OpenCloud

Rhino Scalability and Performance on Blade hardware

Operators preference for blade hardware technology is driven by the benefits that modular technology brings: namely flexibility, scalability, reliability, performance and price.

OpenCloud's Rhino™ event-driven application server has been tested and benchmarked on two blade technologies :

- Alliance Systems ATCA Platforms and
- IBM BladeCentre-T servers

With single system boards, each running a separate operating system image, the solution provider has the challenge of utilising and managing the entire system as a single entity.

Rhino software complements blade hardware by distributing itself across the discrete blades with a highly efficient and fault tolerant architecture. This simplifies and greatly reduces the effort of the solution provider, allowing them to focus on the service logic. Cost is greatly reduced and software re-use greatly increased.

The Modular, Event Driven Application Server

As growth in all networks, including 3G and IMS is driven by other industries such as financial services, defence, retail, transportation etc., and by new technology enablers such as RFID, a next generation of high-performance platforms is required.

A modular, Event-Driven Application Server such as Rhino is well positioned to take advantage of modular blade technology. At the same time, Rhino meets stringent requirements for availability, scalability and open standard programmable interfaces.

Rhino Technology + Blade Technology = Modular, Flexible, Dynamic, Scalable, Fault Tolerant and Extensible solutions.

Rhino + Blade Feature	Benefit
N-way active-active	A highly efficient system that utilises all cards for fault tolerant and scalable event processing (NOT hot standby) – Ideal for IMS, IN, Messaging ...
Fault Tolerant on Standard Hardware	Continue processing events that were being processed on a failed processor card – Delivers not only application continuance (high availability) but also 99.999% session continuance across the system (fault tolerance)
Linear, Dynamic Scaling	Hot-plug processor boards as needed and have them dynamically join the cluster – Matches customer usage requirements
Transparent Single Image System Management	Dynamically add new services, alarms, provisioning etc., across all processing cards – Fault tolerant, consistent and coordinated management
Hot Swap	Operational services and system software – Modular and flexible for convergent services
Open Standard	Event driven application server environment transparent to the programmer. – Enables full advantage to be taken of Rhino and the modular Blade platform

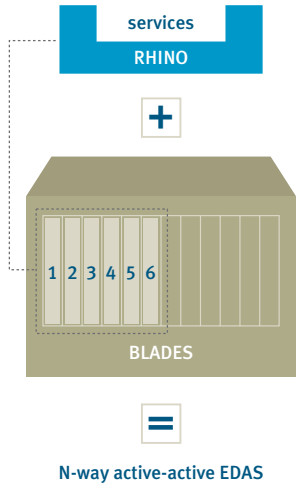

N-way active-active EDAS

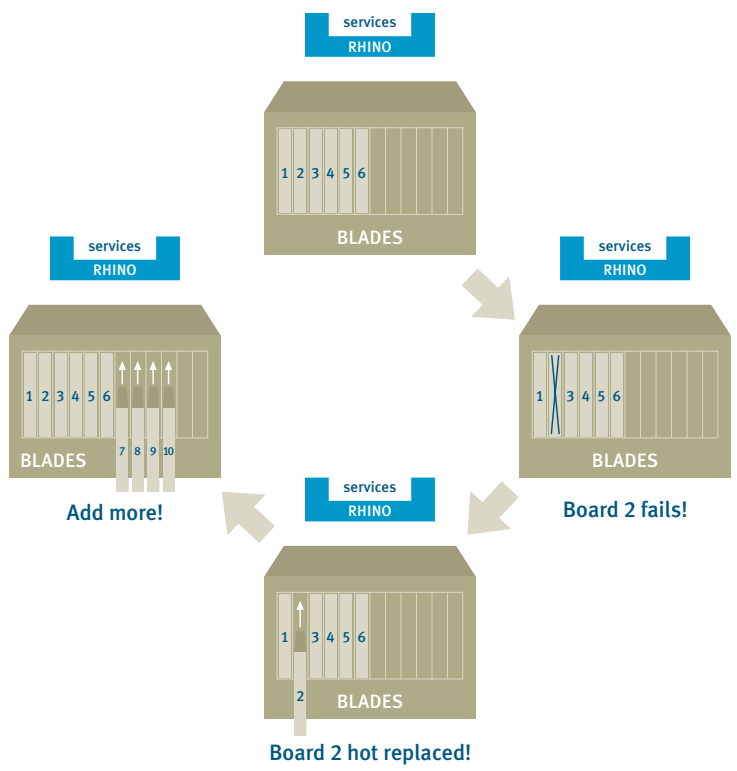

Dynamic Modular Combination of Rhino and Blade Technology

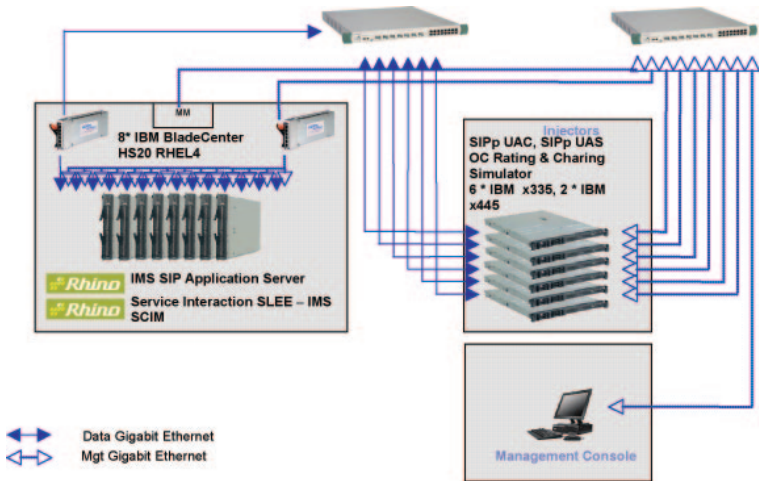

Benchmark deployment architecture

A BOARD FAILS... AND IS REPAIRED... AND ADD MORE

The following diagram clearly demonstrates the dynamic modular combination of Rhino and Blade Technology.

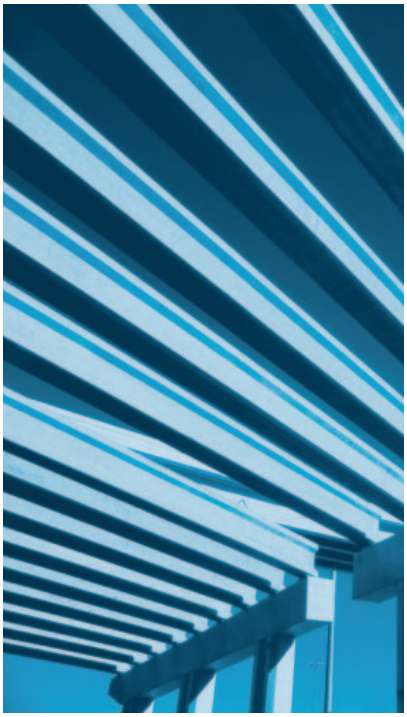

SCALABILITY

Scalability is a property of a system which indicates its ability to either handle growing amounts of work in a graceful manner, or to be readily enlarged. In a scalable system performance improves after adding hardware, proportionally to the capacity added:

- To **scale vertically** or **scale up** means to add resources to a single node in a system, such as adding memory or a faster hard drive to a computer.
- To **scale horizontally** or scale out means to add more nodes to a system, such as adding a new computer to a clustered software application.

A series of tests were carried out on IBM BladeCenter-T hardware to demonstrate the vertical and horizontal scalability profile of Rhino under a range of real-world processing scenarios.

1. Vertical Scalability in a Simple Service Scenario

This scenario is to demonstrate Rhino throughput and vertical scalability. The application demonstrates a single Find Me Follow Me (FMFM) Service implemented on the Rhino Service Interaction Server (SIS).

RESULTS

The results prove the linear vertical scalability of the Rhino platform, and show a throughput of 338 calls/second and 10% CPU utilization

Call Rate	Per Blade	CPU Used	Call Setup Latency (ms)			
(calls/s)	(calls/s)		50%	75%	90%	95%
640	80	20%	3	3	4	4
1200	150	35%	3	4	5	8
1800	225	52%	4	4	10	29
2400	300	72%	5	30	63	84

LinearScalability as call volume increase

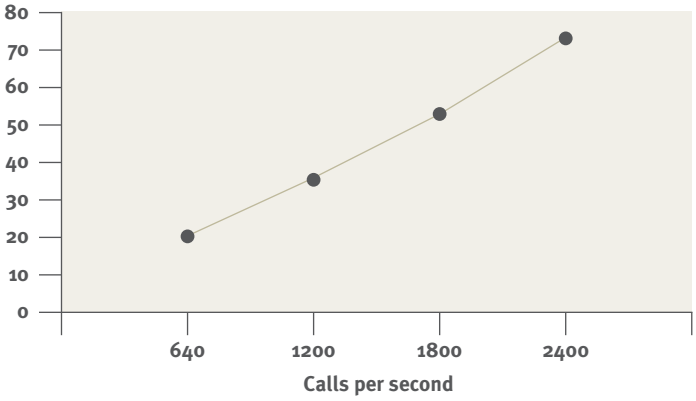

2. Horizontal Scalability in a Complex Service Scenario

This scenario is to demonstrate Rhino throughput in a complex multi-service scenario and also horizontal scalability. The application demonstrates a single Find Me Follow Me (FMFM) Service and a Prepaid charging service, both implemented on the Rhino Service Interaction Server (SIS).

RESULTS

No relative performance degradation is observed as Rhino call throughput increases by adding blades to the cluster (proving linear scalability)

Complex Prepaid scenario: 170 cps/blade

Simple FMFM scenario: 300 cps/blade

The first blade added and the last blade added each deliver the same increase in capacity Rhino Clusters are ideally suited to the BladeCenter-T

Number of Blades	Call Rate	Per Blade	CPU Used	Call Setup Latency (ms)			
	(calls/s)	(call/s)		50%	75%	90%	95%
1	170	170	72%	12	16	37	53
2	340	170	73%	12	18	39	55
4	680	170	75%	13	23	49	63
8	1360	170	72%	11	16	40	60

LinearScalability as call volume increase

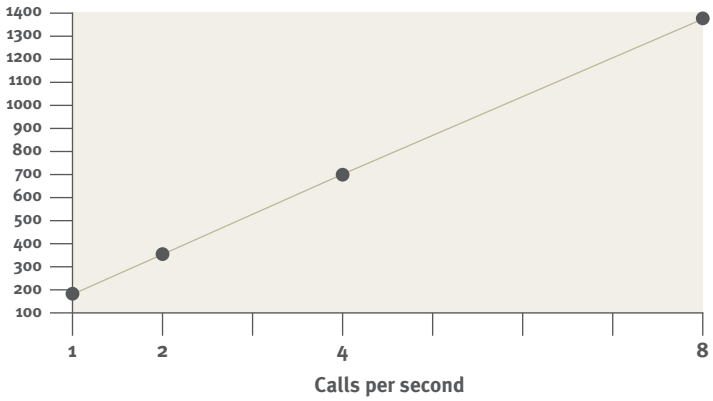

OpenCloud

About OpenCloud

OpenCloud was formed in New Zealand in 2000 to create open standard software technology that would revolutionise the portability and interoperability of services in telecommunications specifically in the evolution to IP and 3G IMS. OpenCloud works with partners to deliver, integrate and support end-to-end solutions incorporating OpenCloud products to network operators and service providers worldwide. OpenCloud has offices in: UK, New Zealand, Madrid, Tokyo, and San Francisco.

More Information

General Enquiries

Email us at info@opencloud.com
or visit us at www.opencloud.com