

SOA-R

CLOUD COMPUTING FOR NATIONAL SECURITY APPLICATIONS

A DATALINE WHITE PAPER

INTRODUCTION

Through collaboration with leading industry partners, Dataline is integrating an end-to-end solution for secure cloud computing. Called Service Oriented Architecture – Real Time (SOA-R), this solution integrates:

- Real-time services oriented architecture;
- Multi-tiered, high-performance storage;
- Event driven automated workflows;
- Mission specific applications and web services;
- Highly reliable, virtualized, hardware infrastructure;
- Efficient wide area networking with accelerate geosynchronous satellite IP throughput; and
- Information access through secure Mobile Adhoc Networking.

SOA-R provides the benefits of cloud computing and secure ad-hoc mobility in a modular, standards-based implementation framework. By exploiting the advantages of a service oriented approach, web services and interfaces specific to the national security business domain are being developed and enhanced for deployment from and interface with a cloud computing infrastructure. By doing this, Dataline, in concert with its partners, can provide leading edge solutions without the risk typically associated with early adoption of leading edge technologies.

Like cloud computing, SOA-R has dual meanings. Whereas cloud computing refers to both a platform and type of application, SOA-R refers to a solution framework and domain specific solutions. To help understand these twin dualities consider the following descriptions:

- A cloud computing platform dynamically provisions, configures, reconfigures, and de-provisions servers as needed. In a SOA-R framework, the provisioning, configuring, reconfiguring and de-provisioning of servers is driven by national security domain specific rules, metrics and mission guidelines
- Cloud applications are applications that are extended to be accessible through the Internet from large datacenters and powerful servers that host Web applications and Web services. SOA-R applications are extended to be accessible through highly secure virtual private networks (VPN) tunnels and secure private networks (e.g. NIPRnet and SIPRnet). SOA-R applications are typically composite applications, served from powerful servers, and provided through Web Services from multiple organizations and datacenters.

The SOA-R framework uses event-driven workflows to provide:

- Intra-workflow component reuse;
- Faster workflow execution;
- Distributed/parallel workflow execution;
- Normalizes workflow design;
- Enhanced business process interoperability;
- Enhanced information sharing;
- Ability to implement more sophisticated IT architectures;
- Reduces ad-hoc workflows; and
- Easier business process redefinition.

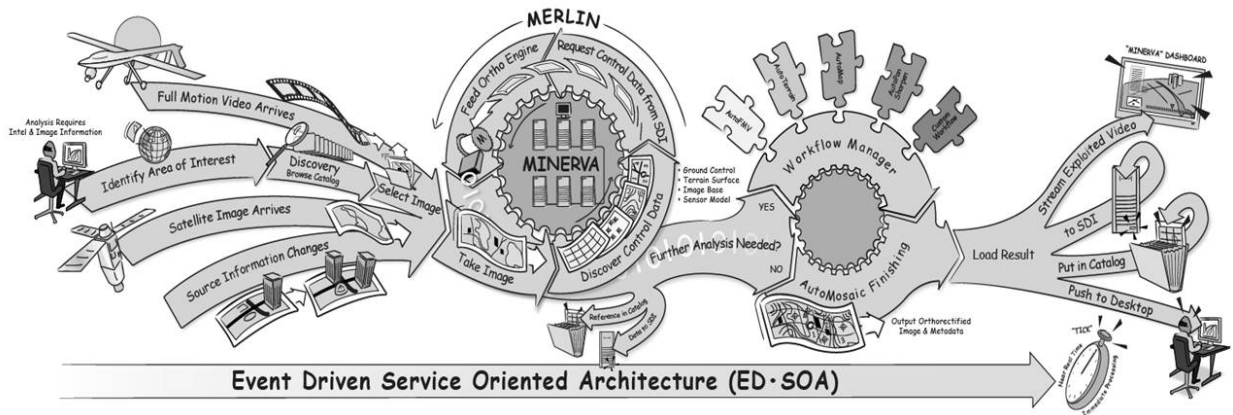


Figure 1 - SOA-R leverages the Event Driven Services Oriented Architecture first developed and deployed by Great-Circle Technologies¹

SOA-R Solution Framework

The deployment of a secure cloud computing platform to support national security missions requires an integrated solution across the following specific functional areas:

- Secure and highly scalable datacenters (Northrop Grumman, Dataline)
- Services Oriented Architecture and Mission Relevant Applications (GCT, IBM, Dataline)
- Hardware and software virtualization (IBM, Google, VMware)²
- Secure Global and regional networks
- Secure wireless ad-hoc networking and network clients (Harris, L-3, Cisco, Dataline)

¹ “Merlin and Minerva Story Map”, Great-Circle Technologies, 2007

² We are including middleware, input/output and network virtualization in this requirement as well.

Northrop Grumman’s Mission Systems sector is one of the premier integrators of complex, mission-critical systems. The sector develops and sustains systems that enable the Department of Defense, Department of Homeland Security and the intelligence community to execute missions at all levels: tactical, operational and strategic. Northrop Grumman is also the industry leader in secure messaging systems, providing accreditable systems that process the most sensitive data of the defense and intelligence communities. Northrop Grumman is the only company that can deliver systems that send secure messages from one source to many addresses located across multiple security enclaves. As an example of their capability, as the prime contractor for the Army’s Force XXI Battle Command Brigade and Below (FBCB2) program, Mission Systems is replacing paper maps and voice radio communications with digital capabilities. As a large-scale integrator, Northrop can provide its ability to design, deliver and operate mission critical end-to-end IT infrastructures to the SOA-R team.

As an IBM Premier Partner, Dataline serves as the server, storage and middleware technology integrator for SOA-R implementations. In this role, Dataline provides the server hardware, storage technology, virtualization technology, SOA middleware and infrastructure management middleware for the core cloud computing platform. In collaboration with Great-Circle Technologies, Dataline designs and implements the enterprise service bus, all required storage tiers and the cloud hardware and application provisioning strategy.

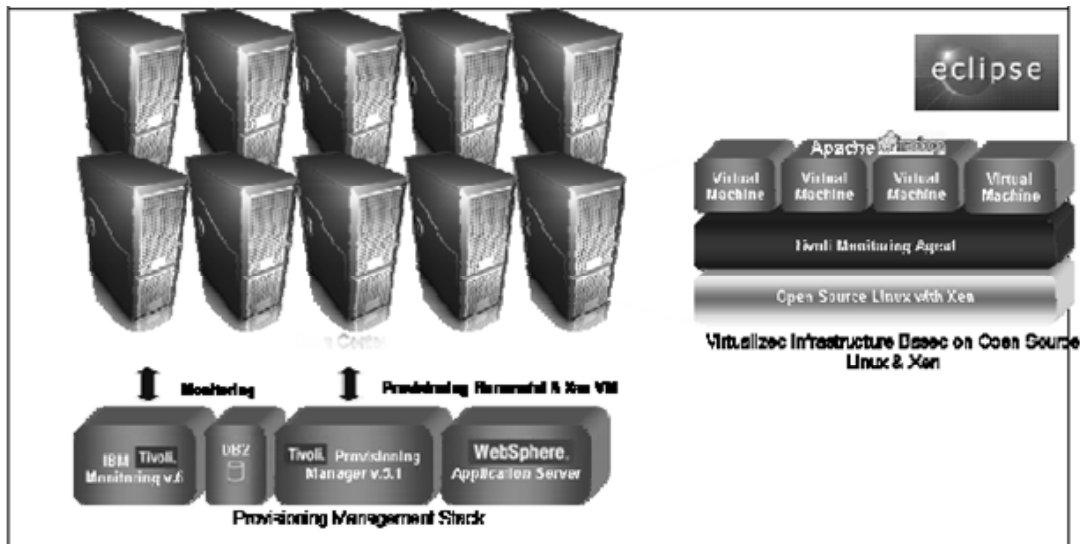


Figure 2 - Example high level cloud architecture

SOA-R also features an IBM WebSphere DataStage based Real Time Integration Architecture (RTIA). This solution component features an event-driven, message-based and trigger based infrastructure for the cost-effective integration of multiple applications. WebSphere DataStage can operate in real-time, capturing messages or extracting data at a moment's notice on the same platform that also integrates bulk data. This provides a key advantage over implementations that require the use of two separate products to achieve the same functionality. This architecture facilitates:

- Real-time performance.
- Scalability (number of nodes, communicating applications and fault tolerance).

- Construction of integration objects across applications, databases and systems without generating or having to manage code separately.
- Elimination of any requirement of altering or compromising the modularity of existing system by decoupling.
- Users and owners of applications / systems retaining control of their systems by not placing Real Time Integration (RTI) functionality in line with their application code.
- Easy maintenance on production front.
- Low cost rework at the end of the development cycle.³

Great-Circle Technologies serves as the subject matter lead on most SOA-R implementations. In this crucial position, GCT designs all domain area workflows and selects specific applications for use within the event-driven SOA architecture. GCT's core expertise is in providing innovative solutions that bridge the "best value" of standards-based commercial off-the-shelf technologies (COTS/SCOTS) with existing government-owned technologies (GOTS). Key to their success has been an ability to tailor pragmatic mission-oriented solutions that utilize advanced geospatial intelligence technologies.

In cloud implementations, Hadoop, an open source implementation of the Google File System, is the SOA-R preferred framework for running large data processing applications. This approach allows the creation and execution of applications using Google's MapReduce programming paradigm, which divides the application into small fragments of work that can be executed on any node in the cloud computing cluster. It also transparently supports reliability and data migration through the use of a distributed file system. Using Hadoop, the cloud can execute parallel applications on a massive data set in a reasonable amount of time, enabling computationally-intensive services such as retrieving information efficiently, customizing user sessions based on past history, or generating results based on Monte Carlo (probabilistic) algorithms. SOA-R Team implementations are also well positioned to leverage any innovations flowing from the IBM-Google cloud computing collaboration.⁴

Since the SOA-R framework is designed to support national security missions, SOA-R solutions typically leverage government provided global and regional communications networks. To do this efficiently, SOA-R uses two key Dataline technologies: SkipWare® and SCPS-M™. These technologies accelerate IP throughput over a geosynchronous satellite by a minimum of 17x the rate of standard TCP. Improvement ratios increase significantly when run over errored and asymmetric links. These protocols also achieve bandwidth efficiencies in excess of 90%, before payload compression. Because SkipWare® and SCPS-M™ are designed to operate in a gateway model, it is not necessary to deploy them on every remote client machine. A single gateway serves the entire LAN, allowing for simple, unobtrusive integration into the network. Smart congestion controls and transparent bridging features make these protocols true plug-and-play technologies. SkipWare and SCPS-M are also designed to interoperate with other common

³ "Real Time Integration / SOA using IBM WebSpher Datastage", Joshy George, IBM, <http://hosteddocs.ittoolbox.com/JG071807.pdf>,

⁴ "Cloud Computing", IBM High Performance On Demand Solutions (HiPODS) Greg Boss, Padma Malladi, Dennis Quan, Linda Legregni, Harold Hall, 8 October 2007

networking protocols, including TCP, DVB and SCPS-TP. Because SkipWare and SCPS-M are based on IETF-standard TCP, all community experience and diagnostic tools associated with TCP can be also applied to SkipWare.

Dataline has years of experience leveraging these global networks through its Data Communication Device. With the addition of NSA approved Type 1 encryption, classified IP connectivity can be provided through Ethernet, serial, wireless, cellular or satellite uplink within an IPSEC VPN tunnel. This Everything on IP (EoIP) infrastructure provides direct email, IP-based voice or video, mapping, network drive access and “behind-the-firewall” Internet access as necessary.⁵ Access to web services also provides mash-up capability as well. Information can also be pushed via subscriptions or event-driven SOA workflow processes.

SOA-R solutions can also leverage the CISCO mobile ad hoc networking solution as a component. These networks are emerging as a means for delivering the benefits of IP networking to users operating beyond the reach of a fixed network. In ad hoc networks mobile nodes associate on an extemporaneous or ad hoc basis. Ad hoc networks have a number of distinguishing characteristics when compared to conventional networking solutions:

- *Self Forming*: nodes that come within radio range of each other can establish a network association without any pre-configuration or manual intervention
- *Self Healing*: nodes can join or leave on a rapid basis without impacting operation of the remaining nodes
- *No Infrastructure*: in an ad hoc network, mobile nodes form their own network, and essentially become their own infrastructure
- *Peer to Peer*: traditional networks typically support end systems operating in client-server mode. In an ad hoc network, mobile nodes can communicate and exchange information without prior arrangement and without reliance on centralized resources
- *Predominantly Wireless*: historically networks have been mostly wired, and enhanced or extended through wireless access. The ad hoc environment is essentially wireless, but can be extended to support wired resources
- *Highly Dynamic*: mobile nodes are in continuous motion and ad hoc networking topologies are constantly changing.

Collectively, these characteristics will enable ad hoc networks to deliver timely information to a new and under-served class of users. Ad hoc networking solutions can be applied to virtually any scenario that involves a cadre of highly mobile users or platforms (which may include stationary devices as well), a strong need to share IP-based information, and an environment in which fixed infrastructure is impractical, impaired, or impossible. SOA-R ad hoc networking uses the Cisco RFC 4938 extensions for Ethernet-based communications between a router and a device such as a mobile radio that operates in a variable bandwidth environment and has limited buffering capabilities. These extensions provide a point-to-point over Ethernet (PPPoE) session-based mechanism for sharing radio network status such as link quality metrics and establishing flow control between a router and an RFC 4938-capable radio. An RFC 4938 radio initiates a Layer 2

⁵ “Data Communications Device - VIP Communications”, Dataline, 2008

PPPoE session with its adjacent RFC 4938 router on behalf of every router/radio neighbor discovered in the network. These Layer 2 sessions are the means by which radio network status for each neighbor link gets reported to the router. The radio establishes the correspondence between each PPPoE session and each link to a neighbor. Cisco's RFC 4938 solution enables network-based applications and information to be quickly and reliably transported over directional radio links. Fast convergence and optimal route selection ensure satisfactory delivery of mission-critical, delay-sensitive traffic.⁶




Figure 3 - AN RFC 4938-based Ad Hoc Network

⁶ "Radio Aware Routing in Mobile Ad Hoc Networks", Cisco, 2008

National Security Mission Relevance

According to Google CEO Eric Schmidt "Cloud computing is the story of our lifetime....Eventually all devices will be on the network." He also, however, stated that with the exception of security requirements, "there's not that much difference between the enterprise cloud and the consumer cloud"⁷. SOA-R addresses that important difference.

Cloud computing advantages include:

- Location of infrastructure in areas with lower costs of real estate and electricity
- Sharing of peak-load capacity among a large pool of users, improving overall utilization
- Separation of infrastructure maintenance duties from domain-specific application development
- Separation of application code from physical resources
- Ability to use external assets to handle peak loads (no requirement to engineer for highest possible load levels)
- No requirement to purchase assets for one-time or infrequent intensive computing tasks⁸

In today's operational, fiscal, and environmentally friendly environment, all national security agencies are being driven to look at the advantages offered by cloud computing. SOA-R answers that need. When the SOA-R framework is engineered into a domain-specific solution, all technical, operational, procedural and environmentally friendly agency information technology requirements can be addressed in a service oriented, event-driven, globally networked, secure infrastructure. Defense, homeland security, justice and intelligence agencies all can benefit by considering SOA-R as the core of their transition to cloud computing.

⁷ "Google, IBM Join Forces To Dominate 'Cloud Computing'", Information Week, May 1, 2008, <http://www.informationweek.com/news/services/data/showArticle.jhtml?articleID=207404265>

⁸ "Cloud Computing", Wikipedia, http://en.wikipedia.org/wiki/Cloud_computing, June 14, 2008

SOA-R Team

Dataline, LLC is a leading transformational technology solutions provider to Department of Defense (DoD) and federal customers. Since 1990, our focus has been to integrate best-in-class technologies to provide robust, cost-effective Information Technology (IT) solutions for organizations of all sizes and missions of all criticalities. In partnership with our customers, we are accelerating the rate of change in the DoD's transformation to Net-Centric Warfare and progressing America's capability for total Information Dominance.

Great-Circle Technologies, Inc. (GCT) specializes in the design, development, and deployment of multi-INT enabled geospatial intelligence solutions. Established in 2000, GCT is headquartered in Chantilly, VA with offices in St. Louis, MO and supporting customers throughout the Washington DC metro area.

Northrop Grumman's Mission Systems sector is one of the premier integrators of complex, mission-critical systems. The sector develops and sustains systems that enable the Department of Defense and the intelligence community to execute missions at all levels of command: tactical, operational and strategic. Internationally, our defense, security and air traffic control solutions offer technology and expertise that few companies can equal.

International Business Machines Corporation (IBM) develops and manufactures information technologies, including computer systems, software, networking systems, storage devices, and microelectronics worldwide. IBM's Software segment primarily offers middleware and operating systems software comprising information management software for database, content management, and information integration; Lotus software for collaboration, messaging, and social networking; rational software, a process automation tool; Tivoli software for infrastructure management, including security and storage management; WebSphere software for Web-enabled applications; and product lifecycle management software.

Google, Inc., a technology company, maintains index of Web sites and other online content for users, advertisers, Google network members, and other content providers. Its automated search technology helps users to obtain instant access to relevant information from its online index. The company's solutions provide targeted advertising and Internet search solutions, as well as intranet solutions via an enterprise search appliance. Google, Inc. was founded in 1998 and is headquartered in Mountain View, California.

Cisco enables people to make powerful connections-whether in business, education, philanthropy, or creativity. Cisco hardware, software, and service offerings are used to create the Internet solutions that make networks possible-providing easy access to information anywhere, at any time. Since the company's inception, Cisco engineers have been leaders in the development of Internet Protocol (IP)-based networking technologies. Today, with more than 65,225 employees worldwide, this tradition of innovation continues with industry-leading products and solutions in the company's core development areas of routing and switching and other advanced technologies.

Additional Information

For further information on cloud computing and SOA-R please register and attend the Dataline educational series entitled “Cloud Computing in a Net-Centric Environment”. The series will feature speakers from Subject Matter Experts from the SOA-R Team companies. Together we are leading the way to true interoperability.

The series is being held from 9:00am – 1:00 pm at:

The Tower Club, Tyson’s Corner
8000 Towers Crescent Drive
Vienna, VA, 22182.

Series topics and dates are as follows:

TOPIC	DATE
What is Cloud Computing?	Wed 7/16/08
Global Information Access	Wed 8/13/08
Increased Efficiency, Reduced Cost	Thu 9/11/08
Event Driven Information	Wed 10/8/08
Information Availability	Wed 10/29/08
Mission Relevance	Wed 11/12/08

Online registration is available at <http://www.dataline.com/soar.htm>

“Cloud Musings”, a blog by Kevin L. Jackson containing personal comments and insight on cloud computing and its relationship to net-centric warfare, is also available at <http://kevinljackson.blogspot.com/>