Capacity Planning: Engineering Infrastructure Performance

META group

Business and Technology Scenario

- Data center consolidation is forcing IT organizations to evaluate overall performance
- Budget pressures are forcing more prudent resource utilization and infrastructure purchasing
- Overall discipline in IT infrastructure planning has been dismal and needs to improve
- IT organizations must minimize capital expenditures while preserving good service
- Growing operational process maturity requires better accountability and information sharing across processes

Critical Issues

- Fine-tune IT services by optimizing infrastructure performance
- Address the challenges of capacity planning in distributed environments
- Technology solutions for capacity planning

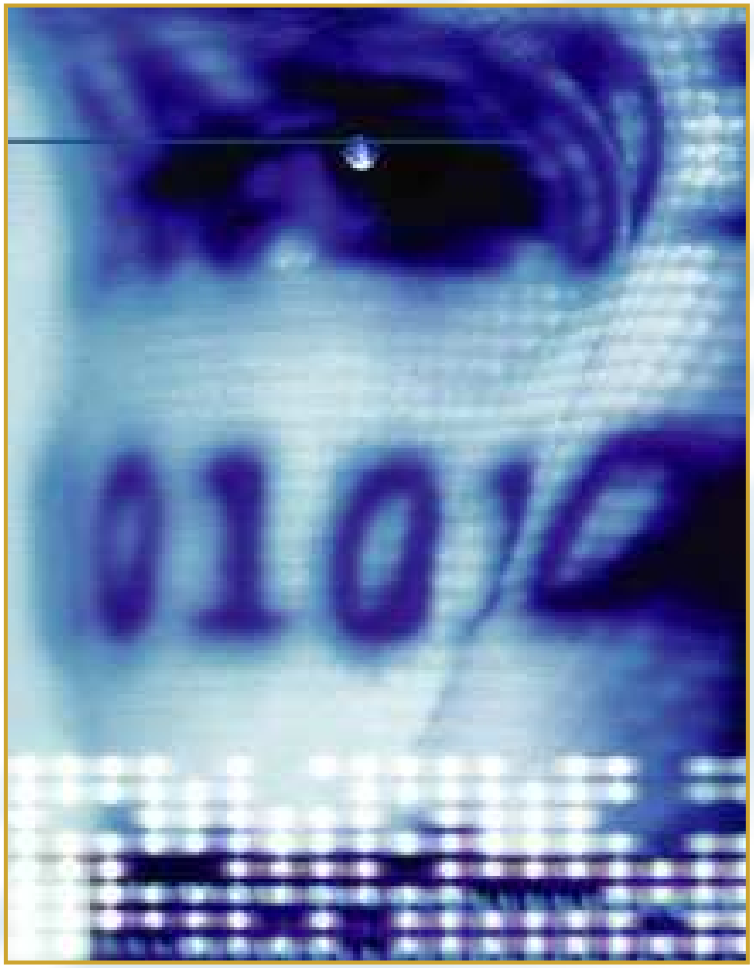

Optimizing Infrastructure Performance

- 2003 IAM study results on capacity planning
- Business drivers for capacity planning
- Preemptive performance planning approaches
- Engineering discipline yields superior quality and efficiency

Balance Capacity for Optimum Performance and Cost Effectiveness

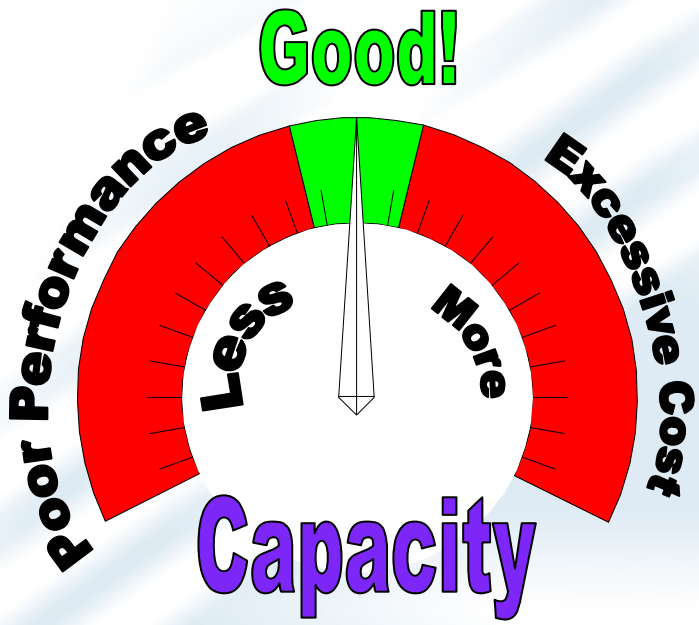

Balance performance and cost through a structured capacity planning process

2003 IAM Study Results

Capacity planning is the #1 critical issue for large enterprises — capacity planning spending plans:

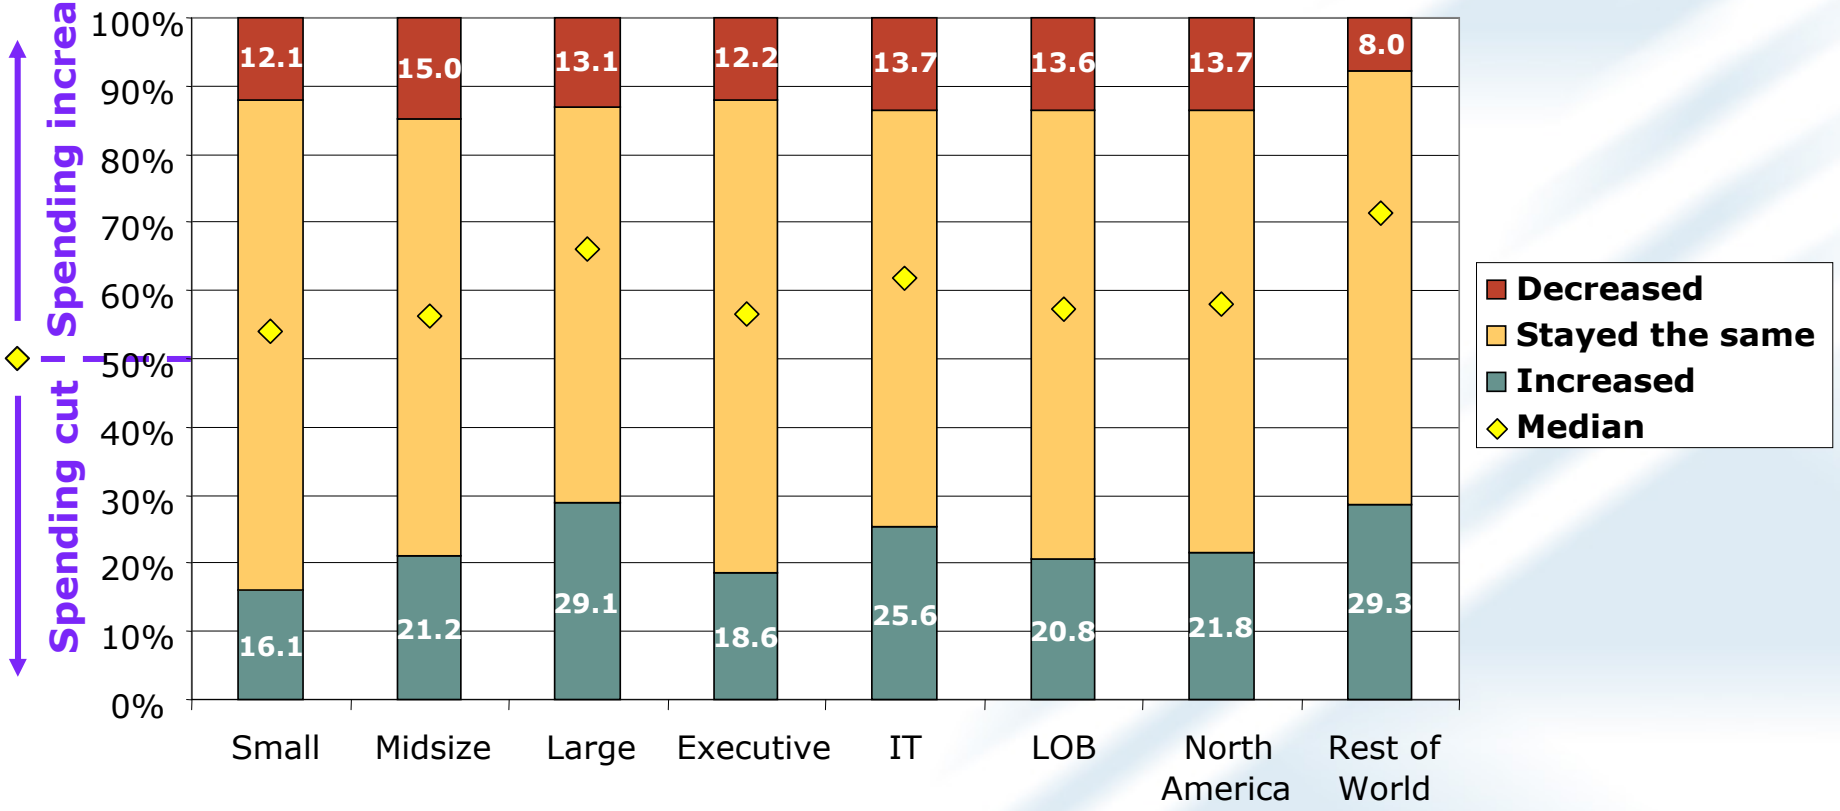

Budget for capacity planning technology and staff to address changing business demands

Capacity Planning Business Drivers

- Costs need to be reduced
  - Operational expenses
  - Capital expenses
- Resource utilization is low (data center higher)
  - Windows average is 10%-15%
  - Unix average is 20%
- Consolidation
- Growing operational maturity spawns strong capacity planning
- Excess capacity can no longer mask demands

Avoid sloppiness in provisioning of excess capacity in favor of engineering discipline

Quality Through Engineering Discipline

- Complex systems follow structured engineering preceding production
  - Autos, aerospace, semiconductors
  - IT systems are complex and must be treated likewise
- Engineering discipline prevents costly redesign
- Capacity planning helps fulfill this discipline

Feedback in Engineered Development Flows

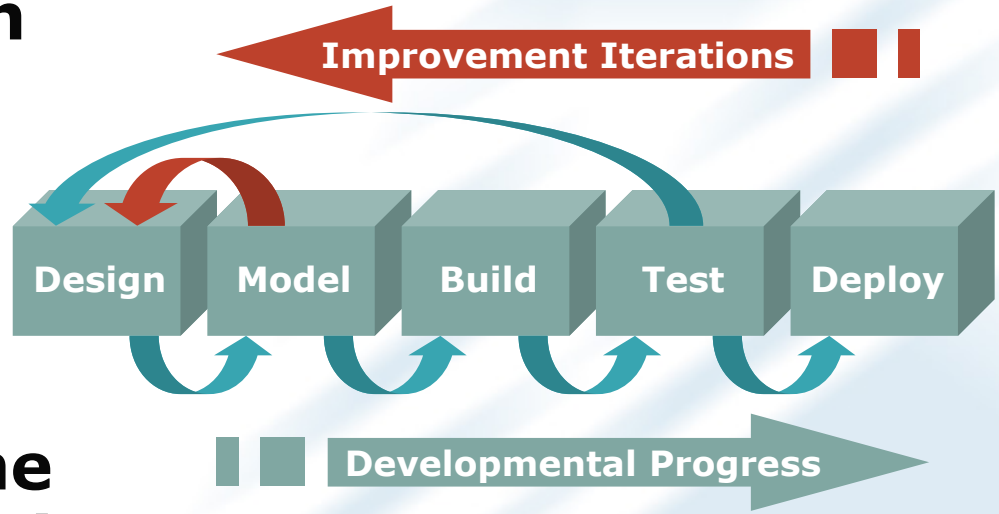

Engineering 101: Employ several evaluation approaches; no one method fulfills all needs

Performance Planning Approaches

- Incident analysis
  - Performance triggers
  - Long-term patterns
- Performance trending
  - Be careful — results can be misleading
- Modeling/simulation
- Evaluate both peak periods and averages

Classify Utilization for Consolidation Candidates

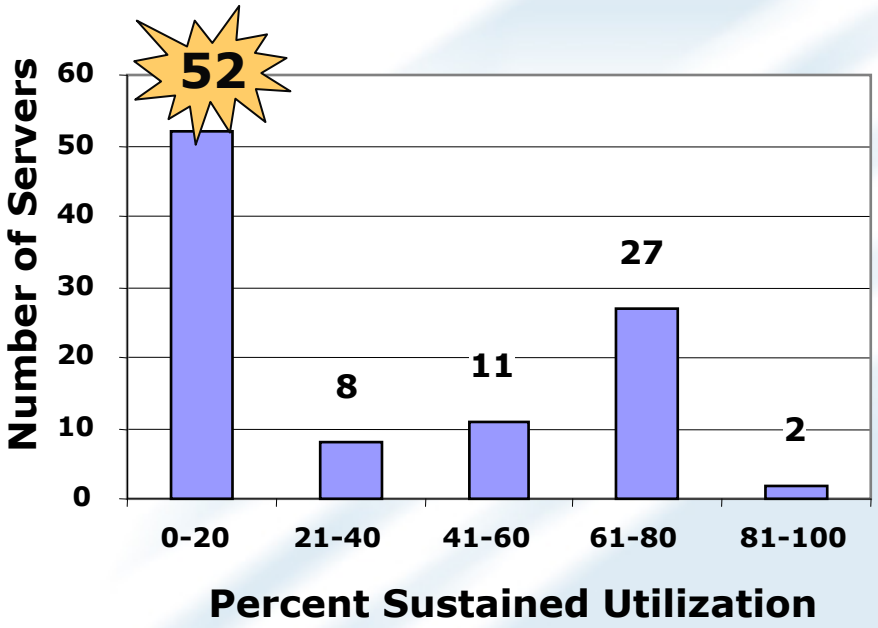

Combine people, processes, and technology to address distributed system complexity

Capacity Planning in Distributed Systems

- Distributed introduces new:
  - Complexity
  - Operational maturity
- Process development
- Server consolidation
- End-to-end efforts

Process Integration

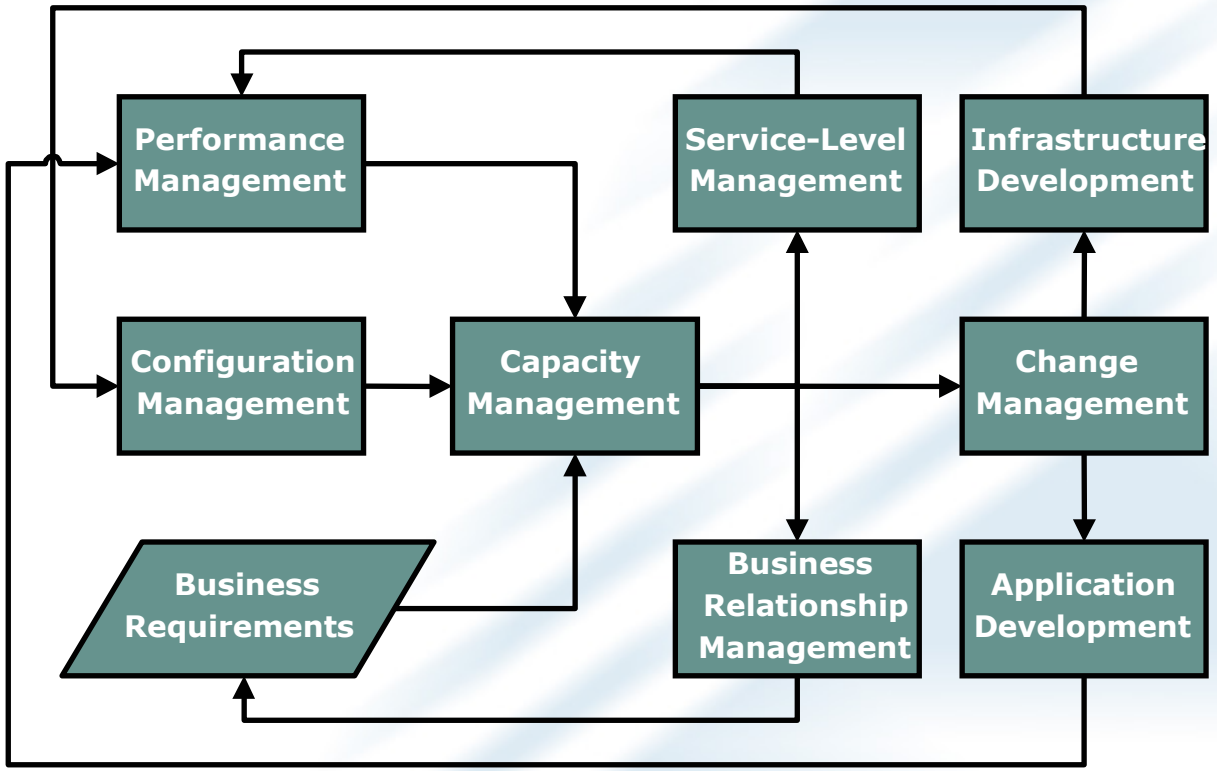

Leverage lessons learned from decades of legacy capacity planning

Borrow Structure From the Mainframe

- Mainframe capacity planning is mature but inflexible
- Far more distributed system complexity
  - Mainframe is more tightly contained and homogenous
  - Mainframe relationships are more static
  - Legacy system capacity provisioning is mature
  - Distributed systems will get more complex
- Legacy operations exhibit strong discipline
  - Distributed operations need to become more conservative

Build on established mainframe processes for distributed capacity planning

Process Development

- All three inputs are necessary
  - Many CP efforts miss business requirements
- Business requirements drive capacity assessments
  - Cycles indefinitely

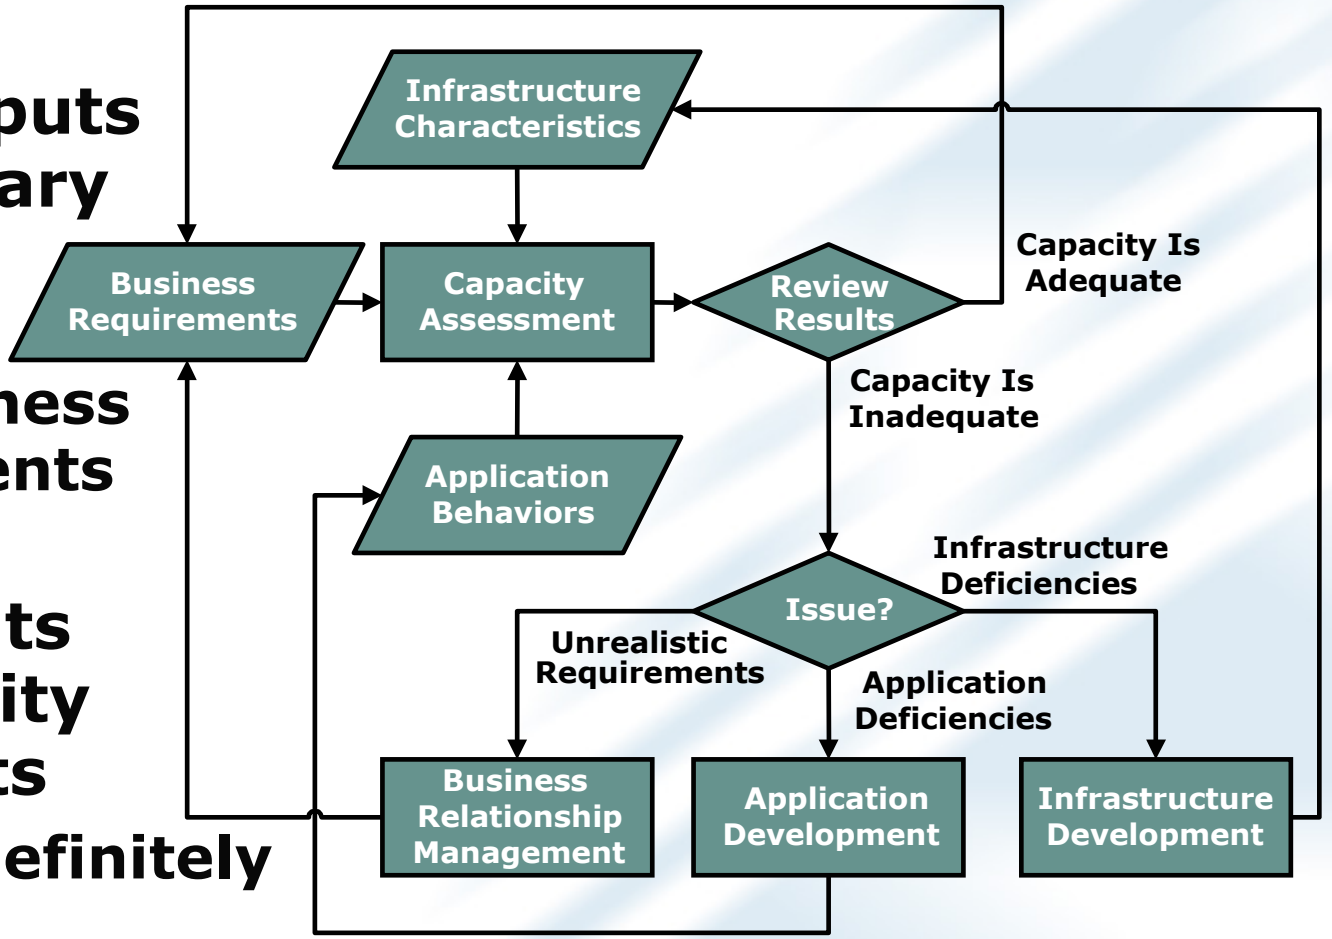

Optimize capacity and performance with a disciplined, consistent engineering process

Server Consolidation

- Provides optimization, but watch usage patterns
  - Geographic and LOB user communities
  - Capacity of network and other supporting factors
  - Application protocols
- Understand user impact
  - Blind consolidation can cause major business disruption

Workload Balancing via Consolidated Loads

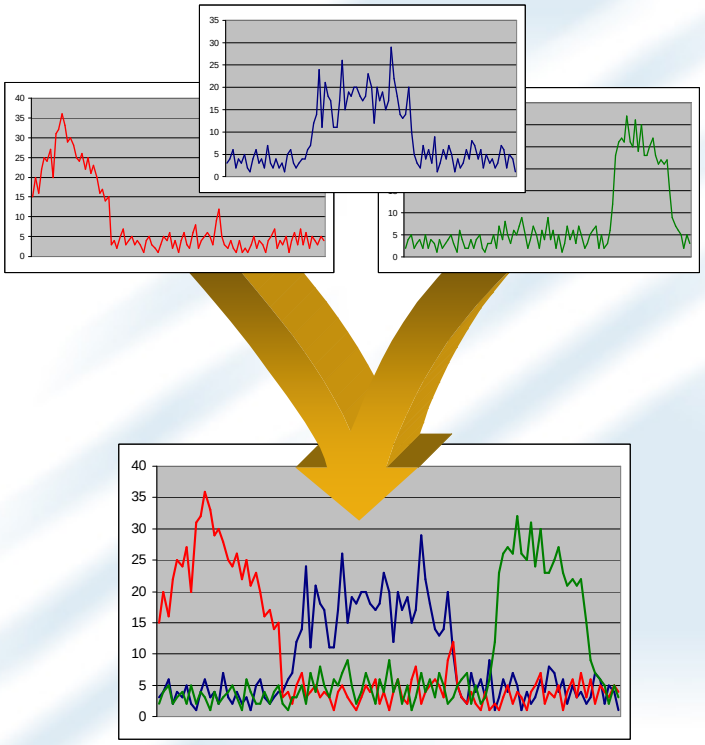

Consider all infrastructure, applications, and users when embarking on consolidation

End-to-End Capacity Planning

- A focus on technology silos is the norm
  - A myopic history of silo-based operations
  - Technology limitations
- End-to-end is the desired state
  - Server, storage, and network are merging
  - Endpoint inclusion will follow
- Adaptive organization and Web services will change everything
  - Highly dynamic relationships
  - Capacity on demand will automate some tasks

Optimize silos until capacity planning technology solutions evolve for end-to-end

Capacity Planning Technology Solutions

- Understand the "magic" within planning tools
- Application-oriented infrastructure tuning
- Modeling and simulation
- Vendor landscape

Capacity Measurement Accuracy Is Imperative

Acquire sufficient platforms for the intensive mathematical processing demands of tools

The Inner Workings of Planning Tools

- **Statistical analysis**
  - **Based on advanced performance management**
  - **New innovations emerging from astrophysics, quantum mechanics, and bioinformatics**
- **Queuing theory models data flows through a system**
- **Discrete simulation is slow, but extremely precise**

Gain Visibility Into Internal Technology Functions

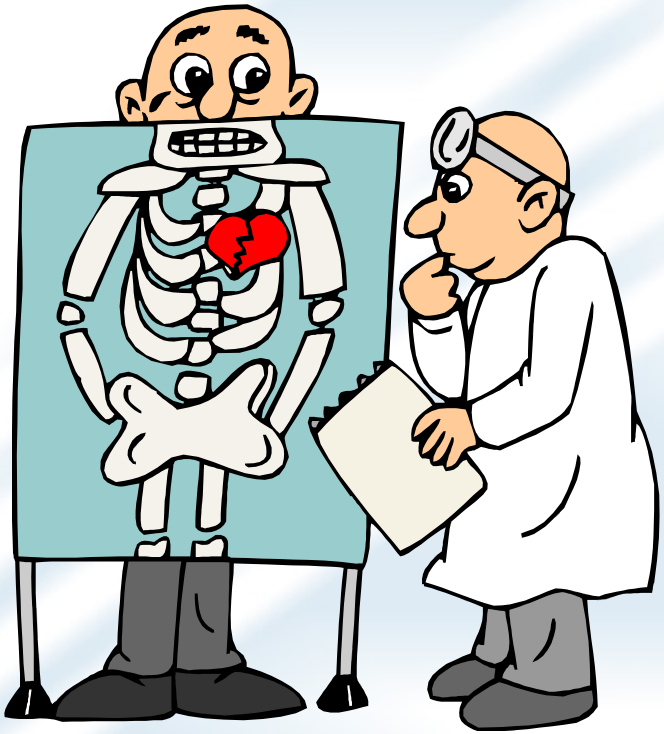

Blend capacity planning skills in operations research, infrastructure, and applications

Application Orientation

- Infrastructure decisions serve applications
  - Interactions within and across infrastructure impact results
- Application patterns must be understood by CP tools
  - Complex applications require the best tools
- Limit efforts to genuine business requirements

Applications Are the Tangible Service to Users

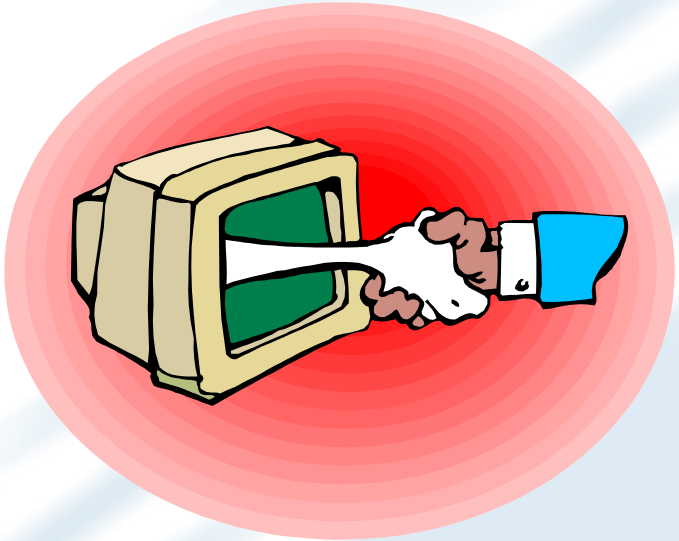

Build software models of application behaviors and infrastructure attributes

Modeling

- Software simulation and emulation of inputs
- Difficult, but high return
  - Getting easier
- Accuracy is critical
  - Infrastructure characteristics
  - Application behaviors
  - Business requirements

Simplified Modeling Process

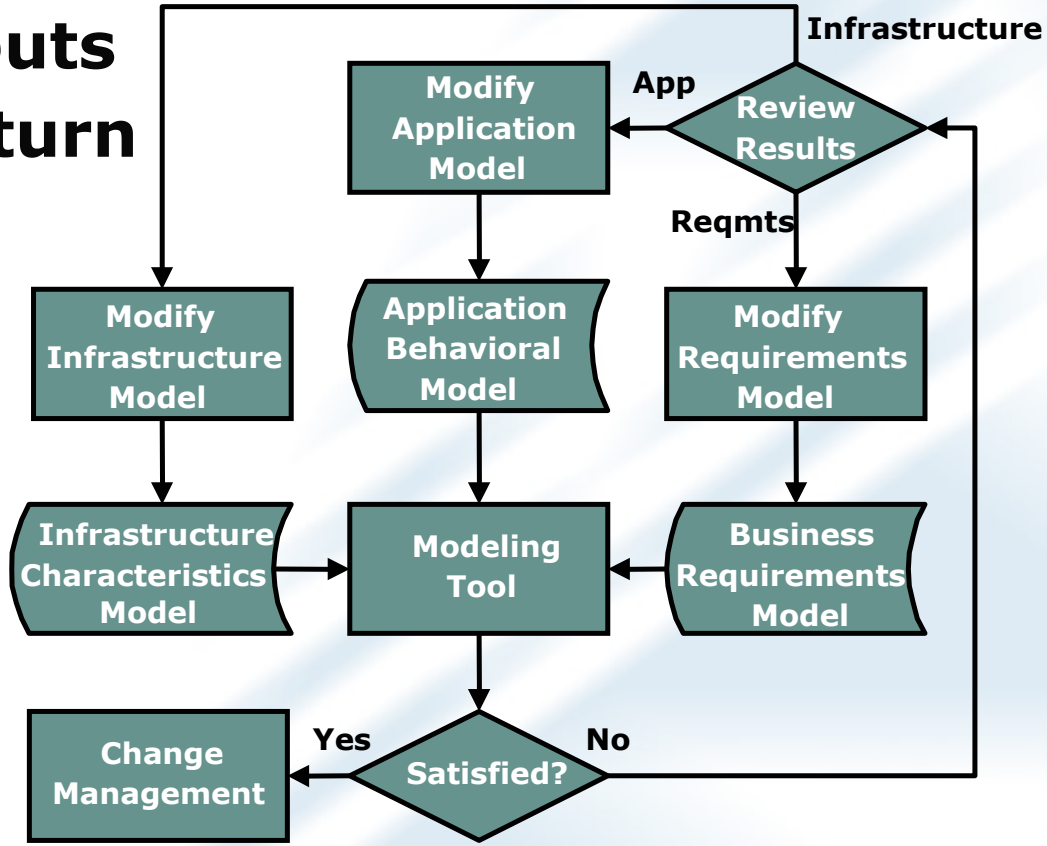

Seek capacity planning vendors with some form of modeling and simulation

Vendor Landscape

- Most vendors are small
  - BMC is the only serious large vendor
  - Expect other large vendors to enter
- Silo-centric offerings will become end-to-end
- M&A activity will increase

Current Vendor Focus on Capacity Planning

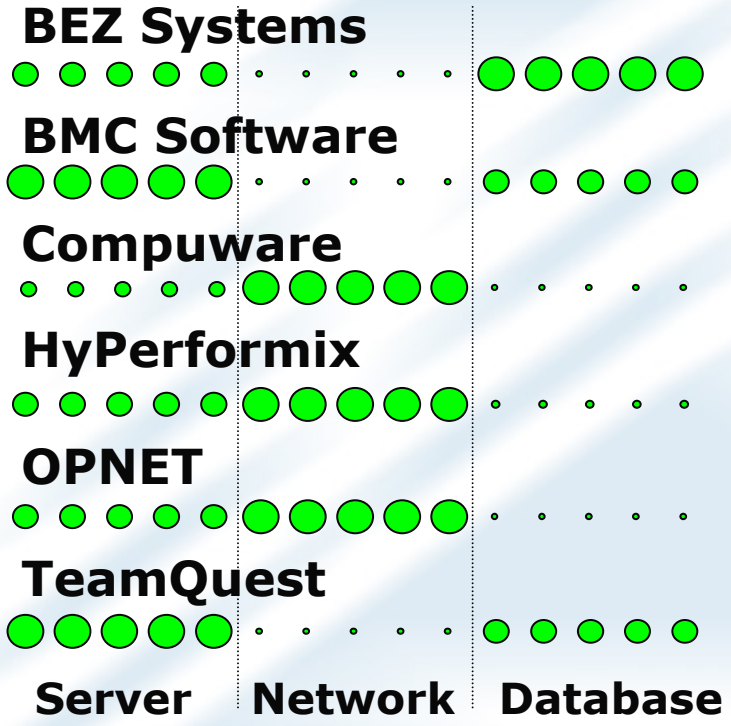

Pressure vendors for end-to-end tools and expect more presence from familiar vendors

Capacity Planning: Engineering Infrastructure Performance

- Optimize infrastructure performance
  - Follow proven structured engineering processes
  - Understand changing business requirements
- Tackle the complexity of distributed systems
  - Leverage decades of mainframe experience
  - Consider all hardware and software components for a complete end-to-end perspective
- Learn about technology details
  - Develop expertise on the advanced methods used by capacity planning products
  - Employ some level of modeling to combine actual infrastructure and application behaviors with realistic business requirements