Solution Profile

August 2020

NOKIA Deepfield

A New Approach to Network Analytics

Author: Grant Lenahan

www.appledoreresearch.com

Appledore RESEARCH

Published by Appledore Research LLC • 44 Summer Street Dover, NH. 03820

Tel: +1 603 969 2125 • Email: grant@appledorerg.com• www.appledoreresearch.com

Figures and projections contained in this report are based on publicly available information only and are produced by the Research Division of Appledore Research LLC independently of any client-specific work within Appledore Research LLC. The opinions expressed are those of the stated authors only.

Appledore Research LLC recognizes that many terms appearing in this report are proprietary; all such trademarks are acknowledged, and every effort has been made to indicate them by the normal USA publishing standards. However, the presence of a term, in whatever form, does not affect its legal status as a trademark.

Appledore Research LLC maintains that all reasonable care and skill have been used in the compilation of this publication. However, Appledore Research LLC shall not be under any liability for loss or damage (including consequential loss) whatsoever or howsoever arising because of the use of this publication by the customer, his servants, agents or any third party.

Publish date: 8/31/2020

Cover image by Francis Haysom

EXECUTIVE SUMMARY

Nokia Deepfield is a network and cloud focused “insight” and analytics platform, acquired by Nokia in 2017. It identifies, maps and analyses network traffic, *with a focus on visibility of applications, their sources and destinations, and classification of traffic types* – whether Netflix servers, CDNs, or “things” of any type. Nokia Deepfield, in Appledore Research’s opinion, represents a unique and highly relevant approach that combines an understanding of network performance with an understanding of how that network is used. This context is increasingly essential in today’s app-driven, dynamic world.

Deepfield is specifically NOT a DPI (Deep Packet Inspection) or traditional data collection system. Rather, it is a *big data platform* that concentrates on multi-dimensional analysis of data gathered from diverse sources – network elements (telemetry, flow, configuration, signaling), performance management systems, DNS servers, AAA servers, etc. In fact, part of Deepfield’s differentiation is that it can be application-aware without requiring costly and equipment-intensive DPI technology. This approach allows Deepfield to extend the ‘dimensions’ of their insight by adding and correlating additional data sets. An example is the Deepfield Subscriber Intelligence module – which extends network/cloud intelligence with subscriber and billing data to provide additional insights on how network services are delivered and consumed.

A key point of Deepfield’s differentiation is achieved through a patented approach that generates the “Deepfield Genome” – essentially, a dynamically updated mapping of the whole internet used to identify applications, associated source and destination addresses “real” DNS names, and how that traffic flows through and interacts with the network. This visibility is becoming increasingly important as many content and application providers use multiple (and shared) CDNs, which makes it harder for service providers to fully understand all the traffic coming from them. For example, Deepfield notes, Disney+ has used 6 different CDNs to deliver their content to European customers.

A subset of this Genome covers the security aspects (e.g., known malicious traffic sources) and while another use of Deepfield Genome is to identify the nature of end-point IP addresses - “things” on the Internet – useful as IoT devices become data sources and open to exploitation.

Deepfield focuses on three applications of this technology: 1) Insight, 2) Performance, and 3) Security. Furthermore, Deepfield’s security applications are pre-configured to support automation (Deepfield’s “auto-mitigation” feature. Out of the box, Deepfield claims to identify (and therefore allow automated mitigation of) the majority of common DDoS attacks. Further customized by customers can extend this to remediate other security threats.

COMPANY

Company Basics

Name	Nokia Networks

Year founded	1865
Headquarters	Espoo, Finland
CEO	Pekka Lundmark
Company Type	Publicly Traded
Revenue 2017 (Deepfield)	Not broken out.
Employees	98,322
Product segment	Telecommunications, full-line NEP (Network Equipment Supplier)
Geographic focus	Global
Primary products	RANs (4G, 5G), Optical, IP, Management Software, Core networking/IMS. **This profile covers Deepfield Only**
Key partners	Many
Key Nokia customers	Verizon, AT&T, BT, Telecom Italia, Telefonica *Deepfield: see below*

Source: Appledore Research

DEEPFIELD OVERVIEW AND APPROACH

A focus on analytics, not DPI or probing

In several significant ways, Nokia Deepfield takes a different, and forward-looking approach to analytics, service assurance, and security/DDoS mitigation. Deepfield is a network and cloud-focused "insight" and big data analytics platform, acquired by Nokia in 2017. It identifies, maps and analyses network traffic and traffic sources, *with a focus on visibility of applications, their sources and destinations and traffic types* – whether Netflix servers, CDNs, or "things" of any type. This visibility is applied to a portfolio that addresses "Insight, Performance and Security", as expanded below.

Deepfield, in Appledore Research's opinion, represents a unique and highly relevant approach that combines an understanding of network performance with an understanding of how that network is used. This context is increasingly essential in today's app-driven, dynamic world.

Furthermore, Deepfield is specifically NOT a DPI (Deep Packet Inspection) or traditional data collection system. Rather, it concentrates on multi-dimensional analysis of data gathered from diverse sources – network elements (telemetry, flow, configuration, signaling), performance management systems, DNS servers, AAA servers, etc. Nokia Deepfield is one of the firms driving a new approach to data collection and processing, which takes advantage of the hugely increased processing and forwarding capabilities of modern network elements – and especially of the emerging generation of packet forwarding silicon, including Nokia's own FP4 network traffic processing chip. These advances are making obsolete the traditional limits on how much data could be polled or streamed from a network element without impacting its performance. That said, data from probes or other data collection software may be inputs to Deepfield and can be integrated through interface "connectors".

Deepfield Genome: "Characterizing the internet"

Deepfield's patented Cloud Genome and Secure Genome technology allows it to be application-aware, device-aware and traffic source location-aware without requiring costly and equipment intensive DPI technology. Deepfield Cloud and Secure Genomes (collectively referred to as "Deepfield Genome") are live data feeds, updated continuously by Deepfield's own cloud-based probes and made available to Deepfield customers in the form of real-time data feeds that enhance data collected from the network. Essentially, Deepfield Genome represents a complete map of internet's cloud applications and services (or, what they refer to as internet 'service supply chain'). Once matched with the data collected from a particular carrier's network, the result is a unique and detailed map of how network's internal (on-net) an internet (off-net) traffic interact with the network.

Furthermore, Deepfield Genome identifies the nature and addresses of "things" on the Internet; increasingly useful as IoT devices become data sources and open to exploitation. Consequently, Deepfield need not inspect packets deeply to identify traffic, and it can also flag when traffic is associated with a new or inappropriate source. Like turning on caller ID for Internet traffic, this provides critical insight on what flows into the network as the majority of traffic originates off-net and flows over-the-top (OTT). Moreover, the use of end-to-end encryption for many services and network protocols renders an increasing portion of internet traffic opaque for DPI techniques.

Separating intelligence (centralized) from policy enforcement (distributed)

Another important and differentiating advantage of the Deepfield approach is related to leveraging big data analytics for security. Traditional vendors (using DPI and probes) implement data inspection and related policy enforcement actions in the data plane. This is a processing-intensive approach which can adversely affect throughput, and therefore cost. And this cost generally scales with volume, which makes it increasingly unsustainable. Deepfield's approach, by way of contrast, is decoupled (and centralized) from network-distributed policy enforcement and traffic cleansing. The primary implication (benefit) is that this cost is shared across many sites and in fact many customers. The secondary implication is that Deepfield's approach scales better with both the volume of threats and the volume of traffic.[1]

License and deployment options

Deepfield's software platform can be provided on either a license basis or as "aaS" basis. It can run on on-premise servers or can be hosted - in a cloud environment, in which case secure connectivity with network elements is required to pull data from a customer's network. It is worth noting that to ingest truly network-wide quantities of data (petabytes), cloud architectures are vastly superior in terms of efficiency, elasticity of scale and therefore cost. Deepfield is typically licensed based on both a) volume of data ingested and analyzed, and b) applications supported.

DEEPFIELD SOLUTION ARCHITECTURE:

Deepfield is fundamentally a big data analytics platform as opposed to a probe or data collection platform with analytics added later. This is a conscious decision designed to ingest and process today's cloud-volumes of data. Nokia Deepfield has concentrated on support for a wide range of real-time, high-data-rate, streaming interfaces. Their primary data sources are the network elements themselves. The solution supports a long list of network element vendors and elements, including Nokia, Cisco, Juniper, Huawei, Brocade, F5, Arista, Arbor, A10, VeriSign, and Radware. Deepfield also supports standard and de-facto interfaces such as vProbe and Splunk.

Deepfield emphasizes that the industry is being driven, especially by the webscale cloud players such as Google, toward much faster, near real-time analytics based on streaming telemetry data, rather than periodic transactions such as FTP transfers and polling of SNMP MIBs. A related trend is the emergence of new chipsets, such as Nokia's own FP4 network processor, that allow for full line-rate operation, while at the same time supporting enhanced packet inspection and header manipulation / steering. Appledore believes that over time, such "network function-native" data collection methods will be highly desirable to increase the frequency, scope and cost efficiency of data collection, and also to simplify collection and identification of streams, as services and VNFs become more dynamic.

[1] Much of the efficiency, and therefore better scaling, comes from the fact that Deepfield's approach does the work to identify the source only once, and from there it is a simple policy rule based on fast, low overhead pattern matching. They also explicitly harness the power of the emerging generation of hyper capable packet forwarding chips that can filter and selectively forward with no capacity penalty.

Deepfield emphasizes the scalability of its approach, while also emphasizing the trend toward more data, and more need for truly real-time insights and resulting actions. Specifically, they deliver scale and near-real-time performance in three ways:

1. A reliance on high volume, native streaming data interfaces rather than costly and bandwidth limited probes
2. The Deepfield Genome approach that constantly inventories the web, identifying and characterizing IP addresses, including "things" – therefore obviating the need for computationally intensive and slow DPI
3. Cloud-native "horizontal" scaling, and the ability to run in cloud environments (initially AWS)

   Deepfield's cloud-native architecture allows dynamic scaling (up and down, as needed) as well as dynamic re-allocation of compute resources among collection, processing, database analysis, etc.

The following diagram illustrates the Nokia Deepfield architecture. A few concepts / corollaries are worth pointing out:

- "Southbound" (at the bottom) we see both streaming and non-streaming interfaces to constantly collect data from network nodes, and things.
- "Southbound" we also see bi-directional interfaces for automation. These are used to configure and communicate with devices that Deepfield may configure, e.g.: to set filters (identify specific information) and proscribe resulting actions on that data (e.g.: remit to Deepfield, send to cleansing center, drop/block).
- "Northbound" we see interfaces to OSS/BSS and other such sources of commercial, subscriber and related data (e.g.; CDRs, subscriber plan info, etc.)
- Westbound (left) is the activity to generate and update Deepfield's patented Genome; both Secure Genome and Cloud Genome. The Genome data set is published as a 'live' data feed to its Deepfield customer deployments.
- Eastbound we see programmable and pre-defined APIs that allows Deepfield data to be utilized in a variety of ways, by users, GUIs, and systems.

Figure 1: Nokia Deepfield High Level Architecture

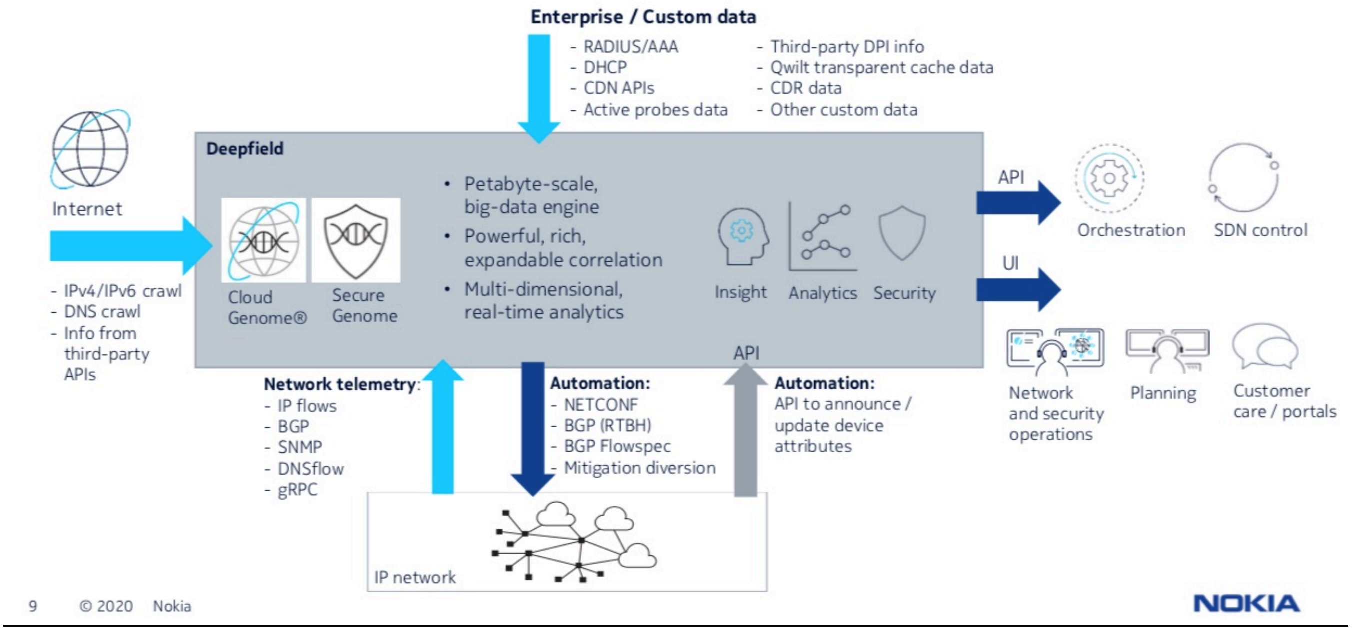

Enterprise / Custom data

- RADIUS/AAA
- DHCP
- CDN APIs
- Active probes data
- Third-party DPI info
- Qwilt transparent cache data
- CDR data
- Other custom data

Internet

- IPv4/IPv6 crawl
- DNS crawl
- Info from third-party APIs

Deepfield

Cloud Genome®

Secure Genome

- Petabyte-scale, big-data engine
- Powerful, rich, expandable correlation
- Multi-dimensional, real-time analytics

Insight

Analytics

Security

API

Orchestration

SDN control

UI

Network and security operations

Planning

Customer care / portals

Network telemetry:

- IP flows
- BGP
- SNMP
- DNSflow
- gRPC

Automation:

- NETCONF
- BGP (RTBH)
- BGP Flowspec
- Mitigation diversion

API

Automation:
API to announce / update device attributes

IP network

9 © 2020 Nokia

NOKIA

Source: Nokia Deepfield

Nokia emphasizes that the Deepfield implementation design is highly configurable and built to support high rates and large volumes of data; in their words “petabytes”. This extends from the inbound (Southbound) APIs, through the internal communications and analytics/filtering architecture, and to the database.

Topology is central to Deepfield, since it essentially maps application data to ingress points, egress points, and segments through the network. Rather than point to an external topology reference, Deepfield builds its own logical topology from various signaling and telemetry feeds – e.g., BGP and others. The topology is stored as part of its multi-dimensional database structure and is constantly updated to reflect the increasingly dynamic network.

CUSTOMERS

Table 3: Market Traction and Evidence

Metric	Evidence
Leading customers	Telefonica Spain (public). Several un-named T1 CSPs. Outside telecom Deepfield claim deployments by several large web-scale platers, IXPs, and digital enterprises.
Number of commercial deployments as of 3Q2020	>50

Number of ongoing trials/PoCs as of 3Q2020	“Several”
Applications supported in commercial service	All applications noted below, including Insight, Performance and Security, with security including both Network Security and Network (DDoS) Firewall

Source: Nokia, Appledore Research

DEEPFIELD APPLICATION PORTFOLIO

Deepfield is both a platform and a set of packaged applications. As a platform it enables interactive use and development; but it is also sold as a set of applications that includes pre-built ML processes, GUIs and integration to data sources and control/throttling devices, primarily routers acting as a form of Policy Enforcement Point.

Deepfield’s portfolio of applications is logically organized along the lines of “Insight”, “Performance” and “Security.” The figure below illustrates how Deepfield views its application portfolio; we comment below.

Figure 1: Deepfield Portfolio (pre-made applications)

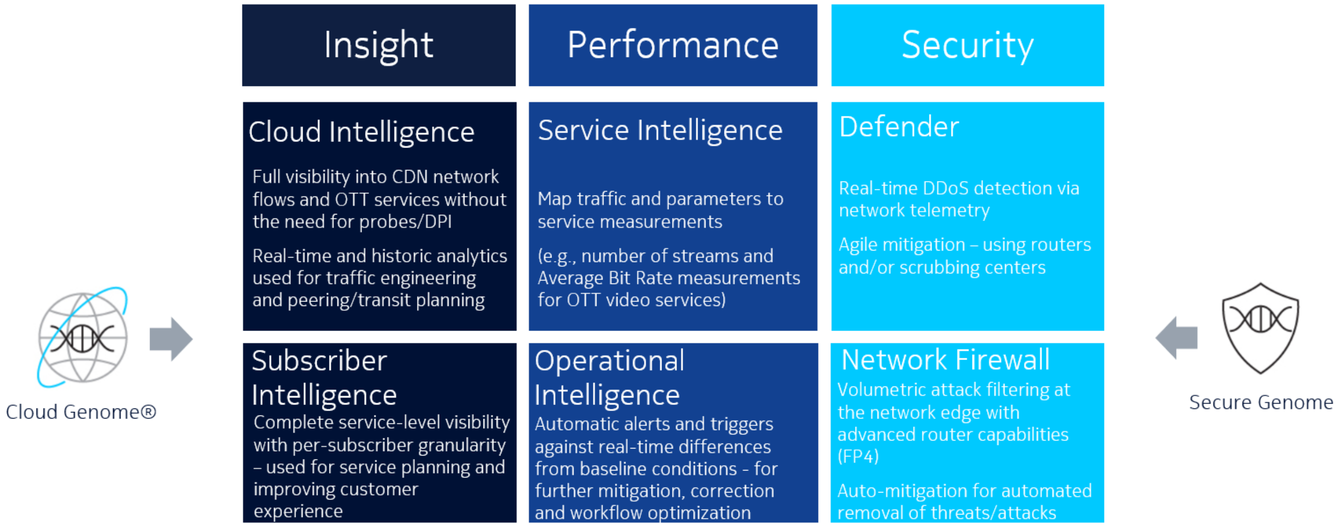

Insight

Performance

Security

Cloud Intelligence

Full visibility into CDN network flows and OTT services without the need for probes/DPI

Real-time and historic analytics used for traffic engineering and peering/transit planning

Service Intelligence

Map traffic and parameters to service measurements

(e.g., number of streams and Average Bit Rate measurements for OTT video services)

Defender

Real-time DDoS detection via network telemetry

Agile mitigation – using routers and/or scrubbing centers

Subscriber Intelligence

Complete service-level visibility with per-subscriber granularity – used for service planning and improving customer experience

Operational Intelligence

Automatic alerts and triggers against real-time differences from baseline conditions - for further mitigation, correction and workflow optimization

Network Firewall

Volumetric attack filtering at the network edge with advanced router capabilities (FP4)

Auto-mitigation for automated removal of threats/attacks

Courtesy: Nokia Deepfield

It is clear from Nokia’s graphic, but also from our discussion above, that Deepfield Genome is central to all applications. The Genome is what allows for insight into the intersection of applications (e.g., Netflix, Apple downloads) and flows (traffic volume, ingress egress). The Genome furthermore allows for such insight at scale, without the cost and processing overhead of DPI which must be used somewhat sparingly.

Insight and Performance are really two analytical sides of the same data corpus – one concentrates on understanding the drivers of traffic, the locations of traffic and how it impacts various subscriber groups (e.g., by service plan, lifetime value, etc.). By having both insight into "what apps" are driving traffic/issues, and "who it is affecting", such context-aware analysis lends itself to actions – whether prioritizing certain flows, re-routing traffic, or targeting capacity expansion to name a few examples.

Security uses a very similar approach – contextual analysis based in part on a pre-built Genome, and in part on network traffic pattern matching – identifying abnormal behavior. It is unique in a few aspects as well, some of which are discussed in more depth below. Specifically, security demands a slightly different Genome. In addition, that genome must characterize not only a limited number of source sites, but "thing" types that may be hijacked in, for example, a DDoS attack. Finally, security demands real-time actions – whether filtering based on source-destination pair rules (network firewall) or throttling/scrubbing/dropping volumes of data once an attack is detected and characterized. Security is also unique in the fact that it has such dire repercussions both monetarily and to enterprise' reputations.

Deepfield has two security apps: "Defender" and "Network Firewall". The Firewall application, added recently, takes DDoS detection and remediation to the next level taking advantage of a new generation of highly capable routing chips such as Nokia's own FP4. Network Firewall dynamically identifies threats and places filters on these routers at the network edge to proactively prevent attacks (and spare the backhaul and cost of scrubbing centers). As DDOS attacks are increasingly coming both from outside and inside a service provider's own network, Deepfield's focus on '360-degree protection' identifies more threats than externally-focused approaches.

DEEPFIELD GENOME

A key differentiator

If we had to choose one unique differentiator to characterize Deepfield, it would be Deepfield Genome. This is because:

1. The Genome is, to our knowledge completely unique to Deepfield (and based on patented technology)
2. It allows classification (including being source and destination aware) of all types of IP flows, including very high-volume flows that might otherwise prove impractical due to throughput or cost.
3. The Genome approach removes the need for costly, capacity-constrained DPI, while being able to identify the nature and origin of traffic.

Nokia summarizes Cloud Genome as "*technology that maps the global service supply chain to provide unprecedented visibility (over 90% of all applications and services) for service providers* *(cable providers, cloud providers and telcos) as well as large digital enterprises*." In Appledore Research's words, Cloud Genome is a patented, proprietary approach to identifying (Nokia claims) *every* endpoint on the Internet and *every* application stream/type - including those endpoint sources down to IoT "things" and is able to identify when redirected to or from CDNs or other relaying devices.

Collectively, Genome's attributes enable lower costs, better scaling, inspection of terabyte-level flows, and faster ability to characterize new flows and ingress points. Critically, it allows characterization of data through the network core based on:

1. Where it originated
2. The type of device it originated on
3. The nature of the application – voice? Conferencing? Streaming Video? Security threat? DDoS?

Genome makeup and creation

Deepfield Genome is sub-divided into two corpuses:

1. *Secure Genome* – which identifies security risks, whether bad actors in their own right, or compromised sites or devices, or other public resources that can be abused (e.g., known servers).
2. *Cloud Genome* – which identifies all sites on the global internet, by originating domain (e.g., Google.com, Zoom.com, Netflix.com) and by application/content type.

Deepfield Genome is assembled by Deepfield's active, cloud-based agents that constantly crawl the global internet, and by algorithms that learn what application specific IP addresses are associated with. The methods use are a combination of patented technology and confidential intellectual property. Deepfield Genome feeds are available to all Deepfield applications – running in customers' networks or in the cloud, in the form of real-time data feeds. In other words, Deepfield Genome is constructed and maintained by Deepfield, and made available to their customers through real-time updates.

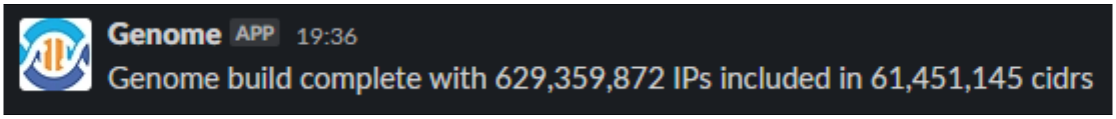
Genome APP 19:36
Genome build complete with 629,359,872 IPs included in 61,451,145 cidrs

What is critical to understand is that the Genomes, once learned, allows fast, low-cost (in money and processing) characterization of ALL data flowing across one's network, at huge scale. No longer is it necessary to focus expensive DPI assets surgically on the most critical subset of your traffic; you can characterize – and act on (groom, block, prioritize) all of it.

Acting in near-real-time on these insights, it also allows Deepfield to become the intelligence behind an active ecosystem that takes proactive steps to remediate problems – whether security, congestion or others. We will cover both applications and automation in some depth below.

Contrast with Traditional Method: DPI

Most systems identify applications via DPI. DPI is a powerful tool, but it has drawbacks; among those it is expensive, has limited capacity/throughput and it is often foiled by encryption. Deepfield re-uses existing network sources of data (reducing duplication and cost) and avoids the need to deeply inspect every (huge) stream to identify it.

Cloud Genome avoids the need for decryption or fingerprinting, by employing software agents to crawl the Internet (somewhat analogous to how search engines catalog the web) and collect publicly-facing data on every "thing" and "application" on the Internet – (and thus, on application traffic traversing SP networks). This data is used to maintain an up-to-date database of applications and their sources. Once the Genome data feed is correlated with the information obtained from one's own network, it provides valuable multi-dimensional insights that extend well beyond the SP's network interfaces / POPs. The resulting network intelligence allows a less processing-intensive linking of data flows to application types, and also allows the identification of anomalies (such as when requests or streams come from the wrong source or source type). Deepfield's real differentiator – beyond a modern and silo-free architecture, is this ability to efficiently map flows to apps and to sources and draw business inferences from these associations. One visualization of what genome can illustrate is shown below.

Figure 6: Traffic sample as mapped by Deepfield using Cloud Genome illustrating depth of insight available

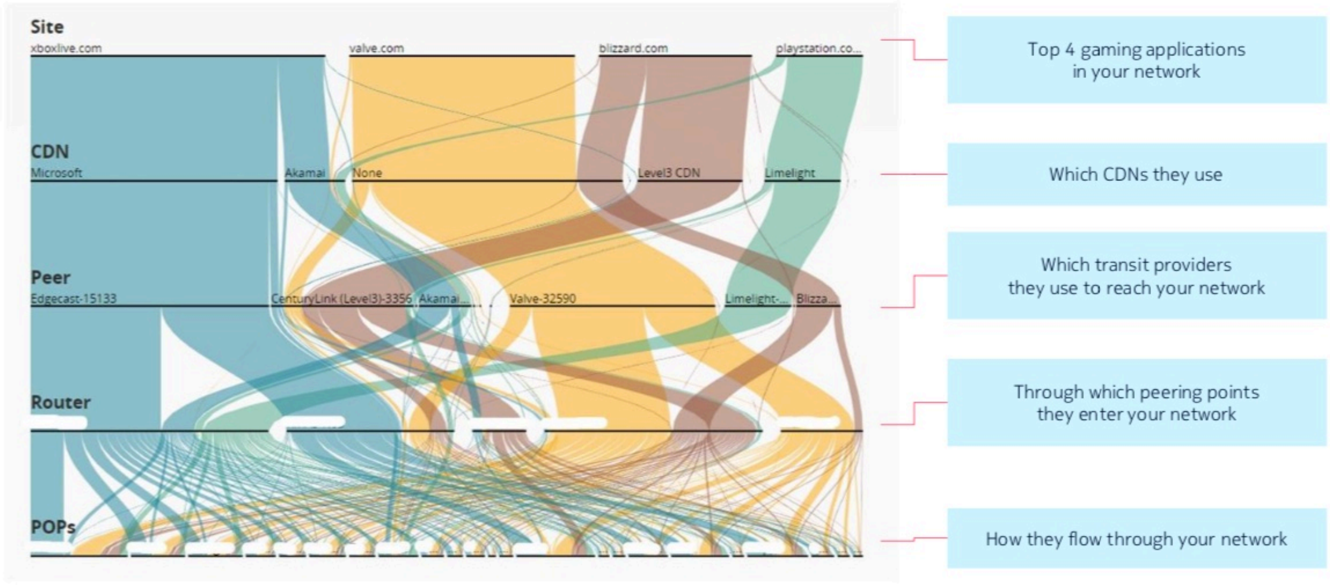

Top 4 gaming applications in your network

Which CDNs they use

Which transit providers they use to reach your network

Through which peering points they enter your network

How they flow through your network

Source: Nokia Deepfield

Metrics

Creating Deepfield Genome is a huge task and must be ongoing to maintain currency. This lends itself well to a shared resource that Deepfield maintains and gives all customers access to – pooling resources and making this feasible. The metrics illustrate the task:

- Over 5B IP addresses (both IPV4 and 6)
- Crawls 200M-1B addresses daily
- Tags >150M of the most relevant IP addresses for fast use (index)
- Maintains 100s of 1000s of "hidden" DNS mappings
- Employs more than 100 Machine Learning rules

In summary, Deepfield Genome is a huge, shared database of IP addresses, classified against useful metrics (bad actors, prominent applications), allowing the hard work of crawling and ML to be accessed by Deepfield customers and used with low overhead and in real time. It is Deepfield's source of much context – and therefore intelligence.

MACHINE LEARNING

Machine learning is a core capability employed by Deepfield, fundamentally in two areas:

1. It is employed to learn "signatures" of network flows, associate these with IP addresses, and therefore build a reliable Genome (of both cloud applications/sources and security threats) that can be queried quickly and with minimal overhead
2. It is also employed to identify "normal" vs. "abnormal" traffic profiles, that enable Deepfield to extend its capabilities to network security – to thwart DDoS and other security threats. Once identified, these are also used to add "bad actor" and/or compromised IP addresses (servers, things) to the Secure Genome.

Deepfield's ability to employ ML is critical to its success. By employing ML, Deepfield avoids the laborious and error-prone manual development of rules and provides the customers with the benefit of having access to the most up-to-date view of internet traffic. And, as noted above, this machine and research-intensive function is shared across customer deployments, making it affordable and spreading benefits widely (a basic concept that Appledore have advocated in everything from VNF models to operational best practices).

AUTOMATION

Deepfield is all about network insight; insight that is real-time, application-aware and source-aware. While such data is inherently useful, the combination of understanding the source and the application in real-time lends itself ideally to automated remediation of congestion, QoE, cost inefficiencies, or security threats or attacks, most notably DDoS attacks.

Automation typically takes the form of a closed loop, employing Deepfield, data sources, the Genomes, and network elements (typically advanced routers that are capable of taking action based on Deepfield insights). The general flow of those loops is as follows:

1. Build Genomes and models of "normal" traffic behavior
   (constant, centralized, real-time activity in the background)
2. Collect traffic data, complete with context for sources, destinations, application type, and whether it meets normal traffic profiles
3. Identify a condition that requires action
   a. E.g.: congestion, DDOS attack, …
   b. Identify unusual patterns
   c. Combine pattern recognition with Genomes for better selectivity and specificity (fewer false positives and negatives)
4. Take action:
   a. Block a flow (using dynamic filters placed on routers)
   b. Redirect a flow (instructing SDN controllers on the intent)
5. Change flow priorities (QoS) to minimize service and customer impacts in times of congestion or failure
6. SDN controllers and/or routers drop or re-route data depending on the situation (there are several methods to achieve this)
7. Back to #1 to deliver constant vigilance.

The diagram below illustrates the general structure of such loops, and where Deepfield fits into the architecture.

Figure 2: Illustrative Closed Loop Employing Deepfield

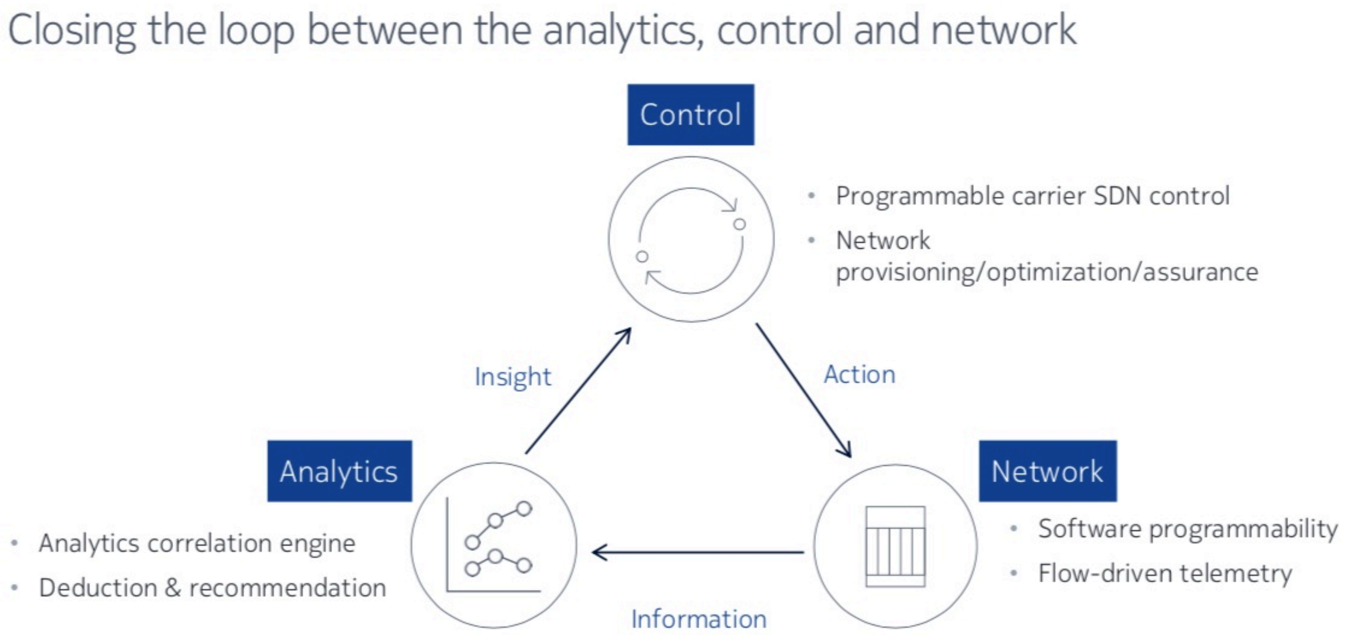

Closing the loop between the analytics, control and network

Control

- Programmable carrier SDN control
- Network provisioning/optimization/assurance

Insight

Action

Analytics

- Analytics correlation engine
- Deduction & recommendation

Network

- Software programmability
- Flow-driven telemetry

Information

Courtesy: Nokia

The case above maps well to the Appledore closed loop reference architecture and therefore also to a traditional control loop based on best practices in control theory. In a generic view, Deepfield provides analytics, ML and context-aware intelligence. SDN controllers may act as an intermediary controller to configure complex flows. And routers, especially next generation of smart edge routers, act as intelligent policy enforcement points, able to execute actions such as redirecting traffic around congestion or dropping malicious traffic. In the context of security, decoupling of security intelligence (control) from the policy enforcement means that network protection capability is not hindered by scaling (as is the case with legacy approaches).

The key takeaway is that Deepfield does not simply identify issues and generate insights. They have gone the extra mile to implement and refine active remediation, working with SDN controllers, Nokia's own Service Routers, and with routers from leading vendors (including Juniper and Cisco, but support exists for any Flowspec or NETCONF router).

An example (and implemented) use case that Nokia uses to demonstrate its approach is that of CDNs dynamically shifting traffic. Dynamic traffic rearrangement by CDNs poses two problems for a network operator, greatly simplified if the source is understood: First, significant traffic will shift ingress points, changing loads and congestion points; Second, these sudden traffic pattern changes are often be mistaken as a DDoS or other nefarious activity. Deepfield can make clear what is happening, allowing manual or automated resolution to take place efficiently.

USE CASES

A Sampling of Leading Deepfield Use Cases / Use Case	Action	Method
CDN visibility	Optimize traffic for best performance or cost-efficiency.	Understand how network traffic is delivered via peering/transit links. Integration with SDN controller
Subscriber insights	Understand service experience and customer / segment experience to drive both commercial and network actions.	Not a traditional closed loop – informs marketing, network and service planning, pricing, traffic / network grooming decisions.
Insight-Driven Network Security (DDoS Prevention and Mitigation)	Real-time detection of network anomalies and DDoS attacks with auto-mitigation (zero-touch auto-provisioning of mitigation).	Integration with advanced routers (Nokia and 3rd party) to both identify and act on malicious traffic. Place filters and if supported, routing actions (using SDN methods)[2]

COMPETITORS

The analytics industry has always been driven by data – more and better data wins. Traditionally this has led to players that are defined by how they acquire data, from where, and what types. On the one hand were analytics platform vendors who could take data from "anywhere", but most often pulled it from sources they either controlled or understood well, often OSS and BSS. Typically, successful players also concentrated on areas in which they had domain expertise (this is equally true for the branch described below). Leaders in this segment include Ericsson, Nokia Software business group (Nokia's software line of business - as distinct from Deepfield), Amdocs and others.

A second branch began as a critical data source, whether fault collection, or most often, specialty probes on signaling and bearer plane data, and subsequently grew "up the value chain" into analytics – but clearly leveraging their unique data or data domain. Many probe and analytics suppliers offer the ability to identify applications, and consequently to understand

[2] Generally, requires both SDN / flow control capabilities as well as a next generation chipset. Nokia's poster child is their own FP4 chipset, but others have competing offerings available or in the works.

application–level performance and impact on the customer experience. Some of the leaders in this segment include Viavi, IBM (again, distinct from Watson), Spirent, Exfo, Huawei, Allot, Cisco, and Rhode & Schwartz.

Recently pure analytics platform players have been rising rapidly. These range from the giants (Microsoft) to focused players (e.g.: Splunk). The rise of these is consistent with the thrust of "data wins" – these players are leveraging the recent feasibility of huge datasets – both the availability of data and the feasibility of storing it and crunching it cost effectively, often utilizing elastic cloud resources.

Deepfield leverages all three methods, while using a unique approach. It is built on its patented Genome (unique data source); it leverages the greater availability of fast, rich, real-time streams directly from network elements, and yet it is fundamentally a data-agnostic analytics platform. Like all successful players so far, it also concentrates on its domain expertise in telecom WAN traffic – along with methods to collect, analyze, integrate and take automated actions.

We discussed the market dynamics in this research note.

Finally, Deepfield is heavily invested in security – both insight and mitigation. Here it competes with a range of players, from virtual firewall firms, to those vending DDoS scrubbers, Arbor Networks[3] (now part of NetScout) plus competing (if not analogous) solutions from Juniper, Cisco and Huawei.

APPLEDORE RESEARCH GROUP ANALYSIS

Appledore Research strongly believes that virtualization technology (at both the NF and service layers), combined with advancements in operations technology will, and must, lead to networks that are largely automated. This automation yields many benefits, from vastly lower costs, to better capital utilization, to pro-active healing and improved customer experiences.

We believe these many changes will bring about new realities in network operations, specifically:

1. The majority of data collection will shift over time from probes to native interfaces
2. The capabilities of routers to be programmed and perform complex routing functions on pattern matched data subsets – such as terminate DDoS data – will grow.
3. Automation will become the norm and become better trusted allowing real-time remediation of myriad ills from congestion to security threats.

[3] The previous venture of Deepfield's founder, Dr. Craig Labovitz

4. Data sources will increase in richness and in diversity – creating a separation between analytics and data sources (including NEs, probes, signaling, and management systems)
5. Security will increase as a threat but become embedded in operations rather than an add one. In other words, “FCAPS” (fault, configuration, accounting, performance, *security*) will finally earn its “S”.

Deepfield is not alone but exploits each of these trends.

It is worth mentioning that Nokia, as a major, end-to-end network and management software supplier, is among the few firms with the ability to effect closed-loop automation solutions in-house. The solutions noted above, work in conjunction with Nokia SR series routers and Network Services Platform SDN software, to not only detect, but automatically correct (cleanse, block, re-route) traffic.

The Deepfield approach provides strong insights to guide automated operations, and is consistent with Appledore Research’s forward-looking “RASA” (Rapid, Automated Service Assurance) architecture in two important ways:

- First: it breaks down silos, integrating multiple technologies, data streams, and service domains into a single query-able environment with both southbound and northbound APIs. This effects a breaking down of narrow silos and their associated integration costs and myopia.
- Second: it is intended to be a common source to feed many other processes; but especially to provide the necessary insight to guide automated scaling, re-arrangement, attack remediation, optimization and healing – among other possibilities.

Insight defines the problem to solve and sets boundaries/constraints. Other systems must formulate a specific action plan and execute it (bearer plan, re-orchestration, etc.) – hopefully by implementing pre-defined “intent” -- solutions to problems. An SDN controller, for instance, can re-balance traffic to eliminate congestion, or prioritize more critical and high value traffic. Similarly, DDoS traffic could be routed to a scrubbing location or in some cases blocked at network ingress, since offending sources are identified by Cloud Genome. See Figure 2 above.

We also believe that Deepfield Genome is a unique and highly valuable asset. DPI has been a highly useful tool for decades, yet it is held back by its fundamentals: it is costly and its cost scales almost linearly with throughput. Deepfield Genome, taking a different approach, scales (with respect to both cost and performance) much better, allowing many of the same benefits.

With regard to use cases, we believe that Deepfield is focused on those that are both important to CSPs and exploit its strengths. It is difficult to manage raw traffic without insight into what applications are being impacted, what users are being impacted, and/or what the original sources are (affecting both peering decisions and security analysis). With the rising profile of security breaches, Deepfield's concentration on various automated security remediation methods is targeted at the triple crown of risk, savings, and complexity.

In summary, we feel that comparing Deepfield with today's solutions, and today's environment in some ways misses the point. Rather Deepfield should be evaluated in the context of the major technology shifts discussed above, and how it fits into that real-time, big data, automated and dynamic world. In that way Deepfield can not only solve emerging problems, but possibly be part of driving the transformation.

SUMMARY AND RECOMMENDATIONS

One of Nokia's marketing and positioning challenges with Deepfield is to simultaneously communicate specific use cases, and also the fundamental breadth of its capabilities. A continual challenge in our industry is that customers buy use cases (specific, tangible solutions) not platforms. And yet, agility, efficiency, and integration cost reduction all depend on elegant, de-siloed platform architectures that span not just the scope of assurance/analytics, but the entire next-generation management environment and its many closed-loops. We encourage readers to refer back to our many Market Outlook reports that delve into best practices across the many parts of a proper control loop (links provided at end).

Appledore believes that Deepfield is an innovative and important approach, with the following summary characteristics; specifically, Deepfield:

1. Supports all the characteristics that Appledore Research suggests for closed-loop, real-time, automated assurance ("RASA").
2. Is built to handle today's high volumes of real-time streaming telemetry data
3. Is application aware, and source/thing aware – which is critical to understand the importance, who is impacted, what is impacted and whether a flow may be nefarious.
4. Collects directly from data sources rather than intermediate probes or other devices, reducing cost and more importantly, simplifying data collection and configuration in dynamic (virtual NFs, virtual flows) networks.
5. Addresses three of the major issues in data networks today: 1) DDoS/security, 2) Customer and application experience, and 3) efficient network utilization.

Deepfield Genome is truly a differentiating technology, allowing Deepfield to identify both applications, and source devices. Moreover, it allows applications to be identified with less cost (compared to DPI), and for more applications to be identified. Deepfield Genome leverages this to three core benefits (put in our terms):

1. Security
2. Traffic / network optimization

3. Customer and service insights

In each case, it is imperative to see both the "forest" of traffic volumes and the individual trees that make up that forest. Absent that data, one cannot identify DDoS or malicious traffic, nor can one understand what trends are driving capital expansion, customer satisfaction, or service quality and QoE. And these are the 'money" questions.

We urge both Nokia and CSPs who deploy Deepfield to embrace automation – not simply "flow-through", but truly hands-off handling of traffic optimization, user experience optimization, and security remediation. The results have the potential to lower costs and improve customer QoE at the same time – a happy result. Without automation there remains a large lag between the advent of "trouble" and its resolution. Moreover, the sheer number of issues generally overwhelms manually-assisted actions, meaning on those best understood or deemed most important get acted upon. With proper automation (and continuous learning and refinement) all can be addressed, and the number of false positives and negatives can fall below the already low levels achieved today.

FURTHER READING & BIBLIOGRAPHY

http://appledoreresearch.com/product/best-practices-orchestration-focus-policy-models-boarding/

http://appledoreresearch.com/product/rapid-automated-service-assurance-nfv-sdn-network/

http://appledoreresearch.com/product/closed-loop-automation-virtualized-networks/

http://appledoreresearch.com/product/nokia-routing-chipset-new-routers-new-operations-methods-research-note/

https://appledoreresearch.com/report/why-intent-matters-the-economic-case-for-the-intent-based-network/

https://appledoreresearch.com/report/sd-wan-drives-new-csp-enterprise-opportunity/

https://appledoreresearch.com/report/nokia-routing-chipset-new-routers-new-operations-methods-research-note/

E: info@appledorerg.com

www.appledoreresearch.com

44 Summer Street Dover, NH. 03820, USA