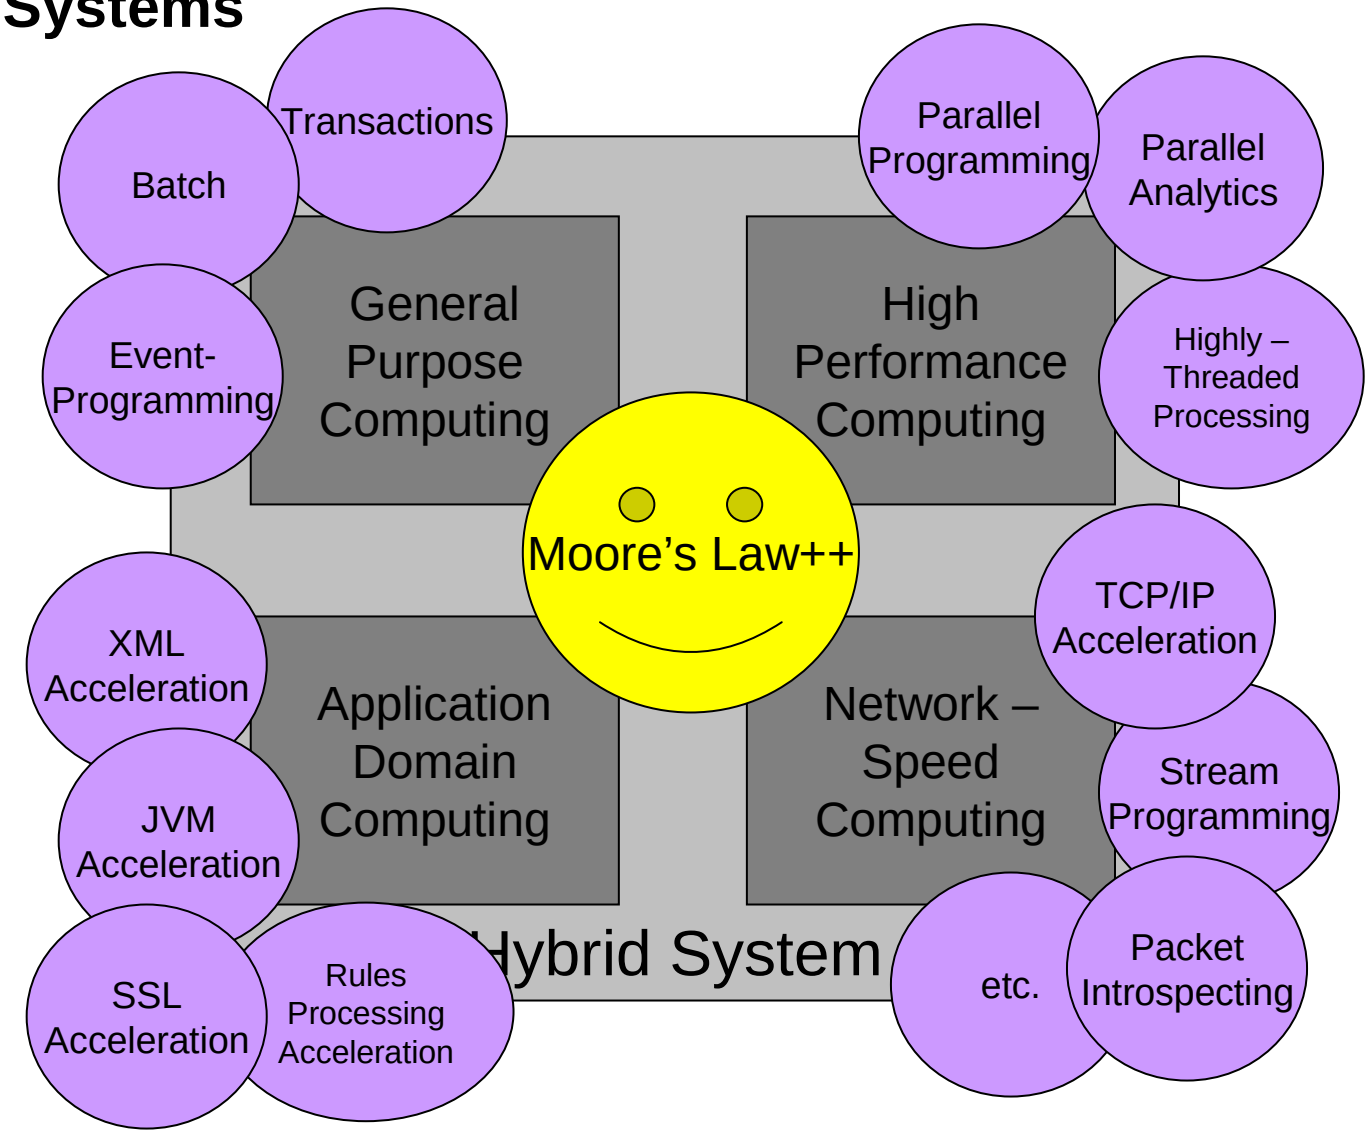

Designed not just to meet Moore's Law, but in the aggregate, to far exceed it!

Topics

Hybrid Computing – a tail of two threading styles

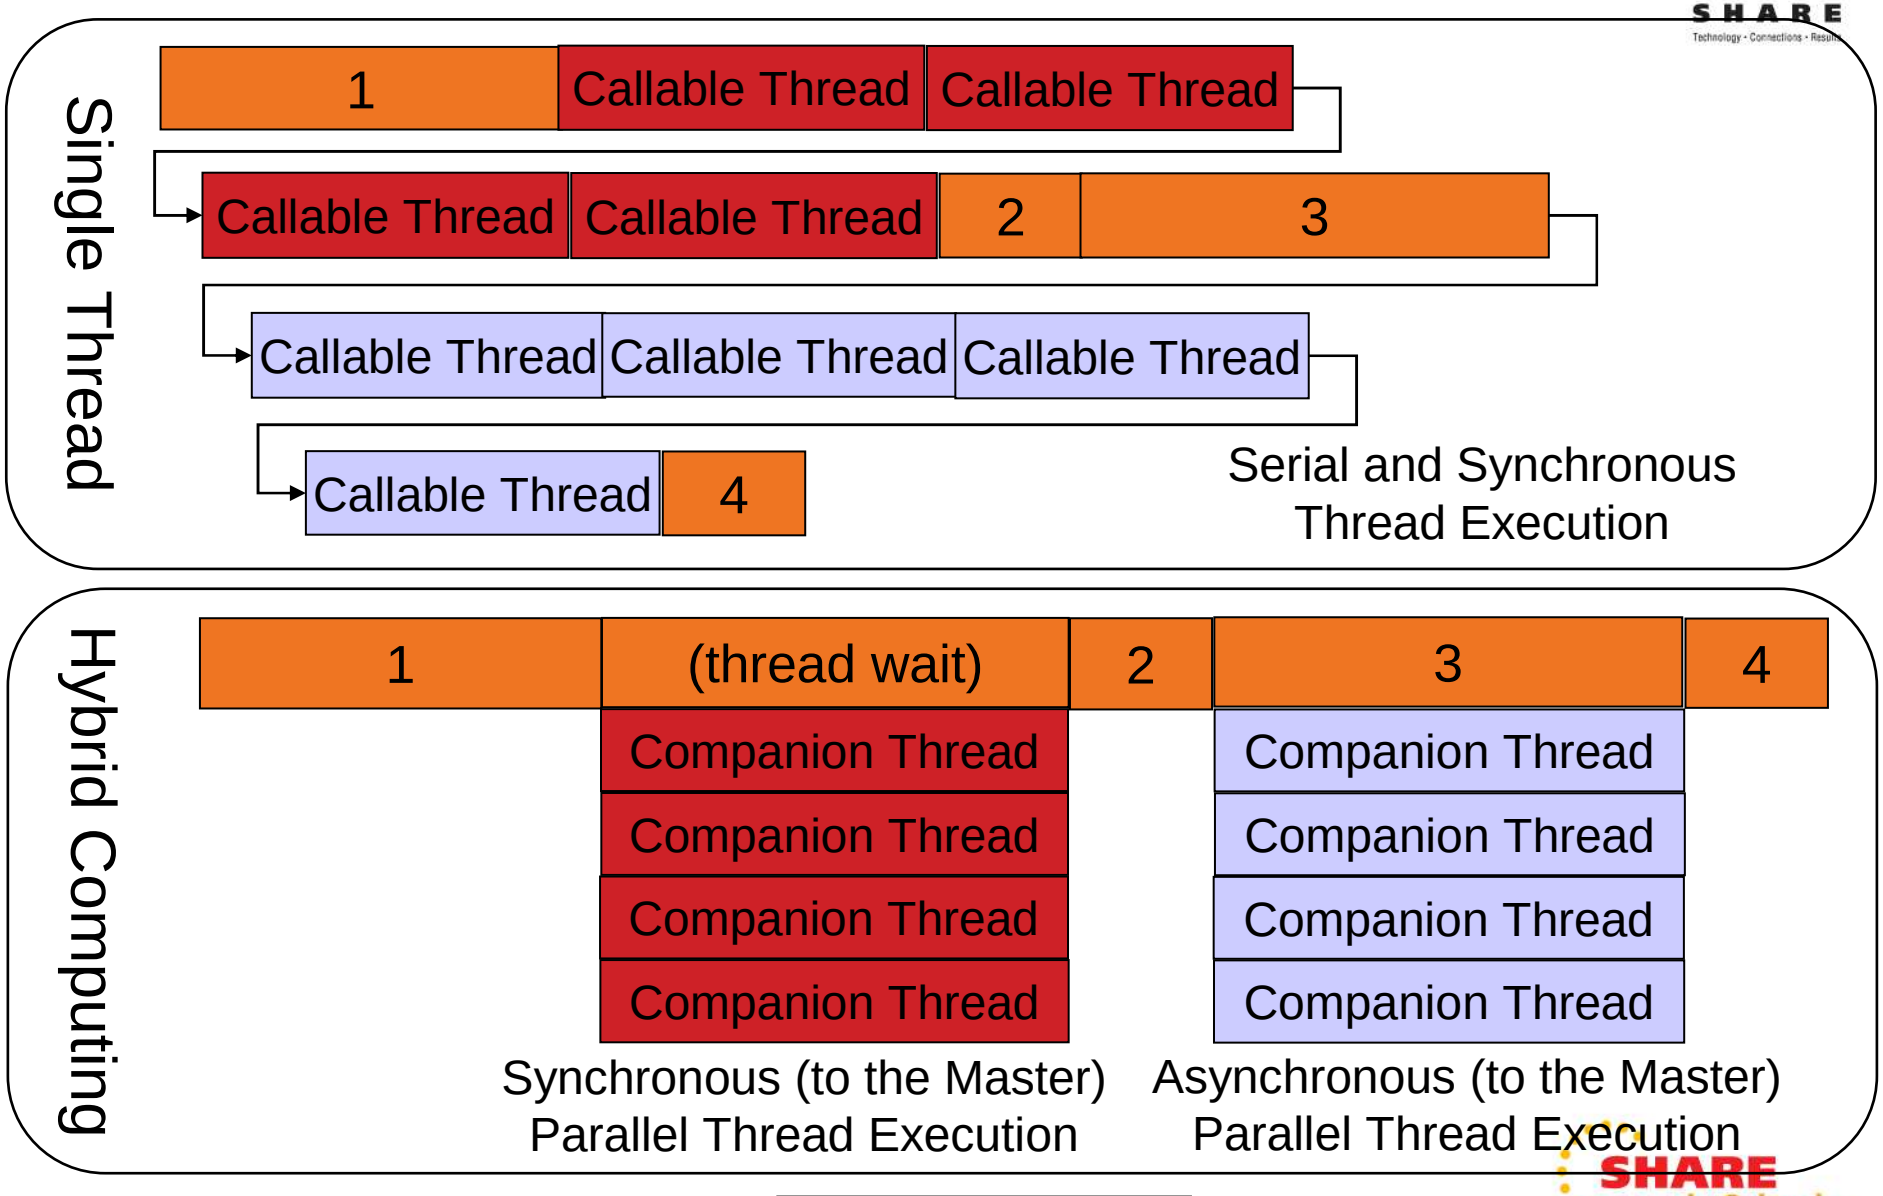

zEnterprise
A “Smarter Server” Design

zEnterprise = z196 or z114 + zBX + zManager

Topics

The z196

- Announced July 22 2010
- Generally Available September 2010
- "Sold Out" in 4Q 2010

(best "z Quarter" in a decade)

- Replacement for the z10 EC family of servers
- Enhances and embellishes on trends started by the z10 EC
- Cornerstone of the zEnterprise

z196

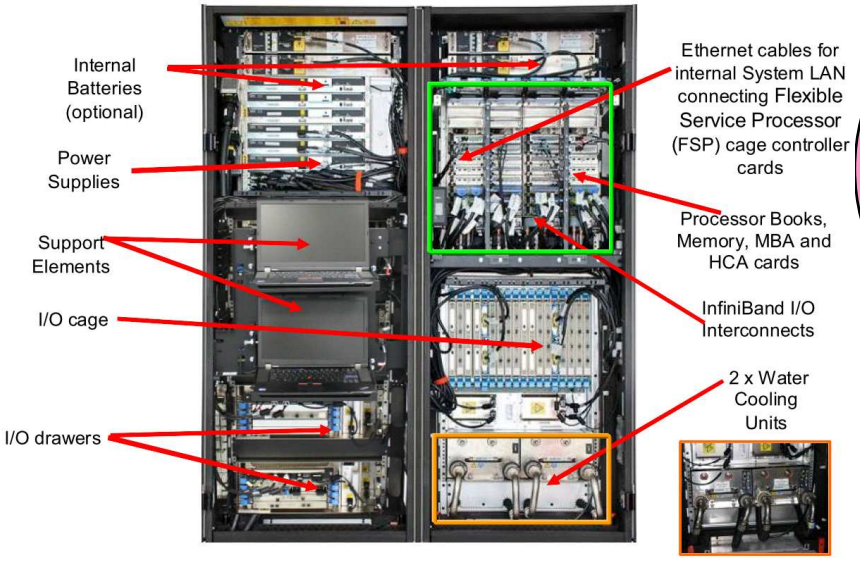

State of the Art In High End General Purpose Computing

Topics

The z114

z114

- Announced July 12 2011
  - Generally Available Sept 09 2011
- Replacement for the z10 BC family of servers
  - Cornerstone of the zEnterprise

State of the Art in Mid Range General Purpose Computing

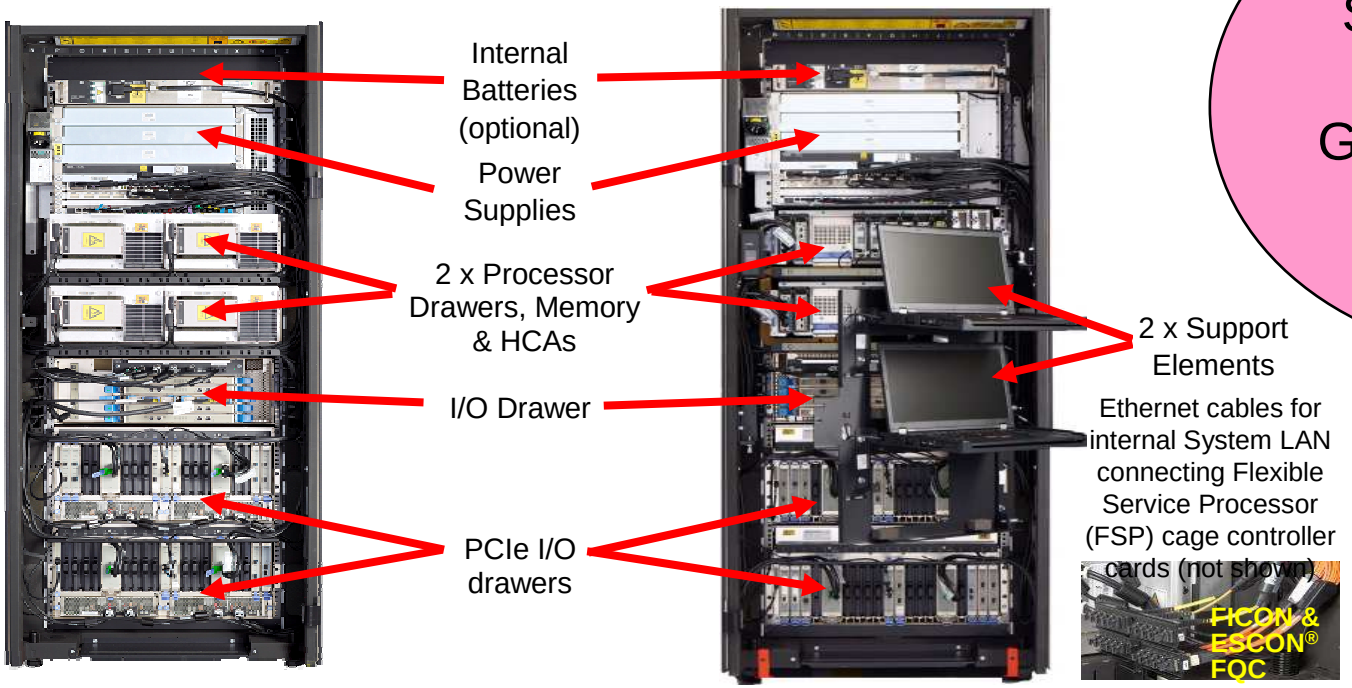

Topics

MICROPROCESSOR:

4.4GHz
Superscalar
In Order
L1 (192KB) and
L1.5 (3MB) per core
(total:12.75MB per chip)
1B transistors
44 new instructions

z10 EC | z196

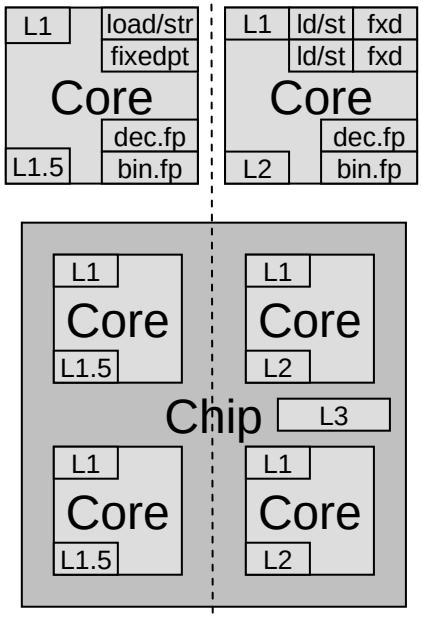

MICROPROCESSOR:

5.2GHz
Super-duper-scalar
Out of Order
L1 (192KB) and
L2 (1.5MB) per core
(new) L3 per chip (24MB)
(total:30.75MB per chip)
1.4B transistors
110 new instructions

Topics

SHARE in Orlando 2011

MICROPROCESSOR:
4.4GHz
Superscalar
In Order
L1 (192KB) and
L1.5 (3MB) per core
(total:12.75MB per chip)
1B transistors
44 new instructions

MCM:
5 “z6” chips
L2 (48MB on 2SCchips)

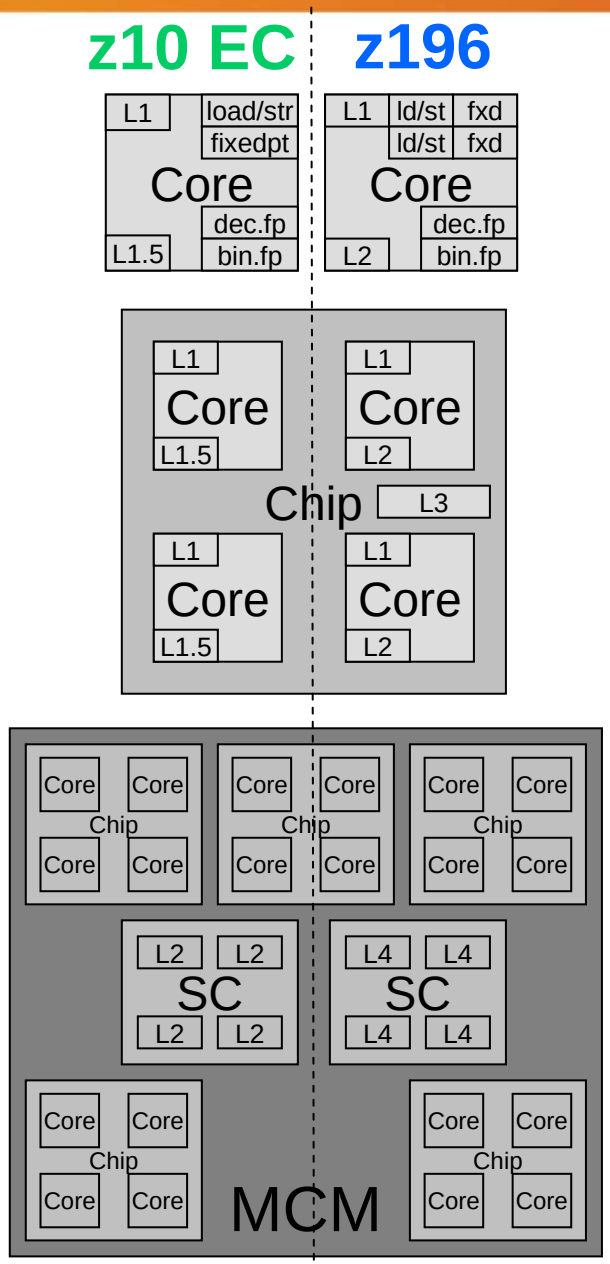

MICROPROCESSOR:
5.2GHz
Super-duper-scalar
Out of Order
L1 (192KB) and
L2 (1.5MB) per core
(new) L3 per chip (24MB)
(total:30.75MB per chip)
1.4B transistors
110 new instructions

MCM:
6 “z7” chips
L4 (192MB on 2SC chips)

Moore’s Law … bent but not broken

Topics

z10 EC

MICROPROCESSOR:
4.4GHz
Superscalar
In Order
L1 (192KB) and
L1.5 (3MB) per core
(total:12.75MB per chip)
1B transistors
44 new instructions

MCM:
5 “z6” chips
L2 (48MB on 2SCchips)

CEC:
1 to 4 books
64/77 cores
up to 1.5TB RAM
~1,000 to ~30,000MIPS
6GBps Infiniband

z196

MICROPROCESSOR:
5.2GHz
Super-duper-scalar
Out of Order
L1 (192KB) and
L2 (1.5MB) per core
(new) L3 per chip (24MB)
(total:30.75MB per chip)
1.4B transistors
110 new instructions

MCM:
6 “z7” chips
L4 (192MB on 2SC chips)

CEC:
1 to 4 books
80/96 cores
up to 3.0TB ***RAIM***
~1,500 to ~50,000MIPS
6GBps Infiniband
Hot plug I/O drawers
Optional Water Cooling
Optional DC Power

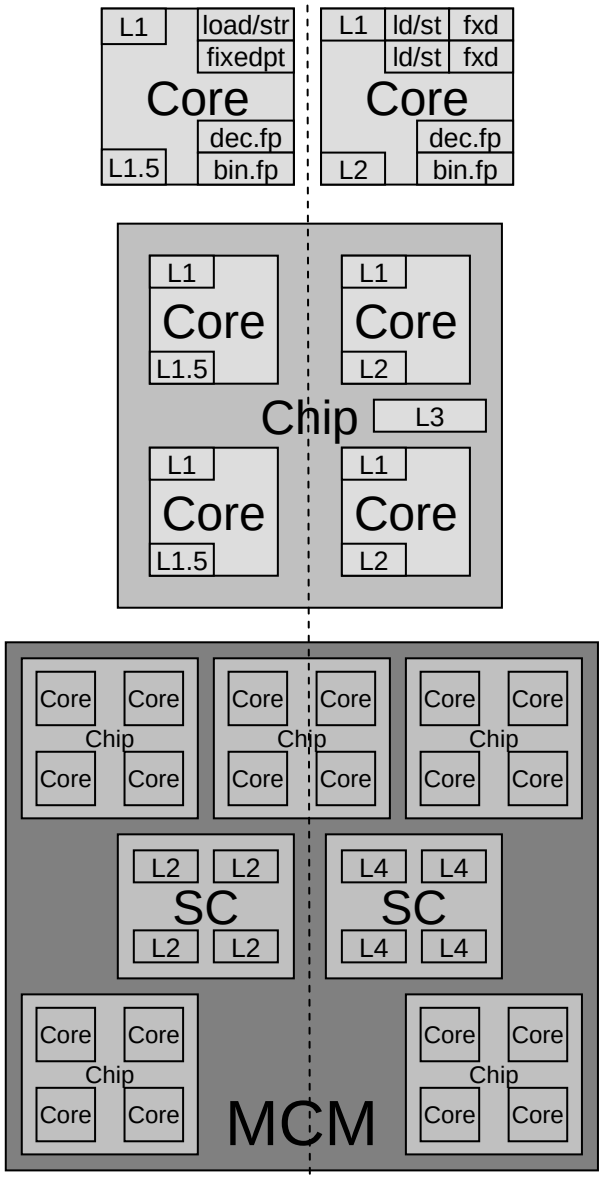

Moore’s Law … bent but not broken

Topics

zEnterprise Server – z196 or z114

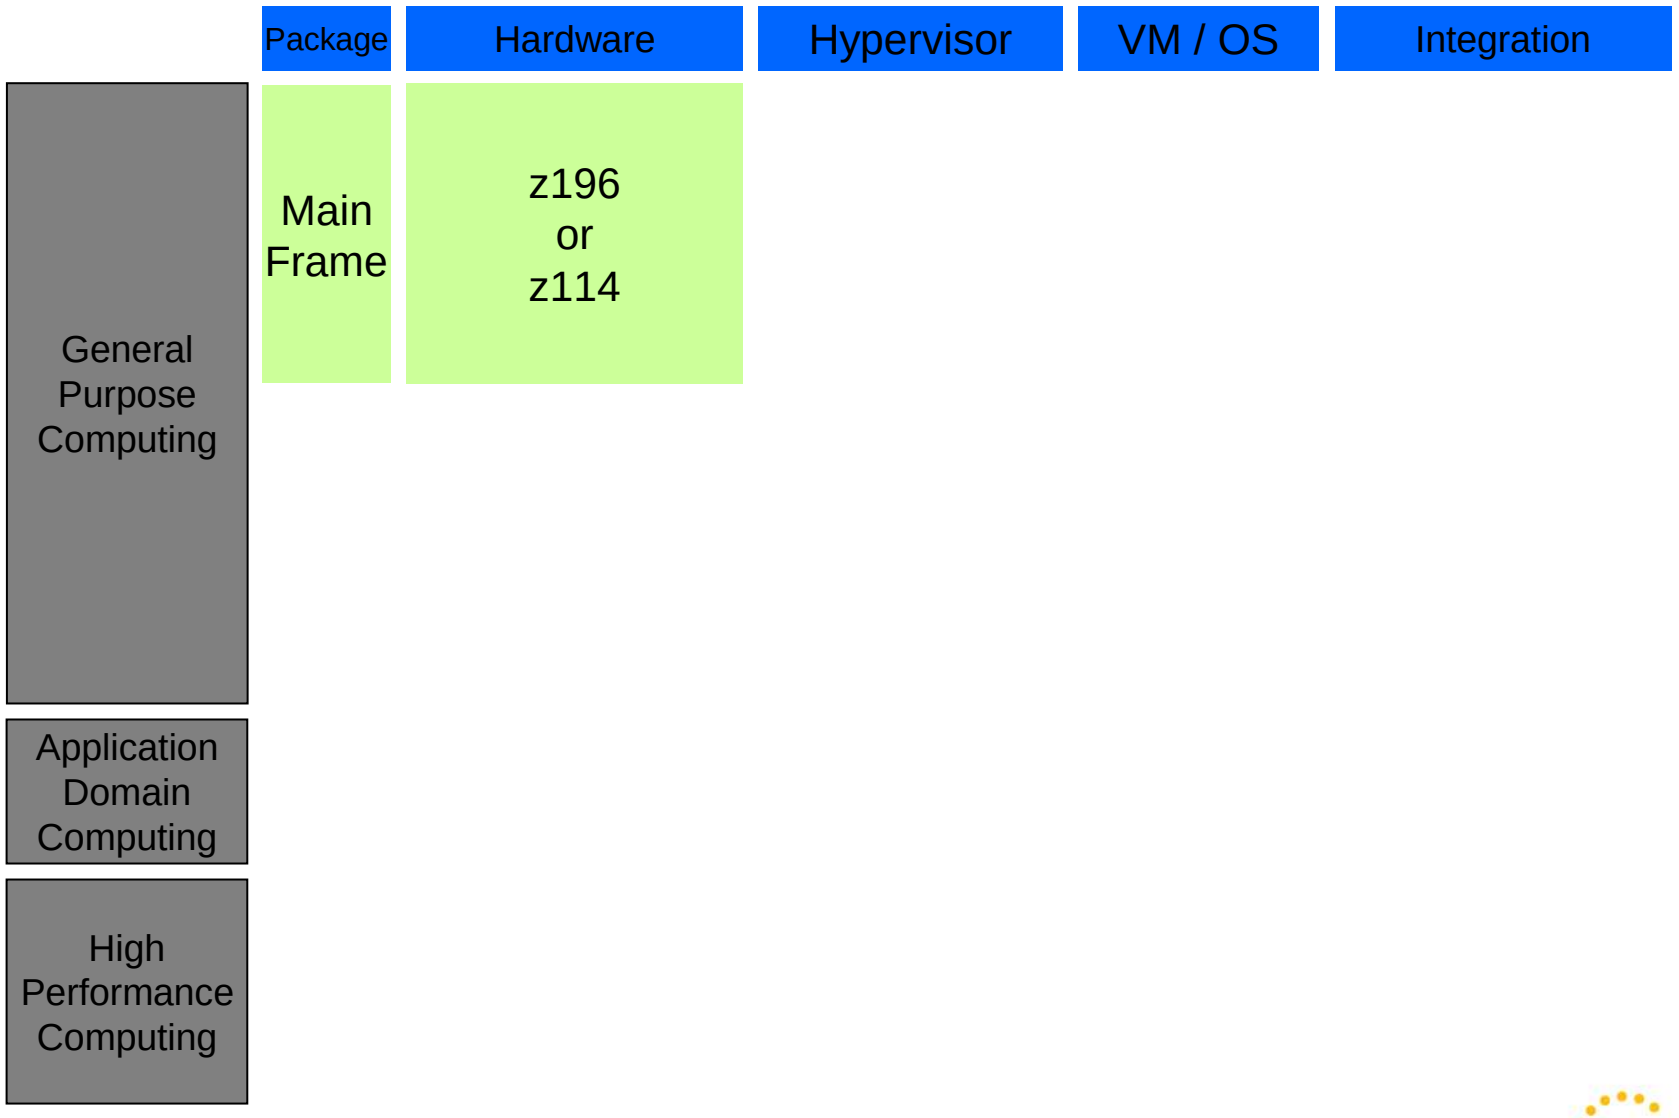

Topics

zEnterprise Server – z196

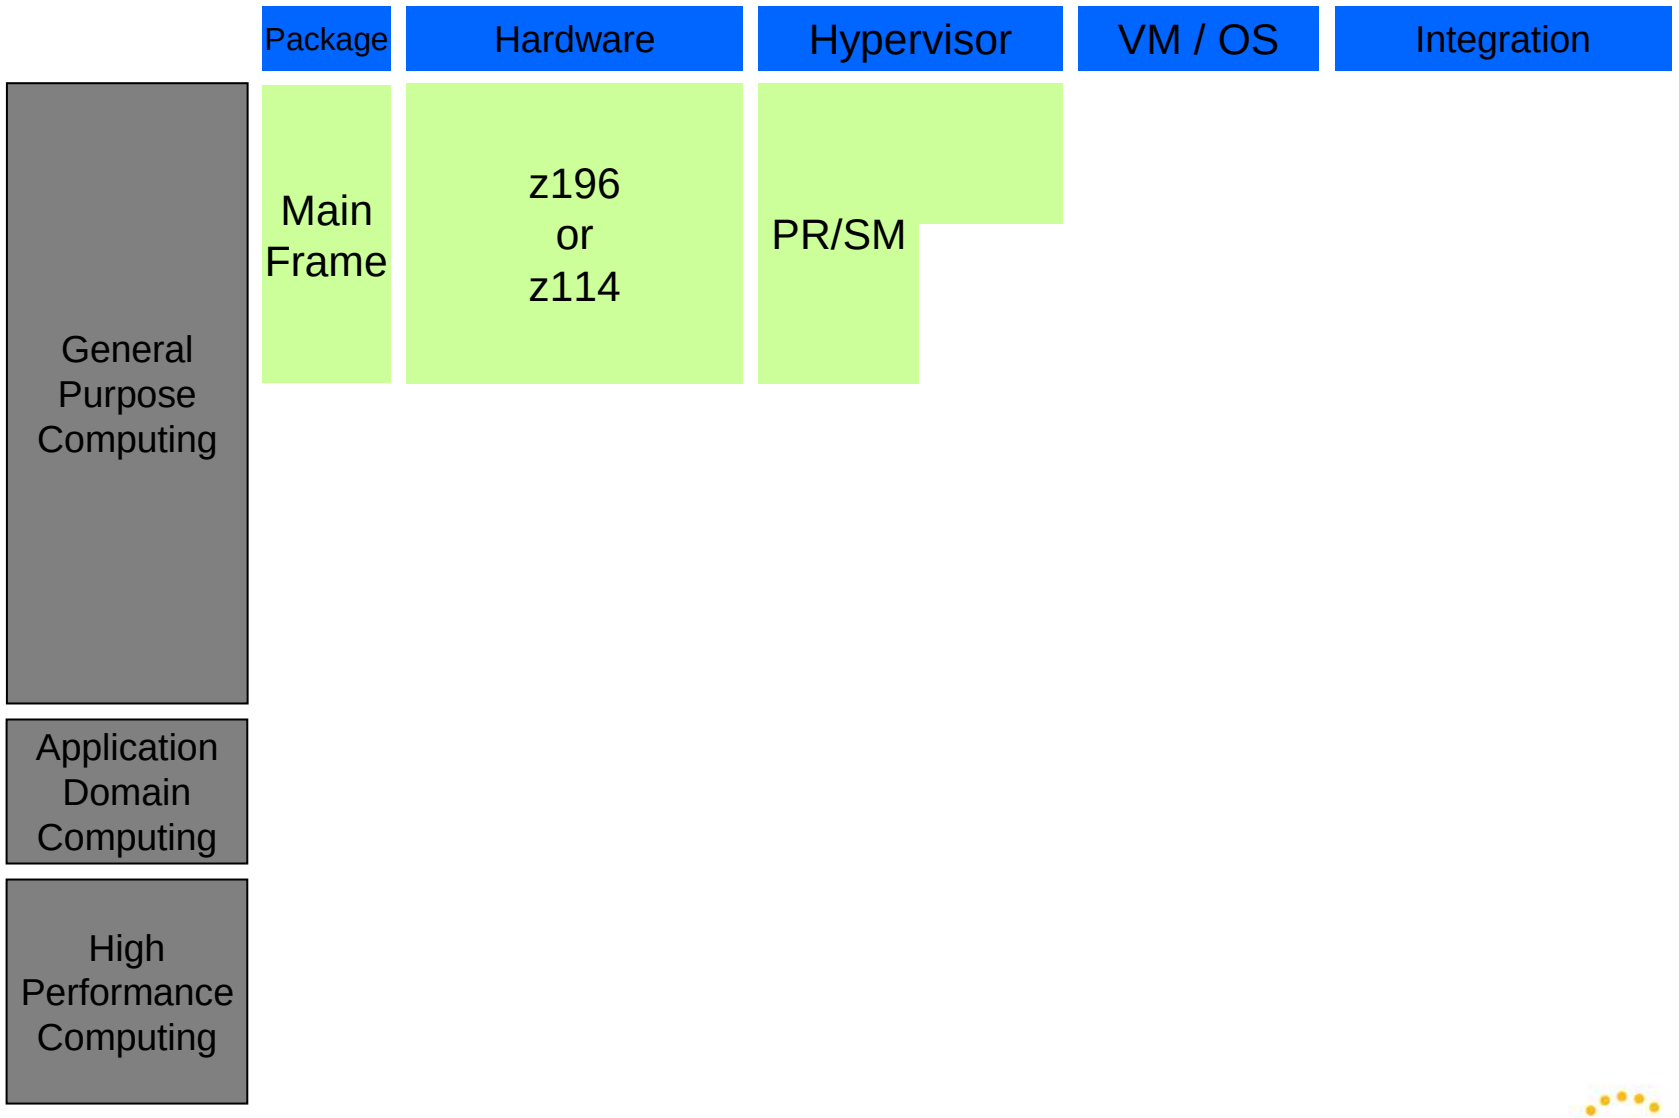

Topics

zEnterprise Server – z196

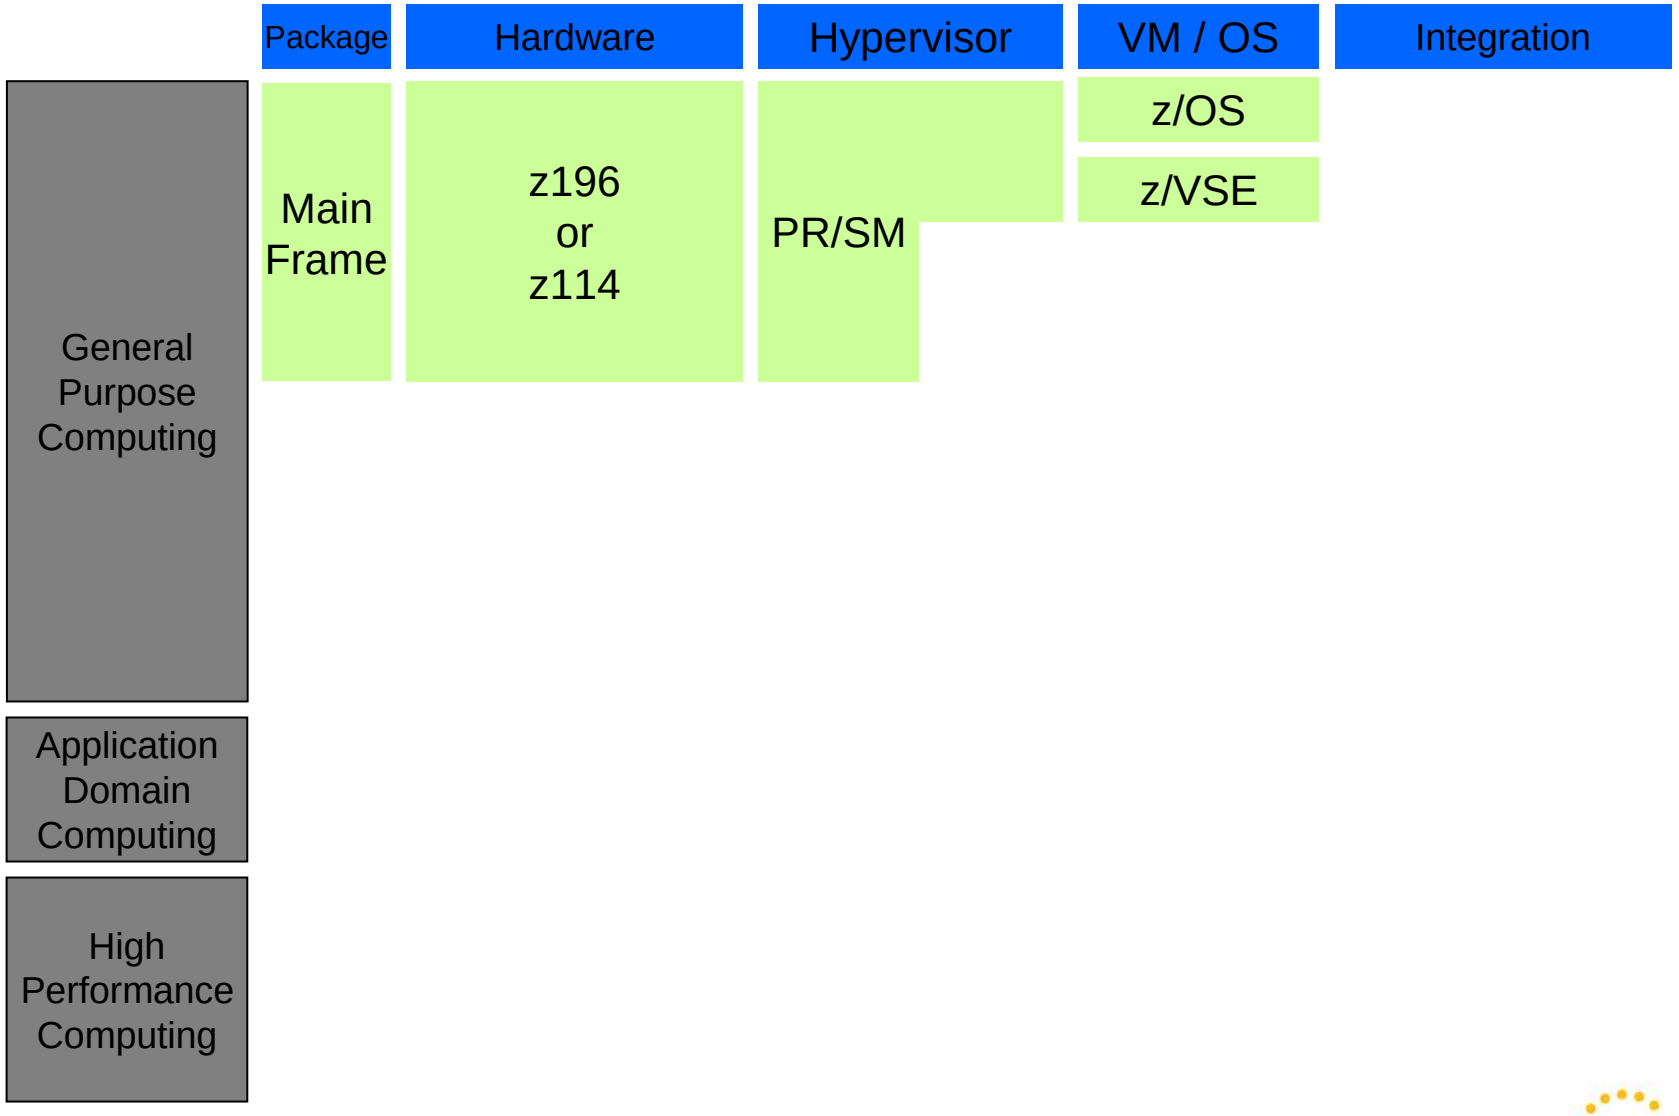

Topics

zEnterprise Server - z196

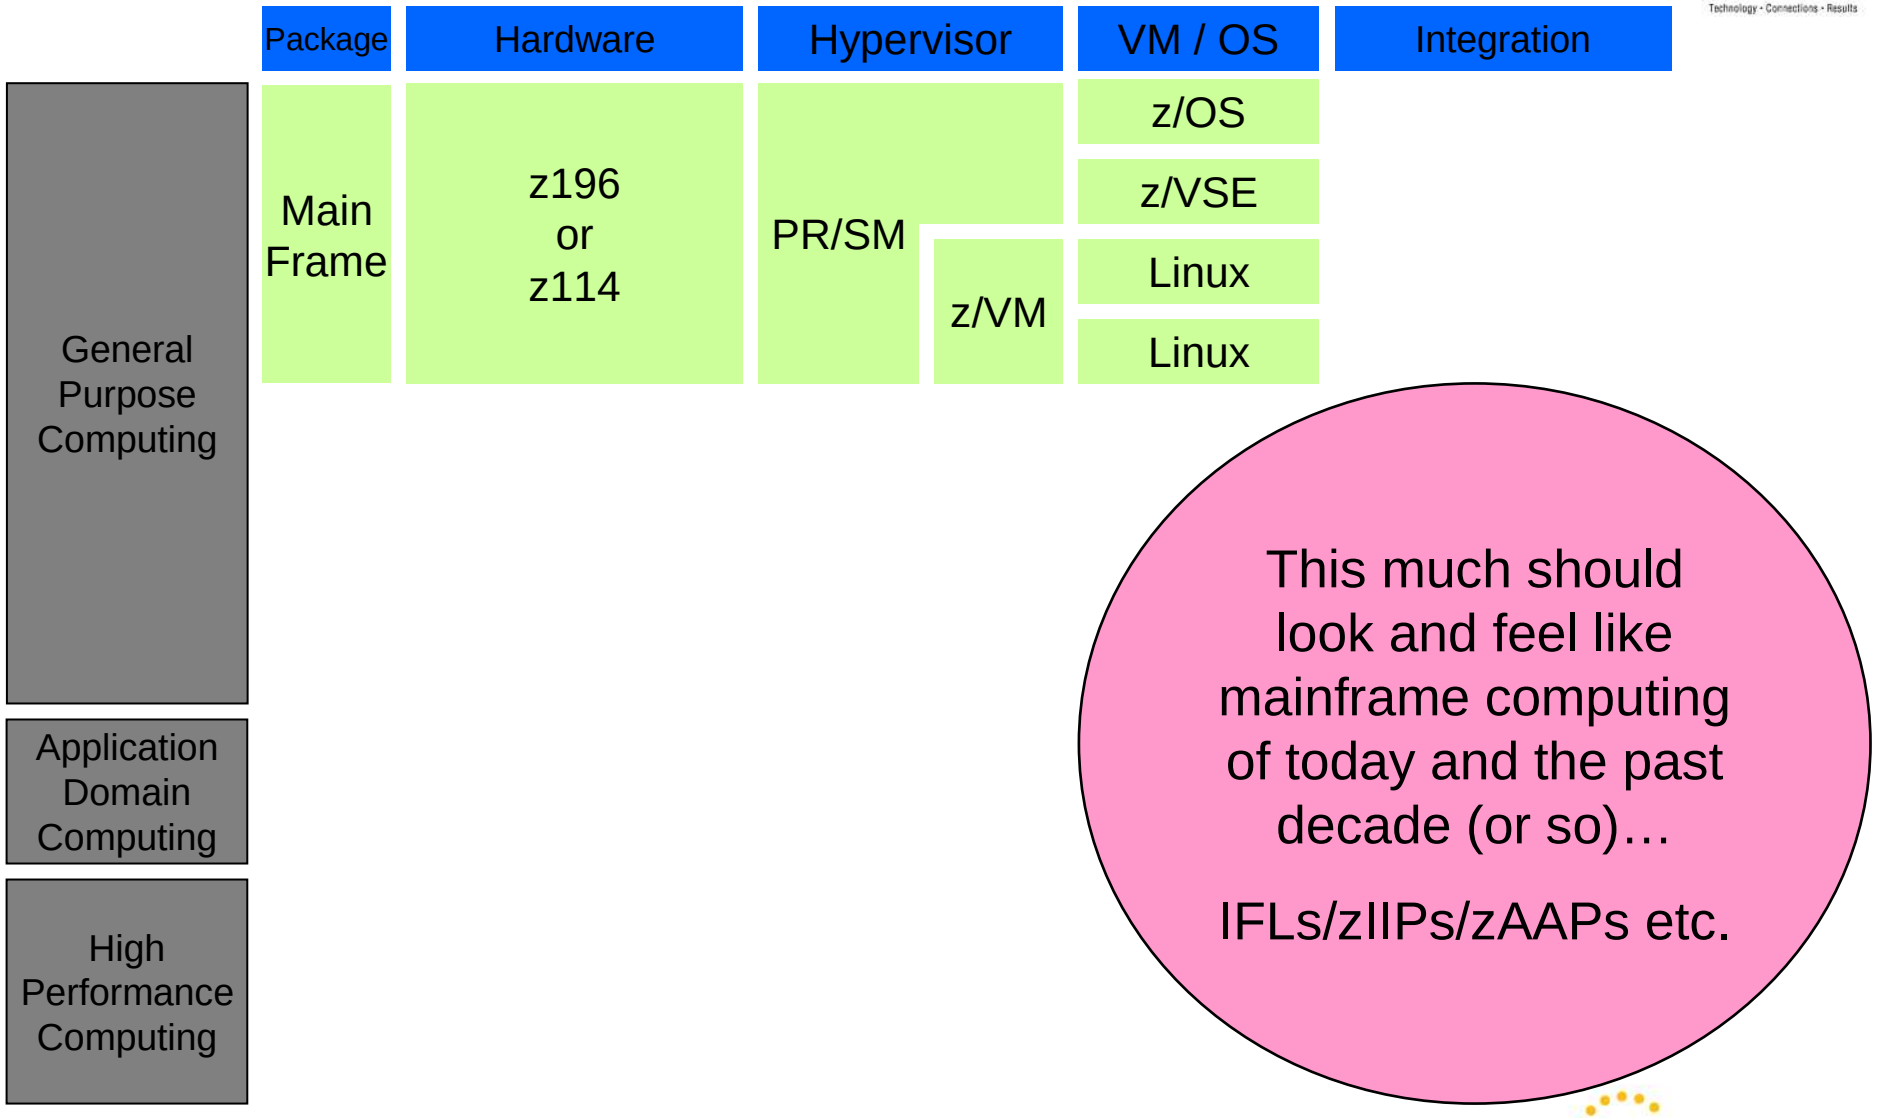

Topics

zEnterprise Server – zBX

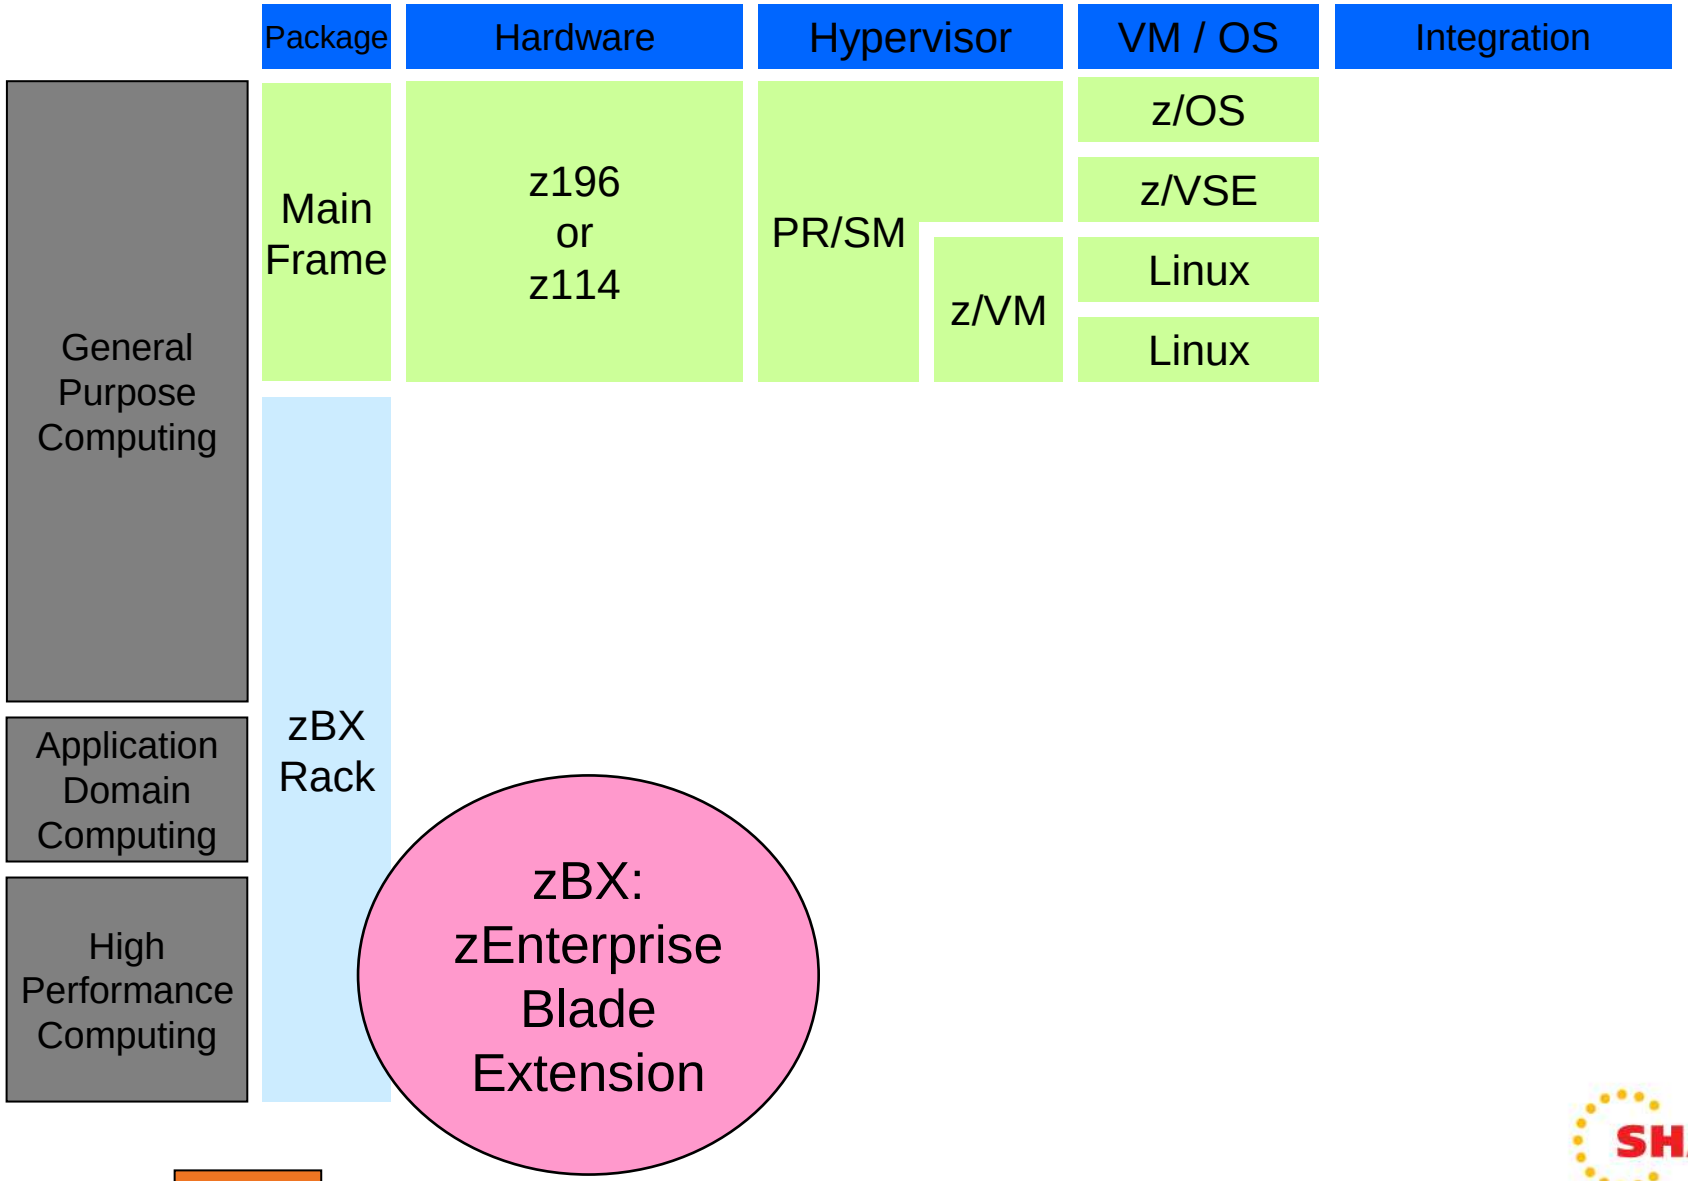

Topics

zEnterprise Server – BladeCenter Chassis

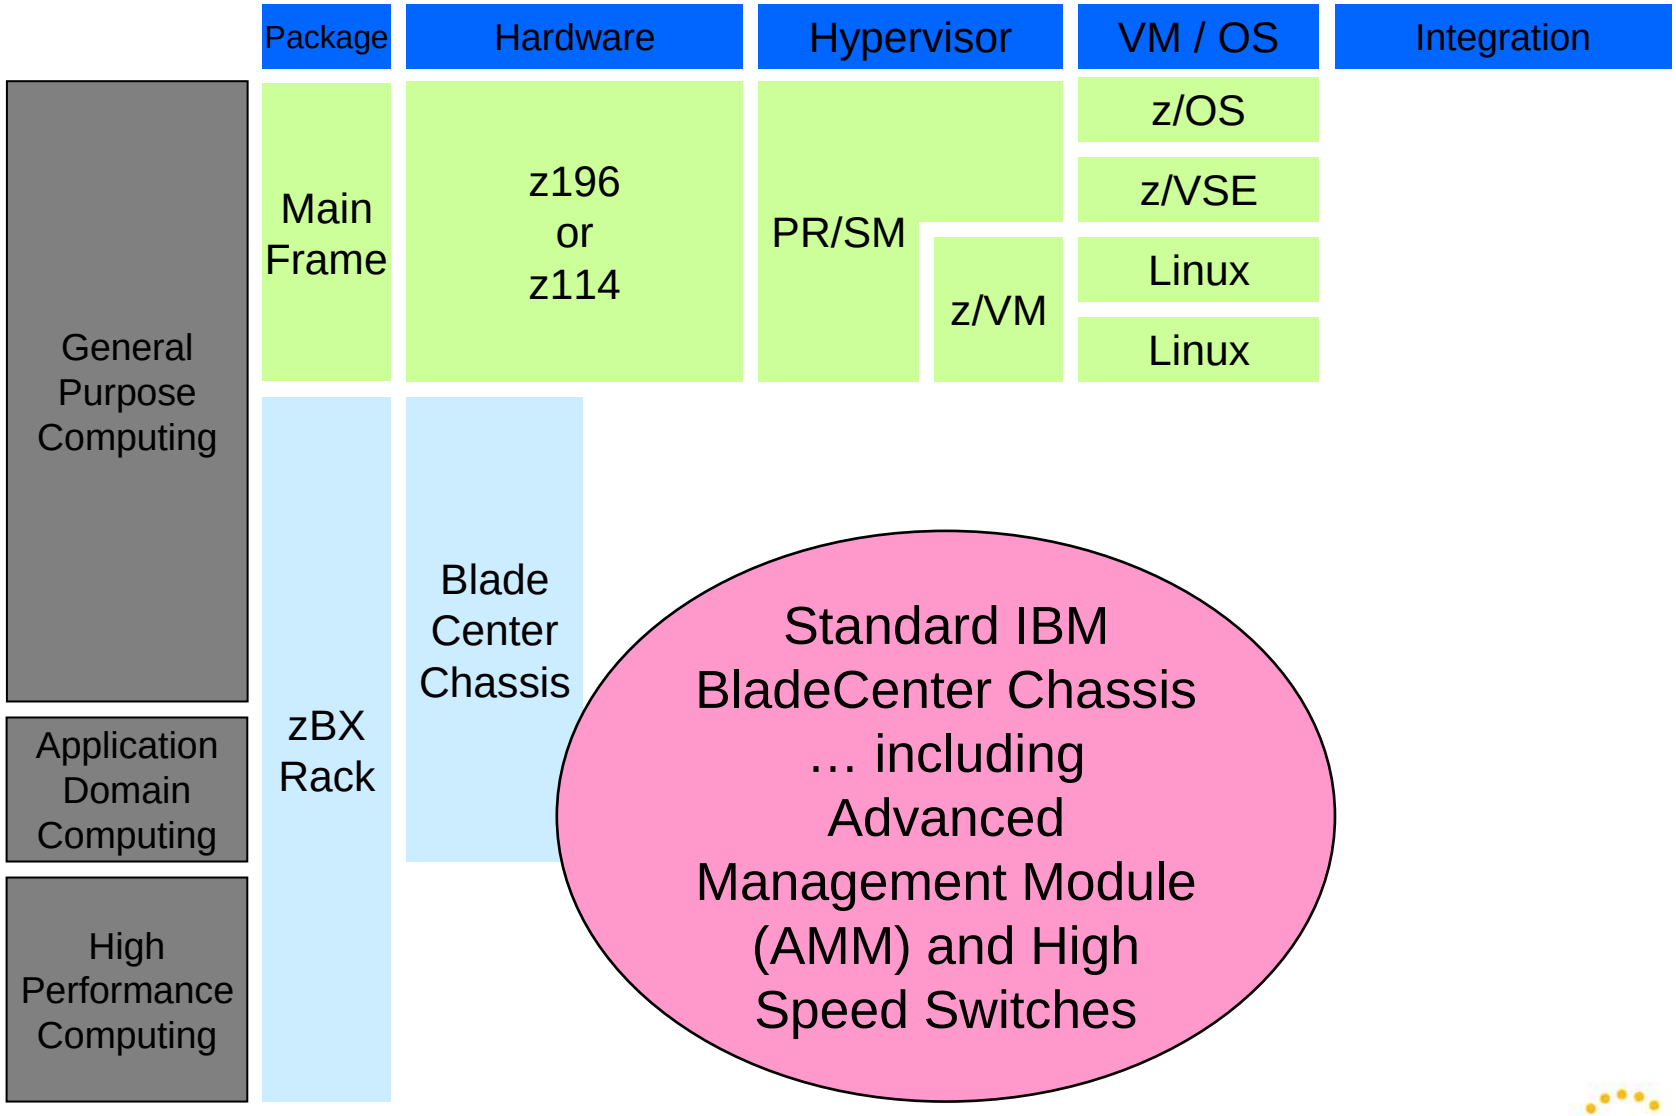

Topics

zEnterprise Server – POWER (PS701)

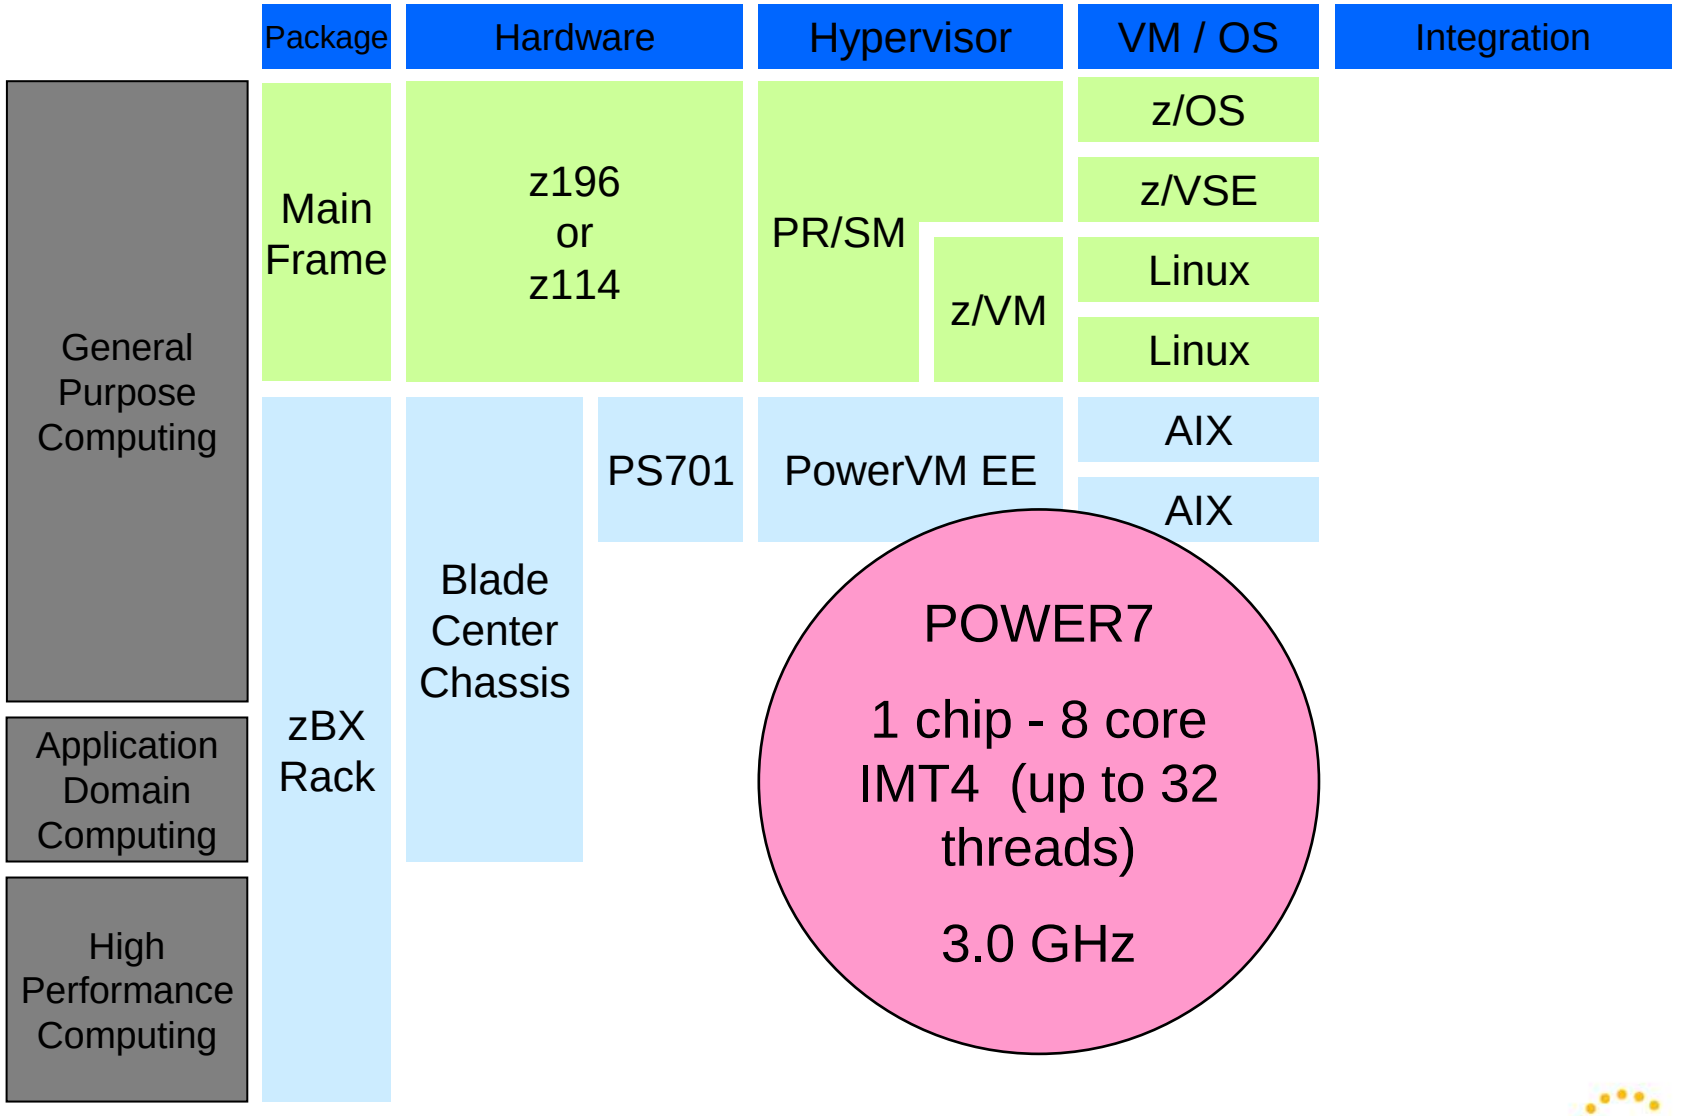

Topics

zEnterprise Server – x86 (HX5)

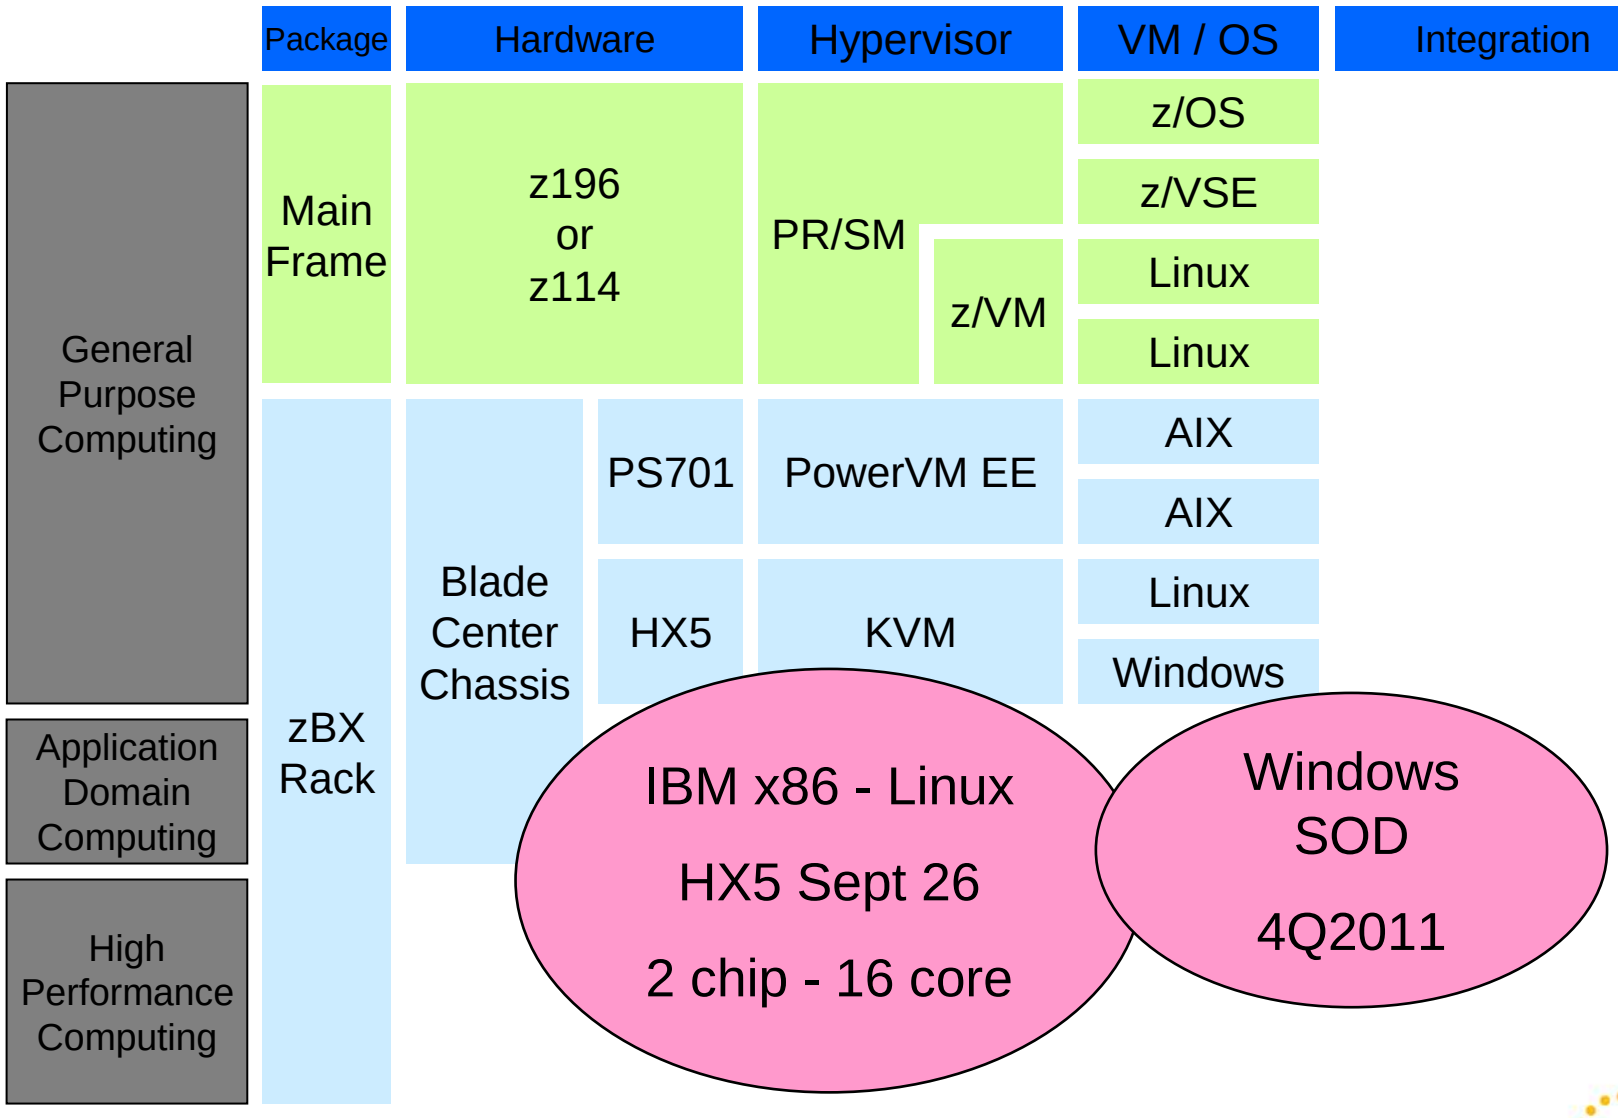

Topics

zEnterprise Server - DataPower

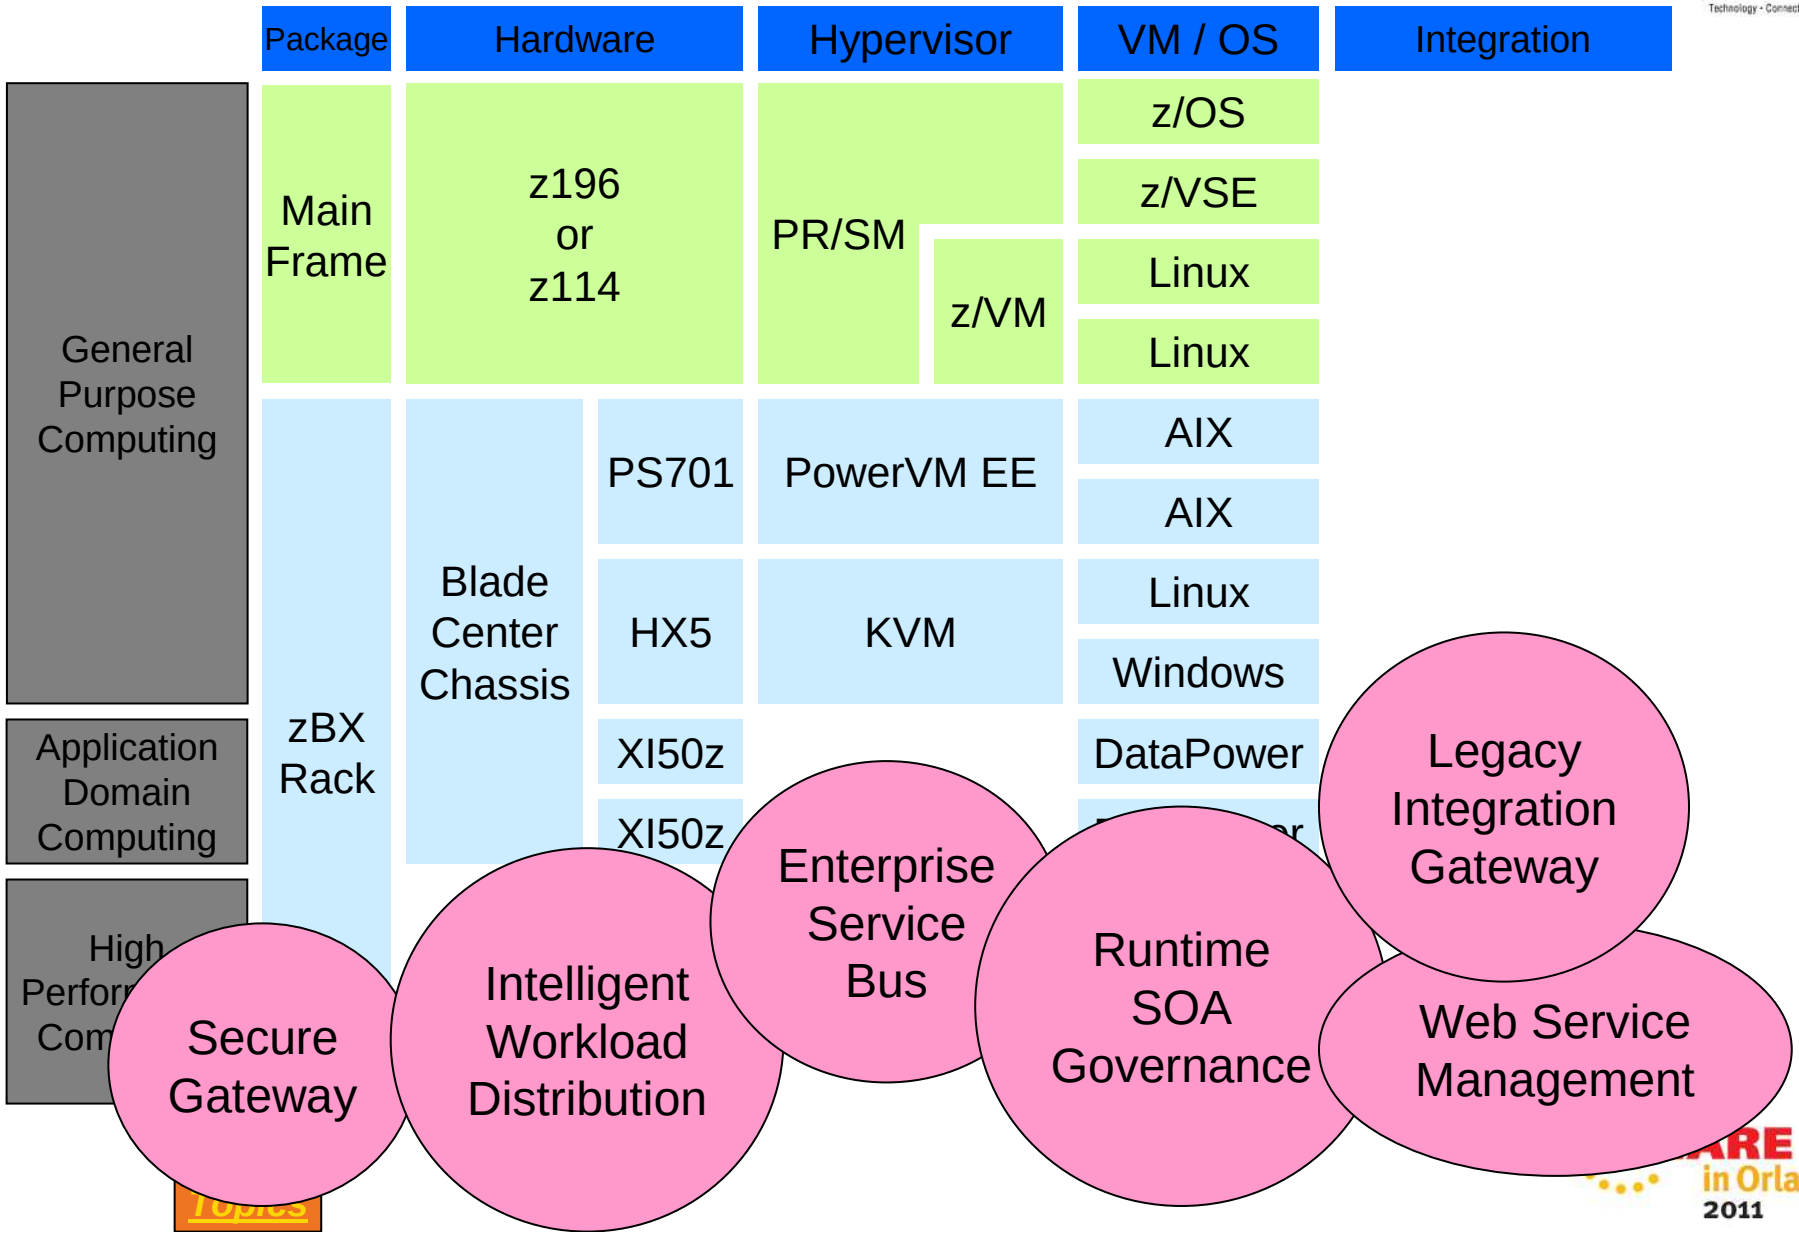

zEnterprise Server - ISAO
IBM Smart Analytics Optimizer

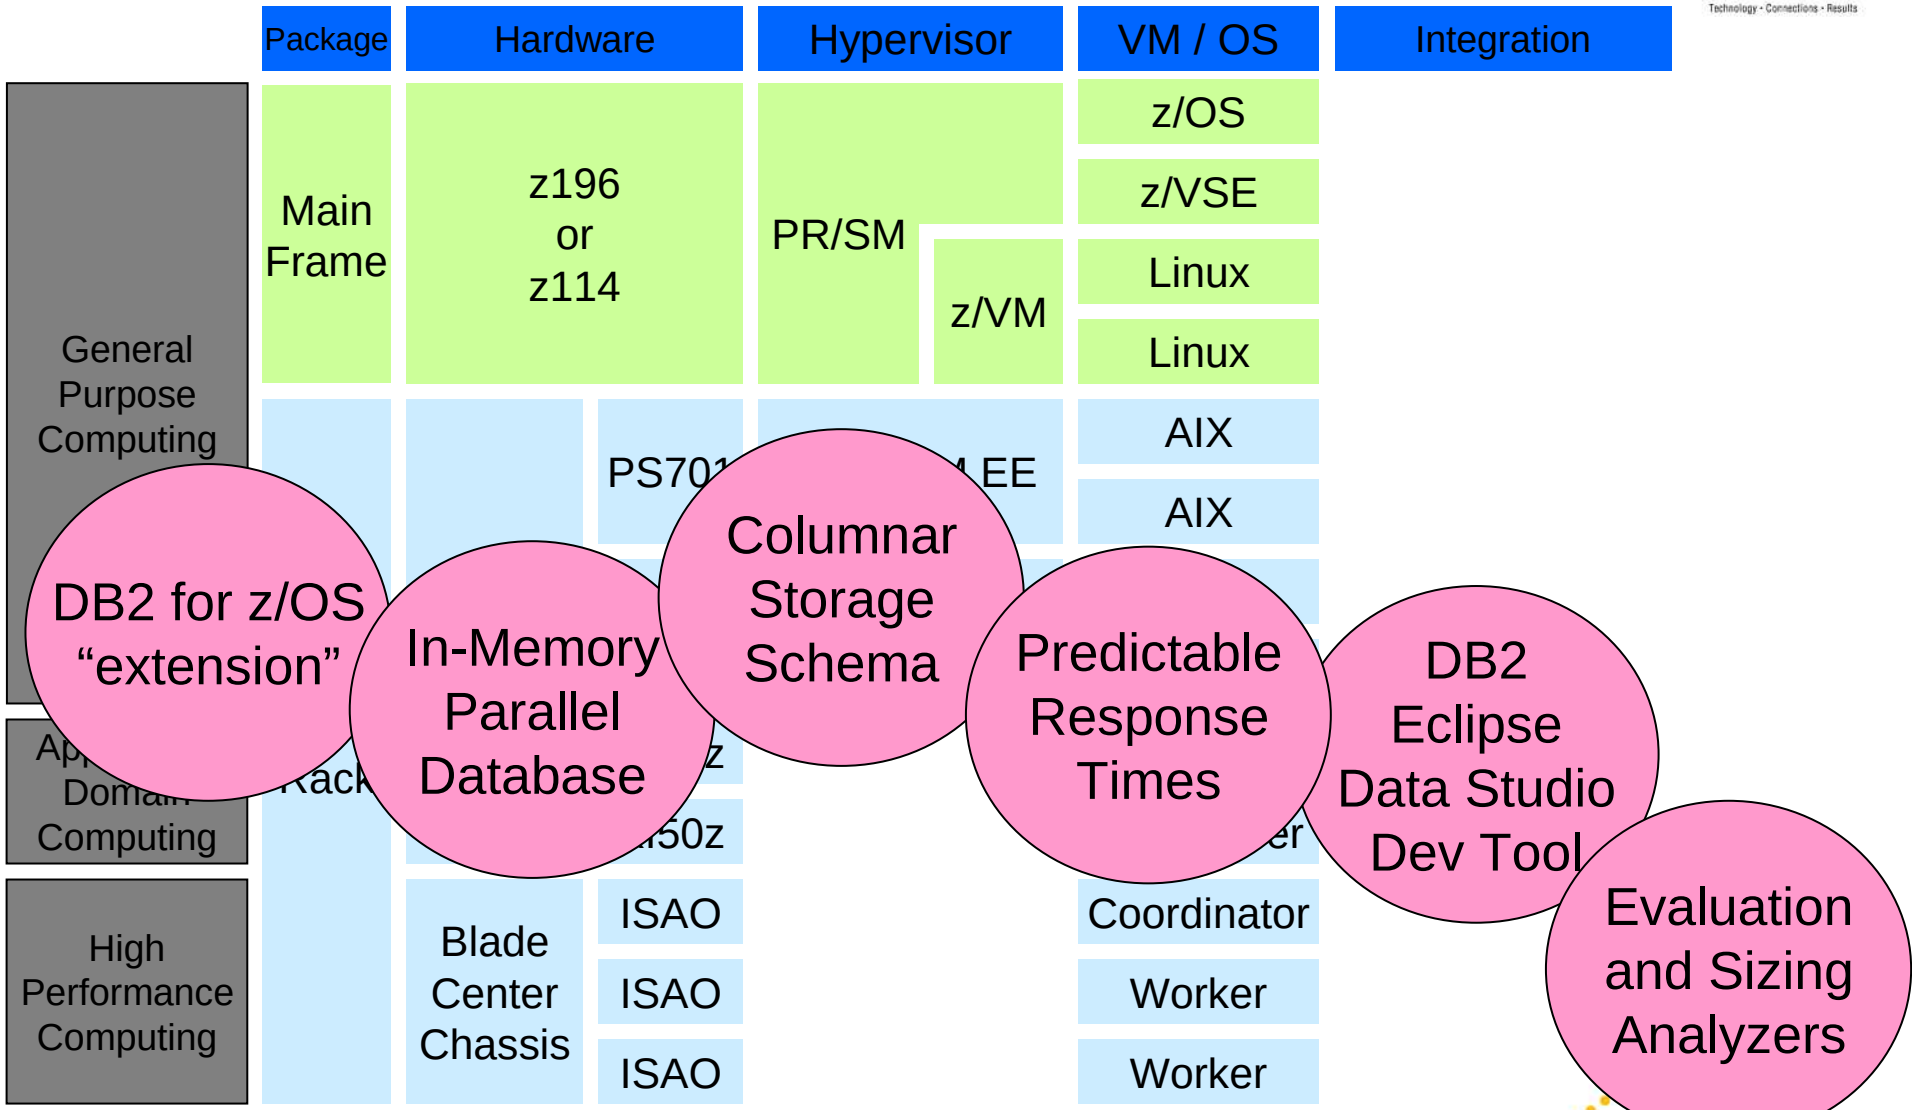

Topics

zEnterprise Server – Storage Area Network

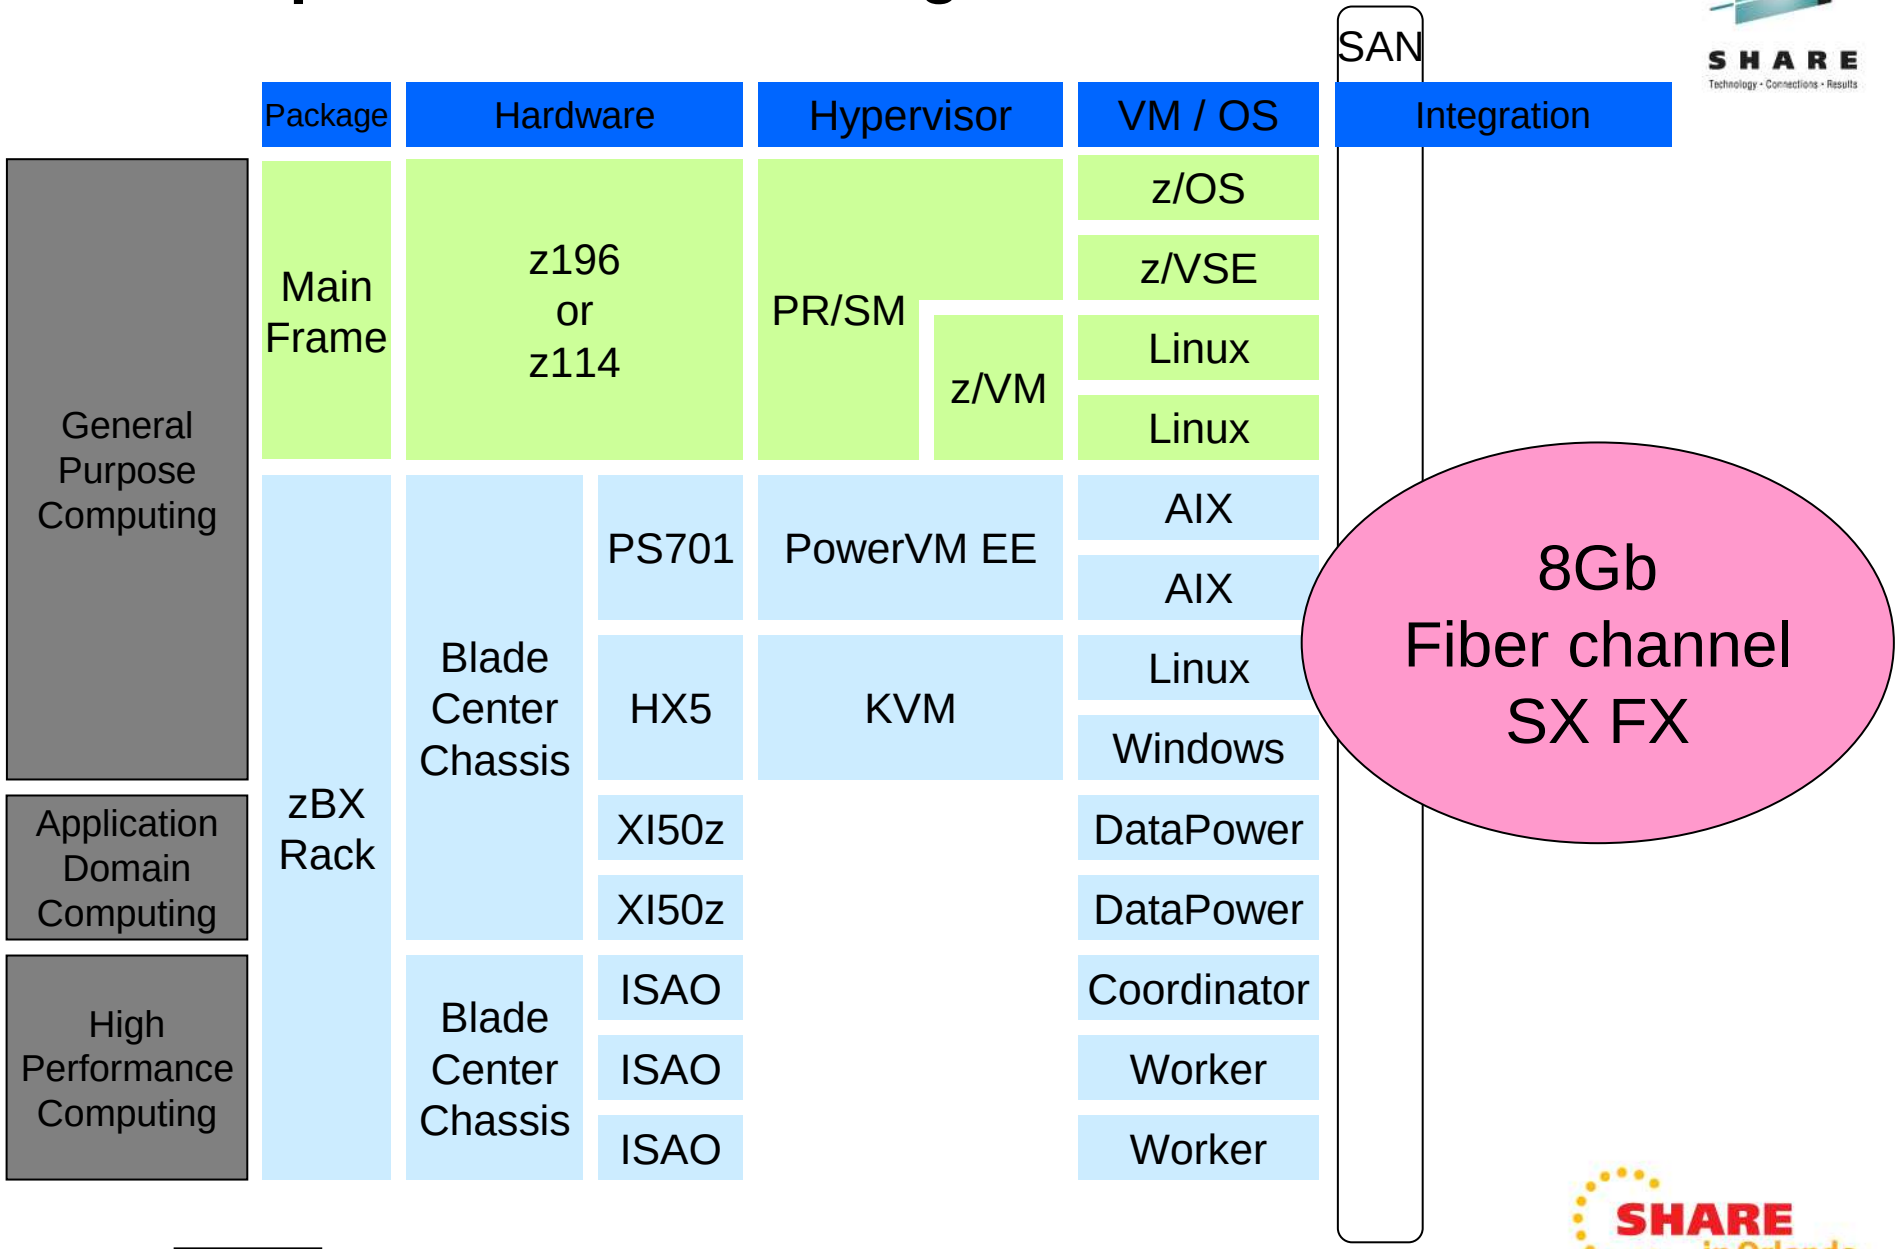

Topics

zEnterprise Server – Data Network

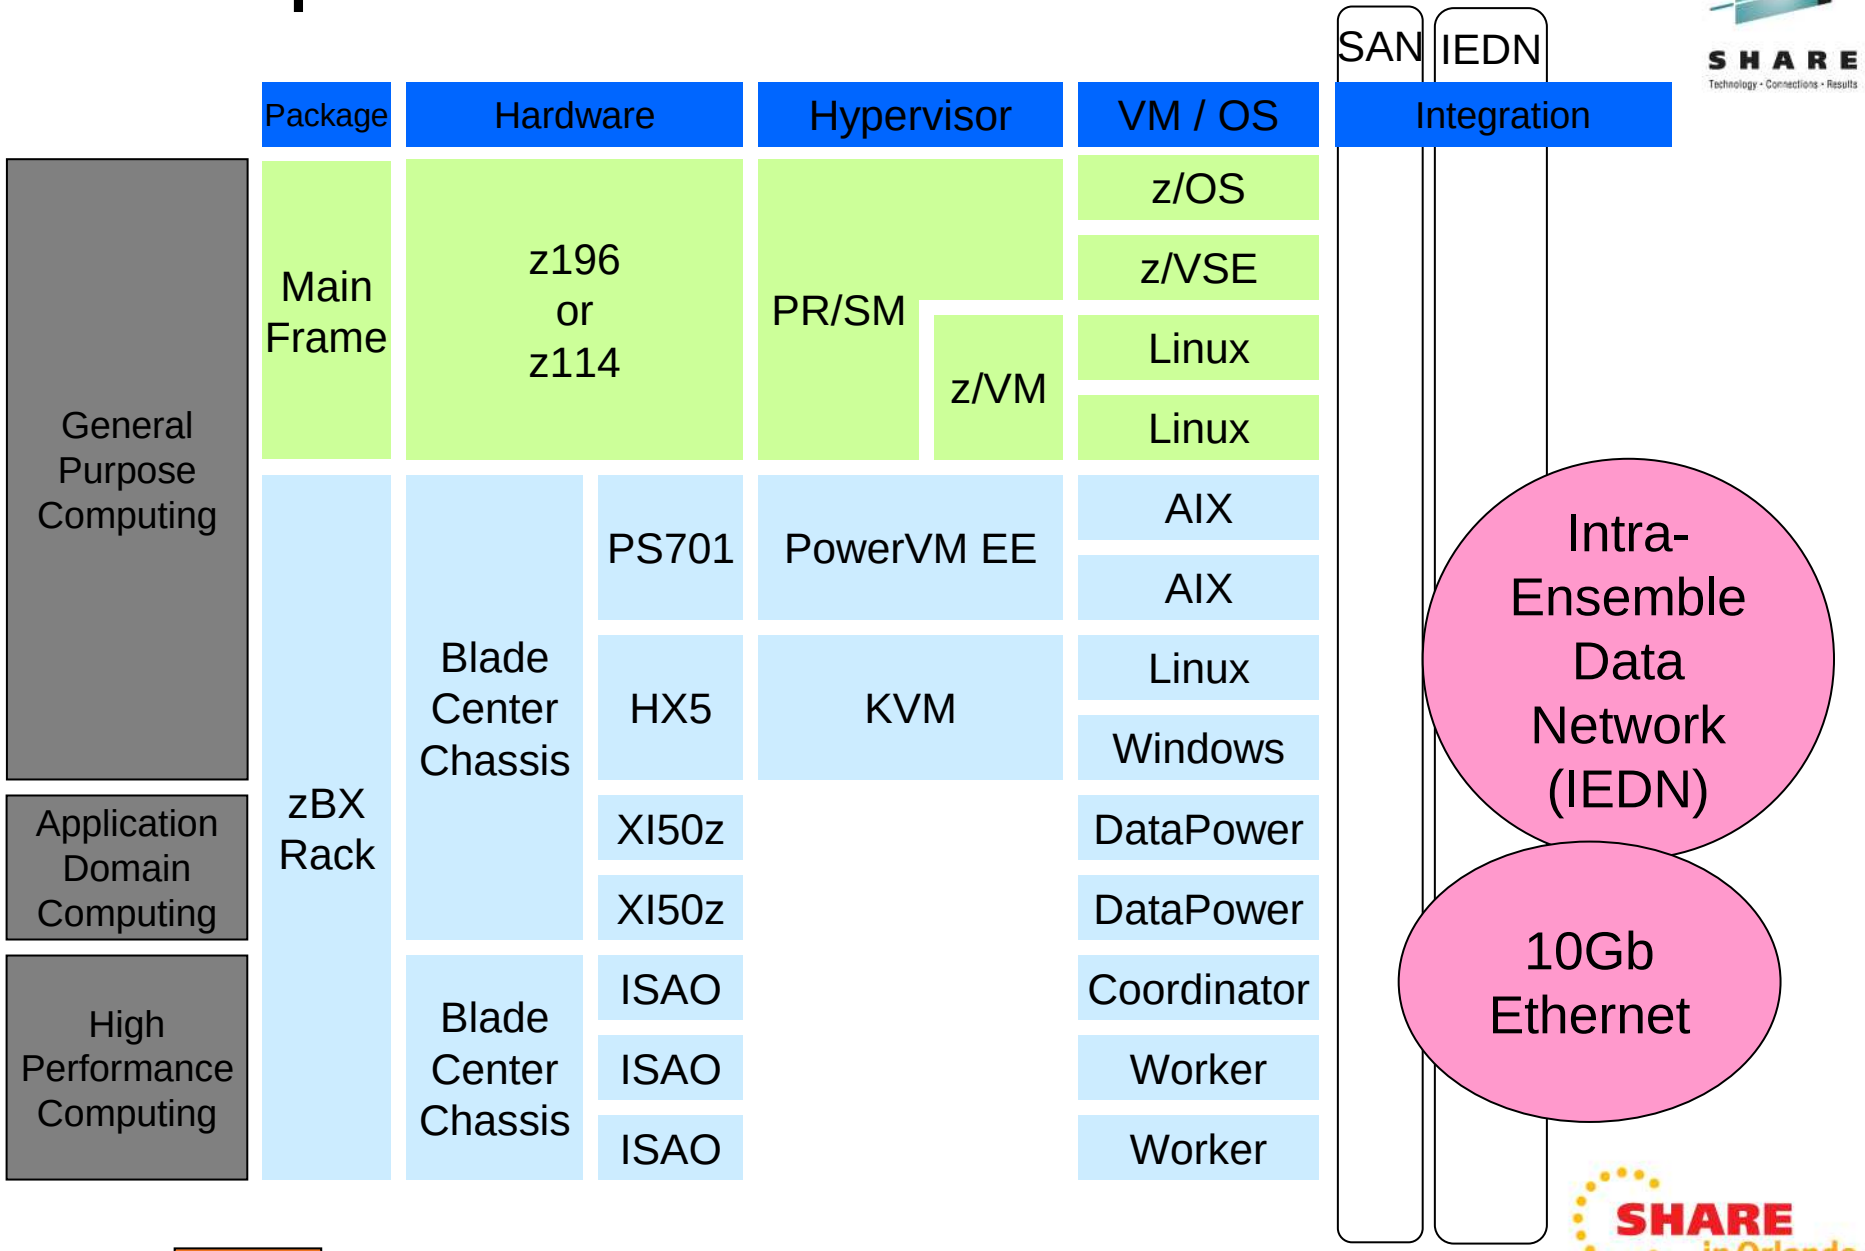

Topics

zEnterprise Server – Server Network

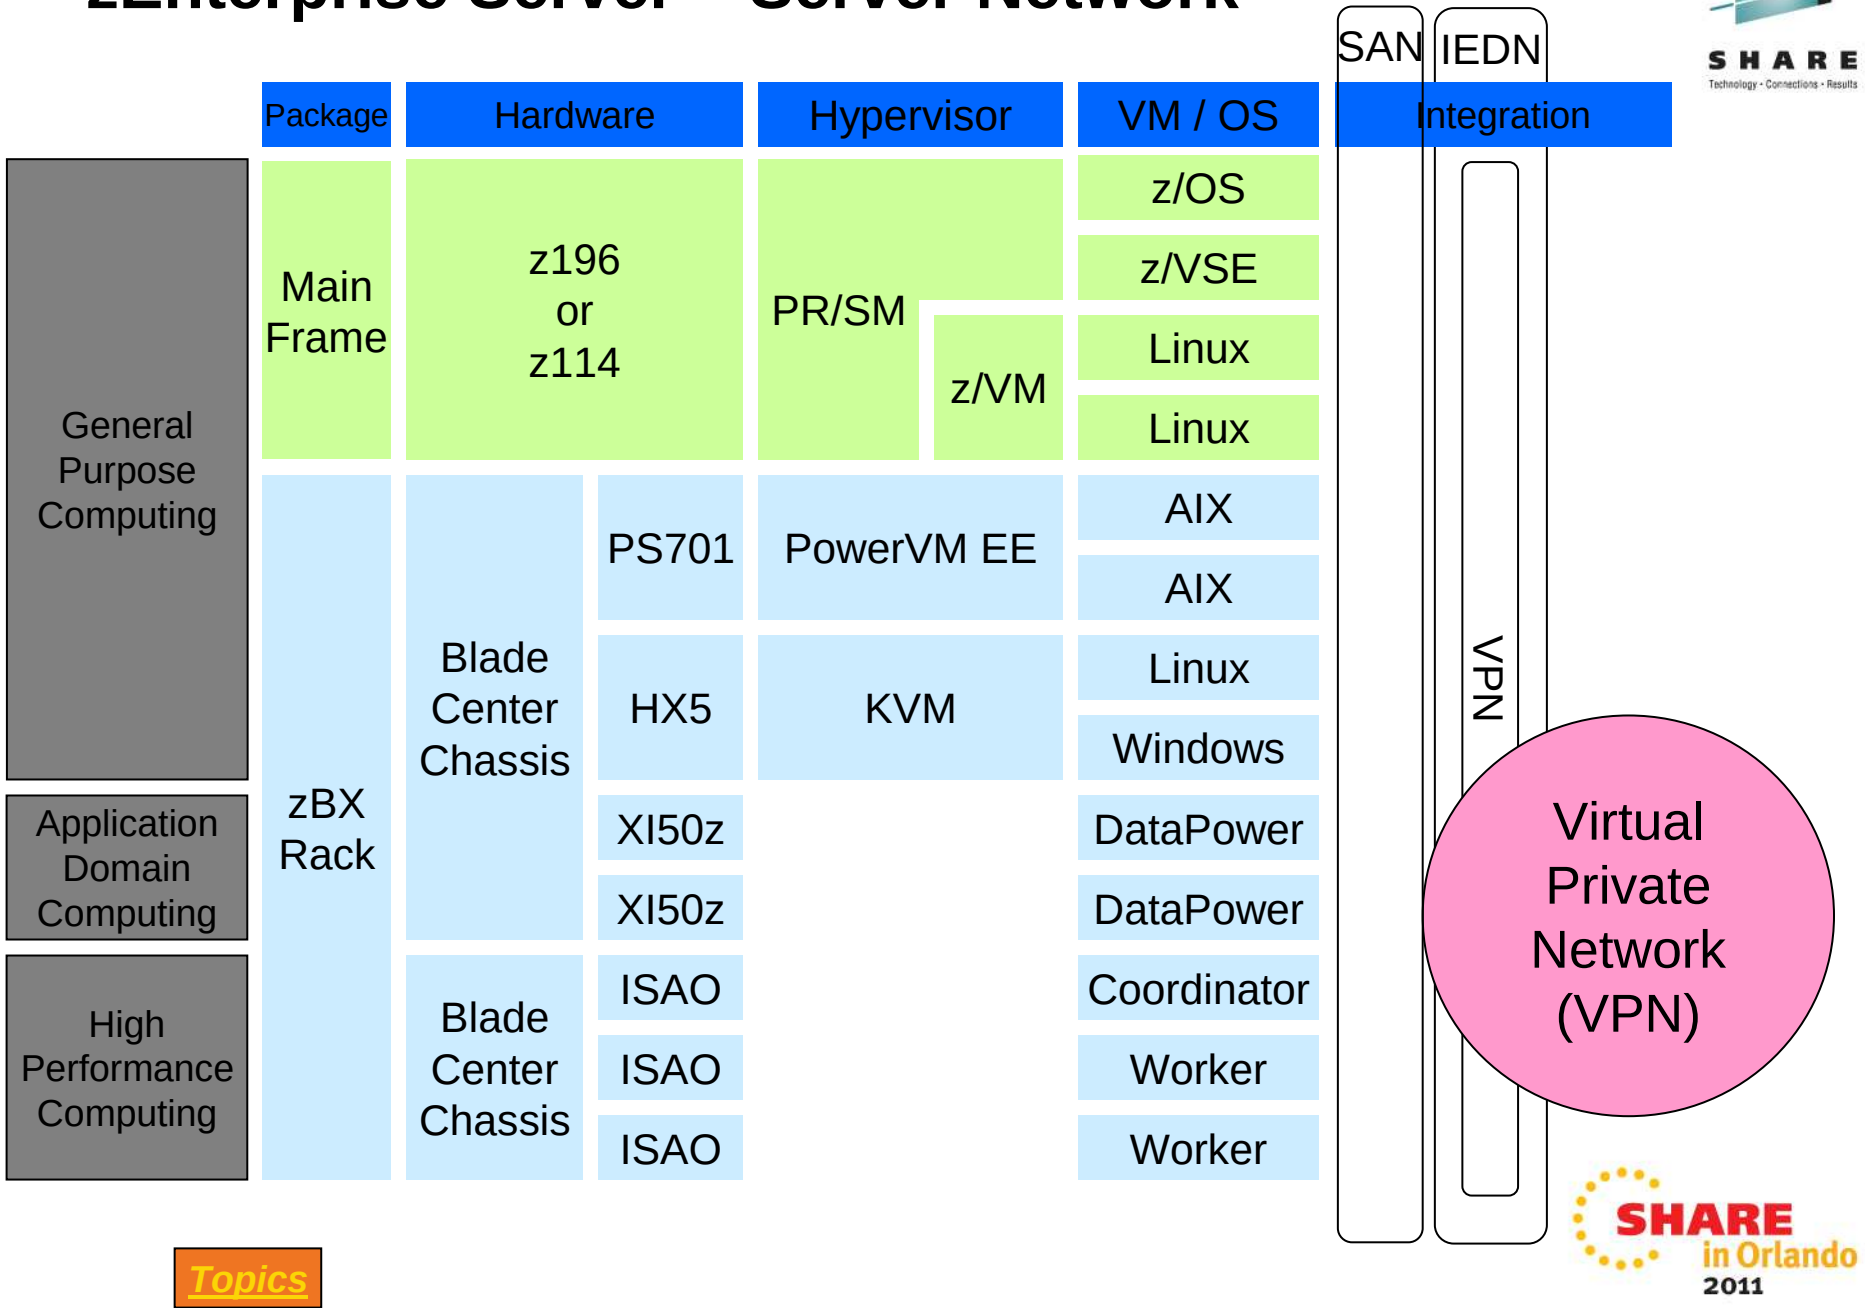

Topics

zEnterprise Server – Management Network

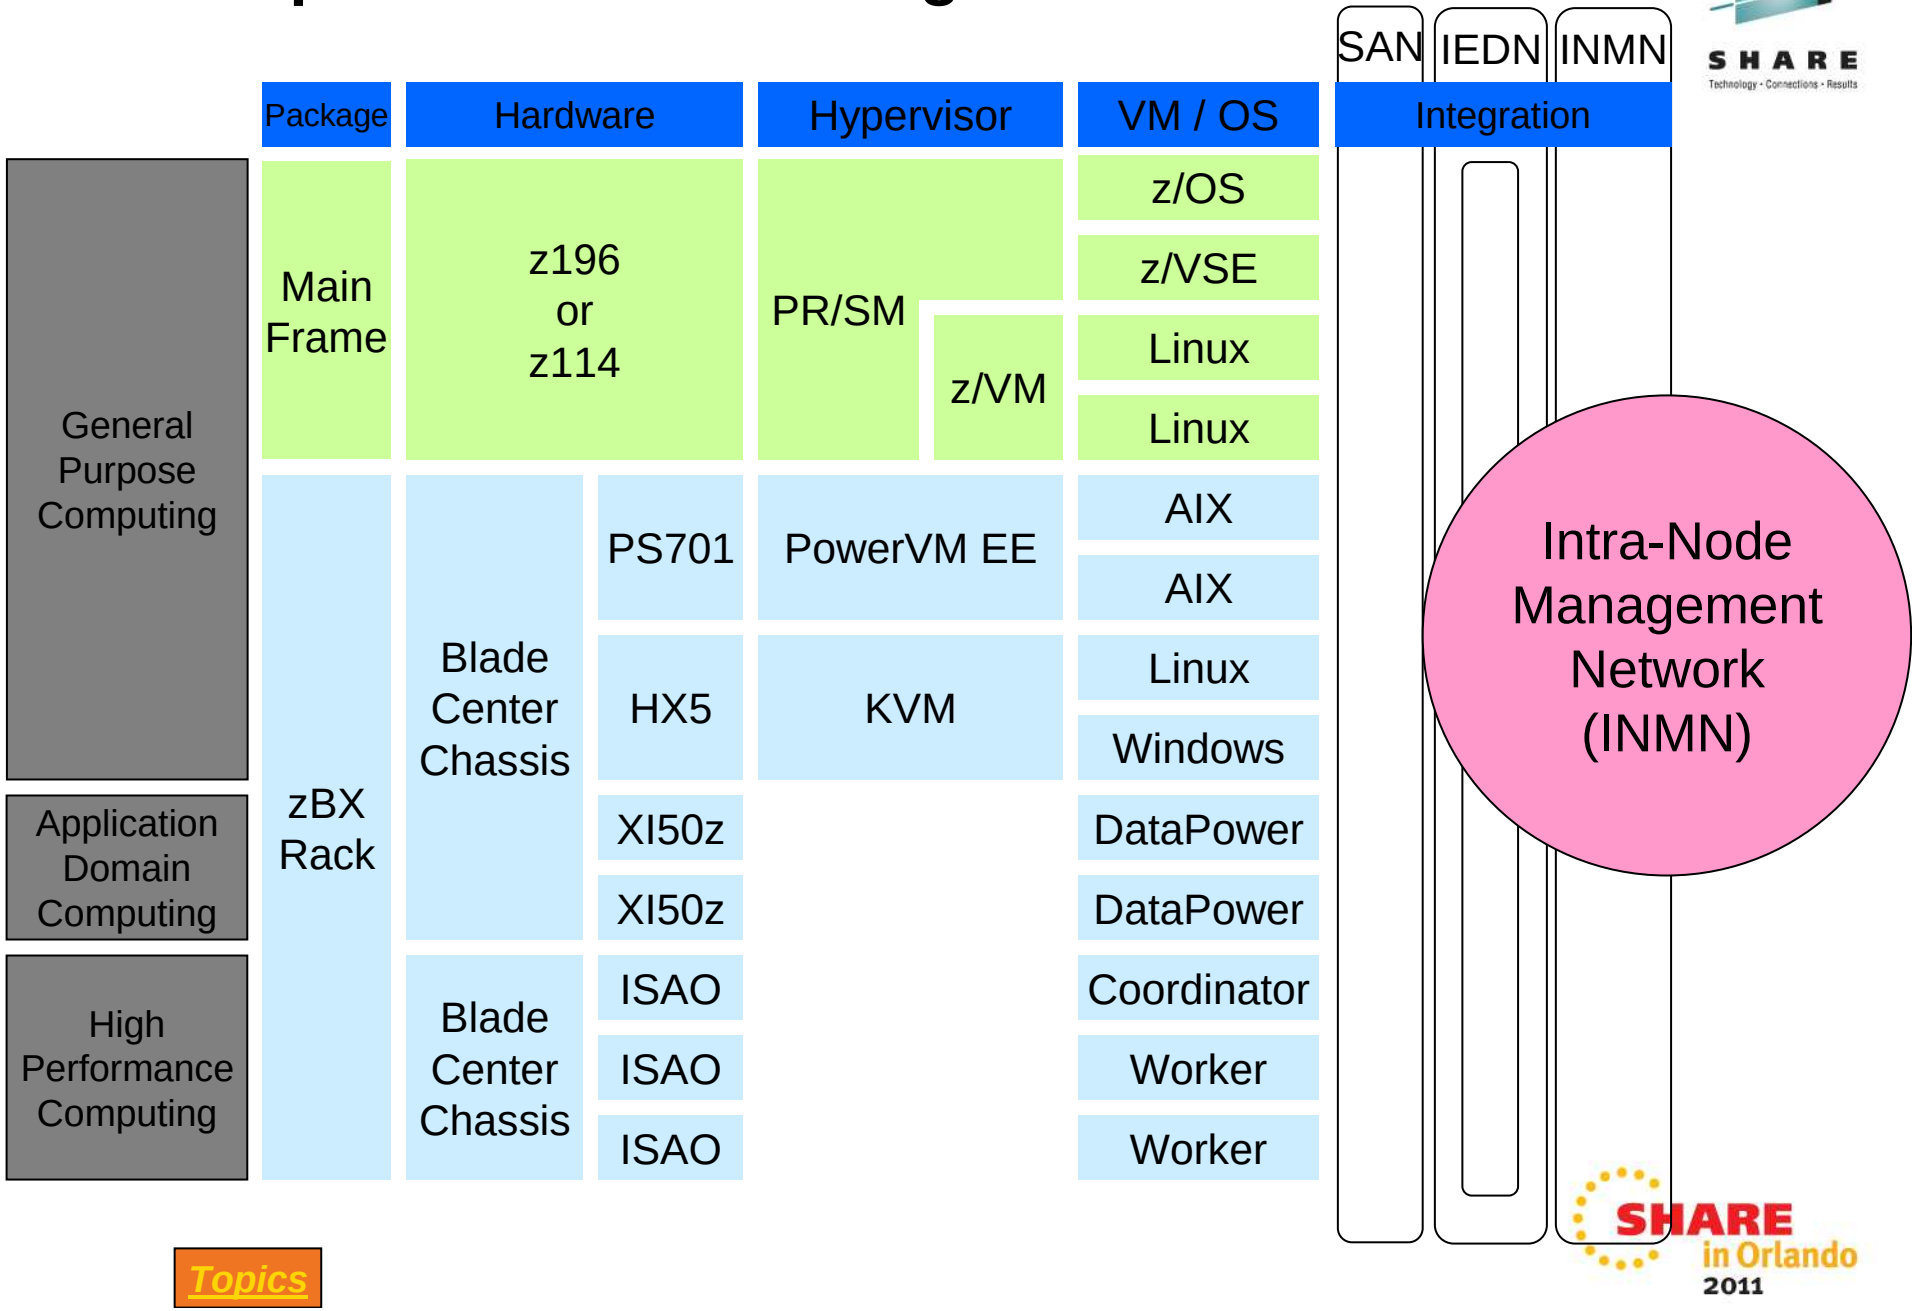

Topics

zEnterprise Server - zManager

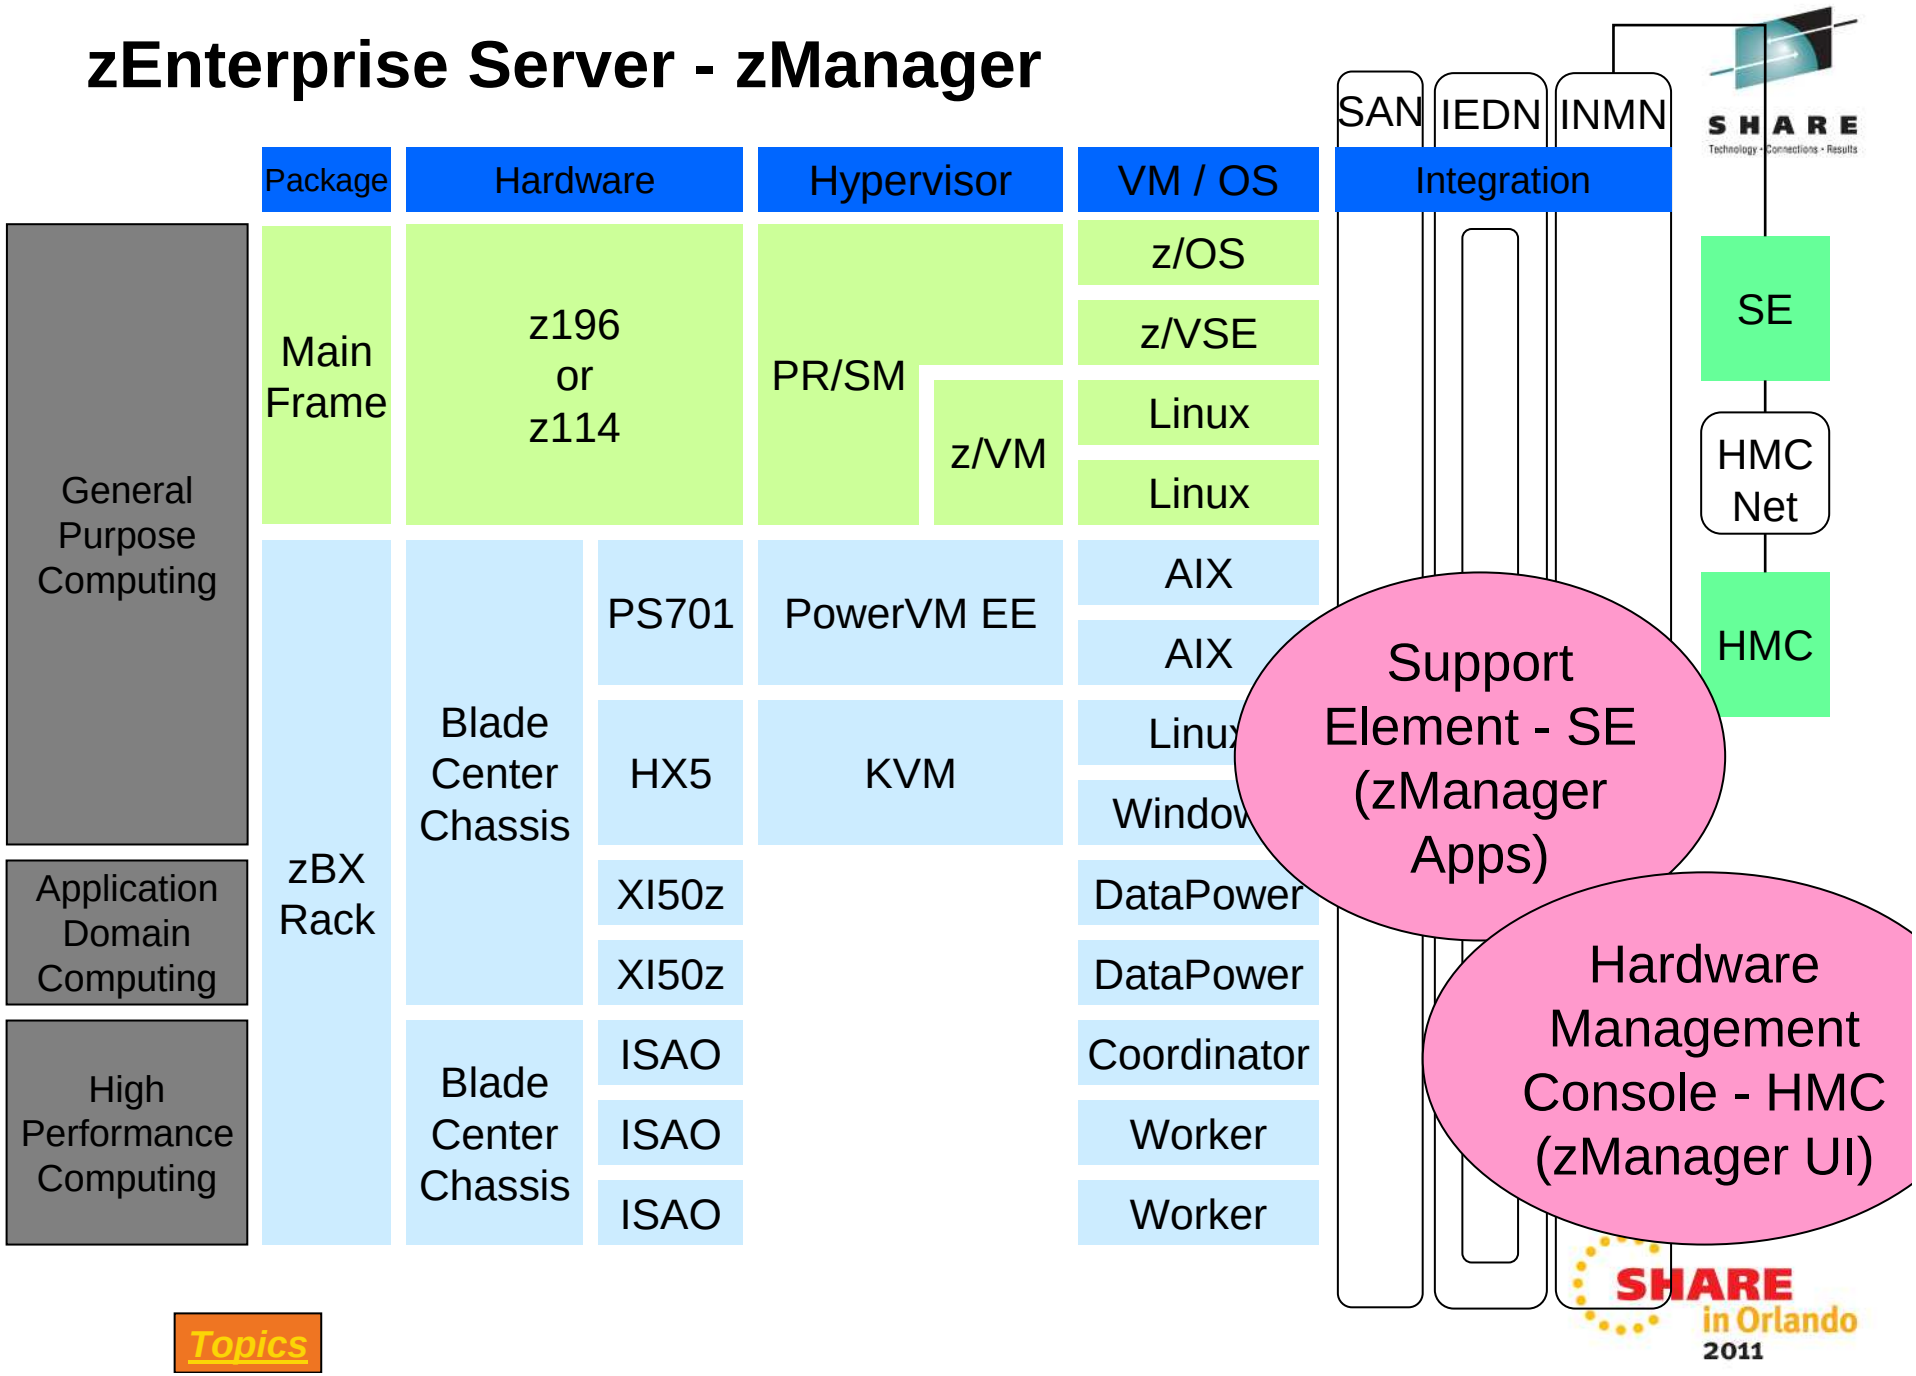

Topics

zEnterprise Server – zManager APIs

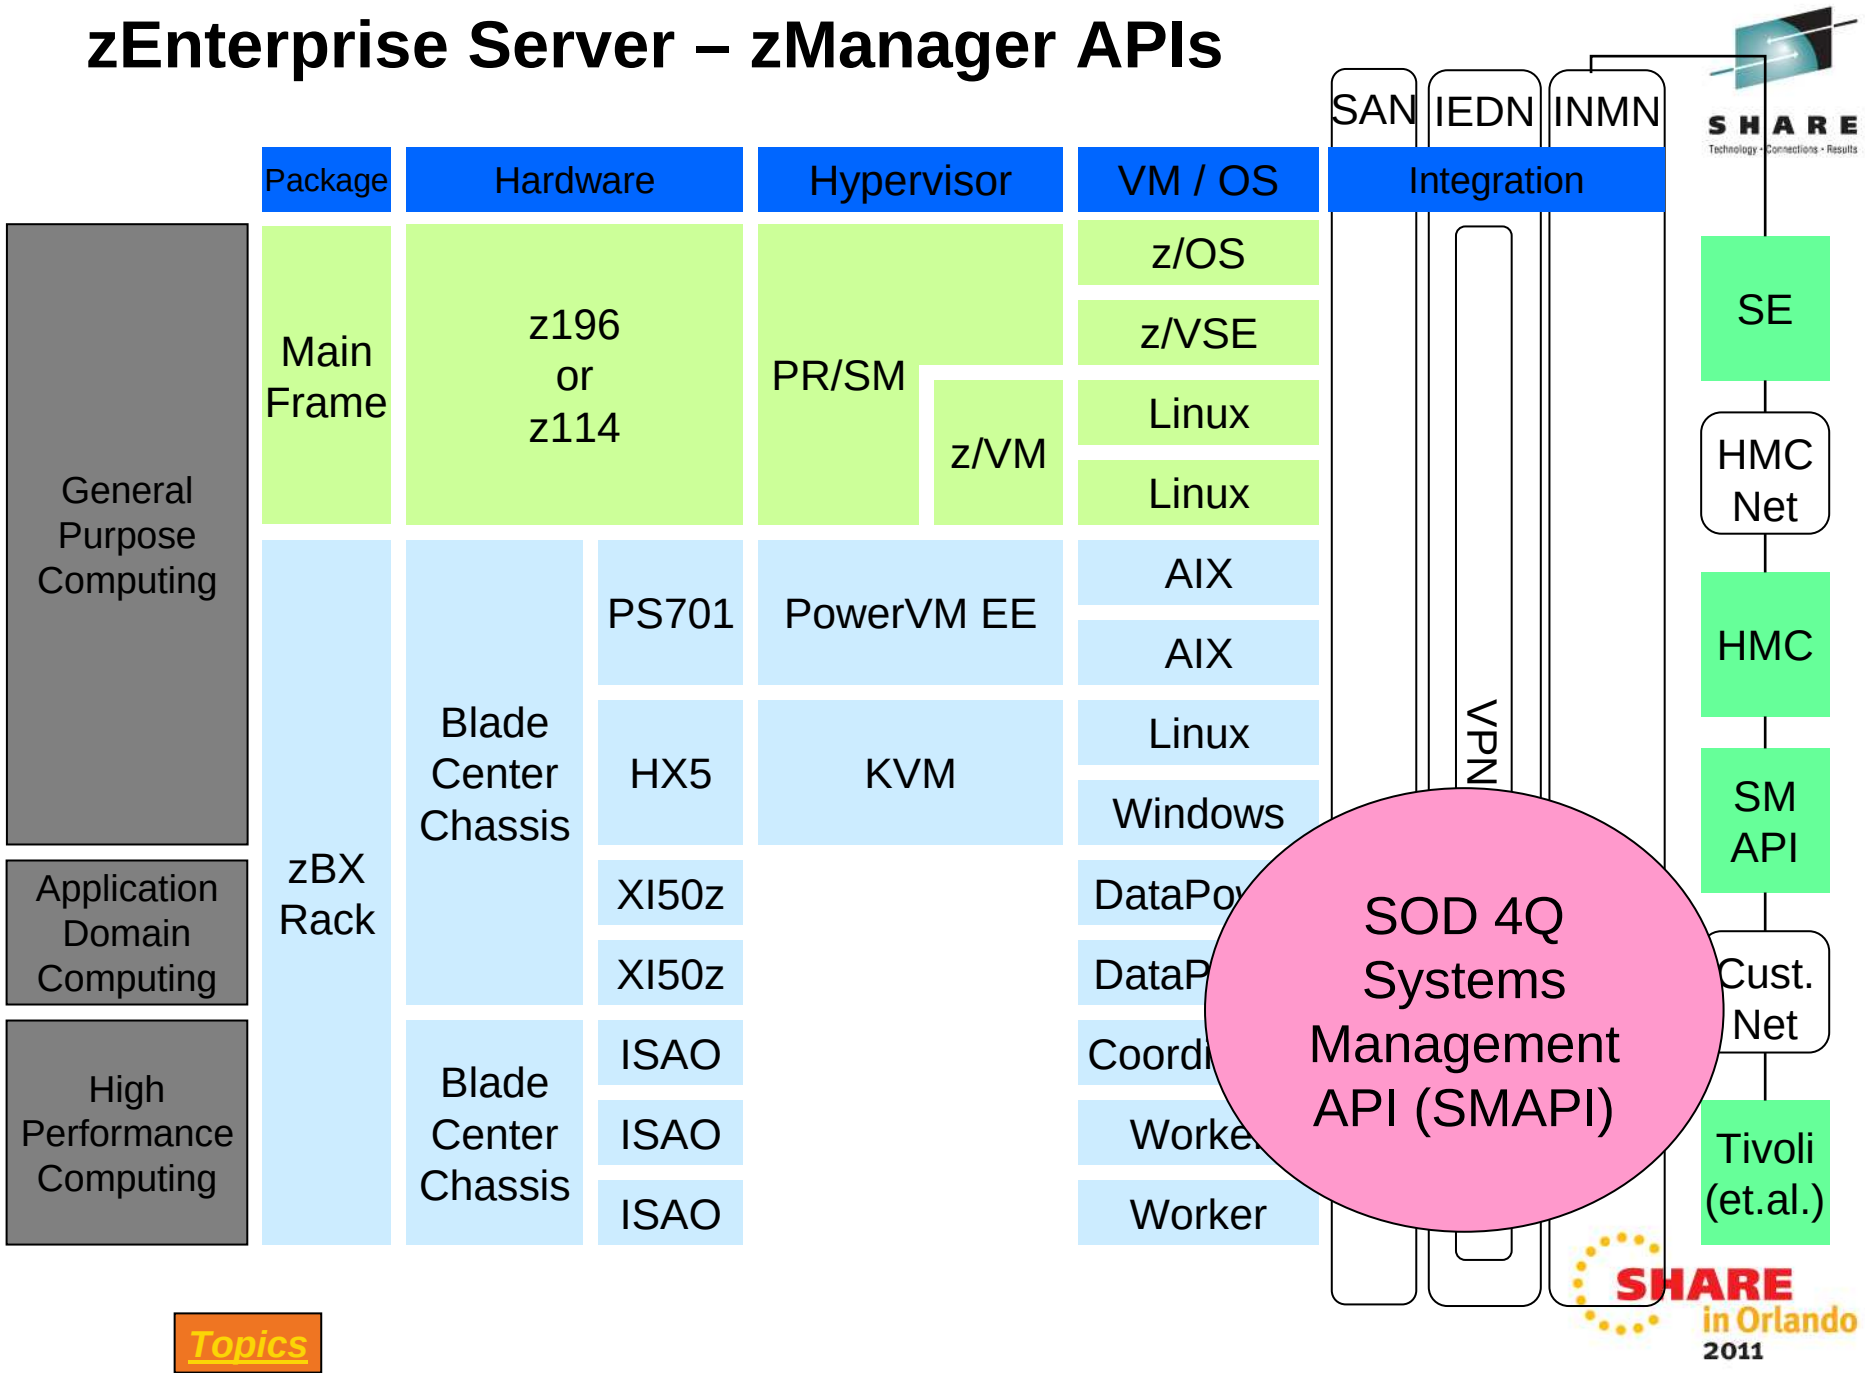

Topics

zEnterprise Server – zManager Functions

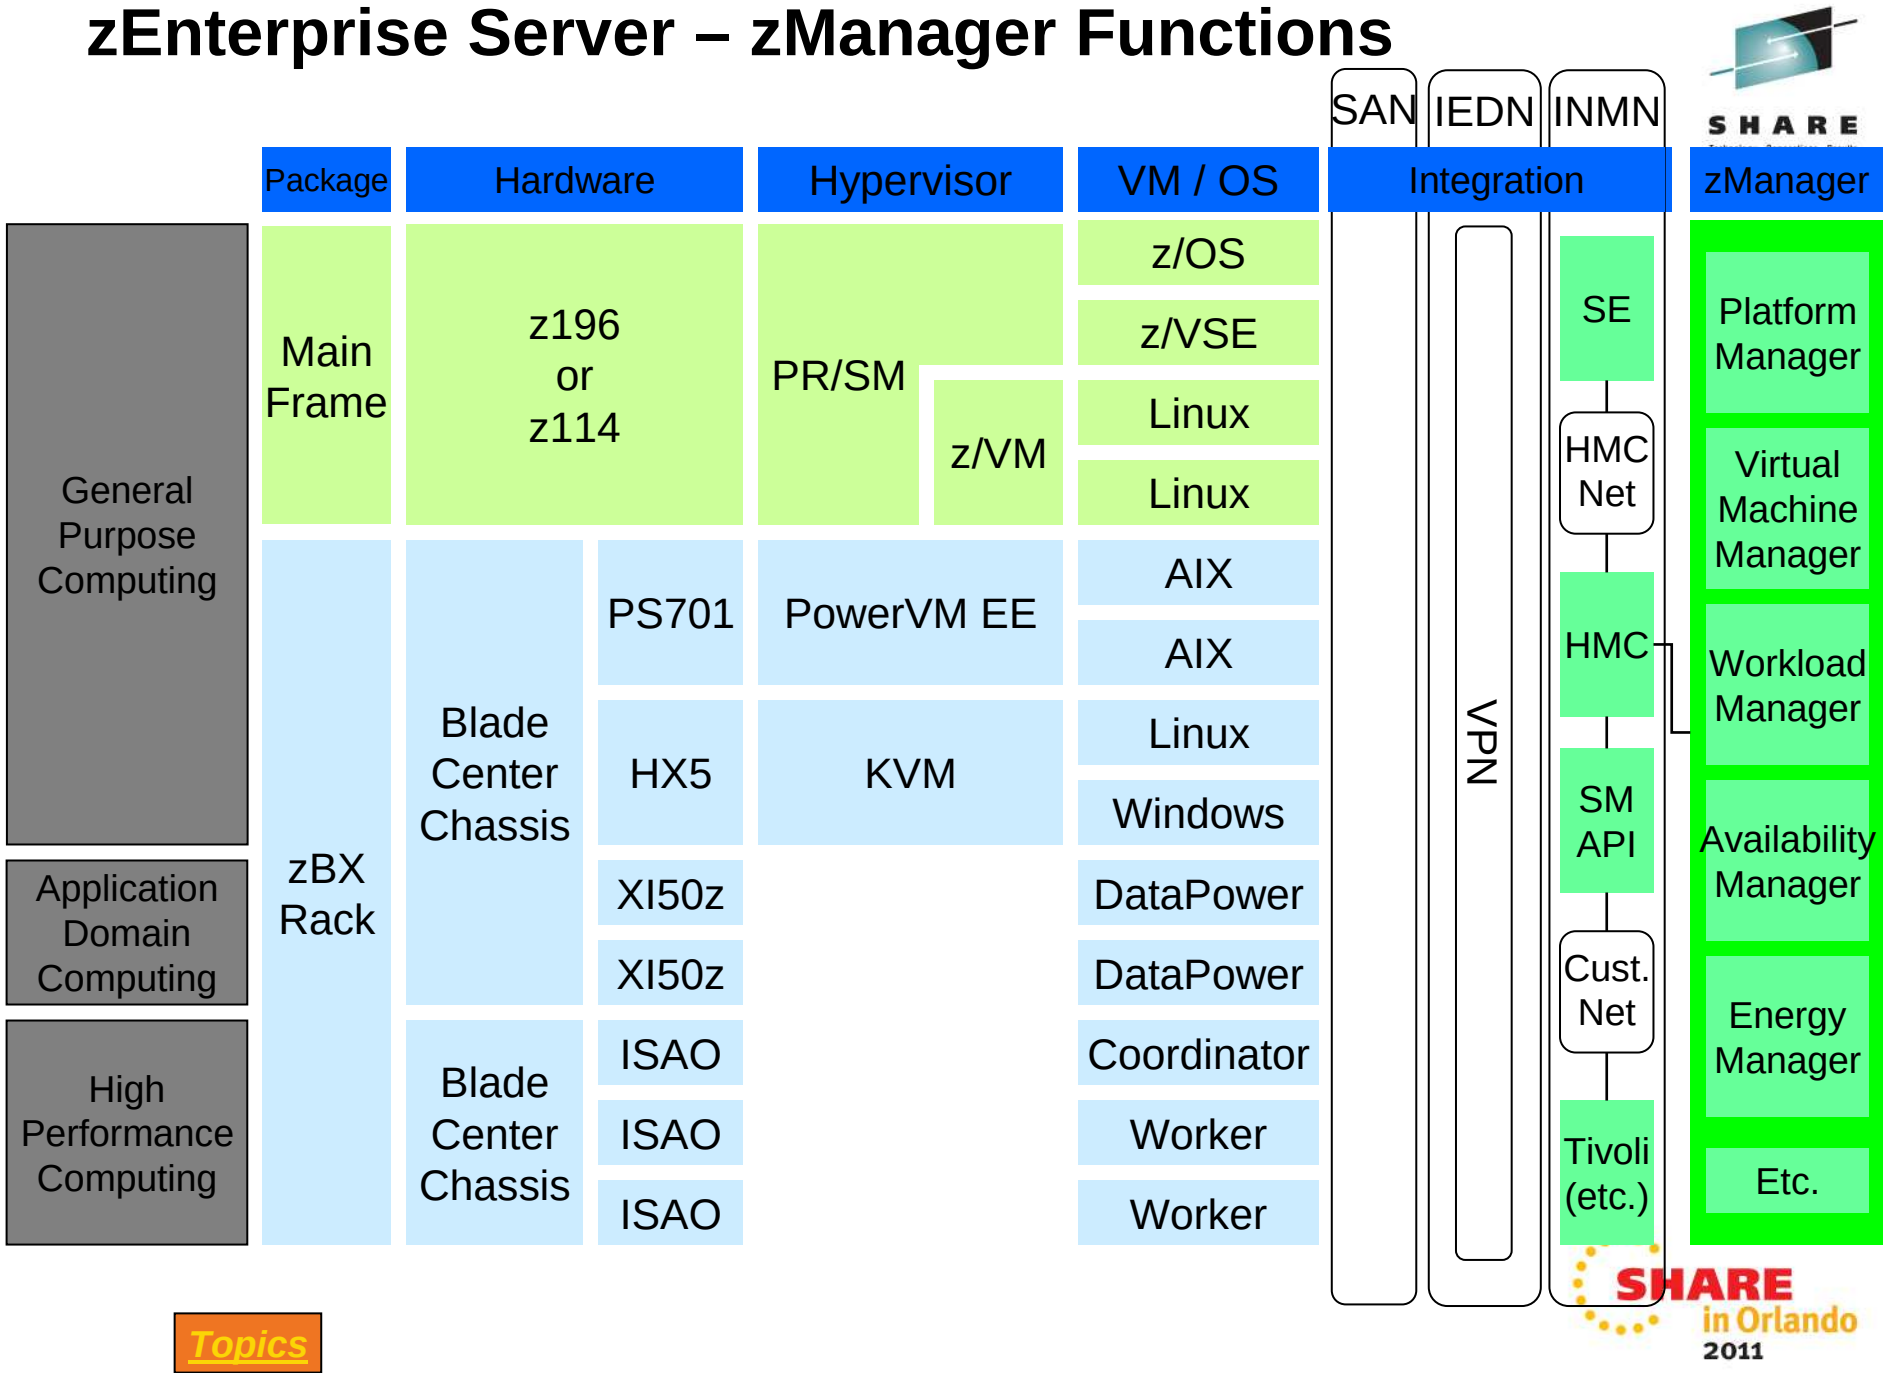

Topics

zEnterprise Node

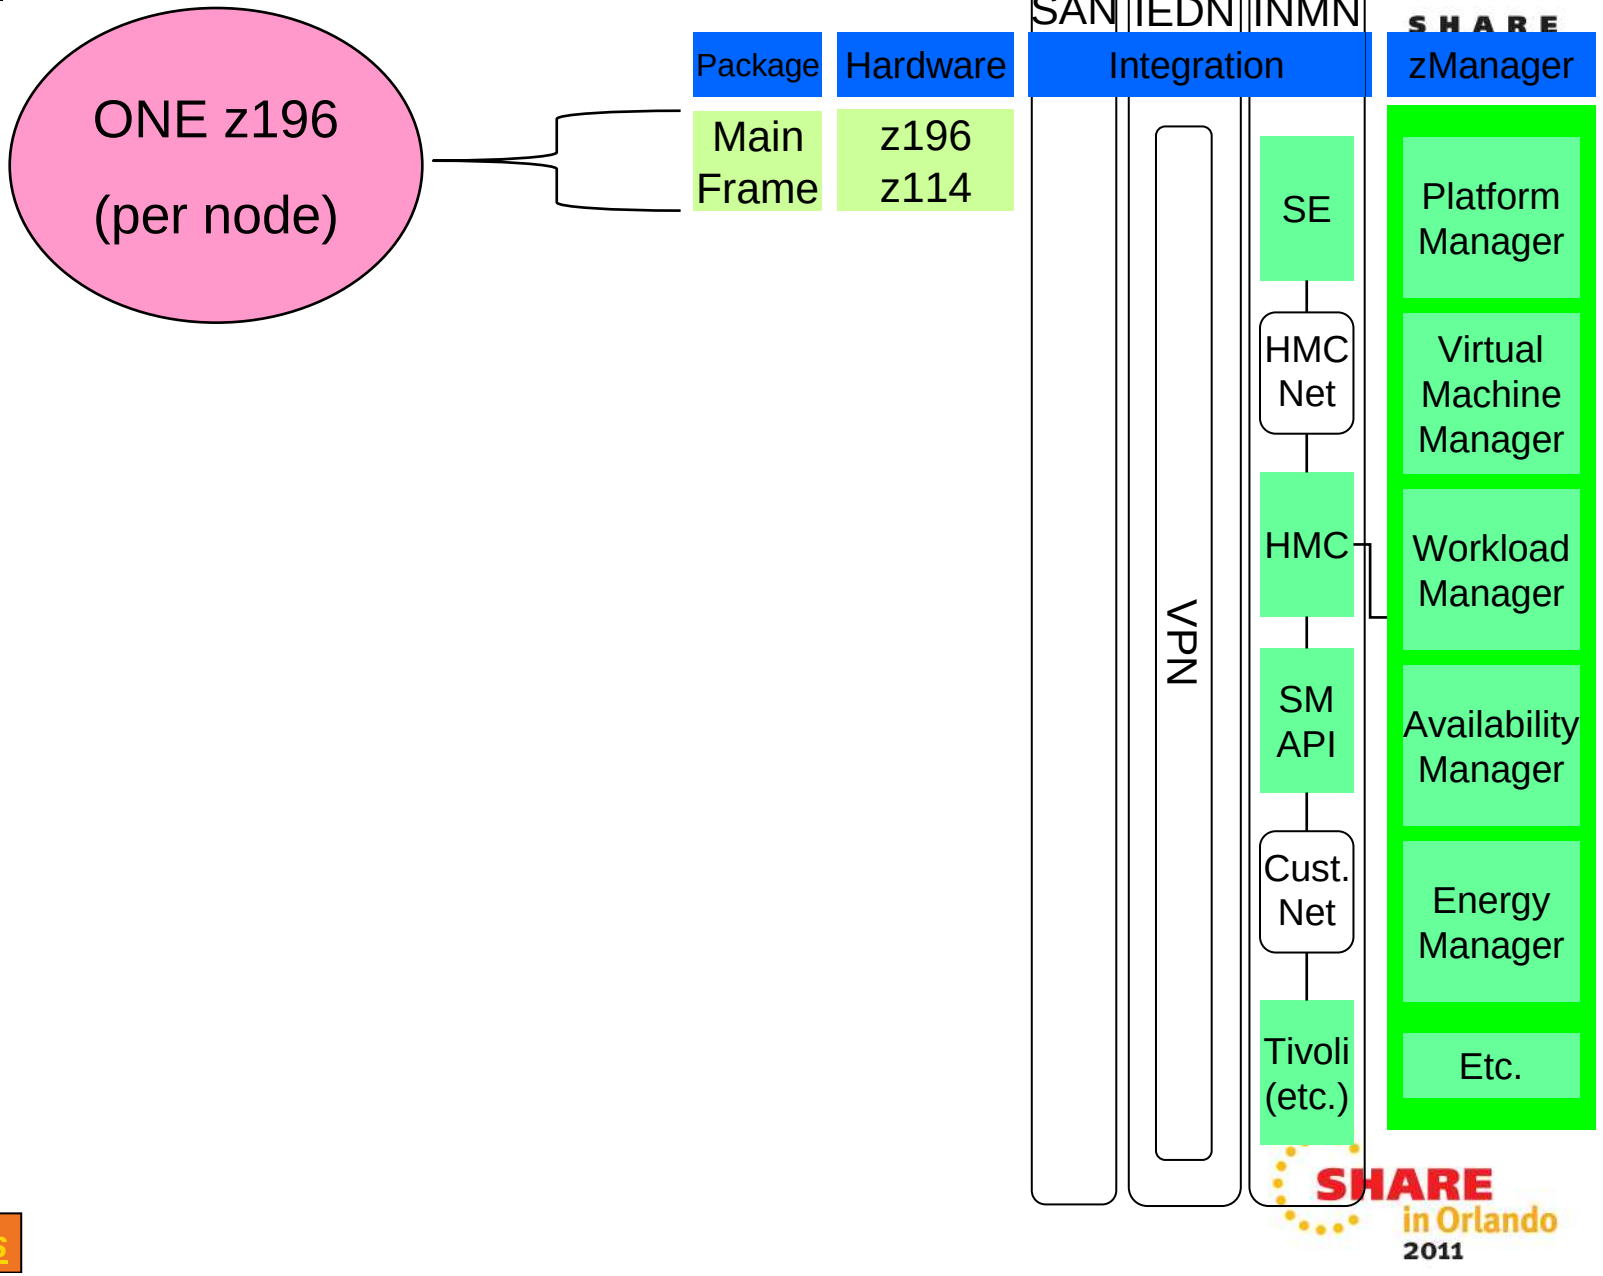

zEnterprise Node

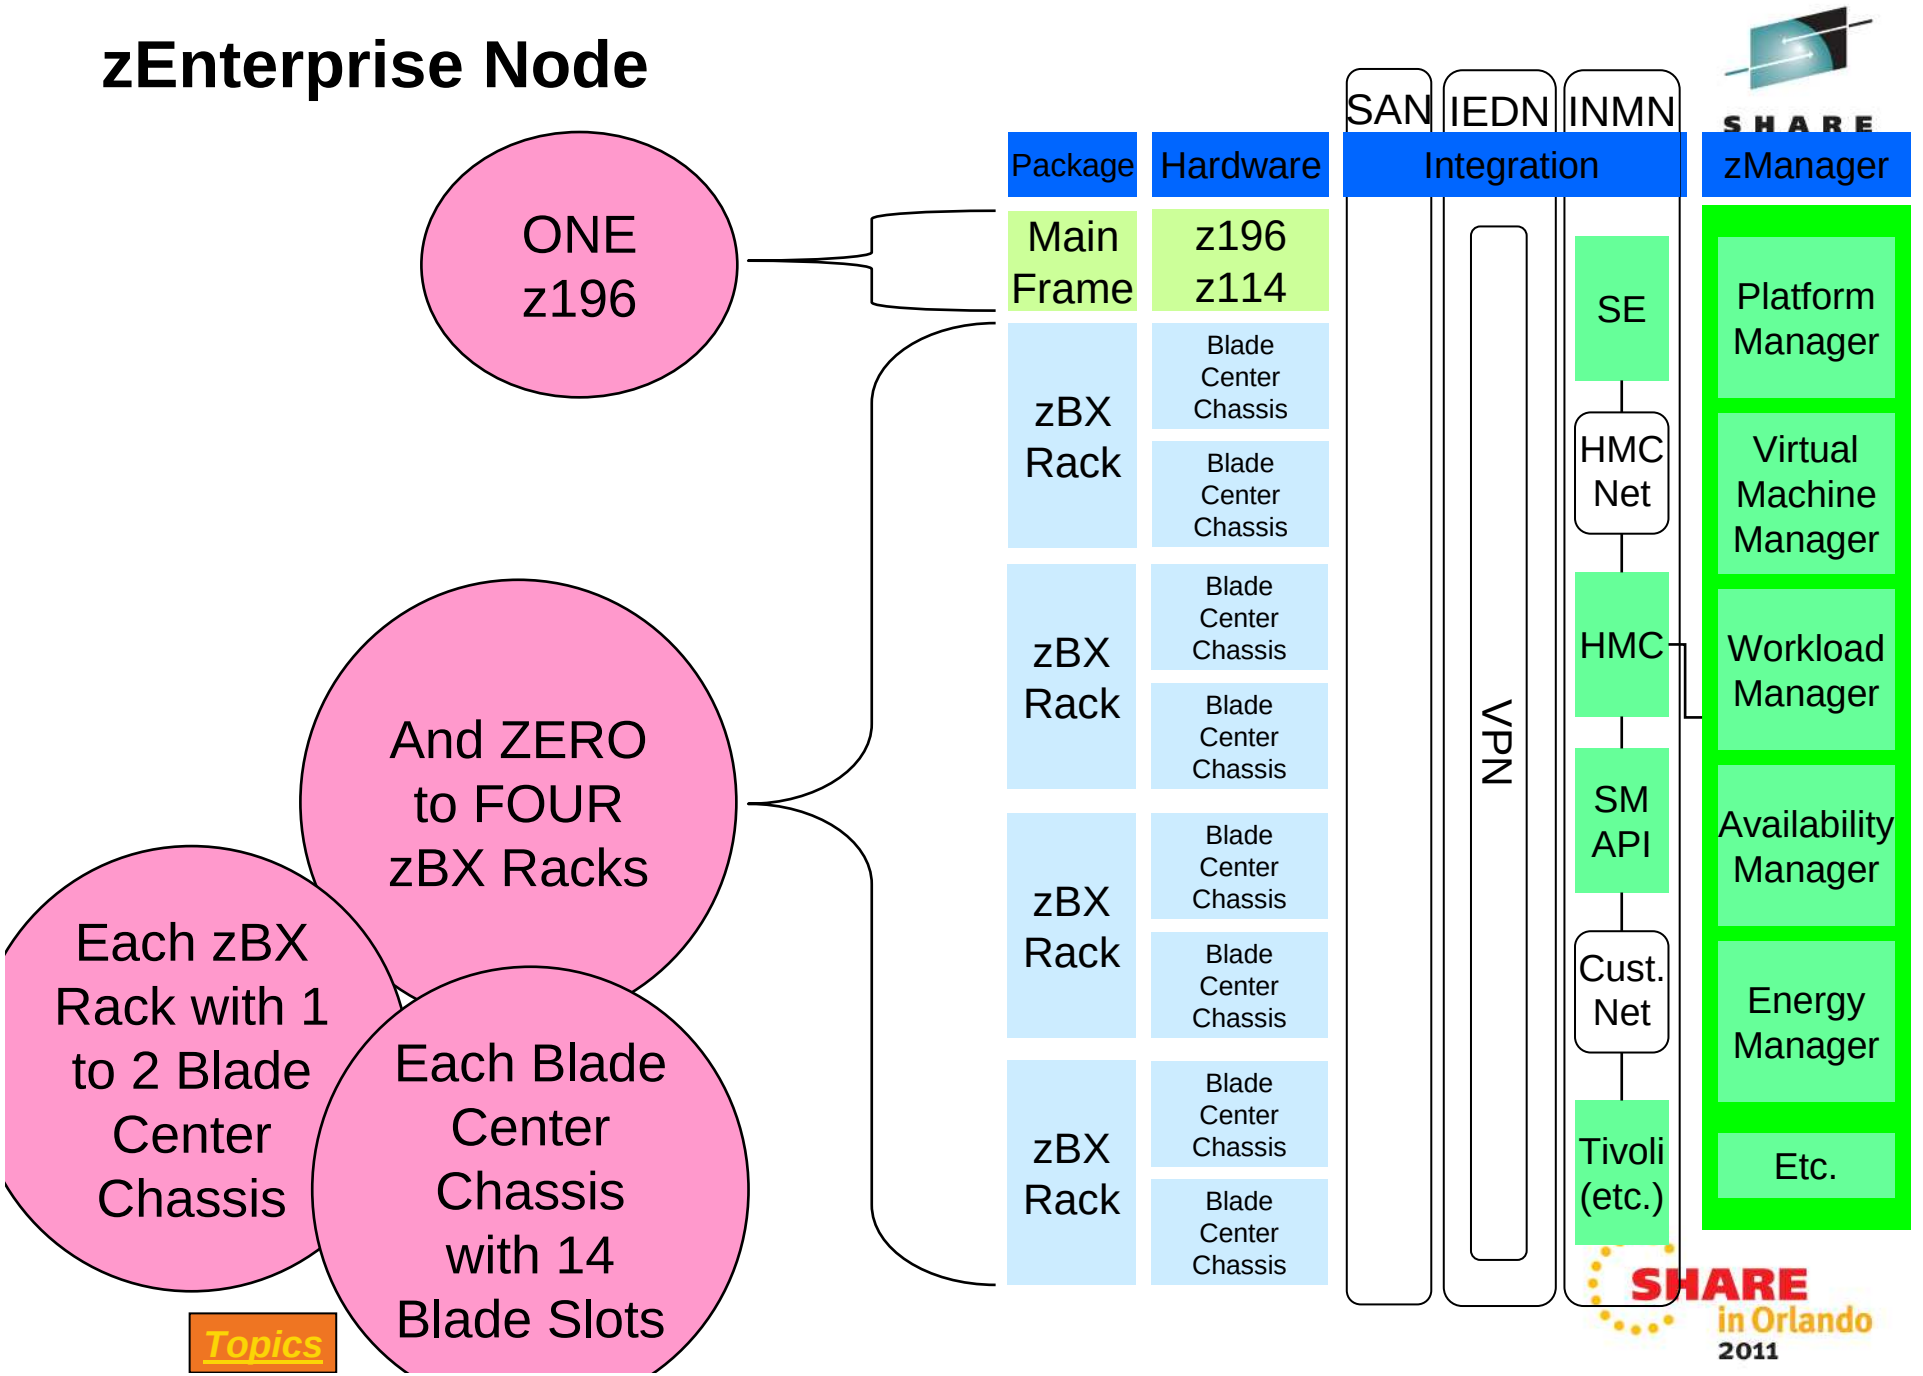

zEnterprise Ensemble

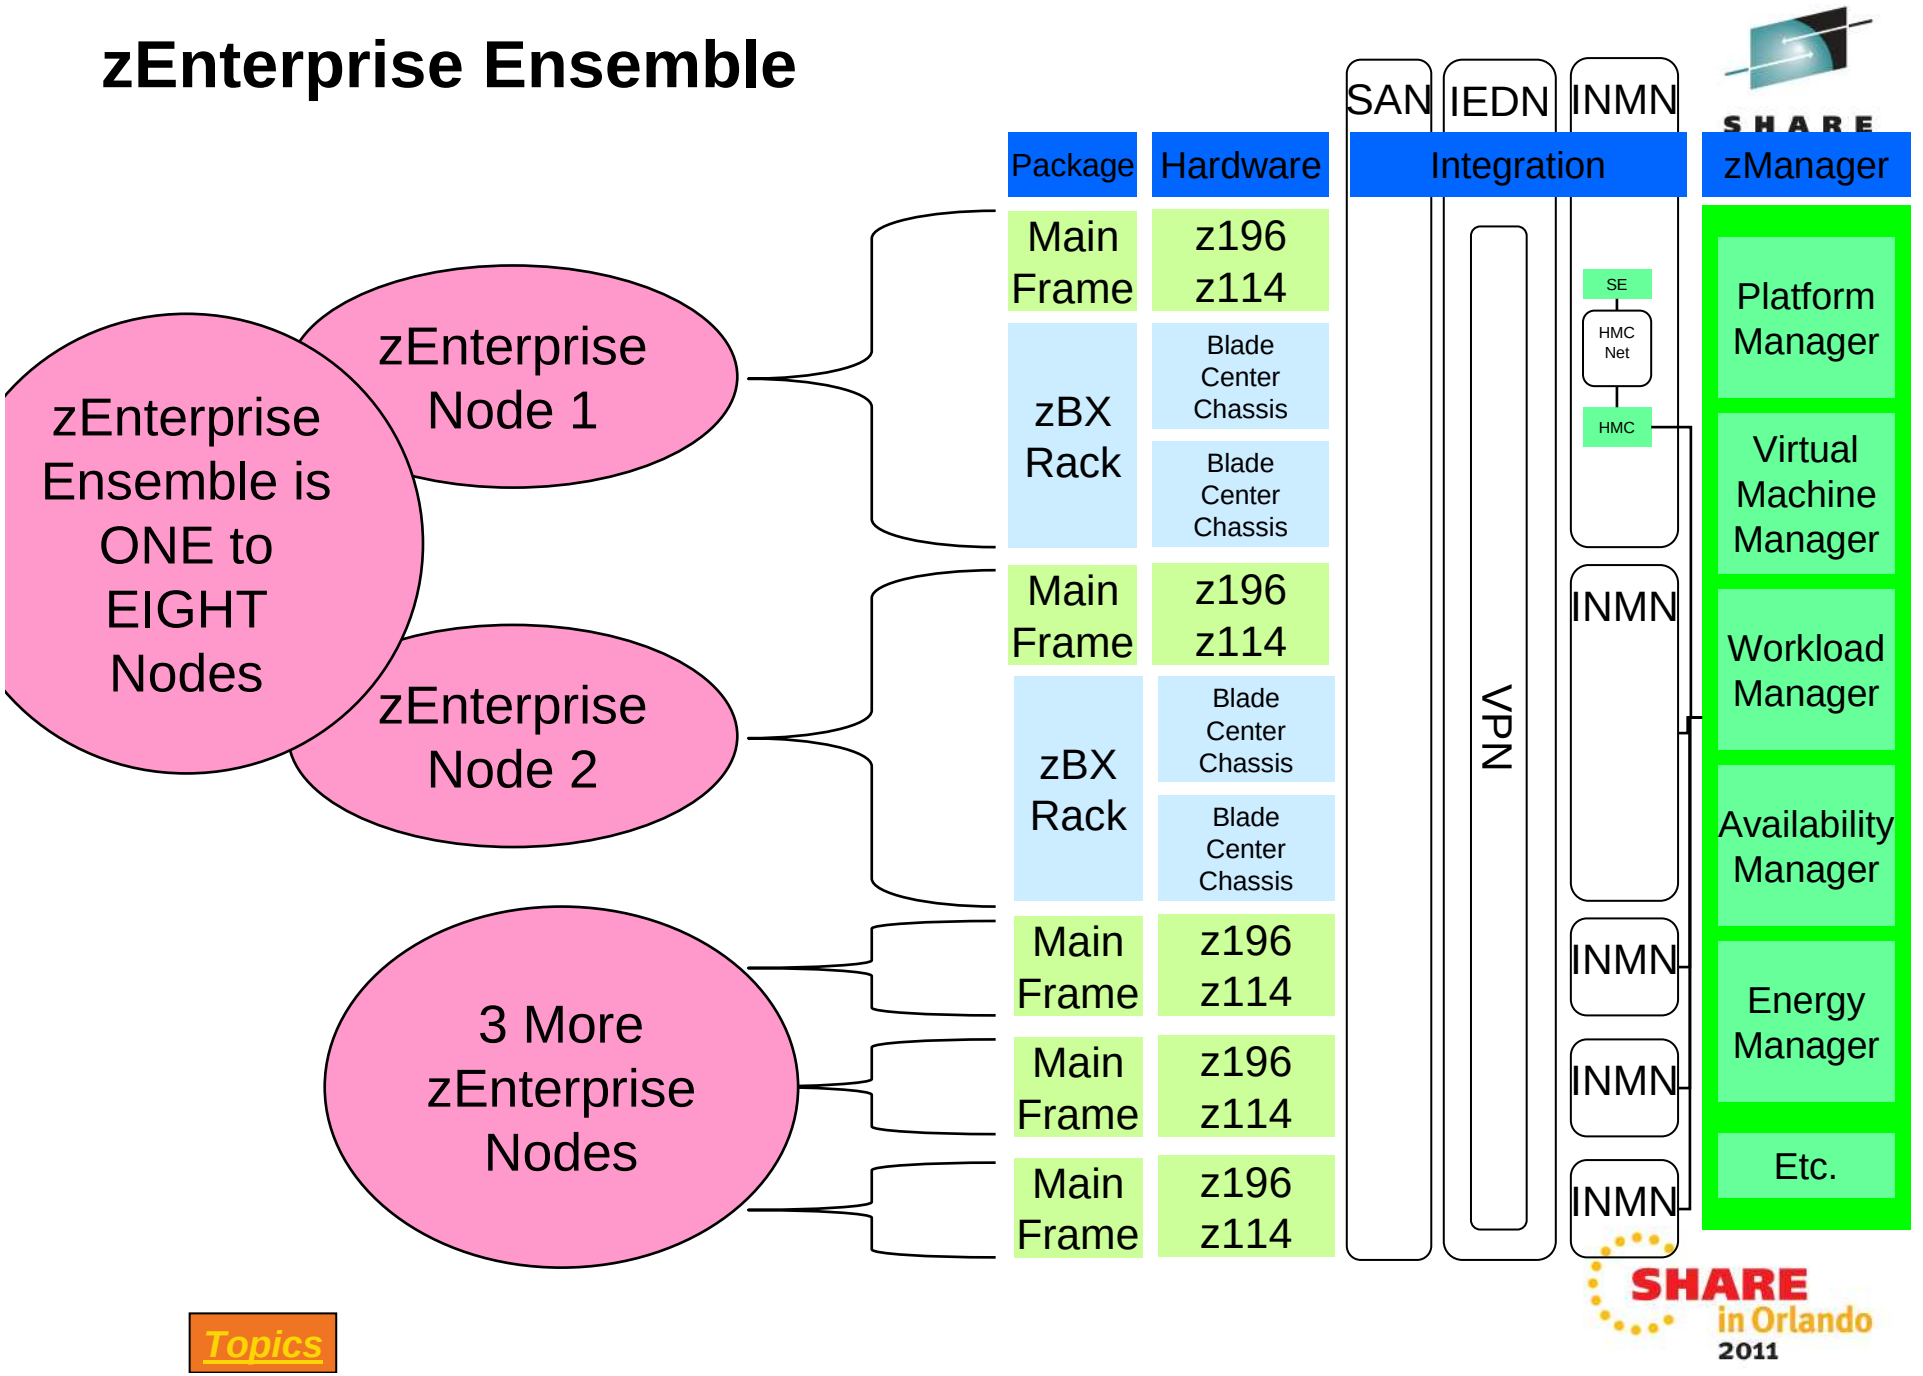

Topics

End of Section

zManager

zManager

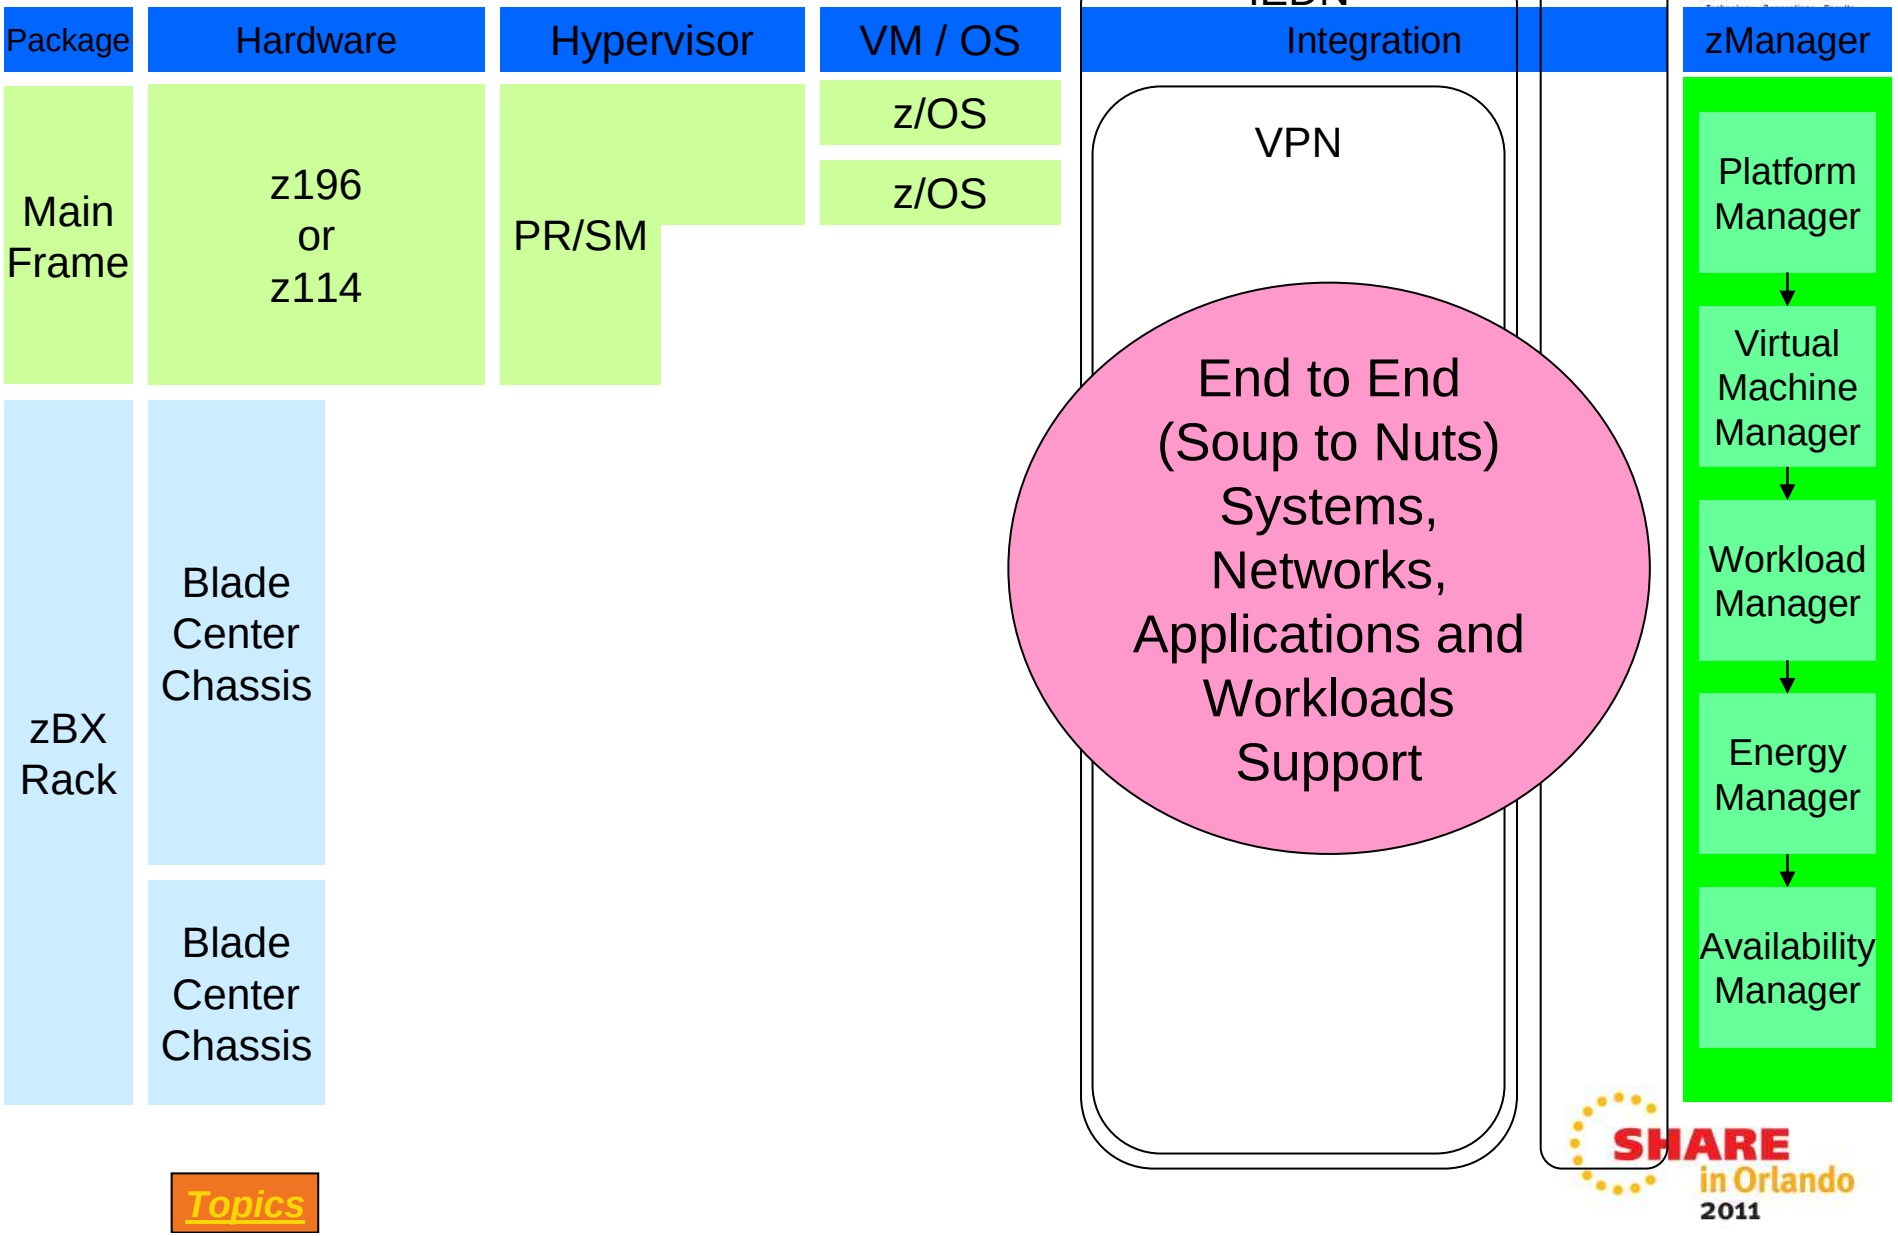

Topics

zManager

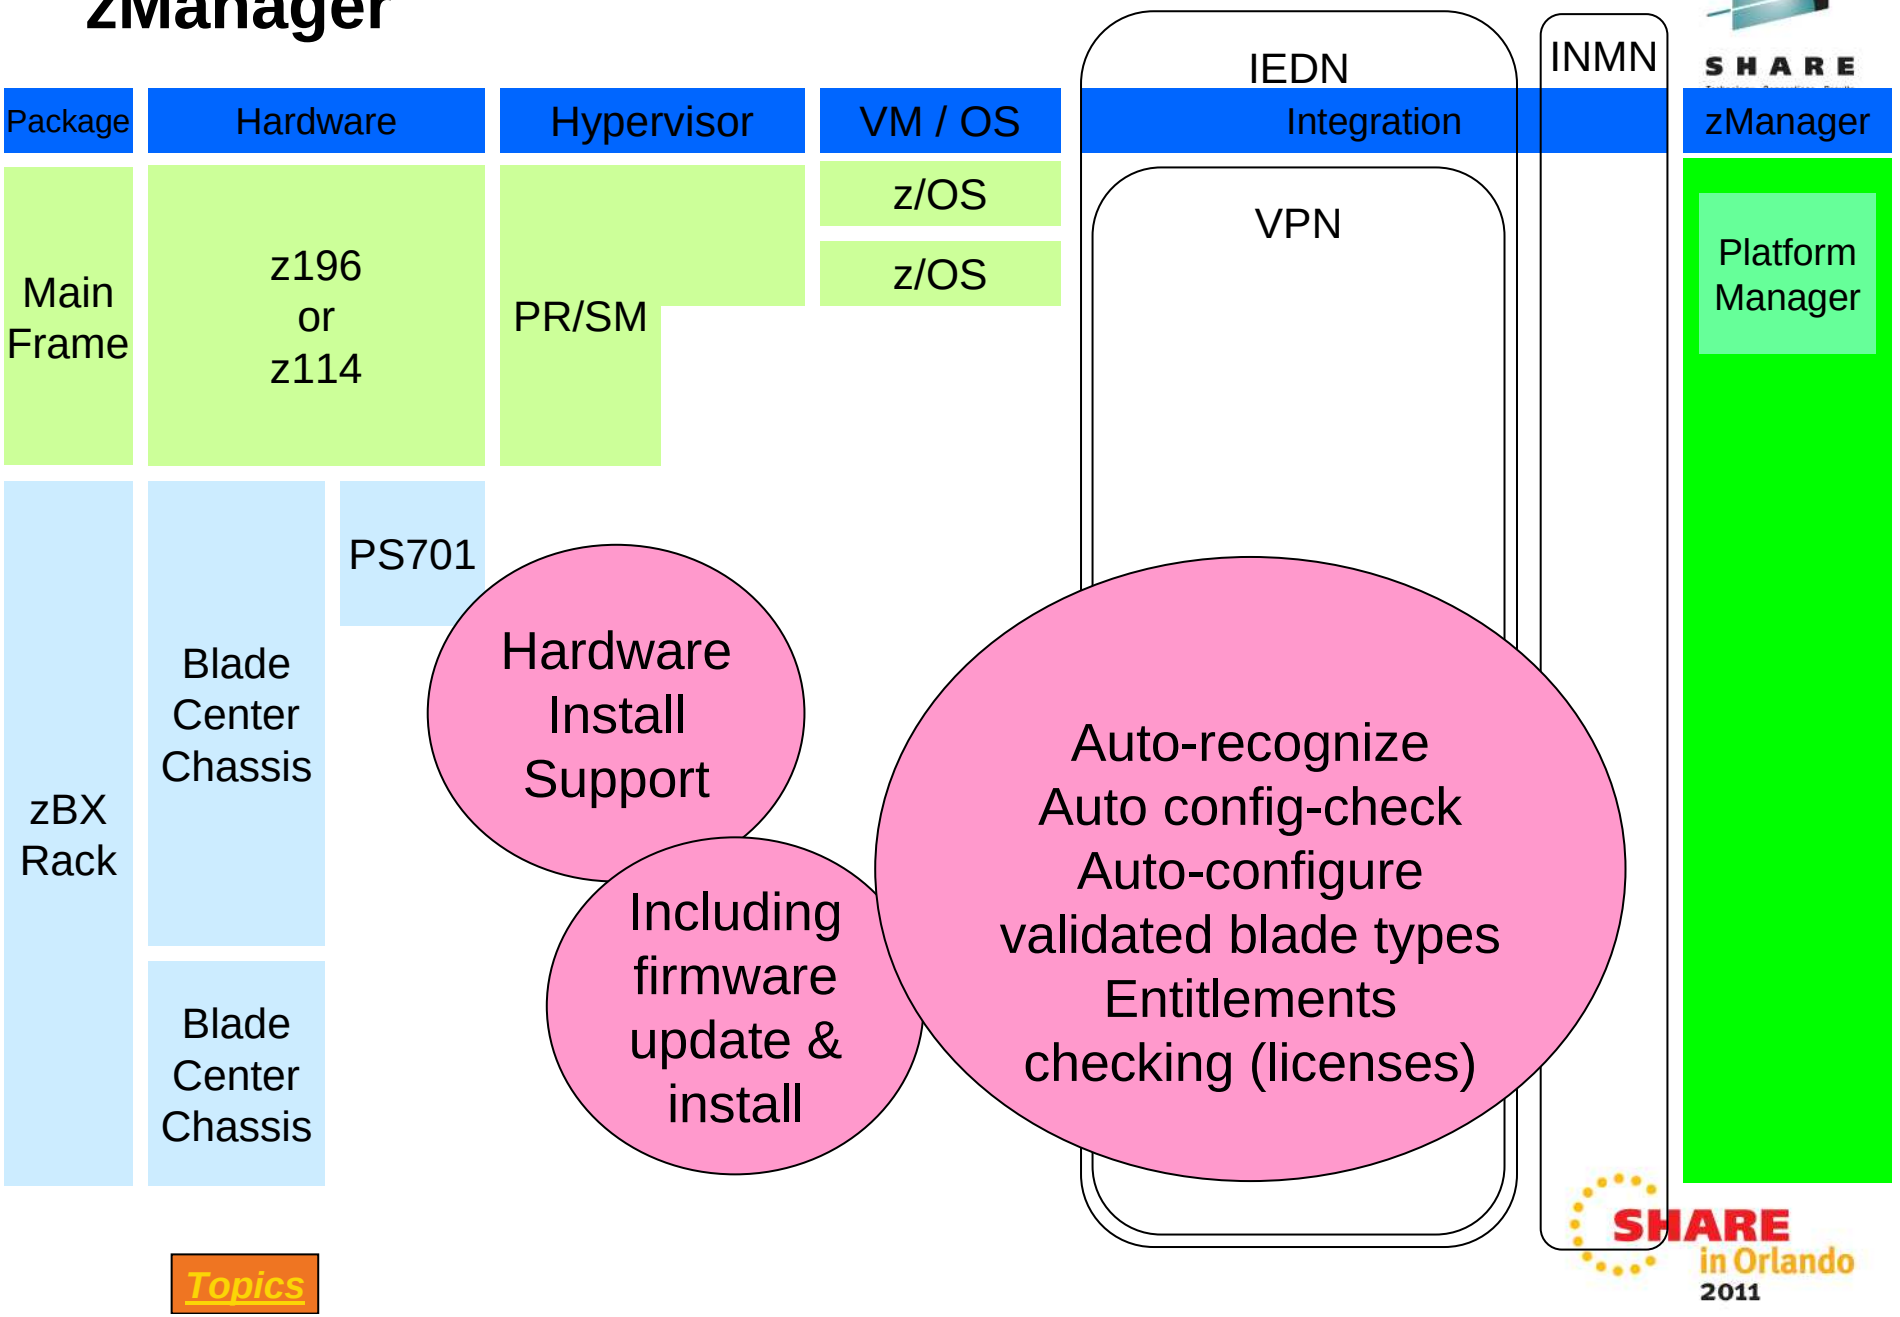

Topics

zManager – Platform Management

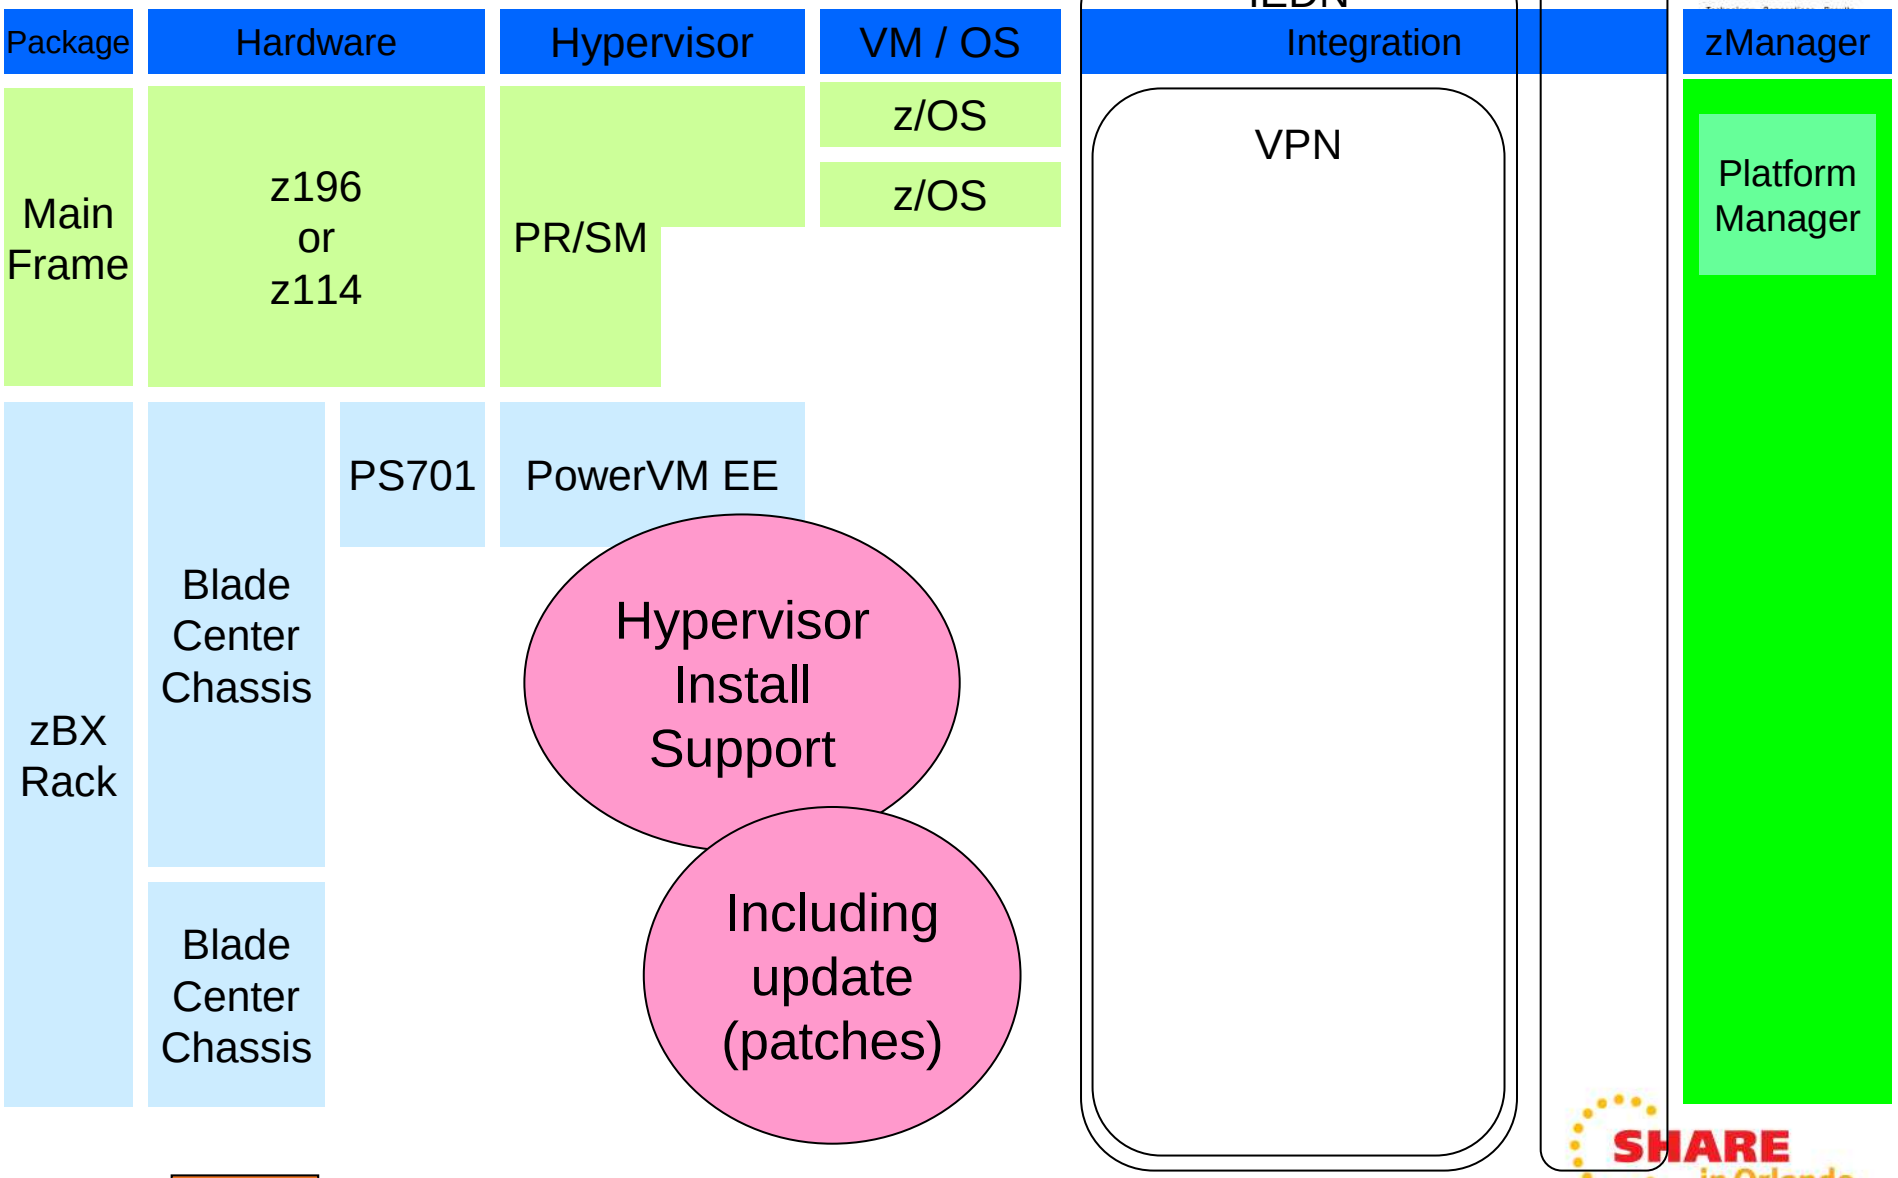

Topics

zManager – Virtual Machine Management

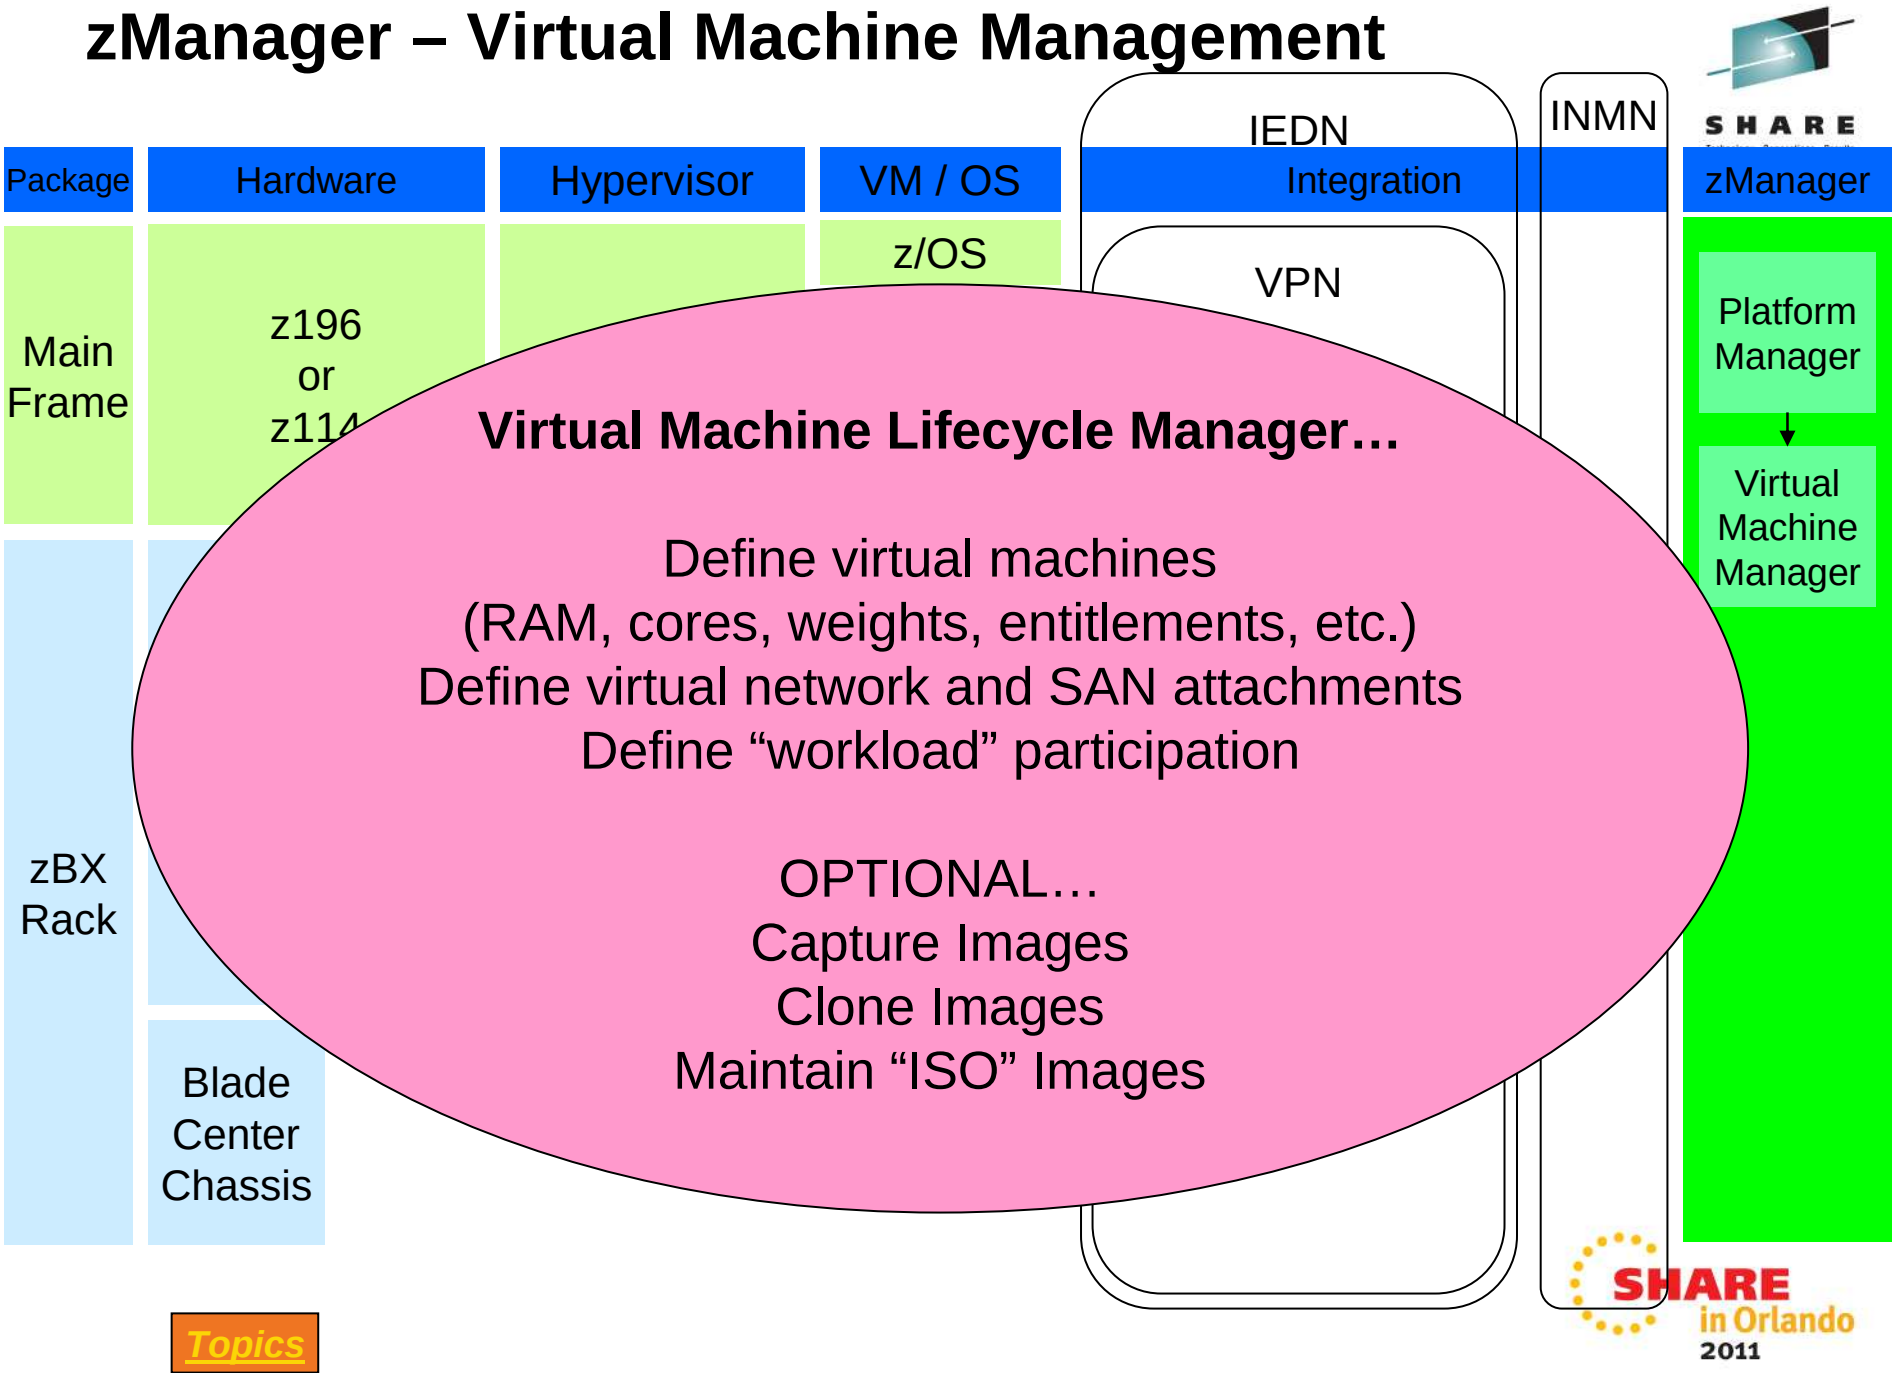

Topics

SHARE in Orlando 2011

zManager – Network Management

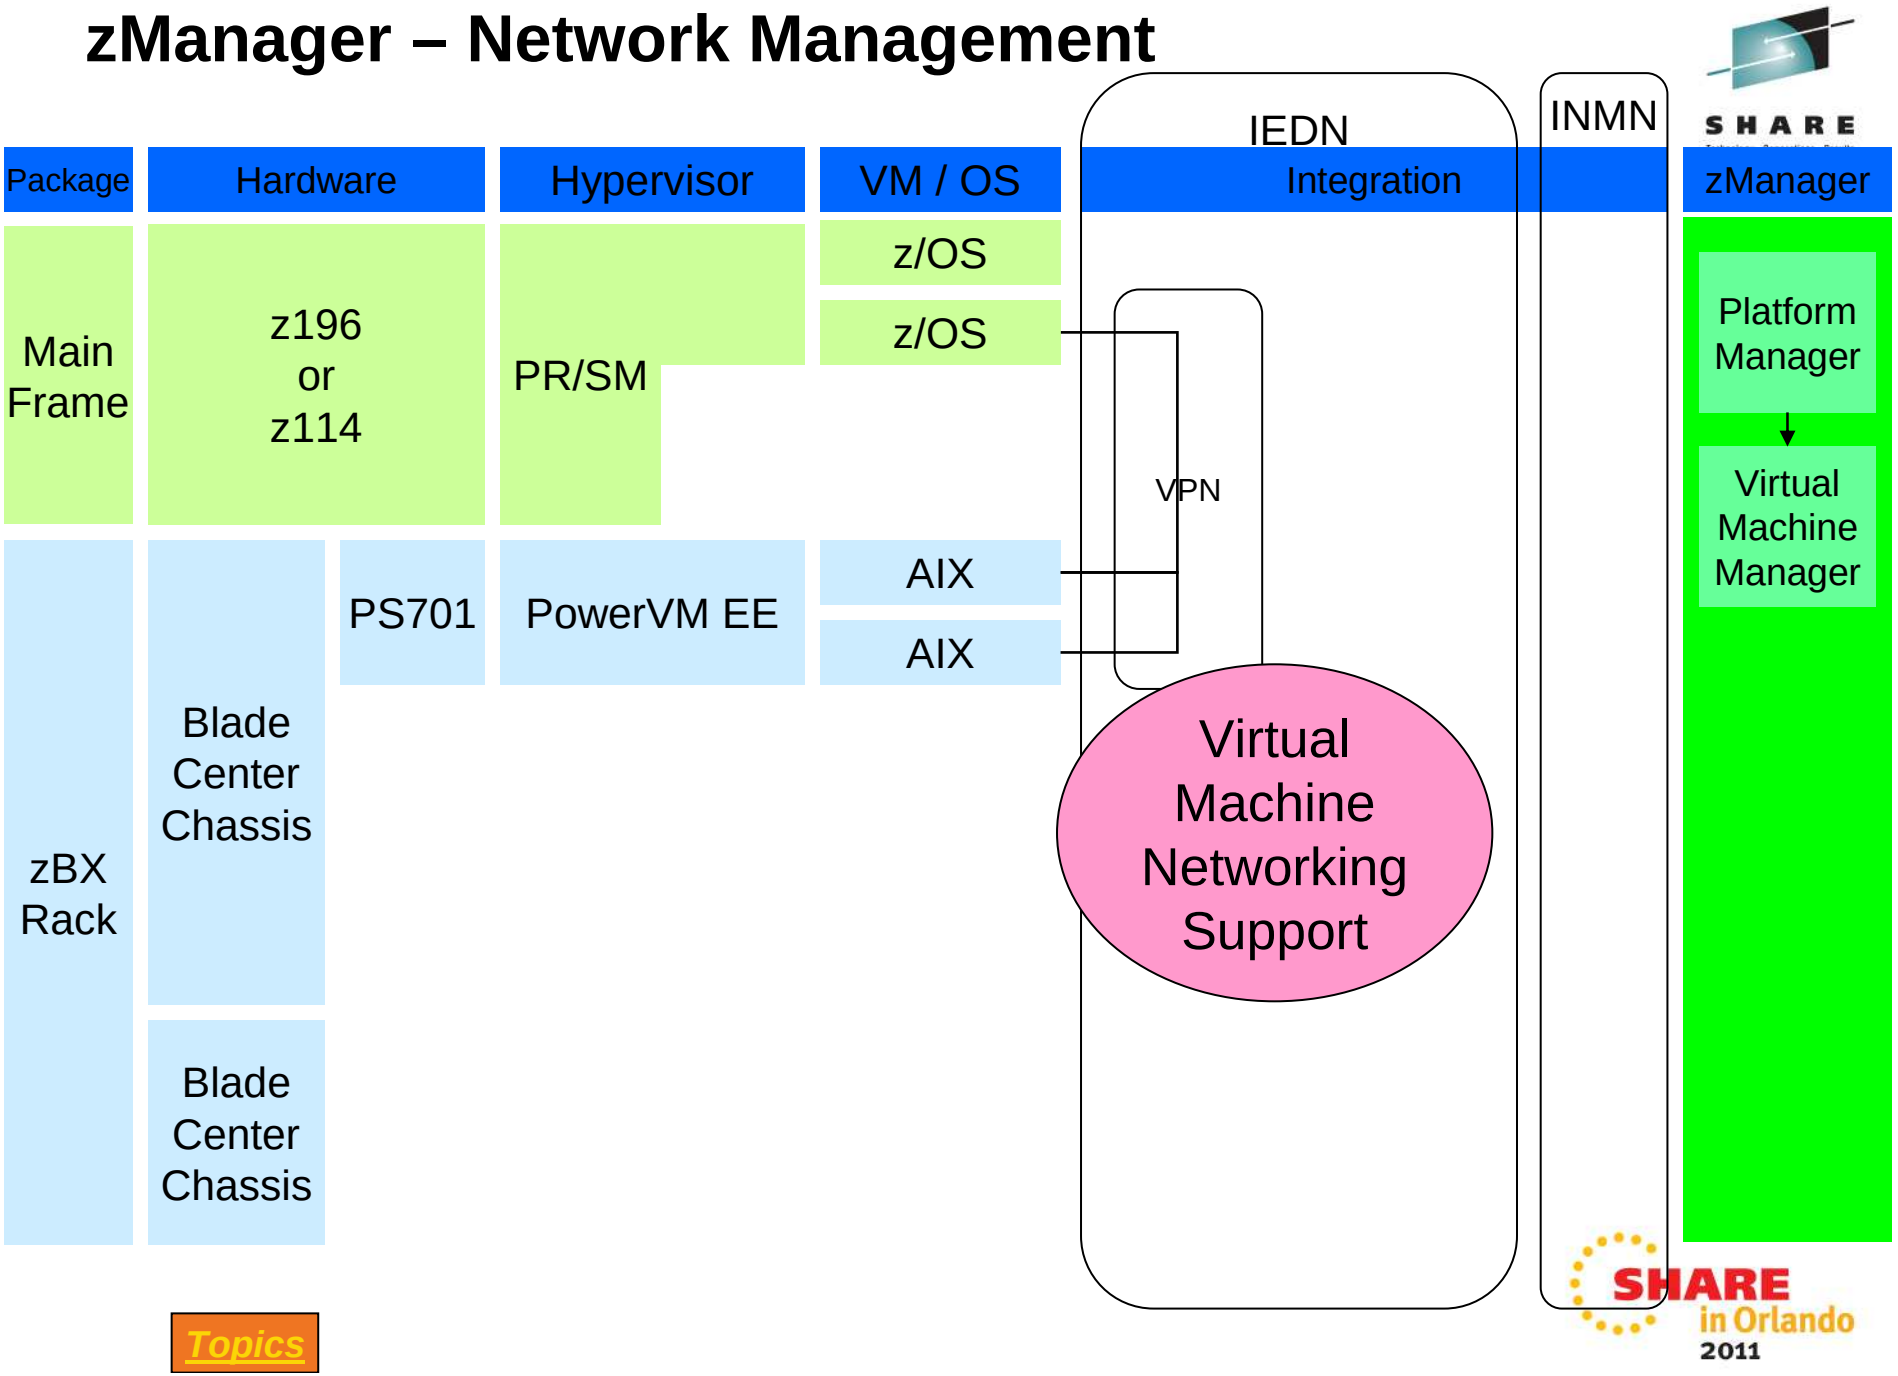

Topics

zManager – Workload Definitions

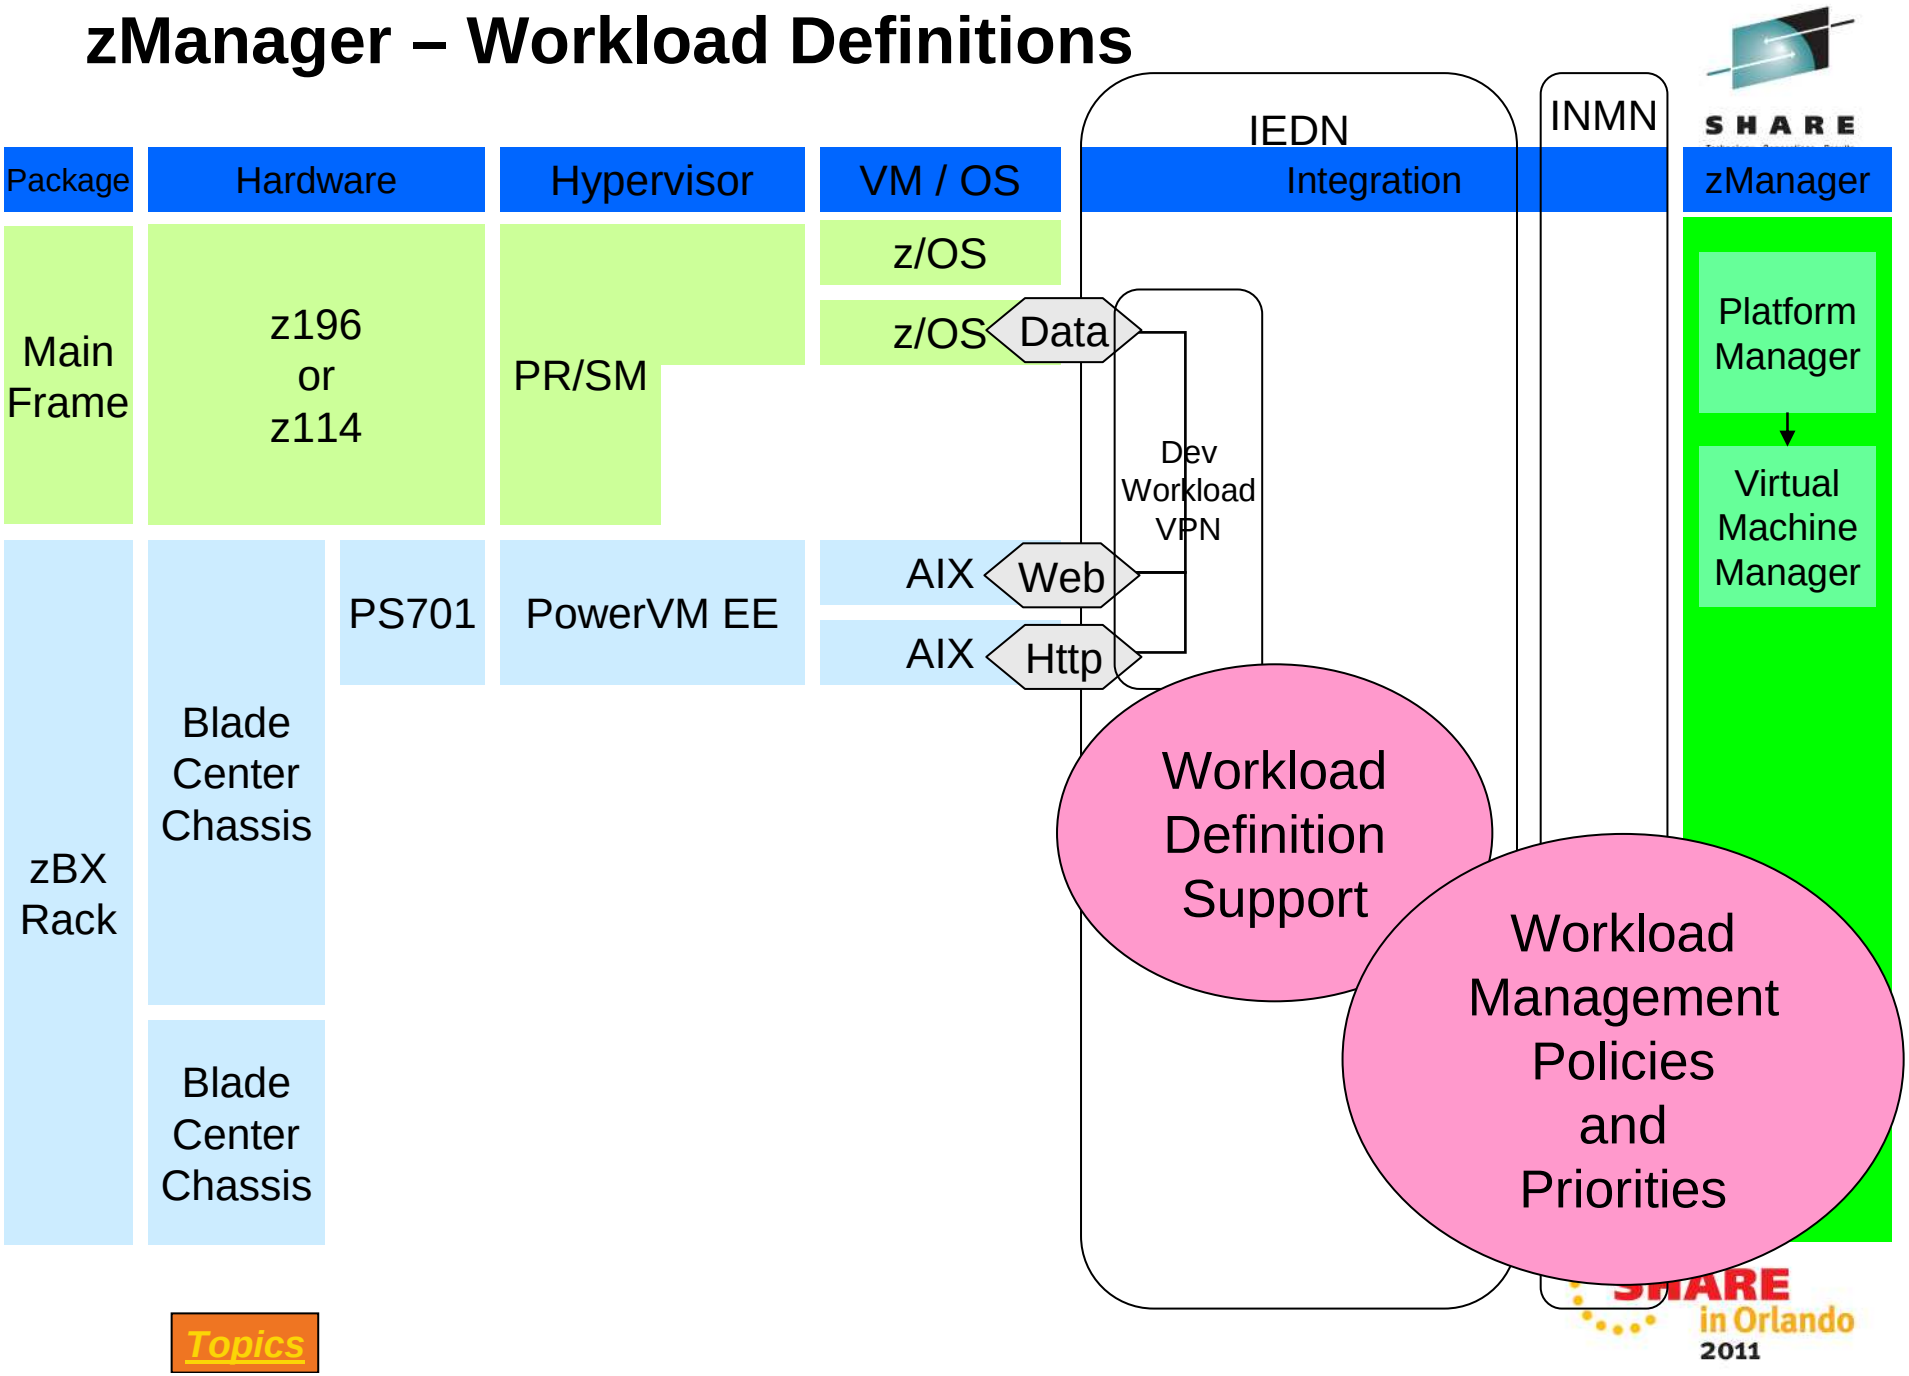

Topics

zManager – Workload Management

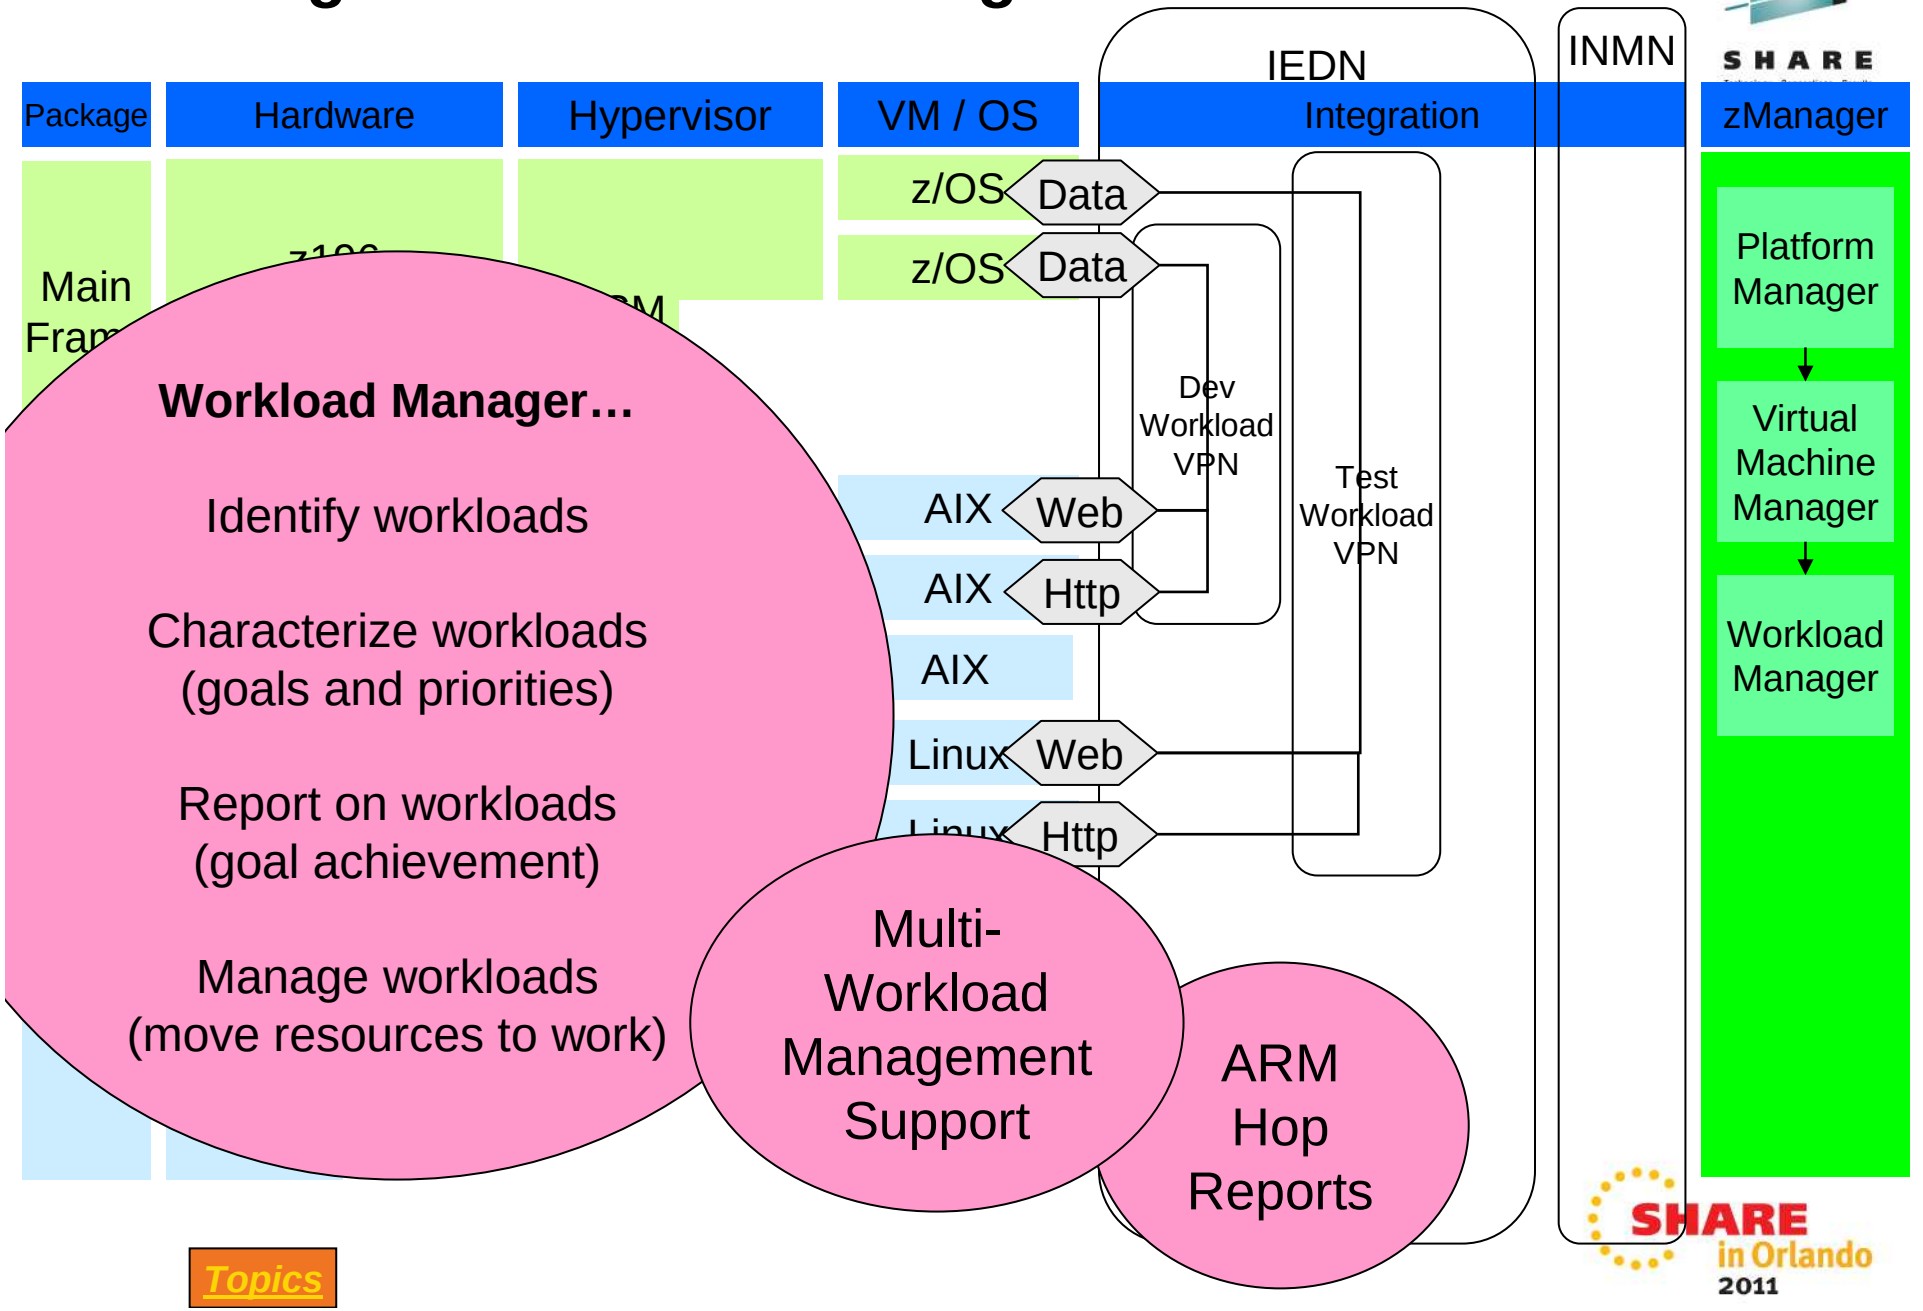

Topics

zManager – Energy Manager

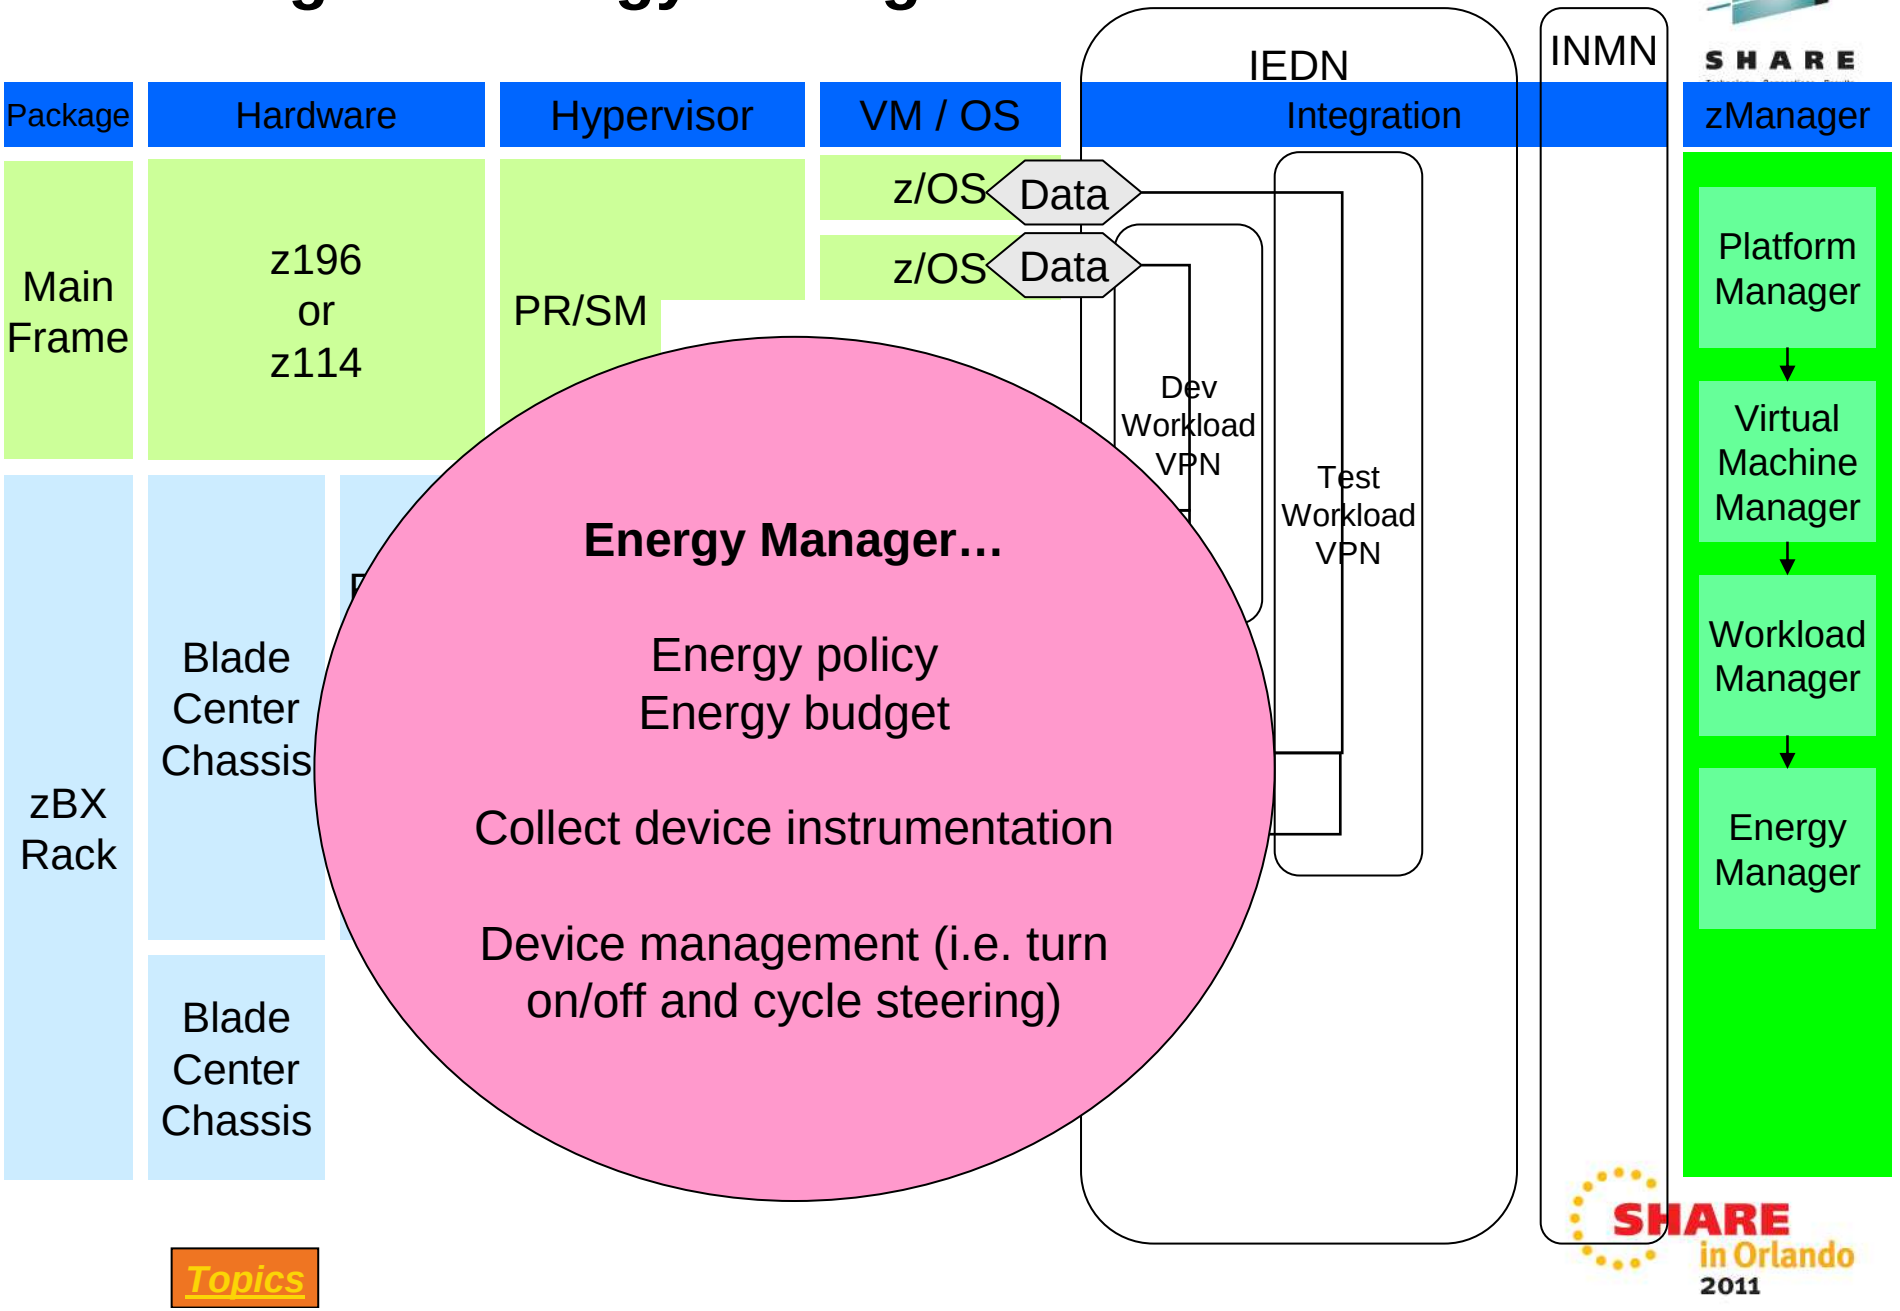

Topics

zManager – Availability Manager

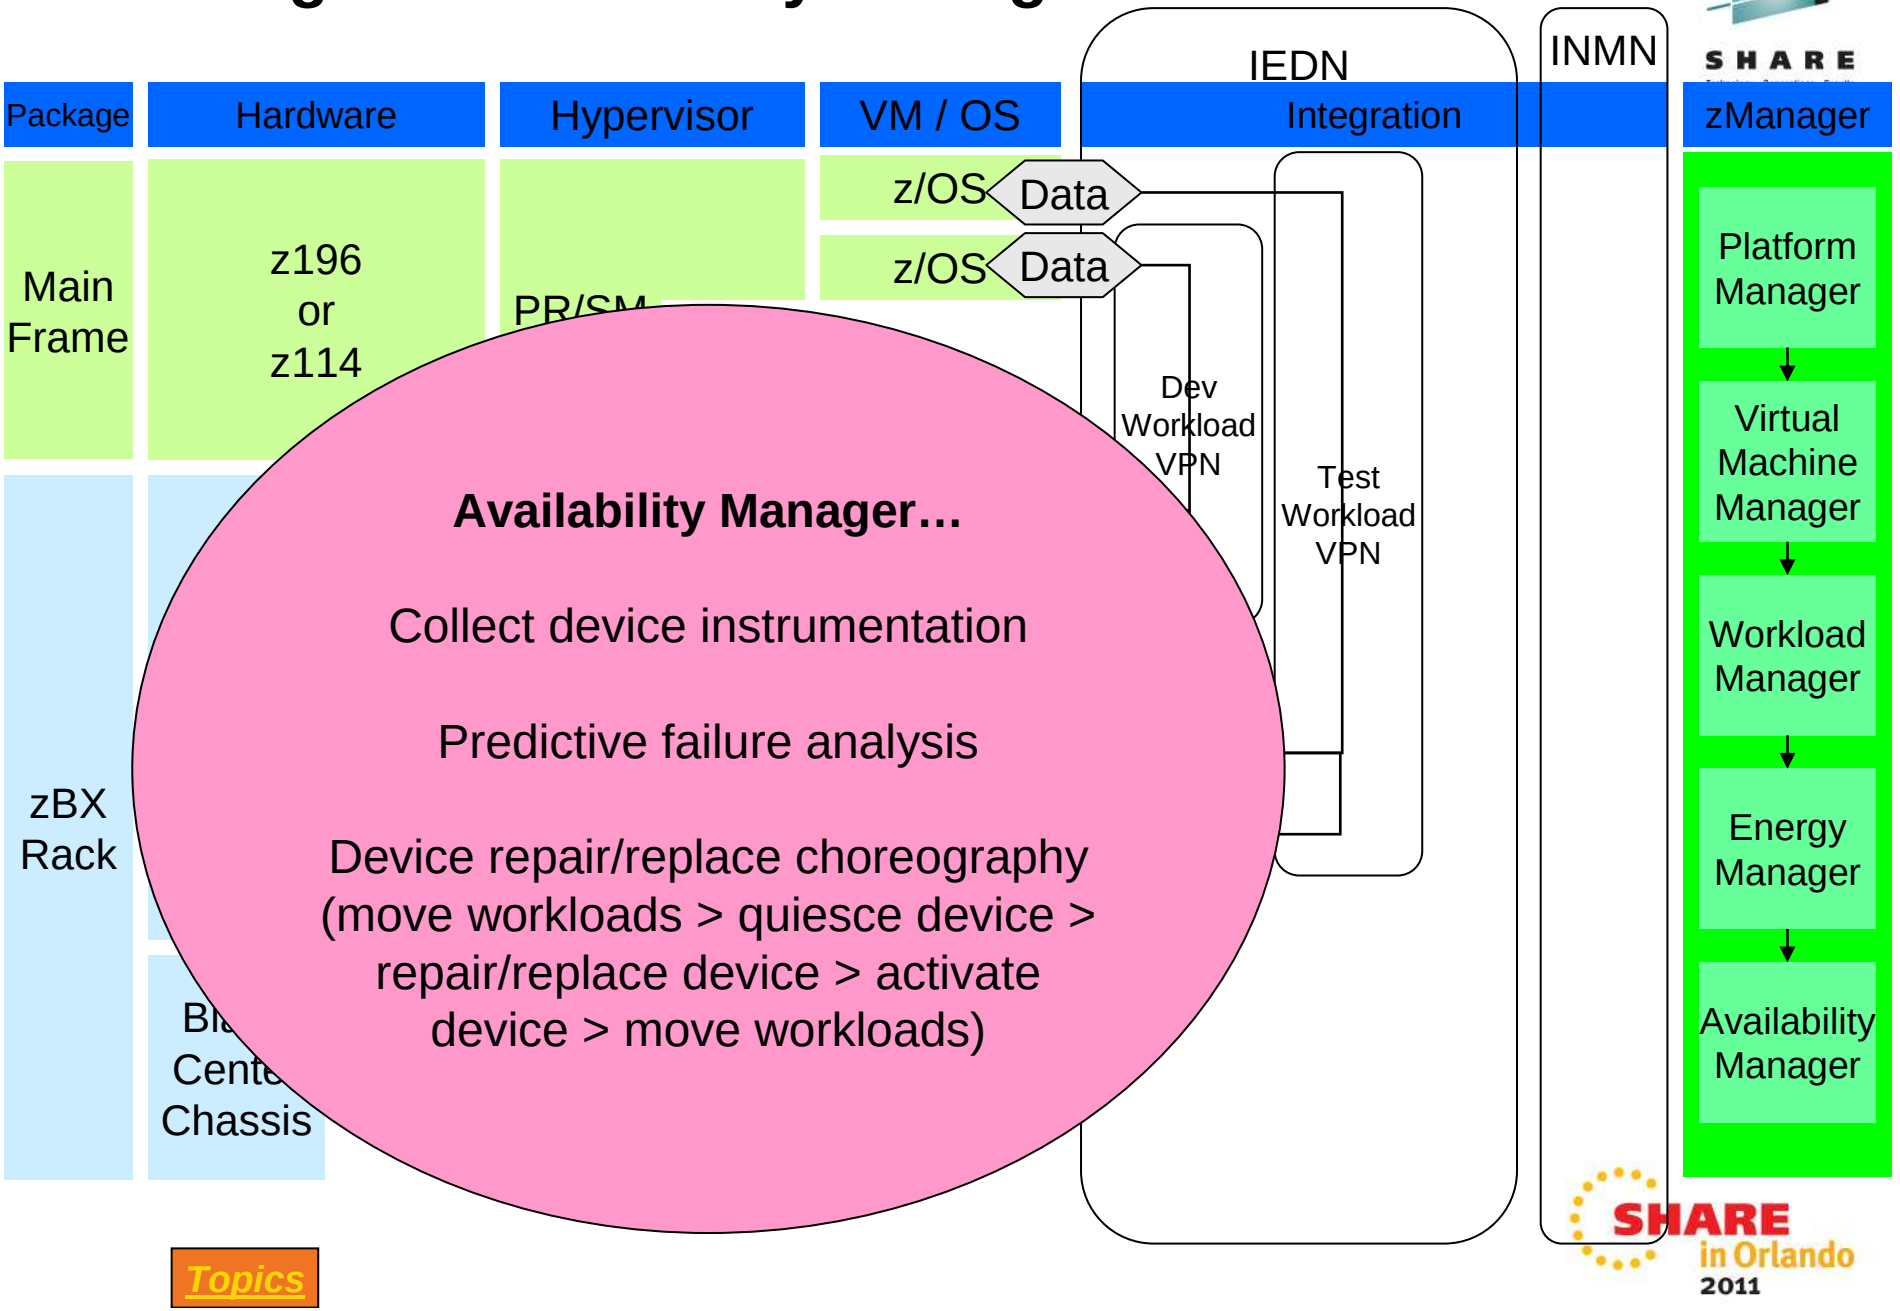

Topics

End of Section

Use Cases

Use Cases

1. Multi-Tier Application Serving with zEnterprise Linux on System z
2. Multi-Tier Application Serving with zEnterprise POWER/AIX and z/OS
3. Real-Time Analytics with zEnterprise and IBM Smart Analytics Optimizer (ISAO)
4. Hardware ESB with zEnterprise and DataPower XI50z
5. Storm Cloud
6. Parallel Batch
7. Combinatorial

Topics

Use Case 1

Use Case: Multi-Tier Application zEnterprise – Linux on z

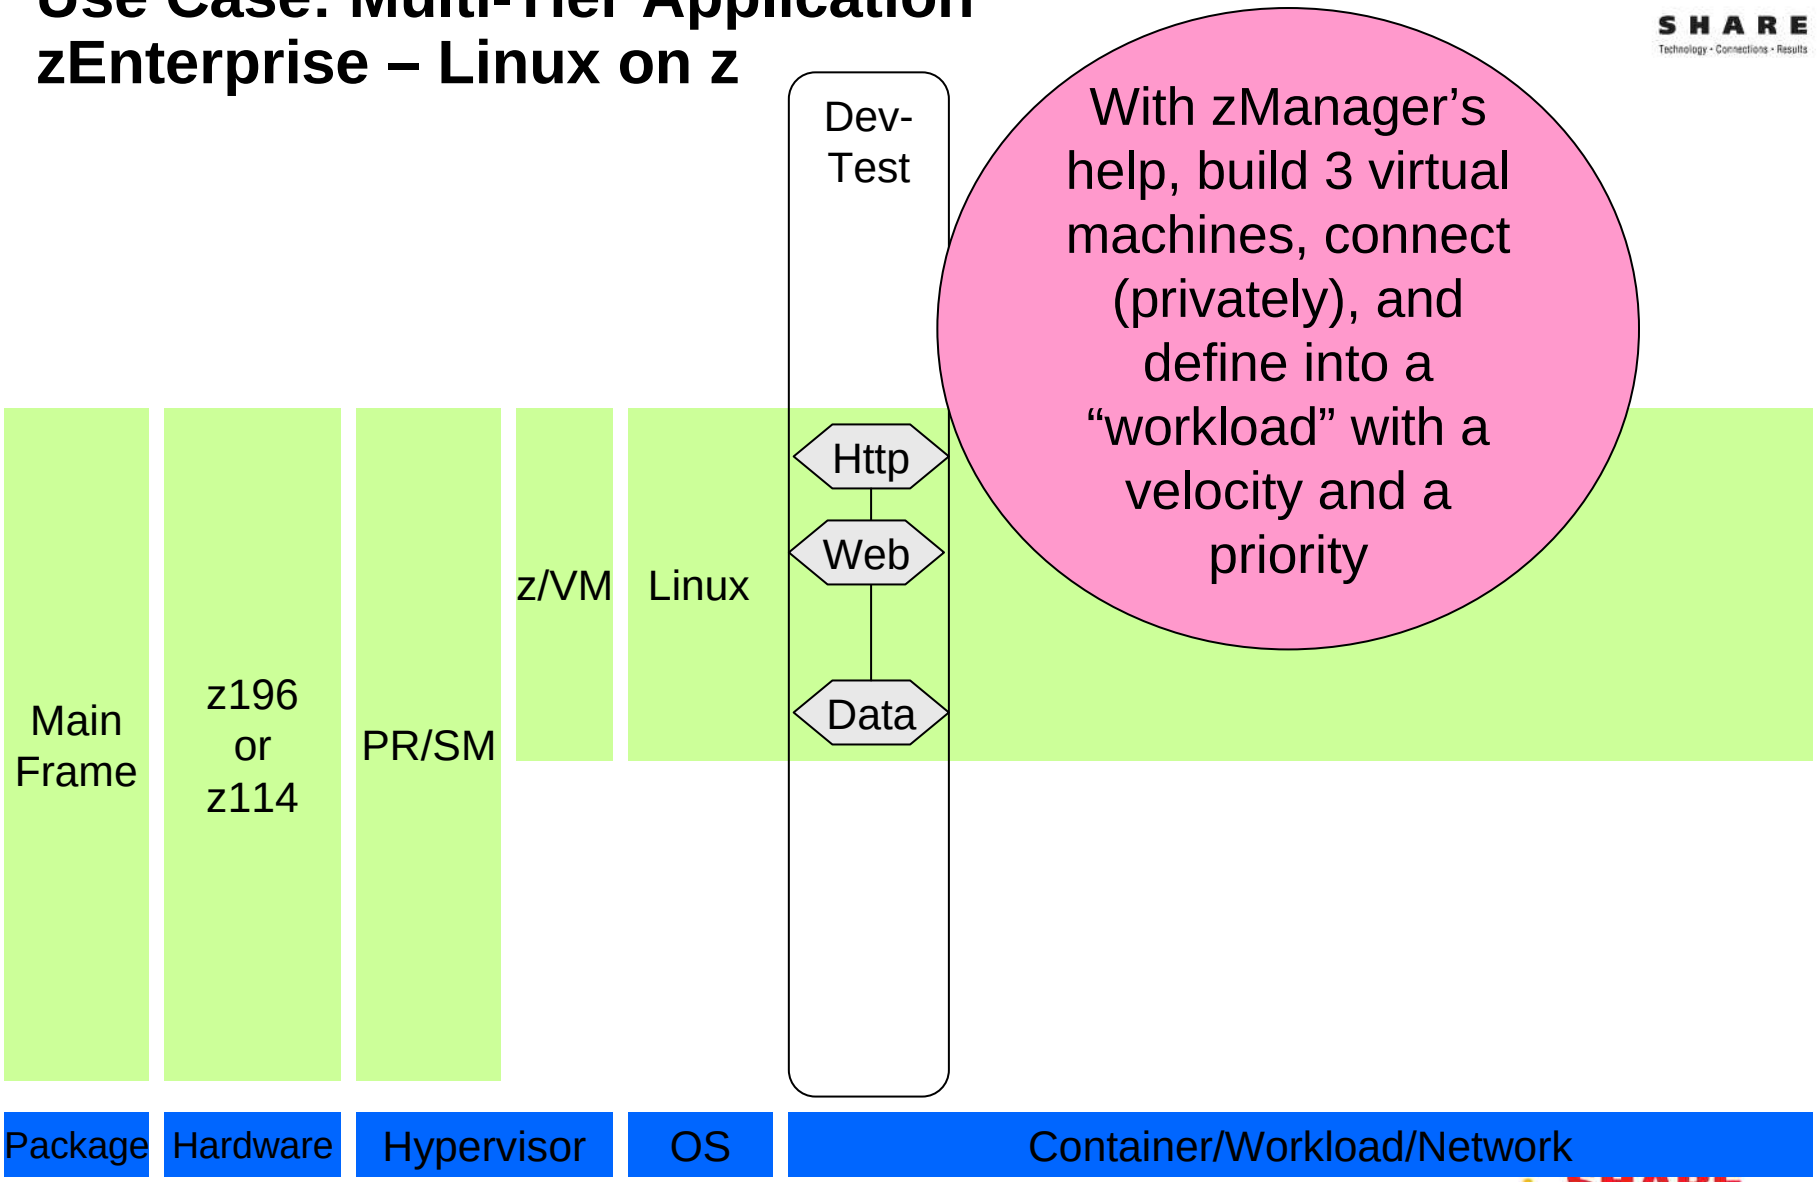

Topics

Use Case: Multi-Tier Application zEnterprise – Linux on z

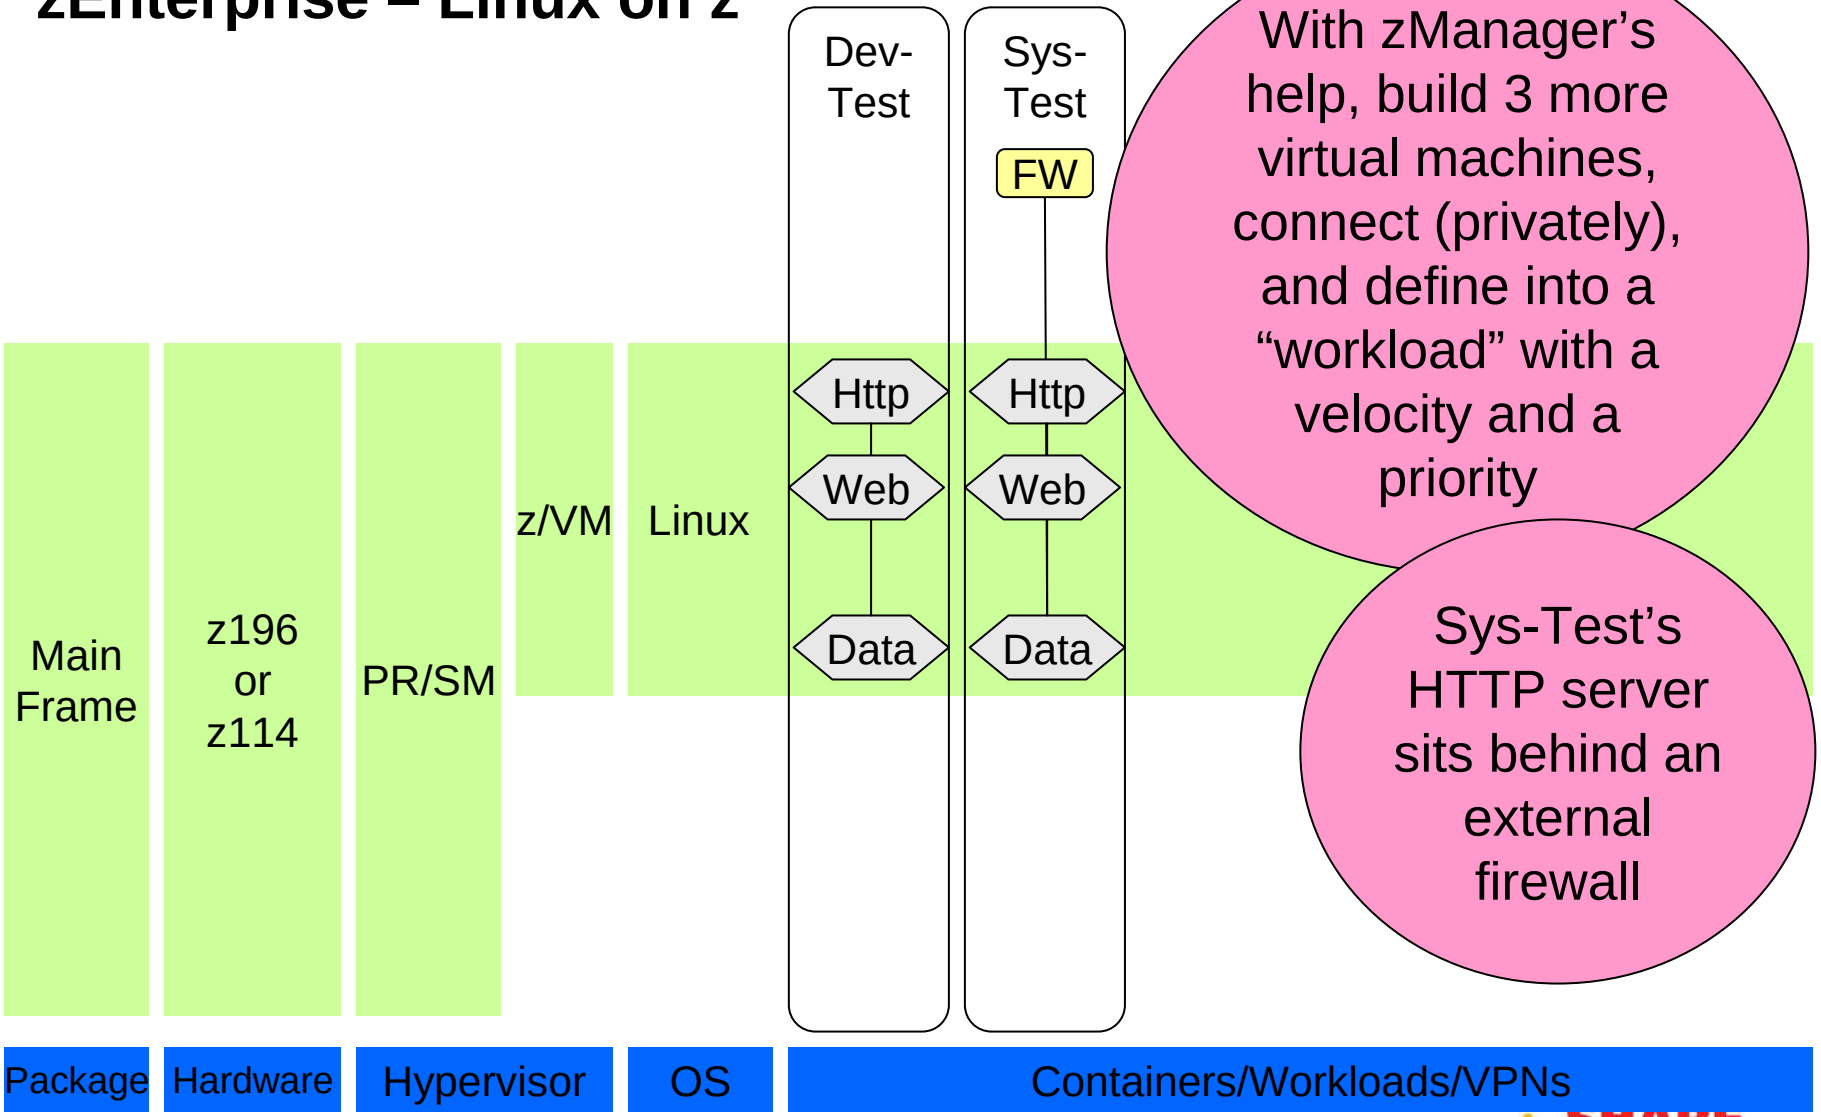

Topics

Use Case: Multi-Tier Application zEnterprise – Linux

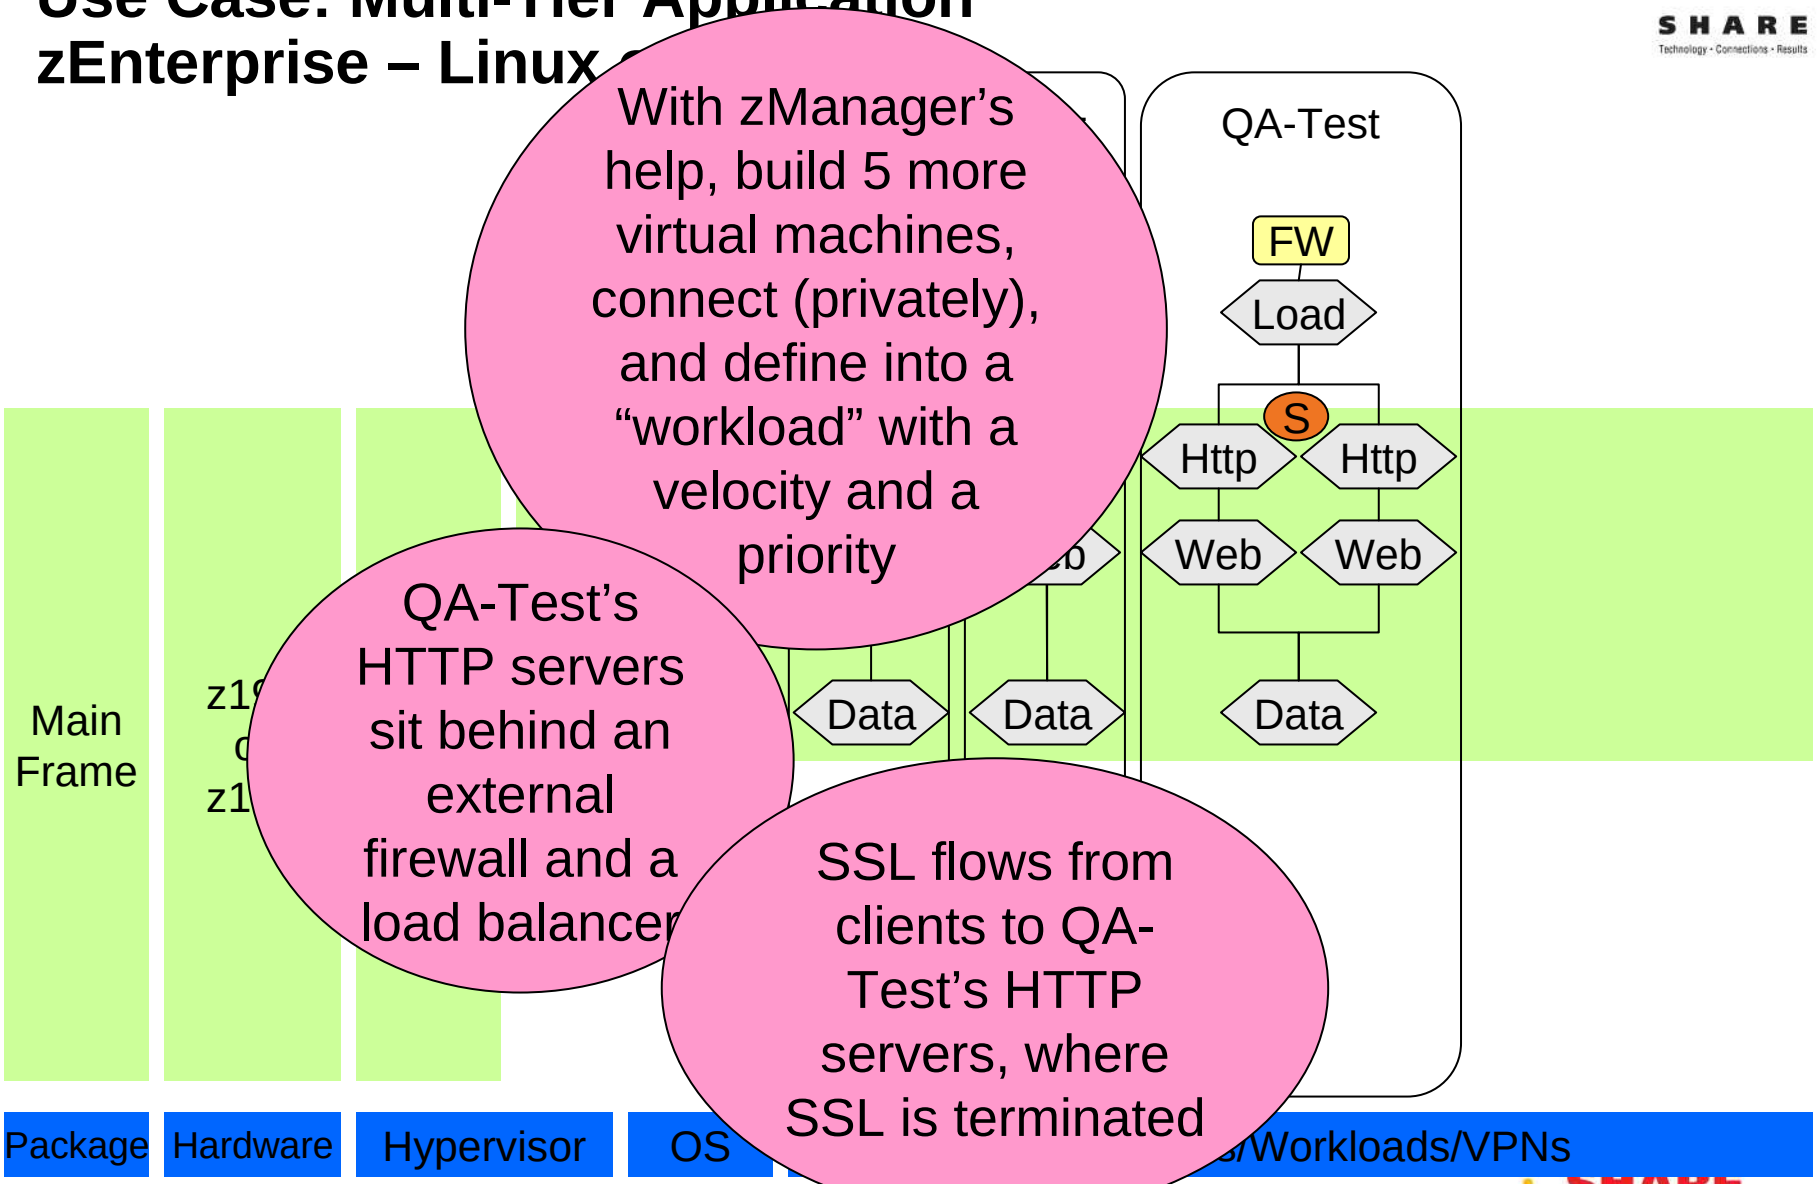

Topics

Use Case: Multi-Tier Application zEnterprise – Linux on z

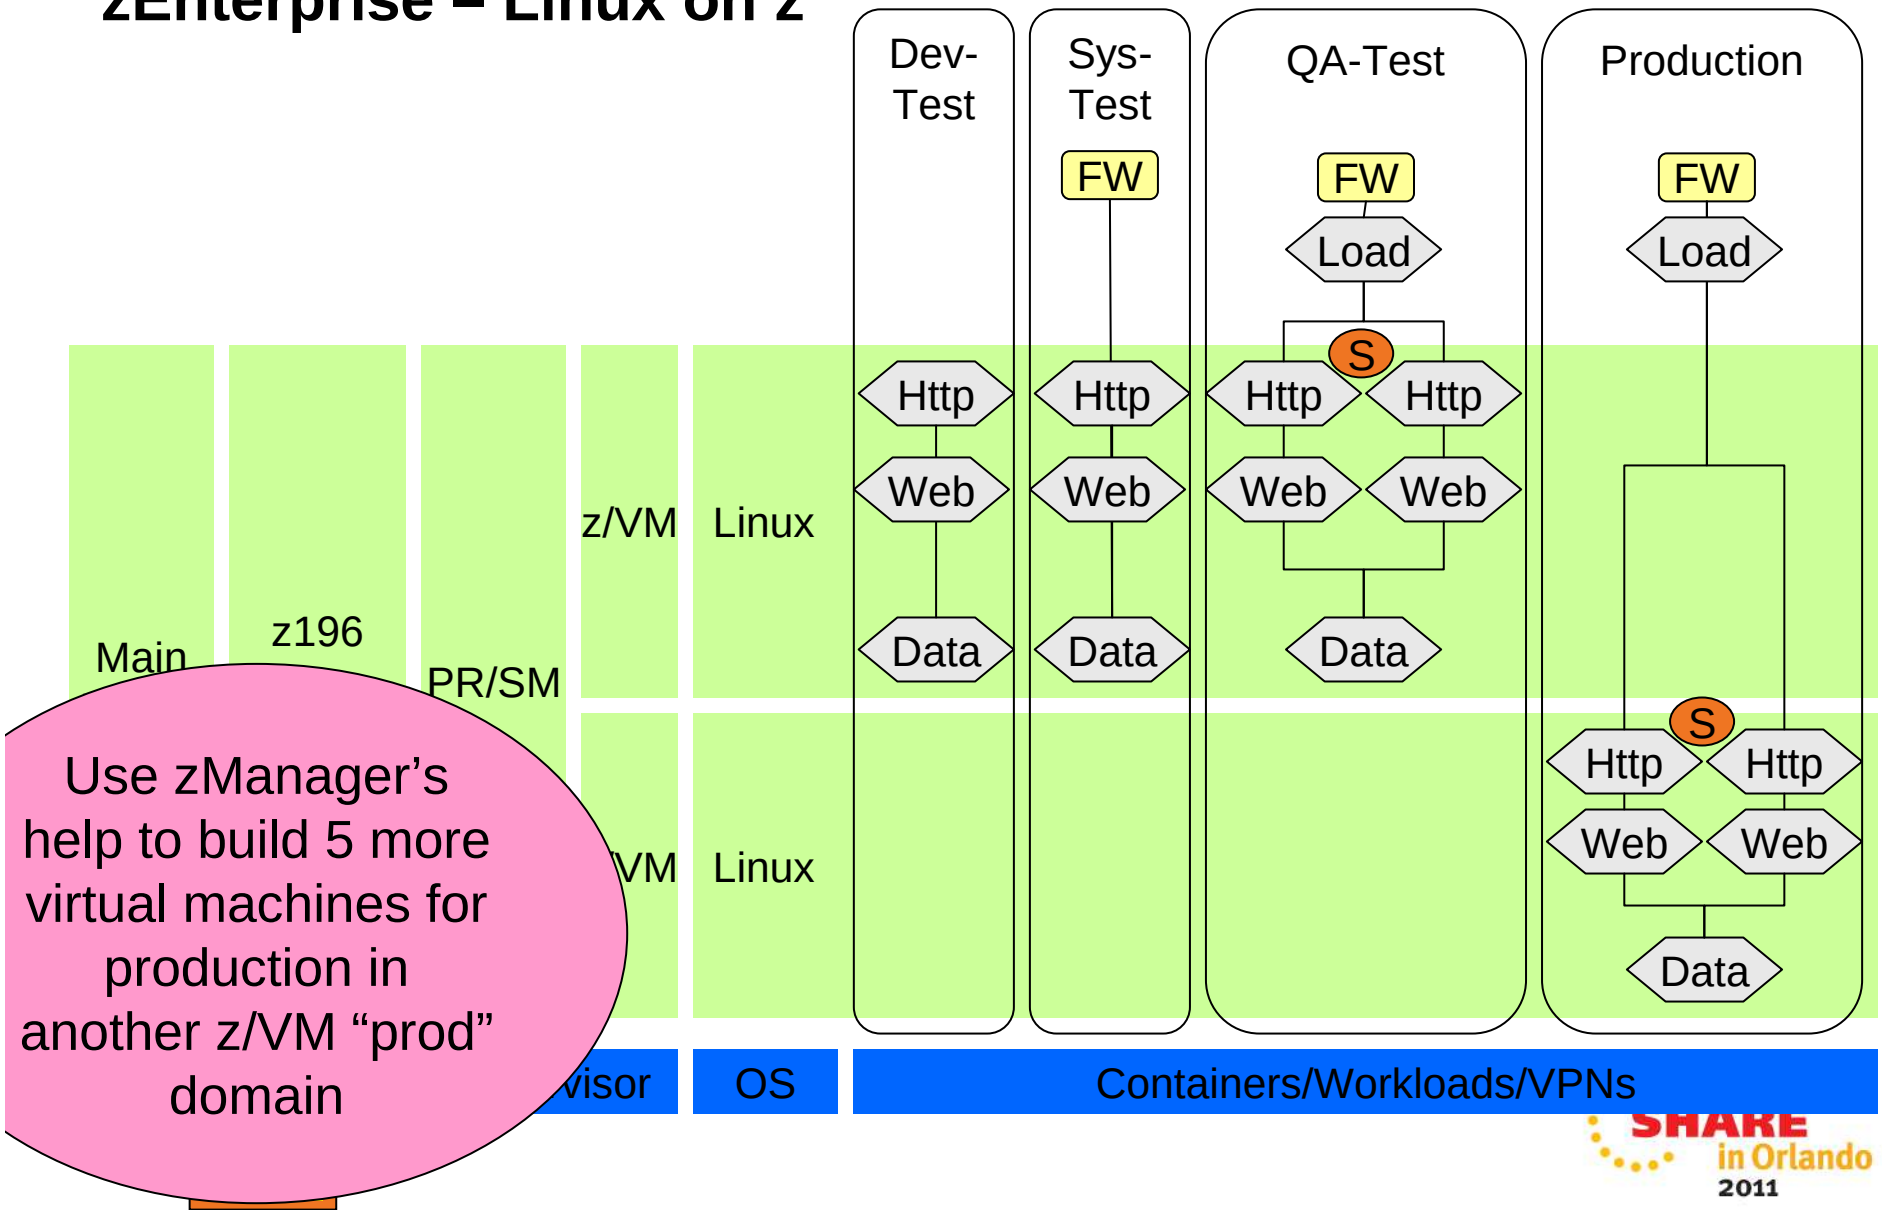

Use Case 2

Use Case: Multi-Tier Application zEnterprise – PS701 & z/OS

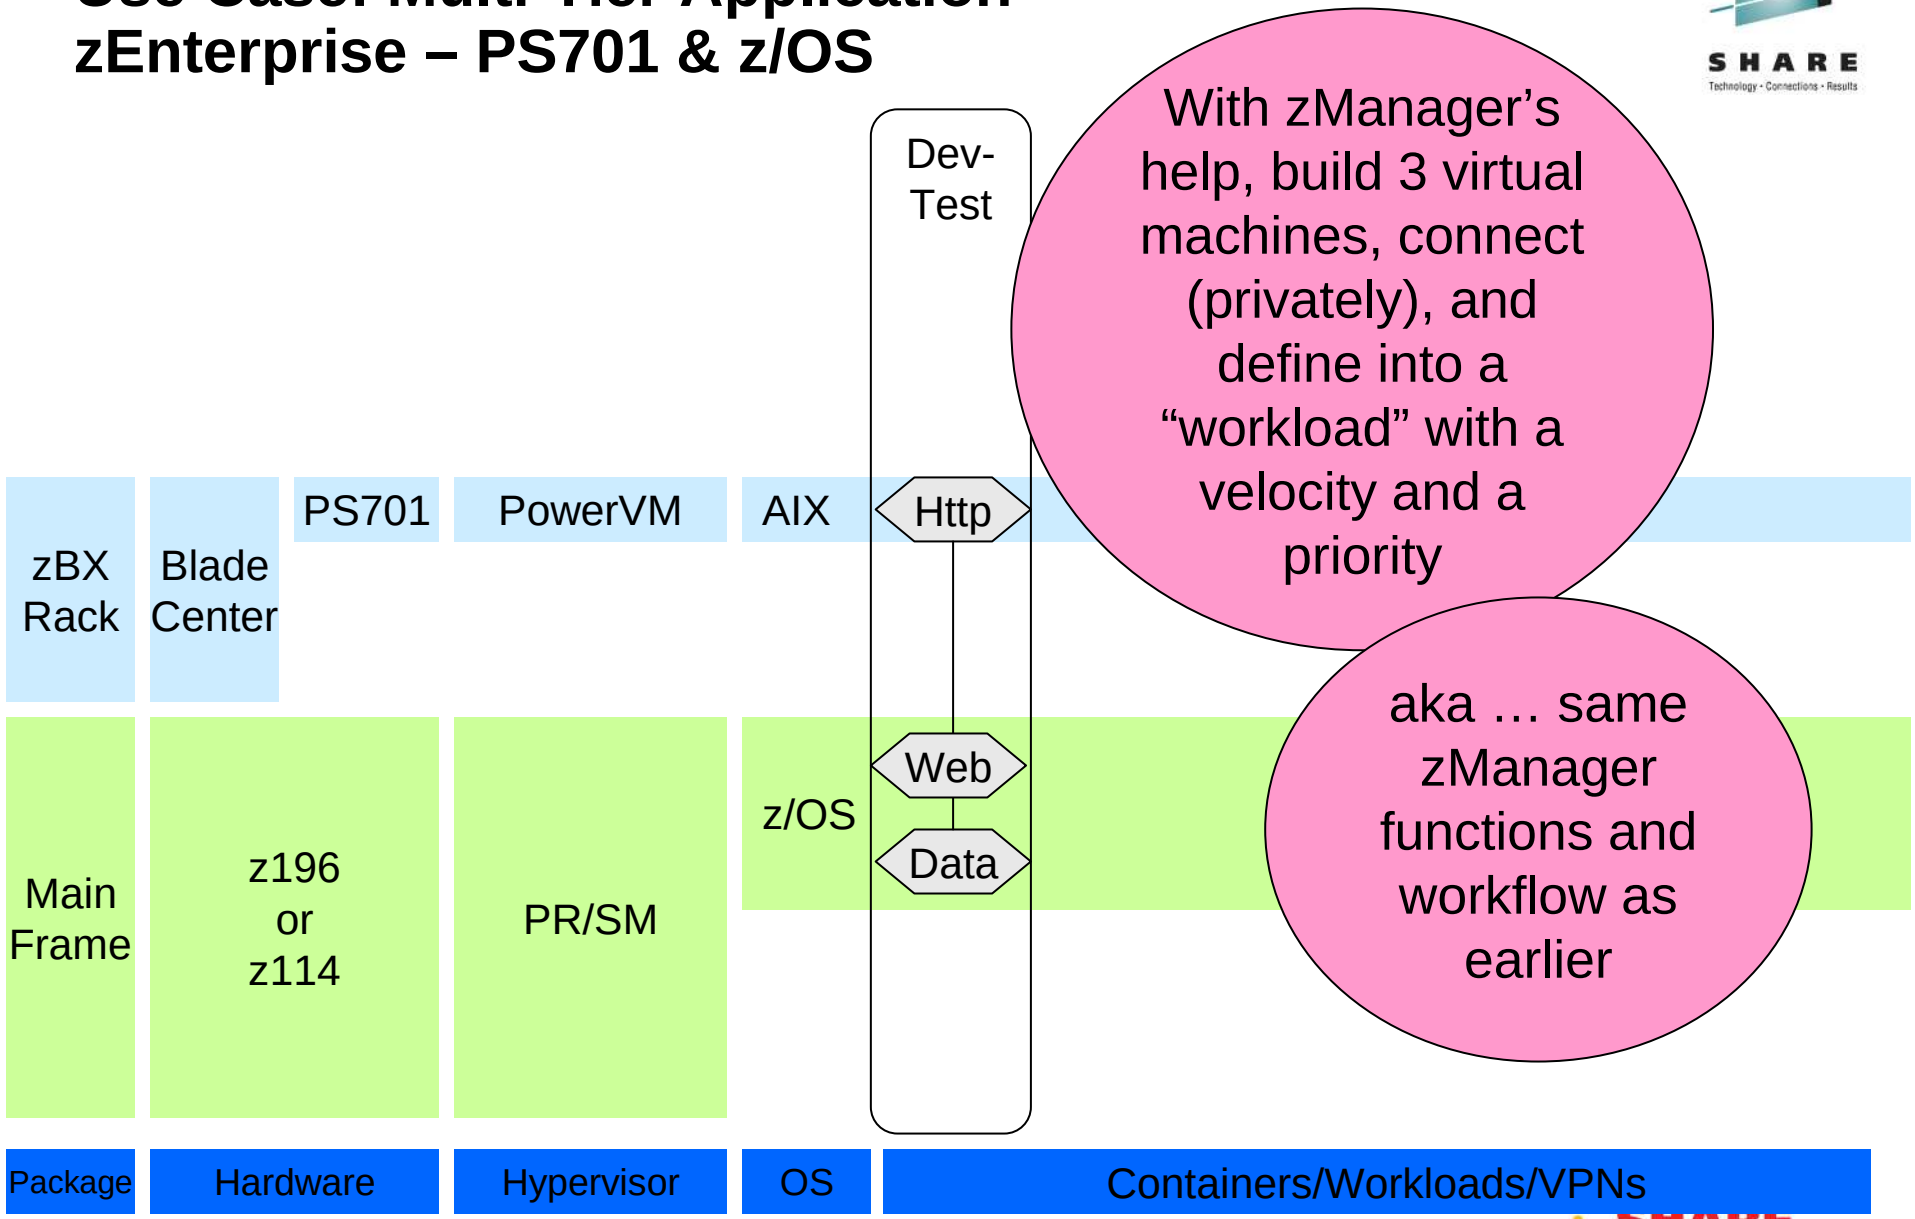

Topics

Use Case: Multi-Tier Application zEnterprise – PS701 & z/OS

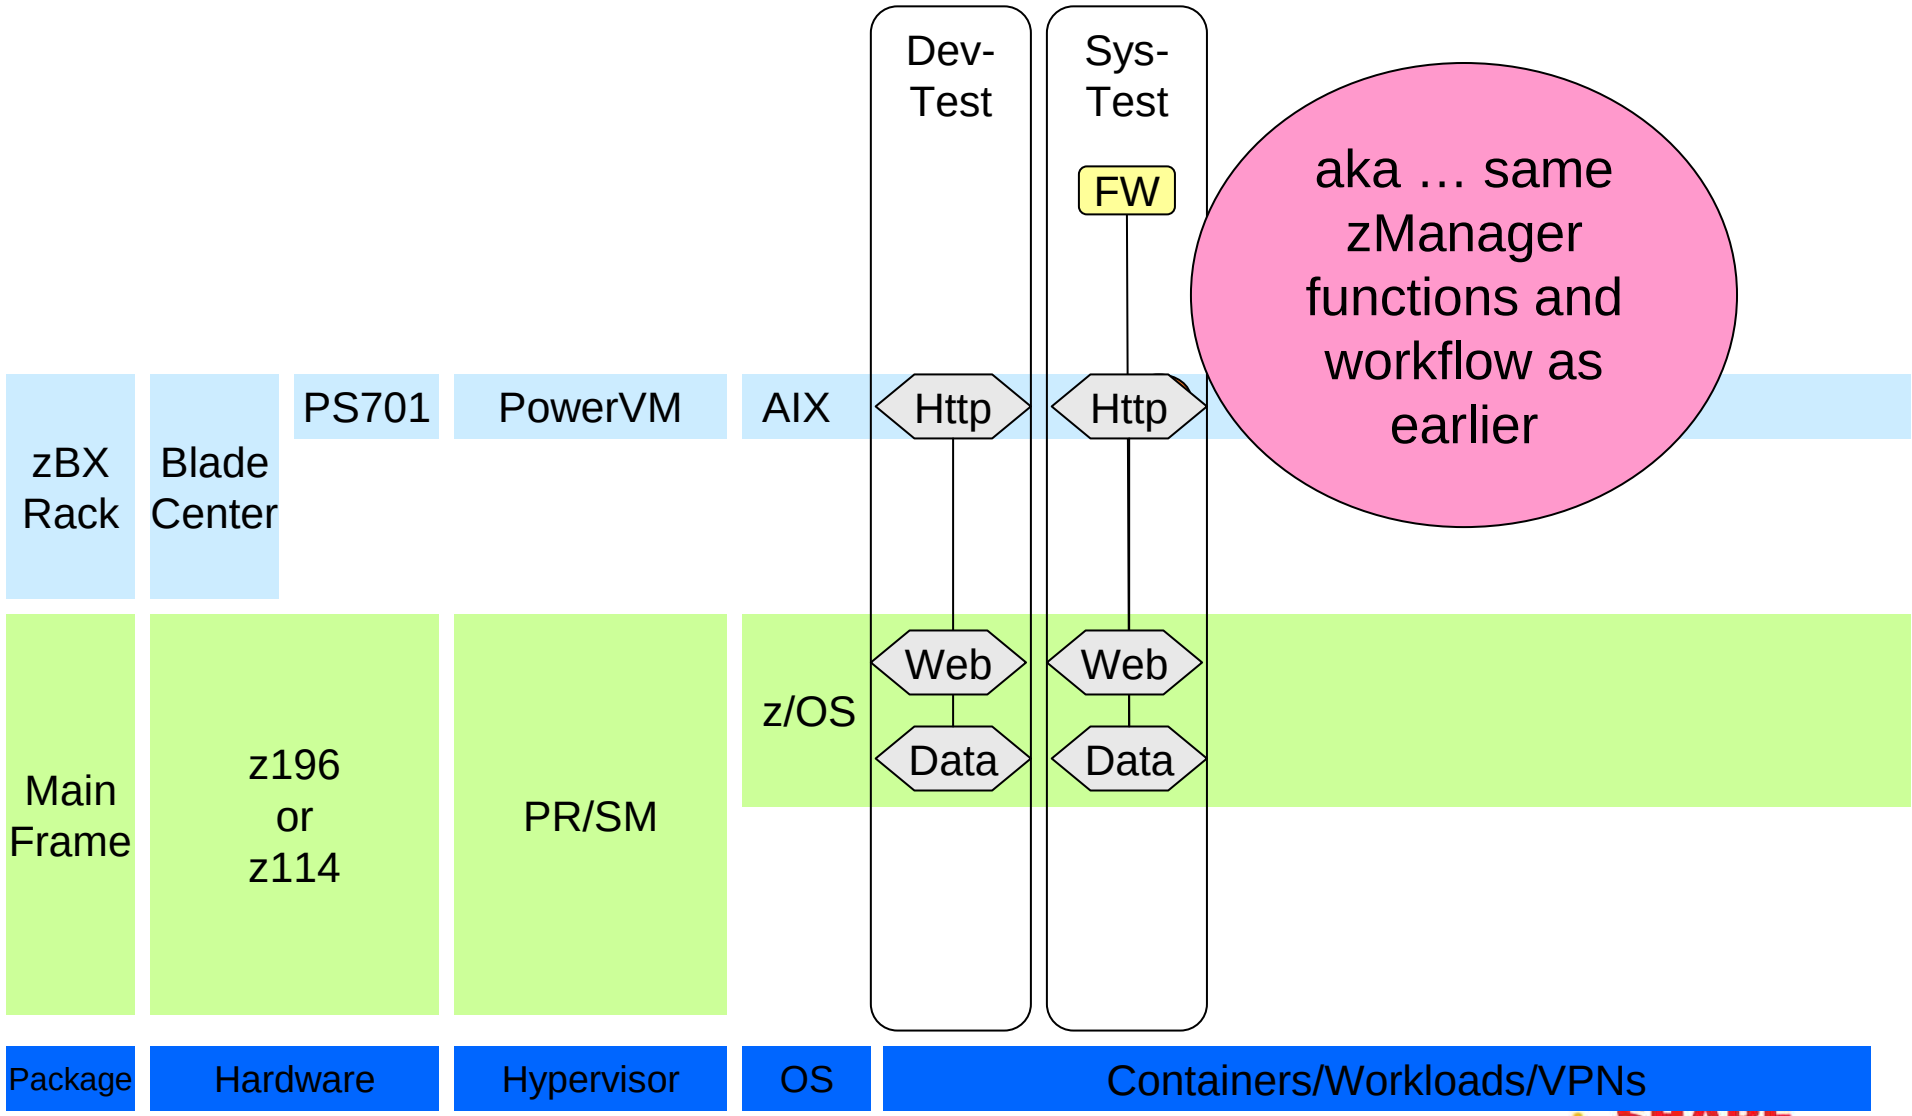

Topics

Use Case: Multi-Tier Application zEnterprise – PS701 & z/OS

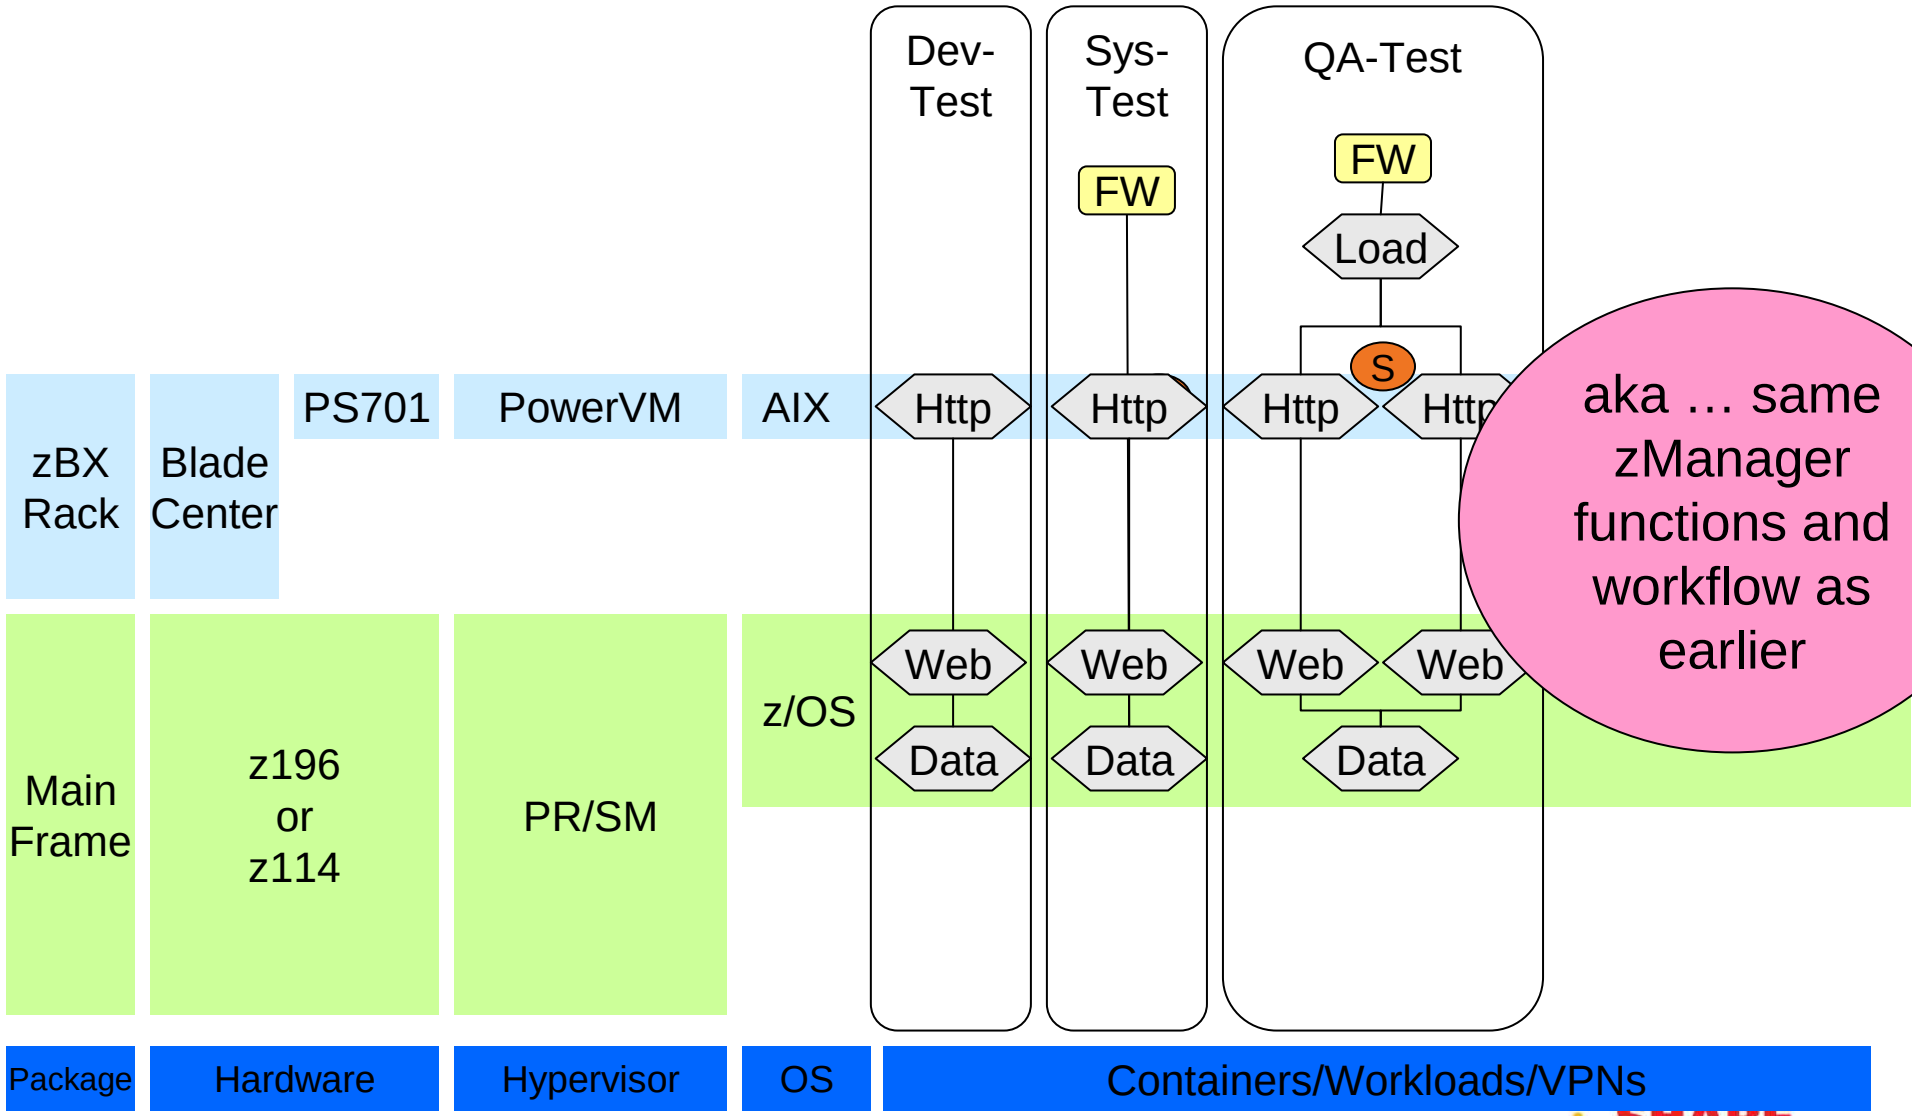

Topics

Use Case: Multi-Tier Application zEnterprise – PS701 & z/OS

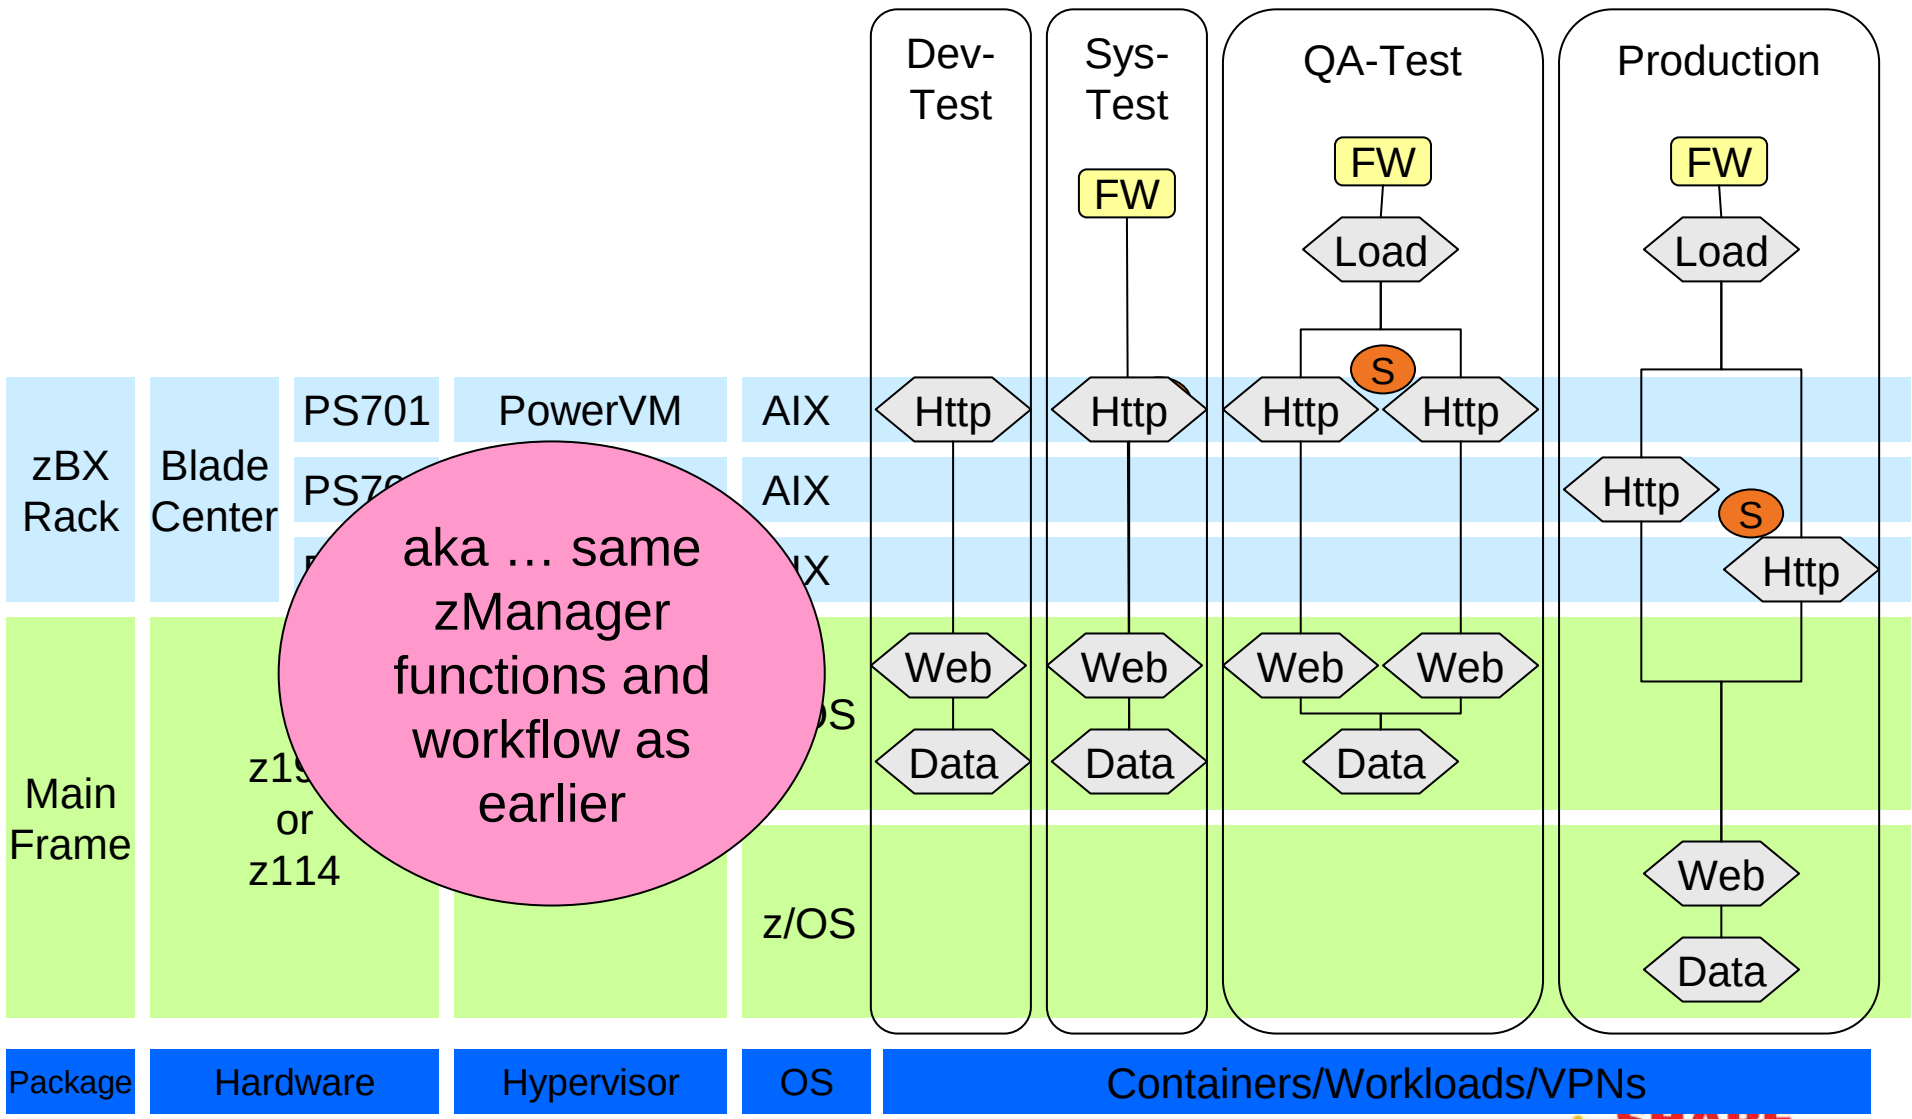

Topics

Use Case 3

Use Case: Real-Time Analytics zEnterprise – ISAO & z/OS

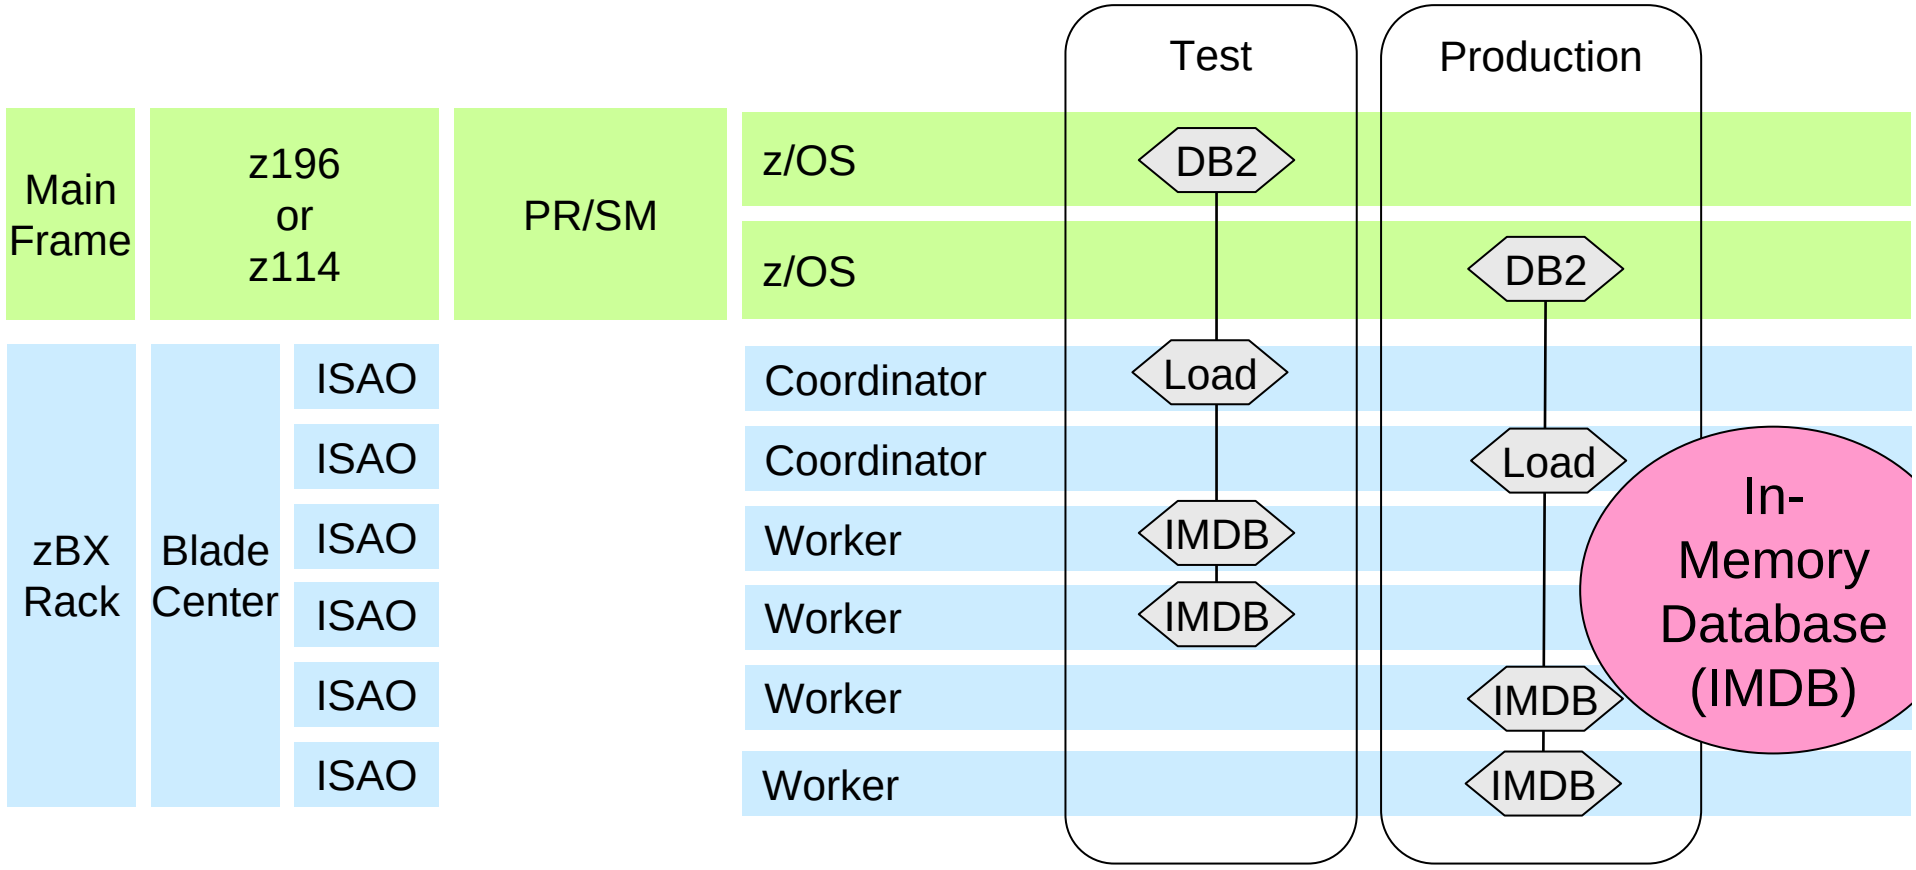

Topics

Use Case 4

Use Case: Hardware ESB
zEnterprise – DataPower XI50z & z/OS

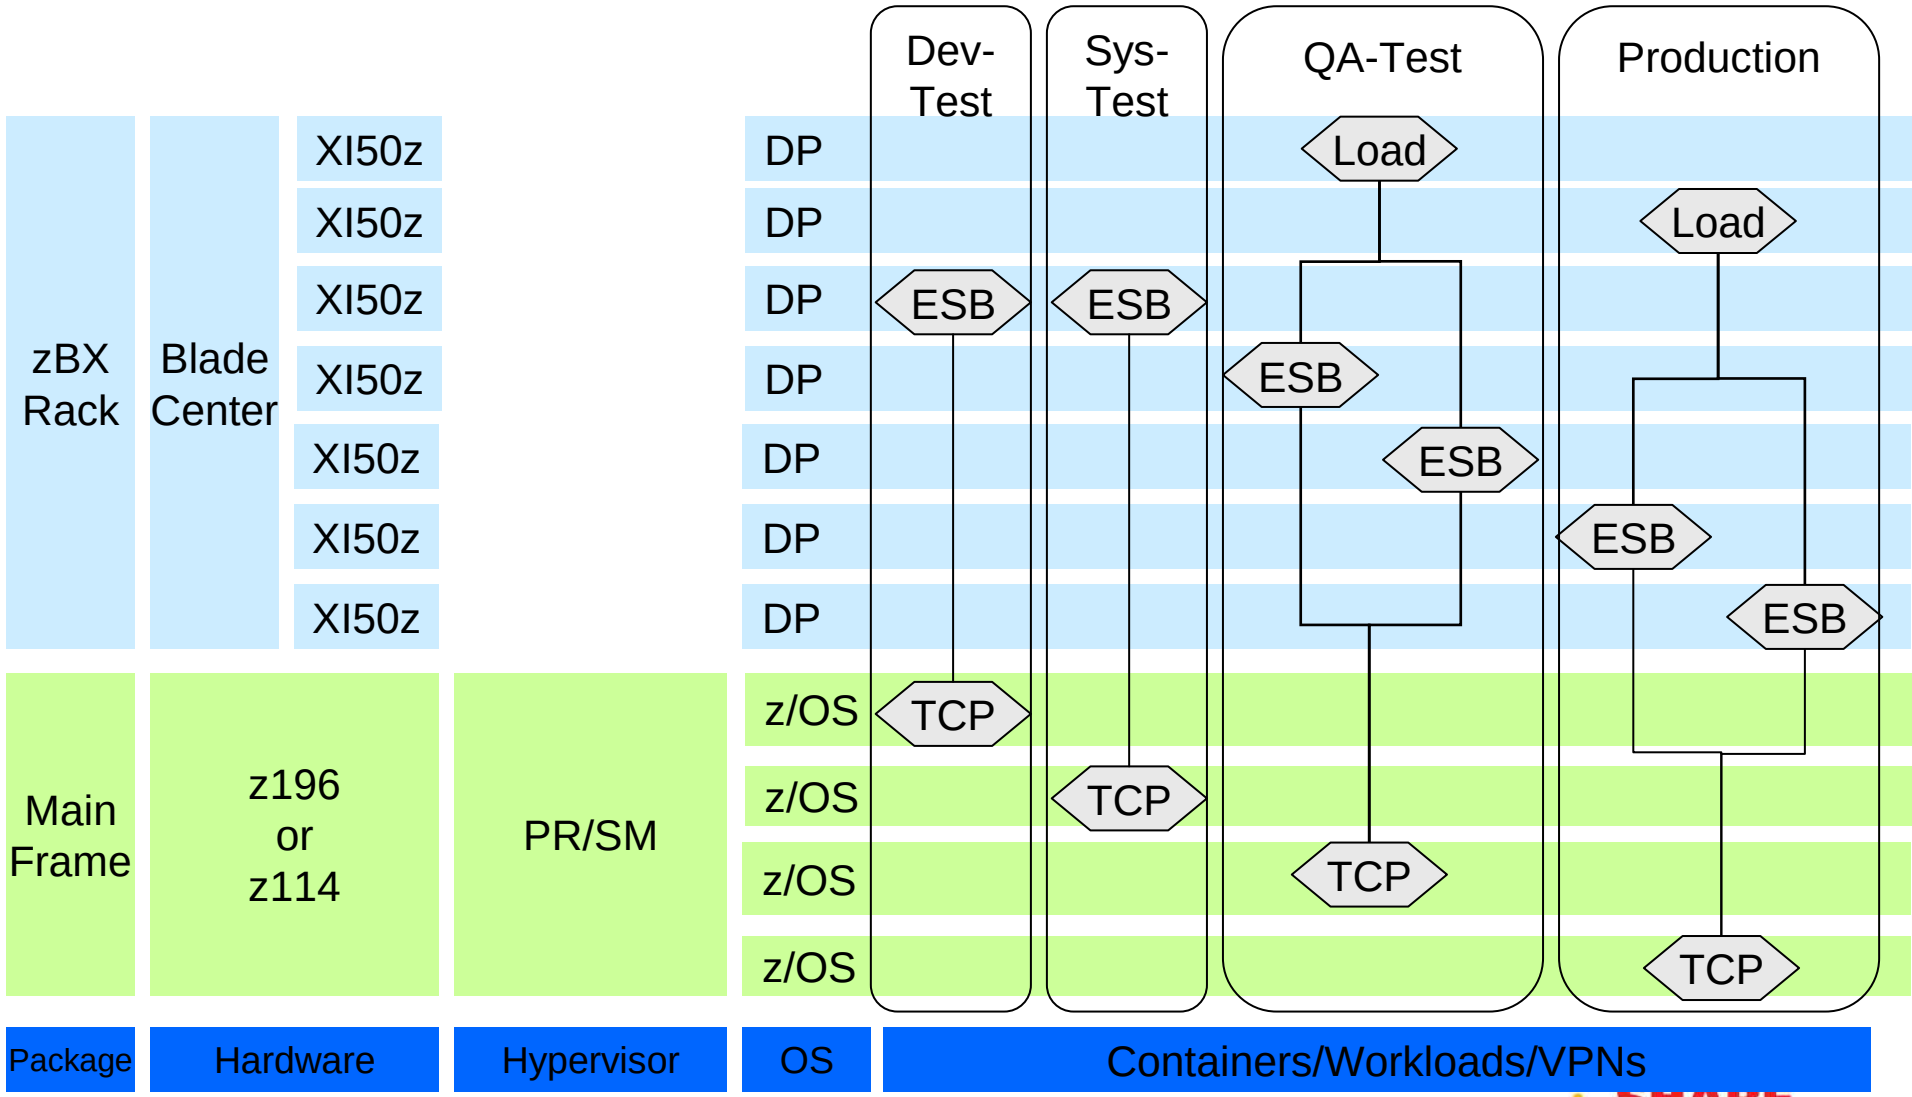

Topics

Use Case 5

Use Case: Storm Cloud

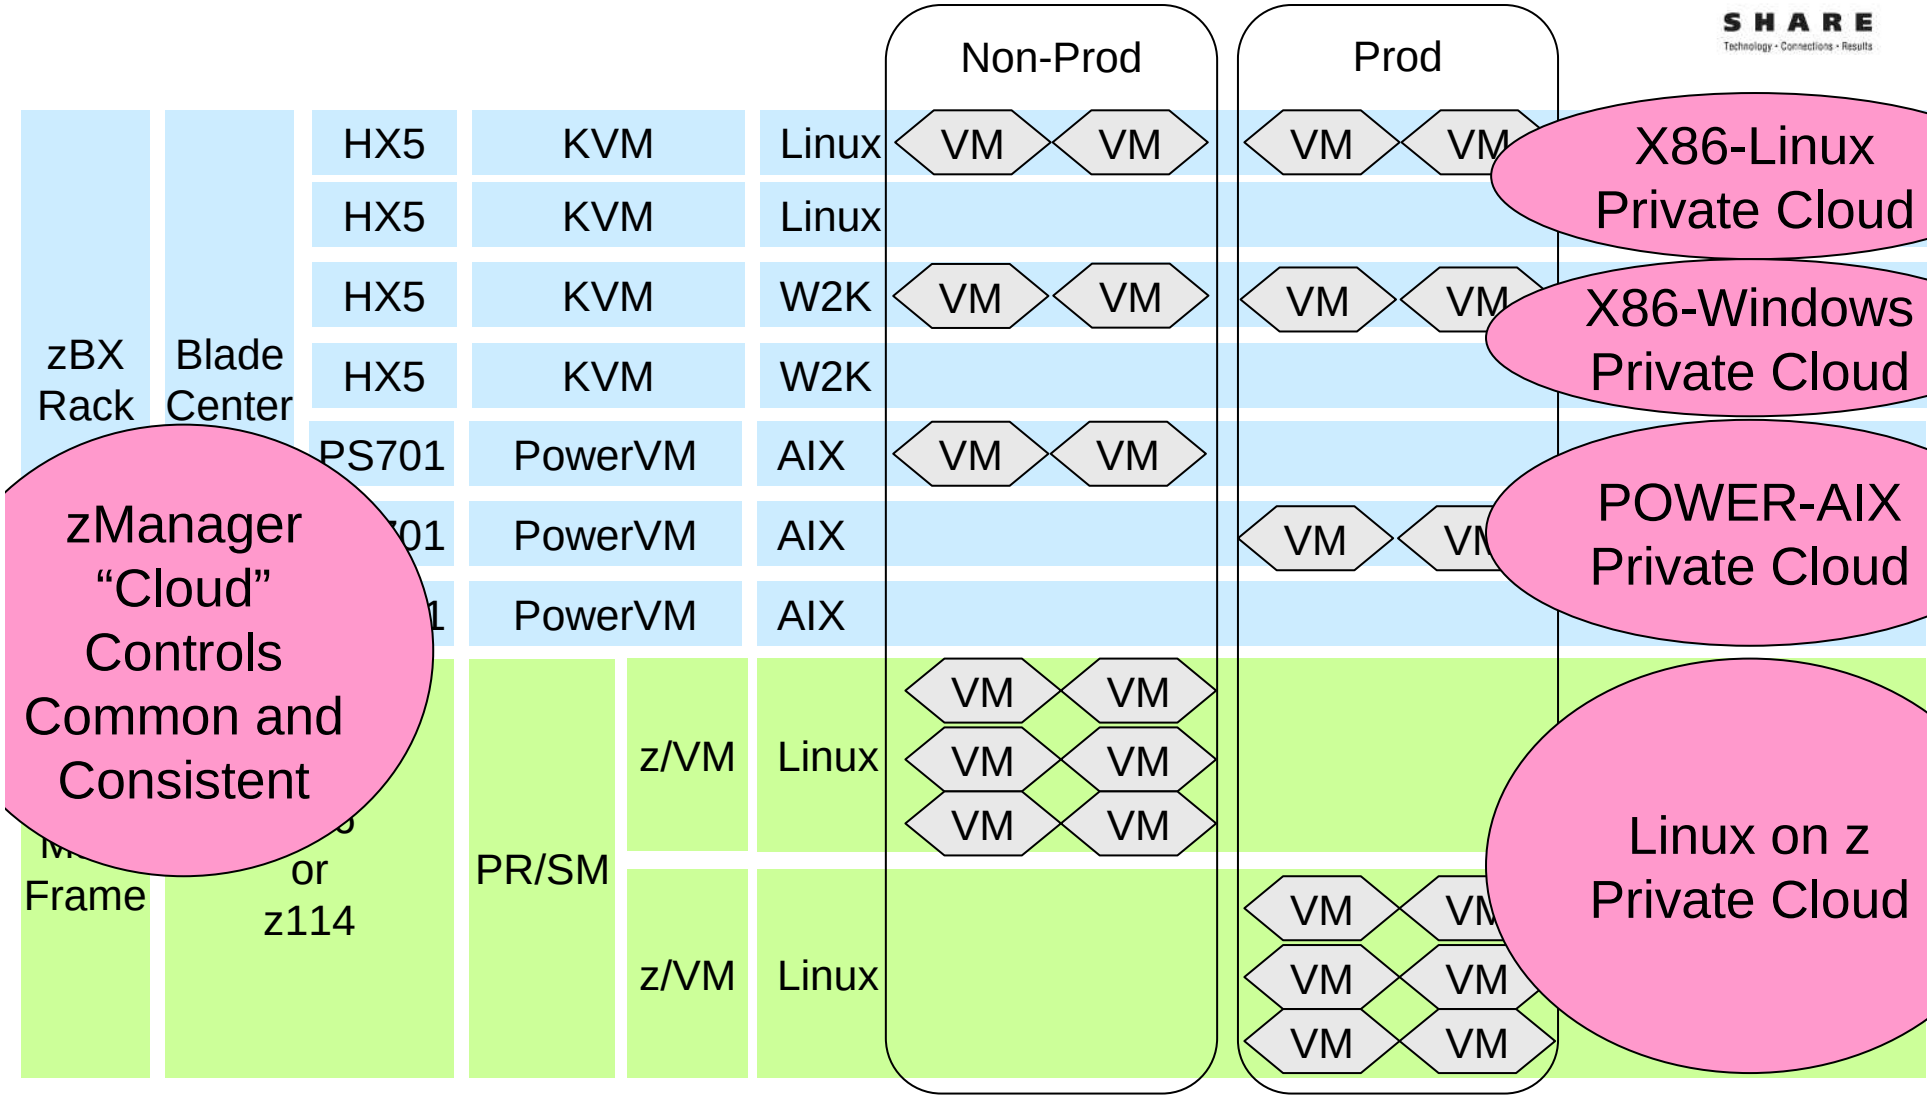

Use Case 6

Use Case: Parallel Batch

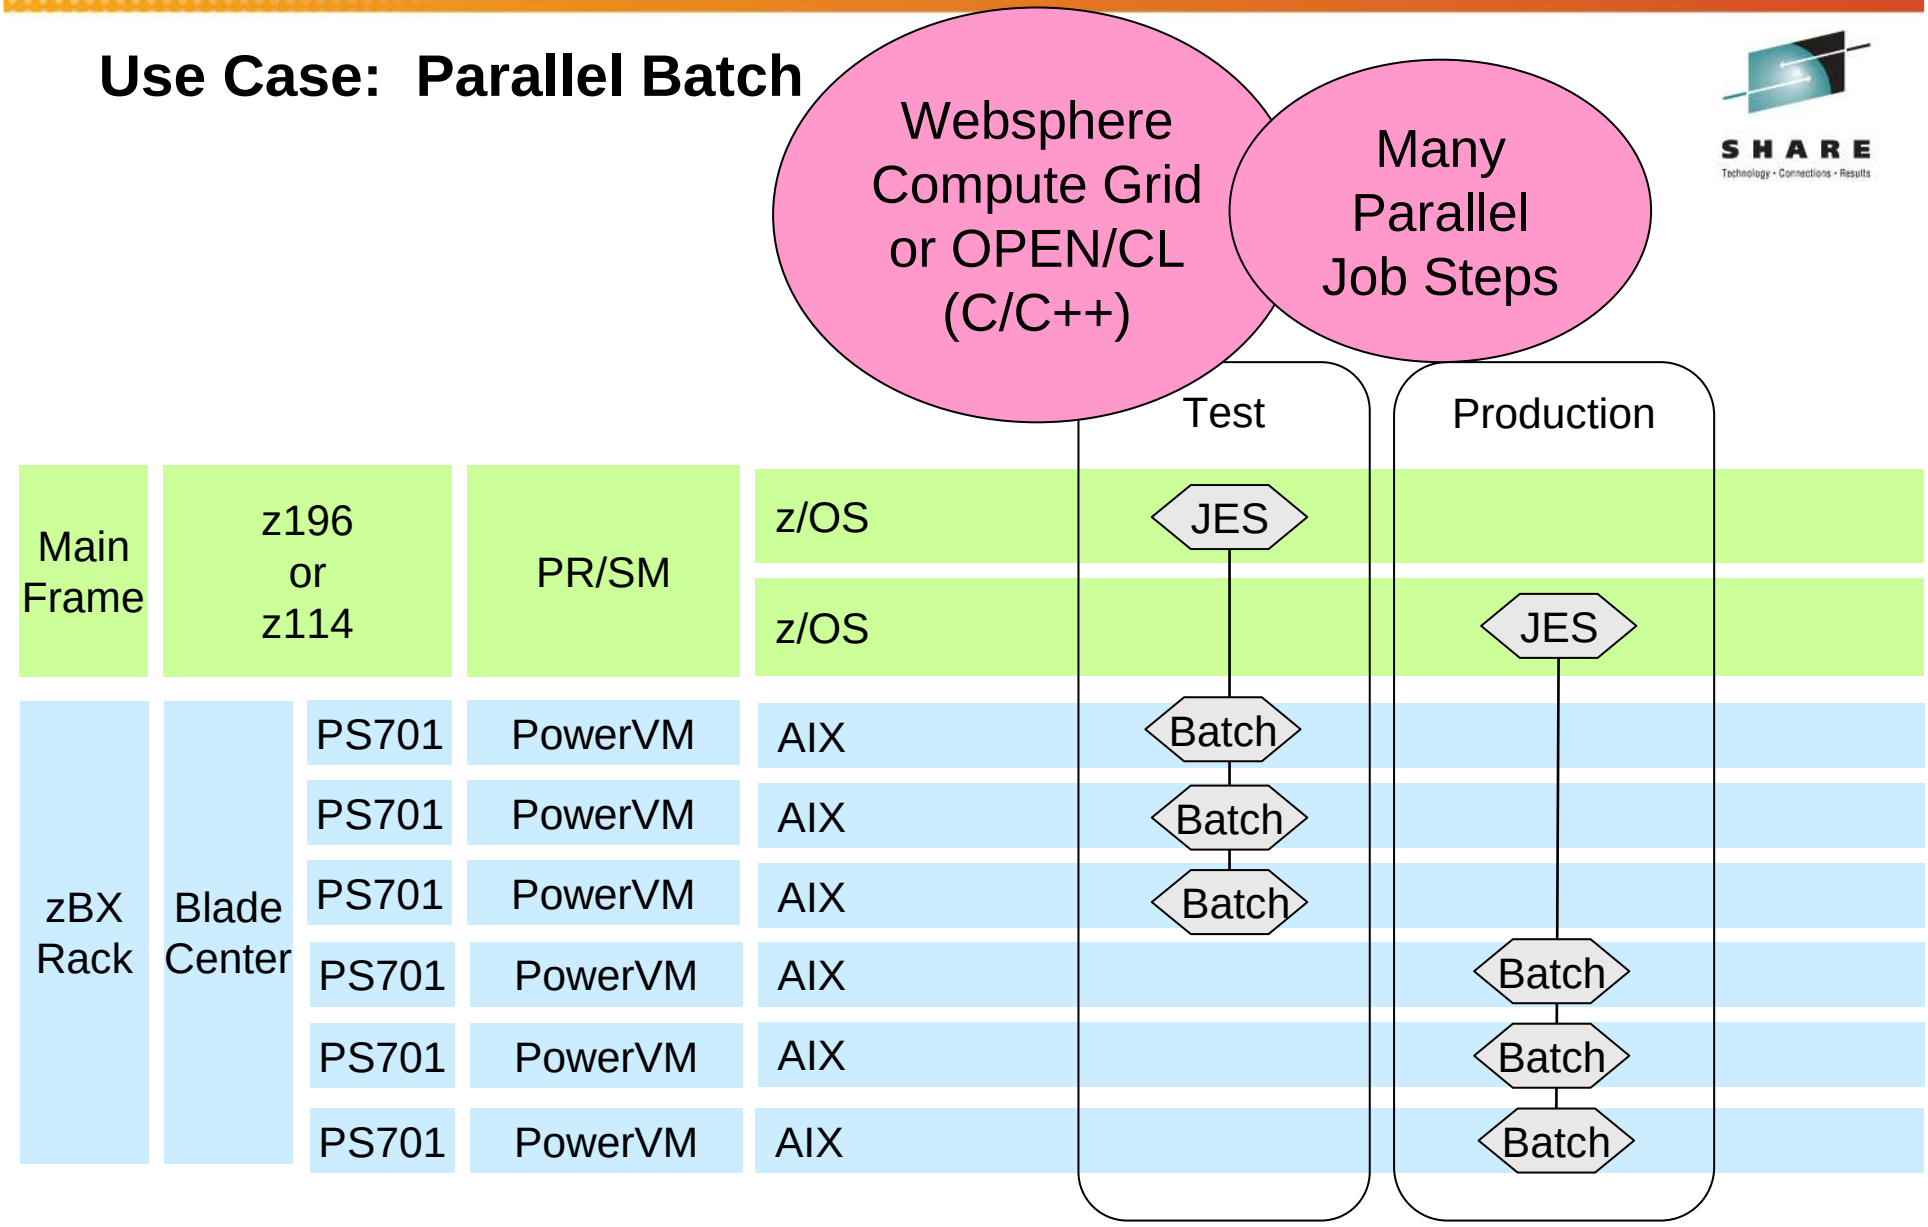

Topics

Use Case 7

Use Case: Combinatorial

Package	Hardware		Hypervisor	OS	Containers/Workloads/VPNs
zBX Rack	Blade Center	XI50z		DP	
		PS701	PowerVM	AIX	
		PS701	PowerVM	AIX	
		PS701	PowerVM	AIX	
Main Frame	z196 or z114		PR/SM	z/OS	
				z/OS	
zBX Rack	Blade Center	ISAO		Coordinator	
		ISAO		Coordinator	
		ISAO		Worker	
		ISAO		Worker	
		ISAO		Worker	
		ISAO		Worker	

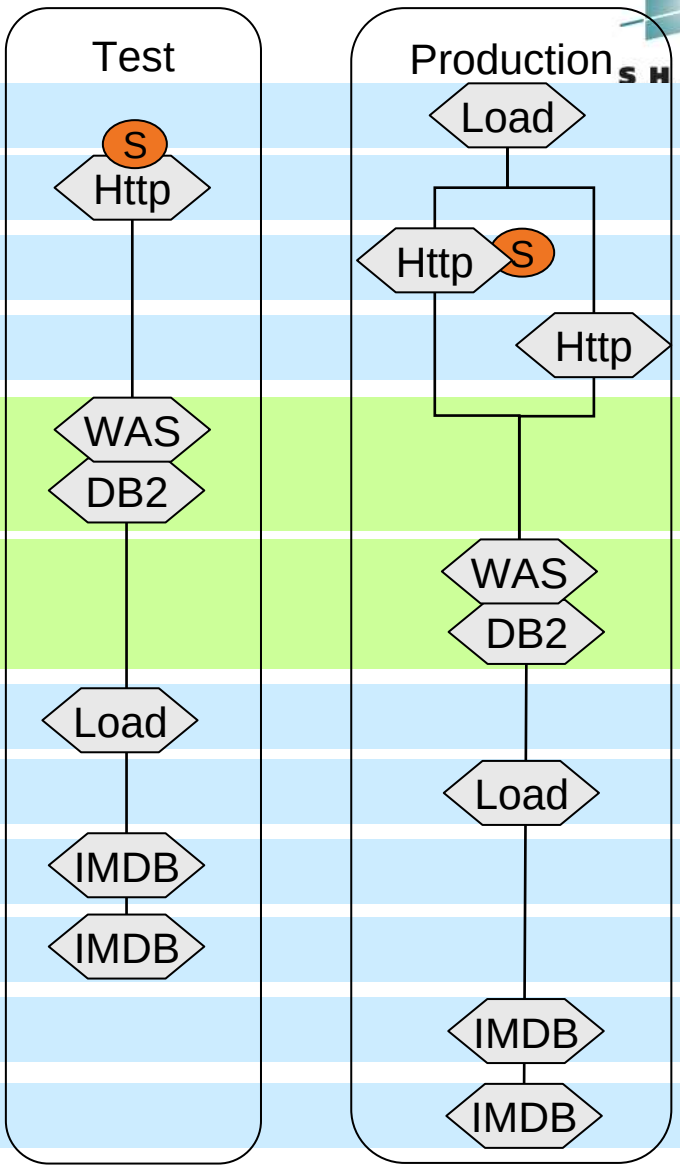

Topics

End of Section

Conclusion

Conclusion – We Talked About… Smarter Applications… Smarter Transactions… and Smarter Business Process Improvement

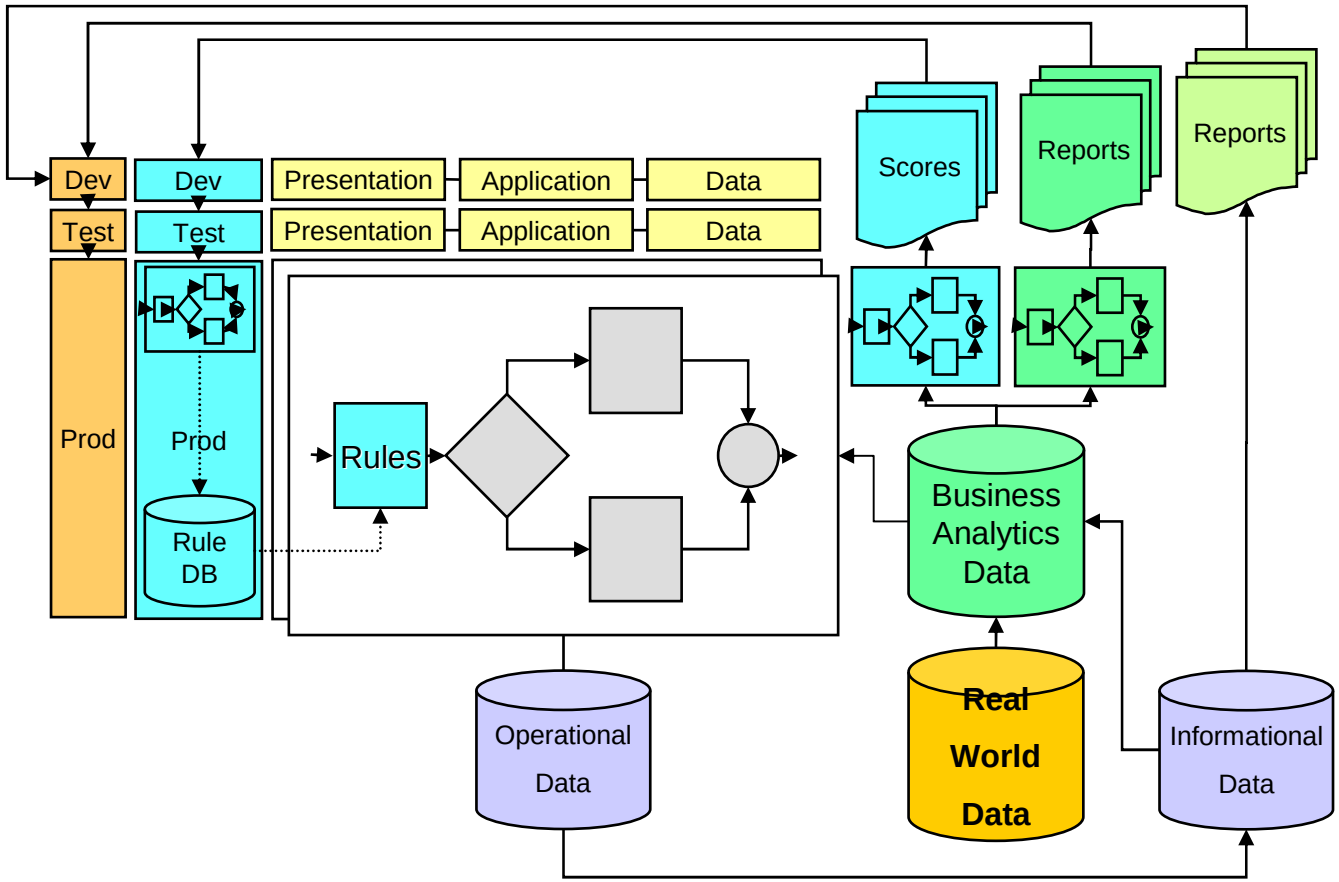

Topics

Conclusion – We Talked About…
The Internet of Things …
Big Data …
Extending the Reach of IT …
and Requirements for Smarter Servers

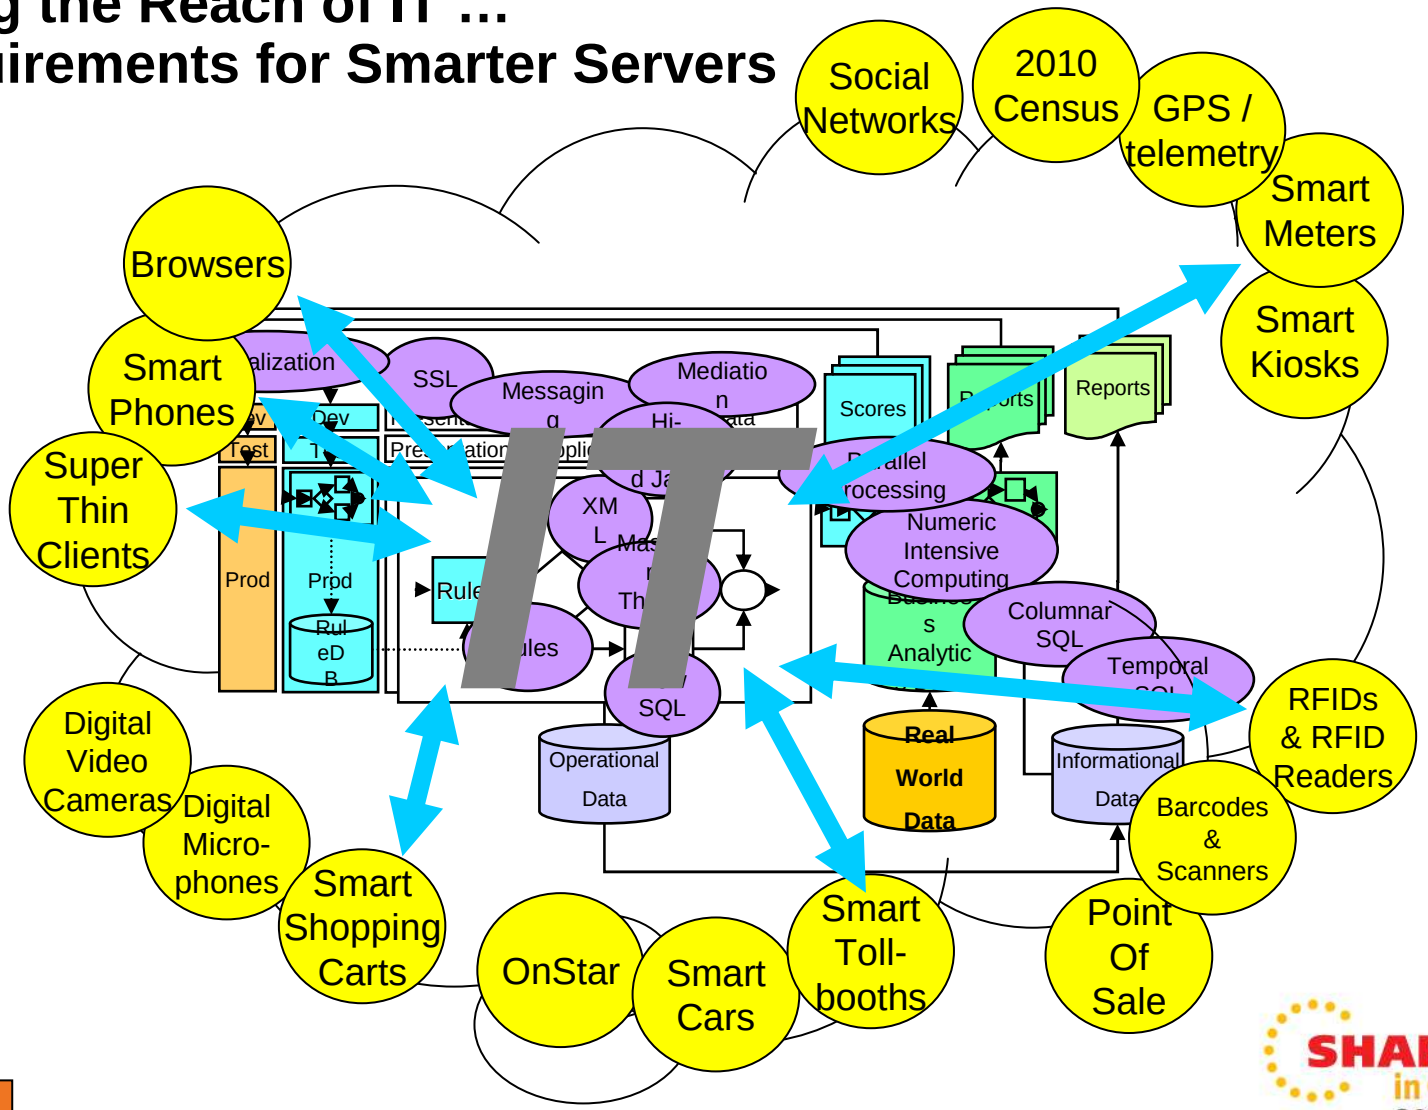

Topics

Conclusion – We Talked About…
Hybrid Computing

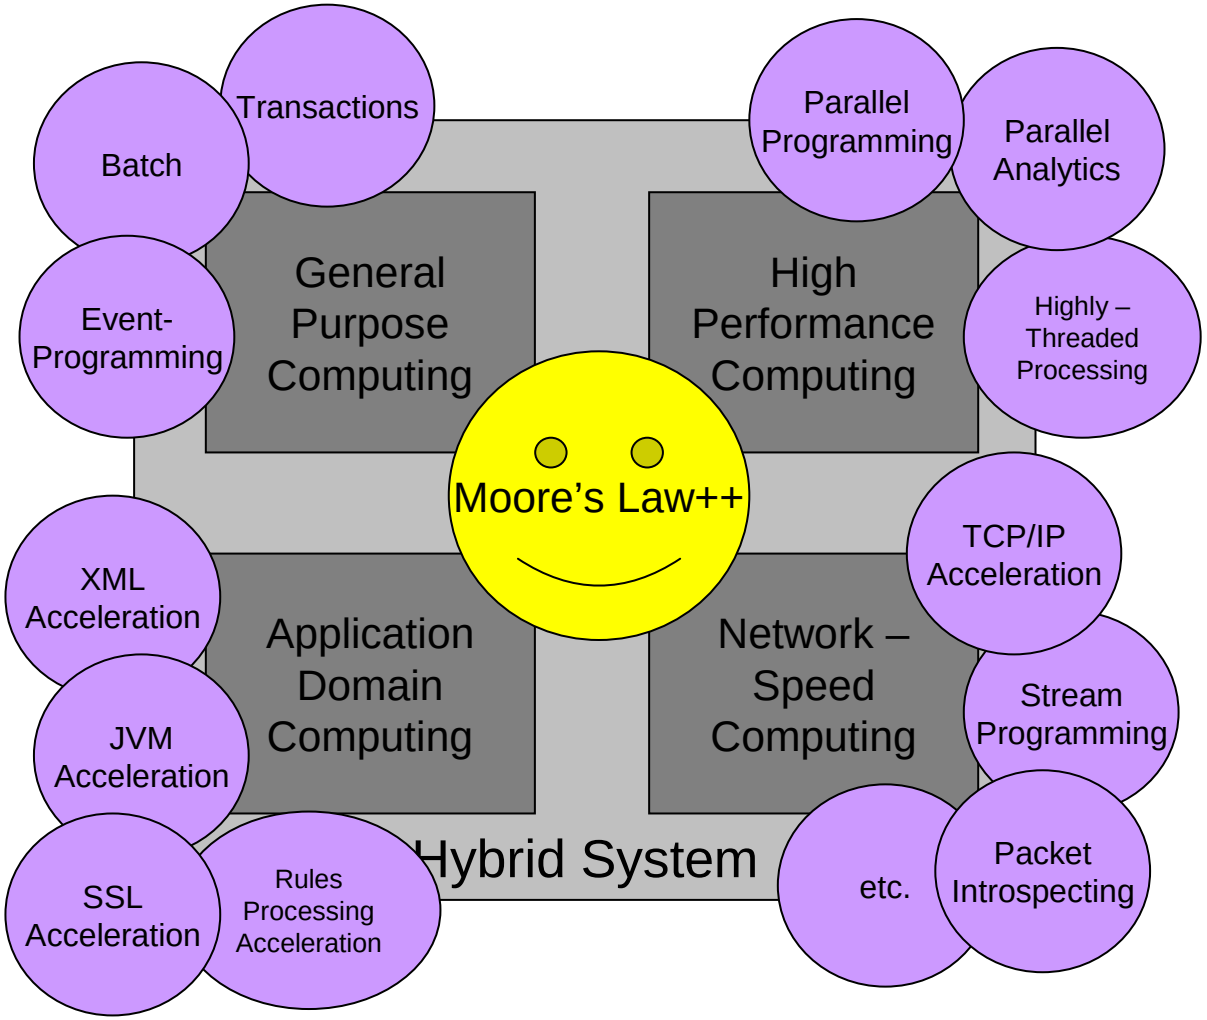

Topics

Conclusion – We Talked About…
State of the Art General Purpose Computing

z10 EC

MICROPROCESSOR:
4.4GHz
Superscalar
In Order
L1 (192KB) and
L1.5 (3MB) per core
(total:12.75MB per chip)
1B transistors
44 new instructions

MCM:
5 "z6" chips
L2 (48MB on 2SCchips)

CEC:
1 to 4 books
64/77 cores
up to 1.5TB RAM
~1,000 to ~30,000MIPS
6GBps Infiniband

z196

MICROPROCESSOR:
5.2GHz
Super-duper-scalar
Out of Order
L1 (192KB) and
L2 (1.5MB) per core
(new) L3 per chip (24MB)
(total:30.75MB per chip)
1.4B transistors
110 new instructions

MCM:
6 "z7" chips
L4 (192MB on 2SC chips)

CEC:
1 to 4 books
80/96 cores
up to 3.0TB ***RAIM***
~1,500 to ~50,000MIPS
6GBps Infiniband
Hot plug I/O drawers
Optional Water Cooling
Optional DC Power

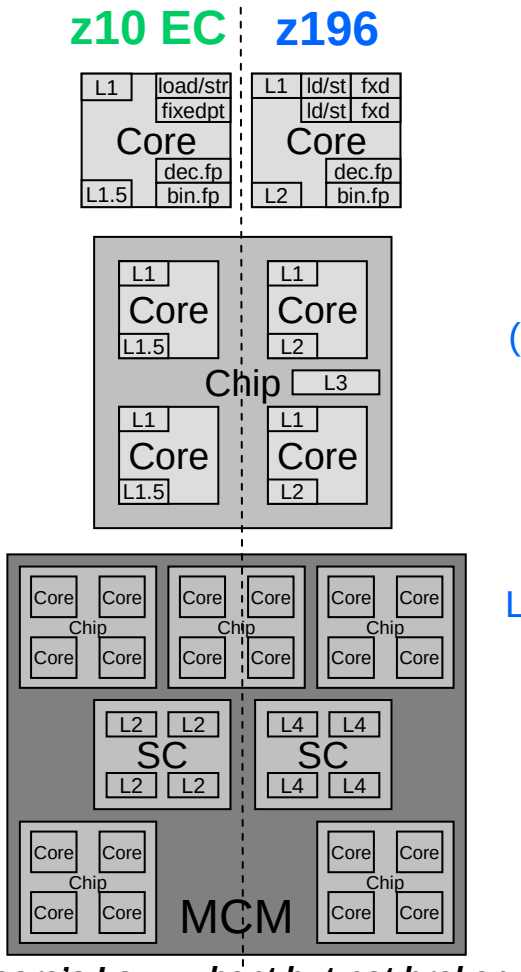

Moore's Law … bent but not broken

Topics

Conclusion – We Talked About…
zEnterprise: the World's First Hybrid Server

zEnterprise = z196/z114 + zBX + zManager

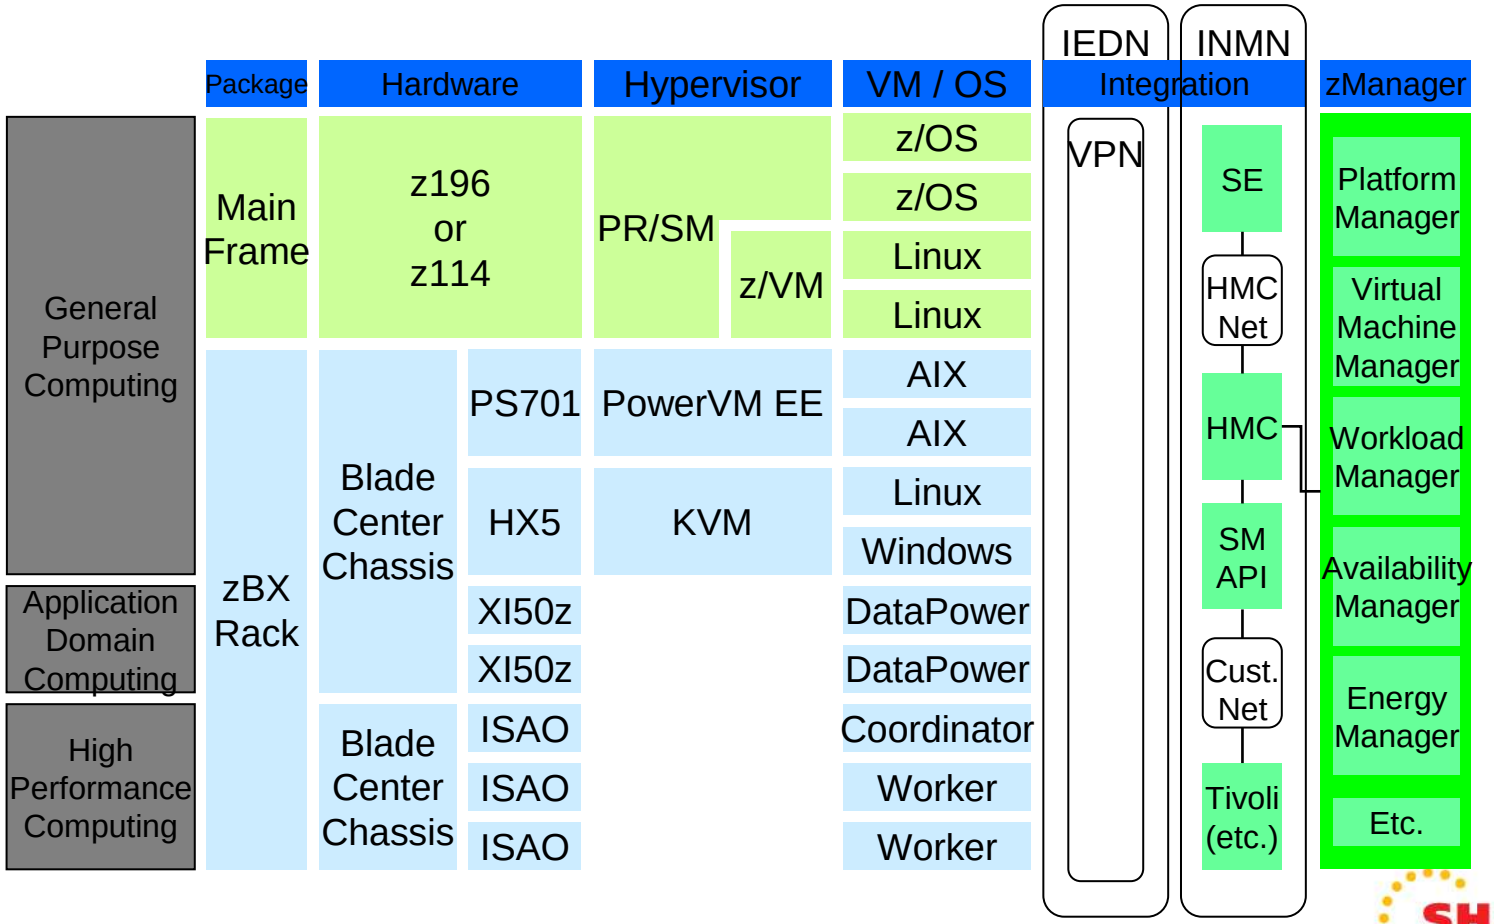

Topics

Conclusion: zEnterprise – Ready for Business

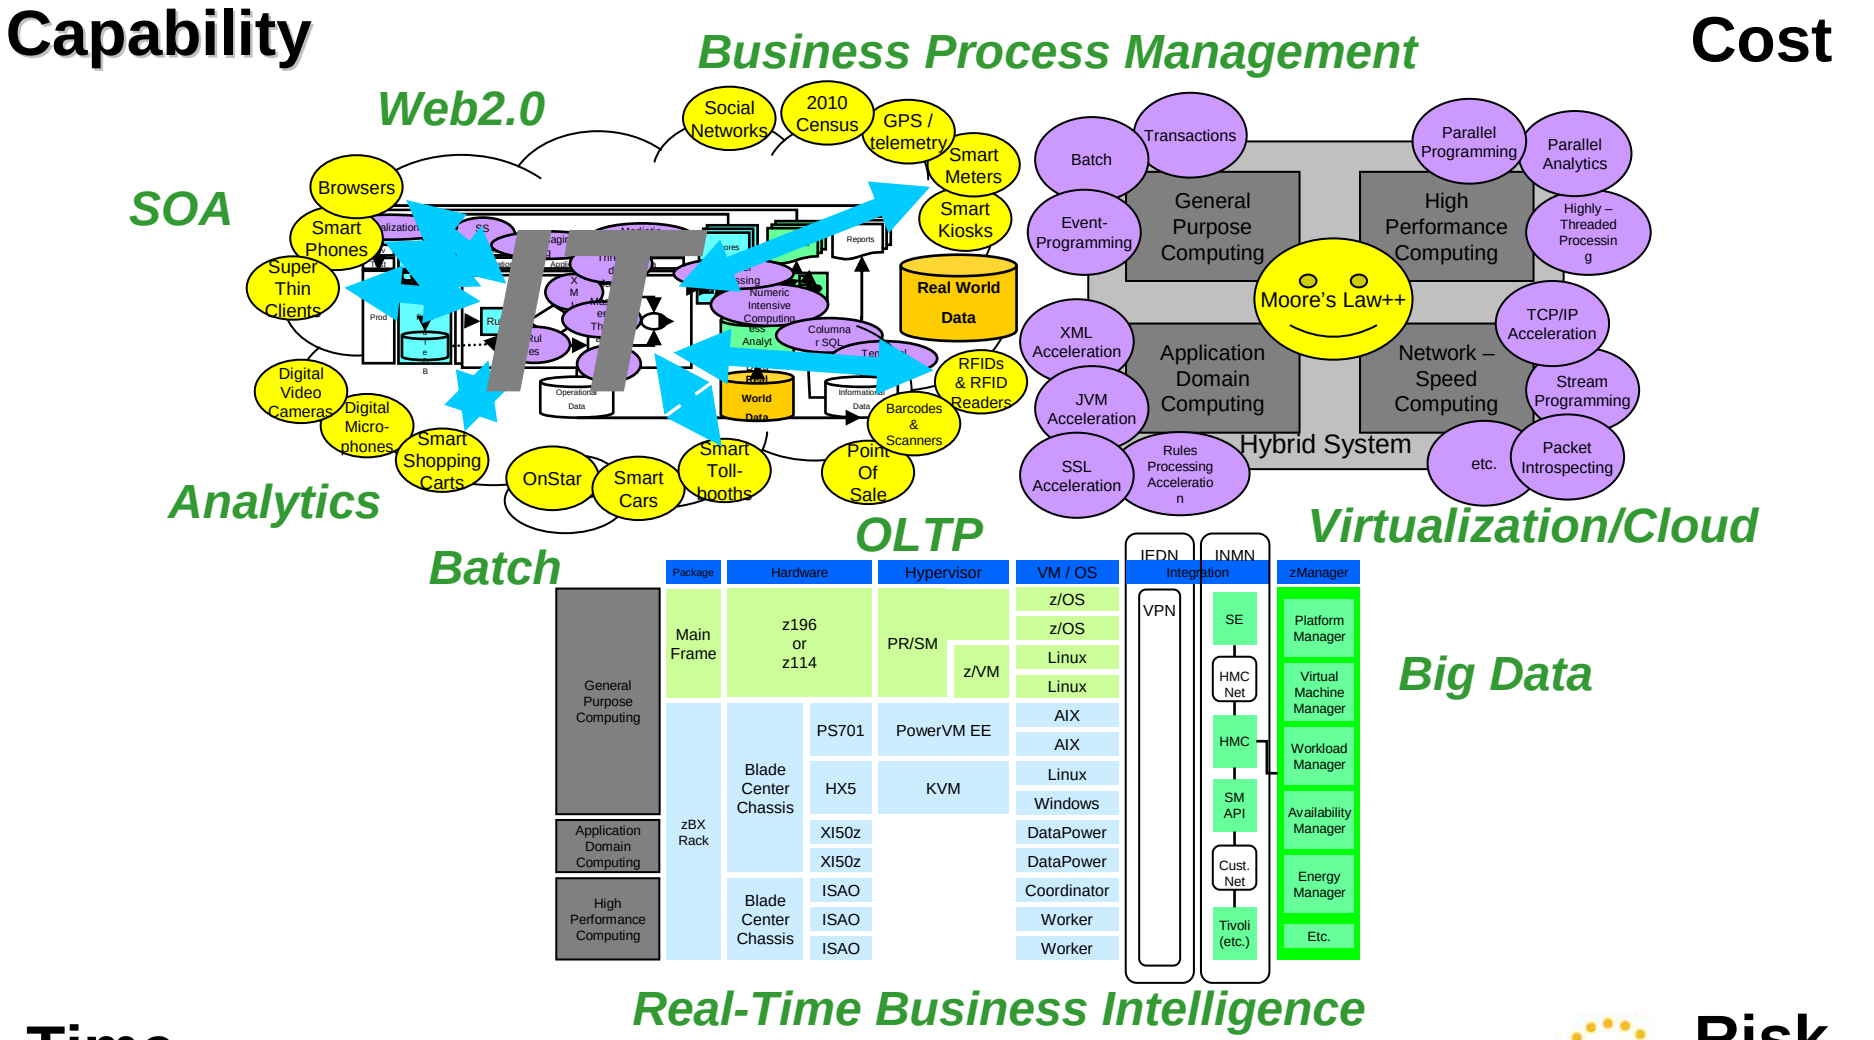

Topics

SHARE in Orlando 2011

Conclusion

- zEnterprise = (z196 or z114) + zBX + zURM
- Hybrid Computing
  - “optimized computing”
- The World’s Fastest “Master Thread”
  - and most highly resilient
- **Manageably** Harnessing Today’s (and Tomorrow’s) Most Advanced Computing Models
- Driving Moore’s Law to New Levels
  - Price AND Performance

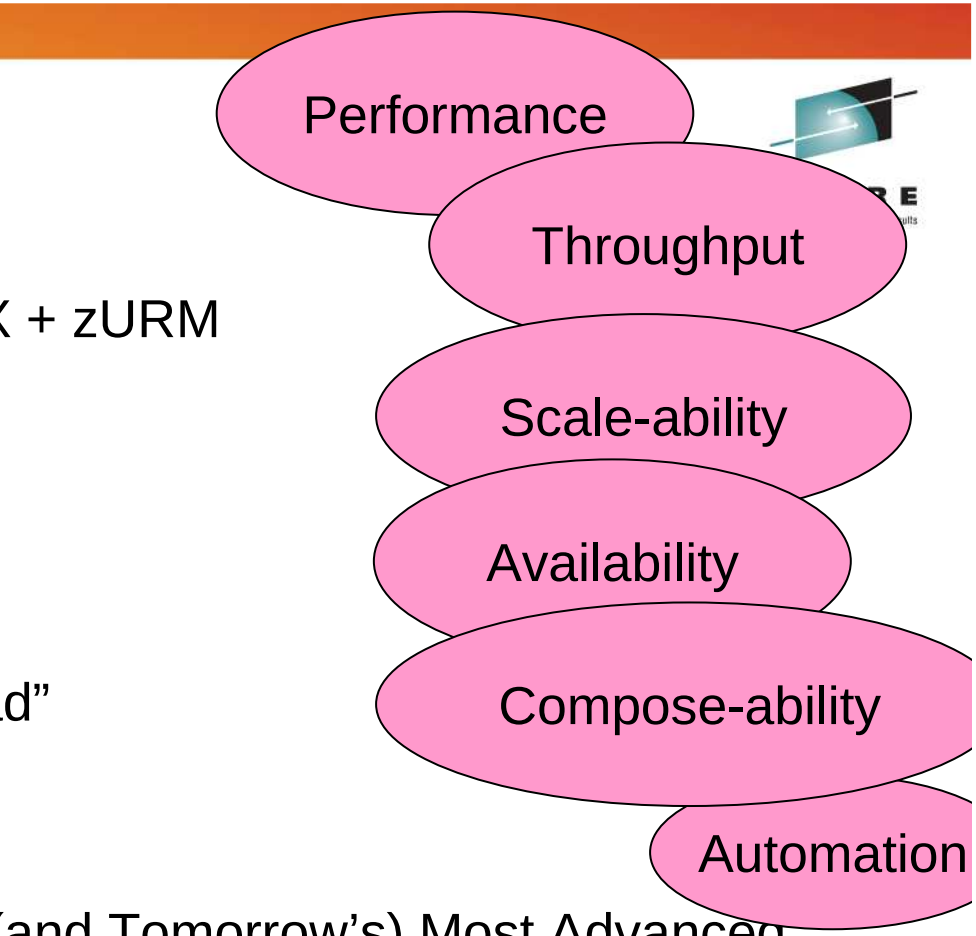

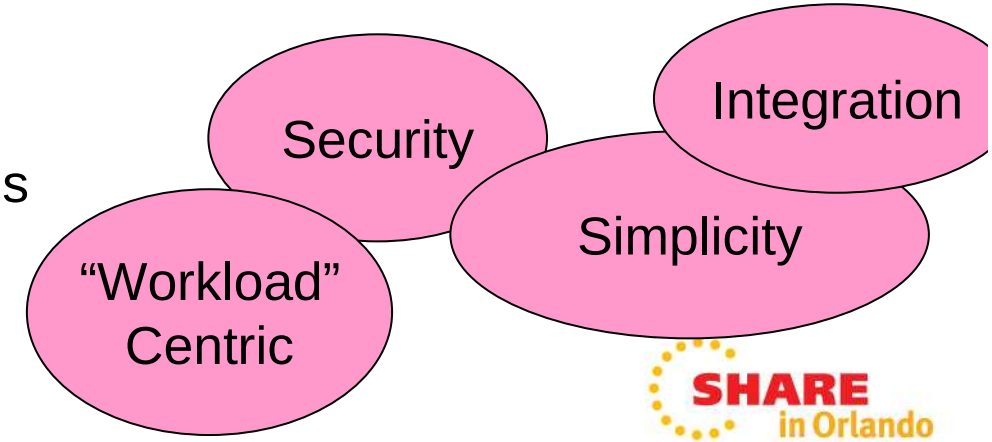

Topics

End of Section

Next Steps

Topics

IBM “Next Step” Workshops and Services

- For the Business Analysts
  - Process Improvement Workshop
- For the Architects
  - Architecture and Integration Workshop
  - Right-Fit Workshop
- For Platform Decision-Making
  - Fit for Purpose Workshop
- For Financial Analysis (TCO)
  - Right-fitting Applications into Consolidated Environments (RACE) Workshop
  - EAGLE Engagement
  - Scorpion Study*
  - Portfolio Review Analysis
- For Missourians (or the Missouri-like)
  - Proof of Concept / Proof of Technology / Benchmarking Centers
  - Design Center
- For the Truly Adventurous
  - IBM Research First of a Kind Joint Development Projects

* Typically provided as a “service” at cost-recovery rates

Topics

Thank You

Questions?

Comments?

Critiques?

Requests??!?!!

Monte Bauman

IBM Columbus

mbauman@us.ibm.com

Topics

Slide Title (Type Size=28) (no more than two lines)

- First Major Topic (Type Size=24)
  - Subtopic One (Type Size=22)
  - Subtopic Two (Type Size=22)
    - Sub-subtopic (Type Size=20)
- Second Major Topic (Type Size=24)
- Third Major Topic (Type Size=24)
- Fourth Major Topic (Type Size=24)

Topics

Two-Column Slide (Type Size=28)

- Topic A (Type Size=24)
  - Subtopic 1 (Type Size=22)
  - Subtopic 2 (Type Size=22)
  - Subtopic 3 (Type Size=22)
  - Subtopic 4 (Type Size=22)
- Topic B (Type Size=24)

- Topic C (Type Size=24)
  - Subtopic 1 (Type Size=22)
  - Subtopic 2 (Type Size=22)
  - Subtopic 3 (Type Size=22)
    - Sub-subtopic 1 (Type Size=20)
    - Sub-subtopic 2(Type Size=20)
- Topic D (Type Size=20)

Topics

Slide with Table

Topics

Slide with Text & Graphic

Topics