



vivit

Discover the independent Micro Focus software user community where you can share, collaborate, exchange, and grow



COMMUNITY



SIG Talk IT Operations: Mainframe to UD, Data Analysis with OpsBridge, Vertica into BVD
June 21, 2018



Welcome to



YOUR INDEPENDENT MICRO FOCUS SOFTWARE USER COMMUNITY



Hosted By



Rocky Pisto
Vivit Engage Coordinator
Vivit Worldwide

YOUR INDEPENDENT MICRO FOCUS SOFTWARE USER COMMUNITY



Today's Presenters



Chip Sutton
CTO
EView Technology



Brian Bowden
Senior Engineer
Greenlight Group



Jay Batson
Solution Architect
Greenlight Group

YOUR INDEPENDENT MICRO FOCUS SOFTWARE USER COMMUNITY



Tweet Your Questions



#vivitsigtalk

YOUR INDEPENDENT MICRO FOCUS SOFTWARE USER COMMUNITY



House Keeping

- This “LIVE” session is being recorded
Recordings are available to all Vivit members
- Session Q&A:
Please type questions in the Questions Pane

YOUR INDEPENDENT MICRO FOCUS SOFTWARE USER COMMUNITY



Webinar Control Panel

The screenshot shows the Vivit Webinar Control Panel. At the top left, there is a Vivit logo and a small thumbnail image. Below the logo, there are two circular icons: one with a download symbol and one with a gear symbol. A blue arrow points from the text "Download Handouts" to the download icon. The main area of the panel is a large white rectangle with the text "No presentation shared" centered inside. At the bottom of this area, there is a blue arrow pointing down to the text "Phone Number Options". At the top right, there are three red circular icons: a question mark, a document, and an information icon. To their right is a language dropdown menu currently set to "ENGLISH". A blue arrow points from the text "Choose Language" to this dropdown. Below the icons and language menu is a vertical sidebar. The top of the sidebar has a Vivit logo and thumbnail. Below that is a section titled "Questions" with a gear icon. At the bottom of the sidebar is a text input field with the placeholder "ask your question" and a right-pointing arrow. A blue arrow points from the text "Questions" to this input field. At the very bottom of the panel, there is a footer with the text: "DIAL IN 415-926-7795 MEETING ID: 238985 USER ID: 6982 INTERNATIONAL NUMBERS".

Choose Language

Questions





How one of the world's largest insurance companies brought their IBM environments into the modern era by extending their Micro Focus Operations Manager i and Universal Discovery (UD) to include the mission-critical IBM systems.

Customer

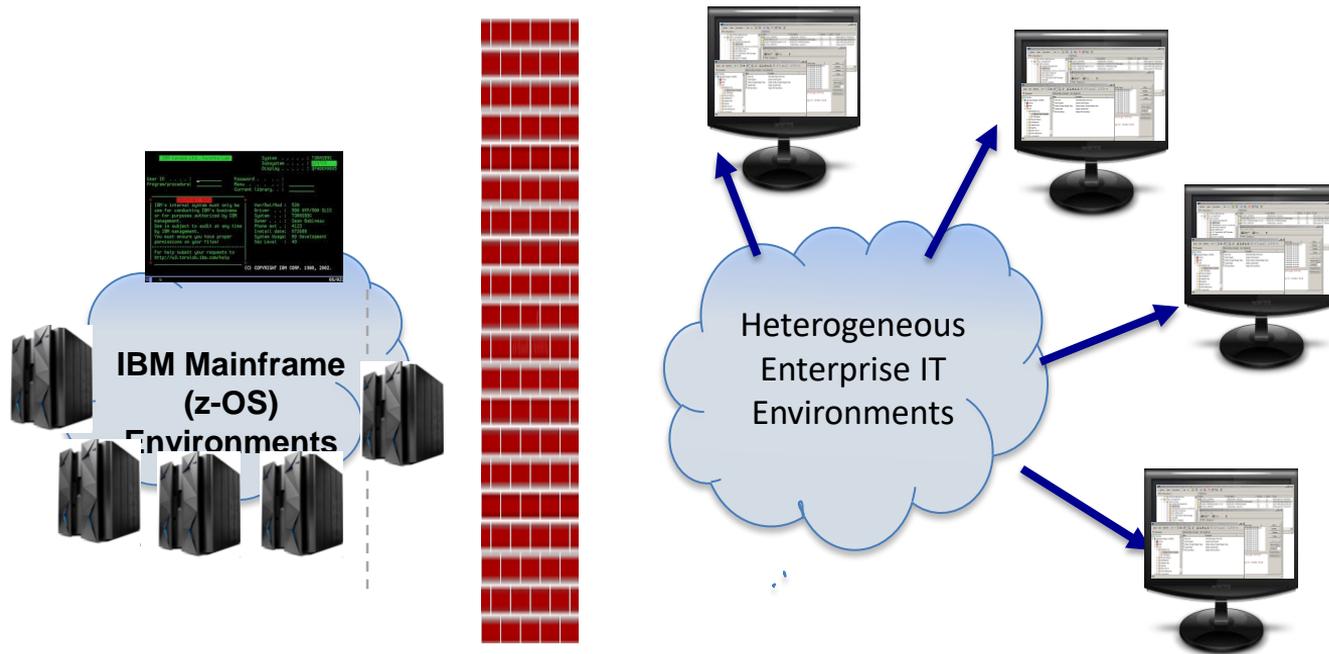
- Large insurance and financial services corporation based in the United States. More than 65,000 employees and more than 18,000 independent contractor agents which service over 83 million policies and accounts throughout the U.S.
- They have 30 operations centers throughout the US, which contain thousands of various types of servers (including over 15 IBM mainframe computers).

Issues

- IT Operations/Configuration Management growing out of control, with different tools to manage different environments
- minimal integration between the various components
- escalating software costs
- skill level required was becoming increasingly diverse and expensive.
- Discovery and tracking of IT assets and CIs was manual
 - Very time intensive
 - Not reliable
 - Often out of date before the was finished

Goal

- Consolidate its vast operations under a common platform - without a corresponding increase in IT infrastructure and costs.
- Integrate its mission-critical IBM mainframe (z/OS) into this system
- Automate IT Discovery processes
- Get an overall, comprehensive, accurate and complete view of corporate mainframe infrastructure for IT Operations, IT Operations analytics, configuration management, service modeling and service impact analysis.

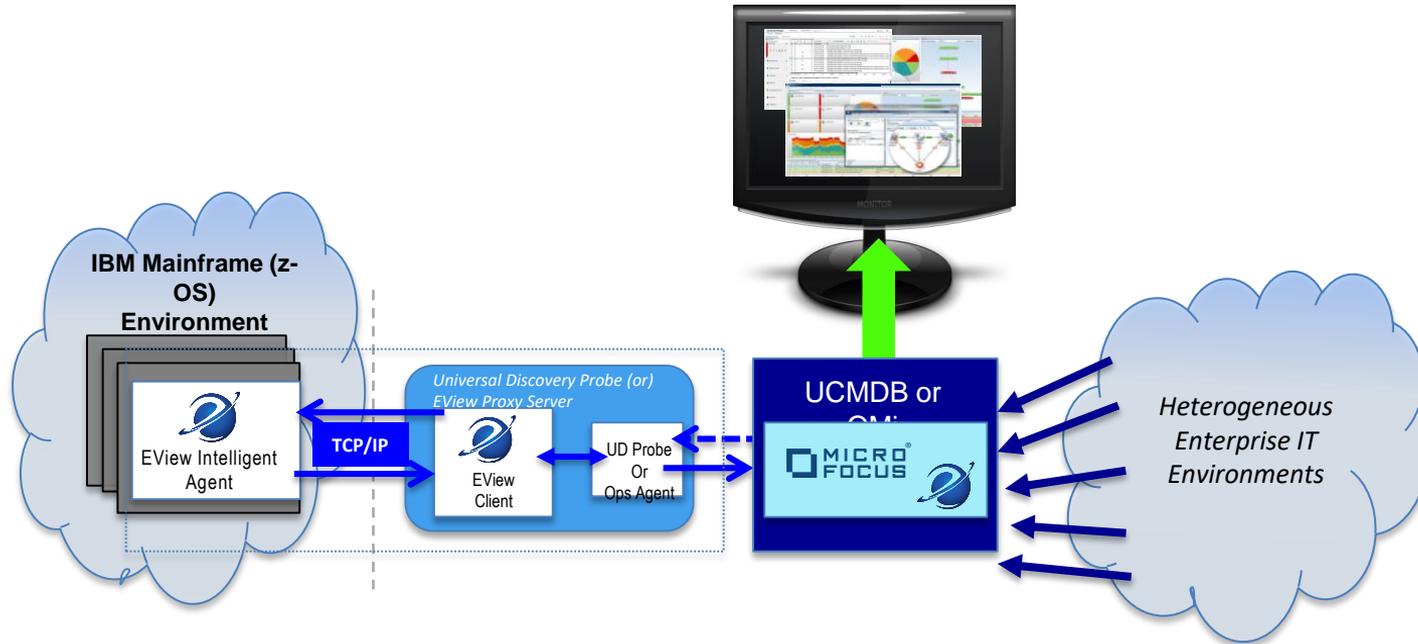


- Little or no visibility into IBM mainframe z/OS environment
- Multiple data sources
- No single source of “truth” of IT Operations Configuration Data

They needed ability to capture the many different CI types which are part of an IBM Mainframe environment

Mainframe (z/OS)

- DB2 - DDF, DSG, databases, table spaces and deep configuration data
- Completed jobs
- DASD Storage and Storage Groups
- CICS - regions, transactions, programs
- IMS - regions, databases, transactions, programs
- MQ (managers, channels, queues)
- Memory
- LPAR
- CPU
- Network Connectivity
- Installed IBM Software



- Seamless Integration with Micro Focus OMi and Universal Discovery - complete end-to-end solution which includes the IBM mainframe/zOS environments.
- Builds on Micro Focus implementations – extending enterprise value to include mission-critical IBM Mainframe data.

- Extends the cross-platform capabilities of the Microfocus ITOM Solutions (UCMDB and Ops Bridge) to include the IBM mainframe-z/OS and iSeries (AS/400) environments.
 - Comprehensive capabilities
 - Seamless integration
 - Consolidated operations
- Real time visibility over the entire environment – 360 degree view.
- Brings corporate IBM legacy environments IT Operations into the modern era.
- The overall cost impact to deploy and maintain EView solutions are predictable and significantly less than a usage based pricing model.
- EView has the valuable, and unique understanding of modern IT Operations platforms and the legacy IBM environment knowledge and support skills.
- Proven Technology – installed in hundreds of customer locations in over 35 countries.

- Creation and optimization of an IT infrastructure that is easier and less expensive to manage, upgrade and run, by having a consolidated, end-to-end management view that includes their mission-critical IBM mainframe environments.
- Immediate return on investment
 - Operational efficiencies gained
 - Elimination of some outdated mainframe tools and licenses
 - Compliance gains
- EView common agent technology
 - Provides a single source of truth for both IT Operations and Universal Discovery



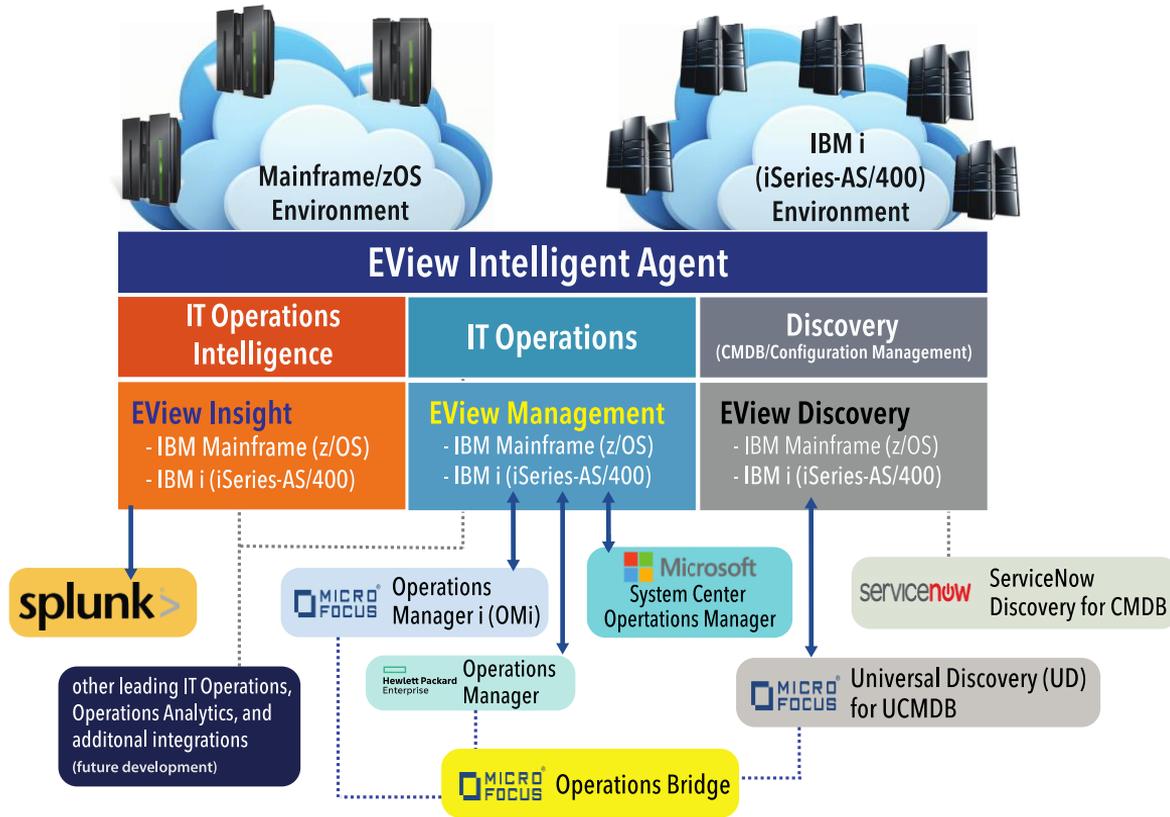
Demo



Gain business insights previously not possible.

Comprehensive and cost-effective solutions which seamlessly integrate, monitor and manage the mission-critical IBM mainframe and IBM i (iSeries-AS/400) environments into today's leading enterprise platforms.

The EView Technology Product Family



Where to go for more information



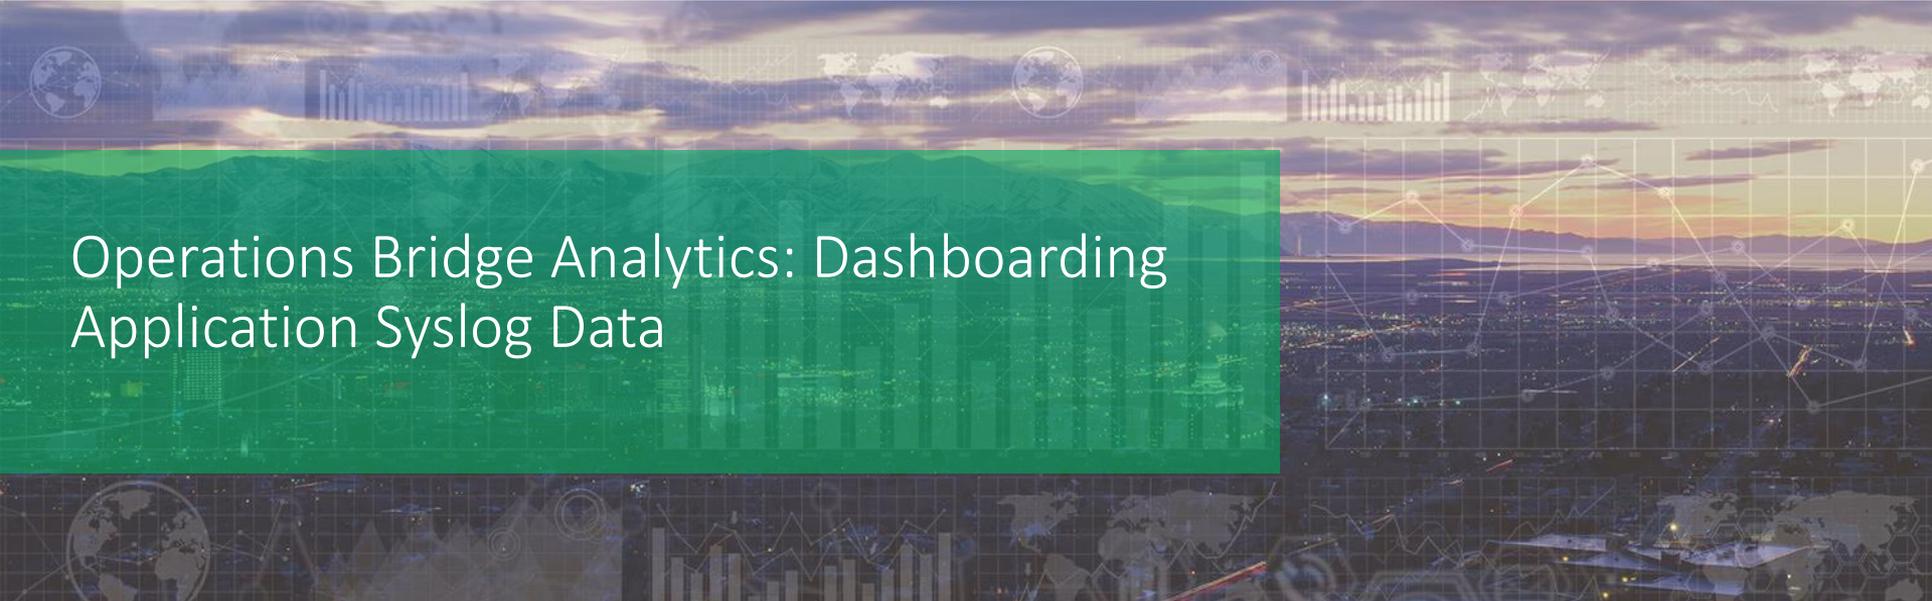
www.eview-tech.com

Or our 24 x 7 General Sales support at:

info@eview-tech.com

Call: +1 919 878 5199





Operations Bridge Analytics: Dashboarding
Application Syslog Data

Brian Bowden – Senior Consulting
Engineer

June 2018

Financial Customer: OBA Use Case



1 | Customer Use Case

- Large Financial Customer
- IBM DataPower data
- No direct integration to OBA
- Need to dashboard Transactions Per Second (TPS), Average Transaction time.
- Log Volume of millions of lines per day (35M in production).
- Required Data is spread across multiple lines of syslog data.

Syslog Data Collection

Documentation:

https://docs.microfocus.com/OMi/10.63/Content/OMi/AdminGuide/Monitoring/PolicyTemplates/om_uc_pt_structured_logfile.html

https://docs.microfocus.com/OMi/10.63/Content/OMi/AdminGuide/Monitoring/log_streaming.htm

- DataPower data forwarded to syslog server. (Modifications made to /etc/rsyslog.conf; port UDP 514 opened)
- MF OA agent 12.04 used for log forwarding to OBA
- MF OMi 10.63 Generic Output from Structured Logfile used to format data for OBA ingestion; Filtering is allowed in 10.63.
- Due to amount of syslog messages being received for each transaction, we filtered on specific variables. 17 lines per transaction was filtered down to 2 lines which contained the URL and transaction Latency.
- ALL lines we selected via Structed Logfile Policy; but only lines with transaction ID variable were forwarded to OBA; this filtering was done on the Log forwarding policy.

Log Forwarding Data & Dashboard Limitation

● OBA Data/ Dashboard Limitations:

- Data from OA log forwarding is inserted in Vertica table `opsa_default.opsa.collection_message`
- AQL and XQL queries can be used to build Dashboards from `opsa_default.opsa.collection_message`; however data needs to be on the same row of the database. No way to join data from two lines using AQL/XQL.
- Customer needs to have specific data from each row joined together to provide need TPS Dashboards.

The Fix:

- Create a script to pull data from both rows of a transaction; join data and insert into another table.

Another Problem discovered:

- Database table was created and data inserted; still no data was received in AQL queries. OBA doesn't allow for AQL queries from tables not created by OBA.

opsa-custom-collection.sh

Documentation:

<https://docs.microfocus.com/OBA/3.03/Content/Collect/customCollections.htm?Highlight=opsa-custom-collection.sh>

/opt/HP/opsa/bin/opsa-custom-collection.sh is used to create a custom Kafka collection table in OBA Vertica that can be AQL\XQL queried for dashboard creation.

#.	Name	Label	Datatype	Datatype	Length	Columnname	Tags	Format	Key	Key seq. number
1.	timestamp	Timestamp	datetime	epoch	0	timestamp		secs	no	
2.	transaction	transaction	float		0	transaction			no	
3.	url	url	string		1065	url			no	
4.	ip	ip	string		1065	ip			no	
5.	hostname	hostname	string		1065	hostname			no	
6.	source	source	string		1065	source			no	
7.	timestamp_orig	timestamp_orig	datetime	date	0	timestamp_orig		yyyy-MM-dd HH:mm:ss	no	
8.	alias	alias	string		1065	alias			no	
9.	transaction_id	transaction_id	integer		0	transaction_id			no	

OPERATIONS BRIDGE ANALYTICS 03.03.2006 New 77A63

» Source Type Manager

Select one of the 17 known data source types below to set connection

Status	Source Type	
✓	BSM RTSM CIs	
✓	Business Process Monitor	
✓	CustomBuiltDatapowerTransaction	N/A
✓	NNM iSPI for Metrics Component	
✓	NNM iSPI for Metrics Interface	
✓	Operations Agent	

Upload Script:

Script used to upload specific column data from `opsa_default.opsa_collection_message` to new custom built table with data received with timestamp within the last 3 minutes. Ran via OMi Scheduled Task policy;
command: `/opt/vertica/bin/vsqli -f "/opt/vertica/bin/db_insert.sql" -U database_user -w database_password`

```
INSERT /*+direct*/ INTO opsa_default.custombuilt_datapower_transaction (transaction_id, hostname, timestamp_utc, ip, source, timestamp_orig, transaction, url, timestamp, alias) select Distinct t.att_int1 as transaction_id, t.hostname as hostname, clock_timestamp() as timestamp_utc, t.att_varchar3 as ip, t.att_varchar4 as source, t.timestamp as timestamp_orig, t.att_int2/1000 as transaction, u.att_varchar1 as url, t.timestamp as timestamp, alias from opsa_default.datapower_ip_alias, opsa_default.opsa_collection_message t inner join opsa_default.opsa_collection_message u on t.att_int1 = u.att_int1 where t.att_varchar3 = datapower_ip_alias.ip_address and u.att_varchar1 is not NULL and t.att_int2 is not NULL and t.timestamp >= TIMESTAMP_ROUND((CLOCK_TIMESTAMP() - INTERVAL '3 minute'), 'MI') and t.timestamp < TIMESTAMP_ROUND((CLOCK_TIMESTAMP() - INTERVAL '2 minute'), 'MI') and t.hostname = 'DataPower';
```

```
INSERT /*+direct*/ INTO opsa_default.custombuilt_datapower_transaction (transaction_id, hostname, timestamp_utc, ip, source, timestamp_orig, transaction, url, timestamp) select Distinct t.att_int1 as transaction_id, t.hostname as hostname, clock_timestamp() as timestamp_utc, t.att_varchar3 as ip, t.att_varchar4 as source, t.timestamp as timestamp_orig, t.att_int2/1000 as transaction, u.att_varchar1 as url, t.timestamp as timestamp from opsa_default.opsa_collection_message t inner join opsa_default.opsa_collection_message u on t.att_int1 = u.att_int1 where (t.att_varchar3 IN (select ip_address from opsa_default.datapower_ip_alias)) is false and u.att_varchar1 is not NULL and t.att_int2 is not NULL and t.timestamp >= TIMESTAMP_ROUND((CLOCK_TIMESTAMP() - INTERVAL '3 minute'), 'MI') and t.timestamp < TIMESTAMP_ROUND((CLOCK_TIMESTAMP() - INTERVAL '2 minute'), 'MI') and t.hostname = 'DataPower';
```

```
update opsa_default.custombuilt_datapower_transaction set alias='Unknown' where alias is null;
```

```
COMMIT;
```

Lessons Learned

1 | Lessons Learned

- OA agent 2GB limitation for logfile size; log forwarding stops with silent fail
 - I am rolling logs with logrotate... <https://www.systutorials.com/docs/linux/man/5-logrotate.conf>
- When creating the custom Kafka collection table do not assign a column with unique data like a transaction number as a key field. OBA will keep track of the key fields unique data causing the entity.index table to grow very large; causing performance problems with OBA.



BVD Integrations

Post Vertica Data to BVD

Jay Batson – Solution Architect

June, 2017

OUR AGENDA

- 1. Vertica Queries
- 2. Connecting to Vertica with Perl
- 3. Posting the data to BVD

Assumptions

- Some scripting knowledge is needed
 - Perl is used in this demo
- Some DB knowledge is needed
 - Vertica is used in this demo
- General knowledge of command line utilities
 - curl is used to post to BVD in this demo

Vertica Query

- Use case: Get data from OBA/Vertica to BVD
 - Most common is getting Operations Agent metric data for BVD dashboards
- Vertica help:
 - <https://my.vertica.com/docs/8.1.x/HTML/index.htm>
- Creating the query
 - DbVisualizer – free download
 - <https://www.dbvis.com/>
 - Has a Vertica driver built in
 - Sample query:
 - `select disk_byte_rate, timestamp, sourceid from oa_sysperf_global where sourceid = 'slcvp-nnmi.greenlightgroup.com' order by timestamp DESC`
- Using the query
 - `select disk_byte_rate, timestamp, sourceid from opsa_default.oa_sysperf_global where sourceid = 'slcvp-nnmi.greenlightgroup.com' order by timestamp DESC`
- Time Selection
 - OBA data is not up to the second – up to the last 1, 2, or 3 minutes
 - `where timestamp > (CLOCK_TIMESTAMP() - INTERVAL '3 minute') and timestamp <= (CLOCK_TIMESTAMP() - INTERVAL '2 minute')`

Perl Script Connect to Vertica

- Connect to the Vertica DB
 - Install the Vertica ODBC driver for your OS
 - <https://my.vertica.com/download/vertica/client-drivers/>
 - Create the DSN: Windows or Linux (we used Windows)
 - Perl uses the system DSN to connect to Vertica
- Perl script

```
use DBI;
use Data::Dumper;
$dbh = DBI->connect('dbi:ODBC:VerticaDSN','userid','password');
unless(defined $dbh){
    die "Failed to connect: $DBI::errstr";}
$stmt = $dbh->prepare("my query here");
$stmt->execute();
while(@row = $stmt->fetchrow_array()){
    push(@dbData, [@row]);}
$dbh->disconnect();
print Dumper \@dbData;
```

Perl Script Connect to Vertica

```
use DBI;
use Data::Dumper;

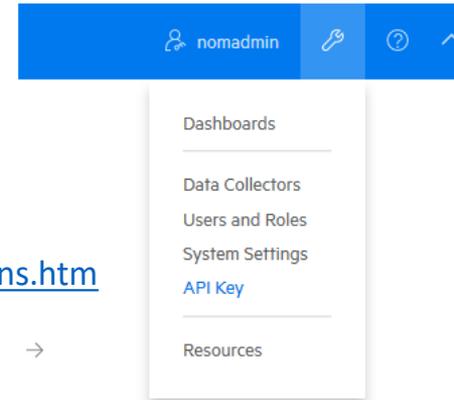
$dbh = DBI->connect('dbi:ODBC:VerticaDSN', 'userid', 'password');
unless(defined $dbh){
    die "Failed to connect: $DBI::errstr";
}

$sth = $dbh->prepare("my query here");
$sth->execute();
while(@row = $sth->fetchrow_array()){
    push(@dbData, [@row]);
}
$dbh->disconnect();

print Dumper \@dbData;
```

Posting Data To BVD

- BVD URL to post data to
 - For BVD in containers:
 - `https://<hostname>/bvd-receiver/api/submit/<API key>`
 - For Classic BVD:
 - `http(s)://<hostname>:<port>/api/submit/<API key>`
- For detailed information:
 - https://docs.microfocus.com/OMi/10.70/OBM_Integrations_Guide/Content/Integrations.htm
- API key
 - f83835880f7749dc906e7e570e4ba426



API Key →

f83835880f7749dc906e7e570e4ba426

Generating a new API key will invalidate your existing one. Therefore you have to adjust all your data senders to use this new API key instead of the old one.

Generate New

BVD URL Post And Data Stream

<http://bvd.domain.com:12224/api/submit/555555/dims/host,metricName,uncleJoe>

“dims” are required in your JSON payload

JSON file:

```
[
  {
    \"host\": \"10.10.11.190\",
    \"metricName\": \"serverStats\",
    \"uncleJoe\": \"AliveAndKicking\",
    \"cpuUtilization\": \"$cpuUtilization\",
    \"responseStatus\": \"200\",
    \"physMemUsed\": \"$physMemUsed\",
    \"pagesPerSec\": \"$pagesPerSec\",
    \"pingResponse\": \"$pingResponse\"
  }
]
```

Widget

group1151 (Line/Area Chart)

Properties

Data Channel:

10.10.11.190 serverStats AliveAndKicking

Data Field:

host
metricName
uncleJoe
cpuUtilization
responseStatus
physMemUsed
pagesPerSec
pingResponse

Chart Period in minutes:

Post to BVD Using curl

- curl
 - Windows and Linux – very widely used; open source
 - Double back slashes are for Perl scripting
 - <https://curl.haxx.se/>

```
C:\curl-7.53.1\src\curl.exe -X POST --url http://bvd.domain.com:12224/api/submit/5555555/dims/host,metricName,TxType -k -v -H "Content-Type:application/json" -d @"$dataFile2" 2>&1
```

Download working scripts from the hand outs section of the webinar.



Questions & Answers

Please type your questions in the questions pane

YOUR INDEPENDENT MICRO FOCUS SOFTWARE USER COMMUNITY



Upcoming Vivit Webinars

July 12, 2018

Integrate your Micro Focus IT Operations into a Single Pane of Glass 4D Solution

8:00 - 9:00 AM PDT (Los Angeles), 11:00 AM- 12:00 PM EDT (New York), 17:00 - 18:00 CET (Frankfurt)

<https://www.vivit-worldwide.org/events/EventDetails.aspx?id=1116786&group=>

YOUR INDEPENDENT MICRO FOCUS SOFTWARE USER COMMUNITY



Thank You

- Complete the short survey
- To speak on a future Vivit SIG Talk or to recommend a speaker, please complete the survey:
<https://www.surveymonkey.com/r/36V5XXH>

YOUR INDEPENDENT MICRO FOCUS SOFTWARE USER COMMUNITY





Thank You
vivit-worldwide.org

