

Intel and Red Hat: Enhancing OpenStack for Enterprise Deployment

Krish Raghuram - Datacenter Software Marketing Manager, Open Source Technology Center, Intel

Sean Cohen, Principal Product Manager, RHEL OpenStack Platform, Red Hat

June 24, 2015

The Vision of a
Software Defined
Infrastructure

Intel and RHEL
OpenStack
Platform

Intel & Red Hat:
Delivering
Solutions
Together

Unprecedented Opportunity for Growth



Services and Devices

Vast Amounts of Data

More Complex Services

Traditional Data Center Models Can't Keep Pace



COMPUTE

Average utilization <50% despite virtualization³



STORAGE

40% data growth CAGR, 90% unstructured²



NETWORK

2-3 weeks to provision new services¹



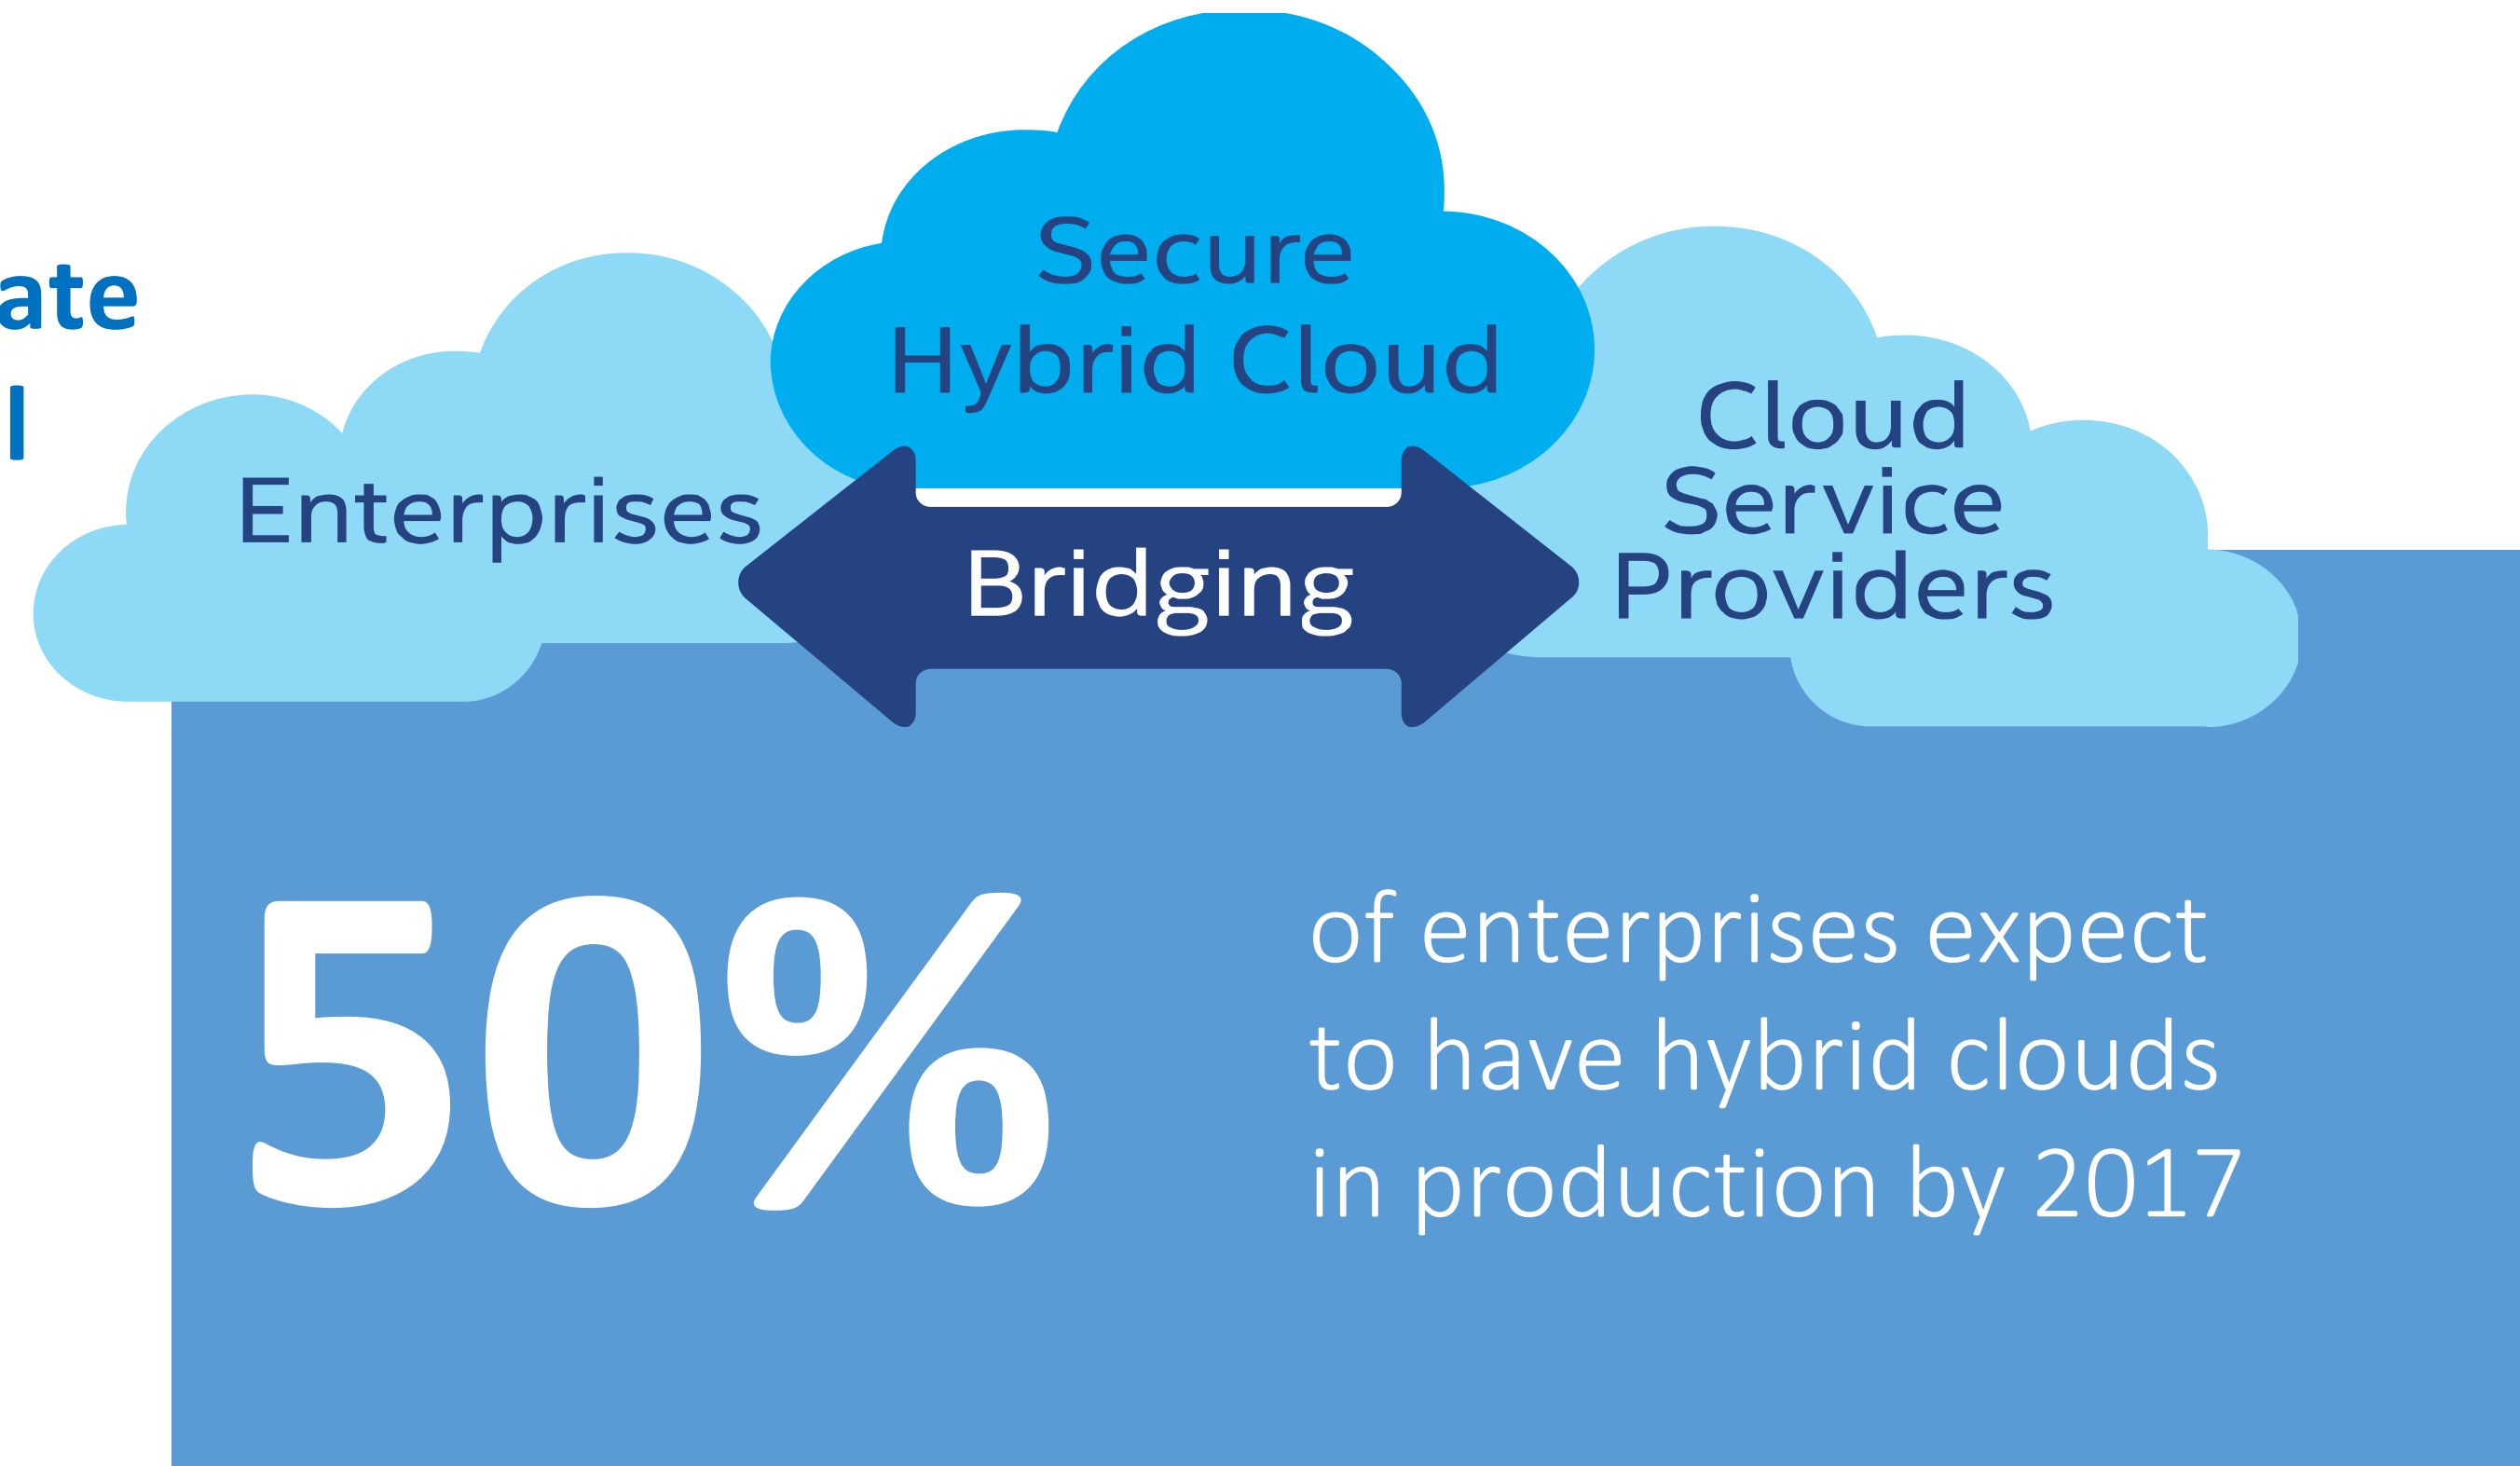
NEW STRATEGIES ARE NEEDED

¹ IDC Server Virtualization and The Cloud 2012; ² IDC's Digital Universe Study, sponsored by EMC, December 2012; ³ Source: Intel IT internal estimate

Hybrid Clouds Answer the Challenge

Agile, Efficient, and Secure Enterprise private cloud for security, compliance, and control

- On-demand compute, network, and storage
- Advanced automation/orchestration
- Assured service levels
- Standards-based hardware and open-source software



Source: "Hybrid Cloud Deployments Rising: Gartner," an eWeek article by Nathan Eddy, October 1, 2013.
www.eweek.com/small-business/hybrid-cloud-deployments-rising-gartner.html

OpenStack Enterprise adoption

75% are planning to use OpenStack for cloud initiatives

52% Reduce IT Operational Costs

The OpenStack enterprise adoption survey of more than 310 IT decision makers and professionals from around the world

77% Enterprise users are looking to OpenStack to reduce time to deployment

44% Want more responsive infrastructure to demand changes

Commissioned by Red Hat through TechValidate

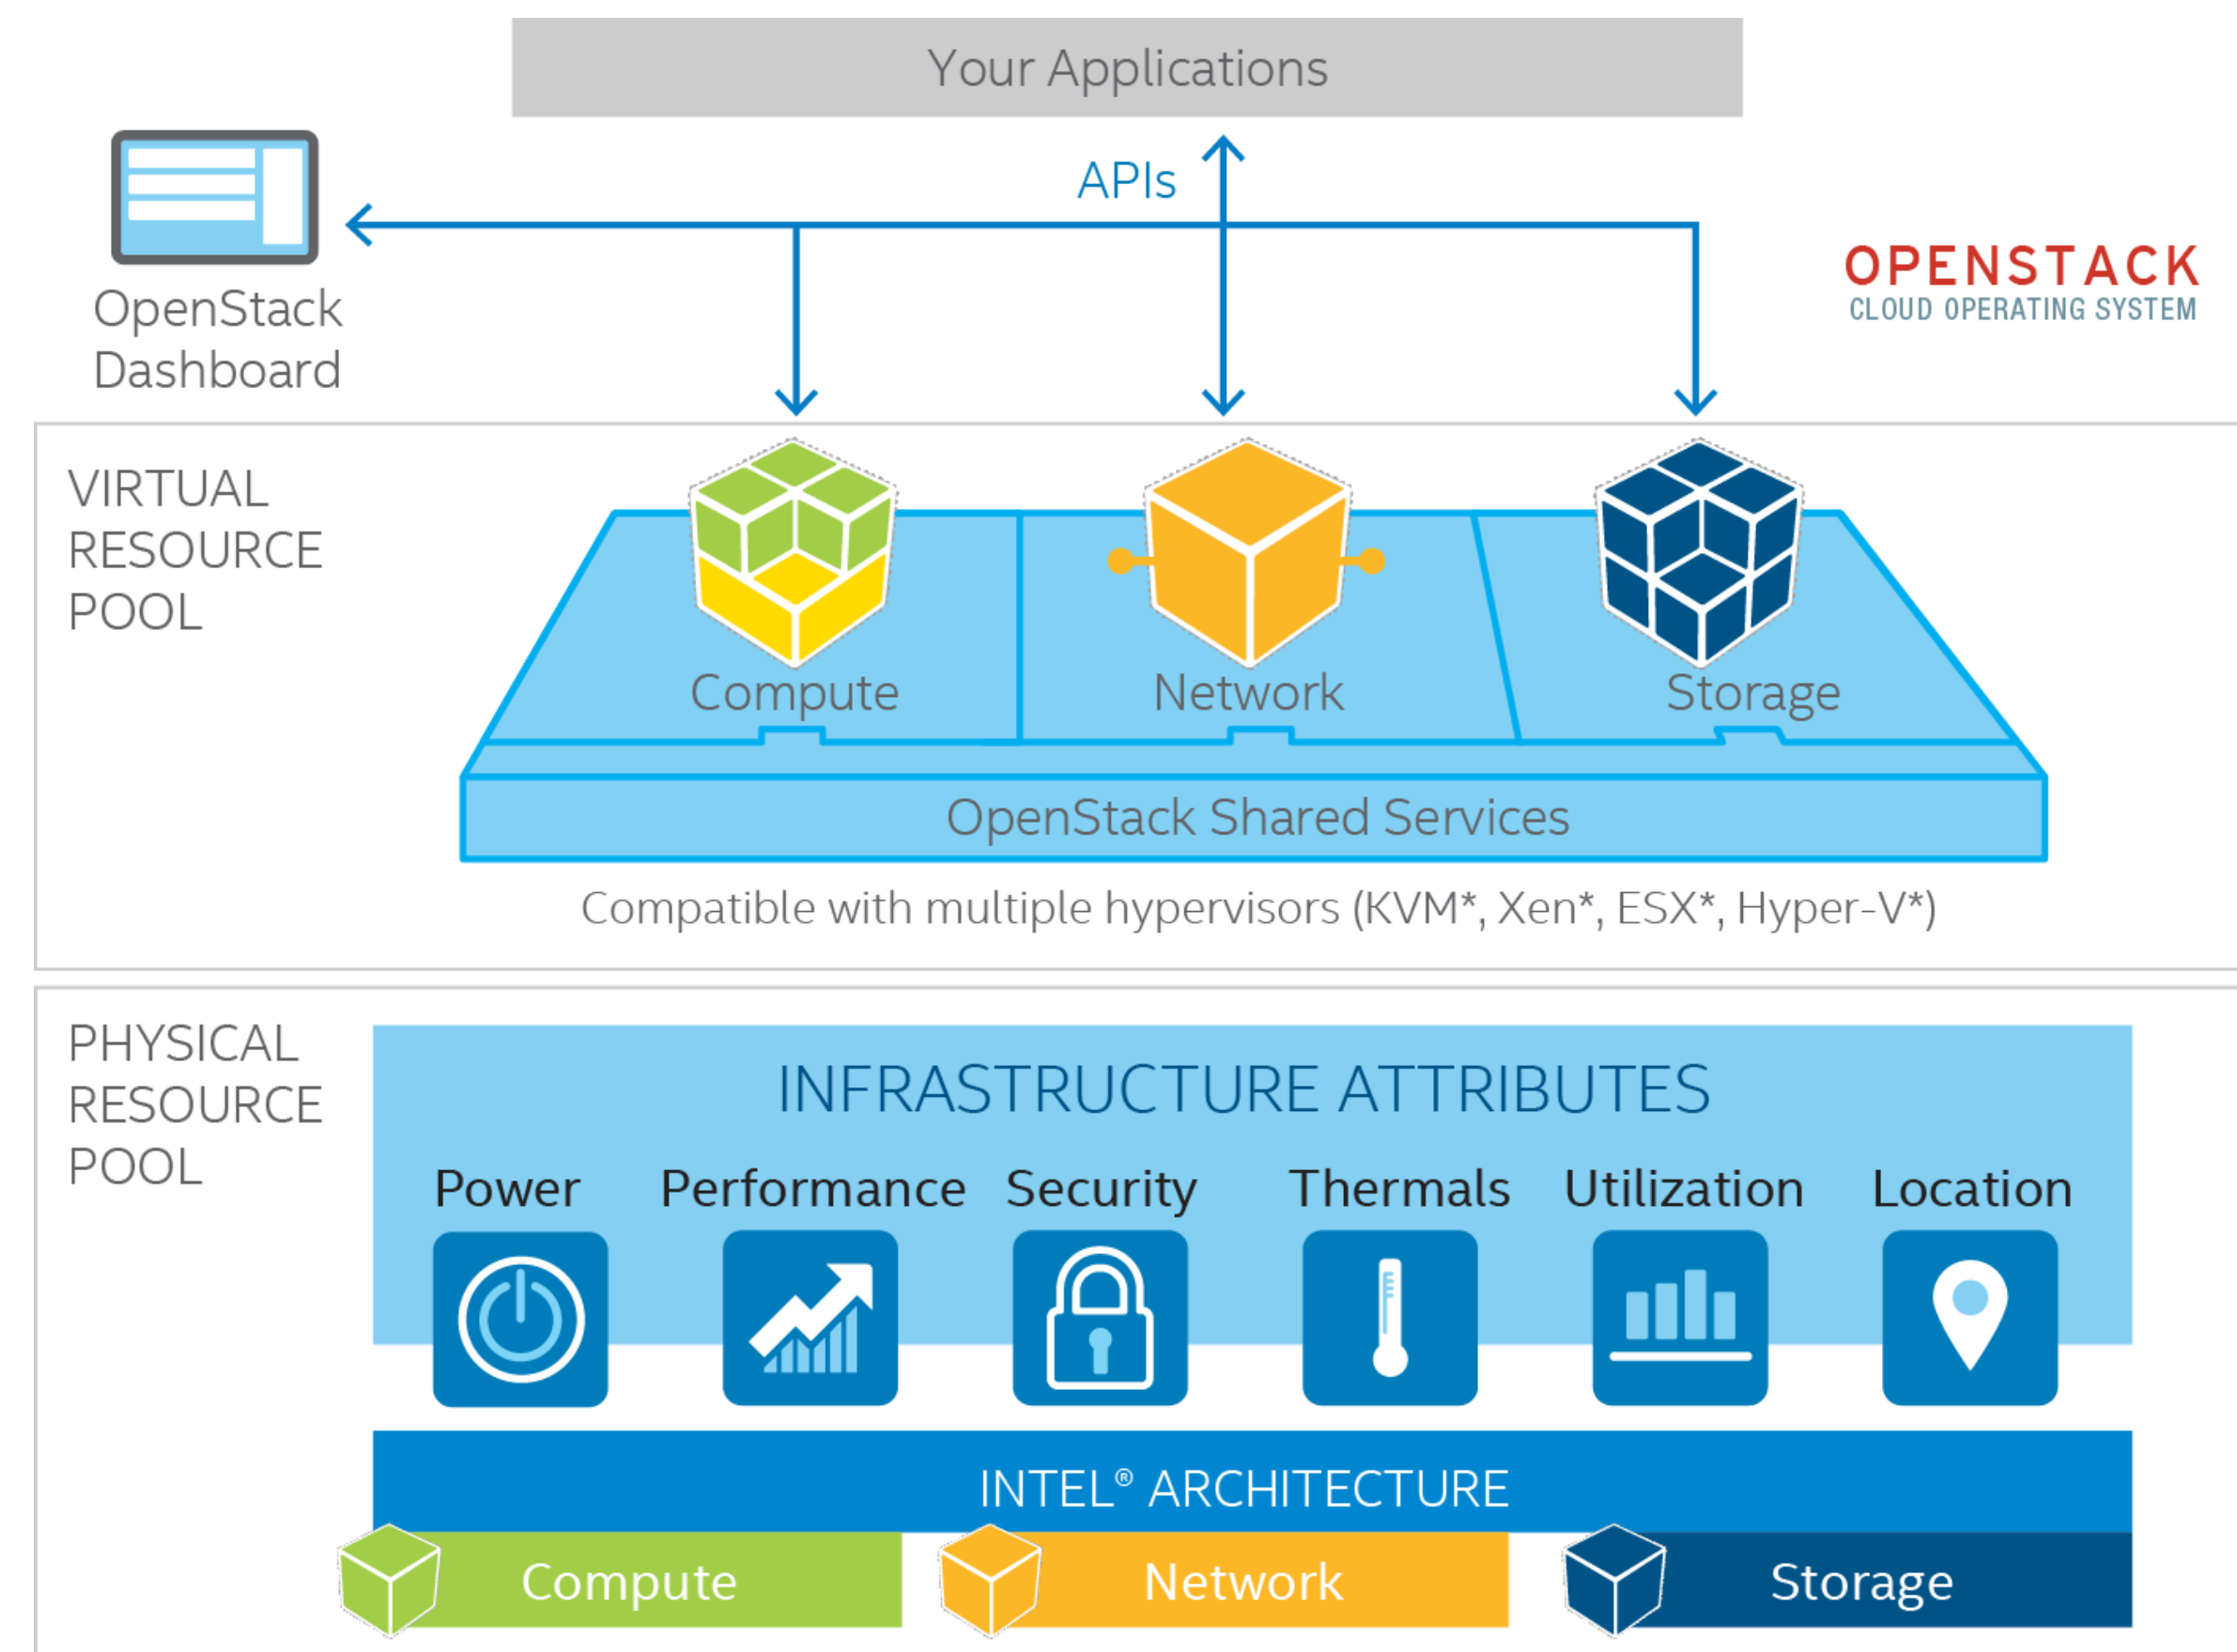
The Power of a Trusted, Private OpenStack* Cloud

Deploy services in minutes with the right service level

Business and IT Agility
through on-demand, self-service delivery

Operational Efficiency
through advanced orchestration

Lower Costs
through efficient resource sharing



*Other names and brands may be claimed as the property of others.

The OpenStack* Advantage

A massively scalable, open cloud computing platform

Flexible and Affordable

- Open-source economics
- Modular components
- Standards-based APIs

Rapid Innovation on an Open Platform

- Global ecosystem
- 25,000+ OpenStack members (519 companies across 166 countries)¹
- 1,300 active contributors¹

“OpenStack is here, and it’s ready. Eleven Fortune 100 firms (eg. Best Buy, BMW, Comcast, Disney, Wal-Mart) are already using OpenStack for production environments, making its viability and presence in the market irrefutable”

- Forrester Research, Inc, 2015

Source: “OpenStack is Ready – Are You?,” by Lauren E. Nelson with Glenn O’Donnell, Andre Kindness and Michael Caputo, Forrester Research, May 18, 2015.
See <https://www.openstack.org/enterprise/forrester-report/>

¹ Source: www.openstack.org

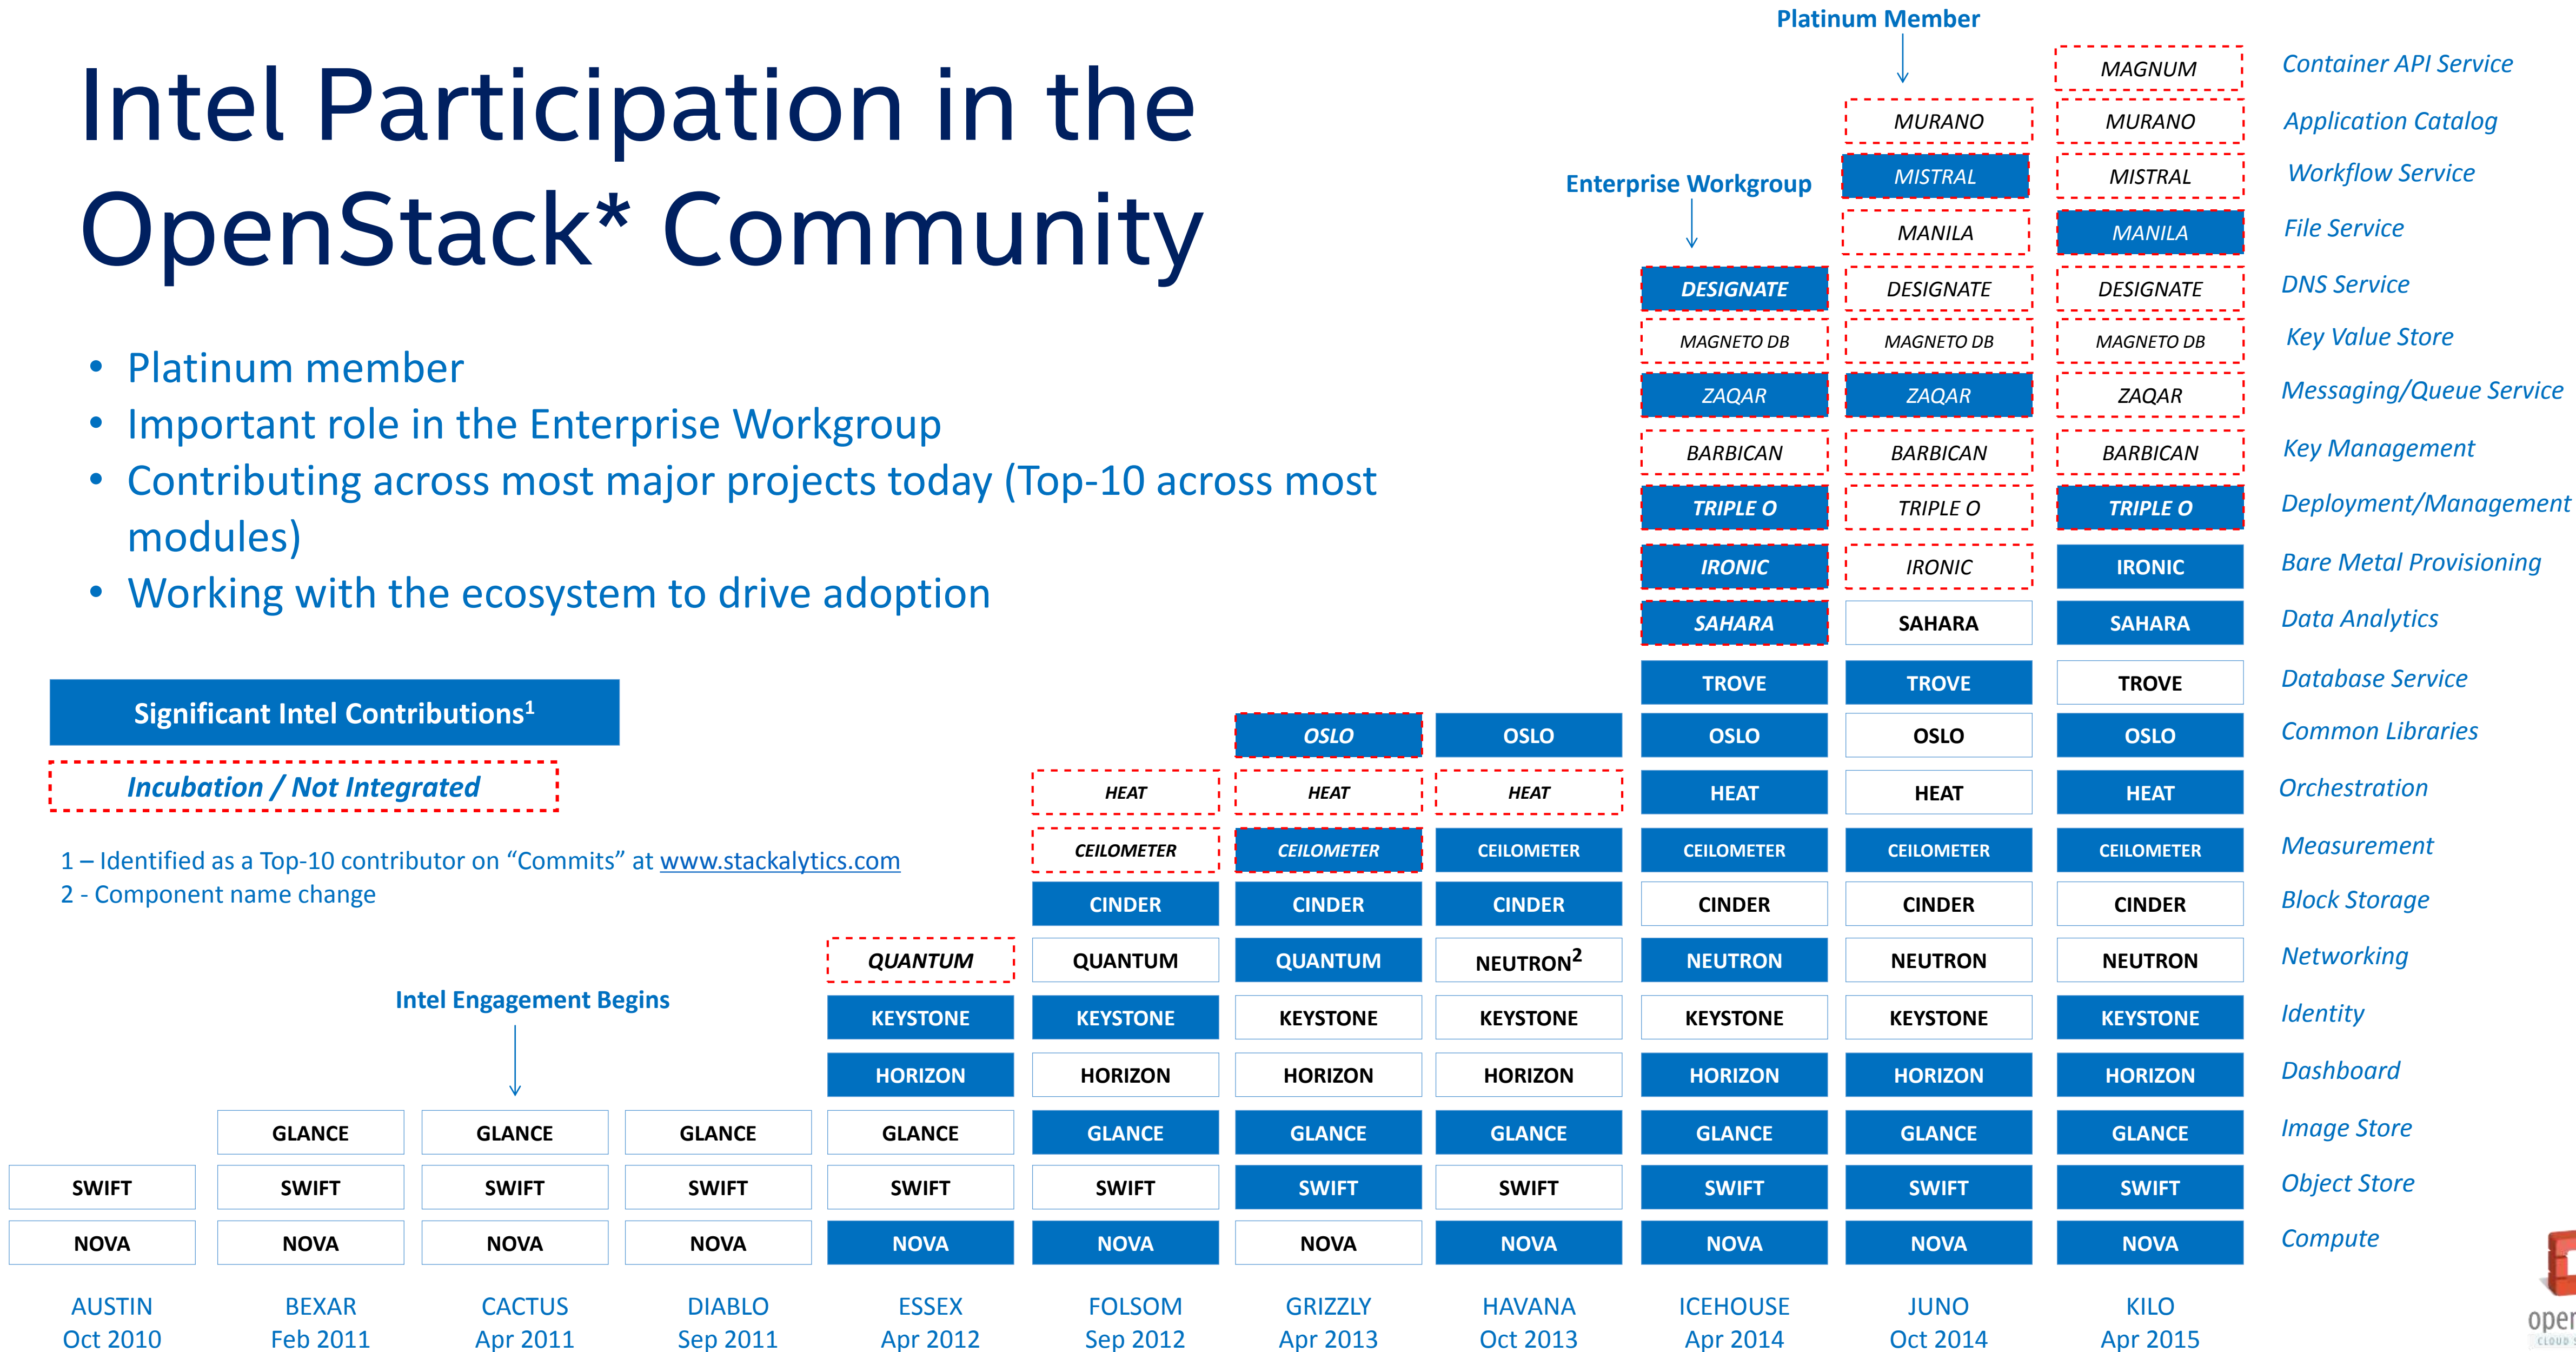
The Vision of a
Software Defined
Infrastructure

Intel and RHEL
OpenStack
Platform

Intel & Red Hat:
Delivering
Solutions
Together

Intel Participation in the OpenStack* Community

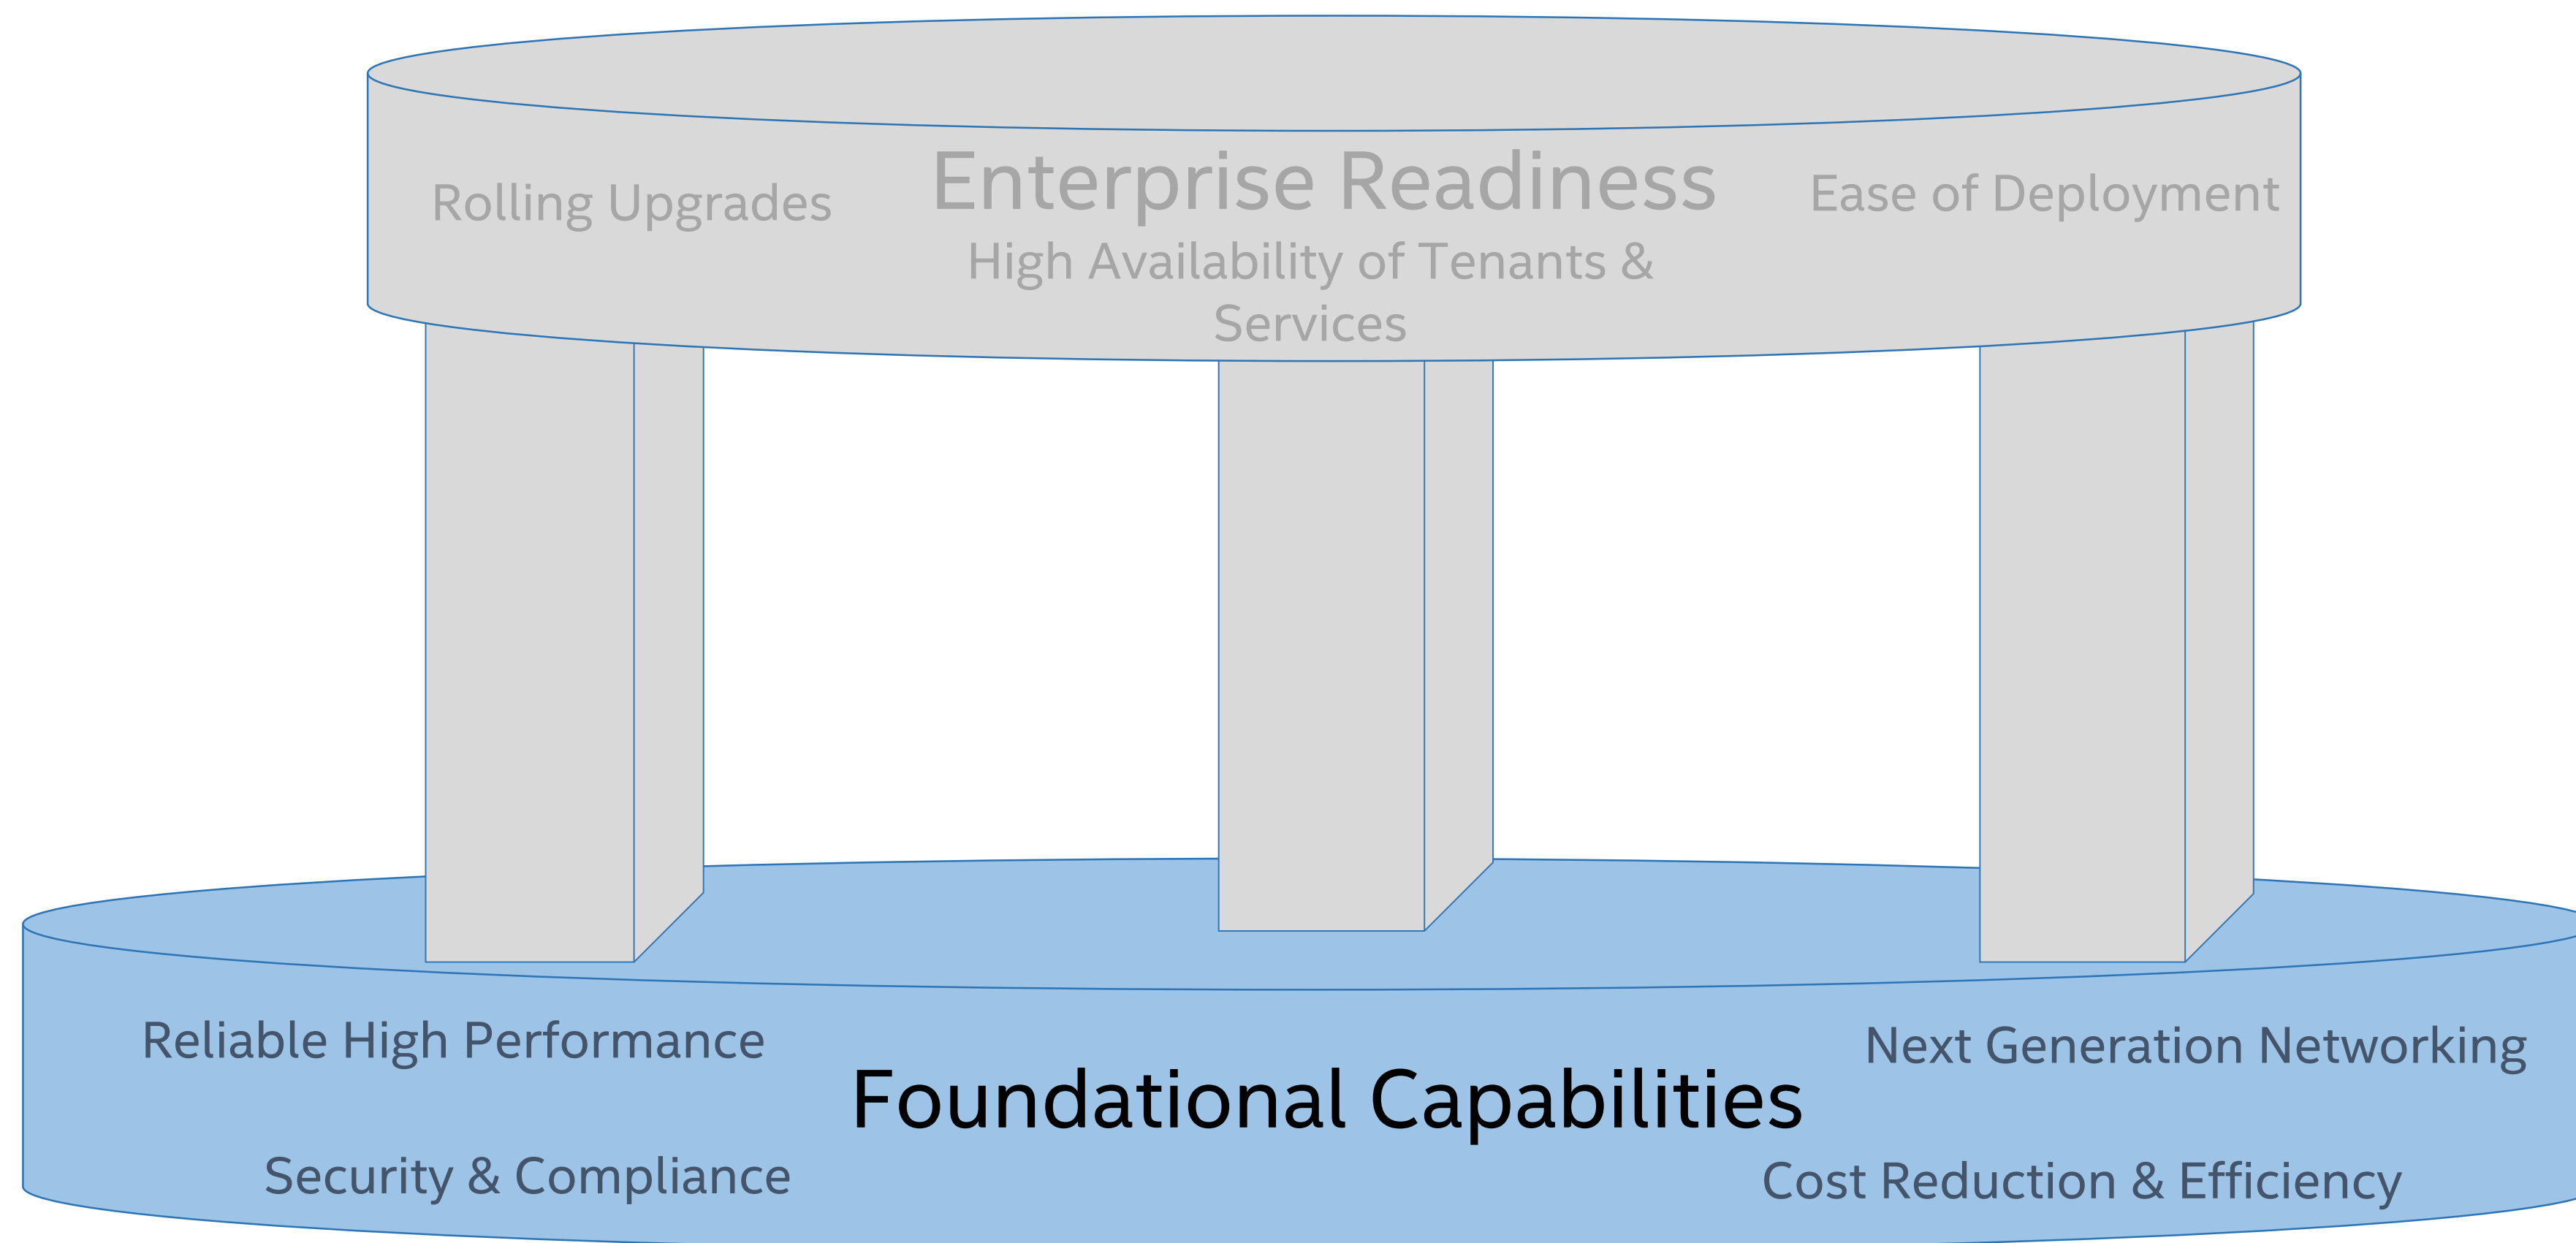
- Platinum member
- Important role in the Enterprise Workgroup
- Contributing across most major projects today (Top-10 across most modules)
- Working with the ecosystem to drive adoption



*Other names and brands may be claimed as the property of others.

Facilitating Enterprise OpenStack Adoption

Intel and Red Hat are key participants in the Enterprise Workgroup at the OpenStack Foundation



Addressing Foundational Capabilities



Strong Security and Compliance

- Trusted compute pools (TCP)
- Encryption and data residency



Efficiency and Low Cost

- Power-aware scheduling and high storage utilization



Reliable High Performance

- Intelligent scheduling and advanced monitoring



Next-Generation Networking

- Software-defined networking (SDN)
- Network functions virtualization (NFV)

Strong Security and Compliance

Available today with RHEL-OSP 6

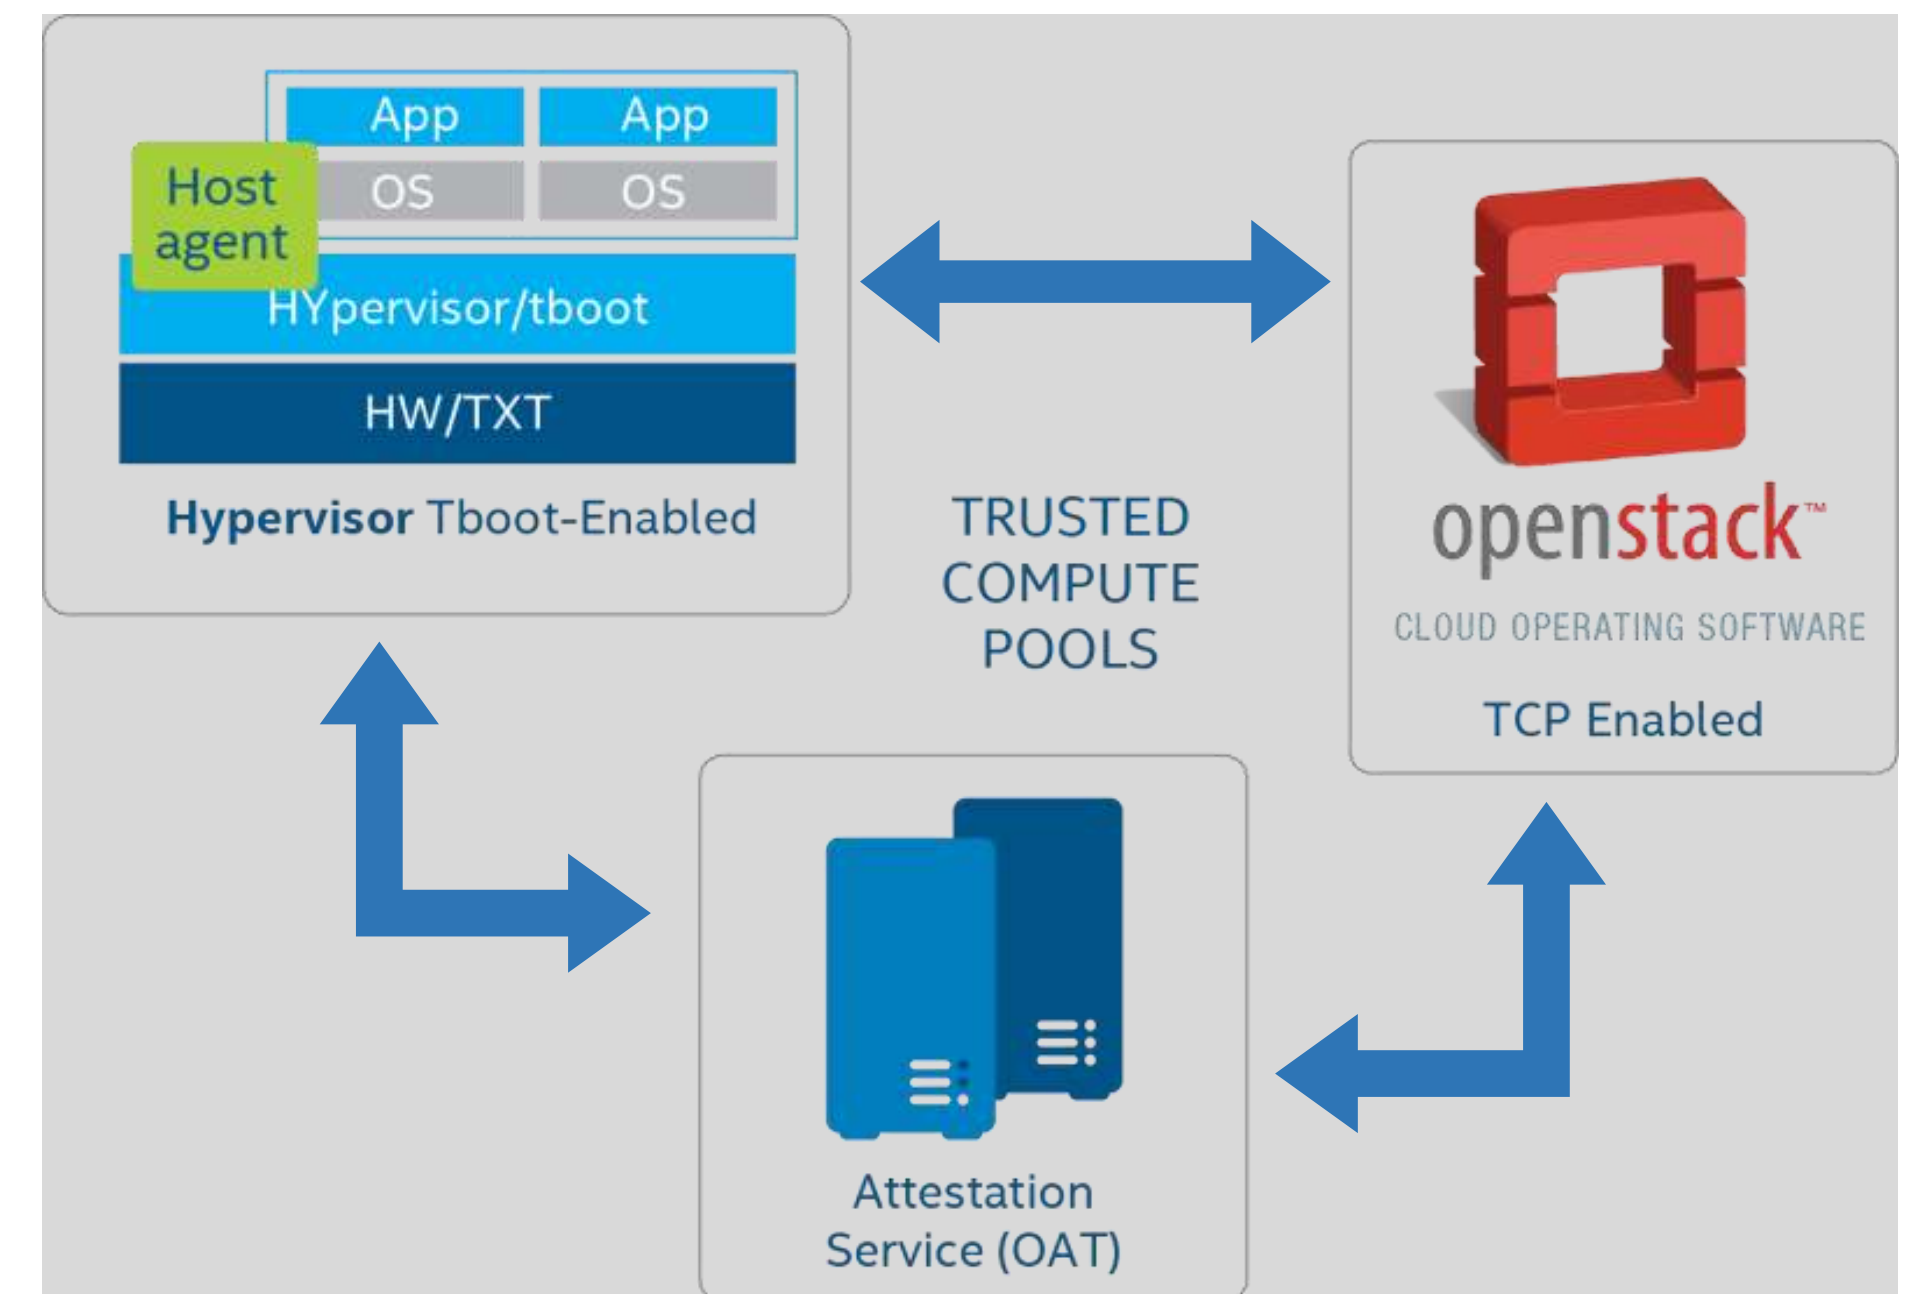
- Launch security-sensitive workloads only on pools of trusted virtualized servers
- Launch encryption-heavy workloads on platforms with appropriate technology (eg. Intel AES-NI, Intel® Quick Assist)

Targeting a future release of RHEL-OSP

Launch sensitive workloads onto trusted bare-metal servers

Encrypt workloads for launch only on targeted servers

Launch workloads only on servers within a geographical or organization boundary



Deploy security-sensitive workloads on trusted servers

¹ Intel technologies' features and benefits depend on system configuration and may require enabled hardware, software or service activation. Performance varies depending on system configuration. **No computer system can be absolutely secure.** Check with your system manufacturer or retailer or learn more at [intel.com](https://www.intel.com).

Reliable High Performance

Available today on RHEL-OSP 6

- Launch applications on servers with the appropriate capabilities and headroom

Coming soon in RHEL-OSP 7


- Schedule workloads based on server CPU usage, and power and thermal status

Targeting a future RHEL-OSP release

- Schedule workloads based on additional parameters like server memory and cache usage

THE OPENSTACK* NOVA SCHEDULER:

1. Evaluates available servers
2. Identifies servers that meet app requirements
3. Schedules new VM on the best-fit server



Deploy demanding applications on best-fit servers

Intel technologies exposed: Intel® AES-NI¹, Intel® AVX 2.0², Intel® QuickAssist, Intel® Quick Sync Video, Intel® Xeon Phi™ co-processors

^{1,2} Intel technologies' features and benefits depend on system configuration and may require enabled hardware, software or service activation. Performance varies depending on system configuration.

No computer system can be absolutely secure. Check with your system manufacturer or retailer or learn more at [intel.com](https://www.intel.com).

*Other names and brands may be claimed as the property of others.

Efficiency and Low Cost

Coming soon in RHEL-OSP 7

- Policy-based control of Swift* object storage for enhanced control and utilization
- Erasure coding for Swift to reduce storage capacity requirements by up to 50 percent¹
- Power- and thermal-aware scheduling for optimizing cost and reliability




Visit us at Intel Booth #905 to learn more about Intel & Red Hat joint work on storage

¹ Erasure coding (EC) algorithms can be configured to provide equal or better data durability than triple RAID data redundancy while using up to 50% less storage. Claim based on internal Intel measurements on usable capacity of 320 drives totaling 960 TB of raw capacity with no single point of failure to store 10 drives worth of data on 14 drives ($14n/10n = 1.4x$ versus $3x$ for a 53% reduction in storage requirements). Tests document performance of components on a particular test, in specific systems. Differences in hardware, software, or configuration will affect actual performance. Consult other sources of information to evaluate performance as you consider your purchase. For more complete information about performance and benchmark results, visit <http://www.intel.com/performance>

High-Performance Network Functions Virtualization (NFV) on Intel® Xeon® Processor-Based Servers

Targeting a future RHEL-OSP release

- Up to ~10x better packet performance and ~4x reduction in CPU utilization with Open vSwitch optimized with DPDK¹
- Low I/O latency using SR-IOV
- CPU pinning, NUMA-awareness, and more

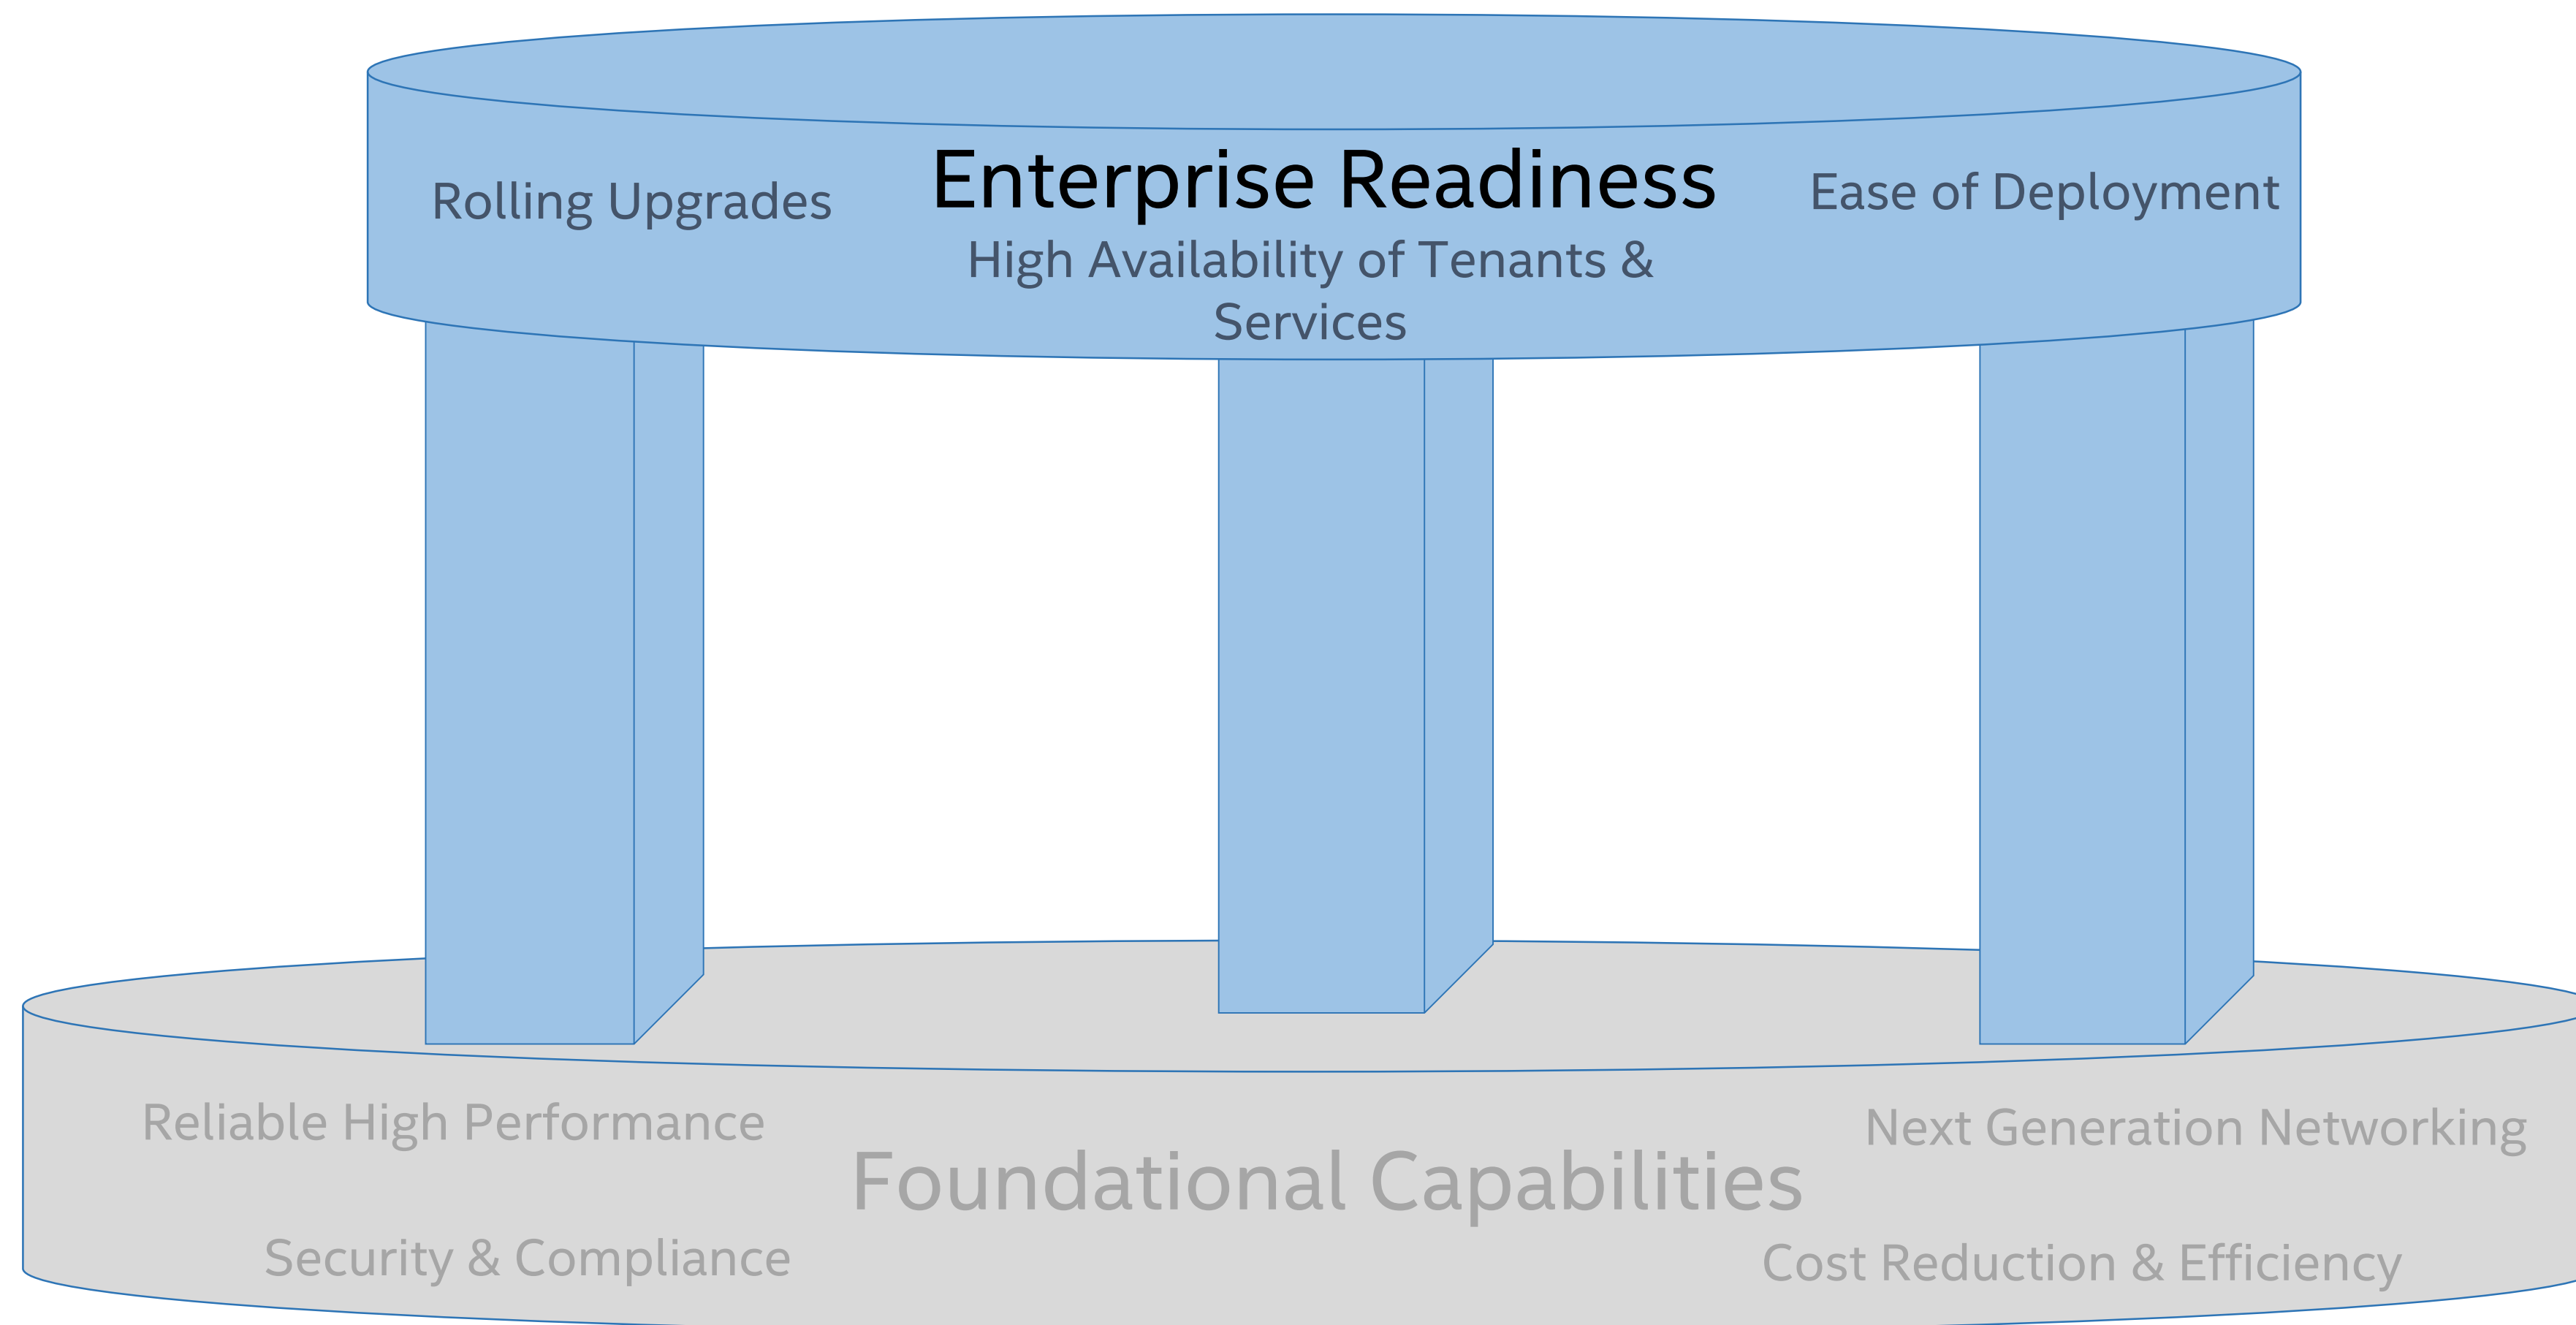


Low-latency and high throughput for OpenStack* networking

¹ Results based on Intel tests on a dual-socket server with Intel® Xeon® E5-2697 V3 processors at 2.6GHz with 32GB of DDR4-1067 RAM, 120GB SSD, Intel® QuickAssist Adapter 8950-SCCP with Intel® Communications Chipset 8950, running Fedora 20 x86-64. Full configuration details are in the document posted at https://01.org/sites/default/files/page/intel_onp_server_release_1.2_benchmark_test_report_v1.0.pdf. Tests document performance of components on a particular test, in specific systems. Differences in hardware, software, or configuration will affect actual performance. Consult other sources of information to evaluate performance as you consider your purchase. For more complete information about performance and benchmark results, visit <http://www.intel.com/performance>

Addressing Enterprise Readiness

Intel and Red Hat are key participants in the Enterprise Workgroup at the OpenStack Foundation

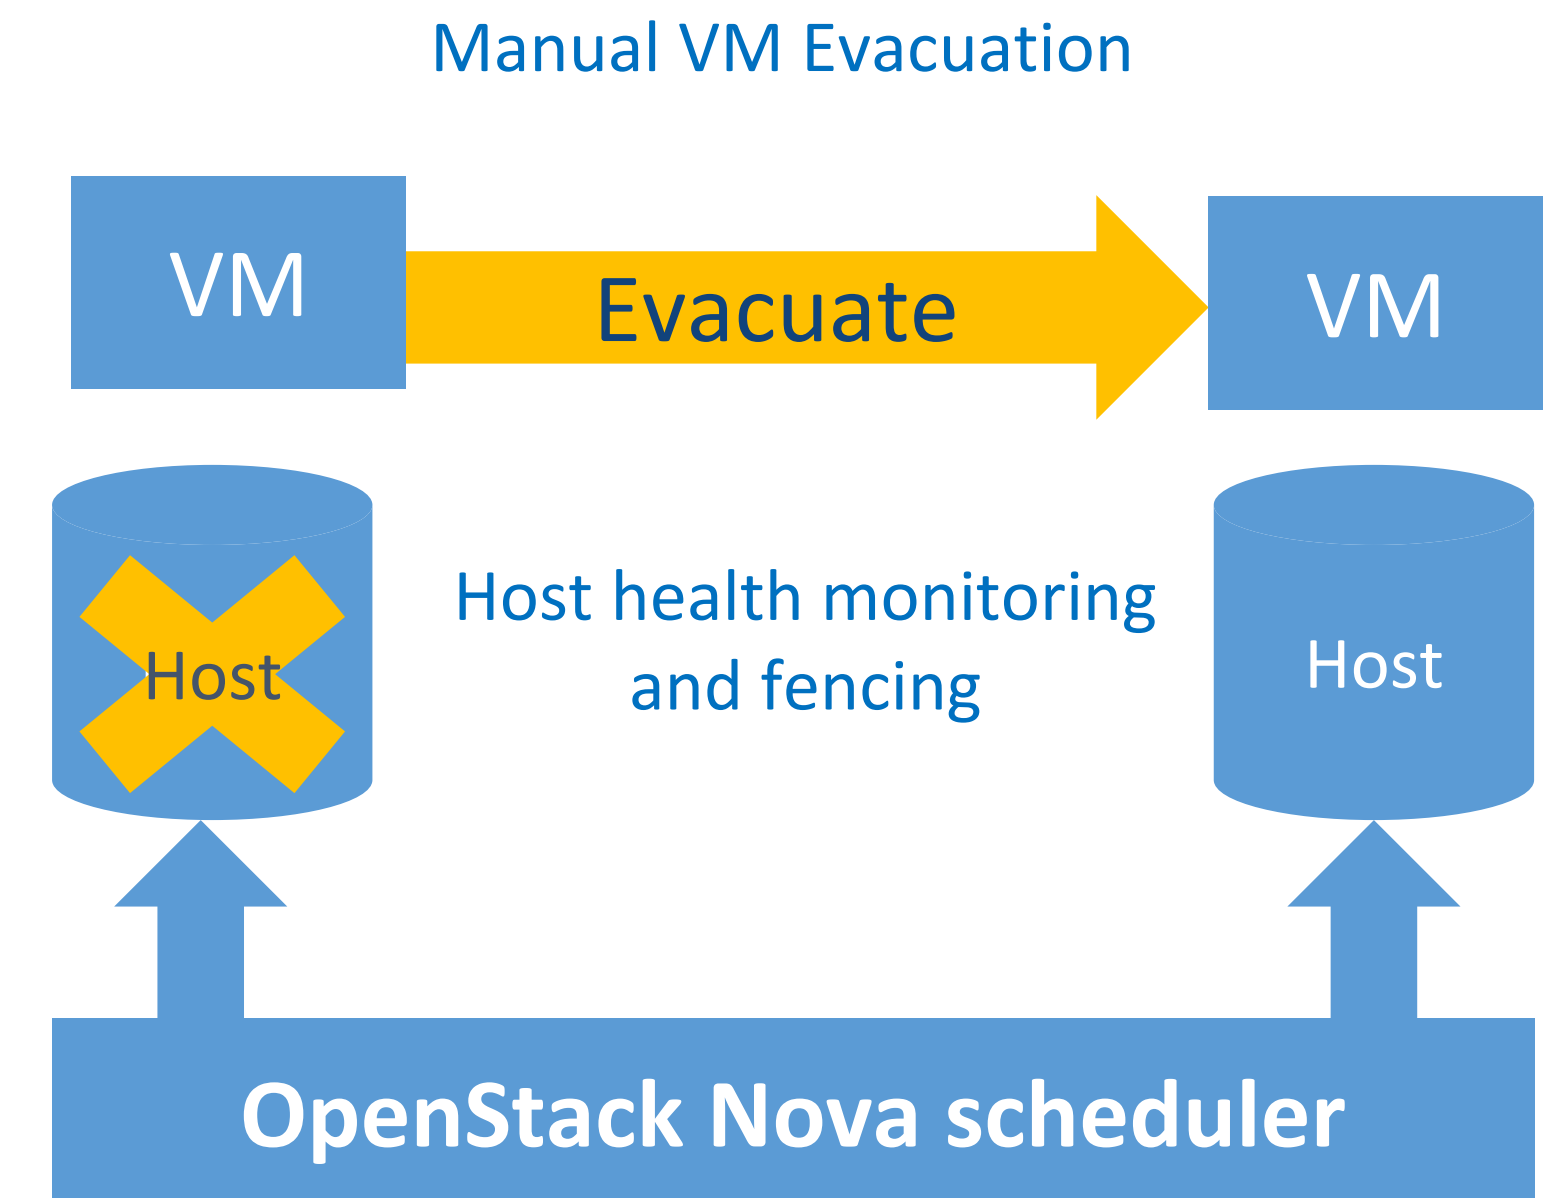


High Availability

Tenant High Availability

Coming in RHEL OpenStack Platform 7

- **Manual VM evacuation** - Ability to move all VMs and Data off of hosts as part of maintenance process.
- The feature implemented as part of the Horizon Dashboard/Console which utilizes a combination of Nova API calls that provide disabling of a Compute Node for scheduling new workloads as well as live and offline migration capabilities.
- Additional work needs to be done to provide a similar functionality directly by a single Nova API call.

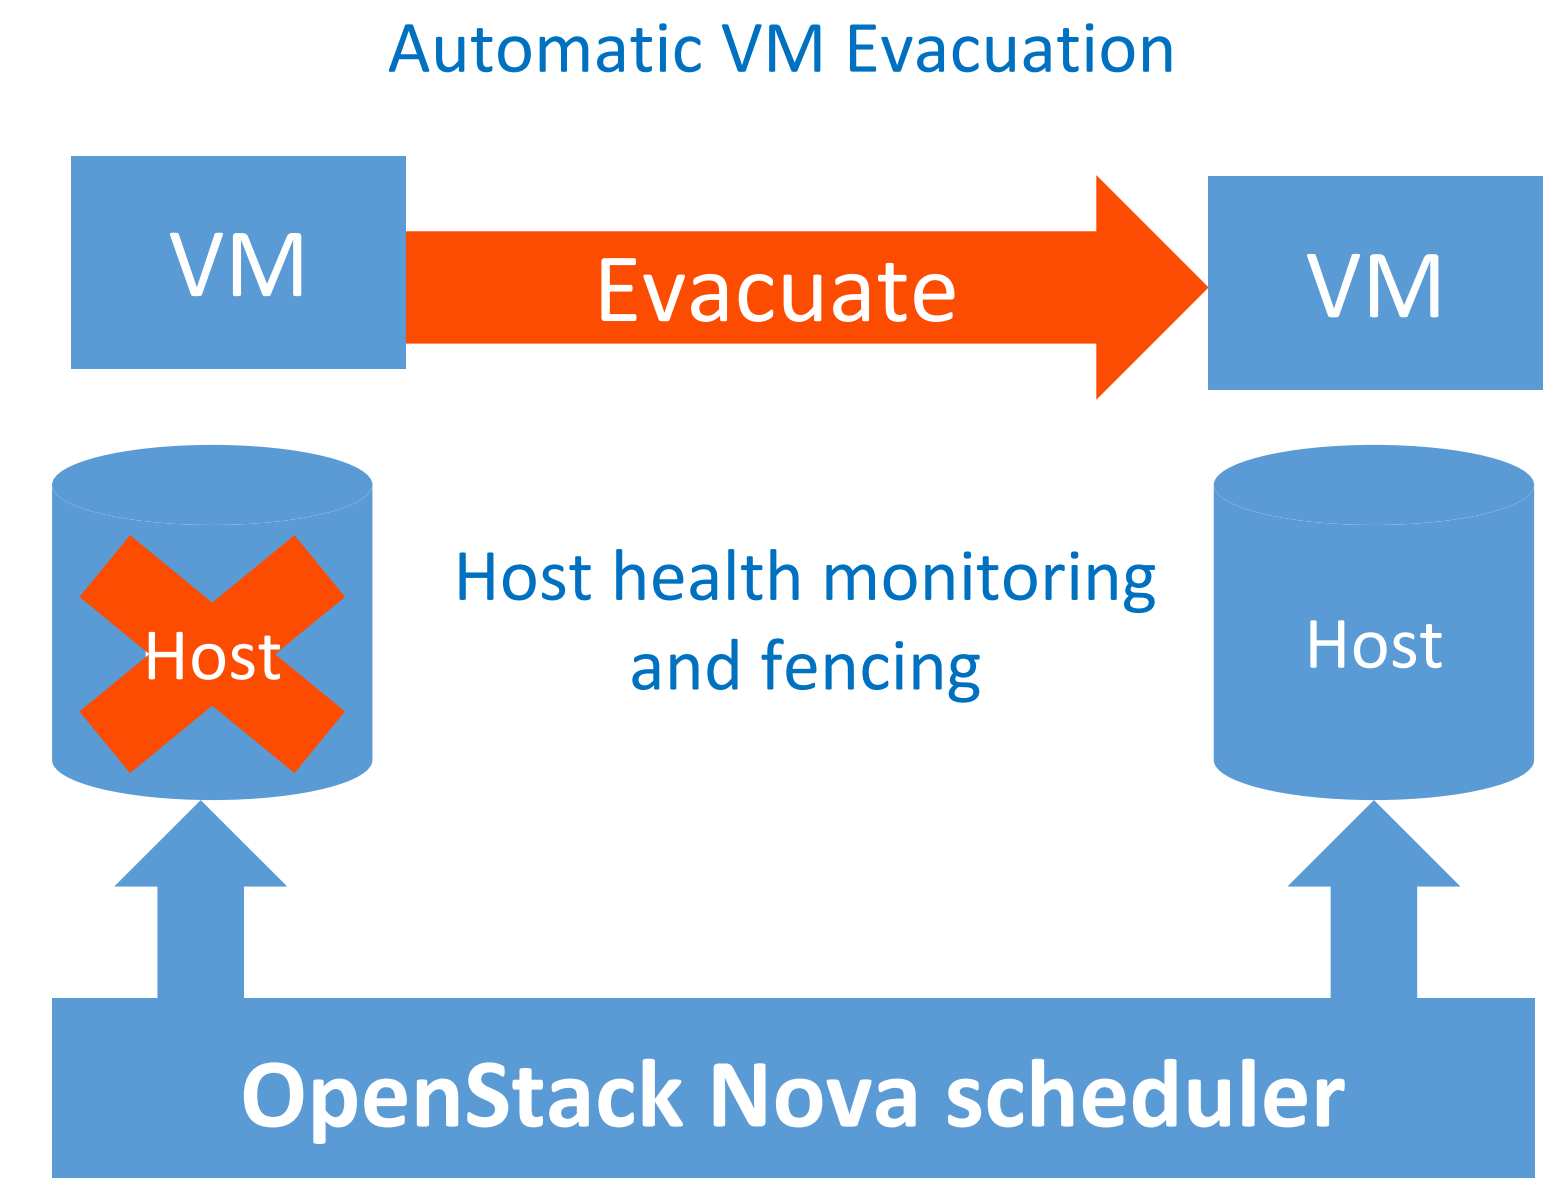


High Availability

Tenant High Availability

Coming in RHEL OpenStack Platform 7

- **Automatic Evacuation** - Added support to detect a failed Hypervisor and migrate the running instances on another Hypervisor (using Pacemaker fence agents)
- On-boarding of pet type workloads which do not monitor their availability themselves, should cater for 80% of the traditional enterprise workloads (pending re-architecting of such workloads to be more cloud-native)



High Availability

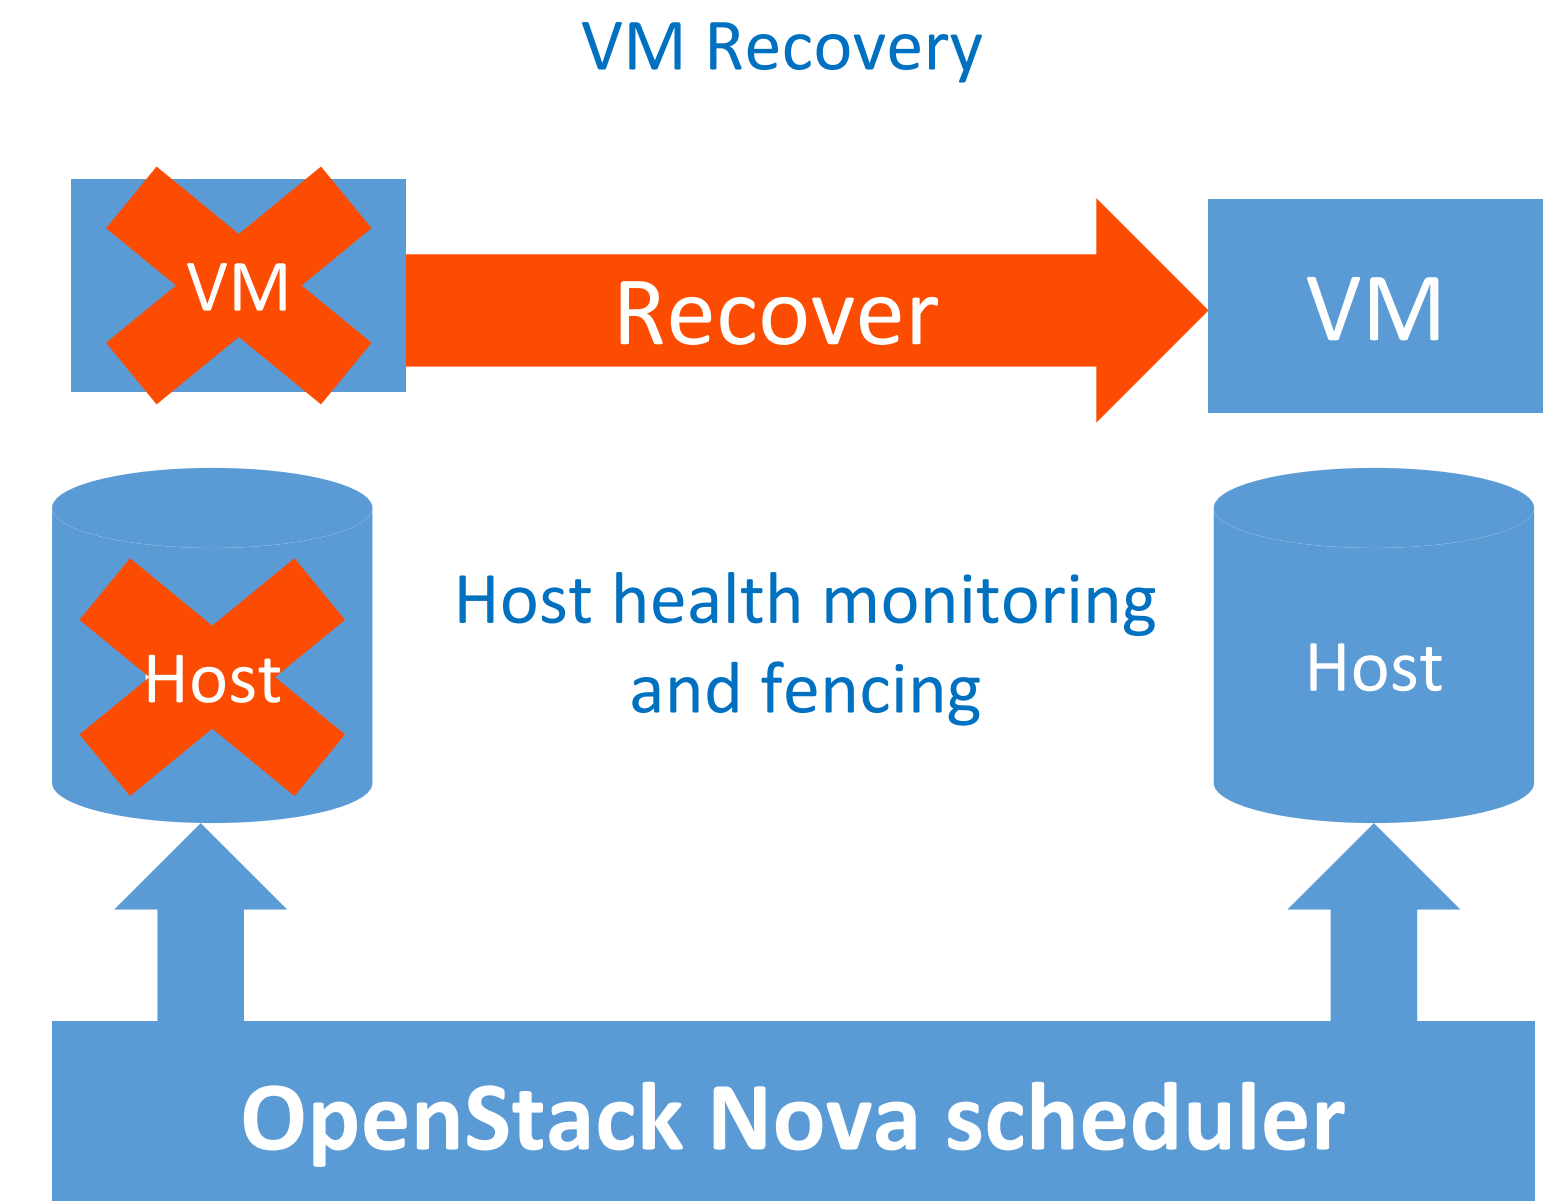
Tenant High Availability

Coming in RHEL OpenStack Platform 7

Automatic VM recovery on Host failure - VM rebuild on a working host due to hypervisor or host issues which have shut down the running VMs

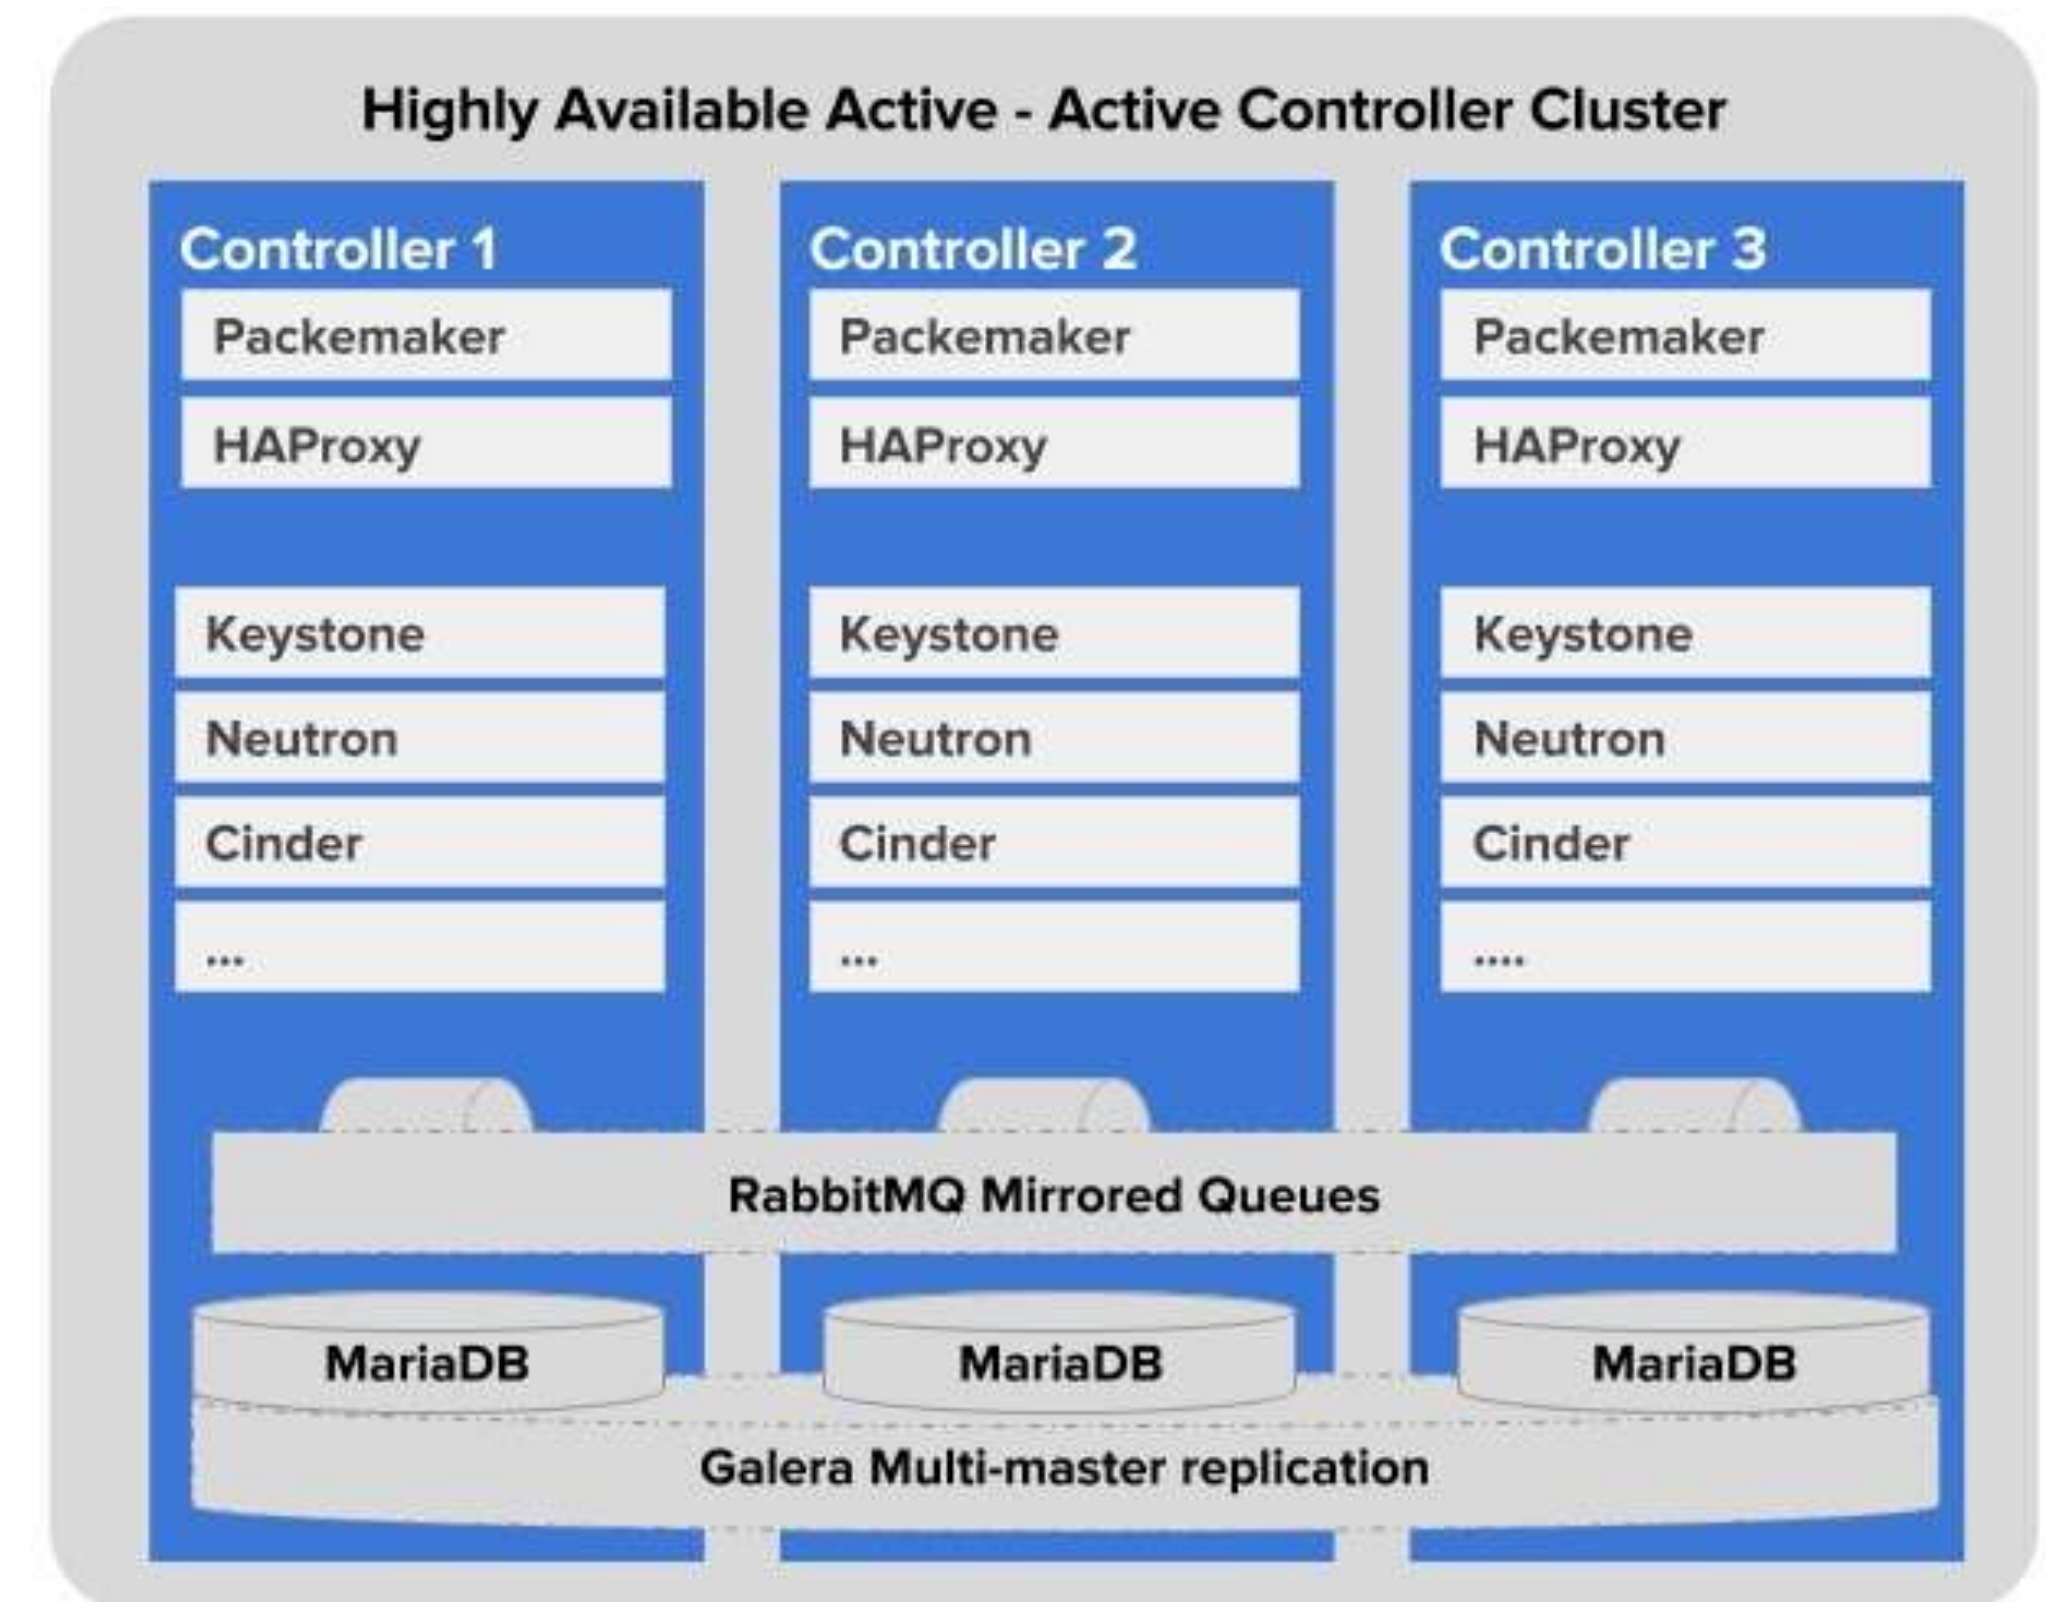
This feature is provided via tight integration between multiple Red Hat products, like RHEL, RHEL High Availability Add-On and RHEL OSP.

Demo info on summit



High Availability of Services

- **Higher availability of OpenStack APIs/Services**
 - At the moment all core OpenStack services can be setup as active/active with the exception cinder-volume.
- **RHEL OpenStack Integration with RHEL HA Add-On**
 - Two architectures: Pacemaker/HAProxy & Keepalived/HAProxy
 - Pacemaker - resource mgr for active/active & active/passive services
 - HAProxy - Load balancing active/active services
 - Keepalived - Load balancer availability



High Availability of Services

Coming in RHEL OpenStack Platform 8 and beyond

- Working upstream on enhancing cinder-volume to support active/active safely
- **Cinder Task Flow for managing create volumes tasks**
 - The improvements of state management can get us a step closer to Active-Active safe operations
 - Goal is to cover Cinder operations beyond volume creation tasks
- **Cinder State Enforcer**
 - Long standing work to improve Cinder volume's states management and reliability, and to improve failure tolerance.
 - In order to mitigate the concurrent resource access problems in Cinder, work was done in the last cycles to refactor the concept of a lock to be a set of allowed and disallowed state transitions (instead of acquiring local filesystem locks in the manager processes) by implementing a new `enforcer` model.

High Availability of Services

Coming in RHEL OpenStack Platform 8 and beyond

- Working upstream on enhancing cinder-volume to support active/active safely
- **Cinder Task Flow for managing create volumes tasks**
 - The improvements of state management can get us a step closer to Active-Active safe operations
 - Goal is to cover Cinder operations beyond volume creation tasks
- **Cinder State Enforcer**
 - Long standing work to improve Cinder volume's states management and reliability, and to improve failure tolerance.
 - In order to mitigate the concurrent resource access problems in Cinder, work was done in the last cycles to refactor the concept of a lock to be a set of allowed and disallowed state transitions (instead of acquiring local filesystem locks in the manager processes) by implementing a new `enforcer` model.

Rolling Upgrades

- **Upgrade a running OpenStack Cloud with no Tenant and API downtime**
- Perform most of the upgrade process of the cloud without downtime of the control interface (API and Web)
 - Versioned Objects - Initial work was done in Kilo/RHEL OpenStack Platform 7 for Nova to assist with rolling upgrades, simplify OpenStack cloud evolution and reduce planned downtime
 - Additional modules to be converted in Liberty and beyond.
 - Rolling upgrades w/ no downtime cannot be really achieved until projects like Neutron and Cinder are at least converted.
 - In addition to versioned object support, online schema migrations are also needed. Nova has converted to expansion only schema model, which supports this.

Rolling Upgrades

Coming in RHEL OpenStack Platform 8 and beyond

- **Cinder Objects**
 - Supporting rolling upgrades by using versioned objects.
 - These objects are isolated from the schema and contain the required information for communications and operations.
 - Work started in Kilo in Cinder (already there in compute with Nova objects)

Ease of Deployment

Deployment is more than the initial OpenStack installation

- It also deals with deployment lifecycle of production clouds that includes both updates, upgrades and capacity adjustment

New with RHEL OpenStack Platform 7


- **RHEL-OSP Director**
 - OpenStack Deployment & Management tool based on OpenStack TripleO project and inspired by SpinalStack project.



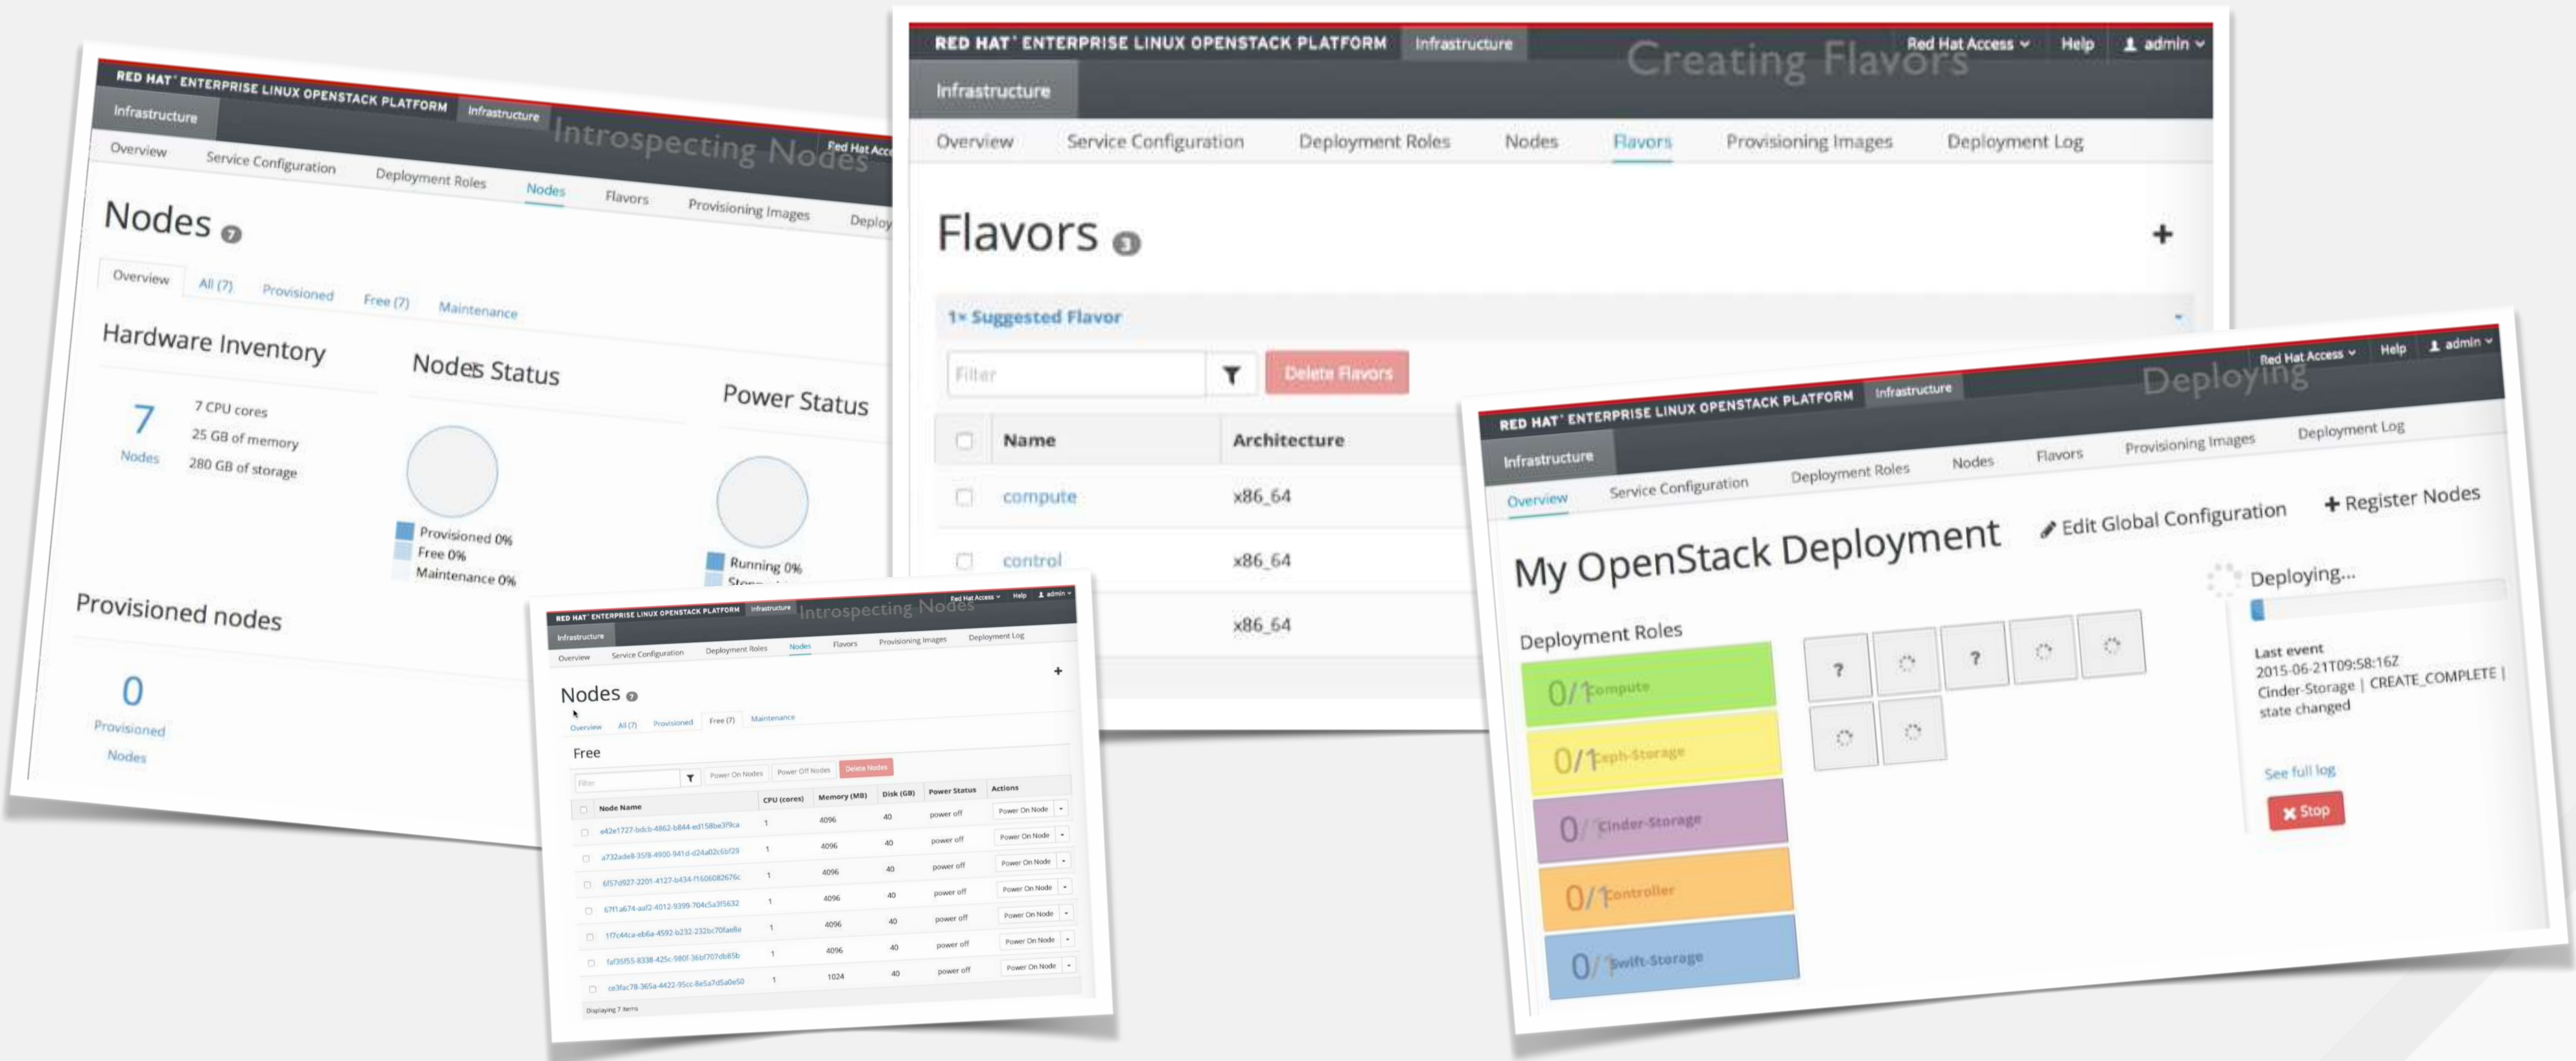
Ease of Deployment

Focus on:

- Provides tools identified by OpenStack operators in production to control and debug
- Provides introspection API to management tools
- Add/Remove capacity (compute/storage nodes)
- Provides control API to any external tool
- Automated orchestrated upgrades from 7.0 onwards

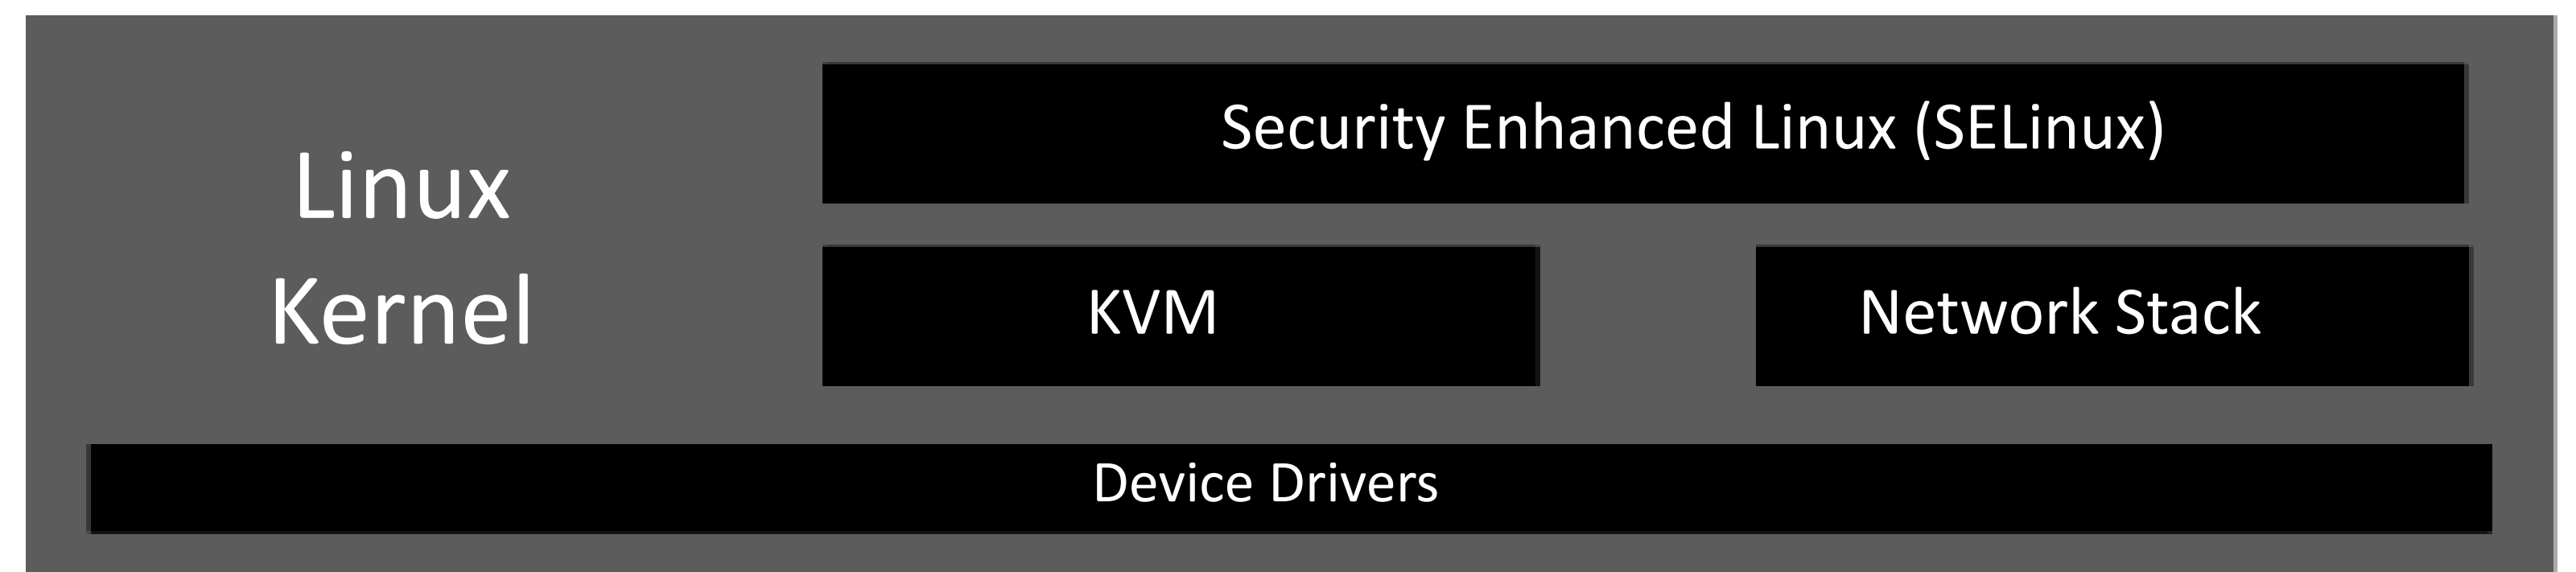
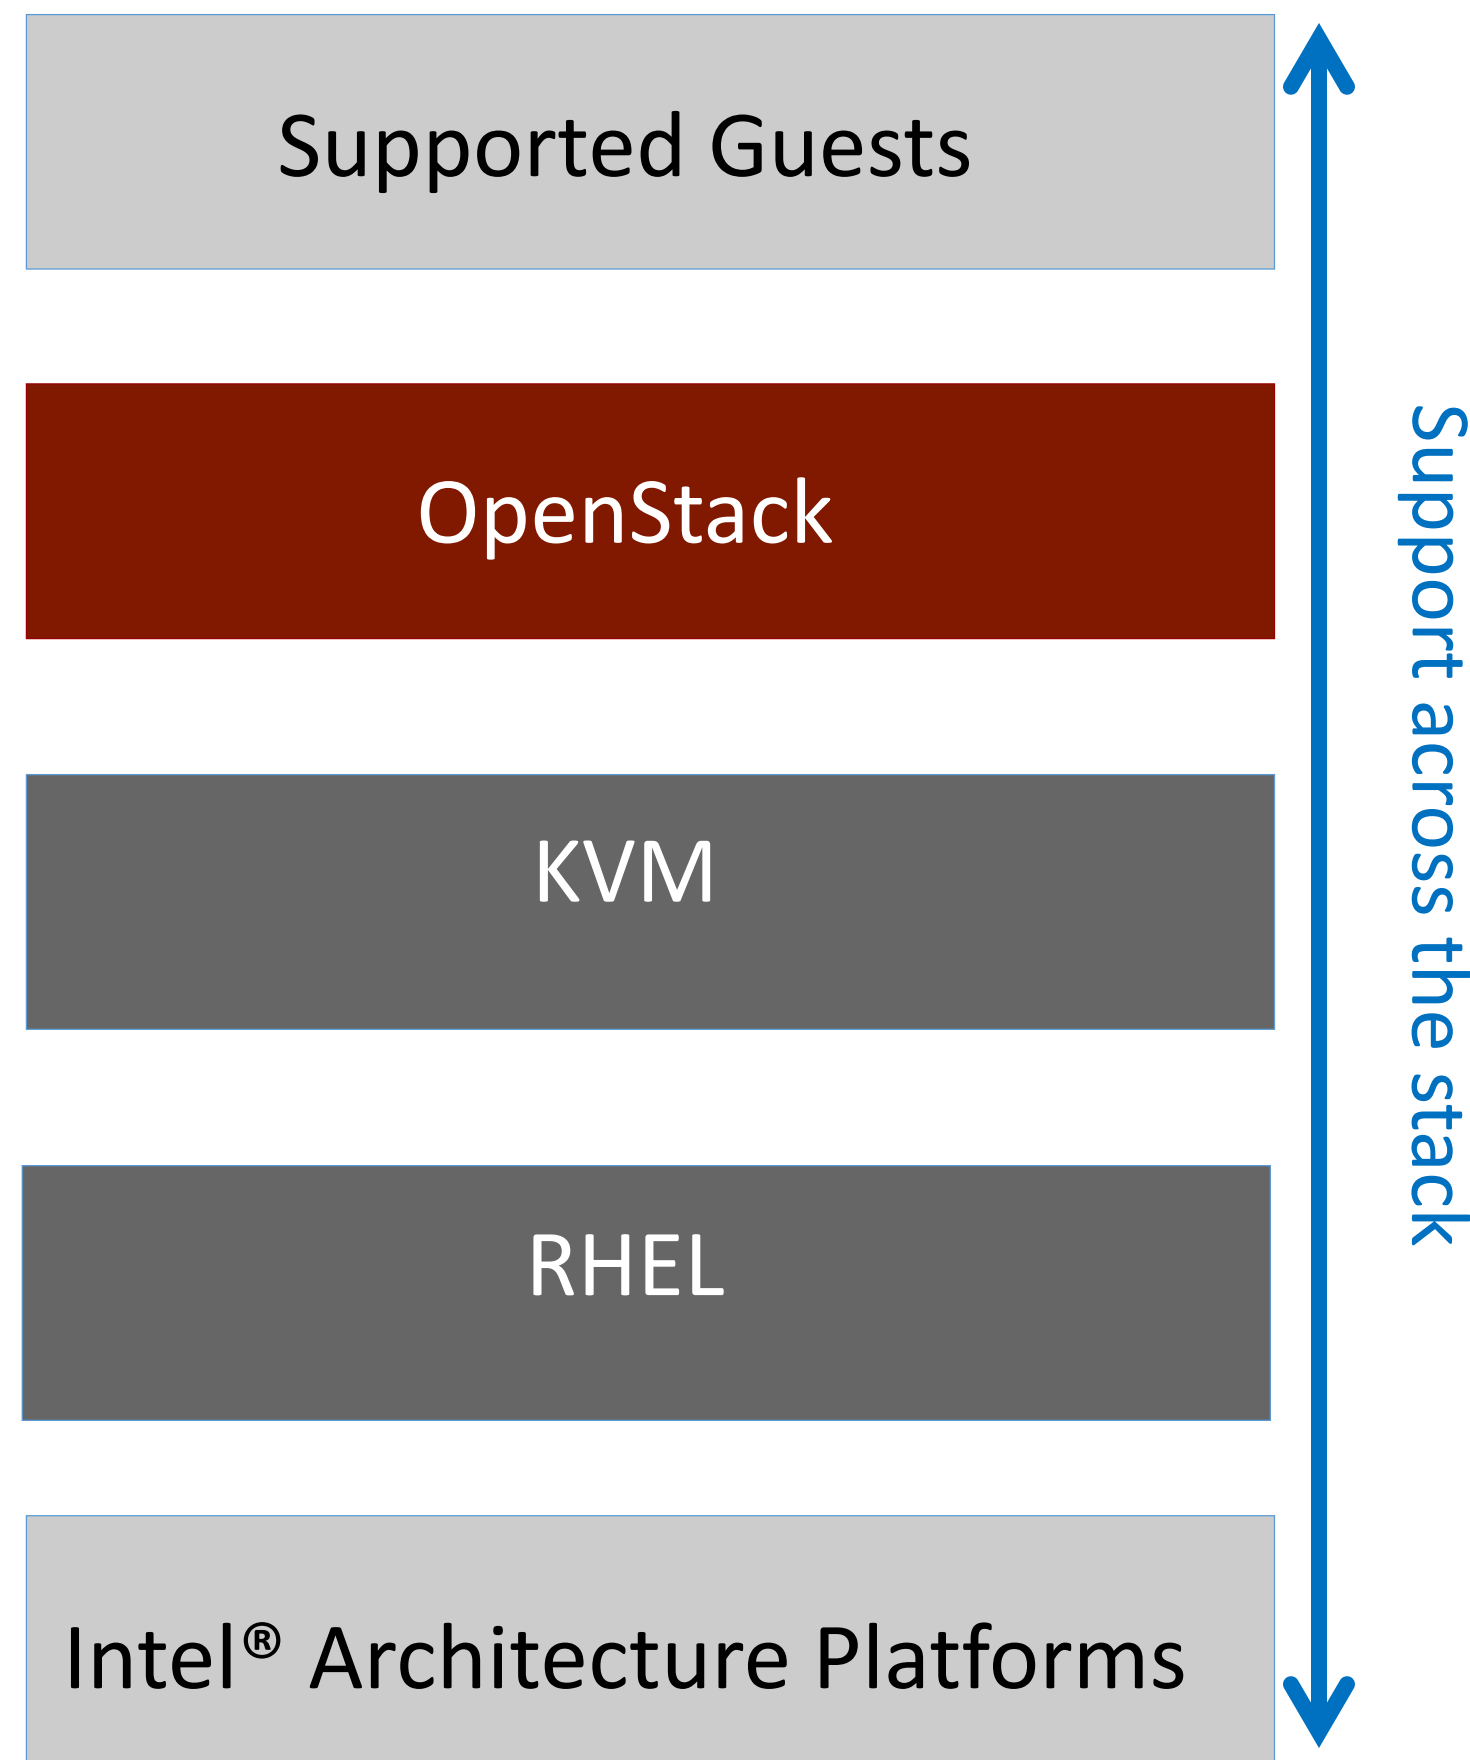


Ease of Deployment



RHEL OpenStack Platform Co-engineered & Integrated

- OpenStack cannot be productized as a stand alone layer!
- Red Hat Enterprise Linux OpenStack Platform is purposely designed with the recognition of the unique dependencies OpenStack has on the underlying Linux it's installed on.
- Engineered together with RHEL, KVM & Intel® Architecture platforms



The Vision of a
Software Defined
Infrastructure

Intel and RHEL
OpenStack
Platform

Intel & Red Hat:
Delivering
Solutions
Together

Intel and Red Hat: Long, Deep History of Innovation and Stability

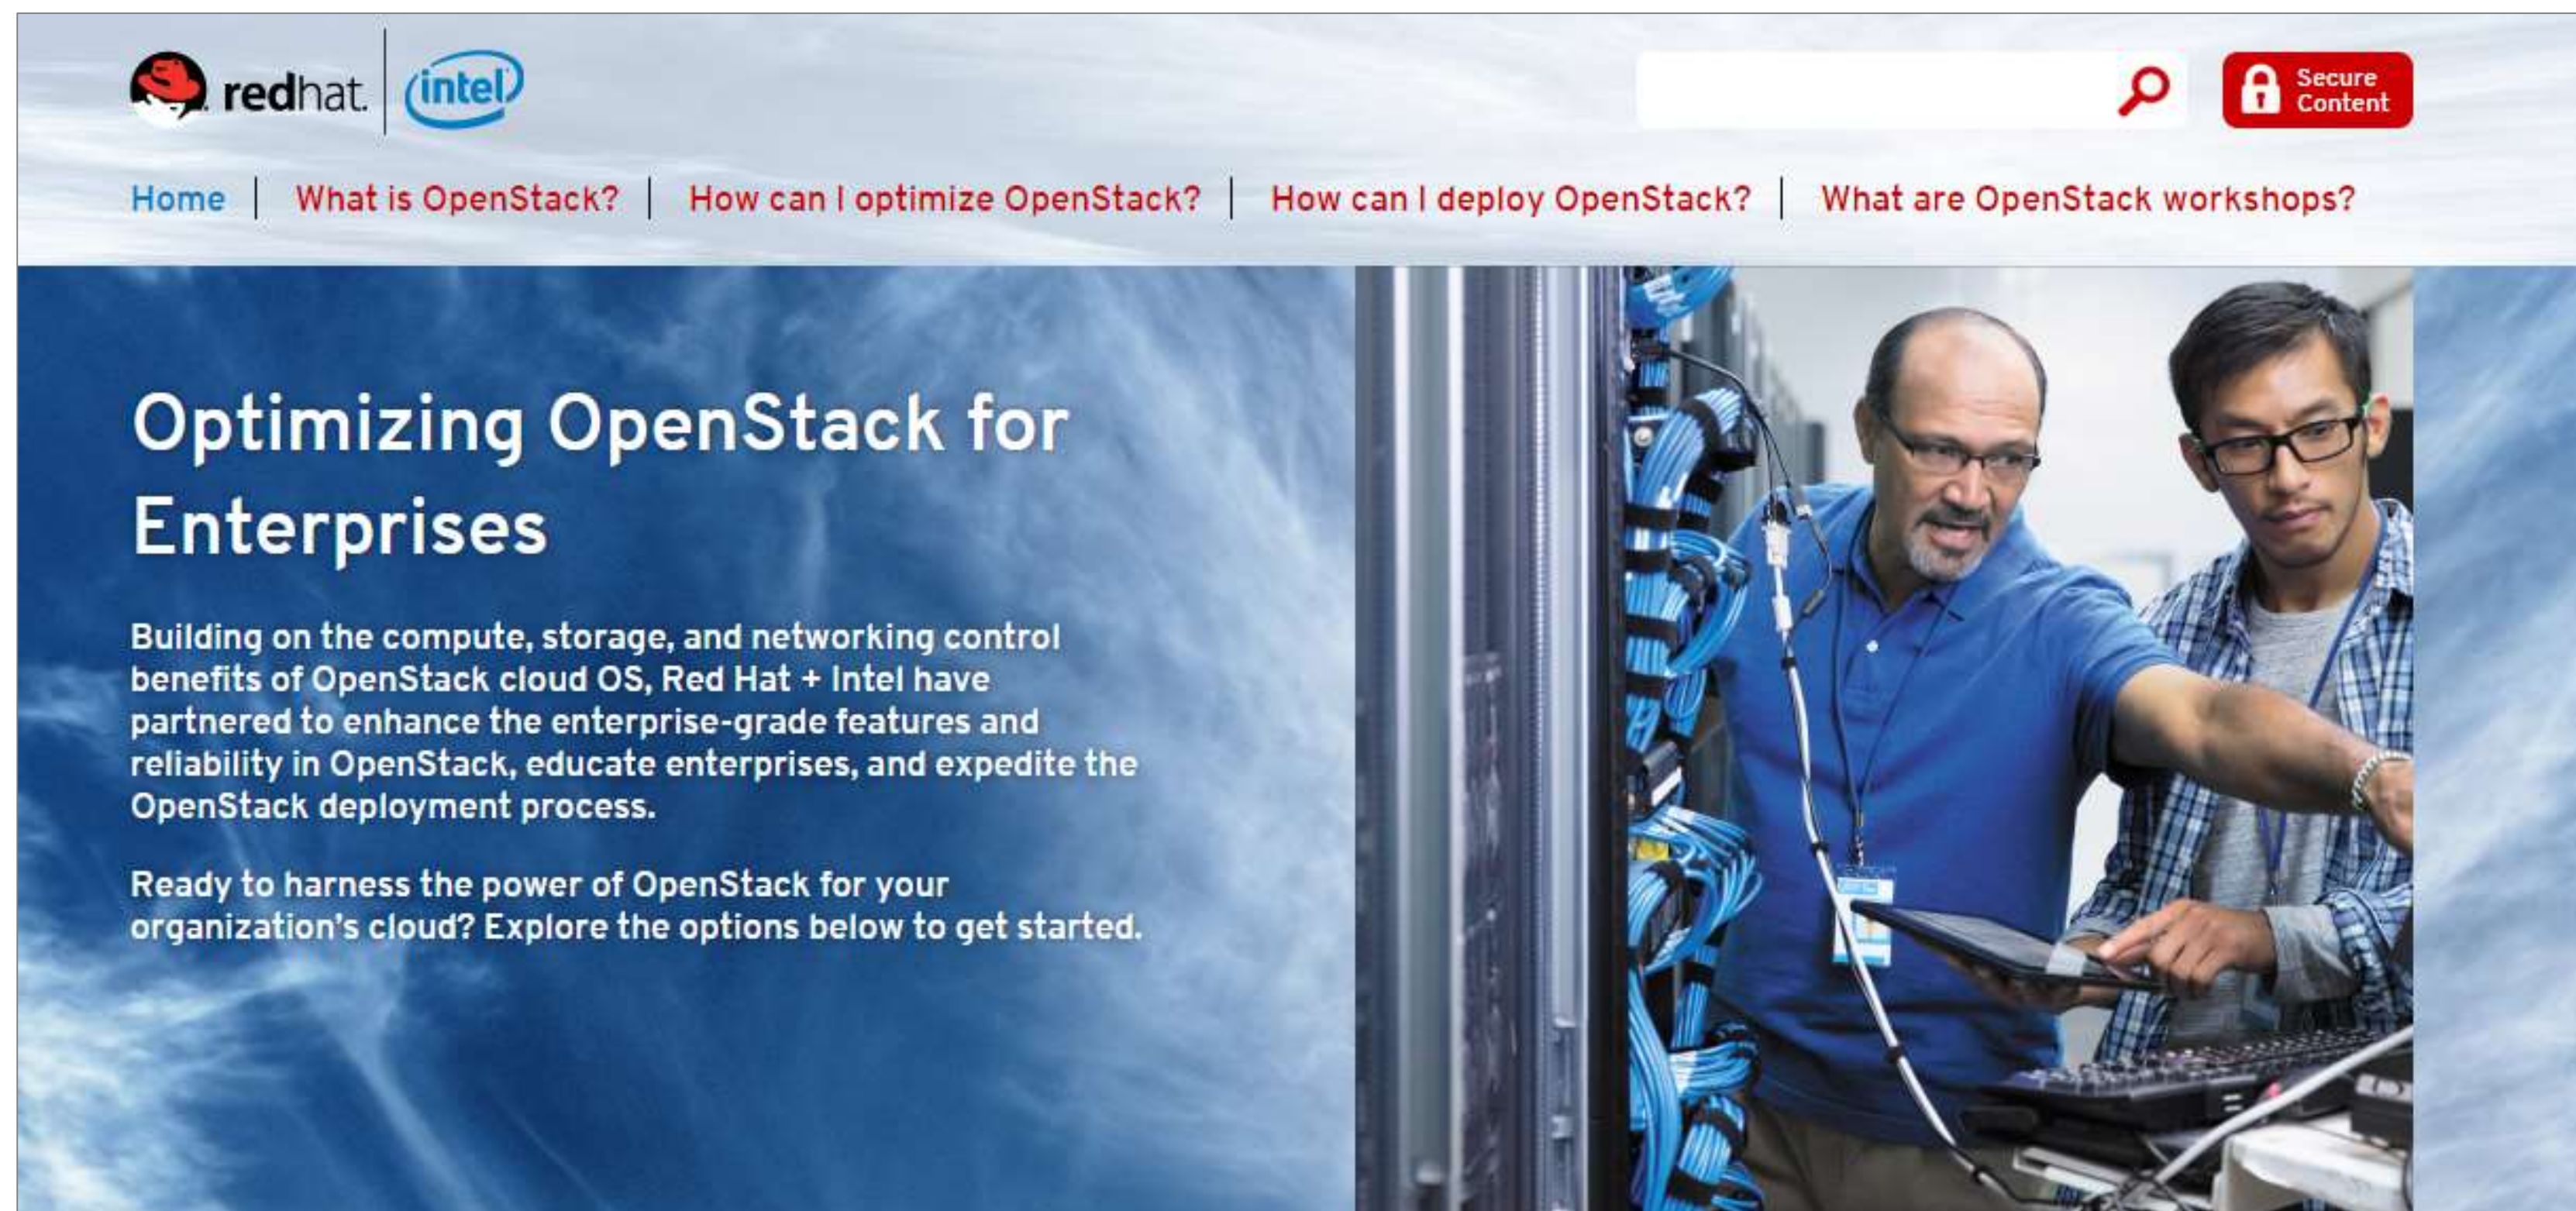
- **20+ years** of collaborative history
- **Focal innovators** at a deep level
- Common **strategic vision**

Driving change and innovation by contributing to open source to deliver the best datacenter platforms for your business

- We make technology safe
- We are trusted innovation engines: #1 and 2 top commercial contributors to Linux*
- We offer more choices with greater savings

OnRamp to Enterprise OpenStack*

Hundreds of customers educated through workshops to speed up the deployment process (see www.onramp2thecloud.com)



The screenshot shows the top section of a website. At the top left are the Red Hat and Intel logos. To their right is a search bar and a red button labeled 'Secure Content'. Below this is a navigation menu with links: 'Home', 'What is OpenStack?', 'How can I optimize OpenStack?', 'How can I deploy OpenStack?', and 'What are OpenStack workshops?'. The main banner features a blue background with a cloud-like pattern. On the left, the text reads: 'Optimizing OpenStack for Enterprises', followed by a paragraph: 'Building on the compute, storage, and networking control benefits of OpenStack cloud OS, Red Hat + Intel have partnered to enhance the enterprise-grade features and reliability in OpenStack, educate enterprises, and expedite the OpenStack deployment process.' Below this is another line: 'Ready to harness the power of OpenStack for your organization's cloud? Explore the options below to get started.' On the right side of the banner is a photograph of two men, one older and one younger, both wearing glasses and looking at a tablet held by the older man. They are standing in front of server racks with blue cables.

Related Red Hat Summit Sessions

- Red Hat Enterprise Linux OpenStack Platform enterprise capabilities, BoF session, Wednesday, June 24 6:00 pm - 7:00 pm
- On-Ramp to Enterprise OpenStack - VM High Availability Live Demo
 - Wednesday, June 24, 10am-2pm
 - Thursday, June 25, 10am-2pm
 - Location: next to Dell Booth#207
- OpenStack with Cisco & Red Hat, Wednesday, June 24 3:40 pm - 4:40 pm
- Dell and Red Hat's OpenStack journey to enterprise, Thursday, June 25 4:50 pm - 5:50 pm



RHEL OpenStack Platform and Intel® Architecture – Better Together

BROAD ENABLED ECOSYSTEM

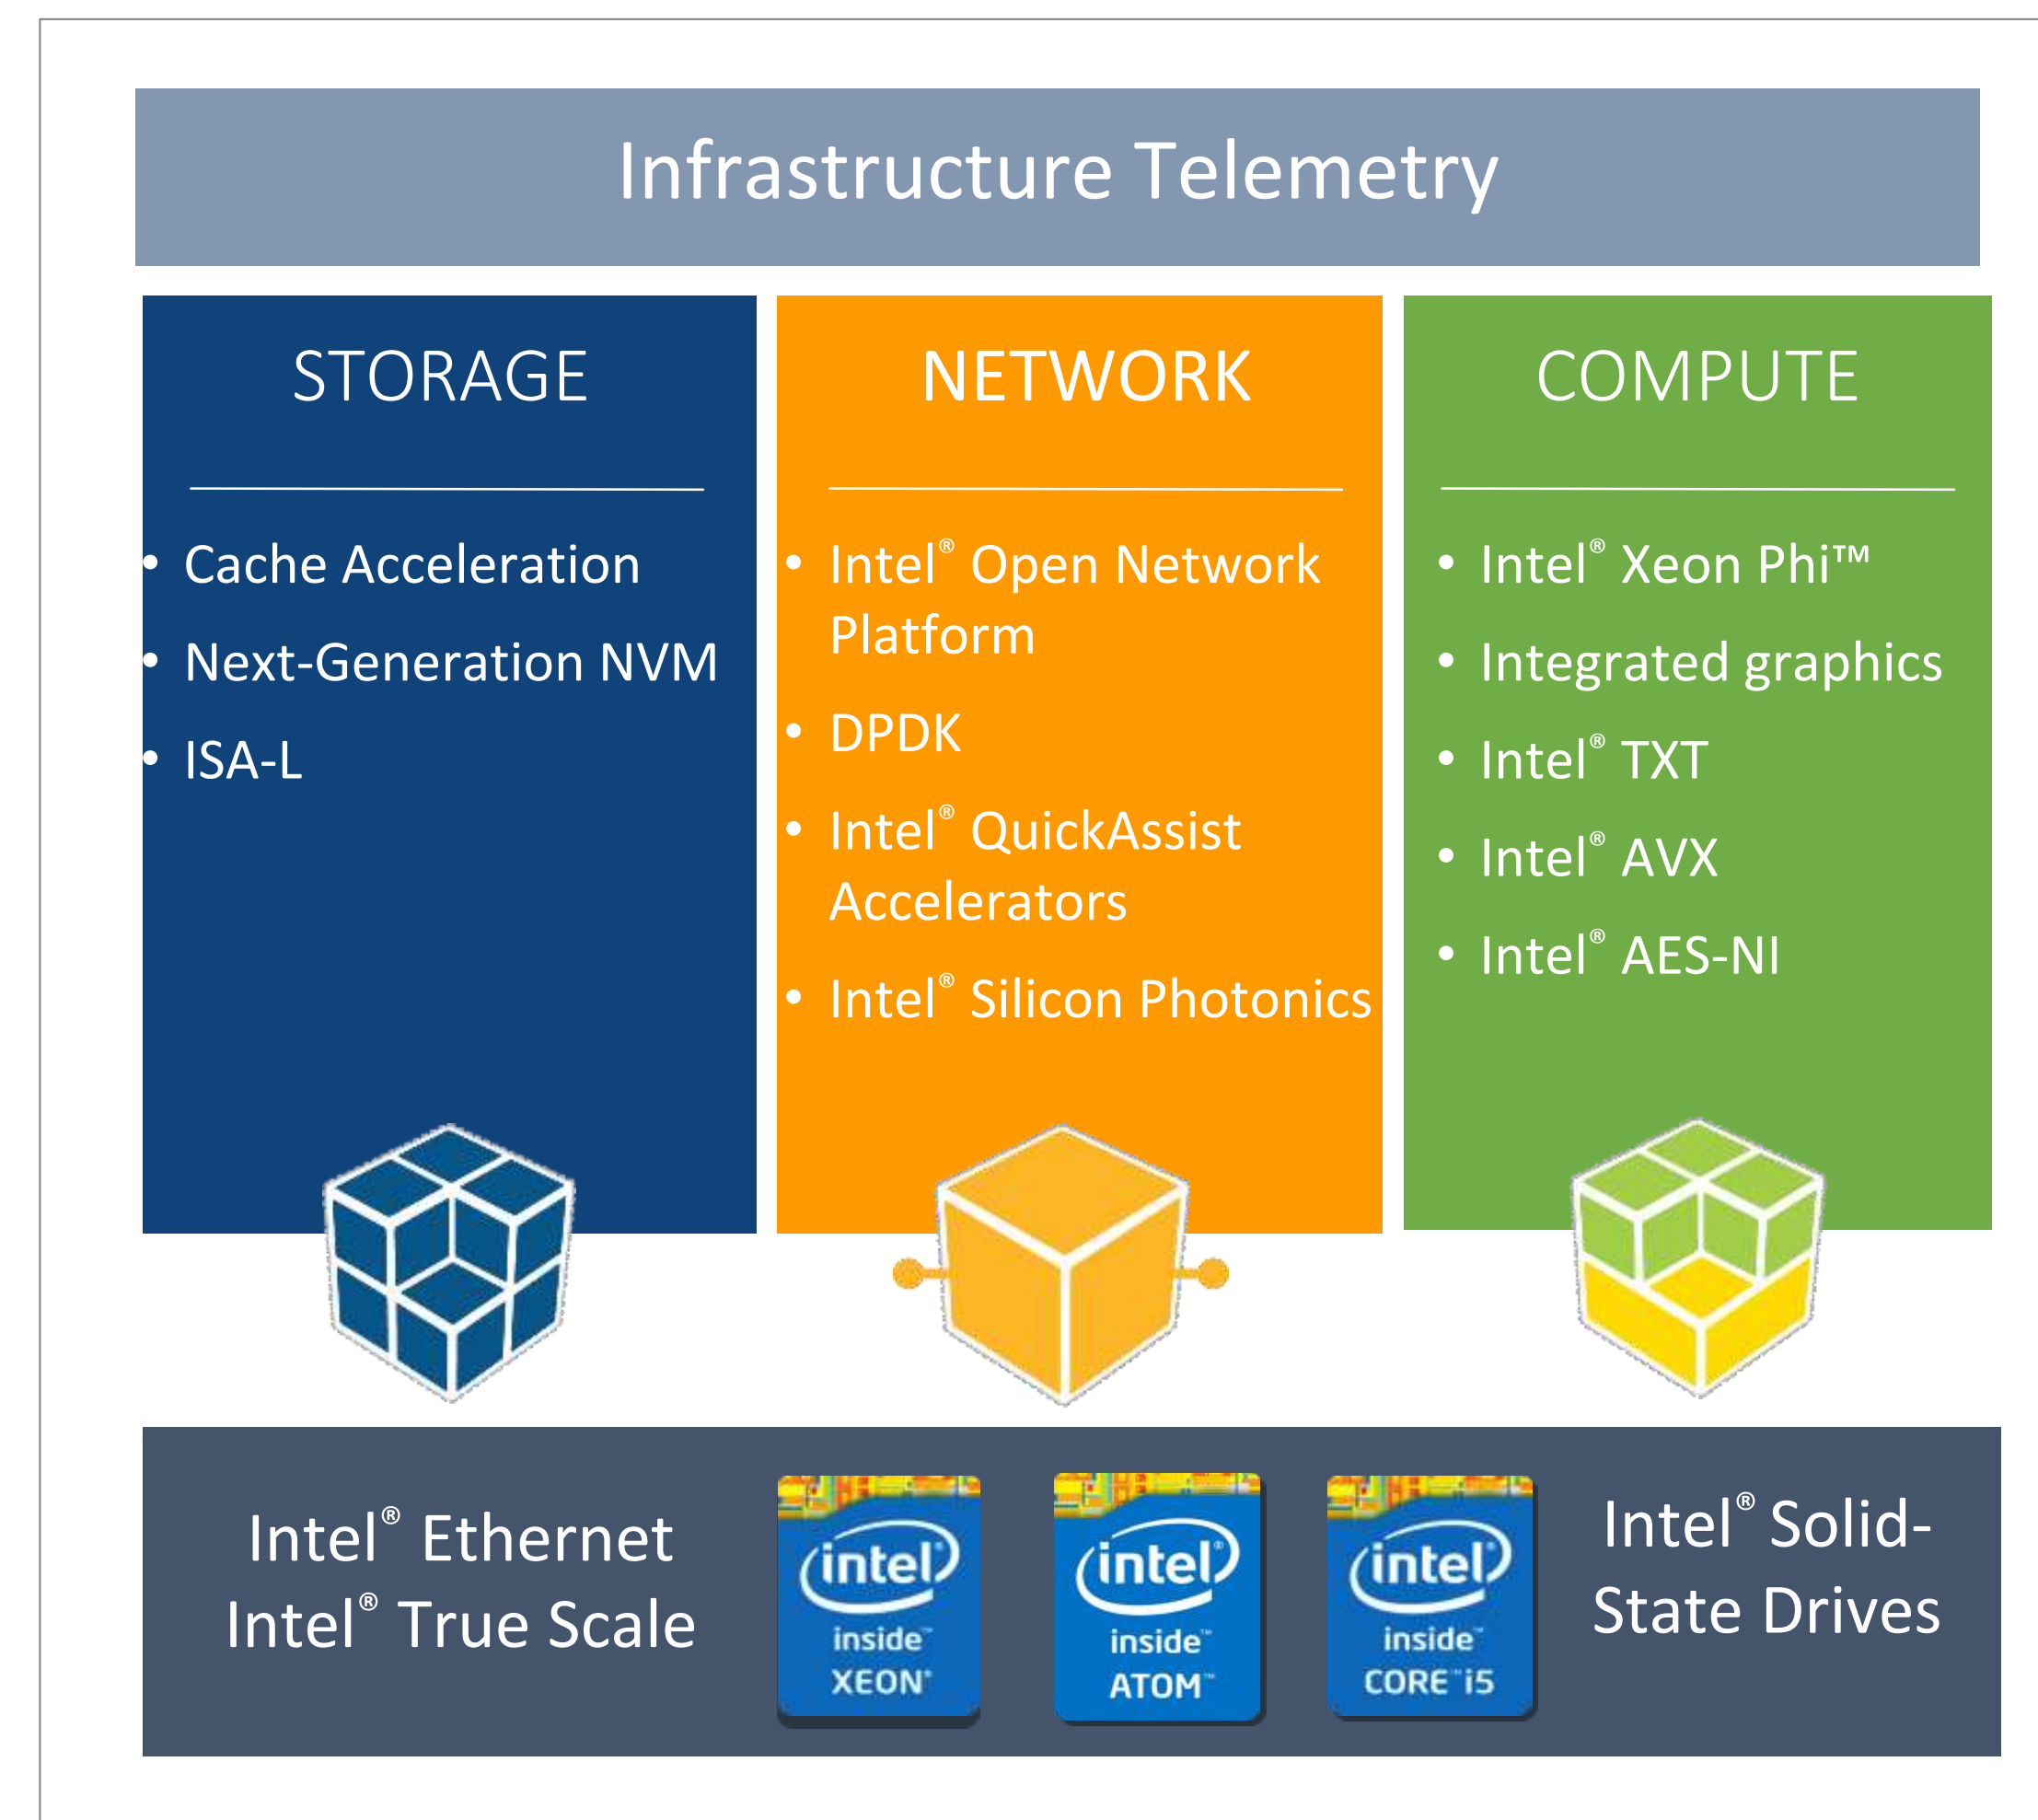
Integrated and optimized for all leading commercial and open-source application environments

EXPOSED AND INTEGRATED TELEMETRY

Hardware and infrastructure attributes are exposed and integrated with orchestration software for deeper insight and optimal provisioning management

PLATFORM AND ARCHITECTURAL LEADERSHIP

Standards-based compute, network, and storage building blocks drive maximum infrastructure efficiency and flexibility



Fulfilling the vision of a Software Defined Infrastructure

In Closing....

- Use RHEL OpenStack Platform on Intel® Architecture to implement on-demand self-service
 - With this combination, you can enjoy the full benefits of resource pooling across compute, network and storage
- Get educated today
 - Take advantage of OnRamp content, best practices and training materials
 - Visit the OnRamp site at www.onramp2thecloud.com and also www.01.org/openstack for more information
- Engage with your Red Hat and Intel account managers or consultants to plan and initiate OpenStack PoCs or pilot deployments

Legal Information

Intel technologies, features and benefits depend on system configuration and may require enabled hardware, software or service activation. Performance varies depending on system configuration. No computer system can be absolutely secure. Check with your system manufacturer or retailer or learn more at intel.com.

Intel's compilers may or may not optimize to the same degree for non-Intel microprocessors for optimizations that are not unique to Intel microprocessors. These optimizations include SSE2, SSE3, and SSSE3 instruction sets and other optimizations. Intel does not guarantee the availability, functionality, or effectiveness of any optimization on microprocessors not manufactured by Intel. Microprocessor-dependent optimizations in this product are intended for use with Intel microprocessors. Certain optimizations not specific to Intel microarchitecture are reserved for Intel microprocessors. Please refer to the applicable product User and Reference Guides for more information regarding the specific instruction sets covered by this notice. Notice Revision #20110804

No license (express or implied, by estoppel or otherwise) to any intellectual property rights is granted by this document.

Intel disclaims all express and implied warranties, including without limitation, the implied warranties of merchantability, fitness for a particular purpose, and non-infringement, as well as any warranty arising from course of performance, course of dealing, or usage in trade.

This document contains information on products, services and/or processes in development. All information provided here is subject to change without notice. Contact your Intel representative to obtain the latest forecast, schedule, specifications and roadmaps.

The products and services described may contain defects or errors known as errata which may cause deviations from published specifications. Current characterized errata are available on request.


Copies of documents which have an order number and are referenced in this document may be obtained by calling 1-800-548-4725 or by visiting www.intel.com/design/literature.htm.

Intel, the Intel logo, Intel vPro, Look Inside., the Look Inside. logo, Intel Xeon Phi, and Xeon are trademarks of Intel Corporation in the U.S. and/or other countries.

*Other names and brands may be claimed as the property of others.

© 2015 Intel Corporation.

THANK YOU
Q&A



experience
what's inside™