



OPEN SOURCE BUSINESS CONFERENCE

Building Your Big Data Future with Open Source

COMPUTERWORLD
OSBC
SAN FRANCISCO



Building Your Big Data Future with Open Source

Innovating at Scale



Phil Robb

Director, Open Source Program Office
Office of Strategy & Technology

Vin Sharma

Manager, Open Source Business Strategy
Industry Standard Servers & Software



Building Your Big Data Future with Open Source





Building Your Big Data Future with Open Source

Information
explosion

Dynamic cloud
services

Content
transformation

Intelligent
infrastructure

Sustainability

Beijing

Tokyo

Palo Alto

Bristol

St. Petersburg

Haifa

Bangalore

5 research themes

7 locations

23 labs

600 researchers

Open Innovation

Technology Transfer





Building Your Big Data Future with Open Source

HP has been helping businesses innovate by providing technology solutions that combine HP intellectual property, standards-based hardware, and open source software.

- | | | |
|---|--|--|
| <ul style="list-style-type: none">• 37,000+ patents worldwide• Top-ranked IT company in patent quality (past year)• Fourth most valuable patent portfolio worldwide | <ul style="list-style-type: none">• ~40% share of Linux & virtualization server market• Top-ranked server vendor for Linux and virtualization since 1998• Infrastructure provider to 9 out of top 10 cloud service providers worldwide | <ul style="list-style-type: none">• Founding member of Linux Foundation• Sponsor Apache Software Foundation• Contribute to ~200 open source projects• Initiated Fossology |
|---|--|--|





Building Your Big Data Future with Open Source

HP is a trusted solution provider to enterprises using open source because we have consistently advocated the use of open source development models to drive innovation.





Building Your Big Data Future with Open Source

HP helps open source software get ready for the enterprise



HP provides support



Partner provides support



Users support themselves

Commercial Linux

- Tests commercial Linux distribution early and often (with beta OS, proto hardware)
- Enables commercial distribution near launch date of new ProLiant hardware
- Obtains Linux vendor certification for representative ProLiant servers
- Provides support for ProLiant Support Pack

Community Linux

- Performs sanity tests of community Linux distributions with several ProLiant servers
- Supports ProLiant server hardware
- Offers the ProLiant Support Pack “as is”
- Directs software issues to community resources and supports the community

Validates reference configurations, publishes best practices, -supports industry events, and contributes to open source projects





Building Your Big Data Future with Open Source

HP helps enterprises get ready for open source software



- Forum to facilitate study of FOSS via free data analysis tools
- Original FOSSology tool-set developed and contributed by HP



- Community for free and open source software governance
- Conduit for developing and sharing information and best practices
- Founded by HP and partners

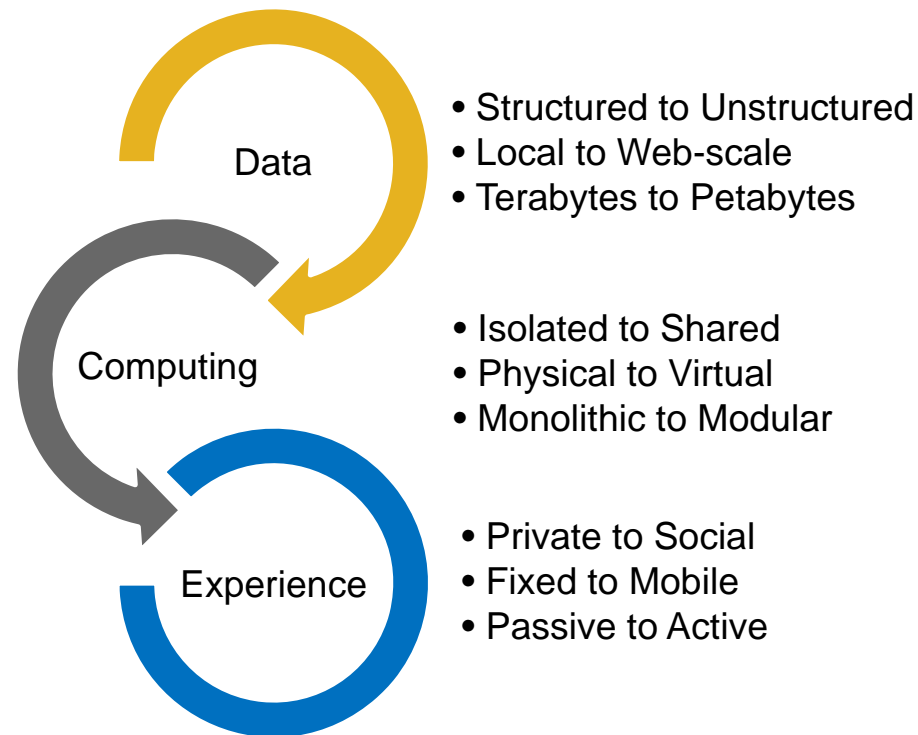
<http://www.fossbazaar.org>





Building Your Big Data Future with Open Source

HP is now helping enterprises, service providers, and web-scale startups tackle the challenges of big data, cloud computing, and “consumerization of IT”.





Building Your Big Data Future with Open Source

Big Data

- Structured to Unstructured
- Local to Planetary
- Terabytes to Exabytes

ENTERPRISE

ERP

Products CRM Suppliers
Customers Partners

FINANCIAL SERVICES

Algorithmic Trading



High-frequency Trading

COMPLIANCE

Sarbanes-Oxley
HIPPA
Basel II

COMMUNICATIONS



Call Detail Records

HEALTHCARE



Electronic Patient Records




Gene Sequencing

Medical Imaging

GEOPHYSICAL EXPLORATION



SENSOR NETWORKS



MOBILITY



SOCIAL MEDIA





Building Your Big Data Future with Open Source

“One Word... Analytics”



The Graduate, © Metro-Goldwyn-Mayer Studios, 1967





Building Your Big Data Future with Open Source


HP Vertica Analytics System



- Low-latency “in-time” analytics
- Up to 10x better load performance
- 50–1000x faster query performance than traditional row-stores
- Hybrid in-memory/on-disk architecture
- Keep data close to the processor/core
- Large scale, multi-use workloads
- Simple install/use with auto setup and tuning
- High scalability and full parallelism
- Runs on Red Hat Enterprise Linux
- Industry-standard ProLiant x86 hardware
- Hadoop Connector

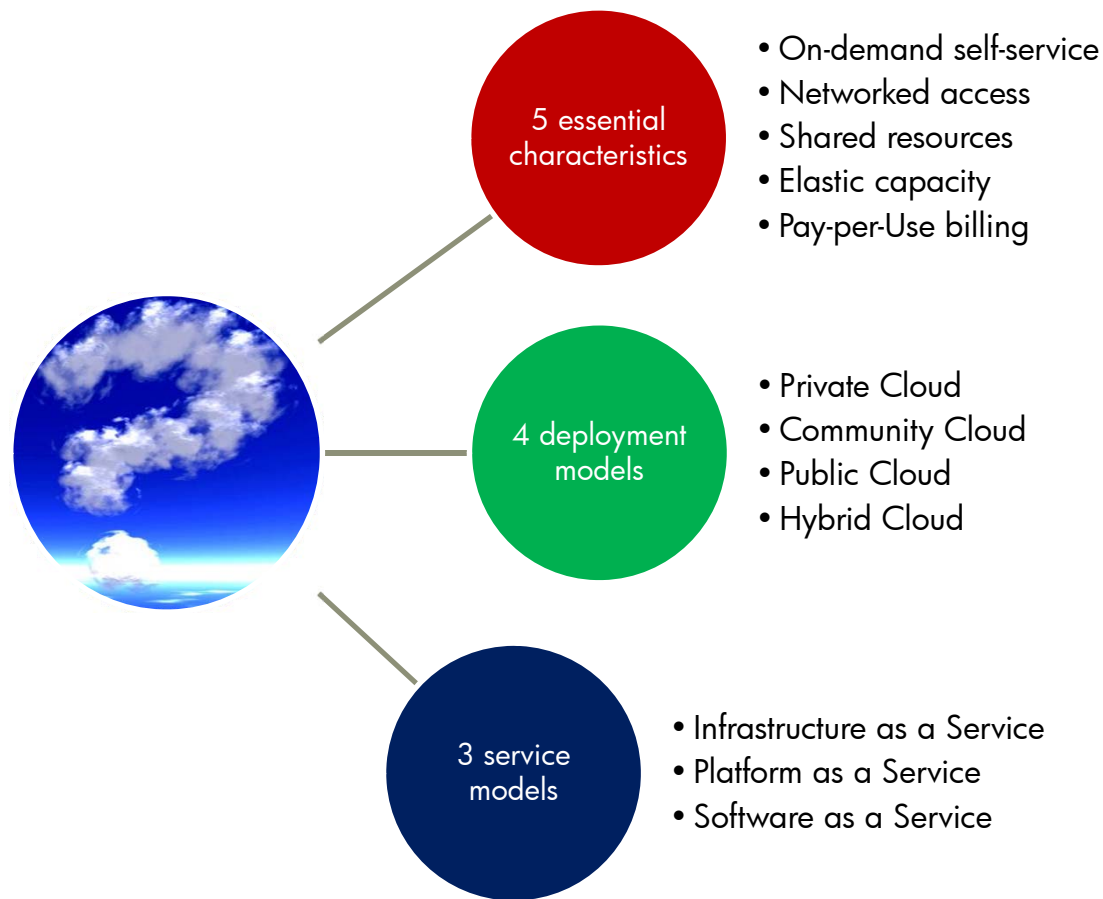


Building Your Big Data Future with Open Source

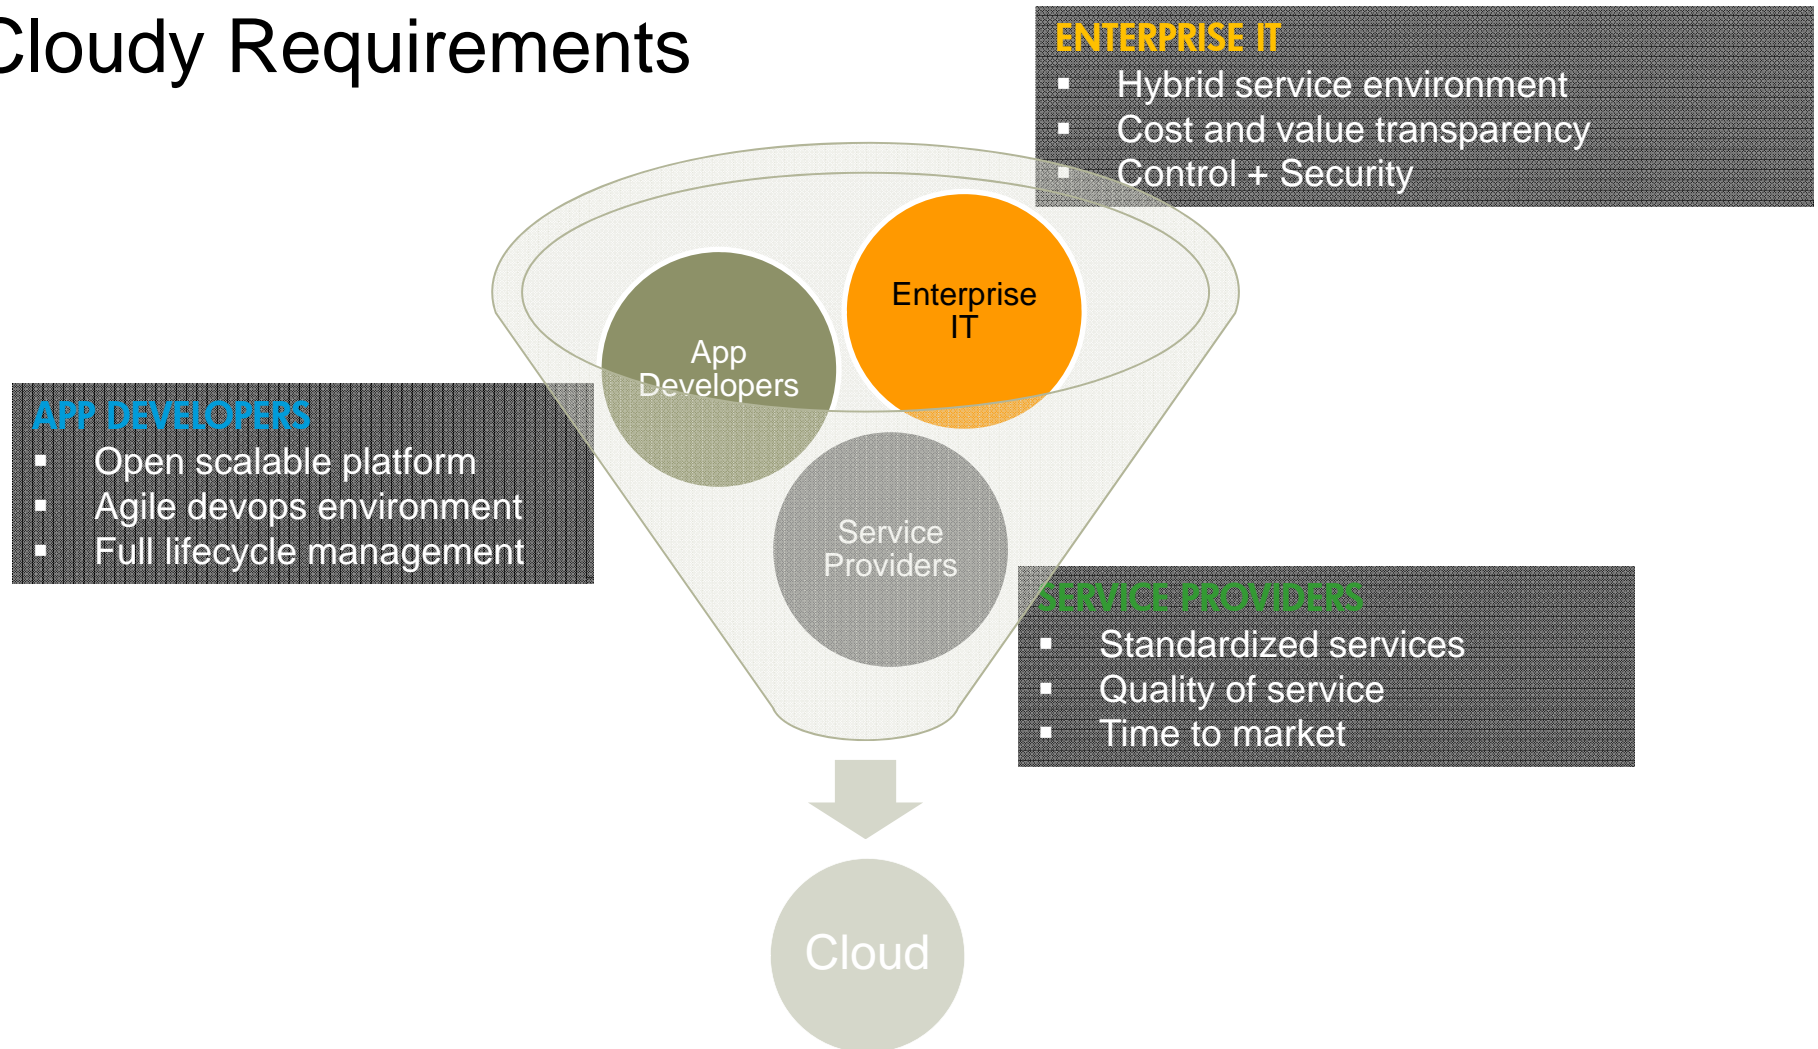


- First analytic DB company to deliver parallel connector for Hadoop
- ~25 customers currently using Vertica with Hadoop
- Flexibility and speed in loading data from Hadoop to Vertica
- Query data from Vertica in Hadoop as part of a MapReduce job
- Enables job scheduling between both environments
- Open Source, supported by HP/Vertica
- <http://www.vertica.com/the-analytics-platform/native-bi-etl-and-hadoop-mapreduce-integration/>

Cloud Computing



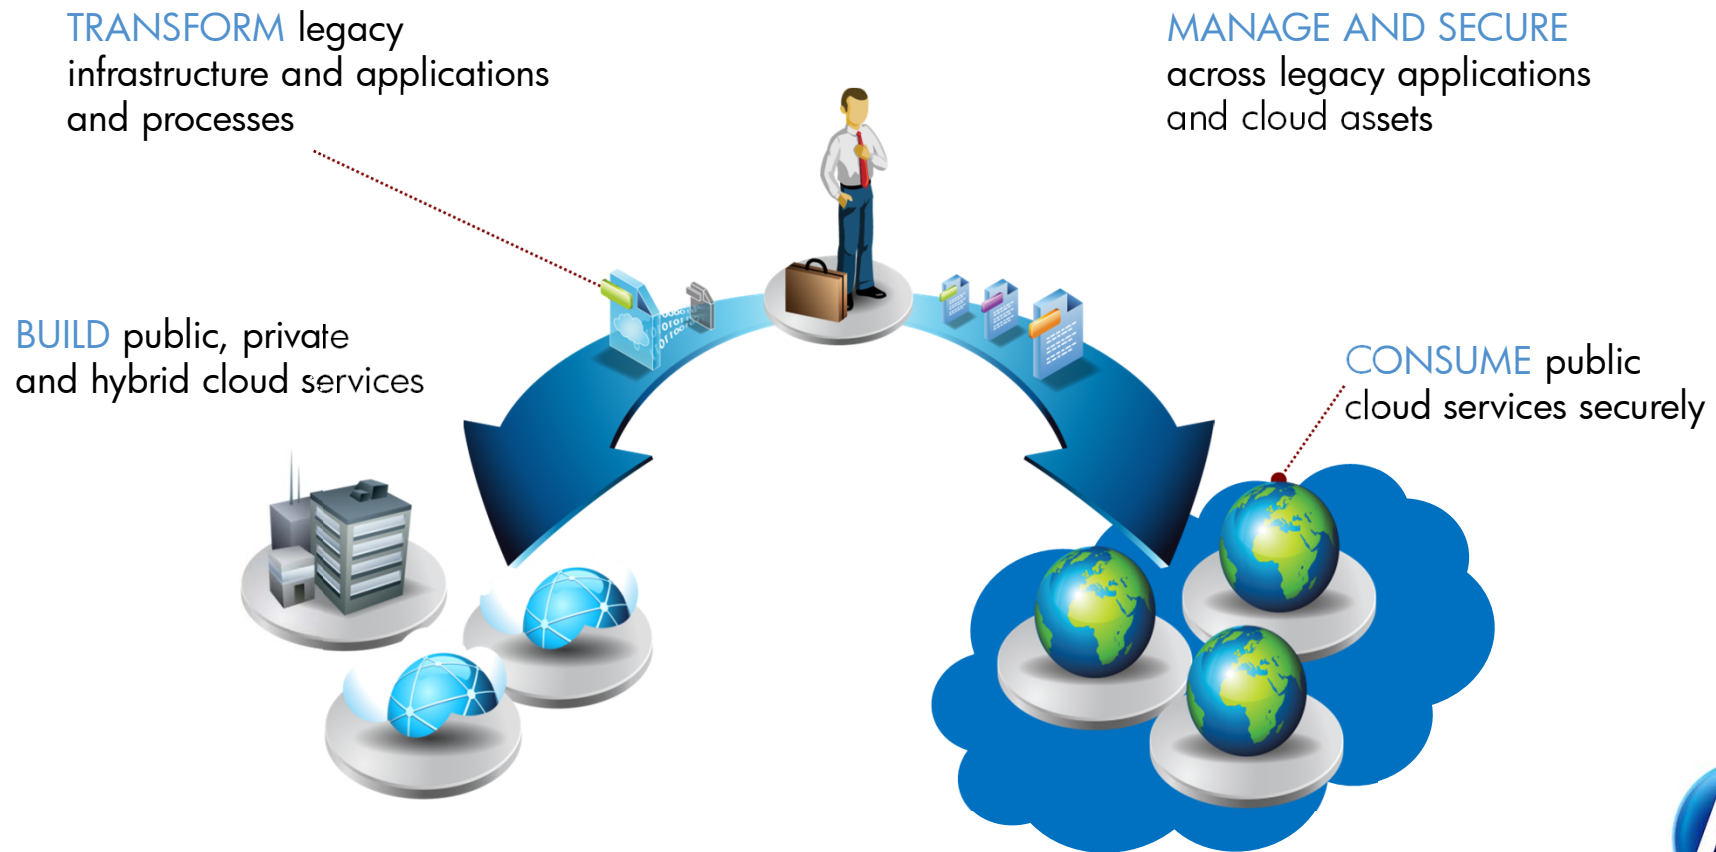
Cloudy Requirements





Building Your Big Data Future with Open Source

Enterprise CIOs need to bridge legacy with cloud






Building Your Big Data Future with Open Source

HP CloudSystem

- Single services view across hybrid cloud
- Multi-hypervisor, multi-OS, heterogeneous infrastructure
- Intelligent automation & orchestration
- Rapid application & infrastructure deployment
- Converged Infrastructure built for the cloud
- Agility for the enterprise, revenue growth for service providers



HP CLOUDSYSTEM





Building Your Big Data Future with Open Source

