NETWITNESS

NEXTGEN ARCHITECTURE

EXECUTIVE SUMMARY

There is no question: It is difficult to efficiently and effectively investigate security breaches. The architecture of NetWitness® NextGen™ is designed to make this crucial process faster and more effective. The NetWitness NextGen system architecture is based on a simple concept: record your network transactions once, and reuse them often. For many organizations, conventional security and networking technologies have proven insufficient. Security, threat analysis, incident response, risk analysis and compliance all compel organizations to move beyond the capabilities of typical network monitoring solutions. NetWitness NextGen is designed to be the vehicle for that move.

NextGen goes far beyond simply capturing and storing packets or flows and providing network statistics. NextGen simultaneously captures and models network layer and application layer traffic in real-time, retaining full packet payload for complete analysis across a secure and flexible framework. NextGen's modular, logically-tiered architecture makes scaling its coverage to handle even the largest networks a straightforward process. Information generated by NextGen's capture capabilities can be aggregated and concentrated in almost limitless ways to allow analysts to review and make sense of the activities occurring across global networks and multi-gigabit total traffic densities.

Global normalization and real-time synchronization of network metadata, supported by deep context and content knowledge into network sessions enables NextGen to solve complex problems not met by other technologies. This delivers a solution with the ability to evolve as changes to your infrastructure and the threat landscape occur.

Five core features characterize NextGen's revolutionary architecture:

- **Scalable and Reliable Capture Infrastructure** – records everything on a network of virtually every size and scope, and supports multiple reuse of data. Unlimited scalability is built into the NextGen architecture with the design of every component and subsystem.
- **Rules & Alerts** – keep, filter, or flag any dimension that NetWitness analyzes. With this flexibility, complex rules can be applied to live network traffic.
- **FlexParse** – instantly customize the processing and modeling behavior of network capture.
- **Threat Feeds** – use external, intelligence, based on IP addresses, to add contextual content to network traffic.
- **API** - enable rapid development of any conceivable application for analysis of raw network traffic.

This paper, intended for those involved in planning, designing, and implementing a NetWitness NextGen deployment, provides a high-level review of the core concepts of the NetWitness NextGen architecture.

NEXTGEN COMPONENTS

Seven key components make up the NextGen system, bringing advanced capabilities and nearly unlimited scalability to bear on network traffic monitoring and analysis problems. Their operation will be described in the next section, but a description of the components will be helpful in understanding the functions.

DECODER

NetWitness Decoder is the cornerstone of the NextGen infrastructure. Decoder is a distributed, highly configurable 64-bit Linux-based network recording appliance that enables users to collect, filter, and analyze full network traffic in an infinite number of dimensions. Decoder fully reassembles and normalizes traffic at every layer for full session analysis. Decoder dynamically builds a complete taxonomy of data across all layers and applications.

CONCENTRATOR

NetWitness Concentrator is a Linux-based network appliance that extends the reach of NetWitness NextGen across the entire enterprise. Comprehensive network and application layer detail can be aggregated and analyzed across multiple Decoder capture locations.

BROKER

NetWitness Broker is a Linux-based network appliance that brokers and distributes queries across multiple concentration points. NetWitness Broker provides a single ubiquitous view across the entire enterprise, including global geographies and multi-gigabit traffic densities. The scalable nature of the Decoder – Concentrator – Broker logical hierarchy makes it possible to scale a NextGen deployment to the largest enterprise and service provider networks.

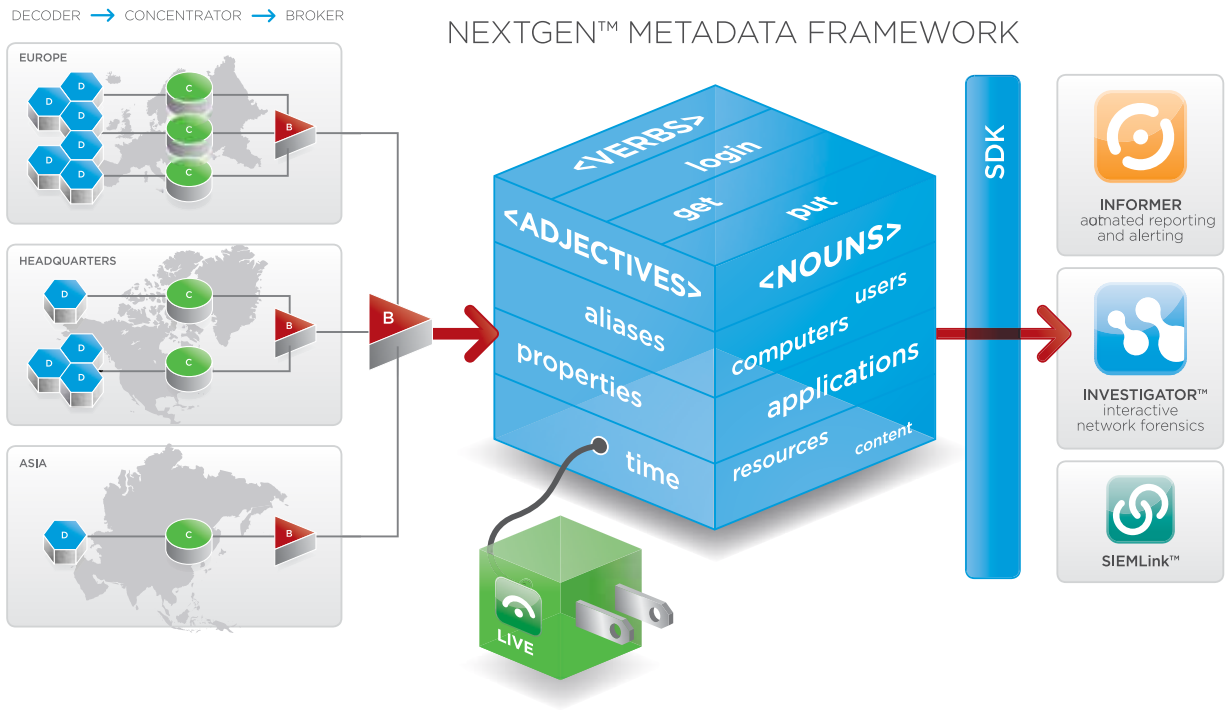

SDK/API

The NetWitness SDK contains an application programming interface (API) that makes it possible to leverage the NetWitness NextGen infrastructure in custom applications. The full featured C API allows for read-only access to query, search and render local and remote data. Its URI-based query language for data retrieval is a powerful tool for integrating NextGen into a workflow/ticketing/incident management framework.

INVESTIGATOR

NetWitness Investigator is a Windows-based application that provides rapid and efficient free-form contextual analysis of raw data captured and reconstructed by NetWitness NextGen. Investigator uses a lexicon of nouns, verbs and adjectives to represent the application layer protocols parsed by NextGen during session reconstruction.

INFORMER

NetWitness Informer uses the comprehensive network traffic captured and reconstructed by the NextGen infrastructure to provide a up-to-the-minute glimpse into incidents, threats, anomalies, mis-configurations, compliance violations, and other activities on the network. Informer is a fully interactive web-based report engine with design features that enable users at virtually any level of technical sophistication to create needed reports without sophisticated programming or outside help.

This detailed information is made available through reports and real-time charts. Every piece of metadata available, in any combination from a network of any size, can be reported on, alerted, plotted over time and presented in interactive formats.

SIEMLINK

SIEMLink enables instant integration of NetWitness NextGen technology with existing enterprise security infrastructures. SIEMLink, a light-weight Windows application, is designed to act as a transparent, translator of critical security event data between Web-based consoles such as security event and information management (SIEM) systems and network and system management (NSM) programs. SIEMLink requires no special coding or systems integration work to link an organization's existing SIEM with NextGen.

THE NEXTGEN ARCHITECTURE

The NextGen system architecture is built on a framework that uses a custom asynchronous protocol to send and retrieve data across a highly scalable system. The system also supports SSL for secure communications among servers and client applications. This approach enables every product in the NextGen infrastructure to interact in real-time.

NetWitness NextGen is made of five logical sub-systems:

- Network Capture performed by Decoder
- Data Processing performed by Decoder
- Data Synchronization performed by Concentrator
- Indexing performed by Concentrator
- Analysis performed by Investigator, Informer or any API-built application.

NETWORK CAPTURE

Network capture technology has existed for years to allow troubleshooting and network performance monitoring. The NetWitness capture subsystem uses existing technology to stream and buffer packets at above-saturated gigabit speeds to hand off to the processing engine. An ability to deploy multiple Decoder modules across a network makes it possible to use NextGen to scale the capture and monitoring functions from enterprise networks up to the networks of global service providers. NextGen takes advantage of open source software and known driver optimizations, and capitalizes on the multi-core processors common in newer computing platforms.

DATA PROCESSING

After capture, NextGen processes full packets to analyze the data. NextGen focuses on processing application-level information, which makes network monitoring applicable to business needs beyond security. NextGen does this by processing sessions, not just packets. This translates to more complete, easily managed and more relevant context around network activity. For example, on an average day, a link may produce 90,000 packets per second, or 7,000 sessions per second. This reduction in data volume creates a compelling reason to look at network activity in sessions versus packets.

NextGen's processing consists of three phases in a process protected under US Patent 7,016,951:

- Packet reassembly reconstitutes packets into sessions, making the evaluation of the entire stream possible. Viewing the session payload provides the full picture needed for today's threats and networking challenges.

- Application parsing via flexible parsing modules enables anyone to define how NextGen processes, categorizes, and organizes network traffic. Parsing enables the definitive identification of session service types, and defines what metadata is extracted for the NextGen data store. FlexParse allows NetWitness operators to configure, in XML, how NextGen identifies applications and what it extracts for analysis.

- Content extraction is based on parser instructions. Metadata is extracted and stored with the full session content. The NetWitness metadata, combined with its native full packet storage, provides a network recording infrastructure capable of providing insight and behavior detail into every network event: internal, external, malicious or benign.

- Default NetWitness parsing provides protocol and application exploitation of: HTTP, FTP, TFTP, TELNET, SMTP, POP3, NNTP, DNS, HTTPS, SSL, SOCKS, SSH, Vcard, PGP, SMIME, REGEX, DHCP, NETBIOS, SMB/CIFS, SNMP, NFS, RIP, MSRPC, Lotus Notes®, TDS(MSSQL), TNS(Oracle®), IRC, Lotus Sametime®, MSN IM, RTP, Gnutella, Yahoo Messenger, AIM, SIP, H.323, Net2Phone®,Yahoo Chat, SCCP (Cisco® Skinny), Bittorrent, GTALK, Hotmail, Yahoo Mail, GMail, TOR Social Networking, Fast Flux and other protocols.

DATA SYNCHRONIZATION

For a network monitoring solution to be effective, there needs to be synchronization across all capture locations. A point solution for capture is inadequate and limits analytical effectiveness. Batch-based systems lack the real-time context demanded by modern network security. NextGen has the ability to aggregate metadata into logical views and to provide global visibility and analysis of the information. This enables global reporting and incident response of massive amounts of network traffic instantly.

INDEXING

NextGen indexing provides rapid access to a massive data store. NetWitness has a proprietary database system to manage its data and has achieved unparalleled performance metrics for network session indexing.

To put some scale on the scope of NextGen's database performance, consider the following:

- In 2006 Visa, on the busiest day of their year, on the busiest hour of that day, processed 6,800 transactions per second.

- The record performance from Oracle stands at above 60,000 transactions per second. To achieve this rate, they use 190 separate processors for a total cost of $4.2 million dollars. For a single system, and $80,000, they can provide 100,000 transactions per minute, or 2,000 per second.

- A single NetWitness Concentrator can process more than 10,000 sessions per second. Hardware I/O, not application design, limits the performance.

The NextGen indexing sub-system is the most technologically sophisticated component of the NetWitness deployment, and the heart of its ability to scale to cover exceptionally large networks. Overall data volumes in the terabyte-plus per day range make effective indexing critical for useful deployments on global high-traffic-volume networks.

ANALYSIS

The final component of NextGen is the NextGen API. This programming interface enables client applications to interact with captured data residing on NextGen servers. The API manages all user access to the data stores, including simple metadata searches, complex full content searches, data export and caching, and everything in between. Every application NetWitness Corporation builds uses the API, and every custom application NetWitness customers create leverages it as well.

DEPLOYMENT

Network administrators or security specialists can install and deploy NextGen appliances in mere moments. The appliances operate by promiscuously listening to a SPAN port on an existing switch or a feed from a network tap.

Defining capture location for NextGen appliances in an enterprise is simple. Place Decoders wherever needed to capture traffic: egress, core, facility, or segment (typically similar to IDS/IPS capture locations). Place Concentrators where needed to consolidate Decoder data for reporting, alerting, and analysis through Investigator or Informer.

A single Broker appliance provides consolidated views across multiple Concentrators, enabling a single unified view into the transactions crossing a network of unlimited complexity.

There are several guidelines critical to a successful NetWitness NextGen deployment:

1. When monitoring at an egress location, place the Decoder inside the firewall. This approach helps manage data characteristics that affect NetWitness database behavior. For example, by design, a denial of service attack will be recorded and modeled by NetWitness appliances. However, an enormous amount of ultimately meaningless, unique network entities will fill the data structures with noise that may ultimately affect speed and effectiveness of analysis.
2. Apply filtering rules to eliminate noisy or analytically useless data, such as backup traffic, internal video, and audio screening.
3. Verify that peak loads do not exceed the rating of the appliances at 1Gbs or packet loss may occur.
4. After deployment, allow for a baseline configuration period to tune the index to the environment. This process is aided by the NetWitness Tune-up service, an offering of the NetWitness Professional Services organization.

NextGen's modular scalable architecture is designed to expand as the network infrastructure changes. As a guideline for storage, expect a saturated 100Mbs link to produce approximately 1TB of data per day. NetWitness database overhead ranges from 5% to 15% of raw data captured. Data retention policies and network monitoring requirements will dictate the size of the storage component associated with a NextGen deployment. NetWitness engineers are available to assist in the process of scoping the optimal size of the initial storage component.

PERFORMANCE

Performance metrics for a NetWitness NextGen solution can vary greatly across implementations and environments especially given the highly-scalable nature of the NexGen architecture. System behavior depends not only on the volume of network traffic, but the NextGen application profile, which is directly related to the enabled parsers and the number of real-time alerts loaded into the Decoder.

The same caveat applies to query response time. Meta and index information cached or present in memory responds in seconds. However, data flushed to disk may take minutes to return. In the case of full content search, this may take minutes or hours depending on the amount of data encompassed in the search. These characteristics depend on the appliance platform purchased and how it is deployed. Typically, the more memory and more resources a Concentrator has, the faster it will perform.

Query response time is highly dependent on the existing volume of data, the profile of traffic, the time frame of query, the uniqueness of query, user navigation, and the configuration of appliances. Upon installation, NetWitness engineers make every effort to optimize query performance.

NetWitness Corporation | 500 Grove Street, Suite 300 | Herndon, VA 20170
T: 703.889.8950 | F: 703.651.3126 | info@netwitness.com