



Moving Analytics Workloads to the Cloud: Agility, Speed and Scale with Managed Costs

DEEPAK RAMANATHAN

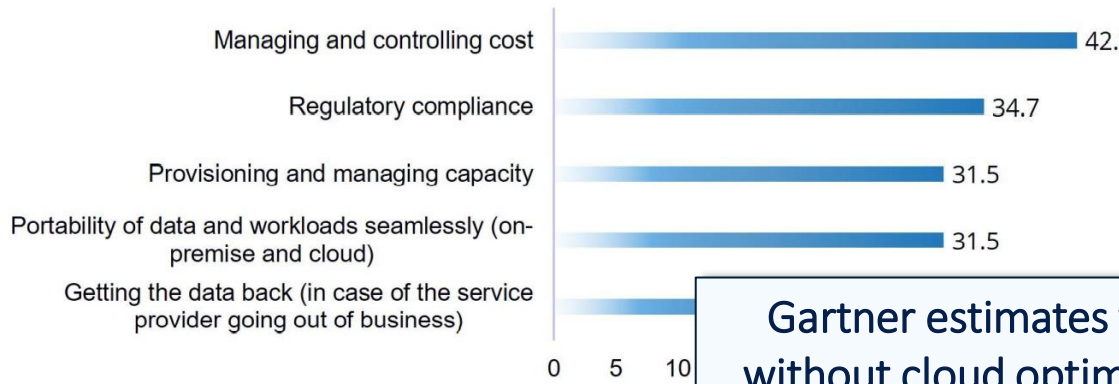
Vice President,
Global Technology Practice, SAS

SAS
INNOVATE

#SASInnovate



Thinking about your current infrastructure, what are the most pressing multi-cloud data management priorities?



Source: IDC European Multicloud Infrastructure Survey, March 2018 (n = 651)

Gartner estimates without cloud optimization, organizations will overspend

Why?


1. Ungoverned costs
2. Unanticipated usage
3. No commitments
4. Dev/Test waste
5. Too much headroom
6. Wongsizing production
7. Suboptimal design & implementation

80% of Customers Report Cloud Repatriation Activities

More customers expect to repatriate workloads next year

Public Cloud Repatriation Rates

Q. In the last year, has your organization migrated any applications or data that were primarily part of a public cloud environment to a private cloud or on-premises environment?



n = 400
Source: IDC's Cloud and AI Adoption Survey, January 2018

Percent of Public Applications Expected to Repatriate Over the Next Two Years (Average)

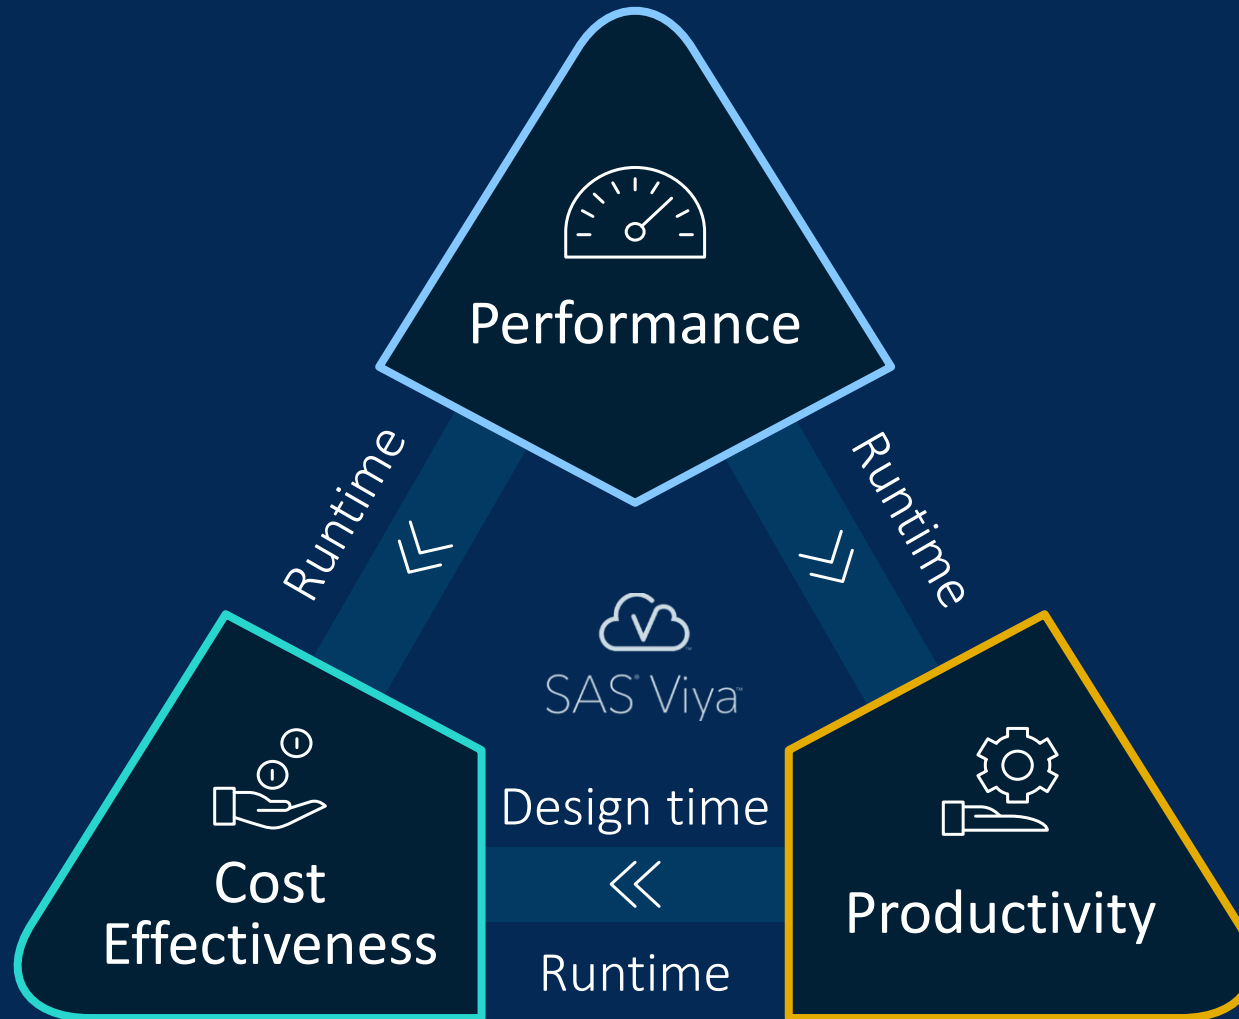
Q. Using your best guess, what proportion of the public cloud applications installed today will move to a private cloud, hosted private cloud or non-cloud environment over the next two years?

50%

Top Repatriation Drivers

Security	19%
Performance	14%
Cost	12%
Control	12%
Centralize/Reduce Shadow IT	11%

Cost Effectiveness in the Cloud



Cloud Cost Optimization Approach



01

Observability

*analyze based on
SAS usage behavior and
resource usage patterns*



02

**Prioritization &
Optimization**

*minimize resource contention
and ensure optimal use of
cloud resources*



03

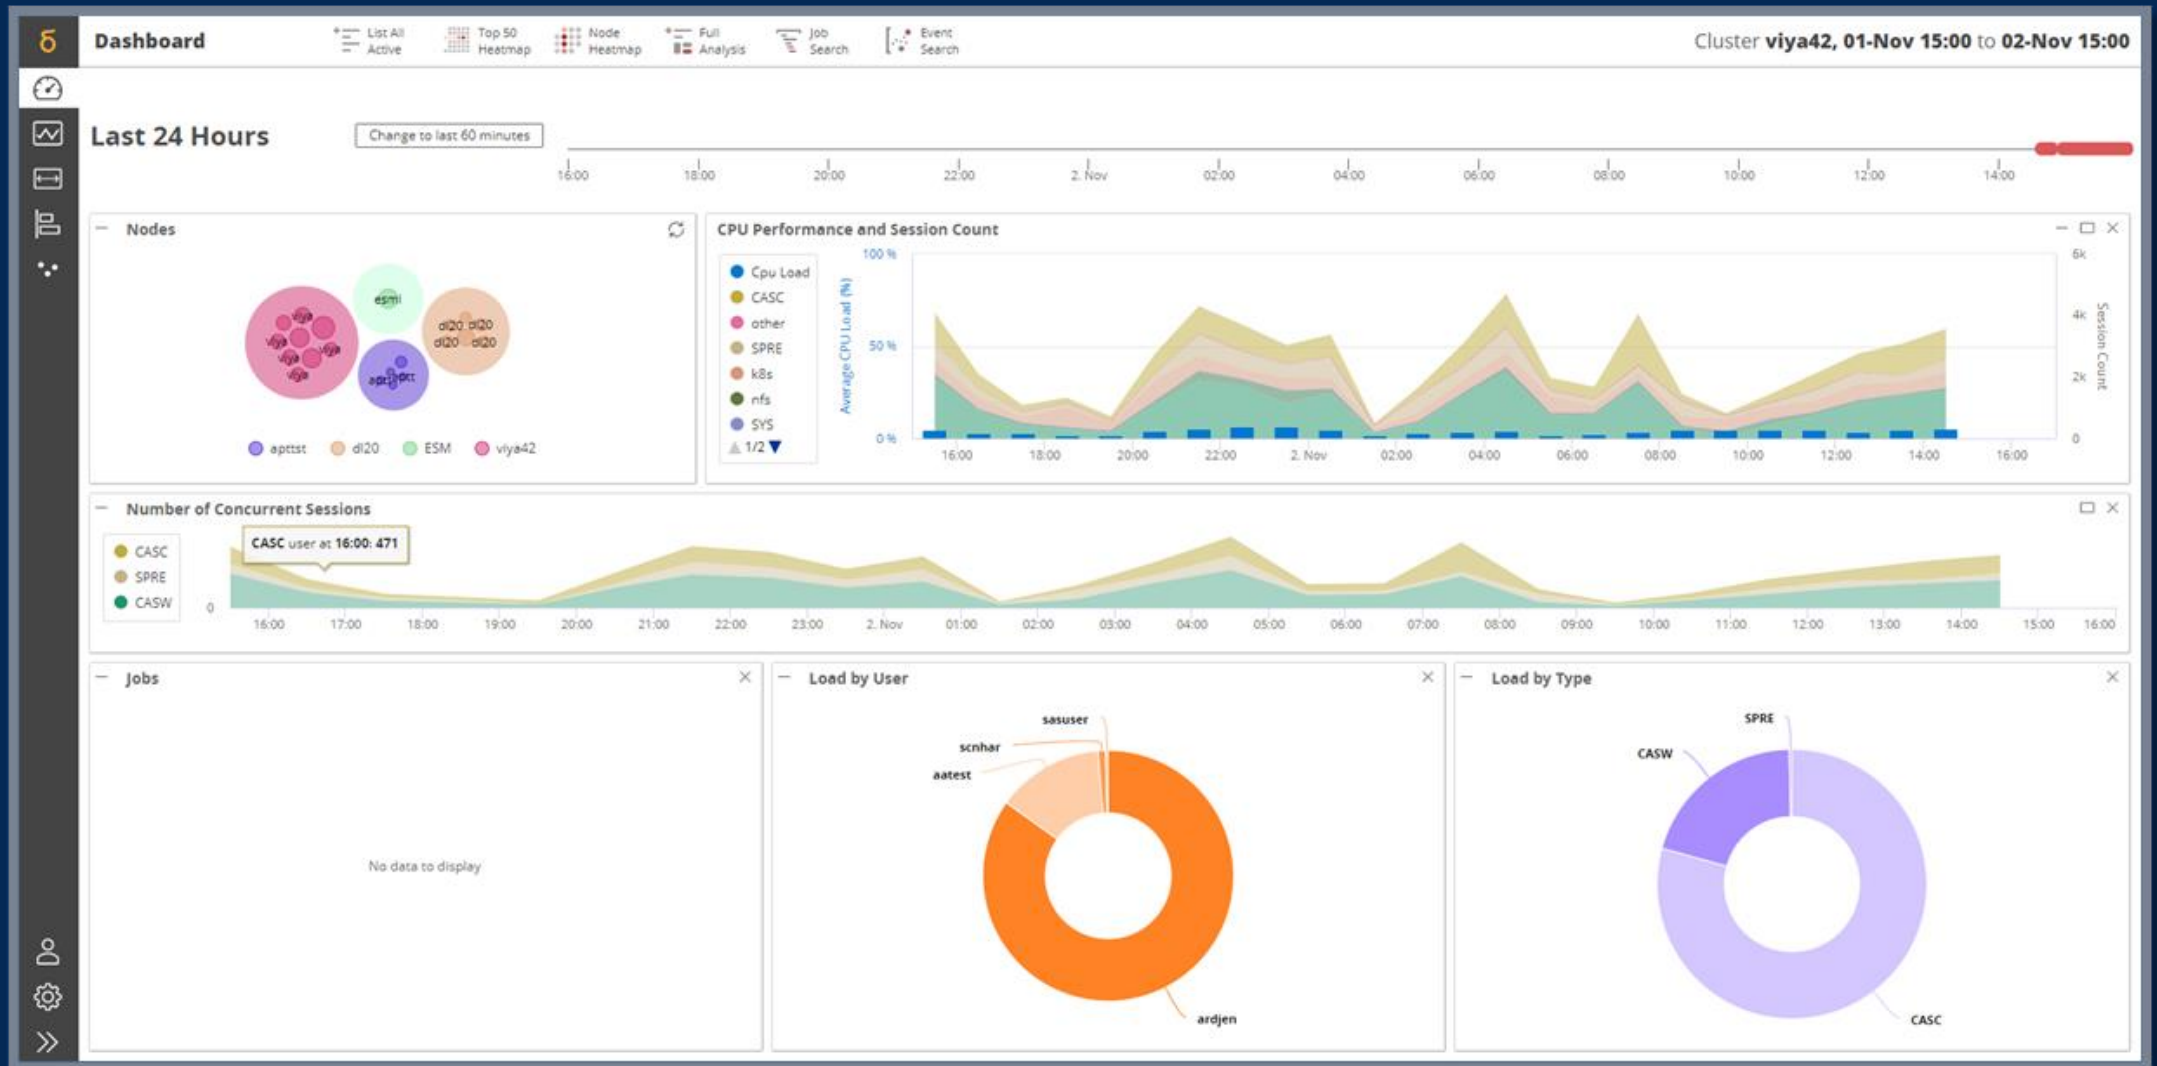
Autoscale

*automatically add/remove
compute instances based on
real-time metrics*

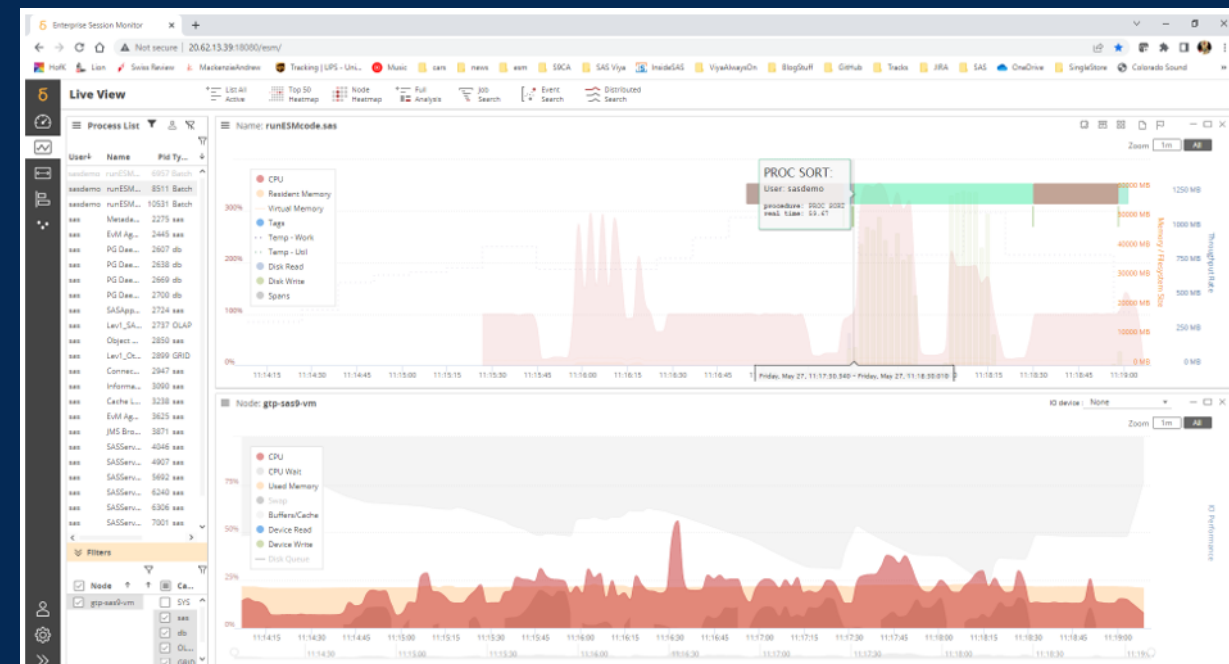
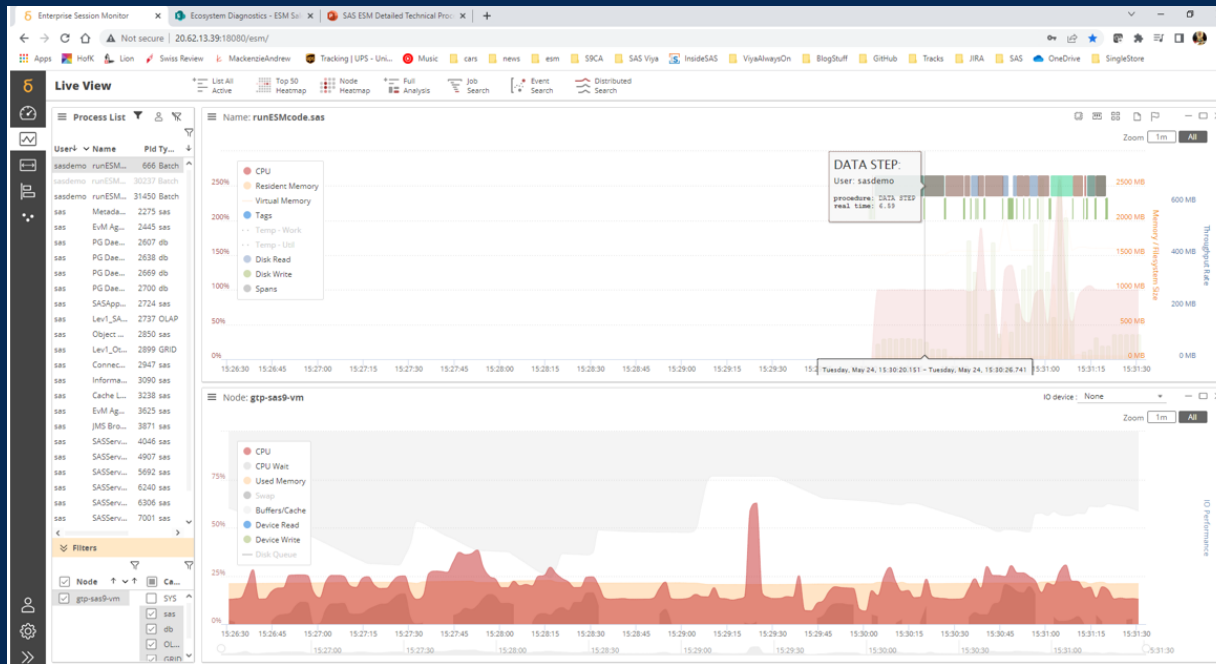
A black and white photograph of a woman with voluminous curly hair, looking upwards and to the right. She is holding a microphone in her right hand, which is raised towards her head. The background is blurred, suggesting an indoor setting with other people. The image is partially obscured by a dark blue diagonal overlay on the right side.

Observability

Monitoring Analytics Consumption



Monitor Users and Batch Jobs



Infra Consumption By User Profile

	Data Analyst	Business Analyst	Citizen Data Scientist	Advanced Data Scientist*	Total
Cluster Consumption Summarized	High Job Count, High Runtime, High Consumption	Low Job Count, Low Runtime, Low Consumption	Moderate Job Count, Moderate Runtime, Moderate Consumption	Low Job Count, Moderate Runtime, High Consumption	
Frequency (# users, % total)	97 (5.0%)	1,052 (53.7%)	680 (34.7%)	129 (6.6%)	1,958
Avg analytic procs per user per session	1.0	1.4	1.9	2.7	2
Sessions per day	21.98	0.49	1.15	0.63	1.80
Runtime hours per day	106.82	3.78	14.88	12.52	13.31
Active cpu hrs per session	32.77	0.71	3.31	4.28	3.44
Jobs per day	16.96	0.13	0.27	0.13	1.01
Workspace=1, Batch=2	1.8	1.0	1.1	1.1	1.1
Count distinct job name	4,042	321	709	65	5,137
Count job name	95,408	8,068	10,593	943	115,012
Avg max mem (mb)	2,024	295	1,606	4,285	1,094
Avg peak cpu (100 = 1 core)	95	39	109	218	78
Sum session runtime hours	600,947	230,482	586,698	91,467	1,509,594

Consumption



Manage Workloads

Know and Understand



SAS Workload Management

Effective workload management – Beyond Autoscaling

Job Overflow
may cause unplanned
costs and delays.

Poor Utilization
gaps of unused infra that
could fit more jobs.

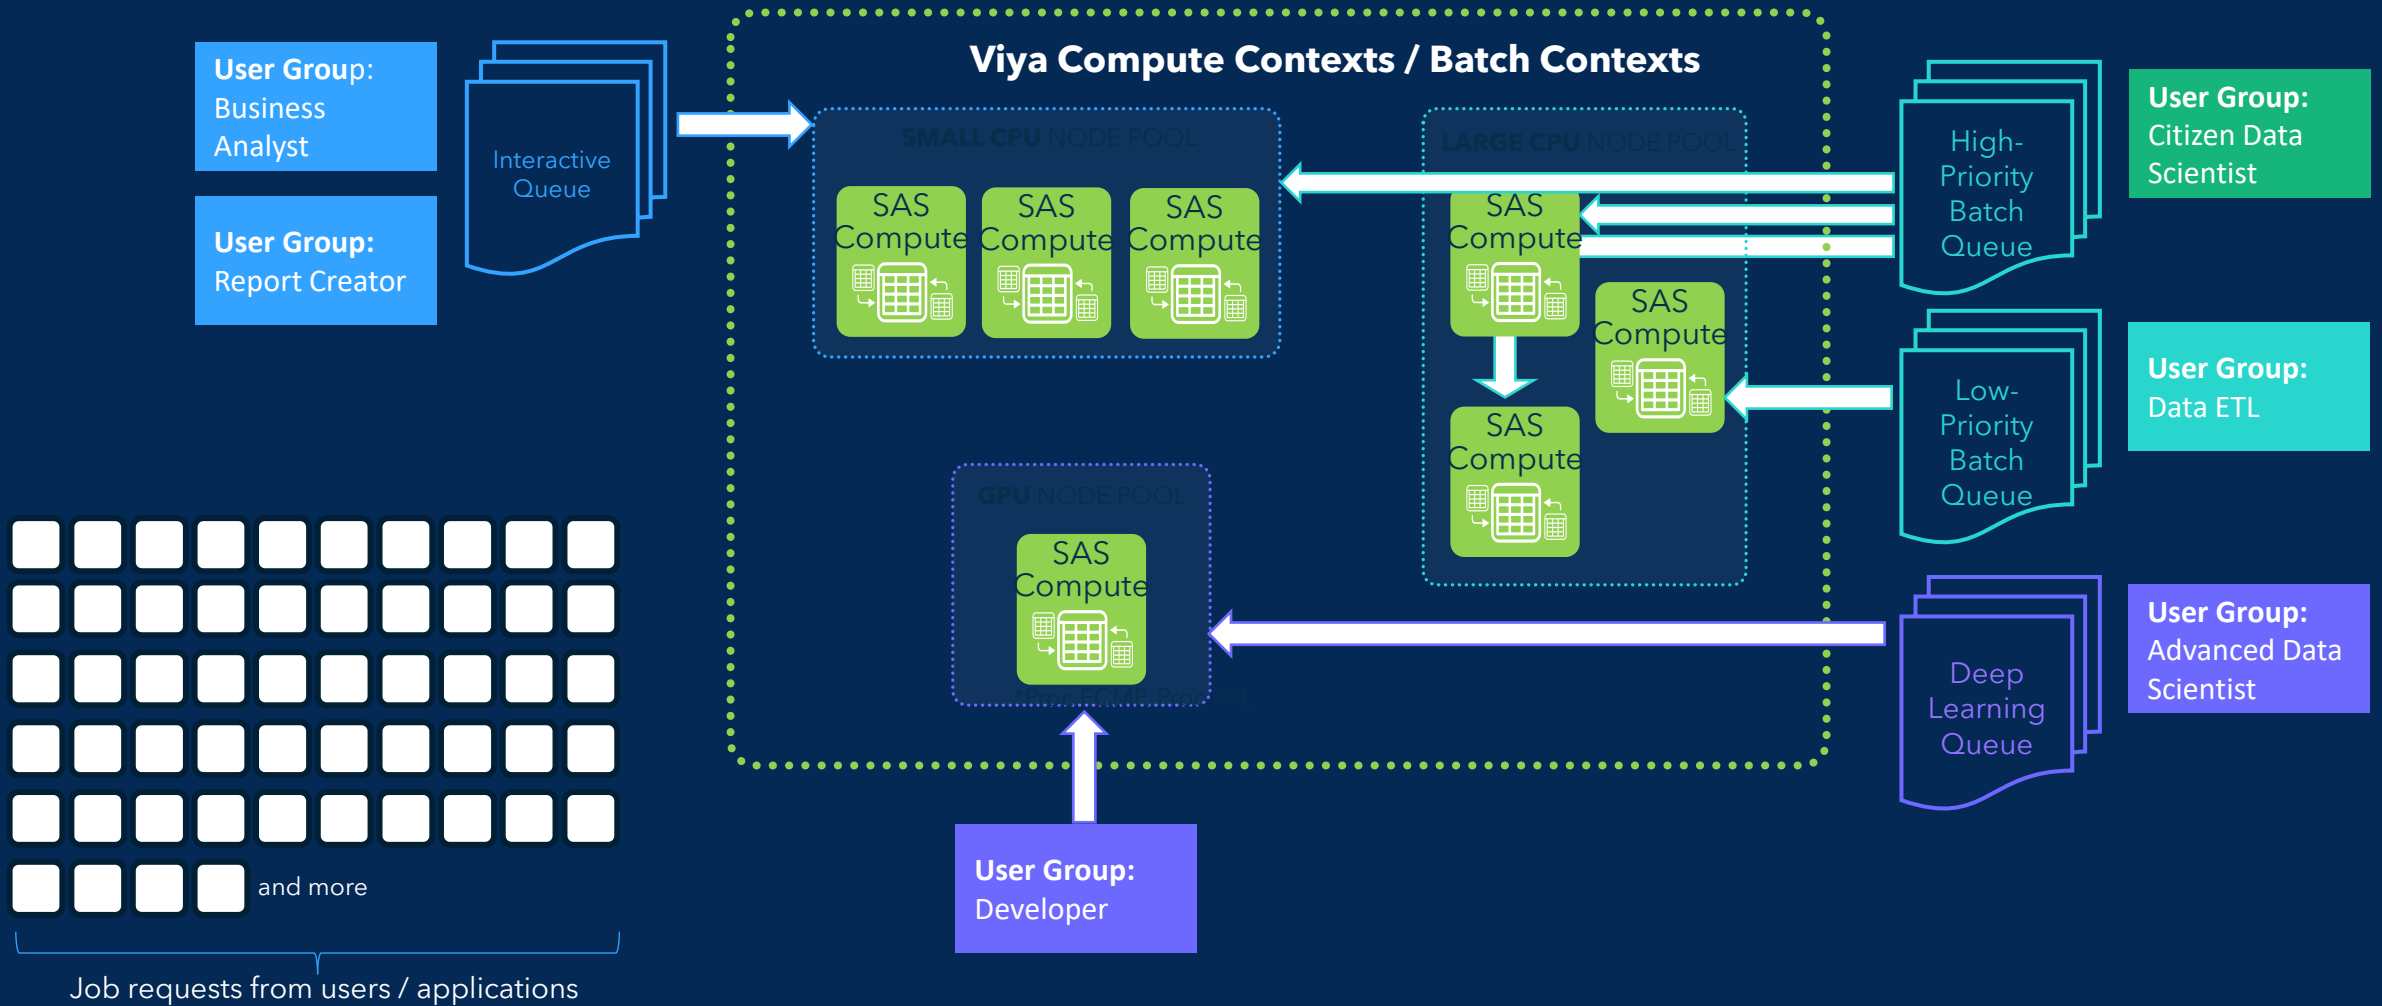
Missed Jobs
may cause major
disruptions and penalties
for missed deadlines



SAS WLM keeps your
compute needs
working smoothly.


Optimal infra spend
to fulfil workload
needs.

WLM Prioritization / Workload Optimization

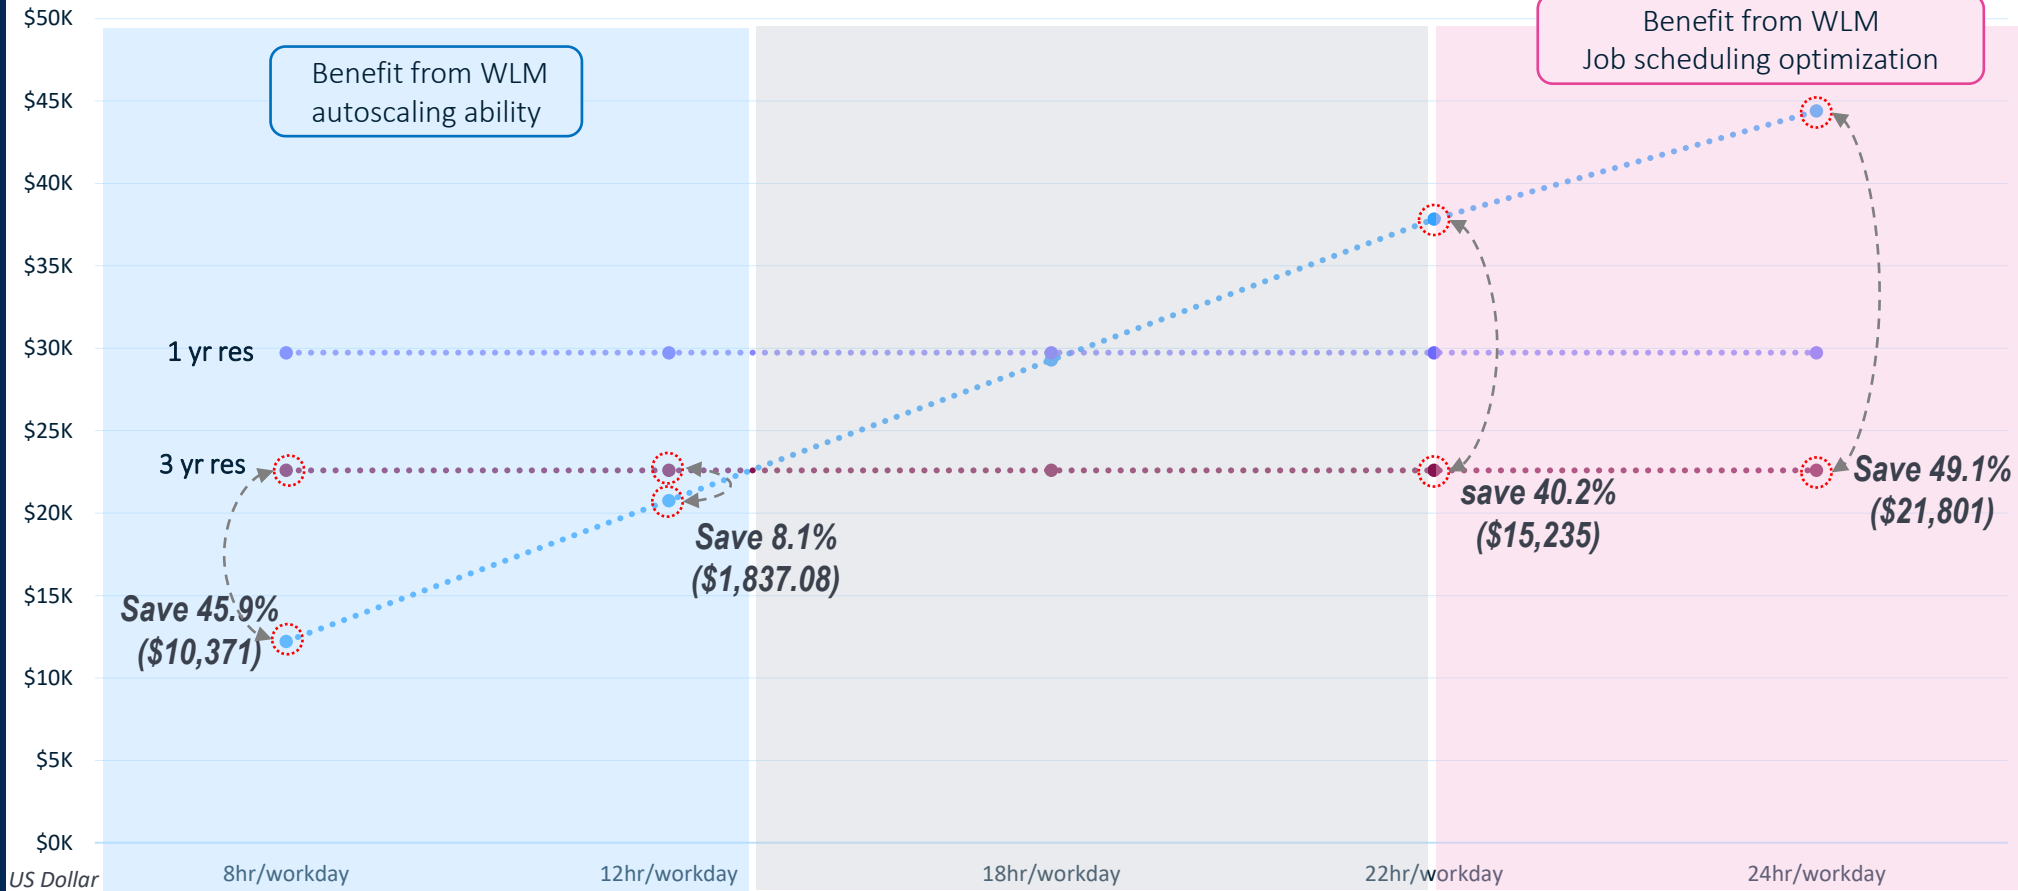


3-year pricing comparison

ex: 5 small compute nodes (E4_v5,128GB standard disk)



SAS Viya4 on Azure AKS

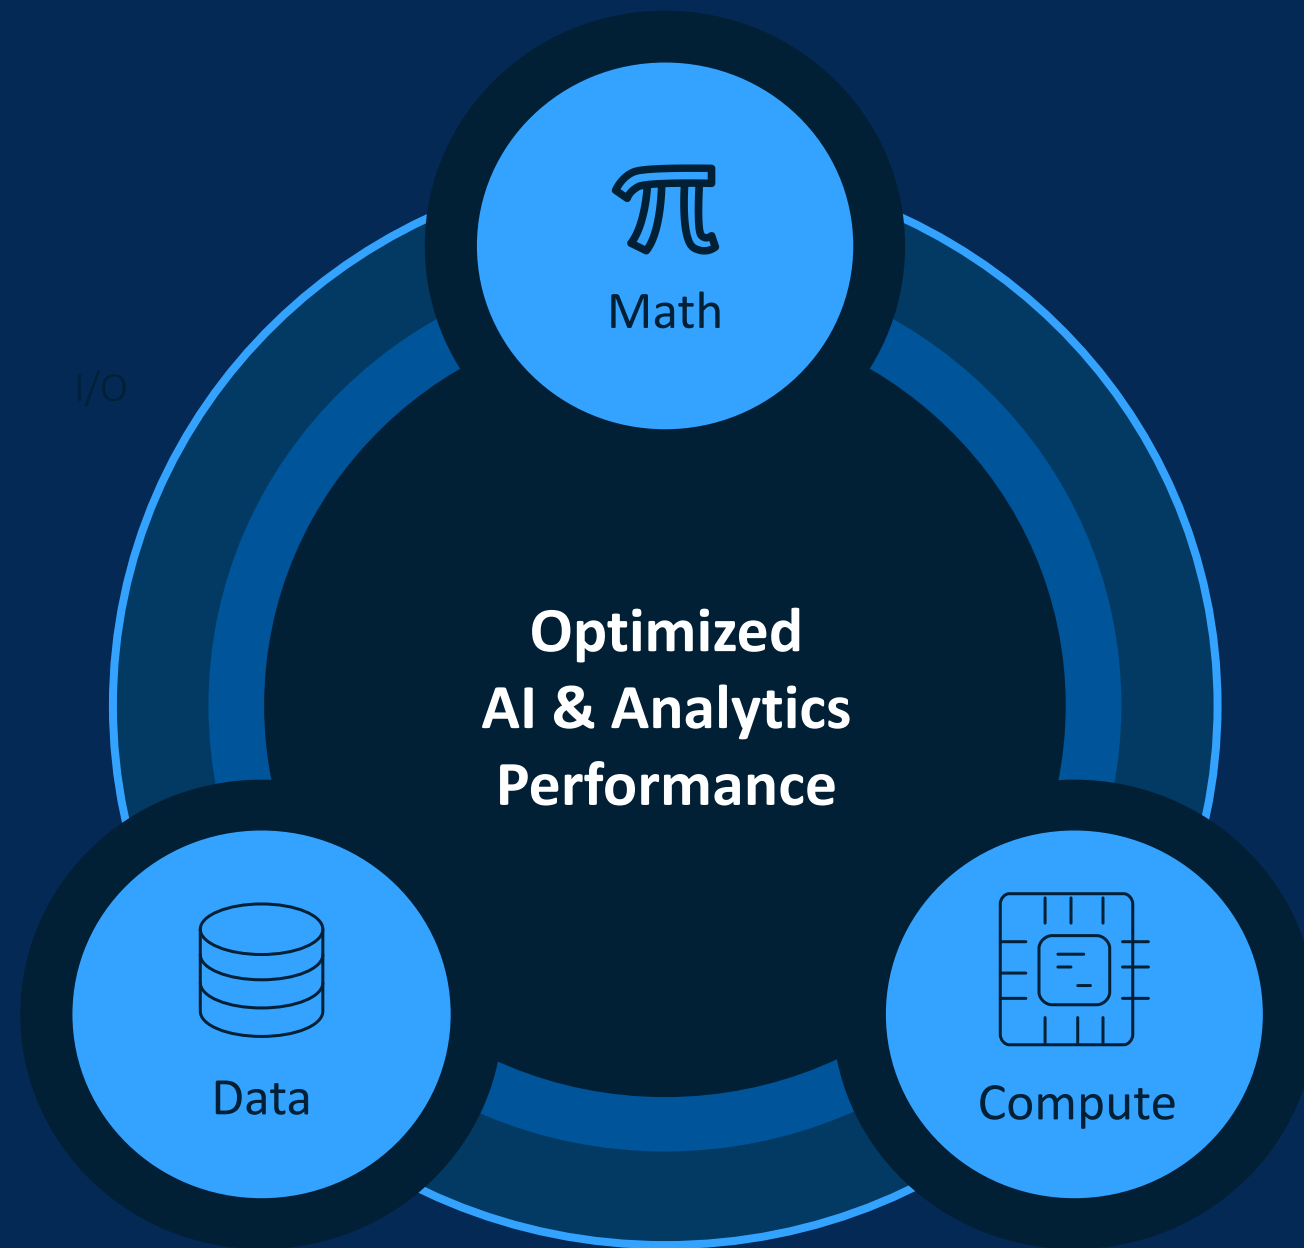


Benefit from WLM
Job scheduling optimization

- Understand workloads, prioritize job queues, right size infrastructure with Workload Manager
- Take Viya4 on Azure AKS as example (with 5 small compute nodes at most for 3 years usage), customer can save \$21,801 USD !!



Performance



SAS Viya's Speed

30x

faster than
**Open Source and
Commercial Offerings**

49x

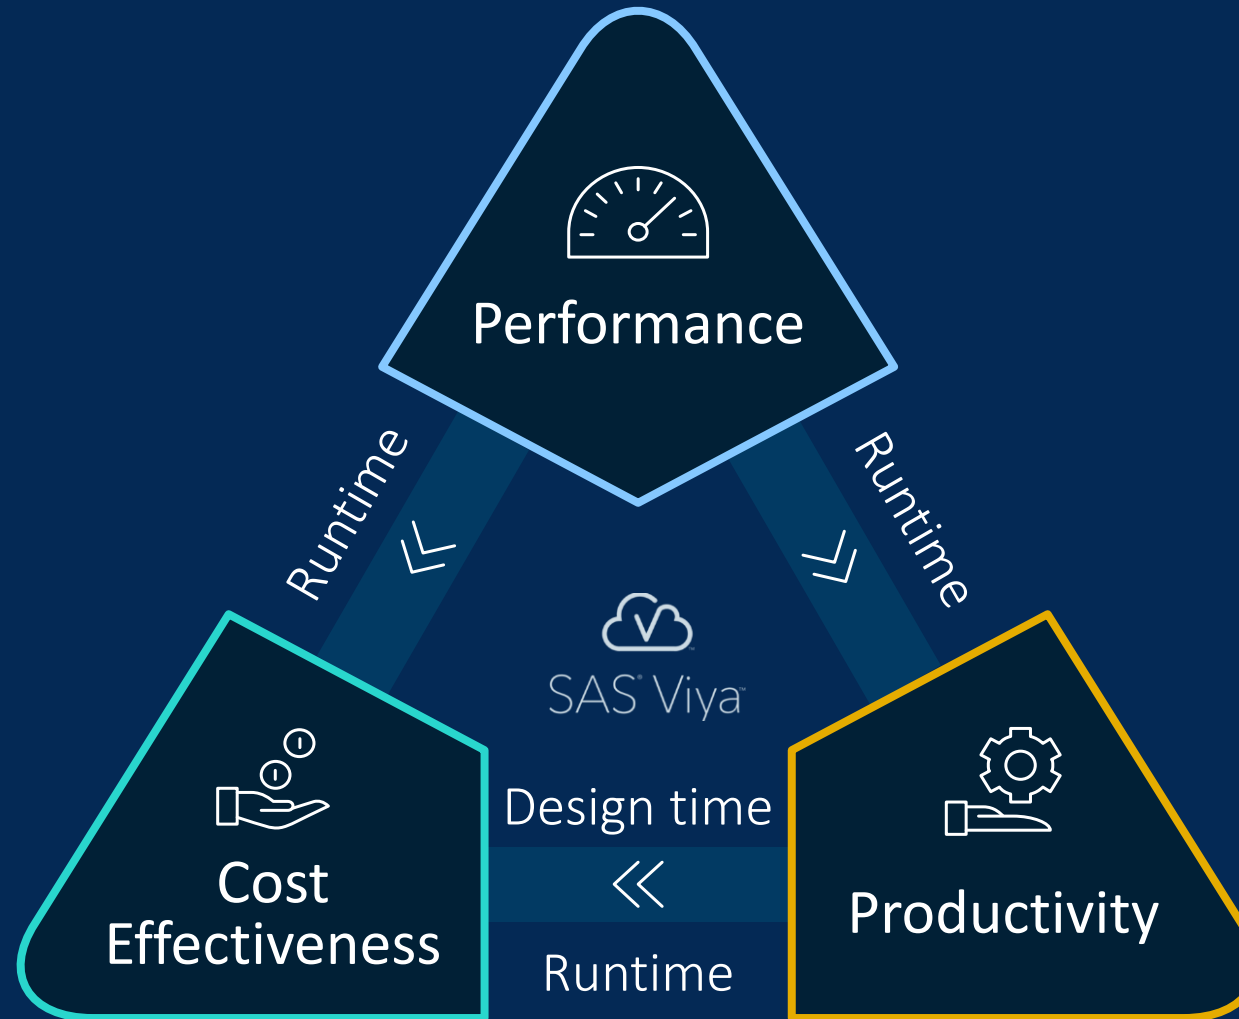
faster than
**Commercial Data
and AI Platform**

326x

faster developing
**Complex Models than
Commercial Data and
AI Platform**



Cost Effectiveness in the Cloud



Thank You