Speaker Panel

Alex Kim, Vicom Infinity
Vincent Terrone, Vicom Infinity

Patrick Kelly, Privakey

Jay Desai, Xtreme Data

Dr Carlos Caicedo, Syracuse University

Panel Moderator

Len Santalucia, Vicom Infinity

Vicom Infinity Voice Assistant

- Business problem

Voice interface gives you the freedom of not touching a device. Current consumer solutions only provide an option of storing your conversation/voice data in a public cloud where you don't own the security controls nor meet the strong requirements for the enterprise.

- Solution

Developed with maximum security in mind, Vicom Infinity Voice Assistant will store your conversation on the IBM Hyper Protect Virtual Server (HPVS) protected by FIPS 140-2 Level 4 HSM device and using IBM Watson, the secure API Gateway for your Enterprise Applications, will enable most applications to be conversational.

- Secure & Enterprise-ready Voice Assistant gives freedom of processing your business conversation securely using Voice User Interface, powered by IBM Hyper Protect Virtual Server, IBM Watson, and Privakey for biometric authentication.
- Our Natural Language Processor API gateway as well as Privakey CX server run on HPVS which makes using voice assistant app for enterprises easy and more secure at a very competitive cost.

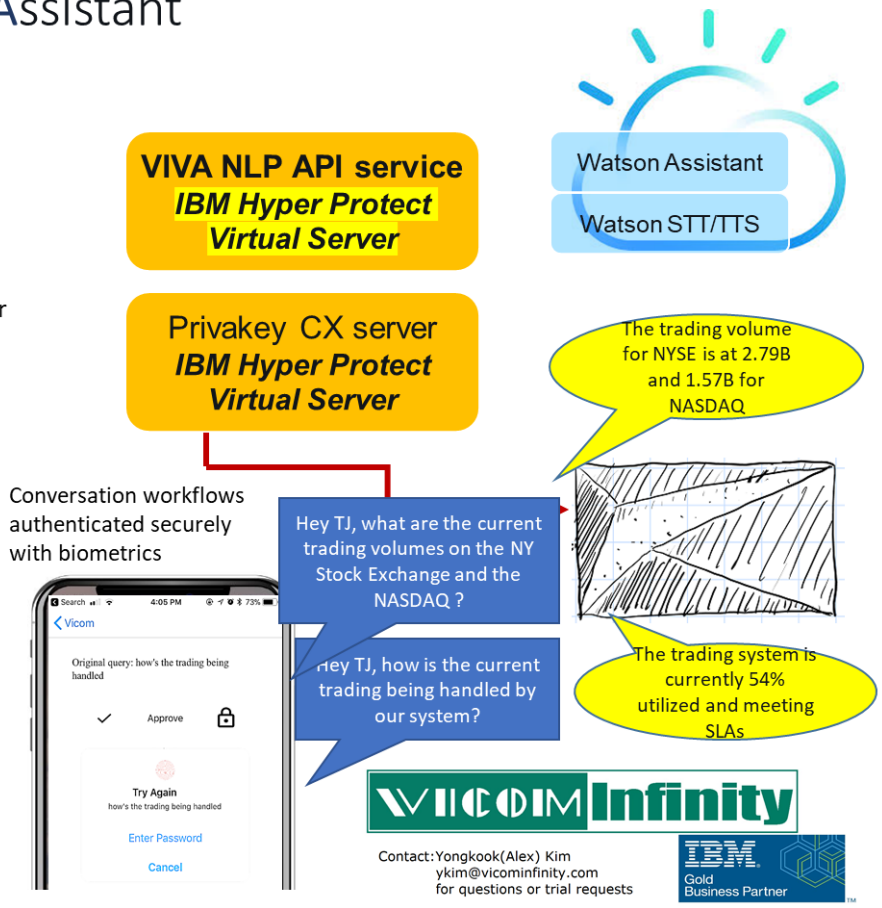

Big Idea: Help enterprise systems clients to use voice user Interface(VUI) for their enterprise applications in a completely private and secure way.

Bob, Managing Director
TBG Corp's
Data Center Ops Grp

As a user, I want to understand health and capacity of our systems by voice, as well as monthly snapshot of estimated HW/SW charges

Hey TJ, what is the current CPU Utilization?

Tom, Sr. BUE
TBG Financials
Financial Advisor Group

As a user, I want to manage my customer groups asset status and daily changes via voice, securely and privately

Hey TJ, how is Dow Jones Index today?

John, IT Manager
TBG Enterprise
Storage Systems Group

As a user, I want to manage and understand overall storage systems QoS via voice command and get reports to my phone

Hey TJ, how much storage space is left?

Dr. Jones, Medical Doctor
TBG Medical Center

As a user, I want to get a patient's medical information via voice securely and verified by my identity

Hey TJ, send me the blood pressure chart


DELIVERING DIGITAL TRUST THROUGH TRANSACTION INTENT VERIFICATION
Securing high value exchanges between a service and its users with human approval
privaKey

Privakey & Vicom Infinity presented at VOICE Summit

July 2019

Building a Secure Voice Solution (VIVA) for Enterprise Applications in Financial Services and Healthcare

PRIVAKEY TRANSACTION FLOW

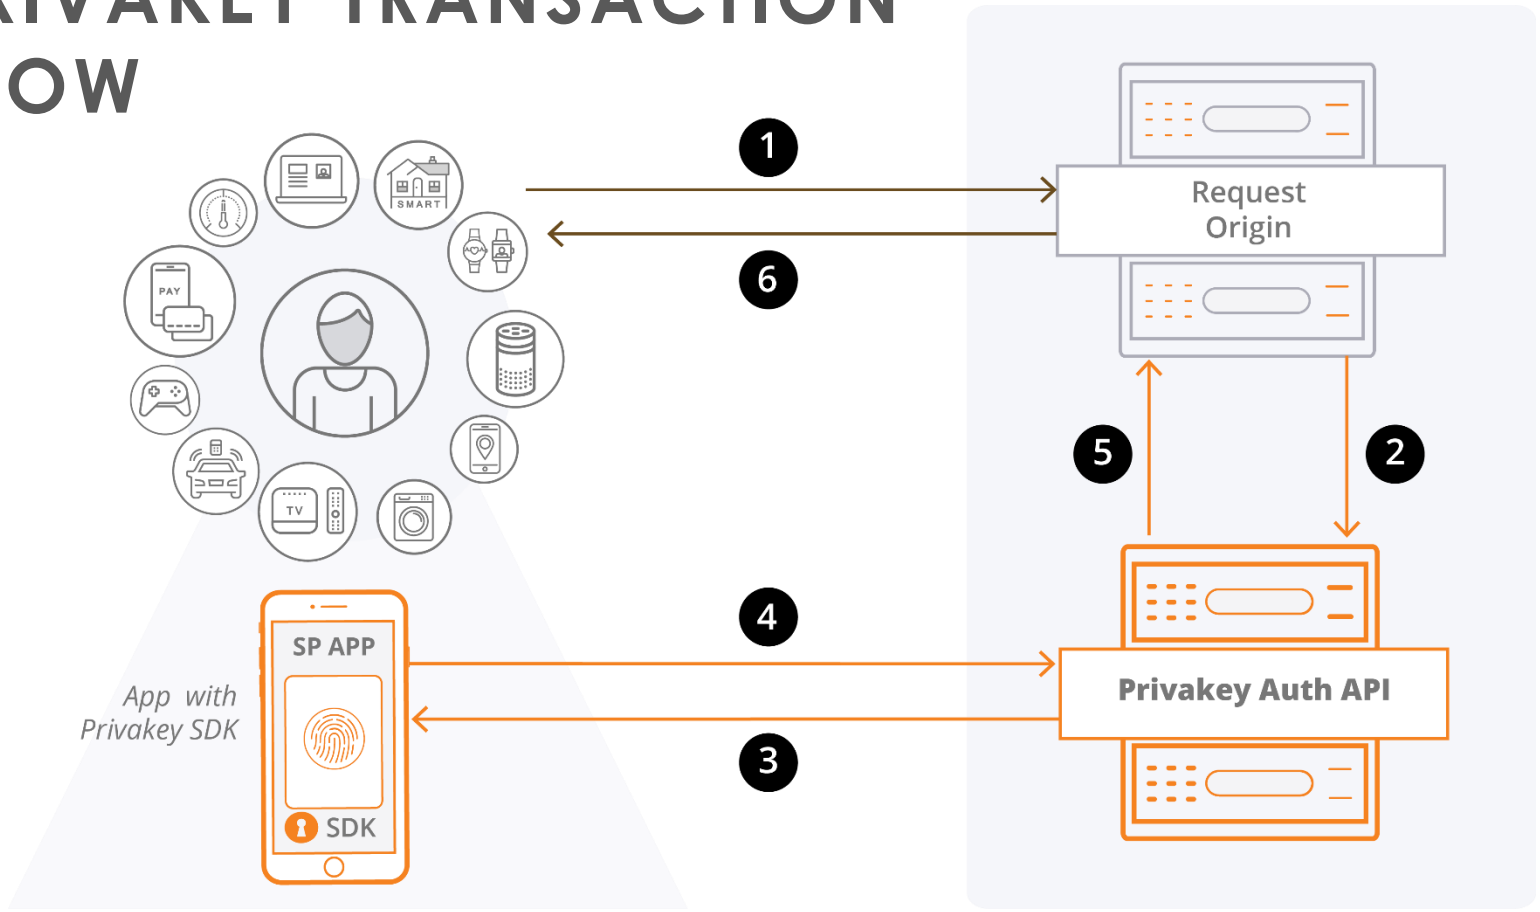

privaKey

MANY POSSIBILITIES

- VOICE interfaces are a great demonstration of VIVA and Privakey technology
- Privakey enhances any workflow requiring strong identity assurance and user consent
- Consumer / retail uses for step-up approvals
  - PSD2 SCA, Fraud Alerts, Transfer Confirmations
- Enterprise use cases including workflow approvals
- Solves the security and experience paradox

IBM WATSON

VICOM Infinity VIVA privaKey

VICOM Infinity

Overcome Data Wrangling for AI

Engine for zLinux

Plug-in Enabler: Turnkey Distributed SQL Engine

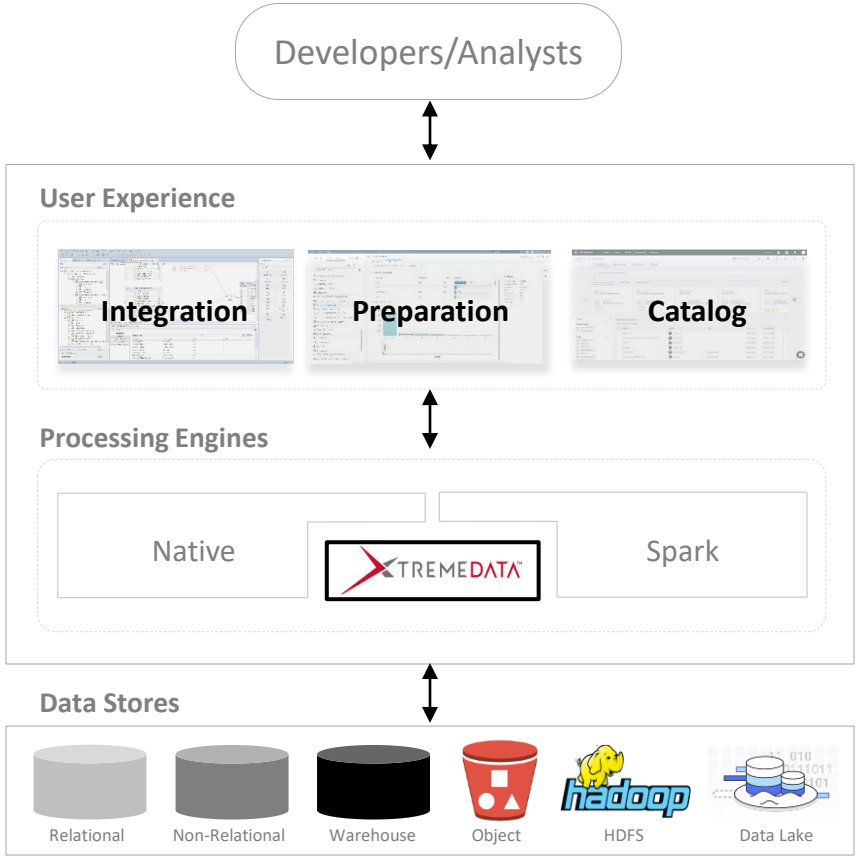

- Disruptive simplicity and performance
- No tuning, “load-and-go”
- Any complexity, size, latency …
- MAINFRAME, cloud, desktop …

Automated data management insights from profiling: data quality stats, schema inferencing, pattern recognition, …

Ushers AI with *Trust*

… and overcomes “80-20” data wrangling issue

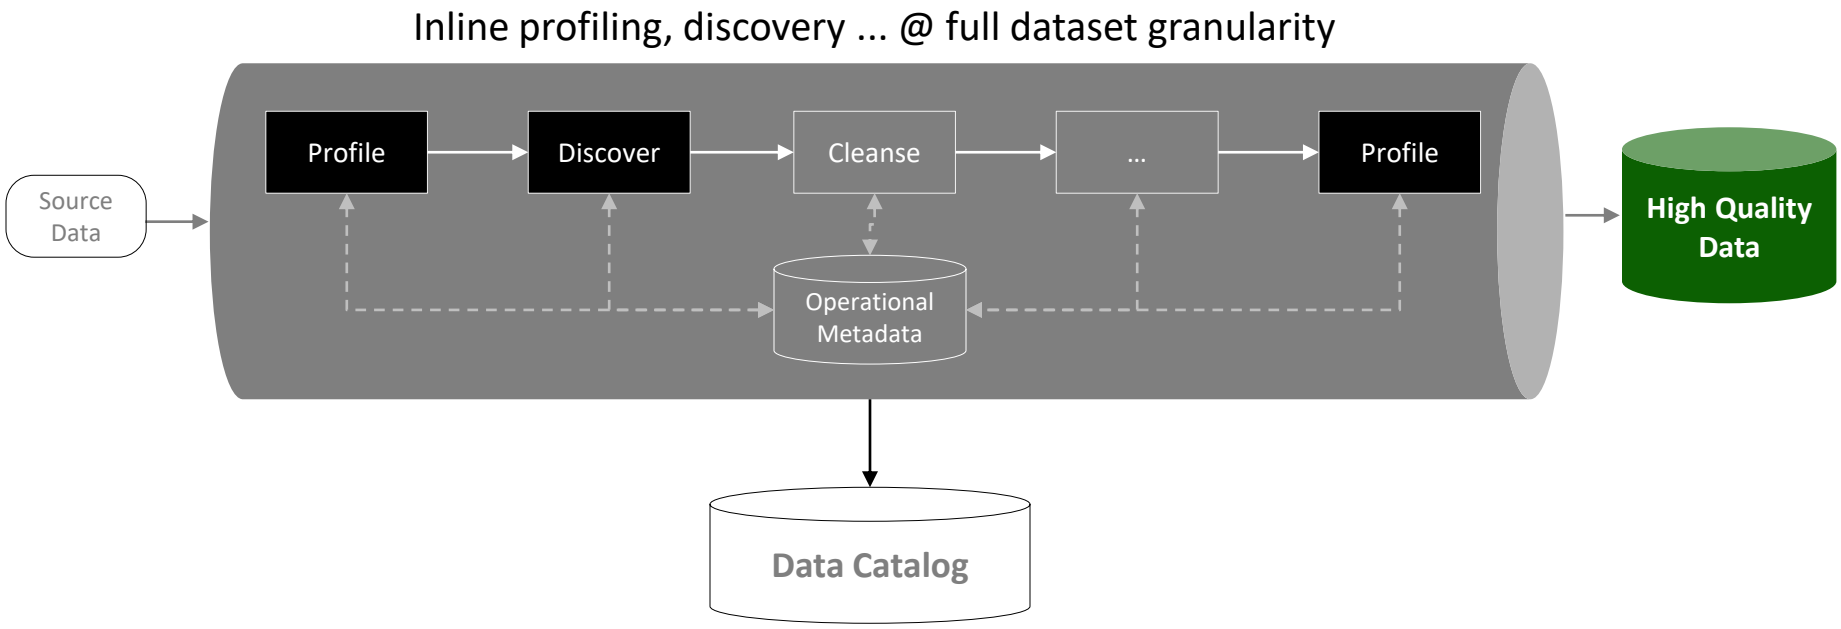

Accelerates Information Architecture

- Complete insights versus sampled data
- Infer schema (data type, precision, scale)
- Detect schema changes/drift, data anomalies …

Augments governance

- Freedom from institutional/tribal knowledge
- Faster development, debugging & testing
- Continuous & proactive data quality…

Enables AI/automation flywheel

- Processing telemetry from every job
- Deep audit trails
- Simplifies validation, reconciliation …

Profile Results: Model & Metadata

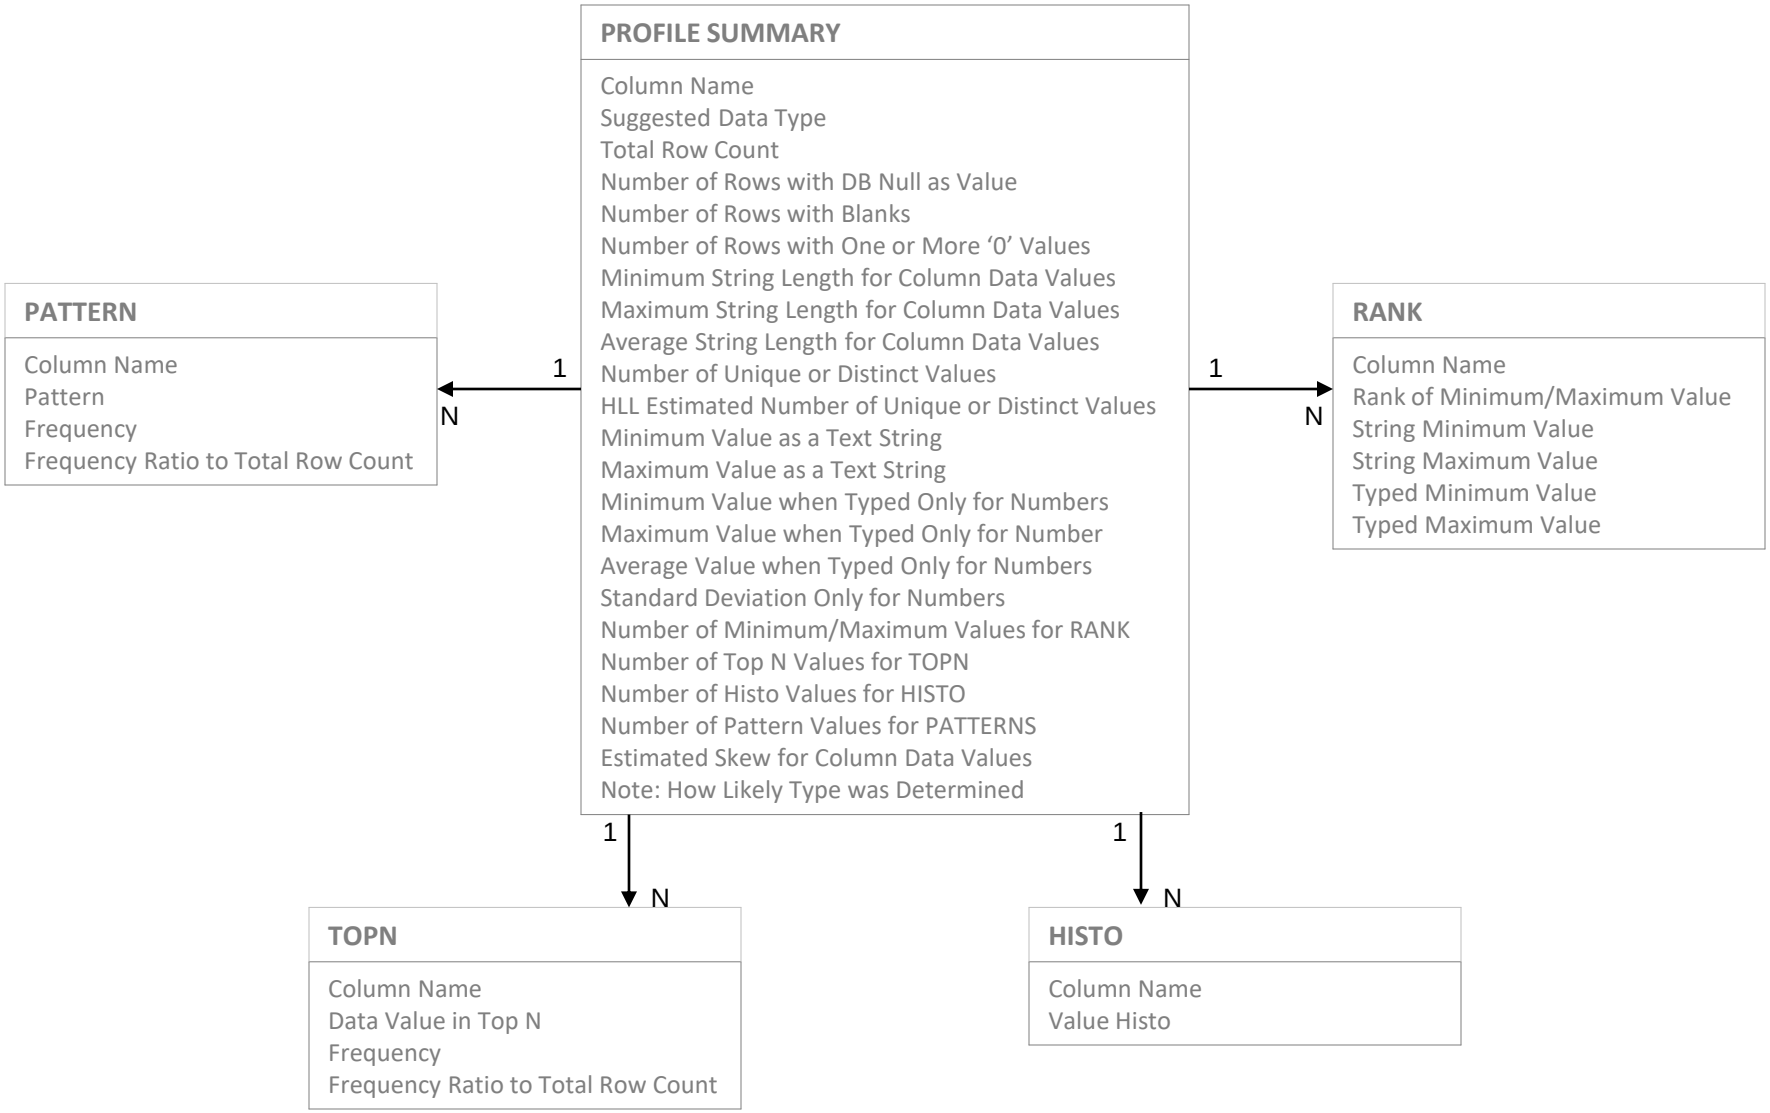

Use Case: Banking, FSI, HLS … Regulatory, Compliance, Audit …

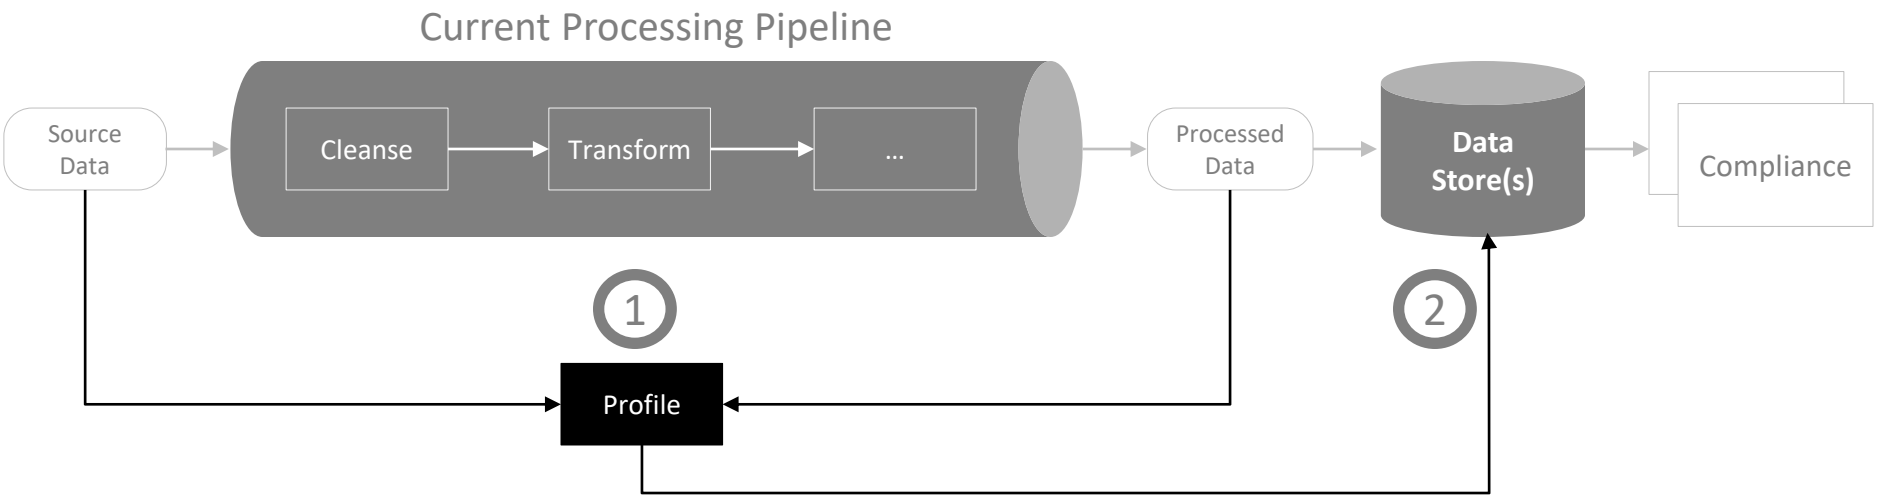

1. Profile source & processed data
2. Upload operational metadata

Benefits

- Automate validation, reconciliation …
- Real-time trace-ability, audit-ability …
- Reduce business risks, costs …

IBM LinuxONE at SU

- IBM LinuxONE (Rockhopper II)
  - Fall / 2019
- Main objective / purpose
  - Support hands-on (lab) activities for campus and **online** students of the new MS Degree in **Enterprise Data Systems (MS-EDS)** at the School of Information Studies
    - MS-EDS focus areas:
      - Cloud-based environments and technologies
      - Containers, Virtual Machines, Automation, System Management
    - MS-EDS industry partner collaboration/input
      - IBM + Red Hat
      - Microsoft
      - VMware, Cisco, …
- Additional uses:
  - Support lab activities of the MS in Applied Data Science (Big Data / AI)

IBM LinuxONE at SU

- OpenShift Cluster (v.4.2) on z/VM
- Several network support services setup
- Future tasks:
  - IBM Cloud Pak for Data

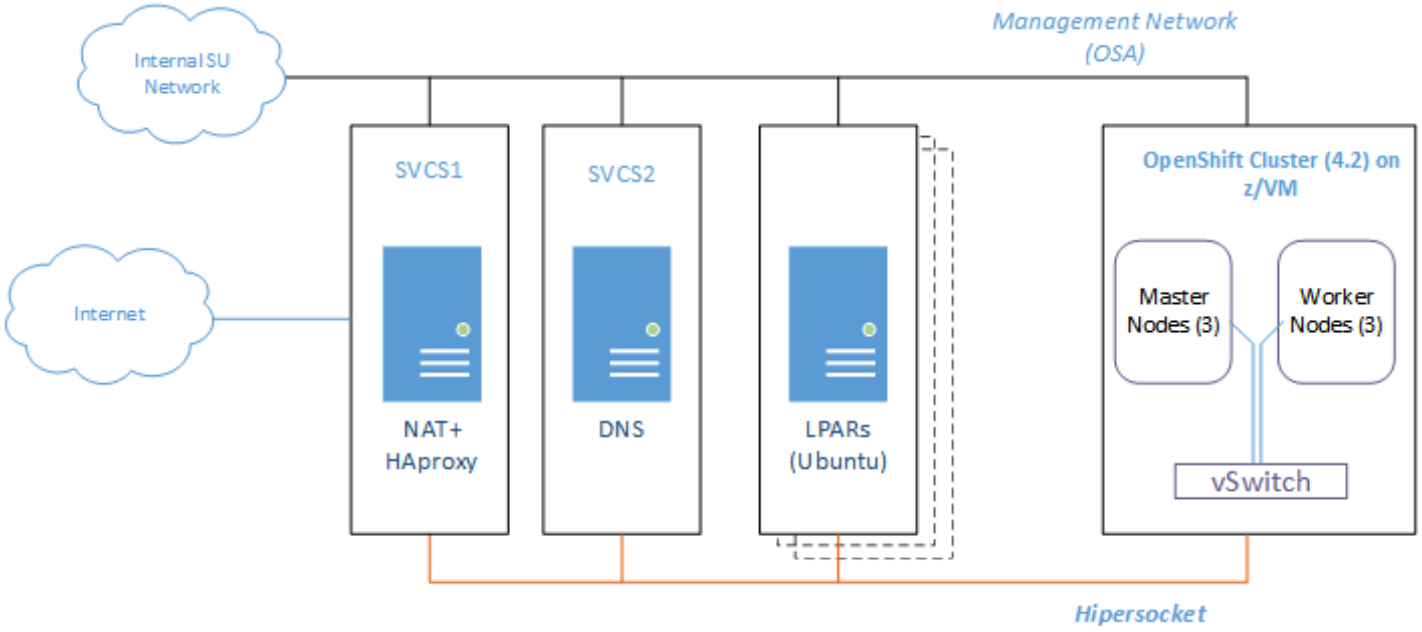

OpenShift workloads on IBM Z run on fewer cores and cost less than on x86

Comparison of 8 IFLs added to existing z15 T01 system versus 7 x86 servers

*When driving **workload based on an SLA,** which platform delivers the best cost?*

Total 136 cores
across 7 x86 servers

$1,507,685
(3 Year Cost Analysis)
Cost per Request meeting SLA: 1,507, 685 ÷ 14792 = $102
Cost per Request measured: 1,507, 685 ÷ 15571 = $97

OpenShift
Banking
workload

Cost per request: x86 is 1.7x more than on Z (meeting SLA)
(102 ÷ 61 = 1.7x)

Cost per request: x86 is 1.7x more than on Z (measured)
(97 ÷ 58 = 1.7x)

SLA: 64 ms Response Time
Workload: 400 users

IBM z15 T01

IFLs 8
across 3 LPARs

$947,067
(3 Year Cost Analysis)
Cost per Request meeting SLA: $947,067 ÷ 15458 = $61
Cost per Request measured: $947,067 ÷ 16272 = $58

17X Core reduction

37% Lower cost

Disclaimer: This is an IBM internal study designed to replicate banking OLTP workload usage in the marketplace on an IBM z15 T01 using eight IFLs across three LPARs. Three IFLs and a total of 512 GB memory were allocated to one LPAR for two OpenShift masters and two worker nodes. Another four IFLs and a total of 512 GB memory were allocated to a second LPAR for one OpenShift master and two workers. One IFL and a total of 128 GB memory were allocated to a third LPAR for the OpenShift load balancer. IBM Storage DS8886 was used to create eight 250 GB DASD minidisks for each of the eight z/VM guests running in the LPARs. The OpenShift cluster version 4.2.20, using Red Hat Enterprise Linux CoreOS (RHCOS) for IBM Z, was running across seven z/VM guests and the remaining eighth z/VM guest was running the OpenShift load balancer. SMT was enabled across all IFLs. The x86 configuration was comprised of six servers running KVM with 15 guests (three masters and twelve workers) for the OpenShift cluster version 4.3.5 with RHCOS and a seventh server was used for the load balancer on RHEL 7.6. For x86 storage each guest operating system was configured with a 100 GB of virtual disk. Each guest had access to all vCPUs of the KVM server on which it was running. Compared x86 models for the cluster were all 2-socket servers containing a mix of 6-core, 8-core, 12-core and 16-core Haswell, Skylake and Ivy Bridge x86 processors using a total of 136 cores with a total of 2,304 GB memory. The load balancer was a 2-socket 8-core server with a total of 384 GB memory. Both environments used jMeter to drive maximum throughput against two OLTP workload instances and were sized to deliver comparable results (15,456 responses per second (RPS) with IBM Z and 14,848 RPS with x86). The results were obtained under laboratory conditions, not in an actual customer environment. IBM's internal workload studies are not benchmark applications. Prices, where applicable, are based on U.S. prices as of 02/12/2020 from our website and x86 hardware pricing is based on IBM analysis of U.S. prices as of 03/01/2020 from IDC. Price comparison is based on a three year total cost of ownership including HW, SW, networking, floor space, people, energy/cooling costs and three years of service & support.

OpenShift workloads on LinuxONE run on fewer cores and cost less than on x86

Comparison of new LinuxONE III system versus 7 x86 servers

When driving ***workload based on an SLA,*** *which platform delivers the best cost?*

OpenShift Banking workload

SLA: 55 ms Response Time
Workload: 400 users

Total 138 cores across 7 x86 servers

$1,507,684
(3 Year Cost Analysis)
Cost per Request meeting SLA: 1,507, 685 ÷ 14325 = $102
Cost per Request measured: 1,507, 685 ÷ 15571 = $97

Cost per request: x86 is 2.1x more than on LinuxONE (meeting SLA)
(102 ÷ 50 = 2.1x)

Cost per request: x86 is 2.0x more than on LinuxONE (measured)
(97 ÷ 48 = 2.0x)

IBM LinuxONE III LT2

IFLs 8
Across 2 LPARs

$779,841
(3 Year Cost Analysis)
Cost per Request meeting SLA: $779,841 ÷ 15487 = $50
Cost per Request measured: $779,841 ÷ 16302 = $48

17X Core Reduction

48% Lower cost

Disclaimer This is an IBM internal study designed to replicate banking OLTP workload usage in the marketplace on an IBM LinuxONE III T02 using eight IFLs across two LPARs. Seven IFLs and a total of 640 GB memory were allocated to one LPAR for three OpenShift masters and four worker nodes. One IFL and a total of 128 GB memory were allocated to the second LPAR for the OpenShift load balancer. IBM Storage DS8886 was used to create eight 250 GB DASD minidisks for each of the eight z/VM guests running in the LPARs. The OpenShift cluster version 4.2.20, using Red Hat Enterprise Linux CoreOS (RHCOS) for IBM Z, was running across seven z/VM guests and the remaining eighth z/VM guest was running the OpenShift load balancer. SMT was enabled across all IFLs. The x86 configuration was comprised of six servers running KVM with 15 guests (three masters and twelve workers) for the OpenShift cluster version 4.3.5 with RHCOS and a seventh server was used for the load balancer on RHEL 7.6. For x86 storage each guest operating system was configured with a 100 GB of virtual disk. Each guest had access to all vCPUs of the KVM server on which it was running. Compared x86 models for the cluster were all 2-socket servers containing a mix of 6-core, 8-core, 12-core and 16-core Haswell, Skylake and Ivy Bridge x86 processors using a total of 136 cores with a total of 2,304 GB memory. The load balancer was a 2-socket 8-core server with a total of 384 GB memory. Both environments used jMeter to drive maximum throughput against two OLTP workload instances and were sized to deliver comparable results (15,487 responses per second (RPS) with IBM Z and 14,325 RPS with x86). The results were obtained under laboratory conditions, not in an actual customer environment. IBM's internal workload studies are not benchmark applications. Prices, where applicable, are based on U.S. prices as of 02/12/2020 from our website and x86 hardware pricing is based on IBM analysis of U.S. prices as of 03/01/2020 from IDC. Price comparison is based on a three-year total cost of ownership including HW, SW, networking, floor space, people, energy/cooling costs and three years of service & support.:

For More Information please contact…

Len Santalucia
CTO & Business Development Manager
Vicom Infinity, Inc.
New York, NY 10001
917-856-4493 mobile
LSantalucia@vicominfinity.com

About Vicom Infinity
Account Presence Since 1990's
IBM Gold Business Partner
Reseller of IBM Z and Storage Hardware, Software, and Maintenance
Vendor Source for the Last 18 Generations of Mainframes/IBM Storage
Professional IT Architectural Services and IBM Tier1 Services Provider
Vicom Family of Companies Also Offer Leasing & Financing and IT Staffing & IT Project Management
Linux Foundation Open Mainframe Project – Chair
IBM Z Champion, Academic Initiative Leader, Council Sponsor, Ecosystem Advocate, Beta Tester

Recipient of *The North America IBM Z Business Partner Sales Excellence Award*