

Profiling and Modeling Resource Usage of Virtualized Applications

Timothy Wood (University of Massachusetts, Amherst)

Joint work with Lucy Cherkasova, Kivanc Ozonat, and Prashant Shenoy




Virtualized Data Centers

- Benefits
 - Lower hardware and energy costs
 - Capacity on demand, agile and dynamic IT
- Challenges
 - Apps have different resource requirements
 - Virtualization overheads
 - Effects of consolidating multiple VMs to one host
- Important for capacity planning and consolidation

Understanding resource needs is
difficult but important

Predicting Resource Requirements



- Most overhead caused by I/O
 - Network and Disk activity
- Xen I/O Model
 - 2 components
 - Dom0 handles I/O
- Must predict CPU needs of:
 1. Virtual machine running the application
 2. Domain 0 performing I/O on behalf of the app



Requires several prediction models based on multiple resources

Scaling Traces


Native Application Trace



Virtualized Application Trace

Why Bother?

- More accurate cost/benefit analysis
 - Capacity planning and VM placement
- Impossible to pre-test some critical services
- Hypervisor comparisons
 - Different platforms or versions



Our Approach

- Automated robust model generation
- Run benchmark set on native and virtual platforms
 - Performs a range of I/O and CPU intensive tasks
 - Gather resource traces
- Build model of Native --> Virtual relationship
 - Use linear regression techniques
 - Model is specific to platform, but not applications

Can apply this general model to any application's traces to predict its requirements

Benchmark Suite

- Focus on CPU and I/O intensive client-server apps
- Benchmark activities:
 - Downloading files
 - Uploading files
 - Computation
 - Reading Files
 - Writing Files
- Need to break correlations between resources
 - High correlation between packets/sec and CPU time
- Cover different types of I/O
 - Support different virtualization and hardware platforms

Building Robust Models

- Outliers can considerably impact regression models
 - Creates model that minimizes absolute error
 - Must use robust regression techniques to eliminate outliers
- Not all metrics are equally significant
 - Starts with 11 metrics: 3 CPU, 4 Network, and 4 Disk
 - Use stepwise regression to find most significant metrics

Want to build a *general* model which
can be used on any application

Sample Model

- Native system metrics:
 - CPU: Kernel, User, IO wait
 - Net: RX pkts/sec, TX pkts/sec, RX B/sec, TX B/sec
 - Disk: Reads/sec, Writes/sec, Read B/sec, Write B/sec
- Virtual System CPU
 - Domain 0 and Virtual Machine
- Model:

Set of equations to solve:



$$\begin{cases} c_0 + c_1 * M_1^1 + c_2 * M_2^1 + \dots + c_{11} * M_{11}^1 = U_{vm}^1 \\ c_0 + c_1 * M_1^2 + c_2 * M_2^2 + \dots + c_{11} * M_{11}^2 = U_{vm}^2 \\ \dots \dots \dots \end{cases}$$

Model VM:

$$c_0^{vm}, c_1^{vm}, \dots, c_n^{vm}$$

Model Performance


- Train the model using simple benchmarks
- Apply to RUBiS web application



90% of Dom0 predictions within 4% error
90% of VM predictions within 11% error

CPU Scaling



- Is it necessary to look at resources other than just total CPU?



Definitely need multiple resources!

Second Hardware Platform

- AMD64 dual CPU 2.6Ghz



Produces different model parameters
Predictions are just as accurate

Further Refinements

- Use model to evaluate of virtual device capacity and detect potential bottlenecks
- Define composition rules for estimating resource requirements of colocated virtualized applications
- Analyze impact of hardware virtualization on resource consumption



Questions?

... or comments?

twood@cs.umass.edu