




OPENSTACK AND THE SKA



OPENSTACK AND THE SKA

- ▶ The overhead of virtualisation
 - ▶ Strategies for remediation
- ▶ OpenStack's emerging capabilities for HPC
 - ▶ How they might support the SKA use case
- ▶ What gaps remain
 - ▶ What can be done to address them



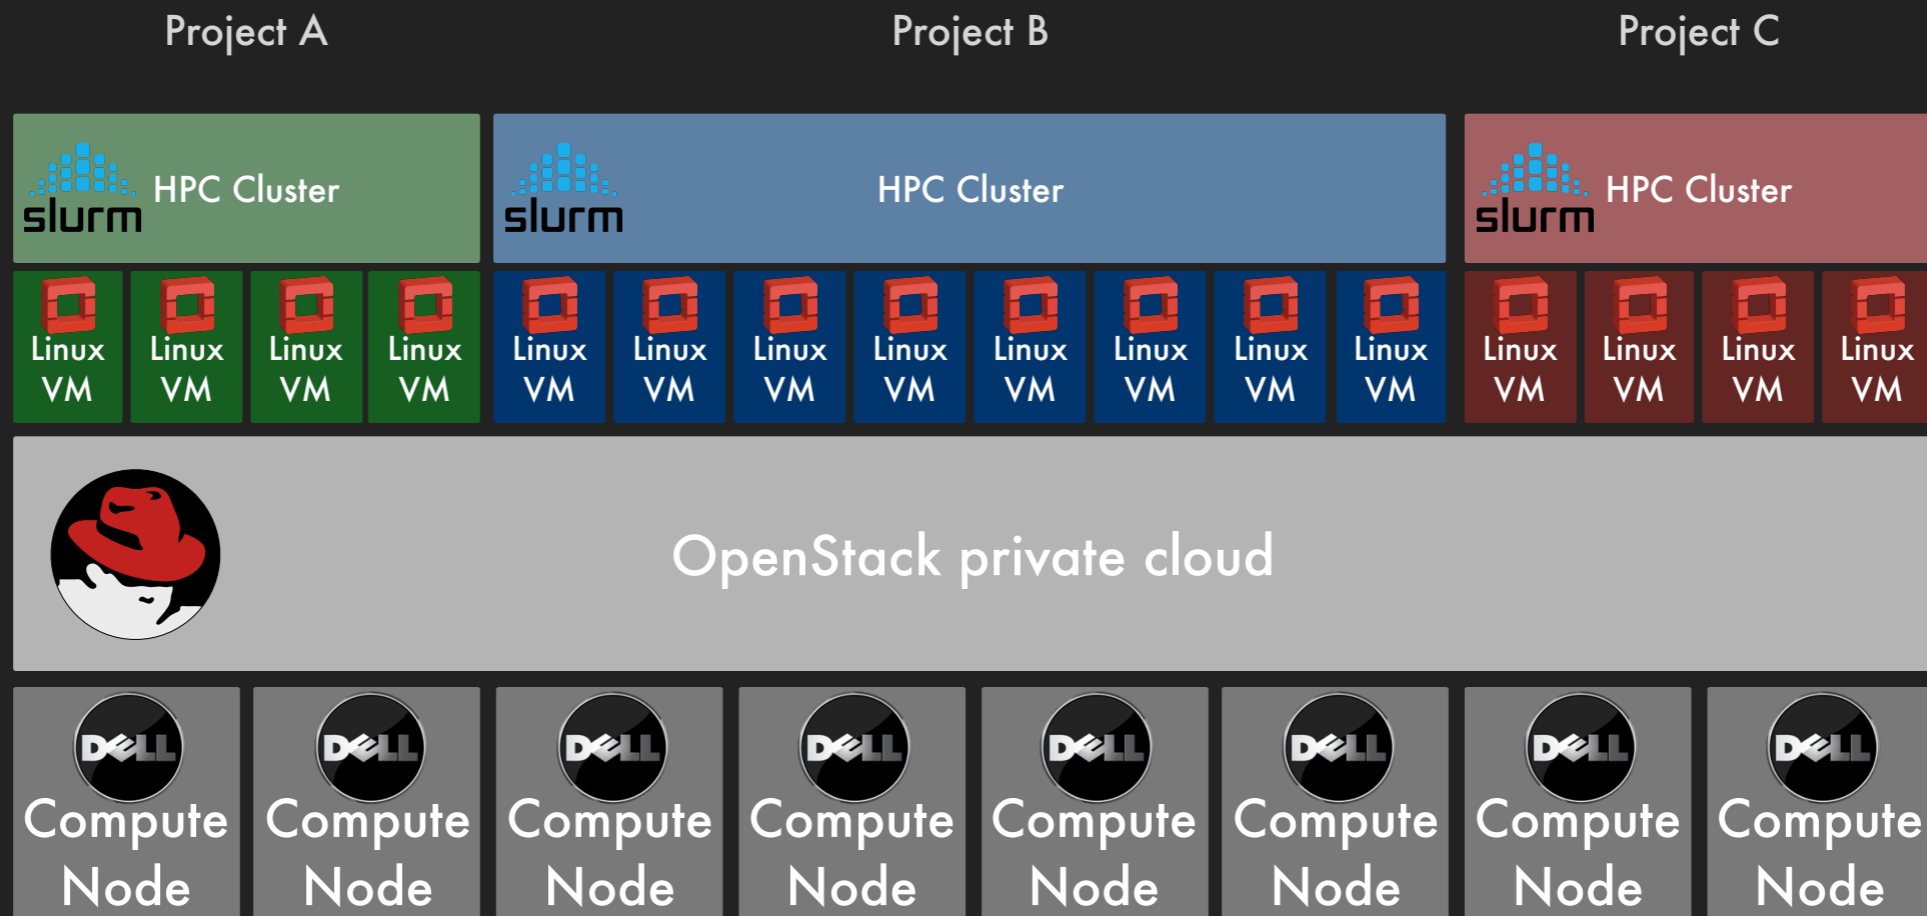


OPENSTACK'S ROLE

- ▶ Software-defined Infrastructure
- ▶ Software-defined Platform
- ▶ Management of users, groups and projects
- ▶ Self-service environment for users

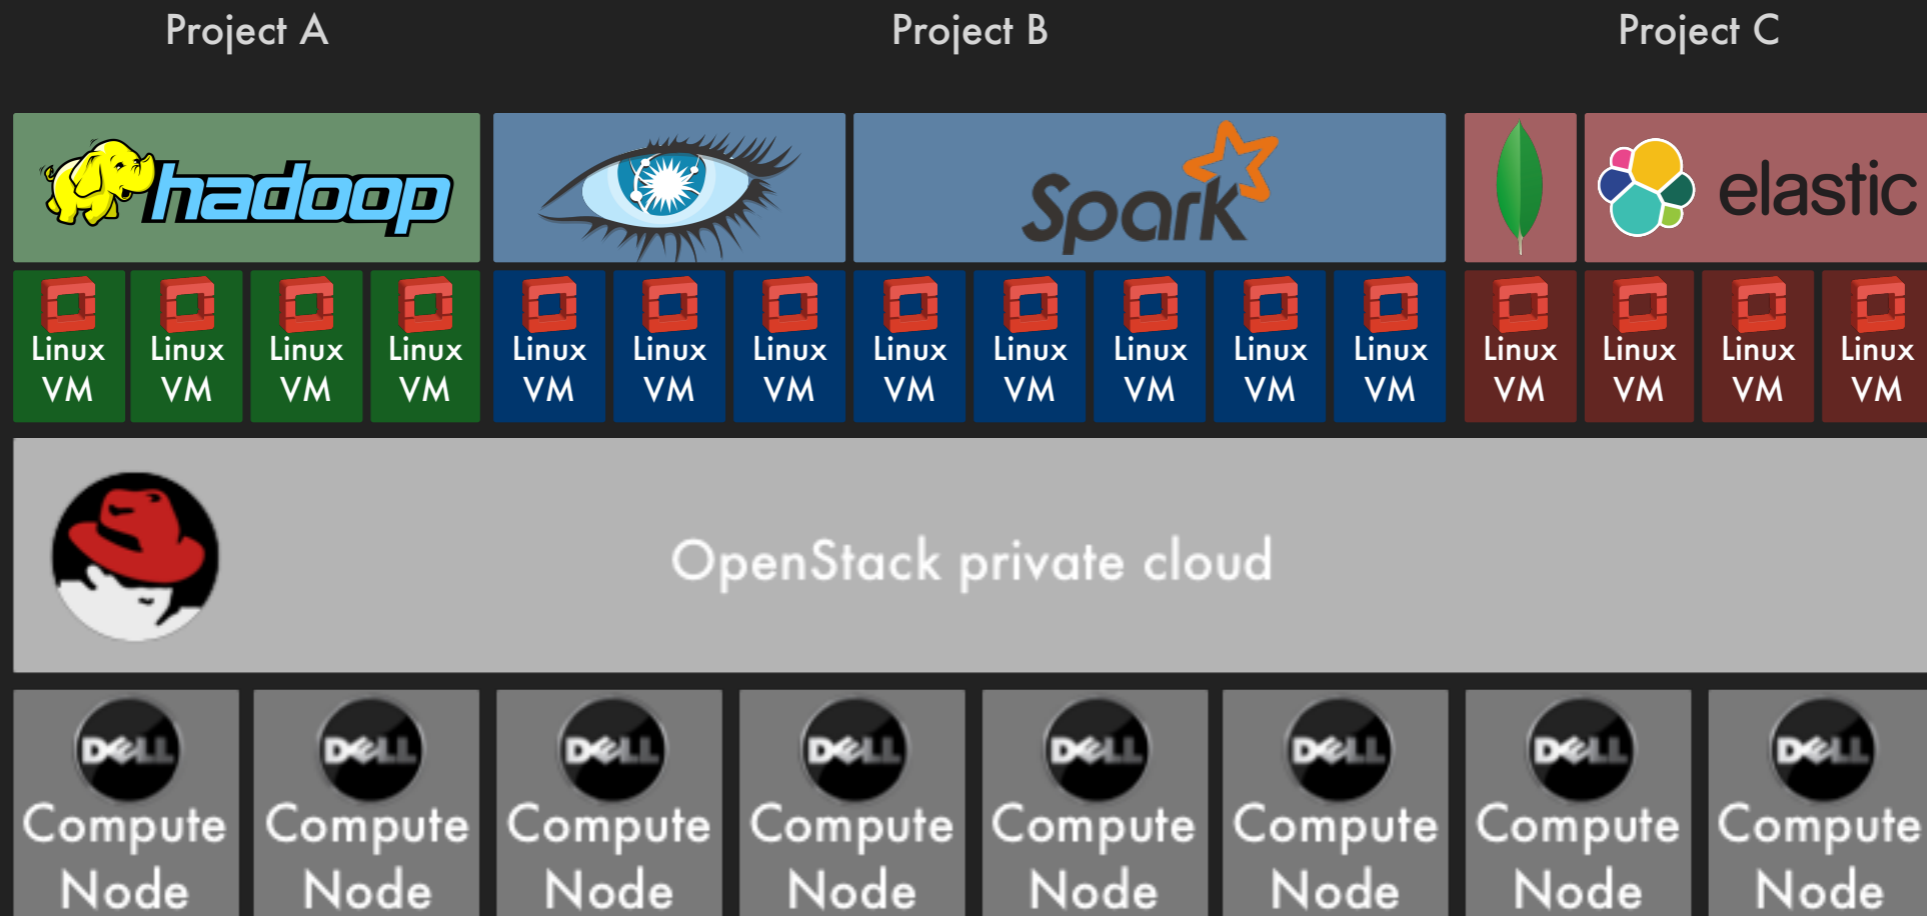


SOFTWARE DEFINED INFRASTRUCTURE





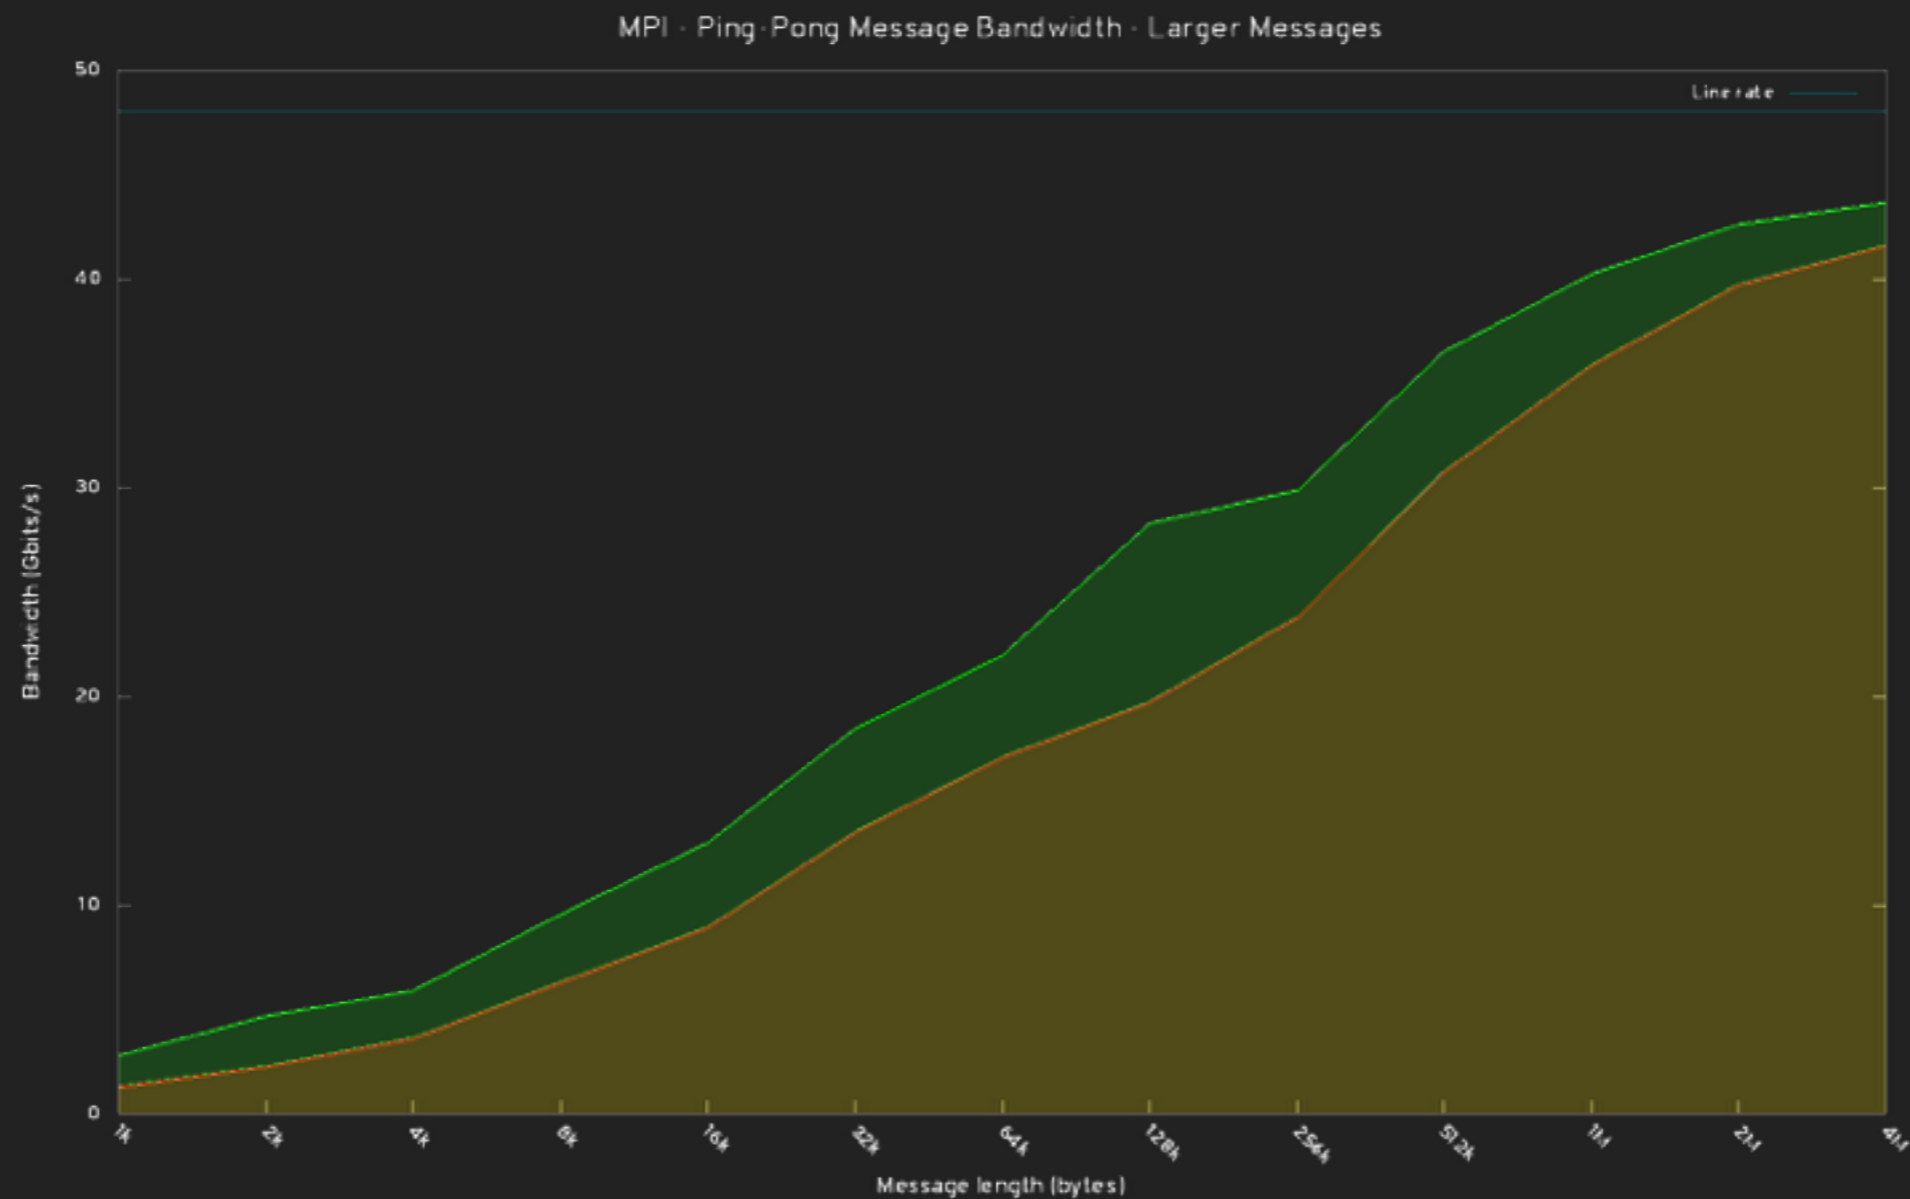
SOFTWARE DEFINED INFRASTRUCTURE





THE OVERHEAD OF VIRTUALISATION

- ▶ Is non-zero - but can be small ... if you work at it!





THE ADVANTAGES OF VIRTUALISATION

- ▶ Migration of long-lived services
- ▶ Sophisticated network configuration
- ▶ Resource over commitment
- ▶ Strong security model
- ▶ Dynamic device management
- ▶ Arbitrary Guest OS



THE RISE OF BARE METAL

- ▶ Capabilities delivered since we last spoke:
 - ▶ Serial consoles
 - ▶ Multi-tenant networking
 - ▶ Improvements in volume attachment



YOU SHOULD USE BARE METAL IF...

- ▶ Utmost performance is needed
- ▶ Flexible on flexibility
- ▶ Running Linux or Linux-like OS
- ▶ Limited use of software defined infrastructure
- ▶ Hardware that doesn't support virtualisation



DON'T USE BARE METAL IF YOU NEED ...

- ▶ High levels of tenancy
- ▶ Strong security guarantees
- ▶ Rapid infrastructure redeployment
- ▶ Rich network configuration
- ▶ Flexible storage device management




BARE METAL CONTAINERISATION STRIKES A BALANCE

- ▶ Easier development model for users
- ▶ Greater utilisation of resource
- ▶ Rapid application startup
- ▶ HPC network technologies can be troublesome




A COMMUNITY TAKES OFF

- ▶ Written with help from the OpenStack Scientific WG
- ▶ Current best practice for OpenStack and HPC
- ▶ Five subject overviews with case studies contributed by WG members





HPC ON OPENSTACK TODAY





HPC ON OPENSTACK IN 2018





OPENSTACK OPTIONS FOR SDP INFRASTRUCTURE

- ▶ Stateless bare metal compute nodes managed using Ironic
 - ▶ Cinder read-only volume multi-attach and ephemeral overlay / Giant initrd
- ▶ SDN network management for multiple RDMA-enabled Ethernet NICs
- ▶ High resolution performance telemetry using Monasca or Prometheus
- ▶ Sysadmins and operators are OpenStack users
- ▶ Application deployment options:
 - ▶ Deployed through selection/configuration of root software image
 - ▶ Containerised application model - Singularity or Kubernetes
 - ▶ Apache software stack (Spark etc.)



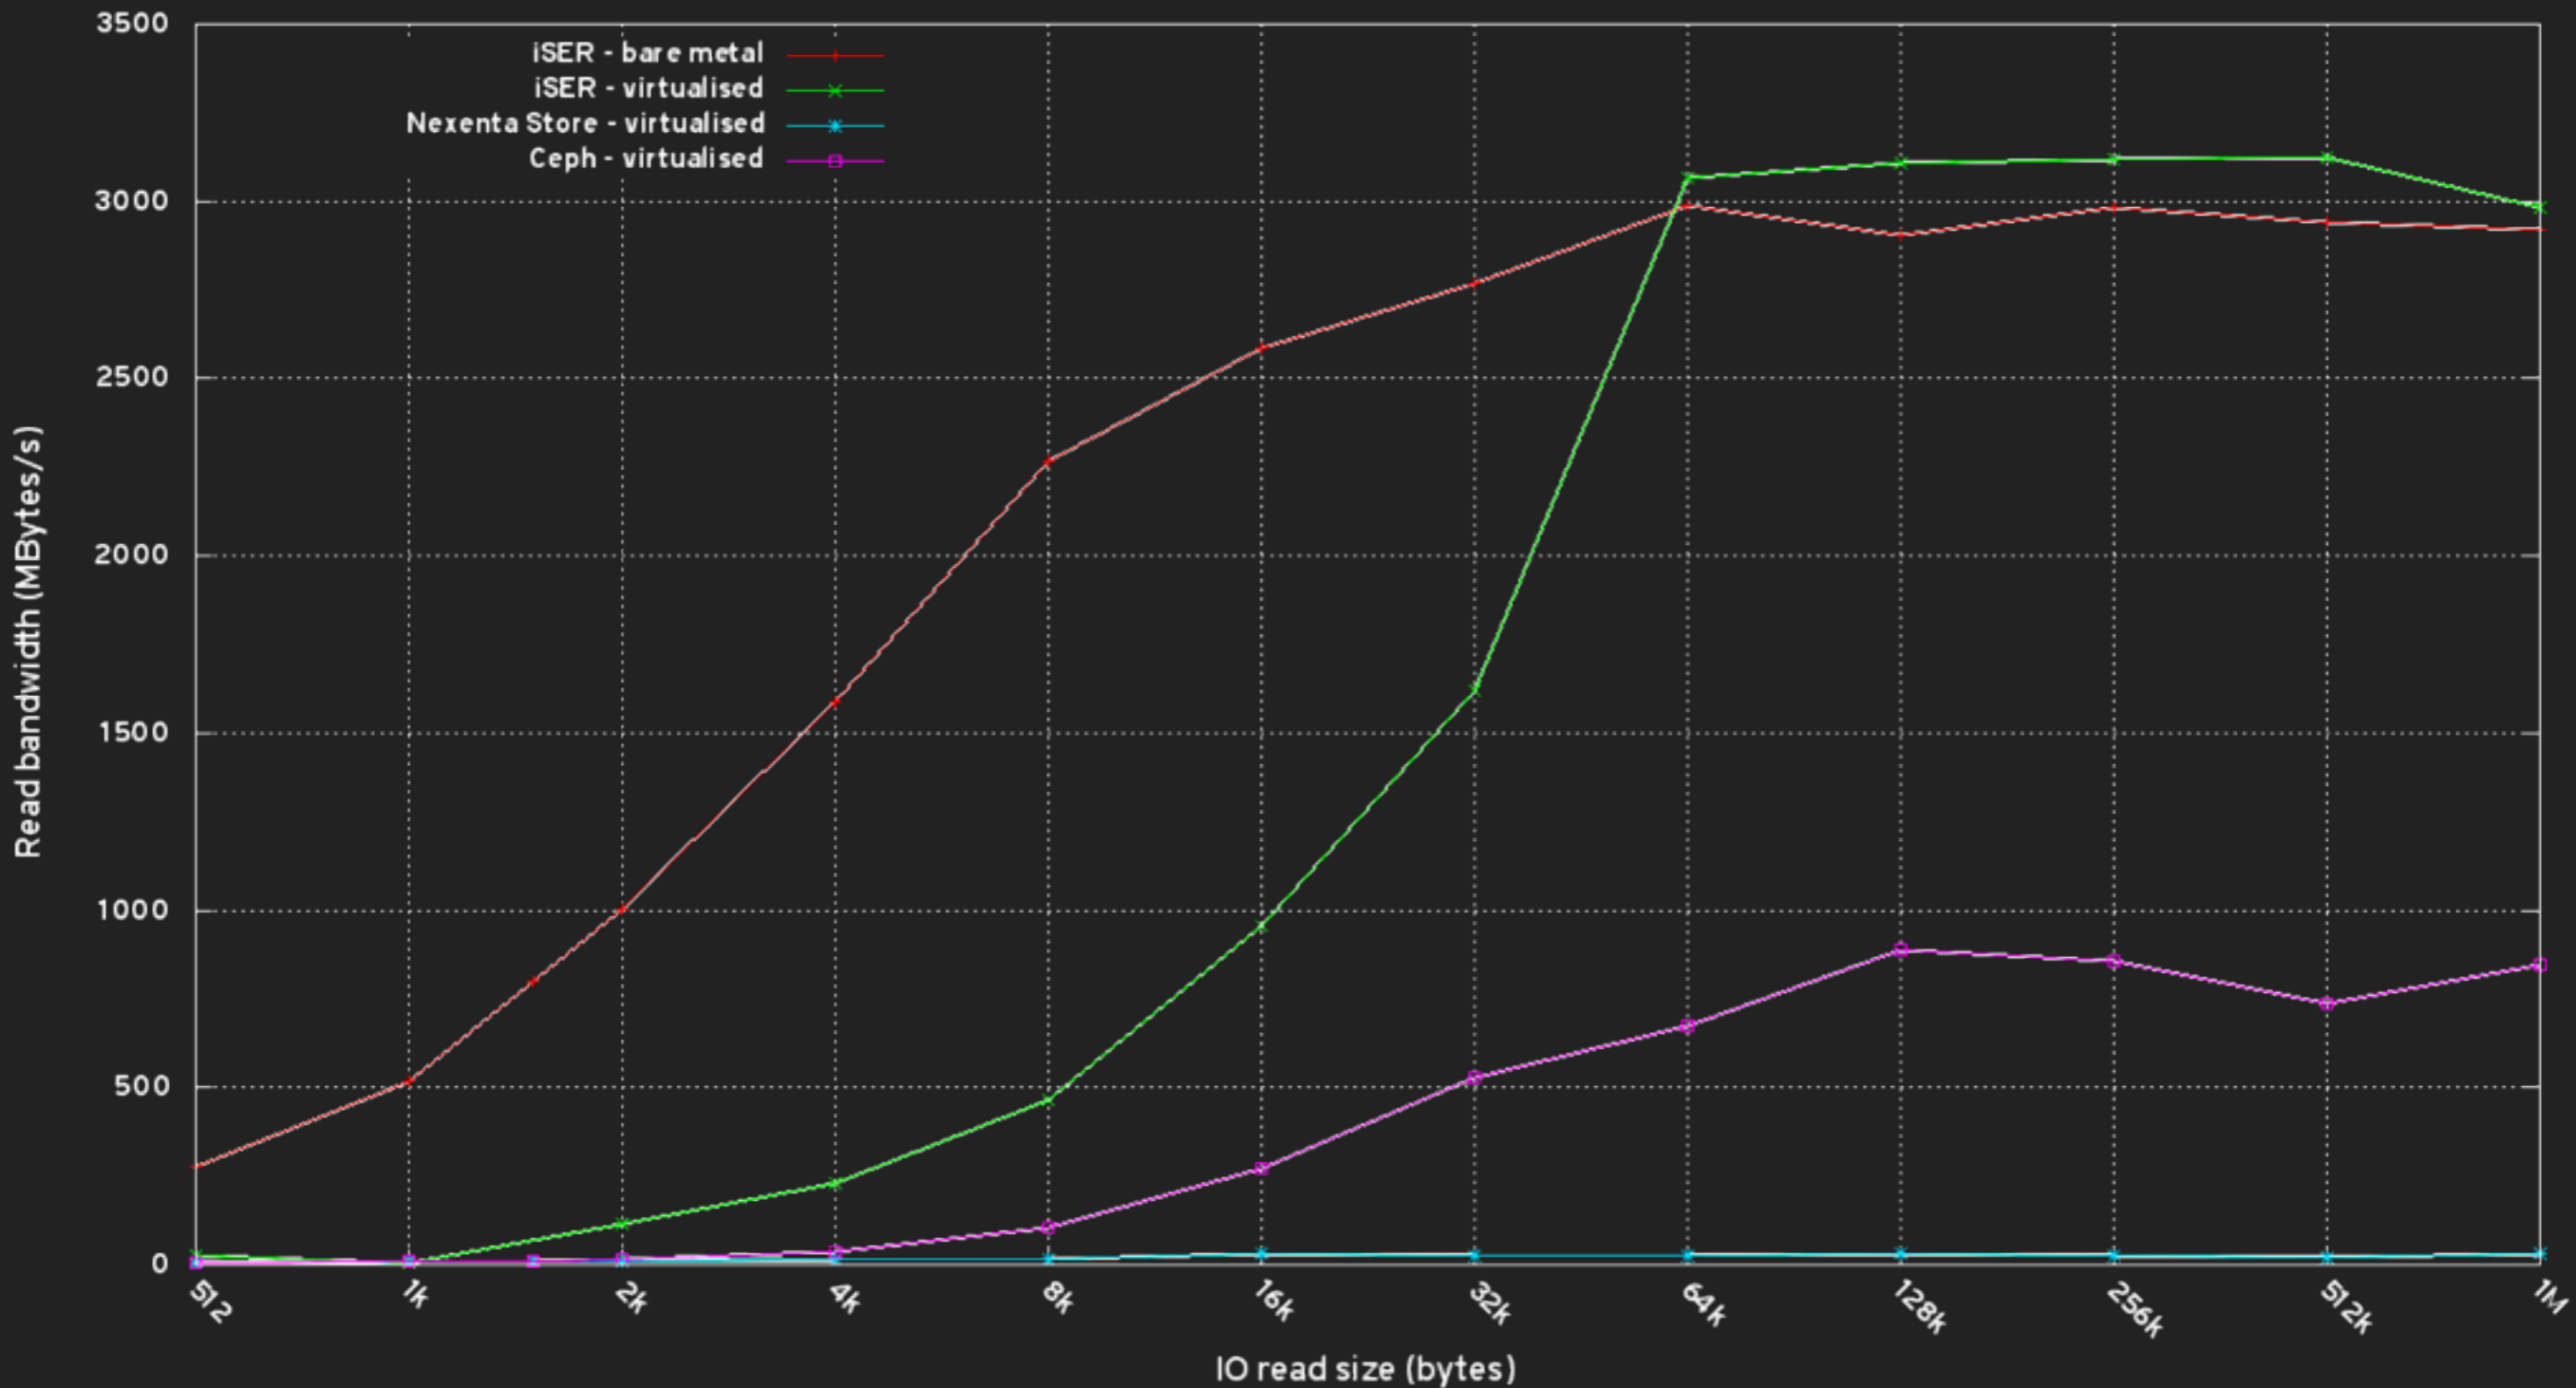
OPENSTACK OPTIONS FOR REGIONAL CENTRES

- ▶ Federated identity management
- ▶ Self-service user interface
- ▶ Virtualised compute
- ▶ Rich environment of application images (containers or VMs)
- ▶ Multi-node application topologies
- ▶ RDMA-centric storage architecture
- ▶ Resource accounting to federated project groups



OPENSTACK - WITH HPC INSIDE

fiio - Block IO Read Bandwidth





WHAT GAPS REMAIN

- ▶ Capabilities at scale
- ▶ Infrastructure complexity
- ▶ Workload management
 - ▶ Future resource reservation
 - ▶ Pre-emptible instances



stig@stackhpc.com
<http://www.stackhpc.com/>
@oneswig

THANK YOU