NetApp®

Go further, faster®

Italian for Beginners: *The Next Steps for SLO-based Management*

Lakshmi N. Bairavasundaram
Gokul Soundararajan
Vipul Mathur
Kaladhar Voruganti
Steven Kleiman

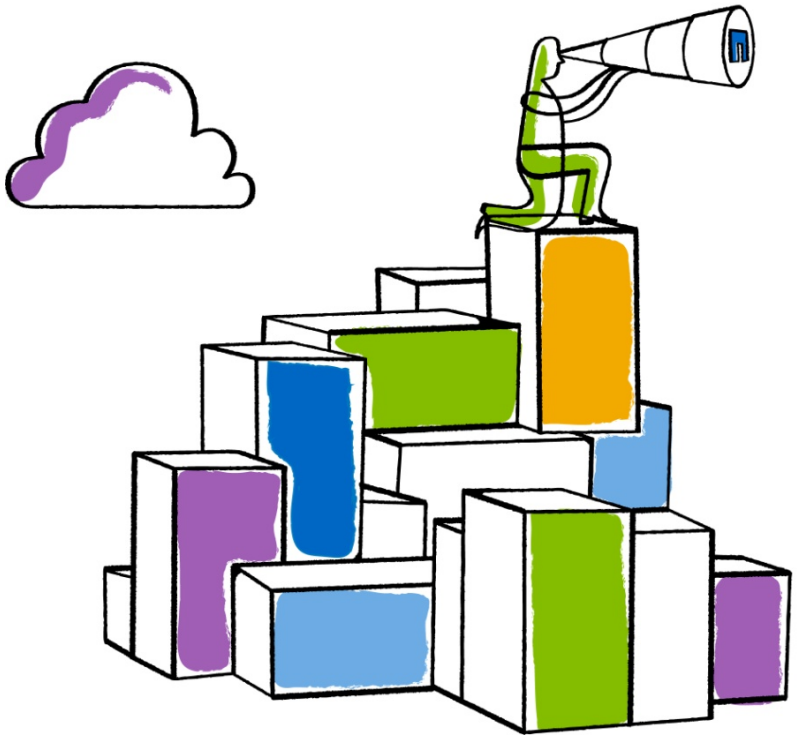

We need Service Automation!

- Systems are complex, therefore
  - low operational efficiency
  - low management efficiency
- Research has shown that
  - Service automation can help, but
  - Done a poor job in getting them into products
- We show some directions to fix the problem

Today's Data Center Trends

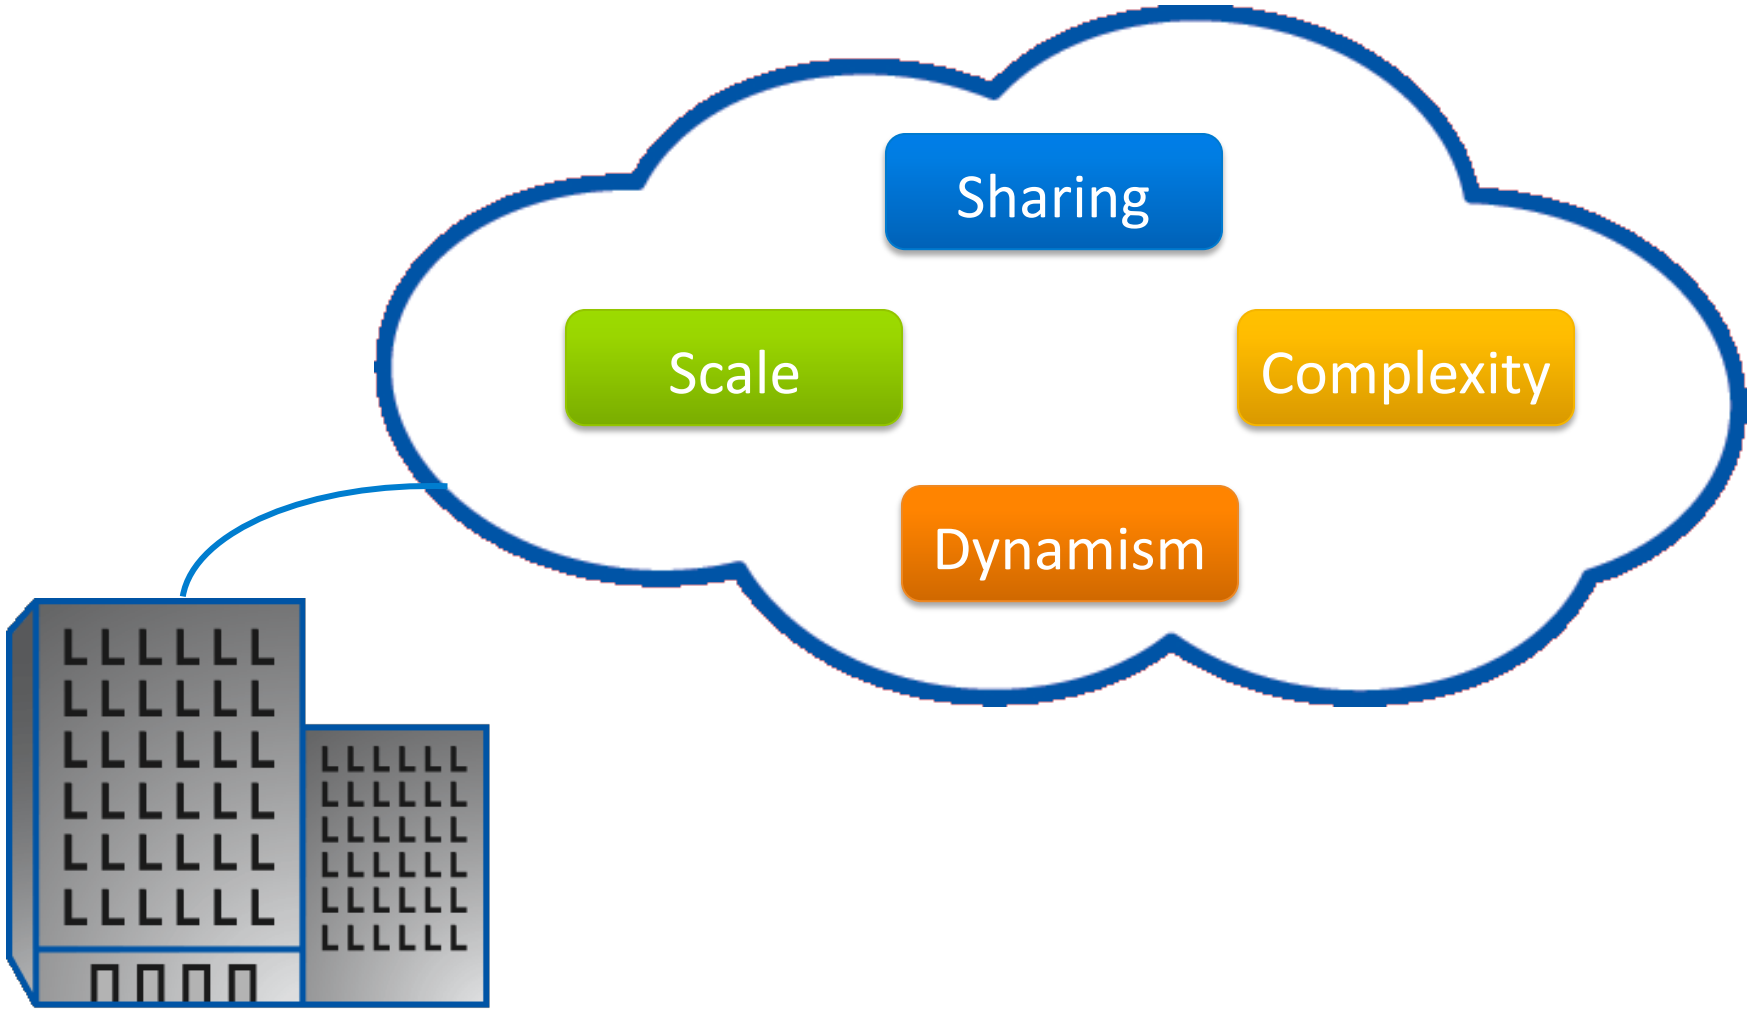

Trend 1: Siloed to Shared World

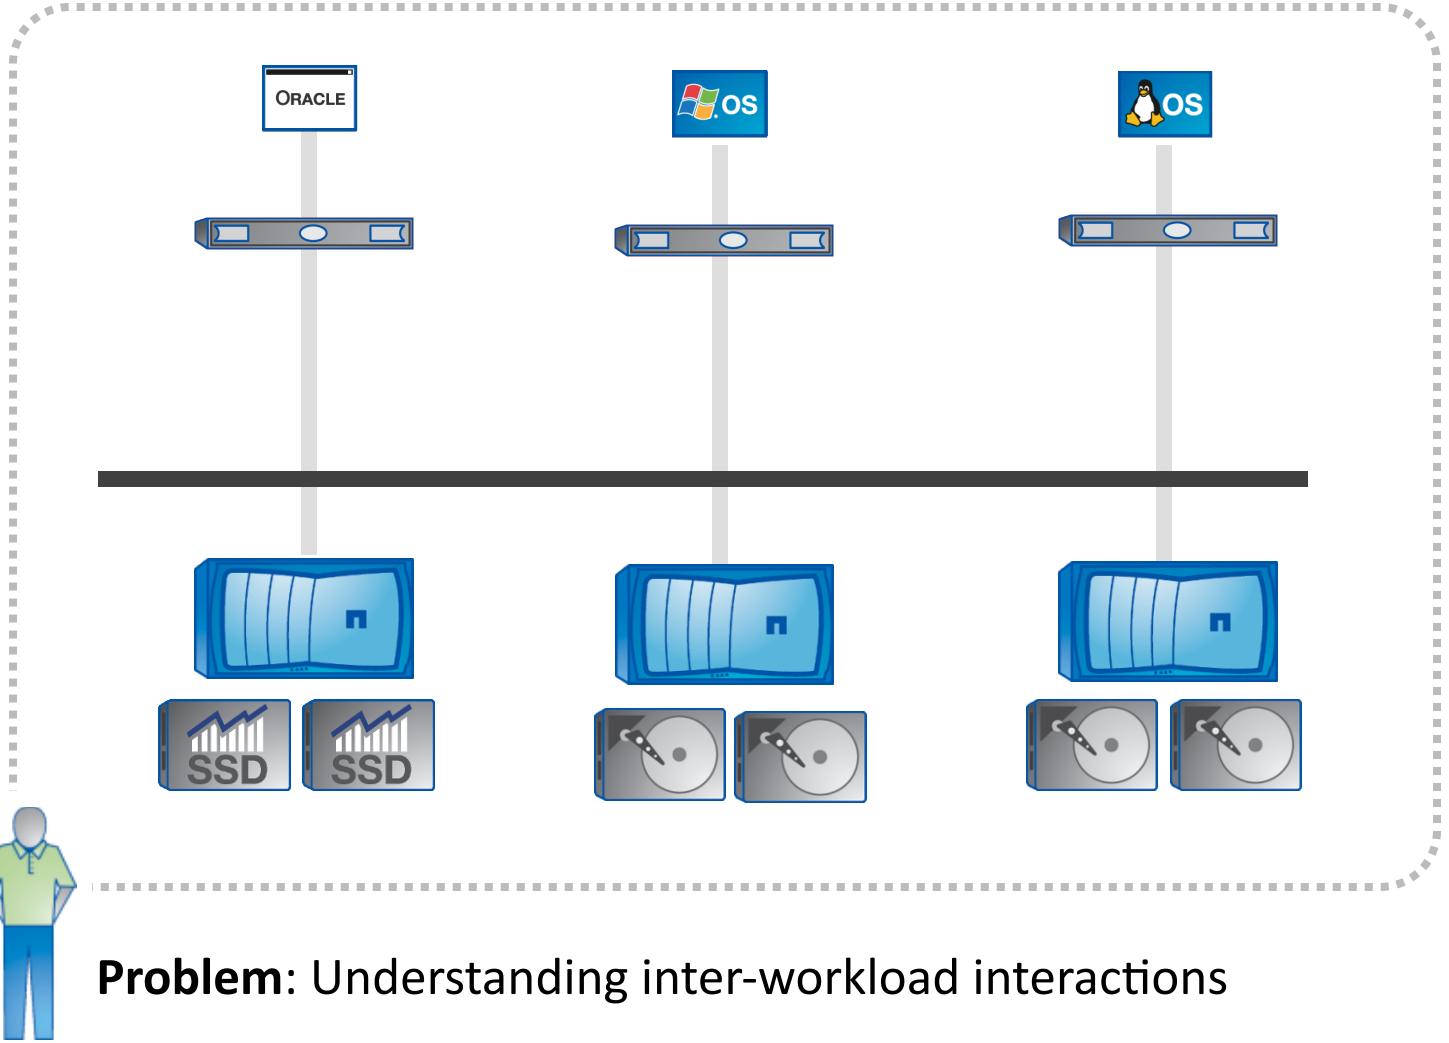

Problem: Understanding inter-workload interactions

Trend 2: Configuration Complexity

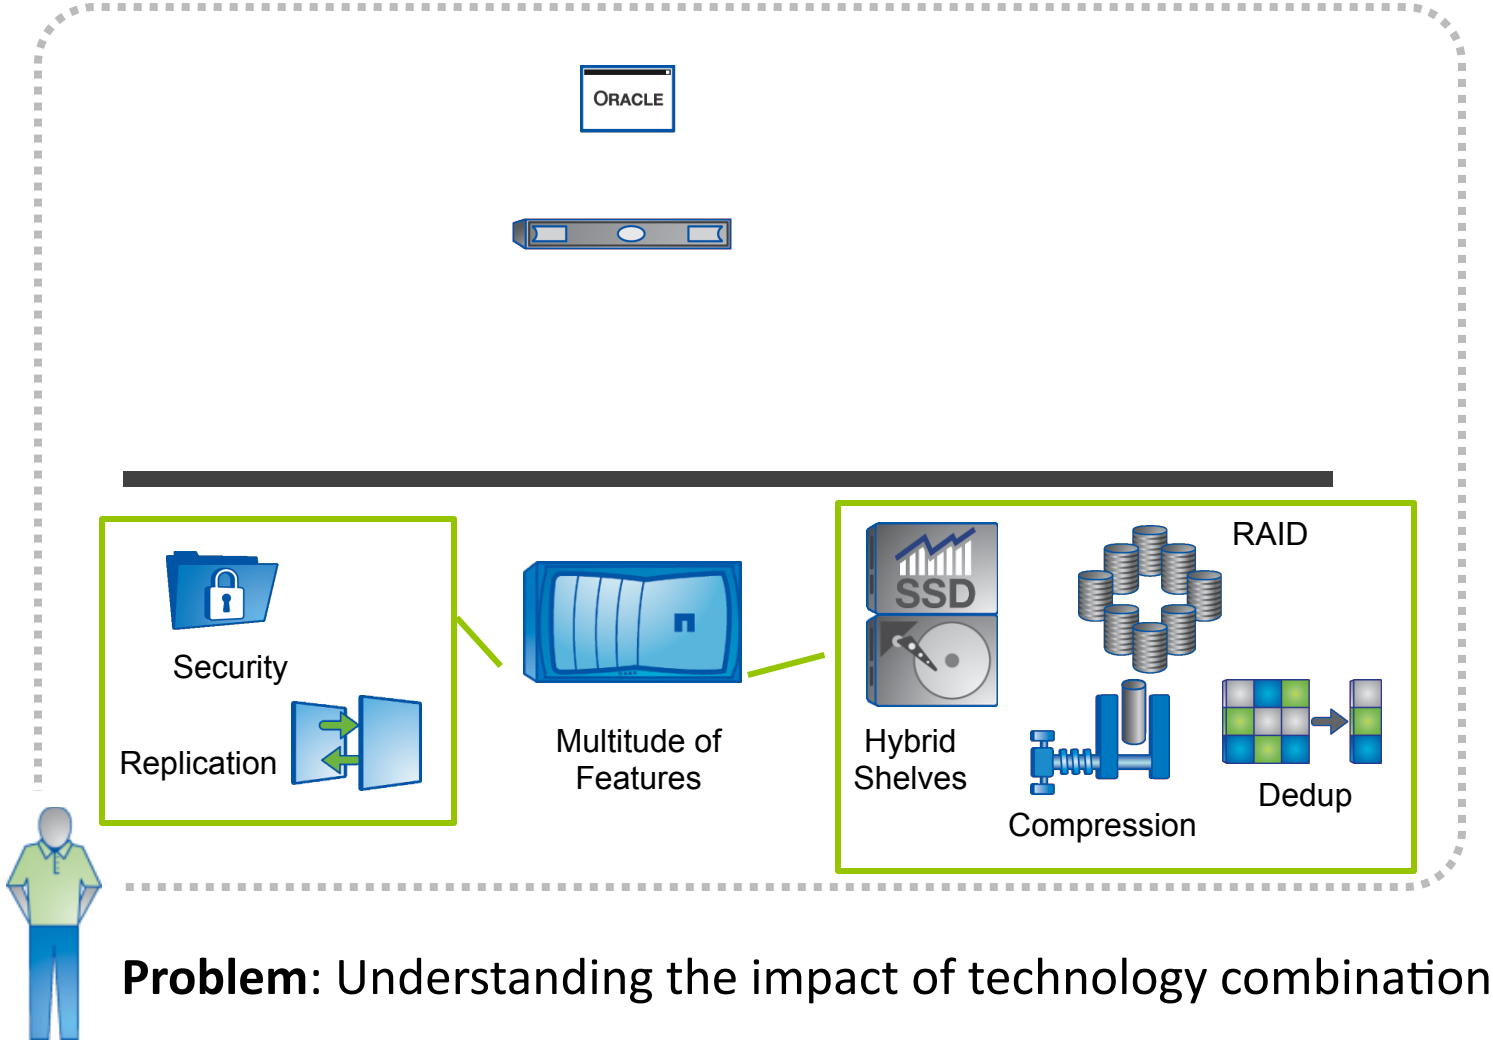

Problem: Understanding the impact of technology combinations

Trend 3: Huge Increase in Scale

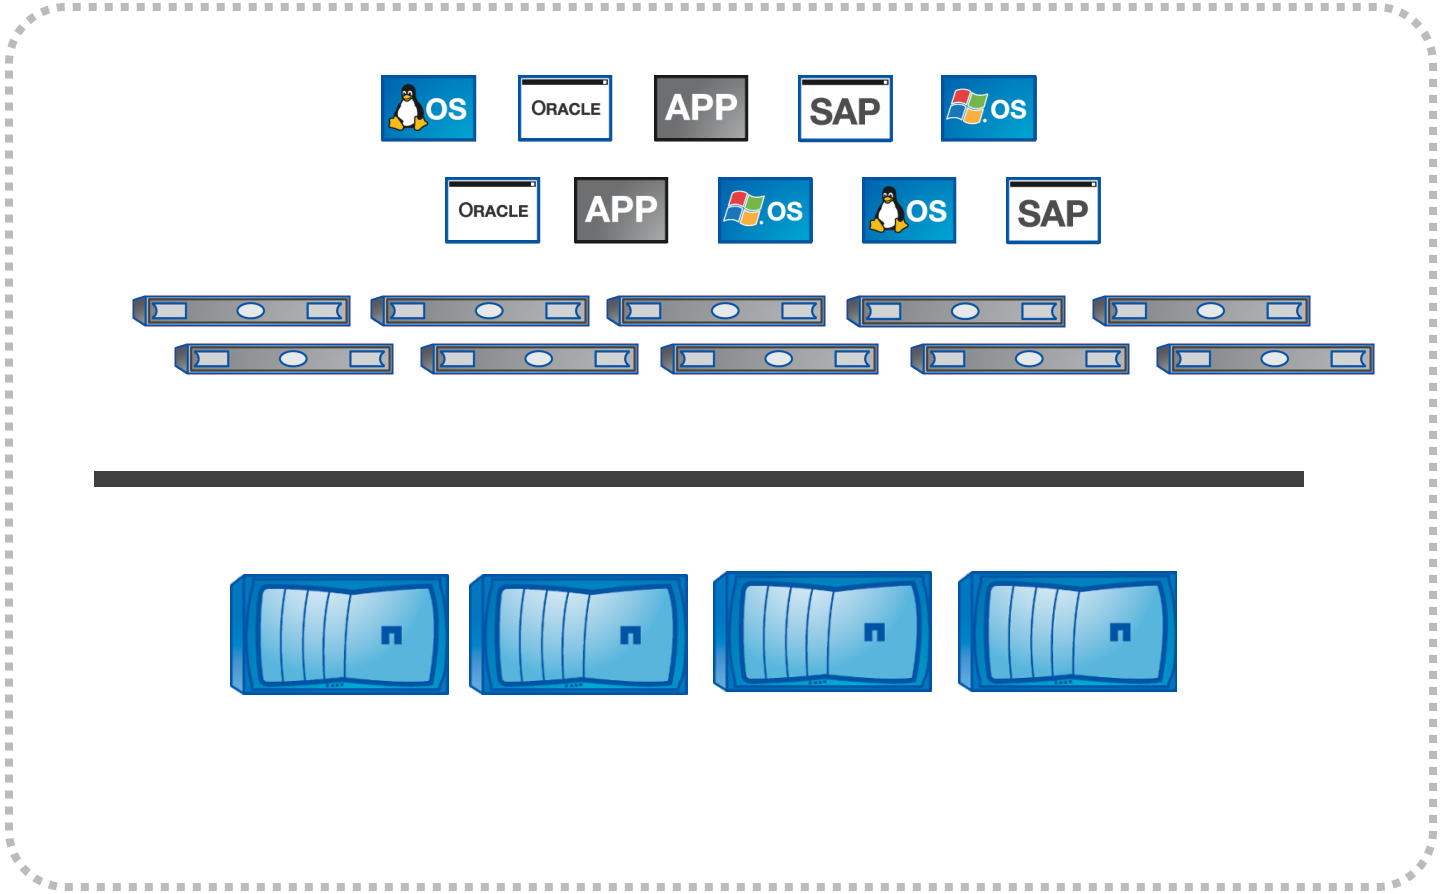

Problem: A large number of administrators needed

Trend 4: Applications are Dynamic

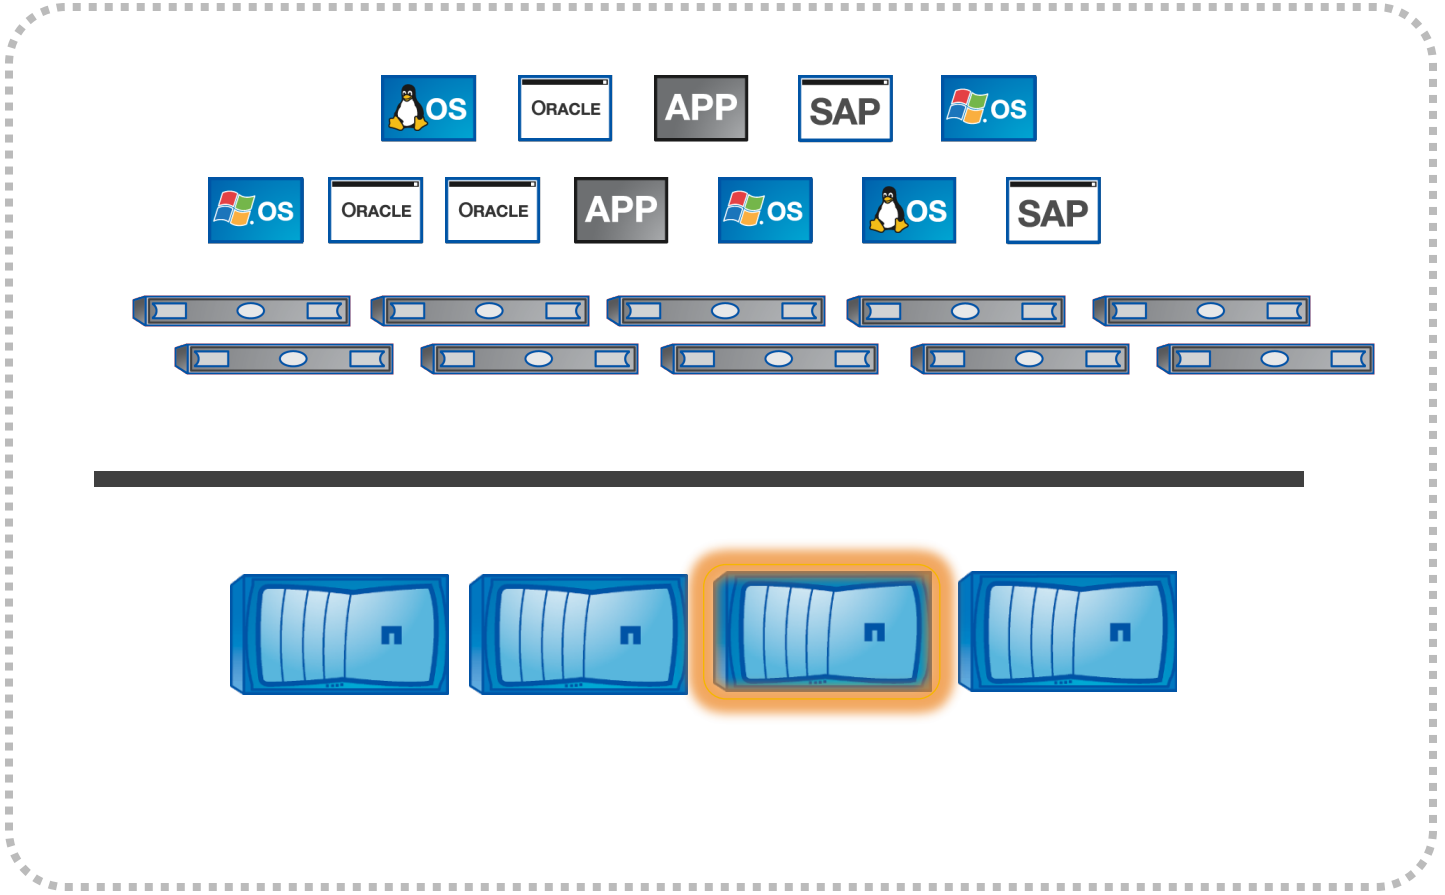

Problem: Need to handle sudden changes in resource requirements

What happens in industry?

- How do datacenters approach the problem today?
  - Provision for peak demand
  - Do not use all features
  - Have many administrators

- They spend more!

Solution: Automated Management with Service Level Objectives (SLOs)

- Specification of application's requirements in technology-independent terms
  - Describe application needs at different levels of the software stack
- Attributes
  - Performance, e.g., avg. I/O latency, IOPs
  - Capacity
  - Reliability and Availability e.g., RPO/RTO
  - Security and Compliance

Solution: Service Automation (MAPE Loop)

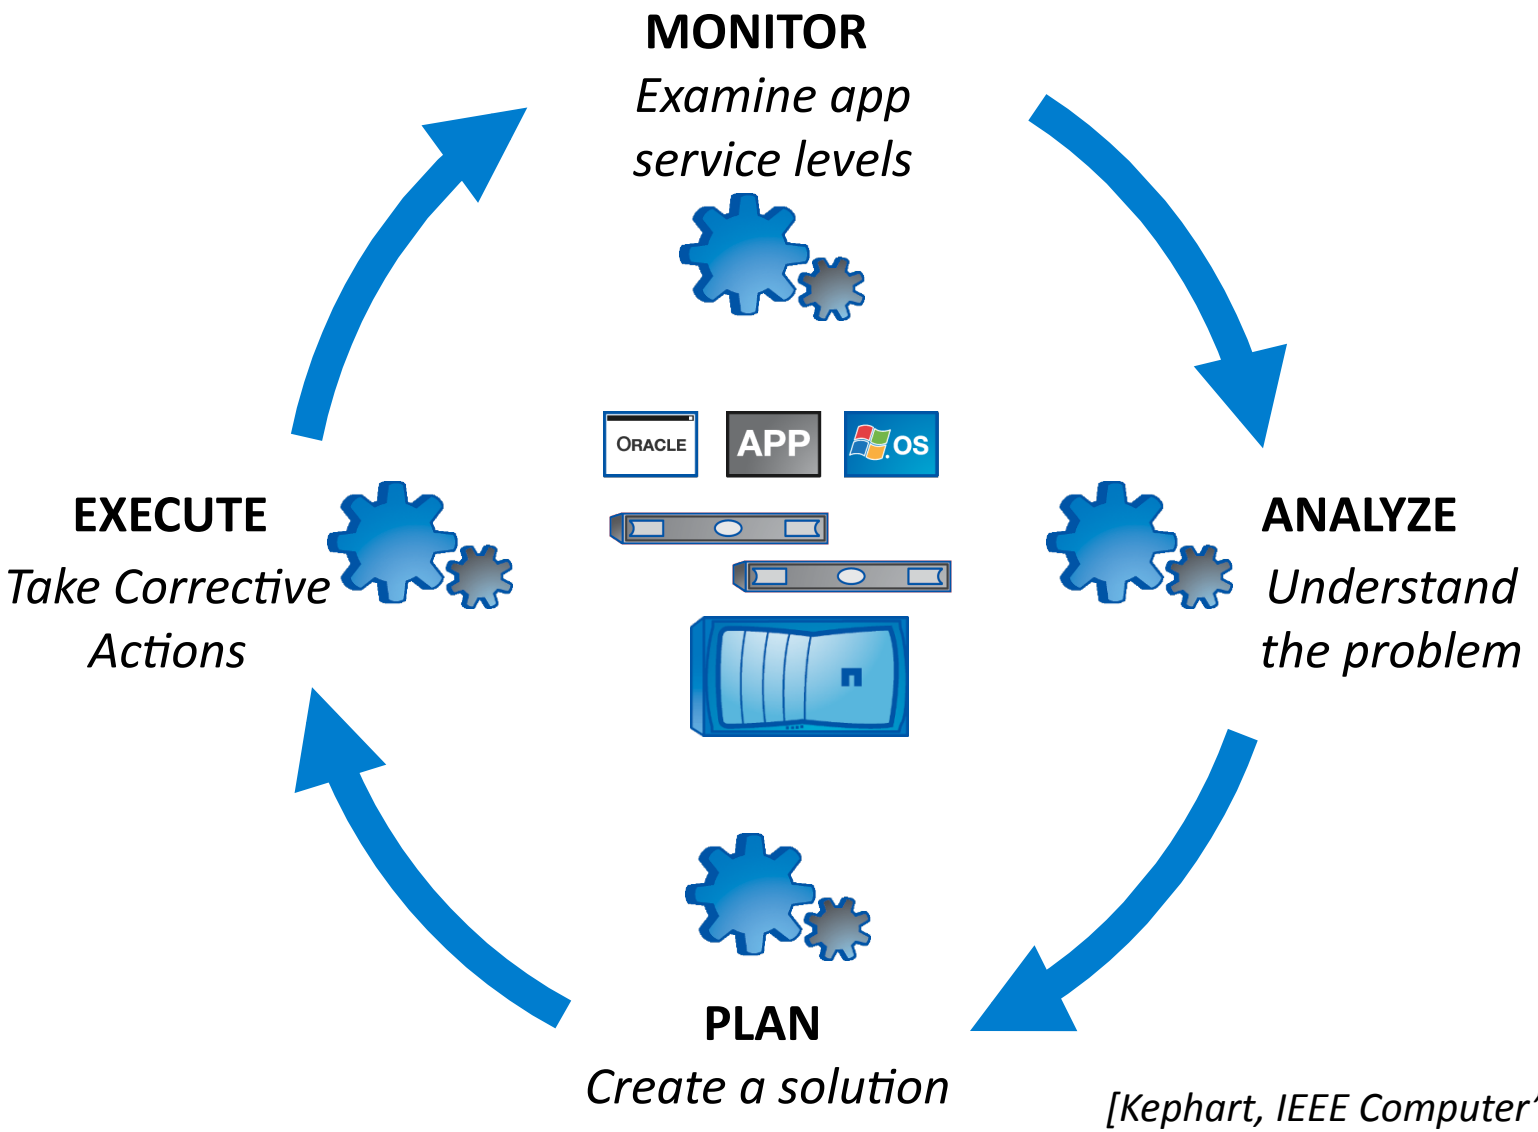

[Kephart, IEEE Computer'03]

Difference between Research and Product

- Evaluated along multiple dimensions
  - SLO Specification
  - Monitoring and reporting
  - Impact analysis
  - Techniques to handle SLO violations
- Products lacking SLO-based features
  - Use technology dependent SLO attributes
  - Lack of multi-dimensional SLO specification
  - Lack of impact analysis and automated planner

Reasons for Slow Adoption

- John Wilkes' conclusion *[Wilkes,OSR-43,2009]*
  - Need to convince that the system can be trusted
  - Simple-to-use, predictable
  - Open about decision making process

- There are additional reasons too
  - Difficult to specify SLO requirements
  - Need to build performance/reliability models
  - High cost of correcting modeling errors

Reason #1: SLO Specification Complexity

- Administrators may not know application's requirements precisely
- Need to combine: performance, protection, security, cost, etc. to provide final specification

- Simplicity wins
  - Easy to specify bandwidth shares, or priorities

Reason #2: Modeling Errors

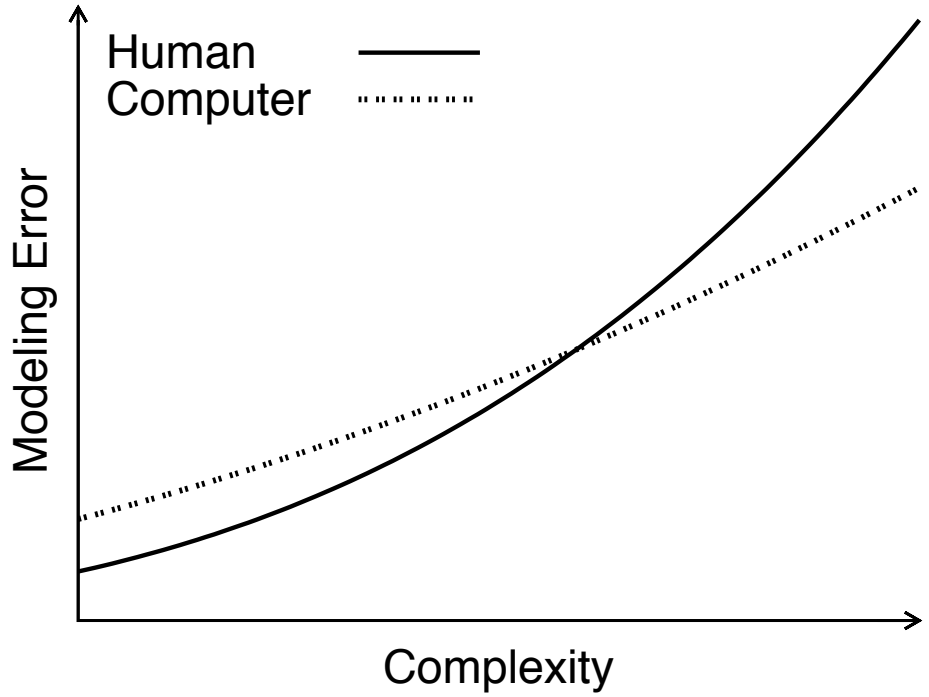

- Need time to build performance models, e.g., white-box models or black-box models
- Human expertise may be better that the moment

Reason #3: High Reconfiguration Cost

- Buy new storage system, or
- Migrate dataset
- Simpler to provision for peak load

Improving Adoption of SLO-based Management

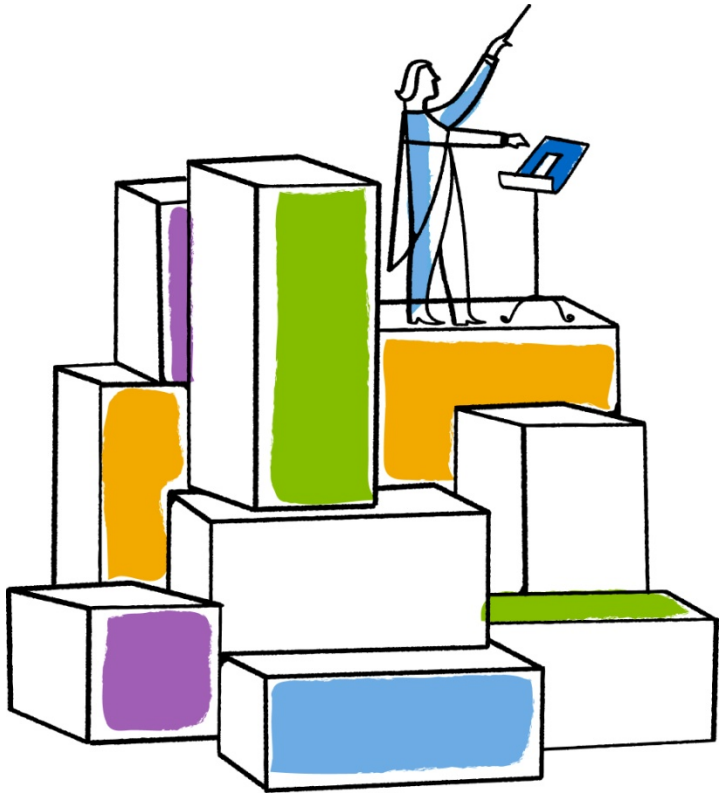

Direction #1: Process, Not Product

- More than product development
  - Easing requirement specification by users
  - Building and updating system models
  - Post-sales support for products

Direction #1: Qualified SLOs

- Identify SLOs for popular applications
  - Leverage expertise of customers/partners
  - Translate application requirements to storage requirements

- For example,
  - Microsoft Exchange with N mailboxes
  - Configured using best practices documents as an atomic unit

Direction #1: Without Qualified SLOs

Corporate Mail

Fast, highly-reliable, and user error protection

Latency=*10ms*
Thgpt=*100MB/s*
RTO=*0 for up to 5 days old, 4 hours for older*
…

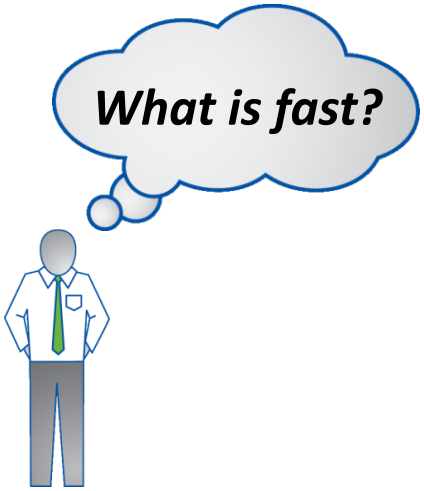

FC Disks
HA
RAID
Snapshots

RAID
Snapshots
Dedup
Compression
SATA Disks
HA
FC Disks

Direction #1: Qualified SLOs

Corporate Mail

Fast, highly-reliable, and user error protection

QSLO-1
Latency=*10-20ms*
Thgpt=*100-200MB/s*
RTO=*0 for up to 5 days old, 4 hours for older*
…

QSLO-2
Latency=*20-50ms*
Thgpt=100-*200MB/s*
RTO=*4 hours*
…

RAID | Snapshot | Dedup | Compression

SATA Disks | HA | FC Disks

Direction #2: Low-Impact Reconfiguration

- Non-disruptive reconfiguration of resources
- Migrating small amounts of data
- Dynamic storage layout
  - Flash+HDD dynamic tiering *[Guerra et al., FAST'11]*
  - pNFS *[Shepler et al, RFC 5661]*
- Deploy caches
  - E.g., host-side caches, storage-side caches

Direction #3: Community Wisdom

- Leverage Qualified SLOs for comparisons
  - Augment internal models with additional data
  - Share improvements across multiple deployments

- Proactively advise customers
  - Alert on misconfigurations *[Wang et al., OSDI'04]*
  - Notify of faults and errors *[Bairavasundaram et al., SIGMETRICS'07]*
  - Guide towards best practices

Direction #3: Guide towards Best Practices

Corporate Mail

Corporate Mail

QSLO-1
Latency=*10-20ms*
Thgpt=*100-200MB/s*
RTO=*0 for up to 5 days old, 4 hours for older*
…

Latency=*12ms*
Thgpt=*150MB/s*
RTO=*0 for up to 5 days old, 4 hours for older*
…

RAID | Snapshot | Dedup | Compression

SATA Disks | HA | FC Disks

Conclusions

- We need SLO-based management
  - Moving towards shared infrastructure
  - Configuration complexity is increasing
  - Scale and dynamism are hard to manage

- The next steps to do
  - Develop Qualified SLOs
  - Build dynamic reconfiguration techniques
  - Leverage community wisdom

NetApp

Thank you