Cyber Resiliency

Technology, Architecture, and Practical Validation

This paper proposes an approach to creating pervasive cyber resiliency. It reviews key technologies, identifies and illustrates architecture components, and proposes to use academic infrastructure to validate cyber infrastructure.

Author: Steve Goeringer, Director and Senior Network Engineer

14900 Conference Center Drive
Suite 280
Chantilly, VA 20151
703-955-7770

Introduction

Traditional information assurance techniques have relied on defense-in-depth and a variety of anomaly detection techniques to identify attacks where threats are expected to achieve Cyber Resilience. These techniques have been expensive in terms of both recurring (operational) and non-recurring (development and capital) costs. Moreover, they have not kept pace with the rapid development of IT infrastructure enabling critical mission environments where networks and information resources are highly fluid and continually evolving. Moreover, current threat mitigation and attack management techniques tend to be mission disruptive – the resources under attack are often made unavailable until defensive mechanisms are placed. The result is contrary to the goal – rather than achieving Cyber Resilience, the architectures supporting our missions remain brittle to determined, sustained attacks.

Cyber Resilience needs to be redefined to include the ability for critical infrastructures to continue to operate even while under attack and essential information to remain available to those in need. This is possible using emerging technologies:

- Pervasive data collection from servers, clients, network elements, and application aware sensors at line rates up to even 100Gbps.
- Strong packet manipulation capabilities at rates up to 10Gpbs that leverage deep packet inspection to assess data flows in real-time and filter, modify, and route as appropriate.
- Dynamic multi-layer and multi-technology centralized control to manage threat exposure as attacks are detected.
- Intelligent network architecture that focuses on mission effectiveness rather than basic resource availability.

This paper proposes the application of these Commercial-Off-The-Shelf (COTS) based technologies at scale leveraging distributed resources in combination with traditional lab-based sandboxes. This will provide an environment where mission specific Cyber Resilience solutions (COTS, proprietary, and government provided methods and techniques) can be researched, designed, developed, tested, evaluated, implemented and demonstrated.

The ideas are presented in two sections. The first addresses technologies and architectures that can be investigated to explore Cyber Resilience in depth, achieving the goal to greatly enhance the cyber resilience and performance posture of mission networks. The second part addresses how Cyber Resilience solutions can be developed at scale leveraging national resources while also still enabling development and investigation in sandbox environments.

14900 Conference Center Drive
Suite 280
Chantilly, VA 20151
703-955-7770

Cyber Resilience Technology and Architecture

Over the past three years, several new technologies have become generally available that provide new opportunities for enhanced Cyber Resilience. Several of these may not be obvious to most technologists as related to robust cyber design for IT architecture.

Application performance management – You can't manage something you can't measure. Several suites of COTS software and hardware are providing a breadth and depth of Application Performance Management (APM) capabilities not possible just a few years ago. The result is the ability to develop pervasive visibility throughout an IT enterprise.

- ***Deep packet inspection*** – One of the enabling technologies for APM is deep packet inspection that is integrated into highly available, high performance sensors. However, the benefits of Deep Packet Inspection (DPI) go far beyond APM. DPI enables the ability to inspect, manipulate (change), and route packet in new and innovative ways. This enables custom network design that leverages COTS DPI solutions.

- ***Control plane architectures*** – Telecommunications equipment and software manufactures have worked with academia for over a decade to develop new solutions for managing today's advanced networks. Several components have been developed throughout the industry including the IETF's Generalized Multi-Protocol Label Switching protocol (GMPLS) and the ITU's Autonomously Switched Optical Network (ASON) standard. The result is a wide range of standards and commercial capabilities, each designed to satisfy a specific niche need. BBN's PHAROS, developed initially for DARPA, extends these capabilities to support multi-layer and multi-technology architectures as appropriate for a given mission environment.

- ***Application delivery controllers*** – Traditional optical transport and packet routers and switches provide robust capabilities to get bits across networks. However, they are fairly limited when network services need to be application aware. Traditional load balancers have evolved and are now a new breed of network element – the application delivery controller. Available from many manufacturers with a wide range of features, the Application Delivery Controller (ADC) incorporates traditional routing and switching capabilities and new application aware IT features that enable new service architectures.

- ***OpenFlow*** – As stated on the OpenFlow website (www.openflow.org), "OpenFlow enables networks to evolve, by giving a remote controller the power to modify the behavior of network devices, through a well-defined 'forwarding instruction set." The growing

14900 Conference Center Drive
Suite 280
Chantilly, VA 20151
703-955-7770

OpenFlow ecosystem now includes routers, switches, virtual switches, and access points from a range of vendors." Enterprise IT operators can leverage OpenFlow to enable routing and switching behaviors as they desire rather than the general practices enabled by traditional routing and switching equipment manufacturers.

- ***Flow optimization*** – The suite of Internet protocols has existed for decades. A wide range of implementations have existed for some time to optimize IP performance. However, these have not been generally incorporated into network elements, servers, and application software. Consequently, data networks continue to show poor performance and little resilience to network disruption, even to moderate changes in latency due to path fault recovery. WAN optimization solutions provide the opportunity to increase network "good put" by utilizing existing bandwidth. Moreover, the same mechanisms used to improve network utilization, throughput, and responsiveness often provide improved resilience against network disruption.

Common enablers through all of these technologies include high speed, low power packaged as ASICs and FPGAs, faster and larger storage (high performance memory and Solid State Drives), standardized processors (such as network processors and systems on a chip) that provide modularity beneficial to open source software development, and continued evolution of networking and application algorithms and best practices.

These new technologies and components can be leveraged collectively to create innovative IT enterprise architectures enable new techniques and establish new best practices for Cyber Resilience. This enables a transformation of IT environments to flow based, application aware, converged architectures that leverage distributed architecture and intelligent control to dynamically dilate exposure the cyber threats while managing mission effectiveness. This potential architecture is illustrated on the following page.

14900 Conference Center Drive
Suite 280
Chantilly, VA 20151
703-955-7770

POLAR STAR Consulting, LLC
14900 Conference Center Drive
Suite 280
Chantilly, VA 20151
703-955-7770

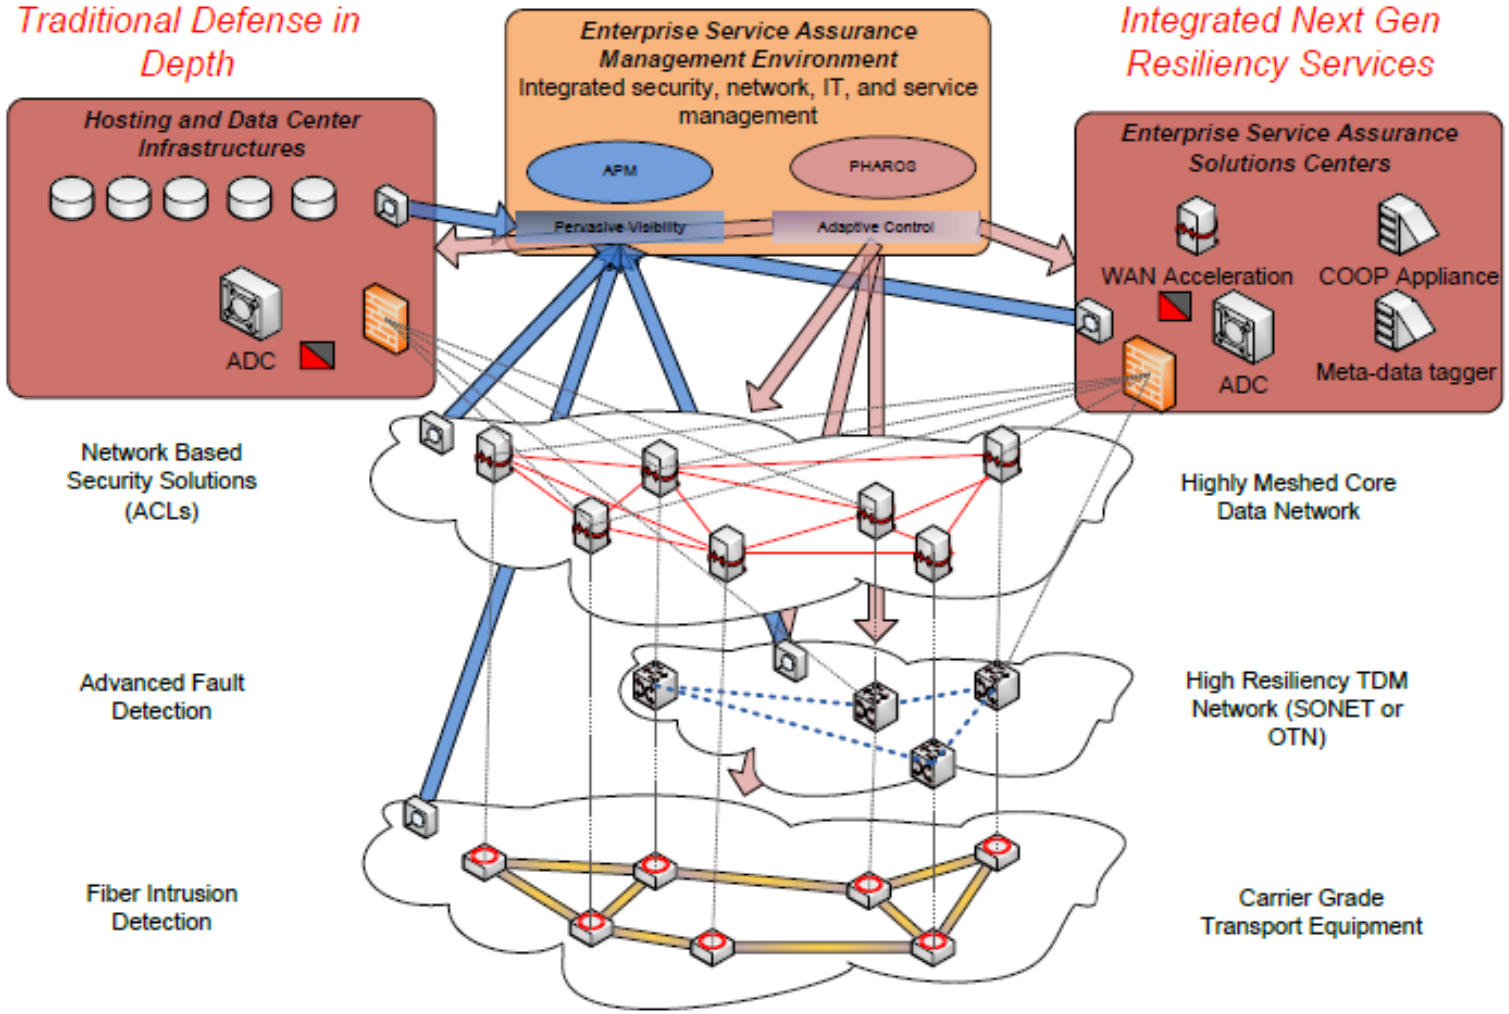

Cyber Resilience Validation

Validation of Cyber Resilience can be performed in three progressive stages. The first stage should be application of traditional operation research techniques and analysis such as Reliability, Availability, Maintainability, and Survivability (RAMS) assessment. Unfortunately, the entire area of information assurance has not received adequate attention in these circles, so techniques of analytically treating Cyber Resilience must be developed.

Given the principles developed through operations research, the next progressive stage is verification in controlled environments (loosely referred to as "labs"). Verification of cyber resilience in a lab environment must be approached as a multidisciplinary endeavor. Traditional functional and performance testing techniques must be integrated with information assurance testing practices. In this way, performance of equipment in various configurations (architectures) can be evaluated while under attack. Most resources for such an environment are expensive, but generally available as COTS.

Unfortunately, the nature of cyber threats is continually evolving. Moreover, it can be difficult to simulate real network behavior at scale in a laboratory environment. In addition, access to researchers can be difficult as the national subject matter experts are distributed geographically among different stakeholders. This leads to the need for a final stage that Cyber Resilience validation – testing and evaluation of Cyber Resilience architectures on real networks.

This must be more than the typical "honey pot". Real networks used for non-sensitive missions should be used. This is readily achievable using networks such GENI, DREN, and other NSF or academic networks. Selection of which resource is appropriate for a given implementation is technology and mission dependent.

GENI provides the ability to virtualize a network environment in a way as to not interfere with other network users. OpenFlow can be leveraged to provide advanced network routing and switching while GENI "stacks" can be used to create application clusters. As GENI is a real world network, performance will be more indicative of actual expected performance in some mission environments. Additional resources and attack vectors can be applied as necessary.

If the nature of research is sensitive or requires technologies not readily integrated into GENI, DREN may provide a more suitable validation architecture. Other network architectures can be approached as well. National and international scale networks, such as Internet2, provide the capability to explore optical networking resilience in combination with routed networks.

Several techniques and methods will be necessary. Exposure to real world (aka, "zero day") threats can be managed using "walled gardens". Network degradations can be managed using laboratory grade devices such as Apposite's Linktrop-10G to create a variety of fault modes. Attacks can be "solicited" using honey pot techniques, or injected using COTS and GOTS exploitation tools.

14900 Conference Center Drive
Suite 280
Chantilly, VA 20151
703-955-7770

Finally, intelligent control can be applied to both develop techniques to avoid threat or attack points while also mitigating exposure to attacks using dynamic filtering and rating limiting on a per flow basis. This may be the essential definition of what it means to be Cyber threat resilient.

Citations

PHAROS: I. Baldine, A. Jackson, J. Jacob, W. Leland, J. Lowry, W. Miliken, P.Pal, R. Ramanathan, K. Rauschenbach, C. Santivanez, and D. Wood, "PHAROS: An Architecture for Next-Generation Core Optical Networks," pp. 154-179, *Next-Generation Internet Architectures and Protocols,* Ed. by Byrav Ramamurthy, George N. Rouskas, and Krishna Moorthy Sivalingam, Cambridge University Press, 2011.

GENI: www.geni.net

14900 Conference Center Drive
Suite 280
Chantilly, VA 20151
703-955-7770