Oracle Networking for Cloud

Ola Torudbakken
Chief Architect

Databases in the Cloud
February 10-11, 2016
TU Munich, Garning

ORACLE

Safe Harbor Statement

The preceding is intended to outline our general product direction. It is intended for information purposes only, and may not be incorporated into any contract. It is not a commitment to deliver any material, code, or functionality, and should not be relied upon in making purchasing decisions. The development, release, and timing of any features or functionality described for Oracle's products remains at the sole discretion of Oracle.

End of an Era

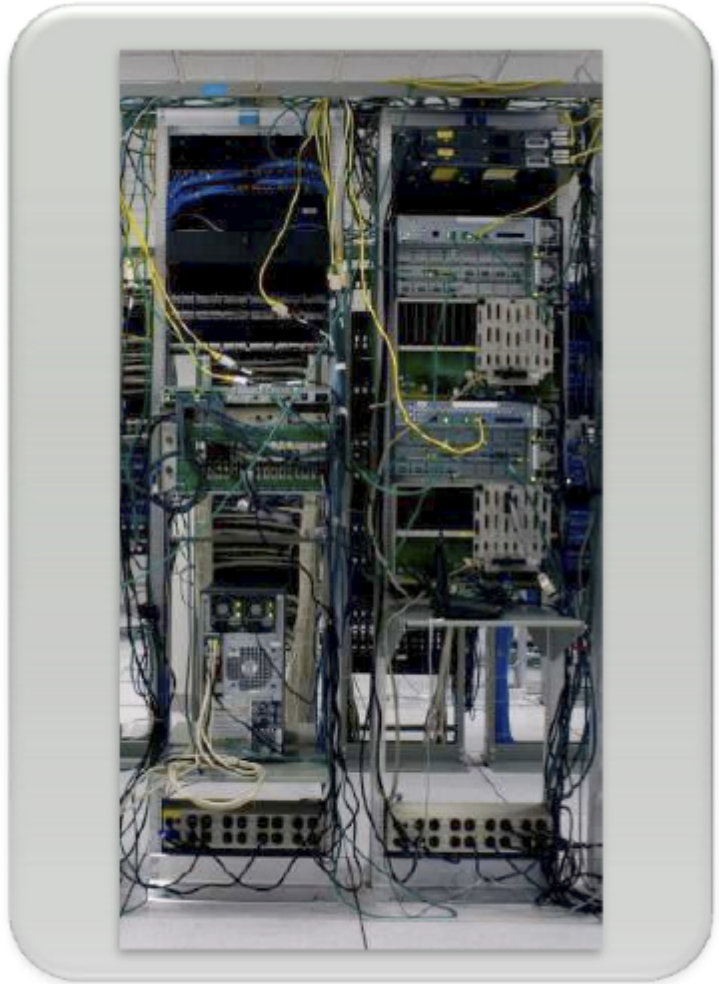

Networking in Cloud Era

"..80%+ of Google´s traffic internal facing.."

(Bikash Koley – Google)

"Every 1kB of external traffic entering the DC generates 93kB of internal traffic"

(Nathan Farrington – Facebook)

- Majority of Communication is Internal
  - Accelerated fast-path for VM-VM and VM-Storage (IP, RDMA and Verbs API)
- Cloud for Enterprise – Secure Multi-Tenant
  - Security Model that covers all bare-metal IaaS to virtualized PaaS/SaaS

Engineered Systems | The Family

Exadata

Exalogic

SuperCluster

Database Appliance

Exalytics

Big Data

PCA

ZFS Storage

Backup

- Expedited time to value
- Easier to manage and upgrade
- Lower cost of ownership
- Reduced change management risk
- Single Vendor Accountability
- Extreme performance

A New Kind of a Computer

Building an elastic cloud with InfiniBand

- Grid Architecture
  - Simple to scale
  - Open, standard
- Lossless Switched Fabric
  - Quality of Service
  - Secure Multi-Tenancy
  - Fault tolerance and Failover
  - Extreme Scalability → 10ks of Nodes

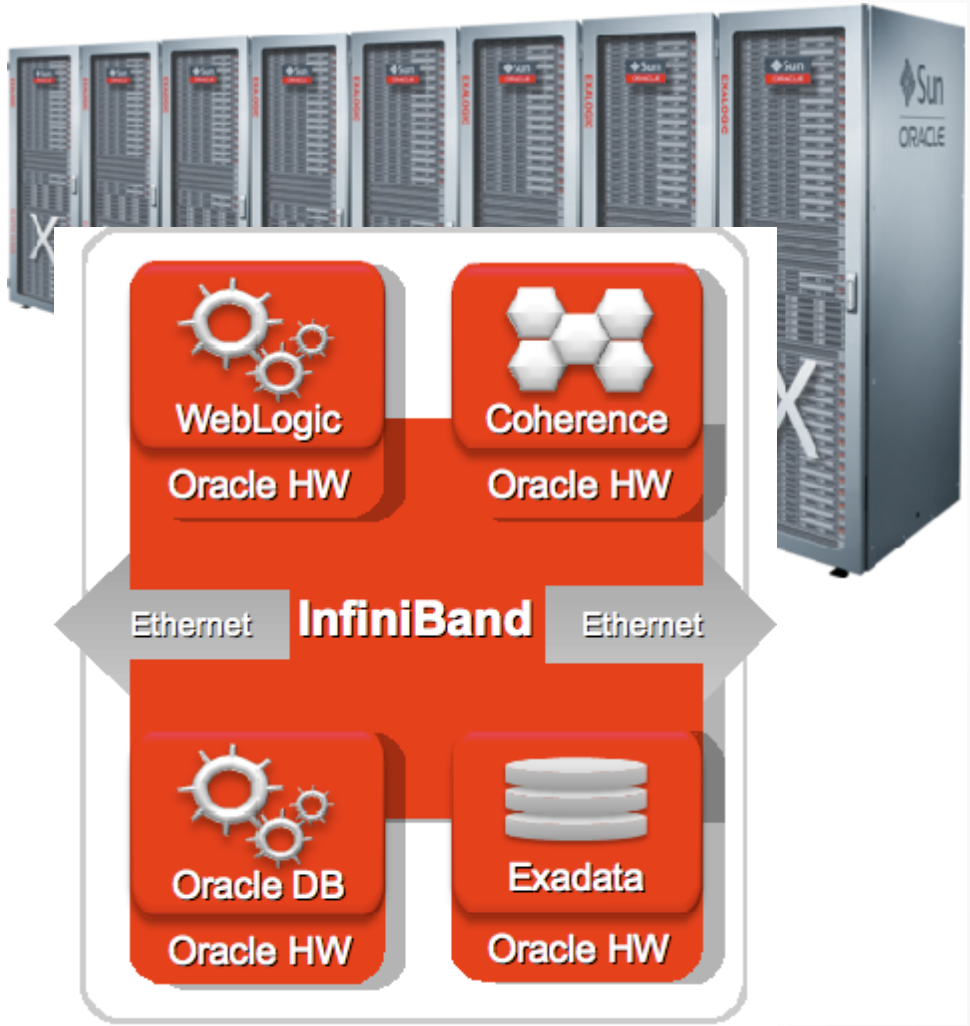

Traditional vs Oracle Exa* Architecture

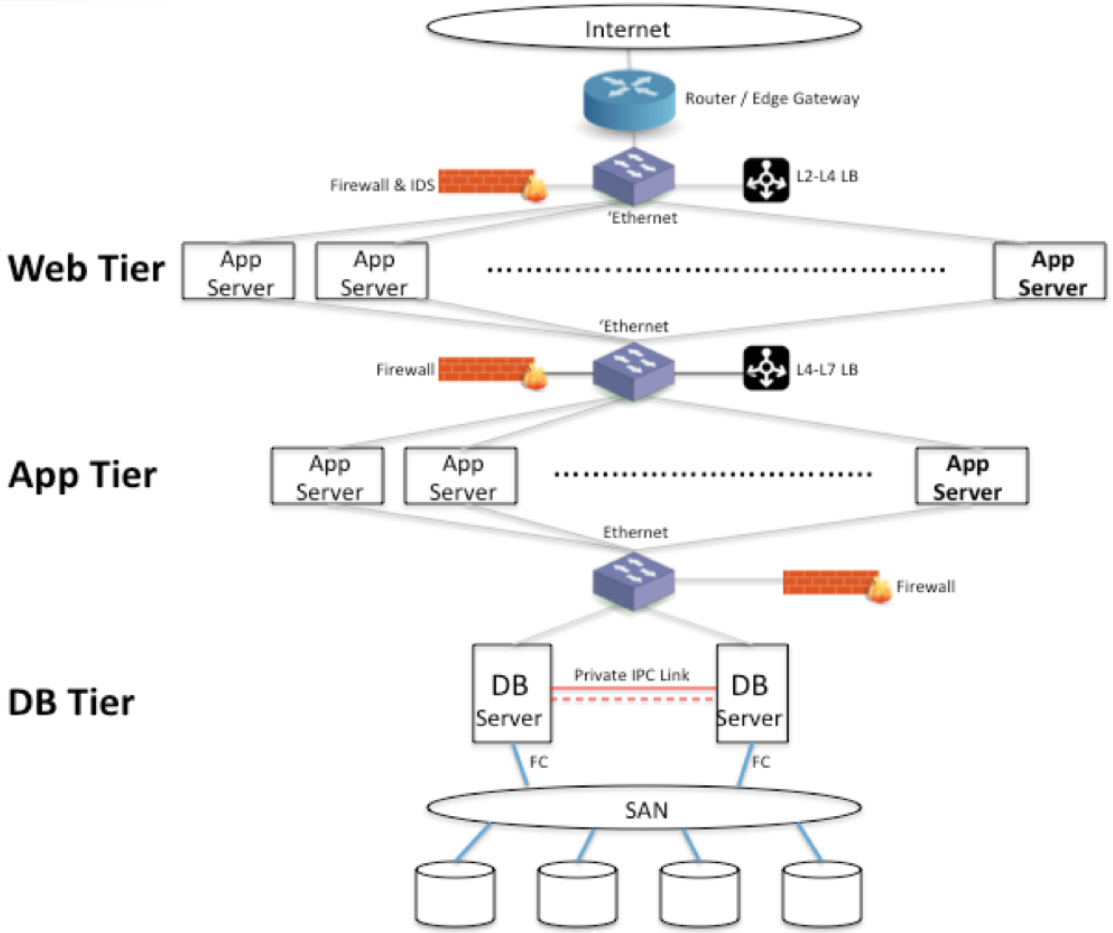

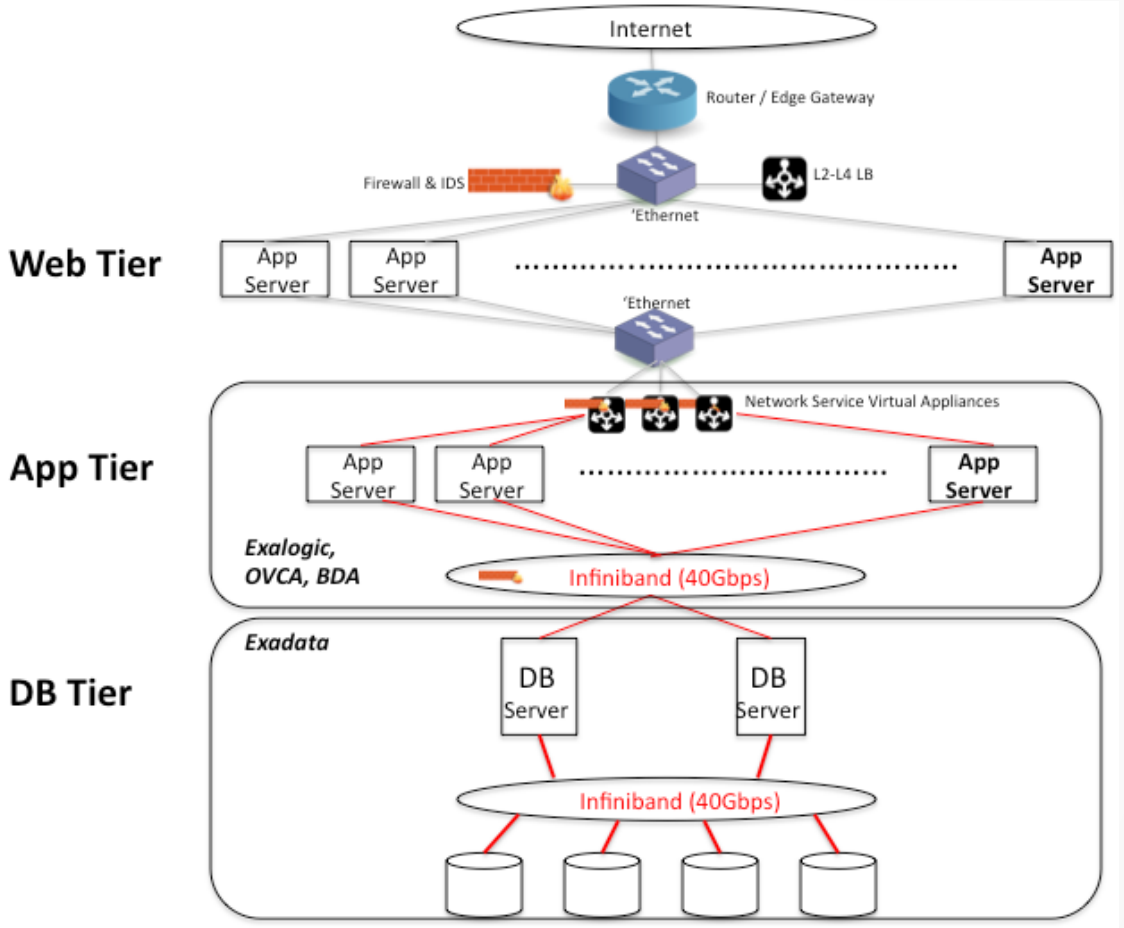

InfiniBand as Oracle´s Private Cloud Fabric

- Highest Bandwidth and Lowest Latency: 100Gbps / 1us
- Constant Bandwidth Network – Predictable Performance
- “Zero” network administration – Pre-configured and automated
- Secure Multi-Tenancy for Bare-Metal and Virtualization
- Verbs API & RDMA
  → Oracle Red Stack Application Acceleration (RDMBS, RAC, FMW, Coh*, OVM)
- Seamless support of legacy Network & Storage apps
- Loss-Less fabric
- OPEN
  - www.infinibandta.org, www.openfabrics.org, kernel.org
  - Rich OS Support (Linux, Solaris, Windows, ESX)

Infiniband Performance Advantage

Up to 12x Performance

- **Network IO is critical**
  - #1 limiting factor for application performance and scale
- **Eliminate Buffer Copies**
  - From four to zero, less CPU
- **Use larger message sizes to reduce network overhead**
  - 64K instead of 4K packets
- **Oracle Red Stack optimized for InfiniBand and Verbs**
  - 3x-10x throughput over 10 GbE & 90% less latency
- **RDMA and Persistent Memory**
- **Full Virtualization Support: SR-IOV and PV**

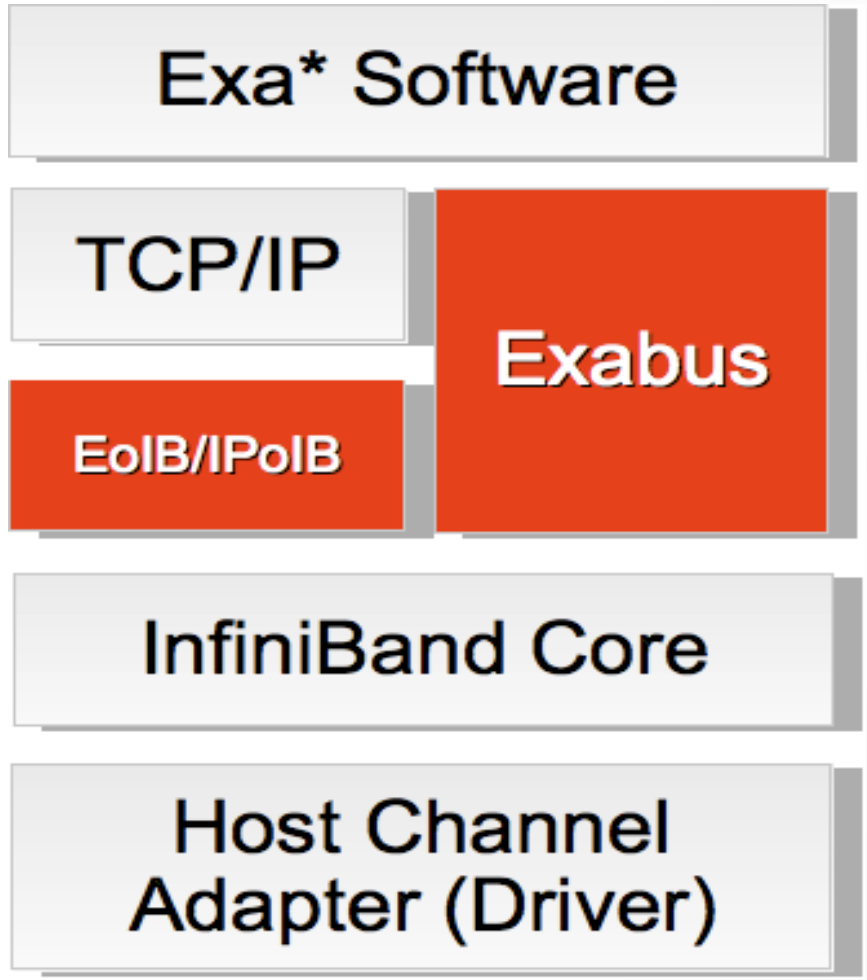

Solid Roadmap

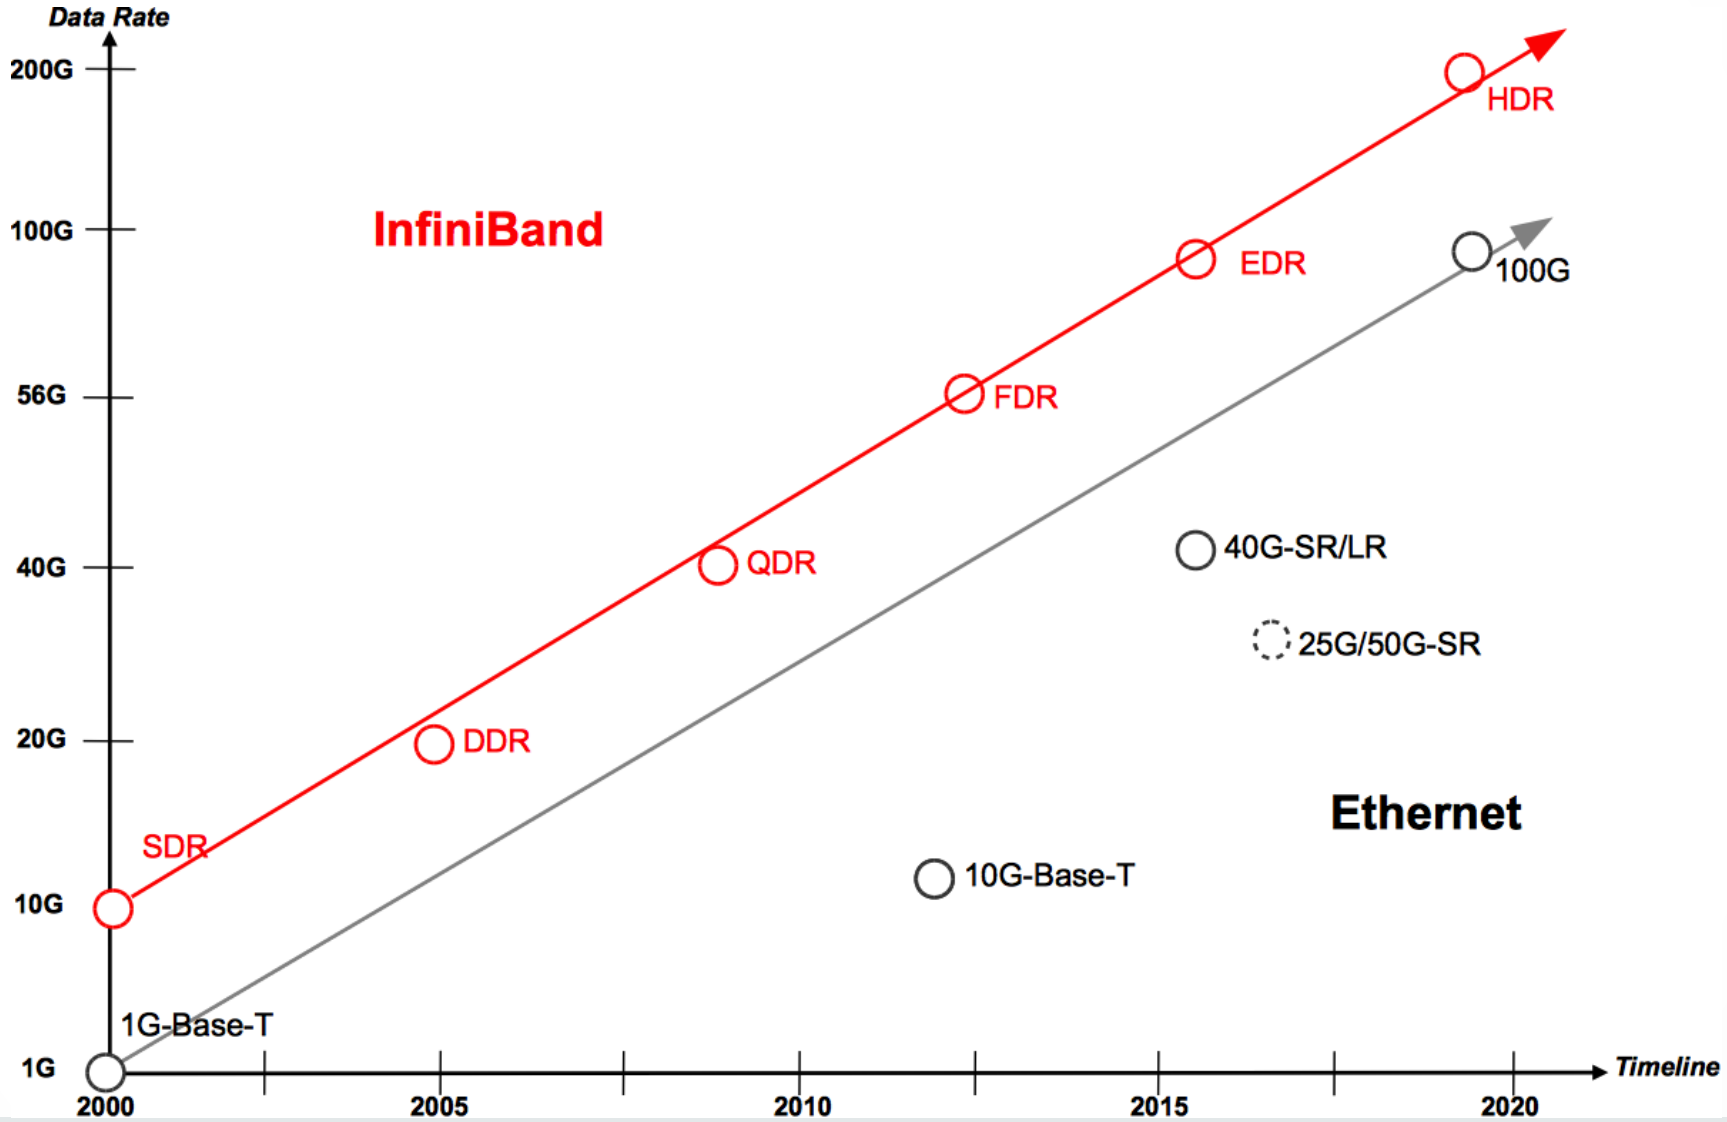

Proven Scalability

ORACLE

Oracle's Networking

Engineered and Integrated to enhance the Oracle Red Stack

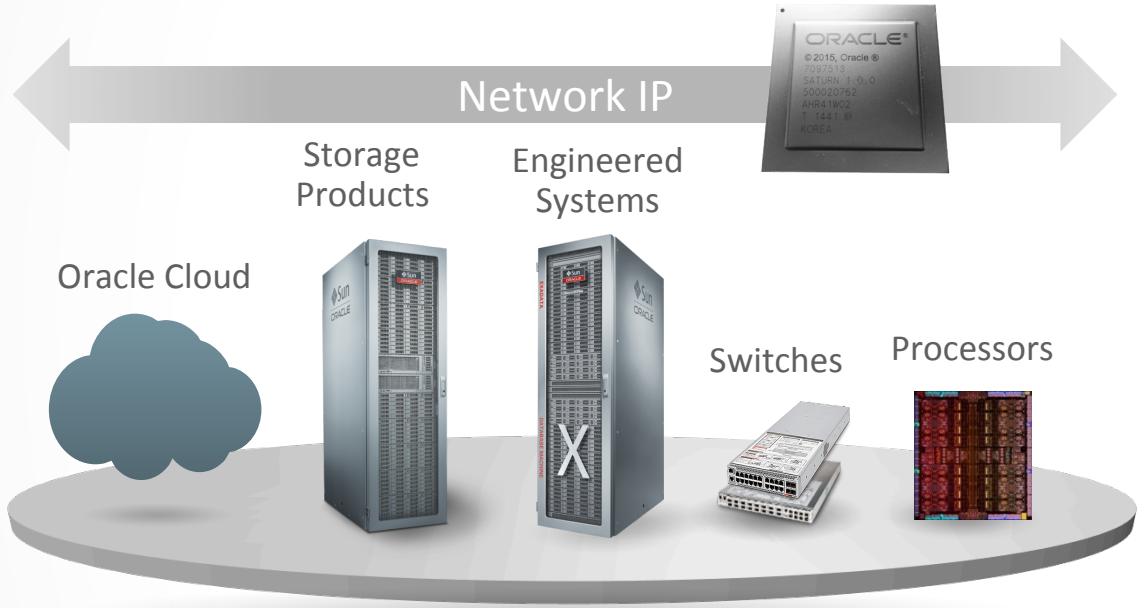

High Performance, Low Latency InfiniBand and Ethernet Network Fabrics

Wire-Once Software Defined Networking with Oracle SDN for your cloud

Unified Management across network fabrics for IaaS/Paas clouds

Deployed in Engineered Systems and in the Oracle Cloud

44K

Shipped over 44K switches

Oracle Networking : Shipping Product Portfolio

Oracle InfiniBand

InfiniBand Switch 36
Versatile Leaf Switch

Infiniband Gateway Switch
32-port Infiniband and 8 10GbE

Oracle Ethernet

Oracle Switch ES2-64 & ES2-72
10/40 GbE switches

Oracle Switch ES1-24
24-port 10GbE switch

10GbE Switch 72p
72 Port 10GbE switch

Oracle Virtual Networking

Fabric Interconnect F1-15
Converged Fabric

Oracle Fabric Interconnect F1-4

Oracle EM

Oracle Fabric Manager

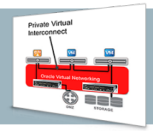

Oracle SDN

OpenStack Plug-in

OVN – Converged Fabric

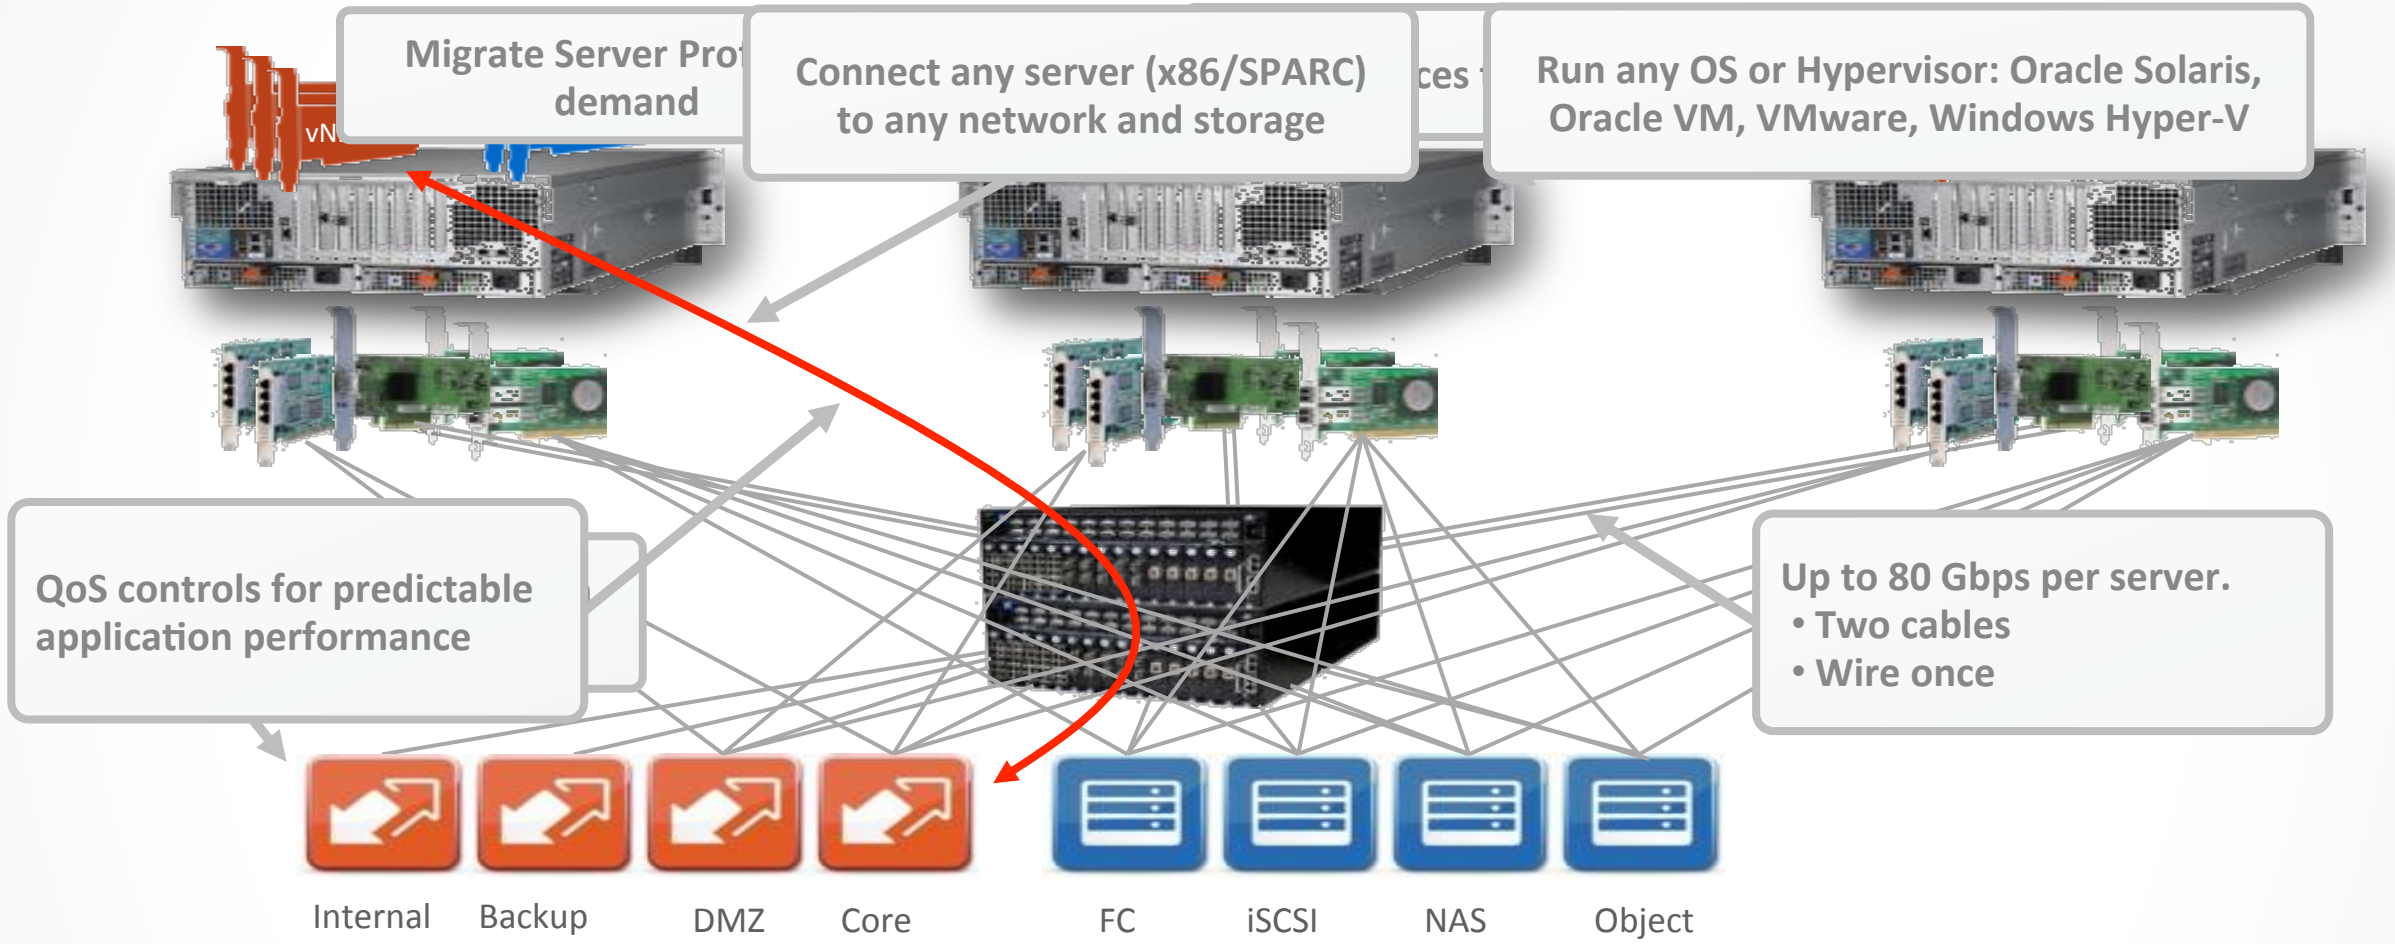

Inflexible system configuration
Low resource utilization

Oracle SDN

- Ethernet Overlay across IB
  - Enable Ethernet L2 with IB and Ethernet hosts
- Server-to-server connectivity
  - 64M isolated virtual Ethernet networks between servers
- Accelerates throughput
  - Up to 100Gb/s server-to-server
  - Traffic remains entirely on fabric
- Flexible connectivity management
  - Fully software defined
  - Isolation without reliance on VLANs

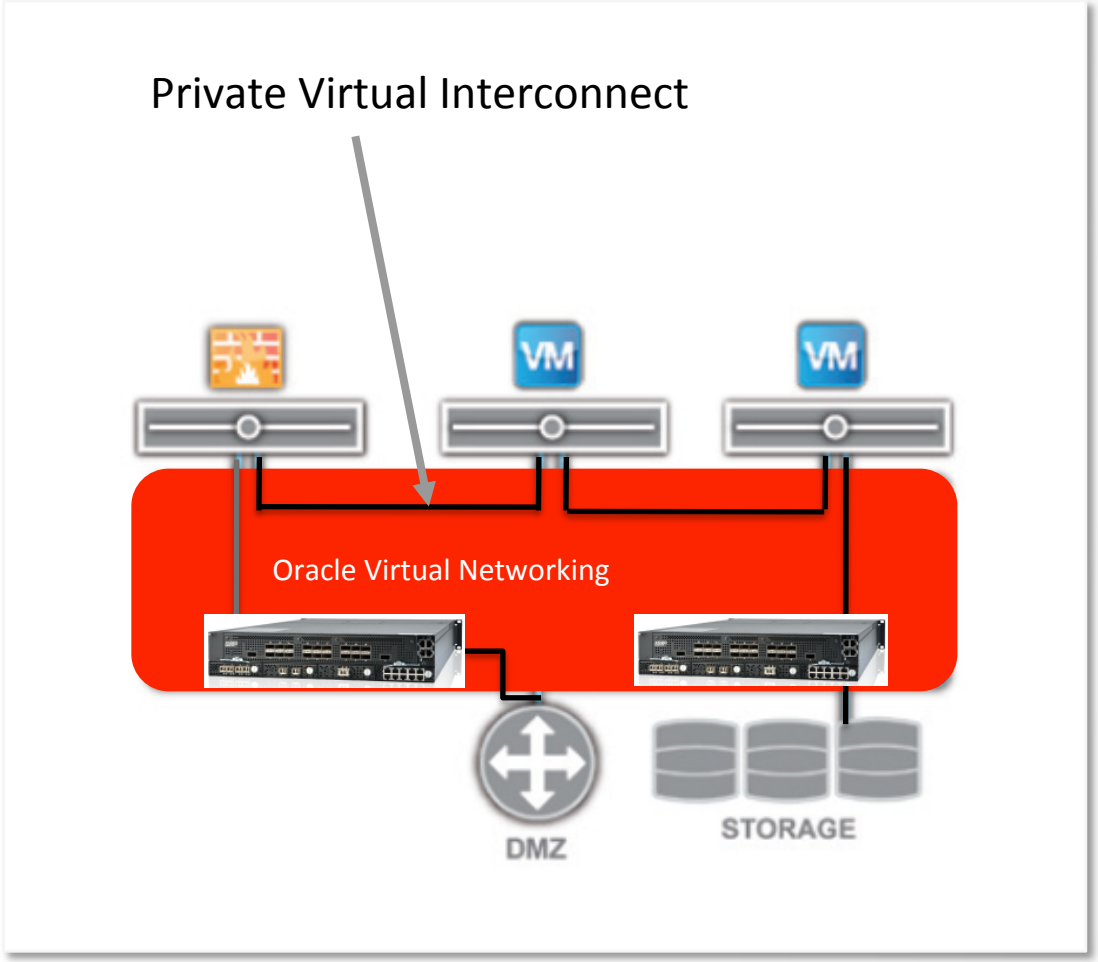

Oracle SDN + Virtual Network Services

Extend Oracle SDN with Virtual Network Services to achieve a fully Software Defined Cloud Enabled Datacenter

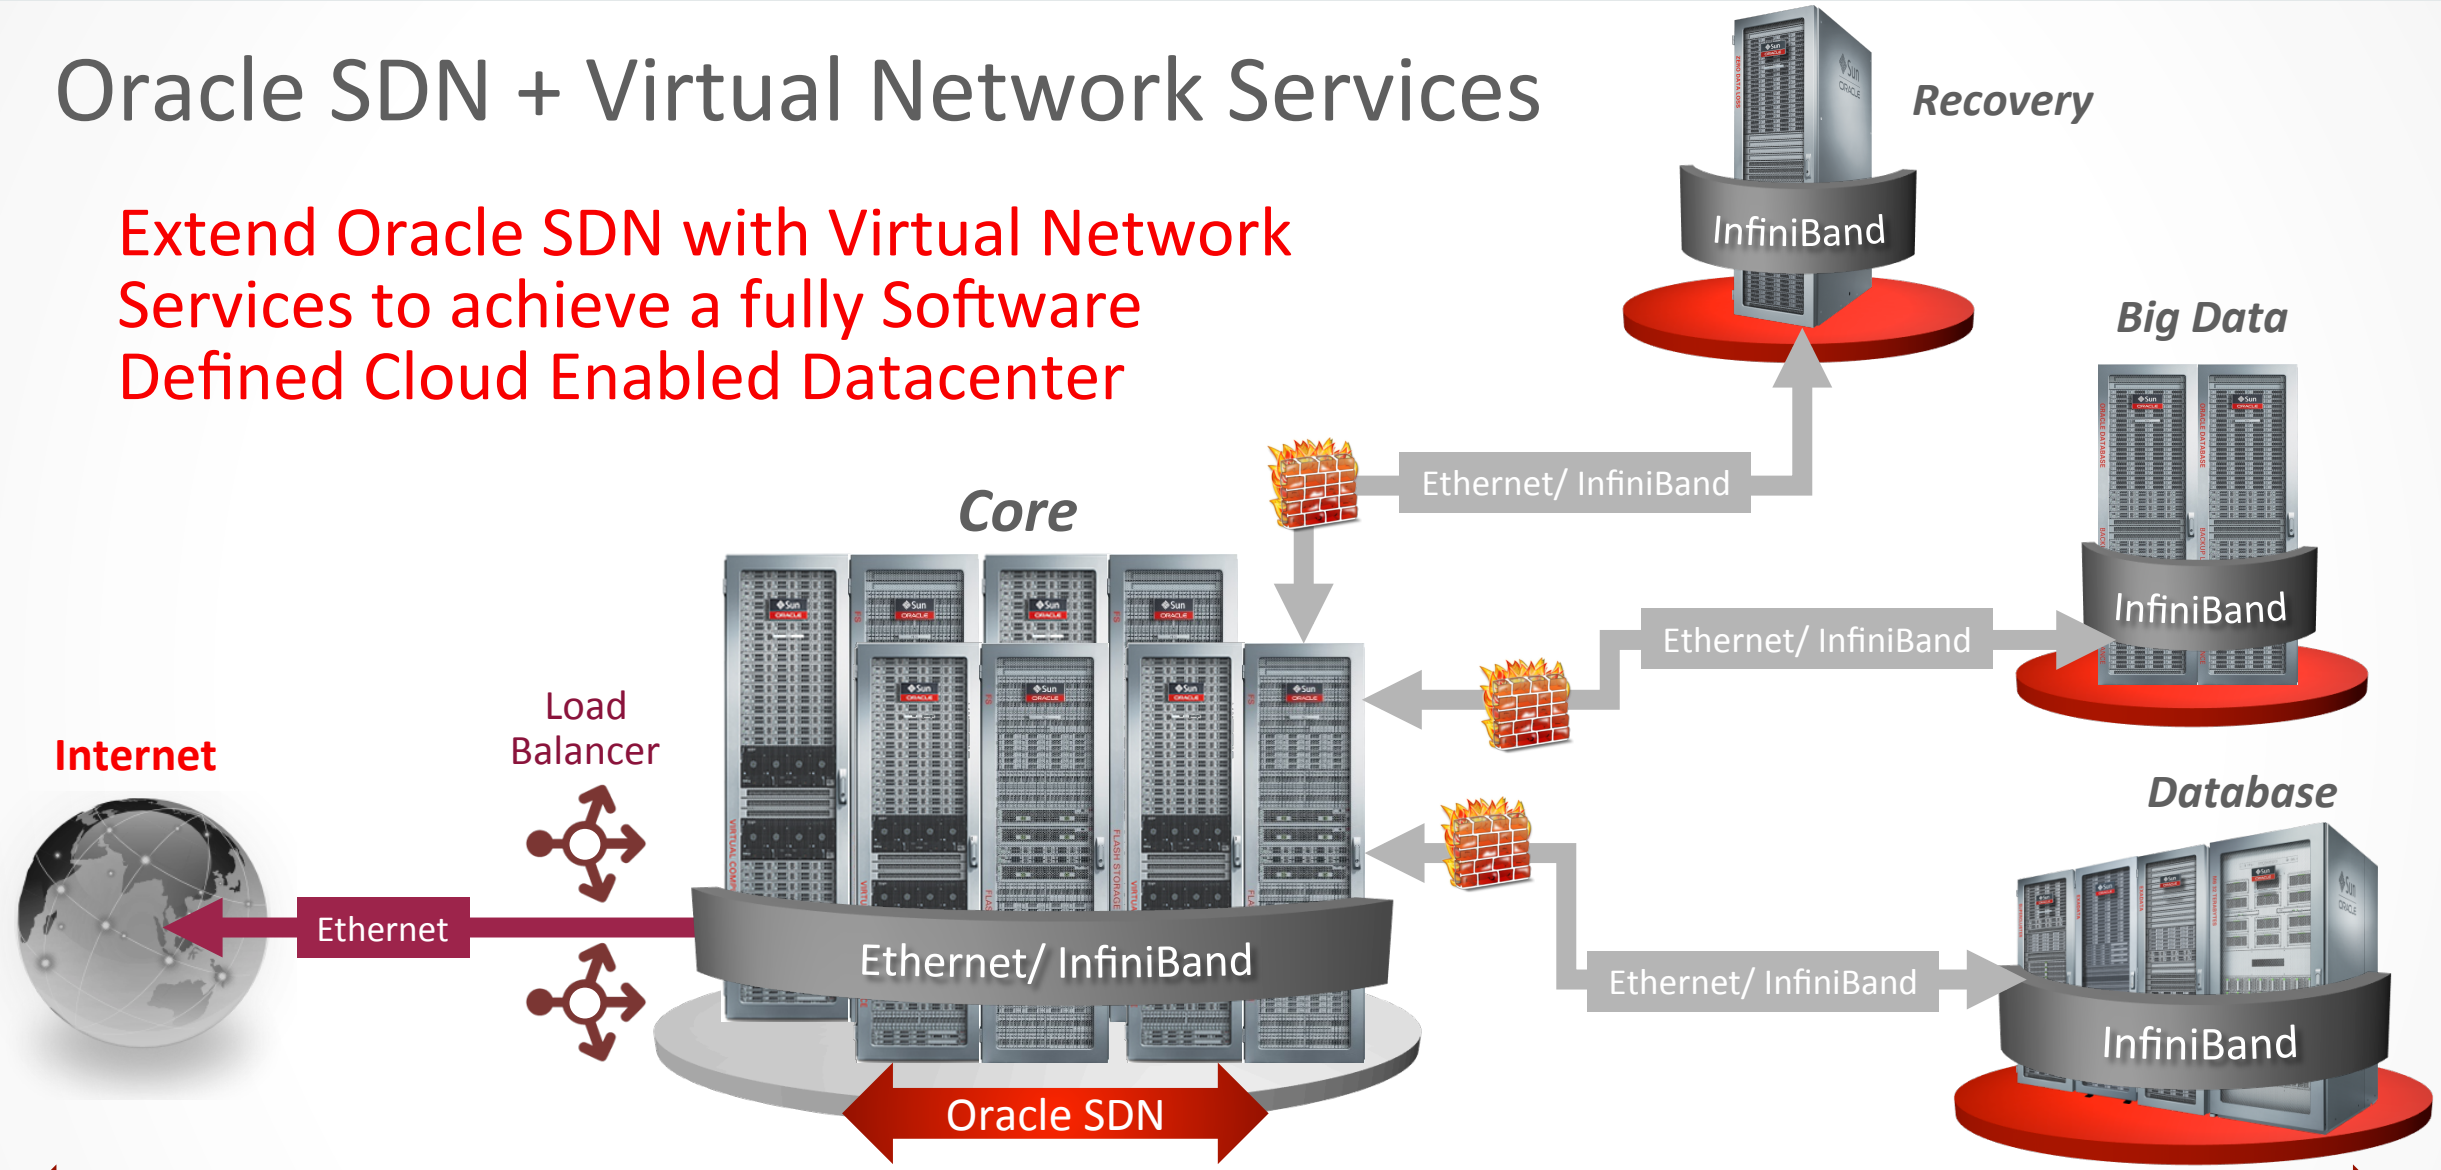

Oracle VNS (Virtual Network Services)

Define and Deploy Network Services Per Tenant

- Feature highlights
  - Stateful firewall with reflexive ACLs
  - Static and dynamic NAT mapping
  - Layer 4 server Load balancing
  - IPSEC VPN
  - Static routing & gateway services across subnets
- Single Virtual Instance per Tenant
- Redundancy built in with VRRP
- Managed from Oracle Fabric Manager

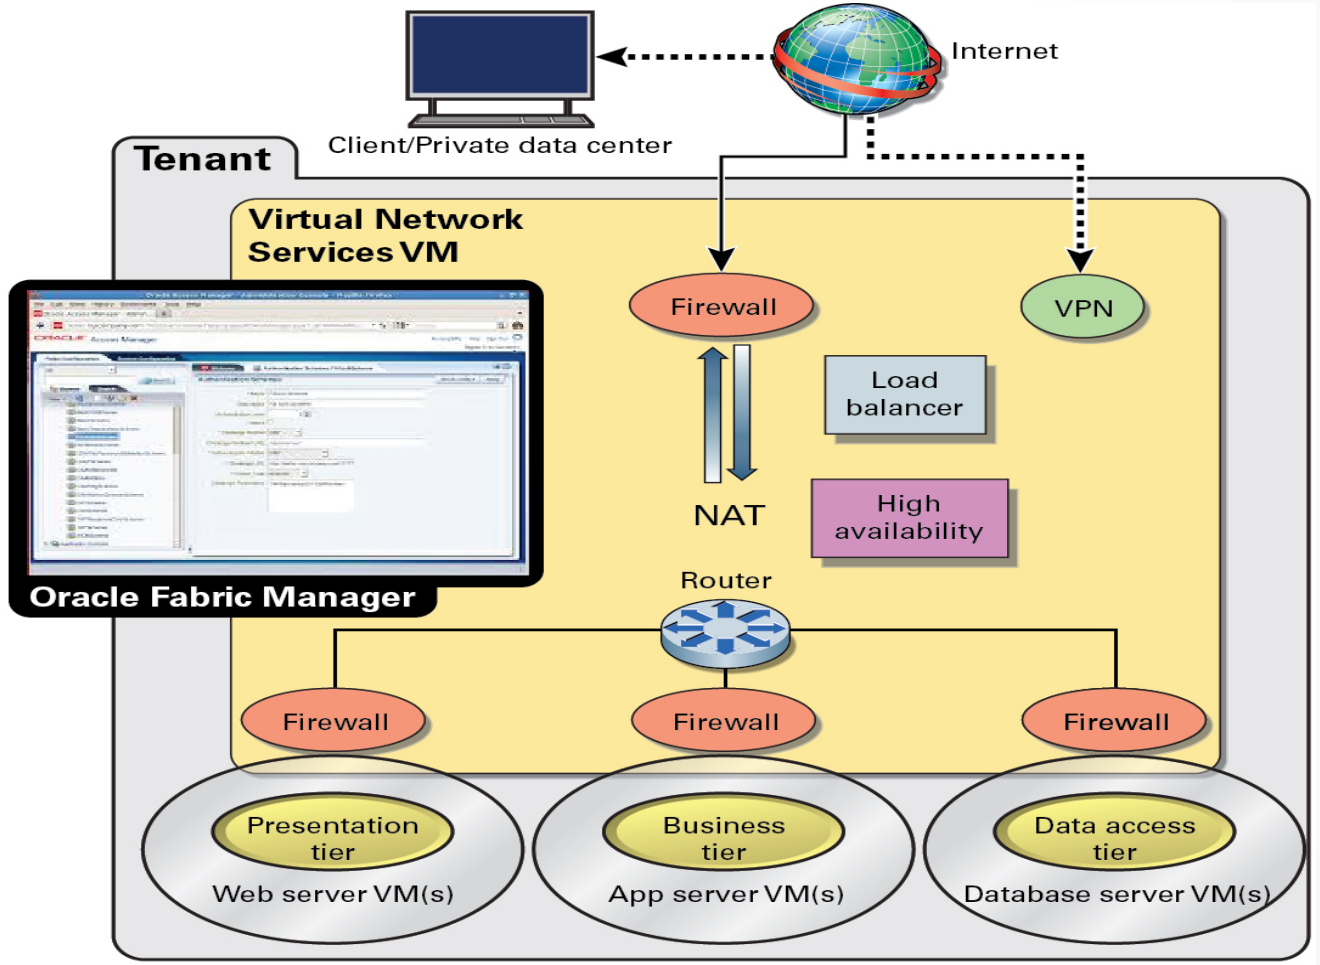

Breakthrough Silicon and Network Systems Design

Speed. Security. Efficiency. Open.

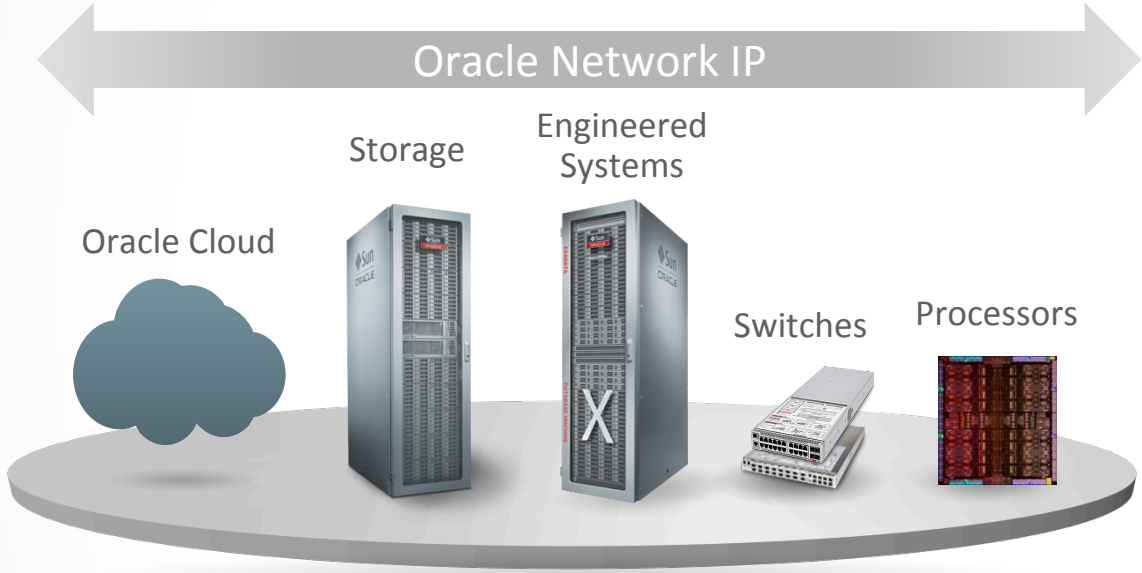

- 100G Converged Fabric
  Oracle Application Acceleration in Fabric
- Engineered for deploying secure multi-tenant enterprise clouds
- Wire-once Software Defined Networking
  Resilient and Self-Healing
- Unified management across network fabrics for public and private clouds

Oracle Cloud, Your Data Center, Engineered Systems: Same Technology and Operating Model

ORACLE

Sillicon Innovation - SPARC SoC with InfiniBand

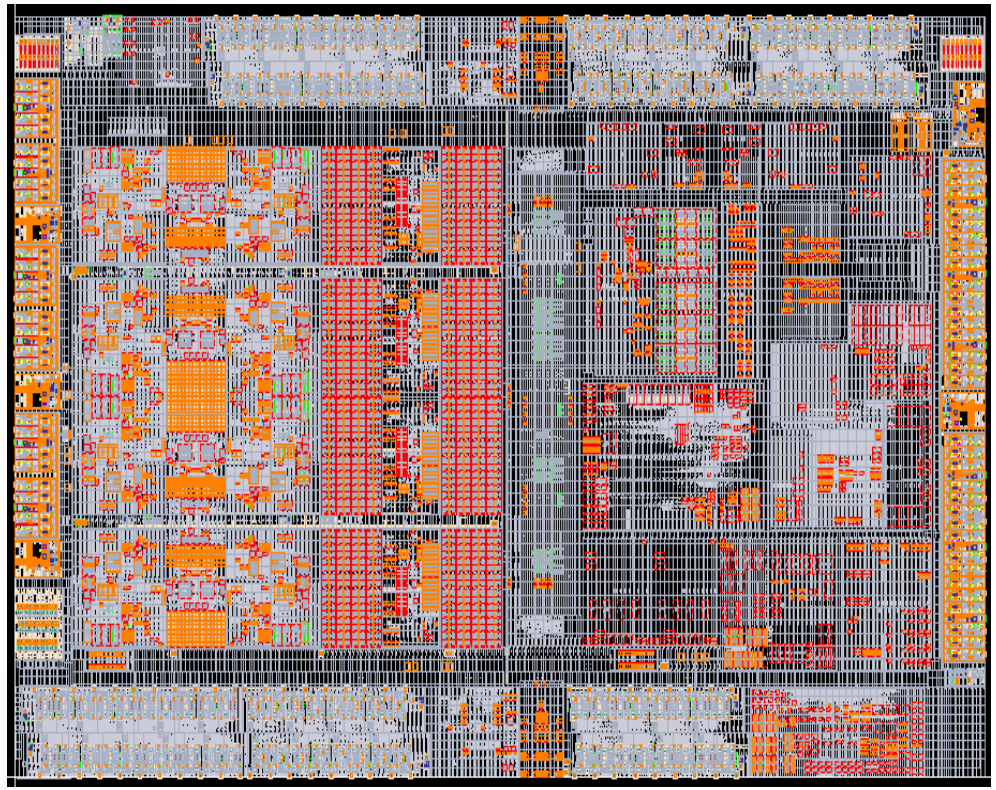

- 8 SPARC 4th generation cores
- Optimized Cache Organization
- Advanced Software in Silicon features
  - Real-time Application Data Integrity (ADI)
  - Concurrent Memory Migration and VA Masking
  - DB query offload engines
- Integrated DDR4 memory
- Integrated PCIe Gen3
- Integrated InfiniBand HCA
- Scale-out IB interconnect

Oracle 100G SDN Fabric for Enterprise Clouds

Speed

100G Converged Fabric - Extreme Performance for Apps and Cloud

Security

Ground-Up Engineered for Secure Multi-Tenant Cloud:
64M Virtual Networks
Secure End-Point Authentication
Isolated Administrative Domains

Efficiency

SDN and virtualized IO services with a single network interface – Extreme Scale (500k Servers) – Breakthrough Oracle Integration & Efficiency

Open

OpenFabrics, OpenStack, InfiniBand, RDMA, LAN & SAN Interoperability

ORACLE

Oracle 100G Fabric - Fastest Converged Fabric

One Fabric for Network, Storage, and Inter-process Communication

- 100 Gb/s Server and Storage connections
  - InfiniBand and Ethernet & Fibre Channel gateways
- Faster Fabric
  - 100G RDMA and Oracle SDN virtual Ethernet networks
  - 89% utilization (routing efficiency) at max load
  - Deterministic QoS and SLA independent of Scale
- Faster Provisioning of Network Infrastructure
  - Networks, Network Services, and Server I/O on-demand
- Faster Server, Storage and Applications
  - Application acceleration built-in to Network

ORACLE

Oracle 100G Fabric

Product Portfolio

- Switch and Virtual I/O Systems
  - Leaf Switch: Oracle InfiniBand Switch IS2-46
  - Spine Switch: Oracle InfiniBand Switch IS2-254
  - Fabric Virtualized I/O: Oracle Fabric Interconnect F2-12
- Server and Storage Interface
  - Oracle Dual Port EDR InfiniBand Adapter
- Fabric and I/O Management
  - Oracle Fabric Manager
- Virtual Networks and Network Services
  - Oracle SDN

Oracle 100G Fabric - Cloud Scale (*500k Servers*)

500k Nodes
3-Tier Clos Fat-Tree

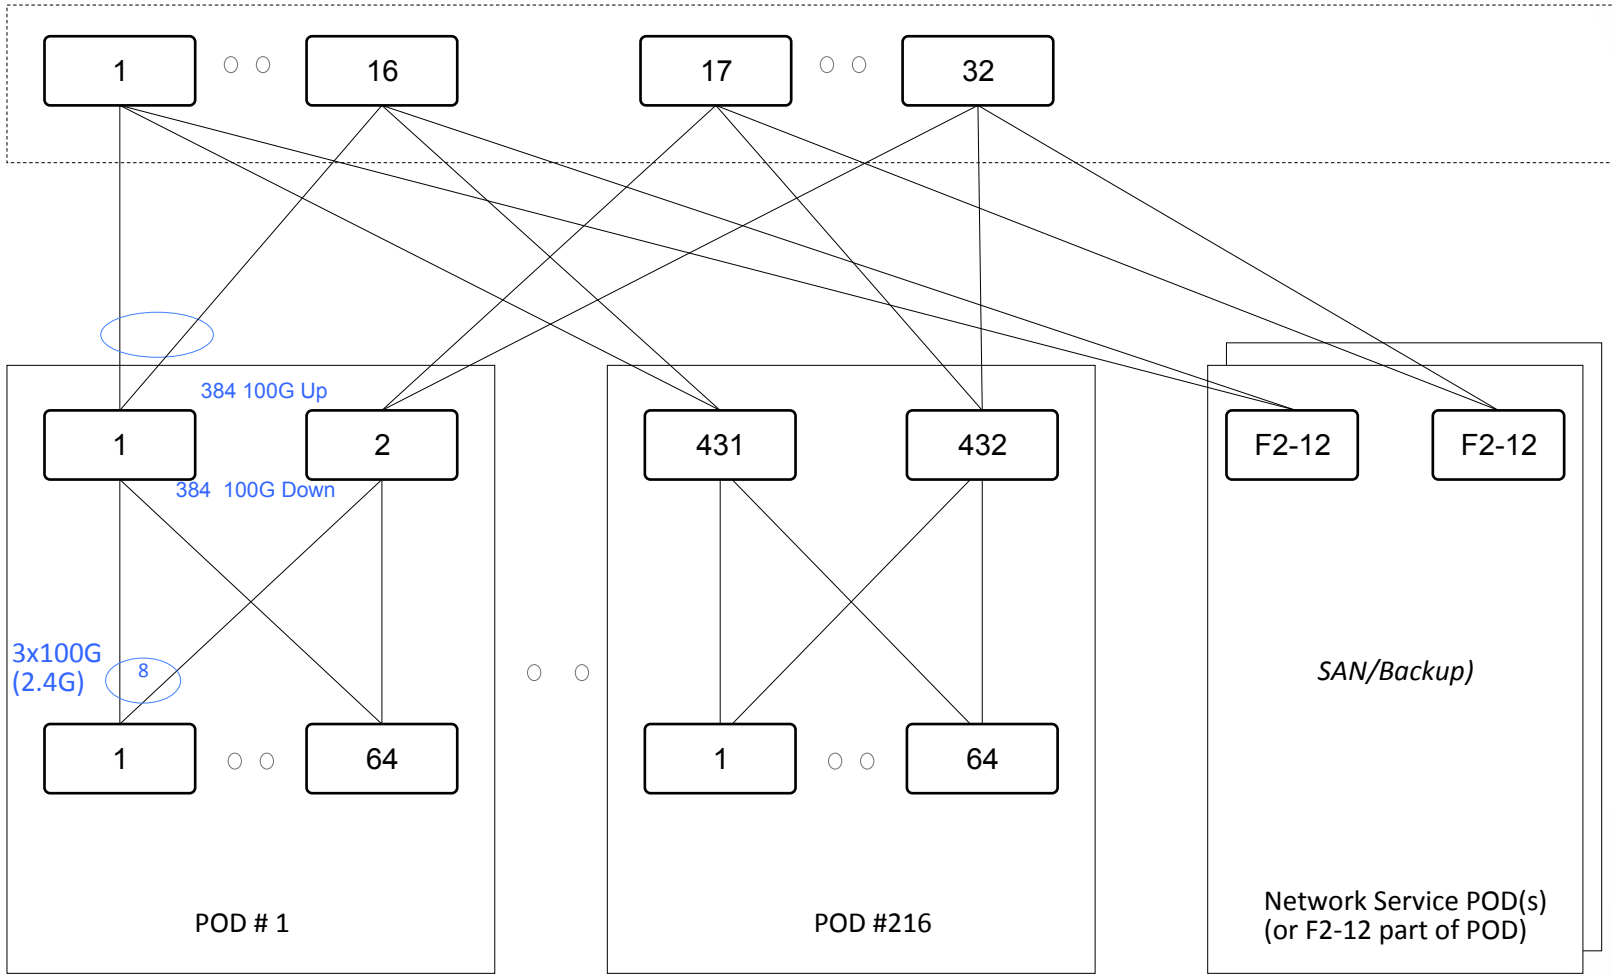

Oracle Networking: Best Network Fabric for your Cloud!

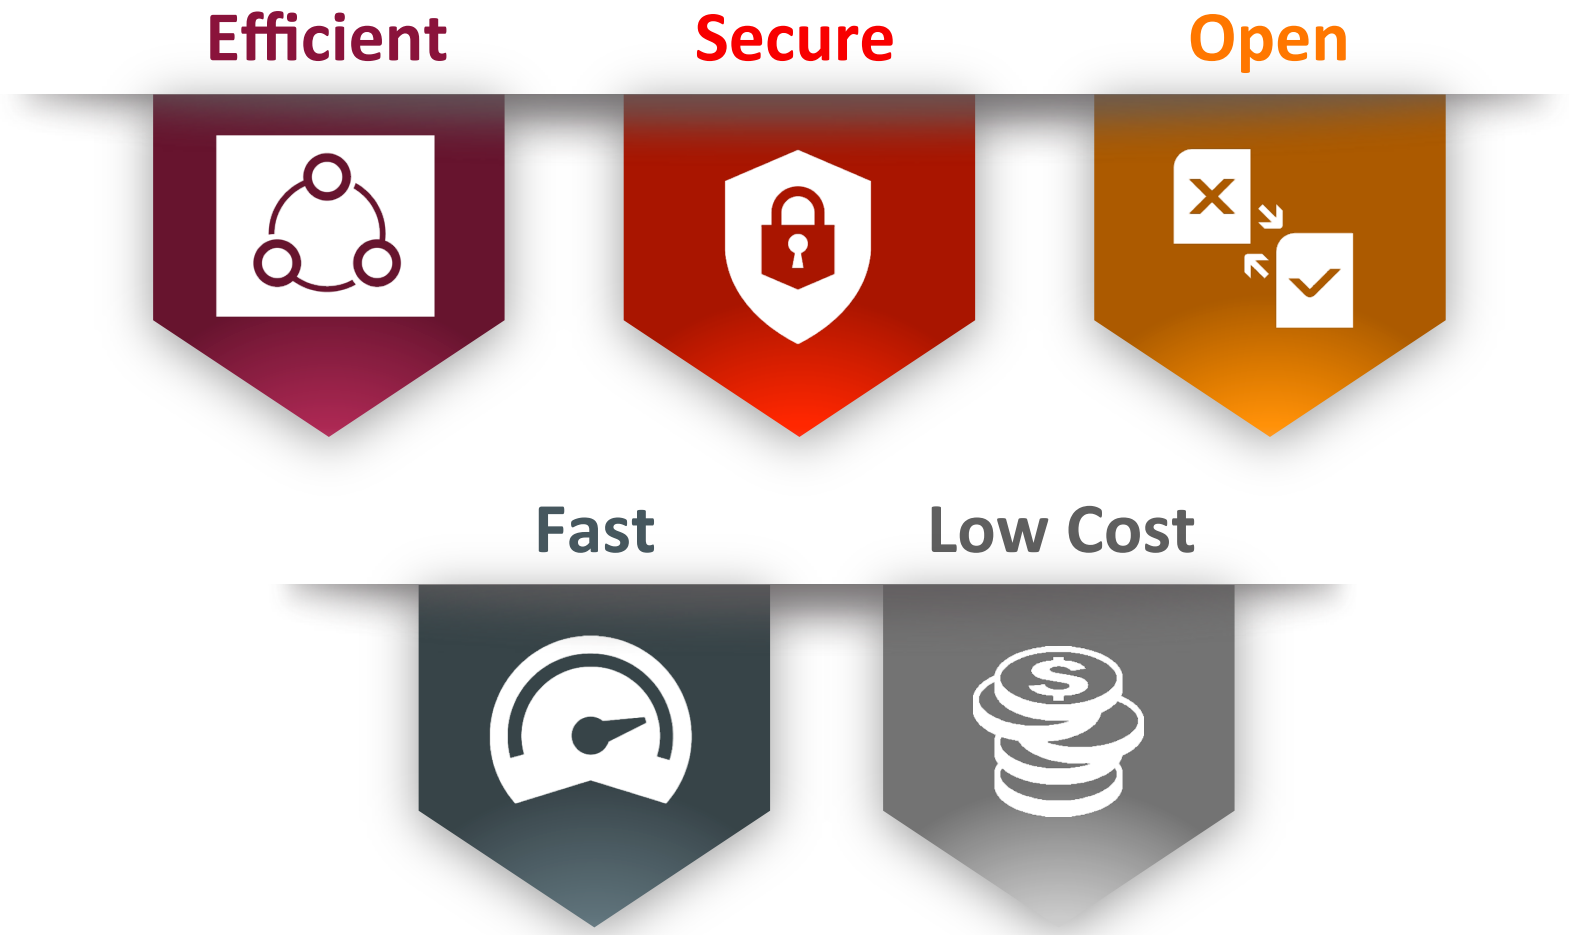

Hardware and Software
Engineered to Work Together