

# SAS on IBM Power Systems

---

**Camilla Rosendal Madsen**  
Cognitive Solutions  
IBM Systems DK

**Preben Jacobsen**  
HPC/AI Architect  
IBM Systems Nordics



# Agenda

- Experiences with SAS on cloud implementation
- Challenges and successes
- Private cloud vs Public cloud
  
- SAS and IBM partnership
- What does IBM infrastructure for SAS vs alternatives look like?
- TCO
- Why is IBM better for SAS?
  
- Q&A

# Customer Case

## - *Moving SAS to the Cloud*

- Customer Scenario
- The Challenge
- Why IBM?
- The Root Cause
- The Solution
- Current Status
- Others do it...

# Customer Scenario

Who: Large Pharmaceuticals Company

- Global reach
- Large SAS User
- New Implementation in isolated environment
- External Consultants for architecture and implementation
- Microsoft Consultants Extensive Assistance

# The Challenge

Customer wanted to deploy SAS on a public Cloud

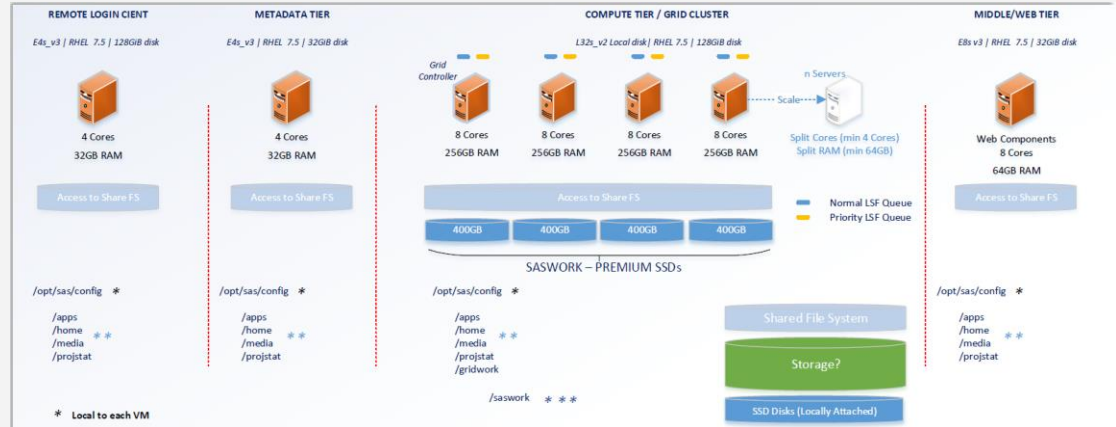
- Target was Microsoft Azure

- PoC

- Performance issues

- Availability

- Data Control



# Selecting Deployment Target

	Private Cloud	Hybrid	Public Cloud
Computational Workload Control	\$\$   ♥♥♥	\$\$   ♥♥	\$   ♥
Data Access Performance	\$\$   ♥♥♥	\$\$   ♥	\$   ♥
Data Preparation/Upload	\$   ♥♥♥	\$\$   ♥♥	\$   ♥♥
Data PostProces/Download	\$   ♥♥♥	\$\$   ♥♥	\$\$\$   ♥♥
Security / Compliance	\$   ♥♥♥	\$\$   ♥♥	\$\$\$   ♥♥
Housing	\$\$\$   ♥♥	\$   ♥♥♥	\$   ♥♥
Power / Cooling	\$\$\$   ♥♥	\$\$   ♥♥♥	\$   ♥♥
Network Connectivity	\$\$   ♥♥	\$   ♥♥♥	\$   ♥
Application Environment	\$\$\$   ♥♥	\$\$   ♥♥	\$   ♥♥
Availability and Service	\$\$   ♥♥♥	\$\$   ♥♥	\$   ♥♥
Data Protection (HA, DR, Backup)	\$\$\$   ♥♥	\$\$   ♥	\$   ♥

# Why IBM?

- Long Term Partnership w SAS
- Knows Application Deployment
- Knows Infrastructure and Data Interaction
- Has Track Record
  
- Has the 'medication'

# The Root Cause

## Infrastructure components

Compute Node Specification: - 4 nodes allocated and installed with Spectrum Scale Client

Size	vCPU	Memory (GiB)	Temp disk1 (GiB)	NVMe Disks <sup>2</sup>	NVMe Disk throughput <sup>3</sup> (Read IOPS / MBps)	Host Cache Size <sup>4</sup>	Max Data Disks	Max NICs / Expected network bandwidth (Mbps)
Standard_L32s_v2	32	256	320	4x1.92 TB	1.4M / 9,000 MBps	N/A	32	8 / 12,800

Storage Node Specification - 9 to be used – Spectrum Scale NSD node:

	vCPU	Memory: GiB	Temp storage (SSD) GiB	Max data disks	Max cached and temp storage throughput: IOPS / MBps (cache size in GiB)	Max uncached disk throughput: IOPS / MBps	Max NICs / Expected network bandwidth (Mbps)
Standard_E32s_v3 2	32	256	512	32	64,000 / 512 (800)	51,200 / 768	8 / 16,000

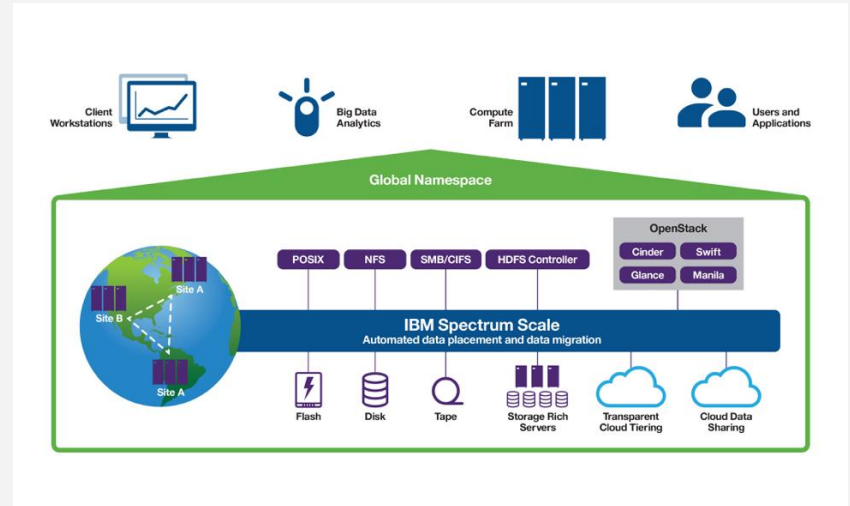



# The Solution

## Architecture design with Spectrum Scale



Aggregated Network Bandwidth – 12,8GB/s ~ 128 GbE



# Current Status

- PoC became a success
- High increase in performance
- Scalability with performance
- Data/storage control

Now moved to AWS....

- Pricing

# Others Do It

**Finance Data Warehouse: The New High-Performance Architecture**

**NVMesh**

**Combined Performance Initiative**

Application	Past Approach	Hardware
<b>Database Layer</b> One large dataset	Transaction Data	Monthly, Daily, Update Window
<b>Portion Layer</b> Business logic split	101, 102, 103	JBOD Storage Layer
<b>Parallelization</b> technical split	↓	↓
<b>Database Layer</b> One large dataset	Transaction Data	Global IP Deutsche Bank Network and Local Networks: IP 25 Gbit / InfiniBand 100 Gbit
<b>Sub-Portion Layer</b> Technical split	101, 102, 103, 180, 181	GPFS (Global Parallel File System) SAS Grid Solution ... scaling ...
		Hyper-Converged Cloud Solution including Excellero NVMesh

**Deutsche Bank**

**CASE STUDY**


**Benefits**

- High performance due to double parallelization
- High utilization based on distributed storage and CPU utilization
- Reduced cost compared to pure 1:1 hardware refresh

# 40+ years of partnership between SAS and IBM

IBM offerings are **co-optimized** for SAS workloads:

- SAS 9.4 for AIX on IBM Power Systems
- SAS Viya on Linux on IBM Power Systems
- IBM Spectrum LSF is the base for SAS Grid

A profile photograph of Ken Gahagan, a man with short dark hair, wearing a light blue collared shirt. He is looking towards the left of the frame.

Ken Gahagan  
Director, R&D  
SAS

“POWER9 has this high-throughput capability that other processors do not have.”

“[POWER9 is] a perfect fit for our machine learning and deep learning capability in SAS Viya”

Ken Gahagan  
Director, R&D  
SAS

[Video: Driving innovation with SAS® Software and IBM Power Systems](#)

# TCO / Value for x86 to Power

Scaling **up** vs. scaling **out**

- Other vendors use scale-out
- IBM can do both

Enterprise servers can be more interesting than scale-out when addressing SAS workload.

Beside the cost savings from choosing a Power platform rather than an x86-based platform comes that SAS users will be able to do **DOUBLE their performance and amount of work** with their SAS licenses on a Power platform

	Intel-based Environment (on-premise)	Intel-based Environment (Public Cloud)	IBM Environment (on- premise)
Servers	6	2	2 (Power E950)
Cores/vCPUs	336	192	80
Processors	12	N/A	8
Compute Performance Metric (CPM)	13,788 40% utilization	N/A	12,570 70% utilization
Total Memory (GB)	6,144	6,192	6,192
CPU Architecture	x86	x86	POWER9
Processor Speed (GHz)	2.7	2.3	3.4
Total 3-year costs (including HW, maintenance, OS, VMware/Power VM, Power/cooling cost)	<b>\$660,000</b>	<b>\$700,000</b>	<b>\$480,000</b>

# SAS on Power TCO

*Power Systems vs x86 On-Premises*

<b>6 Intel-based Servers</b>	<b>VS</b>	<b>2 Power E950 Servers</b>
<b>3 Year TCO Reduction</b>		<b>Expense Reduction</b>
\$180,000		27%

Savings from:

- Reduced HW maintenance and SW support costs
- Reduced power and cooling costs
- Reduced virtualization licensing costs
- Reduced FTE (full time employee) for operation costs

# Reliability

## Power Systems vs x86 On-Premises

**16x to 28x less downtime  
compared to alternatives**

**Cost of Downtime: \$4,998 per server/minute<sup>1</sup>**

**Downtime per year:**

**IBM Power Systems: 1.54 minutes/server per year<sup>1</sup>**

**Cost of Downtime Estimates:**

→ **\$7,696.92 per server/year**


**X86-based Infrastructure for SAS High-Performance Visual  
Analytics: 26-43 minutes/server per year**

→ **\$129,948 - \$214,914 per server/year**

1. ITIC 2020 Global Server Hardware, Server OS Reliability Report

# SAS on Power TCO

*Power Systems On-Premises vs x86 Cloud*



Savings from:

- Network costs
- Outbound bandwidth fees
- Storage costs



# SAS on Power vs Cloud

## Power Systems vs x86 Cloud

Faster insights at lower costs

### SAS on Power Systems

*On-Premises or Private Cloud*

- Easily achieve I/O bandwidth required for SAS workloads
- Enjoy cloud-like pricing and on-demand scalability with Capacity on Demand
- Flexibility to tune deployment to achieve best possible performance

### x86 Cloud

- I/O throughput is crucial and a limiting factor for successful SAS public cloud deployments
- Most public clouds cannot reach the SAS recommended minimum I/O throughput of 100 MB/s per core
- Limited HW and Storage tuning in public cloud Cost in movement of data

## **██████ sees COVID-19, hard-to-control public cloud costs driving more on-premises engagements**

“The costs are so much harder to control [in public cloud],” he said. “You blink, and you pay something to a hyperscaler. You move data for two seconds or you turn a bit on a hard disk with a hyperscaler, and you pay for it. If you have a very steady, even usage, [on-premises is a good option].”

██████ who took helm of the cloud business at ██████ 15 months ago, said he saw the public cloud conundrum first-hand when he was heading up infrastructure and data management services for US\$13.6 billion systems integration behemoth Atos. In one case, a customer moving to the public cloud received a 6,500-page bill. “They had no idea what was in there,” he said.

# Q&A

# Contact information

## IBM SAS Alliance team

—

Camilla Rosendal Madsen

[Camilla.rosendal.madsen@ibm.com](mailto:Camilla.rosendal.madsen@ibm.com)

+45 41 20 37 43

Preben Jacobsen

[prebenj@dk.ibm.com](mailto:prebenj@dk.ibm.com)

+45 40 60 15 89