



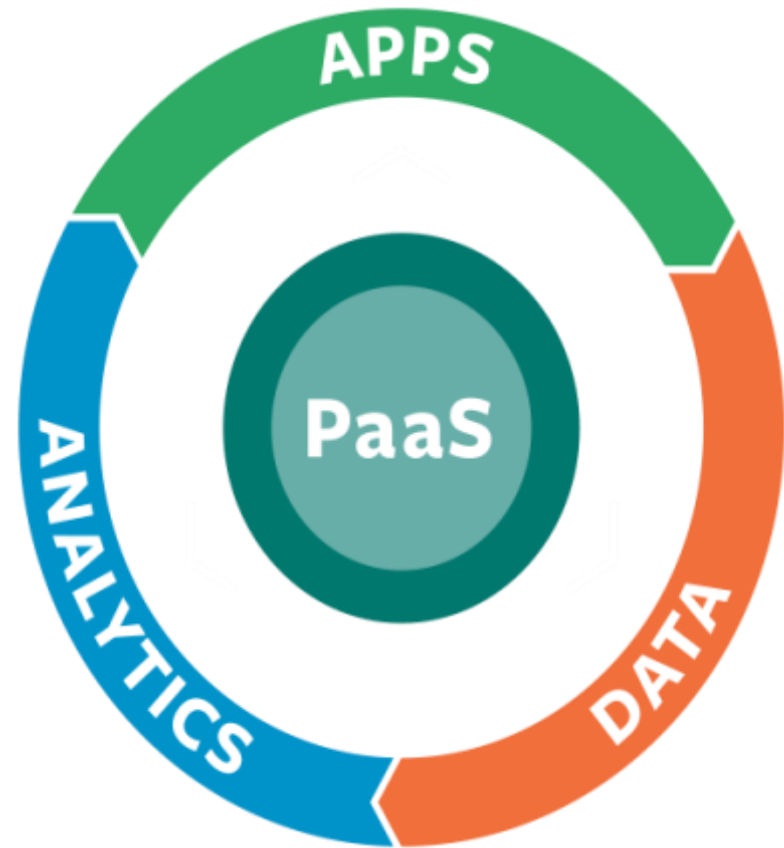
The future of Unstructured Data Workloads

#OFADevWorkshop
Tasneem Maistry Pivotal



What Matters: Apps. Data. Analytics.

- ✓ *Apps power businesses, and those apps generate data*
- ✓ *Analytic insights from that data drive new app functionality, which in-turn drives new data*
- ✓ *The faster you can move around that cycle, the faster you learn, innovate & pull away from the competition*



Pivotal's Opportunity

- 

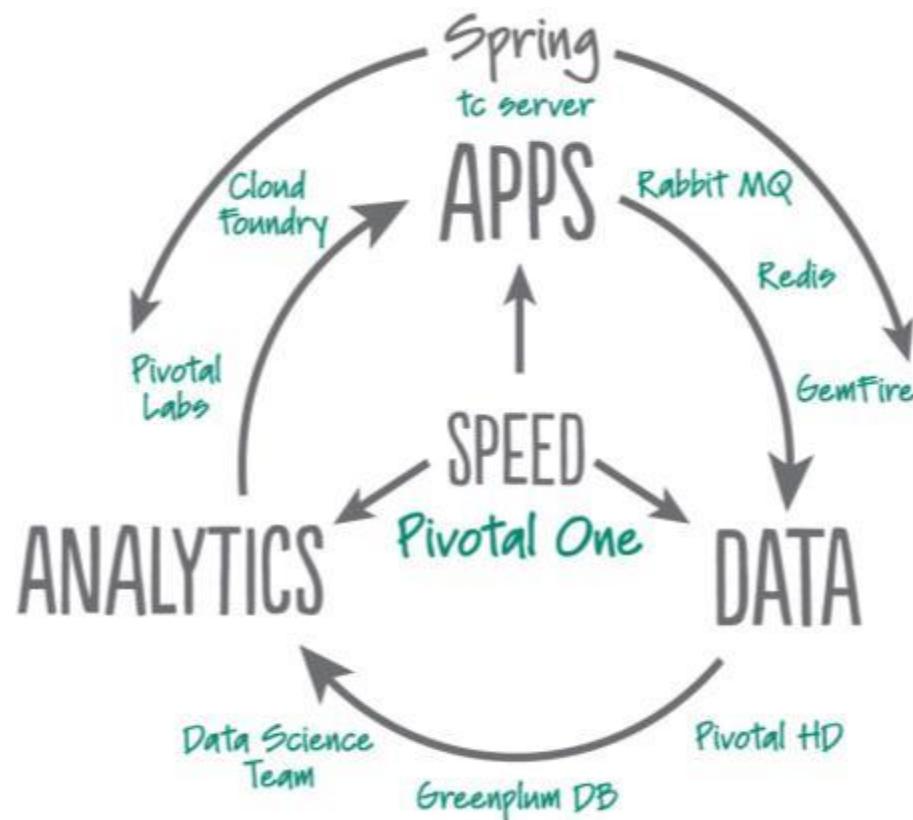
Uniquely positioned to help enterprises modernize each facet of this cycle today

- 

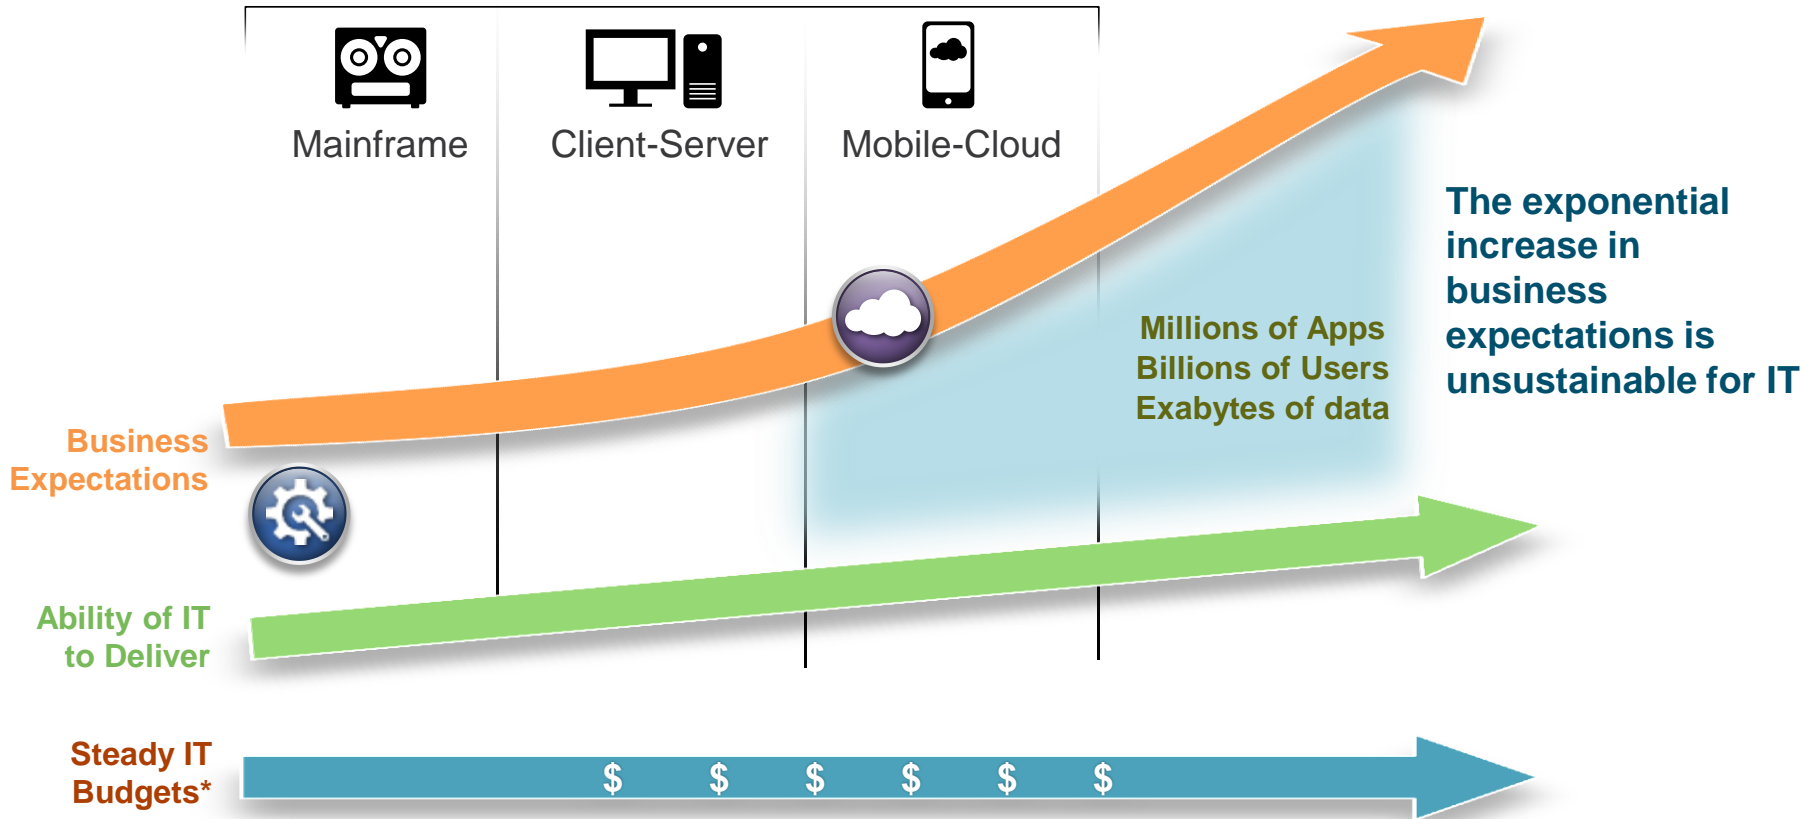
Comprehensive portfolio of products spanning Apps, Data & Analytics

- 

Converging these technologies into a coherent, next-gen Enterprise PaaS platform



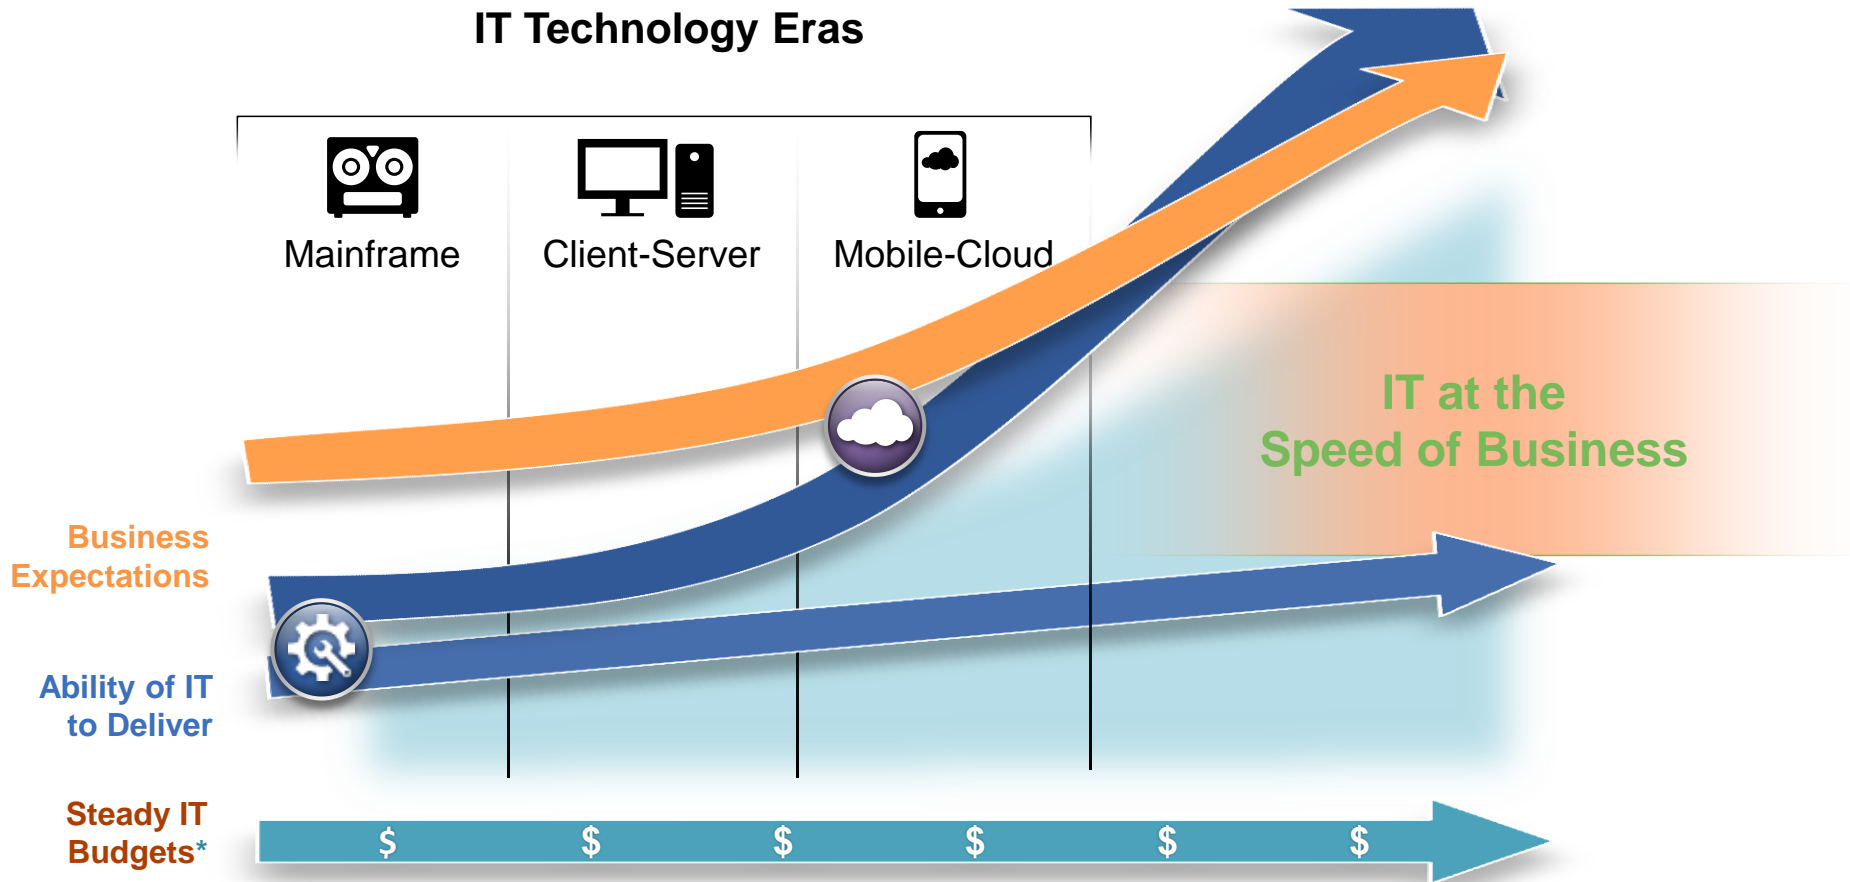
IT Technology Eras



* Source: Gartner, 2013: "Hunting and Harvesting in a Digital World: The 2013 CIO Agenda"

Pivotal Can Help Bridge the Gap

IT Technology Eras



* Source: Gartner, 2013: "Hunting and Harvesting in a Digital World: The 2013 CIO Agenda"

IT is Constrained by Technology Silos

72% of IT operating budget spent on maintaining old technologies

Legacy Silos



Modern Silos



Liberating Resources Serves Customers

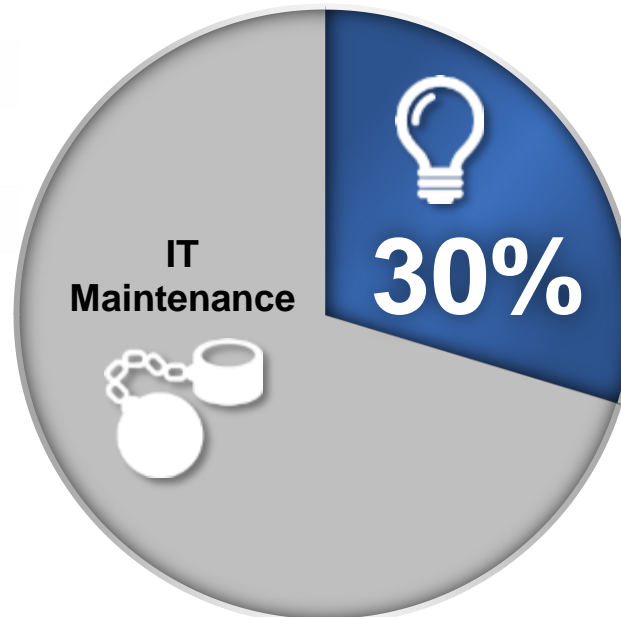
Data isn't available.

Analysis takes too long.

App dev takes too long.

Provisioning
takes weeks.

Production workloads
are managed via email.



Industry Average:
**IT Spend
on Innovation**

* Source: Gartner, 2013: "Hunting and Harvesting in a Digital World: The 2013 CIO Agenda"

Liberating Resources Serves Customers

Any user, any device,
any app... automatically.

Apps get rolled out at
“the speed of business.”

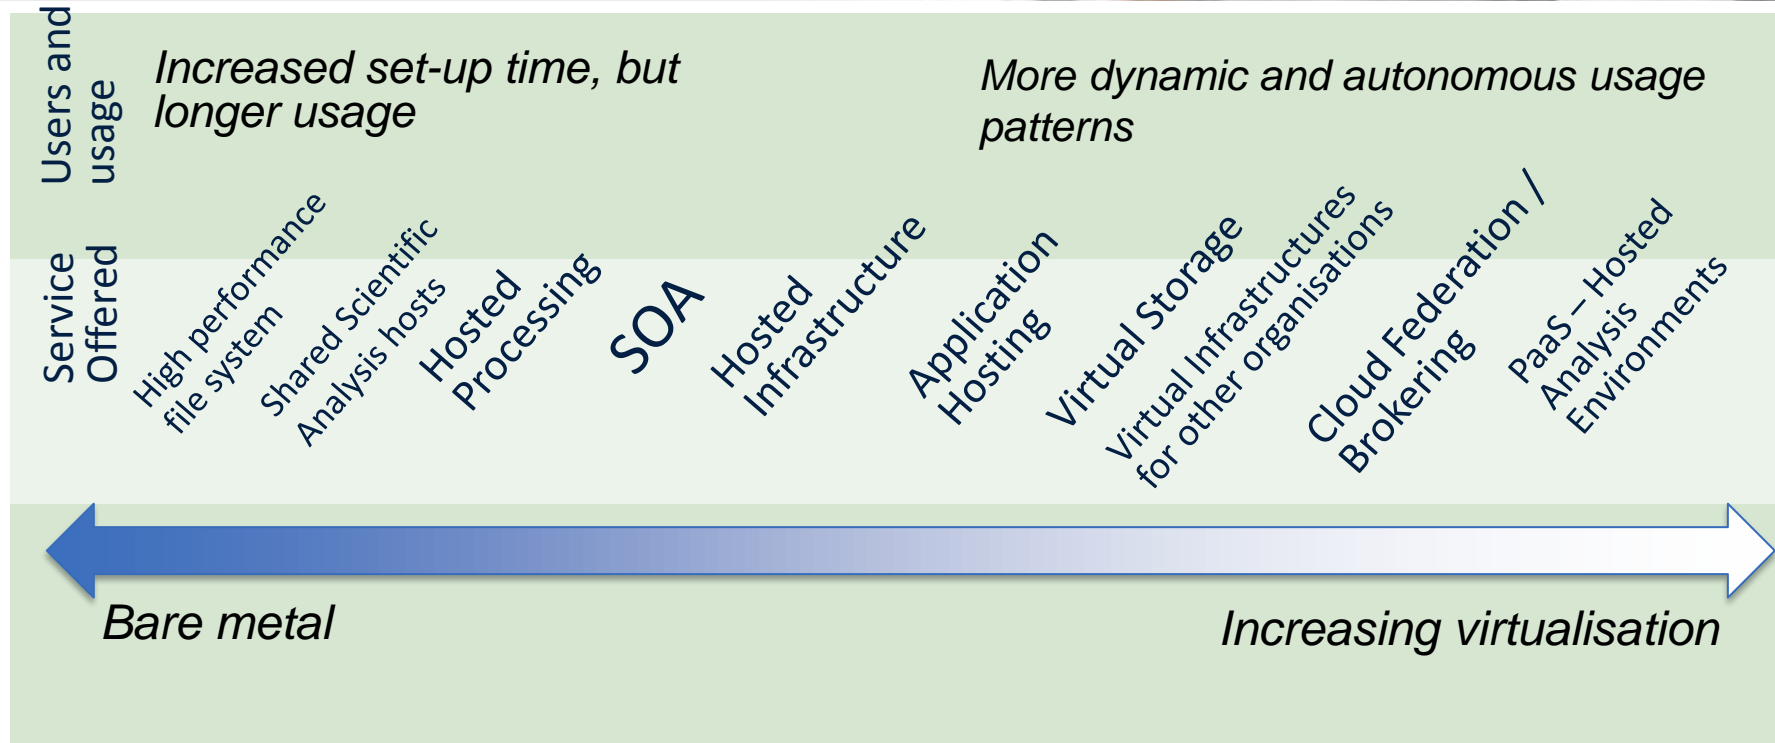
Provisioning a
production environment
takes minutes.

You can routinely deploy
any workload...anywhere.




Goal:
**IT Spend on
Innovation**

Interfaces and Usage Patterns vs. Hosting Solutions

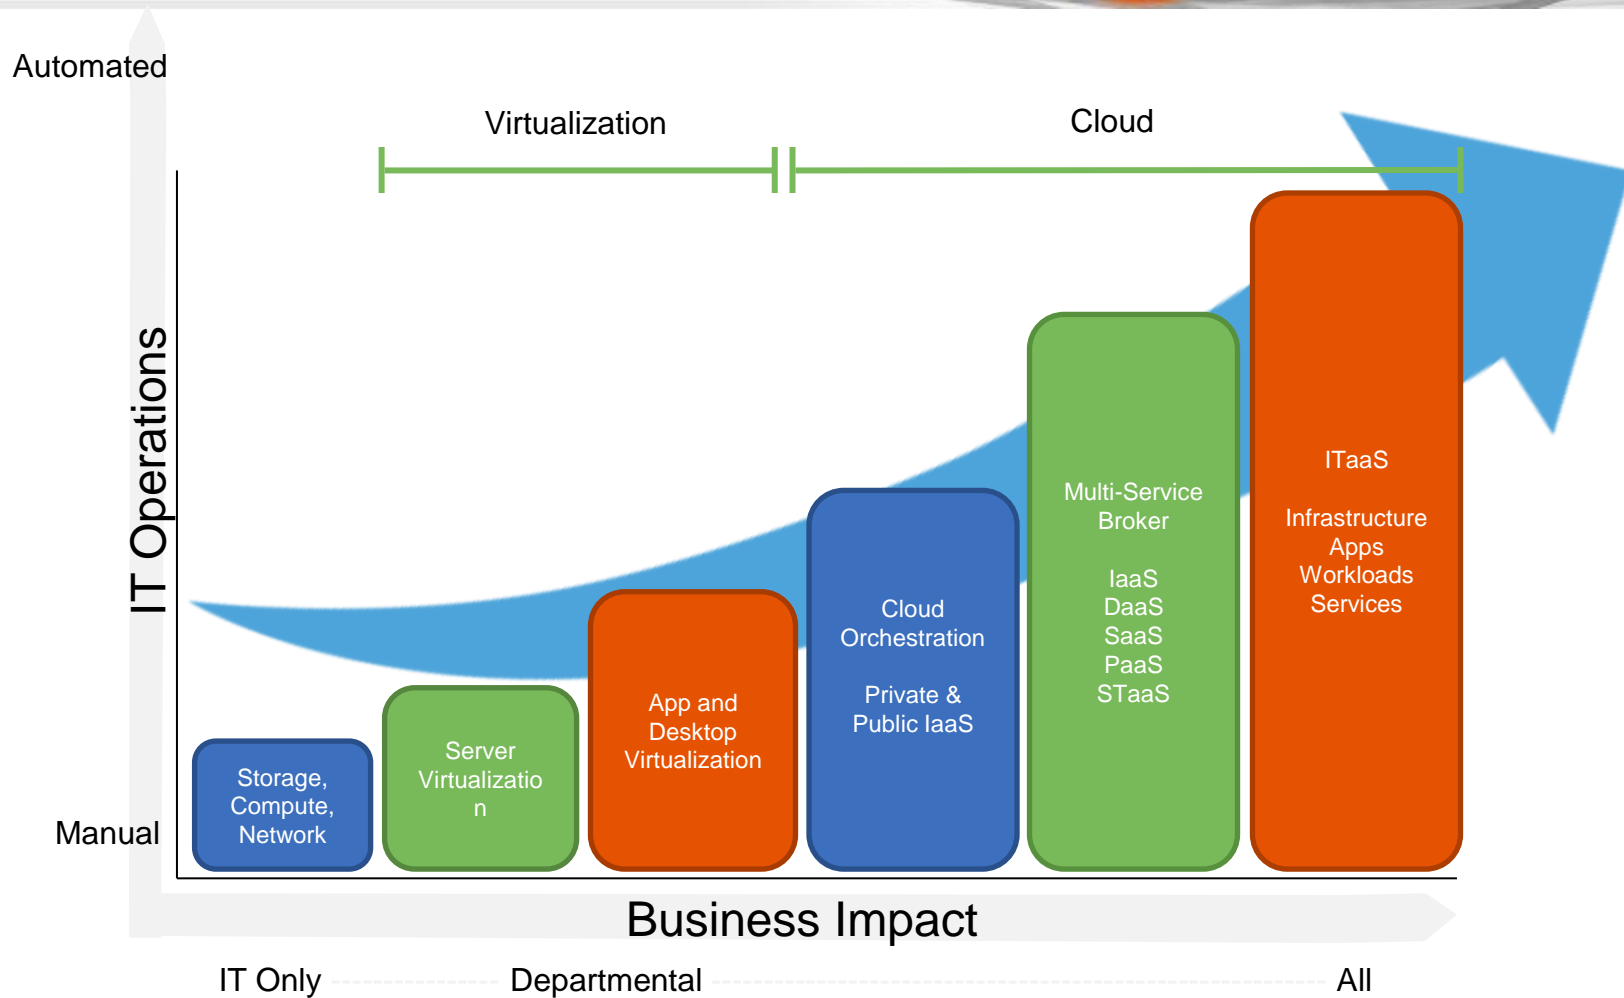


The Need for an Application Centric Layer

- ✓ Developers can focus on development and not infrastructure plumbing
- ✓ Separate the concerns of AppDev and Operations
- ✓ Eliminate the bottleneck of provisioning and deployment processes



Build your cloud for today and tomorrow



Why is this relevant and important

- Sensitivity analysis of each platform component on the outcomes (QoS, analysis etc)
 - What does it need .. Logs, users, processing, and processes
 - What is missing -- processes to capture and collect these
 - What is missing -- people who can analyze this and come up with learnings
 - What can AWB provide -- lots, users, processing
 - **NEED VOLUNTEERS TO** do research on the cluster

Pivotal Analytics Workbench (AWB)



- 1000 node cluster
- Collaborative project with industry leaders (Mellanox, Intel, Seagate)
- Provides mixed mode environment (GPDB, GemXD, PCF, PHD)
- Contains entire Hadoop stack (HDFS, HBASE, PIG, HIVE)

Partners



- Intel
 - Contributed 2,000 6-core CPUs.
- Mellanox
 - Contributed >1,000 network cards and 72 switches.
- Micron
 - Contributed 6,000 8GB DRAM modules.
- Seagate
 - Contributed 12,000 2TB Drives

Cluster



Cluster Size

- Physical Hosts - More than 1,000 nodes
- Processors - Over 24,000 CPU's
- RAM – Over 48TB for memory
- Disk capacity – More than 24PB of raw storage.
 - “Equivalent to nearly half of the entire written works of mankind from the beginning of recorded history”

Typical Target Use Cases

- Pivotal demonstration
- Partner engagements
(Intel, Mellanox, Seagate)
- Industry and academia collaboration
(Alpine, Informatica)



INFORMATICA



MAD



lib

Submitting your Proposal



- Request invite @ www.analyticsworkbench.com
- Project determination made within 2 weeks of submission

Q & A





Thank You



#OFADevWorkshop