OPENFABRICS ALLIANCE

ENKI

ENKI | PrimaCloud: InfiniBand Enables Delivery of a New Class of Service

ENKI | PrimaCloud Overview

- **Mission:** Outsourced Operations delivered to SLAs in a enterprise-grade managed Cloud
- **Vision:**
  - Independently virtualized storage, compute, and transfer
  - Leading performance levels
  - Orchestration that enables 5-nines reliability
  - Secure, Easy provisioning, Built for compliance
- **Product:**
  - Public and hosted private clouds
  - Stand alone product, service franchising, and onsite private clouds.
- **Delivery:** Currently in full deployment.

Why Build PrimaCloud?

What Do CIOs Need From Cloud Computing?	Where Today's Infrastructure Clouds Fall Short
Service Levels (SLAs)	• 99.95% availability at best …only for hardware • Little or no consequences for downtime
Performance similar to a purpose-built datacenter	• Highly variable performance due to I/O cross-coupling • Commodity hardware leads to commodity results
Security / Compliance	• Little to no visibility to what's going on
Reduction in total cost of operations	• Self-service means additional staffing • Doesn't address *application* reliability • Software must be modified to run in the Cloud
Service	• Is there someone to call?
Trust in Providers	• Lack of transparency creates unpredictability

IB Satisfies The Need For Speed

Requirement	IB Delivers in ENKI | PrimaCloud
Service Levels (SLAs)	• End-to-End QOS • Pipe exceeds demand from multicore servers
Performance comparable to a traditional datacenter	• Infiniband has best in class performance • QDR 40Gbits per server available NOW!
Security / Compliance	• Infiniband Virtual Lanes for traffic isolation and QOS
Reduction in total cost of operations	• Low cost per port for switching • High port densities in switches
Oracle-Grade Storage	• Directly connect storage to compute via IB fabric • SRP over IB is very fast • Takes advantage of **Oracle Sun 73xx/74xx** *SSD-accelerated* storage

ENKI Cloud Version 1
Our First Try At The Cloud

Architecture:

- **Multiple Commodity Servers**
- **Xen hypervisor**
- **Gigabit network**
- **Separated control/utility and public networks**
- **Storage in the compute nodes**

Advantages	Problems and Disadvantages
• Inexpensive • Rapid Deployment • Low Barriers To Entry	• Virtualized workloads interfere with one another • Non-deterministic results prevented delivering believable SLAs • Poor Network Performance • Two switching fabrics with multiple single points of failure

ENKI Cloud Version 1.5
Attacking the I/O Bottleneck

Changes to Version 1.0:

- **Replace gigabit network fabric with Xsigo I/O Director and DDR IB interconnect**

Advantages	Problems and Disadvantages
• Much improved network performance • More consistent results enabled offering higher SLAs • Vastly improved stability as Xen matured	• Very expensive! • Xsigo capabilities exceeded requirements of use case • Performance of a virtual lanes limited to slowest attached virtual device • Limitation on total number of vlans

ENKI Cloud Version 2 (PrimaCloud)
Increasing uptime & performance

Architecture:

- **VMWare ESXi**
- **Large Servers**
- **QDR IB Interconnects**
- **IB Connected SAN**

Advantages	Problems and Disadvantages
• Extreme performance • Improved stability • Proven to run enterprise loads • Ability to guarantee *application* performance • Access to best-in-class management suites • Enterprise compatibility • Market leading technology	IB not fully supported • Known bugs cannot be fixed • Support on best-effort basis only • Features delayed • Full performance not realized

PrimaCloud

4 Years of Lessons Learned

- Storage, compute and network as separate independently scalable fabrics
- Scale and manage each fabric independently
- Orchestration is key to delivering quality services at scale
  - Simplify, simplify, simplify
  - Keep humans out of the mix
  - Autonomous remediation

PrimaCloud

4 Years of Lessons Learned

IB is critical to the success of PrimaCloud by enabling best-in-class SLAs

- Keep storage suppliers and storage consumers out of IP networking
- Multipath I/O is the rule rather than the exception
- Network fabric I/O capability must exceed the server's ability to saturate it
- Bridge to IP only when necessary

PrimaCloud Reduce Costs of Meeting SLAs in a Managed Cloud

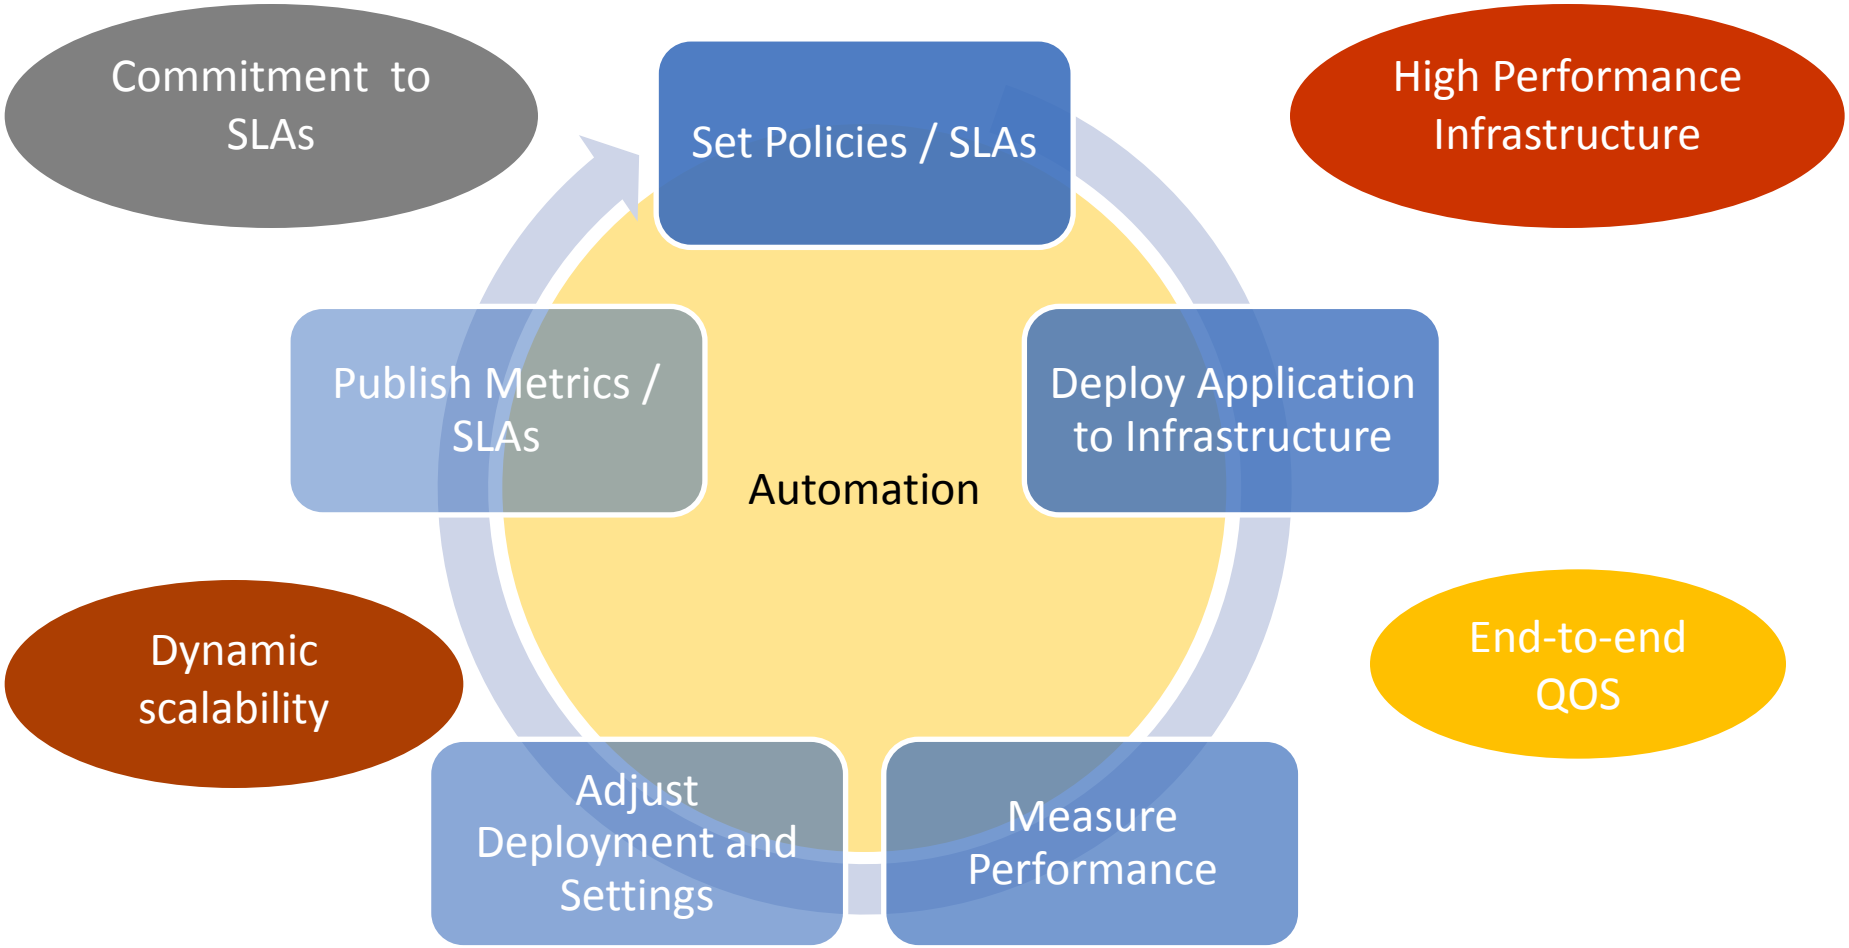

What is Automation?

- Measurement (monitoring) of system parameters
- Decisionmaking based on measurements and goals (SLAs)
- Take action based on decisions
- Optional user output (status, alerts, logging)
- Entered as a flowchart (workflow)
- Automatically executed

Application Deployment Process

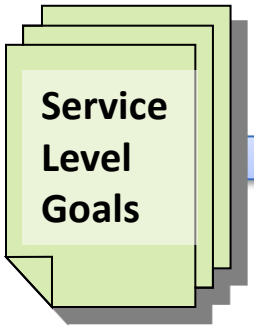

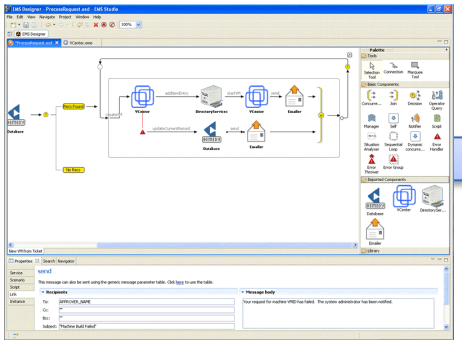

Create Workflow

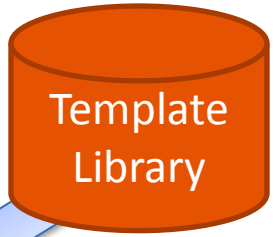

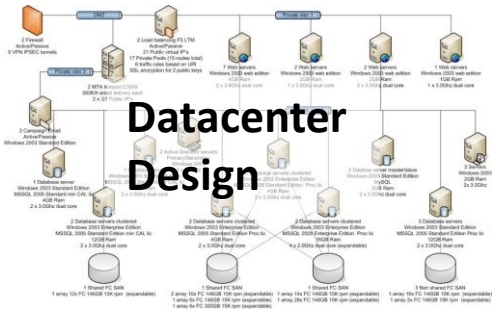

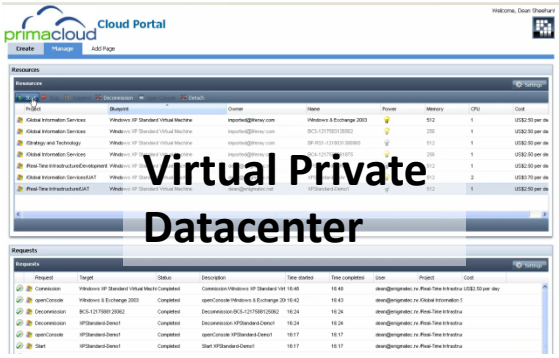

ENKI Total Cost of Operations Savings

… a real-world customer example of automation

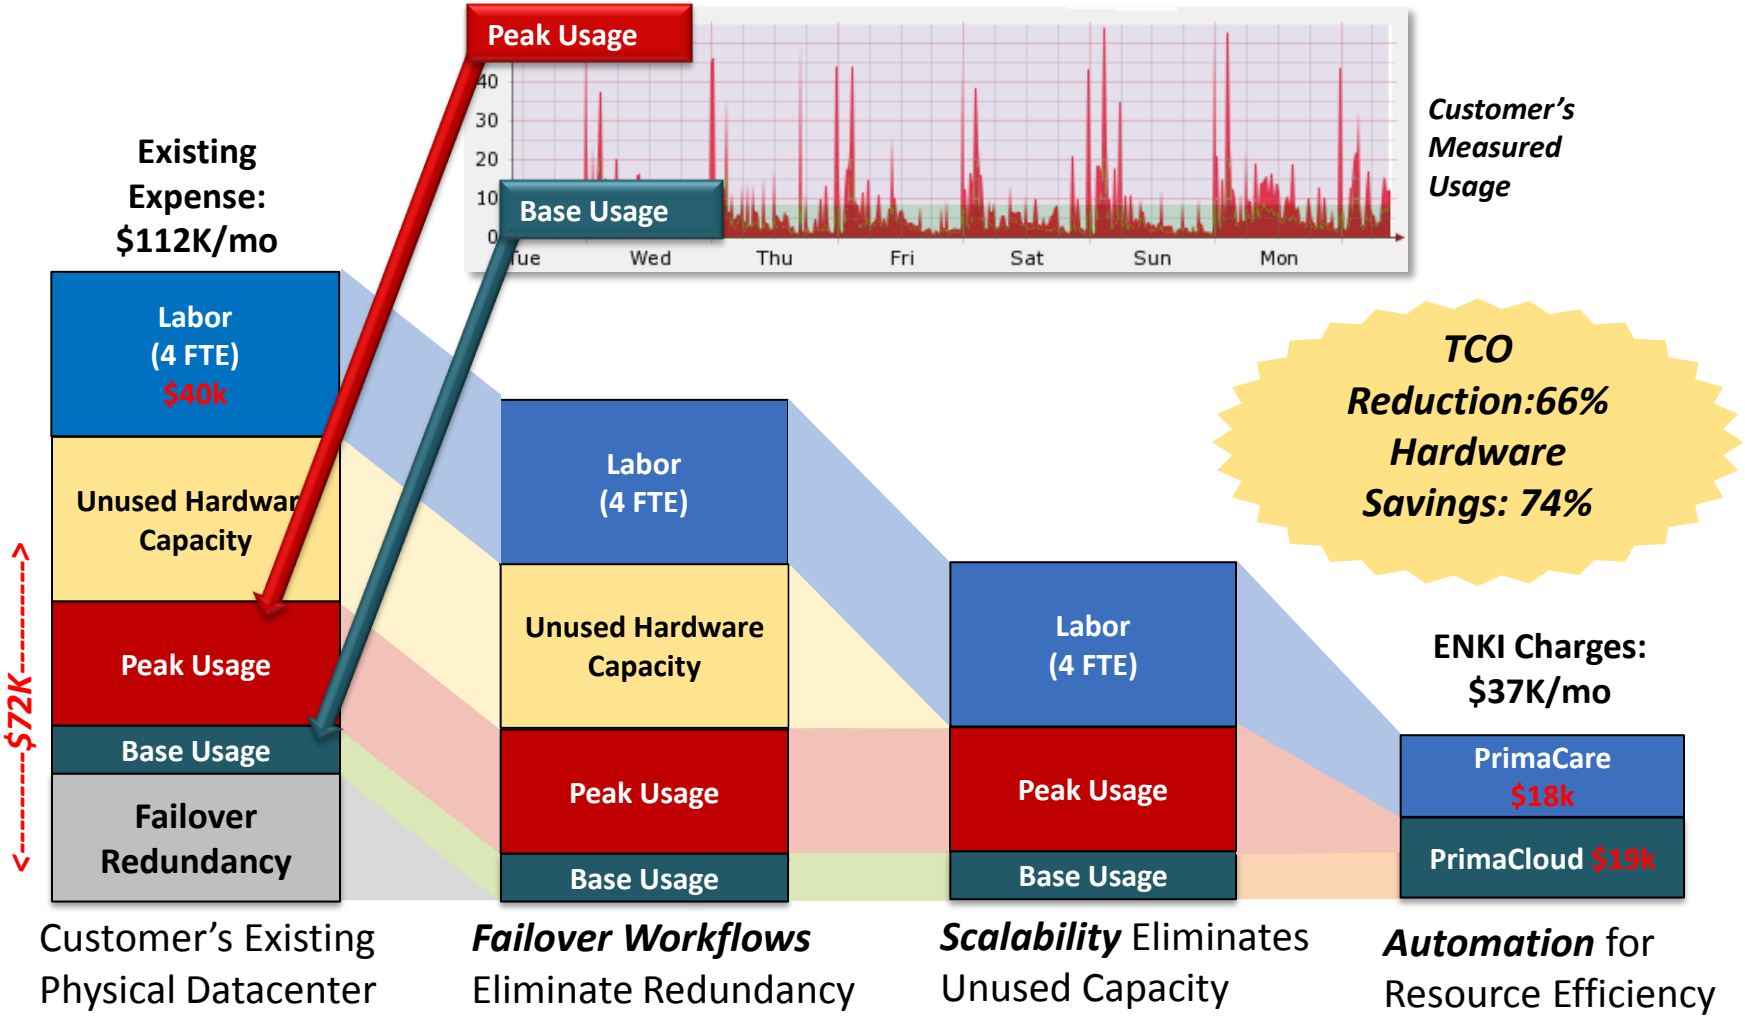

How OFA Can Drive Adoption of the IB-enabled Datacenter

- Migrate VMWare IB drivers into OFED
- Key IB stakeholders must solve the ‘enabled’ status of IB in VSphere
- Work with VMWare to support full IB wire speed to hosts and VMs
- Single point of contact for IB drivers
- Consistency across platforms
- Improve IPoIB performance