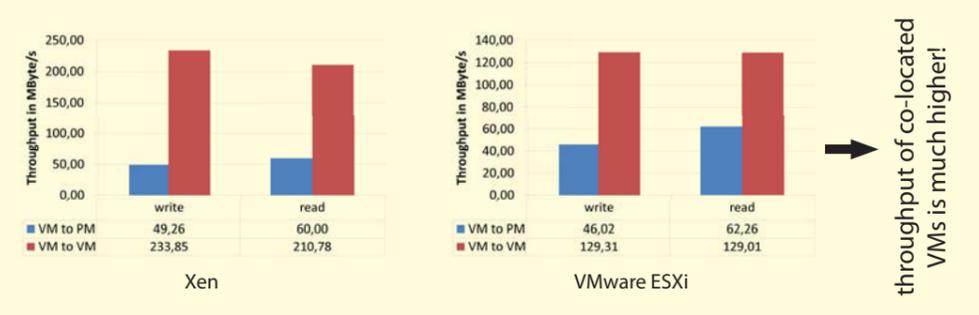
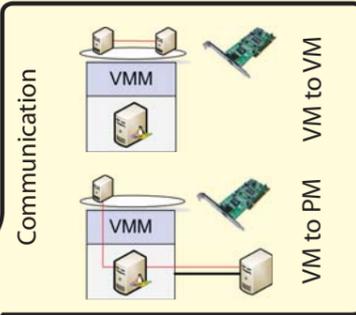
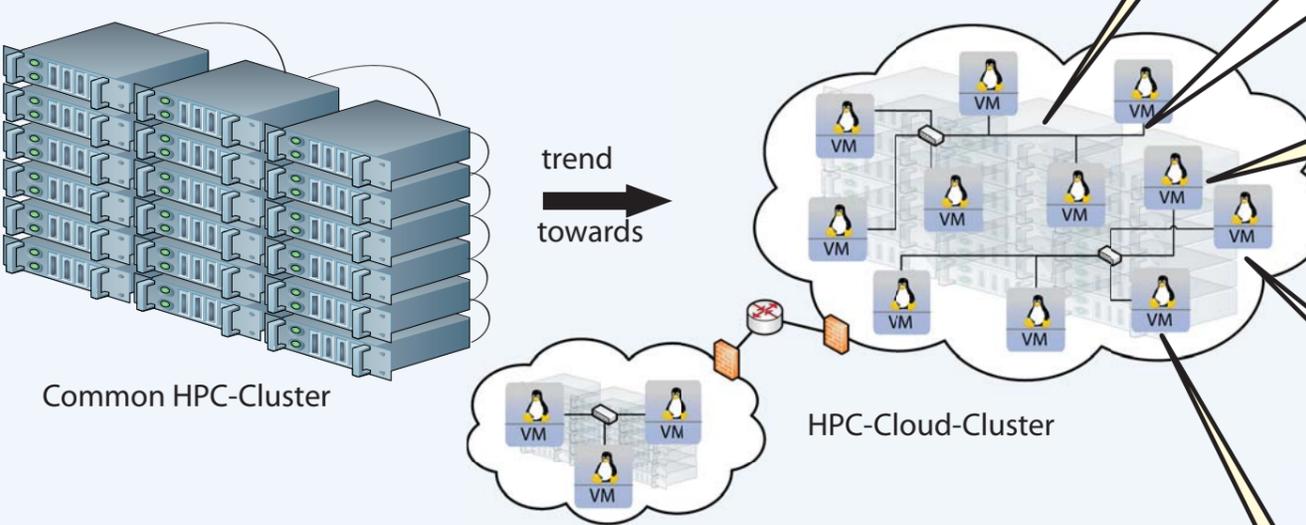


High Performance Cloud Computing - Why not?!

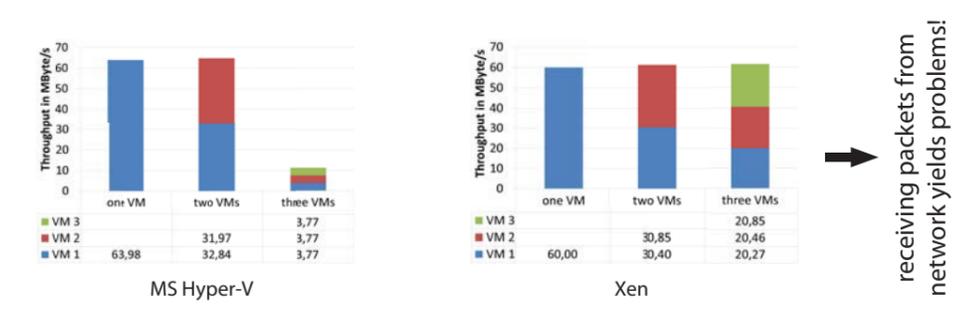
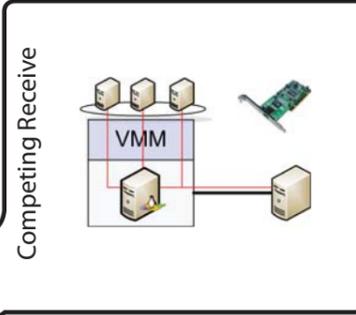
Dieter Kranzlmüller, Tobias Lindinger, Nils Gentschen Felde
Ludwig-Maximilians-Universität München
{kranzlm, lindinger, felde}@mnm-team.org



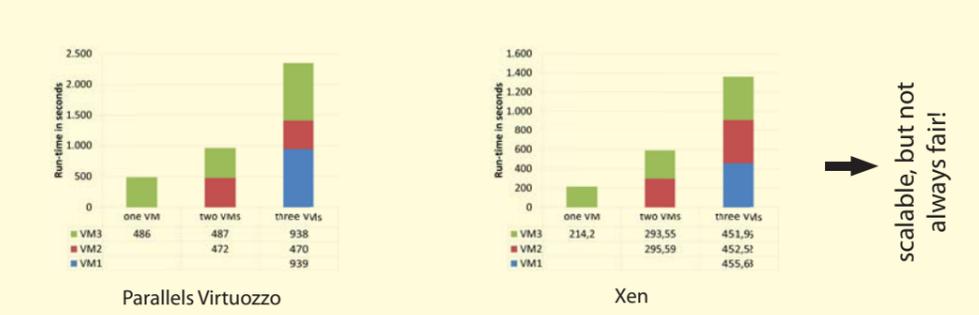
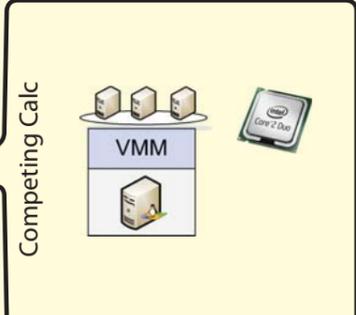
HPC-Cloud-Cluster		
	PROs	CONS
Fault	<ul style="list-style-type: none"> Failures of physical machines can be compensated without resetting the whole cluster to a checkpoint 	<ul style="list-style-type: none"> Faults and their causes may be "cloudy"
Configuration	<ul style="list-style-type: none"> Easy, dynamic and fine grained partitioning of a cluster Management and configuration of the system can be simplified and automatized 	<ul style="list-style-type: none"> Very dynamic changes of configuration expected CMDBs for virtual infrastructures do not yet exist
Accounting	<ul style="list-style-type: none"> New accounting methods possible 	<ul style="list-style-type: none"> Accounting measures are still subject of research
Performance	<ul style="list-style-type: none"> Loadbalancing is possible due to Live Migration <ul style="list-style-type: none"> intra computing center inter computing center 	<ul style="list-style-type: none"> Virtualization overhead decreases performance VMs may influence each other Changing the network topology by the use of VM migrations will decrease MPI performance
Security	<ul style="list-style-type: none"> Better security because sandboxing of processes is possible 	<ul style="list-style-type: none"> Virtualization layer may yield access to entire cluster



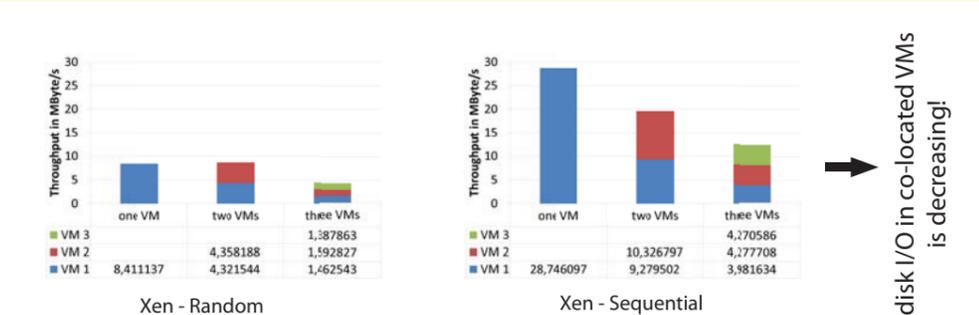
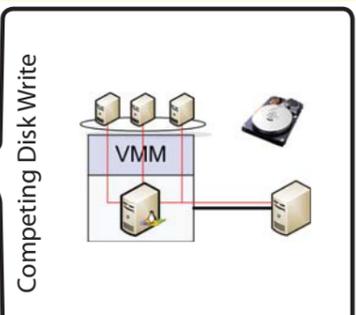
throughput of co-located VMs is much higher!



receiving packets from co-located VMs yields problems!



scalable, but not always fair!

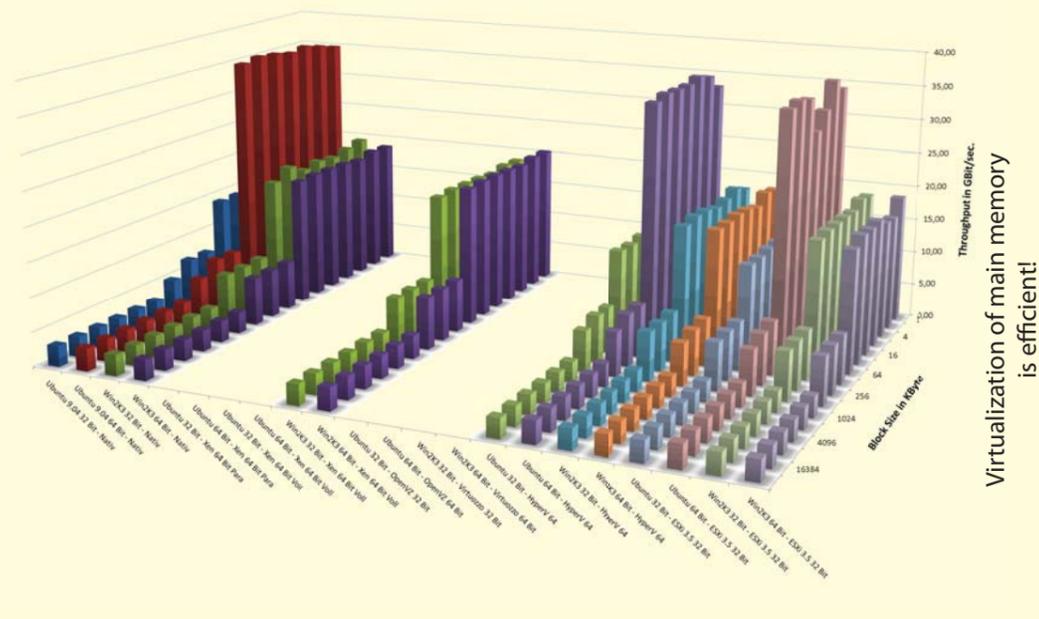
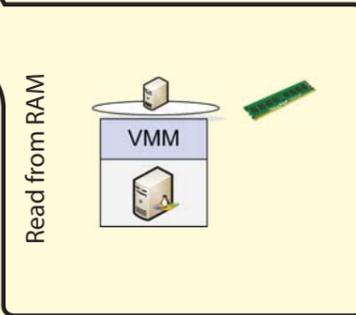


disk I/O in co-located VMs is decreasing!

Recommendations	
static	dynamic
<ul style="list-style-type: none"> I/O of VMs is highly CPU intense 	<ul style="list-style-type: none"> Group communicating VMs on a physical server instance
<ul style="list-style-type: none"> Disk-I/O of VMs tends to be inefficient, sequential access to a virtual disk may be mapped to non-sequential access on physical disks 	<ul style="list-style-type: none"> Avoid migration of VMs due to minor loadbalancing issues <ul style="list-style-type: none"> during migration response time of migrated VM may increase network optimized MPI implementations will get in trouble Separate VMs with high I/O (network, disk, ...)
<ul style="list-style-type: none"> Check your Hypervisor / VMM for fair scheduling implementations Don't overcommit CPU resources - attach VMs to CPUs/cores exclusively Binding VMs to physical CPUs/cores will preserve caches Allocate CPU quota to Hypervisor / VMM 	

Future Work / Trends
<ul style="list-style-type: none"> More CPU-cores in parallel will cause high failure rates due to huge amount of components. New techniques for resilience and fault tolerance will be necessary! MPI must be extended in order to support virtual networks Intelligent distributions of virtual machines must be calculated and realized automatically

Conclusions / Observations
<ul style="list-style-type: none"> Trend towards virtualization also in HPC -> exascale-roadmap But: No HPC-Cloud-Cluster in (upper) TOP 500 Virtualization is efficient and effective But: Coevally running VMs on same hardware reveal side-effects Virtualization may offer great additional benefits: <ul style="list-style-type: none"> management, especially fault mgmt, resilience, disaster recovery performance by "intelligent" VM distributions A lot more research needed!



Virtualization of main memory is efficient!