



GCHQ Analytic Cloud Challenges

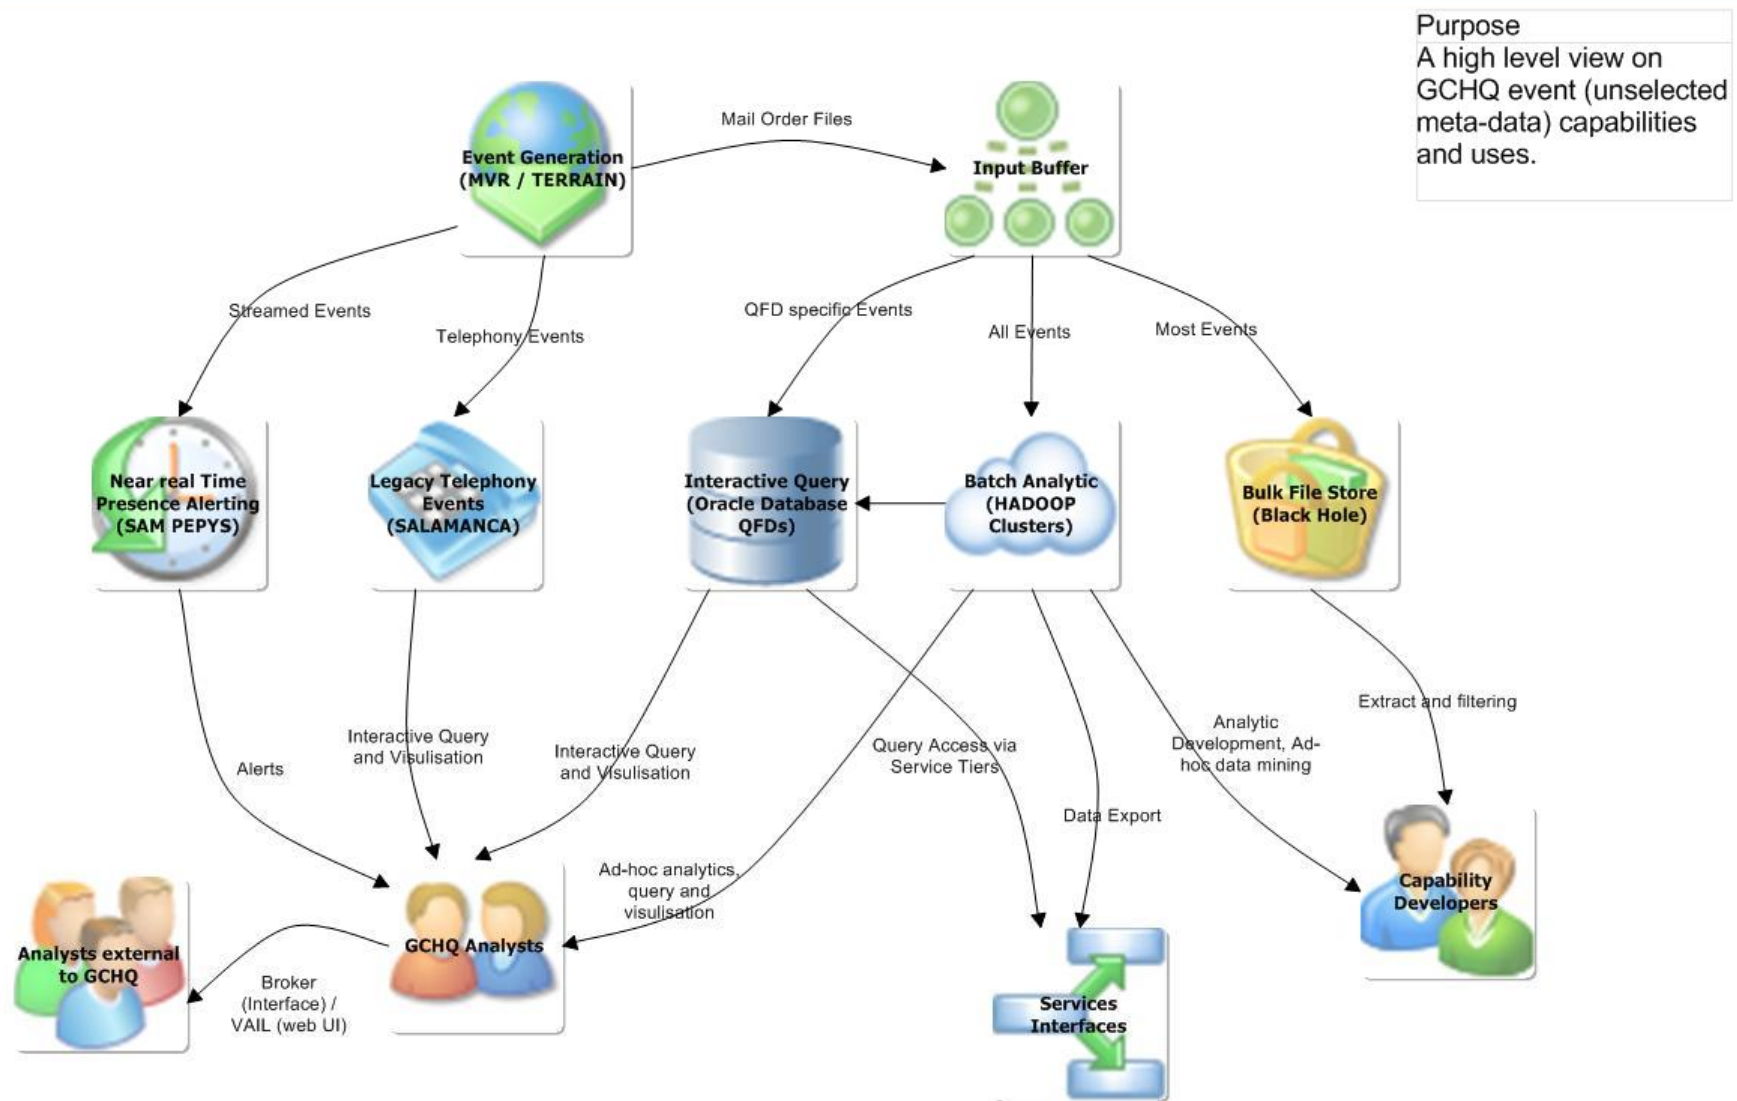
11/11/2016

Approximately 50 Billion Events Per Day from 250+ x10G bearers



- Scale Target for May 2012 is approximately 320 x 10G with a capacity for 100 BEPD
- Typical retention is six months, sometimes reduced to cope with volume increase

TDB Events PC OV-1 - Events High Level Concept Graphic



Query Focused Dataset



22U Rack

What are QFDs

- A database designed for answering a single (or few) analytic question
- Initially adopted from research, engineering tasked to mature and scale ASAP
- The engineering has been incremental based on priority
 - Library of components and design patterns developed
 - Scale using appropriately sized database instances, federation using service tiers
 - Deployment of new instances and updates streamlined
 - Corporate support
- New QFDs still being produced by engineering and others using the standard components
- Some now provided interactive access to datasets generated by HADOOPS batch analytics

QFD Facts and Figures

- Approximately 100 QFD instances deployed for 16 different QFDs + more for UI etc
- Each with events (or results) appropriate for the questions they answer
- Hardware shared depending on QFD, all storage is shared
- Driven by the need for a flexible platform at large scale with low overall cost
- Redhat Enterprise Linux, Oracle DB (DB needs are simple)
- Latest generation HP BL456c G7 blades, 24 core, 64GB RAM, EVA storage ~200TB usable
- Most QFD instances have 70TB usable capacity
- Total storage 15PB raw, 11PB usable

A selection of the major QFDs

Name	Description	Questions Answered	Physical Bude+Chelt
AUTOASSOC	Bulk unselected TDI-TDI correlations with confidence scores.	What other TDIs belong to your target ? What technologies your target is using ?	2+1 instances, each 50-70TB storage
Evolved Mutant Broth	Identify when certain TDIs appear in traffic which indicate target usage and their location. Telephony and C2C data provide a converged view.	Where has my target been? What kind of communications devices has my target been using?	10+5 instances, each 70TB storage
Hard Assoc	Provide strongly correlated selectors for both C2C and Telephony traffic taken from TDIs appearing in the same packet	Are there any alternative C2C or Telephony selectors for my target?	3+2 instances, each 70TB storage
HRMap	Host-referrer relationships - information about how people get to websites, including links followed and direct accesses.	How do people get to my website of interest and where do they go to next? What websites have been visited from a given IP?	5+3 instances, each 70TB storage
KARMA POLICE	Which TDIs have been seen at approximately the same time, and from the same computer, as visits to websites.	Which websites your target visits, and when/where those visits occurred. Who visits suspicious websites, and when/where those visits occurred. Which other websites are visited by people who visit a suspicious website. Which IP address and web browser were being used by your target when they visited a website.	11+7 instances, each 70TB storage, 3+1 correlator instances
SOCIAL ANTHROPOID	Converged comms events allowing you to see who your targets have communicated with via phone, over the internet, or using converged channels (e.g. sending emails from a phone or making voice calls over the internet).	What communications your target is engaged in. Who has your target been communicating with. What communications have occurred using a particular locator (IP address, cell tower, etc).	6+3 instances, each 70TB storage

HADOOP Clusters

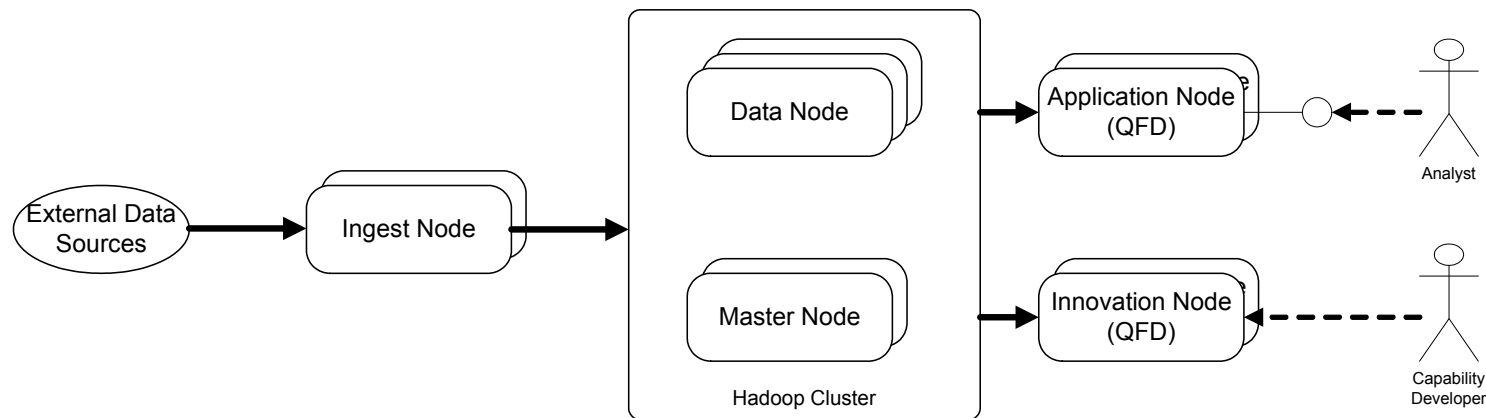
Facts & Figures

- Three batch analytic clusters (HADOOP clouds)
 - Each contains unique, unselected events, reference and results data
 - Redhat Enterprise Linux, Apache HADOOP (Map Reduce)
 - Each with ~900 data nodes, plus UI/ingest/egress servers
 - Each node 8 core, 64GB RAM, 6x1TB
 - Each 6PB raw storage, 1.5PB for events (3x replication and results)
- The three clouds could, in theory, store 24 Trillion events @140 BEPD and 6 months retention period

Usage

- 200 capability developer accounts
- Many more users able to access applications
 - Guiding light > 150 users
 - First Contact, Cloudy Cobra > 50
- Jobs during working hours
 - Operationally ad-hoc (experimental, search, one time use)
 - Development of new capability
- Schedules Jobs run overnight
 - Sustained, operational applications
- Increasing expectation for fully supported, automated jobs such as Rumour Mill

Physical Overview



Private Cloud

- Data Nodes
- HADOOP Master Nodes (Job Track & Name Node)

Edge Nodes connect to Computer Hall Network

- Ingest Nodes
- Application Nodes (provide user interface – web)
- Innovation Nodes (cap dev login to servers)

Data Ingest

Ingest Design Goals

- Store all data received, i.e. don't risk losing important fields by normalising to single file format.
- Partition by directory structure
 - Partition by data type to simplify common queries.
 - Partition by ingest date to allow incremental analytics.
 - Partition by security compartment.
- Horizontal scaling by adding new hardware.

Data Formats

- TLV – Tag Length Value
- Comma Delimited
- Fix Position
- Single Line & Multi-Line
- Multiple Variations
- Actor Action
- Many sources and many types

Analytic use of data

Security

- Data access controlled across 2 dimensions:
 - Usability: Operational (standard events) or Controlled (special purpose e.g. test)
 - Classification: Multiple buckets - Open (TSS2), a few compartments, general CIOs Sensitive
- Capability developer access restricted by file/directory permissions
- User facing applications apply fine grained security on per record basis

Silver Library

- Abstraction layer to separate analytic development from storage dependencies
 - Pluggable parsers for new formats
- Record parsers for event data held in HDFS
 - Applications still need to know field labels (function calls)
 - Applications often add additional abstractions
- Input and Output handlers to read and write data based on type
- Utility functions to aid development
- Simple search and filter routines
- Events Product Centre maintained

Example Analytics

Every Assoc

- User /machine correlations from C2C presence

Every Police

- User /machine website visits

Every Creature

- User /machine search terms

Every eAD

- User /machine electronic attack patterns

Every Cipher

- User /machine cipher events

Grey Fox / Silver Fox

- Country level summary of where identifier observed

First Contact

- 1 & 2 hop contact chains between seeds and targets

Cloudy Cobra

- Glorified grep driven by GUI – find events that contain user search term

Guiding Light

- MI information types/volumes of traffic on bearer

Golden Axe

- Generated list of suspected clone mobile phones (IMEI grey list)

Tribal Carnem

- Uses Radius logs to identify & collect activity for IP session

Public Anemone

- Geolocation based on web-based map searches

Epic Fail

- Identifies careless use of TOR networks

Sterling Moth

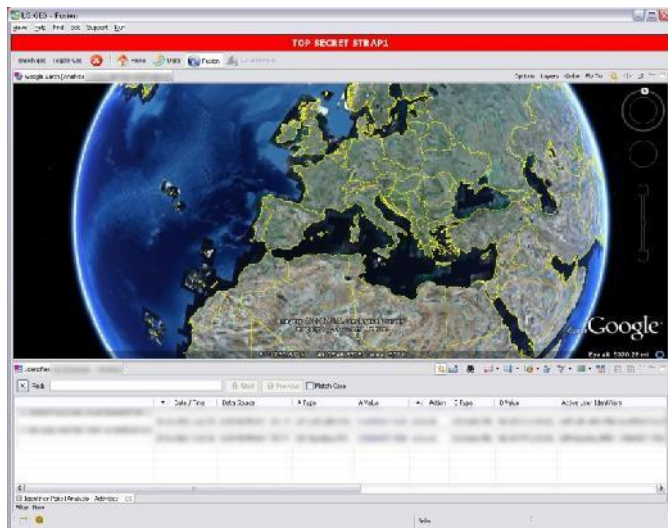
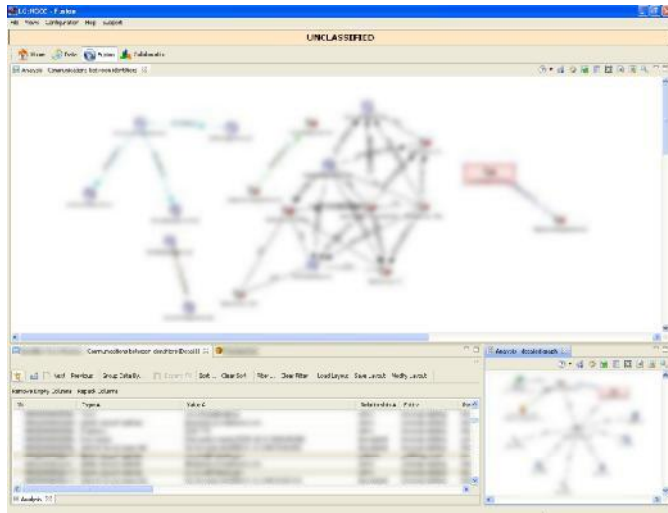
- IP summarisation tool using c2C presence events

Foghorn

- Find non-targets using targets machines

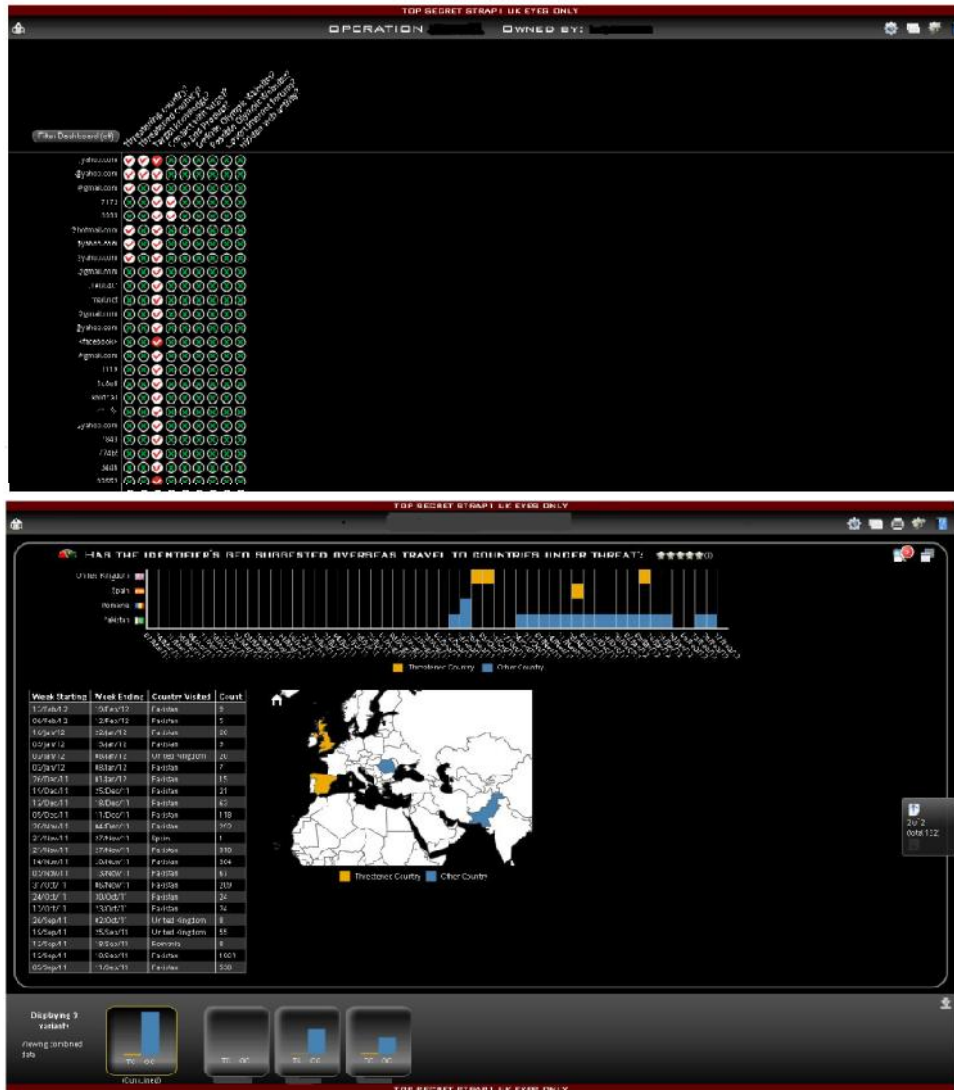
And many more....

Rich Visualisation



- Desktop applications for rich visualisation and analysis - Eclipse RCP
- LOOKING GLASS client platform, FIRE ENGINE question-based federated access to events and reference data sources
- Pilots
 - NGCC Contact chaining with C2C/converged event data
 - QFD Federator
 - NG Geo: Converged geo events analysis
- Next stage to learn from pilots and build the tool for everyone - ALPHA CENTAURI

Rumour Mill – Push Analytics



Sharing & Collaboration

Other SIA and foreign partners

Data (bulk & query) and technology exchange

Two major components:

- Web user interfaces (VAIL) on GCHQ servers but accessible from the partner site. Interactive query of QFDs. Allow exposure of GCHQ tradecraft.
- Brokering services. Sustained access for interactive query of GCHQ data integrated into partner tools.

Future Options for Event Processing

Key Challenges

Affordable, continuing scale of our capabilities

- All dimensions, cost, power, cooling, space, storage, bandwidth, processing etc
- As we share with more partners, the access demands increase
- Enable more delivery by collaborating with others

Enabling our analysts to cope with complexity and volume of data

- Need to streamline understood workflows to save analysts time
- Need to push data to analysts, sometimes with low latency
- Enable next level analysis, behavioural, pattern of life, noticing change, predictions

Complexity and pace of analytic development

- Suitably skilled people are hard to find in-house and externally
- Must be able to cope with mixed maturity capabilities
- Workload separation, processing contention, cross site analytics

Agility, resilience and maturity of platforms

- Keep pace with development, respond to community demands
- Understand and match the evolving business expectations of maturity
- Support appropriately, agreed strategy for resilience

Approach to challenges

Deploy more of the same while maturing

- Suitably refreshed, will answer immediate scaling challenge
- Reduce support burden and streamline new capability development
- Implement improvements to existing architecture

Explore new technology

- Non-relational, distributed databases; QFD consolidation & convergence
- Research initiatives (A)SEM/HAKIM/STREAMING, and commercial offerings
- Collaboration opportunities CLOUDBASE/ACCUMULO

Know the value of our data

- Don't generate/ingest data we don't use, filter or de-dupe upstream
- Be smart about retention, be smart about need for interactive availability
- Distil the raw data to generate rich information sets for analysts

Understand usage of capabilities

- Automate the simple workflows with summaries and push analytics
- Optimise capabilities for their use and access load
- Mature, provide resilience and support as appropriate

Stay flexible and increase collaboration

- Enable collaboration and lower the entry bar for capability developers
- Stay vendor neutral, continue to use open standards, open software
- Enhance the benefit for both research & engineering of working together

Understand usage of capabilities

Use Case Class Mapping



- Developed jointly between GCHQ and NSA to understand
 - the benefits of our current capabilities
 - where respective strengths and weaknesses exist
- Provides a clear set of drivers for architectural evolution
 - Missing capabilities
 - Suboptimal use of capabilities
 - Opportunities for collaboration and reuse

Existing coverage of use cases

Known Target, Known Query

- Streaming, Rumour Mill and QFD capabilities (possibly assisted by cloud analytic)
- Could select data subset based on target, could pre-calculate results
- Large usage and automation suggests need for optimised capability

Known Target, Unknown Query

- Not possible with existing QFD capabilities
- Needs new analytic or more indexes on target selected data
- Large usage and needs suggest a new capability

Unknown Target, Known Query

- Core QFD capability, possibly populated by cloud analytic
- Full unselected data set required but what level of interactive query?
- Is it acceptable that older data be made available non-interactively?

Unknown Target, Unknown Query

- Use batch analytic platform for low level search or new analytic
- Could a heavily indexed store provide a responsive capability?
 - Possibly only over a subset of data - recent data only?

Value assessment

What measures of value?

- Duplicated over time or across bearers
- Value decreases with age
- Value decreases after processing (SPAM, summaries)
- Use by analysts and analytics

What actions to take?

- Do filtering upstream or don't generated
- Investigate use of streaming capabilities for filtering or sampling
- Discard immediately after first stage processing
- Age off and discard more selectively
- Periodically evaluate data types to understand usage

Need to agree possibilities with operations and experiment

Technology candidates

QFD consolidation and HAKIM

Existing QFD drawbacks

- Some duplication between QFDs (~10%)
- Need new QFDs to answer new questions

Need a consolidated database with multiple indexes and flexible additions

- HAKIM is a research prototype to do just that
 - Unification of data, associated data kept together
 - Quick and flexible addition of new data types and indexes
 - Scalable and cost effective

Other candidates exist and some HAKIM components could be replaced

- Oracle DB or distributed, non-relational database?
- Convergence with HADOOP stack HBASE/ACCUMULO

Engineering is working with research to develop to the next stage

Technology candidates

The “skinny” cloud (Laurel)

What is it?

- A batch analytic HADOOP cluster
- Contains all the data but for a short retention period

What would it be used for?

- An ideal place to do incremental analytics
- Answers the cross-site analytic problem
- Could be dedicated to sustained usage
- Provides some resilience by duplicating recent data

What are the challenges?

- Can we transfer and ingest all the data

Engineering could potential build this using on-order kit

Technology candidates

The GCHQ “core” (Hardy)

What is it?

- A HADOOP cluster with Map/Reduce and interactive query/analytics capabilities
- Probably using CLOUDBASE/ACCUMULO and reusing NSA knowledge

What would it be used for?

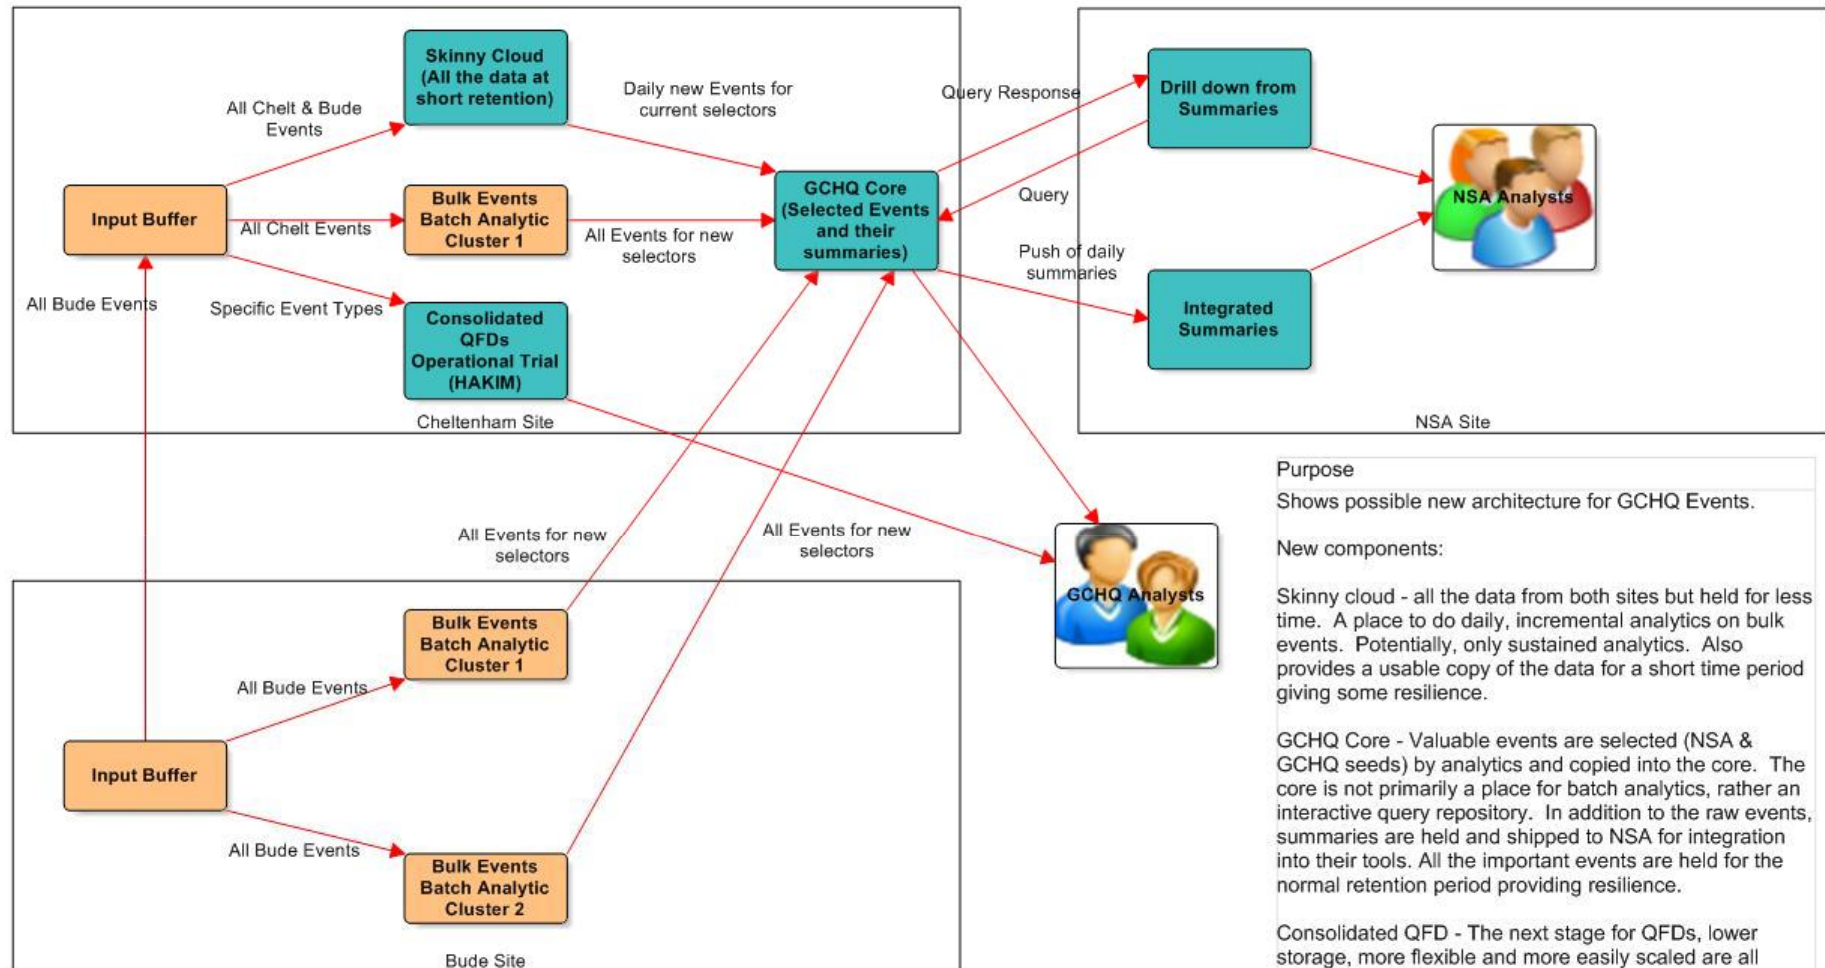
- A place for data and summaries promoted from the bulk stores
- Known Target, Know Query and some Known Target, Unknown Query
- Optimised for major use case and suitable for data sharing
- Provides some resilience by duplicating important data

What are the challenges?

- GCHQ has limited expertise with CLOUDBASE/ACCUMULO this technology
- The promotion analytics and criteria are not developed

Engineering could potential build this using on-order kit

TDB Events PC SV-1 - Skinny cloud, GCHQ Core and QFD consolidation



Purpose

Shows possible new architecture for GCHQ Events.

New components:

Skippy cloud - all the data from both sites but held for less time. A place to do daily, incremental analytics on bulk events. Potentially, only sustained analytics. Also provides a usable copy of the data for a short time period giving some resilience.

GCHQ Core - Valuable events are selected (NSA & GCHQ seeds) by analytics and copied into the core. The core is not primarily a place for batch analytics, rather an interactive query repository. In addition to the raw events, summaries are held and shipped to NSA for integration into their tools. All the important events are held for the normal retention period providing resilience.

Consolidated QFD - The next stage for QFDs, lower storage, more flexible and more easily scaled are all possible benefits.