



Infinetics Backgrounder

Infinetics™ Technologies, Inc is a Boston-based networking software company that has developed a network architecture and implementation for building truly scalable data centers. The company was founded two years ago by a veteran team of entrepreneurs to commercialize theoretical concepts developed over the past 5 years.

Our philosophy is grounded in the general awareness in industry and academia that improvements are possible in data center performance. We believe that while many groups have developed better solutions, none of them come close to theoretical maximum efficiency, and while performance is a slippery word, encompassing bisection bandwidth, latency, fault tolerance, cost per port and other parameters, ultimately fundamental laws of physics and information theory dictate the true limits of network performance. From this foundation, we have developed novel and proprietary methods for constructing and operating networks that achieve the best possible performance for any given set of physical constraints.

Our software runs on network switches, in hypervisors, and in a central control and management system. Its three major components are a wiring manager, distributed packet forwarding logic, and a central system that manages security and quality of service (QoS) policies. The wiring manager uses fundamental discoveries that we have made in the mathematics of constructing physical networks to build networks that have the best possible tradeoffs between cost, bandwidth and latency, far exceeding the characteristics of even the best commercial and research implementations proposed to date. Our packet forwarding logic running on switches and hypervisors makes maximum use of the underlying physical network, while remaining highly scalable due to a careful allocation of effort between the central and distributed components. The tight coupling between the physical and routing layers creates a self-healing network that reallocates resources automatically based upon network topology changes. The hypervisor implementation gives data center operators an agile network to match the agile assignment of virtual computing tasks to physical resources that has become the standard method of operating highly scaled data centers, and the central control system allows the entire network to be managed as a single entity, allowing security policies and bandwidth allocation to take advantage of a global perspective on network usage patterns.

We have a number of patent applications filed that cover all aspects of the architecture, from the methods used to create optimal wiring, to the control and management systems that allow the base architecture to solve pressing data center business problems. Due to the highly general nature of our discoveries, the patents also cover using our discoveries in various hardware contexts, including ASIC switching hardware, switch backplane wiring, and solid state memory architecture.

In summary, our network architecture represents the “end game” for performance improvement, achieving theoretical maximum performance. It is a complete solution, having optimized wiring methods, a truly scalable control plane, strong hypervisor support with virtual machine support, which is also self-configuring. It has a versatile deployment model that can run on any switch, from basic commodity to advanced fabric switches and hypervisors. It offers a far lower cost per port and wiring costs than all existing solutions, and the hybrid centralized-localized control plane supports advanced security, traffic management and QoS policies.

An alpha product version has been completed and is on target for Q4 beta with 1st half 2012 commercial deployment, and we have engaged a number of technology partners include Arista, Fulcrum, KVM, Citrix, and Pronto. We are seeking a strategic relationship to leverage an industry leaders’ market presence to bring our broadly applicable “end game” solution to market.