Extended Metric Support for WF/EAP

Heiko Braun <hbraun@redhat.com>

Jan 2014

Further Information

- The full proposal with more in-depth information can be found here:

http://goo.gl/IxTkLs

Tag Line

- Storing and retrieving time series data
- Events
- Metrics

Use Cases

- View System Health
- Identify Hotspots
- Correlate application usage and resource allocation
- Identify trends and changes

Ingestion

- Get the data into the system
  - HTTP, UDP
- Normalization
- Collectors

Events/Metrics

- Type, Timestamp, Value
- Additional Meta Data:
  - i.e. Host, Tenant, Server, etc
- Measurements and Probes/Traits:
  - Numerical values (dynamic, trending)
  - Strings, Booleans (Resource state)

Querying

- Retrieve data to make sense of it
  - Get to the data: HTTP
  - Separate from making use of it (i.e. visualization)

Aggregation

- Apply aggregation methods to large amounts of events:
  - Min, max, avg, count, etc
  - Basically any reducer, combinator

Aggregation Strategy

- When to apply aggregation methods?
- Write time:
  - Precomputed results, higher resolution:
    i.e. “last 12 hours, 15 min slice, Avg, Count from t1 to t2”
- Read time:
  - Supports tailored, custom queries:
    i.e. “Request to `/path` from `192.168.0.12`, `UA Chrome` from t1 to t2”
  - Query language support:
    `“sum(requests).from(t1).to(t2).eq(path, "/account"))”`

Roll up's, Retention

- Roll Up:
  From fine grained to coarse grained data sets
  - Pyramid aggregation
- Retention: Expiration of data
  - Along the roll up strategy

Data Storage

- Column Oriented
- Ranged queries
- Wide rows
- Write vs. Read Performance
- Multi-node support

Non-Functional Req's

- Grows with users needs
  - Single server, multiple nodes
    - Collocated with applications or separate VM's
- Break up according to service utilization:
  - i.e. some Ingestion Nodes, few query nodes, many rollup nodes

Big Picture

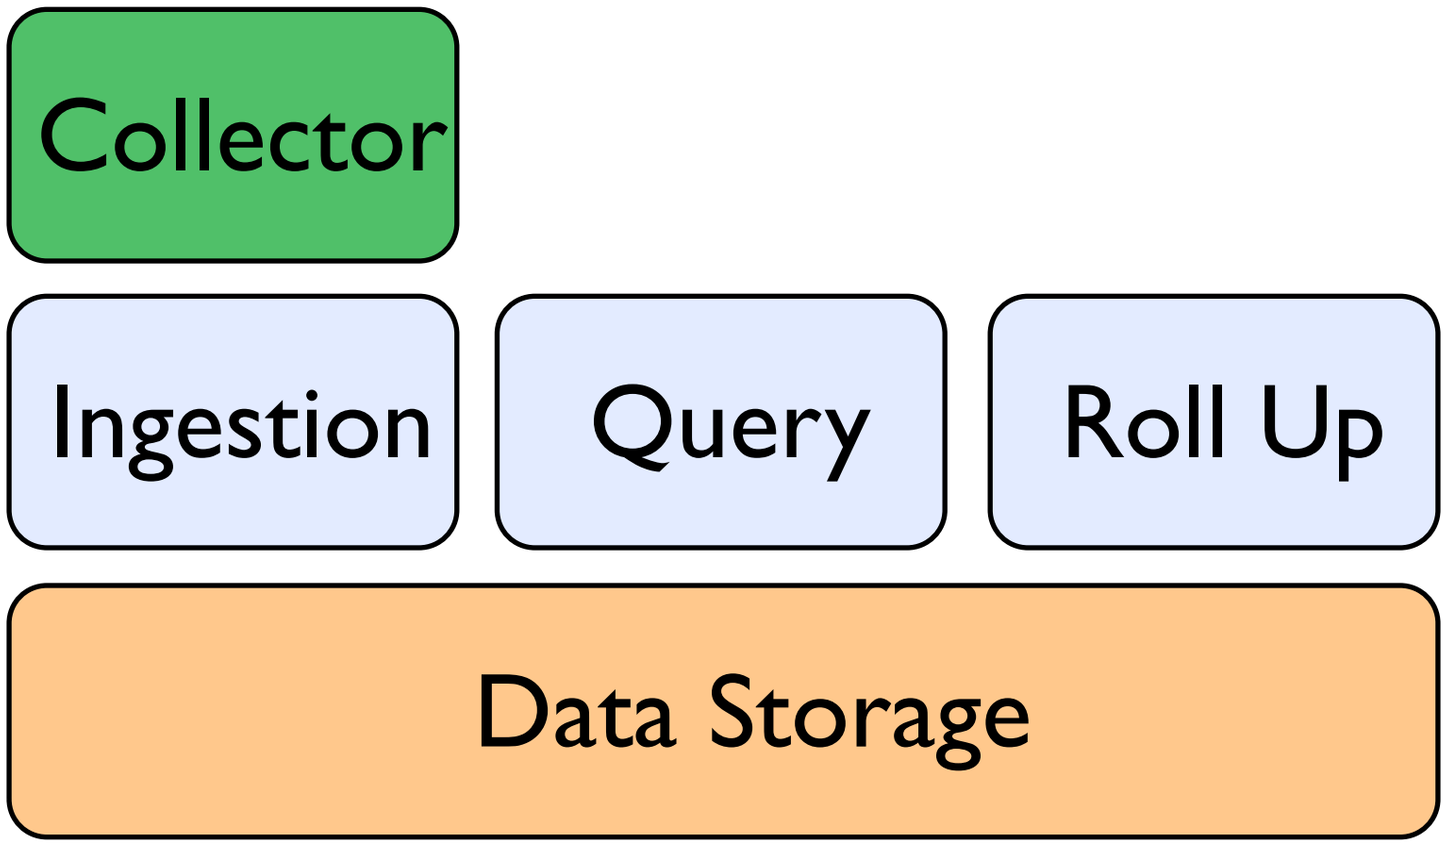

Architecture

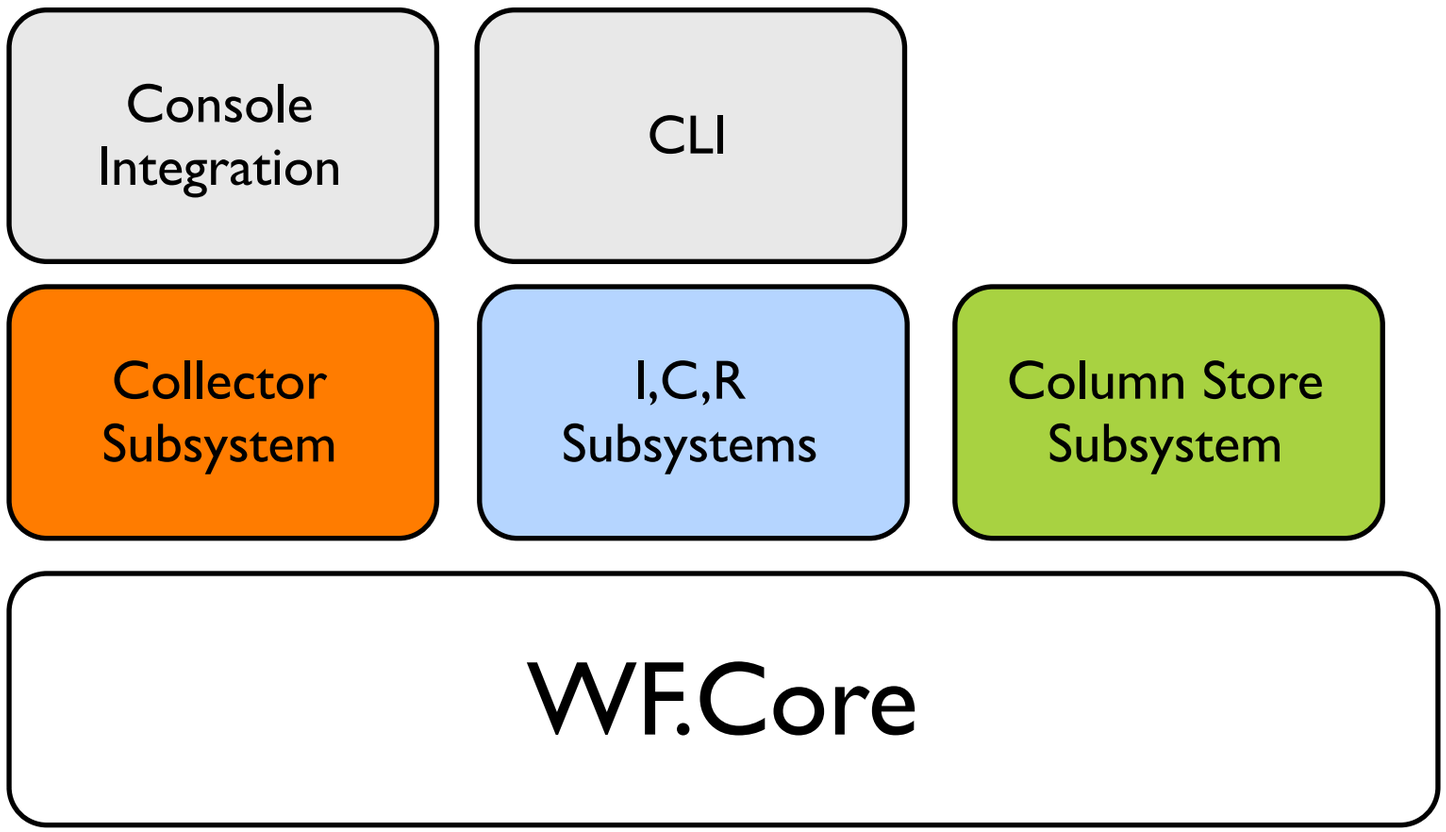

EAP 7 Alignment

- Can be used to provide:
  - LogViewer ✓
  - Extended Metrics ✓
  - NoSQL Support ✓