Diagnostic Capabilities of the Red Storm Compliance Test Suite

Mike Davis
Cray Inc.
http://www.cray.com
CUG Spring 2007

Overview

- Red Storm program initiated mid-2002
- Cray XT3 product introduced late 2004
  - http://www.cray.com/products/xt3/index.html
- Red Storm qualities
  - Size: 27x20x24 dual-core nodes
  - Dual Service Partitions (red, black)
  - Reconfigurable Compute Partitions

Red Storm Statement of Work (SOW)

- 96 Requirements
- 7 major categories
  - Architecture
  - Aggregate System performance
  - Compute node, backplane performance
  - Service node performance
  - RAS
  - Software
  - Secure Computing
- 20+ Software tests
  - Red Storm Compliance Test Suite (CTS)

Red Storm CTS Terminology

- ***Key metric:*** What the test measures, reports
- ***Component-level metric:*** The performance of individual components (e.g., compute nodes)
- ***Performance target:*** The value that the key metric is to meet or exceed
- ***Nominal reference value:*** The “better” of the component-level metric and the performance target (scaled to a component level)
- ***Deviation tolerance:*** A decimal fraction of the nominal reference value

Red Storm CTS Terminology

- ***Key assessment****:* The comparison of the key metric with the performance target
- ***Deviation assessment****:* The comparison of the deviations from nominal reference value with the deviation tolerance
- ***Noncompliance****:* An unfavorable result of either key assessment or deviation assessment
- Scaling prefixes (mega, giga, etc.) are all power of ten
- Compliance targets are not necessarily the same as those specified in the SOW

CTS Test Categories

- Scaled single-component test (SC)
- Scaled component group test (CG)
- Single metric test (SM)

Scaled Single-Component Test

- Can be run on a single component
- Has been designed/adapted to run at (any) scale
- Each component does equal work
- Key metric: performance of slowest component
- No communication between components

Scaled Component-Group Test

- Can be run on a small group of related components
  - Topological: e.g., nodes sharing a common link
  - Conformal: e.g., nodes serving a common FS
- Scaling is constrained so as to maintain relationship across groups
- Each group does equal work
- Key metric: performance of slowest group
- Communication within groups only

Scaled Component-Group Test

- Additional metric: aggregate performance
  - Based on time between first-in and last-out
  - Can constrain the scaling (“LOFI scaling”)
    - Synchronization across groups around timed portion of code
    - Notion of “global time” or “time-keeper”
    - Summary-reduction of group results
    - Selection of “group leader” to gather/report results

Single Metric Test

- Runs on all available components
- Produces a single result metric
  - Performance (single aggregate number)
  - Functionality (output compares with baseline)
- Measurement of individual component performance either not possible or not interesting

Test	Description	Type	Units	Target	Dev. Tol.
104	CPU ID, frequency	SC	GHz	2.4	0.0001
202	HPL	SM	TF	0.0036M	N/A
205	Bisection Bandwidth	CG	TB/s	0.0062M	0.05
206	Link Bandwidth	CG	GB/s	3.8M	0.03
208	Aggregate I/O Bandwidth	CG	GB/s	0.157M	0.1
209	Aggregate NW Bandwidth	CG	GB/s	0.25M	0.1
307	Memory Bandwidth	SC	GB/s	4.0	0.005
607	Single file size	SM	TB	50	N/A
615	Load/launch	SM	s	60	N/A

Test	Description	Type	Units	Target	Dev. Tol.
105	Memory size	SC	GB	1.9	0.005
204	MPI latency	CG	us	11.5	0.01
211	Bisection Bandwidth, compute/service	CG	GB/s	2.5M	0.2
302	IEEE-754 compliance	SM	N/A	N/A	N/A
303	Performance Counters	SM	Events +/-	0	N/A
305	Memory latency	SC	ns	80	0.005
405	Aggregate I/O BW svc	CG	GB/s	0.625M	0.2
605	MPI-2 functionality	SM	N/A	N/A	N/A
617	TotalView capability	SM	N/A	N/A	N/A

AMD Opteron™ Processor

- Scaled single-component test
  - Component = processor
- Key metrics
  - Processor signature (model, family, stepping)
  - Processor speed (gigahertz)
- Target values
  - 33/15/2 for signature
  - 2.4 for speed
- Deviation tolerance
  - 0 for signature
  - 0.0001 for speed (100 clocks per million)

Memory Bandwidth

- Scaled single-component test
  - Component = processor
- Key metric
  - Bandwidth between processor and memory (gigabytes/second)
  - Using STREAM triad kernel
    - http://www.cs.virginia.edu/stream
- Target = 4.0, 4.2 (depending on location)
- Deviation Tolerance = 0.005

Link Bandwidth

- Scaled component-group test
  - Component group = a pair of compute nodes
  - Relationship = sharing a network link
- Key metric
  - The bidirectional bandwidth when exchanging MPI messages of 1 megabyte or less (gigabytes/second)
- Target = 3.8
- Deviation tolerance = 0.04

Link Bandwidth

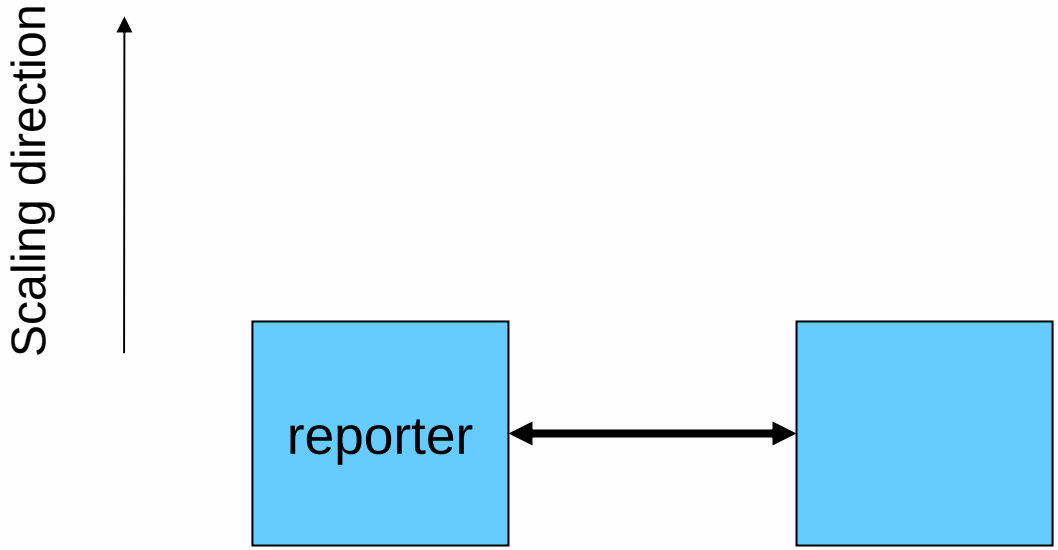

Bisection Bandwidth

- Scaled component-group test
  - Component group = an even number of compute nodes
  - Relationship = topologically contiguous and collinear
- Key metric
  - Bidirectional bandwidth across the bisection link (aggregated over M component groups) when exchanging messages of 1 megabyte or less between paired nodes (terabytes/second)
- Target = 0.0062M
- Deviation tolerance = 0.05

Bisection Bandwidth

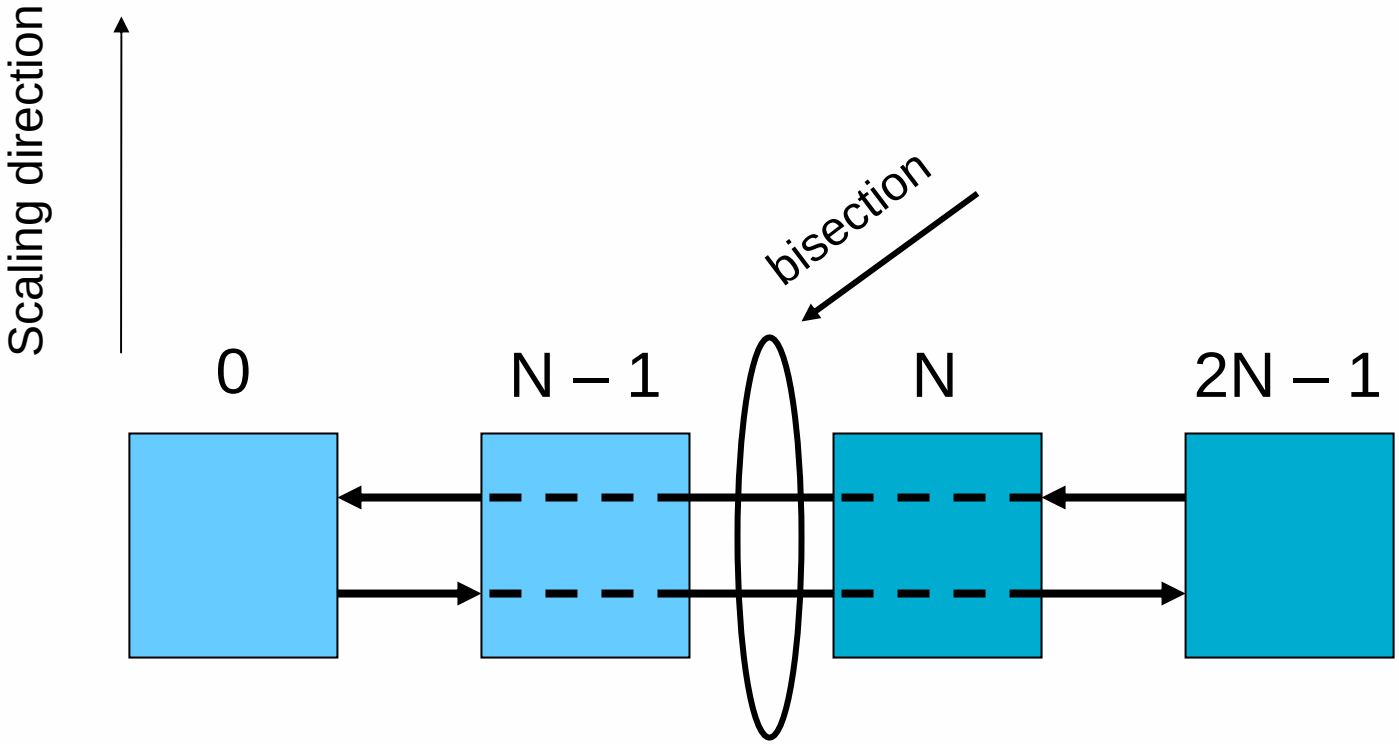

I/O Bandwidth

- Scaled component-group test
  - Component group = a small number of compute nodes and 1 Lustre OST
  - Relationship = topologically “close” and “distinct”
- Key metric
  - I/O bandwidth achieved on the OST (aggregated over M component groups) for read and write operations from a real-world application (gigabytes/second)
- Target = 0.157M
- Deviation tolerance = 0.1

I/O Bandwidth

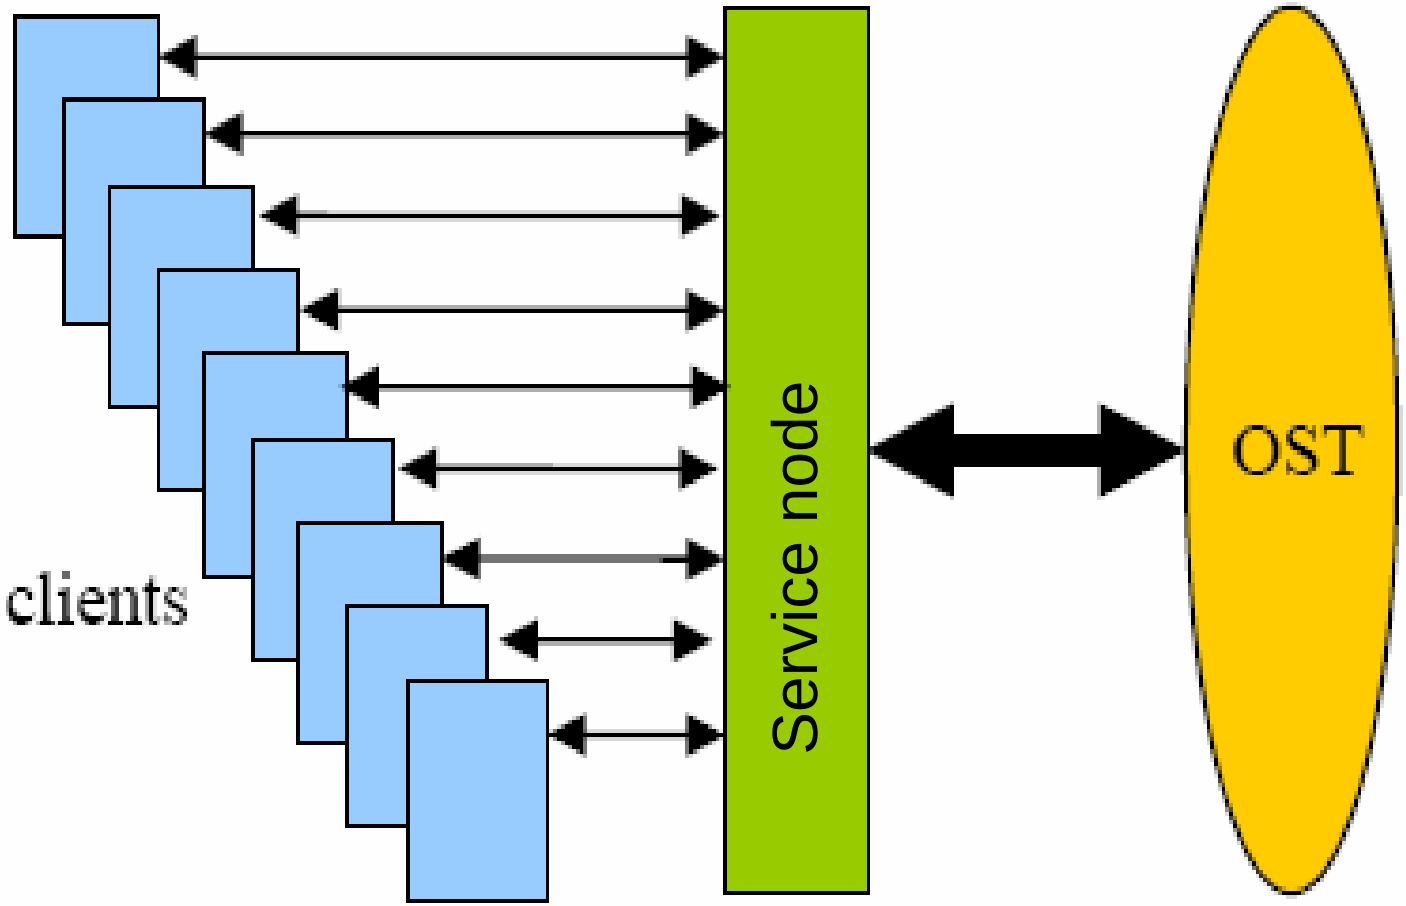

Single File Size and Accessibility

- Scaled component-group test
  - Component group = a small number of compute nodes (clients) and 1 OST
  - Relationship = topologically “close” and “distinct”
- Key metrics
  - The size of a single file generated by M component groups (terabytes)
  - The number of miscompares from the write/read/compare sequence
- Target values
  - 50 for size
  - 0 for miscompares

Aggregate Network Bandwidth

- Scaled component-group test
  - Component group = a service node with attached 10GigE riser (client), a remote dedicated server, and N OSTs
- Key metric
  - I/O bandwidth through the client (aggregated over M component groups) when moving data from files striped across the OSTs to the remote server using iperf (gigabytes/second)
  - http://dast.planr.net/Projects/Iperf
- Target = 0.25M
- Deviation tolerance = 0.1

Aggregate Network Bandwidth

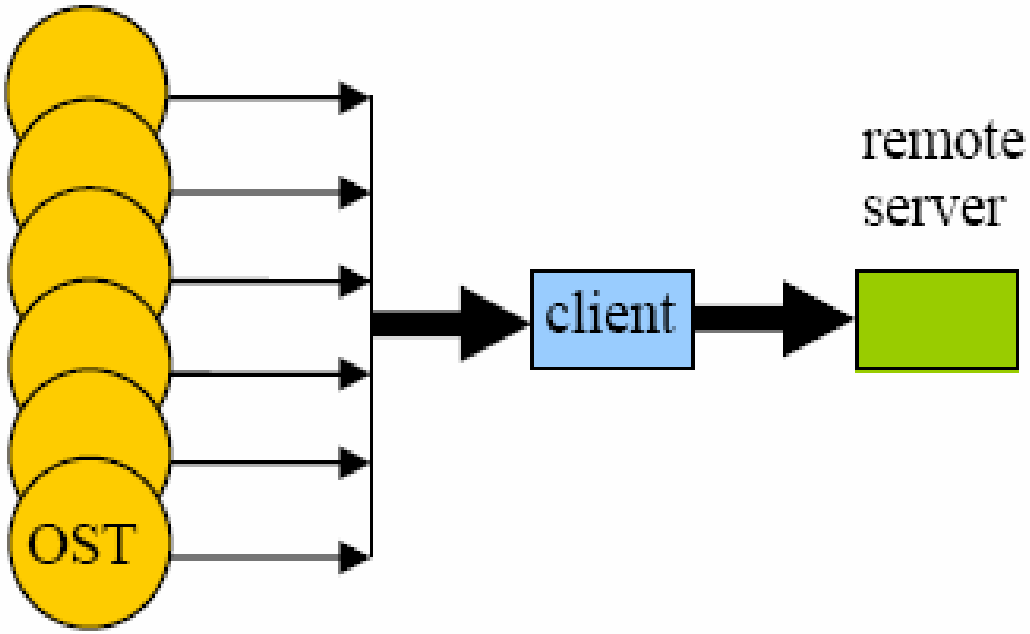

High-Performance LINPACK

- Full system test
  - http://www.netlib.org/benchmark/hpl
  - Interconnect network
  - Environmental monitoring/control
- Software test
  - Compilers
  - ACML (http://developer.amd.com/acml.jsp)
- Scripted to allow:
  - Running a specified time/size
  - Running multiple concurrent copies / filling the mesh

High-Performance LINPACK

- Key metric
  - Performance of the matrix solver (teraflops/second)
- Target
  - 0.0036M, M = number of processor cores

Job Load/Launch Time

- Full system test
- Key metric
  - Time to load and launch a heterogeneous real-world application onto the full system (seconds)
    - Load and launch = time from yod to MPI_Init
    - Heterogeneous = at least three distinct executables, each at least 1 megabyte in size
    - Full system = all available compute nodes plus all available service nodes that are configured to run applications
- Target = 60

CTS In Action

- Initial Operations (Jan – May 2005)
- Memory Upgrade (May – Jul 2005)
- Cray SeaStar™ Voltage Tuning (Aug – Sep 2005)
- 5th Row Upgrade (Jun – Sep 2006)
- UNICOS/lc™ 1.5 Upgrade (Apr 2007)
- Ongoing testing

Initial Operations (Jan – May 2005)

- Identified by Compute node tests
  - Opteron processors with incorrect frequency, incorrect stepping
  - Memory components with incorrect size, high memory error rates
- Identified by HPL test
  - Locations of faulty Seastar processors
- Identified by I/O Bandwidth test
  - Inconsistently configured Lustre nodes
- Identified by Network Bandwidth test
  - Inconsistently configured 10GigE nodes

Memory Upgrade (May – Jul 2005)

- Identified by Memory bandwidth test
  - Effects of differences in speed between Micron™ and Samsung™ parts

Cray SeaStar Voltage Tuning (Aug – Sep 2005)

- Identified by HPL, Bisection bandwidth, and Link bandwidth tests
  - Behavior of links at various voltages
- Identified by HPL test
  - Metrics for maximum cabinet power draw and heat output

5th Row Upgrade (Jun – Sep 2006)

- Added a 5th row to the system
- Upgraded AMD Opteron processors
- Upgraded Cray SeaStar processors
- Reconfigured Lustre file systems
- Upgraded OS to UNICOS/lc 1.4

5th Row Upgrade (Jun – Sep 2006)

- Identified by Memory bandwidth test
  - Effects of mixed-memory parts (and faster AMD Opteron processors) on memory bandwidth
    - Also affects link bandwidth
- Identified by IOR, confirmed by Link bandwidth test
  - Problems in algorithms that compute the aging of network packets

Ongoing Testing

- CTS is run after significant system changes:
  - Hardware upgrades
  - Software upgrades
  - Reconfigurations
  - Significant Maintenance Events

CTS-Generated SPRs

Compilers	17
Catamount	9
Tools	8
Lustre	7
MPICH2	6
Libc	4
Pubs	2
Linux	1

The Future of CTS

- Tests will be adapted as new features are introduced
- SMP Linux
  - I/O Bandwidth – service partition
  - Aggregate network bandwidth
- Accelerated Portals
  - MPI Latency test
- Lustre enhancements
  - Wide file (320 OSTs)
    - Single file size and accessibility test
  - Linux client overhead reduction
    - I/O Bandwidth – service partition
    - Aggregate network bandwidth

The Future of CTS

- Performance tools
  - Integer math operation counters
    - CPU performance counter accessibility test
- Heterogeneous applications
  - Job load/launch time test
  - TotalView capability test

Acknowledgements

- Cray Inc.
  - Bob Alverson
  - Gail Alverson
  - Sarah Anderson
  - Luiz DeRose
  - Dick Dimock
  - Dennis Dinge
  - Mark Pagel
  - Howard Pritchard
  - Kevin Thomas
  - Kevin Welton
- Sandia National Labs
  - Doug Doerfler
  - Sue Goudy
  - Sue Kelly
  - Kevin Pedretti
  - Jim Tomkins
  - John Vandyke
  - Courtenay Vaughan
  - Keith Underwood

Questions?